PART I

Building a VMware vSphere Environment

IN THIS PART

1

Introduction to vSphere

IN THIS CHAPTER, YOU WILL LEARN TO:

VSphere is here! Administrators who have been around for a while may think of the new product as the fourth generation, or simply VMware Infrastructure 4. However, the name vSphere better aligns the new product with the direction that virtualization is taking, and this version of VMware introduces many new and promising features. Not since the introduction of VMware Infrastructure 3 (VI3) has there been this much excitement surrounding new features that promise to continue to revolutionize the infrastructure of the modern and evolving datacenter. The most sought-after three features—VMotion, Distributed Resource Scheduler (DRS), High Availability (HA)—have been improved and are better than ever. And of course a new version wouldn't be complete without additional new features—VMware has worked hard to make this release even bigger than the VI3 release.

Understand the Legacy Features of vSphere

Welcome to the legacy features of vSphere. Why devote an entire section to them? Because they serve as the foundation that brings tremendous flexibility to managing an x86 environment. One of the best is vMotion, which offers the ability to relocate a running virtual machine or server from one physical location to another without any downtime. Another legacy feature is Dynamic Resource Scheduler, which you use to make sure that your servers are getting all the resources they deserve. Finally, with High Availability you'll never have to rush into the office to address bad hardware. There are other legacy features, but we'll cover only the best!

VMotion

The last major version release of VMotion was VI3, released in 2006. VMotion remains one of the most powerful features of virtualization today. With VI3, you can perform work on underlying hosts during business hours rather than having to wait until the wee hours of the morning or weekends to upgrade BIOSs or firmware or do something as simple as add more memory to a host. VMotion requires that each underlying host have a CPU that uses the same instruction set, because,

after all, moving a running virtual machine (VM) from one physical host to another physical host without any downtime is a phenomenal feat. VMware runs on top of the Virtual Machine File System (VMFS); Windows still runs on New Technology File System (NTFS), but the underlying file system is VMFS3. VMFS3 allows for multiple access, and that is how one host can pass a running VM to another host without downtime or interruptions. It is important to realize that even momentary downtime can be critical for applications and databases. Zero downtime when moving a VM from one host to another physical host is crucial.

Unfortunately, there is no way to move from Intel to AMD, or vice versa. In the past there were even issues going from an older Intel CPU to a newer Intel CPU.

VMware has several years of experience mastering virtualization while the competitors are playing catch-up. Furthermore, VMware has explored many approaches to virtualization and has seen firsthand where some approaches fall short and where some excel.

One downside to VMotion is that the technology does require shared storage, but the virtual machine files do not move from that shared storage during the logical transition. If, for example, you have to change the virtual machine's physical location, you must first power down the VM and then "migrate" it from one logical unit number (LUN) or hard drive to another LUN or hard drive.

Another possible downside to VMotion is that traditional Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS) may not work as originally designed. Part of the reason for this is that the traffic of VMs that are communicating with one another inside a host never leaves the host and therefore cannot be inspected. Virtual appliances are being developed to address this concern as they have the ability to run side by side VMs.

Since uptime is important, VMware developed Storage VMotion so that the physical location of a running VM can be changed, again without any downtime and without losing any transactional information. Obviously Storage VMotion is very exciting because one of the reasons that virtualization is the hottest technology in IT today is the flexibility it brings to the datacenter (compared to running servers the old-fashioned way in a physical environment).

There are other ways to leverage the technology. Virtual machines can be moved on the fly from shared storage to local storage if you need

to perform maintenance on shared storage or if LUNs have to be moved to other hosts. Imagine moving a physical server by simply clicking on a new rack in your datacenter and then clicking OK; the move then takes place without any downtime or sweat on your part—wouldn't that ability be useful for a variety of needs and tasks every day?

