


Introduction

ESX Server 3.0.1 and VirtualCenter 2.0.1



Introduction

Revision: 20060925

Item: VI-ENG-Q206-213

You can find the most up-to-date technical documentation on our 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

© 2006-2008 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,961,941, 6,961,806 and 6,944,699; patents pending.

VMware, the VMware “boxes” logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.

3145 Porter Drive
Palo Alto, CA 94304
www.vmware.com

Contents

Preface	5
1 VMware Infrastructure Introduction	9
Physical Topology of VI Datacenter	12
Computing Servers	12
Storage Networks and Arrays	13
IP Networks	13
Management Server	13
Desktop Clients	13
Virtual Datacenter Architecture	14
Hosts, Clusters, and Resource Pools	15
VMware VMotion, VMware DRS, and VMware HA	17
Network Architecture	19
Storage Architecture	22
VMware Consolidated Backup	25
VirtualCenter Management Server	26
Communication Between VirtualCenter and ESX Server	28
Accessing the Virtual Datacenter	29
Conclusion	30
Glossary	31

Preface

This preface describes the contents of this *Introduction* and provides pointers to VMware® technical and educational resources.

This preface contains the following topics:

- [“About This Book”](#) on page 6
- [“Technical Support and Education Resources”](#) on page 8

About This Book

This manual, the *Introduction*, describes the VMware Infrastructure components.

Revision History

This manual is revised with each release of the product or when necessary. A revised version can contain minor or major changes. [Table P-1](#) provides you with the revision history of this manual.

Table P-1. Revision History

Revision	Description
20060615	ESX Server 3.0 and VirtualCenter 2.0 version of the VMware Infrastructure 3 <i>Introduction</i> . This is the first edition of this manual.
20060925	ESX Server 3.0.1 and VirtualCenter 2.0.1 version of the VMware Infrastructure 3 <i>Introduction</i> . This edition contains minor changes.

Intended Audience

The information presented in this manual is written for experienced Windows or Linux system administrators and who are familiar with virtual machine technology datacenter operations.

Document Feedback

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

docfeedback@vmware.com

VMware Infrastructure Documentation

The VMware Infrastructure documentation consists of the combined VirtualCenter and ESX Server documentation set.

You can access the most current versions of this manual and other books by going to:

<http://www.vmware.com/support/pubs>

Conventions

[Table P-2](#) illustrates the typographic conventions used in this manual.

Table P-2. Conventions Used in This Manual

Style	Elements
Blue (online only)	Cross-references and email addresses
Blue boldface (online only)	Links
Black boldface	User interface elements such as button names and menu items
Monospace	Commands, filenames, directories, and paths
Monospace bold	User input
<i>Italic</i>	Document titles, glossary terms, and occasional emphasis
< Name >	Variable and parameter names

Abbreviations Used in Graphics

The graphics in this manual use the abbreviations listed in [Table P-3](#).

Table P-3. Abbreviations

Abbreviation	Description
VC	VirtualCenter
VI	Virtual Infrastructure Client
server	VirtualCenter Server
database	VirtualCenter database
host <i>n</i>	VirtualCenter managed hosts
VM#	Virtual machines on a managed host
user#	User with access permissions
dsk#	Storage disk for the managed host
datastore	Storage for the managed host
SAN	Storage area network type datastore shared between managed hosts
tplt	Template

Technical Support and Education Resources

The following sections describe the technical support resources available to you.

Self-Service Support

Use the VMware Technology Network (VMTN) for self-help tools and technical information:

- Product information – <http://www.vmware.com/products/>
- Technology information – <http://www.vmware.com/vcommunity/technology>
- Documentation – <http://www.vmware.com/support/pubs>
- VMTN Knowledge Base – <http://www.vmware.com/support/kb>
- Discussion forums – <http://www.vmware.com/community>
- User groups – <http://www.vmware.com/vcommunity/usergroups.html>

For more information about the VMware Technology Network, go to <http://www.vmtn.net>.

Online and Telephone Support

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 Education Services

VMware courses offer extensive hands-on labs, case study examples, and course materials designed to be used as on-the-job reference tools. For more information about VMware Education Services, go to <http://mylearn1.vmware.com/mgrreg/index.cfm>.

VMware Infrastructure Introduction

1

VMware Infrastructure is a full infrastructure virtualization suite that provides comprehensive virtualization, management, resource optimization, application availability, and operational automation capabilities in an integrated offering. VMware Infrastructure virtualizes and aggregates the underlying physical hardware resources across multiple systems and provides pools of virtual resources to datacenter in the virtual environment.

In addition, VMware Infrastructure brings about a set of distributed services that enables fine-grain, policy-driven resource allocation, high availability, and consolidated backup of the entire virtual datacenter. These distributed services enable an IT organization to establish and meet their production Service Level Agreements with their customers in a cost effective manner.

The relationships among the various components of the VMware Infrastructure are shown in [Figure 1-1](#).

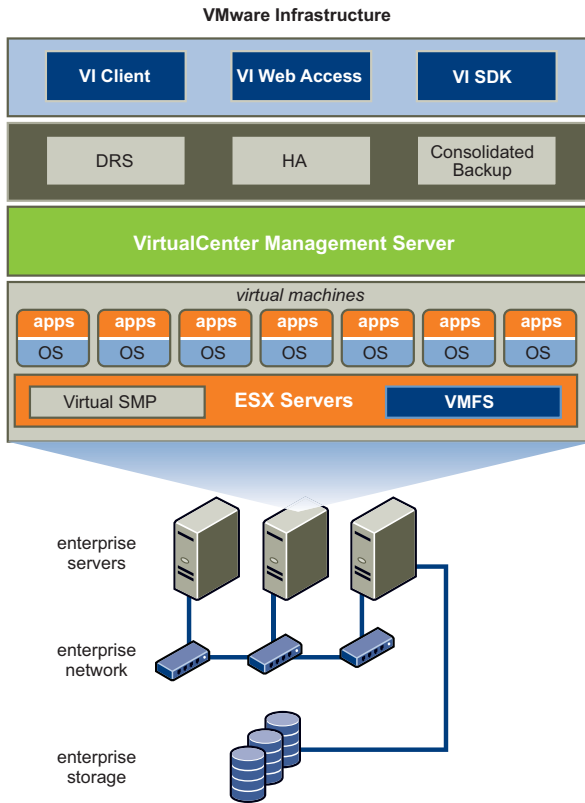


Figure 1-1. VMware Infrastructure

VMware Infrastructure includes the following components shown in [Figure 1-1](#):

VMware ESX Server – A robust, production-proven virtualization layer run on physical servers that abstracts processor, memory, storage, and networking resources into multiple virtual machines.

