

Chapter 1

Introducing VMware vSphere 4

VMware vSphere 4 builds upon previous generations of VMware virtualization products, becoming an even more robust, scalable, and reliable server virtualization product. With dynamic resource controls, high availability, unprecedented fault tolerance features, 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 ranging from a few servers up to thousands of servers.

In this chapter, you will learn to:

- ◆ Identify the role of each product in the vSphere product suite
- ◆ Recognize the interaction and dependencies between the products in the vSphere suite
- ◆ Understand how vSphere differs from other virtualization products

Exploring VMware vSphere 4

The VMware vSphere product suite includes a number of products and features that together provide a full array of enterprise virtualization functionality. These products and features in the vSphere product suite include the following:

- ◆ VMware ESX and ESXi
- ◆ VMware Virtual Symmetric Multi-Processing
- ◆ VMware vCenter Server
- ◆ VMware vCenter Update Manager
- ◆ VMware vSphere Client
- ◆ VMware VMotion and Storage VMotion
- ◆ VMware Distributed Resource Scheduler
- ◆ VMware High Availability
- ◆ VMware Fault Tolerance
- ◆ VMware Consolidated Backup
- ◆ VMware vShield Zones
- ◆ VMware vCenter Orchestrator

Rather than waiting to introduce these products and features in their own chapters, I'll introduce each product or feature in the following sections. This will allow me to explain how each

product or feature affects the design, installation, and configuration of your virtual infrastructure. After I cover the features and products in the vSphere suite, you'll have a better grasp of how each of them fits into the design and the big picture of virtualization.

Certain products outside the vSphere product suite extend the vSphere product line with new functionality. Examples of these additional products include VMware vCenter Lifecycle Manager, VMware vCenter Lab Manager, VMware vCenter Stage Manager, and VMware vCenter Site Recovery Manager. Because of the size and scope of these products and because they are developed and released on a schedule separate from VMware vSphere, they will not be covered in this book.

VMware ESX and ESXi

The core of the vSphere product suite is the hypervisor, which is the virtualization layer that serves as the foundation for the rest of the product line. In vSphere, the hypervisor comes in two different forms: VMware ESX and VMware ESXi. Both of these products share the same core virtualization engine, both can support the same set of virtualization features, and both are considered bare-metal installations. VMware ESX and ESXi differ in how they are packaged.

TYPE 1 AND TYPE 2 HYPERVISORS

Hypervisors are generally grouped into two classes: type 1 hypervisors and type 2 hypervisors. Type 1 hypervisors run directly on the system hardware and thus are often referred to as *bare-metal* hypervisors. Type 2 hypervisors require a host operating system, and the host operating system provides I/O device support and memory management. VMware ESX and ESXi are both type 1 bare-metal hypervisors. Other type 1 bare-metal hypervisors include Microsoft Hyper-V and products based on the open source Xen hypervisor like Citrix XenServer and Oracle VM.


VMware ESX consists of 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 interact with VMware ESX and the virtual machines that run on the server. The Linux-derived Service Console includes services found in traditional 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 traditional operating systems offer. This is not a deficiency, though. In this particular case, the Service Console has been intentionally stripped down to include only those services necessary to support virtualization, making the Service Console a lean, mean virtualization machine.

The second installed component is the VMkernel. While the Service Console gives you access to the VMkernel, the VMkernel 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 VMware ESX.

VMware ESXi, on the other hand, is the next generation of the VMware virtualization foundation. Unlike VMware ESX, ESXi installs and runs without the Service Console. This gives ESXi an ultralight footprint of only 32MB. ESXi shares the same underlying VMkernel as VMware ESX and supports the same set of virtualization features that will be described shortly, but it does not rely upon the Service Console.

FIGURE 1.1
Installing VMware ESX
installs two interoper-
able components: the
Linux-derived Service
Console and the vir-
tual machine-managing
VMkernel.



Compared to previous versions of ESX/ESXi, VMware has expanded the limits of what the hypervisor is capable of supporting. Table 1.1 shows the configuration maximums for this version of ESX/ESXi as compared to the previous release.

