



## Virtualization with VMware ESX Server

By Al Muller, Seburn Wilson

.....  
Publisher: **Syngress**

Pub Date: **August 2005**

Print ISBN-10: **1-59749-019-9**

Print ISBN-13: **978-1-59-749019-1**

Pages: **608**

[Table of Contents](#) | [Index](#)

## Overview

A virtual evolution in IT organizations throughout the world has begun. It is estimated that currently 3% of all servers run virtually and that number is expected to grow rapidly over the next 5 years. VMware's ESX server is the enterprise tool that provides the platform on which a complete virtual infrastructure can be designed, implemented, and managed. ESX is the most powerful, resilient and customizable of VMware's three virtual platforms and this book explores many of the possibilities that a virtual infrastructure running on ESX Server provides.

Server Sprawl and escalating IT costs have managers and system administrators scrambling to find ways to cut costs and reduce Total Cost of Ownership of their physical infrastructure. Combining software applications onto a single server, even if those applications are from the same software vendor, can be dangerous and problems hard to troubleshoot. VMware ESX Server allows you to consolidate 15 to 20 or even more servers onto a single physical server reducing hardware, electrical, cooling, and administrative costs. These virtual servers run completely independent of each other so if one crashes the other are not affected. Planning and implementing a server consolidation is a complex process.

This book details the requirements for such a project, includes sample forms and templates, and delivers several physical to virtual migration strategies which will save both time and costs. You will easily be able to plan and deploy VMware's ESX Server and begin down the path of an evolved, virtual infrastructure in which costs, administration overhead, and complexity are reduced. VMware has provided the technology for a virtual infrastructure and this book shows you how to build it.



## Virtualization with VMware ESX Server

By Al Muller, Seburn Wilson

.....  
 Publisher: **Syngress**  
 Pub Date: **August 2005**  
 Print ISBN-10: **1-59749-019-9**  
 Print ISBN-13: **978-1-59-749019-1**  
 Pages: **608**

[Table of Contents](#) | [Index](#)

[Copyright](#)

[Acknowledgments](#)

[Coauthors](#)

[Contributing Authors](#)

[Foreword](#)

[About This Book](#)

[About Callisma](#)

[Chapter 1. The Joy of ESX](#)

[Section 1.1. Introduction](#)

[Section 1.2. Target Audience](#)

[Section 1.3. VMware Terms and Concepts](#)

[Section 1.4. Summary](#)

[Section 1.5. Solutions Fast Track](#)

[Section 1.6. Frequently Asked Questions](#)

[Chapter 2. Installation](#)

[Section 2.1. Introduction](#)

[Section 2.2. Pre-installation Tasks](#)

[Section 2.3. Performing the Installation](#)

[Section 2.4. Summary](#)

[Section 2.5. Solutions Fast Track](#)

[Section 2.6. Frequently Asked Questions](#)

[Chapter 3. Configuring a Virtual Machine, Gold Builds, and Clones](#)

[Section 3.1. Introduction](#)

[Section 3.2. Virtual Machine Basics](#)

[Section 3.3. Configuring a Virtual Machine](#)

[Section 3.4. Gold Builds and Clones](#)

[Section 3.5. Modify Your Master](#)

[Section 3.6. Summary](#)

[Section 3.7. Solutions Fast Track](#)

[Section 3.8. Frequently Asked Questions](#)

[Chapter 4. Advanced Configurations](#)

[Section 4.1. Introduction](#)

[Section 4.2. Network Traffic Shaping](#)

[Section 4.3. Disk Bandwidth](#)

[Section 4.4. CPU Management, Affinity](#)

[Section 4.5. NUMA and Processor/Memory Affinity](#)

[Section 4.6. Raw Device Mapping](#)

[Section 4.7. Clusters](#)

[Section 4.8. Summary](#)

[Section 4.9. Solutions Fast Track](#)

[Section 4.10. Frequently Asked Questions](#)



## [Chapter 5. Virtual Networking](#)

### [Section 5.1. Introduction](#)

### [Section 5.2. Virtual Switches](#)

### [Section 5.3. Physical and Virtual NICs](#)

### [Section 5.4. MAC Addresses](#)

### [Section 5.5. Port Groups and VLANs](#)

### [Section 5.6. Networking Tools](#)

### [Section 5.7. Summary](#)

### [Section 5.8. Solutions Fast Track](#)

### [Section 5.9. Frequently Asked Questions](#)

## [Chapter 6. Physical-to-Virtual Migrations](#)

### [Section 6.1. Introduction](#)

### [Section 6.2. P-to-V Terms](#)

### [Section 6.3. P-to-V for Free](#)

### [Section 6.4. Format Your Hard Drive](#)

### [Section 6.5. BartPE](#)

### [Section 6.6. VMware P2V Assistant 2](#)

### [Section 6.7. Platespin PowerP2V](#)

### [Section 6.8. P-to-V Conclusion](#)

### [Section 6.9. Summary](#)

### [Section 6.10. Solutions Fast Track](#)

### [Section 6.11. Frequently Asked Questions](#)

## [Chapter 7. Server Consolidation](#)

### [Section 7.1. Introduction](#)

### [Section 7.2. Server Consolidation and Virtualization Defined](#)

### [Section 7.3. Business Case and Technical Requirements](#)

### [Section 7.4. Developing Your Project Methodology](#)

### [Section 7.5. Summary](#)

### [Section 7.6. Solutions Fast Track](#)

### [Section 7.7. Frequently Asked Questions](#)

## [Chapter 8. Cool Tools for a Virtual Infrastructure](#)

### [Section 8.1. Introduction](#)

### [Section 8.2. ISO Images and WinISO](#)

### [Section 8.3. WinSCP](#)

### [Section 8.4. PuTTY](#)

### [Section 8.5. NewSID](#)

### [Section 8.6. Platespin's PowerP2V and PowerRecon](#)

### [Section 8.7. VirtualCenterVmotion](#)

### [Section 8.8. vmkfstools and vmware-cmd](#)

### [Section 8.9. VMware Tools](#)

### [Section 8.10. Summary](#)

### [Section 8.11. Solutions Fast Track](#)

### [Section 8.12. Frequently Asked Questions](#)

## [Chapter 9. Administering a Virtual Infrastructure](#)

### [Section 9.1. Introduction](#)

### [Section 9.2. VMware Tools for Managing the Virtual Infrastructure](#)

### [Section 9.3. Understanding VirtualCenter Terms and Concepts](#)

### [Section 9.4. Installation and Configuration of VirtualCenter 1.2](#)

### [Section 9.5. Using VirtualCenter to Administer Your Virtual Infrastructure](#)

### [Section 9.6. Using the Management Interface and Service Console to Administer the Virtual Infrastructure](#)

### [Section 9.7. Summary](#)

### [Section 9.8. Solutions Fast Track](#)

### [Section 9.9. Frequently Asked Questions](#)

## [Chapter 10. Best Practices, Gotchas, and Common Problems](#)

### [Section 10.1. Introduction](#)

### [Section 10.2. Best Practices](#)

### [Section 10.3. Gotchas](#)

### [Section 10.4. Common Problems](#)

### [Section 10.5. Summary](#)

### [Section 10.6. Solutions Fast Track](#)

### [Section 10.7. Frequently Asked Questions](#)

## [Chapter 11. VMware Under the Covers](#)

### [Section 11.1. Introduction](#)

### [Section 11.2. Understanding the ESX Server File and Directory Structure](#)

### [Section 11.3. The ESX Server Boot Process](#)

### [Section 11.4. Interaction between the Service Console, VMkernel, and Virtual Machines](#)

### [Section 11.5. Summary](#)

### [Section 11.6. Solutions Fast Track](#)

### [Section 11.7. Frequently Asked Questions](#)

## [Chapter 12. Backup Strategies for ESX Server and Virtual Machines](#)

### [Section 12.1. Introduction](#)

### [Section 12.2. Backing Up the Service Console](#)

### [Section 12.3. Backing Up Virtual Machines](#)

### [Section 12.4. Best Practices](#)

### [Section 12.5. Summary](#)

### [Section 12.6. Solutions Fast Track](#)

### [Section 12.7. Frequently Asked Questions](#)

## [Chapter 13. Common and Necessary Command-Line Interface \(CLI\) Commands](#)

### [Section 13.1. Introduction](#)

### [Section 13.2. Using the Command-Line Interface](#)

### [Section 13.3. Configuring and Managing ESX Server Service Console from the Command Line](#)

### [Section 13.4. Managing Virtual Machines from the Command Line](#)

### [Section 13.5. Summary](#)

### [Section 13.6. Solutions Fast Track](#)

### [Section 13.7. Frequently Asked Questions](#)

## [Chapter 14. A Peek at the Future: It's Totally Cool](#)

### [Section 14.1. Introduction](#)

### [Section 14.2. Evolved Virtually You Have Become](#)

### [Section 14.3. Bigger and Badder Brains](#)

### [Section 14.4. ESX Server 3.0...Better the Beta](#)

### [Section 14.5. Faultless Fault Tolerance](#)

### [Section 14.6. The Rise of Utility Computing](#)

### [Section 14.7. Your Moment of Xen](#)

### [Section 14.8. Microsoft Virtual Server...Watch Your Six](#)

### [Section 14.9. Virtual Storage Is the Place to Live](#)

### [Section 14.10. Summary](#)

### [Section 14.11. Solutions Fast Track](#)

### [Section 14.12. Frequently Asked Questions](#)

## [Appendix A. Recommended Reading](#)

### [Section A.1. Guides](#)

### [Section A.2. Books](#)

### [Section A.3. White Papers](#)

### [Section A.4. Web Sites](#)

## [Index](#)





Syngress Publishing, Inc., the author(s), and any person or firm involved in the writing, editing, or production (collectively "Makers") of this book ("the Work") do not guarantee or warrant the results to be obtained from the Work.

There is no guarantee of any kind, expressed or implied, regarding the Work or its contents. The Work is sold AS IS and WITHOUT WARRANTY. You may have other legal rights, which vary from state to state.

In no event will Makers be liable to you for damages, including any loss of profits, lost savings, or other incidental or consequential damages arising out from the Work or its contents. Because some states do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

You should always use reasonable care, including backup and other appropriate precautions, when working with computers, networks, data, and files.

Syngress Media®, Syngress®, "Career Advancement Through Skill Enhancement®," "Ask the Author UPDATE®," and "Hack Proofing®," are registered trademarks of Syngress Publishing, Inc. "Syngress: The Definition of a Serious Security Library"™, "Mission Critical™," and "The Only Way to Stop a Hacker is to Think Like One™" are trademarks of Syngress Publishing, Inc. Brands and product names mentioned in this book are trademarks or service marks of their respective companies.

KEY	SERIAL NUMBER
001	HJIRTCV764
002	PO9873D5FG
003	829KM8NJH2
004	BBBCCC5467
005	CVPLQ6WQ23
006	VBP965T5T5
007	HJJJ863WD3E
008	2987GVTWMK
009	629MP5SDJT
010	IMWQ295T6T

PUBLISHED BY  
Syngress Publishing, Inc.  
800 Hingham Street  
Rockland, MA 02370

## Virtualization with VMware™ ESX Server™

Copyright © 2005 by Syngress Publishing, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.



# Acknowledgments

Syngress would like to acknowledge the following people for their kindness and support in making this book possible.

A special thank-you to Ralph Troupe of Callisma. You have been a great business partner and friend of the Syngress family for many years.

To all of the talented and hardworking people at Callisma and SBC who helped make this book possible: Al Muller, Seburn Wilson, Don Happe, Gary J. Humphrey, Bob Gamiel, Allison Alli, Ruth Anne Renaud, Carrie Chikahisa, and Leslye Harter.

Syngress books are now distributed in the United States and Canada by O'Reilly Media, Inc. The enthusiasm and work ethic at O'Reilly are incredible, and we would like to thank everyone there for their time and efforts to bring Syngress books to market: Tim O'Reilly, Laura Baldwin, Mark Brokering, Mike Leonard, Donna Selenko, Bonnie Sheehan, Cindy Davis, Grant Kikkert, Opol Matsutaro, Steve Hazelwood, Mark Wilson, Rick Brown, Leslie Becker, Jill Lothrop, Tim Hinton, Kyle Hart, Sara Winge, C. J. Rayhill, Peter Pardo, Leslie Crandell, Regina Aggio, Pascal Honscher, Preston Paull, Susan Thompson, Bruce Stewart, Laura Schmier, Sue Willing, Mark Jacobsen, Betsy Waliszewski, Dawn Mann, Kathryn Barrett, John Chodacki, Rob Bullington, and Aileen Berg.

The incredibly hardworking team at Elsevier Science, including Jonathan Bunkell, Ian Seager, Duncan Enright, David Burton, Rosanna Ramacciotti, Robert Fairbrother, Miguel Sanchez, Klaus Beran, Emma Wyatt, Chris Hossack, Krista Leppiko, Marcel Koppes, Judy Chappell, Radek Janousek, and Chris Reinders for making certain that our vision remains worldwide in scope.

David Buckland, Marie Chieng, Lucy Chong, Leslie Lim, Audrey Gan, Pang Ai Hua, Joseph Chan, and Siti Zuraidah Ahmad of STP Distributors for the enthusiasm with which they receive our books.

David Scott, Tricia Wilden, Marilla Burgess, Annette Scott, Andrew Swaffer, Stephen O'Donoghue, Bec Lowe, Mark Langley, and Anyo Geddes of Woodslane for distributing our books throughout Australia, New Zealand, Papua New Guinea, Fiji, Tonga, Solomon Islands, and the Cook Islands.

## Coauthors

**A1 Muller** has been in the IT field since 1995, getting his start as a database administrator in the Navy. In 2002 he began using VMware's GSX Server and within a year was involved in his first virtualization project. Since then, he has been an eager proponent of virtualization technology and has worked on a number of different server consolidation and virtualization projects.

He holds a bachelor's degree in English and plans on writing a series of books on the virtual evolution taking place in data centers worldwide and the tools required to leverage and support a virtual infrastructure.

A native Californian, he lives in San Diego with his wife, Sara, and their dog, Grace.

**Seburn Wilson** is a VMware Certified Professional, MCSE, and HPASE with more than nine years of industry experience in network operating systems, messaging, and directory services. As a Senior Consultant at Callisma, he has focused on Active Directory design, migrations, and server consolidation for large enterprise customers while serving as the technical lead and project manager for all phases of client engagements. Prior to joining Callisma, he worked for Lucent Technologies as a Network Systems Consultant. Seburn holds a BA in Business Administration from St. Leo University with a specialization in Computer Information Systems. Seburn is married to his wife, Christine, and has four children Ashley, Connor, Seth, and Faith and would like to thank all of them for tolerating him while writing this book.



## Contributing Authors

**Don Happe** has more than 17 years of industry experience in internetworking, server operating systems, messaging, and directory services consulting. At Callisma Don founded the Enterprise Operating Systems practice and is the National Practice Director. The Enterprise Operating Systems practice focuses on delivering strategic and project-based solutions to Callisma's clients in the areas of server consolidation and virtualization; directory services; identity management; provisioning; and messaging-related projects. Don has extensive experience in the planning, design, and implementation of enterprise technologies. He is also proficient at program/project management, strategic planning, architecture and design, and pilot/full-scale deployments. Before joining Callisma, Don worked for Computer Sciences Corporation, where he led various large-scale data center, network operating systems, and internetworking projects.

**Gary J. Humphrey** has more than 25 years of industry experience in networking, security, and data communications. Before joining Callisma, he held positions as senior director, architect, vice president, and president with Fortune 50/100/250 companies, along with foreign military security communications design companies. Previous design and implementations of GSX and ESX include the consolidation of more than 2,500 servers into VMware environments. He holds degrees in Petroleum Land Management and Business, with professional certifications of CISSP, MCSE2K3, CISSM, CISA, CCNA, CCDA, and CCA. Currently, he is involved in SOX security audits and the design of role-based access control models for Fortune 100 companies.

# Foreword

Server virtualization has the potential of transforming data centers in almost every major enterprise. This compelling, proven technology comes to market with the promise of providing significant cost savings to enterprises that use it.

