

1

VMware Infrastructure Installation

Before we can begin serving the main recipes of our cookbook, we have to make sure that all our readers understand the terms we'll use in the rest of the book, and that you have all the necessary components installed. This book aims to be useful for both new and seasoned VMware ESX users. Because of the nature of this cookbook, we will assume from time to time that you have advanced knowledge and understanding of how the products work.

Virtualization provides a way for multiple operating systems to be installed on the same physical hardware. By using virtual technology, we can consolidate hardware and instantly build quality assurance and test environments. This is a tremendous breakthrough, as it allows underutilized equipment to do more than sit around idle, as well as allowing developers and administrators to test and use multiple software configurations and packages that require different operating systems on the same piece of equipment, without having to purchase, set up, and maintain multiple computers. This savings makes the accounting department and managers happy and gives the technology lovers an opportunity to do all the things they want or need to do at the same time.

A typical IT scenario goes something like this: you have lots of servers in your rack or collection of racks. Most of them only run at about 10–15% of their capacity the majority of the time, but you let that happen because you want to keep their various functions and operating systems isolated from one another. This provides some security, both because if one server is compromised it does not necessarily mean that access has been granted to others, and because a problem with one piece of software will not cause other parts of your IT infrastructure to go down as you deal with it.

Virtualization helps you make better use of those physical resources, without compromising the original intent of keeping services isolated. Instead of installing your operating system directly on your hardware, you instead begin with a *virtualization layer*: a stripped-down OS designed to schedule access to network, disk, memory, and CPU resources for guest OSs, the same way that those guest OSs control that scheduling for their applications. Most virtualization platforms limit themselves to specific hardware

and present a specific set of virtual components to the operating systems installed on top of them. This provides a very stable and consistent presentation to the operating systems you install and allows them to be moved much more easily.

Once you've installed and configured the virtualization layer is installed and configured, you can partition the physical hardware and assign it to discrete, separate operating system instances that you install on top of the virtualization layer. These virtual installations operate exactly as they would normally. They are not aware of the presence of other virtual installations that exist on the same hardware. Each acts as if it is installed on a predefined piece of equipment by itself, with the virtualization layer controlling what the virtual installation sees and how it interacts with other equipment outside its control. Basically, the virtual installation looks, feels, acts, and is administered exactly the same as a standard installation from the inside, but may be manipulated and configured easily and alongside others from the outside.

Here is the coolest part: virtualization, as provided by sophisticated systems like VMware, also allows you to pool the resources of several physical machines and then divide them up however you want or need. If you have 10 physical servers, each with four processors, 4GB of RAM, and an 80GB hard drive, you can segment those resources to provide a small and low-powered server for your in-house email, a powerful and high-memory processor for the number crunchers in accounting, multiple load-balanced servers for your web server, a separate server for your database, and so on, each with a configuration of memory, disk space, processor power, and so on specific for its needs. Then, if you discover that one virtual server has more resources than it really needs and another doesn't have enough, you can change the configuration quickly, easily, and without taking your servers offline! You can take a physical server offline for maintenance without losing access to any of your virtual servers and its functions. You can move resources in and out and around your pool as needed, and even automatically. These are the sorts of things we will discuss in this book.

There are many companies that provide powerful and stable virtualization platforms, but we have found VMware's offerings to be wonderfully stable, flexible, easy to set up and maintain, and well supported. We like VMware. If you are reading this book, you probably do too, but chances are you want to make better use of its potential than you are doing now—either that, or you have been asked to set it up or maintain it and you are wondering how to get started. Whatever your reason, if you use VMware at all, we hope you will find this book useful and practical.

What Is VMware Infrastructure 3?

VMware, Inc. is a company headquartered in Palo Alto, California with over 7,500 employees and about 120,000 customers, including 100% of the Fortune 100. In 2007, it had revenues of over \$1.33 billion. VMware is a rapidly growing company that began in 1998 and now has over 20,000 partnerships with companies ranging from somewhat small to extremely large.

VMware Infrastructure 3 (VI3) is easily the most widely used virtualization platform today. It is well tested and has been used in applications ranging from very small, localized installations with just a handful of servers to exceptionally large server farms in major corporations. It is robust, scalable, easy to administer, and flexible. It is also small

and fast, which means the virtual installations running on top of it have more processor power and other resources available to them than they would if they were using some of the more resource-heavy virtualization software available.