TABLE 1.1: VMware ESX/ESXi 4.0 Maximums

COMPONENT	VMWARE ESX 4 MAXIMUM	VMWARE ESX 3.5 MAXIMUM
Number of virtual CPUs per host	256	128
Number of cores per host	64	32
Number of logical CPUs (hyperthreading enabled)	64	32
Number of virtual CPUs per core	20	8 (increased to 20 in Update 3)
Amount of RAM per host	1TB	128GB (increased to 256GB in Update 3)

Where appropriate, each chapter will include additional values for VMware ESX/ESXi 4 maximums for NICs, storage, virtual machines, and so forth.


Because VMware ESX and ESXi form the foundation of the vSphere product suite, I'll touch on various aspects of ESX/ESXi throughout the book. I'll go into more detail about the installation of both VMware ESX and ESXi in Chapter 2, "Planning and Installing VMware ESX and VMware ESXi." In Chapter 5, "Creating and Managing Virtual Networks," I'll more closely examine the networking capabilities of ESX/ESXi. Chapter 6, "Creating and Managing Storage Devices," describes the selection, configuration, and management of the storage technologies supported by ESX/ESXi, including the configuration of VMware vStorage VMFS datastores.

VMware Virtual Symmetric Multi-Processing

The VMware Virtual Symmetric Multi-Processing (vSMP, or Virtual 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/ESXi to be installed on servers with multiple processors; it is the technology that allows the use of multiple processors *inside* a

virtual machine. Figure 1.2 identifies the differences between multiple processors in the ESX/ESXi host system and multiple virtual processors.

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



With VMware Virtual SMP, applications that require and can actually use multiple CPUs can be run in virtual machines configured with multiple virtual CPUs. This allows organizations to virtualize even more applications without negatively impacting performance or being unable to meet service-level agreements (SLAs).

In Chapter 7, “Creating and Managing Virtual Machines,” I’ll discuss how to build virtual machines with multiple virtual processors.

VMware vCenter Server

Stop for a moment to think about your current network. Does it include Active Directory? There is a good chance it does. Now imagine your 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 VMware ESX/ESXi hosts would be like without using VMware vCenter Server. Now calm yourself down, take a deep breath, and know that vCenter Server, like Active Directory, is meant to provide a centralized management utility for all ESX/ESXi hosts and their respective virtual machines. vCenter Server 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 (Microsoft SQL Server or Oracle) that vCenter Server uses stores all the data about the hosts and virtual machines.

VCENTER SERVER FOR LINUX

At the time this book was written, VMware had just released a technology preview of a Linux version of vCenter Server. A Linux version of vCenter Server would remove the requirement to have a Windows-based server present in the environment in order to support VMware vSphere, something that Linux- and UNIX-heavy organizations have long desired.

In addition to its configuration and management capabilities—which include features such as virtual machine templates, virtual machine customization, rapid provisioning and deployment of virtual machines, role-based access controls, and fine-grained resource allocation controls—vCenter Server provides the tools for the more advanced features of VMware VMotion, VMware Distributed Resource Scheduler, VMware High Availability, and VMware Fault Tolerance.

In addition to VMware VMotion, VMware Distributed Resource Scheduler, VMware High Availability, and VMware Fault Tolerance, using vCenter Server to manage ESX/ESXi hosts also enables a number of other features:

- ◆ Enhanced VMotion Compatibility (EVC), which leverages hardware functionality from Intel and AMD to enable greater CPU compatibility between servers grouped into VMware DRS clusters
- ◆ Host profiles, which allow administrators to bring greater consistency to host configurations across larger environments and to identify missing or incorrect configurations
- ◆ vNetwork Distributed Switches, which provide the foundation for cluster-wide networking settings and third-party virtual switches

vCenter Server plays a central role in any sizable VMware vSphere implementation. Because of vCenter Server's central role, I'll touch on aspects of vCenter Server's functionality throughout the book. For example, in Chapter 3, "Installing and Configuring vCenter Server," I discuss planning and installing vCenter Server, as well as look at ways to ensure its availability. Chapters 5 through 12 all cover various aspects of vCenter Server's role in managing your VMware vSphere environment. As an integral part of your VMware vSphere installation, it's quite natural that I discuss vCenter Server in such detail.

vCenter Server is available in three editions:

- ◆ vCenter Server Essentials is integrated into the vSphere Essentials edition for small office deployment.
- ◆ vCenter Server Standard provides all the functionality of vCenter Server, including provisioning, management, monitoring, and automation.
- ◆ vCenter Server Foundation is like vCenter Server Standard but is limited to managing three ESX/ESXi hosts.