In a nutshell, server virtualization allows multiple instances of a software application (including its underlying operating system environment) to run concurrently on one physical server. The benefits are realit frees IT professionals from the shackles of the "one server, one application" vendor mandates of the past. The result is significant cost savings from reduced equipment, administrative, and facilities expenses. Server virtualization provides important "soft" benefits as well, including improved utilization of servers and load balancing, faster speed to production for new applications, greater flexibility to support a mixture of production and test environments, and simplified storage and disaster recovery strategies.

IT groups are seeking strategies to deliver more reliable, adaptable, and scalable infrastructure solutions at a predictable cost. Virtualization brings us one step closer to utility computing as more of our enterprise clients view themselves as utilities, or internal "service providers." I am confident that you will find this book not only valuable in describing the area of VMware technology but also effective in articulating the people and process components of server virtualization. We hope you enjoy reading this book as much as we've enjoyed writing it!

*Ralph Troupe,*

*President and CEO,*

*Callisma*

## About This Book

According to Gartner, Inc., enterprises that do not capitalize on server virtualization will spend 25 percent more annually for hardware, software, labor, and space for Intel-based servers.

The goal of this book is to provide you with a thorough understanding of the VMware ESX Server. The target audience includes IT professionals who manage sprawling physical server environments and IT leaders looking to better leverage this technology. Whether you are an IT manager, an administrator, or a student studying to pass the VMware Certified Professional (VCP) exam, this comprehensive guide will take you from the basics to advanced topics and answer important questions such as:

- What are your goals and strategies for virtualization?
- How do you develop a detailed plan identifying target servers for physical-to-virtual migrations?
- How should your architecture leverage server virtualization technologies?
- What are the best practice design options for mitigating the risks of running multiple applications on one server?
- How and when should you consider new processes or automated procedures for effective server provisioning and management?
- What licensing issues need to be considered, and how might they affect chargeback strategies?

Advanced topics include virtual networking and operational tools, along with best practices for developing your project methodology for the evolution of your virtual infrastructure.

## About Callisma

Callisma, Inc., a wholly owned subsidiary of SBC Communications, Inc., provides world-class infrastructure consulting services. Through Callisma's skilled team of professionals, we enable today's corporations to design and deploy infrastructure that will enhance their service and business performance. Our portfolio spans strategic consulting and planning through design, integration, and optimization. By providing our clients with a broad base of technical services; a flexible, results-oriented engagement style; and the highest quality documentation, Callisma delivers superior solutions on time and on budget. Callisma's expertise includes Enterprise Operating Systems, Advanced Infrastructure, Convergence, Network and Systems Operations Management, Security, Data Center, and Project Management.

# Chapter 1. The Joy of ESX

Solutions in this chapter:

- Target Audience
- VMware Terms and Concepts
- Author's Note
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions



## 1.1. Introduction

VMware produces some of the coolest software available at least that's what I think and that's why I wanted to write a book about it. Virtualization has changed the way all of us think about servers and services, backup and restores, development and testing, and particularly operating systems and applications. Many in the industry believe that virtualization will change the way most, if not all, IT shops operate and that in a few years a virtualized platform will be the platform upon which most new x86-based operating systems are deployed.

Before we get started with VMware's ESX Server, let's first review the history of VMware, the company. VMware was founded in 1998. It quickly delivered VMware Workstation in 1999. A few years later, in 2001, VMware delivered GSX Server and ESX Server, both of which brought VMware into the server world. In 2003, VMware began offering Virtual Center, which included VMotion. VMotion is an incredible awesome technology that allows you to move running virtual machines from one physical ESX Server to another without powering down the virtual server. We will discuss VMotion and Virtual Center in greater detail in [Chapter 8](#).

VMware claims nearly 3 million registered users, and at the company's VMWorld user conference, VMware stated that 5 percent of x86-based servers now run on a VMware platform.

Like many, I first started using VMware as a test platform, running virtual machines to test patches, configuration changes, service packs, and so on, but then I attended the Microsoft 2003 TechED in Dallas where many of the hands-on labs Microsoft had set up were running on VMware. Some of these labs had two domain controllers and Exchange 2003 running on a small desktop (it's funny how Microsoft now doesn't support VMware we'll get into that later in the book in the "Gotchas" chapter). Anyway, that got me thinking of the production network that I managed and soon I started architecting and piloting GSX and ESX server in our test environment and production environment with fantastic success. It wasn't long before we decided to cut the number of physical servers nearly in half. That made some management heads turn. IT was actually going to be saving money.

So what is ESX Server? Like all VMware's platform products, Workstation, GSX, and ESX, ESX Server allows you to run multiple virtualized servers or virtualized machines on one physical computer called a host. ESX Server is VMware's enterprise virtualization product. Like its brethren, GSX and Workstation, ESX allows you to run operating systems and applications utilizing idealized hardware in a virtualized environment. Each operating system runs independently of the other, so if one operating system crashes, the other virtualized operating systems are not affected. Unlike GSX and Workstation, both of which are loaded and run on a physical computer with a preinstalled host operating system (either Windows or Linux), ESX has no preinstalled operating system and is loaded directly on the physical hardware running a highly specialized Linux kernel, the vmkernel. VMware describes ESX as its own operating system, and [Chapter 9](#) discusses in greater length the vmkernel. ESX is much more customizable than GSX or Workstation and supports advanced configurations such as physical-to-virtual clustering, throttling of CPU, network bandwidth, and much more.

For systems administrators, VMware and ESX Server represent a truly revolutionary jump in the way you think about servers and services. VMware turns your server into a file. It's hard to comprehend that your server is a file it was for me anyway. But as a file, you can move servers around as simply as moving files around, cutting and pasting, using scp, FTP, or whatever tool you wish. I often burn zipped-up virtual servers to DVDs and have a library of different builds (we'll discuss this issue in depth in [Chapter 8](#) titled "Cool Tools for a Virtual Infrastructure").

Again, what VMware allows is for multiple operating systems such as Windows and Linux to run simultaneously on the same physical computer. So instead of one operating system per computer, you can run multiple and differing operating systems simultaneously, completely independent of one another. These operating systems then run the various applications and services. Why is this important? Well, for years now it has been common practice in medium- to large-sized companies to have one or maybe two applications or services running on one physical server. This, of course, has led to an enormous proliferation of servers throughout many data centers and IT departments. A majority of Callisma's customer engagements involve cost savings initiatives such as server, application, or e-mail consolidation. This proliferation, or server sprawl, adds increased IT costs in many areas. An obvious cost is the physical hardware itself. Another associated cost is administrative, which has also





## 1.2. Target Audience

This book's target audience is IT professionals who are managing a sprawling physical server environment and are looking to control both costs and headaches, architects and consultants interested in server consolidation methodology, as well as anyone who is interested in becoming more familiar with VMware and the new world of a virtual infrastructure. Historically, developers have used VMware for testing; however, VMware and ESX Server have taken hold in production data centers and have growing support in all major hardware manufacturers, such as IBM, Dell, and HP. Oracle and VMware recently announced a new development relationship, and I have worked with many third-party application companies that have designed, or are currently formalizing the design for their products to run in a virtualized environment.

Included within many of the book's following chapters are a lot of step-by-step instructions with associated screenshots that will make your mental transition from the physical world to the virtual world that much easier. By the end of the book, I hope you'll be as excited about virtualization as I am, and think it as cool as I do.



## 1.3. VMware Terms and Concepts

Conceptualizing a virtual server or virtual environment is for some the most difficult step toward moving to a virtualized environment. Virtualization has actually been around for decades in the mainframe space, but most Windows administrators have not had the opportunity or inclination to dabble in it. In the x86 world, where many IT professionals currently live and breathe, a virtual machine is a relatively new concept. As with any newer technology, there are some fundamental terms and concepts that should be defined, and we'll go over those in this chapter. Understanding each of these terms and concepts will make the rest of the book that much more accessible, so it is essential to grasp what is presented in this opening chapter. Terms beget more terms, which beget even more terms, so although there is some rationale to the order in which the vocabulary and concepts are presented here, their order often hinges more or less on how they are connected within the mind of the writers so be warned.

### 1.3.1.

#### 1.3.1.1. Host

A host is the physical computer on which you load your VMware product. ESX has an approved systems compatibility document that can be found at [http://www.vmware.com/pdf/esx\\_systems\\_guide.pdf](http://www.vmware.com/pdf/esx_systems_guide.pdf).

Prior to purchasing ESX Server or any hardware, read this document. You should refer to this guide to ensure compliance. VMware is supported on the major platforms. Nothing is more frustrating than trying to troubleshoot a hardware issue only to find out that the hardware you've purchased is not supported. Avoid this by reading the compatibility guide before acquiring any hardware.

The number of virtual machines you want to run simultaneously, as well as the types of loads these virtual machines will place on your host, will determine the physical requirements of your host. The host is what the ESX product is installed on.

When talking with virtual-newbies about virtualization, a common mistake is to equate a virtual machine with a dual-boot system which it is definitely not.

Our second term is thus virtual machine.

#### 1.3.1.2. Virtual Machine (VM)

Virtual machines are also referred to as virtual servers (not to be mistaken for the Microsoft product of the same name), guests, or the guest OS in other documentation you might read. Throughout this book, we'll use these terms interchangeably to mean the same thing. (In a Webinar for Microsoft's Virtual Server 2005 product, the tech giving the Webinar even referred to his virtual servers as VMs. I thought that was telling.)

We'll discuss the long list of guest OSes supported by VMware later in [Chapter 3](#).

For now, let's do a quick recap. The host is the physical machine on which you load your VMware platform product (ESX, GSX, or Workstation) and run your VMs or guests.

So what makes up a VM? Well, what makes up the core of any computer, really? Most computers must have at least a single processor, some physical memory, and a hard drive. And if you want it on a network, it needs a network interface card (NIC). These are four basic, or core, components for all VMs as well.

The processor(s) of your physical host are shared between your host and your VMs. We will discuss how this works in [Chapter 3](#).

There are different default levels of memory for the various supported guest operating systems when you create the configuration file ([Chapter 3](#) as well) for your VMs. These levels can be changed when the VM is built or powered off.

Memory is very important in the VMware world, and if you over-spec your host at all, do it with memory.



## 1.4. Summary

The virtual evolution of physical hardware and infrastructure has begun. It is an exciting time for IT. System administrators, architects, consultants and managers will all greatly benefit from the logical progression from the physical to the virtual. Portability, encapsulation and idealized virtual hardware make transitioning to the virtual critical. Service response, uptime, stability, testing, fault tolerance, disaster recovery are all positively impacted by moving to a virtual environment. For systems administrators this evolution is huge for meeting Service Level Agreements or simply better uptime and response. For management, virtualization will be an enormous cost-cutting tool while producing better results for services for which they are responsible.

Understanding some the concepts of a virtual world will allow you to better communicate and understand the enhancements that a virtual infrastructure offers.

We hope that readers of this book will become as excited about this, dare I say it, paradigm shift, as the writers. Because of the stability, flexibility, and power of VMware's ESX Server, virtualization of x86 servers will only proliferate and rightfully so. The technology is proven. One day the accepted model of putting a single OS running a single service will seem ridiculous. With VMware, that day is today.

## 1.5. Solutions Fast Track

### 1.5.1.

#### 1.5.1.1. Target Audience

- ☑ This book's target audience is IT professionals who are managing a sprawling physical server environment and are looking to control both costs and headaches, architects and consultants interested in server consolidation methodology, as well as anyone who is interested in becoming more familiar with VMware and the new world of a virtual infrastructure.
- ☑ Historically, developers have used VMware for testing; however, VMware and ESX Server have taken hold in production data centers and have growing support in all major hardware manufacturers, such as IBM, Dell, and HP.
- ☑ The book includes a lot of step-by-step instructions with associated screenshots that will make your mental transition from the physical world to the virtual world that much easier.

#### 1.5.1.2. VMware Terms and Concepts

- ☑ In the x86 world, where all IT professionals currently live and breathe, a virtual machine is a relatively new concept.
- ☑ An understanding of introductory terms concerning virtualization with VMware is essential.
- ☑ Since ESX Server runs on a modified Linux kernel, some very basic Linux terms and concepts are also provided.

## 1.6. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** How is virtualization like a dual-boot system?

**A:** It is nothing like a dual-boot system. A dual-boot system still runs only one OS at a time, has no idealized hardware environment, and provides no encapsulation or portability.

**Q:** What do you mean virtual machines are portable?

**A:** Because virtual machines are encapsulated in their .vmdk files and VMware provides an idealized hardware environment, virtual machines can be moved from one ESX Server to the next either using Vmotion or importing the .vmdks.

**Q:** How is VMware more stable than, say, a new physical server?

**A:** The idealized hardware provided by VMware provides a consistent virtual hardware environment that increases that stability of your virtual machines independent of the underlying physical hardware.

# Chapter 2. Installation

Solutions in this chapter:

- Pre-installation Tasks
- Performing the Installation
- Post-installation Tasks
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions



## 2.1. Introduction

Installing ESX is actually quite easy. Different installation methods exist (the trickiest being the scripted method), but overall ESX Server installation is straightforward. The pre-installation and post-installation requirements take the most time, however. In this chapter, we'll walk through the GUI-based installation method for loading ESX 2.5 as well as the tasks you need to perform before you install ESX and afterward. We'll also cover what you need to do if you run into problems during or after your installation.

### **Configuring & Implementing... Methods of Installing ESX Server**

There are three main ways to install ESX Server. We'll be discussing the graphical-based installation method in detail, but there are also the command-line and scripted methods for installing ESX Server. For more information on these methods, see the VMware ESX Server 2.5 Installation Guide at [http://www.vmware.com/pdf/esx/esx25\\_install.pdf](http://www.vmware.com/pdf/esx/esx25_install.pdf).

You should also review and become familiar with the Discussion Forum for setting up ESX Server at <http://www.vmware.com/community/forum.jspa?forumID=21>.



## 2.2. Pre-installation Tasks

Before you install ESX Server, there are some tasks you need to complete in order to ensure a successful installation. However, if you want to install, blow away, and reinstall ESX Server a few times to get used to the install method or to test certain components of your install method and document your procedures, that's fine, too. We'd highly recommend getting comfortable with the installation of ESX Server. Also, once you're ready for the production build, ensure that the method and options of your installation are well documented.

Before placing the CD in the CD-ROM drive, be sure your server meets at least the minimum specifications and does not exceed the maximum. [Table 2.1](#) outlines both the minimum and maximum specifications supported by ESX Server at the time of this writing.

<b>Table 2-1. Physical Specifications Supported by ESX Server</b>			
<b>Component</b>	<b>Min.</b>	<b>Max.</b>	<b>Notes</b>
Processor			
Intel 700MHz Pentium 3 Xeon or AMD Opteron 32-bit mode  Intel 900MHz Pentium 3 Xeon or AMD Opteron 32-bit mode	2	16	The ESX service console is bound to, and has a higher priority on, Processor 0 or the first processor. (We strongly recommend you have at least two processors.)
RAM	512MB	64GB	Allow enough memory for each virtual server you'll have plus 800MB for the service console for scalability.
Network adapters  Intel Pro/100/1000  Broadcom NetXtreme  570x Gigabit  3Com 9xx adapters	2	16	ESX supports up to eight 1GB Ethernet adapters or 16 10/100MB Ethernet adapters. (We strongly recommend using a separate NIC for the service console; thus, the minimum requirement.)
Storage adapters  Basic SCSI  Adaptec  LSI Logic			





## 2.3. Performing the Installation

This section includes step-by-step descriptions of the installation process of VMware ESX Server along with screenshots. The graphical installer method will be discussed, but two other methods, the text-mode installer and the scripted installer, can also be used. Though these two methods will not be covered in this book, more information about them can be found at the VMware Web site.

### 2.3.1. Installation of ESX Server 2.5

Now it's time to discuss the graphical installer method for setting up ESX Server 2.5. As mentioned earlier, you can also use the text-mode or scripted installation methods. The last two methods will not be covered here, but more information about them can be found on the VMware Web site.