VirtualCenter Management Server (VirtualCenter Server) – The central point for configuring, provisioning, and managing virtualized IT environments.

Virtual Infrastructure Client (VI Client) – An interface that allows users to connect remotely to the VirtualCenter Server or individual ESX Servers from any Windows PC.

Virtual Infrastructure Web Access (VI 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 Server virtual machines.

VMware Virtual Symmetric Multi-Processing (SMP) – Feature that enables a single virtual machine to use multiple physical processors simultaneously.

VMware VMotion – Feature that 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.

VMware HA – Feature that provides easy-to-use, cost-effective high availability for applications running in virtual machines. In the event of server failure, affected virtual machines are automatically 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.

VMware Consolidated Backup (Consolidated Backup) – Feature that provides an easy-to-use, centralized facility for agent-free backup of virtual machines. It simplifies backup administration and reduces the load on ESX Servers.

VMware Infrastructure SDK – Feature that provides a standard interface for VMware and third-party solutions to access the VMware Infrastructure.

Physical Topology of VI Datacenter

As [Figure 1-2](#) shows, a typical VMware Infrastructure datacenter consists of basic physical building blocks such as x86 computing servers, storage networks and arrays, IP networks, a management server, and desktop clients.

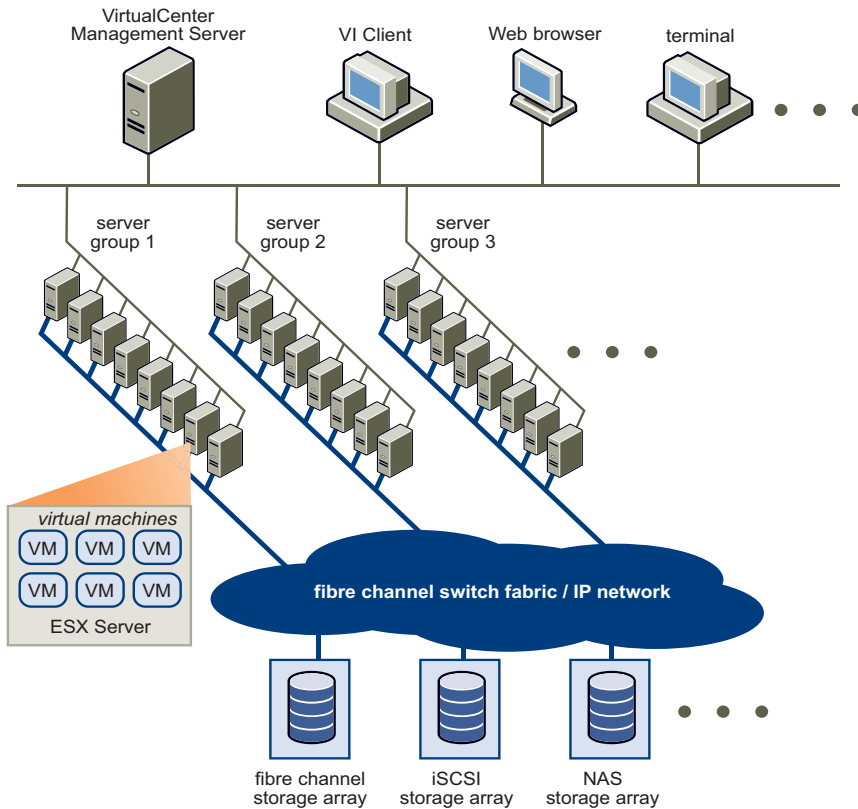


Figure 1-2. VMware Infrastructure Datacenter Physical Topology

Computing Servers

The computing servers are industry standard x86 servers that run VMware ESX Server on the bare metal. ESX Server provides resources for and runs the virtual machines. Each computing server is referred to as a standalone host in the virtual environment. A number of similarly configured x86 servers can be grouped together 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

Fiber Channel SAN arrays, iSCSI SAN arrays, and NAS arrays are widely used storage technologies supported by VMware Infrastructure to meet different datacenter storage needs. Sharing the storage arrays between (by connecting them to) groups of servers via storage area networks allows aggregation of the storage resources and provides more flexibility in provisioning them to virtual machines.

IP Networks

Each computing server can have multiple gigabit Ethernet network interface cards (NICs) to provide high bandwidth and reliable networking to the entire datacenter.

Management Server

The VirtualCenter Management Server provides a convenient single point of control to the datacenter. It runs on top Windows 2003 Server to provide many 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 accomplishes 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 set by the system administrator.

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

The architecture of VirtualCenter Management Server will be described in detail in later sections.

Desktop Clients

VMware Infrastructure provides a selection of interfaces for datacenter management and virtual machine access. Users can choose the interface that best meets their needs: Virtual Infrastructure Client (VI Client), Web Access through a Web browser, or terminal services (such as Windows Terminal Services or Xterm).

Virtual Datacenter Architecture

VMware Infrastructure virtualizes the entire IT infrastructure including servers, storage, and networks. It aggregates these heterogeneous resources and presents a simple and uniform set of elements in the virtual environment. With VMware Infrastructure, IT resources can be managed like a shared utility and dynamically provisioned to different business units and projects without worrying about the underlying hardware differences and limitations.

