



Covers
ESX & ESXi 3.5
and vCenter
Server 2.5,
Update 3

VMware® VI3

IMPLEMENTATION AND ADMINISTRATION

ERIC SIEBERT



Foreword by JOHN TROYER, VMware

Praise for *VMware VI3 Implementation and Administration*

"With this book, Eric has created a roadmap through the intricacies of VMware Infrastructure 3. From initial requirements gathering to advanced configuration topics, the logical progression of the chapters allows readers to jump in at whatever point makes the most sense in their own environment. Beginners and veterans alike will find themselves using it as both a guide and a reference."

—Mark Chandler, Infrastructure Engineer Principal,
VMware Certified Professional and vExpert 2009

"Through his work at Boston Market, his involvement as a community moderator on the VMTN forums, his management of his own vmware-land.com website, and his active blogging, Eric Siebert has established himself as an expert in the field of virtualization. In this book, he shares his knowledge of VMware virtualization in a manner that will benefit the novice and the expert alike."

—Ken Cline, Technical Director, Virtualization, Wells Landers

"Eric's book takes the new VMware administrator by the hand and slaps the seasoned administrator upside the head. He does a great job of covering everything from capacity planning to infrastructure monitoring, all the while explaining the pros and cons that only an experienced virtualization implementer could provide."

—Rich Brambley, Virtualization Consultant,
owner and sole contributor of <http://vmetc.com> blog

"This is another great chance to learn about VMware virtualization techniques from a well-informed seasoned professional. This book is great for someone who is starting out with a VMware infrastructure as well as a seasoned professional looking to reinforce preexisting knowledge."

—Stephen Beaver, Virtualization Evangelist, Tripwire

"I have read hundreds of articles by Eric Siebert and always wished that I could have them in a single book. That wish has finally come true. In Eric's *VMware VI3 Implementation and Administration* book, he covers VI up and down, inside and out, like only Eric can. It's the ultimate VMware VI manual."

—David Davis, VCP, vExpert, and author of the
VMware ESX Video Training Course from www.TrainSignal.com

"*VMware VI3 Implementation and Administration* educates even the most experienced VI administrator and should be part of every (virtual) toolkit!"

—Duncan Epping, Senior Consultant, VMware; virtualization blogger, Yellow-Bricks.com

"This book is a very well-written, comprehensive field guide that does a fantastic job of taking the reader across the entire lifecycle of a virtualization deployment—walking through the background information, the planning stages, and then the actual implementation itself—and it then wraps things up nicely by providing the help and the knowledge necessary to properly maintain and troubleshoot the virtual environment once it goes live. The advantage of reading a book like this from an author who is a subject matter expert and who obviously has years of hands-on experience with the technology is that you can quickly and easily identify the problem areas and technology gaps before you have to experience them first hand; and you can avoid those common pitfalls and gotchas that others have had to go through in their own journeys."

—David Marshall, owner, VMBlog.com

"As an experienced VMware certified instructor, I'm involved in delivering VMware courses on a daily basis. I know how important it is to keep your students involved in a course. When you start reading this book, I can assure you, you will read it from A to Z. In the past three years I've seen lots of books. Eric Siebert did a great job with this one. He was able to write up topics varying from 'installing' and 'advanced configuration' in a clear and understandable way and put them together in this book. *VMWare VI3 Implementation and Administration* is...a must-read for every VMware believer."

—Eric Sloof, NTPRO.NL

"*VMware VI3 Implementation and Administration* is a practical and informative guide on how to get started building a VMware Virtual Infrastructure. Eric Siebert has done a wonderful job explaining in a step-by-step approach (including screen shots) how to accomplish administrative tasks on the latest VMware virtualization software. This book will be a handy guide for beginning VMware administrators to implement VMware best practices to avoid problems and enhance the performance and security of their virtual machines."

—George Vish, Senior Education Consultant

VMware VI3 IMPLEMENTATION AND ADMINISTRATION

ERIC SIEBERT



Upper Saddle River, NJ • Boston • Indianapolis • San Francisco
New York • Toronto • Montreal • London • Munich • Paris • Madrid
Cape Town • Sydney • Tokyo • Singapore • Mexico City

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U.S. Corporate and Government Sales
(800) 382-3419
corpsales@pearsontechgroup.com

For sales outside the United States please contact:

International Sales
international@pearson.com

Visit us on the web: informit.com/ph

Library of Congress Cataloging-in-Publication Data:

Siebert, Eric, 1966-

VMware VI3 implementation and administration / Eric Siebert.
p. cm.
ISBN 978-0-13-700703-5 (pbk. : alk. paper) 1. VMware. 2. Virtual computer systems. I. Title.
QA76.9.V5S48 2009
005.4'3—dc22

2009007980

Copyright © 2009 Pearson Education, Inc.

All rights reserved. Printed in the United States of America. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, write to:

Pearson Education, Inc.
Rights and Contracts Department
501 Boylston Street, Suite 900
Boston, MA 02116
Fax (617) 671-3447

VMware is a trademark of VMware, Inc. Used with permission.

ISBN-13: 978-0-13-700703-5

ISBN-10: 0-13-700703-5

Text printed in the United States on recycled paper at Courier in Stoughton, Massachusetts.

First printing May 2009

Editor-in-Chief

Mark Taub

Acquisitions Editor

Trina MacDonald

Development Editor

Songlin Qiu

Managing Editor

Kristy Hart

Project Editor

Betsy Harris

Copy Editor

Keith Cline

Indexer

Brad Herriman

Proofreader

Williams Woods

Publishing

Technical Reviewers

George Vish

Ken Cline

Steve Beaver

Dave Mishchenko

Publishing Coordinator

Olivia Basegio

Cover Designer

Chuti Prasertsith

Compositor

Nonie Ratcliff

This book is dedicated to all the great members of the VMTN community who share their knowledge and experience with others on a daily basis. Their selfless teaching and dedication to helping others is a credit to themselves and a big reason for the continued success and popularity of VMware.

I also dedicate this book to my family: my mother and father, Yolanda and Edward; my siblings, Jason, Elizabeth, and Jennifer; my wife, Kimberly; and my children, Logan, Zechariah, and my precious daughter Sophia.

This page intentionally left blank

Contents

FOREWORD IX

INTRODUCTION 1

1 ASSESSING YOUR CURRENT
ENVIRONMENT 3

2 PLANNING YOUR VIRTUAL
ENVIRONMENT 25

3 BUILDING YOUR VIRTUAL
ENVIRONMENT 75

4 CONFIGURING YOUR VIRTUAL
ENVIRONMENT 119

5 SECURING YOUR VIRTUAL
ENVIRONMENT 219

6 POPULATING YOUR VIRTUAL
ENVIRONMENT 263

7 MONITORING YOUR VIRTUAL
ENVIRONMENT 335

8 MAINTAINING YOUR VIRTUAL
ENVIRONMENT 365

9 BACKING UP YOUR VIRTUAL
ENVIRONMENT 399

10 TROUBLESHOOTING YOUR VIRTUAL
ENVIRONMENT 409

11 ADVANCED TOPICS 435
INDEX 467

This page intentionally left blank

Foreword

Virtualization is fundamentally changing the data center and how we approach computing. A friend recently showed me a demo of his latest project, a 3D simulation of a data center. In the demo, your avatar would walk around the rows of racks in a virtual representation of your physical facility. With a click of the mouse, a server in a rack would open like a drawer and a little status screen would emerge and hover in the air in front of your virtual data center administrator. “That’s a really fun demo,” I told my friend, “but it would have actually been useful even five years ago. Now? Not so much.”

My friend, who had been in the software industry for 30 years, had missed the wildfire impact that virtualization has had on the IT industry. The hardware is still there, racked up and plugged in, but the physical server is no longer the unit of work or the unit of management in the data center. We no longer look at one physical machine at a time; we manage entire data centers in a single pane of glass. And just as I don’t care exactly which disk sector my bits are stored on, in most cases these days I don’t care which physical server my applications are running at the moment—or even my desktop. (I’m happy to report my friend has seen the light and is now building a virtual representation of his virtual infrastructure.)

The wave of x86 virtualization technology from VMware and others has enabled this transformation, but the benefits of the technology are what have supercharged the velocity of the change. My employer, VMware, reports its 130,000 customers can reduce hardware and operating costs by as much as 50%, reduce energy costs by 80%, reduce the time it takes to provision new servers by up to 70%, and save more than \$3,000 per year for every server workload virtualized. Usually the savvy IT professional takes marketing numbers from a vendor with a grain of salt, but just talk to your fellow IT professionals for confirmation. Every day on blogs and social networks, I see messages pass by like “I Love VMware,” “Snapshots saved my bacon once again,” and “Entering maintenance mode, watching production servers VMotion while eating lunch at my desk.”

One of the paradoxes of virtual infrastructure is that most things are the same as physical infrastructure, while at the same time being completely different. You’re running

the same applications in the same topologies. You can treat the applications basically the same as you had been doing previously—after all, from inside the virtual machine, the workloads think they’re still back in a physical box. But at the same time, virtualization touches and transforms every part of the enterprise software stack: networking, storage, security, disaster recovery, management, provisioning, and other business processes, and even how you handle the financials. As a result, the syllabus for a virtualization training class often looks like a complete university curriculum in information technology because it has to touch on all these areas. Virtualization experts are the mixed martial artists of IT—experts in kicks, punches, throws, and wrestling moves.

Because of this breadth of impact, virtualization can be daunting. VMware has had a robust set of user forums for years, where the storage expert and the networking expert can come together and help each other with their virtual infrastructure projects—and both become virtualization experts in the process. By now, the VMware Communities is the best enterprise software online community I’ve ever seen, and by far the best place to ask a quick or not-so-quick question about your virtual infrastructure. There you’ll often see esiebert 7625, the author of this book, answering questions and adding to the conversation.

Eric Siebert is a natural encyclopedist and cataloguer. You can now find more information about VMware online than most of us can comfortably digest—documentation, white papers, presentations, blogs, wikis, magazines, and community sites abound. On the Communities, Eric has a knack for not only answering your question, but also always seeming to have a set of links to resources that explain the answer and give you a mini-course in why the answer is the way it is. That’s why on the VMware Communities, you’ll see a little brain icon with “Guru” beside Eric’s handle.

Eric is a prime example of the Roman philosopher Seneca’s maxim *docendo discimus*—by teaching we are learning. Eric is a working, hands-on VMware administrator and 25-year IT veteran, but with his work helping literally thousands on the VMware Communities and through his website, freelance writing, and blogging, Eric truly has become an expert educator. Eric was recently one of the first recipients of the VMware vExpert award for his contributions giving back to other virtualization users.

This book, written by an expert educator and hands-on practitioner, takes you through the full lifecycle of a virtual infrastructure implementation, clearly lays out both the concepts and the steps required for someone new to virtualization, and can serve as a quick, clear review of best practices for the more experienced virtualization administrator.

I wish you luck in your journey in virtualization. The first time you come into work in the morning and realize that one of your servers restarted during the night because of VMware’s High Availability feature, but neither your monitoring systems nor your end users noticed anything amiss, I encourage you to march into your boss’s office and ask for a promotion and a raise. You will have earned it.

John Troyer

VMware Communities
Palo Alto, CA

Acknowledgments

First and foremost, I want to acknowledge some of the members of the VMTN community who are my peers and, more important, my friends. This includes Ken Cline, whose knowledge and experience continues to astound me; Dave Mishchenko, whose dedication to helping others is equal to mine; Edward Haletky, who knows more about security than anyone I know and is always happy to share his knowledge with others; Oliver Reeh, who inspired me to start my website; Steve Beaver, whose passion for virtualization is very infectious; Jason Boche, who is the biggest VMware geek I know; Robert Dell'Imagine, who runs the VMTN communities; and John Troyer, who is VMware's head blogger and virtualization evangelist.

I also want to acknowledge a few fellow bloggers from the thriving VMware user community who inspire me on a daily basis to help others and be the best that I can be, including Eric Sloof from ntpro.nl, Scott Lowe from blog.scottlowe.org, Duncan Epping from yellowbricks.com, Rich Brambley from vmetc.com, and my fellow bloggers at Tech Target. There are many other bloggers and members of the VMTN community whom I haven't mentioned, but you know who you are and you are all an inspiration to me.

I also want to thank Jan Stafford from Tech Target, who got me started in writing, and all the great editorial staff at Tech Target, including Hannah Drake, Adam Trujillo, Mark Gallagher, Matthew McDonough, and Jo Maitland.

Finally, I want to thank my editor at Pearson, Trina MacDonald, who gave me this opportunity and believed in my ability to write this book. Trina is a wonderful person who helped me through the many behind-the-scenes things that go into writing a book, including formatting, layouts, and all the other various steps it takes to write a book. In addition, thanks to my development editor, Songlin Qiu, for correcting my many formatting mistakes. And, a huge thank you to my technical editors, Ken Cline, Dave Mishchenko, Steve Beaver, and George Vish, who all kept me honest, and whose great knowledge and experience helped make this book even better.

About the Author

Eric Siebert has been working with computers for more than 25 years. For the past 15 years, his main focus has been the administration of Windows server environments and supporting enterprise applications such as web, email, and database servers. Four years ago, he discovered virtualization technology in the form of VMware ESX and has been hooked ever since, spending countless hours learning the product inside and out. He now spends much of his time helping others in VMware's community support forums, having achieved Guru status and become one of the forum's moderators. His own website, <http://vmware-land.com>, is where he shares his tips and experiences and hundreds of links to technical information and top ten lists on a variety of subjects. In addition, Eric writes tips and blogs for Tech Target on sites such as searchvmware.com, searchservvirtualization.com, and searchdatabackup.com. Most recently, he was a speaker, panelist, and Best of VMworld awards judge at VMworld 2008 and is a regular guest on VMware's weekly VMTN roundtable podcast. In 2009, Eric was awarded the VMware vExpert award, one of only 300 people worldwide recognized for their knowledge, experience, and contributions to the VMware community.

Introduction

Virtualization is not a new technology, but it has gained popularity in recent years and is used to some degree in almost all datacenters. For most companies, it's not a question of *when* they are going to virtualize their infrastructure but *how much* they are going to virtualize. Virtualization has many benefits over traditional physical servers, and the technology is constantly evolving and improving to further make the decision of whether to virtualize an easy one. In addition to the many vendors offering virtualization hypervisors, such as VMware, Microsoft, and Citrix, the physical hardware technology vendors, such as AMD and Intel, are changing their products to optimize them to work with virtual hosts. In addition, a large amount of vendors have written virtualization-specific applications or modified their applications and hardware to work in virtual environments because they recognize that virtualization is here to stay. Virtualization is so popular now that almost all software vendors support it. So, you don't have to worry about issues with a vendor supporting their application running in a virtual environment. In addition, most vendors have changed their licensing policies so that they are friendly to running on virtual machines.

If you are looking at this book, you must be interested in learning more about virtualization and how to implement it. This book was written to cover all the various phases of implementing a virtualization project using VMware Virtual Infrastructure 3, from the initial planning stages all the way through designing, building, configuring, maintaining, troubleshooting, and more. VMware has made it fairly easy to install and use their virtualization products, but there's a lot you need to know to properly set up your environment and to understand how everything works, including the differences between

physical and virtual servers. This book walks you through the various stages and provides information and tips to guide you through them so that you make informed decisions, use best practices, and avoid common mistakes.

This book was written using ESX version 3.5 Update 3 and vCenter Server version 2.5 Update 3. VMware is continuously updating their products to further improve them and provide more features and functionality, and as of the writing of this book, VMware is close to releasing their next-generation product (VI4), tentatively named vSphere. This new version should be an exciting release as VMware continues to build on their Virtual Datacenter OS vision and to head toward their vCloud initiative. Look for a follow-up to this book that will help you understand the new version and how to implement and administrate it and upgrade to it from the VI3 release.

A friend of mine likes to say that virtualization is a journey, not a project. The journey begins with learning about virtualization, but does not end after you implement it. Virtualization is an enabling technology that will change the way you do things in your datacenter and provide you with greater flexibility and more options for administering your servers. My own virtualization journey began many years ago as part of a server consolidation project and continues today; I am still learning and adapting to all the new technologies and features that come with each new release. So head on over to the first chapter and let your own journey into virtualization begin!

Chapter 1

Assessing Your Current Environment

Before starting on your virtualization journey, it is important to thoroughly understand your current IT infrastructure. By introducing virtualization into your environment, you are making a big change that will have a ripple effect on all parts of your environment. Standard procedures such as monitoring, backups, patching, and administration will all be affected by this. In addition, all components of your infrastructure will most likely be affected in some way by this big change. As a result, you need to assess all parts of your infrastructure, not just the servers you plan to virtualize, to uncover any potential problems or hurdles that may impact your project. The old woodworking rule “measure twice, cut once” also applies with computers. You can save yourself from making costly mistakes by making sure you get accurate measurements before you begin.

An Important Note

In December 2008, VMware announced that they were changing the product names of some of the components in VI3. This was done to better align the product naming with their Virtual Datacenter OS vision. The main change involved introducing vCenter as the new name for their many automation and management products. This affected VirtualCenter, which was renamed as vCenter Server, and so subsequently this book has been updated to use this new name. However, although the name of the product has changed, the application and documentation for VI3 has not, and you will still see the old name used in both. This name change applies only to VirtualCenter Server 2.5, and the terms vCenter Server and VirtualCenter refer

to the exact same product. Older 2.0.x versions of VirtualCenter continue to use that name and are not considered part of the change. It is expected that the new name will be fully used in the application and documentation when VMware introduces their next major release of ESX and vCenter Server (VI4 or vSphere) sometime in 2009. In addition, some of the other products have had the name vCenter added to the beginning of their names, such as vCenter Converter and vCenter Update Manager, and that has also been reflected in this book.

Documenting Your Current Server Environment

Most virtualization projects will involve migrating your current physical servers to virtual machines (VMs). Therefore, it is important that you thoroughly understand your current environment before attempting to migrate it to virtual servers. By doing this, you can ensure that you purchase properly sized server hardware and the right number of VMware licenses. It's a good idea to do a thorough inventory of all your current physical servers so that you know exactly what you have before you start virtualizing. Also identify what you intend to do with the old physical hardware after it has been virtualized. Often, you may end up reusing newer physical hardware as ESX hosts. It's best to decide what you will do with your old servers as part of your planning so that you will know exactly what hardware you will be discarding, reusing, and leaving alone.

Watch Out!

Don't virtualize known problems; make sure your current server environment is healthy before attempting to virtualize it. For example, if you have existing performance or application problems in a physical environment, attempt to resolve those before moving them to a virtual environment.

Measuring Your Current Performance Usage

Measuring your current performance is necessary so that you can get a good idea of how your current environment is performing. By doing this, you can ensure that you properly size your virtual hardware and can avoid any bottlenecks on your ESX hosts. Doing this before you start your project is important so that you do not run into any surprises that can cause problems during your deployment phase.

What to Measure

You should focus on four general performance categories: CPU, memory, disk, and network. You should gather these metrics for a minimum of one week, and preferably over a one-month period of time. Gathering these metrics for a longer period of time gives you a better understanding of any performance trends that you may be experiencing that might not happen on a regular basis. It is also important to gather metrics during critical business cycles (for example, weekly payroll processing or a monthly reporting process) where performance may spike. The combined results of these metrics will help determine your overall consolidation ratio (number of VMs per ESX host) and how many ESX servers you will need for the number of physical servers that you want to virtualize. Consolidation ratios can vary from as little as 2:1 to as high as 50:1 based on the total amount of resources that your VMs will require and the size of your ESX host servers.

Let's go over the categories, some important metrics, and some guidelines on each one.

Measuring CPU Usage

Typically, most Windows servers have very low overall CPU usage (< 10%), which is why virtualization is a great solution to maximize your hardware resources and reduce the number of physical servers in your environment. Average processor utilization is the best metric to use to measure how busy a server actually is. It will give you an overall indication of how much processor the physical server is using, which you can use to help plan your ESX host size. Most servers will peak near 100% at various times, but the peaks are not as important as the overall average utilization. High processor queue lengths can indicate a bottleneck on a physical server, which may disappear in a virtual environment because of the way the ESX hypervisor handles the scheduling and processing of CPU requests. Table 1.1 lists the CPU metrics that you will want to watch to determine the amount of CPU usage on your servers.

Table 1.1

Important CPU Metrics

Statistic	Description	Why This Is Important
Processor queue length (average and maximum)	Processor queue length is the number of threads in the processor queue. There is a single queue for processor time even on computers with multiple processors or cores. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload.	A sustained processor queue length of ten or more threads typically indicates a processor bottleneck.

continues...

Table 1.1 continued

Statistic	Description	Why This Is Important
% processor time (average and maximum)	% processor time is the percentage of elapsed time that the processor spends to execute a non-idle thread. It is calculated by measuring the duration the idle thread is active in the sample interval and subtracting that time from the interval duration. (Each processor has an idle thread that consumes cycles when no other threads are ready to run.) This counter is the primary indicator of processor activity, and displays the average percentage of busy time observed during the sample interval. It is calculated by monitoring the time that the service is inactive and subtracting that value from 100%.	This value indicates how much CPU that your server is actually using, which can be used to plan the amount of CPU needed on a virtual host.

Measuring Memory Usage

The actual amount of physical memory that a server uses will determine how much memory your ESX hosts will need to be able to support all the VMs on it. It is possible to overcommit an ESX host (assigning VMs more memory than the host physically has), but it is not recommended in most cases because it will degrade the performance of your VMs once your host's physical memory has been used up. Table 1.2 lists the memory metrics that you will want to watch to determine the amount of memory usage on your servers.

Table 1.2

Important Memory Metrics		
Statistic	Description	Why This Is Important
Available free memory (average and least)	Available MBytes is the amount of physical memory, in megabytes, immediately available for allocation to a process or for system use. It is equal to the sum of memory assigned to the standby (cached), free, and zero page lists.	This value indicates how much physical memory is not being used by your server. If you have excessive free memory then consider reducing the amount of RAM assigned to the server when moving it to a virtual host.

Measuring Your Current Performance Usage

Statistic	Description	Why This Is Important
Pages/sec (average and maximum)	Pages/sec is the rate at which pages are read from or written to disk to resolve hard page faults. This counter is a primary indicator of the kinds of faults that cause systemwide delays.	This value counts the number of times per second that the computer must access virtual memory rather than physical memory. This number normally increases as available memory decreases. Too many pages/sec can cause excessive disk activity and create a disk bottleneck. This often indicates that a system does not have enough physical memory.

Measuring Disk Usage

The important things to know about a disk are how much you are using (disk space) and how much reading and writing to the disk that each server does (transfer rate). Disk is the slowest of the resources because it relies on a mechanical device and is usually the first bottleneck to performance in most systems. Therefore, it is important to understand how much disk activity your servers will be doing so that you can select a proper storage solution for your virtual hosts. It's also important to factor in the number of spindles (hard disks) in your redundant array of inexpensive disk (RAID) groups on your physical servers. A RAID group with more spindles will have better disk performance than one with fewer spindles. If you were to virtualize a physical server with a ten-spindle RAID group, you may not get the same performance if your ESX host is configured with only a five-spindle RAID group. Table 1.3 lists the disk metrics that you will want to watch to determine the amount of disk usage on your servers.

Table 1.3

Important Disk Metrics

Statistic	Description	Why This Is Important
% disk time	% disk time is the percentage of elapsed time that the selected disk drive was busy servicing read or write requests.	Similar to % processor time, this can be useful in characterizing the workload and gives a general indication of how busy the disk is.

continues...

Table 1.3 continued

Average disk queue length	Average disk queue length is the average number of both read and write requests that were queued for the selected disk during the sample interval.	This tells you how many I/O operations are waiting for the hard disk to become available. This number should be as low as possible; a high number (> 5) can indicate an I/O bottleneck depending on the number of spindles (hard disks) in your RAID group. It's best to divide your average queue length by the number of spindles in your RAID group to get a more accurate number.
Disk bytes/sec	Disk bytes/sec is the rate bytes are transferred to or from the disk during write or read operations.	This provides information about the throughput of the disk system and how busy it is.
Physical disk transfers/sec	Disk transfers/sec is the rate of read and write operations on the disk.	This is the total number of read and write requests processed per second (commonly known as I/O operations per second or IOPS). Like disk bytes/sec, this also measures the throughput of the system. The difference is that this counter does not consider the size of the request, just the fact that it is a request.

Measuring Network Usage

Network is a resource that typically is plentiful in virtual hosts because you can easily put many multiport high-speed network interface cards (NICs) in your ESX servers. You should still identify any servers that generate a large amount of network traffic so that you can add extra NICs if needed, and it will also help you when you build and configure your virtual switches (vSwitches). Also, network traffic between VMs that are on the same vSwitch does not go over the physical network (it travels along the host bus), which could reduce the amount of network traffic generated by your servers after they are virtual. Table 1.4 lists the network metrics that you will want to watch to determine the amount of network usage on your servers.

Table 1.4

Important Network Metrics

Statistic	Description	Why This Is Important
Bytes total/sec	Bytes total/sec is the rate at which bytes are sent and received over each network adapter, including framing characters. Network interface\\bytes received/sec is a sum of network interface\\bytes received/sec and network interface\\bytes sent/sec.	This counter shows the amount of traffic through your network interface in byte per second.

How to Measure It

You can determine your current performance usage in a number of ways:

- Use existing enterprise monitoring systems.
- Use operating system built-in performance monitoring tools (PerfMon).
- Use third-party analysis tools such as PlateSpin PowerRecon or Tek-Tools Profiler for VMware.
- Contact a VMware business partner and have them install the Capacity Planner tool in your environment.

VMware's Capacity Planner

Capacity Planner is a powerful tool that automatically collects all the relevant performance metrics on each Windows server in your environment and prepares a report that you can use to determine your hardware requirements for your virtual environment. It can identify trends in your environment and make recommendations for grouping physical servers on virtual hosts. It uses the built-in Microsoft performance counters and does not require that an agent be installed on each server that will be analyzed (it uses the WMI and the Remote Registry service). The Enterprise dashboard screen from Capacity Planner, along with all the other available options, is shown in Figure 1.1.

Currently, Capacity Planner is provided by VMware to its business partners only and is not available to the general public. Most business partners will install and configure it in your environment for you for free as long as you plan on buying software/hardware and professional services from them for your virtualization project. Using Capacity Planner is the best method for collecting data from your servers and reporting on it, because it was developed specifically for infrastructure assessment and data analysis and will provide consolidation estimates, recommendations, and capacity assessments.




Figure 1.1 Sample screen from VMware’s Capacity Planner tool

Beginning with vCenter Server version 2.5, a “lite” version of Capacity Planner was integrated into vCenter Server as a feature called Guided Consolidation. This utility uses a built-in wizard to discover physical systems and analyze them to prepare them to be converted into VMs. Once these systems have been analyzed, they can be converted into VMs by the built-in VMware Converter feature of vCenter Server 2.5. The data gathered by this utility is basic and does not use some of the more advanced metrics that the full version of Capacity Planner uses. It can analyze up to 100 systems simultaneously and reports only on average CPU and memory utilization. Because of its limitations, it is recommended that you use a more robust performance monitor for your initial implementation. We discuss this feature in detail in a later chapter.

Using Built-In Operating System Tools to Gather Server Performance Statistics

For your Windows servers, you can use the Windows built-in performance monitor utility (PerfMon) to measure your server’s statistics. The downside to this method is that you will have to set up, collect, and review the statistics for each server individually, which can be time-consuming if you have many servers. Alternatively, you can set up a dedicated workstation or server to centrally monitor and collect statistics from each server.

Most Linux servers have only built-in real-time statistic reporting tools. You may look at some free tools that provide historical performance reporting for Linux servers, like Sysstat (<http://pagesperso-orange.fr/sebastien.godard/>).

If you do choose to use PerfMon to gather your statistics, the following steps will help you set up and configure it. Before you begin, if you are going to use a central workstation to collect statistics, make sure the Performance Logs and Alerts service on the workstation is

configured to start with a domain account that has access to every server that you want to monitor:

1. Load the PerfMon utility on a workstation or server (Administrative Tools > Performance).
2. In the left pane, select Counter Logs (located under Performance Logs and Alerts).
3. Select Action from the top menu (or right-click Counter Logs) and choose New Log Settings.
4. Enter a descriptive name for your log settings.
5. Click the Add Counters button.
6. Choose the Select Counters from Computer option, and type in the name of one of the servers you are going to monitor below it. Be sure and include the \\ before the Windows server name.
7. After you enter your server name, it will connect to it and display a list of available counters below it, as shown in Figure 1.2.



Figure 1.2 PerfMon Add Counters window

8. Select the performance object that you want to display counters for (for example, processor, memory, network interface), and then select the individual counter (for instance, Pages/sec), select All Instances if it is applicable (except for Network Interfaces, you do not want to select the Loopback interface) and not grayed out, and then click the Add button.
9. Repeat this for every performance counter that you want to monitor on the server. The recommended counters you will want to add are listed here:

Memory: Available MBytes

Memory: Pages/sec

Processor: % processor time

System: Processor queue length

Network Interface: Bytes total/sec

Physical Disk: % Disk Time

Physical Disk: Avg. disk queue length

Physical Disk: Disk bytes/sec

Physical Disk: Disk transfers/sec

10. After you have added all counters for a particular server, you can type in a new server name to continue adding counters for other servers.
11. Click the Close button after you have added all counters.
12. Select the data sample interval, as shown in Figure 1.3; the default is 15 seconds, which is an aggressive interval and will result in more peak instances because of the shorter sampling period. You may want to consider changing this to a high interval between one and five minutes so that you do not overwhelm the workstation and cause it to miss data from some of the servers.
13. Click OK to save your custom log settings.




Figure 1.3 PerfMon Log Settings window

14. Collection will automatically begin (as indicated when the icon turns green). The results will be written to a log file (for example, C:\PerfLogs\MyServers000001.blg). You can stop it at any time by selecting your log settings and selecting Action, Stop (or by right-clicking it and selecting Stop). When you stop a collection, the log file it has written to is no longer used; a new log file is created once you start it again.
15. If you have stopped your collection, you can review it by selecting System Monitor in the left pane, and then clicking the Disk icon (View Log Data). Then, on the Source tab, select your log file that was created; optionally, you can change the time range. On the Data tab, add your performance counters for each server. On the General tab, select your view type (Graph, Histogram, or Report) and click OK. Your counter will be displayed, and you can see the minimum, maximum, and average results for each one, as shown in Figure 1.4.

Last		1138.940	Average	1179.118	Minimum	1116.139	Maximum	1363.724
		Duration						3:59
Color	Scale	Counter	Instance	Parent	Object	Computer		
	1,000	Available MBytes	---	---	Memory	\\lcefyre		
	1,000	Pages/sec	---	---	Memory	\\lcefyre		
	0.000...	Bytes Total/sec	VMware Accel...	---	Network Interface	\\lcefyre		
	1,000	% Processor Time	_Total	---	Processor	\\lcefyre		

Figure 1.4 PerfMon resulting historical data for each counter

16. It's a good idea to test this for a short period (for example, one hour) and review the results to make sure it is working before you leave it running for a longer period of time.

Using Enterprise Monitoring Systems

If you are using an existing monitoring system, try to report on only the appropriate statistics that will be relevant to determining your needs to size your virtual hosts. Too many statistics can make it more difficult to determine how busy a host is in each of the categories. Also, remember that when you convert your physical servers to VMs your enterprise monitoring system may not report accurate statistics because of the differences inherent with virtual environments.

What to Do with the Data You Collect

After you have gathered your performance statistics, you should group your servers into three categories:

- High overall resource utilization
- Medium overall resource utilization
- Low overall resource utilization

Then identify the servers that have the highest resource utilizations in specific areas: CPU, memory, disk, and network. You should then review the servers in the high overall resource utilization category to make sure that virtualizing them makes sense. Also, do the same for the top few servers in each of the specific resource areas. When you've determined which servers you want to virtualize, you can move on to sizing your hardware to match your expected workload.

It is helpful to put together a spreadsheet that contains the following information about your physical servers to help you add up the amount of CPU, memory, and disk needed for your ESX hosts:

- Server name
- Model
- Operating system
- Function
- Number of CPUs
- Speed of CPUs
- Total disk space
- Total disk space used
- Physical memory

Next, add your performance measurements to it:

- Average CPU usage (% processor time)
- Maximum CPU usage (% processor time)
- Average processor queue length
- Maximum processor queue length
- Average available free memory
- Minimum available free memory
- Average memory pages/sec
- Maximum memory pages/sec
- Average % disk time
- Maximum % disk time
- Average disk queue length

- Maximum disk queue length
- Average disk bytes/sec
- Maximum disk bytes/sec
- Average disk transfers/sec
- Maximum disk transfers/sec

Finally, add a ranking for each resource using a scale of one (least) to five (most) based on the averages for the measurements of each category. This ranking will help give you an idea of where each server ranks in usage for each of the resource areas. A server that has high ranking in more than two of the following categories may not be a good virtualization candidate:

- CPU resource usage
- Memory resource usage
- Disk resource usage
- Network resource usage

When you are done, you will have a spreadsheet that contains an inventory of all your physical servers and the resource usage statistics that you can use to help size your ESX hosts properly. In the next chapter, we discuss sizing hardware for your ESX hosts.

Which Servers Are Not Good Virtualization Candidates?

Almost all servers and workloads can be virtualized, but in some cases you may not want to virtualize certain servers because of high-resource utilization, licensing issues, and application support issues. Let's cover some reasons you might not want to virtualize certain servers and some reasons that you may consider virtualizing these types of servers:

- High-resource utilization servers

Why you might not want to virtualize. A server that has very high resource requirements may not always be as good a fit as a virtual server. Typically, these types of servers have very high CPU and memory usage and high disk and network I/O, and on a virtual host where multiple servers are competing for resources they might not perform as well.

Why you should consider virtualizing. When virtualizing these types of servers, you may be able to have only one or two VMs on a host server. You might wonder why anyone would want to put just a single VM on an ESX host. The reason for this is to take advantage of some of the powerful features that virtualization offers such as snapshots, VMotion, and high availability (HA) that are more difficult and costly to implement in a physical environment. Also, virtualizing these servers can make for easier disaster recovery implementation and simplified hardware upgrades.

- Vendor licensing models

Why you might not want to virtualize. Some applications, such as Oracle, do not have virtualization-friendly licensing and require you to license their software based on the number of physical CPUs in the host server and not the number of virtual CPUs assigned to the VM that is running the application. So, a VM with two virtual CPUs on a four-CPU host server would require you to purchase a license for four CPUs regardless of the fact that the VM that is running the application has only two virtual CPUs.

Why you should consider virtualizing. Thankfully, most vendors today license their software based on the number of CPUs assigned to the server regardless of whether they are physical or virtual. Other vendors are changing their licensing models to meet the growing demand for using virtual servers. Check with your vendor to see if they have changed their licensing model or have any plans to do so in the future.

- Application support

Why you might not want to virtualize. Some vendors will not provide support for their application if it is running on a virtual server. We discuss this in more detail later on.

Why you should consider virtualizing. Very few vendors do not support running their applications on virtual servers. Consider alternative support options. (For example, HP provides support for both VMware and Microsoft products.)

- Specific licensing requirements:

Why you might not want to virtualize. Certain applications use stricter licensing methods to prevent piracy and illegal use of their software. Examples of this are hardware dongles (parallel/serial port/USB device keys) that plug into the server and must be present at all times and specific MAC address or hard drive serial number licensing. If a VM moves from one host to another as a result of a failure or due to resource constraints on a host, then the hardware dongle will no longer be present and the application will no longer work.

Why you should consider virtualizing. There are ways to accommodate these types of licensing schemes on virtual servers. Digi makes a device called AnywhereUSB that works with ESX servers and provides IP-based connections to USB devices.

- Hardware that cannot be virtualized

Why you might not want to virtualize. Some servers might have nonstandard hardware like fax and modem boards that are not supported in ESX which supports only a limited, very specific set of hardware.

Why you should consider virtualizing. Solutions are available for faxing and using modems via network connections over IP.

Should I Consider a 100% Virtualized Environment?

Although you will find that having a 100% virtualized environment is certainly achievable, there are a few reasons why you might want to maintain a few physical servers:

- **Support issues.** Some application vendors may require you to reproduce the problem on a physical server if they suspect the virtual host might be causing the problem. For this reason, it is a good idea to keep a physical server around for certain infrastructure components like Active Directory. If you have eight domain controllers in your environment, you might consider virtualizing six of them and leaving two of them as physical servers. Same with database servers; if you leave one or two Oracle or SQL database servers as physical servers, it gives you the flexibility to move a database hosted on a VM to a physical server if the vendor requests it.
- **Infrastructure issues.** If your environment suffers a major failure (for example, a storage area network [SAN] goes down or you experience a major network failure), you may lose most of your VMs. For this reason, you may want to copy at least one DNS/DHCP physical server (because many functions rely on DNS to work properly).
- **All the eggs in one basket.** You will most likely be using shared storage with your ESX hosts to take advantage of all the features that require it. However, if something happens to that storage, it will affect all your hosts and VMs that utilize it. To offset this risk, consider running a few of your key infrastructure VMs (for example, domain controllers, authentication servers, database servers, and DNS servers) on local disk rather than shared storage. That way the VM will not be affected if something happens to your shared storage device.

Application Compatibility

The assessment of your environment should also include software applications in addition to hardware. You should do a complete inventory of applications running on your servers that you plan to virtualize so that you can ensure there will be no support or licensing issues when running them on virtual servers. You do not want to find out after you are done with your project that the application vendor will not provide support to you because the project is running on a virtual server. In addition, there might be special licensing considerations or configuration changes that need to be made to an application that has been virtualized:

- **Support issues.** One of the first steps that you should complete when considering virtualization is to determine whether the applications you use are supported by the vendor in a virtualization environment. Almost all applications will run properly on virtual

servers, but you will find that vendors typically have varying levels of support for virtual servers. The levels you will see will include the following:

Complete support for it. The vendor has certified their application to run on virtual servers and will support it without question. You will find most major applications will fall into this category, with a few notable exceptions, such as Microsoft.

Best effort support. The vendor will make an effort to support their application on a virtual server but may ask you to reproduce the problem on a physical server if they determine that the virtual environment is at fault. Microsoft falls into this category; if you have premier-level support with them, they will make more of an effort to help you.

No support. The vendor will provide no support for their application in a virtual environment. Typically, this is either because of known issues when running the application on a virtual server or that the vendor has not yet tested their product on virtual servers. If this is the case, you need to decide whether you want to risk virtualizing the application. If you do, plan for the times that you do need to contact support for help with problems (such as having a physical server available for reproducing the problem).

- **Licensing issues.** Some vendors have non-virtualization-friendly licensing models when you run their products on virtual servers; Oracle is a good example of this. Typically, they will still license based on the physical number of processors in the host server regardless of how many processors the VM has assigned to it. So for a VM running on a four-CPU host that has only a single virtual CPU, you may be required to license for four CPUs. Other vendors will change their models slightly for virtual servers. For example, IBM has a new processor value unit formula that makes licensing on virtual servers much more difficult to calculate compared to physical servers. Other vendors will license differently based on clusters of servers with pooled resources. If this is the case, you may need to create a smaller cluster just for the specific application to keep costs down. It's also common in virtual environments for a VM to not always be on the same host server because of features such as Distributed Resource Scheduler (DRS), HA, and VMotion. This can also cause headaches when licensing applications that are tied to specific hosts and hardware resources. It's best to contact all your software application vendors and find out their virtualization support policy in the early stages of your project. You might find that it may cost more to run their applications on virtual servers, but often the advantages that virtualization provides outweigh the increased costs.

Did You Know?

Put together a spreadsheet of all your applications and check with each vendor to find out their virtualization support policy. Also check to see how they license their products in a virtual environment. Often, you can find this information on their website or in their knowledge base. Include columns for support level, licensing model, and the URLs to their policies. After you complete your spreadsheet, meet with the application owners to discuss the results and make sure they understand the vendor's policy toward virtualization.

Getting Everyone On Board with Virtualization

Virtualization introduces many unique and new concepts into your environment, and as a result many groups within IT often put up resistance to it. This is typically a result of the fear and mistrust of a nonstandard technology like VMware and is most often caused by the lack of understanding of how VMware works. Therefore, plan to educate everyone who will be involved in your virtualization project early on so that they have a good understanding about what VMware is and its capabilities and features. Once they learn more about it and discover the great benefits of virtualization, they will be better prepared and more willing to help you implement it. This section provides examples of the types of resistance you will experience from each group and how you can best deal with each of them. You will often find that most people who are initially negative toward VMware eventually become supporters after they learn more about it and experience it for themselves.

Did You Know?

Before you attempt to explain virtualization concepts to other groups, make sure you understand the product thoroughly yourself. First download an evaluation copy, install and configure it, and make sure to read through the documentation. Also contact VMware or a business partner and have them assist you by presenting their product to your company.

The following subsections provide some tips for dealing with specific groups within IT to help them better understand virtualization concepts.

Dealing with Network Administrators

Traditionally, most network groups manage the physical network connection of a server from the switch all the way to the NIC. Virtualization changes that with vSwitches, which

effectively extend the physical network from the NIC in an ESX host to a vSwitch that is managed by the ESX server and a virtual NIC that connects a VM to the vSwitch. This vSwitch is usually managed by ESX administrators and not network administrators, which can cause some concern among network administrators because they can no longer control and manage part of the network that connects a virtual server to a physical network.

802.1Q VLAN tagging is a network technology commonly used when virtualizing servers. It enables you to use multiple VLANs on a single vSwitch and is a must-have in large environments. Without it, you would have to create a separate vSwitch for each VLAN and dedicate at least one NIC to it. This technology is not used that often with physical servers, and some network people might not have much experience with it. It's fairly simple to set up and configure, and we cover more on this in a later chapter.

Another networking area that is often a concern with virtualization is connecting VMs to your public demilitarized zone (DMZ) while keeping your ESX service console on your private internal network. The concern with this is that the ESX server is straddling the DMZ, because it has connections to both the private and public networks, and a potential attacker could compromise a VM in the DMZ and gain access to your internal network. The design of ESX does not allow for this to occur, and the only scenario in which this could potentially happen is if someone mistakenly configured a VM with two virtual NICs (vNICs), one being on an internal network vSwitch and the other on an external network vSwitch, which you would never want to do (unless the VM is acting as a firewall or proxy server).

Here's what you should tell your network administrators:

- Explain the concept of vSwitches and vNICs and how they interact with physical switches and physical NICs.
- Show them how to set up and configure a vSwitch and how to install a vNIC in a VM and connect it to a vSwitch.
- Explain to them how ESX uses trunked network ports and how 802.1Q VLAN tagging works in a virtual networking environment.
- Explain virtual network security principles and how vSwitches are isolated from each other so that traffic cannot leak between them.
- Demonstrate NIC teaming and failover in a virtual switch.

By the Way

Putting together a pilot project is a great way to learn and experience virtualization and demonstrate its capabilities and potential. Consider a small-scale project using the 60-day evaluation licenses with one to two ESX or ESXi hosts on some existing hardware. You might also try using the free VMware Server, which can run on a wide variety of hardware and operating systems as a good introduction to virtualization.

Dealing with Developers

Many developers will be concerned that their applications may not run properly on virtual servers. Another concern may be that software vendors will not support their products running in virtual environments. Early on in your project, gather support statements from software vendors that show their level of support for virtualization. Demonstrate the snapshot and cloning features of VMware that will be a great benefit to them. Also explain what virtual hardware is and how VMs see the same hardware regardless of the underlying physical hardware (except for the CPU). By having consistent hardware on all servers, you can eliminate any potential problems that may be caused by using different hardware on different servers running the same applications.

Here's what you should tell your developers:

- Show them statements of support for VMware from software vendors.
- Show them a VM's hardware configuration.
- Explain how VM hardware can easily be modified (more RAM, more disk space, and so on).
- Tell them about VMware's capability to rapidly provision new servers and to have dedicated, isolated development sandboxes.
- Show them information about the Lab Manager and Stage Manager automation products that VMware offers as additional components to VI3.
- Demonstrate creating snapshots and reverting back and cloning existing VMs and creating new ones from templates.

Dealing with Security Administrators

This is the group that tends to put up the most resistance to VMware because of the fear that if a VM is compromised it will allow access to the host server and the other VMs on that host. This is commonly known as "escaping the cave," and is more an issue with hosted products such as VMware Workstation and Server and less an issue with ESX, which is a more secure platform.

By the Way

The term *escaping the cave* comes from the analogy that a VM is trapped inside a cave on the host server. Every time it tries to escape from the cave, it gets pushed back in, and no matter what it does, it cannot escape from the cave to get outside. To date, there has never been an instance of a VM escaping the cave on an ESX server.

ESX has a securely designed architecture, and the risk level of this happening is greatly reduced compared to hosted virtual products such as Server and Workstation. This doesn't mean it can't happen, but as long as you keep your host patched and properly secured, the chances of it happening are almost nonexistent. Historically, ESX has a good record when it comes to security and vulnerabilities, and in May 2008, ESX version 3.0.2 and VirtualCenter 2.0.2¹ received the Common Criteria certification at EAL4+ under the Communications Security Establishment Canada (CSEC) Common Criteria Evaluation and Certification Scheme (CCS). EAL4+ is the highest assurance level that is recognized globally by all signatories under the Common Criteria Recognition Agreement (CCRA).

Another concern when it comes to security is with storage logical unit numbers (LUNs). The concern is that a VM that has its virtual disk on a SAN LUN that is shared with other VMs may allow for an attacker to access other data on that LUN or on the SAN fabric itself. Again, the secure design of ESX specifically prevents this from being possible. A VM cannot directly access the Fibre Channel cards in a host system and therefore cannot see anything beyond the virtual disk assigned to it.

Here's what you should tell your security administrators:

- Show them the industry security certifications that ESX has achieved.
- Explain how the design of ESX does not allow for VMs to directly access host hardware without going through the hypervisor.
- Show them CIS ESX Host and Virtual Machine Benchmarks and Guidelines.
- Allow them access to a VM so that they can verify its security for themselves.
- Show them the vCenter Server roles and very granular permissions that control access to the ESX hosts and VMs.
- Explain the ESX Service Console is not just a Linux operating system but a modified and more secure version based on Red Hat Linux. In addition, explain how ESXi no longer has a Service Console and is less vulnerable than ESX.
- Explain that the guest operating system on a VM is subject to the same security risks as a physical system and if compromised does not allow access to the ESX host.

Dealing with Management

IT management groups are usually the ones that get your funding approved and are typically the ones that sponsor your project. It's important that they understand the technology and its benefits so that they can support you, ensure you get the appropriate funding, and promote your project within the rest of your company.

Here's what you should tell your management:

- Demonstrate some of the cool features that virtualization provides, such as snapshots, VMotion, NIC teaming, and HA.

- Explain the cost-savings benefits and ROI that virtualization can provide (greatly reduced power and cooling costs, for instance).
- Show them the many customer success stories that VMware provides on its website.
- Explain how virtualization can greatly simplify disaster recovery.
- Provide a high-level executive overview of the technology, its features, and how it works.

Dealing with Storage Administrators

Many storage administrators have their own ideas about designing and configuring storage and do not like deviating from them. The most frequent area of contention when deploying ESX is the size of the SAN LUNs. Some old-school storage administrators like creating smaller LUNs (for example, 20GB) and do not like creating the larger LUN sizes that work best with ESX. In addition, assigning storage to ESX servers is a bit different from traditional methods because the same LUNs must all be presented to every ESX server with the same LUN IDs because ESX servers must all see the same storage for features such as VMotion to work.

Here's what you should tell your storage administrators:

- Explain how the VMFS file system is a cluster file system that leverages shared storage to allow multiple instances of ESX Server concurrent read and write access to the same storage resources.
- Explain what a virtual disk file (vmdk) is and how they are used on VMFS volumes.
- Explain how VMFS volumes work best with larger LUNs, and how using extents to enlarge VMFS volumes across multiple LUNs is not a best practice.
- Show them the SAN Configuration guide and the SAN Design and Deployment guide that VMware provides.
- Explain the reduced SAN administration workload that results because there is no need to configure storage for each server (VM) individually; it's only necessary to configure it for each ESX host.
- Explain how ESX servers use multipathing to connect to the SAN fabric.

Dealing with Operating System Administrators

This group will typically be concerned with performance, compatibility, security, and manageability of their servers running in a virtual environment. The biggest concerns are typically resource contention and not knowing how their servers will perform on a virtual host.

Here's what you should tell your operating system administrators:

- Show them how templates work and will allow them to quickly and easily deploy new servers.
- Explain how resource pools work and how resource shares, limits, and reservations can help control the amount of host resources that a VM can access.
- Explain how the VI client and vCenter Server are used to administer ESX hosts and VMs and how roles and permissions are used to assign specific privileges to access both host servers and VMs.
- Explain the key differences between virtual and physical hardware and how the ESX scheduler handles CPU requests.

Did You Know?

Seeing is believing. Download and set up an evaluation version, and then demonstrate some of the advanced features, such as HA and VMotion. Also, unplug a NIC cable on a multi-NIC vSwitch to demonstrate NIC teaming. To get a visual indication of these features working, ping a VM from a separate workstation while you demonstrate. Seeing these features in action is a sure way to impress any naysayers.

Summary

Assessing your current environment is important if you are planning on migrating your current physical servers to VMs on your ESX hosts. You might be tempted to do some guessing in your haste to get started, but a proper assessment ensures that you do not run into surprises later on and that you plan on adequate hardware to support your needs. So take the time and see where you are today, before you progress to where you want to be tomorrow. In the next chapter, we cover all the many things that you will want to plan for when architecting your new VI3 environment.

Endnotes

1. The use of the term VirtualCenter here refers to the old name, which still applies to VirtualCenter version 2.0.x and is also still present inside the application in vCenter Server 2.5 (because the software has not yet been updated to reflect the new name).

Chapter 2

Planning Your Virtual Environment

Proper planning is critical to ensuring a successful virtualization project. To plan properly, you need to understand the many features and components of VI3. Try not to rush into your project and make quick decisions, because if you do so you might regret them later on, and they might be difficult to change afterward. When you have a good understanding of this, you can make the proper choices and decisions to ensure your success.

Selecting the Right Edition and Features

Choosing an ESX Edition

Several different versions of ESX are available. Currently, these versions are ESXi, Foundation, Standard, and Enterprise. The differences between the versions are the standard features that they come with. You can also purchase certain features à la carte (Distributed Resource Scheduler [DRS], High Availability [HA], and Storage VMotion) if you want a particular feature that exists in a different version but do not want all the features that version contains. Upgrade paths do exist if you decide to upgrade to a different version later (for example, Foundation to Standard or ESXi to Enterprise). Figure 2.1 illustrates the main differences between each version.

Resource Management Power Management		DRS DPM	
Live VM Migration Live VM Disk File Migration		VMotion™ Storage VMotion	
Availability		High Availability	High Availability
Backup	Consolidated Backup	Consolidated Backup	Consolidated Backup
Patch Management	Update Manager	Update Manager	Update Manager
Central Management	VC Agent	VC Agent	VC Agent
Storage Virtualization	VMFS Virtual SMP** VMware ESXi	VMFS Virtual SMP** VMware ESX 3i or VMware ESX	VMFS Virtual SMP** VMware ESXi or VMware ESX
Enterprise VMs			
Next Generation Hypervisor	VMware ESXi	VI Foundation	VI Standard

Figure 2.1 VI3 editions and features summary

Determining which features that you want to use in your environment will help you decide which ESX edition is right for you. Also, whether you will use a shared storage solution (for example, a storage area network [SAN]) will be a big factor in which edition you choose. If you intend to use only local storage on your ESX hosts, features such as VMotion, DRS, and HA will be useless to you. HA is a great feature that will benefit any environment that uses shared storage. Other features, such as DRS, Distributed Power Management (DPM), and Storage VMotion, are ones that you might be able to live without and are more geared for larger enterprise environments.

The Embedded Version, ESXi

ESXi is the embedded version of VMware that typically comes preloaded on certain models of servers. In July 2008, VMware announced that the base edition of ESXi would be available as a free download beginning with ESXi 3.5 Update 2. It's functionally the same as ESX, but the Service Console has been removed, reducing its total footprint to 32MB. ESXi is usually embedded via a bootable USB device located inside the server. This embedded version enables you to boot directly to the hypervisor instead of the traditional method that ESX uses booting to the Service Console first and then loading the hypervisor. ESXi eliminates the need for the full Service Console, which was based on Red Hat Linux and is instead replaced by a remote command-line interface (RCLI) and a POSIX environment called Busybox that provides a very limited remote interface. There is also an installable version of ESXi available that can be installed directly onto the hard drive of an existing system that does not have an internal USB port that is capable of being booted from. Figure 2.2 shows the ESXi startup screen compared to the traditional ESX startup screen shown in Figure 2.3.




Figure 2.2 ESXi startup screen



Figure 2.3 ESX startup screen

Common Confusion about ESXi

Many people are often confused as to which edition they should choose. VMware currently has four editions available; ESXi is considered the entry-level edition with limited features. The editions above that are Foundation, Standard, and Enterprise. If you already have ESXi, you must upgrade to one of these other editions to gain the additional features that they provide. If you do upgrade, you're not technically using the ESXi edition anymore. However, the

Foundation, Standard, and Enterprise editions all allow you to use either ESXi or the full-fledged ESX server. ESXi is available as an option with any of these other editions; so ESXi is not really an edition, it's just the same ESX hypervisor with a different architecture that doesn't include a Service Console. Table 2.1 compares some features and characteristics of ESX and ESXi.

Table 2.1

ESX Versus ESXi Comparison

	ESX	ESXi
Thick client access (VI Client)	Supported for VM (VM) or host server management by connecting to an ESX host or vCenter Server.	Supported for VM or host server management by connecting to an ESX host or vCenter Server.
RCLI	Can connect to the ESX Service Console using an SSH client. ESX 3.5 Update 2 also supports RCLI.	Can connect to ESXi using the RCLI application or virtual appliance. The RCLI is limited to read only access for configuration changes (read/write is allowed for other operations) unless the host is licensed for Foundation, Standard or Enterprise.
Web client access	Supported for VM management by connecting to an ESX host or vCenter Server.	Not supported.
Service Console	The ESX Service Console is based on a Red Hat Linux installation. Default access is through the root user and provides a method for accessing the file system, running scripts, and executing standard Linux commands.	The traditional ESX Service Console is removed; instead, ESXi uses a management interface based on the very small Busybox shell that combines multiple utilities into a single executable. VMware has provisioned RCLI to allow the execution of scripts.
Scriptable installations	VMware ESX supports scriptable installations through utilities like KickStart.	VMware ESXi Installable currently does not support scriptable installations like ESX does. VMware ESXi does provide support for post-installation configuration scripts using RCLI-based configuration scripts and taking advantage of the VMware PowerShell toolkit.

Selecting the Right Edition and Features

	ESX	ESXi
Patches and updates	VMware ESX software patches and updates behave like traditional Linux-based patches and updates. The installation of a software patch or upgrade might require multiple system boots as the patch or upgrade may have dependencies on earlier patches or upgrades.	VMware ESXi patches and updates behave like firmware patches and updates. Any given patch or update is all-inclusive of earlier patches and updates. Because ESXi is smaller and does not have a Service Console, it also requires fewer patches than ESX.
Built-in firewall	A built-in firewall protects the Service Console and is more complex, with more than a dozen inbound and outbound connection types allowed by default.	The built-in firewall is much simpler because there is no Service Console to protect. Only two connection types are allowed by default.
Hardware monitoring	Vendor-specific agents are installed on the Service Console agents to provide hardware statistics and status. Currently, these agents do not integrate into vCenter Server and must be viewed and configured through the vendor's interface.	VMware ESXi provides hardware instrumentation through Common Information Model (CIM) providers. Standards-based CIM providers are distributed with all versions of VMware ESXi. Server hardware health monitoring is built in to vCenter Server.

Why choose ESXi over ESX?

- The base version of ESXi is available for free.
- It's OEM installed on many server models from vendors such as HP, IBM, and Dell.
- The lack of the Linux-based Service Console makes it more secure and require less patching than ESX.
- It boots faster than ESX because of its reduced footprint.
- Hardware management agents are preinstalled with ESXi, and the data is displayed directly in the VI Client; no need to install them separately and use a separate management console to display hardware information. These agents can also integrate with third-party tools such as HP's Systems Insight Manager.
- The Service Console can increase administration and require additional training because it's based on Linux. Windows administrators may not have experience working with Linux.

Why choose ESX over ESXi?

- The Service Console can be very useful for performing command-line operations and running scripts.

- Additional components (for example, backup agents) can be installed inside the Service Console. This is not recommended, but is sometimes necessary.
- Currently, ESXi supports only a limited range of hardware vendors/models for its CIM extensions, which monitor the host's hardware. This support is expected to increase and expand in the future.
- Both ESX and vCenter Server support integration with Active Directory. The free stand-alone version of ESXi does not support this.
- ESX supports booting from a SAN, ESXi currently does not.
- Centralized management through vCenter Server. Entry-level ESXi cannot be managed by vCenter Server because it does not include a virtual circuit (VC) agent and does not include a web interface for VM management.

By the Way

The following features all require shared storage for them to function: HA, DRS, DPM, VMotion, Storage VMotion, and Consolidated Backup (although the latest version of Consolidated Backup now supports backing up VMs on local disk).

vCenter Server

All ESX hosts can be managed individually through the provided Windows-based Virtual Infrastructure Client (VI Client). vCenter Server provides centralized management of ESX hosts by allowing the VI Client to connect to the vCenter Server rather than to individual ESX hosts, as shown in Figure 2.4.




Figure 2.4 Centralized ESX host management with vCenter Server

In addition, vCenter Server is necessary to use many of the advanced features, such as VMotion, HA, and DRS. Because of its centralized management, vCenter Server is a must-have in large environments when you have many ESX hosts. Without vCenter Server, you would have to connect to each host separately, which would be very time-consuming and non-productive. Although aimed at larger enterprises, vCenter Server is just as good for environments with as few as two ESX hosts because of the advanced features and functionality that it offers. As of version 2.5, vCenter Server can manage up to 200 ESX hosts and 2,000 VMs.

vCenter Server is not included with ESX and must be purchased separately. VMware does offer some bundles aimed at the small business market, which makes vCenter Server more affordable in a package with some ESX host licenses. vCenter Server provides many additional advantages and features to ESX, as listed here. It's important to note that vCenter Server does not automatically provide all of these features to your ESX hosts. Depending on the edition of ESX that you are running, you may or may not be licensed for all the features:

- Ability to use templates to allow fast creation of new VMs
- Centralized management of ESX hosts
- Scheduling to automate performing certain tasks
- Alerts and notifications via alarm triggers that can perform actions such as sending emails, Simple Network Management Protocol (SNMP) traps, running scripts, and more
- Customized security roles and permissions
- Centralized licensing of hosts and features
- Integration with Active Directory for secure access to VMs and ESX hosts
- HA for VMs
- Resource optimization using the DRS
- Migration of live VMs across ESX hosts using VMotion
- Migration of live VMs across storage datastores using storage VMotion (new to vCenter Server 2.5)
- Automated VM and ESX host patching through Update Manager (new to vCenter Server 2.5)
- Automated energy efficiency using the DPM feature (new to vCenter Server 2.5)
- Ability to clone existing VMs
- Historical performance monitoring of VMs and ESX hosts
- Ability to analyze physical servers in your environment, plan for consolidation of good virtualization candidates, and convert physical machines to VMs (new to vCenter Server 2.5)
- Integrated physical to VM conversions (P2V) (new to vCenter Server 2.5)
- Cold migration of VMs between ESX hosts

ESX by itself has very limited historical performance monitoring; it can monitor statistics only in real time and going back for the last 60 minutes. vCenter Server expands on this and enables you to retain statistics for much longer time periods. In vCenter Server, you can configure the daily, weekly, monthly, and yearly sample rates and retention periods of the performance data. In addition, vCenter Server provides you with basic alarms to send alerts on certain events. These events are limited to CPU/memory/disk/network usage being above or below a certain percentage and to host/VM state.

An additional requirement for vCenter Server is a database in which to store its configuration and performance data. This database can be either a Microsoft SQL Server or an Oracle database. Chapter 3, “Building Your Virtual Environment,” discusses the vCenter Server database in detail. vCenter Server can be installed either on a physical system or on a VM. Each method has pros and cons (as discussed in Chapter 3). In addition, vCenter Server can be installed only on a Microsoft Windows Server platform. The minimum requirements for vCenter Server 2.5 are as follows:

■ Hardware

- ◆ Processor: 2.0 GHz Intel or AMD or higher.
- ◆ Memory: 2GB RAM.
- ◆ Disk: 560MB minimum, 2GB recommended. Additional disk will be needed if you choose to run the database on the same physical server as vCenter Server.
- ◆ Networking: Gigabit Ethernet recommended

Note: These are just minimums. For larger deployments, consider using a two-processor server, 4GB of RAM, and 20GB of disk.

■ Operating systems

- ◆ Windows 2000 SP4 with Update Rollup 1
- ◆ Windows XP Pro SP2
- ◆ Windows Server 2003 SP1/SP2/R2

Note: vCenter Server will run on only 32-bit versions of these operating systems. 64-bit versions are currently not supported.

■ Databases

- ◆ Microsoft SQL Server 2000 Standard SP4
- ◆ Microsoft SQL Server 2000 Enterprise SP4
- ◆ Microsoft SQL Server 2005 Standard SP1 or SP2
- ◆ Microsoft SQL Server 2005 Enterprise SP1 or SP2
- ◆ Microsoft SQL Server 2005 Express SP2
- ◆ Oracle 9i Release 2 Standard or Enterprise (9.2.0.8.0)

- ◆ Oracle 10g Release 1 Standard or Enterprise (10.1.0.3.0)
- ◆ Oracle 10g Release 2 Standard or Enterprise (10.2.0.1.0)

Note: SQL Server Express is recommended only for small deployments of up to 5 hosts and 50 VMs.

The VI Client that comes with both ESX stand-alone, and vCenter Server runs on only the following 32-bit versions of these Windows operating systems:

- Windows 2000 Pro SP4
- Windows 2000 Server SP4 with Update Rollup 1
- Windows XP Pro SP2
- Windows 2003 Server SP1 and SP2 (all releases except 64-bit)
- Windows 2003 Server R2
- Windows Vista Business and Enterprise

In addition, the VI Client requires the Microsoft .NET 2.0 Framework. ESX and vCenter Server can also be managed using a web access UI that runs on both the ESX host servers and vCenter Servers. However, this web access UI is limited and allows you to perform only basic VM operations (for example, power on, power off, remote console) and not manage the ESX hosts themselves or create, edit, or delete VMs. The following web browsers are supported with the web access UI:

- **Windows.** Internet Explorer 6.0 or later, Netscape Navigator 7.0, Mozilla 1.x, Firefox 1.0.7 and later
- **Linux.** Netscape Navigator 7.0, Mozilla 1.x, Firefox 1.0.7 and later

Did You Know?

A Linux version of the VI Client has been a much-requested feature for quite some time, but as of ESX 3.5/vCenter Server 2.5 it still is not available. Therefore, a Windows machine is required to manage both ESX hosts and vCenter Server.

VMotion

VMotion is another powerful feature that enables you to quickly move an entire running VM from one ESX host to another without any downtime or interruption to the VM, as shown in Figure 2.5. This is also known as a “hot” or “live” migration.




Figure 2.5 VM migrating from one ESX host to another through VMotion

The entire state of a VM is encapsulated, and the VMFS file system allows both the source and the target ESX host to access the VM files concurrently. The active memory and precise execution state of a VM can then be rapidly transmitted over a high-speed network. The VM retains its network identity and connections, ensuring a seamless migration process, as outlined here:

1. The migration request is made to move the VM from ESX1 to ESX2.
2. vCenter Server verifies that the VM is in a stable state on ESX1.
3. vCenter Server checks the compatibility of ESX2 (CPU, networking, and so on) to ensure that it matches that of ESX1.
4. The VM is registered on ESX2.
5. The VM state information (including memory, registers, and network connections) is copied to ESX2. Additional changes are copied to a memory bitmap on ESX1.
6. The VM is quiesced on ESX1, and the memory bitmap is copied to ESX2.
7. The VM is started on ESX2, and all requests for the VM are now directed to ESX2.
8. A final copy of the VM's memory is done from ESX1 to ESX2.
9. The VM is unregistered from ESX1.
10. The VM resumes operation on ESX2.

VMotion requires shared storage for it to function and has additional requirements to ensure compatibility of a VM moving from one ESX host to another, as outlined here:

- Both the source ESX host and destination ESX host must be able to access the same shared storage that the VM is located on. The shared storage can be either a Fibre Channel, iSCSI, or NFS.
- ESX hosts must have a Gigabit Ethernet network adapter to be configured on the VMkernel virtual switch (vSwitch) used by VMotion. For best results and because VMotion traffic is sent clear text, it is best to have an isolated network for VMotion traffic.

- VMotion will also work with RDMs (raw device mappings) as long as they are configured to work in virtual compatibility mode.
- ESX hosts must have processors that can execute the equivalent instructions of each other. Processor clock speeds, cache sizes, and number of cores can differ between ESX hosts, but they must have the same processor vendor class (Intel or AMD) and compatible feature sets.

By the Way

It is possible to override these restrictions, but by doing so it is possible to cause a VM to crash (due to it accessing a CPU feature or instruction that the new ESX host does not support). You can also hide certain CPU-specific features (for example, NX flag for AMD processors or the XD flag for Intel processors) from a VM to increase VMotion compatibility between ESX hosts. In addition, ESX 3.5 Update 2 introduces Enhanced VMotion Compatibility (EVC) to further ensure compatibility between ESX hosts. EVC leverages the Intel FlexMigration technology and the AMD AMD-V Extended Migration technology to present the same feature set as the baseline processors. This feature will still not allow you to migrate VMs from an Intel CPU host to an AMD host. Therefore, you should create clusters with ESX hosts of the same processor family only or choose a processor vendor and stick with it.

Did You Know?

Although you can't hot migrate (VMotion) a running VM from an Intel host to an AMD host or vice versa, you can cold migrate it by shutting it down first.

Additional requirements for VMotion to function properly include the following:

- The vSwitch network labels (port groups) must match *exactly* (including case) on each ESX host.
- A VM cannot be using CPU affinity, which pins a VM to run on a specific processor on an ESX host.
- A VM cannot be connected to an internal-only (no NICs assigned to it) vSwitch.
- A VM cannot have its virtual CD-ROM and floppy drive mapped to either a host device or a local datastore ISO file.

Storage VMotion

Storage VMotion is a new feature introduced in ESX 3.5. It enables you to migrate a running VM and its disk files from one datastore to another on the same ESX host, as shown in Figure 2.6.




Figure 2.6 VM migrating from one ESX datastore to another through storage VMotion

The difference between VMotion and Storage VMotion is that VMotion simply moves a VM from one ESX host to another but keeps the storage location of the VM the same. Storage VMotion, on the other hand, changes the storage location of the VM while it is running and moves it to another datastore on the same ESX host. The VM can be moved to any datastore on the ESX host, which includes local and shared storage. The storage VMotion process is as follows:

1. A new VM directory is created on the target datastore. VM configuration files and all nonvirtual disk files are copied to the target directory.
2. ESX host does a “self” VMotion to the target directory.
3. A snapshot (without memory) is taken of the VMs disks in the source directory.
4. VM disk files are copied to the target directory.
5. The snapshot that is located in the source directory is consolidated into the VM disk files located in the target directory.
6. The source disk files and directory are deleted.

As of ESX 3.5 Update 2, Storage VMotion is not built in to the VI Client but must instead be run using the RCLI and executing the `svmotion` command with the required parameters. Alternatively, you can do this using the VI Client by downloading one of the third-party

vCenter Server plug-ins that have been created specifically for performing storage VMotions. These plug-ins make the process much simpler than using the RCLI. VMware will most likely integrate storage VMotion into the VI Client in a future release. You should be aware of several requirements and limitations for using storage VMotion, including the following:

- VM disks must be in persistent mode or be an RDM that is in virtual compatibility mode.
- If a VM has any snapshots, it cannot be migrated.
- The ESX host that the VM is running on must be licensed for VMotion and must be configured to use VMotion.
- The ESX host that the VM is running on must have access to the source and target datastores.
- The ESX host that the VM is running on must have enough resources available to support two instances of the VM running at the same time.

High Availability

High Availability (HA) is one of ESX's best features and is a low-cost alternative to traditional server clustering. HA does not provide 100% availability of VMs, but instead provides higher availability by rapidly recovering VMs on failed hosts on to other ESX hosts. The HA feature continuously monitors all ESX Server hosts in a cluster and detects failures and will automatically restart VMs on other host servers in an ESX cluster in case of a failure of a host. An agent placed on each host maintains a "heartbeat" with the other hosts in the cluster, and loss of a heartbeat initiates the process of restarting all affected VMs on other hosts. The vCenter Server does not provide a single point of failure for this feature, and it will continue to work even if the vCenter Server is unavailable. If the vCenter Server goes down, HA functionality changes as follows: HA clusters can still restart VMs on other hosts in case of failure; however, the information about which extra resources are available will be based on the state of the cluster before the vCenter Server went down.

HA monitors whether sufficient resources are available in the cluster at all times to be able to restart VMs on different physical host machines in the event of host failure. Safe restart of VMs, as illustrated in Figure 2.7, is made possible by the locking technology in the ESX Server storage stack, which allows multiple ESX Servers to have access to the same VMs file simultaneously.

With earlier versions of ESX 3.0.x, the failure detection interval was fixed at 15 seconds. Starting with ESX 3.0.2/vCenter Server 2.0.2, this interval can be modified via a custom setting in vCenter Server. A host failure is detected after the HA service on a host has stopped sending heartbeats to the other hosts in the cluster. A host stops sending heartbeats if it is either isolated from the network or if the host crashes or is completely down due to a hardware failure. When a failure is detected, other hosts in the cluster treat this host as failed,

while this host declares itself as isolated from the network. By default, the isolated host powers off its VMs, but the isolation response for each VM is configurable on a per-VM basis. These VMs can then successfully fail over to other hosts in the cluster. HA also has a restart priority that can be set for each VM so that certain VMs are started before others. This priority can be set to low, medium, or high; it can also be disabled so that VMs are not automatically restarted on other hosts.




Figure 2.7 Host failover using VMware HA

The HA feature was enhanced starting with ESX 3.5 and now provides VM failure monitoring in case of operating system failures (such as Windows blue screens). If an operating system failure is detected, the VM will automatically be reset on the same host so that its operating system is restarted. This new functionality allows HA to also monitor VMs via a heartbeat that is sent every second when using VMware Tools and further enhances HA's capability to recover from failures in your environment.

Distributed Resource Scheduler

The Distributed Resource Scheduler (DRS) is a powerful feature that enables your virtual environment to automatically balance itself across your ESX host servers in an effort to eliminate resource contention. It utilizes the VMotion feature to provide automated resource optimization and automatic migration of VMs across hosts in a cluster, as illustrated in Figure 2.8. DRS also provides automatic initial VM placement on any of the hosts in the cluster and makes automatic resource relocation and optimization decisions as hosts or VMs are added or removed from the cluster. You can also configure DRS for manual control so that it provides recommendations that only you can review and carry out.

DRS works by utilizing resource pools and clusters that combine the resources of multiple hosts into a single entity, as shown in Figure 2.9. Multiple resource pools can also be created so that you can divide the resources of a single or multiple hosts into separate entities. Currently, DRS will migrate VMs based only on the availability and utilization of the CPU and memory resources. It does not take into account high disk or network utilization to load balance VMs across hosts.




Figure 2.8 DRS global scheduler of clusters in vCenter Server




Figure 2.9 Resource aggregation in VMware clusters

Resource pools allow you to partition a host's or multiple hosts' CPU and memory resources into pools so that VMs can access only the resources assigned to the specific pool that they are a member of. Resource pools can also have a hierarchy with parent pools that can have child pools beneath them, to provide a more granular level of resource allocation. VMs that run in a resource pool are not tied to the particular ESX host on which they are running at any given point in time. When a VM experiences increased load, DRS first evaluates its priority against the established resource allocation rules and then, if justified, redistributes VMs among the physical servers to try to eliminate contention for resources. VMotion will then handle the live migration of the VM to a different ESX host, with complete transparency to end users. The dynamic resource allocation ensures that capacity is preferentially dedicated to the highest-priority applications, while at the same time maximizing overall resource utilization. Unlike the HA feature, which will

still operate when vCenter Server is unavailable, DRS requires that vCenter Server be running for it to function.

DRS makes its recommendations by applying stars to indicate how much the recommendation would improve the cluster's performance. The lowest rating (one star) indicates a slight improvement, whereas the highest rating (four stars) indicates significant improvement. Five stars, the maximum, indicates a mandatory move because of a host entering maintenance mode or affinity rule violations. If DRS is set to work in fully automated mode, you have the option to set a migration threshold based on the number of stars of the recommendations. The lowest threshold, which is the most conservative, applies resource allocation five-star recommendations; the highest threshold, which is very aggressive, applies all recommendations. There are also settings in between to only apply two-, three-, or four-star recommendations.

Affinity rules can also be created in the DRS settings to keep certain VMs together on the same ESX host or to ensure that certain VMs are kept on separate ESX hosts from each other. Automation levels can also be set on individual VMs: Manual (initial placement is the host it was last located on, migration recommendations require approval), Partially Automated (initial placement is automated, migration recommendations require approval), Fully Automated (initial placement is automated, migration recommendations automatically executed), or Disabled (the VM is never relocated by DRS).

Distributed Power Management

Distributed Power Management (DPM) is a subcomponent of DRS and is a new green feature that was introduced in ESX 3.5 that will power down ESX hosts during periods of inactivity to help save power. All the VMs on a host that will be powered down are first relocated to other ESX hosts before the host is powered down. When activity increases on the other hosts and DPM deems that additional capacity is needed, it will automatically power the host back on and move VMs back on to it using DRS. DPM requires that an ESX host has network cards that support the wake-on-LAN feature that enables a magic packet to be sent to a host to automatically power it on. DPM can be configured in either manual or automatic mode. In manual mode, it will simply make recommendations similar to DRS, and you will have to manually approve them before they are applied. You can also configure DPM so that certain host servers are excluded from DPM and you can specify that certain hosts are always automatic or always manual.

Virtual SMP

The virtual SMP (vSMP) feature enables you to assign more than one virtual CPU to a VM. Up to four virtual CPUs can be assigned to any VMs. Be careful when using this feature, however, because vSMP can actually slow down your VMs. (Read more about when to use vSMP at the end of this chapter.) Starting with ESX 3.5, the vSMP feature comes with all editions of

ESX and ESXi. Previously, with ESX 3.0.x, the vSMP feature was not available in the Foundation edition (previously known as Started edition).

Just because you can assign a VM more than one virtual CPU (vCPU) doesn't mean you always should. Bigger isn't always better when it comes to VMs. The decision to use more than one virtual processor in a VM should be based on an actual requirement by the applications installed on the VM and not simply because two or more processors sound better than one. Many physical servers commonly had multiple CPUs regardless of whether the applications running on them needed them. Besides being wasteful, this does not negatively impact a physical server, but most VMs will usually run better with one virtual processor and can actually run slower when more than one is assigned to it.

The reason for this is the hypervisor's CPU scheduler must find simultaneous cores available equal to the number assigned to the VM. So a four-vCPU VM will need to have four free cores available on the host for every CPU request that is made by the VM. If four cores are not available because other VMs are using them, the VM must wait until the cores become available. Single-vCPU VMs have a much easier time because there only needs to be a single core available for the scheduler to process CPU requests for it. Beginning with ESX 3.x, VMware introduced a new "relaxed" co-scheduling scheme that provides more flexibility and increases the number of scheduling opportunities available to the scheduler and thereby increasing overall system throughput. Regardless of this, you should still take care when creating multiple-vCPU VMs.

So when it comes down to deciding whether to use vSMP, use these rules as guidelines:

- It's best to start with one vCPU when configuring a VM and add more later if needed. If you know ahead of time that the VM will benefit from more than one CPU, start with two vCPUs first and add additional if needed.
- Limit the number of vSMP VMs on your hosts. The fewer you have, the better all your VMs will perform.
- Only assign a VM multiple vCPUs if you are running an application that requires it and will make use of them.
- Don't assign a VM the same number of vCPUs as your host system has total cores available.
- If you are going to use vSMP, have at least twice (preferably three or four times) the number of cores available on your host system than that of your VM with the most vCPUs. So if you have a 4-vCPU VM, have at least 8 cores available on your host server, and preferably 16.
- If you are converting a multi-CPU physical Windows server to a single-vCPU VM, make sure you change the hardware abstraction layer (HAL) from multiprocessor to uniprocessor. Likewise, if you are adding a vCPU to a single-vCPU VM, make sure you change the HAL of a Windows server from uniprocessor to multiprocessor.

- Don't use CPU affinity, because it restricts the scheduler and makes it harder to process CPU requests. The scheduler is good at what it does, and so it's best to just let it do its job. In addition, when you use CPU affinity on your VMs, they can no longer be VMotioned to another host server.

Update Manager

Update Manager is a new feature introduced in ESX 3.5 that provides automated patching of ESX hosts and select Microsoft Windows and Linux VM operating systems and applications. Update Manager can scan ESX hosts and certain guest operating systems and compare them to a baseline and then apply updates and patches to them. Update Manager is integrated with DRS, and so ESX hosts can be patched without affecting VMs (as long as the VMs are on shared storage). It can also scan and remediate both online and offline VMs, as illustrated in Figure 2.10. An additional feature will automatically snapshot a VM before patching it (in case the patch causes a problem with the guest operating system and needs to be rolled back).




Figure 2.10 Patching ESX hosts and online and offline VMs with Update Manager

Update Manager is installed as a plug-in component to vCenter Server. After it has been installed, you can begin creating baselines and configuring Update Manager to patch your virtual environment. The process for using Update Manager is as follows:

1. The latest patches are gathered via the Internet according to the schedule defined in the Update Manager settings.
2. Baselines are created by collecting information about service packs, patches, and updates. These baselines can be either static, which are manually defined, or dynamic, which are set based on the urgency of the vendor patch.

3. The ESX hosts and VMs are then compared to the baselines. Scans can be either scheduled or run on demand and can be run against individual or groups of hosts and VMs. When the scan is complete, ESX hosts and VMs that do not match the baseline are flagged for patch updates.
4. Finally, the ESX hosts and VMs are remediated as needed so that they match the baseline. Patching can be done manually or by a schedule. For patches that require reboots, you can configure the settings to reboot immediately or up to 60 minutes after the patch has been installed.

Consolidated Backup

VMware Consolidated Backup (VCB) is a Windows-based application that provides a centralized backup facility to back up VMs through a proxy server without affecting the VM itself. It is an alternative to traditional agent-based backup methods, and was designed specifically for virtual environments to minimize host server impact during backup operations. VCB is an enablement technology and cannot back up VMs by itself but instead works with third-party backup products to help offload backup overhead from the VMs. It integrates with most major software backup providers and eliminates the need to back up VMs over the network when using a Fibre Channel SAN. It works by taking a snapshot of a VM to suspend writes to the original disk and then copying the original disk to a proxy server, where the VMs disk is mounted on the backup proxy server. It then backs up the contents of the VM as a virtual disk image or as a set of files and directories and then unmounts the virtual disk and deletes the snapshot. Figure 2.11 shows this process.




Figure 2.11 VMware Consolidated Backup enables off-host backup of VMs.

Some key features of Consolidated Backup are as follows:

- Supports both file-level full and incremental backups for VMs running Windows and image-level backups for VMs running any operating system (Windows and Linux)
- Eliminates the need for installing a software backup agent on every VM that you want to back up
- Reads virtual disks directly over the SAN when using Fibre Channel or hardware iSCSI or through the ESX server I/O stack for network attached storage (NAS)
- Supported by most major backup software

Consolidated Backup had some major enhancements in ESX version 3.5, which greatly increased its usability, including the following:

- Support for backing up VMs that are located on iSCSI, NAS, and local storage on an ESX host. Previously, only VMs located on a Fibre Channel SAN could be backed up.
- Support for copying the backup snapshots of VMs over the LAN. Previously, the snapshots were copied over the SAN fabric.
- Support for running the proxy server inside a VM. Previously, the proxy server had to be run on a physical server. This support is only for backing up VMs on iSCSI, NAS, or local storage. VMs located on Fibre Channel SAN storage still require VCB to be installed on a physical server that has access to the SAN network.
- Easier restores when utilizing VMware Converter to restore Consolidated Backup images.
- Better operating system support for the proxy server, including 64-bit Windows Server 2003.

Consolidated Backup is basically a set of utilities and scripts that work together with third-party backup software. Most backup software vendors provide integration modules containing pre-backup and post-backup scripts. These integration modules, scripts, backup software, and utilities all run on the Windows server that is used as the VCB proxy. The backup process when using Consolidated Backup is as follows:

1. The backup software calls a pre-backup script that quiesces the VM, puts the VM into snapshot mode, and then un-quiesces the VM. Optionally, pre-freeze and post-thaw custom scripts can be run to further prepare the VM for backup.
2. The VM's original disk snapshot is then made available to backup software, while all new writes occur to the snapshot file that was created. The backup software performs a normal backup of the VM to its physical media.
3. The backup software then calls a post-backup script that unmounts the VM's original disk snapshot from the backup proxy and then deletes the VM snapshot, which writes all changes that were made since the snapshot was taken back to the original disk.

Consolidated Backup is not required to back up your VMs, but is a great alternative to help save time and host resources. If you choose not to use Consolidated Backup, you can still use alternative backup solutions such as traditional backup agents installed on your VMs, custom scripts to copy virtual disk files to host local disk or other servers, or even install a tape device on the ESX server (not recommended, but you can do this).

ESX, vCenter Server, and VM Configuration Maximums

Many configuration limits exist in VMware Infrastructure that you should be aware of when planning an implementation. These limits typically change when new versions of vCenter Server and ESX are released. Tables 2.2 to 2.6 cover some of the limits as of ESX release 3.5 and vCenter Server 2.5. It is always best to check the latest configuration maximum guide that VMware publishes, in case the maximums have been increased with a newer release.

Table 2.2

VM Maximums		
Item	ESX 3.0.x and VirtualCenter ¹ 2.0.x	ESX 3.5.x and vCenter Server 2.5.x
Number of virtual CPUs per VM	4	4
Size of RAM per VM	16GB	64GB
Number of NICs per VM	4	4
SCSI controllers per VM	4	4
Devices per SCSI controller	15	15
Size of SCSI disk	2TB	2TB
Number of IDE devices per VM	4	4
Number of floppy devices per VM	2	2
Number of parallel ports per VM	2	3
Number of serial ports per VM	2	4
Number of remote consoles to a VM	10	10

Note: Can have a combined total of only 4 NIC/SCSI controllers per VM (for example, 3 NICs, 1 SCSI)

Table 2.3

ESX Server Host Maximums: Storage

Item	ESX 3.0.x and VirtualCenter ¹ 2.0.x	ESX 3.5.x and vCenter Server 2.5.x
VMFS block size (MB)	8MB	8MB
RDM size (TB)	2TB	2TB
Recommended number of hosts that can share a VMFS volume while running VMs against that volume	32	32
Number of hosts per cluster	32	32
Number of VMFS volumes configured per server	256	256
Number of extents per VMFS volume	32	32
Number of host bus adapters (HBAs) of any type	16	16
Number of targets per HBA (iSCSI HBA)	15 (64)	15 (64)
VMFS3 extent size	2TB	2TB
VMFS3 volume size	64TB (2TB × 32 extents)	64TB (2TB × 32 extents)
Fibre Channel LUNs per server	256	256
Fibre Channel LUN size	2TB	2TB
Fibre Channel number of paths to a LUN	32	32
Number of NAS datastores	32	32
iSCSI LUNs per server	256	256
Hardware iSCSI initiators per server	2	2
iSCSI targets	64	64

Note: The number of hosts per cluster was 16 in versions earlier than ESX 3.0.2.

Table 2.4

ESX Server Host Maximums: CPU and Memory

Item	ESX 3.0.x and VirtualCenter 2.0.x ¹	ESX 3.5.x and vCenter Server 2.5.x
Number of virtual CPUs per server	128	128
Number of cores per server	32	32
Number of (hyperthreaded) logical processors per server	32	32
Number of virtual CPUs per core	8	8
Size of RAM per server	64GB	256GB
RAM allocated to service console	800MB	800MB

Note: Starting with ESX 3.5 Update 2, support for up to 192 virtual CPUs per host server as long as the maximum number of VMs per host do not exceed 170.

Table 2.5

ESX Server Host Maximums: Networking

Item	ESX 3.0.x and VirtualCenter 2.0.x ¹	ESX 3.5.x and vCenter Server 2.5.x
Number of e100 Physical NICs	26	26
Number of e1000 Physical NICs	32	32
Number of Broadcom physical NICs	20	20
Number of port groups	512	512
Number of NICs in a team	32	32
Number of Ethernet ports	32	32
Number of virtual NICs per virtual switch	1,016	1,016
Number of vSwitches	127	127
Number of port groups (VLANs)	4,096	4,096

Table 2.6

vCenter Server Maximums

Item	ESX 3.0.x and VirtualCenter 2.0.x ¹	ESX 3.5.x and vCenter Server 2.5.x
Number of VMs	1,500	2,000
Number of hosts per DRS cluster	32	32
Number of hosts per HA cluster	16	32
Number of hosts per vCenter Server	100	200

Note: The number of hosts per DRS cluster was 16 in versions earlier than ESX 3.0.2.

Choosing Server Hardware

You will have to make many choices when selecting the hardware that you will use for your virtual infrastructure. You'll need to choose a server brand and models, decide between blades or traditional servers, select a storage type to use, and much more. In this section, we discuss the many hardware choices you will face and provide you with the information that you need to make an informed decision.

Determining How Many ESX Hosts You Will Need

Based on your analysis of your current environment, you should have a general idea about your current resource requirements. If you choose to use Capacity Planner, as discussed in the previous chapter, it will provide you with hardware estimates based on the performance data that it collected. If you do not use Capacity Planner, you can use the rules listed here to estimate your hardware needs:

- **Host servers.** CPU and memory requirements of your existing physical servers will be the primary determining factor for how many ESX host servers you will need in your environment. A secondary factor will be your I/O utilization requirements. Having high I/O requirements of a particular server resource can limit the number of VMs that can be placed on a host server.
- **CPU.** How busy each physical system's CPU is will help determine how many VMs you can put on an ESX host based on the total number of processor cores that it has. A general rule of thumb is that a single core can support four single-vCPU VMs. So, a dual-socket/core host system can generally support about 16 VMs. This number will vary based on the applications that your VMs will run and how busy they will be. On

servers that run CPU-hungry applications, like Exchange and SQL Server, that require multiple vCPUs, you may be able to support only one or two VMs per core; and on servers that run CPU-light applications, like web/file/print servers, you may be able to support up to eight single-vCPU VMs per core. Just because a physical server has multiple CPUs does not mean that you should assign it multiple CPUs in a virtual environment.

- **Memory.** The actual amount of physical memory that a server uses will determine how much memory your ESX hosts will need to be able to support all the VMs on it. Memory is typically the first resource to be completely utilized on ESX hosts because a lot of it is typically needed by the operating systems and applications running on the VMs. Therefore, you should try to get the maximum amount of RAM that you can in your ESX hosts. It is possible to overcommit an ESX host (assigning VMs more memory than the host physically has), but it is not recommended in most cases because it will degrade the performance of your VMs if you exhaust all physical host memory and the VMs start swapping to their VSWP file. ESX uses advanced memory features to help reduce the total amount of physical host memory used by VMs by implementing such features as Transparent Page Sharing (TPS). TPS looks for common memory pages among VMs and reduces the amount of physical memory that is needed. For example, Windows servers load many of the same common operating system files in memory. Rather than allocate a physical memory page for each VM, an ESX host will allocate a single memory page for all the VMs and thus reduce the amount of physical memory that the ESX host uses.
- **Disk.** Based on the current amount of disk usage, you should be able to calculate the amount you will need on your ESX hosts and storage subsystems. It's best to look at the total amount of disk in use on your servers that you want to migrate to ESX and not the combined total of free and used space. When you migrate physical servers to VMs, you can specify the size of the virtual hard disk for each VM. It's best to only give the VMs the disk they need and not excessively more. So, if you have a physical Exchange server with a total of 160GB of disk but are using only 40GB, you should consider reducing the size of the disk when moving it to a VM. If you do end up needing more disk on that VM later, it is simple to add more to it.

Besides a VM's disk file, you should also allow extra space on your ESX VMFS volumes for additional files like logs, VSWP files, and snapshots. You can find information about these additional files in Chapter 11, "Advanced Topics." How much space you will need for this will depend mainly on how many snapshots you plan on taking and how long you need to keep them. New snapshots start out small (16MB) and grow as changes to the VM's disk are made; they cannot grow larger than the original disk, however. Virtual machine log files are typically small and will not take more than a couple megabytes of space. VSWP files by default will take out as much disk space as you have assigned to the VM; so if you assign 4GB of RAM to it, a 4GB VSWP file will be created when the VM is started.

Leaving Enough Room for Future Potential Growth

When you are sizing your ESX hosts to fit your current needs based on the current workload that you want to virtualize, you should consider leaving room for any potential new VMs that you may need. You typically want to get about 75% to 80% resource usage on your ESX host servers, which is the reason for implementing virtualization, to consolidate and maximize resource utilization. It's not always easy to predict future needs and requirements, but you should make an attempt to find out what additional servers may be needed within the next year. VM sprawl is also common in a virtual environment, so you may find that the unused resources on your ESX hosts disappear fairly quickly after they are put into operation. You can read more about VM sprawl and how to prevent it at the end of this chapter.

If you plan to use the HA feature, leave some extra room on your ESX hosts for additional capacity (in case of a host failure). If you do experience host failure, the remaining ESX hosts will need to be able to have resources available to pick up the load as VMs are started on them. If you are also running DRS, then it will help distribute the load of the VMs from the failed hosts evenly among the ESX hosts that remain.

You should also consider using resource pools to split up, limit, and protect resources from being used by all VMs. Having multiple resource pools is a good way to dedicate a limited amount of resources to a group of VMs running on one or multiple ESX hosts.

Hardware Compatibility Lists

Before you purchase new hardware or use existing hardware for your ESX hosts and storage solutions, ensure that the brand and model of the hardware is on the current hardware compatibility list (HCL) for the version of ESX that you are implementing. ESX has a limited set of device drivers installed with it, and therefore a lot of hardware will not work with it. Most of the latest server models from all the major hardware vendors are on the list. The HCL contains information about all the hardware that is supported to use with ESX. It includes servers, I/O adapters, storage devices, network adapters, backup software, and storage adapters:

- **I/O compatibility guide.** Lists all supported network and storage adapters
- **Systems compatibility guide.** Lists all supported server brands and models
- **Storage/SAN compatibility guide.** Lists all supported Fibre Channel SAN, iSCSI, and NAS devices
- **Backup software compatibility guide.** Lists all supported backup software that is used to back up ESX hosts and VMs

Anything on the HCL is officially supported by VMware to use with ESX. This list is constantly updated, and new hardware is added and older hardware is removed. However, just because something is not on the HCL does not mean it will not work with ESX; many older

server models that have been removed from the list will still work just fine with ESX. What this means is if you are using hardware that is not on the HCL for the version of ESX you are running and you call VMware for support, there is a chance that they will not provide you with support. Also, because the HCL is always changing, you should ensure that your hardware is listed on the HCL for any versions of ESX that you may be considering upgrading to.

Blades Versus Traditional Servers

Another hardware decision to be made when purchasing ESX host servers is whether to use blade servers or traditional rack-mount servers. Often, this decision is dictated by whatever type you currently use in your datacenter. Each type has its advantages and disadvantages, with blade servers providing overall better rack density and a smaller footprint. Let's cover some reasons why you might choose one over another.

Why choose blade servers over traditional servers?

- Rack density is better for datacenters where space is a concern. Up to 50% more servers can be installed in a standard 42U rack compared to traditional servers.
- Blade servers provide easier cable management because they simply connect to a chassis and need no additional cable connections.

Why choose traditional servers over blade servers?

- Traditional servers have more internal capacity for local disk storage. Blade servers typically have very limited local disk storage capacity.
- Traditional servers have more expansion slots available for network and storage adapters. Blade servers typically have very few or no expansion slots. ESX servers are often configured with many NICs to support the Service Console network, VMkernel network, NAS, and VMs networks. Also, to provide failover and load balancing, additional network adapters are needed.
- Once a chassis is full, the cost of a new chassis to add a single new additional server can be costly. Traditional servers can be installed without any additional infrastructure components.
- Traditional servers are often less complicated to set up and manage than blade servers.
- Traditional servers have multiple USB ports (although the use of USB ports is not currently supported in ESX) for connecting external devices, and an optical drive for loading software on the host. They also have serial and parallel ports, which are sometimes used for hardware dongles that are used for licensing software. In addition, tape backup devices can be installed in them. Blade servers make use of virtual devices that are managed through the embedded hardware management interfaces.

Blade servers have evolved in recent years and now offer comparable hardware options as traditional servers. Older blade servers were typically limited to only one or two single-core CPUs, two NICs, and no support for Fibre Channel storage. Some modern blade servers can support up to 16 NICs, 4 quad core processors, and multiple Fibre Channel or iSCSI HBA adapters.

The choice between blade and traditional servers often comes down to personal preference and what type of server is already in use in your datacenter. Regardless of which server type you choose, they both work equally well as ESX hosts.

Multicore CPUs

In the old days, CPUs had only a single core, which made choosing a CPU relatively easy. You just ordered a server with however many CPUs you needed. These days, multicore CPU servers are very popular and are standard on many server models. Currently, dual-core and quad-core CPU servers are available, with eight or more cores available sometime in the future.

Simply put, a multicore CPU combines multiple independent cores on a single physical CPU. Depending on the CPU brand and model, these cores sometimes share a single cache or have separate Level 2 caches for each core, as depicted in Figure 2.12.




Figure 2.12 Diagram of a generic dual-core processor, with CPU-local Level 1 caches, and a shared, on-die Level 2 cache

Each multiple independent core appears as a separate processor to the host computer. So a computer that has two quad-core processors will actually have eight different processors to use. You may also hear the term *socket*, which is just the connector on a computer motherboard that a physical CPU plugs into. The term socket is often used to specify the number of physical CPUs that a computer has instead of the amount of total cores.

Multicore servers are ideal for virtual hosts because VMware licenses ESX by the socket regardless of how many cores the CPU in that socket has. This licensing model may change

in the future as CPU vendors increase the number of cores on each CPU. When purchasing a server to be used as an ESX host, you must decide between dual-core and quad-core CPUs. Your first impulse may be to go for the quad-core CPUs over the dual-core CPUs because you might think that more cores is better. However, there are some differences between the two that you should know about before making your decision.

Core increases do not scale the same as CPU clock speed increases. A 2.8GHz CPU will be twice as fast as a 1.4GHz CPU, but a quad-core CPU is not four times as fast as a single core. A dual-core CPU is roughly 50% faster than a single-core CPU (not 100% as you would think), and a quad-core CPU is only about 25% faster than a dual-core CPU. Also, dual-core CPUs typically have higher clock speeds than quad-core CPUs. The reason for this is that quad-core CPUs generate excessive heat and they cannot be clocked as high as single- and dual-core CPUs.

So which should you choose? In general, quad-core CPUs are recommended for ESX hosts. There are two reasons for this. The first reason is that because ESX is licensed by the number of sockets in a computer and not the number of cores, you get more bang for your buck when it comes to buying licenses. The second reason is that having more cores in a host server gives the hypervisor's CPU scheduler more flexibility when trying to schedule CPU requests that are made by the VMs. Having more cores available to choose from makes the CPU scheduler's job much easier and will generally improve the performance of all VMs on a host.

When would you want to use dual-core CPUs instead of quad-core CPUs? You might consider dual-core CPUs for ESX hosts that you plan to run less than eight VMs on, and on hosts that you plan to run only single-vCPU VMs on. Single-vCPU VMs make are easier for the hypervisor to schedule than multiple vCPU VMs.

Choosing a CPU Vendor

You have two vendor choices when it comes to choosing a CPU for your ESX hosts: Intel and AMD. There have been many performance studies and debates comparing the two over the years, and with the constantly changing processor architectures sometimes AMD is ahead of Intel, and vice versa. Both Intel and AMD have integrated virtualization extensions (Intel-VT and AMD-V) into their latest processors in an attempt to speed up instruction execution for virtual servers. The main difference between Intel and AMD processors is their physical architecture. Intel uses a front side bus model to connect the processors to the memory controller, and AMD uses an integrated memory controller for each processor that is interconnected through a hyper transport. In addition, they tend to have different power-consumption levels depending on the processor family.

To further help virtual hosts handle large workloads, a new technology called Nested Page Technology (NPT) was developed by AMD to help reduce the performance impact of virtualizing large applications such as databases. AMD's Rapid Virtualization Indexing (RVI)

feature includes NPT and is currently available for their third-generation Opteron processors. Intel is developing a similar solution that is named Extended Page Tables; it will be available in its next-generation eight-core micro-architecture, which is code named Nehalem and due to be released sometime in 2009.

In September 2008, Intel announced their new Xeon 7400 processor family (previously known as Dunnington) at VMworld 2008, with VT FlexMigration, a larger 16MB Level 3 cache and six cores to dramatically increase performance. Because ESX is licensed by processor socket and not processor core, having a six-core processor is definitely an advantage over a four-core processor. As technology is constantly changing, you should analyze both offerings to see which vendor has the best technology when it comes time for you to choose a processor for your ESX hosts.

When it comes to performance, they tend to be fairly equal when you compare processors of similar speeds, features, and number of cores. Some performance studies will show Intel processors with a performance advantage, and others will show AMD with a performance advantage. Both Intel and AMD processors work well in ESX hosts, and it tends to be more a matter of brand preference when it comes time to choosing one. The important things to keep in mind when choosing a CPU vendor are as follows:

- Ensure that the processor model that you choose for your system is supported on the ESX HCL.
- Look for processors that support the Intel-VT and AMD-V virtualization extensions, which are optimized for virtual environments. In addition, look for processors that support the new NPT technology from AMD (RVI) and Intel (EPT).
- Keep VMotion compatibility in mind. By default, you can't VMotion a VM from an ESX host with an Intel processor to another ESX host with an AMD processor. You also cannot VMotion between certain processor families within each brand. Check the CPU compatibility matrixes to ensure your processors are compatible with each other.
- If you do end up with ESX hosts with different processor brands, you should create separate clusters for each. For example, create one cluster for all your Intel processor ESX hosts and another for all your AMD processor hosts.
- ESX version 3.5 Update 2 introduced a new feature called Enhanced VMotion compatibility that helps ensure compatibility of all hosts in a cluster. However, this feature requires all hosts to be of the same processor brand and they must also have hardware live migration support (either Intel FlexMigration or AMD-V Extended Migration).

Looking ahead, AMD plans to release a six-core processor, code named Istanbul, sometime in 2009. In addition, they are looking to release a 12-core processor sometime in 2010. Intel will most likely follow suit with this but have currently not announced this in their processor roadmap.

Storage Adapters

You may use different types of storage adapters in your ESX hosts: local storage adapters (SCSI/SAS/SATA), Fibre Channel HBA adapters, and iSCSI HBA adapters. Again, make sure whatever adapters you end up using are on the HCL. When you are choosing storage adapters for your ESX host, keep the following in mind:

- For local storage adapters, it is best to get adapters that have the largest read/write caches on them, especially if you plan to exclusively use local disk on your ESX hosts. In addition, having a battery-backed write cache (BBWC) on your array controller improves performance and reliability. BBWCs add some additional memory that is used to cache disk writes and have a battery backup to protect data that has not yet been written to disk (in case of a power failure).
- For Fibre Channel adapters, you will typically want two of them for maximum reliability. The default Fibre Channel drivers that are included with ESX support the QLogic and Emulex FC adapters, which are usually rebranded by server manufacturers (for example, HP and IBM).
- If you plan to use iSCSI hardware initiators, for iSCSI adapters you will again want two of them for maximum reliability. The default iSCSI drivers that are included with ESX support only the QLogic adapters, which are usually rebranded by server manufacturers (for example, HP and IBM).

Network Adapters

Most servers will have at least two network adapters built in. When you are choosing additional network adapters for your ESX hosts, keep the following in mind:

- ESX supports only a limited number of NIC brands and chipsets. The two most popular NIC chipsets that ESX supports are from Intel and Broadcom. Make sure you check the HCL to see whether your NIC brand and model are on it.
- Consider quad-port NICs that provide four NICs on a single adapter and take up only one expansion slot on your server. However, you should use more than one adapter for maximum reliability so that in case one fails you still have an active NIC on your vSwitches. Having two onboard NIC ports and a separate four-port adapter provides you this flexibility; adding an additional four-port adapter gives you even more flexibility and reliability.

Unsupported Hardware

ESX will run on many hardware platforms that are not included on the HCL, including generic servers that are called “white boxes.” However, you may run into issues when you use unsupported hardware, and you may be unable to obtain support for technical problems. Typically, VMware provides “best effort” support for hardware that is not on the HCL. Unsupported hardware is good for labs and for learning and playing around with ESX in home and nonproduction environments. Many older servers from HP, IBM, and Dell that are not on the current HCL will load and run ESX without issue. ESX will not work with all hardware, however, because it includes only a limited driver set in its base installation. For example, IDE drives are not supported with ESX. You can usually install ESX on an IDE drive, but you will be unable to create any VMFS volumes on them. If you plan to use white box hardware, it is best to look through some of the published lists of white box hardware that is known to work with ESX.

Storage Options

Storage is probably the most important hardware choice you will make when selecting hardware for your host servers. Because storage relies on a mechanical device, it is often the first resource bottleneck encountered on host servers. Consequently, choosing a proper storage solution is critical to ensuring a successful virtualization project. Many variables will factor into which storage option you choose for your virtual environment. Ultimately, cost and I/O requirements tend to be the two biggest factors that will influence your storage choice. What you can afford will play a large part in determining which storage option you choose, and the disk I/O requirements for the applications that you plan to run will also be a big factor. Therefore, analyze the storage choices carefully and make sure you understand the differences between the various choices available to you.