Unlike some of the other hosted virtualization products you may be familiar with, including the company's well-known VMware Server, VMware Infrastructure 3 does not require any other operating system. Most virtualization platforms begin with a Linux/Unix, Mac OS X, or Windows platform, install their product on top of it, and then begin segmenting the resources from there. This is how a developer may run a copy of Windows on top of her laptop's base installation of Linux, perhaps using a product like VMware Server, Xen, or VirtualBox. VI3 is designed to be installed on bare metal, as the base operating system. This design choice eliminates a layer of software between the virtual installations and the hardware and results in faster, smoother performance.

The platform is comprised of several major products, including ESX, ESXi, vCenter Server, and vCenter Converter. VMware recently changed the names of its VirtualCenter Server (now vCenter Server) and VMware Infrastructure client (now vCenter client); however, the products themselves haven't been changed to reflect the new names. We will use the new terminology within the book and reference the versions when appropriate. The remainder of this chapter will introduce these key pieces of technology within the virtualization environment.

What Is VMware vSphere 4.0?

In June 2009 VMware launched ESX 4.0, which now falls under the vSphere 4.0 suite of products. vSphere 4.0 encapsulates both vCenter and ESX 4.0, along with other products that we don't discuss in this book. This new release not only brings incremental improvements but takes the virtualization platform to a whole new level by letting you think in terms of complete installations or deployments instead of managing your site server by server. Throughout the book we'll lightly reference ESX 4.0 in relevant chapters to explain the differences between versions 3.x and 4.x. However, it's worth mentioning that ESX 4.0 is a completely new product:

- ESX4 is 64-bit, as compared to the 32-bit versions of ESX 3.x.
- ESX4 supports more hardware and more virtual machines.
- The new version sets the bar for commercial virtualization products even higher, with new features such as fault tolerance, vShield, distributed switches, and much more.

For a complete list of new features, please visit the following website:

http://vmware.com/support/vsphere4/doc/vsp_40_new_feat.html

Here are 10 features we feel are worth highlighting:

vCenter server linked mode

Allows you to connect multiple vCenter servers in a “linked” fashion. Doing this allows them to share licenses and roles and provides them with an “overview” of the entire virtual environment.

Host profiles

Provided in Enterprise Plus to make it easy to push a template of an already existing ESX host to a new server. This feature is a huge benefit, as it means you no longer need to manually configure networks, storage, and so on in your new ESX servers.

Performance, graphs, and events

CPU, memory, storage, and networking graphs are now displayed on one central page. In addition to improved graphs, customers will benefit from improved event messages, alarm settings, and error reporting.

Storage VMotion within vCenter

You can initiate a storage VMotion within the vCenter client. This takes away the need for the command-line management or third-party plug-ins that were popular with the ESX 3.x versions.

Enhanced VMotion

Enhanced VMotion Compatibility (EVC) allows customers to initiate VMotion between Intel FlexMigration and AMD-V Extended Migration technologies. In essence, this means you can now move a running server between hosts using different CPU technologies.

Virtual machine hot add support

You can add CPU and memory resources to an existing virtual machine without having to reboot.

Virtual machine hardware improvements

Virtual machines can now scale up to eight virtual SMP processors and up to 255GB of RAM. vSphere supports IDE devices, Serial Attached SCSI (SAS), and VMXNET Generation 3 network support.

Physical node improvements

Physical hosts can now have up to 1TB of memory, 64 logical CPUs, and 320 virtual machines. The 64-bit architecture of ESX4 allows greater memory capacity, better performance, and seamless hardware support compared to previous releases.

Storage improvements

ESX 4.0 brings a huge set of improvements to the realm of storage within ESX. Thin provisioning for virtual disks enables virtual machines to utilize storage on an as-needed basis, eliminating a major source of wasted space on a storage area network (SAN). The VMFS Volume Grow feature allows you to dynamically grow a Virtual Machine File System without interrupting the running virtual machines. There is also enhanced support for NFS and iSCSI software initiators, which now supply support for jumbo frames on 1GB or 10GB local networks.

Distributed switch

Provided on Enterprise Plus to let customers create a virtual switch for all their ESX hosts to connect to and utilize. This helps reduce network maintenance and allows virtual machines to be moved to any host using VMotion without having to worry about network connections.

As we mentioned earlier, this is only a small subset of the new features available in vSphere 4.0.

