VMware vSphere: Install, Configure, Manage [V8]

Lecture Manual

VMware® Education Services
VMware, Inc.
www.vmware.com/education

VMware vSphere: Install, Configure, Manage [V8]

Lecture Manual

ESXi 7 and vCenter 8

Part Number EDU-EN-VSICM8-LECT (11-NOV-2022)

Contents

Module 1

Course Introduction

1-2 Course Introduction

1-3 Importance

As a vSphere administrator, you require knowledge about vSphere components and resources and how they work together in your environment. You also require practical skills in installing, deploying, and managing your vSphere environment. By developing your knowledge and skills, you can build and run a highly scalable vSphere virtual infrastructure.

1-4 Learner Objectives (1)

- Install and configure ESXi hosts
- Deploy and configure vCenter
- Use the vSphere Client to create the vCenter inventory and assign roles to vCenter users
- Create and configure virtual networks using vSphere standard switches and distributed switches
- Create and configure datastores using storage technologies supported by vSphere
- Use the vSphere Client to create virtual machines, templates, clones, and snapshots
- Create a content library for deploying virtual machines
- Manage virtual machine resource use

1-5 Learner Objectives (2)

- Configure ESXi hosts to support vSphere vMotion and vSphere Storage vMotion migrations
- Create and configure a vSphere cluster that is enabled with vSphere High Availability and vSphere Distributed Resource Scheduler
- Manage the life cycle of vSphere to keep vCenter and ESXi hosts up to date

1-6 Course Outline

1. Course Introduction
2. vSphere and Virtualization Overview
3. Installing and Configuring vSphere ESXi
4. Deploying and Configuring vCenter
5. Configuring vSphere Networking
6. Configuring vSphere Storage
7. Deploying Virtual Machines
8. Managing Virtual Machines
9. Deploying and Configuring vSphere Clusters
10. Managing the vSphere Lifecycle

1-7 Typographical Conventions

The following typographical conventions are used in this course.

Conventions	Use and Examples
`Monospace`	Identifies command names, command options, parameters, code fragments, error messages, filenames, folder names, directory names, and path names: • Run the `esxtop` command. • … found in the `/var/log/messages` file.
`Monospace Bold`	Identifies user inputs: • Enter **`ip a`**.
Boldface	Identifies user interface controls: • Click the **Configuration** tab.
Italic	Identifies book titles: • *vSphere Virtual Machine Administration*
< >	Indicates placeholder variables: • <ESXi_host_name> • … the `Settings/<Your_Name>.txt` file

1-8 References (1)

Title	Location
VMware ESXi Installation and Setup	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-8.0-installation-setup-guide.pdf
vCenter Server Installation and Setup	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-vcenter-80-installation-guide.pdf
vCenter Server and Host Management	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-management-guide.pdf
vSphere Virtual Machine Administration	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-server-80-vm-administration.pdf
vSphere Networking	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-networking-guide.pdf
vSphere Storage	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-storage-guide.pdf

1-9 References (2)

Title	Location
vSphere Resource Management	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-resource-management-guide.pdf
vSphere Availability	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-availability-guide.pdf
Managing Host and Cluster Lifecycle	https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-esxi-vcenter-80-vsphere-lifecycle-manager.pdf
VMware Compatibility Guide	https://vmware.com/resources/compatibility
VMware Product Lifecycle Matrix	https://lifecycle.vmware.com
VMware Configuration Maximums	https://configmax.vmware.com

1-10 VMware Online Resources

Documentation for vSphere: https://docs.vmware.com/

VMware Communities: https://communities.vmware.com

- Start a discussion.
- Access the knowledge base.
- Access documentation, technical papers, and compatibility guides.
- Access communities.
- Access user groups.

VMware Support: https://www.vmware.com/support

VMware Hands-on Labs: https://hol.vmware.com

VMware Learning: https://www.vmware.com/learning

- Access course catalog.
- Access worldwide course schedule.

1-11 VMware Learning Overview

You can access the following Education Services:

- VMware Learning Paths:
  - Help you find the course that you need based on the product, your role, and your level of experience
  - Can be accessed at https://vmware.com/learning
- VMware Customer Connect Learning, which is the official source of digital training, includes the following options:
  - On Demand Courses: Self-paced learning that combines lecture modules with hands-on practice labs
  - VMware Lab Connect: Self-paced, technical lab environment where you can practice skills learned during instructor-led training
  - Certification Exam Prep: Comprehensive video-based reviews of exam topics and objectives to help you take your certification exam
- For more information, see https://vmware.com/learning/connect-learning.

VMware Customer Connect Learning offers different subscriptions. For a comparison of the subscriptions, go to https://vmware.com/learning/connect-learning and click **WHICH SUBSCRIPTION IS RIGHT FOR YOU?**.

1-12 VMware Certification Overview

VMware certifications validate your expertise and recognize your technical knowledge and skills with VMware technology.

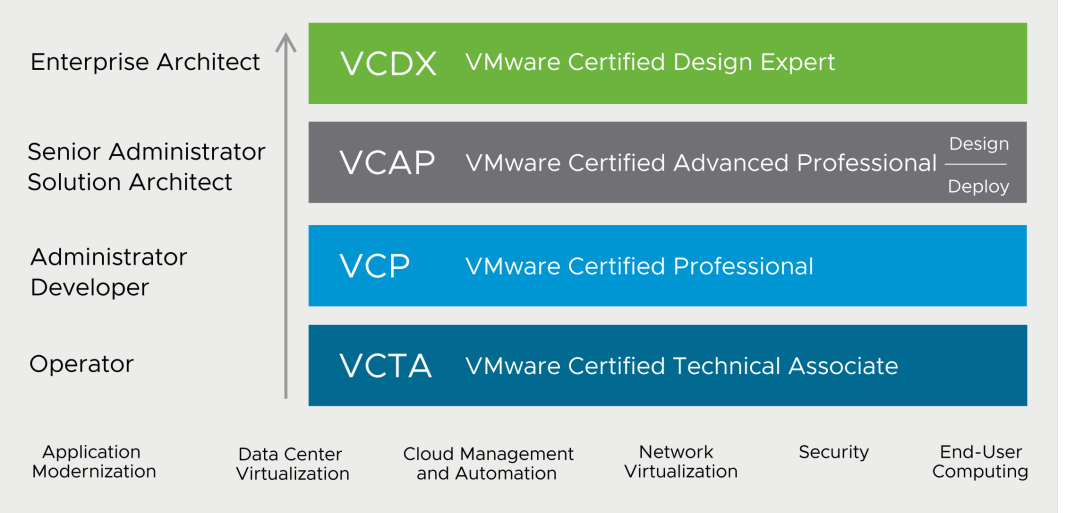

VMware certification sets the standards for IT professionals who work with VMware technology. Certifications are grouped into technology tracks. Each track offers one or more levels of certification (up to four levels).

For the complete list of certifications and details about how to attain these certifications, see https://vmware.com/certification.

1-13 VMware Credentials Overview

VMware badges are digital emblems of skills and achievements. Career certifications align to job roles and validate expertise across a solution domain. Certifications can cover multiple products in the same certification.

Specialist certifications and skills badges align to products and verticals and show expanded expertise.

Digital badges have the following features:

- Easy to share in social media
- Validate and verify achievement
- Contain metadata with skill tags and accomplishments
- Based on Mozilla's Open Badges standard

For the complete list of digital badges, see https://www.pearsonvue.com/vmware/badging.

Module 2

vSphere and Virtualization Overview

2-2 Importance

As a vSphere administrator, you must be familiar with the components on which vSphere is based.

You must also understand the following concepts:

- Virtualization, the role of the ESXi hypervisor in virtualization, and virtual machines
- Fundamental vSphere components and the use of vSphere in the software-defined data center
- When to use each one of the user interfaces to administer and manage vSphere environments

2-3 Module Lessons

1. vSphere Virtualization Overview
2. vSphere Virtualization of Resources

2-4 Lesson 1: vSphere Virtualization Overview

2-5 Learner Objectives

- Explain basic virtualization concepts
- Describe vSphere
- Describe how vSphere fits in the software-defined data center
- Describe vSphere+
- Recognize the user interfaces for accessing vSphere

2-6 Terminology (1)

Virtualization is associated with several key concepts, products, and features.

Term	Definition	Examples
Operating system	Software designed to allocate physical resources to applications	Microsoft Windows, Linux
Application	Software that runs on an operating system, consuming physical resources	Microsoft Office, Chrome
Hypervisor	Specialized operating system designed to run VMs	ESXi, Workstation, Fusion
Virtual machine	Specialized application that abstracts hardware resources into software	
Guest	The operating system that runs in a VM (also called the guest operating system)	Microsoft Windows, Linux
Host	Physical computer that provides resources to the ESXi hypervisor	

2-7 Terminology (2)

Term	Definition
vSphere	Server virtualization product of VMware that combines the ESXi hypervisor and the vCenter Server management platform
Cluster	Group of ESXi hosts whose resources are shared by VMs
vSphere vMotion	Feature that supports the migration of powered-on VMs from host to host without service interruption
vSphere HA	Cluster feature that protects against host hardware failures by restarting VMs on hosts that are running normally
vSphere DRS	Cluster feature that uses vSphere vMotion to place VMs on hosts and ensure that each VM receives the resources that it needs

2-8 Data Center Physical Infrastructure

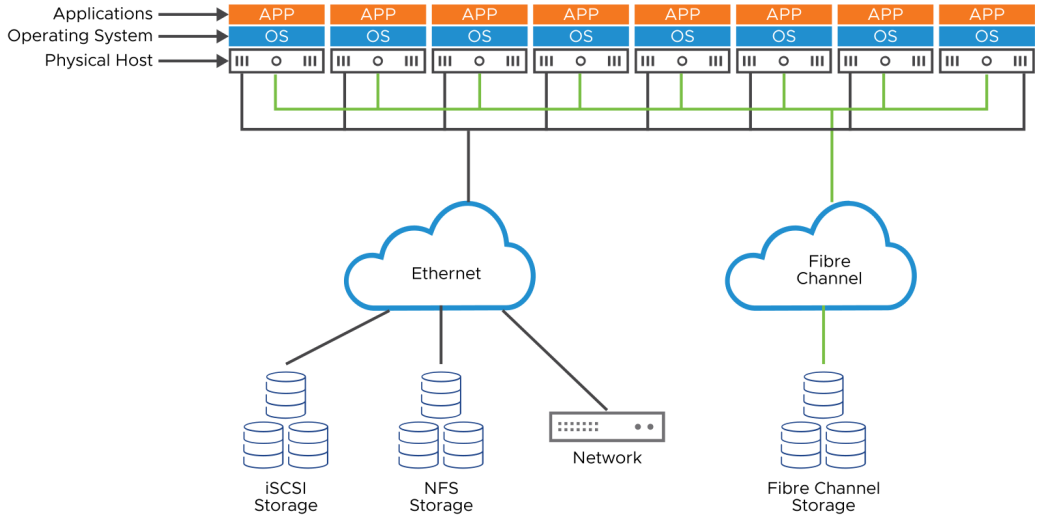

Traditionally, operating systems and software run on a physical computer. Several challenges are posed to running a large number of physical servers in a data center. The model is not flexible and can be inefficient. The planning and cost of proper infrastructure (square footage, rack space, power, cooling, cabling, and server provisioning) are but a few of the problems that IT staff must address.

Typically, a one-to-one relationship exists between a physical computer and the software that it runs. This relationship leaves most computers vastly underused. The cost of the space and power required to house, run, and keep these systems cool can be expensive.

Provisioning physical servers is a time-consuming process. In nonvirtualized environments, time must be allotted to procure new hardware, place it in the data center, install an operating system, and patch the operating system. Installing and configuring the required applications can take weeks. This process also includes a myriad of other tasks to integrate the system into the infrastructure, for example, configuring firewall rules, enabling switch ports, and provisioning storage.

2-9 About the Virtual Infrastructure

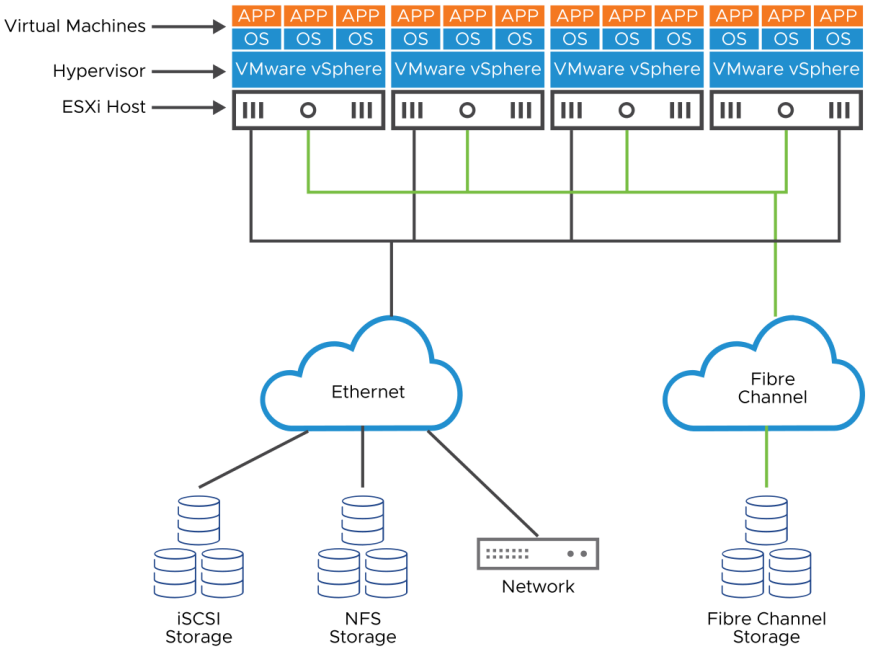

Virtualization enables you to run more workloads on a single server by consolidating the environment so that your applications run on virtual machines. Converting to a virtualized data center reduces the required data center square footage, rack space, power, cooling, cabling, storage, and network components by reducing the sheer number of physical machines.

The reduction of physical machines can be realized by converting physical machines to virtual machines and consolidating the converted machines onto a single host.

Using virtualization technology also changes the way servers are provisioned. You do not need to wait for the hardware to be procured or cabling to be installed. Virtual machine provisioning is performed using an intuitive graphical user interface. In contrast to the long process of deploying physical servers, deploying virtual machines can be deployed in a matter of minutes.

2-10 About Virtual Machines

A virtual machine (VM) is a software representation of a physical computer and its components.

Virtual Machine

The virtual machine includes the following components:

- Guest operating system
- VMware Tools
- Virtual resources, such as:
  - CPU and memory
  - Network adapters
  - Disks and controllers
  - GPUs

A virtual machine (VM) includes a set of specification and configuration files and is supported by the physical resources of a host. Every VM has virtual devices that provide the same functionality as physical hardware but are more portable, more secure, and easier to manage.

VMs include an operating system, applications, VMware Tools, virtual resources, and hardware that you manage in the same way that you manage a physical computer.

VMware Tools is a bundle of drivers. Using these drivers, the guest operating system can interact efficiently with the virtual hardware. VMware Tools also adds extra functionality so that ESXi can better manage the VMs use of physical hardware.

2-11 Benefits of Using Virtual Machines

Physical machines have the following constraints:

- Difficult to move or copy
- Bound to a specific set of hardware components
- Often have a short life cycle
- Require personal contact to upgrade hardware

Virtual machines provide the following benefits:

- Easy to move or copy
- Independent of physical hardware because VMs are encapsulated into files
- Isolated from other VMs running on the same physical hardware
- Insulated from physical hardware changes

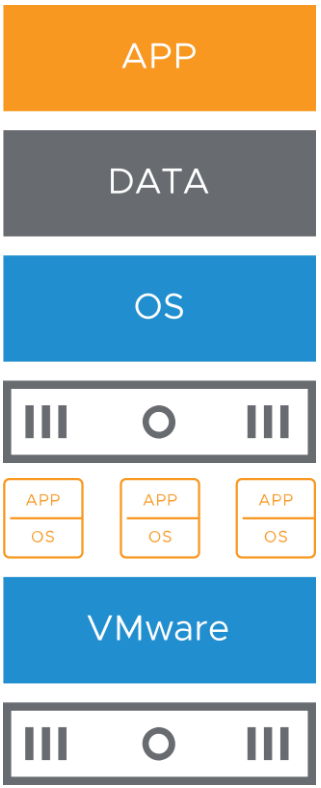

In a physical machine, the operating system (for example, Windows or Linux) is installed directly on the hardware. The operating system requires specific device drivers to support specific hardware. If the computer is upgraded with new hardware, new device drivers are required.

If applications interface directly with hardware drivers, an upgrade to the hardware, drivers, or both can have significant repercussions if incompatibilities exist. Because of these potential repercussions, hands-on technical support personnel must test hardware upgrades against a wide variety of application suites and operating systems. Such testing costs time and money.

Virtualizing these systems saves on such costs because VMs are 100 percent software.

Multiple VMs are isolated from one another. You can have a database server and an email server running on the same physical computer. The isolation between the VMs means that software-dependency conflicts are not a problem. Even users with system administrator privileges on a VM guest operating system cannot breach the layer of isolation to access another VM. The users must explicitly be granted access by the ESXi system administrator. As a result of VM isolation, if a guest operating system running in a VM fails, other VMs on the same host are unaffected and continue to run.

A guest operating system failure does not affect access and performance:

- Users can still access the other VMs.
- The operational VMs can access the resources that they need.
- The other VMs can still perform.

With VMs, you can consolidate your physical servers and make more efficient use of your hardware. Because a VM is a set of files, features that are not available or not as efficient on physical architectures are available to you, for example:

- You can rapidly and consistently provision VMs.
- With VMs, you can use live migration, fault tolerance, high availability, and disaster recovery scenarios to increase the uptime and reduce recovery time from failures.
- You can use multitenancy to mix VMs to specialized configurations, such as a demilitarized zone (DMZ).

With VMs, you can support legacy applications and operating systems on newer hardware when maintenance contracts on the existing hardware expire.

2-12 About vSphere

vSphere is the virtualization platform that includes two core administrative components for running virtual machines:

- ESXi: Hypervisor on which you run virtual machines
- vCenter: Central administration platform for ESXi hosts, virtual machines, storage, and networking

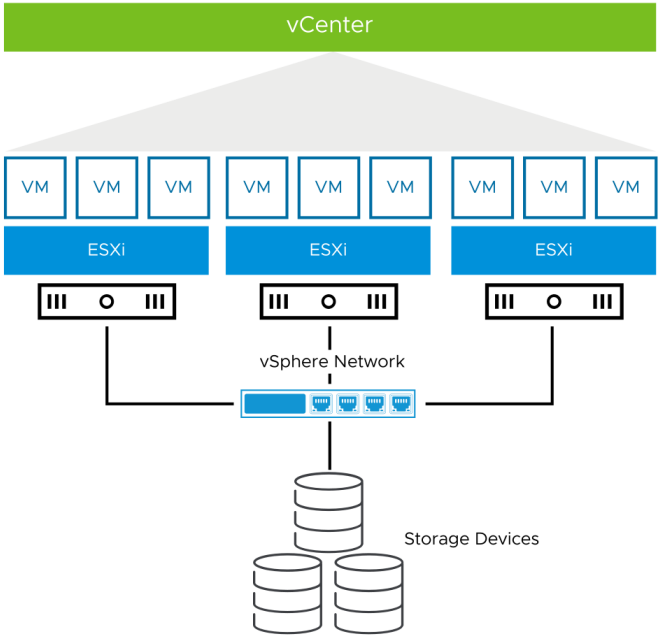

vCenter Server and ESXi are core to run a vSphere environment.

vSphere is a virtualization platform that provides virtualization, management, resource optimization, application availability, and operational automation capabilities.

vSphere virtualizes and aggregates the underlying physical hardware resources across multiple systems and provides pools of virtual resources to the data center.

In addition, vSphere provides a set of distributed services that enable detailed, policy-driven resource allocation, high availability, and scalability of the entire virtual data center.

2-13 Types of Virtualization

Virtualization is the process of creating a software-based representation of a physical unit, such as a server, desktop, network, or storage device.

Virtualization is the single most effective way to reduce IT expenses while boosting efficiency and agility for all business sizes.

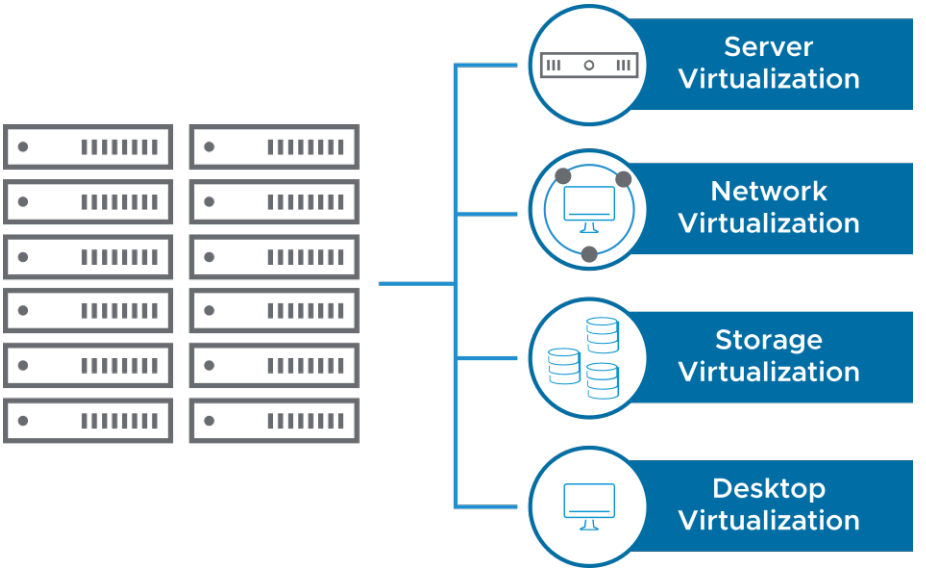

Server virtualization allows companies to create multiple operating systems that run on a single physical server, called VMs. Each VM has access to the underlying server's computing resources.

Network virtualization is the complete reproduction of a physical network in software. Applications run on the virtual network exactly as if on a physical network.

Storage virtualization is the process of creating a software-based representation of network storage devices into what appears to be a single unit.

By deploying desktops as a managed service, you can respond more quickly to changing needs and opportunities.

2-14 About the Software-Defined Data Center

In a software-defined data center (SDDC), all infrastructure is virtualized, and the control of the data center is automated by software. vSphere is the foundation of the SDDC.

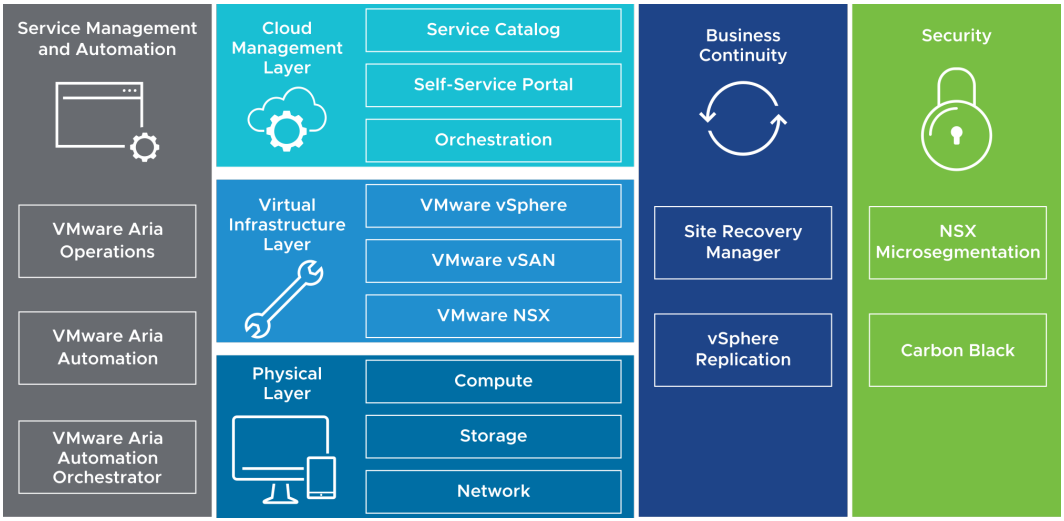

A software-defined virtual data center (SDDC) is deployed with isolated computing, storage, networking, and security resources that are more efficient, more manageable, and more scalable than the traditional, hardware-based data center.

All the resources (CPU, memory, disk, and network) of a software-defined data center are an abstraction in software of the physical resources. The abstraction provides the benefits of virtualization at all levels of the infrastructure, independent of the physical infrastructure.

An SDDC can include the following components:

- Service management and automation: Use service management and automation to track and analyze the operation of multiple data sources in the multiregional SDDC. Deploy VMware Aria Operations and vRealize Log Insight across multiple nodes for continued availability and increased log ingestion rates.
- Cloud management layer: The layer includes the service catalog, which has the facilities to deploy. The cloud management layer also includes orchestration, which provides the workflows to deploy catalog items, and the self-service portal for end users to access and use the SDDC.
- Virtual infrastructure layer: The layer establishes a robust virtualized environment that all the other solutions can integrate. The virtual infrastructure layer includes the virtualization platform for the hypervisor, pools of resources, and virtualization control. Additional

processes and technologies build on the infrastructure to support Infrastructure as a Service (IaaS) and Platform as a Service (PaaS).

- Physical layer: The lowest layer of the solution includes compute, storage, and network components.
- Security: Customers use the layer of the platform to meet the demanding compliance requirements for virtualized workloads and manage business risk.

2-15 About vSphere+

VMware vSphere+ is a subscription-based offering that brings the benefits of cloud to on-premises workloads.

vSphere+ consists of on-premises and cloud components that interact with each other.

vSphere+ lets you centrally manage your on-premises workloads from a cloud console, with access to cloud services.

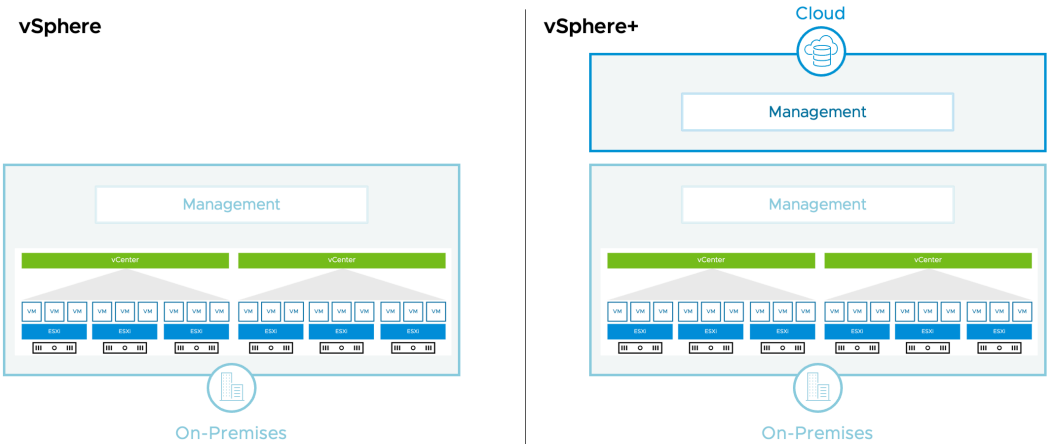

vSphere+ supports traditional workloads based on VMs, plus modern workloads based on containers. There is no difference between vSphere+ and vSphere in the types of workloads supported. Your workloads run on-premises on ESXi hosts. None of your workloads or on-premises vSphere infrastructure are moved to the cloud.

vSphere+ consists of both on-premises and cloud components:

- On-premises components:
  - vCenter instances and ESXi hosts
  - Cloud gateway that connects vCenter instances to the VMware Cloud Console
- Cloud components:
  - VMware Cloud Console, where you can centrally manage on-premises infrastructure and access cloud services
  - Cloud services for admins (or IT operations) and developers (or DevOps) that augment and enhance on-premises capabilities

vSphere+ has no special hardware requirements beyond the basic vSphere requirements.

For more details on vSphere+, see https://vsphereplus.com.

2-16 vSphere+: Accessing Cloud Services

vSphere+ lets you access cloud services to augment and enhance on-premises capabilities:

- Admin Services
  - Inventory management
  - Events and alerts management
  - VM provisioning
  - Lifecycle management
  - Configuration management
- Developer Services
  - Tanzu Kubernetes Grid
  - Tanzu integrated services
- Add-On Services
  - Disaster recovery

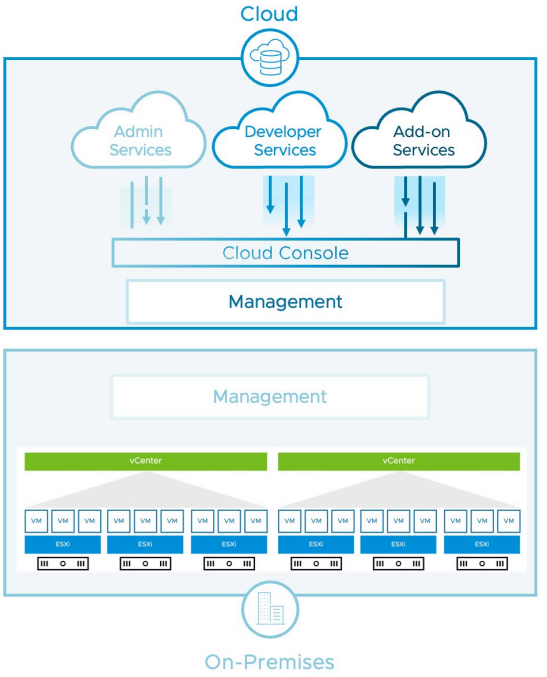

VMware is developing vSphere+ add-on services to expedite disaster recovery, ransomware protection, capacity planning and more. For a list of add-on services currently available for purchase or in development, contact your VMware partner or sales representative.

vSphere+ includes extensive developer services, including VM service, Storage service, Network service, Registry service, Tanzu Kubernetes Grid service, Tanzu integrated services, Tanzu Mission-Control Essentials and more. These services are included with vSphere+ at no additional charge.

vSphere+ includes extensive admin services, including Global Inventory service, Event View service, Security Health Check service, VM Provisioning service, Lifecycle Management service, Configuration Management service and more. These services are included with vSphere+ at no additional charge.

2-17 vSphere User Interfaces

You can use the vSphere Client, PowerCLI, VMware Host Client, vSphere ESXi Shell and ESXCLI to interact with the vSphere environment.

For information about ports and protocols, see https://ports.vmware.com.

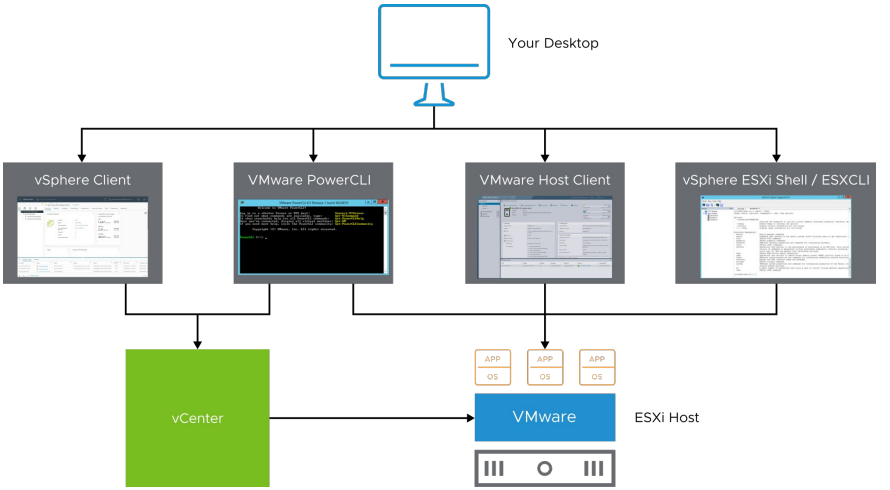

VMware Host Client is an HTML5-based user interface that you can use to manage individual ESXi hosts directly when vCenter Server is unavailable. VMware Host Client is provided from ESXi. You access it from a supported browser at https://<ESXi_FQDN_or_IP_Address>/ui.

The vSphere Client is an HTML5-based client. You manage the vSphere environment with the vSphere Client by connecting to vCenter Server and managing the vCenter Server object inventory. You access the vSphere Client from a supported browser at https://<vCenter_FQDN_or_IP_Address>/ui.

PowerCLI is a command-line and scripting tool that is built on Windows PowerShell. The tool provides a PowerShell interface to the vSphere API. PowerCLI provides more than 700 cmdlets for managing and automating vSphere. For more information about PowerCLI, see https://code.vmware.com/web/tool/12.0.0/vmware-powercli.

vSphere ESXi Shell provides a command-line interface for running essential maintenance commands. You use the vSphere ESXi Shell mainly for troubleshooting purposes.

From the vSphere ESXi Shell, you can run ESXCLI commands. You can use the ESXCLI command set lets you remotely manage ESXi hosts. ESXCLI commands can be run against a vCenter system and target any ESXi system.

You can download ESXCLI from the VMware {code} page at https://code.vmware.com/web/tool/7.0/esxcli. You can install ESXCLI on a Windows or Linux system.

2-18 Lab 1: Accessing the Lab Environment

Log in to the Console VM and access the vSphere Client and VMware Host Client:

1. Access the Student Console
2. Log In to an ESXi Host with VMware Host Client
3. Log In to vCenter with the vSphere Client

2-19 Review of Learner Objectives

- Explain basic virtualization concepts
- Describe vSphere
- Describe how vSphere fits in the software-defined data center
- Describe vSphere+
- Recognize the user interfaces for accessing vSphere

2-20 Lesson 2: vSphere Virtualization of Resources

2-21 Learner Objectives

- Explain how ESXi interacts with resources:
  - CPUs
  - Memory
  - Networks
  - Storage
  - GPUs

2-22 Virtual Machine: Guest and Consumer of ESXi Host

Any application in any supported OS can run in a VM (guest) and use CPU, memory, disk, and network from host-based resources.

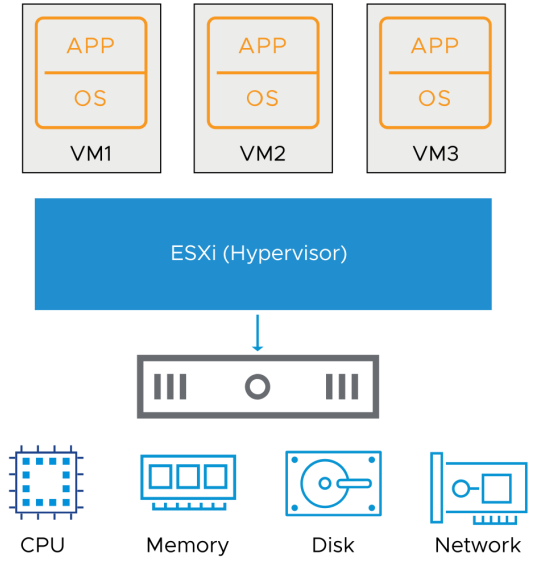

A virtual machine is an abstraction in software of a physical machine. ESXi manages the physical resources that are used by virtual machines as virtual resources.

For the list of all supported operating systems, see *VMware Compatibility Guide* at https://www.vmware.com/resources/compatibility.

2-23 Physical and Virtual Architecture

Virtualization technology abstracts physical components to software components and provides solutions for many IT problems.

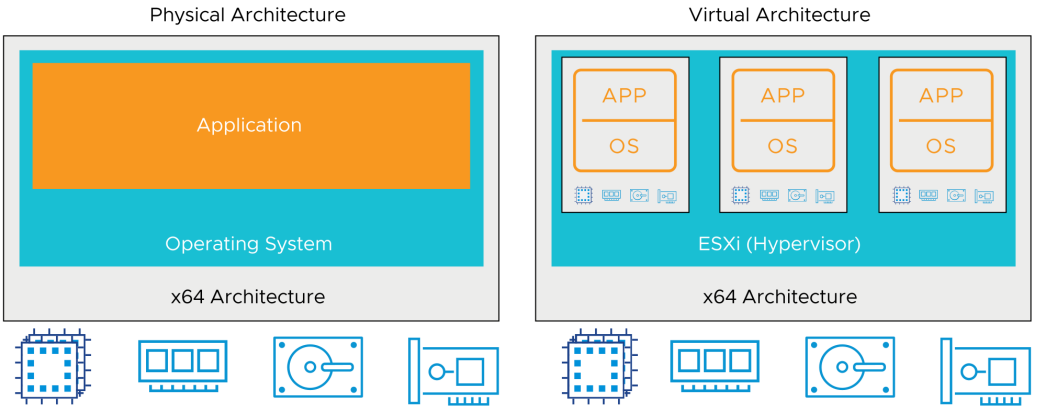

You can use virtualization to consolidate and run multiple workloads as VMs on a single computer.

The slide shows the differences between a virtualized and a non-virtualized host. In traditional architectures, the operating system interacts directly with the installed hardware. The operating system schedules processes to run, allocates memory to applications, sends and receives data on network interfaces, and reads from and writes to attached storage devices. In comparison, a virtualized host interacts with the installed hardware through a thin layer of software called the virtualization layer or hypervisor.

The hypervisor provides physical hardware resources dynamically to VMs as needed to support the operation of the VMs. With the hypervisor, VMs can operate with a degree of independence from the underlying physical hardware. For example, a VM can be moved from one physical host to another. In addition, its virtual disks can be moved from one type of storage to another without affecting the functioning of the VM.

2-24 Physical Resource Sharing

Multiple VMs, running on a physical host, share computing, memory, network, and storage resources of the host.

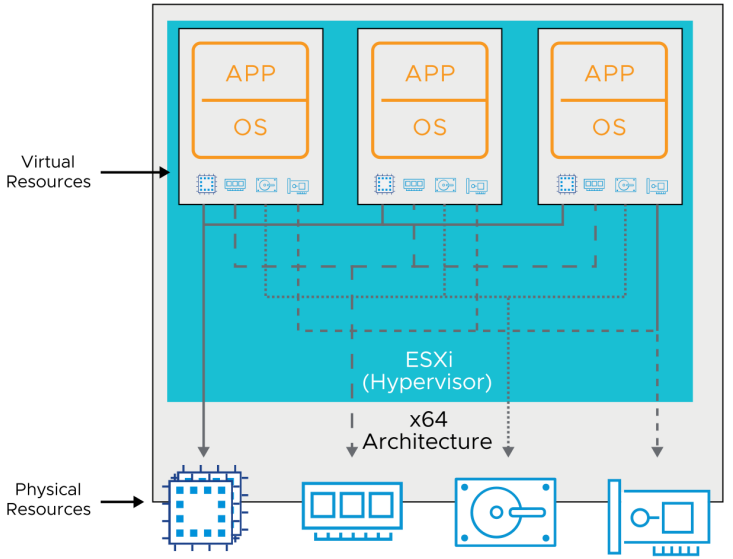

With virtualization, you can run multiple VMs on a single physical host, with each VM sharing the resources of one physical computer across multiple environments. VMs share access to CPUs and are scheduled to run by the hypervisor. In addition, VMs are assigned their own region of memory to use and share access to the physical network cards and disk controllers. Different VMs can run different operating systems and applications on the same physical computer.

When multiple VMs run on an ESXi host, each VM is allocated a portion of the physical resources. The hypervisor manages VM schedules, similarly to traditional operating system management of memory allocation and application scheduling. These VMs run on various CPUs. The ESXi hypervisor can also over commit memory. Memory is over committed when the combined virtual RAM of your powered-on VMs is larger than the physical RAM available on the host.

VMs, like applications, use network and disk bandwidth. However, VMs are managed with elaborate control mechanisms to manage how much access is available for each VM. With the default resource allocation settings, all VMs associated with the same ESXi host receive an equal share of available resources.

2-25 CPU Virtualization

In a physical environment, the operating system assumes the ownership of all the physical CPUs in the system.

CPU virtualization emphasizes performance and runs directly on the available CPUs.

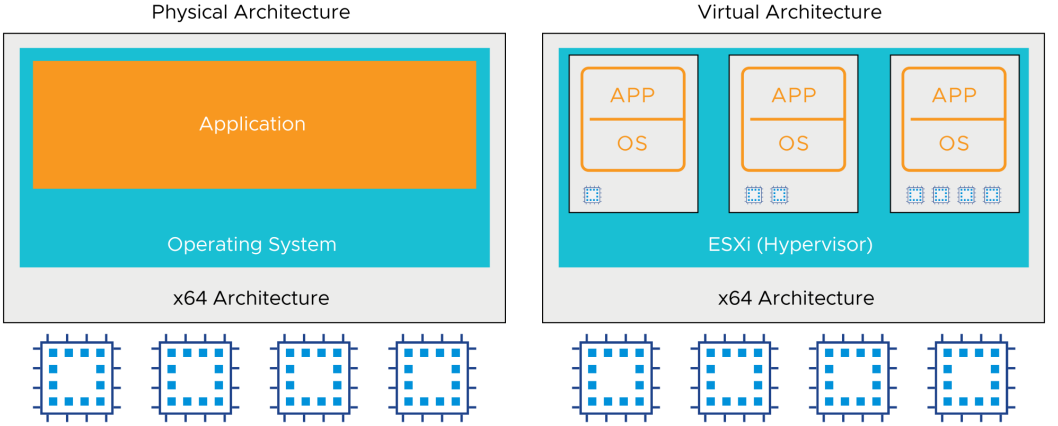

The virtualization layer runs instructions only when needed to make VMs operate as if they were running directly on a physical machine. CPU virtualization is not emulation. With a software emulator, programs can run on a computer system other than the one for which they were originally written. Emulation provides portability, but might negatively affect performance. CPU virtualization is not emulation because the supported guest operating systems are designed for x64 processors. Using the hypervisor, the operating systems can run natively on the hosts' physical x64 processors.

When many virtual VMs run on an ESXi host, the VMs might compete for CPU resources. When CPU contention occurs, the ESXi host time slices the physical processors across all virtual machines so that each VM runs as if it has a specified number of virtual processors.

2-26 Physical and Virtualized Host Memory Usage

In a physical environment, the operating system assumes the ownership of all the physical memory in the system.

Memory pages are allocated to virtual machines on first access.

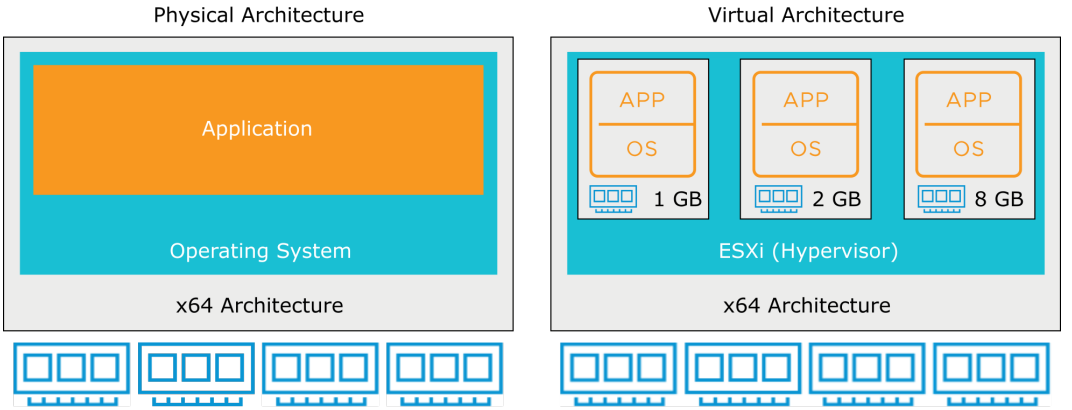

When an application starts, it uses the interfaces provided by the operating system to allocate or release virtual memory pages during the execution. Virtual memory is an old technique used in most general-purpose operating systems. Operating systems use virtual memory to present more memory to applications than to which they have physical access. Almost all modern processors have hardware to support virtual memory.

Virtual memory creates a uniform virtual address space for applications. With the operating system and hardware, virtual memory can handle the address translation between the virtual address space and the physical address space. The technique adapts the execution environment to support large address spaces, process protection, file mapping, and swapping in modern computer systems.

In a virtualized environment, the VMware virtualization layer creates a contiguous addressable memory space for the VM when it is started. The allocated memory space is configured when the VM is created and has the same properties as the virtual address space. With this configuration, the hypervisor can run multiple VMs simultaneously while protecting the memory of each VM from being accessed by others.

2-27 Physical and Virtual Networking

Virtual Ethernet adapters and virtual switches are key virtual networking components.

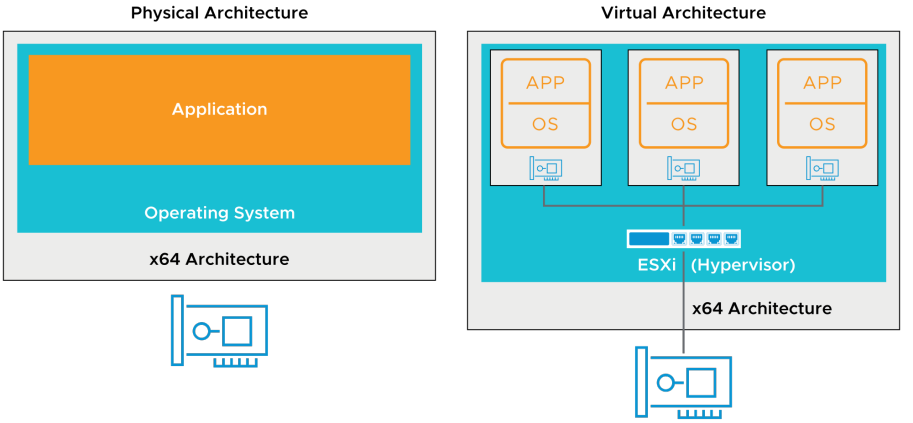

A VM can be configured with one or more virtual Ethernet adapters. VMs use virtual switches on the same ESXi host to communicate with one another by using the same protocols that are used over physical switches, without the need for additional hardware. Virtual switches also support VLANs that are compatible with standard VLAN implementations from other networking equipment vendors. With VMware virtual networking, you can link local VMs together and link local VMs to the external network through a virtual switch.

A virtual switch, like a physical Ethernet switch, forwards frames at the data link layer. An ESXi host might contain multiple virtual switches. The virtual switch connects to the external network through outbound Ethernet adapters, called vmnics. The virtual switch can bind multiple vmnics together, like NIC teaming on a traditional server, offering greater availability and bandwidth to the VMs using the virtual switch.

Virtual switches are similar to modern physical Ethernet switches in many ways. Like a physical switch, each virtual switch is isolated and has its own forwarding table. Every destination that the switch looks up can match only ports on the same virtual switch where the frame originated. The feature improves security and makes it difficult for hackers to break virtual switch isolation.

Virtual switches also support VLAN segmentation at the port level, so that each port can be configured as an access or trunk port, providing access to either single or multiple VLANs.

However, unlike physical switches, virtual switches do not require the spanning tree protocol (STP) because a single-tier networking topology is enforced. Multiple virtual switches cannot be interconnected, and network traffic cannot flow directly from one virtual switch to another virtual switch on the same host. Virtual switches provide all the ports that you need in one switch. Virtual switches do not need to be cascaded because virtual switches do not share physical Ethernet adapters and leaks do not occur between virtual switches.

2-28 Physical File Systems and Datastores

vSphere datastores provide a distributed storage architecture, where multiple ESXi hosts can read or write to the shared storage concurrently.

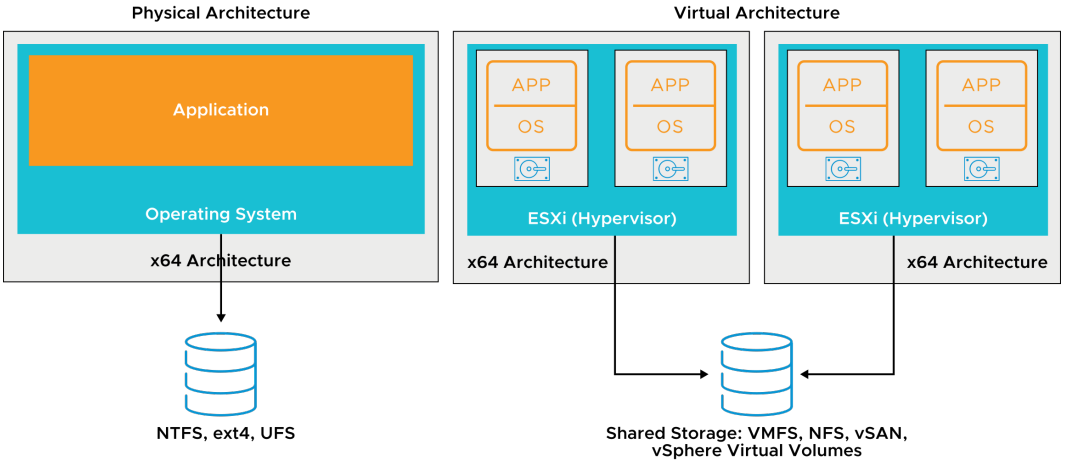

To store virtual disks, ESXi uses datastores, which are logical containers that hide the specifics of physical storage from VMs and provide a uniform model for storing VM files.

vSphere supports the following types of datastores:

- VMFS: VMware Virtual Machine File System (VMFS) is a clustered file system that provides storage virtualization optimized for virtual machine. Multiple ESXi hosts can read and write to the same VMFS datastore simultaneously.
- NFS: Network File System (NFS) is a file-sharing protocol that ESXi hosts use to communicate with a network-attached storage (NAS) device.
- vSAN: vSAN is a software-defined storage solution that provides shared storage for virtual machines. vSAN virtualizes local physical storage in the form of HDD or SSD devices on ESXi hosts in a cluster, turning them into a unified datastore.
- vSphere Virtual Volumes: The datastore virtualizes SAN and NAS devices by abstracting physical hardware resources to logical pools of capacity. Storage arrays or servers are designed to manage all aspects of vSphere Virtual Volumes datastores and ESXi hosts have no direct access to a vSphere Virtual Volumes storage.

2-29 GPU Virtualization

GPU graphics devices optimize complex graphics operations. These operations can run at high performance without overloading the CPU.

Virtual GPUs can be added to VMs for the following use cases:

- Rich 2D and 3D graphics
- VMware Horizon virtual desktops
- Graphics-intensive applications
- Scientific computation applications
- Artificial intelligence (AI) and machine learning (ML) workloads

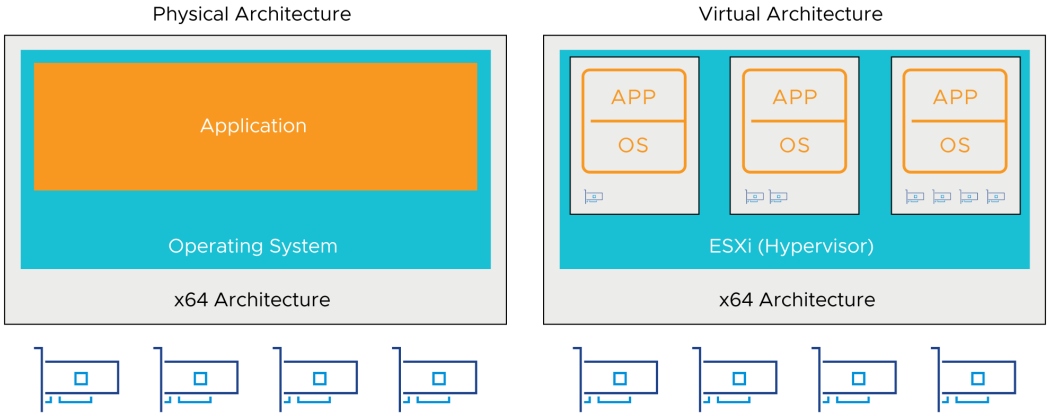

GPUs (graphical processing units) can be used by developers of server applications. Although servers do not usually have monitors, GPU support is important and relevant to server virtualization. You can configure VMs with up to four vGPU devices to cover use cases requiring multiple GPU accelerators. VMware supports AMD and NVIDIA graphics cards.

2-30 Sharing GPUs with vSphere Bitfusion

vSphere Bitfusion virtualizes hardware accelerators such as GPUs to provide a pool of shared, network-accessible resources that support AI and ML workloads.

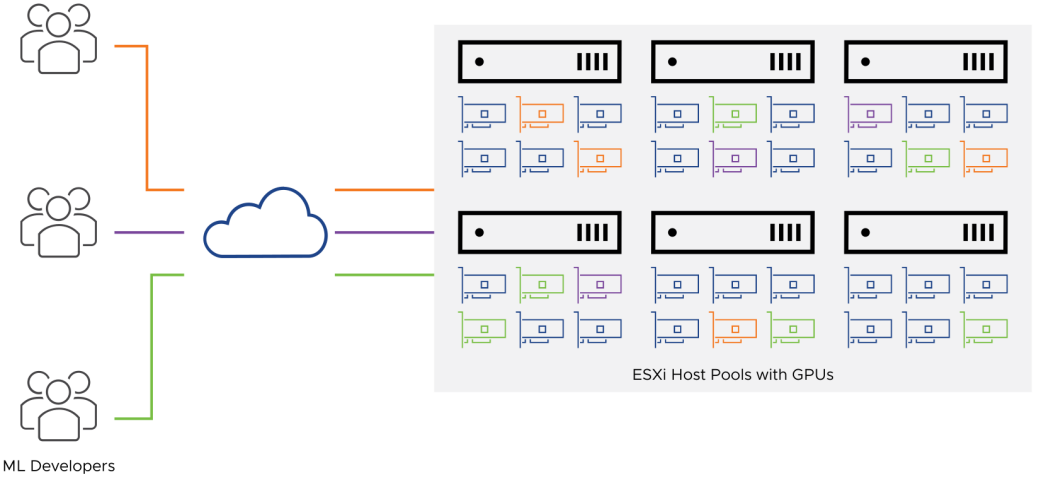

With vSphere Bitfusion, GPUs on one host can be used by virtual machines running on other hosts.

Business use cases for sharing GPUs include the following areas:

- Entertainment and visualizations, such as rendering of complex animations and 3D graphics
- Transportation and government, such as autonomous vehicles and smart city projects
- Manufacturing and shipping, for example, optimizing factory workflows and supply chain logistics
- Infectious disease and epidemiology, for example, vaccine research and modeling how viruses spread
- Higher education, such as allocating GPU resources for research both in and out of the classroom
- Retail, such as inventory management, analyzing buyer behavior, and helping to detect fraud
- Robotics, such as creating models for performing mundane or dangerous tasks
- Financial services, such as risk-analysis

2-31 Review of Learner Objectives

- Explain how ESXi interacts with resources:
  - CPUs
  - Memory
  - Networks
  - Storage
  - GPUs

2-32 Key Points

- Virtual machines are hardware independent.
- vSphere consists of two core components, ESXi and vCenter Server.
- The ESXi hypervisor runs directly on the host.
- VMs share the physical resources of the ESXi host on which they reside.
- vSphere abstracts CPUs, GPUs, memory, storage, and networking for VM use.

Questions?

Module 3

Installing and Configuring ESXi

3-2 Importance

ESXi is the virtualization platform on which you can create and run virtual machines. Proper configuration of the ESXi host ensures that virtual machines run in an environment that is reliable, secure, and performant.

3-3 Module Lessons

1. Installing and Configuring ESXi

3-4 Lesson 1: Installing and Configuring ESXi

3-5 Learner Objectives

- Describe the ESXi host architecture
- Navigate the Direct Console User Interface (DCUI) to configure an ESXi host
- Recognize the user account best practices
- Install an ESXi host
- Configure the ESXi host settings

3-6 About ESXi

ESXi is a bare-metal hypervisor that is licensed as a part of vSphere. A free version is also available as a standalone server.

ESXi has the following features:

- High security:
  - Host-based firewall
  - Memory hardening
  - Kernel module integrity
  - Trusted Platform Module (TPM 2.0)
  - UEFI secure boot
  - Encrypted core dumps
- Small disk footprint
- Quick boot for faster patching and upgrades
- Installable on hard disks, SAN LUNs, SSD, SATADOM, and diskless hosts

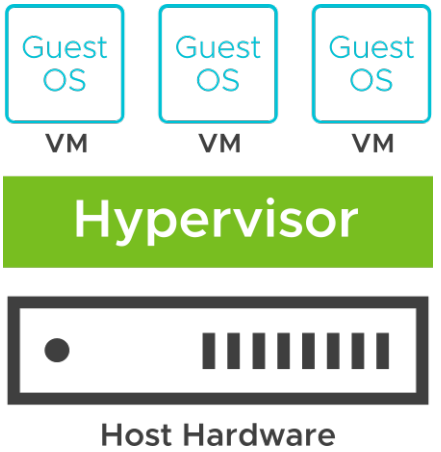

To verify that your physical servers are supported by ESXi 8.0, check *VMware Compatibility Guide* at https://www.vmware.com/resources/compatibility.

You can obtain the free version of ESXi, called vSphere Hypervisor, or you can purchase a licensed version with vSphere. ESXi can be installed on a hard disk, SAN LUNs, SSD, and SATADOM. ESXi can also be run on diskless hosts (directly into memory) with vSphere Auto Deploy.

You can use USB-based boot devices however, VMware encourages using SSD devices instead for improved reliability. For more information, see VMware knowledge base article 82515 at http://kb.vmware.com/kb/82515.

ESXi has a small disk footprint for added security and reliability. ESXi provides additional protection with the following features:

- Host-based firewall: To minimize the risk of an attack through the management interface, ESXi includes a firewall between the management interface and the network.
- Memory hardening: The ESXi kernel, user-mode applications, and executable components, such as drivers and libraries, are located at random, unpredictable memory addresses. Combined with the non-executable memory protections made available by microprocessors, memory hardening provides protection that makes it difficult for malicious code to use memory exploits to take advantage of vulnerabilities.
- Kernel module integrity: Digital signing ensures the integrity and authenticity of modules, drivers, and applications as they are loaded by the VMkernel.
- Trusted Platform Module: Is a hardware element that creates a trusted platform. This element affirms that the boot process and all drivers loaded are genuine.
- UEFI secure boot: The feature is for systems that support UEFI secure boot firmware, which contains a digital certificate to which the VMware infrastructure bundles (VIBs) chain. At boot time, a verifier is started before other processes to check the VIB's chain to the certificate in the firmware.
- ESXi Quick Boot: With this feature, ESXi can reboot without reinitializing the physical server BIOS. Quick Boot reduces remediation time during host patch or host upgrade operations. Quick Boot is activated by default on the supported hardware.

3-7 ESXi Installation Requirements

Ensure that the host meets the minimum hardware configurations supported by ESXi 8.0:

- Supported server platform
- At least two CPU cores
- At least 8 GB of physical RAM, 12 GB for a production environment
- One or more Gigabit or faster Ethernet controllers
- Boot disk of at least 32 GB of persistent storage

For a list of supported server platforms, supported processors and supported network adapter models, see the *VMware Compatibility Guide* at https://www.vmware.com/resources/compatibility.

3-8 Interactive ESXi Installation

An interactive installation is recommended for small deployments of fewer than five hosts.

You boot from the installer and follow the prompts in the installation wizard:

- Start at the Welcome page.
- Accept the EULA.
- Select the disk.
- Select the keyboard layout.
- Enter the root password.
- Start the installation.

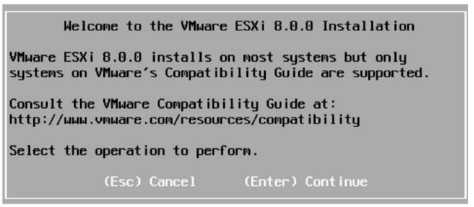

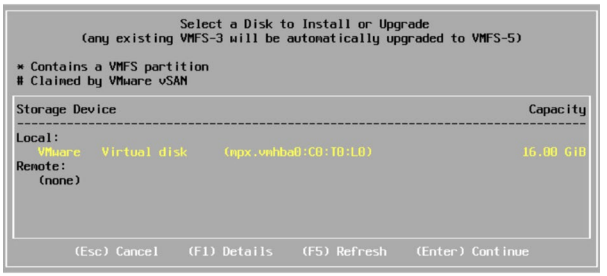

You can boot the ESXi installer from a CD or DVD, bootable USB device, or by PXE booting over the network.

For a list of supported server platforms, supported processors and supported network adapter models, see the *VMware Compatibility Guide* at https://www.vmware.com/resources/compatibility.

3-9 Configuring an ESXi Host

During the ESXi installation, the ESXi host is given a DHCP-assigned IP address.

You use the ESXi host's DCUI to configure certain settings, such as the host's network settings.

The DCUI is a text-based user interface with keyboard-only interaction.

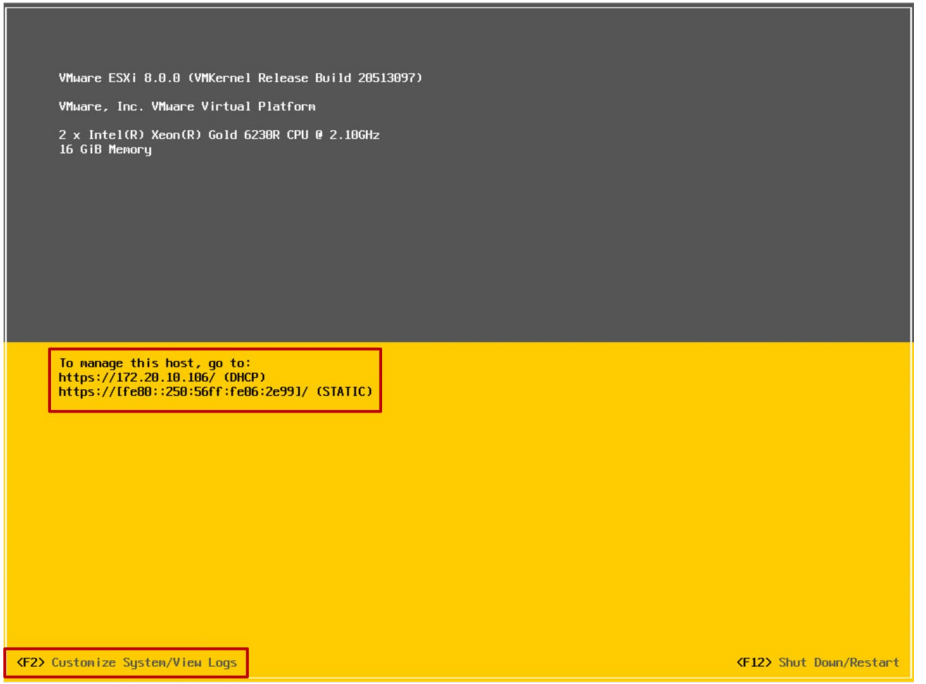

You use the Direct Console User Interface (DCUI) to configure certain settings for ESXi hosts. The DCUI is a low-level configuration and management interface, accessible through the console of the server, that is used primarily for initial basic configuration. You press F2 to start customizing system settings.

3-10 Configuring an ESXi Host: Management Network

You must configure management network settings before your ESXi host is operational. By default, a DHCP-assigned IP address is configured for the ESXi host.

Use the DCUI to configure management network settings:

- Network adapter selection
- VLAN ID
- IPv4 and IPv6 configuration (IP address, subnet mask, default gateway)
- Host name
- DNS servers and suffixes

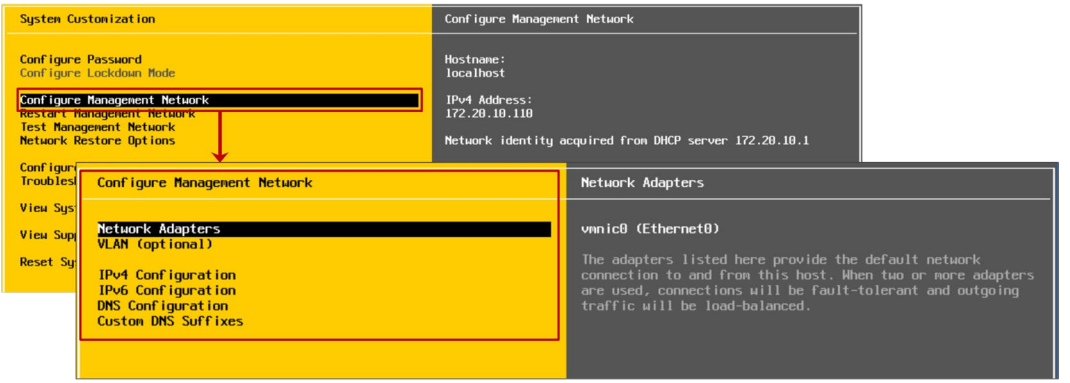

You can perform the following management network configuration tasks from the DCUI:

- Configure VLAN settings
- Configure IPv4 addressing
- Configure IPv6 addressing
- Set custom DNS suffixes
- Restart the management network (without rebooting the system)
- Test the management network (using ping and DNS requests)
- Restore the original network configuration (useful if you misconfigure something)

3-11 Configuring an ESXi Host: Root Access

Administrators use the DCUI to configure the root access settings:

- Change the root password (complex passwords only)
- Activate or deactivate the lockdown mode:
  - Limits the management of the host to vCenter
  - Can be configured only for hosts managed by a vCenter instance

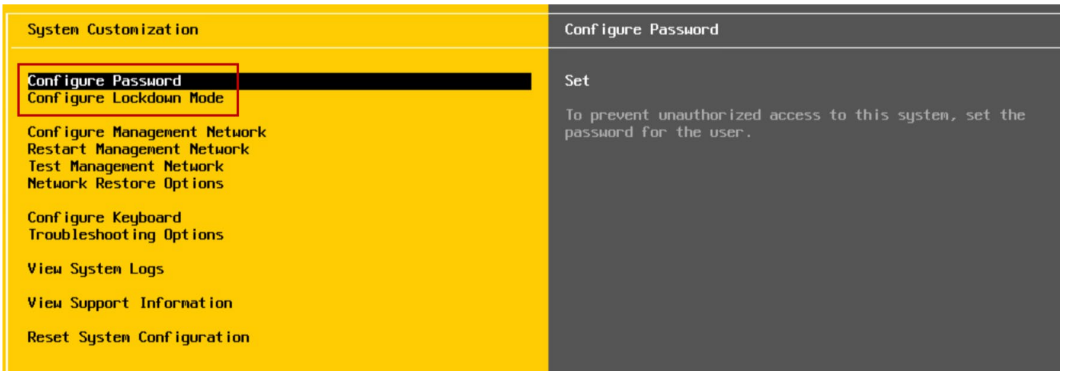

The administrative username for the ESXi host is root. The root password must be configured during the ESXi installation process, but can be changed from the DCUI.

3-12 Configuring an ESXi Host: Other Settings

Using the DCUI, you can configure the keyboard layout, activate troubleshooting services, view support information, and view system logs.

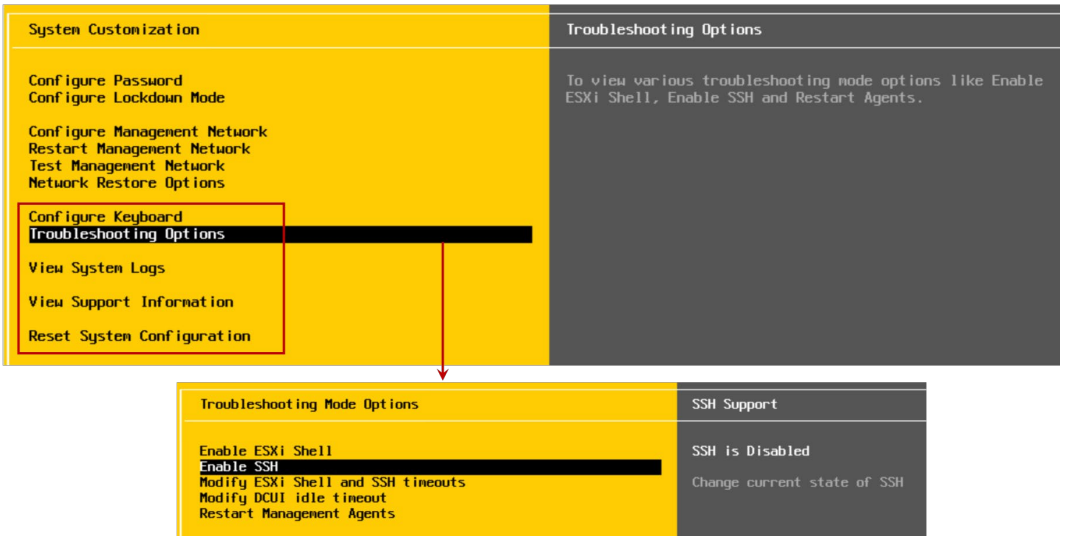

From the DCUI, you can change the keyboard layout, view support information, such as the host's license serial number, and view system logs. The default keyboard layout is U.S. English.

You can use the troubleshooting options, which are deactivated by default, to activate or deactivate troubleshooting services:

- vSphere ESXi Shell: For troubleshooting problems locally
- SSH: For troubleshooting problems remotely by using an SSH client, for example, PuTTY

The best practice is to keep the troubleshooting services deactivated until they are necessary. For example, when you work with VMware technical support to resolve a problem.

By selecting the Reset System Configuration option, you can reset the system configuration to its software defaults and remove custom extensions or packages that you added to the host. This option also deletes the root password, which might allow unauthorized access to the system.

3-13 Time Synchronization for the ESXi Host

To ensure precise timekeeping and synchronization between the ESXi host and the other components of the vSphere network, you can synchronize an ESXi host's clock to a time reference.

Time synchronization is important:

- For accurate performance graphs
- For accurate time stamps in log messages
- So that virtual machines have a source to synchronize

Synchronizing an ESXi host’s time includes the following benefits:

- Performance data can be displayed and interpreted properly.
- Accurate time stamps appear in log messages, which make audit logs and troubleshooting meaningful.
- VMs can synchronize their time with the ESXi host. Time synchronization benefits applications, for example database applications running on VMs.

3-14 Methods for Synchronizing Time

You can use the time synchronization options:

- Manual configuration
- NTP, Network Time Protocol, providing millisecond timing accuracy
- PTP, Precision Time Protocol, providing microsecond timing accuracy

You can configure NTP or PTP using VMware Host Client or the vSphere Client.

The NTP and PTP services cannot run simultaneously.

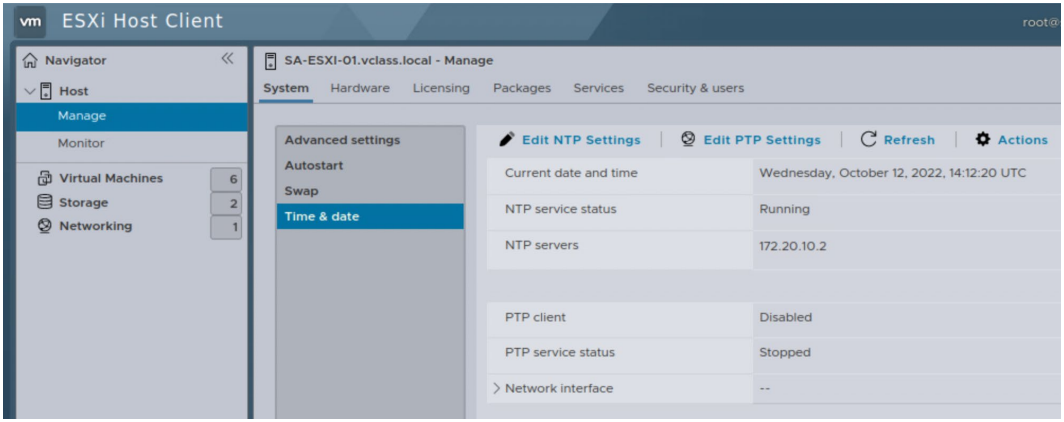

Configuring time and date in VMware Host Client

For information about using NTP and PTP on ESXi hosts, see the Synchronizing Clocks on the vSphere Network section in *vCenter Server and Host Management* at https://docs.vmware.com/en/VMware-vSphere/index.html.

3-15 Configuring NTP

An ESXi host can be configured as an NTP client. It can synchronize time with an NTP server on the Internet or your corporate NTP server.

NTP client uses UDP over port 123.

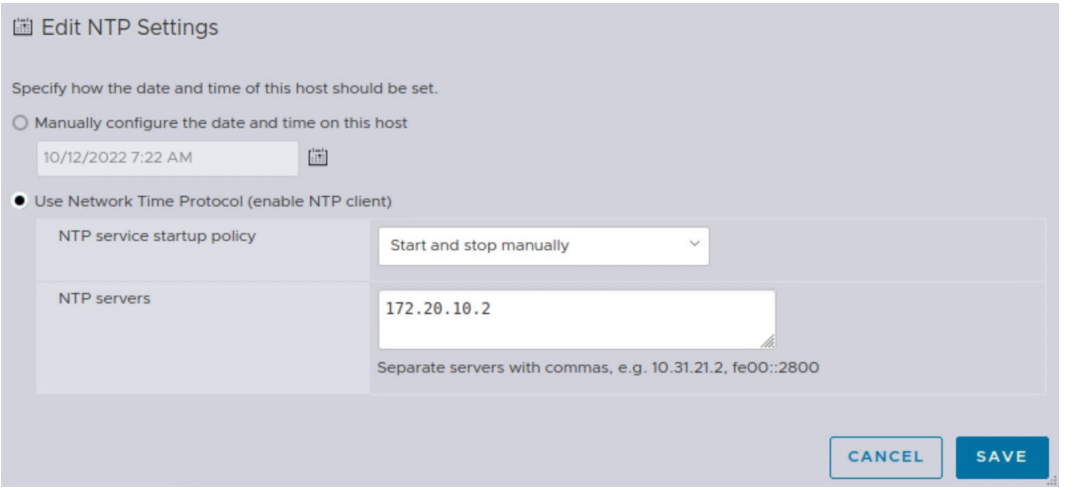

Configuring NTP using VMware Host Client

NTP is a client-server protocol. When you configure the ESXi host to be an NTP client, the host synchronizes its time with an NTP server, which is a server on the Internet or your corporate NTP server.

For general information about NTP, see http://www.ntp.org.

For details about configuring NTP on an ESXi host, see the *vCenter Server and Host Management*, section Use NTP Servers for Time and Date Synchronization of a Host at https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vcenterhost.doc/GUID-8756D419-A878-4AE0-9183-C6D5A91A8FB1.html.

3-16 Configuring PTP

PTP provisions hardware-based timestamping for the virtual machines and the hosts within a network.

PTP client uses UDP over ports 319 and 320.

You can use NTP as a fallback if the PTP service does not work.

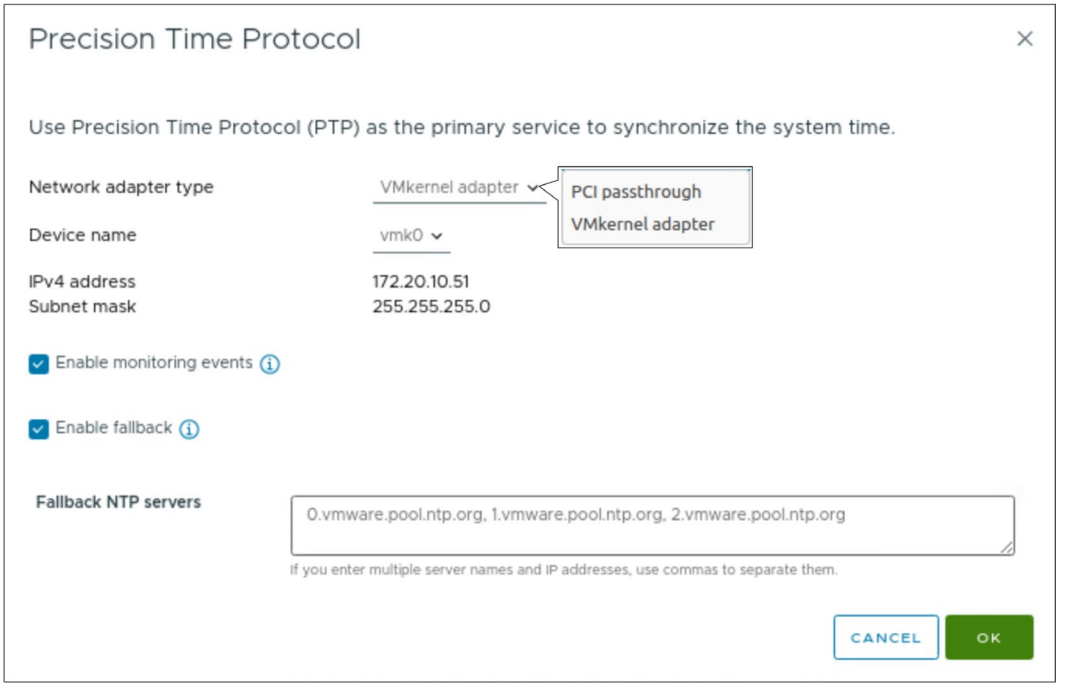

Configuring PTP using the vSphere Client

PTP provides highly accurate time synchronization and allows both software-based and hardware-based timestamping on ESXi hosts:

- For configuring hardware timestamping, select **PCI passthrough** as the network adapter type.
- For configuring software timestamping, select **VMkernel Adapter** as the network adapter type.

For details about configuring PTP on an ESXi host, see the Use PTP for Time and Date Synchronization of a host section in *vCenter Server and Host Management* at https://docs.vmware.com/en/VMware-vSphere/index.html.

3-17 Controlling Remote Access to an ESXi Host

You can use the vSphere Client to customize the essential security settings that control remote access to an ESXi host:

- The ESXi firewall is activated by default.

  The firewall blocks incoming and outgoing traffic, except for the traffic activated in the host's firewall settings.

- Services, such as the NTP client and the SSH client, are managed by users with administrator privileges.

- Lockdown mode prevents remote users from logging in to the host directly. The host is accessible only through the DCUI or vCenter.

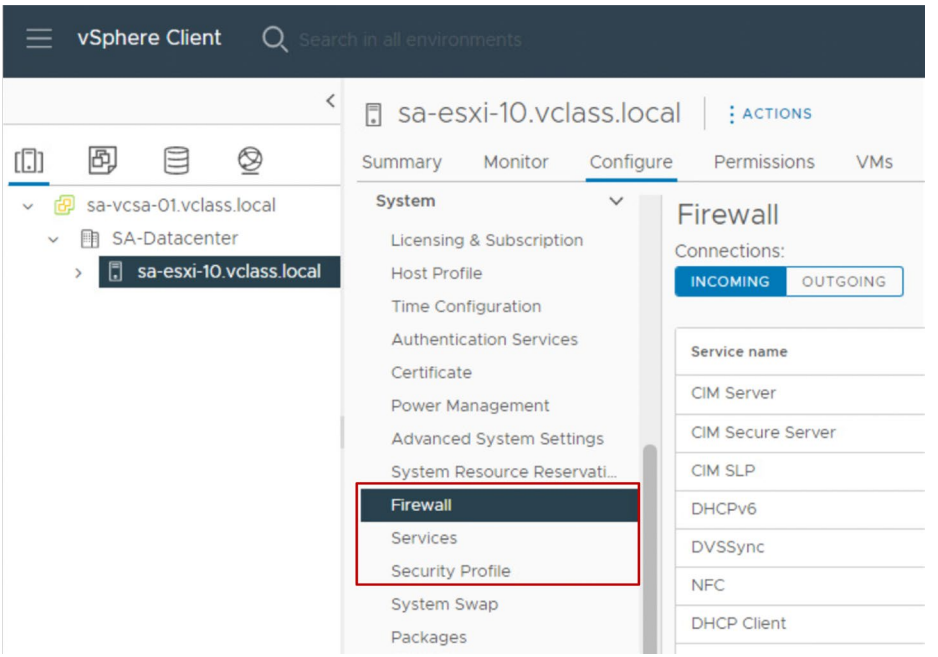

Every ESXi host includes a firewall. On ESXi hosts, remote clients are prevented from accessing services on the host. Similarly, local clients are prevented from accessing services on the remote hosts. To verify the integrity of the host, few ports are open by default. To enable or prevent access to certain services or clients, you must modify the properties of the firewall.

You can configure the firewall settings for incoming and outgoing connections for a service or a management agent. For some services, you can manage the service details. For example, you

can use START, STOP, or RESTART to change the status of a service temporarily. Alternatively, you can change the startup policy so that the service starts with the host or with port use. For some services, you can specify IP addresses, ranges of IP addresses, or entire subnets from which connections are allowed.

3-18 Managing User Accounts: Best Practices

When assigning user accounts to access ESXi hosts or vCenter systems, you must follow these security guidelines:

- Strictly control root access to the ESXi hosts.
- Create strong root account passwords that have at least eight characters. Use special characters, case changes, and numbers. Change passwords periodically.
- Manage the ESXi hosts centrally through vCenter Server by using the vSphere Client.
- Minimize the use of local users on the ESXi hosts:
    - Add the ESXi hosts to a domain and add the relevant administrator users to the ESX Admins domain group. Users in the domain group have root privileges on the ESXi hosts.

On an ESXi host, the root user account is the most powerful user account on the system. The root user can access all files and all commands. Securing this account is the most important step that you can take to secure an ESXi host.

When possible, use the vSphere Client to log in to the vCenter Server system and manage your ESXi hosts. In some circumstances, for example, when the vCenter Server system is down, you use VMware Host Client to connect directly to the ESXi host. Although you can log in to your ESXi host through the vSphere ESXi Shell, these access methods must be reserved for troubleshooting or configuration that cannot be accomplished by using VMware Host Client.

If a host must be managed directly, avoid creating local users on the host. If possible, join the host to an Active Directory domain and log in with domain credentials instead.

3-19 Demonstration: Installing and Configuring ESXi Hosts

Your instructor will run a demonstration.

3-20 Lab 2: Configuring an ESXi Host

Use VMware Host Client to configure an ESXi host:

1. Add an ESXi Host to an LDAP Server
2. Log In to the ESXi Host as an LDAP User
3. Activate the SSH and vSphere ESXi Shell Services
4. Configure the ESXi Host as an NTP Client

3-21 Review of Learner Objectives

- Describe the ESXi host architecture
- Navigate the Direct Console User Interface (DCUI) to configure an ESXi host
- Recognize the user account best practices
- Install an ESXi host
- Configure the ESXi host settings

3-22 Key Points

- The Direct Console User Interface (DCUI) allows you to configure certain settings for ESXi hosts.
- Securing the root user account is very important to secure an ESXi host, as the user is the host's most powerful user.
- NTP provides millisecond timing accuracy and PTP provides microsecond timing accuracy for ESXi hosts.

Questions?

Module 4

Deploying and Configuring vCenter

4-2 Importance

vCenter helps you centrally manage multiple ESXi hosts and their virtual machines. If you do not properly deploy, configure, and manage vCenter, your environment might experience reduced administrative manageability of the ESXi hosts and virtual machines.

4-3 Module Lessons

1. Centralized Management with vCenter
2. Deploying vCenter Server Appliance
3. vSphere Licenses
4. Managing vCenter Inventory
5. vCenter Roles and Permissions
6. Monitoring vCenter Events

4-4 Lesson 1: Centralized Management with vCenter

4-5 Learner Objectives

- Describe the vCenter architecture
- Recognize ESXi hosts communication with vCenter
- Identify vCenter services

4-6 About the vCenter Management Platform

vCenter acts as a central administration point for ESXi hosts and virtual machines. The ESXi hosts and virtual machines connected in a network:

- Directs the actions of VMs and hosts
- Runs on a Linux-based appliance

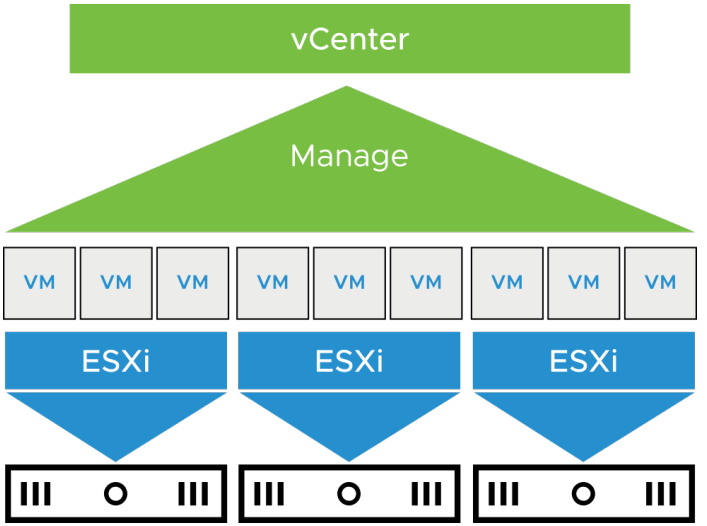

With vCenter, you can pool and manage the resources of multiple hosts. vCenter provides advanced features, such as vSphere DRS, vSphere HA, vSphere Fault Tolerance, vSphere vMotion, and vSphere Storage vMotion.

vCenter is deployed as a virtual appliance. You deploy vCenter Server Appliance on an ESXi host in your infrastructure. vCenter Server Appliance is a preconfigured Linux-based virtual machine, which is optimized for running vCenter and the vCenter components.

4-7 About vCenter Server Appliance

vCenter Server Appliance is a prepackaged Linux-based VM, optimized for running vCenter and associated services.

The vCenter Server Appliance package contains the following software:

- Photon OS
- PostgreSQL database
- vCenter services

During deployment, you can select the vCenter Server Appliance size for your vSphere environment and the storage size for your database requirements.

vCenter consists of a collection of services that run in vCenter Server Appliance. vCenter acts as a central administration point for ESXi hosts that are connected to a network.

4-8 vCenter Services

vCenter services include:

- vCenter Server
- vSphere Client
- License service
- Content Library
- vSphere Lifecycle Manager

When you deploy vCenter Server Appliance, all these services are included.

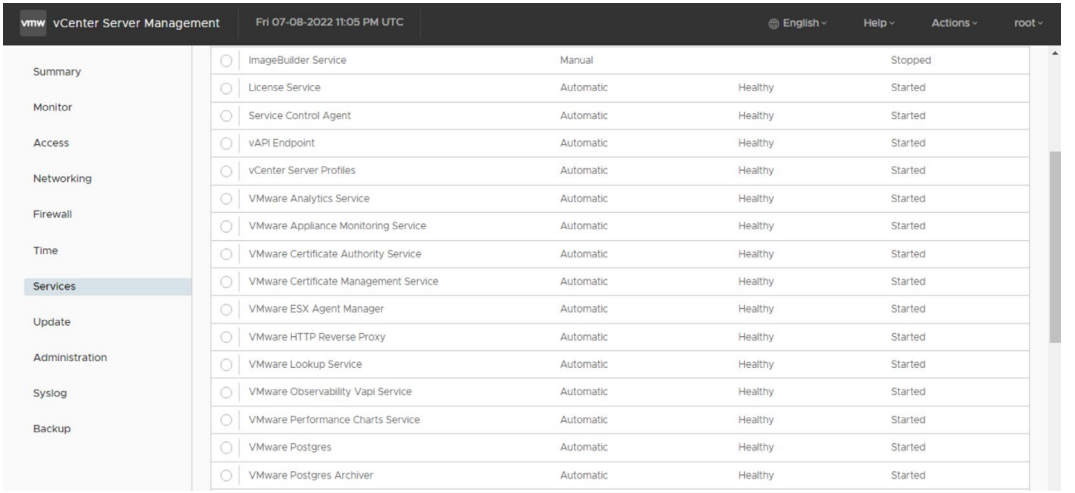

All vCenter services are installed on a single VM.

4-9 vCenter Architecture

vSphere Client, vCenter database, and managed hosts supports vCenter.

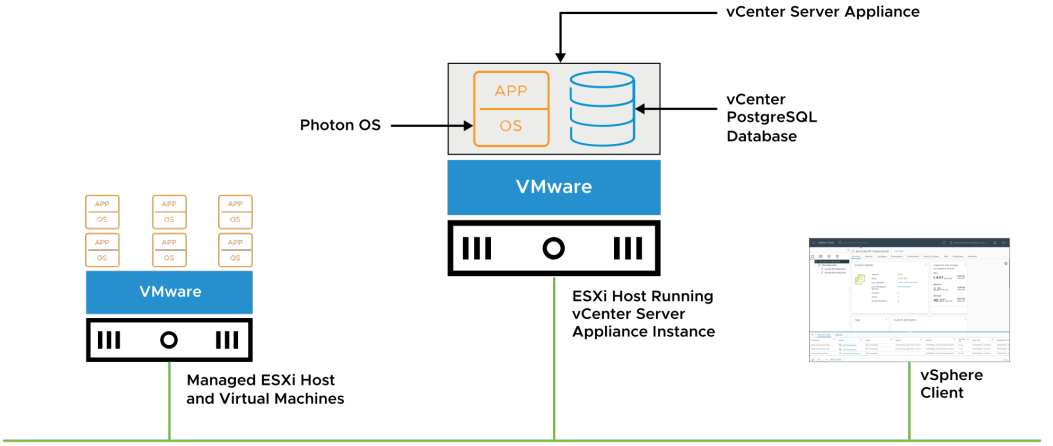

The vCenter architecture relies on the following components:

- vSphere Client: You use this client to connect to vCenter and manage your ESXi hosts centrally. When vCenter manages an ESXi host, you should always use vCenter and the vSphere Client to manage that host.
- vCenter database: The vCenter database is a critical component. The database stores inventory items, security roles, performance data, and other critical information for vCenter.
- Managed hosts: You can use vCenter to manage ESXi hosts and the VMs running on ESXi hosts.

4-10 About vCenter Single Sign-On

vCenter Single Sign-On allows vSphere components to communicate with each other through a secure token mechanism.

vCenter Single Sign-On can authenticate users using built-in or external identity providers.

Built-in identity providers:

- By default, vCenter uses the vsphere.local domain as the identity source.
- You can configure vCenter to use Active Directory as the identity source using LDAP, LDAPS, OpenLDAP, or OpenLDAPS.

External identity provider using federated authentication:

- vSphere supports Active Directory Federation Services (AD FS).

Although you can still configure Integrated Windows Authentication (IWA), VMware recommends using Active Directory over LDAP or Federated Identity with AD FS for authentication for vCenter Server and ESXi. For more details, see VMware knowledge base article 78506 at https://kb.vmware.com/kb/78506.

For details about configuring vCenter Single Sign-On and identity providers, see *vSphere Authentication* at https://docs.vmware.com/en/VMware-vSphere/index.html.

4-11 vCenter Single Sign-On with Built-In Identity Provider

The following is the user login flow when vCenter acts as the identity provider:

1. User logs in to the vSphere Client.
2. vCenter Single Sign-On authenticates credentials against a directory service (for example, Active Directory).
3. A SAML token is sent back to the user's browser.
4. The SAML token is sent to vCenter, and the user is granted access.

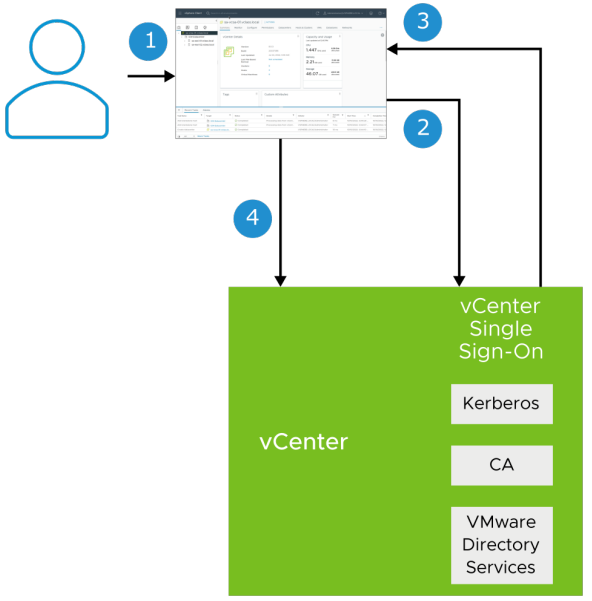

For more details about user login flow, see *vSphere Authentication* at https://docs.vmware.com/en/VMware-vSphere/index.html.

4-12 About Enhanced Linked Mode

With Enhanced linked mode, you can log in to the vSphere Client and manage the inventories of all the vCenter instances in the group:

- You can link up to 15 vCenter instances in one vCenter Single Sign-On domain.
- You can create an enhanced linked mode group during the deployment of vCenter Server Appliance.

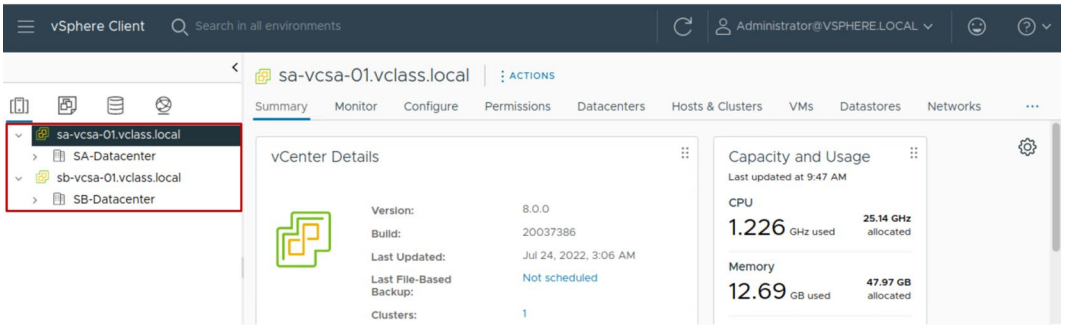

Enhanced linked mode provides the following features:

- You can log in to all linked vCenter instances simultaneously with a single username and password.
- You can view and search the inventories of all linked vCenter instances in the vSphere Client.
- You can replicate roles, permission, licenses, tags, and policies (such as storage policies) across linked vCenter instances.

To join vCenter instances in Enhanced Linked Mode, connect the vCenter instances to the same vCenter Single Sign-On domain.

You can create a vCenter enhanced linked mode group during the deployment of vCenter Server Appliance. You can also join a vCenter enhanced linked mode group by moving, or repointing, a vCenter instance from one vSphere domain to another existing domain. For more information on repointing vCenter instances, see *vCenter Server Installation and Setup* at https://docs.vmware.com/en/VMware-vSphere/index.html.

Enhanced linked mode requires the vCenter Standard licensing level. Enhanced linked mode is not supported with vCenter Foundation or vCenter for Essentials.

4-13 ESXi and vCenter Communication

The vSphere Client is the primary method to manage ESXi hosts. vSphere Client communicates directly with vCenter.

If vCenter is not available, you use VMware Host Client to communicate directly with the ESXi host.

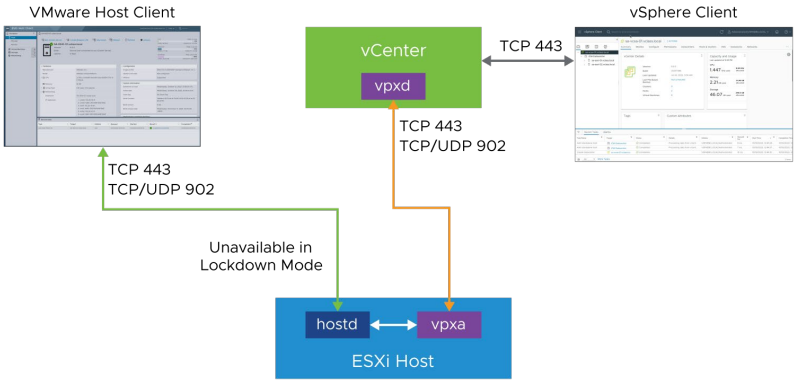

vCenter provides direct access to the ESXi host through a vCenter agent called virtual provisioning X agent (vpxa). The vpxa process is automatically installed on the host and started when the host is added to the vCenter inventory. The vCenter service (vpxd) communicates with the ESXi host daemon (hostd) through the vCenter agent (vpxa).

Clients that communicate directly with the host, and bypass vCenter, converse with hostd. The hostd process runs directly on the ESXi host and manages most of the operations on the ESXi host. The hostd process is aware of all VMs that are registered on the ESXi host, the storage volumes visible to the ESXi host, and the status of all VMs.

Most commands or operations come from vCenter through vpxa. Examples include creating, migrating, and powering on virtual machines. Acting as an intermediary between the vpxd process, which runs on vCenter, and the hostd process, vpxa relays the tasks to perform on the host.

When you are logged in to the vCenter system through the vSphere Client, vCenter passes commands to the ESXi host through the vpxa.

The vCenter database is also updated. If you use VMware Host Client to communicate directly with an ESXi host, communications go directly to the hostd process and the vCenter database is not updated.

4-14 vCenter Scalability

Metric	vCenter 8.0
Hosts per vCenter instance	2,500
Powered-on VMs per vCenter instance	40,000
Registered VMs per vCenter instance	45,000
Hosts per cluster	96
VMs per cluster	8,000

You can scale vCenter to support large, enterprise environments. For the recommended configuration limits, see *VMware Configuration Maximums* at https://configmax.vmware.com.

4-15 Review of Learner Objectives

- Describe the vCenter architecture
- Recognize ESXi hosts communication with vCenter
- Identify vCenter services

4-16 Lesson 2: Deploying vCenter Server Appliance

4-17 Learner Objectives

- Deploy vCenter Server Appliance into an infrastructure
- Configure vCenter settings

4-18 Preparing for vCenter Server Appliance Deployment

Before deploying vCenter Server Appliance, you must complete several tasks:

- Verify that all vCenter Server Appliance system requirements are met.
- Get the fully qualified domain name (FQDN) or the static IP of the host machine on which you install vCenter Server Appliance.
- Get FQDN and IP address to assign to vCenter Server Appliance.
- Ensure that date and time on all VMs in the vSphere network are synchronized.

For more information on system requirements for vCenter Server Appliance, see *vCenter Server Installation and Setup* at https://docs.vmware.com/en/VMware-vSphere/index.html.

4-19 vCenter Server Appliance Native GUI Installer

The vCenter Server Appliance Native GUI installer has several features:

- With the GUI installer, you can perform an interactive deployment of vCenter Server Appliance.
- The GUI installer is a native application for Windows, Linux, and macOS.
- The GUI installer performs validations and prechecks during the deployment.

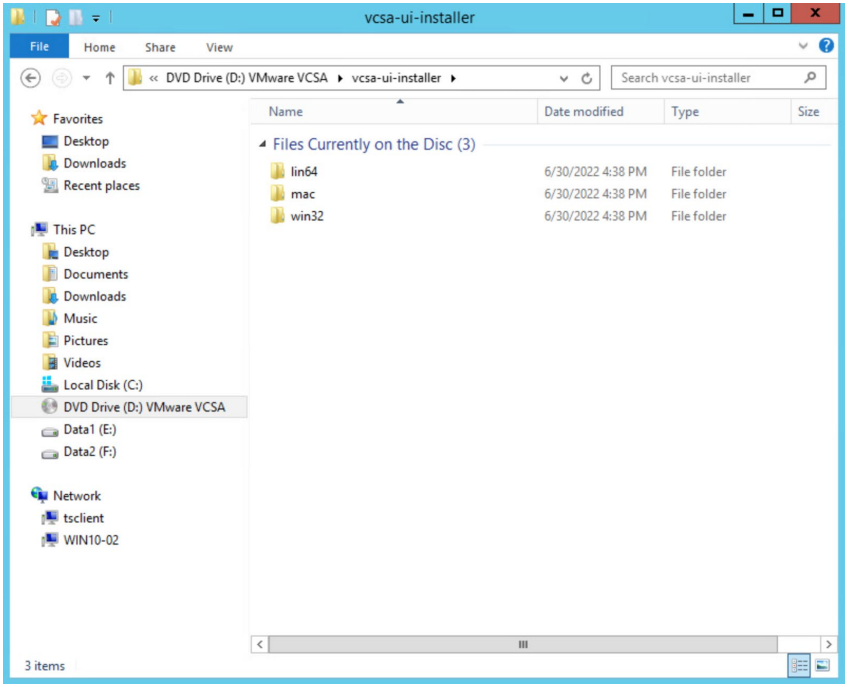

The GUI installer performs validations and prechecks during the deployment phase to ensure that no mistakes are made and that a compatible environment is created.

4-20 vCenter Server Appliance Installation

The vCenter Server Appliance installation is a two-stage process:

- Stage 1: Deployment of OVF
- Stage 2: Configuration

The deployment can be fully automated by using JSON templates with the CLI installer on Windows, Linux, or macOS.

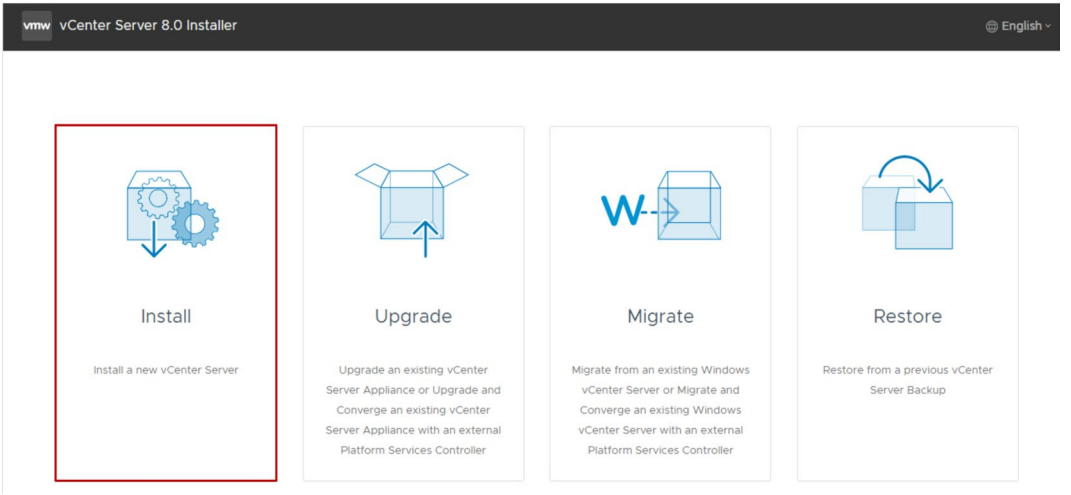

The **Install** option installs a new vCenter Server Appliance.

The **Upgrade** option upgrades an existing vCenter Server Appliance instance, or upgrades and converges an existing vCenter Server Appliance instance with an external Platform Services Controller.

The **Migrate** option migrates from an existing Windows vCenter instance, or migrates and converges an existing Windows vCenter instance with an external Platform Services Controller.

The **Restore** option restores from a previous vCenter Server Appliance backup.

4-21 vCenter Server Appliance Installation: Stage 1

Stage 1 begins with the UI phase:

1. Accept the EULA.
2. Connect to the target ESXi host or vCenter system.
3. Define the vCenter Server Appliance name and root password.
4. Select compute size, storage size, and datastore location (thin disk).
5. Define networking settings.

Stage 1 continues with the deployment phase:

6. OVF is deployed to the ESXi host.
7. Disks and networking are configured.

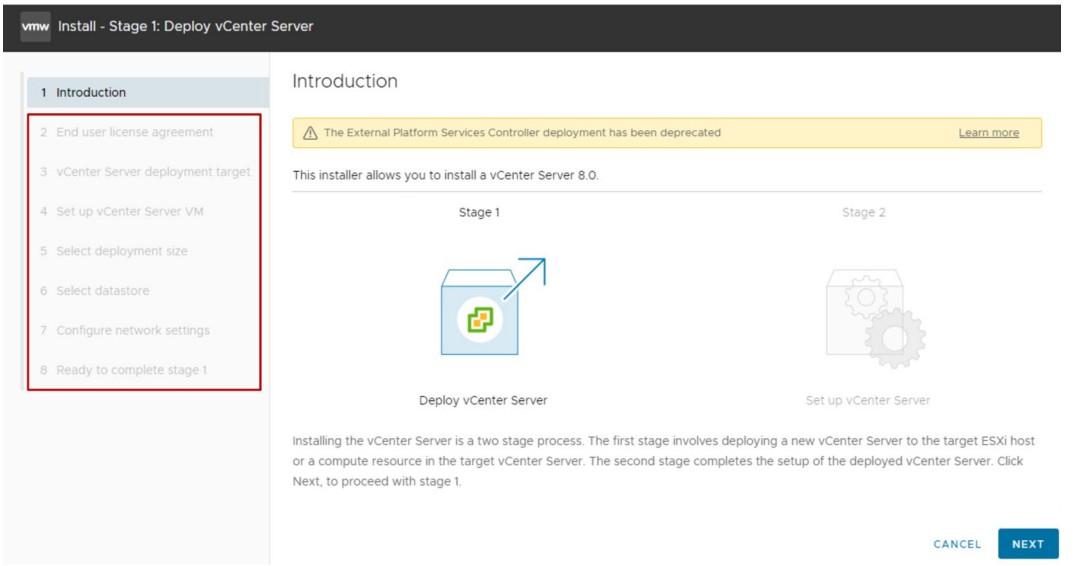

4-22 vCenter Server Appliance Installation: Stage 2

Stage 2 is the configuration phase:

- Configure time synchronization mode and SSH access.
- Create a vCenter Single Sign-On domain or join an existing SSO domain.
- Join the Customer Experience Improvement Program (CEIP).

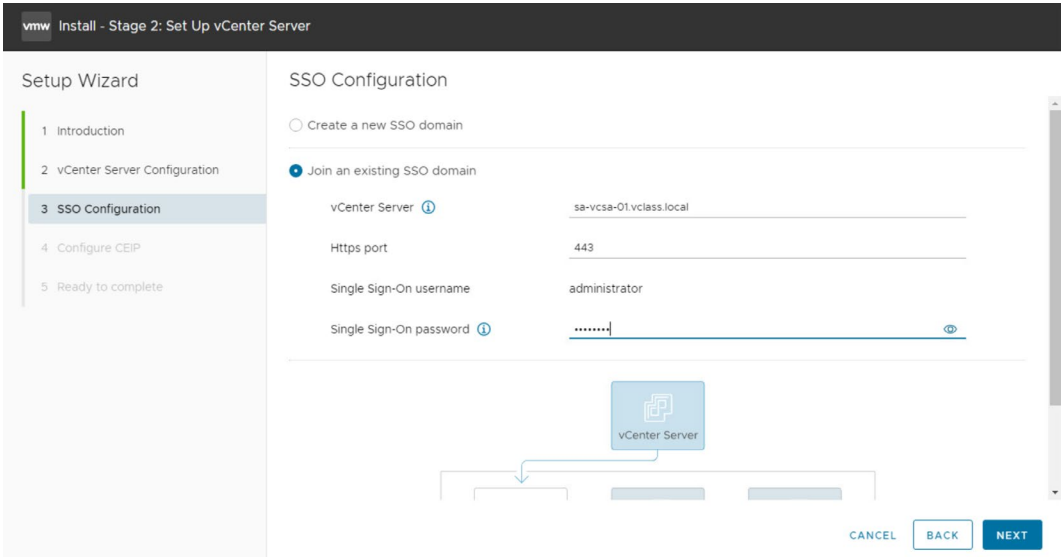

In stage 2, you configure whether to use the ESXi host or NTP servers as the time synchronization source. You can also activate SSH access. SSH access is deactivated by default.

4-23 Getting Started with vCenter

After you deploy vCenter Server Appliance, use the vSphere Client to log in and manage your vCenter inventory: https://<vCenter_FQDN_or_IP_address>/ui.