#### 2.3.1.1. Graphical Installer

This is VMware's recommended installation method. It is by far the easiest of the three installation methods, especially for those with limited time at the command line. The graphical installer looks familiar to all of us point and clickers. It provides a proper graphical user interface (GUI) within which you can use your mouse.

There are two locations where you can install ESX Server 2.5: to the local disk or on the SAN. The installation method described in this book will be to the local disk. For information on installing ESX Server 2.5 on a SAN, please refer to the documentation at the following Web addresses:

- [www.vmware.com/pdf/esx25\\_san\\_cfg.pdf](http://www.vmware.com/pdf/esx25_san_cfg.pdf)
- [www.vmware.com/pdf/esx/esx25\\_install.pdf](http://www.vmware.com/pdf/esx/esx25_install.pdf)
- [www.vmware.com/pdf/esx/esx25\\_admin.pdf](http://www.vmware.com/pdf/esx/esx25_admin.pdf)

The reason we're not going into great detail about this method is that the method itself and some of its ramifications and caveats are worth several chapters, if not a small book.

#### 2.3.1.2. Installation Process

As we step through the graphical installer installation process, it's a good idea to write down the configuration options you choose on a piece of paper. This way you can incorporate the options you select into your server build script and have a record of how you built your ESX server. So, whatever you select, write it down.

1. As you can see in [Figure 2.1](#), when you insert the ESX Server CD into the CD-ROM and power on your server, you're presented with a number of installation choices. The default is the graphical installer, or standard installation method. (If you do not see this screen, you may need to change the BIOS of your server to boot from a CD. See your server manufacturer's documentation for details, but commonly you can do this by pressing the F2 key or Delete key.) It's a good idea to fully understand each of these choices since they are all applicable to your environment. Press **Enter**.

**Figure 2-1. The ESX Server 2.5 Installer Screen**





## 2.4. Summary

This chapter primarily covered the default and recommended method for installing ESX Server: the graphical installer method. This is a very good method for getting used to the requirements of an ESX Server build as well as the outcome of what a successful build should be. However, if you are building a large number of ESX Servers all on the same physical platform, you may want to script your installation method. If you are a command-line junkie, feel free to use the command-line method.

The important issue is getting a build process that you're comfortable with and are able to repeat, which produces a stable ESX Server build. Documenting this process is essential to the repeatability, even if you're scripting your installation. Once you're comfortable with your build process, we suggest repeating it until you have it down. By stepping through it numerous times, you may find certain changes or tweaks you may want to include. Document those and the reason for the change. This will prove handy later when reviewing the build document.

Building a solid ESX Server allows you stability, which is a main component of your new virtual infrastructure. Document the procedure and take the time to build it right.



## 2.5. Solutions Fast Track

### 2.5.1.

#### 2.5.1.1. Pre-installation Tasks

- ☒ Obtain all necessary IP information for your ESX Server
- ☒ Gather all installation media and licensing information for your ESX Server
- ☒ Set the hardware time of your ESX Server in its BIOS

#### 2.5.1.2. Performing the Installation

- ☒ We recommend going through the installation process repeatedly, changing certain default selections to become familiar with the outcome.
- ☒ Documentation of your ESX Server build for future reference is essential.
- ☒ Make your changes, break the ESX Server, and then rebuild it.
- ☒ Report errors during your install to VMware using its tools.

#### 2.5.1.3. Post-installation Tasks

- ☒ Update your ESX Server build documentation
- ☒ Ensure network connectivity of your ESX Server (using ping to another node will work)
- ☒ Remove all install media from your ESX Server

## 2.6. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** Is any installation method better than any other?

**A:** No. Whichever installation method you're comfortable with and that produces a good build is the one for you.

**Q:** I'm having problems with my installation. What should I do?

**A:** Ensure that the hardware you're loading ESX Server on is supported. The same goes for the configuration. If you are certain the hardware and configuration are supported, then run the vm-support script mentioned previously.

**Q:** Why does this chapter include only the graphical installer method?

**A:** It's the recommended installation method from VMware and one that many in the x86-world are the most comfortable with. If you wish to try it another way, VMware offers a very complete installation document covering the various installation methodologies.

# Chapter 3. Configuring a Virtual Machine, Gold Builds, and Clones

Solutions in this chapter:

- Virtual Machine Basics
- Configuring a Virtual Machine
- Gold Builds and Clones
- Modify Your Master
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions

## 3.1. Introduction

In this chapter, we'll walk through building a virtual machine, discussing its basic requirements, as well as some issues to consider when building virtuals. We'll also discuss the concept of the Gold Master, why this concept is important in building your virtual infrastructure, and how to build your VMlibrary to leverage the power and flexibility of VMware. Along the way, the beauty of the clone will be discussed. Finally, we'll explain how to modify your Gold Master to your specific environmental needs.



## 3.2. Virtual Machine Basics

Building a virtual machine is really quite similar to building a regular computer. All computers, virtual or physical, need a bare minimum of specific hardware such as a CPU (an important one), memory, disk space, and if you want it to talk with other computers on your LAN/MAN/WAN, a network interface card. Of course, to load software for instance, the operating system you may want a CD-ROM or floppy drive, at least initially.

### Designing & Planning... Hardware Virtualization Overview

The following is a brief overview (per VMware) of how virtualization of the CPU, memory, disk, and network works:

**CPU virtualization** Each virtual machine appears to run on its own CPU, or set of CPUs, fully isolated from other virtual machines, with its own registers, translation look-aside buffer, and other control structures. Most instructions are directly executed on the physical CPU, allowing compute-intensive workloads to run at near-native speed. Privileged instructions are performed safely by the patented and patent-pending technology in the virtualization layer.

**Memory virtualization** Although a contiguous memory space is visible to each virtual machine, the physical memory allocated may not be contiguous. Instead, noncontiguous physical pages are remapped efficiently and presented to each virtual machine. Some of the physical memory of a virtual machine may in fact be mapped to shared pages or to pages that are unmapped or swapped out. This virtual memory management is performed by ESX Server without the knowledge of the guest operating system and without interfering with its memory management subsystem.

**Disk virtualization** Support of disk devices in ESX Server is an example of the product's hardware independence. Each virtual disk is presented as a SCSI drive connected to a SCSI adapter. This device is the only disk storage controller used by the guest operating system, despite the wide variety of SCSI, RAID and Fibre Channel adapters that might actually be used in the system.

This abstraction makes virtual machines at once more robust and transportable. There is no need to worry about the variety of potentially destabilizing drivers that may need to be installed on guest operating systems, and the file that encapsulates a virtual disk is identical no matter what underlying controller or disk drive is used. VMware ESX Server can be used effectively with storage area networks (SANs). ESX Server supports QLogic and Emulex host bus adapters, which allow an ESX Server computer to be connected to a SAN and to see the disk arrays on the SAN.

**Network virtualization** You may define up to four virtual network cards within each virtual machine. Every virtual network card has its own MAC address and may have its own IP address (or multiple addresses) as well. Virtual network interfaces from multiple virtual machines may be connected to a virtual switch. Each virtual switch may be configured as a purely virtual network with no connection to a physical LAN, or it may be bridged to a physical LAN via one or more of the physical NICs on the host machine.

The following step-by-step process outlines how to build your virtual machine and goes into detail about the many options available when building your virtual, as well as some things you may want to consider when making your selections.

Before proceeding, the assumption here is that your ESX Server is built, powered on, and stable. See [Chapter 2](#) for more information.





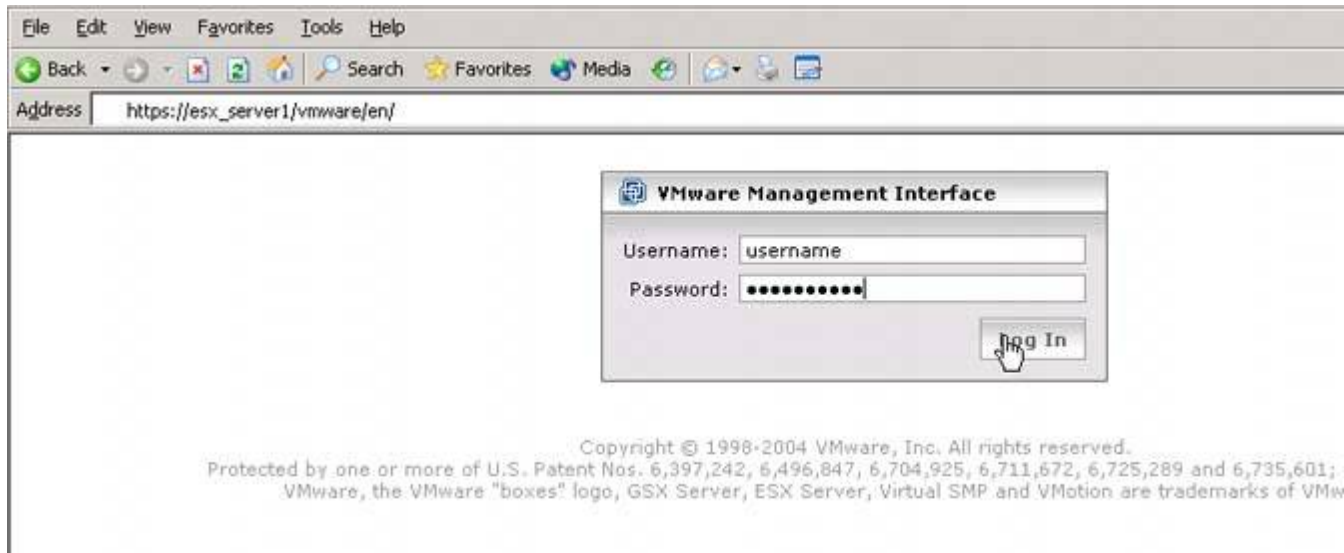


## 3.3. Configuring a Virtual Machine

Follow these steps to configure your virtual machine:

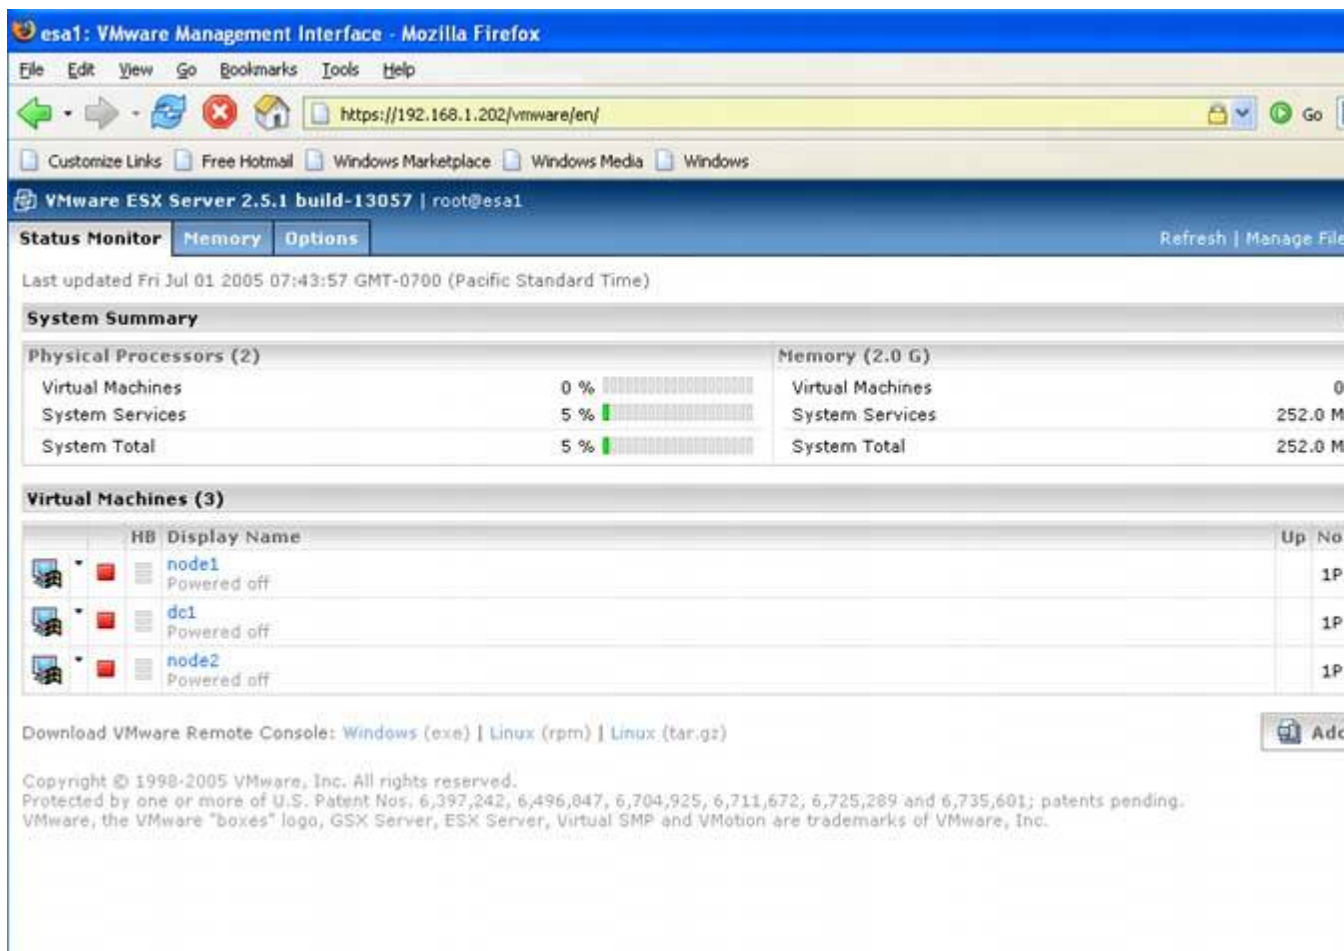
1. To begin building a virtual machine, log on to the MUI through your Web browser. You will be prompted for a username and password, as shown in [Figure 3.1](#)

**Figure 3-1. Logging on to the Management User Interface through a Web**



2. In the lower right-hand corner, you should see the **Add Virtual Machine** button. Click it (see [Figure 3.2](#))

**3. Figure 3-2. Adding the Virtual Machine Button**



As shown in [Figure 3.3](#), the **Add Virtual Machine** window should appear. Here you can choose which operating system you want for the virtual machine, its display name and, most importantly, where the configuration files are located. Every virtual machine needs its own directory. It's good practice to give the directory a name that matches the virtual machine itself.



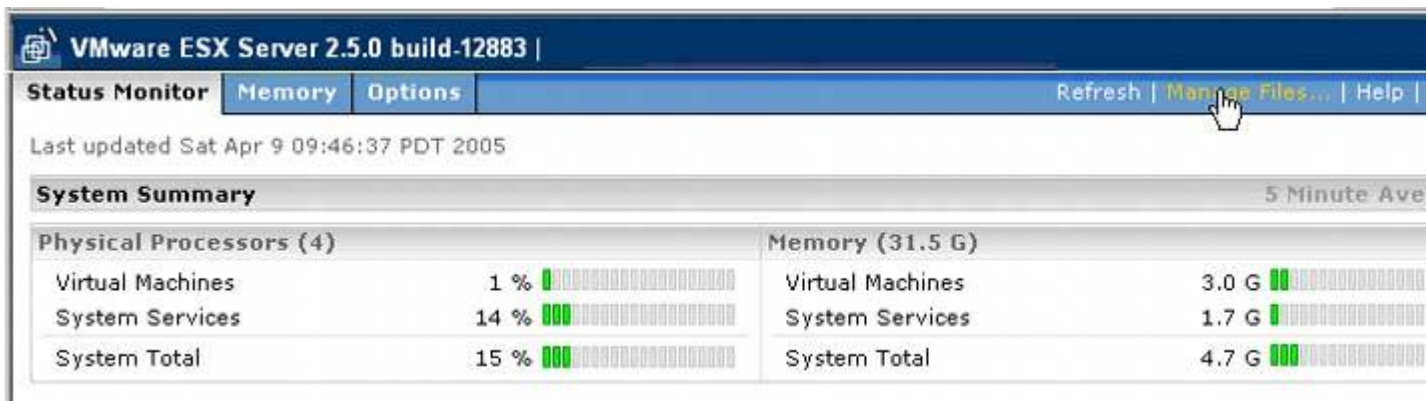


## 3.4. Gold Builds and Clones

Now that you have completed the first of your virtual machine builds, this could theoretically be the last major build you have to perform (at least for the particular operating system you just built.). This virtual machine is your *Gold Master*. Because your virtual machine is just a small collection of files, creating another virtual machine on the build of your Gold Master is as easy as copying your virtual disk file, creating a configuration file, and making some small changes. While it used to take hours, if not days or weeks, to provision a new server, you now reduced this process down to minutes. Thus, ROI begins almost immediately. Let's walk through provisioning a new server based on your Gold Master.

