

WHITE PAPER

VMware Infrastructure Architecture Overview



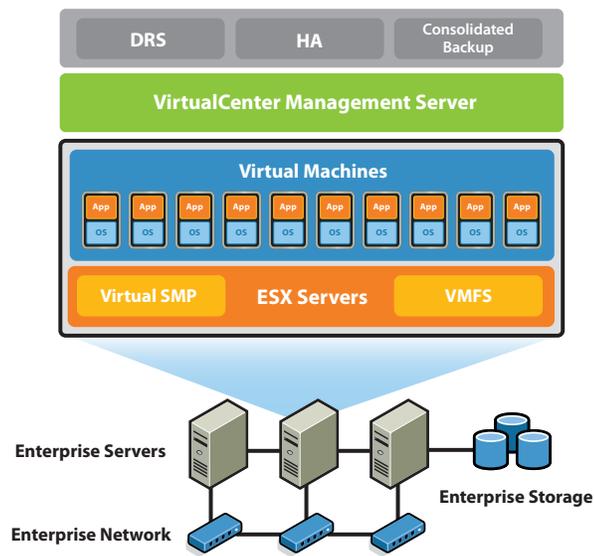
Table of Contents

Physical Topology of the VMware Infrastructure Data Center4
Virtual Data Center Architecture5
Hosts, Clusters and Resource Pools6
VMware VMotion, VMware DRS and VMware HA7
Networking Architecture8
Storage Architecture9
VMware Consolidated Backup 10
ESX Server External Interfacing Components 10
VirtualCenter Management Server Architecture 11
Conclusion 13

VMware Infrastructure Architecture Overview

VMware® Infrastructure is the industry's first full infrastructure virtualization suite that allows enterprises and small businesses alike to transform, manage and optimize their IT systems infrastructure through virtualization. VMware Infrastructure delivers comprehensive virtualization, management, resource optimization, application availability and operational automation capabilities in an integrated offering.

Figure 1-1: VMware Infrastructure



VMware Infrastructure includes the following components as shown in Figure 1-1:

- VMware ESX Server – A production-proven virtualization layer run on physical servers that abstract processor, memory, storage and networking resources to be provisioned to multiple virtual machines
- VMware Virtual Machine File System (VMFS) – A high-performance cluster file system for virtual machines
- VMware Virtual Symmetric Multi-Processing (SMP) – Enables a single virtual machine to use multiple physical processors simultaneously
- VirtualCenter Management Server – The central point for configuring, provisioning and managing virtualized IT infrastructure
- Virtual Infrastructure Client (VI Client) – An interface that allows administrators and users to connect remotely to the VirtualCenter Management Server or individual ESX Server installations from any Windows PC

What is Virtualization and What are Virtual Machines?

Virtualization is an abstraction layer that decouples the physical hardware from the operating system to deliver greater IT resource utilization and flexibility. Virtualization allows multiple virtual machines, with heterogeneous operating systems (e.g., Windows 2003 Server and Linux) and applications to run in isolation, side-by-side on the same physical machine.

A virtual machine is the representation of a physical machine by software. It has its own set of virtual hardware (e.g., RAM, CPU, NIC, hard disks, etc.) upon which an operating system and applications are loaded. The operating system sees a consistent, normalized set of hardware regardless of the actual physical hardware components. VMware virtual machines contain advanced hardware features such as 64-bit computing and virtual symmetric multiprocessing.

For more information on virtualization, please read the Virtualization Overview VMware white paper.

