

Chapter 1

Introducing VMware Infrastructure 3

VMware Infrastructure 3 (VI3) is the most widely used virtualization platform available today. The lineup of products included in VI3 makes it the most robust, scalable, and reliable server virtualization product on the market. With dynamic resource controls, high availability, distributed resource management, and backup tools included as part of the suite, IT administrators have all the tools they need to run an enterprise environment consisting of anywhere from ten to thousands of servers.

In this chapter you will learn to:

- Identify the role of each product in the VI3 suite
- Discriminate between the different products in the V13 suite
- Understand how V13 differs from other virtualization products

Exploring VMware Infrastructure 3

The VI3 product suite includes several products that make up the full feature set of enterprise virtualization. The products in the VI3 suite include:

- ◆ VMware ESX Server
- ◆ VMware Virtual SMP
- ◆ VMware VirtualCenter
- ◆ Virtual Infrastructure Client
- ◆ VMware VMotion
- ◆ VMware Distributed Resource Scheduler (DRS)
- ◆ VMware High Availability (HA)
- ◆ VMware Consolidated Backup (VCB)

Rather than wait to introduce the individual products in their own chapters, I'll introduce each product so I can refer to the products and explain how they affect each piece of the design, installation, and configuration of your virtual infrastructure. Once you understand the basic functions and features of each product in the suite, you'll have a better grasp of how that product fits into

the big picture of virtualization, and you'll more clearly understand how each of the products fits into the design.


VMware ESX Server

VMware ESX Server 3.5 and ESXi are the core of the VI3 product suite. They function as the hypervisor, or virtualization layer, that serves as the foundation for the whole VI3 package. Unlike some virtualization products that require a host operating system, ESX Server is a bare metal installation, which means no host operating system (Windows or Linux) is required. ESX Server is a leaner installation than products requiring a host operating system, which allows more of its hardware resources to be utilized by virtual machines rather than by processes required to run the host. The installation process for ESX Server installs two components that interact with each other to provide a dynamic and robust virtualization environment: the Service Console and the VMkernel.

The Service Console, for all intents and purposes, is the operating system used to manage ESX Server and the virtual machines that run on the server. The console includes services found in other operating systems, such as a firewall, Simple Network Management Protocol (SNMP) agents, and a web server. At the same time, the Service Console lacks many of the features and benefits that other operating systems offer. This deficiency, however, serves as a true advantage in making the Service Console a lean, mean, virtualization machine.

The other installed component is the VMkernel. While the Service Console gives you access to the VMkernel, it is the VMkernel that is the real foundation of the virtualization process. The VMkernel manages the virtual machines' access to the underlying physical hardware by providing CPU scheduling, memory management, and virtual switch data processing. Figure 1.1 shows the structure of ESX Server.

FIGURE 1.1
Installing ESX Server installs two interoperable components: 1) the Linux-derived Service Console, and 2) the virtual machine-managing VMkernel.



ESXi is the next generation of the VMware virtualization foundation in that it lightens the load to a 32MB footprint as installation of a hypervisor only. ESXi is *only* a hypervisor and does not have any reliance on an accompanying Service Console.

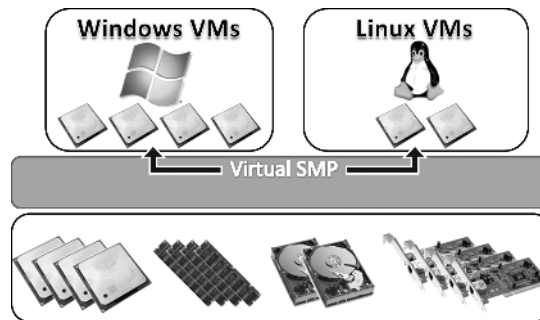
I'll go into much more detail about the installation of ESX Server in Chapter 2. The installation procedure of ESX Server also allows for the configuration of VMware File System (VMFS) datastores. Chapter 4 will provide an in-depth look at the various storage technologies. Once your core product, ESX Server, is installed, you can build off this product with the rest of the product suite.