10K Versus 15K rpm Hard Drives

Most SCSI hard drives are available in two speeds: 10,000 rpm (10K) and 15,000 rpm (15K). This speed is how fast the hard drive’s platter spins, and is otherwise known as its rotational speed. With the drive platter spinning faster, data can be read and written faster, and overall latency is reduced. However, even though the drive platter is spinning faster, the head actuator that moves across the drive to access data does not move any faster. Therefore, just because the drive is spinning 50% faster doesn’t mean overall drive performance is increased by 50%. Typical performance increases for 15K drives over 10K drives are about 30%, which result in higher I/O operations per second (IOPS) and decreased average access times.

There are generally two deciding factors when it comes to choosing between 10K rpm and 15K rpm drives. The first is whether you will be using applications that will have heavy disk utilization and could benefit from the extra speed that 15K rpm drives provide. The second is

whether you can afford the more expensive drives. If you can afford them, it is generally best to get them. The only downside to 15K drives is their additional expense over 10K drives. If you plan to run disk I/O-intensive applications on your VMs, definitely consider them.

Boot from SAN

The boot from SAN option allows ESX hosts (currently ESXi hosts cannot use this feature) to be configured without any internal disk and instead boot directly from a SAN device. This can prove particularly advantageous when using blade servers that have very limited internal storage. When you use boot from SAN, the ESX Service Console is installed on a LUN that is located on the SAN storage. Then when the ESX host starts up, it boots directly from the LUN on the SAN. When you enable boot from SAN, you configure your HBA adapter to boot from the specific LUN for that ESX host. Each host must have its own separate LUN on the SAN. Booting from SAN has advantages and disadvantages.

■ Boot from SAN advantages

Lower server hardware cost because local disks are not needed.

Easier server hardware replacement. You can replace old servers with new ones and simply point to the old boot location on the SAN.

Very useful for disaster recovery (DR). By having the boot images of the ESX hosts stored on the SAN array, you can easily replicate them to remote sites where standby ESX hosts can quickly boot in case of a disaster.

SAN-level snapshots can be used to create point-in-time backups of the Service Console when patching ESX hosts.

Faster deployment of ESX hosts by cloning existing boot images.

■ Boot from SAN disadvantages

SAN disk can be expensive, and having local disk on your ESX hosts available for VMs can be a low-cost alternative to storing all your VMs on SAN disk.

If something happens to the SAN or the connection to the SAN, the whole ESX host and all of its VMs will go down. Having local disk enables VMs that are not on the SAN to continue running when the SAN is unavailable.

Adds more load on the SAN because of the additional load from the disk I/O from the Service Console partitions.

Will not work with Microsoft Cluster Services (MSCS), which requires that the VMs in the cluster have their boot disks on local disk.

Bootting from SAN can be more complex to configure and more difficult to troubleshoot.

Requires a separate LUN for each ESX, which can be complicated to manage and more expensive if your SAN is licensed by zones.

As stated, there are definitely certain situations when booting from SAN makes sense and should be considered. Booting from local disk is used in most cases because many people prefer to have local disk and do not want to waste SAN disk for the Service Console.

Local Storage

Local storage is typically cheap and reliable and is good to have on ESX hosts even if you plan to run all your VMs on shared storage.

■ Advantages of using local storage

Cheap and reliable storage.

Good choice for running development and test VMs (so that you can save your shared storage for production servers).

Good for backing up VMs that are on shared storage. You can run scripts that periodically snapshot a VM and then copy its disk files from shared storage to local storage as an additional backup option.

The Storage VMotion feature enables you to move running VMs from one ESX host to another by first copying the VM to shared storage and then back to local storage on another ESX host.

■ Disadvantages of using local storage

Cannot use advanced features (HA/DRS/VMotion) that require shared storage.

Not available for other ESX hosts to use. Only the local ESX host can access it.

Unless you are using the boot from SAN feature, you should consider getting at least two local disks that are using RAID on your ESX host, and more if you can afford it.

Did You Know?

You can turn local storage into shared storage for your ESX hosts using a third-party product like Left Hand Network's Virtual SAN Appliance (VSA). This application provides you a lower-cost alternative to using one of the traditional shared storage solutions.

Fibre Channel

Fibre Channel (FC) traditionally uses fiber-optic cables to connect FC host bus adapters (HBAs) to SAN devices through special FC switches. Speeds can vary from 1Gbps to 8Gbps, with 4Gbps being the most popular speed today. All components in an FC network must support whichever speed you use; this includes the FC HBA, the FC switch, and the FC controller on the storage device. Each FC HBA has a unique World Wide Name (WWN), which is similar

to a network MAC address. FC networks typically have multiple redundant paths from the host servers to the storage devices that include multiple HBAs, switches, and controllers.

FC or SAN storage is generally the most popular storage choice for ESX in larger environments. The reasons for this are its speed (4Gbps has been around for years, and 8Gbps is now supported) and reliability. FC SANs are isolated and more secure.

■ Advantages of using Fibre Channel storage

Typically, the best performing and most secure storage.

ESX is able to boot from Fibre Channel storage.

Block-level storage.

■ Disadvantages of using Fibre Channel storage

Typically, the most expensive storage option to implement.

Can be complex to implement and manage.

If you already have an FC SAN solution in your environment, using it with ESX just makes sense. Expanding an existing SAN is much easier and cheaper than implementing a new SAN from scratch. Also, designing a SAN architecture and administering it usually requires specialized training, which can further add to the expense of implementing it. Many larger environments that use SANs have dedicated storage administrators to do this.

If you plan to have many high disk I/O VMs running on your ESX hosts, seriously consider using SAN storage to achieve maximum performance. Ultimately, cost is usually the factor that determines whether you will use SAN storage or choose a less-expensive alternative.

iSCSI

iSCSI is a relatively new type of network storage that was supported beginning with ESX 3.0. It works by using a client called an initiator to send SCSI commands over a LAN to SCSI devices (called targets) located on a remote storage device. iSCSI utilizes traditional networking components and TCP/IP and does not require special cables and switches like FC storage does. iSCSI can still be considered a type of SAN storage device because it writes using a block-level method rather than a file-level method that NFS uses.

iSCSI initiators can be either software based or hardware based. An initiator is the client that replaces the traditional SCSI adapter that servers typically use to access SCSI storage. Software initiators utilize device drivers that are built in to the VMkernel to use existing network adapters and protocols to write to remote SCSI devices, as shown in Figure 2.13. This can cause additional CPU and network overhead on the host server. Some characteristics of software initiators are listed here:

- Uses existing NICs and the native VMkernel stack.
- Good choice for blade servers and servers with limited expansion slots.

- Cheaper solution than hardware initiators.
- Can be CPU intensive due to the additional overhead of protocol processing.
- The ESX server is unable to boot from it.




Figure 2.13 iSCSI software initiator

Hardware initiators use a dedicated iSCSI HBA, which combines a network adapter, a TCP/IP offload engine (TOE), and a SCSI adapter into one device to help improve the performance of the host server, as shown in Figure 2.14. Some characteristics of hardware initiators are listed here:

- Uses dedicated iSCSI HBAs.
- Moderately better I/O performance than software initiators.
- Uses less ESX server host resources, especially CPU.
- The ESX server is able to boot from it.

Advantages of using iSCSI storage include the following:

- Lower cost to implement than FC.
- Software or hardware initiators can be used.
- Can create VMFS volumes on it.
- Block-level storage.
- Speed and performance are increased with 10Gbps Ethernet.

Disadvantages of using iSCSI storage include the following:

- Currently, no support for jumbo frames
- CPU overhead when using a software initiator




Figure 2.14 iSCSI hardware initiator

iSCSI is a great alternative to using FC storage because it is cheaper to implement and provides very good performance. Starting with the release of ESX Server 3.5 Update 2, it now supports using 10Gbps Ethernet, which provides a big performance boost. Previously, it supported only up to 1Gbps Ethernet, which made FC and its 4Gbps (now 8Gbps) data transfer rate a better option. The biggest risks to using iSCSI are the additional CPU overhead when using software initiators, which can be mitigated by using hardware initiators, and the more fragile and volatile network infrastructure that it relies on, which can be mitigated by completely isolating iSCSI traffic from other network traffic.

NAS/NFS

Network attached storage (NAS) utilizes the Network File System (NFS) protocol to allow hosts to mount partitions on a remote file system and access them as if they were local disks. NAS storage has performance characteristics similar to software iSCSI; however, the performance is heavily dependent on the speed and health of the network connection between the host and the remote storage and the type of NAS device you are connecting to. Using a dedicated NAS appliance will provide much better performance than using a Linux or Windows server that is running NFS services.

Advantages of using NAS/NFS storage include the following:

- No substantial performance difference compared to iSCSI.
- Can be the lowest cost shared storage option.
- Can use existing infrastructure components.

- No single disk I/O queue. Performance is strictly dependent on the size of the network connection and the speed of the disk array.
- Smallest storage footprint because it uses thin provisioned disks by default.

Disadvantages of using NAS/NFS storage include the following:

- Cannot boot ESX from it
- Does not support VMFS volumes or RDMs
- CPU overhead

NAS storage has a few disadvantages compared to iSCSI and FC SAN storage, mainly in the features that it supports, but is a viable alternative to use with ESX and should be considered. If you choose to use NAS storage, consider using a dedicated NAS device from a vendor such as NetApp. Table 2.7 provides a summary of the features that each storage option supports.

Table 2.7

Storage Option Characteristics

	Boot from SAN	Supports RDMs	Supports VMFS	Type
Local Storage	N/A	N/A	Yes	Block level
Fibre Channel	Yes	Yes	Yes	Block level
iSCSI (software)	No	Yes	Yes	Block level
iSCSI (hardware)	Yes	Yes	Yes	Block level
NAS/NFS	No	No	No	File level

Mixing Storage Types

With so many storage types to choose from, it's common to mix them on your ESX hosts by configuring separate datastores for each type. You cannot, however, mix storage types on a single datastore. You may have one datastore that is configured to use local disk on your ESX host, another that uses shared storage on an FC SAN, and one more that uses storage on an NFS device. You can then choose the placement of your VMs based on the characteristics of each storage type. Most likely, you will want your most critical and disk I/O-intensive VMs on your best performing shared storage. Local disk is good for nonessential VMs, and NFS datastores can be set up relatively easily on Windows servers and used to store ISO files and templates. You can also easily move VMs from one storage type to another as needed by cold migrating them while powered off or while they are running by using the new Storage VMotion feature. The important thing to know when using multiple storage options in your environment is to recognize the performance and feature differences between each option.

Virtual Networking

Networking in a virtual environment is much different from networking in a physical environment. Therefore, it is important to understand the differences so that you can properly plan and configure your virtual networking environment. When it comes to networking, not too many choices need to be made when deciding on what is needed for your host servers. Your main decision will be how many physical NICs are needed for your host servers. This number will be influenced largely by how much redundancy you want to provide and whether you plan to use network-based storage. Although it is important to understand your requirements so that you can make the proper networking choices early on, if you do find that you need to add additional networking later, in most cases this is easy to do.

Physical NICs

Physical NICs (pNICs) are not configured using traditional methods in ESX hosts. You do not assign an IP address to each pNIC as you would in a standard operating system. Instead, pNICs connect to uplink ports that provide a connection between the vSwitches and physical networks. When it comes to the VM networks, however, the IP addresses are assigned to them through their operating systems as usual, and the pNIC just acts as an uplink that connects it from the vNIC, through the vSwitch, to the pNIC, and finally to the physical network switch that it is plugged in to. Even the ESX Service Console management interface (vswif), which is configured when installing ESX, does not have an IP address directly assigned to a pNIC; because the Service Console runs as a privileged VM, it connects through a vSwitch just like other VMs.

When it comes to pNICs in your ESX hosts, you will typically want at least two, preferably four, and sometimes more depending on your storage choice, if you plan to use VMotion, and how you plan to configure your virtual networking. Also, additional pNICs provide you with the ability to isolate sensitive traffic for the Service Console and VMkernel and provide fault tolerance in your vSwitch configurations. When planning the number of pNICs for your ESX host, consider the following:

- It's not necessary, but it's a best practice to isolate the Service Console traffic for security reasons. Therefore, you should isolate the Service Console on its own vSwitch with its own pNIC. For maximum reliability, it is recommended to add a second pNIC in case of a failure. The second pNIC should be plugged into a separate physical switch so that a single switch failure cannot affect both pNICs assigned to the Service Console vSwitch. Optionally, you can add a second Service Console interface (vswif) to another vSwitch to provide redundancy. The Service Console network is critical for a lot of the ESX host functionality, and it is important to ensure it does not lose network connectivity. For example, the HA feature relies on a heartbeat that is broadcast over the Service Console network. Without redundancy, if a pNIC on it failed it would

trigger an HA event and would shut down the VMs and move them to other hosts. In addition, if the Service Console network loses connectivity, the vCenter Server will be unable to communicate with the host, and you won't be able to manage or monitor it.

- If you plan to use VMotion, you will need to configure a VMkernel network on a vSwitch. The best practice is to also isolate this traffic to only the ESX hosts that will be using VMotion. As with the Service Console, you will want to add a second pNIC to this vSwitch that is connected to a separate physical switch.
- For your VM networks, how many pNICs you will need will depend on how many VLANs that your VMs will need to connect to and whether you plan to use 802.1Q VLAN tagging. If you plan to use only one VLAN or use tagging, you need only one pNIC in your vSwitch. However, you should add a second pNIC for redundancy purposes for load balancing and failover. If you plan to plug many VMs into the vSwitch and expect heavy network I/O, you should add additional pNICs to the vSwitch as needed.
- You may also consider creating additional vSwitches for your VMs so that you can isolate traffic on to separate pNICs. Another reason you might want to do this is to isolate Internet DMZ traffic from your internal network traffic. It's best to isolate your DMZ traffic on a separate vSwitch with separate pNICs, for maximum security.
- If you plan to use network storage like software iSCSI or NFS, you should also isolate this traffic onto its own vSwitch with redundant pNICs.

You need to take all these factors into account when trying to determine how many pNICs you need for each ESX host. It's better to have too many pNICs in your ESX hosts instead of too few. These days you can purchase pNIC cards that contain up to four pNICs on a single card. Having six NICs in your ESX host servers will give you a lot of flexibility and redundancy when configuring your vSwitches. ESX does support up to 32 pNICs in a single host. It's best to get up front as many as you think you will need. If you do find yourself running into a network bottleneck later or need to add additional capacity to your ESX host, you can always add additional pNICs later.

Use these guidelines when trying to decide on the number of pNICs you will need in your ESX hosts:

- **If you are using local, FC, or hardware iSCSI storage**
Use at least six physical NICs: two for the Service Console network, two for VMotion, and at least two for your VM networks.
- **If you are using software iSCSI or NFS network-based storage**
Use at least eight physical NICs: two for the Service Console network, two for VMotion, two for your IP storage, and at least two for your VM network.

If you have smaller ESX hosts and do not plan to use features like VMotion or network-based storage, plan to have at least two NICs, but preferably four. In addition, if you plan to physically segment your virtual networks for cases like setting up a DMZ environment, consider additional NICs so that you can set up separate vSwitches for them and provide network redundancy on all vSwitches.

Virtual NICs

A virtual NIC (vNIC) is what you assign to a VM and is connected to a vSwitch on the ESX host server. You can assign up to four vNICs to a single VM. A vNIC assigned to a VM does not use the MAC address of the physical NIC. Instead, it uses a MAC address generated by the ESX host and within the 00-50-56 range that has been assigned to VMware.

In most cases, you assign only a single vNIC to a VM unless you need to connect the VM to multiple vSwitches to access separate network VLANs. Assigning multiple vNICs to a VM to try to team them at the operating system layer can prove problematic in a virtual environment.

There are many types of vNICs to choose from when configuring a VM. The types that you have to choose from will vary based on the host version and the operating system type that you select when setting up a new VM. The various types listed here are also described in the VMware knowledge base article 1001805:

- **Vlance.** Vlance (also called PCNet32) is a faithful virtual implementation of a common, if now somewhat aging, physical network adapter. Most 32-bit guest operating systems, except for Windows Vista, have built-in support for this card, so a VM configured with this network adapter can use its network immediately.
- **Vmxnet.** The vmxnet virtual network adapter has no physical counterpart. VMware makes vmxnet available because Vlance, a faithful implementation of a physical card, is far from optimal for network performance in a VM. Vmxnet is highly optimized for performance in a VM. Because there is no physical card of type vmxnet, operating system vendors do not provide built-in drivers for this card. You must install VMware Tools to have a driver for the vmxnet network adapter available.
- **Flexible.** The Flexible network adapter identifies itself as a Vlance adapter when a VM boots, but initializes itself and functions as either a Vlance or a vmxnet adapter, depending on which driver initializes it. VMware Tools versions recent enough to know about the Flexible network adapter include the vmxnet driver but identify it as an updated Vlance driver, so the guest operating system uses that driver. When using the Flexible network adapter, you can have vmxnet performance when sufficiently

recent VMware Tools is installed. When an older version of VMware Tools is installed, the Flexible adapter uses the Vlan adapter (with Vlan performance) rather than giving no network capability at all when it can't find the vmxnet adapter.

- **e1000.** e1000 is a faithful virtual implementation of a physical network adapter that is broadly supported by newer operating systems, specifically most 64-bit operating systems and both 32- and 64-bit Windows Vista. e1000 performance is intermediate between Vlan and vmxnet.
- **Enhanced vmxnet.** The enhanced vmxnet adapter is based on the vmxnet adapter but provides some high-performance features commonly used on modern networks, such as jumbo frames. This virtual network adapter is the current state-of-the-art device in virtual network adapter performance, but it is available only for some guest operating systems on ESX Server 3.5. This network adapter will become available for additional guest operating systems in the future.

Enhanced vmxnet is supported only for a limited set of guest operating systems:

- ◆ 32/64-bit versions of Microsoft Windows 2003 (Enterprise and Datacenter Editions)
- ◆ 32/64-bit versions Red Hat Enterprise Linux 5.0
- ◆ 32/64-bit versions SUSE Linux Enterprise Server 10
- ◆ 64-bit versions Red Hat Enterprise Linux 4.0

The VM setup wizard automatically hides vNIC types based on your operating system choice and will display only the types appropriate for your VM operating system. In most cases, you will have only one type available to you. Here's a general overview of what types will be available to you:

- For VMs native to VMware ESX Server 2.x, you can explicitly choose between Vlan and vmxnet.
- For most 32-bit VMs native to VMware ESX Server 3, only the Flexible adapter is available.
- For most 64-bit VMs and for 32-bit Microsoft Windows Vista VMs, only the e1000 adapter is available.
- For certain guest operating systems on VMware ESX Server 3.5 and later, you can choose the enhanced vmxnet adapter in addition to the Flexible or e1000 adapter mentioned for that guest type in the previous list.

Virtual Switches

Virtual switches are the core networking component on an ESX host. A vSwitch is built and contained in the RAM of an ESX host and connects the physical NICs in the host server to

the vNICs in VMs, as shown in Figure 2.15. vSwitches emulate many of the traits of traditional Ethernet switches and can perform many of the same functions, such as forwarding frames at the data link layer, VLAN segmentation, and supporting the copying of packets to a mirror port for use with a sniffer or intrusion detection system (IDS).




Figure 2.15 vSwitches in ESX Server 3 connect VMs and the Service Console to each other and to external networks.

Each ESX host may have up to 127 vSwitches configured on it with up to 1,016 ports per vSwitch and a maximum of 4,096 total vSwitch ports per host. Depending on the pNIC type that you use, you may use up to 32 pNICs in a single ESX host, which you can then assign to your vSwitches. You can assign multiple pNICs to a vSwitch (up to 32) to team them for load-balancing and failover purposes, but a pNIC can be assigned to only a single vSwitch at any one time.

vSwitches support 802.1Q VLAN tagging, which allows for multiple VLANs to be used on a single physical switch port. This capability can greatly reduce the number of pNICs needed in an ESX host. Instead of needing a separate pNIC for each VLAN that you need to connect to on an ESX host, you can use a single NIC to connect to multiple VLANs. The way tagging works is that tags are applied to all network frames to identify them as belonging to a particular VLAN. There are several methods for doing this in ESX, as outlined in the following list. The main differences between the modes are where the tags are applied. VGT mode does this at the guest operating system layer, EST mode does this on the external physical switch, and VST mode does this inside the VMkernel. The differences between the VLAN tagging modes are outlined here:

- **VM guest tagging (VGT mode).** With this mode, the 802.1Q VLAN trunking driver is installed inside the VM. Tags are preserved between the VM networking stack and external switch when frames are passed to and from vSwitches.

- **External switch tagging (EST mode).** With this mode, you use external switches for VLAN tagging. This is similar to a physical network, and VLAN configuration is normally transparent to each individual physical server. The tag is appended when a packet arrives at a switch port and stripped away when a packet leaves a switch port toward the server.
- **vSwitch tagging (VST mode).** With this mode, you configure port groups on a vSwitch for each VLAN and connect the vNIC of the VM to the port group instead of the vSwitch directly, as shown in Figure 2.16. The vSwitch port group tags all outbound frames and removes tags for all inbound frames. It also ensures that frames on one VLAN do not leak into a different VLAN.




Figure 2.16 vSwitch tagging mode (VST)

The VST mode is the one that is most commonly used with ESX because it's easier to configure and manage. In addition, it eliminates the need to install a specific VLAN driver inside a VM, and there is almost no performance impact by doing the tagging inside the vSwitches.

When you create a vSwitch, the first thing you do is assign one or more physical network adapters to it. If you assign more than one, the additional NICs can be configured as a team

to support load balancing and failover. It is also possible to not assign a physical NIC to a vSwitch, thereby creating an internal-only or private vSwitch that is isolated from the rest of the network. This is useful for testing purposes or when building a new VM to prepare it before joining it to the rest of the network. After you have assigned NICs to a vSwitch, you can then create port groups for individual VLANs.

Port groups are a unique networking feature to VMware vSwitches. Even though port groups can be assigned VLAN IDs, they should not specifically be considered to be a VLAN itself. They are more like configuration templates for the vNIC ports on the vSwitch. It is possible to have different port groups using the same VLAN ID on the same vSwitch. In addition to the VLAN ID, port groups allow you to set different networking characteristics, such as traffic shaping (quality of service, QoS) and security settings like promiscuous mode or allowing MAC address changes. These settings are discussed in detail in the next chapter.

Licensing

There is not all that much to plan for when it comes to licensing. Just be sure to understand VMware's licensing model. It is always subject to change, and so you should check VMware's website for the most up-to-date information.

Centralized Versus Stand-Alone

ESX hosts have two licensing options: individual stand-alone licensing using a license file on each host, and centralized licensing using a separate license server for all your ESX hosts. The licensing server uses FLEXnet from Macrovision and is included with vCenter Server and runs as either a separate application on the vCenter Server itself or on a different server. The formats of server-based and host-based licensing files are different and cannot be intermixed. You can change your licensing model at any time by going to the VMware licensing page and regenerating your licensing key.

If you have a vCenter Server, you will almost always want to use the centralized licensing model. This model uses a single license file and automatically decrements available licenses when new ESX hosts are added to the environment. You never have to install licenses on each individual host when using this model. The license server does not provide a single point of failure because the ESX hosts stay fully functional for up to 14 days if a license server is unavailable.

The stand-alone host-based licensing is typically used in environments for ESX hosts that are not managed by a vCenter Server. With this model, license files are uploaded to each ESX host using the VI client, as depicted in Figure 2.17. Each license file is different, so this model can be more complex to manage if you have many ESX host servers and is more suitable for smaller environments.




Figure 2.17 Stand-alone host-based licensing model

Licensing Considerations

All ESX licenses are sold in two-processor (socket) increments. Previously, VMware's licensing agreement prohibited you from splitting a license so that you could use it on a single-processor server. VMware recently changed that policy, and you can now split a processor license to use on a single-processor system (with up to four cores); this also includes servers that have two or more processor sockets but are only populated with a single processor. However, if you do want to split a license, you must use the centralized licensing model.

Additional Considerations

You should take into account a few additional considerations when planning your virtual infrastructure. If you currently have a disaster recovery (DR) strategy or are looking to implement one in the future, you should consider this when planning your virtual infrastructure. Depending on your DR goals, this could impact the decisions and choices that you need to make as part of the planning process. Having remote datacenters could also impact the design of your virtual architecture, and special consideration should be given to them. In addition, be aware of VM sprawl (which commonly occurs in virtual environments) so that you can try to control it in your own environment.

Remote Datacenters

If you plan to install ESX hosts in remote datacenters, you should also plan on how you will manage these hosts. Your options for this may be dictated by the speed and quality of the network link between remote datacenters. vCenter Server is capable of managing remote ESX hosts, but depending on the speed of the connection between the vCenter Server and the

ESX host, you may experience occasional timeouts and slowness, especially when opening up remote console sessions to the VMs running on the hosts. If you plan to have a large number of ESX hosts in a remote datacenter, you might consider a separate vCenter Server installation located at that datacenter, especially if you plan to use some of the advanced features like HA, DRS, and VMotion.

Disaster Recovery

ESX is frequently used when implementing DR models because of its architecture and capabilities. VMs are ideal for DR because of their hardware independence, the encapsulation of their disks into a single file, and because they can run on shared storage. There are many ways to implement DR based on requirements and expected recovery times. Storage is typically the focal point when implementing DR. There are many ways to move data on your servers from one site to another, and which one you choose will largely be dictated by your available network bandwidth between your main site and your DR site, how much money you have to spend on a solution, and your minimum recovery times. ESX provides many options for implementing DR. Before deciding on one, you need to take into account some architectural considerations:

- What method will be used to back up VMs? Traditional agent-based backups at the operating system level take longer and are more complex to restore. Backing up the whole virtual disk file outside of the operating system provides an easier and quicker recovery option. Consolidated Backup is also a good option to use for backing up VMs for DR purposes. You might consider mixing backup methods, doing daily backups through traditional methods and weekly backups of VM disk files.
- Will you need quick or real-time recovery of your VMs? If so, consider SAN replication between sites or using some additional products that provide this capability.
- Will your hardware at the DR location be equal to your main site? If you will have limited hardware at your DR site, ensure it will be able to handle whatever load will be placed on it when a DR event occurs. Consider which VMs at your main site actually need to be included in your DR plan.
- Don't just plan on recovery for ESX hosts and VMs. You should also take into account recovering vCenter Server, its database, and the licensing server.

As mentioned, there are many ways to move your data from a main site to a DR site, including the following:

- **Traditional agent backups.** This is typically the slowest method to recover from. You will need to create new VMs at your DR site, install operating systems and backup agents on them, and perform a restore to recover them.

- **VM image backups.** Quicker recovery than traditional backups. Just restore the VM disk files to an ESX host at the DR site and create a new VM and configure it to use the restored disk file. Optionally, you can also back up the VM configuration files and just register the VM on the new ESX host.
- **Physical machine cloning.** This method creates a clone of an existing physical server to be used as a virtual server at a DR site through a product that supports physical to virtual (P2V) machine conversions. P2V products that provide this capability include VMware vCenter Converter, Vizioncore vConverter, and Platespin PowerConvert.
- **VM cloning.** This method creates a clone of an existing VM to be used at a DR site. You can schedule this to occur, and an exact copy of the VM will be created on the destination host. There are several ways to do this, including using vCenter Server, VMware vCenter Converter, and using scripts.
- **VM replication.** This method provides very quick or even real-time recovery for your VMs. A number of products will do this, including VMware's Site Recovery Manager, Double-Take for VI3, and VizionCore vReplicator.
- **SAN replication.** This method provides replication at the SAN layer to a redundant SAN at the DR site. This method also provides very quick or real-time recovery for your VMs. This method can be used in conjunction with VMware's Site Recovery Manager product.

ESX provides many additional options for DR execution that are not available in a traditional physical server environment. Choose the option that works best for you and that fits within your budget and satisfies your recovery time requirements.

VM Sprawl

Virtual environments are often susceptible to a condition that is called VM sprawl. VM sprawl can be defined as the uncontrolled growth of VMs in a virtual environment. *Star Trek* fans can relate to this from the popular "Trouble with Tribbles" episode, where the Tribbles reproduce so quickly that they threaten to overwhelm the host ship's food supplies. VM sprawl is similar because administrators often create VMs without any regard for the resources that they consume and the possibility of overwhelming the host server's resources. Many consider VMs as "free" servers because they do not have a physical presence and creating them is simple and easy.

You can deal with VM sprawl in various ways. For example, you can implement a charge-back system like Vkernel's capacity and chargeback appliance, or you can use Vizioncore's vFoglight products for VMware, which create reports on resource usage. In addition, limiting who can create VMs and implementing a formal process for requesting new VMs is a more effective method than allowing too many people to have access to create VMs at will. You should think about requiring justification for requests for any new VMs and have an approval

process to get users to think twice about whether they really need another one. Finally, creating resource pools is a helpful way to limit the amount of resources available on your host servers for new VMs.

One solution to protect your resources on your ESX hosts is to set up a few host servers using the free hosted product VMware Server or ESXi (now that it is also free). By doing this, you can put nonessential VMs on a lower-cost host server and conserve your ESX host resources. However, this is not a solution to controlling sprawl and should be considered as a way to divert requests for new VMs to alternative hosts.

It's important to control sprawl early on; otherwise, you might use up all your host resources before you know it and create bottlenecks that may reduce the performance of all your VMs. Make people aware that VMs are not free and that there is an associated cost for them regardless of how they are configured. Having tight controls on your virtual environment is the key to limiting the unnecessary growth of VMs on your host servers.

Summary

In this chapter, we covered the many things that you need to consider when planning your VI3 environment. You have learned about various features and about decisions that need to be made when planning your project. Take some time to think about this material and understand everything, and be sure to research anything that you do not understand. Choose carefully before you start assembling the various pieces of your VI3 environment. Confirm that you have all the correct pieces and that they will all fit together properly. In the next chapter, we cover the installation of the various components of your VI3 environment.

Endnotes

1. The use of the term *VirtualCenter* here refers to the old name, which still applies to VirtualCenter version 2.0.x and is still present inside the application in vCenter Server 2.5 (because the software has not yet been updated to reflect the new name).

This page intentionally left blank

Chapter 3

Building Your Virtual Environment

Now is it's time to start constructing your virtual environment. There are several different components that you will be building for this, including a licensing server, ESX and ESXi hosts, and your vCenter Server. This phase is one of the most exciting because you can finally see your virtual environment take shape and materialize. You might be anxious to get started, but before you rush in and start building servers it is best to make sure you are prepared and fully understand all the many steps that you need to complete as part of building your environment.

Preparing to Build Your Virtual Environment

Before you start with building your environment, prepare everything you will need to install and configure all the various components (as listed here). Doing this ahead of time will help ensure that your installation goes smoothly. It's always a good idea to read through the documentation before you install the software so that you can get a good idea of what to expect when installing it. Here are some steps that you should take to prepare for building your virtual environment:

1. Download the latest version of the ESX or ESXi ISO files from VMware's website.
2. Download the latest version of the vCenter Server installation application from VMware's website.
3. Generate and download your license files from the VMware licensing website.
4. Make sure you have a database server set up to be used with vCenter Server.

- Download the Release Notes and Installation documentation from VMware's website for the version of ESX or ESXi and vCenter Server that you are installing.

About MD5 Checksums

All installation files downloaded from VMware's website contain MD5 checksums (md5sum) to ensure that the integrity of the downloaded file has not been compromised. The md5sum is derived by taking the input file and producing a message digest that serves as an almost unique digital fingerprint and consists of a 128-bit MD5 hash that appears as a series of numbers and letters. When you download a file from VMware's website, the md5sum will be displayed next to the file to be downloaded, as shown in Figure 3.1.

Download VMware ESX 3.5 Server To use this product, you will need to redeem your license activation codes that you should have received by email during your purchase of VMware Infrastructure 3. For questions about license redemption and activation, please email vi-hotline or view the License Redemption Workflow Overview . Latest Released Version: 3.5 Update 3 11/06/08 119826		
Localized language pack for ESX Server	74 MB	.zip (md5sum:28f3cdfb07af4a6e8a5eca13471296)
ESX Server 3.5 Update 3 CD image Refresh Boot your server with this CD in order to install ESX Server 3.5 Update 3	600 MB	.iso (md5sum:e9bdaad2d37872820a4cad8e8dbde536)
Upgrade package from ESX Server 2.x to ESX Server 3.5 Update 3 Refresh This package provides expert users with remote, headless upgrade capability from ESX Server 2.x to ESX Server 3.5 Update 3. Most normal upgrades can be accomplished using the standard CD image.	482 MB	.tar.gz (md5sum:2da08fed15bd4b1ed5b19433e837591c)
Upgrade package from ESX Server 3.0.x to ESX Server 3.5 Update 3 Refresh This package provides expert users with remote, headless upgrade capability from ESX Server 3.0.x to ESX Server 3.5 Update 3. Most normal upgrades can be accomplished using the standard CD image. NOTE: Need Patch ESX-1003519 for ESX 3.0.1 and ESX-1003525 for ESX 3.0.2/3.0.2 U1.	843 MB	.zip (md5sum:d631aa8418d99fce4280fc3905ac4c37)
Upgrade package from ESX Server 3.5 to ESX Server 3.5 Update 3 Refresh This package provides expert users with remote, headless upgrade capability from ESX Server 3.5 to ESX Server 3.5 Update 3. This package requires the esxupdate utility. Most normal upgrades can be accomplished using the standard CD image.	418 MB	.zip (md5sum:4dea5d943d0c0469c397b6520dfef0fb)

Click [here](#) for information about using MD5 checksums.

Figure 3.1 Md5sum listed on the download page

Note the original md5sum. When the download completes, you can run a utility to check the md5sum on the downloaded file to make sure it matches the original. Most UNIX and Linux systems have a built-in md5sum utility to do this. A Windows version of this utility can be downloaded from www.fourmilab.ch/md5/. To validate the md5sum, just enter `md5sum filename` on a Linux system. If you are using the Windows version, enter `md5 filename`. After a few seconds, a checksum will be returned that should match the one provided on the download page, as shown in Figure 3.2.

```
C:\vmware>md5 esx-3.5.0_Update_2-110268.iso  
71F61A69289514377EB3025ABB36CD1F esx-3.5.0_Update_2-110268.iso  
C:\vmware>
```

Figure 3.2 Computing the md5sum on a downloaded file

Comparing md5sums is not mandatory, but does provide a good verification check to ensure the integrity of any files that you download and is a good best practice.

Choosing and Configuring a Database Server for vCenter Server

The vCenter Server database contains a variety of information, including host and virtual machine (VM) configurations, alarm and event data, historical performance data, HA/DRS/resource pool/cluster data, and task information. When ESX hosts are added to vCenter Server, the host and VM configuration is automatically imported into the database. vCenter Server continually updates its database by polling the ESX servers; so if changes are made to the host directly, vCenter Server stays up-to-date.

The database is not critical to the operation of ESX servers or their VMs; they would continue to function normally if vCenter Server or its database were unavailable (except for DRS and VMotion, which would not work; HA would still work, but you would not be able to make HA configuration changes to it). If the database were to crash and a new one were created, you could add your ESX servers back in, and vCenter Server would repopulate the configuration information. The only data unique to the database that would be lost in the event that the database was rebuilt is performance statistics, alarms, events, tasks, clusters, HA and DRS configurations, resource pools, roles, permissions, and custom attributes.

The size of the database will vary based on the number of hosts and VMs managed by the vCenter Server and the level and interval of performance statistics being collected. Most of the tables in the database are smaller and do not grow that much, with the exception of the tasks, events, and performance statistics tables, which can grow quite large over time. More detailed information about the database tables that vCenter Server uses is provided in Chapter 11, “Advanced Topics.” The database will initially start out small and grow as more and more performance statistics are collected. Therefore, it is important to plan to have adequate disk space available on your database server for your database to grow. You can change the amount, interval, and level of statistics that are collected in the vCenter Server configuration settings, as shown in Figure 3.3.

As the settings are decreased or increased, the total size of the database will also change. Starting with vCenter Server 2.5, a database size estimator was built in to the vCenter Server that displays the total projected database size based on the current settings. Changing the statistic level from the default Level 1 can drastically increase the size of the database, as shown in Figure 3.4, and is recommended only for short intervals when troubleshooting problems in your environment.




Figure 3.3 vCenter Server statistics collection settings




Figure 3.4 vCenter Server statistics collection settings changed to Level 2

By the Way

You have four statistic levels to choose from. The higher levels provide increasingly more statistics, as outlined here:

Level 1: Includes the basic metrics: average usage for CPU, memory, disk, and network; system uptime, system heartbeat, and DRS metrics. Statistics for devices are not included in this level.

Level 2: Includes all metrics for CPU, memory, disk, and network counters (average, summation, and latest rollup types [maximum and minimum rollup types excluded]); system uptime, system heartbeat, and DRS metrics. Statistics for devices are not included in this level.

Level 3: Includes all metrics (including devices) for all counter groups (average, summation, and latest rollup types [maximum and minimum rollup types excluded]).

Level 4: Includes all metrics supported by vCenter Server.

Also, changing the collection intervals and the length of time that samples are kept can drastically increase the size of the database, as shown in Figure 3.5.




Figure 3.5 vCenter Server statistics collection settings interval and retentions changed

For medium to large production environments, it is recommended to install either the Standard or Enterprise versions of SQL Server 2000 or 2005 or an Oracle Standard or Enterprise 9i or 10g database. SQL Server 2005 Express is recommended only for small installations of up to 5 ESX hosts and 50 VMs. The database requirements for vCenter Server are listed here. Because these can change, it is best to check the installation documentation and release notes for the version of vCenter Server that you are installing:

■ **VirtualCenter 2.0.x¹**

Microsoft SQL Server 2000 SP4

Microsoft MSDE (not supported for production environments)

Microsoft SQL Server 2005 SP2 (starting with VirtualCenter 2.0.2)

Oracle 9i Standard or Enterprise Release 2

Oracle 10g Standard or Enterprise Release 1 (10.1.0.3)

Oracle 10g Standard or Enterprise Release 2

■ **VCenter Server 2.5.x**

Microsoft SQL Server 2000 Standard or Enterprise SP4

Microsoft SQL Server 2005 Enterprise SP1 or SP2

Microsoft SQL Server 2005 Express SP2

Oracle 9i Standard or Enterprise Release 2 (9.2.0.8)

Oracle 10g Standard or Enterprise Release 2 (10.2.0.1)

Which database you choose will mostly depend on what already exists in your environment and which one you have expertise in administering. If you do not have either database currently running in your environment, SQL Server is recommended because of its lower cost and because it is easier to set up, configure, and maintain than Oracle. It's possible to install the database server on the same server as vCenter Server, but this is not recommended because of resource contention between the two applications. You can also consider installing your database server on a VM if you do not have a physical server available for it.

After you have chosen your database platform, read through the latest release notes and installation guide for the version of vCenter Server that you are installing for specific information about setting up a database to be used with vCenter Server. If you are going to be running your database on a separate server from your vCenter Server, which is recommended, you need to set up an ODBC connection on the vCenter Server with information to connect to the database server, as described later.

By the Way

For maximum reliability and performance, it is recommended that your vCenter Server and database server reside on the same network LAN segment.

Oracle Considerations

If you plan to use Oracle for your vCenter Server database, be aware of the following considerations:

- You need to install the Oracle client software, which includes the ODBC driver on the vCenter Server. Do not use the built-in Microsoft ODBC driver for Oracle. Make sure the client version is one supported by vCenter Server.
- Edit the TNSNAMES.ORA file that is installed so that it contains the configuration information (database, host, and port name) for the Oracle server that you are connecting to. Optionally, you can copy an existing TNSNAMES.ORA file that contains the database information from another server. After you have done this, you can test connectivity to the database by using the TNSPING command.
- When creating the tablespace, ensure that you size it correctly for the number of hosts and anticipated statistics. Use the size calculator that is built in to vCenter Server or the one that is available on VMware's website (www.vmware.com/support/vi3/doc/vc_db_calculator.xls). It's also a good idea to have the table auto-extend so that it automatically grows when it reaches the initial size that was specified when it was created.
- When you create the vpxadmin user in Oracle, make sure that it has the proper permissions that it needs. The required permissions are listed in the installation documentation. You can use either the DBA permission, which provides all the needed permissions, or assign them individually to the user.

Microsoft SQL Server Considerations

If you plan to use Microsoft SQL Server for your database, be aware of the following considerations:

- Do not use the SQL Server Master database for your vCenter Server tables. Create a separate database for vCenter Server to use.
- Windows NT authentication is not supported for remote SQL servers. Use SQL Server authentication instead.
- When you create a SQL user, make sure you give it database operator (DBO) rights and set the default database to the one that you created for vCenter Server. It is recommended that you do not use the SA user for vCenter Server.
- Check the installation documentation for the specific roles needed for the SQL user account for vCenter Server. Usually, the db_owner role is needed for both the vCenter Server database that was created and the existing MSDB database. The db_owner role can be removed from the MSDB database after vCenter Server has been installed.

because it is used only to create scheduled jobs as part of the installation. It is recommended to not assign the sysadmin role to the vCenter Server SQL user and to instead only give it the access that it needs.

- For SQL 2000, use the SQL Server ODBC driver. For SQL 2005, use the SQL Native Client driver.
- It's best to modify the options for the vCenter Server database that is created to use the Simple Recovery model rather than the Full Recovery or Bulk-Logged models to ensure that the transaction logs do not grow excessively and consume all the database server's disk space.

Setting Up an ODBC Connection

Before setting up your ODBC connection, first make sure the new database and user account are created on your database server. The information you need includes the database name, database username and password, and the database hostname or IP address. If you are using Oracle for your database, you should also install the Oracle client, which also installs the Oracle ODBC driver. If you are using SQL Server 2005 as your database server, first download the SQL Server Native Client from Microsoft (sqlncli.msi) and install it. If you will be using the SQL Server 2005 Express database that comes with vCenter Server, there is no need to perform this step because the vCenter Server installation will automatically install the database, configure it, and create the ODBC data source.

Setting Up an ODBC Connection to a SQL Server Database

To set up an ODBC connection on your vCenter Server to a SQL Server database, follow these steps:

1. On the vCenter Server, open the ODBC Data Source Administrator utility, which is usually located under Administrative Tools.
2. Click the System DSN tab.
3. Click the Add button.
4. Select your driver, as shown in Figure 3.6. For SQL Server 2000, choose SQL Server, and for SQL Server 2005, choose SQL Native Client.
5. Enter a name for your data source. The default when installing vCenter Server is VMware VirtualCenter,¹ so using that name is easiest. You can also optionally enter a description for the data source, and finally enter the hostname or IP address of your database server or select it from the drop down list, as shown in Figure 3.7. Click the Next button after you enter or select your database server.




Figure 3.6 Setting up a SQL Server ODBC connection: choosing a driver




Figure 3.7 Setting up a SQL Server ODBC connection: entering a name and choosing a database server

6. Select an authentication method. If the SQL server is on a separate physical server, select SQL Server Authentication; otherwise, select Windows NT Authentication, as shown in Figure 3.8. If you select SQL Server, enter your SQL server username and password in the bottom fields.




Figure 3.8 Setting up a SQL Server ODBC connection: selecting an authentication method

7. Check the Change the Default Database To box, and select the VCenter Server database name that was created on your database server, as shown in Figure 3.9. Then click Next to continue.




Figure 3.9 Setting up a SQL Server ODBC connection: selecting a default database

8. Click the Finish button, and a summary window will display. It's a good idea to test your data source first, which you can do by clicking the Test Data Source button, which will display a Test Results window, as shown in Figure 3.10.



Figure 3.10 Setting up a SQL Server ODBC connection: Test Results window

Setting Up an ODBC Connection to an Oracle Database

To set up an ODBC connection on your vCenter Server to an Oracle database, follow these steps:

1. On the vCenter Server, open the ODBC Data Source Administrator utility, which is usually located under Administrative Tools.
2. Click the System DSN tab.
3. Click the Add button.
4. Select your driver, as shown in Figure 3.11. Do not choose the Microsoft ODBC driver for Oracle. Select the Oracle ODBC driver that was installed with the Oracle client and is usually called Oracle in OraClient. Click the Finish button after you select your driver.
5. Enter a name for your data source, as shown in Figure 3.12. The default when installing vCenter Server is VMware VirtualCenter,¹ so using that name is easiest. You can also optionally enter a description for the data source. Enter the name of your TNS data source. This is the name that is defined in your TNSNAMES.ORA file that was installed as part of the Oracle client installation. Finally, enter the username for the Oracle user account that was created for vCenter Server. You do not need a password here because this is set as part of the vCenter Server installation.




Figure 3.11 Setting up an Oracle ODBC connection: choosing a driver




Figure 3.12 Setting up an Oracle ODBC connection: entering a name and choosing a TNS data source

6. You can test the connection before you save it by clicking the Test Connection button and entering the password for the Oracle user account that you entered in the preceding step. A small window will be returned indicating whether the connection was successful. If the connection fails, check your TNSNAMES.ORA file to see whether it has the correct database name, hostname, and port. To verify that your TNSNAMES.ORA is correct, use the TNSPING command-line utility (for example, `TNSPING database name`).

After you have set up and configured your database and created your ODBC connection, you are ready to start installing vCenter Server.

Installing vCenter Server

Before installing vCenter Server, make sure you meet all the requirements as outlined in the installation documentation available on VMware's website. Also, be sure to read through the most current release notes and installation guide for the version of vCenter Server that you are installing.

Deciding Whether to Install vCenter Server on a Physical Server or on a Virtual Machine

Before you install vCenter Server, you need to decide whether you want to run it on a physical server or on a VM. Both options are fully supported by VMware, but there are definite advantages and disadvantages to running vCenter Server on a VM.

Here are some of the advantages of running vCenter Server on a VM:

- **High availability.** The HA feature on your ESX host servers provides high availability for the vCenter Server. A hardware failure on a physical server normally would make vCenter Server unavailable until it was repaired or reinstalled on another server.
- **Cost savings.** One less physical server needed to dedicate to vCenter Server. This can be particularly advantageous for small installations that have very few physical servers.
- **Snapshots.** The snapshot feature can prove useful when patching or installing upgrades to the vCenter Server.

Here are some of the disadvantages of running vCenter Server on a VM:

- **Licensing issues.** If you plan to run vCenter Server on a VM, that usually means you will be running the licensing server on it, too. If you have issues with your license server and it is past your 14-day grace period for the licensing server being unavailable to ESX hosts, any VMs you try to power on will not start. If your licensing server is on a VM that cannot be started, you are stuck until you configure an alternative licensing server on a physical machine.
- **Cold migration.** Because cold migration is a vCenter Server function, you cannot cold migrate the vCenter Server because it has to be down to perform a cold migration to another ESX host. (Alternatively, you can use a manual command-line process to accomplish this.)
- **Recovery.** If a major outage strikes that affects all ESX hosts (for example, a SAN failure), you lose centralized management capabilities of ESX hosts while vCenter Server is down. When it comes time to restart VMs, you will need to log in to each host individually to locate which host the vCenter Server VM is on so that it can be started.

None of the disadvantages listed here are big concerns and should not stop you from considering virtualizing vCenter Server. Many administrators are not comfortable with virtualizing vCenter Server and stick to running it on a physical server. There are some definite

advantages to it, however, that should be considered before making the choice. If you do decide to do it, it is a good idea to set up a backup license server on a physical workstation or server to be used in the event of a licensing issue (as described earlier). Alternatively, you can create a host-based license for the ESX host running vCenter Server to be used in case the license server is not available so that the VM can be started. After the vCenter Server and the license server start up, you can switch the host server back to using the license server.

vCenter Server Installation

The vCenter Server installation consists of a few different components. Which ones are installed will be based on the choices you make within the Installation Wizard:

- **VMware vCenter Server.** The main application that manages ESX hosts and installs as a Windows service. This also installs a Tomcat web server to be used for web access to vCenter Server.
- **Virtual Infrastructure Client (VI Client).** The client application that installs on the vCenter Server and can also be installed on stand-alone workstations to connect to an ESX host or vCenter Server.
- **Microsoft .NET Framework.** Version 2.0 of the .NET Framework is required for many of the vCenter Server components, including the VI Client as is automatically installed on the vCenter Server.
- **Microsoft SQL Server 2005 Express.** Automatically installed and configured if you choose to not use an existing database when installing vCenter Server. The Express version is free to use and does not require a license.
- **Update Manager Service.** A component of vCenter Server that provides patching capabilities through vCenter Server for ESX hosts and VMs. This feature is available only if you have an edition that includes the license for it.
- **Converter Enterprise Service.** A component of vCenter Server that provides P2V and V2V conversions of servers through vCenter Server.
- **License Server.** A separate licensing application based on Macrovision's FLEXlm licensing technology that runs as a Windows service and provides centralized licensing for ESX hosts.

Before you begin with the installation, make sure you have the database set up and configured as described earlier. If you choose to instead use the SQL Server Express 2005 database that is included with vCenter Server, there is no need to do this, because the installation program will set up and configure it for you. The steps listed here detail the installation for vCenter Server version 2.5:

1. Download the vCenter Server installation software from VMware's website. It is available in both ISO and zip format. Copy it to the server you want to install it on

- or burn a CD and load it on the server. Be sure to check the md5sum of the downloaded file to ensure its integrity.
2. Click Next past the first few welcome screens. Accept the licensing agreement, and then enter your customer information and click Next.
 3. At the Installation Type screen, as shown in Figure 3.13, you have the option to choose from three different types. The first type, which is the VMware Infrastructure Client (VI Client), is used for installing the client software that connects to either ESX hosts or vCenter Servers. This type is used only if you want to install the client software. The second type is the VMware VirtualCenter¹ Server installation, which installs both the Server and VI Client software and is the one you want to choose. Optionally, the last type is a Custom installation that lets you choose which components you want to install (VI Client/VC Server/Update Manager/Converter). Click Next to continue.



Figure 3.13 vCenter Server installation: Installation Type screen

4. At the Database Selection screen, choose between using Microsoft SQL Server 2005 Express or using an existing database server, as shown in Figure 3.14. If you choose to use an existing database server, enter the name of the ODBC Data Source Name (DSN) that you created earlier, along with the database username and password. Click Next to continue.
5. At the Licensing Server screen, you can select to either use the 60-day evaluation license that vCenter Server comes with or to use an existing licensing server, as shown in Figure 3.15. You would choose the existing license server option if you have an existing FLEXlm server in your environment that you want to use instead of the

one provided with vCenter Server or if you have manually installed the license server before installing vCenter Server. If you choose to use an evaluation license, the licensing server is also installed as part of the vCenter Server installation. If you do choose to use an existing server, enter the port and hostname of it. If you choose this option, a licensing server will not be installed on the vCenter Server. Also, select your vCenter Server edition, which is either the VirtualCenter¹ Management Server or the VirtualCenter¹ Foundation Management Server. Click Next to continue.




Figure 3.14 vCenter Server installation: Database Selection screen



Figure 3.15 vCenter Server installation: Licensing Server screen

- At the Server Authorization screen, enter the vCenter Server hostname (recommended) or IP address and the username and password that you use to log in to the operating system of the vCenter Server. Any authorized extensions (plug-ins) to vCenter Server are listed on this screen, too, for informational purposes, as shown in Figure 3.16. Click Next to continue.



Figure 3.16 vCenter Server installation: Server Authorization screen

- Finally, click the Install button to begin the installation. A warning message may appear stating that the Microsoft .NET Framework is not installed and will be installed first before the vCenter Server installation begins. If you have any Internet Explorer windows open on the vCenter Server, close them so that the .NET Framework can be installed. The various components will be installed, and a status message will update you on which one is currently being installed. The installation will also create and populate all the database elements in whatever database you choose. When the installation has completed, you will see a window with an option to Finish and optionally launch the VI Client.

When the installation of vCenter Server has completed, you should see that the following Windows services have been added and started:

- **VMware VirtualCenter¹ Server.** The main component for vCenter Server
- **VMware License Server.** The licensing server component
- **VMware Infrastructure Web Access.** The web server (Tomcat) component for accessing vCenter Server using a web browser
- **VMware Capacity Planner Service.** The built-in capacity planner for analyzing and converting physical systems to VMs

- **VMware Update Manager Service.** The built-in patching component for ESX hosts and VMs
- **VMware Converter Enterprise Service.** The built-in component to convert physical servers to VMs

The following applications will also have been installed under the VMware Start menu folder, as shown in Figure 3.17:

- **VMware Infrastructure Client (VI Client).** The client application for connecting to ESX hosts and vCenter Servers.
- **VMware Infrastructure Update.** An application that can be used to patch ESXi servers if you do not use Update Manager. This application cannot patch ESX hosts.
- **Generate VirtualCenter¹ log bundle.** A diagnostic utility to collect all the vCenter Server log files to send to VMware support when troubleshooting problems.
- **Generate Update Manager log bundle.** A diagnostic utility to collect all the Update Manager log files to send to VMware support when troubleshooting problems. (Optional)
- In the VMware Web Access subfolder are various applications for configuring and monitoring Tomcat. Normally, this is something you will not need to do.
- In the VMware License Server subfolder is the VMware License Server Tools application, which is used to configure the license server and some documentation.




Figure 3.17 VMware Start menu folder after vCenter Server has been installed

That concludes the installation of vCenter Server. Next we cover installing the VI Client, which is used to connect to vCenter Server and ESX/ESXi hosts from Windows workstations.

Installing the Virtual Infrastructure Client

You will want to install the VI Client on any workstations that you will be using to administer ESX hosts and vCenter Servers. The VI Client can be installed on any Windows system that meets the software requirements and hardware requirements listed here. It cannot be installed on Linux systems.

- Windows 2000 Professional or Server with SP4
- Windows XP Professional with SP2
- Windows Server 2003 with SP1, SP2, or R2
- Windows Vista Business or Enterprise
- Microsoft .NET Framework 2.0

By the Way

The VI Client supported only 32-bit versions of the operating systems listed here until version 2.5 Update 1 was released, which added support for 64-bit versions.

VI Client hardware requirements include the following

- **Processor.** 266MHz or higher Intel or AMD processor, 500MHz recommended
- **Memory.** 256MB RAM minimum, 512MB recommended
- **Disk Storage.** 150MB free disk space for basic installation
- **Networking.** 10/100 minimum, Gigabit recommended

The web access UI that installs on ESX hosts and vCenter Servers provides only limited administration capabilities. It can be used only to manage VMs, not ESX hosts. The VI Client provides full administration capabilities for your whole VI3 environment. With it, you can connect to ESX hosts individually and manage them or connect to a vCenter Server so that you can centrally manage your ESX hosts. Once installed, you can have multiple copies of the VI Client running if you need to connect to different hosts at the same time.

To install the VI Client, you first need to get the installation program. The easiest way to do this is to access either a vCenter Server or an ESX host using a web browser and typing in the hostname or IP address of the server. On the default welcome screen is a link for downloading the client software to your PC, as shown in Figure 3.18.




Figure 3.18 Web access default welcome screen

Because ESX hosts and vCenter Servers may be running different versions of software, it is best to download it from one that you know is running the latest version in your environment. You can also copy the VI Client installation application from the vCenter Server installation CD, from the unpacked installation zip file (located in the vpx subdirectory), or from the C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\docRoot\client¹ directory on the vCenter Server. Optionally, you can run the vCenter Server installation program and select the VMware Infrastructure Client option to install it on a workstation.

Installing the client is pretty straightforward. Just follow these steps:

1. If you do not have the .NET Framework 2.0 installed, the installation will fail with an error message to that effect. Currently, the .NET Framework is not installed by default on any Windows operating systems. You can check to see whether it is installed on your PC by going to the Add or Remove Programs option in Control Panel and seeing whether Microsoft .NET Framework 2.0 exists. If it does, you are all set and can begin the installation. If it doesn't, you can download it by searching for it on Microsoft's website. If you already have a version 1.x of the .NET Framework installed, you do not have to remove it first (because both the 1.x version and 2.0 version can be installed concurrently on your PC).
2. Run the VMware-viclient.exe installation program that you downloaded from your server to begin the installation.

3. When the installation welcome screen appears, click Next, and then accept the licensing agreement and click Next. Enter your customer information and click Next. Either accept or change the default destination folder and click Next. Then, click Install to begin the installation.
4. Click Finish when the installation completes. You will see a VMware folder in your Start menu that contains both the VMware Infrastructure Client and the VMware Infrastructure Update application.

When you start the VI Client, a login screen will appear, as shown in Figure 3.19, where you can enter the name or IP address of the ESX/ESXi host or vCenter Server that you want to connect to and your username and password information. (This is typically either “root” or any local accounts that you configured for ESX/ESXi and the local “administrator” account or any Windows domain accounts that you configured for vCenter Server.)



Figure 3.19 Virtual Infrastructure Client login screen

You may receive an untrusted SSL certificate warning when connecting to a host, which you can ignore and suppress. Optionally, you can install a trusted SSL certificate on your hosts to prevent the messages from occurring. (See VMware knowledge base article 4646606 for information about how to do this.) After you have logged in to the VI Client, you can edit the client settings to customize it for your PC. By default, with version 2.5 the Getting Started tabs are displayed to assist new users with using the client. To edit the client settings, select Edit from the top menu and then Client Settings, as shown in Figure 3.20.




Figure 3.20 Virtual Infrastructure Client settings

You have the option to change the following client settings:

- **Remote Command Timeout.** This is the length of time that the VI Client waits when a command (for example, power on a VM) is selected before timing out if it does not receive a response from the vCenter Server. Can be set to a default value (30 seconds), or a custom value in seconds can be set (usually 60 to 120 seconds). Normally, you will not need to change this timeout unless you are connected over a lower-speed or high-latency network to your vCenter Server.
- **Virtual Machine Consoles.** The maximum number of remote console sessions that can be opened for VMs at one time. The default value of 25 is usually sufficient but can be increased if necessary.
- **Hint Messages.** Hint messages are displayed by default in certain situations, but they can be suppressed when responding to them. The Restore All Hints button restores all hint messages and discards previous responses to hint message questions.
- **Getting Started Tabs.** A new feature that was introduced on vCenter Server 2.5, these additional tabs are displayed when selecting ESX hosts or VMs and offer help and assistance to those who are new to using the VI Client. The Remove All Tabs button removes them for more experienced users, and the Restore All Tabs restores them if you want them back.
- **Tasks and Events (Lists Tab).** The page size defines how many tasks and events will be displayed when clicking the Tasks and Events tab located on the ESX hosts, VMs, clusters, and datacenter pages. The default value is 100, which means only the most recent 100 tasks or events will be displayed. (There is no button to display additional pages.)

You can change this value to a number between 10 and 1,000 to display more events. Even though more than 1,000 tasks or events may exist, there is no way to see more than 1,000 in the VI Client, and you must instead query the SQL or Oracle database directly.

- **Virtual Machine and Hosts (Lists Tab).** The Retrieve Contents before Filtering and Sorting option is useful for lower-bandwidth connections. When this option is selected, VM and host lists are sorted after the information is retrieved from the server rather than before.

In the next chapter, we discuss the VI Client in more detail because it is the application you will use the most while administering your virtual environment.

Installing the Licensing Server

Normally, the license server is installed when installing the vCenter Server application, but there may be situations when you want to install it separately from vCenter Server or on another server. The licensing server runs as a Windows service and can be installed anywhere, but it is usually recommended to install it on the vCenter Server. If you do want to install this separately, the following steps describe how to install it. Before you begin, have a copy of your license file (.lic file); you will be prompted for this during the installation. Copy this to a directory on the server that you are installing it on. If you do not have a copy, you can go to VMware's licensing website and download your licenses into a file.

1. The installation program for the license server is bundled with the vCenter Server ISO and zip file. Browse through the extracted zip file or the install CD and the license server installation application, which is named VMware-licenseserver.exe and located in the vpx subdirectory. Run the application to begin the installation.
2. When the installation welcome screen appears, click Next. Then accept the licensing agreement and click Next. Either accept or change the default destination folder and click Next.
3. At the Licensing screen, enter the full path and filename of your licensing file or click the Browse button to select it, as shown in Figure 3.21, and then click Next.
4. Click the Install button to begin the installation. When the installation is complete, click the Finish button. A Windows service called VMware License Server will have been installed and started, and under the VMware start menu folder in the VMware License Server subfolder is the VMware License Server Tools application, which is used to configure the license server and some documentation. We cover how to configure the license server in the next chapter.




Figure 3.21 Selecting a license file when installing the license server

Installing ESX and ESXi

Before you install ESX or ESXi, download the ISO file from VMware's website for the version that you want to install and burn the ISO to a CD that you can mount on the server.

Optionally, if you are using a remote console utility like HP's Integrated Lights Out (ILO) or Dell's Remote Access Card (DRAC), you can mount the ISO or map a workstation CD-ROM as a virtual CD-ROM on the server from a remote workstation. This is not the preferred method, however, because data corruption may occur if the server encounters too much load while the installation is running. If you do choose this method, it is recommended that you run the media test that is built in to the ESX installer and that you are on the same network LAN segment as the target host, to ensure adequate bandwidth and lower latency. If you are not using the boot from SAN feature, it is best to disconnect the server that you are installing ESX or ESXi on from the SAN fabric before installing. When the install completes, you can reconnect it before you reboot the server.

Preparing the Server for Installation

Before you proceed with installing ESX or ESXi on your server hardware, always test or "burn in" the memory to ensure that there are no defects with it. Defective memory will usually not be noticeable when you first deploy a new server, and it might be months before you see signs of defective memory, which can cause hard server crashes or the dreaded purple screen of death (PSoD).

Most servers perform a brief memory test when they first start up as part of their POST procedure. This is not a very good test and will detect only the most obvious of memory problems. A much more thorough test checks the interaction of adjacent memory cells to ensure that writing to one cell does not overwrite an adjacent cell. There is a good free memory test utility available called Memtest86+ that performs many different tests to thoroughly test your memory. You can download it from www.memtest.org as a small 2MB ISO file that you can burn to a CD and then boot from it on your new server and let the memory burn in for at least 24 hours or at least until one complete pass finishes. However, it is best to let it run for several passes to more thoroughly test your server's memory.

Memtest86+ will run indefinitely, and the pass counter will increment as all the tests are run. The more RAM you have in your system, the longer it will take to complete one pass. A system with 32GB will generally take about one day to complete. Memtest86+ not only tests your system's RAM but also the CPU L1 and L2 caches. If it detects an error, the easiest way to identify the memory module that caused it is to just remove a DIMM and run the test again and repeat until it passes. There is good documentation on Memtest86+ that includes troubleshooting methods, detailed test descriptions, and the causes of errors.

Performing this extra step ahead of time is worthwhile to ensure that your server is not affected by memory defects later.

Boot from SAN Considerations

If you are installing ESX to a LUN to utilize the boot from SAN, you should be aware of several considerations:

- ESX does not support booting from a shared LUN. If you want an ESX Server host to boot from a SAN, you should allocate an entire LUN to each ESX Server host. Make sure that other boot LUNs for other ESX servers are masked so that each ESX host only sees its own boot LUN.
- It is best to present only the LUN to which ESX is being installed onto until after the installation is complete. Do not present any data LUNs for VMFS storage until after ESX is correctly installed.
- Configure the HBA BIOS to enable the boot from SAN option and select the LUN to be used for booting before installing ESX.
- Make note of the World Wide Name (WWN, the SAN equivalent of a network MAC address, which consists of 16 hexadecimal digits grouped into 8 pairs) for the HBA adapter. You will need it when configuring the SAN.

ESX Partition Considerations

Partitions let you segment a physical drive into separate logical drives. Partitions are treated as independent disks with their own defined sizes and are assigned mount points so that

they can be accessed from the default root partition. Partitions are used to ensure that unchecked log growth does not affect other critical operating system partitions. The default partitions on a Linux system are the `/boot` partition, which stores files used to boot the computer, and the `/` partition, which is the root partition from which all mount points derive. The additional partitions used by ESX are set to mount automatically when the server boots, and so they are accessible from the root partition.

When you install ESX, you have the option to use the default partition sizes or change them. The default sizes will typically work for many environments, but there are changes you can make to ensure that you do not run into problems later if your partitions run out of space or you try to increase the memory allocated to the Service Console. First, let's cover the partitions created by default:

- **/boot.** The boot partition contains files used to boot the ESX server. By default, this partition is 100MB.
- **/.** The root partition contains the ESX server operating system files. By default, this partition is 5GB.
- **swap.** This partition does not have a mount point and is used by the Service Console as swap space for virtual memory. By default, this partition is 544MB.
- **/var/log.** This partition is where the ESX server log files are stored. By default, this partition is 2GB.
- **vmkcore.** This partition serves as a repository for VMkernel core dump files in the event that a core dump occurs. By default, this partition is 100MB.
- **vmfs.** The remaining free space after all the other partitions are created on the drive ESX is installed on is used to create a VMFS volume.

The following changes are recommended for the default partitions:

- **/boot.** Change from 100MB to 384MB to allow extra space for any future ESX upgrades.
- **/.** Okay to leave the root partition at 5GB.
- **swap.** Change from 544MB to 1600MB. This should be twice the amount of memory that is dedicated to the Service Console. The default amount of memory devoted to the Service Console is 272MB, which is why the default swap size is 544MB. The recommended amount of memory for the Service Console is the maximum of 800MB, which would require a 1600MB swap partition.
- **/var/log.** Change from 2GB to 5GB to allow for extra space for log files.
- **vmkcore.** Okay to leave this partition at 100MB.
- **vmfs.** This will decrease in size as the other partitions increase in size.

By the Way

It is recommended to not create a VMFS partition during the ESX installation and to instead create it afterward using the VI Client to ensure that it is properly aligned. You can do this by deleting the default vmfs3 partition during the ESX installation at the Partition Disks screen.

In addition, you can create the following partitions to further segregate your drive to help protect the root directory from filling up:

- **/home.** Create a partition of 2GB for any home directories created for local users on the ESX host.
- **/tmp.** Create a partition of 2GB for the directory used to store temporary files.
- **/var.** Create a partition of 4GB for the directory used to hold administrative log and configuration files.

We cover how to make these changes during the ESX installation steps later in this chapter. Table 3.1 provides a summary of the default sizes compared to the recommended partition size changes.

Table 3.1

ESX Default and Recommended Partition Size Comparison

Partition Name	Default Size	Recommended Size
/boot	100MB	384MB
/ (root)	5GB	5GB
Swap	544MB	1600MB
/var/log	2GB	5GB
Vmkcore	100MB	100MB
/home	N/A	2GB
/tmp	N/A	2GB
/var	N/A	4GB

ESX Installation Steps

When you are ready to install ESX, just follow these steps to complete the installation:

1. When you boot from the ESX installation CD or ISO file, the ESX installer loads, and you are presented with a welcome screen that provides you with the option of using either the graphical installer or a text-based installer, as shown in Figure 3.22. If you do not choose an option, the graphical installer automatically starts after a period of time. The text-based installer can be used if you experience mouse or video problems when using the graphical installer. To access the text mode, type in `esx text` at the welcome screen prompt. These instructions are based on the graphical installer, but most of the options are similar with the text-based installer.



Figure 3.22 ESX Server 3.5 initial welcome screen

2. The first screen to appear will be the CD Media Test screen, as shown in Figure 3.23, which is recommended if you have burned a CD using your own media instead of using media supplied by VMware. Skipping this test can cause your ESX installation to become corrupted if the CD burn process encounters errors or the physical media has defects. When the media test completes and you receive confirmation that the media is okay to use, click OK.



Figure 3.23 ESX Server 3.5 Media Test screen

By the Way

You only need to do a media test the first time you use an ESX installation CD that has been burned. You can skip this test for any subsequent ESX installations.

3. At the welcome screen, click Next.
4. Select the appropriate keyboard that applies to you and click Next.
5. Select the appropriate mouse that applies to you and click Next.
6. A message may appear warning that the partition table is unreadable for any of the drives that the server has access to and asking whether you want to initialize the drive to create new partitions, which will erase all data on it, as shown in Figure 3.24. Select Yes for the drive that you are installing ESX on.
7. At the Licensing Agreement screen, check the box to accept the license agreement and click Next.
8. At the Partitioning Options screen, you have the option to use the Recommended partitioning options or to select the Advanced option to modify the default partition sizes, as shown in Figure 3.25. If you choose Recommended, you also have the option to select which disk ESX is installed on. The disk selections will include all disks that the ESX host can see, including SAN and local disks. Make sure you have

the correct disk selected that you want ESX installed on, because the ESX installer may not always default to the local disk on the server. (It might sometimes default to a SAN volume.) If you choose the Advanced option, you must manually create all the required and optional ESX partitions. After you make a selection and click Next, a warning message will appear confirming that all non-VMFS partitions on the disk you selected will be removed and asking whether you are sure. Click Yes. The partition sizes will be displayed in the next window, which will give you an opportunity to change the sizes.



Figure 3.24 ESX Server 3.5 Partition Warning screen




Figure 3.25 ESX Server 3.5 Partitioning Options screen

By the Way

Selecting the option to keep VMs and the VMFS that contains them will preserve any existing VMFS partitions and not wipe them out as part of the ESX installation. This is useful if you are reinstalling ESX and do not want to lose your existing VMFS partitions and VMs.

- At the Partition Disks screen, you can either accept the default recommendations, as shown in Figure 3.26, or make changes to them, as outlined in the previous section. By default, all extra unused space is assigned to a VMFS volume that is created. To change the existing partitions that are displayed, select them and click the Edit button. When the partition is displayed, type in a new size in megabytes and select the Fixed Size option. To add an additional partition, click the New button. Then enter a mount point (for example, /home), enter a size in megabytes, and select the Fixed Size option. Any size changes to existing partitions or newly added partitions are automatically subtracted from the VMFS partition so that the total size of all partitions is equal to the total size of the physical drive. If at any time you want to start over with the default partition sizes, just click the Reset button, which sets all the partitions back to the defaults. After you have made the recommended changes, your partitions should look similar to those in Figure 3.27. After you have finished changing your partition sizes, click Next to continue.




Figure 3.26 ESX Server 3.5 Partition Disks screen with default settings




Figure 3.27 ESX Server 3.5 Partition Disks screen with recommended settings

10. The Advanced Options screen displays next. Here, you change the default boot-loader configuration, as shown in Figure 3.28. In most cases, you will not need to change anything on this screen. Click Next to continue.
11. At the Network Configuration screen, the server's network interface cards (NICs) will all be displayed in a drop-down list, and you have the option to configure them to use DHCP to automatically obtain an IP address or to manually set the IP configuration, as shown in Figure 3.29. This network configuration is for Service Console networking. Select a NIC that you would like to use for the Service Console vSwitch and enter the IP configuration for that NIC. It is recommended to not use DHCP for this because if the IP address of the ESX Service Console were to change later it can cause problems. If you do need to use DHCP, make sure you set an IP address reservation for the ESX host on the DHCP server. If you are using 802.1Q VLAN trunking on the NIC that the Service Console will be using, set the appropriate VLAN ID; if not, leave it blank. This information is used to create a vSwitch for the Service

Console network and to set up a vswif interface for it with the IP address you enter. If you also select the option to create a VM network, a port group is also created on the vSwitch for your VMs. All of these settings can be changed later after the ESX host is started. When you are finished, click Next to continue.

By the Way

The Service Console vswif interface uses a generated MAC address and not the default MAC address of the physical NIC. If you do use a DHCP reservation to obtain an IP address for the Service Console interface, you have to do it after the ESX host is installed because you will not know the MAC address of the vswif interface until after ESX has booted the first time. You can find the MAC address after ESX has booted by connecting to the ESX host with the VI Client and selecting Configuration, and then Networking, and selecting the vSwitch properties for one with the vswif interface (usually vSwitch0). Next in the Port List, select the Service Console, and the MAC address that has been assigned to it will be displayed on the right.




Figure 3.28 ESX Server 3.5 Bootloader configuration screen

ESX Server 3.5

Network Configuration
Select and configure the network interface card that is used for console communication.

Network Interface Card
Device: 1:2:0 - tg3 - NetXtreme BCM5703 Gigabit Ethernet

Network Address and Host Name
☐ Set automatically using DHCP
☒ Use the following network information:

IP Address	172	20	11	66
Subnet mask	255	255	255	0
Gateway	172	20	11	1
Primary DNS	172	20	3	50
Secondary DNS	172	20	3	51

Host name: esx1.xyz.com Enter a fully qualified host name (e.g. host.vmware.com)

VLAN Settings
VLAN ID: (Leave blank if you are unsure whether your network requires a VLAN ID)

☒ Create a default network for virtual machines

Back Next Cancel

Figure 3.29 ESX Server 3.5 Network Configuration screen

12. At the Time Zone Selection screen, select the appropriate time zone that the ESX host is located in. You can either select it from the map that is displayed or click the Location tab and select a time zone from the list, as shown in Figure 3.30. When you are finished, click Next to continue.
13. At the Account Configuration screen, enter a password for the root user account, as shown in Figure 3.31. Optionally, you can add additional user accounts, which can also be added later after the ESX host is started. It is recommended that you create at least one nonroot user account that can be used to remotely connect to the host using SSH and then using the `su -` command to elevate access to it so that you can perform administration commands. When you are finished, click Next to continue.
14. A summary screen will be displayed listing all the options that you selected, as shown in Figure 3.32. Review the options. If necessary, go back and make changes, and then click Next to begin the installation.




Figure 3.30 ESX Server 3.5 Time Zone Selection screen




Figure 3.31 ESX Server 3.5 Account Configuration screen




Figure 3.32 ESX Server 3.5 installation summary screen

15. The installation will begin and will display a status as it completes, as shown in Figure 3.33.
16. When the installation completes, you will see a screen that says that the installation is complete, as shown in Figure 3.34.
17. Click the Finish button to have the ESX host reboot, which will display the console login screen, as shown in Figure 3.35. After the ESX host reboots, you can log in to the physical console using the root username and password that you set.
18. You can also connect to it using the VI Client. If you do not have the VI Client installed on a PC yet, you can download it by entering the hostname or IP address of the server into a browser and clicking the download link that displays on the welcome page. Optionally, if you are using vCenter Server, you can connect to your vCenter Server and add it by selecting a datacenter or cluster and clicking the Add Host button and entering the ESX hostname and root username/password, as shown in Figure 3.36. When vCenter Server connects to the ESX host, it creates a user account called vpxuser that is used to log in to the host from that point forward.




Figure 3.33 ESX Server 3.5 Installing Packages screen



Figure 3.34 ESX Server 3.5 Installer Complete screen




Figure 3.35 ESX Server 3.5 ESX console screen after installation completes

Did You Know?

The `vpuser` account that vCenter Server creates on every ESX host for management operations has a pseudo randomly generated password. A unique 32-character 256-bit password is generated for each ESX host, which is then encrypted using a 1024-bit RSA encryption key to ensure that it cannot be compromised.




Figure 3.36 ESX Server 3.5 Adding a Host to vCenter Server screen

At this point, the installation is complete. In the next chapter, we cover configuring your ESX host in detail.

ESXi (Installable Version) Installation Steps

As previously mentioned, ESXi comes in two versions: the embedded version, which is built in to some new servers; and the installable version, which can be installed on any server that meets the ESXi requirements. When you are ready to begin the installation, just follow these steps:

1. When you boot from the ESXi installation CD or ISO file, a welcome screen will be displayed, as shown in Figure 3.37. Press Enter to begin the installation.




Figure 3.37 ESXi Server 3.5 welcome screen

2. At the EULA screen, press F11 to accept the licensing agreement, as shown in Figure 3.38.
3. At the Disk Selection screen, select a disk to install ESXi on and press Enter to continue, as shown in Figure 3.39. You may receive a warning message that the disk contains an existing partition and will be overwritten. If so, make sure you have selected the correct disk and press Enter to continue.




Figure 3.38 ESXi Server 3.5 EULA screen




Figure 3.39 ESXi Server 3.5 Disk Selection screen

4. At the Confirmation screen, press F11 to begin the installation, as shown in Figure 3.40.
5. The installation will begin, and the progress bar will update, as shown in Figure 3.41.

- When the installation completes, a screen will display confirming that the installation was successful, as shown in Figure 3.42. Remove the CD from the server, and press Enter to reboot.




Figure 3.40 ESXi Server 3.5 Confirmation screen




Figure 3.41 ESXi Server 3.5 Installing screen




Figure 3.42 ESXi Server 3.5 Installation Complete screen

The ESXi installation is now complete. When the server reboots, it will display its login screen, as shown in Figure 3.43, with options for configuring it from the console (F2) or to shut down and restart the server (F12). The server will initially attempt to use DHCP to get an IP address for the Management network. If you are not using DHCP, it will fail, and you can manually configure it. After you have the Management network configured, you will be able to connect to the ESXi host using the VI Client. The next chapter covers configuring your ESXi host in detail.




Figure 3.43 ESXi login screen

Summary

In this chapter, we covered how to install the various components that make up VMware VI3. If you make a mistake while installing something or want to configure something different afterward (for example, ESX partitions), you can easily reinstall it later. I recommend running through the installations several times before you install to your production environment so that you can gain experience with it and play around with the various configuration options. After you have your installation down pat, make sure to document it so that your other hosts are consistent and so that others who may be installing hosts in your environment choose the same options that you do. In the next chapter, we cover how to configure the many options and resources in your environment so that you can prepare it to create VMs.

Endnotes

1. The use of the term *VirtualCenter* here refers to the old name, which still applies to VirtualCenter version 2.0.x and is still present inside the application in vCenter Server 2.5 (because the software has not yet been updated to reflect the new name).

This page intentionally left blank

Chapter 4

Configuring Your Virtual Environment

After the various components of your new environment have been installed, it is time to configure them to prepare them for VMs. The biggest and most important steps involve configuring your virtual networking, storage, and resources. Beyond that, many other items need to be configured in your environment to fine-tune it from the default parameters.

In this chapter, we cover how to take the default installation and configure it so that it meets your needs and requirements. There many different parts of your virtual environment to configure, from the big things such as licensing, High Availability (HA), Distributed Resource Scheduler (DRS), storage, and networking, to the little things such as Network Time Protocol (NTP) time and Domain Name System (DNS). All of this configuration is done via the Virtual Infrastructure Client. So before we get started, we first give you a tour of it.

A Tour of the Virtual Infrastructure Client

The Virtual Infrastructure Client (VI Client) is the Windows application that is provided with ESX/ESXi and vCenter Server and is the main tool that you will use to configure and manage your VI3 environment. Therefore, it is important to understand how to properly use it so that you can best manage your host servers and VMs. In the preceding chapter, we discussed how to install the VI Client. Now we cover how to use it in detail. Before you begin configuring your environment, let's take a tour of the VI Client, which is used to connect either directly to ESX hosts and manage them individually or to a vCenter Server to centrally manage your ESX hosts.

VI Client Layout

The VI Client is made up of three window panes, as shown in Figure 4.1.




Figure 4.1 VI Client

The left pane is the main pane where all objects are displayed. It supports dragging and dropping of objects. So, if you create a new cluster and want to move an ESX host into it, you just drag the ESX host object to the cluster object, and it will be moved. The information displayed in the left pane can be changed by right-clicking in the pane below all of your objects that are listed and selecting one of the four view options (Hosts and Clusters, Virtual Machines and Templates, Networks, Datastores). Optionally, you can also change the view by clicking the Inventory button or selecting the View option from the top menu.

The Hosts and Clusters view displays datacenter, cluster, host, and virtual machine (VM) information in a hierarchical format. You can suppress the display of VM objects in the left pane by right-clicking in the pane below all of your objects that are listed and unchecking the Show Virtual Machines option (also from the View menu). When an object is selected in the left pane, information about that object will display in the right pane.

The Virtual Machines and Templates view displays VM and template objects only (no hosts). It also allows you to put your VMs into folders so that they display in a more organized form. The folders will display only in the Virtual Machines and Templates view. They will not display in the Hosts and Clusters view. Folders are a great way to categorize your VMs.

For example you can create folders to organize your VMs by their operating system type or by their business classification (dev/test/production). You can also create alarms on folders so that they generate alerts on all the VMs in a particular folder.

To create folders, you just right-click the top-level Virtual Machines and Templates object (or your Datacenter object) and select New Folder. You can also create subfolders under folders if you want to nest your folders for better organization. After you've created folders, you can simply drag VMs into the folders to keep them organized. A VM can only ever be in one folder at a time. You can suppress the display of VM objects and templates in the left pane by right-clicking in the pane below all of your objects that are listed and unchecking the Show Virtual Machines or Show Templates options (also from the View menu).

The Networks view displays all the defined networks in your datacenter on all your host servers, as shown in Figure 4.2. The networks displayed are taken from all the virtual switch (vSwitch) configurations on your hosts and are a combination of all the network names that have been assigned to vSwitch port groups. When you select a network in the left pane, it then displays information in the right pane. The Virtual Machine tab displays all the VMs on that network, and the Hosts tab displays all the hosts that have the selected network defined.




Figure 4.2 Networks view

Finally, the Datastores view displays all the VMFS datastores that are configured on your ESX hosts, as shown in Figure 4.3. This includes both local and shared storage VMFS datastores, but does not include any raw device mappings (RDMs) that may be attached to your VMs. The reason for this is that RDMs are attached directly to VMs; they can be displayed in the Datastores view in the VI Client. When you select a datastore in the left pane, it then displays information in the right pane. The Virtual Machine tab displays all the VMs that use that datastore, and the Hosts tab displays all the hosts that have the selected datastore defined.




Figure 4.3 Datastores view

The right pane displays information based on what is selected in the left pane. If you select a datacenter object in the left pane, the right pane will display information about the whole datacenter. If you select a cluster in the left pane, the right pane will only display information about hosts in the cluster, and if you select an ESX host in the left pane, the right pane will only display information about that ESX host. The right pane is also tabbed, so it can display various information depending on which tab you select. Table 4.1 describes the various tabs shown in the right pane, along with the parent object that must be selected for it to show and whether it displays when you connect to ESX/ESXi or vCenter Server. What tabs, actions, and objects you see inside the VI Client depends on your access permissions. You might not see all tabs if you do not have the required access for it.

Table 4.1

VI Client Tabs and Objects

Tab Name	Description	Parent Object	ESX	VC
Getting Started	This tab was introduced with vCenter Server 2.5/ESX 3.5 and contains helpful information for those new to using the VI Client. This tab can be closed and hidden.	All	Yes	Yes
Summary	Provides general summary configuration information and a list of commands for the item selected. Also, provides resource usage information, and when selecting a cluster provides information about HA and DRS configurations.	Datacenter Cluster Host Virtual Machine	Yes	Yes
Datacenters	Provides a list of all datacenters managed by the vCenter Server with the total number of hosts and VMs in each.	Hosts and Clusters	No	Yes
Virtual Machines	Provides a list of all VMs on the object selected, along with their power state, status, host CPU, memory usage, and other information. Additional columns can be displayed by right-clicking the column headings and selecting more from the list that is displayed.	Hosts and Clusters Datacenter Cluster Host	Yes	Yes
Hosts	Provides a list of all host servers on the object selected along with the state, status, host CPU, memory usage, and other information.	Hosts and Clusters Datacenter Cluster	No	Yes

Tab Name	Description	Parent Object	ESX	VC
DRS Recommendations	Displays all recommendations when DRS is configured for manual or partially automated mode. Also, displays DRS Action History, which shows all the DRS actions that have taken place.	Cluster	No	Yes
Resource Allocation	Shows cluster resource configuration settings for VMs (CPU and memory shares/limits/reservations). Will also show for ESX host servers if connected to them directly.	Cluster	Yes	Yes
Performance	Displays real-time or historical performance statistics for the object selected. Configurable so that you can display a certain time period and various different counters such as CPU, disk, memory, network, or system. When connected directly to an ESX host will only display statistics real time and for the past 24 hours.	Cluster Host Virtual Machine	Yes	Yes
Configuration	Allows configuration of host hardware (storage, network, memory, and so on) and software (license, time, DNS, and so forth) settings.	Host	Yes	Yes
Users and Groups	Displays all local host users and groups on ESX and ESXi servers when connected directly to them with the VI Client. Users and groups can be added, edited, or deleted.	Host	Yes	No
Tasks and Events	Displays all tasks and events that have taken place for the selected object. The view can be switched between the two by clicking the button at the top. Columns can be sorted by clicking the column headings. By default, only the last 100 tasks or events display. To see more or fewer, you can edit the client settings and increase the Tasks and Events page size.	All	No	Yes


continues...

Table 4.1 continued

Tab Name	Description	Parent Object	ESX	VC
Events	Displays all events that have taken place for the selected object. Columns can be sorted by clicking the column headings. By default, only the last 100 tasks or events display. To see more or fewer, you can edit the client settings and increase the Tasks and Events page size.	Host Virtual Machine	Yes	No
Alarms	Allows you to view triggered alarms and to configure alarm settings for the selected object. The view can be switched by clicking the Triggered Alarms and Definitions buttons at the top. Alarm definitions inherit in a hierarchical fashion. If you set an alarm at the cluster level, all objects in that cluster are covered by that alarm.	All	No	Yes
Console	Opens up a remote console connection to the selected VM. Console connections can be opened within the tabbed page or in a separate window by right-clicking a VM and selecting the Open Console option, or from the Summary tab of the VM.	Virtual Machine	Yes	Yes
Permissions	Enables you to assign permissions to users and groups. Permissions allow for granular access and control of all objects on the vCenter Server or ESX host. Permissions set inside vCenter Server use Active Directory users and groups. Permissions set on ESX hosts directly use local users and groups.	All	Yes	Yes
Maps	Displays a graphical representation of the relationships between inventory objects. Can be configured to display only certain relationships (for example, host to VM, VM to network). Maps can be viewed, printed, and exported and are accessible by clicking the Maps button at the top of the VI Client. The Maps tab for a VM displays a VMotion map, where you can see potential migration targets for the selected VM.	All	No	Yes

Tab Name	Description	Parent Object	ESX	VC
Datstores	Only displayed when the Datstores inventory object is selected. Displays information about all configured VMFS and NFS datstores (no RDMS), including device path, total capacity, and free space.	Datacenter	No	Yes

Finally, the bottom pane is where recent task information and alarms are displayed for a short period of time, as shown in Figure 4.4. Switching between task and alarm information is accomplished by clicking either the Tasks or Alarms link in the lower-left corner. You can resize any of the three panes by selecting and dragging the pane border. You can close the bottom pane by clicking the X in the right corner of the pane. After you close it, you can display it again by clicking either the Tasks or Alarms link.



Object	Status	Name	Triggered
esxi1.xyz.com	●○○	Host down	2/11/2009 4:53:06 PM
Win2003-1	●○○	VM Down	2/11/2009 4:56:43 PM
green	●○○	VM Down	2/11/2009 4:56:43 PM
orange	●○○	VM Down	2/11/2009 4:56:43 PM

Figure 4.4 Bottom pane with task and alarm information

Connecting to an ESX Host Directly Versus Connecting to a vCenter Server

When you connect to an ESX host directly with the VI Client rather than through vCenter Server, some features and functionality are not present. Let's first go over some of the things that are missing when you connect to an ESX host directly:

- **VMotion.** You cannot VMotion (hot migrate) a VM from one ESX to another.
- **Cold migration.** You cannot cold migrate a VM from one ESX host to another.
- **Users and roles.** You cannot view or manage custom roles and can only manage ESX local user accounts (not domain accounts).
- **Performance statistics.** You can only view real time and the past 60 minutes of ESX host and VM performance statistics.
- **Cluster settings.** You will only see clusters when connecting to a vCenter Server.

- **Scheduled tasks.** You cannot view or manage any scheduled tasks for the ESX host.
- **Templates.** You cannot view or use templates that exist on the ESX host.

In addition, if you connect to an ESX or ESXi host directly with the VI Client, you will see a tab in the right pane that you will not see when connecting to a vCenter Server. This tab is the Users and Groups tab which contains a list of all the local users and groups on the ESX host's Linux-based Service Console or the ESXi Busybox shell.

The VI Client does not show this tab when connected to a vCenter Server because it uses Active Directory for all of its user and group management. The Users and Group tab is used when managing stand-alone ESX and ESXi hosts and you want to configure user accounts to log in to the host with instead of using the root account. When you add an ESX or ESXi host to vCenter Server, it automatically adds a local user to the host server called vpxuser, which it uses to connect to the host server and manage it. We talk more about managing users and groups later.

You can make configuration changes using the VI Client directly on the ESX host instead of using vCenter Server, but it is recommended that you use vCenter Server if possible. If you do make changes directly to an ESX host, the vCenter Server agent that runs on the ESX host will ensure that the changes are also updated in vCenter Server. However, this sometimes may not happen right away, and in rare cases restarting the VC agent is necessary so that vCenter Server picks up on the change. For information about how to restart the VC agent, see the “Troubleshooting vCenter Server” section in Chapter 10, “Troubleshooting Your Virtual Environment.” Also, if you delete a VM directly on an ESX host that is being managed by vCenter Server, the VM will show up as orphaned in vCenter Server and will have to manually be removed from it. For the purposes of this tour, we will be connected to a vCenter Server. Not all features shown here will be available if you connect to an ESX host instead. In addition, we will be connecting to a vCenter Server 2.5 Update 2 server managing ESX 3.5 Update 2 hosts.

Additional VI Client Features

The Scheduled Tasks feature comes with vCenter Server and enables you to schedule certain tasks to take place at a set date and time. The types of tasks that can be scheduled include making a snapshot of a VM, changing the power state of a VM, and moving a VM with VMotion. Tasks can be scheduled to run once at a set date or time or at a repeat interval. Scheduled tasks are available only on ESX hosts managed by vCenter Server and are not a feature when connecting to an ESX host directly with the VI Client.

vCenter Server 2.5 introduced a new feature to the VI Client called plug-ins, which are add-on components that extend the power of the VI Client. Two default plug-ins are included with vCenter Server: the Converter Enterprise plug-in and the Update Manager plug-in. In

addition, some third-party plug-ins can be installed, including a graphical interface for Storage VMotion from Andrew Kutz. Plug-ins can be managed by selecting the top menu Plug-ins option.

Did You Know?

As of VMware vCenter Server 2.5 Update 2, the VI Client can use Windows Single Sign-on (SSO) to log on to a vCenter Server. The Windows Single Sign-on feature uses the currently logged-on user credentials to log on to the vCenter Server. To set it up, just follow these steps:

1. On the PC that you want to set up SSO for vCenter Server, create a new shortcut on your desktop.
2. In the Create Shortcut Wizard, click Browse and navigate to the location of the VpxClient.exe program and click OK. (By default, it is located in C:\Program Files\VMware\Infrastructure\Virtual Infrastructure Client\Launcher\.)
3. After the full path is in the Location field, append `-passthroughAuth -s vCenter Server hostname` to the end of the line, where *vCenter Server hostname* is the host-name or IP address of the vCenter Server instance you want to connect to.
4. Click Next and name the shortcut. Then click Finish. When you double-click the newly created shortcut, you will be logged on to vCenter Server using your currently logged-on Windows credentials. For more information about this, see VMware knowledge base article 1006611.

Configuring Licensing

In Chapter 2, “Planning Your Virtual Environment,” we discussed the two licensing options (centralized versus stand-alone or host based) for licensing your ESX and ESXi hosts and why you would choose one over the other. In Chapter 3, “Building Your Virtual Environment,” we discussed installing the license server as part of the vCenter Server installation. In this chapter, we cover the steps to configure it for use with your ESX hosts and how to configure stand-alone licensing.

Getting Your License Files

When you purchase your VMware product, you are entitled to download the license files to install to license your software. You will typically receive an email from whoever you purchased it from (HP, IBM, Dell, and so on) containing information about how to obtain

your licenses. After you follow the instructions and register your product, you can then download the license file. A license file for VI3 is just a text file in UNIX format (it uses the UNIX newline character LF rather than the Windows one CR-LF) that contains license types and authorization keys for ESX hosts, vCenter Servers, and the various features that you can license, such as VMotion. The formats of the stand-alone license files are different from those of the centralized license files, and they cannot be intermixed. The contents of a license file look similar to what is shown here. Centralized license files must always contain the first three lines in the header, whereas stand-alone files do not. The license file should always have a .lic extension.

Note

When a line of code is too long to fit on one printed line, we've broken it and added a code-continuation arrow ➤ in front of the continuation.

```
SERVER this_host ANY 27000
VENDOR VMWARELM port=27010
USE_SERVER

vCenter Server Management Server Required for operation of Virtual Center 2.
INCREMENT PROD_VC VMWARELM 2005.05 permanent 1 \
__VENDOR_STRING=licenseType=Server ISSUED=06-Jul-2006 \
__NOTICE=FulfillmentId=1204 SIGN="0C31 E631 5A09 0BA6 A7C7 B722 \
__0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 \
__0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000"
vCenter Server Host Agent One license per CPU of ESX Server hosts
➤managed by Virtual Center.
INCREMENT VC_ESXHOST VMWARELM 2005.05 permanent 32 \
__VENDOR_STRING=licenseType=Server;capacityType=cpuPackage \
__ISSUED=06-Jul-2006 NOTICE=FulfillmentId=1203 SIGN="15DB C49D \
__ 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 \
__0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 \
__7804 8880 3043 A2FC"
VMotion One license per CPU of ESX Server hosts using VMotion.
INCREMENT VC_VMOTION VMWARELM 2005.05 permanent 32 \
__VENDOR_STRING=licenseType=Server;capacityType=cpuPackage \
```

```
___ISSUED=06-Jul-2006 NOTICE=FulfillmentId=1203 SIGN="1234 EECE \  
___0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 \  
___0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 \  
___A5A0 2739 7061 4BEB"
```

VMware HA One license per ESX Server CPU.

```
INCREMENT VC_DAS VMWARELM 2005.05 permanent 32 \  
___VENDOR_STRING=licenseType=Server;capacityType=cpuPackage \  
___ISSUED=06-Jul-2006 NOTICE=FulfillmentId=1203 SIGN="0137 B2E4 \  
___0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 \  
___0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 \  
___BB34 B457 72D1 2BD2"
```

ESX Server One license per CPU.

```
INCREMENT PROD_ESX_FULL VMWARELM 2005.05 permanent 32 \  
___VENDOR_STRING=licenseType=Server;capacityType=cpuPackage;gp=14;  
➔exclude=BACKUP \  
___ISSUED=06-Jul-2006 NOTICE=FulfillmentId=1203 SIGN="0778 0B1A \  
___0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 \  
___0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 \  
EF5E 2A57 C6EE 641E"
```

When you go to VMware's license management portal web page (www.vmware.com/download/licensing.html), you can log on to manage and view your licenses. You will see a list of all the available licenses associated with your VMware store account and the licensing files available for download. When you first go to the licensing web page and log on, you will see all the products listed that you have purchased licenses for and how many are currently activated and how many are available to activate. There are buttons on the licensing web page that perform various functions, including creating a new license file, editing the numbers of an existing license file, combining license files, and deleting license files.

Combining license files is useful if you are already using an existing license file and you purchase additional products with new licenses. Normally, you would have to copy and paste the information from the new license file into the existing one, but the combine feature takes care of that for you. If you do manually edit the license file, make sure you do not use Notepad for this, because it does not support UNIX newlines, which the license file uses. You should instead use WordPad or a third-party editor that supports UNIX newlines. When you create a new license file, you have the option of choosing either stand-alone or centralized. Which one you choose is not permanent, and you can switch between stand-alone and centralized or vice versa at any time by going back to the licensing page.

Watch Out!

ESX hosts using stand-alone mode cannot be managed by vCenter Server, which requires the ESX host to be using centralized licensing. Also, if you plan to use features that require vCenter Server (such as VMotion, HA, and DRS), you must also use centralized licensing.

VMware provides a free license checker utility that validates your license file and can reformat or repair it if necessary. It also displays the license information in a more readable form and tells you the number of licenses that you have for each feature. To use this utility, simply go to their web page (www.vmware.com/checklicense/), paste in all the text from your license file, and click the Validate button. It will parse your license file and display some useful information about the contents and any warnings on possible problems.

Configuring Stand-Alone/Host-Based Licensing

As discussed in Chapter 2, stand-alone licensing is used when you do not have a vCenter Server in your environment or for ESX/ESXi hosts that are not being managed by a vCenter Server. This type of licensing relies on uploading a license text file to each host. To configure stand-alone licensing on an ESX host, follow these steps:

1. Install ESX Server software on the host server.
2. Power on the host server. It will initially run in evaluation mode for up to 60 days.
3. Connect to your ESX host using the VI Client. Once connected, select the ESX host, and then on the Configuration tab select Licensed Features, and the current evaluation license will be displayed.
4. Click the Edit link next to License Source, and a window will display, as shown in Figure 4.5. Change the licensing source from Use Evaluation Mode to Use Host License File. Then browse to your license file location (for example, C:\vmware\VMware.lic) and select it. It will automatically be uploaded to the host server in the /etc/vmware directory.
5. When you click OK, the ESX host will be switched over to use the newly uploaded license file, and the new licensing information will display, as shown in Figure 4.6.

Configuring Centralized Licensing

When using the centralized licensing server, make sure the VMware License Server service is running and set to start automatically. This service is the actual licensing server component and is what vCenter Server communicates with to gather information about the installed license files. Then make sure that the default ports (27000 and 27010) that vCenter Server uses open between your vCenter Server and ESX.




Figure 4.5 Choosing a license type in the Licensing Sources window




Figure 4.6 Licensing information for the ESX host is displayed after you upload a valid license file.

By the Way

Before VirtualCenter¹ 2.0.1 Patch 2, server-based licenses for each VI3 component that requires a license had to be concatenated in a single file (vmware.lic) for use with the VMware license server. Beginning with VirtualCenter¹ 2.0.1 Patch 2, the license server obtains license information from a license directory (rather than the single vmware.lic file), so new licenses can be added to the directory. This directory is specified in the License Server Tools application that runs on the vCenter Server.

To configure centralized licensing using the licensing server packaged with vCenter Server, follow these steps:

1. Start vCenter Server, which will initially run in evaluation mode for up to 60 days.
2. Copy your license file to the directory of the license server. This is typically the C:\Program Files\VMware\VMware License Server\Licenses directory. After you copy the file, it is best to restart the VMware License Server Windows service to ensure it reads the file.
3. Connect to the vCenter Server using the VI Client. Once connected, select Administration from the top menu, and then vCenter Server Management Server Configuration, which will bring up a new window, as shown in Figure 4.7




Figure 4.7 License Server settings for vCenter Server

4. Uncheck the Evaluate VirtualCenter¹ Server option, and then select either the Use License Services on this VirtualCenter¹ Server option or the Use the Following Licensing Server option. If you choose the latter option, enter the hostname or IP address of the licensing server that you want to use.
5. Then click the Administration button, and select the Licenses tab, which will display all the licensed features and quantities available with the license files being used, as shown in Figure 4.8.

License Feature	Cost	Remaining	Total
ESX Server Standard	1/CPU Package	10	12
VMware Consolidated Backup	1/CPU Package	12	14
VirtualCenter Agent for ESX Server	1/CPU Package	10	12
VMotion	1/CPU Package	12	12
VMware HA	1/CPU Package	14	14
VMware DRS	1/CPU Package	12	12
VirtualCenter Management Server	1/Server	0	1

Figure 4.8 Licensed features are displayed once you select a license server.

As you add ESX hosts to your vCenter Server, the number of remaining licenses for each feature will be reduced until they are all used up. For example, adding a two-CPU ESX host to vCenter Server will result in a reduction of two of the ESX Server Standard one-CPU licenses and two of each additional feature that the server will use (for example, VMotion, HA, DRS).

Did You Know?

If you have purchased ESX Enterprise licenses, which include all the extra features, it will still show up listed as only ESX Server Standard. This is normal behavior, with the difference being that all the extra features will show up separately in the feature list.

Configuring Free ESXi Licensing

The free version of ESXi (starting with ESXi 3.5 Update 2 or later) is licensed differently than ESX and the other versions of ESXi. It uses a 25-character license key or serial number composed of numbers and letters in the format #####-#####-#####-#####-##### rather than a license file or a license server. When you first install ESXi, it runs using a 60-day evaluation license. When you register to download ESXi, you are sent an email that contains a link to activate your free license. When you log on and activate it, you are given the license key, which you can enter using the VI Client. To enter the license key, follow these steps:

1. Connect to the ESXi server using the VI Client.
2. Click the Configuration tab, and then select Licensed Features in the Software area.

3. Click the Edit button next to License Source.
4. Select the Use Serial Number option, as shown in Figure 4.9, and then enter the license key that was given to you. Then click OK.




Figure 4.9 Entering a serial number to license the free version of ESXi

5. Click OK if you get a warning about certain features not being available after disabling the evaluation license. (The evaluation license comes with some features, such as Consolidated Backup, that are not available with the ESXi free license.) Also, click Yes if you receive a confirmation message about releasing the current license.
6. Once completed, your license source will be changed from 60-day Evaluation to Serial Number and will display your serial number.

If you have purchased Foundation, Standard, and Enterprise licenses and you want to use ESXi rather than ESX, you license ESXi the same way as you would ESX. The license file for ESXi is no different than the ESX license file, and you can use either with your license. To change an ESXi server from using either the evaluation or free license, you just edit the license source and select either the Use License Server or Use Host License File option.

Additional Considerations

This section covers some considerations that you should be aware of when setting up and configuring licensing in your environment.

License Server Ports

If you use a license server, by default it uses ports 27000 (TCP) for inbound requests and 27010 (TCP) for outbound requests to communicate with ESX/ESXi hosts and the vCenter Server. Therefore, for licensing to function properly, you need to ensure these ports are not being blocked by a firewall. The ports need to be open between each ESX/ESXi host that uses centralized licensing and the license server. In addition, if you are not running the license server on the vCenter Server, the ports also need to be open between the license server and the vCenter Server.

Single CPU Licensing

ESX/ESXi is licensed in two-CPU increments on host servers, and until recently it was against the VMware EULA to split the license and license a single CPU server. In April 2008, VMware changed their policy and allowed customers to license a single-CPU host server. It is still sold in only two-CPU increments, but this change in their policy is beneficial for customers and provides greater flexibility when deploying ESX/ESXi hosts. This can be particularly advantageous for smaller customers who can buy two single-CPU multicore servers, split an ESX/ESXi license between them, and take advantage of features like HA and VMotion. The details of VMware's single CPU policy are as follows:

- To install licenses in single-processor increments, a customer must use the centralized license server model for installing and managing VMware license files.
- Includes servers with two sockets that are populated with a single processor.
- Each processor may contain up to four cores.
- Licenses of VMware ESX and VI are still sold in minimum increments of two processors.