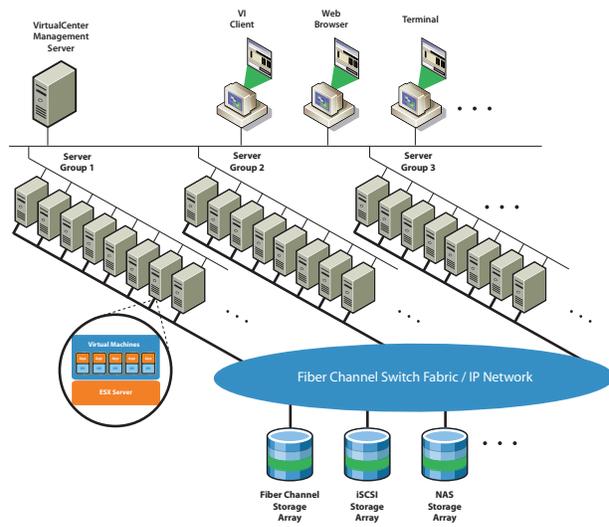
- Virtual Infrastructure Web Access – A Web interface for virtual machine management and remote consoles access
- VMware VMotion™ – Enables the live migration of running virtual machines from one physical server to another with zero downtime, continuous service availability and complete transaction integrity
- VMware High Availability (HA) – 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) – Intelligently allocates and balances computing capacity dynamically across collections of hardware resources for virtual machines
- VMware Consolidated Backup – Provides an easy to use, centralized facility for agent-free backup of virtual machines. It simplifies backup administration and reduces the load on ESX Server installations
- VMware Infrastructure SDK – Provides a standard interface for VMware and third-party solutions to access VMware Infrastructure

The following sections describe the architecture of VMware Infrastructure, beginning with the elements that make up its physical topology, followed by the virtual, or logical, view of VMware Infrastructure where the relationships between the virtual architectural elements and the physical world are explored. Lastly, the architectures of two core VMware Infrastructure components are discussed in further detail.

Physical Topology of the VMware Infrastructure Data Center

With VMware Infrastructure, IT departments can build a virtual data center using their existing industry standard technology and hardware. There is no need to purchase specialized hardware. In addition, VMware Infrastructure allows users to create a virtual data center that is centrally managed by management servers and can be controlled through a wide selection of interfaces.

Figure 1-2: VMware Infrastructure Data Center Physical Building Blocks



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

Computing Servers

The computing servers are industry standard x86 servers that run VMware ESX Server on the “bare metal.” 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 data center 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.

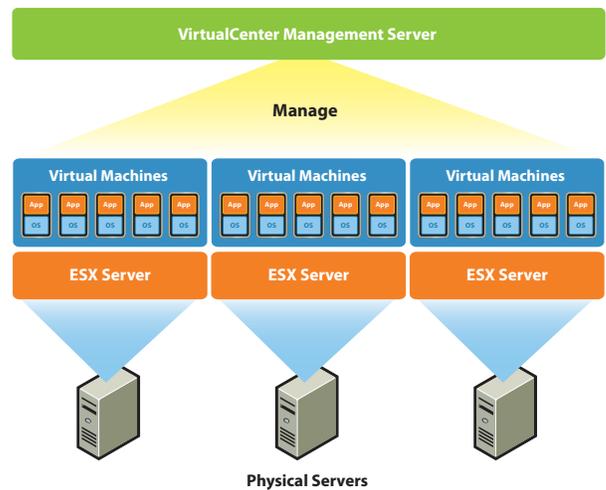
IP Networks

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

Management Server

The VirtualCenter Management Server provides a convenient single point of control to the data center. It runs on Windows 2003 Server to provide many essential data center 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 data center. As shown in Figure 1-3, VirtualCenter Management Server accomplishes this by managing the assignment of virtual machines to the computing servers. VirtualCenter Management Server also manages the assignment of resources to the virtual machines within a given computing server based on the policies set by the system administrator.

Figure 1-3: VirtualCenter Management Server centrally manages the assignment of virtual machines to physical servers



Computing servers will continue to function even in unlikely event that VirtualCenter Management Server became unreachable (e.g., the network is severed). Computing servers can be managed separately and will continue to run their assigned virtual machines based on the resource assignments that were last set. Once the VirtualCenter Management Server becomes available, it can manage the data center as a whole again.