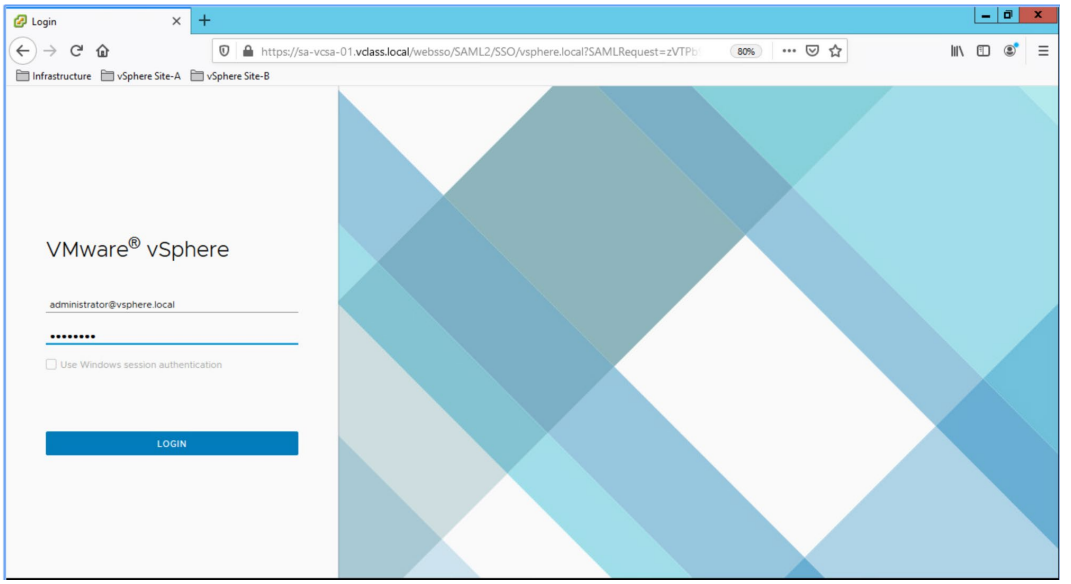

4-24 Configuring vCenter Using the vSphere Client

Using the vSphere Client, you can configure vCenter, including settings such as licensing, statistics collection, and logging.

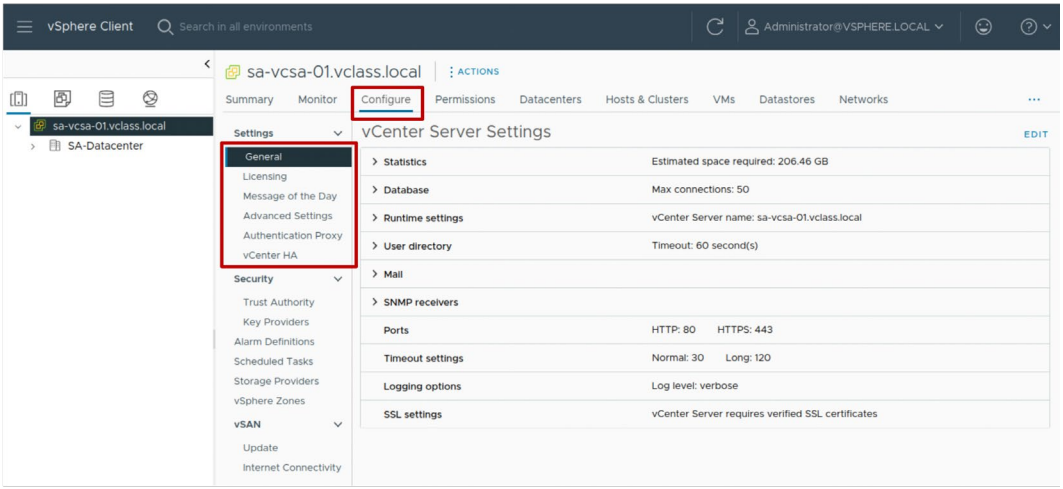

To access the vCenter system settings by using the vSphere Client, select the vCenter system in the navigation pane, click the **Configure** tab, and expand **Settings**.

4-25 vCenter Management Interface

Using the vCenter Management Interface, you can configure and monitor your vCenter instance.

Tasks include:

- Monitoring resource use by the appliance
- Backing up the appliance
- Monitoring vCenter services
- Adding additional network adapters

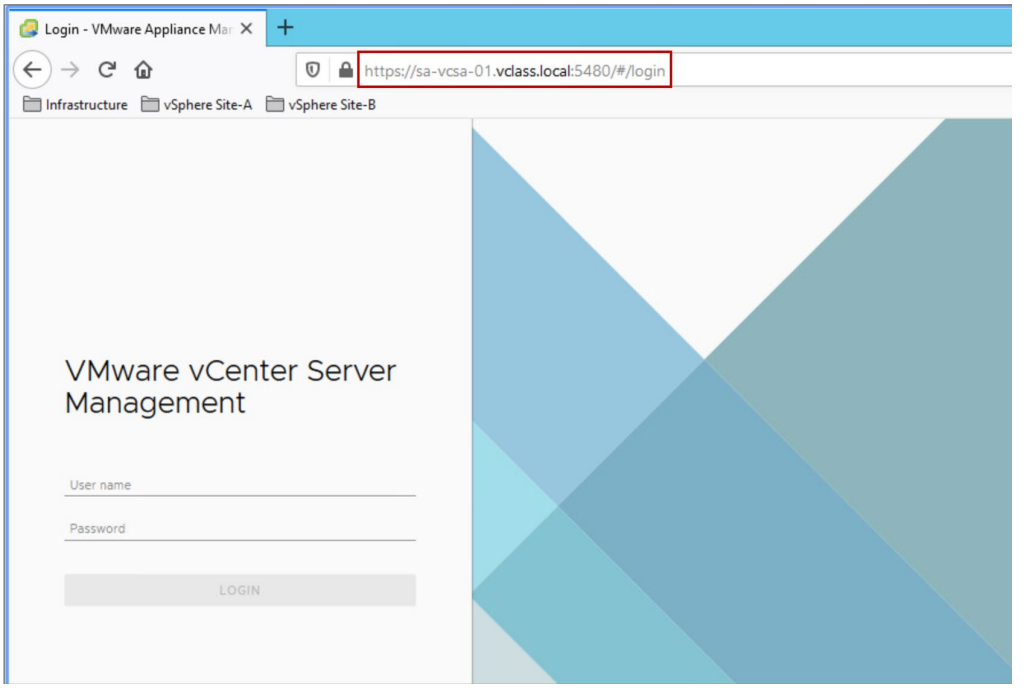

The vCenter Management Interface is an HTML client designed to configure and monitor a vCenter instance.

The vCenter Management Interface connects directly to port 5480. Use the URL https://<FQDN_or_IP_address>:5480.

4-26 Multi-homing the vCenter Server Appliance

With vCenter Server Appliance multi-homing, you can configure multiple NICs to manage network traffic.

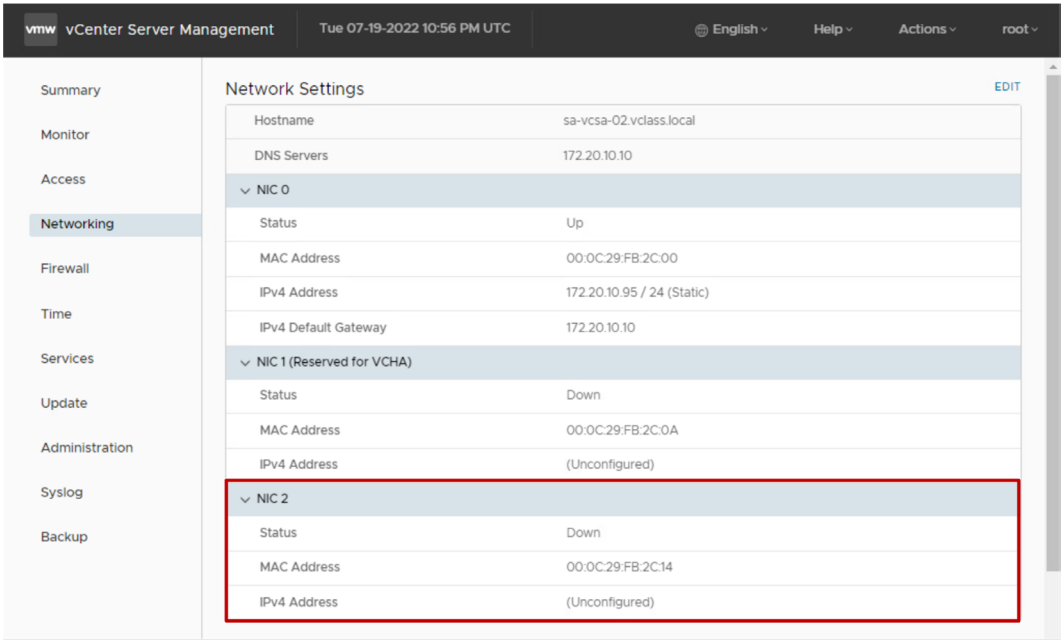

A maximum of four NICs are supported for multi-homing. All four multi-homing NIC configurations are preserved during upgrade, backup, and restore processes.

4-27 Demonstration: Deploying vCenter Server Appliance

Your instructor will run a demonstration.

4-28 Review of Learner Objectives

- Deploy vCenter Server Appliance into an infrastructure
- Configure vCenter settings

4-29 Lesson 3: vSphere Licensing

4-30 Learner Objectives

- View licensed features for vCenter or an ESXi host
- Add license keys to vCenter

4-31 About vSphere Licenses

VMware provides a number of vSphere products to suit your needs.

vSphere Essential Kit	**vSphere Essential Plus Kit**	**vSphere Standard**	**vSphere Enterprise Plus**
For small businesses (up to three hosts with up to two CPUs each)	For small businesses (up to three hosts with up to two CPUs each)	Entry-level solution for basic server consolidation	Full range of features for transforming your data center into a simplified cloud infrastructure
vCenter and ESXi	vCenter and ESXi	vCenter and ESXi	vCenter and ESXi
	vSphere vMotion, vSphere Storage vMotion, vSphere HA, vSphere Data Protection, vSphere Replication	vSphere vMotion, vSphere Storage vMotion, vSphere HA, vSphere Replication	vSphere vMotion, vSphere Storage vMotion, vSphere HA, vSphere Trust Authority, VM encryption, vSphere Replication

For details, go to the vSphere product on https://store.vmware.com.

4-32 vSphere Licensing Overview

Licensing vSphere components is a two-step process:

1. Add a license to the vCenter License Service
2. Assign the license to the ESXi hosts, vCenter instances, and other vSphere components

4-33 vSphere License Service

The vSphere License Service runs on vCenter.

The vSphere License Service performs the following functions:

- Provides centralized license management
- Provides an inventory of vSphere licenses
- Manages the license assignments for products that integrate with vSphere, such as Site Recovery Manager.

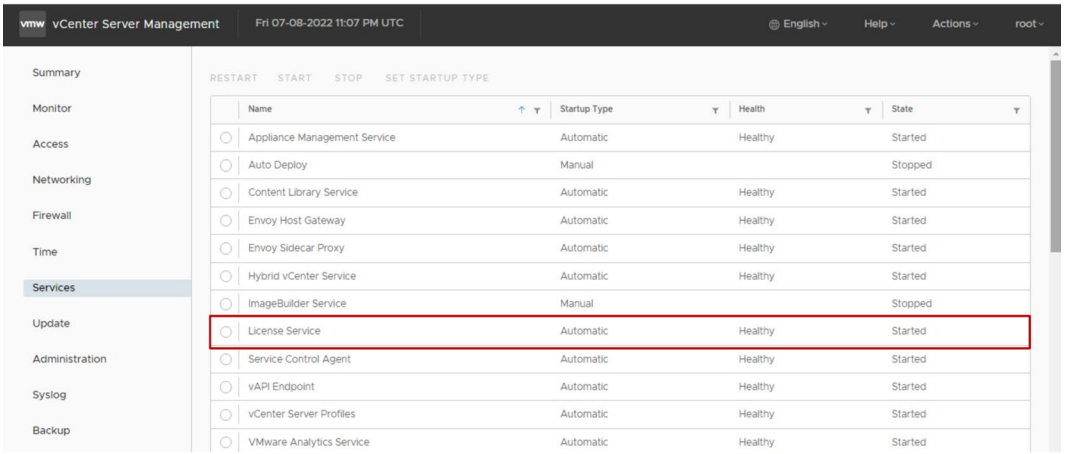

The vSphere License Service manages the license assignments for ESXi hosts, vCenter systems, and clusters with vSAN activated.

You can monitor the health and status of the vSphere License Service by using the vCenter Management interface.

4-34 Adding License Keys to vCenter

You must assign a license to vCenter before its 60-day evaluation period expires.

In the vSphere Client from the main menu, select **Administration** > **Licenses** to open the Licenses pane.

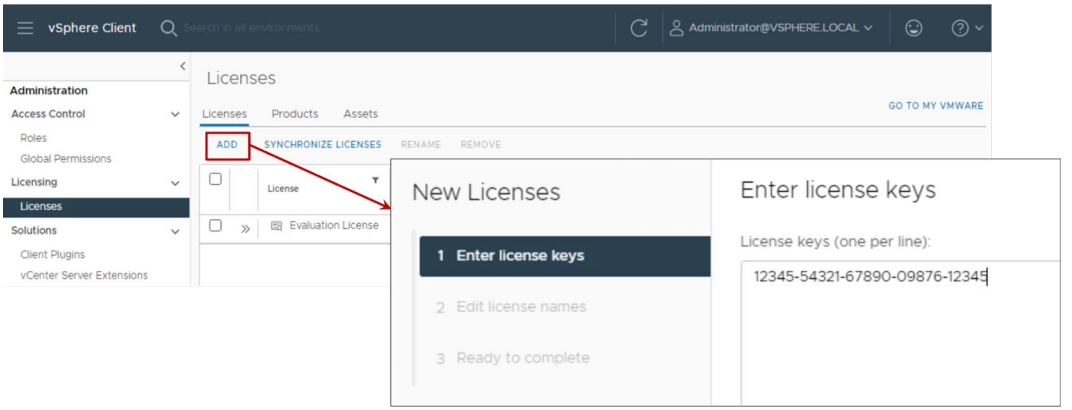

In vSphere, license reporting and management are centralized. All product and feature licenses are encapsulated in 25-character license keys that you can manage and monitor from vCenter.

You can view license information by product, license key, or asset:

- Product: A license to use a vSphere software component or feature, for example, vCenter or vSphere Enterprise Plus.
- License key: The serial number that corresponds to a product.
- Asset: A component that has been assigned a product license. For an asset to run certain software legally, the asset must be licensed.

4-35 Assigning a License to a vSphere Component

You can assign a license to an asset, such as vCenter.

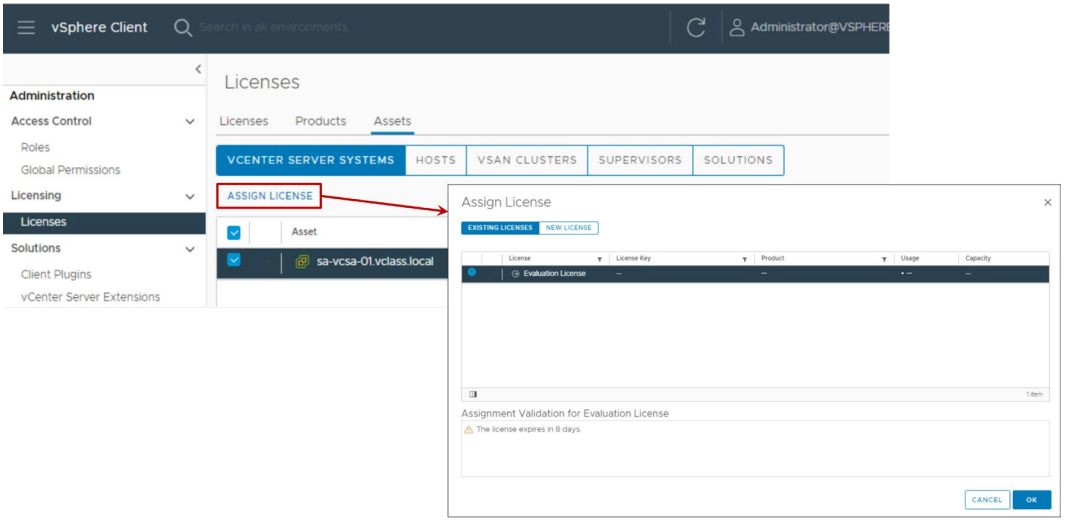

4-36 Viewing Licensed Features

You manage licenses using the License pane in vCenter's **Configure** tab. This pane shows the type of license and available features.

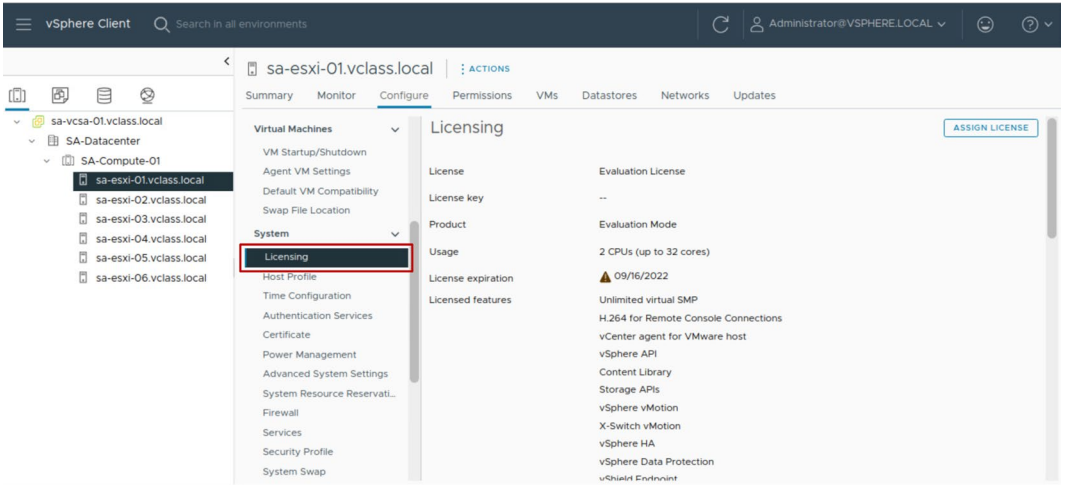

Before purchasing and activating licenses for ESXi and vCenter, you can install the software and run it in evaluation mode. Evaluation mode is intended for demonstrating the software or evaluating its features. During the evaluation period, the software is fully operational.

The evaluation period is 60 days from the time of installation. During this period, the software notifies you of the time remaining until expiration. The 60-day evaluation period cannot be paused or restarted. After the evaluation period expires, you can no longer perform some operations in vCenter and ESXi. For example, you cannot power on or reset your virtual machines. In addition, all hosts are disconnected from the vCenter system. To continue to have full use of ESXi and vCenter operations, you must acquire, install and assign license keys.

4-37 Lab 3: Adding vSphere Licenses

Use the vSphere Client to add vSphere licenses to vCenter and assign a license to vCenter:

1. Add vSphere Licenses to vCenter
2. Assign a License to the vCenter Instance

4-38 Review of Learner Objectives

- View licensed features for vCenter or an ESXi host
- Add license keys to vCenter

4-39 Lesson 4: Managing vCenter Inventory

4-40 Learner Objectives

- Use the vSphere Client to manage the vCenter inventory
- Create and organize vCenter inventory objects
- Add data center and organizational objects to vCenter
- Add ESXi hosts to the inventory
- Create custom inventory tags for inventory objects

4-41 vSphere Client Main Menu

From the vSphere Client main menu, you can manage your vCenter system inventory, manage your infrastructure environment, and complete system administration tasks.

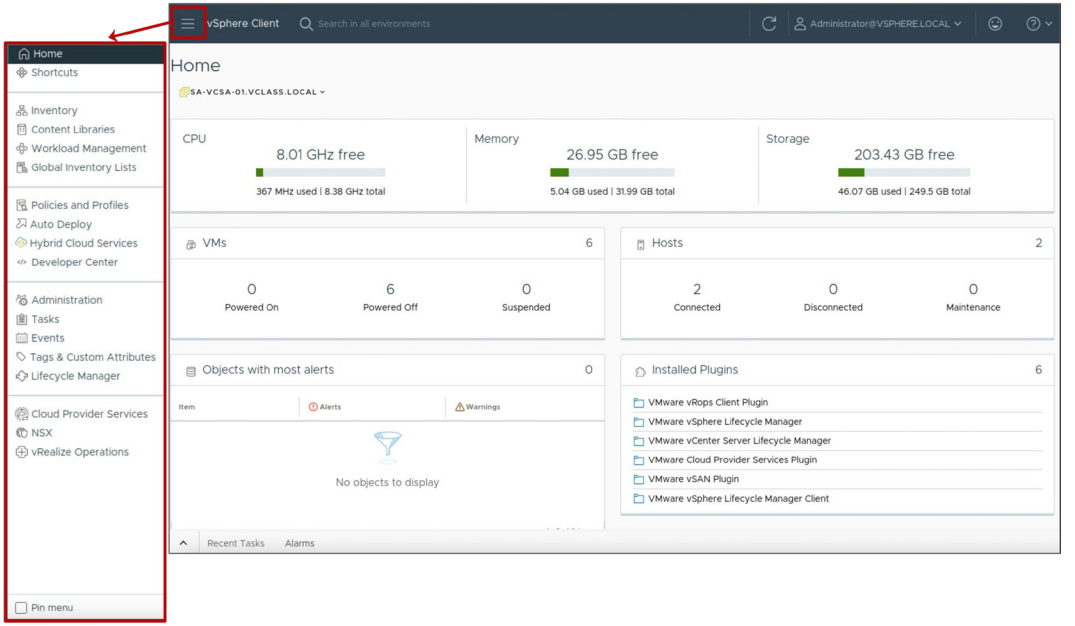

The vSphere Client main menu is indicated by a three-lined icon, located in the upper left corner of the vSphere Client window.

4-42 Navigating the Inventory

You can use the navigation pane to browse and select objects in the vCenter inventory.

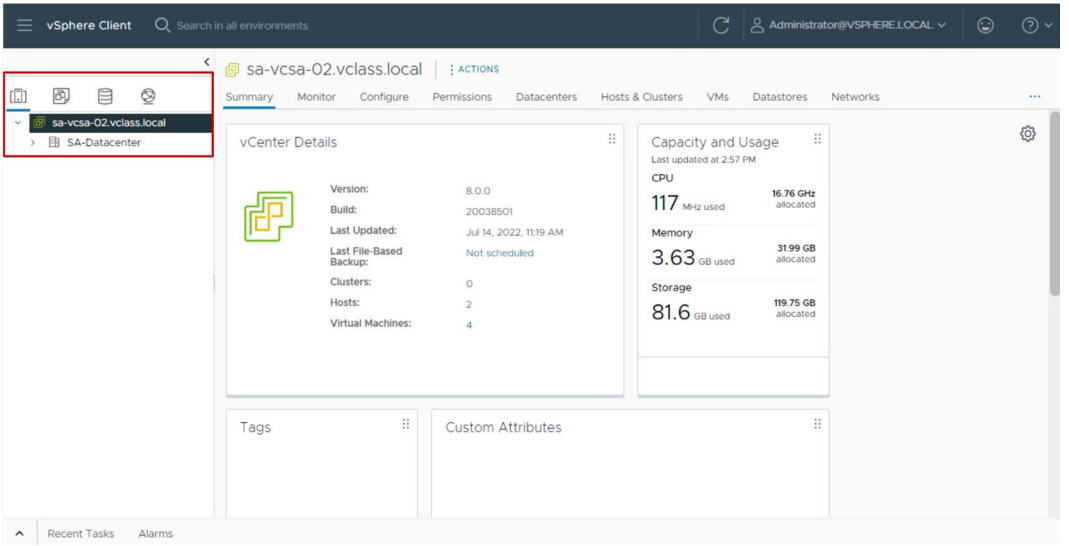

4-43 Views for Hosts, Clusters, VMs, and Templates

Host and cluster objects appear in one view, and VM and template objects are displayed in another view.

Host and Clusters Inventory View

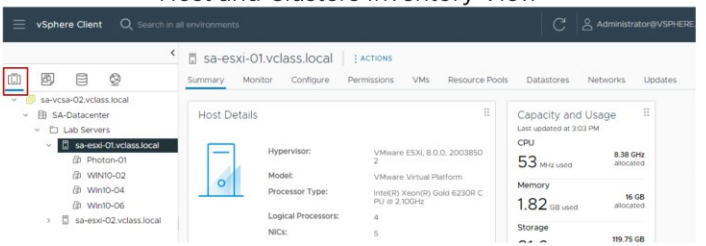

VMs and Templates Inventory View

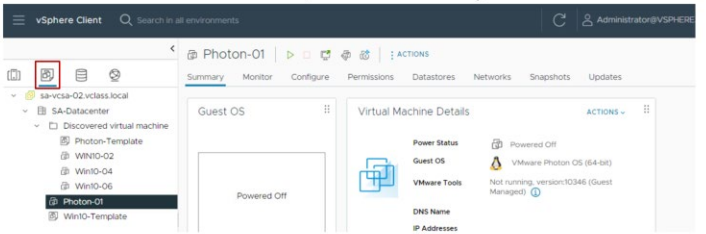

The Hosts and Clusters inventory view shows all host and cluster objects in a data center. You can further organize the hosts and clusters into folders.

The VMs and Templates inventory view shows all VM and template objects in a data center. You can also organize the VMs and templates into folders.

4-44 Views for Storage and Networks

The storage inventory view shows all the details for datastores in the data center.

The networking inventory view shows all the port groups on standard switches and distributed switches.

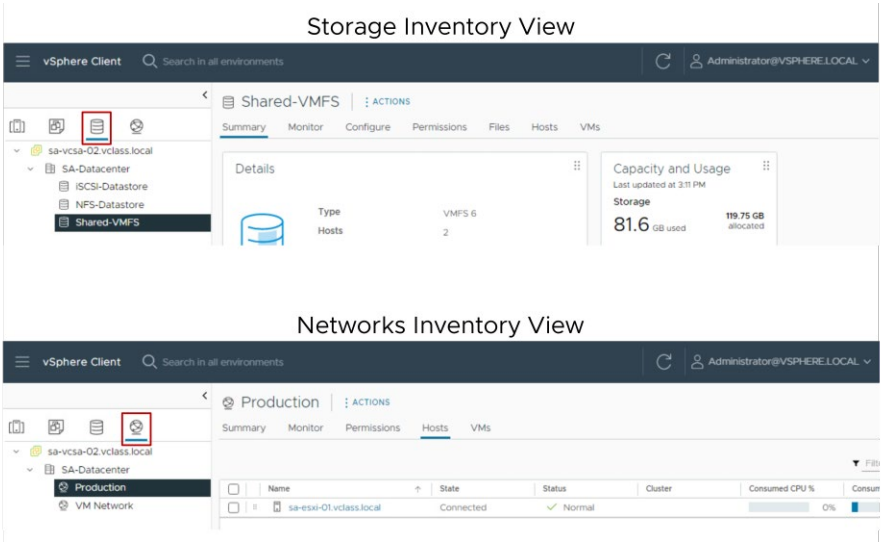

As with the other inventory views, you can organize your datastore and network objects into folders.

4-45 Viewing Object Information

Because you can view object information and access related objects, monitoring and managing object properties is easy.

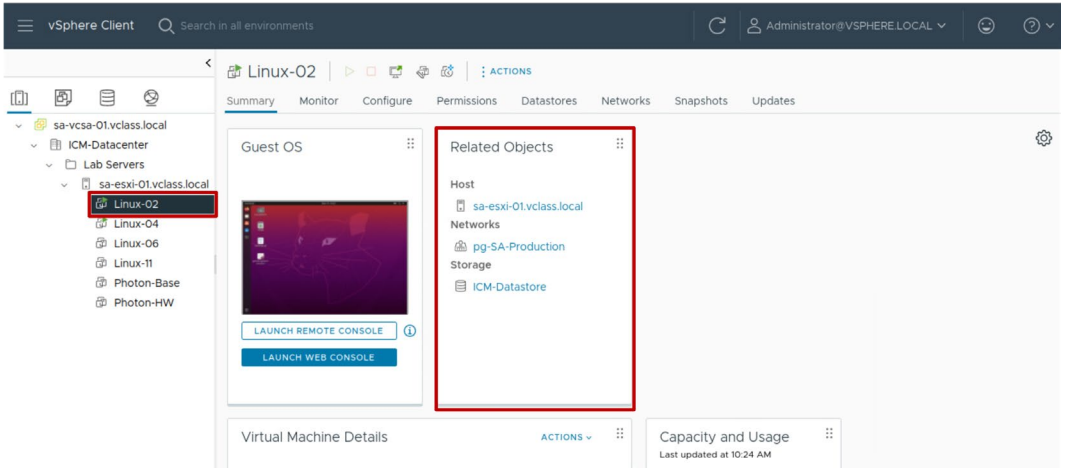

4-46 About Data Center Objects

A virtual data center is a logical organization of all the inventory objects. Those inventory objects are required to complete a fully functional environment for operating VMs:

- You can create multiple data centers to organize sets of environments.
- Each data center has its own hosts, VMs, templates, datastores, and networks.

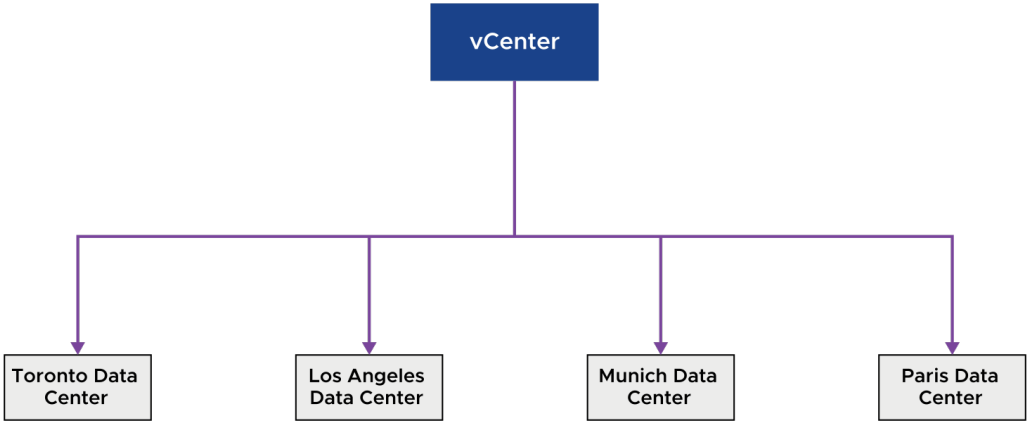

You might create a data center object for each data center geographical location. Or, you might create a data center object for each organizational unit in your enterprise.

4-47 Organizing Inventory Objects into Folders

You can place Objects in a data center in folders. You can create folders and subfolders to better organize systems.

Each of the four inventory views has its own folder structure.

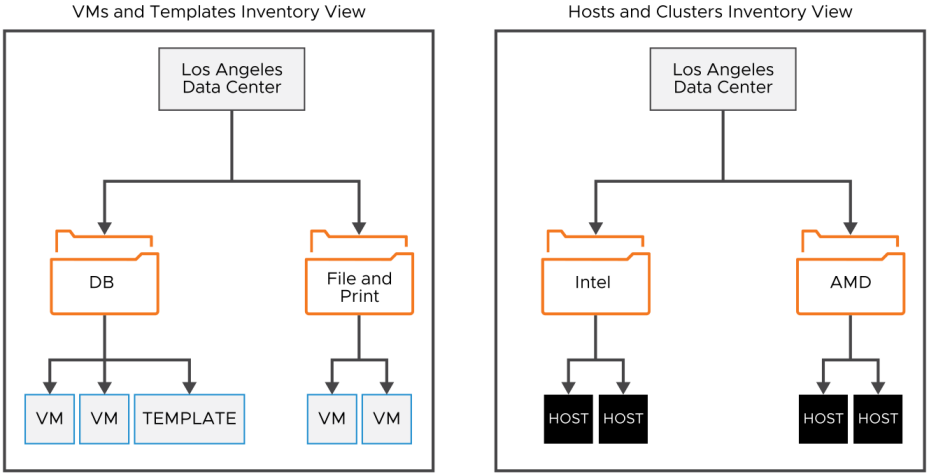

You plan the setup of your virtual environment depending on your requirements.

A large vSphere implementation might contain several virtual data centers with a complex arrangement of hosts, clusters, resource pools, and networks. vSphere implementation might include multiple vCenter systems.

Smaller implementations might require a single virtual data center with a less complex topology.

Regardless of the scale of your virtual environment, consider how the VMs that it supports are used and administered.

Populating and organizing your inventory involves the following tasks:

- Creating data centers
- Creating clusters to consolidate the resources of multiple hosts and VMs
- Adding hosts to the clusters or to the data centers
- Organizing inventory objects in folders
- Setting up networking by using vSphere standard switches or vSphere distributed switches
- Configuring storage systems and creating datastore inventory objects to provide logical containers for storage devices in your inventory

4-48 Adding a Data Center and Organizational Objects to vCenter

You can add a data center, a host, a cluster, and folders to vCenter.

You can use folders to group objects of the same type for easier management.

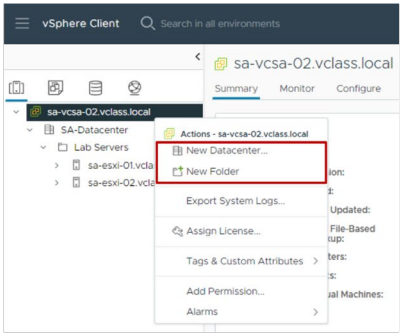

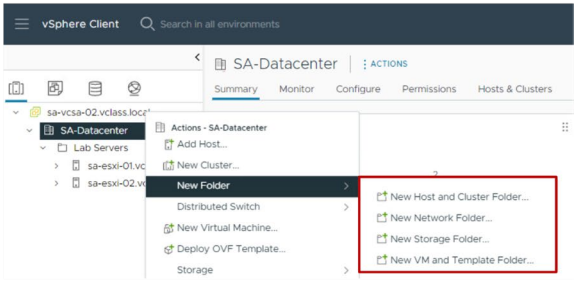

4-49 Adding ESXi Hosts to vCenter

You can add ESXi hosts to vCenter using the vSphere Client.

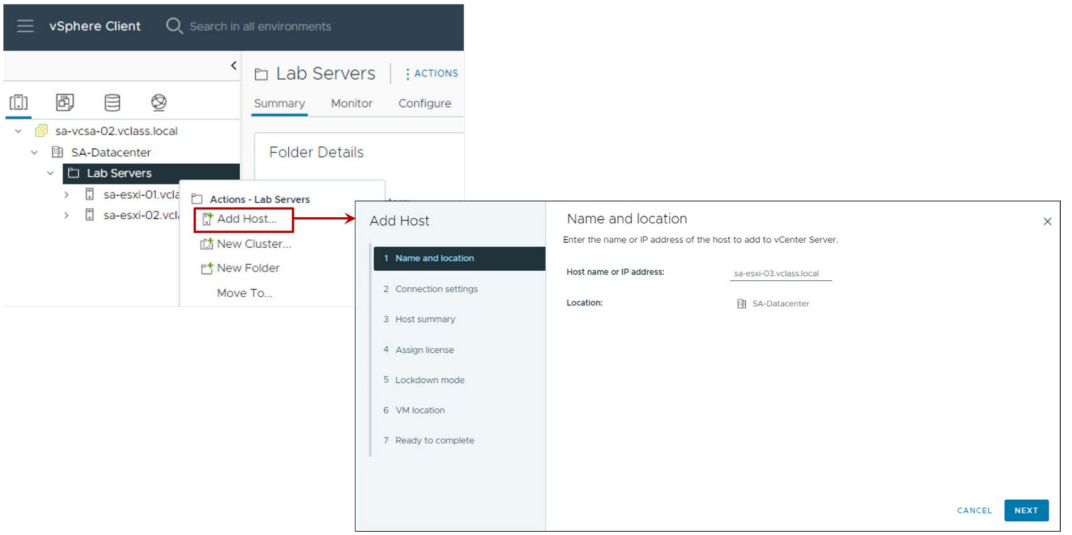

4-50 Creating Custom Tags for Inventory Objects

You can use tags to attach metadata to objects in the vCenter inventory. Tags help make these objects more sortable.

You can associate a set of objects of the same type by searching for objects by a given tag.

You can use tags to group and manage VMs, clusters, and datastores, for example:

- Tag VMs that run production workloads.
- Tag VMs based on their guest operating system.

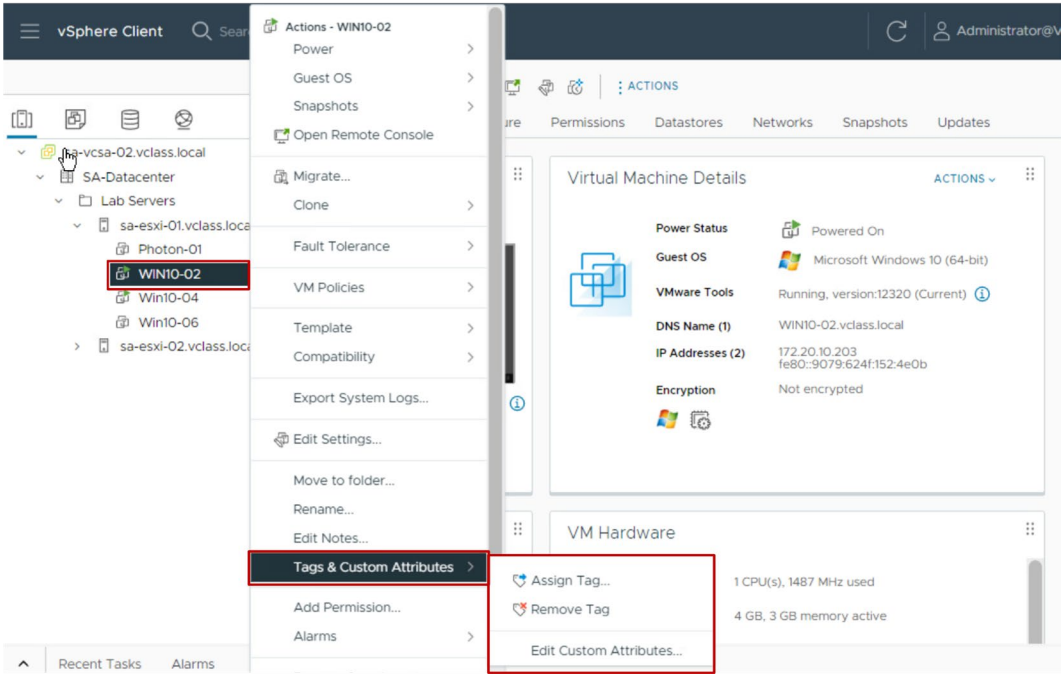

4-51 Lab 4: Creating and Managing the vCenter Inventory

Use the vSphere Client to create and configure objects in the vCenter inventory:

1. Create a Data Center Object
2. Add Two ESXi Hosts to the Inventory
3. View Information About the ESXi Hosts
4. Configure an ESXi Host as an NTP Client
5. Create a Folder for the ESXi Hosts
6. Create Folders for VMs and VM Templates

4-52 Review of Learner Objectives

- Use the vSphere Client to manage the vCenter inventory
- Create and organize vCenter inventory objects
- Add data center and organizational objects to vCenter
- Add ESXi hosts to the inventory
- Create custom inventory tags for inventory objects

4-53 Lesson 5: vCenter Roles and Permissions

4-54 Learner Objectives

- Define the term permission in the context of vCenter
- Recognize the rules for applying permissions
- Create a custom role
- Assign global permission to a user

4-55 About vCenter Permissions

Using the access control system, the vCenter administrator can define user privileges to access objects in the inventory.

The following concepts are important:

- Privilege: An action that can be performed
- Role: A set of privileges
- Object: The target of the action
- User or group: Indication of who can perform the action
- Permission: Gives one user or group a role (set of privileges) for the selected object

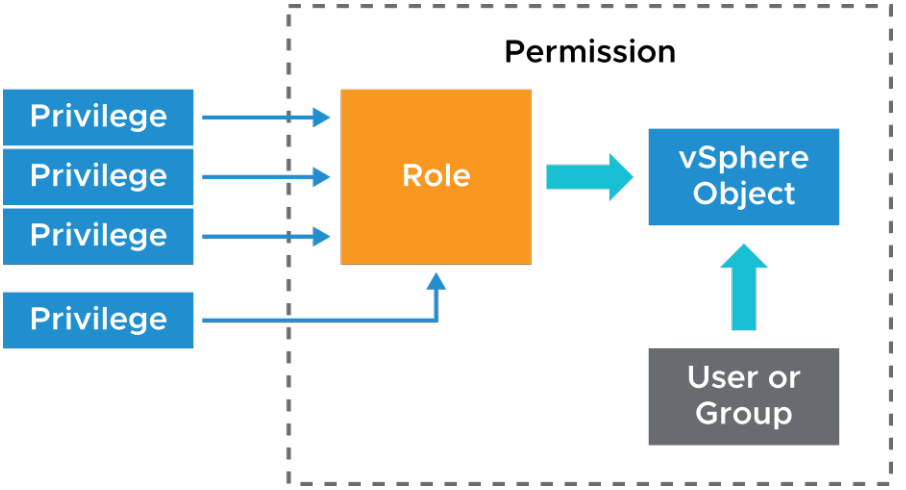

The authorization to perform tasks in vCenter is governed by an access control system. Through this access control system, the vCenter administrator can specify in detail which users or groups can perform which tasks on which objects.

A permission is set on an object in the vCenter object inventory. Each permission associates the object with a group or user and the group or user access roles. For example, you can select a VM object, add one permission that gives the Read-only role to group 1, and add a second permission that gives the Administrator role to user 2.

By assigning a different role to a group of users on different objects, you control the tasks that those users can perform in your vSphere environment. For example, to allow a group to configure memory for the host, select that host and add a permission that grants a role to that group that includes the Host.Configuration.Memory Configuration privilege.

4-56 About Roles

Privileges are grouped into roles:

- A privilege allows access to a specific task and is grouped with other privileges related to it.
- Roles allow users to perform tasks.

vCenter provides a few system roles, which you cannot modify.

Sample roles are also provided. You can clone them to create custom roles.

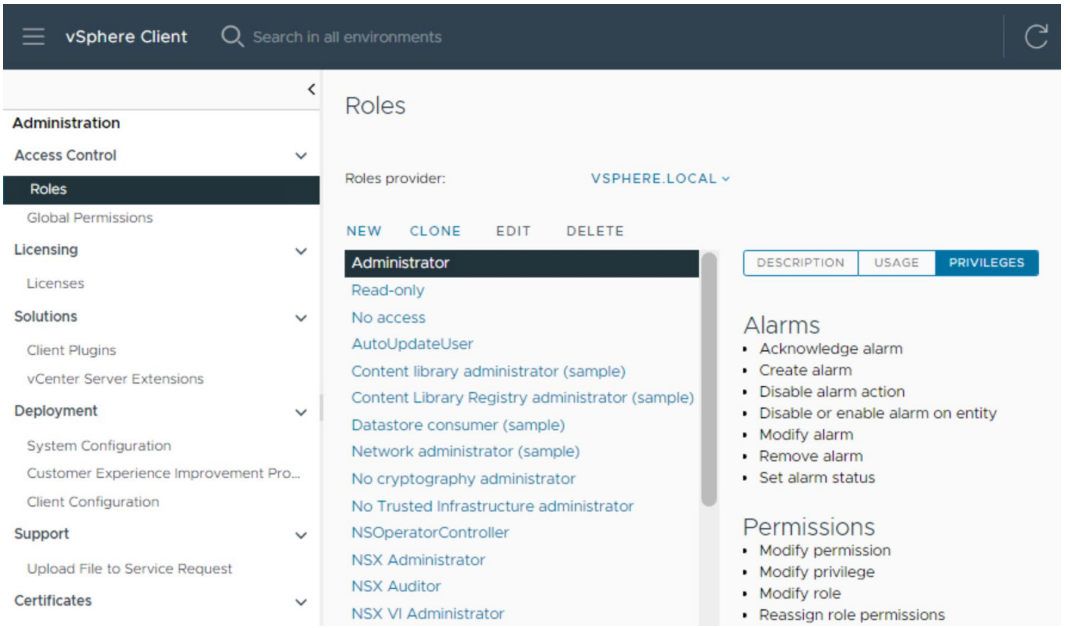

A role is a set of one or more privileges. For example, the Virtual Machine Power User sample role consists of several privileges in categories such as Datastore and Global. A role is assigned to a user or group and determines the level of access of that user or group.

You cannot change the privileges associated with the system roles:

- Administrator role: Users with this role for an object may view and perform all actions on the object.
- Read-only role: Users with this role for an object may view the state of the object and details about the object.

- No access role: Users with this role for an object may not view or change the object in any way.
- No cryptography administrator role: Users with this role for an object have the same privileges as users with the Administrator role, except for privileges in the Cryptographic operations category.

All roles are independent of each other. There is no hierarchy or inheritance between roles.

4-57 About Objects

Objects are entities on which actions are performed. Objects include data centers, folders, clusters, hosts, datastores, networks, and virtual machines.

All objects have a **Permissions** tab. The **Permissions** tab shows which user or group and role are associated with the selected object.

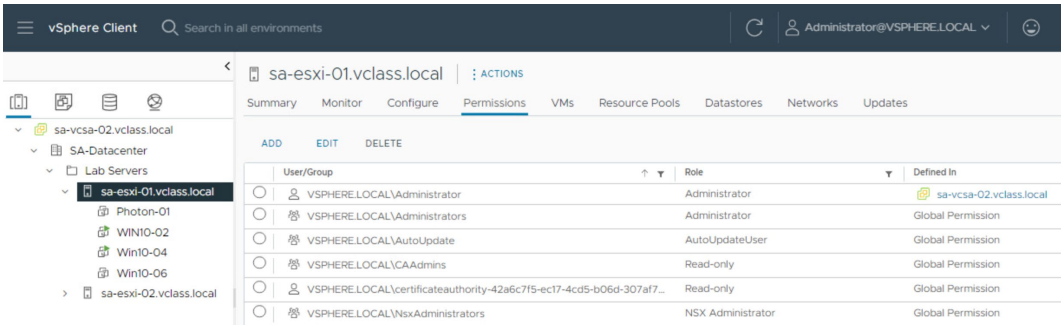

4-58 Assigning Permissions

To assign a permission:

1. Select an object
2. Select a **Domain**
3. Select a **User/Group**
4. Select a **Role**
5. Propagate the permission to the child objects

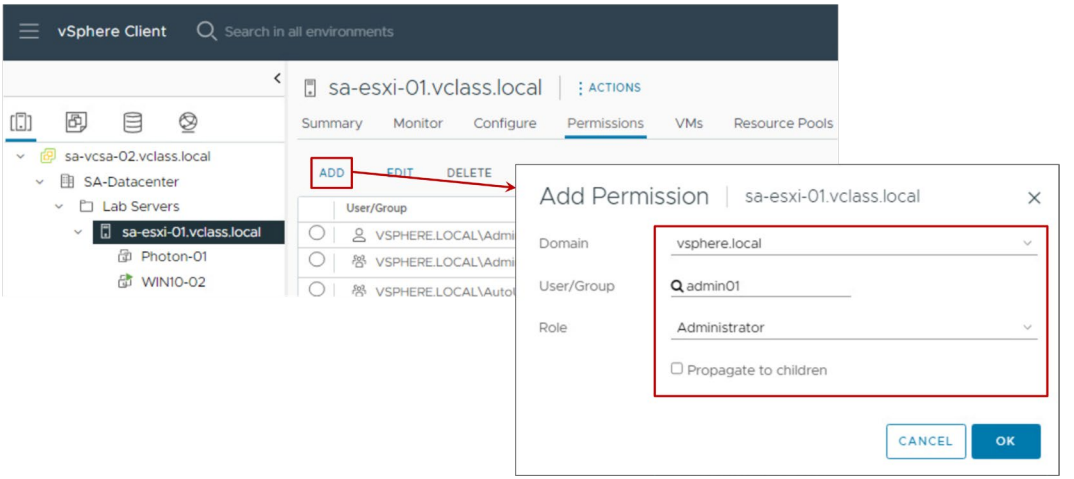

You can assign permissions to objects at different levels of the hierarchy. For example, you can assign permissions to a host object or to a folder object that includes all host objects. You can also assign permissions to the global root object to apply the permissions to all objects in all solutions.

For information about hierarchical inheritance of permissions and global permissions, see *vSphere Security* at https://docs.vmware.com/en/VMware-vSphere/index.html

4-59 Viewing Roles and User Assignments

The **Roles** pane shows which users are assigned the selected role on a particular object.

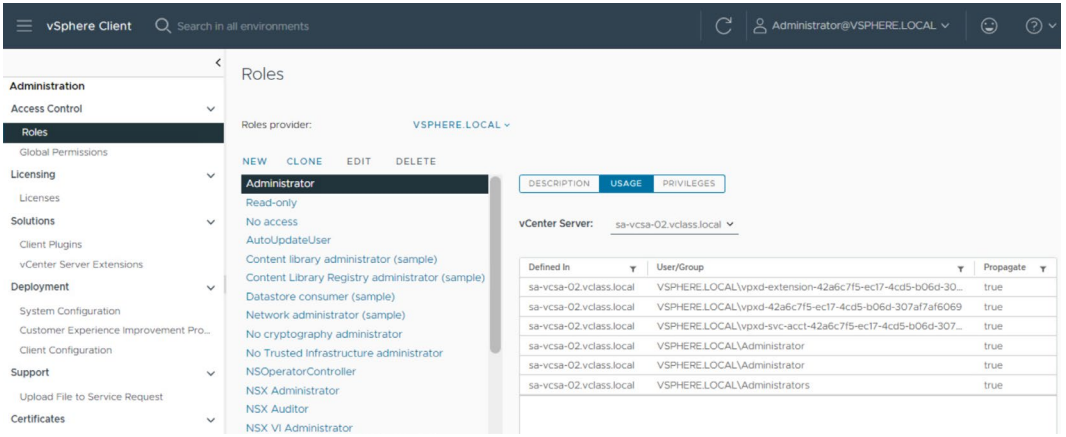

You can view all the objects to which a role is assigned and all the users or groups who are granted the role.

To view information about a role, click **Usage** in the Roles pane and select a role from the Roles list. The information provided to the right shows each object to which the role is assigned and the users and groups who were granted the role.

4-60 Applying Permissions: Scenario 1

A permission can propagate down the object hierarchy to all sub-objects, or a permission can apply only to a specific object.

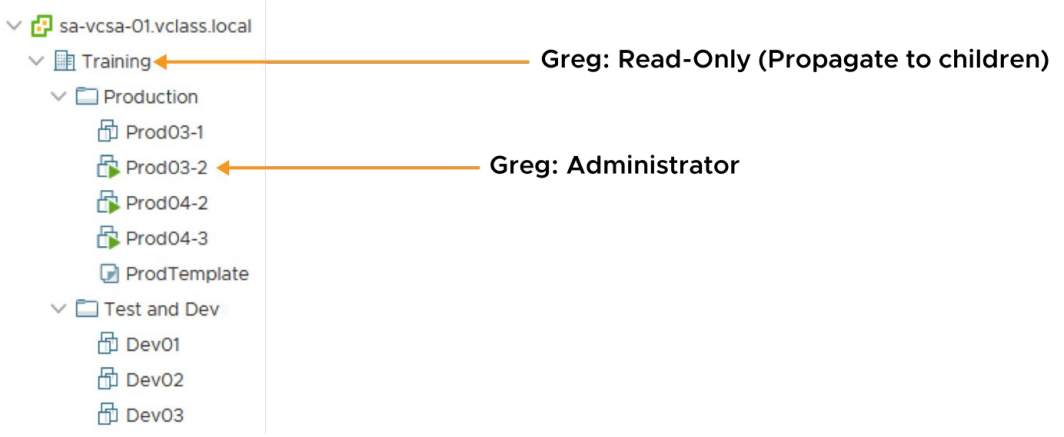

In addition to specifying whether permissions propagate downward, you can override permissions to set at a higher level by explicitly setting different permissions for a lower-level object.

On the slide, user “Greg,” is given Read-only access to the Training data center. This role is propagated to all child objects except one, the Prod03-2 VM. For this VM, Greg is an administrator.

4-61 Applying Permissions: Scenario 2

When a user is a member of multiple groups with permissions on the same object, the user is assigned the union of privileges assigned to the groups for that object.

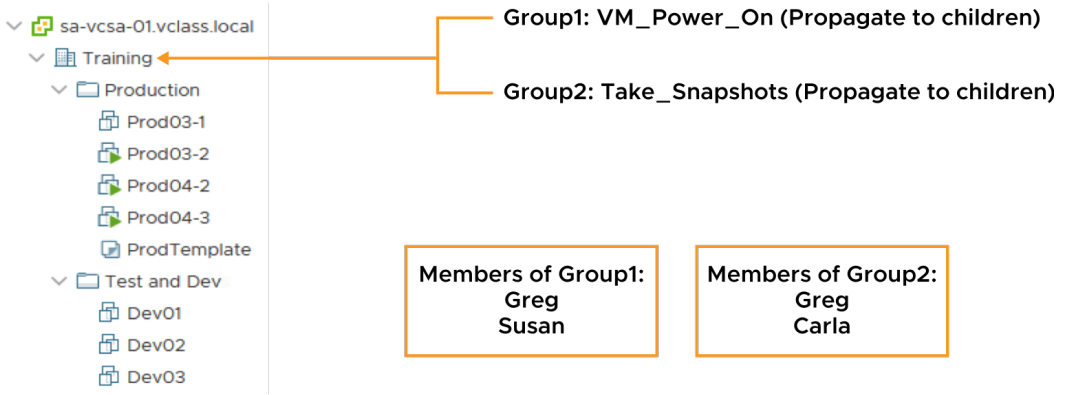

On the slide, Group1 is assigned the VM_Power_On role, a custom role that contains only one privilege: the ability to power on a VM. Group2 is assigned the Take_Snapshots role, another custom role that contains the privileges to create and remove snapshots. Both roles propagate to the child objects.

Because Greg belongs to both Group1 and Group2, he gets both VM_Power_On and Take_Snapshots privileges for all objects in the Training data center.

4-62 Activity: Applying Group Permissions (1)

If Group1 has the Administrator role and Group2 has the No Access role, what permissions does Greg have?

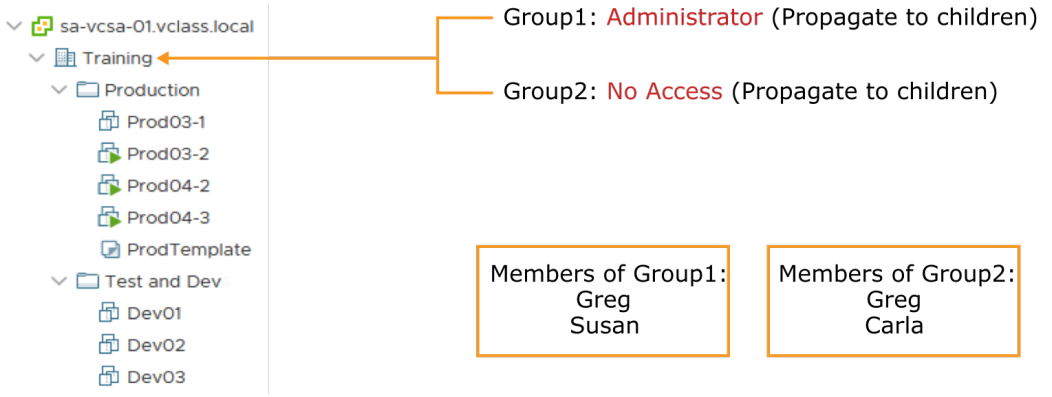

4-63 Activity: Applying Group Permissions (2)

Greg has Administrator privileges.

Greg is assigned the union of privileges assigned to Group1 and Group2.

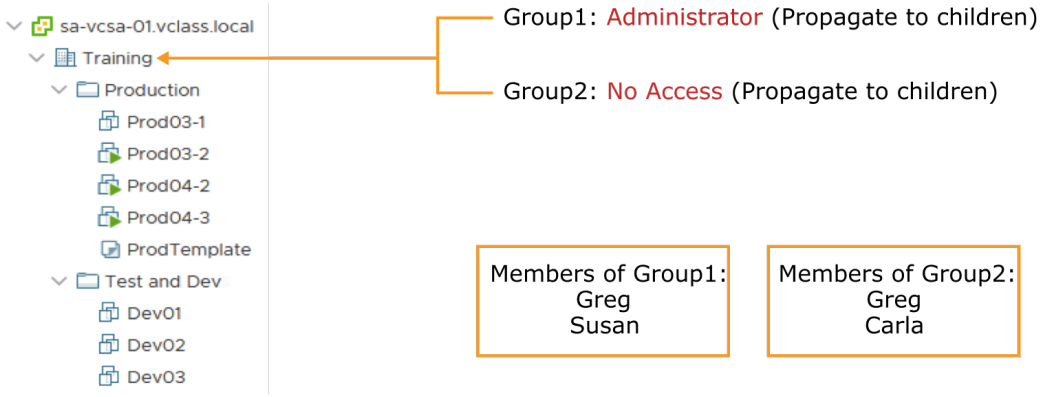

4-64 Applying Permissions: Scenario 3

A user can be a member of multiple groups with permissions on different objects. In this case, the same permissions apply for each object on which the group has permissions, as though the permissions were granted directly to the user.

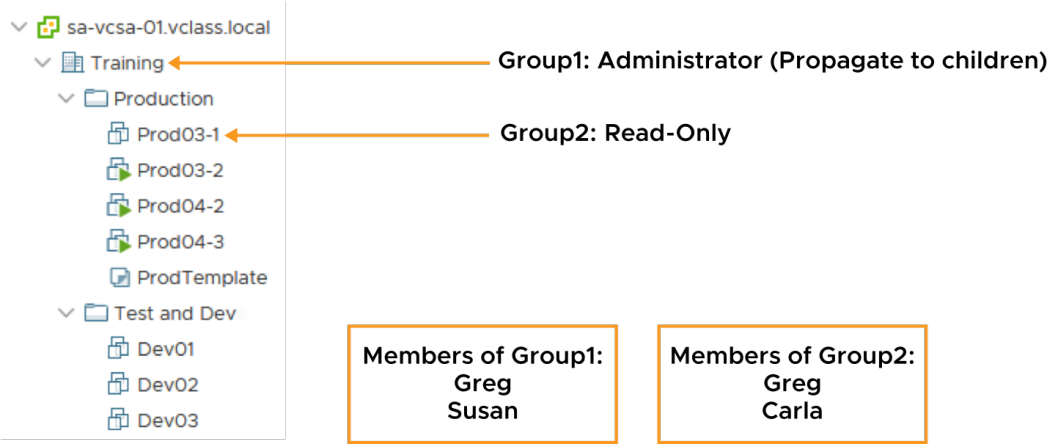

You can override permissions set for a higher-level object by explicitly setting different permissions for a lower-level object.

On the slide, Group1 is assigned the Administrator role at the Training data center and Group2 is assigned the Read-only role on the VM object, Prod03-1. The permission granted to Group1 is propagated to child objects.

Because Greg is a member of both Group1 and Group2, he gets administrator privileges on the entire Training data center (the higher-level object), except for the VM called Prod03-1 (the lower-level object). For this VM, Greg gets read-only access.

4-65 Applying Permissions: Scenario 4

A user (or group) is given only one role for any given object.

Permissions defined explicitly for the user on an object take precedence over all group permissions on that same object.

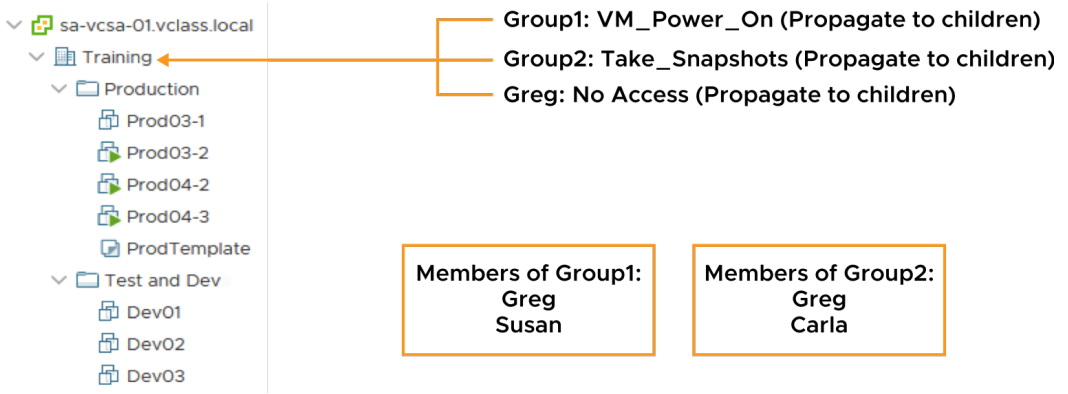

On the slide, three permissions are assigned to the Training data center:

- Group1 is assigned the VM_Power_On role.
- Group2 is assigned the Take_Snapshots role.
- Greg is assigned the No Access role.

Greg is a member of both Group1 and Group2. Assume that propagation to child objects is selected on all roles. Although Greg is a member of both Group1 and Group2, Greg gets the No Access privilege to the Training data center and all objects under it. Greg gets the No Access privilege because explicit user permissions on an object take precedence over all group permissions on that same object.

4-66 Creating a Role

Create roles with only the necessary privileges.

For example, you can create a Provision VMs role that allows a user to deploy VMs from a template.

Use folders to contain the scope of permissions. For instance, you can assign the Provision VMs role to user nancy@company.com and apply it to the Production VMs folder.

Provision VMs role:

Datastore > Allocate space

Resource > Assign virtual machine to resource pool

Virtual machine >

- Edit Inventory > Create from existing
- Interaction > Power on

Provisioning >

- Allow read-only disk access
- Customize guest
- Deploy template
- Read customization specifications

The Provision VMs role is one of many examples of roles that you can create.

Define a role using the smallest number of privileges possible to maximize security and control over your environment. Give the roles names that explicitly indicate what each role allows, to make its purpose clear.

4-67 About Global Permissions

Global permissions support assigning privileges across solutions from the global root object:

- Span solutions, such as vRealize Orchestrator, and multiple vCenter instances
- Give a user or group privileges for all objects in all vCenter hierarchies

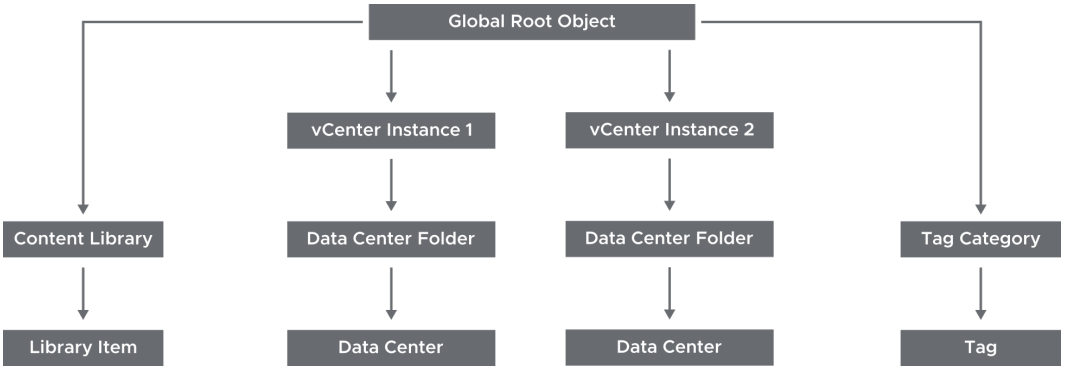

Often, you apply a permission to a vCenter inventory object, such as an ESXi host or a VM. When you apply a permission, you specify that a user or group has a set of privileges, called a role, on the object.

Global permissions give a user or group privileges to view or manage all objects in each of the inventory hierarchies in your deployment.

The example shows that the global root object has permissions over all vCenter objects, including content libraries, vCenter instances, and tags. Global permissions allow access across vCenter instances. vCenter permissions, however, Global permissions are effective only on objects in a particular vCenter instance.

4-68 Lab 5: Adding an Identity Source

Add vclass.local as an LDAP identity source:

1. Add vclass.local as an LDAP Identity Source

4-69 Lab 6: Users, Groups, and Permissions

Assign roles and permissions so that an LDAP user can perform functions in vCenter:

1. View LDAP Users
2. Assign Root-Level Global Permission to an LDAP User
3. Assign Object Permission to an LDAP User
4. Verify that the cladmin User Can Access Content Library
5. Verify that the studentadmin User Can Create a Virtual Machine

4-70 Review of Learner Objectives

- Define the term permission in the context of vCenter
- Recognize the rules for applying permissions
- Create a custom role
- Assign global permission to a user

4-71 Lesson 6: Monitoring vSphere Events

4-72 Learner Objectives

- Monitor tasks and events that occur on objects in the vCenter inventory
- Recognize the vCenter log levels for controlling the amount of data collected in the vCenter database

4-73 About vSphere Tasks

Every action that you perform in vSphere as a part of your day-to-day operations is called a task, for example:

- Powering on a virtual machine
- Updating the network configuration
- Modifying the configuration of hosts and virtual machines

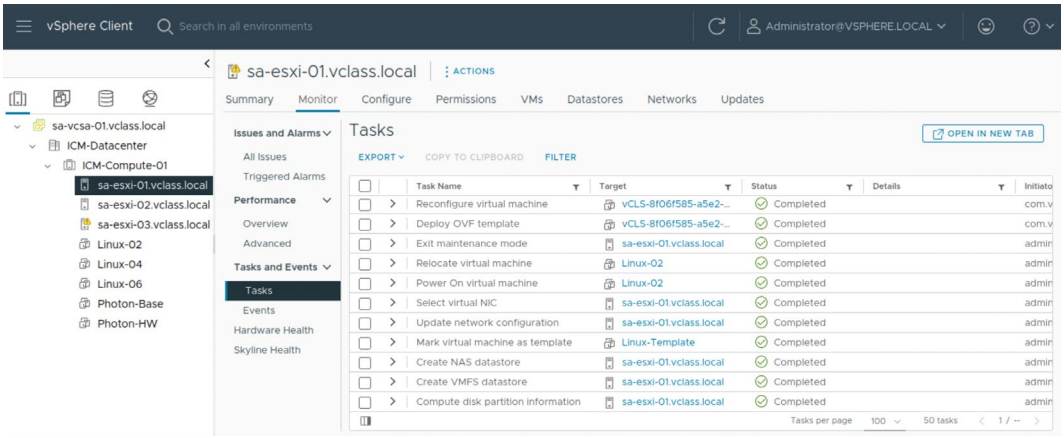

As an administrator, you can monitor who performs tasks in the vSphere Client. Task information is useful when you troubleshoot problems because the tasks list shows the actions performed in the vSphere environment.

4-74 About vSphere Events

vSphere events are records of user actions or system actions that occur on objects in the vCenter inventory:

- User-action information includes the user's account and specific event details.
- Event details are reported, such as the event's date and time, type, description, and the object on which the event occurred.
- Events and alarms are displayed to alert the user to changes in the vCenter service health or when a service fails.

The vCenter Tasks and Events panes provide an audit trail, maintaining a 30-day history, by default.

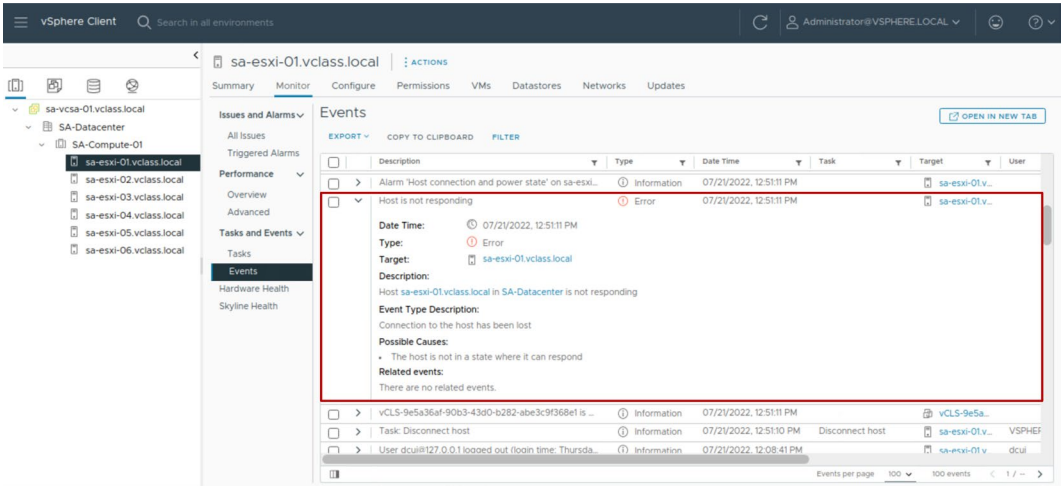

Actions that might be recorded as events include, but are not limited to, the following examples:

- A license key expires.
- A virtual machine is powered on.
- A user logs in to a virtual machine.
- A host connection is lost.

4-75 About vCenter Log Levels

vCenter services create their own log files, which can be used for troubleshooting purposes.

You can set log levels to control the quantity and type of information stored by vCenter.

Examples of when to set log levels:

- When troubleshooting complex issues, set the log level to verbose or trivia.
- For controlling the amount of information being stored in the log files.

Option	Description
None	Turns off logging
Error (errors only)	Displays only error log entries
Warning (errors and warnings)	Displays warning and error log entries
Info (normal logging)	Displays information, error, and warning log entries
Verbose	Displays information, error, warning, and verbose log entries
Trivia (extended verbose)	Displays information, error, warning, verbose, and trivia log entries

Changes to the logging settings take effect immediately. You do not have to restart the vCenter system.

When you set the log level to verbose or trivia to troubleshoot problems, always remember to set the log level back to info (normal logging) when you are done.

4-76 Setting Log Levels

You can configure the amount of log information detail that vCenter collects in log files:

- Edit the log levels in the vSphere Client.
- More verbose logging requires more space on your vCenter system.

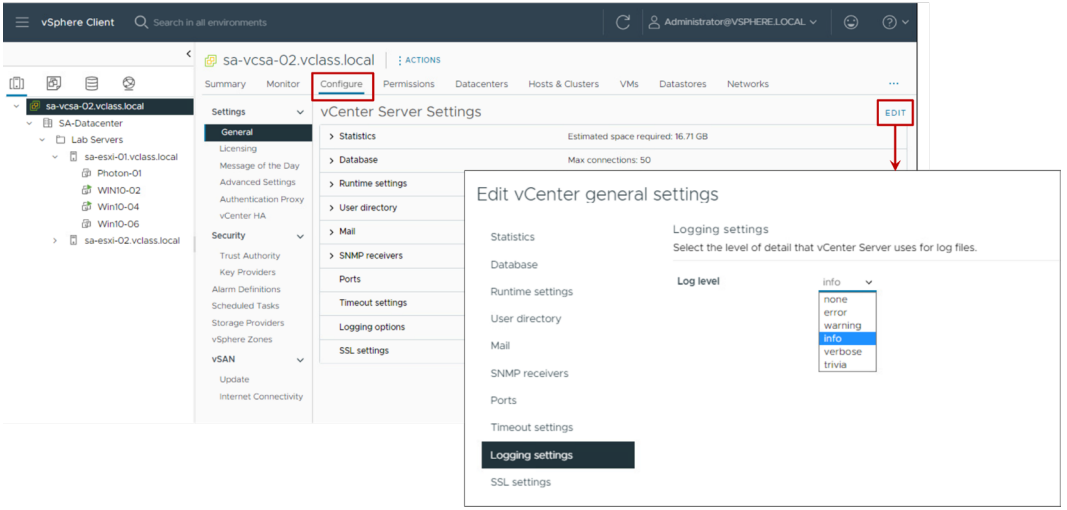

To configure logging levels, follow these steps:

1. In the vSphere Client, select the vCenter instance in the navigation pane
2. Click the **Configure** tab
3. Under Settings, select **General**
4. Click **EDIT**
5. Select **Logging settings** in the left pane
6. Select an option from the **Log level** drop-down menu

4-77 Forwarding vCenter Log Files to a Remote Host

vCenter can stream its log information to a remote Syslog server.

You can activate this feature in the vCenter Management Interface.

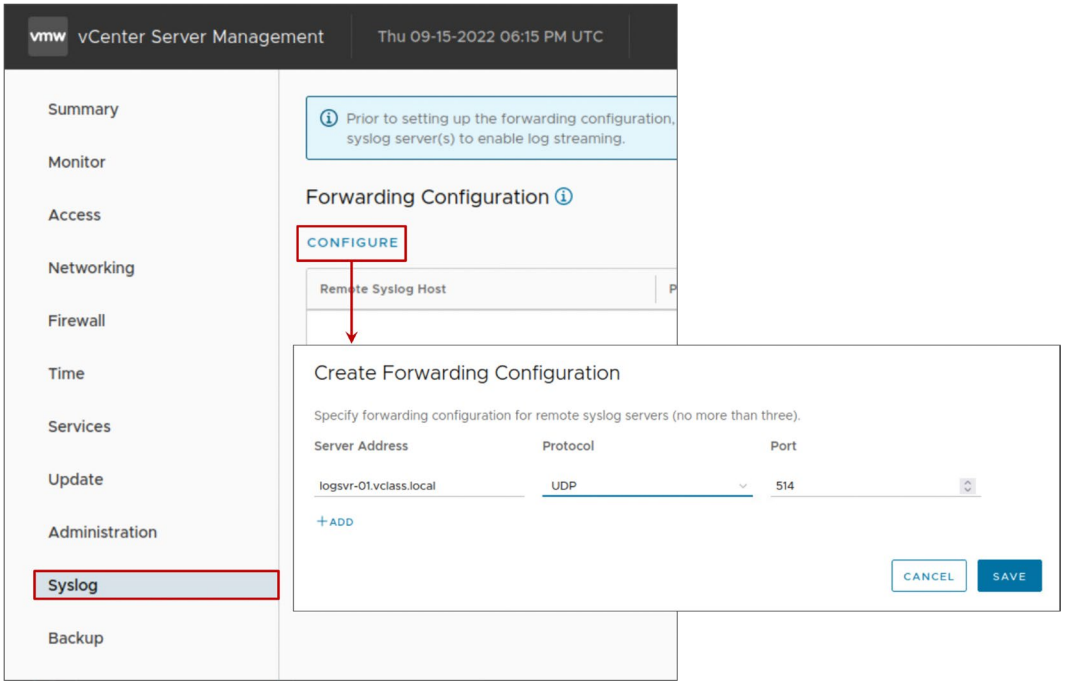

4-78 Forwarding ESXi Host Log Files to a Remote Host

For ESXi hosts, specify the remote Syslog server name in the Advanced System Settings pane in the vSphere Client.

You can further analyze ESXi host log files with log analysis products, such as vRealize Log Insight.

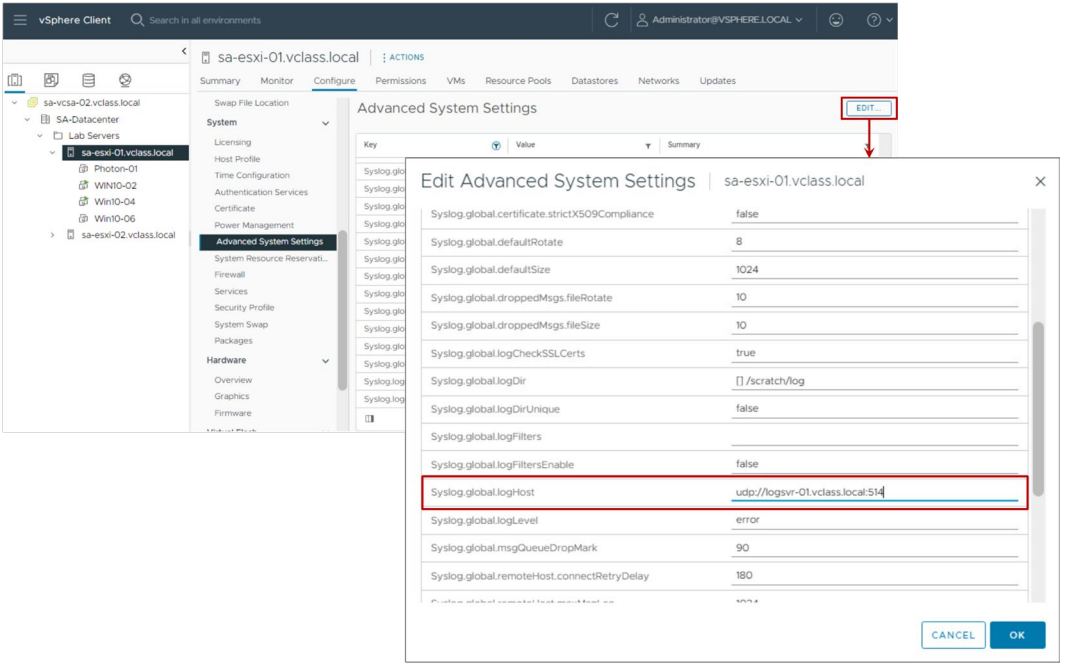

4-79 Review of Learner Objectives

- Monitor tasks and events that occur on objects in the vCenter inventory
- Recognize the vCenter log levels for controlling the amount of data collected in the vCenter database

4-80 Key Points

- vCenter Server Appliance uses the Photon operating system and the PostgreSQL database.
- You can use the vCenter Management Interface to manage vCenter, including vCenter networking and vCenter services.
- You use the vSphere Client to connect to vCenter instances and manage vCenter inventory objects.
- A permission, defined in vCenter, gives one user or group a role (set of privileges) for a selected object.
- Global permission allows access to all vCenter objects, including content libraries, vCenter instances, and tags.
- You can control the vCenter logging level. Changing the logging level affects the vCenter's filesystem usage.

Questions?

Module 5

Configuring vSphere Networking

5-2 Importance

When you successfully configure ESXi networking, virtual machines can communicate with other machines, both virtual and physical. Additionally, a successfully configured ESXi network allows the VMkernel to operate remote host management and IP-based storage effectively.

vSphere standard switches provide effective networking for small environments. As you scale your vSphere environment, the built-in features and functions of vSphere distributed switches can help you manage networking in larger environments.

5-3 Module Lessons

1. vSphere Standard Switches
2. Configuring Virtual Switch Policies
3. vSphere Distributed Switches

5-4 Lesson 1: vSphere Standard Switches

5-5 Learner Objectives

- Identify virtual switch connection types
- Configure and view standard switch configurations

5-6 About Virtual Switches

Virtual switches connect VMs to the physical network.

They provide connectivity between VMs on the same ESXi host or on different ESXi hosts.

They also support VMkernel services, such as vSphere vMotion migration, iSCSI, NFS, and access to the management network.

5-7 Types of Virtual Switches

A virtual network supports standard and distributed switches. Both switch types are elastic, ports are created and removed automatically:

- Standard switch:
  - Virtual switch that is configured for a single host.
- Distributed switch:
  - Virtual switch that is configured for an entire data center.
  - Up to 2,000 hosts can be attached to the same distributed switch.
  - The configuration is consistent across all attached hosts.
  - Hosts must either have an Enterprise Plus license or belong to a vSAN cluster.

5-8 Types of Virtual Switch Connections

A virtual switch has specific connection types:

- VM ports
- VMkernel ports
  - IP storage, vSphere vMotion migration, vSphere Fault Tolerance, vSAN, vSphere Replication, and the ESXi management network
- Uplink ports

VM ports and VMkernel ports exist in port groups.

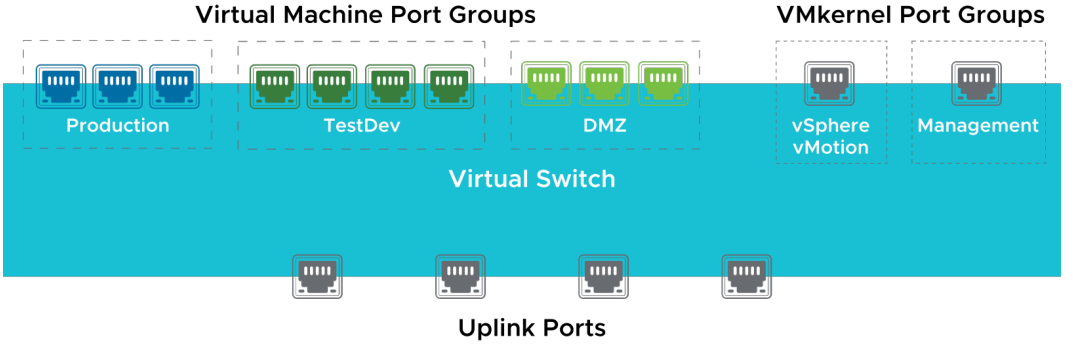

The ESXi management network port is a VMkernel port that connects to network or remote services, including vpxd on vCenter and VMware Host Client.

Each ESXi management network port as well as all other VMkernel ports must be configured with their own IP address, netmask, and gateway.

To configure virtual switches, you create port groups. A port group is a template that stores configuration information to create virtual switch ports on a virtual switch. Port groups can contain VM ports, which connect VMs to one another with common networking properties. Port groups can also contain VMkernel ports.

VM ports and VMkernel ports connect to the outside world through the physical Ethernet adapters that are connected to the virtual switch uplink ports.

5-9 Virtual Switch Connection Examples

Networks (port groups) can coexist on the same virtual switch or on separate virtual switches.

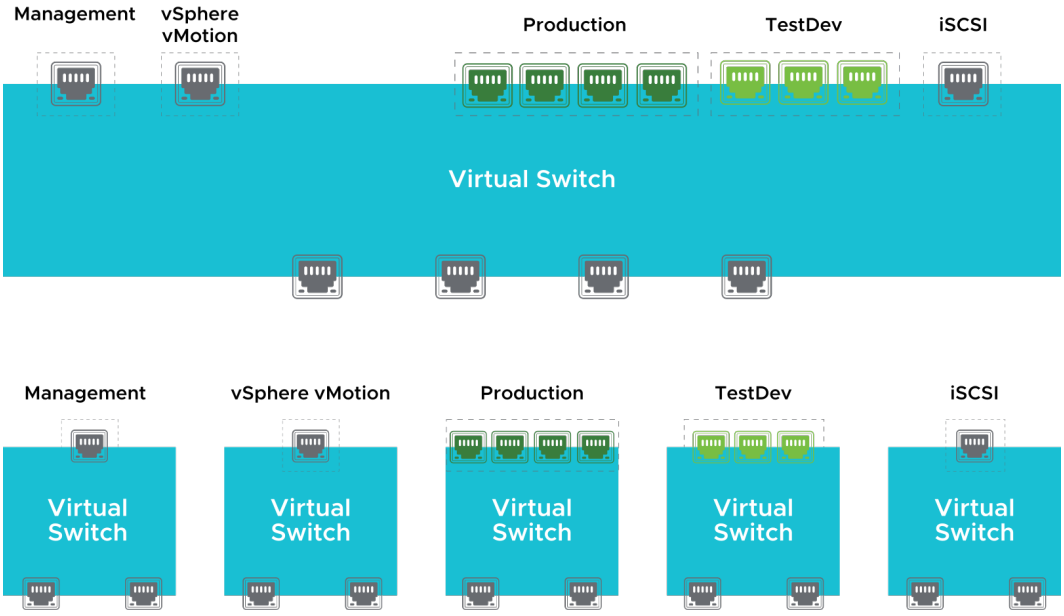

When you design your networking environment, you can group all your networks on a single virtual switch. Alternatively, you can opt for multiple virtual switches, each with a separate network. The decision partly depends on the layout of your physical networks.

For example, you might not have enough network adapters to create a separate virtual switch for each network. Instead, you might place your network adapters in a single virtual switch and isolate the networks by using VLANs.

Because physical NICs are assigned at the virtual switch level, all ports and port groups that are defined for a particular switch share the same hardware.

5-10 About VLANs and Virtual Switch Tagging

ESXi provides VLAN support by assigning a VLAN ID to a port group. ESXi supports 802.1Q VLAN tagging.

Virtual switch tagging is one of the supported tagging policies:

- Frames from a VM are tagged as they exit the virtual switch.
- Tagged frames arriving at a virtual switch are untagged before they are sent to the destination VM.
- The effect on performance is minimal.

Physical switch ports must be configured as trunk ports.

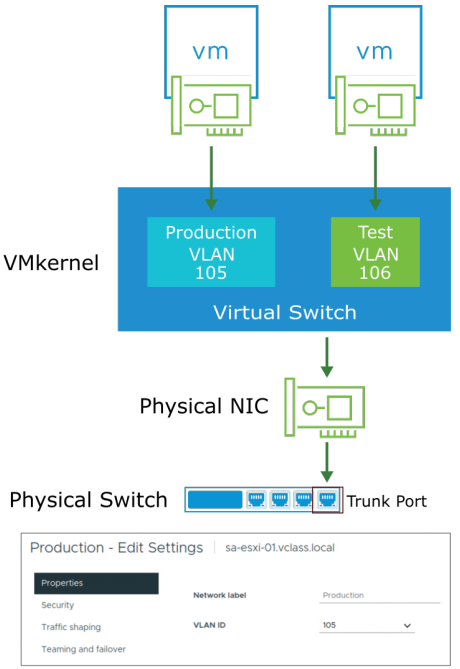

VLANs provide for logical groupings of switch ports. All virtual machines or ports in a VLAN communicate as if they are on the same physical LAN segment. A VLAN is a software-configured broadcast domain. Using a VLAN provides the following benefits:

- Creation of logical networks that are not based on the physical topology
- Improved performance by confining broadcast traffic to a subset of ports on a switch
- Cost savings by partitioning the network without the overhead of deploying new routers

VLANs are configured at the port group level. The ESXi host provides VLAN support through virtual switch tagging, which is provided by giving a port group a VLAN ID. By default, the VLAN ID is 0. The VMkernel takes care of all tagging and untagging as the packets pass through the virtual switch.

The port on a physical switch to which an ESXi host is connected must be defined as a static trunk port. A trunk port is a port on a physical Ethernet switch that is configured to send and receive packets tagged with a VLAN ID. No VLAN configuration is required in the VM. In fact, the VM does not know that it is connected to a VLAN.

For more information about how VLANs are implemented, see VMware knowledge base article 1003806 at http://kb.vmware.com/kb/1003806.

5-11 Viewing Standard Switches

In the vSphere Client, you can view a host's standard switch configuration by selecting **Virtual Switches** on the **Configure** tab.

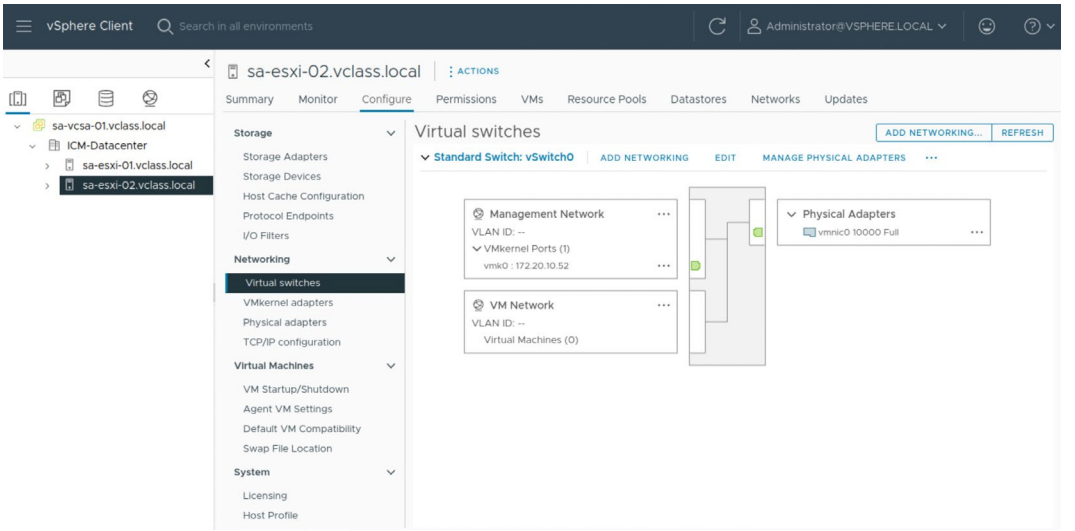

The slide shows the standard switch, vSwitch0, on the sa-esxi-01.vclass.local ESXi host. By default, the ESXi installation creates a virtual machine port group named VM Network and a port group named Management Network that contains a VMkernel port for management traffic. You can create additional port groups for VMs and VMkernel ports. For example, you can create an IP Storage port group that contains a VMkernel port for accessing iSCSI storage.

For performance and security, you should remove the VM Network virtual machine port group and keep VM networks and management networks separated on different physical networks or VLANs.

5-12 Adding Standard Switches

You can add new standard switches to an ESXi host or configure existing ones using the vSphere Client or VMware Host Client.

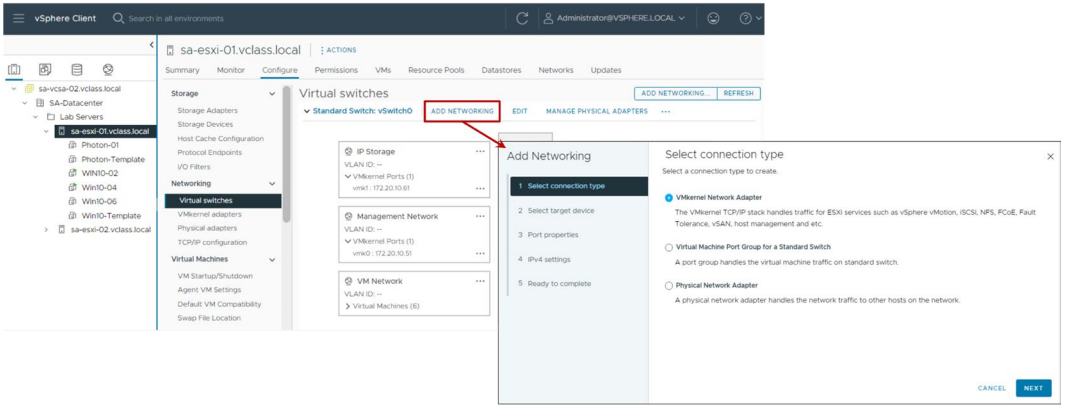

5-13 VMkernel Adapter Properties

The VMkernel adapters pane shows details about the VMkernel interfaces, such as its name, the switch on which it is located, the IP address, and enabled services.

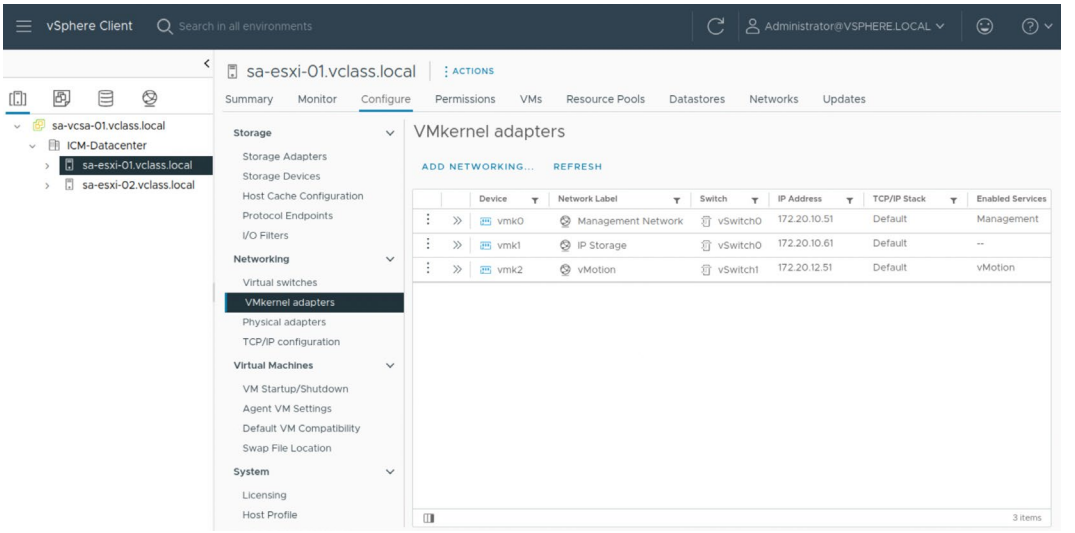

5-14 VMkernel Adapter Properties: Enabled Services

You can activate services for the VMkernel adapter.

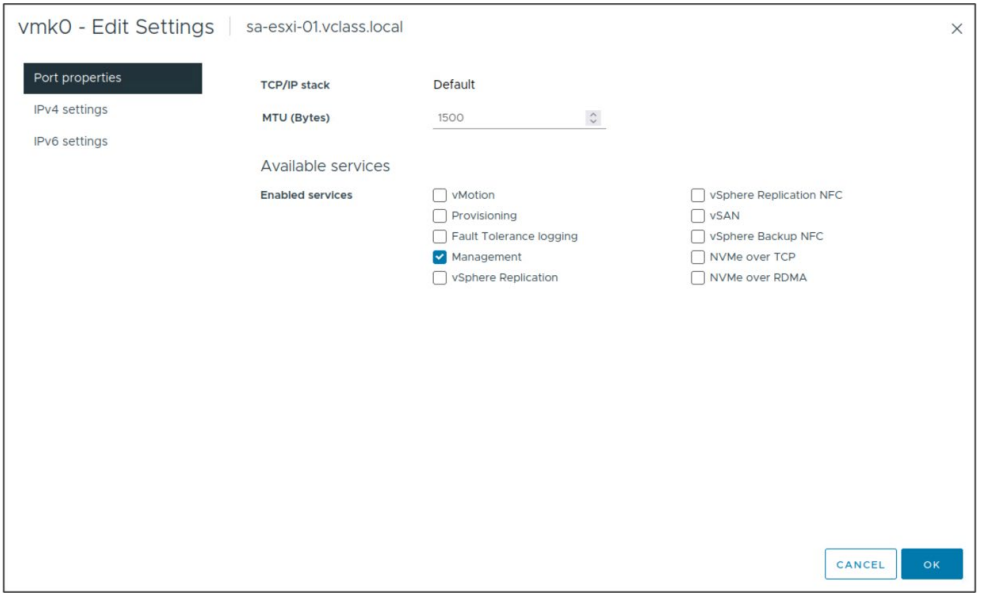

Select from the available services:

- vMotion: Allows the VMkernel adapter to advertise itself to another host as the network connection where vSphere vMotion traffic is sent.
- Provisioning: Handles the data transferred for virtual machine cold migration, cloning, and snapshot migration.
- Fault Tolerance logging: Activates Fault Tolerance logging on the host.
- Management: Activates the management traffic for the host and vCenter.
- vSphere Replication: Handles the outgoing replication data that is sent from the source ESXi host to the vSphere Replication server.
- vSphere Replication NFC: Handles the incoming replication data on the target replication site.
- vSAN: Activates the vSAN traffic on the host.
- vSphere Backup NFC: VMkernel port setting for dedicated backup NFC traffic.

- NVMe over TCP: VMkernel port setting for dedicated NVMe over TCP storage traffic. NVMe over TCP storage traffic goes through the VMkernel Adapter when NVMe over TCP adapter is enabled.
- NVMe over RDMA: VMkernel port setting for dedicated NVMe over RDMA storage traffic. NVMe over RDMA storage traffic goes through the VMkernel Adapter when NVMe over RDMA adapter is enabled.

5-15 Physical Adapter Properties

The Physical adapters pane shows adapter details such as speed, duplex, and networks.

Although the speed and duplex settings are configurable, the best practice is to leave the settings to auto negotiation.

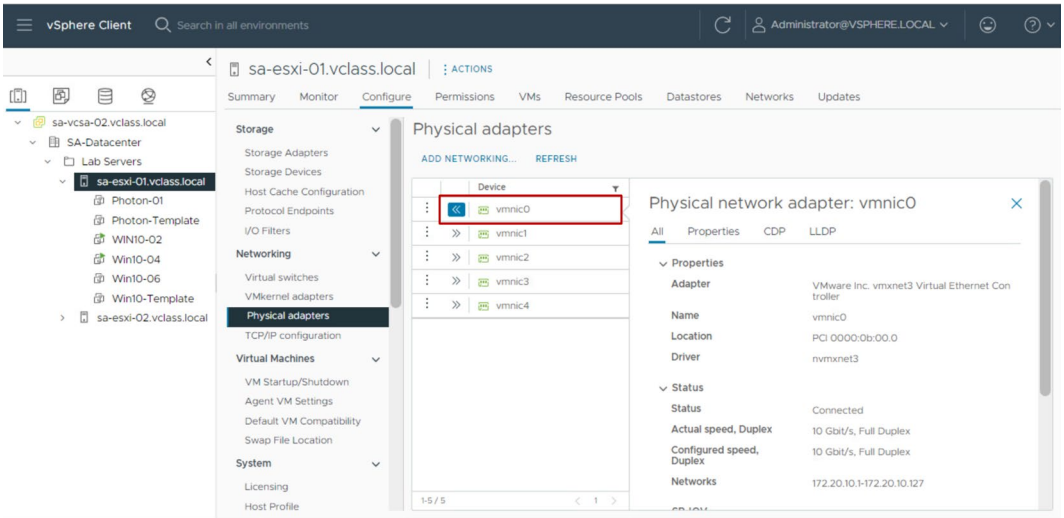

You can change the connection speed and duplex of a physical adapter to transfer data in compliance with the traffic rate.

5-16 Lab 7: Creating Standard Switches

Create a standard switch and a port group for virtual machines:

1. View the Standard Switch Configuration
2. Create a Standard Switch with a Virtual Machine Port Group
3. Attach Virtual Machines to the Virtual Machine Port Group

5-17 Review of Learner Objectives

- Identify virtual switch connection types
- Configure and view standard switch configurations

5-18 Lesson 2: Virtual Switch Networking Policies

5-19 Learner Objectives

- Explain how to set security policies for a virtual switch
- Explain how to set traffic shaping policies for a virtual switch
- Explain how to set NIC teaming and failover policies for a virtual switch

5-20 About Networking Policies

As an administrator, you set networking policies on virtual switches to configure virtual network properties, such as security, performance, and availability.

Depending on the virtual switch type, networking policies can be applied at different levels of the virtual switch.

Virtual Switch Type	Set Default Policy At	Override Default Policy At
vSphere Standard Switch	Standard switch level	Port group level
vSphere Distributed Switch	Distributed port group level	Individual port level

The networking security policy provides protection against MAC address impersonation and unwanted port scanning.

The traffic shaping policy is useful when you want to limit the amount of traffic to a VM or a group of VMs.

Use the teaming and failover policy to determine the following information:

- How the network traffic of VMs and VMkernel adapters connected to the switch is distributed between physical adapters.
- How the traffic should be rerouted if an adapter fails.

5-21 Configuring Security Policies

As an administrator, you can define security policies at both the standard switch level and the port group level:

- **Promiscuous mode**: Allow or disallow all traffic to be forwarded, regardless of the destination.
- **MAC address changes**: Accept or reject inbound traffic when the MAC address is altered by the guest.
- **Forged transmits**: Accept or reject outbound traffic when the MAC address is altered by the guest.

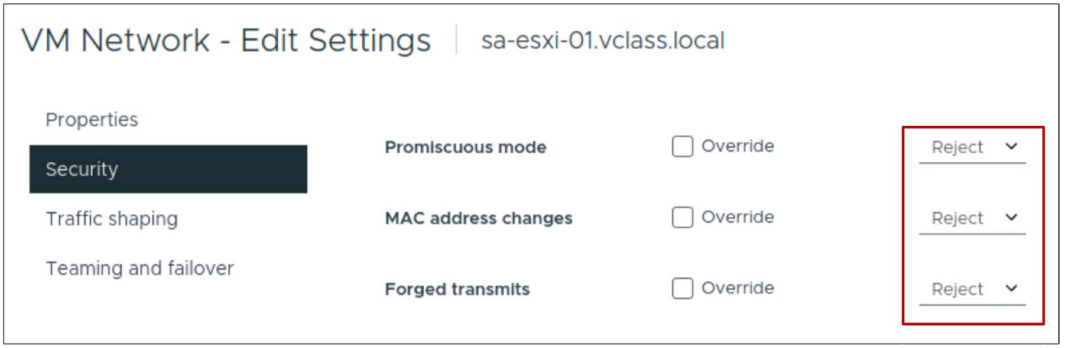

For most cases, the recommended setting for all policies is **Reject**.

The network security policy contains the following exceptions:

- **Promiscuous mode**: **Promiscuous mode** allows a virtual switch or port group to forward all traffic, regardless of their destinations. The default is **Reject**.
- **MAC address changes**: If this option is set to **Reject** and the guest attempts to change the MAC address assigned to the virtual NIC, it stops receiving frames. The default is **Reject**. Keep the default setting to help protect against attacks launched by a rogue guest operating system.
- **Forged transmits**: A frame's source address field may become altered by the guest and contain a MAC address other than the assigned virtual NIC MAC address. You can set the **Forged Transmits** parameter to accept or reject such frames. The default is **Reject**. Keep the default setting to help protect against attacks launched by a rogue guest operating system.

In general, these policies give you the option of disallowing certain behaviors that might compromise security. For example, a hacker might use a promiscuous mode device to capture network traffic for unscrupulous activities. Or, someone might impersonate a node and gain unauthorized access by spoofing its MAC address.

Change the setting to **Accept** only for specialized use cases. Examples:

- Set **Promiscuous mode** to **Accept** to use an application in a VM that analyzes or sniffs packets, such as a network-based intrusion detection system.
- Set **MAC address changes** and **Forged transmits** to **Accept** if your applications change the mapped MAC address, as do some guest operating system-based firewalls.

5-22 Traffic-Shaping Policies

Network traffic shaping is a mechanism for limiting a virtual machine's consumption of available network bandwidth.

Average rate, peak rate, and burst size are configurable.

Network traffic shaping is deactivated by default.

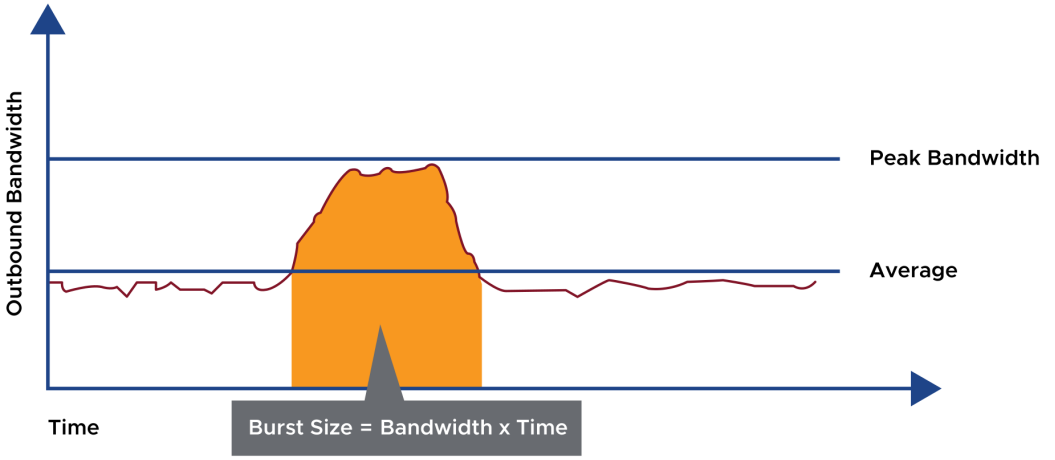

The ESXi host shapes traffic by establishing parameters for the following traffic characteristics:

- Average bandwidth (Kbps): Establishes the number of kilobits per second to allow across a port, averaged over time. The average bandwidth is the allowed average load.
- Peak bandwidth (Kbps): The maximum number of kilobits per second to allow across a port when it is sending a burst of traffic. This number tops the bandwidth that is used by a port whenever the port is using the burst bonus that is configured using the Burst size parameter.
- Burst size (KB): The maximum number of kilobytes to allow in a burst. If this parameter is set, a port might gain a burst bonus if it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than the average bandwidth, the port might be allowed to temporarily transmit data at a faster speed if a burst bonus is available. This parameter tops the number of kilobytes that have accumulated in the burst bonus and so transfers at a faster speed.

5-23 Configuring Outbound Traffic Shaping

A traffic-shaping policy is defined by average bandwidth, peak bandwidth, and burst size.

Parameters apply to each virtual NIC in the standard switch.

On a standard switch, traffic shaping controls only outbound traffic. Outbound traffic travels from the VMs to the virtual switch and out onto the physical network.

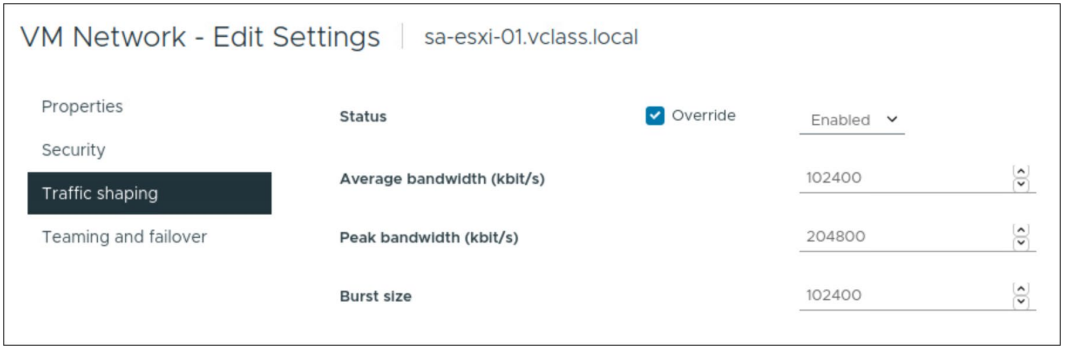

A virtual machine's network bandwidth can be controlled by activating the network traffic shaper.

The slide shows activating traffic shaping on a standard switch. The network traffic shaper, when used on a standard switch, shapes only outbound network traffic. To control inbound traffic, use a load-balancing system or turn on rate-limiting features on your physical router.

5-24 Configuring NIC Teaming and Failover

With NIC teaming, you can increase the network capacity of a port group by including two or more physical NICs in a team.

Add the physical NICs (or uplinks) to the Active uplinks group.

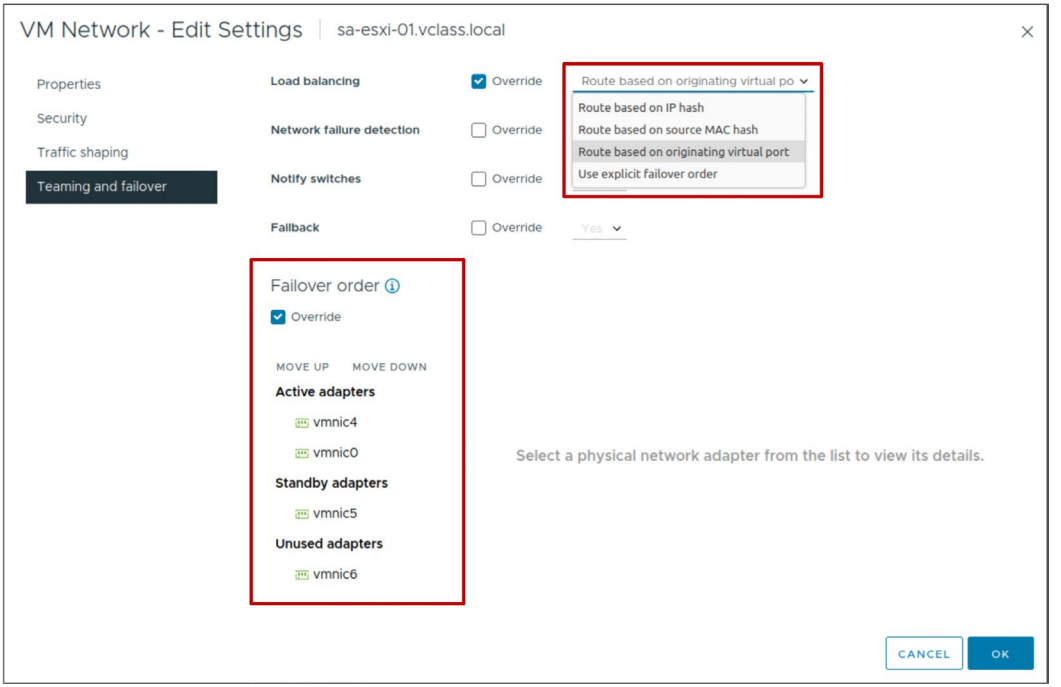

VM traffic is load balanced across the Active uplinks using the selected Load balancing option.