VMware Virtual SMP

The VMware Virtual Symmetric Multi-Processing (SMP) product allows virtual infrastructure administrators to construct virtual machines with multiple virtual processors. VMware Virtual

SMP is *not* the licensing product that allows ESX Server to be installed on servers with multiple processors; it is the configuration of multiple processors *inside* a virtual machine. Figure 1.2 identifies the differences between multiple processors in the ESX Server host system and multiple virtual processors.

FIGURE 1.2
VMware Virtual SMP
allows virtual machines
to be created with two
or four processors.



In Chapter 6 we'll look at how, why, and when to build virtual machines with multiple virtual processors.

ESX Server includes a host of new features and support for additional hardware and storage devices. At the urging of the virtualization community, ESX Server now boasts support for Internet Small Computer Systems Interface (iSCSI) storage and network attached storage (NAS) in addition to Fibre Channel storage technologies. Chapter 4 describes the selection, configuration, and management of all three storage technologies supported by ESX Server.

VMware VirtualCenter

Stop for a moment and think about your current Windows network. Does it include Active Directory? There is a good chance it does. Now imagine your Windows network without Active Directory, without the ease of a centralized management database, without the single sign-on capabilities, and without the simplicity of groups. That is what managing ESX Server computers would be like without using VMware VirtualCenter 2.0. Now calm yourself down, take a deep breath, and know that VirtualCenter, like Active Directory, is meant to provide a centralized management utility for all ESX Server hosts and their respective virtual machines. VirtualCenter is a Windows-based, database-driven application that allows IT administrators to deploy, manage, monitor, automate, and secure a virtual infrastructure in an almost effortless fashion. The back-end database (SQL or Oracle) used by VirtualCenter stores all the data about the hosts and virtual machines. In addition to its configuration and management capabilities, VirtualCenter provides the tools for the more advanced features of VMware VMotion, VMware DRS, and VMware HA. Figure 1.3 details the VirtualCenter features provided for the ESX Server hosts it manages.

In Chapter 5, you'll learn the details of the VirtualCenter implementation, configuration, and management, as well as look at ways to ensure its availability.

Virtual Infrastructure Client

The Virtual Infrastructure (VI) Client is a Windows-based application that allows you to connect to and manage an ESX Server or a VirtualCenter Server. You can install the VI Client by browsing to the URL of an ESX Server or VirtualCenter and selecting the appropriate installation link. The VI Client is a graphical user interface (GUI) used for all the day-to-day management tasks and for the advanced configuration of a virtual infrastructure. Using the client to connect

directly to an ESX Server requires that you use a user account residing in the Service Console (a Linux account), while using the client to connect to a VirtualCenter Server requires you to use a Windows account. Figure 1.4 shows the account authentication for each connection type.

FIGURE 1.3

VirtualCenter 2.0 is a Windows-based application used for the centralization of authentication, accounting, and management of ESX Server hosts and their corresponding virtual machines.





FIGURE 1.4

The Virtual Infrastructure Client can be used to manage an individual ESX Server by authenticating with a Linux account that resides in the Service Console; however, it can also be used to manage an entire enterprise by authenticating to a VirtualCenter Server using a Windows account.



Almost all the management tasks available when you're connected directly to an ESX Server are available when you're connected to a VirtualCenter Server, but the opposite is not true. The management capabilities available through VirtualCenter Server are more significant and outnumber the capabilities of connecting directly to an ESX Server.

VMware VMotion and Storage VMotion

If you have read anything about VMware, you have most likely read about the extremely unique and innovative feature called VMotion. VMotion is a feature of ESX Server and VirtualCenter that allows a running virtual machine to be moved from one ESX Server host to another without having to power off the virtual machine. Figure 1.5 illustrates the VMotion feature of VirtualCenter.

FIGURE 1.5

The VMotion feature of VirtualCenter allows a running virtual machine to be transitioned from one ESX Server host to another.



VMotion satisfies an organization's need for maintaining service-level agreements (SLAs) that guarantee server availability. Administrators can easily instantiate a VMotion to remove all virtual machines from an ESX Server host that is to undergo scheduled maintenance. Once the maintenance is complete and the server is brought back online, VMotion can once again be utilized to return the virtual machines to the original server.

