# Introduction to VMware vSphere

ESX 4.1

ESXi 4.1

vCenter Server 4.1

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-000284-00

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com



**VMware, Inc.**
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

# Contents

# About This Book

*Introduction to VMware vSphere* provides information about the features and functionality of VMware® vSphere. *Introduction to VMware vSphere* describes ESX, ESXi, and vCenter Server.

## Intended Audience

This information is for those who need to familiarize themselves with the components and capabilities of VMware vSphere. This information is for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacenter operations.

## VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to http://www.vmware.com/support/pubs.

## Document Feedback

VMware welcomes your suggestions for improving our documentation. If you have comments, send your feedback to:

docfeedback@vmware.com

## VMware vSphere Documentation

The VMware vSphere documentation consists of the combined VMware vCenter Server and ESX/ESXi documentation set.

## Abbreviations Used in Figures

The figures in this manual use the abbreviations listed in Table 1.

**Table 1.** Abbreviations

| Abbreviation | Description |
|---|---|
| database | vCenter Server database |
| datastore | Storage for the managed host |
| dsk# | Storage disk for the managed host |
| hostn | vCenter Server managed hosts |
| SAN | Storage area network type datastore shared between managed hosts |

**Table 1.** Abbreviations (Continued)

| Abbreviation | Description |
|---|---|
| tmplt | Template |
| user# | User with access permissions |
| VC | vCenter Server |
| VM# | Virtual machines on a managed host |

# Technical Support and Education Resources

The following technical support resources are available to you. To access the current version of this book and other books, go to http://www.vmware.com/support/pubs.

**Online and Telephone Support**

To use online support to submit technical support requests, view your product and contract information, and register your products, go to http://www.vmware.com/support.

Customers with appropriate support contracts should use telephone support for the fastest response on priority 1 issues. Go to http://www.vmware.com/support/phone_support.html.

**Support Offerings**

Find out how VMware support offerings can help meet your business needs. Go to http://www.vmware.com/support/services.

**VMware Professional Services**

VMware Education Services courses offer extensive hands-on labs, case study examples, and course materials designed to be used as on-the-job reference tools. Courses are available onsite, in the classroom, and live online. For onsite pilot programs and implementation best practices, VMware Consulting Services provides offerings to help you assess, plan, build, and manage your virtual environment. To access information about education classes, certification programs, and consulting services, go to http://www.vmware.com/services.

# VMware vSphere Introduction

VMware vSphere leverages the power of virtualization to transform datacenters into simplified cloud computing infrastructures and enables IT organizations to deliver flexible and reliable IT services. VMware vSphere virtualizes and aggregates the underlying physical hardware resources across multiple systems and provides pools of virtual resources to the datacenter.

As a cloud operating system, VMware vSphere manages large collections of infrastructure (such as CPUs, storage, and networking) as a seamless and dynamic operating environment, and also manages the complexity of a datacenter. The following component layers make up VMware vSphere.

**Infrastructure Services**

Infrastructure Services are the set of services provided to abstract, aggregate, and allocate hardware or infrastructure resources. Infrastructure Services are categorized into several types.

- VMware vCompute, which includes the VMware capabilities that abstract away from underlying disparate server resources. vCompute services aggregate these resources across many discrete servers and assign them to applications.
- VMware vStorage, which is the set of technologies that enables the most efficient use and management of storage in virtual environments.
- VMware vNetwork, which is the set of technologies that simplify and enhance networking in virtual environments.

**Application Services**

Application Services are the set of services provided to ensure availability, security, and scalability for applications. Examples include High Availability and Fault Tolerance.

**VMware vCenter Server**

VMware vCenter Server provides a single point of control of the datacenter. It provides essential datacenter services such as access control, performance monitoring, and configuration.

**Clients**

Users can access the VMware vSphere datacenter through clients such as the vSphere Client or Web Access through a Web browser.

Figure 1 shows the relationships between the component layers of VMware vSphere.

**Figure 1.** VMware vSphere Component Layers

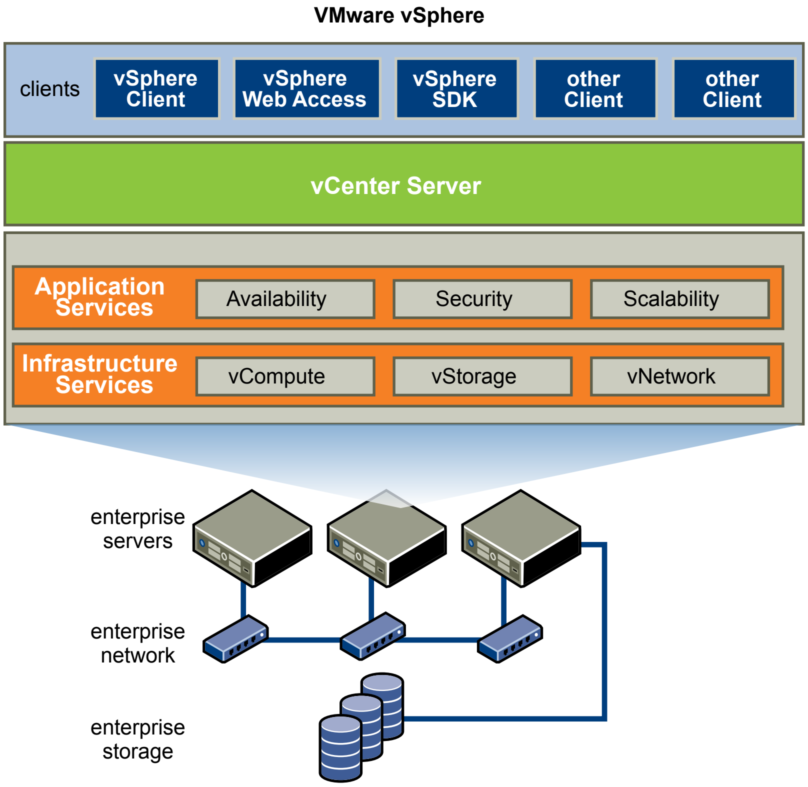


# VMware vSphere Components

An introduction to the components of VMware vSphere helps you to understand the parts and how they interact.

VMware vSphere includes the following components.

**VMware® ESX and VMware® ESXi**

A virtualization layer run on physical servers that abstracts processor, memory, storage, and resources into multiple virtual machines.

Two versions of ESX are available:

- VMware ESX 4.1 contains a built-in service console. It is available as an installable CD-ROM boot image.
- VMware ESXi 4.1 does not contain a service console. It is available in two forms: VMware ESXi 4.1 Embedded and VMware ESXi 4.1 Installable. ESXi 4.1 Embedded is firmware that is built into a server's physical hardware. ESXi 4.1 Installable is software that is available as an installable CD-ROM boot image. You install the ESXi 4.1 Installable software onto a server's hard drive.

**VMware® vCenter Server**

The central point for configuring, provisioning, and managing virtualized IT environments.

**VMware® vSphere Client**

An interface that allows users to connect remotely to vCenter Server or ESX/ESXi from any Windows PC.