Figure 1-3 shows the key elements in virtual datacenter. You can view, configure, and manage these key elements using VirtualCenter Server. These elements include:

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

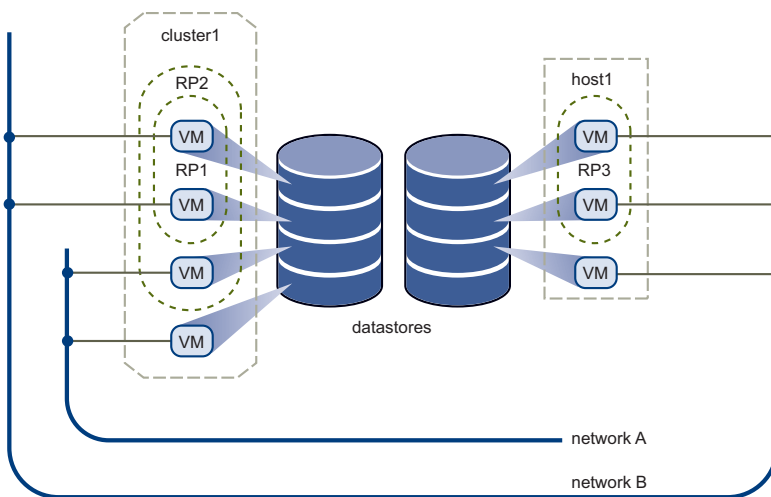


Figure 1-3. Virtual Datacenter Architecture

A host is the virtual representation of the computing and memory resources of a physical machine running ESX Server. When one or more physical machines are grouped together 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 local SCSI disks of the server, the Fiber Channel SAN disk arrays, the iSCSI SAN disk arrays, or Network Attached Storage (NAS) arrays.

Networks in the virtual environment connect virtual machines to each other or to the physical network outside of the virtual datacenter.

Virtual machines are designated to a particular host, cluster or resource pool, and a datastore when they are created. A virtual machine consumes resources like an appliance consuming electricity. While in powered-off, suspended, or idle state, they consume no resources. After being powered-on they 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, no purchase order is required, no waiting, and no physical constraints to worry about. 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 just as though they were being installed on a physical machine. To make things easier, a virtual machine can even be provisioned with the operating system and applications already installed and configured.

Resources are provisioned to virtual machines based on the policies 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 will be prevented from being powered-on (to consume resources) if doing so would violate the resource allocation policies. For more information on resource management, see the *Resource Management Guide*.

The following sections examine in detail the virtual elements of the datacenter.

Hosts, Clusters, and Resource Pools (SEE UPDATE)

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 4 gigahertz each and 32 gigabytes of system memory, the host will have 32 gigahertz of computing power and 32 gigabytes of memory available for running virtual machines that are assigned to it.

A cluster acts and can be managed much like a host. 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, each server has four dual-core CPUs running at 4 gigahertz each and 32 gigabytes of memory. The cluster will then have 256 gigahertz of computing power and 256 gigabytes of memory available for the running virtual machines assigned to it.

Resource pools are partitions of computing and memory resources from a single host or a cluster. Any resource pool can be partitioned into smaller resource pools to further divide and assign resources to different groups or for different purposes. In other words, resource pools can be hierarchical and nested.

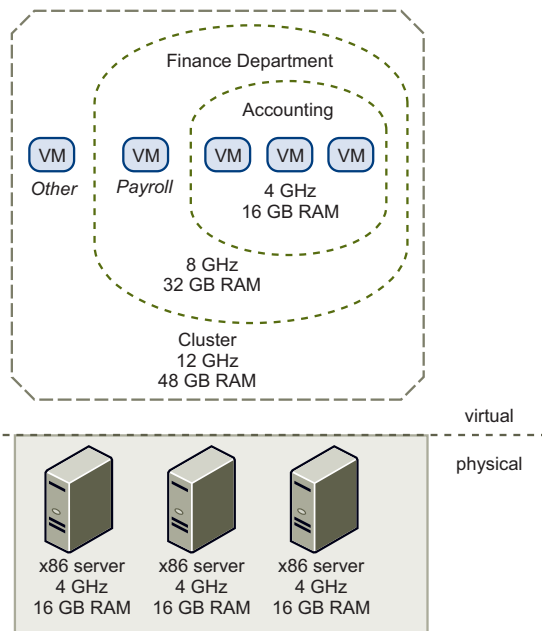


Figure 1-4. Hosts, Clusters, and Resource Pools

Figure 1-4 illustrates the use of resource pools. Three x86 servers with 4 gigahertz computing power and 16 gigabyte of memory each are aggregated to form a cluster of 12 gigahertz computing power and 48 gigahertz of memory. A resource pool (“Finance Department”) of reserves 8 gigahertz computing power and 32 gigabytes of from the cluster, leaving 4 gigahertz computing power and 16 gigabytes of memory reserved for the virtual machine “Others.” From the “Finance Department” resource pool, a smaller resource pool (“Accounting”) reserves 4 gigahertz computing power and 16 gigabytes of memory for the virtual machines from the accounting department. That leaves 4 gigahertz and 16 gigabytes of memory for the virtual machine called “Payroll.”

Resources reserved can be dynamically changed. Imagine that at year end, Accounting's workload increases, we want to bump up the resource pool "Accounting" from reserving 4 gigahertz of computing power to 6 gigahertz. We can make the change to the resource pool dynamically without shutting down the associated virtual machines. Resources reserved for a resource pool or virtual machine are not taken away immediately. They are taken on demand dynamically.

For example, if the 4 gigahertz of computing resources reserved for the Accounting department are not being used, virtual machine "Payroll" can make use of those gigahertz during its peak time. When Accounting demands those resources, "Payroll" will dynamically give them back. As a result, even though resources are reserved for different resource pools, they are not being wasted if not used by their owner.

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 infrastructure while still profiting from the efficiency of resource pooling.

VMware VMotion, VMware DRS, and VMware HA

VMware VMotion, DRS, and HA are distributed services that enable efficient and automated resource management and high virtual machine availability.

Virtual machines run on and consume resources from ESX Server. VMotion enables the migration of running virtual machines from one physical server to another without service interruption, as shown in [Figure 1-5](#). This allows virtual machines to move from a heavily loaded server to a lightly loaded one. The effect is a more efficient assignment of resources. With VMotion, resources can be dynamically reallocated to virtual machines across physical servers.




Figure 1-5. Migration with VMotion

VMware DRS aids in resource control and management capability in the virtual datacenter. A cluster can be viewed as an aggregation of the computing and memory resources of the underlying physical hosts put together in a single pool. Virtual

machines can be assigned to that pool. DRS monitors the workload of the running virtual machines and the resource utilization of the hosts to assign resources.