VMware Cluster

A VMware cluster allows you to pool the resources of several physical hosts and create logical and physical boundaries within a virtual infrastructure or datacenter. Some organizations may want to create several clusters in their vCenter (formerly VMware Virtual Center) based on functionality—for example, a demilitarized zone (DMZ) cluster or an application cluster. You may want to create a VMware cluster based on the type of LUNs, their speed, maybe their size, or the type of appliance they represent—for example, EMC DMX versus Left Hand Networks. Networking teams may not always want to present all networks to all hosts in the cluster. By creating pools of resources, you can manage these assets and work may be performed on individual clusters rather than the entire infrastructure.

What kind of resources are pooled? CPU, memory, networking bandwidth, storage, and physical hosts are all shared by the VMs that are defined on the specific cluster. A good rule of thumb is to have sufficient capacity to run extra VMs in the event that one or more ESX hosts go down. For example, a cluster with three hosts that runs at 50 percent of resources on each host could probably handle one host failure; the remaining two hosts would then take on 25 percent of the load from the failed host and the cluster would be running at 75 percent of resources. In this scenario, a second host failure would overwhelm the remaining host and some virtual machines would not be able to start on the last host. Obviously, you should plan for extra capacity when designing clusters if failover is important.

Distributed Resource Scheduler

Another feature introduced with VI3 is DRS, which helps you load-balance workloads across a VMware cluster. Advanced algorithms constantly analyze the cluster environment and even use VMotion to move a running server or VM from one host to another without any downtime. You can specify that DRS perform these actions automatically. Say, for instance, that a VM needs more CPU or memory and the host

it is running on lacks those resources. With the automatic settings you specify, DRS will use VMotion to move the VM to another host that has more resources available. DRS can be set to automatically make needed adjustments any time of the day or night or to issue recommendations instead. Two circumstances that often trigger such events are when an Active Directory server is used a lot in the morning for logins and when backups are run. A DRS-enabled cluster shares all the CPU and memory bandwidth as one unified unit for the VMs to use.

DRS is extremely important because in the past VMware administrators had to do their best to analyze the needs of their VMs, often without a lot of quantitative information. DRS changed the way the virtualization game was played and revolutionized the datacenter; you can load VMs onto a cluster and the technology will sort out all the variables in real time and make necessary adjustments. DRS is easy to use, and many administrators boast about how many VMotions their environments have completed since inception.

For example, let's say an admin virtualizes a Microsoft Exchange server, a SQL Server, an Active Directory server, and a couple of heavily used application servers and puts all of them on one host in a cluster. The week before, another admin virtualized several older NT servers that were very lightweight; because those servers didn't use very much CPU, memory, network, or disk input/output (I/O), the admin put those servers on another host. At this point, the two hosts are off balance based on their workloads: one host has too little to do because its servers have low utilization, and the other host is getting killed with heavily used applications. Before DRS, a third admin would have to look at all the servers running on these two hosts and determine how to distribute the VMs evenly across the hosts. On average, administrators would have to use a bit of ingenuity—along with trial and error—to figure out how to balance the needs of each server with the underlying hardware. DRS analyzes these needs and moves VMs when they need more resources so that you can attend to other, more pressing issues.

High Availability

When CIOs and management types begin learning about virtualization, one of their most common fears is "putting all their eggs in one basket." "If all our servers are on one host, what happens if that host fails?" This is a smart question to ask, and one that VMware prepared for when they revealed The HA, or High Availability, feature of VI3. A virtual infrastructure is controlled by vCenter, which is aware of all the hosts that

are in its control and all the VMs that are on those hosts. Each host has a "heartbeat," and if that heartbeat is lost, vCenter believes the host is down and tasks other hosts with restarting the VMs that were affected. VMware recommends a strategy referred to as N+1. This simply means that your cluster should include enough hosts (N) so that if one fails there is enough capacity to restart the VMs on the other host(s). Shared storage among the hosts is a requirement of HA. When a host fails and HA starts, there is a small window of downtime, roughly the same amount that you might expect from a reboot. If the organization has alerting software, a page or email message might be sent indicating a problem, but at other times, this happens so quickly that no alerts are triggered. The goal of virtualization is to keep the uptime of production servers high: hosts can go down, but if servers keep running, you can address the challenge during business hours.