NIC teaming increases the network bandwidth of the switch and provides redundancy. To determine how the traffic is rerouted when an adapter fails, you include physical NICs in a failover order.

The load-balancing policy determines how network traffic is distributed between the network adapters in a NIC team. Depending on the needs and capabilities of your environment, select load-balancing algorithms to have the virtual switch distribute the network traffic between the physical NICs in a team. Virtual switches only load balance outgoing traffic. Incoming traffic is controlled by the load-balancing policy on the physical switch.

A failover order determines which links are active during normal operations, and which links are active in the event of a failover. You can customize the following adapters in a failover order list:

- Active: Use the NICs in this group whenever the NIC connectivity is up and active.
- Standby: Use a NIC in this group if one of the NICs is down.
- Unused: Do not use this NIC. NICs are placed in this group to reserve them for emergencies. They can be moved to the Active group when needed.

5-25 Load Balancing Method: Originating Virtual Port ID

With the load balancing method that is based on the originating virtual port ID, a virtual machine's outbound traffic is mapped to a specific physical NIC.

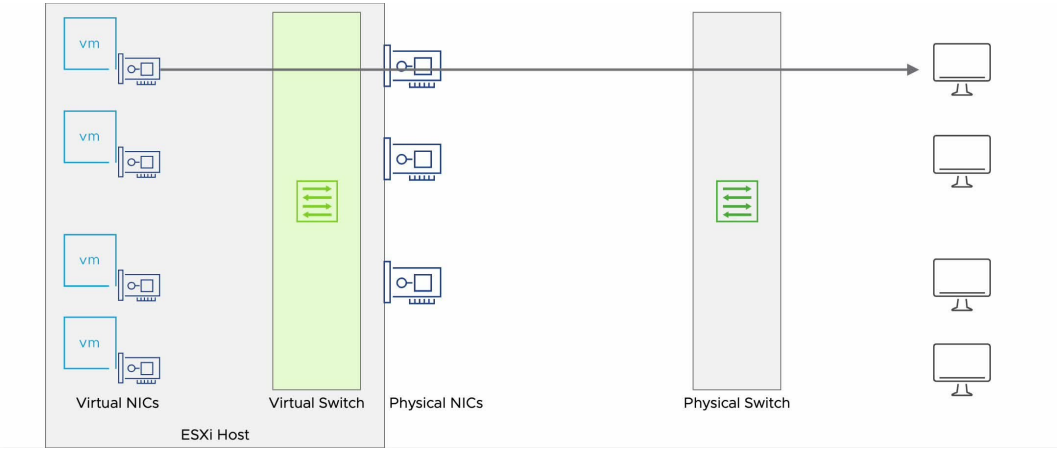

To play the animation, go to https://vmware.bravais.com/s/JRjAEOiEo0jRe4mD7IWd

The load-balancing method that uses the originating virtual port ID is simple and fast and does not require the VMkernel to examine the frame for the necessary information. The NIC is determined by the ID of the virtual port to which the VM is connected. With this method, no single-NIC VM gets more bandwidth than what can be provided by a single physical adapter.

This method has advantages:

- Traffic is evenly distributed if the number of virtual NICs is greater than the number of physical NICs in the team.
- Algorithm overhead is low because, in most cases, the virtual switch calculates uplinks for the VM only once.
- No changes on the physical switch are required.

This method also has disadvantages:

- The virtual switch does not take into account the traffic load on the uplinks. The virtual switch does not load balance the traffic to uplinks that are used less.
- The bandwidth that is available to a VM is limited to the speed of the uplink that is associated with the relevant port ID, unless the VM has more than one virtual NIC.

5-26 Load Balancing Method: Source MAC Hash

A virtual machine's outbound traffic, when load balanced using the source MAC hash method, is mapped to a specific physical NIC based on the virtual NIC's MAC address.

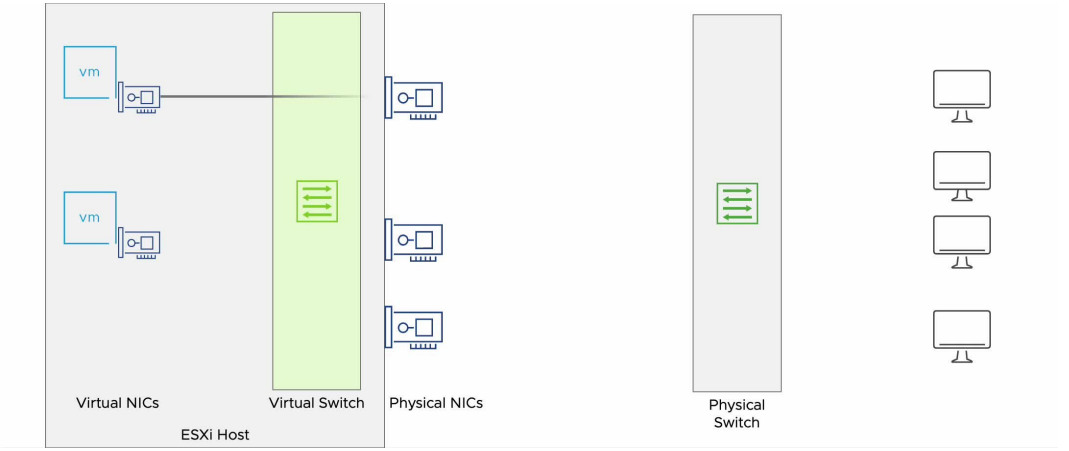

To play the animation, go to https://vmware.bravais.com/s/5kWE2uiGtaCVLFwz8S3z

The load balancing method based on source MAC hash has low overhead and is compatible with all switches, but it might not spread traffic evenly across all the physical NICs. In addition, no single-NIC virtual machine gets more bandwidth than a single physical adapter can provide.

This method has advantages:

- VMs use the same uplink because the MAC address is static. Powering a VM on or off does not change the uplink that the VM uses.
- No changes on the physical switch are required.

This method has disadvantages:

- The bandwidth that is available to a VM is limited to the speed of the uplink that is associated with the relevant port ID, unless the VM uses multiple source MAC addresses.
- Algorithm overhead is higher than with a route based on the originating virtual port because the virtual switch calculates an uplink for every packet.
- The virtual switch is not aware of the load of the uplinks, so uplinks might become overloaded.

5-27 Load Balancing Method: Source and Destination IP Hash

With the IP-based load balancing method, a NIC for each outbound packet is selected based on its source and destination IP addresses.

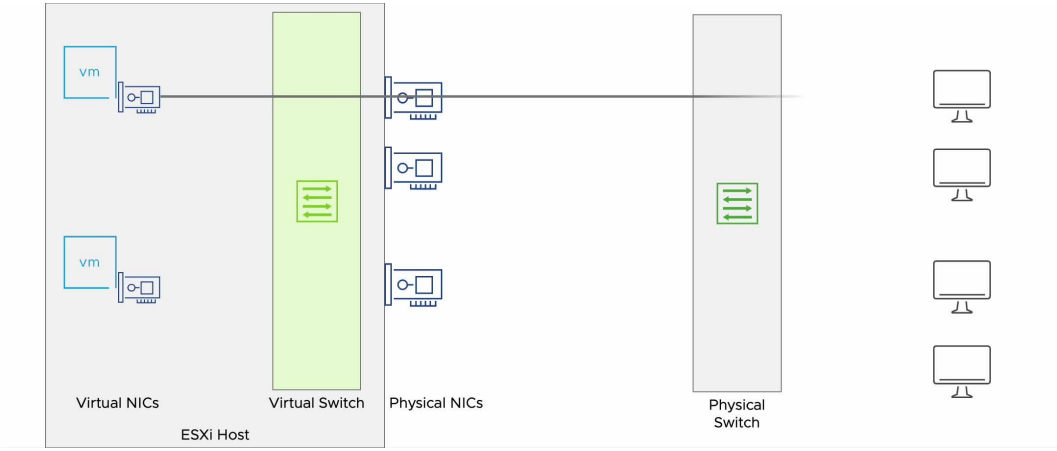

To play the animation, go to https://vmware.bravais.com/s/Z4td6qx8jCjFtYLycPl0

The IP-based method requires 802.3ad link aggregation support or an EtherChannel on the switch. The Link Aggregation Control Protocol (LACP) is a method to control the bundling of several physical ports to form a single logical channel. LACP is part of the IEEE 802.3ad specification.

EtherChannel is a port trunking technology that is used primarily on Cisco switches. With this technology, you can group several physical Ethernet links to create one logical Ethernet link for providing fault tolerance and high-speed links between switches, routers, and servers.

With this method, a single-NIC virtual machine might use the bandwidth of multiple physical adapters.

The IP-based load balancing method only affects outbound traffic. For example, a VM might select a particular NIC to communicate with a particular destination VM. The return traffic might not be handled on the same NIC as the outbound traffic, but by another NIC in the same NIC team.

This method has advantages:

- The load is more evenly distributed compared to the route based on the originating virtual port and the route based on source MAC hash because the virtual switch calculates the uplink for every packet.
- VMs that communicate with multiple IP addresses have a potentially higher throughput.

This method has disadvantages:

- Algorithm overhead is the highest compared to the other load balancing algorithms.
- The virtual switch is not aware of the actual load of the uplinks.
- Changes on the physical network are required.
- The method is complex to troubleshoot.

5-28 Detecting and Handling Network Failure

Network failures are monitored and detected by the VMkernel. The VMkernel monitors the link state and performs beacon probing (if selected) on one second intervals to ensure network uptime.

If the VMkernel determines a network failure, the VMkernel notifies physical switches of changes in the physical location of a MAC address.

Failover is implemented by the VMkernel based on configurable parameters:

- Failback: How the physical adapter is returned to active duty after recovering from failure.
- Load-balancing option: Use explicit failover order. Always use the vmnic uplink at the top of the active adapter list.

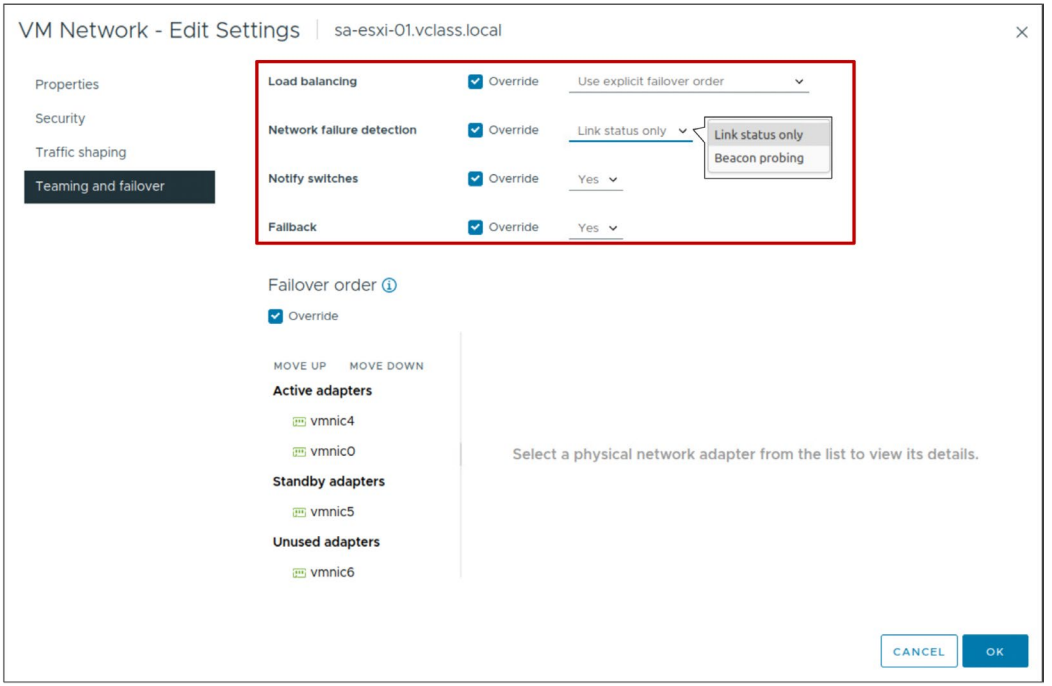

Monitoring the link status that is provided by the network adapter detects failures such as cable pulls and physical switch power failures. This monitoring does not detect configuration errors, such as a physical switch port being blocked by the Spanning Tree Protocol or misconfigured VLAN membership. This method cannot detect upstream, nonphysically connected switch or cable failures.

Beacon probing introduces a 62-byte packet load approximately every 1 second per physical NIC. When beacon probing is activated, the VMkernel sends out and listens for probe packets

on all NICs that are configured as part of the team. This technique can detect failures that link-status monitoring alone cannot. A specific network topology is required for beacon probing to work. Consult your switch manufacturer to verify the support of beacon probing in your environment. For information on beacon probing, see VMware knowledge base article 1005577 at http://kb.vmware.com/kb/1005577.

A physical switch can be notified by the VMkernel whenever a virtual NIC is connected to a virtual switch. A physical switch can also be notified whenever a failover event causes a virtual NIC's traffic to be routed over a different physical NIC. The notification is sent over the network to update the lookup tables on physical switches. In most cases, this notification process is beneficial because, without it, VMs experience greater latency after failovers and vSphere vMotion operation.

Do not set this option when the VMs connected to the port group are running Microsoft Network Load Balancing (NLB) in unicast mode. NLB in multicast mode is unaffected. For more information about the NLB issue, see VMware knowledge base article 1556 at http://kb.vmware.com/kb/1556.

When using explicit failover order, always use the highest order uplink from the list of active adapters that pass failover-detection criteria.

The failback option determines how a physical adapter is returned to active duty after recovering from a failure:

- If **Failback** is set to **Yes**, the failed adapter is returned to active duty immediately on recovery, displacing the standby adapter that took its place at the time of failure.
- If **Failback** is set to **No**, a failed adapter is left inactive even after recovery, until another currently active adapter fails, requiring its replacement.

5-29 Physical Network Considerations

Your virtual networking environment relies on the physical network infrastructure. As a vSphere administrator, you should discuss your vSphere networking needs with your network administration team.

The following issues are topics for discussion:

- Number of physical switches
- Network bandwidth required
- Physical switch configuration support for 802.1Q, for VLAN tagging
- Physical switch configuration support for NIC teaming: 802.3ad, Link Aggregation Control Protocol (LACP), or EtherChannel
- Network port security
- Link Layer Discovery Protocol (LLDP) and Cisco Discovery Protocol (CDP) and their operation modes, such as listen, broadcast, listen and broadcast, and disabled

5-30 Activity: Networking Security Policy (1)

Which statement accurately describes Promiscuous mode when it is set to Accept? (Choose one.)

◊ The ESXi host is allowed to drop network packets that seem suspicious.

◊ An administrator provides enhanced security to the virtual switch.

◊ An administrator can use a network-based intrusion detection system in a VM.

◊ The guest OS is given permission to change the VM's MAC address.

5-31 Activity: Networking Security Policy (2)

Which statement accurately describes Promiscuous mode when it is set to Accept? (Choose one.)

◊ The ESXi host is allowed to drop network packets that seem suspicious.

◊ An administrator provides enhanced security to the virtual switch.

✓ An administrator can use a network-based intrusion detection system in a VM.

◊ The guest OS is given permission to change the VM's MAC address.

When promiscuous mode is set to Accept, an administrator can run legitimate software such as a network-based intrusion detection system. However, when this mode is enabled, the virtual switch is vulnerable to security breaches.

5-32 Activity: Traffic Shaping Policy (1)

Which statement accurately describes the traffic shaping policy configuration? (Choose one.)

◊ Traffic shaping is activated on the entire standard switch.

◊ The traffic shaping policy for the TestDev port group overrides the policy defined on the standard switch.

◊ The bandwidth used for normal operation by VMs on the TestDev port group is 100 Mbps.

◊ The Accounting port group is subject to traffic shaping, with an average bandwidth of 1 Gbps.

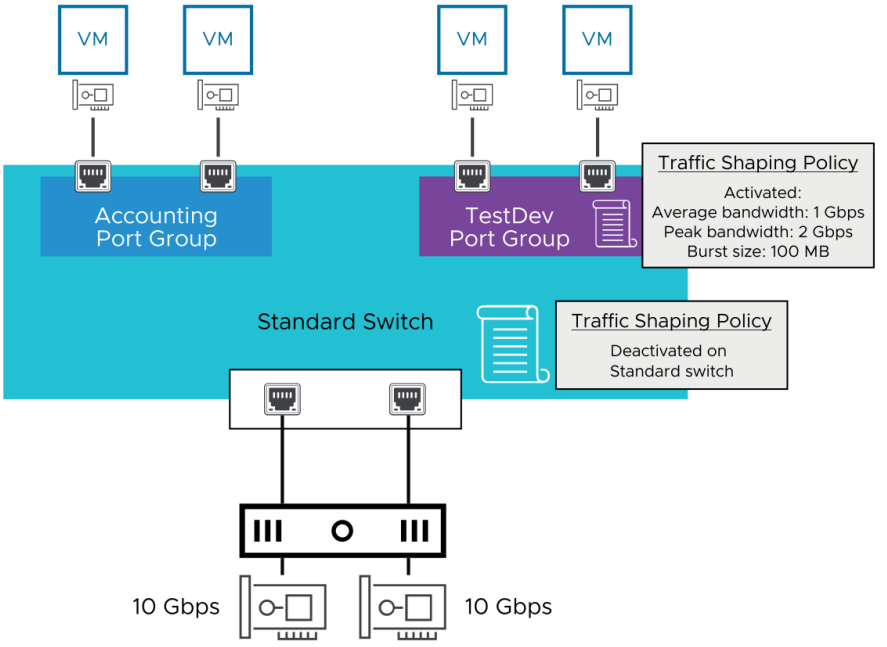

5-33 Activity: Traffic Shaping Policy (2)

Which statement accurately describes the traffic shaping policy configuration? (Choose one.)

◊ Traffic shaping is activated on the entire standard switch.

✓ The traffic shaping policy for the TestDev port group overrides the policy defined on the standard switch.

◊ The bandwidth used for normal operation by VMs on the TestDev port group is 100 Mbps.

◊ The Accounting port group is subject to traffic shaping, with an average bandwidth of 1 Gbps.

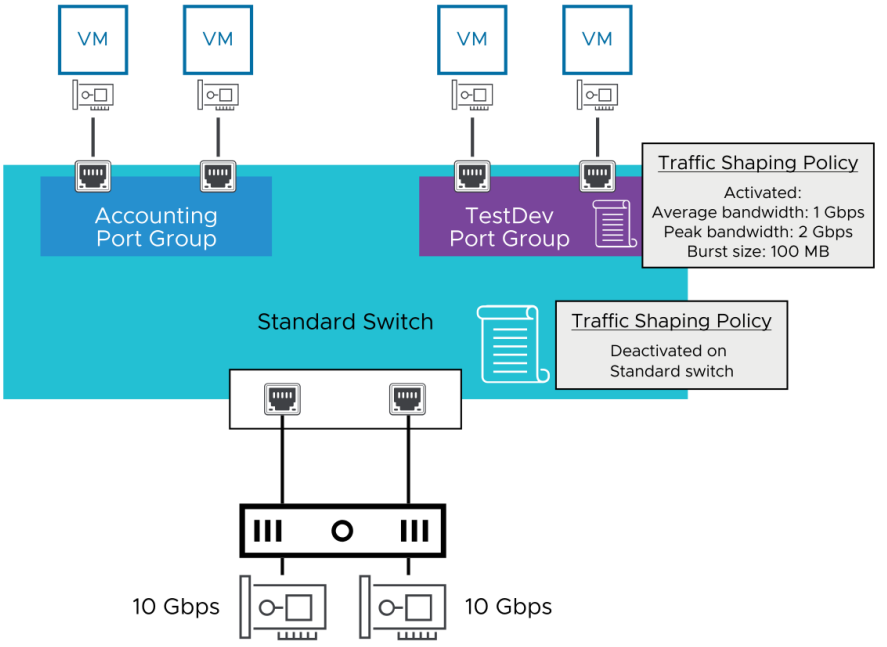

5-34 Activity: NIC Teaming and Failover Policy (1)

The load balancing method called Originating Virtual Port ID is only available on distributed switches.

◊ True

◊ False

5-35 Activity: NIC Teaming and Failover Policy (2)

The load balancing method based on the originating virtual port ID is only available on distributed switches.

◊ True

✓ False

The load balancing method based on physical NIC load is the only method supported on distributed switches.

The load balancing method that is only available on distributed switches is the **Route based on physical NIC load** option. This method ensures that physical NIC capacity in a NIC team is optimized.

5-36 Review of Learner Objectives

- Explain how to set security policies for a virtual switch
- Explain how to set traffic shaping policies for a virtual switch
- Explain how to set NIC teaming and failover policies for a virtual switch

5-37 Lesson 3: vSphere Distributed Switches

5-38 Learner Objectives

- Recognize the difference between standard switches and distributed switches
- Identify the benefits and features of distributed switches
- Create a distributed switch

5-39 About Distributed Switches

A distributed switch functions as a single virtual switch across all associated hosts. Distributed switches have several benefits over standard switches:

- Distributed switches centralize the virtual network administration, and simplifies the data center administration.
- Distributed switch ports are statically assigned by vCenter and offer more granular control over network statistics and policies.

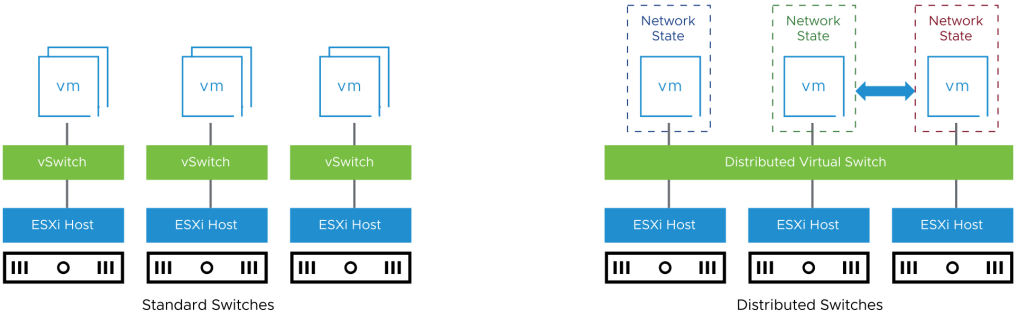

Standard switches are configured at the host level. Distributed switches are configured at the data center level, which gives distributed switches the following advantages:

- Data center setup and administration are simplified through this centralized network configuration. For example, adding a host to a cluster and making it compatible with vSphere vMotion is much easier than with a standard switch.
- Distributed ports migrate with their VMs. For example, when you migrate a VM with vSphere vMotion, the distributed port statistics and policies move with the VM, which simplifies debugging and troubleshooting.

5-40 Distributed Switch Architecture

Managed by vCenter, a distributed switch is a logical entity that you can use to create and maintain a consistent virtual networking configuration throughout all your ESXi hosts.

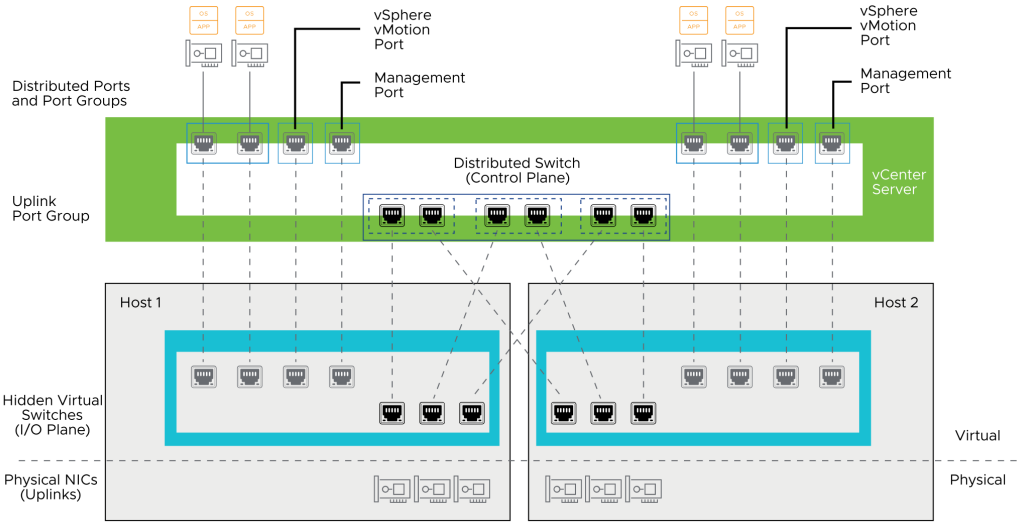

A distributed switch moves network management components to the data center level.

The distributed switch architecture consists of the control plane and the I/O plane.

- The control plane resides in vCenter. The control plane configures distributed switches, distributed port groups, distributed ports, uplinks, NIC teaming, and so on. The control plane also coordinates the migration of the ports and is responsible for the switch configuration.
- The I/O plane is implemented as a hidden virtual switch in the VMkernel of each ESXi host. The I/O plane manages the I/O hardware on the host and is responsible for forwarding packets. vCenter oversees the creation of these hidden virtual switches.

Each distributed switch includes distributed ports. You can connect any networking entity, such as a VM or a VMkernel interface, to a distributed port. vCenter stores the state of distributed ports in the vCenter database.

With a distributed port group, you can logically group distributed ports to simplify configuration. A distributed port group specifies port configuration options for each member port on a distributed switch. Ports can also have their own unique configuration.

Uplinks are abstractions of vmnics from multiple hosts to a single distributed switch. An uplink is to a distributed switch what a vmnic is to a standard switch. Two VMs on different hosts can communicate with each other only if both VMs have uplinks in the same broadcast domain.

5-41 Standard and Distributed Switches: Shared Features

Standard and distributed switches share some features.

Feature	Standard Switch	Distributed Switch
Layer 2 switch	✓	✓
VLAN segmentation	✓	✓
802.1Q tagging	✓	✓
IPv4 and IPv6 support	✓	✓
NIC teaming	✓	✓
Outbound traffic shaping	✓	✓

5-42 Distributed Switch Features

Distributed switches have several features that standard switches do not have.

Feature	Standard Switch	Distributed Switch
Inbound traffic shaping		✓
Configuration backup and restore		✓
Private VLANs		✓
Link Aggregation Control Protocol		✓
Data center level management		✓
vSphere vMotion migration of network state		✓
Network I/O Control		✓
Per-port policy settings		✓
Port state monitoring		✓
NetFlow		✓
Port mirroring		✓
Support for NSX		✓

During a vSphere vMotion migration, a distributed switch tracks the virtual networking state (for example, counters and port statistics) VMs moving between hosts. This tracking provides a consistent view of a virtual network interface, regardless of the VM location or vSphere vMotion migration history.

Tracking simplifies network monitoring and troubleshooting activities when migrating VMs between hosts using vSphere vMotion.

5-43 Viewing Distributed Switches

In the vSphere Client, you can view the distributed switch configuration using the Topology pane in the **Configure** tab.

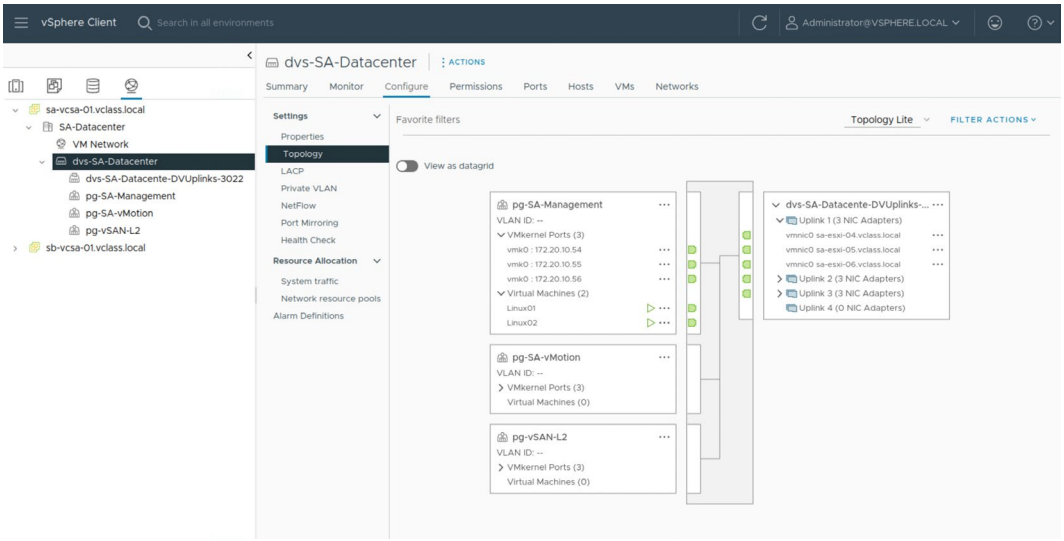

5-44 Discovery Protocols

Switch discovery protocols help network administrators gather configuration and connection information about physical or virtual switches.

vSphere supports the following discovery protocols:

- Cisco Discovery Protocol (CDP): For vSphere standard switches and distributed switches connected to Cisco physical switches
- Link Layer Discovery Protocol (LLDP): A vendor-neutral protocol for distributed switches only

Standard switches can be configured to use CDP.

Distributed switches can use CDP or LLDP.

Switch discovery protocols help network administrators determine the capabilities of a network device. Such information might help in troubleshooting network problems.

Cisco Discovery Protocol (CDP) was developed by Cisco Systems to broadcast connected device information at network layer 2. CDP is supported in vSphere since version 4.0.

Link Layer Display Protocol (LLDP) supports the standards-based IEEE 802.1AB discovery protocol and is available on distributed switches only.

Network devices use CDP or LLDP to advertise their identity, capabilities, and neighbors on a network.

5-45 Configuring CDP or LLDP

With CDP or LLDP enabled, you can configure a virtual switch for different modes of operation:

- Listen: Information is received from the physical switches.
- Advertise: Information is sent to the physical switches.
- Both: Information is both sent to and received from the physical switches.

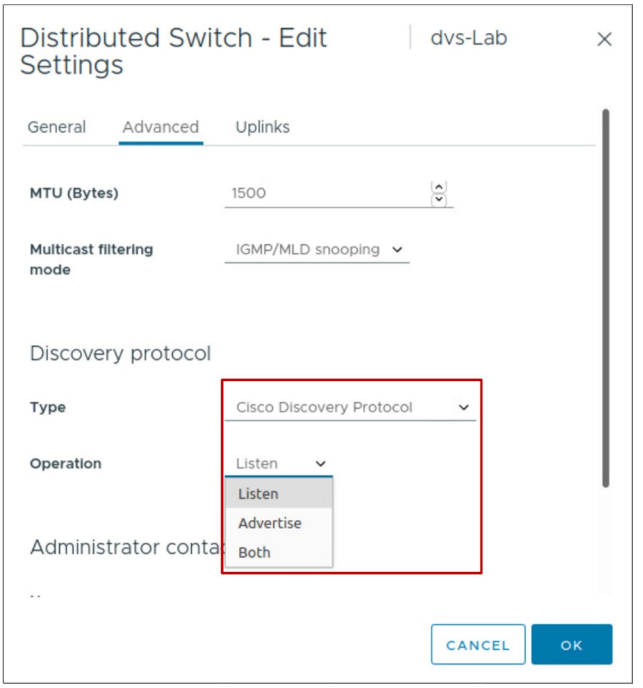

With CDP and LLDP, the vSphere Client can identify properties of a physical switch, such as switch name, port number, and port speed or duplex settings. You can also configure CDP or LLDP so that information about physical adapters and ESXi host names is passed to the CDP or LLDP compatible switches.

You can configure the discovery protocol to use one of the following modes of operation:

- Listen (default): The ESXi host detects and displays information about the associated physical switch port, but information about the virtual switch is not available to the physical switch administrator.
- Advertise: The ESXi host provides information about the virtual switch to the physical switch administrator but does not detect and display information about the physical switch.
- Both: The ESXi host detects and displays information about the associated physical switch and provides information about the virtual switch to the physical switch administrator.

You can use the `esxcli` command to enable CDP on a standard switch.

5-46 About Port Binding

◊ Port binding determines when and how a VM virtual NIC is assigned to a virtual switch port. Port binding is configured at the distributed port group level, and binding options include:

- Static binding (default): vCenter assigns a permanent port for the VM or VMkernel interface.
- Ephemeral: ESXi (not vCenter) assigns the port to the VM. The assigned port changes when the VM reboots.

◊ Port allocation options for static binding:

- Elastic (default): When all ports are assigned, a new set of eight ports is created.
- Fixed: No additional ports are created when all ports are assigned.

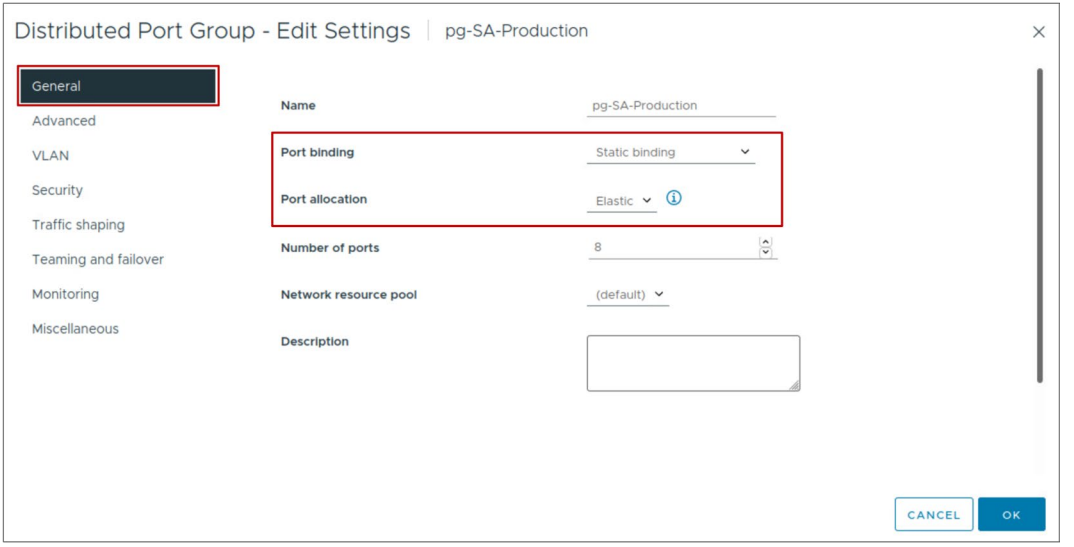

When you connect a VM to a port group that is configured with static binding, a port is immediately assigned and reserved for the VM, guaranteeing connectivity at all times. The port is disconnected only when the VM is removed from the port group. Static binding is recommended for general use.

If you select static binding, the default number of ports is set to eight. Elastic is the default port allocation setting.

With ephemeral binding, a port is created and assigned to a VM when the VM is powered on and its NIC is in a connected state. The port is deleted when the VM is powered off or the VM NIC is disconnected.

You can make ephemeral port assignments through the ESXi host and vCenter, providing flexibility in managing VM connections through the host when vCenter is down. Only an ephemeral binding allows you to modify VM network connections when vCenter is down. However, network traffic is unaffected by a vCenter failure, regardless of the port binding type.

Use ephemeral port groups only for recovery purposes when you want to provision ports directly on an ESXi host, bypassing vCenter, and not in any other case.

5-47 Configuring Inbound Traffic Shaping

Distributed switches support inbound traffic shaping and outbound traffic shaping.

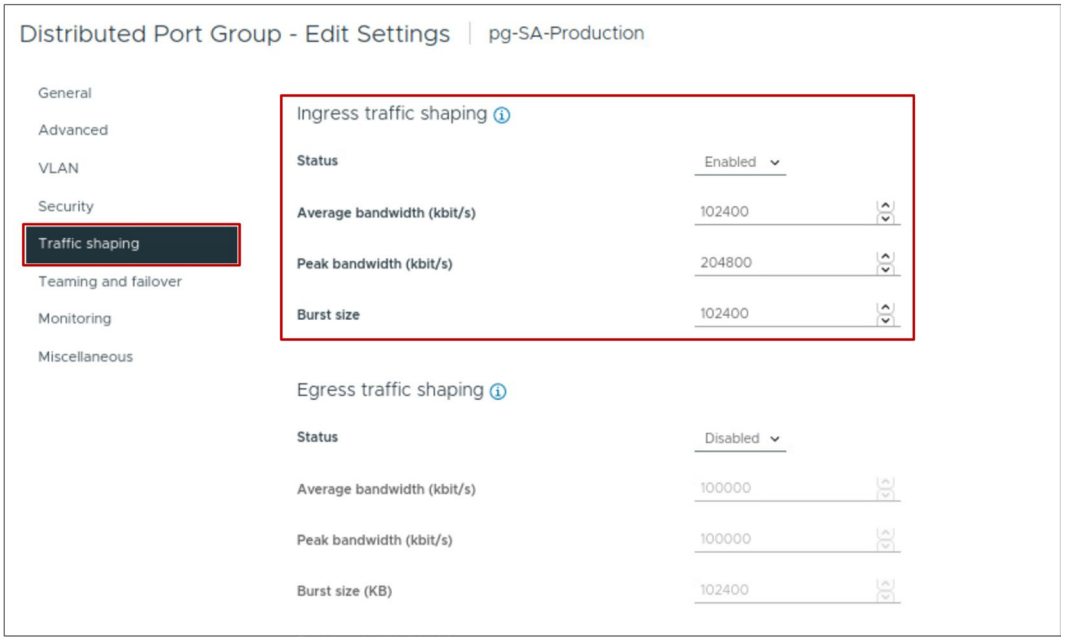

Where outbound (or egress) traffic shaping is supported by both standard switches and distributed switches, inbound (or ingress) traffic shaping is supported only by distributed switches.

Inbound traffic is traffic traveling from the physical network to the virtual switch to the VMs.

5-48 Load Balancing Method: Physical NIC Load

This method is supported only on distributed switches and is the recommended policy for distributed port groups.

Load balancing based on physical NIC load ensures that physical NIC capacity in a NIC team is optimized. This feature works in the following ways:

- It moves I/O flows among uplinks.
- A flow is moved only when the mean send or receive utilization of an uplink exceeds 75 percent of the capacity over a 30-second period.

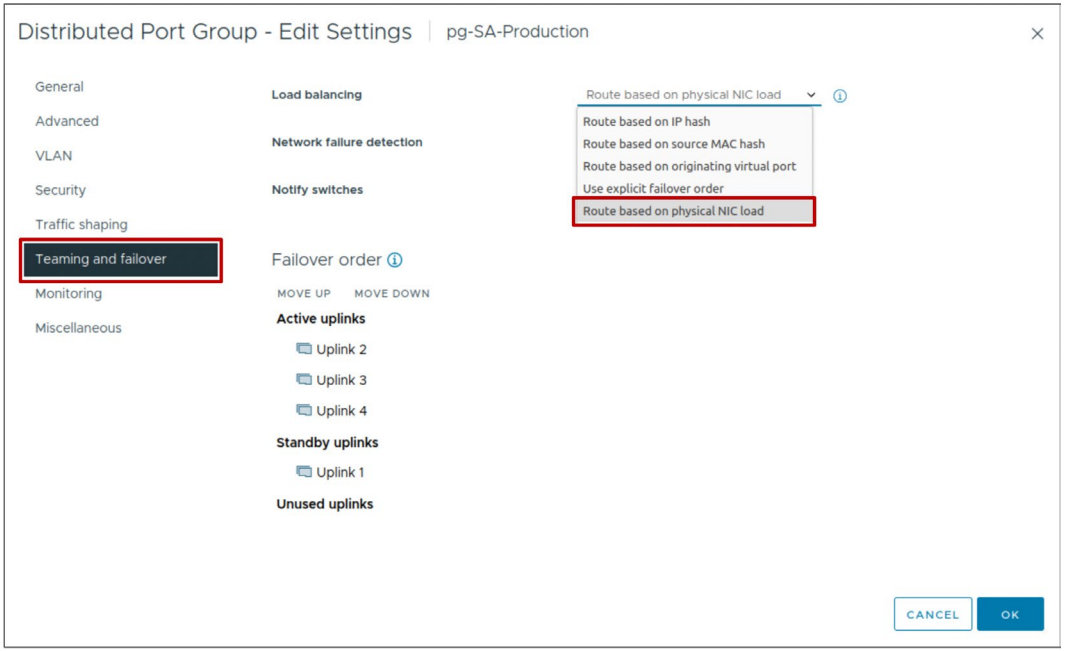

To use this method, edit the distributed port group settings, and select **Route based on physical NIC load**.

This NIC teaming method checks the real load of the uplinks and reduces the load on overloaded uplinks. No changes on the physical switch are required.

The distributed switch calculates uplinks for VMs by checking the VM port ID and the number of uplinks in the NIC team. The distributed switch tests the uplinks every 30 seconds. If the load of an uplink exceeds 75 percent of usage, the port ID of the VM with the highest I/O is moved to a different uplink.

Route based on physical NIC load is not the default teaming policy. You must configure the policy to use it on a distributed switch.

This method has advantages:

- Low algorithm overhead because the distributed switch calculates uplinks for virtual machines only once and checking the uplinks has minimal impact.
- The distributed switch is aware of the load of uplinks and takes care to reduce it if needed.
- No changes on the physical switch are required.

This method has disadvantages:

- The bandwidth that is available to virtual machines is limited to the uplinks that are connected to the distributed switch.

5-49 Lab 8: Configuring vSphere Distributed Switches

Create and configure a distributed switch:

1. Create a Distributed Switch
2. Add ESXi Hosts to the Distributed Switch
3. Verify Your Distributed Switch Configuration

5-50 Review of Learner Objectives

- Recognize the difference between standard switches and distributed switches
- Identify the benefits and features of distributed switches
- Create a distributed switch

5-51 Key Points

- Virtual switches can have the following connection types: VM ports, VMkernel port, and physical uplinks.
- A standard switch is a virtual switch configuration for a single host.
- A distributed switch provides functions that are similar to a standard switch. But the distributed switch defines a single configuration that is managed by vCenter and is shared across all associated hosts.
- You set networking policies on virtual switches to configure properties for security, performance, and availability.
- Network policies set at the standard switch level can be overridden at the port group level. Network policies set at the distributed switch port group level can be overridden at the individual port level.

Questions?

Module 6

Configuring vSphere Storage

6-2 Importance

Understanding the available storage options helps you set up your storage according to your cost, performance, and manageability requirements.

You can use shared storage for disaster recovery, high availability, and moving virtual machines between hosts.

6-3 Module Lessons

1. Storage Concepts
2. Fibre Channel Storage
3. iSCSI Storage
4. VMFS Datastores
5. NFS Datastores

6-4 Lesson 1: Storage Concepts

6-5 Learner Objectives

- Recognize vSphere storage technologies
- Identify types of datastores
- Recognize storage device naming conventions

6-6 About Datastores

A datastore is a logical storage unit that can use space on one or more physical storage devices.

Datastores are used to hold data such as VMs, VM templates, and ISO images.

vSphere supports the following types of datastores:

- VMFS
- NFS
- vSAN
- vSphere Virtual Volumes

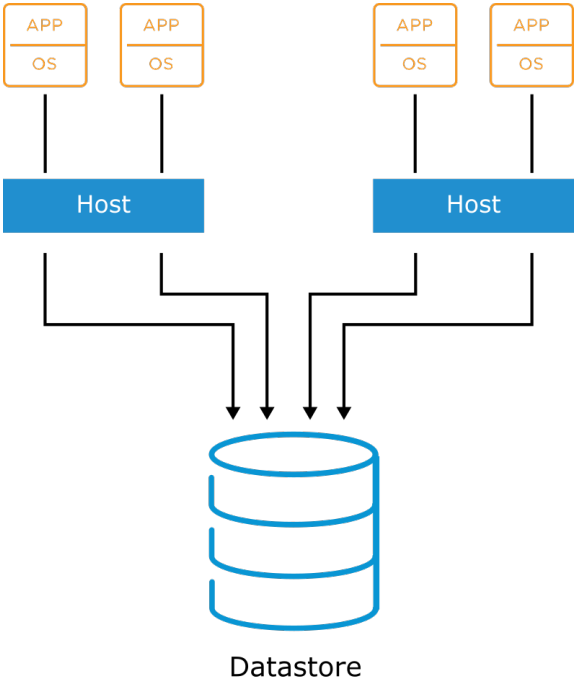

A datastore is a generic term for a container that holds files and objects. Datastores are logical containers, analogous to file systems, that hide the specifics of the underlying storage device and provide a uniform model for storing virtual machine files. A VM is stored as a set of files in its own directory or as a group of objects in a datastore.

You can display all datastores that are available to your hosts and analyze their properties.

6-7 Datastore Access Methods

vSphere datastores store and access data as blocks or files:

Block-backed storage:

- Stores data as blocks (a sequence of bytes)
- Used on local storage
- Used on Storage Area Networks (SANs) and accessed through either iSCSI or Fibre Channel
- Used by VMFS, vSAN, and vSphere Virtual Volumes datastores

File-backed storage:

- Stores data hierarchically in files and folders
- Used on network-attached storage (NAS)
- Used by NFS and vSphere Virtual Volumes datastores

6-8 Datastore Contents

Depending on the datastore type, contents can be stored in the form of files or objects.

File-based datastores:

- A VM consists of a set of files.
- Each VM has its own directory.
- VMFS and NFS datastores hold files.

Object-based datastores:

- A VM consists of a set of data containers called objects.
- vSAN and vSphere Virtual Volumes datastores hold objects.

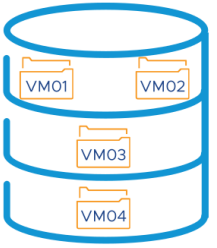

Datastore with Files

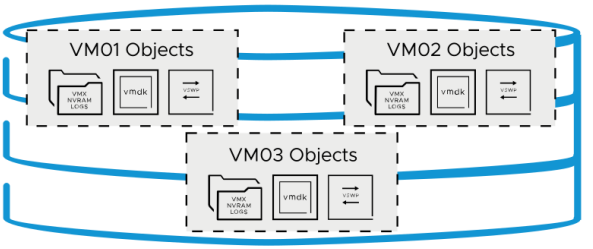

Datastore with Objects

In file-based datastores, a directory contains the VM's files, such as the configuration file, one or more virtual disk files, swap files, and so on.

In an object-based datastore, each VM consists of a VM configuration object, one or more virtual disk objects, a swap space object, and so on. An object is a data container. Each object on the datastore includes data, some metadata, and a unique ID.

6-9 Datastore Summary

To summarize, vSphere datastores can be categorized by its access method and its contents.

Datastore Type	Datastore Access Method	Datastore Contents
VMFS	Block access	Files
NFS	File access	Files
vSAN	Block access	Objects
vSphere Virtual Volumes	Block or file access	Objects

6-10 Storage Overview

ESXi hosts should be configured with shared access to datastores.

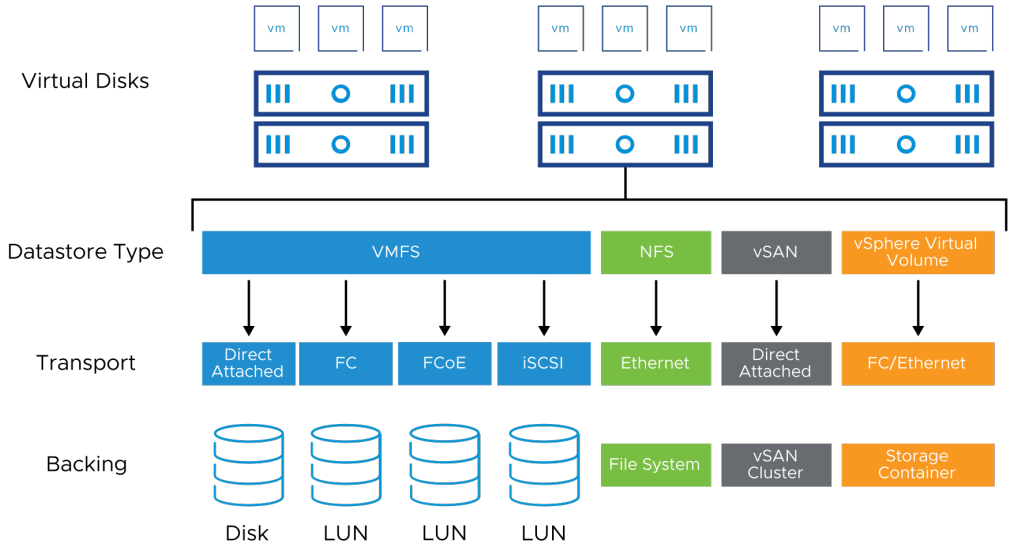

Depending on the type of storage that you use, datastores can be formatted with VMFS or NFS.

In the vSphere environment, ESXi hosts support several storage technologies:

- Direct-attached storage: Internal or external storage disks or arrays attached to the host through a direct connection instead of a network connection.
- Fibre Channel (FC): A high-speed transport protocol used for SANs. Fibre Channel encapsulates SCSI commands, which are transmitted between Fibre Channel nodes. In general, a Fibre Channel node is a server, a storage system, or a tape drive. A Fibre Channel switch interconnects multiple nodes, forming the fabric in a Fibre Channel network.
- FCoE: The Fibre Channel traffic is encapsulated into Fibre Channel over Ethernet (FCoE) frames. These FCoE frames are converged with other types of traffic on the Ethernet network.
- iSCSI: A SCSI transport protocol, providing access to storage devices and cabling over standard TCP/IP networks. iSCSI maps SCSI block-oriented storage over TCP/IP. Initiators, such as an iSCSI host bus adapter (HBA) in an ESXi host, send SCSI commands to targets, located in iSCSI storage systems.

- NAS: Storage shared over standard TCP/IP networks at the file system level. NAS storage is used to hold NFS datastores. The NFS protocol does not support SCSI commands.
- iSCSI, network-attached storage (NAS), and FCoE can run over high-speed networks providing increased storage performance levels and ensuring sufficient bandwidth. With sufficient bandwidth, multiple types of high-bandwidth protocol traffic can coexist on the same network. For more information about physical NIC support and maximum ports supported, see VMware Configuration Maximums at https://configmax.vmware.com.

6-11 Storage Device Naming Conventions

Storage devices are identified in several ways:

- Runtime name: Uses the vmhbaN:C:T:L convention. This name is not persistent through reboots.
- Target: Identifies the target address and port.
- LUN: A unique identifier designated to individual or collections of storage devices.

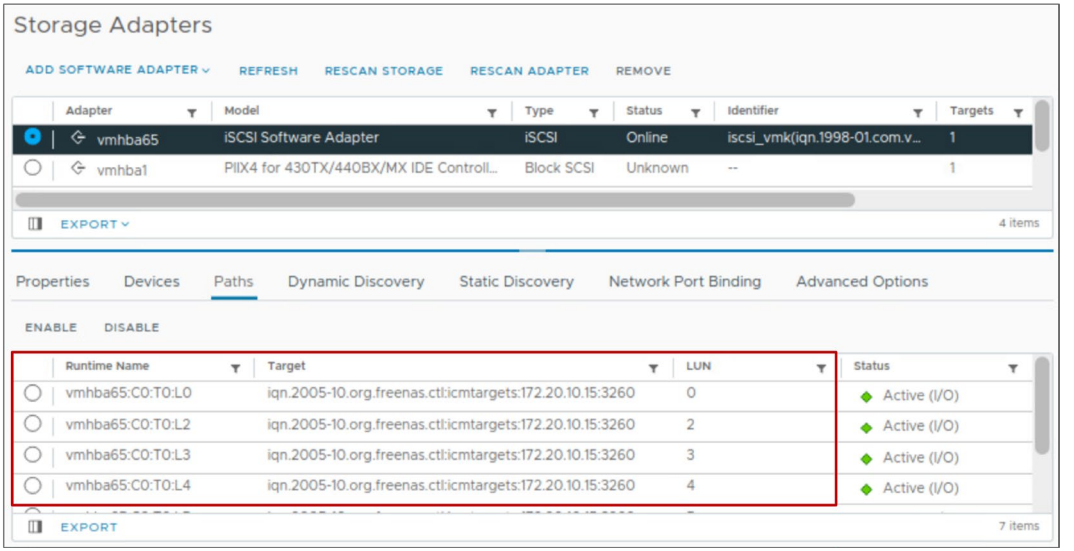

This slide shows examples of software iSCSI device names.

On the ESXi hosts, SCSI storage devices use various identifiers. Each identifier serves a specific purpose. For example, the VMkernel requires an identifier, generated by the storage device, which is guaranteed to be unique to each LUN. If the storage device cannot provide a unique identifier, the VMkernel must generate a unique identifier to represent each LUN or disk.

The following SCSI storage device identifiers are available:

- Runtime name: The name of the first path to the device. The runtime name is a user-friendly name that is created by the host after each reboot. It is not a reliable identifier for the disk device because it is not persistent. The runtime name might change if you add HBAs to the ESXi host. However, you can use this name when you use the command-line utilities to interact with the storage that an ESXi host recognizes.
- Target: A worldwide unique name for identifying the node. For example, iSCSI uses the IQN and EUI. IQN uses the iqn.yyyy-mm.naming-authority:unique name format.
- LUN: A logical unit number is addressed by the SCSI protocol or SAN protocols that encapsulate SCSI, such as iSCSI or Fibre Channel.

Storage device names appear in various panels in the vSphere Client.

6-12 Storage Protocol Overview

Each datastore uses a protocol with varying support features.

Datastore Type	Storage Protocol	Boot from SAN Support	vSphere vMotion Support	vSphere HA Support	vSphere DRS Support
VMFS	Fibre Channel	Yes	Yes	Yes	Yes
	FCoE	Yes	Yes	Yes	Yes
	iSCSI	Yes	Yes	Yes	Yes
	iSER/NVMe-oF (RDMA)	No	Yes	Yes	Yes
	DAS (SAS, SATA, NVMe)	N/A	Yes*	No	No
NFS	NFS	No	Yes	Yes	Yes
vSphere Virtual Volumes	FC/Ethernet (iSCSI, NFS)	No	Yes	Yes	Yes
vSAN Datastore	vSAN	No	Yes	Yes	Yes

Direct-attached storage (DAS) supports vSphere vMotion when combined with vSphere Storage vMotion.

Direct-attached storage, as opposed to SAN storage, is where many administrators install ESXi. Direct-attached storage is also ideal for small environments because of the cost savings associated with purchasing and managing a SAN. The drawback is that you lose many of the features that make virtualization a worthwhile investment, for example, balancing the workload on a specific ESXi host. Direct-attached storage can also be used to store noncritical data:

- CD/DVD ISO images
- Decommissioned VMs
- VM templates

In comparison, storage LUNs must be pooled and shared so that all ESXi hosts can access them. Shared storage provides the following vSphere features:

- vSphere vMotion
- vSphere HA
- vSphere DRS

Using shared SAN storage also provides robust features in vSphere:

- Central repositories for VM files and templates
- Clustering of VMs across ESXi hosts
- Allocation of large amounts (terabytes) of storage to your ESXi hosts

ESXi supports different methods of booting from the SAN to avoid handling the maintenance of additional direct-attached storage or if you have diskless hardware configurations, such as blade systems. If you set up your host to boot from a SAN, your host's boot image is stored on one or more LUNs in the SAN storage system. When the host starts, it boots from the LUN on the SAN rather than from its direct-attached disk.

For ESXi hosts, you can boot from software iSCSI, a supported independent hardware SCSI adapter, and a supported dependent hardware iSCSI adapter. The network adapter must support only the iSCSI Boot Firmware Table (iBFT) format, which is a method of communicating parameters about the iSCSI boot device to an operating system.

6-13 About vSphere Virtual Machine File System

ESXi hosts support vSphere Virtual Machine File System (VMFS) VMFS5 and VMFS6:

- Features supported by both VMFS5 and VMFS6:
  - Concurrent access to shared storage
  - Dynamic expansion
  - On-disk locking
- Features supported by VMFS6:
  - 4K native storage devices
  - Automatic space reclamation
  - 128 hosts per datastore

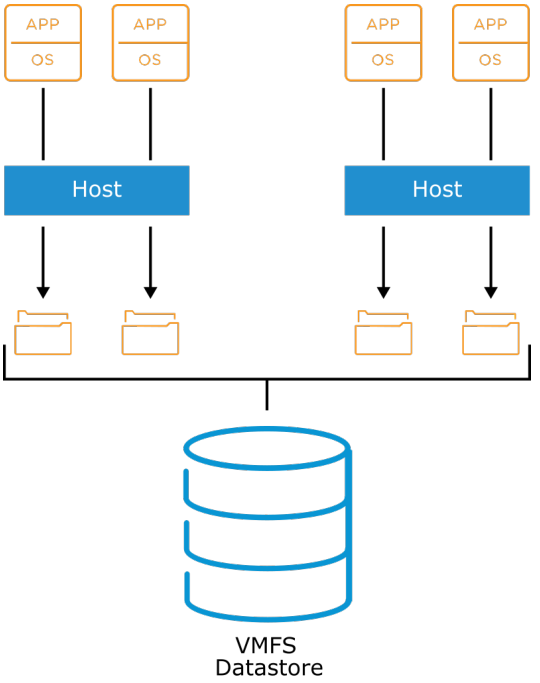

VMFS is a clustered file system where multiple ESXi hosts can read and write to the same storage device simultaneously. The clustered file system provides unique, virtualization-based services:

- Migration of running VMs from one ESXi host to another without downtime
- Automatic restarting of a failed VM on a separate ESXi host
- Clustering of VMs across various physical servers

Using VMFS, IT organizations can simplify VM provisioning by efficiently storing the entire VM state in a central location. Multiple ESXi hosts can access shared VM storage concurrently.

The size of a VMFS datastore can be increased dynamically when VMs residing on the VMFS datastore are powered on and running. A VMFS datastore efficiently stores both large and small files belonging to a VM. A VMFS datastore can support virtual disk files. A virtual disk file has a maximum of 62 TB. A VMFS datastore uses sub-block addressing to make efficient use of storage for small files.

VMFS provides block-level distributed locking to ensure that the same VM is not powered on by multiple servers at the same time. If an ESXi host fails, the on-disk lock for each VM is released and VMs can be restarted on other ESXi hosts.

On the slide, each ESXi host has two VMs running on it. The lines connecting the VMs to the VM disks (VMDKs) are logical representations of the association and allocation of the larger VMFS datastore. The VMFS datastore includes one or more LUNs. The VMs see the assigned storage volume only as a SCSI target from within the guest operating system. The VM contents are only files on the VMFS volume.

VMFS can be deployed on three kinds of SCSI-based storage devices:

- Direct-attached storage
- Fibre Channel storage
- iSCSI storage

A virtual disk stored on a VMFS datastore always appears to the VM as a mounted SCSI device. The virtual disk hides the physical storage layer from the VM's operating system.

For the operating system in the VM, VMFS preserves the internal file system semantics. As a result, the operating system running in the VM sees a native file system, not VMFS. These semantics ensure correct behavior and data integrity for applications running on the VMs.

6-14 About NFS

A Network File System (NFS) is a file-sharing protocol that ESXi hosts use to communicate with a network-attached storage (NAS) device.

NFS supports NFS 3 and 4.1 over TCP/IP.

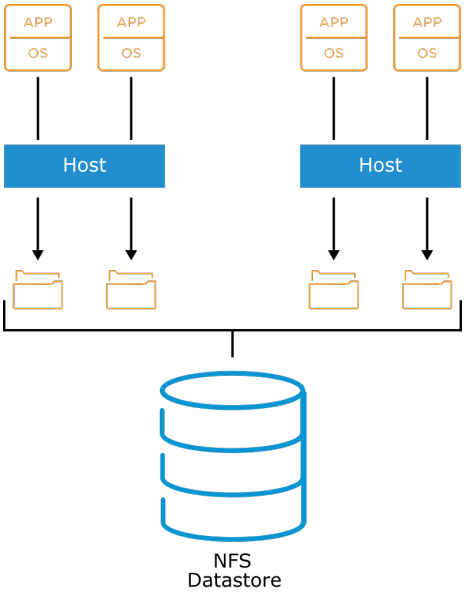

NAS is a specialized storage device that connects to a network and can provide file access services to ESXi hosts.

NFS datastores are treated like VMFS datastores because they can hold VM files, templates, and ISO images. In addition, like a VMFS datastore, an NFS volume allows the vSphere vMotion migration of VMs whose files reside on an NFS datastore. The NFS client built in to ESXi uses NFS protocol versions 3 and 4.1 to communicate with the NAS or NFS servers.

NFS 3 locking on ESXi does not use the Network Lock Manager (NLM) protocol, which is a standard protocol that is used to support the file locking of NFS-mounted files. Instead, VMware provides its own locking protocol. NFS 3 locks are implemented by creating lock files on the NFS server. NFS 4.1 uses server-side file locking.

Because NFS 3 and NFS 4.1 clients do not use the same locking protocol, you cannot use different NFS versions to mount the same datastore on multiple hosts. Accessing the same virtual disks from two incompatible clients might result in incorrect behavior and cause data corruption.

6-15 About vSAN

vSAN is a hypervisor-converged, software-defined storage solution for virtual environments that does not use traditional external storage.

By clustering host-attached solid-state drives (SSDs) and hard disk drives (HDDs), vSAN creates an aggregated datastore that is accessible to all the ESXi hosts in the vSAN cluster.

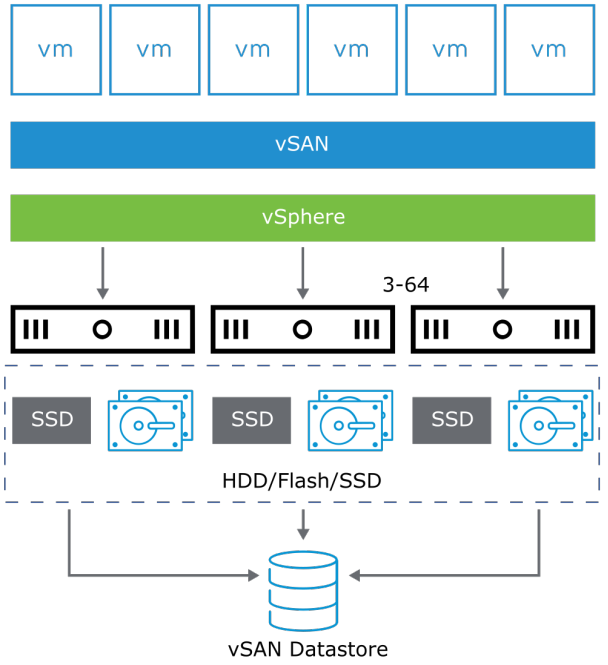

When vSAN is activated on a cluster, a single vSAN datastore is created. This datastore uses the storage components of each host in the cluster.

vSAN can be configured as hybrid or all-flash storage.

In a hybrid storage architecture, vSAN pools server-attached HDDs and SSDs to create a distributed shared datastore. This datastore abstracts the storage hardware to provide a software-defined storage tier for VMs. Flash is used as a read cache/write buffer to accelerate performance, and magnetic disks provide capacity and persistent data storage.

Alternately, vSAN can be deployed as an all-flash storage architecture in which flash devices are used as a write cache. SSDs provide capacity, data persistence, and consistent, fast response times. In the all-flash architecture, the tiering of SSDs results in a cost-effective implementation: a write-intensive, enterprise-grade SSD cache tier and a read-intensive, lower-cost SSD capacity tier.

6-16 About vSphere Virtual Volumes

vSphere Virtual Volumes provides several functionalities:

- Native representation of VMDKs on SAN/NAS: No LUNs or volume management
- Works with existing SAN/NAS systems
- A new control path for data operations at the VM and VMDK level
- Snapshots, replications, and other operations at the VM level on external storage
- Automates control of per-VM service levels by using storage policies
- Standard access to storage with the vSphere API for Storage Awareness protocol endpoint
- Storage containers that span an entire array

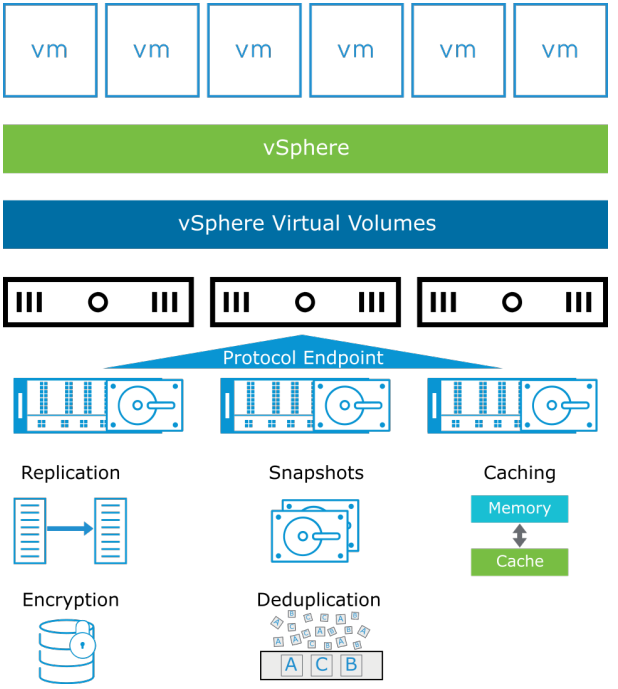

vSphere Virtual Volumes virtualizes SAN and NAS devices by abstracting physical hardware resources into logical pools of capacity.

vSphere Virtual Volumes provides the following benefits:

- Lower storage cost
- Reduced storage management overhead
- Greater scalability
- Better response to data access and analytical requirements

6-17 About Raw Device Mapping

Although not a datastore, raw device mapping (RDM) gives a VM direct access to a physical LUN.

The mapping file (`-rdm.vmdk`) that points a VM to a LUN must be stored on a VMFS datastore.

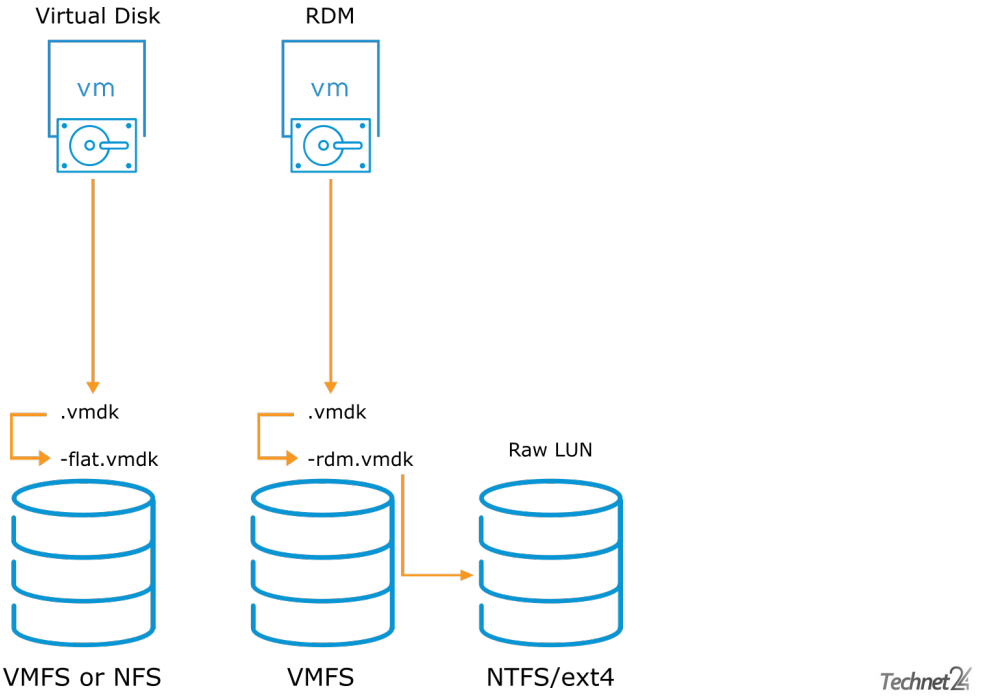

Raw device mapping (RDM) is a file stored in a VMFS volume that acts as a proxy for a raw physical device.

Instead of storing VM data in a virtual disk file that is stored on a datastore, you can store the guest operating system data directly on a raw LUN. Storing the data in this way is useful if you run applications in your VMs that must know the physical characteristics of the storage device. By mapping a raw LUN, you can use existing SCSI commands to manage storage for the disk.

Use RDM when a VM must interact with a real disk on the SAN. This condition occurs when you make disk array snapshots or have a large amount of data that you do not want to move onto a virtual disk as a part of a physical-to-virtual conversion.

You can use RDMs in virtual compatibility or physical compatibility modes. Virtual compatibility mode allows the RDM to behave as if it were a virtual disk, so you can use such features as taking snapshots, cloning, and so on. The RDM pointer file uses the extension `-rdm.vmdk`.

Physical compatibility mode allows the guest operating system to access the hardware directly. Physical compatibility is useful if you are using SAN-aware applications on the virtual machine. The RDM pointer file uses the extension –`rdmp.vmdk`.

6-18 Physical Storage Considerations

Before implementing your vSphere environment, discuss the storage needs with your storage administration team. Consider the following factors:

- LUN sizes
- I/O bandwidth required by your applications
- I/O requests per second that a LUN is capable of
- Disk cache parameters
- Zoning and masking
- Multipathing setting for your storage arrays (active-active or active-passive)
- Export properties for NFS datastores

For information to help you plan for your storage needs, see *vSphere Storage* at https://docs.vmware.com/en/VMware-vSphere/index.html.

https://docs.vmware.com/en/VMware-vSphere/index.htmlAnother good source of information is the vSphere Storage page at https://core.vmware.com/.

6-19 Review of Learner Objectives

- Recognize vSphere storage technologies
- Identify types of datastores
- Recognize storage device naming conventions

6-20 Lesson 2: Fibre Channel Storage

6-21 Learner Objectives

- Describe uses of Fibre Channel with ESXi
- Describe Fibre Channel components and addressing
- Explain how multipathing with Fibre Channel works

6-22 About Fibre Channel

Fibre Channel is a protocol used for accessing storage devices across a network.

A Fibre Channel SAN is a specialized high-speed network that connects your hosts to high-performance storage devices.

The network uses the Fibre Channel protocol to transport SCSI traffic from VMs to the Fibre Channel SAN devices.

ESXi supports:

- 32 Gbps Fibre Channel
- Fibre Channel over Ethernet (FCoE)

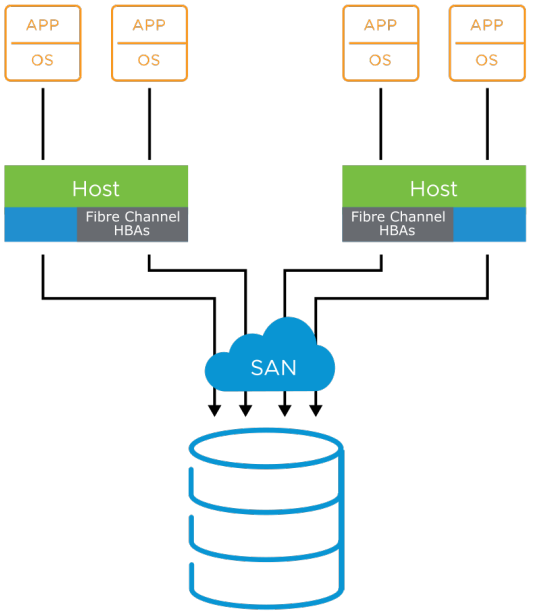

VMFS or vSphere Virtual Volumes Datastore on Fibre Array

To connect to the Fibre Channel SAN, your host should be equipped with Fibre Channel host bus adapters (HBAs).

Unless you use Fibre Channel direct connect storage, you need Fibre Channel switches to route storage traffic. If your host contains FCoE adapters, you can connect to your shared Fibre Channel devices by using an Ethernet network.

In this configuration, a host connects to a SAN fabric, which consists of Fibre Channel switches and storage arrays, using a Fibre Channel adapter. LUNs from a storage array become available to the host. You can access the LUNs and create datastores for your storage needs. These datastores use the VMFS format.

Alternatively, you can access a storage array that supports vSphere Virtual Volumes and create vSphere Virtual Volumes datastores on the array's storage containers.

6-23 Fibre Channel SAN Components

A Fibre Channel SAN consists of one or more servers that are attached to a storage array using one or more Fibre Channel switches.

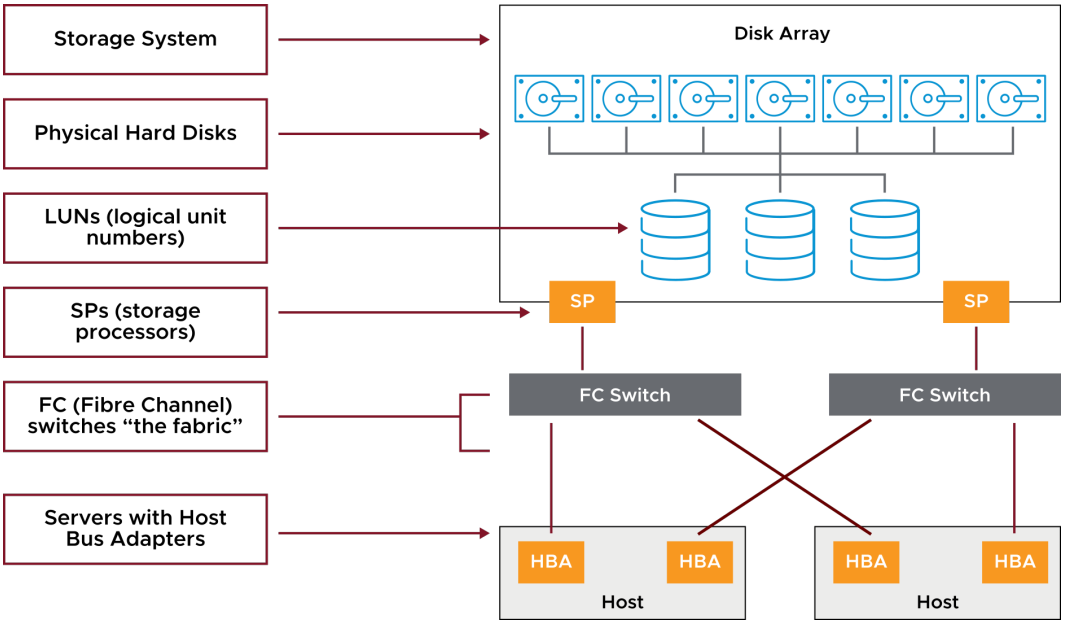

Each Fibre Channel SAN server might host numerous applications that require dedicated storage for applications processing.

The following components are involved:

- Fibre Channel SAN switches: SAN switches connect various elements of the SAN. SAN switches might connect hosts to storage arrays. Using SAN switches, you can set up path redundancy to address any path failures from host server to switch, or from storage array to switch.
- Fabric: The SAN fabric is the network portion of the SAN. When one or more SAN switches are connected, a fabric is created. The Fibre Channel (FC) protocol is used to communicate over the entire network. A SAN can consist of multiple fabrics that are interconnected. Even a simple SAN often consists of two fabrics for redundancy.
- Connections (HBAs and storage processors): Host servers and storage systems are connected to the SAN fabric through ports in the fabric:
  - A host connects to a fabric port through an HBA.
  - Storage devices connect to the fabric ports through their storage processors.

6-24 Fibre Channel Addressing and Access Control

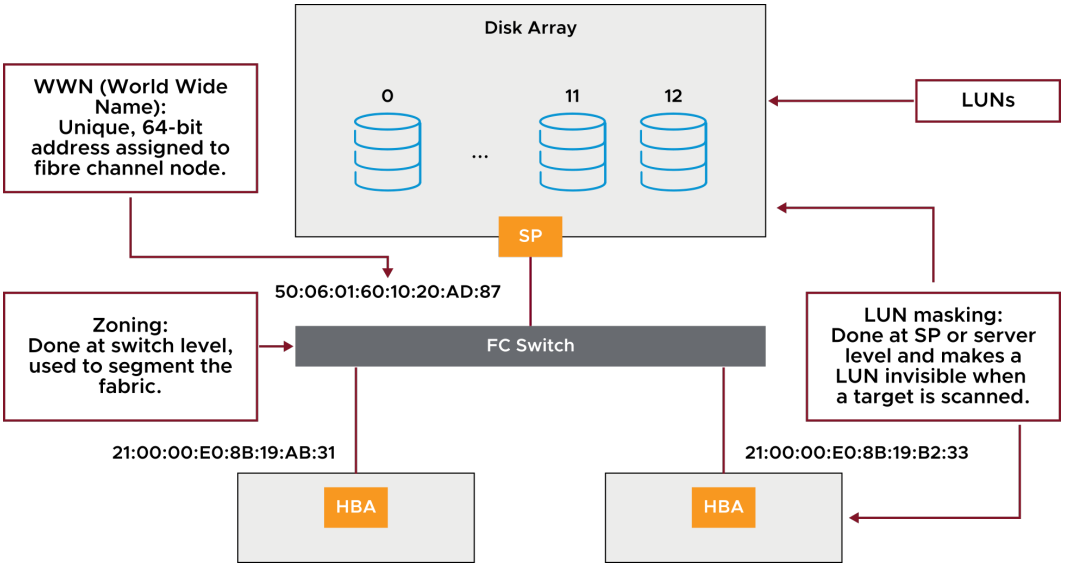

A port connects from a device into the SAN. Each node in the SAN includes each host, storage device, and fabric component (router or switch). Each node in the SAN has one or more ports that connect it to the SAN. Ports can be identified by their Worldwide Port Name (WWPN). WWPN is a globally unique identifier for a port that allows certain applications to access the port. The Fibre Channel switches discover the WWPN of a device or host and assign a port address to the device.

You can use zoning and LUN masking to separate SAN activity and restrict access to storage devices.

You can protect access to storage in your vSphere environment by using zoning and LUN masking with your SAN resources. For example, you might manage zones defined for testing independently within the SAN so that they do not interfere with activity in the production zones. Similarly, you might set up different zones for different departments.

When you set up zones, consider host groups that are set up on the SAN device.

Zoning and masking capabilities for each SAN switch and disk array, and the tools for managing LUN masking, are vendor-specific.

See your SAN vendor's documentation and *vSphere Storage* at https://docs.vmware.com/en/VMware-vSphere/index.html.

6-25 Multipathing with Fibre Channel

Multipathing is having more than one path from a host to a LUN. Multipathing provides the following functions:

- Continued access to SAN LUNs if hardware fails
- Load balancing

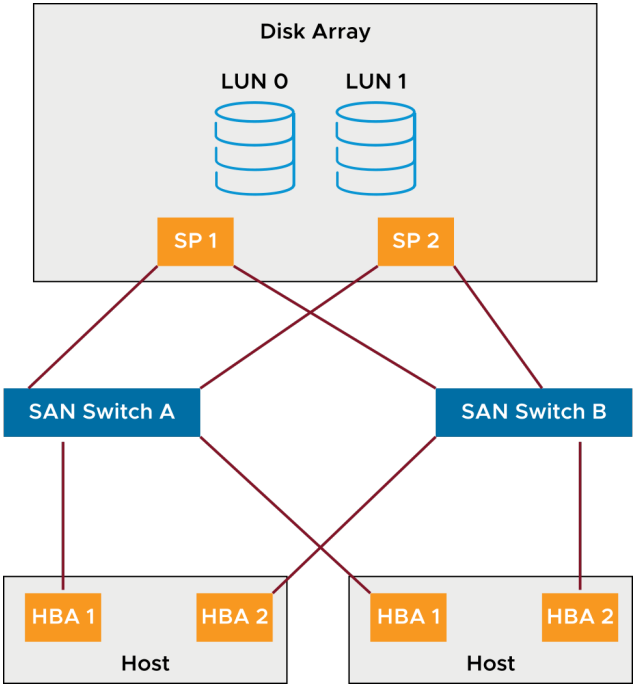

A Fibre Channel path describes a route:

- From a specific HBA port in the host
- Through the switches in the fabric
- Into a specific storage port on the storage array

By default, ESXi hosts use only one path from a host to a given LUN at any one time. If the path actively being used by the ESXi host fails, the server selects another available path.

The process of detecting a failed path and switching to another is called path failover. A path fails if any of the components along the path (HBA, cable, switch port, or storage processor) fail.

Distinguishing between active-active and active-passive disk arrays can be useful:

- An active-active disk array allows access to the LUNs simultaneously through the available storage processors without significant performance degradation. All the paths are active at all times (unless a path fails).
- In an active-passive disk array, one storage processor is actively servicing a given LUN. The other storage processor acts as a backup for the LUN and might be actively servicing other LUN I/O.

  I/O can be sent only to an active processor. If the primary storage processor fails, one of the secondary storage processors becomes active, either automatically or through administrative intervention.

6-26 Review of Learner Objectives

- Describe uses of Fibre Channel with ESXi
- Describe Fibre Channel components and addressing
- Explain how multipathing with Fibre Channel works

6-27 **Lesson 3: iSCSI Storage**

6-28 Learner Objectives

- Describe iSCSI components and addressing
- Configure iSCSI initiators

6-29 iSCSI Components

An iSCSI SAN consists of an iSCSI storage system, which contains LUNs and storage processors. Communication between the host and storage array occurs over an Ethernet network.

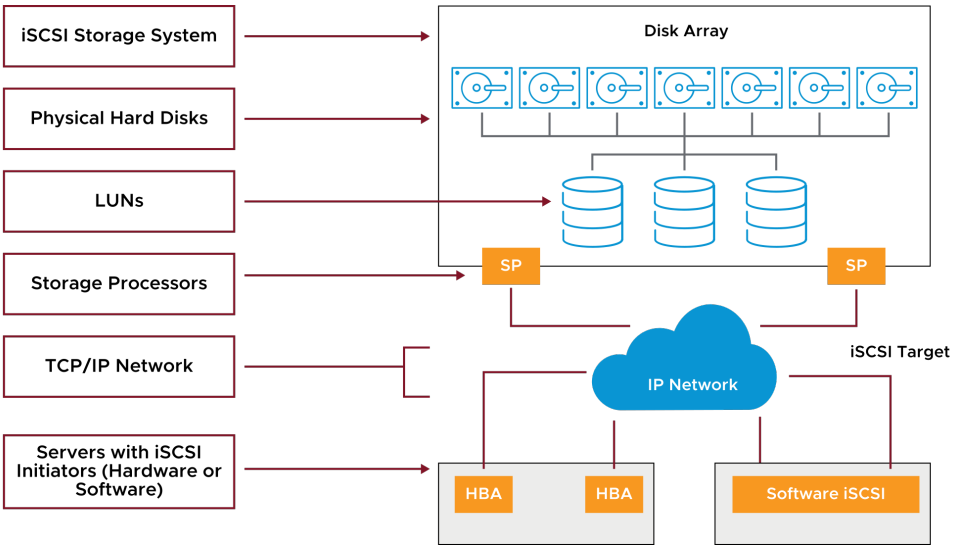

An iSCSI SAN consists of an iSCSI storage system, which contains one or more LUNs and one or more storage processors. Communication between the host and the storage array occurs over a TCP/IP network.

The ESXi host is configured with an iSCSI initiator. An initiator can be hardware-based, where the initiator is an iSCSI HBA. Or the initiator can be software-based, known as the iSCSI software initiator.

An initiator transmits SCSI commands over the IP network. A target receives SCSI commands from the IP network. Your iSCSI network can include multiple initiators and targets. iSCSI is SAN-oriented for the following reasons:

- The initiator finds one or more targets.
- A target presents LUNs to the initiator.
- The initiator sends SCSI commands to a target.

An initiator resides in the ESXi host. Targets reside in the storage arrays that are supported by the ESXi host.

To restrict access to targets from hosts, iSCSI arrays can use various mechanisms, including IP address, subnets, and authentication requirements.

6-30 iSCSI Addressing

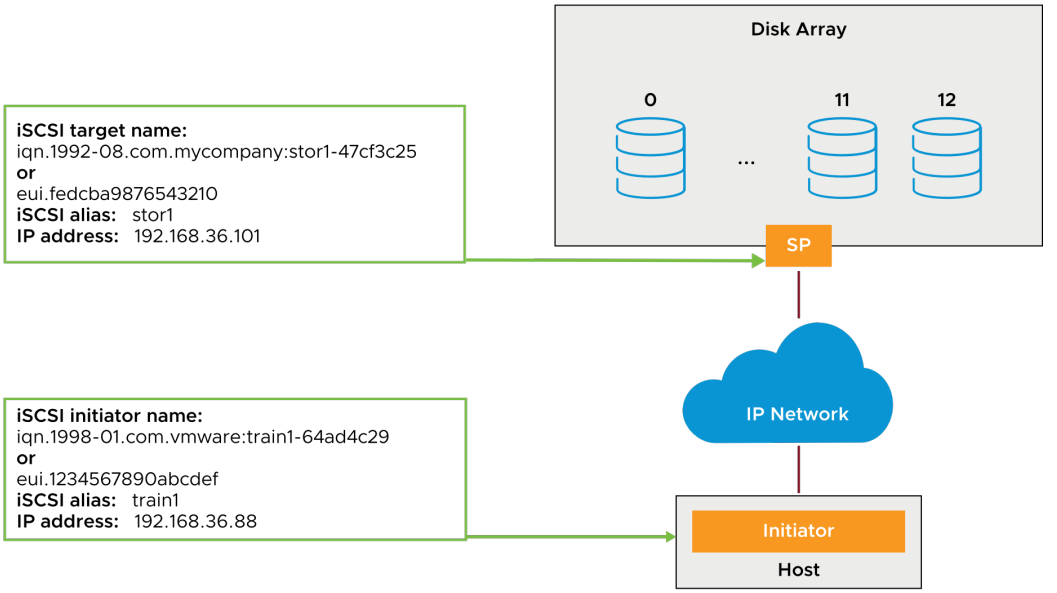

The main addressable, discoverable entity is an iSCSI node. An iSCSI node can be an initiator or a target. An iSCSI node requires a name so that storage can be managed regardless of address.

The iSCSI name can use one of the following formats: The iSCSI qualified name (IQN) or the extended unique identifier (EUI).

The IQN can be up to 255 characters long. Several naming conventions are used:

- Prefix `iqn`
- Date code specifying the year and month in which the organization registered the domain or subdomain name that is used as the naming authority string
- Organizational naming authority string, which consists of a valid, reversed domain or subdomain name
- (Optional) Colon (:) followed by a string of the assigning organization's choosing, which must make each assigned iSCSI name unique

EUI naming conventions are as follows:

- Prefix is `eui`.
- A 16-character name follows the prefix.

  The name includes 24 bits for a company name that is assigned by the IEEE and 40 bits for a unique ID, such as a serial number.

6-31 iSCSI Adapters

You must set up software or hardware iSCSI adapters before an ESXi host can work with iSCSI storage. To access iSCSI targets, your host uses iSCSI initiators.

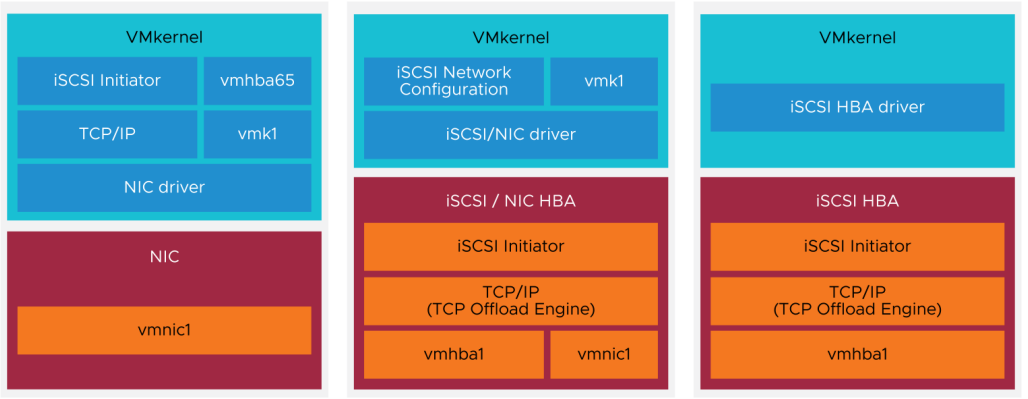

Software iSCSI | Dependent HW iSCSI | Independent HW iSCSI

The iSCSI initiators transport SCSI requests and responses, encapsulated in the iSCSI protocol, between the host and the iSCSI target. Your host supports two types of initiators: software iSCSI and hardware iSCSI.

A software iSCSI initiator is VMware code built in to the VMkernel. Using the initiator, your host can connect to the iSCSI storage device through standard network adapters. The software iSCSI initiator handles iSCSI processing while communicating with the network adapter. With the software iSCSI initiator, you can use iSCSI technology without purchasing specialized hardware.

A hardware iSCSI initiator is a specialized third-party adapter capable of accessing iSCSI storage over TCP/IP. Hardware iSCSI initiators are divided into two categories: dependent hardware iSCSI and independent hardware iSCSI.

A dependent hardware iSCSI initiator, also known as an iSCSI host bus adapter, is a standard network adapter that includes the iSCSI offload function. To use this type of adapter, you must configure networking for the iSCSI traffic and bind the adapter to an appropriate VMkernel iSCSI port.

An independent hardware iSCSI adapter handles all iSCSI and network processing and management for your ESXi host. In this case, a VMkernel iSCSI port is not required.

For configuration information, see *vSphere Storage* at https://docs.vmware.com/en/VMware-vSphere/index.html.

6-32 ESXi Network Configuration for Software iSCSI

A VMkernel port must be created before enabling the software iSCSI initiator.

To optimize your vSphere networking setup, you separate iSCSI networks from NAS and NFS networks:

- Physical separation is preferred
- If physical separation is not possible, use VLANs

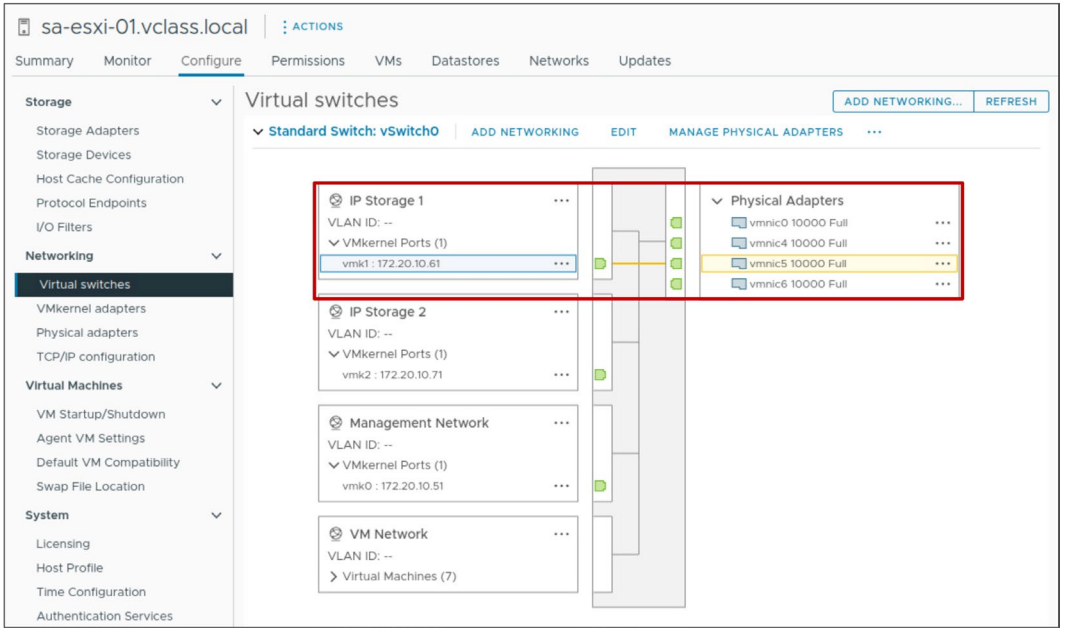

Networking configuration for software iSCSI involves creating a VMkernel port on a virtual switch to handle your iSCSI traffic.

Depending on the number of physical adapters that you want to use for the iSCSI traffic, the networking setup can be different:

- If you have one physical network adapter, you need a VMkernel port on a virtual switch.
- If you have two or more physical network adapters for iSCSI, you can use these adapters for host-based multipathing.

VMkernel ports are also used to handle the connection to dependent hardware iSCSI.

For performance and security, isolate your iSCSI network from other networks. Physically separate the networks. If physically separating the networks is impossible, logically separate the networks from one another on a single virtual switch by configuring a separate VLAN for each network.

6-33 Activating the Software iSCSI Adapter

To add the software iSCSI adapter:

1. Select the host and click the **Configure** tab.
2. Select **Storage Adapters** and click **ADD SOFTWARE ADAPTER > Add iSCSI adapter**.

The software iSCSI adapter (vmhba65) appears in the list.

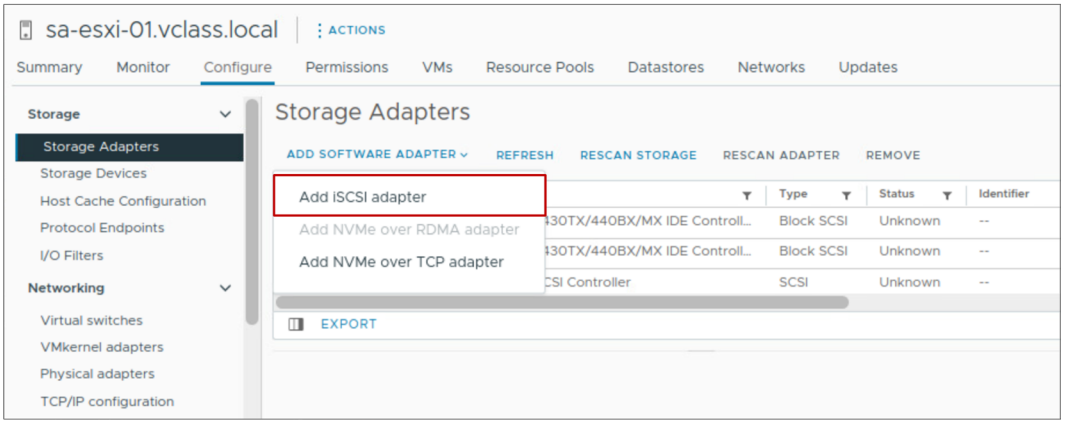

You must activate your software iSCSI adapter so that your host can use it to access iSCSI storage.

You can activate only one software iSCSI adapter.

If you boot from iSCSI using the software iSCSI adapter, the adapter is active, and the network configuration is created at the first boot. If you deactivate the adapter, it is reactivated each time you boot the host.

6-34 Discovering iSCSI Targets

The iSCSI adapter discovers storage resources on the network and determines which resources are available for access.

An ESXi host supports the following discovery methods:

- Static
- Dynamic or SendTargets

The SendTargets response returns the IQN and all available IP addresses.

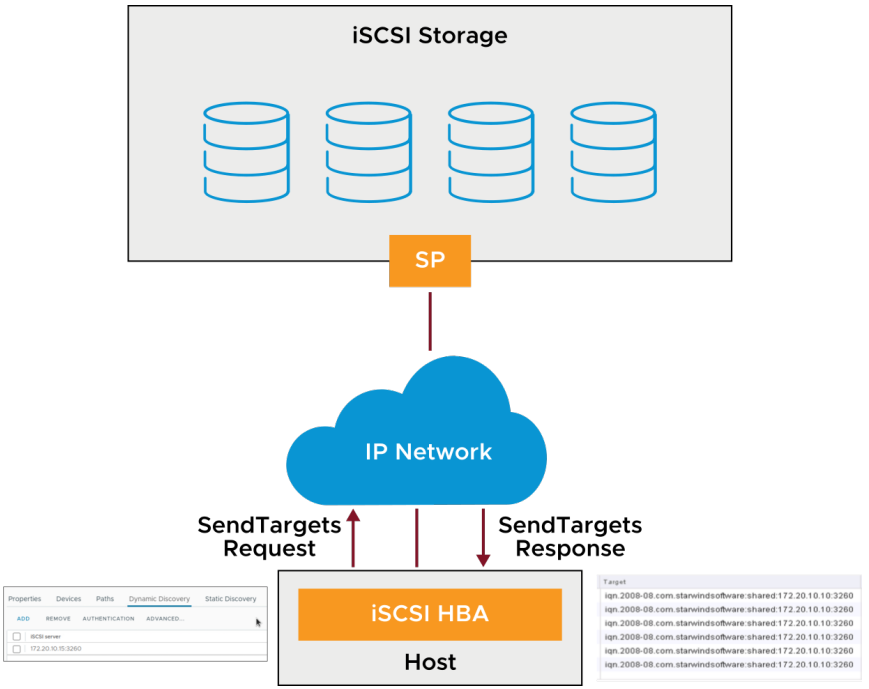

The host does not see iSCSI LUNs until it can communicate with the storage array. You do this by providing the host with the IP address of the array's storage processors.

The ESXi host supports the following iSCSI target-discovery methods:

- Static discovery: The initiator does not have to perform discovery. The initiator knows in advance all the targets that it will contact. It uses their IP addresses and domain names to communicate with them.
- Dynamic discovery or SendTargets discovery: Each time the initiator contacts a specified iSCSI server, it sends the SendTargets request to the server. The server responds by supplying a list of available targets to the initiator.

The names and IP addresses of these targets appear as static targets in the vSphere Client. You can remove a static target that is added by dynamic discovery. If you remove the target, the target might be returned to the list during the next rescan operation. The target might also be returned to the list if the HBA is reset or the host is rebooted.

6-35 iSCSI Security: CHAP

iSCSI initiators use CHAP for authentication purposes.

By default, CHAP is not configured.

ESXi supports two types of CHAP authentication:

- Unidirectional
- Bidirectional

ESXi also supports per-target CHAP authentication.

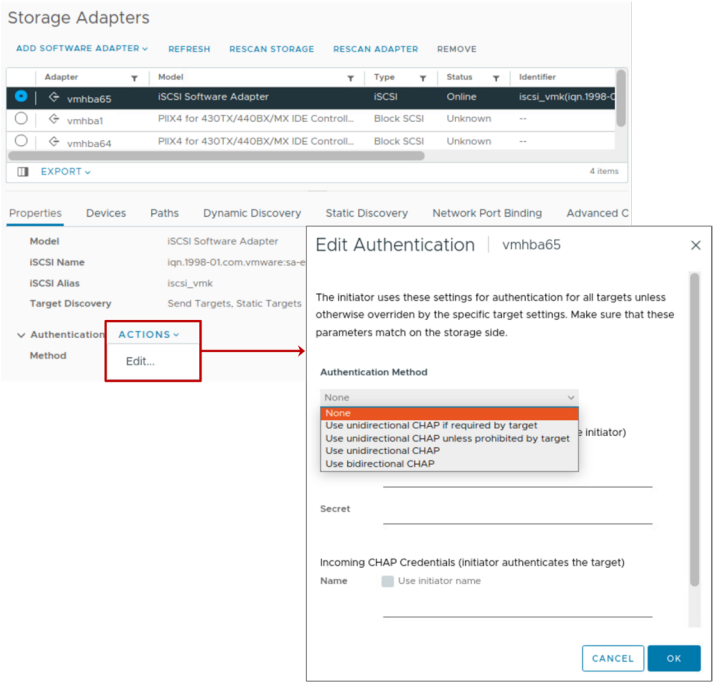

You can implement CHAP to provide authentication between iSCSI initiators and targets.

ESXi supports the following CHAP authentication methods:

- Unidirectional or one-way CHAP: The target authenticates the initiator, but the initiator does not authenticate the target. You must specify the CHAP secret so that your initiators can access the target.
- Bidirectional or mutual CHAP: With an extra level of security, the initiator can authenticate the target. You must specify different target and initiator secrets.

CHAP uses a three-way handshake algorithm to verify the identity of your host and, if applicable, of the iSCSI target when the host and target establish a connection. The verification is based on a predefined private value, or CHAP secret, that the initiator and target share. ESXi implements CHAP as defined in RFC 1994.

ESXi supports CHAP authentication at the adapter level. All targets receive the same CHAP secret from the iSCSI initiator. For both software iSCSI and dependent hardware iSCSI initiators, ESXi also supports per-target CHAP authentication.

Before configuring CHAP, check whether CHAP is activated at the iSCSI storage system and check the CHAP authentication method that the system supports. If CHAP is activated, you must activate it for your initiators, verifying that the CHAP authentication credentials match the credentials on the iSCSI storage.

Using CHAP in your iSCSI SAN implementation is recommended, but consult with your storage vendor to ensure that best practices are followed.

You can protect your data in additional ways. For example, you might protect your iSCSI SAN by giving it a dedicated standard switch. You might also configure the iSCSI SAN on its own VLAN to improve performance and security. Some inline network devices might be implemented to provide encryption and further data protection.

6-36 Multipathing with Software iSCSI

Software iSCSI uses multiple NICs:

- Each NIC is connected to a separate VMkernel port.
- Each VMkernel port binds with the iSCSI initiator.

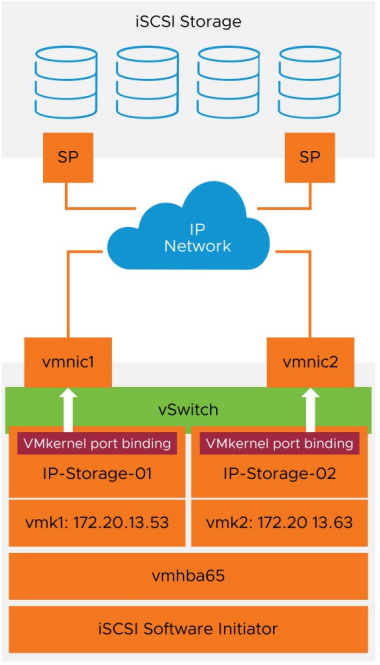

When setting up your ESXi host for multipathing and failover, you can use multiple hardware iSCSI adapters or multiple NICs. The choice depends on the type of iSCSI initiators on your host.

With software iSCSI, you can use multiple NICs that provide failover for iSCSI connections between your host and iSCSI storage systems.

After iSCSI multipathing is set up, each port on the ESXi host has its own IP address, but the ports share the same iSCSI initiator IQN. When iSCSI multipathing is configured, the VMkernel routing table is not consulted for identifying the outbound NIC to use. Instead, iSCSI multipathing is managed using vSphere multipathing modules. Because of the latency that can be incurred, routing iSCSI traffic is not recommended.

6-37 Multipathing with Dependent Hardware iSCSI

Dependent hardware iSCSI uses multiple NICs and iSCSI HBAs:

- Each NIC is connected to a separate VMkernel port.
- Each VMkernel port binds with the iSCSI initiator.

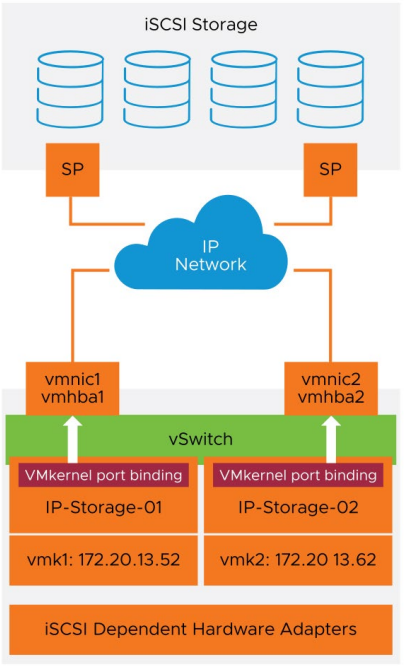

A dependent hardware iSCSI adapter is a third-party adapter that depends on vSphere networking and iSCSI configuration.

When a dependent hardware iSCSI adapter is installed on an ESXi host, it presents its two components, a standard network adapter (vmnic) and an iSCSI engine, to the same port. The iSCSI engine appears on the list of storage adapters as an iSCSI adapter (vmhba).

The iSCSI adapter is enabled by default. To make it functional, you must connect it, through a VMkernel adapter (vmk), to a physical network adapter (vmnic) associated with it. You can then configure the iSCSI adapter.

6-38 Multipathing with Independent Hardware iSCSI

Independent hardware iSCSI uses two or more hardware iSCSI adapters.

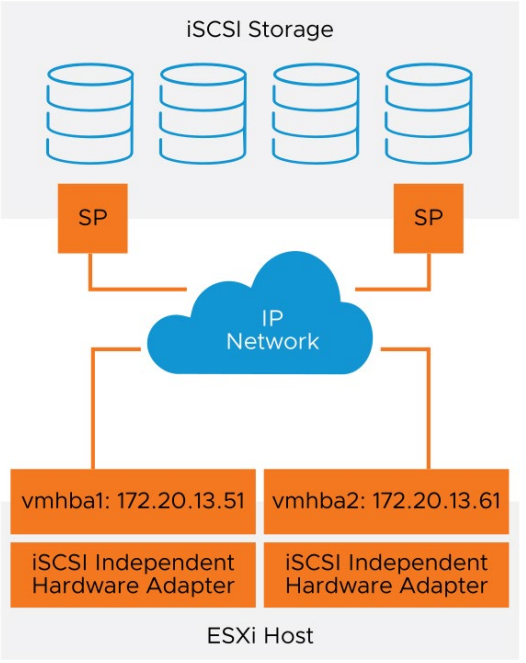

With independent hardware iSCSI, the host typically has two or more available hardware iSCSI adapters, from which the storage system can be reached by using one or more switches.

Alternatively, the setup might include one adapter and two storage processors so that the adapter can use a different path to reach the storage system.

6-39 Binding VMkernel Ports with the iSCSI Initiator

With port binding, each VMkernel port that is connected to a separate NIC becomes a different path that the iSCSI storage stack can use.

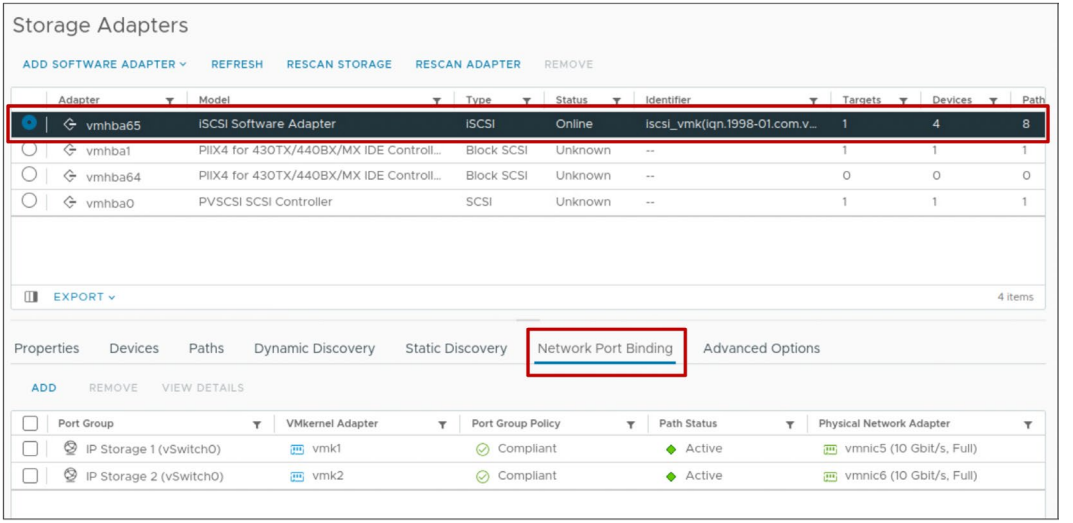

With software iSCSI and dependent hardware iSCSI, multipathing plug-ins do not have direct access to physical NICs on your host. For this reason, you must first connect each physical NIC to a separate VMkernel port. Then you use a port-binding technique to associate all VMkernel ports with the iSCSI initiator.