VMware ESX 3.x/4.x Configuration Maximums

VMware's vSphere (ESX 4.x) and Virtual Infrastructure (ESX 3.x) products have limits within which they can operate. We feel it is important to include this information so that you have it at your disposal prior to installing ESX/ESXi 3.x/4.x or vCenter. These values can be crucial when planning your virtual environment, and we suggest you read through them to become familiar with the different maximums. We've tried to include all relevant, publicly provided values from VMware, but you may notice a "-" when we could not find a matching reference between the two versions of ESX.

Table 1-1. Virtual machine maximums

Value/Item	Maximum	
	ESX 3.5	ESX 4.0
Number of virtual CPUs per virtual machine	4	8
SCSI adapters per virtual machine	4	4
Devices per SCSI adapter	15	15
SCSI targets per virtual machine	60	60
SCSI disk size	2TB	2TB
Size of RAM per virtual machine	64GB	255GB
Virtual machine swap file size	64GB	255GB
Number of NICs per virtual machine	4	10
Number of IDE devices per virtual machine	4	4
Number of IDE controllers per virtual machine	1	1
Number of floppy devices per virtual machine	2	2
Number of parallel ports per virtual machine	3	3
Number of serial ports per virtual machine	4	4
Size of a virtual machine swap file	65532MB	255GB
Number of remote consoles to a virtual machine	10	40

VMDirectPath PCI/PCIe devices per virtual machine	-	2
VMDirectPath SCSI targets per virtual machine	-	60

[...content deleted for brevity...]

VMware ESX 3.x Server Overview

VMware ESX Server is the foundation for every other piece of the virtualization package. It is the *hypervisor*, or main software layer that installs on the bare metal and allows everything above it to communicate with the hardware, to allow virtualization. When you install VMware ESX Server, you are actually installing two main components: the VMkernel and the Service Console.

The VMkernel is the base on which all other software in the package is built: the operating system. For those familiar with Linux, this would be the equivalent of (and is built from) the Linux kernel, without any other software.

The Service Console (or COS) is an alternative means of communicating with and configuring the VMkernel using standard Linux and VMware-specific commands to modify and adjust parameters. Typically, management will be done via the vCenter client; however, there may be cases where you find you can accomplish things more easily with the command line or need to use it because you can't access vCenter.

VMware has designed ESX Server to run only on specific pieces of hardware and has removed support for any kinds of devices it is not interested in, thereby reducing the kernel code. What remains is a stripped-down, fast kernel and tool package with little to no extra overhead. This is one of the things that gives VMware an advantage over other virtualization technologies that require installation on top of a standard operating system, which will be filled with drivers and features you won't need.

It's important to verify the hardware on which you will be running your virtualized environment, as VMware doesn't directly support smaller desktop-related hardware. However, everything that a server needs is well supported. This is a sleek operating system designed to put as little as possible between the virtual machines and the hardware.

VMware ESX 3.x Installation

Installation of ESX 3.5 is pretty straightforward if you are familiar with installing Linux, and especially Red Hat, as VMware has taken Red Hat's default installer and made a few modifications.

ESX 4.0 and ESXi 4.0 can now be installed and virtualized within VMware Workstation and VMware Fusion. This allows you to install, test, and get a feel for the product before implementing it on physical hardware.

Before you get started installing ESX, you should verify that your hardware is fully compatible by visiting the following URL—you will notice that ESX 3.5 has a wider range of hardware support than its newer relative, ESXi:

[http://www.vmware.com/resources/compatibility/search.php?
action=base&deviceCategory=server](http://www.vmware.com/resources/compatibility/search.php?action=base&deviceCategory=server)

The ESXi ISO image is available for download at <http://www.vmware.com/download/vi/>.

ESX Server has the following hardware requirements:

- At least two processors chosen from among the 1500MHz Intel Xeon and later, AMD Opteron (32-bit), 1500MHz Intel ViiV, or AMD A64 x2 dual-core processors (be sure to check the latest specs to confirm)
- 1GB RAM minimum
- One or more Ethernet controllers (we recommend a minimum of four ports for the Service Console, VMkernel, Virtual Machine Network, etc.)
- Direct attached storage (DAS) or network attached storage (NAS)

To avoid possible data loss make sure to remove any attached storage, such as Fibre, DAS, or iSCSI, before the initial installation of ESX.

The installer has a couple of modes:

Graphical mode

The recommended method for installation. It uses a mouse and a graphical interface to guide you through the installation. This is the method we will be using for our installation as we continue in this section.