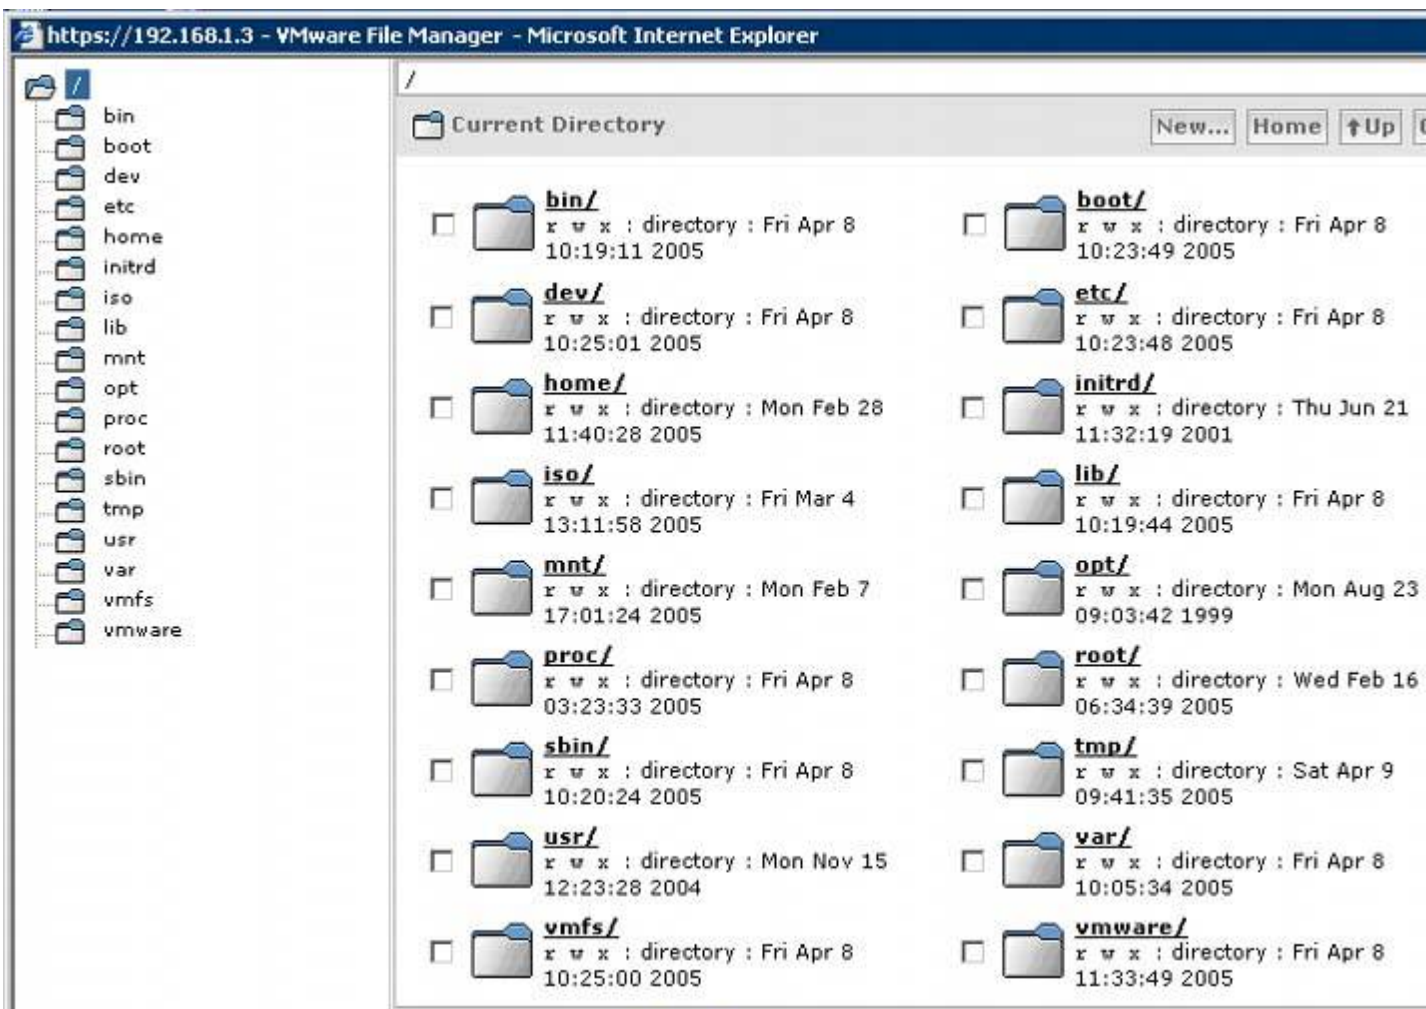
Consider your Gold Master as the parent for the rest of your virtual machines that use the same operating system. To make a copy of your Gold Master, open the MUI. Click **Manage Files** (see [Figure 3.22](#)).

**Figure 3-22. Managing Files via MUI**



Here, you're accessing the file system of your ESX server, something which you'll soon become very familiar with (see [Figure 3.23](#)).

**Figure 3-23. The ESX File System**





## 3.5. Modify Your Master

Once you have a well-tested and trusted Gold Master, guard it well. As mentioned earlier in this chapter, it's a good idea to make a copy of it for safekeeping. You can use another copy of your Gold Master to test security patches, hotfixes, and service packs. Make sure you test any software package or configuration change on a copy of your Gold Master before applying the change to it. Of course, over time your Gold Master will have to change to keep up with operating system improvements as well as VMware virtual hardware improvements, but VMware allows for enormous flexibility when it comes to testing.

When making changes to your Gold Master, we recommend first putting the virtual disk in Undoable Mode. If you remember, this is the disk mode option that allows you to apply or discard changes when the virtual machine is powered down. If for some reason something goes wrong with your testing and you need to get back to the state your virtual machine was in prior to applying the software or configuration change, Undoable Mode allows you to revert to the state before the change. However, if the changes you make to your virtual machine are successful and you want to apply them, Undoable Mode allows for this as well. This snapshot-like functionality is a tremendous time-saver and provides excellent and flexible testing and restore possibilities.

## 3.6. Summary

Building virtual machines is an interesting process that you should practice, test, and document. Once you're confident you have a good virtual machine build for a specific operating system, make a copy of it and save the copy as your Gold Master. You can use this for all future-based builds of the operating system. It will increase your turnaround time for provisioning servers, ensure the quality of the build is always the same, and immediately start realizing some of the ROI of a virtual infrastructure.

Changes to your Gold Master should be thoroughly tested and documented prior to committing the changes permanently. The easiest way to do this is to place the virtual disk of the Gold Master in Undoable mode prior to your testing so you can discard any changes that go bad.

Create a library of builds with the different operating systems you support or are testing. You can also create Gold Master virtual machines that are application specific. Store your virtual in a VMlibrary which should house all the tools, Gold Masters, ISO images, and other materials used within your virtual infrastructure.

## 3.7. Solutions Fast Track

### 3.7.1.

#### 3.7.1.1. Virtual Machine Basics

- ☑ Building a virtual machine is really quite similar to building a regular computer.
- ☑ All computers, virtual or physical, need a bare minimum of specific hardware such as a CPU (an important one), memory, disk space, and if you want it to talk with other computers on your LAN/MAN/WAN, a network interface card.
- ☑ Although a contiguous memory space is visible to each virtual machine, the physical memory allocated may not be contiguous.

#### 3.7.1.2. Configuring a Virtual Machine

- ☑ Every virtual machine needs its own directory. It's good practice to give the directory a name similar to the virtual machine itself.
- ☑ All running virtual machines need to be on a VMFS volume. The VMFS is a specialized file system used especially for virtual machines.
- ☑ Disk mode options are definitely one of the more flexible and powerful elements of a virtual machine. Options include Persistent, Nonpersistent, Undoable, and Append Modes.

#### 3.7.1.3. Gold Builds and Clones

- ☑ The first virtual machine that you build is your Gold Master.
- ☑ Consider your Gold Master as the parent for the rest of your virtual machines that use the same operating system.
- ☑ Because your virtual machine is just a small collection of files, creating another virtual machine based on the build of your Gold Master is as easy as copying your virtual disk file, creating a configuration file, and making some small changes.

#### 3.7.1.4. Modify Your Master

- ☑ Make sure you test any software package or configuration change on a copy of your Gold Master before applying the change to it.
- ☑ Over time your Gold Master will have to change to keep up with operating system improvements as well as VMware virtual hardware improvements, but VMware allows for enormous flexibility when it comes to testing.
- ☑ When making changes to your Gold Master, we recommend first putting the virtual disk in Undoable Mode.



## 3.8. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** When I copy my virtual machine and try to run it on my network, I get hostname already exists and IP address already exists errors. I thought I could simply copy my virtual machine. What is the problem?

**A:** You *can* copy your virtual machine, which is why you're running into this problem. The copy is an exact copy. Thus, you need to change the hostname, IP address, and even the computer's SID (if it has one).

**Q:** Can I move my Gold Master to my other ESX Servers?

**A:** Yes, it's a recommended time-saver. If your ESX Servers share a LUN, that too is a way to access your Gold Master or your VMlibrary for that matter.

**Q:** Are there any limitations to the number of virtual machines I can run on my ESX Server?

**A:** Yes. Each ESX Server allows for only 80 virtual CPUs and 200 registered virtual machines.

**Q:** Can I take a virtual machine running on VMware Workstation or GSX and run it in ESX?

**A:** It depends. The process is not as simple as cutting and pasting. There is a specific import process required for migrating virtual machines from either Workstation or GSX into ESX. Access the following link for the exact instructions:  
[http://www.vmware.com/pdf/mobility\\_guide.pdf](http://www.vmware.com/pdf/mobility_guide.pdf).

# Chapter 4. Advanced Configurations

Solutions in this chapter:

- Network Traffic Shaping
- Disk Bandwidth
- CPU Management, Affinity
- NUMA and Processor/Memory Affinity
- Raw Device Mapping
- Clusters
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions

## 4.1. Introduction

In this chapter we review some of the advanced configuration possibilities for both your ESX Server and your virtual machines. VMware provides tremendous flexibility with configuration possibilities. We look at network bandwidth shaping and disk shaping as well as different memory configuration possibilities. We also examine your ability to set CPU affinity and finally explore clustering options. We end with an explanation of how to set up a cluster in a box and a complete walkthrough of setting up a Windows 2003 Cluster Service.

## 4.2. Network Traffic Shaping

ESX Server allows you to control certain aspects of network traffic on your virtual machines. ESX Server can control three areas of outgoing network traffic:

- Average bandwidth
- Peak bandwidth
- Burst size

To enable network traffic shaping, open the details page of the virtual server you want to configure. Click the **Network** tab. Mark the check box to **Enable Traffic Shaping** and place the values you want in the appropriate boxes (see [Figure 4.1](#)).

**Figure 4-1. Enabling Network Traffic Shaping**



The values you choose can be bits per second, kilobits per second, or megabits per second. Behind the scenes, this selection essentially enables the nfshaper module. To read more about nfshaper, at the command line of your ESX Server type:


```
#man nfshaper
```

## 4.3. Disk Bandwidth

If you are running virtual servers that require heavy access to diskssuch as a file, messaging, or database serveryou may want to modify the bandwidth allocation values for that virtual machine. VMware ESX Server uses an "allocation policy" based on proportional shares. The more shares a virtual machine has on a disk, the more access to the disk the virtual machine will have. Shares can be set numerically, such as 100, 200, and so on, or by setting the share value to Low, Normal, or High.

To set the share value, open the MUI and right-click the virtual machine you want to modify and then select **Properties**. On the **Disk** tab of your virtual machine, click **Edit**. Type the value of disk shares you want this virtual machine to have. By default, it's set to 1,000 (see [Figure 4.2](#)).

**Figure 4-2. Disk Shares**



Depending on the type of virtual machine you are running, you may or may not want to adjust this value. A file server, for example, may need greater access to the disk then say a Web server. Whatever change you make here, ensure it's documented.



## 4.4. CPU Management, Affinity

VMware ESX Server allows you to control access to the CPUs on which the Service Console and your virtual machines run. This management is in the form of minimum and maximum percentage values, as well as shares.

A minimum percentage value is the minimum amount of time that either the Service Console or the virtual machine is guaranteed on a particular processor. The maximum percentage value is the maximum amount of time that the Service Console of a virtual machine is guaranteed on a particular processor or processors if you are running virtual SMP.

As shown in [Figure 4.3](#), a default build of one of our ESX Servers gave the Service Console the following access to Processor0.

**Figure 4-3. Default CPU Access for the Service Console**



The values located under the Resources heading show that the Service Console will receive a minimum of 8 percent availability of the processor capacity while the most, or the maximum, it can utilize is 100 percent.

### Designing & Planning... CPU Percentage Values

Minimum CPU percentage establishes the lowest limit of a CPU that a virtual machine will always be able to utilize. If the virtual machine cannot reserve its minimum CPU percentage, it will not be able to power on. The admission control policy enforces this, guaranteeing the minimum CPU allocation.

Maximum CPU percentage establishes the highest limit of consumption the virtual machine can have on a CPU. The virtual may never exceed this usage since it is a fixed allocation, and can not be exceeded even if there are idle CPU cycles available.

The Service Console is also configured with 2,000 shares. Shares represent a relative priority. This priority denotes which virtual machines will have priority or preferential access to the CPU.

VMware describes shares in the following manner: "the shares value represents a relative metric for allocating processor capacity, where this value is compared to the sum of all shares of all virtual machines on the server and the Service Console."

In our example of the Service Console, it has been allotted 2,000 shares on processor0. Since the Service Console has been allocated 2,000 shares, it has priority access on processor0 over a virtual machine, which default is configured to Normal.





## 4.5. NUMA and Processor/Memory Affinity

ESX Server supports physical hosts that are of Non-Uniform Memory Access (NUMA) architecture. Each NUMA machine can be comprised of multiple nodes that contain the processors, memory cache, and so on. Running a virtual machine's resources locally on one of these nodes without having to access resources remotely or on a different NUMA node increases performance, and that's where affinity comes in.

By default, ESX Server utilizes all NUMA nodes equally, balancing virtual machines between each node by using local memory when possible for VMs running on the same node. ESX Server designates every virtual machine as a NUMA node the CPUs of which the virtual machine will only run on. According to VMware documentation, ESX Server "periodically...compares the utilization levels of all NUMA nodes and attempts to 'rebalance' the nodes if one node has a higher utilization level than the other nodes. ESX Server rebalances the nodes by changing a virtual machine's 'home' NUMA node from the overutilized node to the underutilized node. "This is an automatic process. However, you can also tweak NUMA settings yourself.

Since ESX automatically tries to balance virtual machines on NUMA nodes, you may never have to manually set NUMA options. However, if you find it necessary, you can set both CPU and memory affinity. Under the CPU and Memory tabs of your virtual machine, select the affinity for the virtual machine, similar to what was shown previously in this chapter.

## 4.6. Raw Device Mapping

New in ESX Server 2.5, raw device mapping allows you to access a raw physical device such as a SAN LUN. Through the use of a symbolic link in the VMFS partition, a mapping is created to the raw device. A common use for mapping raw devices is to support Microsoft Cluster Service, which we will discuss in the next section. But first let's drill down a little bit more into raw device mapping.

