

Chapter 1

Introducing VMware vSphere 6.7

VMware vSphere 6.7 builds on previous generations of VMware's enterprise-grade virtualization products that have been leading the industry since 2001. vSphere 6.7 gives greater control, performance, and extensibility with a focus on enabling workload security and mobility. With dynamic resource controls, high availability, and fault-tolerance features along with distributed resource management and operational 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 to tens of thousands of servers distributed among multiple clouds.

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 6.7

VMware vSphere is a comprehensive collection of products and features that together provide a full array of enterprise virtualization functionality. The vSphere product suite includes the following products and main features:

- ◆ VMware ESXi
- ◆ VMware vCenter Server
- ◆ vSphere Update Manager (VUM)
- ◆ vSphere Virtual Symmetric Multi-Processing
- ◆ vSphere vMotion and Storage vMotion
- ◆ vSphere Distributed Resource Scheduler (DRS)
- ◆ vSphere Storage DRS (SDRS)
- ◆ Storage I/O Control (SIOC) and Network I/O Control (NIOC)
- ◆ Storage-Based Policy Management (SBPM)
- ◆ vSphere High Availability (HA)

- ◆ vSphere Fault Tolerance (FT)
- ◆ vSphere Storage APIs
- ◆ VMware Virtual SAN (vSAN)
- ◆ vSphere Replication
- ◆ vSphere Content Library

Rather than waiting to introduce these products and features in their own chapters, we will introduce each product or feature in the following sections. This will allow us to explain how each one affects the design, installation, and configuration of your virtual infrastructure. After we cover the features and products in vSphere, 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. These additional products include VMware Horizon View, VMware vRealize Automation, and VMware vCenter Site Recovery Manager, just to name a few. VMware even offers bundles of vSphere and these other products in the vCloud Suite to make it easier for users to purchase and consume the products in their environments. However, because of the size and scope of these products, they are not covered in this book.

As of this writing, VMware vSphere 6.7 is the latest release of the VMware vSphere product family. This book covers functionality found in version 6.7. Where possible, we've tried to note differences between vSphere versions. For detailed information on other vSphere versions, refer to the previous books in the *Mastering VMware vSphere* series, also published by Sybex.

To help simplify navigation and to help you find information on the breadth of products and features in the vSphere product suite, we've prepared Table 1.1, which contains cross-references to where you can find more information about a particular product or feature elsewhere in the book.

TABLE 1.1: Product and Feature Cross-References

VMWARE vSPHERE PRODUCT OR FEATURE	CHAPTERS WHERE THIS IS COVERED
VMware ESXi	Installation:—Chapter 2 Networking:—Chapter 5 Storage:—Chapter 6
VMware vCenter Server	Installation:—Chapter 3 Networking:—Chapter 5 Storage:—Chapter 6 Security:—Chapter 8
vSphere Update Manager	Chapter 4
vSphere Host Client and vSphere Web Client	vSphere Host Client: Chapter 2 vSphere Web Client: Chapter 3
VMware vRealize Orchestrator and PowerCLI	Chapter 14

TABLE 1.1: Product and Feature Cross-References (CONTINUED)

VMWARE vSPHERE PRODUCT OR FEATURE	CHAPTERS WHERE THIS IS COVERED
vSphere Virtual Symmetric Multi-Processing	Chapter 9
vSphere vMotion and Storage vMotion	Chapter 12
vSphere Distributed Resource Scheduler	Chapter 12
vSphere Storage DRS	Chapter 12
Storage I/O Control and Network I/O Control	Chapter 11
Profile-driven storage	Chapter 6
vSphere High Availability	Chapter 7
vSphere Fault Tolerance	Chapter 7
vSphere Storage APIs for Data Protection	Chapter 7
VMware Data Protection	Chapter 7
VMware Virtual SAN	Chapter 6
vSphere Replication	Chapter 7
vSphere Flash Read Cache	Installation:—Chapter 6 Usage:—Chapter 11
vSphere Content Library	Chapter 9

First we'll look at the products that make up the VMware vSphere suite, and then we'll examine the major features. Let's start with the products in the suite, beginning with VMware ESXi.

Examining the Products in the vSphere Suite

In the following sections, we'll describe and review the products found in the vSphere product suite.

VMWARE 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 5 and later, including vSphere 6.7, the hypervisor comes solely in the form of VMware ESXi.

Longtime users of VMware vSphere will remember this as a shift in the way VMware provides the hypervisor. Prior to vSphere 5, the hypervisor was available in two forms: VMware ESX and VMware ESXi. Although both products shared the same core virtualization engine,

supported the same set of virtualization features, leveraged the same licenses, and were considered bare-metal installation hypervisors (also referred to as Type 1 hypervisors; see the sidebar “Type 1 and Type 2 Hypervisors”), there were still notable architectural differences. In VMware ESX, VMware used a Red Hat Enterprise Linux (RHEL)-derived Service Console to provide an interactive environment through which users could interact with the hypervisor. The Linux-based Service Console also included services found in traditional operating systems, such as a firewall, Simple Network Management Protocol (SNMP) agents, and a web server.

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 ESXi is a Type 1 bare-metal hypervisor. (In earlier versions of vSphere, VMware ESX was also considered a Type 1 bare-metal hypervisor.) Other Type 1 bare-metal hypervisors include KVM (part of the open source Linux kernel), Microsoft Hyper-V, and products based on the open source Xen hypervisor like Citrix Hypervisor (formerly XenServer) and Oracle VM.