Managing the License Server

The license server installs as part of the vCenter Server installation, but as a separate application it can optionally be installed separately by running the vCenter Server installation executable. Once the application is installed, a separate management application is also installed to directly manage the license server. This application (LMTools) is typically located in the Start menu under the VMware folder in a subfolder called VMware License Server and is called VMware License Server Tools. In most cases, you will not have to use LMTools, but it is available if you need it for things like troubleshooting, changing the license server configuration, and rereading the license file. When you run LMTools, you will see a tabbed window, as shown in Figure 4.10.




Figure 4.10 Licensing Server Tools application (LMTools)

Let's cover a few of the more commonly used tabs in the LMTools application:

- On the Start/Stop/Reread tab, there are buttons to either start or stop the license server or force it to reread the license file if changes have been made to it. Optionally, these tasks can also be accomplished by stopping and starting the VMware License Server Windows service that is installed when the license server is installed. It is a good idea to also restart the vCenter Server Windows service after you restart the license server or reread the file so that it picks up the changes.
- On the Server Status tab, you can click the Perform Status Enquiry button to show the current status of the licensing server, including which licenses are in use.
- On the Config Services tab, you can specify the path to the license file. This is where you place the .lic license files that VMware uses for centralized licensing. With VirtualCenter¹ 2.0.x, you specified the path and name of the license file to be used. Starting with vCenter Server 2.5, you specify only the path, and any .lic files in the directory specified are automatically read and used.

Configuring vCenter Server

Before you add ESX and ESXi hosts to your vCenter Server, you need to configure the containers to put them inside. This begins with creating a datacenter object and then folder and cluster objects below it to put your hosts in. From there, you configure some of the advanced features that vCenter Server provides (for example, HA, DRS, and VMotion), and finally you can divide your host resources up into resource pools that you can put your VMs inside.

Creating a Datacenter

When you start vCenter Server for the first time, it will be empty and not have any inventory objects configured. Your first step should be to configure a top-level datacenter object so that you can add ESX hosts and clusters beneath it. A datacenter object is a top-level object that contains hosts and clusters. Within a datacenter, no two networks and datastores may have the same name. Before you configure any objects in vCenter Server, take a moment and consider how you want to lay out your environment inside vCenter Server. You can organize your ESX hosts and VMs using datacenters, clusters, and resource pools. You may consider creating multiple datacenter objects based on the geographical locations or functional purpose of your ESX hosts.

By the Way

Before you start naming objects in vCenter Server, it is a good idea to give some thought to naming standards so that your objects have descriptive and useful names that can easily be understood by anyone who will be using vCenter Server.

To add a datacenter object, simply log on to vCenter Server and right-click the top-level Hosts and Clusters folder and select the New Datacenter option and provide a name for it. If you plan to use multiple datacenters, consider giving it a descriptive name that has meaning.

Now that you have a datacenter created, you're ready to start adding ESX hosts to vCenter Server. If you plan to use clusters, it's a good idea to create them first before you start adding your hosts to vCenter Server.

Creating Folders

You can also create folders in vCenter Server to help you organize your virtual environment however you like. Two different types of folders can be created: yellow folders that contain datacenters/clusters/hosts, as shown in Figure 4.11, and blue folders that contain VMs/templates, as shown in Figure 4.12.

By the Way

Because this book is not in color, you will not see the blue and yellow folders that are mentioned here in the figures shown in this book, but you will see them on your own computer when you are using vCenter Server.




Figure 4.11 Yellow datacenter/cluster/host folders in vCenter Server




Figure 4.12 Blue VM/template folders in vCenter Server

The yellow folders are displayed in all the different vCenter Server views, but the blue folders are displayed only in the Virtual Machines and Templates view. Folders can also be nested to create an organizational hierarchy inside vCenter Server. For example, you might create yellow folders to group your production host together or create blue folders to group all your Linux or Microsoft Exchange VMs together.

A folder can contain other folders, datacenters, clusters, VMs, templates, or hosts, but any given folder hierarchy can contain only one type of object. For example, you can place hosts and folders containing hosts into the same folder, but you cannot put a folder containing hosts and a folder containing VMs into the same folder. Also, VM names must be unique within a given VM folder.

Folders make for a great way for assigning permissions inside of vCenter Server for users to manage hosts and VMs. Assigning permissions at the folder level on groups of hosts or VMs is much easier than assigning them individually. Once you create folders, you can just drag and drop any object in vCenter Server to them to start organizing your environment. We cover permissions in detail later in Chapter 5, “Securing Your Virtual Environment.”

Watch Out!

Before deleting folders, it is best to remove all objects from inside them. If you delete a yellow folder that contains host objects, you will have to re-add the hosts to vCenter Server. If you delete a yellow folder with clusters, you will have to re-create the cluster and add your hosts back into it. If you delete a blue folder containing VMs, it will completely delete the VMs from your hosts.

Creating Clusters

After you have a datacenter created, you can optionally create clusters. You don't have to create a cluster, but they are necessary to use some of the advanced features (such as HA and DRS). A cluster is a collection of ESX hosts and their associated VMs that share resources and a common heartbeat that is used to detect ESX host failures so that an appropriate action can be taken. Once you add an ESX host to a cluster, its individual resources become part of the cluster's total resources. You also have the option to enable the HA and DRS feature when you configure a cluster. You can enable only one or both of these features. For example, if you want to use only HA, you can enable it and disable the DRS feature. You can always create a cluster at a later time and then add any ESX hosts that you already have in vCenter Server to it.

Before you create a cluster, take into account the hardware of your ESX host servers. If you plan to use a mix of Intel or AMD processors, you will probably want to create separate clusters for the Intel and AMD servers because of the incompatibility of VMotion between the two-processor brands. If you mix Intel and AMD hosts in the same cluster, you cannot hot migrate (VMotion) servers between Intel and AMD hosts. Furthermore, you limit the effectiveness of DRS because it cannot move VMs from AMD hosts to Intel hosts. In addition, consider the size of your hosts in the cluster. For example, if you have a cluster with four smaller two-CPU, 8GB hosts and two larger four-CPU, 32GB hosts, if one of the larger hosts fails the smaller hosts might not be able to accept the additional VM load from the failed larger host. As an alternative, you might consider creating separate clusters, one for your smaller hosts and another for your larger hosts.

To create a cluster, right-click a datacenter that you have created and select the Create Cluster option. A window will display where you can enter a name for your cluster and optionally enable either HA or DRS or both features. If you choose to enable these features, additional options appear for you to configure each feature. We cover the HA and DRS configurations in a bit.

After you give the cluster a name and click Next, you have the option to either keep the VM's swapfile in the same directory as the VM or store to store the swapfile in the datastore specified by the host server that the VM is located on. A VM swapfile (VSWP) file is the file that is created when a VM is powered on, which it will use if a host server's physical memory is used up due to overcommitment. More detail on this file is provided in Chapter 11, "Advanced Topics." It is recommended to store the VSWP file with the VM because if you store it elsewhere it may cause issues with VMotion if a destination host server cannot access the datastore that the VSWP file is located on. When you select a VSWP file location and click Next and then click the Finish button, your cluster will be set up and ready for use.

After you have created a cluster, you can view a summary of the cluster by selecting the cluster in the left pane of the VI Client and then selecting the Summary tab in the right pane.

The Summary tab displays general information about the cluster, including the total number of hosts and VMs, total CPU and memory resources, current HA and DRS configurations, and DRS resource distribution statistics. If you select the DRS Recommendations tab,

you will see a list of recommendations for DRS and an action history for DRS-related events. Also, if you select the Resource Allocation tab, you will see the resource settings (shares/limits/reservations) for all your VMs in the cluster. You can toggle between CPU and memory resources by clicking the two buttons. In addition, you can select the other tabs in the right pane to see VM, host, performance, alarms, task, and event information for the VMs and hosts in the cluster.

Adding ESX and ESXi Hosts to vCenter Server

After you have your datacenter and clusters set up in vCenter Server, you can start adding ESX and ESXi hosts so that you can begin managing them with vCenter Server. When you add a host to a cluster, the following things happen:

- All the resources for the host become available to the cluster for use in the cluster's root resource pool. Resource pools are discussed in detail later in this chapter.
- All VMs on the hosts that are added to the cluster have a default restart priority for HA of medium; this can be changed, as detailed in the next section.
- Unless the cluster is also enabled for VMware DRS, all resource pools are collapsed into the cluster's top-level (invisible) resource pool.
- Any capacity on the host beyond what is required or guaranteed for each running VM becomes available as spare capacity in the cluster pool. This spare capacity can be used for starting VMs from other hosts in case of a host failure.
- If you add a host with several running VMs, and the cluster no longer fulfills its failover requirements because of that addition, a warning appears, and the cluster's status is changed to invalid (red). This applies to HA-enabled clusters only.

Did You Know?

An ESX and ESXi host can only be managed by a single vCenter Server Management server. If you try to add a host that is already being managed by a VC server to another VC server, you will receive a warning message. If you continue despite the warning, the host will no longer be managed by the previous VC server and will instead be managed by the new VC server.

To add an ESX or ESXi host to vCenter Server, follow these steps:

1. Right-click the cluster, datacenter, or folder where you want the host to be and select the Add Host option. A window will appear asking you for the hostname or IP address and the username and password to log on to the host with (typically root). As mentioned in Chapter 3, this logon is used once for the initial logon and then a

separate vpxuser account is created to be used by vCenter Server for subsequent logons to the host. After you enter the host and username information, click Next to continue.

2. vCenter Server will connect to the host and display a host summary, which includes the server model and version, ESX version, and a list of any VMs that already exist on the host. Click Next to continue.
3. If you are adding a host to a cluster and it has DRS enabled, a window will let you choose a resource pool for the host's resources to be added to. You have the option to put the host's resources in the cluster's root resource pool (the default top-level invisible resource pool) or to create a new child resource pool. Choose an option and click Next to continue.
4. At the Ready to Complete window, a summary is displayed. Click Finish to begin the process of adding the host to the cluster. The process will begin, and VC will first install a management agent on the ESX host and then add the host to the cluster. Once the host has been added, you will see it displayed in vCenter Server.

Configuring High Availability

You can set up the High Availability (HA) feature either during your initial cluster setup or afterward. To configure it, just select the cluster that you want to enable HA on and right-click it and edit the settings for it. Put a check mark next to the Enable VMware HA field on the General page and it will be enabled for the cluster. You can optionally configure some additional settings to change the way HA functions. To access the settings, click the VMware HA item in the cluster settings window, as shown in Figure 4.13.

The first part of the HA settings is for Admission Control, which controls how many host failures a cluster can tolerate and whether to allow VMs to be powered on if the cluster violates the availability constraints.

This is used to ensure that there is sufficient capacity among the remaining host servers to be able to handle the additional load from the VMs on failed host servers. Setting the number of host failures allowed will cause the cluster to continuously monitor that sufficient resources are available to power on additional VMs on other hosts in case of a failure. Specifically, only CPU and memory resources are factored in when determining resource availability; the disk and network resources are not. You should set the number of host failures allowed based on the total number of hosts in your cluster, the size of the hosts in your cluster, and how busy your hosts in the cluster are. vCenter Server supports up to a maximum of four host failures per cluster.

For example, if you have four ESX hosts in your cluster, you probably want to allow for only one host failure. If you have eight ESX hosts in your cluster, you might want to allow for two host failures. And if you have a larger cluster with 20 ESX hosts, you might want to allow for up to 4 host failures. This number will vary based on how busy your hosts are and the size

of your hosts. If you have a high VM to host ratio and have high CPU and memory usage on your hosts, you will probably need to allow for a smaller number of host failures, because there will not be the necessary spare capacity on your host servers. Also, if you have both larger and smaller hosts in your cluster, you need to ensure that the remaining hosts in your cluster can tolerate the additional load (in case some of the larger hosts fail).




Figure 4.13 HA settings

Also, in the Admission Control settings, you can specify whether to allow VMs to be powered on if they violate the availability constraints. If you configure HA to allow this, your VMs will be started on other hosts even if they do not have the resources available to accept the additional VM load.

One other important thing to note is that when a host fails all of its VMs will be restarted on a single ESX host that has the lightest workload and can quickly overload the host. When this occurs, DRS then kicks in to spread the load across the remaining hosts in the cluster. If you plan to use HA without DRS, ensure that you have plenty of extra capacity on your ESX hosts to handle the load from any one failed host. In addition, you can set restart priorities so that you can specify which VMs get restarted first and even prevent some VMs from being restarted in case of a failure.

The next group of HA settings are for the cluster settings that will apply to all VMs in this cluster by default.

This means each VM created in or moved to this cluster will default to whatever this setting is set to unless it is later individually changed. The first part is the VM restart priority, which is the priority given to a VM when it is restarted on another host in case of a host failure. This setting can be set to High, Medium, Low, or Disabled. Any VM set to Disabled will not be restarted in case of a host failure. The second setting is the Host Isolation Response, which is used to determine what action the failed host that is isolated should take with the VMs that are running on it. This is used in the case that a host is still running but has a failure in a particular subsystem, (for example, NIC or host bus adapter [HBA] failure) or a connectivity problem (for example, cable or switch) and is not completely down. When a host declares itself isolated and the VMs are restarted on other hosts, this setting dictates what happens on the failed host. The options include to leave the VM powered on, power off the VM (hard shutdown), and shut down the VM (graceful shutdown). If you choose the shut down the VM option, HA will wait five minutes for the VM to shut down gracefully before it forcefully shuts it down. The time period that it waits for a graceful shutdown can be modified, as outlined in the advanced configuration settings later. It is usually desirable to have the VM on the failed host powered off or shut down so that it releases its lock on its disk file and does not cause any conflicts with the new host that powers on the VM. One reason you might choose to instead leave the VM powered on is if you do not have network redundancy or do not have a reliable network. In this case, you might experience false triggers to HA, where the ESX host is okay but has just lost network connectivity. If you have proper network redundancy on your ESX hosts, HA events should be very rare. This setting will not come into play if the failed host has a disastrous event like completely losing power, because all the VMs will be immediately powered off anyway. In two-host configurations, you almost always want to set this to leave powered on.

The last group of HA settings are for Virtual Machine Monitoring, a new feature that was introduced in ESX 3.5 that enhances HA by providing support for the failure of individual VMs due to failures inside the operating system of the VM (for example, Windows blue screen of death).

Prior to this feature being introduced, HA would handle ESX host failures only by restarting VMs on other ESX hosts, in case of a problem with the host server. The new feature extends HA to also monitor individual VMs via a heartbeat that is sent to the VM every second through VMware Tools. To turn on this feature in ESX 3.5 Update 2 or later, simply check the Enable Virtual Machine Monitoring field and then select a sensitivity using the slider from Low to High, which uses preset settings. If you want to customize your settings, you can click the Advanced Options button and change the settings. VMs that fail as a result of Virtual Machine Monitoring are always restarted on the same host.

Before ESX 3.5 Update 2, the only way to configure this feature was using the following HA Advanced Option settings:

- **das.vmFailoverEnabled. True.** True or false. Enabled through the GUI on ESX 3.5 Update 2 and later.
- **das.FailureInterval. 30.** Declare VM failure if no heartbeat is received for the specified number of seconds.
- **das.minUptime. 120.** After a VM has been powered on, its heartbeats are allowed to stabilize for the specified number of seconds. This time should include the guest operating system boot time.
- **das.maxFailures. 3.** Maximum number of failures and automated resets allowed for the time that das.maxFailureWindow specifies. If das.maxFailureWindow is -1 (no window), das.maxFailures represents the absolute number of failures after which the automated response is stopped and further investigation is necessary.
- **das.maxFailureWindow. 3600.** Either -1 or a value in seconds. If das.maxFailures is set to a number, and that many automated resets have occurred within that specified failure window, automated restarts stop and further investigation is necessary.

The final HA section, Virtual Machine Options, is for setting individual VM options for restart priority and host-isolation response, as shown in Figure 4.14.




Figure 4.14 HA settings: Setting Virtual Machine Options

As mentioned earlier, the default for each VM is to use whatever the default cluster settings are set to. In this section, you change the options for individual VMs. This is useful if you have VMs on an ESX host that you do not want to restart in case of an HA event. For example, if you have dev, test, and prod VMs mixed on an ESX host and you want only the more important prod VMs to be restarted, you set the restart priority for the dev and test VMs to disabled so that they will not be restarted on other hosts in case of an HA event.

Split-Brain Condition

Split brain is the name for the condition where an ESX host loses its network connection to its default gateway on the Service Console network but the ESX host still continues to run along with all of its VMs. As a result, one half of the brain (the ESX host that is isolated) thinks it is isolated, and the other half of the brain (the remaining ESX hosts) thinks the isolated ESX host has failed and attempts to power on the VMs that are on that host. If a VM is configured to stay powered on, then when the host is isolated the remaining ESX hosts will try to power on the VM, which could cause issues because of the locks on the VM disks that the isolated ESX host will still have. To help minimize the occurrence of this situation, it is best to configure maximum redundancy on your ESX hosts using any of the following methods to protect against situations when network connections are lost but the host continues to run:

- Use multiple physical NICs (pNICs) in your Service Console vSwitch to provide redundancy and protect against a single NIC or network connection failure. In addition, you should connect each NIC to a separate physical switch to protect against a physical switch failure.
- Configure a secondary Service Console on a separate vSwitch. The primary Service Console network is still used for network and management purposes. But with a secondary Service Console network, HA will send heartbeats over both the primary and secondary Service Console networks. If one of the Service Console networks fails, HA can still send and receive heartbeats over the other Service Console network.
- Use some of the advanced configuration parameters (as discussed next).
- Use some of the HA advanced parameters, including increasing the `das.failedetectiontime` setting to protect against brief network outages and using the `das.isolationaddress` setting to provide an additional check to see whether the host is isolated.

Advanced Configuration

Many advanced configuration options can be set to tweak how HA functions. These options can be set through the Advanced Options button in the HA settings, but you have to know the setting name and its values to be able to set it. These options are not displayed by default except for the ones that are set if you enable Virtual Machine Monitoring, and you must manually add them if you want to use them. Listed here are some of the more common

HA settings that will work as of vCenter Server 2.5 Update 3 and ESX 3.5 Update 2. Not all settings will work in earlier versions of ESX.

- **das.failedetectiontime.** This setting was introduced in VirtualCenter¹ 2.0.2 to allow the default interval for HA detection to be changed. Previously, this was hard-coded so that if a host did not respond in 15 seconds, it would be considered failed. This setting is in milliseconds and can be changed from the default of 15,000 milliseconds. For example, if you want to increase this to 60 seconds, you set it to 60,000 milliseconds. (Don't use commas in the number.)
- **das.poweroffonisolation.** This setting is the default cluster setting for VM isolation response that is set through the VI Client. The default is true, which powers off VMs in case of an HA event. Setting this to false leaves the VM still running on the isolated host when an HA event occurs. Note: As of vCenter Server 2.5 Update 3, VMware changed this so the default is false and so the VMs are left powered on.
- **das.isolationaddress.** This setting was introduced in VirtualCenter¹ 2.0.2 and is the IP address that HA will ping to determine whether a host is isolated from the network. If this option is not specified, the default gateway of the console network is used. This default gateway has to be some reliable address that is known to be available, so that the host can determine whether it is isolated from the network. Using multiple isolation-response addresses gives HA a potentially more accurate picture of the network connectivity of a host. There may be situations in which a single isolation address indicates that a host is in a state of complete isolation from the network, but access to additional isolation addresses would show that only a partial network failure has occurred. You can have a total of ten isolation addresses set using the format `das.isolationaddressX`, where *X* is a number from one to ten. If you use this setting, you should change the default failure detection time (`das.failedetectiontime`) to 20 seconds or greater. In general, the more isolation-response addresses configured, the longer you should make the timeout, to ensure that proper failure detection can occur.
- **das.usedefaultisolationaddress.** By default, HA uses the default gateway of the console network as an isolation address. This attribute specifies whether that should be used or not; the values are true or false. If you are using the `das.isolationaddress` setting, this should be set to false.
- **das.isolationShutdownTimeout.** The amount of time that HA waits to gracefully shut down a VM if the Shutdown VM option is selected for isolation response. The default is 300 seconds.
- **das.defaultfailoverhost.** If this is set, HA first tries to fail over hosts to the host specified by this option. This is useful if you want to use one host as a spare failover host, but is not usually recommended, because HA tries to use all available spare capacity among all hosts in the cluster. If the specified host does not have enough spare capac-

ity, VMware HA tries to fail over the VM to any other host in the cluster that has enough capacity.

- **das.failedetectioninterval.** This setting was introduced in vCenter Server 2.5 Update 2 and is the interval used for heartbeat detection among ESX hosts. The default is that a host will check for a heartbeat every second (1,000 milliseconds). You might want to increase this if your hosts are on remote or high-latency networks.
- **das.allowVmotionNetworks.** This setting was introduced in vCenter Server 2.5 Update 2 specifically for ESXi, which does not use a vswif Service Console network like ESX hosts and instead uses a special management network and will allow a NIC that is used for VMotion networks to be used for HA also. This is to allow a host that has only a single NIC configured for the both the Service Console and VMotion combined to be used for HA. By default, VMotion networks are ignored.
- **das.allowNetwork.** This setting was introduced in vCenter Server 2.5 Update 2 and allows the use of port group names to control the networks used for HA. Starting with vCenter Server 2.5 Update 2, HA has an enhanced network compliance check to increase cluster reliability. This enhanced network compliance check helps to ensure correct clusterwide heartbeat network paths. This also helps prevent delayed failure detection and “split-brain” conditions in certain scenarios. When configured, the HA cluster uses only the specified networks for HA communication. You can set the value to be Service Console 2 or Management Network, to use the networks associated with those port group names in the networking configuration. You can set this using the following format `das.allowNetworkX`, where *X* is a number, starting with zero (for example, `das.allowNetwork0` = Service Console or `das.allowNetwork1` = Service Console 2). You can find more information about this setting in the following VMware knowledge base articles: <http://kb.vmware.com/kb/1006606> and <http://kb.vmware.com/kb/1006541>.
- **das.vmMemoryMinMB.** This setting specifies the minimum amount of memory (in megabytes) sufficient for any VM in the cluster to be usable. This value is used only if the memory reservation is not specified for the VM and is used for HA admission control and calculating the current failover level. If no value is specified, the default is 256MB. Reducing this can help prevent the warning about insufficient resources, and the red exclamation mark indicates that a cluster does not have sufficient failover capacity.
- **das.vmCpuMinMHz.** This setting specifies the minimum amount of CPU (in megahertz) sufficient for any VM in the cluster to be usable. This value is used only if the CPU reservation is not specified for the VM, and is used for HA admission control and calculating the current failover level. If no value is specified, the default is 256MHz. Reducing this can help prevent the warning about insufficient resources, and the red exclamation mark indicates that a cluster does not have sufficient failover capacity.

- **das.bypassNetCompatCheck.** This setting was introduced in vCenter Server Update 3 to provide the ability to disable the HA enhanced network compliance check that was introduced in vCenter Server 2.5 Update 2. The enhanced network compliance check helps to ensure correct clusterwide heartbeat network paths. This setting enables you to bypass this check to prevent HA configuration problems. To bypass the check, add `das.bypassNetCompatCheck=true` to the HA advanced settings.

To configure HA with any of these advanced settings, just follow these steps:

1. Select your cluster in the VI Client and right-click it and select Edit Settings.
2. Select VMware HA in the left pane.
3. In the right pane, click the Advanced Options button.
4. When the Advanced Options window is displayed, double-click one of the blank fields under the Options column to edit it.
5. Enter an option name (for example, `das.failedetectiontime`) and press Enter.
6. Double-click the field next to it in the Value column and edit a value (for example, `60000`) and press Enter.
7. Repeat this procedure for any additional options that you want to add and click the OK button.

Configuring Distributed Resource Scheduler

Similar to HA, the Distributed Resource Scheduler (DRS) feature can be set up in a cluster either during its initial setup or afterward. To configure it, just select the cluster that you want to enable DRS on and right-click it and edit the settings for it or select the cluster and in the summary pane click the Edit Settings link. Put a check mark next to the Enable VMware DRS field on the General page, and it will be enabled for the cluster. You can optionally configure some additional settings to change the way DRS functions. To access the settings, click the VMware DRS item in the cluster settings window, as shown in Figure 4.15.

Choosing an Automation Level

When you enable DRS, you must select an automation level that sets the level of automation and controls how DRS will function. You have three levels to choose from:

- **Manual.** vCenter Server will only suggest migration recommendations for VMs. The suggested recommendations must be approved by an administrator before they are carried out. This level requires constant user interaction for DRS to function and is useful in smaller environments or when first enabling DRS to test out how it works before automating it.




Figure 4.15 DRS settings

- **Partially Automated.** VMs will automatically be placed onto hosts when they are powered on. vCenter Server will suggest migration recommendations for VMs. This mode is basically like Manual mode except for the fact that newly powered-on VMs are automatically placed onto hosts.
- **Fully Automated.** VMs will automatically be placed onto hosts when they are powered on and will automatically be migrated from one host to another to optimize resource usage. This mode has a migration threshold setting that controls how aggressive DRS acts when deciding when to migrate VMs to other hosts. The migration threshold ranges from conservative (5 stars) to aggressive (1 star).

When considering an automation level, it is usually best to choose Fully Automated and let DRS handle everything. However, when first enabling DRS, you might want to set it to Manual or Partially Automated so that you can observe its recommendations for a while before turning it loose on Fully Automated. Even when selecting Fully Automated, you can still configure individual VM automation levels so that you can specify certain VMs to not be migrated at all (disabled) or be set to Manual or Partially Automated. To configure individual VM automation levels, click Virtual Machine Options located under DRS, as shown in Figure 4.16.




Figure 4.16 DRS settings: Setting individual VM automation levels

Choosing a Migration Threshold

The five migration thresholds for DRS that you can choose from are as follows:

- **5 stars.** vCenter Server will only apply recommendations that must be taken to satisfy cluster constraints, such as DRS affinity rules and host maintenance.
- **4 stars.** vCenter Server will apply recommendations that promise a significant improvement to the cluster's load balance.
- **3 stars.** vCenter Server will apply recommendations that promise at least good improvement to the cluster's load balance.
- **2 stars.** vCenter Server will apply recommendations that promise even a moderate improvement to the cluster's load balance.
- **1 star.** vCenter Server will apply recommendations that promise even a slight improvement to the cluster's load balance.

Usually, the default three-star level is a good starting point and works well for most environments. Take care when choosing more aggressive levels because you could potentially have VMs moving very frequently between hosts (VM pong), which can create performance issues because of the constant VMotions, which causes an entire LUN to be locked during the operation (SCSI reservations). If you do choose a more aggressive level, ensure that the following conditions are all met:

- The VM's resource utilizations remain fairly even. If you have VMs that have frequent large CPU and memory fluctuations, this can cause constant DRS migrations.
- You do not have too many affinity rules configured. If you have a lot of affinity rules configured, it can make it difficult for DRS to move VMs around and can even cause more migrations to balance the resource load and satisfy the rules that are in place.
- The hosts in the cluster have similar hardware configurations. If you have a mix of hardware that is widely different (for example, two-CPU hosts and eight-CPU hosts), it can make it more difficult for DRS to balance the resource load to avoid contention, causing more migrations to take place.

Using DRS Affinity Rules

Rules can be configured in DRS to keep certain VMs either on the same host or separate hosts when DRS migrates VMs from host to host. These rules are known as affinity rules (not to be confused with CPU affinity) and are useful for ensuring that when DRS moves VMs around it has some limits on where it can place VMs. To create a rule, follow these steps:

1. In the DRS section of the cluster properties, click Rules.
2. Click the Add button to add a new rule. Enter a descriptive name for the rule. Choose the rule type ("Keep Virtual Machines Together" or "Separate Virtual Machines") and select the VMs that you want to be part of the rule, as shown in Figure 4.17. If you choose the "Separate Virtual Machines" rule, you can choose only two VMs for the rule. You can choose more than two for the "Keep Virtual Machines Together" rule. Click OK when you have finished to save the rule.




Figure 4.17 DRS affinity rules: Adding a new rule

3. Your rules will display in the Rules section. You can check/uncheck them to enable or disable the rules. You can also select the rule and click the Details button to see information about the rule and whether it conflicts with another rule. Rule conflicts are when two rules for a VM cannot both be satisfied, and will be displayed as a small triangle icon with a red exclamation point.

Let's look at some reasons why you might want to keep specific VMs on the same or different host servers.

You might want to keep VMs on the same host under the following circumstances:

- If the VMs are on the same vSwitch and port group, all network traffic between them never leaves the host server. Therefore, you might want to keep them on the same host to improve network response between them and minimize traffic on the physical network. An example of this is a tiered application that runs on multiple VMs, such as a web, application, and database server.

You might want to keep VMs on different hosts under the following circumstances:

- For servers that are clustered or redundant, like Active Directory, DNS, or web servers, you should keep them on separate hosts so that a single ESX failure does not affect both redundant or clustered servers at the same time. Doing this ensures that at least one stays up and remains available while the other recovers from a host failure.
- For servers that have high I/O workloads, you might want to separate them so that you do not overburden a specific host with too many high-workload servers. Because DRS does not take into account network and disk workloads when moving VMs, creating a rule for servers that are known to have high workloads in those areas can help avoid disk and network I/O bottlenecks on your hosts. For example, if you have two Exchange servers that have very high disk I/O, ensuring they stay on separate hosts ensures that any one host does not experience the combined high disk I/O of multiple servers.

In general, try to limit the number of rules that you create and only create ones that are necessary. Having too many rules makes it more difficult for DRS to try to place VMs on hosts to balance the resource load. Also, watch out for conflicts between multiple rules, which can cause problems. An example of this is if you have a rule that says VM1 and VM2 should be on different hosts and a rule that says VM2 and VM3 should be on the same host and another that says VM1 and VM3 should be on different hosts, you will run into problems.

Monitoring DRS and Applying Recommendations

After you have DRS enabled, you can monitor it by selecting the cluster in vCenter Server and choosing the Summary tab, as shown in Figure 4.18.




Figure 4.18 Cluster Summary tab in vCenter Server

Here you can see the following items:

- The total number of migrations that have taken place so far.
- The current failover capacity (hosts available to the cluster for failover) and configured failover capacity (number of hosts allowed to be used for failover).
- Whether the EVC (Enhanced VMotion Compatibility) mode is enabled.
- Total amount of CPU and memory resources from all the hosts in the cluster.
- The current DRS migration automation level (Manual, Partially Automated, or Fully Automated).
- The number of current recommendations awaiting approval.
- The current migration threshold (number of stars).
- DRS resource distribution, which has two graphs: Utilization percent shows the number of hosts on one axis and CPU (blue bars) and memory (yellow bars) utilization on the other axis. If a cluster is unbalanced, you will see multiple bars across different utilization levels. The Percent of Entitled Resources Delivered graph shows the number of hosts on one axis and the percentage of entitled resources delivered for each host on the other axis. DRS computes a resource entitlement for each VM based on

share/limit/reservation settings and current demand, and then computes a resource entitlement for each host by adding the entitlements for all VMs running on that host. The percentage is equal to the host's capacity divided by the entitlements of the VMs. So in a cluster where all entitled resources are being delivered, the graph should have a single bar for each resource in the 90% to 100% range. If this is not the case, your DRS level may be set too conservatively, or you may have affinity rules in place that are preventing the cluster from balancing properly.

In addition, you can select the DRS Recommendations tab in vCenter Server to display any pending recommendations and the DRS action history, as shown in Figure 4.19.




Figure 4.19 DRS Recommendations tab in vCenter Server

By default, DRS recommendations are generated every five minutes. If you do not want to wait, you can click the Generate Recommendations button to generate them immediately. You can also click Apply Recommendations to automatically apply the pending recommendations. All DRS recommendations that are displayed include the priority (number of stars), recommendation (which VM and to/from host server), and the reason (for example, satisfy affinity rule, balance average CPU load, balance average memory load). The DRS action history displays a list of all migrations that have taken place and which host that newly powered-on VMs were placed on.

Additional DRS Best Practices

You should consider some additional best practices when configuring DRS:

- DRS relies on VMotion, so make sure all your hosts in the cluster have compatible CPUs and similarly configured network labels. Remember that you can't VMotion between Intel and AMD processors, so it's best not to put them in the same cluster. Create separate clusters for your Intel and AMD hosts.
- The default DRS cluster evaluation frequency is every 5 minutes, which is configurable between 1 and 60 minutes. (The value for this setting is in number of seconds.) The default value works well in most cases, and it is not recommended to lower it, because doing so can cause increased overhead as the DRS algorithm is computed more frequently. To change the frequency, you must edit the `vpdx.cfg` file, which is usually located in the `C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter` directory on the vCenter Server, and add the following lines:

```
<config>
  <drm>
    <pollPeriodSec>
      300
    </pollPeriodSec>
  </drm>
</config>
```

- Try to have your VMs in DRS automatic mode whenever possible, because they are considered for cluster load-balance migrations across the ESX hosts before VMs that are not in automatic mode.
- Only provide your VMs with the number of vCPUs and memory that they need. Assigning VMs excessive resources adds more constraints when migrating VMs. VMs with less memory and vCPUs make it easier for DRS to migrate them when balancing them across the cluster.
- DRS works best when hosts in a cluster have CPU and memory configurations that are as similar as possible. If you have widely different hardware between your hosts, consider splitting them up into multiple clusters of similar hardware configurations. If hosts have similar CPUs but different frequencies and memory sizes, DRS will typically locate VMs on the hosts with high CPU frequencies and memory sizes because they have more capacity to handle peak loads.
- If HA strict admission control is enabled, DRS does not evacuate VMs from a host for the purpose of placing it in maintenance mode or standby mode if placing the host in this state would violate failover requirements. You can still manually evacuate VMs to

2. Check that your host servers' NICs support WOL. You can do this by clicking Configuration, Network Adapters in the VI Client, as shown in Figure 4.21. In the Network Adapter list, look at the last column to see whether WOL is supported. The NIC that DPM uses is the one assigned to the VMkernel vSwitch, so you might need to rearrange your NICs so that you have one that supports WOL in the VMkernel vSwitch.

Network Adapters					
Device	Speed	Configured	vSwitch	Observed IP ranges	Wake on LAN Supported
Broadcom NetXtreme II BCM5708 1000Base-T					
vmnic1	1000 Full	Negotiate	vSwitch3		Yes
vmnic0	100 Full	Negotiate	vSwitch0	172.20.11.16-172.20.11.23	Yes
Intel PRO/1000 PT Low Profile Quad Port PCIe Gigabit Ethernet Controller					
vmnic5	100 Full	100 Full	vSwitch4		No
vmnic4	1000 Full	Negotiate	vSwitch1	172.20.1.60-172.20.1.60 ...	No
vmnic3	1000 Full	Negotiate	vSwitch1	172.20.1.60-172.20.1.60 ...	No
vmnic2	1000 Full	Negotiate	vSwitch1	172.20.1.60-172.20.1.60 ...	Yes

Figure 4.21 Checking the host NICs in the VI Client to see whether they support WOL

- To test the WOL feature, select your host and choose the Enter Standby Mode option, which powers down the host. Your host will begin shutting down and should power off after a few minutes. You can wake it back up by selecting the host in the VI Client and right-clicking it and choosing the Power On option.
- After you have verified that the WOL feature works properly, you need to enable DPM. To do this, edit the settings for your cluster. Then, under the DRS category, select Power Management, as shown in Figure 4.22. You can then select either Off, Manual, or Automatic options. The Manual option will only make recommendations for powering off hosts. The Automatic option will enable vCenter Server to automatically execute power management recommendations. You can change settings on individual hosts to have them use the cluster default, always use manual, or always use automatic. Or, you can disable the feature for them by changing the setting in the Power Management column.

DPM Considerations

You should consider a few things when enabling DPM:

- If you are using a monitoring system to monitor that your ESX servers are up or down, you will trigger an alert whenever a host is shut down. Having servers going up and down automatically can generate a lot of confusion and false alarms. In addition, many datacenters measure uptime statistics on all the servers. Having hosts going down during periods of inactivity can significantly skew those numbers. If you use DPM, consider adjusting your operational procedures to exclude the ESX hosts from monitoring and instead only monitor the VMs.
- If you are using HA in your cluster, be aware that DRS and DPM maintain excess powered-on capacity to meet the HA settings. Therefore, the cluster may not allow additional VMs to be powered on or some hosts may not be powered down even though the cluster may appear to be sufficiently idle. In addition, if HA strict admission control is enabled, DPM will not power off hosts if doing so would violate failover requirements. If HA strict admission control is disabled, DPM will power off hosts even if doing so violates failover requirements.




Figure 4.22 Enabling the DPM feature

- DPM can work only on hosts that have WOL capabilities and are running ESX 3.5.
- In its initial release (ESX 3.5), DPM is considered experimental and not to be used for production use. As of ESX Update 3, it is still considered experimental, but this should change in a later release or update of ESX.
- Having more hosts in your cluster will provide you the opportunity to maximize power savings when using this feature. If you have a large environment, this feature can definitely help save you a significant amount of money.
- Similar to DRS, the DPM feature works best if all hosts are in automatic mode, which gives DPM more flexibility in selecting hosts to power on and off. If hosts are in manual DPM mode, DPM must wait for approval to perform its action, which can limit its effectiveness. DPM is more biased toward hosts in automatic mode than manual mode because of this. Consider using automatic mode whenever possible and disabling DPM on hosts that you do not want to be powered off.
- DPM will consider historical demand when determining how much capacity to keep available and keeps some extra capacity available for changes in demand.

Configuring VMotion

Before configuring VMotion on your host servers, confirm they meet the requirements for using it, which are listed here:

- You must have shared storage connected to your host server; this can be either Fibre Channel SAN, iSCSI, or NFS.
- Gigabit Ethernet connection for the VMkernel network.
- Host must be licensed for the VMotion feature.

To set up VMotion, you must first set up the VMkernel networking stack, which is used for VMotion, and connections to NFS and iSCSI storage by creating a port group on a vSwitch. To set it up, follow these steps:

1. Select the host server in vCenter Server, select the Configuration tab, and then choose Networking.
2. If you have an existing vSwitch that you want to add the VMkernel to, click the Properties link for that vSwitch, and then click the Add button. Otherwise, click the Add Networking link to add a new vSwitch to be used by the VMkernel.
3. At the Connection Type screen, select VMkernel and then click Next.
4. At the Network Access screen (new vSwitch only), select the Create a Virtual Switch option and either one or two NICs to be used with the vSwitch, as shown in Figure 4.23, and then click Next. By selecting two NICs, you provide redundancy (in case a single NIC fails or loses connectivity).
5. At the Connection Settings screen, enter a network label or use the default VMkernel name. (Whichever name you use needs to be consistent across all hosts in a cluster.) Add a VLAN tag if you are using VLAN tagging on the NIC that you are using; otherwise, leave it blank. Put a check mark by the Use This Port Group for VMotion option. Enter a unique IP address for the VMkernel (must be different from the Service Console IP address). Finally, enter a network subnet mask and click Next, as shown in Figure 4.24.
6. At the Summary screen, verify everything is correct and either press Back and make any changes or click Finish. In most cases, a message will display that a default gateway is not configured and asking you whether you want to configure one now.




Figure 4.23 Configuring the VMkernel: Network Access screen




Figure 4.24 Configuring the VMkernel: Connection Settings screen

7. Select Yes, and a DNS and Routing Configuration screen will appear where you can enter the gateway IP address for the IP address that you assigned to the VMkernel. The VMkernel IP address is in a separate field from the Service Console (vswif) IP address, but can potentially use the same default gateway if they are both on the same VLAN. You can change this later by selecting the Configuration tab and then DNS and Routing and selecting the Routing tab. When you save your settings, try to ping the VMkernel gateway with the `vmkping` command inside the Service Console to make sure it responds. If the VMkernel is on a nonroutable network, confirm you ping it from another host on that same network.
8. You should now have a vSwitch that is configured similarly to the one in Figure 4.25 and are ready to start using VMotion. Make sure your other hosts are also configured similarly.




Figure 4.25 VSwitch configured with the VMkernel for VMotion

VMotion Considerations

Setting up VMotion can be tricky because you need to be aware of several requirements and compatibility issues. Here are some considerations that you should know about when using and configuring VMotion:

- In versions earlier than ESX 3.5, VMs that had their swapfile (VSWP) file not located on shared storage could not be moved with VMotion. The reason for this is that the destination host would not be able to access the VSWP file that is located on the source host's local disk. Beginning with ESX 3.5, support for using VMotion on VMs that have local VSWP files was added. If a VM with a local VSWP file is VMotioned, the VSWP file is re-created on the destination host, and the nonzero contents of the VSWP file are copied over as part of the VMotion operation. This can cause the VMotion operation to take slightly longer than normal because of the added VSWP copy operation in addition to the normal CPU and memory state copy operation. Using a local swapfile datastore can be advantageous because it frees up valuable and expensive shared disk to be used for other things such as snapshots and virtual disks.
- If your VMs have their CD-ROM drive mapped to either a host device or a local ISO datastore, they cannot be VMotioned, because the destination server will not have access to it. In addition, if the CD-ROM is mapped to a shared ISO datastore, make sure all ESX hosts are able to see that shared ISO datastore. Consider using a shared ISO datastore on a VMFS volume or alternately on an NFS or Samba share instead.

- If you try to VMotion a VM with running snapshots from one host to another, you will receive a warning saying that “Reverting to snapshot would generate errors (warnings) on the destination host.” This is simply a warning that reverting to a snapshot after migration with VMotion might cause the VM to crash, because the Migration Wizard cannot verify the compatibility of the VM state in the snapshot with the destination host. In addition, if you have changed the default locations for snapshot files of the VM (or VSWP files on 3.0.x hosts), the VM will crash when the migration is complete if the destination host cannot access the same storage that the files are located on as the source host. If your VM has all its files on shared storage that is accessible to all ESX hosts, you will not have any issues.
- It’s very important to ensure that vSwitch network labels are identical (case sensitive) across all hosts. If they are not, you cannot VMotion a VM between two hosts that do not have the same network labels configured on their vSwitches.
- CPU compatibility is one of the biggest headaches when dealing with VMotion because VMotion transfers the running architectural state of a VM between host systems. To ensure a successful migration, it is required that the processor of the destination host be able to execute the equivalent instructions as that of the source host. Processor speeds, cache sizes, and number of cores can vary between the source and destination hosts, but the processors must come from the same vendor (either Intel or AMD) and use compatible feature sets to be compatible with VMotion. When a VM is first powered on, it determines its available CPU feature set based on the host’s CPU feature set. It is possible to mask some of the host’s CPU features using a CPU compatibility mask to allow VMotions between hosts that have slightly dissimilar feature sets. See VMware knowledge base articles 1991, 1992, and 1993 for more information about how to set up these masks. In addition, you can use the Enhanced VMotion feature that was introduced in vCenter Server 2.5 Update 2 to help deal with CPU incompatibilities between hosts.
- It is a recommended security practice to put your VMotion network traffic onto its own isolated network so that it is accessible only to the host servers. The reason for this is twofold. First, VMotion traffic is sent in clear text and is not encrypted, so isolating it ensures that sensitive data cannot be sniffed out on the network. Second, it ensures that VMotion traffic experiences minimal latency and is not affected by other network traffic (because it is a time-sensitive operation).
- Migration compatibility checks are done using a number of criteria before a VMotion is attempted to ensure that it will be successful. When you select a destination host in the Migrate Virtual Machine Wizard, the Compatibility panel at the bottom will display any information about the compatibility of the selected host or cluster as it pertains to the VM’s configuration. If the VM is compatible, a “Validation succeeded” message will display. If the VM is not compatible with either the host or cluster, a warning or error message will display, as shown in Figure 4.26. Warning messages do

not prevent migration, and you can usually continue despite them. Errors message can prevent migration if there are no error-free hosts among the selected destination hosts. Network and datastore configurations are taken into account for clusters, and individual host configuration is used for hosts. In their knowledgebase article 1003684, VMware provides a complete overview of the various messages that appear during the compatibility check.




Figure 4.26 Warnings and errors displayed during VMotion compatibility check

Configuring Enhanced VMotion Compatibility

Enhanced VMotion Compatibility (EVC) was introduced in ESX 3.5 Update 2 to help ensure VMotion compatibility for hosts in a cluster. EVC ensures that all hosts in a cluster present the same CPU feature set to every VM, even if the actual CPUs differ on the host servers. EVC works only among the same CPU vendor, and it will not support compatibility between Intel and AMD CPUs.

Before you enable EVC, make sure your hosts meet the requirements for it:

- You must be running vCenter Server 2.5 Update 2 or later.
- All hosts in the cluster must be running ESX 3.5 Update 2 or later.
- All hosts in the cluster must have the CPUs from the same vendor (either Intel or AMD).

- All VMs in the cluster must be powered off or migrated out of the cluster when EVC is enabled.
- All hosts in the cluster must either have hardware live migration support (Intel FlexMigration or AMD-V Extended Migration) or have the CPU whose baseline feature set you intend to enable for the cluster. See VMware knowledge base article 1003212 for a list of supported processors.
- Host servers must have the following enabled in their BIOS settings: For AMD systems, enable AMD-V and No Execute (NX), and for Intel systems enable Intel VT and Execute Disable (XD).

When you are sure your hosts meet the requirements, you are ready to enable EVC. You can use two methods for doing this. (EVC cannot be enabled on existing clusters unless all VMs are shut down.) The first method is to create a new cluster that is enabled for EVC and then move your ESX hosts into the cluster. The second method is to shut down all the VMs in your current cluster or migrate them out of the cluster to enable it. The first method tends to be the easier way because it does not require any VM downtime. If you choose the first method, you can simply create a new cluster and then move your hosts one by one to it by first putting it in maintenance mode to migrate the VMs to other hosts. Then once the host is moved to the new cluster, you can VMotion the VMs back to the host from the old cluster to the new one. The downside to this method is that you have to re-set up your cluster HA and DRS settings on the new cluster, and you lose your cluster performance and migration history. When you are ready to set up EVC, just follow these instructions.

To set up EVC on a new cluster, follow these steps:

1. Create a new cluster, provide it a name, and enable HA and DRS if necessary, and set up any additional HA and DRS configuration. Note: You will not see the option to enable EVC until after you save the cluster.
2. After you have saved the cluster, go back and edit its settings, and you will see the EVC option, as shown in Figure 4.27.
3. Select the appropriate processor vendor (either AMD or Intel) for your hosts, and then click OK to save the settings.
4. Before you move hosts into the new cluster, power off its VMs or put it into maintenance mode so that the VMs migrate to other hosts. Once the host is in maintenance mode, you can move it into your new cluster.
5. Edit your cluster's settings again and select the EVC option. If your host's CPU is not compatible, you will see an error message. Make sure the appropriate BIOS settings are enabled.
6. Move the VMs back to your host using VMotion, and repeat the process for the remaining hosts.




Figure 4.27 EVC option in cluster settings

To set up EVC on an existing cluster, follow these steps:

1. Ensure that all VMs in the cluster are powered off or moved to other hosts outside of the cluster. Also, ensure that all hosts in the cluster have either Intel or AMD processors.
2. Edit the cluster settings and select the appropriate processor vendor (either AMD or Intel) for your hosts, and then click OK to save the settings. If your host's CPU is not compatible, you will see an error message. Make sure the appropriate BIOS settings are enabled.
3. Power your VMs back on or move them back into the cluster.

Listed here are some considerations that you should be aware of when using EVC:

- If you receive errors that EVC cannot be enabled because one or more hosts in the cluster are not compatible, verify that all of your hosts have CPUs that are capable of 64 bit. Most enterprise class server hardware purchased in the past few years should be capable of 64 bit. Next, make sure you have the proper settings enabled in your server's BIOS, as previously detailed in the requirements. If you do not, you will also receive this error.

- With EVC, all hosts in the cluster must have compatible CPUs or you cannot enable EVC. If you have some that are not compatible, consider creating a separate cluster for those hosts.

Configuring Resources

Resources are what virtualization is all about, so you should configure your clusters and hosts to optimize and protect each of the four valuable resources that it has (CPU, memory, network, and disk). Because you have multiple guest VMs all competing for the resources of a single host, you may want to set up controls so that you can prioritize and guarantee resources for certain VMs. There are many ways to accomplish this: using shares, reservations, limits, and resource pools.

Understanding Memory Usage on Virtual Hosts

Memory is typically one of the most wasted and needed resources on a host server. Guest VMs are typically given more memory than they actually need, and the amount of memory used by a VM may vary at any given time based on the applications that it runs. ESX uses some advanced techniques to try to control the amount of wasted physical memory that is in use by the VM. All physical memory on an ESX host is managed by the VMkernel, which allocates a small amount for the Service Console and the remainder for the VMs.

Memory Overcommitment

It is possible to overcommit memory on a host so that more memory is assigned to the VMs than the host physically has. For example, on a host with 8GB of RAM, it is possible to assign six VMs 2GB of RAM, each for a total of 12GB. However, just because a VM has 2GB assigned to it doesn't mean it is using 2GB of physical host RAM. ESX presents 2GB of RAM to a VM to make it think it has 2GB, but in reality the amount of physical RAM in use on the ESX host depends on how much memory is actually in use on the guest VM. So a VM that is assigned 2GB of RAM may actually only be using less than 1GB of physical host memory.

Now you may be wondering how the host can give the guest VMs more memory than it physically has. It can achieve this by using memory swap disk files (.vswp) for each VM that it can use in lieu of physical memory when that becomes exhausted, similar to how Windows uses a paging swapfile for virtual memory. There is a price to pay for using swapfiles, however, because it is significantly slower than physical memory because it relies on a slower mechanical device. However, memory overcommitment is not necessarily a bad thing. Because VMs typically do not use all the memory assigned to them, overcommitment can benefit the hosts by allowing more VMs to run on the host than you could normally run. If you do overcommit memory, however, it is important to monitor your hosts to make sure that

their physical memory is not exhausted as a result performance degrades resulting from the VMs swapping to disk to access additional memory. ESX has a few tricks up its sleeve to maximize its memory usage and allow for more memory overcommitment.

Transparent Page Sharing

Memory page sharing helps reduce the amount of physical memory in use by VMs that are running, similar to operating systems and applications that have some of the same common data loaded into the memory of each. For example, all Windows Server 2003 servers typically load many of the same services and DLLs into memory when they start up. Instead of having ten VMs with the same redundant copy of that memory page using physical memory on the host, the nine redundant copies are eliminated, and the one remaining copy is shared among all ten VMs. This allows for the physical memory used by the nine redundant copies of that memory page to be freed up on the host server to be used for other things.

Using a special hashing algorithm to scan for memory pages that have similar content, transparent page sharing (TPS) works by comparing pages of memory. When it finds pages that the hash algorithm has identified as potential matches, it performs a binary comparison of the two to ensure they are identical. ESX then frees up one of the pages by updating the memory mappings for both VMs to point to the same memory page. If for some reason the page is later modified by one of the VMs, ESX first makes a copy of the memory page so that the VM that needs to write to it has its own page and will not affect the remaining VMs that are still using the shared page. Page sharing can occur both within a single VM or across multiple VMs.

TPS does take time to work, and you will not immediately experience the effects of it on newly powered-on VMs. TPS works during periods of idleness on the host server as to not affect the performance of the host server. The memory savings by using TPS can vary depending on the diversity of the VMs running on the host, but can be anywhere from 5% to 30% memory reduction or in some cases even much higher. TPS is enabled by default on all hosts and is a very beneficial feature to help conserve physical memory on your hosts and achieve higher levels of overcommitment.

Memory Balloon Driver

ESX uses a memory balloon driver that is part of the VMware Tools application (vmmemctl) that should be installed on all VMs to help encourage the guest VM's operating system to reclaim memory. The balloon driver is basically an application that runs on the guest operating system and inflates (increases its memory footprint) or deflates (decreases its memory footprint) as needed to help trigger the guest operating system's own memory management algorithms. Thus, when memory is tight on the host server, the guest operating system will decide which of its pages to reclaim and swap them to its own paging file thereby freeing up more physical memory on the host server, as shown in Figure 4.28.




Figure 4.28 Memory balloon driver in action

Understanding Shares, Reservations, and Limits

Shares, reservations, and limits can be set on VMs and resource pools to control the amount of host resources that a VM or resource pool can use. You can use any one or a combination of these controls to customize resource usage in your environment. Let's talk about how each one works.

Shares

Shares are used to assign an importance to VMs and resource pools to control the amount of resources (CPU, memory, or disk) that each VM is entitled to. Shares are useful to assign a general priority for resource usage to VMs without having to set specific reservations and limits on them. A VM that has more shares than another VM is entitled to more resources. Shares can be set to Low, Normal (default), High, and Custom for each VM, and each setting has a numeric value associated, as listed here:

- **Low.** 500 shares per vCPU and 5 shares per megabyte of memory assigned to the VM
- **Normal.** 1,000 shares per vCPU and 10 shares per megabyte of memory assigned to the VM
- **High.** 2,000 shares per vCPU and 20 shares per megabyte of memory assigned to the VM
- **Custom.** Can be set to any value

You can also assign shares on VM disk usage, but they do not have numeric values associated with them. The amount of resources that each share represents changes based on the number of VMs that are running on the host server. When a VM is assigned a share level, it is entitled to a certain amount of resources, but it does not limit the VM to have more resources if they are available. Resource entitlements from shares come into play only when the host's CPU resources become constrained, because they are designed to prioritize resources for certain VMs.

Using shares has the following benefits:

- Easier to use than reservations and limits because they are more flexible and are not set to a fixed amount of resources.
- Ensure that critical VMs get priority access to host resources.

The drawback of using shares is as follows:

- Shares do not guarantee a fixed amount of resources for VMs, just the fact that they get priority access to available resources.

Reservations

Reservations are used to guarantee a certain amount of host resources (CPU or memory) to a VM or resource pool. When you set a reservation on a VM, the host will ensure that the VM always gets that amount of resources if it needs it. Reservations are set in the unit of the specific resource. CPU reservations are set in megahertz, and memory reservations are set in megabytes. By default, reservations are set to zero for all VMs. Because reservations guarantee a certain amount of resources, a VM will not be allowed to be powered on if the host does not have enough resources to guarantee a reservation. For example, if you have a host with 8GB of memory and 6GB is in use by other VMs and you try to power on a VM with a 3GB reservation, it will not power on.

Did You Know?

Memory reservations affect the size of the VSWP file that is created when a VM is powered on, to allow for memory overcommitment on the host server. A VSWP file will be created equal to the amount of memory that the VM is assigned minus the memory reservation configured for the VM. So, a VM with 4GB of memory and no memory reservation will have a 4GB VSWP file created on disk. If you set a 1GB memory reservation on the VM, a 3GB VSWP file will be created instead. If you make a change to a memory reservation while a VM is powered on, the VSWP file size will not change until the VM is powered off and back on.

The benefits of using reservations are as follows:

- The reduced size of the VSWP file frees up more disk space on your VMFS volumes.
- Guarantee of resources for critical VMs.

The drawbacks of using reservations are as follows:

- If enough resources are not available for a VM with a reservation, it will not start.
- Reservations can be more difficult to manage than shares because they are based on numeric resource amount rather than percentages.

Limits

Limits specify the maximum amount of host resources (CPU or memory) that a VM can use. Limits are also set in the unit of the specific resource. CPU reservations are set in megahertz, and memory reservations are set in megabytes. By default, VMs are set to use unlimited host resources if they are available on the host (although a VM cannot use more memory than it has been assigned). Setting a CPU limit on a VM will always limit the VM to the amount set regardless of the amount of resources available. Setting a limit on memory will limit the amount of physical host memory that a VM can use; once that limit is reached, the VM will instead use its VSWP file for memory.

The benefit of using limits is as follows:

- Ensures that noncritical VMs do not hog resources on a host.

The drawbacks of using limits are as follows:

- Holds back a VM by allowing only a certain amount of resources, even when the host has plenty of unused resources.
- Limits can be more difficult to manage than shares because they are based on numeric resource amount rather than percentages.

Understanding Resource Pools

Resource pools enable you to split a host's or cluster's resources (CPU and memory only) into pools so that they can be assigned to VMs. They can also be used for delegation of administration and to simply help organize VMs. Every stand-alone host and DRS cluster has an invisible root resource pool containing all the resources of the host or cluster. You can optionally create child resource pools to split the host's or cluster's resource into separate pools. Child resource pools own part of the parent's resources and can contain additional child resource pools to form a resource hierarchy. Resource pools can contain additional resource pools and VMs and are referred to as siblings if they are at the same level, as shown in Figure 4.29.




Figure 4.29 Resource pool hierarchy

When you create resource pools, you allocate CPU and memory to them using shares, limits, and reservations. Reservations in resource pools can be set to expandable so that they can draw more resources from the parent pool if they are available. If a resource pool is set to not be expandable, it is unable to draw additional resources from the parent even if they are available. Some reasons you might want to use resource pools are as follows:

- Splitting a host's or cluster's resources up into separate pools to be used by different departments within your organization, as shown in Figure 4.30.




Figure 4.30 ESX host configured with resource pools for different departments

- Segregating resources for development, test, and production VMs on your hosts so that your production VMs have access to more resources.
- You can delegate control of resource pools to certain individuals or groups so that they have access to create VMs inside their own pool.

Resource pools can get a bit complicated if you have many of them with multiple children. Some things to know about using resource pools are as follows:

- Be aware of admission controls when powering on VMs in resource pools. If a pool has a nonexpandable reservation and there are not enough resources available in the pool,

the VM will not be powered on. If the pool has an expandable reservation, the host will check to see whether enough resources are available in the parent resource pool or to other resource pools above the parent. If they are, the VM will be powered on. If not, the VM will not be powered on.

- If you add a host to a cluster that does not have DRS enabled, the host's resource pool hierarchy is discarded, and no resource pool hierarchy can be created.
- If you add a host to a cluster that does have DRS enabled, you will be prompted to choose what to do with the host's resources. The default is to discard the host's resource pool and add the host's resource to the root resource pool. You can optionally choose to graft the host's resource pools into the cluster's resource pool hierarchy.
- If you change your cluster's capacity by either upgrading existing hosts or adding/removing hosts, revisit your existing resource pools and change their allocations accordingly.

Configuring Shares, Reservations, and Limits

Now that you understand what they are, let's look at how to use them. To access the resource settings on a VM, select it and choose Edit Settings, and then select the Resources tab, as shown in Figure 4.31.




Figure 4.31 VM resource settings

Here you can change the resource settings for CPU, memory, and disk for the VM. For shares, you choose a level or set a custom level. For reservations, you can choose an amount between zero and either the combined processor speed of a host or the total memory assigned to a VM. This will be the amount that will be reserved for the VM. This can be a bit misleading, because even though you can select an amount for CPU that is higher than an individual CPU clock speed, the most a single-vCPU VM will be able to use is the maximum clock speed of a single-host CPU. (On a host with two dual-core 2.4GHz processors, the max speed is 2.4GHz for a single-vCPU VM, not 9.6GHz.) For memory, you are limited to setting the reservation up to the amount of memory the VM has.

To use limits, you have to uncheck the Unlimited setting, and you can then set a limit on the VM. Again, limits are a bit misleading. You can set a memory limit higher than the amount of memory a VM has assigned to it, but the VM cannot use more memory than it has. The same is true with CPU limits. A single-vCPU VM cannot use more CPU clock speed than the maximum clock speed of a single host CPU.

Configuring Resource Pools

Resource pools can be created on stand-alone ESX hosts, on clusters, and on other resource pools. You cannot create a resource pool on an individual ESX host if it is a member of a DRS-enabled cluster. To create a resource pool, follow these steps:

1. Select the object where you want to create the resource pool and choose the New Resource Pool option. A window will appear, as shown in Figure 4.32, where you can set the attributes for that resource pool.




Figure 4.32 New resource pool settings

2. Enter a name for your resource pool.
3. In the top section, you specify CPU resources. In the bottom section, specify memory resources. You can select a share setting and specify a reservation and limit for your pool. Reservations cannot exceed the combined total of all resource pools; a yellow triangle will show you the limit of the available resources. For example, if you have 10GHz in your root resource pool and already have one resource pool with a 4GHz reservation, the most you can allocate to a sibling resource pool is 6GHz. To specify a limit, uncheck the Unlimited field. Reservations can be made expandable so that the child resource pool can borrow resources from the parent if they are available.
4. After you have set the resources for the pool, click OK. The pool will be created and displayed in the Inventory panel.

Viewing Resource Pools

To view information about your resource pools, select a resource pool, and then select the Summary tab, which will display various information about your resource pool, as shown in Figure 4.33.

PoolA	
Getting Started Summary Virtual Machines Resource Allocation Performance Tasks & Events Alarms Permissions Maps	
General	
Number of Virtual Machines:	2/2
Number of Child Resource Pools:	1/1
CPU	
Shares:	High (8000)
Reservation:	4000 MHz
Type:	Expandable
Limit:	Unlimited
Unreserved:	4000 MHz
Commands	
New Virtual Machine New Resource Pool Edit Settings	
Resources	
CPU usage:	0 MHz
Memory usage:	0 MB
Memory	
Shares:	Normal (163840)
Reservation:	2048 MB
Type:	Expandable
Limit:	Unlimited
Unreserved:	6034 MB

Figure 4.33 Resource pool Summary page

To view more detailed information about your resource pools, select either a stand-alone ESX host, cluster, or resource pool, and select the Resource Allocation tab, as shown in Figure 4.34. To view information about child resource pools, select them, and then select their Resource Allocation tab.




Figure 4.34 Resource pool Allocation page

You can toggle between viewing CPU and memory in your pool by clicking the buttons on the page to change the view. On this page, you can view the CPU and memory reservation information, including the total amount reserved, used, or unreserved and whether the reservation is expandable. You can also view detailed information about your VMs that are part of the pool and also child pools.

Adding VMs to Resource Pools

To add new VMs to your resource pools, simply create a VM on your host server, and one of the sections of the New VM Wizard will enable you to select a resource pool for your VM.

To add an existing VM to a resource pool, select the VM and drag it to the resource pool that you want it to be located in. The VM will be moved, and any existing resource settings on the VM will be unchanged. You may receive a resource settings warning when moving a VM into a resource pool stating that the VM will not be entitled to many resources because of the current resource settings. You can change either the resource pool settings or the VM's resource settings to correct this. After the VM is moved into the resource pool, the % Shares value will be updated to show the total number of shares in use in the resource pool.

By the Way

VMware has written a whole guide (Resource Management Guide) on resource management that is updated with each new release of ESX and vCenter Server. It is a must-read if you want to know more detailed information about understanding and configuring resources. You can access the latest version by going to http://vmware.com/support/pubs/vi_pubs.html and selecting the version that you are using.

Configuring ESX and ESXi Hosts

Now it's time to configure the various components of your host so that you can start adding VMs to it. Networking and storage are the two most critical configuration pieces of your host servers, and there are many different ways to configure to meet your requirements. Let's get started with configuring the networking on your hosts.

Configuring Networking

Be sure you understand your requirements and environment before configuring your networking on your host servers. It's a good idea to consult with your network group to ensure that they understand the needs of your host servers when connecting them to the physical network. It is common to use 802.1Q VLAN tagging with ESX hosts because it allows you to use multiple VLANs on a single pNIC. You may also want to involve your security group because they may have concerns when it comes to how your host connects to the network. It's important to make sure both groups have a firm grasp of how virtual hosts integrate into the environment because some of the traditional concepts do not work well with virtual hosts.

Network Design Considerations

You should first figure out your requirements for your host servers. This includes what VLANs they will need to connect to, how much redundancy you want to have, if you will be using VMotion or not, and whether you will be using IP-based storage. You should also consider how much bandwidth you will need to support your VMs and whether you will need to connect to external networks. If you plan to use VLAN tagging, you need to decide which method you will use.

You should start by totaling up the number of NICs that you have in your host servers because this will ultimately affect your network design. The more NICs you have available, the more redundancy you can provide on your vSwitches. After you have this number, you need to define your VLAN requirements for the VMs that will be running on the host. If you will be using many VLANs, you will most likely have to use VLAN tagging; otherwise, you will have to use one NIC for each VLAN that you want to configure on your host server. Let's take a minute to talk about VLAN tagging.

802.1Q VLAN tagging enables a single NIC and physical switch port to support multiple VLANs, and is commonly referred to as a trunk port. A VLAN tag is added to all traffic that goes over the NIC that signifies what network VLAN the traffic belongs to. The physical switch that is connected to the NIC reads this tag and routes the traffic appropriately. There are several different methods for tagging the network traffic, as follows:

- **Virtual switch tagging (VST).** This method is the most commonly used with ESX and uses port groups configured on vSwitches with a specific VLAN ID (1 to 4094). The

VLAN ID is applied to all network traffic by the vSwitch so that the physical switch that it is connected to knows which VLAN the traffic is from. The vSwitch port group tags all outbound frames and removes tags for all inbound frames and is transparent to the guest operating system of the VMs. This method is fairly simple to set up and use and requires no special VLAN drivers to be installed on the VMs.

- **Virtual machine guest tagging (VGT).** This method requires a special VLAN driver to be installed inside the operating system of the guest, which is responsible for tagging all network traffic for the VM. This method also uses port groups, with one specific VLAN ID (4095) configured on the vSwitches, and traffic passing through the vSwitch already has a tag applied by the VM's operating system. The downside of this method is that it can be more complicated because you have to install and configure drivers in all of your VMs that need to use VGT. In addition, drivers that support this method may not be available for all operating systems. Also, it might cause increased CPU usage on your VMs because of the extra overhead of tagging all the frames by the operating system. Use this method if your VM needs access to more than three VLANs.
- **External switch tagging (EST).** This method uses external switches to do the tagging and is similar to traditional physical network using VLANs. The tag is appended when a packet arrives at the physical switch port and stripped away when a packet leaves the physical switch port toward the server. The drawback to this method is that it requires a pNIC in your ESX host for each VLAN that you need to connect.

So basically, the three methods come down to this: VST tags at the vSwitch, VGT tags at the OS, and EST tags at the physical switch. Which method you use depends on your network environment. VST is typically the best and most flexible method to use, but you must have physical network switches that can support it. EST is the default for ESX, and if you choose to use it, there is no setup required on the ESX host. VGT requires installing operating system drivers and is a bit more difficult to manage and support. If you plan to use very few VLANs on your hosts, you can use EST. If you have the capability to use VST, however, I strongly recommend using it.

By the Way

VMware has a white paper that provides further detail on the three 802.1Q tagging methods and describes how to set them up. It is available at http://www.vmware.com/pdf/esx3_vlan_wp.pdf.

To help visualize how to lay out your network configuration, you can draw it out on paper. List out your individual NICs, and then map them to the network connections (Service Console, VMkernel, VM VLANs, and so on) that you need, as shown in Figure 4.35.




Figure 4.35 Laying out your virtual networking on your host server

You should also take into consideration redundancy on your vSwitches to provide maximum reliability and throughput. For the best redundancy, make sure that you use NICs from separate physical multiport NIC cards in your vSwitches so that a single multiport card failure does not take out all the NICs on a vSwitch. Likewise, ensure that your pNICs on your vSwitches plug into separate physical switches if possible so that a single switch failure does not take out all the NICs on a vSwitch.

In addition, consider isolating your Service Console and VMkernel networks from your user and VM networks for both security and performance reasons. VMotion traffic, which uses the VMkernel network, is sent in clear text over the network. Also, isolating them protects them from other network traffic that may cause latency on the network. Consider isolating them on separate physical switches rather than using VLANs.

Network DMZ Considerations

If you plan to have your host server run VMs that will be connected to a public-facing external network (demilitarized zone, or DMZ), consider some security precautions to protect your host servers and internal network. Many network and security groups become concerned with this because the ESX host can potentially be straddling the internal and external networks, thereby creating a potential path for malicious activity into the less-secure internal network from the outside. You might be pressured to put the ESX host entirely in the DMZ, which is generally not a good idea because it exposes the host management network (Service Console) to the outside network, where it can be exploited and attacked. If the ESX Service

Console is compromised, an attacker could easily get access to all the VMs on that host. What is typically done for deploying VMs to the DMZ is to have the host on the internal network and then create a separate isolated vSwitch to be used for the DMZ network. In this configuration, even if a VM connected to the external network is compromised, your internal network would still be secure because an attacker would be limited to accessing devices on your external network only and would not be able to break through the VMkernel to access VMs on your internal network. (This doesn't mean it can't happen; but as long as you keep your host patched and properly secured, the chances of it happening are almost nonexistent. Historically, ESX has a good record when it comes to security and vulnerabilities.) If you use this method, take care that you do not have VMs plugged into both your external vSwitch and internal vSwitches, because if the VM is compromised an attacker will have access to your internal network. We discuss DMZ and ESX security more in the next chapter.

Configuring the Host Management Network

The host management network is referred to as the Service Console on ESX hosts and simply the management network on ESXi hosts. This network is what is used for administration and vCenter Server connections and is critical to the well-being of the host.

Configuring the ESX Service Console

As part of your ESX installation, you should already have a default Service Console set up, which you may want to modify to add redundancy to it. ESX networking can be configured using two different methods: using the VI Client and using the Service Console command-line interface (CLI). You can use the VI Client as long as you have a working network connection already configured on your ESX Service Console. If you do not, you must log on to the ESX host from the physical console and set up or correct your networking configuration using the CLI commands.

Here's how to configure networking using the VI Client:

1. Connect to the host (or vCenter Server) with the VI Client, and then select the Configuration tab and then the Network Adapters link. Here you can see all the pNICs that your ESX has in it, as shown in Figure 4.36. Ensure that all the NICs that your host has are all showing up. You will see them grouped by the network device type (NIC model), and it will show their name (vmnic#), speed, how they are configured (auto-negotiate or not), which vSwitch they are currently connected to, observed IP ranges (which VLANs the NIC sees traffic on), and whether WOL is supported.
2. On the Configuration tab, select the Networking link. This is where you configure your vSwitches. Your default network should look similar to the one in Figure 4.37 (you might not have a VM network if you did not select the option when installing ESX) with a Service Console network (vswif0) already configured on vSwitch0.




Figure 4.36 Physical NICs displayed on an ESX host




Figure 4.37 Default network on an ESX host

3. To add pNIC redundancy to the Service Console network, click the Properties link for the vSwitch, and then click the Network Adapters tab. Then click the Add button, select an unused pNIC, and click Next. As mentioned earlier, you should try to use a pNIC that is not on the same physical multiport card and is plugged into a different physical switch than your other Service Console NIC.
4. At the NIC Order window, you can change the order of the NICs and can make one of the NICs a standby NIC. (By default, they are both active.) A standby NIC is not used until the active has failed. This is useful in certain vSwitch configurations where you have the Service Console and VMkernel running on the same vSwitch. In that situation, you can specify NIC1 as active for the Service Console and NIC2 as a standby and NIC2 as active for the VMkernel and NIC1 as a standby.
5. After you click the Finish button, your vSwitch should look similar to the one in Figure 4.38, with two pNICs attached to it rather than one, thereby providing protection in case one NIC fails.




Figure 4.38 Service Console vSwitch configured with two pNICs

By the Way

Seeing is believing. Once you configure a Service Console vSwitch with a second pNIC so it is redundant, try continuously pinging the Service Console from a workstation. Then unplug one of the network cables and see what happens. Then plug it back in and unplug the other one. While you do this, look at your ping results. You might miss one or two pings while the other NIC takes over for the one that was unplugged, but your Service Console network will stay up.

6. Instead of adding a second NIC to your Service Console vSwitch, you may instead create a second Service Console interface (vswif1) on another vSwitch to provide redundancy. This can be done if you do not have extra pNICs to spare and can instead use an existing vSwitch. To do this, either create a new vSwitch or edit an existing vSwitch and add a Service Console network type. The name should default to Service Console 2. Then add a unique IP address for it and a subnet mask and click Finish. You should now have a second Service Console configured, as shown in Figure 4.39. The first Service Console is vswif0, and the second is vswif1.

Here's how to configure networking using the ESX Service Console CLI commands:

1. Log on to the ESX physical console as root or another user. (Use `su` to elevate privileges to root.)

By the Way

By default, remote access using the root account via Secure Shell (SSH) is disabled, and you can only log on using an alternate account. To enable root access via SSH (not recommended for security purposes), log on to the physical console as root and switch to the `/etc/ssh` directory. Then type `nano sshd_config` and change the line that says `PermitRootLogin` from `no` to `yes`. Press `Ctrl-O` to save the file, and then `Ctrl-X` to exit the nano editor. Finally, type `service sshd restart` and you will be able to log on using SSH with the root account.




Figure 4.39 Second Service Console interface configured on an alternate vSwitch

2. At the command prompt, type `esxcfg-nics -l` to see a list of all pNICs displayed. You should see output similar to that in Figure 4.40, which lists all the NICs that ESX can see and their names (vmnic#).

```
[root@esx1 root]# esxcfg-nics -l
Name      PCI      Driver      Link Speed  Duplex MTU      Description
vmnic1    04:02.00 tg3         Up    1000Mbps Full  1500    Broadcom Corporation NC778
1 Gigabit Server Adapter (PCI-X, 10,100,1000-T)
vmnic2    01:01.00 e1000      Up    1000Mbps Full  1500    Intel Corporation 82546EB
Gigabit Ethernet Controller (Copper)
vmnic0    01:02.00 tg3         Up    1000Mbps Full  1500    Broadcom Corporation NC778
1 Gigabit Server Adapter (PCI-X, 10,100,1000-T)
vmnic3    01:01.01 e1000      Up    1000Mbps Full  1500    Intel Corporation 82546EB
Gigabit Ethernet Controller (Copper)
```

Figure 4.40 Output of the `esxcfg-nics -l` command

3. To see your vSwitch configuration, type `esxcfg-vswitch -l`. You will see output similar to that in Figure 4.41, with the vSwitch configuration displayed followed by the port group configuration for the vSwitch.

```
[root@esx1 root]# esxcfg-vswitch -l
Switch Name  Num Ports  Used Ports  Configured Ports  MTU      Uplinks
vSwitch0    64         4           64               1500     vmnic0

PortGroup Name  VLAN ID  Used Ports  Uplinks
Virtual Machine Network0  0        0           vmnic0
Service Console  0        1           vmnic0
```

Figure 4.41 Output of the `esxcfg-vswitch -l` command

Did You Know?

All commands and names (including directory/file/option names) that you use when typing in the Service Console are case sensitive, and in many cases are lowercase. For example typing `Esxcfg-nics` rather than `esxcfg-nics` will display an error that the command is not found.

4. To add an additional pNIC (called an uplink) to vSwitch0, we can use the `esxcfg-vswitch` command. To see all the available options, type `esxcfg-vswitch` with no options. To add a pNIC, we will use the `-L (link)` option. Type the following to add an unused NIC to the vSwitch: `esxcfg-vswitch -L vmnic3 vswitch0`. If the command completes successfully, you should be brought back to the command prompt without any errors being displayed, as shown in Figure 4.42.

```
[root@esx1 root]# esxcfg-vswitch -L vmnic3 vSwitch0
[root@esx1 root]#
```

Figure 4.42 Adding a pNIC to a vSwitch using the `esxcfg-vswitch` command

5. To verify that the command completed successfully, type `esxcfg-vswitch -l` and look at the new vSwitch configuration. You should see in the Uplink column our new NIC listed, as shown in Figure 4.43.

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch0	64	5	64	1500	vmnic3,vmnic0
PortGroup Name	VLAN ID	Used Ports	Uplinks		
Virtual Machine Network0		0	vmnic0,vmnic3		
Service Console	0	1	vmnic0,vmnic3		

Figure 4.43 Output of the `esxcfg-vswitch -l` command showing the new uplink

6. Alternatively, to add an additional vSwitch and then add a second Service Console interface to it, we would first type `esxcfg-vswitch -a vSwitch1` to add a new vSwitch, and then `esxcfg-vswitch -L vmnic3 vSwitch1` to add a NIC to it, and finally `esxcfg-vswitch -A "Service Console 2" vSwitch1` to add a new port group for the second Service Console. (Lowercase `-a` adds a new vSwitch, and uppercase `-A` adds a port group.) You can next type `esxcfg-vswitch -l` to see your new vSwitch displayed, as shown in Figure 4.44.

```
[root@esx1 root]# esxcfg-vswitch -a vSwitch1
[root@esx1 root]# esxcfg-vswitch -L vmnic3 vSwitch1
[root@esx1 root]# esxcfg-vswitch -A "Service Console 2" vSwitch1
[root@esx1 root]# esxcfg-vswitch -l
```

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch0	64	4	64	1500	vmnic0
PortGroup Name	VLAN ID	Used Ports	Uplinks		
Virtual Machine Network0	0	0	vmnic0		
Service Console	0	1	vmnic0		
Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch1	64	3	64	1500	vmnic3
PortGroup Name	VLAN ID	Used Ports	Uplinks		
Service Console 2	0	0	vmnic3		

Figure 4.44 Adding a vSwitch for a second Service Console interface using the `esxcfg-switch` command

- Now that we have the new vSwitch created, it's time to use another command to add a Service Console interface to it. We'll use the `esxcfg-vswif` command, which is used to add, edit, and delete Service Console interfaces. The first Service Console interface is `vswif0`, so we'll add a `vswif1` now. Type `esxcfg-vswif -a vswif1 -p "Service Console 2" -i 172.20.11.69 -n 255.255.255.0` to add a new `vswif` interface, specify the port group name, and set an IP address and netmask for it. When you have finished, you can type `esxcfg-vswif -l` to verify that it was created, as shown in Figure 4.45.

```
[root@esx1 root]# esxcfg-vswif -a vswif1 -p "Service Console 2" -i 172.20.11.69 -n 255.255.255.0
[2008-10-30 05:28:02 'Vnic' warning] Generated New MAC address, 00:50:56:4d:41:62 for vswif1
Nothing to flush.
[root@esx1 root]# esxcfg-vswif -l
```

Name	Port Group	IP Address	Netmask	Broadcast	Enabled	DHCP
vswif0	Service Console	172.20.11.66	255.255.255.0	172.20.11.255	true	false
vswif1	Service Console 2	172.20.11.69	255.255.255.0	172.20.11.255	true	false

Figure 4.45 Adding a second Service Console interface using the `esxcfg-vswif` command

These are just a few of the many CLI commands that enable you to configure many aspects of your ESX hosts. Almost all the functionality of the CLI commands can also be done using the VI Client, but there are times when you may need to use the CLI instead, especially if you have networking problems. It's a good idea to familiarize yourself with all the commands and try them out so that you get experience using them. You can find more information about these commands in Chapter 11.

Configuring the ESXi Management Network

Because ESXi does not have a Service Console like ESX does, configuring its management network is a bit different. When you install ESXi, it does not prompt you to configure the networking for the management network; instead, it automatically configures it to use DHCP to

set an IP address for it. If you do not have a DHCP server available, you will have to configure networking using the ESXi console. If you do have DHCP enabled, your ESXi server will automatically obtain an IP address, and you can connect to it using the VI Client. Be warned, however, that if you make network changes using the VI Client you may get disconnected from the ESXi server and have to reconnect.

To configure the management network for ESXi using the console, follow these steps:

1. Press F2 and log on to the ESXi console, and then select the Configure Management Network option. A menu will display where you will have multiple options available. Network Adapters lets you select which NIC to be used for your management network. VLAN lets you set a VLAN ID (tag) to be used. IP Configuration lets you choose either DHCP or setting a static IP address, subnet mask, and gateway. DNS Configuration lets you choose to either use DHCP to obtain DNS servers or to manually set them. In addition, you can set the DNS name of the ESXi host. Custom DNS Suffixes lets you specify multiple DNS suffixes to be used.
2. When you have finished setting the information, press Esc to exit. You will be prompted to save the changes and restart the management network.

To configure the management network for ESXi using the VI Client, follow these steps:

1. Connect to the ESXi host using the VI Client and log on.
2. Select the Configuration tab, and then the Networking link. Here you can see and configure vSwitches for ESXi similar to ESX. The difference is the management network uses the VMkernel network, as shown in Figure 4.46, and you cannot configure a second management network like you can with ESX and the Service Console interface. However, you can add an additional NIC to the vSwitch that is used for the management network for redundancy, the same way as you would for ESX.




Figure 4.46 ESXi default network configuration

Configuring the VMkernel Network

The VMkernel network is a special network that is configured on your vSwitches to support VMotion traffic and connections to IP storage (iSCSI and NFS). In addition, for ESXi hosts, it

is used for its management network. If possible, try to isolate this network from other network traffic because VMotion traffic is sent in clear text (not encrypted) and to ensure other network traffic does not affect this critical network. This includes both isolating the VMkernel on its own vSwitch and physically isolating the pNICs on the vSwitch from the rest of the network. If you do isolate the NICs, make sure the network is configured so that all ESX hosts will have access to each other's VMkernel network. Also try to ensure that the vSwitch that the VMkernel is located on has network redundancy with two pNICs.

To configure the VMkernel network on ESX hosts, follow the steps that were listed in the “Configuring VMotion” section earlier in this chapter. ESXi hosts will usually already have this configured because it is part of the management network configuration.

Configuring vSwitches

After you have your host management network configured, it is time to configure additional vSwitches to be used for your VM networks.

vSwitch Settings

Let's take an in-depth look at a vSwitch and explore all the various settings and information that can be found inside it. If you select a host and then click the Configuration tab and choose the Networking option under Hardware and then select a vSwitch and click the Properties link, you will see a window displayed similar to the one in Figure 4.47. Each port group on a vSwitch can have its own custom settings.




Figure 4.47 vSwitch Properties window

The first tab is the Ports tab, where you configure the vSwitch settings and the ports and networks that are configured on it. When you select items in the list on the left, information about them is displayed on the right. If you select the first item in the Ports tab, which is the vSwitch, and click the Edit link at the bottom, you can edit the properties of the vSwitch. When you edit the vSwitch properties, a new window will be opened to the General tab, where you can change the number of ports on the vSwitch, as shown in Figure 4.48.




Figure 4.48 vSwitch properties: General tab

The number of ports a vSwitch has determines how many VMs can be connected to it. You can choose a number from the selection list. The values are all multiplies of 8, and are 8, 24, 56, 120, 248, 504, and 1,016. Do not set this higher than it needs to be, because doing so will cause the ESX host to use more memory. Also, note that other interfaces (Service Console, VMkernel) connected to the vSwitch will also use up ports on it. When you change the number of ports, you must restart the host for the change to take effect.

The next tab is the Security tab, which controls the security policies that can be enabled or disabled on the vSwitch, as shown in Figure 4.49.




Figure 4.49 vSwitch properties: Security tab

The first policy is for promiscuous mode, which controls what network traffic a guest virtual NIC can listen to. When this is set to Reject, a guest VM can listen to traffic on only its own MAC address of the port group that it is connected to. If this is set to Accept, a guest VM can listen to all traffic on a port group instead of just its own. It is best to leave this disabled unless you need to monitor the network via an intrusion detection/prevention system (IDS/IPS) device.

The next policy is for MAC address changes of the virtual NICs in a VM. When this policy is set to Reject, if a MAC address of a VM is changed through the operating system and does not match the one set in the VM's configuration (VMX) file, it drops all inbound traffic destined to the MAC set inside the operating system. When it is set to Accept, it allows inbound traffic to the MAC address set inside the operating system to pass. By default, this is set to Accept. There are usually not many cases when you would need to allow this, so consider changing it to Reject for maximum security. Note that this policy is required to be set to Accept for some clustering products to work properly.

The last policy is for forged transmits, which protect against outbound MAC address changes. When it is set to Reject, outbound traffic from a VM that does not match the MAC address of the VM's virtual NIC are dropped. When it is set to Accept, all outbound traffic is passed even if the MAC address is different from the one of the VM's NIC. By default, this is set to Accept. Again, there are usually not many cases when you would need to allow this, so consider changing it to Reject for maximum security. Note that this policy is required to be set to Accept for the Microsoft Network Load Balancing (NLB) product to work properly.

The next tab is Traffic Shaping, which controls the amount of network traffic that a vSwitch can pass through it, as shown in Figure 4.50.




Figure 4.50 vSwitch properties: Traffic Shaping tab

By default, this is disabled, and no traffic restrictions are in place on a vSwitch. If you enable it, you can set traffic shaping policies based on average bandwidth (measured over time), peak bandwidth (max bandwidth during a burst, must be equal to or larger than average bandwidth), and burst size (amount of data that can be sent in one burst). All the policies are set in kilobytes (KB). Consider enabling this if you want to limit outbound network traffic on a particular port group so that it does not affect network traffic on other port groups on a vSwitch.

The last tab is NIC Teaming, which controls how network traffic is distributed between NICs on a vSwitch and how to reroute traffic in the event of a NIC failure, as shown in Figure 4.51.




Figure 4.51 vSwitch properties: NIC Teaming tab

The first setting is for the NIC load-balancing policy, which controls how outbound traffic is distributed among the NICs assigned to the vSwitch. This setting does not affect inbound traffic (because that is controlled by the load-balancing policy of the physical switch that the vSwitch is connected to). You can select the following options for the load-balancing policy:

- **Route Based on the Originating Port ID (default).** Choose an uplink based on the virtual port where the traffic entered the vSwitch. With this option, traffic from a VM's virtual NIC is consistently sent using the same pNIC unless there is a failover to another NIC in the NIC team. Replies are received on the same pNIC as the physical switch learns the port association. This setting provides an even distribution of virtual NICs over pNICs if the number of virtual NICs is greater than the number of pNICs. Also, note that this setting places slightly less load on the ESX host than the MAC hash option.
- **Route Based on Source MAC Hash.** Choose an uplink based on a hash of the source Ethernet MAC address. This option is almost identical to the port ID option where a VM's virtual NIC is consistently sent using the same pNIC unless there is a failover to another NIC in the NIC team. Replies are received on the same pNIC as the physical

switch learns the port association. This setting provides an even distribution of traffic if the number of virtual NICs is greater than the number of pNICs. Because the hash is based on source MAC, it is possible that you could have idle pNICs, and therefore this is not a recommended load-balancing option.

- **Route Based on IP Hash.** Choose an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets such as NetBIOS or NetWare IPX, whatever is at those offsets is used to compute the hash. This option offers the best load balancing among pNICs but is dependent on the number of unique IP destinations. If there is a high number of unique IP destinations, this method will work much better than having a lower number of IP destinations. You must use link aggregation with this method, which uses multiple pNICs to create a larger pipe for a VM with a single virtual NIC. For a VM to achieve more bandwidth than is available on a single pNIC, you must use aggregation and have multiple IP conversations that hash to different pNICs on the vSwitch. This method relies on 802.3ad teaming (also known as EtherChannel), which must be configured on your physical switches.
- **Use Explicit Failover Order.** Always use the highest-order uplink from the list of active adapters that passes failover detection criteria. This option provides the least amount of load balancing on your vSwitch.

The default port ID load balancing works fine in most cases, but if you have vSwitches that are very busy and configured with more than three pNICs, you should consider using the IP hash option instead, which provides the most balanced load balancing of the available options.

The next setting is for Network Failover Detection, which specifies the method that is used for pNIC failure detection. The two options available are as follows:

- **Link Status Only (default).** This method relies only on the link status that is provided by the NIC and will detect such things as disconnected cables and physical switch port failures. This method will not typically detect things such as physical switch configuration errors or misconfigurations.
- **Beacon Probing.** This method is more reliable than link status and works by sending out and listening for beacon probes on every NIC in the team on a vSwitch and additionally uses link status to detect failures. Beacon probing packets are 62 bytes in size and sent every ten seconds.

You should consider using beacon probing instead of link status because it provides additional detection methods. However, if you have a Cisco switch that supports link-state tracking, it is best to use that instead of beacon probing because it allows the Cisco switch to detect failures and automatically disable the port and is a more advanced method providing

better protection. You can read more about how to set up this method in VMware knowledge base article 10098 (<http://kb.vmware.com/kb/10098>).

The next setting is Notify Switches. When set to Yes (default), a notification is sent out over the network to update the lookup (RARP) tables on physical switches whenever a virtual NIC's traffic would be routed over a different pNIC in a vSwitch because of a NIC failure. You will almost always want to leave this setting set to Yes unless you are using Microsoft Network Load Balancing in unicast mode on your VMs.

The last setting is Failback. When set to Yes (default), any active pNIC in a vSwitch that fails will be returned to active duty as soon as the failure has been corrected and thereby replace any standby adapters. When set to No, a pNIC that fails is left in standby mode until another active pNIC fails, which it then replaces. In most cases, you will also want to leave this set to Yes.

The bottom section of the NIC Teaming tab is for NIC Failover Order and is used to set which pNICs will be active and which will be used as standby NICs to take over in case of a failover. In most cases, you will want to have all your pNICs active in a vSwitch to provide maximum bandwidth and performance. However, in some instances you might want to specify standby NICs, which can also be active for other port groups in a vSwitch (for example, if you have the Service Console and VMkernel both set up on a single vSwitch with two pNICs to provide redundancy). Because it would be best to separate the traffic for each onto its own NIC, you could edit the port group for each and specify that NIC1 be active for the Service Console and standby for the VMkernel and NIC2 be active for the VMkernel and standby for the Service Console. If a NIC failure occurs, the other NIC becomes active for both the Service Console and VMkernel. This way you have both redundancy and isolated traffic between the two. You can use the Move Up and Move Down buttons to move NICs and specify which are active and standby.

That's it for the Ports tab on your vSwitch settings. The other tab is the Network Adapters tab, which lists all the pNICs configured for that vSwitch, as shown in Figure 4.52.

The left pane in this tab displays all the pNICs in the vSwitch. When you select one, it displays information about the NIC in the right pane, which includes the physical slot location, driver name, link status, configured and actual speed/duplex, and the networks (VLANs) that the NIC can see. You can use the buttons at the bottom of the window to add or remove a pNIC to the vSwitch. In addition, you can edit a NIC and change its configured speed and duplex (for example, Auto/Auto, 100/Full). You should make sure that whatever speed and duplex that you configure for your NICs on your host server match what the physical switch port that it is connected to is set to. This is typically Auto/Auto for gigabit NICs and 100/Full for 100Mbps NICs. Check with your network team to see what the switch is set to so that you can configure your NICs the same; if you fail to do so, you may experience network performance problems if they are set differently.




Figure 4.52 vSwitch properties: Network Adapters tab

Sample vSwitch Configurations

Let's look at some sample vSwitch configurations that you might use based on the number of NICs in your host server and some different requirements.

Scenario 1: Two pNICs in your host, one VLAN needed for VMs, and non-IP-based shared storage

- **Sample configuration 1.** Provides isolation of Service Console/VMkernel traffic from VM traffic but no redundancy
 - ◆ NIC1, vSwitch0, Service Console and VMkernel port groups
 - ◆ NIC2, vSwitch1, single VM port group
- **Sample configuration 2.** Provides no isolation of Service Console/VMkernel traffic from VM traffic but does provide redundancy
 - ◆ NIC1 and NIC2, vSwitch0, Service Console, VMkernel and VM port groups

Scenario 2: Four pNICs in your host, multiple VLANs needed for VMs, and non-IP-based shared storage

- **Sample configuration 1.** Provides no redundancy for the Service Console/VMkernel traffic, but does separate them on different vSwitches; provides redundancy for your VM vSwitch
 - ◆ NIC1, vSwitch0, Service Console
 - ◆ NIC2, vSwitch1, VMkernel
 - ◆ NIC3 and NIC4, vSwitch2, VM port groups
- **Sample configuration 2.** Provides redundancy for Service Console/VMkernel by using active/standby NICs; provides redundancy for your VM vSwitch
 - ◆ NIC1, vSwitch0, Service Console (active), VMkernel (standby)
 - ◆ NIC2, vSwitch0, VMkernel (active), Service Console (standby)
 - ◆ NIC3 and NIC4, vSwitch1, VM port groups

Scenario 3: Six pNICs in your host, multiple VLANs needed for VMs (including an external DMZ), and non-IP-based shared storage

- **Sample configuration 1.** Provides no redundancy for the Service Console/VMkernel traffic and isolates them on different vSwitches; provides redundancy for your VM vSwitch and provides redundancy for a second vSwitch to isolate DMZ traffic on for your VMs
 - ◆ NIC1, vSwitch0, Service Console
 - ◆ NIC2, vSwitch1, VMkernel
 - ◆ NIC3 and NIC4, vSwitch1, VM port groups (internal)
 - ◆ NIC5 and NIC6, vSwitch2, VM port groups (external DMZ)
- **Sample configuration 2.** Provides redundancy for Service Console/VMkernel by using active/standby NICs; provides redundancy for your VM vSwitch and provides redundancy for a second vSwitch to isolate DMZ traffic on for your VMs
 - ◆ NIC1, vSwitch0, Service Console (active), VMkernel (standby)
 - ◆ NIC2, vSwitch0, VMkernel (active), Service Console (standby)
 - ◆ NIC3 and NIC4, vSwitch1, VM port groups (internal)
 - ◆ NIC5 and NIC6., vSwitch2, VM port groups (external DMZ)

Scenario 4: Six pNICs in your host, multiple VLANs needed for VMs, and using IP-based shared storage

- **Sample configuration 1.** Provides redundancy for the Service Console/VMkernel traffic and isolates them on different vSwitches; provides redundancy for your VM vSwitch
 - ◆ NIC1 and NIC2, vSwitch0, Service Console
 - ◆ NIC3 and NIC4, vSwitch1, VMkernel (IP storage)
 - ◆ NIC5 and NIC6, vSwitch2, VM port groups

- **Sample configuration 2.** Provides no redundancy for Service Console but does provide redundancy for the VMkernel; provides redundancy and extra bandwidth for your VM vSwitch
 - ◆ NIC1, vSwitch0, Service Console
 - ◆ NIC2 and NIC3, vSwitch1, VMkernel (IP storage)
 - ◆ NIC4, NIC5, and NIC6, vSwitch2, VM port groups (internal)

Additional Networking Considerations and Tips

Here are some additional things you should know about VI3 networking.

Configuring Cisco Discovery Protocol

The Cisco Discovery Protocol (CDP) provides information for VI3 administrators on the physical Cisco switch that it is connected to. In addition, it provides network administrators with information about the vSwitch that is connected to the physical switch. There are multiple mode settings for CDP, and they can be set only by using the ESX Service Console `esxcfg-vswitch` command or the ESXi `vicfg-vswitch.pl` command. To check the current mode for CDP on a vSwitch, type `esxcfg-vswitch -b <vswitch name>`. To change the mode, type `esxcfg-vswitch -B <mode> <vswitch name>`. The available mode options are as follows:

- **Down.** CDP is disabled.
- **Listen (default).** Information is detected about the physical switch port and displayed in the VI Client, but information about the vSwitch is not available to network administrators.
- **Advertise.** Information about the vSwitch is made available to network administrators, but information is not detected about the physical switch port and is not displayed in the VI Client.
- **Both.** Information is detected about the physical switch port and displayed in the VI Client, and information about the vSwitch is also made available to network administrators.

To view CDP information, go to Configuration, Networking, and then click the little dialog icon to the right of the Physical Adapters section of your vSwitch.

How to Map a VM NIC to a Physical NIC

There is no easy method for being able to tell which `vmnic#` (the name that ESX gives a pNIC) is mapped to which pNIC. Sometimes you need to know this information when configuring or troubleshooting your networking on your hosts. For example, you may have `vmnic1` to `vmnic6` in your host and you need to know which pNICs they are mapped to so that you can configure them properly. Here are two ways to find this out:

- The easiest method is to unplug the NIC's cable temporarily and log on to the Service Console and type `esxcfg-nics -l` and see which NIC shows its link status as down.
- Another method is to temporarily change the speed/duplex of a NIC by logging on to the Service Console and typing `esxcfg-nics -s 10 -d half`, which sets a NIC to 10/half. You can then have your network administrators check the port status to find the NIC that is running at 10/half. When you have finished, you can change it back by using the `-a` option to set it to `auto/auto` or enter the appropriate values for `-s` and `-d`.

Internal-Only vSwitches

vSwitches that have no pNICs assigned to them are like having an isolated, internal-only switch that is completely isolated from the network and any other vSwitches on an ESX host. These types of vSwitches are very useful for isolating VMs that you do not want communicating on your regular network. Some examples of where you may use this are as follows:

- Building a new VM and isolating it until you get it fully patched and secure
- Penetration testing or running security scans on a VM
- Cloning an existing VM to troubleshoot it so that the original server and clone can be running at the same time and will be isolated from one another
- Performing a P2V conversion of a physical server so that the physical server and VM can be running with the identical IP address at the same time until the VM is ready and the physical server can be shut down
- Building an isolated test environment that has no network connectivity to other servers

Because of the usefulness of internal vSwitches, it's always a good idea to configure one on every ESX server. Just be aware that a VM connected to an internal vSwitch cannot be VMotioned to another host. You can also configure port groups and VLAN tags so that you can set up an isolated environment using VLANs with multiple servers in them. After you have them configured, it is easy to switch a VM to another vSwitch by editing the VM's settings and changing the network in the NIC properties. There is no need to power off the VM to do this. To create an internal vSwitch, follow these steps:

1. Select your ESX host in the VI Client, and on the Configuration tab select Networking.
2. Click the Add Networking link.
3. Select Virtual Machine as the connection type.
4. Select Create a VSwitch and do not select any network adapters.
5. Optionally, assign a VLAN ID. If you need additional VLANs, create more port groups.

Changing a vNIC MAC Address

Sometimes you might need to change a MAC address of the virtual NIC in your VM because of software that is licensed by the MAC address of the server it is running on or if you migrate a physical server to virtual and want to keep the same MAC address. A VM's MAC address will change if the location of the VM changes (for example, different path on the same host). During a hot migration (VMotion), the location of the VM does not change. This is also true of a VM that changes hosts because of HA/DRS migration. During a cold migration, the location of the VM does change, so the MAC address of the VM will change.

MAC addresses are hex values and consist of six groups of hex numbers. The first 3 octets of the MAC address comprise a unique code assigned to each NIC vendor. This is also called the organizationally unique identifier (OUI). The VMware OUI used for VI3 VMs is 00:50:56. The last 3 octets are assigned to all the NICs for that vendor. Each MAC address must be unique to avoid conflicts with other network devices. You do have to use the VMware assigned range (00:50:56) for your NICs when changing the MAC address using the VI Client or by editing the VMX file of the VM. However, if you change the MAC address through your guest operating system, you can use any valid MAC address that you want. This is handy if you are coming from a physical server and you want to use a MAC address using the previous NIC's first 3 octet range instead (for example, HP/Compaq NICs use 00:08:02). Just make sure the pNIC and virtual NIC with the same MAC address are not active on the same physical network at the same time. There are several ways to set a static MAC address on a VM, which are listed here. A new feature in ESX 3.5 enables you to change a MAC address of a VM NIC using the VI Client, so there is no need to edit the VM's configuration (VMX) file to do this. If you use the two methods listed first, you must stay inside VMware's allowed MAC addresses (00:50:56:00:00:00 to 00:50:56:3F:FF:FF) or the VM will not start.

Here's how to change the MAC address by editing the VMX file of the VM:

1. Edit the VMX file of the VM using a text editor such as nano or vi or by opening it using a file manager such as WinSCP. The VMX file is usually located in the VM's working directory on the VMFS volume that it is on.
2. Change the line `ethernetN.addressType="vpx"` to `ethernetN.addressType="static"`. (*N* is the number of your Ethernet adapter, usually 0.)
3. Change the line `ethernetN.GeneratedAddress` to `ethernetN.address`, and then change the current MAC address to `00:50:56:XX:YY:ZZ`. (Again, *N* is the number of your Ethernet adapter, and *XX* is a valid hex number between 00 and 3F, and *YY* and *ZZ* are valid hex numbers between 00 and FF.) The value for *XX* must not be greater than 3F to avoid conflict with MAC addresses that are generated by the VMware Workstation and VMware Server products.
4. Power your VM back on. Log on to the OS, go to the CMD prompt, and type `ipconfig /all`. Your manually assigned MAC address should be listed for the NIC that you changed.

Here's how to change the MAC address using the VI Client (ESX version 3.5 or later only):

1. Power off the VM that you want to change the MAC on.
2. Edit the settings of the VM and select the NIC that you want to change.
3. In the right pane, change the MAC address from automatic to manual and enter a MAC address in the text field, as shown in Figure 4.53. The MAC address must be in the range of 00:50:56:00:00:00 to 00:50:56:3f:ff:ff or you will receive an error message.




Figure 4.53 Manually setting a MAC address using the VI Client

4. After you enter your MAC address, click OK to save it, and the VMX file will be updated.

Here's how to change the MAC address by setting the NIC properties in Windows:

1. Edit the local area connection properties for the NIC you want to change.
2. Click the Configure button next to the NIC name.
3. On the Advanced tab, select NetworkAddress.
4. In the Value field, enter a new value for the MAC address; enter only numbers or letters, no spaces, dashes, or colons. This MAC address can be any valid hex numbers between 00 and FF for any of the octets.
5. Click OK, and that's it. The new MAC address takes effect immediately and will override any MAC address set by VMware.

Adding Additional NICs to an ESX Host

When adding, changing, or removing hardware (for example, NICs, fiber cards) to your ESX host, you may come across a situation where ESX rennumbers your NICs, which can cause your networking to no longer work. This can also happen by simply replacing a defective card with the same model card. What happens is that the BIOS in your server will sometimes change the physical topology of its hardware devices, which can subsequently change the order in which devices are discovered. When this occurs, ESX will see the NICs as new devices and assign them new NIC numbers (for example, vmnic#).

An example of this is if you replace a two-port NIC with a four-port NIC. The original NIC numbering is typically vmnic0 and vmnic1. When ESX is started after the NIC has been

changed, it will see the new NIC and number them as vmnic2 to vmnic5. This happens because even though vmnic0 and vmnic1 are no longer physically present they still exist in the ESX configuration. Because your current vSwitch configuration relies on vmnic0 and vmnic1, your networking will no longer work when the ESX host starts up. You will also not be able to connect to the Service Console remotely via SSH or the VI Client.

