Chapter 1

Exploring VMware Infrastructure 3 as Your Virtual Solution

In This Chapter

- How virtualization works
- Benefits of virtualization
- Parts of VMware Infrastructure 3
- Planning your virtualization

VMware Infrastructure 3 (VI3) is a robust, feature-rich, fault-tolerant, and highly reliable platform for virtualization. In fact, VMware created x86-based virtualization. (Now that I got that out of my system, it's time to explore the possibilities of saving time and money with your virtualization project.)

Most x86 computers don't use their hardware to the fullest capacity because *de facto* system design often dictates that you use a single server for a single purpose. As I mention in the book's Introduction, virtualization simply uses "smoke and mirrors" to separate your programs and operating systems from the hardware on which they run. That way, many virtual machines use common hardware, and the hardware is far better utilized.

In this chapter, virtual machines, ESX hosts, the benefits of virtulization, and VI3 are covered. Additionally, the last section describes the major steps in implementing your virtual infrastructure.

Knowing What You Must About Virtual Machines

Making as many of your physical machines as possible become virtual is the end game of virtualization. The more you *virtualize* (convert physical to virtual machines), the more benefits you see, and the more you realize how

reliable and stable a platform VMware Infrastructure 3 is. Being skeptical by nature, I went slowly at first, but I quickly began to trust VMware the more I worked with it.

Virtual machines: The non-physical workhorses

For all intents and purposes, virtual machines are just like physical machines. You can log on to them; and they have BIOS, hard disks, memory, CPUs, operating systems, and applications. In fact, if you connect remotely to a machine, you'll never know that it's virtual unless someone tells you. Virtual machines work and behave just like physical machines. Even the machines themselves don't even know they are virtual!

Virtual machine pluses and minuses

Aside from all the general benefits of virtualization (covered in the "Benefitting from VMware Infrastructure 3" section, later in this chapter), here are the major pluses for using virtual machines:

- **They can be rebooted much faster than physical machines.** I was able to reboot one server in 15 seconds!
- **They are more simpler than their physical counterparts.** For instance, there are no array controllers to configure in the virtual machines.
- **They are easy to back up and restore.** The entire machine is stored in a set of files.

Unfortunately, there are also downsides. However, you can easily minimize the downsides:

- **Support can be a gray area.** Some software vendors might tell you, "Hey, it's on a virtual machine. We don't support that." Pay for VMware support and know your own operating systems to mitigate this risk.
- **Troubleshooting can be a little tricky.** People troubleshooting a problem on a virtual machine might incorrectly deduce the problem is caused because the machine is virtual. As long as you hone your own troubleshooting skills, this risk is minimized.

  I had a consultant tell me that the most likely problem with a report running slowly and having connectivity issues was because the server was virtual. I had to prove to her that it was the report — not the virtual machine — before we could resolve the issue.

- **Inherent flaws are transferable.** Flaws in a virtual infrastructure design can affect all virtual machines. This, in turn, will affect all users of those machines. Watch where you make your trade-offs, and design your system with ample capacity to avoid this risk.
- **Some apps can be troublesome for time syncs.** For example, I have one application that throws off the virtual machine's time sync. You can compensate for this by synching the virtual machine's time to the *ESX (Engagement Simulation Exercise)* host. (An *ESX host* is a server that your virtual machines run on. It provides access to all the hardware resources your virtual machines share.) In turn, sync your ESX host to your network's time source.

Symmetrical multiprocessing and why you care

Commonly, physical machines use multiple processors. In fact, you'd be hard pressed to find a server that doesn't come with at least a dual-core processor. Each core is treated as a separate CPU, so a machine using a single, dual-core CPU is taking advantage of *symmetrical multiprocessing (SMP).* In the physical world, multiple CPUs can greatly aid in processing speed. Things are a little different in the virtual world, however.

Your ESX host will most likely have multiple processors with multiple cores. Whenever your virtual machine needs the CPU, the VMkernel (covered in the next section) can send the work to any CPU in your ESX host. Your single CPU virtual machine is, therefore, getting benefits similar to SMP without even knowing it.

If you have a license for SMP, you can assign multiple processors to a virtual machine. However, just because you *can* do something doesn't mean that you *should.* Dual, virtual CPU machines force co-scheduling of physical processors.