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

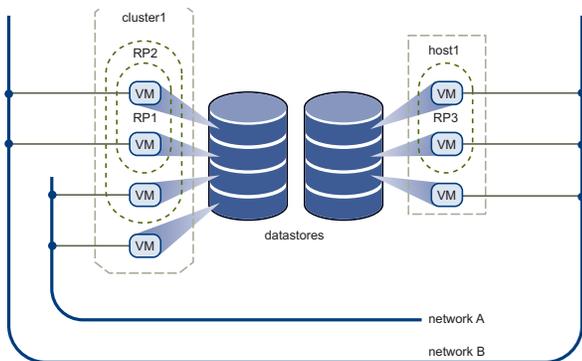
Desktop Clients

VMware Infrastructure provides a selection of interfaces for data center 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 Data Center 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 1-4: Virtual Data Center Architecture



As shown in Figure 1-4, VMware Infrastructure presents a simple set of virtual elements used to build a virtual data center:

- 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 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 data center. 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 data center.

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 a physical appliance consumes electricity. While in powered-off, suspended, or idle state, it consumes no resources. Once powered-on, it consumes resources dynamically, using more 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, no physical constraints to worry about. Once 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, please see the Resource Management Guide.

The following sections examine in detail the virtual elements of the data center.

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 4 GHz each and 32 GB of system memory, then the Host will have 32 GHz of computing power and 32 GBs of memory available for running virtual machines that are assigned to it.

A Cluster 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 8 servers, each server has 4 dual-core CPUs running at 4 GHz each and 32 GB of memory. The Cluster will then have 256 GHz of computing power and 256 GB of memory available for the running virtual machines assigned to it.

The virtual resource owners do not need to be concerned with the physical composition (number servers, quantity and type of CPUs—Multi-core, Hyperthreading, etc) of the underlying Cluster to provision resources. They simply set up the resource provisioning policies based on the aggregate available resource. VMware Infrastructure will automatically assign the appropriate resources dynamically to the virtual machines within the boundaries of those policies.

Figure 1-5: Hosts, Clusters and Resource Pools

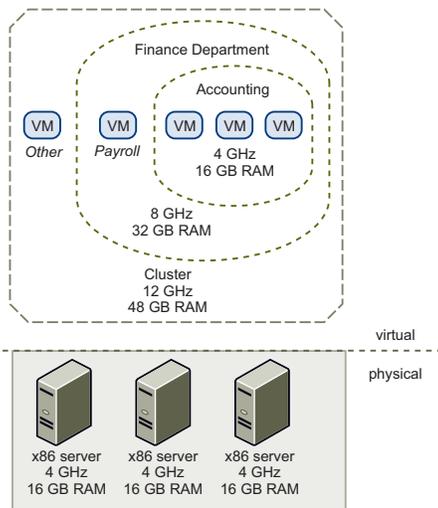


Figure 1-5 illustrates the concept of Resource Pools. Three x86 servers with 4 GHz computing power and 16 GB of memory each are aggregated to form a Cluster of 12 GHz computing power and 48 GHz of memory. A Resource Pool (“Finance Department”) reserves 8 GHz computing power and 32 GBs of from the Cluster, leaving 4 GHz computing power and 16 GBs of memory for the virtual machine “Others.” From the “Finance Department” Resource Pool, a smaller Resource Pool (“Accounting”) reserves 4 GHz computing power and 16 GBs for the virtual machines from the accounting department. That leaves 4 GHz and 16 GBs of memory for the virtual machine called “Payroll”. **Resources reserved can be dynamically changed.** Imagine that at year end, Accounting’s workload increases, the user wants to increase the Resource Pool “Accounting” from reserving 4 GHz of computing power to 6 GHz. We can simply **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 dynamically respond to the demand. For example, if the 4 GHz of computing resources reserved for the Accounting department are not being used, virtual machine “Payroll” can make use of the remaining processing capacity during its peak time. When Accounting again requires the processing capacity, “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 use dedicated infrastructure resources while still benefiting from the efficiency of resource pooling.

Resources Pools provide a flexible and dynamic way to divide and organize computing and memory resources from a Host or a Cluster. Any Resource Pools can be partitioned into smaller Resource Pools at a fine-grain level to further divide and assign resources to different groups or for different purposes.