You can find more information on licensing and product editions for VMware vSphere in the section "Licensing VMware vSphere."

VMware vCenter Update Manager

vCenter Update Manager is a plug-in for vCenter Server that helps users keep their ESX/ESXi hosts and select virtual machines patched with the latest updates. vCenter Update Manager provides the following functionality:

- ◆ Scans to identify systems that are not compliant with the latest updates
- ◆ User-defined rules for identifying out-of-date systems
- ◆ Automated installation of patches for ESX/ESXi hosts
- ◆ Full integration with other vSphere features like Distributed Resource Scheduler
- ◆ Support for patching Windows and Linux operating systems
- ◆ Support for patching select Windows applications inside virtual machines


Chapter 4, “Installing and Configuring vCenter Update Manager,” features more extensive coverage of vCenter Update Manager.

VMware vSphere Client

The VMware vSphere Client is a Windows-based application that allows you to manage ESX/ESXi hosts, either directly or through a vCenter Server. You can install the vSphere Client by browsing to the URL of an ESX/ESXi host or vCenter Server and selecting the appropriate installation link. The vSphere 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/ESXi host requires that you use a user account residing on that host, while using the client to connect to a vCenter Server requires that you use a Windows account. Figure 1.3 shows the account authentication for each connection type.

FIGURE 1.3

The vSphere Client manages an individual ESX/ESXi host by authenticating with an account local to that host; however, it manages an entire enterprise by authenticating to a vCenter Server using a Windows account.



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

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, also known as *live migration*, is a feature of ESX/ESXi and vCenter Server that allows a running virtual machine to be moved from one physical host to another physical host without having to power off the virtual machine. This migration between two physical hosts occurs with no downtime and with no loss of network connectivity to the virtual machine.

VMotion satisfies an organization’s need for maintaining SLAs that guarantee server availability. Administrators can easily initiate VMotion to remove all virtual machines from an ESX/ESXi host that is to undergo scheduled maintenance. After the maintenance is complete and the server is brought back online, VMotion can again be utilized to return the virtual machines to the original server.

Even in normal day-to-day operations, 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/ESXi 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/ESXi host that

has more available CPU. More details on the VMware VMotion feature and its requirements are provided in Chapter 10, “Managing Resource Allocation.”

Storage VMotion builds on the idea and principle of VMotion, further reducing planned downtime with the ability to move a virtual machine’s storage while the virtual machine is still running. Deploying VMware vSphere in your environment generally means that lots of shared storage—Fibre Channel or iSCSI SAN or NFS—is needed. What happens when you need to migrate from an older storage array to a newer storage array? What kind of downtime would be required?

Storage VMotion directly addresses this concern. Storage VMotion moves the storage for a running virtual machine between datastores. Much like VMotion, Storage VMotion works without downtime to the virtual machine. This feature ensures that outgrowing datastores or moving to a new SAN does not force an outage for the affected virtual machines and provides administrators with yet another tool to increase their flexibility in responding to changing business needs.

VMware Distributed Resource Scheduler

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/ESXi hosts that are configured in a cluster. The use of the term *cluster* often draws IT professionals into thoughts of Microsoft Windows Server clusters. However, ESX/ESXi 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 different between ESX/ESXi clusters and Windows Server clusters.

AGGREGATE CAPACITY AND SINGLE HOST CAPACITY

Although I say that a DRS cluster is an implicit aggregation of CPU and memory capacity, it’s important to keep in mind that a virtual machine is limited to using the CPU and RAM of a single physical host at any given time. If you have two ESX/ESXi servers with 32GB of RAM each in a DRS cluster, the cluster will correctly report 64GB of aggregate RAM available, but any given virtual machine will not be able to use more than approximately 32GB of RAM at a time.