Even in a normal day-to-day operation, VMotion can be used when multiple virtual machines on the same host are in contention for the same resource (which ultimately is causing poor performance across all the virtual machines). VMotion can solve the problem by allowing an administrator to migrate any of the running virtual machines that are facing contention to another ESX host with greater availability for the resource in demand. For example, when two virtual machines are in contention with each other for CPU power, an administrator can eliminate the contention by performing a VMotion of one of the virtual machines to an ESX host that has more available CPU. More details on the VMware VMotion feature and its requirements will be provided in Chapter 9.

Storage VMotion builds on the idea and principle of VMotion in that downtime can be reduced when running virtual machines can be migrated to different physical environments. Storage VMotion, however, allows running virtual machines to be moved between datastores. This feature ensures that outgrowing datastores or moving to a new SAN does not force an outage for the effected virtual machines.


VMware Distributed Resource Scheduler (DRS)

Now that I've piqued your interest with the introduction of VMotion, let me introduce VMware Distributed Resource Scheduler (DRS). If you think that VMotion sounds exciting, your anticipation will only grow after learning about DRS. DRS, simply put, is a feature that aims to provide automatic distribution of resource utilization across multiple ESX hosts that are configured in a cluster. An ESX Server cluster is a new feature in VMware Infrastructure 3. The use of the term *cluster* often draws IT professionals into thoughts of Microsoft Windows Server clusters. However, ESX Server clusters are not the same. The underlying concept of aggregating physical hardware to serve a common goal is the same, but the technology, configuration, and feature sets are very different between ESX Server clusters and Windows Server clusters.

An ESX Server cluster is an implicit aggregation of the CPU power and memory of all hosts involved in the cluster. Once two or more hosts have been assigned to a cluster, they work in unison to provide CPU and memory to the virtual machines assigned to the cluster. The goal of DRS is to provide virtual machines with the required hardware resources while minimizing the amount of contention for those resources in an effort to maintain good performance levels.

DRS has the ability to move running virtual machines from one ESX Server host to another when resources from another host can enhance a virtual machine's performance. Does that sound familiar? It should, because the behind-the-scenes technology for DRS is VMware VMotion. DRS can be configured to automate the placement of each virtual machine as it is powered on as well as to manage the virtual machine's location once it is running. For example, let's say three servers have been configured in an ESX Server cluster with DRS enabled. When one of those servers begins to experience a high contention for CPU utilization, DRS will use an internal algorithm to determine which virtual machine(s) will experience the greatest performance boost by being moved to another server with less CPU contention. Figure 1.6 outlines the automated feature of DRS.

FIGURE 1.6
VMware Distributed Resource Scheduler (DRS) aims to maintain balance and fairness of resource utilization for virtual machines running within an ESX Server cluster.




Chapter 9 dives deeper into the configuration and management of DRS on an ESX Server cluster.

VMware High Availability (HA)

With the introduction of the ESX Server cluster, VMware has also introduced a new feature called VMware High Availability (HA). Once again, by nature of the naming conventions (clusters, high availability), many traditional Windows administrators will have preconceived notions about this feature. Those notions, however, are premature in that VMware HA does not function like a high-availability configuration in Windows. The VMware HA feature provides an automated process for restarting virtual machines that were running on an ESX Server at a time of complete server failure. Figure 1.7 depicts the virtual machine migration that occurs when an ESX Server that is part of an HA-enabled cluster experiences failure.

FIGURE 1.7
The VMware High Availability (HA) feature will power on any virtual machines that were previously running on an ESX Server that has experienced server failure.



The VMware HA feature, unlike DRS, does not use the VMotion technology as a means of migrating servers to another host. In a VMware HA failover situation, there is no anticipation of failure; it is not a planned outage and therefore there is no time to perform a VMotion. VMware HA does not provide failover in the event of a single virtual machine failure. It provides an automated restart of virtual machines during an ESX Server failure.

Chapter 10 will explore the configuration and working details of VMware High Availability.