**VMware® vSphere Web Access**

A Web interface that allows virtual machine management and access to remote consoles.

| | |
|---|---|
| **VMware® Virtual Machine File System (VMFS)** | A high performance cluster file system for ESX/ESXi virtual machines. |
| **VMware® Virtual SMP** | Feature that enables a single virtual machine to use multiple physical processors simultaneously. |
| **VMware® vMotion and Storage vMotion** | VMware vMotion enables the live migration of running virtual machines from one physical server to another with zero down time, continuous service availability, and complete transaction integrity. Storage vMotion enables the migration of virtual machine files from one datastore to another without service interruption. You can choose to place the virtual machine and all its disks in a single location, or select separate locations for the virtual machine configuration file and each virtual disk. The virtual machine remains on the same host during Storage vMotion.<br><br>Migration with vMotion lets you move a powered-on virtual machine to a new host. Migration with vMotion allows you to move a virtual machine to a new host without any interruption in the availability of the virtual machine. Migration with vMotion cannot be used to move virtual machines from one datacenter to another.<br><br>Migration with Storage vMotion lets you move the virtual disks or configuration file of a powered-on virtual machine to a new datastore. Migration with Storage vMotion allows you to move a virtual machine's storage without any interruption in the availability of the virtual machine. |
| **VMware® High Availability (HA)** | Feature that provides high availability for virtual machines. If a server fails, affected virtual machines are restarted on other production servers that have spare capacity. |
| **VMware® Distributed Resource Scheduler (DRS)** | Feature that allocates and balances computing capacity dynamically across collections of hardware resources for virtual machines. This feature includes distributed power management (DPM) capabilities that enable a datacenter to significantly reduce its power consumption. |
| **VMware® vSphere SDK** | Feature that provides a standard interface for VMware and third-party solutions to access the VMware vSphere. |
| **VMware® Fault Tolerance** | When Fault Tolerance is enabled for a virtual machine, a secondary copy of the original (or primary) virtual machine is created. All actions completed on the primary virtual machine are also applied to the secondary virtual machine. If the primary virtual machine becomes unavailable, the secondary machine becomes active, providing continuous availability. |
| **vNetwork Distributed Switch (vDS)** | Feature that includes a distributed virtual switch (vDS), which spans many ESX/ESXi hosts enabling significant reduction of on-going network maintenance activities and increasing network capacity. This allows virtual machines to maintain consistent network configuration as they migrate across multiple hosts. |

**Host Profiles**

Feature that simplifies host configuration management through user-defined configuration policies. The host profile policies capture the blueprint of a known, validated host configuration and use this to configure networking, storage, security, and other settings across multiple hosts. The host profile policies also monitor compliance to standard host configuration settings across the datacenter. Host profiles reduce manual steps involved in configuring a host and can help maintain consistency and correctness across the datacenter.

**Pluggable Storage Architecture (PSA)**

A storage partner plug-in framework that enables greater array certification flexibility and improved array-optimized performance. PSA is a multipath I/O framework allowing storage partners to enable their array asynchronously to ESX release schedules. VMware partners can deliver performance-enhancing multipath load-balancing behaviors that are optimized for each array.

## Physical Topology of vSphere Datacenter

A typical VMware vSphere datacenter consists of basic physical building blocks such as x86 virtualization servers, storage networks and arrays, IP networks, a management server, and desktop clients.

This physical topology of the vSphere datacenter is illustrated in Figure 2.

**Figure 2.** VMware vSphere Datacenter Physical Topology

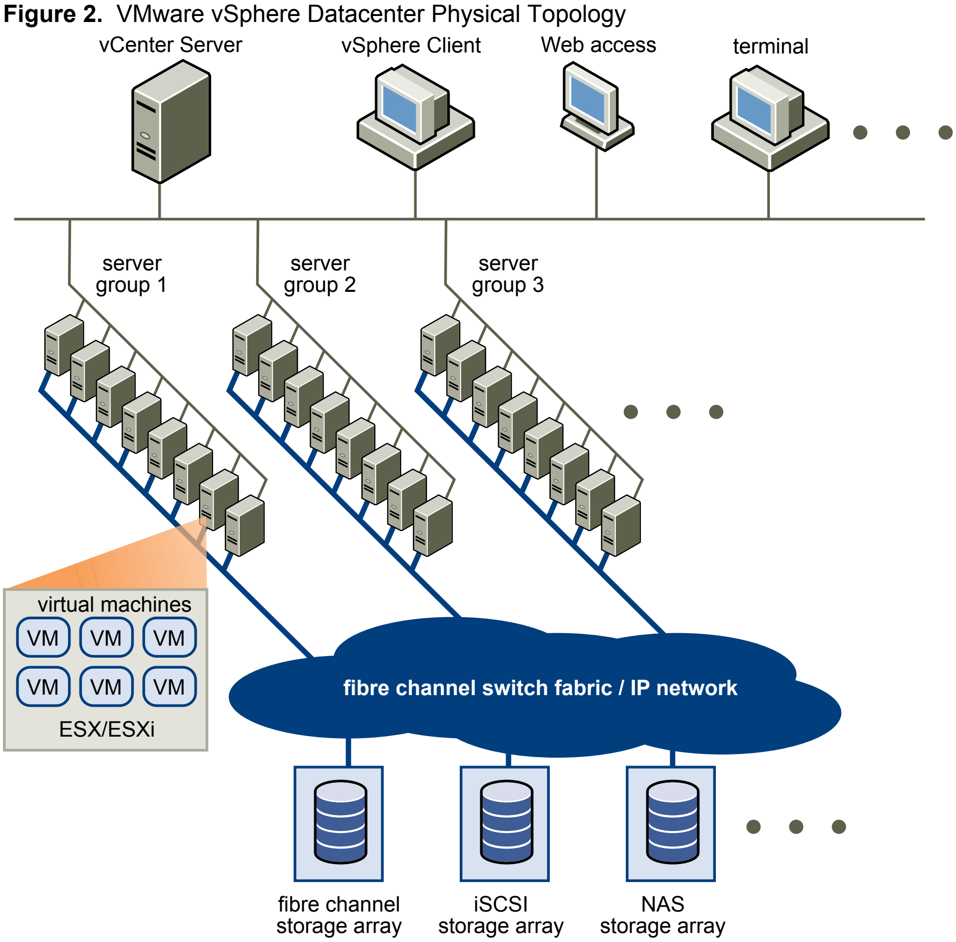

The vSphere datacenter topology includes the following components.

**Computing servers**
Industry standard x86 servers that run ESX/ESXi on the bare metal. ESX/ESXi software provides resources for and runs the virtual machines. Each computing server is referred to as a standalone host in the virtual environment. You can group a number of similarly configured x86 servers with connections to the same network and storage subsystems to provide an aggregate set of resources in the virtual environment, called a cluster.

**Storage networks and arrays**
Fibre Channel SAN arrays, iSCSI SAN arrays, and NAS arrays are widely used storage technologies supported by VMware vSphere to meet different datacenter storage needs. The storage arrays are connected to and shared between groups of servers through storage area networks. This arrangement allows aggregation of the storage resources and provides more flexibility in provisioning them to virtual machines.

**IP networks**
Each computing server can have multiple NICs to provide high bandwidth and reliable networking to the entire VMware vSphere datacenter.

**vCenter Server**
vCenter Server provides a single point of control to the datacenter. It provides essential datacenter services such as access control, performance monitoring, and configuration. It unifies the resources from the individual computing servers to be shared among virtual machines in the entire datacenter. It does this by managing the assignment of virtual machines to the computing servers and the assignment of resources to the virtual machines within a given computing server based on the policies that the system administrator sets.

Computing servers continue to function even in the unlikely event that vCenter Server becomes unreachable (for example, if the network is severed). Servers can be managed separately and continue to run the virtual machines assigned to them based on the resource assignment that was last set. After connection to vCenter Server is restored, it can manage the datacenter as a whole again.

**Management clients**
VMware vSphere provides several interfaces for datacenter management and virtual machine access. These interfaces include VMware vSphere Client (vSphere Client), web access through a web browser, vSphere Command-Line Interface (vSphere CLI), or vSphere Management Assistant (vMA).

# Virtual Datacenter Architecture

VMware vSphere virtualizes the entire IT infrastructure including servers, storage, and networks.

VMware vSphere aggregates these resources and presents a uniform set of elements in the virtual environment. With VMware vSphere, you can manage IT resources like a shared utility and dynamically provision resources to different business units and projects.

Figure 3 shows the key elements in a virtual datacenter.