Text mode

A text-only mode for installing ESX; this option is normally used when you have mouse or video problems in the graphical installer.

To begin the installation, make sure your server is set to boot off the CD-ROM. This can be accomplished by changing the settings in your server's BIOS. Next, follow these steps:

1. Upon booting from the CD-ROM, you will be presented with a screen that allows you to choose your installation method. Pressing the Enter key accepts the default graphical installer. By typing **esx text**, you can invoke the text-mode installer (Figure 1-1).



Figure 1-1. ESX installation screen

The installer will begin to load, displaying another screen that will allow you to test the CD-ROM media for errors. If you want to test the media, use the Tab key to highlight the Test button; otherwise, select the Skip button to continue the installation (Figure 1-2).



Figure 1-2. CD-ROM media test

2. A welcome screen will be presented. Click the Next button to proceed to the keyboard configuration (Figure 1-3), where you will select the default keyboard layout. Click Next to continue.

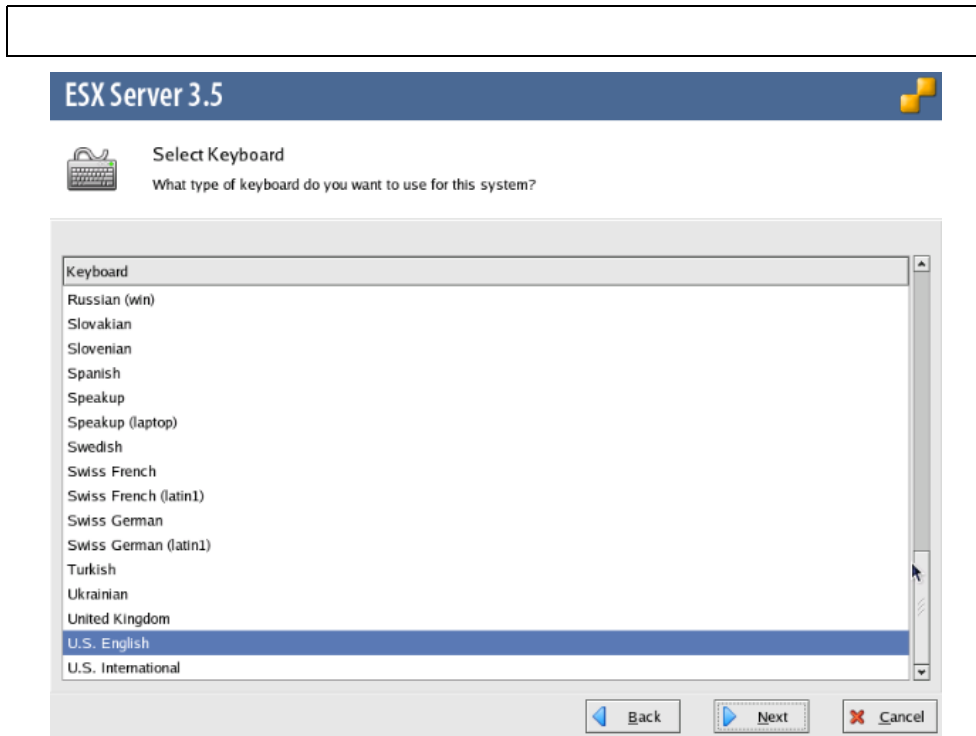


Figure 1-3. Keyboard selection during ESX installation

3. The mouse configuration menu now appears, allowing you to select the mouse type for your ESX server (Figure 1-4). Accepting the default here is okay because the mouse will not be used at any point after the initial installation. Proceed by clicking the Next button.

[...content deleted for brevity...]

VMware ESXi 3.5 Overview

ESXi is a newer form of the technology that strips down the ESX Server even further, to the absolute bare minimum possible. It does not include or rely upon a service console and can perform its hypervisor duties with an installation of a mere 32MB of disk space. In ESX Server, the service console runs on top of the hypervisor alongside the installed virtual machines. In ESXi, the hypervisor is all you install on your hardware. This allows for very fast hardware additions to existing pools. Want to add a new server to your rack? No problem. Hook it up, power it on, install ESXi, set your root password and networking details, take a minute or two to configure the virtual network, and you're all set. This is an amazingly quick way to create or expand an existing ESX hardware cluster.

Even more exciting than the installable version of ESXi is the embedded version. It is now possible to buy hardware that is preconfigured with ESXi installed on it. This completely eliminates the installation step and speeds up the configuration process.