Prior releases of ESX Server supported accessing raw disk, but if changes were made, say, adding an additional HBA or SCSI card and the order in which ESX Server read these cards changed, disk resources were made unavailable. Through tools like `vmkpcidivyi` you could generally recover, but it was a pain, and you suffered obvious downtime.

In ESX Server 2.5, raw device mapping helps with this. So what is a raw device mapping? As mentioned already in this chapter, it is basically a symbolic link (or mapping file) that is in the VMFS partition and has a `.vmdk` extension. The mapping file includes metadata that is used to "manage and redirect disk accesses to the physical device," according to VMware documentation.

If you are clustering using Microsoft Cluster Service or taking SAN snapshots in virtual machines, VMware recommends using raw device mapping for those VMs.

Two compatibility modes exist for utilizing device mapping: virtual and physical.

Virtual mode makes raw device mapping act as a traditional `.vmdk`. This mode allows options such as redo logs and advanced file locking, and allows for greater portability over SAN hardware.

Physical mode allows greater flexibility with SAN software, allowing for SAN snaps in the virtual machine. This mode also provides the option of virtual to physical clustering.

See the following white paper for more details on raw device mapping at [www.vmware.com/pdf/esx25\\_rawdevicemapping.pdf](http://www.vmware.com/pdf/esx25_rawdevicemapping.pdf).



## 4.7. Clusters

It's possible to cluster your virtual servers, which run applications that are cluster-aware such as Exchange or MSSQL.

This functionality has been enhanced with the use of raw device mapping (see the previous section for details). Clusters provide a high-availability configuration between two or more servers working together as one. Redundancy on the software and/or hardware level allows for the high availability.

### 4.7.1. Cluster Types

There are a number of clustering solutions you can consider, such as the following:

- Cluster in a Box
- Cluster across Boxes
- Consolidated Cluster
- Cost-Effective Standby Host
- Network Load Balancing
- Geo-clusters

#### 4.7.1.1. Cluster in a Box

Cluster in a box allows for two or more virtual machines running on the same ESX Server to run a cluster-aware application. This type of cluster provides redundancy for software failures and administrative errors.

#### 4.7.1.2. Cluster across Boxes

A cluster across boxes allows for two or more virtual machines running on separate ESX Servers and provides software and hardware redundancy.

#### 4.7.1.3. Consolidated Cluster

A consolidated cluster configuration basically allows you to consolidate a number of hardware clusters onto a single physical ESX Server. Let's say you have three clusters with two nodes each. By consolidating your clusters onto a single ESX Server, you reduce the number of physical servers you need by two-thirds, going from six to two.

#### 4.7.1.4. Cost-Effective Standby Host

A cost-effective standby cluster configuration lets you place all of your failover nodes on a single ESX Server while maintaining your primary nodes on physical servers. This allows for a physical server consolidation rate of nearly 50 percent.

#### 4.7.1.5. Network Load Balancing

VMware permits the deployment of parallelization clustering utilizing Microsoft's Network Load Balancing.

This type of clustering helps you scale a service, such as VPN or Web, by balancing and distributing the traffic amongst the virtual servers that are a part of the cluster so that the load is shared equally.

#### 4.7.1.6. Geo-Clustering

Geo-clustering, or geographically distributed clustering, allows for cluster nodes to exist in remote sites. VMware has published a whitepaper discussing how this is implemented with the EMC Symmetrix Remote Data Facility (SRDF). Neither writer of this book has implemented this solution yet, but we strongly suggest reading the whitepaper to see what is available now and to imagine what the future holds. You can find the white paper at [www.vmware.com/pdf/SRDF\\_wp\\_eng.pdf](http://www.vmware.com/pdf/SRDF_wp_eng.pdf).

### 4.7.2. Building a Two-Node Cluster in a Box with Windows 2003 Server, Enterprise Edition

We will build the most basic of clusters: a cluster in a box. This is the cluster type you build when you go to



## 4.8. Summary

For most servers, the default configuration VMware places on your virtual machines is adequate. However, for virtual machines that need their own CPU or greater access to disks, you may find it necessary to configure more advanced configurations for your virtual machines. It's recommended that you become familiar and comfortable with making advanced configuration changes so you can troubleshoot performance issues if you experience them.

For those who need high availability for their virtual machine, VMware allows for myriad clustering configurations. For software and hardware redundancy, VMware makes the setting up and deployment of such solutions easy.



## 4.9. Solutions Fast Track

### 4.9.1.

#### 4.9.1.1. Network Traffic Shaping

- ☒ ESX Server controls three areas of outgoing network traffic: average bandwidth, peak bandwidth, and burst size.
- ☒ To control network traffic, VMware allows you to control the bandwidth of your virtual machines. This is easily configured and simple to modify.
- ☒ To enable network traffic shaping, open the details page of the virtual server you want to configure.

#### 4.9.1.2. Disk Bandwidth

- ☒ Disk bandwidth allows you to control disk access by your virtual machines. Based on proportional shares, you can manage the amount of time your machines have on disk.
- ☒ VMware ESX Server uses an "allocation policy" based on proportional shares. The more shares a virtual machine has on a disk, the more access to the disk the virtual machine will have.
- ☒ Shares can be set numerically such as 100, 200, and so on or by setting the share value to Low, Normal, or High.

#### 4.9.1.3. CPU Management, Affinity

- ☒ CPU Affinity allows you to control which CPU your virtual runs on. If your virtual machine needs its own CPU(s), VMware easily allows for this.
- ☒ This CPU management is in the form of minimum and maximum percentage values as well as shares.
- ☒ The Service Console is also configured with 2,000 shares. Shares represent a relative priority.

#### 4.9.1.4. NUMA and Processor/Memory Affinity

- ☒ ESX Server supports physical hosts that are of Non-Uniform Memory Access (NUMA) architecture.
- ☒ NUMA and Processor/Memory Affinity allow for NUMA systems to be managed easily. ESX Server automatically balances resource allocation between NUMA nodes and allows you to manually control on which NUMA nodes your virtual machines live.
- ☒ By default, ESX Server utilizes all NUMA nodes equally, balancing virtual machines between each node using local memory when possible for VMs running on the same node.

#### 4.9.1.5. Raw Device Mapping

- ☒ Raw device mapping lets you access a raw physical device such as a SAN LUN.
- ☒ Raw device mapping allows for more reliable and easily configurable access to raw disks. Through the use of symbolic links, VMware has made the management and use of raw devices easier for the administrator.
- ☒ A common use for mapping raw devices is to support Microsoft Cluster Service.

#### 4.9.1.6. Clusters

- ☒ Clusters provide a high-availability configuration between two or more servers working together as one.
- ☒ VMware provides a number of ways to configure highly available and fault-tolerant virtual machines utilizing clustering technologies.
- ☒ VMware allows for the deployment of parallelization clustering utilizing Microsoft's Network Load Balancing.





## 4.10. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the **"Ask the Author"** form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** Do you need to tweak the settings of your virtual machines after you've built them?

**A:** No. VMware has configured the defaults of your build to suit most server loads. However, depending on the profile of your server and the resource intensiveness of the service it's providing, you may want to consider tweaking some of the default settings to better ensure service stability and reliability.

**Q:** Are there more advanced features than the ones detailed in this chapter?

**A:** Yes... a lot more.

**Q:** Is clustering more stable in VMware?

**A:** We find that the idealized hardware, in combination with VMware's improvements in clustering support, has made virtual clustering very stable.

**Q:** Can you Vmotion a cluster?

**A:** That depends on the cluster type and how it is configured. Clusters using Raw Device Mapping in virtual mode can be moved with Vmotion according to VMware documentation.

# Chapter 5. Virtual Networking

Solutions in this chapter:

- Virtual Switches
- Physical and Virtual NICs
- MAC Addresses
- Port Groups and VLANs
- Networking Tools
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions

## 5.1. Introduction

This chapter expands on the basics of virtual networking that were introduced in [Chapters 1](#) and [2](#). We'll discuss virtual switches, physical and virtual NICs, bonded physical NICs, MAC addressing, Port Groups and VLANs, networking tools, and firewall possibilities. We'll go beyond the basics to discuss advanced network configurations and possible solutions. ESX provides remarkable flexibility when it comes to networking. This chapter is essential for the beginner because it offers the basic knowledge needed to set up ESX Server and your virtual machines. We'll build upon the basics to more advanced topics that both the beginner and possibly the advanced ESX administrator will find useful. Before moving on to the next chapter, though, we recommend you fully understand all of the information in this chapter because it will be referenced in later chapters, thus testing your understanding of the pertinent information detailed here.

After you've installed ESX Server on your Host, one of the first things you'll need to do when you log into the MUI is set up a virtual switch. You can think of a virtual switch as a software hub controlled by ESX that routes the traffic of your virtual machines both internally, between virtual machines on the same physical host, as well as externally, to your production network or the Internet.

Many possible networking configurations can enhance network communication, provide redundancy, and increase security for your production of virtual machines.



## 5.2. Virtual Switches

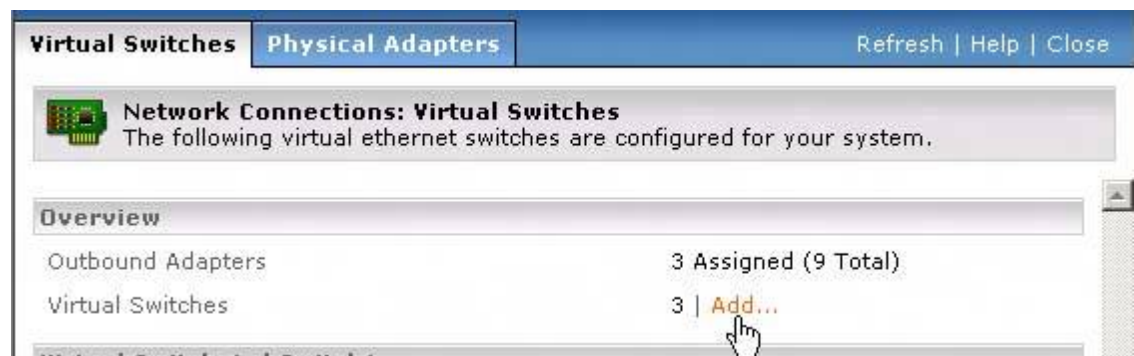
Like a physical switch or hub, a virtual switch has a certain number of ports into which you can plug the virtual machines. A virtual switch has 32 logical ports; thus, it can support a maximum of 32 virtual machines connected to it.

### 5.2.1. VMnets

You may set up a virtual switch called a VMnet exclusively for routing high-speed network traffic between virtual machines on the same ESX Server (see [Figure 5.1](#)). VMnets can provide added security and isolation of network traffic and provide a more flexible and cost-effective network topology. VMnets can be used for advanced configurations such as parallelization (NLB) or redundant (MSCS) cluster solutions, as well as interesting firewall possibilities. Creating a VMnet is quite simple. Just follow these steps:

1. Open the MUI.
2. Click the **Options** tab.
3. Click **Network Connections...**
4. In the **Overview** section, click **Add**.

**Figure 5-1. Configuring a VMnet**



5. In the **Network Label** property (see the details later in this section), type in a label such as VMnet0 or something more descriptive (perhaps heartbeat if this VMnet will be used for the heartbeat traffic of a cluster).
6. Do not check any of the network adapters to bind to this virtual switch. Click **Create Switch**.
7. Find the newly created virtual switch. It should look something like the image in [Figure 5.2](#)

**Figure 5-2. A Virtual Switch**



8. Notice the text: **No outbound adapters. Traffic will be routed locally.** Virtual machines connecting to this VMnet virtual switch will be able to route network traffic only within the ESX Server itself to those virtual machines that connect to it.

If you were to look at the configuration file of a virtual machine that was connected to a VMnet, the Ethernet configuration would look like the screenshot in [Figure 5.3](#)

**Figure 5-3. The Configuration File Information of a Virtual Machine**

```
Ethernet1.present = "TRUE"
Ethernet1.connectionType = "monitor_dev"
Ethernet1.virtualDev = "vmxnet"
Ethernet1.devName = "vmnet_0"
Ethernet1.networkName = "VMnet0"
```








## 5.3. Physical and Virtual NICs

Within a virtual environment physical NICs belong to your ESX server and can be bound to virtual switches service console. There are also virtual NICs which attaché virtual machines to virtual switches.

### 5.3.1. Physical NICs

Your ESX Server should have a minimum of two physical NICs. During the ESX Server installation process, one of the physical NICs will be allocated exclusively to the ESX console. It is through this ESX Server Service Console NIC that all connections to the MUI are made, as well as SCP, SSH, or any other tool to access the ESX Service console. The other physical NIC will be dedicated to your virtual machines. [Figure 5.12](#) depicts this simple configuration.

**Figure 5-12. The Configuration of an ESX Server with Two NICs**



[Figure 5.12](#) depicts an ESX server with two NICs: one dedicated to the ESX Server's Service Console and the other dedicated to the virtual machines on the ESX Server itself. How ESX Server sees and names these NICs from the operating system perspective is important to understand.

When you install ESX Server, the first NIC in the first PCI bus/slot that ESX detects will be the one it assigns to the ESX console. Physical NICs that are not allocated to the Service Console but are for use by the VMkernel are labeled *vmnic0* and incremented from there as *vmnic1*, *vmnic2*, and so on. The name of the NIC will be derived from the PCI bus/slot location. If your ESX Server has different physical NICs, say Broadcom and Intel, the name of your physical NICs will depend on which NIC driver is loaded first. If you have a multifunction card, say a dual- or quad-port NIC, then the function is also numbered. For instance:

- 1 one port NIC might be named *bus1. slot1. function1*
- 1 dual port NIC might be enumerated with the following names:

*bus1.slot1.function0*  
*bus1.slot1.function1*





## 5.4. MAC Addresses

As you probably know, every physical NIC has its own unique MAC address. This is burned into the NIC by its manufacturer prior to shipping. What you may or may not know is that every virtual NIC has its own unique MAC address as well. A virtual NIC's MAC address can be changed manually if necessary.

### 5.4.1. MAC Address Overview

The MAC address is a six-byte hexadecimal number that looks something like this: 00-09-6B-BF-6D-31.

The first three bytes are a unique Vendor ID for the manufacturer of the network adapter. This is called the organizationally unique identifier or OUI. In the MAC address example, the OUI would be 00-09-6B, which is IBM.

Although VMware does not manufacture a physical network adapter, it does provide virtual ones. Thus, VMware has its own OUI; in fact, it has two.

The first VMware OUI is for automatically generated MAC addresses, and in most cases these are the MAC addresses your virtual machines will use. The second VMware OUI is used for MAC addresses that are manually set.

If you were to vi (see [Chapter 13](#)) one of your virtual machine's configuration files, it might look something like the image in [Figure 5.22](#)

**Figure 5-22. A vi of a Virtual Machine's Configuration Files**

```
floppy0.fileType = "file"

ide0:0.startConnected = "TRUE"
Ethernet0.addressType = "generated"
uuid.location = "56 4d 08 50 15 ba b5 cb-98 52 ba 9d fa c8 7f 88"
uuid.bios = "56 4d 08 50 15 ba b5 cb-98 52 ba 9d fa c8 7f 88"
ethernet0.generatedAddress = "00:0c:29:c8:7f:88"
ethernet0.generatedAddressOffset = "0"
tools.syncTime = "FALSE"
```

The third line down, "Ethernet0.addressType =" shows that the MAC address, also known as the Ethernet address or physical address, has been automatically generated. The sixth line down, "ethernet0.generatedAddress =" shows that the MAC address in this example is 00:0c:29:c8:7f:88. The first three bytes, 00:0c:29, represent VMware's OUI for automatically generated MAC addresses. Every MAC address that is automatically generated in ESX (as well as GSX or Workstation) will have this OUI value. However, for the last three bytes, ESX Server uses a specific algorithm to generate MAC addresses.

### Configuring & Implementing... VMware's Algorithm for MAC Address Generation

The following is taken from the ESX Server 2.5 Administration Guide and describes how VMware's algorithm for MAC address generation works:

"We use the VMware UUID (Universally Unique Identifier) to generate MAC addresses. We then check for any conflicts. If there is a conflict, we add an offset and check again, until there is no conflict. (The VMware UUID is based on the path to the virtual machine and the host's SMBIOS UUID.)