With *co-scheduling,* if one physical processor is scheduled, a second one is as well. This can take resources away from your other virtual machines. Additionally, if Process 1 on virtual CPU1 is waiting too long for Process 2 on virtual CPU2 to finish, both processes might get *scheduled out* (finish their allotted share of CPU time and lose the processor until their next turn) before completion. This can negate the benefits of using multiple CPUs.

Best practices dictate adding multiple CPUs to a virtual machine *only* if you can prove an increase in performance. This is because multiple virtual CPUs can have some negative side effects:

- **Potential performance hit:** If you give a virtual machine multiple CPUs, you remove some of the scheduling flexibility available to the VMkernel. This trade-off might (no guarantee) give a few machines a performance boost, but at the cost of your other virtual machines losing performance. If you really need one machine to have a performance boost, you might want to try using a resource pool instead of SMP. Resource pools are covered in Chapter 14.
- **Prevent processor fragmentation:** If you have two, single-core processors or two, dual-core processors, you can create a processor fragmentation scenario using SMP. If you create a dual-CPU virtual machine, it might experience *processor starvation.*

  Say a single CPU virtual machine is running on one processor, and the dual CPU machine has two active threads ready to run. Both threads need to be co-scheduled, but one physical CPU is in use, so neither thread is scheduled. Meanwhile, only one physical CPU is in use, and the other is just sitting there: doing nothing. Enabling *hyperthreading* (making a single processor appear as two processors to an SMP-aware operating system to make the CPU run more efficiently) allows a single-core CPU to act like two physical CPUs for a virtual machine. Enabling it alleviates the processor fragmentation problem in systems that have more than one processor.

Understanding the Role That VMkernel Plays

Even though each of your virtual machines thinks that it has its own dedicated hardware, the machines actually share a common pool of hardware. The magician creating this illusion is the VMkernel that runs on each ESX host. The simplest way to think of the *VMkernel* is as a scheduler: It schedules virtual machines access to resources. The VMkernel even schedules the management console that you use to configure and operate your ESX hosts. The VMkernel virtualization model is shown in Figure 1-1.

As shown in Figure 1-1, there are three layers to the VMkernel virtualization model:

- **Virtual Machine layer:** Here all your operating systems and applications are neatly housed in separate virtual machines. Each virtual machine thinks it is a physical machine with its own hardware, but this is not the case. The ESX tricks all the virtual machines.
- **ESX layer:** The VMkernel is the boss of this layer. The kernel schedules hardware for the virtual machines as well as the management interface. The VMkernel is a liar. It tells all the virtual machines that they are physical entities and have access to their own hardware.

- **Shared Resources layer:** This layer consists of all your hardware subsystems. It includes your physical and storage area networks.

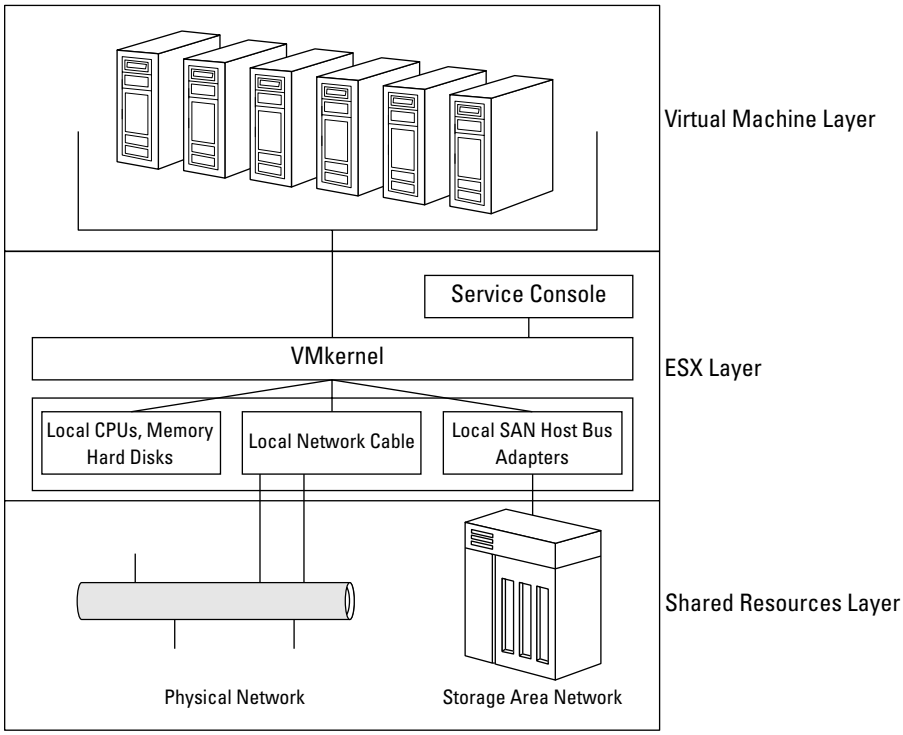

Figure 1-1: The VMkernel virtualization model.