**Figure 3.** Virtual Datacenter Architecture

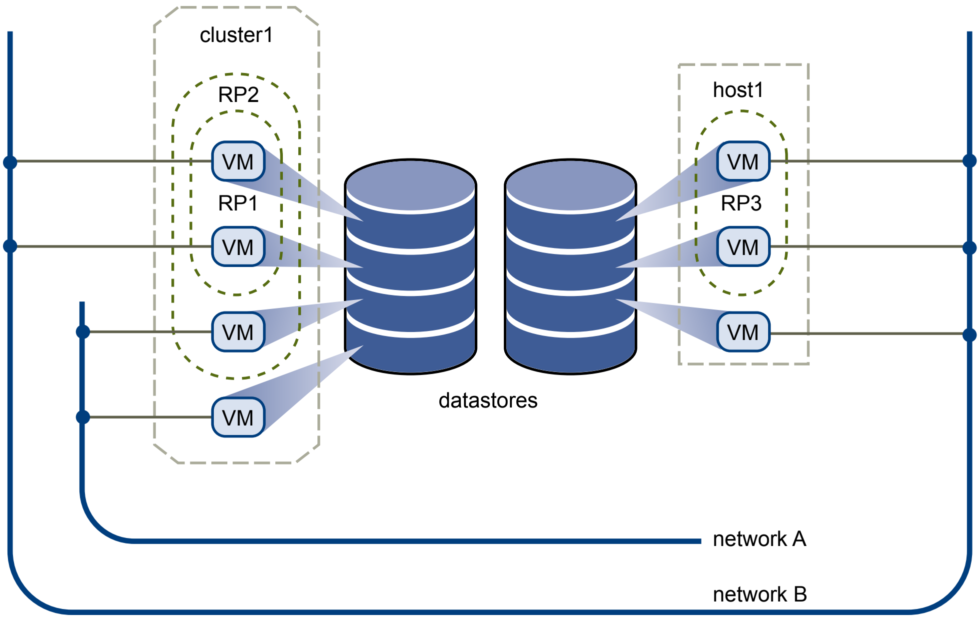


You can use vSphere to view, configure, and manage these key elements. The following is a list of the key elements:

- Computing and memory resources called hosts, clusters, and resource pools
- Storage resources called datastores
- Networking resources called networks
- Virtual machines

A host is the virtual representation of the computing and memory resources of a physical machine running ESX/ESXi. When two or more physical machines are grouped to work and be managed as a whole, the aggregate computing and memory resources form a cluster. Machines can be dynamically added or removed from a cluster. Computing and memory resources from hosts and clusters can be finely partitioned into a hierarchy of resource pools.

Datastores are virtual representations of combinations of underlying physical storage resources in the datacenter. These physical storage resources can come from the following sources:

- Local SCSI, SAS, or SATA disks of the server
- Fibre Channel SAN disk arrays
- iSCSI SAN disk arrays
- Network Attached Storage (NAS) arrays

Networks in the virtual environment connect virtual machines to one another and to the physical network outside of the virtual datacenter.

Virtual machines can be designated to a particular host, cluster or resource pool, and a datastore when they are created. After they are powered-on, virtual machines consume resources dynamically as the workload increases or give back resources dynamically as the workload decreases.

Provisioning of virtual machines is much faster and easier than physical machines. New virtual machines can be created in seconds. When a virtual machine is provisioned, the appropriate operating system and applications can be installed unaltered on the virtual machine to handle a particular workload as though they were being installed on a physical machine. A virtual machine can be provisioned with the operating system and applications installed and configured.

Resources get provisioned to virtual machines based on the policies that are set by the system administrator who owns the resources. The policies can reserve a set of resources for a particular virtual machine to guarantee its performance. The policies can also prioritize and set a variable portion of the total resources to each virtual machine. A virtual machine is prevented from being powered-on and consuming resources if doing so violates the resource allocation policies. For more information on resource and power management, see the *Resource Management Guide*.

## Hosts, Clusters, and Resource Pools

Hosts, clusters, and resources pools provide flexible and dynamic ways to organize the aggregated computing and memory resources in the virtual environment and link them back to the underlying physical resources.

A host represents the aggregate computing and memory resources of a physical x86 server. For example, if the physical x86 server has four dual-core CPUs running at 4GHz each and 32GB of system memory, the host has 32GHz of computing power and 32GB of memory available for running virtual machines that are assigned to it.

A cluster acts and can be managed as a single entity. It represents the aggregate computing and memory resources of a group of physical x86 servers sharing the same network and storage arrays. For example, if the group contains eight servers with four dual-core CPUs each running at 4GHz and 32GB of memory, the cluster has an aggregate 256GHz of computing power and 256GB of memory available for running virtual machines.

Resource pools are partitions of computing and memory resources from a single host or a cluster. Resource pools can be hierarchical and nested. You can partition any resource pool into smaller resource pools to divide and assign resources to different groups or for different purposes.

Figure 4 illustrates the use of resource pools. Three x86 servers with 4GHz computing power and 16GB of memory each are aggregated to form a cluster of 12GHz computing power and 48GB of memory. The Finance Department resource pool reserves 8GHz of computing power and 32GB of memory from the cluster. The remaining 4GHz computing power and 16GB of memory are reserved for the other virtual machine. From the Finance Department resource pool, the smaller Accounting resource pool reserves 4GHz computing power and 16GB of memory for the virtual machines from the accounting department. That leaves 4GHz of computing power and 16GB of memory for the virtual machine called Payroll.

**Figure 4.** Hosts, Clusters, and Resource Pools

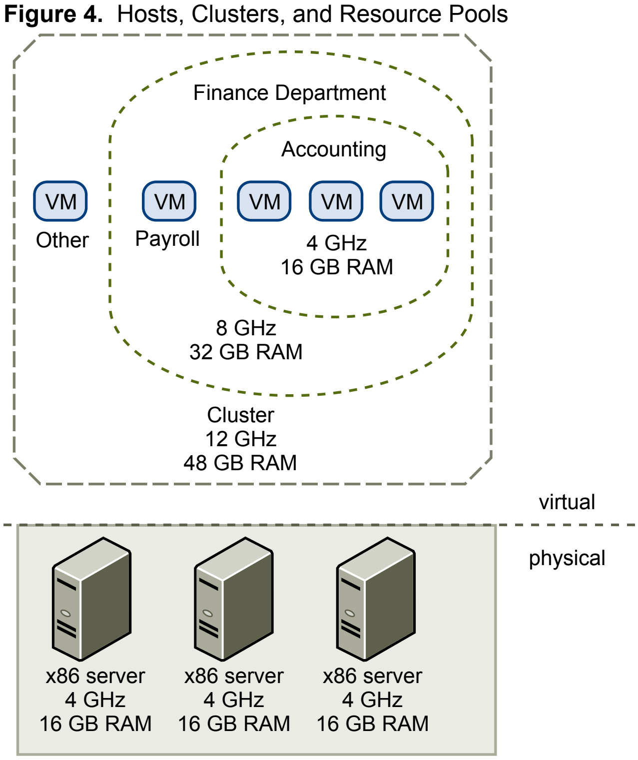

You can dynamically change resource allocation policies. For example, at year end, the workload on Accounting increases, and which requires an increase in the Accounting resource pool reserve of 4GHz of power to 6GHz. You can make the change to the resource pool dynamically without shutting down the associated virtual machines.

When reserved resources are not being used by a resource pool or a virtual machine, the resources can be shared. In the example, if the 4GHz of resources reserved for the Accounting department are not being used, the Payroll virtual machine can use those gigahertz during its peak time. When Accounting resource demands increase, Payroll dynamically returns them. Resources are reserved for different resource pools, but resources are not wasted if an owner does not use them. This capability helps to maximize resource use while also ensuring that reservations are met and resource policies enforced.

As demonstrated by the example, resource pools can be nested, organized hierarchically, and dynamically reconfigured so that the IT environment matches the company organization. Individual business units can receive dedicated resources while still exploiting from the efficiency of resource pooling.

ESX/ESXi provides a memory compression cache to improve virtual machine performance when you use memory overcommitment. Memory compression is enabled by default. When a host's memory becomes overcommitted, ESX/ESXi compresses virtual pages and stores them in memory.

Because accessing compressed memory is faster than accessing memory that has been swapped to disk, memory compression in ESX/ESXi allows you to overcommit memory without hindering performance. When a virtual page needs to be swapped, ESX/ESXi first attempts to compress the page. Pages that can be compressed to 2 KB or smaller are stored in the virtual machine's compression cache, increasing the capacity of the host.