VMware ESXi, on the other hand, is the next generation of the VMware virtualization foundation. Unlike VMware ESX, ESXi installs and runs without the Linux-based Service Console. This gives ESXi an ultralight footprint of approximately 150 MB. Despite the lack of the Service Console, ESXi provides all the same virtualization features that VMware ESX supported in earlier versions. Of course, ESXi 6.7 has been enhanced from earlier versions to support even more functionality, as you’ll see in this and future chapters.

The key reason that VMware ESXi is able to support the same extensive set of virtualization functionality as VMware ESX but without the Service Console is that the core of the virtualization functionality wasn’t found in the Service Console. It’s the *VMkernel* that is the foundation of the virtualization process. It’s the VMkernel that manages the virtual machines’ access to the underlying physical hardware by providing CPU scheduling, memory management, and virtual switch data processing. The section “VMware ESXi Architecture” in Chapter 2 will go into more detail on how the VMkernel supports and interacts with the rest of the hypervisor. Figure 1.1 shows the high level structure of VMware ESXi.

FIGURE 1.1
The VMkernel is the foundation of the virtualization functionality found in VMware ESXi.




TABLE 1.2: VMware ESXi Maximums

COMPONENT	VMWARE ESXi 6.7	VMWARE ESXi 6.5	VMWARE ESXi 6.0	VMWARE ESXi 5.5	VMWARE ESXi 5.0
Number of virtual CPUs per host	4,096	4,096	4,096	4,096	2,048
Number of logical CPUs (hyperthreading enabled)	768	576	480	320	160
Number of virtual CPUs per core	32	32	32	32	25
Amount of RAM per host	16 TB	12 TB	6 TB	4 TB	2 TB
Number of virtual machines per host	1,024	1,024	1,024	512	512
Number of virtual CPUs per virtual machine	128	128	128	64	32
Amount of RAM per virtual machine	6 TB	6 TB	4 TB	1 TB	1 TB

We mentioned earlier that VMware ESXi 6.7 is enhanced, and one such area of enhancement is in the configuration limits of what the hypervisor can support. Table 1.2 shows the configuration maximums for the last few versions of VMware ESXi.

These are just some of the configuration maximums. Where appropriate, future chapters will include additional values for VMware ESXi maximums for network interface cards (NICs), storage, virtual machines (VMs), and so forth.

Given that VMware ESXi is the foundation of virtualization within the vSphere product suite, you'll see content for VMware ESXi throughout the book. Table 1.1, earlier in this chapter, tells you where you can find more information about specific features of VMware ESXi.

VMWARE vCENTER SERVER

Stop for a moment to think about your current IT environment. Does it include Active Directory? There is a good chance it does. Now imagine your environment without Active Directory, without the ease of a centralized management database, without the single sign-on capabilities, and without the simplicity of groups. That's what managing VMware ESXi hosts would be like without using VMware vCenter Server. Not a very pleasant thought, is it? Now calm yourself down, take a deep breath, and know that vCenter Server, like Active Directory, is meant to provide a centralized management platform and framework for all ESXi hosts and their respective VMs. vCenter Server allows IT administrators to deploy, manage, monitor, automate, and secure a virtual infrastructure in a centralized fashion. To help provide scalability, vCenter Server leverages a backend database that stores all the data about the hosts and VMs.

In previous versions of VMware vSphere, vCenter Server was a Windows-only application. Version 6.7 of vSphere still offers this Windows-based installation of vCenter Server, but this will be the last release available for Windows. VMware offers a prebuilt vCenter Server Appliance (a virtual appliance, in fact, something you'll learn about in Chapter 10, "Using Templates and vApps") that is based on Photon, a thin and lightweight Linux distribution. The Linux-based vCenter Server appliance, or vCSA, is now a more feature-rich version of vCenter since development of new features has ceased on a Windows version. Chapter 3, "Installing and Configuring vCenter Server," will include more details on what is missing from the Windows version of vCenter Server. But for now, unless you already have an existing Windows-based installation, all new installations should use the Linux-based vCenter Server Appliance to ensure a supported future.

vCenter Server not only provides configuration and management capabilities—which include features such as VM templates, VM customization, rapid provisioning and deployment of VMs, role-based access controls, and fine-grained resource allocation controls—it also provides the tools for the more advanced features of vSphere vMotion, vSphere Distributed Resource Scheduler, vSphere High Availability, and vSphere Fault Tolerance. All of these features are described briefly in this chapter and in more detail in later chapters.

In addition to vSphere vMotion, vSphere Distributed Resource Scheduler, vSphere High Availability, and vSphere Fault Tolerance, using vCenter Server to manage ESXi hosts 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
- ◆ Host profiles, which allow you to bring greater consistency to host configurations across larger environments and to identify missing or incorrect configurations
- ◆ Storage I/O Control, which provides cluster-wide quality of service (QoS) controls so you can ensure critical applications receive sufficient storage I/O resources even during times of congestion
- ◆ vSphere Distributed Switches, which provide the foundation for networking settings and third-party virtual switches that span multiple hosts and multiple clusters
- ◆ Network I/O Control, which allows you to flexibly partition physical NIC bandwidth and provide QoS for different types of traffic
- ◆ vSphere Storage DRS, which enables VMware vSphere to dynamically migrate storage resources to meet demand, much in the same way that DRS balances CPU and memory utilization

vCenter Server plays a central role in any sizable VMware vSphere implementation. In Chapter 3, we discuss planning and installing vCenter Server as well as look at ways to ensure its availability. As previously mentioned, Chapter 3 will examine the differences between the Windows-based version of vCenter Server and the Linux-based vCenter Server virtual appliance. Because of vCenter Server's central role in a VMware vSphere deployment, we'll touch on vCenter Server in almost every chapter throughout the rest of the book. Refer to Table 1.1, earlier in this chapter, for specific cross-references.

vCenter Server is available in three packages:

- ◆ vCenter Server Essentials is integrated into the vSphere Essentials kits for small office deployment.
- ◆ vCenter Server Foundation provides all the functionality of vCenter Server, but for a limited number of ESXi hosts.
- ◆ vCenter Server Standard provides all the functionality of vCenter Server, including provisioning, management, monitoring, and automation.