VMware Infrastructure 3

Using the VMware Infrastructure product suite enables you to virtualize servers, storage, and networks. In addition, the suite offers you ways to add extreme fault tolerance as well as centralized management, load leveling, and centralized backup.

VMware always offers several different ways to purchase their products. Below are the offerings for version 3.5 (the offerings for 3.0 were entirely different):

- **Single ESX:** You can buy the plain ESX that you install on a hard drive or ESXi, which has the operating system on a chip instead of a hard drive.
- **VMware Infrastructure Foundation:** This is ESX with a few bells and whistles and offers you automated updates, a VirtualCenter Agent, and Update Manager.

- **VMware Infrastructure Standard:** This is ESX with more bells and whistles and offers everything in the Foundation level plus High Availability. This allows you to essentially create an active-passive cluster and would be the minimum level of fault tolerance that is acceptable for production systems.
- **VMware Infrastructure Enterprise:** Every available option — If you want it all, then this is for you. This package provides every possible feature to enable active-active fault tolerance and dynamic load balancing across servers. If you are virtualizing datacenters, you want the Enterprise package.

Benefitting from VMware Infrastructure 3

Virtualization simply makes life easier from a technical and administrative viewpoint. Fortunately, it also makes life easier from an economic viewpoint, so everyone can agree that virtualization is a good thing. Before you decide to virtualize, consider some of its many benefits:

- **Better hardware usage rates:** This translates to needing less hardware to do the same amount of work.
- **Lower hardware-maintenance costs:** You need fewer physical servers, which means less maintenance contracts to pay for.
- **Lower cooling costs:** Less heat is generated, so less cooling is needed.
- **Lower electric costs:** You have fewer physical servers so your electric bill drops. VMware is a very green technology.
- **Lower space costs:** Your server room can be much smaller, which leaves more room for offices. And face it, it is the people in the office that produce your company's income. The server room is an expense that you can help minimize.
- **Longer infrastructure run time from UPSes:** If you virtualize all your physical servers and keep the same UPS system, think how much longer it can run during an outage — especially if you are condensing twenty or thirty virtual machines onto one physical machine.
- **Faster server deployment:** You can deploy a new server in as little time as about 15 minutes. And you don't need to spec-out hardware and wait for delivery. In fact, deploying a server from a template is as easy as right-clicking and answering a few simple questions.
- **Simplified management:** All your virtual servers use the same drivers, and servers are just a collection of files on a hard disk. Whether you're installing a program or adding virtual hardware to a server, all servers are managed through the same client.

I've added "hardware" to a virtual server in New Jersey remotely from a beach in Sarasota. Nice!

- **Easy backup and fast recovery:** Again, your servers are just a bunch of files.
- **The ability to freeze your server in time through snapshots:** You can take a snapshot before applying a Service Pack. If you have problems after the update, you can go back in time to before the service pack was applied instead of rebuilding your server and restoring your data. Fixing a bad update takes only minutes instead of hours.
- **Quality-of-life improvement:** All the time and effort saved makes your IT life much better! All the money you save makes management extremely happy as well. Everyone benefits from virtualization.

After reading this list of benefits, you likely think that using virtualization can prevent many IT headaches — and you're right! Time to meet the components of VI3.

After you start to virtualize machines, it can become very addictive. You might even get the overwhelming urge to create far more virtual machines than you actually need. You should resist that urge! In fact, this is called *virtual machine sprawl.* While it does not take up as much space as physical server sprawl, it can be detrimental from an efficiency, resource, and management point of view. As a rule of thumb, only create a virtual machine if you would have created a physical machine to accomplish your goal in the past.