There are two ways to recover from this when it happens. The first method is easier and saves you from changing your vSwitch configuration but involves editing the `esx.conf` file (so you need to be careful). In addition, the first method preserves your NIC numbering if you do not want your NICs starting at vmnic2.

Here's how to recover from this situation by editing the `esx.conf` file:

1. Log on to the Service Console.
2. Check your existing NIC numbering by typing `esxcfg-nics -l`.
3. Type `cd /etc/vmware` to change to the correct directory.
4. Type `cp esx.conf esx.conf.bak` to make a backup of this file (because it is a critical configuration file for ESX).
5. Type `nano esx.conf` to open the file for editing.
6. Press Ctrl-W, and then enter vmnic2 to search for the new first NIC.
7. Change vmnic2 to vmnic0.
8. Change the subsequent NICs from vmnic3 to vmnic1, vmnic4 to vmnic2, and vmnic5 to vmnic3.
9. Press Ctrl-O to save the file.
10. Press Ctrl-X to exit the nano editor.
11. Shut down and restart the ESX server. When the server comes back up, the NICs will be numbered vmnic0 to vmnic3. Verify this by typing `esxcfg-nics -l`.

Here's how to recover from this situation by modifying your vSwitch configuration:

1. Log on to the Service Console.
2. Check your existing NIC numbering by typing `esxcfg-nics -l`.
3. Check your current vSwitch configuration by typing `esxcfg-vswitch -l`. Note which NICs are assigned to which vSwitches (Uplink column).
4. Remove the old NICs that have been renamed by typing `esxcfg-vswitch -U vmnic# vswitch name` (for example, `esxcfg-vswitch -U vmnic0 vswitch1`).
5. Add the new NICs with the correct names by typing `esxcfg-vswitch -L vmnic# vswitch name` (for example, `esxcfg-vswitch -L vmnic2 vswitch1`).
6. Repeat this process for any additional NICs. After you have the vSwitch that contains the Service Console corrected, you can also log in via the VI Client and correct the other vSwitches that way.

7. Your newly renamed NICs should now be assigned to the original vSwitches and your networking should now work again.

How Traffic Routes on vSwitches

A vSwitch is basically just software that is contained in the memory of the host servers that connect VMs with pNICs. You might wonder how the network traffic between two VMs that are both located on the same host server routes and whether it actually leaves the host server and goes out over the network. Here are a few scenarios that cover how network traffic is routed in different situations between two VMs on the same host server:

- **Different vSwitches, same port group and VLAN.** VM A is connected to vSwitch1 and VM B is connected to vSwitch2. In this example, the VMs are plugged into separate vSwitches on the same host server. Network traffic between VM A and VM B goes from a pNIC on vSwitch1 to a physical switch that it is connected to and then back to a pNIC on vSwitch2 and then to VM B, as shown in Figure 4.54.




Figure 4.54 Network traffic routing example with different vSwitches, same port group and VLAN

- **Same vSwitch, different port group and VLAN.** VM A is connected to vSwitch1, port group 20, and VM B is connected to vSwitch1, port group 17. In this example, the VMs are plugged into the same vSwitch on the same host server. Network traffic between VM A and VM B goes from a pNIC on vSwitch1 to a physical switch that it is connected to and then back to a pNIC on vSwitch1 and then to VM B, as shown in Figure 4.55.




Figure 4.55 Network traffic routing example with same vSwitch, different port group and VLAN

- **Same vSwitch, same port group and VLAN.** VM A is connected to vSwitch1, port group 20, and VM B is connected to vSwitch1, port group 20. In this example, the VMs are plugged into the same vSwitch and the same port group on the same host server. Network traffic between VM A and VM B never leaves the host server and does not go to the pNICs on the host server and therefore never travels on the physical network, as shown in Figure 4.56.




Figure 4.56 Network traffic routing example with same vSwitch, same port group and VLAN

Because network traffic between VMs on the same host, same vSwitch, and same port group does not leave the host, it can be advantageous to configure VMs that have a lot of network traffic between them in this manner (for example a web server and an application server or an application server and a database server). Doing this will result in increased network speed and reduced network latency between the VMs. If you use DRS, you might also consider creating a rule to ensure that the VMs stay on the same host.

Understanding Network Hints

Network hints, also called observed IP ranges, are useful for seeing which VLANs are configured on the physical switch that a NIC is connected to. This is particularly useful when using VLAN tagging, where many VLANs are configured for a single physical switch port. They are not 100% accurate, but are a best guess on the traffic seen so far by the pNIC, which is why they are called hints and are not authoritative. Some of the IP ranges for certain VLANs may be missing from the list if there has not been any traffic on that VLAN yet. They can be seen in the VI Client by selecting the Configuration tab and then choosing Network Adapters. (You can also see them by selecting the properties of a vSwitch and clicking the Network Adapters tab.) A pop-up window will display when you place your cursor in the Observed IP Ranges column of a NIC in the VI Client, as shown in Figure 4.57. You can also see them by increasing the size of the Observed IP Ranges column, but it is more difficult that way.

A more accurate way to do this via the Service Console is with the `esxcfg-info` command. The VLANs and networks are listed under the network hint for each adapter. You can type `esxcfg-info -n | grep -E -i "_name|Hint"` to see a list of all the network hints for each adapter.

Network Adapters					
Device	Speed	Configured	vSwitch	Observed IP ranges	Wake on LAN Supported
Broadcom NetXtreme II BCM5708 1000Base-T					
vmnic1	1000 Full	Negotiate	vSwitch3		Yes
vmnic0	100 Full	Negotiate	vSwitch0	172.20.11.16-172.20.11...	Yes
Intel PRO/1000 PT Low Profile Quad Port PCIe Gigabit Ethernet Controller					
vmnic5	100 Full	100 Full	vSwitch4		No
vmnic4	1000 Full	Negotiate	vSwitch1	172.20.1.60-172.20.1.60...	No
vmnic3	1000 Full	Negotiate	vSwitch1	172.20.1.60-172.20.1.60...	No
vmnic2	1000 Full	Negotiate	vSwitch1	172.20.1.60-172.20.1.60... 172.20.11.16-172.20.11.23 (VLAN 11), 172.20.20.40-172.20.20.43 (VLAN 20), 172.20.3.32-172.20.3.63 (VLAN 3), 172.20.8.1-172.20.8.127 (VLAN 8), 172.20.12.20-172.20.12.20 (VLAN 12), 172.20.4.11-172.20.4.11 (VLAN 4), 172.20.6.1-172.20.6.7 (VLAN 6),	

Figure 4.57 Observed IP ranges displayed in the VI Client

Configuring Storage

Storage is one of the configuration tasks that you want to make sure you get right the first time. It can be difficult and time-consuming to change storage configurations later after they have been configured. If you make the wrong choices when configuring storage, it can affect the performance of your hosts and VMs. So take some time to understand your needs and requirements and ensure that you thoroughly understand your options before configuring storage on your hosts.

LUN Size Considerations

This is always a common question among those new to VMware: How big should I make my LUNs to use with my ESX hosts. ESX supports using up to 2TB LUNs, but there is no magic number when it comes to selecting an optimal LUN size. However, it is generally recommended to keep them around 500 to 600GB in size. Having many, smaller LUNs is more difficult to manage and limits the size of your VMFS volumes. It is possible to put multiple LUNs together to make one VMFS volume using extents, but it is not recommended because a LUN failure at the root extent can take down the whole VMFS volume. Also, VMs can end up spanning LUNs when using extents, which can cause performance problems. You also do not want to have too many VMs running on a single LUN. In general, you should plan to have about 14 to 16 VMs per LUN. Having too many VMs on a single LUN can cause problems with metadata locking issues (SCSI reservations). You should consider fewer VMs per LUN (8–10) if you will have very disk I/O intensive applications running on them, and you should consider more VMs per LUN (20–22) if you have low disk I/O applications running on them. In some cases, you might put more VMs on a LUN, such as VDI implementations where you have many low disk I/O virtual desktops running on your host servers.

It is best to look at your disk storage requirements per VM and then allow for additional storage for things such as VSWP, log, and snapshot files that will also be created in your VM's working directory. How much extra storage you will need for your VMs will be heavily influenced by how much you plan to use snapshots with your VMs. A general rule of thumb is to allow for about 25% additional disk space in addition to the size of the VM's virtual disks. For example, if you have a VM with a 20GB disk, you should allow for at least 4GB of additional space. Here are some things to consider about the extra files that your VMs will have:

- A VSWP file will be created equal in size to the amount of memory assigned to a VM minus any memory reservations. So a VM with 4GB of RAM and no memory reservation will have a 4GB VSWP file created on your VMFS volume.
- Snapshots will vary in size based on the number of changes made to a VM's disk after the snapshot is created. They are initially small (16MB) and can grow up to the total size of the original VM disk (but it cannot exceed it). How much and how fast they change depends on your applications. Static servers like web and app servers will grow slowly, dynamic servers like database and email servers will grow faster. So how much space you should allow for snapshots will depend on how many snapshots you plan on taking and how long you need to keep them. You should also allow for additional disk space so that you can delete multilevel snapshots, which require extra disk space to commit. If you plan to use a lot of snapshots, allow for plenty of extra disk space on your VMFS volumes.
- In addition, if you plan to save memory state with your snapshots (VMSN files), this will require extra disk space equal to the size of the amount of memory a VM has for each one.
- VM log files do not take up that much room and generally total fewer than 50MB of disk space.
- If you suspend a VM, a VMSS file is created that holds the contents of the VM's memory and is equal in size to the amount of memory a VM has. You typically don't have to allow for extra space for these because the VSWP file is deleted when the VMSS file is created, so they balance each other out.

When you add up the amount of disk space your 12 to 14 VMs will need based on these considerations, you can calculate the size of the LUNs that will work best for you. You can then create your VMFS volumes to be equal in size to that of your LUNs. If you do have a need for VMs with very large virtual disks (greater than 1TB), consider using RDMs rather than VMFS volumes, which we talk about in a moment.

Another consideration when you create your VMFS volumes is the maximum size of the virtual disks that you will be assigning to your VMs. By default, VMFS volumes are created with a 1MB block size, which allows for a single virtual disk (vmdk) to be created up to 256GB. If you try to create a virtual disk larger than 256GB, it will not allow you to. Once you

set a block size on a VMFS volume, you cannot change it later. Instead, you need to move all the VMs from the volume and delete it and re-create it with a new block size. Therefore, confirm you choose a block size that works for you based on your current and future needs. The block size choices and related maximum virtual disk sizes are listed here:

- **1MB block size.** 256GB maximum virtual disk file size (default)
- **2MB block size.** 512GB maximum virtual disk file size
- **4MB block size.** 1024GB maximum virtual disk file size
- **8MB block size.** 2048GB maximum virtual disk file size

Choosing a larger block size will not impact disk performance and will only affect the minimum amount of disk space that files will take up on your VMFS volumes. Block size is the amount of space that a single block of data takes up on the disk, and subsequently the amount of disk space a file takes up will be based on a multiple of the block size. For example, with a 1MB block size, a 16KB file will take up 1MB of actual disk space, and a 2.1MB file will take up 3MB of disk space, and a 50MB file will take up 50MB of disk space. With an 8MB block size, a 16KB file will take up 8MB, and a 2.1MB file will also take up 8MB, and a 50MB file will take up 56MB of disk space. However, VMFS does use subblock allocation, so small files do not take up an entire block.

So as you can see, there is some wasted disk space that becomes even larger when using larger block sizes. Smaller files are affected more than larger files are. However, this is not really a problem on a VMFS volume because there are not many files on them and most of the files are all very large and are not affected all that much by having a larger block size. So if you need to increase your block size because you think you might use larger virtual disks, by all means go ahead and do it.

Choosing Between Using a VMFS Volume and Raw Device Mappings

When you add a virtual disk to a VM, you have the option of using a traditional virtual disk on a VMFS volume or using a raw device mapping (RDM) instead. Let's talk about the difference between the two and why you might want to choose one type over another. A RDM is a mapping from a VM directly to a raw LUN on a SAN, and the VM's disk data file does not reside on a VMFS volume. A small disk descriptor file is still created for the virtual disk in the VM's working directory on the VMFS volume, however. A RDM provides a VM direct access to a LUN, which provides some additional benefits over VMFS volumes.

The performance characteristics between an RDM and a VMFS volume are very similar, and there is only a slight difference between them in certain situations. For random workloads, they both produce similar I/O throughput. For sequential workloads with small I/O block sizes, RDMs provide a small increase in throughput compared to VMFS. For all workloads, RDMs have slightly better CPU costs because of the filename metadata that VMFS maintains. Both VMFS and RDM provide many of the same features, like file locking and VMotion.

RDMs can be configured in two different modes: virtual compatibility mode and physical compatibility mode. Virtual compatibility mode virtualizes the mapped device and is mostly transparent to the guest operating system. This mode also provides some advantages of VMFS volumes, like the ability to create snapshots. Physical compatibility mode provides minimal SCSI virtualization of the mapped device, and the VMkernel passes almost all SCSI commands directly to the device, thereby allowing closer integration between the VM and the LUN.

Here are some reasons why you might choose to use an RDM rather than a VMFS volume:

- RDMs can easily be grown in place to increase their size via SAN utilities. Growing a vmdk file on a VMFS volume is a bit more difficult and requires more steps.
- Useful for advanced SAN functions and when running a SAN-aware application inside a VM.
- If you plan to use Microsoft Cluster Services on your VMs, you will need to use an RDM if they are going to be on separate hosts.
- If you need to create a virtual disk for a VM that is larger than 1TB, consider using an RDM (just a best practice, not a limitation).
- RDMs are formatted by the guest operating system and use whatever file system the guest chooses (for example, NTFS, ext3, reiserfs).
- RDMs can be easily disconnected from a VM and connected to another VM or physical server.

Here are some reasons why you might not choose to use an RDM rather than a VMFS volume:

- They can be a bit more difficult to manage than VMFS volumes. RDMs cannot be seen in the VI Client datastore list.
- Physical mode RDMs cannot use snapshots, which can cause problems for backup applications (VCB) that snapshot a VM before backing it up. This can be advantageous, however, if you want to snapshot a VM with multiple disks and do not want the RDMs to be included in the snapshot.
- If you cold migrate a VM and change the disk location of the VM, any RDMs configured on the VM will be converted to VMFS volumes as part of the migration. To prevent this, you need to temporarily remove the RDMs before you cold migrate and re-add them after the VM has been migrated.
- You must assign a whole LUN to be used as an RDM for a virtual disk. You cannot use part of a LUN for this.

In general, it is recommended to use VMFS volumes for your VM's virtual disks unless you are using MSCS, plan to use disks larger than 1TB, or need to use SAN functionality with

an application. If you do choose to use an RDM on a VM, it's a good idea to create a small virtual disk file on a VMFS volume for the VM's operating system and then add one or more RDMs to the VM as additional virtual disks. To create an RDM, follow these steps:

1. Ensure that your host can see the LUN that you want to use. (Perform a LUN rescan if necessary.)
2. Add a new virtual disk to a VM, and at the Select a Disk screen choose the Raw Device Mapping volume option and click Next.
3. At the next screen, you will be presented with a list of unused LUNs that you can select to use as your target LUN.
4. Select a location for the mapping file that is created for the RDM. The default is to store it with the VM, but you can choose an alternative location. If you do choose an alternative location, make sure it is on a shared volume that all your hosts can see so that features such as DRS and VMotion will work.
5. Select a mode for the RDM, either virtual or physical compatibility. (Remember that snapshots on the RDM will not work in physical mode.)
6. At the Advanced Options screen, you can select the SCSI ID that you want the RDM to appear as to the VM.
7. Click Finish and your virtual disk will be created.

Configuring Local Storage

In most cases, you will have local storage on your host servers (unless you are using boot from SAN), which you can use to create nonshared VMFS volumes on your hosts. Even if you plan to use only shared storage for your VMs, you should still configure local VMFS volumes because they can be useful in many situations.

To create a local VMFS volume, follow these steps:

1. On your host server, select the Configuration tab and then Storage Adapters. Here you can select your local storage adapter and see the target disks that it has available.
2. Click the Configuration tab, then Storage, and then click the Add Storage link. Select Disk/LUN as your storage type, and click Next to continue.
3. At the Device Location screen, choose a device from the list of unused devices that are displayed, and click Next to continue.
4. At the next screen, the current disk layout is displayed. If it is a new device, it should show as blank. Click Next to continue.
5. At the Properties screen, give your new datastore a name. This is the friendly name (actual name is the UUID that is automatically assigned) that will be associated with the device and can be easily renamed later.

6. At the Formatting screen, choose a block size (as discussed earlier) and select whether to use the maximum capacity of the drive (in most cases, you will want to), and click Next to continue.
7. At the Ready to Complete screen, review the information. Click Back if you need to change anything, and when ready click Finish to create the VMFS volume.

Did You Know?

When you create a VMFS volume, you are prompted for a name for it. This name is not what the Service Console uses to reference the volume; it is purely to make it easier for the user to identify the volume. The Service Console actually uses a unique identifier called a UUID to reference the volume. The name you specify when you create it is a user-defined device name, which is a symbolic link to the UUID. This is done to solve the problem of changing the device name. When you change the name, you are only changing the user-defined device name and not the UUID of the volume. So when you look in your `/vmfs/volumes` directory, you will see both a UUID (for example, 4404e8b4-bcfd52fc-1e4b-0017a4a91076) and the symbolic link (for example, `ServerA-Local`). Clicking the symbolic link simply takes you to the UUID directory. It is a lot easier to remember this than the long UUID of a volume.

Configuring Fibre Channel Storage

Configuring Fibre Channel (FC) storage can be complicated and usually involves working with your SAN administrator to get everything set up for your host servers. Proper preparation is the key for proper configuration, and you should work closely with your SAN administrators to ensure that they understand your needs and configure the LUNs that you will use properly. VMware has created two very large white papers dedicated to SAN configuration, and I recommend that both you and your SAN administrators read through both before proceeding. The first white paper is titled “Fibre Channel Configuration Guide” and is available at www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_san_cfg.pdf. The second white paper is called “SAN System Design and Deployment Guide” and is available at www.vmware.com/pdf/vi3_san_design_deploy.pdf. When you are ready to configure FC storage on your host servers, follow these steps:

1. Design and create your SAN LUNs. When you present SAN LUNs to ESX hosts, ensure that you present each LUN as the same LUN ID to every ESX host. This is essential because every ESX host needs to see each LUN the same to work properly and for features such as DRS, HA, and VMotion to work. Many SAN administrators may be resistant to setting up the LUNs this way because they might think that if multiple hosts see the same LUNs it could cause data corruption. This is not the

case with ESX, however. By design, the VMFS volume system is made to work this way and prevents multiple hosts from writing to the same file at the same time with its unique file locking technique. It is helpful if you explain that from a storage perspective you can consider an ESX cluster to be a single host with multiple HBAs. This will often help storage administrators understand how to configure things correctly.

2. Set up the FC HBAs on your ESX hosts. Once your HBAs are connected to the FC switch, you can perform a rescan to look for new LUNs. Sometimes, for it to see the LUNs, you might have to reboot the host server if the LUNs were presented after the host was started. If you are using more than one FC adapter in your host server for redundancy, make sure both your adapters can see all the same LUNs and that all your paths to the SAN are showing up. You can use the ESX Service Console command `esxcfg -mpath` to see and verify your paths. Select your ESX host, and on the Configuration tab select Storage Adapters. Next select one of your storage adapters and you should see all the LUNs it sees displayed, as shown in Figure 4.58. You may see the same LUNs displayed twice because of the multipaths that are set up for redundancy. You'll notice the different path numbers under each target. For example, the same LUN may show up as `vmhba1:1:8` and as `vmhba1:0:8` (8 is the LUN number). You can click the Rescan link to initiate a rescan of new storage devices and VMFS volumes.

The screenshot shows the 'Storage Adapters' configuration window in the ESX Service Console. The 'Details' tab is selected for the 'vmhba1' adapter. The adapter is a QLA231x/2340 Fibre Channel adapter with WWPN 21:00:00:e0:8b:8f:73:1d and 4 targets. Below, two SCSI targets are listed, each with five paths. The paths are labeled with their canonical paths, type (disk), capacity (500.00 GB), and LUN ID (8, 9, 10, 11, 12, 13).

Storage Adapters					
Device	Type	SAN Identifier			
QLA231x/2340					
vmhba1	Fibre Channel	21:00:00:e0:8b:8f:73:1d			
vmhba3	Fibre Channel	21:00:00:e0:8b:8f:e3:1d			
Smart Array 6i					
vmhba0	Block SCSI				
iSCSI Software Adapter					
iSCSI Software Adapter	iSCSI				

Details					
vmhba1					
Model: QLA231x/2340					
WWPN: 21:00:00:e0:8b:8f:73:1d					
Targets: 4					
SCSI Target 0					
Path	Canonical Path	Type	Capacity	LUN ID	
vmhba1:0:8	vmhba1:0:8	disk	500.00 GB	8	
vmhba1:0:9	vmhba1:0:9	disk	600.00 GB	9	
vmhba1:0:10	vmhba1:0:10	disk	500.00 GB	10	
vmhba1:0:11	vmhba1:0:11	disk	600.00 GB	11	
vmhba1:0:12	vmhba1:0:12	disk	500.00 GB	12	
vmhba1:0:13	vmhba1:0:13	disk	600.00 GB	13	
SCSI Target 1					
Path	Canonical Path	Type	Capacity	LUN ID	
vmhba1:1:8	vmhba1:0:8	disk	500.00 GB	8	
vmhba1:1:9	vmhba1:0:9	disk	600.00 GB	9	
vmhba1:1:10	vmhba1:0:10	disk	500.00 GB	10	

Figure 4.58 FC adapter targets are displayed.

3. Once your host can see the SAN LUNs, you can go to the Configuration tab and then Storage and then create VMFS volumes by following the same steps as you use to set up local storage. When you initially create a VMFS volume on shared storage, there is no need to create it again on other host servers. You only need to do this one time. To configure your other host servers to see this new VMFS volume, go to that host and select the Configuration tab, Storage Adapters, click the Rescan link, and make sure Scan for New VMFS volumes is selected. After a rescan is performed, go to the Configuration tab, and then Storage, and the new VMFS volume should be displayed and ready for use.

Here are some additional considerations when using FC storage:

- If for some reason you want to hide certain LUNs from specific hosts, you can do this one of two ways. The first way is to mask the LUN at the SAN switch level so that the host cannot see it. The second way is by changing a setting in ESX. To access the settings, select your ESX host, select the Configuration tab, and then click the Advanced Settings link. Then choose Disk in the left pane of the Settings window. In the right pane, find the Disk.MaskLUNs setting and enter the LUN numbers that you want to hide from that host server in the format *HBA adapter# : target #: LUN #*, using commas to separate multiple entries. For example, if you want to mask IDs 10 to 20 and 28 and 30, enter *vmhba1 : 1 : 10-20,28,30*.
- Check your multipathing by selecting your host server and going to the Configuration tab and then selecting Storage. Then choose a storage volume and click the Properties link. In the lower-right window, you will see your paths displayed, as shown in Figure 4.59. The current active path will show as a solid green diamond. You can manage your paths by clicking the Manage Paths button, which leads you to where you can select a path policy. If ESX sees the storage as active/passive, the default path policy is Most Recently Used (which means the same path that was most recently used continues to be used until a failure occurs). If it sees it as active/active, Fixed is used. Besides Most Recently Used and Fixed, there is also a Round Robin option, which is considered experimental in ESX 3.5 and not recommended for production use. For more information about Round Robin load balancing, see VMware's tech note at www.vmware.com/pdf/vi3_35_25_roundrobin.pdf.
- By default, active/passive storage arrays use the Most Recently Used path policy. To avoid LUN thrashing, do not use the Fixed path policy for active/passive storage arrays. By default, active/active storage arrays use the Fixed path policy. When using this policy, you can maximize the utilization of your bandwidth to the storage array by designating preferred paths to each LUN through different storage controllers.
- Create separate LUNs for development, test, and production VMs. This prevents development and test VMs from negatively impacting the performance of production VMs.




Figure 4.59 Multiple paths are displayed for a VMFS volume.

- Consider separating VM disks onto different LUNs (for example, system disks on one LUN, data disks on another LUN).
- The default ESX SCSI command queue depth is 32. Consider raising this to 64 using the `esxcfg-module service console` command. (See VMware's "FC SAN Configuration Guide" at www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_san_cfg.pdf for instructions.) Also, change the `Disk.SchedNumReqOutstanding` setting in the VI Client to 64 under Advanced Settings, Disk.

Configuring iSCSI Storage

iSCSI storage can be configured with your ESX hosts using either a software or hardware initiator. If you are using hardware initiators, you will be using an iSCSI HBA installed in your host server to offload all the iSCSI processing and management from the host server's CPU and NICs. Software initiators are already part of the VMkernel and let you use your existing NICs to connect to iSCSI storage targets.

Here's how to configure your host to use a hardware initiator:

1. On your host, select the Configuration tab, and then Storage Adapters. Your iSCSI HBAs will display in the Adapter list. (It is normal for iSCSI HBAs to have high numbers.) Select the one you want to configure and click the Properties link; a new window will open, as shown in Figure 4.60.




Figure 4.60 Hardware iSCSI initiator Properties window

2. Click the Configure button on the General tab to set up the IP address information for the HBA and to change the default iSCSI name (if you want) and set the alias for the HBA. Click OK when finished, and reboot the host for the changes to take effect.

Here's how to configure your host to use a software initiator:

1. You must first have a VMkernel port group configured on one of your vSwitches to be able to use the iSCSI software initiator that is built in to the VMkernel. If you do not have this set up already, see the "Configuring VMotion" section earlier in this chapter.
2. Enable the port used by iSCSI on the built-in ESX firewall. To do this, select the Configuration tab, and then the Security Profile option, and click the Properties link. Find the Software iSCSI Client entry (port 3260) and click so that a check mark is next to it. Click OK to save it. (Later versions of ESX will automatically open this port when you enable the iSCSI software adapter.)
3. Go to the Configuration tab and select the Storage Adapters option. Select the iSCSI Software Adapter, and then click the Properties link. On the General tab, click the Configure button, and then check the Enabled box and click OK. You can also optionally change the default iSCSI name.

After you have your initiator set up and enabled, you can connect to an iSCSI device and look for storage targets by following these steps:

1. Go to the Configuration tab and select the Storage Adapters option. Select your initiator and click the Properties link.
2. If you are using a hardware initiator, you can set up discovery addresses using dynamic discovery. To locate iSCSI storage devices on your network, select the Dynamic Discovery tab. Click the Add button and enter the IP address of an iSCSI server. Once your host establishes a dynamic discovery session with the iSCSI server, a list of targets on that server will display. You can also add additional targets and edit or remove existing ones.
3. For both hardware and software initiators, you click the Static Discovery tab to enter specific iSCSI host and target names. If you have already done a dynamic discovery, those targets that were found will already be displayed in the list of targets.
4. To set up CHAP authentication if it is enabled on your iSCSI device, click the CHAP Authentication tab and then click the Configure button. Select the Use the Following CHAP Credentials option and enter a CHAP name (or use the initiator name) and secret.

Now that you have found storage targets, it's time to add them to your ESX host and create VMFS volumes by following these steps:

1. First you should rescan your iSCSI device to make sure it sees all the targets available. Go to the Configuration tab and select the Storage Adapters option. Select your initiator and click the Rescan link to have the initiator look for new targets on the iSCSI device. You should see your targets displayed in the bottom pane, as shown in Figure 4.61.

The screenshot shows the 'Storage Adapters' configuration page in vSphere. At the top, there's a 'Rescan...' button. Below it is a table with columns 'Device', 'Type', and 'SAN Identifier'. The table lists two adapters: 'vmhba0' (Block SCSI) and 'vmhba32' (iSCSI). The 'vmhba32' adapter is selected, and its details are shown in the 'Details' pane below. The details pane shows the 'Model' as 'iSCSI Software Adapter', the 'iSCSI Name' as 'iqn.1998-01.com.vmware:esx1-50bc812b', and the 'iSCSI Alias' as 'esx1.xyz.com'. It also shows the 'IP Address' as '192.168.1.100', 'Discovery Methods' as 'Send Targets', and 'Targets' as '1'. The 'SCSI Target 0' details show the 'iSCSI Name' as 'iqn.2006-01.com.openfiler:iscsi-esx1', 'iSCSI Alias' as 'esx1', and 'Target LUNs' as '1'. At the bottom, there's a table with columns 'Path', 'Canonical Path', 'Type', 'Capacity', and 'LUN ID'. The table shows one entry: 'vmhba32:0:0', 'vmhba32:0:0', 'disk', '79.97 GB', and '0'.

Device	Type	SAN Identifier
Smart Array 5i		
vmhba0	Block SCSI	
iSCSI Software Adapter		
vmhba32	iSCSI	iqn.1998-01.com.vmware:...

Details		Properties...		
Model:	iSCSI Software Adapter	IP Address:		
iSCSI Name:	iqn.1998-01.com.vmware:esx1-50bc812b	Discovery Methods: Send Targets		
iSCSI Alias:	esx1.xyz.com	Targets: 1		
SCSI Target 0				
iSCSI Name:	iqn.2006-01.com.openfiler:iscsi-esx1			
iSCSI Alias:	esx1			
Target LUNs:	1			
Hide LUNs				
Path	Canonical Path	Type	Capacity	LUN ID
vmhba32:0:0	vmhba32:0:0	disk	79.97 GB	0

Figure 4.61 iSCSI initiator targets are displayed.

2. Click the Storage link. If there are already VMFS volumes created on the iSCSI device, they will automatically appear on your host after you do a rescan, and you do not need to do anything else to use them.
3. If you need to set up a new VMFS volume with unused iSCSI storage, click the Add Storage link and configure your new VMFS volume using the same steps that were discussed earlier.

Here are some additional considerations when using iSCSI storage:

- The performance of iSCSI storage is highly dependent on network health and utilization. For best performance and security, isolate on a dedicated network.
- You can configure only one software initiator on an ESX Server host. When configuring a vSwitch that will provide iSCSI connectivity, use multiple pNICs to provide redundancy.
- Ensure the NICs used in your iSCSI vSwitch connect to separate network switches to eliminate any single points of failure.
- It is not recommended to extend a VMFS volume using extents to another iSCSI target. A failure of any one of your extents will bring down the whole VMFS volume.
- Always use gigabit NICs; 100MBps NICs are not supported for iSCSI use.
- Always use static IP addresses for initiators. ESX Server does not support DHCP for iSCSI connections.

Configuring NFS Storage

ESX also has a built-in NFS client to connect to network attached storage devices that are running the NFS version 3 protocol. NFS storage is different because it is a file-level protocol and you do not create VMFS volumes on your NFS storage devices. In addition, many NFS storage devices use thin-provisioned virtual disks by default, unlike VMFS volumes, which use thick-provisioned virtual disks. Thin-provisioned disks do not allocate all the space assigned to a VM's virtual disk at once as thick-provisioned disks do. They will be created smaller and then grown in increments as needed up to the maximum size of the assigned disk. NFS storage volumes will appear in the VI Client as VMFS volumes do, and you can use them to store your VMs the same way as you would a VMFS volume.

Watch Out!

When a host accesses a VM disk file on an NFS-based datastore, a special .lck-XXX lock file is generated in the same directory where the disk file resides to prevent other hosts from accessing this virtual disk file. Do not remove the .lck-XXX lock file, because without it, the running VM cannot access its virtual disk file.

To configure your host to use an NFS datastore, follow these steps:

1. You must first have a VMkernel port group configured on one of your vSwitches to be able to connect to an NFS storage device. If you do not have this set up already, see the “Configuring VMotion” section earlier in this chapter.
2. NFS does not use any special storage adapters because it simply accesses storage using the NICs configured on the vSwitch that has the VMkernel configured on it. To add an NFS volume to your host, go to Configuration, select Storage, and then click the Add Storage button.
3. Select Network File System and click Next.
4. Enter the name or IP address of the NFS server and the name of the mount point folder and finally a name that you want to call the datastore. Click Next.
5. Review your entries. Press Back if you need to change one, and then click Finish. The datastore will be added to your host.

Here are some additional considerations when using NFS storage:

- Just like iSCSI, the performance of NFS storage is highly dependent on network health and utilization. For best performance and security, isolate on a dedicated network.
- Mount your datastores exactly the same way on all hosts to prevent the datastore UUID from changing if your hostname is not consistent. (For example, if you are using `nfs1.xyz.com` on one host, don’t use just `nfs1` or the IP address on the others.)
- Always use gigabit NICs; 100MBps NICs are not supported for NFS use.
- Ensure the NICs that are used in your VMkernel vSwitch connect to separate network switches to eliminate any single points of failure.
- Increase your max NFS mounts from 8 to 32 using the advanced setting `Nfs.MaxVolumes` and increase your TCP/IP heap size from 6MB to 30MB using the `Net.TcpipHeapSize`.

Additional Configuration

You can complete some additional configuration steps on your hosts to ensure they are properly configured and ready for operation. Here are some additional things that you should configure on your host servers.

Using NTP Time Synchronization with ESX

It’s best to ensure that your ESX hosts all have the same time, and the best way to do this is to set them up to sync from a reliable time source. In ESX 3.0.x, you had to manually configure this using the Service Console. Beginning with ESX 3.5, you can configure synchronization of time on your ESX host using the VI Client instead. You can read additional

information about timekeeping in a white paper that VMware has published titled “Timekeeping in Virtual Machines” available on their website at www.vmware.com/pdf/vmware_timekeeping.pdf.

To configure NTP time synchronization using the Service Console (ESX 3.0.x), follow these steps:

1. Log on to the Service Console.
2. Edit `/etc/ntp.conf` with `nano` or `vi`.
3. Add the following lines to the `OUR TIMESERVERS` section:

```
restrict default kod nomodify notrap
server 0.pool.ntp.org
server 1.pool.ntp.org
server 2.pool.ntp.org
```
4. Save the file and exit the editor.
5. Edit `/etc/ntp/step-tickers` with `nano` or `vi`.
6. Add the following lines to the file:

```
0.pool.ntp.org
1.pool.ntp.org
2.pool.ntp.org
pool.ntp.org
```
7. Save the file and exit the editor.
8. Enable the NTP time port (123) on the built-in ESX firewall. Type `esxcfg-firewall -enableService ntpClient` or use VI Client to enable it on the Configuration tab and then under Security Profile.
9. Type `service ntpd restart`. It may say that it failed for stop if it is not currently running.
10. Type `chkconfig --level 345 ntpd` to enable the NTP daemon to autostart.
11. Type `hwclock --systohc` to set hardware clock to system clock.
12. Type `ntpdate -q 0.pool.ntp.org` to see the offset between local and NTP clock.

To configure NTP time synchronization using the VI Client (ESX 3.5.x), follow these steps:

1. Log on to vCenter Server or your ESX host with the VI Client, select your host, and click the Configuration tab. Then select Time Configuration under Software.
2. Click the Properties link to edit the settings. In the top section, you can manually edit the date/time. In the bottom section, you can configure the NTP client. Place a check mark by the NTP Client Enabled box, and then click the Options button.

3. On the General tab, select the option to start automatically if ports are opened (default).
4. On the NTP Settings tab, you can add NTP servers by clicking the Add button.
5. When you have finished adding NTP servers, click OK. If you put a check mark by the Restart NTP Service box and the service is not running, you will receive an error message.
6. Click OK, and the settings will be saved. If the NTP service was not running, it will be automatically started. The NTP firewall port (123) will be automatically opened up if it is not already.

Configuring Service Console Memory

By default, an ESX host is configured to use 272MB, which is usually sufficient. However, it is highly recommended to increase this to the maximum size of 800MB, especially if any of the following are true:

- If backup agents are installed on the Service Console
- If third-party management agents (for example, hardware) are installed on the Service Console
- If any additional software at all is installed on the Service Console
- If heavy swapfile utilization occurs on the Service Console or if the Service Console becomes starved for resources (see VMware knowledge base article 1003496 at <http://kb.vmware.com/kb/1003496>)

You can change the amount of RAM assigned to the Service Console, but you must also increase the size of the swap partition in the ESX Service Console, which should be twice the size of the amount of RAM assigned to the Service Console. The default size of the swap partition is 544MB. So if you increase the amount of RAM to 800MB, you need to increase the swap partition size to 1600MB. If you followed the recommended partitioning guidelines from Chapter 3 and already increased the size of the swap partition to 1600MB when you installed ESX, you will be okay to change the Service Console memory to 800MB, which you can do by following these steps:

1. Log on to vCenter Server or your ESX host with the VI Client, select your host, and click the Configuration tab. Then select Memory under Hardware.
2. The current memory settings are displayed. Click the Properties link to edit the settings.
3. You can enter a value between 256MB and 800MB. Enter a new value and click OK.
4. Restart the ESX host for the new memory setting to take effect.

If your partition size is at the default size of 544MB, chances are you will have no room left on the disk that ESX is installed on. So if you want to increase the partition size, you should reinstall ESX and make the swap partition size 1600MB. After you have done this, you can change the Service Console memory to 800MB.

Configuring DNS and Routing

DNS is critical for the proper operation of your ESX host. If you need to change any of the DNS settings (for example, DNS servers, default gateways, search domains) for your host, you can do this by following these steps:

1. Log on to vCenter Server or your ESX host with the VI Client, select your host, and click the Configuration tab. Then select DNS and Routing under Software.
2. The current DNS and routing settings are displayed. Click the Properties link to edit the settings.
3. On the DNS Configuration tab, you can change the host DNS name and domain (must reboot host for changes to take effect), specify a fixed DNS server or use DHCP to obtain them, and set search domains (the DNS domains that are searched when a DNS hostname is entered without a DNS suffix).
4. On the Routing tab, you can change the gateway IP address for the Service Console interface and the VMkernel interface.

Configuring Virtual Machine Startup/Shutdown

The VM startup and shutdown settings control what happens with VMs when a host server is powered on or off. By default, they are disabled, and VMs are not automatically started when a host is powered on, and VMs are not powered off automatically if a host server is powered off. Because these settings are host specific, if you are using a cluster with HA or DRS enabled you will want to leave these settings disabled, because VMs may migrate to other hosts and the settings will no longer work. Even if you do not use a cluster, you may want to leave these settings disabled because there may be times that you do not want VMs to automatically start up on your host.

Watch Out!

There was a bug in ESX 3.0.x that was patched in 3.0.2 and was later reintroduced in ESX 3.5.x, where if you had VM startups enabled and you restarted the management service on an ESX host (mgmt - vmware) all your VMs would be shut down (see VMware knowledge base article 1003312). You might consider temporarily disabling automatic VM startups if you have to restart that service to resolve a problem.

If you do want to use these settings, you can configure them as follows:

1. Log on to vCenter Server or your ESX host with the VI Client, select your host, and click the Configuration tab. Then select Virtual Machine Startup/Shutdown under Software.
2. The current settings and startup order are displayed. Click the Properties link to edit the settings.
3. At the Virtual Machine Startup and Shutdown window check the Allow Virtual Machines to Start and Stop Automatically with the System Box so that you can edit the settings.
4. The top sections let you set a default startup and shutdown delay to be used by all VMs. This delay is how long the host will wait after it starts before it starts up the VM. In addition, you can configure a default shutdown action for your VMs (power off, guest shutdown, or suspend).
5. The bottom section lets you configure an automatic startup order (or any order), for which you can configure VM settings individually. For example, you may want to set a longer startup delay on an application server so that you can ensure that a database server that it relies on has already been started. You can use the Move Up and Move Down buttons to set your VM's startup and shutdown order. When you have finished, click OK to save your settings.
6. If you want to temporarily disable this later, simply uncheck the top box. Your current VM settings will be preserved, and you can go back and check it again later when you want to enable it again.

Changing the Root Password

If you need to change the root password of an ESX host server, there are two ways to do this. It's a good idea to do this periodically for security purposes because the root account is the most important account on an ESX host and if it is compromised someone could easily gain access to all the VMs. Changing this password will not affect vCenter Server because it uses a separate account called vpxuser that it creates after it connects to an ESX host for the first time.

To change the root password using the Service Console, follow these steps:

1. Log on to the Service Console as the root user.
2. Type `passwd` and press Enter.
3. You will be prompted to enter a new password. You will receive an error message if the password is too short, based on a dictionary word, or not complex enough. However, it will still let you change it.
4. You will be prompted a second time to enter the password, and then the password will be changed.

To change the root password using the VI Client, follow these steps:

1. You need to connect directly to the ESX host with the VI Client. If you connect to it through vCenter Server, you will not see the Users and Groups tab. Connect to the VI Client and log on as the root user.
2. Click the Users and Group tab, and then in the Users view double-click the root account.
3. Once the account Properties window displays, check the Change Password field and enter a new password and confirm it.
4. Click OK, and the new root password is set.

How to Reset a Forgotten Root Password

If you have forgotten the root password of an ESX server and have no other way to change it, there is a way to change it if you have physical access to the ESX server:

1. Shut down ESX and restart it. At the first menu, type `a`.
2. At the next prompt, type `single`.
3. ESX will now go through the boot process, and you will end up at a `#` prompt.
4. Type `passwd`, enter a new password for root, and retype it when prompted.
5. Type `reboot`, and you're done.

Summary

This chapter has covered a lot of configuration steps to get your environment ready to be populated with VMs. For further information about configuring the different components of your environment, be sure to also check the VMware documentation for the release that you are installing. Now that we have our environment configured, it is time to make sure it is properly secured before we can start putting VMs on our host servers.

Endnotes

1. The use of the term VirtualCenter here refers to the old name which still applies to VirtualCenter version 2.0.x and is also still present inside the application in vCenter Server 2.5 as the software has not yet been updated to reflect the new name.

Chapter 5

Securing Your Virtual Environment

You might be tempted to skip over this chapter and start populating your new virtual environment instead. Although ESX is fairly secure out of the box, it is important that you understand what makes it secure and the various configuration settings that can change its security. If you do not, you may inadvertently change a setting that decreases the security of your host and its VMs. Security is amplified in a virtual environment, where a single physical server is running many virtual servers and inadequate security on the physical server can directly affect the security of all the virtual servers running on that host.

So take some time to understand security in virtual environments and apply recommended security settings to all the components that make up your virtual environment, including hosts, VMs, networks, and vCenter Servers. Protecting the host server cannot be emphasized enough. Think of a host server as a castle with the VMs all protected inside of it. If someone compromises your castle's defenses, he will have free access to everything inside. Therefore, you want to do everything you can to make sure your castle's defenses are adequate and that you do not forget to put water in the moat and raise the castle drawbridge. So let's get started securing our castle.

Network Connectivity and Ports

The different components that make up a VI3 infrastructure operate on a variety of network ports, so you should know which ports are used so that you can ensure that they are not being blocked by any firewalls. Table 5.1 summarizes the most commonly used ports with VI3 as of version 3.5.

Table 5.1

Commonly Used Network Ports in VI3

Port Number	Usage	Traffic Type
22	Used by ESX for SSH server access.	Incoming TCP
25	Used by vCenter Server to send email alerts via SMTP.	Outgoing TCP
80	HTTP access. The default nonsecure TCP web port typically used in conjunction with port 443 as a front end for web access to ESX hosts and vCenter Server. Port 80 is used for the default page, which then redirects to port 443 when logging on.	Incoming TCP
123	Used by ESX if you set up NTP to sync from a time source.	Outgoing UDP
443	HTTPS access. The default SSL web port used for the following: Connection to VI web access from the web VI web access and third-party network management client connections to the vCenter Server Direct VI web access and third-party network management clients access to ESX hosts VI Client access to the vCenter Server Direct VI Client access to ESX hosts VMware Update Manager VMware Converter	Incoming TCP
902	Authentication traffic for the ESX host and VM configuration. Used for VI Client access to the vCenter Server, vCenter Server access to ESX hosts, direct VI Client access to ESX hosts, and ESX Server host access to other ESX hosts for migration and provisioning.	Incoming TCP Outgoing UDP
903	Remote console traffic generated by user access to VMs on a specific ESX host. Used for VI Client and VI Web Client access to VM consoles.	Incoming TCP
2049	Transactions from your NFS storage devices. Used on the VMkernel interface rather than the Service Console interface	Incoming TCP Outgoing TCP
2050– 2250	Traffic between ESX hosts for VMware High Availability (HA) and EMC Autostart Manager.	Outgoing TCP Incoming UDP Outgoing UDP
3260	Transactions from your iSCSI storage devices. Used on the VMkernel interface rather than the Service Console interface.	Outgoing TCP

Port Number	Usage	Traffic Type
8000	Incoming requests from VMotion. Used on the VMkernel interface rather than the Service Console interface	Incoming TCP Outgoing TCP
8042-8045	Traffic between ESX hosts for HA and EMC Autostart Manager.	Outgoing TCP Incoming UDP Outgoing UDP
27000	License transactions from ESX to the license server.	Outgoing TCP
27010	License transactions to ESX from the license server.	Incoming TCP

By the Way

Earlier versions of ESX 3.0 used port 902 for VI Client access to ESX hosts and vCenter Servers instead of port 443.

Once you know the ports that are used by the various components, make sure both operating system firewalls and network firewalls are configured to allow the ports that you will use. This includes the operating system firewall on the vCenter Server and any workstations that will be using the VI Client or web UI. The connections and port numbers used in VI3 are outlined in Figure 5.1.




Figure 5.1 Network connections and port numbers used in VI3

One important thing to note is how remote console connections to VMs are handled. When you connect to vCenter Server and manage ESX hosts, all network traffic is between your workstation and the vCenter Server on ports 902 and 443. However, if you select a VM and open a remote console to it, you are connecting to the ESX host that the VM is located on directly on port 903 and not through vCenter Server. Therefore, you must be sure that the ESX hostname is resolvable through DNS from your workstation and that you can connect directly to the ESX on port 903. If your network is configured so that you have network access only to vCenter Server and not the ESX hosts, you will need to RDP to the vCenter Server to open a remote console session to a VM. In addition, you can add a setting to ESX hosts to allow remote console traffic to be proxied on port 902 rather than port 903 (see VMware knowledge base article 749640).

Securing the ESX Service Console

The ESX Service Console should be protected at all costs because it can provide backdoor access to the VMs regardless of the operating system security that is in place. The ESX Service Console is based on a modified Red Hat Linux distribution, and likewise many of the commands and security practices are the same. A standard unmodified ESX installation is fairly secure by default, but can definitely be improved on to further enhance its security. At the same time, it is possible to weaken the security by making certain configuration changes that make administrative tasks easier.

The Root User Account

The root account is the default administrator account for the ESX Service Console and should not be directly used for standard administration tasks. It has access to all the commands and files on the Service Console and can be considered a super user account. The root account password should be complex and known only by a limited number of administrators. In addition, the password should be changed periodically to protect the account from anyone who may have known the password but has left the company or inadvertently found it out.

ESX hosts do not allow remote Telnet access to the Service Console because it is inherently unsecure because Telnet uses unencrypted plain-text communication. Instead, SSH (Secure Shell) is used, which encrypts the communication between the client and the host server. By default, ESX is configured to not allow the root account to be able to log on using remote SSH connections to the Service Console. Only nonroot accounts are allowed to do this. This is done intentionally to discourage administrators from using the root account. However, the first thing many administrators do to new ESX installations is to enable this access to make it easier to log on to the host using SSH instead of setting up a separate account.

So first let's go over the proper and more secure method for accessing the Service Console, both remotely via SSH and locally. You should create at least one and preferably more additional accounts to be used in lieu of the root account. Instead of using a generic account, create a separate account for each user who will be accessing the Service Console. These accounts will not have the privileges of the root account, but you can use the `su` command after you log on with a nonroot account to temporarily grant the root super user privileges to your account. The `su` command by itself only grants you root privileges. You will be prompted for the password for the root account after you enter the command. If you instead type a dash after the `su` command (for example, `su -`), you will also inherit the root user's path and user environment, which makes it easier to enter in commands. If you omit the dash, you will have to either switch to specific directories to enter commands or include the path to the command. By using the dash, you can enter the commands from any directory. For even more security, consider limiting access to the `su` command or use `sudo` instead. Let's go over the methods for creating a nonroot user account.

Creating a nonroot user account using the VI Client

1. Connect to an ESX host directly using the VI Client. You cannot create ESX local accounts through vCenter Server.
2. Click the Users and Groups tab, as shown in Figure 5.2. This tab displays only when directly connected to ESX hosts and will not appear in vCenter Server. This tab is for managing all the Service Console local users and groups. Be careful here and do not delete or change any users or groups unless you know what they are. Deleting some of the default system users and groups can cause problems with your host and vCenter Server. You can delete user accounts simply by selecting them and right-clicking them and choosing the Remove option.
3. Right-click inside the Users pane and select the Add option.
4. Enter a logon name and password and check the Grant Shell Access to This User box, as shown in Figure 5.3. If you want to enter a descriptive name, you can enter one in the Username field. This name is not used to log on with. The UID field will automatically be populated when you save the account. When you have finished, click the OK button and that's it.

Creating a nonroot user account using the Service Console

1. Log on to the ESX Service Console and use `su` to gain root privileges.
2. Type in `useradd username to create` and press Enter to create the account.
3. Next type in `passwd username from previous step` and you will be prompted to enter a password for the user. You will be prompted a second time to confirm the password. You may receive warnings if the password is not complex enough or based on a dictionary word, but the password will still be allowed.

esx1.xyz.com VMware ESX Server, 3.5.0, 110268

Summary Virtual Machines Resource Allocation Performance Configuration **Users & Groups** Events Permissions

View: Users Groups

UID	User	Name
8	mail	mail
99	nobody	Nobody
0	root	root
11	operator	operator
7	halt	halt
13	gopher	gopher
500	vpuser	VMware VirtualCenter administration account
12	vimuser	vimuser
29	rpcuser	RPC Service User
5	sync	sync
32	rpc	Portmapper RPC user
69	vcsa	virtual console memory owner
37	rpm	
1	bin	bin
74	sshd	Privilege-separated SSH
65534	nfsnobody	Anonymous NFS User
6	shutdown	shutdown
77	pcap	
38	ntp	
3	adm	adm
14	ftp	FTP User
28	nscd	NSCD Daemon
10	uucp	uucp
4	lp	lp
9	news	news
2	daemon	daemon

Figure 5.2 Users and Groups tab when connecting directly to an ESX host server

Add New User

User Information

Login: UID:

User Name:

User name and UID are optional

Enter password

Password:

Confirm:

Shell Access

☒ Grant shell access to this user

Group membership

Group:

Figure 5.3 Creating a new user account directly on an ESX host server

After you have created an account, you can log on to the Service Console with it to test it out. After you log on, you can type `su -` to grant root privileges to your account and then try executing any of the ESX command-line utilities to make sure they work okay. You can delete user accounts by using the `userdel` command (for example, `userdel username`).

Now let's cover the less-secure method, which changes the ESX Service Console configuration to allow the root account to log on remotely using SSH. This is not recommended, but if you do want to do this, follow these steps to change the configuration to allow it:

1. Log on to the ESX Service Console as the root user.
2. Change to the `/etc/ssh` directory.
3. Edit the `sshd_config` file by typing `nano sshd_config`. (Optionally, you can use the `vi` editor.)
4. Page down and find the entry called `PermitRootLogin` and change the `No` to `Yes`.
5. Save the file by pressing `Ctrl-O` (WriteOut).
6. Exit the nano editor by pressing `Ctrl-X` (Exit).
7. Restart the `sshd` service by typing `service sshd restart`.

Watch Out!

Allowing SSH access to the root account weakens the security of the ESX Service Console, and it is highly recommended not to do this in production environments.

Configuring Sudo


Sudo, which is a more secure method for allowing access to the Service Console, lets users log on with their own non-super-user account and run certain commands that you define as other users. Sudo access is controlled by a special configuration file called `/etc/sudoers`. Although this is a text file, it is highly recommended that you use the special `visudo` application to modify it. `Visudo` uses the Linux `vi` editor to open the file, so all the standard `vi` commands work when editing it. `Vi` can be fairly complicated to use for someone new to it because it relies on many different commands to perform operations. We'll show you how to use some of the basic commands as we show you how to set it up. To set up `sudo` on your ESX host, follow these steps:

1. Log on to the Service Console as the root user.
2. If you have not already created your additional user accounts to use with `sudo`, set them up using the `useradd` command as previously described.

3. Type `visudo` to open the sudo configuration editor.
4. The `sudoers` file has several sections that you need to set up, as shown in Figure 5.4 and as listed here:

Creating a new user account directly on an ESX host server

- **Host_Alias.** This section specifies the hosts that the privileges apply to. You will not need this section if you set up privileges individually on hosts. If you do want to sync a single `sudoers` file among multiple hosts, you should look up information about using `rsync`.
- **User_Alias.** This section is for setting up an alias that you can use for groups of users in lieu of having to set them up individually. Once you create an alias, you can use it in the Privileges section to assign commands to users. For example, you could use `VIADMINS` as an alias for multiple user account like this: `User_Alias VIADMINS = dduck, mmouse, goofy`.
- **Cmnd_Alias.** This section is for specifying an alias for a specific command, which makes it easier to assign commands to users so that you don't have to type out the full path and command name each time you assign it to a user. For example, you could use `NICS` for the command `/sbin/vmware/esxcfg-nics` like this:
`Cmnd_Alias NICS = /sbin/vmware/esxcfg-nics`.
- **User privilege.** This section is for assigning privileges to users. Here you specify which users can use which commands using either the user and command aliases or the actual user and command names. For example, to assign the `VIADMINS` user alias to be able to use the `esxcfg-nics` command with an alias of `NICS`, you would add a line like this: `VIADMINS ALL=NOPASSWD: NICS`.



```
# Host alias specification
# User alias specification
# Cmnd alias specification
# Defaults specification
# User privilege specification
root    ALL=(ALL) ALL
```

Figure 5.4 The default `sudoers` file is shown using the `visudo` editor.

5. `Vi` is different from text editors like Windows Notepad and can be a bit tricky to use if you are not familiar with the commands to use. Here is a list of some of the most common commands that you might use. You will initially be in command mode, and once you enter insert mode you can press the `Esc` key to go back to command mode:
- `i` Insert text before cursor
 - `a` Append text after cursor

- o Open and put text in a new line below current line
 - r Replace single character under cursor
 - R Replace characters starting with current position
 - x Delete the single character under cursor
 - D Delete the remainder of the line starting with current position
 - :wq or :x Save the file and quit
 - :q Quit
 - :q! Quit without saving
6. After you set up your user and command aliases and privileges, you can save the file, and they will take effect immediately. See Figure 5.5 for a sample `sudoers` configuration file. If you do not want to specify individual command privileges, you can specify a directory instead (for example, `/usr/sbin/`), which will provide access to all the commands in that directory.

```
# Host alias specification

# User alias specification
User_Alias    VIADMINS = dduck,mmouse,goofy


# Cmnd alias specification
Cmnd_Alias    NETWORK = /usr/sbin/esxcfg-nics, /usr/sbin/esxcfg-vswitch, /usr/
/sbin/esxcfg-vswif
Cmnd_Alias    FIREWALL = /usr/sbin/esxcfg-firewall
Cmnd_Alias    VMKD = /usr/sbin/vmkfstools

# Defaults specification

# User privilege specification
VIADMINS      ALL=NOPASSWD:  NETWORK, FIREWALL, VMKD
```

Figure 5.5 A sample `sudoers` file is shown with user and command aliases and privileges.

7. One last step is to add the path to the commands to users' profiles so that they will not have to type the path before the command each time they want to run one. To do this, go into their home directory (for example, `/home/dduck`) that is created when the user was created and edit the `.bash_profile` file using `vi` or `nano`, and add the path to the commands to it. To separate multiple entries in the path, use a colon. The dot in front of the filename hides the file from directory listings so that you won't see it if you use the `ls` command but you can still edit it by specifying the filename (for example, `nano .bash_profile`). In Figure 5.6, we have added `/usr/sbin`, which is the location of many of the ESX commands to the path, so that the user will not have to type the path when entering them.



```
root@esx2:/home/dduck
GNU nano 1.2.1      File: .bash_profile

# .bash_profile

# Get the aliases and functions
if [ -f ~/.bashrc ]; then
    . ~/.bashrc
fi

# User specific environment and startup programs


PATH=$PATH:$HOME/bin:/usr/sbin

export PATH
unset USERNAME

```

Figure 5.6 Editing a user profile file to add a path to it

Now that you have sudo set up, when users want to use one of the commands, they type `sudo command name` (for example, `sudo vmkfstools` or `sudo esxcfg-nics`), and they can use all the parameters and options at the end of them like they would normally use for the command. If they try to run the command without typing `sudo` before it, they will receive a permission denied error, as shown in Figure 5.7.



```
[dduck@esx2 dduck]$ sudo esxcfg-nics -l
Name      PCI      Driver      Link Speed  Duplex MTU      Description
vmnic1    0a:02:00 tg3         Up 1000Mbps Full 1500      Broadcom Corporation NC
7770 Gigabit Server Adapter (PCI-X, 10/100/1000-T)
[dduck@esx2 dduck]$
[dduck@esx2 dduck]$ esxcfg-nics -l
-bash: /usr/sbin/esxcfg-nics: Permission denied
[dduck@esx2 dduck]$
```

Figure 5.7 Using sudo to run the `esxcfg-nics` command and running the command without sudo

There are a lot more options that you can use in your `sudoers` file for things such as controlling whether users are prompted to enter their password again when running `sudo` commands, to send emails for wrong passwords or unallowed attempts, `sudo` logging options, and more. For more information about using `sudo`, check out the online reference at www.sudo.ws. For more information about using the `vi` editor, check out the online reference at www.computerhope.com/unix/uvi.htm.

Configuring the ESX Service Console to Authenticate Using Active Directory

You can optionally configure the ESX Service Console to use Active Directory (AD) as an authentication mechanism in addition to using local user accounts. This is useful if you have many people who need to have access to the Service Console and saves you from having to manage separate passwords for them on the ESX hosts. PAM (Pluggable Authentication Modules) authentication enables you to use your AD passwords to log on to the ESX Service Console. You still need to create an identical account on the ESX host, but once you do that you can log on using your AD password after PAM is configured.

The advantage of doing this is password management. You do not have to manage ESX passwords separately and can leverage AD's powerful account management features. Before setting this up, you want to ensure that your ESX servers sync time from the same sources as your AD servers. If you use NTP on both systems, this should not be an issue. Kerberos is very sensitive to time variations, and by default an AD environment requires computer clocks to be synchronized within five minutes. If your ESX servers are not time synced with AD, you will probably run into logon issues. (You will see "clock skew too great" errors in `/var/log/messages`.) Here are the steps to set this up on an ESX host:

1. Log on to the ESX Service Console.
2. If you're adding a user that does not exist in the Service Console, you must first add it using the techniques that we covered earlier. You can view the existing user accounts by using the VI Client and connecting directly to an ESX host and looking at the Users and Groups tab or by opening the `passwd` file located in the `/etc` directory, which is where the accounts are stored on Linux systems. The username must exactly match the same account name of the AD user. Also be advised that usernames that begin with a number are not supported on Linux systems and cannot be used.
3. You will need to open the required ports on the built-in ESX firewall. To do this, just enter the following commands:
 - `esxcfg-firewall -e activeDirectoryKerberos` (Note the missing `y` in Directory is correct.)
 - `esxcfg-firewall -e LDAP`
 - `esxcfg-firewall -e kerberos`
4. Next you need to enable PAM authentication in ESX using the `esxcfg-auth` command, which writes the configuration information to various files, including `/etc/krb5.conf`. Use `esxcfg-auth` as outlined here, where DN is your fully qualified AD domain name (for example, `xyz.com`) and DC is the fully qualified name of a domain controller (for example, `dc1.xyz.com`). Optionally, you can use just the

domain name rather than the domain controller name if your DNS is set up correctly and you can ping the domain name. DNS will automatically forward from the domain name to the domain controller name:

- `esxcfg-auth --enablead --addomain=DN --addc=DC`
- `esxcfg-auth --enablekrb5 --krb5realm=DN --krb5kdc=DC`
`--krb5adminserver=DC`

The options for the `esxcfg-auth` command are described as follows:

- `--enablead` enables AD authentication.
 - `--addomain` sets the AD domain.
 - `--addc` sets the AD domain controller. Use this multiple times to add domain controllers for redundancy.
 - `--enablead` enables Kerberos authentication.
 - `--krb5realm` sets the Kerberos realm (AD domain).
 - `--rb5kdc` sets the Kerberos Key Distribution Center (domain controller).
 - `--krb5adminserver` sets the Kerberos Admin Server (domain controller).
5. Once you enable this, it takes effect immediately. Log back on to the console, and you should be able to use your AD password for that account. If you run into difficulty, you can check the `/var/log/messages` file for errors. Also, ensure that your ESX host can resolve the FQDN of the AD server you entered previously.
 6. If you plan to log on using the VI Client directly to the ESX server with the local user account, you will first need to click the Permissions tab in the VI Client and add a permission and role for the local user account.
 7. If you want to disable this, type `esxcfg-auth --disablead` and `esxcfg-auth --disablekrb5` and disable the firewall ports by doing Step 3 with a `-d` rather than `-e`.

Installing Antivirus Software on the ESX Service Console

The subject of whether to install antivirus software on the ESX Service Console comes up fairly often because most administrators are used to installing it on all servers to protect them. Although the ESX Service Console is an operating system, it is less susceptible to viruses and other problems that affect general-purpose operating systems like Microsoft Windows. ESX runs a customized, locked-down version of Linux based on Red Hat Enterprise 3 Update 6, and subsequently there is much less likelihood of security exploits than in a standard Linux distribution.

Despite this, it does not mean that the ESX Service Console is completely impervious to a virus; the chances of it happening are just extremely low. If you follow the standard security best practices for hardening the Service Console, including isolating the Service Console network, it is recommended to not install antivirus software on the Service Console.

Another reason not to install antivirus software is that it can impact the performance of the ESX host and subsequently all the VMs that reside on it because of the extra CPU, memory, and disk resources that antivirus software typically uses. Antivirus software can be particularly resource intensive and can draw resources away from the VMs and potentially have a very negative impact on them. In addition, there is the possibility that the software may cause other issues on the Service Console, as is true with any additional third-party software installed on the Service Console.

Despite the reasons for not installing it, some enterprises mandate that antivirus software be installed on all systems regardless of how secure they are. You might try to exempt your ESX hosts from this by explaining the following points:

- The design of ESX does not allow for a VM that is infected by a virus to spread it to the ESX host through the VMkernel. It can spread only through traditional means, which is typically over the network by leveraging open ports and OS vulnerabilities. If you isolate your Service Console network, it is protected from any VMs that may be infected. Reference the white paper at www.vmware.com/pdf/vi3_security_architecture_wp.pdf from VMware for more information.
- Most viruses are written to exploit Windows systems. *Very few* are written specifically for Linux systems.
- ESX has a built-in firewall that protects the Service Console and blocks all ports except those few required by ESX and vCenter Server.
- ESX has a good historical track record. I've never heard of a virus infecting the ESX Service Console. This doesn't mean it can't happen, just that the chances are extremely low.
- Explain the negative performance impact the antivirus software will have on an ESX host, which can affect all the VMs on the host.

If you are forced to install antivirus software on the Service Console because of security requirements in your environment, you should follow these guidelines:

1. Make sure that you run a version designed for the version of Linux that the Service Console uses (Red Hat Enterprise 3 Update 6).
2. Try to configure the antivirus software as minimal as possible to help minimize the impact on the server's resources. It's best to exclude all your VMFS volumes from scanning or at a minimum exclude specific VM files such as VMDK and VSWP files that are frequently written to.

3. Make sure to keep a close eye on the performance of the host to see how much impact the antivirus software is having on it.
4. If you have to run on-demand scans, make sure they are performed in off-peak hours when activity on the ESX host is low.

The ESX Built-In Firewall

Beginning with version 3.x, a firewall was built in to ESX to help protect the Service Console. This firewall is based on the iptables firewall that is used by Linux systems and runs as a service on the ESX host. However, you should not use standard iptables commands to administer it and should use the VI Client and `esxcfg-firewall` command instead. It's important to note that it only protects the Service Console and does not protect any of the VMs, even if they are on the same vSwitch and network as the Service Console. It is enabled by default, and the rules are configurable using the VI Client or through the Service Console. The services configuration file that is used for the firewall is called `services.xml` and is located in the `/etc/vmware/firewall` directory on the ESX host. This file contains a list of the well-known services and ports that are defined by name. It does not contain information about which rules are enabled and disabled. To change the firewall rules that control which network ports are allowed and blocked, you can use either the VI Client or ESX Service Console.

You can manage the firewall through the VI Client as described here:

1. Select your ESX host, and then choose the Configuration tab, and then the Security Profile option under Software, as shown in Figure 5.8.




Figure 5.8 ESX firewall under Security Profile

- Click the Properties link to bring up the firewall properties, as shown in Figure 5.9. You can check and uncheck the services that will open the various ports related to that service. Toggling the state of the check box next to a service will enable/disable the ports associated with the service. When a check mark is present, the ports are open, allowing traffic to flow through. You cannot add new services using the VI Client. If you need to open a port that is not listed, use the Service Console method instead. Any changes to the firewall properties will take effect immediately.




Figure 5.9 ESX Firewall configuration properties

You can manage the firewall using the Service Console `esxcfg-firewall` command as described here:

- Log on to the ESX Service Console. The syntax for the command is `esxcfg-firewall options`. The command can optionally use well-known services that are defined by name in the `services.xml` file. The `esxcfg-firewall` command can be used with the following options:
 - `-q` Lists the current settings.
 - `-q <service>` Lists settings for the specified service.
 - `-q incoming | outgoing` Lists settings for non-required incoming/outgoing ports.
 - `-s` Lists all the well-known services defined in the `services.xml` file

- `-l` Loads current settings contained in the `services.xml` file
 - `-r` Resets all options to the defaults
 - `-e <service>` Allows specified service through the firewall (enables)
 - `-d <service>` Blocks specified service (disabled)
 - `-o <port, tcp|udp, in|out, name>` Opens a port
 - `-c <port, tcp|udp, in|out>` Closes a port previously opened by `-o`
 - `-h` Displays command help
 - `-allowIncoming` Allows all incoming ports
 - `-allowOutgoing` Allows all outgoing ports
 - `-blockIncoming` Blocks all nonrequired incoming ports (default)
 - `-blockOutgoing` Blocks all nonrequired outgoing ports (default)
2. In addition, the services defined (not all are enabled, just defined) by default are listed here. These are current as of ESX 3.5 Update 3. Some are new, and the names have changed on others since ESX 3.0.x. If you want to add new ones, edit the `/etc/vmware/firewall/services.xml` file and add them at the end. Just make sure to increment the service ID properly so that it is unique for the new service.
- `aam` Traffic between ESX Server hosts for VMware HA and EMC Autostart Manager. Inbound and outbound TCP and UDP ports 2050–2250 and 8042–8045 (previously known as AAMClient)
 - `activeDirectorKerberos` Active Directory Kerberos (outbound TCP ports 88 and 464)
 - `caARCserve` Backup agent: CA Arcserve (incoming TCP port 6051)
 - `CIMHttpServer` First-party optional service: CIM HTTP server (inbound TCP port 5988)
 - `CIMHttpsServer` First-party optional service: CIM HTTPS server (inbound TCP Port 5989)
 - `CIMSLP` First-party optional service: CIM SLP (inbound and outbound TCP and UDP ports 427)
 - `commvaultDynamic` Backup agent: Commvault dynamic (inbound and outbound TCP ports 8600–8619)
 - `commvaultStatic` Backup agent: Commvault static (inbound and outbound TCP ports 8400–8403)
 - `ftpClient` FTP client (outbound TCP port 21)
 - `ftpServer` FTP server (inbound TCP port 21)
 - `kerberos` Kerberos (outbound TCPs port 88 and 749)
 - `LDAP` Non-encrypted LDAP. outbound TCP and UDP port 389

- **LDAPS** Encrypted LDAP - outbound TCP and UDP port 636
- **legatoNetWorker** Backup agent Legato (EMC) NetWorker (inbound and outbound TCP ports 7937–9936)
- **LicenseClient** FLEXlm license server client (outbound TCP ports 27000 and 27010)
- **nfsClient** NFS client (outbound TCP and UDP ports 111 and 2049 [0–65535])
- **nisClient** NIS client (outbound TCP and UDP ports 111 [0–65535])
- **ntpClient** NTP client (outbound UDP port 123)
- **smbClient** SMB client (outbound TCP ports 137–139 and 445)
- **snmpd** SNMP services (inbound TCP port 161 and outbound TCP port 162)
- **sshClient** SSH client (outbound TCP port 22)
- **sshServer** SSH server (inbound TCP port 22)
- **swISCSIClient** First-party optional service: Software iSCSI client (outbound TCP port 3260)
- **telnetClient** Telnet client (outbound TCP port 23)
- **TSM** Backup agent: IBM Tivoli Storage Manager (inbound and outbound TCP ports 1500)
- **symantecBackupExec** Backup agent: Symantec BackupExec (inbound TCP ports 10000–10200, previously known as veritasBackupExec)
- **symantecNetBackup** Backup agent: Symantec NetBackup (inbound TCP ports 13720, 13732, 13734, and 13783; previously known as veritasNetBackup)
- **updateManager** VMware Update Manager (outbound TCP ports 80, 9000–9100)
- **VCB** VMware Consolidated Backup (outbound TCP ports 443, 902)
- **vncServer** VNC server: Allow VNC sessions 0–64 (inbound TCP ports 5900–5964)
- **vpXHeartbeats** VPX heartbeats (outbound UDP port 902)

Some examples of using the `esxcfg-firewall` command are listed here:

- Enable SSH client connections from the Service Console: `esxcfg-firewall -e sshClient`
- Disable the Samba client connections: `esxcfg-firewall -d smbClient`
- Allow syslog outgoing traffic: `esxcfg-firewall -o 514,udp,out,syslog`
- Turn off the firewall (for troubleshooting purposes): `esxcfg-firewall -allowIncoming, esxcfg-firewall -allowOutgoing`
- Reenable the firewall: `esxcfg-firewall -blockIncoming, esxcfg-firewall -blockOutgoing`

If you need to stop and start the firewall for troubleshooting purposes, you can also use the following command to stop it: `service firewall stop`. Then to restart it, you use this command: `service firewall start`. It is not recommended to disable your firewall for any length of time because it provides critical protection for your ESX host.

Security Best Practices for the ESX

Let's go over some of the recommended security best practices for ESX as recommended by VMware and the Center for Internet Security (CIS). These best practices are guidelines for increasing the security of your ESX host, and the full documents for these are available in the "Additional Security Resources" section at the end of this chapter.

Server Hardware

Here are recommended security best practices for physical server hardware:

- Consider password protecting access to your host server's BIOS to prevent unauthorized changes being made by anyone who has physical access to the host.
- Try to prevent physical access to the host system by ensuring that it is located in a locked room that has controlled access. If the host is in a datacenter, ensure that the rack it is located in is locked to prevent unauthorized access.
- Configure the host server's BIOS to not be allowed to boot from removable storage devices (USB, floppy, CD-ROM).

Service Console

Here are recommended security best practices for the ESX Service Console:

- Avoid using it unless absolutely necessary. Most tasks can be accomplished using the VI Client instead.
- Avoid installing any software on the Service Console unless absolutely required. The only software that you may consider installing are hardware management agents (although later versions of ESX 3.5 no longer require this because they are built in using CIM SMASH providers), backup software (typically not needed, however, because there is no need to back up the Service Console), antivirus software (only if you have to), and sometimes third-party applications that are designed to work with ESX (for example, backup and replication).
- Do not create shared user accounts on the Service Console; create unique accounts for anybody who needs to log on to it. In addition, configure AD authentication to simplify password management and control.
- Do not enable remote access (SSH) for the root account; leave it disabled like it is by default.

- Consider disabling su access, which provides root access and instead set up sudo.
- If you are not using AD authentication on the Service Console, configure password policies on the ESX host. To configure this, follow these steps:
 1. Set the maximum number of days a password can be used by using the `esxcfg-auth --passmaxdays=days` command (for example, `esxcfg-auth --passmaxdays=60`). This updates the `PASS_MAX_DAYS` field in the `/etc/login.defs` file.
 2. Set the minimum number of days before a password can be changed by using the `esxcfg-auth --passmindays=days` command (for example, `esxcfg-auth --passmindays=3`). This updates the `PASS_MIN_DAYS` field in the `/etc/login.defs` file.
 3. Set the number of days before a warning is given before a password expires by using the `esxcfg-auth --passwarnage=days` command (for example, `esxcfg-auth --passmindays=10`). This updates the `PASS_WARN_AGE` field in the `/etc/login.defs` file.
 4. Set the number of failed logins before an account is locked out by using the `esxcfg-auth --maxfailedlogins=count` command (for example, `esxcfg-auth --maxfailedlogins=5`). This updates the `deny=` field in the `/etc/pam.d/system-auth` file.
 5. Enable password complexity by using the `esxcfg-auth --usepamqc=values` command. Six values are required: `N0`, `N1`, `N2`, `N3`, `N4`, `match` (for example, `esxcfg-auth --usepamqc -1 -1 -1 8 7 4`). This updates the `min=` and `match=` fields in the `/etc/pam.d/system-auth` file. The values define the length of the password based on the number of character classes used in the password; `-1` disables the class. There are four character classes: digits, lowercase letters, uppercase letters, and all others. The values are as follows:
 - ◆ `N0` This value represents a password that consists of one character class, such as all lowercase letters. This is usually disabled (`-1`).
 - ◆ `N1` This value represents a password that consists of two character classes and does not meet the requirements for a passphrase. This is usually set to 24. It is recommended to disable this, too.
 - ◆ `N2` This value is reserved for passphrases, which must contain a minimum number of words. This is usually set to 12. It is recommended to disable this, too.
 - ◆ `N3` This value represents a password that consists of three character classes. This is usually set to 8, which is a good value to use.
 - ◆ `N4` This value represents a password that contains four different character classes. This is usually set to 7, which is a good value to use.

- ◆ **Match** This value is a case and direction insensitive substring-matching function used to determine whether a substring is common. This determines the number of characters from an old password that can be reused in a new password. This is usually set to 4, which is a good value to use.
- Set a GRUB password so that the boot process cannot be modified by passing commands to GRUB (GRand Unified Bootloader). The purpose of the GRUB is to recognize file systems and load boot images, and it provides both menu-driven and command-line interfaces to perform these functions. Setting a password protects the ESX host from unauthorized modifications to the boot process. Follow these steps to set a GRUB password:
 1. Log on to the Service Console as root and edit the `/boot/grub/grub.conf` file.
 2. Add the line *password the password that you want to set* after the first uncommented (`#`) line and save the file.
 3. Type the following commands:
 - ◆ `chown root:root /boot/grub/grub.conf`
 - ◆ `chmod 600 /boot/grub/grub.conf`
- Protect against the root file system filling up, which can cause a denial-of-service attack on your ESX host. The `/home`, `/tmp`, and `/var/log` directories all have the potential for filling up the root partition. By creating separate partitions for them, as outlined in Chapter 3, “Building Your Virtual Environment,” you can prevent those partitions from affecting the root partition.
- Disable automatic mounting of USB devices when they are plugged into an ESX host to prevent users from running malicious code on the host. By default, automatic USB drive mounting is enabled. You can disable it by editing the `/etc/modules.conf` file and commenting out the `alias usb-controller` line by placing a `#` sign before it.

Did You Know?

The `esxcfg-auth` command is a front end for setting authentication and password information that is contained in various files on the Service Console. You can also change settings by editing those files directly, but it is easier with `esxcfg-auth`. You can also view authentication and password settings by using the `esxcfg-auth -p` (or `-p`) command. This displays the contents of various configuration files, including `/etc/pam.d/system-auth`, `/etc/nsswitch.conf`, `/etc/login.defs`, `/etc/yp.conf`, `/etc/nscd.conf`, `/etc/openldap/ldap.conf`, `/etc/ldap.conf`, `/etc/krb.realms`, and `/etc/krb.conf`.

Firewall

Here is a recommended security best practice for the built-in ESX firewall:

- By default, the built-in ESX firewall that protects the Service Console has very few ports opened up. Avoid opening any unnecessary ports unless they are needed for certain features and functionality to work properly (for example, NTP, iSCSI, NFS).

Networking

Here are recommended security best practices for virtual networking:

- Isolate the Service Console and VMkernel networks on a separate vSwitch. In addition, isolate the VLANs for the Service Console and VMkernel vSwitches so that they are segregated from the rest of the network. Physical segregation is preferred over virtual segregation. If this is not possible, consider deploying a network firewall to limit access to the Service Console and VMkernel networks.
- Make sure your vSwitch network labels are clearly and properly labeled. If you do not, someone may inadvertently connect a VM to a vSwitch on a network that it should not be on and thereby expose sensitive information. An example of this is someone connecting a VM to a public DMZ network that was meant to be on an internal network only.
- Configure all your vSwitches and port groups so that promiscuous mode, MAC address spoofing, and forged transmissions are disabled, as covered in Chapter 4, “Configuring Your Virtual Environment.” Only enable these settings if they are absolutely needed by an application running on your ESX hosts.
- Unfortunately, no security or auditing of vSwitches are built in to ESX and vCenter Server. As a result, anyone who has the appropriate privileges can move a VM to any vSwitch or port group at will. This can be a big security risk, and therefore you should carefully assign permissions to users and user roles to ensure they do not have access to do this. To prevent someone from changing vSwitch configurations, remove the Host, Configuration, Network Configuration permission; and to prevent someone from changing the NIC settings for a VM, remove the Virtual Machine, Configuration, Add or Remove Device, and Modify Device Settings from them. We talk more about permissions later in this chapter. In addition, if you want to provide security and audit logging on your vSwitches, look at some of the third-party applications that can provide this (as discussed later in this chapter).

ESXi Security Considerations

The design of ESXi is different from ESX. For example, there is no full Service Console (ESXi instead uses a small console based on Busybox), and subsequently there are some different

security recommendations for ESXi. Many of the security recommendations that apply to ESX will also apply to ESXi, except for those that have to do with the Service Console. Even though ESXi does not have a Service Console, it still has a small management component that should also be protected. Here are some of the security considerations that you should know about for ESXi:

- An ESXi host still has local accounts just like an ESX host, but the management of them must be done using the VI Client, because there is no Service Console to do it from. You should protect the root account just like with ESX and create separate user accounts for users who need to connect directly to the ESXi host. ESXi hosts cannot be configured to use AD to authenticate users.
- ESXi has a special tech support mode, which is an interactive command-line interface that is enabled by default and intended to be used for troubleshooting purposes. This mode should be disabled and not used unless you actually need it for troubleshooting and you know what you are doing. Although an attempt to access the mode is logged, most activity that occurs when using it is not logged. In addition, you can render your ESXi host unusable if you do not know how to use the commands properly. To disable this mode, follow these steps:
 1. Connect to the ESXi host with the VI Client and select the Configuration tab and then select Advanced Settings under Software.
 2. Select VMkernel in the left pane.
 3. Scroll down and find the VMkernel.Boot.techSupportMode setting and deselect it.
 4. Restart the ESXi host for the setting to take effect.

Securing Virtual Machines

It's been said that you should secure VMs as you would physical servers. Although this is true to a degree, you still need to take additional steps to protect them because a virtual environment is a lot different from a physical environment. So be sure to implement all security practices to the operating system and applications as you would a physical server, and then apply the best practices listed here for maximum security:

- Remove all unnecessary virtual hardware from the VM. This includes floppy drives and serial and parallel ports. In addition, go into the VM's BIOS (F2 at boot) and disable serial, parallel, and floppy devices in Advanced, I/O Device Configuration.
- Disable all unnecessary services and applications inside the VM's operating system. In addition to being potential security risks, these all consume memory and resources, both of which are in demand on an ESX host server.
- Do not enable any fancy screen savers on your guest OS. If you must use a screen saver, use one that is simple and effective, like the Windows login screen saver.

- Minimize the use of the remote console feature in the VI and web clients. If you do use it, close it as soon as you are done with it. The remote console provides power management and removable device controls, which could potentially allow a malicious user to bring down a VM. In addition, there is a performance impact on the ESX host when a remote console session is opened, which grows as more sessions are opened. If you can, use operating system native tools such as Windows Remote Desktop or SSH consoles when connecting to VMs.
- Disable copy and paste operations between the guest operating system and the VI Client remote console. When the remote console is used, nonprivileged users and processes running inside the VM can access the contents of the clipboard from the remote console window. If a user copies sensitive data into her workstation's clipboard before using the remote console, that information will be exposed to the VM. To disable this behavior, follow these steps:
 1. Power off the VM to be able to edit the configuration settings.
 2. Edit the settings for the VM, and then select the Options tab.
 3. In the left pane, under Advanced, select General, and then click the Configuration Parameters button.
 4. Click the Add Row button. A blank row will be added to the list of parameters.
 5. In the Name field, enter `isolation.tools.copy.disable`, and then enter `true` in the Value field.
 6. In the Name field, enter `isolation.tools.paste.disable`, and then enter `true` in the Value field.
 7. In the Name field, enter `isolation.tools.setGUIOptions.enable`, and then enter `false` in the Value field.
 8. When you have finished adding the rows, it should look similar to Figure 5.10. Click OK to save the parameters, and then click OK again to save the VM settings.
- Limit the amount of logging that a VM can output to avoid a denial-of-service attack in which a VMFS volume is filled up so that no more writes can occur. VMs write information to their `vmware.log` files that are located in the same directory as the VM. It is possible for users and processes inside the VM to be configured in such a manner that large amounts of data are written to the log file. As it grows, it consumes disk space on the VMFS volume, which can eventually fill it up and cause problems, especially if you have snapshots running. To avoid this situation, you can modify the log settings for the VMs to limit the size and number of the log files. Once a log file reaches the defined limit, it automatically writes to a new log file and preserves the old one. If the maximum number of log files reaches the defined limit, older ones are automatically deleted when new ones are created. To limit the total size and number of log files, add the following parameters to the VM's settings, as described previously:




Figure 5.10 Advanced configuration parameters being set to disable copy and paste operations

1. In the Name field, enter `log.rotateSize`, and then enter a numeric value in bytes in the Value field. For example, 100000 would set the maximum size of a log file at 100KB.
 2. In the Name field, enter `log.keepOld`, and then enter a numeric value in the Value field. This value is the number of old log files to keep. A recommended value is 10, and when combined with a 100KB max size will limit the VM to a maximum of 1MB of disk space that its logs can take up.
- Be sure to disconnect removable devices (CD-ROM and floppy) from the VM when you have finished with them, especially operating system installation media. Avoid using the host devices and instead use client devices and ISO datastores. If you use host devices, you risk multiple VMs trying to access the device when booting, and if one VM locks it, the others will have to wait, which can delay them from booting.
 - Use vCenter Server roles and permissions to limit the access and functions that users and administrators have to ESX hosts and VMs. Do not provide full access to users if it is not necessary. Only provide access to specific VMs and ESX hosts as required. In addition, go through the VM privileges and assign only what is needed. Avoid assigning configuration, provisioning, and inventory privileges, and instead assign just the ones needed to operate the VM (state and interaction).

Securing vCenter Server

vCenter Server is a powerful tool to centrally manage ESX/ESXi hosts and subsequently must be carefully configured to prevent unauthorized and unnecessary access. A user who has full access to vCenter Server also has full access to all the hosts and VMs managed by it.

Subsequently, an inexperienced user may inadvertently do or change something that could harm your environment. Therefore, you should never assign full access to vCenter Server if a user does not require it. vCenter Server has many built-in privileges and roles that can provide very granular access to VMs and hosts. In addition, vCenter Server integrates with AD to provide centralized user account management and to leverage its strong authentication mechanism.

Active Directory Integration

By default, vCenter Server is configured to use the local Windows Administrators group for its administrator role. Any user in that group is also a full administrator of VC and typically included in the local administrator account of the Windows OS. vCenter Server will automatically integrate with AD if the VC server is part of an AD domain, but no access from AD is provided by default unless a user in the AD is in the local Administrators group of the VC server. You can assign access to vCenter Server using either local users and groups or domain users and groups. vCenter Server also supports multiple AD domains as long as a trust exists between them or they are in the same forest.

If you have a large AD domain or forest, you can change some configuration settings to change the behavior of how vCenter Server searches AD. These settings are accessed inside vCenter Server by selecting Administration from the top menu and then selecting VirtualCenter¹ Server Management Server Configuration. The settings that can be changed, as shown in Figure 5.11, are listed here:

- **Active Directory Timeout.** This is the number of seconds that vCenter Server will spend searching a domain before it times out. By default, this is set to 60 seconds, which is usually sufficient unless you have a very large domain with many objects in it. If you find that your searches time out, increase this to a larger value.
- **Enable Query Limit.** This can either be enabled or disabled and is enabled by default with a value of 5000. This limits the number of users and groups that will be returned from a domain when setting permissions. This is limited to prevent the display of users and groups from taking too long. You can increase this number or disable the feature if you have a large domain and are not seeing all your users and groups. Alternatively, if you do not see it in the listing of users and groups, you can search for a user or group.




Figure 5.11 Active Directory configuration setting in vCenter Server

- **Enable Validation.** This is the time interval in minutes that VC checks AD to validate all permissions that are set in AD for domain users and groups are still valid. The default is 1 day (1440 minutes), which you may consider decreasing if you are using many users and groups and have frequent changes to your AD environment. When VC performs the validation, it verifies that every user and group used in vCenter Server permissions still exists in AD, and if it does not it automatically removes the permission. It is not recommended to set this value too high or to disable the validation entirely.

One important thing to note is that VC does not use the unique AD SIDs (system identifiers) when it assigns access to an AD user or group. All AD permissions are assigned using the following format: *AD domain\user or group name* (for example, *acme\mjones*). Therefore, if a user or group is deleted from AD, a new user or group could be created with that same name, and even though it would be completely different in AD because it will have a new SID, it will still appear to be the same user or group to VC. Therefore, it is important to make sure to have the Enable Validation setting configured to ensure that access to VC is not granted to a user and group that should not have it.

Privileges and Permissions

Permissions allow you to assign privileges to users or groups to perform certain actions and make changes to objects inside of vCenter Server. vCenter Server permissions affect only users who log on to vCenter Server, and do not apply when logging on to an ESX host directly. Permissions can be very granular, and there are dozens of them that can be assigned to control specific access to the various objects inside vCenter Server. Permissions can be set on many different levels, from the very top datacenter level down to an individual VM. You can also control whether they propagate down from the level they are set to other objects below in the hierarchy. You can also set multiple permissions for a user or group to control different access levels for different levels of objects inside VC. For example, you could grant a user read access at the top datacenter level so that the user can read all the objects in the datacenter but set higher access to certain specific objects like ESX hosts or VMs. Permissions work with roles, which allow you to create a set of privileges as a role to assign to a user and group. You cannot assign privileges directly to a user or group, but must instead assign privileges to a role and then assign a user to that role. Users can be assigned multiple roles with different privileges on different objects in vCenter Server. However, if a user has different levels of permissions on the same object, the most restrictive permissions apply. For example, if a user has read-only permission on a host and full administrator permissions, the end result is that he will have only read-only permissions.

Privileges are grouped into categories, which are further broken down into subcategories and then down to the individual privilege. You can select privileges at any level of the category hierarchy and the privileges below it are included with whatever you select, as shown in Figure 5.12.




Figure 5.12 Selecting privileges when configuring a role

For example, selecting the Virtual Machine category selects all the subcategories and privileges under it. If you want to be more granular, you can uncheck the items under the category that you selected to achieve the desired privileges for the user or group. The privileges and categories available to choose from are listed in Table 5.2.

Table 5.2

All vCenter Server 2.5 Privileges

Top-level Category	Subcategory	Privileges
Global	None	Manage Custom Attributes Set Custom Attribute Log Event Cancel Task Licenses Diagnostics Settings VC Server Capacity Planning Script Action Proxy Disable Methods Enable Methods Service Managers
Folder	None	Create Folder Delete Folder Rename Folder Move Folder
Datacenter	None	Create Datacenter Remove Datacenter Rename Datacenter Move Datacenter
Datastore	None	Rename Datastore Remove Datastore Browse Datastore Remove File File Management
Network	None	Remove

Top-level Category	Subcategory	Privileges
Host	Inventory	Add Stand-Alone Host Create Cluster Add Host to Cluster Remove Host Move Cluster/Stand-Alone Host Rename Cluster Remove Cluster Modify Cluster Move Cluster
Host	Configuration	Systems Management Connection Maintenance Virtual Machine Auto-start Configuration HyperThreading Storage Partition Configuration Security Profile and Firewall Memory Configuration Network Configuration Advanced Settings System Resources Change SNMP Settings Change Date Time Settings Change Settings Query Patch Firmware
Host	Local Operations	Add Host to vCenter Server Manage User Groups Create Virtual Machine Delete Virtual Machine
Host	CIM	CIM Interaction
Virtual Machine	Inventory	Create Remove Move

continues...

Table 5.2 continued

Top-level Category	Subcategory	Privileges
Virtual Machine	Interaction	Power On Power Off Suspend Reset Answer Question Console Interaction Device Connection Configure CD Media Configure Floppy Media Tools Install Defragment All Disks
Virtual Machine	Configuration	Rename Add Existing Disk Add New Disk Remove Disk Raw Device Host USB Device Change CPU Count Memory Add or Remove Device Modify Device Settings Settings Change Resource Upgrade Virtual Hardware Reset Guest Information Advanced Disk Lease Swap Placement DiskExtend
Virtual Machine	State	Create Snapshot Revert to Snapshot Remove Snapshot Rename Snapshot

Top-level Category	Subcategory	Privileges
Virtual Machine	Provisioning	Customize Clone Create Template from Virtual Machine Deploy Template Clone Template Mark as Template Mark as Virtual Machine Read Customization Specifications Modify Customization Specification Allow Disk Access Allow Read-Only Disk Access Allow Virtual Machine Download Allow Virtual Machine Files Upload
Resource	None	Assign Virtual Machine to Resource Pool Apply Recommendation Create Pool Rename Pool Modify Pool Move Pool Remove Pool Migrate Relocate Query VMotion
Alarms	None	Create Alarm Remove Alarm Modify Alarm
Tasks	None	Create Update
Scheduled Task	None	Create Task Remove Task Run Task Modify Task

continues...

Table 5.2 continued

Top-level Category	Subcategory	Privileges
Sessions	None	View and Terminate Sessions Validate Session Global Message Impersonate User
Performance	None	Modify Intervals
Permissions	None	Modify Role Reassign Role Permissions Modify Permission
Extension	None	Register Update Unregister
VMware Update Manager	Manage Baseline	Assign Baseline Manage Baseline
VMware Update Manager	Configure	Configure Service
VMware Update Manager	Manage Updates	Remediate Updates Scan Updates View Compliance Status

You'll notice that a lot of privileges can be assigned, especially on a host and VM. Some actions in vCenter Server require multiple privileges to be assigned for the action to be completed. For example, to create a VM requires the following privileges:

- Virtual Machine, Inventory, Create
- Virtual Machine, Inventory, Add New Disk (or Add Existing Disk or Raw Device)
- Resource, Assign Virtual Machine to Resource Pool (optional)
- Read-only access to the datacenter that the VM will reside in

Many of the privileges are self-explanatory, but trying to figure out which ones are needed for certain actions can be confusing. Create a test user account and assign/remove privileges to it and then log on as that user and see whether you can complete the action. If you cannot complete the action, go back and add a privilege and try again. It is helpful sometimes to select a whole category of privileges and then remove them one by one until you achieve the desired permission level.

Permissions can be set on many different objects in vCenter Server, including the following:

- Hosts and Clusters
- Datacenter
- Cluster
- Resource pool
- Host
- Folder
- Virtual machine
- Templates

You can view the permissions of a particular object in vCenter Server by clicking it and selecting the Permissions tab, as shown in Figure 5.13.




Figure 5.13 Viewing the permissions on a VM object

To assign permissions to users or groups, follow these steps:

1. Select an object in vCenter Server (for example, Virtual Machine) in the left pane.
2. In the right pane, select the Permissions tab.
3. Right-click in the right pane and select the Add Permission option.
4. In the Users and Groups section, click the Add button.
5. In the Select Users window, select an AD domain or (server) to use local users and groups, as shown in Figure 5.14. Once you select a domain, it will list the users and groups for you to select from. You can choose an individual user/group or multiple ones by holding the Ctrl or Shift keys when selecting. As you select users and groups, they will appear in the Users and Groups fields separated by a semicolon. A drop-down selection opens to display users or groups and displays everything in alphabetical order. You can also search for users/groups by entering text into the input box and clicking the Search button. After you have selected the users/groups, click the Add button to add them.
6. Select a role from the Assign Roles section and choose whether you want to propagate the permissions to child objects. After you have selected the users/groups and role, as shown in Figure 5.15, click OK to save the permissions.




Figure 5.14 Selecting users and groups to configure a permission




Figure 5.15 A completed permission with an assigned user and role

Configuring and Using Roles

A role is a set of privileges that can be assigned to a user or group. Roles can be customized to include or exclude any of the many privileges that exist in vCenter Server. Privileges provide the right to perform a specific action. Roles are a collection of privileges, and permissions are the application of a role to a user or group for an object. If a role is granted on an object for a user (for example, ESX host), the user is able to see all information about that object but will have permission to perform only whatever privileges he is granted and nothing more. Many of the actions that will show for a full administrator will be masked or grayed out if the user does not have permissions for that action on that object. In addition, if you apply permissions at a low level of a hierarchy (for example, VM), users will not see the other objects (VMs) that they do not have permissions to see unless those permissions are specifically applied at a higher level. For example, if you assign a user the permission to manage a specific VM and that is all, when he logs on to vCenter Server he will see the datacenter object and the VM. That user will not see any other VMs, the ESX host that the VM is on, clusters, resource pools, and so forth.

vCenter Server comes with some preconfigured roles that you can use, or you can create your own custom roles. These roles fall into two categories: system roles that are permanent and cannot be deleted or changed, and sample roles that are configured for certain types of access and can be modified and deleted. To access roles, click the Administration button in vCenter Server and then select Roles. Table 5.3 lists the preconfigured roles in vCenter Server.

Table 5.3

vCenter Server Preconfigured Roles

Role	Type	Capabilities
No Access	System	Cannot view or change the assigned object. VI Client tabs associated with an object display without content. This is the default role for all users except those users in the Administrators group.
Read-only	System	View the state and details about the object. View all the tab panels in the VI Client except the Console tab. Cannot perform any actions through the menus and toolbars.
Administrator	System	All privileges for all objects. Add, remove, and set access rights and privileges for all the vCenter Server users and all the virtual objects in the VMware Infrastructure environment. This is the default role for all members of the Administrators group.

continues...

Table 5.3 **continued**

Role	Type	Capabilities
Virtual Machine User	Sample	Perform actions on VMs only. Interact with VMs, but not change the VM configuration. This includes the following: ■ All privileges for the Scheduled Task privileges group ■ Selected privileges for the Global Items and VM privileges groups ■ No privileges for the Folder, Datacenter, Datastore, Network, Host, Resource, Alarms, Sessions, Performance, and Permissions privileges groups
Virtual Machine Power User	Sample	Perform actions on the VM and resource objects. Interact and change most VM configuration settings, take snapshots, and schedule tasks. This includes the following: ■ All privileges for Scheduled Task privileges group ■ Selected privileges for Global Items, Datastore, and VM privileges groups ■ No privileges for Folder, Datacenter, Network, Host, Resource, Alarms, Sessions, Performance, and Permissions privileges groups
Virtual Machine Administrator	Sample	Perform actions on global items, folders, datacenters, datastores, hosts, VMs, resources, alarms, and sessions. This includes the following: ■ All privileges for all privilege groups, except Permissions
Datacenter Administrator	Sample	Perform actions on global items, folders, datacenters, datastores, hosts, VMs, resources, and alarms. Set up datacenters, but with limited ability to interact with VMs. This includes the following: ■ All privileges for Folder, Datacenter, Datastore, Network, Resource, Alarms, and Scheduled Task privileges groups ■ Selected privileges for Global Items, Host, and VM privileges groups ■ No privileges for Session, Performance, and Permission privileges groups

Role	Type	Capabilities
Resource Pool Administrator	Sample	Perform actions on datastores, hosts, VMs, resources, and alarms. Provides resource delegation and is assigned to resource pool inventory objects. This includes the following: ■ All privileges for Folder, VM, Alarms, and Scheduled Task privileges groups ■ Selected privileges for Global Items, Datastore, Resource, and Permissions privileges groups ■ No privileges for Datacenter, Network, Host, Sessions, or Performance privileges groups
Consolidated Backup User	Sample	Perform actions on VMs that are needed to use Consolidated Backup to back up the VM. This includes the following: ■ Disk lease configuration ■ State to create or remove snapshots ■ Provisioning to allow read-only disk access and allow VM download

When you select the Roles tab, you will see a list of configured roles. If you click a role, you will see the users associated with that role and the level that the access is assigned at, as shown in Figure 5.16.




Figure 5.16 Viewing users and groups assigned to a role

You can add, rename, remove, and edit roles, and clone them when creating new roles so that you can make a copy of an existing role and modify it rather than starting from scratch. To create a role, follow these steps:

1. In vCenter Server, click the Administration button, and then select the Roles tab.
2. Click the Add Role button (or optionally select an existing role and click the Clone Role button).

3. Provide a name for the role, and then select the privileges to assign to that role. If you select a category, everything under that category is automatically selected. You can deselect privileges if you do not want to include them.
4. After you have selected all the privileges that you want to grant to that role, click OK to save it.
5. To assign users to the role to create a permission, click a Permission tab for an object and add a permission and assign a user to the role.

Roles are a great way for assigning access inside vCenter Server and should definitely be used to ensure that more access than is needed is not assigned to vCenter Server users. Some best practices for using roles are listed here:

- Never assign more privileges than are needed to a user. Create roles that are customized to a user's specific requirements. For example, to create a role for an operations team that is responsible for monitoring VMs, create a role that allows VM interactions only (for example, power on, power off, reset, and console interaction). That way they can look at the console of a VM to see what is happening and power cycle a VM if needed.
- Never assign more permissions than are needed to a user. Assign permissions at the required level and avoid assigning them at the highest level, if possible. For example, if a network group needs access to only one or two VMs on a host, only assign permissions at the specific VM level that they need access to.
- Certain privileges can be very harmful to your hosts and should be assigned to users only if absolutely required. This includes any privilege that allows a user to delete, rename, remove, or create anything that might cause data loss or datastores to be filled up (which can cause a denial-of-service attack on your VMs; for example, snapshot creation).

vCenter Server Security Best Practices

Here are some security best practices that you should follow for vCenter Server to ensure it is configured for maximum security:

- Don't forget about the database that is used to store configuration information for vCenter Server. If you do not properly secure and protect your database, that lack of security may allow someone to insert data into the tables and authorize himself to log on to vCenter Server as a full administrator or cause problems in your environment. Ensure you use strong database passwords and restrict access to the tables that vCenter Server uses. For maximum security, do not use a shared database server; instead, isolate the database on a dedicated database server.

- If you do set up Single Sign-on to log on to vCenter Server, as outlined in Chapter 4, take precautions to prevent someone from accessing an unlocked workstation and connecting to vCenter Server using the credentials of the currently logged-on user. If you log on to shared workstations, make sure you log off. If you do not, all someone has to do is open a browser and access the default page on any ESX or vCenter Server to download and install the VI Client on that workstation, and then he can log on to vCenter Server as you and access anything you have access to. Anytime you leave your workstation, make sure you lock it. You can make a shortcut for doing this in one mouse click or use the Windows and L keystroke combination. Also, make sure you have your workstation set to automatically lock at a maximum of 15 minutes of idle time as protection in case you forget to lock it when you walk away.
- Avoid assigning the Console Interaction privilege to users and instead have them access as guests via traditional remote access methods such as Remote Desktop and SSH. Only allow administrators and server operations staff to access the VM's remote console.
- Remove the local Windows Administrators group from the vCenter Server Administrator role and instead create a local account and assign it the Administrator role. By doing this, you limit the access granted to vCenter Server by users who do not require it. Typically, a local Administrators group also includes a Domain Administrators group, and in most cases you do not want your domain administrators to be full administrators of vCenter Server.
- Disable the web access service on your vCenter Server if you do not plan to use it. To do this, disable the VMware Infrastructure Web Access Windows service on the vCenter Server.

Third-Party Security Products and Security Appliances

Many third-party security products can enhance the security of your VI3 environment and provide better control, reporting, and auditing. Some are free products, and others require that you purchase a license for them (but most have evaluation versions available that you can download and try out in your environment before deciding to purchase them). If your environment requires more advanced security or you just want to be more aware of what is happening in your environment, check out some of the products listed here:

- **Reflex Systems Virtual Management Center (not free).** Reflex's Virtualization Management Center (VMC) enables next-generation datacenters to enforce IT policies, ensure compliance with government mandates, and manage and protect virtual servers, desktops, and networks across multiple platforms. VMC provides a single

authoritative visual interface, central management, and security for heterogeneous virtual environments. By combining the centralized event database, virtual infrastructure integration, and analysis engines with a robust visual interface, it is possible to administer, audit, secure, and monitor complex, dynamic, virtual infrastructures. This results in better network and event visibility for a faster and more efficient management and security response. Through extensive real-time and historical visual reporting and revision control, Reflex VMC gives administrators the tools they need to efficiently meet stringent compliance requirements and provide secure management for virtualized datacenters.