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

VSPHERE UPDATE MANAGER

vSphere Update Manager is a component of vCenter Server that helps users keep their ESXi hosts and select VMs patched with the latest updates. vSphere 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 ESXi hosts
- ◆ Full integration with other vSphere features like Distributed Resource Scheduler

vSphere Update Manager works as an installable package with the Windows-based installation of vCenter Server as well as the prepackaged feature pre-installed in the vCenter Server virtual appliance. Refer to Table 1.1 for more information on where vSphere Update Manager is described in this book.

VMWARE VSPHERE CLIENT AND VSPHERE HOST CLIENT

vCenter Server provides a centralized management framework for VMware ESXi hosts, but it’s the web-based vSphere Client (like its predecessor, the Windows-based vSphere Desktop Client) where you will spend most of your time.

With the release of vSphere 5, VMware shifted its primary administrative interface to a web-based vSphere Client built on Adobe Flash. The “vSphere Web Client” provided a web-based user interface for managing a virtual infrastructure and enabled you to manage your infrastructure without needing to install the Windows-based vSphere Desktop Client on a system. Unfortunately, the Flash-based client was not well received and ultimately VMware decided to move to the HTML5 web standard. This transition took a number of releases, and as a result, multiple clients could be used to do some (but not all) administrative tasks.

Initially, the HTML5-based vSphere Web Client (simply known as the “vSphere Client”) offered only a subset of the functionality available to the “Flash” vSphere Web Client. However, in subsequent releases—including the 6.7 release—the vSphere Client has been enhanced and expanded to include most of the functionality you need to manage a vSphere environment.

Further, VMware has stated that the Flash-based vSphere Web Client and the Windows-based vSphere Desktop Client are now end-of-life. Luckily, the step-by-step procedures for the Flash-based vSphere Web Client and the HTML5-based vSphere client are usually identical. For this reason, we'll use Flash-based vSphere Web Client screen shots and step-by-step guidance throughout this book to ensure each instruction can be completed with the same client.

Administering hosts without vCenter has also changed. You now access the user interface by browsing to the URL of each ESXi host. This loads an HTML5-based user interface (UI) but only for that particular host. No client installation is needed.

This can be a little confusing if this is your first foray into the VMware landscape, so let us recap. The vSphere Web Client, based on Flash, has been deprecated. The Windows-installable vSphere Desktop Client (for connecting to vCenter and hosts) has been deprecated. To administer vCenter, and hosts attached to a vCenter Server, use the new HTML5-based vSphere Client or the Flash-based vSphere Web Client. To administer ESXi hosts directly, without vCenter, use the HTML5-based vSphere Host Client.


Examining the Features in VMware vSphere

In the following sections, we'll take a closer look at some of the features available in the vSphere product suite. We'll start with Virtual SMP.

vSPHERE VIRTUAL SYMMETRIC MULTI-PROCESSING

The vSphere Virtual Symmetric Multi-Processing (vSMP or Virtual SMP) product allows you to construct VMs with multiple virtual processor cores and/or sockets. vSphere Virtual SMP is *not* the licensing product that allows ESXi to be installed on servers with multiple processors; it is the technology that allows the use of multiple processors *inside* a VM. Figure 1.2 identifies the differences between multiple processors in the ESXi host system and multiple virtual processors.

FIGURE 1.2
vSphere Virtual SMP allows VMs to be created with more than one virtual CPU.



With vSphere Virtual SMP, applications that require and can actually use multiple CPUs can be run in VMs 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).

This functionality also allows users to specify multiple virtual cores per virtual CPU. Using this feature, a user could provision a dual “socket” VM with two cores per “socket” for a total of four virtual cores. This approach gives users tremendous flexibility in carving up CPU processing power among the VMs.

VSPHERE VMOTION AND VSPHERE STORAGE VMOTION

If you have read anything about VMware, you have most likely read about the extremely useful feature called vMotion. vSphere vMotion, also known as *live migration*, is a feature of ESXi and vCenter Server that allows you to move a running VM from one physical host to another physical host without having to power off the VM. This migration between two physical hosts occurs with no downtime and with no loss of network connectivity to the VM. The ability to manually move a running VM between physical hosts on an as-needed basis is a powerful feature that has a number of use cases in today’s datacenters.

Suppose a physical machine has experienced a nonfatal hardware failure and needs to be repaired. You can easily initiate a series of vMotion operations to remove all VMs from an ESXi host that is to undergo scheduled maintenance. After the maintenance is complete and the server is brought back online, you can use vMotion to return the VMs to the original server.

Alternately, consider a situation in which you are migrating from one set of physical servers to a new set of physical servers. Assuming that the details have been addressed—and we’ll discuss the details of vMotion in Chapter 12, “Balancing Resource Utilization”—you can use vMotion to move the VMs from the old servers to the newer servers, making quick work of a server migration with no interruption of service.