For dependent hardware iSCSI, you must correctly install the physical network card, which should appear on the host's **Configure** tab in the Virtual Switches view.

6-40 Lab 9: Accessing iSCSI Storage

Configure access to an iSCSI datastore:

1. View an Existing ESXi Host iSCSI Configuration
2. Add a VMkernel Port for IP Storage
3. Add a Second VMkernel Port for IP Storage
4. Add the iSCSI Software Adapter to an ESXi Host
5. Discover LUNs on the iSCSI Target Server

6-41 Review of Learner Objectives

- Describe iSCSI components and addressing
- Configure iSCSI initiators

6-42 Lesson 4: VMFS Datastores

6-43 Learner Objectives

- Create a VMFS datastore
- Increase the size of a VMFS datastore
- Delete a VMFS datastore

6-44 About VMFS Datastores

VMFS is a high-performance, cluster file system that serves as a repository for files such as VM files, VM templates and ISO images.

A VMFS datastore is optimized for storing and accessing large files, such as virtual disks and memory images of suspended VMs.

A VMFS datastore can have a maximum volume size of 64 TB.

6-45 Creating a VMFS Datastore

You can create VMFS datastores on any SCSI-based storage devices that the host discovers, including Fibre Channel, iSCSI, and local storage devices.

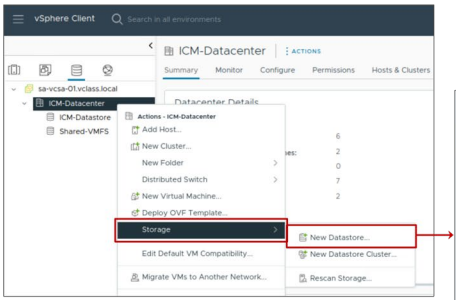

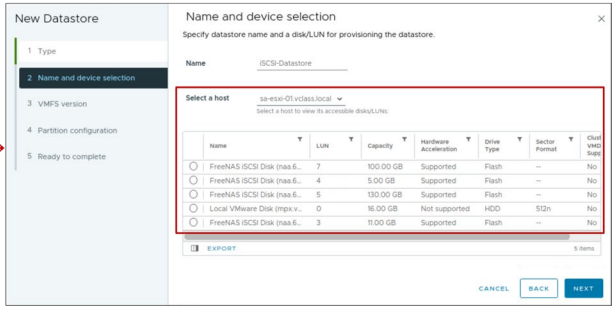

6-46 Browsing Datastore Contents

You use the datastore file browser to manage the contents of your datastores.

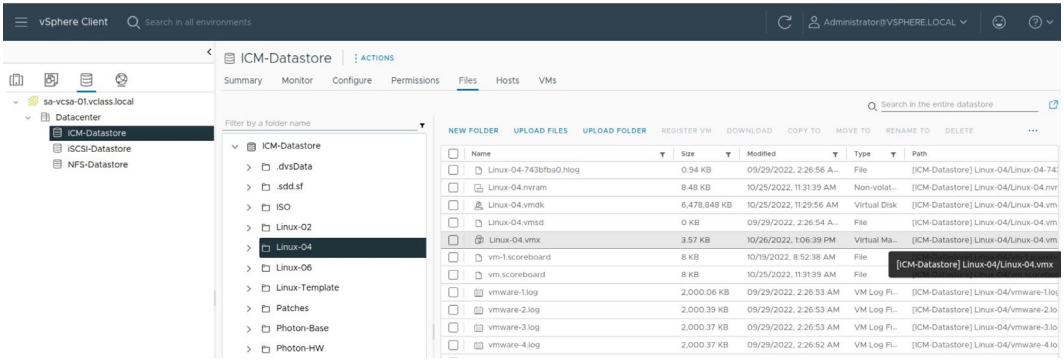

The Datastores pane lists all datastores currently configured for all managed ESXi hosts.

The example shows the contents of the VMFS datastore named ICM-Datastore. The contents of the datastore are folders that contain the files for virtual machines or templates.

6-47 Increasing the Size of VMFS Datastores

Increase a VMFS datastore's size to give it more space or to possibly improve performance.

In general, before changing your storage allocation:

- Perform a rescan to ensure that all hosts see the most current storage.
- Record the unique identifier of the volume that you want to expand.

To dynamically increase the size of a VMFS datastore, use one of the following methods:

- Add an extent (LUN).
- Expand the datastore within its extent.

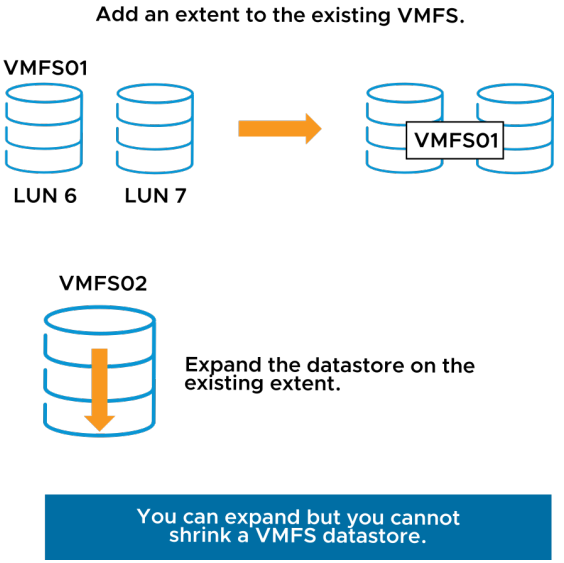

An example of the unique identifier of a volume is the NAA ID. You require this information to identify the VMFS datastore that must be increased.

You can dynamically increase the capacity of a VMFS datastore if the datastore has insufficient disk space. You discover whether insufficient disk space is an issue when you create a VM or you try to add more disk space to a VM.

Use one of the following methods:

- Add an extent to the VMFS datastore: An extent is a partition on a LUN. You can add an extent to any VMFS datastore. The datastore can stretch over multiple extents, up to 32.
- Expand the VMFS datastore: You expand the size of the VMFS datastore by expanding its underlying extent first.

6-48 Datastore Maintenance Mode

Before taking a datastore out of service, place the datastore in maintenance mode.

Before placing a datastore in maintenance mode, you must first move all VMs (powered on and powered off) and templates to a different datastore.

The datastore enters maintenance mode after all VMs and templates are moved off the datastore.

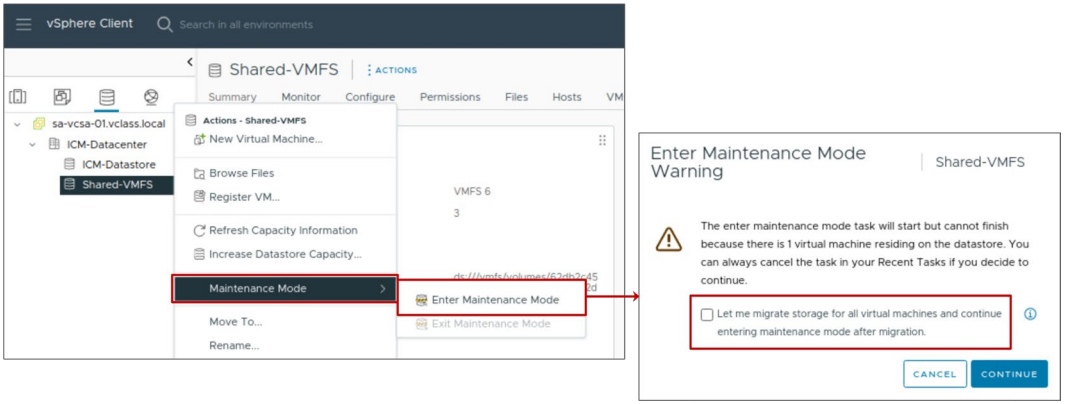

By selecting the **Let me migrate storage for all virtual machines and continue entering maintenance mode after migration** check box, the Migrate wizard starts, giving you the opportunity to migrate VMs to another datastore. If VM templates exist on the datastore, you can convert the templates to VMs, migrate the VMs to another datastore, then convert the VMs back to VM templates.

Datastore maintenance mode is a function of the vSphere Storage DRS feature, but you can use maintenance mode without activating vSphere Storage DRS. If vSphere Storage DRS is configured and fully-automated, the VM migrations happen automatically. For more information on vSphere Storage DRS, see *vSphere Resource Management* at https://docs.vmware.com/en/VMware-vSphere/index.html.

6-49 Deleting or Unmounting a VMFS Datastore

An unmounted datastore remains intact, but cannot be seen from the hosts that you specify.

It continues to appear on other hosts, where it remains mounted.

A deleted datastore is destroyed and disappears from all hosts that have access to it.

The deleted datastore permanently removes all files on the datastore.

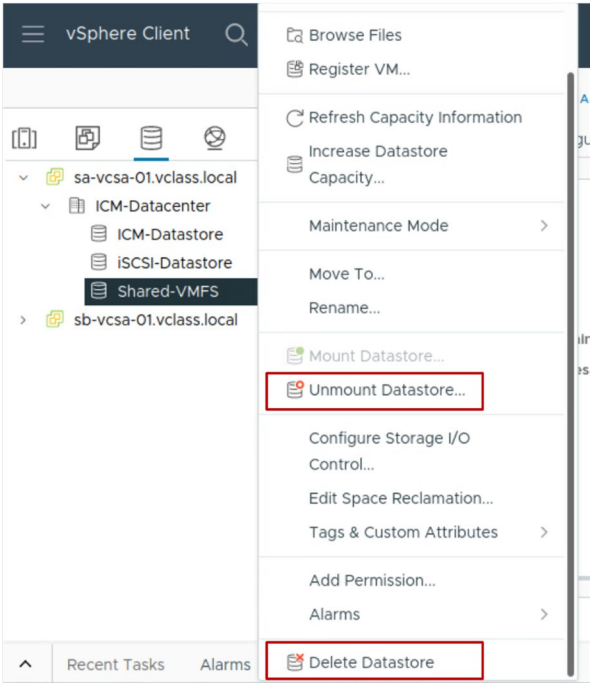

Unmounting a VMFS datastore preserves the files on the datastore but makes the datastore inaccessible to the ESXi host.

Do not perform any configuration operations that might result in I/O to the datastore while the unmounting is in progress.

You can delete any type of VMFS datastore, including copies that you mounted without re-signing. Although you can delete the datastore without unmounting, you should unmount the datastore first. Deleting a VMFS datastore destroys the pointers to the files on the datastore, so the files disappear from all hosts that have access to the datastore.

Before you delete or unmount a VMFS datastore, power off all VMs whose disks reside on the datastore. If you do not power off the VMs and you try to continue, an error message tells you that the resource is busy. Before you unmount a VMFS datastore, use the vSphere Client to verify the following conditions:

- No virtual machines reside on the datastore.
- The datastore is not part of a datastore cluster.
- The datastore is not managed by vSphere Storage DRS.
- vSphere Storage I/O Control is deactivated.
- The datastore is not used for vSphere HA heartbeat.

To keep your data, back up the contents of your VMFS datastore before you delete the datastore.

6-50 Multipathing Algorithms

Arrays provide active-active and active-passive storage processors. Multipathing algorithms interact with these storage arrays:

- vSphere offers native path selection, load-balancing, and failover mechanisms.
- Third-party vendors can create software for ESXi hosts to properly interact with the storage arrays.

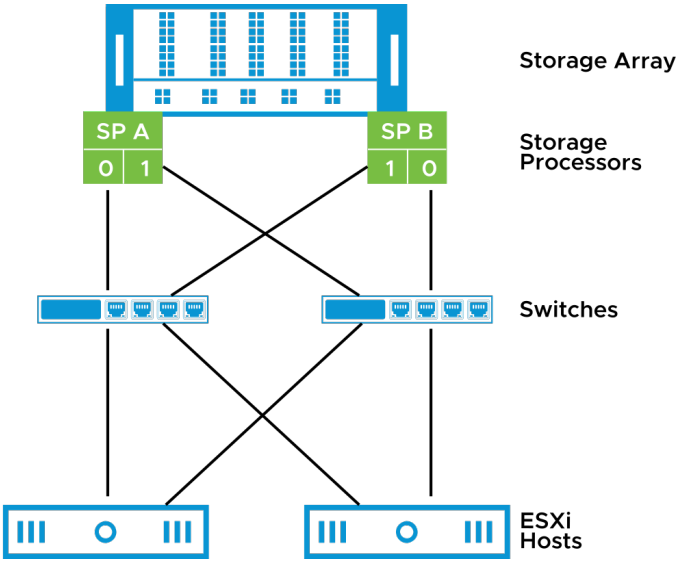

The Pluggable Storage Architecture is a VMkernel layer responsible for managing multiple storage paths and providing load balancing. An ESXi host can be attached to storage arrays with either active-active or active-passive storage processor configurations.

VMware offers native load-balancing and failover mechanisms. VMware path selection policies include the following examples:

- Round Robin
- Most Recently Used (MRU)
- Fixed

Third-party vendors can design their own load-balancing techniques and failover mechanisms for particular storage array types to add support for new arrays. Third-party vendors do not need to provide internal information or intellectual property about the array to VMware.

6-51 Configuring Storage Load Balancing

Path selection policies provide:

- Scalability:
  - Round Robin
- Availability:
  - Most Recently Used
  - Fixed

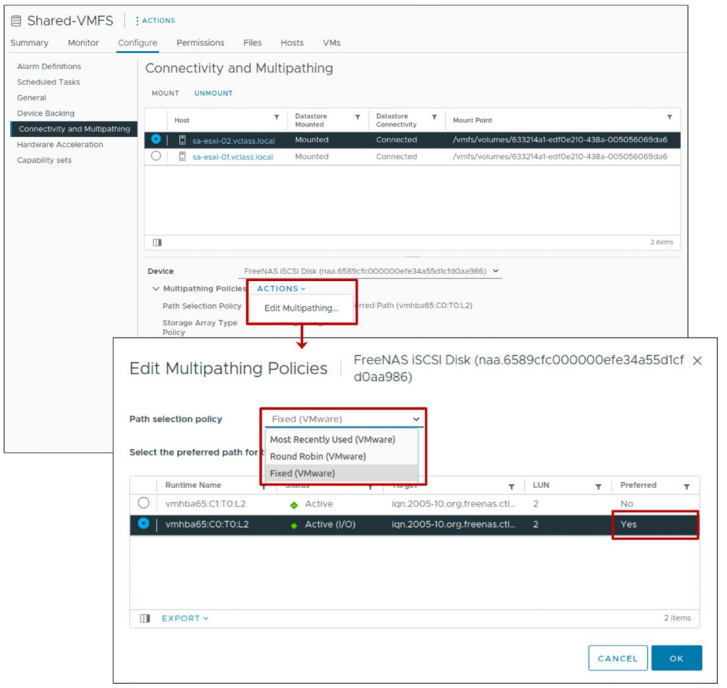

Multiple paths from an ESXi host to a datastore are possible.

For multipathing with Fibre Channel or iSCSI, the following path selection policies are supported:

- Fixed: The host always uses the preferred path to the disk when that path is available. If the host cannot access the disk through the preferred path, it tries the alternative paths. This policy is the default policy for active-active storage devices.
- Most Recently Used: The host selects the first working path discovered at system boot time. When the path becomes unavailable, the host selects an alternative path. The host does not revert to the original path when that path becomes available. The Most Recently

Used policy does not use the preferred path setting. This policy is the default policy for active-passive storage devices and is required for those devices.

- Round Robin: In addition to path failover, this policy supports load balancing across the paths. By default, the latency mechanism is activated on the host. The mechanism considers I/O bandwidth and path latency to select an optimal path for I/O. When using the latency mechanism, the Round Robin policy can dynamically select the optimal path and achieve better load balancing results. Before using this policy, check with storage vendors to find out whether a Round Robin configuration is supported on their storage.

6-52 Lab 10: Managing VMFS Datastores

Create VMFS datastores, increase the size of these datastores, and share datastores between ESXi hosts:

1. Create VMFS Datastores for the ESXi Hosts
2. Expand a VMFS Datastore to Consume Unused Space on a LUN
3. Remove a VMFS Datastore
4. Extend a VMFS Datastore by Adding a LUN

6-53 Review of Learner Objectives

- Create a VMFS datastore
- Increase the size of a VMFS datastore
- Delete a VMFS datastore

6-54 Lesson 5: NFS Datastores

6-55 Learner Objectives

- Identify NFS components
- Recognize the differences between NFS 3 and NFS 4.1
- Configure and manage NFS datastores

6-56 NFS Components

An NFS file system is on a NAS device that is called the NFS server.

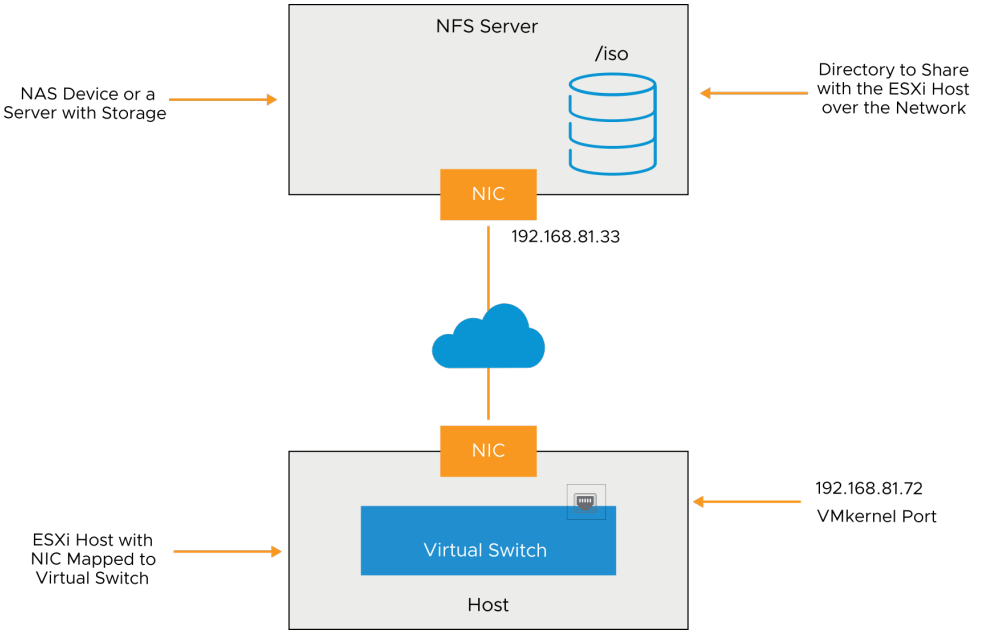

The NFS server contains one or more directories that are shared with the ESXi host over a TCP/IP network. An ESXi host accesses the NFS server through a VMkernel port that is defined on a virtual switch.

6-57 NFS 3 and NFS 4.1

An NFS datastore can be created as either NFS 3 or NFS 4.1.

NFS 3	NFS 4.1
ESXi managed multipathing	Native multipathing and session trunking
AUTH_SYS (root) authentication	Optional Kerberos authentication
VMware proprietary client-side file locking	Server-side file locking
Client-side error tracking	Server-side error tracking

Compatibility issues between the two NFS versions prevent access to datastores using both protocols at the same time from different hosts. If a datastore is configured as NFS 4.1, all hosts that access that datastore must mount the share as NFS 4.1. Data corruption can occur if hosts access a datastore with the wrong NFS version.

6-58 NFS Version Compatibility with Other vSphere Technologies

vSphere supports NFS 4.1 to overcome many limitations when using NFS 3. Both NFS 3 and NFS 4.1 shares can be used, but you must consider important constraints when designing a vSphere environment in which both versions are used.

vSphere Technology	NFS 3	NFS 4.1
vSphere vMotion and vSphere Storage vMotion	Yes	Yes
vSphere HA and vSphere Fault Tolerance	Yes	Yes
vSphere DRS and vSphere DPM	Yes	Yes
Stateless ESXi and Host Profiles	Yes	Yes
vSphere Storage DRS and Storage I/O Control	Yes	No
Site Recovery Manager	Yes	Partial*
vSphere Virtual Volumes and vSphere Replication	Yes	Yes
vRealize Operations Manager	Yes	Yes
Host Profiles	Yes	Yes

* Site Recovery Manager does not support NFS 4.1 datastores for array-based replication and vSphere Virtual Volumes replication. You can use Site Recovery Manager with NFS v 4.1 datastores for vSphere Replication.

NFS 4.1 provides the following enhancements:

- Native multipathing and session trunking: NFS 4.1 provides multipathing for servers that support session trunking. When trunking is available, you can use multiple IP addresses to access a single NFS volume. Client ID trunking is not supported.
- Kerberos authentication: NFS 4.1 introduces Kerberos authentication in addition to the traditional AUTH_SYS method used by NFS 3.
- Improved built-in file locking.
- Enhanced error recovery using server-side tracking of open files and delegations.
- Many general efficiency improvements including session leases and less protocol overhead.

The NFS 4.1 client offers the following new features:

- Stateful locks with share reservation using a mandatory locking semantic
- Protocol integration, side-band (auxiliary) protocol no longer required to lock and mount
- Trunking (true NFS multipathing), where multiple paths (sessions) to the NAS array can be created and load-distributed across those sessions
- Enhanced error recovery to mitigate server failure and loss of connectivity

6-59 Configuring NFS Datastores

To configure an NFS datastore:

1. Create a VMkernel port:
    - For better performance and security, separate your NFS network from the iSCSI network.
2. Create the NFS datastore by providing the following information:
    - NFS version: 3 or 4.1
    - Datastore name
    - NFS server names or IP addresses
    - Folder on the NFS server, for example, `/templates` or `/nfs_share`
    - Whether to mount the NFS file system as read only
    - Hosts that mount the datastore
    - Authentication parameters

For each ESXi host that accesses an NFS datastore over the network, a VMkernel port must be configured on a virtual switch. The name of this port can be anything that you want.

For performance and security reasons, isolate your NFS networks from the other networks, such as your iSCSI network and your virtual machine networks.

6-60 Configuring ESXi Host Authentication and NFS Kerberos Credentials

As a requirement of Kerberos authentication, you must add each ESXi host to the Active Directory domain. Then you configure NFS Kerberos credentials.

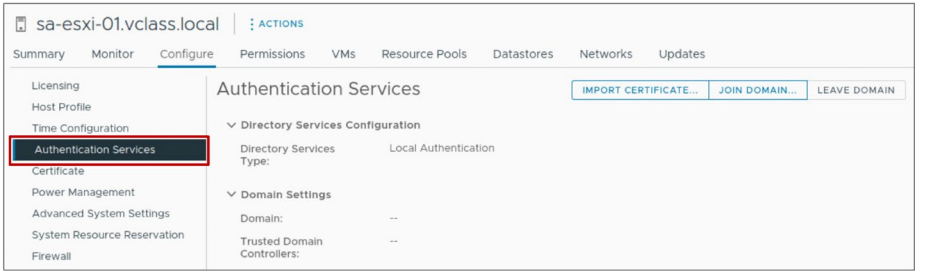

You must take several configuration steps to prepare each ESXi host to use Kerberos authentication.

Kerberos authentication requires that all nodes involved (the Active Directory server, the NFS servers, and the ESXi hosts) be synchronized so that little to no time drift exists. Kerberos authentication fails if any significant drift exists between the nodes.

To prepare your ESXi host to use Kerberos authentication, configure the NTP client settings to reference a common NTP server (or the domain controller, if applicable).

When planning to use NFS Kerberos, consider the following points:

- NFS 3 and 4.1 use different authentication credentials, resulting in incompatible UID and GID on files.
- Using different Active Directory users on different hosts that access the same NFS share can cause the vSphere vMotion migration to fail.
- NFS Kerberos configuration can be automated by using host profiles to reduce configuration conflicts.
- Time must be synchronized between all participating components.

For details on configuring ESXi hosts for Kerberos authentication, see *vSphere Storage* at https://docs.vmware.com/en/VMware-vSphere/index.html.

6-61 Configuring the NFS Datastore to Use Kerberos

When creating each NFS 4.1 datastore, you activate Kerberos authentication by selecting one of the security modes:

- Kerberos5 authentication
- Kerberos5i authentication and data integrity

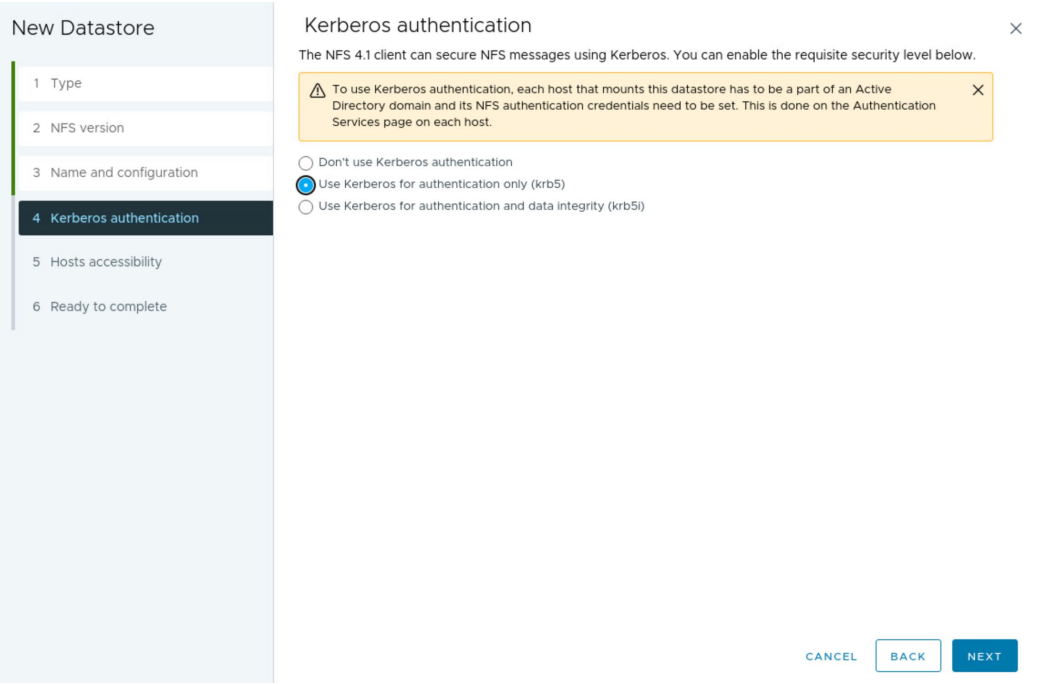

After performing the initial configuration steps, you can configure the datastore to use Kerberos authentication.

The screenshot shows a choice of Kerberos authentication only (krb5) or authentication with data integrity (krb5i). KRB5i ensures that man-in-the-middle attacks that modify data can be detected. With KRB5, these attacks cannot be detected.

For more information about how to configure the ESXi hosts for Kerberos authentication, see *vSphere Storage* at https://docs.vmware.com/en/VMware-vSphere/index.html.

6-62 Unmounting an NFS Datastore

Unmounting an NFS datastore causes the files on the datastore to become inaccessible to the selected ESXi hosts.

Before unmounting an NFS datastore, you must power off all VMs whose disks reside on the datastore.

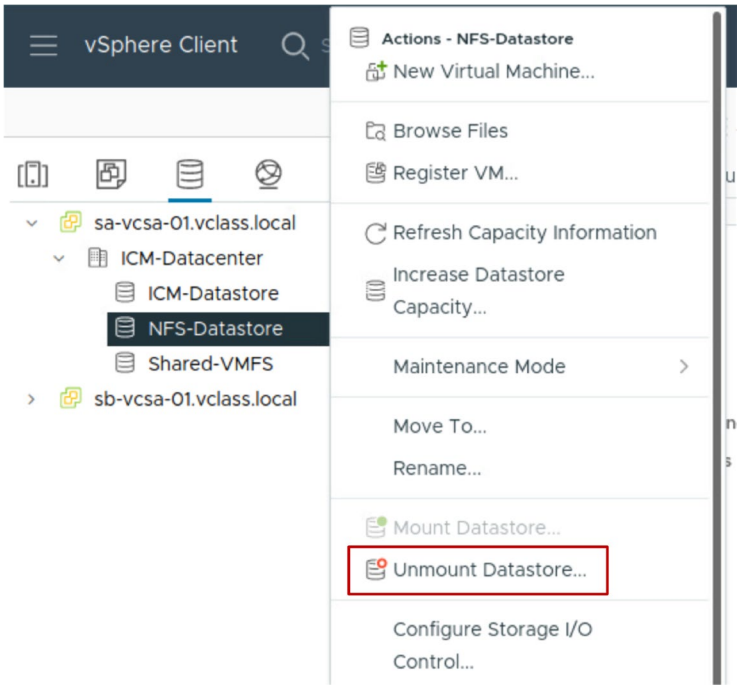

6-63 Multipathing and NFS Storage

For a highly available NAS architecture, configure NFS multipathing to avoid single points of failure.

Example of a multipathing configuration:

- Configure one VMkernel port.
- Attach NICs to the same physical switch to configure NIC teaming.
- Configure the NFS server with multiple IP addresses (same subnet is OK).
- To better use multiple links, configure NIC teams with the IP hash load-balancing policy.

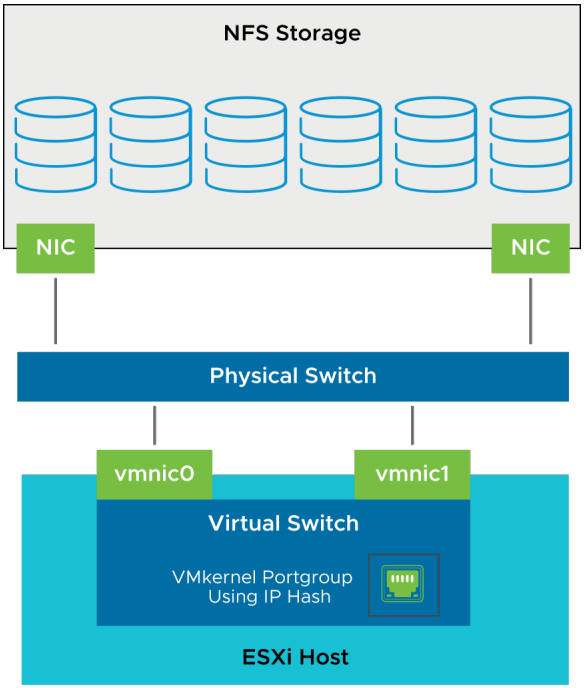

Examples of a single point of failure in the NAS architecture include the NIC card in an ESXi host, and the cable between the NIC card and the switch. To avoid single points of failure and to create a highly available NAS architecture, configure the ESXi host with redundant NIC cards and redundant physical switches.

The best approach is to install multiple NICs on an ESXi host and configure them in NIC teams. NIC teams should be configured on separate external switches, with each NIC pair configured as a team on the respective external switch.

In addition, you might apply a load-balancing algorithm, based on the link aggregation protocol type supported on the external switch, such as 802.3ad or EtherChannel.

An even higher level of performance and high availability can be achieved with cross-stack, EtherChannel-capable switches. With certain network switches, you can team ports across two or more separate physical switches that are managed as one logical switch.

NIC teaming across virtual switches provides additional resilience and some performance optimization. Having more paths available to the ESXi host can improve performance by activating distributed load sharing.

Only one active path is available for the connection between the ESXi host and a single storage target (LUN or mount point). Although alternative connections might be available for failover, the bandwidth for a single datastore and the underlying storage is limited to what a single connection can provide.

To use more available bandwidth, an ESXi host requires multiple connections from the ESXi host to the storage targets. You might need to configure multiple datastores, each using separate connections between the ESXi host and the storage.

The table shows the recommended configuration for NFS multipathing.

External Switches Support Cross-Stack EtherChannel	**External Switches Do Not Support Cross-Stack EtherChannel**
Configure one VMkernel port.	Configure two or more VMkernel ports on different virtual switches on different subnets.
Configure NIC teaming by using adapters attached to separate physical switches.	Configure NIC teaming with adapters attached to the same physical switch.
Configure the NFS server with multiple IP addresses. IP addresses can be on the same subnet.	Configure the NFS server with multiple IP addresses. IP addresses can be on the same subnet.
To use multiple links, configure NIC teams with the IP hash load-balancing policy.	To use multiple links, allow the VMkernel routing table to decide which link to send packets (requires multiple datastores).

6-64 Configuring Multipathing for NFS 4.1

NFS 4.1 supports native multipathing and session trunking.

To configure multipathing, enter multiple server IP addresses when configuring the datastore.

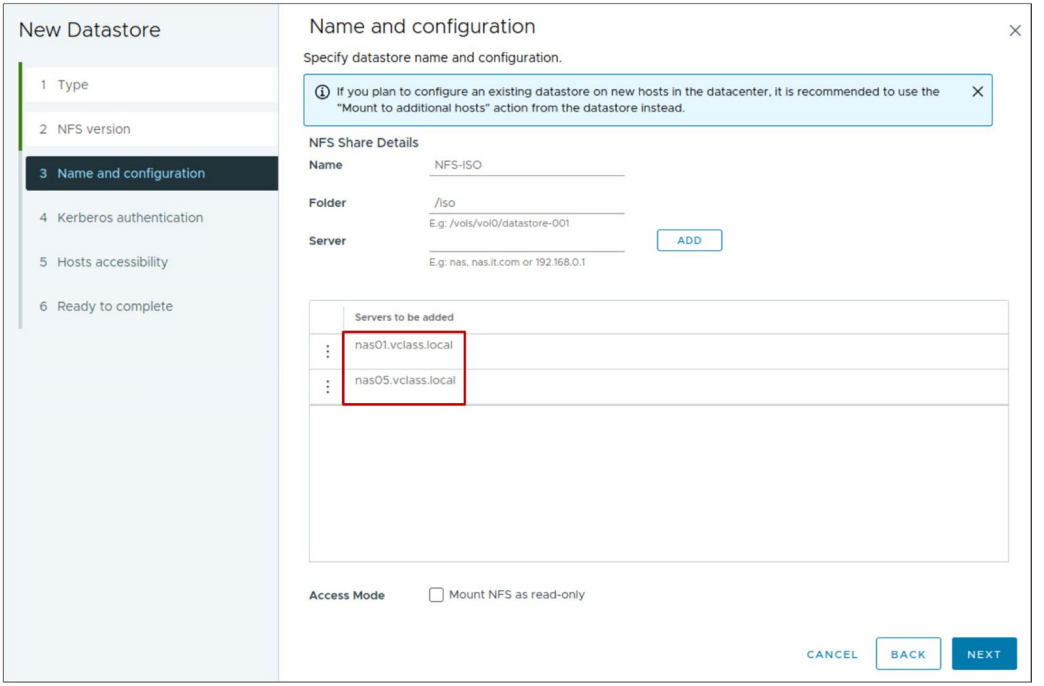

NFS 4.1 provides multipathing for servers that support the session trunking. When trunking is available, you can use multiple IP addresses to access a single NFS volume. Client ID trunking is not supported.

6-65 Lab 11: Accessing NFS Storage

Create an NFS datastore and record its storage information:

1. Configure Access to an NFS Datastore
2. View NFS Storage Information

6-66 Review of Learner Objectives

- Identify NFS components
- Recognize the differences between NFS 3 and NFS 4.1
- Configure and manage NFS datastores

6-67 Key Points

- ESXi hosts support various storage technologies: Direct-attached storage, Fibre Channel, FCoE, iSCSI, and NAS.
- VMFS and NFS datastores hold VM files.
- vSAN and vSphere Virtual Volumes hold VM objects.
- With port binding, each VMkernel port that is connected to a separate NIC becomes a different path that the iSCSI storage can use.
- Shared storage is integral to vSphere features such as vSphere vMotion, vSphere HA, and vSphere DRS.

Questions?

Module 7

Deploying Virtual Machines

7-2 Importance

Virtual machines are the foundation of your virtual infrastructure. Deploying VMs effectively involves recognizing the different types of virtual hardware. It also requires skills in creating, cloning and managing VMs and templates, modifying VMs, and updating templates.

7-3 Module Lessons

1. Creating Virtual Machines
2. Virtual Machine Hardware Deep Dive
3. Modifying Virtual Machines
4. Creating Templates and Cloning VMs
5. Introduction to Content Libraries
6. Subscribing to Content Libraries
7. Managing Templates in a Content Library

7-4 Lesson 1: Creating Virtual Machines

7-5 Learner Objectives

- Create and provision a virtual machine
- Explain the importance of VMware Tools
- Install VMware Tools
- Remove a virtual machine

7-6 About Provisioning Virtual Machines

You can create VMs in several ways.

Provisioning Method	Use vSphere Client	Use VMware Host Client
Use the New Virtual Machine wizard.	Yes	Yes
Deploy VMs from existing templates or clones.	Yes	No
Deploy VMs from OVF templates.	Yes	Yes

The optimal method for provisioning VMs for your environment depends on factors such as the size and type of your infrastructure and the goals that you want to achieve.

You can use the New Virtual Machine wizard to create a single VM if no other VMs in your environment meet your requirements, such as a particular operating system or hardware configuration. For example, you might need a VM that is configured only for testing purposes. You can also create a single VM, install an operating system on it, and use that VM as a template from which to clone other VMs.

Deploy VMs, virtual appliances, and vApps stored in Open Virtual Machine Format (OVF) to use a preconfigured VM. A virtual appliance is a VM that typically has an operating system and other software preinstalled. You can deploy VMs from OVF templates that are on local file systems (for example, local disks such as C:), removable media (for example, CDs or USB keychain drives), shared network drives, or URLs.

In addition to using the vSphere Client, you can also use VMware Host Client to create a VM by using OVF files. However, several limitations apply when you use VMware Host Client for this deployment method. For information about OVF and OVA limitations for the VMware Host Client, see *vSphere Single Host Management - VMware Host Client* at https://docs.vmware.com/en/VMware-vSphere/index.html.

7-7 Creating VMs with the New Virtual Machine Wizard

In the vSphere Client, you can use the New Virtual Machine wizard to create a VM.

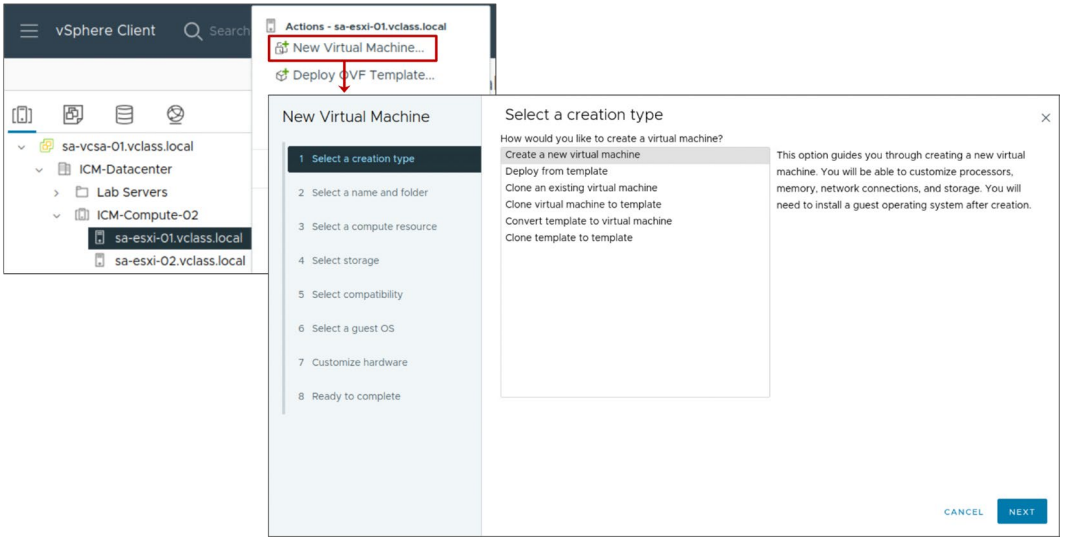

7-8 New Virtual Machine Wizard: Name, Folder, Compute Resource

You can use the New Virtual Machine wizard in the vSphere Client to create a VM.

The New Virtual Machine wizard prompts you for standard information:

- The VM name
- Folder in which to place the VM
- Resource on which the VM runs

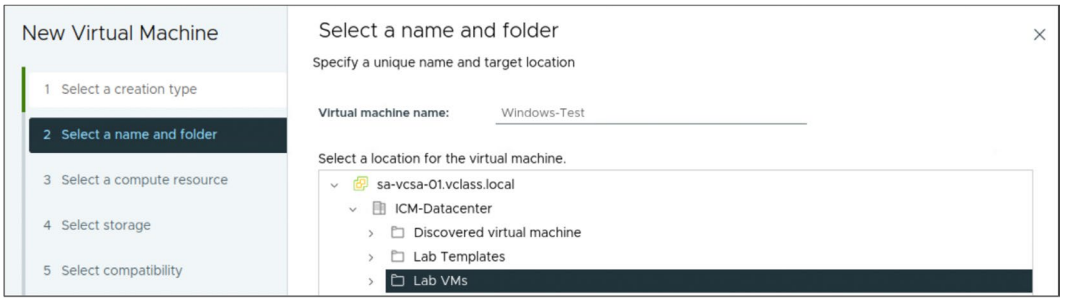

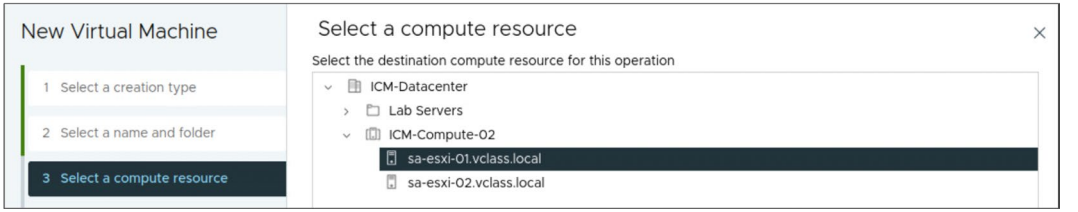

When specifying the compute resource, you can specify a host, a cluster, a vApp, or a resource pool. The VM can access the resources of the selected object.

7-9 New Virtual Machine Wizard: Storage, Compatibility

You select the datastore on which to store the VM's files.

You select the ESXi version that this virtual machine will be compatible with.

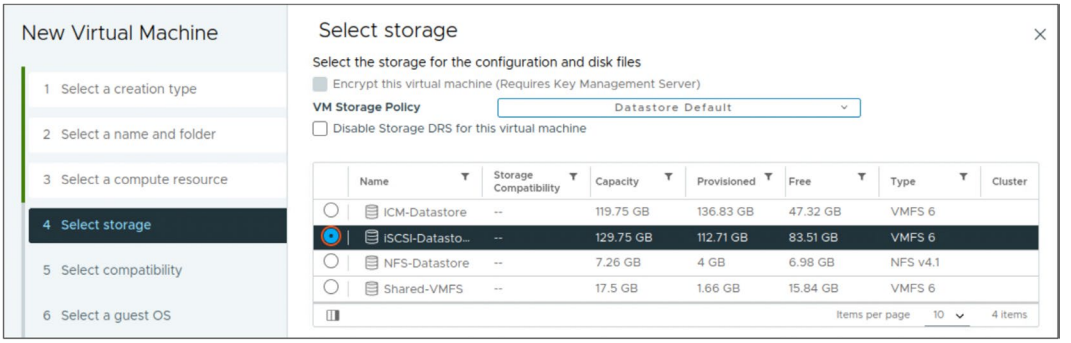

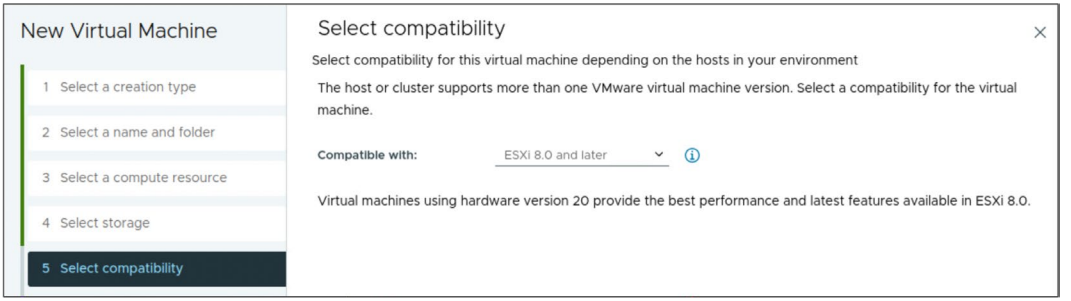

When selecting storage, you can choose the storage policy and datastore. Each datastore might have a different size, speed, availability, and other properties. The available datastores are accessible from compute resource that you selected.

When selecting compatibility, you choose the ESXi host version that you want the virtual machine to be compatible with. To give your VM access to the latest hardware features, select the latest ESXi host version.

7-10 New Virtual Machine Wizard: Guest Operating System

You select the guest OS to be installed in the VM.

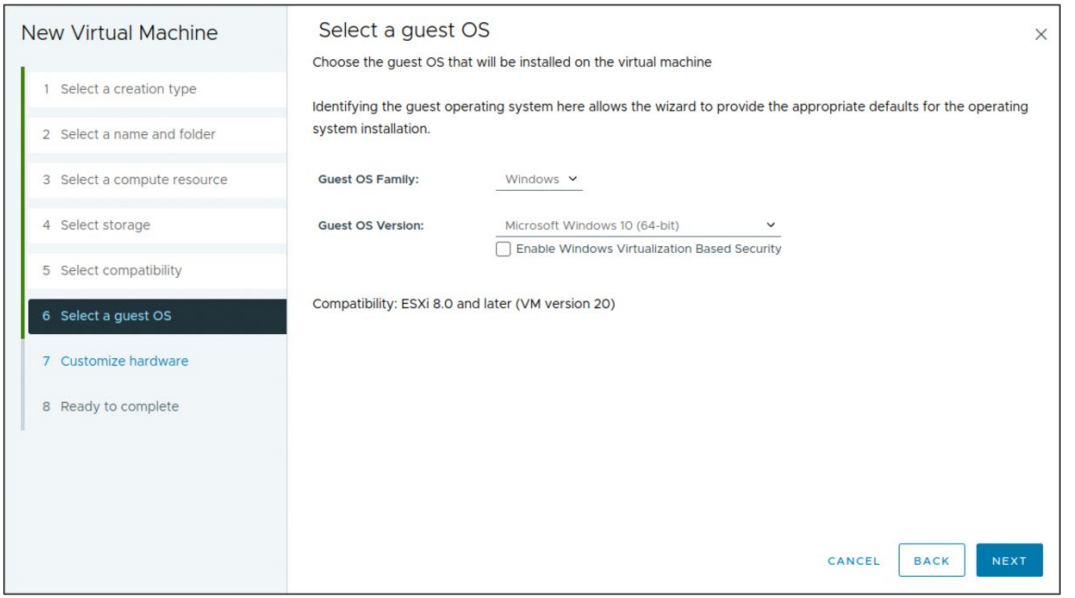

When you select a guest operating system BIOS or EFI is selected by default, depending on the firmware supported by the operating system. If the operating system supports BIOS and EFI, you can change the default by editing the VM after you create it and before you install the guest operating system. If you select EFI, you cannot boot an operating system that supports only BIOS, and the reverse.

7-11 New Virtual Machine Wizard: Virtual Hardware

You can configure the virtual machine hardware. The default values for CPU, memory and hard disk size are based on the guest OS that you selected.

You can also mount the ISO image containing the guest operating system installation files.

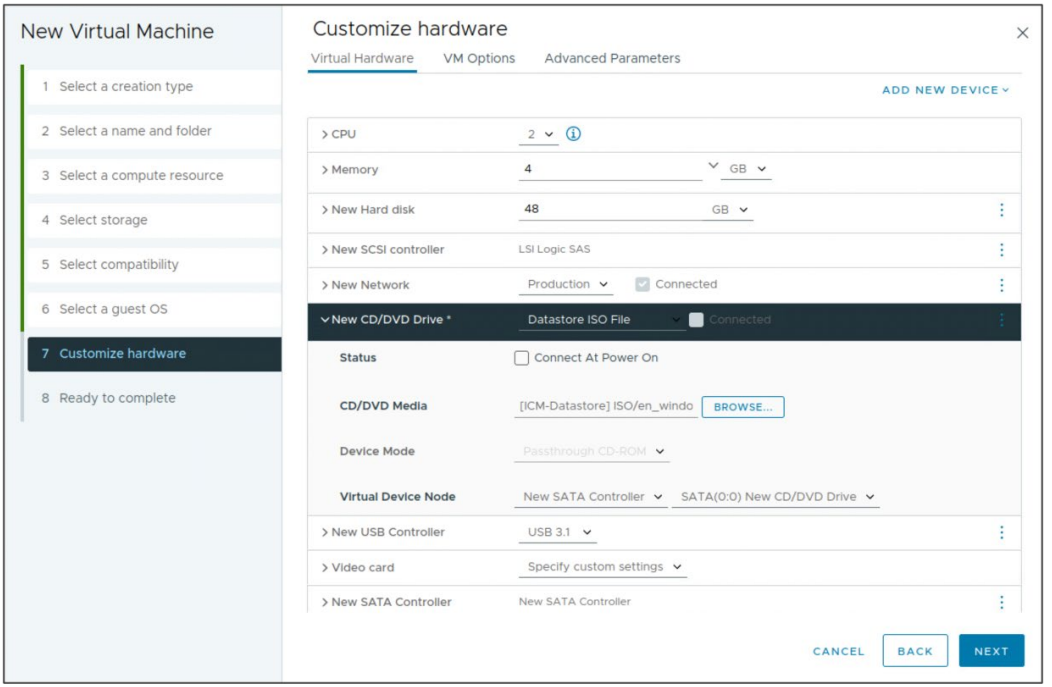

7-12 Installing the Guest Operating System

Installing a guest operating system in your VM is similar to installing it on a physical computer.

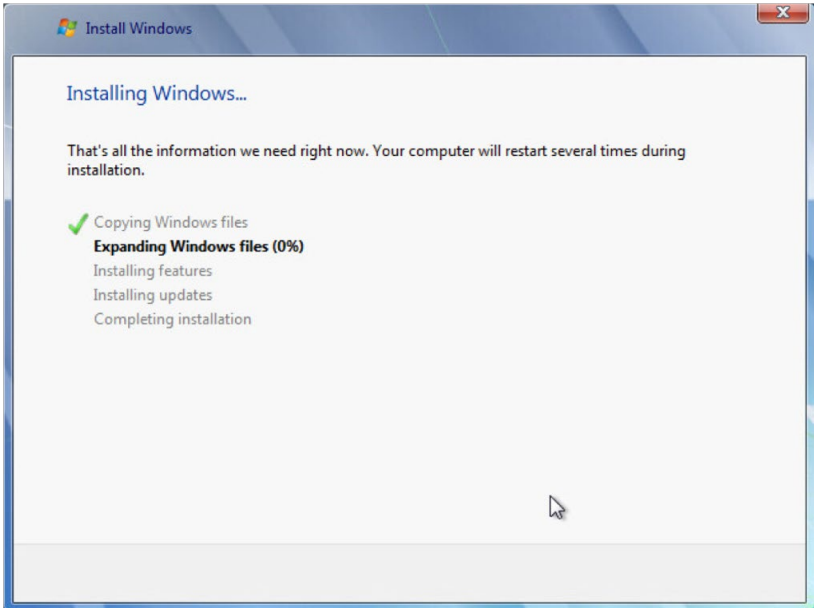

To install the guest operating system, you interact with the VM through the VM console. Using the vSphere Client, you can attach a CD, DVD, or ISO image containing the installation image to the virtual CD/DVD drive.

On the slide, the Windows Server 2008 guest operating system is being installed. You can use the vSphere Client to install a guest operating system. You can also install a guest operating system from an ISO image or a CD. Installing from an ISO image is typically faster and more convenient than a CD installation.

For more information about installing guest operating systems, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/en/VMware-vSphere/index.html.

For more about the supported guest operating systems, see *VMware Compatibility Guide* at https://www.vmware.com/resources/compatibility.

7-13 About VMware Tools

VMware Tools is a set of features that enhance the performance of a VM's guest operating system.

Benefits and features include:

- Device drivers
  - SVGA display
  - VMXNET/VMXNET3
  - Balloon driver for memory management
  - Sync driver for quiescing I/O
  - Paravirtual SCSI controller
- Increased graphical performance
- Improved mouse performance
- Guest OS heartbeat service
- Time synchronization
- Ability to shut down the VM remotely

VMware Tools improves management of the VM by replacing generic operating system drivers with VMware drivers tuned for virtual hardware. You install VMware Tools into the guest operating system. When you install VMware Tools, you install these items:

- The VMware Tools service: This service synchronizes the time in the guest operating system with the time in the host operating system.
- A set of VMware device drivers, with additional Perfmon monitoring options.
- A set of scripts that helps you automate guest operating system operations: You can configure the scripts to run when the VM's power state changes.

VMware Tools enhances the performance of a VM and makes many of the ease-of-use features in VMware products possible:

- Faster graphical performance
- Shared folders between host and guest file systems
- Copying and pasting text, graphics, and files between the virtual machine and the host or client desktop
- Scripting that helps automate guest operating system operations

Although the guest operating system can run without VMware Tools, many VMware features are not available until you install VMware Tools. For example, if VMware Tools is not installed in your VM, you cannot use the shutdown or restart options from the toolbar. You can use only the power options.

7-14 Installing VMware Tools

Ensure that you select the latest version of VMware Tools for your guest operating system.

To find out which VMware Tools ISO images are bundled with vSphere 8, see the vSphere 8 Release Notes.

The method for installing VMware Tools depends on the guest operating system type.

Guest Operating System Type	VMware Tools Installation Method
Microsoft Windows	Install from `windows.iso` for Vista and later guests
Linux	Use one of the following methods: • Install from `linux.iso`. • For later Linux distributions, use `open-vm-tools`, available in various Linux package management systems, such as `yum`, `apt`, or `rpm`.

For details on installing VMware Tools in a Windows or Linux guest operating system, more information about using Open VM tools, see *VMware Tools Administration* at https://docs.vmware.com/en/VMware-Tools/index.html.

7-15 Downloading VMware Tools

You can download a specific version of VMware Tools from the VMware Tools product download page.

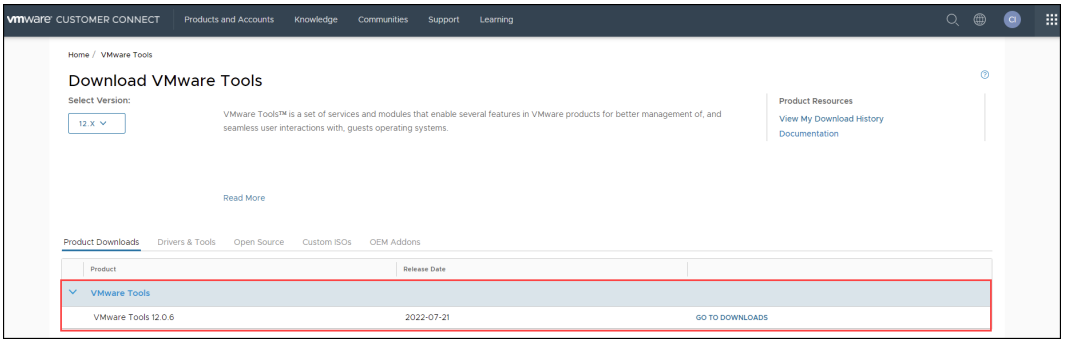

For information about installing VMware Tools in specific guest operating systems, see *VMware Tools Administration* at https://docs.vmware.com/en/VMware-Tools/index.html.

7-16 Deploying OVF Templates

You can deploy any VM or virtual appliance stored in OVF format.

Virtual appliances are preconfigured VMs:

- They are usually designed for a single purpose.
- They are also available from VMware Marketplace.

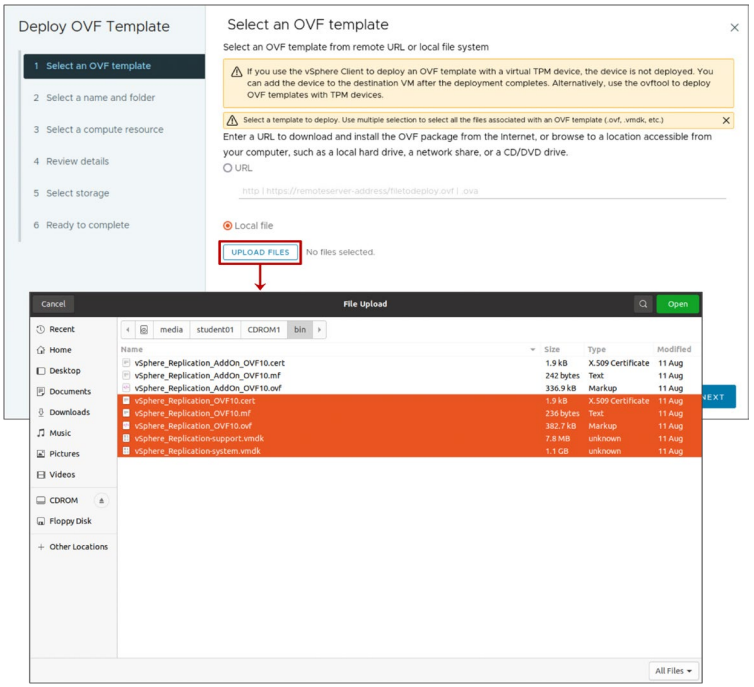

A virtual appliance is a preconfigured VM that typically includes a preinstalled guest operating system and other software. A virtual appliance is usually designed for a specific purpose, for example, to provide a secure web browser, a firewall, or a backup and recovery utility.

A virtual appliance can be added or imported to your vCenter Server system inventory or ESXi inventory. Virtual appliances can be imported from websites such as the VMware Marketplace at https://marketplace.cloud.vmware.com.

Virtual appliances are deployed as OVF templates. OVF is a platform-independent, efficient, extensible, and open packaging and distribution format for VMs. OVF files are compressed, resulting in faster downloads.

To deploy an OVF template using the vSphere Client, right-click the appropriate object, such as a data center or an ESXi host, and select **Deploy OVF Template**. The Deploy OVF Template wizard appears. Click **UPLOAD FILES** to select the template files to deploy.

The vSphere Client validates an OVF file before importing it and ensures that it is compatible with the intended destination server. If the appliance is incompatible with the selected host, you cannot import it.

7-17 Removing VMs

You can remove a VM in the following ways:

- Remove from the inventory:
  - The VM is unregistered from the ESXi host and vCenter.
  - The VM's files remain on the disk.
  - The VM can later be registered (added) to the inventory.
- Delete from disk:
  - All VM files are permanently deleted from the datastore.
  - The VM is unregistered from the ESXi host and vCenter.

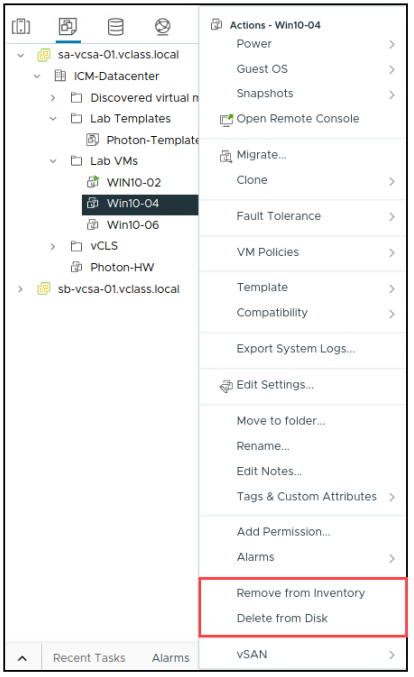

When a VM is removed from the inventory, its files remain at the same storage location, and the VM can be re-registered in the datastore browser.

7-18 Lab 12: Creating and Removing a Virtual Machine

Use the vSphere Client to create a VM, remove a VM from the inventory, and delete a VM from the datastore:

1. Create a Virtual Machine
2. Remove the Virtual Machine from the vCenter Inventory
3. Register the Virtual Machine to Re-Add it to the vCenter Inventory
4. Delete the Virtual Machine from the Datastore

7-19 Lab 13: (Simulation) Installing VMware Tools

Use the vSphere Client to install VMware Tools to an existing Windows VM:

1. Mount the VMware Tools Image to the VM's DVD Drive
2. Install VMware Tools with the VMware Tools Setup Wizard
3. Verify that VMware Tools Is Running in the Guest OS

7-20 Review of Learner Objectives

- Create and provision a virtual machine
- Explain the importance of VMware Tools
- Install VMware Tools
- Remove a virtual machine

7-21 Lesson 2: Virtual Machine Hardware Deep Dive

7-22 Learner Objectives

- Identify the files that make up a VM
- Compare VM hardware versions
- Recognize the virtual hardware components of a VM
- Navigate the vSphere Client and examine VM settings
- Identify methods for accessing a VM console

7-23 Virtual Machine Encapsulation

Each VM is stored either as a collection of files or objects:

- Files in a directory on a VMFS or NFS datastore
- Objects on a vSAN or vSphere Virtual Volumes datastore

Each virtual disk is encapsulated into a single file or object.

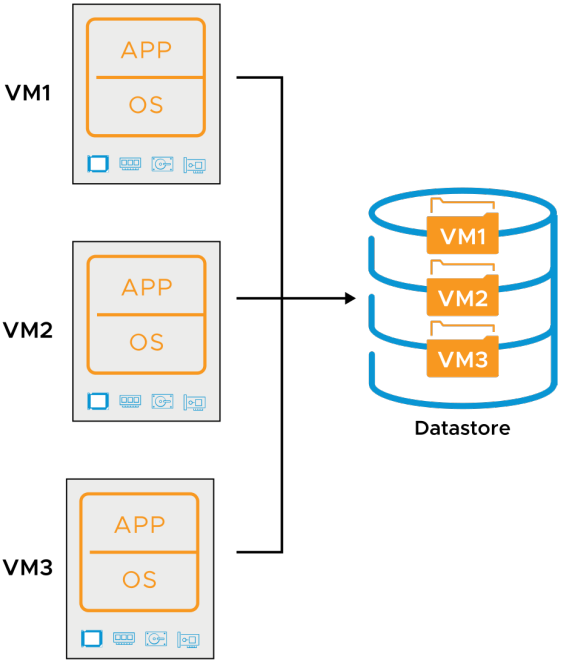

vSphere encapsulates each VM into a few files or objects, making VMs easier to manage and migrate. The files and objects for each VM are stored in a separate folder on a datastore.

7-24 About Virtual Machine Files

A VM includes a set of related files.

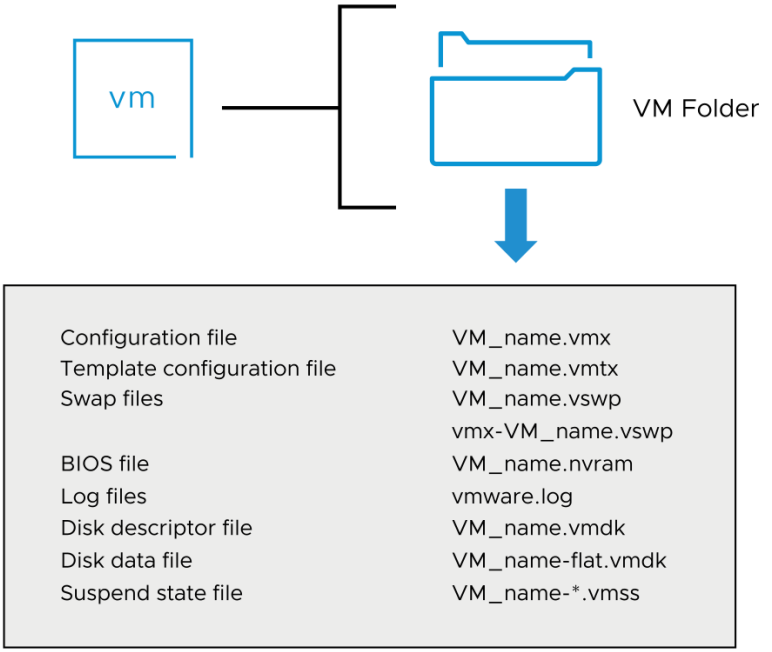

The slide lists some of the files that make up a VM. Except for the log files, the name of each file contains the VM's name <VM_name>. A VM consists of the following files:

- A configuration file (`.vmx`).
- If the VM is converted to a template, a VM template configuration file (`.vmtx`) replaces the VM configuration file (`.vmx`).
- Swap files (`.vswp`) used to reclaim memory during periods of contention.
- A file containing the VM's BIOS or EFI settings (`.nvram`).
- A VM's current log file (`.log`) and a set of files used to archive old log entries (`-#.log`).

  In addition to the current log file, `vmware.log`, up to six archive log files are maintained at one time. For example, `-1.log` to `-6.log` might exist at first.

The next time an archive log file is created, for example, when the VM is powered off and powered back on, the following actions occur: The `-6.log` is deleted, the `-5.log` is recalled to `-6.log`, and so on. Finally, the previous `vmware.log` is recalled to the `-1.log`.

- One or more virtual disk files. The first virtual disk has files `VM_name.vmdk` and `VM_name-flat.vmdk`.

  If the VM has more than one disk file, the file pair for the subsequent disk files is called `VM_name_#.vmdk` and `VM_name_#-flat.vmdk`. # is the next number in the sequence, starting with 1. For example, if the VM called Test01 has two virtual disks, this VM has the `Test01.vmdk`, `Test01-flat.vmdk`, `Test01_1.vmdk`, and `Test01_1-flat.vmdk` files.

The list of files shown on the slide is not comprehensive. For a complete list of all the types of VM files, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/8.0/TBD.

7-25 About VM Virtual Hardware

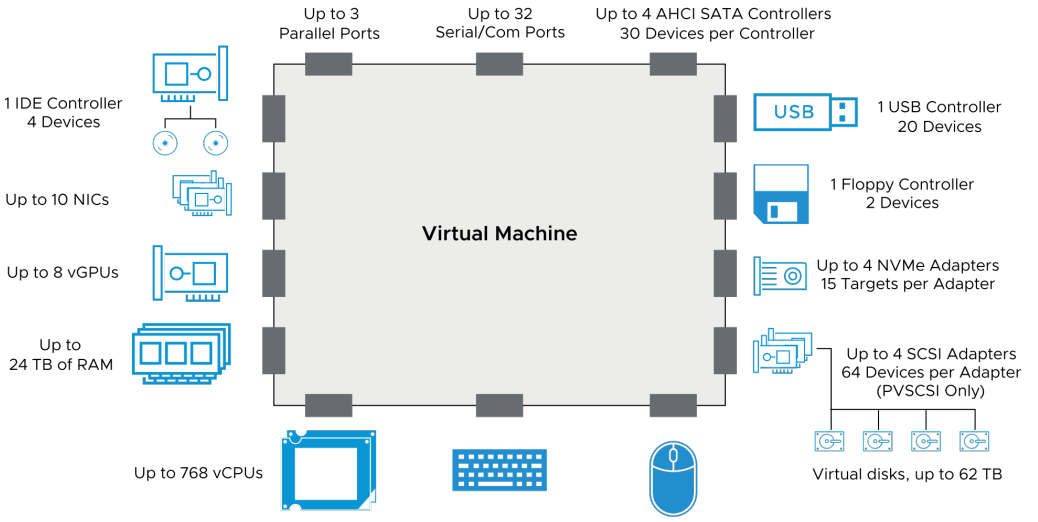

A guest OS accesses hardware devices. The guest OS does not know that these devices are virtual, not physical. All VMs have uniform hardware, except for a few variations that the system administrator can apply. Uniform hardware makes VMs portable across VMware virtualization platforms.

You can configure the VM memory and CPU settings. vSphere supports several latest CPU features, including virtual CPU performance counters. You can add virtual hard disks and NICs. You can also add and configure virtual hardware, such as CD/DVD drives, and SCSI devices. Not all devices are available to add and configure. For example, you cannot add video devices, but you can configure the available video devices and video cards.

You can add multiple USB devices, such as security dongles and mass storage devices, to a VM that exists on an ESXi host to which the devices are physically attached. When you attach a USB device to a physical host, the device is available only to VMs that are located on that host. Those VMs cannot connect to a device on another host in the data center. A USB device is available to only one VM at a time. When you remove a USB device from a VM, it is available to the other VMs that are located on the host.

You can have 64 virtual SCSI targets for a virtual SCSI adapter if you use a PVSCSI driver. However, you can have only 15 virtual SCSI targets for a virtual SCSI adapter other than PVSCSI.

The SATA controller provides access to virtual disks, CD/DVD devices, and USB devices. The SATA virtual controller is presented to a virtual machine as an AHCI SATA controller.

The Virtual Machine Communication Interface (VMCI) is an infrastructure that provides a high-speed communication channel between a VM and the hypervisor. You cannot add or remove VMCI devices.

The VMCI SDK facilitates the development of applications that use the VMCI infrastructure. Without VMCI, VMs communicate with the host using the network layer. Using the network layer adds overhead to the communication. With VMCI, communication overhead is minimal and tasks that require communication can be optimized.

The following types of communication are available:

- Datagrams: Connectionless and similar to UDP queue pairs
- Connection oriented: Similar to TCP

VMCI provides socket APIs that are similar to APIs that are used for TCP/UDP applications. IP addresses are replaced with VMCI ID numbers. For example, you can port netperf to use VMCI sockets instead of TCP/UDP. VMCI is deactivated by default.

For more information about virtual hardware, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/en/VMware-vSphere/index.html.

For a complete list of virtual machine configuration maximums, see VMware Configuration Maximums at https://configmax.vmware.com.

7-26 Virtual Hardware Versions

The virtual hardware version, or VM compatibility level, determines the operating system functions that a VM supports.

Do not use a later version that is not supported by the VMware product.

Compatibility	Virtual Hardware Version
ESXi 8.0	20
ESXi 7.0 U2 and later	19
ESXi 7.0 U1 and later	18
ESXi 7.0 and later	17
ESXi 6.7 U2 and later	15
ESXi 6.7 and later	14
Virtual hardware version 16 is specific to Workstation and Fusion Pro.	

Each release of a VMware product has a corresponding VM hardware version included. The table shows the latest hardware version that each ESXi version supports. Each VM compatibility level supports at least five major or minor vSphere releases.

For more information on virtual machine compatibility, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/en/VMware-vSphere/index.html.

For a complete list of virtual machine hardware versions, see Virtual machine hardware versions KB article at https://kb.vmware.com/s/article/1003746.

7-27 About CPU and Memory

You can add, change, or configure CPU and memory resources to improve VM performance.

The maximum number of virtual CPUs (vCPUs) that you can assign to a VM depends on the following factors:

- The number of logical CPUs on the host
- The type and version of installed guest operating system

A VM running on an ESXi 8.0 host can have up to 768 vCPUs.

The maximum memory size of a VM depends on the VM's compatibility setting.

The maximum memory size of a VM with ESXi 8.0 compatibility running on ESXi 8.0 is 24 TB.

You size the VM's CPU and memory according to the applications and the guest operating system.

You can use the multicore vCPU feature to control the number of cores per virtual socket in a VM. With this capability, operating systems with socket restrictions can use more of the host CPU's cores, increasing overall performance.

A VM cannot have more virtual CPUs than the number of logical CPUs on the host. The number of logical CPUs is the number of physical processor cores, or twice that number if hyperthreading is activated. For example, if a host has 128 logical CPUs, you can configure the VM for 128 vCPUs.

You can set most of the memory parameters during VM creation or after the guest operating system is installed. Some actions require that you power off the VM before changing the settings.

The memory resource settings for a VM determine how much of the host's memory is allocated to the VM.

The virtual hardware memory size determines how much memory is available to applications that run in the VM. A VM cannot benefit from more memory resources than its configured virtual hardware memory size.

You can reconfigure the amount of memory allocated to a VM to enhance performance. Maximum memory size for a VM depends on the VM's compatibility setting.

7-28 Compute Maximums

vSphere provides compute maximums, which are available at https://configmax.vmware.com.

	vSphere 8
Virtual CPUs per VM	768
Memory per VM	24 TB
CPUs per host	896
Memory per host	24 TB
Hosts per cluster	96

For all configuration maximums, see https://configmax.vmware.com/

7-29 About Virtual Storage

Virtual disks are connected to virtual storage adapters.

The ESXi host offers VMs several choices in storage adapters:

- BusLogic Parallel
- LSI Logic Parallel
- LSI Logic SAS
- VMware Paravirtual SCSI
- AHCI SATA controller
- Virtual NVMe

Storage adapters provide connectivity for your ESXi host to a specific storage unit or network.

ESXi supports different classes of adapters, including SCSI, iSCSI, RAID, Fibre Channel, Fibre Channel over Ethernet (FCoE), and Ethernet. ESXi accesses the adapters directly through device drivers in the VMkernel:

- BusLogic Parallel: The Mylex (BusLogic) BT/KT-958 compatible host bus adapter.
- LSI Logic Parallel: The LSI Logic LSI53C10xx Ultra320 SCSI I/O controller is supported.
- LSI Logic SAS: The LSI Logic SAS adapter has a serial interface.
- VMware Paravirtual SCSI: A high-performance storage adapter that can provide greater throughput and lower CPU use.
- AHCI SATA controller: Provides access to virtual disks and CD/DVD devices. The SATA virtual controller appears to a VM as an AHCI SATA controller.
- Virtual NVMe: NVMe is a protocol for attaching and accessing flash storage devices to the PCI Express bus. NVMe is an alternative to existing block-based server storage I/O access protocols.

7-30 About Thick-Provisioned Virtual Disks

Thick provisioning uses all the defined disk space at the creation of the virtual disk, regardless of how much disk space is actually used by the guest operating system file system.

Thick-provisioned disk types are either eager zeroed or lazy zeroed:

- In an eager-zeroed thick-provisioned disk, every block is prefilled with a zero.
- In a lazy-zeroed thick-provisioned disk, a block is filled with zeroes before data is written for the first time.

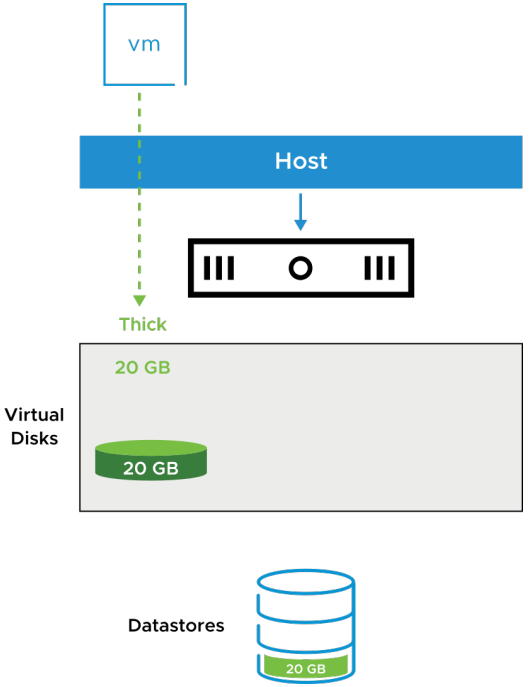

In a lazy-zeroed thick-provisioned disk, space required for the virtual disk is allocated during creation. Data remaining on the physical device is not erased during creation. Later, the data is zeroed out on demand on first write from the VM. This disk type is the default.

In an eager-zeroed thick-provisioned disk, the space required for the virtual disk is allocated during creation. Data remaining on the physical device is zeroed out when the disk is created.

7-31 About Thin-Provisioned Virtual Disks

With thin provisioning, VMs use the disk space as needed:

- Virtual disks use only the capacity needed to hold the current files.
- The VM always sees the full allocated disk size.

Run the `unmap` command to reclaim the unused space from the virtual disks.

Reporting and alerts help manage allocations and capacity.

You can mix thick and thin formats.

The following examples show efficient use of storage:

- Provisioned space for virtual disks: 140 GB
- Available datastore capacity: 100 GB
- Used datastore capacity: 80 GB

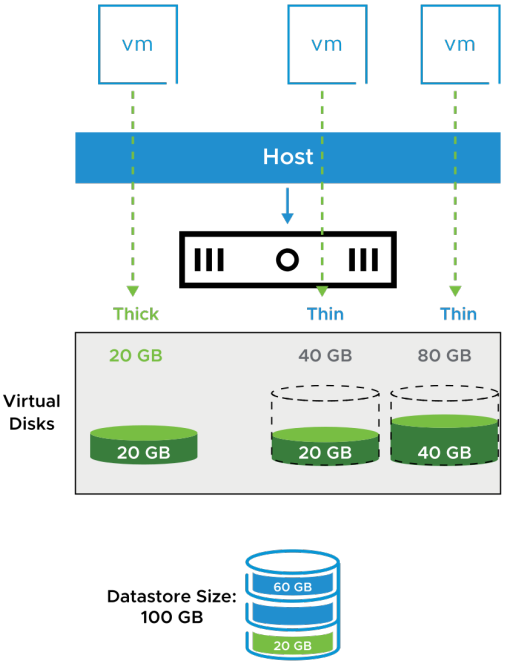

A thin-provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can expand to the maximum capacity allocated to it.

Thin provisioning provides alarms and reports that track allocation versus the current use of storage capacity. Storage administrators can use thin provisioning to optimize the allocation of storage for virtual environments. With thin provisioning, users can optimally but safely use the available disk space through overallocation.

ESXi supports reclamation of free space, which is also called the `unmap` command, which helps the storage array reclaim unused free space. For more information about space reclamation and how to run the SCSI `unmap` command, see *vSphere Storage* at https://docs.vmware.com/en/VMware-vSphere/index.html.

7-32 Managing Datastores Containing Thin-Provisioned Disks

When the total provisioned space of thin-provisioned disks is greater than the size of the datastore, the datastore becomes overcommitted.

To actively monitor datastore capacity:

- Set alarms to send notifications about:
  - Datastore disk overallocation
  - VM disk use
- Use reporting to view space use.

To actively manage datastore capacity:

- Increase datastore capacity when necessary.
- Use vSphere Storage vMotion to mitigate space use problems on a particular datastore.

Using thin-provisioned virtual disks for your VMs is a way to make the most of your datastore capacity. But if your datastore is not sized properly, it can become overcommitted. A datastore becomes overcommitted when the full capacity of its thin-provisioned virtual disks is greater than the datastore's capacity.

When a datastore reaches capacity, the vSphere Client prompts you to provide more space on the underlying VMFS datastore. VMs will pause if they try to write data to the datastore.

Monitor your datastore capacity by setting alarms to alert you about how much a datastore's disks are fully allocated or how much disk space a VM is using.

Manage your datastore capacity by dynamically increasing the size of your datastore when necessary. You can also use vSphere Storage vMotion to mitigate space use issues.

For example, with vSphere Storage vMotion, you can migrate a VM off a datastore. The migration can be done by changing from virtual disks of thick format to thin format at the target datastore.

7-33 Thick-Provisioned and Thin-Provisioned Disks

Virtual disk options differ in terms of creation time, block allocation, layout, and zeroing out of allocated file blocks.

	Thick Provisioned Lazy-Zeroed	**Thick Provisioned Eager-Zeroed**	**Thin Provisioned**
Creation time	Fast	Slow and proportional to disk size	Fastest
Block allocation	Fully preallocated	Fully preallocated	Allocated and zeroed out on demand at first write to block
Virtual disk layout	Higher chance of contiguous file blocks	Higher chance of contiguous file blocks	Layout varies according to the dynamic state of the volume at time of block allocation
Zeroing out of allocated file blocks	File blocks are zeroed out when each block is first written to	File blocks are allocated and zeroed out when disk is created	File blocks are zeroed out when blocks are allocated

7-34 About Virtual Networks

VMs and physical machines communicate through a virtual network.

When you configure networking for a VM, you select or change the following settings:

- Network adapter type
- Port group to connect to
- Network connection state
- Whether to connect to the network when the VM powers on

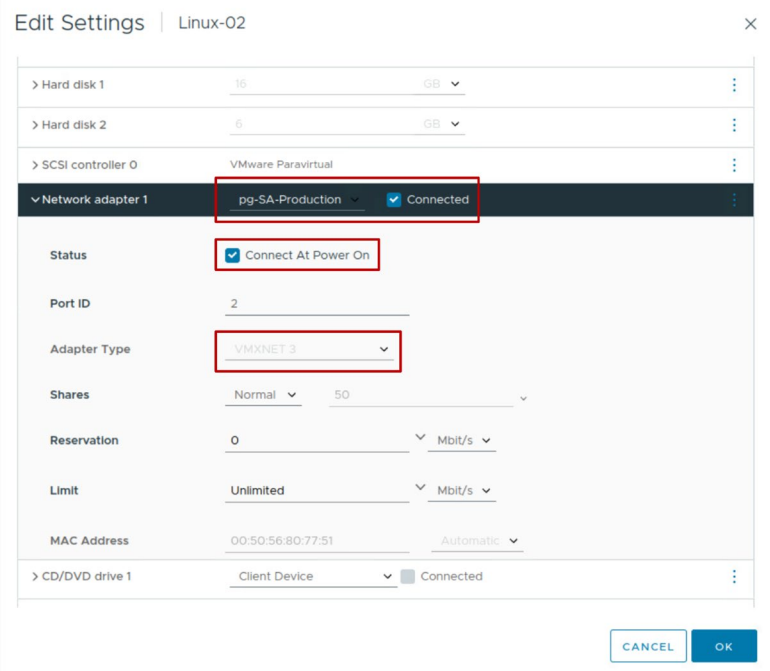

For more information about virtual networks, see *vSphere Networking* at https://docs.vmware.com/en/VMware-vSphere/index.html.

7-35 About Virtual Network Adapters

When you configure a VM, you can add network adapters (NICs) and specify the adapter type. Whenever possible, select VMXNET3.

Network Adapter Type	Description
E1000-E1000E	Emulated version of an Intel Gigabit Ethernet NIC, with drivers available in newer guest operating systems.
VMXNET3	Available only with VMware Tools.
Flexible	Can function as either a Vlance or VMXNET adapter.
PVRDMA	Paravirtualized device that provides improved virtual device performance. It provides an RDMA-like interface for vSphere guests.
SR-IOV pass-through	Allows VM and physical adapter to exchange data without using the VMkernel as an intermediary.

The types of network adapters that are available depend on the following factors:

- VM compatibility level (or hardware version), which depends on the host that created or most recently updated it. For example, the VMXNET3 virtual NIC requires hardware version 14 or later.
- Whether the VM compatibility is updated to the latest version for the current host.
- Guest operating system

The following NIC types are supported:

- E1000E: Emulated version of the Intel 82574 Gigabit Ethernet NIC. E1000E is the default adapter for Windows 8 and Windows Server 2012.
- E1000: Emulated version of the Intel 82545EM Gigabit Ethernet NIC, with drivers available in newer guest operating systems, including Windows XP and later and Linux versions 2.4.19 and later.
- Flexible: Identifies itself as a Vlance adapter when a VM starts, but initializes itself and functions as either a Vlance or a VMXNET adapter, depending on which driver initializes it. With VMware Tools installed, the VMXNET driver changes the Vlance adapter to the higher performance VMXNET adapter.
- Vlance: Emulated version of the AMD 79C970 PCnet32 LANCE NIC, an older 10 Mbps NIC with drivers available in 32-bit legacy guest operating systems. A VM configured with this network adapter can use its network immediately.

- VMXNET3: A paravirtualized NIC designed for performance. VMXNET3 offers all the features available in VMXNET2 and adds several new features, such as multiqueue support (also known as Receive Side Scaling in Windows), IPv6 offloads, and MSI/MSI-X interrupt delivery.
- With PVRDMA, multiple guests can access the RDMA device by using verbs API, an industry-standard interface. A set of these verbs was implemented to expose an RDMA-capable guest device (PVRDMA) to applications. The applications can use the PVRDMA guest driver to communicate with the underlying physical device. PVRDMA supports RDMA, providing the following functions:
  - OS bypass
  - Zero-copy
  - Low latency and high bandwidth
  - Less power use and faster data access
- SR-IOV pass-through: Representation of a virtual function on a physical NIC with SR-IOV support. This adapter type is suitable for VMs that require more CPU resources or where latency might cause failure. If VMs are sensitive to network delay, SR-IOV can provide direct access to the virtual functions of supported physical NICs, bypassing the virtual switches and reducing overhead.

  SR-IOV pass-through is available for Red Hat Enterprise Linux 6 and later, and Windows Server 2008 R2 with SP2. An operating system release might contain a default virtual function driver for certain NICs. For others, you must download and install it from a location provided by the NIC or host vendor.

7-36 About PCI Passthrough Devices

Passthrough devices help your environment use resources efficiently and improve performance.

You connect a VM's guest OS to PCI or PCIe passthrough devices that are configured on an ESXi host.

PCI Passthrough Device	Description
vSphere DirectPath I/O	• VM accesses directly the physical PCI or PCIe device on a specific host. • VM is restricted to that particular host.
vSphere Dynamic DirectPath I/O	• PCI or PCIe passthrough device is not directly mapped to the VM. • Allows vSphere DRS to place a VM on any ESXi host in the cluster that provides the assigned passthrough device
NVIDIA GRID GPU	• Graphics device that uses the NVIDIA GRID vGPU technology • Lets VMs use partial, full, or multiple GPU allocations

With vSphere DirectPath I/O, you can directly access devices, such as high-performance graphics or sound cards. A VM specifies a PCI or PCIe device by using its hardware address. When vSphere DirectPath I/O devices are made available to a virtual machine, you cannot perform certain operations on the virtual machine. These operations include suspending, migration with vMotion, and taking or restoring snapshots of the virtual machine.

With vSphere Dynamic DirectPath I/O, the hardware address of the passthrough device is no longer directly mapped to the VM. vSphere Dynamic DirectPath I/O is useful on hosts that have PCI passthrough devices and for virtualized devices that require a directly assigned hardware device to back it. vSphere Dynamic DirectPath I/O allows vSphere DRS to identify a host within the cluster that has an available device with the same vendor and model name. vSphere DRS takes action only during VM power on and does not perform any load balancing actions.

NVIDIA GRID vGPU devices optimize complex graphics operations and make them run at high performance without overloading the CPU. NVIDIA GRID vGPU provides unparalleled graphics performance and scalability by sharing a single physical GPU among multiple virtual machines as separate vGPU-enabled passthrough devices.

For more information on vSphere Dynamic DirectPath I/O, see *vSphere 7 - Assignable Hardware* at https://blogs.vmware.com/vsphere/2020/03/vsphere-7-assignable-hardware.html.

7-37 Other Virtual Devices

A VM must have a vCPU and virtual memory. The addition of other virtual devices makes the VM more useful:

- CD/DVD drive: For connecting to a CD, DVD, or ISO image.
- USB 3.0 and 3.1: Supported with host-connected and client-connected devices.
- Floppy drive: For connecting a VM to a floppy drive or a floppy image.
- Generic SCSI devices: A VM can be connected to additional SCSI adapters.
- vGPUs: A VM can use GPUs on the physical host for high-computation activities.
- Precision Clock: Provides a virtual machine with access to the system time of the primary ESXi host.
- vTPM: Trusted Platform Module 2.0 virtual cryptoprocessor, providing hardware-based security-related functions.

Virtual CPU (vCPU) and virtual memory are the minimum required virtual hardware. Having a virtual hard disk, virtual NICs, and other virtual devices make the VM more useful.

For information about available virtual machine hardware and about adding virtual hardware to a VM, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/en/VMware-vSphere/8.0/TBD.

7-38 About the Virtual Machine Console

The VM console provides the mouse, keyboard, and screen features to control the VM.

You can use the remote console or the Web console to connect to client devices.

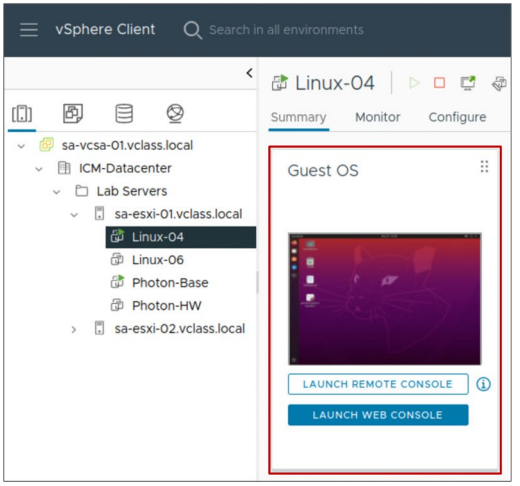

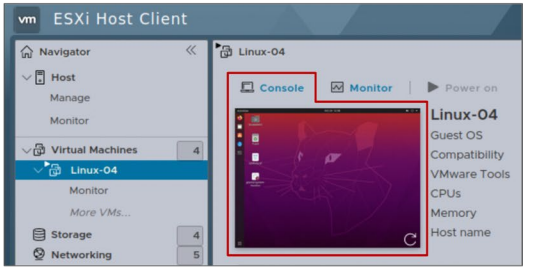

You can open the VM console from the vSphere Client:

- The Web console displays the VM console in a separate browser tab.
- The remote console requires downloading the VMware Remote Console (VMRC) standalone application, which opens in a separate window. The VMware Remote Console standalone application lets you connect to client devices and launch virtual machine consoles on remote hosts.

You use the VM console to access the BIOS or EFI of the VM, install an operating system on a VM, power the VM on and off, and reset the VM.

The VM console is normally not used to connect to the VM for daily tasks. Remote Desktop Connection, Virtual Network Connection, or other options are normally used to connect to the virtual desktop. The VM console is used for tasks such as power cycling, configuring hardware, and troubleshooting network issues.

7-39 Lab 14: Adding Virtual Hardware

Use the vSphere Client to examine a virtual machine's configuration and add virtual hardware to the virtual machine:

1. Examine a Virtual Machine's Configuration
2. Add Virtual Hard Disks to the Virtual Machine
3. Compare Thin-Provisioned and Thick-Provisioned Disks

7-40 Review of Learner Objectives

- Identify the files that make up a VM
- Compare VM hardware versions
- Recognize the virtual hardware components of a VM
- Navigate the vSphere Client and examine VM settings
- Identify methods for accessing a VM console

7-41 Lesson 3: Modifying Virtual Machines

7-42 Learner Objectives

- Describe virtual machine settings and options
- Add a hot-pluggable device
- Dynamically increase the size of a virtual disk

7-43 Modifying Virtual Machine Settings

You can modify a VM's configuration by editing the VM's settings:

- Add virtual hardware:
  - You can add some hardware while the VM is powered on.
- Remove virtual hardware:
  - You can remove some hardware only when the VM is powered off.
- Set VM options.
- Control a VM's CPU and memory resources.

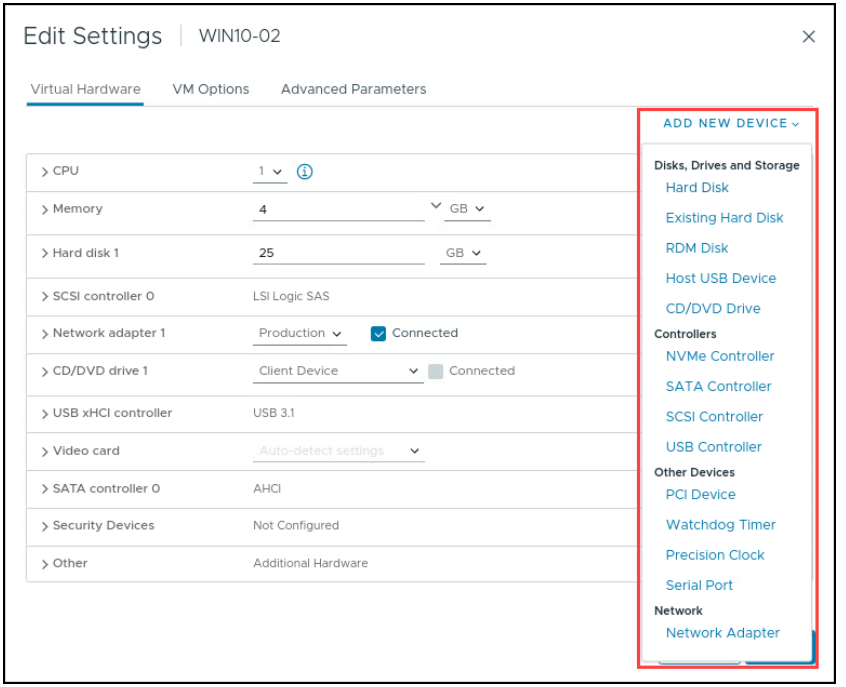

You might have to modify a VM's configuration, for example, to add a network adapter or a virtual disk. You can make all VM changes while the VM is powered off. Some VM hardware changes can be made while the VM is powered on.

7-44 Hot-Pluggable Devices

With the hot plug option, you can add resources to a running VM.

Examples of hot-pluggable devices:

- USB controllers
- Ethernet adapters
- Hard disk devices

With supported guest operating systems, you can also add CPU and memory while the VM is powered on.

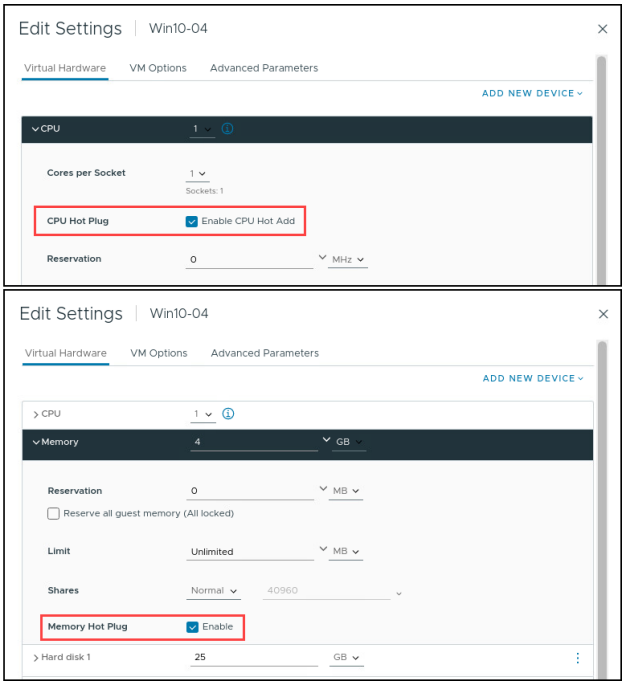

Adding devices to a physical server or removing devices from a physical server requires that you physically interact with the server in the data center. When you use VMs, resources can be added dynamically without a disruption in service. You must shut down a VM to remove hardware, but you can reconfigure the VM without entering the data center.

You can add CPU and memory while the VM is powered on. These features are called the CPU Hot Add and Memory Hot Plug, which are supported only on guest operating systems that support hot-pluggable functionality. These features are deactivated by default. To use these hot-plug features, the following requirements must be satisfied:

- You must install VMware Tools.
- The VM must use hardware version 11 or later.
- The guest operating system in the VM must support CPU and memory hot-plug features.
- The hot-plug features must be activated in the **CPU** or **Memory** settings on the **Virtual Hardware** tab.

If virtual NUMA is configured with virtual CPU hot-plug settings, the VM is started without virtual NUMA. Instead, the VM uses UMA (Uniform Memory Access).

7-45 Dynamically Increasing Virtual Disk Size

You can increase the size of a virtual disk that belongs to a powered-on VM.

- It must not have snapshots attached.
- It might require system tools to make the new space usable.

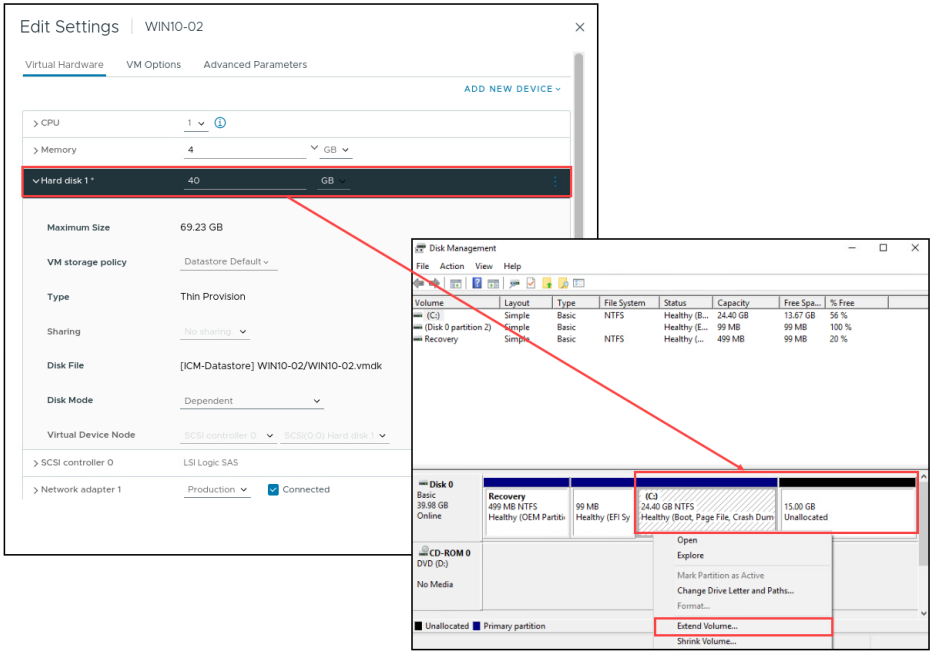

After you increase the size of a virtual disk, you might need to increase the size of the file system on this disk. Use the appropriate tool in the guest OS to configure the file system to use the newly allocated disk space.

7-46 Inflating Thin-Provisioned Disks

Thin-provisioned virtual disks can be converted to a thick, eager-zeroed format.

Choose one of the following methods to inflate a thin-provisioned disk on a VM that is either powered on or off:

- Select the VM's file with the `.vmdk` extension and select **Inflate**.
- Select thick-provisioned when you use vSphere Storage vMotion to migrate the VM to a different datastore.

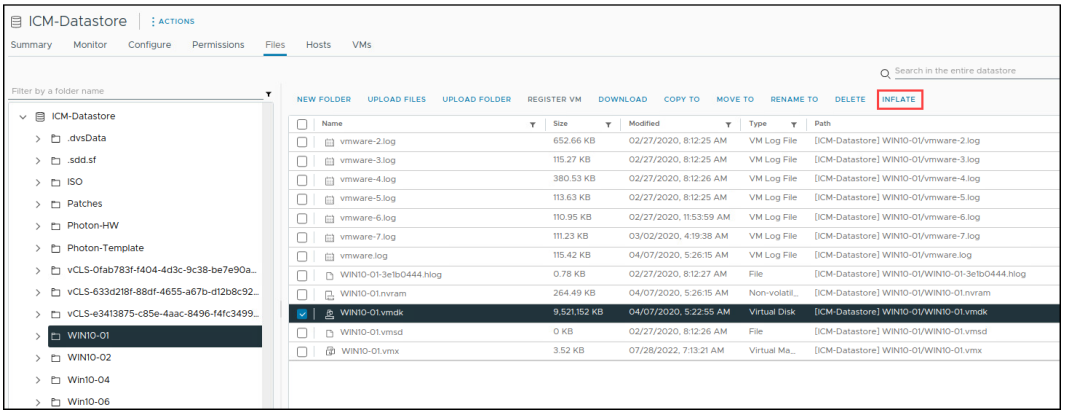

When you inflate a thin-provisioned disk, the inflated virtual disk occupies the entire datastore space originally provisioned to it. Inflating a thin-provisioned disk converts a thin disk to a virtual disk in thick-provisioned, eager-zeroed format.

7-47 VM Options: General Settings

You can use the **VM Options** tab to modify properties such as the display name for the VM and the type of guest operating system that is installed.

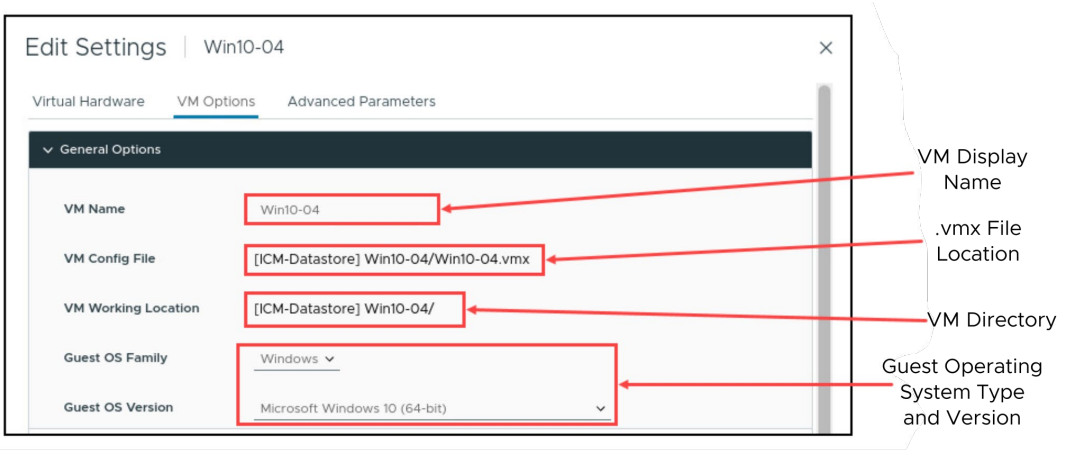

Under General Options, you can view the location and name of the configuration file (with the `.vmx` extension) and the location of the VM's directory.

You can select the text for the configuration file and the working location to copy and paste them into a document. However, only the display name and the guest operating system type can be modified.

Changing the display name does not change the names of all the VM files or the directory that the VM is stored in. When a VM is created, the filenames and the directory name associated with the VM are based on its display name. But changing the display name later does not modify the filename and the directory name.

7-48 VM Options: VMware Tools Settings

You can use the VMware Tools controls to customize the power buttons on the VM.

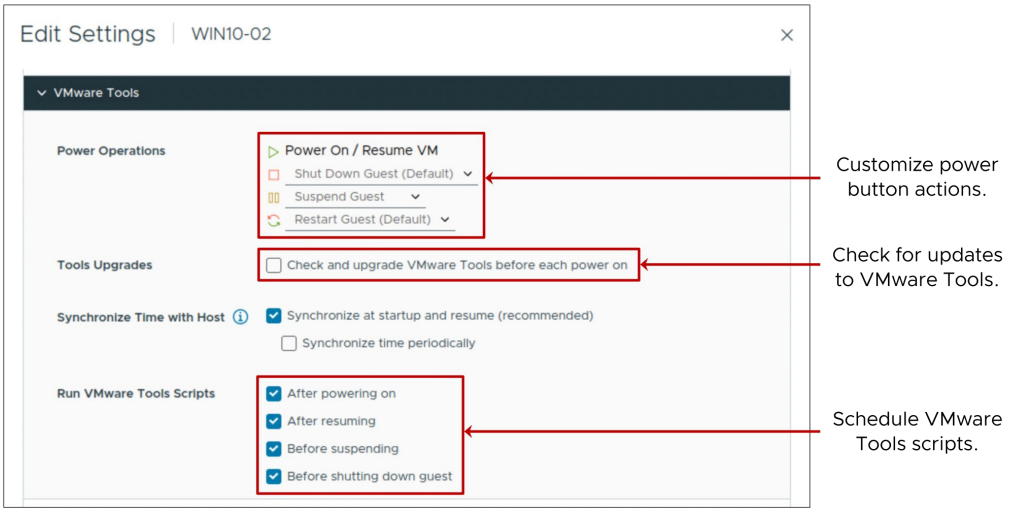

When you use the VMware Tools controls to customize the power buttons on the VM, the VM must be powered off.

You can select the **Check and upgrade VMware Tools before each power on** check box to check for a newer version of VMware Tools. If an ISO image containing a newer version is found, VMware Tools is upgraded when the VM is power cycled.

When you select the **Synchronize guest time with host** check box, the guest operating system's clock synchronizes with the host.

For information about time keeping best practices for the guest operating systems that you use, see VMware knowledge base articles at http://kb.vmware.com/kb/1318 and http://kb.vmware.com/kb/1006427.

7-49 VM Options: VM Boot Settings

Occasionally, you might need to set the VM boot options.

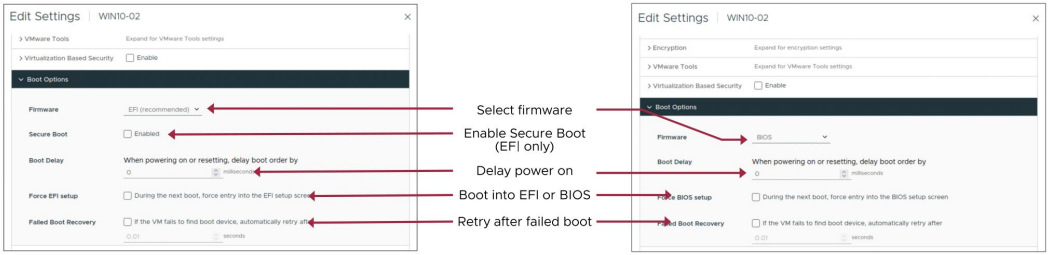

When you build a VM and select a guest operating system, **BIOS** or **EFI** is selected automatically, depending on the firmware supported by the operating system. If the operating system supports BIOS and EFI, you can change the boot option as needed. However, you must change the option before installing the guest OS.

UEFI Secure Boot is a security standard that helps ensure that only manufacturer-trusted software is used when booting your PC. In an OS that supports UEFI Secure Boot, each piece of boot software is signed, including the bootloader, the operating system kernel, and operating system drivers. If you configure Secure Boot for a VM, you can load only signed drivers into that VM.

With the **Boot Delay** value, you can set a delay between the time when a VM is turned on and the guest OS starts to boot. A delayed boot can help stagger VM start-ups when several VMs are powered on.

You can change the BIOS or EFI settings. For example, you might want to force a VM to start from a CD/DVD. The next time the VM powers on, it goes straight into the BIOS. A forced entry into the firmware setup is much easier than powering on the VM, opening a console, and quickly trying to press the F2 key.

With the Failed Boot Recovery setting, you can configure the VM to retry booting after 10 seconds (the default) if the VM fails to find a boot device.

7-50 Lab 15: Modifying Virtual Machines

Modify a VM's memory size, increase a VM's storage size, and rename a VM:

1. Adjust Memory Allocation on a Powered-On Virtual Machine
2. Increase the Size of a Virtual Disk
3. Configure the Guest OS to Recognize the Additional Disk Space
4. Rename a Virtual Machine in the vCenter Inventory

7-51 Review of Learner Objectives

- Describe virtual machine settings and options
- Add a hot-pluggable device
- Dynamically increase the size of a virtual disk

7-52 Lesson 4: Creating Templates and Cloning VMs

7-53 Learner Objectives

- Create a template of a virtual machine
- Deploy a virtual machine from a template
- Clone a virtual machine
- Create customization specifications for guest operating systems

7-54 About Templates

A template is a static image of a virtual machine. You use templates to create and provision new VMs.

A template typically includes:

- A guest operating system
- One or more applications
- A specific VM hardware configuration
- VMware Tools

To use templates, you must be connected to vCenter.