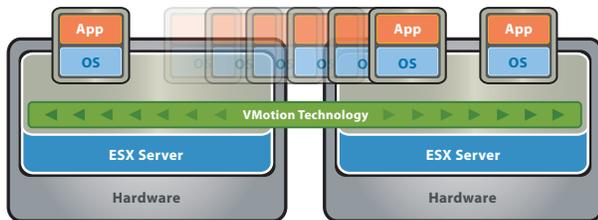
VMware VMotion, VMware DRS and VMware HA

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

VMware VMotion

As mentioned earlier, virtual machines run on and consume resources from individual physical x86 servers through VMware ESX Server. VMotion enables the migration of running virtual machines from one physical server to another without service interruption as shown in Figure 1-6. 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. Hence, with VMotion, resources can be dynamically reallocated to virtual machines across physical servers.

Figure 1-6: VMware VMotion

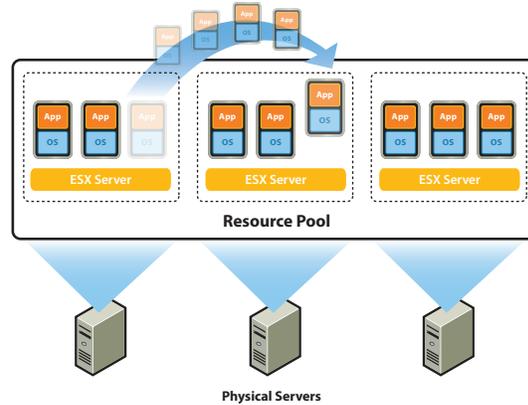


VMware DRS

Taking the VMotion capability one step further by adding an intelligent scheduler, **VMware DRS enables the system administrator to set resource assignment policies that reflect business needs and let VMware DRS do the calculation and automatically handle the detail physical resource assignments.** VMware DRS dynamically monitors the workload of the running virtual machines and the resource utilization of the physical servers within a Cluster. It checks those results against the resource assignment policies, if there is a potential for violation or improvement, it utilizes VMotion and dynamically reassigns virtual machines to different physical servers, as shown in Figure 1-7, to ensure that the policies are complied and resource allocation is optimal.

If a new physical server is made available, VMware DRS automatically redistributes the virtual machines to take advantage of it. Conversely, if a physical server needs to be taken down for any reason, VMware DRS redistributes its virtual machines to other servers automatically. For more information, see the VMware DRS white paper.

Figure 1-7: VMware DRS



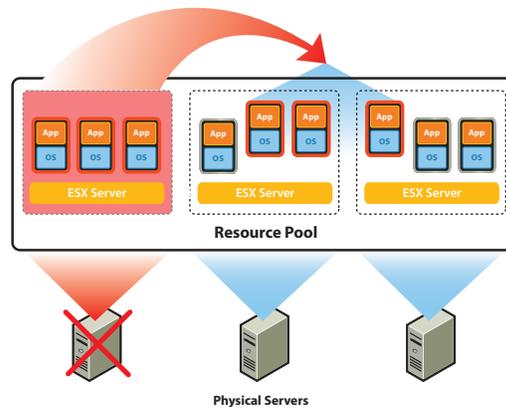
VMware HA

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 should the hosting server fails. All applications within the virtual machines will benefit from high availability, not just one (via application clustering).

VMware HA works by placing an agent on each physical server to maintain a "heartbeat" with the other servers in the Cluster. As shown in Figure 1-8, loss of a "heartbeat" from one server automatically initiates the restarting all affected virtual machines on other servers.

Setting up VMware HA can be done simply by designating the priority order of virtual machines to be restarted in the Cluster. This is very simple when compared to the set up and configuration effort required for application clustering. Furthermore, even though VMware HA requires certain amount of non-reserved resources to be maintained at all times to ensure that the remaining live servers can handle the total workload, it does not require doubling the amount of resources like application clustering. For more information, see the VMware HA white paper.