- **Catbird V-Security (not free).** Via its V-Agent virtual appliance, Catbird secures the hypervisors and VMs from vulnerabilities and attacks. A stateless agent residing on the virtual network itself, the noninvasive, self-contained V-Agent has a view of the network impossible to see from traditional devices outside of the host. Catbird's sweeping set of services brings hypervisor security, VM/guest OS monitoring, policy compliance, and VM sprawl management under one single web-based management interface. Catbird delivers VMware security as a service, not via an expensive appliance. Security-as-a-Service exploits the true value of the service model, with Catbird's Security Operations Center (SOC) continuously and vigilantly monitoring the security posture of the virtual deployment and taking instant action against weakness or hostile activity.
- **Tripwire ConfigCheck (free).** Tripwire ConfigCheck is a free utility that rapidly assesses the security of VMware ESX 3.0 and 3.5 hypervisor configurations compared to the VMware Infrastructure 3 security hardening guidelines. Developed by Tripwire in cooperation with VMware, Tripwire ConfigCheck ensures ESX environments are properly configured and provides the necessary steps toward full remediation when they are not.
- **Tripwire Enterprise 7.5V (not free).** Tripwire Enterprise for VMware ESX provides visibility across the VMware virtual infrastructure, enabling continuous configuration control for your virtual environment. In the rush to virtualize, organizations often forget to hold these virtual environments accountable to the same regulatory and security standards as their physical counterparts. Tripwire closes this gap by integrating with VMware's Virtual Center to automatically discover virtual infrastructure objects. Tripwire then performs proactive assessments, using the most comprehensive set of configuration assessment policies, to ensure virtual objects are configured securely. For remediation of misconfigured virtual infrastructure objects, Tripwire provides out-of-the-box prescriptive guidance. Finally, by monitoring for unauthorized changes that could negatively affect system integrity, regulatory compliance, and IT service availability, Tripwire's detection of unauthorized VMs created in the production environment helps ensure continuous configuration control. With Tripwire's policy-based security

and control of the virtual infrastructure, organizations increase their security posture, rapidly attain compliance, and reduce the inherent risks of virtualization.

- **Check Point VPN-1-VE (not free).** VPN-1 VE (Virtual Edition) protects your virtual systems with the same best-of-class security that you rely on for your physical network. It enables you to segregate virtual systems from each other and from external threats. VPN-1 VE extends your open choice to encompass Check Point appliances and software, “Secured by Check Point” appliances, and virtual systems (all managed by a single interface to ensure consistent, efficient security management).
- **Configuresoft Compliance Checker (free).** Compliance Checker is a free, downloadable tool that provides a real-time compliance check for multiple VMware ESX host servers at a time. This product is a robust utility providing detailed compliance checks against both the VMware hardening guidelines and the CIS benchmarks for ESX. Unlike other free tools in the market, Compliance Checker is a fully functional product. For example, you can print the reports that Compliance Checker produces, and can run compliance checks across multiple ESX servers at once.
- **Configuresoft Enterprise Configuration Manager (ECM) (not free).** ECM for Virtualization extends the value of ECM to virtual environments. ECM for Virtualization helps organizations to effectively manage and secure virtual environments by providing enterprise visibility, control, and compliance across the entire VMware environment. ECM for Virtualization provides a central console to view the security posture of their virtual environments. Organizations can ensure compliance with various standards such as CIS benchmarks for VMware ESX server, VMware hardening guidelines, GLBA, HIPAA, and the Sarbanes-Oxley Act. ECM for Virtualization provides at a glance a view of all guests and their build details to ensure that all the deployed virtual servers are consistent and follow current IT standards. IT operations can prevent virtual server sprawl by having control over virtual environments. ECM for Virtualization helps IT to automate their common IT tasks and achieve a more efficient IT operational state.
- **Altor Network’s Virtual Network Security Analyzer (not free).** Altor Network’s Virtual Network Security Analyzer (VNSA) is the only solution on the market that delivers real-time visibility, troubleshooting, and analysis of ongoing virtual network activity. Unlike physical network security and monitoring solutions that detect only activity on physical networks, Altor’s VNSA detects and analyzes virtual network traffic to troubleshoot problems, baseline network activity, audit for inappropriate protocols, and surface anomalies. With the VNSA, datacenter administrators can keep their virtual networks running smoothly, gain insight into virtual network activity, and develop and apply compliance-related security policies.
- **Altor Network’s Virtual Firewall (not free).** Altor Network’s Altor VF virtual firewall is purpose built to meet the unique security challenges of the virtual environment. Administrators can now control the virtual network by enforcing a rule-based policy

for each VM. Because the Altor VF was designed with virtualization features in mind, it synchronizes automatically with VMware's vCenter Server and secures VMotion while maintaining the open connections of VMs "in flight."

- **Third Brigade Deep Security (not free).** Third Brigade's intrusion defense system brings proven security approaches, including deep packet inspection, stateful firewall, application firewall, and intrusion detection and prevention, down to individual networked hosts and VMs within virtual environments. This software solution has been architected for enterprises that recognize the need to further enhance their security posture to protect mission-critical IT assets, including physical and virtual servers, from known and zero-day attacks.
- **Apani EpiForce VM (not free).** Apani EpiForce VM helps large companies accelerate adoption of VMware ESX Server by securing both VMs and physical servers and endpoints from a single platform. EpiForce VM is a software-based, cross-platform, server-isolation solution that provides two critical security services: logical security zoning and policy-based encryption of data in motion. EpiForce VM is a distributed, centrally managed solution that is transparent to users, applications, and infrastructure, making it quicker to deploy and less costly to manage than hardware- or appliance-based solutions.
- **X-M0n0wall virtual firewall appliance (free).** A small (5MB) and effective Linux-based firewall that is available from VMware's Virtual Appliance Marketplace that can be used to create firewalls between your vSwitches on your host servers.

Additional Security Resources

Here are some additional security resources that are definitely worth reading to find out more about VI3 security. They consist of security benchmarks, white papers, VMworld presentations, and VMware web pages that focus on security and compliance:

- CIS ESX Server 3.x Security Benchmark
www.cisecurity.org/tools2/vm/CIS_VMware_ESX_Server_Benchmark_v1.0.pdf
- CIS Virtual Machine Security Benchmark
www.cisecurity.org/tools2/vm/CIS_VM_Benchmark_v1.0.pdf
- ESX Server Security Technical Implementation Guide
http://iase.disa.mil/stigs/stig/esx_server_stig_v1r1_final.pdf
- Managing VMware vCenter Server Roles and Permissions
www.vmware.com/pdf/vi3_vc_roles.pdf

- DMZ Virtualization with VMware Infrastructure
www.vmware.com/files/pdf/dmz_virtualization_vmware_infra_wp.pdf
- Security Design of the VMware Infrastructure 3 Architecture
www.vmware.com/pdf/vi3_security_architecture_wp.pdf
- VMware Infrastructure 3 Security Hardening
www.vmware.com/files/pdf/vi35_security_hardening_wp.pdf
- Using the Secure Technical Implementation Guide (STIG) with VI3 (VMworld 2007 presentation, free registration required)
<http://vmworld.com/vmworld/mylearn?classID=11680>
- Proven Practice: 20 Questions from IT Security Professionals
<http://viops.vmware.com/home/docs/DOC-1141>
- VMware ESX Server - Providing LUN Security
www.vmware.com/pdf/esx_lun_security.pdf
- Security in a Virtualized Environment (VMworld 2007 presentation, free registration required)
www.vmworld.com/vmworld/mylearn?classID=11276
- Security Architecture Design and Hardening VI3 (VMworld 2007 presentation, free registration required)
www.vmworld.com/vmworld/mylearn?classID=11047
- How to Achieve Security and Satisfy Compliance (VMworld 2007 presentation, free registration required)
www.vmworld.com/vmworld/mylearn?classID=11461
- Best Practices for Surviving Regulatory Compliance (VMworld 2007 presentation, free registration required)
www.vmworld.com/vmworld/mylearn?classID=11450
- Achieving Compliance in a Virtualized Environment
www.vmware.com/files/pdf/technology/compliance_virtualized_environment_wp.pdf
- VMware Compliance Center
www.vmware.com/technology/security/compliance/resources.html
- VMware Security Center
www.vmware.com/security/
- VMware's Security Response Policy
www.vmware.com/support/policies/security_response.html

Summary

Security always comes at a cost, and that cost is usually convenience. When you apply strict security standards in your environment, administration tasks are sometimes more difficult to complete and may require more steps. But at the same time, it also makes it more difficult for an unauthorized person to try to access the data on your virtual hosts. In a virtual environment, protecting the host is very important because the risk is increased exponentially because of the many VMs running on that host that can potentially be compromised if the host is compromised. Because of this, many security groups are usually wary of virtual hosts because of the increased risk. Therefore, take the time to sit down with them and explain how security works in a virtual environment and tell them about the extra steps you took to further protect your virtual hosts and VMs.

This chapter covered many different ways to enhance the security of VI3. It's important to address all areas and not just one or two when tightening your security. By doing this, you form a tight shield to protect your whole environment. If you skip one area, it may (no matter how secure the other areas) be the single chink in your armor that could compromise your environment.

Endnotes

1. The use of the term *VirtualCenter* here refers to the old name, which still applies to VirtualCenter version 2.0.x and is still present inside the application in vCenter Server 2.5 (because the software has not yet been updated to reflect the new name).

Chapter 6

Populating Your Virtual Environment

Now that you've installed, configured, and secured your virtual environment, it's time to start populating it with VMs. In many cases, you will want to convert your current physical servers to VMs instead of creating new VMs and then migrating applications and data to them. This is possible using a process called P2V (physical to virtual) migration, which will allow you to quickly and easily move your physical servers into your virtual environment. In addition, you can use some time-saving methods to quickly create new VMs using templates and ISO files. In this chapter, we cover everything you need to know to start populating your environment and some best practices that you can use for creating VMs and converting physical servers.

Creating Virtual Machines

Before you get started creating VMs, take care to ensure that you know the resource requirements needed for each VM. Otherwise, you might find that your host resources may be used up quicker than you expect. When doing this, look at the actual requirements that the application on the VM needs to function properly. Rather than give a VM two vCPUs because that's what it had as a physical server, you might consider giving it one instead and see how it performs. With memory, you should assign only what a VM is going to use. Look at the memory statistics for your existing servers to see how much memory they are using at peak usage and assign memory to the VM accordingly. The same thing goes for creating virtual disks. If your physical server had a 72GB disk because that was the smallest drive you could get and you are using only 8GB of disk space on it, create your virtual disks smaller. Doing this will conserve valuable and expensive disk

space, especially if you are using shared storage. And don't worry if you find that you need more later; an advantage of a virtual environment is that you can easily add more resources to a VM as needed.

Creating a New Virtual Machine

To create a new VM, follow these steps:

1. Right-click a host or resource pool or cluster if you are using vCenter Server and select the New Virtual Machine option, which will launch the New VM Wizard.
2. You can choose between Typical and Custom for the wizard type. Most of the options are the same between them except that Custom gives you more options for configuring the VM's virtual disk. With Typical, you get to choose only the virtual disk size, and that's it. Custom gives you the following additional options:
 - a. Choosing a SCSI adapter type (BusLogic or LSI Logic). LSI Logic is most commonly used with Linux and Windows XP/2003, and BusLogic is used with older operating systems like Windows NT/2000 that do not have a driver for LSI Logic.
 - b. Instead of creating a virtual disk, you can choose to use an existing virtual disk, use a raw device mapping (RDM), or not to create a virtual disk at all. (Earlier versions of ESX required you to create a virtual disk.)
 - c. Choosing a location for the virtual disk file, which can be either stored with the VM or on a different datastore.
 - d. Choosing a virtual device node. This is the SCSI ID of the disk, which consists of the controller number and device number (for example, 0:0).
 - e. Choosing the disk mode, you can optionally select for the disk to be independent, which means it will not be included in snapshots. You can also choose if the disk is persistent or nonpersistent. You would typically want to use persistent disks because all changes to nonpersistent disks are discarded when a VM is powered off, which is useful for kiosks, training workstations, and call centers.
3. After you have chosen your wizard type and clicked Next, you will need to select a name and location for your VM, as shown in Figure 6.1. The name you choose is used to create the VM's directory, and most of the files within that directory will use that name with different file extensions (for example, VM1.vmx, VM1.vswp). This name is purely used by ESX and does not have to be the same name used in the guest operating system (although it is recommended to avoid confusion). It's a good idea to not use spaces in your VM names, which can make scripting and CLI commands more difficult. The location is the folder name that you want the VM to be in, and you can always move the VM to a different folder later. After you have entered a name and selected a location, click Next.




Figure 6.1 Creating a new VM: Choosing a name and location

4. Optionally, if you have a resource pool configured, you will be prompted to select one for the VM to be in. You can select either the host server (root) or one of the child resource pools, and then click Next.
5. Next choose a datastore that you want to store the various files that make up the VM, as shown in Figure 6.2, and click Next. If you selected Custom earlier, this step is specifying where your VMX file and other metadata will be stored. You will have the opportunity to specify a separate location for the virtual disk file later in the wizard.
6. Choose a guest operating system for your VM, as shown in Figure 6.3, and click Next. This will not install a guest operating system, but will instead change some of the defaults in the wizard depending on which one is selected. In addition, the guest operating system that you choose will change how the host handles the VM because it will optimize scheduling and the handling of the VM according to the guest operating system chosen. Therefore, choose the exact operating system that you plan to install on the VM, including the correct edition/version and 32 or 64 bits. If you choose to install a different guest operating system later, you can always change this setting by editing the VM's settings. If your operating system is not listed, choose the closest match or choose Other.




Figure 6.2 Creating a new VM: Selecting a datastore




Figure 6.3 Creating a new VM: Choosing a guest operating system

7. Choose the number of virtual CPUs that the VM should have, as shown in Figure 6.4, and click Next. Don't give the VM more vCPUs than it needs. Most servers work best with just one on a virtual host.




Figure 6.4 Creating a new VM: Choosing the number of virtual CPUs

8. Choose the amount of memory (RAM) to assign to the VM in 4MB increments, as shown in Figure 6.5, and click Next. The default amount is set based on the guest operating system that you choose in Step 6. This setting can range from 4MB up to 64GB; it is possible to assign more memory to a VM than the host physically has (not recommended though; the difference is made up by disk swapfiles). It's best to keep this at a minimum also. Only give the VM what it will actually need for the application that it is running. You can always change it later if needed.
9. Choose the number of virtual NICS (0-4) to assign to the VM, the networks that they should be connected to, the NIC type, and whether it should be connected when the VM is powered on, as shown in Figure 6.6, and then click Next. The networks displayed are the ones that have been configured on your vSwitches. The adapter types that are displayed are based on the guest operating system selected in Step 6. In most cases, you will have only one type to choose from (Flexible). More detail on the various types is in the "Virtual Networking" section in Chapter 2, "Planning Your Virtual Environment."




Figure 6.5 Creating a new VM: Choosing the amount of memory

10. If you chose the Typical wizard type, you will see a Disk Capacity screen that lets you choose the size of the virtual disk to create on the datastore that you selected in Step 5. After you select a capacity, click Next to continue; you can skip to Step 14.
11. If you chose the Custom wizard type, a Storage Adapter Type screen will display, as shown in Figure 6.7. Select between the BusLogic and LSI Logic SCSI adapter types and click Next. In general, the LSI Logic adapter is preferred because it sometimes offers slightly better performance depending on the guest operating system and disk I/O patterns. Some older operating systems (for example, Windows NT/2000) did not include the driver for the LSI Logic adapter, so it was easier to use BusLogic instead.
12. If you chose the Custom type, a screen will display letting you choose a virtual disk option for the VM, as shown in Figure 6.8. You can choose from one of the four options:
 - a. Create a New Virtual Disk. This option simply creates a new virtual disk to be used with the VM. If you choose this option and click Next, you will get a Disk Capacity screen that lets you choose the size of the virtual disk to create and choose to store the virtual disk file on the datastore that you selected for the VM in Step 5 or on another datastore.

Creating Virtual Machines




Figure 6.6 Creating a new VM: Choosing virtual NICs




Figure 6.7 Creating a new VM: Choosing a storage adapter (Custom only)

- b. Use an Existing Virtual Disk. This option lets you choose an existing virtual disk file to use with the VM. If you choose this option and click Next, you will get a Select an Existing Disk screen where you can browse through datastores and directories to select an existing virtual disk file.
 - c. Raw Device Mappings. This option lets you choose to use an RDM as your virtual disk file. If you choose this option, you will be presented with several different screens to select an RDM. The first screen is the Select Target LUN screen, which displays a list of all unused LUNs for you to choose from. If you do not see your LUN displayed, go back and do a rescan of your HBA controllers. Next a Select Datastore screen is displayed, which gives you two options for storing the LUN mapping file. You can choose to store it on the datastore that you selected for the VM in Step 5 or on an alternative datastore. Next a Compatibility Mode screen is displayed that lets you choose between physical or virtual compatibility modes. Physical will allow the guest to access the SAN directly, but does not allow snapshots to be taken of the RDM. If you choose Virtual, snapshots are allowed.
 - d. Do Not Create Disk. This option will not create a virtual disk for your VM. If you choose this option, you can always create a virtual disk later if needed.
13. If you chose the Custom type, a screen will display (if you chose not to create a virtual disk, this screen will not display) letting you choose a virtual device node and the disk mode, as shown in Figure 6.8. The virtual device node is the SCSI ID of the disk, which consists of the controller number and device number (for example, 0:0). The disk mode can be set to Independent, which means it will not be included in snapshots. You can also choose whether the disk is persistent or nonpersistent. You would typically want to use persistent disks because all changes to nonpersistent disks are discarded when a VM is powered off. After you select the virtual device node and disk mode, click Next.
14. At the Ready to Complete screen, a summary of all the VM options that were selected displays, as shown in Figure 6.9. You can select the check box to edit the VM settings before you continue, which allows you to edit the settings and remove or add devices as needed. Click Continue to edit the VM's settings or click Finish to create the VM.
15. When you click finish, you will see the Create Virtual Machine task in the VI Client. When it completes, your VM is ready to be powered on.

Now that you have created your VM, you can power it on and install an operating system on it. Next we cover using ISO files and how to map host and client CD-ROM drives to your VM so that you can install your operating system on your VM.

Creating Virtual Machines




Figure 6.8 Creating a new VM: Choosing a virtual device node and disk mode (Custom only)




Figure 6.9 Creating a new VM: Ready to Complete summary screen

Using ISO Files and Other Removable Devices with Your VM

ISO files are an image of an optical disk (CD or DVD) that are in the conventional ISO 9660 format that can be used with your VM's virtual CD/DVD drive instead of using traditional physical media. An ISO file is basically an archive file similar to a zip file but without the file compression and can be any size, from as small as a few megabytes to several gigabytes. Installing operating systems and applications on your virtual servers is much simpler and faster when using ISO files. ISO files are easy to mount and use on your VMs, and you do not have to worry about physical media imperfections or trying to find a specific physical CD/DVD when you need it.

Many operating systems and applications can be downloaded in ISO format, but you can also create your own ISO files from physical media or by using an application that will build an ISO file from whatever files you select. Some of the many applications that will let you create and build ISO files are listed here:

- dd (Linux command, which is included in the ESX Service Console)
- LC ISO Creator (free utility, www.lucersoft.com/files/free/LCISOCreator.zip)
- Imgburn (free utility, www.imgburn.com/)
- ISO Recorder (free utility, <http://isorecorder.alexfeinman.com/isorecorder.htm>)
- Win Image (shareware, www.winimage.com/)
- Ultra ISO (paid application, www.ezbsystems.com/ultraiso/)
- Magic ISO Maker (paid application, www.magiciso.com/)
- Nero (paid application, www.nero.com)

Let's go over creating an ISO file from a physical CD/DVD using dd.exe:

1. Log in to the ESX Service Console.
2. Insert the physical CD/DVD into the ESX host.
3. Create a temp directory on a VMFS volume and switch to it so that you do not fill up the disk on your root file system.
4. The syntax for the dd command is `dd if=input file to read from of=output file to write to bs=byte size to copy at once`. So, to create an ISO image from the disc in your CD/DVD drive, type in `dd if=/dev/cdrom of=/tmp/win2003-install.iso bs=32k`.
5. This will create an ISO file that you can mount on your VM's CD/DVD drive.

Now that you know how to create ISO files, let's cover how to use them. It's a good idea to create a central ISO repository that can be used by all of your ESX/ESXi hosts to map your VM's CD/DVD drive to when needed. The easiest way to do this is to create an ISO directory on a shared VMFS volume that all your hosts can access. If you have many ISO files, this might be undesirable, however, because it uses up valuable disk space on your shared storage

device. An alternative is to store them on a server that can provide NFS services so that all your hosts can access it. Here are the steps to set this up on a Windows Server 2003 R2 server, which has built-in support for NFS:

1. On the Windows 2003 R2 server, make sure Microsoft Services for NFS is installed. If it is not, you can add it by going into Add/Remove Programs, selecting Windows Components, and then selecting Other Network File and Print Services.
2. Go to the folder you want to share and right-click it and select Properties.
3. Select the Security tab and add the local group ANONYMOUS LOGON with read/write access.
4. Select the NFS Sharing tab, and select the Share this folder option. Put a check mark by Allow Anonymous Access.
5. Click the Permissions button, and for ALL MACHINES change the type of access to read-write and put a check mark by Allow Root Access.

After you have set up an ISO repository, you need to add it to your host servers. If you created a directory on an existing VMFS volume, there is nothing more you need to do because your hosts will already be able to see it. If you are using an NFS share, make sure you have a VMkernel network set up first (required for NFS). Then all you need to do is add it to your hosts as a storage device by doing the following:

1. Select a host in the VI Client and select the Configuration tab.
2. Under Hardware, select Storage.
3. Click the Add Storage link.
4. For Storage Type, select Network File System.
5. Enter your NFS server name or IP address. If you enter a hostname, it must be able to be resolved by the ESX/ESXi host.
6. Enter the folder (NFS share) name on the NFS server.
7. Enter the datastore name (that is, ISO) that you want it to be called on the ESX/ESXi host. After you have all the information populated, click the Next button.
8. Click the Finish button, and once the Create NAS Datastore task completes it is ready to be used. Repeat this for your other hosts to add it to them, too.

After you have an ISO repository set up, you are ready to start copying your ISO files to it so that you can use them with your VM's virtual CD/DVD drives. After you have some ISO files available to use, you can map your VM's CD/DVD drive to them by completing the following steps:

1. Edit the settings for your VM and select the CD/DVD drive, and on the right side change the device type to Datastore ISO File.
2. Click the Browse button and select your ISO repository from the list of datastores.

3. Select an ISO file and click OK.
4. On the VM Properties screen, as shown in Figure 6.10, make sure you put a check mark next to Connected (only if the VM is powered on) so that the CD/DVD drive is connected to the ISO file, and put one by Connect at Power On so that the CD/DVD drive will be connected when the VM is powered on so that you can boot from it if needed.




Figure 6.10 Selecting an ISO file to use with your VM’s virtual CD/DVD drive

You can also map the physical CD/DVD drive of your local workstation or an ISO file that is accessible on your local workstation or an ESX/ESXi host server’s CD/DVD drive to a VM’s CD/DVD drive by doing the following:

1. Edit the settings for your VM and select the CD/DVD drive, and on the right side change the device type to Client Device.
2. Open a remote console session to the VM in the VI Client.
3. At the top is a button(s) to connect your local workstation’s CD/DVD drive to the VM. When you click the button, it will give you an option to connect a specific drive letter that belongs to a CD/DVD drive on your PC to the VM or connect an ISO that is available on your PC to a VM.

Finally, you can map an ESX/ESXi host server's physical CD/DVD drive to a VM's CD/DVD drive by doing the following:

1. Edit the settings for your VM and select the CD/DVD drive, and on the right side change the device type to Host Device.
2. Select a host device to be used from the drop-down list. If there is only one CD/DVD drive in the host system, this will typically be `/dev/hda`, because most hosts use their IDE controllers only for their CD/DVD drives.

With any of these modes, you can change the virtual device node of the CD/DVD drive if needed by editing the VM's settings and selecting the CD/DVD drive. The number for the IDE device can be set to either 0:0, 0:1, 1:0, or 0:1. If you connect a CD/DVD drive on a VM, make sure to disconnect it when you have finished (even with datastore ISO files), because this can cause problems when moving the VM to another host. To disconnect your local CD/DVD drive, just click the Disconnect button that will appear above the remote console window when you connect it to the VM. To disconnect a host device or ISO datastore, just uncheck the Connect/Connected at Power On fields or select Client Device instead.

Did You Know?

A VM's floppy drive works much the same way as its CD/DVD drive, in that you can map it to a client or host device or a special floppy image file in FLP format.

After you have connected your VM's CD/DVD drive to a client/host device or an ISO image, you might want to boot from it so that you can install an operating system from it. To do this, you need to press the Esc key while the black VMware POST screen is displayed very briefly when a VM is powered on. Because a VM boots so fast, it is often difficult to press the Esc key during the time the POST screen is displayed. By default, this screen is displayed for only a second or two; but with ESX 3.5, you can now configure the delay that it is displayed for. Before ESX 3.5, this was not configurable, and you often had to power the VM on and off several times until you pressed the Esc key at just the right time. To change this delay, follow these steps:

1. Edit the settings for your VM and select the Options tab.
2. On the left side under Advanced, select Boot Options.
3. On the right side, you can set the Power-On Boot Delay in milliseconds (for example, 5,000 milliseconds = 5 seconds). Be careful not to set this too high because it will delay your VM from booting in situations where you might want it to boot quickly, such as an HA event in case of a host failure.

4. Optionally, if you want to force a one-time boot to the BIOS screen, you can check the option for the VM to Force BIOS Set Up the Next Time It Is Powered On. This option works only once unless you go back and check it again. Click OK when you are done.

When you are ready to boot your VM from its CD/DVD drive, just follow these steps:

1. Open a remote console session to the VM using the VI Client or web UI.
2. Power on or reboot/restart your VM. Click inside the remote console window with your mouse pointer to give it focus.
3. At the black POST screen that appears press Esc. Be careful; if you press it multiple times, you will dismiss the boot menu and initiate the default boot option instead.
4. A boot menu will appear, as shown in Figure 6.11. Select the CD-ROM drive option, and your VM will boot from its CD-ROM drive. This only affects the current boot process; subsequent boots attempt to boot from the hard disk first.




Figure 6.11 Selecting the CD-ROM device from the VM Boot menu

Now that we have covered how to create and build a VM using ISO files, let's cover how to use prebuilt templates as a quicker method for building new VMs.

Using Templates to Build Virtual Machines

If you plan to create many VMs, templates can be a real time saver. A template is essentially a preconfigured VM that has been turned into a master template to create other VMs from. Templates usually have a guest operating system already installed and configured, so a new VM can be quickly created and customized to make it different from the template and give it its own identity. Templates are a feature of VCenter Server and cannot be used without it; however, there are other ways to quickly create VMs if you do not have VCenter Server (as covered later).

Templates are created from existing VMs that are either cloned or converted into a template. If you clone an existing VM into a template, the original VM is left intact, and a copy of it is made for the template. However, any updates made to the VM will not be reflected in the template. Existing templates can also be cloned to create new templates. If you instead convert an existing VM into a template, the original VM becomes the template and is no longer considered a VM. VMs that are converted into templates can be converted back to VMs at any time. This is usually done to keep the template up-to-date with the latest patches and drivers. Although templates may be similar to VMs, they cannot be managed like VMs, and therefore cannot be powered on or migrated to other hosts. An alternative to using templates is to not convert a VM to a template and just clone it instead when you need to create a new VM.

Creating Templates

Templates will appear in both the Virtual Machines and Templates and Hosts and Clusters inventory view and can be distinguished from VMs by their different icon, as shown in Figure 6.12. To create a new template by converting an existing VM into a template, simply follow these steps:

1. It's best if the VM that you want to convert is powered off. Earlier versions of vCenter Server required this.
2. Select a VM in the VI Client and right-click it and select the Convert to Template option.
3. Once the task completes, the selected VM will now be a template and appear as such.




Figure 6.12 Templates are displayed in the VI Client.

To create a new template by cloning an existing VM, follow these steps:

1. The VM can be either powered on or off.
2. Select a VM in the VI Client and select the Clone to Template option.
3. A wizard will appear, and on the first screen you can provide a name for the template and a data folder location. Click Next when done. This wizard is similar to the one for creating a new VM that was covered earlier.
4. Select a host/cluster to put the template on, and click Next. If you select a cluster, you will get another screen to select a specific host server to put the template on. If you are using HA or DRS, each template must be assigned to a specific host. Once you select a host, click Next.
5. Select a datastore to store the template on and click Next.
6. Select a disk format for your template. Your options are Normal or Compact. If you select Normal, the template's disk file will be the same size as the VM that you are cloning (thick disk type). If you choose Compact, the disk is compacted to save space on your VMFS volume (thin disk type). So with Normal, a VM that is cloned with a 20GB virtual disk will have a 20GB disk as a template even if there is only 2GB of space in use on the disk. With Compact, a VM that is cloned with a 20GB virtual disk will have a disk roughly equal in size to the amount of disk space that is actually in use on the virtual disk. This can save a lot of space on your VMFS volumes if you plan on storing many templates.

7. Click Finish. When the task completes, you will have a new template that is based on the VM that you selected.

As mentioned earlier, you can also clone templates to create new templates. To do this, select a template and choose the Clone option. If you need to update a template to install the latest patches on it, all you need to do is select a template and choose the Convert to Virtual Machine option. This converts the template back to a VM that can be powered on and used like any other VM. After you have made the necessary updates, just power it off and select the Convert to Template option to convert it back into a template.

Using Templates

After you have your templates created, you can use them to start creating VMs. A template located on a specific ESX host server can be used to create a VM on any host server managed by vCenter Server. To use a template to create a VM, follow these steps:

1. Select a template in VCenter Server and right-click it and select the Deploy Virtual Machine from this Template option.
2. When the wizard launches, select a name for the VM and an inventory location and click Next.
3. Select a host/cluster to put the VM on and click Next. If you select a cluster, you will get another screen to select a specific host server to put the template on.
4. If a resource pool exists, you are prompted to select one. Then click Next.
5. Select a datastore for the VM and click Next.
6. Select to customize the guest operating system if it supports this. Not all operating systems are supported for this. To be able to do this, you must configure your vCenter Server to support this. Guest customization allows for changing the settings on a VM to make it unique from the template and other VMs. This includes things such as the server name, network configuration, Windows SID, and domain information. For information about how to set up guest customization, see the VMware documentation (Basic System Administration, at http://vmware.com/pdf/vi3_35/esx_3/r35u2/vi3_35_25_u2_admin_guide.pdf).
7. At the Ready to Complete screen, review the configuration options for your VM, go back and make any changes if needed, optionally select to power on the VM after it is created or edit the VM's virtual hardware, and then click Finish to create your VM.

As you can see, templates can be a real time saver (so that you do not have to build VMs from scratch whenever you need to deploy a new one).

Did You Know?

Because templates are associated with specific host servers, if you ever remove a host from vCenter Server and add it back in the templates will not automatically register like VMs do. To add an existing template back into vCenter Server, just browse the datastore that the template is located on using the VI Client and find the VMTX template file and select the option to add it to the inventory.

Template Best Practices

You should follow some best practices to ensure that you are making optimal use of your templates:

- Ensure that you keep your templates up-to-date with the latest patches and application updates. This includes operating system patches, antivirus pattern files, and driver updates. Doing this will ensure that your VM is less vulnerable to known vulnerabilities that may exist and will have already been patched.
- Install all the applications that you will need on your templates so that it is completely ready to go once you create a VM from it. This includes things such as backup agents, endpoint and system management applications, antivirus software, and whatever other applications that you install on all of your servers to manage them.
- Make sure to install VMware Tools on your templates and update it any time the version changes because of updates and patches to ESX.
- Apply any hardening and additional configuration that is used in your environment. This includes things such as disabling unneeded services, setting the local security policy, and configuring any software firewalls and any other settings that you commonly use in your environment.
- Give your templates descriptive names so that everyone can figure out what they are, and put a detailed description in the Notes field of the template. It's also a good idea to keep track of any template changes in the Notes field and document when the changes took place.
- Use a fresh installation of your operating system for your template rather than a preexisting one to ensure a more consistent and cleaner operating system. When you have finished installing and configuring everything, run an OS-level disk defragmentation on it before converting it to a template.

Templates are one convenient method for creating new VMs, but there are also other ways to do this, such as cloning an existing VM (as covered in the next section).

Cloning Virtual Machines

Cloning provides you with a way to make a copy of an existing VM to create a new VM, which you can then customize to make it unique. The clone option is a feature of vCenter Server and is not available when connecting directly to an ESX host. However, there are several other ways to create clones without VCenter Server (as covered shortly).

Creating Clones

If you are running vCenter Server 2.5 or later, you can clone a VM either when the VM is powered on or off. Earlier versions of vCenter Server required that the VM be powered off to clone it. As part of the cloning process, you will have the option to customize the guest operating system to make it unique, just like with templates. The process of creating a VM by cloning an existing VM is similar to that of creating a VM from a template. To create a clone of an existing VM, follow these steps:

1. Select a VM in vCenter Server and right-click it and select the Clone option.
2. When the wizard launches, enter a name for the VM and an inventory location and click Next.
3. Select a host/cluster to put the VM on and click Next. If you select a cluster, you will get another screen to select a specific host server to put the template on.
4. If a resource pool exists, you will be prompted to select one. Then click Next.
5. Select a datastore for the VM and click Next.
6. Select to customize the guest operating system if it supports this. Not all operating systems are supported for this. To be able to do this, you must configure your vCenter Server to support this. Guest customization allows for changing the settings on a VM to make it unique from the template and other VMs. This includes things such as the server name, network configuration, Windows SID, and domain information. For information about how to set up guest customization see the VMware documentation (Basic System Administration).
7. At the Ready to Complete screen, review the configuration options for your VM, go back and make any changes if needed, optionally select to power on the VM after it is created or edit the VM's virtual hardware, and then click Finish to create your VM.

As you can see, cloning is similar to using templates, but there may be some specific situations where you want to use clones rather than templates. One situation is for troubleshooting purposes. If you have a VM that is having application or operating system problems, you might create a clone of it so that you can try different things to resolve the problem without affecting the original VM. After you have determined the necessary changes, just apply them to the original VM, and you can delete the clone. Just be careful when creating the clone and

put it on an isolated vSwitch so that you don't experience network conflicts between the original VM and the clone. Another situation is creating separate environments (for example, development, test, UAT) from an existing environment that is already built and configured. By creating clones, you can ensure that each environment is consistent with each other to ensure the integrity of your application through its life cycle across the environments.

Alternative Methods for Creating Clones

There are several ways to create clones of existing VMs on hosts not managed by vCenter Server. These methods require a few more steps than that of the vCenter Server clone feature, but can still save you time when creating new VMs. The three alternative methods are detailed in the following sections.

Method 1: Use VMware vCenter Converter Starter

Converter is a free utility that can convert physical servers into VMs (P2V), but can also convert VMs into VMs (V2V). Officially, the free version of Converter (Starter) does not support doing remote conversions to an ESX host; you must instead install it on the VM that you want to clone and run it from there. However, there is a way to do remote conversions by doing a custom installation and installing only the Converter Agent on the source server and then installing the Converter Manager on a workstation and running it. Also, Converter is a Windows application and can only be installed on VMs running Windows. It is possible to convert Linux machines, but this requires the Enterprise version of Converter, which includes some Linux utilities and a cold clone option. We cover using Converter in more detail later in this chapter, but these steps show how to use it to clone a Windows VM:

1. Install Converter directly on the VM that you want to clone. The process for this is fairly easy and straightforward and is outlined here:
 - a. Copy the Converter application to the VM and start the installation program.
 - b. Click Next at the welcome screen, and then accept the license agreement and click Next.
 - c. Accept the default installation directory or change it and click Next.
 - d. Select a Typical installation type and click Next, and then click the Install button to begin the installation.
 - e. When the installation completes, you can click the Finish button, which will launch the Converter Manager application.
2. Once the Converter Manager is running, click the Convert Machine button to begin the conversion process.
3. The Conversion Wizard will start. Click Next at the welcome screen.
4. Click Next at the first Source screen, and at the Source Type screen select Physical Computer (don't worry, Converter doesn't care if it is physical or virtual), and then click Next. For the Source Login screen, select This Local Machine and click Next.

At the Source Data screen, you have the option to select drives to include in the conversion process and to resize them if desired. You also can specify whether to include the paging file (in most cases, you do not). Click Next when you have finished.

5. Click Next at the first Destination screen, and at the Destination Type screen select VMware Infrastructure Virtual Machine. At the Destination Login, enter a hostname and username/password to connect to it, and then click Next. At the VM Name screen, enter a name for your new VM and click Next. At the Host Name screen, select a host and optionally a resource pool for your VM and click Next. At the Datastore screen, select a datastore for your VM and click Next. At the Networks screen, select a network for your VM. It is best to not connect it at power on so that you can change the IP information for the VM so that it does not conflict with the source VM. Click Next when you have finished.
6. At the Customization screen, you can select to install VMware Tools on the VM, do guest customizations, and remove any system restore checkpoints (which is recommended). If you select to customize the guest operating system, you must have the Microsoft Sysprep files installed on the VM. Click Next when you have finished.
7. At the Ready to Complete screen, review your selections, click back if needed to make changes, and click Finish to begin the process.
8. When the process is completed, you will have a clone of your original VM that you can power on and customize if needed.

Method 2: Use the CLI Command `vmkfstools`

This method makes use of the `vmkfstools` command-line utility that is available in the ESX Service Console and with ESXi's Remote Command-Line Interface (RCLI). With this method, you are making a copy of the raw virtual disk of a source VM and using it with a new VM that you create by following these steps:

1. In the VI Client, create a new VM without a virtual hard disk (using the Custom option). Later versions of ESX 3.5 allow you to create a VM without a virtual hard disk. If you are running an earlier version and are forced to create a virtual hard disk, just create one of any size, and then go back and edit the VM's settings and delete the hard disk.
2. Log on to the Service Console. You will be using the `vmkfstools` command with the `-i` parameter, which will make a copy of the virtual disk file using the following format: `vmkfstools -i source virtual disk filename destination virtual disk filename`. Note that a virtual disk consists of two different virtual disk files: a small descriptor VMDK file and a large data VMDK file. Both files are automatically copied by `vmkfstools`. To find the name of the source disk filename, you can edit the settings of the VM and select the virtual disk, as shown in Figure 6.13.




Figure 6.13 Looking up a VM's virtual disk filename

3. It's best to power off the VM that you want to clone. If this is not possible, take a snapshot of the source VM, which makes the VM's disk read-only; otherwise, you will get a resource busy error. Run the `vmkfstools` command with the appropriate source and destination virtual disk filenames. (Watch your case because it is case sensitive.) The destination should include the path to the VM that you created in Step 1 (for example, `vmkfstools -i /vmfs/volumes/storage1/win2003-1/win2003-1.vmdk /vmfs/volumes/storage1/win2003-new/win2003-new.vmdk`). This will create a copy of the virtual disk on the win2003-1 VM into the directory for the new win2003-new VM. If you have more than one virtual disk in your VM, you will need to do this for each one.
4. Now that you have created a copy of the virtual disk, you can add it to the new VM by editing its settings and adding a new virtual disk. Instead of creating a new disk, just select the option to use an existing disk and browse to the new disk file that was created in your VM's directory. Once you add the disk and save the settings, you can power on your VM and customize it if needed. Don't forget to delete the snapshot on your source VM if you created one.

Method 3: Use the Datastore Browser

This method is similar to method two, but makes use of the datastore browser built in to the VI Client that was enhanced in ESX 3.5 to copy the virtual disk file instead of doing it using the CLI `vmkfstools` utility. It uses the same method of creating a new VM, as outlined here:

1. In the VI Client, create a new VM without a virtual hard disk (using the Custom option). Later versions of ESX 3.5 allow you to create a VM without a virtual hard disk. If you are running an earlier version and are forced to create a virtual hard disk, just create one of any size, and then go back and edit the VM's settings and delete the hard disk.
2. It's best to power off the VM that you want to clone. If this is not possible, take a snapshot of the source VM, which makes the VM's disk read-only; otherwise, you will get a resource busy error.

3. In the VI Client, select your host server, and then on the Summary tab select the datastore that your source VM is located on, and right-click it and choose the Browse Datastore option.
4. When the Datastore Browser window appears, select your source VM directory in the left pane, as shown in Figure 6.14, which will display all the files for the VM in the right pane. Even though there are two VMDK files that make up every virtual disk, the datastore browser abstracts them into one when displaying them.




Figure 6.14 File listing when selecting a source VM directory

5. Right-click the virtual disk file for the VM (type will be Virtual Disk) and select the Copy option.
6. In the left pane, select the VM that you created in Step 1, and then right-click in the right pane and select Paste. Note that the new disk that is created has the same name as the source VM, which is okay because it is in a separate directory from the source. However, it is a good idea to rename it to avoid confusion. Unfortunately, you cannot rename virtual disk files with the datastore browser. To do this, you must use the `vmkfstools` CLI utility with the `-E` option, which renames a virtual disk. The syntax for this command is `vmkfstools -E source virtual disk filename destination virtual disk filename`. Don't be tempted to use the Linux `mv` command to rename the files. There is metadata inside the file that the `vmkfstools -E` command updates, but that the `mv` command does not.
7. When the copy operation completes, you can close the Datastore Browser window. You can now add the copied disk to the new VM by editing its settings and adding a new virtual disk. Instead of creating a new disk, just select the option to use an

existing disk and browse to the new disk file that was created in your VM's directory. When you add the disk and save the settings, you can power on your VM and customize it if needed. Don't forget to delete the snapshot on your source VM if you created one.

These methods all provide a way to use existing VMs as templates if you do not have a vCenter Server in your environment. By using these methods, you can duplicate the functionality of templates and cloning and save yourself some time and ensure that your new VMs have consistent operating systems (because they are all created from your master templates). Just be sure that you stay in compliance with the OS and application vendor licensing guidelines when cloning and making templates.

Virtual Machine Configuration Settings

There are many configuration options to choose from when setting up your VMs. It's important that you understand these settings so that you can set them correctly and not cause problems with your VMs. These settings are all available by editing the settings of a VM. The settings fall into three categories: hardware, options, and resources that appear in tabs of the same names in the VM Properties window and are detailed here.

Virtual Machine Hardware

This part of the VM's settings lets you add and remove virtual hardware to and from your VMs and configure some of the various hardware components installed in your VM, as shown in Figure 6.15.

Listed here are the various hardware components that can be used with your VM:

- **Memory.** Can be set in 4MB intervals from 4MB up to 64GB. The upper limit is not defined by the amount of memory a host physically has, but instead by how much is supported by ESX. It's possible to assign more memory than a host physically has because of the VSWP file that is created and used if a host runs out of physical memory. The amount of memory a VM is assigned can be changed only when a VM is powered off.
- **CPUs.** Can be set to 1, 2, or 4 depending on the number of physical CPUs and cores that your host server has present and whether it supports hyperthreading. The amount that can be selected will change based on your host server. For example, you will only be able to select 1 or 2 on a host that only has two physical single-core, non-hyperthreading processors. The number of CPUs can be changed only when a VM is powered off.
- **Floppy drive.** A VM can have up to two floppy drives, which can only be added or removed when it is powered off. Configurable settings include the device status (connected or not) and device type (use client/host device or an image), which can be changed while the VM is either powered off or on.




Figure 6.15 VM hardware component settings

- **CD/DVD drive.** A VM can have up to four CD/DVD drives, which can only be added or removed when it is powered off. Configurable settings include the device status (connected or not), device type (use client/host device or an image), and mode (Passthrough or Emulate IDE), which can be changed while the VM is either powered off or on. In addition, the virtual device node (either 0:0, 0:1, 1:0 or 1:1) can be changed while the VM is powered off.
- **Network adapter.** A VM can have up to four network adapters, which can only be added or removed when it is powered off. The number of network adapters that can be added to a VM will vary based on the number of SCSI controllers that your VM has present. A VM can have up to six PCI devices, of which one is taken up by the video adapter, which cannot be removed from a VM. That leaves five, of which typically one is always used by the SCSI adapter, which leaves room for up to four network adapters. If you increase the number of SCSI adapters in a VM, this will reduce the number of network adapters that you can have. Configurable settings include device status (connected or not) and network label, which can be configured while the VM is either powered on or off. In addition, the adapter type (for example, flex or e1000) and MAC address can be changed while the VM is powered off.

- **SCSI controller.** A VM can have up to four SCSI controllers, which can only be added or removed when it is powered off. Similar to network adapters, the number of SCSI controllers that can be added to a VM will vary based on the number of network adapters that your VM has present. You cannot manually add a SCSI controller to a VM because they are automatically added or removed when you add hard disks and assign them a virtual device node ID. When you add hard disks to a VM, they will default to use the existing SCSI controller unless you select a virtual device node on another SCSI controller. In addition, the SCSI controller type (BusLogic or LSI Logic) and the SCSI bus sharing option can be changed only when the VM is powered off. SCSI bus sharing allows for the sharing of a disk between VMs (although only one VM can be powered on and using it at any given time).
- **Hard disk.** A VM can have up to 60 hard disks, 15 per SCSI controller. The virtual device node ID (for example, 0:0) that you can set is a combination of the controller number followed by the hard disk device number. The device number will be incremented when new hard disks are added until all the devices are used up on it (up to 15). It will then add another SCSI controller (1–3) as needed for additional disks. Optionally, you can choose to manually set a virtual device node ID to force it to add another controller. This is done in the Advanced Options section when adding a hard disk. If you choose a virtual device node ID that is not on an existing controller (for example, 1:0), a new controller will be added. You will only see IDs available for the number of SCSI controllers that can be added based on the number of free PCI slots. Hard disks can be added while a VM is either powered on or off, but they can be removed only while the VM is powered off. Virtual device node IDs can be changed on existing drives only when a VM is powered off. Take care, though, because it may change your disks around in your operating system and cause the VM to not boot. Additional options that can be changed when the VM is powered off are the disk size (ESX 3.5.x only), which can only be increased, and the disk mode (independent or not, persistent or nonpersistent).
- **Serial port.** A VM can have up to four serial ports that can only be added or removed when it is powered off. The device status (connected or not), connection (host physical port, output file, or named pipe), and I/O mode can be changed while the VM is powered either on or off.
- **Parallel port.** A VM can have up to three parallel ports that can only be added or removed when it is powered off. The device status (connected or not) and connection (host physical port or output file) can be changed while the VM is powered either on or off.

Additional hardware like audio adapters and USB ports are not supported with VMs running on ESX/ESXi hosts because they are on some of the hosted virtualization software (such as VMware Workstation and Server).

Virtual Machine Options

This part of the VM's settings lets you set advanced configuration options that affect the operation of the VM, such as power and boot options, as shown in Figure 6.16.




Figure 6.16 VM option settings

Listed here are the various options and settings that can be used with your VM:

- **General Options.** This section contains the VM name, configuration filename, working directory location, and guest operating system. The VM name is the display name for the VM, and if this is changed the VM's files that are all named beginning with the VM's name are not updated (for example, VMX and VMDK files). The working directory location and configuration filename cannot be changed in the VI Client. The VM name and guest operating system can both be changed when the VM is powered either off or on.
- **VMware Tools.** This section contains various settings that affect the VMware Tools application that is installed on the guest operating system of a VM. The first part is power controls, where you can set the behavior of the various VM power controls, such as the power off button, which you can configure to either perform a hard power off or a graceful operating system shutdown. The second part is for running VMware

Tools scripts when a VM changes certain power states. These scripts are located on the guest operating system (typically in C:\Program Files\VMware\VMware Tools), where they can be modified as needed. You can also set script options inside the VMware Tools application that runs on the VM. The last part is advanced settings, which control whether a VM checks for the latest version of VMware Tools and automatically updates it if needed and whether the VM syncs its time from the host server. These options can also be set on the VMware Tools application on the VM. Each of these settings can be changed only when a VM is powered off.

- **Power Management.** This section controls the behavior of a VM when the guest OS is placed in standby mode; the default is to leave the VM powered on. This can be changed to suspend the VM instead. In addition, you can set the vNIC used for Wake on LAN traffic. Each of these settings can be changed only when a VM is powered off.
- **Advanced, General.** This section contains settings to enable or disable video hardware acceleration and logging to the vmware.log file, which can be changed while the VM is either powered on or off. The Debugging and Statistics option enables you to change the amount of information logged about the VM, which can be useful for troubleshooting purposes and can only be changed while the VM is powered off. The Configuration Parameters button can be selected only while the VM is powered off and opens a new window where you can set advanced configuration parameters that are written to the VM's VMX configuration file. Many of the VM options (for example, CPU IDs, VMware Tools options) that you set are displayed in this area in their actual setting form. It's best to not add, remove or modify these settings unless you are sure you know what you are doing.
- **Advanced, CPUID Mask.** This section enables you to control what specific CPU features are presented to the VM. You can choose to either hide or expose the NX flag to the VM and set advanced CPU masks to hide certain CPU capabilities. The NX flag is an advanced CPU feature used by both Intel and AMD processors that stands for No Execute and provides the ability to segregate areas of memory so that any code in them is not allowed to execute. Intel refers to this feature as XD (Execute Disabled). By hiding this flag from the VM and preventing it from using this feature, you can increase VMotion compatibility between hosts that may not support certain CPU features. The Advanced button lets you set CPU identification masks on your VM to further mask CPU features and increase compatibility. CPU masks can be fairly complicated to figure out and set properly. For more information about how to set them, see VMware knowledge base articles 1991, 1992, and 1993 (<http://kb.vmware.com/kb/1991> and <http://kb.vmware.com/kb/1992> and <http://kb.vmware.com/kb/1993>). Each of these settings can be changed only when a VM is powered off.
- **Advanced, Boot Options.** This section enables you to set the length of time that the BIOS screen will display when a VM is powered on or rebooted. The default is 0, which causes the BIOS screen to be only very briefly displayed. Setting this to a value

in milliseconds will increase the delay (for example, 5,000 ms = 5 seconds). Setting this to a higher value can be useful if you are trying to get to the boot option menu to boot from a CD-ROM and need more time to press the Esc key. In addition, you can force the VM to go to the BIOS setup screen at the next boot by checking the Force BIOS Setup field. These settings can be changed while the VM is powered either on or off.

- **Advanced, Paravirtualization.** This section is for enabling or disabling paravirtualization support for the VM if the guest operating system supports it. (Currently, only certain Linux operating systems support it.) Paravirtualization allows the operating system to communicate directly with the virtualization layer hypervisor (VMware Tools is a form of paravirtualization) for certain instructions. Enabling this feature will utilize one of the six PCI slots of a VM and will reduce the number of NICs and SCSI adapters that can be used with your VM. Also, enabling this feature will limit the VMotion of a VM, because it cannot be moved while running to a host that does not support paravirtualization. This setting can be changed only while the VM is powered off.
- **Advanced, Fibre Channel NPIV.** This section is for assigning World Wide Numbers (WWNs) to VMs running on hosts with Fibre Channel adapters that support NPIV (N-port ID virtualization). NPIV allows for sharing a single physical Fibre Channel port among multiple virtual ports using unique identifiers. NPIV is supported only for VMs that have RDMs; otherwise, VMs use the WWNs of the host's physical Fibre Channel adapters. This setting can be changed while the VM is powered either on or off.
- **Advanced, Virtualized MMU.** This section controls whether the memory management unit (MMU) is virtualized. Newer CPUs support this feature and can improve the performance of a VM by decreasing latencies of certain memory operations. By default, the host will automatically determine whether this feature can be used, but you can also choose to either force or forbid that the feature be used. This setting can be changed while the VM is powered either on or off.
- **Advanced, Swapfile Location.** This section controls where the VMs swapfile (VSWP) is stored. By default, it is stored based on the setting of the cluster or host that the VM is on, which is typically in the same directory as the VM is located in. If you do not want to use the default setting for a particular VM, you can change this setting to force the VM to store it either in its directory or on a separate host swapfile datastore. If you choose to store it on a separate datastore, this can greatly increase the length of time that it takes to VMotion a VM between hosts if both hosts do not have access to the datastore that it is located on (for example, a local datastore). Because swapfiles can take up a lot of valuable disk space, it is sometimes desirable to store them on alternative cheaper storage locations (for example, local disk or iSCSI; NFS is not recommended because it is not a block-level protocol and can perform poorly when used for swapfiles).

Take care when changing these options because they can affect the behavior and performance of your VMs in both positive and negative ways. Make sure you understand each setting and what the effect will be on your VMs before you change them.

Virtual Machine Resources

This part of the VM's settings lets you control the hardware resource settings (CPU, memory, and disk) for your VMs.

These settings can both limit and guarantee the amount of resources that each VM will have available to it and prioritize a VM's access to resources against other VMs.

Listed here are the various resource settings that can be used with your VM:

- **CPU.** This section enables you to set shares, reservations, and limits for a VM to control the amount of host CPU resources that the VM has access to, as shown in Figure 6.17. You can either use one of these settings or use them together in combination to specifically tailor the amount of CPU resources that a VM can use (for example, set a reservation of 500MHz and a limit of 1GHz). Shares are generally easier to configure and use than setting up individual reservations and limits on your VMs. It is often better to manage resource allocations on resource pools rather than individual VMs. For more detailed information about this, see the "Configuring Resources" section in Chapter 4, "Configuring Your Virtual Environment."




Figure 6.17 VM CPU resource settings

- ◆ Shares define a VM's priority (low, normal, high, or custom) to access CPU resources compared to other VMs.

- ◆ Reservations guarantee that a VM has access to a certain amount of CPU resources and can be set between zero and the total amount of CPU megahertz that a host has. This can be misleading, however, because a VM can never use more CPU megahertz than that of the total megahertz of the number of vCPUs assigned to it. For example, a host with two dual-core 2.6GHz processors will have a total of 9768MHz. A VM with a single vCPU can only ever use 2.6GHz, and not the full 9768MHz, even if you set a reservation higher than 2.6GHz.
- ◆ Limits control the maximum amount of CPU resources that a VM can use regardless of how much is available. By default, this is set to unlimited and can be set between zero and the total amount of CPU megahertz that a host has. Again, similar to reservations, this setting can be misleading because a VM can never use more CPU megahertz than the total megahertz of the number of vCPUs assigned to it.
- **Memory.** This section enables you to set shares, reservations, and limits for a VM to control the amount of host memory resources that the VM has access to, as shown in Figure 6.18. You can either use one of these settings or use them together in combination to specifically tailor the amount of memory resources that a VM can use (for example, set a reservation of 512MB and a limit of 1GB). Similar to using shares on CPU, using shares is generally easier to configure and use than setting up individual reservations and limits on your VMs.




Figure 6.18 VM memory resource settings

- ◆ Shares define a VM's priority (low, normal, high, or custom) to access a host's physical host memory compared to other VMs.
 - ◆ Reservations guarantee that a VM has access to a set amount of physical host memory and can be set between zero and the total amount of memory that has been assigned to a VM. If a memory reservation is set on a VM and the host server does not have enough free physical memory to meet that reservation, the VM will not be allowed to be powered on. When memory reservations are set on a VM, the size of a VM's disk swapfile (VSWP) is reduced by the amount of RAM reserved for the VM.
 - ◆ Limits control the maximum amount of physical host memory that a VM can use, regardless of how much is available. By default, this is set to Unlimited and can be set between zero and the total amount of physical host memory that a host has. The Unlimited setting is also misleading because a VM can never use more memory than it has been assigned, regardless of the limit that is set. When a VM reaches the limit of physical host memory that has been set on it, it begins using its disk swapfile (VSWP) for additional memory.
- **Disk.** This section enables you to allocate disk I/O bandwidth used by a VM for its virtual disks, as shown in Figure 6.19. Because disk I/O is a different type of resource than CPU and memory, you can only set a priority to a VM using shares to control the I/O bandwidth that it has available compared to other VMs. By default, all VMs have equal priority, but you can change the priority of individual VMs to low, normal, high, or a custom priority.




Figure 6.19 VM disk resource settings

- **Advanced CPU.** This section enables you to set some additional CPU options such as hyperthreaded core sharing and CPU affinity, as shown in Figure 6.20. This section will not show up if you are using DRS or if a host has only one processor and does not support hyperthreading. Be careful with these settings (because the host CPU scheduler does a good job of scheduling CPU resources among VMs, and these settings can disrupt that and slow down your VMs if not set properly). Only use these settings if you must fine-tune CPU control for critical VMs.




Figure 6.20 VM advanced CPU settings

Watch Out!

If you set affinity and then join a cluster with DRS, the option to change it goes away and is not shown anymore, but affinity is still set.

- ◆ Hyperthreaded core sharing controls whether a VM can share a physical processor core or not and is applicable only if the host server supports hyperthreading. The default option of Any allows a VM to freely share CPU cores with vCPUs within the VM or with other VMs. Changing this to None will ensure that the vCPUs of the

VM have exclusive use of a processor core when they schedule it and that the other hyperthread of the core is halted so that nothing else can use it while the VM is using it (which essentially wastes CPU resources). You can also set this to Internal, which means a VM can share a core with another vCPU assigned to a VM if a VM has only two vCPUs. Otherwise, the VM will never share a core with any other VM. If a VM has one or four vCPUs rather than two, this setting is treated the same as the None setting.

- ◆ Setting a CPU scheduled affinity defines which CPU core that a VM is forced to always use. By default, affinity is not defined, and VMs are free to use any CPU core that is available as controlled by the host CPU scheduler. If you enable this, you can select specific CPU cores that a VM can only use. You can select more than one core even for a single-vCPU VM, but it can access only a single core at any given time. The check boxes are based on CPU cores and not sockets, so a host with two quad-core CPUs will have eight check boxes. Likewise, a host with two CPUs that support hyperthreading will have four check boxes. If you check all boxes, it is the same as not using affinity at all, because the VM can access all CPU cores. If you enable this option, you must select at least one check box per the number of vCPUs that a VM has been assigned. You cannot set affinity if a VM resides on a host in a DRS cluster and affinity is specific to a single host and cleared if the VM is migrated to a new host. For this reason, it is not recommended to use affinity unless you have a situation that requires it.
- **Advanced Memory.** This section enables you to select whether to use NUMA memory affinity for your VM, as shown in Figure 6.21. NUMA is a way of linking several small memory nodes together as a single larger node and can offer some performance benefits in certain situations. For more information about the NUMA architecture and how it relates to VMware, read the “VMware Resource Management Guide” located in the Documentation section of their website. This section is not displayed if your host does not use a NUMA memory architecture. In most situations, you should not change this, and should instead let ESX manage this. Some exceptions to this are if you have a small number of VMs that use a large amount of memory or if your VMs run applications that have memory-intensive workloads. Just like CPU affinity, this section does not display if a VM resides on a host in a DRS cluster and the affinity is specific to a single host and cleared if the VM is migrated to a new host. Checking all the check boxes displayed is the same as running no affinity.

That covers all the various options that you can set on your VMs to customize how it operates and functions. Before you try changing these options, make sure you understand their effects and try them out on noncritical systems so that you can get more familiar with them.




Figure 6.21 VM advanced memory settings

Installing VMware Tools

VMware Tools is a suite of utilities that can be installed on VMs and contains drivers and applications that help optimize the guest operating system to run on a VMware host. VMware Tools is bundled with all ESX/ESXi hosts, and there are different versions of it that are specific to different operating system types. VMware Tools is not required for a VM to run, but it is highly recommended that you install it for the added functionality (for example, time synchronization) and the benefits that it provides. VMware Tools contains the following components:

- Enhanced and optimized device drivers, including video, network, SCSI, memory, mouse, and a sync driver designed to work with Consolidated Backup
- A Control Panel inside the guest operating system that enables you to change certain settings and connect/disconnect virtual drives
- A memory balloon driver that inflates and deflates to optimize the memory management of the host server

- A service that can synchronize the guest's clock with the host server's clock and controls the seamless grab and release of the mouse cursor on Windows VMs
- A set of scripts that run when a VM's power state changes and which can help automate operations on the guest operating system
- A user process that controls copying and pasting between a VM and a remote console session

Installing VMware Tools is pretty straightforward and simple on Windows operating systems, as outlined here:

1. Open a remote console session to the VM with the VI Client and power the VM on.
2. When the VM boots, log on to the operating system.
3. From the VM top menu of the remote, choose VM, and then select the Install/Upgrade VMware Tools option. Optionally, you can select this option by right-clicking the VM in the VI Client view.
4. Click OK when the initial message displays. The Installation Wizard will then launch. Click Next on the welcome screen to continue.
5. Choose from one of the three setup types (Typical, Complete, or Custom). In most cases, you will want to select Typical. Custom lets you choose from the various components to install, and Complete installs some additional components that are not included with Typical that are not used very often, such as Shared Folders (which are not supported in ESX). Click Next after you select a setup type.
6. Optionally, if you choose the Custom option, you can select from the various components to install, as shown in Figure 6.22. Click Next.




Figure 6.22 Choosing a VMware Tools custom installation

7. Click Install to begin the installation. When the installation completes, a window may appear on Windows systems stating that hardware acceleration is not enabled and asking whether you want to enable it. It is recommended to do this to improve mouse and video performance. If you select Yes, a Display Properties window will open (as will a Notepad window containing instructions on how to enable it). To enable it, click the Advanced button in the Settings tab of the Display Properties window. A new window will open. Click the Troubleshoot tab of that window. Next, move the hardware acceleration slider from None to Full, as shown in Figure 6.23, and click OK when finished, and then OK again to close the Display Properties window.




Figure 6.23 Setting Windows hardware acceleration

8. When you click Finish on the VMware Tools installation, you will be prompted to reboot the guest VM, which you can choose to do immediately or later. A few of the drivers take effect without a reboot, but it is recommended to reboot the VM so that all the new drivers and applications will be loaded.

Installing VMware Tools on other operating systems (Linux, Solaris, and NetWare) is a bit more complicated. Linux guests can use either a tar or rpm installer that can be downloaded to the VM and launched with an included Perl script. For detailed information about installing VMware Tools on non-Windows operating system, see the Basic System Administration documentation on VMware's website or VMware knowledge base article 5242329 (kb.vmware.com/kb/5242329).


Once VMware Tools is installed, you can launch a Control Panel application that lets you configure a few options (such as time synchronization, connecting virtual devices, and setting script options for power events). To open the application on Windows systems, just double-click the icon that appears in the icon bar. If the icon has been hidden, you can access it using the Windows Control Panel. To open the application on Linux and Solaris guests, open a terminal window and enter the command `/usr/bin/vmware-toolbox &`.

When the application launches, you can select various tabs to set different options, as follows:

- **Options tab.** Here you can choose to sync a VM's clock with the clock of the host server. If you set this option, make sure you have no other time sync mechanisms in use on your guest operating system, such as Windows guests syncing time from the domain that they are part of. You can also choose to hide the VMware Tools icon in the icon tray and have a pop-up displayed if an update to VMware Tools is available because of a patch or new release that was installed on the host server.
- **Devices tab.** Here you can choose whether various virtual devices are connected or disconnected from the VM so that it can no longer access them. This includes virtual NICs, floppy drives, and CD/DVD drives.
- **Scripts tab.** Here you can set various script options for power events. The default scripts (batch files) are located in the default VMware Tools installation directory. You can edit the default scripts or create your own to use. The default scripts are empty except the suspend and resume scripts, which launch the `vmip.exe` command to release and renew a VM's IP address if it is using DHCP.
- **Shared Folders.** This tab is not used with ESX, and is intended for use with the hosted VMware products such as Server and Workstation.
- **Shrink.** This tab is not used with ESX, and is intended for use with the hosted VMware products such as Server and Workstation.
- **About.** Gives the version and build number of VMware Tools.

You can use the VI Client to keep track of your VMs and whether VMware Tools is installed and up-to-date. Just select the Virtual Machine tab on any object (for example, host, cluster, datacenter) to see a list of all VMs in that object, and you will see a VMware Tools Status column, as shown in Figure 6.24.

If you don't see the VMware Tools column, right-click the column headings. A selection menu will appear, from which you can select which columns are displayed.



Name, State, Host or Guest OS contains: ▾					
MHz	Host Mem - MB	Guest Mem - %	VMware Tools Status		
837		11			ToolsOk
938		66			ToolsOld
527		8			ToolsOk
1000		0			ToolsOld
3091		6			ToolsOk
404		7			ToolsOk
1127		9			ToolsOk
376		10			ToolsOk
471		9			ToolsOk
471		8			ToolsOk
0		0			Not Installed
390		7			ToolsOk

Figure 6.24 VMware Tools Status column shown in the VI Client VM listing

Best Practices for Creating Virtual Machines

It's easy to create a VM with a few clicks of the mouse. To ensure that your VM is configured properly, however, consider the following best practices when creating and configuring them:

- Always install VMware Tools on all of your VMs and make sure they stay up-to-date. VMware Tools contains drivers and applications that can benefit the performance of your VM. As patches and updates are installed on your host servers, the VMware Tools version may change and will need to be updated.
- Only give a VM the resources it needs as required by the applications that it is running. Don't go off the recommendations given to you by the application vendor; these are general in nature and are different in every environment. Use an application to monitor CPU and memory usage to see how much is actually used, and make sure you assign enough memory to cover any peaks that may occur. When creating virtual disks, don't go overboard on the amount of disk that you assign to a VM, and only assign as much disk space as the VM will actually use. VM resources can easily be increased, so start conservatively and increase them as needed.
- With physical Windows servers, the \I386 directory from the installation media is typically copied to the server's hard disk so that it can be quickly accessed when changes are made without trying to find the CD and putting it in the server. This directory is usually 500MB+; so to save valuable disk space on your VMs, it is a good idea not to copy it to their virtual hard disks. If you ever need to access it when OS changes are made, you can quickly and easily mount the ISO as needed to access it and then unmount it when finished.

- It's a good idea to create partial memory reservations for your VMs to reduce the amount of disk space used on your VMFS volumes for VSWP files. So for a VM with 2GB of memory assigned to it, create a reservation for 1GB, and that way only a 1GB VSWP file is created rather than a 2GB VSWP file. Just be careful that the total of your reservations does not exceed the amount of physical memory that a host has in it. You might also consider putting your VSWP files on a local VMFS volume instead to conserve space on your shared VMFS volumes.
- The optimal number of VMs per LUN/VMFS volume is typically 14 to 16, to avoid disk I/O contention to the physical disks that make up the LUN or RAID group and to keep down the amount of SCSI reservations that lock a LUN for certain VM operations. If your VMs have relatively low disk I/O (web and app servers), you can increase the number to 18 to 24; and if your VMs have high disk I/O (database and email servers), you should reduce the number to 8 to 12.
- Remove any unneeded virtual hardware from your VM, including floppy drives and serial and parallel ports. These devices consume extra resources on your host servers even if they are not being used.
- Many VMs will run better with a single vCPU rather than with multiple vCPUs. If the application running on a VM will not benefit from multiple vCPUs, use only a single vCPU. Having too many vCPU VMs on a host can make it more difficult for the CPU scheduler to schedule time among the VMs and can reduce the performance of all the VMs on a host.
- When creating VMs on your VMFS volumes, allow some extra disk space for a VM's additional files. A VM with a 20GB virtual disk uses more than 20GB. Additional space is needed for its various files, including the swapfile, suspend files, snapshots, and logs. Factor at least 20% more disk space to be used by these files. So for a 20GB VM, you should allow for an additional 4GB. If you plan to use snapshots and keep them around for an extended period of time, allow for even more disk space. Likewise, if you plan to use multiple snapshots, allow for even more disk space. Individual snapshots can vary in size depending on how much information changes after they are taken and can grow up to the size of the original disk.

This covers creating new VMs on your host servers. In the next section, we cover migrating existing physical servers and converting them into VMs.

Migrating Physical Servers

In many instances, you will want to migrate your existing physical servers to your new hosts as VMs. This is done through a process called P2V (physical to virtual) migration, which copies all the data from a physical server to a new VM that is created. When the P2V process

completes, you shut down the physical server, power on the new VM, and do some post-conversion cleanup and configuration. The length of time that the P2V process takes to complete can vary depending on the speed of the network connection and the amount of data to be copied from the physical server to the VM. For servers that have a small amount of data (less than 10GB) on a gigabit network, the conversion can take less than 30 minutes. For servers with larger amounts of data (greater than 40GB), this can take several hours to complete. Also, the type of data to be copied and the copy method (block versus file) can greatly affect the amount of time it takes to complete.

Conversion Methods: Hot Versus Cold

There are many different P2V methods and tools that you can use for this (which we cover shortly), but the most commonly used tool is a free application that VMware provides called vCenter Converter. These tools typically support two conversion methods: hot cloning and cold cloning. Cold cloning is done while the operating system of the VM is not active by booting from a live CD that contains the P2V application. This method ensures the most reliable conversion process because there are no applications running on the server while the data is being copied to the new VM. Hot cloning is done while the operating system is running and can sometimes be risky because files are open and in use while the copy is being performed, which can cause data corruption on the VM in some instances. Many conversion tools will quiesce the disk before copying the data, which reduces the risk of corruption by forcing the operating system to write any data held in memory to disk before copying the data. The advantages and disadvantages of each method are listed here.

Advantages and disadvantages of cold cloning include the following:

- More reliable than hot cloning because operating systems and applications are not running.
- Less convenient than hot cloning because downtime is required while the cloning process is running.
- Requires physical access to the server (unless you have a remote management card that supports virtual CD-ROMs) to boot from a cold-clone CD.
- Can be quicker than hot cloning because the operating systems and applications are not running, which can cause contention with the cloning process.
- Supported for more operating systems because the cloning process does not involve the operating system.

Advantages and disadvantages of hot cloning include the following:

- More convenient than cold cloning because downtime is not required while the cloning process is running.
- Slight chance of data corruption on destination VM because files and applications are open while the cloning process is running.

- No physical access to the server is required because the cloning application can be installed and run on the VM remotely.
- Can take longer than cold cloning because of contention with the operating system and applications that are running on the source physical server while the cloning process is running.
- Not supported for all operating systems.

In general, choosing the cold-clone method is best because it is the more reliable because the operating system and applications are not running while the cloning process is running. However, there may be times that you want to hot clone to avoid all the downtime required for cold cloning. The chances of experiencing a data-corruption problem while hot cloning are small and will vary based on the type of application running on the server that you are cloning. Servers that are running applications like web and application servers that are fairly static and do not have much information that changes have less risk. Servers that are running applications like database and email servers have more risk because they frequently change a lot of data on the server's hard disk. Therefore, your decision to choose a method should be based on the risk factor and the type of data that is on the server that you want to clone.

The bottom line is that I highly recommend cold cloning if you are able to do it. If you have to use hot cloning, make sure to use some of the best practices found later in this chapter to reduce your risk of problems and data corruption. You should also be aware that with hot cloning any data that changes during the cloning process is not copied to the target VM, so it's a good idea to shut down all applications and services when hot cloning.

When to Use P2V and When to Build New

Cloning a server is not always a problem-free process, and certain types of applications are just as easy to migrate to a VM by building a new server and installing the applications fresh and then moving the application data to it instead of using a P2V process. Doing it this way minimizes any potential problems that can occur during the P2V process when data is copied from the source server to the destination VM. Some examples of these types of servers are listed here:

- **Active Directory domain controllers.** Active Directory (AD) uses its own databases to store information, in which it replicates among other domain controllers in a domain. AD is often finicky and complicated, especially in large domains, and trying to P2V a domain controller can sometime cause problems within the domain after the process has been completed and the new VM brought up. The types of problems that you may encounter after doing a P2V conversion of a domain controller include USN rollback problems, tombstoning, domain database corruption (ntds.dit), and replication problems. As an alternative to doing a P2V conversion on a domain controller, do a fresh install of Windows on a VM and dcpromo it, transfer any FSMO roles to it, and then

dcpromo the physical server so that it is no longer a domain controller and shut it down.

- **Database servers.** Databases are very sensitive to any disk operations, and copying a database while it is running can easily corrupt it. Compare this to changing your car's oil while it is running. It can be messy and cause big problems. As an alternative to doing a P2V conversion on a database server, do a fresh install of your operating system on a VM and install the database software and then detach the database from the physical server, copy it to the VM, and attach it there.
- **File servers.** File servers typically have a lot of data, which can take a very long time to copy, and because of the amount of data there is more potential for data corruption. As an alternative to doing a P2V conversion on a file server, do a fresh install of your operating system on a VM and then use a copy utility or application that can copy the file security and attributes and do scheduled and incremental updates of the data. When you have finished, you can just shut down the old file server and start using the new one. Microsoft's Robocopy application is an example of a utility that can do this.

The added advantage of building a new VM and migrating to it rather than doing a P2V conversion on it is that you get a fresh build without all the built-up junk that often occurs on Windows servers over time as applications are added, removed, and upgraded and settings and configurations are changed. So before you do a P2V conversion, stop and consider whether it might be just as easy to build a new VM and copy the data instead.

P2V Methods and Tools

You can use several methods and tools to perform P2V conversions. Which method or tool you choose will depend on several factors, including the number of conversions, how fast you need to do the conversions, how easy you want the conversions to be, and the amount of money you have to spend on tools that are not free. Conversion methods go from simple ones like using an imaging application like Ghost and boot CDs to full applications specifically written to perform P2V conversions. In addition to P2C conversions, some of these tools support doing V2V and V2P conversions. The following sections detail several methods and tools that you might use depending on your situation.

VMware vCenter Converter

This utility, which was developed by VMware, comes in two editions: Starter and Enterprise. The Starter edition is free for anyone to download and use, and the Enterprise edition comes with vCenter Server. The Starter edition has a few limitations compared to the Enterprise edition:

- The Starter edition supports only hot cloning. The Enterprise edition supports both hot and cold cloning.

- The Starter edition supports only a single conversion at one time. The Enterprise edition supports multiple simultaneous conversions.
- The Starter edition officially supports only local conversions, and so the Converter must be installed on the server you are converting. The Enterprise edition supports both local and remote conversions, and so conversions can be run using a remote workstation. However, there is a way to do remote conversions by doing a custom installation and installing only the Converter Agent on the source server and then installing the Converter Manager on a workstation and running it.

Converter is a great tool but has some limitations and is not as robust as some of the other applications, such as Platespin's PowerConvert or Vizioncore's vConverter. It is a good tool for converting Windows servers, but has limited support for Linux and other operating systems. It can perform both P2V and V2V conversions, but cannot do V2P conversions, which is sometimes necessary when you need to reproduce a problem for a vendor on a physical server. In addition to supporting Windows operating systems as source systems, it supports importing many common image formats and converting them to VMs, including Norton Ghost, Symantec LiveState and Backup Exec System Recovery, StorageCraft ShadowProtect, and Acronis True Image.

I suggest first trying Converter to determine whether it meets your needs. If it doesn't, look at some of the more robust applications instead. For more information about vCenter Converter, visit VMware's website at www.vmware.com/products/converter/ and check out the documentation at http://vmware.com/support/pubs/converter_pubs.html.

Platespin Migrate (Previously Known as PowerConvert)

This application is very robust, with many advanced features, and supports many operating systems and conversion methods (P2V, V2P, and V2V). It has full support for converting all 32-bit and 64-bit versions of Windows and Suse and Red Hat Linux operating systems. It includes a number of advanced features, including the following:

- Live transfer, which ensures that data changes made while migrating are also moved
- Server sync, which ensures that the source and destination stay in sync before cutting over
- High migration speed and scalability to ensure fast and many concurrent transfers
- Remote control to perform migrations of remote servers with no agents, live CDs, or physical contact with servers

If you are looking for a tool that supports non-Windows operating systems, has high scalability and speed, and offers many advanced features and functionality, check out this tool by visiting their website at <http://platespin.com/products/plateSpinMigrate/>.

Vizioncore vConverter

Vizioncore vConverter is another very robust and feature-rich product, but is limited to supporting only Windows systems. It has great conversion speeds and scalability and includes the following features:

- Extensive automation for pre- and post-conversion tasks to simplify and speed up conversions
- Support for remote cold cloning
- Synchronized cutover to test and synchronize source and target systems before cutting them over
- One of the fastest cloning engines due to streamlined file-transfer mechanisms and advanced intelligence handling of cloning

If you are looking for an upgrade from VMware's vCenter Converter and are converting only Windows systems, Vizioncore's vConverter is a great choice. For more information about their product, visit their website at <http://vizioncore.com/vConverter.html>.

Additional P2V Methods

There are some other ways to migrate physical servers to VMs, such as using one of the many imaging programs that are available, such as Norton Ghost or Acronis True Image. By creating a disk image of a server, you can manually create a VM, boot it from a floppy or CD image, and then restore the disk image file to it. This is essentially what most of the P2V applications do when they clone a physical server into a VM. Doing this manually, however, can be a bit more complicated and add more steps for pre- and post-conversion tasks. These additional methods were popular back when VMware did not offer a free conversion tool but now are mainly used for operating systems that are not supported by VMware's free vCenter Converter tool. For more information about these methods, check out the following web pages:

- Ultimate P2V web page, www.rtfm-ed.co.uk/?page_id=174
- MOA, <http://sanbarrow.com/moa.html>
- Linux P2V wiki, http://conshell.net/wiki/index.php/Linux_P2V

Using VMware vCenter Converter

As previously mentioned, vCenter Converter comes in two editions. The free Starter edition is a stand-alone application that can install on any Windows system to convert it to a VM. There is also an Enterprise edition, which has two components: a service that is installed on the vCenter Server that is included in the vCenter installation and a VI Client plug-in component

that is installed on any workstations that are running the VI Client and want to use Converter. The Enterprise edition is integrated into both the vCenter Server and the VI Client and provides easy P2V and V2V conversions with just a few simple steps.

Both editions support a variety of Windows source operating systems, as follows:

- Windows NT 4 SP6+
- Windows 2000
- Windows 2003 32-bit and 64-bit
- Windows XP Professional 32-bit and 64-bit
- Windows Vista 32-bit and 64-bit

In addition, Converter has experimental support for additional source operating systems, including Linux, Windows 98, Windows 95, Windows Me, Windows NT 3.x, and MS-DOS. This means that Converter can clone these systems but the destination VM may not work until additional configuration is completed on it, which typically involves installing drivers and modifying boot partitions.

Both editions also support a variety of different source import formats, as follows:

- VMware Workstation 4.x, 5.x, and 6.x
- VMware Fusion 1.x
- VMware ESX and ESXi 3.x
- VMware ESX 2.5.x (if managed by vCenter Server 2.x)
- VMware Server 1.x
- vCenter Server 2.x
- Microsoft Virtual PC version 7 or later
- Microsoft Virtual Server (any version)
- VMware Consolidated Backup images
- Acronis True Image 9
- Symantec Backup Exec System/LiveState Recovery 6.5 and 7.0, LiveState Recovery 3.0 and 6.0, and Norton Ghost 9.0, 10.0 and 12.0
- StorageCraft ShadowProtect

Both editions also support a variety of destination formats for your source systems, including the following:

- VMware Workstation 4.x, 5.x, and 6.x
- VMware Fusion 1.x
- VMware ACE 1.x and 2.x
- VMware Player 1.x and 2.x

- VMware ESX and ESXi 3.x
- VMware ESX 2.5.x (if managed by vCenter Server 2.x)
- VMware Server 1.x
- vCenter Server 2.x
- OVF and OVA virtual appliance formats

Converter can be used in two ways: either by installing it directly on a Windows operating system and then running it inside the OS to convert it to a VM or by selecting an image format directly like an existing workstation VM or Norton Ghost image file. In most cases, when converting physical Windows servers you will be installing it on the operating system either locally or remotely and then converting the server.

Installing and Using Converter Starter Edition

Converter Starter edition supports only hot cloning a source server into a new VM and can be installed directly on the operating system of the server that you want to convert or by doing a custom installation and installing only the Converter Agent on the source server and then installing the Converter Manager on a workstation and running it. Before using Converter, read through the release notes and the latest documentation for the version that you are using. These are both available on VMware's website at http://vmware.com/support/pubs/converter_pubs.html. Also, be sure to read the best practices at the end of this chapter to help ensure that you have a successful conversion. To install Converter Starter edition, follow these steps:

1. Download the latest version of Converter at www.vmware.com/download/converter/. Free registration is required to download it.
2. Copy the installation file to the server that you want to convert and run it.
3. Once the installation program loads, click Next at the welcome screen.
4. Accept the license agreement and click Next.
5. Select the default installation folder or change it and click Next.
6. Select Typical for the installation type, which installs both the Converter Agent and Manager application, and click Next.
7. Review the options. Click Back if you need to make changes, and click Install to begin the installation.
8. Once the installation completes, click the Finish button and the Manager application will launch.

Now that Converter is installed, you can run the Manager application to begin the conversion. The Manager lets you create and monitor conversion tasks and must be run on the

server that you are converting with the Starter edition. Once you launch the Manager application, follow these steps to hot clone your server into a VM:

1. It's best to log on to the server that you are running Converter on using the local administrator account. If not, make sure the account that you are using has local administrator privileges.
2. The first time you launch the Manager on a server, you will get a license screen on which you can enter a license key to run it as the Enterprise edition or you can continue and run it as the Starter edition. Adding an Enterprise license key simply adds some features to the stand-alone application, which is different from the Enterprise edition that is built in to vCenter Server. When you click the Continue in Starter Mode button, you can start using the Manager.
3. To start the Conversion Wizard, click the Convert Machine button. The Conversion Wizard will launch, and you can click Next at the welcome screen.
4. The first step is to select a source computer. Click Next at the first Source screen.
5. At the Source Type screen, select Physical Computer to choose the computer that you are running the application on, as shown in Figure 6.25. This option can be used whether the server is physical or virtual, because it simply means that we are choosing a server to convert by using the operating system to clone it instead of choosing a VM or image file. The other options include VM, Virtual Appliance, or Other (image file). Click Next to continue.




Figure 6.25 Running Converter Starter edition: Selecting a source type

- At the Source Login screen, select This Local Machine, as shown in Figure 6.26. You can optionally select a remote server as your source if you entered an Enterprise license code (not supported in Starter). If you do select a remote server, you must enter the credentials to log on to that server. When you select This Local Machine, the credentials of the currently logged-on user are used.




Figure 6.26 Running Converter Starter edition: Setting the source login information

- At the Source Data screen, all hard disk partitions configured on the server display, as shown in Figure 6.27. You can select which disks you want to be cloned to the new VM and resize them smaller or larger if you want. Most physical servers have more disk space than they need, so it is recommended to reduce the size of the disks to what is actually needed. You can also choose to ignore the paging and hibernation files, which is recommended because they will be created fresh on the new VM and copying them will add more time to the cloning process. Once you choose your hard disk options, click Next.
- Now you must choose a destination for your cloned server. Click Next at the first Destination screen.
- At the Destination Type screen, select VMware Infrastructure Virtual Machine, as shown in Figure 6.28, which will let you later choose an ESX/ESXi host or vCenter Server. Optionally, you can select Other Virtual Machine if you want to convert to one of the stand-alone VMware products like Workstation or Server. Click Next to continue.




Figure 6.27 Running Converter Starter edition: Selecting the hard disks for the source data




Figure 6.28 Running Converter Starter edition: Selecting a destination type

- At the Destination Login screen, enter the hostname or IP address of an ESX/ESXi host or vCenter Server and the login credentials, as shown in Figure 6.29. If your destination is a host that is managed by vCenter, it is highly recommended to enter a vCenter Server name rather than the host server name. Once you enter a vCenter Server name, you will have the option to select a host in a later step.




Figure 6.29 Running Converter Starter edition: Setting the destination server and login information

- If you selected a vCenter Server, the next screen will allow you to enter a VM name and choose a folder for the VM, as shown in Figure 6.30. If you selected an ESX/ESXi host instead, you will only be able to enter a VM name. Once you enter a name and folder (if applicable), click Next.
- If you selected a vCenter Server, the next screen allows you to choose a host and cluster that your VM will reside on, as shown in Figure 6.31. If you selected an ESX/ESXi host instead, you will only be able to choose a host server. Once you make a selection, click Next.




Figure 6.30 Running Converter Starter edition: Setting a VM name and folder




Figure 6.31 Running Converter Starter edition: Choosing a destination host server and cluster

13. At the Datastore screen, you can select one of the datastores that is configured on the host that you selected for your VM to reside on, as shown in Figure 6.32. Optionally, you can click the Advanced button to select a separate datastore for the VM's configuration file. Once you select a datastore, click Next.




Figure 6.32 Running Converter Starter edition: Choosing a destination datastore

14. At the Networks screen, you can configure the vNICs for your VM and choose which networks they are connected to, as shown in Figure 6.33. Once you configure your vNIC, click Next.
15. Now you're ready for guest customization. You can customize your operating system to be different from the VM that you are cloning, as shown in Figure 6.34. To use this, you must have Sysprep configured on your vCenter Server. You can also choose to install VMware Tools on the new VM and remove any system restore checkpoints. After you have you made your selections, click Next. If you choose to do guest customization, additional screens will appear to customize the new VM.
16. Finally, at the Ready to Complete screen, you can review all the selections that you have made and go back if needed to change them. When you are ready to begin the conversion process, click the Finish button.




Figure 6.33 Running Converter Starter edition: Setting vNIC information




Figure 6.34 Running Converter Starter edition: Setting guest customization options

17. The status of the conversion will be updated as it progresses, as shown in Figure 6.35. Once the conversion completes, you will have a new VM located on the host that you chose.




Figure 6.35 Running Converter Starter edition: Conversion in progress

If your conversion does not successfully complete, you can check the log files that it creates located in the C:\Windows\temp\vmware-temp for clues as to why the conversion process failed. After you have finished cloning your physical server to a VM, make sure you power off the physical server that you cloned before powering on and configuring your new VM so that they do not conflict with each other. You can also uninstall the Converter application; it is not needed on the new VM that was created. Make sure you follow the post-conversion best practices listed later in this chapter.

Installing and Using Converter Enterprise Edition

Before vCenter Server version 2.5 (2.0.x), the Enterprise edition was the same stand-alone application (now called Starter), with more features unlocked by entering an Enterprise license code. With the release of vCenter Server 2.5, the stand-alone edition still exists, but

Converter was integrated into the vCenter Server and the VI Client and labeled the Enterprise edition. The Enterprise edition of Converter that integrates with vCenter Server consists of two components: a Windows service that installs on the vCenter Server and a plug-in component to the VI Client. In addition, the Enterprise edition includes an ISO image so that you can use the cold-clone feature to boot from a CD-ROM and clone a server while the operating system is not running.

Installing Converter Enterprise Edition

Because Converter Enterprise edition is integrated into vCenter Server 2.5.x and VI Client, there is nothing that you need to download to install it. There is an optional zip file that contains the ISO boot CD that you can download from VMware's website so that you can use the cold-clone feature. The vCenter Server component installs automatically as part of the vCenter Server installation, which you can verify by looking at the Windows services and making sure the VMware Converter Enterprise Service exists and is running. If for some reason the service is not listed in your Windows services, you can manually install it by completing the following steps:

1. Locate the vCenter Server installation CD or the unzipped installation directory.
2. In the converter directory is the installation application for Converter. Run the VMware-Converter executable program to install it.
3. Choose a setup language and click OK, and then click Next at the welcome screen.
4. Accept the license agreement and click Next.
5. Either accept the default installation directory or change it, and then click Next.
6. Choose the Typical installation type, which installs the Converter Enterprise application and service and the command-line interface (CLI), and then click Next.
7. Enter the vCenter Server information, including the hostname or IP address, port (which defaults to 443), and a username/password to connect to vCenter Server (defaults to local administrator), and then click Next.
8. At the Port Configuration screen, accept the default ports of 9085 (SOAP) and 9086 (web) and click Next.
9. At the Identification screen, choose a hostname (recommended) or IP address to identify the Converter Enterprise server and click Next.
10. Click Install to begin the installation. When it completes, the VMware Converter Enterprise Service will be listed in the Windows services and will be started.

Once the VMware Converter Enterprise Service is running on the vCenter Server, you need to install the plug-in component on any workstation that uses the VI Client that you want to do conversions from. vCenter Server plug-ins were introduced in vCenter Server 2.5 as additional components that could be bolted on to vCenter Server to provide more features and functionality to it. VMware first released two plug-ins: the Converter Enterprise plug-in

and the Update Manager plug-in. Although plug-ins can be used with the vCenter Server, they actually install into the VI Client. By default, no plug-ins are installed with the VI Client, and they must be manually installed afterward. To install the Converter Enterprise plug-in component, follow these steps:

1. Launch the VI Client and connect to your vCenter Server.
2. From the top menu, choose Plug-Ins, and then choose the Manage Plug-Ins option.
3. The plug-in manager will load, as shown in Figure 6.36, and has two tabs, one for available plug-ins and the other for installed plug-ins. There should be two available plug-ins listed, Converter Enterprise and Update Manager. Click the Download and Install button to begin the installation of the Converter Enterprise plug-in.




Figure 6.36 VI Client plug-in manager: Available tab

4. When you click the Download and Install button, a Setup Wizard will launch. Choose a language and click OK, and then click Next at the welcome screen.
5. Click the Install button to begin the installation, and then click the Finish button when it completes.
6. When the installation completes, the Converter Enterprise plug-in will show as installed.
7. Click the Installed tab and the Converter Enterprise plug-in will be listed. Put a check mark next to the Enabled field to enable the plug-in in the VI Client, and then click OK to close the plug-in manager.

Now that the plug-in is installed, you will notice a new Import Machine option in the menus displayed when you right-click a host, resource pool, or cluster. This option launches the Converter Enterprise Wizard to hot clone a physical server, image file, or VM into a new VM.

Hot Cloning with Converter Enterprise

The hot-cloning process for Converter Enterprise edition is similar to that of the Starter edition, with the main difference being that the Conversion Wizard is built in to the VI Client and launched with the Import Machine option instead of installing and running a separate application. To hot clone a server with Converter Enterprise, follow these steps:

1. To begin, select a host, resource pool, or cluster in vCenter Server and right-click it and choose the Import Machine option.
2. The Conversion Import Wizard will launch, and the first step is to choose a source server. Because much of this wizard is identical to that of the Starter edition, which we previously covered, we will skip over most of it and point out the differences. For the source server, the difference is that you cannot choose the This Local Machine option for the physical server choice because Converter is not installed directly on the server that we want to convert. You must instead choose the name or IP address of the server that you want to convert.
3. The destination and customization steps are the same as previously covered with the Starter edition. The Enterprise edition includes a fourth step, which allows you to schedule the conversion for a later date/time and schedule it as a repeating task, as shown in Figure 6.37 (good for copying a VM to a DR site on a regular basis).




Figure 6.37 Scheduling a conversion with Converter Enterprise

4. When you complete the wizard, the conversion will appear as a task within vCenter Server. When it completes, your new VM will appear and be ready for use.

Cold Cloning with Converter Enterprise

Cold cloning is done by using the special boot CD that comes with Converter Enterprise to boot from a live CD that contains a self-contained operating system so that the operating system on the hard disk of the server that you are cloning is not loaded. This method ensures that no files are in use on the server that you are cloning and is the most reliable and safest cloning method. Cold cloning can be used with any operating system that is supported by Converter, and is the only method available to clone non-Windows servers (because hot cloning is not available for them). The live CD contains a slim Windows PE operating system, which loads the same Converter application that is used in the Starter edition. To cold clone a server, follow these steps:

1. Download the zip file from VMware's website and extract the ISO file (usually called coldclone.iso) from it.
2. If the server that you are converting has a management board that supports virtual CD-ROM devices, you can mount the ISO file to it without burning it to a physical CD-ROM. Otherwise, burn the ISO file to a physical CD-ROM using a burner and an application like Nero.
3. Power off the server that you want to clone, insert the CD-ROM, and power the server on. Most servers will try to boot from the hard disk first, so you might need to press a particular key or go into the BIOS and change the boot device order so that you can boot from the CD-ROM instead of the operating system on the hard disk of the server.
4. When the live CD operating system loads, the Converter application will start, and the first screen you will see is a licensing agreement. Accept it and click OK to continue.
5. Next the Converter application will try to configure a network connection using the NICs that are found and DHCP. If DHCP is not available, you might receive an error that no network connection is available and asking whether you want to update the network parameters. If a DHCP server is found, it will automatically configure the IP address information from it, and you will receive a prompt asking whether you want to update the network parameters to make any changes.
6. If you selected Yes to configure or update the network parameters, a new screen will appear letting you choose a NIC from the ones detected and manually setting IP address and DNS/WINS server information, as shown in Figure 6.38. In addition, on the Advanced tab, you can set some other network settings, such as the adapter speed and duplex. It is important that you set the adapter speed and duplex to

match the settings on the physical switch that your NIC is connected to. If you have a speed/duplex mismatch, it can drastically slow down your conversion speed and time. The default for Converter is Auto/Auto. If your switch is set the same, you do not need to change anything. It's best to check with your network team to see what the physical switch port is set to so that you can set it to match appropriately in Converter. If there are no NICs available to select from, the driver for the NIC in your server is not included in the live CD. Only a limited number of the most common NIC drivers are included in the live CD. If you are using a NIC that has a driver that is not included, you will be unable to connect to the network and will have to manually add the driver to the ISO file using the Petool application that is included with the ISO file download (more on this later). Also on the Advanced tab you can configure a network drive in case you want to write the new VM files there instead of directly to an ESX/ESXi host or vCenter Server.




Figure 6.38 Cold cloning with Converter Enterprise: Configuring NIC settings

7. The Converter application will load, which is basically the same stand-alone Converter application that is used in the Starter edition, as shown in Figure 6.39. Click the Import Machine button to begin the cloning process.




Figure 6.39 Cold cloning with Converter Enterprise: Manager application

8. The first step is to choose the source. Because we are cold cloning, the only option is to clone the server that you are running the application on. Click Next at the first Source screen, and then on the next screen you can choose to import all the disks and maintain their sizes or you can choose the option to select individual disks to include and resize them if necessary, as shown in Figure 6.40. If no disks are found, the live CD does not have the driver for your storage controller, and you will need to manually add it to the live CD using the Petool application (more on this later). Click Next after you have selected your source disks.
9. The second step is to choose a destination. Click Next on the first Destination screen, and on the next screen choose a destination type. If you choose Other VMware Virtual Machine, you can select the network drive that you can configure in the network parameters step as the destination. Otherwise, choose your ESX/ESXi host or vCenter Server and VM name, datastore, and network the same way as described earlier with the Starter edition.




Figure 6.40 Cold cloning with Converter Enterprise: Choosing source data

10. The third step is customization, which is set the same way as described earlier with the Starter edition. When you complete this, you will be at the Ready to Complete screen, where you can review all your options and change anything by going back. You then click the Finish button to begin the cloning process.
11. The conversion will progress with task updates being displayed similar to the Starter edition. When the conversion completes, you can power off the physical server and power on the new VM to configure it and clean it up.

As previously mentioned, the Converter Enterprise live CD has many of the most common NIC and storage controller drivers for all the major hardware brands, but you might come across a server that has a NIC or storage controller whose driver is not on the live CD. This is most common with older hardware, and if this happens you will need to manually add the driver to the live CD image. To do this, you must first locate the Windows driver for the model of NIC/storage controller that is in your server and then use the command-line utility called Petool to inject it into the ISO image. Once the driver is injected into the ISO image, you can burn it again to a physical disc and try the cloning process again. For instructions on using the Petool utility, see the “Converter Enterprise Administration Guide” on VMware’s website.

Best Practices for Using Converter

You should use many best practices when using Converter, to both prepare your server for conversion and after the conversion to ensure that you have successful and quick

conversions. The speed of your conversions can vary greatly depending on your server's configuration and the options you choose when you convert it. Therefore, it is highly recommended that you follow these best practices when you do your conversions.

Preconversion Best Practices

Here are some best practices that you should follow to prepare a server for conversion:

- When providing logon credentials for Converter to use for a remote conversion on a server, use a local administrator account if possible. Avoid using Windows domain accounts that may have restrictive group policies associated with them. Likewise, if you are installing and running Converter locally on a server, use a local administrator account.
- If hot cloning, the preferred method is to install and run Converter locally on the server to be converted instead of running it from a remote workstation.
- If hot cloning, stop or disable any running applications, if possible, to avoid problems with open files or data being changed while the conversion is running. This is especially important on transaction-based applications such as database and email servers.
- Disable antivirus and backup software before you begin the conversion process.
- Run `chkdsk` on the disks to be converted and defragment them.
- Clean up all temporary files and any unneeded data to reduce the amount of time that the conversion process takes.
- Ensure that you have adequate free disk space on the server that you are converting because Converter needs 200MB of space when it quiesces the disk before the conversion begins.
- Check your physical network switch port to see what speed/duplex it is set to. This should match whatever is configured in your operating system NIC settings if hot cloning. If cold cloning, make sure you set Converter to use the same speed/duplex setting as your physical switch port.
- Ensure that ports 902 and 443 (including software and hardware firewalls) are open between the server that you are converting and the ESX/ESXi host and vCenter Server. Also, ensure that the host server and vCenter Server names are resolvable in DNS (both common name and fully qualified) by the server that is being converted. Use the IP address of the host server or vCenter Server if you have problems with name resolution.
- Make sure that the TCP/IP NetBIOS Helper service is running on Windows NT/2000 servers. On Windows XP/2003 servers, make sure the Volume Shadow Copy service is not disabled (set to Manual is okay).
- On Windows NT/2000 servers, reboot the server after installing Converter and before running the conversion, because the special driver that is installed is not started until after the server is rebooted.

- When hot cloning, keep users off the server being cloned, if possible; this includes accessing any network shares, remote desktop sessions, and accessing any applications running on the server.
- If you have any UPS power-monitoring drivers/applications installed, remove them; they can cause the new VM to not boot properly or to shut down immediately.

Running Converter Best Practices

Here are some best practices that you should follow while running the Converter application to convert a server:

- Do not select any small hardware utility partitions that may exist on the disk to be included in the conversion process, because this can cause your VM to not boot. If a boot partition exists and you do not select it, you might need to modify the boot.ini on destination Windows systems so that they boot properly.
- If your destination host server is managed by a vCenter Server, make sure you choose the vCenter Server as the destination instead of the ESX/ESXi host directly.
- If the server you are converting has a hard disk greater than 256GB, ensure that your destination VMFS datastore is not configured with the default 1MB block size, which limits a single virtual disk to 256GB in size. If you have no VMFS datastores with a larger block size, resize the disk so that it is less than 256GB.
- If your destination is an ESX/ESXi server and not a vCenter Server, make sure you use the root account to log on to the host server.
- Deselect the options to install VMware Tools and customize the guest operating system. Install VMware Tools manually after you power on the new VM.
- Do not modify the default number of vNICs; these can be changed after the conversion completes.
- Resizing your disks is a great way to conserve disk space on your destination datastore. However, you might occasionally experience problems when resizing because a different cloning process is used when the option to resize disks is chosen. When you choose to reduce the size of a disk, a file-level copy operation is done as part of the cloning process instead of the normal block-level copy operation. This method can cause problems, especially when copying a very large amount of small files or very large files. If you experience problems with the conversion failing, try not resizing the disk instead. This only happens when decreasing the size of a disk; increasing the size of a disk uses the block-level copy operation.
- Make sure you choose a unique name for your new VM that is not already in use.

Post-Conversion Best Practices

Here are some best practices that you should follow after you convert a server into a VM:

- Clean up the VM's virtual hardware before powering it on. This cleanup includes adjusting the RAM size, number of CPUs and NICs, and removing any unneeded hardware (for example, serial/parallel/USB ports, floppy drive).
- Boot first into Windows Safe mode and remove all specific drivers and applications that were specific to the original server hardware, including network teaming/management applications, array configuration software, and video and sound drivers. Do not restart if prompted when uninstalling; wait until everything is uninstalled first, and then reboot to Normal mode.
- If you went from multiple CPUs to a single vCPU or vice versa, make sure you change the Windows HAL appropriately. For example, on a Windows 2003 server, you should first take a snapshot of your VM (to recover in case of a blue screen of death), and then go to Device Manager and the Computer section to see what HAL is loaded (for example, ACPI Uniprocessor PC). Select whatever is listed and choose Properties, and then select the Driver tab and click the Update Driver button. Select No for connecting to Windows Update, and then click Next, and then select Install from a list or specific location and click Next. Choose Don't Search I will Choose the Driver to Install and click Next. Make sure Show Compatible Hardware is checked, and choose the appropriate CPU driver for your VM (usually ACPI Multiprocessor PC or ACPI Uniprocessor PC), click Next, and then click Finish.
- Install VMware Tools, and then reboot when finished.
- Remove all nonpresent hardware from the operating system. There will be a lot of hardware devices that are no longer present on the new VM after a conversion from different physical hardware. Although they are no longer present, they still consume Windows resources and should be completely removed. Also, if your old NICs used a static IP address, you will get a warning error when trying to set your new vNIC to the same IP address because they are still configured as nonpresent devices. Removing them will resolve this. To do this, follow these steps:
 1. Open a Windows command (CMD) prompt.
 2. Type the following: `SET DEVMGR_SHOW_NONPRESENT_DEVICES=1` (not case sensitive, make sure you spell it exactly right) and press Enter.
 3. In the same CMD prompt window, type `DEVMGMT.MSC` to open the Device Manager application. You cannot open it through the GUI because the `SET` command only affects applications run from the same CMD window.
 4. Once Device Manager opens, select View from the top menu, and then the Show Hidden Devices option, which will display all hidden and nonpresent hardware devices.

5. Expand each section and remove all hardware that is grayed out (lighter in appearance than the other devices), which means it is no longer present, by right-clicking it and choosing Uninstall. There will be a lot of it under the Non-Plug and Play Drivers and System Devices sections. If you see no grayed-out hardware, you might have spelled the command in Step 2 incorrectly. If so, go back and try again.
6. When you have finished, reboot.
 - Check your Windows services and make sure there are no vendor hardware-specific services running. If the uninstall applications did not remove them, disable them.
 - Check your vNIC properties to make sure there are no vendor-specific clients or services loaded (for example, NIC teaming drivers). If there are any, uninstall them.

Following these best practices will help ensure a successful conversion. If you do find that you have trouble powering on your new VM (for example, blue screen), this is typically caused by some of the vendor hardware-specific applications/drivers or problems with the boot.ini file. To resolve this, try adding the new VM's virtual hard disk to another working (helper) VM as an additional hard disk. You can then edit the boot.ini file from the helper VM and make sure the partition information is correct and delete directories that contain vendor-specific hardware drivers and applications. When you have finished, just remove (not delete) the hard disk from the helper VM and try powering on the new VM again.