NOTE VMotion is not utilized in a High Availability scenario.

VMware vCenter Converter

If your organization is new to virtualization, VMware vCenter Converter is handy. It's a plug-in to vCenter in the Enterprise version, but there is also a free stand-alone download that lets you convert physical servers into the virtual infrastructure without downtime—thanks to a technology that enables incremental changes to be captured during the physical to virtual (P2V) conversion process. This application works extremely well, and you can use it to convert a single server or multiple servers, move a VM from a workstation or another virtual infrastructure, resize hard drives, or work with partitions. Organizations are choosing virtualization because they have the flexibility to convert already-built working servers, or if the need exists, rebuild a server from a known good build. Both Windows and Linux servers can be virtualized, and there is interoperability with other third-party formats, such as Norton Ghost, Acronis, and Windows Virtual PC (a feature of Windows 7).

VMware vCenter Update Manager

A second plug-in that has proved to be invaluable is VMware vCenter Update Manager. This feature allows for a baseline creation that represents a security standard. A baseline, for example, would be one host or virtual machine that has been configured to be the golden image; it has all

the right patches and all other hosts or VMs should have this level of configuration. You can then apply this baseline to all hosts or select Microsoft and Linux virtual machines, and the technology will remediate updates and apply them to the infrastructure, thus saving you valuable time. The technology will automatically place one host in a cluster in maintenance mode, migrate the VMs to another host, update the host, reboot, exit maintenance mode, and move to the next host to continue the process. You can remediate one host at a time to achieve a fine level of control in environments or organizations that have high visibility or special needs, or you can remediate an entire cluster and sit back to watch it happen.

Another outstanding feature of VMware vCenter Update Manager is its ability to patch offline virtual machines. Obviously not possible with physical servers, this feature offers a level of security compliance far superior to datacenters without virtual infrastructure.

Although this may not sound like a breakthrough in technology, in the old days administrators would have to go to the VMware site; download several patches with long, cryptic names; copy the patches to each host; clear off the virtual machines; open the command line; run an even more cryptic command to work with each file on each ESX host; reboot—and do the same thing over on each and every host. Now with the click of a mouse, VMware vCenter Update Manager does all these steps quickly and efficiently. Furthermore, since a baseline is utilized, each host receives the exact same build. No longer will you need to worry whether you applied every patch to every host; the technology handles this task for you.

VMware Capacity Planner

A third plug-in we'll look at is VMware Capacity Planner. When tasked with virtualizing a physical datacenter, this tool enables you to gather quantitative data from physical servers to better understand which servers are the best candidates for virtualization. Of course, the underlying premise of virtualizing is that on average, most physical servers use significantly less than 10 percent of the resources available on a server; it only makes sense to take a handful of those servers and place them on a host to achieve better utilization of the company's hardware. VMware Capacity Planner analyzes over time how much CPU, memory, disk I/O, and network bandwidth a server uses, and how much it isn't using. As any tenured VMware administrator will attest, VMs need less CPU and less memory than their physical counterparts. Capacity Planner is the quantitative friend in the virtualization journey.

VMware Consolidated Backup

A virtual infrastructure would not be complete without a backup solution, and so VMware introduced its Consolidated Backup feature back in 2006. The traditional approach to backing up virtual machines (agent-based backups) utilizes a lot of system resources from a host; Consolidated Backup relieves that unnecessary pressure by providing a centralized and agent-free process to back up virtual machines. A separate physical server or servers that have access to the shared storage will pull the data through their own network instead of pulling the data through the ESX host, thereby reducing the load on the ESX/ESXi servers and allowing the host to provide its resources for performance instead of backups.