Using VMotion and an intelligent resource scheduler, VMware DRS automates the task of assigning virtual machines to servers within the cluster to use the computing and memory resources of that server as shown in [Figure 1-6](#). DRS does the calculation and automates the pairing.

If a new physical server is made available, DRS automatically redistributes the virtual machines using VMotion to balance the workloads. If a physical server must be taken down for any reason, DRS automatically reassigns its virtual machines to other servers.




Figure 1-6. VMware DRS

VMware HA offers a simple and low cost high availability alternative to application clustering. It enables quick restart of virtual machines on a different physical server within a cluster automatically if the hosting server fails. All applications within the virtual machines enjoy the high availability benefit, not just one (through application clustering).

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, and loss of a heartbeat initiates the process of restarting all affected virtual machines on other hosts. See [Figure 1-7](#). HA 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 1-7. VMware HA

Network Architecture

VMware Infrastructure is the only solution that brings about a rich set of virtual networking elements that makes networking the virtual machines in the data center as easy and simple as in the physical environment. Furthermore, it enables a new set of capabilities not possible in the physical environment because many of the limitations in the physical world don't apply.

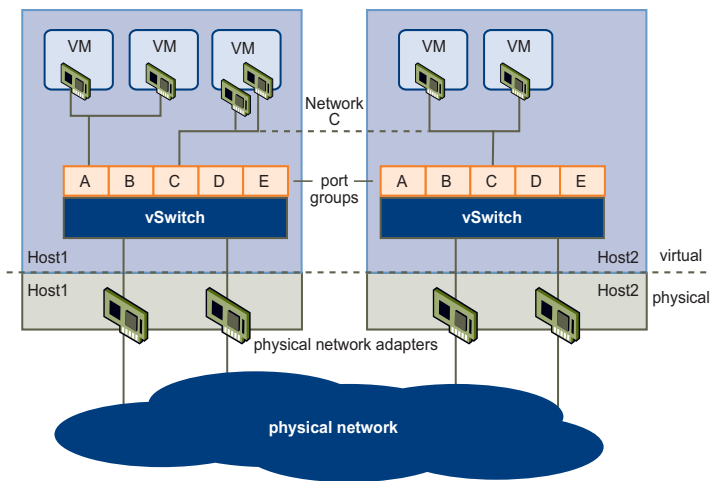


Figure 1-8. Networking

Figure 1-8 shows the relationship between the networks inside and outside the virtual environment. The virtual environment provides similar networking elements as the physical world. They are virtual network interface cards (vNIC), virtual switches (vSwitch), and port groups.

Like a physical machine, each virtual machine has its own vNIC. The operating system and applications talk to the vNIC through a standard device driver or a VMware optimized device driver just as though the vNIC is a physical NIC. To the outside world, the vNIC has its own MAC address and one or more IP addresses, and responds to the standard Ethernet protocol exactly as a physical NIC would. In fact, an outside agent does not know 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 on NIC teaming, see the *Server 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. All 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.

Port groups can be configured to enforce a number of policies that provide enhanced networking security, network segmentation, better performance, higher availability, and traffic management:

- **Layer 2 security options** – Enforces what vNICs in a virtual machine can do by controlling promiscuous mode, MAC address change, or forged transmits.
- **VLAN support** – Allows virtual networks to join a physical VLANs or support QOS policies.
- **Traffic shaping** – Defines average bandwidth, peak bandwidth, and burst size. These are policies that can be set 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 (SEE UPDATE)

The VMware Infrastructure storage architecture, shown in [Figure 1-9](#), consists of layers of abstraction that hide and manage the complexity and differences among physical storage subsystems.




Figure 1-9. Storage Architecture

To the applications and guest operating systems inside each virtual machine, the storage subsystem is a simple virtual Bus Logic or LSI SCSI host bus adapter connected to one or more virtual SCSI disks as shown in [Figure 1-9](#).

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

The datastore provides a simple model to allocate storage space to the individual virtual machines without exposing them to the complexity of the variety of physical storage technologies available, such as Fibre Channel SAN, iSCSI SAN, direct attached storage, and NAS.

A virtual machine is stored as a set of files in a directory in the datastore. A virtual disk inside each virtual machine is one or more files in the directory. As a result, you can operate on a virtual disk (copy, move, back up, and so on) just like a file. New virtual disks can be “hot-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 hot-added virtual disk or an existing virtual disk file is associated with a virtual machine.

Each datastore is physically a VMFS volume on a NAS device. Datastores can span multiple physical storage subsystems. As shown in [Figure 1-9](#), 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 automatically discovered and made available to all existing or new datastores. Storage capacity on a previously created VMFS volume (datastore) can be hot-extended without powering down physical hosts or storage subsystems by adding a new physical LUN from any of the storage subsystems that are visible to it. Conversely, if any of the LUNs within a VMFS volume (datastore) fails or becomes unavailable, only those virtual machines that touch that LUN are affected. All other virtual machines with virtual disks residing in other LUNs continue to function as normal.

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 enterprise-class crash consistency and recovery mechanisms, such as distributed journaling, a crash consistent virtual machine I/O path, and machine state snapshots. These mechanisms can aid quick root-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 is useful for supporting 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.

- Any use of Microsoft Clustering Services (MSCS) that spans physical hosts: virtual-to-virtual clusters as well as physical-to-virtual clusters. Cluster data and quorum disks should be configured as RDMs rather than as files on a shared VMFS.




Figure 1-10. Raw Device Mapping

An RDM can be thought of as a symbolic link from a VMFS volume to a raw LUN (see [Figure 1-10](#)). 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 Consolidated Backup

VMware Infrastructure's storage architecture enables a simple virtual machine backup solution: *VMware Consolidated Backup*. Consolidated Backup provides a centralized facility for LAN-free backup of virtual machines.

As shown in [Figure 1-11](#), Consolidated Backup works in conjunction with a third-party backup agent residing on a separate backup proxy server (not on the server running ESX Server) but does not require an agent inside the virtual machines.