An ESX/ESXi cluster is an implicit aggregation of the CPU power and memory of all hosts involved in the cluster. After 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 twofold:

- ◆ At startup, DRS attempts to place each virtual machine on the host that is best suited to run that virtual machine at that time.
- ◆ While a virtual machine is running, DRS seeks to provide that virtual machine with the required hardware resources while minimizing the amount of contention for those resources in an effort to maintain good performance levels.

The first part of DRS is often referred to as *intelligent placement*. DRS can automate the placement of each virtual machine as it is powered on within a cluster, placing it on the host in the cluster that it deems to be best suited to run that virtual machine at that moment.

DRS isn't limited to operating only at virtual machine startup, though. DRS also manages the virtual machine's location while it is running. For example, let's say three servers have been configured in an ESX/ESXi cluster with DRS enabled. When one of those servers begins to experience a high contention for CPU utilization, DRS uses 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.

DRS performs these on-the-fly adjustments without any downtime or loss of network connectivity to the virtual machines. Does that sound familiar? It should, because the behind-the-scenes technology used by DRS is VMware VMotion, which I described previously. In Chapter 10, "Managing Resource Allocation," I'll dive deeper into the configuration and management of DRS on an ESX/ESXi cluster.

FEWER BIGGER SERVERS OR MORE SMALLER SERVERS?

Remember from Table 1.1 that VMware ESX/ESXi supports servers with up to 64 CPU cores and up to 1TB of RAM. With VMware DRS, though, you can combine multiple smaller servers together for the purpose of managing aggregate capacity. This means that bigger, more powerful servers may not be better servers for virtualization projects. These larger servers are generally significantly more expensive than smaller servers, and using a greater number of smaller servers may provide greater flexibility than a smaller number of larger servers. The key thing to remember here is that a bigger server isn't necessarily a better server!


VMware High Availability

In many cases, high availability (HA)—or the lack of high availability—is the key argument used against virtualization. The most common form of this argument more or less sounds like this: "Before virtualization, the failure of a physical server affected only one application or workload. After virtualization, the failure of a physical server will affect many more applications or workloads running on that server at the same time. We can't put all our eggs in one basket!"

VMware addresses this concern with another feature present in ESX/ESXi clusters 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/ESXi host at a time of complete server failure. Figure 1.4 depicts the virtual machine migration that occurs when an ESX/ESXi host that is part of an HA-enabled cluster experiences failure.

FIGURE 1.4

The VMware 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 is intended to address unplanned downtime because of the failure of a physical ESX/ESXi host.

By default VMware HA does not provide failover in the event of a guest operating system failure, although you can configure VMware HA to monitor virtual machines and restart them automatically if they fail to respond to an internal heartbeat. For users who need even higher levels of availability, VMware Fault Tolerance (FT), which is described in the next section, can satisfy that need.


Chapter 11, “Ensuring High Availability and Business Continuity,” explores the configuration and working details of VMware High Availability and VMware Fault Tolerance.

VMware Fault Tolerance

For users who require even greater levels of high availability than VMware HA can provide, VMware vSphere introduces a new feature known as VMware Fault Tolerance (FT).

VMware HA protects against unplanned physical server failure by providing a way to automatically restart virtual machines upon physical host failure. This need to restart a virtual machine in the event of a physical host failure means that some downtime—generally less than three minutes—is incurred. VMware FT goes even further and eliminates any downtime in the event of a physical host failure. Using vLockstep technology, VMware FT maintains a mirrored secondary VM on a separate physical host that is kept in lockstep with the primary VM. Everything that occurs on the primary (protected) VM also occurs simultaneously on the secondary (mirrored) VM, so that if the physical host on which the primary VM is running fails, the secondary VM can immediately step in and take over without any loss of connectivity. VMware FT will also automatically re-create the secondary (mirrored) VM on another host if the physical host on which the secondary VM is running fails, as illustrated in Figure 1.5. This ensures protection for the primary VM at all times.

FIGURE 1.5
VMware FT provides protection against host failures with no downtime to the virtual machines.



In the event of multiple host failures—say, the hosts running both the primary and secondary VMs failed—VMware HA will reboot the primary VM on another available server, and VMware FT will automatically create a new secondary VM. Again, this ensures protection for the primary VM at all times.


VMware FT can work in conjunction with VMotion, but it cannot work with DRS, so DRS must be manually disabled on VMs that are protected with VMware FT.