NOTE ESXi hosts are both bare-metal enterprise-class hypervisors that function as hosts for virtual machines.

In earlier releases the backup tool did not enjoyed widespread adoption. Backing up was done mostly via the command line, and several physical servers were required to back up the virtual infrastructure. In addition, the backup tool wasn't terribly easy to use for file-level restores. However, VMware Consolidated Backup can be integrated with backup tools and technologies that already exist as part of your organization's datacenter. Full and incremental file level backups are much easier to perform in this release.

vSphere Client

One of many areas where VMware excels over its competitors is management features. There are several ways to interact with a VMware infrastructure, chief among them vSphere Client. (The other methods, vSphere Web Access and VMware Service Console [vSphere's command-line interface], are discussed in upcoming sections.) You can use vSphere Client to connect directly to vCenter or directly to a host; however, we recommend that you use vCenter as the central administrative unit for the infrastructure. vSphere Client is installed on a Windows machine and with it you can do the following:

- Configure vCenter
- Create virtual machines
- Monitor, manage, and adjust settings for hosts, VMs, and vCenter

vSphere Client is not a tool for end users; it's intended for VMware administrators only. As tools go, this is a great one. The user interface is intuitive and the features are easy to navigate. You can open up your favorite browser, enter the server name of vCenter or any ESX hostname, and download the client.

VMware Service Console

The command-line interface, or VMware Service Console, is an interface used to configure, manage, and monitor an ESX host from a granular level. The Service Console is in essence the first virtual machine on an ESX, and it serves as a communication device between the administrator and the hypervisor. ESXi uses a similar command line called the vSphere Command-Line Interface. For all intents and purposes it performs all the same functions that the Service Console does. An interesting recent development is that PowerShell is starting to be used to configure and manage the infrastructure as well, and this will open up many powerful opportunities for scripting.

NOTE A hypervisor is a high speed scheduler. It hands out resources (CPU, memory, network, disk) to the virtual machines asking for them, very quickly.

---

vSphere Web Access

The final interface is the vSphere Web Access; this access method allows for basic VM management and configuration and is often given to end users that need access to their respective virtual machines but not beyond their sphere of influence. This method is useful because there is a limit on how many vCenter clients can connect to vCenter without impacting performance.

Understand the New Features of vSphere

The list of new features is a lot longer than what we will introduce here, but the following sections detail some of the most exciting ones.

64 Bit

First and foremost, ESX has gone to 64-bit technology for both the Service Console and the VMkernel. Because of this, an ESX host can now handle greater workloads: 64 logical CPUs and 320 VMs per host, with a total of up to 512 virtual CPUs per host and 1TB of host memory. The difference between 32 and 64 bit is that with 64 bit you can achieve higher consolidation ratios (VMs per host) and a better return on investment (ROI) on your hardware. Most likely, only the largest organizations will approach those top ends, but either way, the infrastructure just became more robust.

The downside is that many organizations will have no choice but to purchase new hardware for their infrastructure. With prior versions of VMware, organizations had the option of virtualizing some of their newer servers, and then turning around and using that hardware as their next ESX host. This is still a possibility; however, since 64-bit servers are relatively new compared to their 32-bit counterparts, the 64-bit variety has not yet saturated datacenters to the same degree.

Distributed Power Management

Distributed Power Management (DPM) is the ability of the system to identify when there is enough extra capacity to either automatically shut down hosts or make recommendations to reduce power consumption (think holidays, evenings, and weekends!).

You don't have to be "green" to appreciate this feature. Studies by many industry analysts are available that point out how fast energy costs have gone up in the past few years, and those costs now account for a significant portion of operating costs. The ability to reduce unneeded capacity during the course of a fiscal year can add up to a bigger bonus at year-end. And as any business student will point out, cutting costs adds directly to the bottom line.