## VMware vSphere Distributed Services

VMware vMotion, VMware Storage vMotion, VMware DRS, Storage I/O Control, VMware HA, and Fault Tolerance are distributed services that enable efficient and automated resource management and high availability for virtual machines.

Virtual machines run on and consume resources from ESX/ESXi. vMotion enables the migration of running virtual machines from one physical server to another without service interruption, as shown in Figure 5. The effect is a more efficient assignment of resources. With vMotion, resources can be dynamically reallocated to virtual machines across physical servers.

**Figure 5.** Migration with vMotion

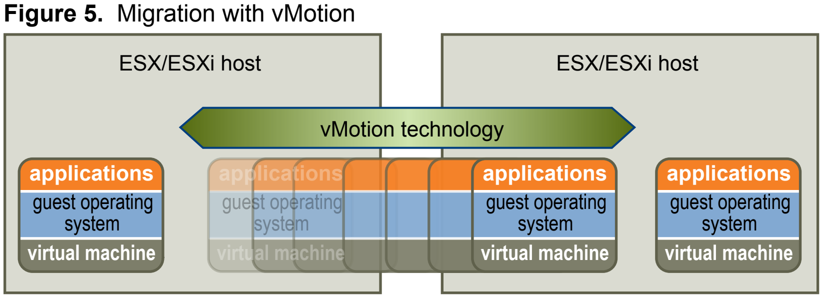


Storage vMotion enables the migration of virtual machines from one datastore to another datastore without service interruption. This allows administrators, for example, to off-load virtual machines from one storage array to another to perform maintenance, reconfigure LUNs, resolve out-of-space issues, and upgrade VMFS volumes. Administrators can also use Storage vMotion to optimize the storage environment for improved performance by seamlessly migrating virtual machine disks.

VMware DRS helps you manage a cluster of physical hosts as a single compute resource. You can assign a virtual machine to a cluster and DRS finds an appropriate host on which to run the virtual machine. DRS places virtual machines in such a way as to ensure that load across the cluster is balanced, and cluster-wide resource allocation policies (for example, reservations, priorities, and limits) are enforced. When a virtual machine is powered on, DRS performs an initial placement of the virtual machine on a host. As cluster conditions change (for example, load and available resources), DRS migrates (using vMotion) virtual machines to other hosts as necessary.

**Figure 6.** VMware DRS

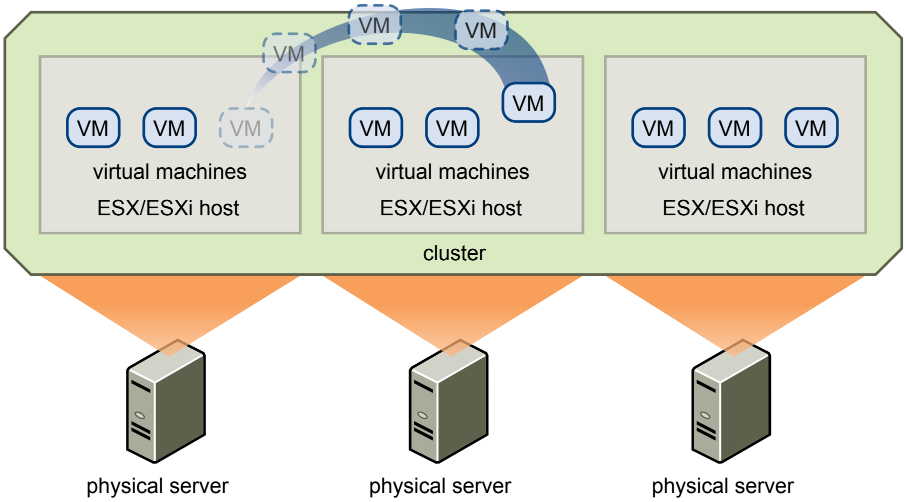


When you add a new physical server to a cluster, DRS enables virtual machines to immediately take advantage of the new resources because it distributes the running virtual machines.

When DPM is enabled, the system compares cluster-level and host-level capacity to the demands of virtual machines running in the cluster. If the resource demands of the running virtual machines can be met by a subset of hosts in the cluster, DPM migrates the virtual machines to this subset and powers down the hosts that are not needed. When resource demands increase, DPM powers these hosts back on and migrates the virtual machines to them. This dynamic cluster right-sizing that DPM performs reduces the power consumption of the cluster without sacrificing virtual machine performance or availability.

You can configure DRS to execute virtual machine placement, virtual machine migration, and host power actions, or to provide recommendations which the datacenter administrator can assess and manually act upon.

Storage I/O Control congestion management allows cluster-wide storage I/O prioritization and enables administrators to set congestion thresholds for I/O shares.

VMware HA enables quick automated restart of virtual machines on a different physical server within a cluster if a host fails. All applications within the virtual machines have the high availability benefit.

HA monitors all physical hosts in a cluster and detects host failures. An agent placed on each physical host maintains a heartbeat with the other hosts in the resource pool. Loss of a heartbeat initiates the process of restarting all affected virtual machines on other hosts. See Figure 7 for an example of VMware HA. HA admission control ensures that sufficient resources are available in the cluster at all times to restart virtual machines on different physical hosts in the event of host failure.

**Figure 7.** VMware HA

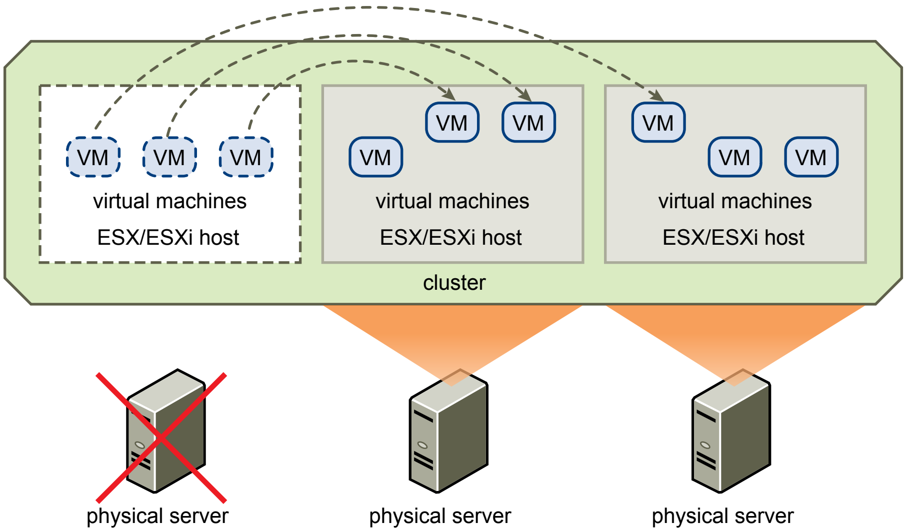


HA also provides a Virtual Machine Monitoring feature that monitors the status of virtual machines in an HA cluster. If a virtual machine does not generate heartbeats within a specified time, Virtual Machine Monitoring identifies it as having failed and restarts it. If restarts occur, policies can control the number of restarts.

HA is configured centrally through vCenter Server. After HA is configured, it operates continuously and in a distributed manner on every ESX host without needing vCenter Server. Even if vCenter Server fails, HA failovers can still successfully restart virtual machines.

Using VMware vLockstep technology, VMware Fault Tolerance (FT) on the ESX/ESXi host platform provides continuous availability by protecting a virtual machine (the primary virtual machine) with a shadow copy (secondary virtual machine) that runs in virtual lockstep on a separate host. Inputs and events performed on the primary virtual machine are recorded and replayed on the Secondary virtual machine ensuring that the two remain in an identical state. For example, mouse-clicks and keystrokes are recorded on the primary virtual machine and replayed on the secondary virtual machine. Because the secondary virtual machine is in virtual lockstep with the primary virtual machine, it can take over execution at any point without service interruption or loss of data.

## Network Architecture

VMware vSphere has a set of virtual networking elements that lets you network the virtual machines in the datacenter like a physical environment.