The third-party backup agent manages the backup schedule. It starts Consolidated Backup when it is time to do a back up. When started, Consolidated Backup runs a set of pre-backup scripts to quiesce the virtual disks to take their snapshots. It then runs a set of post-thaw scripts to restore the virtual machine back to normal operation. At the same time, it mounts the disk snapshot to the backup proxy server. Finally, the third-party backup agent backs up the files on the mounted snapshot to its backup targets. By taking snapshots of the virtual disks and backing them up through a separate backup proxy server, Consolidated Backup provides a simple, less intrusive, and low-overhead backup solution for the virtual environment.




Figure 1-11. VMware Consolidated Backup

VirtualCenter Management Server

VirtualCenter Management Server provides a convenient centralized management cockpit for the datacenters. It aggregates physical resources from multiple ESX Servers and presents a central collection of simple and flexible resources for the system administrator to provision to virtual machines in the virtual environment.

The VirtualCenter Management Server components are user access control, core services, distributed services, and various interfaces.

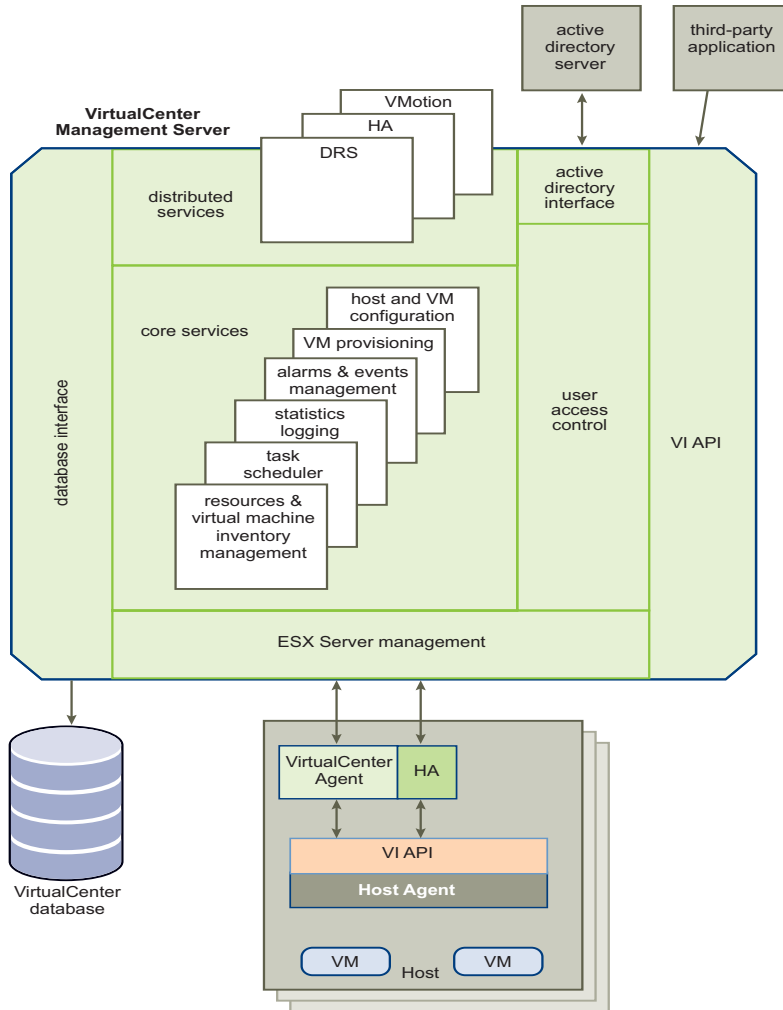


Figure 1-12. VirtualCenter Management Server Components

Figure 1-12 shows the following key components of the VirtualCenter Server.

The User Access Control allows the system administrator to create and manage different levels of access to the VirtualCenter for different users.

For example, there might be a user class that manages configuring the physical servers in the datacenter and there might be a different user class that manages only virtual resources within a particular resource pool.

Core Services are basic management services for a virtual datacenter. They include services such as:

- **VM Provisioning** – Guides and automates the provisioning of virtual machines
- **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 utilization statistics of datacenter elements, such as virtual machines, hosts, and clusters
- **Alarms and Event Management** – Tracks and warns users on potential resource over-utilization or event conditions.
- **Task Scheduler** – Schedules actions such as VMotion to happen at a given time

Distributed Services are solutions that extend VMware Infrastructure's capabilities to the next level such as VMware DRS, VMware HA, and VMware VMotion. Distributed Services allow the configuration and management of these solutions centrally from VirtualCenter Management Server.

VirtualCenter Server has four key interfaces:

- **ESX Server management** – Interfaces with the VirtualCenter agent to manage each physical server in the datacenter.
- **VMware Infrastructure API** – Interfaces with VMware management clients and third-party solutions.
- **Database interface** – Connects to Oracle or Microsoft SQL Server to store information, such as virtual machine configurations, host configurations,

Accessing the Virtual Datacenter

Users can manage the VMware Infrastructure datacenter or access the virtual machine console through three different means: the VI Client, Web Access through a Web browser, or terminal services (such as Windows Terminal Services or Xterm), as shown in [Figure 1-14](#). Accessing hosts should be done only by physical host administrators in special circumstances. All relevant functionality that can be done on the host can also be done in VirtualCenter Server.