Figure 1-8. VMware HA



Networking 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 1-9: Networking Architecture

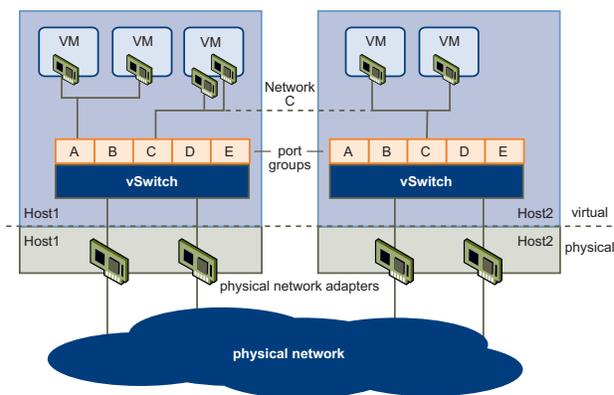


Figure 1-9 shows the relationship between the networks inside and outside the virtual environment. The virtual environment provides similar networking elements as the physical world. They are virtual network interface cards (vNIC), virtual switches (vSwitch) and Port Groups.

Like a physical machine, each virtual machine has its own vNICs. The operating system and applications talk to the vNICs through a standard networking device driver or a VMware optimized networking device driver just as though the vNIC is a physical NIC. To the outside world also, each vNIC appears just like a physical NIC – it has its own MAC address, one or more IP addresses and it responds to the standard Ethernet protocol exactly as a physical NIC would.

A vSwitch works like a Layer 2 physical switch. Each physical server has its own vSwitches. On one side of the vSwitch are Port Groups which connect to virtual machines. On the other side are uplink connections to physical Ethernet adapters on the server where the vSwitch resides. Virtual machines connect to the outside world through the physical Ethernet adapters that are connected to the vSwitch uplinks.

A virtual switch can connect its uplinks to more than one physical Ethernet adapter to enable **NIC teaming—two or more physical adapters used to share the traffic load**

or provide failover in the event of an adapter hardware failure. The failover is transparent to all virtual machines without the need to configure any of them specifically for NIC teaming. 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 as shown in Figure 1-9. Network C is the same on both Host1 and Host2.

Even if all other conditions are met, a virtual machine can VMotion from one physical server to another only if both servers have the same vSwitch (with the same Port Groups). Network connection is maintained after VMotion because the virtual machine is automatically connected to the same Port Group on the same vSwitch on new hosting server.

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 can be set for a Port Group to **isolate compromised or malicious virtual machines** and prevent them from potentially doing harm to other machines in the network.
- VLAN support can be configured for a Port Group to allow segmentation of the network
- Specific NIC teaming policies can be set for an individual Port Group (Network) to share traffic load or provide failover in case of hardware failure
- Traffic Shaping policies can be set to improve traffic management

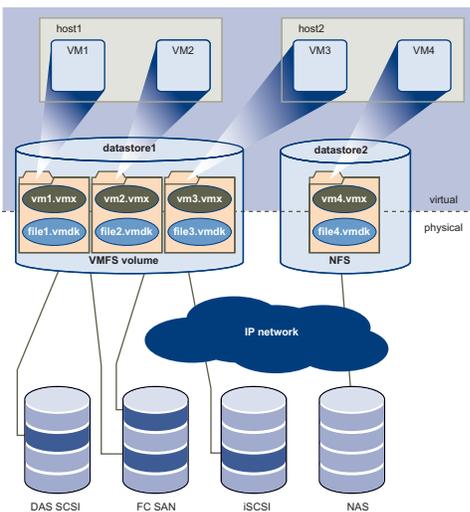
For more information on Port Group configuration, please see the Server Configuration Guide.

Storage Architecture

The VMware Infrastructure enables enterprise-class storage performance, functionality and availability without adding complexity to the user applications and guest operating systems.