Chapter 10 provides more information on how to disable DRS for specific VMs. Chapter 11 provides more information on VMware FT.

VMware Consolidated Backup

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 set of tools and interfaces that provide both LAN-free and LAN-based backup functionality to third-party backup solutions. VCB offloads the backup processing to a dedicated physical or virtual server and provides ways of integrating third-party backup solutions like Backup Exec, TSM, NetBackup, or others. VCB takes advantage of the snapshot functionality in ESX/ESXi to mount the snapshots into the file system of the dedicated VCB server. After 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 1.6 details a VCB implementation.

FIGURE 1.6
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 11, you'll learn how to use VCB to provide a solid backup and restore process for your virtualized infrastructure.

VMware vShield Zones

VMware vSphere offers some compelling virtual networking functionality, and vShield Zones builds upon vSphere's virtual networking functionality to add virtual firewall functionality. vShield Zones allows vSphere administrators to see and manage the network traffic flows occurring on the virtual network switches. You can apply network security policies across entire groups of machines, ensuring that these policies are maintained properly even though virtual machines may move from host to host using VMware VMotion and VMware DRS.

You can find more information on VMware vShield Zones in Chapter 13, "Securing VMware vSphere."

VMware vCenter Orchestrator

VMware vCenter Orchestrator is a workflow automation engine that is automatically installed with every instance of vCenter Server. Using vCenter Orchestrator, vSphere administrators can build automated workflows to automate a wide variety of tasks available within vCenter Server. The automated workflows you build using vCenter Orchestrator range from simple to complex. To get an idea of the kind of power available with vCenter Orchestrator, it might help to know that VMware vCenter Lifecycle Manager, a separate product in the vCenter virtualization management family of products, is completely built on top of vCenter Orchestrator.

Chapter 14, “Automating VMware vSphere,” provides more information on vCenter Orchestrator and other automation technologies and tools.



Real World Scenario

VMWARE VSPHERE COMPARED TO HYPER-V AND XENSERVER

It’s not really possible to compare some virtualization solutions to other virtualization solutions because they are fundamentally different in approach and purpose. Such is the case with VMware ESX/ESXi and some of the other virtualization solutions on the market.

To make accurate comparisons between vSphere and other virtualization solutions, one must include only type 1 (“bare-metal”) virtualization solutions. This would include ESX/ESXi, of course, and Microsoft Hyper-V and Citrix XenServer. It would not include products such as VMware Server or Microsoft Virtual Server, both of which are type 2 (“hosted”) virtualization products. Even within the type 1 hypervisors, there are architectural differences that make direct comparisons difficult.

For example, both Microsoft Hyper-V and Citrix XenServer route all the virtual machine I/O through the “parent partition” or “domo.” This typically provides greater hardware compatibility with a wider range of products. In the case of Hyper-V, for example, as soon as Windows Server 2008—the general-purpose operating system running in the parent partition—supports a particular type of hardware, then Hyper-V supports it also. Hyper-V “piggybacks” on Windows’ hardware drivers and the I/O stack. The same can be said for XenServer, although its “domo” runs Linux and not Windows.

VMware ESX/ESXi, on the other hand, handles I/O within the hypervisor itself. This typically provides greater throughput and lower overhead at the expense of slightly more limited hardware compatibility. In order to add more hardware support or updated drivers, the hypervisor must be updated because the I/O stack and device drivers are in the hypervisor.

This architectural difference is fundamental. Nowhere is this architectural difference more greatly demonstrated than in ESXi, which has a very small footprint yet provides a full-featured virtualization solution. Both Citrix XenServer and Microsoft Hyper-V require a full installation of a general-purpose operating system (Windows Server 2008 for Hyper-V, Linux for XenServer) in the parent partition/domo in order to operate.

In the end, each of the virtualization products has its own set of advantages and disadvantages, and large organizations may end up using multiple products. For example, VMware vSphere might be best suited in the large corporate datacenter, while Microsoft Hyper-V or Citrix XenServer might be acceptable for test, development, or branch-office deployment. Organizations that don’t require VMware vSphere’s advanced features like VMware DRS, VMware FT, or Storage VMotion may also find that Microsoft Hyper-V or Citrix XenServer is a better fit for their needs.