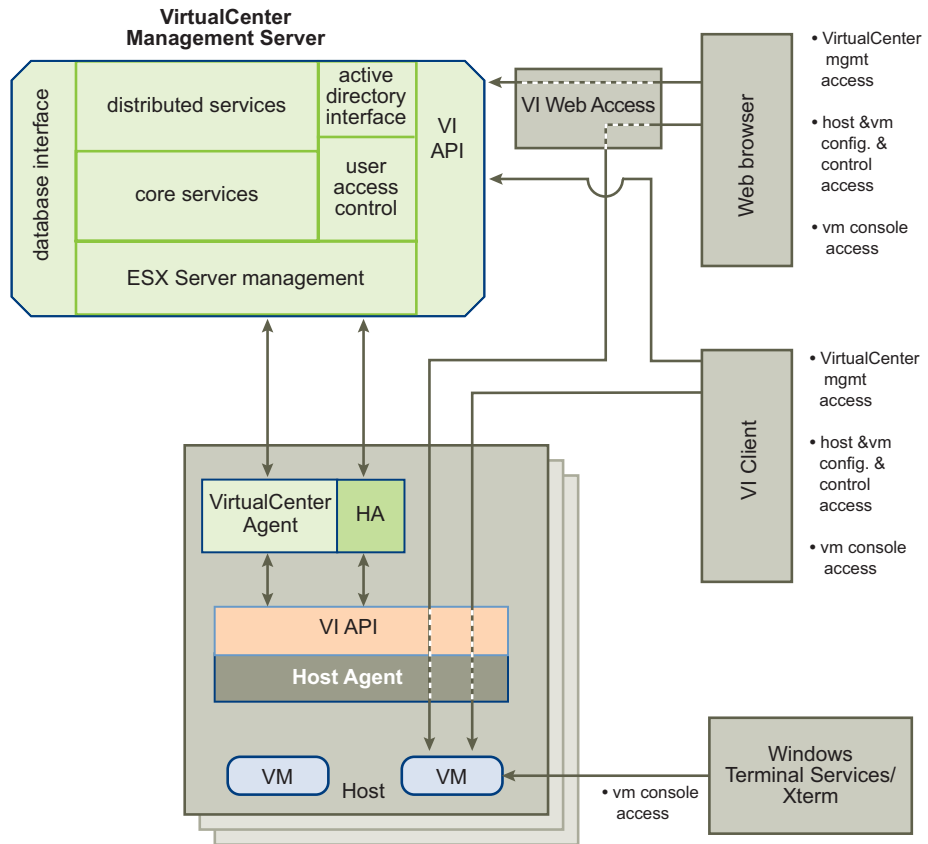


Figure 1-14. VMware Infrastructure Access and Control

The VI Client accesses VirtualCenter through the VMware API. After the user is authenticated, a session starts in VirtualCenter, and the user sees the resources and virtual machines that are assigned to the user. For virtual machine console access, the VI Client first gets the virtual machine location from VirtualCenter through the

VMware API. It then connects to the appropriate host and provides access to the virtual machine console.

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

To access the virtual machine consoles through the Web browser, users can make use of the bookmark that is created by VirtualCenter Server. The bookmark first points to the VI Web Access.

VI Web Access resolves the physical location of the virtual machine and redirects the Web browser to the ESX Server 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 also access the virtual machine console using standard tools, such as Windows Terminal Services or Xterm.

Conclusion

VMware Infrastructure provides a simple architecture in the virtual environment to allow companies to manage computing, storage, and networking resources without worrying about the underlying physical hardware. VI architecture allows enterprises to create and configure their datacenters and reallocate resources to different priorities without the time delay and cost of reconfiguring their physical hardware infrastructure.

With a suite of complementary virtualization and management services, such as VMware VMotion, VMware DRS, VMware HA, and VMware Consolidated Backup, VMware Infrastructure is the only product that provides a complete solution rather than a piecemeal approach to building datacenters in the virtual environment.

Glossary

A **administrative lockout**

A global setting providing password protection for Windows hosts. Administrative lockout restricts users from creating new virtual machines, editing virtual machine configurations, and changing network settings.

alarm

An entity that monitors one or more properties of a virtual machine, such as CPU load. Alarms use green, red, and yellow color coding to issue notifications as directed by the configurable alarm definition.

allocated disk

A type of virtual disk in which all disk space for the virtual machine is allocated at the time the disk is created. This is the default type of virtual disk created by VirtualCenter.

append mode

When software running in the virtual machine writes to a disk used in append mode, the changes appear to be written to the disk. In fact, however, they are stored in a temporary file (.REDO). If a system administrator deletes this redo-log file, the virtual machine returns to the state it was in the last time it was used in persistent mode. See also disk mode.

authorization role

A set of privileges grouped for convenient identification under names such as "Administrator."

C child

A managed entity grouped by a folder object or other managed entity. See also folder.

clone

(n.) A duplicate of a virtual machine.

(v.) To make a copy of a virtual machine. This process includes the option to customize the guest operating system of the new virtual machine. When a clone is created, VirtualCenter provides an option to customize the guest operating system of that virtual machine. Clones can be stored on any host within the same farm as the original virtual machine.

Clone Virtual Machine Wizard

A point-and-click interface for convenient, easy duplication of a virtual machine.

cluster

Server group in the virtual environment.

cluster compute resource

An extended compute resource that represents a cluster of hosts available for backing virtual machines.

compute resource

A managed object that represents either a single host or a cluster of hosts available for backing virtual machines.

customization

The process of customizing a guest operating system in a virtual machine as it is being deployed from a template or cloned from another existing virtual machine. Customization options include changing the new virtual machine identification and network information.

custom networking

Any type of network connection between virtual machines and the host that does not use the default bridged, host-only, or network address translation (NAT) configurations. For instance, different virtual machines can be connected to the host by separate networks or connected to each other and not to the host. Any network topology is possible. See also host-only networking.

D **data object**

A composite object that is passed by value between the client and the Web service. A data object has properties associated with it but does not have any operations of its own. See also managed object.

data source name

An ODBC object that you must configure to enable VirtualCenter Server to access a database.

datastore

Virtual representations of combinations of underlying physical storage resources in the datacenter. The datastore is the storage location for the virtual machine files. This can be a physical disk, a RAID, a SAN, or a partition on any of these.

disk mode

A property of a virtual disk that defines its external behavior but is completely invisible to the guest operating system. There are four modes: persistent (changes to the disk are always preserved across sessions), nonpersistent (changes are never preserved), undoable (changes are preserved at the user's discretion), and append (similar to undoable, but the changes are preserved until a system administrator deletes the redo-log file).

drag-and-drop

A feature of VMware VirtualCenter that allows you to move virtual machines easily between groups.

E **event**

An action that is of interest to VirtualCenter. Each event triggers an event message. Event messages are archived in the VirtualCenter database and appear in two locations: the Events option in the navigation bar and the Events tab for an object under the Inventory button.

F **fault**

A data object containing information about an exceptional condition encountered by an operation.

folder

A managed entity used to group other managed entities. The contents of a group are child entities with respect to the Folder object. Folders are typed by the kinds of child entities they can contain. See also *child*.