Even in normal day-to-day operations, vMotion can be used when multiple VMs on the same host are in contention for the same resource (which ultimately causes poor performance across all the VMs). With vMotion, you can migrate any VMs facing contention to another ESXi host with greater availability for the resource in demand. For example, when two VMs contend with each other for CPU resources, you can eliminate the contention by using vMotion to move one VM to an ESXi host with more available CPU resources.

vMotion moves the execution of a VM, relocating the CPU and memory footprint between physical servers but leaving the storage untouched. Storage vMotion builds on the idea and principle of vMotion: you can leave the CPU and memory footprint untouched on a physical server but migrate a VM’s storage while the VM is still running.

Deploying vSphere in your environment generally means that lots of shared storage—Fibre Channel or FCoE or iSCSI SAN or NFS—is needed. What happens when you need to migrate from an older storage array to newer storage hardware based on vSAN? What kind of downtime would be required? Or what about a situation where you need to rebalance utilization of the array, either from a capacity or performance perspective?

With the ability to move storage for a running VM between datastores, Storage vMotion lets you address all of these situations without downtime. This feature ensures that outgrowing datastores or moving to new storage hardware does not force an outage for the affected VMs and provides you with yet another tool to increase your flexibility in responding to changing business needs.

VSPHERE DISTRIBUTED RESOURCE SCHEDULER

vMotion is a manual operation, meaning that you must initiate the vMotion operation. What if VMware vSphere could perform vMotion operations automatically? That is the basic idea behind vSphere Distributed Resource Scheduler (DRS). If you think that vMotion sounds exciting, your anticipation will only grow after learning about DRS. DRS, simply put, leverages vMotion to provide automatic distribution of resource utilization across multiple ESXi hosts that are configured in a cluster.

Given the prevalence of Microsoft Windows Server in today's datacenters, the use of the term *cluster* often draws IT professionals into thoughts of Microsoft Windows Server Failover Clusters. Windows Server clusters are often active-passive or active-active-passive clusters. However, ESXi clusters are fundamentally different, operating in an active-active mode to aggregate and combine resources into a shared pool. Although the underlying concept of aggregating physical hardware to serve a common goal is the same, the technology, configuration, and feature sets are quite different between VMware ESXi clusters and Windows Server clusters.

AGGREGATE CAPACITY AND SINGLE HOST CAPACITY

Although we say that a DRS cluster is an implicit aggregation of CPU and memory capacity, it's important to keep in mind that a VM is limited to using the CPU and RAM of a single physical host at any given time. If you have two small ESXi servers with 64 GB of RAM each in a DRS cluster, the cluster will correctly report 128 GB of aggregate RAM available, but any given VM will not be able to use more than approximately 64 GB of RAM at a time.

An 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 VMs assigned to the cluster (keeping in mind that any given VM can only use resources from one host; see the sidebar "Aggregate Capacity and Single Host Capacity"). The goal of DRS is twofold:

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

The first part of DRS is often referred to as *intelligent placement*. DRS can automate the placement of each VM 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 VM at that moment.

DRS isn't limited to operating only at VM startup, though. DRS also manages the VM's location while it is running. For example, let's say three hosts have been configured in an ESXi cluster with DRS enabled. When one of those hosts begins to experience a high contention for CPU utilization, DRS detects that the cluster is imbalanced in its resource usage and uses an internal algorithm to determine which VM(s) should be moved in order to create the least imbalanced cluster. For every VM, DRS will simulate a migration to each host and the results will

be compared. The migrations that create the least imbalanced cluster will be recommended or automatically performed, depending on the DRS configuration.

DRS performs these on-the-fly migrations without any downtime or loss of network connectivity to the VMs by leveraging vMotion, the live migration functionality we described earlier. This makes DRS extremely powerful because it allows clusters of ESXi hosts to dynamically rebalance their resource utilization based on the changing demands of the VMs running on that cluster.

FEWER BIGGER SERVERS OR MORE SMALLER SERVERS?

Recall from Table 1.2 that VMware ESXi supports servers with up to 768 logical CPU cores and up to 16 TB of RAM. With vSphere DRS, though, you can combine multiple smaller servers for the purpose of managing aggregate capacity. This means that bigger, more powerful servers might not be better servers for virtualization projects. These larger servers, in general, are significantly more expensive than smaller servers, and using a greater number of smaller servers (often referred to as “scaling out”) may provide greater flexibility than a smaller number of larger servers (often referred to as “scaling up”). The key thing to remember is that a bigger server isn’t necessarily a better server.

vSPHERE STORAGE DRS

vSphere Storage DRS takes the idea of vSphere DRS and applies it to storage. Just as vSphere DRS helps to balance CPU and memory utilization across a cluster of ESXi hosts, Storage DRS helps balance storage capacity and storage performance across a cluster of datastores using mechanisms that echo those used by vSphere DRS.

Earlier, we described vSphere DRS’s feature called intelligent placement, which automates the placement of new VMs based on resource usage within an ESXi cluster. In the same fashion, Storage DRS has an intelligent placement function that automates the placement of VM virtual disks based on storage utilization. Storage DRS does this through the use of datastore clusters. When you create a new VM, you simply point it to a datastore cluster, and Storage DRS automatically places the VM’s virtual disks on an appropriate datastore within that datastore cluster.

Likewise, just as vSphere DRS uses vMotion to balance resource utilization dynamically, Storage DRS uses Storage vMotion to rebalance storage utilization based on capacity and/or latency thresholds. Because Storage vMotion operations are typically much more resource-intensive than vMotion operations, vSphere provides extensive controls over the thresholds, timing, and other guidelines that will trigger a Storage DRS automatic migration via Storage vMotion.