As you can see, VMware vSphere offers some pretty powerful features that will change the way you view the resources in your datacenter. Some of these features, though, might not be applicable to all organizations, which is why VMware has crafted a flexible licensing scheme for organizations of all sizes.

Licensing VMware vSphere

With the introduction of VMware vSphere, VMware introduces a number of new licensing tiers and bundles that are intended to provide a good fit for every market segment. In this section, I'll explain the different licensing tiers—called *editions*—and how the various features that I've discussed so far fit into these editions.

Six editions of VMware vSphere are available:

- ◆ VMware vSphere Essentials
- ◆ VMware vSphere Essentials Plus
- ◆ VMware vSphere Standard
- ◆ VMware vSphere Advanced
- ◆ VMware vSphere Enterprise
- ◆ VMware vSphere Enterprise Plus

Each of these six editions features a different combination of features and products included. Table 1.2 lists the different editions and which features and products are included in each edition.

The vSphere Essentials and vSphere Essentials Plus editions are not licensed per CPU; these two editions are licensed for up to three physical servers. The cost for VMware vSphere Essentials is \$995 and includes one year of subscription; support is optional and available on a per-incident basis. vSphere Essentials Plus is \$2,995.

On all editions of VMware vSphere except Essentials, Support and Subscription (SnS) is sold separately. At least one year of SnS is required for each license. Subscription—but not support—is bundled with vSphere Essentials.

Also, it's important to note that all editions of VMware vSphere include support for thin provisioning, vCenter Update Manager, the VMsafe APIs, and the vStorage APIs. I did not include them in Table 1.2 because they are supported in all editions. I've specified only the list price for all editions; discounts may be available through a value-added reseller or VMware partner.

Looking carefully at Table 1.2, you may also note that VMware has moved away from licensing per pair of CPUs to licensing per CPU. With the advent of multicore processors, it's far more common to see physical servers with only a single physical CPU. To reflect this trend, VMware is licensing vSphere on a per-CPU basis.

In addition to all the editions listed in Table 1.2, VMware also offers a free edition of ESXi, named ESXi Free. ESXi Free includes the ESXi hypervisor and support for thin provisioning, but it cannot be managed by vCenter Server and does not support any of the other advanced features listed in Table 1.2. Customers can, though, go from ESXi Free to vSphere Standard, Advanced, or Enterprise simply by applying the appropriate license to an ESXi Free installation. This provides a great upgrade path for smaller organizations as they grow.

Now that you have an idea of how VMware licenses vSphere, I'd like to briefly review why an organization might choose to use vSphere and what benefits that organization could see as a result.

TABLE 1.2: Overview of VMware vSphere Product Editions

	ESSENTIALS	ESSENTIALS PLUS	STANDARD	ADVANCED	ENTERPRISE	ENTERPRISE PLUS
vCenter Server compatibility	vCenter Server for Essentials	vCenter Server for Essentials	vCenter Server Foundation and Standard	vCenter Server Foundation and Standard	vCenter Server Foundation and Standard	vCenter Server Foundation and Standard
Cores/CPU	6	6	6	12	6	12
vSMP support	4-way	4-way	4-way	4-way	4-way	8-way
RAM/Server	256GB	256GB	256GB	256GB	256GB	No license limit
VMware HA		✓	✓	✓	✓	✓
vCenter Data Recovery		✓		✓	✓	✓
VM Hot Add				✓	✓	✓
VMware FT				✓	✓	✓
VMotion				✓	✓	✓
Storage VMotion					✓	✓
VMware DRS					✓	✓
vNetwork Distributed Switch						✓
Host profiles						✓
Third-party multipathing						✓
List price per CPU			\$795	\$2245	\$2875	\$3495

Source: "VMware vSphere Pricing, Packaging and Licensing Overview" white paper published by VMware, available at www.vmware.com

Why Choose vSphere?

Much has been said and written about the total cost of ownership (TCO) and return on investment (ROI) for virtualization projects involving VMware virtualization solutions. Rather than rehashing that material here, I'll instead focus, very briefly, on why an organization should choose VMware vSphere as their virtualization platform.