Meeting the pieces and parts of VMware Infrastructure 3

Many pieces make up VI3, and each has a specific purpose to help create a seamless whole. Although you can purchase different parts separately, buying them as a package costs less. If you're virtualizing your infrastructure, you will want the entire product suite.

Here's a list of what's included in the VI3Enterprise Suite:

- **VMware ESX:** This comprises the operating system that you put on your server hardware that allows you to create virtual machines and share hardware resources between them. Your physical servers are referred to as *hosts.* The virtual machines run *guest operating systems.*

  A new version — ESXi — is also available. This preinstalled version can be configured by non-technical people via menus at boot-up. This version supports everything that the ESX supports, but it lacks a Service Console. This version is designed for remote deployment and management. And did I mention that it runs on a chip? You don't even need any hard disks in your ESXi server.

- **VMware Virtual SMP:** *Virtual SMP (symmetric multiprocessing)* enables a virtual machine to use up to four physical processors simultaneously. To benefit from multiple virtual CPUs, your operating system and application need to support SMP. However, VMware is very good at scheduling resources and you should only use multiple virtual processors if you can prove a performance increase.
- **VMware VMFS:** Virtual Machine File System is a file system that allows multiple ESX hosts to access the same data storage concurrently. This allows any host to run any virtual machine and provides the ability to switch between hosts on the fly using VMotion.
- **VMware VMotion:** *VMotion* is the resource that actually moves running machines from one host to another with no loss of connectivity. In version 3.0.x, if you shut down a virtual machine, you can also change where its files are stored.

- **VMware Storage VMotion:** While VMotion allows you to move your virtual machine from on ESX to another to better utilize hardware resources, Storage VMotion allows you to move the virtual machine files from one storage location to another to better utilize storage resources. The virtual machine stays on the same ESX host while its files are moved to a new location. Again, this can be done while the virtual machine is running.
- **VMware Distributed Resource Scheduler (DRS):** Use *DRS* monitors your resources and decides which host is best to run a virtual machine on. It provides system wide load leveling. DRS uses VMotion to move virtual machines off hosts that are under heavy loads and onto hosts that have more resources available.
- **VMware High Availability (HA):** This high-availability resource can restart virtual machines on a new host if the host on which they were running fails.
- **VMware Consolidated Backup (VCB):** A centralized way to backup virtual machines through a backup proxy server.
- **VMware Update Manager:** This feature is designed to manage patches for your ESX hosts as well as the guest operating systems running on them.
- **Distributed Power Management (DPM):** DPM attempts to consolidate virtual machines onto the least number of hosts in a cluster as possible so the remaining hosts can sleep and save power. The hosts automatically wake up again if they are needed later.
- **VMware VirtualCenter:** *VirtualCenter* (VC) is a centralized management framework that lets you create fault tolerant clusters (clusters are covered in Chapter 12). It controls HA, DRS, and VMotion for an entire cluster. Additionally, VirtualCenter provides one spot for you to configure all hosts and virtual machines in your virtual infrastructure.

You need VirtualCenter to take advantage of load-leveling and fault tolerance. VirtualCenter is a separate purchase and is not included with VI3.

Two other useful products fall outside the realm of VI3, but you might want them anyway.

- **VMware Converter:** You will use this product over and over again. With it, you can easily convert physical Windows machines into virtual machines. Additionally, if you use Virtual Consolidated Backup to back up images of your virtual machines, you can restore those images by using VMware Converter. This is also the product to use to convert between various virtual machine formats from VMware and third-party products.

  The two versions of VMware Converter are

  - *Starter Edition:* Use this free version to convert physical machines into virtual machines on an ESX from the physical machine itself while it's running. It also enables you to convert physical machines to other VMware formats remotely.
  - *Enterprise Edition:* Use this version to convert physical machines to virtual machines remotely and run multiple conversions simultaneously. You can also schedule migrations if you want to run them off-hours. Another benefit is migrating a machine that is powered off, which allows you to create a boot CD for cold cloning. The cold-clone command line interface (CLI) can be used to convert SUSE and Red Hat Linux machines.
- **Capacity Planner:** This product is designed for consultants to quickly gather comprehensive data about your IT infrastructure. It is a hosted application service that gathers data without the use of agents, and it's designed to get all the information you need to correctly design your virtual environment.

Planning Your VMware Solution

After you decide to virtualize, there are four distinct stages to bringing your plan to fruition. You start by defining your capacity requirements and figuring your return on investment (ROI). Then it's time to design and build your virtual infrastructure. Then, after you virtualize your physical machines, back up your hard work and investment.