STORAGE I/O CONTROL AND NETWORK I/O CONTROL

VMware vSphere has always had extensive controls for modifying or controlling the allocation of CPU and memory resources to VMs. Before the release of vSphere 4.1, however, vSphere could not apply extensive controls to storage I/O and network I/O. Storage I/O Control and Network I/O Control address that shortcoming.

Storage I/O Control (SIOC) allows you to assign relative priority to storage I/O as well as assign storage I/O limits to VMs. These settings are enforced cluster-wide; when an ESXi host detects storage congestion through an increase of latency beyond a user-configured threshold, it will apply the settings configured for that VM. The result is that you can help the VMs that need

priority access to storage resources get more of the resources they need. In vSphere 4.1, Storage I/O Control applied only to VMFS storage; vSphere 5 extended that functionality to NFS datastores.

The same goes for Network I/O Control (NIOC), which provides you with more granular controls over how VMs use network bandwidth provided by the physical NICs. As the widespread adoption of 10 Gigabit Ethernet (10GbE) and faster continues, Network I/O Control provides you with a way to more reliably ensure that network bandwidth is properly allocated to VMs based on priority and limits.

POLICY-BASED STORAGE

With profile-driven storage, vSphere administrators can use storage capabilities and VM storage profiles to ensure VMs reside on storage that provides the necessary levels of capacity, performance, availability, and redundancy. Profile-driven storage is built on two key components:

- ◆ Storage capabilities, leveraging vSphere APIs for storage awareness (VASA)
- ◆ VM storage profiles

Storage capabilities are either provided by the storage array itself (if the array can use VASA and/or defined by a vSphere administrator. These storage capabilities represent various attributes of the storage solution.

VM storage profiles define the storage requirements for a VM and its virtual disks. You create VM storage profiles by selecting the storage capabilities that must be present for the VM to run. Datastores that have all the capabilities defined in the VM storage profile are compliant with the VM storage profile and represent possible locations where the VM could be stored.

This functionality gives you much greater visibility into storage capabilities and helps ensure that the appropriate functionality for each VM is indeed being provided by the underlying storage. These storage capabilities can be explored extensively by using VVOLs or vSAN.

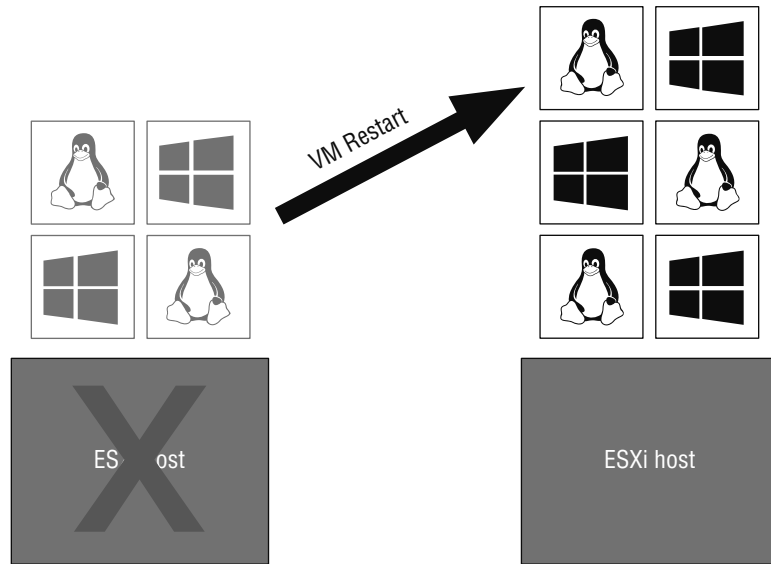
Refer to Table 1.1 to find out which chapter discusses profile-driven storage in more detail.

VSPHERE HIGH AVAILABILITY

In many cases, high availability—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 ESXi clusters called vSphere 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 incorrect in that vSphere HA does not function like a high-availability configuration in Windows. The vSphere HA feature provides an automated process for moving and restarting VMs that were running on an ESXi host at a time of server failure (or other qualifying infrastructure failure, as we’ll describe in Chapter 7, “Ensuring High Availability and Business Continuity”). Figure 1.3 depicts the VM migration that occurs when an ESXi host that is part of an HA-enabled cluster experiences failure.

FIGURE 1.3
The vSphere HA feature will restart any VMs that were previously running on an ESXi host that experiences server or storage path failure.



The vSphere HA feature, unlike DRS, does not always use the vMotion technology as a means of migrating servers to another host. vMotion applies only to planned migrations, where both the source and destination ESXi host are running and functioning. Let us explain what we mean. In a vSphere HA failover situation, there is no anticipation of failure; it is not a planned outage, which means there is no time to perform a vMotion operation. vSphere HA is intended to minimize unplanned downtime because of the failure of a physical ESXi host or other infrastructure components. We'll go into more detail in Chapter 7 on what kinds of failures vSphere HA helps protect against.

vSPHERE HA IMPROVEMENTS FROM vSPHERE 5

vSphere HA received a few notable improvements over the last few releases. Scalability was significantly improved, and it was closely integrated with the intelligent placement functionality of vSphere DRS, giving vSphere HA greater ability to restart VMs in the event of a host failure. However, perhaps the most significant improvement is the complete rewrite of the underlying architecture for vSphere HA; this newer architecture, known as Fault Domain Manager (FDM), eliminated many of the constraints found in earlier versions of VMware vSphere (before version 5.0).