Additional Converter Resources

If you would like additional information about using Converter, check out the following additional resources:

- VMware Converter Best Practices (VMworld 2007 presentation, free registration required)
www.vmworld.com/vmworld/mylearn?classID=11610
- Checking networking connections between Converter endpoints
<http://kb.vmware.com/kb/1006607>
- Best practices for using and troubleshooting VMware Converter
<http://kb.vmware.com/kb/1004588>
- Using VMware Converter for P2V migration
http://searchvmware.techtarget.com/tip/0,289483,sid179_gci1304471,00.html
- Conducting hot P2V migrations with VMware Converter
http://searchvmware.techtarget.com/tip/0,289483,sid179_gci1304484,00.html

- Troubleshooting failed VMware Converter P2V migrations
http://searchvmware.techtarget.com/tip/0,289483,sid179_gci1304490,00.html
- Remove nonpresent devices from converted system
<http://support.microsoft.com/kb/315539>

Using Guided Consolidation

vCenter Server 2.5 included a new feature called Guided Consolidation, which uses a built-in wizard to discover physical servers, analyze them to determine whether they are good candidates for conversion into VMs, and convert them into VMs using the Converter application. This feature is based on the Capacity Planner tool that is used by VMware business partners to analyze customer environments and only works with servers running Microsoft Windows. The process used by Guided Consolidation consists of the following steps:

1. **Find servers.** Searches for physical systems on your network to analyze as candidates for virtualization.
2. **Analyze servers.** Servers that you select are analyzed, and performance data on each system is collected. The longer the duration of the analysis, the higher the confidence in vCenter Server's recommendations.
3. **Convert servers.** Performance data is compared to the resources available on the host systems. The selected physical systems are converted to VMs and imported into vCenter Server on the recommended hosts.

To be able to use this feature, the following prerequisites must be met:

- At least one ESX/ESXi host must be registered with vCenter Server.
- At least one datacenter inventory object must exist on the vCenter Server.
- Consolidation services require local administrator privileges on the vCenter Server. Specifically, the Collector service must be run with local administrator privileges. In addition, the account used must also be granted the Logon as Service privilege. If Active Directory is deployed on your network, the credentials used to run Consolidation services must also have sufficient privileges to query the Active Directory database.
- Consolidation services also require administrator access to the systems selected for analysis. Specifically, the Collector service uses these credentials to connect to and retrieve configuration and performance data from the physical systems under analysis.

- Remote Registry service must be running on the servers that you want to analyze because this is used by vCenter Server to collect data from them.

Guided Consolidation uses a Windows service that is installed on the vCenter Server called VMware Capacity Planner Service. Similar to the Converter service, this service installs as part of the default vCenter Server installation. If it is not installed on your server, you can install it manually by locating the vCenter Server installation CD or the unzipped installation directory, and then in the vpx directory you can launch the capacityplanner.exe application. Follow the prompts after you launch it, and it should appear in your Windows services list afterward.

This feature was introduced in vCenter Server 2.5, but the UI was enhanced in the Upgrade 2 version to make it easier to search and analyze servers. As a result, if you are using an earlier version, you may see slightly different screens than shown here. To use this feature, follow these steps:

1. Connect to your vCenter Server with the VI Client, and then on the top icon bar click the Consolidation button.
2. Before you begin, you should configure your consolidation settings, as shown in Figure 6.41, by choosing Administration from the top menu and then selecting the Consolidation Settings option. Here you can set both the service and default credentials used by the Consolidation service. The user account that the VMware Capacity Planner service will run as serves as the service credentials. After you enter the credentials, it will automatically update the service so that it starts as the user account that you specify rather than Local System. Whichever user you specify must be a member of the Local Administrators group on the vCenter Server, be granted the Log On as a Service user right in the local security policy, and must be a domain account so that it can read information from Active Directory (does not need to be a domain administrator). The user account that is used when connecting to systems that you want to analyze serves as the default credentials. This account must have administrator access to the systems to be analyzed. It's easiest to use a domain administrator account for the default credentials. Optionally, you can override the default credentials when analyzing a system by clicking the Set Authentication button when selecting it.
3. When you click the Consolidation button, an Analysis tab will display, as shown in Figure 6.42, where you can click the Start Analysis button to begin.




Figure 6.41 Guided Consolidation: Configuring credential settings




Figure 6.42 Guided Consolidation: Starting an analysis

4. An Analysis Selection screen will display, as shown in Figure 6.43, where you can choose which servers that you want to include. You can select by server name, IP addresses by selecting a filename containing a list of server names/IP addresses, or by browsing through your available Active Directory domains. After you have made your selection, click the Add to Analysis button.




Figure 6.43 Guided Consolidation: Adding servers to be analyzed

5. A window will display giving you the option to either use the default credentials that were previously configured or use different credentials to connect to the servers. After you have made a selection and entered the credential information, click OK to begin the analysis.
6. The analysis will begin, and performance data will be gathered from the selected servers, as shown in Figure 6.44. The data gathered from the servers is basic and does not use some of the more advanced metrics that the full version of Capacity Planner uses. Up to 100 systems can be analyzed simultaneously. The Analysis page displays CPU and memory info on the computers, along with a status, confidence, and average CPU/memory usage over time. Confidence is the degree to which vCenter Server can gather performance data about the system and how good a candidate the system is based on the available data.




Figure 6.44 Guided Consolidation: Analysis in progress

- Once you decide to consolidate a physical server into a VM, you can click the Plan Consolidation button (or right-click the server and select the Convert to Virtual Machine with Recommendations option), which displays a window, as shown in Figure 6.45, with the selected computers, a host destination (which you can change), and a destination rating that indicates the degree to which the host can comfortably accommodate the estimated resource needs of the resultant VM (the more stars, the better). VMware recommends that you consolidate only a single computer at a time.
- When you click the Consolidate button, it automatically begins the process of converting the physical system into a VM by hot cloning the system while it is running, as shown in Figure 6.46. It first creates a new VM on the selected host with the same number of CPUs and memory of the physical system and then imports the server's hard disks using the embedded Converter plug-in.




Figure 6.45 Guided Consolidation: Server consolidation recommendations displayed




Figure 6.46 Guided Consolidation: Converting the selected server into a VM

Although the Guided Consolidation feature is not as robust as the full Capacity Planner tool that VMware business partners can use to analyze your environment to prepare it for virtualization, it is nonetheless a handy tool to use when you have a small number of servers that you want to virtualize.

Summary

Populating your virtual environment is probably the most fun and exciting step because you get to create and work with VMs and see your virtual environment start to function. Just be careful that you keep tabs on your host's resources because it is very easy to quickly use them up by creating too many new VMs. Running servers in virtual environments has many advantages, and you will find that your normal server administration routines will change as a result: things such as being able to create snapshots before making changes to your server's operating system and applications, easily adding more memory or disk space to a server, using ISO files for virtual CD-ROM drives, and being able to remotely power a nonresponsive VM on and off. Now that your environment is built, it is time to unleash its power and to take advantage of the many features and functionality that it provides. Once you get some experience working with VMs, you will find it hard to work with physical servers ever again.

Chapter 7

Monitoring Your Virtual Environment

Now that your virtual environment is configured and populated, you will want to monitor it both for performance and for problems. The root cause of performance issues is often not very obvious and can have a ripple effect on many VMs and host servers. Therefore, you should both configure monitoring in your environment and understand the metrics and data that are being reported so that you can proactively take steps to eliminate bottlenecks and problems. You may also want to look into some of the many third-party monitoring and reporting tools that are available. These tools are much more robust and powerful than the tools that are built in to ESX and vCenter Server and will greatly enhance your monitoring abilities.

Monitoring Your ESX and ESXi Hosts

Monitoring your hosts is very important because things such as a hardware component failure can be disruptive to your environment. The HA feature is a good recovery feature for hardware failures, but still causes some disruption while a VM is being powered up on another host server. Knowing when a fan, drive, or memory module is failing so that you can take action can help minimize any disruption to your environment. When managing ESX and ESXi hosts without a vCenter Server, no alarm capabilities are configurable through the VI Client, and you must instead use third-party applications and SNMP to provide alerting capabilities for your host servers. In this section, we cover how to monitor your host hardware and configure SNMP (which can be used to send alerts and statistics to a monitoring system).

ESX Hardware Monitoring

Beginning with ESX version 3.5 Update 2, VMware added the ability to monitor server hardware health status via the VI Client using hardware providers that they included in ESX. The level of health detail that will display will vary depending on the server brand/model and how new it is. Newer hardware has built-in Common Information Model (CIM) providers that show more health information (for example, fan status, temperature, voltages) than older hardware, which will only show some basic hardware health information, as shown in Figure 7.1.




Figure 7.1 Hardware health status displayed in the VI Client

To get more detailed information, you can install special vendor-specific hardware management agents on your ESX hosts inside the Service Console Linux-based operating system to monitor the hardware health of a server. These agents can be used with enterprise management applications such as HP's Systems Insight Manager and Dell's OpenManage for hardware alerting and reporting. In addition, the information that they provide can be accessed by using a web browser and accessing a specific port on the server.

The installation procedure for the hardware management agents will vary depending on the manufacturer of the server. The following example is based on installing them on an HP server. Not all vendors provide these agents for use with ESX, although all the major manufacturers do (HP, IBM, Dell). Visit your vendor's support website and search for ESX drivers to determine whether they are available. Here are the steps for installing the HP agents on to an ESX host:

1. Download the latest drivers from the vendor's support website. For HP servers, select your server model, and then select VMware ESX Server 3.5 (or whatever your version is) as your operating system. Then choose to download the HP Management Agents for VMware ESX Server 3.x.

2. After you have downloaded the agent, which is usually in a TGZ tar format archive file, you need to copy it to the Service Console. To do this, use an application such as WinSCP or FastSCP, which lets you copy files using a Windows Explorer-like interface to and from ESX hosts. Copy the file to the /tmp directory on the ESX host.
3. Next you need to log on to the Service Console and unarchive the file using the `tar` command with the `-zxvf` parameters, which extracts all the files (for example, `tar -zxvf hpmgmt-8.1.0-vmware3x.tgz`).
4. The files will usually be extracted to a subdirectory. Switch to that directory and do a directory listing, and you should see a file with a `.sh` extension, which is a Linux shell script file, which you can run using the `--install` option (for example, `./installvm811.sh --install`).
5. Follow the prompts, and the installation should begin. You will be prompted as to whether you want to enable the port (2381) that is used to access the web-based management page on the ESX built-in firewall. Choose Yes and the port will be automatically opened for you. You will also be prompted as to whether you want to enable the port (2301) used by the Systems Insight Manager (SIM) application so that it can be managed by it if you are running it in your environment. Choose Yes and the port will be automatically opened for you. If you are not running SIM, choose No.
6. Another prompt will display asking whether you want to enable the SNMPD port in the firewall. SNMP is used to send alert messages using traps. If you do not have a management system that can receive these traps in your environment, select No for this. If you do, select Yes.
7. The next prompt will be for enabling the port (280) used to add the SIM certificate to the web-based Systems Management Homepage (SMH). Choose Yes if you are using this feature.
8. The next prompt asks whether you want to install the Insight Manager Agent. Choose Yes.
9. The next prompt asks whether you want to load the HP modules even though they will taint the kernel. This simply triggers a kernel taint flag, which will not affect ESX. Choose Yes.
10. The next prompt asks whether you want to start the SNMP agents. If you are using SIM or another SNMP application, choose Yes. Otherwise, choose No.
11. The next prompt asks whether you want to start the storage agents, which provide detailed information about Fibre Channel-attached storage devices. If you choose Yes for this, you will also have to manually install the libraries for the FC driver (QLogic or Emulex) that you are using. See the readme file that is included in the .tar archive for instructions on how to do this.

12. The next prompt asks whether you want to start the NIC agents, which provide detailed information about HP NICs in your host. Choose Yes if you want this additional detail. Otherwise, choose No.
13. The next prompt asks whether you want to use an existing `snmpd.conf` (configuration file). If you select Yes, no additional information will be written to this file by the installation script. If you choose No, the file will be updated by the script. Choose Yes if you have already set this up or want to set it up yourself. Otherwise, choose No.
14. The next prompt asks whether you want to disable SMH support. If you select Yes, an `hpsmh` local user account and group will not be created. Selecting No will allow the script to create these to be used by SIM and SMH. Choose Yes if you are not using SIM or SMH.
15. The HP agents and services will be loaded, and the installation will complete, as shown in Figure 7.2.

```
Starting SNMP stack:
Starting snmpd: [ OK ]

Proliant System Health Monitor loading
*****
* System Management Homepage installed successfully with *
* default configuration values. To change the default *
* configuration values, type the following command at *
* the root prompt: *
* *
* perl /usr/local/hp/hpSMHSetup.pl *
* *
*****
Please read the License Agreement for this software at
/opt/compaq/hpsmh/hpsmh.license

By not removing this package, you are accepting the terms
of the "License for HP Value Added Software".

HP Insight Manager agents have been installed successfully!
[root@esx1 011]#
```

Figure 7.2 HP management installation completed

16. You can now add the ESX host to SIM or SMH for centralized management or access the host directly from a workstation with a web browser by going to the URL `https://hostname or IP:2381`, and a login page will display, as shown in Figure 7.3.
17. You can log in using the host's root account, and once you successfully authenticate you will be at the management home page, where you can select from various options to see your hardware's health status, configuration information, statistics, and alerts, as shown in Figure 7.4.



HP System Management Homepage for esx1.xyz.com

Account Login

This is a monitored private system. Do not attempt to login unless you are an authorized user.


Note: This version of the HP System Management Homepage uses host operating system authentication.

User:

Password:

HP System Management Homepage v2.1.15.210 ©2004-2008 Hewlett-Packard Development Company, L.P.

Figure 7.3 System Management login page is displayed



HP System Management Homepage for esx1.xyz.com

Support | Forums | Help

System Model: NULL
Current User: root
logout

Home Settings Tasks Logs

Home

Integrated Agents
HP Foundation Agent
HP NIC Agent
HP Server Agent
HP Storage Agent

Other Agents
none

Management Processor
none

Other Software
HP Essentials Software

KEY: ✓ OK
▲ Degraded
▼ Failed
? Unknown

Friday, January 02, 2009 10:39:36 AM
refresh: manual

System Status Summary
no failed/degraded items

Storage
File System Space Used

Utilization
✓ Processor Utilization

System
Software Version Info
System Info
System Board
Memory Subsystem
Exp Buses
... 2 items not shown

Recovery
Environment
Power Supply
Remote Communications

Figure 7.4 System Management home page is displayed

Your installation may vary depending on the version of agents that you install. Be sure to check with your vendor because the agents are usually updated periodically and you want to make sure you have the most up-to-date versions on your ESX hosts. Also, be sure that the version of the agent you are using is supported with the version of ESX that you are using as sometimes the latest versions are only supported on certain versions of ESX.

ESXi Hardware Monitoring

Vendor-specific hardware management agents cannot be installed on ESXi hosts because there is no Service Console to install them on. ESXi uses built-in hardware providers that leverage the CIM framework to provide basic hardware information to the VI Client using the Health Status feature that was introduced in ESXi version 3.5. HP integrated a special hardware management agent in ESXi because they are not following industry standards. That agent is available as a separate ESXi version that can be installed to obtain more detailed hardware information and to integrate with HP SIM. For other hardware vendors, you can configure ESXi to use SNMP to send alerts for hardware events. Here are some sites from Dell and HP on configuring ESXi to work with HP SIM and Dell OpenManage:

- Dell

http://support.dell.com/support/edocs/software/eslvmwre/EN/VES_3i/Systems%20Management%20Document/PDF/MSMPA01.pdf

- HP

http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01570108/c01570108.pdf?jumpid=reg_R1002_USEN

SNMP Monitoring

SNMP stands for Simple Network Management Protocol and is used to gather statistics and alerts from a wide variety of application and hardware devices, information that can be sent to a management application such as HP SIM and OpenView, Dell OpenManage, and Microsoft System Center. These applications can receive this information and monitor the device and trigger actions such as sending an alert email for certain events. Because ESX/ESXi and vCenter Server do not have very robust monitoring and alerting mechanisms, using SNMP is a good way to use a more powerful enterprise monitoring system to monitor your hosts and your whole virtual environment using a centralized monitoring system.

Both ESX and ESXi support using SNMP, but ESX has a more robust implementation because of its Service Console, which has a Linux-based net-snmp agent. This agent can interact with a management application using the GET and SET commands and can send alerts using the TRAP command. The simpler agent that is used with ESXi (and can also be used with ESX) only supports sending alerts via traps and not gathering statistics and other advanced functions that require use of the GET and SET commands. SNMP has evolved into

several different versions over the years (v1, v2, and v3). The first version, v1, had very weak security and was improved in v2. The latest version is v3, which provides the best security because it added support for using authentication and access control. The version of SNMP used with ESX supports all versions of SNMP, but all traps are sent using v1 format. ESX can be configured to use the better v3 security for GET and SET operations (traps are still sent in v1 format, however) where a management application reads or writes information to the host server. ESXi is limited to using the v1 version because of size constraints, and so it can only send traps and not read or write information with SNMP.

To configure SNMP to use with ESX, follow these steps:

By the Way

If you have already installed one of the hardware management agents (for example, Dell or HP), SNMP may already be configured for you. You can further configure it using some of the steps included here.

1. To allow SNMP traffic through the built-in ESX firewall, first connect to the server with the VI Client and select the Configuration tab and then the Security Profile option under Software. Click the Properties link. When the window is displayed, put a check mark next to SNMP Server to enable it, as shown in Figure 7.5, and then click OK.




Figure 7.5 Enabling SNMP traffic in the ESX firewall

2. Edit the SNMP configuration file so that you can configure the various options. Log in to the ESX Service Console and type `nano /etc/snmp/snmpd.conf -w`, which opens the file in the nano editor in word-wrap mode (you can also use the vi editor), as shown in Figure 7.6.

```

root@esx1:/etc/snmp
GNU nano 1.2.1      File: snmpd.conf

# Sample snmpd.conf containing VMware MIB module entries.
# This is a simple snmpd.conf that may help you test SNMP.
# It is not recommended for production use. Consult the
# snmpd.conf(5) man pages to set up a secure installation.

syscontact root@localhost (edit snmpd.conf)
syslocation room1 (edit snmpd.conf)
rocommunity public
trapcommunity public
trapsink localhost

# VMware MIB modules. To enable/disable VMware MIB items
# add/remove the following entries.
dlmod SNMPESX      /usr/lib/vmware/snmp/libSNMPESX.so

[ Read 15 lines ]
^G Get Help  ^O WriteOut  ^R Read File ^Y Prev Page ^K Cut Text  ^C Cur Pos
^X Exit      ^J Justify   ^W Where Is  ^V Next Page ^U UnCut Txt ^T To Spell

```

Figure 7.6 Editing the `snmpd.conf` file with the nano editor

3. You can set many options in this file. The most basic ones are listed here:
 - a. **rocommunity**. This is the name of the string that is used for read-only access to SNMP data from the host. This is essentially a password that is needed by any management application to read data from the host. The default of public is very well known and should be changed to a unique, stronger string instead.
 - b. **trapcommunity**. This is used to send alerts (traps) to SNMP receivers, and the string used must match whatever is configured on the receiver for the device sending it. This also defaults to public and should be changed to something more complex.
 - c. **trapsink**. This is the destination device for all SNMP trap messages and is typically set to that of a management or monitoring application server. This can be a hostname or IP address, and you can have multiple devices set by having multiple trapsink lines on separate lines. The default is localhost, which means traps are sent to the ESX host server and nowhere else. Change this to whatever device you want to receive the trap messages.
 - d. **rwcommunity**. This is the name of the string that is used for read access to SNMP data from the host and for devices to write data to the host. By default, this is not

set, so no write access is allowed. If you set this, make sure you set it to something complex, because this string not only allows read access to your host but also write access.

4. When you have finished editing the file, save it by pressing Ctrl-O (WriteOut), and then exit by pressing Ctrl-X (Exit).
5. Start the SNMP daemon, which does not run by default. To do this, type the following command: `service snmpd start` (or optionally use the `/etc/rc.d/init.d/snmpd start` command).
6. This only starts the SNMP daemon until the host is rebooted. To set the SNMP daemon to start automatically when the host starts, enter the `chkconfig snmpd on` command.

At this point, SNMP is configured for your ESX host. Traps will be sent to whatever device you configured as trapsinks, and you can add the host to any applications that support SNMP monitoring using the strings that you defined in the `rocommunity` and `rwcommunity`. You can use many additional SNMP commands for testing and advanced functionality, including configuring v3 security. These commands are located in the `/usr/bin` directory, and all begin with `snmp` (for example, `snmpconf`, `snmpptest`, and so on). For more information about using these commands and configuring advanced options, visit the `net-snmp` website at www.net-snmp.org/.

To configure SNMP to use with ESXi, follow these steps (can also be used with ESX to configure traps only):

1. You need to first install the Remote Command-Line Interface (RCLI) utility that is available for download from VMware's website (www.vmware.com/download/vi/drivers_tools.html) to configure SNMP with ESXi. This utility is available as an installable application for Windows or Linux computers or as a virtual appliance that can run on your hosts. The RCLI was initially developed for ESXi to provide some of the advanced command functionality that the Service Console provided for ESX. The RCLI works with both ESX and ESXi, and is basically a collection of Perl scripts that leverage the API that VMware developed to use command-line utilities to administer ESX and ESXi. After you have downloaded and installed the RCLI, you are ready to proceed.
2. Using the RCLI's `vicfg-snmp.pl` command, you can configure SNMP. The syntax for it is `vicfg-snmp.pl --server ESXi hostname/IP address --username root or another equiv. local account --password password [options]`.

The options that can be used with this command include the following:

- a. `-c community name` This is the name of the trap community string. You can enter multiple strings separated by a comma.

- b. `-t target hostname/IP address@port/community name` This sets the target host to receive the traps for the specified community name and the port number (for example, `-t 172.13.20.55@162/mycommunity`).
 - c. `-D` This disables the SNMP service.
 - d. `-E` This enables the SNMP service.
 - e. `-p port number` This sets the port number used by SNMP. If not specified, the default port of UDP 162 is used.
 - f. `-s` This shows the current SNMP configuration.
 - g. `-r` This resets all the currently set community names and targets.
 - h. `-T` This sends a test notification that can be used to validate your SNMP configuration.
3. After you have used `vicfg-snmp.pl` to set your community name and targets, you are all set. You can use the `-T` option to test out your notifications and make sure they are received by your targets.

SNMP is a great way to integrate your virtual environment into an existing monitoring system and is also useful if you do not have a vCenter Server to use for sending alerts. If you do have a vCenter Server, you can configure it to also send SNMP alerts in addition to its other alerting mechanisms, which we cover in the next section. You may also check out some of the free SNMP monitoring applications, such as Nagios, which can be used as a virtual appliance (www.vmware.com/appliances/directory/330), and Solarwinds VM Monitor (www.solarwinds.com/products/freetools/vm_monitor.aspx).

Monitoring with vCenter Server

vCenter Server enables you to configure alarms for certain conditions and take action when a condition is met. The alarms are fairly basic and cover only a few conditions, but are useful nonetheless to alert you when things happen in your environment. In addition to using the alarms to monitor your ESX and ESXi hosts, you should use any existing monitoring systems that you might have to monitor the vCenter Server itself and its database.

Using vCenter Server Alarms

You can configure alarms in vCenter Server that will trigger based on a defined condition and perform the action that is specified with the alarm. These alarms can be set at many different levels, all the way from the top datacenter level down to the individual VM level. Alarms inherit from the level that they are set at down to all child levels under them, as shown in Figure 7.7.




Figure 7.7 Alarm definitions inherited from the datacenter level to the host level

This inheritance cannot be changed at the child levels, and any additional alarms that are set at lower levels combine with the ones that are set at higher levels. Therefore, take care and set your alarms at a low enough level so that you do not have unwanted alarms on different objects. Even though alarms can be set at many levels, they apply only to host and VM objects (not clusters or resource pools).

Configuring SMTP for Email Alerts

Before you configure alarms, you should set up your vCenter Server so that it can send alerts via email. SMTP alerts are just one of the alert mechanisms that can be used when an alarm is triggered. To configure SMTP, follow these steps:

1. First make sure that the SMTP port (25) is open on any firewalls between your vCenter Server and your email server. In addition, some antivirus applications block SMTP email because it is used to spread viruses, so check this, too. You can test this by going to a CMD prompt and typing `telnet` to launch the Telnet application and then typing the following commands:
 - a. `open SMTP server hostname/IP address 25` Opens a connection to your SMTP server on port 25. Depending on your SMTP server configuration, you may see a welcome message.
 - b. `helo your PC hostname/IP address` Sends hello command to SMTP server identifying yourself.
 - c. `mail from: email address` This is the email address that you want to show as it originated from.
 - d. `rcpt to: email address` This is the destination email address that you want to send to.
 - e. `data text of email` This is the body of your email. When you have finished, press Enter and then a period and then Enter again to end text entry, and your email will be sent.
 - f. `quit` This ends your session with the SMTP server.

2. Connect to your vCenter Server with the VI Client and select Administration from the top menu and then VirtualCenter¹ Management Server Configuration.
3. Select Mail in the left pane and enter your SMTP server hostname/IP address and an email address that you want to show as being sent from, as shown in Figure 7.8. Click OK when you have entered this.




Figure 7.8 Configuring an SMTP mail server in vCenter Server

Your vCenter Server is now configured to be able to send email alerts for any alarms that you create that are configured to send an email.

Configuring SNMP Alerts

You can also configure vCenter Server to send alerts as SNMP traps to an external monitoring system as an action to a trigger. These alerts are sent to a host or device name/IP address and to the community string that you specify, which is the password that is used to read and write data. If you plan to use SNMP alerts with vCenter Server, configure it to use SNMP as outlined here:

1. Connect to your vCenter Server with the VI Client and select Administration from the top menu and then VirtualCenter¹ Management Server Configuration.
2. Select SNMP in the left pane and enter your SNMP server hostname/IP address in the Receiver field followed by the port number (162). Below that enter the

community string name to be used for your receiver, as shown in Figure 7.9. You can also enter up to four additional receivers if needed. Click OK after you have entered this.




Figure 7.9 Configuring SNMP receivers in vCenter Server

Your vCenter Server is now configured to be able to send SNMP traps for any alarms that you create that are configured to use SNMP.

Configuring Alarms

Now you are ready to start configuring alarms. There are a few default alarms configured at the hosts and cluster level that you can choose to use or delete and create your own.

Remember that because these default alarms are at the highest level, they also apply to all objects at lower levels. To create an alarm, follow these steps:

1. First decide which level/object that you want to place the alarm on and select the object in vCenter Server and click the Alarm tab.
2. Click the Definitions button, and all alarms for that object will display, as will any that are inherited from higher levels that are displayed in the Defined In column, as shown in Figure 7.10.




Figure 7.10 Alarms displayed in the VI Client on the Alarm tab

3. Right-click in the right pane and select the New Alarm option to create a new alarm.
4. The first tab is the General tab, as shown in Figure 7.11. Here you can enter a name and description for the alarm, the alarm type (host or VM), the trigger priority, and whether it is enabled. The type you choose will change the triggers that are available on the other tabs. The trigger priority is the order in which the alarm is reported. Red alarms are reported before yellow alarms are.




Figure 7.11 Creating a new alarm: General tab

5. The second tab, the Triggers tab, is for defining the triggers for the alarm. Clicking the Add button will add a new trigger, which you can customize. Clicking the default trigger type that appears displays a drop-down list of other available triggers (resource usage, state, and health), as shown in Figure 7.12. When you change the trigger type, the Condition, Warning, and Alert columns will also change to be specific to the type that was selected. You can select a value for both the yellow (warning) and red (alert) conditions, and you can then set an action on each on the Actions tab. You can also define multiple triggers that must all be met for the alarm to trigger.




Figure 7.12 Creating a new alarm: Triggers tab

6. The third tab is the Reporting tab and is for defining the reporting behavior of the trigger, as shown in Figure 7.13. You can define both a tolerance and frequency for the trigger. The tolerance is a percentage for those triggers that rely on numeric values. When an alarm is triggered, it will not trigger again unless the value is above or below the tolerance range. For example, if you have a red trigger for Host CPU usage set for 90% and you set the tolerance to 10%, it will not trigger again unless the CPU usage goes below 81% or above 99%. When the CPU usage goes outside of the tolerance range, it triggers again at the defined value of the trigger. The frequency is the amount of time that must pass after an alarm has triggered before it can trigger again. Both of these settings are designed to limit the reporting actions of the alarms as to not generate too many alarms for a single event.




Figure 7.13 Creating a new alarm: Reporting tab

7. The last tab is the Actions tab and is where you define the actions that will take place when an alarm is triggered. The available actions, as shown in Figure 7.14, include sending an email, sending an SNMP trap, and running a script. When you select an action, you define a value for that action. For an email action, the value is an email address to send to. For SNMP traps, no value is needed because it just sends to the receivers defined in the vCenter Server settings. For a script, you define a path (relative to the vCenter Server) and filename to an application or script that you want to run. You can use parameters when defining a script to run, and the alarm values are also passed, including the entity name, alarm name, and alarm state. Also on the Actions tab you define the condition for the action to fire. The condition is the change of state (for example, green to red). You can create multiple actions for an alarm. For example, you might have one that sends an email and sends an SNMP trap to a monitoring application. For VM alarms, you can also perform actions such as changing the power state of the VM.
8. After you have defined all the settings for the alarm, click OK to save it, and it will be listed for the object that you set it on and all objects below it.

You might note that there are not many alarm triggers and the settings are not that robust. For example, if you set an alarm for VM CPU usage at 90%, you will probably not want to be alerted if it reaches that level for only a few seconds. Unfortunately, you cannot set

an alarm trigger that it must be above 90% for a certain length of time, so you will experience many nuisance alarms. If you want to have more robust and powerful monitoring capabilities, check out some of the third-party products listed later in this chapter as an alternative to using vCenter Server for monitoring.




Figure 7.14 Creating a new alarm: Actions tab

Monitoring the vCenter Server

vCenter Server can monitor your ESX and ESXi hosts for you, but you should also monitor the vCenter Server and its components so that you can be alerted when problems happen. Although a problem with vCenter Server will not directly impact your hosts and VMs, it can still disrupt your environment because of the many advanced functions (for example, VMotion, DRS) that rely on your vCenter Server to be functioning properly. Therefore, if you have an enterprise monitoring system, use it to monitor your vCenter Server and its components, including the following:

- Monitor all the vCenter Server Windows services to ensure that they are running.
- Monitor the health and performance of the vCenter Server, including CPU and memory usage and that adequate disk space is available.
- Monitor the database that vCenter Server relies on to ensure it has adequate space and is performing properly.

- Monitoring the license server (Windows service) that is used by vCenter Server for centralized licensing. If the license server stops functioning, there is a 14-day grace period before certain host operations are no longer permitted (for example, powering on VMs). Because there are no obvious signs when the license server stops, you may find yourself past the grace period before you find out that it has stopped, which could cause problems. Monitor the Windows service to ensure it is running and the port (27000) that it uses to ensure it is listening on that port.

Performance Monitoring

Monitoring the performance of your hosts and VMs is very important in virtual environments because many VMs are competing for host resources and a single bottleneck can have a great impact on the performance of your VMs. Resource bottlenecks are not always obvious, and monitoring the performance of your hosts can help identify them so that you can take action to correct them. When monitoring performance of VMs, you should rely on tools designed for virtual environments because many operating system tools, such as Windows PerfMon, are not aware of the underlying virtualization layer and will often provide inaccurate results on certain counters and measurements. ESX/ESXi and vCenter Server have built-in performance monitoring tools that you can use to more accurately monitor the performance of your hosts and VMs and are the best choice when monitoring VM performance. vCenter Server provides configurable performance statistic collections so that you can view both real-time and historical statistics going back as far as you want. ESX and ESXi hosts that are not managed by a vCenter Server are limited to providing real-time statistics and up to 60 minutes of historical statistics.

To view real-time and historical statistics, you select the object in the left pane of the VI Client that you want to view statistics on, which can be a cluster, host, VM, or resource pool, and then select the Performance tab in the right pane, as shown in Figure 7.15.

Here you can switch between resource types to be displayed and customize the chart options to include specific individual statistics for the resource that is selected. You can also save the chart to an image file or export the data to an Excel spreadsheet. The charts are customizable by clicking the Change Chart Options link, and so you can choose the chart type (line or stacked graph), the date range, and the objects and counters that are included, as shown in Figure 7.16.

Now that you know how to display the performance statistics, let's talk about how to interpret them.

Performance Monitoring




Figure 7.15 Performance tab in the VI Client




Figure 7.16 Customizing performance chart display options

Understanding Virtual Machine Performance Metrics

Many different performance statistics are gathered for your VMs for its CPU, memory, disk, and network resources. Each resource category has many different types of statistics that are collected and can be used to analyze the performance and resource usage of your VMs. Let's cover some of the key statistics for each resource category that you should be aware of:

VM CPU Statistics

- **CPU Usage (average).** This is a basic measurement of CPU usage and is a percentage of the amount of CPU a VM is using compared to the total of all megahertz from all CPUs that are available on a host, which is 100%.
- **CPU Usage in MHz (average).** Another basic measurement of CPU usage, but displayed in megahertz rather than as a percentage.
- **CPU Ready.** This is the amount of time that is spent waiting for a CPU to become available. High ready times (> 10%) can indicate a CPU bottleneck or too many vSMP VMs on a host.
- **CPU Used.** The amount of CPU time that is used, in milliseconds. A high value can indicate that a VM is saturating its CPU and may benefit from an additional vCPU, especially if this value is high and CPU Ready is low.

VM Disk Statistics

- **Disk Read Rate.** This is a basic measurement of disk usage by a VM and is displayed as a kilobytes per second (KBps) measurement for the rate of disk reads by the VM.
- **Disk Write Rate.** This is a basic measurement of disk usage by a VM and is displayed as a kilobytes per second measurement for the rate of disk writes by the VM.
- **Disk Usage (average).** Another basic measurement that shows the average kilobytes per second of disk usage used by a VM.
- **Disk Command Aborts.** The number of disk command aborts, which is a result of the disk subsystem having not responded in an acceptable amount of time. If a VM is waiting too long to access its virtual disk because of contention with other VMs, this can indicate a disk I/O bottleneck or poorly configured disk (for example, too many VMs per LUN).

VM Memory Statistics

- **Memory Usage (average).** This is a basic measurement of what percentage of the memory allocated to the VM that it is actually using.

- **Memory Active (average).** The amount of “true” memory used by the VM in the past small window of time. Any unused memory may be swapped out or ballooned with no impact to the guest’s performance.
- **Memory Swap In (average).** The rate at which memory is being swapped in from a VM’s virtual disk swapfile. A large number here represents a problem with lack of memory and is a clear indication that performance is suffering as a result.
- **Memory Swap Out (average).** The rate at which memory is being swapped out to a VM’s virtual disk swapfile. A large number here represents a problem with lack of memory and is a clear indication that performance is suffering as a result.
- **Memory Consumed (average).** The amount of actual host memory that is in use by the VM, which does not include any shared memory that a VM may be using. If a VM has 2GB of RAM assigned to it and is using only 512MB of consumed memory, this can be an indication that you can reduce the amount of memory assigned to the VM.
- **Memory Balloon (average).** The amount of memory that is in use by the balloon driver inside the guest operating system of the VM. A value greater than zero indicates the host is starting to take memory from less-needful VMs for those with large amounts of active memory.

VM Network Statistics

- **Network Usage (average).** This is a basic measurement of the amount of network bandwidth a VM is using measured in kilobytes per second.

Besides these statistics, many additional statistics can be monitored for each resource, with some being more useful than others. Take a look at the other statistics for a specific resource if you suspect a performance issue with your VMs.

Understanding Host Server Performance Metrics

Many of the resource statistics used for monitoring host performance are similar to the ones used by VMs, but they can be interpreted differently because they apply to a whole host and not just a single VM. Here are some of the key resource statistics to look at for your host servers.

Host CPU Statistics

- **CPU Usage (average).** This is a basic measurement of CPU usage and is the total amount of CPU usage by all the VMs on a host (including the Service Console but excluding the VMkernel).

- **CPU Usage in MHz (average).** Another basic measurement of CPU usage, but displayed in megahertz rather than as a percentage.

Host Disk Statistics

- **Disk Usage (average).** This is a basic measurement of the total size of data read and written for all disk instances of the host.
- **Disk Command Aborts.** The number of disk command aborts, which is a result of the disk subsystem having not responded in an acceptable amount of time. If a host is waiting too long to access a disk because of contention with other hosts, this can indicate a disk I/O bottleneck or poorly configured disk (for example, too many hosts accessing a LUN).
- **Disk Read Latency.** Average time taken for a read by a guest OS, the sum of kernel read latency and physical device read latency. High latency times can indicate a disk I/O bottleneck or poorly configured disk (for example, too many VMs per LUN).
- **Disk Write Latency.** Average time taken for a write by a guest OS, the sum of kernel write latency and physical device write latency. High latency times can indicate a disk I/O bottleneck or poorly configured disk (for example, too many VMs per LUN).
- **Disk Command Latency.** Average time taken for disk commands by a guest OS, the sum of kernel command latency and physical device command latency.

Host Memory Statistics

- **Memory Usage (average).** This is a basic measurement of what percentage of the host's total memory is in use by all the VMs, including the Service Console and VMkernel memory usage.
- **Memory Active (average).** The amount of memory that is actively being used by the all the VMs on a host. This is just the total amount that has been recently used and is not the total amount being consumed by all the VMs on a host, including the Service Console and VMkernel memory usage.
- **Memory Consumed (average).** The amount of actual host memory that is in use by all the VMs on a host plus the overhead memory and memory used by the Service Console. This is the memory usage total that is displayed on the Summary tab of a host.
- **Memory Granted (average).** The amount of memory that was granted to all VMs on a host. Memory is not granted by the host until it is touched one time by a VM, and granted memory may be swapped out or ballooned away if the VMkernel needs the memory.

- **Memory Swap Used (average).** The amount of swap memory (VSWP files) that is in use by all the VMs on a host. A high number can indicate the host is running low on physical memory or that VMs with memory limits set below their allocated memory are exceeding their limits and using swap memory rather than physical host memory.
- **Memory Overhead (average).** This is the amount of additional host memory used by the VMkernel to maintain and execute the VM. This memory is separate from the memory that is assigned to a VM for its use.
- **Memory Balloon (average).** The amount of memory ballooning that is occurring on all the VMs on a host. A value greater than zero indicates the host is starting to take memory from less-needful VMs to allocate to those with large amounts of active memory that the VMkernel cannot satisfy with the current free memory available.
- **Memory Shared (average).** The amount of memory pages shared by all VMs on a host as a result of the transparent page sharing (TPS) feature. A high number here is good because it results in physical memory savings on the host.

Host Network Statistics

- **Network Usage (average).** This is a basic measurement of the amount of network bandwidth being used by all VMs on a host and is measured in kilobytes per second. This includes all NICs in a host and VMotion and Service Console traffic.

When looking at host statistics, remember that you can look at both individual component statistics or a combined total of all of them. For example, you can look at statistics for individual NICs, disk controllers, or CPUs, or you can select the host to view all of them at once.

Configuring vCenter Server Performance Statistics

As previously mentioned, ESX and ESXi hosts that are not managed by vCenter Server provide only real-time statistics and historical statistics for the past 60 minutes. vCenter Server enables you to save historical statistics for your VMs and hosts for up to 5 years, depending on how you configure the statistic collection settings. This historical statistic information is stored in one of the database tables that are used by vCenter Server. As a result, this table can grow very large depending on the number of hosts and VMs that are managed by vCenter Server and what you set the statistic intervals to. To access the configuration settings for statistics, connect to your vCenter Server and select Administration from the top menu and then choose VirtualCenter¹ Management Server Configuration, and then select Statistics in the right pane, as shown in Figure 7.17.




Figure 7.17 Changing statistics collection settings in vCenter Server

Here you can choose which statistic intervals that you want to use and for how long and how much statistic detail is kept. Changing these settings can drastically increase the size of your vCenter Server database. To help with this, a database size estimator tool was introduced in one of the vCenter Server 2.0.x updates so that you can see the approximate size of the database based on the settings you choose. For example, using the default settings of statistic level 1 and keeping just 1 day/week/month/year of each interval for 20 hosts and 500 VMs, the resulting database would be approximately 810MB in size. Changing the settings to keep the 5-minute interval statistics from 1 day to 5 days increases the size to 1.5GB. Further changing the statistics level from 1 to 4 will result in a database that will end up approximately 13.2GB in size. Increasing all the settings to their maximums would result in a database that would be approximately 85GB in size. So as you can see, you should plan accordingly if you want to keep more statistics so that you have enough disk space available on your database server. In addition, as the database grows larger, its performance will become reduced, which will result in vCenter Server being slower when viewing performance statistics. Therefore, try to keep the amount of data that you want to retain to a minimum based on your needs.

The statistic collection settings consist of three items: the sample interval duration, how long to keep the samples for, and the level of detail of the statistic:

■ Sample interval duration

Four durations can be used to gather statistics. The shortest level can be set from 1 to 5 minutes and is used for daily statistics. The next level is set at 30 minutes and cannot be changed and is used for weekly statistics. The next level is set at 2 hours and cannot be changed and is used for monthly statistics. The final level is set at 1 day and cannot be changed and is used for yearly statistics. These interval durations specify the length of time during which statistics are collected for each level. The weekly, monthly, and yearly statistics are rolled up from the previous levels from an average of the short duration samples taken from the previous level and are not the result of samples taken at the interval duration. You can choose to disable statistics collections for any level if you do not want to collect statistics for that period.

■ Amount of time to keep samples for

This is the amount of time to keep a sample collection for. The 1- to 5-minute samples can be kept for up to 5 days (default is 1). The 30-minute samples are kept for 1 week and cannot be changed. The 2-hour samples are kept for 1 month cannot be changed. The 1-day samples can be kept for up to 5 years (default is 1). Older samples are automatically deleted when this interval is reached so that you never have more samples than defined by this setting. The statistic data is aggregated and rolled up into the subsequent collection frequency if that interval is enabled. For example, if the Day interval is set to a 5-minute collection frequency and the Week interval a 30-minute collection frequency, when the Day interval completes, it aggregates the 5-minute queries into groups of 6 (equaling 30 minutes) and rolls the 30 minute data block into the Week interval database table. The day-old data is then purged from the database to make room for new data.

■ Statistics level

This determines how much of the many statistics that are available are stored in the database. It can be set from the most basic level of 1 (default) all the way up to 4, with each incremental level saving more statistics. The statistics level for a collection interval cannot be greater than the statistics level set for the preceding collection interval. For example, if the Month interval is set to collection level 3, the Year interval can be set to collection level 1, 2, or 3, but not to collection level 4. Setting these levels higher can dramatically increase the size of the vCenter Server database and decrease the performance of the vCenter Server. Therefore, take care when changing these levels and only increase them as needed, such as when troubleshooting performance problems. If you do increase these levels, make sure you remember to decrease them when you are done; otherwise, you might end up filling up your database server. The detail that each level provides is described here:

- ◆ **Level 1.** Basic metrics, including average usage for CPU/memory/disk/network counters, system uptime, and heartbeat and DRS metrics

- ◆ **Level 2.** All metrics for CPU/memory/disk/network counters, average, summation, and latest rollout types (does not include maximum and minimum rollout types), system uptime, and heartbeat and DRS metrics
- ◆ **Level 3.** All metrics for all counter groups, excluding those for maximum and minimum rollout types
- ◆ **Level 4.** All metrics for all counter groups supported by vCenter Server

The combination of these three settings will influence the number and size of the statistics kept. Existing statistic data is reset (lost) when you change the interval configuration. However, only the data for that specific interval is reset. For example, if you change only the weekly time interval, the daily and monthly data is retained.

Using the ESXTOP Utility

For ESX hosts (not ESXi), you can run a command-line utility that is included in the Service Console called ESXTOP to monitor and collect data from CPU, memory, disk, and network resources. For ESXi hosts, you can use the Linux version of the RCLI (not available in the Windows version), which contains a utility called RESXTOP that works just like ESXTOP does for ESX. In addition, you can download and use the VMware Infrastructure Management Assistant (VIMA), a virtual appliance that can be used with both ESX and ESXi hosts and that includes the ESXTOP utility. The ESXTOP utility provides real-time reporting of ESX worlds, which are processes running on the VMkernel. There are different types of worlds, including system worlds for running various system functions, a Service Console world that runs on CPU0, and a world for VMs. When you first start ESXTOP, it might look a little intimidating, but once you know how to use it and interpret the information displayed, it can be a valuable tool for troubleshooting performance problems. To access ESXTOP, log on to the Service Console of an ESX host and type `esxtop` (lowercase). The initial screen will display, as shown in Figure 7.18.

```

9:41:20am up 5 days 22:54, 67 worlds; CPU load average: 0.03, 0.03, 0.03
PCPU(%):  2.27,  3.86 ;  used total:  3.06
LCPU(%):  1.72,  0.55,  3.30,  0.55
CCPU(%):  0 us,  0 sy,  99 id,  0 wa ;      cs/sec:  184

```

ID	GID NAME	NWLD	%USED	%RUN	%SYS	%WAIT	%RDY
1	1 idle	4	192.89	199.87	0.00	0.00	198.04
2	2 system	6	0.00	0.00	0.00	596.87	0.00
6	6 helper	23	0.00	0.01	0.00	2287.98	0.01
7	7 drivers	15	0.04	0.01	0.00	1492.15	0.00
8	8 vmotion	1	0.00	0.00	0.00	99.47	0.00
9	9 console	1	1.23	1.61	0.01	97.73	0.14
19	19 vmware-vmkauthd	1	0.00	0.00	0.00	99.47	0.00
20	20 Win2003-1	6	0.72	0.96	0.00	595.67	0.19
21	21 Win2003-new	5	3.15	1.02	0.00	496.17	0.19
22	22 Win2000-1	5	0.22	0.40	0.00	496.89	0.09

Figure 7.18 ESXTOP utility running on an ESX host

The default display usually only shows a few columns, but if you are logged on via Secure Shell (SSH), you can expand the window to display more columns. There are no menus, and all controls are done with single keystrokes. Here is a list of some of the most common keystrokes used with ESXTOP:

- **h** Displays a help menu for the current screen giving a brief summary of the available commands
- **f** Displays a screen that lets you add or remove columns from the current screen
- **o** Displays a screen that lets you change the order on the current screen
- **s** Change the delay between screen updates
- **q** Quit from ESXTOP
- **c** Switch to CPU Resource Utilization screen
- **m** Switch to Memory Resource Utilization screen
- **d** Switch to Storage Adapter Resource Utilization screen
- **u** Switch to Storage Device Resource Utilization screen
- **v** Switch to Storage VM Resource Utilization screen
- **n** Switch to Network Resource Utilization screen

ESXTOP is commonly used to investigate CPU performance and scheduling problems. The columns show statistics, and the rows contain VMs on the host and other system objects. The %RDY (Ready Time) column is one of the key indicators of a CPU scheduling problem. Ready Time indicates the percent of time that the VM is ready to execute a command but is waiting to be scheduled CPU time. The percentage is based on the number of CPUs a VM has (100% for each). So a VM with a single vCPU can have a %RDY value from 0% to 100%, a two vCPU can have a value from 0% to 200%. In general, you want to see %RDY values of less than 10% for a healthy system. Higher values can indicate performance and scheduling problems as VMs are competing for CPU time and having to wait for it. Having too many multiple vCPU (vSMP) VMs can be one cause of this. The Idle column is an exception to this and having high values there is okay.

There are many ways to use this handy utility, and for additional information about using ESXTOP, you can type `man esxtop` when logged on to the Service Console to display a help manual for the command. In addition, VMware has published some good information about using this utility that is available at the following websites:

- Using ESXTOP to Troubleshoot Performance Problems
www.vmware.com/pdf/esx2_using_esxtop.pdf
- Ready Time Observations
www.vmware.com/pdf/esx3_ready_time.pdf

- Interpreting ESXTOP Statistics
<http://communities.vmware.com/docs/DOC-9279>
- ESXTOP Performance Counters
<http://communities.vmware.com/docs/DOC-5240>
- Ready Time
<http://communities.vmware.com/docs/DOC-7390>
- CPU Performance Analysis and Monitoring
<http://communities.vmware.com/docs/DOC-5420>

Third-Party Reporting and Monitoring Tools

Although vCenter Server can provide some great statistics on your environment, its reporting and monitoring capabilities are not all that robust. Fortunately, a number of third-party products can provide this for you. If you are looking for a better application to provide this function in your environment, check out some of the following products that provide enhanced reporting and monitoring for your virtual environment.

Vizioncore vFoglight (Formerly vCharter Pro)

vFoglight is a mature and powerful reporting and monitoring application that is specifically for monitoring and analyzing ESX hosts. It integrates with vCenter Server and provides enhanced reporting capabilities, configurable dashboards, intelligent rules, and alarms and custom reporting to identify trends and bottlenecks. It has built-in intelligence that goes beyond simply reporting statistics; it analyzes them to help you better understand what is happening with your ESX hosts. vFoglight is a good choice for use in any ESX environment.

Veeam Management Suite

This product suite includes the Reporter, Monitor, and Configurator products, which can also be purchased individually. Reporter automatically discovers and collects information about VI3 environments and provides reports for analysis and to document the environment. Monitor integrates with vCenter Server and provides enhanced monitoring and alerting of the health and performance of your environment. It can also be used without vCenter Server to monitor individual ESX hosts. Finally, Configurator helps you more easily configure ESX hosts and provides a GUI to change settings that normally are changed from the ESX command line.

eG VM Monitor

Another powerful reporting and monitoring application designed specifically for ESX hosts. eG VM Monitor is a web-based application specifically for VMware environments and is part of the larger eG Enterprise Suite. It includes support for both agent-based and agentless (for ESXi) monitoring of servers and only requires installation on the ESX host, not on the individual VMs. This product is rich with features and uses an In-N-Out monitoring approach to provide an outside view of a VM's performance. It has extensive reporting capabilities and can analyze the performance across applications hosted in the VMware environments to help discover VM dependencies and to identify performance bottlenecks.

Summary

Monitoring your environment is important to ensure that it stays healthy and functions properly. Often, problems might not be all that obvious, and having a good monitoring system in place will alert you to potential or existing problems so that you can resolve them. In virtual environments, even small problems can have big effects because so many VMs are running on a single host and all are contending for the resources that the host provides. Therefore, it's important to not ignore monitoring. Without it, your virtual environment may be trying to tell you something but you will not hear it because you are not listening.

Endnotes

1. The use of the term *VirtualCenter* here refers to the old name, which still applies to VirtualCenter version 2.0.x and is still present inside the application in vCenter Server 2.5 (because the software has not yet been updated to reflect the new name).

This page intentionally left blank

Chapter 8

Maintaining Your Virtual Environment

With any software application, patching and upgrades are a necessity to ensure that bugs and security vulnerabilities are resolved and new features are added. With virtual hosts, it is even more important to apply patches in a timely manner because the effect of a bug or vulnerability can affect many VMs running on that host. VMware usually releases patches for ESX and ESXi every month, although not on a set schedule or day. In addition to patching your host servers, many patches require updates to the VMware Tools application that runs on your VMs. In addition, VMware releases periodic updates to vCenter Server to fix bugs and occasionally to add new features and functionality. When it comes to patching and upgrading your whole virtual environment, there are often dependencies as to what you should patch first. In most cases, you should first patch vCenter Server, then your host servers, and finally your VMs (VMware Tools). In this chapter, we cover the methods for patching each one of these.

Patching and Updating ESX and ESXi Hosts

Two types of updates can be applied to ESX hosts: patches and updates. Patches are specific fixes for bugs and vulnerabilities and are released monthly. Updates are new versions of ESX that contain all the patches released since the earlier version and sometimes new features and are released every three to four months. VMware has no set schedule for either patches or updates and does not communicate release dates in advance:

- **Patches.** You should periodically check VMware's website so that you will know when new patches are released. Patches can be downloaded at <http://support.vmware.com/selfsupport/download/>. Patches are specific to certain versions of ESX (for example, ESX 3.5 Update 3) and sometimes have dependencies that other patches be installed before a specific patch is applied. Patches are released under three classifications: General, Security, and Critical. General patches are the least critical and usually fix minor issues and bugs. Security patches should be considered critical and fix known vulnerabilities. Critical patches fix issues and bugs that can cause issues and crashing with hosts and VMs. VMware's patch web page lists each patch, the version of ESX that it applies to, a description of the patch, the system impact, a link to an appropriate knowledge base article that discusses the problem that the patch resolves, and the patch requirements. The system impact of each patch is important because some patches require that the VMs on a host be shut down when applying the patch and that the host be rebooted after the patch is installed. Patches come packaged as either bundles or roll-ups. A bundle is a set of metadata files and RPM packages for a specific fix or update to a host. A roll-up is a collection of bundles that groups together a set of single bundles. Roll-ups provide you with a convenient way to download and install multiple bundles at one time.
- **Updates.** Updates can be downloaded at <https://www.vmware.com/download/vi/> and differ from patches because they are new releases of the product. With ESX 3.0, VMware released updates as new version numbers (for example, 3.0.1, 3.0.2), but with ESX 3.5 they instead left the version number the same and referred to the new releases as updates (for example, 3.5 Update 1, 3.5 Update 2). However, even though the version number stays the same with each new release, the build number of the product is changed for each update. Updates are much larger than patches (usually 400MB+), and they install over any previous releases. Several different download files are available for updates: an ISO file that contains the full version of ESX, an upgrade package (tar file) that can be used to upgrade ESX 2.x servers, and upgrade packages (zip file) that can be used to upgrade ESX 3.0.x and ESX 3.5 hosts. Before installing any updates, always read the release notes for it.

Methods for Patching

Because of their architectural differences, there are different methods used for patching ESX and ESXi hosts and one method that can be used to patch both. ESX hosts can be patched using a Service Console command-line utility called `esxupdate`. ESXi hosts can be patched using either a remote command-line utility called `vihostupdate` or using a Windows application that is installed with the VI Client called Infrastructure Update. Both ESX and ESXi hosts can also be patched using the Update Manager plug-in that is included with vCenter Server.

Update Manager was released as part of vCenter Server 2.5 and can be used to patch both host servers and VMs running Windows and Linux operating systems. It works by establishing a baseline of your host and VMs and then scanning them to identify variances from the baseline, which you can then remediate as needed by applying any patches that are available that are not included in the baseline.

Patching ESX Hosts with Esxupdate

ESX hosts can either be patched using the Service Console `esxupdate` utility or using the vCenter Server Update Manager plug-in. Update Manager is the easier way to patch your hosts because it uses a graphical interface using the VI Client, but if you do not have this available, `esxupdate` will do the job. In this section, we cover only using `esxupdate` to patch ESX hosts and will cover how to use Update Manager later in this chapter.

`Esxupdate` is included in the ESX Service Console, but you do need to set up a patch repository to store your patches to be used with `esxupdate`. This repository can be a web, FTP, or NFS server, and you can also set it up on an ESX server, but this is not recommended because of the limited disk space on the ESX Service Console disk partitions. The repository can act as a central location for your patches that all your ESX hosts can use as a source for their patches. If you have an existing web, FTP, or NFS server, you can use that; otherwise, you can set up one of the free servers such as Apache to act as your repository. When you set up a repository, you then create depots, which are directories that contain the extracted patch zip files. You can set up multiple depots that are specific to certain versions of ESX (for example, 3.0.3, 3.5).

The `esxupdate` command can operate in four different modes depending on the options that are used with the command (for example, `esxupdate option`), which are listed here:

- **Inspection.** This mode is used to query for installed bundle information and the options for this mode are `query` (displays installed bundles) and `info` (displays bundle information).
- **Scan.** This mode scans the host to determine which bundles are applicable to it by querying the bundles in a depot and the bundles installed on the host for bundle and system dependencies using the `scan` option of `esxupdate`.
- **Test.** This mode lets you simulate a bundle installation without actually installing the bundle by downloading the files, checking for bundle and system dependencies, and determining the bundle order and RPMs to be installed using the `test` option of `esxupdate`.
- **Update.** This mode does the actual installation of the individual bundles and roll-ups and scans the depot for dependencies and resolves them if possible before installing them. This mode uses the `update` option of `esxupdate`.

Before we cover how to set up and use `esxupdate`, let's go over what a bundle actually is. A bundle is a zip file that consists of multiple files, including the following:

- **Contents.xml file.** This file contains a reference to the bundle's `descriptor.xml` file and a list of the RPM files and their signatures.
- **Contents.xml.sig file.** This file is used to validate the contents file to ensure its integrity.
- **Descriptor.xml file.** This file contains various information about the bundle, including a summary of the patch, dependency information, and RPM details.
- **RPM files.** These files are the packaged installation files for the patches.
- **Header files.** These files are kept in a separate subdirectory, and there is one associated to each RPM file.

Did You Know?

RPM stands for Red Hat Package Manager and is used by Linux servers as a software package file format to install applications. These files are roughly equivalent to MSI files that are used by Windows servers to install applications.

Let's now cover how to set up and use the `esxupdate` utility to patch your ESX hosts. This includes the preparation steps for setting up your repository, downloading patch bundles, scanning your hosts, running a test, and finally installing the bundles.

Step 1: Setting Up a Patch Repository

Your patch repository can be a web (HTTP) server, an FTP server, or an NFS server. If you have an existing IIS, Apache, or another type of web or FTP server, you can leverage that to use for ESX or you can set up a VM running a free web or FTP server instead. You will need to make sure that whatever server that you use has access to download patches from the Internet and that all of your ESX hosts can access it. Create a directory on your web/FTP/NFS server to be used as your repository and allow anonymous read access to it. You should create a directory for each version of ESX that you are running (for example, `\ESX303`, `\ESX35`). These directories are considered depots. Depending on the protocol (`http-80`, `ftp-21`, `NFS-111`, and `2049`) you use for your depot, you will need to open those outgoing ports on the built-in ESX firewall. Although not recommended, you can also temporarily allow all outgoing connections by logging on to the Service Console and typing `esxcfg-firewall --AllowOutgoing`. When you are done, you can block them again by typing `esxcfg-firewall --blockOutgoing`.

Step 2: Downloading Patches

After you have set up your repository and depot, you can go to VMware's website (<http://support.vmware.com/selfsupport/download/>) and download the patch bundles to the appropriate directory on your repository server and then extract the bundle zip files into it. These files will have a name like ESX350-200810201-UG.zip. You need to also be sure to download any other patches that the patches you select may require. In addition, you need to download the latest contents.zip file and extract it to your depot because this is the catalog file for the patches. After you have your files downloaded, your depot should look similar to the one in Figure 8.1. If you choose to not use a depot, you can download the files to a temporary directory on your ESX host instead.




Figure 8.1 A patch depot set up on an FTP server

Step 3: Scanning for Applicable Bundles

Log on to the Service Console as root on the ESX host that you want to patch and type the command `esxupdate -d depot URL scan`. The depot URL should be the hostname/IP address of the repository server, including the protocol and path (for example, `http://172.20.20.111/depots`). If you are not using a depot, omit the `-d` option and make sure you are running `esxupdate` in the directory that the patches are located in. The results of the scan will display and show any applicable bundles, including the bundle names and summary. You can display further information by using the `--explain` option, as shown in Figure 8.2 (for example, `esxupdate -d depot URL --explain scan`).

```
[root@esx1 root]# esxupdate -d ftp://172.20.20.111/depots/ESX35 --explain scan
- Bundle Name - AppFlags --- Summary --- iFlags
ESX350-200810201-UG -----v Updates VMkernel, Service Console, hostd rmh
[ESX350-200810201-UG] cannot be installed when VMs are still running.
ESX350-200810202-UG -----o--- Updates ESX Scripts ---
[ESX350-200810202-UG] superseded by ['ESX350-200811402-SG'].
ESX350-200811401-SG -----v Updates VMkernel, hostd, and Other RPMs rm-
[ESX350-200811401-SG] cannot be installed when VMs are still running.
ESX350-200811402-SG ----- Updates ESX Scripts ---
```

Figure 8.2 Using esxupdate to scan a depot for patches

Step 4: Retrieving Bundle Information

You can view which bundles are already installed on your host by using the `esxupdate query` command. To see detailed information about all bundles in the depot, use the `esxupdate -d depot URL info` command. To see detailed information about specific bundles in the depot, type `esxupdate -d depot URL -b bundleID info`, as shown in Figure 8.3.

```
[root@esx1 root]# esxupdate -d ftp://172.20.20.111/depots/ESX35 -b ESX350-200811402-SG info
Product       : VMware ESX Server
Vendor        : VMware, Inc. (support@vmware.com)
Bundle ID:    : ESX350-200811402-SG
Release Date  : 2008-11-26T11:34:30-08:00
Summary       : Updates ESX Scripts
Description   :
This patch fixes the following issue:

    * The output of vdf -P on a LUN greater than 1TB is wrapped. Because the line spacing allows only 10 digits, some tools or scripts cannot use the value. This change increases the spacing to 15 digits.

Refer to KB 1007502 for more details
Requires      :
Conflicts with:
Obsoletes     : ESX350-200802403-BG, ESX350-200803208-UG, ESX350-200804408-BG, ESX350-200805502-BG, ESX350-200806202-UG, ESX350-200808202-UG, ESX350-200808408-BG, ESX350-200810202-UG
Will reboot after install : False
Maintenance Mode required : False
Bundle URL    : ftp://172.20.20.111/depots/ESX35/ESX350-200811402-SG
```

Figure 8.3 Using esxupdate to obtain detailed patch information

Step 5: Verifying Sufficient Free Disk Space

Before you can deploy patches to your ESX host, you need to ensure that sufficient disk space exists in the root partition. You should have at least twice the space free as the size of the bundles that you are installing. If you are installing a 250MB bundle, you need at least 500MB free. You can check free disk space by typing the `df -h` command, as shown in Figure 8.4.


```
[root@esxi root]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/cciss/c0d0p7 4.9G  1.3G  3.4G  27% /
/dev/cciss/c0d0p1 373M   30M  324M   9% /boot
/dev/cciss/c0d0p5  2.0G   33M   1.9G   2% /home
none            132M    0   132M   0% /dev/shm
/dev/cciss/c0d0p6  2.0G   33M   1.9G   2% /tmp
/dev/cciss/c0d0p3  4.0G  207M   3.6G   6% /var
/dev/cciss/c0d0p2  5.0G  120M   4.6G   3% /var/log
```

Figure 8.4 Using the `df` command to check space on Service Console disk partitions

Step 6: Running a Test Installation

A test install downloads the bundles to the host, checks for bundle and system dependencies, and determines the bundle install order. It also verifies that you have sufficient disk space free on your host. If the bundles included in the test require the host to be in maintenance mode and the VMs shut down, you will need to do this before running the test. To run a test on all bundles, type `esxupdate -d depot URL --test update`. To run a test on a specific bundle, type `esxupdate -d depot URL -b bundleID --test update`. The output will show whether each bundle will be installed and if not the reason why not, as shown in Figure 8.5.

```
[root@esxi root]# esxupdate -d ftp://172.20.20.111/depots/ESX35 --test update
INFO: No -b specified, selecting all bundles in depot.
INFO: Configuring...
INFO: The following bundles are being skipped.
INFO: [ESX350-200810202-UG] superseded by ['ESX350-200811402-SG'].
INFO: Preparing to install ['ESX350-200810201-UG', 'ESX350-200811401-SG', 'ESX350-200811402-SG']...
INFO: Downloading VMware-esx-vmkctl-3.5.0-130756.1386.rpm...
INFO: Downloading VMware-esx-tools-3.5.0-130756.1386.rpm...
INFO: Downloading VMware-esx-ima-qla4xxx-3.5.0-123630.1386.rpm...
INFO: Downloading kernel-vmnix-350.2.4.21-57.EL.123630.1386.rpm...
INFO: Downloading VMware-esx-vmkernel-3.5.0-130756.1386.rpm...
INFO: Downloading kernel-source-350.2.4.21-57.EL.123630.1386.rpm...
INFO: Downloading VMware-esx-iscsi-3.5.0-123630.1386.rpm...
INFO: Downloading VMware-esx-apps-3.5.0-130756.1386.rpm...
INFO: Downloading VMware-esx-scripts-3.5.0-130756.1386.rpm...
INFO: Downloading VMware-hostd-esx-3.5.0-130756.1386.rpm...
INFO: Downloading VMware-esx-srvrvmgmt-3.5.0-123630.1386.rpm...
INFO: Downloading VMware-esx-perftools-3.5.0-123630.1386.rpm...
INFO: Downloading VMware-esx-vmx-3.5.0-130756.1386.rpm...
INFO: Checking disk space and running test transaction...
INFO: --- TOTALS: 0 packages installed, 13 pending or failed, 0 removed, 0 excluded ---
INFO: Test install of ['ESX350-200810201-UG', 'ESX350-200811401-SG', 'ESX350-200811402-SG'] succeeded.
== The esxupdate --test report ==
----- ESX350-200811401-SG -----: installed
----- ESX350-200810201-UG -----: installed
----- ESX350-200810202-UG -----: not installed
[ESX350-200810202-UG] superseded by ['ESX350-200811402-SG'].
Package VMware-esx-scripts.1386 superseded by bundle ESX350-200811402-SG
----- ESX350-200811402-SG -----: installed
```

Figure 8.5 Using `esxupdate` to do a test installation of patch bundles

Step 7: Installing Bundles

When you are ready to install the bundles, make sure the host is in maintenance mode (if required), and then type `esxupdate -d depot URL update`. To install a specific bundle, type `esxupdate -d depot URL -b bundleID update`. You will see detailed installation information about the screen as the bundles are installed. When they complete, you will be notified, and the host will reboot if needed, as shown in Figure 8.6. (You can specify that the host not reboot when complete by using the `-n` option.) When the bundles have been installed, you can verify this by using the `esxupdate query` command. You will also notice when you select the host in the VI Client or use the `vmware -v Service Console` command that the build number will change.

```
INFO: | Transaction(s) Complete
INFO: Shutting down hostd...
INFO: Running esxcfg-boot to regenerate initrds...
INFO: Restarting hostd...
INFO: --- TOTALS: 13 packages installed, 0 pending or failed, 0 removed, 0 excluded ---
INFO: Install succeeded - please standby for reboot.
INFO: Rebooting in 5 seconds...

Broadcast message from root (pts/0) (Tue Jan 13 12:16:45 2009):
The system is going down for reboot NOW!
```

Figure 8.6 Patch installation with `esxupdate` is completed, and the host is rebooting.

By the Way

When the `esxupdate` utility is used, it logs all activity to `/var/log/vmware/esxupdate.log`. You can use this log to troubleshoot patch installation problems or to get detailed information about everything that happened when patching.

A lot more options can be used with the `esxupdate` command. For a complete guide on using the command to update your ESX hosts, see VMware's "Patch Management Guide" located on their website at http://vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_esxupdate.pdf.

Updating ESX Hosts

To update an ESX host to a new release (for example, Version 3.5 Update 3) without using Update Manager, you have two options. You can download the ISO file for the new release and use it to create a CD, which you can use to upgrade your host, or you can download the upgrade zip file and use it with the `esxupdate` utility to install the update as you would patch bundles. Which method you use is more personal preference; they both get the job done. However, using the ISO file can be easier if you have physical access to the server.

To update an ESX host using the ISO file, follow these steps:

1. Download the ISO file for the version that you want to install from VMware's website and burn it to a CD (or you can use it with a virtual CD-ROM drive if your host has a remote management board).
2. Place the host in maintenance mode to evacuate the VMs or manually move them to another host or optionally shut them down.
3. Place the install CD in the host and shut down the ESX host and reboot it. It will automatically boot from the CD.
4. The install is similar to a new install of ESX. Go through the first screens, and you will see a screen asking whether you want to upgrade the host or do a fresh installation, as shown in Figure 8.7. Choose the Upgrade option and go through the next screens (make sure you choose the correct disk to install on) until you get to the About to Upgrade screen. Review the few options listed, and then click Next to begin the upgrade.




Figure 8.7 Upgrading an ESX host using the install CD and choosing the Upgrade installation type

5. When the installer completes, click the Finish button. The host will reboot using the newly installed version of ESX.

To update a host using `esxupdate`, follow these steps:

1. The process for this is the same as installing patch bundles with `esxupdate`. Download the zip file for the update that you want to install from VMware's website. Make sure you choose the correct one for the version of ESX that you are running before the update is applied.
2. It's easier to create a new depot for the update rather than put it in with all the existing patches that you may have in your depot. You can create a new depot directory

(for example, ESX35U3) and unzip the file there. The file usually unzips to a very long subdirectory name, so you may want to rename it to something shorter (for example, update-123630 [build #]). Optionally, you can copy the contents of the unzipped file to a directory on the ESX host if you do not want to use a depot. (Just make sure you have enough disk space available.)

3. Type `esxupdate -d depot URL scan` to make sure your host can see the update. It will return a number of bundles and the update (for example, ESX350-Update03). You can also retrieve detailed information about the update by typing `esxupdate -d depot URL info`.
4. Put the host in maintenance mode, and then run a test before doing the actual install by typing `esxupdate -d depot URL --test update`.
5. When the test completes, you install the update by typing `esxupdate -d depot URL update`. (If you are installing it locally without a depot, just type `esxupdate update` instead.)

After you update your ESX hosts and reboot them, you can check the version number in the VI Client or by using the `vmware -v` command. There is no way to uninstall an update after you have applied it to an ESX host. If you do need to roll back to an earlier version, you will need to install the earlier version over the newer version. Also, you should download the VI Client to your workstation and reinstall it because an updated VI Client is included with ESX host updates.

Patching and Updating ESXi Hosts

Unlike ESX, where patches and updates replace only the specific files that have been modified since the earlier version, every ESXi patch is a full replacement of the hypervisor image and essentially a full update, including the firmware image, the VI Client install, and the VMware Tools install. When you apply a patch or update, a new firmware image is uploaded to the host, and the old image is kept as a backup in case you need to fall back to it. This is done while the host is running, without taking it down, because the host continues to run on the old image until it is rebooted. When you reboot the host, the new image is loaded, and the old image is saved as a backup.

To update ESXi, there are two methods that you can use if you do not have Update Manager. You can use the Windows-based application called VMware Infrastructure Update or the Remote Command-Line Utility (RCLI) called `vihostupdate`. Infrastructure Update is sometimes confused with Update Manager, so let's consider the difference between the two. Infrastructure Update is installed as a separate application with the VI Client 3.5 and is used only for patching ESXi installable and embedded versions and is unable to patch ESX hosts. Update Manager is a plug-in for vCenter Server that is more robust and is used to patch ESX and ESXi hosts and VMs.

To update an ESXi host using Infrastructure Update, follow these steps:

1. Ensure you have the VI Client installed on a workstation and launch the Infrastructure Update utility, which should be located under the Start menu in Programs\VMware.
2. When the application loads, click the Hosts tab, as shown in Figure 8.8. Make sure your host is checked, and then click the Apply button to check the host to determine whether it needs updating. Optionally, you can elect to download the patches yourself and then add them to the Package Cache by clicking the Add Files button on the Updates tab. You can then browse to the zip files containing the latest updates. You can also set it to automatically check for updates on a certain day and time.




Figure 8.8 Selecting an ESXi host to update in the Infrastructure Update utility

3. When it is done scanning the host, it will list all available updates for that host, as shown in Figure 8.9. You can see in this screen that the firmware (hypervisor) update is needed, an update to the VI Client is needed, and the VMware Tools images that are on the host server need updating.
4. When you click the Install Updates button, it will begin downloading the updates listed. When those have finished, you will be prompted to log on to the host before it installs the updates. The updates will then be installed, as shown in Figure 8.10, and the utility will install the updated firmware, VI Client, and VMware Tools images.




Figure 8.9 Available updates shown for an ESXi host




Figure 8.10 Updates being installed on an ESXi host

5. When the installation has completed, as shown in Figure 8.11, click Continue to finish, and then close the Infrastructure Update application.
6. If you connect to the ESXi host using the VI Client, you will see the old build number and a message that a reboot is required, as shown in Figure 8.12.




Figure 8.11 Completed update install on an ESXi host



Figure 8.12 Old build number displayed and reboot message after updating an ESXi host

7. Shut down or migrate any VMs that are running on the host, and right-click the host in the VI Client and select Reboot. When you restart and log back on to the VI Client, you will see the new build number, as shown in Figure 8.13.



Figure 8.13 New build number displayed after rebooting an ESXi host after it has been updated

To update an ESXi host using the `vihostupdate` RCLI utility, follow these steps:

1. This method requires that the RCLI utility be installed on a Windows or Linux workstation or by using the preconfigured virtual appliance. Both can be downloaded from VMware's website at www.vmware.com/download/vi/drivers_tools.html. The RCLI is a collection of Perl scripts that connect to the ESXi host using APIs and executes commands. It is a substitute for the actual program files that were located on the Service Console with ESX hosts. When you install the RCLI for Windows, it also installs the Active Perl application necessary to execute the Perl scripts, which can't be run in Windows natively. This method also requires you to manually download the updates and place them in a directory that is accessible by the `vihostupdate` utility.

- When you install the RCLI, you can access it through the Start menu under the VMware folder or by simply going to a command prompt and going to the RCLI directory. When you are in the directory, switch to the bin subdirectory. If you do a `dir`, you can see a list of all the Perl files that you can use, as shown in Figure 8.14.

```

C:\Program Files\VMware\VMware UI Remote CLI\bin>dir
07/12/2008 08:52 PM      8,209 vicfg-cfgbackup.pl
07/12/2008 08:52 PM      7,457 vicfg-dns.pl
07/12/2008 08:52 PM     19,437 vicfg-dumppart.pl
07/12/2008 08:52 PM      9,989 vicfg-module.pl
07/12/2008 08:52 PM     17,534 vicfg-mpath.pl
07/12/2008 08:52 PM      6,919 vicfg-nas.pl
07/12/2008 08:52 PM      8,938 vicfg-nics.pl
07/12/2008 08:52 PM      6,495 vicfg-ntp.pl
07/12/2008 08:52 PM      1,646 vicfg-rescan.pl
07/12/2008 08:52 PM      3,143 vicfg-route.pl
07/12/2008 08:52 PM      9,006 vicfg-snap.pl
07/12/2008 08:52 PM      4,166 vicfg-syslog.pl
07/12/2008 08:52 PM     31,497 vicfg-user.pl
07/12/2008 08:52 PM      4,309 vicfg-vmhbaudev.pl
07/12/2008 08:52 PM      8,815 vicfg-vmknic.pl
07/12/2008 08:52 PM     18,326 vicfg-vswitch.pl
07/12/2008 08:52 PM     13,445 vifs.pl
07/12/2008 08:52 PM     13,101 vihostupdate.pl
07/12/2008 08:52 PM     30,149 vmkfstools.pl
07/12/2008 08:52 PM         72 vmkuptime.pl
07/12/2008 08:52 PM     54,087 vmware-cmd.pl
         42 File(s)      638,554 bytes
         2 Dir(s)      9,249,935,360 bytes free

```

Figure 8.14 RCLI Perl scripts

- Download the update that you want to install as a zip file from the VMware ESXi web page at <http://support.vmware.com/selfsupport/download/> to the workstation that you are running the RCLI on. There is no need to unzip the file.
- Shut down or migrate your VMs on your ESXi host. If you do not, the `vihostupdate` command will be unable to place the host in maintenance mode. The syntax to run the command is `vihostupdate.pl -server server name or ip -username root or other user -password user password -i -b update filename`. The username and password are optional. If you omit one or both, it will prompt you to enter them when you run the script. The `-i` parameter specifies to install a patch bundle. The `-b` parameter specifies the bundle filename to install. When you enter the command, the install will begin, as shown in Figure 8.15.
- When the script runs, it first unpacks the zip file, then copies it to the ESXi host, installs it, and finally prompts you to reboot the host when complete, as shown in Figure 8.16.
- After the host reboots, you can check the host version by typing `vihostupdate.pl -server server name or ip -username root -q`, which will display the version and build numbers, as shown in Figure 8.17.

In addition to these two methods, you can patch and update ESXi hosts using the Update Manager component of vCenter Server, which we cover shortly.


```
C:\Program Files\VMware\VMware UI Remote CLI\bin>vihostupdate.pl -server 172.21.28.22 -username root -i -b c:\vmware\patches\ESX350-200810401-0-UG.zip
Enter password:
unpacking c:\vmware\patches\ESX350-200810401-0-UG.zip ...
  < skipping verification : ESX350-200810401-0-UG/ESX350-200810401-I-UG.zip.sig >
unpacking ESX350-200810401-0-UG/ESX350-200810401-I-UG.zip ...
  < skipping verification : ESX350-200810401-0-UG/ESX350-200810402-T-UG.zip.sig >
unpacking ESX350-200810401-0-UG/ESX350-200810402-T-UG.zip ...
  < skipping verification : ESX350-200810401-0-UG/ESX350-200810403-C-UG.zip.sig >
unpacking ESX350-200810401-0-UG/ESX350-200810403-C-UG.zip ...
Installing : ESX350-200810401-I-UG
Copy to server : VMware-image.tar.gz ...
Copy to server : VMware-OEM-image.tar.gz ...
Copy to server : descriptor.xml ...
Copy to server : install.sh ...
Copy to server : contents.xml.sig ...
Copy to server : contents.xml ...
```

Figure 8.15 ESXi host being updated using the RCLI vihostupdate utility

```
Installing : ESX350-200810401-I-UG
Copy to server : VMware-image.tar.gz ...
Copy to server : VMware-OEM-image.tar.gz ...
Copy to server : descriptor.xml ...
Copy to server : install.sh ...
Copy to server : contents.xml.sig ...
Copy to server : contents.xml ...
Removed ESX350-200810401-I-UG Success</I>Installing : ESX350-200810402-T-UG
Copy to server : VMware-tools.tar.gz ...
Copy to server : descriptor.xml ...
Copy to server : install.sh ...
Copy to server : contents.xml.sig ...
Copy to server : contents.xml ...
Removed ESX350-200810402-T-UG Success</I>Installing : ESX350-200810403-C-UG
Copy to server : VMware-viclient.tar.gz ...
Copy to server : descriptor.xml ...
Copy to server : install.sh ...
Copy to server : contents.xml.sig ...
Copy to server : contents.xml ...
Removed ESX350-200810403-C-UG Success</I>
The host needs to be rebooted for the new firmware to take effect.
Type 'yes' to continue:
```

Figure 8.16 ESXi host update completed and prompting to reboot the host

```
C:\Program Files\VMware\VMware UI Remote CLI\bin>vihostupdate.pl -server 172.21.28.22 -username root -q
Enter password:
VMware ESX Server 3i 3.5.0 build-123629

Installed packages:
firmware      123629
tools         123629
viclient      119801
```

Figure 8.17 Vihostupdate utility used to query the ESXi host version

Rolling Back to an Earlier Version with ESXi

Sometimes it's necessary to roll back to an earlier version because of bugs and incompatibilities, and fortunately with ESXi it's fairly easy and straightforward to do this. ESXi saves a copy of the old firmware image when you do upgrades, so you can easily roll back to it if necessary. To roll back to an earlier version, follow these steps:

1. First you need to reboot the ESXi host. During the boot process, when the white bar is displayed across the screen, as shown in Figure 8.18, hold down the Shift key and press the R key.



Figure 8.18 ESXi host boot process

2. A screen warning that the hypervisor image will be replaced with a previous build will display, as shown in Figure 8.19. Press Y to continue with the process.




Figure 8.19 Warning when replacing the current ESXi build with a previous one

3. Press Enter at the next screen to boot with the previous build, as shown in Figure 8.20. When you do this, the older build is what ESXi will continue to use until you update it to a later build again.



Figure 8.20 Selecting a previous build to use with an ESXi host

This quick-and-easy rollback can be invaluable when problems are encountered after upgrading. This feature gives ESXi a big advantage over using ESX. Also, note that when rolling back to an earlier version, all existing host configuration settings are retained.

Patching ESX and ESXi Hosts Using Update Manager

Update Manager is a new plug-in for vCenter Server that was introduced in version 2.5 to add additional patching capabilities to VI3. It can be used to patch both ESX and ESXi hosts and VMs running Windows or Linux operating systems. Update Manager automates patching and works by using a baseline and then installing patches and updates on hosts and VMs that need them. Update Manager can scan and remediate VMs and templates in any power state (on/off/suspended) and any host that is powered on. Snapshots are taken of VMs before upgrading them so that they can be reverted back in case there is a problem with the patches applied. (Be careful not to leave snapshots running any longer than needed or having too many running on a host at once. See Chapter 11, "Advanced Topics," for more information about snapshots.)

Update Manager can be run on Windows workstations and consists of several components, the plug-in component to the VI Client, a Windows service that can run on the vCenter Server, and a database that is used to store patch metadata and other information. This database is different from the one used by vCenter Server, and you can use any of the formats that are also supported by vCenter Server (SQL 2000, SQL 2005 and Express, Oracle 9i and 10g). You can use SQL 2005 Express, which can be installed with Update Manager, but it is recommended to use one of the other databases for larger production environments.

To get started with Update Manager, follow these steps:

1. The Update Manager Windows service is usually installed on the vCenter Server if you chose the default options when installing it. You can verify that it is installed and running by checking the Service list; there should be one called VMware Update Manager Service. If it is already installed, you can skip to Step 6. If it is not already installed, you can install it by either running the `autorun.exe` application from the vCenter Server install CD or download and select to install the Update Manager component or by running `VMware-UpdateManager.exe`, which is usually located in the `updateManager` directory of the vCenter Server CD or download. When installing Update Manager, once you accept the license agreement you will be prompted to choose a directory to install it to and a directory to store the downloaded patches in, as shown in Figure 8.21. You should have at least 18GB of space in whatever drive/directory you choose.




Figure 8.21 Installing Update Manager: Choosing a destination and repository folder

2. You will then be prompted for vCenter Server credentials and port number, as shown in Figure 8.22. These are used to connect to vCenter Server and register the extensions used by Update Manager. Enter a username/password and the IP address/hostname of your vCenter Server and the default port number of 443. You will be prompted for this even if you are installing Update Manager on the vCenter Server itself.
3. Next you will be prompted to choose a database server to be used with Update Manager, as shown in Figure 8.23. If you choose to install and use a dedicated SQL Server 2005 instance, SQL Express will be installed on the server and automatically

configured. If you choose to use an existing database server, you will be prompted to enter an ODBC data source name (DSN) and username/password to connect to the database with. You should make sure that this is already configured on your database server and that you have an ODBC connection configured on the Update Manager server before proceeding.



Figure 8.22 Installing Update Manager: Entering vCenter Server information




Figure 8.23 Installing Update Manager: Choosing a database server

4. Next you will be prompted for port information for Update Manager to use, as shown in Figure 8.24. The default settings are for the Update Manager server to

listen on port 9084 for web and 8084 for SOAP requests and using ports 80 and 443 as a web proxy to the backend app server. These ports should work unless you have an existing web server on the server you are installing Update Manager on. If so, you should change port 80 and 443 to something else so that they do not conflict with the other web server. If you are installing it on the vCenter Server, you cannot change those ports because it uses the existing web server installed as part of vCenter Server.




Figure 8.24 Installing Update Manager: Setting port and proxy information

5. When you enter your port information and click Next, you can begin the installation by clicking the Install button. If you elected to install SQL Express, it will be installed and configured first, and then Update Manager will be installed. Click Finish when the install completes, and you can then check your Windows Service list; the VMware Update Manager Service should be present and running.
6. Load the VI Client on your workstation and connect to your vCenter Server. You will first need to install the plug-in if you do not already have it. Select Plug-Ins from the top menu, and then Manager Plug-Ins. On the Available tab, click the Download and Install button for the Update Manager Extension. When it is installed, click the Installed tab and enable it, and then click OK and restart the VI Client.
7. When you restart the VI Client and connect to the vCenter Server, you will have a new icon called Update Manager on the icon bar, as shown in Figure 8.25. Click the icon to get started. The first thing you will need to do is to create a baseline, which is a set of updates against which the compliance of hosts and VMs is measured. There are a few default baselines that exist, or you can create your own.




Figure 8.25 Installing Update Manager: Update Manage button shown on VI Client icon bar

8. There are four main tabs or pages that are used with Update Manager, as shown in Figure 8.26. The Getting Started tab is optional, and you can close it if you do not want to see it. The Baselines tab is where you create and manage your host and VM baselines. The Configuration tab is where you set options for guest and host settings, how frequently updates are downloaded, and Internet connection and port settings. The Events tab shows you all the events that have taken place that are associated with Update Manager. The Update Repository tab shows you all the host and VM updates that have been downloaded that can be applied.




Figure 8.26 Update Manager tabs and pages displayed

9. Select the Baselines tab and click the New Baseline button, and the New Baseline Wizard will appear, as shown in Figure 8.27. Enter a name and description for your baseline, and then choose whether it is a host or VM update and click Next.
10. Select the type of baseline that you want to use, which is either fixed or dynamic, as shown in Figure 8.28, and click Next. Fixed baselines let you choose specific updates from the repository to apply, and dynamic lets you specify criteria to use to download the latest updates.

The screenshot shows the 'New Baseline Wizard' dialog box with the 'Name and Description' step selected. The left sidebar contains 'Baseline Name', 'Baseline Type', 'Criteria', and 'Ready to Complete'. The main area has two sections: 'Baseline Name and Description' with 'Name' (containing 'New Baseline') and 'Description' (empty text box), and 'Baseline Target' with two radio buttons. The first radio button, 'Virtual Machine / Guest OS Updates', is selected and includes the text 'Applies to virtual machines and includes all vendor supplied guest OS updates and security patches'. The second radio button is 'ESX Server Updates' with the text 'Applies to ESX servers and includes all VMware supplied server updates and security patches'.

Figure 8.27 Creating a baseline: Choosing a name and target

The screenshot shows the 'New Baseline Wizard' dialog box with the 'Type' step selected. The left sidebar contains 'Baseline Name', 'Baseline Type', 'Criteria', and 'Ready to Complete'. The main area has a section 'Type' with the instruction 'Select the type of updates that you would like to use in this baseline.' and a sub-instruction 'Choose from a dynamic baseline that has the latest available updates or a fixed baseline that has specific updates.' Below this are two radio buttons: 'Fixed' (with text 'Select specific updates for this baseline') and 'Dynamic' (with text 'The baseline will be populated with the latest available updates'). The 'Dynamic' radio button is selected. At the bottom are buttons for 'Help', '< Back', 'Next >', and 'Cancel'.

Figure 8.28 Creating a baseline: Selecting a baseline type

11. If you select a fixed baseline, the next screen will allow you to pick updates that are specific to either a host or VM, as shown in Figure 8.29. After you have selected the updates to include in the baseline, click Next.




Figure 8.29 Creating a baseline: Configuring a fixed baseline

12. If you select a dynamic baseline, you will instead see a screen that will allow you to set criteria for the updates to include in the baseline, as shown in Figure 8.30. You can set specific criteria, including the product, severity, vendor, language, and release date. The release date can be set to ensure that updates from older versions are not included. You can also check the Add or Remove Specific Updates field to add exclusions and inclusions from the patch repository. After you have set your criteria, click Next, and then click Finish to save the baseline.
13. Next you need to attach the baseline you created to an object in vCenter Server. This can be a cluster, host, datacenter, or VM. Select the object in the left pane, and then click the Update Manager tab in the right pane, as shown in Figure 8.31.




Figure 8.30 Creating a baseline: Configuring a dynamic baseline




Figure 8.31 Update Manager tab displayed when selecting a host

14. Click the Attach Baseline link, and a window will appear where you can choose one or more baselines to attach to the object that you selected, as shown in Figure 8.32. After you have selected your baselines, click OK, and they will appear listed in the Update Manager tab.
15. Initially, they will show as status unknown, and you will need to scan the object to determine whether it has all the updates included in the baseline. Select the object in the left pane, and right-click it and select the Scan for Updates option, as shown in Figure 8.33.




Figure 8.32 Attaching a baseline to a host server





Figure 8.33 Scanning a host for updates to determine whether it is compliant with its baselines

- When the scan completes, it will show the results and whether the object is compliant with the baseline, as shown in Figure 8.34. It will list the status, how many updates are compliant and how many are not, and how many are unknown.
- If the object is not compliant, you can click the number link in the Not Compliant column to see a list of the updates that are missing, as shown in Figure 8.35.



esx1.xyz.com VMware ESX Server, 3.5.0, 110268

Summary Virtual Machines Performance Configuration Tasks & Events Alarms Permissions Maps Update Manager

Type: **Host System**


Compliant Baselines: **0** Compliant Updates: **32** Last Scan: **1/15/2009 11:14 AM**
 Not Compliant Baselines: **1** Not Compliant Updates: **12**

Baselines:

Name, Attached At or Description contains: ▼

Name	Attached At	Status	Compliant Updates	Not Compliant Updates	Unknown Updates	Description
My Host Updates	esx1.xyz.com	Not Compliant	32	12	0	

Figure 8.34 The results of a host being scanned



Baseline Details - My Host Updates - esx1.xyz.com - Not Compliant Updates

Update Name, Type, Severity, Compliance, Impact, Release Date or Vendor ID contains: ▼ Clear

Update Name	Type	Severity	Compliance	Impact	Release Date	Vendor ID	Patch A
Critical Updates for VMker...	Patch	Security (Host)	Missing	reboot, mmode	9/18/2008 1:00:...	ESX350-20080...	true
Updates the Service Cons...	Patch	Critical (Host or ...	Missing	reboot, mmode	9/18/2008 1:00:...	ESX350-20080...	true
Updates the Software iSC...	Patch	General (Host)	Missing	reboot, mmode	9/18/2008 1:00:...	ESX350-20080...	true
Security Update to Net-S...	Patch	Security (Host)	Missing		9/18/2008 1:00:...	ESX350-20080...	true
Security Update to Perl	Patch	Security (Host)	Missing		9/18/2008 1:00:...	ESX350-20080...	true
Updates Software QLogic ...	Patch	General (Host)	Missing	reboot, mmode	9/18/2008 1:00:...	ESX350-20080...	true
Updates ESX Scripts	Patch	General (Host)	Missing		9/18/2008 1:00:...	ESX350-20080...	true
Security Update to BIND	Patch	Security (Host)	Missing		9/18/2008 1:00:...	ESX350-20080...	true
Updates Software Broadc...	Patch	General (Host)	Missing	reboot, mmode	9/18/2008 1:00:...	ESX350-20080...	true
Updates Performance Tools	Patch	General (Host)	Missing		9/18/2008 1:00:...	ESX350-20080...	true
Updates Invcfg	Patch	General (Host)	Missing		9/18/2008 1:00:...	ESX350-20080...	true
Security Update to Openw...	Patch	Security (Host)	Missing		9/18/2008 1:00:...	ESX350-20080...	true

Close

Figure 8.35 List of missing updates displayed from a noncompliant host

18. To apply the missing patches and updates to make the object compliant, select the object and right-click it and select the Remediate option. When you select this option, a window will appear and you can select the baselines to use to remediate the object, as shown in Figure 8.36. After you select your baselines, click Next.
19. The next screen allows you to select which updates to apply to the object, as shown in Figure 8.37. By default, all updates are selected. After you choose your updates, click Next.

Patching ESX and ESXi Hosts Using Update Manager




Figure 8.36 Remediating a host: Selecting a baseline to use




Figure 8.37 Remediating a host: Selecting the missing updates to install

20. The next screen lets you choose when to remediate the object selected, as shown in Figure 8.38. You can choose to do this immediately or schedule it for a later date/time. You can also set failure options to choose what happens if the server cannot be placed in maintenance mode when the update is attempted. After you have set a start time and failure options, click **Next**, and then click **Finish** to complete the remediation task.





Figure 8.38 Remediating a host: Scheduling a remediation and setting failure options

21. When the remediation begins, you will see task progress in the VI Client, as shown in Figure 8.39. There will be multiple task entries as each update is installed. The host will enter maintenance mode and reboot if needed.

Recent Tasks			
Name	Target	Status	Initiated by
Install	esx1.xyz.com	39%	com.vmware.vcIntegrity
Install	esx1.xyz.com	Completed	com.vmware.vcIntegrity
Enter Maintenance Mode	esx1.xyz.com	Completed	com.vmware.vcIntegrity
Install	esx1.xyz.com	Completed	com.vmware.vcIntegrity
Remediate Entity	esx1.xyz.com	43%	Administrator

Figure 8.39 Remediating a host: Remediation in progress

22. When the remediation task has been completed, the host will exit maintenance mode and will be compliant with the baseline, as shown in Figure 8.40.



The screenshot shows the VMware vCenter Update Manager interface for a host system. The top navigation bar includes tabs for Summary, Virtual Machines, Performance, Configuration, Tasks & Events, Alarms, Permissions, Maps, and Update Manager. The main content area displays the following information:

- Type: Host System
- Compliant Baselines: 1
- Compliant Updates: 44
- Last Scan: 1/15/2009 12:12 PM
- Not Compliant Baselines: 0
- Not Compliant Updates: 0

Below this information is a section titled "Baselines:" with a sub-header "Name, Attached At or Description contains:". A table lists the updates:

Name	Attached At	Status	Compliant Updates	Not Compliant Updates	Unknown Updates	Description
My Host Updates	esx1.wyz.com	Compliant	44	0	0	

Figure 8.40 Remediation of a host complete and showing as compliant

Update Manager greatly simplifies the patching process and addresses a pain point that previously existed for patching your hosts using `esxupdate`. It adds a graphical interface and more features, and as an added benefit can patch your VMs. Update Manager is fairly new to vCenter Server and should get more robust with future releases. For more information about using Update Manager, check out the additional links at the end of this chapter.

Upgrading vCenter Server

There are no individual patches released for vCenter Server, only periodic maintenance releases (for example, 2.5 Update 2) that contain fixes and sometimes new features and functionality. Each maintenance release of vCenter Server is a full installation that installs over the earlier version. One very important thing that you should know when performing upgrades to your virtual environment is that you should always upgrade your vCenter Server first before you upgrade any hosts because newer versions of ESX and ESXi are often not compatible with older releases of vCenter Server. The upgrade process for vCenter Server is almost identical to that of a new installation, so we will not repeat it here. However, you should consider a few things before upgrading your vCenter Server:

- Always read the release notes for the version of vCenter Server that you are installing because they will list any known issues, new features, version incompatibilities, and other information that you may need to know about it.
- Make sure to back up your database before upgrading. During the upgrade, you will be prompted as to whether you want to reinitialize the database and whether you want to clear old historical performance data and tasks and events. If you choose these options, it will wipe out all existing data, and you will have to re-add your hosts and

VMs and re-set up clusters, resource pools, and much more. Having a backup if you inadvertently choose to reinitialize the database is good insurance. If you do need to restore your database, first reinstall the version of vCenter Server that was used with the backed-up database, and then restore it. vCenter Server upgrades often make changes to the design and data in the database, so you cannot use an old database with a new version of vCenter Server.

- If you are running vCenter Server on a VM, take a snapshot of it before upgrading. Otherwise, make sure you have a recent backup of it.
- Make sure you have sufficient disk space on your vCenter Server to install the upgrade.
- It's best to edit your clusters and disable DRS temporarily when performing the upgrade.
- The upgrade usually pushes out new vpx agents to all the hosts that are managed by the vCenter Server and being upgraded. Occasionally, this fails to some hosts due to various reasons, and they will show as disconnected after the vCenter Server is upgraded. If this happens, try logging on to the Service Console of the ESX host that is showing disconnected and type the command `service mgmt-vmware restart`, which restarts the management agent on the host. If that does not work, you can also try using the command `service vmware-vpxa restart`, which restarts the vCenter Server (vpx) agent.

When you upgrade your vCenter Server, you will want to install the newer VI Client on your workstations because often the older client cannot be used with a newer version of vCenter Server. Some upgrades also upgrade the license server on the vCenter Server, so check and make sure it is running properly after the upgrade. Also, make sure you enable HA and DRS in your clusters if you disabled them for the upgrade.

Upgrading Virtual Machines

When you upgrade your hosts, whether it is just patches or full updates, often the VMware Tools application that runs inside a VM's guest operating system will also need to be updated. If you log on to a host or VirtualCenter after you have upgraded a host, you will see the status of VMware Tools listed in the Virtual Machine view in the right pane, as shown in Figure 8.41.

If VMware Tools on a VM does not need to be upgraded, the status will say ToolsOK. If it does need to be upgraded, it will say ToolsOld. A VM will typically run without problem with older versions of VMware Tools, but you should plan to upgrade your VMs as soon as possible after upgrading a host if necessary, because it is always best to run the latest version that the host is running, for security and stability reasons.




Figure 8.41 VMware Tools status shown in the VI Client

To upgrade VMware Tools on your VMs, follow these steps:

1. Select a single or multiple VMs in the VI Client, and then right-click and choose the Install/Upgrade VMware Tools option.
2. A window will appear asking whether you want to perform an Interactive or Automatic upgrade, as shown in Figure 8.42. The Interactive option starts the VMware Tools installation application inside the guest operating system. For this to work, you must have a remote console session open with the VM and be logged on to the operating system. You will get several prompts when the wizard launches, and once the wizard completes and the VMware Tools is upgraded you will be prompted to restart the VM. The Automatic upgrade will upgrade VMware Tools on the selected VMs without any user interaction and automatically reboot the VMs when completed, if necessary.




Figure 8.42 Choosing an option when upgrading VMware Tools on a VM

After you have upgraded VMware Tools and the VM restarts, the status will be updated in the VI Client to ToolsOK. Upgrading VMware Tools on many VMs can be a real pain, but it is often a necessary function after upgrading a host. An alternative method that you can use is to use an endpoint management system or Active Directory group policy to push out the updates using the VMwareTools.msi application. You can access the VMware Tools installation program by editing the settings of a VM and selecting the CD/DVD-ROM and selecting

the Datastore ISO option and browsing to the /vmimages directory and selecting the ISO for your operating system. When you connect the CD/DVD-ROM drive to the VM, you can copy the VMware Tools installation files to a network share, where you can use it to push out to your VMs. Optionally, you can copy the ISO file for the VMware Tools install for your guest operating system from an ESX host, which is located in /vmimages/tools-isoimages/ on the host.

One other alternative for performing mass upgrades of VMware Tools on your VMs is to use a command-line utility that comes with vCenter Server called `vmware-vmupgrade.exe`. This utility is located in the vCenter Server program directory, and the syntax for running it is listed here:

```
vmware-vmupgrade.exe -u user [-p password] [-n vmname] [-h host]
[-m maxpowerons] [-o port] [-t maxpowerontime] [-s] [-q]
```

The various options that can be used with this command are listed here:

- `-u user` Specifies a user with sufficient privileges on the target VM.
- `-p password` Specifies a password on the command line. If this is omitted, the tool immediately prompts for a password.
- `-n vmname` Specifies the name of the VM to upgrade. This name corresponds to the display name of a VM. Specify multiple VMs using multiple `-n` parameters. The `-n` option is ignored if `-h` is specified.
- `-h host` Attempts to upgrade all the VMs on a particular host. Fails if the specified host is not version ESX Server 3.0 or later.
- `-m maxpowerons` On a particular host, powers on only this number of VMs at a time.
- `-o port` Specifies the VirtualCenter Server port, if a port other than the default port 902 has been configured.
- `-t maxpowerontime` Allows a user to set the maximum amount of time for a VM to be powered on in case the guest is unable to power off the machine itself. After the tools upgrade is scheduled on a VM, the VM is powered on and allowed to run through the tools installation process. In most cases, the guest powers down the machine automatically when the process completes.
- `-s` Skips the tools and does only the virtual hardware upgrade.
- `-q` Works quietly. Doesn't produce status or completion messages on stdout.

To specify a host or a VM name for the `vmware-vmupgrade.exe` command, you must specify a path to the host or VM. The path corresponds to the location of the host or VM displayed in the VI Client inventory. To determine host paths, display the Hosts and Clusters view in the Inventory panel. To determine VM paths, display the Virtual Machines and

Templates view in the Inventory panel. Some usage examples of this command are shown here:

- To upgrade a single VM named myvm, in datacenter DC, in the root VM folder
`vmware-vmupgrade -u user -n /DC/myvm`
- To upgrade all (powered-off) VMs on host myhost.vmware.com in the root host folder of datacenter DC, while powering on at most two VMs at a time on the host
`vmware-vmupgrade -u user -h /DC/myhost.vmware.com -m 2`

Summary

Patching is a necessary evil of any software application, and you might be tempted to put it off because it can be time-consuming and disruptive to your hosts and VMs. Patching is one of the key elements of the security of your environment and should not be ignored. You should try to establish a regular patch routine and not wait too long after patches and updates are released before you apply them in your environment. The new Update Manager component greatly simplifies and automates the task of patching your environment. If you do not have the Update Manager component, you may instead explore other automation alternatives using some of the free tools and scripts available from websites such as <http://vmprofessional.com> and www.vmts.net/.

Additional Resources

Here are some additional resources that you can use for further information about maintaining and patching your virtual environment:

- Update Manager Administration Guide
http://vmware.com/pdf/vi3_vum_10u2_admin_guide.pdf
- Best Practices for Patching VMware ESX/ESXi
www.vmware.com/files/pdf/esx_patching_best_practices.pdf
- VMware Update Manager Performance and Best Practices
www.vmware.com/pdf/vum_1.0_performance.pdf
- Patch Management for ESX Server 3.5
http://vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_esxupdate.pdf
- VMware VI3 Upgrade Guide
http://vmware.com/pdf/vi3_35/esx_3/r35u2/vi3_35_25_u2_upgrade_guide.pdf

- Managing Patches and Updates for Hosts and VMs (VMworld 2007 presentation, free registration required)
www.vmworld.com/mylearn?classID=11028
- VMware ESX Server 3 Patch Management (VMworld 2007 presentation, free registration required)
www.vmworld.com/mylearn?classID=11128
- Strategies for ESX Server Update Management (VMworld 2007 presentation, free registration required)
www.vmworld.com/mylearn?classID=11474

Chapter 9

Backing Up Your Virtual Environment

Good backups are a must for any environment, physical or virtual. Traditional backup methods that are used in physical environments can still be used in virtual environments, but there are alternatives available that can be more efficient and faster for backing up your VMs. VMware introduced a product called Consolidated Backup in VI3 that was a new approach to backing up ESX hosts. Consolidated Backup leverages a proxy server to back up virtual disks to offload the burden from the VM and to provide networkless backups by backing up the VM's virtual disks over the SAN fabric. Virtual environments allow for some different approaches to backups because of their snapshot abilities and unique architecture. Taking a snapshot of a VM's virtual disk suspends writes to it and allows for backups to occur without the possibility of files being modified during the backup. In addition, some third-party backup products have been specifically developed for ESX hosts and are worth considering for even better integration, efficiency, and backup speed.