Once the MAC address has been generated, it does not change, unless the virtual machine is moved to a different location for example, a different path on the same server or a different ESX Server machine. We save the MAC address in the configuration file of the virtual machine.





## 5.5. Port Groups and VLANs


When you go to configure a virtual switch, you'll see in the same window the area in which you can configure a port group. So, what is a port group?

Port groups are VLANs basically. VLANs, or virtual local area networks, are defined by the IEEE 802.1Q standard.

A VLAN is a group of network devices on a LAN that is configured to communicate as if they were attached to the same LAN segment and were in the same broadcast domain. In reality, however, they can be geographically dispersed and connected to a number of different LAN segments. This is accomplished through VLAN tagging. VLAN tags, though, can be switch port based, MAC address based, protocol based, or policy based. A VLAN tag is an extra four bytes inserted after the source and destination address in an Ethernet frame.

For more information on 802.1Q and VLAN tags, go to <http://standards.ieee.org/getieee802/download/802.1Q-2004.pdf>. Under the USER TYPE drop-down, we input Systems Administrator, clicked Accept/Begin Download, and got the result shown in [Figure 5.24](#). See section 9 of this document for VLAN tag format information.

**Figure 5-24. User Type Drop-Down Menu**




In ESX Server, there exist three options for VLANs:

- External Switch Tagging (EST)
- Virtual Switch Tagging (VST)
- Virtual Guest Tagging (VGT)

Exterior Switch Tagging (EST) mode relies exclusively on the tagging of network traffic by the physical switch (see [Figure 5.25](#)). All traffic goes through one of your ESX Server's physical NICs, which is attached to a physical switch providing VLAN tagging on outgoing traffic and strips the tag for incoming traffic. The limitation to this method is that it's switch port based so the number of VLANs your ESX Server can support is limited to the number of physical NICs you have in your ESX Server.

**Figure 5-25. Exterior Switch Tagging**









## 5.6. Networking Tools

To assist with network troubleshooting and configuration, there are a number of tools, mostly run from the command line, that you can use.

### 5.6.1. service network restart

Use the *service network restart* command if you're having problems with the Service Console NIC, perhaps regarding MUI responsiveness (service *servicename* restart can be used to restart any service, just substitute *servicename* for the actual service name). This restarts eth0, your service console NIC, as well as the lo, which is the local loopback. From the command line of your ESX Server, type **service network restart**.

The output should look something like the image in [Figure 5.33](#)

**Figure 5-33. The Output of *service network restart***

```
[root@esx1 root]# service network restart
Shutting down interface eth0: [ OK ]
Setting network parameters: [ OK ]
Bringing up interface lo: [ OK ]
Bringing up interface eth0: [ OK ]
[root@esx1 root]#
```

This command bounces the network service.

### 5.6.2. findnic

The VMkernel labels the physical adapters it's given for virtual machines vmnic0, vmnic1, vmnic2, and so on. The lower the number, the lower the physical network adapter is in the PCI bus/slot hierarchy. If you have many physical NICs in your ESX Server, it may be difficult for you to map where physically, say, vmnic9 resides. This can be especially true if you have dual- or quad-port network adapters. *findnic* can assist you in finding and labeling the physical adapter from the VMkernel nic name.

The *findnic* command allows you to bind an IP address to a vmnic for instance, vmnic1 and then ping another IP address. From either the success of the ping or from the blinking lights on the physical adapter, you should be able to determine what physical NIC is vmnic1.

From the command line of your ESX Server and type **findnichelp**.

The following are your options:

Usage: findnic [-f --flood] [-i --interval seconds] nicname local\_ip remote\_ip.

Thus, an example would be **findnic vmnic1 192.168.1.203 192.168.1.202**.

The syntax in the previous example will bind IP address 192.168.1.203 to vmnic1 and then constantly ping 192.168.1.202.

The output should look like the image in [Figure 5.34](#)

**Figure 5-34. The Output of *findnic vmnic1 192.168.1.203 192.168.1.202***

```
[root@esx1 root]# findnic vmnic0 192.168.1.203 192.168.1.202
Warning: loading /tmp/vmxnet_nic_test will taint the kernel: no license
Checking eth1
PING 192.168.1.202 (192.168.1.202) from 192.168.1.202 : 56(84) bytes of data.
64 bytes from 192.168.1.202: icmp_seq=0 ttl=255 time=14 usec
64 bytes from 192.168.1.202: icmp_seq=1 ttl=255 time=19 usec
64 bytes from 192.168.1.202: icmp_seq=2 ttl=255 time=13 usec

--- 192.168.1.202 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/mdev = 0.013/0.015/0.019/0.004 ms
```



## 5.7. Summary

VMware provides incredible flexibility when it comes to networking. VMnets provides network communication between virtual machines on the same ESX Server and when used with VMnics can provide a much more secure and cost-effective solution for network access to services. By allowing the administrator to easily control every aspect of network configurations down to customizing MAC addresses, virtual networking is another very cool component that comes from evolving into a virtual environment. Learn it well.

## 5.8. Solutions Fast Track

### 5.8.1.

#### 5.8.1.1. Virtual Switches

- ☑ Like a physical switch or hub, a virtual switch has a certain number of ports into which you can plug the virtual NIC of your virtual machines.
- ☑ Virtual switches allow you to create VMnets and VMnics, both of which allow your virtual machines to pass network traffic. VMnets pass this traffic only within the ESX Server itself, while VMnics are bound to at least one physical NIC, which provides the ability to pass network traffic to your corporate intranet, the Internet, and other virtual machines.
- ☑ A virtual switch called a VMnet may be set up exclusively for routing highspeed network traffic between virtual machines on the same ESX Server.

#### 5.8.1.2. Physical and Virtual NICs

- ☑ You should have a minimum of two NICs for your ESX Server: one for the Service Console and the other dedicated to your virtual machines. It's through the physical NIC dedicated to your virtual machines that a virtual NIC can be attached to pass network traffic.
- ☑ When you install ESX Server, the first NIC in the first PCI bus/slot that ESX detects will be the one it assigns to the ESX console.
- ☑ Each virtual machine can have four virtual NICs.
- ☑ Like any physical NIC, a virtual NIC can have one or more IP addresses and a MAC address.

#### 5.8.1.3. MAC Addresses

- ☑ MAC addresses are, by default, automatically assigned; however, VMware lets you manually assign them as well.
- ☑ The MAC address is a six-byte hexadecimal number.
- ☑ A recommendation you might want to follow if you are going to set MAC addresses manually is to pick a range of MAC addresses, say, 00:55:56:01:a1:zz, and only increment the last bit until you run out of values

#### 5.8.1.4. Port Groups and VLANs

- ☑ VMware provides for three types of VLAN or Port Groups: External Switch Tagging (EST), Virtual Switch Tagging (VST), and Virtual Guest Tagging (VGT).
- ☑ Each port group represents one VLAN.

#### 5.8.1.5. Networking Tools

- ☑ A few command-line tools can assist you in managing the network infrastructure of your virtual environment, such as findnic, mii-tool, and others.
- ☑ Use the service network restart command if you're having problems with the Service Console NIC
- ☑ The findnic command allows you to bind an IP address to a vmnic.
- ☑ mii-tool can be used to modify the NIC settings of your virtual server



## 5.9. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the **"Ask the Author"** form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** Does VMware support NIC teaming?

**A:** Yes, and it's a very good idea to configure your ESX Server, especially if it's for production, with teamed NICs. This will provide hardware fault tolerance in case one NIC fails.

**Q:** How many physical NICs do I need on my ESX Server?

**A:** That depends on the number of virtual machines and the network traffic they produce. You should have a minimum of two: one for your Service Console and one dedicated to your virtual machines.

**Q:** If I'm building a cluster using Microsoft Clustering Service, what's the best configuration for my heartbeat NIC?

**A:** For a Cluster in a Box, create a VMnet that your nodes of your cluster can attach their heartbeat NICs.

**Q:** Can you attach virtual machines on any ESX Server to a VMnet?

**A:** No. Only virtual machines on the ESX Server that the VMnet resides can attach to it.

# Chapter 6. Physical-to-Virtual Migrations

Solutions in this chapter:

- P-to-V Terms
- cat Your Hard Drive
- BartPE
- VMware P2V Assistant 2
- Platespin PowerP2V
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions



## 6.1. Introduction

In this chapter, we discuss and walk through the concepts and methodologies for migrating your physical machines into a virtual environment. We also discuss the pros and cons of each methodology as well as the requirements and considerations for each.

We examine both free methods to perform a p-to-v migration as well as tools you can purchase to assist you. The essential component of the method or product you choose is that you are comfortable in performing it and attain the result you set out for which, of course, is a smooth and successful migration.

## 6.2. P-to-V Terms

The following are P-to-V terms that we'll be using throughout this chapter:

- **Source Server** The source server is the physical server you migrate into your virtual infrastructure.
- **Target Server** The virtual server that migrate your source server into.
- **Swing Server** Where you place the images of your source servers.

**Note:** We are in no way endorsing any method or product in this chapter, but want to provide you with a number of options for your p-to-v migration.

## 6.3. P-to-V for Free

The following examines p-to-v migration processes that are either completely free of cost or cost very little. These methods may require manual modifications to files like Hal and Ntoskrnl.



## 6.4. cat Your Hard Drive


I'll begin with the cat p-to-v process because it was the first p-to-v migration process shown to me (Ron McKelvey developed it and instructed me). More importantly, it works 100 percent of the time with any physical server running any operating system. This is the most difficult p-to-v method by far, but, again, it's worked perfectly every time I've seen someone proficient with Linux use it. If you're not a Linux guru and want to skip ahead to a more "friendly" method, such as the BartPE or other methodologies, that's fine. Just the same, I recommend that everyoneLinux expert or noloop through this methodology and understand its components.

The concept of this p-to-v process, or any p-to-v process, is really quite simple: you are making a copy of your hard drive and piping it eventually into a virtual server. With this method, you make a copy of your hard drive at the block level, and then pipe the standard output of the cat command into another file located on a swing server. From the swing server, you pipe it back into a target virtual server.

The cat command can be used to display the contents of a file. Linux views everything, including hard drives as files, so running cat against your hard drive generates an output of your hard drive's contents. I don't recommend you do this, but I have, and it's a fast and furious output. However, if you cat your hard drive and redirect the output to somewhere else, say another file, then you have a replica or image of your hard drive. That's what this and (really) all p-to-v methods do create an image of your hard drive. So let's get to it.

Initially, this is a two-step process: cat the hard drive of your source server, pipe it to the swing server, and then from your target virtual server you pipe that copy from the swing server onto the virtual disk of your target vm. The process flow is illustrated in [Figure 6.1](#)

**Figure 6-1. Source Server Image Flow**



### 6.4.1. cat P-to-V Prerequisites

Some prerequisites to this method are as follows:

- Have an NFS share mounted and accessible on the network or can use Windows Services for Unix (if you really want)
- Have a Red Hat 9, disc 1

### 6.4.2. cat Process

The following is the basic method for catting your hard drive and creating a virtual machine from its standard output:

1. Boot the server with the Red Hat disc, and then type **linux rescue** at the boot prompt.
2. Choose **No** for network.
3. At the prompt that next appears, enable networking by typing the following commands:
- 4.
5. `ifconfig eth0 up`
6. `ifconfig eth0 x.x.x.x`

Try and ping the gateway. if you can't see it, try a different eth device, such as eth1





## 6.5. BartPE

BartPE is very cool. Whether you are going through a p-to-v migration or not, you should become familiar with this tool if you aren't already.

BartPE allows you to create a bootable Windows CD-ROM or DVD that creates a Win32 environment in which applications such as Ghost can be run. BartPE includes network support, a GUI, and file system support for FAT/NTFS/CDFS.

If you haven't had a chance to familiarize yourself with BartPE, which is free, do so at <http://nu2.nu>. You can use BartPE to migrate all Windows platforms ESX Server supports, as well as many of the Linux varieties. See [Table 6.1](#) near the end of this chapter.

The way you can use BartPE to migrate your physical machines into your virtual environment is really straightforward, although there are some legal issues to be aware of.

### 6.5.1. Legal Issues

BartPE requires its own Windows XP or Windows 2003 Server license. If you have an XP license for a workstation and create a BartPE CD-ROM and use it on another workstation, you have violated the End User License Agreement (EULA). Buy a copy of Windows XP specifically for BartPE use but don't expect Microsoft to support you if you have issues with BartPE because they won't. Microsoft does not support BartPE.

### 6.5.2. P-to-V with BartPE

Before jumping into the p-to-v process, the following are some premigration requirements that must be carried out (and which are covered in the subsequent sections):

- Create your BartPE disc
- Include Ghost 8.0
- Include network drivers for your physical server
- Create a share or swing virtual machine for the image

### 6.5.3. Creating Your BartPE Disc

To create your BartPE disc, follow the instructions provided by Bart for his PEBuilder. They are straightforward and very accurate.

I recommend creating both a CD-ROM and keeping its ISO image. Store the ISO image in your VMlibrary when it's complete.

#### 6.5.3.1. Including Ghost 8.0

In your BartPE CD-ROM, include the Ghost plug-in. This is a key plug-in for the p-to-v process. Of course, you will have to own Ghost, so include the cost of a copy in your budget.

#### 6.5.3.2. Including Network Driver for Your Physical Server

You will want to include network interface card drivers in your BartPE CD-ROM. Search the servers you'll be migrating for their NICs and ensure you include them in your BartPE build.

#### 6.5.3.3. Including Disk Drivers for Your Physical Server

You'll want to include SCSI drivers in your BartPE CD-ROM. Search the servers you'll be migrating for the SCSI cards they use, if any, and ensure you include them in your BartPE build.

For NICs and disk drivers, Bart offers a number of Help options to assist you in creating the correct plug-in. Just do a search on the help page for "plug-ins."

#### 6.5.3.4. Creating a Share or Swing Virtual Machine

Create a share large enough to hold an image of the server or servers you will be p-to-v'ing. You may also





## 6.6. VMware P2V Assistant 2

The P2V Assistant 2 tool, provided by VMware, works well enough for the price. Although earlier versions required service packs and hotfixes be reinstalled after the p-to-v process, the most recent version states that this is no longer necessary, with Windows NT 4.0 being the only possible caveat.

The tool works on Windows platforms only and is fairly straightforward. However, there is mandatory training that VMware requires their customers to go through prior to using this tool.

The tool works well with third-party imaging products such as Symantec's Ghost and can take a Ghost image and turn it into a virtual machine replacing any files, such as Hal and Ntoskrnl, with the correct versions needed for the virtual machine to run optimally.

The P2V Assistant is licensed on a per-server basis as well since they offer an enterprise license which allows you an unlimited number of migrations. The perserver license decrements whether or not the migration was successful, which isn't exactly great.

For more information on this migration tool, go to [www.vmware.com/products/vtools/p2v\\_features.html](http://www.vmware.com/products/vtools/p2v_features.html).

## 6.7. Platespin PowerP2V

Another option for p-to-v'ing servers is a third-party application by a company called Platespin. This tool is called PowerP2V and looks very promising. It provides all of the basic functionalities, along with some very advanced features that some might find useful.

PowerP2V runs from a management console. From there, you can perform drag-and-drop physical-to-virtual migrations. Simply click your physical source server and drag it to your ESX Server to initiate the process. There's a lot you can customize, including keeping the source server turned on or off after the migration is complete. You can also schedule the migrations, which can be very handy for larger servers. Like VMware's P2V Assistant, Platespin handles adding the correct Hal and Ntoskrnl files as part of the migration process. Unlike the VMware's tool, Platespin can migrate Linux servers as well. In addition to basic p-to-v migrations, Platespin offers virtual-to-virtual or v-to-v migrations. For instance, you can migrate a virtual machine running in Microsoft's Virtual Server product into a virtual machine on an ESX host, and vice versa.



## 6.8. P-to-V Conclusion

[Table 6.1](#) is a matrix that examines the four p-to-v methods discussed in this chapter. The pros and cons of each method are listed, which may help you decide which method is right for you.