VMware states on their website that power and cooling costs can be cut by up to 20 percent in the datacenter during low utilization time periods. How this works is simple: the technology utilizes Distributed Resource Scheduler (DRS) to accomplish the task. During a weekend or holiday, vCenter recognizes extra CPU capacity in a cluster and uses DRS to migrate VMs off a designated ESX host. Once all systems are off the host, that host can be powered off or put in standby mode to conserve energy and lower costs. If the need for capacity starts to

increase, that host will be powered back up and VMs will migrate back onto it to take advantage of all cluster resources.

Enhanced VMotion Compatibility

Although VMotion is not new to vSphere, there is a new feature called Enhanced VMotion Compatibility (EVC) that will add more flexibility when you are configuring VMotion between the same chip manufacturer. As noted earlier, VMotion is not always compatible between older and newer CPU generations. However, the EVC feature allows the hypervisor to mask or hide certain differences (CPU instruction sets) so that compatibility between generations is more relaxed—and this works for both Intel and AMD.

Enhanced Storage VMotion

vSphere includes a new feature called Enhanced Storage VMotion (ESV) that works across iSCSI, Fibre Channel, and Network File System (NFS). ESV can also be used to:

- Convert from thick disk formats to thin virtual disk formats during a move.
- Reduce CPU and memory usage of the host.
- Provide a more efficient block copy method from prior versions.

NOTE A thin disk is presented to a VM as the total size, but it only consumes what is needed by the server, thus saving disk resources.

Storage VMotions are faster, take fewer resources to perform, and have fewer problems completing tasks compared to the earlier version.

Why is ESV important? Storage area network (SAN) storage is expensive and administrators often face a need to dole out small amounts of resources—in this case shared storage—while at the same time not running out of disk space inside the operating system. ESV addresses both sides of the equation. Now you can overprovision storage space by using thin disks. ESV allows you to convert the disk without downtime to the thin format and then use the overprovision feature. If most of your organization's VMs have free space on their hard drives, you can reclaim valuable shared storage with these two technologies.

As you can imagine, this new feature requires additional viewpoints into the datastores so that overprovisioning will not hurt your infrastructure. To that point VMware has introduced several new events and alerting within vCenter to address this new need.

Enhanced High Availability

While VMware High Availability is not a new feature, vSphere adds many enhancements to make it better, smarter, and more efficient than ever before. As we noted earlier, HA is a feature that guards against hardware failure in a VMware cluster. If the hardware goes down, vCenter recognizes this and restarts VMs on the remaining hosts in that cluster.

What has been enhanced in this version? vSphere High Availability now guards against operating system failures as well as hardware failures, and it does so through heartbeats detected via VMware tools. As an added benefit, if DRS is also enabled on a cluster, vCenter will restart VMs, choosing which host to restart a VM based on its resource needs. In the older version, VMs were simply thrown at the remaining hosts and restarted without taking resource needs into consideration; clearly we are moving forward.

vCenter does a much better job of positively detecting high availability events rather than false positives because vCenter monitors the heartbeats of multiple VMs on each ESX host (versus just the Service Console in the older version). One of many reasons administrators will love this feature is because it singlehandedly eliminates the need for them to rush into the office at night, during holidays, or any other time hardware has an issue; because everything keeps running, the problem most likely can be dealt with during regular business hours or at the next convenient opportunity.

Another major reason VMware High Availability will be one of your favorite features is because this technology rivals physical clustering—all it takes to set up High Availability is a single mouse click and all your VMs and applications in that cluster will be protected. There's no complicated process of setting up duplicates of hardware—it's ready to go on the spot!

Virtual Machine Scalability

Virtual machine scalability has been increased. There's now eight-way virtual symmetric multiprocessing (SMP) for even more demanding workloads; 256GB of RAM can be assigned to VMs; and more support

is available for hot-pluggable devices such as memory and virtual CPUs. All this is accomplished with new virtual hardware, which diminishes downtime for compliant guest operating systems. If a VM is running low on hard drive space, you will be able to "extend" the hard drives on a running VM instead of having to shut down the VM and use other tools to extend the hard drive.