You can use several methods to back up your VMs, and they all have advantages and disadvantages. The simplest and least efficient is to use traditional backup agents running inside the guest operating system. Another method is to use backup scripts that run inside the Service Console that snapshot a VM and make a copy of its virtual disk file to an alternative storage device. You can also use VMware's Consolidated Backup product or one of the third-party backup solutions. Let's take a look at each method so that you can decide which one will work best in your environment.

Traditional Backups

Traditional backups are where a backup agent is installed on the guest operating system, and the backup system communicates with the agent to back up the files on the server. This method can cause high disk and network I/O and high CPU utilizations on the servers that are being backed up. Although this might be fine on a physical server, when done in a virtual environment it can affect other VMs running on the same host because they are all competing for the host's resources and a backup on one VM can leave less resources for the others.

Advantages of this method

- Simple to deploy because most environments are already using this method.
- Easier restores of individual files than other methods.
- No additional software is needed and no process changes are required if you are already using this method.

Disadvantages of this method

- Causes excessive resource usage on hosts.
- Slower than other methods.
- No bare-metal restores, requires OS to be installed first.

If your VMs do not see that much usage during your backup windows and the increased resource utilization is not a problem in your environment, you may consider this method. If you do use this method, be sure to monitor your host's performance during backups to ensure you are not encountering resource bottlenecks that could be affecting other VMs on the host. Also, try to ensure that you only back up a single VM concurrently on a particular host, if possible, to avoid bottlenecks, or stagger your backups so that you are not backing up too many VMs on the same host or LUN at the same time.

Backup Scripts

Backup scripts typically work by taking a snapshot of a VM and then copying a VM's large VMDK file to another storage device where it can be backed up to tape or just left there to use in case a file or the VM needs to be restored. These scripts run inside the ESX Service Console and can use the Perl language or other scripting methods (such as bash) that are supported by ESX. This method can be slow and inefficient, and is useful as a complementary method to another backup method. Because the VMDK file as a whole is being backed up, it makes individual file restores more difficult. To restore individual files, the VMDK file needs to be mounted by another VM or utility, and then the file copied to the original server. This method is best suited for restoring a whole VM image, and is useful for disaster recovery scenarios. You can read detailed information about how snapshots work in Chapter 11, "Advanced Topics."

Advantages of this method

- Simple to set up and use.
- No additional software needed.
- Good for whole VM image restores (bare metal).

Disadvantages of this method

- Individual file restores are difficult.
- Requires access to the ESX Service Console.
- May not work with ESXi.
- Can be slower than other methods.

If you choose to use this method either as your primary backup solution or in conjunction with another method, you might want to check out some of the user-developed scripts available from <http://vmts.net> and www.xtravirt.com.

Consolidated Backup

Chapter 2, “Planning Your Virtual Environment,” provided a detailed summary of the Consolidated Backup application that comes with the VI3 Enterprise license, and rather than repeat everything that was written there, this chapter contains just a brief description of the feature and how it works.

VMware Consolidated Backup (VCB) is a Windows-based application that provides a centralized backup facility to back up VMs through a proxy server without affecting the VM itself. It is an alternative to traditional agent-based backup methods and was designed specifically for virtual environments to minimize host server impact during backup operations. VCB is an enablement technology and cannot back up VMs by itself, but instead works with third-party backup products to help offload backup overhead from the host servers. It integrates with most major software backup providers and eliminates the need to back up VMs over the network when using a Fibre Channel SAN. It works by taking a snapshot of a VM, which suspends the writes to the original disk, and then copies the original disk to a proxy server, where the VM’s disk is mounted on the backup proxy server. It then backs up the contents of the VM as a virtual disk image or as a set of files and directories, and then unmounts the virtual disk and deletes the snapshot. This backup is performed over the SAN fabric and not over the network, which results in greater backup speeds.

Advantages of using VCB

- Eliminates the need for installing a software backup agent on every VM that you want to back up.
- New version now supports backing up VMs on SAN, iSCSI, NFS, and local disk.

- Supported by most major backup software, including Symantec NetBackup and Backup Exec, CA Arcserve, CommVault Simpana, HP Data Protector, and IBM Tivoli Storage Manager.
- Supports both file-level full and incremental backups for VMs running Windows and image-level backups for VMs running any operating system (Windows and Linux).
- Provides OS-level quiescing of file systems, yielding a file system-consistent backup.
- Integrates with VSS to provide application-level consistency for applications that are VSS aware.
- Offloads CPU, disk, and network resource usage from the host to the proxy server.

Disadvantages of using VCB

- More costly, requires a separate physical server and Fibre Channel card to back up VMs on a SAN. (On newer versions of VCB, you can optionally run VCB inside a VM to avoid this.)
- Limited OS support (Windows and Linux only).
- May require excessive additional space on the SAN to back up VMs.
- Individual file restores require more steps and are more difficult.
- Can be complicated to set up, configure, and maintain.

Consolidated Backup is a great a solution to remove the backup burden from your hosts, and is more suited for larger environments. The first release of Consolidated Backup was rather limited, but VCB has evolved with each new release, with more support and better features, and you should definitely consider using it in your environment if you are using one of the backup applications supported by it. VMware has released several guides that you can use to help implement VCB. These are listed in the “Additional Resources” section at the end of this chapter.

Third-Party VI3-Specific Backup Products

Several third-party backup products have been developed specifically for backing up VI3 environments and are a great alternative to traditional backup methods and provide more options and greater flexibility when backing up and restoring your VMs. These products all work by copying VM virtual disks over the network from a source datastore to a destination disk-based storage device (local, SAN, NFS, iSCSI, and CIFS). In addition, most of these products can integrate with VCB and can do full and differential backups and both whole VM (bare-metal) and individual file restores.

These products all vary in price, performance, and features, and you should evaluate each one to see which one integrates best into your environment and satisfies your requirements.

You should also make sure that any product you look at supports everything in your environment, including VCB, ESXi, guest operating systems, and your storage devices. Make sure to request an evaluation copy of each product and install and use it to make sure it will work properly for you before selecting a backup product.

Vizioncore vRanger Pro

vRanger Pro provides image-level hot backups while VMs are still running. It has an easy to use graphical interface, startup wizards, and VirtualCenter integration. Advanced functionality, such as VCB integration, VSS quiescing, and differential backups, is enabled through its GUI, with no manual RPM installs or complicated scripting. A differential engine backs up only the changes made since the last full backup image, reducing the size of the backup files on disk. Differential characteristics and retention policies can easily be configured to suit customer needs. vRanger Pro supports differential backups via the VCB framework. It also supports backing up and restoring both ESX and ESXi host servers. You can find more information about vRanger Pro at <http://vizioncore.com/vRangerPro-features.html>.

Veeam Backup

Veeam Backup provides enterprise-ready functionality, including an intuitive user interface with simple backup and replication job management; restoration of a full VM or a single file with just a few mouse clicks; comprehensive statistics, reporting, and email alerting; and integrated ESX file management through Veeam FastSCP. Veeam Backup can work with the VCB framework for backup and replication purposes. Using its proprietary “VCB on-the-fly” technology, Veeam Backup does not require storing VM images on the VCB proxy during backup, allowing for faster backup directly to the target without additional disk space requirements. Veeam Backup also supports incremental backup and replication through VCB, which is not supported natively by VCB. Veeam Backup features limited support for ESXi servers through VCB. You can now back up ESXi servers using the VCB option in the Backup Wizard. File-level recovery is fully supported for guests running on ESXi, whereas full image restore is only supported to ESX 3.x servers. After the guest has been restored to an ESX 3.x server, you can then VMotion the restored machine back to the ESXi server. Replication is also supported with ESXi as a source and ESX 3.x as a target. You can find more information about Veeam Backup at <http://veeam.com/vmware-esx-backup.html>.

esXpress

esXpress by PHD Technologies makes use of virtual backup appliances (VBAs) to quickly and securely back up VMs without affecting console resources or VM availability. VBAs are just tiny VMs that back up VMs. Unlike dedicated hardware solutions, VBAs are cost-effective

and scale with your environment. And because of their distributed nature, VBAs are completely fault tolerant, very low impact, and yet yield incredibly fast speeds. esXpress allows for multiple transport protocols and targets, including FTP, SSH, SMB/CIFS, iSCSI, SAN, local storage, and more. Additional esXpress features include support for both file- and image-level backups, compression for increased backup speeds, and integration with the VI Client. You can find more information about esXpress at <http://esxpress.com/index.php>.

Did You Know?

Many backup applications use a process called quiescing before backing up servers. This process ensures that the data on the disk is in a state suitable for backups and includes operations such as flushing dirty buffers from the OS memory cache to disk. Different types of quiescing can be done. When it is done at the operating system level, it may not be aware of applications that are running and consequently might not work properly. A better form of quiescing is done at the application level (for example, Exchange, SQL) and is aware of the application and so it can properly write data before backing up the server.

What to Back Up

When it comes to backing up VMs, you have two options: back up the large VMDK virtual disk file (known as image-level backups) at the virtualization datastore level or back up the individual files (known as file-level backups) at the guest operating system level. Both options have advantages and can be used jointly to provide different restoration alternatives for your VMs.

File-level backups allow for individual file restores, which are useful when you have just a few files to restore to a VM and do not need to restore all the data just to restore the individual files. File-level backups can be done using traditional backup methods by running a backup agent inside the guest operating system of a VM. In addition, many backup applications that have support for virtualization can perform file-level backups through nontraditional methods such as using VMware's Consolidated Backup application. File-level backups are great for restoring a small number of files, but restoring a whole VM can be more time-consuming because you typically need to install a guest operating system and backup agent before you can restore the rest of the VM data.

Image-level backups allow for whole VM or bare-metal restores and are useful when you need to quickly restore a VM to a previous state or for disaster recovery purposes, which require bare-metal restores rather than individual file restores. To perform an image-level backup, you need to use a product or method that supports it rather than traditional backup agents, which work inside the guest operating system and do not have access to the virtualization layer and the VMDK virtual disk file. You can still do individual file restores with

image-level backups, it just typically takes a few more steps to do it. To do this, you can restore the disk image and mount it on another VM to access the file and then copy it back to the source VM. In addition, VCB can make use of vCenter Converter to restore individual files.

So, which method should you use? It really depends on the backup method you plan to use. If you are going to use traditional backup methods that will enable you to restore individual files, you might also look to use one of the scripts or third-party products that can copy the image to disk storage so that you have that available as another restore option. This is a good way to make use of local disk on your hosts that you might otherwise not use or a low-cost NFS or iSCSI system or server. You might back up your VMs every day using the file-level method and back them up once a week using the image-level method. If you instead plan to use VCB or a third-party backup product, many of them will already provide you with the ability to do both file- and image-level backups. Image-level backups are great when you want to retire a physical or virtual server so that you have a bare-metal copy of it if you ever need it later.

Another question that is often asked is whether you should install a backup agent inside the ESX Service Console to back it up and the VM files from the VMFS datastores. Generally this is not a good idea for several reasons. The first is that it can have a big impact on the performance of the host when backups are running, which will also affect the VMs on that host. Second, you should avoid installing any software or utilities on the Service Console (with the exception of hardware management agents) because it could cause instability, which could lead to your host crashing. The last reason is that it is easier to rebuild an ESX host if there is a problem with it rather than try to restore it. When you rebuild an ESX host, you have the option to leave all the VMFS datastores intact so that your VMs can be re-registered with the host after it has been rebuilt. You will lose some configuration information (for example, vSwitches), but you can back up certain key files from your host to restore this. If your configuration is fairly simple, you can reconfigure your host instead. It is also a good idea to periodically run the vm-support script on your ESX hosts, which will provide detailed information about the host configuration and much more that you can later use as a guide to reconfigure it.

To back up the key ESX Service Console configuration files, just back up the `/etc/vmware` directory. You can do this with the `tar` command by typing `tar -cvf esxhostnamebackup.date.tar /etc/vmware`. This will create a single tar file with the contents of that directory that you can copy to another location for safekeeping. You can do this periodically for each of your hosts to use if you ever need to restore them. VMware has a knowledge base article that covers how to back up and restore ESX host configurations (KB article 1000761, <http://kb.vmware.com/kb/1000761>). To back up the configuration of ESXi hosts, you can use the `vicfg-cfgbackup.pl` script that is part of the RCLI.

Attaching Tape Drives to ESX Hosts

Although it is best to avoid doing this, sometimes it is necessary, especially in smaller environments that do not have a dedicated tape backup server. By installing a tape drive in your ESX host, you can use a VM to access it to back up the VMs on your host or other hosts. If you want to do this, you should follow a few guidelines to ensure success:

- Your tape device must use an Adaptec SCSI I/O adapter. (A dedicated one is recommended.) You cannot use an IDE or FC tape device, because even though the host may see it the VM will be unable to because it supports only virtual SCSI controllers.
- Make sure that the Adaptec SCSI controller that you are using is listed on the ESX hardware compatibility list for I/O devices.
- To check whether ESX can see the tape drive, use the following command in the Service Console: `cat /proc/scsi/scsi`. This command will list all attached SCSI devices. When you can see the tape device in the ESX Service Console, install a backup agent that supports ESX. If one is not available, use a backup agent that supports Red Hat Linux Enterprise version 3.
- When selecting a SCSI controller for your VM, make sure to choose LSI Logic rather than BusLogic.

For more information about using a tape device in your ESX host, see VMware knowledge base article 1000024 (<http://kb.vmware.com/kb/1000024>).

Watch Out!

One risk of attaching a tape drive to an ESX host is that if the tape drive gets hung up, which is not uncommon, you will need to reboot the ESX host to clear it so that you can use it again.

Summary

Backups are important, so make sure you choose a solution that will work for the needs of your environment. Backups are like having good insurance policies: You pay your premium each month, and you hope you never need to use them; but if something happens, you are real glad that you have them. Restoring your data is even more important than backing it up. You need to make sure the solution that you use can easily restore files as needed in your environment. You should also plan your backup strategy around your current or future disaster recovery strategy because good backups are the foundation for any disaster recovery plan. If you plan on doing disaster recovery, look for a backup strategy that supports your disaster

recovery requirements for recovering from an event. In addition, you may look into some of the products that can replicate your VMs from your production environment to a disaster recovery environment. When you implement a backup solution in your virtual environment, make sure you test the restore capability so that you don't run into any surprises later when trying to restore critical data.

Additional Resources

For more information about backing up and restoring your environment, check out these resources:

- Virtual Machine Backup Guide
http://vmware.com/pdf/vi3_35/esx_3/r35u2/vi3_35_25_u2_vm_backup.pdf
- Consolidated Backup in VMware Infrastructure 3
www.vmware.com/pdf/vi3_consolidated_backup.pdf
- Using VMware Infrastructure for Backup and Restore
www.vmware.com/pdf/esx3_backup_wp.pdf
- VMware Consolidated Backup
www.vmware.com/files/pdf/vcb_best_practices.pdf
- VMware Consolidated Backup: Improvements in Version 3.5
www.vmware.com/files/pdf/vcb_35_new.pdf
- VMware Consolidated Backup—Partner Integration Guide
www.vmware.com/files/pdf/vcb_partner_integration_guide.pdf
- Best Practices for Architecting VCB Enabled Solutions (VMworld 2007 presentation, free registration required)
<http://vmworld.com/vmworld/mylearn?classID=11154>

This page intentionally left blank

Chapter 10

Troubleshooting Your Virtual Environment

When problems occur in your virtual environment, you need to know where to look for clues to the cause and what to do to resolve them. Often, just trying to figure out the exact cause is the most difficult part, because virtual servers are more complicated than physical servers and there are more potential causes of problems. Troubleshooting problems can often be frustrating, and you should take your time and not make any rash decisions that might worsen the problem. When you know where to look to find the cause of a problem, the process becomes a lot easier.

Troubleshooting ESX and ESXi Hosts

Some differences exist between ESX and ESXi hosts, so we cover them separately. ESX hosts are more complex than ESXi hosts due to the Service Console that they have, and therefore can be a bit more challenging to troubleshoot. There are also many more log files that can be viewed on an ESX host compared to an ESXi host, and therefore more things to check when you are troubleshooting an ESX host.

Troubleshooting ESX Hosts

There are many common problems that happen with ESX hosts. We cover some of them and resolutions and what log files to check, how to check the version of ESX components, and what resources you can use to help you troubleshoot your problem.

Log Files

Let's begin with the log files that you can use to troubleshoot your ESX host. There are numerous log files that you can look at depending on the nature of your problem. These log files are located on the ESX Service Console and can be viewed by opening the logs with a text editor, via SCP with a tool such as WinSCP, using the VI Client, or using the `tail` Linux command. Here is a summary on how to use the various methods for viewing log files:

- **Text editor.** Log on to the Service Console and use the built-in text editors to open a log file. You can use `nano`, which is a bit easier to use, or `vi` to open a log file. Type `nano` or `vi log file path/name` to open one (for example, `nano /var/log/vmware/hostd.log`).
- **SCP tool.** If you install an SCP tool such as WinSCP or FastSCP on your workstation, you can connect to an ESX host and browse through the file system using a GUI similar to Windows Explorer. Then just select the file you want to view and open it, which downloads it to your workstation for you to view.
- **tail command.** The Linux `tail` command enables you to view the last part of a text file. By default, without any options, it will show you the last ten lines of a file, but you can specify more lines or have it follow the file so that it continually reads the file and displays new additions to the file. To use this command, log on to the Service Console and type `tail log file path/name`. You can use the `-f` option to follow the file or the `-n #` command to specify the number of lines to display (for example, `tail /var/log/vmware/hostd.log -f` or `tail /var/log/vmware/hostd.log 50`).
- **VI Client.** If you connect to an ESX host with the VI Client, you can view the logs by clicking the Administration button at the top and then selecting the System Logs tab, as shown in Figure 10.1. Here you can select between the various log files, search for log entries, and control how many of the log entries display.

The following log files are available on an ESX host:

- **VMkernel, /var/log/vmkernel.** Records activities related to the VMs and ESX hosts. Rotated with a numeric extension. The current log has no extension, and the next newest one has a .1 extension
- **VMkernel Warnings, /var/log/vmkwarning.** Records activities with the VMs. It is a subset of the VMkernel log and uses the same rotation scheme.
- **VMkernel Summary, /var/log/vmksummary.** Used to determine uptime and availability statistics for ESX hosts; a human-readable summary can be found in `/var/log/vmksummary.txt`.
- **ESX Server host agent, /var/log/vmware/hostd.log.** Contains information about the agent that manages and configures the ESX host and its VMs. (Search the file date/time stamps to find the log file it is currently outputting to or open `hostd.log`, which is linked to the current log file.)




Figure 10.1 Using the VI Client to view ESX host log files

- **ESX Firewall, /var/log/vmware/esxcfg-firewall.log.** Logs all firewall rule events. Rotated with a numeric extension. The current log has no extension, and the next newest one has a .1 extension
- **ESX Boot, /var/log/vmware/esxcfg-boot.log.** Logs all ESX boot events. Rotated with a numeric extension. The current log has no extension, and the next newest one has a .1 extension.
- **ESX Update, /var/log/vmware/esxupdate.log.** Logs all updates done through the esx-update tool.
- **Service Console, /var/log/messages.** Contains all general log messages used to troubleshoot VMs or an ESX host.
- **Web Access, /var/log/vmware/webAccess.** Records information about web-based access to an ESX host.
- **High Availability (HA), /var/log/vmware/aam.** Logs information related to the High Availability service. In ESX 3.0.x, these logs were in /opt/LGTOaam512/ instead.
- **Authentication, /var/log/secure.** Contains records of connections that require authentication, such as VMware daemons and actions initiated by the xinetd daemon.
- **Vpxa, /var/log/vmware/vpxa.log.** Contains information about the agent that communicates with vCenter Server. Search the file date/time stamps to find the log file it is currently outputting to or open vpxa.log file, which is linked to the current log file.

Which log you use will depend on the problem you are troubleshooting. In general, the VMkernel and hosted log files are the ones you will use for most problems. For other problems that lie in specific areas, you may check logs related to those areas (for example, checking the `esxupdate.log` for problems with Update Manager or the `vpix.log` for problems with connections to vCenter Server). Viewing some of these log files, you might be overwhelmed by the number of entries and the somewhat cryptic entries in the log files. You might not understand everything in the log file, but you can look for obvious error messages that may relate to your problem. Log files contain a lot of information, and it can sometimes be difficult to sort out what is normal and what is not. If you open a support case with VMware, they will request the log files, and will be able to analyze them and tell you what they find.

In addition to the methods to view individual log files, you can use the `vm-support` command (it's actually a script) that you can run on the ESX Service Console that will bundle together all the log files, configuration files, and output from various commands into a single TGZ file. After the file has been created, you can copy it to your workstation and extract it using the Linux `tar` command or WinZip. There can be hundreds or thousands of little files that are extracted. These files are very useful to VMware support to help troubleshoot your problem.

You can run the `vm-support` command without any parameters, which will collect all the files and bundle them together, or you can specify parameters to collect specific information or do certain tasks. Here are some of the most commonly used parameters that can be used with `vm-support`:

- `-n` Causes no core files to be included in the tar file
- `-s` Takes performance snapshots in addition to other data
- `-S` Takes only performance snapshots
- `-x` Lists world IDs (`wid`) for running VMs
- `-X wid` Grabs debug info for a hung VM
- `-w dir` Sets working directory for output files
- `-h` Displays help for command usage and available options

To run `vm-support`, log on to the Service Console and type the command either by itself or with options. The script will run and perform all the data collection and create the tar file, as shown in Figure 10.2.

Just remember to delete these files from your ESX host when you have copied them off to avoid filling up your disk partitions.

The VI Client also has an option to export diagnostic data (`vm-support` output and more) to files on your local workstation, which you can open and review the various files inside or send to VMware support. If you connect to a vCenter Server, you can select multiple hosts simultaneously and include VI Client and vCenter Server logs and files in the output. When you use the VI Client and select hosts, it calls the `vm-support` script on the host and then copies the output file to your local PC. To do this, connect to a host or vCenter Server with

the VI Client, and then select File from the top menu, and then Export, and then the Export Diagnostic Data option. A window will appear, as shown in Figure 10.3, in which you can select what you want to include and a directory on your workstation to download the data to.

```
[root@esx2 root]# vm-support
VMware ESX Server Support Script 1.29

Preparing files: \
Waiting up to 300 seconds for background commands to complete:

Waiting for background commands: \
Creating tar archive ...

File: /root/esx-2009-01-22--08.20.24895.tgz
Please attach this file when submitting an incident report.
To file a support incident, go to http://www.vmware.com/support/sr/sr\_login.jsp

To see the files collected, run: tar -tzf /root/esx-2009-01-22--08.20.24895.tgz
Done
[root@esx2 root]#
```

Figure 10.2 Running vm-support on an ESX host




Figure 10.3 Exporting diagnostic data from hosts

When it begins, a task will appear in the Recent Tasks pane of the VI Client, and a new window will open showing the download progress of each bundle, as shown in Figure 10.4.

When it completes, it will create a new subdirectory in the directory that you specified with a date/time stamp and containing the TGZ file from the hosts, a zip file from the vCenter Server, and a VI Client subdirectory containing log files created by the local VI Client.




Figure 10.4 Log bundles are written to a workstation.

Determining Versions

Many times, you need to know the versions of the various components of an ESX host and what patches are applied because VMware support will need this information to help determine whether certain known issues affect your host. Listed here are some Service Console commands that you can use to determine this:

- `vmware -v` will display the ESX host version and build number.
- `/opt/vmware/vpxa/sbin/vpxa -v` will display the vCenter Server agent version and build number.
- `rpm -qa | grep VMware-esx-tools` will display the version and build number of the VMware Tools installation bundle on the ESX host.
- `esxupdate -l query` will display all patches installed on the ESX host.

Sample output from some of these commands is shown in Figure 10.5.

```
[root@esx2 root]# vmware -v
VMware ESX Server 3.5.0 build-123630
[root@esx2 root]#
[root@esx2 root]# /opt/vmware/vpxa/sbin/vpxa -v
VMware VirtualCenter Agent Agent Daemon 2.5.0 build-119598
[root@esx2 root]#
[root@esx2 root]# rpm -qa | grep VMware-esx-tools
VMware-esx-tools-3.5.0-123630
```

Figure 10.5 Determining versions of an ESX host components

You can also obtain the ESX host and build number by selecting the host in the VI Client. Version numbers usually change only when major product upgrades are installed (for example, ESX 3.0, 3.5) and build numbers change frequently when any type of upgrade is applied, whether it is a patch, update, or major upgrade.

Common Problems and Resolutions

There are some problems that occur with ESX hosts that are fairly common, and we cover those here. Many times, you can take simple steps to correct them, but some require more in-depth troubleshooting and resolutions.

Purple Screen of Death

One occurrence that can happen on both ESX and ESXi hosts is called the purple screen of death (PSoD), which is VMware's version of Microsoft's infamous blue screen of death (BSoD) and the result of a host crashing and becoming inoperative. A PSoD, as shown in Figure 10.6, is definitely not something you want to see on your hosts.

```
VMware ESX Server [ReleaseBuild-32039]
Exception type 13 in world 1193:vmware-vmx @ 0x6c4f5a
gate=0xd frame=0x36a7e58 eip=0x6c4f5a cr2=0x524ad000 cr3=0x96b3a000 cr4=0x168
eax=0x0 ebx=0xe ecx=0x68000000 edx=0xba0ced es=0xba04028 ds=0x4028
fs=0x0 gs=0x0 ebp=0x36a7e5c esi=0x36a7ed8 edi=0xba00b48 err=0 ef=0x10246
cpu 0 1024 console: cpu 1 1140 mks:OMHQ5: cpu 2 1138 vmm0:OMHQ: cpu 3 1083 vmm0:
OMHQ:
cpu 4 1267 mks:OMHQ5: cpu 5 1107 vmm0:OMHQ: cpu 6 1069 vmm0:OMHQ: cpu 7 1078 mks:
OMHQ5:
cpu 8 1032 idle8: cpu 9 1158 vmm0:OMHQ: cpu 10 1102 vmm0:OMHQ: cpu 11 1165 vmm0:
OMHQ:
cpu 12 1095 vmm0:OMHQ: CPU 13 1193 vmware-vm: cpu 14 1232 mks:OMHQ5: cpu 15 1039
idle15:
0x36a7e5c:[0x6c4f5a]UserObjFDPoll+0x5e(0xba00b48, 0x7, 0x0)
0x36a7ee8:[0x6c4fc7]UserObjPollCleanupWaiters+0x4b(0xba15928, 0xe, 0x0)
0x36a7f40:[0x6c52b8]UserObjPoll+0x1ec(0xba15928, 0xe, 0x12)
0x36a7f74:[0x6daf7d]LinuxFileDescPoll+0xad(0xbf5feb84, 0xe, 0x12)
0x36a7fa8:[0x6bf314]UserLinuxSyscallHandler+0x6c(0x36a7fe0, 0x10000023, 0xd2002
3)
0xbf5fab38:[0x667cfc]CommonTrap+0xc(0x0, 0x0, 0x0)
VMK uptime: 9:10:36:17.943 TSC: 1630838606661150
TSC: 4742056872 cpu0:0)Chipset: 665: Make sure that if 'noapic' is used, it is o
n purpose
0:00:00:31.779 cpu0:1024)PCI: 1650: failed for 000:15.2
9:07:02:41.693 cpu0:1158)APIC: 1265: Lint1 interrupt on pcpu 0 (port x61 contain
s 0xb0)
Starting core dump to disk using slot 1 of 1... 98766666543210 Disk dump successf
ul.
Debugger is listening on serial port ...
Remote debugger activated, Local debugger no longer available
```

Figure 10.6 ESX host purple screen of death

When one occurs, it will completely halt the host and cause it to become unresponsive. The causes of PSoDs are typically hardware related (defective memory is the most common cause) or a bug in the ESX application. Your only recourse when it happens is to power off the host and power it back on. The information that displays on the screen is useful, however, and you should attempt to capture it by writing it down, using a camera phone to take a picture of it, or taking a screen capture from a remote management board if present. You might not be able to make much sense out of the information, but it will be very useful to VMware support. What is displayed consists of the ESX version and build number, the exception type, register dump, what was running on each CPU when the crash occurred, back-trace, server uptime, error messages, and memory core information.

When you reboot a host after experiencing a PSoD, a file beginning with the name `vmkernel-zdump` should be present in your host's `/root` directory. This file will be useful to VMware support, and you can also use it to help determine the cause by using the `vmkdump` utility to extract the VMkernel log file and look for any clues as to the cause of the PSoD. To use this command, type `vmkdump -l dump filename`. As previously mentioned, defective memory is a common cause of a PSoD. You can use the dump file to help identify the memory module that caused the problem so that it can be replaced.

If you suspect defective memory is the cause, you can test your host's memory by using a utility application that burns in memory. These utilities require you to shut down your host and boot from a CD to run the memory tests. One commonly used utility is Memtest86+, which does extensive memory testing, including checking the interaction of adjacent memory cells to ensure that writing to one cell does not overwrite an adjacent cell. You can download this utility at www.memtest.org/.

Did You Know?

ESX 3.0.x servers had a built-in memory test utility called Ramcheck that was removed in ESX 3.5. This nondisruptive utility could be run by typing the command `service ramcheck start` and would test unused memory in the background with little resource overhead. Ramcheck wrote its result output to the `/var/log/vmware/` directory in files named `ramcheck.log` and `ramcheck-err.log` and would run until the host was restarted or it was stopped with the `service ramcheck stop` command.

It is a good idea to burn in and test your host's memory when you are first building it to avoid disruptions later. Most memory problems are not obvious and will not be detected by the simple memory test a server does as part of its POST boot procedure. You can download the free Memtest86+ utility as a small 2MB ISO file, which you can burn to a CD to boot from and let it run for at least 24 hours to run various memory tests on your host. The more RAM you have in the system, the longer it will take to complete one pass. A server with 32GB of RAM will generally take about one day to complete. Besides the system memory, Memtest86+ will also test your CPU's L1 and L2 cache memory. Memtest86+ will run indefinitely, and the pass counter will increment as all the tests are run.

Service Console Problems

You might sometimes experience a problem with your Service Console where it hangs and will not allow you to log on locally. The condition, which can be caused by hardware lockups or a deadlocked condition, will usually not affect the operation of the VMs running on the host. Rebooting is often the only recovery for this condition. Before doing this, however, you should shut down or VMotion the VMs to other hosts. You do this by whatever method works, such as using the VI Client, connecting to the Service Console remotely via SSH, or trying to use an alternative/emergency console, which can be accessed by pressing Alt-F2

through Alt-F6. When you move the VMs or shut them down, you can reboot the host with the `reboot` command. If all the console methods are unresponsive, you will need to cold boot the host instead.

Networking Problems

You may also experience a condition that causes you to lose all or part of your networking configuration or where a configuration change causes the Service Console to lose network connectivity. If this happens, you will not be able to connect to the host by any remote method, including the VI Client or SSH. Your only recourse will be to rebuild or fix the network configuration from the local Service Console using the `esxcfg-` command-line utilities. Here are some of the commands that you can use to configure networking from the ESX CLI:

- **esxcfg-nics** This command displays a list of physical network adapters along with information about the driver, PCI device, and link state of each NIC. You can also use this command to control a physical network adapter's speed and duplexing. Type `esxcfg-nics -l` to display NIC information and `esxcfg-nics -h` to display available options for this command. Here are some examples:
 - ◆ Set the speed and duplex of a NIC (vmnic2) to 100/Full:
`esxcfg-nics -s 100 -d full vmnic2`
 - ◆ Set the speed and duplex of a NIC (vmnic2) to autonegotiate:
`esxcfg-nics -a vmnic2`
- **esxcfg-vswif** Creates and updates Service Console network settings, including IP address and port group. Type `esxcfg-vswif -l` to display current settings and `esxcfg-vswif -h` to display all available options for changing settings. Here are some examples:
 - ◆ Change your Service Console (vswif0) IP and subnet mask:
`esxcfg-vswif -i 172.20.20.5 -n 255.255.255.0 vswif0`
 - ◆ Add a Service Console (vswif0):
`esxcfg-vswif -a vswif0 -p "Service Console"`
`➡ -i 172.20.20.40 -n 255.255.255.0`
- **esxcfg-vswitch** Creates and updates VM (vSwitch) network settings, including uplink NICs, port groups, and VLAN IDs. Type `esxcfg-vswitch -l` to display current vSwitch configurations and `esxcfg-vswitch -h` to display all available options for changing settings. Here are some examples:
 - ◆ Add a physical NIC (vmnic2) to a vSwitch (vSwitch1):
`esxcfg-vswitch -L vmnic2 vswitch1`
 - ◆ Remove a pNIC (vmnic3) from a vSwitch (vSwitch0):
`esxcfg-vswitch -U vmnic3 vswitch0`

- ◆ Create a port group (VM Network3) on a vSwitch (vSwitch1):
`esxcfg-vswitch -A "VM Network 3" vSwitch1`
- ◆ Assign a VLAN ID (3) to a port group (VM Network 3) on a vSwitch (vSwitch1):
`esxcfg-vswitch -v 3 -p "VM Network 3" vSwitch1`
- `esxcfg-route` Sets or retrieves the default VMkernel gateway route. Type `esxcfg-route -l` to display current routes and `esxcfg-route -h` to display all available options for changing settings. Here are some examples:
 - ◆ Set the VMkernel default gateway route:
`esxcfg-route 172.20.20.1`
 - ◆ Add a route to the VMkernel:
`esxcfg-route -a default 255.255.255.0 172.20.20.1`
- `esxcfg-vmknics` Creates and updates VMkernel TCP/IP settings for VMotion, NAS, and iSCSI. Type `esxcfg-vmknics -l` to display VMkernel NICs and `esxcfg-vmknics -h` to display all available options for changing settings. Here is an example:
 - ◆ Add a VMkernel NIC and set the IP and subnet mask:
`esxcfg-vmknics -a "VM Kernel" -i 172.20.20.19 -n 255.255.255.0`

In addition, you can restart your Service Console network by using the command `service network restart`.

Other Problems

Sometimes just restarting some of the ESX services will resolve problems and not affect the VMs running on the host. Two services that can be restarted and often fix many problems are the `hostd` service and the `vpixd` service. The `hostd` service runs in the Service Console and is responsible for managing most of the operations on the ESX host. To restart the `hostd` service, log on to the Service Console and type `service mgmt-vmware restart`.

Watch Out!

Be aware that a bug has existed in many versions of ESX 3.x that causes VMs to reboot when restarting the `hostd` service. (See VMware knowledge base article 1003312 for more info.) This only happens if your host is configured to use automatic startups. This bug was fixed in a patch for 3.0.1 and in 3.0.2, but appeared again in ESX 3.5, and another patch was released to fix it. It's best to temporarily disable auto-startups before you run this command.

The `vpixd` service is the management agent that handles communication between the ESX host and its clients, including the vCenter Server and anyone who connects to the host using the VI Client. If you experience problems with vCenter Server showing a host as

disconnected, not showing current information, or any other strange problems involving vCenter Server and a host, restarting this service may resolve it. To restart the vpxa service, log on to the Service Console and type `service vmware-vpxa restart`. It is recommended to try to restart these two services when you are experiencing problems because this will often resolve many problems.

Troubleshooting ESXi Hosts

Troubleshooting ESXi hosts is a bit different because there is no Service Console. However, ESXi still has a console that you can use for troubleshooting purposes. In addition, ESXi has many of the same log files as an ESX host, but the methods to access them are a bit different.

Log Files

ESXi uses some of the same core log files as ESX, with a few differences. There are four main log files on an ESXi host: the `hostd` log, messages that combine the `VMkernel` and `vmkwarnings` log files (the `vmksummary` log file does not exist on ESXi hosts), the `config` log file, and the `vpxa` log file if the ESXi host is managed by a vCenter Server. You can access the log files in various ways, as follows:

- **Local console.** If you log on to the local ESXi console and select the View System Logs option, you will be able to choose from the log files to view, as shown in Figure 10.7. Option 4, which is the `vpxa` log, will display only if the host is managed by a vCenter Server.





Figure 10.7 Viewing system logs of an ESXi host

- **VI Client.** Similar to that of an ESX host, if you connect to an ESXi host using the VI Client and select the Administration tab and then the System Logs tab, you can select between the various log files, search for log entries, and control how many of the log entries display.

Determining Versions

You can determine the version of the various components of an ESXi host using the following methods:

- **RCLI.** Type the command `vihostupdate --server hostname or IP address --username root -q` to display the ESXi version and build number and the versions of the installed packages, including the VI Client and VMware Tools, as shown in Figure 10.8.



```
C:\Program Files\VMware\VMware VI Remote CLI\bin>vihostupdate.pl --server esxi1.x
yz.com --username root -q
Enter password:
VMware ESX Server 3i 3.5.0 build-110271

Installed packages:
viclient      119801
tools        130755
firmware      110271
```

Figure 10.8 Determining versions of an ESXi host components

- **ESXi console.** The local ESXi console will display the version and build number on the main screen and at the bottom of every screen.

Just such as with ESX, if you select an ESXi host in the VI Client, it will also show you the version and build number of the ESXi host.

Accessing Tech Support Mode

Tech Support Mode is the name for a hidden console (Busybox shell) that can be accessed to troubleshoot and diagnose an ESXi host. This mode is enabled by default, but can also be disabled for security purposes using the Advanced Settings configuration. It is generally not recommended to use this console unless you know what you are doing and are advised by VMware support to use it to resolve a problem. However, if you want to access it, be aware of the following caveats:

- The console is not audited or logged, so any commands issued inside of it are not recorded.
- Some of the commands if used incorrectly can result in an unusable system.
- You can only log on to the console using the root account.

To access this mode, follow these steps:

1. At the local ESXi console, as shown in Figure 10.9, press Alt-F1. It does not matter whether you are at the main screen or in the Customize System screen.




Figure 10.9 ESX console screen

2. This will display a screen that shows various system messages, as shown in Figure 10.10. You will not receive a prompt, and anything you type on this screen will not be displayed. Type the word unsupported to enter the Tech Support Mode.




Figure 10.10 Hidden console screen displayed after you press Alt-F1

3. A warning message will display, and you will be prompted for the root password, as shown in Figure 10.11.

```
You have activated Tech Support Mode.  
The time and date of this activation have been sent to the system logs.  
  
WARNING - Tech Support Mode is not supported unless used in  
consultation with VMware Tech Support. Tech Support Mode may be  
disabled by an administrative user. Disabling requires a reboot of  
the system. Please consult the ESX Server 3i Configuration Guide  
for important additional information.  
  
Password: _
```

Figure 10.11 Logging on to Tech Support Mode

4. When you enter the root password, you will be at a prompt where you can enter commands, as shown in Figure 10.12.

```
Tech Support Mode successfully accessed.  
The time and date of this access have been sent to the system logs.  
  
WARNING - Tech Support Mode is not supported unless used in  
consultation with VMware Tech Support.  
  
_ # _
```

Figure 10.12 Prompt shown after you successfully log in to Tech Support Mode

5. A limited number of Linux commands work in the console. Commands such as `cd`, `ls`, `vi`, `reboot`, and `cp` will work, but others will not. In addition, you can use many of the `esxcfg` commands to view or change configurations, as shown in Figure 10.13, and use the `esxtop` application to view performance statistics. These commands are all located in the `/sbin` directory, and you can view a list of them all using the `ls` command.

When you have finished using the console, you can log off by typing the `exit` command. To get back to the regular ESXi console at any time, just press `Alt-F2`.

Common Problems and Resolutions

Some of the common problems with ESX also apply to ESXi. The PSoD screen, which will lock up a host, can also occur on an ESXi host, and dealing with it is similar to dealing with it in ESX. Rebuilding or correcting networking configurations is also similar in ESXi. You can configure the management network by using the configuration options in the local console (which are accessible by pressing `F2`). You can also use the hidden console to access the `esxcfg` commands to view and configure the rest of your networking. In addition, you can restart the `vpix` management agents by selecting the option from the local console menu that is displayed by pressing `F2`. You can also view the VMkernel log from the local console at any time by pressing `Alt-F12`.

```
dcui
decodeSel
decodeSel.sh
dosfsck
dropbearmulti
esxcfg-advcfg
esxcfg-dhcp
esxcfg-dumppart
esxcfg-hwiscsi
esxcfg-info
esxcfg-init-esx
esxcfg-locker
esxcfg-loglevel
esxcfg-module
esxcfg-mpath
esxcfg-nas
esxcfg-nas
generateSLPReg.sh
groupadd
groupdel
hostd
hwclock
hwinfo
init
lspci
ntpd
openwsmand
partedUtil
poweroff
reboot
scantools
services.sh
sfcbd
vmkdump
vmkerrcode
vmkfstools
vmkgdbd
vmkiscsi-tool
vmkiscsi-util
vmkiscsid
vmkload_mod
vmklogger
vmkperf
vmkping
vmksystemswap
vmksitools
vsi_traverse
vsish
watchdog.sh

/sbin # esxcfg-vswitch -l
Switch Name      Num Ports      Used Ports      Configured Ports      MTU      Uplinks
vSwitch0         64              4                64                  1500     vmnic0

  PortGroup Name      VLAN ID      Used Ports      Uplinks
  VM Network          0            0              vmnic0
  Management Network  0            1              vmnic0

/sbin # _
```

Figure 10.13 Command listing and running the esxcfg-vswitch command in Tech Support Mode

Troubleshooting vCenter Server

A vCenter Server that is down will not affect hosts or running VMs, but certain features will no longer work while it is unavailable. The HA feature will continue to work, but features such as DRS and VMotion that rely on vCenter Server will not. In addition, you lose centralized management of your hosts while it is down, which can make managing large environments difficult. Therefore, it is important to troubleshoot and resolve any problems with vCenter Server as quickly as possible.

Log Files

The primary log file used by vCenter Server is the vpxd-#.log file (the # sign represents a number), which is located in the %allusersprofile%\Application Data\VMware\VMware VirtualCenter\Logs directory on the vCenter Server 2.5. The %allusersprofile% variable on most systems is the C:\Document and Settings\All Users directory. On older VirtualCenter 2.0 servers, this file was located in the C:\Windows\Temp directory. New log files are automatically created when the logs grow to a certain size, and the number in their filename is incremented. You can see which one is currently being used by sorting by modified date or by checking the vpxd-index file, which lists the number of the log file currently in use.

In addition to viewing the files on the vCenter Server's file system, you can view them using the VI Client when you connect to a vCenter Server. Just click the Administration button and select the System Logs tab, and you can then choose from the various log files listed. If you select vpxd-, you will see the number of the one currently in use.

You can control the number of logs kept, the maximum size in bytes, and the amount of information written to the log by editing the `vpzd.cfg` file, which is in the `%allusersprofile%\Application Data\VMware\VMware VirtualCenter` directory on the vCenter Server. There is a section for `<log>`, as shown here, where you can configure this.

```
<log>
  <level>info</level>
  <maxFileSize>10485760</maxFileSize>
  <maxFileNum>10</maxFileNum>
</log>
```

Here you can change the logging level to either None, Error, Warning, Info, Verbose, or Trivia, with each level displaying more information. These levels can also be changed using the VI Client under Administration, VirtualCenter Management Server Configuration, and then choosing Logging Options. You can also change the maximum file size and number of logs. When you change these settings, you must restart the vCenter Server Windows service for them to take effect.

If you are experiencing database-related problems with vCenter Server, you can also enable database logging in the `vpzd` log by adding the following section beneath the `</log>` line in the `vpzd.cfg` file:

```
<trace>
  <db>
    <verbose>true</verbose>
  </db>
</trace>
```

In addition to the vCenter Server log, there are logs for some of the other components, including Update Manager, Capacity Planner, and Converter Enterprise. The logs are located in the `\logs` subdirectory under the various product directories, which are located in `%allusersprofile%\Application Data\VMware\`.

If you do increase your log levels or enable database logging to troubleshoot a problem, just remember to change them back to avoid quickly filling up your logs and causing increased utilization on your vCenter Server.

You can also generate diagnostic log bundles for your vCenter Server that you can use to send to VMware support, who will usually request them when you open a case with them. To do this, connect to your vCenter Server with the VI Client and select File, Export, and then Export Diagnostic Data. When the selection window is displayed, don't select any hosts if you want to include only vCenter Server in the output, and make sure the Include Information from VirtualCenter Server and VI Client is selected. Then enter a directory on your workstation to save it in, and it will create the bundle. The bundle will contain log files,

configuration files, dump files, and files that contain the output of various commands that are run as part of generating the bundle.

Common Problems and Resolutions

The most common problems with vCenter Server include database problems and communication problems between the vCenter Server and the hosts that it manages. When you learn to recognize the symptoms of these problems, resolving them is often fairly simple.

Database Problems

Many problems that the vCenter Server may experience are caused by issues with its database. These types of problems can be caused by different things, and in most cases you will want to work with your database administrators to resolve them. Common occurrences include the following:

- **Authentication problems.** If your `vpdx.cfg` log shows problems authenticating with the database server, first verify that you can successfully authenticate with the database server. Open the ODBC Administration Control Panel on the vCenter Server and select the ODBC connection for the database server and click the Test Connection button. Enter the database username and password that is used by the vCenter Server and determine whether it can authenticate. If it cannot, check with your database administrator to make sure they are correct or that the account is not locked or modified. If you do need to change the database username and password, you must run the vCenter Server installation again and choose to do a Repair; the wizard will ask you for a database username and password. Just be sure you select to use an existing database server and you choose to *not* reinitialize the database (which would wipe out your existing data). The reason that you need to run the installation again is that the database username and password is stored in encrypted form in the registry of the vCenter Server and can only be set when installing vCenter Server or by accessing the settings using the VI Client once the vCenter Server is running. See VMware knowledge base article 1003928 for more information about this.
- **Connectivity problems.** The vCenter Server relies on its database constantly, and losing connectivity to the database server even for short periods can cause it to stop running. The `vpdx.cfg` log will show any connectivity errors in it. Have your network administrators make sure there are no connectivity or latency problems between the vCenter Server and the database server. Also, check for mismatched NIC speed and duplex and make sure both the physical NIC and physical switch port are set identically (for example, 100/Full). In addition, if your vCenter Server or database server is running on VMs, check that they are getting sufficient resources and are not having bottleneck problems.

- **Tablespace problems.** For Oracle databases, the tablespace assigned to the database that is used by vCenter Server may be set to not auto-extend automatically once it reaches the maximum size that was configured when the database was created. Check with your database administrators and have them extend the tablespace if needed. See VMware knowledge base article 1003982 for more information about this.
- **Disk space problems.** A common problem when using a SQL Server database is the transaction logs filling up and subsequently the database server runs out of disk space. Have your database administrator check the recovery model that you are using. The default Full Recovery Model causes the transaction logs to grow very large. Consider switching to the Simple Recovery Model instead to greatly reduce the size of the transaction logs. You can also shrink the transaction log to reduce its size. See VMware knowledge base article 1003980 for more information about this. You may also want to purge some of the older task and event and performance data from the database; this data may be responsible for taking up most of the space in the database. Task and event data is not automatically purged and can grow quite large over time. You can also adjust the performance data-retention intervals to decrease the amount of data kept in the database. You can also manually purge old data from the database by using some SQL scripts that VMware provides. See VMware knowledge base article 1000125 for more information about this.

Host Problems

Sometimes a host will show as disconnected in vCenter Server or will not be displaying up-to-date information or statistics. This can also happen when upgrading a vCenter Server, which also updates the management agents on the hosts. If for some reason the management agent cannot be upgraded, the host will show disconnected in vCenter Server. Restarting the management agent or the hostd agent will usually resolve these types of problems. To restart both services, type `service mgmt-vmware restart` and `service vmware-vpxa restart` from the Service Console. See VMware knowledge base article 4478241 for more information about this. As mentioned earlier, restarting these services will often resolve many problems, so it is recommended to do this first before trying other options.

Troubleshooting Virtual Machines

VM problems can sometimes be challenging to troubleshoot compared to physical servers because more things can affect your VM in virtual environments. When a server becomes stuck in a physical environment, you can always unplug the power cord. This is a last resort, but in a virtual environment doing this to a host will also affect all the other VMs running on the host. At the same time, however, virtual environments have some benefits that can make troubleshooting easier, such as being able to quickly mount bootable ISO files and mounting a disk from a VM with a problem onto another VM so that it can be fixed.

Log Files

There is only one log file for VMs on a host, and it is located in the working directory on the VMFS volume for the VM. The current log file is always `vmware.log`, and older log files are incremented numerically (for example, `vmware-5.log`). You can access this log file using some of the same methods as you would for other ESX logs, including using text editors, SCP utilities, the `tail` command, and you can also use the datastore browser built in to the VI Client to access it. A lot of information is written to this log, and much of it will not be relevant to any problems you are experiencing, but you should look for any obvious error messages. In addition, check the `VMkernel` and `hostd` log files on the host that the VM is located on for any errors that may relate to the problem that your VM is experiencing.

Common Problems and Resolutions

Many common VM problems relate to the state of the VM and changing it from one state to another. Sometimes a VM will get stuck and you will not be able to easily power it off, or you may encounter difficulty powering a VM on. Other times, you might have problems with a VM not booting properly or with its virtual disk file.

Power-State Problems

These types of problems can include issues with either powering off or on a VM. Occasionally a VM will not power off when using the power controls in the VI Client. Fortunately when this happens, there are several ways to forcibly shut down the VM so that you do not have to reboot the host to fix the VM. Listed here are a few ways to kill a stuck VM that will not power off. These methods are all done using the ESX Service Console and are listed in the order of usage preference:

- **Method 1.** You can use the `vmware-cmd` command, which is the CLI equivalent to using the VI Client. Here's how:
 1. Log on to the Service Console.
 2. Type `vmware-cmd -l` to get a list of all VMs running on the host and the path to their configuration file, as shown in Figure 10.14. Note the path uses the UUID or long name of the datastore. You can optionally use the friendly name instead.

```
[root@esx2 root]# vmware-cmd -l
/vmfs/volumes/4979f8f9-a7e0850c-0dd0-000802edb060/green/green.vmx
/vmfs/volumes/4979f8f9-a7e0850c-0dd0-000802edb060/yellow/yellow.vmx
/vmfs/volumes/4979f8f9-a7e0850c-0dd0-000802edb060/orange/orange.vmx
/vmfs/volumes/4979f8f9-a7e0850c-0dd0-000802edb060/purple/purple.vmx
```

Figure 10.14 Running the `vmware-cmd` command to list running VMs

3. You can first check the power state of the VM by typing `vmware-cmd VM config file path and name getstate`, as shown in Figure 10.15.

```
[root@esx2 root]# vmware-cmd /vmfs/volumes/esx2-local/purple/purple.vmx getstate
getstate() = on
[root@esx2 root]#
```

Figure 10.15 Running the `vmware-cmd` command to display the power state of a VM

4. To forcibly shut down a VM, type `vmware-cmd VM config file path and name stop hard`, as shown in Figure 10.16.

```
[root@esx2 root]# vmware-cmd /vmfs/volumes/esx2-local/purple/purple.vmx stop hard
stop(hard) = 1
[root@esx2 root]#
```

Figure 10.16 Running the `vmware-cmd` command to forcibly stop a VM

5. You can check the state again to determine whether it worked; if it did, the state should now be off, as shown in Figure 10.17.

```
[root@esx2 root]# vmware-cmd /vmfs/volumes/esx2-local/purple/purple.vmx getstate
getstate() = off
[root@esx2 root]#
```

Figure 10.17 Running the `vmware-cmd` command to confirm the VM has been powered off

- **Method 2.** You can use the `vm-support` command to shut down the VM by first finding the VM ID and then using the `vm-support` command to forcibly terminate it. This method does a lot more than shutting down the VM; it also produces debug information that you can use to troubleshoot an unresponsive VM.
 1. Log on to the Service Console.
 2. Type `vm-support -x` to list the VM IDs (VMID) of your running VMs. Optionally, you can also type `cat /proc/vmware/vm/*/names` to display VMIDs, as shown in Figure 10.18.
 3. To forcibly shut down the VM and generate core dumps and log files, type `vm-support -X VMID`. You will receive prompts asking whether you want to take a screen shot of the VM, which can be useful to determine whether there are any error messages, and you will also be prompted as to whether you want to send an NMI and an ABORT to the VM, which can aid in debugging, as shown in Figure 10.19. When the process completes, which can take five to ten minutes, a TGZ file will be created in the directory in which you ran the command that you can use for troubleshooting purposes. To avoid filling up your file system, when the file is created, switch to the `/tmp` directory when you run the command.

```
[root@esx2 root]# vm-support -x
VMware ESX Server Support Script 1.29

Available worlds to debug:

vmid=1080      orange
vmid=1085      green
vmid=1095      yellow
vmid=1100      purple

[root@esx2 root]#
[root@esx2 root]# cat /proc/vmware/vm/*/names
vmid=1080 pid=-1   cfgFile="/vmfs/volumes/4979f8f9-a7e0850c-0dd0-000802edb060/orange/orange.vmx"
x" uid="50 23 53 6b 63 f7 b5 76-ea 96 8a 31 59 81 94 d5" displayName="orange"
vmid=1085 pid=-1   cfgFile="/vmfs/volumes/4979f8f9-a7e0850c-0dd0-000802edb060/green/green.vmx"
x" uid="50 23 bf f8 62 9b 6b 80-38 2c 4b 91 6a 6d 32 d1" displayName="green"
vmid=1095 pid=-1   cfgFile="/vmfs/volumes/4979f8f9-a7e0850c-0dd0-000802edb060/yellow/yellow.vmx"
x" uid="50 23 b4 ab 4b 5e ad 14-f6 e4 05 39 b5 0c a2 4b" displayName="yellow"
vmid=1100 pid=-1   cfgFile="/vmfs/volumes/4979f8f9-a7e0850c-0dd0-000802edb060/purple/purple.vmx"
x" uid="50 23 47 07 0f be f0 b7-77 eb c4 56 40 ab 6f b9" displayName="purple"
[root@esx2 root]#
```

Figure 10.18 Running `vm-support` and `cat` to obtain a list of running VMs and their VMIDs

```
[root@esx2 root]# vm-support -X 1095
VMware ESX Server Support Script 1.29

Preparing files: |

Can I include a screenshot of the VM 1095? [y/n]: n
Can I send an NMI (non-maskable interrupt) to the VM 1095? This might crash the VM, but could aid
in debugging [y/n]: y
Can I send an ABORT to the VM 1095? This will crash the VM, but could aid in debugging [y/n]: y
Preparing files: \
Grabbing data & core files for world 1095. This will take 5 - 10 minutes.

Taking performance snapshots. This will take about 300 seconds.

Starting detailed scheduler stats.

Starting vscsiStats.
Snapping 7: 207 seconds left.[]
```

Figure 10.19 Running `vm-support` to forcibly stop a VM

4. You can check the state of the VM again either by using the `vmware-cmd` command or by typing `vm-support -x`, and you should not see the VMID for that VM listed anymore.
- **Method 3.** You can use the `kill` command by first finding the process identifier (PID) of the VM and then using the `kill` command to forcibly terminate it.
 1. Log on to the Service Console.
 2. Type `vmware-cmd -l` to get a list of all VMs running on the host and the path to their configuration file. Note the path uses the UUID or long name of the datastore. Optionally, you can use the friendly name instead.

3. You can first check the power state of the VM by typing `vmware-cmd VM config file path and name getstate`.
4. Type `ps auxfww | grep virtualmachinename` to get the PID of the VM, as shown in Figure 10.20. You will have two entries returned. The longer entry that ends in the config filename of the VM is the correct one. The number in the second column of that entry is the PID of the VM.

```
[root@esx2 root]# ps auxfww | grep purple
root      1050  0.0  0.2  3700  676 pts/0    S      10:58   0:00      \_ grep purple
root      31253 0.0  0.3  2012  908 ?        S<     10:23   0:00 /usr/lib/vmware/bin/vmkload_app /usr/lib/vmware/bin/vmware-vmx -ssched.group=host/user -# name=VMware ESX Server;version=3.5.0;license=VMware ESX Server;licenseversion=2.0 build-123630; -@ pipe=/tmp/vmhsdaemon-0/vmx1d7a697443d6e5f0; /vmfs/volumes/4979f8f9-a7e0850c-0dd0-000802edb060/purple/purple.vmx
[root@esx2 root]#
```

Figure 10.20 Running the `ps` command to get the PID of a running VM

5. Type `kill -9 PID` to forcibly terminate the VM, as shown in Figure 10.21.

```
[root@esx2 root]# ps auxfww | grep purple
root      1372  0.0  0.2  3692  672 pts/0    S      11:03   0:00      \_ grep purple
root      31253 0.0  0.3  2012  908 ?        S<     10:23   0:00 /usr/lib/vmware/bin/vmkload_app /usr/lib/vmware/bin/vmware-vmx -ssched.group=host/user -# name=VMware ESX Server;version=3.5.0;license=VMware ESX Server;licenseversion=2.0 build-123630; -@ pipe=/tmp/vmhsdaemon-0/vmx1d7a697443d6e5f0; /vmfs/volumes/4979f8f9-a7e0850c-0dd0-000802edb060/purple/purple.vmx
[root@esx2 root]#
[root@esx2 root]# kill -9 31253
[root@esx2 root]#
```

Figure 10.21 Using the `kill` command to forcibly stop a VM

6. You can check the state again to determine whether it worked. If it did, the state should now be off.

Additional power problems may occur when attempting to power on a VM. If this happens, check that there are enough resources available for the VM to meet any reservations that the VM may have assigned to it. If a VM has a 2GB memory reservation and the host does not have 2GB of free physical RAM available, the VM will not be allowed to power on. Instead, you will need to remove the reservation, move the VM to another host with enough free memory, or shut down or move other VMs to free up memory on the host. In addition, ensure that you have enough free disk space for the VSWP file that is created when the VM is powered on and removed when it is powered off. The VSWP file needs enough disk space equal to the amount of memory assigned to the VM minus the VM's memory reservations. You can set a memory reservation equal to the amount of memory assigned to the VM, which will result in a 0-byte VSWP file being created, but you should always take care to leave sufficient disk space free on your VMFS volumes for logs, VSWP files, and snapshots.

Other Problems

Sometimes you might encounter problems with your VMs not booting properly because of configuration changes or operating system problems. Virtual environments make these types of problems relatively easy to fix. You can easily mount an ISO file to the VM's CD/DVD-ROM drive and boot from a live CD (for example, Ultimate Boot CD or Knoppix) to fix the problem. In addition, you can mount the VM's disk file onto another VM, where you can access the file system and fix any issues with it. Just add a new virtual disk to another working VM and choose an existing disk and browse to the problem VM's virtual disk. Then power on the working VM and you can access the problem VM's disk as an additional drive. When you have finished, remove (not delete) the problem VM's disk from the working VM and power on the problem VM to determine whether the problem has been resolved.

Summary

Troubleshooting can often be frustrating and challenging, and knowing where to look and what to do is the key to quickly finding and resolving problems. You shouldn't just look through log files when you are experiencing known problems, however. Often, many problems might not be that obvious, and the log files are a good place to look for signs of them happening. You should keep a list of all the log files handy so that you can quickly access them if needed and so not have to waste time when a problem is happening trying to remember their path and filenames. You might not know how to resolve or troubleshoot every problem you encounter, so be sure to rely on the resources available to you, including documentation, support forums, knowledge base, and VMware's technical support. Being properly prepared to handle problems when they occur is one of the best troubleshooting skills that you can have.

Additional Resources

Listed here are some additional resources that will help you troubleshoot and resolve problems with your virtual environment.

VMware Knowledge Base

VMware puts considerable time and effort into their online knowledge base (<http://kb.vmware.com>), as shown in Figure 10.22, and it is full of fantastic information.

The knowledge base contains a lot more than solutions to known problems and issues. It also contains many how-to and troubleshooting documents. I urge you to periodically browse through the knowledge base and sign up for the weekly update emails that list all the new entries to the knowledge base. The knowledge base should be the first place you start when troubleshooting any problems. You can search on the keywords of the symptom of your problem and narrow the results by choosing your product name.




Figure 10.22 VMware’s knowledge base website

VMTN Forums

VMware has one of the best technical support forums around: the VMTN support forums (VMware Technology Network, <http://communities.vmware.com/community/vmtn>), as shown in Figure 10.23.

The discussion forums are filled with other VMware users from all over the world who use their knowledge and experience to assist others with problems and questions. The level of participation in these forums is very high; it is not uncommon to get an answer to your question in minutes and usually by multiple people. The forums are unofficially known as VMware’s tier 1 support. You can usually get your problem/question resolved in the forums faster than it would take you to call VMware. There are many VMware experts in these forums who help others just for the feeling of satisfaction that comes with helping someone else solve a problem.

A point system also exists that allows the author of a topic to reward helpful answers with points as a form of recognition. The forums tend to be very competitive, with users competing to collect points for posting responses to questions. Six points can be awarded by the person who asked the question for up to two helpful responses, and ten points can be awarded to one response that is deemed correct. The point system allows people to gain status levels as their points increase. There are nine status levels, ranging from Lurker (0–5 points) all the way up to the elite Guru level (20,000+ points). There are more than 500,000

forum members, with an average of about 20,000 new members being added each month. Figure 10.24 shows the multiple status levels that users can achieve in the forums.




Figure 10.23 VMware's VMTN support forums

Status Level Points	
	Guru 20,001 - 50,000 points
	Champion 10,001 - 20,000 points
	Virtuoso 5,001 - 10,000 points
	Master 2,001 - 5,000 points
	Expert 751 - 2,000 points
	Hot Shot 251 - 750 points
	Enthusiast 51 - 250 points
	Novice 6 - 50 points
	Lurker 0 - 5 points

Figure 10.24 VMTN forum status levels

The forums are also a good place for users of all levels to discuss advanced subjects and share methods and experiences with each other. There is also a fair amount of participation from VMware employees who do their best to also lend a hand in the forums. I spent a lot of

times in the forums when I was trying to learn more about VMware. I've made countless friends there, and would like to mention some of the most active VMTN contributors who consistently go out of their way to help others in the forums. These are a very bright, experienced bunch of guys, and if you are fortunate enough to have one of them answer your question, be assured that you are in very capable hands:

- Dave Mishchenko (forum name Dave.Mishchenko)
- Ken Cline (forum name Ken.Cline)
- Oliver Reeh (forum name Oreeh)
- Steve Beaver (forum name Sbeaver)
- Jason Boche (forum name Jasonboche)
- Edward Haletky (forum name Texiwill)
- Tom Howarth (forum name Tom Howarth)

So don't be afraid to ask your questions in the forums, and be sure to search on anything that you need help with. Chances are your question may have been already asked and answered. If you use any VMware product, I highly recommend that you head on over to the forums and check them out, even if you do not have a question of your own. Reading through some of the posts from others is a great way to expand your own knowledge as you benefit from the experiences of others. I also find that helping others in the forums is a great way to help yourself, because teaching and helping others is a great way to help yourself learn, too.

Chapter 11

Advanced Topics

This chapter covers some advanced topics to give you more insight into the technical workings of VMware VI3 and provide you with information about scripting, utilities, and additional resources that you can use to continue learning more about VMware VI3. You should be familiar with many commands when using the command-line interfaces for ESX and ESXi. In addition, scripting is a great way to help automate many of the common administration tasks that you will frequently perform.

Snapshot Technical Overview

When you start using snapshots with VI3, you'll wonder how you ever lived without them. Snapshots can be invaluable when making changes or updates to VMs, such as applying patches or upgrading applications. Although snapshots are a great feature, they can also cause big problems in your environment if they are not properly managed.

What Is a Snapshot?

A snapshot is a point-in-time picture of a VM that preserves the disk file system and optionally the memory of your VM. When a snapshot of a VM is taken, all writes to the original VM disk file are suspended, and it becomes read-only, and all new writes to the disk are done to a newly created delta virtual disk file, as shown in Figure 11.1.

The snapshot can then be deleted, which merges the delta file back into the original disk file, and then the delta file is deleted or you can revert to a snapshot, which discards the changes in the delta file and brings the VM back to the point of when the snapshot was taken. You can also create multiple snapshots that provide you with various restore points for a VM to revert to. This can be useful during a multistage upgrade to a server where one stage may go wrong and you don't have to start completely over.




Figure 11.1 Visual representation of a snapshot

A few different files are used with snapshots, as follows:

- **base filename-delta.vmdk file.** This is the differential file created when you take a snapshot of a VM, and such as all virtual disks consists of two files, a small disk descriptor file and the larger data file. The delta file is a bitmap of the disk block changes to the original VMDK and thus can never grow larger than the base VMDK. A delta file will be created for each snapshot that you create for a VM. These files are automatically deleted when the snapshot is deleted or reverted in Snapshot Manager.
- **base filename.vmsd file.** This file is used to store metadata and other information about snapshots. It is in text format and will contain information such as the snapshot display name, UID, and disk filename. It is initially a 0-byte file until you create your first snapshot of a VM, and from that point it will populate the file and continue to update it whenever new snapshots are taken. This file does not clean up completely after snapshots are taken. When you delete a snapshot, it still leaves the fields in the file for each snapshot and just increments the UID and sets the name to Consolidate Helper to be used with Consolidated Backups.
- **base filename.vmsn file.** This is the snapshot state file, which stores the exact running state of a VM at the time you take that snapshot. This file will either be small or large depending on whether you select to preserve the VM's memory as part of the snapshot. If you do choose to preserve the VM's memory, this file will be a few megabytes larger than the maximum RAM memory allocated to the VM. This file is similar to the VMSS (Suspend) file. A VMSN file will be created for each snapshot taken on the VM. These files are automatically deleted when the snapshot is.

Snapshot Growth

Snapshots start out small and can grow up to the maximum size of the original disk file and can never exceed the size of the original disk file. The reason for this is that any time a disk block is written to after a snapshot is taken it is created once in the delta file and simply updated if changed later. If you were to change every single disk block on your server after taking a snapshot, your snapshot would be the same size as your original disk file. Snapshot files initially start out small in size at 16MB and then grow in 16MB increments to help

reduce SCSI reservation conflicts. Whenever a request is made to change a block on the original disk, it is instead changed in the snapshot delta file. If the previously changed disk block in a delta file is changed again, it will not increase the size of the delta file because it simply updates the existing block in the delta file. The rate of growth of a snapshot will be determined by how much data is changed after the snapshot is taken. Servers that run applications such as SQL and Exchange that change data on the disk frequently will have their snapshot files grow rapidly. Servers with mostly static content that do not change much data on the disk such as web and application servers will grow at a much slower rate. When you create multiple snapshots, new delta files are created, and the previous delta files become read-only. With multiple snapshots, each delta file can potentially grow as large as the original disk file.

By default, snapshot files are created in the VM's home directory. You can change this behavior on a per-VM basis, which we cover shortly. You should plan on extra space on your VMFS datastores for snapshot files. How much space will depend on how many snapshots you plan to have simultaneously and how long you want to keep them. A general guideline is to leave free space equal to at least 25% of the size of a VM's disk file. So for a 20GB VM, you should leave at least 5GB of free space. This number will increase if you plan to keep multiple snapshots and for longer periods of time and if your VM is running a disk I/O-intensive application. Also, if you plan to include the memory state of your VM in the snapshot, you need to allow for extra disk space equal to the amount of memory assigned to the VM.

One other thing you should be aware of if you plan to use multiple snapshots is that you will need extra disk space available when you delete all the snapshots. Unlike a VM with a single snapshot, which requires no extra disk space to delete the snapshot, a VM with multiple snapshots has a multistep process for merging snapshots back into the original disk that requires extra disk space. An example of this is if you have a VM with three snapshots, called Snap1, Snap2 and Snap3, and you delete them all, first Snap3 is merged into Snap2, which causes Snap2 to grow larger. Next, Snap2 is merged into Snap1, which causes it to grow even larger, and finally Snap1 is merged into the original disk file, which requires no extra disk space. Snap1, Snap2, and Snap3 are not deleted as they are merged. Once they have all been merged into the original disk file, they are all deleted. This can require up to twice the free disk space to be able to merge them back into the original disk file and delete them because each snapshot must grow as it is merged into the next one. If you are low on space on your VMFS volume, you will be unable to delete them to free up space. You also want to be careful that you never run out of space when using or deleting snapshots because this can cause disk corruption of your VMs. Alternatively, you can delete multiple snapshots one by one, starting with the one further down the snapshot tree, which requires much less disk space than if you were to delete them all simultaneously.

Using and Committing Snapshots

Committing a snapshot is the same as deleting it because the delta files are committed back into the original virtual disk file when a snapshot is deleted. Although snapshots are very

useful, take care to not leave them running any longer than needed. The reason for this is besides the extra disk space they consume they can also have a negative effect on the performance of your VM and hosts in several different ways, including the following:

- When you create a snapshot, it will cause a brief pause of your VM as the delta file is created. This can cause issues with applications that are very sensitive to timely disk and network I/O.
- Creating a snapshot requires a lock on the LUN that the VM is on, which is also known as a SCSI reservation, because the host needs exclusive access to the VM's disk file. Excessive SCSI reservations can cause disk performance issues as other hosts need to wait to write to the disk.
- While a snapshot is active, the disk performance of the VM will be slightly reduced because ESX writes to snapshot delta files differently than it does to standard virtual disks. Delta files grow in 16MB increments, which causes slightly reduced performance because the delta file needs to be grown as disk writes occur. In addition, every time the delta file is grown in a 16MB increment, it causes a SCSI reservation that briefly locks the LUN, which can affect other VMs and hosts that write to the LUN.
- Deleting a snapshot also creates a SCSI reservation, and the delta files being committed back into the original disk on the VM can greatly reduce its disk I/O performance. You should consider deleting snapshots during periods of inactivity when the VM and host are not that busy.

When deleting a snapshot, be aware that the snapshot can take quite a long time to commit depending on how long it has been running. Snapshots that have been left running for weeks with a lot of data changes after the snapshot was taken can sometimes take hours to commit. This time will vary based on how busy the host is. A snapshot will commit faster if a VM is powered off. When deleting a snapshot, you will also notice that the status of the Delete Snapshot task in the VI Client is misleading because it jumps to 95% fairly quickly and stays there until the snapshot is deleted. In addition, vCenter Server has a 15-minute default timeout on tasks, so it may show as timed out even though the task is still running. You can increase this timeout by adding the following code to the <config> section of the vpxd.cfg file on the vCenter Server. The `timeout_value` must be a number in seconds with a maximum value of 1800.

```
<task>
  <timeout>timeout_value</timeout>
</task>
```

In addition, you can browse the datastore to see when the snapshot delta files disappear so that you know the snapshot commit has finished. So, be sure to remove any snapshots as soon as possible to ensure you do not experience problems and to ensure the optimal performance of your hosts and VMs.

By the Way

Committing large or deeply nested snapshots in ESX 3.5 can take longer than 3.0 because of a change in the consolidation algorithm in ESX 3.5, which was made to better handle VMs running with high I/O loads. This delay affects only the removal of snapshots and not other snapshot operations.

Using VMotion with Snapshots

You can VMotion a VM that has active snapshots, but you will receive a warning message stating that reverting to snapshots would generate error warnings on the destination host. What this means is that if you change the default locations for any files of the VM (for example, working directory or VSWP files), the VM will crash when the migration has completed if the destination host cannot access the same storage that the files are located on as the source host. VMware does recommend that you commit all snapshots before performing VMotion of a VM, but deleting and reverting will work just fine if you do not.

Excluding VM Disks from Snapshots

If you have a VM with more than one disk and you want to exclude a disk from being included in a snapshot, you must edit the VM's settings and change the disk mode to Independent, and then make sure you select Persistent. Independent mode provides you the means to control how each disk functions independently, and there is no difference to the disk file or structure. When a disk is independent, it will not be included in any snapshots. One thing to note is you will not be able to include memory snapshots on a VM that has independent disks. This is done to protect the independent disk in case you revert back to a previous snapshot with a memory state that may have an application running that was writing to the independent disk. Because the independent disk is not reverted when the other disks are, it could potentially corrupt data on it.

Changing Snapshot File Locations

When you create a snapshot, the default is to store it in the home directory of the VM. You can change this behavior, however, if you want to store it on a separate VMFS datastore. Having all your snapshots on one shared datastore will limit the disk locks from SCSI reservations to that datastore and not affect your VMs that are not running snapshots. Keep in mind that if you use a local datastore for your snapshots, you will not be able to use features such as HA, DRS, and VMotion because all hosts will not have access to that local datastore. To specify a new location for the snapshots of a VM, you have to edit the configuration file (VMX) of the VM. This is only a per-VM setting and not a host setting, so you will have to set

it on each VM for which you want to store the snapshots elsewhere. When you specify a new working directory for snapshots on each VM, however, both snapshot and VSWP files are written to this directory. You can optionally use another setting to have your VSWP file reside in the VM's home directory instead. To do this, follow these steps:

1. Make sure the VM is powered off.
2. Edit the settings for the VM, and on the Options tab select General, and then click the Configuration Parameters button.
3. Click Add Row, and in the first column put `workingDir`, and in the second column put the path to the directory in which you want to store the snapshots (for example, `/vmfs/volumes/ESX1-Local/Snapshots`), and then click OK to save it. Note that this setting will be added to the VMX file, but if you go back into the Configuration Parameters you will not see it. Optionally, you can add the setting directly to the VMX file by adding the line `workingDir = "/vmfs/volumes/volume name/directory name/"` to it. If you want your VSWP files to be in a different directory, add or edit the line `sched.swap.dir = "/vmfs/volumes/VM-Volume1/MyVM/"`.
4. Power on the VM, and all snapshots will go into the directory you specified.

Locating VM Snapshots

Sometimes you take snapshots and forget about them, which is not good (for the previously mentioned reasons). Unfortunately, vCenter Server does not have a good reporting tool to show snapshots that are running in your environment. The only way you can determine whether you have running snapshots in vCenter Server is to check each VM individually. Until VMware improves this, there are a few easy ways you can check for running snapshots in your whole environment:

- **PowerShell script.** You can use simple PowerShell commands to check this. You can find a few sample scripts to do this at the following websites:
 - ◆ www.peetersonline.nl/index.php/vmware/powershell-oneliner-4/
 - ◆ <http://communities.vmware.com/docs/DOC-6980>
- **Service Console commands.** You can use the Linux `find` command to find snapshot files. Just log on to the Service Console and change to the root of your VMFS volume and type the command `find -iname "*-delta.vmdk"` to find all snapshots, or to find orphaned snapshots that have not been modified in a number of days type `find -iname "*-delta.vmdk" -mtime +7 ls`.

- **Perl scripts.** You can use some free Perl scripts on the Service Console that make use of the VI Perl Toolkit. You can schedule these scripts to run and generate email reports on running snapshots. You can download these scripts from the following websites:
 - ◆ Snapalert from Dominic Rivera
<http://vmprofessional.com/index.php?content=snapalert>
 - ◆ Snaphunter from Xtravirt
<http://engineering.xtravirt.com/products/phd-technologies/snaphunter.html>

Virtual Machine Technical Overview

In this section, we cover the various virtual hardware components and files that make up a VM. A *virtual machine* is defined as the representation of a physical machine by software that has its own set of virtual hardware upon which an operating system and applications can be loaded. VMs all use the same virtual hardware regardless of the physical hardware that the host is running on. When you create a VM and assign virtual hardware components to it, such as CD/DVD drives and SCSI adapters, these components will all be seen by the operating system as specific hardware components. The exception to this is virtual CPUs, which are seen by the VM as whatever brand and type is running on the host, which is the reason why you cannot VMotion between hosts with different physical CPU families and brands. The standard virtual hardware that a VM sees is listed here:

- **System manufacturer.** VMware.
- **BIOS.** Phoenix 6.0.
- **Motherboard.** Intel 440BX.
- **CPU.** Will vary based on whatever CPU brand and model is in the host servers. VMs running on AMD hosts will see whatever AMD processor is in the host, and VMs running on Intel hosts will see whatever Intel processor is in the host. A VM will only see the number of vCPUs assigned to it, not the actual physical number that the host has. All CPUs presented to VMs are seen as single-core CPUs regardless of whether the host has multicore CPUs.
- **Memory.** VMs will see four memory slots, which will be configured with virtual DIMMs automatically based on the amount of memory assigned to the VM. A VM with 512MB of memory will typically see one slot configured with a 512MB DIMM, and a VM with 4GB of memory will typically see two slots configured with 2048MB DIMMs.
- **Video controller.** VMware Standard VGA Graphics Adapter with 4MB video memory.

- **CD/DVD-ROM drive.** NEC VMware IDE CDR00.
- **Network controller.** Will vary depending on the operating system that you choose when configuring the VM. The most commonly used NIC in most 32-bit VMs is based on the AMD PCnet physical NIC and is used with the flexible or Vlane adapter types. Additional NICs include ones based on Intel's e1000 (64-bit VMs and 32-bit Vista VMs) and VMware's own vmxnet (no physical counterpart, used in ESX 2.x).
- **IDE controller.** Intel 82371 AB/EB PCI Bus Master IDE Controller.
- **SCSI controller.** Can be either an LSI Logic PCI-X Ultra320 or a BusLogic BA80c30 PCI-SCSI MultiMaster depending on the operating system chosen when creating a VM. LSI Logic is the preferred choice and offers better performance. The BusLogic is typically used by older operating systems. You can manually select the SCSI controller type if you choose the Custom wizard type rather than Typical when creating a VM. In addition, you can change the type later by editing the VM configuration; be careful, however, because some operating systems do not have both drivers and might not boot when changed.

Note that USB and audio devices (sound cards) are not supported at all on ESX hosts. The total IDE devices in a VM cannot exceed four, which is the limit of CD/DVD drives that you can add to a VM. Also, the total number of PCI devices in a VM cannot exceed six. Included in this total are NICs, SCSI controllers, and the video adapter. Because you can't remove the video adapter from a VM, that leaves five PCI devices to be divided up between NICs and SCSI controllers. You will notice that you cannot add a SCSI controller to a VM; controllers are automatically added or removed when you add hard disks and assign them a virtual device node ID.

VMs are made up of many different files, as shown in Figure 11.2, which are kept in the VM's home directory.

Each file type has a different purpose, such as the VM's BIOS, virtual disk, configuration, and so forth. Most of these files will be named as whatever the name of the VM is, and will have different extensions based on the type of file it is. Depending on the state of your VM, you might not see certain files unless the state changes. For example, the VSWP file is present only when a VM is powered on, and the VMSS file is present only when a VM is suspended. Listed here are the various file types that make up a VM and a description of what each one is:

- **NVRAM file.** This small binary format file (less than 10K) contains the Phoenix BIOS that is used as part of the boot process of the VM. This is the same type of BIOS that is used in a physical server that lets you set hardware configuration options. You can access the BIOS by pressing the F2 key at the black splash screen that displays when a VM boots. Any changes made to the hardware configuration of the VM are saved in the NVRAM file. If this file is deleted or missing from a VM's directory, it is automatically re-created when the VM is powered on.




Figure 11.2 The many file types that make up a VM

- **VMX file.** This small, text format file contains all the configuration information and hardware settings of the VM. The data in this file is written when you make changes to the configuration of a VM using the VI Client. You can also edit this file directly when a VM is powered off, but you should be careful and make a backup copy beforehand. The type of data in this file includes things such as specific hardware configuration (for example, RAM size, NIC, and hard drive configurations), resource settings, VMware tools, and power management options.
- **VSWP file.** This is the memory swapfile created when a VM is powered on that is used if a host is overcommitted and uses all of its physical memory or when a memory limit is placed on a VM that limits the amount of physical host memory that a VM can use. You can also assign more virtual memory to a VM than a host physically has, and the VSWP file makes up for the physical memory that the host does not have. The size of this file is equal in size to the amount of memory assigned to a VM minus any memory reservations (default is zero) that a VM may have set on it. (For example, a 4GB VM with a 1GB reservation will have a 3GB VSWP file created.) Although these files are always created for VMs when they are powered on, they are used only if a host exhausts all of its physical memory. These types of files can take up a large amount of disk space on your VMFS volumes, so ensure that you have adequate space available for them because a VM will not power on if there is not enough free space available to create this file. These files are deleted when a VM is powered off or suspended.

- **VMSS file.** This file is used when VMs are suspended and is used to preserve the memory contents of the VM so that it can start up again where it left off. This file will be approximately the same size as the amount of memory assigned to a VM because even empty memory contents are written to it. When a VM is brought out of a suspended state, the contents of this file are written back into the physical memory of a host server. After this file is created, it is not automatically deleted until a VM is powered off (not rebooted). If a previous suspend file exists when a VM is suspended again, this file is reused rather than deleted and re-created. If this file is deleted while the VM is suspended, the VM will start normally and not from a suspended state.
- **VMSD file.** This text file is used with snapshots to store metadata and other information about each snapshot that is active on a VM. It is initially 0 bytes in size, and then once a snapshot is created it is updated with information as snapshots are created or deleted. There will be only one of these files present regardless of the number of snapshots running, because they all update this single file. The information in this file consists of the name of the VMDK and VMSN file used by each snapshot, the display name, description, and UID of each snapshot. When you delete snapshots, this file still retains old snapshot information, but increments the snapshot UID to be used with future snapshots. It also renames the first snapshot to Consolidate Helper, which is used if you ever use Consolidated Backup, which relies on snapshots.
- **VMSN file.** This file is similar to the VMSS file and is used with snapshots to store the state of a VM when a snapshot is taken. One of these files is created for every snapshot that is created for a VM, and they are automatically deleted when the snapshot is deleted. This file will vary in size depending on whether you choose to include the memory state of the VM with the snapshot. If you choose to store the memory state, this file will be slightly larger than the amount of memory assigned to the VM, because the entire memory contents, including empty memory, is copied to this file. If you do not choose to store the memory state of the snapshot, this file will be fairly small (under 32KB).
- **VMXF file.** This file is a supplemental configuration file that is not used with ESX but is retained for compatibility purposes with Workstation. It is in text format, and is used by Workstation for VM teaming, by which multiple VMs can be assigned to a team so that they can be powered on/off or suspended and resumed as a single object.
- **VMDK file.** These are the virtual hard disks files, and two different types of these files make up a single virtual disk. The first type is a large data file equal to the size of the virtual disk and contains the raw data of the virtual disk. The second type is a small (less than 2K) text disk descriptor file that describes the size and geometry of the virtual disk file and contains a pointer to the large data file and information about the virtual disk's drive sectors, heads, cylinders, and disk adapter type. In most cases, these files will have the same name as the data file that it is associated with (for

example, orangevm.vmdk and orangevm-flat.vmdk). You can tell which descriptor file is tied to which data file by checking the Extent Description field in this file to see which -flat, -rdm, or -delta file is linked to it. Three different types of virtual disk data files can be used with VMs:

- ◆ **-flat.vmdk file.** This is the default large virtual disk data file that is created when you add a virtual hard drive to your VM that is not a RDM. If you are using thick disks, this file will be approximately the same size as what you specify when you create your virtual hard drive. One of these files is created for each virtual hard drive that a VM has configured.
- ◆ **-delta.vmdk file.** These virtual disk data files are used only when snapshots are created of a VM. When a snapshot is created, all writes to the original -flat.vmdk are halted, and it becomes read-only; changes to the virtual disk are then written to these delta files instead. The initial size of these files is 16MB, and they are grown as needed in 16MB increments as changes are made to the VM's virtual hard disk. Because these files are a bitmap of the changes made to a virtual disk a single -delta.vmdk file cannot exceed the size of the original -flat.vmdk file. A delta file will be created for each snapshot that you create for a VM, and their filenames will be incremented numerically (for example, orangevm-000001-delta.vmdk, orangevm-000002-delta.vmdk). These files are automatically deleted when the snapshot is deleted after they have been merged back into the original flat.vmdk file.
- ◆ **-rdm.vmdk file.** This is the mapping file for a raw device mapping (RDM) that contains information for the mapping to a SAN LUN. The mapping file is presented to the ESX host as an ordinary disk file, and is available for the usual file system operations. However, to the VM, the storage virtualization layer presents the mapped device as a virtual SCSI device. The metadata in the mapping file includes the location of the mapped device (name resolution) and the locking state of the mapped device. If you do a directory listing, you will see that these files will appear to be taking up the same amount of disk space on the VMFS volume as the actual size of the LUN that it is mapped to, but in reality they just appear that way and their size is very small. One of these files is created for each RDM that is created for a VM.
- **LOG file.** These are the files that are created to log various types of information about the VM. There will be a number of these files present in a VM's directory. The current log file is always named vmware.log, and up to six older log files will also be retained with a number at the end of their name (for example, vmware-2.log). A new log file is created either when a VM is powered off and back on, or if the log file reaches the maximum defined size limit. The number of log files that are retained and the maximum size limits are both defined as VM advanced configuration parameters (log.rotateSize and log.keepOld).

Virtual Disks

In the preceding section, we covered the different types of VMDK files that can be present in your VM's directory. Now we cover the different disk formats that virtual disks can be. A virtual disk can be one of three general format types: raw, thick, or thin. This section describes the differences between each format type.

Raw Disks

Raw disks are known as raw device mappings (RDMs) and enable a VM to have a virtual disk that is mapped directly to a LUN on a SAN. As mentioned previously, RDMs have a small mapping file on a VMFS volume in the VM's home directory that contains mapping data to the SAN LUN. When you create an RDM for a VM, you can choose one of two modes that it can operate in. The first mode is virtual compatibility mode, which virtualizes the mapped device and is mostly transparent to the guest operating system. The other mode is physical compatibility mode, which provides minimal SCSI virtualization of the mapped device and the VMkernel passes most SCSI commands directly to the device, which allows for closer integration between the VM and the LUN.

Thick Disks

When you create a thick disk, all space on a volume is allocated at the time that it is created, so the resulting VMDK file will take up as much room on the volume as the size of the disk that you create. There are several types of thick disks, with the main difference between them being how disk blocks are zeroed before they are written to. The different types of thick disks are as follows:

- **Thick disks.** All space is allocated at creation time and may contain stale data on the physical media. Disk blocks are not zeroed before being written to the first time, which makes these types of disk less secure because any previously written data has not been cleared from the disk blocks.
- **Lazy-zeroed thick disks.** All space is allocated at creation time and wiped clean of any previous data on the physical media. Disk blocks are zeroed out on demand as they are written, but only for the first write to a disk block. This type of disk is the default used when creating virtual disks on VMFS volumes using the VI Client. These disks have slightly less I/O performance on the first write to a disk block because it must be zeroed before writing to it. Subsequent writes to a disk block have the same optimal performance as the other disk types.
- **Eager-zeroed thick disks.** All space is allocated at creation time and wiped clean of any previous data on the physical media. All disk blocks are zeroed out when the disk is created, which increases the time it takes to create these disks compared to the other

types. These types of disk are the most secure and offer slightly better performance only on the first write to a disk block because it has been already zeroed. Subsequent writes to a disk block have the same optimal performance as the other disk types.

Thin Disks

When you create a thin disk, its initial size is very small (1MB), and it grows as data is written to it up to the maximum size that was defined when the disk was created. Thin disks have a slight performance penalty both as they grow and space is allocated on demand and when the first write to a disk block is zeroed on demand. When the disk has grown and its blocks have been zeroed, they have the same performance as the other disk types. Thin disks are good for conserving disk space on a VMFS volume but can cause problems if you do not monitor their growth to make sure you do not run out of disk space on your VMFS volume. Thin disks are often used by default on NFS datastores based on the allocation policy of the NFS server and not the ESX server. You can view the current actual size of a thin disk using the VI Client datastore browser. Other methods such as the Service Console command line `ls` command will show the maximum size of the disk and not the actual size. You can also use the `du -a -h` command to see the actual size of a thin disk.

2GBSparse Disks

This is a special format disk that cannot be used with a running VM on an ESX/ESXi host and is instead often used to import/export virtual disks to and from ESX hosts. This format divides a virtual disk in up to multiple 2GB pieces, which is handy when you need to copy a virtual disk to another host or so that it can fit on physical media such as DVD-ROMs. The 2GB file sizes are used because some older file systems do not support file sizes larger than 2GB. Other VMware products such as Workstation and Server can use this format for running VMs, but for ESX hosts you must first import these type of disks into a thick or thin format using the `vmkfstools` utility.

Using Different Formats

In most cases, you will want to use the default lazy-zeroed thick disks because administration is simpler, but in some situations you may want to use the other disk formats:

- If you are concerned about security and want slightly better performance on initial disk writes, you might use eager-zeroed thick disks.
- If you have limited disk space on your VMFS volumes, you might use thin disks.

Whatever format you use, just remember that the disk performance is equal on all the formats once the initial write to a previously unwritten disk block is made. One exception to

this is with thin disks, which can become fragmented because the disk file is not allocated all at once, and if this happens, you may experience slightly reduced disk performance. Lazy-zeroed thick disks are the default format using the VI Client, and if you want to use the other formats, you need to create the disk using the Service Console utility `vmkfstools`. To create another format thick disk or a thin disk to use with your VM, follow these steps:

1. You first need to create a VM without a hard disk, which is an option in ESX 3.5 if you choose the Custom configuration when creating a VM rather than the Typical configuration. If you are using ESX 3.0, you can instead create a VM with any size hard disk and then delete it afterward.
2. Log on to the Service Console and change to the directory of the VM and use the `vmkfstools` command to create the hard disk by typing `vmkfstools -c disk size -d disk format disk filename`. For example, to create a 20GB thin disk for a VM called orange, use the command `vmkfstools -c 20G -d thin orange.vmdk`. If you use the `ls -lh` command afterward, it will show the disk as its maximum size (20GB) and not the true current size (1MB), because only the datastore browser shows the true current size.
3. Edit the VM and add a virtual hard disk and choose to use an existing disk and then browse to the newly created virtual disk in the VM's directory. You can now power on your VM and install an OS. If you check your VM disk size after you install your operating system, you will find that it has grown larger as more data is written to the disk.

If you want to change the format of existing disks, you can also do this with the `vmkfstools` command. Before you do this, however, confirm that your VM is powered off and does not have any running snapshots. Here are some examples of changing disk formats:

- `vmkfstools -w disk filename` This will completely zero out a virtual disk file; all existing data will be wiped out. This parameter should be used only on newly created virtual disks that have no data.
- `vmkfstools -j disk filename` This will inflate a thin disk to its maximum size and zero it out, which will change it to an eager-zeroed thick disk. Existing data will remain intact because only the new disk blocks are zeroed out and not the existing ones.
- `vmkfstools -i source disk filename destination disk filename -d thin` This will convert an existing thick disk to a thin disk. When the conversion completes, you need to remove the existing disk, add a new disk, and browse to the destination disk filename, and then when you are sure the VM boots okay you can delete the source disk file using the datastore browser. Note that your thin disk size will vary and might not match the amount of disk space in use by the VM. If data has been written in the OS and later deleted, the ESX host will see it as written to and not as free space.

You can read more about how to use the `vmkfstools` command in Appendix B of VMware's "Server Configuration Guide" in the Documentation section of their website.

Management Console Technical Overview

At some point, you will usually end up having to use either the ESX Service Console or the Remote Command-Line Utility (RCLI) for ESXi (can also be used with ESX). Both management consoles require the use of command-line commands to be able to perform functions inside of them. In addition, because the ESX Service Console is based on Red Hat Linux, it is useful to know some Linux commands, too.

Remote Command-Line Utility

The RCLI utility can be installed on a Windows or Linux workstation or by using the preconfigured virtual appliance, which can run on your host. Both can be downloaded from VMware's website at www.vmware.com/download/vi/drivers_tools.html. The RCLI is a collection of Perl scripts that connect to either an ESX or ESXi host using API interfaces and executes commands. It is a substitute for the actual program files that were located on the Service Console with ESX hosts. On Windows systems, when you install the RCLI it also installs the Active Perl application necessary to execute the Perl scripts, which can't be run in Windows natively.

When you install the RCLI, you can access it through the Start menu under the VMware folder or by simply going to a command prompt and going to the RCLI directory. When you are in the directory, switch to the bin subdirectory. You can do a directory listing to see a list of all the Perl files that can be used. The commands available as of ESX/ESXi 3.5 Update 2 are as follows:

- **resxstop** Similar to `esxstop` in the Service Console and currently only included in the Linux version of the RCLI. Monitors in real time how ESX Server hosts use resources. Runs in interactive or batch mode.
- **svmotion** Used to perform Storage VMotions, which move a VM's configuration file and optionally its disks while the VM is running. Runs on a vCenter Server.
- **vicfg-advcfg** Performs advanced configuration. Use this command as instructed by VMware.
- **vicfg-cfgbackup** Backs up the configuration data of an ESXi system and restores previously saved configuration data.
- **vicfg-dns** This command can list and specify your ESX Server host's DNS configuration.
- **vicfg-dumppart** This command can query, set, and scan an ESX Server host's diagnostic partitions.

- **vicfg-module** This command enables you to set and get VMkernel options. The command is usually used when VMware Technical Support, a knowledge base article, or VMware documentation instructs you to do so. This command is not supported on vCenter Server.
- **vicfg-mpath** This command can configure multipath settings for Fibre Channel or iSCSI LUNs.
- **vicfg-nas** This command can manipulate NAS file systems associated with your ESX Server host.
- **vicfg-nics** This command can manage physical NICs (that is, the Ethernet switches used by an ESX Server host).
- **vicfg-ntp** This command enables you to specify the NTP (Network Time Protocol) server for an ESX Server host.
- **vicfg-rescan** Rescans the storage configuration.
- **vicfg-route** This command can list or set the default IP gateway.
- **vicfg-snmp** Sets up SNMP (Simple Network Management Protocol).
- **vicfg-syslog** Specifies the syslog server and the port to connect to that server for ESXi hosts.
- **vicfg-user** Creates, modifies, deletes, and lists local direct access users and groups of users.
- **vicfg-vmhbadevs** This command can display information about available LUNs on ESX Server hosts.
- **vicfg-vmknic** Adds, deletes, and modifies virtual network adapters (VMkernel NICs).
- **vicfg-vswitch** Adds or removes vSwitches or modifies vSwitch settings.
- **vifs** This command performs common operations such as copy, remove, get, and put on files and directories.
- **vihostupdate-** This command enables you to perform maintenance of your ESXi hosts. The command can install software updates, enforce software update policies, and track installed software.
- **vmkfstools** Creates and manipulates virtual disks, file systems, logical volumes, and physical storage devices on an ESX Server host.
- **vmware-cmd** Performs VM operations remotely. This includes, for example, creating a snapshot, powering the VM on or off, and getting information about the VM.

For more information about using the RCLI, check out the “RCLI Installation and References Guide” on VMware’s website. The RCLI is updated periodically with new commands and functionality added to it, so make sure you check back with VMware’s website to determine whether newer versions are released.

ESX Service Console

When you log on to the ESX Service Console, you can use most traditional Linux commands and a number of commands that perform ESX- and VM-specific functions. You might not use these commands that often, but they are good to know for when you do need to use them, so you should get familiar with them.

ESX Commands

These commands enable you to configure and administer your host and VMs from the Service Console command line. Most of these commands duplicate functionality that is done in the VI Client, and are very useful when you have problems and are unable to connect to a host using the VI Client:

- **esxcfg-advcfg** Configures advanced options for ESX Server 3. To configure advanced options in VI Client, click Advanced Settings. When the Advanced Settings dialog box opens, use the list on the left to select the device type or activity you want to work with, and then enter the appropriate settings.
- **esxcfg-auth** Configures authentication. You can use this command to switch between the `pam_cracklib.so` and `pam_passwdqc.so` plug-ins for password-change rule enforcement. You also use this command to reset options for these two plug-ins. There is no means of configuring these functions in VI Client.
- **esxcfg-boot** Configures bootstrap settings. This command is used for the bootstrap process and is intended for VMware Technical Support use only. You should not issue this command unless instructed to do so by a VMware Technical Support representative. There is no means of configuring these functions in VI Client.
- **esxcfg-dumppart** Configures a diagnostic partition or searches for existing diagnostic partitions. When you install ESX, a diagnostic partition is created to store debugging information in the event of a system fault. You don't need to create this partition manually unless you determine that there is no diagnostic partition for the host. You can perform the following management activities for diagnostic partitions in VI Client:
 - ◆ **Determine whether there is a diagnostic partition.** Click Storage > Add, and check the first page of the Add Storage Wizard to see whether it includes the Diagnostic option. If Diagnostic is not one of the options, ESX already has a diagnostic partition.
 - ◆ **Configure a diagnostic partition.** Click Storage > Add > Diagnostic and step through the wizard.
- **esxcfg-firewall** Configures the Service Console firewall ports. To configure firewall ports for supported services and agents in VI Client, click Security Profile > Firewall > Properties and use the Firewall Properties dialog box to add services. You cannot configure unsupported services through the VI Client. For these services, use the `esxcfg-firewall` command instead.

- **esxcfg-info** Prints information about the state of the Service Console, VMkernel, various subsystems in the virtual network, and storage resource hardware. The VI Client doesn't provide a method for printing this information, but you can obtain much of it through different tabs and functions in the user interface. For example, you can check the status of your VMs by reviewing the information about the Virtual Machines tab.
- **esxcfg-init** Performs internal initialization routines. This command is used for the bootstrap process, and you should not use it under any circumstances. Using this command can cause problems for your ESX host. There is no VI Client equivalent for this command.
- **esxcfg-linuxnet** Converts vswif to eth when booting ESX into Service Console-only mode rather than into ESX mode. This command is used for the bootstrap process and is intended for VMware Technical Support use only. You should not issue this command unless instructed to do so by a VMware Technical Support representative. There is no VI Client equivalent for this command.
- **esxcfg-module** Sets driver parameters and modifies which drivers are loaded during startup. This command is used for the bootstrap process and is intended for VMware Technical Support use only. You should not issue this command unless instructed to do so by a VMware Technical Support representative. There is no VI Client equivalent for this command.
- **esxcfg-mpath** Configures multipath settings for your Fibre Channel or iSCSI disks. To configure multipath settings for your storage in VI Client, click Storage. Select a datastore or mapped LUN, and click Properties. When the Properties dialog box opens, select the desired extent if necessary. Then, click Extent Device > Manage Paths and use the Manage Path dialog box to configure the paths.
- **esxcfg-nas** Manages NAS mounts. You use this command to add, delete, list, and change the attributes of NAS devices. To view NAS devices in VI Client, click Storage and scroll through the Storage list.
- **esxcfg-nics** Prints a list of physical network adapters along with information about the driver, PCI device, and link state of each NIC. You can also use this command to control a physical network adapter's speed and duplex settings. To view information about the physical network adapters for the host in VI Client, click Network Adapters. To change the speed and duplexing for a physical network adapter in the VI Client, click Networking > Properties for any of the vSwitches associated with the physical network adapter.
- **esxcfg-resgrp** Restores resource group settings and lets you perform basic resource group management. In the VI Client, select a resource pool from the Inventory panel and click Edit Settings on the Summary tab to change the resource group settings.

- **esxcfg-route** Sets or retrieves the default VMkernel gateway route and adds, removes, or lists static routes. To view the default VMkernel gateway route in VI Client, click DNS and Routing. To change the default routing, click Properties and update the information in both tabs of the DNS and Routing Configuration dialog box.
- **esxcfg-swiscsi** Configures your software iSCSI software adapter. To configure your software iSCSI system in VI Client, click Storage Adapters, select the iSCSI adapter you want to configure, and click Properties. Use the iSCSI Initiator Properties dialog box to configure the adapter.
- **esxcfg-upgrade** Upgrades ESX Server from ESX Server 2.x to ESX Server 3. This command is not for general use. You complete the following three tasks when upgrading from 2.x to 3.x. Some of these can be performed in VI Client:
 - ◆ **Upgrade the host.** You upgrade the binaries, converting from ESX Server 2.x to ESX Server 3. You cannot perform this step from VI Client. For information about performing this upgrade, see the “Installation and Upgrade Guide.”
 - ◆ **Upgrade the file system.** To upgrade VMFS 2 to VMFS 3, suspend or power off your VMs and then click Inventory > Host > Enter Maintenance Mode. Click Storage, select a storage device, and click Upgrade to VMFS 3. You must perform this step for each storage device you want to upgrade.
 - ◆ **Upgrade the VMs.** To upgrade a VM from VMS 2 to VMS 3, right-click the VM in the Inventory panel and choose Upgrade Virtual Machine.
- **esxcfg-vmhbadevs** Prints a map of VMkernel storage devices to Service Console devices. There is no VI Client equivalent for this command.
- **esxcfg-vmknics** Creates and updates VMkernel TCP/IP settings for VMotion, NAS, and iSCSI. To set up VMotion, NFS, or iSCSI network connections in VI Client, click Networking > Add Networking, select VMkernel, and step through the Add Network Wizard.
- **esxcfg-vswif** Creates and updates Service Console network settings. This command is used if you cannot manage the ESX host through the VI Client because of network configuration issues. To set up connections for the Service Console in VI Client, click Networking > Add Networking, select Service Console, and step through the Add Network Wizard.
- **esxcfg-vswitch** Creates and updates VM network settings. To set up connections for a VM in VI Client, click Networking > Add Networking, select Virtual Machine, and step through the Add Network Wizard.
- **vmkfstools** You use the vmkfstools utility to create and manipulate virtual disks, file systems, logical volumes, and physical storage devices on the VMware ESX Server hosts. Using vmkfstools, you can create and manage a VM file system (VMFS) on a

physical partition of a disk. You can also use the command to manipulate files, such as virtual disk files, stored on VMFS 2, VMFS 3, and NFS. You can perform most `vmkfstools` operations using the VI Client.

For information about how to use these commands and the various options that you can use, just type the name of the command without any options to see a list of syntaxes and options.