**Figure 8.** Networking with vNetwork Standard Switches

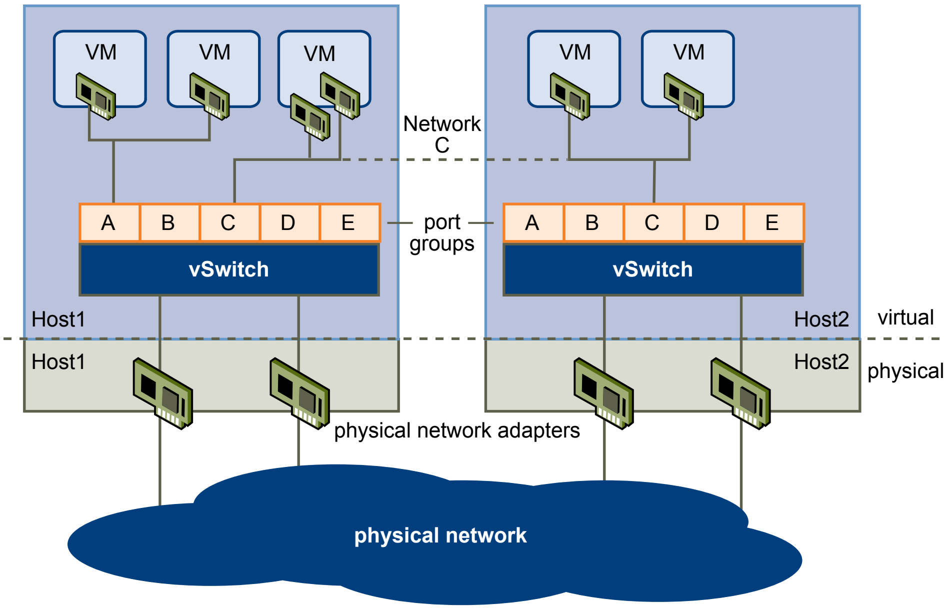


Figure 8 shows the relationship between the networks inside and outside the virtual environment for vSwitches. The virtual environment provides networking elements similar to the physical world. They are virtual network interface cards (vNIC), vNetwork Standard Switches (vSwitch), vNetwork Distributed Switches (vDS), and port groups. vDS networking is shown in Figure 9.

Like a physical machine, each virtual machine has one or more vNICs. The guest operating system and application programs communicate with a vNIC through either a commonly available device driver or a VMware device driver optimized for the virtual environment. In either case, communication in the guest operating system occurs just as it would with a physical device. Outside the virtual machine, the vNIC has its own MAC address and one or more IP addresses. It responds to the standard Ethernet protocol as would a physical NIC. An outside agent does not detect that it is communicating with a virtual machine.

A virtual switch works like a layer 2 physical switch. Each server has its own virtual switches. On one side of the virtual switch are port groups that connect to virtual machines. On the other side are uplink connections to physical Ethernet adapters on the server where the virtual switch resides. Virtual machines connect to the outside world through the physical Ethernet adapters that are connected to the virtual switch uplinks.

A virtual switch can connect its uplinks to more than one physical Ethernet adapter to enable NIC teaming. With NIC teaming, two or more physical adapters can be used to share the traffic load or provide passive failover in the event of a physical adapter hardware failure or a network outage. For information about NIC teaming, see the *ESX Configuration Guide* or *ESXi Configuration Guide*.

Port group is a unique concept in the virtual environment. A port group is a mechanism for setting policies that govern the network connected to it. A vSwitch can have multiple port groups. Instead of connecting to a particular port on the vSwitch, a virtual machine connects its vNIC to a port group. Virtual machines that connect to the same port group belong to the same network inside the virtual environment even if they are on different physical servers.

You can configure port groups to enforce policies that provide enhanced networking security, network segmentation, better performance, high availability, and traffic management.

A vNetwork Distributed Switch (vDs) functions as a single virtual switch across all associated hosts. This functionality allows virtual machines to maintain consistent network configuration as they migrate across multiple hosts. Like a vSwitch, each vDS is a network hub that virtual machines can use. A vDS can route traffic internally between virtual machines or link to an external network by connecting to physical Ethernet adapters. Each vDS can also have one or more dvPort groups assigned to it. dvPort groups aggregate multiple ports under a common configuration and provide a stable anchor point for virtual machines connecting to labeled networks.

**Figure 9.** Networking with vNetwork Distributed Switches

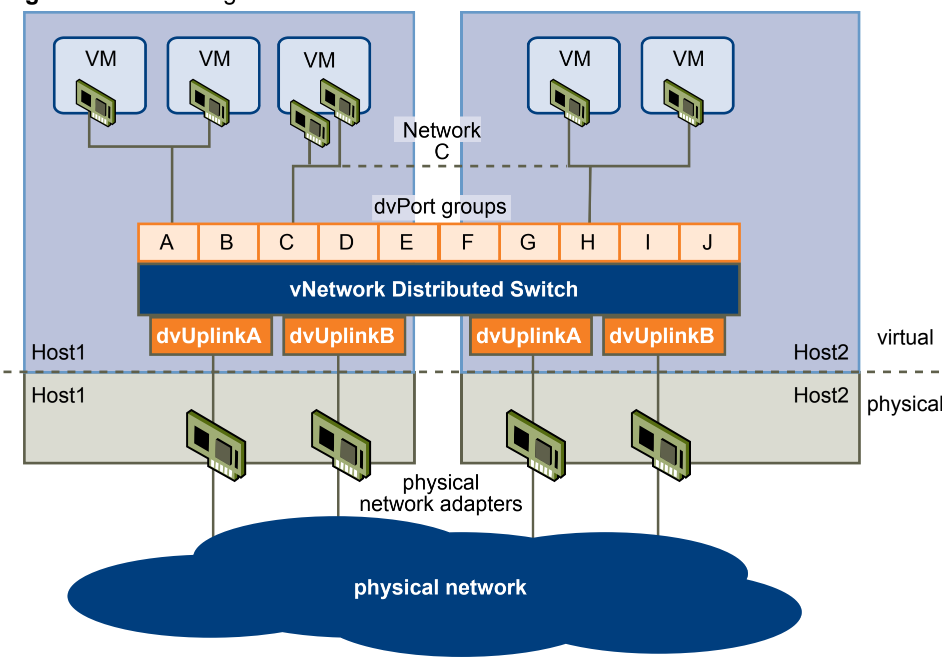


Network resource pools determine the priority different network traffic types are given on a vDS. When network resource management is enabled, vDS traffic is divided into six network resource pools: FT traffic, iSCSI traffic, vMotion traffic, management traffic, NFS traffic, and virtual machine traffic. You can control the priority for the traffic from each of these network resource pools by setting the physical adapter shares and host limits for each network resource pool.

**Layer 2 security options** — Enforces what vNICs attached to a port group in a virtual machine can do by controlling capabilities for a promiscuous mode, MAC address changes, or forged transmissions.

**VLAN support** — Integrates virtual networks with physical network VLANs.

**Private VLAN** — Solves VLAN ID limitations and avoids using up VLAN ids in certain deployment scenarios.

**Traffic shaping** — Defines QOS policies for average and peak bandwidth, and traffic burst size. You set policies to improve traffic management.

**NIC teaming** — Sets the NIC teaming policies for an individual port group or network to share traffic load or provide failover in case of hardware failure.

# Storage Architecture

The VMware vSphere storage architecture consists of layers of abstraction that hide and manage the complexity and differences among physical storage subsystems.

This storage architecture is shown in Figure 10.