By default, vSphere HA does not provide failover in the event of a guest OS failure, although you can configure vSphere HA to monitor VMs and restart them automatically if they fail to respond to an internal heartbeat. This feature is called VM Failure Monitoring, and it uses a combination of internal heartbeats and I/O activity to attempt to detect if the guest OS inside a VM has stopped functioning. If the guest OS has stopped functioning, the VM can be restarted automatically.

With vSphere HA in a failure scenario, it's important to understand that there will be an interruption of service. If a physical host or storage device fails, vSphere HA restarts the VM, and while the VM is restarting, the applications or services provided by that VM are unavailable. The only time that this is not true is if Proactive HA is enabled on the host. Proactive HA uses hardware monitoring to proactively move VMs from a host that is suffering from hardware issues.


For users who need even higher levels of availability than can be provided using vSphere HA, vSphere Fault Tolerance (FT), which is described in the next section, can help.

VSPHERE FAULT TOLERANCE

Although vSphere HA provides a certain level of availability for VMs in the event of physical host failure, this might not be good enough for some workloads. vSphere FT might help in these situations.

As we described in the previous section, vSphere HA protects against unplanned physical server failure by providing a way to automatically restart VMs upon physical host failure. This need to restart a VM in the event of a physical host failure means that some downtime—generally less than three minutes—is incurred. vSphere FT goes even further and eliminates any downtime in the event of a physical host failure. vSphere FT maintains a mirrored secondary VM on a separate physical host that is kept in lockstep with the primary VM. vSphere's newer Fast Checkpointing technology supports FT of VMs with one to four vCPUs. Everything that occurs on the primary (protected) VM also occurs simultaneously on the secondary (mirrored) VM, so that if the physical host for the primary VM fails, the secondary VM can immediately step in and take over without any loss of connectivity. vSphere FT will also automatically re-create the secondary (mirrored) VM on another host if the physical host for the secondary VM fails, as illustrated in Figure 1.4. This ensures protection for the primary VM at all times.

FIGURE 1.4
vSphere FT provides protection against host failures with no downtime experienced by the VMs.



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

vSphere FT can work in conjunction with vMotion. As of vSphere 5.0, vSphere FT is also integrated with vSphere DRS, although this feature does require Enhanced vMotion Compatibility (EVC). VMware recommends that multiple FT virtual machines with multiple vCPUs have 10GbE networks between hosts.

VSPHERE STORAGE APIS FOR DATA PROTECTION AND VMWARE DATA PROTECTION

One of the most critical aspects of any IT infrastructure, not just virtualized infrastructure, is a solid backup strategy as defined by a company's disaster recovery and business continuity plan. To help address organizational backup needs, VMware vSphere has a key component: the vSphere Storage APIs for Data Protection (VADP).

VADP is a set of application programming interfaces (APIs) that backup vendors leverage in order to provide enhanced backup functionality of virtualized environments. VADP enables functionality like file-level backup and restore; support for incremental, differential, and full-image backups; native integration with backup software; and support for multiple storage protocols.

On its own, though, VADP is just a set of interfaces, like a framework for making backups possible. You can't actually back up VMs with VADP. You'll need a VADP-enabled backup application. There are a growing number of third-party backup applications that are designed to work with VADP from vendors such as CommVault, DelleMC, and Veritas.

VSPHERE DATA PROTECTION

In vSphere 5.1, VMware phased out its earlier data protection tool, VMware Data Recovery (VDR), in favor of vSphere Data Protection (VDP). Although VDR was provided with vSphere 5.0, VDR is not supported with vSphere 5.1 and later. Subsequently, VMware has also discontinued VDP from vSphere 6.5. Backups of your vSphere environment now need to be handled by another vendor.

VIRTUAL SAN (vSAN)

vSAN was a major new feature included with, but licensed separately from, vSphere 5.5 and later. It is the evolution of work that VMware has been doing for a number of years now. vSAN lets organizations leverage the internal local storage found in individual compute nodes and turn it into a *virtual SAN*.

vSAN requires a minimum of two ESXi hosts (or nodes) for some limited configurations, but it will scale to as many as 64. vSAN also requires solid-state (flash) storage in each of the compute nodes providing vSAN storage; this is done to help improve I/O performance given that most compute nodes have a limited number of physical drives present. vSAN pools the aggregate storage across the compute nodes, allowing you to create a datastore that spans multiple compute nodes. vSAN employs policies and algorithms to ensure performance or to help protect

against data loss, such as ensuring that the data exists on multiple participating vSAN nodes at the same time.

There's more information on vSAN in Chapter 6, "Creating and Configuring Storage Devices."

VSPHERE REPLICATION

vSphere Replication brings data replication, which is a feature typically found in hardware storage platforms, into vSphere itself. It's been around since vSphere 5.0, when it was only enabled for use in conjunction with VMware Site Recovery Manager (SRM) 5.0. In vSphere 5.1, vSphere Replication was decoupled from SRM and enabled for independent use without VMware SRM.

vSphere Replication enables customers to replicate VMs from one vSphere environment to another vSphere environment. Typically, this means from one data center (often referred to as the primary or production data center) to another datacenter (typically the secondary, backup, or disaster recovery [DR] site). Unlike hardware-based solutions, vSphere Replication operates on a per-VM basis, so it gives customers very granular control over which workloads will be replicated and which workloads won't be replicated.

You can find more information about vSphere Replication in Chapter 7.

VSPHERE FLASH READ CACHE

Since the release of vSphere 5.0 in 2011, the industry has seen tremendous uptake in the use of solid-state or "flash" storage across a wide variety of use cases. Because solid-state storage can provide massive numbers of I/O operations per second (IOPS) and very large bandwidth (Mbps) it can handle the increasing I/O demands of virtual workloads. However, depending on the performance, solid-state storage is still typically more expensive on a per-gigabyte basis than traditional, magnetic-disk-based storage and therefore is often first deployed as a caching mechanism to help speed up frequently accessed data.