Linux Commands

As previously mentioned, most Linux commands work in the ESX Service Console, and you can use them for troubleshooting and configuration purposes. Many new VMware administrators without Linux experience are intimidated by the ESX Service Console, which is based on Red Hat Linux. The reality is that unless you want to do a lot of the configuration and administration of ESX the hard way, you will not have to log on to the Service Console and use the CLI that much. Almost all the configuration and administration of the ESX server can also be done using the VI Client that is provided with ESX. If you do want to get your hands dirty and use the CLI, it's best to learn some basic Linux commands. In certain situations, using the CLI is necessary, such as restarting some of the ESX services, reconfiguring networking, or killing a stuck VM. Listed here are some common Linux commands that you might use in the Service Console:

- `cp` Copy a file (not recommended to use for VMDK files; use `vmkfstools` instead)
- `rm` Delete a file
- `mv` Move or rename a file (not recommended to use for VMDK files; use `vmkfstools` instead)
- `nano` Edit a file (simple editor)
- `vi` Edit a file (more complex editor)
- `cat` Type a file to the screen
- `find` Find a file
- `ls` Lists files
- `rmdir` Remove an empty directory
- `mkdir` Create a directory
- `cd` Change to a directory
- `rm -rf` Delete a directory (`-rf` are options)
- `pwd` Check current path
- `chmod` Change file permissions
- `chown` Changes file ownership
- `df` Shows drive partition space

- `vdf` Shows VMFS volume space
- `du` Reports the size of files in a directory

Some usage examples of these commands are listed here:

- `df -h` will show drive and space information (no VMFS volumes, though). The `-h` will tell it to show the disk information in human-readable format (for example, 1.9G rather than 1967156) and can also be used with other commands, such as `ls`.
- `vdf -h` is a VMware-specific command and will show VMFS volumes and space. The `-h` will tell it to show the disk information in human-readable format (1.9G rather than 1967156).
- `cat /proc/scsi/scsi` will show SCSI device information.
- `cat /proc/vmware/vm/*/names` will show running VMs and their VMIDs
- `ls -l` will show detailed file information, including size, permissions, owner, and last modified.
- `ls -sh` will show basic file and size information.
- `du -a -h` will show total disk usage for a directory.
- `find -name '*.vmx' -type f -print0 | xargs -0 grep "tools.syncTime"` will find text inside a file (that is, search for the string `tool.syncTime` in all the VMX files).
- `find -iname "*-flat.vmdk" -ls` will find files (that is, search for all virtual disk files that have `-flat.vmdk` in the filename).

Watch Out!

Just remember that unlike Windows, all Linux commands are case sensitive and usually lowercase, so make sure you get the case right when typing them; otherwise, they will not work. This also is true about directory and filenames on the file system.

VMware Infrastructure Management Assistant

VMware Infrastructure Management Assistant (VIMA) is a Red Hat Linux 64-bit appliance that you can download in OVF format and import into ESX. It runs as a VM and includes prepackaged software to help you manage your VMware environment. It includes the latest version of the RCLI, VI Perl Toolkit, VMware Tools, a centralized logging component, and an authentication component that supports noninteractive login. VIMA acts as a centralized Service Console for all your hosts, and you can also use VIMA to run scripts on both ESX and ESXi hosts. When you add hosts to VIMA, you only need to log on to the VIMA appliance and not the ESX hosts to manage them. VIMA is strictly a command-line utility and has no

GUI and is a great alternative to using the various interfaces and tools that VMware provides for managing ESX and ESXi hosts. You can download and learn more about VIMA at VMware's website www.vmware.com/download/vi/drivers_tools.html.

Scripting Options

VMware has two powerful scripting toolkits that you can use to write your own scripts to help automate administration tasks. Both the Perl and PowerShell toolkits contain libraries that can be used to connect to the API on ESX/ESXi host and vCenter Servers to do reporting, configuration, and much more. Scripting is a great way to automate repetitive tasks, overcome shortcomings in the product, and configure new hosts and VMs.

Perl

Perl is a scripting language originally developed for UNIX/Linux systems and is installed inside the ESX Service Console, which is based on Red Hat Linux. There is also a version available for Windows called ActivePerl that is included in the Windows version of the RCLI and can also be downloaded and installed separately. Perl scripts typically have a .pl extension and can be written with any text editor. After you have Perl installed, you can download and install VMware's VI Perl Toolkit, which is available at www.vmware.com/support/developer/viperltoolkit/ and will add specific commands that can be used in your VMware environment. Be sure to read the documentation to learn the many VMware-related methods and properties that you can use in your scripts. You might also check out the various prebuilt Perl scripts that make up the commands that are used in the RCLI so that you can see how to write scripts of your own.

PowerShell

PowerShell is an extensible command-line shell and associated scripting language developed by Microsoft that can be used to help automate common administration tasks and provide information about your VMware environment. PowerShell can be used for many different things in Windows environments, but can also be used with VMware environments since VMware released their VI Toolkit, which provides PowerShell with access to the VMware API. PowerShell is fairly easy to install and use, and many great scripts have been written that work with VMware environments.

After you download PowerShell from the Microsoft website at www.microsoft.com/windowsserver2003/technologies/management/powershell/download.mspx and install it, you then need to download and install VMware's VI Toolkit for Windows at www.vmware.com/support/developer/windowstoolkit/, which contains the many PowerShell cmdlets that are specific to VMware that you can use. You can run the individual cmdlets from the special PowerShell command prompt or put them together into a script to perform

multiple tasks. PowerShell scripts typically have a .ps1 extension and can be written with any text editor.

Using PowerShell with VMware has grown in popularity, and there are currently many freely available scripts written by other users that you can use in your environment. For more information about using PowerShell and sample scripts, check out the following resources:

- Getting Started with VMware's PowerShell Toolkit
www.rtfm-ed.co.uk/docs/vmwdocs/whitepaper-powershell.pdf
- Using PowerShell for Virtually Everything
www.peetersonline.nl
- VMware Community Sample Code
<http://communities.vmware.com/community/developer/utilities>
- Managing VI3 with Windows PowerShell (VMworld 2007 presentation, free registration required)
www.vmworld.com/vmworld/mylearn?classID=11465
- Managing VMware with PowerShell FAQ
<http://communities.vmware.com/docs/DOC-4210>
- Virtu-al
<http://virtu-al.net/>
- VI Toolkit Quick Reference Guide
<http://virtu-al.net/Downloads/VIToolkitQuickReferenceGuide.pdf>
- VI Toolkit Developer Community
http://communities.vmware.com/community/developer/windows_toolkit

Looking at sample code is often the best way to learn and understand how to script, so be sure to check out the many scripts that other users have written to get you started writing your own. In addition, check out VMware's Developer Community at <http://communities.vmware.com/community/developer> for more scripting resources and to post any questions you may have.

You can also add a graphical interface to PowerShell using Quest Software's PowerGUI and their PowerPack for VMware, which makes PowerShell easier to use and adds additional functionality. These are available at www.powergui.org/downloads.jspa and www.powergui.org/entry.jspa?externalID=1802andcategoryID=290.

Third-Party Utilities

Many VMware-provided and third-party utilities can help you administer your virtual environment. Often, these utilities make up for shortcomings in VMware's product and make for

easier administration of your hosts and VMs. The following subsections discuss some free and paid utilities that you might want to use in your environment.

SSH Console Utilities

These utilities enable you to remotely connect to your ESX host Service Console using the secure SSH protocol for remote administration.

- **Putty.** Almost every admin has this one installed on their workstation because it is one of the best utilities that you can use to connect to your ESX hosts remotely using SSH. Available for free at www.chiark.greenend.org.uk/~sgtatham/putty/download.html.

SCP File-Transfer Utilities

These utilities enable you to transfer files to and from ESX hosts using the secure SCP protocol:

- **Veeam FastSCP.** One of the fastest file-transfer utilities because it does not encrypt data as it is copied. Available for free at <http://veeam.com/vmware-esx-fastscp.html>.
- **WinSCP.** A good file-transfer utility that has a Windows Explorer-like interface. Available for free at <http://winscp.net/eng/index.php>.

Miscellaneous Utilities

These utilities will help you with a variety of administration tasks.

- **RV Tools.** A very handy .NET application that uses the VI SDK to display information about your VMs. It can display all sorts of information about VMs and datastores, disconnect CD-ROM drives from VMs, list VMware Tools versions, and update them to the latest version. Available for free at <http://rvtools.deveij.com/>.
- **KS QuickConfig.** Is designed to reduce the time to deploy and configure VMware ESX 3 servers and eliminate inconsistencies that can arise with manual operations. Available for free at <http://engineering.xtravirt.com/products/ks-quickconfig.html>.
- **VP Snapper.** A utility that lets you revert multiple VM snapshots at one time rather than having to do them one by one. Available for free at http://virtualizeplanet.com/mambo/index.php?option=com_contentandtask=blogsectionandid=5andItemid=40.
- **VMDK Recovery Tool.** VMware's utility that can help recover VMDK files if they are deleted and help if a VMFS volume becomes corrupted or deleted. Available for free at <http://kb.vmware.com/kb/1007243>.

- **VMware Converter.** VMware's great application that lets you perform physical to virtual (P2V) and virtual to virtual (V2V) operations. Available for free at <http://vmware.com/products/converter/>.
- **vmCDconnected.** A handy utility that scans all VMs and shows whether they have a CD connected to it. After scanning, you can disconnect all the CDs with a click of a button. Available for free at www.ntpro.nl/blog/archives/172-Software.html.
- **CPU Identification Utility.** VMware's utility that displays CPU features for VMotion compatibility, EVC, and indicates 64-bit VMware support. Available for free at https://www.vmware.com/download/shared_utilities.html.
- **VMTS Patch Manager.** A great ESX host patching application for those who do not have Update Manager. Available for free at www.vmts.net/VMTSPatchManager.htm.
- **SnapHunter and SnapAlert.** Two utilities that can report all running snapshots on ESX hosts, including name, size, and date. They can also automatically email reports and optionally commit snapshots. Available for free at <http://engineering.xtravirt.com/products/snaphunter.html> and <http://vmprofessional.com/index.php?content=snapalert>.
- **Visio Stencils.** Some Visio stencils to help you document your environment. Available for free at <http://veeam.com/vmware-esx-stencils.html> and www.vmguru.com/content/view/33/63/ and www.visiocafe.com/index.htm.
- **VMotion Info.** A utility that gathers the system and CPU information from your hosts and puts this in one single overview to determine whether they are compatible for VMotion. Available for free at www.run-virtual.com/?page_id=155.
- **Tripwire OpsCheck.** A utility that helps ensure your systems are configured to support VMotion by rapidly analyzing ESX 3.0, 3.5, and ESXi hosts and providing troubleshooting guidance for VMotion compatibility. Available for free at <http://www.vwire.com/free-tools/opscheck/>.
- **ITQ VLAN and Port Group Manager.** A network utility that enables you to set up VLANs and port groups on multiple ESX hosts at once. Available for free at www.run-virtual.com/?page_id=160.

Monitoring Utilities

These utilities can be used to help enhance your monitoring of your virtual infrastructure and are more robust than the monitoring built in to vCenter Server:

- **Solarwinds VM Monitor.** A nice graphical VM monitor that shows statistics for host and VMs and their power state and other information. Available for free at https://www.solarwinds.com/products/freetools/vm_monitor.aspx.

- **Veeam Monitor Free Edition.** An easy-to-use VMware monitoring solution that provides real-time performance monitoring and alerting. Available for free at <http://veeam.com/esxi-monitoring-free.html>.
- **Veeam Monitor.** The full paid version of Monitor that integrates with vCenter Server and provides enhanced monitoring and alerting of the health and performance of your environment. It can also be used without vCenter Server to monitor individual ESX hosts. Available at <http://veeam.com/vmware-monitoring.html>.
- **eG VM Monitor.** A powerful reporting and monitoring paid application specifically for ESX hosts. eG VM Monitor is a web-based application specifically for VMware environments and is part of the larger eG Enterprise Suite. Available at www.eginnovations.com/web/vmware.htm.
- **Vizioncore vFoglight.** A powerful reporting and monitoring paid application that is specifically for monitoring and analyzing ESX hosts. It integrates with vCenter Server and provides enhanced reporting capabilities, configurable dashboards, intelligent rules and alarms, and custom reporting to identify trends and bottlenecks. Available at <http://vizioncore.com/vFoglight/index.html>.
- **Nagios.** Although not ESX specific, Nagios is a free, open source host, service, and network monitoring application that can be installed and configured on a Linux server and is also available as a preconfigured virtual appliance that can be installed on an ESX host. Nagios can be configured to monitor a wide range of devices, including Windows servers, Linux servers, UNIX servers, network printers, routers/switches, and services such as HTTP, SSH, FTP, and so on. Nagios can also be configured to receive SNMP traps from other applications, such as vCenter Server and hardware agents. Available for free at www.vmware.com/appliances/directory/372.

Backup Utilities

These utilities are specifically designed to work with virtual environments and can be used to either complement or in lieu of your current backup methods:

- **VISBU.** A backup utility that is run from the Service Console that provides VMDK-level backups of any VM on storage accessible by the host. Available for free at <http://engineering.xtravirt.com/products/visbu.html>.
- **VM Backup Script.** A free backup script to perform hot backups of your VMs. Available for free at www.vmts.net/vmbk3.htm.
- **PHD esXpress.** A paid disk-to-disk backup application that can integrate with VMware's VCB and offers many powerful backup and restore features. Available at <http://esxpress.com/>.

- **Vizioncore vRanger.** A paid disk-to-disk backup application that can integrate with VMware's VCB and offers many powerful backup and restore features. Available at <http://vizioncore.com/vRangerPro.html>.
- **VM Explorer.** A management tool that eases management, backup, and disaster recovery tasks in your VMware ESX Server environment. Available for free at www.trilead.com/Products/VM_Explorer.
- **Veeam Backup.** A paid disk-to-disk backup application that can integrate with VMware's VCB and offers many powerful backup and restore features. Available at <http://veeam.com/vmware-esx-backup.html>.

Storage Utilities

These utilities enable you to use local disk as shared storage and are a cheap alternative to more expensive shared storage solutions:

- **Openfiler.** An open source, browser-based storage appliance that supports NFS and iSCSI. It can be downloaded as an ISO to install on a server or as a VMware appliance to import to an ESX host. A great way to get more shared disk in your environment by turning physical servers into NAS servers or turning your local disk on your ESX hosts into shared disk when using the appliance. Available for free at www.openfiler.com/.
- **Xtravirt Virtual SAN.** A solution to turn local disk on your ESX hosts into shared VMFS volumes to avoid purchasing costly SAN disk. Available for free at <http://engineering.xtravirt.com/products/xtravirt-virtual-san.html>.
- **Lefthand Virtual SAN Appliance.** A paid application that can virtualize the internal storage of VMware ESX hosts and then cluster that across other ESX hosts to provide redundancy, eliminating the need for a more expensive physical SAN. VSA has an evaluation period during which the advanced features work, and after the evaluation expires the basic iSCSI target continues to work but without the advanced functions such as replication. Available at www.lefthandnetworks.com/vsa.aspx.

Security Utilities

These utilities are designed specifically for virtual environments to provide enhanced security, auditing, and reporting capabilities:

- **Tripwire ConfigCheck.** A utility that rapidly assesses the security of VMware ESX 3.0 and 3.5 hypervisor configurations compared to the VMware Infrastructure 3 Security Hardening guidelines. Available for free at <http://tripwire.com/configcheck/>.

- **Configuresoft Compliance Checker.** A free tool that provides a real-time compliance check for multiple VMware ESX host servers at a time and provides detailed compliance checks against both the VMware Hardening Guidelines and the CIS benchmarks for ESX. Available for free at www.configuresoft.com/compliance-checker.aspx.
- **Altor Networks Virtual Firewall.** A paid security product that enables you to monitor your virtual network security and create rule-based policies to protect VMs. It has many advanced networking features and can integrate into physical IDS and IPS systems. Available at <http://altornetworks.com/products/vnf/index.php>.
- **Reflex Systems Virtualization Management Center.** VMC is a paid security product that provides a single authoritative visual interface, central management, and security for virtual environments. VMC enables you to administer, audit, secure, and monitor complex, dynamic, virtual infrastructures. It has many network security features and uses rule-based policies to protect VMs. In addition, it can provide performance monitoring and reporting and event correlation for troubleshooting purposes. Available at <http://reflexsystems.com/Section/Products/VMC>.
- **Catbird V-Security.** A paid security product that use a V-Agent virtual appliance to secure the hypervisor and VMs from vulnerabilities and attacks. A stateless agent residing on the virtual network itself, the noninvasive, self-contained V-Agent has a view of the network impossible to see from traditional devices outside of the host. This product brings hypervisor security, VM/guest OS monitoring, policy compliance, and VM sprawl management under one single web-based management interface. Available at www2.catbird.com/our_services/vmware.php.

Reporting and Search Utilities

These utilities provide greatly enhanced reporting capabilities for your virtual environment and make up for some of the reporting shortcomings in vCenter Server:

- **Embotics v-Scout.** An agentless tool for tracking and reporting on VMs in VMware vCenter Server-enabled environments. Available for free at <http://embotics.com/solutions/v-scout>.
- **Hyper-9.** Their search-based reporting tool is a fresh approach to reporting and a great addition to every administrator's toolbox. It uses a query engine to search your virtual environment and compare results from various points in time. It can monitor your entire virtual environment from the operating system to the VMs and hosts. It also has a unique VM DNA feature that lets you compare a VM at different points in time or compare different VMs to each other. Available at www.hyper9.com.
- **MCS StorageView.** A utility that displays all the logical partitions, operating system, capacity, free space, and percent free of all VMs on ESX hosts or vCenter Servers. Available for free at www.mightycare.de/downloads/task_cat_view/gid,16/.

- **ESX HealthCheck.** A script that collects configuration and other data for ESX hosts and generates a report in HTML format. Available for free at <http://sourceforge.net/projects/esxhealthscript>.
- **Veeam Reporter.** Collects information about your VI3 environment, its components, and configuration settings on an ongoing basis. Using this data, it provides comprehensive and easy-to-understand reports for analysis, documentation, and change control. It's fully integrated with VMware vCenter Server and supports change control for both ESX and ESXi servers. Available at http://veeam.com/vmware-esx-reporting_enterprise.html
- **Vizioncore vFoglight.** A powerful reporting and monitoring paid application that is specifically for monitoring and analyzing ESX hosts. It integrates with vCenter Server and provides enhanced reporting capabilities, configurable dashboards, intelligent rules and alarms, and custom reporting to identify trends and bottlenecks. Available at <http://vizioncore.com/vFoglight/index.html>.

There are many more third-party utilities, both free and paid, available from many different vendors, so be sure to check around for them. Also, be sure to check out the many great virtual appliances that can be used in your environment at VMware's Virtual Appliance Marketplace.

Summary

This chapter covered some different topics to help you gain a better understanding of VMware VI3 and to help you administer your environment. The better you understand how everything works, the better equipped you will be to administer and troubleshoot your environment. In addition, we provided some information about the many third-party products available that you can use in your environment to help you automate and enhance your virtual environment. The many additional add-on products that are available are what make VMware VI3 the strongest virtualization solution available, so make sure you take a look at some of them to see which may be a good fit for your environment. We also listed some resources that you should check out so that you can learn even more and stay current on the ever-changing virtualization industry.

Additional Resources

A number of other resources are available to help you learn more about VMware so that you can continue gaining more knowledge and experience. These resources consist of some great VMware-related blogs and websites that you can follow and other great technical resources.

Blogs

Blogs are a great source for tips, news, experiences, and information from other VMware users, and there are dozens devoted to VMware. Listed here are some of the best ones that you should check out frequently or subscribe to their RSS feeds.

- Yellow Bricks by Duncan Epping
www.yellow-bricks.com/
- Scott Lowe's blog
<http://blog.scottlowe.org/>
- Mike DePetrillo's blog
<http://mikedatl.typepad.com/mikedvirtualization/>
- NTPro.nl by Eric Sloof
www.ntpro.nl/blog/
- Tech Target's Virtualization Pro blog
<http://itknowledgeexchange.techtarget.com/virtualization-pro/>
- Tech Target's Search Server Virtualization blog
<http://servervirtualization.blogs.techtarget.com/>
- VM/ETC by Richard Brambley
<http://vmetc.com/>
- RTFM Education by Mike Laverick
www.rtfm-ed.co.uk/
- Rational Survivability by Christofer Hoff
<http://rationalsecurity.typepad.com/blog/>
- Virtual Geek by Chad Sakac
http://virtualgeek.typepad.com/virtual_geek/
- VMware Virtualization Evangelist by Jason Boche
www.boche.net/blog/
- Planet VM by Tom Howarth
<http://planetvm.net/blog/>
- Gabe's Virtual World by Gabe Van Zanten
www.gabesvirtualworld.com/
- VMware Tips by Rick Scherer
<http://vmwaretips.com/wp/>
- Vroom! by VMware's Performance Team
<http://blogs.vmware.com/performance/>

VMware has a web page and RSS feed that aggregates all these blogs and more into a centralized page. Check out Planet V12n at www.vmware.com/vmtn/planet/v12n/ for all virtualization-related websites and Planet VMware at www.vmware.com/vmtn/planet/vmware/ for all VMware-related websites.

Websites

There are also some great websites that provide tips, news, and more that you should check out. Some of the best and most popular are listed here:

- VMware-land
<http://vmware-land.com/>
- Virtualization.info
www.virtualization.info/
- Tech Target's Search VMware
<http://searchvmware.techtarget.com/>
- Tech Target's Search Server Virtualization
<http://searchservervirtualization.techtarget.com/>
- Virtual Strategy Magazine
www.virtual-strategy.com/en/home
- Virtualization Admin
<http://virtualizationadmin.com/>
- Virtualization Review
<http://virtualizationreview.com/>
- X86 Virtualization
<http://x86virtualization.com/>
- VM Blog
<http://vmblog.com/>
- Virtualization
<http://virtualization.com/>
- Dabcc
www.dabcc.com/
- Virtualization Sys-con
<http://virtualization.sys-con.com/>

Many, many more good blogs and websites have not been listed here. For a full listing of all the most current blogs and websites, check out the Launchpad page on my website at

<http://vlp.vmware-land.com>. Also be sure to check out my Top 10 lists at http://vmware-land.com/Top_10_Lists.html, which cover the top ten things you should read on a wide variety of subjects.

More Stuff

For even more VMware-related resources on the web, check out these websites:

- www.vmworld.com

Here nonattendees of VMware's annual technical conference can view the hundreds of VMworld presentations from years past for free, except for the most recent year, for which you can purchase a subscription so that you can access it.

- www.vmware.com/events/webinars/

VMware presents many technical webinars. You can view recordings of old ones and see upcoming ones.

- <http://twitter.com>

Many virtualization professionals converse on Twitter daily. Check out the virtualization group at <http://twittgroups.com/group/virtualization> and follow some of the best minds in the business. You can also follow me at <http://twitter.com/ericsiebert> and check out the many virtualization folks I am following.

- <http://viops.vmware.com/home/index.jspa>

VMware's proven-practice website that contains many great documents from both VMware and their customers.

- <http://vmware.com/vmtn/resources/>

VMware frequently publishes many good white papers on a wide variety of topics.

- http://vmware.com/support/pubs/vi_pubs.html

VMware has fantastic documentation that is updated with each new release.

- <http://vmware.com/resources/communities/usergroup/events.html>

There are VMware user groups that meet periodically in most metropolitan areas. These meetings are hosted by VMware and are a great way to learn, meet other users, and share ideas and experiences.

Index

Symbols

- delta.vmdk file (VMs), 445
- flat.vmdk file (VMs), 445
- rdm.vmdk file (VMs), 445
- .NET Framework, vCenter Server, 88
- / partition (ESX), 100
- /boot partition (ESX), 100
- /home partition (ESX), 101
- /tmp partition (ESX), 101
- /var partition (ESX), 101
- /var/log partition (ESX), 100
- /var/log/messages log file, 411
- /var/log/secure log file, 411
- /var/log/vmkernel log file, 410
- /var/log/vmksummary log file, 410
- /var/log/vmkwarning log file, 410
- /var/log/vmware/aam log file, 411
- /var/log/vmware/esxcfg-boot.log log file, 411
- /var/log/vmware/esxcfg-firewall.log log file, 411
- /var/log/vmware/esxupdate.log log file, 411
- /var/log/vmware/hostd.log log file, 410
- /var/log/vmware/vpxa.log log file, 411
- /var/log/vmware/webAccess log file, 411
- 100% virtualized environments, 17
- 10K rpm hard drives, 56-57
- 15K rpm hard drives, 56-57
- 2GBSparse disks, 447

A

- About tab (VMware Tools), 300
- Account Configuration screen (ESX), 109
- accounts, ESX Service Console, 222-225
- Acronis True Image, P2V (physical to virtual) migration, 307
- Action tab, alarms, creating, 350-351
- AD (Active Directory)
 - ESX Service Console, authentication, 229-230
 - vCenter Server, security, 243-244
- AD (Active Directory) domain controllers, P2V (physical to virtual) migration, problems, 304
- Add Counters window (PerfMon), 11
- Adding a Host to vCenter Server screen (ESX), 112
- Administrator role (vCenter Server), 253
- administrators
 - network administrators, educating, 19-20
 - operating system administrators, educating, 23-24
 - security administrators, educating, 21-22
 - storage administrators, educating, 23
- Admission Controls settings, HA (High Availability), 141-142
- advanced configuration options (VMs), 289-291
- Advanced CPU resource settings (VMs), 295-296

- Advanced Memory resource settings (VMs), 296-297
- Advanced Option settings, HA (High Availability), 143-144
- Advanced, Boot Options (VMs), 290
- Advanced, CPUID Mask (VMs), 290
- Advanced, Fibre Channel NPIV (VMs), 291
- Advanced, General (VMs), 290
- Advanced, Paravirtualization (VMs), 291
- Advanced, Swapfile Location (VMs), 291
- Advanced, Virtualized MMU (VMs), 291
- affinity rules, DSR (Distributed Resource Scheduler), 151-152
- agent backups, DR sites, moving to, 71
- Alarm tab, VI Client, 348
- alarms, vCenter Server
 - configuring, 344-351
 - inheritance, 345
- Alarms privilege (vCenter Server), 249
- Alarms tab (VI Client), 124
- Altor Network's Virtual Firewall, 259
- Altor Network's Virtual Network Security Analyzer, 259
- Altor Network's Virtual Firewall, 462
- AMD CPUs, 53-54
- antivirus software, ESX Service Console, installing on, 230-232
- Apani EpiForce VM, 260
- application compatibility, virtualization, 17-19
- application support, virtualization, 16
- applications, vendor virtualization support
 - policies, 19
- authentication
 - ESX Service Console, configuring, 229-230
 - vCenter Server, troubleshooting, 425
- authentication log file, 411
- automatic mode, DSR (Distributed Resource Scheduler), 155
- automation levels, DSR (Distributed Resource Scheduler), 148-150

B

- backup scripts, 399-401
- backup software compatibility guide, 50
- backups, 399, 406-407
 - backup scripts, 399-401
 - ESX hosts, tape drives, 406
 - esXpress, 403-404
 - file-level backups, 404-405
 - image-level backups, 404-405
 - quiescing process, 404
 - traditional backup agents, 399-400
 - VBAs (virtual backup appliances), 403
 - VCB (Consolidated Backup), 43-45, 399-402
 - Veeam Backup, 403
 - VMs, utilities, 460-461
 - vRanger Pro, 403
- base filename-delta.vmdk file, snapshots, 436
- base filename-delta.vmsd file, snapshots, 436
- base filename-delta.vmsn file, snapshots, 436
- baselines, Update Manager
 - attaching, 387-389
 - creating, 385-388
- Beacon Probing Network Failover Detection, ESX hosts, 190
- Beaver, Steve, 434
- best practices
 - Converter (vCenter Server)
 - post-conversion, 327-328
 - preconversion, 325-326
 - running, 326
 - ESX
 - security, 236-239
 - server hardware security, 236
 - Service Console security, 236-238
 - vCenter Server
 - DSR (Distributed Resource Scheduler), 155-156
 - security, 256-257
 - VMs
 - creating, 301-302
 - templates, 280

- “Best Practices for Architecting VCB Enabled Solutions,” 407
 - “Best Practices for Patching VMware ESX/ESXi,” 397
 - BIOS, VMs, 441
 - blade servers, traditional servers, compared, 51-52
 - blogs, resources, 464-465
 - blue screen of death (BSoD), 415
 - blue VM/template folders (vCenter Server), 138
 - Boche, Jason, 434, 464
 - boot from SAN, 57-58
 - ESX, 99
 - boot processes, ESXi hosts, 380
 - Bootloader configuration screen (ESX), 107
 - Brambley, Richard, 464
 - BSoD (blue screen of death), 415
 - building virtual environments, 75
 - database servers, 77-86
 - ESX, 98-116
 - ESXi, 98-116
 - licensing servers, 97-98
 - preparation, 75-77
 - vCenter Server, 87-93
 - VI Client, 93-97
 - built-in firewalls
 - ESX Service Console, 232-236
 - ESXi, 29
 - bundles, 368
 - information, retrieving, 370
 - installing, esxupdate, 372
 - patches, scanning for, 369-370
 - test installations, esxupdate, 371
- ### C
- cable management, blade servers, 51
 - Capacity Planner, performance usage, measuring, 9-10
 - cat Linux command, 454
 - Catbird V-Security, 258, 462
 - cd Linux command, 454
 - CD-ROM devices, VMs, selecting, 276
 - CD/DVD-ROM drives, VMs, 442
 - configuration settings, 287
 - CDP (Cisco Discovery Protocol), ESX hosts, 194
 - centralized licensing, 69-70
 - configuring, 130-133
 - Check Point VPN 1-VE, 259
 - chmod Linux command, 454
 - chown Linux command, 454
 - Cline, Ken, 434
 - cloning
 - cold cloning
 - Converter Enterprise (vCenter Server), 321-324
 - P2V (physical to virtual) migration, 303-304
 - hot cloning
 - Converter Enterprise (vCenter Server), 320
 - P2V (physical to virtual) migration, 303-304
 - VMs, 281-286
 - cluster evaluation frequency, DSR (Distributed Resource Scheduler), 155
 - cluster settings, HA (High Availability), 143
 - Cluster Summary tab (vCenter Server), 153
 - clusters
 - Converter Starter (vCenter), choosing, 314
 - vCenter Server, creating, 139-140
 - cold cloning
 - Converter Enterprise (vCenter Server), 321-324
 - P2V (physical to virtual) migration, 303-304
 - commands
 - esxcfg-advcfg, 451
 - esxcfg-auth, 238, 451
 - esxcfg-boot, 451
 - esxcfg-dumppart, 451
 - esxcfg-firewall, 232-233, 451
 - esxcfg-info, 452
 - esxcfg-init, 452

- esxcfg-linuxnet, 452
- esxcfg-module, 452
- esxcfg-mpath, 452
- esxcfg-nas, 452
- esxcfg-nics, 228, 417, 452
- esxcfg-resgrp, 452
- esxcfg-route, 418, 453
- esxcfg-swiscsi, 453
- esxcfg-upgrade, 453
- esxcfg-vmhbadevs, 453
- esxcfg-vmknic, 418, 453
- esxcfg-vswif, 417, 453
- esxcfg-vswitch, 417, 453
- esxupdate, 367-372
- GET, 340
- kill, 429-430
- RCLI (Remote Command-Line Utility),
 - 449-450
- reboot, 417
- resxtp, 449
- Service Console, 451-454
 - Linux commands, 454-455
- service firewall stop, 236
- Set, 340
- svmotion, 449
- tail, 410
- useradd, 225
- vicfg-advcfg, 449
- vicfg-cfgbackup, 449
- vicfg-dns, 449
- vicfg-dumppart, 449
- vicfg-module, 450
- vicfg-mpath, 450
- vicfg-nas, 450
- vicfg-nics, 450
- vicfg-ntp, 450
- vicfg-rescan, 450
- vicfg-route, 450
- vicfg-snmp, 450
- vicfg-syslog, 450
- vicfg-user, 450
- vicfg-vmhbadevs, 450
- vicfg-vmknic, 450
- vicfg-vswitch, 450
- vifs, 450
- vihostupdate-, 450
- vm-support, 412-413
- vmkfstools, 450, 453
- vmware-cmd, 427-428, 450
- vmware-support, 428-429
- vmware-vmupgrade.exe, 396-397
- committing snapshots, 437-439
- compatibility lists, hardware, 50-51
- Compliance Checker, 259
- configuration
 - database servers, 77-86
 - ESX hosts, 45-48, 176
 - DNS, 216
 - networking, 176-201
 - NTP time synchronization, 213-215
 - routing, 216
 - Service Console memory, 215-216
 - storage, 201-213
 - VM (virtual machine)
 - startup/shutdown, 216-217
 - ESX Service Console
 - authentication, 229-230
 - sudo, 225-228
 - ESXi hosts, 176
 - networking, 176-201
 - SNMP, 343-344
 - storage, 201-213
 - FC (Fibre Channel) storage, 206-209
 - iSCSI storage, 209-212
 - licensing
 - centralized licensing, 130-133
 - ESXi, 133-134
 - host-based licensing, 130-131
 - single-CPU licensing, 135
 - stand-alone licensing, 130-131
 - NFS storage, 212-213
 - Oracle, ODBC connections, 85-86

- SQL Server, ODBC connections, 82-85
- vCenter Server, 45-48, 136-175
 - alarms, 344-345, 347-351
 - DPM (Distributed Power Management), 156-158
 - DSR (Distributed Resource Scheduler), 148-156
 - email alerts, 345-347
 - EVC (Enhanced VMotion Compatibility), 163-166
 - HA (High Availability), 141-148
 - limits, 172
 - performance monitoring, 357-360
 - reservations, 172
 - resource pools, 173-174
 - resources, 166-175
 - roles, 253, 255-256
 - shares, 172
 - VMotion, 159-163
- VMs, 286
 - advanced configuration options, 289-291
 - hardware, 286-288
 - resources, 292-297
- Configuration tab (VI Client), 123
- Configuresoft Compliance Checker, 259, 462
- Configuresoft ECM (Enterprise Configuration Manager), 259
- Confirmation screen (ESXi), 115
- Connection Settings screen (VMotion), 160
- connectivity, ports, 219-222
- connectivity problems, vCenter Server, troubleshooting, 425
- console screen (ESX), 112
- Console tab (VI Client), 124
- consoles, log files, viewing, 419
- "Consolidated Backup in VMware Infrastructure 3," 407
- Consolidated Backup User role (vCenter Server), 255
- Consolidated Backup (VMware), 43-45, 399-402
- contents.xml files, 368
- contents.xml.sig files, 368
- Converter (vCenter Server)
 - destination formats, 308-309
 - operating system support, 308
 - P2V (physical to virtual) migration, 305-329
 - source import formats, 308
 - VMs, cloning, 282-283
- Converter Enterprise (vCenter Server), 88, 126, 317
 - best practices, preconversion, 325-328
 - destination formats, 308-309
 - hot cloning, 320-324
 - installing, 318-320
 - operating system support, 308
 - P2V (physical to virtual) migration, 317-329
 - source import formats, 308
- Converter Starter (vCenter Server)
 - best practices
 - post-conversion, 327-328
 - preconversion, 325-326
 - running, 326
 - clusters, choosing, 314
 - destination datascores, choosing, 315
 - destination formats, 308-309
 - destination host servers, choosing, 314
 - destination servers, setting, 313
 - destination types, selecting, 312
 - guest customization information, setting, 316
 - hard disks, selecting for source data, 312
 - installing, 309-310
 - login information, setting, 313
 - operation system support, 308
 - P2V (physical to virtual) migration, 309-317
 - running, 310-317
 - source import formats, 308
 - source login information, setting, 311
 - source types, selecting, 310
 - VMs, naming, 314
 - vNIC information, setting, 316

- cp Linux command, 454
- CPU Identification Utility, 459
- CPU resource settings (VMs), 292-293
- CPU usage, measuring, 5
- CPUs (central processing units)
 - AMD, 53-54
 - ESX hosts, 47-48
 - performance monitoring, 355
 - Intel, 53-54
 - multicore CPUs, 52-53
 - single-vCPU VMs, 48
 - vCPUs, VMs, 267
 - vendors, choosing, 53-54
 - VMs, 441
 - assigning, 286
 - performance monitoring, 354

D

- Dabcc.com, 465
- das.allowNetwork setting (HA), 147
- das.allowVmotionNetworks setting (HA), 147
- das.bypassNetCompatCheck setting (HA), 148
- das.defaultfailoverhost setting (HA), 146
- das.failedetectioninterval setting (HA), 147
- das.failedetectiontime setting (HA), 146
- das.FailureInterval. 30 setting (HA), 144
- das.isolationaddress setting (HA), 146
- das.isolationShutdownTimeout setting (HA), 146
- das.maxFailures. 3 setting (HA), 144
- das.maxFailureWindow. 3600 setting (HA), 144
- das.minUptime. 120 setting (HA), 144
- das.poweroffonisolation setting (HA), 146
- das.usedefaultisolationaddress setting (HA), 146
- das.vmCpuMinMHz setting (HA), 147
- das.vmFailoverEnabled Advanced Option setting (HA), 144
- das.vmMemoryMinMB setting (HA), 147
- database, vCenter Server, troubleshooting, 425-426

- Database Selection screen (vCenter Server), 90
- database servers
 - P2V (physical to virtual) migration, problems, 305
 - vCenter Server, choosing and configuring, 77-86
- databases, vCenter Server, 77
 - ODBC connections, 82-86
 - Oracle, 80-81
 - SQL Server, 80-82
 - statistics collection settings, 77-79
- Datacenter Administrator role (vCenter Server), 254
- Datacenter privilege (vCenter Server), 246
- datacenter/cluster/host folders (vCenter Server), 138
- datacenters
 - remote datacenters, 70
 - vCenter Server, creating, 137
- Datacenters tab (VI Client), 122
- datascores, Converter Starter (vCenter), choosing, 315
- datastore browser (VI Client), VMs, cloning, 284-286
- Datastore privilege (vCenter Server), 246
- datastores, VMs, choosing, 265-266
- Datastores tab (VI Client), 125
- Datastores view (VI Client), 121
- dd command (Service Console), ISO files, creating, 272
- deeply nested snapshots, committing, 439
- default networks, ESX hosts, 180
- deleting snapshots, 435
- Dell OpenManage, ESX hosts, hardware monitoring, 336
- DePetrillo, Mike, 464
- depots, applicable bundles, scanning for, 369-370
- descriptor.xml files, 368
- destination datascores, Converter Starter (vCenter), choosing, 315

- destination host servers, Converter Starter (vCenter), choosing, 314
- destination servers, Converter Starter (vCenter), selecting, 313
- destination types, Converter Starter (vCenter), selecting, 312
- developers, virtualization, educating, 21
- Devices tab (VMware Tools), 300
- df Linux command, 454
- disaster recovery (DR), ESX, 71-72
- Disk resource settings (VMs), 294
- Disk Selection screen (ESXi), 114
- disk space, vCenter Server, troubleshooting, 426
- disk statistics
 - hosts, performance monitoring, 356
 - VMs, performance monitoring, 354
- disk usage
 - ESX hosts, 49
 - measuring, 7
- disks, snapshots, excluding from, 439
- displaying physical NICs, ESX hosts, 180
- Distributed Power Management (DPM). *See* DPM (Distributed Power Management)
- Distributed Resource Scheduler (DRS). *See* DRS (Distributed Resource Scheduler)
- DMZ (demilitarized zone), networks, ESX hosts, 178-179
- DNS (domain name system), ESX hosts, configuring, 216
- downloading
 - license files, 127-130
 - patches, esxupdate, 369
- DPM (Distributed Power Management), 26, 40
 - vCenter Server, configuring, 156-158
- DRS Recommendations tab (VI Client), 123
- DRS (Distributed Resource Scheduler), 25-26, 38-40
 - four-star migration threshold, 150
 - Fully Automated level, 149

- vCenter Server
 - affinity rules, 151-152
 - automatic mode, 155
 - automation levels, 148-150
 - best practices, 155-156
 - cluster evaluation frequency, 155
 - configuring, 148-156
 - DSR Recommendations tab, 154
 - migration threshold, 150-151
 - monitoring, 152-154
 - VMotion, 155
- du Linux command, 455
- dual-core processors, 52

E

- e1000 virtual network adapter, 66
- eager-zeroed thick disks, 446
- ECM (Enterprise Configuration Manager), 259
- editions, ESX, choosing, 25-48
- eG VM Monitor, 363, 460
- email alerts, vCenter Server, configuring, 345-347
- Embotics v-Scout, 462
- employees, virtualization, education, 19-24
- enhanced vmxnet virtual network adapter, 66
- enterprise monitoring systems, performance usage, measuring, 13
- environments
 - assessing, 3-15
 - documenting, 4
 - performance usage, measuring, 4-15
 - virtualization, 75-116
 - 100% virtualized environments, 17
 - application compatibility, 17-19
 - educating staff, 19-24
- Epping, Duncan, 464
- "escaping the cave," 21
- ESH hosts
 - patching, 365-366
 - esxupdate, 367-372
 - methods, 366-367
 - updating, 365-366

- EST (external switch tagging), 177
- ESX
 - Account Configuration screen, 109
 - Adding a Host to vCenter Server screen, 112
 - Bootloader configuration screen, 107
 - built-in firewall, 29
 - configuration, 45-48
 - console screen, 112
 - downloading, 75
 - DPM (Distributed Power Management), 26
 - DR (disaster recovery), 71-72
 - DRS (Distributed Resource Scheduler), 25-26, 38-40
 - editions, choosing, 25-48
 - ESXi, compared, 27, 30
 - firewall, security best practices, 239
 - HA (High Availability), 25
 - hardware
 - blade servers, 51-52
 - compatibility lists, 50-51
 - CPUs, 53-54
 - multicore CPUs, 52-53
 - network adapters, 55
 - storage adapters, 55
 - storage options, 56-62
 - traditional servers, 51-52
 - unsupported hardware, 56
 - hardware monitoring, 29
 - installation, 98-116
 - boot from SAN, 99
 - partitions, 99-101
 - preparation, 98-99
 - Installation Complete screen, 111
 - installation summary screen, 110
 - Installing Packages screen, 111
 - licensing, 69-70
 - Media Test screen, 103
 - media tests, 103
 - Network Configuration screen, 108
 - networking, security best practices, 239
 - Partition Disks screen, 105-106
 - Partition Warning screen, 104
 - Partitioning Options screen, 104
 - partitions, 103-106
 - patches, 29
 - physical server hardware, security, 236
 - RCLI, 28
 - remote datacenters, 70
 - scriptable installations, 28
 - security, best practices, 236-239
 - Service Console, 28
 - antivirus software, 230-232
 - authentication, 229-230
 - built-in firewall, 232-236
 - root user account, 222-225
 - securing, 222-240
 - security best practices, 236-238
 - sudo, 225-228
 - startup screen, 27
 - Storage VMotion, 25
 - thick client access, 28
 - Time Zone Selection screen, 109
 - updates, 29
 - web client access, 28
 - welcome screen, 102
- ESX Enterprise licenses, 133
- ESX HealthCheck, 463
- ESX hosts
 - backups, 399
 - backup scripts, 400-401
 - tape drives, 406
 - CDP (Cisco Discovery Protocol), 194
 - configuring, 176
 - DNS, 216
 - networking, 176-201
 - NTP time synchronization, 213-215
 - routing, 216
 - Service Console memory, 215-216
 - storage, 201-213
 - VM (virtual machine)
 - startup/shutdown, 216-217
 - CPUs, 47-48
 - disk usage, 49
 - esxcfg-nics -l command, 182-184

- firewalls, SNMP traffic, 341
- HA (High Availability), 37-38
- host management networks, configuring, 179-185
- host servers, 48
- licensing, 131
- memory, 47, 49
- memory usage, 166-168
- monitoring, 335
 - hardware, 336-340
 - performance, 352-353
 - SNMP (Simple Network Management Protocol), 340-344
 - vCenter Server, 344-351
- network hints, 200-201
- networking, 47
- networks
 - default network, 180
 - design considerations, 176-178
 - DMZ (demilitarized zone), 178-179
- NICs, adding, 197, 199
- number, determining, 48-49
- patching, 365
 - Update Manager, 381-393
- performance monitoring
 - configuring, 357-360
 - CPU statistics, 355
 - disk statistics, 356
 - eG VM Monitor, 363
 - ESXTOP command-line utility, 360-362
 - memory statistics, 356-357
 - network statistics, 357
 - Veeam Management suite, 362
 - vFoglight, 362
- physical NICs, displaying, 180
- RCLI (Remote Command-Line Utility), 449-450
- root passwords
 - changing, 217-218
 - resetting, 218
- SCP file-transfer utilities, 458
- scripting
 - Perl, 456
 - PowerShell, 456-457
- Service Console, 451
 - backing up configuration files, 405
 - backup agents, 405
 - commands, 451-454
- service console, configuring, 179-184
- Service Console, Linux commands, 454-455
- split-brain condition, 145
- SSH console utilities, 458
- storage, 46
- Storage VMotion, 36-37
- troubleshooting, 409
 - determining versions, 414-415
 - esxcfg commands, 417
 - hostd service, 418
 - log files, 410-414
 - networking problems, 417-418
 - PSoD (purple screen of death), 415-416
 - Service Console problems, 416-417
 - vxpa service, 418-419
- Update Manager, 42-43
- updating, 372
 - ISO file, 373-374
- vCenter Server, 30-33
 - adding to, 140-141
- VI Client, connecting to directly, 125-126
- VM NICs, mapping to pNICs, 194
- VMkernel networks, configuring, 185-186
- VMotion, 33-35
- vNICs, changing MAC addresses, 196-197
- vSwitches
 - configuring, 186-193
 - internal-only vSwitches, 195
 - load-balancing policies, 189-190
 - Network Failover Detection, 190
- ESX servers, Ramcheck memory test utility, 416

- ESX version 3.5 Update 3, 2
- esxcfg commands
 - ESX hosts, troubleshooting, 417
 - ESXi hosts, troubleshooting, 422
- esxcfg-advcfg command, 451
- esxcfg-auth command, 238, 451
- esxcfg-boot command, 451
- esxcfg-dumppart command, 451
- esxcfg-firewall command, 232-233, 451
- esxcfg-info command, 452
- esxcfg-init command, 452
- esxcfg-linuxnet command, 452
- esxcfg-module command, 452
- esxcfg-mpath command, 452
- esxcfg-nas command, 452
- esxcfg-nics -l command, 182-184
- esxcfg-nics command, 228, 417, 452
- esxcfg-resgrp command, 452
- esxcfg-route command, 418, 453
- esxcfg-swiscsi command, 453
- esxcfg-upgrade command, 453
- esxcfg-vmhbadevs command, 453
- esxcfg-vmknic command, 418, 453
- esxcfg-vswif command, 417, 453
- esxcfg-vswitch command, 417, 453
- ESXi, 26
 - built-in firewall, 29
 - Confirmation screen, 115
 - Disk Selection screen, 114
 - downloading, 75
 - ESX, compared, 27, 30
 - EULA screen, 114
 - hardware monitoring, 29
 - installation, 98-116
 - boot from SAN, 99
 - partitions, 99-101
 - preparation, 98-99
 - Installation Complete screen, 116
 - Installing screen, 115
 - licensing, configuring, 133-134
 - login screen, 116
 - management network, configuring, 185
 - patches, 29
 - RCLI, 28
 - scriptable installations, 28
 - Service Console, 28
 - security, 239-240
 - startup screen, 27
 - thick client access, 28
 - updates, 29
 - web client access, 28
 - welcome screen, 113
- ESXi hosts
 - boot process, 380
 - CDP (Cisco Discovery Protocol), 194
 - configuring, 176
 - networking, 176-201
 - storage, 201-213
 - connecting to directly, VI Client, 125-126
 - EULA screen, 114
 - host management networks, configuring, 179-185
 - monitoring, 335
 - hardware, 340
 - performance, 352-353
 - SNMP (Simple Network Management Protocol), 340-344
 - vCenter Server, 344-351
 - networks
 - design considerations, 176-178
 - DMZ (demilitarized zone), 178-179
 - patching, 365-366, 374-379
 - methods, 366-367
 - Update Manager, 381-393
 - performance monitoring, configuring, 357-360
 - RCLI (Remote Command-Line Utility), 449-450
 - rolling back, 379-381
 - scripting
 - Perl, 456
 - PowerShell, 456-457
 - SNMP, configuring, 343-344
 - troubleshooting, 419
 - determining versions, 420
 - esxcfg commands, 422

- log files, 419-420
- PSoD (purple screen of death), 422
- Tech Support Mode, 420-423
- updating, 365-366, 374-379
 - Infrastructure, 374-377
 - vihostupdate, 377-379
- vCenter Server, adding to, 140-141
- VM NICs, mapping to pNICs, 194
- VMkernel networks, configuring, 185-186
- vNICs, changing MAC addresses, 196-197
- vSwitches
 - configuring, 186-193
 - internal-only vSwitches, 195
 - load-balancing policies, 189-190
 - Network Failover Detection, 190
- esXpress, 403-404
- ESXTOP command-line utility, performance monitoring, 360-362
- esxupdate, 367
 - activity logs, retrieving, 372
 - applicable bundles, scanning for, 369-370
 - bundle information, retrieving, 370
 - bundles
 - installing, 372
 - test installations, 371
 - disk space, verifying, 370-371
 - modes, 367
 - patch repositories, setting up, 368
 - patches, downloading, 369
- EULA screen (ESXi), 114
- EVC (Enhanced VMotion Compatibility), vCenter Server, configuring, 163-166
- Events tab (VI Client), 124
- Extension privilege (vCenter Server), 250
- external switch tagging (EST mode), vSwitches, 68
- external switch tagging (EST), 177

F

- FastSCP, log files, viewing, 410
- FC (Fibre Channel), 58-59
 - storage, configuring, 206-209

- file locations, snapshots, changing, 439-440
- file types, VMs, 442-445
- file-level backups, 404-405
- files, bundles, 368
- files servers, P2V (physical to virtual)
 - migration, problems, 305
- find Linux command, 454
- firewalls
 - ESX
 - security best practices, 239
 - SNMP traffic, 341
 - ESX Server Console, built-in firewall, 232-236
 - Virtual Firewall, 259
 - X-M0n0wall virtual firewall appliance, 260
- five-star migration threshold (DSR), 150
- Flexible network adapter, 65
- floppy disks, VMs, configuration settings, 286
- floppy drives, VMs, 275
- Folder privilege (vCenter Server), 246
- folders, creating
 - vCenter Server, 137-138
 - VI Client, 121
- four-star migration threshold (DSR), 150
- Fully Automated level (DSR), 149

G

- Gabe's Virtual World blog, 464
- General Options (VMs), 289
- Generate Update Manager log bundle, vCenter Server, 92
- Generate VirtualCenter1 log bundle, vCenter Server, 92
- GET command, net-snmp agent, interaction, 340
- Getting Started tab (VI Client), 122
- Global privilege (vCenter Server), 246
- growth
 - leaving room for, 50
 - snapshots, 436-437
- guest customization information, Converter Starter (vCenter), 316

guest operating systems, VMs, choosing,
265-266

Guided Consolidation (vCenter Server), P2V
(physical to virtual) migration, 329-334

H

HA (High Availability), 25, 37-38

log file, 411

vCenter Server

Admission Control settings, 141-142

Advanced Option settings, 143-144

cluster settings, 143

configuring, 141-148

split-brain condition, 145

Virtual Machine Monitoring settings,
143-144

Haletky, Edward, 434

hard disks

10K rpm hard drives, 56-57

15K rpm hard drives, 56-57

Converter Starter (vCenter), selecting for
source data, 312

VMs, configuration settings, 288

hardware

blade servers, 51-52

compatibility lists, 50-51

CPUs, vendors, 53-54

ESX hosts, monitoring, 336-340

ESXi hosts, monitoring, 340

growth, leaving room for, 50

multicore CPUs, 52-53

network adapters, 55

requirements, VI Client, 93

storage, 56

10K rpm hard drives, 56-57

15K rpm hard drives, 56-57

boot from SAN, 57-58

FC (Fibre Channel), 58-59

iSCSI, 59-61

local storage, 58

mixing types, 62

NAS/NFS, 61-62

storage adapters, 55

traditional servers, 51-52

unsupported hardware, 56

VMs

configuration settings, 286-288

virtual hardware, 441-442

hardware initiator, iSCSI, 61

hardware management agents, HP agents,
336, 340

installing, 336-338

System Management home page, 339

System Management login page, 339

hardware monitoring, ESXi, 29

header files, 368

High Availability (HA). *See* HA (High
Availability)

high-resource utilization servers,
virtualization, 15

Hoff, Christopher, 464

Host, 247

host management networks, ESX hosts,
configuring, 179-185

Host privilege (vCenter Server), 247

host servers, ESX, 48

host-based licensing, configuring, 130-131

hostd service, ESX hosts, troubleshooting, 418

hosts. *See also* ESX hosts; ESXi hosts

resource pools, 170-172

vCenter Server, troubleshooting, 426

Hosts and Clusters view (VI Client), 120

Hosts tab (VI Client), 122

hot cloning

Converter Enterprise (vCenter
Server), 320

P2V (physical to virtual) migration,
303-304

Howarth, Tom, 434, 464

HP hardware management agents, 336

ESX hosts, 340

installing on, 336-338

System Management home page, 339

System Management login page, 339

HP Systems Insight Manager, ESX hosts,
hardware monitoring, 336

Hyper-9, 462

hypervisors, vendors, 1

I

I/O compatibility guide, 50

IDE controllers, VMs, 442

image backups, DR sites, moving to, 72

image-level backups, 404-405

Imgburn, ISO files, creating, 272

infrastructure

 assessing, 3-15

 performance usage, measuring, 4-15

Infrastructure Update, ESXi hosts, 374-377

Inheritance, vCenter Server alarms, 345

Inspection mode (esxupdate), 367

installation

 bundles, test installations, 371

 Converter Enterprise (vCenter Server),
 318-320

 Converter Starter (vCenter Server),
 309-310

 ESX, 98-116

 boot from SAN, 99

 partitions, 99-101

 preparation, 98-99

 ESXi, 98-116

 boot from SAN, 99

 partitions, 99-101

 preparation, 98-99

 HP hardware management agents, ESX

 hosts, 336-339

 licensing servers, 97-98

 patches, esxupdate, 372

 RCLI (Remote Command-Line Interface)
 utility, 343

 Update Manager, 382-385

 vCenter Server, 87-93

 Database Selection screen, 90

 Installation Type screen, 89

 Licensing Server screen, 90

 physical servers, 87-88

 Server Authorization screen, 91

 virtual machines, 87-88

 VI Client, 93-97

 VMware Tools, VMs, 297-301

Installation Complete screen (ESX), 111

Installation Complete screen (ESXi), 116

Installation documentation, downloading, 76

installation files, MD5 checksums, 76-77

installation summary screen (ESX), 110

Installation Type screen, vCenter Server, 89

Installing Packages screen (ESX), 111

Installing screen (ESXi), 115

Intel CPUs, 53-54

internal-only vSwitches, ESX hosts, 195

intervals, statistics collection settings,
 vCenter Server, 79

iSCSI storage, 59-61

 configuring, 209-212

 hardware initiator, 61

 software initiator, 60

ISO files

 creating, applications, 272

 ESX hosts, updating, 373-374

 virtual CD/DVD drives, selecting for, 274

 VMs, 272-277

ISO Recorder, ISO files, creating, 272

IT infrastructure

 assessing, 3-15

 performance usage, measuring, 4-15

IT management, virtualization, educating,
 22-23

ITQ VLAN and Portgroup Manager, 459

J-K

kill command, VM power-state problems,
 troubleshooting, 429-430

knowledge base website (VMware), 431-432

known problems, virtualization, 4

KS QuickConfig, 458

L

large snapshots, committing, 439

Laverick, Mike, 464

Layout, VI Client, 120-125

lazy-zeroed thick disks, 446

LC ISO Creator, ISO files, creating, 272

- Lefthand Virtual SAN Appliance, 461
- levels, statistics collection settings, vCenter Server, 78-79
- license files, obtaining, 127-130
- license server, managing, 135-136
- License Server (vCenter Server), 88
 - settings, 132
- licenses, downloading, 75
- licensing, 127
 - centralized licensing, configuring, 130, 132-133
 - ESX Enterprise licenses, 133
 - ESXi licensing, configuring, 133-134
 - host-based licensing, configuring, 130-131
 - license files, obtaining, 127-130
 - license server, managing, 135-136
 - Licensing Sources window, 131
 - server ports, 135
 - single-CPU licensing, 135
 - stand-alone licensing, configuring, 130-131
 - VMware, 69-70
- Licensing Server screen (vCenter Server), 90
- Licensing Server Tools application (LMTools), 136
- licensing servers, installation, 97-98
- Licensing Sources window, 131
- limits, vCenter Server, 170
 - configuring, 172
- Link Status Only (default) Network Failover Detection, ESX hosts, 190
- Linux tail command, log files viewing, 410
- Linux-based net-snmp agent, 340
- LMTools, Licensing Server Tools application, 136
- load-balancing policies, vSwitches, ESX hosts, 189-190
- local storage, 58
 - configuring, 205-206
- locating VM snapshots, 440-441
- locations, VMs, choosing, 265
- LOG file type (VMs), 445

- log files
 - ESX hosts, troubleshooting, 410-414
 - ESXi hosts, troubleshooting, 419-420
 - vCenter Server, troubleshooting, 423-424
 - VMs, troubleshooting, 427
- Log Settings window (PerfMon), 12
- login information, Converter Starter (vCenter), setting, 313
- login screen (ESXi), 116
- ls Linux command, 454
- LUN/VMFS volumes, VMs, optimal number, 302
- LUNs, size considerations, 201-203

M

- MAC addresses, vNICs, changing, 196-197
- machine cloning, DR sites, moving to, 72
- Magic ISO Maker, ISO files, creating, 272
- maintenance releases, vCenter Server, 393-394
- management, virtualization, educating, 22-23
- management networks, ESXi, configuring, 185
- "Managing Patches and Updates for Hosts and VMs," 398
- Manual automation level (DSR), 148
- Maps tab (VI Client), 124
- MCS StorageView, 462
- MD5 checksums, 76-77
- measuring performance usage, 4-15
 - Capacity Planner, 9-10
 - enterprise monitoring systems, 13
 - PerfMon (performance monitor), 10-13
- Media Test screen (ESX), 103
- media tests, ESX, 103
- memory
 - ESX hosts, 47-49
 - Service Console, configuring, 215-216
 - VMs, 441
 - assigning, 263, 267-268, 286
- memory balloon drivers, hosts, 167-168
- memory overcommitment, hosts, 166-167
- Memory resource settings (VMs), 293-294

- memory statistics
 - hosts, performance monitoring, 356-357
 - VMs, performance monitoring, 354-355
 - memory usage
 - measuring, 6
 - vCenter Server, virtual hosts, 166-168
 - Microsoft .NET Framework, vCenter Server, 88
 - Microsoft SQL Server 2005 Express, vCenter Server, 88
 - migration, P2V (physical to virtual)
 - migration, 263, 302
 - choosing, 304-305
 - cold cloning, 303-304
 - hot cloning, 303-304
 - Platespin Migrate, 306
 - vCenter Server Converter, 305-329
 - vCenter Server Guided Consolidation, 329-334
 - Vizioncore vConverter, 307
 - migration threshold, DSR (Distributed Resource Scheduler), 150-151
 - Mishchenko, Dave, 434
 - mixing storage types, 62
 - mkdir Linux command, 454
 - modes, esxupdate, 367
 - monitoring
 - DSR (Distributed Resource Scheduler), 152-154
 - ESX hosts, 335
 - hardware, 336-340
 - performance, 352-357
 - SNMP (Simple Network Management Protocol), 340-344
 - vCenter Server, 344-351
 - ESXi hosts, hardware, 340
 - vCenter Server, 351-352
 - VMs
 - performance, 352-355
 - utilities, 459-460
 - motherboards, VMs, 441
 - multicore CPUs, 52-53
 - AMD CPUs, 53-54
 - Intel processors, 53-54
 - vendors, choosing, 53-54
 - mv Linux command, 454
- ## N
- Nagios, 460
 - naming
 - VMFS volumes, 206
 - VMs, 265
 - nano editor, snmp.conf file, editing, 342
 - nano Linux command, 454
 - NAS (network attached storage), 61
 - NAS/NFS storage, 61-62
 - Nero, ISO files, creating, 272
 - Nested Page Technology (NPT), 53
 - net-snmp agent, 340
 - Network Access screen (VMotion), 160
 - network adapters, 55
 - e1000, 66
 - Flexible, 65
 - virtual network adapters, 65
 - VMs, configuration settings, 287
 - network administrators, virtualization, educating, 19-20
 - network attached storage (NAS), 61
 - Network Configuration screen (ESX), 108
 - network controllers, VMs, 442
 - Network Failover Detection, ESX hosts, vSwitches, 190
 - Network File System (NFS) protocol, 61
 - network hints, 200-201
 - Network privilege (vCenter Server), 246
 - network statistics
 - hosts, performance monitoring, 357
 - VMs, performance monitoring, 355
 - network traffic
 - tagging, 176-177
 - vSwitches, routing, 199-200
 - network usage, measuring, 8-9

- networking
 - ESX, security best practices, 239
 - ESX hosts, 47
 - configuring, 176-201
 - troubleshooting, 417-418
 - virtual networking, 63
 - pNICs (physical NICs), 63-65
 - virtual switches, 66-69
 - vNICs (virtual NICs), 65-66
 - networks
 - ESX hosts
 - default network, 180
 - design considerations, 176-178
 - DMZ (demilitarized zone), 178-179
 - network ports, 219, 221-222
 - Networks view (VI Client), 121
 - New VM Wizard, 264-271
 - NFS (Network File System) protocol, 61
 - NFS storage, configuring, 212-213
 - NICs (network interface cards)
 - ESX hosts, adding, 197-199
 - virtual NICs, VMs, 267-269
 - No Access preconfigured role (vCenter Server), 253
 - nonroot user account, creating, ESX Service Console, 223
 - Norton Ghost, P2V (physical to virtual) migration, 307
 - NPT (Nested Page Technology), 53
 - NTP time synchronization, ESX hosts, configuring, 213-215
 - NTPro.nl blog, 464
 - NVRAM file type (VMs), 442
- O**
- objects, VI Client, 122-125
 - observed IP ranges, 200-201
 - ODBC connections
 - Oracle, configuring, 85-86
 - SQL Server, configuring, 82-85
 - vCenter Server, 82-86
 - one-star migration threshold (DSR), 150
 - Openfiler, 461
 - OpenManage, ESX hosts, hardware monitoring, 336
 - operating system administrators, virtualization, educating, 23-24
 - Options tab (VMware Tools), 300
 - Oracle
 - ODBC connections, configuring, 85-86
 - vCenter Server, 80-81
- P**
- P2V (physical to virtual) migration, 263, 302
 - choosing, 304-305
 - cold cloning, 303-304
 - hot cloning, 303-304
 - Platespin Migrate, 306
 - vCenter Server Converter, 305-329
 - best practices, 325-328
 - Converter Enterprise, 317-324
 - Converter Starter, 309-317
 - destination formats, 308-309
 - operating system support, 308
 - source import formats, 308
 - vCenter Server Guided Consolidation, 329-334
 - Vizioncore vConverter, 307
 - parallel ports, VMs, configuration settings, 288
 - partial memory reservations, VMs, creating for, 302
 - Partially Automated level (DSR), 149
 - Partition Disks screen (ESX), 105-106
 - Partition Warning screen (ESX), 104
 - Partitioning Options screen (ESX), 104
 - partitions, ESX, 99-106
 - passwords, root passwords, ESX hosts, 217-218
 - "Patch Management for ESX Server 3.5," 397
 - patch repositories, setting up, esxupdate, 368

- patches, 397
 - downloading, 369
 - ESX hosts, 365-366
 - esxupdate, 367-372
 - methods, 366-367
 - Update Manager, 381-393
 - ESXi, 29
 - ESXi hosts, 374-379
 - Update Manager, 381-393
- patching vCenter Server, 365
- PCNet32 virtual network adapter, 65
- PerfMon (performance monitor)
 - Add Counters window, 11
 - Log Settings window, 12
 - performance usage, measuring, 10-13
- performance monitoring, 352-353
 - eG VM Monitor, 363
 - ESXTOP command-line utility, 360-362
- hosts
 - CPU statistics, 355
 - disk statistics, 356
 - memory statistics, 356-357
 - network statistics, 357
- PerfMon (performance monitor), 10-13
- vCenter Server, configuring, 357-360
- Veeam Management suite, 362
- vFoglight, 362
- VMs, 354
 - CPU statistics, 354
 - disk statistics, 354
 - memory statistics, 354-355
 - network statistics, 355
- Performance privilege (vCenter Server), 250
- Performance tab (VI Client), 123
- performance usage
 - measuring, 4-15
 - Capacity Planner, 9-10
 - enterprise monitoring systems, 13
 - PerfMon (performance monitor), 10-13
 - statistics, analyzing, 13-15
- Perl scripts, 456
 - running snapshots, checking for, 441
- permissions, vCenter Server, 245-246, 250-252
- Permissions privilege (vCenter Server), 250
- Permissions tab (VI Client), 124
- PHD esXpress, 460
- PHD Technologies esXpress, 403-404
- physical machine cloning, DR sites, moving to, 72
- physical NICs, ESX hosts, displaying, 180
- physical NICs (pNICs), 63-65
- physical server hardware, ESX, security, 236
- physical servers
 - P2V (physical to virtual) migration, 302
 - choosing, 304-305
 - cold cloning, 303-304
 - hot cloning, 303-304
 - Platespin Migrate, 306
 - vCenter Server Converter, 305-329
 - vCenter Server Guided
 - Consolidation, 329-334
 - Vizioncore vConverter, 307
 - vCenter Server, installation, 87-88
 - virtualization, compared, 1
- physical to virtual (P2V) migration. *See* P2V (physical to virtual) migration
- pilot projects, virtualization, 20
- Planet VM blog, 464
- Platespin Migrate, P2V (physical to virtual) migration, 306
- plug-ins, VI Client, 126
- pNICs (physical NICs), 63-65
 - VM NICs, mapping from, 194
- ports, network connectivity, 219-222
- post-conversion best practices, Converter (vCenter Server), 327-328
- Power Management (VMs), 290
- power-state problems (VMs), troubleshooting, 427-430
- PowerShell scripts, 456-457
 - running snapshots, checking for, 440
- preconversion best practices, Converter (vCenter Server), 325-326
- privileges, vCenter Server, 245-246, 250-252

- PSoD (purple screen of death)
 - ESX hosts, troubleshooting, 415-416
 - ESXi hosts, troubleshooting, 422
- purple screen of death (PSoD), ESX hosts, troubleshooting, 415-416
- Putty SSH console utility, 458
- pwd Linux command, 454
- Q–R**
- quiescing process (backups), 404
- rack density, blade servers, 51
- Ramcheck memory test utility, 416
- Rapid Virtualization Indexing (RVI), 53
- Rational Survivability blog, 464
- raw disks, 446
- RCLI (Remote Command-Line Interface)
 - utility, 449-450
 - commands, 449-450
 - ESXi, 28
 - installing, 343
 - vihostupdate, ESXi hosts, 377-379
- RDMs (raw device mappings), 446
 - VMs, choosing, 270
- RDP (raw device mapping), VMFS volumes, compared, 203-205
- Read-only role (vCenter Server), 253
- reboot command, 417
- recommendations, DSR (Distributed Resource Scheduler), 154
- Reeh, Oliver, 434
- Reflex Systems Virtualization Management Center, 257, 462
- Release Notes, downloading, 76
- remediating hosts, Update Manager, 390-393
- Remote Command-Line Interface (RCLI)
 - utility. *See* RCLI (Remote Command-Line Interface) utility
- remote datacenters, 70
- Reporting tab, alarms, creating, 349-350
- reporting utilities, 462-463
- requirements, VI Client, 93
- reservations, vCenter Server, 169-170
 - configuring, 172
- resetting root passwords, ESX hosts, 218
- Resource Allocation tab (VI Client), 123
- Resource Management Guide, 175
- Resource Pool Administrator role (vCenter Server), 255
- resource pools, vCenter Server, 170-172
 - adding VMs to, 175
 - configuring, 173-174
 - viewing, 174-175
- Resource privilege (vCenter Server), 249
- resource requirements, VMs, 263
- resources
 - blogs, 464-465
 - vCenter Server, configuring, 166-175
 - VMs
 - allocating, 301
 - configuration settings, 292-297
 - websites, 465-466
- resxtop command, 449
- retentions, statistics collection settings, vCenter Server, 79
- rm -rf Linux command, 454
- rm Linux command, 454
- rmdir Linux command, 454
- rocommunity option (snmp.conf file), 342
- roles, vCenter Server, 253-256
- rollbacks, ESXi hosts, 379-381
- root passwords, ESX hosts
 - changing, 217-218
 - resetting, 218
- root user account, ESX Service Console, 222-225
- Route Based on IP Hash policy, vSwitches, 190
- Route Based on Source MAC Hash policy, vSwitches, 189
- Route Based on the Originating Port ID (default) load-balancing policy, vSwitches, 189
- routing
 - ESX hosts, configuring, 216
 - traffic, vSwitches, 199-200
- RPM (Red Hat Package Manager) files, 368

- RTFM Education blog, 464
- running Converter (vCenter Server), best practices, 326
- running snapshots, checking for, 440-441
- RV Tools, 458
- RVI (Rapid Virtualization Indexing), 53
- rwcommunity option (snmp.conf file), 342

S

- Sakac, Chad, 464
- SAN replication, DR sites, moving to, 72
- Scan mode (esxupdate), 367
- scanning hosts, Update Manager, 388-390
- Scheduled Task privilege (vCenter Server), 249
- Scheduled Tasks feature, VI Client, 126
- Scherer, Rick, 464
- SCP file transfer utilities, 458
- SCP tools, log files, viewing, 410
- scriptable installations, ESXi, 28
- Scripts tab (VMware Tools), 300
- SCSI controllers, VMs, 442
 - configuration settings, 288
- search utilities, 462-463
- security
 - Apani EpiForce VM, 260
 - Catbird V-Security, 258
 - Check Point VPN 1-VE, 259
 - Configuresoft Compliance Checker, 259
 - Configuresoft ECM (Enterprise Configuration Manager), 259
- ESX
 - best practices, 236-239
 - firewalls, 239
 - networking, 239
 - physical server hardware, 236
 - Service Console, 222-240
- ESX Service Console
 - antivirus software, 230-232
 - authentication, 229-230
 - best practices, 236-238
 - built-in firewall, 232-236
 - root user account, 222-225
 - ESXi, 239-240
 - ports, network connectivity, 219-222
- Reflex Systems Virtual Management Center, 257
- Third Brigade Deep Security, 260
- Tripwire ConfigCheck, 258
- Tripwire Enterprise 7.5V, 258
- utilities, 461-462
- vCenter Server, 243
 - AD (Active Directory) integration, 243-244
 - best practices, 256-257
 - permissions, 245-252
 - privileges, 245-252
 - roles, 253-256
- Virtual Firewall, 259
- Virtual Network Security Analyzer, 259
- VMs (virtual machines), 240-242
- X-M0n0wall virtual firewall appliance, 260
- security administrators, virtualization, educating, 21-22
- serial ports, VMs, configuration settings, 288
- Server Authorization screen, vCenter Server, 91
- server environments, documenting, 4
- server hardware, ESX, security, 236
- server ports, licensing, 135
- servers
 - application support, virtualization, 16
 - blade servers, 51-52
 - high-resource utilization servers, virtualization, 15
 - P2V (physical to virtual) migration, 302-304
 - choosing, 304-305
 - Platespin Migrate, 306
 - vCenter Server Converter, 305-329
 - vCenter Server Guided Consolidation, 329-334
 - Vizioncore vConverter, 307
 - traditional servers, 51-52

- vendor licensing models, 16
- virtualization candidates, 16
- Service Console (ESX), 451
 - antivirus software, installing on, 230-232
 - authentication, configuring, 229-230
 - backup agents, installing inside, 405
 - backup scripts, 399-401
 - built-in firewall, 232-236
 - commands, 451-454
 - Linux commands, 454-455
 - configuration files, backing up, 405
 - configuring, 179-184
 - ESX hosts
 - memory configuration, 215-216
 - split-brain condition, 145
 - hostd service, 418
 - log files, 410-414
 - nonroot user accounts, creating, 223
 - root user account, 222-225
 - running snapshots, checking for, 440
 - securing, 222-240
 - security, 222-240
 - best practices, 236-239
 - sudo, configuring, 225-228
 - troubleshooting, 416-417
- Service Console (ESXi), security, 239-240
- service firewall stop command, 236
- Sessions privilege (vCenter Server), 250
- SET command, net-snmp agent,
 - interaction, 340
- Shared Folders tab (VMware Tools), 300
- shares, vCenter Server, 168-169
 - configuring, 172
- Show Virtual Machines option (VI Client), 120
- Shrink tab (VMware Tools), 300
- single-CPU licensing, 135
- single-vCPU VMs, 48
- size considerations, LUNs, 201-203
- Sloof, Eric, 464
- SMTP (Simple Mail Transfer Protocol),
 - vCenter Server, email alerts, 345-347
- SnapHunter and SnapAlert, 459
- snapshots, 435-436
 - base filename-delta.vmdk file, 436
 - base filename-delta.vmsd file, 436
 - base filename-delta.vmsn file, 436
 - committing, 437-439
 - deleting, 435
 - file locations, changing, 439-440
 - growth, 436-437
 - locating, 441
 - multiple snapshots, creating, 435
 - running snapshots, checking for, 440-441
 - VM disks, excluding, 439
 - VMotion, 439
- SNMP (Simple Network Management Protocol)
 - ESX hosts, monitoring, 340-344
 - ESXi hosts, configuring, 343-344
- snmpd.conf file, editing, nano editor, 342
- software
 - patches, 365
 - requirements, VI Client, 93
 - upgrades, 365
- software initiator, iSCSI, 60
- Solarwinds VM Monitor, 459
- source login information, Converter Starter (vCenter), setting, 311
- source types, Converter Starter (vCenter),
 - selecting, 310
- split-brain condition, ESX hosts, 145
- SQL Server
 - ODBC connections, configuring, 82-85
 - vCenter Server, 80-82
- SQL Server 2005 Express, vCenter Server, 88
- SSH console utilities, ESX hosts, 458
- SSO (Single Sign-on), VI Client, 127
- staff, virtualization, education, 19-24
- stand-alone licensing, 69-70
 - configuring, 130-131

- Star Trek*, 72
 - Start menu folder, vCenter Server, 92
 - startup screen
 - ESX, 27
 - ESXi, 27
 - statistics collection settings, vCenter Server
 - database, 77-78
 - intervals, 79
 - levels, 78-79
 - retentions, 79
 - Status column (VMware Tools), 301
 - storage, 56
 - 10K hard drives, 56-57
 - 15K rpm hard drives, 56-57
 - boot from SAN, 57-58
 - ESX hosts, 46
 - configuring, 201-213
 - FC (Fibre Channel), 58-59
 - FC (Fibre Channel) storage, configuring, 206-209
 - iSCSI storage, 59-61
 - configuring, 209-212
 - local storage, 58
 - configuring, 205-206
 - LUNs, size considerations, 201-203
 - mixing types, 62
 - NAS/NFS, 61-62
 - NFS storage, configuring, 212-213
 - RDP (raw device mappings), 203-205
 - VMFS volumes, 203-205
 - VMS volumes, naming, 206
 - storage adapters, 55
 - VMs, choosing, 268-269
 - storage administrators, virtualization, educating, 23
 - storage utilities, 461
 - Storage VMotion, 25, 36-37
 - storage/SAN compatibility guide, 50
 - "Strategies for ESX Server Update Management," 398
 - Summary tab (VI Client), 122
 - support forums, VMTN (VMware Technology Network), 432-434
 - svmotion command, 449
 - swap partition (ESX), 100
 - switches. *See* vSwitches
 - systems compatibility guide, 50
 - Systems Insight Manager, ESX hosts,
 - hardware monitoring, 336
- ## T
- tablespace problems, vCenter Server,
 - troubleshooting, 426
 - tabs, VI Client, 122, 124-125
 - tagging (VST mode), vSwitches, 68
 - tail command, log files, viewing, 410
 - tape drives, ESX hosts, attaching to, 406
 - Tasks and Events tab (VI Client), 123
 - Tasks privilege (vCenter Server), 249
 - Tech Support Mode, ESXi hosts,
 - troubleshooting, 420-423
 - Tech Target Search Server Virtualization, 464-465
 - Search Server Virtualization, 464-465
 - VMware, 465
 - Virtualization Pro blog, 464
 - templates (VMs), 277
 - best practices, 280
 - creating, 277-279
 - displaying, 278
 - using, 279-280
 - Test mode (esxupdate), 367
 - text editor, log files, viewing, 410
 - thick client access, ESXi, 28
 - thick disks, 446-447
 - thin disks, 447
 - Third Brigade Deep Security, 260
 - third-party security products, 257, 259-260
 - three-star migration threshold (DSR), 150
 - Time Zone Selection screen (ESX), 109
 - TPS (Transparent Page Sharing), 49
 - traditional agent backups, DR sites, moving to, 71
 - traditional backup agents, 399-400
 - traditional servers, blade servers, compared, 51-52

traffic, vSwitches, routing, 199-200
 transparent page sharing, hosts, 167
 Transparent Page Sharing (TPS), 49
 trapcommunity option (snmp.conf file), 342
 trapsink option (snmp.conf file), 342
 Triggers tab, alarms, creating, 349
 Tripwire ConfigCheck, 258, 461
 Tripwire Enterprise 7.5V, 258
 Tripwire OpsCheck, 459
 troubleshooting, 409, 431
 ESX hosts, 409
 determining versions, 414-415
 esxcfg commands, 417
 hostd service, 418
 log files, 410-414
 networking problems, 417-418
 PSoD (purple screen of death), 415-416
 Service Console problems, 416-417
 vxpa service, 418-419
 ESXi hosts, 419
 determining versions, 420
 esxcfg commands, 422
 log files, 419-420
 PSoD (purple screen of death), 422
 Tech Support Mode, 420-423
 resources, 431-434
 vCenter Server, 423
 database problems, 425-426
 host problems, 426
 log files, 423-424
 VMs, 426, 431
 log files, 427
 power-state problems, 427-430
 two-star migration threshold (DSR), 150

U

Ultra ISO, ISO files, creating, 272
 unsupported hardware, 56
 Update Manager, 42-43, 397
 baselines
 attaching, 387-389
 creating, 385-388
 ESX hosts, patching, 381-393

 hosts
 remediating, 390-393
 scanning, 388-390
 installing, 382-385
 tabs, 385
 “Update Manager Administration Guide,” 397
 Update Manager plug-in, VI Client, 126
 Update Manager Service, vCenter Server, 88
 Update mode (esxupdate), 367
 updates
 ESX hosts, 365-366, 372
 ISO file, 373-374
 ESXi, 29
 ESXi hosts, 374-379
 Infrastructure, 374-377
 vhostupdate, 377-379
 VMs, 365
 upgrading
 vCenter Server, 393-394
 VMs, 394-397
 Use Explicit Failover Order policy, vSwitches, 190
 useradd command, 225
 Users and Groups tab (VI Client), 123
 “Using VMware Infrastructure for Backup and Restore,” 407

V

Van Zanten, Gabe, 464
 VBAs (virtual backup appliances), 403
 VCB (VMware Consolidated Backup), 43-45, 399-402
 vCenter Server, 3, 30-33, 48
 alarms
 configuring, 344-351
 inheritance, 345
 Cluster Summary tab, 153
 clusters, creating, 139-140
 configuration, 45-48, 136-175
 connections, VI Client, 125-126
 Converter, 307
 best practices, 325-328
 cloning VMs, 282-283
 Converter Enterprise, 317-324

- Converter Starter, 309-317
 - destination formats, 308-309
 - operating system support, 308
 - P2V (physical to virtual) migration, 305-329
 - source import formats, 308
- Converter Enterprise Service, 88
- database, 77
 - ODBC connections, 82-86
 - Oracle, 80-81
 - SQL Server, 80-82
 - statistics collection settings, 77-79
- database servers, choosing and
 - configuring, 77-86
- datacenter, creating, 137
- DPM (Distributed Power Management),
 - configuring, 156-158
- DSR (Distributed Resource Scheduler)
 - affinity rules, 151-152
 - automatic mode, 155
 - automation levels, 148-150
 - best practices, 155-156
 - cluster evaluation frequency, 155
 - configuring, 148-156
 - DSR Recommendations tab, 154
 - migration threshold, 150-151
 - monitoring, 152-154
 - VMotion, 155
- email alerts, configuring, 345-347
- ESX hosts
 - adding, 140-141
 - monitoring, 344-351
- ESXi hosts, adding, 140-141
- EVC (Enhanced VMotion Compatibility),
 - configuring, 163-166
- folders
 - blue VM/template folders, 138
 - creating, 137-138
 - datacenter/cluster/host folders, 138
- Generate Update Manager log bundle, 92
- Generate VirtualCenter1 log bundle, 92
- Guided Consolidation, P2V (physical to virtual) migration, 329-334
- HA (High Availability)
 - Admission Controls settings, 141-142
 - Advanced Option settings, 143-144
 - cluster settings, 143
 - configuring, 141-148
 - split-brain condition, 145
 - Virtual Machine Monitoring settings, 143-144
- installation
 - Database Selection screen, 90
 - Installation Type screen, 89
 - Licensing Server screen, 90
 - physical servers, 87-88
 - Server Authorization screen, 91
 - virtual machines, 87-88
- installation application, downloading, 75
- installing, 87-93
- License Server, 88
- License Server settings, 132
- limits, 170
 - configuring, 172
- Microsoft .NET Framework, 88
- Microsoft SQL Server 2005 Express, 88
- monitoring, 351-352
- patching, 365
- performance monitoring, 352-353
 - configuring, 357-360
 - ESX hosts, 355-357
 - VMs, 354-355
- reservations, 169-170
 - configuring, 172
- resource pools, 170-172
 - adding VMs to, 175
 - configuring, 173-174
 - viewing, 174-175
- resources, configuring, 166-175
- Scheduled Tasks feature, 126
- security, 243
 - AD (Active Directory) integration, 243-244
 - best practices, 256-257
 - permissions, 245-246, 250-252

- privileges, 245-246, 250-252
- roles, 253-256
- shares, 168-169
 - configuring, 172
- Start menu folder, 92
- troubleshooting, 423
 - database problems, 425-426
 - host problems, 426
 - log files, 423-424
- Update Manager, 381-393, 397
 - baselines, 385-389
 - installing, 382-385
 - remediating hosts, 390-393
 - scanning hosts, 388-390
 - tabs, 385
- Update Manager Service, 88
- upgrading, 393-394
- VI Client (Virtual Infrastructure Client), 88, 92
 - SSO (Single Sign-on), 127
- virtual hosts, memory usage, 166-168
- VM templates, adding existing templates back into, 280
- VMotion, configuring, 159-163
- VMware Capacity Planner Service, 91
- VMware Converter Enterprise Service, 92
- VMware Infrastructure Update, 92
- VMware Infrastructure Web Access, 91
- VMware License Server, 91
- VMware Update Manager Service, 92
- VMware vCenter Server, 88
- VMware VirtualCenter1 Server, 91
- vCenter Server version 2.5 Update 3, 2
- vCloud, 2
- vCPUs, VMs, assigning, 263, 267
- vdf Linux command, 455
- Veeam Backup, 403, 461
- Veeam FastSCP file transfer utility, 458
- Veeam Management suite, performance
 - monitoring, 362
- Veeam Monitor, 460
- Veeam Monitor Free Edition, 460
- Veeam Reporter, 463
- vendor licensing models, virtualization, 16
- vendors
 - CPUs, choosing, 53-54
 - virtualization, 1
 - virtualization support policies, 19
- versions
 - ESX host components, determining, 414-415
 - ESXi host components, determining, 420
- vFoglight, 460
 - performance monitoring, 362
- VGT (virtual machine guest tagging), 177
- VI Client, 119
 - Alarm tab, 348
 - Converter Enterprise plug-in, 126
 - datastore browser, cloning VMs, 284-286
 - Datastores view, 121
 - ESX hosts, connecting to directly, 125-126
 - folders, creating, 121
 - hardware health status, displaying, 336
 - Hosts and Clusters view, 120
 - installation, 93-97
 - hardware requirements, 93
 - software requirements, 93
 - layout, 120-125
 - log files, viewing, 410, 420
 - Networks view, 121
 - nonroot user accounts, creating, 223
 - objects, 122-125
 - performance monitoring, 352-353
 - ESX hosts, 355-357
 - VMs, 354-355
 - plug-ins, 126
 - Scheduled Tasks feature, 126
 - Show Virtual Machines option, 120
 - SSO (Single Sign-on), 127
 - tabs, 122, 124-125
 - Update Manager plug-in, 126
 - Users and Groups tab, 224
- vCenter Server, 88, 92
 - connecting to, 125-126
- Virtual Infrastructure Client login screen, 95

- Virtual Infrastructure Client settings, 96-97
- Virtual Machines and Templates view, 120-121
- VM templates, displaying, 278
- Web access default welcome screen, 94
- window panes, 120
- vi Linux command, 454
- VI3, components, name changes, 3
- vicfg-advcfg command, 449
- vicfg-cfgbackup command, 449
- vicfg-dns command, 449
- vicfg-dumppart command, 449
- vicfg-module command, 450
- vicfg-mpath command, 450
- vicfg-nas command, 450
- vicfg-nics command, 450
- vicfg-ntp command, 450
- vicfg-rescan command, 450
- vicfg-route command, 450
- vicfg-snmpp command, 450
- vicfg-syslog command, 450
- vicfg-user command, 450
- vicfg-vmhbadevs command, 450
- vicfg-vmknic command, 450
- vicfg-vswitch command, 450
- video controllers, VMs, 441
- viewing vCenter Server resource pools, 174-175
- vifs command, 450
- vihostupdate, ESXi hosts, updating, 377-379
- vihostupdate- command, 420, 450
- VIMA (VMware Infrastructure Management Assistant), 455-456
- virtual backup appliances (VBAs), 403
- virtual CD/DVD drives, ISO files, selecting for, 274
- Virtual Datacenter OS, 2
- virtual device nodes, VMs, choosing, 271
- virtual disks, 446
 - 2GBSparse disks, 447
 - creating, 263, 268
 - multiple formats, 447-449
 - nodes, choosing, 271
 - RDMs (raw device mappings), 446
 - thick disks, 446-447
 - thin disks, 447
- virtual environments, building, 75
 - database servers, 77-86
 - ESX, 98-116
 - ESXi, 98-116
 - licensing servers, 97-98
 - preparation, 75-77
 - vCenter Server, 87-93
 - VI Client, 93-97
- Virtual Firewall, 259
- Virtual Geek blog, 464
- virtual hardware
 - virtual disks, 446
 - 2GBSparse disks, 447
 - multiple formats, 447-449
 - RDMs (raw device mappings), 446
 - thick disks, 446-447
 - thin disks, 447
 - VMs, 441-442
- virtual hosts, memory usage, 166-168
- Virtual Infrastructure 3, 1
- Virtual Infrastructure Client settings, VI Client, 96-97
- Virtual Infrastructure Client login screen, 95
- Virtual Machine Administrator role (vCenter Server), 254
- "Virtual Machine Backup Guide," 407
- virtual machine guest tagging (VGT), 177
- Virtual Machine Monitoring settings, HA (High Availability), 143-144
- Virtual Machine Power User role (vCenter Server), 254
- Virtual Machine privilege (vCenter Server), 247-249
- Virtual Machine User role (vCenter Server), 254
- Virtual Machines and Templates view (VI Client), 120-121
- Virtual Machines tab (VI Client), 122

- virtual network adapters
 - enhanced vmxnet, 66
 - PCNet32, 65
 - vmxnet, 65
- virtual network adapters, PCNet32, 65
- Virtual Network Security Analyzer, 259
- virtual networking, 63
 - pNICs (physical NICs), 63-65
 - virtual switches, 66-69
 - vNICs (virtual NICs), 65-66
- virtual NICs (vNICs), 65-66
 - VMs, choosing, 267, 269
- Virtual SAN Appliance (VSA), 58
- Virtual Strategy Magazine, 465
- virtual switch tagging (VST), 176
- virtual switches, 66-69
- VirtualCenter, name change, 3
- virtualization, 1-2
 - 100% virtualized environments, 17
 - application compatibility, 17-19
 - application support, 16
 - benefits, 1
 - high-resource utilization servers, 15
 - known problems, 4
 - physical servers, compared, 1
 - pilot projects, 20
 - servers, good candidates, 16
 - staff, education, 19-24
 - vendor licensing models, 16
 - vendors, 1
- Virtualization Admin, 465
- Virtualization Review, 465
- Virtualization Sys-con, 465
- Virtualization.com, 465
- Virtualization.info, 465
- VISBU, 460
- Visio Stencils, 459
- Vizioncore vConverter, P2V (physical to virtual) migration, 307
- Vizioncore vFoglight, 460, 463
- Vizioncore vRanger, 461
- Vizioncore vRanger Pro, 403
- VLAN tagging, 176-177
- Vlance network adapters, 65
- VM (virtual machine)
- VM Backup Script, 460
- VM Blog, 465
- VM cloning, DR sites, moving to, 72
- VM Explorer, 461
- VM guest tagging (VGT mode), vSwitches, 67
- VM image backups, DR sites, moving to, 72
- VM NICs, pNICs, mapping to, 194
- VM replication, DR sites, moving to, 72
- VM sprawl, 72-73
- vm-support command, log files, viewing, 412-413
- VM/ETC blog, 464
- VMC (Virtualization Management Center), 257
- vmCDconnected, 459
- VMDK file type (VMs), 444
- VMDK Recovery Tool, 458
- vmfs partition (ESX), 100-101
- VMFS volumes
 - naming, 206
 - RDP (raw device mapping), compared, 203-205
 - VMs, creating on, 302
- vmkernel partition (ESX), 100
- VMkernel
 - ESX hosts, configuring, 185-186
 - log files, 410
 - VMotion, configuring, 160-161
- vmkfstools command-line utility, 450, 453
 - VMs, cloning, 283-284
- VMotion, 33-35, 459
 - Connection Settings screen, 160
 - DSR (Distributed Resource Scheduler), 155
 - Network Access screen, 160
 - pNICs (physical NICs), 64
 - snapshots, 439
 - vCenter Server, configuring, 159-163
 - VMkernel, configuring, 160-161

- VMs (virtual machines), 441
 - backups, 399, 406-407
 - backup scripts, 400-401
 - esXpress, 403-404
 - file-level backups, 404-405
 - image-level backups, 404-405
 - traditional backup agents, 400
 - utilities, 460-461
 - VCB (Consolidated Backup), 401-402
 - Veeam Backup, 403
 - vRanger Pro, 403
- CD-ROM devices, selecting, 276
- CD/DVD drives, configuring, 287
- cloning, 281-286
- configuration, 45-48
- configuration settings, 286
 - advanced configuration options, 289-291
 - hardware, 286-288
 - resources, 292-297
- CPUs, assigning, 286
- creating, 263-302
 - best practices, 301-302
- datastores, choosing, 265-266
- ESX hosts, configuring
 - startup/shutdown, 216-217
- file types, 442-445
- floppy disks, configuring, 286
- floppy drives, 275
- guest operating systems, choosing, 265-266
- hard disks, configuring, 288
- ISO files, 272-277
- locations, choosing, 265
- LUN/VMFS volumes, optimal
 - number, 302
- memory, assigning, 263, 267-268, 286
- monitoring utilities, 459-460
- naming, 265
 - Converter Starter (vCenter), 314
- network adapters, configuring, 287
- P2V (physical to virtual) migration, 263, 302
 - choosing, 304-305
 - cold cloning, 303-304
 - hot cloning, 303-304
 - Platespin Migrate, 306
 - vCenter Server Converter, 305-329
 - vCenter Server Guided
 - Consolidation, 329-334
 - Vizioncore vConverter, 307
- parallel ports, configuring, 288
- partial memory reservations, creating for, 302
- performance monitoring, 352-354
 - configuring, 357-360
 - CPU statistics, 354
 - disk statistics, 354
 - memory statistics, 354-355
 - network statistics, 355
- RDMs, choosing, 270
- resource requirements, 263
- resources, allocating, 301
- scripting
 - Perl, 456
 - PowerShell, 456-457
- SCSI controllers, configuring, 288
- securing, 240-242
- serial ports, configuring, 288
- snapshots, 435-436
 - base filename-delta.vmdk file, 436
 - base filename-delta.vmsd file, 436
 - base filename-delta.vmsn file, 436
 - changing file locations, 439-440
 - committing, 437-439
 - deleting, 435
 - excluding VM disks from, 439
 - growth, 436-437
 - locating, 440-441
 - multiple snapshots, 435
 - VMotion, 439
- storage adapters, choosing, 268-269

- templates, 277
 - best practices, 280
 - creating, 277-279
 - displaying, 278
 - using, 279-280
- troubleshooting, 426, 431
 - log files, 427
 - power-state problems, 427-430
- unnneeded hardware, removing, 302
- updating, 365
- upgrading, 394-397
- vCenter Server
 - adding resource pools to, 175
 - installation, 87-88
- vCPUs, assigning, 263, 267
- virtual CD/DVD drives, ISO files, 274
- virtual device nodes, choosing, 271
- virtual disk nodes, choosing, 271
- virtual disks, 446
 - 2GBSparse disks, 447
 - creating, 263, 268
 - multiple formats, 447-449
 - RDMs (raw device mappings), 446
 - thick disks, 446-447
 - thin disks, 447
- virtual hardware components, 441-442
- virtual NICs, choosing, 267, 269
- VMFS volumes, creating on, 302
- VMware Tools, installing on, 297-301
- VMSD file type (VMs), 444
- VMSN file type (VMs), 444
- VMSS file type (VMs), 444
- VMTN (VMware Technology Network),
 - support forums, 432-434
- VMTS Patch Manager, 459
- VMware, 2
 - knowledge base website, 431-432
 - licenses, downloading, 75
 - licensing, 69-70
 - Virtual Infrastructure 3, 1
 - VMware Capacity Planner Service, vCenter Server, 91
 - "VMware Consolidated Backup," 407
 - "VMware Consolidated Backup: Improvements in Version 3.5," 407
 - "VMware Consolidated Backup—Partner Integration Guide," 407
 - VMware Converter, 459
 - VMware Converter Enterprise Service, vCenter Server, 92
 - "VMware ESX Server 3 Patch Management," 398
 - VMware Infrastructure Management Assistant (VIMA), 455-456
 - VMware Infrastructure Update, vCenter Server, 92
 - VMware Infrastructure Web Access, vCenter Server, 91
 - VMware License Server, vCenter Server, 91
 - VMware Technology Network (VMTN),
 - support forums, 432-434
 - VMware Tips blog, 464
 - VMware Tools
 - About tab, 300
 - Devices tab, 300
 - Options tab, 300
 - Scripts tab, 300
 - Shared Folders tab, 300
 - Shrink tab, 300
 - Status column, 301
 - upgrading, 394-397
 - VMs, installing on, 297-301
 - VMware Tools (VMs), 289
 - "VMware Update Manager Performance and Best Practices," 397
 - VMware Update Manager privilege (vCenter Server), 250
 - VMware Update Manager Service, vCenter Server, 92
 - "VMware VI3 Upgrade Guide," 397

- VMware VirtualCenter1 Server, vCenter Server, 91
- VMware Virtualization Evangelist blog, 464
- vmware-cmd command, 450
 - VM power-state problems, troubleshooting, 427-428
- VMware-land, 465
- vmware-support command, VM power-state problems, troubleshooting, 428-429
- vmware-vmupgrade.exe command, 396-397
- VMX file type (VMs), 443
- VMXF file type (VMs), 444
- vmxnet virtual network adapter, 65
- vNICs
 - Converter Starter (vCenter), setting information, 316
 - ESX hosts, changing MAC addresses, 196-197
- vNICs (virtual NICs), 65-66
- VP Snapper, 458
- VPN-1 VE (Virtual Edition), 259
- vpaa log file, 411
- vRanger Pro, 403
- Vroom! by VMware's Performance Team, 464
- VSA (Virtual SAN Appliance), 58
- vSMP, 40-42
- vSphere, 2
- VST (virtual switch tagging), 176
- vSwitch Properties window (ESX), 186-192
- vSwitches, 67
 - ESX hosts
 - configuring, 186-193
 - internal-only vSwitches, 195
 - load-balancing policies, 189-190
 - Network Failover Detection, 190
 - External switch tagging (EST mode), 68
 - traffic, routing, 199-200
 - VM guest tagging (VGT mode), 67
 - vSwitch tagging (VST mode), 68
- VSWP file type (VMs), 443
- VT FlexMigration, 54
- vxpa service, ESX hosts, troubleshooting, 418-419

W-X-Y-Z

- web access, log file, 411
- Web access default welcome screen (VI Client), 94
- web client access, ESXi, 28
- websites, resources, 465-466
- Win Image, ISO files, creating, 272
- window panes, VI Client, 120
- WinSCP, log files, viewing, 410
- WinSCP file transfer utility, 458
- wizards, New VM Wizard, 264-271
- X-M0n0wall virtual firewall appliance, 260
- X86 Virtualization, 465
- Xtravirt Virtual SAN, 461
- Yellow Bricks blog, 464

LEADING VMWARE TITLES

from Prentice Hall and Pearson Certification



VMware ESX Server in the Enterprise

Author: Edward L. Haletky

December 2007

ISBN 10: 0132302071

ISBN 13: 9780132302074

Master VMware ESX best practices and make the most of the powerful virtualization tools that solve IT problems



VMware Virtual Infrastructure Security: Securing ESX and the Virtual Environment


Author: Edward L. Haletky

August 2009

ISBN 10: 0137158009

ISBN 13: 9780137158003

Let VMware best-selling author Edward Haletky help you master VMware security issues



VMware VI SDK


Author: Steve Jin

October 2009

ISBN 10: 0137153635

ISBN 13: 9780137153633

Fully unlock the power of virtualization by writing VMware applications that automate virtual infrastructure management



VCP Exam Cram

Author: Elias Khnaser

December 2008

ISBN 10: 0789738058

ISBN 13: 9780789738059

Prepare for the VMware Certified Professional exam with proven Exam Cram preparation and practice tools on virtualization technologies



VMware VCP 310 Video Mentor

Author: Chris McCain

September 2009

ISBN 10: 0789740427

ISBN 13: 9780789740427

Expert, personal DVD video training: the fastest, easiest, most cost-effective way to prepare for VMware's VCP-310 exam

To find out more about these and other networking titles, visit www.informit.com



REGISTER



THIS PRODUCT

informit.com/register

Register the Addison-Wesley, Exam Cram, Prentice Hall, Que, and Sams products you own to unlock great benefits.

To begin the registration process, simply go to **informit.com/register** to sign in or create an account.

You will then be prompted to enter the 10- or 13-digit ISBN that appears on the back cover of your product.

Registering your products can unlock the following benefits:

- Access to supplemental content, including bonus chapters, source code, or project files.
- A coupon to be used on your next purchase.

Registration benefits vary by product. Benefits will be listed on your Account page under Registered Products.

About InformIT — THE TRUSTED TECHNOLOGY LEARNING SOURCE

INFORMIT IS HOME TO THE LEADING TECHNOLOGY PUBLISHING IMPRINTS Addison-Wesley Professional, Cisco Press, Exam Cram, IBM Press, Prentice Hall Professional, Que, and Sams. Here you will gain access to quality and trusted content and resources from the authors, creators, innovators, and leaders of technology. Whether you're looking for a book on a new technology, a helpful article, timely newsletters, or access to the Safari Books Online digital library, InformIT has a solution for you.

informIT.com


THE TRUSTED TECHNOLOGY LEARNING SOURCE

Addison-Wesley | Cisco Press | Exam Cram
IBM Press | Que | Prentice Hall | Sams

SAFARI BOOKS ONLINE

LearnIT at InformIT

Go Beyond the Book



11 WAYS TO LEARN IT at www.informIT.com/learn

The online portal of the information technology
publishing imprints of Pearson Education

Try Safari Books Online FREE

Get online access to 5,000+ Books and Videos



Safari[®]
Books Online

FREE TRIAL—GET STARTED TODAY!
www.informit.com/safaritrial



Find trusted answers, fast

Only Safari lets you search across thousands of best-selling books from the top technology publishers, including Addison-Wesley Professional, Cisco Press, O'Reilly, Prentice Hall, Que, and Sams.



Master the latest tools and techniques

In addition to gaining access to an incredible inventory of technical books, Safari's extensive collection of video tutorials lets you learn from the leading video training experts.

WAIT, THERE'S MORE!



Keep your competitive edge

With Rough Cuts, get access to the developing manuscript and be among the first to learn the newest technologies.



Stay current with emerging technologies

Short Cuts and Quick Reference Sheets are short, concise, focused content created to get you up-to-speed quickly on new and cutting-edge technologies.

Addison
Wesley

Adobe Press

ALPHA

Cisco Press

FT Press
FINANCIAL TIMES

IBM
Press

lynda.com

Microsoft
Press

New
Riders

O'REILLY

Peachpit
Press

PRENTICE
HALL

QUE

Redbooks


SAMS

Sas
Publishing

Sun
microsystems

Wiley
Publishing

WILEY



FREE Online Edition

Your purchase of **VMware® VI3 Implementation and Administration** includes access to a free online edition for 45 days through the Safari Books Online subscription service. Nearly every Prentice Hall book is available online through Safari Books Online, along with more than 5,000 other technical books and videos from publishers such as Addison-Wesley Professional, Cisco Press, Exam Cram, IBM Press, O'Reilly, Que, and Sams.

SAFARI BOOKS ONLINE allows you to search for a specific answer, cut and paste code, download chapters, and stay current with emerging technologies.

Activate your FREE Online Edition at
www.informit.com/safarifree



STEP 1: Enter the coupon code: ZGA0ZBI.



STEP 2: New Safari users, complete the brief registration form. Safari subscribers, just log in.

If you have difficulty registering on Safari or accessing the online edition, please e-mail customer-service@safaribooksonline.com

Safari 
 Books Online

Addison
Wesley

Adobe Press

ALPHA

Cisco Press

FT Press
FINANCIAL TIMES

IBM
Press

lynda.com

Microsoft
Press

New
Riders

O'REILLY

Peachpit
Press

PRENTICE
HALL

que

Redbooks

SAMS

sas
Publishing

Sun
microsystems

WILEY
Publishing

WILEY