<b>Table 6-1. The Four P-to-V Methods</b>				
	<b>cat Your Hard Drive</b>	<b>BartPE</b>	<b>VMware P2V Assistant</b>	<b>Platespin Power P2V</b>
Costs money	N	N	Y	Y
Physical contact with source or target server required	Y	Y	Y	N
Multiple simultaneous conversions possible	Y	Y	N	Y
<b>All ESX guest operating system support</b>				
Windows NT 4.0	Y	Y	Y	Y
Windows 2000	Y	Y	Y	Y
Windows 2003	Y	Y	Y	Y
Windows XP	Y	Y	Y	Y
Red Hat Enterprise Linux 3.0	Y	Y	N	N
Red Hat Enterprise Linux 2.1	Y	Y	N	N
Red Hat Linux Advanced Server 2.1	Y	Y	N	N
Red Hat Linux 9.0	Y	Y	N	N
Red Hat Linux 8.0	Y	Y	N	Y
Red Hat Linux 7.3	Y	Y	N	Y
Red Hat Linux 7.2	Y	Y	N	N
SUSE Linux Enterprise Server 9	Y	N	N	N
SUSE Linux Enterprise	Y	N	N	N



## 6.9. Summary

This chapter discusses four methods for migrating your physical servers into a virtual environment. Each p-to-v method is unique in certain ways and similar in others. Using native Linux commands such as *cat* to image your machine or utilizing BartPE and Ghost to image your machine produces the same output, and the same image of your physical server that you will then virtualize. It's best to try each of these methods and become familiar with the plusses and pitfalls of each. If you're a Linux guru, *cat* is a great methodology, which can be used on any server that ESX supports. (And it's free!) If you're strictly a Windows person, then BartPE is a fantastic solution, even if you're asked to migrate Linux servers.

If you prefer utilizing a tool and the support that comes with it, both the P2V Assistant and PowerP2V will get the job done. Of course, these will cost you but if you plan ahead and include the costs within your budget you'll definitely be able to accomplish your p-to-v process quickly. The P2V Assistant allows for basic migrations while the PowerP2V allows for scheduling multiple p-to-vs at the same time. Plus, Platespin has announced that it will soon be able to support v-to-p migrations in case you need to put a server back on a physical platform.

## 6.10. Solutions Fast Track

### 6.10.1.

#### 6.10.1.1. P-to-V Terms

- ☑ The source server is the physical server that you'll be migrating into your virtual infrastructure.
- ☑ The target server is the target virtual server that you'll be migrating your source server into.
- ☑ The swing server is where you'll place the images of your source servers.

#### 6.10.1.2. cat Your Hard Drive

- ☑ With the cat p-to-v process, you make a copy of your hard drive at the block level, and then pipe the standard output of the cat command into another file located on a swing server.
- ☑ Using the native Linux command *cat*, you can image your physical server. That image can then be pulled into a virtual machine using a process very similar to the one you used to extract it from the physical server.
- ☑ The cat p-to-v process works 100 percent of the time with any physical server running any operating system.

#### 6.10.1.3. BartPE

- ☑ If you're more comfortable with Windows, then use BartPE and Ghost 8.0 to create the image of your physical server, which you can then Ghost into a virtual.
- ☑ BartPE allows you to create a bootable Windows CD-ROM or DVD that generates a Win32 environment where applications such as Ghost can be run.
- ☑ BartPE requires its own Windows XP or Windows 2003 Server license.

#### 6.10.1.4. VMware P2V Assistant 2

- ☑ VMware's P2V Assistant is a tool that creates virtual machines from physical machines and takes care of changing out the Hal and Ntoskrnl files to correct versions.
- ☑ The tool works on Windows platforms only and is fairly straightforward.
- ☑ The tool can take a Ghost image and turn it into a virtual machine.

#### 6.10.1.5. Platespin PowerP2V

- ☑ The tool with the most bells and whistles is Platespin PowerP2V, which allows for multiple p-to-vs running at the same time, schedules p-to-v jobs, allows for disk resizing on the fly, and, of course, changes out Hal and Ntoskrnl files for you.
- ☑ Like VMware's P2V Assistant, Platespin handles adding the correct Hal and Ntoskrnl files as part of the migration process.
- ☑ Unlike VMware's tool, Platespin can migrate Linux servers as well.



## 6.11. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** Why should you pay for a p-to-v toll when there are a number of ways to do it for free?

**A:** If you don't have the time or inclination to learn the manual process or need the assurance of a vendor for your physical-to-virtual migrations, then a tool that comes with support is a very reasonable option.

**Q:** What takes the longest in the p-to-v process?

**A:** Transferring the data of the physical server into either an image file or the virtual server itself. That's why tweaking your network settings for optimal throughput is essential when p-to-v'ing a server.

**Q:** Can you create complete backups of your physical servers and then recover them into virtuals?

**A:** I've used NTBackup to back up Windows 2000 Professional and Windows XP, and created virtuals out of .bkf files. It works but there's a bit of futzing around with the virtual once it is running. Try it.

**Q:** What's the most essential part of creating a smooth migration plan?

**A:** Practice. No matter what the tool or process, practice it and learn its gotchas and the workarounds. No matter what method you use, there will be times when it won't go as planned. Practicing your p-to-v process on a number of physical platforms and disk configurations allows you to become more adept once you're doing it for real.

# Chapter 7. Server Consolidation

Solutions in this chapter:

- Server Consolidation and Virtualization Defined
- Business Case and Technical Requirements
- Developing Your Project Methodology
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions

## 7.1. Introduction

Although virtual machines are not a new concept for the mainframe world, the idea is fairly new to the production x86 server realm. Virtualization in an x86 environment is a complete paradigm shift for IT, which has for many years been accustomed to deploying one application per server. Hardware and software vendors have also historically recommended one application instance per server. While this model was financially attractive for operating systems and server hardware vendors, it's increasingly a suboptimal way to run a company's IT infrastructure. Companies are realizing that the proliferation of servers is increasingly difficult to manage and can actually hinder IT from meeting the needs of business. CIOs' agendas for cost performance, mergers and acquisitions, and compliance risk reduction are driving organizations toward server consolidation and virtualization projects. Furthermore, consolidation and virtualization projects involve complexities beyond the technology. Consequently, this chapter deals with the people and process components that make up a server consolidation and virtualization project, and leverages many of the best practices and lessons learned from our customer engagements in the area of server consolidation.

One of the primary messages that Callisma delivers to our clients in our *Server Virtualization Lessons Learned* presentations is the fact that these projects require a significant amount of planning. Specifically, the planning phase should focus not only on architecture or technology components but also on the people and process components. All three components are necessary to consolidate servers and build operational efficiencies into the environment.

### Designing & Planning... People, Process, and Technology

Addressing the people and process components of your server consolidation project will dramatically increase the acceptance of the technology. In our experience, companies that develop a comprehensive and detailed plan, and that engage the business early in the project, can consolidate a considerably higher percentage of application servers than companies that take an *ad hoc* approach.

Many organizations have difficulty making proactive decisions when it comes to server provisioning. This can be due to rapid growth, resource constraints, mergers and acquisitions, or a lack of appropriate tools to proactively manage the environment. As a result, these organizations often enlist the assistance of professional services firms with specialized knowledge and experience in server consolidation and virtualization. Through Callisma's services for virtualization and consolidation, we have helped a number of organizations achieve a proactive posture for server provisioning. We do this by helping our clients develop a comprehensive plan for server virtualization that encompasses people, processes, and technology. The previous chapters addressed VMware and virtualization technology. Now, in this chapter, we share many of our "lessons learned" in server virtualization, especially as they relate to a project approach, capacity planning, and server provisioning. We focus on how to achieve a successful server consolidation project, concentrating mainly on the people and process aspects of the project.



## 7.2. Server Consolidation and Virtualization Defined

"Our data centers are full!" When we hear our customers say this and we hear it frequently we know that virtualization may provide the optimal environment for them to consolidate their servers. But what do we mean by consolidation? In general terms, consolidation means joining multiple entities together into a single entity. For example, merging multiple remote data centers into a single location, streamlining the amount of space used within a single data center location, eliminating multiple overlapping or redundant applications into a single unified application, or reducing multiple physical servers into a single physical server that supports multiple virtual servers. The last of these examples serves as our primary focus: reducing the number of hardware server platforms in your environment.

Consolidation has many benefits, including better utilization of company assets, better control of resources, and the overall reduction of hardware capital costs. Specifically, server consolidation has positive ripple effects on other types of consolidation. For example, reducing the number of physical servers results in reduced requirements for floor and rack space, power consumption, air conditioning, and data wiring. This allows a given data center to reclaim physical space. Or consider an entire data center consolidation. If the servers are virtualized before decommissioning the remote data center, the server hardware can be re-tasked, sold, or otherwise disposed of locally, thereby saving the cost of shipping. Additionally, server consolidation can provide a cost-effective alternative to data center expansion or relocation.

Consolidating the number of physical servers also reduces the hardware required to operate a service or group of applications. This, in turn, reduces the operational costs of maintaining this hardware, which can contribute to improving the overall quality of the service. The organization as a whole benefits from standardization of operating systems, better utilization of systems resources, a reduction in physical space needed to contain the servers, easier manageability, and an increase in security. The net result is a more proactive environment, and more time available to focus on providing value back to the business.

Now that we've defined consolidation, let's narrow our focus to virtualization. As with consolidation, there are several forms of virtualization such as storage virtualization, server virtualization, and application virtualization. We have all witnessed the rapid adoption of SANs (storage area networks), the practical implementation of storage virtualization. Storage virtualization basically took the direct attached (unshared) storage devices and created a generally available, shared storage device. A greater amount of storage becomes available in a shared environment, making it easier and more cost-effective to manage.

Server virtualization does for server hardware what storage virtualization did for disks: it makes the server hardware a generally available, shared hardware platform. By server virtualization, we mean the logical partitioning of a single hardware platform into multiple virtual hardware platforms through a software layer. Simply put, virtualization software allows you to run many operating systems on a single hardware platform. This is what VMware's ESX Server does: it abstracts the server hardware from the operating systems that run on it. VMware pioneered x86 server virtualization software, and its stability has brought the use of virtualization to the production-distributed systems environment. Building upon server virtualization, we are now seeing application virtualization: that is, encapsulating individual applications to run in their own isolated environment, one that isn't tied to a physical server platform.

For the purposes of this chapter, when we refer to virtualization we are focusing on server virtualization, which enables physical hardware consolidation. It's important to note that even with server virtualization you still need to purchase the individual server operating systems and third-party software agents that you plan to run in each virtual machine. So it's still important to achieve results in the area of server consolidation by stacking multiple applications on one virtual machine and running multiple virtual machines on one physical server.

### 7.2.1. Adoption and Support of Server Virtualization

Sounds great, but is it industry-proven and do vendors support it? This is one of the first questions an organization wisely asks about any new technology. As we have already mentioned, mainframe-based virtualization has been well established for decades. UNIX- and AIX-based virtualization is somewhat newer but also accepted and deployed.

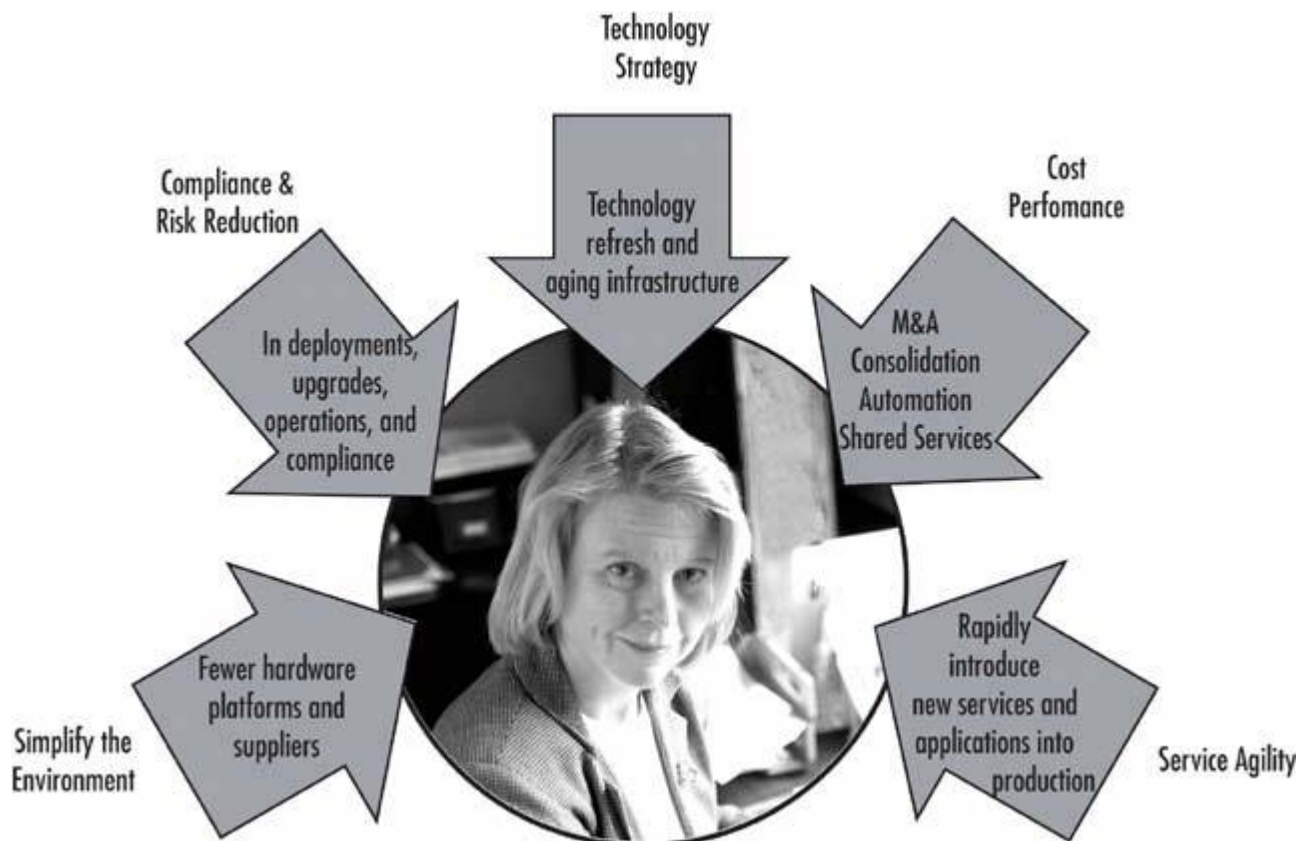




## 7.3. Business Case and Technical Requirements

The business and technical requirements need to be defined and documented for the project. This data is not only necessary for subsequent phases to define the objectives and motivations for the project but to assist in identifying the technology components that will be incorporated into your architecture or design phase. In addition, you may need to develop the business case to attain approval and funding for the project. In other cases, the server consolidation initiative may be a subproject in an overall technology refresh project. In all of these cases, by understanding the business, technical, and operational objectives, the project team has a much better opportunity to influence cost savings, create automation, and more specifically meet your organizational objectives. Callisma utilizes [Figure 7.1](#) to illustrate and map back the benefits of our project to executives and customer project sponsors.

**Figure 7-1. The CIO's Agenda**



Utilizing the "CIO Agenda" graphic in [Figure 7.1](#), let's walk through and explore potential areas that can address some of the things on the minds of CIOs today.

### 7.3.1. Technology Strategy: Technology Refresh and Aging Infrastructure

One of the best opportunities for leveraging server consolidation or virtualization is when a large percentage of the server hardware infrastructure is aging and a potential technology refresh project is necessary. The equipment may be coming off a lease, and your company may want to use this as an opportunity to reduce the number of hardware platforms in the environment. Other times, our server consolidation and virtualization projects are tied closely to other initiatives such as a technology refresh. A very common example of this is when we're assisting our clients with Active Directory or Exchange 2003 design and deployment projects. Quite frequently a subproject includes server consolidation and virtualization. In other cases, we see virtualization projects as good follow-on projects to clean up the remaining legacy servers that housed applications that possibly did not support the latest operating system or where the business was unwilling to perform remediation on the application.