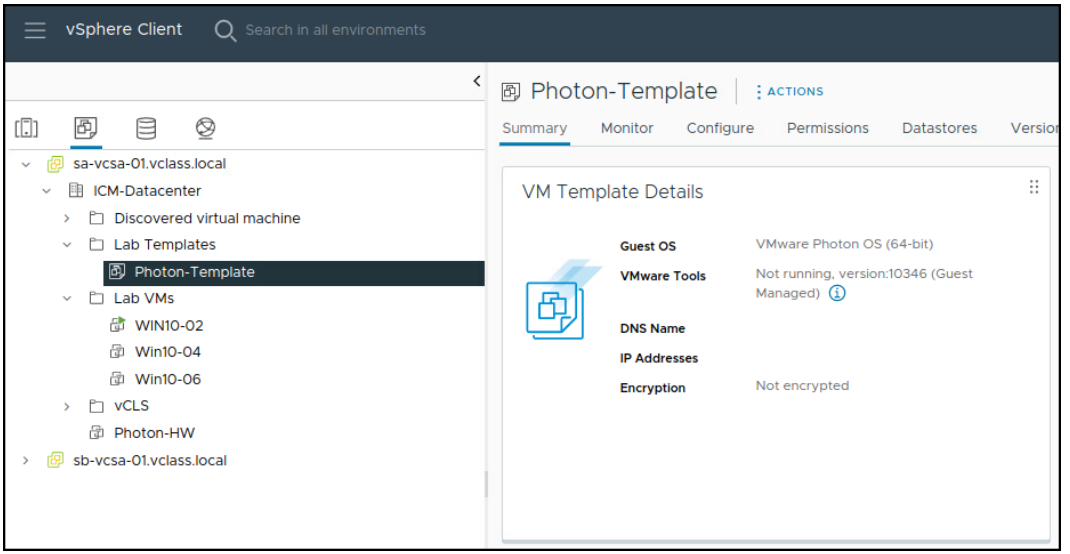

Creating templates makes the provisioning of virtual machines much faster and less error-prone than provisioning physical machines and creating a VM by using the New Virtual Machine wizard.

Templates coexist with VMs in the inventory. You can organize collections of VMs and templates into arbitrary folders and apply permissions to VMs and templates. You can change VMs into templates without having to make a full copy of the VM files and create an object.

You can deploy a VM from a template. The deployed VM is added to the folder that you selected when creating the template.

To create and use templates, you must be connected to vCenter. You cannot use templates if you use VMware Host Client to manage a host directly.

7-55 Creating a Template: Clone VM to Template

You can create templates using different methods. One method is to clone the VM to a template. The VM can be powered on or off.

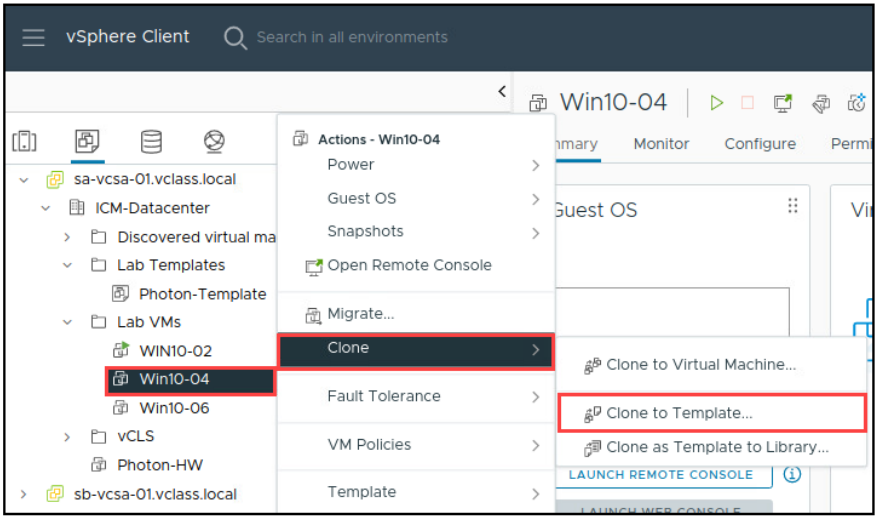

One of the choices that the **Clone to Template** option offers is the type of format for storing the VM's virtual disks:

- Same format as source
- Thin-provisioned format
- Thick-provisioned lazy-zeroed format
- Thick-provisioned eager-zeroed format

7-56 Creating a Template: Convert VM to Template

You can create a template by converting a VM to a template. In this case, the VM must be powered off.

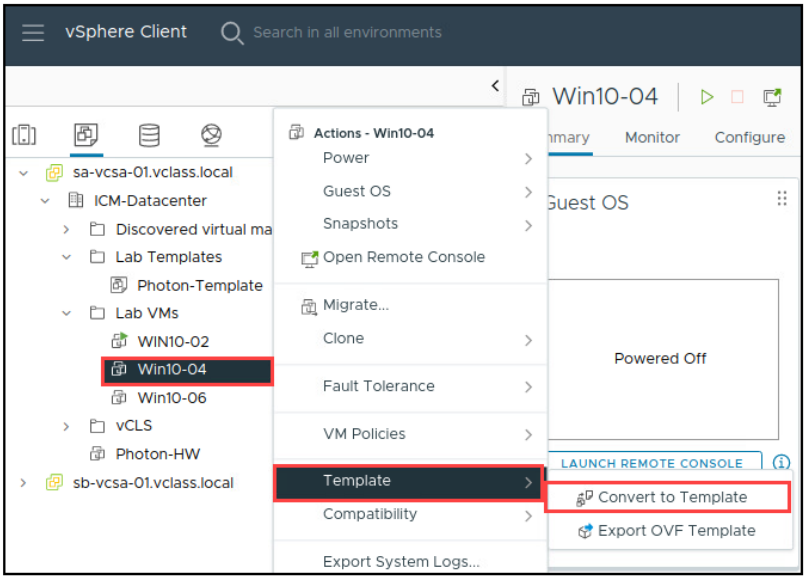

The **Convert to Template** option does not offer a choice of format and leaves the VM's disk file unchanged. The VM's configuration file (`.vmx`) is replaced with a template configuration file (`.vmtx`).

7-57 Creating a Template: Clone a Template

You can create a template from an existing template.

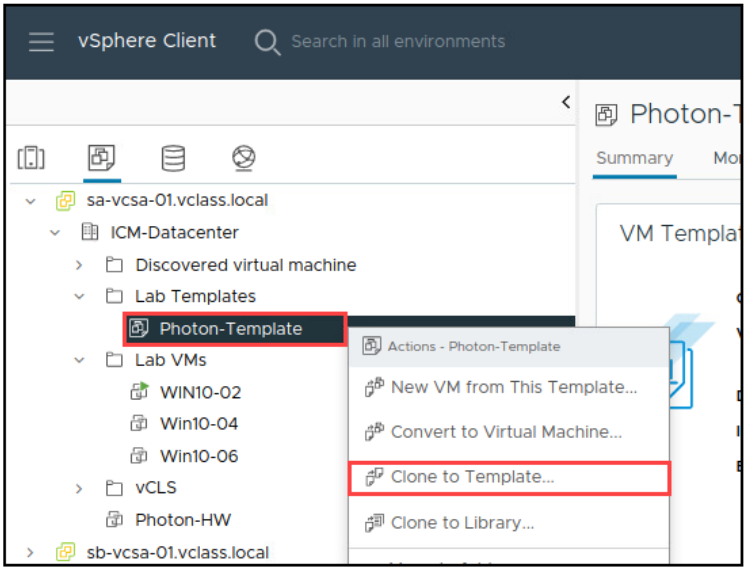

7-58 Updating Templates

You update a template to include new patches, add and remove virtual hardware, upgrade VMware Tools, update the VM hardware version, and install new applications.

To update a template:

1. Convert the template to a VM.
   - VMs cannot be deployed from this template while it is a VM.
2. Place the VM on an isolated network to prevent user access.
3. Make appropriate changes to the VM.
4. Convert the VM to a template.

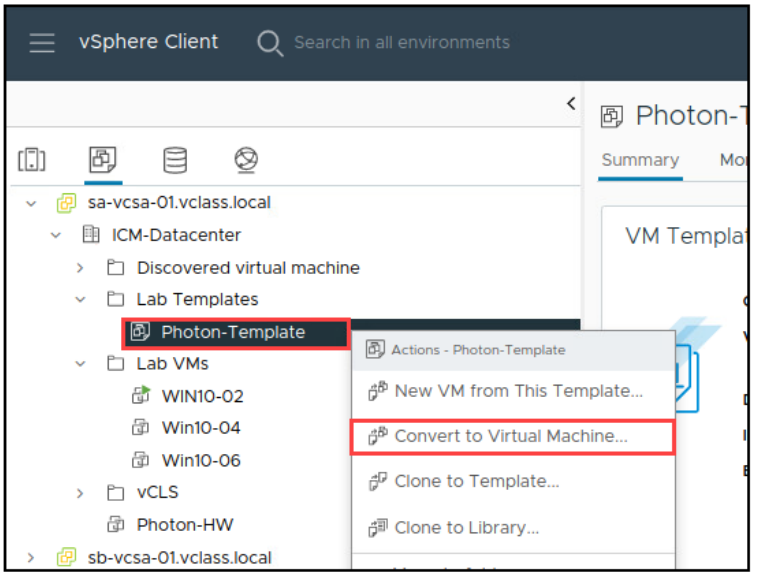

To update your template to include new patches or software, you do not need to create another template. Instead, you convert the template to a VM. You can then power on the VM and make changes to the VM.

For added security, you might want to prevent users from accessing the VM while you update it. To prevent access, either disconnect the VM from the network or place it on an isolated network.

Log in to the VM's guest operating system and apply the patch or install the software. When you finish, power off the VM and convert it to a template again.

7-59 Deploying VMs from a Template

To deploy a VM, you must provide information such as the VM name, inventory location, host, datastore, and guest operating system customization data.

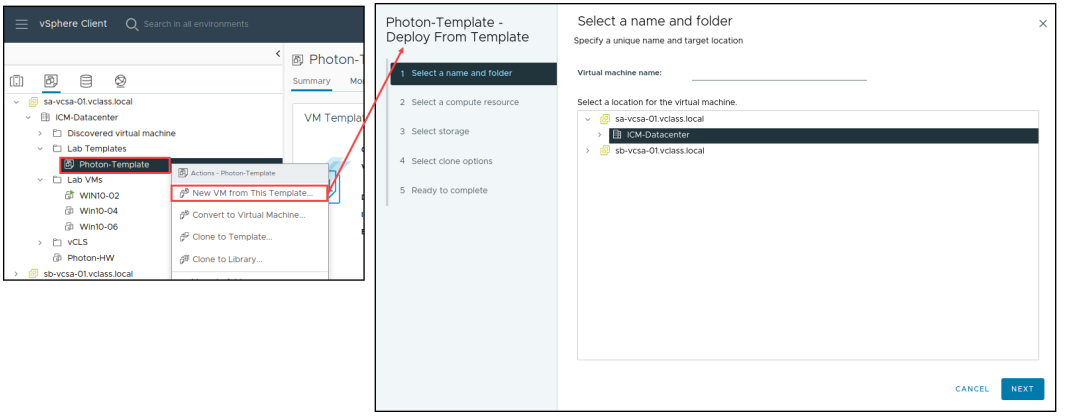

7-60 Cloning Virtual Machines

Cloning a VM creates a VM that is an exact copy of the original:

- Cloning is an alternative to deploying a VM from a template.
- The source VM can be powered on or off.

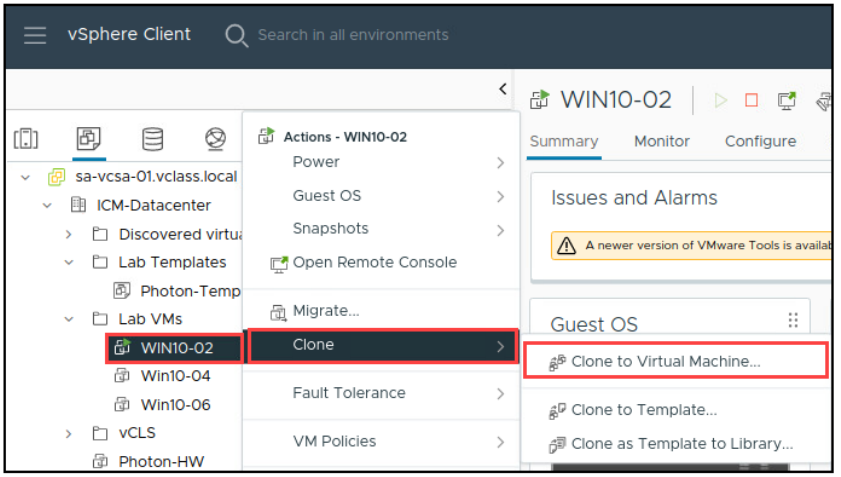

To clone a VM, you must be connected to vCenter. You cannot clone VMs if you use VMware Host Client to manage a host directly.

When you clone a VM that is powered on, services and applications are not automatically quiesced when the VM is cloned.

When deciding whether to clone a VM or deploy a VM from a template, consider the following points:

- When you deploy many VMs from a template, all the VMs start with the same base image. Cloning many VMs from a running VM might not create identical VMs, depending on the activity happening within the VM when the VM is cloned.
- VM templates use disk space, so you must plan your disk space requirements accordingly.
- Deploying a VM from a template is quicker than cloning a running VM, especially when you must deploy many VMs at a time.

7-61 Guest Operating System Customization

When you deploy a VM from a template or clone a VM, you can customize some aspects of the guest operating system.

By customizing a guest operating system, you can change information, including the following details:

- Computer name
- Network settings
- License settings
- Time zone
- Administrator or root password
- Windows Security Identifier

Customizing the guest operating system prevents conflicts that might occur when you deploy a VM and a clone with identical guest OS settings. Without customization, VMs retain the host name, IP address and so on, of the source VM or template.

7-62 About Customization Specifications

You can create a customization specification to prepare the guest operating system:

- Specifications are stored in the vCenter database.
- Windows and Linux guests are supported.

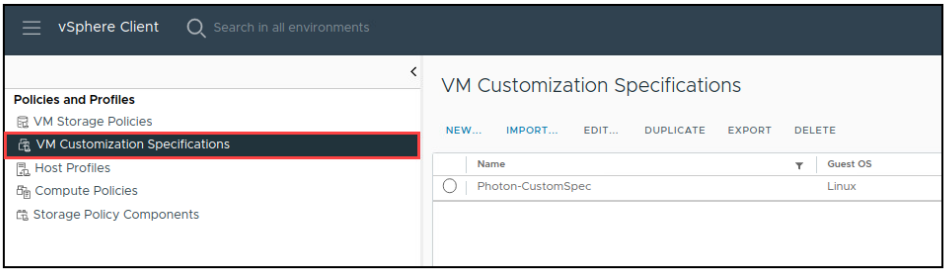

To manage customization specifications, select **Policies and Profiles** from the main menu.

On the VM Customization Specifications pane, you can create specifications or manage existing ones.

7-63 Customizing the Guest Operating System

When cloning a VM or deploying a VM from a template, you can use a customization specification to prepare the guest operating system.

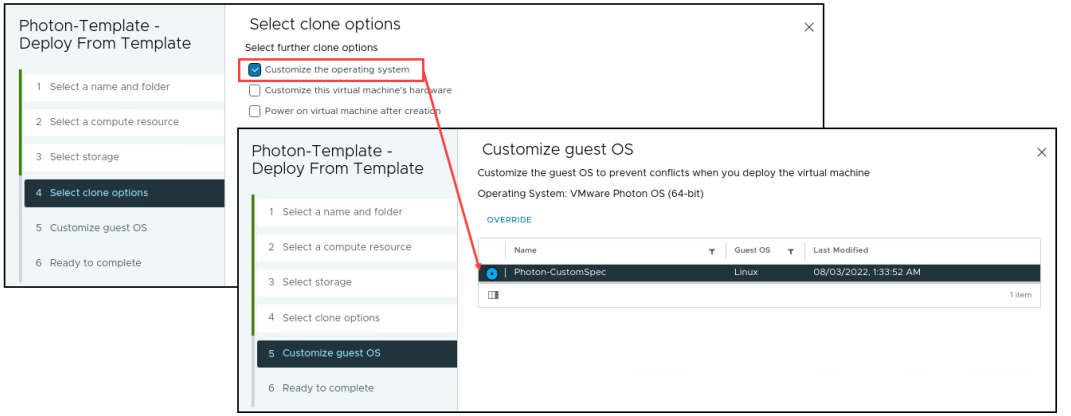

You can define the customization settings by using an existing customization specification during cloning or deployment. You create the specification ahead of time. During cloning or deployment, you can select the customization specification to apply to the new VM.

VMware Tools must be installed on the guest operating system that you want to customize.

The guest operating system must be installed on a disk attached to SCSI node 0:0 in the VM configuration.

For more about guest operating system customization, see *vSphere Virtual Machine Administration* at https://docs.vmware.com/en/VMware-vSphere/index.html.

7-64 Lab 16: Creating Templates and Deploying VMs

Create a VM template, create a customization specification, and deploy VMs from a template:

1. Create a Virtual Machine Template
2. Create Customization Specifications
3. Deploy Virtual Machines from a Template

7-65 Review of Learner Objectives

- Create a template of a virtual machine
- Deploy a virtual machine from a template
- Clone a virtual machine
- Create customization specifications for guest operating systems

7-66 Lesson 5: Introduction to Content Libraries

7-67 Learner Objectives

- Identify the benefits of a content library
- Recognize types of content libraries
- Create a local content library
- Deploy a virtual machine from a content library

7-68 About Content Libraries

Content libraries are repositories of OVF templates and other file types that can be shared and synchronized across vCenter systems globally.

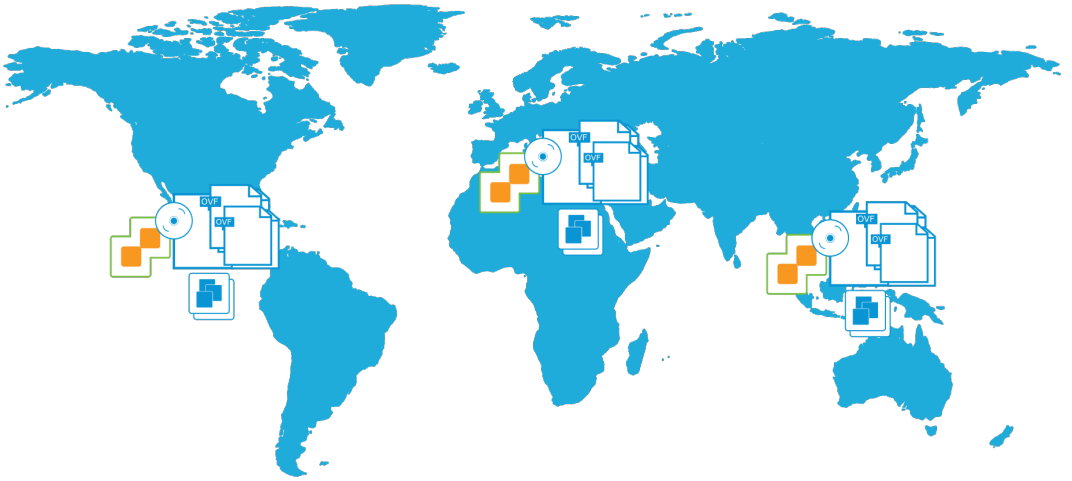

Organizations might have multiple vCenter instances in data centers around the globe. On these vCenter instances, organizations might have a collection of templates, ISO images, and so on. The challenge is that all these items are independent of one another, with different versions of these files and templates on various vCenter instances.

The content library is the solution to this challenge. It lets you store OVF templates, ISO images, or any other file types in a central location. The templates, images, and files can be published, and other content libraries can subscribe to and download content. The content library keeps content up to date by periodically synchronizing with the publisher, ensuring that the latest version is available.

7-69 Benefits of Content Libraries

Storage efficiency and consistency are key reasons to install and use a content library.

Using content libraries, administrators can perform the following functions:

- Store and share content, such as templates, ISO images, scripts.
- Perform distributed file management.
- Synchronize content libraries across sites and vCenter instances.
- Mount an ISO file directly from a content library.
- Maintain versions of VM templates.

Content libraries are stored on vSphere datastores.

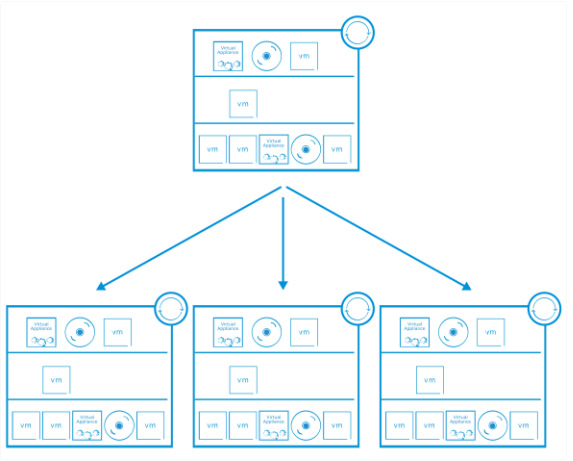

Sharing content and ensuring that the content is kept up to date are major tasks.

For example, for a main vCenter instance, you create a central content library to store the original copies of OVF templates, ISO images, scripts, and other file types. When you publish this content library, other libraries, which might be located anywhere in the world, can subscribe and download an exact copy of the data.

When an OVF template is added, modified, or deleted from the published catalog, the subscriber synchronizes with the publisher, and the libraries are updated with the latest content.

Starting with vSphere 7, you can update a template while simultaneously deploying VMs from the template. In addition, the content library keeps two copies of the VM template, the previous and current versions. You can roll back the template to reverse changes made to the template.

7-70 Content Library Types

Content is stored in one of the content library types:

- Local: Content that is controlled by the administrator
- Published: A local library that is available for subscription
- Subscribed: A library that synchronizes with a published library

Administrators can change content in a local or published content library.

Users cannot change content in a subscribed content library.

A subscribed content library cannot be published.

A local library is the simplest form of library. A local library is available for use in data centers that are objects to the local vCenter.

A published library is a local library that is available for subscription. Using the vCenter database, the content library tracks version changes. No option to use previous versions of content is available.

A subscribed library is configured to subscribe to a published library. An administrator cannot directly change the contents of the subscribed library. But the administrator can control how the data is synchronized to the subscribed content library.

7-71 Content Library Interface

To create and manage your content libraries, from the main menu, select **Content Libraries**.

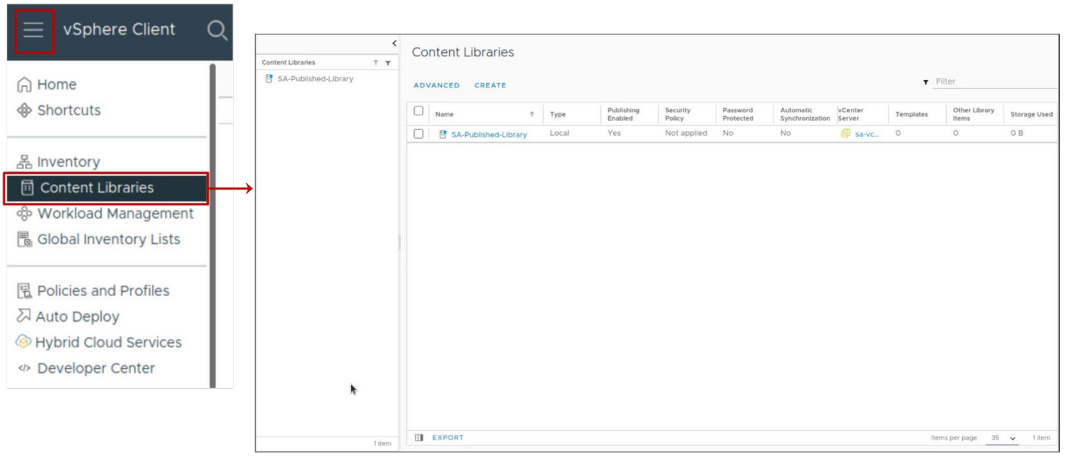

7-72 Creating a Local Content Library

When you create a content library, you select the content library type, for example, **Local content library**.

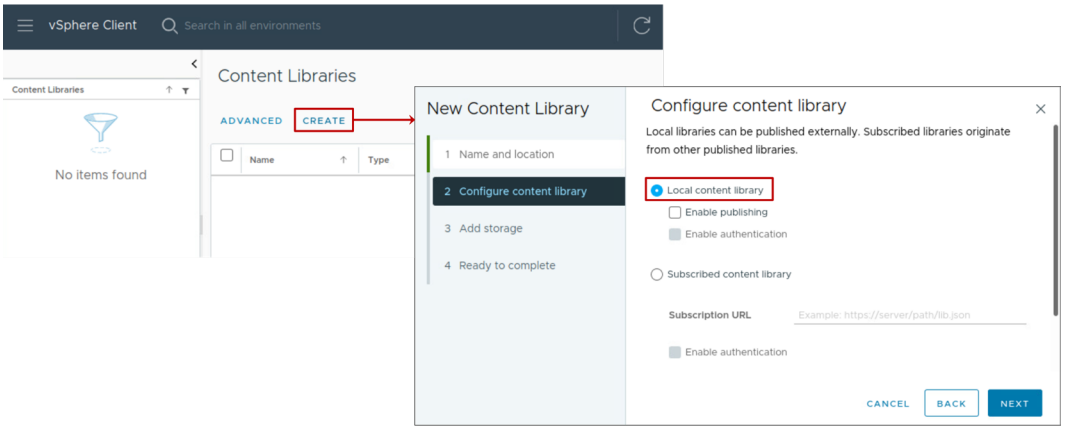

7-73 Populating the Content Library with Templates

You can populate the content library with the following template types:

VM Templates:

- Can be stored on any datastore type, except NFS
- Stored in the default disk format of the datastore (for example, thick-provisioned eager-zeroed)
- Are associated with a host
- Appear in the vCenter inventory

OVF Templates:

- Must be stored on the datastore (of any type) that is associated with the content library
- Stored in thin-provisioned format
- Are not associated with a host
- Do not appear in the vCenter inventory

When you create a VM template (instead of an OVF template) in a content library, the library item is backed by a VM template in the vCenter inventory.

The content library item and the corresponding inventory object are related in the following ways.

- If you convert the VM template in the vCenter inventory to a VM, the corresponding VM template library item is deleted.
- If you rename the VM template in vCenter inventory, the corresponding VM template library item is also renamed, and vice versa.
- If you delete the VM template in the vCenter inventory, the corresponding VM template library item is also deleted.

For the complete list of differences between VM templates and OVF templates in a content library, see vSphere Virtual Machine Administration at https://docs.vmware.com/en/VMware-vSphere/index.html.

7-74 Adding VM or OVF Templates to a Content Library

When you clone a virtual machine into a template in a content library, you can select whether to create a VM template or an OVF template.

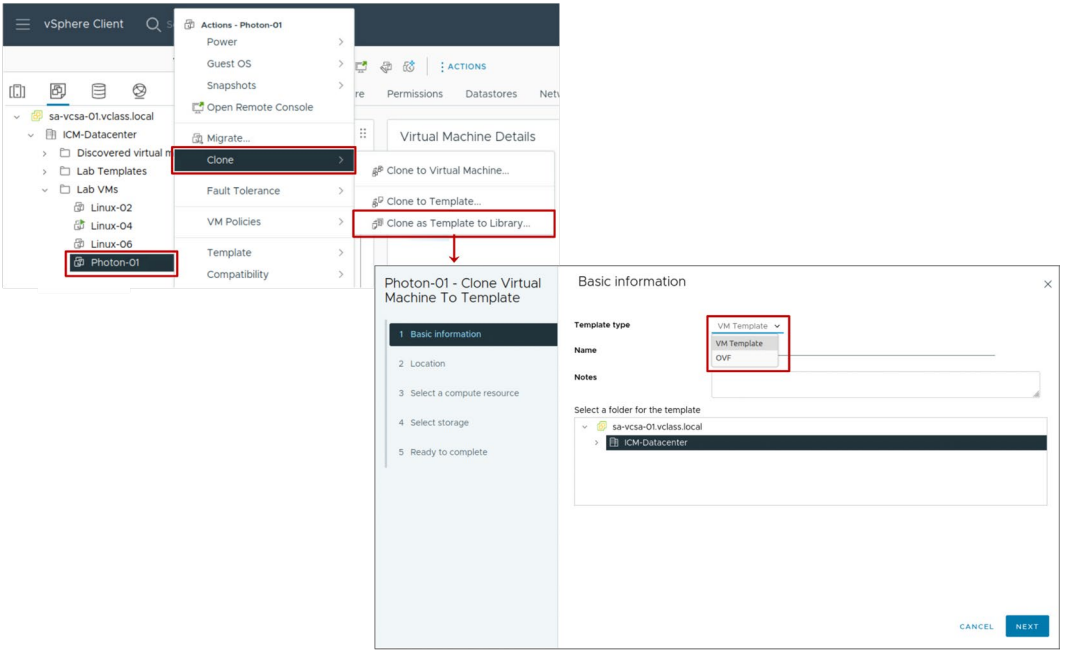

In this example, you clone the Photon-01 VM as a VM template to a content library.

7-75 Adding OVF Templates to a Content Library

When you clone a template from the vCenter inventory into a template in a content library, the template is stored as an OVF template in the library.

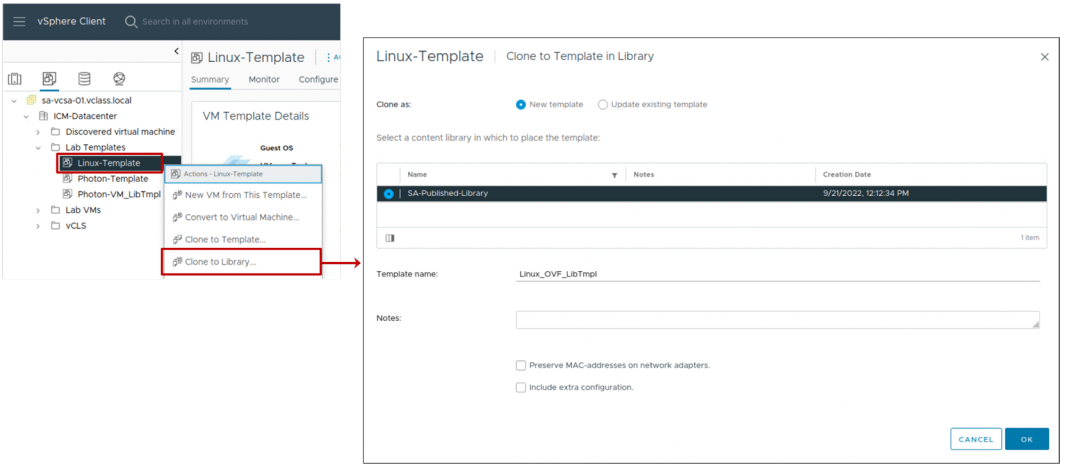

7-76 Viewing Content Library Items

The content in the content library is divided into categories:

- Templates:
  - VM templates
  - OVF templates
- Other Types:
  - All other file types, such as ISO images

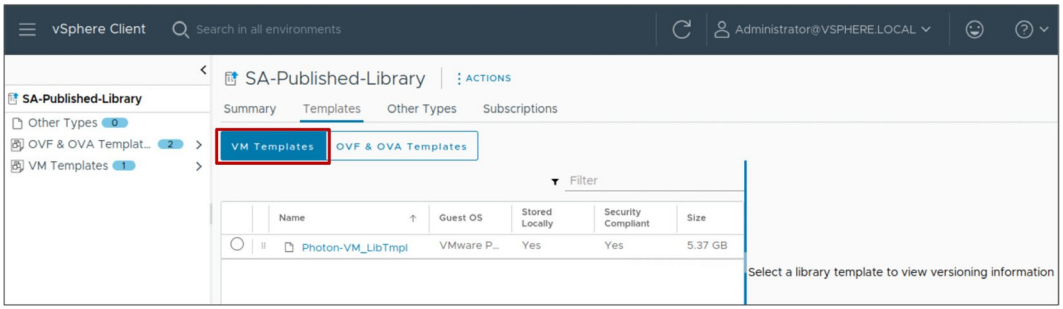

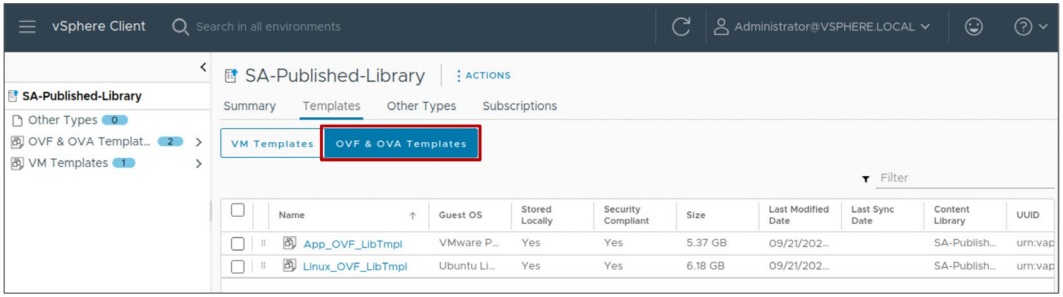

You can mount an ISO file directly from a content library. You connect the CD/DVD device to an ISO file that is stored in the content library, for example, to install a guest OS on a new VM.

The ISO files are available only to VMs that are registered on an ESXi host that can access the datastore where the content library is located. These ISO files are not available to VMs on hosts that cannot see the datastore on which the content library is located.

7-77 Deploying VMs from a Content Library

You can deploy VMs from templates in a content library by using the New Virtual Machine wizard.

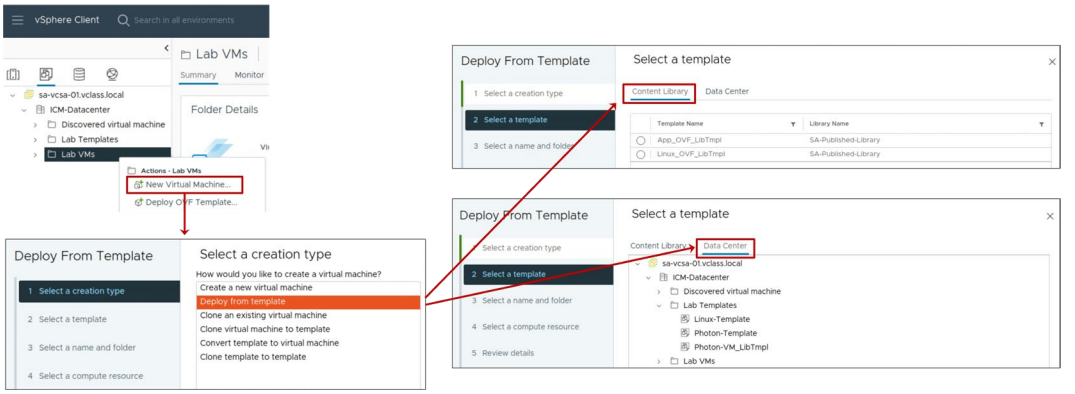

On the Select a template page, you can choose a template from a content library or from the vCenter inventory. The **Content Library** tab lists OVF templates and the **Data Center** tab lists VM templates, including VM templates that you added to the content library.

7-78 Lab 17: Using Local Content Libraries

Create a local content library to clone and deploy virtual machines:

1. Create a Local Content Library
2. Create an OVF Template in the Content Library
3. Create a VM Template in the Content Library
4. View the Content Library Templates
5. Deploy a VM from a Template in the Content Library

7-79 Review of Learner Objectives

- Identify the benefits of a content library
- Recognize types of content libraries
- Create a local content library
- Deploy a virtual machine from a content library

7-80 Lesson 6: Subscribing to Content Libraries

7-81 Learner Objectives

- Publish content libraries
- Subscribe to a published content library

7-82 Content Libraries: Local, Published, and Subscribed

You can publish a local content library so that other libraries can subscribe and download a copy of the data.

After synchronization, both published and subscribed libraries contain the same items, or the subscribed library contains the metadata for the items.

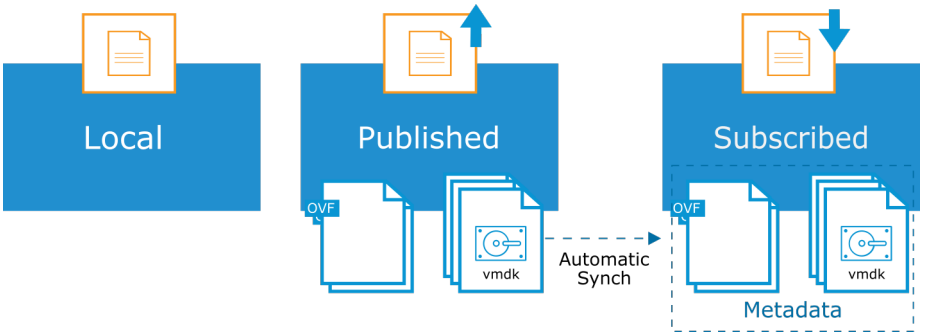

You can create a local library as the source for content that you want to save or share. You create the local library on a single vCenter instance. Furthermore, you can then add or remove items to and from the local library.

You can publish a local library, and this content library service endpoint can be accessed by other vCenter instances in your virtual environment, whether or not they are in the same enhanced linked mode group. When you publish a library, you can configure the authentication method, which a subscribed library must use to authenticate to it.

You can create a subscribed library and populate its content by synchronizing it to a published library. You can choose whether a subscribed library contains copies of the published library files or only the metadata of the library items.

The published library can be on the same vCenter instance as the subscribed library, or more likely, the subscribed library can reference a published library on a different vCenter instance.

You cannot add library items to a subscribed library. You can add items only to a local or published library.

After synchronization, both libraries contain the same items, or the subscribed library contains the metadata for the items.

You can publish this content library so that other libraries that are located across the world can subscribe and download an exact copy of the data. If an OVF template is added, modified, or deleted from the published library, the subscribed library synchronizes with the published library and the libraries are updated with the latest content.

7-83 Publishing a Content Library

You can enable publishing on a local content library by editing its settings.

You can add password protection to the library.

Subscribed libraries use the subscription URL to access the published library.

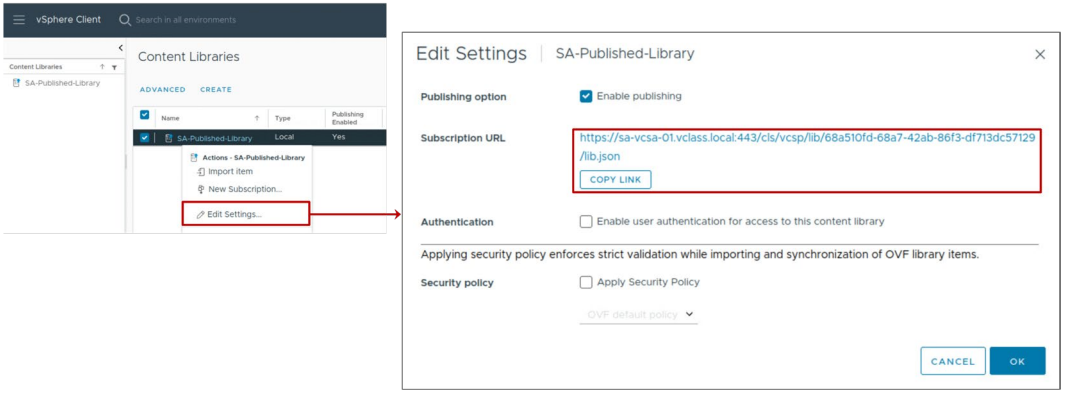

7-84 Subscribing to a Content Library

You subscribe a content library to a published content library by configuring the path to the subscription URL:

- You can immediately download all library content to the storage location that you configure.
- To save space, you can store only metadata for items until they are used.

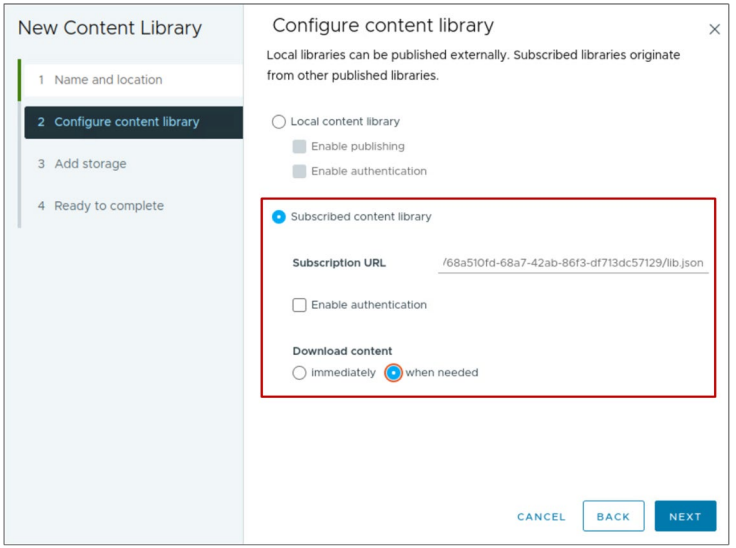

When a content library subscription is created, the administrator selects how the content synchronizes with the published content library. Content can be downloaded immediately if space for the content is not a concern.

Synchronization can be set to on-demand so that only the metadata is copied when the subscription is created. The full payload of the content library is downloaded the first time the content is used. So, space is saved because some templates and ISO images are not yet accessed and therefore, not downloaded.

7-85 Viewing Content Libraries

The Content Libraries pane shows all local and subscribed libraries and whether publishing is activated on a local library.

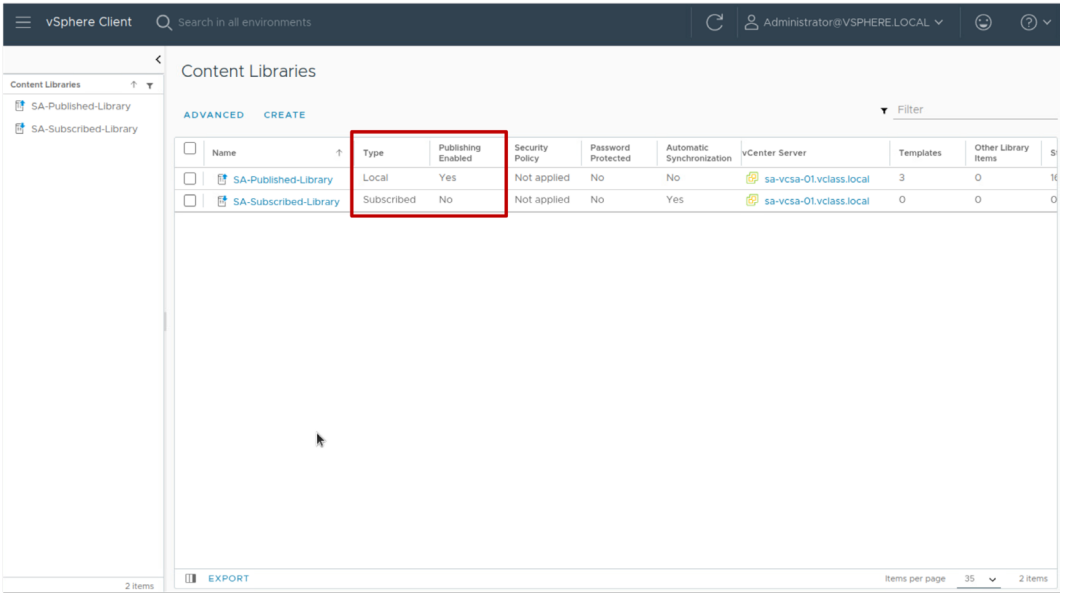

7-86 Viewing Subscribed Content Library Templates

The OVF & OVA Templates pane lists OVF templates. To update the list, you select **ACTIONS** > **Synchronize**.

The VM Templates pane lists VM templates. To update the list, you must create and publish subscription in the published library to push VM templates to the subscribed library.

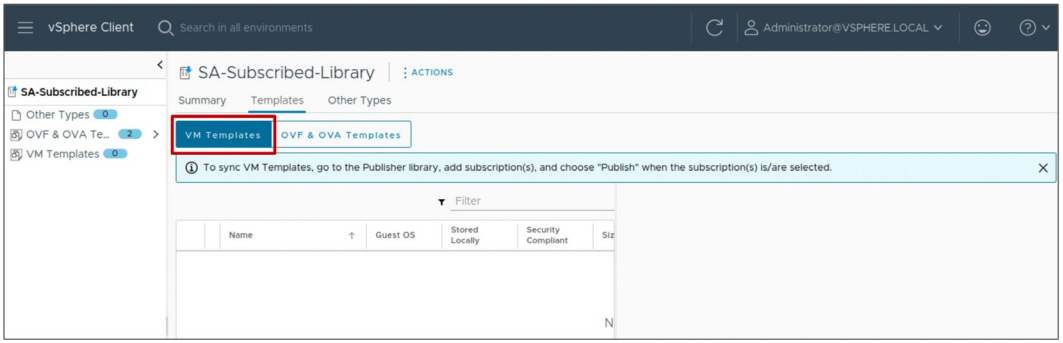

You can synchronize only OVF templates. If a subscriber initiates synchronization with a published library that contains both VM templates and OVF templates, only the OVF templates are synchronized in the subscribed library. VM templates are synchronized when a publisher library publishes them to its subscribers.

You can publish only VM templates. If you publish an entire library that contains both VM templates and OVF templates, only the VM templates are replicated to the subscriber. To synchronize OVF templates and other types of files, the subscriber must initiate the synchronization.

7-87 Creating a Subscription to Publish VM Templates

1. From the published library pane, select **ACTIONS** > **New Subscription** and create a subscription.
2. After the subscription is created, select the subscription in the list and click **PUBLISH**.

If you add new VM templates to the published library, you must publish them so that the subscribed library receives them.

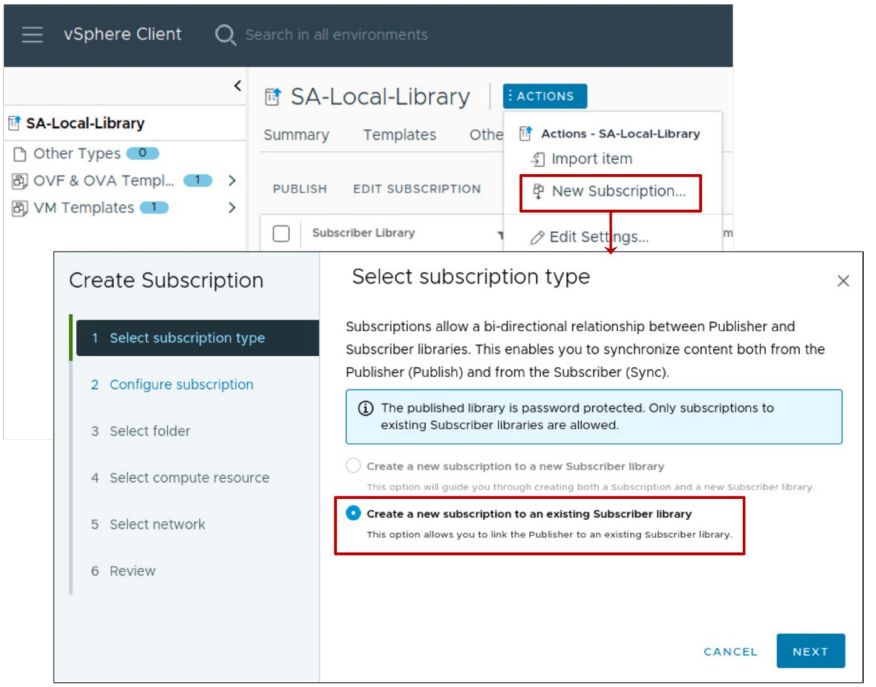

Subscriptions enable you to publish library items to a subscriber whenever you want. Create a subscription for a publisher library to control the distribution of templates to the subscriber.

7-88 Synchronizing Libraries With or Without Enhanced Linked Mode

Transfer speeds are optimized during the synchronization of published and subscribed content libraries in the same vCenter Single Sign-On domain.

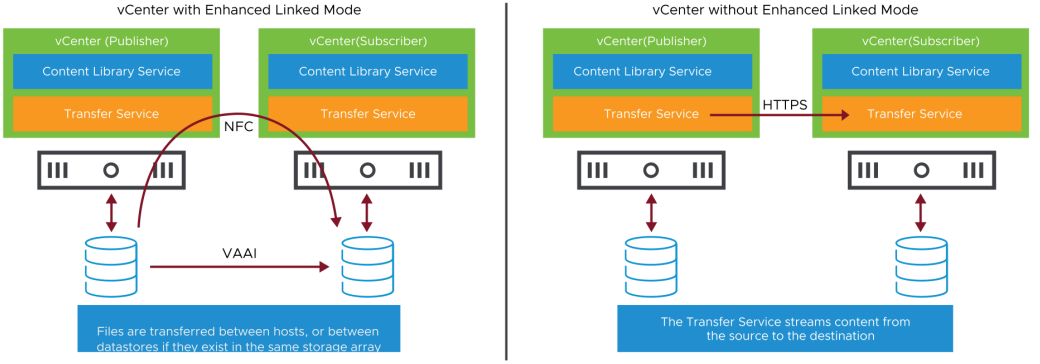

If Enhanced Linked Mode is configured across vCenter instances and the ESXi hosts can communicate with each other, then replication takes place between ESXi hosts using NFC. If the datastores used by the published and subscribed libraries exist in the same storage array, then VAAI is used for more efficient transfers.

When Enhanced Linked Mode is not available, the contents of a published library must be streamed over HTTPS through the Transfer Service component. In such a case, three scenarios are possible, based on the storage configuration:

- Both published and subscribed libraries reside on a datastore.
- Both published and subscribed libraries reside on an NFS file system mounted on the vCenter instances.
- The published library resides on an NFS file system, whereas the subscribed library resides on a datastore.

When any item is changed in a content library, the item number is incremented to note the modification. The version number for the content library as a whole is also incremented. This process is called simple versioning. The Content Library Service determines whether synchronization should occur by using simple versioning. Simple versioning does not store multiple versions or offer rollback features, it is purely a numerical value that is assigned to the content library and its content.

The content library's version number is checked first.

- If this number matches the one in the published and subscribed libraries, no synchronization is required.
- If a discrepancy is found, the version number of each item in the library is checked.
- If discrepancies are found, only items with such discrepancies are synchronized.

7-89 Advanced Configuration

You can use the Advanced Configuration page to control and optimize how a published library stores and synchronizes content.

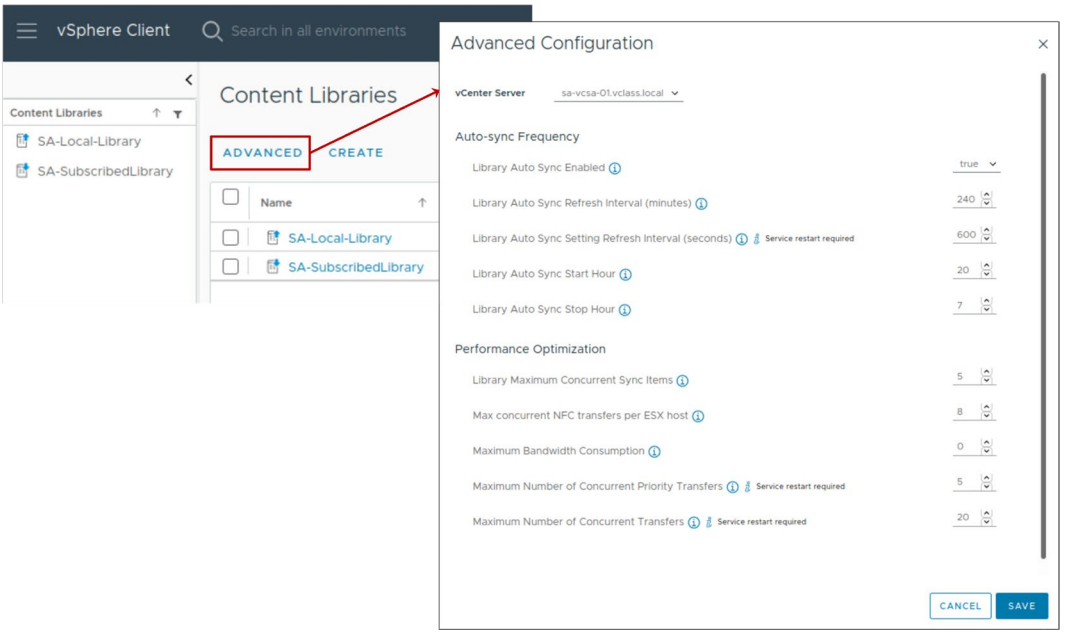

7-90 Content Library Maximums

The content library has the following configuration maximums.

Configuration Item	Maximum
Content library item size	1 TB
Total items per library	1,000
Total library items per vCenter instance (across all libraries)	2,000
Maximum number of concurrent sync operations on the published library's vCenter instance	16
Total number of libraries per vCenter instance	1,000

The content library can be supported by a datastore or stored to available storage on vCenter.

Regardless of the option selected, the content library can be supported only by a single file system or datastore.

The maximum size of a library item is 1 TB. A content library can hold a maximum of 1,000 items and a total of 2,000 items across all libraries in a vCenter instance.

The maximum number of concurrent synchronization operations on the published library's vCenter instance is 16.

Automatic synchronization occurs once every 240 minutes by default, but the time and frequency can be configured through the content library advanced configuration settings. The administrator can synchronize an entire content library or an individual item at any time through the vSphere Client.

7-91 Lab 18: Using Subscribed Content Libraries

Publish a local content library and create a second library that subscribes to it:

1. Publish a Local Content Library
2. Create a Subscribed Content Library
3. Create a Subscription for VM Templates
4. Deploy a VM from the Subscribed Content Library

7-92 Review of Learner Objectives

- Publish content libraries
- Subscribe to a published content library

7-93 Lesson 7: Managing Templates in a Content Library

7-94 Learner Objectives

- Manage multiple versions of VM templates in content libraries

7-95 Benefits of Using a Content Library to Manage VM Templates

When managed by a content library, you can work with VM templates in the following ways:

- Make changes to VM templates by checking them out.
- Deploy VMs while simultaneously patching the VM template.
- Get a history of changes made to the VM template.
- Access two copies of the VM template, the previous and current versions.
- Revert to a previous version of the VM template.

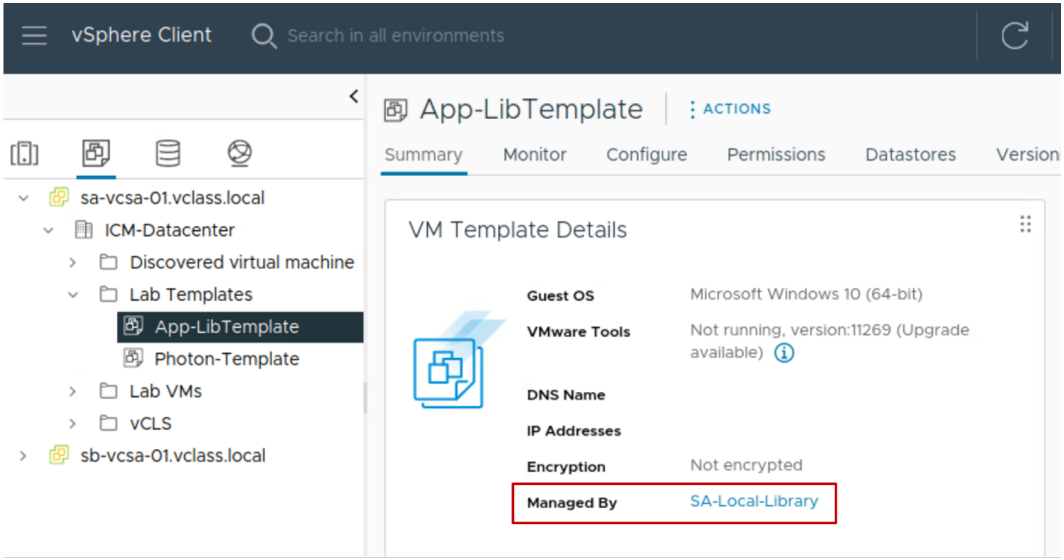

App-LibTemplate is managed by content library SA-Local-Library.

Managing VM templates with the content library has several benefits:

- You can monitor the changes that have been made by privileged users.
- You can patch or modify VMs while the template is being used to deploy VMs.
- You can view the history of changes made to a VM.

7-96 Overview of the Template Versioning Process

To create a new version of a template, you perform the following steps:

1. Check out a VM from the template.
2. Make changes to the VM.
3. Check in the VM to the template.

7-97 Checking Out a VM from the Template

To update a VM template, you use the **Versioning** tab to check out a VM from the template.

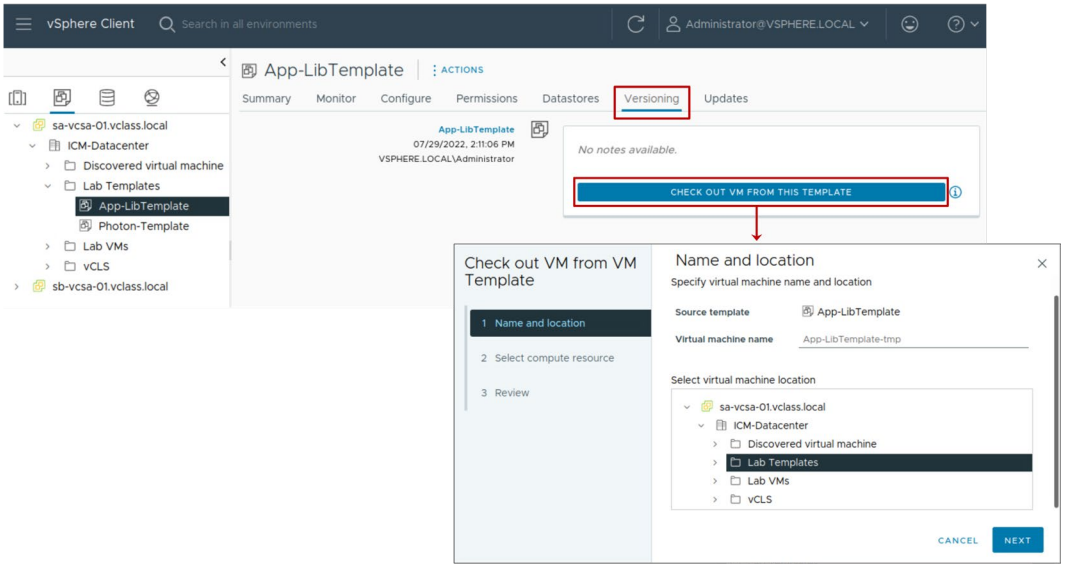

You can perform the checkout operation to update a virtual machine from the VM template. During this process, other users cannot check out the VM template, but they can deploy a virtual machine from the VM template without any disruptions.

When you check out a VM template, you cannot convert the virtual machine to a template or migrate the virtual machine to a different vCenter inventory.

7-98 Making Changes to the VM

After checking out the template to a VM, you can make hardware or software changes to the VM.

You can change the VM while the VM template continues to be available for VM deployments.

If you do not want to keep the changes you made to the VM, you can discard the checked out VM.

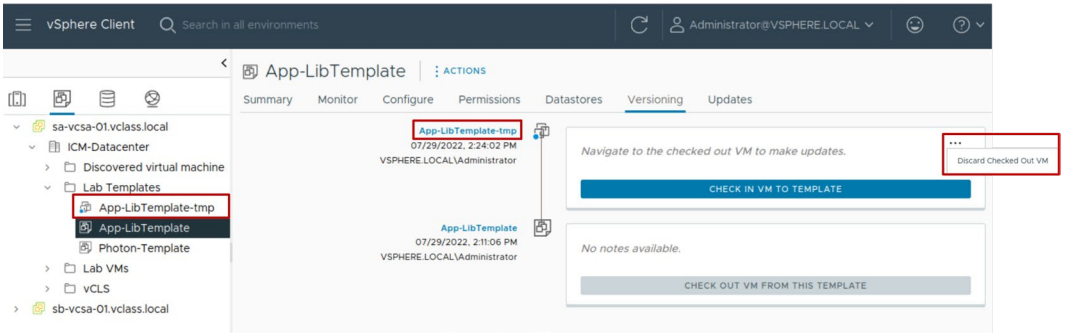

vSphere uses linked clone technology to clone a VM template on checkout. This VM clone is used for changes and updates, whereas the original VM template remains available in the content library to deploy VMs. When the changes to the cloned VM are complete, the VM clone is re-merged with the VM template, the cloned VM is destroyed, and the VM template is updated.

You can discard the checked out virtual machine to avoid creating new versions or to prevent other users from using a faulty version. To discard the checked out VM, click the ellipsis next to the checked out VM and click **Discard Checked Out VM**.

7-99 Checking In the VM to the Template

After making changes to the VM, you check in the VM to the template.

The **Check in notes** box is required.

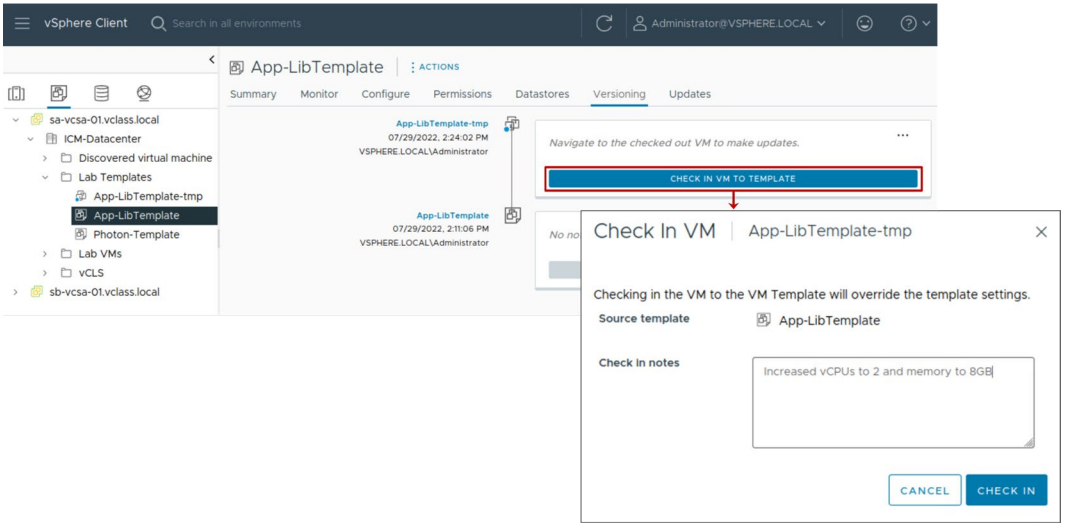

When you check in the VM to a template, you create a new version of the VM template containing the updated state of the virtual machine.

7-100 Viewing Template Versions

The version information appears in the **Versioning** tab.

Each time you check in the VM back to the template, you create a new version of the VM template.

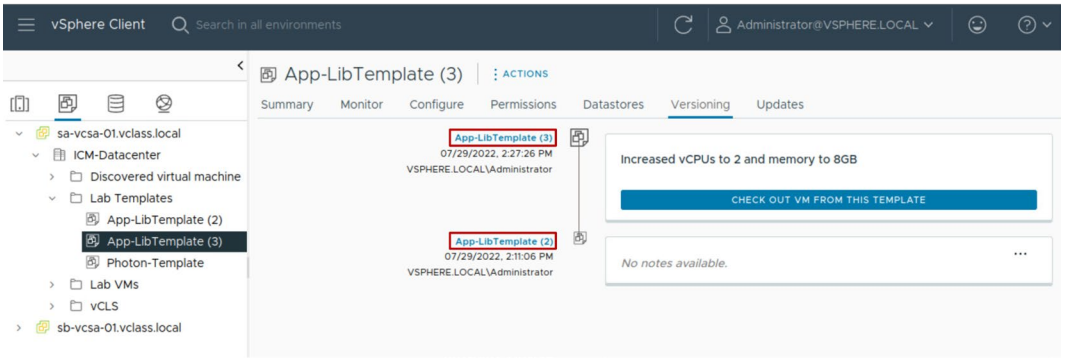

The versioning information highlights the following details:

- Who made the updates
- When the updates were made
- Notes or information added during the updates

The quality of the notes on the **Versioning** tab can vary, depending on who enters them. Whatever the quality, the notes are captured and retained in the vCenter database until the VM template is deleted.

7-101 Deleting and Reverting to Template Versions

Delete a previous version of a VM template if you no longer want to allow the use of the template.

Revert to a VM template version if:

- The latest VM template contains changes that you no longer need.
- You made a mistake during your last edit.

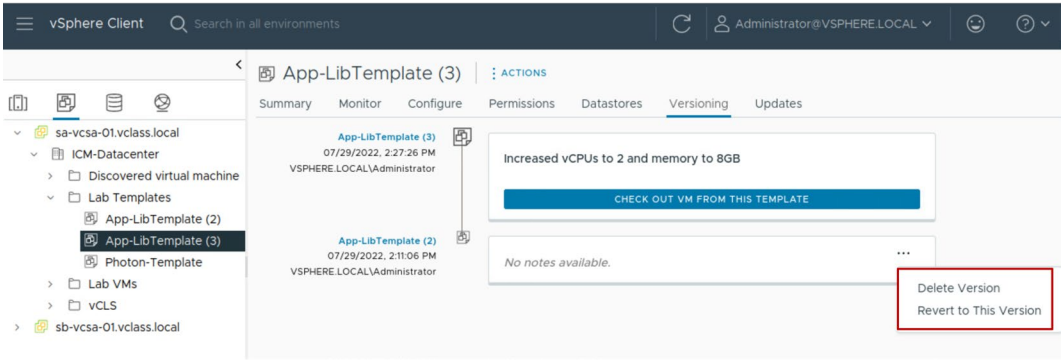

Deleting a VM template removes the template and its content from the inventory.

7-102 Lab 19: Versioning Templates in the Content Library

Manage multiple versions of a VM template versioning in the local content library:

1. Check Out a VM Template in the Content Library
2. Make Changes to the VM Template
3. Check In the VM Template to the Content Library
4. Revert to a Previous Version of the VM Template

7-103 Review of Learner Objectives

- Manage multiple versions of VM templates in content libraries

7-104 Key Points

- vCenter provides features for provisioning virtual machines, such as templates, cloning, and content libraries.
- By deploying VMs from a template, you can create many VMs easily and quickly.
- You can dynamically manage a VM's configuration by adding hot-pluggable devices and increasing the size of a VM's virtual disk.
- You can publish a local content library so that other libraries can subscribe and download an exact copy of the data.
- You can update a VM template managed by content library while the same VM template is being used to deploy VMs.

Questions?

Module 8

Managing Virtual Machines

8-2 Importance

Managing VMs effectively requires skills in migrating VMs, taking snapshots, and managing the resources of the VMs.

8-3 Module Lessons

1. Migrating VMs with vSphere vMotion
2. Configuring Enhanced vMotion Compatibility
3. Migrating VMs with vSphere Storage vMotion
4. Cross vCenter Migrations
5. Creating Virtual Machine Snapshots
6. Virtual CPU and Memory Concepts
7. Resource Controls

8-4 Lesson 1: Migrating VMs with vSphere vMotion

8-5 Learner Objectives

- Recognize the types of VM migrations that you can perform within a vCenter instance
- Explain how vSphere vMotion works
- Verify vSphere vMotion requirements
- Migrate virtual machines using vSphere vMotion

8-6 About VM Migration

Migration means moving a VM from one host, datastore, or vCenter instance to another host, datastore, or vCenter instance.

Migration can be cold or hot:

- A cold migration moves a powered-off or suspended VM.
- A hot migration moves a powered-on VM.

vCenter performs compatibility checks before migrating suspended or powered-on VMs to ensure that the VM is compatible with the target host.

8-7 Migration Types

The type of migration that you perform depends on the power state of the VM that you select in the inventory and the migration type that you select in the Migrate wizard.

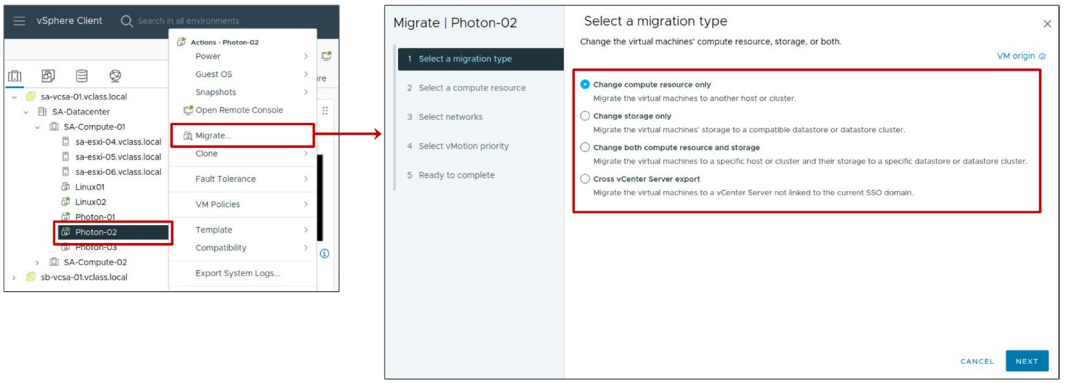

The Migrate wizard provides the following migration options:

- Compute resource only:
  - Move a VM, but not its storage, to another host.
  - For a hot migration, vSphere vMotion is used to move the VM.
- Storage only:
  - Move a VM's files or objects to a new datastore.
  - For a hot migration, vSphere Storage vMotion is used to move the VM.
- Both compute resource and storage:
  - Move a VM to another host and datastore.
  - For a hot migration, vSphere vMotion and vSphere Storage vMotion are used to move the VM.
- Cross vCenter Server export:
  - Move the VM to a host and datastore managed by a different vCenter instance that is not linked to the current SSO domain.

The purpose of the migration determines which migration technique to use. For example, you might need to shut down a host for maintenance but keep the VMs running. Use vSphere vMotion to migrate the VMs instead of performing a cold or suspended VM migration. If you must move a VM's files to another datastore to better balance the disk load or transition to another storage array, you use vSphere Storage vMotion.

Some migration techniques, such as vSphere vMotion migration, have special hardware requirements that must be met. Other techniques, such as a cold migration, do not have special hardware requirements.

8-8 About vSphere vMotion

A vSphere vMotion migration moves a powered-on VM from one host (compute resource) to another.

vSphere vMotion provides the following capabilities:

- Improvement in overall hardware use
- Continuous VM operation while accommodating scheduled ESXi host downtime
- vSphere DRS to balance VMs across hosts

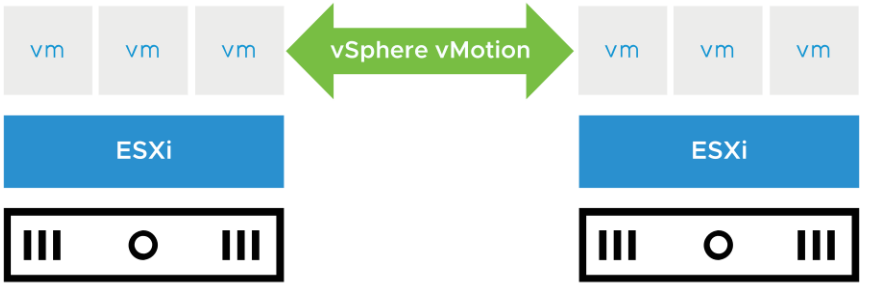

Using vSphere vMotion, you can migrate running VMs from one ESXi host to another ESXi host with no disruption or downtime. vSphere DRS uses vSphere vMotion to migrate running VMs from one host to another to ensure that the VMs have the resources that they require.

With vSphere vMotion, the entire state of the VM is moved from one host to another, but the data storage remains in the same datastore.

The state information includes the current memory content and all the information that defines and identifies the VM. The memory content includes transaction data and whatever bits of the operating system and applications are in memory. The definition and identification information stored in the state includes all the data that maps to the VM hardware elements, such as the BIOS or EFI, devices, CPU, and MAC addresses for the Ethernet cards.

8-9 Configuring vSphere vMotion Networks

vSphere vMotion migrations require correctly configured VMkernel adapters on the source and destination hosts.

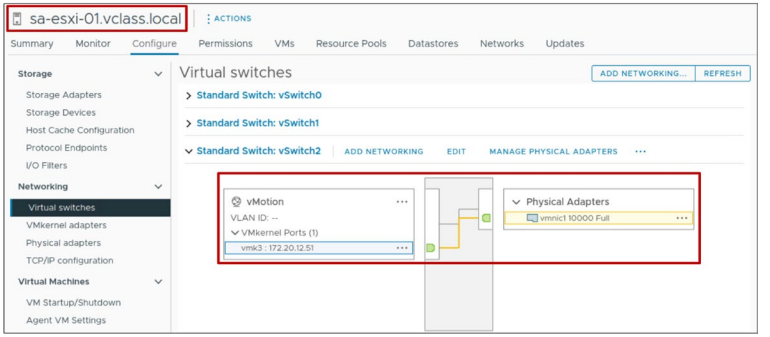

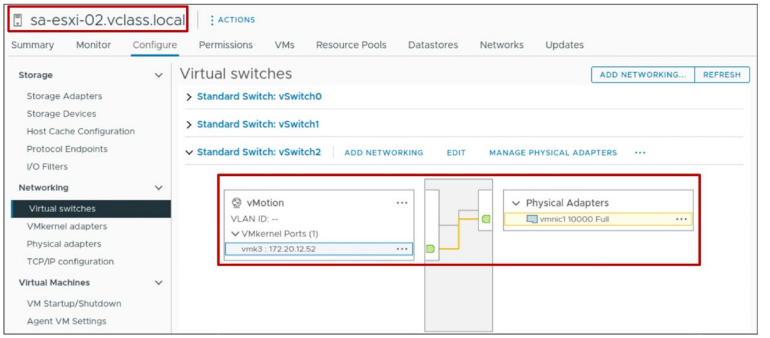

To configure vSphere vMotion networking, you must set up a VMkernel adapter with the vSphere vMotion service activated on the source and destination host.

8-10 vSphere vMotion Migration Workflow

The source host (ESXi01) and the destination host (ESXi02) can access the shared datastore that holds the VM's files.

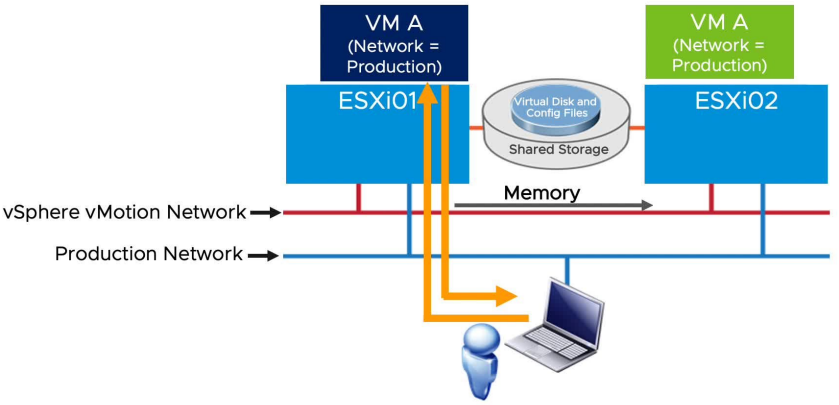

To play the animation, go to https://vmware.bravais.com/s/FbzaDb6owpSMKyKc940F.

A vSphere vMotion migration consists of the following steps:

1. A shadow VM is created on the destination host.
2. The VM's memory state is copied over the vSphere vMotion network from the source host to the target host through the vSphere vMotion network. Users continue to access the VM and, potentially, update pages in memory. A list of modified pages in memory is kept in a memory bitmap on the source host.
3. After the first pass of memory state copy completes, another pass of memory copy is performed to copy any pages that changed during the last iteration. This iterative memory copying continues until the number of changed pages is small enough to copy across within 500 milliseconds.
4. After most of the VM's memory is copied from the source host to the target host, the VM is quiesced. No additional activity occurs on the VM. In the quiesce period, vSphere vMotion transfers the VM device state and memory bitmap to the destination host.
5. Immediately after the VM is quiesced on the source host, the VM is initialized and starts running on the target host. A Gratuitous Address Resolution Protocol (GARP) request notifies the subnet that VM A's MAC address is now accessible through a different physical switch port. The VM's files are unlocked by the source host and locked by the destination host.

6. Users access the VM on the target host instead of the source host.
7. The memory pages that the VM was using on the source host are marked as free.

For a detailed discussion of how vSphere vMotion works, see *The vMotion Process Under the Hood* at https://blogs.vmware.com/vsphere/2019/07/the-vmotion-process-under-the-hood.html.

8-11 VM Requirements for vSphere vMotion Migration

For migration with vSphere vMotion, a VM must meet these requirements:

- If it uses an RDM disk, the RDM file and the LUN to which it maps must be accessible by the destination host.
- It must not have a connection to a virtual device, such as a CD/DVD or floppy drive, with a host-local image mounted.

You can use vSphere vMotion to migrate a VM with a device that is attached through a remote console (such as a physical device or disk image).

For the complete list of vSphere vMotion migration requirements, see *vCenter Server and Host Management* at https://docs.vmware.com/en/VMware-vSphere/index.html.

8-12 Host Requirements for vSphere vMotion Migration (1)

Source and destination hosts must have the following characteristics:

- Accessibility to all the VM's storage:
  - 128 concurrent migrations are possible per datastore.
  - If the swap file location on the destination host differs from the swap file location on the source host, the swap file is copied to the new location.
- VMkernel port with vSphere vMotion activated
- Matching management network IP address families (IPv4 or IPv6) between the source and destination hosts

You cannot migrate a VM from a host that is registered to vCenter with an IPv4 address to a host that is registered with an IPv6 address.

Copying a swap file to a new location can result in slower migrations. If the destination host cannot access the specified swap file location, it stores the swap file with the VM configuration file.

8-13 Host Requirements for vSphere vMotion Migration (2)

- The number of concurrent, active vSphere vMotion tasks is a function of the vSphere vMotion network speed:
  - Each active vSphere vMotion process requires a minimum throughput of 250 Mbit/second (Mbps) on the vSphere vMotion network.
  - Concurrent migrations are limited to four on a 1 Gbps network.
  - Concurrent migrations are limited to eight on a 10 Gbps (or faster) network.
  - For better performance, dedicate at least two VMkernel port groups to the vSphere vMotion traffic.
- Compatible CPUs:
  - The CPU feature sets of both the source host and the destination host must be compatible.
  - Some features can be hidden by using Enhanced vMotion Compatibility or compatibility masks.

For more details on host requirements for vSphere vMotion, see *vCenter Server and Host Management* at https://docs.vmware.com/en/VMware-vSphere/index.html.

8-14 Performing a vSphere vMotion Migration

vSphere vMotion is used to move a powered-on VM to a different host.

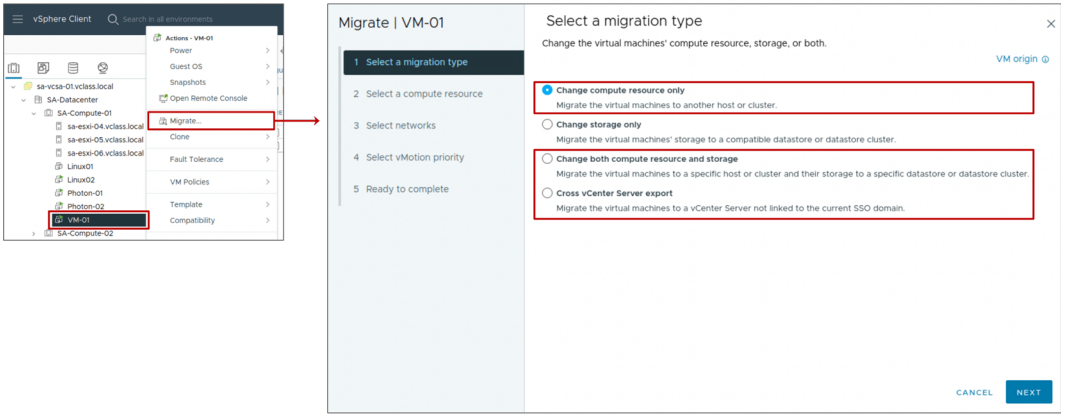

Whenever you migrate a powered-on VM, vSphere vMotion performs the work of moving the VM from one host to another. The host is also known as the compute resource.

The Migrate wizard has a few options for moving a VM from one compute resource to another.

- Change compute resource
- Change both compute resource and storage
- Cross vCenter Server export

In all of these cases, if the VM being migrated is powered on, then vSphere vMotion performs the migration.

8-15 Checking Migration Errors

When you select the host or cluster, validation checks are performed to verify that most vSphere vMotion requirements are met.

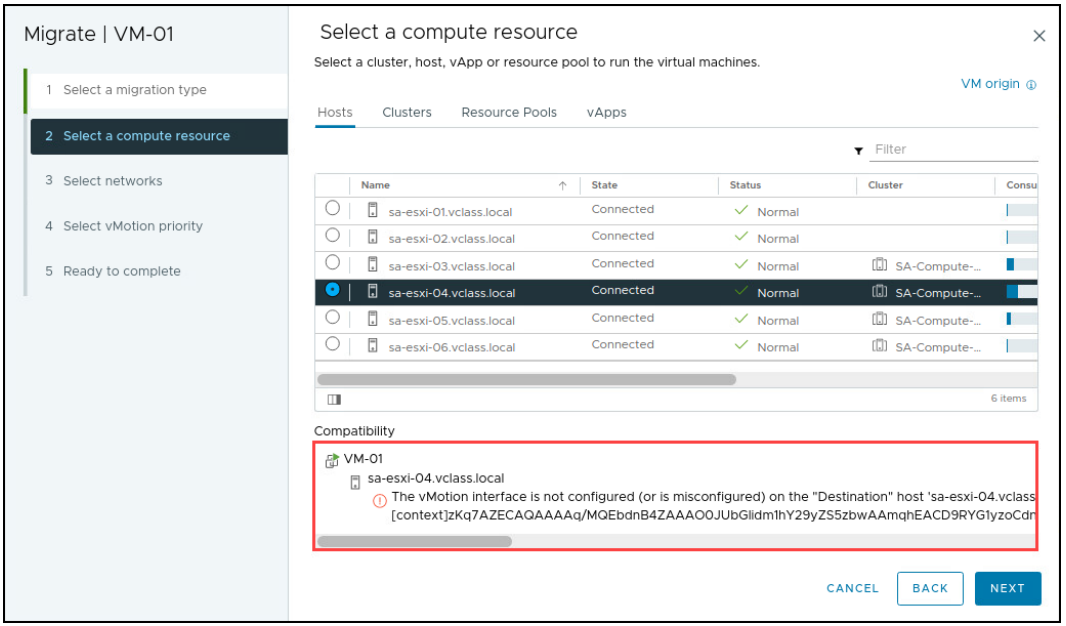

If validation succeeds, you can continue in the wizard. If validation does not succeed, a list of vSphere vMotion errors and warnings displays in the Compatibility pane.

With warnings, you can still perform a vSphere vMotion migration. But with errors, you cannot continue. You must exit the wizard and fix all errors before retrying the migration.

If a failure occurs during the vSphere vMotion migration, the VM is not migrated and continues to run on the source host.

8-16 Migrating Encrypted VMs

When powered-on, encrypted VMs are migrated, encrypted vSphere vMotion is automatically used.

For VMs that are not encrypted, select one of the following encrypted vSphere vMotion menu items:

- **Disabled**
- **Opportunistic** (default): Encrypted vSphere vMotion is used if the source and destination hosts support it
- **Required**: If the source or destination host does not support encrypted vSphere vMotion, the migration fails

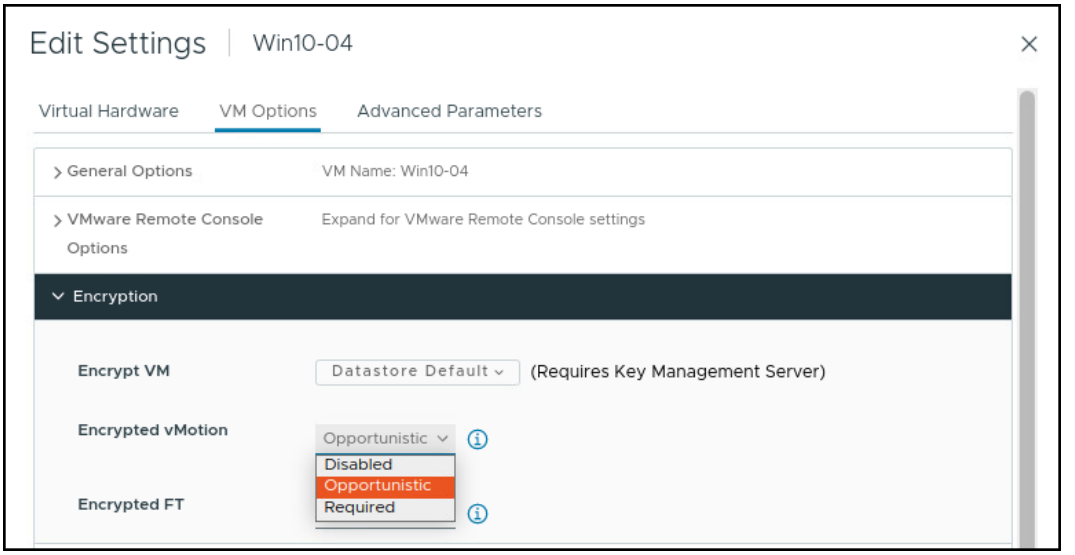

Encrypted vSphere vMotion secures confidentiality, integrity, and authenticity of data that is transferred using vSphere vMotion. Encrypted vSphere vMotion supports all variants of vSphere vMotion, including migration across vCenter systems. Encrypted vSphere Storage vMotion is not supported.

You cannot turn off encrypted vSphere vMotion for encrypted VMs.

8-17 Lab 20: vSphere vMotion Migrations

Configure vSphere vMotion networking and migrate virtual machines using vSphere vMotion:

1. Configure vSphere vMotion Networking on sa-esxi-01.vclass.local
2. Configure vSphere vMotion Networking on sa-esxi-02.vclass.local
3. Prepare Virtual Machines for vSphere vMotion Migration
4. Migrate Virtual Machines Using vSphere vMotion

8-18 Review of Learner Objectives

- Recognize the types of VM migrations that you can perform within a vCenter instance
- Explain how vSphere vMotion works
- Verify vSphere vMotion requirements
- Migrate virtual machines using vSphere vMotion

8-19 Lesson 2: Configuring Enhanced vMotion Compatibility

8-20 Learner Objectives

- Describe the role of Enhanced vMotion Compatibility in migrations
- Configure EVC CPU mode on a vSphere cluster
- Explain how per-VM EVC CPU mode works with vSphere vMotion
- Configure EVC Graphics mode on a vSphere cluster or a VM

8-21 CPU Constraints on vSphere vMotion Migration

CPU compatibility between source and target hosts is a vSphere vMotion requirement that must be met.

CPU Characteristics	Exact Match Required By Source Host and Target Host	Reason
Clock speeds, cache sizes, hyperthreading, and number of cores	No	The VMkernel virtualizes these characteristics.
Manufacturer (Intel or AMD) family and generation (Opteron4, Intel Westmere)	Yes	Instruction sets contain many small differences.
Presence or absence of SSE3, SSSE3, or SSE4.1 instructions	Yes	Multimedia instructions are usable directly by applications.
Virtualization hardware assist	For 32-bit VMs: No	The VMkernel virtualizes this characteristic.
	For 64-bit VMs on Intel: Yes	Intel 64-bit with VMware implementation uses Intel VT.

Depending on the CPU characteristic, an exact match between the source and target host might or might not be required. For example, if hyperthreading is activated on the source host and deactivated on the destination host, the vSphere vMotion migration continues because the VMkernel handles this difference in characteristics.

But, if the source host processor supports SSE4.1 instructions and the destination host processor does not support them, the hosts are considered incompatible and the vSphere vMotion migration fails. SSE4.1 instructions are application-level instructions that bypass the virtualization layer and might cause application instability if mismatched after a migration with vSphere vMotion.

8-22 About Enhanced vMotion Compatibility

Enhanced vMotion Compatibility is a cluster feature that enables vSphere vMotion migrations between hosts without identical feature sets.

The feature uses CPU baselines to configure all the processors in the cluster that are activated for Enhanced vMotion Compatibility.

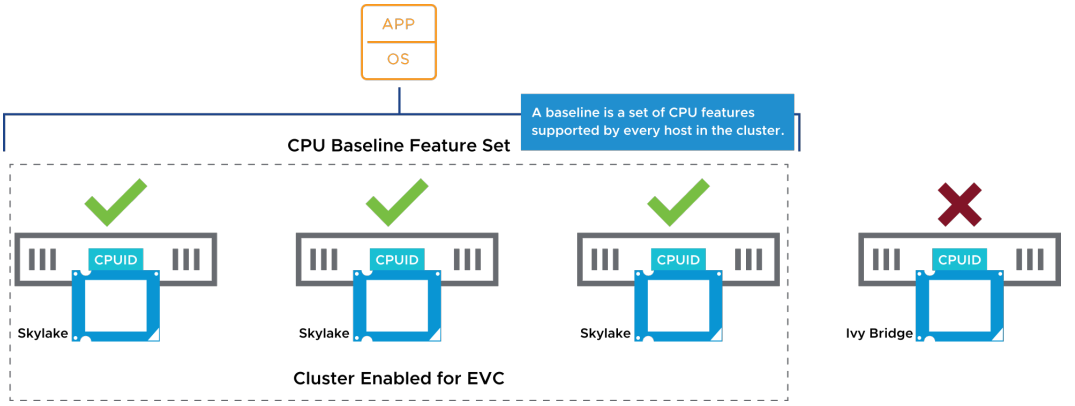

Enhanced vMotion Compatibility verifies that all hosts in a cluster present the same CPU feature set to VMs, even if the CPUs on the hosts differ.

Enhanced vMotion Compatibility facilitates safe vSphere vMotion migration across a range of CPU generations. With Enhanced vMotion Compatibility, you can use vSphere vMotion to migrate VMs among CPUs that otherwise are considered incompatible.

With Enhanced vMotion Compatibility, vCenter can enforce vSphere vMotion compatibility among all hosts in a cluster by forcing hosts to expose a common set of CPU features (baseline) to VMs. A baseline is a set of CPU features that are supported by every host in the cluster. When you configure Enhanced vMotion Compatibility, you set all host processors in the cluster to present the features of a baseline processor. After the features are activated for a cluster, hosts that are added to the cluster are automatically configured to the CPU baseline.

Hosts that cannot be configured to the baseline are not permitted to join the cluster. VMs in the cluster always see an identical CPU feature set, no matter on which host they run. Because the process is automatic, Enhanced vMotion Compatibility is easy to use and requires no specialized knowledge of CPU features and masks.

8-23 EVC Cluster Requirements for CPU Mode

All hosts in the cluster must meet several CPU-based requirements:

- Use CPUs from a single vendor, either Intel or AMD.
- Be activated for hardware virtualization: AMD-V or Intel VT.
- Be activated for execution-disable technology: AMD No eXecute (NX) or Intel eXecute Disable (XD).
- Be configured for vSphere vMotion migration.

Applications in VMs must be CPU ID compatible.

Before you create an Enhanced vMotion Compatibility cluster, ensure that the hosts that you intend to add to the cluster meet the requirements.

Enhanced vMotion Compatibility automatically configures hosts whose CPUs have Intel FlexMigration and AMD-V Extended Migration technologies to be compatible with vSphere vMotion with hosts that use older CPUs.

For Enhanced vMotion Compatibility to function properly, the applications on the VMs must be written to use the CPU ID machine instruction for discovering CPU features, as recommended by the CPU vendors. vSphere cannot support Enhanced vMotion Compatibility with applications that do not follow the CPU vendor recommendations for discovering CPU features.

To determine which EVC modes are compatible with your CPU, search the VMware Compatibility Guide at http://www.vmware.com/resources/compatibility. Search for the server model or CPU family, and click the entry in the CPU Series column to display the compatible EVC modes.

8-24 Configuring EVC CPU Mode on an Existing Cluster

You configure EVC CPU mode on an existing cluster to ensure vSphere vMotion CPU compatibility between the hosts in the cluster.

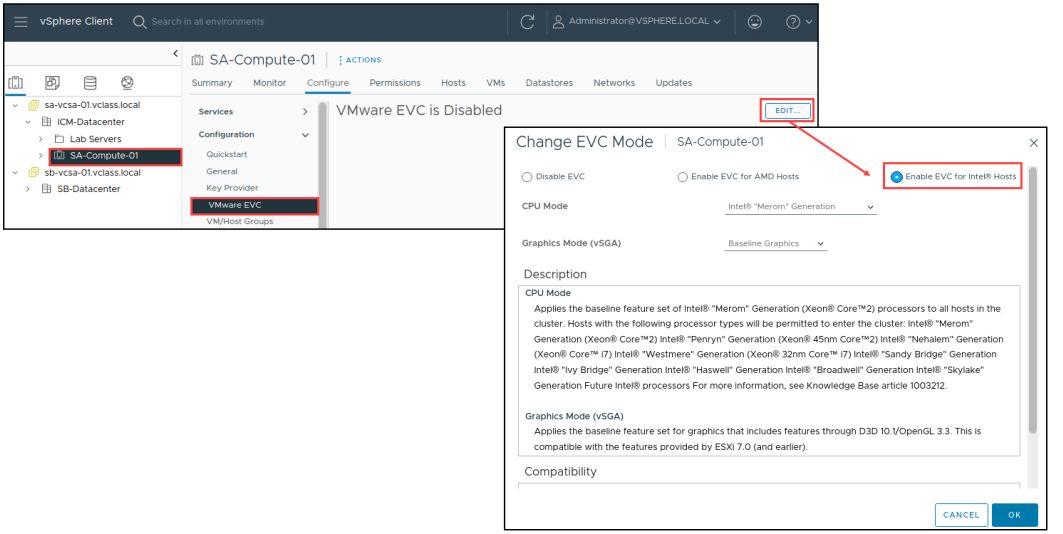

You can use one of the following methods to create an Enhanced vMotion Compatibility cluster:

- Create an empty cluster with EVC mode configured and move hosts into the cluster.
- Configure EVC mode on an existing cluster.

To activate EVC on an existing cluster, all hosts must be placed in maintenance mode, and therefore all VMs in the existing cluster must be powered off or suspended. To avoid VM downtime, the recommended method for activating EVC is to activate this feature on a new, empty cluster.

For information about Enhanced vMotion Compatibility processor support, see VMware knowledge base article 1003212 at http://kb.vmware.com/kb/1003212.

8-25 Changing the EVC CPU Mode for a Cluster

Several EVC mode approaches are available to ensure CPU compatibility:

- If all the hosts in a cluster are compatible with a newer EVC mode, you can change the EVC mode of an existing Enhanced vMotion Compatibility cluster.
- You can configure EVC mode for a cluster that does not have EVC mode configured.

You can raise or lower the EVC mode, but the VMs must be in the correct power state to do so.

EVC Mode	VM Power Action
Raise the EVC mode to a CPU baseline with more features.	• Running VMs can remain powered on. • New EVC mode features are not available to the VMs until they are powered off and powered back on again (Suspending and resuming the VM is not sufficient.)
Lower the EVC mode to a CPU baseline with fewer features.	• Power off VMs if they are powered on and running at a higher EVC mode than the one you intend to configure.

8-26 Virtual Machine EVC CPU Mode

EVC mode can be applied to some or all VMs in a cluster:

- At the VM level, EVC mode facilitates the migration of VMs beyond the cluster and across vCenter systems and data centers.
- You can apply more granular definitions of Enhanced vMotion Compatibility for specific VMs.
- VM EVC mode is independent of the EVC mode defined at the cluster level.

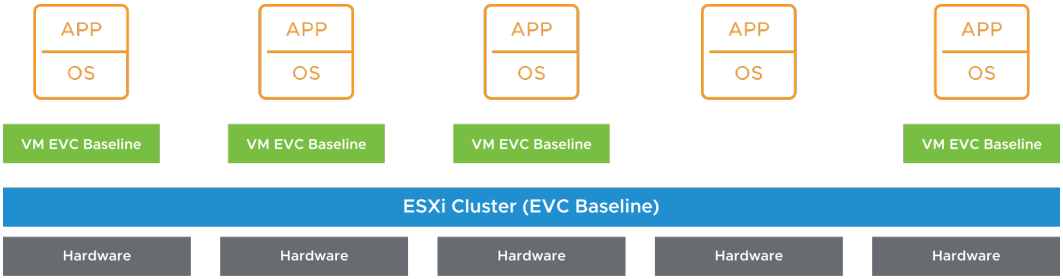

With per-VM EVC mode, the EVC mode becomes an attribute of the VM rather than the specific processor generation it happens to be booted on in the cluster. This feature supports seamless migration between two data centers that have different processors. Further, the feature is persisted per VM and does not lose the EVC mode during migrations across clusters or during power cycles.

In this diagram, EVC mode is not configured on the cluster. The cluster consists of differing CPU models with different feature sets. The VMs with per-VM EVC mode can run on any ESXi host that can satisfy the defined EVC mode.

8-27 Enhanced vMotion Compatibility for vSGA GPUs

EVC can also define a common baseline of GPU feature sets in a cluster.

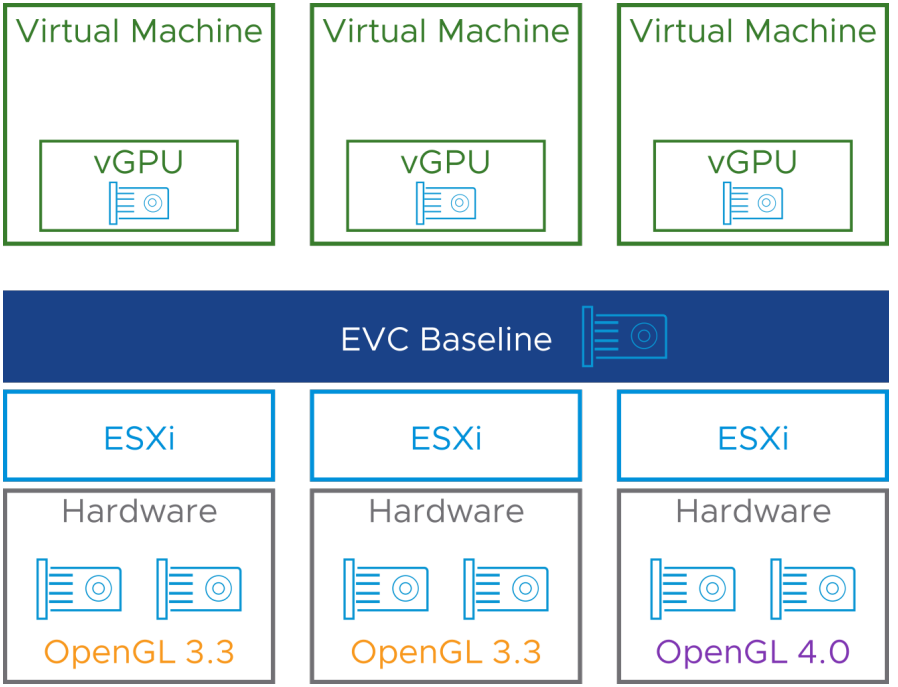

Features not included in the applied baseline are masked and not exposed to VMs.

Enhanced vMotion Compatibility for vSGA is supported with hardware GPUs and also software GPU renderers.

8-28 EVC Cluster Requirements for Graphics Mode

EVC for vSGA GPUs is configured at the ESXi cluster level:

- All ESXi hosts must satisfy GPU requirements of the defined baseline.
- Additional hosts cannot join the cluster if they cannot satisfy the baseline requirements.
- A mixed cluster of ESXi 6.7 and ESXi 7.0 hosts is supported when using Enhanced vMotion Compatibility at a cluster level.

EVC for vSGA GPUs is configured at a virtual machine level:

- VM compatibility for ESXi 7.0 Update 1 is required (virtual machine hardware version 18).
- VMs using GPU Enhanced vMotion Compatibility at a VM level must run on ESXi 7.0 Update 1.

8-29 Configuring EVC Graphics Mode on an Existing Cluster

At the cluster level, you configure EVC for vSGA in the same EVC settings as EVC for CPU.

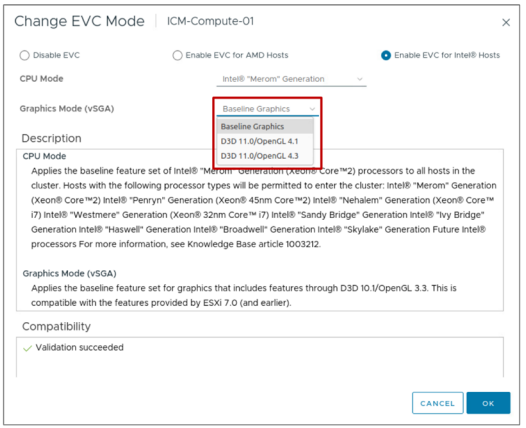

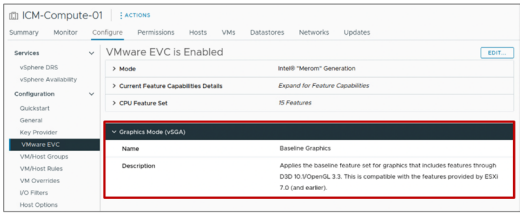

8-30 Virtual Machine EVC Graphics Mode

At the VM level, you configure EVC for vSGA in the same VM EVC settings as EVC for CPU.

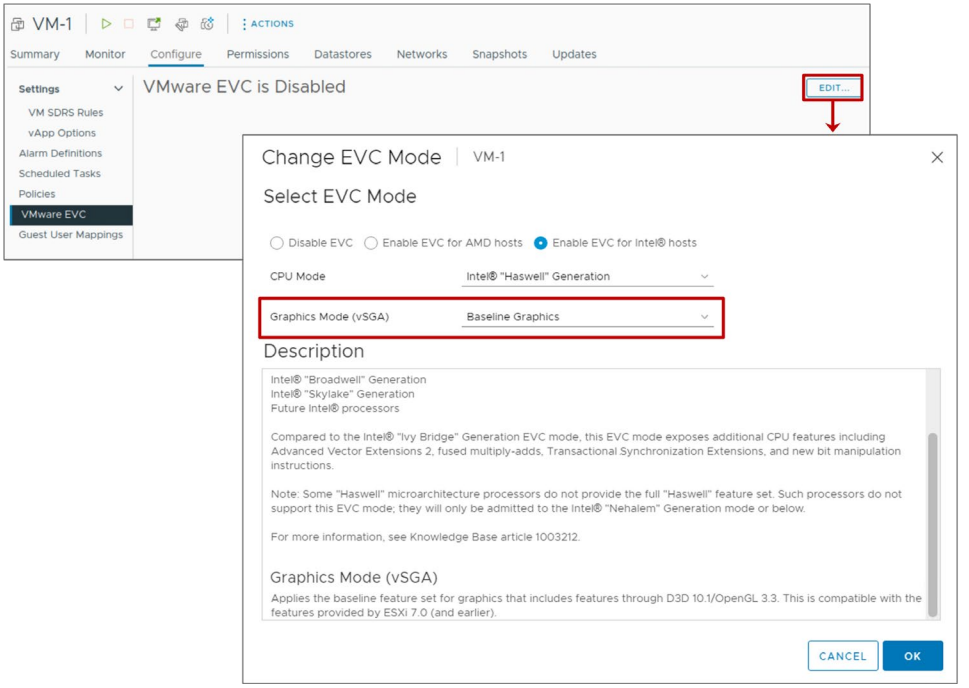

8-31 Review of Learner Objectives

- Describe the role of Enhanced vMotion Compatibility in migrations
- Configure EVC CPU mode on a vSphere cluster
- Explain how per-VM EVC CPU mode works with vSphere vMotion
- Configure EVC Graphics mode on a vSphere cluster or a VM

8-32 Lesson 3: Migrating VMs with vSphere Storage vMotion

8-33 Learner Objectives

- Explain how vSphere Storage vMotion works
- Recognize guidelines for using vSphere Storage vMotion
- Migrate virtual machines using vSphere Storage vMotion
- Migrate both the compute resource and storage of a virtual machine

8-34 About vSphere Storage vMotion

With vSphere Storage vMotion, you can migrate a powered-on VM from one datastore to another datastore of any type.

Using vSphere Storage vMotion, you can perform the following tasks:

- Move VMs off arrays for maintenance or to upgrade.
- Change the disk provisioning type.
- Change VM files on the destination datastore to match the inventory name of the VM.
- Migrate between datastores to balance traffic across storage paths and reduce latencies.
- Redistribute VMs or virtual disks to different storage volumes to balance capacity or improve performance.

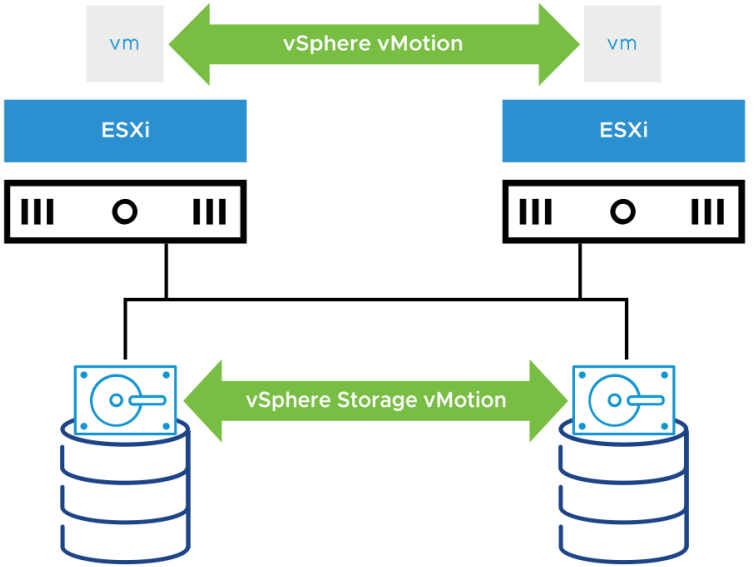

vSphere Storage vMotion provides flexibility to relocate virtual disks for performance or transform disk types, which you can use to reclaim space.

You can place the VM and all its disks in a single location, or you can select separate locations for the VM configuration file and each virtual disk. During a migration with vSphere Storage vMotion, the VM does not change the host that it runs on.

Whether performing a vSphere Storage vMotion migration or a cold migration, you can rename a VM's files on the destination datastore. The migration renames all virtual disk, configuration, snapshot, and `.nvram` files.

8-35 vSphere Storage vMotion In Action

vSphere Storage vMotion uses an I/O mirroring architecture to copy disk blocks between the source and destination.

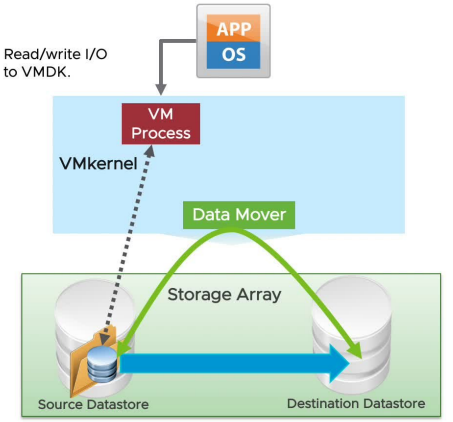

To play the animation, go to https://vmware.bravais.com/s/FnHZwq043PJ8dV3ZRV7p.

The vSphere Storage vMotion migration process includes the following steps:

1. Initiate storage migration.
2. Use the VMkernel data mover or vSphere Storage APIs – Array Integration to copy data.
3. Start a new VM process.
4. Mirror I/O calls to file blocks that are already copied to the virtual disk on the destination datastore.
5. Transition to the destination VM process to begin accessing the virtual disk copy.

The storage migration process does a single pass of the disk, copying all the blocks to the destination disk. If blocks are changed after they are copied, the blocks are synchronized from the source to the destination through the mirror driver, with no need for recursive passes.