The VMware Infrastructure Storage Architecture consists of layers of abstraction that hide and manage the complexity and differences between physical storage subsystems and present simple standard storage elements to the virtual environment (see Figure 1-10). To the applications and guest operating systems inside each virtual machine, storage is presented simply as SCSI disks connected to a virtual Bus Logic or LSI SCSI Host Bus Adapter.

Figure 1-10: Storage Architecture



The virtual SCSI disks inside the virtual machines are provisioned from Datastore elements in the data center. A Datastore is like a storage appliance that serves up storage space for virtual disks inside the virtual machines as well as storing the virtual machines themselves. As shown in Figure 1-10, a virtual machine is stored as a set of files in its own directory in the Datastore. A virtual disk inside each virtual machine is located on or more files inside the directory. As a result, a virtual disk can be easily manipulated (copied, moved, backed-up, and so on) just like a file. **Virtual disks can be “hot-added” to a virtual machine without powering it down.** In which case, a new virtual disk file is created or an existing virtual disk file is associated with the virtual machine.

The Datastore provides a simple model to allocate storage space for the individual virtual machines without exposing them to the complexity of the variety of physical storage technologies available, such as Fiber Channel SAN, iSCSI SAN, Direct Attached Storage, and NAS.

A Datastore is physically a VMFS file system volume or a directory on a NAS device. Each Datastore can span multiple physical storage subsystems. As shown in Figure 1-10, a single VMFS volume can contain one or more LUNs from a direct attached SCSI disk array on a physical server, a Fiber 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. They can be added to extend a previously created Datastore without powering down physical servers or storage subsystems. Conversely, if any of the LUNs within a Datastore fails or becomes unavailable, only those virtual machines that reside in that LUN are affected. All other virtual machines residing in other LUNs continue to function as normal.

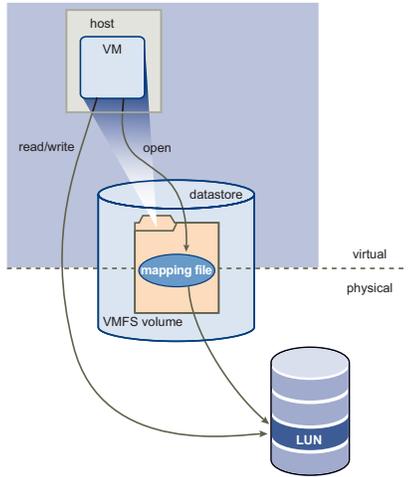
VMFS is a clustered file system that leverages shared storage to allow multiple physical servers to read and write to the same storage simultaneously. VMFS provides on-disk distributed locking to ensure that the same virtual machine is not powered on by multiple servers at the same time. If a physical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

VMFS also features enterprise class crash consistency and recovery mechanisms, such as distributed journaling, crash consistent virtual machine IO path, and machine state snapshots. These mechanisms can aid quick root-cause analysis and recovery from virtual machine, physical servers, 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 (Fiber 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) spans physical servers: 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 1-11: Raw Device Mapping



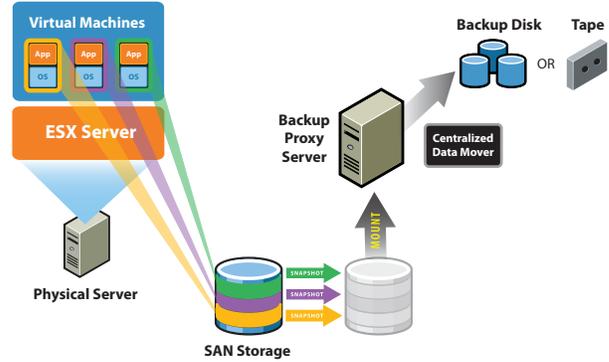
An RDM can be thought of as a symbolic link from a VMFS volume to a raw LUN (see Figure 1-11). The mapping makes LUNs appear as files in a VMFS volume. The mapping file—not the raw LUN—is referenced in the virtual machine configuration.