However, the caveat is that the operating system must support these features (hot add memory or hard drives). You will also have the ability to virtualize bigger or more important systems, and organizations will do this to get the added flexibility for their most important applications. The organization I work for is now virtualizing Microsoft Exchange because we like how easy it is to be able to recover VMs as opposed to recovering a physical Microsoft Exchange server. The flexibility to add hard drive space or memory with only a few clicks of a mouse, fault tolerance, and HA clusters—all of these reasons are driving organizations to virtualize more and more of their physical infrastructure.

vCenter Improvements

The management platform has been significantly expanded in vSphere:

- Licensing is no longer done via a licensing server; it is all included in a 25-character license key.
- Performance charts have been revised so that you can view all metrics in the same window.
- Enhancements have been made to alarms and notifications, including storage awareness enhancements.
- A resource usage statistics feature has been introduced for VMs, and at the resource pool layer, think chargeback for utilized resources! (For example, if your organization charges for servers or resources used, this feature will help you out.)
- The new vSphere Management Assistant functions as a VM and includes software that you can use via the command line to run scripts and agents against ESX and ESXi hosts.

The licensing change is a special gift from VMware since the old licensing server in version 3 was sometimes challenging. Usually after it was working everything went well, but many administrators, myself included, pulled out some of their hair during that first install.

The ability to see all four main metrics—memory, CPU, disk I/O, and network utilization—in the same window in performance graphs may not sound worthy of mention; however, due to the nature of virtualization, troubleshooting is greatly simplified when an administrator can see them all at the same time. You can quickly assess what is and is not happening on a virtual machine or ESX host. When viewing these four metrics, for example, you can see that the issue is an isolated memory issue, or that CPU is out of whack and the network is not doing anything. And since these graphs can be viewed at different time intervals, you can pinpoint when the problem started.

Fault Tolerance

Fault tolerance is one of vSphere's exciting new features, and it promises both zero downtime and zero data loss from hardware-based failures. The technology is able to accomplish this task using the playback feature that was perfected in VMware Workstation (a desktop virtualization product). A second VM on a different host has every action from the original VM replicated, so setup and maintenance are easier compared to a physical cluster. If hardware on one host fails, the other VM takes over on the spot—there's no need for HA (to restart a VM on another host) because the VM is already running live on a separate host. Additionally, the technology is savvy enough to spawn another fault-tolerant VM after an event like this, so the redundancy can be maintained after a failure automatically.

This is not a feature that you will want to use for all VMs, however, because it will require twice the amount of hardware. For example, if an organization wants to provide fault tolerance for all its existing virtual infrastructure, it will need to acquire twice the hardware to match what it currently runs on. Obviously this technology should be reserved for your most important servers.

vNetwork Distributed Switch

VMware has created vNetwork Distributed Switch (DVS) so that you can set up networking once and then everything else becomes plug and play. DVS essentially spans multiple ESX hosts so there is less need to set up each and every host with identical virtual networking. This would be similar to downloading a configuration file for a new router—downloading it, running, and saving the changes instead of

reconfiguring settings every time. There will also be support for third-party virtual switches, the ability to track VMs as they migrate around a cluster, and better support for private virtual local area networks (VLANs) into existing networking environments.

The goal of this feature (and host profiles, discussed next) is to eliminate extra work and to make ESX hosts more like devices and less like operating systems. You can simply plug in a device and associate it in vCenter with a cluster; then all the settings that need to be specified are done automatically. Under those conditions, you could add and subtract additional CPU and memory (do you even need to think of it as a host at that point?) anywhere in the infrastructure.

Host Profiles

Host profiles are similar in notion to a template or a golden image used to consistently replicate new desktops or virtual machines. Prior to this new feature, you either rebuilt each host from scratch or used some kind of automated build process and then did your best to create consistency across all hosts. Host profiles greatly reduce configuration management by allowing you to build a golden image once and then "plug it" into any new hosts, thus ensuring standards across the infrastructure.