This approach guarantees complete transactional integrity and is fast enough to be unnoticeable to the end user. The mirror driver uses the VMkernel data mover to copy blocks of data from the source disk to the destination disk. The mirror driver synchronously mirrors writes to both disks during the vSphere Storage vMotion operation.

Finally, vSphere Storage vMotion operations are performed either internally on a single ESXi host or offloaded to the storage array. Operations performed internally on the ESXi host use a data mover built into the VMkernel. Operations are offloaded to the storage array if the array supports vSphere Storage APIs - Array Integration, also called hardware acceleration.

8-36 Identifying Storage Arrays That Support vSphere Storage APIs – Array Integration

vSphere Storage vMotion offloads its operations to the storage array if:

- The array supports VMware vSphere Storage APIs – Array Integration, also called hardware acceleration.
- Both datastores are located within the same storage array.

Use the vSphere Client to determine whether your storage array supports hardware acceleration.

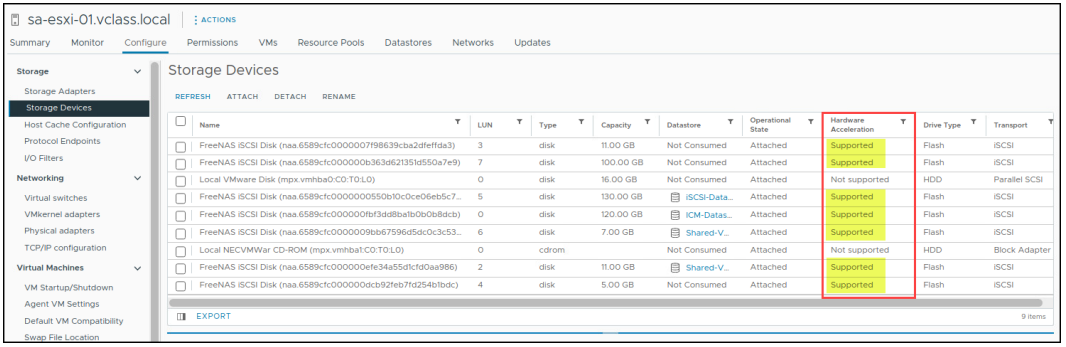

8-37 vSphere Storage vMotion Guidelines and Limitations

Guidelines:

- Plan the migration and coordinate with administrators
- Perform migrations during off-peak hours

Limitation:

- Independent virtual machine disks must be in persistent mode

A VM and its host must meet certain resource and configuration requirements for the virtual machine disks (VMDKs) to be migrated with vSphere Storage vMotion. One of the requirements is that the host on which the VM runs must have access both to the source datastore and to the target datastore.

During a migration with vSphere Storage vMotion, you can change the disk provisioning type. Migration with vSphere Storage vMotion changes VM files on the destination datastore to match the inventory name of the VM. The migration renames all virtual disk, configuration, snapshot, and `.nvram` files. If the new names exceed the maximum filename length, the migration does not succeed.

8-38 Changing Both Compute Resource and Storage During Migration

When you change both compute resource and storage during migration, a VM changes its host and datastore and optionally its network, either within or across vCenter instances.

- This technique combines vSphere vMotion and vSphere Storage vMotion into a single operation.
- You can migrate VMs across clusters, data centers, and vCenter instances.

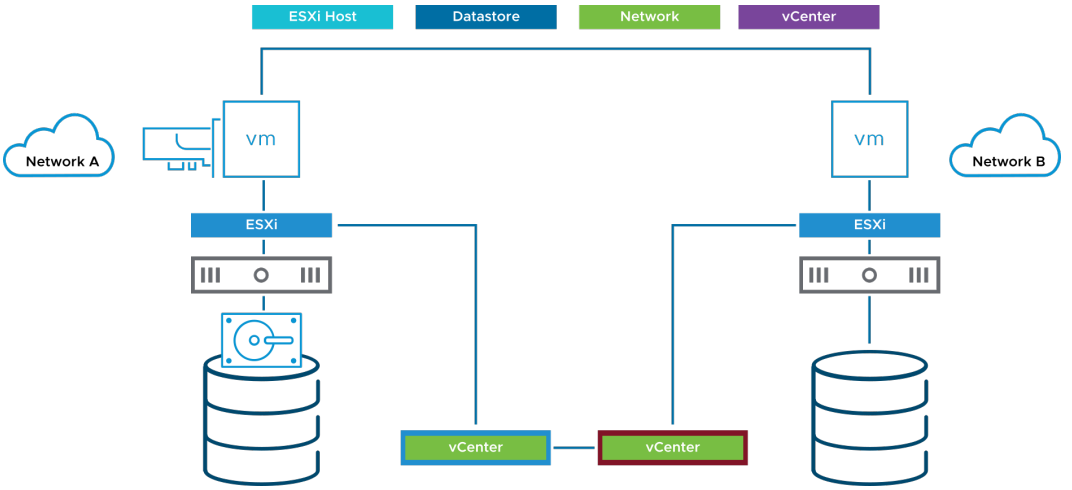

You can migrate VMs beyond storage accessibility boundaries and between hosts, within and across clusters, data centers, and vCenter instances.

8-39 Use Cases for Changing Both Compute Resource and Storage

Compute resource and storage migration is useful in the following cases:

- Migrating a VM located on a host's local storage to a host that uses shared storage
- Migrating a VM to a new cluster, when the target cluster does not have access to the source cluster's storage
- Migrating a VM to a different data center, whose storage is not shared with the source data center
- Migrating a VM to a different vCenter instance, whose storage is not shared with the source vCenter instance

8-40 Lab 21: vSphere Storage vMotion Migrations

Use vSphere Storage vMotion to migrate virtual machines:

1. Migrate Virtual Machine Files from One Datastore to Another
2. Migrate Both the Compute Resource and Storage of a Virtual Machine

8-41 Review of Learner Objectives

- Explain how vSphere Storage vMotion works
- Recognize guidelines for using vSphere Storage vMotion
- Migrate virtual machines using vSphere Storage vMotion
- Migrate both the compute resource and storage of a virtual machine

8-42 Lesson 4: Cross vCenter Migrations

8-43 Learner Objectives

- Recognize the types of VM migrations that you can perform across vCenter instances

8-44 About Cross vCenter Migrations

With vSphere vMotion, you can migrate VMs between vCenter instances, whether or not they are in the same Enhanced Linked Mode group.

Cross vCenter migrations are helpful in the following cases:

- Balancing workloads across clusters and vCenter instances that are in the same site or in another geographical area.
- Moving VMs between environments that have different purposes, for example, from a development environment to production environment.
- Moving VMs to meet different service level agreements (SLAs) for storage space, performance, and so on.
- Moving VMs from an on-premises vSphere data center to a data center in the public cloud, such as VMware Cloud on AWS.

8-45 Cross vCenter Migration Requirements

Cross vCenter migrations have the following requirements:

- Hosts must be time-synchronized.
- Both vCenter instances can be in the same or different vCenter Single Sign-On domain.

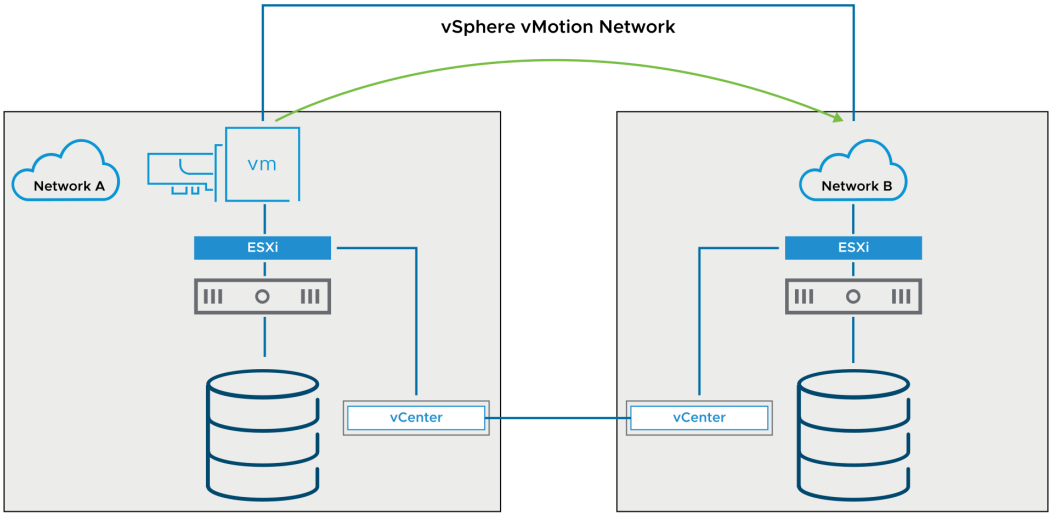

Cross vCenter migrations must meet the following requirements:

- Both vCenter instances must be time-synchronized with each other for correct vCenter Single Sign-On token verification.
- Both vCenter instances do not need to be in the same Enhanced Linked Mode group.
- For migration of compute resources only, both vCenter instances must be connected to the shared storage on which the VM is located.

You can perform cross vCenter migrations between vCenter instances of different versions. For information on the supported versions, see VMware knowledge base article 2106952 at http://kb.vmware.com/kb/2106952.

8-46 Performing a Cross vCenter vMotion in Same SSO Domain

In the Migrate wizard, select a compute resource at the destination vCenter instance.

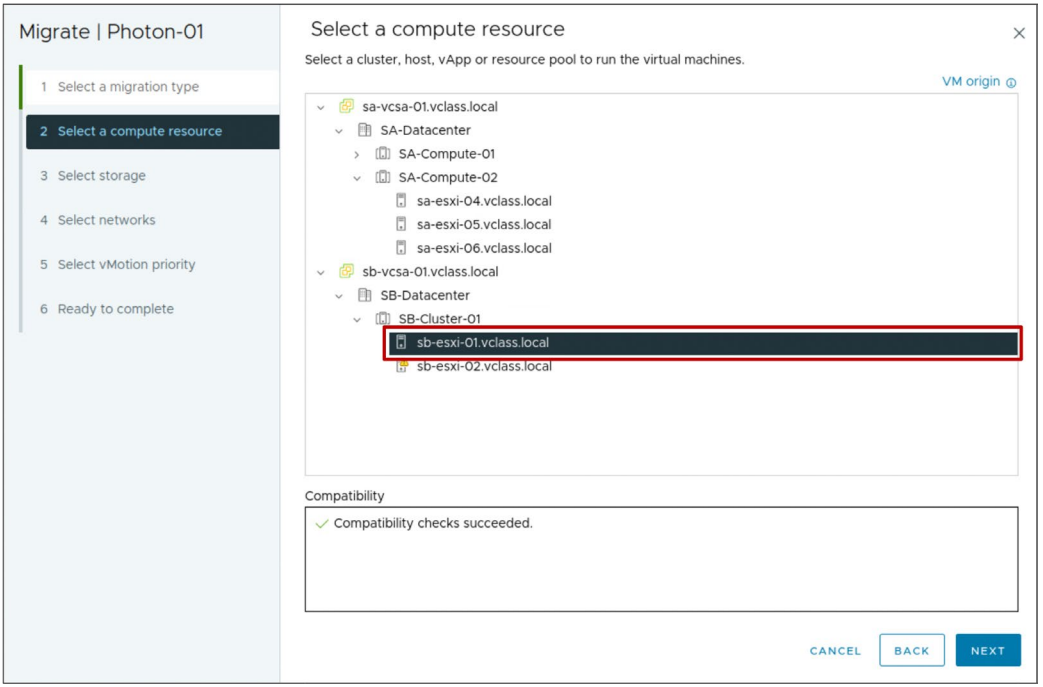

8-47 Performing a Cross vCenter vMotion in Different SSO Domain (1)

From the Migrate wizard, select **Cross vCenter Server export** as the migration type.

Optionally, you can clone the VM rather than migrate it.

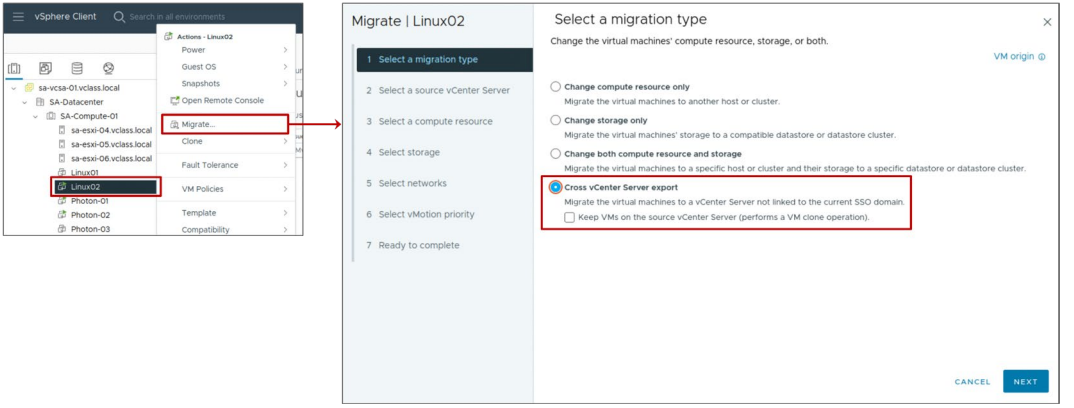

In the Migrate wizard, you specify the FQDN or IP address of the target vCenter instance and the vCenter credentials. The other wizard options are similar to the option for changing both compute resource and storage: you select the compute resource, datastore, and VM network to use in the target vCenter instance.

8-48 Performing a Cross vCenter vMotion in Different SSO Domain (2)

Then, you specify the FQDN or IP address of the target vCenter instance and the vCenter credentials.

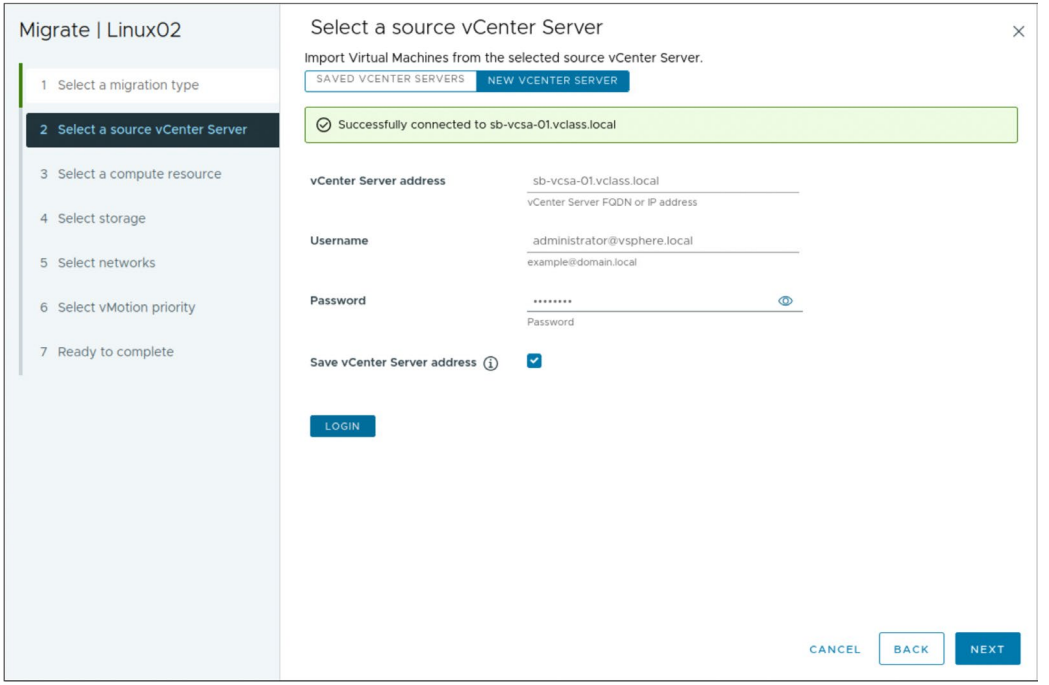

The other wizard options are similar to the option for changing both compute resource and storage: you select the compute resource, datastore, and VM network to use in the target vCenter instance.

8-49 Network Checks for Cross vCenter Migrations

vCenter performs several network compatibility checks to prevent the following configuration problems:

- MAC address incompatibility on the destination host
- vSphere vMotion migration from a distributed switch to a standard switch
- vSphere vMotion migration between distributed switches of different versions

8-50 VMkernel Networking Layer and TCP/IP Stacks

The VMkernel networking layer provides connectivity to hosts and handles the standard system traffic of vSphere vMotion, IP storage, vSphere Fault Tolerance, vSAN, and others.

TCP/IP stacks at the VMkernel level that are configured by default:

- Default TCP/IP stack
- vSphere vMotion TCP/IP stack
- Provisioning TCP/IP stack

You can also create a custom TCP/IP stack.

Consider the following key points about TCP/IP stacks at the VMkernel level:

- Default TCP/IP stack: Provides networking support for the management traffic between vCenter and ESXi hosts and for system traffic such as vSphere vMotion, IP storage, and vSphere Fault Tolerance.
- vSphere vMotion TCP/IP stack: Supports the traffic for hot migrations of VMs.
- Provisioning TCP/IP stack: Supports the traffic for VM cold migration, cloning, and snapshot creation. You can use the provisioning TPC/IP stack to handle NFC traffic during long-distance vSphere vMotion migration. VMkernel adapters configured with the provisioning TCP/IP stack handle the traffic from cloning the virtual disks of the migrated VMs in long-distance vSphere vMotion.

  By using the provisioning TCP/IP stack, you can isolate the traffic from the cloning operations on a separate gateway. After you configure a VMkernel adapter with the provisioning TCP/IP stack, all adapters on the default TCP/IP stack are deactivated for the provisioning traffic.
- Custom TCP/IP stacks: You can create a custom TCP/IP stack on a host to forward networking traffic through a custom application. Open an SSH connection to the host and run the vSphere CLI command:

  ```
  esxcli network ip netstack add -N="stack_name"
  ```

Take appropriate security measures to prevent unauthorized access to the management and system traffic in your vSphere environment. For example, isolate the vSphere vMotion traffic in a separate network that includes only the ESXi hosts that participate in the migration. Isolate the management traffic in a network that only network and security administrators can access.

8-51 vSphere vMotion TCP/IP Stacks

Each ESXi host has a second TCP/IP stack that is dedicated to vSphere vMotion migration.

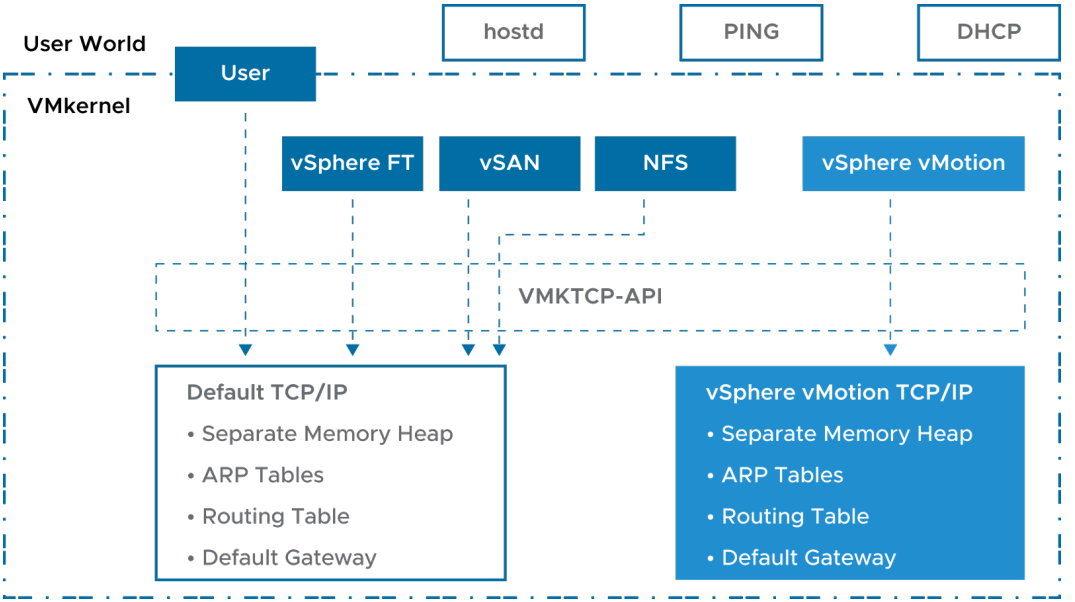

vSphere vMotion TCP/IP stacks support the traffic for hot migrations of VMs. Use the vSphere vMotion TCP/IP stack to provide better isolation for the vSphere vMotion traffic. After you create a VMkernel adapter on the vSphere vMotion TCP/IP stack, you can use only this stack for vSphere vMotion migration on this host.

The VMkernel adapters on the default TCP/IP stack are deactivated for the vSphere vMotion service after you create a VMkernel adapter on the vSphere vMotion TCP/IP stack. If a hot migration uses the default TCP/IP stack while you configure VMkernel adapters with the vMotion TCP/IP stack, the migration completes successfully. However, these VMkernel adapters on the default TCP/IP stack are deactivated for future vSphere vMotion sessions.

8-52 About Long-Distance vSphere vMotion Migration

Long-distance vSphere vMotion migration is an extension of cross vCenter migration.

vCenter instances are spread across large geographic distances and where the latency across sites is high.

Use cases for long-distance vSphere vMotion migration:

- Permanent migrations
- Disaster avoidance
- Site Recovery Manager and disaster avoidance testing
- Multisite load balancing
- Follow-the-sun scenario support

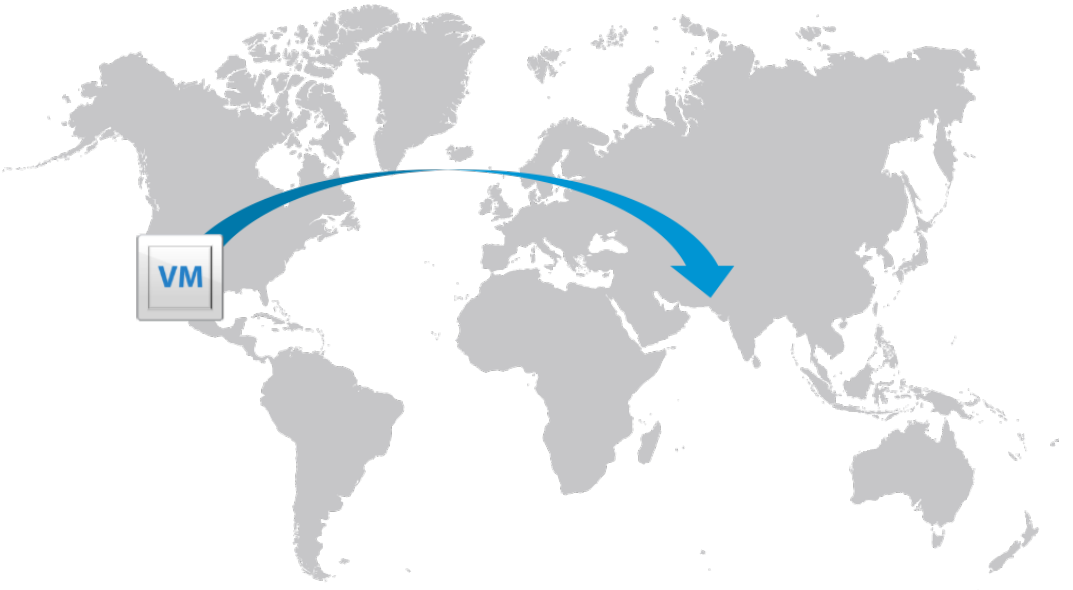

In the follow-the-sun scenario, a global support team might support a certain set of VMs. As one support team ends their workday, another support team in a different timezone takes over support duty. The VMs being supported can be moved from one geographical location to another so that the support team on duty can access those VMs locally instead of long distance.

8-53 Networking Prerequisites for Long-Distance vSphere vMotion

Long-distance vSphere vMotion migrations must connect over layer 3 connections:

- Virtual machine network:
  - L2 connection
  - The same VM IP address is available at the destination
- vSphere vMotion network:
  - L3 connection
  - 250 Mbps per vSphere vMotion operation
  - Round-trip time between hosts can take up to 150 milliseconds.

8-54 Review of Learner Objectives

- Recognize the types of VM migrations that you can perform across vCenter instances

8-55 Lesson 5: Creating Virtual Machine Snapshots

8-56 Learner Objectives

- Take a snapshot of a virtual machine
- Manage multiple snapshots
- Delete virtual machine snapshots
- Consolidate snapshots

8-57 About VM Snapshots

With snapshots, you can preserve the state of the VM so that you can repeatedly return to the same state.

For example, if problems occur during the patching or upgrading process, you can revert to the previous state.

VM snapshots are not recommended as a VM backup strategy.

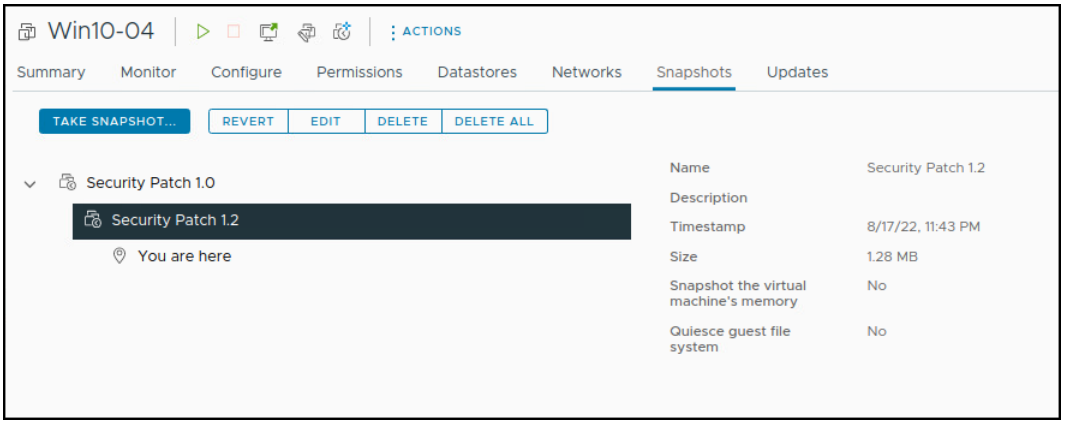

Snapshots are useful when you want to revert repeatedly to the same state but do not want to create multiple VMs. Examples include patching or upgrading the guest operating system in a VM.

The relationship between snapshots is like the relationship between a parent and a child. Snapshots are organized in a snapshot tree. In a snapshot tree, each snapshot has one parent and one or more children, except for the last snapshot, which has no children.

8-58 Taking Snapshots

You can take a snapshot while a VM is powered on, powered off, or suspended.

A snapshot captures the following items:

- VM configuration
- VM memory state (optional)
- Virtual disks

A snapshot capture does not include Independent virtual disks (persistent and nonpersistent).

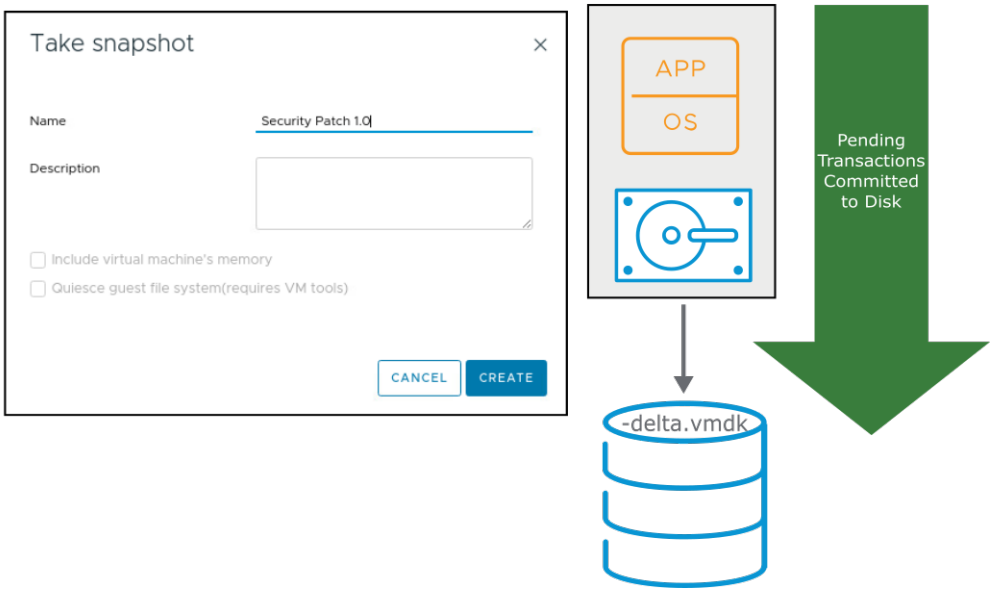

A snapshot captures the entire state of the VM at the time that you take the snapshot, including the following states:

- Memory state: The contents of the VM's memory. The memory state is captured only if the VM is powered on and if you select the **Snapshot the virtual machine's memory** check box (selected by default).
- Settings state: The VM settings.
- Disk state: The state of all the VM's virtual disks.

At the time that you take the snapshot, you can also quiesce the guest operating system. This action quiesces the file system of the guest operating system. This option is available only if you do not capture the memory state as part of the snapshot.

8-59 Types of Snapshots

A delta or child disk is created when you create a snapshot:

- On the datastore, the delta disk is a sparse disk.
- Delta disks use different sparse formats depending on the type of datastore.

Snapshot Type	Datastore Type	Filename	Block Size
VMFSsparse	• VMFS5 with virtual disks smaller than 2 TB	`#-delta.vmdk`	512 bytes
SEsparse	• VMFS6 • VMFS5 with virtual disks larger than 2 TB • Space efficient (thin provisioned) • Supports disk reclamation (unmap)	`#-sesparse.vmdk`	4 KB
vsanSparse	• vSAN	`Delta object`	4 MB

Delta disks use different sparse formats depending on the type of datastore.

- VMFSsparse: VMFS5 uses the VMFSsparse format for virtual disks smaller than 2 TB. VMFSsparse is implemented on top of VMFS. The VMFSsparse layer processes I/O operations issued to a snapshot VM. Technically, VMFSsparse is a redo log that starts empty, immediately after a VM snapshot is taken. The redo log expands to the size of its base VMDK, when the entire VMDK is rewritten with new data after the VM snapshot. This redo log is a file in the VMFS datastore. On snapshot creation, the base VMDK attached to the VM is changed to the newly created sparse VMDK.
- SEsparse: SEsparse is a default format for all delta disks on the VMFS6 datastores. On VMFS5, SEsparse is used for virtual disks of the size 2 TB and larger. SEsparse is a format that is like VMFSsparse with some enhancements. This format is space efficient and supports the space-reclamation technique. With space reclamation, blocks that the guest OS deletes are marked. The system sends commands to the SEsparse layer in the hypervisor to unmap those blocks. The unmapping helps to reclaim space allocated by SEsparse after the guest operating system deletes the data.

8-60 VM Snapshot Files

A snapshot consists of a set of files:

- `-Snapshot#.vmsn`: Configuration state
- `-Snapshot#.vmem`: Memory state (optional)
- `-00000#.vmdk`: Disk descriptor
- `-00000#-delta.vmdk`: VMFS5 delta
- `-00000#-sesparse.vmdk`: VMFS6 delta
- `.vmsd`: Stores names, descriptions, and relationships for all the VM's snapshots

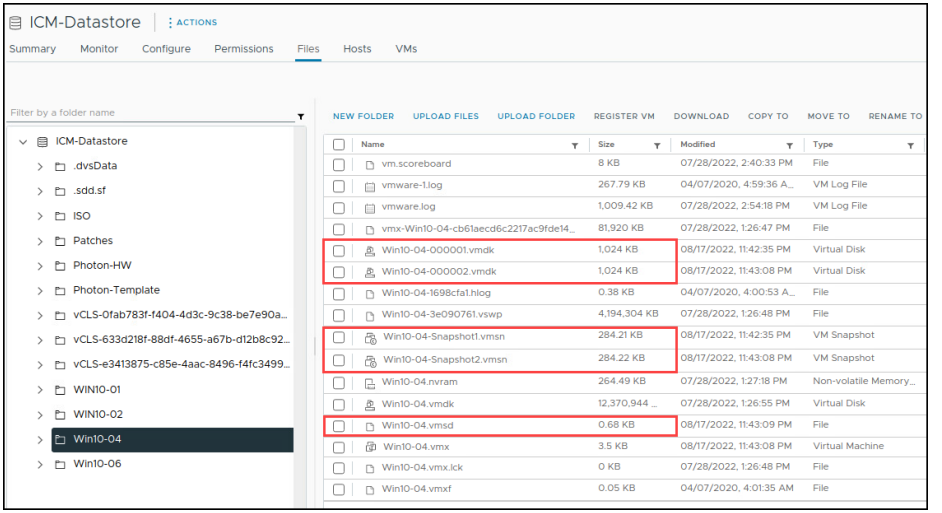

A VM can have one or more snapshots. For each snapshot, the following files are created:

- Snapshot delta file: This file contains the changes to the virtual disk's data since the snapshot was taken. When you take a snapshot of a VM, the state of each virtual disk is preserved. The VM stops writing to its `-flat.vmdk` file. Writes are redirected to `>-######-delta.vmdk` (or `-######-sesparse.vmdk`) instead (for which `######` is the next number in the sequence). You can exclude one or more virtual disks from a snapshot by designating them as independent disks. Configuring a virtual disk as independent is typically done when the virtual disk is created, but this option can be changed whenever the VM is powered off.
- Disk descriptor file: `-00000#.vmdk`. This file is a small text file that contains information about the snapshot.

- Configuration state file: –`.vmsn`. # is the next number in the sequence, starting with 1. This file holds the active memory state of the VM at the point that the snapshot was taken, including virtual hardware, power state, and hardware version.
- Memory state file: –`.vmem`. This file is created if the option to include memory state was selected during the creation of the snapshot. It contains the entire contents of the VMs at the time that the snapshot of the VM was taken.
- Snapshot active memory file: –`.vmem`. This file contains the contents of the VM memory if the option to include memory is selected during the creation of the snapshot.
- The `.vmsd` file is the snapshot list file and is created at the time that the VM is created. It maintains snapshot information for a VM so that it can create a snapshot list in the vSphere Client. This information includes the name of the snapshot `.vmsn` file and the name of the virtual disk file.
- The snapshot state file has a `.vmsn` extension and is used to store the state of a VM when a snapshot is taken. A new `.vmsn` file is created for every snapshot that is created on a VM and is deleted when the snapshot is deleted. The size of this file varies, based on the options selected when the snapshot is created. For example, including the memory state of the VM in the snapshot increases the size of the `.vmsn` file.

You can exclude one or more of the VMDKs from a snapshot by designating a virtual disk in the VM as an independent disk. Placing a virtual disk in independent mode is typically done when the virtual disk is created. If the virtual disk was created without activating independent mode, you must power off the VM to activate it.

Other files might also exist, depending on the VM hardware version. For example, each snapshot of a VM that is powered on has an associated `_.vmem` file, which contains the guest operating system main memory, saved as part of the snapshot.

8-61 VM Snapshot Files Example (1)

VM with
no snapshots


```
Win10-01.vmsd
Win10-01-flat.vmdk
Win10-01.vmdk
```

This example shows the snapshot and virtual disk files that are created when a VM has no snapshots.

8-62 VM Snapshot Files Example (2)

VM with
no snapshots


```
Win10-01.vmsd
Win10-01-flat.vmdk
Win10-01.vmdk
```

First snapshot taken
(with memory state)


```
Win10-01-Snapshot1.vmem
Win10-01-Snapshot1.vmsn
Win10-01-000001-sesparse.vmdk
Win10-01-000001.vmdk
```

This example shows the snapshot and virtual disk files that are created when a VM has one snapshot.

8-63 VM Snapshot Files Example (3)

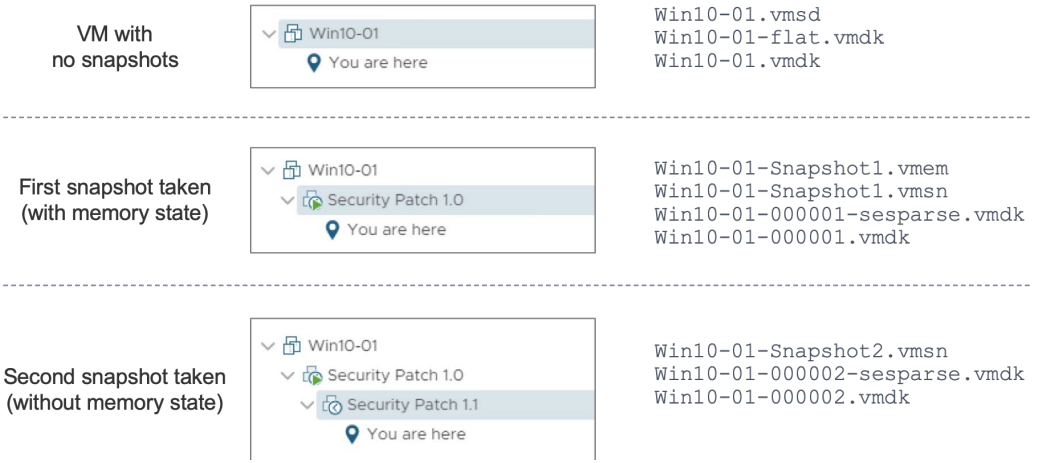

This example shows the snapshot and virtual disk files that are created when a VM has two snapshots.

8-64 Managing Snapshots

In the vSphere Client, you can view snapshots for the active VM and take edit, delete, and revert to actions.

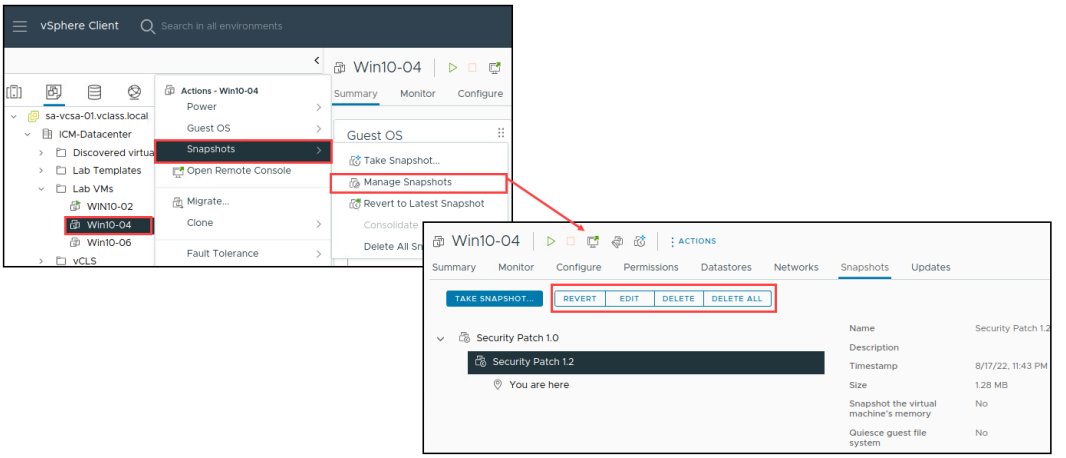

You can perform the following actions from the Manage Snapshots window:

- Edit the snapshot: Edit the snapshot name and description.
- Delete the snapshot: Remove the snapshot from the Snapshot Manager, consolidate the snapshot files to the parent snapshot disk, and merge with the VM base disk.
- Delete all snapshots: Commit all the intermediate snapshots before the current-state icon (You are here) to the VM and remove all snapshots for that VM.
- Revert to a snapshot: Restore, or revert to, a particular snapshot. The snapshot that you restore becomes the current snapshot.

When you revert to a snapshot, you return all these items to the state that they were in at the time that you took the snapshot. If you want the VM to be suspended, powered on, or powered off when you start it, ensure that the VM is in the correct state when you take the snapshot.

Deleting a snapshot (**DELETE** or **DELETE ALL**) consolidates the changes between snapshots and previous disk states. Deleting a snapshot also writes to the parent disk all data from the delta disk that contains the information about the deleted snapshot. When you delete the base parent snapshot, all changes merge with the base VMDK.

8-65 Deleting VM Snapshots (1)

If you delete a snapshot one or more levels above the, `You are here` level, the snapshot state is deleted. In this example, the snap01 data is committed into the parent (base disk), and the foundation for snap02 is retained.

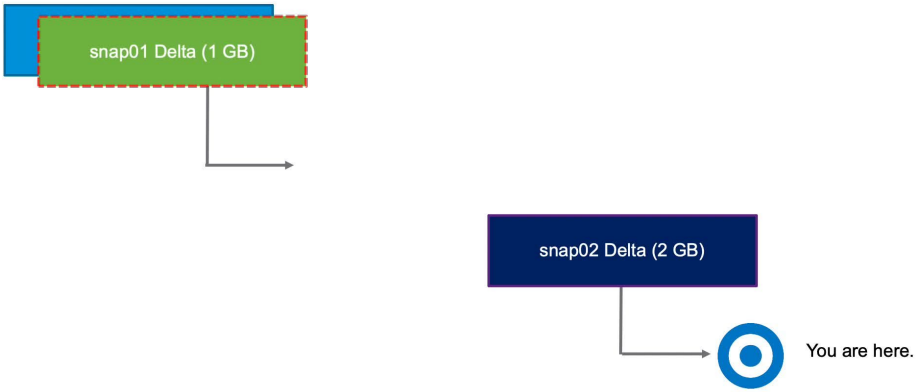

To play the animation, go to https://vmware.bravais.com/s/WhbcXR4sSwk2Vl7MeaXD.

8-66 Deleting VM Snapshots (2)

If you delete the latest snapshot, the changes are committed to its parent. The snap02 data is committed into snap01 data, and the snap02 `-delta.vmdk` file is deleted.

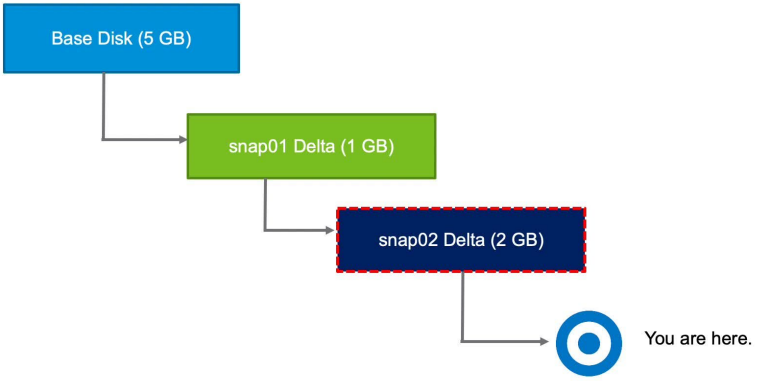

To play the animation, go to https://vmware.bravais.com/s/l0JYYQzMTv7pvxBqNcQp.

8-67 Deleting VM Snapshots (3)

If you delete a snapshot one or more levels below the `You are here level`, subsequent snapshots are deleted, and you can no longer return to those states. The snap02 data is deleted.

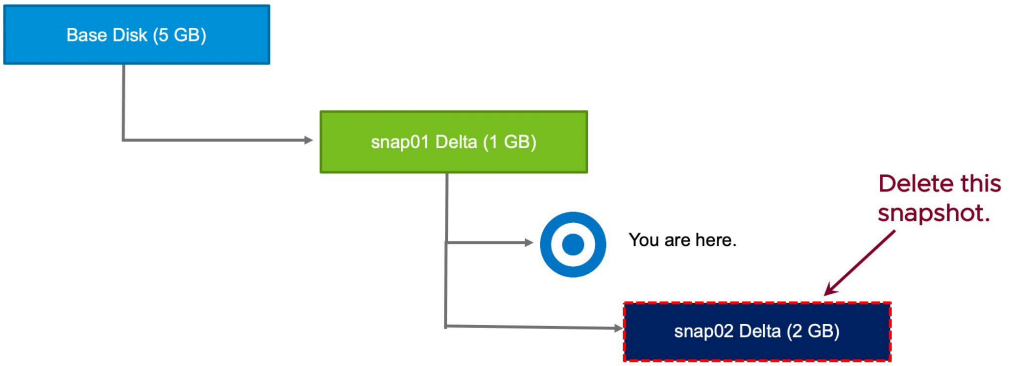

To play the animation, go to https://vmware.bravais.com/s/NiQxPT3iycemQ8WYXKom.

8-68 Deleting All VM Snapshots

The delete-all-snapshots mechanism uses storage space efficiently. The size of the base disk does not increase. Snap01 is committed to the base disk before snap02 is committed.

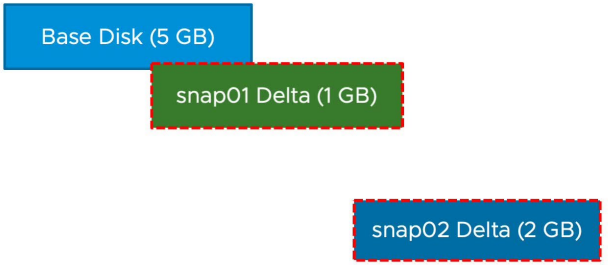

To play the animation, go to https://vmware.bravais.com/s/L3ilQHlrywEhlgr5p7RP.

All snapshots before the `You are here` point are committed all the way up to the base disk. All snapshots after `You are here` are discarded.

Like a single snapshot deletion, changed blocks in the snapshot overwrite their counterparts in the base disk.

8-69 About Snapshot Consolidation

Snapshot consolidation is a method for committing a chain of delta disks to the base disks when the Snapshot Manager shows that no snapshots exist, but the delta disk files remain on the datastore.

Snapshot consolidation resolves problems that might occur with snapshots:

- The snapshot descriptor file is committed correctly, and the Snapshot window shows that all the snapshots are deleted.
- The snapshot files (`-delta.vmdk` or `-sesparse.vmdk`) still exist in the VM's folder on the datastore.
- Snapshot files can continue to expand until they reach the size of the `-flat.vmdk` file or until the datastore runs out of space.

Snapshot consolidation is a way to clean unneeded delta disk files from a datastore. If no snapshots are registered for a VM, but delta disk files exist, snapshot consolidation commits the chain of the delta disk files and removes them.

If consolidation is not performed, the delta disk files might expand to the point of consuming all the remaining space on the VM’s datastore or the delta disk file reaches its configured size. The delta disk cannot be larger than the size configured for the base disk.

8-70 Discovering When to Consolidate Snapshots

On the **Monitor** tab under **All Issues** for the VM, a warning notifies you that a consolidation is required.

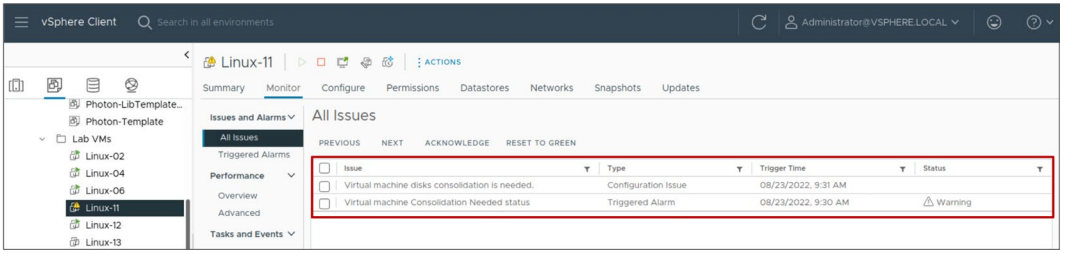

With snapshot consolidation, vCenter displays a warning when the descriptor and the snapshot files do not match. After the warning displays, you can use the vSphere Client to commit the snapshots.

8-71 Consolidating Snapshots

After the snapshot consolidation warning appears, you can use the vSphere Client to consolidate the snapshots.

All snapshot delta disks are committed to the base disks.

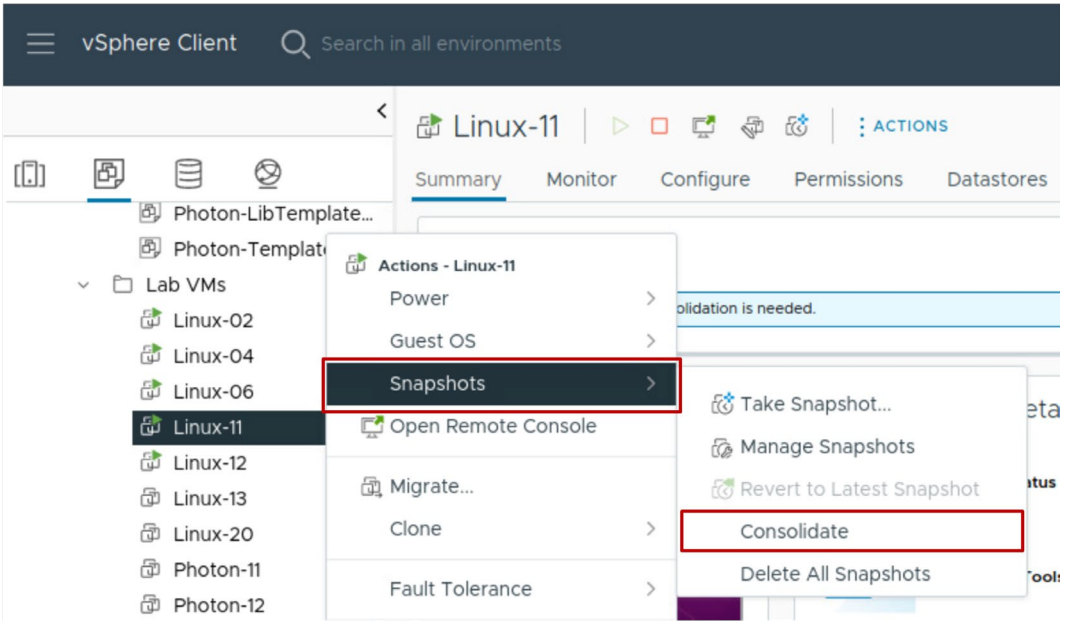

For a list of best practices for using snapshots in a vSphere environment, see VMware knowledge base article 1025279 at http://kb.vmware.com/kb/1025279.

8-72 Lab 22: Working with Snapshots

Take VM snapshots, revert a VM to a different snapshot, and delete snapshots:

1. Take Snapshots of a Virtual Machine
2. Add Files and Take Another Snapshot of a Virtual Machine
3. Revert the Virtual Machine to a Snapshot
4. Delete a Snapshot
5. Delete All Snapshots

8-73 Review of Learner Objectives

- Take a snapshot of a virtual machine
- Manage multiple snapshots
- Delete virtual machine snapshots
- Consolidate snapshots

8-74 Lesson 6: Virtual CPU and Memory Concepts

8-75 Learner Objectives

- Describe CPU and memory concepts in relation to a virtualized environment
- Recognize techniques for addressing memory resource overcommitment
- Identify additional technologies that improve memory use
- Describe how VMware Virtual SMP works
- Explain how the VMkernel uses hyperthreading

8-76 Memory Virtualization Basics

vSphere has the following layers of memory:

- Guest OS virtual memory is presented to applications by the operating system
- Guest OS physical memory is presented to the virtual machine by the VMkernel
- Host machine memory that is managed by the VMkernel provides a contiguous, addressable memory space that is used by the VM

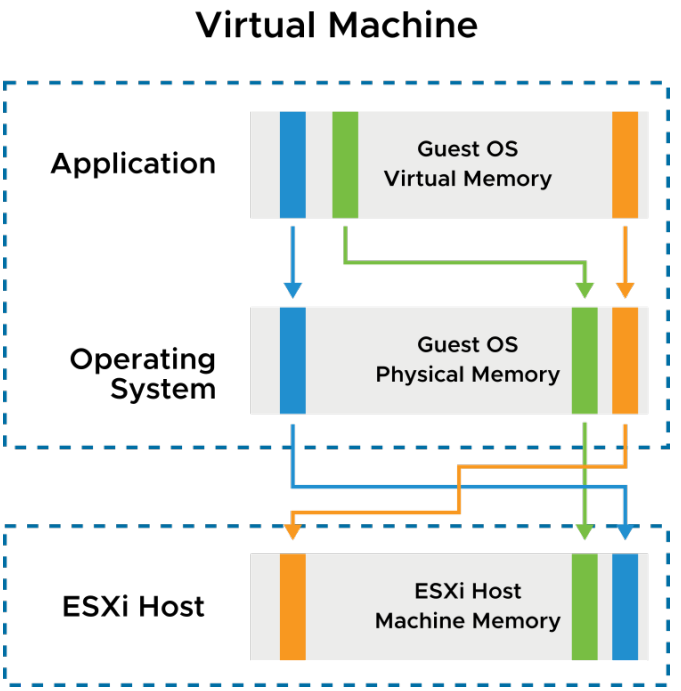

When running a virtual machine, the VMkernel creates a contiguous addressable memory space for the VM. This memory space has the same properties as the virtual memory address space presented to applications by the guest operating system. This memory space allows the VMkernel to run multiple VMs simultaneously while protecting the memory of each VM from being accessed by others. From the perspective of an application running in the VM, the VMkernel adds an extra level of address translation that maps the guest physical address to the host physical address.

8-77 VM Memory Overcommitment

Memory is overcommitted when the combined configured memory footprint of all powered-on VMs exceeds that of the host memory sizes.

When memory is overcommitted:

- VMs do not always use their full allocated memory
- To improve memory use, an ESXi host reclaims memory from idle VMs to allocate to VMs that need more memory
- VM memory can be swapped out to the `.vswp` file
- VM memory overhead can be swapped out to the `vmx-*.vswp` file

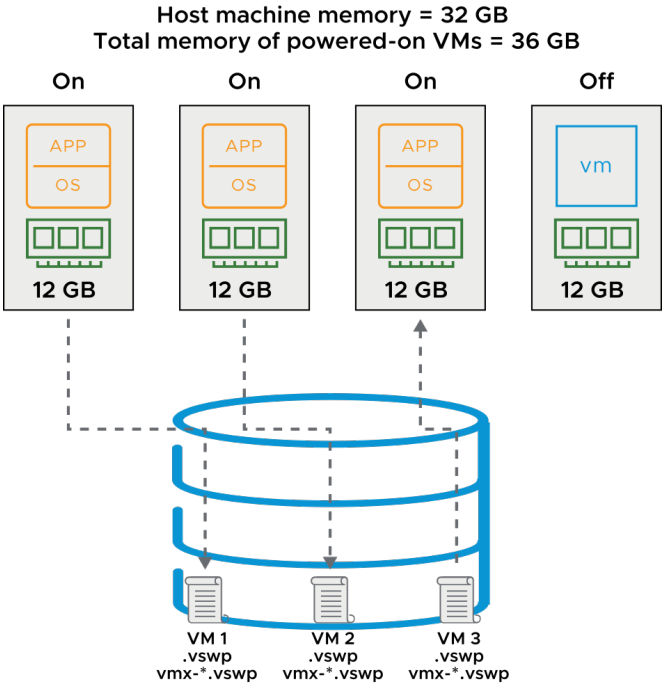

The total configured memory sizes of all VMs might exceed the amount of available physical memory on the host. However, this condition does not necessarily mean that memory is overcommitted. Memory is overcommitted when the working memory size of all VMs exceeds that of the ESXi host's physical memory size.

Because of the memory management techniques used by the ESXi host, your VMs can use more virtual RAM than the available physical RAM on the host. For example, you can have a host with 32 GB of memory and three VMs running with 12 GB of memory each. In that case, the

memory is overcommitted. If all three VMs are idle, the combined consumed memory is below 32 GB. However, if all VMs are actively consuming memory, then their memory footprint might exceed 32 GB and the ESXi host becomes overcommitted. An ESXi host can run out of memory if VMs consume all reservable memory in an overcommitted memory environment. Although the powered-on VMs are not affected, a new VM might fail to power on because of lack of memory. Overcommitment makes sense because, typically, some VMs are lightly loaded whereas others are more heavily loaded, and relative activity levels vary over time.

VM memory from this file is swapped out to disk when host machine memory is overcommitted. Extra memory from a VM is gathered into a swap file with the `.vswp` extension. The host uses the `vmx-*.vswp` swap file to gather and track memory overhead. Memory overhead refers to memory used by the VMX (VM Executable) process.

8-78 Memory Overcommit Techniques

An ESXi host uses memory overcommit techniques to allow the overcommitment of memory while possibly avoiding the need to page memory out to disk.

Methods Used by the ESXi Host	Details
Transparent page sharing	This method economizes the use of physical memory pages. In this method, pages with identical contents are stored only once.
Ballooning	This method uses the VMware Tools balloon driver to deallocate memory from virtual machines. The ballooning mechanism becomes active when memory is scarce, sometimes forcing VMs to use their own paging areas.
Memory compression	This method reduces a VM's memory footprint by storing memory in a compressed format.
Host-level SSD swapping	The ESXi host can swap out memory to locally-attached solid-state drives.
VM memory paging to disk	Using VMkernel swap space is the last resort because of poor performance.

The VMkernel uses various techniques to dynamically reduce the amount of physical RAM that is required for each VM. Each technique is described in the order that the VMkernel uses it:

- Page sharing: ESXi can use a proprietary technique to transparently share memory pages between VMs, eliminating redundant copies of memory pages. Although pages are shared by default within VMs, as of vSphere 6.0, pages are no longer shared by default among VMs.
- Ballooning: If the host memory begins to get low and the VM's memory use approaches its memory target, ESXi uses ballooning to reduce that VM's memory demands. Using the VMware-supplied `vmmemctl` module installed in the guest operating system as part of VMware Tools, ESXi can cause the guest operating system to relinquish the memory pages it considers least valuable. Ballooning provides performance closely matching that of a native system under similar memory constraints. To use ballooning, the guest operating system must be configured with sufficient swap space.
- Memory compression: If the VM's memory use approaches the level at which host-level swapping is required, ESXi uses memory compression to reduce the number of memory pages that it must swap out. Because the decompression latency is much smaller than the

swap-in latency, compressing memory pages has significantly less impact on performance than swapping out those pages.

- Swap to host cache: Host swap cache is an optional memory reclamation technique that uses local flash storage to cache a virtual machine's memory pages. By using local flash storage, the virtual machine avoids the latency associated with a storage network that might be used if it swapped memory pages to the virtual swap (`.vswp`) file.
- Regular host-level swapping: When memory pressure is severe and the hypervisor must swap memory pages to disk, the hypervisor swaps to a host swap cache rather than to a `.vswp` file. When a host runs out of space on the host cache, a virtual machine's cached memory is migrated to a virtual machine's regular `.vswp` file. Each host must have its own host swap cache configured.

8-79 Configuring Multicore VMs

You can build VMs with multiple virtual CPUs (vCPUs). The number of vCPUs that you configure for a single VM depends on the physical architecture of the ESXi host.

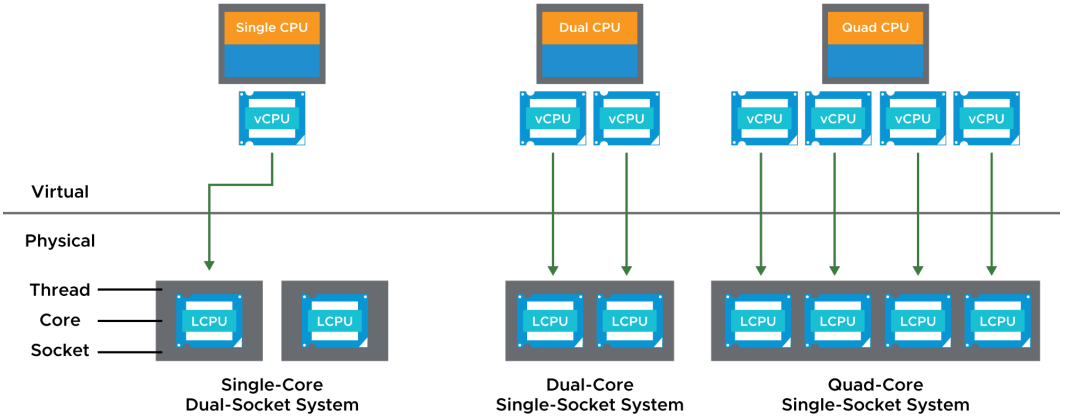

In addition to the physical host configuration, the number of vCPUs configured for a VM also depends on the guest operating system, the needs and scalability limits of the applications, and the specific use case for the VM itself.

The VMkernel includes a CPU scheduler that dynamically schedules vCPUs on the physical CPUs of the host system.

The VMkernel scheduler considers socket-core-thread topology when making scheduling decisions. Intel and AMD processors combine multiple processor cores into a single integrated circuit, called a socket in this discussion.

A socket is a single package with one or more physical CPUs. Each core has one or more logical CPUs (LCPU in the diagram) or threads. With logical CPUs, the core can schedule one thread of execution.

On the slide, the first system is a single-core, dual-socket system with two cores and, therefore, two logical CPUs.

When a vCPU of a single-vCPU or multi-vCPU VM must be scheduled, the VMkernel maps the vCPU to an available logical processor.

8-80 About Hyperthreading

With hyperthreading, a core can execute two threads or sets of instructions at the same time:

- Hyperthreading provides more logical CPUs on which vCPUs can be scheduled
- Hyperthreading is activated by default

To activate hyperthreading:

- Verify that the host system supports hyperthreading
- Activate hyperthreading in the system BIOS
- Ensure that hyperthreading for the ESXi host is turned on

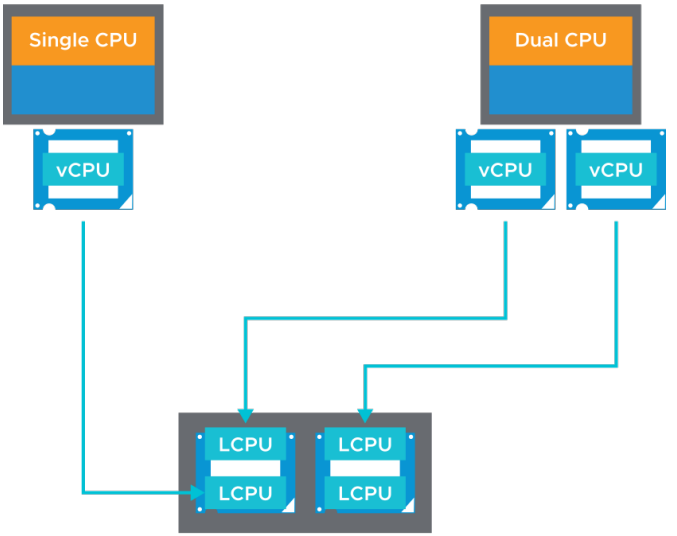

Dual-Core Single-Socket System with Hyperthreading

If hyperthreading is activated, ESXi can schedule two threads at the same time on each processor core (physical CPU). The drawback of hyperthreading is that it does not double the power of a core. So, if both threads of execution need the same on-chip resources at the same time, one thread has to wait. Still, on systems that use hyperthreading technology, performance is improved.

An ESXi host that is activated for hyperthreading should behave almost exactly like a standard system. Logical processors on the same core have adjacent CPU numbers. Logical processors 0 and 1 are on the first core, logical processors 2 and 3 are on the second core, and so on.

Consult the host system hardware documentation to verify whether the BIOS includes support for hyperthreading. Then, activate hyperthreading in the system BIOS. Some manufacturers call this option Logical Processor and others call it Enable Hyperthreading.

Use the vSphere Client to ensure that hyperthreading for your host is turned on. To access the hyperthreading option, go to the host's **Summary** tab and select **CPUs** under Hardware.

8-81 CPU Load Balancing

The VMkernel balances processor time to guarantee that the load is spread smoothly across processor cores in the system.

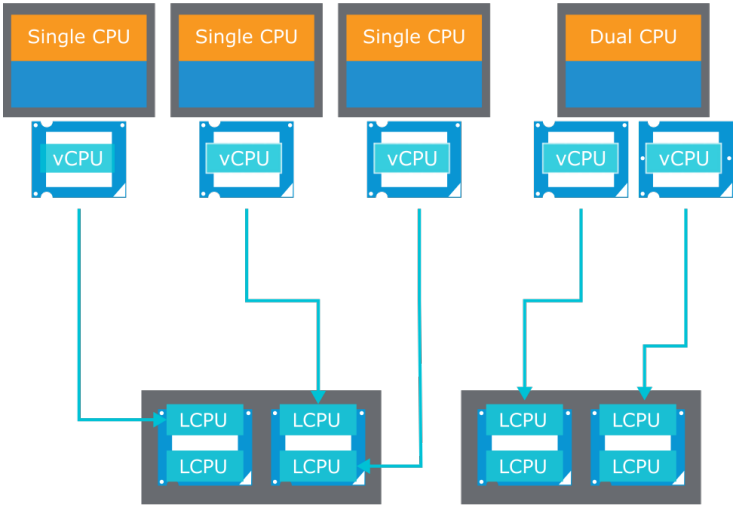

Hyperthreaded Dual-Core
Dual-Socket System

The CPU scheduler can use each logical processor independently to execute VMs, providing capabilities that are similar to traditional symmetric multiprocessing (SMP) systems. The VMkernel intelligently manages processor time to guarantee that the load is spread smoothly across processor cores in the system. Every 2 milliseconds to 40 milliseconds (depending on the socket-core-thread topology), the VMkernel seeks to migrate vCPUs from one logical processor to another to keep the load balanced.

The VMkernel does its best to schedule VMs with multiple vCPUs on two different cores, rather than on two logical processors on the same core. But, if necessary, the VMkernel can map two vCPUs from the same VM to threads on the same core.

If a logical processor has no work, it is put into a halted state. This action frees its execution resources, and the VM running on the other logical processor on the same core can use the full execution resources of the core. Because the VMkernel scheduler accounts for this halt time, a VM running with the full resources of a core is charged more than a VM running on a half core. This approach to processor management ensures that the server does not violate the ESXi resource allocation rules.

8-82 Review of Learner Objectives

- Describe CPU and memory concepts in relation to a virtualized environment
- Recognize techniques for addressing memory resource overcommitment
- Identify additional technologies that improve memory use
- Describe how VMware Virtual SMP works
- Explain how the VMkernel uses hyperthreading

8-83 Lesson 7: Resource Controls

8-84 Learner Objectives

- Assign share values for CPU and memory resources
- Describe how virtual machines compete for resources
- Define CPU and memory reservations and limits

8-85 Reservations, Limits, and Shares

Beyond the CPU and memory configured for a VM, you can apply resource allocation settings to a VM to control the amount of resources granted:

- A reservation specifies the guaranteed minimum allocation for a VM
- A limit specifies an upper bound for CPU or memory that can be allocated to a VM
- A share is a value that specifies the relative priority or importance of a VM's access to a given resource

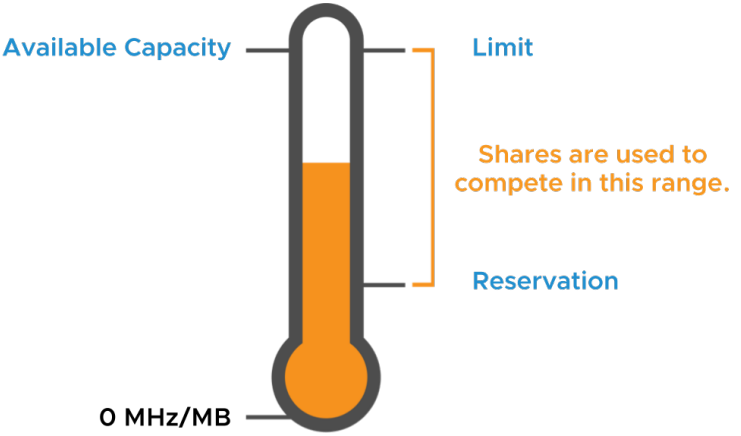

Because VMs simultaneously use the resources of an ESXi host, resource contention can occur.

To manage resources efficiently, vSphere provides mechanisms to allow less, more, or an equal amount of access to a defined resource. vSphere also prevents a VM from consuming large amounts of a resource. vSphere grants a guaranteed amount of a resource to a VM whose performance is not adequate or that requires a certain amount of a resource to run properly.

When host memory or CPU is overcommitted, a VM’s allocation target is somewhere between its specified reservation and specified limit, depending on the VM’s shares and the system load. vSphere uses a share-based allocation algorithm to achieve efficient resource use for all VMs and to guarantee a given resource to the VMs that need it most.

8-86 Resource Allocation Reservations: RAM

RAM reservations:

- Memory reserved to a VM is guaranteed never to swap or balloon.
- If an ESXi host does not have enough unreserved RAM to support a VM with a reservation, the VM does not power on.
- Reservations are measured in MB, GB, or TB. The default is 0 MB.

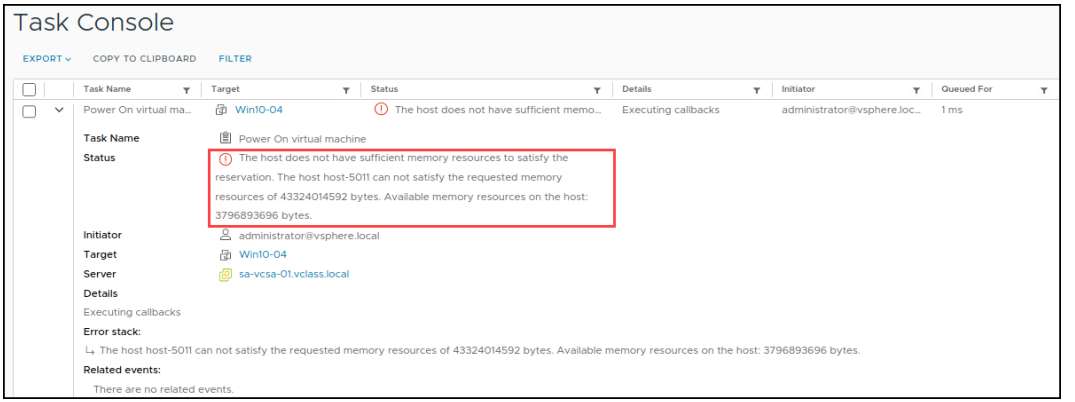

When configuring a memory reservation for a VM, you can specify the VM's configured amount of memory to reserve all the VM's memory. For example, if a VM is configured with 4 GB of memory, you can set a memory reservation of 4 GB for the VM. You might configure such a memory reservation for a critical VM that must maintain a high level of performance.

The ESXi host must have enough RAM to support a VM with a reservation plus a certain amount of overhead memory. VMs require a certain amount of available overhead memory to power on. For example, a VM with 4 GB of memory and two virtual CPUs requires approximately 53 MB of overhead memory. For information on sample overhead memory on VMs, see *vSphere Resource Management* at https://docs.vmware.com/en/VMware-vSphere/index.html.

8-87 Resource Allocation Reservations: CPU

CPU reservations:

- CPU that is reserved for a VM is guaranteed to be immediately scheduled on physical cores. The VM is never placed in a CPU ready state.
- If an ESXi host does not have enough unreserved CPU to support a VM with a reservation, the VM does not power on.
- Reservations are measured in MHz or GHz.
- The default is 0 MHz.

8-88 Resource Allocation Limits

RAM limits:

- VMs never consume more physical RAM than is specified by the memory allocation limit.
- VMs use the VM swap mechanism (`.vswp`) if the guest OS attempts to consume more RAM than is specified by the limit.

CPU limits:

- VMs never consume more physical CPU than is specified by the CPU allocation limit.
- CPU threads are placed in a ready state if the guest OS attempts to schedule threads faster than the limit allows.

Usually, specifying a limit is not necessary.

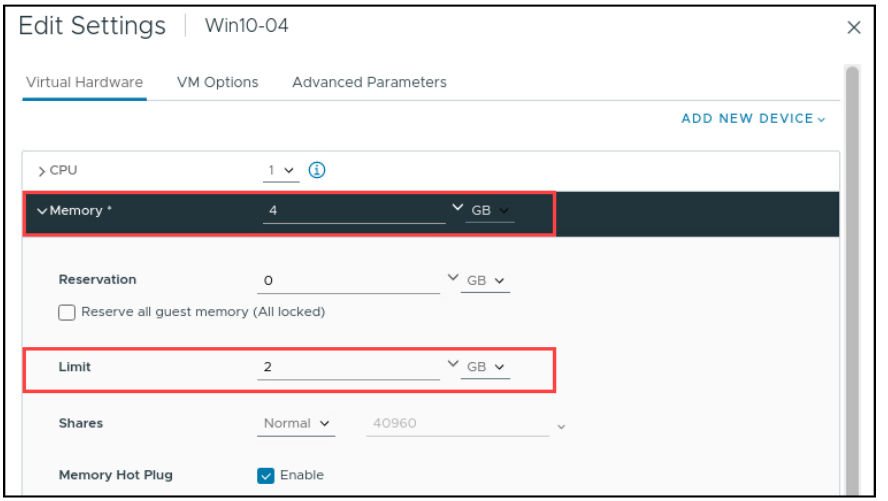

Specifying limits has the following benefits and drawbacks:

- Benefits: Assigning a limit is useful if you start with a few VMs and want to manage user expectations. The performance deteriorates as you add more VMs. You can simulate having fewer resources available by specifying a limit.
- Drawbacks: You might waste idle resources if you specify a limit. The system does not allow VMs to use more resources than the limit, even when the system is underused and idle resources are available. Specify the limit only if you have good reasons for doing so.

8-89 Resource Allocation Shares

Shares define the relative importance of a VM:

- If a VM has twice as many shares of a resource as another VM, the VM is entitled to consume twice as much of that resource when these two VMs compete for resources.
- Share values apply only if an ESXi host experiences contention for a resource.

You can set shares to high, normal, or low. You can also select the custom setting to assign a specific number of shares to each VM.

Setting	CPU Share Values	Memory Share Values
High	2,000 shares per vCPU	20 shares per MB of configured VM memory
Normal	1,000 shares per vCPU	10 shares per MB of configured VM memory
Low	500 shares per vCPU	5 shares per MB of configured VM memory

High, normal, and low settings represent share values with a 4:2:1 ratio, respectively. A custom value of shares assigns a specific number of shares (which expresses a proportional weight) to each VM.

8-90 Resource Shares Example (1)

VMs are resource consumers. The default resource settings that you assign during VM creation work well for most VMs.

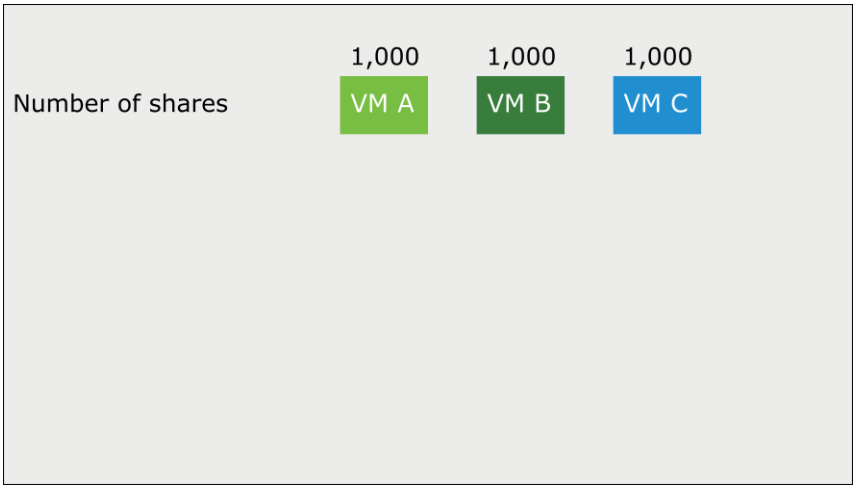

The proportional share mechanism applies to CPU, memory, storage I/O, and network I/O allocation. The mechanism operates only when VMs contend for the same resource.

8-91 Resource Shares Example (2)

You can add shares to a virtual machine while it is running.

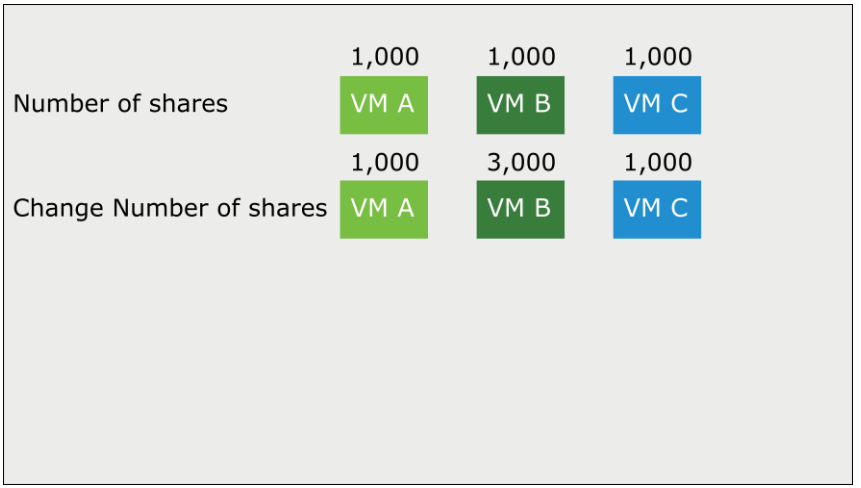

You can add shares to a VM while it is running, and the VM gets more access to that resource (assuming competition for the resource). When you add a VM, it gets shares too. The VM's share amount factors into the total number of shares, but existing VMs are guaranteed not to be starved for the resource.

8-92 Resource Shares Example (3)

Shares guarantee that a VM is given a certain amount of a resource.

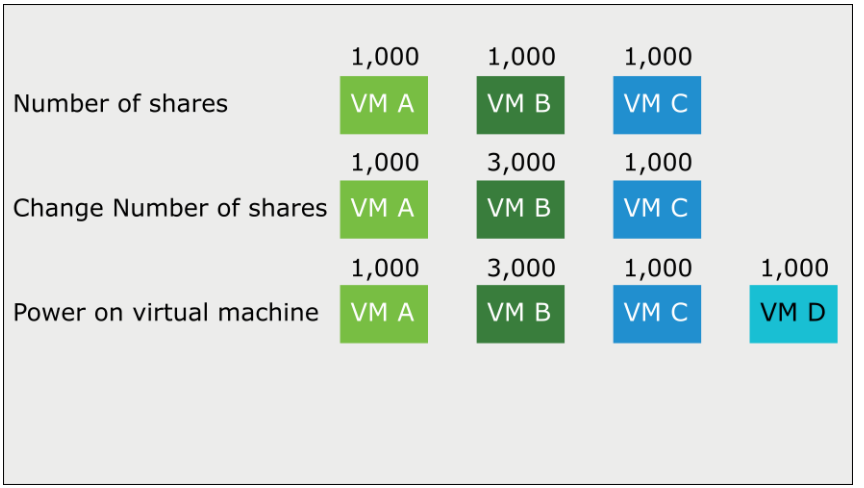

Shares guarantee that a VM is given a certain amount of a resource (CPU, RAM, storage I/O, or network I/O).

For example, consider the third row of VMs on the slide:

- VM D is powered on with 1,000 shares.
- Before VM D was powered on, a total of 5,000 shares were available, but VM D's addition increases the total shares to 6,000.
- The result is that the other VMs' shares decline in value. But VM A is still allocated one-sixth of the resource because it owns one-sixth of the shares.

8-93 Resource Shares Example (4)

When you delete or power off a VM, fewer total shares remain, so the surviving VMs get more access.

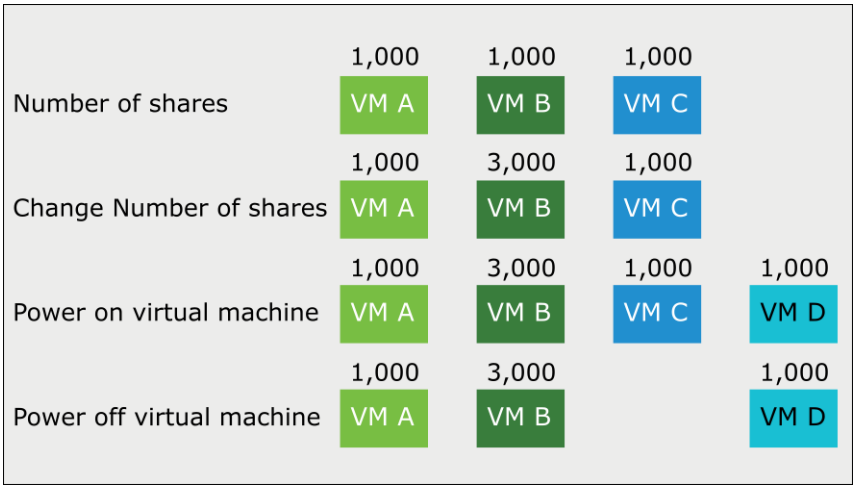

8-94 Defining Resource Allocation Settings for a VM

You can edit a VM's settings to configure CPU and memory resource allocations.

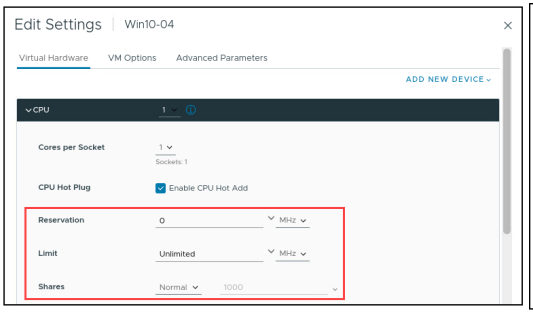

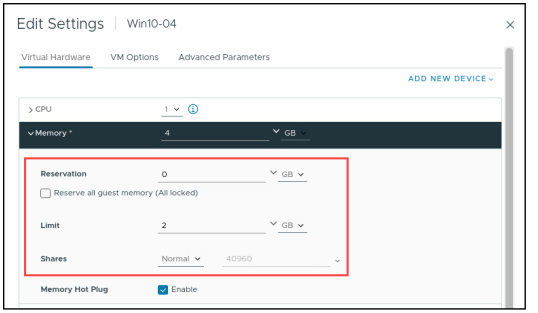

When reserving memory, you can select the **Reserve all guest memory (All locked)** check box. Selecting this check box ensures that all the VM's memory gets reserved, even if you change the total amount of memory for the VM. The memory reservation is immediately readjusted when the VM's memory configuration changes.

8-95 Viewing VM Resource Allocation Settings

You can view reservations, limits, and shares settings for all VMs in a cluster.

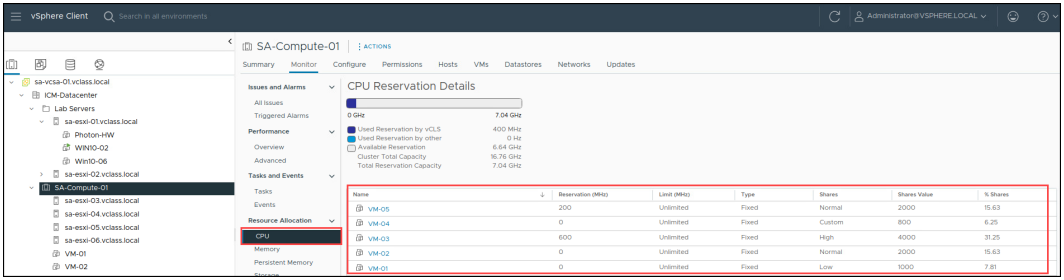

8-96 Lab 23: Controlling VM Resources

Observe the behavior of VMs with different CPU share values:

1. Create CPU Contention
2. Verify the CPU Share Functionality

8-97 Review of Learner Objectives

- Assign share values for CPU and memory resources
- Describe how virtual machines compete for resources
- Define CPU and memory reservations and limits

8-98 Key Points

- Hot migrations use vSphere vMotion, vSphere Storage vMotion, or both.
- You can migrate VMs between vCenter instances, whether they are in the same SSO domain, different SSO domain, or geographically far apart.
- Enhanced vMotion Compatibility prevents vSphere vMotion migrations from failing because of incompatible CPUs or incompatible vSGA GPUs.
- You can use VM snapshots to preserve the state of the VM so that you can return repeatedly to the same state.
- You can apply reservations, limits, and shares against a VM to control the amount of CPU and memory resources granted.

Questions?

Module 9

Deploying and Configuring vSphere Clusters

9-2 Importance

Most organizations rely on computer-based services such as email, databases, and web-based applications. The failure of these services can mean lost productivity and revenue.

By understanding and using vSphere HA, you can configure highly available, computer-based services, which are important for an organization to remain competitive in contemporary business environments. And by developing skills in using vSphere DRS, you can improve service levels by guaranteeing appropriate resources to virtual machines.

9-3 Module Lessons

1. vSphere Clusters Overview
2. vSphere Distributed Resource Scheduler
3. Introduction to vSphere High Availability
4. vSphere High Availability Architecture
5. Configuring vSphere High Availability

9-4 Lesson 1: vSphere Clusters Overview

9-5 Learner Objectives

- Create a vSphere cluster
- Recognize cluster options that you can configure with Cluster Quickstart
- View information about a vSphere cluster

9-6 About vSphere Clusters

A cluster is used in vSphere to share physical resources between a group of ESXi hosts. vCenter manages cluster resources as a single pool of resources.

You can create one or more clusters based on the purpose each cluster must fulfill, for example:

- Management
- Production
- Compute

A cluster can contain up to 96 ESXi hosts.

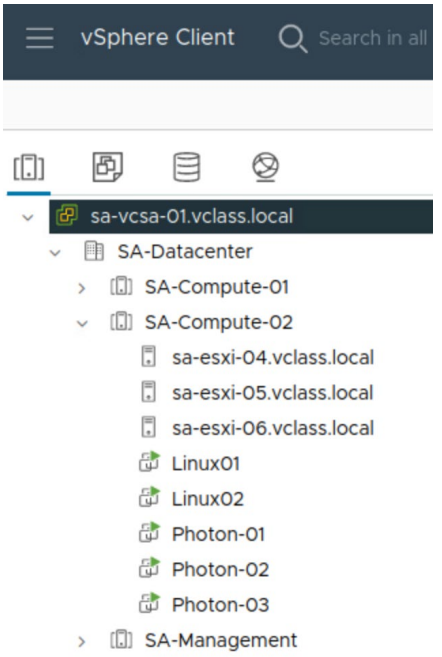

In this screenshot, five clusters are shown: SA-Compute-01, SA-Compute-02, SA-Management, SB-Development, and SB-Management.

To support 96 hosts per cluster, vSphere 7 Update 1 or later is required.

9-7 Creating a vSphere Cluster

You can create a cluster by giving it a name and selecting the relevant cluster services.

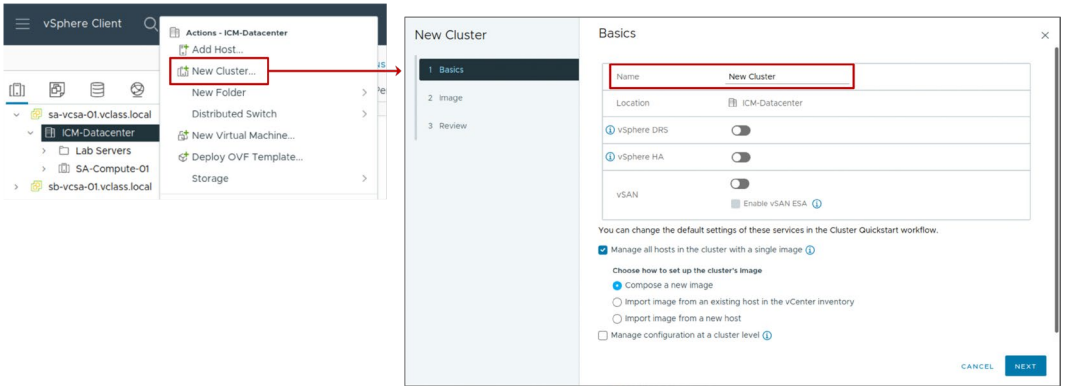

You can activate the following services in a vSphere cluster:

- vSphere HA, for high availability
- vSphere DRS, for VM placement and load balancing
- vSAN, for software-defined storage

You can also manage host updates using images. With vSphere Lifecycle Manager, you can update all hosts in the cluster collectively, using a specified ESXi image.

9-8 About Cluster Quickstart

After you create a cluster, you can use the Cluster Quickstart workflow to configure the cluster.

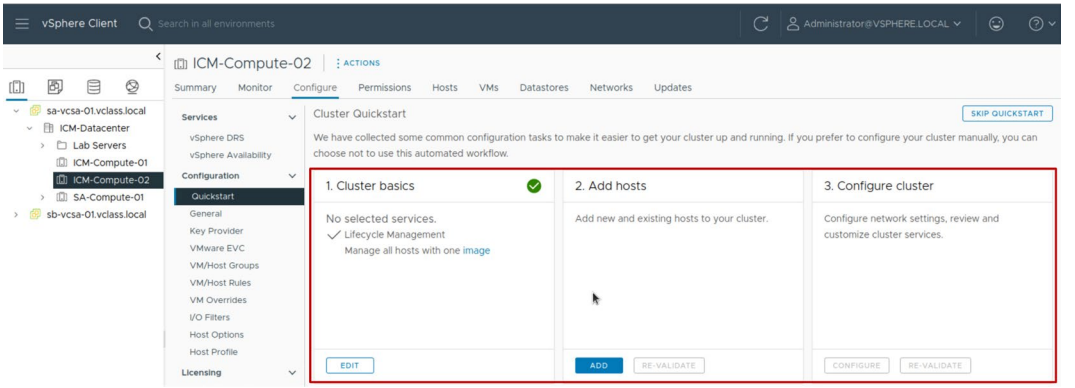

The Cluster Quickstart workflow guides you through the deployment process for clusters. It covers every aspect of the initial configuration, such as host, network, and vSphere settings. With Cluster Quickstart, you can also add additional hosts to a cluster as part of the ongoing expansion of clusters.

Cluster Quickstart reduces the time it takes to configure a cluster.

The workflow includes the following tasks:

- Setting up services such as vSphere HA, vSphere DRS, and vSAN
- Verifying hardware and software compatibility
- Deploying vSphere Distributed Switches
- Configuring network settings for vSphere vMotion and vSAN
- Creating a vSAN stretched cluster or vSAN fault domains
- Ensuring consistent NTP configuration across the cluster

For more information about creating clusters, see *vCenter Server and Host Management* at https://docs.vmware.com/en/VMware-vSphere/index.html.

9-9 Cluster Quickstart: Activating Services

The first step in the Cluster Quickstart workflow is to check that the correct cluster services are selected.

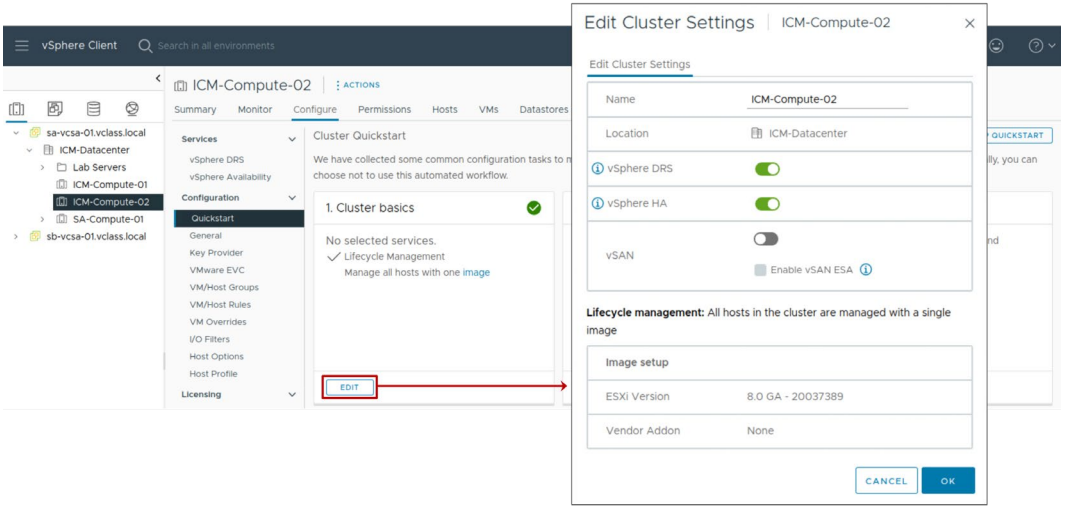

The Cluster basics pane lets you configure the following information:

- Edit the cluster's name.
- Activate or deactivate the vSphere DRS, vSphere HA, and vSAN services.
- Select the image for vSphere Lifecycle Manager to use to manage hosts in the cluster.

For more information about creating clusters, see *vCenter Server and Host Management* at https://docs.vmware.com/en/VMware-vSphere/index.html.

9-10 Cluster Quickstart: Adding Hosts

The second step in the Cluster Quickstart workflow is to add new or existing hosts to the cluster.

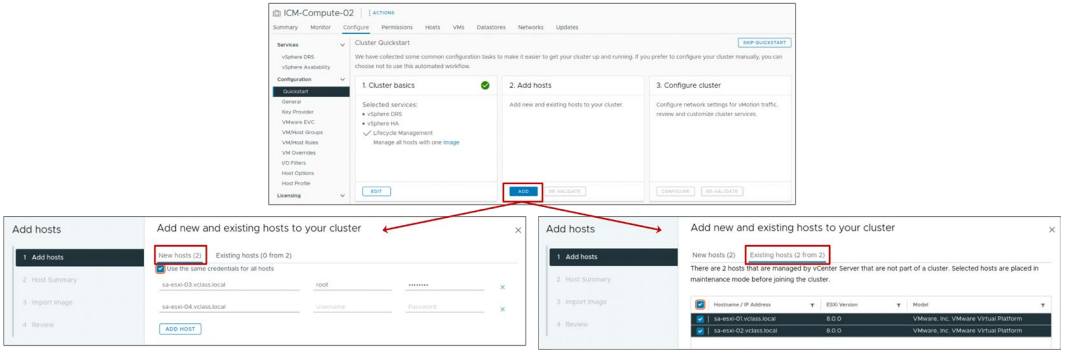

The Add hosts pane lets you add ESXi hosts to the cluster. Select **New hosts** to add hosts to the inventory and cluster simultaneously. For each host that you add to the cluster, you enter the IP address or FQDN of each host. Select **Existing hosts** to add hosts to the cluster that are already present in the inventory.

After hosts are added, the workflow shows the total number of hosts that are present in the cluster and provides a health check validation for those hosts.

9-11 Cluster Quickstart: Configuring the Cluster

The third step is to configure the host networking settings and customize cluster services.

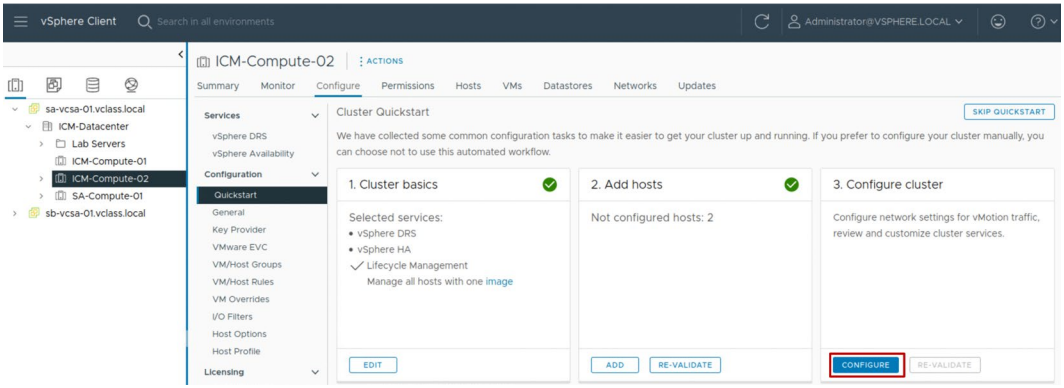

If you click **SKIP QUICKSTART**, you can manually configure your cluster by using the menus in the vSphere Client. By skipping the Cluster Quickstart workflow, you cannot restore it for the current cluster. Any hosts added to this cluster must be configured manually.

9-12 Configuring a Cluster: Distributed Switches

To configure a cluster, you can:

- Select up to three distributed switches.
- Select a network for vSphere vMotion.
- Select at least one physical adapter.

You can also configure networking settings later.

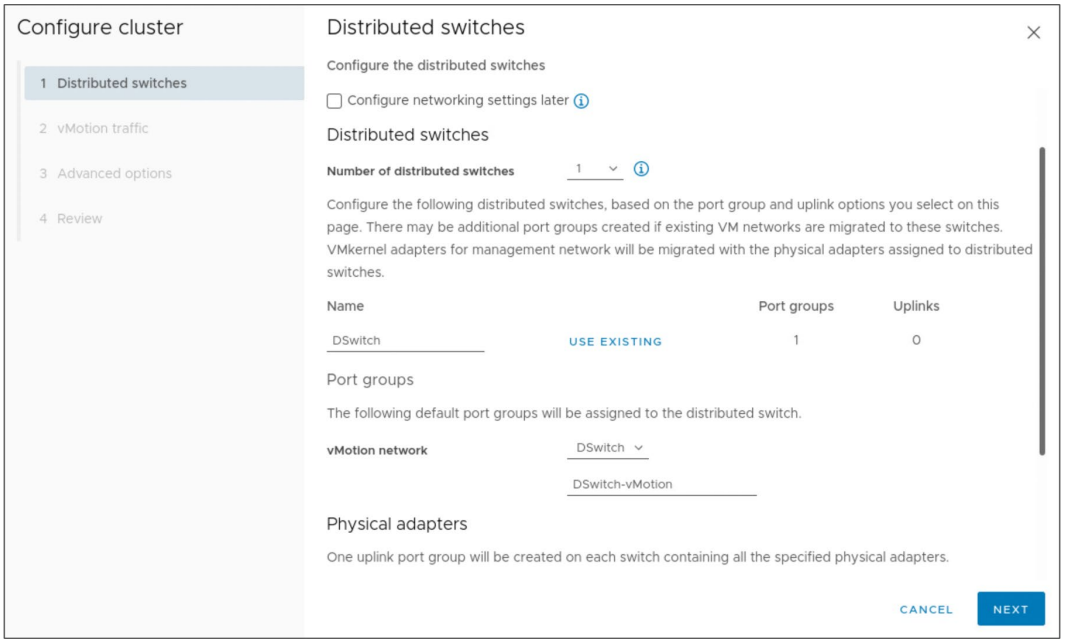

You can select the **Configure networking settings later** check box to configure the default settings only for the cluster services and to hide all options that are related to host networking. But, by selecting this option, you cannot perform the networking configuration by using the Configure cluster wizard.

9-13 Configuring a Cluster: vSAN and vMotion Traffic

If you selected vSphere DRS, you are prompted to enter the IP address information for the vSphere vMotion network.

If you selected vSAN, you are prompted to enter the IP information for the vSAN network.

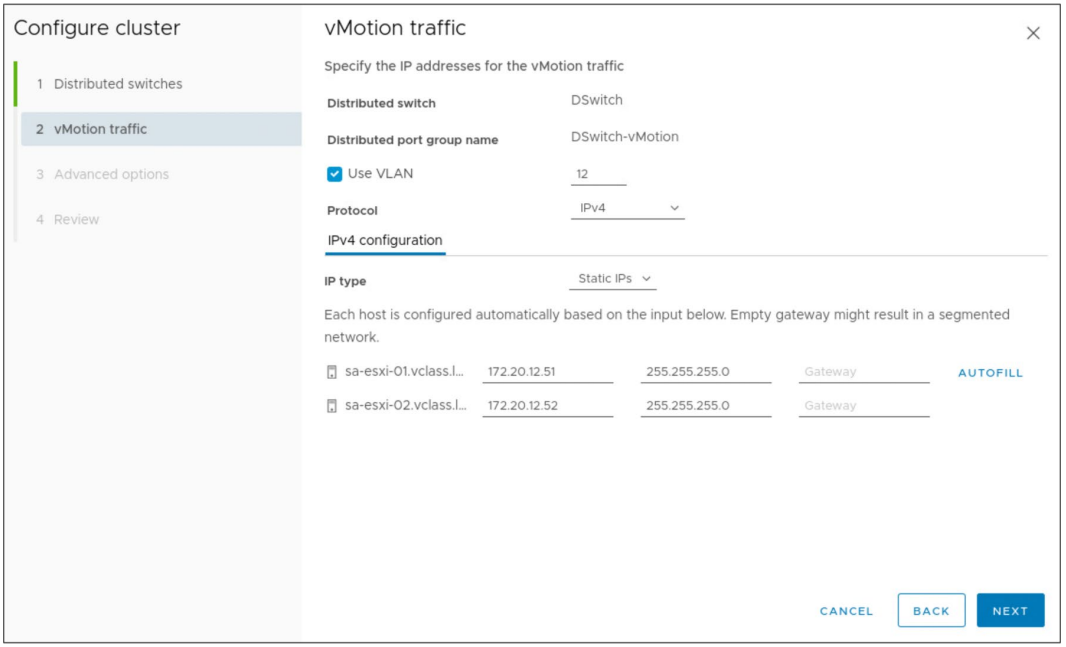

On the vMotion traffic page, if the IP information for ESXi hosts is similar, you can enter the information for the first host and click **AUTOFILL** to complete the fields for the following hosts.

9-14 Configuring a Cluster: Advanced Options

You get different settings depending on the cluster services that are enabled:

- High Availability (optional)
- Distributed Resource Scheduler (optional)
- Host Options
- Enhanced vMotion Compatibility

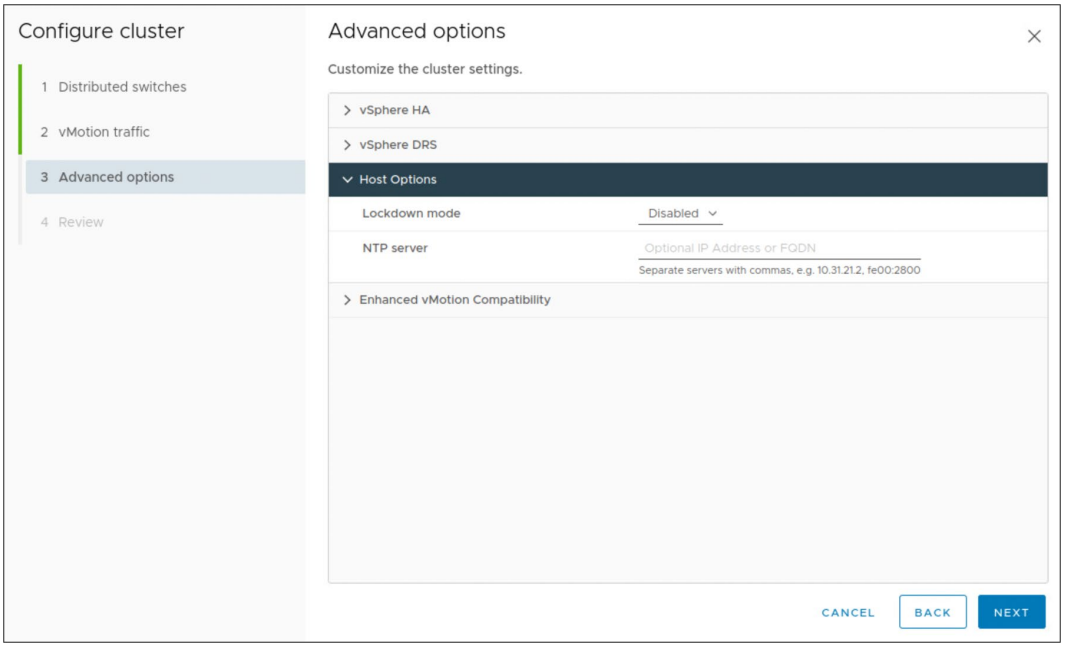

9-15 Viewing Cluster Summary Information

The **Summary** tab provides a quick view of information about a cluster's resources and its consumers.

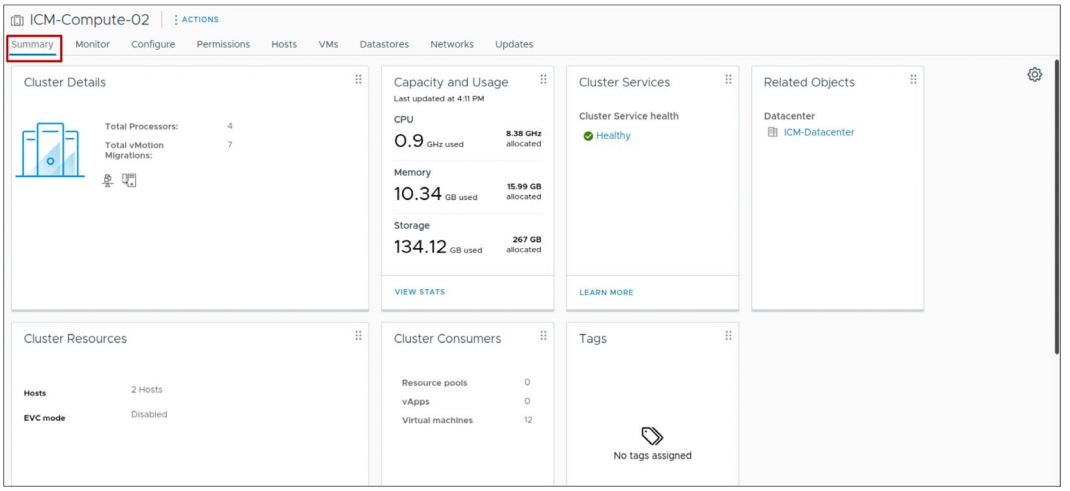

9-16 Monitoring Cluster Resources

You can view CPU and memory allocation details.

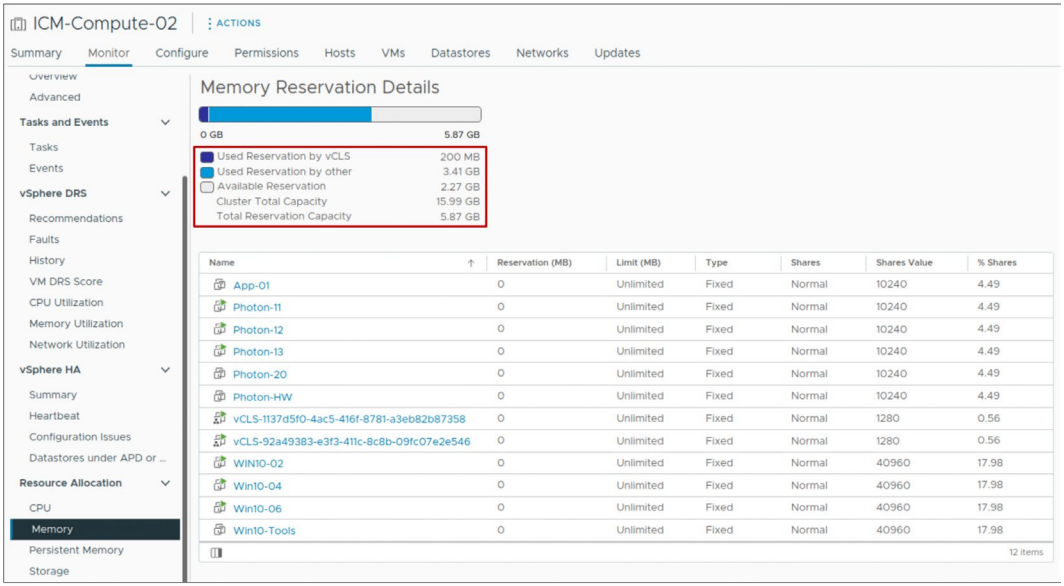

Memory statistics for the cluster

You can view a report of total cluster CPU, memory, memory overhead, storage capacity, the capacity reserved by VMs, and how much capacity remains available.

9-17 vSphere Cluster Services VMs

Up to three vSphere Cluster Service VMs are present in each vSphere cluster.

These VMs are required for maintaining the health and availability of cluster services such as vSphere DRS and vSphere HA.