VMware Consolidated Backup (VCB)

One of the most critical aspects to any network, not just a virtualized infrastructure, is a solid backup strategy as defined by a company's disaster recovery and business continuity plan. VMware Consolidated Backup (VCB) is a Windows application that provides a LAN-free Fibre Channel or iSCSI-based backup solution that offloads the backup processing to a dedicated physical server. VCB takes advantage of the snapshot functionality in ESX Server to mount the snapshots into the file system of the dedicated VCB server. Once the respective virtual machine files are mounted, entire virtual machines or individual files can be backed up using third-party

backup tools. VCB scripts integrate with several major third-party backup solutions to provide a means of automating the backup process. Figure 1.8 details a VCB implementation.

FIGURE 1.8

VMware Consolidated Backup (VCB) is a LAN-free online backup solution that uses a Fibre Channel or iSCSI connection to expedite and simplify the backup process.



In Chapter 10 you'll learn how to use VCB to provide a solid backup and restore practice for your virtual infrastructure.



Real World Scenario

VIRTUAL INFRASTRUCTURE 3 VS. VMWARE SERVER (AND THE OTHERS)

The Virtual Infrastructure 3 (VI3) product holds a significant advantage over most other virtualization products because virtualization on VI3 does not require a host operating system. Products like VMware Server and Microsoft Virtual Server 2005 both require an underlying operating system to host the hypervisor.

The lack of the host operating system in VI3 offers additional stability and security. Without an underlying operating system like Windows, there is less concern for viruses, spyware, and unnecessary exposure to vulnerabilities.


With products like VMware Server (which require a host operating system), limitations from the host operating systems spill into the virtualization deployment. For example, installing VMware Server on Windows Server 2003 Web edition would establish two processors and 2GB of RAM limitations on VMware Server, despite its ability to use up to 16 processors and 64GB of RAM. At the same time, however, there's the advantage that hosted products have over the bare metal install of ESX Server. The existence of the host operating system greatly extends the level of hardware support on which the hypervisor will run. If the host operating system offers support, then the virtual machine will too. A great example of this hardware support is to look at the use of USB. ESX Server does not support USB, while VMware Server (and Workstation) includes support. Since the underlying host understands the USB technology, the virtual machines will also offer support.

In all, each of the virtualization products has its place in a network infrastructure. The Virtual Infrastructure 3 product is more suited to the mission-critical enterprise data center virtualization scenario, while the VMware Server product is best for noncritical test or branch office scenarios. And of course you cannot forget the best part of VMware Server: it's *free!*

The Bottom Line

Identify the role of each product in the VI3 suite. Now that you've been introduced to the products included in the VMware Infrastructure 3 suite, we can begin discussing the technical details, best practices, and how-tos that will make your life as a virtual infrastructure administrator a whole lot easier. This chapter has shown that each of the products in the VI3 suite plays an integral part in the overall process of creating, managing, and maintaining a virtual enterprise. Figure 1.9 highlights the VI3 product suite and how it integrates and interoperates to provide a robust set of tools upon which a scalable, reliable, and redundant virtual enterprise can be built.

FIGURE 1.9
The products in the VMware Infrastructure suite work together to provide a scalable, robust, and reliable framework for creating, managing, and monitoring a virtual enterprise.



The next chapter will begin a start-to-finish look at designing, implementing, managing, monitoring, and troubleshooting a virtual enterprise built on VI3. I'll dive into much greater detail on each of the products I introduced in this chapter. This introduction should provide you with a solid foundation so we can discuss the different products beginning with the next chapter. You can use this introduction as a reference throughout the remaining chapters if you want to refresh your base knowledge for each of the products in the suite.

Master It You want to centralize the management of ESX Server hosts and all virtual machines.

Master It You want to minimize the occurrence of system downtime during periods of planned maintenance.

Master It You want to provide an automated method of maintaining fairness and balance of resource utilization.

Master It You want to provide an automated restart of virtual machines when an ESX Server fails.

Master It You want to institute a method of providing disaster recovery and business continuity in the event of virtual machine failure.