Every organization and administrator chooses how they want to configure their hosts. Some manually build each one; some use scripting to create exact copies. These approaches all have their merits. However, host profiles will allow an administrator to create a golden image and then apply those settings to any new or replaced ESX hosts.

vCenter Linked Mode

vCenter Linked Mode creates a simplified approach to management in large environments by allowing you to use a single interface for multiple vCenter servers. If there is more than one vCenter server in your environment, they can be interconnected in a mode that allows you to share management roles across the infrastructure, licensing, and other related tasks. This yields a further reduction in workloads associated with setting up the same configurations on multiple vCenter servers.

Virtual Disk Thin Provisioning

In the storage realm there is a new feature called *virtual disk thin provisioning*. SAN storage is more expensive than hard drive space; therefore,

administrators are careful to properly provision each new VM with the right amount of GBs on their drives. Virtual disk thin provisioning allows you to overprovision valuable shared storage while at the same time allowing the VM to grow into its allocated hard drives. This technology would not be complete without the underlying reporting and notifications that ensure proper maintenance of the storage occurs, and that is well taken care of on the management side in vCenter Server. This feature will reduce the need for SAN storage, thus helping keep costs low and under control.

Take a moment to imagine a virtual infrastructure that has, say, 255 virtual machines in version 3.5, update 2. If on average each VM has between 1 and 5GB of unused space on just the C drive, that means between 255GB and 1.2TB of SAN storage is unused. If each VM has two hard drives, that could mean almost 2TB of space could be reclaimed if this feature is utilized; now that is valuable.

VMkernel Protection

The new VMkernel Protection technology helps protect the hypervisor by ensuring that the integrity of the VMkernel is not compromised and/or changed by either common attacks or software loaded on the host. The VMkernel modules are now digitally signed and validated during each reboot so that nothing is overwritten, and they use memory integrity for protection from buffer overflow. When you combine this technology with VMware VMsafe (which is used to protect VMs by including an Application Programming Interface (API) for third-party developers to create security products), you'll see that security has been enhanced yet again.

VMkernel Protection is somewhat similar to what Microsoft did to try to eliminate the "Blue Screen of Death": they created digitally signed device drivers. Before this, third-party vendors created all sorts of software that interacted with Windows operating systems, and sometimes that software was coded well and played nicely; other times it blue-screened the operating system. Microsoft did not have control over outside companies, so they did the next best thing by introducing digitally signed drivers. In a similar manner, VMkernel Protection ensures that the kernel is not modified; VMware can thus ensure that their platform is stable and will most likely remain that way.

VMware vShield Zones

VMware vShield Zones is a virtual appliance that acts as a firewall. It can be enabled to monitor, log, block, and allow inter-virtual machine traffic within an ESX host or between hosts in a cluster without having to direct traffic externally through physical devices like routers or switches. If you need to isolate, bridge, or firewall VMs between multiple zones, you can do so with vShield Zones. Successful and denied activities can be logged, graphed, and analyzed.

VMware Data Recovery

vCenter now comes with an easy-to-use plug-in that allows for disk-based de-duplication of all your virtual machines. This technology runs on a VM in either ESX or ESXi and allows you to restore individual files or entire VM images. A wizard-driven interface is used to configure and schedule backup jobs, and multiple restore points can be easily selected for specific points-in-time copies or restores. The technology uses de-duplication to significantly reduce the size of backups while allowing for multiple incremental changes; it even supports the Volume Shadow Copy service to enable backup of Microsoft applications.

NOTE De-duplication is the process of identifying redundant pieces of data and essentially compressing it to save space.

This feature is a huge win for any business, whether it is big or small, as it will have the ability to back up applications and/or servers with relative ease as well as restore those applications and servers using intuitive wizards.