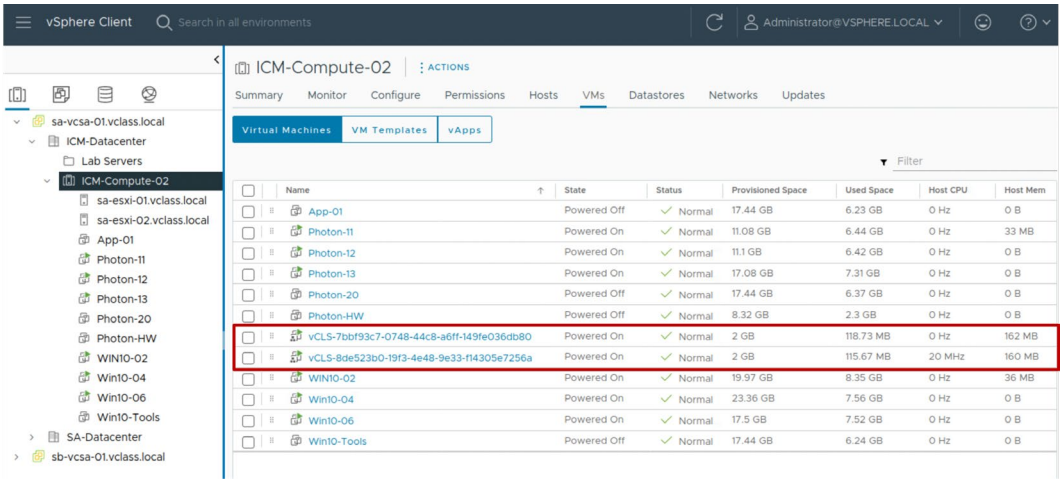

vSphere Cluster Service VMs are deployed to a cluster at cluster creation and after hosts are added to the cluster. A vSphere Cluster Service VM is deployed from an OVA with a minimal installed profile of Photon OS. vSphere Cluster Services Manager is a vCenter service that manages the resources, power state, and availability of these VMs. Any impact on the power state or resources of these VMs might degrade the health of vSphere Cluster Services and cause vSphere DRS to stop working in the cluster.

In the vSphere Client, vSphere Cluster Service VMs are not visible in the Hosts and Clusters inventory view. However, you can view these VMs from the cluster's **VMs** tab. You can also view these VMs from the VMs and Templates inventory view.

The vSphere cluster shows an alert message if healthy vSphere Cluster Service VMs are not available in the cluster.

If vSphere Cluster Service VMs are manually powered off, these VMs are automatically powered on by vCenter.

9-18 Review of Learner Objectives

- Create a vSphere cluster
- Recognize cluster options that you can configure with Cluster Quickstart
- View information about a vSphere cluster

9-19 Lesson 2: vSphere Distributed Resource Scheduler

9-20 Learner Objectives

- Describe the functions of a vSphere DRS cluster
- Explain how vSphere DRS determines VM placement on hosts in the cluster
- Recognize use cases for vSphere DRS settings
- Configure vSphere DRS in a cluster
- Monitor a vSphere DRS cluster

9-21 About vSphere Distributed Resource Scheduler

vSphere Distributed Resource Scheduler (DRS) helps to improve resource allocation across all hosts in a cluster.

vSphere DRS use cases:

- Initial placement when a VM is powered on
- Load balancing
- Migrating VMs when an ESXi host is placed in maintenance mode

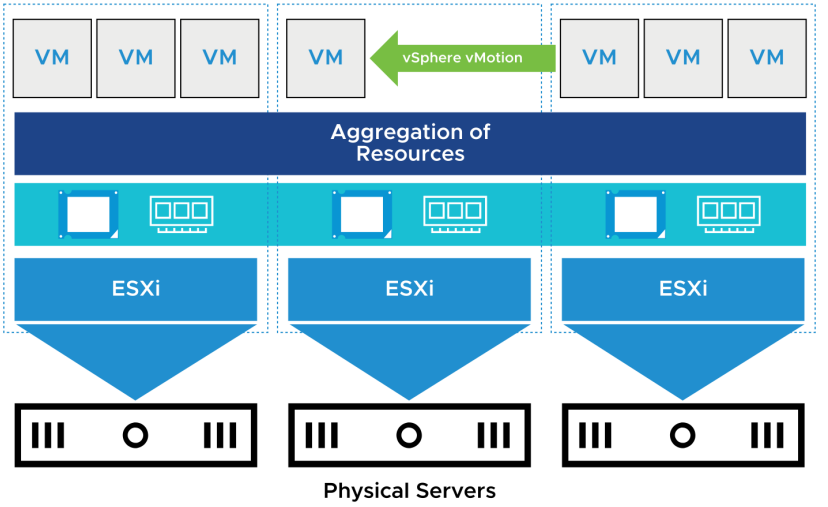

vSphere DRS aggregates computing capacity across a collection of servers into logical resource pools.

When you power on a VM in the cluster, vSphere DRS either places the VM on the host with the most available resources when DRS is in Fully Automated mode, or makes a resource allocation recommendations when DRS is in Partially Automated or Manual mode.

DRS attempts to improve resource use across the cluster by performing automatic migrations of VMs (using vSphere vMotion) or by providing a recommendation for VM migrations.

Before an ESXi host enters maintenance mode, VMs running on the host must be migrated to another host either manually or automatically by DRS, shut down, or suspended.

9-22 vSphere DRS: VM Focused

vSphere DRS is VM focused:

- While the VM is powered on, vSphere DRS operates on an individual VM basis by ensuring that each VM's resource requirements are met.
- vSphere DRS calculates a score for each VM and gives recommendations (or migrates VMs) for meeting VM's resource requirements.

The DRS algorithm recommends where individual VMs should be moved for maximum performance. If the cluster is in fully automated mode, DRS executes the recommendations and migrates VMs to their optimal host based on the underlying calculations performed every minute.

9-23 About the VM DRS Score

The VM DRS score tracks how well a VM's resource requirements are met.

- Scores closer to 0% indicate severe resource contention.
- Scores closer to 100% indicate mild to no resource contention.

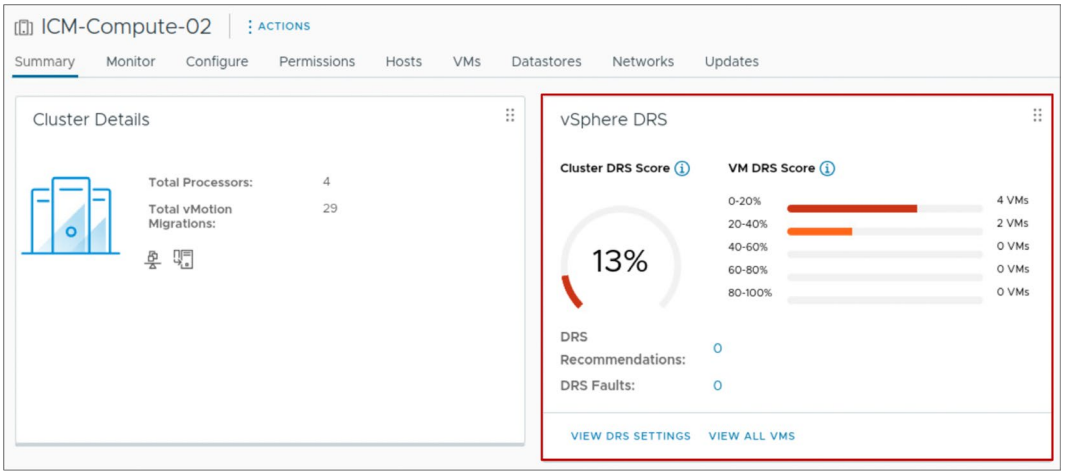

A VM DRS score is computed from an individual VM's CPU, memory, and network metrics. DRS uses these metrics to gauge the goodness or wellness of the VM.

In vSphere 7 and later, the DRS algorithm runs every minute. The Cluster DRS Score is the last result of DRS running and is filed into one of five buckets. These buckets are simply 20 percent ranges: 0-20, 20-40, 40-60, 60-80 and 80-100 percent over the sample period.

9-24 VM DRS Score List

The cluster's **Monitor** tab lists the VM DRS Score and more detailed metrics for all the VMs in the cluster.

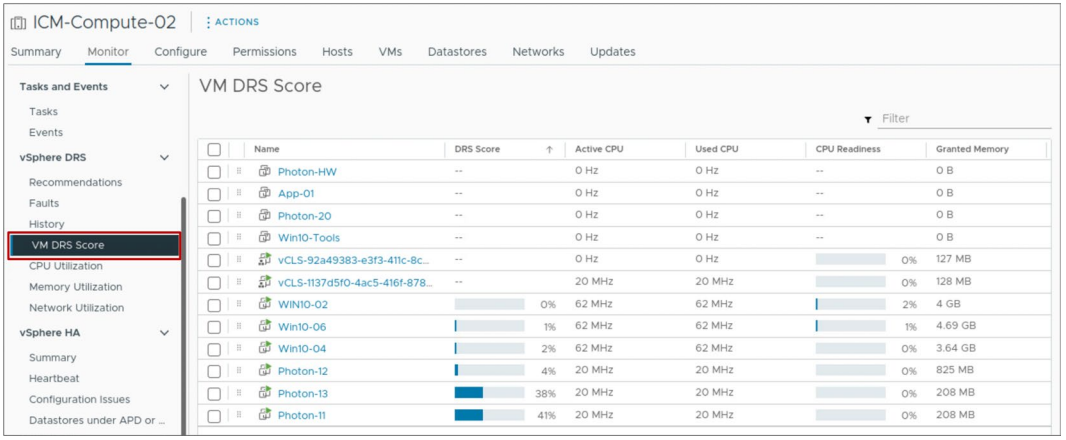

The VM DRS Score page shows the following values for VMs that are powered on:

- DRS Score
- Active CPU
- Used CPU
- CPU Readiness
- Granted Memory
- Swapped Memory
- Ballooned Memory

9-25 vSphere DRS Cluster Requirements

ESXi hosts that are added to a vSphere DRS cluster must meet certain requirements to use cluster features successfully:

- To use vSphere DRS for load balancing, the hosts in your cluster must be part of a vSphere vMotion network:
  - If the hosts are not part of a vSphere vMotion network, vSphere DRS can still make initial placement recommendations.
  - vSphere DRS works best if the VMs meet vSphere vMotion requirements.
- Configure all managed hosts to use shared storage.

9-26 vSphere DRS Settings: Automation Level

You can configure the automation level for the initial placement of VMs and for dynamic balancing while VMs are running.

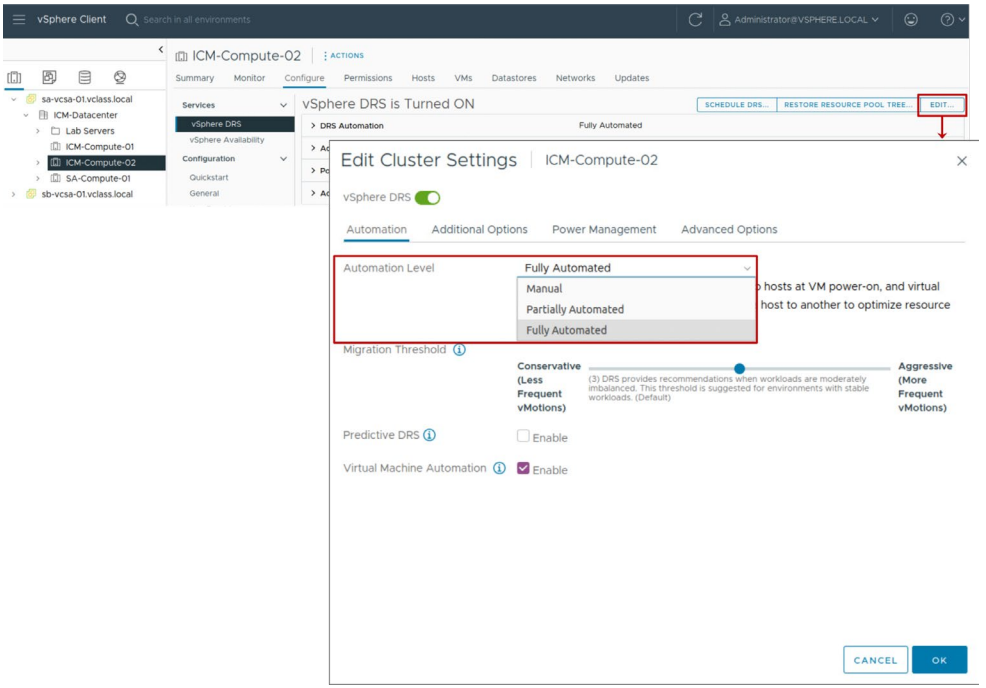

The automation level determines whether vSphere DRS makes migration recommendations or automatically places VMs on hosts. vSphere DRS makes placement decisions when a VM powers on and when VMs must be rebalanced across hosts in the cluster.

The following automation levels are available:

- **Manual**: When you power on a VM, vSphere DRS displays a list of recommended hosts on which to place the VM. To improve the cluster's overall DRS score, vSphere DRS displays recommendations for manual VM migration.
- **Partially automated**: When you power on a VM, vSphere DRS places it on the best-suited host. To improve the cluster's overall DRS score, vSphere DRS displays recommendations for manual VM migration.
- **Fully automated**: When you power on a VM, vSphere DRS places it on the best-suited host. To improve the cluster's overall DRS score, vSphere DRS automatically migrates VMs between hosts.

9-27 vSphere DRS Settings: Migration Threshold

The migration threshold determines how aggressively vSphere DRS selects to migrate VMs.

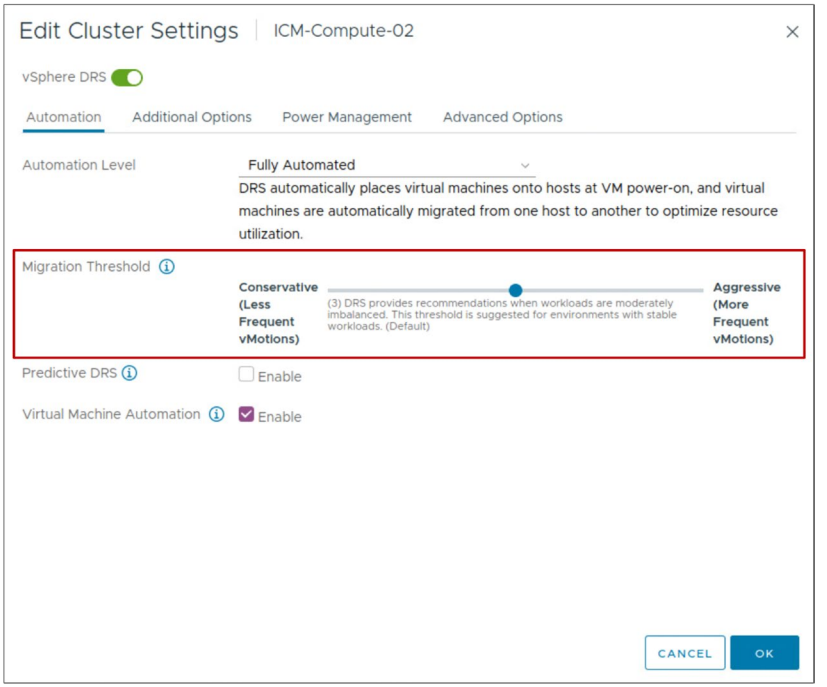

The following migration threshold settings are available:

- Level 1 (Conservative): Applies only priority 1 recommendations. vCenter applies only recommendations that must be taken to satisfy VMs' requirements, such as affinity rules and host maintenance.
- Level 2: Apply priority 1 and priority 2 recommendations. vCenter applies recommendations that promise a significant improvement to the cluster's overall DRS score.
- Level 3 (default): Apply priority 1, priority 2, and priority 3 recommendations. vCenter applies recommendations that promise at least a good improvement to the cluster's overall DRS score.

- Level 4: Apply priority 1, priority 2, priority 3, and priority 4 recommendations. vCenter applies recommendations that promise even a moderate improvement to the cluster's overall DRS score.
- Level 5 (Aggressive): Apply all recommendations. vCenter applies recommendations that promise even a slight improvement to the cluster's overall DRS score.

9-28 vSphere DRS Settings: Predictive DRS

Predictive DRS is used to predict future demand and determine when and where high resource utilization occurs.

To make predictive decisions, the vSphere DRS data collector retrieves the following data:

- Resource usage statistics from ESXi hosts
- Predicted usage statistics from the VMware Aria Operations server

Goals of Predictive DRS:

- Move VMs before their DRS score drops.
- Perform migrations before host resources are in contention.

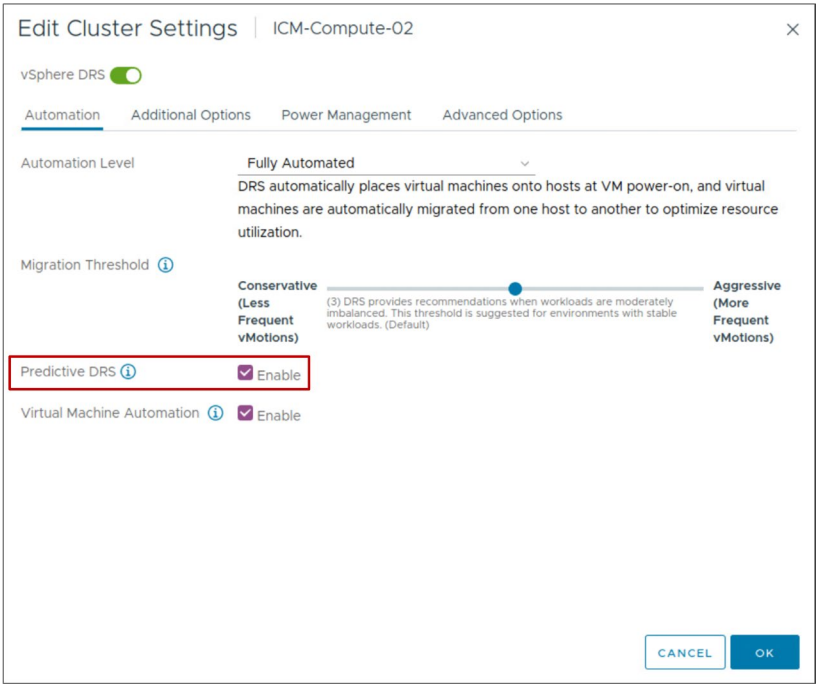

Predicted usage statistics always take precedence over current usage statistics.

9-29 Viewing vSphere DRS Settings

When you click **VIEW DRS SETTINGS**, the main vSphere DRS parameters and their current values are shown.

vSphere DRS settings include:

- Automation level
- Migration threshold

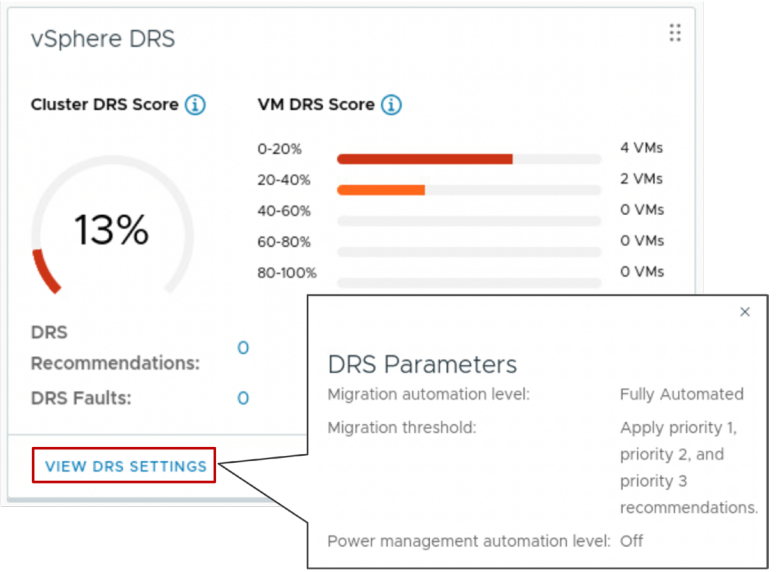

To view the vSphere DRS pane, go to the cluster's **Summary** tab.

9-30 vSphere DRS Settings: VM-Level Automation

You can customize the automation level for individual VMs in a cluster to override the automation level set on the entire cluster.

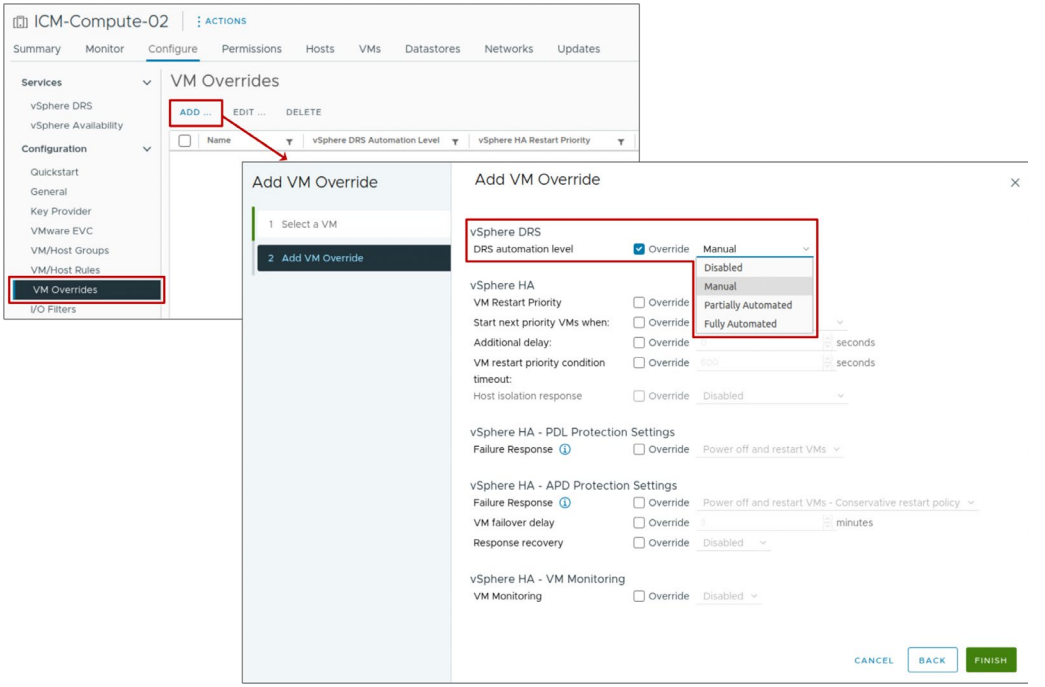

By setting the automation level for individual VMs, you can fine-tune automation to suit your needs. For example, you might have a VM that is especially critical to your business. You want more control over its placement, so you set its automation level to **Manual**.

If a VM's automation level is set to Disabled, vSphere DRS does not migrate the VM, nor does it provide initial placement or migration recommendations.

Select the automation level based on your environment and level of comfort.

For example, if you are new to vSphere DRS clusters, you might select **Partially Automated** because you want control over the movement of VMs.

When you are comfortable with what vSphere DRS does and how it works, you might set the automation level to **Fully Automated**.

You can set the automation level to **Manual** on VMs over which you want more control, such as your business-critical VMs.

9-31 vSphere DRS Settings: VM Swap File Location

ESXi hosts can be configured to place VM swap files on a local datastore.

If vSphere DRS is configured, you should place the VM swap files on a shared datastore.

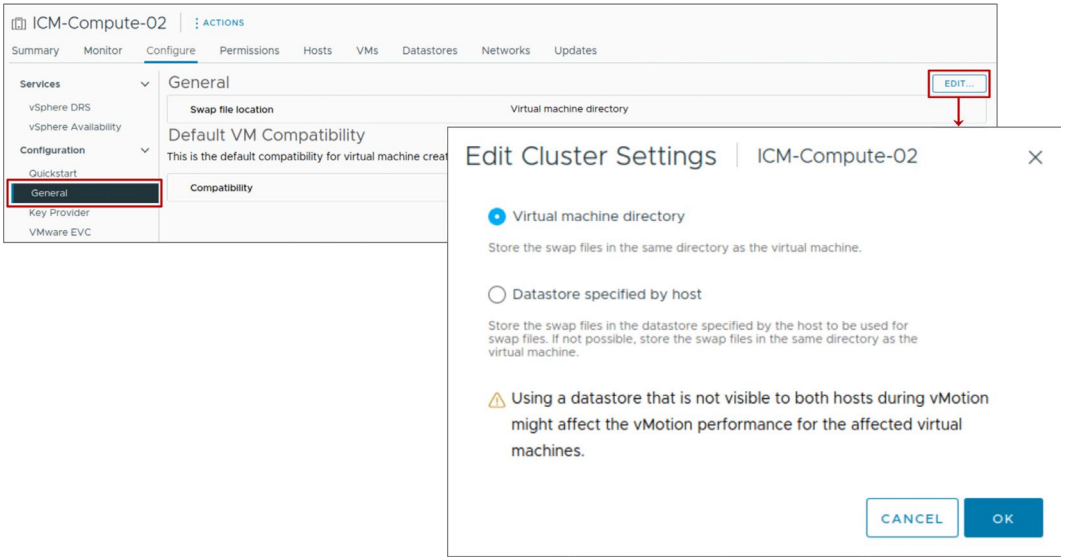

On a VMFS or NFS datastore, the VM's swap space is a set of two swap files. On a vSAN or vSphere Virtual Volumes datastore, the swap files are created as a separate objects.

Two swap files are created by the ESXi host when a VM is powered on. If these files cannot be created, the VM cannot power on.
By default, swap files for a VM are on a datastore in the folder containing the other VM files. Instead of accepting the default, you can also use the following options:

- Use per-VM configuration options to change the datastore to another shared storage location.
- Use host-local swap, which allows you to specify a datastore stored locally on the host. You can swap at a per-host level. However, it can lead to a slight degradation in performance for vSphere vMotion because pages swapped to a local swap file on the source host must be transferred across the network to the destination host. Currently, vSAN and vSphere Virtual Volumes datastores cannot be specified for host-local swap.

9-32 vSphere DRS Settings: VM Affinity

vSphere DRS virtual machine affinity rules specify that selected VMs be placed either on the same host or on separate hosts:

- Affinity rules: For VMs that communicate heavily with one another
- Anti-affinity rules: For VMs where load balancing or high availability is desired.

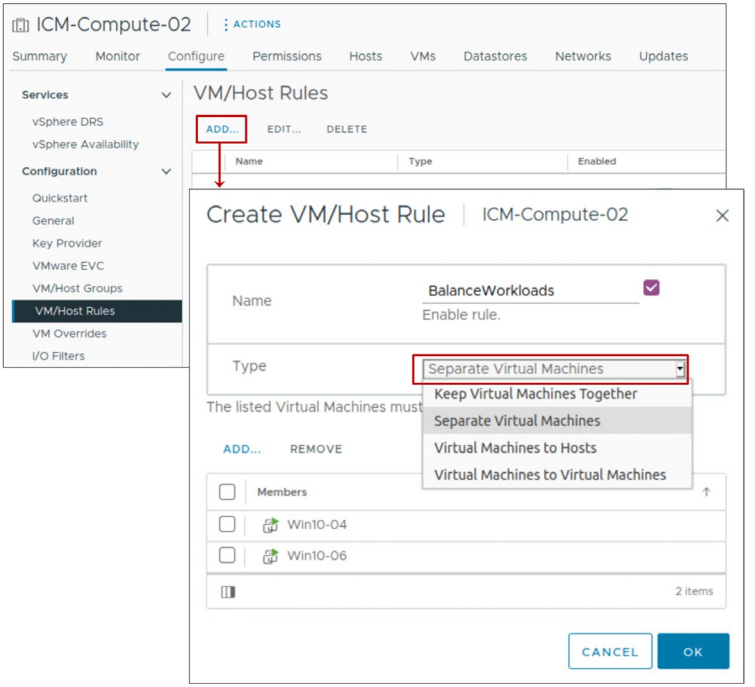

After a vSphere DRS cluster is created, you can edit its properties to create rules that specify affinity. The following types of rules can be created:

- Affinity rules: vSphere DRS keeps certain VMs together on the same host (for example, for performance reasons).
- Anti-affinity rules: vSphere DRS ensures that certain VMs are placed on different hosts (for example, for availability reasons).

If two rules conflict, you are prevented from activating both.

When you add or edit a rule, and the cluster is immediately in violation of that rule, the cluster continues to operate and tries to correct the violation.

For vSphere DRS clusters that have an automation level of manual or partially automated, migration recommendations are based on both rule fulfillment and load balancing.

9-33 vSphere DRS Settings: DRS Groups

VM groups and host groups are used in defining VM-Host affinity rules.

The VM-Host affinity rule specifies whether a VM can or cannot be run on a host.

Types of groups:

- VM group: One or more VMs
- Host group: One or more ESXi hosts

A VM can belong to multiple VM groups.

A host can belong to multiple host groups.

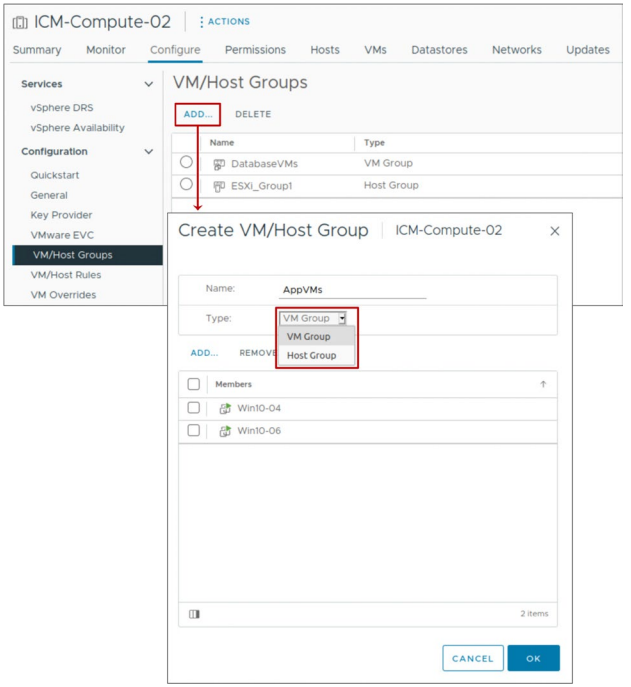

For ease of administration, virtual machines can be placed in VM or host groups. You can create one or more VM groups in a vSphere DRS cluster, each consisting of one or more VMs. A host group consists of one or more ESXi hosts.

The main use of VM groups and host groups is to help in defining the VM-Host affinity rules.

9-34 vSphere DRS Settings: VM-Host Affinity Rules

A VM-Host affinity rule:

- Defines an affinity (or anti-affinity) relationship between a VM group and a host group
- Is either a required rule or a preferential rule

Rule options:

- Must run on hosts in group
- Should run on hosts in group
- Must not run on hosts in group
- Should not run on hosts in group

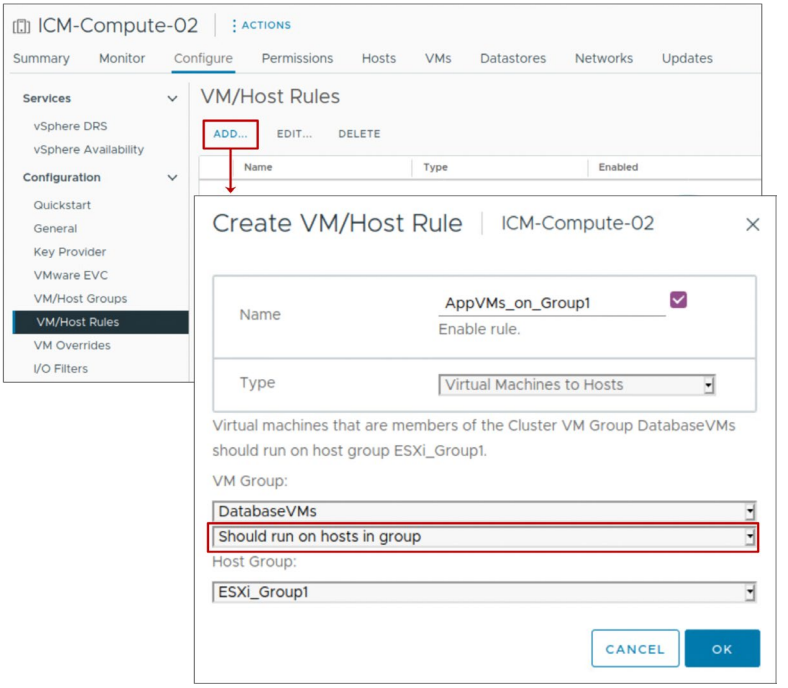

A VM-Host affinity or anti-affinity rule specifies whether the members of a selected VM group can run on the members of a specific host group.

Unlike an affinity rule for VMs, which specifies affinity (or anti-affinity) between individual VMs, a VM-Host affinity rule specifies an affinity relationship between a group of VMs and a group of hosts.

Because VM-Host affinity rules are cluster-based, the VMs and hosts that are included in a rule must all reside in the same cluster. If a VM is removed from the cluster, the VM loses its membership from all VM groups, even if it is later returned to the cluster.

9-35 VM-Host Affinity Preferential Rules

A preferential rule is softly enforced and can be violated if necessary.

Example: Separate VMs on different blade systems for improved performance.

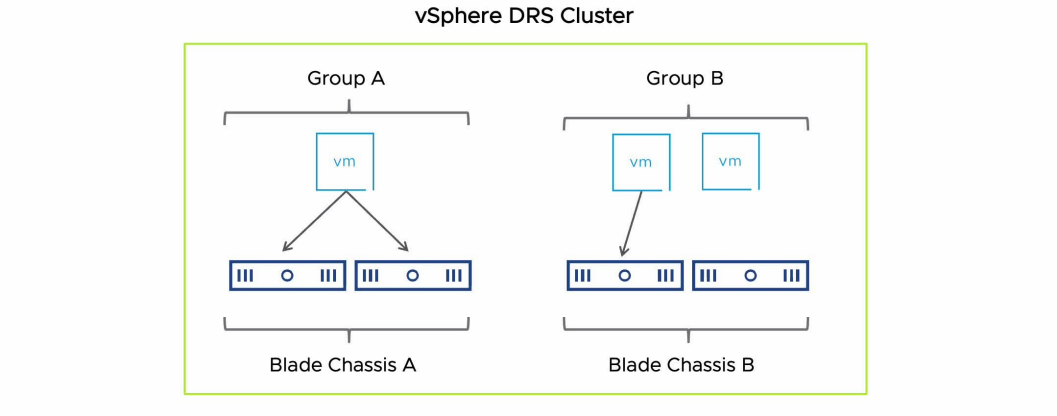

To play the animation, go to https://vmware.bravais.com/s/2V3GkAYpnmAK6d43Oheh

Preferential rules can be violated to allow the proper functioning of vSphere DRS, vSphere HA, and VMware vSphere DPM.

On the slide, Group A and Group B are VM groups. Blade Chassis A and Blade Chassis B are host groups. The goal is to force the VMs in Group A to run on the hosts in Blade Chassis A and to force the VMs in Group B to run on the hosts in Blade Chassis B. If the hosts fail, vSphere HA restarts the VMs on the other hosts in the cluster. vSphere DRS moves the VMs to the other hosts in the cluster if the hosts are put into maintenance mode or hosts are needed to satisfy VMs' resource requirements.

9-36 VM-Host Affinity Required Rules

A required rule is strictly enforced and can never be violated.

Example: Enforce host-based ISV licensing.

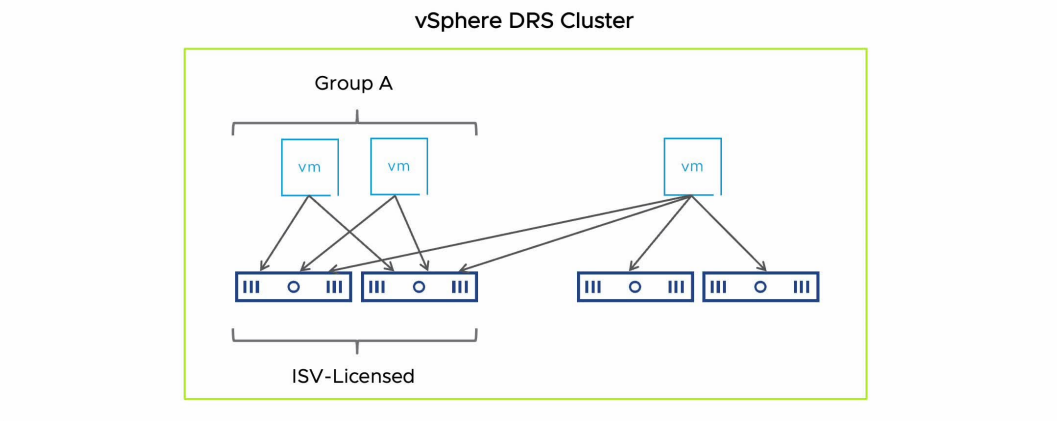

To play the animation, go to https://vmware.bravais.com/s/4UeCSZE8oK4Tc3SJ8kiM

A VM-Host affinity rule that is required, instead of preferential, can be used when the software running in your VMs has licensing restrictions. You can enforce this rule when the software running in your VMs has licensing restrictions. You can place such VMs in a VM group. Then you can create a rule that requires the VMs to run on a host group, which contains hosts with the required licenses.

When you create a VM-Host affinity rule that is based on the licensing or hardware requirements of the software running in your VMs, you are responsible for ensuring that the groups are properly set up. The rule does not monitor the software running in the VMs. Nor does it know which third-party licenses are in place on which ESXi hosts.

On the slide, Group A is a VM group. You can force Group A to run on hosts in the ISV-Licensed group to ensure that the VMs in Group A run on hosts that have the required licenses. But if the hosts in the ISV-Licensed group fail, vSphere HA cannot restart the VMs in Group A on hosts that are not in the group. If the hosts in the ISV-Licensed group are put into maintenance mode or become overused, vSphere DRS cannot move the VMs in Group A to hosts that are not in the group.

9-37 Viewing vSphere DRS Cluster Resource Utilization

From the cluster's **Monitor** tab, you can view CPU, memory, and network utilization per host.

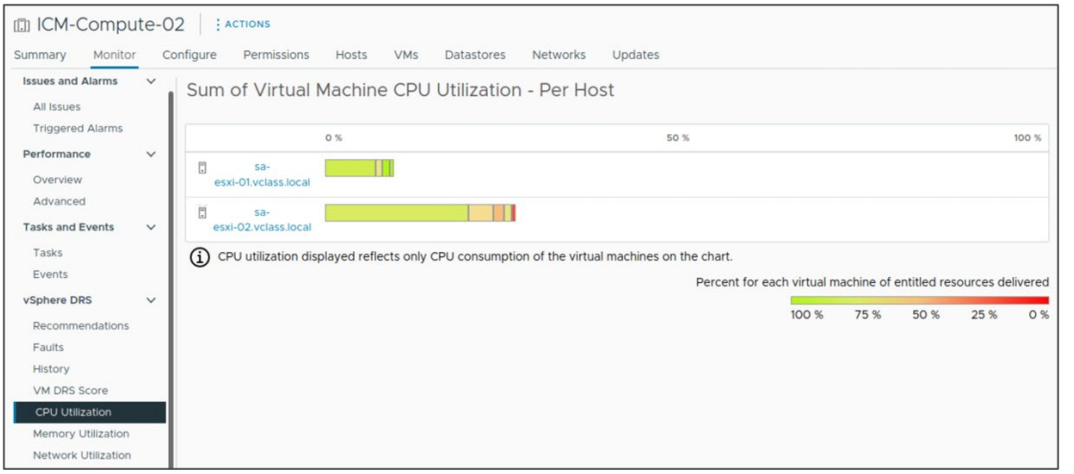

The CPU Utilization chart shows all the hosts in the cluster and how their CPU resources are allocated to each VM. For CPU usage, a colored box represents the VM information. If you point to the colored box, the VM's CPU usage information appears. If the VM is receiving the resources that it is entitled to, the box is green. Green means that 100 percent of the VM's entitled resources are delivered. If the box is not green (for example, entitled resources are 80 percent or less) for an extended time, you might want to investigate what is causing this shortfall (for example, unapplied recommendations).

The Memory Utilization chart shows all the hosts in the cluster and how their memory resources are allocated to each VM. For memory usage, the VM boxes are not color-coded because the relationship between consumed memory and entitlement is often not easily categorized.

The Network Utilization chart displays network data, which reflects all traffic across physical network interfaces on the host.

9-38 Viewing vSphere DRS Recommendations

Select **Recommendations** to display information about the vSphere DRS recommendations made for the cluster.

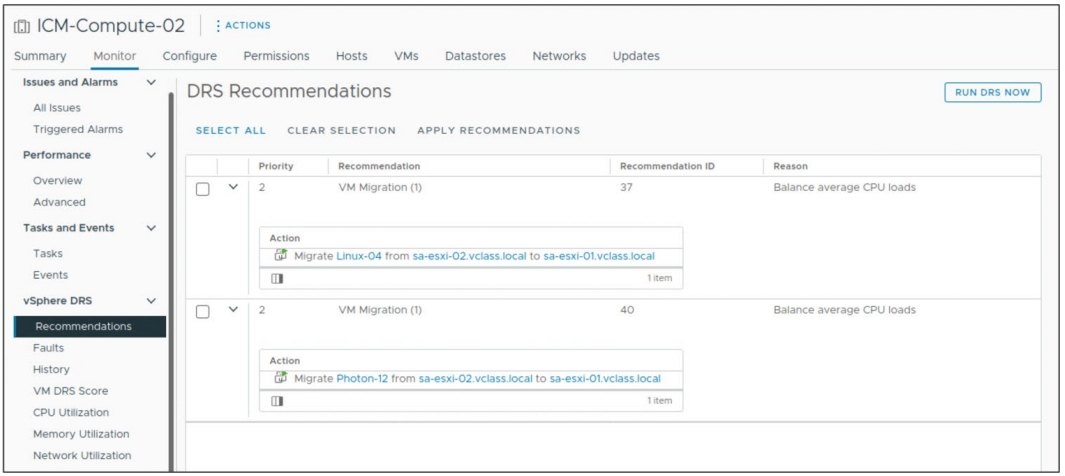

In the DRS Recommendations pane, you can see the current set of recommendations that are generated for optimizing resource use in the cluster through either migrations or power management. Recommendations appear in the list if vSphere DRS is set for Manual or Partially Automated mode.

To refresh the recommendations, click **RUN DRS NOW**.

To apply all recommendations, click **SELECT ALL** to select all recommendations and click **APPLY RECOMMENDATIONS**.

To apply a subset of the recommendations, select the checkbox next to each desired recommendation and click **APPLY RECOMMENDATIONS**.

In addition to the DRS Recommendations pane, you can select **Faults** to view the faults that occurred when the recommendations were applied. You can select **History** to view the history of vSphere DRS actions.

9-39 Maintenance Mode and Standby Mode

Maintenance mode:

- Should be used to service a host in a cluster
- Makes the host's resources unavailable for use

Standby mode:

- Is used by vSphere DPM to optimize power usage. When a host is placed in standby mode, it is powered off.

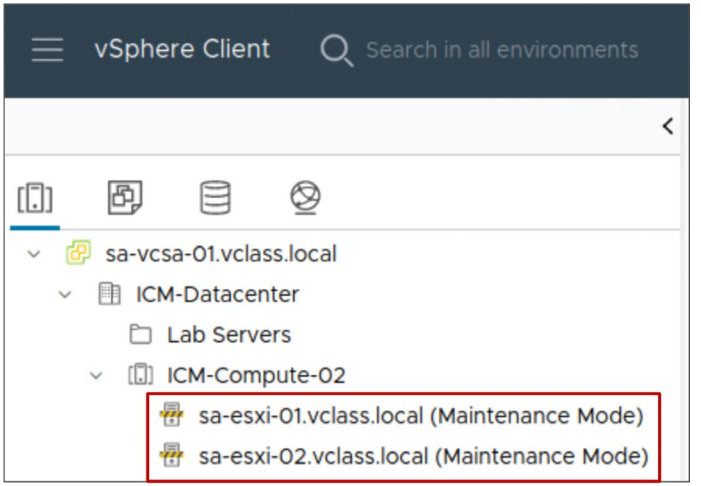

A host enters or leaves maintenance mode as the result of a user request. While in maintenance mode, the host does not allow you to deploy or power on a VM.

VMs that are running on a host entering maintenance mode must be shut down, suspended, or migrated to another host (either manually by a user or automatically by vSphere DRS). The host continues to run the Enter Maintenance Mode task until all VMs are powered down or moved away.

When no more running VMs are on the host, the host's icon indicates that it has entered maintenance mode. The host's **Summary** tab indicates the new state.

Place a host in maintenance mode before servicing the host, in instances such as installing more memory or removing a host from a cluster.

You can place a host in standby mode manually. When a host is placed in standby mode, it is powered off, except for the Baseboard Management Controller (BMC). The BMC is powered on even when the host itself is powered off. For more information on vSphere DPM, see *vSphere Resource Management* at https://docs.vmware.com/en/VMware-vSphere/index.html.

If you place a host in standby mode, the next time that vSphere DRS runs, it might undo your change or recommend that you undo the changes. If you want a host to remain powered off, place it in maintenance mode and turn it off.

9-40 Removing a Host from the vSphere DRS Cluster

To remove a host from a cluster:

1. Place the host in maintenance mode.
   - All running VMs on the host must be migrated to another host, shut down or suspended.
   - If DRS is in fully automated mode, powered-on VMs are automatically migrated from a host that is placed in maintenance mode.
2. Drag the host to a different inventory location, for example, the data center or another cluster.

The resources available for the cluster decrease.

When a host is put into maintenance mode, all its running VMs must be shut down, suspended, or migrated to other hosts by using vSphere vMotion. VMs with disks on local storage must be powered off, suspended, or migrated to another host and datastore.

When you remove the host from the cluster, the VMs that are currently associated with the host are also removed from the cluster. If the cluster still has enough resources to satisfy the reservations of all VMs in the cluster, the cluster adjusts resource allocation to reflect the reduced amount of resources.

9-41 Lab 24: Implementing vSphere DRS Clusters

Create a vSphere DRS cluster, use Cluster Quickstart to perform the basic configuration, and verify proper vSphere DRS functionality:

1. Create a Cluster That Is Configured for vSphere DRS
2. Verify vSphere vMotion Configuration on the ESXi Hosts
3. Add ESXi Hosts to the Cluster
4. Modify vSphere DRS Settings
5. Power On VMs and Review vSphere DRS Recommendations
6. Review vSphere DRS Recommendations When the Cluster Is Imbalanced

9-42 Review of Learner Objectives

- Describe the functions of a vSphere DRS cluster
- Explain how vSphere DRS determines VM placement on hosts in the cluster
- Recognize use cases for vSphere DRS settings
- Configure vSphere DRS in a cluster
- Monitor a vSphere DRS cluster

9-43 Lesson 3: Introduction to vSphere High Availability

9-44 Learner Objectives

- Describe how vSphere HA responds to various types of failures
- Describe how vSphere HA responds to network isolation
- Identify options for configuring network redundancy in a vSphere HA cluster

9-45 Protection at Every Level

With vSphere, you can reduce planned downtime, prevent unplanned downtime, and recover rapidly from outages.

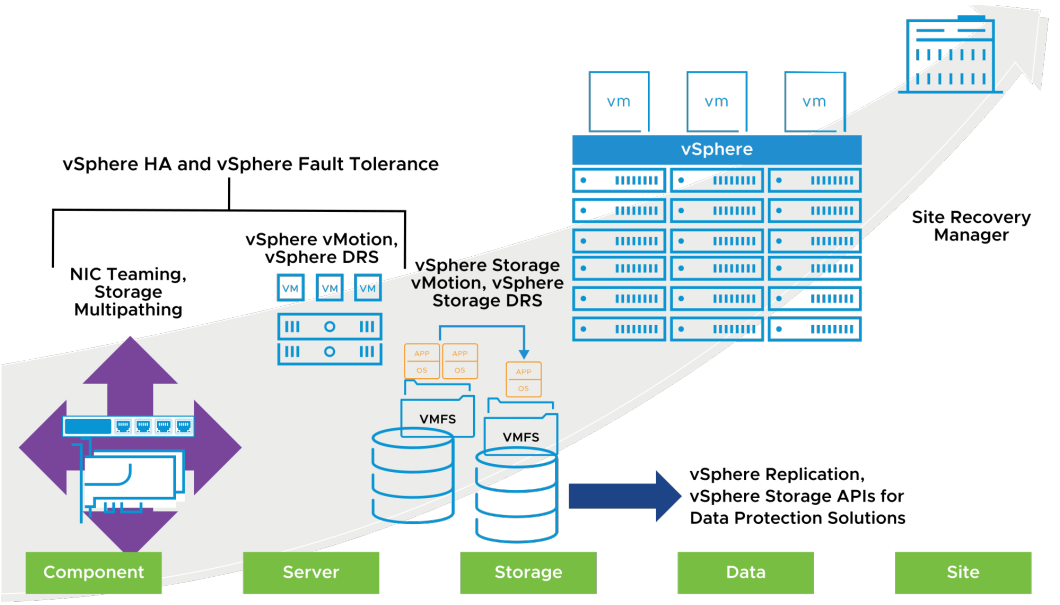

Whether planned or unplanned, downtime brings with it considerable costs. However, solutions to ensure higher levels of availability are traditionally costly, hard to implement, and difficult to manage.

VMware's software makes it simpler and less expensive to provide higher levels of availability for important applications. With vSphere, organizations can easily increase the baseline level of availability provided for all applications and provide higher levels of availability more easily and cost effectively. With vSphere, you can:

- Provide higher availability independent of hardware, operating system, and applications.
- Reduce planned downtime for common maintenance operations.
- Provide automatic recovery in cases of failure.

vSphere HA provides a base level of protection for your VMs by restarting VMs if a host fails. vSphere Fault Tolerance provides a higher level of availability, allowing users to protect any VM from a host failure with no loss of data, transactions, or connections. vSphere Fault Tolerance provides continuous availability by ensuring that the states of the primary and secondary VMs are identical at any point in the instruction execution of the VM.

vSphere vMotion and vSphere Storage vMotion keep VMs available during a planned outage, for example, when hosts or storage must be taken offline for maintenance. System recovery from unexpected storage failures is simple, quick, and reliable with the encapsulation property of VMs. You can use vSphere Storage vMotion to support planned storage outages resulting from upgrades to storage arrays to newer firmware or technology and VMFS upgrades.

With vSphere Replication, a vSphere platform can protect VMs natively by copying their disk files to another location where they are ready to be recovered.

VM encapsulation is used by third-party backup applications that support file and image-level backups using vSphere Storage APIs – Data Protection. Backup solutions play prominent roles in recovering from deleted files or disks and corrupt or infected guest operating systems or file systems.

With Site Recovery Manager, you can quickly restore your organization's IT infrastructure, shortening the time that you experience a business outage. Site Recovery Manager automates setup, failover, and testing of disaster recovery plans. Site Recovery Manager requires that you install vCenter at the protected site and at the recovery site. Site Recovery Manager also requires either host-based replication through vSphere Replication or preconfigured array-based replication between the protected site and the recovery site.

9-46 About vSphere High Availability

vSphere High Availability (HA) provides rapid recovery from outages and cost-effective high availability for applications running in VMs. vSphere HA protects application availability in several ways.

Protects Against	How Does vSphere HA Provide Protection?
ESXi host failure	By restarting the VMs on other hosts within the cluster
VM failure	By restarting the VM when a VMware Tools heartbeat is not received within a set time
Application failure	By restarting the VM when an application heartbeat is not received within a set time
Datastore accessibility failure	By restarting the affected VMs on other hosts that still can access the datastores.
Network isolation	By restarting VMs if their host becomes isolated on the heartbeat network. This protection is provided even if the network becomes partitioned.

Unlike other clustering solutions, vSphere HA protects all workloads by using the infrastructure itself. After you configure vSphere HA, no actions are required to protect new VMs. All workloads are automatically protected by vSphere HA.

9-47 vSphere HA Scenario: ESXi Host Failure

When a host fails, vSphere HA restarts the impacted VMs on other hosts in the cluster.

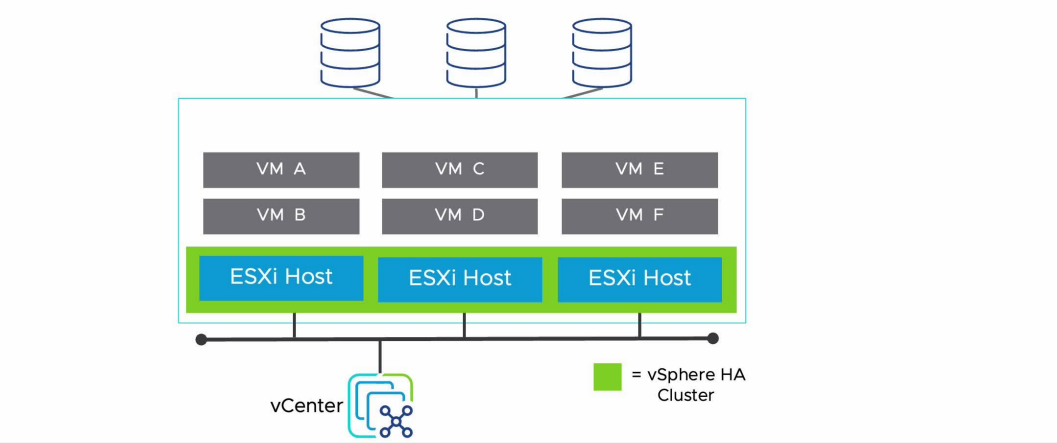

To play the animation, go to https://vmware.bravais.com/s/Ww2W2cOxz1xVpMdinLsQ

vSphere HA can determine whether a ESXi host is isolated or has failed. If an ESXi host fails, vSphere HA attempts to restart all VMs that were running on the failed host by using the remaining hosts in the cluster.

In every cluster, the time to recover depends on how long it takes your guest operating systems and applications to restart when the VM is failed over.

9-48 vSphere HA Scenario: Guest Operating System Failure

When a VM stops sending heartbeats or the VM process (vmx) fails unexpectedly, vSphere HA restarts the VM on the same host.

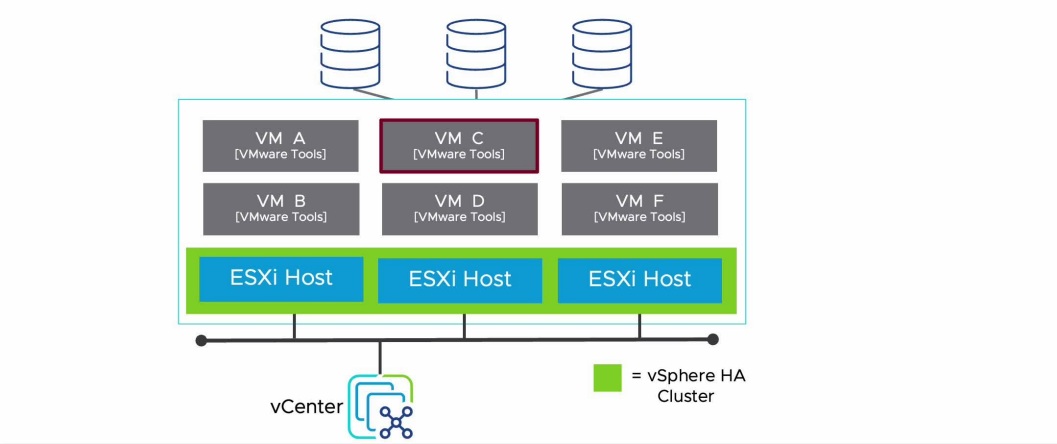

To play the animation, go to https://vmware.bravais.com/s/51nBTCVaUE3JJH0Ou0E2

If VM monitoring is configured, the vSphere HA agent on each individual host monitors VMware Tools in each VM running on the host. When a VM stops sending heartbeats, the VM is restarted on the same host.

9-49 vSphere HA Scenario: Application Failure

When an application stops sending heartbeats, vSphere HA restarts the impacted VM on the same host.

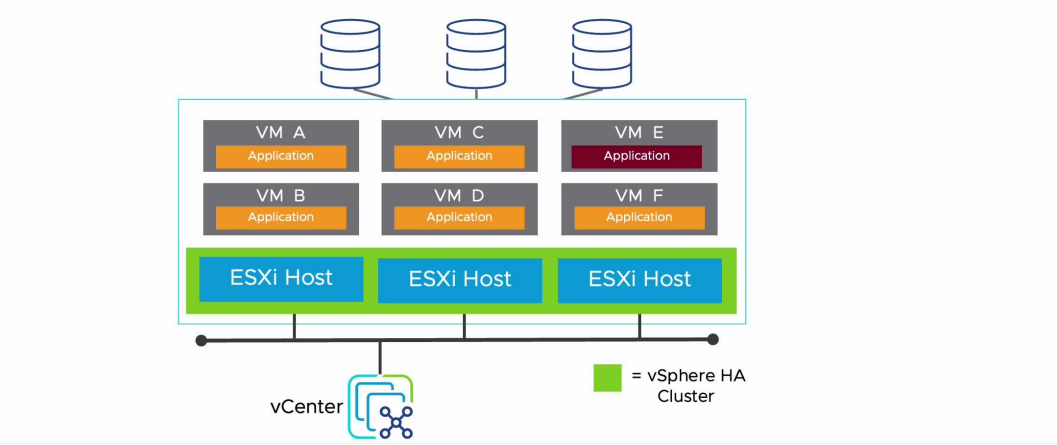

To play the animation, go to https://vmware.bravais.com/s/sYP7PbI2HxVLZxu9pbbO

To enable Application Monitoring, you must obtain the appropriate SDK, or use an application that supports VMware Application Monitoring. Then, use the application to set up customized heartbeats for the applications you want to monitor. After you have done this, Application Monitoring works much the same way that VM Monitoring does. The agent on each host can monitor heartbeats of applications running in each VM. When an application fails (stops sending heartbeats), the VM on which the application was running is restarted on the same host.

9-50 vSphere HA Scenario: Datastore Accessibility Failures

vSphere HA can detect datastore accessibility failures and provide automated recovery for affected VMs.

You can determine the response that vSphere HA makes to such a failure:

- All paths down (APD):
  - Recoverable.
  - Represents a transient or unknown accessibility loss.
  - Response can be either **Issue events**, **Power off and restart VMs – Conservative restart policy**, or **Power off and restart VMs – Aggressive restart policy**.
- Permanent device loss (PDL):
  - Unrecoverable loss of accessibility.
  - Occurs when a storage device reports that the datastore is no longer accessible by the host.
  - Response can be either **Issue events** or **Power off and restart VMs**.

You can determine the response that vSphere HA makes to such a failure, ranging from the creation of event alarms to VM restarts on other hosts.

Power off and restart VMs – Conservative restart policy: vSphere HA does not attempt to restart the affected VMs unless vSphere HA determines that another host can restart the VMs. The host experiencing the all paths down (APD) communicates with the vSphere HA primary host to determine whether the cluster has sufficient capacity to power on the affected VMs. If the primary host determines that sufficient capacity is available, the host experiencing the APD stops the VMs so that the VMs can be restarted on a healthy host. If the host experiencing the APD cannot communicate with the primary host, no action is taken.

Power off and restart VMs – Aggressive restart policy: vSphere HA stops the affected VMs even if it cannot determine that another host can restart the VMs. The host experiencing the APD attempts to communicate with the primary host to determine whether the cluster has sufficient capacity to power on the affected VMs. If the primary host is not reachable, sufficient capacity to restart the VMs is unknown. In this scenario, the host takes the risk and stops the VMs, so they can be restarted on the remaining healthy hosts. However, if sufficient capacity is not available, vSphere HA might not be able to recover all the affected VMs. This result is common in a network partition scenario where a host cannot communicate with the primary host to get a definitive response to the likelihood of a successful recovery.

9-51 Importance of Heartbeat Networks

A heartbeat network is implemented by using a VMkernel port that is marked for management or vSAN traffic.

Heartbeats have the following characteristics:

- They are sent between the primary host and the secondary hosts.
- They are used to determine whether a primary host or a secondary host has failed.
- They are sent over a heartbeat network.

When both vSAN and vSphere HA are activated on the cluster, vSphere HA uses the vSAN network as the heartbeat network instead of the management network.

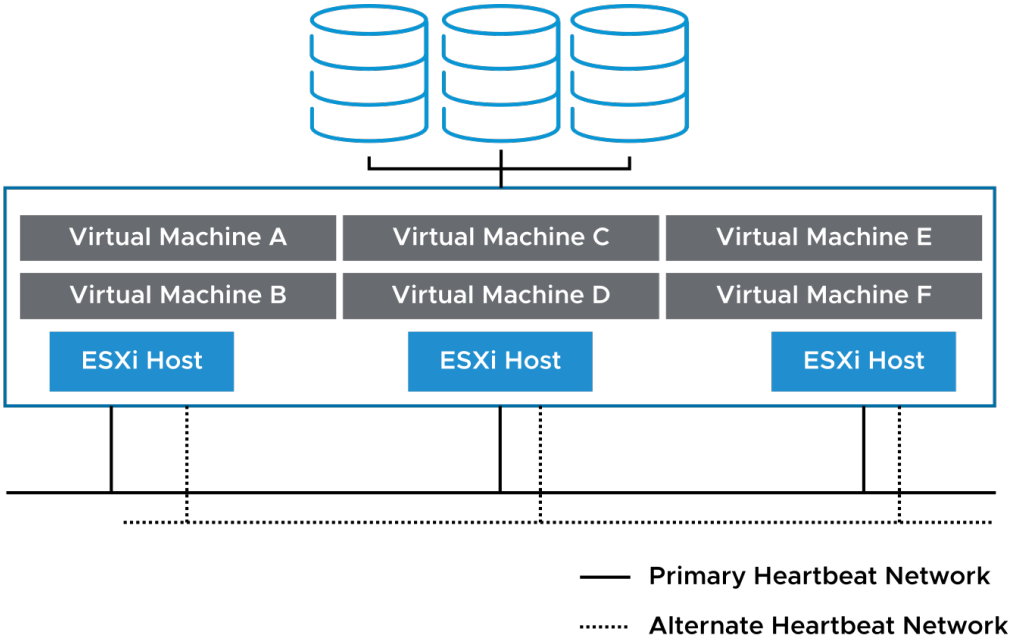

Redundant heartbeat networking supports the reliable detection of failures and prevents isolation or partition conditions from occurring because heartbeats can be sent over multiple networks.

Redundant heartbeat networking is the best approach for your vSphere HA cluster. When a primary host's connection fails, a second connection is still available to send heartbeats to the other hosts. If you do not provide redundancy, your failover setup has a single point of failure.

9-52 vSphere HA Scenario: Protecting VMs Against Network Isolation

vSphere HA restarts VMs if their host is isolated on the management or vSAN network.

Host network isolation occurs when a host is still running, but it can no longer observe traffic from the vSphere HA agents on the heartbeat network.

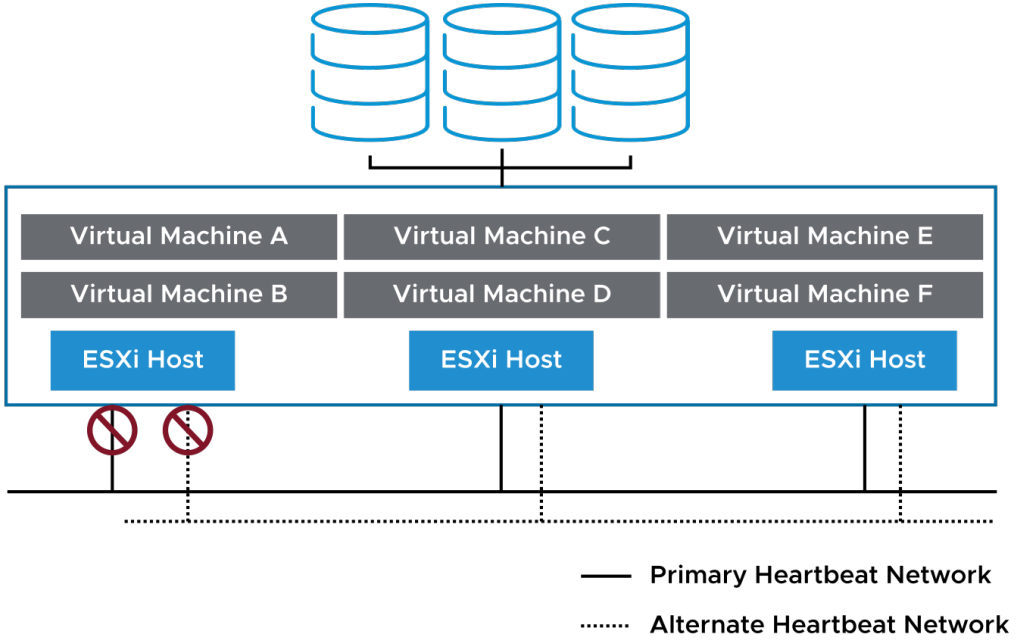

If you verify that the network infrastructure is sufficiently redundant and that at least one network path is always available, host network isolation might not occur.

9-53 Heartbeat Network Redundancy Using NIC Teaming

Redundant heartbeat networks ensure reliable failure detection and minimize the chance of host-isolation scenarios.

You can use NIC teaming to create a redundant heartbeat network on ESXi hosts.

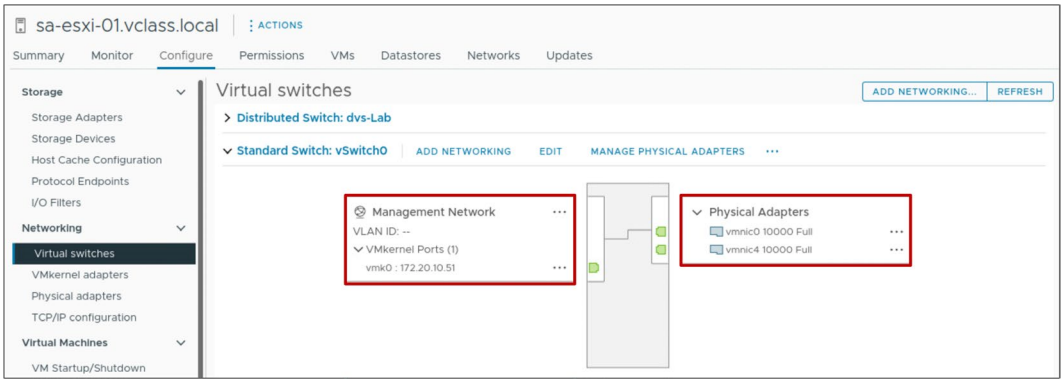

In this example, vmnic0 and vmnic4 form a NIC team in the Management network. The vmk0 VMkernel port has the Management service activated.

9-54 Heartbeat Network Redundancy Using Additional Networks

You can create redundancy by configuring additional heartbeat networks.

For example, create a second VMkernel port in a port group on a separate virtual switch with its own physical adapter.

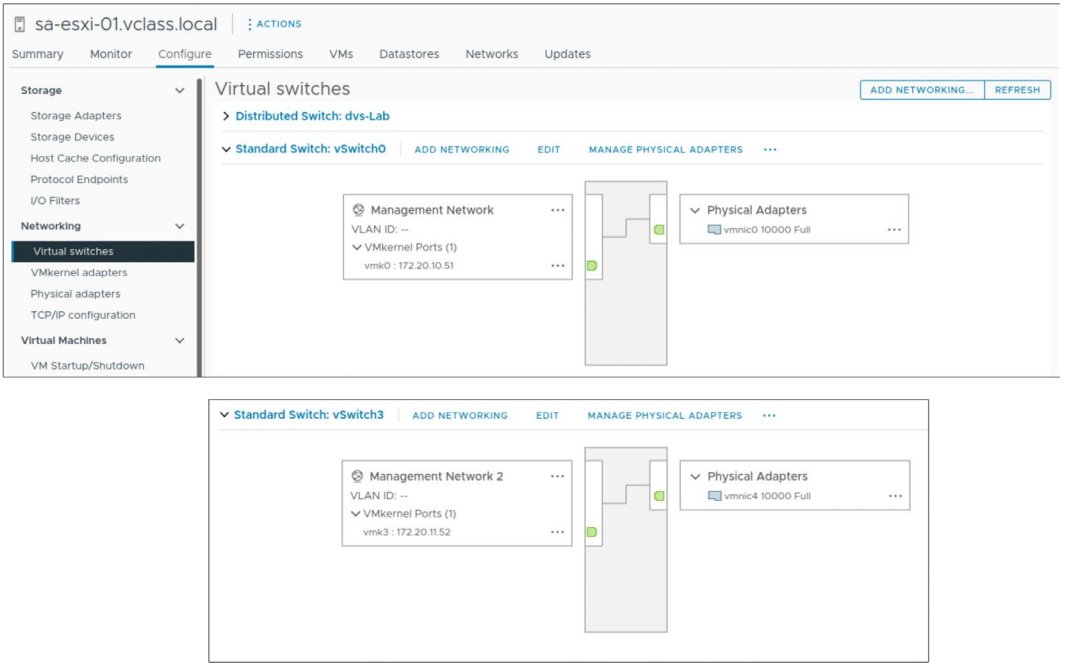

In most implementations, NIC teaming provides sufficient heartbeat redundancy. As an alternative, you can create a second VMkernel port attached to a separate virtual switch. Both VMkernel ports can also be on the same virtual switch, but in different port groups, each port group using a different physical adapter.

In this example, the management network is used as the heartbeat network. The original management network connection is used for network and management purposes. When the second management network connection is created, the primary host sends heartbeats over both management network connections. If one path fails, the primary host still sends and receives heartbeats over the other path.

9-55 Review of Learner Objectives

- Describe how vSphere HA responds to various types of failures
- Describe how vSphere HA responds to network isolation
- Identify options for configuring network redundancy in a vSphere HA cluster

9-56 Lesson 4: vSphere HA Architecture

9-57 Learner Objectives

- Identify the heartbeat mechanisms used by vSphere HA
- Describe failure scenarios
- Recognize vSphere HA design considerations

9-58 vSphere HA Architecture: Agent Communication

When vSphere HA is configured in a cluster, the Fault Domain Manager (FDM) service is uploaded to each host in the cluster and started.

The FDM service is also known as the vSphere HA agent.

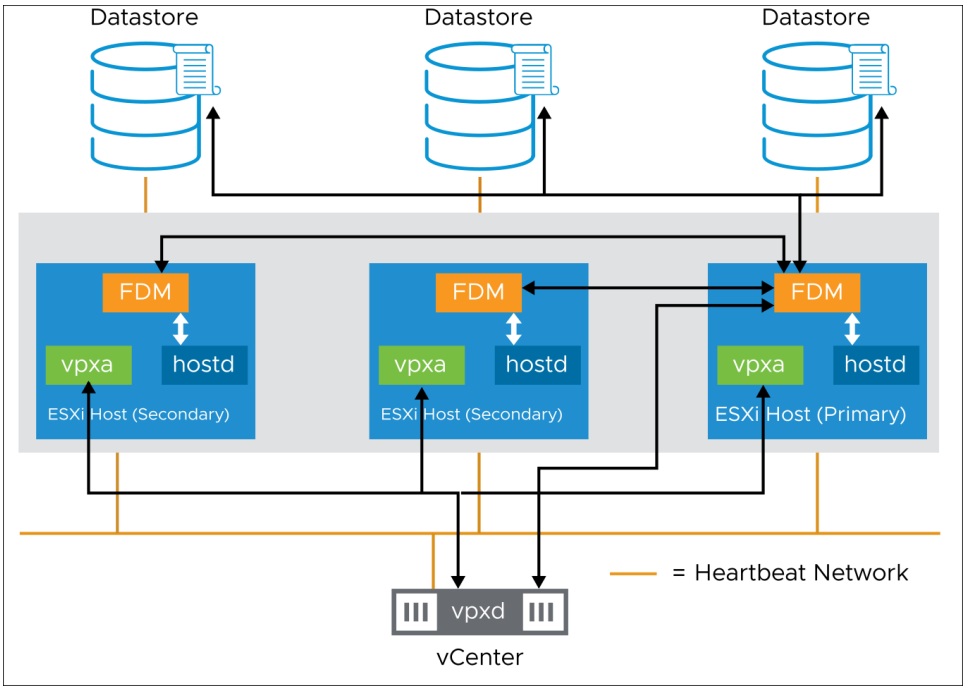

The vSphere HA cluster is managed by a primary host. All other hosts are called secondary hosts. Fault Domain Manager (FDM) services on secondary hosts all communicate with FDM on the primary host.

Hosts cannot participate in a vSphere HA cluster if they are in maintenance mode, in standby mode, or disconnected from vCenter.

To determine which host is the primary host, an election process takes place. The host that can access the greatest number of datastores is elected the primary host. If more than one host sees the same number of datastores, the election process determines the primary host by using the host Managed Object ID (MOID) assigned by vCenter.

The election process for a new primary host completes in approximately 15 seconds and occurs under these circumstances:

- vSphere HA is configured.
- The primary host encounters a system failure because of one of the following factors:
  - The primary host is placed in maintenance mode.
  - The primary host is placed in standby mode.
  - vSphere HA is reconfigured.
- The secondary hosts cannot communicate with the primary host because of a network problem.

During the election process, the candidate vSphere HA agents communicate with each other over the heartbeat network (vSAN or management network) by using User Datagram Protocol (UDP). All network connections are point-to-point. After the primary host is determined, the primary host and secondary hosts communicate using secure TCP. When vSphere HA is started, vCenter contacts the primary host and sends a list of hosts with membership in the cluster with the cluster configuration. That information is saved to local storage on the primary host and then pushed out to the secondary hosts in the cluster. If additional hosts are added to the cluster during normal operation, the primary host sends an update to all hosts in the cluster.

The primary host provides an interface for vCenter to query the state of and report on the health of the fault domain and VM availability. vCenter tells the vSphere HA agent which VMs to protect with their VM-to-host compatibility list. The agent learns about state changes through hostd and vCenter learns through vpxa. The primary host monitors the health of the secondary hosts and takes responsibility for VMs that were running on a failed secondary host.

A secondary host monitors the health of VMs running locally and sends state changes to the primary host. A secondary host also monitors the health of the primary host.

vSphere HA is configured, managed, and monitored through vCenter. The vpxd process, which runs on the vCenter system, maintains the cluster configuration data. The vpxd process reports cluster configuration changes to the primary host. The primary host advertises a new copy of the cluster configuration information, and each secondary host fetches an updated copy. Each secondary host writes the updated configuration information to local storage. A list of protected VMs is stored on each datastore. The VM list is updated after each user-initiated power-on (protected) and power off (unprotected) operation. The VM list is updated after vCenter observes these operations.

A VM becomes protected when an operation results in a power on. Reverting a VM to a snapshot with memory state causes the VM to power on and become protected. Similarly, a user action that causes the VM to power off, for example, reverting to a snapshot without memory state or a standby operation performed in the guest, causes the VM to become unprotected.

9-59 vSphere HA Architecture: Network Heartbeats

The primary host sends periodic heartbeats to the secondary hosts.

In this way, the secondary hosts know that the primary host is alive, and the primary host knows that the secondary hosts are alive.

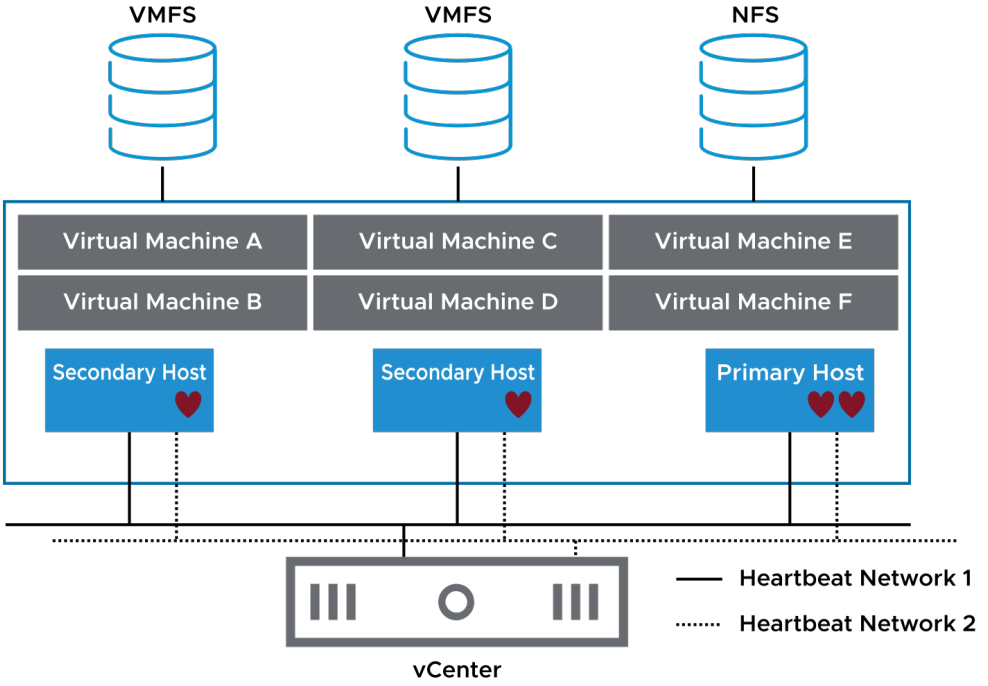

Heartbeats are sent to each secondary host from the primary host over all configured heartbeat networks. However, secondary hosts use only one heartbeat network to communicate with the primary host. If the heartbeat network used to communicate with the primary host fails, the secondary host switches to another heartbeat network to communicate with the primary host.

If the secondary host does not respond within the predefined timeout period, the primary host declares the secondary host as agent unreachable. When a secondary host is not responding, the primary host attempts to determine the cause of the secondary host's inability to respond. The primary host must determine whether the secondary host crashed, is not responding because of a network failure, or the vSphere HA agent is in an unreachable state.

9-60 vSphere HA Architecture: Datastore Heartbeats

When the primary host cannot communicate with a secondary host over the heartbeat network, the primary host uses datastore heartbeating to determine the cause:

- Secondary host failure
- Network partition
- Network isolation

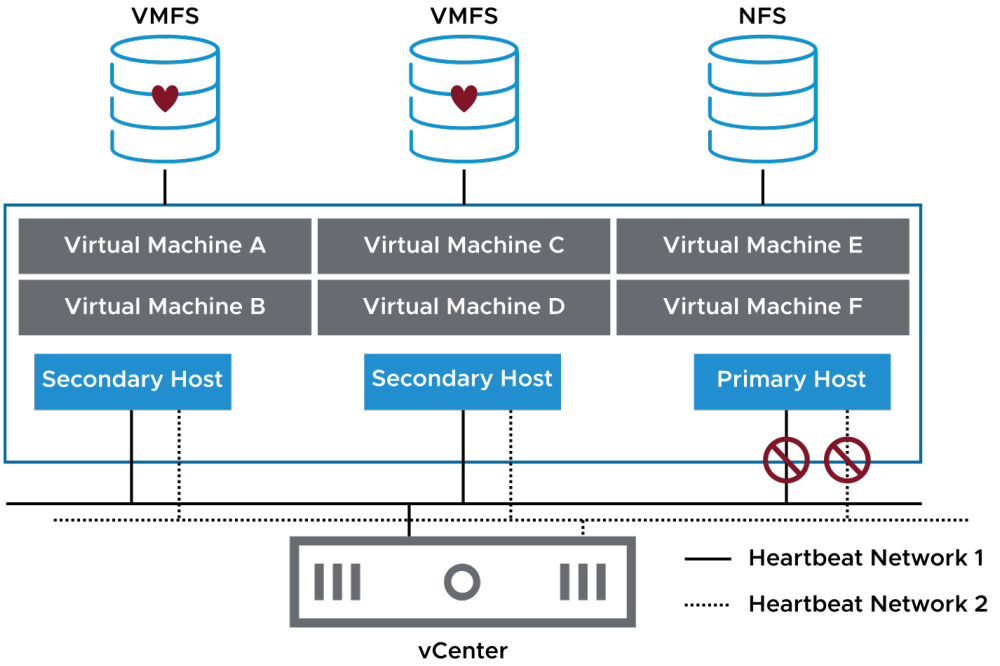

Using datastore heartbeating, the primary host determines whether a host has failed or a network isolation has occurred. If datastore heartbeating from the host stops, the host is considered failed. In this case, the failed host's VMs are started on another host in the vSphere HA cluster.

9-61 vSphere HA Failure Scenarios

vSphere HA can identify and respond to various types of failures:

- Secondary host failure
- Primary host failure
- Network failure (host isolation)
- Datastore accessibility failures:
  - APD
  - PDL

vSphere HA can also determine whether an ESXi host is isolated or has failed. Isolation refers to when an ESXi host cannot see traffic coming from the other hosts in the cluster and cannot ping its configured isolation address. If an ESXi host fails, vSphere HA attempts to restart the VMs that were running on the failed host on the remaining hosts in the cluster. If the ESXi host is isolated because it cannot ping its configured isolation address and sees no heartbeat network traffic, the host executes the Host Isolation Response.

9-62 Failed Secondary Hosts

When a secondary host does not respond to the network heartbeat issued by the primary host, the primary host tries to identify the cause.

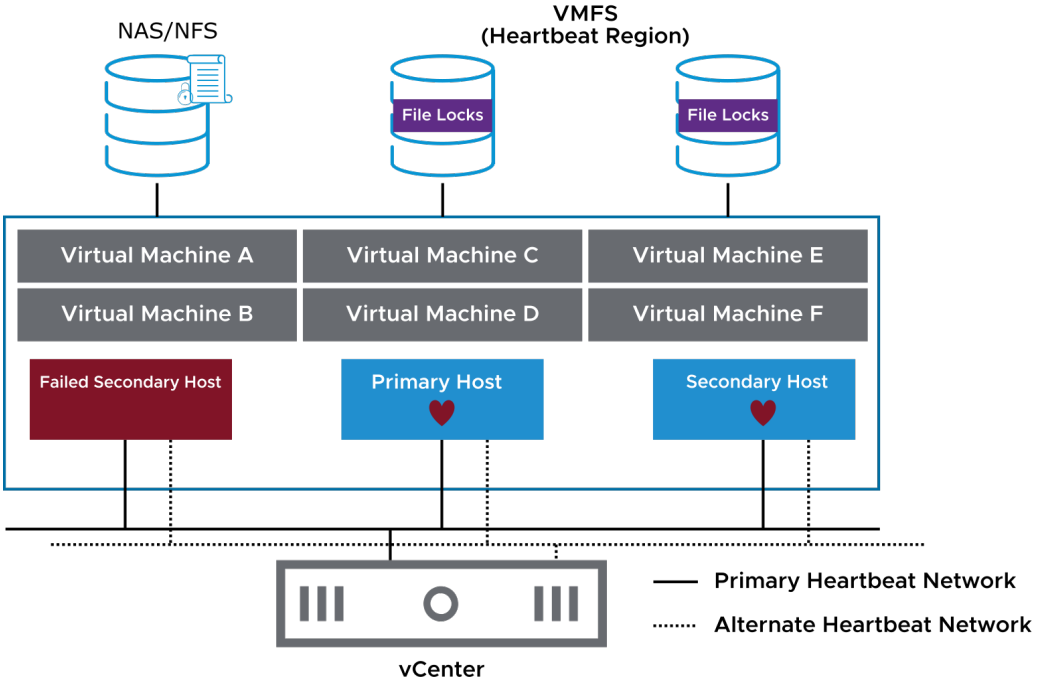

The primary host must determine whether the secondary host is isolated or has failed, for example, because of a misconfigured firewall rule or component failure. The type of failure dictates how vSphere HA responds.

When the primary host cannot communicate with a secondary host over the heartbeat network, the primary host uses datastore heartbeating to determine whether the secondary host failed, is in a network partition, or is network-isolated. If the secondary host stops datastore heartbeating, the secondary host is considered to have failed, and its virtual machines are restarted elsewhere.

For VMFS, a heartbeat region on the datastore is read to find out if the host is still heartbeating to it. For NFS datastores, vSphere HA reads the `host--hb` file, which is locked by the ESXi host accessing the datastore. The file guarantees that the VMkernel is heartbeating to the datastore and periodically updates the lock file.

The lock file time stamp is used by the primary host to determine whether the secondary host is isolated or has failed.

In both storage examples, the vCenter instance selects a small subset of datastores for hosts to heartbeat to. The datastores that are accessed by the greatest number of hosts are selected as candidates. But two datastores are selected (by default) to keep the associated overhead and processing to a minimum.

9-63 Failed Primary Hosts

When the primary host is placed in maintenance mode or fails, the secondary hosts detect that the primary host is no longer issuing heartbeats.

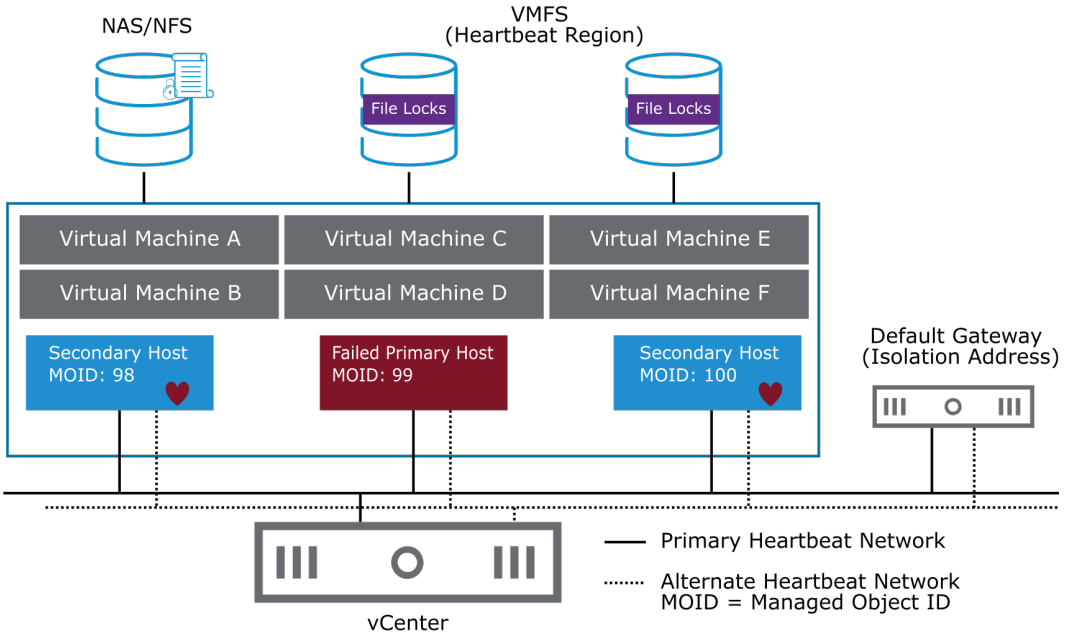

To determine which host is the primary host, an election process takes place. The host that can access the greatest number of datastores is elected the primary host. If more than one host sees the same number of datastores, the election process determines the primary host by using the host Managed Object ID (MOID) assigned by vCenter. If the primary host fails, is shut down, or is placed into maintenance mode, a new election is held.

9-64 Isolated Hosts

A host is declared isolated when both the following conditions occur:

- The host does not receive network heartbeats.
- The host cannot ping its isolation addresses.

When a host identifies itself as isolated, it executes the cluster's isolation response.

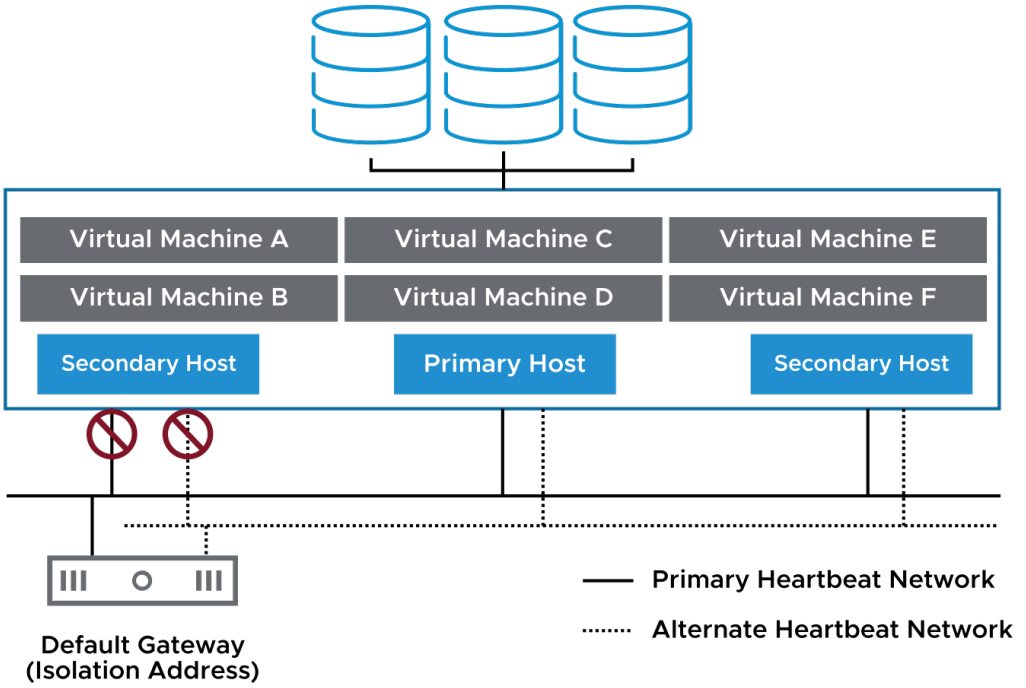

The slide illustrates a scenario that might lead to host isolation. If a host loses connectivity to both the primary heartbeat network and the alternate heartbeat network, the host does not receive network heartbeats from the other hosts in the vSphere HA cluster. The same host cannot ping its isolation address. An isolation address is an IP address or FQDN that can be manually specified (the default is the host's default gateway).

If a host is isolated, the primary host must determine if that host is active and merely isolated, by checking for the datastore heartbeats. Datastore heartbeats are used by vSphere HA only when a host is isolated or partitioned.

Host isolation response determines what happens when a host in a vSphere HA cluster loses its heartbeat network connections but continues to run. You can use the isolation response for vSphere HA to power off virtual machines that run on an isolated host and restart them on a nonisolated host.

9-65 VM Storage Failures

Storage connectivity problems might arise because of:

- Network or switch failure
- Array misconfiguration
- Power outage

Storage connectivity problems affect VM availability:

- VMs on affected hosts are difficult to manage.
- Applications with attached disks fail.

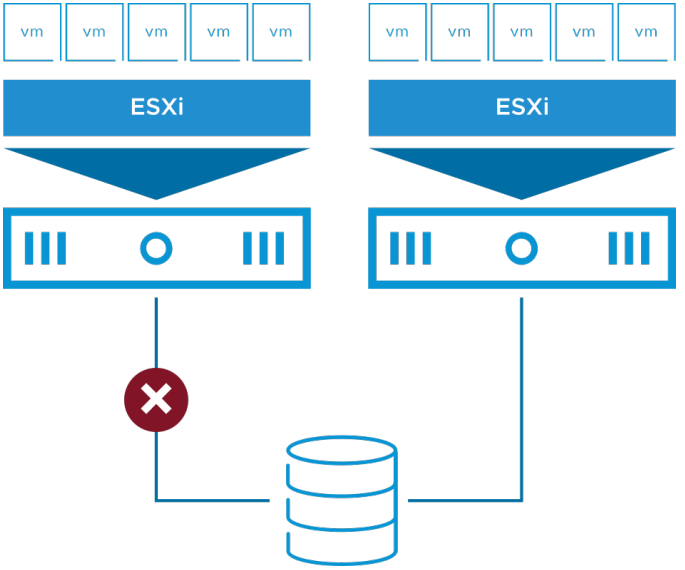

9-66 Protecting Against Storage Failures with VMCP

VM Component Protection protects against storage failures on a VM.

With VMCP, vSphere HA can detect datastore accessibility failures and provide automated recovery for affected VMs.

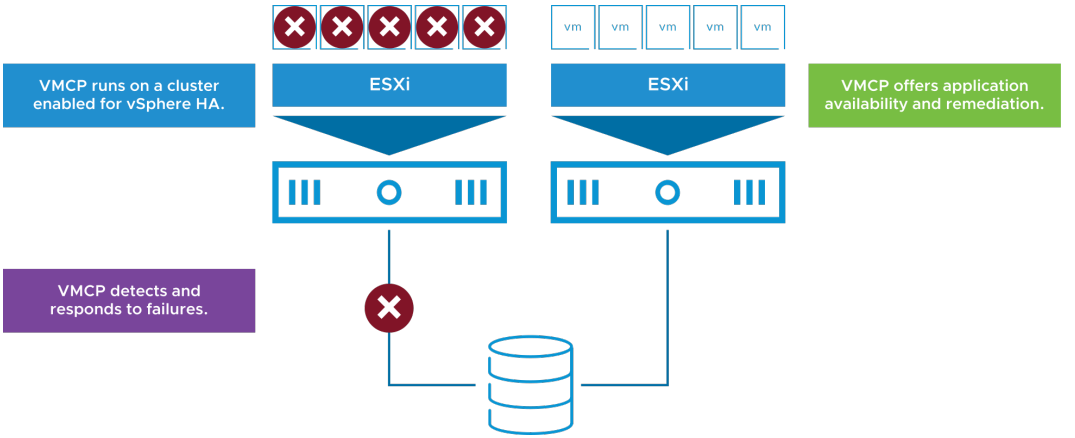

When a datastore accessibility failure occurs, the affected host can no longer access the storage path for a specific datastore. You can determine the response that vSphere HA gives to such a failure, ranging from the creation of event alarms to VM restarts on other hosts.

For more information about VM Component Protection, see https://blogs.vmware.com/vsphere/.

9-67 vSphere HA Design Considerations

When designing your vSphere HA cluster, consider these guidelines:

- Implement redundant heartbeat networks and redundant isolation addresses:
  - Redundancy minimizes host isolation events.
- Physically separate VM networks from the heartbeat networks.
- Implement datastores so that they are separated from the heartbeat network by using one or both of the following approaches:
  - Use Fibre Channel over fiber optic for your datastores.
  - If you use IP storage, physically separate your IP storage network from the heartbeat network.

If a datastore is based on Fibre Channel, an Ethernet network failure does not disrupt datastore access. When using datastores based on IP storage (for example, NFS, iSCSI, or Fibre Channel over Ethernet), you must physically separate the IP storage network and the heartbeat network.

9-68 Review of Learner Objectives

- Identify the heartbeat mechanisms used by vSphere HA
- Describe failure scenarios
- Recognize vSphere HA design considerations

9-69 Lesson 5: Configuring vSphere HA

9-70 Learner Objectives

- Recognize the requirements for creating and using a vSphere HA cluster
- Recognize the use cases for various vSphere HA settings
- Recognize when to use vSphere Fault Tolerance
- Configure a vSphere HA cluster

9-71 vSphere HA Prerequisites

To create a vSphere HA cluster, you must meet several requirements:

- All hosts must be configured with static IP addresses. If you are using DHCP, the address must persist across host reboots.
- All hosts must have at least one heartbeat network in common.
- For VM monitoring to work, VMware Tools must be installed in every VM.
- You must not exceed the maximum number of hosts that are allowed in a cluster.

  See VMware Configuration Maximums at https://configmax.vmware.com.

To determine the maximum number of hosts per cluster, see VMware Configuration Maximums at https://configmax.vmware.com.

9-72 Configuring vSphere HA Settings

When you create or configure a vSphere HA cluster, you must configure settings that determine how the feature works.

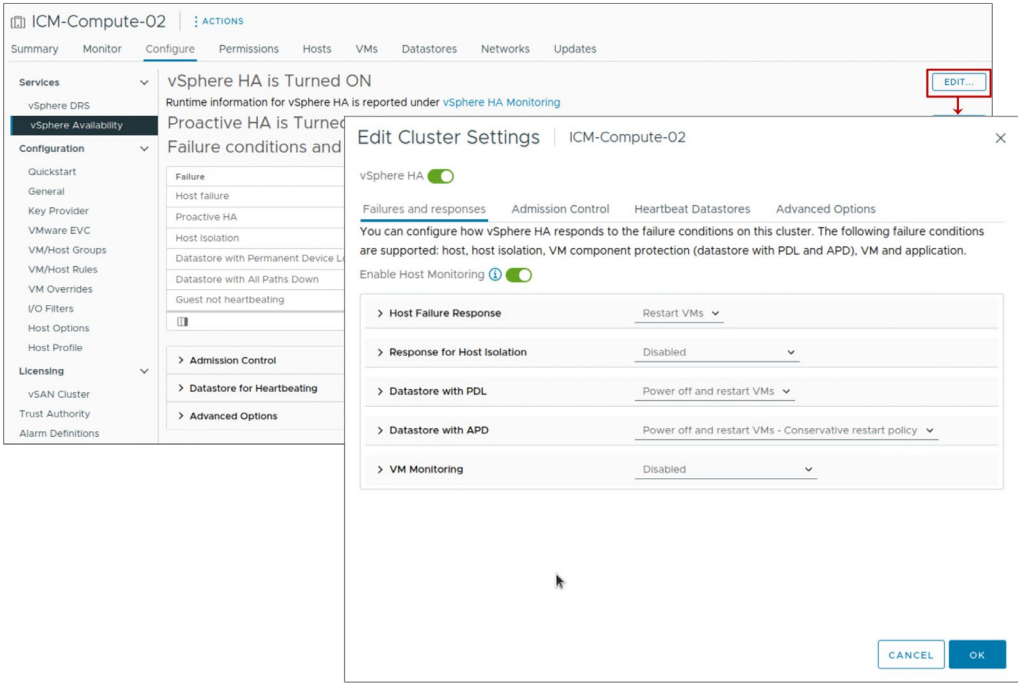

In the vSphere Client, you can configure the following vSphere HA settings:

- Availability failure conditions and responses: Provide settings for host failure responses, host isolation, VM monitoring, and VMCP.
- Admission control: Activate or deactivate admission control for the vSphere HA cluster and select a policy for how it is enforced.
- Heartbeat datastores: Specify preferences for the datastores that vSphere HA uses for datastore heartbeating.
- Advanced options: Customize vSphere HA behavior by setting advanced options.

9-73 vSphere HA Settings: Failures and Responses

You use the Failures and responses pane to configure a cluster's response if a failure occurs.

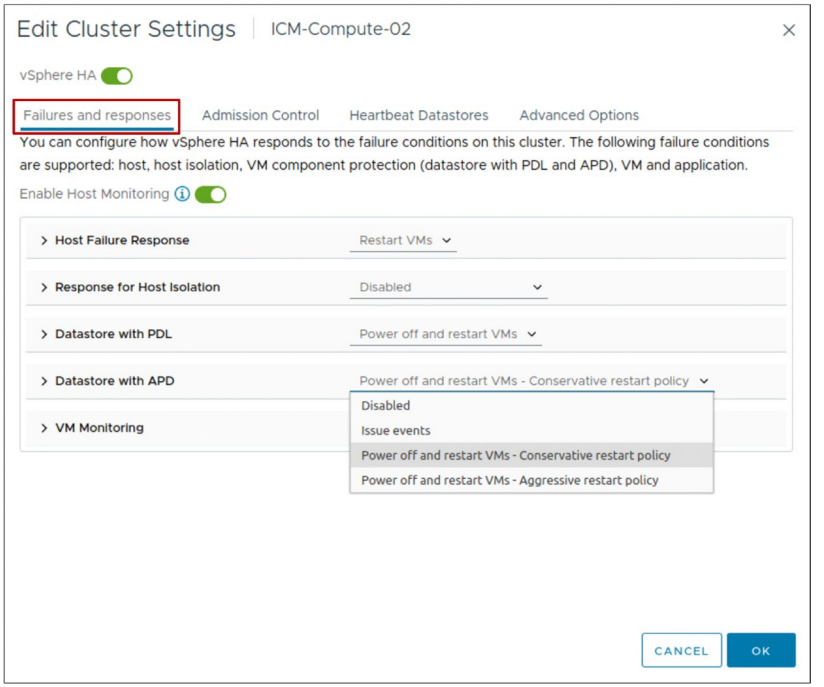

Using the Failures and Responses pane, you can configure how your cluster must function when problems occur. You can specify the vSphere HA cluster's response for host failures and isolation. You can configure the VMCP actions when permanent device loss and all paths down situations occur and activate the VM monitoring.

The following host failure responses are available:

- Disabled: Host monitoring is turned off and VMs are not restarted.
- Restart VMs: VMs are failed over based on their restart priority.

The following responses to host isolation are available:

- Disabled
- Power off and restart VMs
- Shut down and restart VMs

The following responses to a datastore PDL condition are available:

- Disabled
- Issue events: No action is taken against the affected VMs. The administrator is notified when a PDL event occurs.
- Power off and restart VMs

The following responses to a datastore APD condition are available:

- Disabled
- Issue events: No action is taken against the affected VMs. The administrator is notified when an APD event occurs.
- Power off and restart VMs - Conservative restart policy: vSphere HA does not attempt to restart the affected VMs unless vSphere HA determines that another host can restart the VMs.

  The host experiencing the APD communicates with the primary host to determine whether sufficient capacity exists in the cluster to power on the affected VMs. If the primary host determines that sufficient capacity exists, the host experiencing the APD stops the VMs so that the VMs can be restarted on a healthy host. If the host experiencing the APD cannot communicate with the primary host, the VM is not stopped.
- Power off and restart VMs - Aggressive restart policy: vSphere HA stops the affected VMs even if it cannot determine that another host can restart the VMs.

  The host experiencing the APD attempts to communicate with the primary host to determine if sufficient capacity exists in the cluster to power on the affected VMs. If the primary host is not reachable, sufficient capacity for restarting the VMs is unknown. In this scenario, the host takes the risk and stops the VMs so that they can be restarted on the remaining healthy hosts.

  However, if sufficient capacity is not available, vSphere HA might not be able to recover all the affected VMs. This result is common in a network partition scenario in which a host cannot communicate with the primary host.

The following VM monitoring options are available:

- VM monitoring only
- VM and application monitoring

9-74 vSphere HA Setting: Default VM Restart Priority

The VM restart priority determines the order in which vSphere HA restarts VMs on a running host.

VMs are assigned the Medium restart priority by default, unless the restart priority is explicitly set using VM overrides.

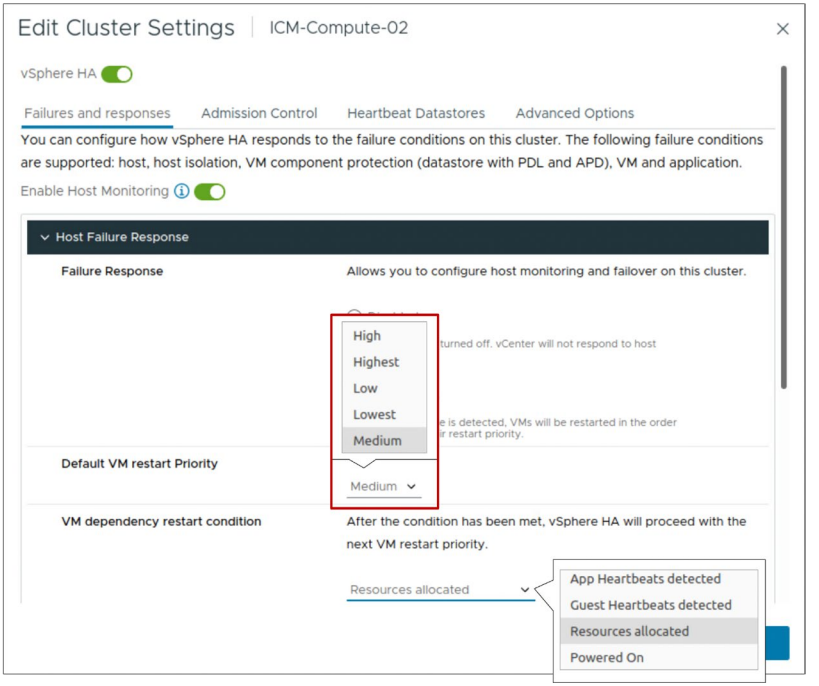

After a host failure, VMs are assigned to other hosts with unreserved capacity, with the highest priority VMs placed first. The process continues to those VMs with lower priority until all have been placed, or no more cluster capacity is available to meet the reservations or memory overhead of the VMs. A host then restarts the VMs assigned to it in priority order.

If insufficient resources exist, vSphere HA waits for more unreserved capacity to become available, for example, because of a host coming back online, and then retries the placement of these VMs. To reduce the chance of this situation occurring, configure vSphere HA admission control to reserve more resources for failures. With admission control, you can control the amount of cluster capacity that is reserved by VMs, which is unavailable to meet the reservations and memory overhead of other VMs if a failure occurs.

Optionally, you can configure a delay when a certain restart condition is met.

The following conditions must be met before a VM is considered ready:

- VM has resources allocated
- VM is powered on
- VMware Tools heartbeat is detected
- VMware Tools application heartbeat is detected

9-75 vSphere HA Settings: VM-Level Restart Priority

You can customize the restart priority for individual VMs in a cluster to override the default level set for the entire cluster.

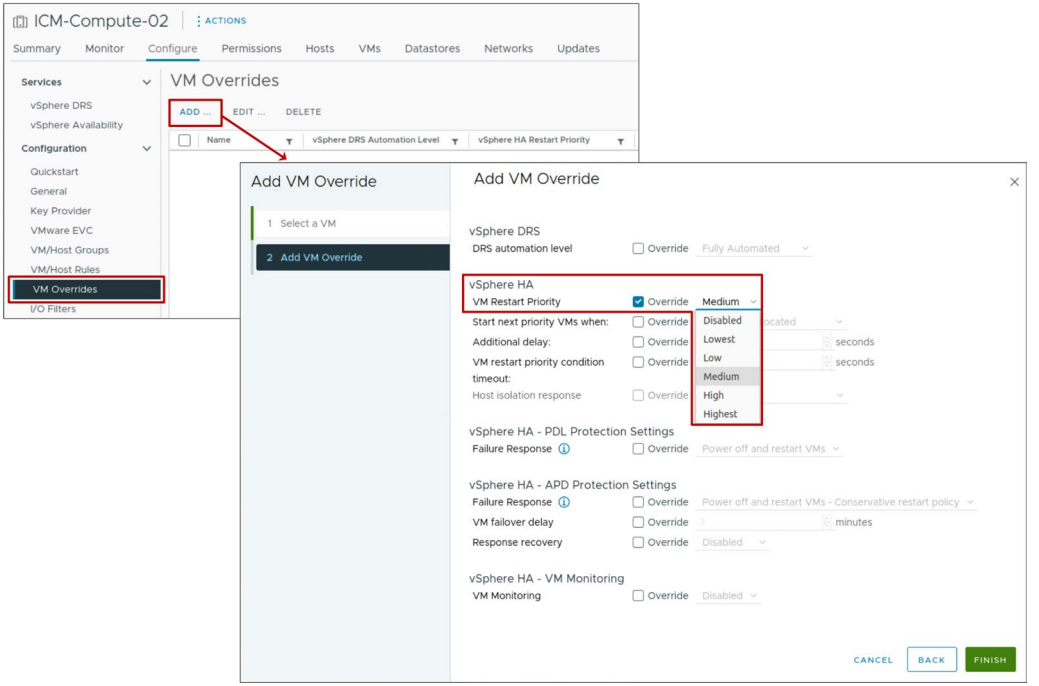

9-76 About vSphere HA Orchestrated Restart

Orchestrated restart is an alternative to using the VM restart priority settings for vSphere HA.

With orchestrated restart, you define the order in which the VMs restart, which is useful when services must be started in a particular order.

A common use case is for restarting a three-tier application.