**Figure 10.** Storage Architecture

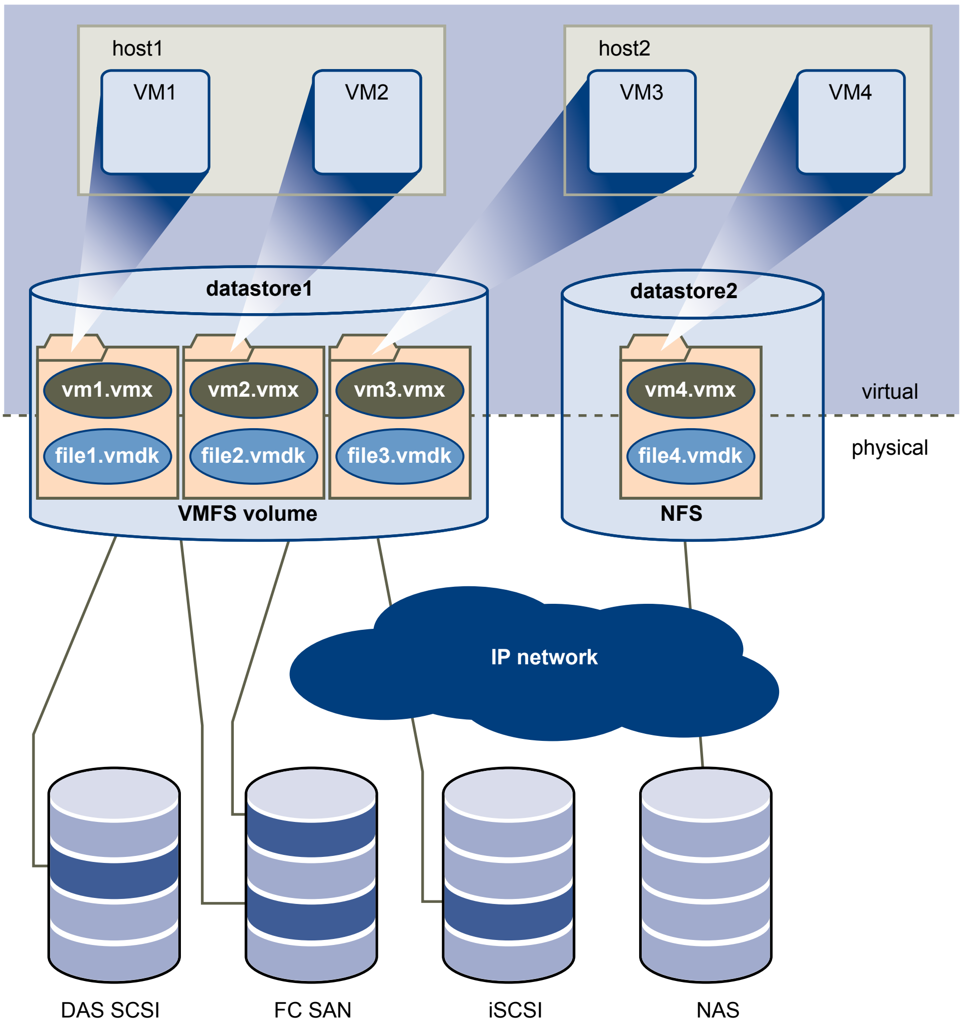


To the applications and guest operating systems inside each virtual machine, the storage subsystem appears as a virtual SCSI controller connected to one or more virtual SCSI disks. These controllers are the only types of SCSI controllers that a virtual machine can see and access. These controllers include BusLogic Parallel, LSI Logic Parallel, LSI Logic SAS, and VMware Paravirtual.

The virtual SCSI disks are provisioned from datastore elements in the datacenter. A datastore is like a storage appliance that delivers storage space for virtual machines across multiple physical hosts.

The datastore abstraction is a model that assigns storage space to virtual machines while insulating the guest from the complexity of the underlying physical storage technology. The guest virtual machine is not exposed to Fibre Channel SAN, iSCSI SAN, direct attached storage, and NAS.

Each datastore is a physical VMFS volume on a storage device. NAS datastores are an NFS volume with VMFS characteristics. Datastores can span multiple physical storage subsystems. A single VMFS volume can contain one or more LUNs from a local SCSI disk array on a physical host, a Fibre Channel SAN disk farm, or iSCSI SAN disk farm. New LUNs added to any of the physical storage subsystems are detected and made available to all existing or new datastores. Storage capacity on a previously created datastore can be extended without powering down physical hosts or storage subsystems. If any of the LUNs within a VMFS volume fails or becomes unavailable, only virtual machines that use that LUN are affected. An exception is the LUN that has the first extent of the spanned volume. All other virtual machines with virtual disks residing in other LUNs continue to function as normal.

Each virtual machine is stored as a set of files in a directory in the datastore. The disk storage associated with each virtual guest is a set of files within the guest's directory. You can operate on the guest disk storage as an ordinary file. The disk storage can be copied, moved, or backed up. New virtual disks can be added to a virtual machine without powering it down. In that case, a virtual disk file (`.vmdk`) is created in VMFS to provide new storage for the added virtual disk or an existing virtual disk file is associated with a virtual machine.

VMFS is a clustered file system that leverages shared storage to allow multiple physical hosts to read and write to the same storage simultaneously. VMFS provides on-disk locking to ensure that the same virtual machine is not powered on by multiple servers at the same time. If a physical host fails, the on-disk lock for each virtual machine is released so that virtual machines can be restarted on other physical hosts.

VMFS also features failure consistency and recovery mechanisms, such as distributed journaling, a failure-consistent virtual machine I/O path, and virtual machine state snapshots. These mechanisms can aid quick identification of the cause and recovery from virtual machine, physical host, and storage subsystem failures.

VMFS also supports raw device mapping (RDM). RDM provides a mechanism for a virtual machine to have direct access to a LUN on the physical storage subsystem (Fibre Channel or iSCSI only). RDM supports two typical types of applications:

- SAN snapshot or other layered applications that run in the virtual machines. RDM better enables scalable backup offloading systems using features inherent to the SAN.
- Microsoft Clustering Services (MSCS) spanning physical hosts and using virtual-to-virtual clusters as well as physical-to-virtual clusters. Cluster data and quorum disks must be configured as RDMs rather than files on a shared VMFS.

**Figure 11.** Raw Device Mapping

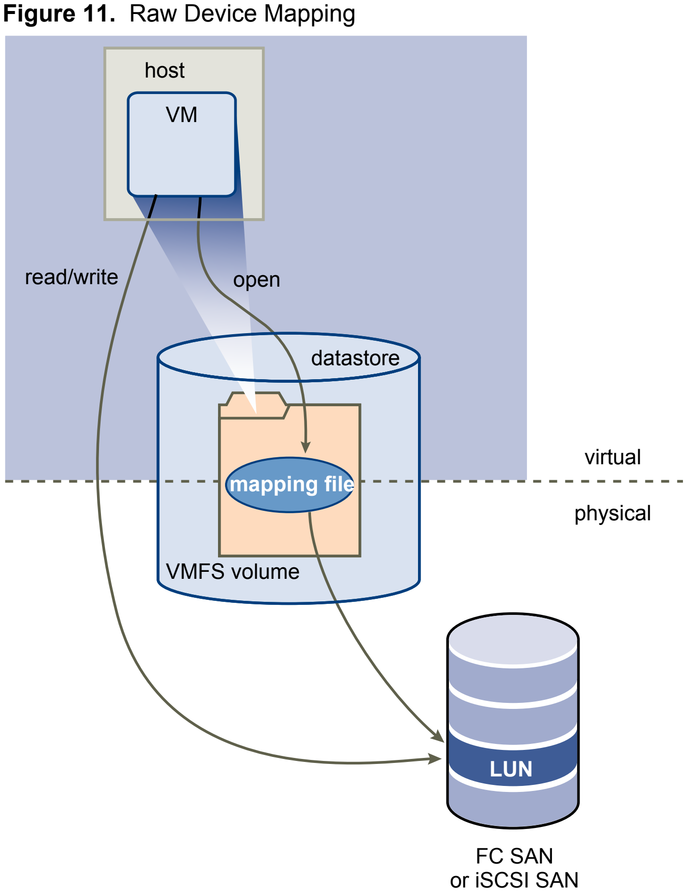

An RDM is a symbolic link from a VMFS volume to a raw LUN. The mapping makes LUNs appear as files in a VMFS volume. The mapping file, not the raw LUN, is referenced in the virtual machine configuration.