Companies like Dell, which graciously allowed us to borrow some of its equipment as we were writing this book, are now selling servers that require only a little more than racking, cabling, and powering on to expand your VMware cluster.

The benefits of using ESXi over the standard ESX product are that it's extremely lightweight and installs quickly, it can be purchased on some servers from Dell and other vendors as an embedded option, it has no service console, and it requires less patching and maintenance. ESXi is also available for free and includes the basic functions needed to start a virtual environment.

VMware ESXi 3.5 Installation

VMware customers who have received a pre-installed OEM version of ESXi may have vendor-specific customizations and drivers. These versions will be different from the downloadable version on VMware's website.

Before you get started installing ESXi, you should verify that your hardware is fully compatible by visiting the following URL, as VMware's ESXi product has a much smaller compatibility list than its older ESX model:

[http://www.vmware.com/resources/compatibility/search.php?
action=base&deviceCategory=server](http://www.vmware.com/resources/compatibility/search.php?action=base&deviceCategory=server)

ESX 4.0 and ESXi 4.0 can now be installed and virtualized within VMware Workstation and VMware Fusion. This allows you to install, test, and get a feel for the products before implementing them on physical hardware.

As of update 3, ESXi has the following hardware requirements:

- 1500MHz Intel Xeon and later models or AMD Opteron (32-bit mode) for ESXi
- 1500MHz Intel Xeon and later models or AMD Opteron (32-bit mode) for Virtual SMP
- 1500MHz Intel Viiv or AMD A64 x2 dual-core processors
- 1GB of RAM minimum
- Broadcom NetXtreme 570x Gigabit controllers or
- Intel PRO 100/1000 adapters
- Basic SCSI adapters—Adaptec Ultra-160, Ultra 320, LSI Logic Fusion-MPT, or most NCR/Symbios SCSI controllers
- RAID adapters—Dell PercRAID (Adaptec RAID and LSI MegaRAID) or IBM Serve RAID controllers

This list represents the minimum requirements to install ESXi, but many more devices are available. To ensure complete compatibility, search for your specific hardware at the abovementioned URL.

The ESXi ISO image is available for download at <http://www.vmware.com/download/esxi/>.

VMware has done a great job of streamlining the installation of ESXi. If you are familiar with installing ESX, this will be a breeze. After booting to the installation CD-ROM, you'll notice a blue screen titled "ThinESX Installer" that looks similar to the bootloader in Linux.

Once the initial boot has completed, the installer will present a couple of options: you can choose to cancel the installation, upgrade an existing installation, or perform a fresh ESXi installation (Figure 1-12). Once you have selected the appropriate installation type, you will be shown the EULA. Press F11 to accept and continue the installation.




Figure 1-12. ESXi initial installation

Select the disk on which you wish to install ESXi. Typically, this will be a local disk. Also, make sure you have disconnected all your SAN and direct-attached storage, to prevent any accidental data removal. Once you've selected the disk, the installation will begin and will take only a few minutes. When it is complete, remove the CD-ROM and reboot the system.

Unlike a normal ESX installation, ESXi requires a little more configuration after the initial install, because the only thing being configured during installation is the disk on which the hypervisor is going to be installed. Once the new ESXi server has been rebooted, you will notice a screen that is similar to ESX but with two new options: F11 to reboot the system and F2 to manage the system.

To begin configuration, make sure you have a keyboard and monitor plugged into the ESXi server. Press F2 to access the menu-based configuration utility, which will be the main screen in which all ESXi configuration will take place (Figure 1-13).

[...content deleted for brevity...]

See Also

[Recipe "VMware ESX 3.5 Installation"], [Recipe "VMware ESX 3.x/4.x Configuration Maximums"]

VMWare vCenter Server 2.x Overview

There comes a point as networks grow when it becomes unwieldy to manage each server individually. For some of us, that happens right around the time the second box is added, because we want to use our time for more interesting things than repetitive procedures

and maintenance. VMware vCenter provides a central location for managing all of the virtual machines on your VMware network. It is licensed and sold separately and requires a dedicated Windows server or a Windows virtual machine and a database (Oracle or Microsoft SQL Server). Once installed and configured, it will make the system administrator's life much easier. vCenter provides a way to distribute resources, manage users, move virtual machines from one piece of physical hardware to another (while still running!), schedule tasks, and much more.

[...content deleted for brevity...]