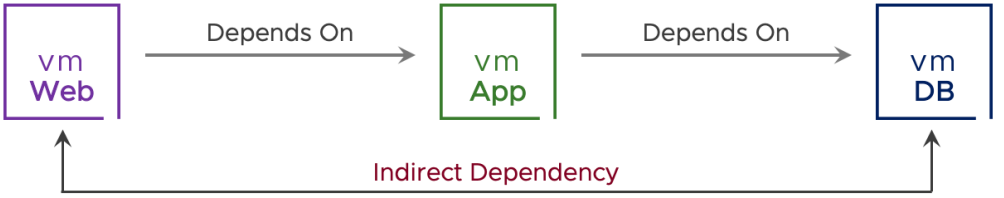

Only direct dependencies are supported. Creating cyclical dependencies causes a VM's restart to fail.

9-77 Orchestrated Restart In Action

vSphere HA restarts VMs only from a failed host. Configure affinity rules to keep VMs on the same host if necessary.

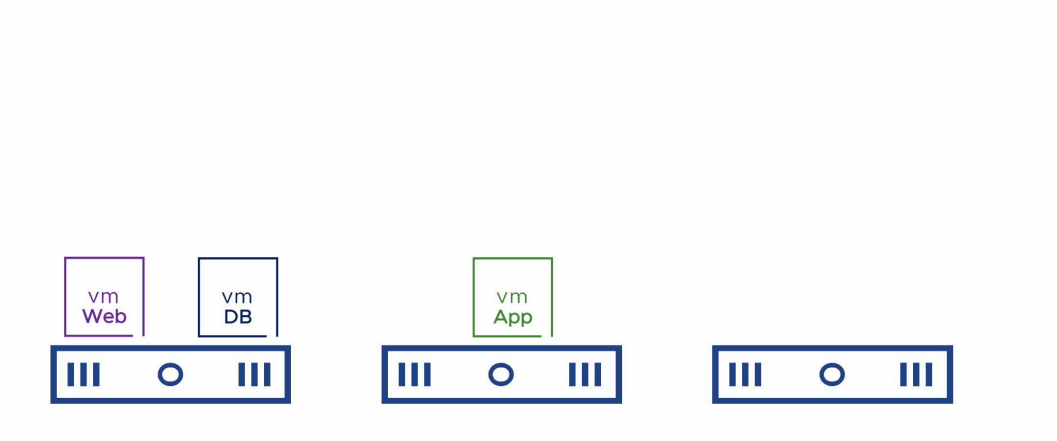

To play the animation, go to https://vmware.bravais.com/s/FPFepiQhbhlhdTfJwW4N

9-78 Configuring Orchestrated Restart

To configure an orchestrated restart:

1. Create VM groups for each of the VMs or VM category.
2. Create VM rules (of type Virtual Machines to Virtual Machines) to define the VM restart dependencies.

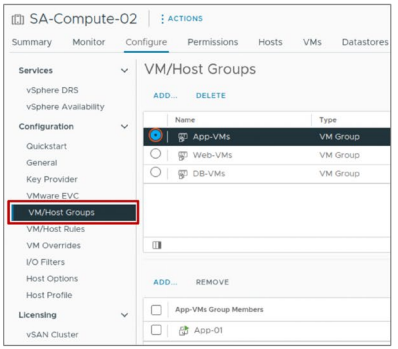

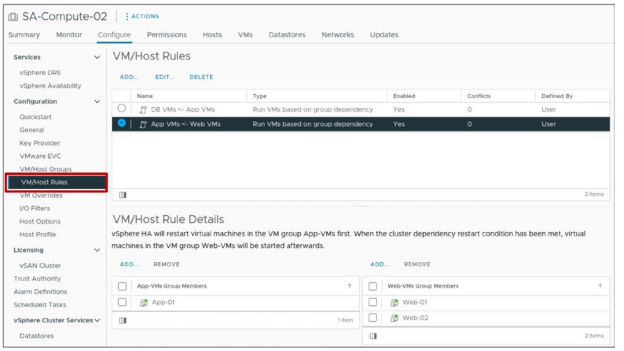

In the orchestrated restart example, the database server VMs restart before the application server VMs and the application server VMs restart before the Web server VMs.

The following VM groups are created:

- App-VMs: Contains the App-01 VM
- Web-VMs: Contains the Web-01 and Web-02 VMs
- DB-VMs: Contains the DB-01 VM

The following VM rules of the Virtual Machines to Virtual Machines type are created:

- DB VMs <- App VMs: Verifies that the VMs in group DB-VMs are restarted before the VMs in group App-VMs
- App VMs <- Web VMs: Verifies that the VMs in the App VMs group are restarted before the VMs in the Web VMs group

To create the DB VMs <- App VMs rule:

1. On the cluster's **Configure** tab, select **VM/Host Rules** and click **ADD**.
2. In the **Name** text box, enter **DB VMs <- App VMs**.
3. From the **Type** drop-down menu, select **Virtual Machines to Virtual Machines**.
4. From the first drop-down menu, select **DB-VMs**.
5. From the second drop-down menu, select **App-VMs**.

The rule must appear as: On restart for VM group DB-VMs, the VM dependency restart condition must be met before continuing to App-VMs.

Create the App-VMs <- Web VMs rule similarly, by selecting **App-VMs** from the first drop-down menu and **Web-VMs** from the second drop-down menu. The rule must appear as: On restart for VM group App-VMs, the VM dependency restart condition must be met before continuing to Web-VMs.

The restart rules that you create enforce the restart order for each VM in the dependency chain. Creating these dependencies increases the likelihood that an impacted application properly recovers when vSphere HA restarts VMs.

9-79 vSphere HA Settings: VM Monitoring

You use VM Monitoring settings to control the monitoring of VMs and applications.

By default, VM and Application Monitoring is set to Disabled.

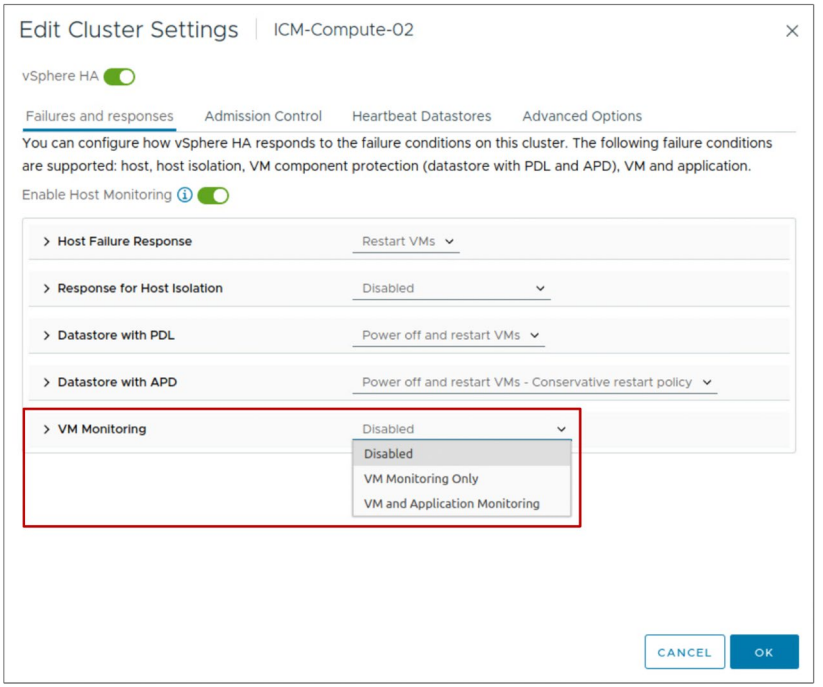

The VM monitoring service determines that the VM has failed if one of the following events occurs:

- VMware Tools heartbeats are not received.
- The guest operating system has not issued an I/O for the last 2 minutes (by default).

If the VM has failed, the VM monitoring service resets the VM to restore services.

You can configure the level of monitoring sensitivity. Highly sensitive monitoring results in a more rapid conclusion that a failure has occurred. Although unlikely, highly sensitive monitoring might lead to falsely identifying failures when the VM or application is still working but heartbeats have not been received because of factors like resource constraints. Low-sensitivity monitoring results in longer interruptions in service between actual failures and VMs being reset. Select an option that is an effective compromise for your needs.

You can select VM and Application Monitoring to activate application monitoring. Application monitoring restarts a virtual machine if the heartbeats for an application it is running are not received.

9-80 vSphere HA Settings: Admission Control

vCenter uses admission control to ensure the following:

- Sufficient resources are available in a cluster to provide failover protection.
- VM resource reservations are respected.

Admission control settings:

- Disabled
- Slot Policy
- Cluster Resource Percentage (default)
- Dedicated Failover Hosts

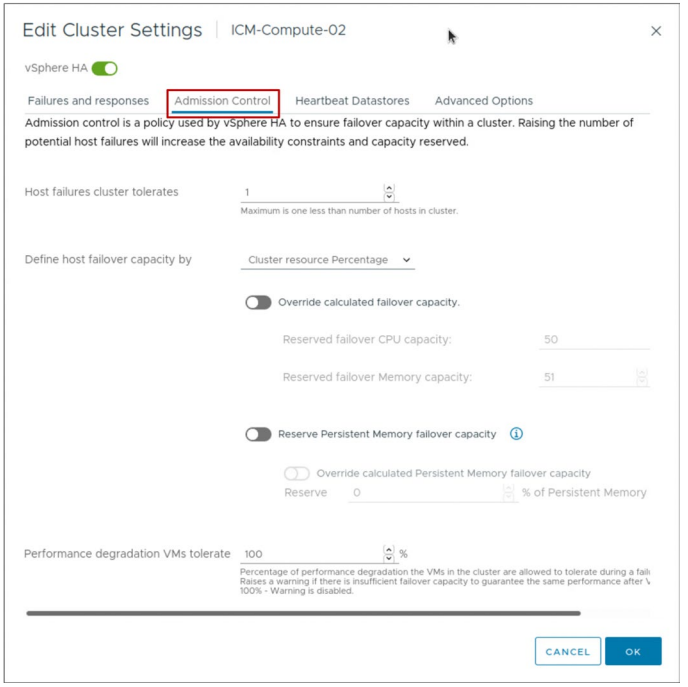

After you create a cluster, you can use admission control to specify whether VMs can be started if they violate availability constraints. The cluster reserves resources to allow failover for all running VMs for a specified number of host failures.

The admission control settings include:

- **Disabled**: This option deactivates admission control, allowing the VMs violating availability constraints to power on (Not recommended).
- **Slot Policy**: A slot is a logical representation of memory and CPU resources. With the slot policy option, vSphere HA calculates the slot size, determines how many slots each host in the cluster can hold, and therefore determines the current failover capacity of the cluster.
- **Cluster Resource Percentage**: (Default) This value specifies a percentage of the cluster's CPU and Memory resources to be reserved as spare capacity to support failovers.
- **Dedicated Failover Hosts**: This option selects hosts to use for failover actions. If a default failover host does not have enough resources, failovers can still occur to other hosts in the cluster.

9-81 Example: Admission Control Using Cluster Resources Percentage

Example of calculating total failover capacity using cluster resource percentages:

- Total cluster capacity:
  - CPU: 18 GHz
  - Memory: 24 GB
- Total VM reservations:
  - CPU: 7 GHz
  - Memory: 6 GB
- Current failover CPU capacity is 61%:
  - ((18 GHz - 7 GHz)/18 GHz) = 61%
- Current failover memory capacity is 75%:
  - ((24 GB - 6 GB)/24 GB) = 75%

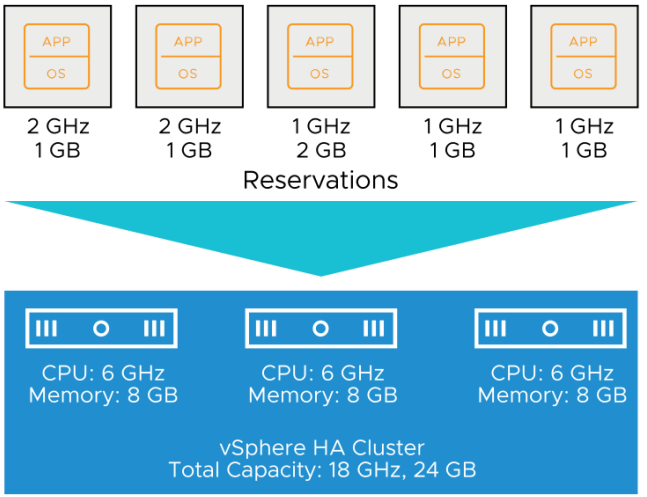

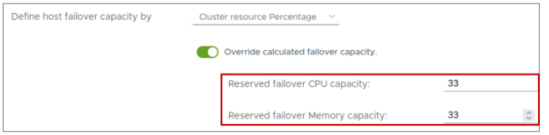

Cluster resource percentage is the default admission control policy. Recalculations occur automatically as the cluster's resources change, for example, when a host is added to or removed from the cluster.

9-82 Example: Admission Control Using Slots (1)

A slot is calculated by combining the largest memory reservation and the largest CPU reservation of any running VM in the cluster.

vSphere HA performs admission control by calculating the following values:

- Slot size:
  - In this example, the slot size is 2 GHz CPU and 2 GB memory.
- Number of slots each host in the cluster can hold:
  - Three
  - The cluster has a total of nine slots (3 + 3 + 3).

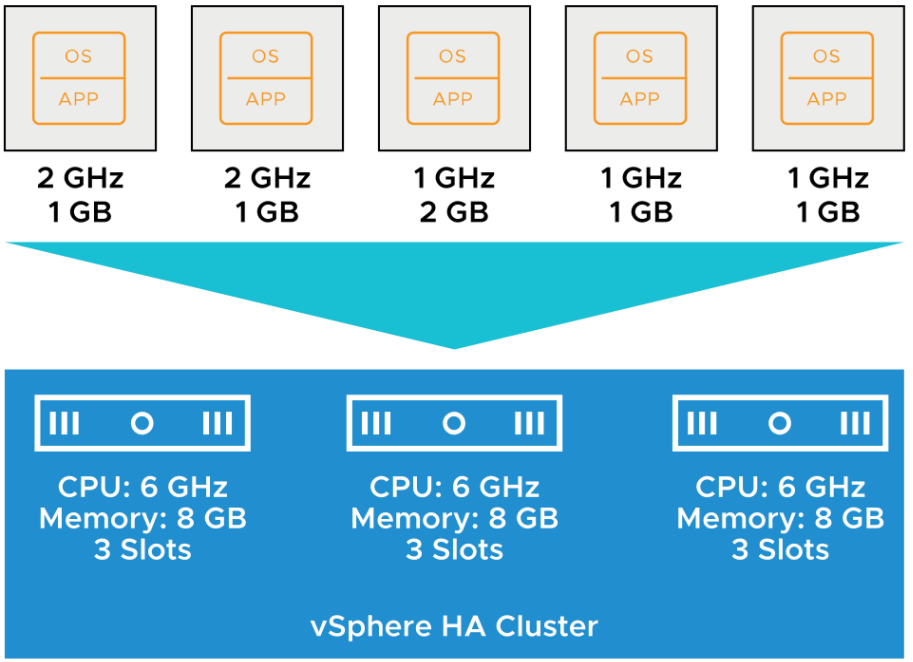

9-83 Example: Admission Control Using Slots (2)

vSphere HA also calculates the current failover capacity. In this example, the failover capacity is one host:

- If the first host fails, six slots remain in the cluster, which is sufficient for all five of the powered-on VMs.
- If the first and second hosts fail, only three slots remain, which is insufficient for all five of the VMs.
- If the current failover capacity is less than the configured failover capacity, vSphere HA does not allow any more VMs to power on.

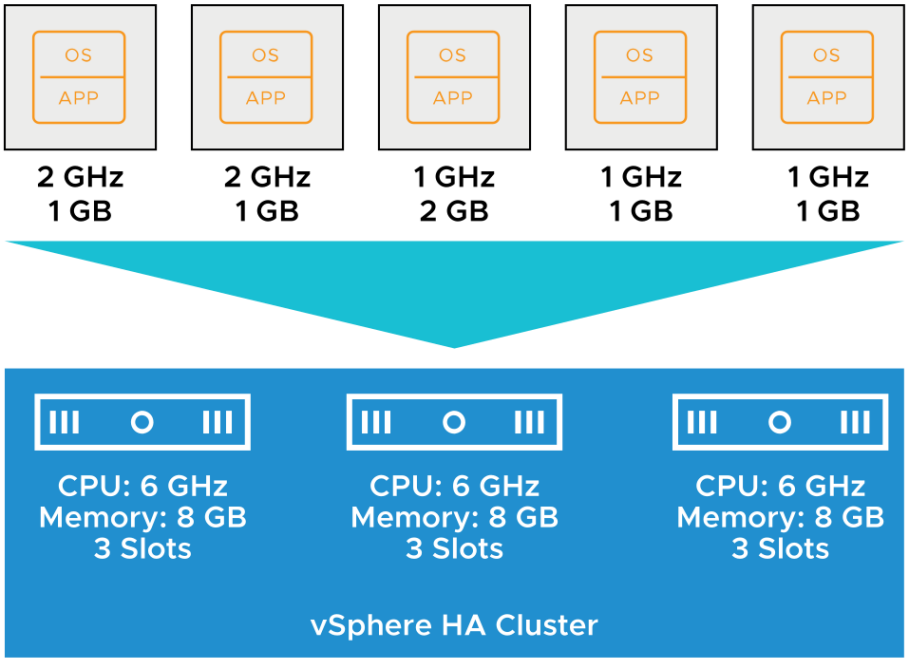

9-84 vSphere HA Settings: Performance Degradation VMs Tolerate

The **Performance degradation VMs tolerate** threshold specifies the percentage of performance degradation that the VMs in the cluster are allowed to tolerate during a failure.

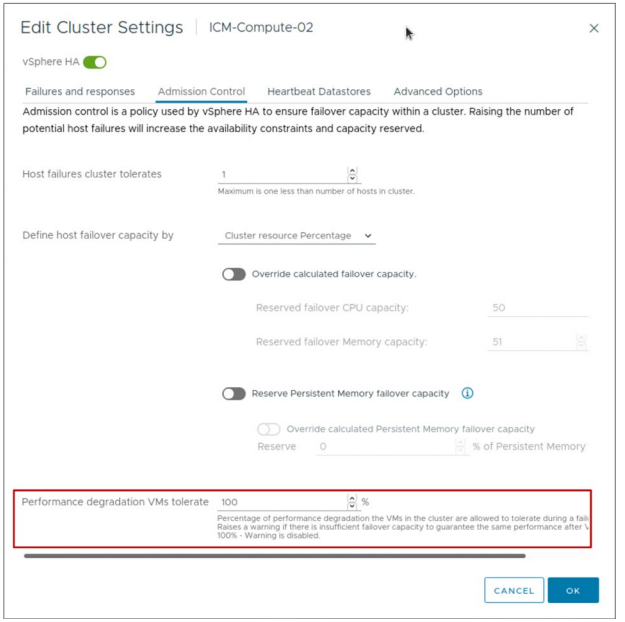

Admission control can also be configured to offer warnings when the actual use exceeds the failover capacity percentage. The resource reduction calculation takes into account a VM's reserved memory and memory overhead.

By setting the **Performance degradation VMs tolerate** threshold, you can specify when a configuration issue should generate a warning or notice. For example:

- The default value is 100 percent, and produces no warnings.
- If you reduce the threshold to 0 percent, a warning is generated when cluster use exceeds the available capacity.
- If you reduce the threshold to 20 percent, the performance reduction that can be tolerated is calculated as performance reduction = current use x 20 percent.

  When the current use minus the performance reduction exceeds the available capacity, a configuration notice is issued.

The **Performance degradation VMs tolerate** threshold is not available unless vSphere DRS is configured.

9-85 vSphere HA Settings: Heartbeat Datastores

A heartbeat file is created on the selected datastores and is used if the heartbeat network fails.

Heartbeat datastores:

- VMFS
- NFS
- vSphere Virtual Volumes

A vSAN datastore cannot be used for datastore heartbeating.

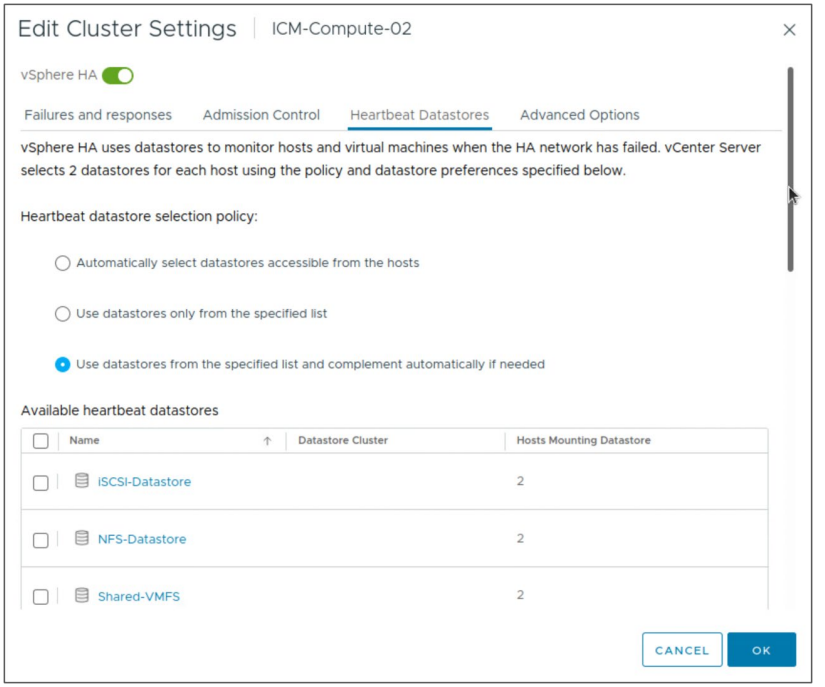

vSphere HA takes checking the health of a host to another level by checking more than the heartbeat network to determine a host's health. You can select which datastores to use for datastore heartbeating, or you can let vSphere HA decide. You can also combine both methods.

9-86 vSphere HA Settings: Advanced Options

You can set advanced vSphere HA options to customize vSphere HA behavior.

Description	Option	Value
Force a cluster not to use the default isolation address (default gateway).	das.usedefaultisolationaddress	`false`
Force a cluster to ping alternate isolation addresses.	das.isolationaddressX	IP address or FQDN
Force a cluster to wait beyond the default 30-second isolation action window.	das.config.fdm.isolationPolicyDelaySec	>=30 seconds
Force maximum bound on the memory slot size.	das.slotmeminmb	100
Force maximum bound on the CPU slot size.	das.slotcpuinmhz	32

You can set advanced options that affect the behavior of your vSphere HA cluster. For more details, see *vSphere Availability* at https://docs.vmware.com/en/VMware-vSphere/index.html.

9-87 Network Configuration and Maintenance

Deactivate host monitoring before modifying virtual networking components that involve the VMkernel ports configured for management or vSAN traffic.

This practice prevents unwanted attempts to fail over VMs.

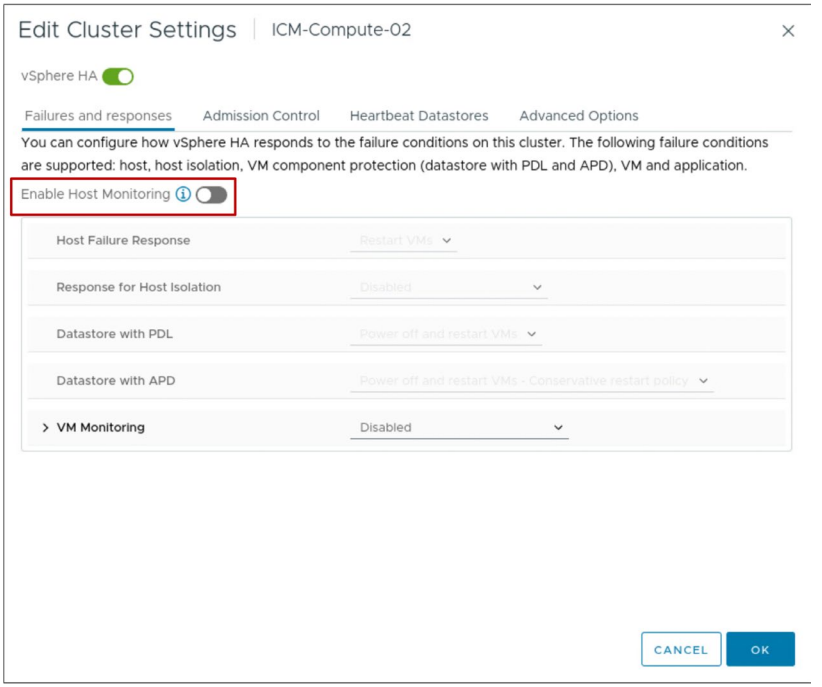

The following network maintenance suggestions can help you avoid the false detection of host failure and network isolation because of dropped vSphere HA heartbeats:

- Changing your network hardware or networking settings can interrupt the heartbeats used by vSphere HA to detect host failures, and might result in unwanted attempts to fail over VMs. When changing the management or vSAN networks of the hosts in the vSphere HA-configured cluster, suspend host monitoring and place the host in maintenance mode.
- Deactivating host monitoring is required only when modifying virtual networking components and properties that involve the VMkernel ports configured for the Management or vSAN traffic, which are used by the vSphere HA networking heartbeat service.
- To recap the previous bullet points, you must reconfigure vSphere HA on all hosts in the cluster if you:
  - Change the networking configuration on ESXi hosts
  - Add port groups
  - Remove virtual switches
  - Suspend host monitoring

  This reconfiguration causes the network information to be reinspected. Then, you must reactivate host monitoring.

9-88 Monitoring vSphere HA Cluster Status

You can monitor the status of a vSphere HA cluster on the Summary page of the **Monitor** tab.

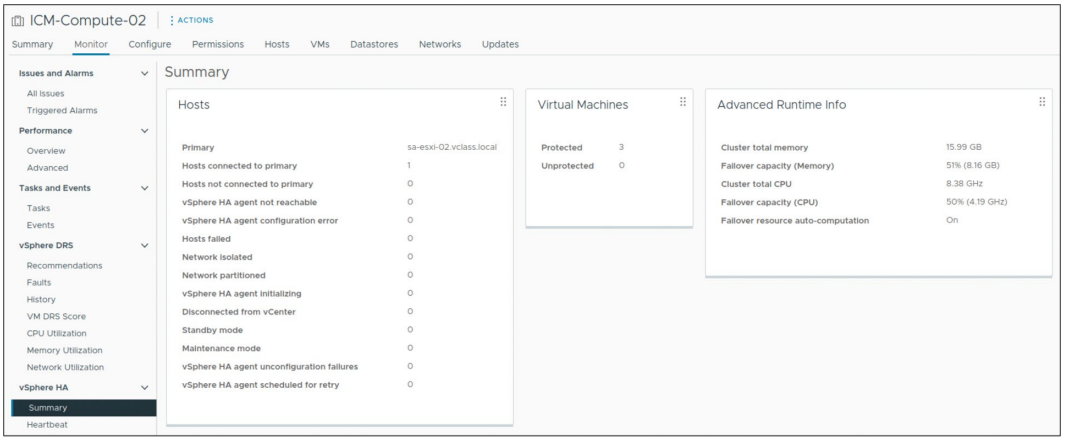

Your cluster or its hosts can experience configuration issues and other errors that adversely affect proper vSphere HA operation.

9-89 Using vSphere HA with vSphere DRS

vSphere HA is closely integrated with vSphere DRS:

- When a failover occurs, vSphere HA checks whether resources are available on each host for the failover.
- If resources are not available, vSphere HA asks vSphere DRS to accommodate for the VMs where possible.

vSphere HA might not be able to fail over VMs for the following reasons:

- vSphere HA admission control is deactivated, and resources are insufficient in the remaining hosts to power on all the failed VMs.

When vSphere HA performs failover and restarts VMs on different hosts, its first priority is the immediate availability of all VMs. After the VMs are restarted, the hosts on which they were restarted are usually heavily loaded, and other hosts are comparatively lightly loaded. vSphere DRS helps to balance the load by migrating VMs between hosts in the cluster.

9-90 About vSphere Fault Tolerance

vSphere Fault Tolerance protects mission-critical, high-performance applications regardless of the operating system used.

vSphere Fault Tolerance provides instantaneous failover and continuous availability:

- Zero downtime
- Zero data loss
- No loss of VM network connectivity

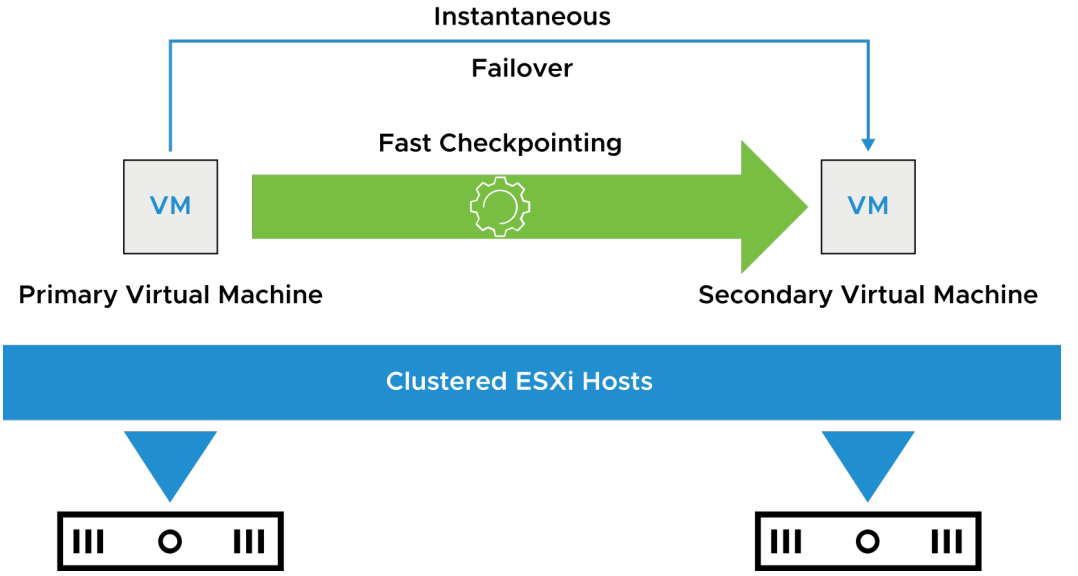

You can use vSphere Fault Tolerance for most mission-critical VMs. vSphere Fault Tolerance is built on the ESXi host platform.

The protected VM is called the primary VM. The duplicate VM is called the secondary VM. The secondary VM is created and runs on a different host to the primary VM. The secondary VM's execution is identical to that of the primary VM. The secondary VM can take over at any point without interruption and provide fault-tolerant protection.

The primary VM and the secondary VM continuously monitor the status of each other to ensure that fault tolerance is maintained. A transparent failover occurs if the host running the primary VM fails, in which case the secondary VM is immediately activated to replace the primary VM. A new secondary VM is created and started, and fault tolerance redundancy is reestablished automatically. If the host running the secondary VM fails, the secondary VM is immediately replaced. In either case, users experience no interruption in service and no loss of data.

9-91 vSphere Fault Tolerance Checkpointing

Changes on the primary VM are note processed on the secondary VM. The memory is updated on the secondary VM.

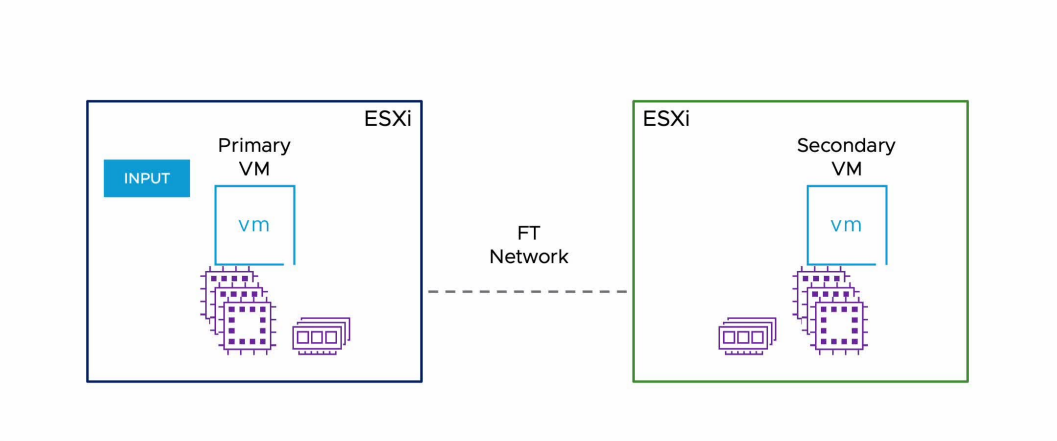

To play the animation, go to https://vmware.bravais.com/s/CAQIyfAA4o5uRM54k6oU.

9-92 vSphere Fault Tolerance with vSphere HA and vSphere DRS

vSphere HA and vSphere DRS are vSphere Fault Tolerance aware:

- vSphere HA is required for vSphere Fault Tolerance.
- vSphere DRS:
  - Selects which hosts run the primary and secondary VM, when a VM is powered on
  - Does not automatically migrate fault-tolerant VMs

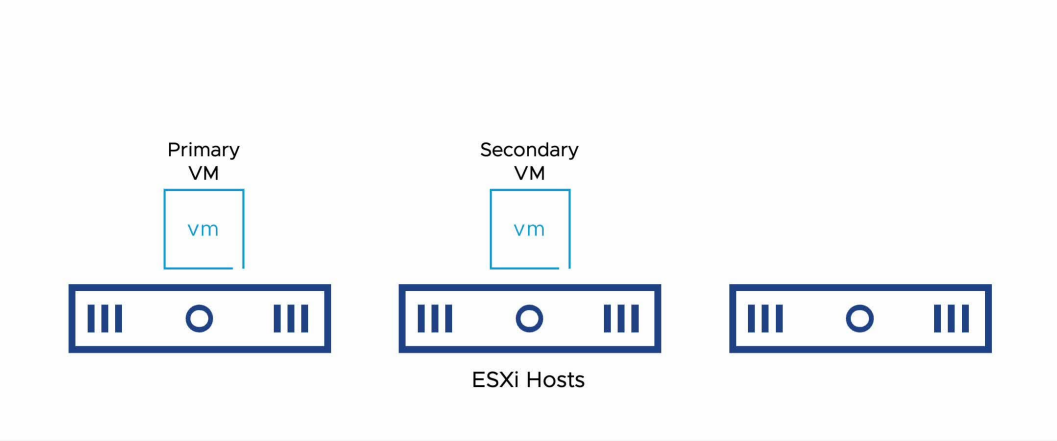

To play the animation, go to https://vmware.bravais.com/s/uZwjd73t46eXD69ZqXdW.

A fault-tolerant VM and its secondary copy are not allowed to run on the same host. The restriction ensures that a host failure cannot result in the loss of both the VMs. In addition, vSphere vMotion cannot be used to migrate the VMs to the same host.

9-93 vSphere Fault Tolerance Features

vSphere Fault Tolerance:

- Supports VMs configured with up to 8 vCPUs and 128 GB memory
- Supports up to four fault-tolerant VMs per host with no more than eight vCPUs between them
- Supports vSphere vMotion migration for primary and secondary VMs
- Creates a secondary copy of all VM files and disks
- Supports multiple VM disk formats:
  - Thin provision
  - Thick provision lazy-zeroed
  - Thick provision eager-zeroed
- Supports interoperability with vSAN
- Provides fast checkpoint copying to keep primary and secondary VMs synchronized

For details on vSphere Fault Tolerance, see *vSphere Availability* at https://docs.vmware.com/en/VMware-vSphere/index.html.

9-94 Configuring vSphere Fault Tolerance on a VM

You can turn on vSphere Fault Tolerance for a VM using the vSphere Client.

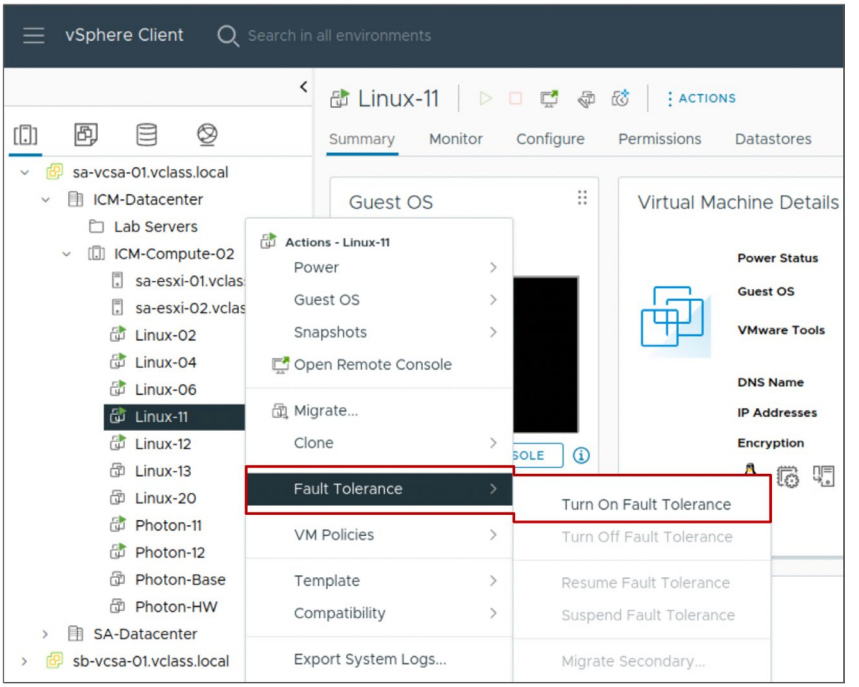

After you perform the required steps for configuring vSphere Fault Tolerance for your cluster, you can use the feature by turning it on for individual VMs.

Before vSphere Fault Tolerance can be turned on, validation checks are performed on a VM. When vSphere Fault Tolerance is turned on, vCenter resets the VM's memory limit to the default (unlimited memory) and sets the memory reservation to the memory size of the VM. When vSphere Fault Tolerance is turned on, you cannot change the memory reservation, size, limit, number of virtual CPUs, or shares. You also cannot add or remove disks for the VM. When vSphere Fault Tolerance is turned off, any parameters that were changed are not reverted to their original values.

In addition to turning on vSphere Fault Tolerance for a VM, additional options are available:

- **Turn Off Fault Tolerance**: Deletes the secondary virtual machine, its configuration, and all history.
- **Suspend Fault Tolerance**: Suspends the vSphere Fault Tolerance protection, but preserves the secondary VM, its configuration, and all history. Use this option if at a future time, you want to resume vSphere Fault Tolerance protection on the VM.
- **Migrate Secondary**: Migrates the secondary VM to the specified host.
- **Test Failover**: Induces a failover situation for the selected primary VM to test your vSphere Fault Tolerance protection. The option is unavailable if the virtual machine is powered off.
- **Test Restart Secondary**: Induces the failure of a secondary VM to test the vSphere Fault Tolerance protection provided for the selected primary VM. The option is unavailable if the virtual machine is powered off.

9-95 Lab 25: Configuring vSphere HA

Configure vSphere HA and test its functionality:

1. Configure vSphere HA in a Cluster
2. View Information About the vSphere HA Cluster
3. Configure Network Management Redundancy
4. Test the vSphere HA Functionality
5. View the vSphere HA Cluster Resource Usage
6. Configure the Percentage of Resource Degradation to Tolerate

9-96 Review of Learner Objectives

- Recognize the requirements for creating and using a vSphere HA cluster
- Recognize the use cases for various vSphere HA settings
- Recognize when to use vSphere Fault Tolerance
- Configure a vSphere HA cluster

9-97 Key Points

- When you create a cluster, you can configure vSphere DRS, vSphere HA, vSAN, and the ability to manage image updates on all hosts collectively.
- vSphere DRS clusters provide automated resource management to ensure that VMs' resource requirements are satisfied.
- vSphere DRS works best when the VMs meet vSphere vMotion migration requirements.
- vSphere HA protects against various types of failures: host, guest OS, application, datastore accessibility, and network isolation.
- You implement redundant heartbeat networks either with NIC teaming or by creating additional heartbeat networks.
- vSphere Fault Tolerance provides zero downtime for applications that must always be available.

Questions?

Module 10

Managing the vSphere Lifecycle

10-2 Importance

Managing the life cycle of vSphere involves keeping vCenter and ESXi hosts up to date and integrated with other VMware and third-party solutions. It also involves keeping VMware Tools and virtual hardware in a VM up to date. To achieve these goals, you must understand how to keep vCenter up-to-date and how to use the features provided by vSphere Lifecycle Manager, namely, cluster-level management of ESXi hosts.

10-3 Module Lessons

1. Managing the vCenter Life Cycle
2. Overview of vSphere Lifecycle Manager
3. Managing the Life Cycle of ESXi Hosts and Clusters
4. Managing the Life Cycle of VMware Tools and VM Hardware

10-4 Lesson 1: Managing the vCenter Life Cycle

10-5 Learner Objectives

- Plan for vCenter updates and upgrades by generating interoperability reports
- Recognize how to update vCenter with the latest patches, updates, and upgrades

10-6 About Upgrades, Updates, and Patches

vSphere products distinguish between upgrades, updates, and patches:

Upgrades:

- Major change to the software
- Either digit in the release number changes, for example, 6.7, 7.0, 8.0

Updates and Patches:

- Smaller changes to the software
- Update number is added to release number, for example, 7.0 Update 3

10-7 Planning for vCenter Updates and Upgrades

To plan for vCenter updates and upgrades, you can produce interoperability reports about VMware products associated with vCenter.

You check VMware products against both the installed version of vCenter and the version to which you plan to upgrade.

To generate interoperability reports, you must first join the VMware Customer Experience Improvement Program (CEIP).

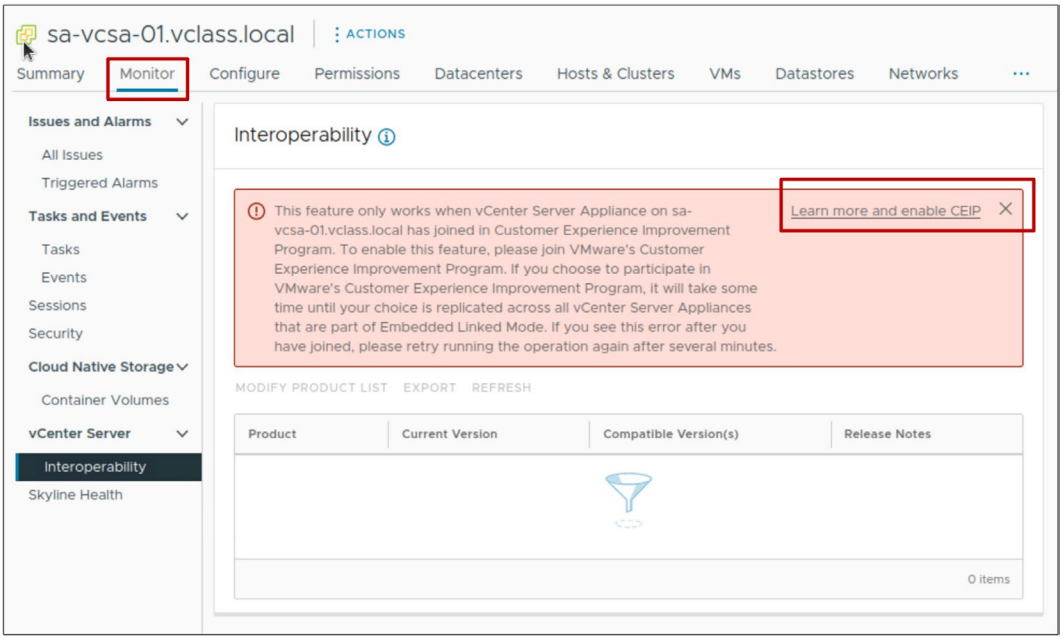

The Customer Experience Improvement Program (CEIP) is a program that provides you with features to help you proactively manage your vSphere environment. The vCenter interoperability report, also referred to as vCenter Server Update Planner, is one of the features that CEIP offers.

To join CEIP, click the **Learn more and enable CEIP** link in the message and click **JOIN PROGRAM**.

Alternatively, you can join or leave the program using the following steps:

1. From the main menu, select **Administration**.
2. Under **Deployment**, select **Customer Experience Improvement Program**.
3. Click **JOIN PROGRAM** (or **LEAVE PROGRAM**).

You can find more information on VMware's CEIP by visiting https://www.vmware.com/solutions/trustvmware/ceip.html

10-8 Generating the Interoperability Report

The Interoperability page on the vCenter instance's **Monitor** tab shows VMware products that are currently registered with vCenter and their compatibility with the current version of vCenter.

You can export report results in CSV format and use the report as a guide to prepare for an update.

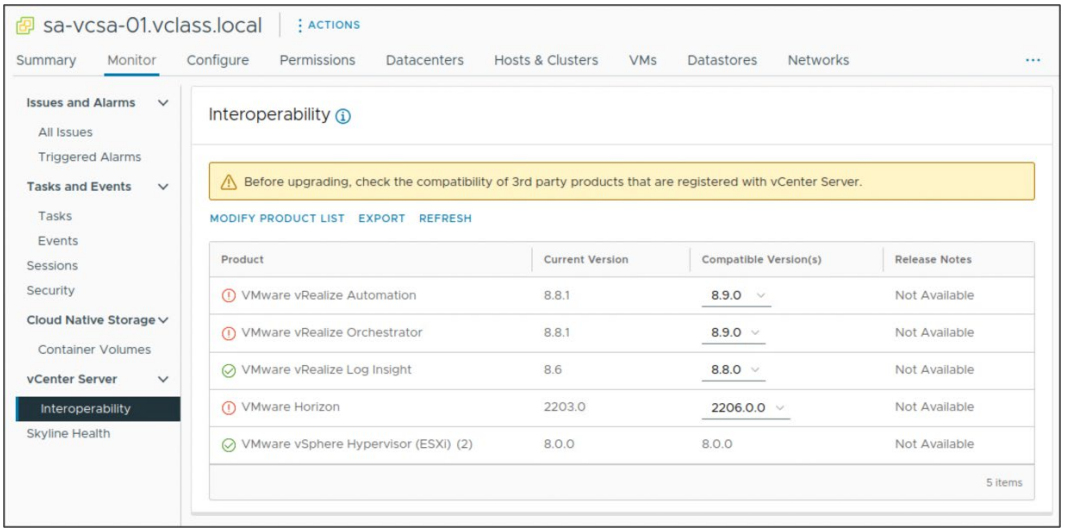

In the vSphere Client, the Interoperability page appears on the **Monitor** tab of vCenter. This page displays VMware products currently registered with vCenter.

This report shows the product name, the current version of the product, and the product version that is compatible with the current version of vCenter.

If VMware products in your environment are undetected, you can click **MODIFY PRODUCT LIST** to manually add them to the list of products to check and regenerate the interoperability report.

10-9 Updating and Patching vCenter

To manage the life cycle of vCenter, use the vCenter Management Interface to update and patch vCenter.

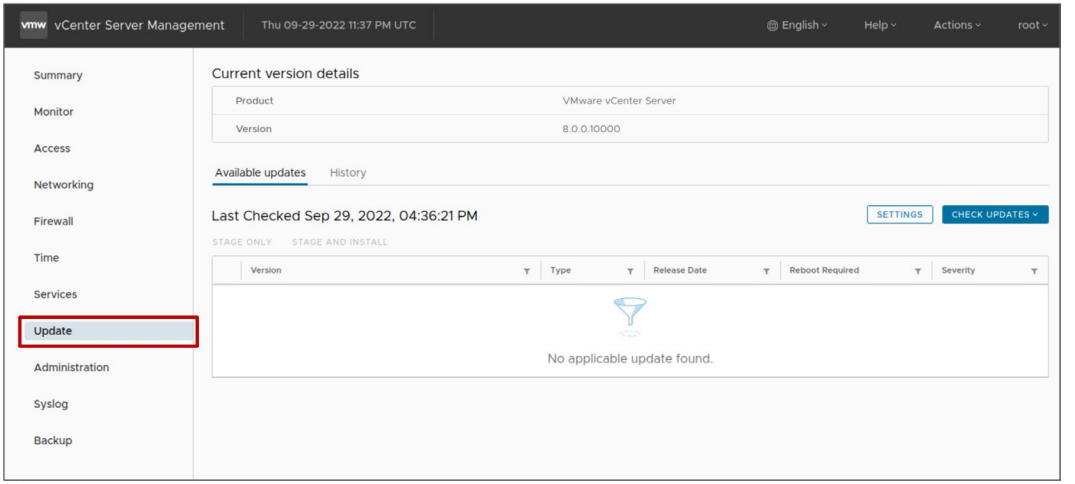

The Update pane lists the vCenter updates that you can select. Details include release date, version, severity, and other information about each vCenter update available. The Type column tells you if the release item is an update, an upgrade, or a patch. If multiple versions appear, the recommended version is preselected.

vCenter updates and patches are located in a default VMware repository URL. You can configure the vCenter Server Appliance to use the default VMware repository URL or a custom repository URL, for example, a repository URL that you previously built on a local Web server running within your data center.

For details about upgrading and patching vCenter, see vCenter Server Upgrade at https://docs.vmware.com/en/VMware-vSphere/index.html.

10-10 Upgrading vCenter Server Appliance

You use the vCenter installer to upgrade vCenter Server Appliance to a newer version.

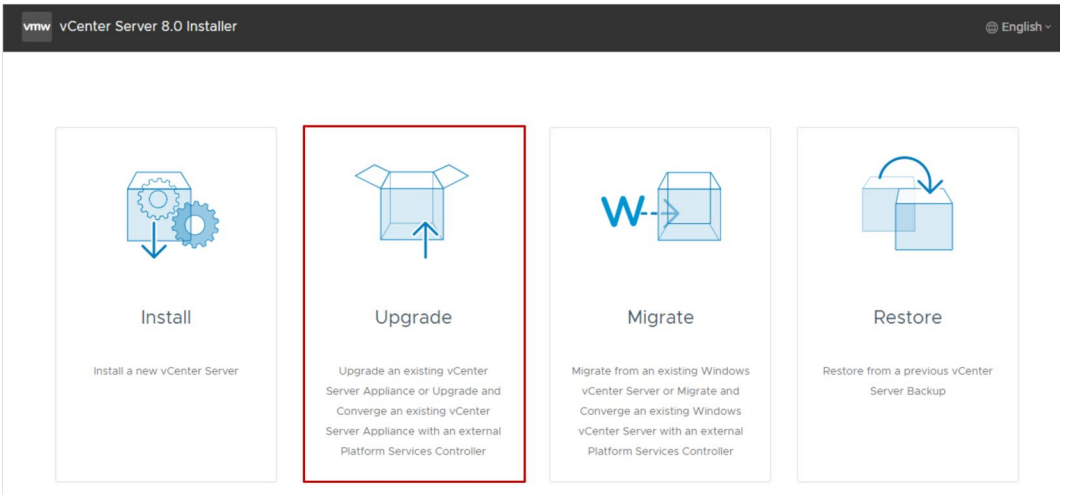

For details about upgrading vCenter, see *vCenter Server Upgrade* at https://docs.vmware.com/en/VMware-vSphere/index.html.

10-11 Overview of the vSphere Upgrade Process

vSphere is a product with multiple components to upgrade.

Knowing the required sequence of tasks is vital for a successful vSphere upgrade.

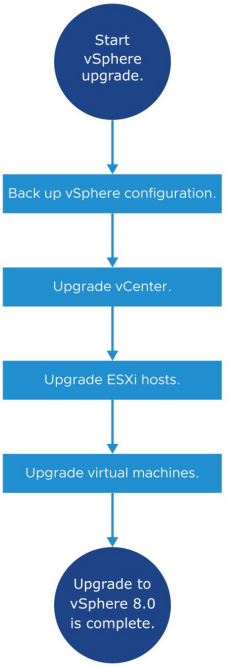

Upgrading vSphere includes the following tasks:

1. Read the vSphere release notes.
2. Verify that you backed up your vSphere configuration.
3. If the vSphere environment includes VMware solutions or plug-ins, verify that they are compatible with the vCenter version to which you upgrade.
4. Upgrade vCenter.
5. Upgrade your ESXi hosts.
6. Upgrade your VMs manually or by using vSphere Lifecycle Manager.

For details about upgrading vCenter, see *vCenter Server Upgrade* at https://docs.vmware.com/en/VMware-vSphere/index.html.

For details about upgrading ESXi, see *VMware ESXi Upgrade* at https://docs.vmware.com/en/VMware-vSphere/index.html.

For details about upgrading VMs with vSphere Lifecycle Manager, see *Managing Host and Cluster Lifecycle* at https://docs.vmware.com/en/VMware-vSphere/index.html.

10-12 Review of Learner Objectives

- Plan for vCenter updates and upgrades by generating interoperability reports
- Recognize how to update vCenter with the latest patches, updates, and upgrades

10-13 Lesson 2: Overview of vSphere Lifecycle Manager

10-14 Learner Objectives

- Recognize features of vSphere Lifecycle Manager
- Import images into the vSphere Lifecycle Manager image depot
- Change the download source for patches and updates

10-15 About vSphere Lifecycle Manager

vSphere Lifecycle Manager enables centralized and simplified life cycle management for ESXi hosts in a cluster through the use of images.

vSphere Lifecycle Manager includes the following tasks:

- Managing VMware Tools and VM hardware upgrades
- Upgrading and patching ESXi hosts
- Installing and updating third-party software on ESXi hosts
- Installing and updating ESXi drivers and firmware
- Standardizing ESXi images across hosts in a cluster

From the main menu, select **Lifecycle Manager**.

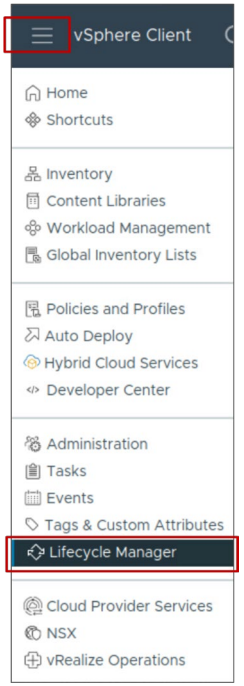

10-16 About Images

Managing clusters with images helps to standardize the software running on your ESXi hosts.

An ESXi image consists of several elements:

- ESXi base image: An update that provides software fixes and enhancements
- Components: A logical grouping of one or more VIBs (vSphere Installation Bundles) that encapsulates functionality in ESXi
- Vendor add-ons: Sets of components that OEMs create and distribute
- Firmware and Drivers Add-On: Firmware and driver bundles that you can define for your cluster image
  - Requires the Hardware Support Manager plug-in for the desired server family

To maintain consistency, you apply a single ESXi image to all hosts in a cluster.

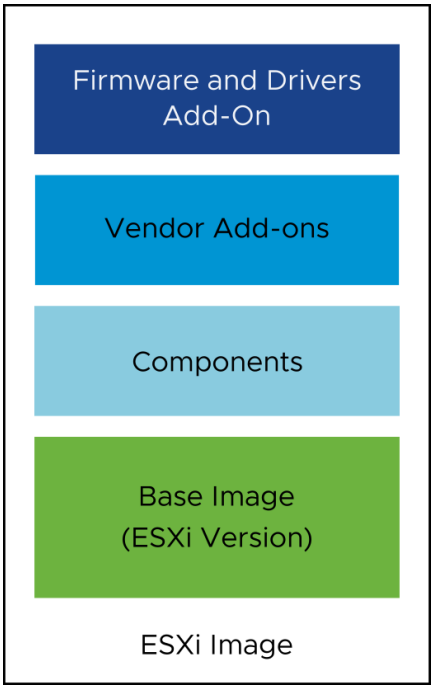

The ESXi base image is a complete ESXi installation package and is enough to start an ESXi host. Only VMware creates and releases ESXi base images.

The ESXi base image is a grouping of components. You must select at least the base image or vSphere version when creating a cluster image.

The component is the smallest unit that is used by vSphere Lifecycle Manager to install VMware and third-party software on ESXi hosts. Components are the basic packaging for VIBs and metadata. The metadata provides the name and version of the component.

On installation, a component provides you with a visible feature. For example, a third-party vendor's network driver is provided as a component. Components are optional elements to add to a cluster image.

Vendor add-ons are custom OEM images. Each add-on is a collection of components customized for a family of servers. OEMs can add, update, or remove components from a base image to create an add-on. Selecting an add-on is optional.

The firmware and drivers add-on is a vendor-provided add-on. It contains the components that encapsulate firmware and driver update packages for a specific server type. To add a firmware and drivers add-on to your image, you must first install the Hardware Support Manager plug-in for the respective family of servers.

10-17 About Image Depots

The vSphere Lifecycle Manager image depot represents all software available for consumption to vSphere Lifecycle Manager:

- The depot is located on the vCenter system.

In the **Image Depot** tab, you can view details about downloaded content:

- ESXi base images
- Vendor add-ons
- Third-party components

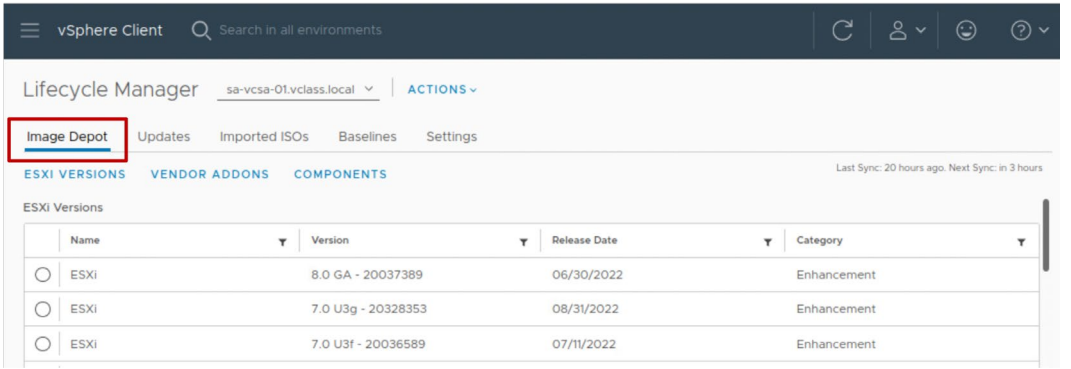

When you select a downloaded file, the details appear to the right:

- When you select an ESXi version, the details include the version name, build number, category, and description, and the list of components that make up the base image.
- When you select a vendor add-on, the details include the add-on name, version, vendor name, release date, category, and the list of added or removed components.
- When you select a component, the details include the component name, version, publisher, release date, category, severity, and contents (VIBs).

Downloading host updates and related metadata is a predefined automatic process that you can modify. The automatic download task is enabled by default and starts immediately after you deploy vCenter. After the initial download, the task runs according to its schedule.

10-18 Importing Content Into the Image Depot from Online Sources

At regular intervals, vSphere Lifecycle Manager downloads updates from configured download sources to the image depot.

Regardless of the download schedule, you can manually initiate synchronization between the image depot and the download sources.

From the **ACTIONS** drop-down menu in the Lifecycle Manager pane, select **Sync Updates**.

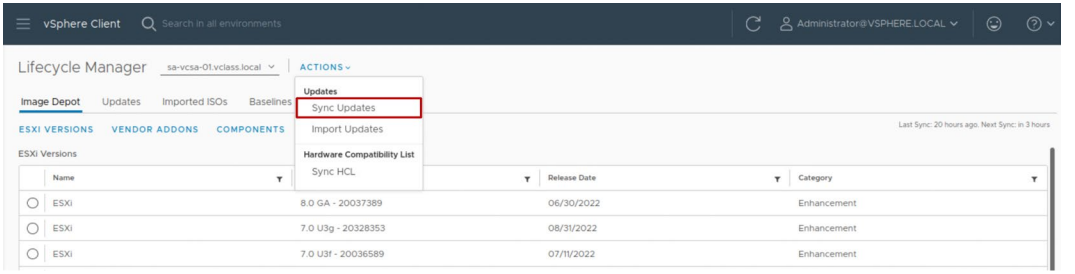

An online depot is the hosted version of the software updates that VMware, OEMs, and third-party software providers ship. You enable vSphere Lifecycle Manager to access an online depot by providing a URL to that depot.

An Internet connection is required to access an online depot.

10-19 Specifying the Download Source

Select **Settings** > **Patch Setup** to view the default download sources.

From this pane, you can change the download source or add a URL to configure a custom download source.

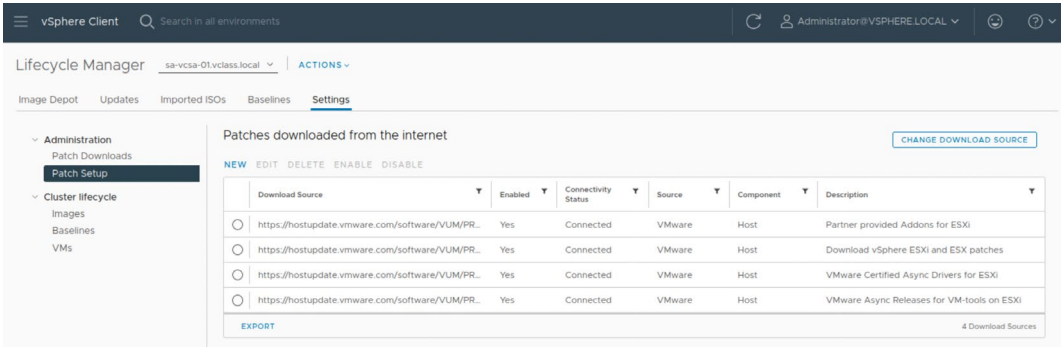

You can click **NEW** to add other download sources besides the VMware download sources. For example, third-party vendors might provide a download source (depot) that contains images for additional components such as CIM modules.

However, working with additional third-party depots and independent components is rarely necessary. In most cases, the vendor add-ons that are available in the official VMware depot provide full OEM customization for ESXi.

vSphere Lifecycle Manager downloads to the local vSphere Lifecycle Manager depot the content from all the online depots that you configure it to use.

10-20 Importing Content Into the Image Depot from Offline Sources

You can also import updates from an offline bundle:

- From the **Actions** drop-down menu, select **Import Updates**.
- Enter a URL or browse for a ZIP file that contains the update (an ESXi image, vendor add-on or component).

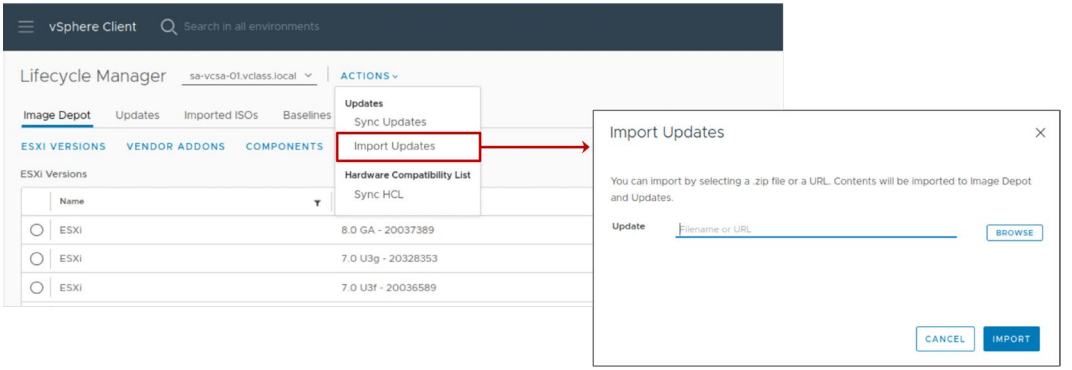

You use the import option to populate the vSphere Lifecycle Manager depot with updates from an offline bundle, for example, one provided by VMware. Apart from the legacy patches and extensions, an offline bundle can also contain an ESXi base image, a vendor add-on, or third-party software, for example, asynchronous drivers specific to the OEM hardware requirements.

10-21 Review of Learner Objectives

- Recognize features of vSphere Lifecycle Manager
- Import images into the vSphere Lifecycle Manager image depot
- Change the download source for patches and updates

10-22 Lesson 3: Managing the Life Cycle of ESXi Hosts and Clusters

10-23 Learner Objectives

- Activate vSphere Lifecycle Manager in a cluster and define a cluster image
- Validate ESXi host compliance against a cluster image
- Remediate ESXi hosts using vSphere Lifecycle Manager
- View recommended images for a cluster

10-24 Creating a Cluster and Specifying an Image

A vSphere Lifecycle Manager cluster is a cluster of ESXi hosts that you can manage with a single image.

When creating a cluster, you can select a cluster image:

1. Create a cluster.
2. Select the **Manage all hosts in the cluster with a single image** check box.
3. Choose the image to use for the cluster.

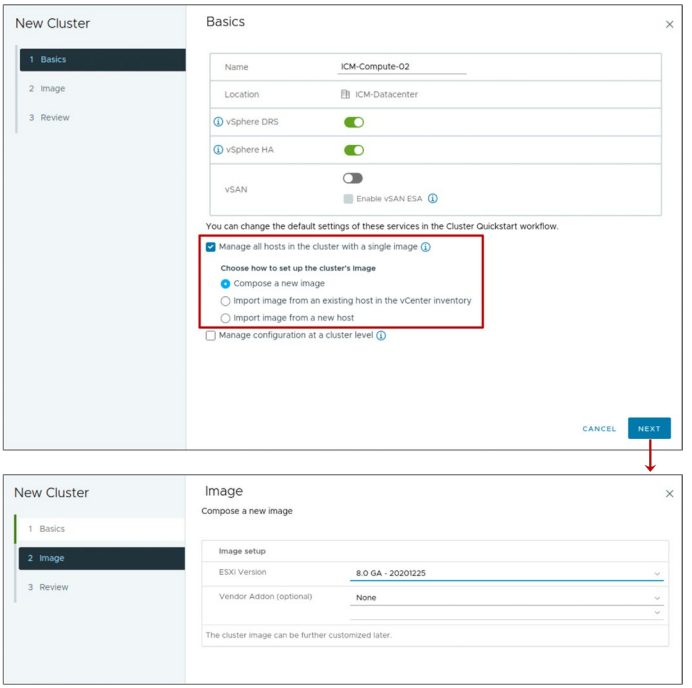

In the Create Cluster wizard, you specify an image to be applied to the hosts. You can import an image from an ESXi host that is in the same or a different vCenter instance. You can also import an image from an ESXi host that is not managed by vCenter:

- **Import image from an existing host in the vCenter inventory**: Select this method to import an image from a host that is in the same vCenter inventory.
- **Import image from a new host**: Select this method to import an image from a host that is in a different vCenter instance or a standalone host that is not added to a vCenter instance. You can use the option to move the selected host to the cluster.

10-25 Viewing Cluster Image Information

The Image pane in the **Updates** tab shows the image for the cluster.

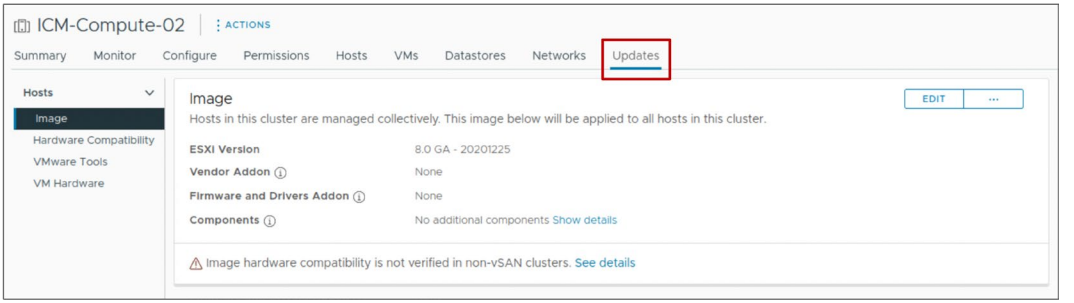

10-26 Overview of Managing Clusters with vSphere Lifecycle Manager

Over time, software updates become available in the image depot.

Using images from the image depot to update ESXi hosts is a multi-stage process:

1. Check the compliance of the ESXi hosts in the cluster against the image specification.
2. Run a remediation pre-check on the hosts in the cluster to ensure software and hardware compatibility with the image.
3. Remediate the non-compliant ESXi hosts in the cluster.

10-27 Checking Compliance of Hosts Against an Image

After you define a cluster image, you can perform a compliance check to compare the cluster image with the software and firmware that runs on the ESXi hosts in the cluster.

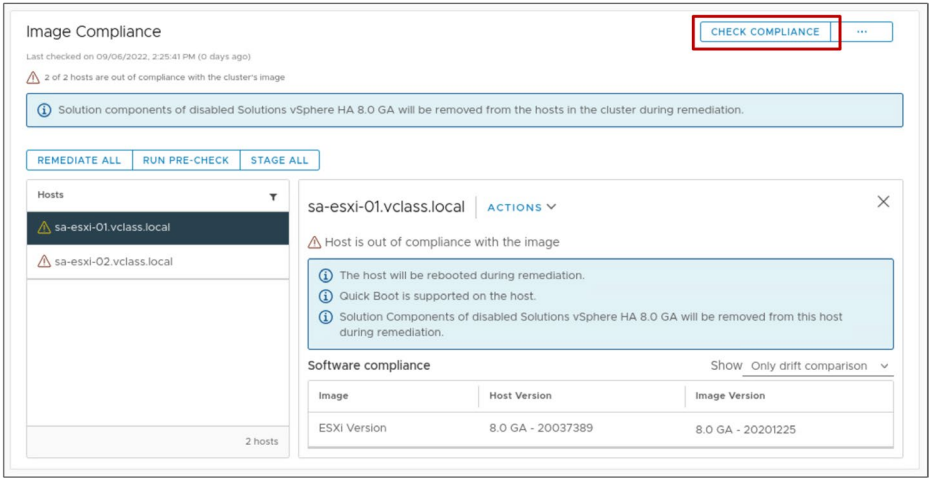

The status of a host can be unknown, compliant, out of compliance, or not compatible with the image.

- A host status is unknown before you check compliance.
- A compliant host is one whose installed software and firmware matches the ESXi image defined for the cluster and with no standalone VIBs or differing components.
- If the host is out of compliance, a message about the impact of remediation appears. In the example, the host must be rebooted as part of the remediation. Another impact that might be reported is the requirement that the host enters maintenance mode.
- A host is not compatible if it is running software and firmware that is newer than the desired cluster image version, or if the host does not meet the installation requirements for the vSphere build.

You can check the image compliance at the level of various vCenter objects:

- At the host level for a specific ESXi host in a cluster
- At the cluster level for all ESXi hosts in the cluster
- At the data center level for all clusters and hosts in the data center
- At the vCenter level for all data centers, clusters, and ESXi hosts in the vCenter inventory.

10-28 Running a Remediation Pre-check

To ensure that the cluster's health is good and that no problems occur during the remediation process of your ESXi hosts, you can perform a remediation pre-check.

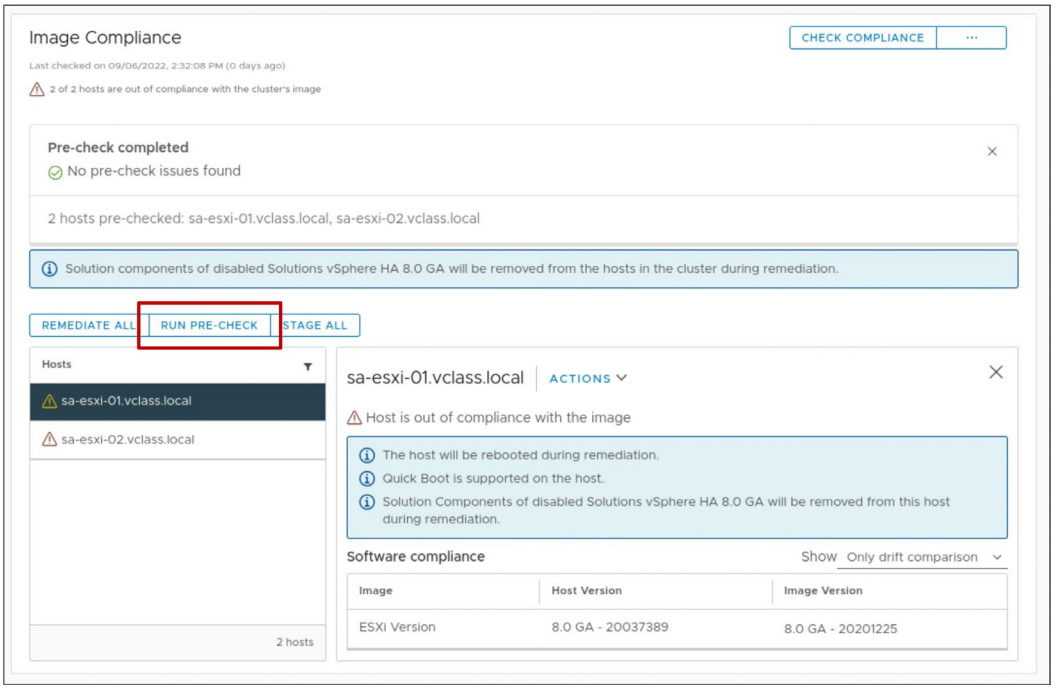

10-29 Staging the Cluster

Staging is the process during which vSphere Lifecycle Manager downloads patches and extensions from the image depot to the ESXi hosts.

During staging, the patches and extensions are not installed on the host.

Staging patches and extensions speeds up the remediation process, because the patches and extensions are already available locally on the hosts.

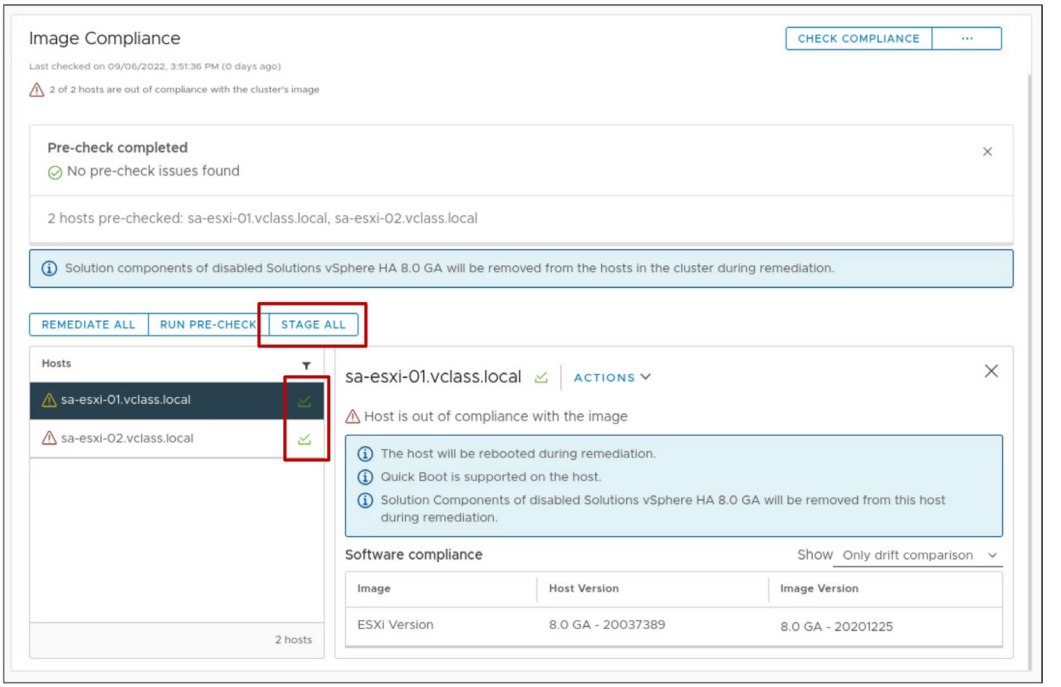

A green icon next to a host indicates that the current image has been staged to the host.

Staging patches and extensions speeds up the remediation process and minimizes the time that each host spends in maintenance mode.

Without staging, the patches and extensions are transferred to the host during maintenance mode. As a result, the host is unavailable to run VMs for a longer period of time than if the updates were staged beforehand.

10-30 Remediating a Cluster Against an Image

Remediation makes the selected hosts compliant with the cluster image.

You can remediate the entire cluster or a single ESXi host, or simply pre-check hosts without updating them.

The goal is to make the entire cluster compliant with the image.

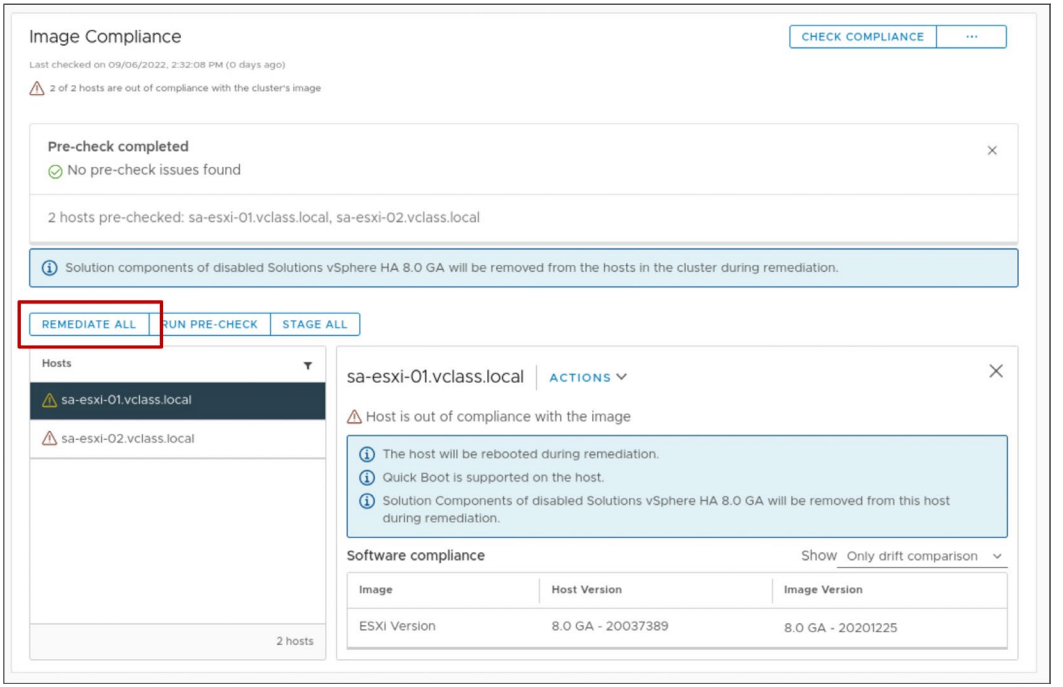

vSphere Lifecycle Manager can perform a remediation pre-check on images. If the pre-check is successful, vSphere Lifecycle Manager applies the image to the hosts.

During each step of a remediation process, vSphere Lifecycle Manager determines the readiness of the host to enter or exit maintenance mode or be rebooted.

When you remediate a cluster that you manage with an image, vSphere Lifecycle Manager applies the following elements to the ESXi hosts:

- ESXi image version
- Optional: vendor addon
- Optional: firmware and driver addons
- Optional: user specified components

10-31 Reviewing Remediation Impact

The Review Remediation Impact dialog box contains information about all changes that remediation performs on the hosts in the cluster.

vSphere Lifecycle Manager performs a remediation pre-check before each remediation.

If the pre-check is successful, vSphere Lifecycle Manager applies the image to the hosts.

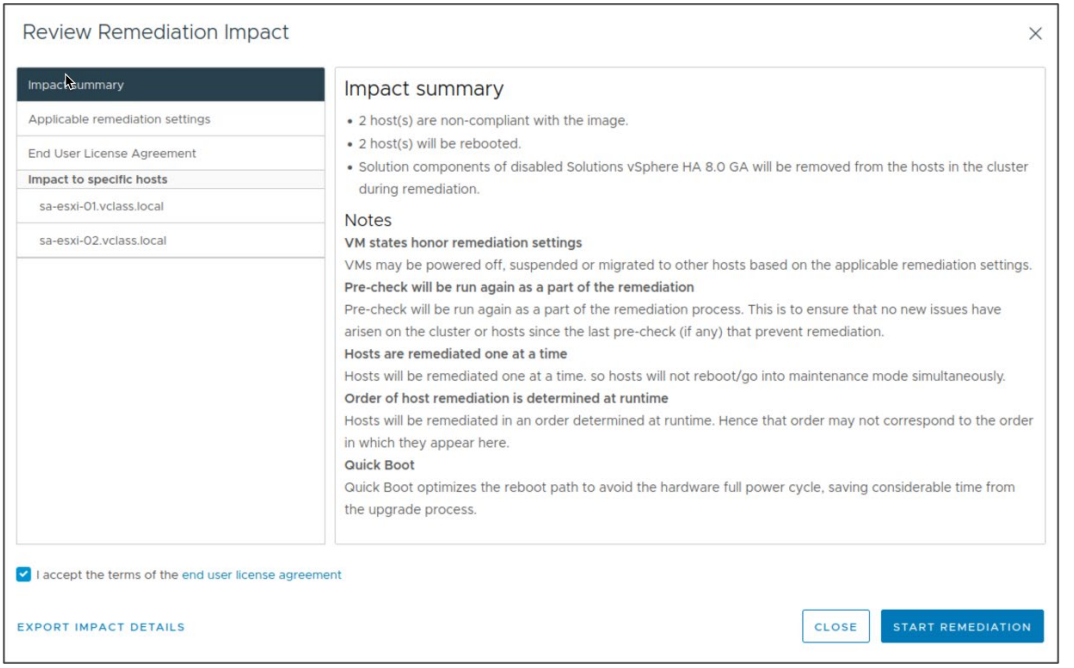

The Review Remediation Impact dialog box includes the following information:

- Impact summary
- Applicable remediation settings
- End User License Agreement
- Impact to specific hosts

10-32 Parallel Remediation

You can configure vSphere Lifecycle Manager to perform parallel remediation in a cluster.

Parallel remediation remediates all hosts that are in maintenance mode in parallel instead of in sequence.

You can also specify the maximum number of concurrent remediations.

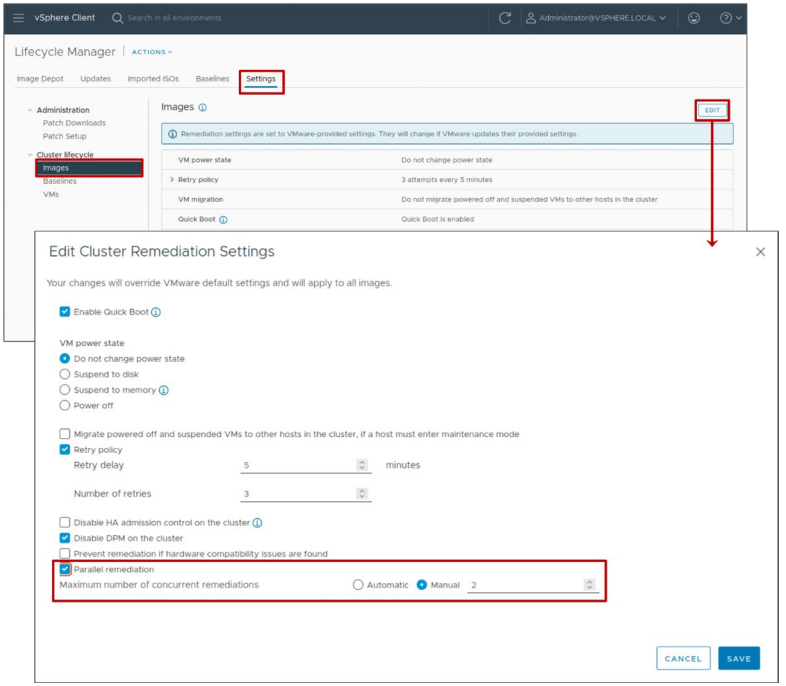

When upgrading a cluster with vSphere Lifecycle Manager, by default, ESXi hosts are remediated sequentially, which can be a time-consuming process. To speed up the cluster upgrade process and reduce the downtime window, you can configure

vSphere Lifecycle Manager to remediate hosts in parallel using the cluster image. The hosts must be in maintenance mode. Hosts that are not in maintenance mode are skipped, but you can upgrade them separately later.

10-33 vSphere Lifecycle Manager Integration with vSphere DRS

When performing remediation operations on a cluster that is configured for vSphere DRS, vSphere Lifecycle Manager automatically integrates with vSphere DRS:

- When vSphere Lifecycle Manager places hosts in the maintenance mode, vSphere DRS evacuates each host before the host is patched.
- When vSphere Lifecycle Manager attempts to place a host in the maintenance mode, certain checks are performed to verify that the ESXi host can enter the maintenance mode.
- The vSphere Client reports any configuration problems that might prevent an ESXi host from entering the maintenance mode.
- After a host is patched and rebooted, vSphere Lifecycle Manager exits the host from the maintenance mode and vSphere DRS migrates some VMs back to the host.

10-34 Recommended Images

Using vSphere Lifecycle Manager, you can check for recommended images for a cluster that you manage with an image.

vSphere Lifecycle Manager checks for compatibility across the image components. The process verifies that the recommended image has no missing dependencies or conflicting components.

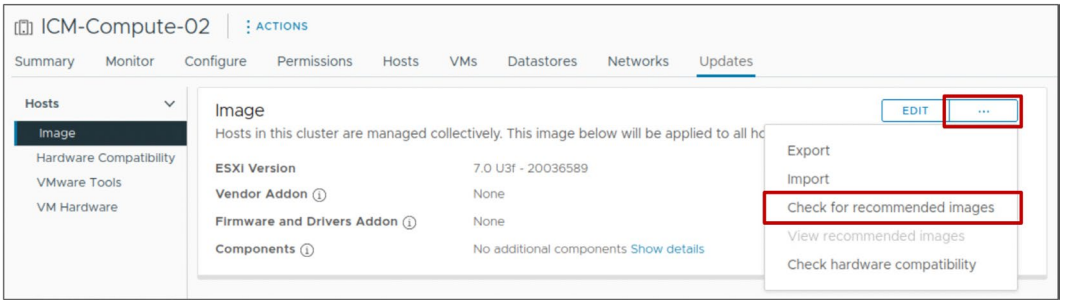

To generate recommendations, vSphere Lifecycle Manager checks what software is available in the vSphere Lifecycle Manager depot. vSphere Lifecycle Manager also checks the depot for what firmware is available from the selected hardware support manager. Based on the available software, firmware, and, for vSAN clusters, the hardware compatibility checks, vSphere Lifecycle Manager provides you with the latest and most appropriate image for your environment.

10-35 Viewing Recommended Images

To view recommended images for a cluster, click **Recommendations** at the top, or select **View recommended images** from the drop-down menu.

The recommended image is always in the major release series of the ESXi version of the image for the cluster.

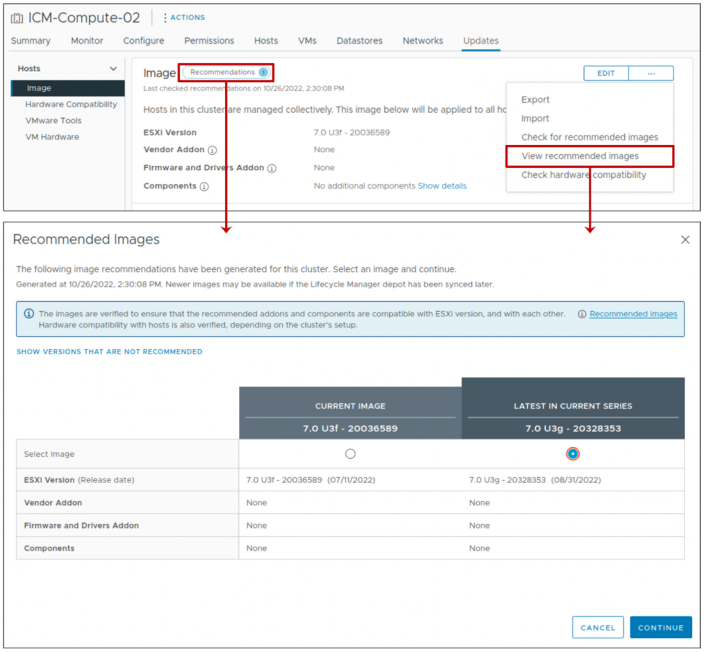

The recommendation that vSphere Lifecycle Manager generates is based on the major ESXi version in the image that the cluster uses. vSphere Lifecycle Manager recommends the most recent ESXi version that causes no hardware compatibility problems or regressions.

For example, if the current image for the cluster is version 7.0 U3f and the versions 7.0U3g and 8.0 are available in the image depot, the latest image in the current series recommendation ESXi version 7.0U3g. Version 8.0 is not considered to be the latest image version because 8.0 starts a new series of a major release.

When you view recommended images, the following types of images appear in vSphere:

- CURRENT IMAGE: The image specification that is running on the cluster.
- LATEST IN CURRENT SERIES: If available, a later version in the same release series appears.

If the current image is the same as the latest release, no recommendations appear.

10-36 Customizing Cluster Images

After you start managing a cluster with an image, you can edit the image:

- Change the base image.
- Change, add, and remove components.

Before saving the image specification, you can validate it:

- Ensures completeness of the image
- Verifies that the image has no missing component dependencies
- Confirms that components do not conflict with one another

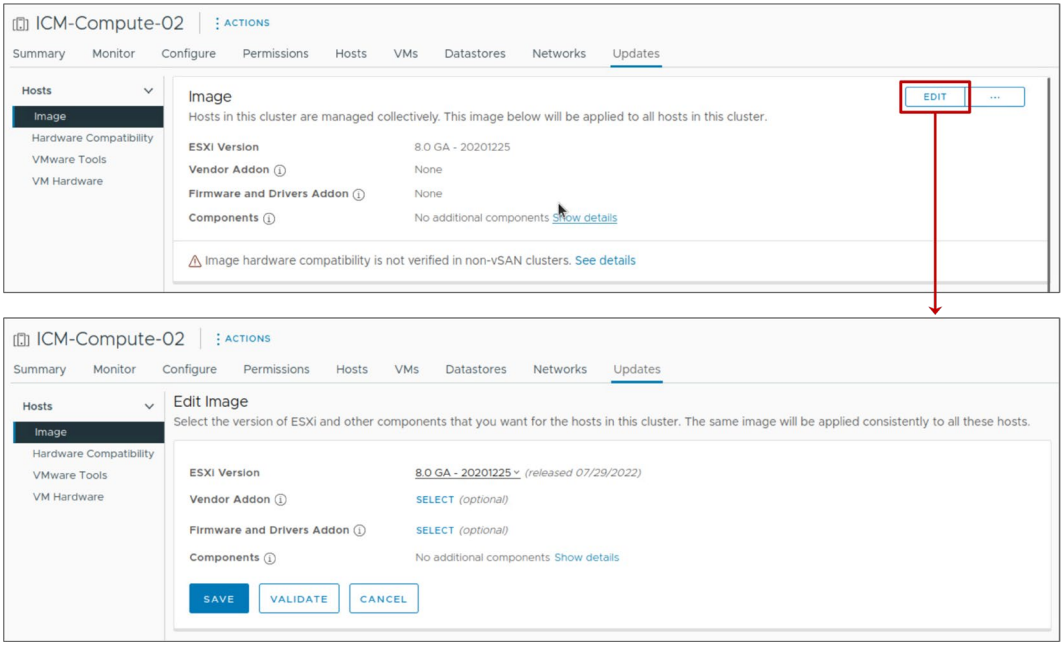

You can edit the image by changing, adding, or removing components, such as the ESXi image version, vendor add-ons, firmware and driver add-ons, and other components.

10-37 Hardware Compatibility

Hardware compatibility checks ensure that the host or cluster hardware is compliant with the VMware Compatibility Guide and vSAN Hardware Compatibility List.

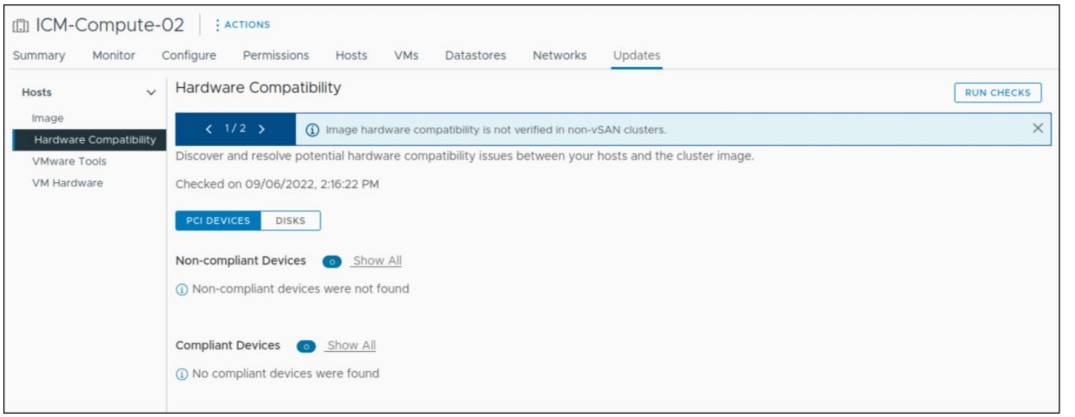

Hardware compatibility lists are lists of hardware certified for use with various vSphere releases.

The VMware Compatibility Guide contains information about server models and I/O devices that are certified for use with particular vSphere releases. You can access this guide at https://www.vmware.com/resources/compatibility/search.php.

Besides VMware Compatibility Guide, vSAN maintains a separate hardware compatibility list that lists all I/O and networking device controller hardware and the respective firmware versions certified for use with vSAN. The vSAN HCL also contains information about the disk drives that a specific vSphere release supports and the earliest disk drive firmware version certified for use with vSAN. You can view information in the vSAN hardware compatibility list at https://www.vmware.com/resources/compatibility/search.php?deviceCategory=vsan.

10-38 Tech Preview: Managing Host Configurations

You can manage the configuration settings of all hosts in a cluster:

- Known as Configuration Manager
- A Tech Preview feature: vSphere 8.0 provides you with early access to test the feature and provide feedback.

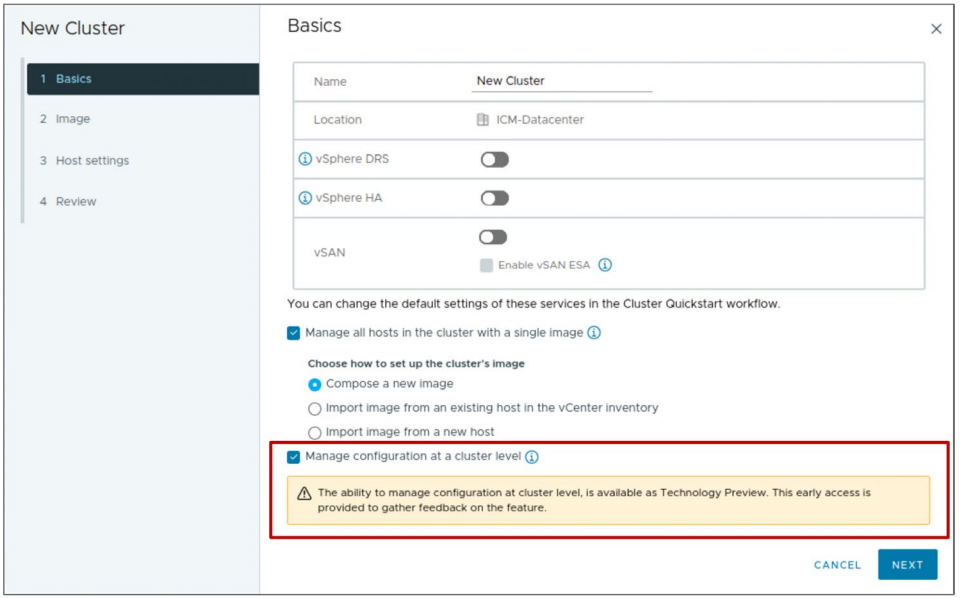

To manage the configuration settings of all hosts in a cluster, you must first manage the cluster with a single image.

Also, all hosts in the cluster must be 8.0 or later.

Configuration Manager does not support the configuration of vSphere Distributed Switches or VMware NSX for ESXi hosts.

Since Configuration Manager is a technology preview (Tech Preview) feature, support requests are serviced. This feature should not be used in a production environment.

For production environments, use host profiles to manage the configuration settings of the hosts in a cluster. For details, see *vSphere Host Profiles* at https://docs.vmware.com/en/VMware-vSphere/index.html.

10-39 Review of Learner Objectives

- Activate vSphere Lifecycle Manager in a cluster and define a cluster image
- Validate ESXi host compliance against a cluster image
- Remediate ESXi hosts using vSphere Lifecycle Manager
- View recommended images for a cluster

10-40 Lesson 4: Managing the Life Cycle of VMware Tools and VM Hardware

10-41 Learner Objectives

- Use vSphere Lifecycle Manager to upgrade VMware Tools and VM hardware

10-42 Keeping VMware Tools Up To Date

VMware can provide a new release of VMware Tools with a release of ESXi.

New releases include:

- Bug fixes
- Security patches
- New driver support for ESXi enhancements
- Performance enhancements for virtual devices

Keeping VMware Tools up to date is an important part of ongoing data center maintenance.

10-43 Upgrading VMware Tools (1)

From a host or cluster's **Updates** tab, select **VMware Tools**.

Step 1: Check the status of VMware Tools running in your VMs:

- Up to Date
- Upgrade Available
- Version Unsupported
- Not Installed
- Guest Managed
- Unknown

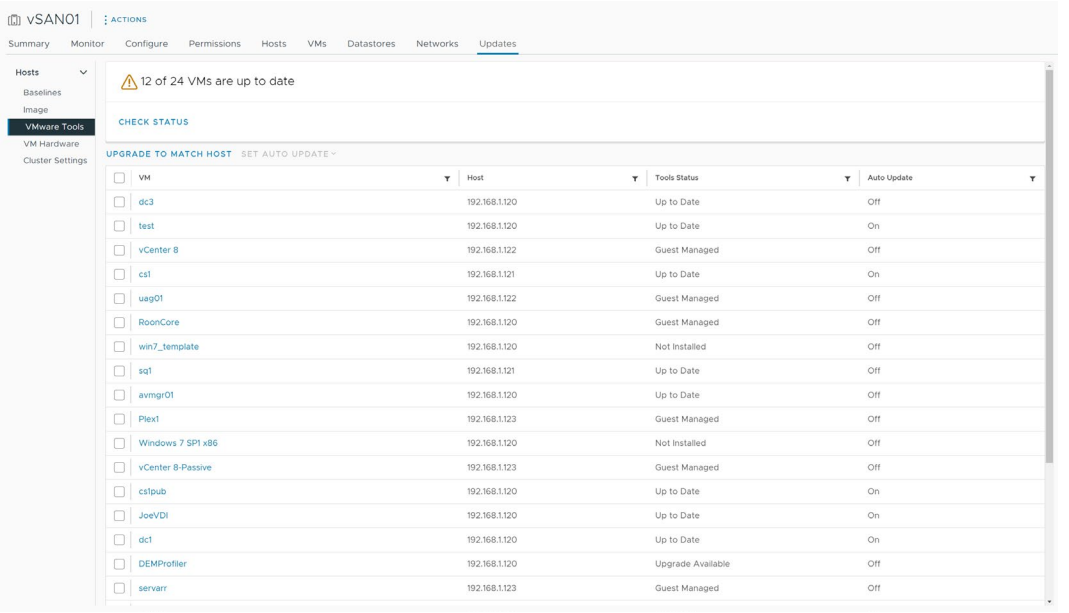

A VM has one of the following VMware Tools status values:

- Up to Date:
  - VMware Tools is installed, supported, and the version is compliant.
  - VMware Tools is installed, supported, and the version is newer than the version that is available on the ESXi host.
- Upgrade Available:
  - VMware Tools is installed, but the version is old.
  - VMware Tools is installed and supported, but a newer version is available on the ESXi host.
- Version Unsupported:
  - VMware Tools is installed, but the version is old.
  - VMware Tools is installed, but the version has a known problem and must be immediately upgraded.
  - VMware Tools is installed, but the version is too new to work correctly with this virtual machine.
- Not Installed: VMware Tools is not installed on this virtual machine. Consider installing VMware Tools in this VM.
- Guest Managed: vSphere does not manage VMware Tools. For example, your VM is running the Linux OpenVMTools package. Use native Linux package management tools to upgrade VMware Tools.
- Unknown: The status of the virtual machine has not been checked. Verify that the VM is powered on before you click **CHECK STATUS**.

10-44 Upgrading VMware Tools (2)

Select the VMs that use VMware Tools whose version you want to upgrade.

Step 2: Click **UPGRADE TO MATCH HOST**.

1. Select the VMs to upgrade.
2. Schedule the upgrade.

   Plan the upgrade during your maintenance window.
3. Select rollback options.

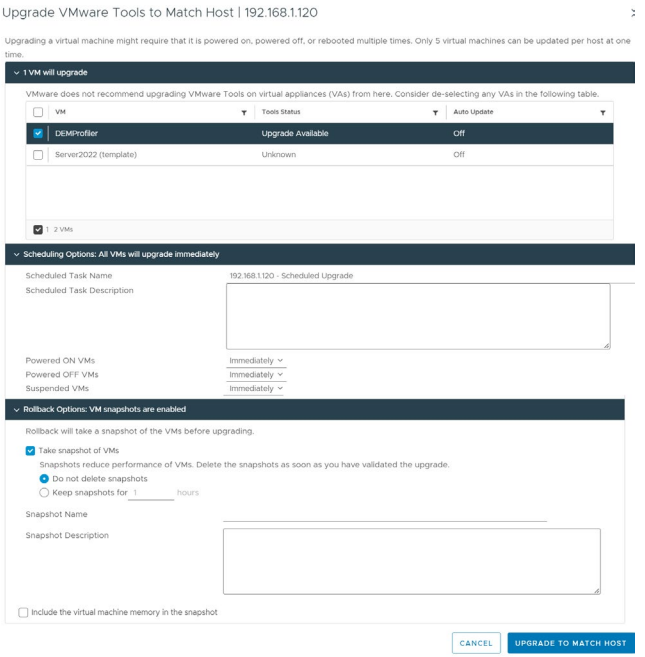

Plan the VMware Tools upgrade during your maintenance window, especially if a reboot of the VM is required after the upgrade completes.

10-45 Keeping VM Hardware Up To Date

With almost every update of ESXi, VMware provides a new release of VM hardware.

As ESXi improves its hardware support, VMware often carries that support into its VMs.

New releases include:

- Greater configuration maximums
- New types of hardware (for example, vGPU, vNVMe, vSGX, vTPM, and so on)

Always upgrade VMware Tools before upgrading VM hardware.

Consider upgrading VM hardware only when new features are required.

10-46 Upgrading VM Hardware (1)

Select VM Hardware to upgrade your VMs' hardware.

Step 1: Check the status of the VM hardware running in your VMs:

- Upgrade Available: Upgrade VM hardware to match the current version available for your ESXi hosts.
- Up to Date: The version of VM hardware configured in the VM matches the latest available version for the ESXi host.

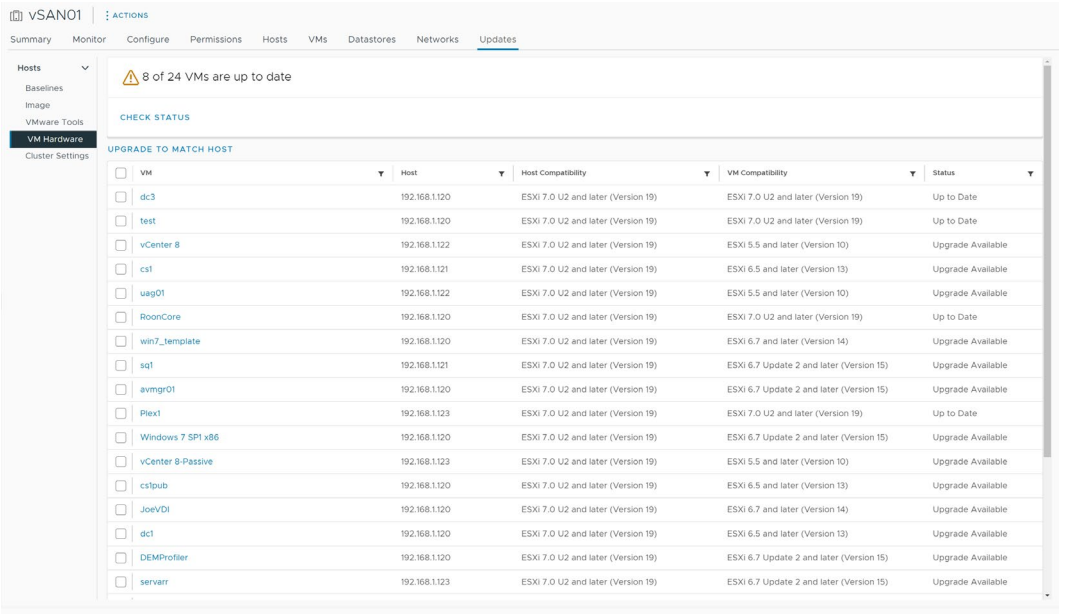

10-47 Upgrading VM Hardware (2)

Step 2: Select the VMs whose hardware version you want to upgrade.

Step 3: Click **UPGRADE TO MATCH HOST**.

1. Schedule the upgrade.

   Plan the upgrade during your maintenance window.

2. Select rollback options.

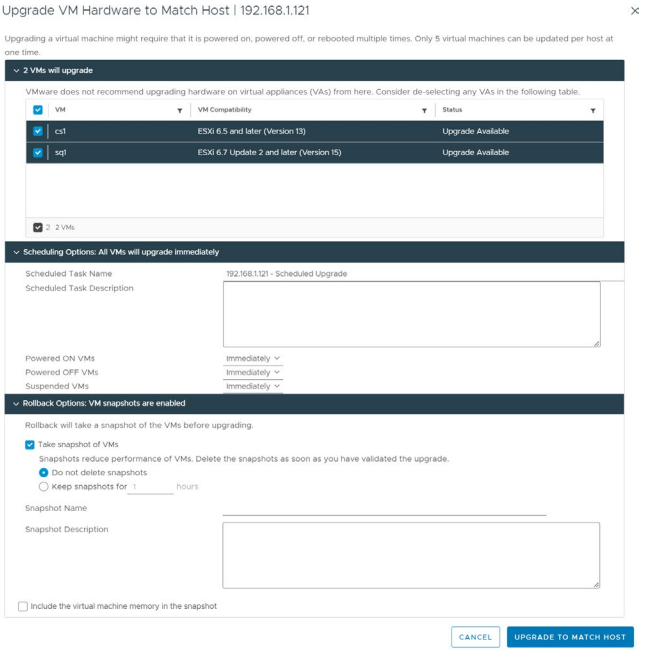

10-48 Lab 26: Using vSphere Lifecycle Manager

Update ESXi hosts using vSphere Lifecycle Manager:

1. Create a Cluster and Select an Image
2. Add ESXi Hosts to the Cluster
3. Check for Host Compliance
4. Remediate Noncompliant Hosts

10-49 Review of Learner Objectives

- Use vSphere Lifecycle Manager to upgrade VMware Tools and VM hardware

10-50 Key Points

- You can generate interoperability reports to verify that your vCenter system meets the minimum requirements for a successful upgrade. It also verifies the order in which vCenter and other products should be upgraded.
- vSphere Lifecycle Manager centralizes automated patch and version management for clusters, ESXi hosts, drivers and firmware, VM hardware, and VMware Tools.
- In vSphere Lifecycle Manager, you can manage ESXi hosts by using images.
- Keeping VMware Tools up to date is an important part of ongoing data center maintenance.
- Consider upgrading VM hardware only when new features are required.

Questions?