When a LUN is opened for access, the mapping file is read to obtain the reference to the raw LUN. Thereafter, reads and writes go directly to the raw LUN rather than going through the mapping file.

## VMware vCenter Server

VMware vCenter Server provides centralized management for datacenters.

vCenter Server aggregates physical resources from multiple ESX/ESXi hosts and presents a central collection of flexible resources for the system administrator to provision to virtual machines in the virtual environment.

vCenter Server components are user access control, core services, distributed services, plug-ins, and various interfaces.

Figure 12 shows the key components of vCenter Server.

**Figure 12.** vCenter Server Components

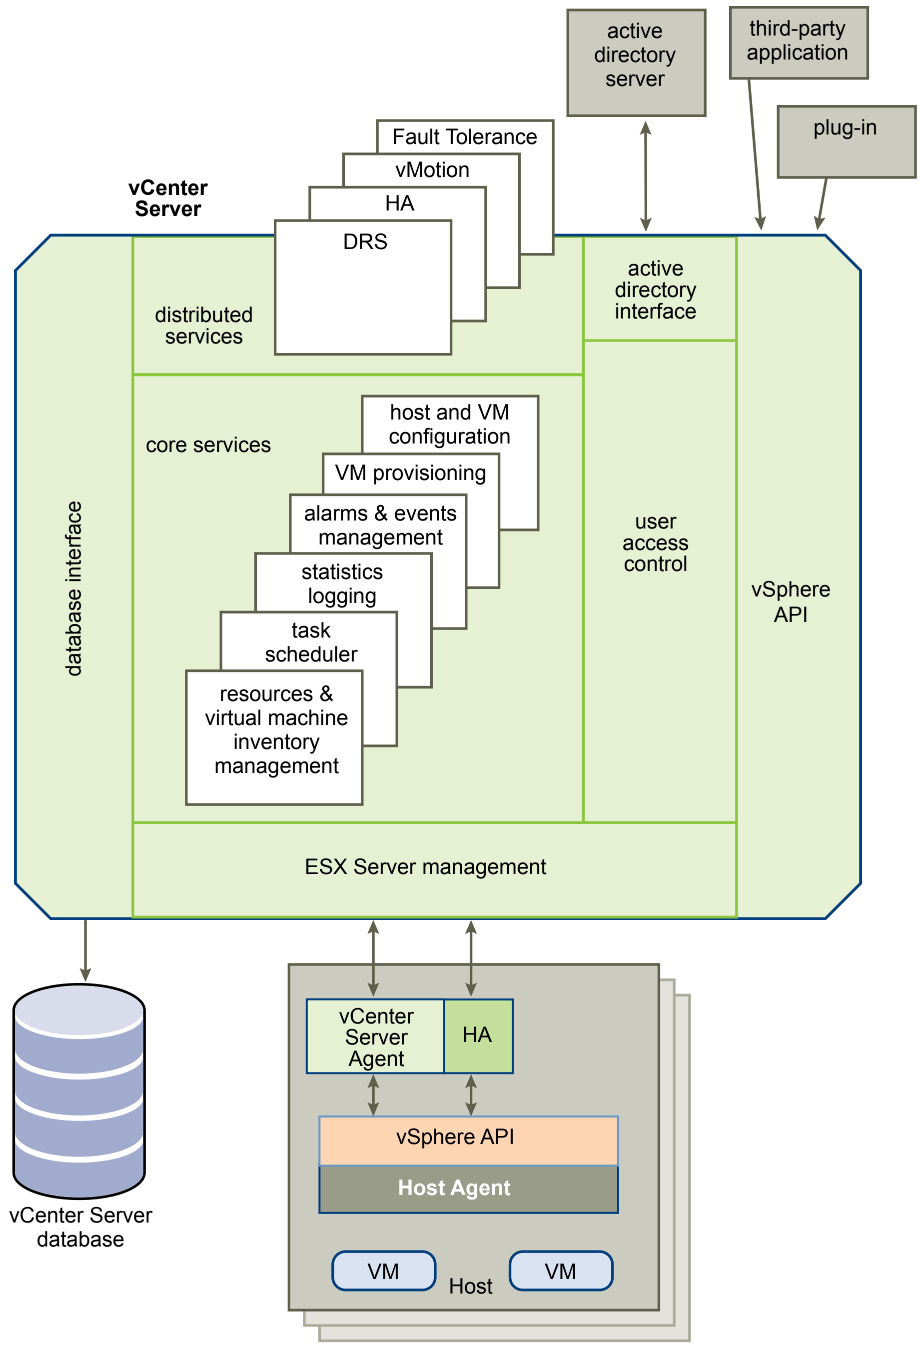


The User Access Control component allows the system administrator to create and manage different levels of access to vCenter Server for different classes of users.

For example, a user class might manage and configure the physical virtualization server hardware in the datacenter. Another user class might manage virtual resources within a particular resource pool in the virtual machine cluster.

## vCenter Server Core Services

Core Services are basic management services for a virtual datacenter.

Core Services include the following services:

| | |
|---|---|
| **Virtual machine provisioning** | Guides and automates the provisioning of virtual machines and their resources. |
| **Host and VM configuration** | Allows the configuration of hosts and virtual machines. |
| **Resources and virtual machine inventory management** | Organizes virtual machines and resources in the virtual environment and facilitates their management. |
| **Statistics and logging** | Logs and reports on the performance and resource use statistics of datacenter elements, such as virtual machines, hosts, storage, and clusters. |
| **Alarms and event management** | Tracks and warns users on potential resource overuse or event conditions. You can set alarms to trigger on events and notify when critical error conditions occur. Alarms are triggered only when they satisfy certain time conditions to minimize the number of false triggers. |
| **Task scheduler** | Schedules actions such as vMotion to occur at a given time. |
| **Consolidation** | Analyzes the capacity and use of a datacenter's physical resources. Provides recommendations for improving use by discovering physical systems that can be converted to virtual machines and consolidated onto ESX/ESXi. Automates the consolidation process, but also provides the user with flexibility in adjusting consolidation parameters. |
| **vApp** | A vApp has the same basic operation as a virtual machine, but can contain multiple virtual machines or appliances. With vApps, you can perform operations on multitier applications as separate entities (for example, clone, power on and off, and monitor). vApps package and manage those applications. |

Distributed Services are solutions that extend VMware vSphere capabilities beyond a single physical server. These solutions include: VMware DRS, VMware HA, and VMware vMotion. Distributed Services allow the configuration and management of these solutions centrally from vCenter Server.

Multiple vCenter Server systems can be combined into a single connected group. When a vCenter Server host is part of a connected group, you can view and manage the inventories of all vCenter Server hosts in that group.

## vCenter Server Plug-Ins

Plug-ins are applications that you can install on top of vCenter Server. Plug-ins add additional features and functionality.

vCenter Server Plug-ins include the following:

**VMware vCenter Converter** Enables users to convert physical machines, and virtual machines in a variety of formats, to ESX/ESXi virtual machines. Converted systems can be imported into any location in the vCenter Server inventory.

**VMware Update Manager** Enables security administrators to enforce security standards across ESX/ESXi hosts and managed virtual machines. This plug-in lets you create user-defined security baselines that represent a set of security standards. Security administrators can compare hosts and virtual machines against these baselines to identify and remediate virtual machines that are not in compliance.

## vCenter Server Interfaces

vCenter Server interfaces integrate vCenter Server with third party products and applications.

vCenter Server has the following key interfaces:

**ESX management** Interfaces with the vCenter Server agent to manage each physical server in the datacenter.

**VMware vSphere API** Interfaces with VMware management clients and third-party solutions.

**Database interface** Connects to Oracle, Microsoft SQL Server, or IBM DB2 to store information, such as virtual machine configurations, host configurations, resources and virtual machine inventory, performance statistics, events, alarms, user permissions, and roles.

**Active Directory interface** Connects to Active Directory to obtain user access control information.

## Communication Between vCenter Server and ESX

vCenter Server communicates with the ESX/ESXi host agent through the VMware vSphere API.