G **Go to snapshot**

To restore a snapshot of the active virtual machine. See also *Revert to snapshot*.

guest operating system

An operating system that runs inside a virtual machine.

H **headless**

Describes a program or application that runs in the background without any interface connected to it. A running virtual machine that has no consoles connected to it is running headless.

host

The physical computer on which the virtual machines managed by VirtualCenter are installed.

host agent

Software that, when installed on a virtual machine host, performs actions on behalf of a remote client.

host computer

The physical computer on which the VirtualCenter software is installed. It hosts the VirtualCenter virtual machines.

host-only networking

A type of network connection between a virtual machine and the host. Under host-only networking, a virtual machine is connected to the host on a private network, which normally is not visible outside the host. Multiple virtual machines configured with host-only networking on the same host are on the same network. See also *custom networking*.

I **inventory**

A hierarchical structure used by the VirtualCenter Server or the host agent to organize managed entities. This hierarchy is presented as a list in the left panel of

the VirtualCenter client window that provides a view of all the monitored objects in VirtualCenter.

L **license server**

A server that stores and allocates licenses.

M **managed entity**

A managed object that is present in the inventory. See also inventory.

managed object

A composite object that resides on a server and is passed between the client and the Web service only by reference. A managed object has operations associated with it but might not have properties. See also data object.

migration

Moving a virtual machine between hosts. Unless VMotion is used, the virtual machine must be powered off when you migrate it. See also migration with VMotion.

migration with VMotion

Moving a virtual machine that is powered on and has met selected requirements, including the activation of VMotion on both the source and target hosts. When you migrate a virtual machine using VMotion, the operations of the virtual machine can continue without interruption.

N **New Virtual Machine Wizard**

A point-and-click interface for convenient creation of a virtual machine configuration. It creates files that define the virtual machine, including a virtual machine configuration file and optionally a virtual disk or physical disk file.

nonpersistent mode

If you configure a virtual disk as an independent disk in nonpersistent mode, all disk writes issued by software running inside a virtual machine with a disk in nonpersistent mode appear to be written to disk but are in fact discarded after the virtual machine is powered off. As a result, a virtual disk or physical disk in independent-nonpersistent mode is not modified by activity in the virtual machine. See also persistent mode.

O operation

A function performed for a client by the Web service.

P parent

The source virtual machine from which you take a snapshot or make a clone. A clone has no continued link to its parent, but a snapshot must have access to the parent's virtual disk files. If you delete the parent virtual machine, any snapshot becomes permanently disabled. To prevent deletion, you can make the parent a template virtual machine. See also snapshot, template.

permission

A data object consisting of an authorization role, a user or group name, and a managed entity reference. Allows a specified user to access the entity with any of the privileges pertaining to the role.

persistent mode

If you configure a virtual disk as an independent disk in persistent mode, all disk writes issued by software running inside a virtual machine are immediately and permanently written to the virtual disk in persistent mode. As a result, a virtual disk or physical disk in independent-persistent mode behaves like a conventional disk drive on a physical computer. See also nonpersistent mode.

physical disk

A hard disk in a virtual machine that is mapped to a physical disk drive or partition on the host machine. A virtual machine's disk can be stored as a file on the host file system or on a local hard disk. When a virtual machine is configured to use a physical disk, VirtualCenter directly accesses the local disk or partition as a raw device (not as a file on a file system). See also virtual disk.

physical network

A network of physical machines that are connected so that they can send data to and receive data from each other. See also virtual network.

port group

A mechanism for setting policies that govern the network connected to it.

privilege

Authorization to perform a specific action or set of actions on a managed object or group of managed objects.

property

An attribute of a managed object or data object. A property can be a nested data object or a managed object reference.

property collector

A managed object used to control the reporting of managed object properties. The primary means of monitoring status on host machines.

R**raw device mapping (RDM)**

A mechanism that enables a virtual machine to have direct access to a LUN on the physical storage subsystem (Fibre Channel or iSCSI only).

read-only user

A role in which the user is allowed to view the inventory but not allowed to perform any tasks.

redo log

The file that stores the changes made to a disk in undoable or nonpersistent mode. You can permanently apply the changes saved in the redo log to a disk in undoable mode so they become part of the main disk files. For a disk in nonpersistent mode, however, the redo-log file is deleted when you power off or reset the virtual machine without writing any changes to the disk. See also disk mode.

resource pool

A division of computing resources used to manage allocations between virtual machines.

resume

To return a virtual machine to operation from its suspended state. When you resume a suspended virtual machine, all applications are in the same state they were when the virtual machine was suspended. See also suspend.

Revert to snapshot

To restore the status of the active virtual machine to its immediate parent snapshot. This parent is represented in the snapshot manager by the snapshot appearing to the immediate left of the You Are Here icon. See also *Go to snapshot*, *Snapshot manager*, *You Are Here icon*.

role

A defined set of access rights and privileges in VirtualCenter.

S **scheduled task**

A VirtualCenter activity that is configured to occur at designated times. The Tasks pane displays a list of selected activities and a means to schedule them.

service console

The command-line interface for an ESX Server system. It allows administrators to configure the ESX Server system. You can open the service console directly on an ESX Server system. If the ESX Server system's configuration allows Telnet or SSH connections, you can also connect remotely to the service console.

service instance

The managed entity at the root of the inventory. Clients must access the service instance to begin a session.

snapshot

A reproduction of the virtual machine just as it was when you took the snapshot, including the state of the data on all the virtual machine's disks and whether the virtual machine was powered on, powered off, or suspended. You can take snapshots of a virtual machine at any time and go to any snapshot at any time. You can take a snapshot when a virtual machine is powered on, powered off, or suspended. You can configure a virtual machine to exclude specified disks from snapshots.

Snapshot manager

A control that allows you to take actions on any of the snapshots associated with the selected virtual machine. See also *snapshot*.

suspend

To save the current state of a running virtual machine. To return a suspended virtual machine to operation, use the resume feature. See also resume.

T **template**

A master image of a virtual machine. This typically includes a specified operating system and a configuration that provides virtual counterparts to hardware components. Optionally, a template can include an installed guest operating system and a set of applications. Templates are used by VirtualCenter to create new virtual machines.