Stage 1: Capacity planning and return on investment

Start off by determining how much capacity you need today as well as how much growth you anticipate. Today's needs, obviously, determine what hardware you have to buy to get started. Your anticipated growth needs determine how much you will save by not purchasing physical servers.

Even without taking the electricity, cooling, and space savings into consideration, you can usually find a ROI just by the savings generated through not buying future servers and their associated hardware maintenance contracts. For example, I designed my system with a minimum virtual-to-physical ratio of 15:1. As I need more capacity, I can add a single physical machine and build at least 15 more virtual servers.

Think about the math behind that for a moment: You can spend $12,000–$15,000 on each server 15 times — or just once. Additionally, you can spend roughly $500–$1,000 per year for maintenance contracts 15 times — or once. Say that you're going to roll out 15 servers ($12,000 each for a physical server) in a year; also assume each server costs $500 per year for a maintenance contract. And, say that the expected lifetime is five years. Look at the following three options:

Plan	*Expenses*	*Total Costs*
Physical infrastructure	Hardware: $12,000 × 15 Maintenance: $500 × 15 × 5	$180,000 $37,500
		$217,500
Adding to an existing virtual infrastructure	Hardware: $12,000 Maintenance: $500 × 5	$12,000 $2,500
		$14,500
Building a two-node virtual infrastructure from scratch	N+1 fault-tolerant server hardware: $12,000 × 2 SAN hardware Maintenance: $500 × 3 × 5 By using N+1 you have enough excess capacity to absorb a single system failure without an outage.	$24,000 $120,000 $7,500 **$151,500**

In these scenarios, regardless of whether you're rolling out a new virtual infrastructure or expanding an existing one, it pays to virtualize because of hardware and maintenance alone. And that's not even taking into account easier management, space savings, or the electrical savings from not having to cool and power as many physical servers.

I talk about capacity planning in depth in Chapter 2.

Stage 2: Designing and building your virtual infrastructure

After you know the hardware you need and can show that it more than pays for itself, you need to design your virtual infrastructure. This can include the following:

- Your ESX, VC, and VCB design
- Your virtual networking design
- Your backup strategy
- Possibly your storage area network (SAN) design
- Possibly your IP network design
- Possibly your Windows domain design

If you carefully plan ahead and design your infrastructure well, building it will be a piece of cake. Building consists of connecting the hardware and installing the software. That part is easy!

Stage 3: Virtualizing your physical machines

Converting your physical Windows servers to virtual machines is easy: Just use either the VMware Converter Starter or Enterprise edition. Converting Linux machines takes a little more work: You need to buy the Enterprise edition, make a boot disk, and convert via the CLI. This is currently only supported experimentally.

Converting Linux machines without VMware Converter Enterprise edition requires making a disk image, applying that disk image to a virtual machine disk, booting to a recovery CD, and modifying configuration files to replace the SCSI drivers with Buslogic drivers. Many different ways exist to convert Linux to a virtual machine. Try googling "P2V Linux" or "Converting Linux to virtual" to see what other people have done.

Stage 4: New ways to protect your data

If you back up your physical machine, you want to continue backing it up after it becomes virtual. I have good news for you: The method you use to back up your physical machines will work exactly the same after you virtualize them. However, you now have a new backup option that uses virtual machine snapshots.

Snapshots take a picture of your machine at a moment in time. Any changes to a machine or its data after that point in time can either be applied or discarded. The choice is yours. This creates a lot of possibilities to prevent IT headaches.

You can backup an image of your machine to tape for rapid restore using snapshots. This is how VMware Consolidated Backup works. You can use VMware Converter to restore machines backed up using snapshots. It works really well and greatly speeds disaster recovery.

Additionally, snapshots can be used to safely and easily test an upgrade without the risk of destroying your machine. Simply take a snapshot and then apply your update. If it works, apply the snapshot. If not, discard the snapshot.

Of course, you can also get into trouble with snapshots:

- If your virtual machine is a database server, shut down the database before doing anything with snapshots to prevent database corruption.
- Snapshots can be used in your backup strategy, but they do not replace your backup strategy. Do not accumulate multiple snapshots over a long period of time. They are designed to be used temporarily. This is covered in more detail in Chapter 16.