Unfortunately, without support in vSphere for managing solid-state storage as a caching mechanism, vSphere architects and administrators have had difficulty fully leveraging solid-state storage in their environments. In vSphere 5.5 and later, VMware addresses that limitation through a feature called *vSphere Flash Read Cache*.

vSphere Flash Read Cache brings full support for using solid-state storage as a caching mechanism to vSphere. Using this feature, you can assign solid-state caching space to VMs in much the same way as you assign CPU cores, RAM, or network connectivity to VMs. vSphere manages how the solid-state caching capacity is allocated and assigned as well as how it is used by the VMs.

VMWARE vSPHERE COMPARED TO MICROSOFT HYPER-V AND CITRIX HYPERVISOR

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

To make accurate comparisons between vSphere and others, you must include only Type 1 ("bare-metal") virtualization solutions. This would include ESXi, Microsoft Hyper-V, and Citrix Hypervisor.

It would not include products such as VMware Fusion or Workstation and Windows Virtual PC, all 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 Hypervisor route all the VM I/O through the “parent partition” or “dom0.” 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—the general-purpose operating system running in the parent partition—supports a particular type of hardware, Hyper-V supports it also. Hyper-V “piggybacks” on Windows’ hardware drivers and the I/O stack. The same can be said for Citrix Hypervisor, although its “dom0” runs Linux and not Windows.

VMware 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. 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, and nowhere is it more greatly demonstrated than in ESXi, which has a small footprint yet provides a full-featured virtualization solution. Both Citrix Hypervisor and Microsoft Hyper-V require a full installation of a general-purpose operating system (Windows Server for Hyper-V, Linux for Citrix Hypervisor) in the parent partition/dom0 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 a large corporate datacenter, whereas Microsoft Hyper-V or Citrix Hypervisor might be acceptable for test, development, or branch office deployment. Organizations that don’t require VMware vSphere’s advanced features like vSphere DRS, vSphere FT, or Storage vMotion may also find that Microsoft Hyper-V or Citrix Hypervisor 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. vSphere also has a wide range of features and functionality. 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 each new version, VMware usually revises the licensing tiers and bundles intended to provide a good fit for every market segment. Introduced with vSphere 5.1 (and continuing on through vSphere 6.7), VMware refined this licensing arrangement with the vCloud Suite—a bundling of products including vSphere, vRealize Automation, vCenter Site Recovery Manager, and vRealize Operations Management Suite.

Although licensing vSphere via the vCloud Suite is likely the preferred way of licensing vSphere moving forward, discussing all the other products included in the vCloud Suite is beyond the scope of this book. Instead, we’ll focus on vSphere and explain how the various features discussed so far fit into vSphere’s licensing model when vSphere is licensed stand-alone.

One thing that you need to be aware of is that VMware may change the licensing tiers and capabilities associated with each tier at any time. You should visit the vSphere products web page (www.vmware.com/products/vsphere.html) or talk to your VMware representative before making any purchasing decisions.

VSPHERE VS VSPHERE WITH OPERATIONS MANAGEMENT

VMware sells “stand-alone” vSphere in one of two ways: as vSphere, with all the various kits and editions, and as vSphere with Operations Management. vSphere with Operations Management is the same as vSphere but adds the vRealize Operations Management product. In this section, we are focused on stand-alone vSphere only, but keep in mind that vSphere with Operations Management would be licensed and packaged in much the same way.

You’ve already seen how VMware packages and licenses VMware vCenter Server, but here’s a quick review:

- ◆ VMware vCenter Server for Essentials is bundled with the vSphere Essentials kits (more on the kits in just a moment).
- ◆ VMware vCenter Server Standard includes all functionality and does not have a preset limit on the number of vSphere hosts it can manage (although normal sizing limits do apply). vRealize Orchestrator is included only in the Standard edition of vCenter Server.

In addition to the two editions of vCenter Server, VMware offers two editions of VMware vSphere:

- ◆ vSphere Standard Edition
- ◆ vSphere Enterprise Plus Edition

These editions are differentiated primarily by the features each edition supports, although there are some capacity limitations with the different editions.

Table 1.3 summarizes the features that are supported for each edition of VMware vSphere 6.7.

TABLE 1.3: Overview of VMware vSphere Product Editions

	ESSENTIALS KIT	ESSENTIALS PLUS KIT	STANDARD	ENTERPRISE PLUS
vCenter Server compatibility	vCenter Server Essentials	vCenter Server Essentials	vCenter Server Standard	vCenter Server Standard
vCPUs per VM	128	128	128	128
Cross-Host / vSwitch vMotion		X	X	X
Cross-vCenter / Long Distance / Cross-Cloud vMotion				X
High Availability		X	X	X
Data Protection		X	X	X
vSphere Replication		X	X	X

TABLE 1.3: Overview of VMware vSphere Product Editions (CONTINUED)

	ESSENTIALS KIT	ESSENTIALS PLUS KIT	STANDARD	ENTERPRISE PLUS
vShield Endpoint		X	X	X
Hot Add			X	X
Fault Tolerance			2 vCPU	4 vCPU
Storage vMotion			X	X
Virtual Volumes and Storage Policy-based Management			X	X
Distributed Resource Scheduler and Distributed Power Management				X
Storage APIs for Array Integration, Multipathing				X
Big Data Extensions				X
Reliable Memory				X
Distributed Switch				X
I/O Controls (Network and Storage) and SR-IOV				X
Host Profiles				X
Auto Deploy				X
Storage DRS				X
Flash Read Cache				X
Content Library				X
Proactive High Availability				X
VM-level Encryption				X

Source: “VMware vSphere 6.7 Licensing, Pricing and Packaging” white paper published by VMware, available at www.vmware.com.