V **VirtualCenter administrator**

A role in which the user is allowed to set the user+role permissions and control the VirtualCenter licensing.

VirtualCenter agent

Installed on each virtual machine host, this software coordinates the actions received from the VirtualCenter Server.

VirtualCenter database

A persistent storage area for maintaining the status of each virtual machine and user managed in the VirtualCenter environment. Located on the same machine as the VirtualCenter Server.

VirtualCenter Server

A service that acts as a central administrator for VMware servers connected on a network. This service directs actions on the virtual machines and the virtual machine hosts. VirtualCenter Server is the working core of VirtualCenter.

virtual disk

A file or set of files that appears as a physical disk drive to a guest operating system. These files can be on the host machine or on a remote file system. See also physical disk.

Virtual Infrastructure

A system of hosts, agents, and clients that communicate to deploy and operate virtual machines. The total VMware solution to managing a data center. See also host, host agent, VirtualCenter Server.

Virtual Infrastructure Client (VI Client)

A user interface that runs locally in a Windows machine and provides access to the virtual machine's display. The Virtual Infrastructure Client runs on a networked machine. This can be on the same machine as the VirtualCenter Server or another networked machine. The Virtual Infrastructure Client requires a monitor for access to the virtual machine's display.

Virtual Infrastructure Web Access Client

A user interface that runs in a Web browser and provides access to the virtual machine's display.

virtual machine

A virtualized x86 PC environment in which a guest operating system and associated application software can run. Multiple virtual machines can operate on the same host system concurrently.

virtual machine configuration

The specification of which virtual devices, such as disks and memory, are present in a virtual machine and how they are mapped to host files and devices.

virtual machine configuration file

A file containing a virtual machine configuration. It is created when you create the virtual machine. It is used by VirtualCenter to identify and run a specific virtual machine.

virtual machine monitor (VMM)

Software that is responsible for virtualizing the CPUs.

Virtual Machine Properties

A point-and-click control panel used to view and modify the resource settings of all the virtual machines on a host.

Virtual Machine Settings Editor

A point-and-click control panel used to view and modify a virtual machine's settings. See also New Virtual Machine Wizard.

virtual network

A network connecting virtual machines that does not depend on physical hardware connections. For example, you can create a virtual network between a virtual machine and a host that has no external network connections. You can also create a LAN segment for communication between virtual machines on a team.

VMkernel

A high-performance operating system that occupies the virtualization layer and manages most of the physical resources on the hardware, including memory, physical processors, storage, and networking controllers.

VMODL

The interface definition language used in the Virtual Infrastructure SDK.

VMotion

A feature that enables you to move running virtual machines from one ESX Server system to another without interrupting service. It requires licensing on both the source and target hosts. VMotion is activated by the VirtualCenter agent. The VirtualCenter Server centrally coordinates all VMotion activities. See also migration with VMotion.

VMware File System (VMFS)

A file system that is optimized for storing virtual machines. One VMFS partition is supported per SCSI storage device or SAN. Each version of ESX Server uses a corresponding version of VMFS. For example, VMFS3 was introduced with ESX Server 3.

VMware HA (HA)

An optional feature that supports distributed availability services in an environment that includes ESX Server and VirtualCenter. If you have configured DRS and one of the hosts managed by VirtualCenter Server goes down, all virtual machines on that host are immediately restarted on another host.

VMware Tools

A suite of utilities and drivers that enhances the performance and functionality of your guest operating system. Key features of VMware Tools include some or all of the following, depending on your guest operating system: an SVGA driver, a mouse driver, the VMware Tools control panel, and support for such features as shared folders, drag-and-drop in Windows guests, shrinking virtual disks, time

synchronization with the host, VMware Tools scripts, and connecting and disconnecting devices while the virtual machine is running.

VMware virtual machine console

An interface to a virtual machine that provides access to one or more virtual machines on the local host or a remote host running VirtualCenter. You can view the virtual machine's display to run programs within it or modify guest operating system settings. In addition, you can change the virtual machine's configuration, install the guest operating system, or run the virtual machine in full screen mode.

W Web service

A programming interface based on SOAP and WSDL.

Y You Are Here icon

An icon in the Snapshot manager that indicates the current status of the active virtual machine. Checking the position of this icon can help you decide whether to revert to a snapshot or go to a snapshot. See also [Go to snapshot](#), [Revert to snapshot](#), [Snapshot manager](#).

Updates for the Introduction

ESX Server 3.0.1 and VirtualCenter 2.0.1

Last Updated: February 22, 2008

This document provides you with updates to ESX Server 3.0.1 and VirtualCenter 2.0.1 version of the *Introduction*. Updated descriptions, procedures, and graphics are organized by page number so you can easily locate the areas of the guide that have changes. If the change spans multiple sequential pages, this document provides the starting page number only.

The following is a list of the *Introduction* page updates in this document:

[“Updates for the Storage Architecture Discussion on Page 22”](#)

[“Updates for the Hosts, Clusters, and Resource Pools Discussion on Page 15”](#)

Updates for the Storage Architecture Section on **Page 22**

In the [“Storage Architecture \(SEE UPDATE\)”](#) section, the first sentence of the fifth paragraph after figure 1-9 inaccurately implies that datastores are VMFS volumes on NAS devices.

The sentence should read as follows:

Each datastore is physically a VMFS volume (or, for NAS datastores, a NFS volume with VMFS characteristics) on a storage device.

Updates for the Hosts, Clusters, and Resource Pools Section on [Page 15](#)

In the “[Hosts, Clusters, and Resource Pools \(SEE UPDATE\)](#)” section, the second paragraph contains several errors and inaccuracies including use of “gigahertz of memory” instead of “gigabytes of memory”.

The paragraph should read as follows:

[Figure 1-4](#) illustrates the use of resource pools. Three x86 servers with 4 gigahertz of computing power and 16 gigabytes of memory each are aggregated to form a cluster of 12 gigahertz computing power and 48 gigabytes of memory. A resource pool (Finance Department) reserves 8 gigahertz of computing power and 32 gigabytes of memory from the cluster, leaving 4 gigahertz of computing power and 16 gigabytes of memory reserved for the virtual machine Others.