ONLINE TCO CALCULATOR

VMware offers a web-based TCO calculator that helps you calculate the total cost of ownership and return on investment for a virtualization project using VMware virtualization solutions. This calculator is available online at www.vmware.com/go/calculator.

You've already read about the various features that VMware vSphere offers. To help understand how these features can benefit your organization, I'll apply them to the fictional XYZ Corporation. I'll walk through several different scenarios and look at how vSphere helps in these scenarios:

Scenario 1 XYZ Corporation's IT team has been asked by senior management to rapidly provision six new servers to support a new business initiative. In the past, this meant ordering hardware, waiting on the hardware to arrive, racking and cabling the equipment once it arrived, installing the operating system and patching it with the latest updates, and then installing the application. The timeframe for all these steps ranged anywhere from a few days to a few months and was typically a couple of weeks. Now, with VMware vSphere in place, the IT team can use vCenter Server's templates functionality to build a virtual machine, install the operating system, and apply the latest updates, and then rapidly clone—or copy—this virtual machine to create additional virtual machines. Now their provisioning time is down to hours, likely even minutes. Chapter 7 will discuss this functionality in more detail.

Scenario 2 Empowered by the IT team's ability to quickly respond to the needs of this new business initiative, XYZ Corporation is moving ahead with deploying updated versions of a line-of-business application. However, the business leaders are a bit concerned about upgrading the current version. Using the snapshot functionality present in ESX/ESXi and vCenter Server, the IT team can take a "point-in-time picture" of the virtual machine so that if something goes wrong during the upgrade, it's a simple rollback to the snapshot for recovery. Chapter 7 discusses snapshots.

Scenario 3 XYZ Corporation is really impressed with the IT team and vSphere's functionality and is now interested in expanding their use of virtualization. In order to do so, however, a hardware upgrade is needed on the servers currently running ESX/ESXi. The business is worried about the downtime that will be necessary to perform the hardware upgrades. The IT team uses VMotion to move virtual machines off one host at a time, upgrading each host in turn without incurring any downtime to the company's end users. Chapter 10 discusses VMotion in more depth.

Scenario 4 After the great success it's had virtualizing its infrastructure with vSphere, XYZ Corporation now finds itself in need of a new, larger shared storage array. vSphere's support for Fibre Channel, iSCSI, and NFS gives XYZ room to choose the most cost-effective storage solution available, and the IT team uses Storage VMotion to migrate the virtual machines without any downtime. Chapter 6 discusses Storage VMotion.

These scenarios begin to provide some idea of the benefits that organizations see when virtualizing with an enterprise-class virtualization solution like VMware vSphere.

The Bottom Line

Identify the role of each product in the vSphere product suite. The VMware vSphere product suite contains ESX and ESXi and vCenter Server. ESX and ESXi provide the base virtualization functionality and enable features like Virtual SMP. vCenter Server provides management for ESX/ESXi and enables functionality like VMotion, Storage VMotion, VMware Distributed Resource Scheduler (DRS), VMware High Availability (HA), and VMware Fault Tolerance (FT). VMware Consolidated Backup is a backup framework that allows for the integration of third-party backup solutions into a vSphere implementation.

Master It Which products are licensed features within the VMware vSphere suite?

Recognize the interaction and dependencies between the products in the vSphere suite. VMware ESX and ESXi form the foundation of the vSphere product suite, but some features require the presence of vCenter Server. Features like VMotion, Storage VMotion, VMware DRS, VMware HA, and VMware FT require both ESX/ESXi as well as vCenter Server.

Master It Name three features that are supported only when using vCenter Server along with ESX/ESXi.

Understand how vSphere differs from other virtualization products. VMware vSphere's hypervisor, ESX/ESXi, uses a type 1 bare-metal hypervisor that handles I/O directly within the hypervisor. This means that a host operating system, like Windows or Linux, is not required in order for ESX/ESXi to function. Although other virtualization solutions are listed as "type 1 bare-metal hypervisors," most other type 1 hypervisors on the market today require the presence of a "parent partition" or "dom0," through which all virtual machine I/O must travel.

Master It One of the administrators on your team asked whether he should install Windows Server on the new servers you purchased for ESX. What should you tell him, and why?