When a LUN is opened for access, VMFS resolves the RDM file to the correct physical device and performs appropriate access checks and locking. 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 an easy to use, centralized facility for agent-less backup of virtual machines.** As shown in Figure 1-12, 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. Once 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 at any time without worrying about backup windows or need to quiesce applications, Consolidated Backup provides a simple, less intrusive and low overhead backup solution for the virtual environment.

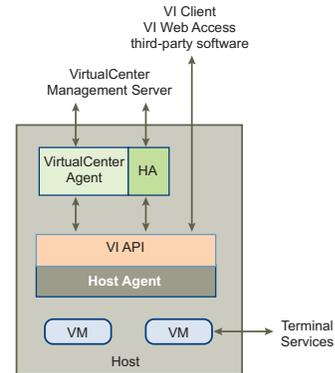
Figure 1-12: How Consolidated Backup Works



ESX Server External Interfacing Components

ESX Server is a virtualization layer that abstracts the processor, memory, storage, and networking resources of a physical server into multiple virtual machines that run side-by-side in a secured and isolated manner. There is one copy of the ESX Server that runs on each physical x86 machine. In this section, only the components of ESX Server that interfaces with the outside world are described, see Figure 1-13. For more detail information on ESX Server, please consult numerous other ESX Server technical documents.

Figure 1-13: VMware ESX Server External Interfacing Components



ESX Server communicates with other resources through its Host Agent using the VMware Infrastructure API (VI API). Using the VMware Infrastructure SDK, both VMware and third-party software can collect information from and manage ESX Server through the VI API.

As discussed earlier, the VMware HA agent provides a “heart-beat” such that when it is absent, it signals other physical servers in the Cluster to restart the virtual machines that have gone down with the hosting server. The VMware HA agent resides on each physical server and communicates with the Host Agent to constantly check on the health of the server.

The VirtualCenter Agent (VC Agent) is sent by the VirtualCenter Management Server when a machine running ESX Server is first added to be managed. VC Agent resides on each machine and acts as the broker between VirtualCenter Management Server and the ESX Server that it manages.

The VC Agent communicates with the Host Agent through the VI API as shown in Figure 1-13. Its main task is to offload some of the VirtualCenter Management Server functions specific to one physical server. This enables VirtualCenter Management Server to scale up and manage a much larger number of physical servers without demanding a proportional amount of computing power and network bandwidth. Some of the key tasks performed by VC Agent are:

- Relaying and enforcing resource allocation decisions made in VirtualCenter Management Server (including those sent by the VMware DRS engine)
- Passing virtual machines provisioning and configuration change commands to the Host Agent
- Passing Host configuration change commands to the Host Agent
- Monitoring and collecting performance statistics, alarms, and events by communicating with the Host Agent via the VI API
- Passing information back to the VirtualCenter Management Server once a certain quantity of information is collected.

VirtualCenter Management Server Architecture

VirtualCenter Management Server provides a convenient centralized management cockpit for the data center. It aggregates physical resources from multiple ESX Server installations and presents a central collection of simple and flexible resources for the system administrator when provisioning virtual machines in the virtual environment.

Figure 1-14: VirtualCenter Management Server Components

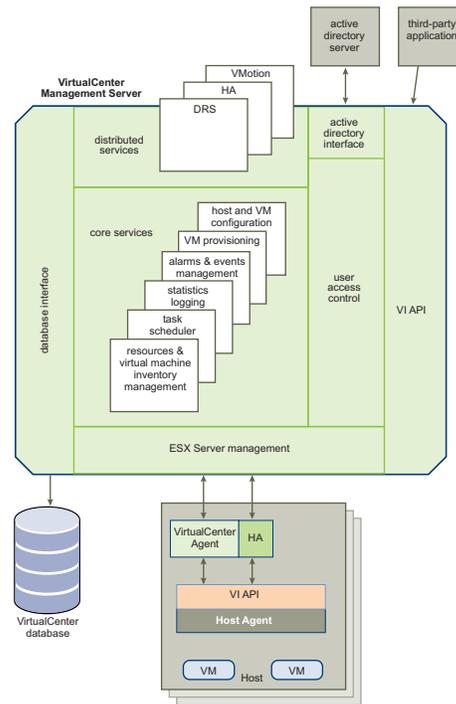


Figure 1-14 shows the key functional components of the VirtualCenter Management Server. They include User Access Control, Core Services, Distributed Services, and interfaces to external resources.