It's important to note that all editions of VMware vSphere 6.7 include support for thin provisioning, vSphere Update Manager, and the vSphere Storage APIs for Data Protection. We did not include them in Table 1.3 because these features are supported in all editions. Because prices change and vary depending on partner, region, and other factors, we have not included any pricing information here. Also, we did not include vSAN in Table 1.3, because it is licensed separately from vSphere.

For all editions of vSphere, VMware requires at least one year of Support and Subscription (SnS). The only exception is the Essential Kits, as we'll explain in a moment.

In addition to the different editions described previously, VMware offers some bundles, referred to as *kits*.

Essentials Kits are all-in-one solutions for small environments, supporting up to three vSphere hosts with two CPUs each. To support three hosts with two CPUs each, the Essentials Kits come with six licenses. All these limits are product-enforced. Two Essentials Kits are available:

- ◆ VMware vSphere Essentials
- ◆ VMware vSphere Essentials Plus

You can't buy these kits on a per-CPU basis; they are bundled solutions for three servers. vSphere Essentials includes one year of subscription; support is optional and available on a per-incident basis. Like other editions, vSphere Essentials Plus requires at least one year of SnS, which must be purchased separately and is not included in the bundle.

The Retail and Branch Offices (ROBO) Kits are differentiated from the "normal" Essentials and Essentials Plus Kits only by the licensing guidelines. These kits are licensed per pack of 25 virtual machines. Central management of all the sites via vCenter Server Standard is possible, though vCenter Server Standard must be purchased separately. vCenter Server Essentials is included.

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

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, we'll instead focus, 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 TCO and ROI for a virtualization project using VMware virtualization solutions. This calculator is available online at <https://www.vmware.com/go/tcocalculator/>.

You've already read about the various features that VMware vSphere offers. To help you understand how these features can benefit your organization, we'll apply them to the fictional XYZ Corporation. We'll walk you through several scenarios and show 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 time frame 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 VM, install the operating system, and apply the latest updates, and then rapidly clone—or copy—this VM to create additional VMs. Now their provisioning time is down to hours, likely even minutes. Chapter 10 discusses this functionality in 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 ESXi and vCenter Server, the IT team can take a “point-in-time picture” of the VM so that if something goes wrong during the upgrade, it’s a simple rollback to the snapshot for recovery. Chapter 9, “Creating and Managing Virtual Machines,” discusses snapshots.

Scenario 3 XYZ Corporation is impressed with the IT team and vSphere’s functionality and is now interested in expanding its use of virtualization. To do so, however, a hardware upgrade is needed on the servers currently running ESXi. The business is worried about the downtime that will be necessary to perform the hardware upgrades. The IT team uses vMotion to move VMs off one host at a time, upgrading each host in turn without incurring any downtime to the company’s end users. Chapter 12 discusses vMotion in more depth.

Scenario 4 After the great success it has 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, NFS, or vSAN gives XYZ room to choose the most cost-effective storage solution available, and the IT team uses Storage vMotion to migrate the VMs without any downtime. Chapter 12 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.

WHAT DO WE VIRTUALIZE WITH VMWARE vSPHERE?

Virtualization, by its very nature, means that you are going to take multiple operating systems—such as Microsoft Windows, Linux, Solaris, or Novell NetWare—and run them on a single physical server. While VMware vSphere offers broad support for virtualizing a wide range of operating systems, it would be almost impossible for us to discuss how virtualization impacts all the different versions of all the operating systems that vSphere supports.

VMware provides in-depth information on all the operating systems it supports and how vSphere interacts with those operating systems on its website at www.vmware.com. Also worth noting is that while VMware lists the operating systems it supports, many older operating systems still work when virtualized—they are just not on the “supported” list anymore, so your mileage may vary. As an example, we have virtualized MS-DOS and Windows 3.1 for the fun of it, but these are not listed on VMware’s website.

The Bottom Line

Identify the role of each product in the vSphere product suite. The VMware vSphere product suite contains VMware ESXi and vCenter Server. ESXi provides the base virtualization functionality and enables features like Virtual SMP. vCenter Server provides management for ESXi and enables functionality like vMotion, Storage vMotion, vSphere Distributed Resource Scheduler (DRS), vSphere High Availability (HA), and vSphere Fault Tolerance (FT). Storage I/O Control and Network I/O Control provide granular resource controls for VMs. The vStorage APIs for Data Protection (VADP) provide 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?

Master It Which two features of VMware ESXi and VMware vCenter Server together aim to reduce or eliminate downtime due to unplanned hardware failures?

Recognize the interaction and dependencies between the products in the vSphere suite. VMware ESXi forms the foundation of the vSphere product suite, but some features require the presence of vCenter Server. Features like vMotion, Storage vMotion, vSphere DRS, vSphere HA, vSphere FT, SIOC, and NIOC require ESXi as well as vCenter Server.

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

Master It Name two features that are supported without vCenter Server but with a licensed installation of ESXi.

Understand how vSphere differs from other virtualization products. VMware vSphere's hypervisor, 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 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 VM I/O must travel.

Master It One of the administrators on your team asked whether he should install the standard Red Hat Linux (RHEL) deployment on the new servers you purchased for ESXi. What should you tell him, and why?