When you first add a host to vCenter Server, vCenter Server sends a vCenter Server agent to run on the host. As Figure 13 shows, that agent communicates with the host agent.

**Figure 13.** Host Agent

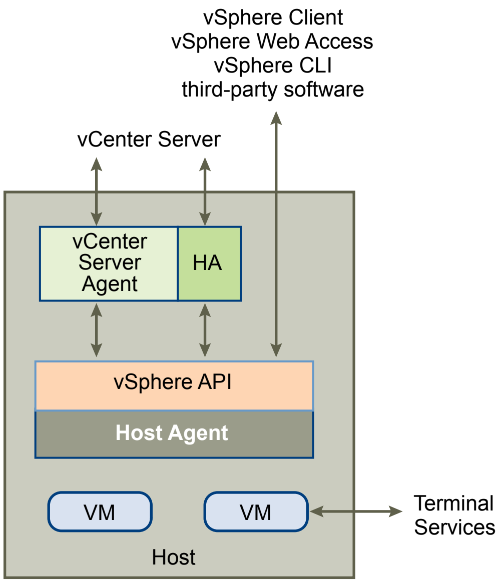


The vCenter Server agent acts as a small vCenter Server to perform the following functions:

- Relays and enforces resource allocation decisions made in vCenter Server, including those that the DRS engine sends.
- Passes virtual machine provisioning and configuration change commands to the host agent.
- Passes host configuration change commands to the host agent.
- Collects performance statistics, alarms, and error conditions from the host agent and sends them to the vCenter Server.
- Allows management of ESX/ESXi hosts at different release versions.

## Accessing the Virtual Datacenter

Users can access the VMware vSphere datacenter through the vSphere Client, Web Access through a Web browser, or terminal services (such as Windows Terminal Services).

Only physical host administrators in special circumstances should access hosts. All relevant functions that can be done on the host can also be done in vCenter Server.

**Figure 14.** VMware vSphere Access and Control

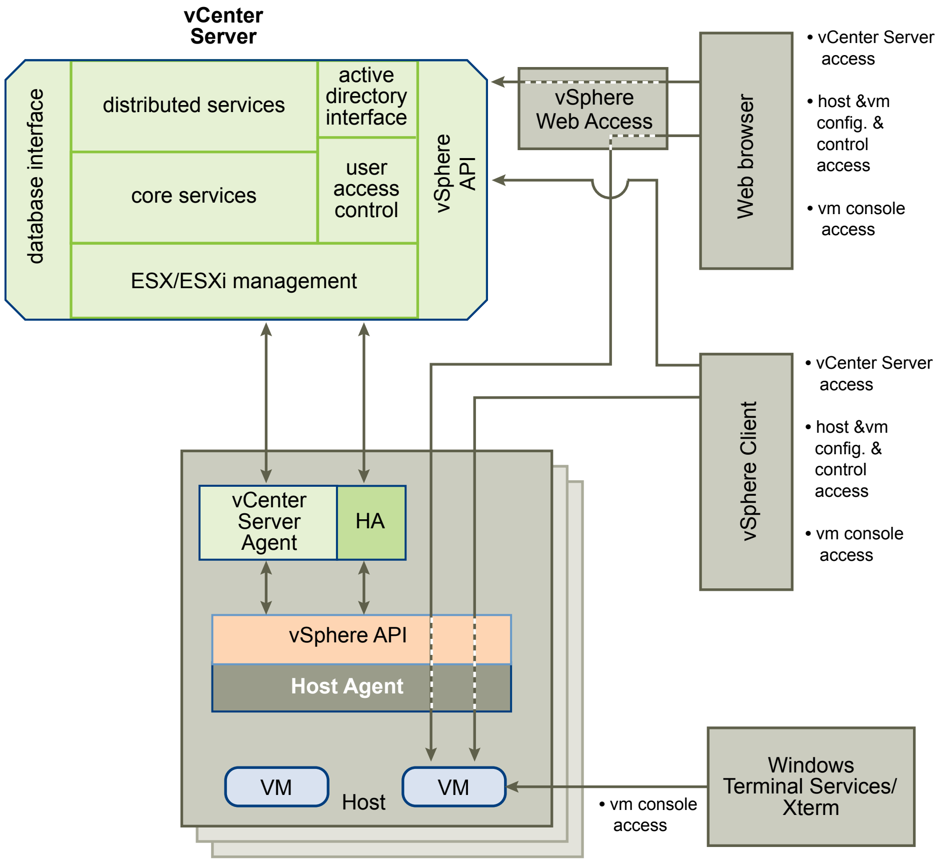


The vSphere Client accesses vCenter Server through the VMware API. After the user is authenticated, a session starts in vCenter Server, and the user sees the resources and virtual machines that are assigned to the user. For virtual machine console access, the vSphere Client first obtains the virtual machine location from vCenter Server through the VMware API. It then connects to the appropriate host and provides access to the virtual machine console.

---

**Note** vSphere Web Access cannot be used to access a host running ESXi 4.1.

---

## First Time Use

The vSphere Client includes embedded assistance that guides users who are new to virtualization concepts through the steps to set up their virtual infrastructure. This embedded assistance is in-line content presented in the vSphere Client GUI and an online tutorial. You can turn off the assistance for experienced users. You can turn on assistance when new users are introduced to the system.

## Web Access

Users can access vCenter Server through the Web browser by first pointing the browser to an Apache Tomcat Server set up by vCenter Server. The Apache Tomcat Server mediates the communication between the browser and vCenter Server through the VMware API.

To access the virtual machine consoles through the Web browser, users can use the bookmark that vCenter Server creates. The bookmark first points to the vSphere Web Access.

vSphere Web Access resolves the physical location of the virtual machine and redirects the Web browser to ESX/ESXi where the virtual machine resides.

If the virtual machine is running and the user knows the IP address of the virtual machine, the user can access the virtual machine console by using standard tools, such as Windows Terminal Services.

**Note** Web Access is turned off by default for ESX hosts.

## Additional Resources

You must perform additional tasks to set up a virtual infrastructure. Each task described has references to the documentation that contains details about the task.

Table 2 lists the tasks and references documentation for setting up VMware vSphere. Documentation also exists for the following topics:

- Documentation road map and quick start
- Virtual machine mobility planning
- VMware SDK and API developer resources
- Configuration maximums and release notes

**Table 2.** Documentation

| Tasks | Documents |
|---|---|
| Install vCenter Server and the vSphere Client | *ESX and vCenter Server Installation Guide*<br>*ESXi Installable and vCenter Server Setup Guide* |
| Install ESX 4.1<br>Install and Configure ESXi 4.1 Installable | *ESX and vCenter Server Installation Guide*<br>*ESXi Installable and vCenter Server Setup Guide* |
| Upgrade vCenter Server, vSphere Clients, ESX, or ESXi | *Upgrade Guide* |
| Obtain and install licenses | *Datacenter Administration Guide*<br>*ESXi Installable and vCenter Server Setup Guide* |
| Configure storage | *iSCSI SAN Configuration Guide*<br>*Fibre Channel SAN Configuration Guide*<br>*ESX Configuration Guide*<br>*ESXi Configuration Guide* |
| Configure networks | *ESX Configuration Guide*<br>*ESXi Configuration Guide* |
| Configure security<br>■ ESX security<br>■ User management<br>■ Virtual machine patch management | *ESX Configuration Guide*<br>*ESXi Configuration Guide*<br>*Virtual Machine Administration Guide* |
| Deploy virtual machines | *Virtual Machine Administration Guide*<br>*Guest Operating System Installation Guide* |
| Import physical systems, virtual machines, virtual appliances, or backup images into the virtual infrastructure | *Datacenter Administration Guide*<br>*Virtual Machine Administration Guide*<br>*VMware Converter Enterprise Administration Guide* |
| Configure distributed services<br>■ VMware HA and Fault Tolerance<br>■ VMware DRS | *VMware Availability Guide*<br>*Resource Management Guide*<br>*Virtual Machine Backup Guide* |

# Index