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

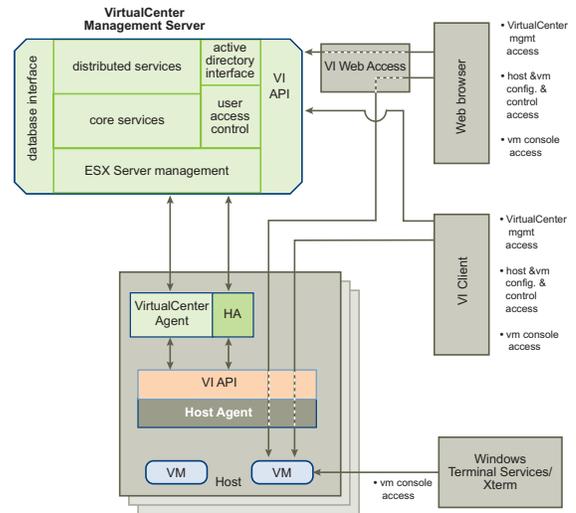
VirtualCenter Management Server has four key interfaces:

- ESX Server Management – Interfaces with the VC Agent to manage each physical server in the data center
- 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, 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

Accessing the Virtual Data Center

Users can manage the VMware Infrastructure data center 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).

Figure 1-15: VMware Infrastructure Access and Control



Both the VirtualCenter Management Server and VMware ESX Server can be managed through the VI API by VMware Infrastructure tools or third-party tools.

Managing ESX Server directly should be done by physical server administrators on special circumstances only. All management functionality that can be done on the ESX Server can also be done in VirtualCenter Management Server.

Using the VI Client, users can manage both the physical and virtual resources of the data center. When accessing VirtualCenter Management Server, a session starts after a user is authenticated. After authentication, a view of the resources and virtual machines that belongs to the user is presented. For accessing a virtual machine console directly, the VI Client first obtains the virtual machine location from the VirtualCenter Management Server through the VI API. It then connects directly to the appropriate ESX Server and presents appropriate virtual machine console.

With a Web browser, users can manage the virtual resources and access the virtual machine console. Users can do so by first pointing the Web browser to the VI Web Access. The VI Web Access mediates the communication between the browser and VirtualCenter Management Server through the VI API. To access the virtual machine consoles through the Web browser, users can make use of the bookmark that is created by VirtualCenter Management Server. The bookmark first points to the VI Web

Access. VI Web Access resolves the physical location of the virtual machine then redirects the Web browser to the appropriate ESX Server where the virtual machine resides.

Users of course can access their virtual machine console exactly the same way that they access the physical machine by using standard tools such as Windows Terminal Services or Xterm.

Conclusion

Of all the virtualization solutions in the market, only VMware Infrastructure is architected from the ground up to enable customers to bring the benefits of virtualization to the entire enterprise platform, including server, storage and networking systems. VMware Infrastructure not only enables customers to build an entire virtual data center that matches the capabilities available in the physical world, but goes one step further by offering the flexibility and control that is physically cost-prohibited or impossible.

With a whole suite of complementary virtualization and management services such as virtual VMware VMotion, VMware DRS, VMware HA and VMware Consolidated Backup, VMware Infrastructure is the only offering that brings about a complete solution rather than piece-meal approach for customers to build a data center in a virtual environment.

In its third generation since being the first-ever to virtualize the x86 platform, VMware Infrastructure has developed a large third-party solutions ecosystem and enjoys industry-wide hardware and software certification support. Since inception in 1998, it has been production tested and deployed in some of most demanding and largest IT production environments. VMware Infrastructure is the only virtualization solution ready for enterprise-wide deployment for corporate IT.



VMware, Inc. 3145 Porter Drive Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com
© 2006 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.