### 7.3.2. Cost Performance

The key issues related to cost performance are merger and acquisition (M&A) activities, consolidation, and hardware refresh. With server consolidation, the goal is to reduce the number of hardware platforms in the environment, which can lead to significant cost savings.







## 7.4. Developing Your Project Methodology

Once executive management has accepted your proposal and approved the funding, you're ready to begin the project. In this section, we provide a high-level outline of the project's phases, and along the way we discuss some best practices and lessons learned. This section is far from a complete methodology, however we discuss the essential steps at a high-level to help your project succeed. Many of the lessons learned come from actual Callisma customer engagements. Our focus is primarily on the project goals to reduce the number of hardware platforms within the production environment.

Server consolidation projects should generally contain the following high-level phases:

- Establish project
- Gather data and application/server inventory
- Assessment
- Capacity planning
- Technology/design
- Testing and validation
- Piloting and full-scale implementation

### 7.4.1. Establishing the Project

The first phase of server consolidation is to establish the project. Server consolidation projects require a highly collaborative effort between many departments within the systems community and lines of business, including the server team, applications team, capacity planning, security, operations, the line of business (LOB) or application owners, and so on. Some organizations take a noncollaborative project approach: they develop the architecture and implement the new technology components, hoping that the lines of business adopt the technology. As a result, the adoption of the new technology tends to be lower. Instead, Callisma recommends that customers develop a comprehensive plan and engage the lines of business early in the process. Callisma is happy to assist customers in either approach; however, based upon our experience we estimate that the collaborative approach yields a higher ratio of consolidation.

The core project team should not only develop a collaborative forum to work with the various departments within the systems community, but it should proactively work with the business to ensure processes are communicated and schedules are maintained. Throughout this chapter, we'll offer best practices and lessons learned that will help you envision how this takes place, what project efficiencies are gained, and how this builds cooperation within an organization.

Server consolidation requires involvement from many departments within an organization. So one of the first steps in the project is to identify your stakeholders and the various teams that will take part in your project. You'll need to define roles and responsibilities, and you'll need to assess resource requirements. We've listed some of the key roles here. Of course, staff titles and roles within your specific organization may vary, and in some cases, it can be one person or a department of people representing the roles. The key roles include the following:

- **Core project team** The core project team consists of the program and project managers and is typically made up of some of the members of the other teams in this list. They have the overall responsibility to establish and successfully complete the project on time and within budget. This team will typically work with the applications team and systems community to determine information needed to develop requirements, hardware configurations, and technical business requirements.
- **Server team** The responsibilities of the server team include the architecture of the VMware solution and in general the overall server architecture for your organization. In many cases, this will also be the team primarily responsible for the VMware project. It's important for this team to avoid a "silo" approach that focuses just on the technology. Server virtualization can have a tremendous operational impact on an organization; it changes the IT chargeback model, server documentation (run books), and the way servers and applications are deployed in the environment. It affects business processes.



## 7.5. Summary

Mastering the technology does not always mean that it makes it easy to deploy in a live production environment and does not guarantee you have support and buy-in from the business. In this chapter, we reviewed best practices and lessons learned from Callisma's customer engagements. There are many "people" and "process" components that have to be assessed and adjusted in order to enable a successful project. A highly collaborative team with a well-established and documented plan is essential to achieving a high ratio of servers consolidated. As documented in this chapter, the following list provides a summary of project enablers that can be leveraged to help make your project a success:

- Develop a business case
- Perform financial analysis
- Establish teams and define roles
- Assess resource requirements and address resource constraints
- Enable capacity planning for all phases of the project
- Utilize forums to establish a coalition, including a weekly status meeting, an application owner forum, and technology briefings
- Conduct structured interviews and develop questionnaires
- Create a centralized repository
- Create a project Web site
- Assess technical and operational gaps
- Manage variance in project support
- Establish a server consolidation/tracking database

## 7.6. Solutions Fast Track

### 7.6.1.

#### 7.6.1.1. Server Consolidation and Virtualization Defined

- ☑ In general terms, consolidation means joining multiple entities together into a single entity
- ☑ Consolidation has many benefits, including better utilization of company assets, better control of resources, and the overall reduction of hardware capital costs.
- ☑ Server virtualization does for server hardware what storage virtualization did for disks: it makes the server hardware a generally available, shared hardware platform.

#### 7.6.1.2. Business Case and Technical Requirements

- ☑ Reducing hardware platforms can lower an organization's total cost of ownership.
- ☑ Coupling technologies such as VirtualCenter and VMotion provide benefits in the area of maintenance and recovery of your servers and help ensure that improved service levels are met.
- ☑ Speed-to-market and agility may become the key business drivers for your virtualization project.

#### 7.6.1.3. Developing Your Project Methodology

- ☑ Develop a complete, accurate, and centralized dashboard for planning data
- ☑ Create a staffing plan and assess resource constraints
- ☑ Ensure rationalization is incorporated into the project
- ☑ Enable the flexibility to create subteams for the remediation of process or technical components

## 7.7. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** Will my software vendor support my application in a VMware environment?

**A:** You'll need to contact your specific software or applications vendor to find out; however, an increasing number of software vendors support virtualized environments. Once you've migrated your application from a physical to virtual environment it's very easy to migrate back to a physical environment for vendor support. An excellent document provided by VMware can be found at the following link: [www.vmware.com/pdf/ms\\_support\\_statement.pdf](http://www.vmware.com/pdf/ms_support_statement.pdf).

**Q:** What are the best enablers to help ensure our server consolidation is a success?

**A:** Engaging with the lines of business and application owners not only assists in attaining an actionable schedule, but also develops a highly collaborative environment that facilitates buy-in and support for your project. In addition it will be difficult to incorporate rationalization into your project unless you are working directly with the business or application owners.

**Q:** What tools do you recommend for a server consolidation or virtualization project?

**A:** It's imperative to have a robust and comprehensive tool for capacity planning, including historical data. This tool can be utilized in every phase of your project. Other tools to support the project include modeling tools used in test consolidation scenarios. Finally, you'll need to evaluate and document tools and processes for each technology you wish to consolidate. The migration tools will need to support all the platforms you wish to use in your migration scenarios.

# Chapter 8. Cool Tools for a Virtual Infrastructure

Solutions in this chapter:

- ISO Images and WinISO
- WinSCP
- PuTTY
- NewSID
- Platespin's PowerP2V and PowerRecon
- VirtualCenterVmotion
- *vmkfstools* and *vmware-cmd*
- VMware Tools
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions



## 8.1. Introduction

In this chapter, we discuss some of the very cool tools you need to manage your virtual infrastructure. We'll also discuss the requirements and considerations for each. Again, your tool choice will depend on what works for you, but knowing these tools is a great start in understanding how a virtual infrastructure can be managed.

Like the construction and maintenance of anything, having the right tool for the job makes all the difference.



## 8.2. ISO Images and WinISO

You'll use ISO images a lot in your virtual infrastructure. Ultimately, you'll have a small or large library of them most likely. WinISO is one of many products that allows you to create ISO images from CDs and DVDs as well as from any single file or collection of files. We chose WinISO solely because we use it and feel it's one of the better products out there for this particular task if Windows is your desktop OS. Just about any similar functioning tool will work. This type of tool is essential for working within a virtual infrastructure since you'll find the use of ISO images within an ESX environment to be critical.

What is an ISO image? (For a definition pulled from Wikipedia, see the Configuring & Implementing sidebar titled "ISO 9660 Defined.") In the context of using ISO images within a virtual infrastructure, an ISO image is a file that's basically a complete copy of the original media from which the ISO image is taken. For example, let's say you want to build a new Windows 2003 Server, Enterprise Edition virtual machine. You can go into your data center, insert your CD-ROM, go back to your desk and then begin the build process. Or, using a tool like WinISO, you can create a bootable ISO image from your Windows 2003 media at your desktop, place the ISO image on your ESX Server with tools like WinSCP (which we'll discuss later in this chapter), and then access the ISO image by simply changing your virtual machine's CD-ROM configuration. This way you never physically have to touch your ESX Server, wherever it may be, and yet you still have access to CD-ROM images from your virtual machines.

### Configuring & Implementing... ISO 9660 Defined

The following is the definition of ISO 9660 pulled from Wikipedia at [http://www.wikisearch.net/en/wikipedia/i/is/iso\\_9660.html](http://www.wikisearch.net/en/wikipedia/i/is/iso_9660.html):

ISO 9660, a standard published by the International Organization for Standardization, defines a file system for CD-ROM media. It aims at supporting different computer operating systems, such as Microsoft Windows, Mac OS, and systems that follow the UNIX specification, so that data may be exchanged.

DVDs may also use the ISO 9660 file system (although the UDF file system is far more commonly used).

#### Levels and Restrictions

There are different levels to this standard:

- **Level 1** File names are restricted to eight characters with a three-character extension, uppercase letters, numbers, and underscores; the maximum depth of the directories is eight.
- **Level 2** File names may be up to 31 characters.
- **Level 3** Files can be fragmented (used on CD-Rs written by the software "Direct CD" for Windows, for example).

All levels restrict names to uppercase letters, digits, and underscores (\_). Some CD-mastering applications allow the user to use almost any ASCII character. While this does not strictly conform to the ISO 9660 standard, most operating systems that can read ISO 9660 file systems support the use of most ASCII characters as an extension.

The restrictions on filename length and directory depth have been seen by many as a more serious limitation of the file system. While in many cases, it's possible to rename files or truncate the filenames in question, and many CD burning applications do this automatically, in certain special cases, this isn't feasible. For example, class files in Java are required to follow strict naming rules, which, especially in the case of nested classes, may require filenames longer than 31 characters.

#### Disk Images





## 8.3. WinSCP

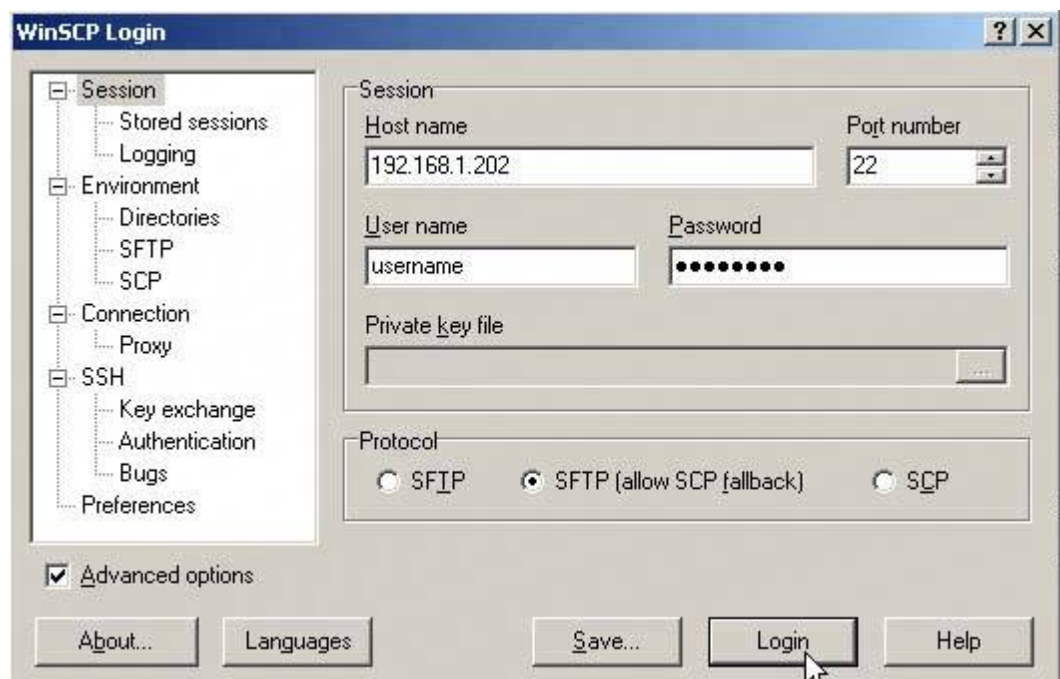
WinSCP is a free open source tool that allows Windows machines SSH access to Linux/Unix systems. The interface for this tool is very intuitive and familiar looking if you use Norton Commander or Windows Explorer. We use this tool for transferring ISO images and other files such as .flp (floppy images) onto our ESX Server.

To get WinSCP, go to [winscp.net/eng/download.php](http://winscp.net/eng/download.php).

To use WinSCP, load it on your client system and start it.

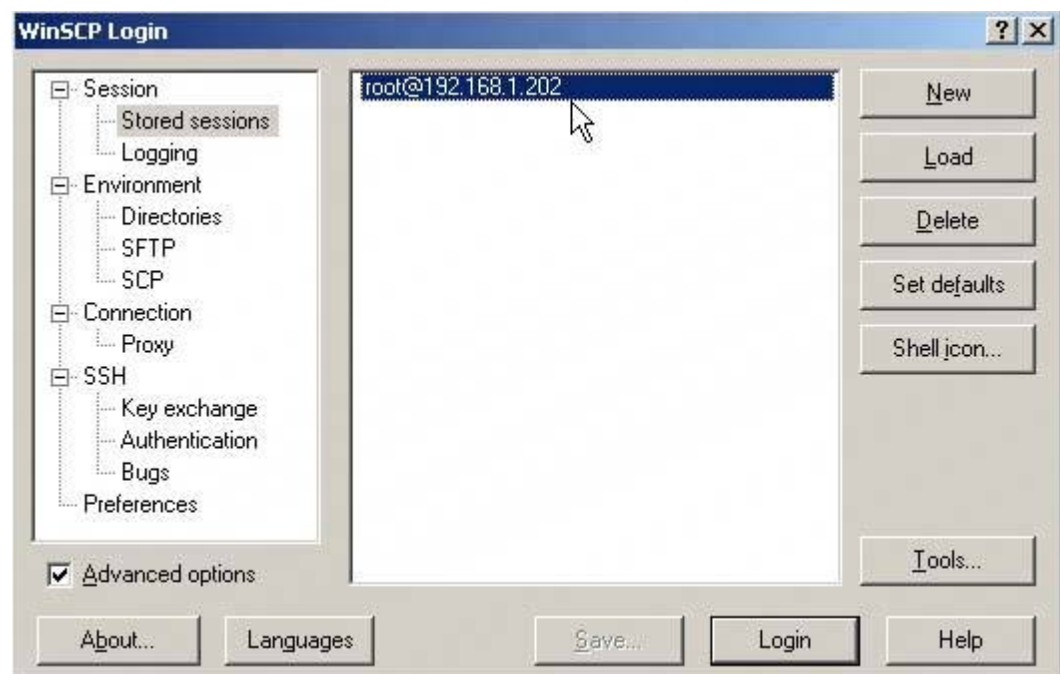
Input the required data to log in to your ESX Server, and then click Login (see [Figure 8.6](#)).

**Figure 8-6. Starting a WinSCP Session**



Save the login information as a Stored Session so you don't have to type it in every time. The next time you log in to your ESX Server, you'll be prompted to select a stored session or create a new connection (see [Figure 8.7](#)).

**Figure 8-7. Selecting a Stored Session**









## 8.4. PuTTY

PuTTY is a free tool produced, maintained, and updated in Cambridge, England, by Simon Tatham and his team. The tool allows you to connect via Secure Shell (SSH), Telnet, and Rlogin to remote Linux/Unix systems. PuTTY is a program that you run from, say, your Windows desktop and, utilizing SSH, Telnet, or Rlogin, allows you to access the file systems of remote Unix/Linux computers (for example, your ESX Server). The ESX MUI allows you to perform the majority of the functions, say 95 percent, that ESX Server provides; accessing the console of your ESX Server via PuTTY provides 100 percent of all ESX Server functionality. PuTTY allows for this access if you're using a Windows desktop.

By default, the most current version of ESX Server, as of this writing, sets the ESX Server security to High. This allows for console access utilizing SSH only. Of course, if your workstation is Linux/Unix-based you can easily shell into your ESX Server. However, if you run Windows on your desktop, as we do, most of the time, PuTTY is used as an SSH tool, so you'll want to become familiar with it.

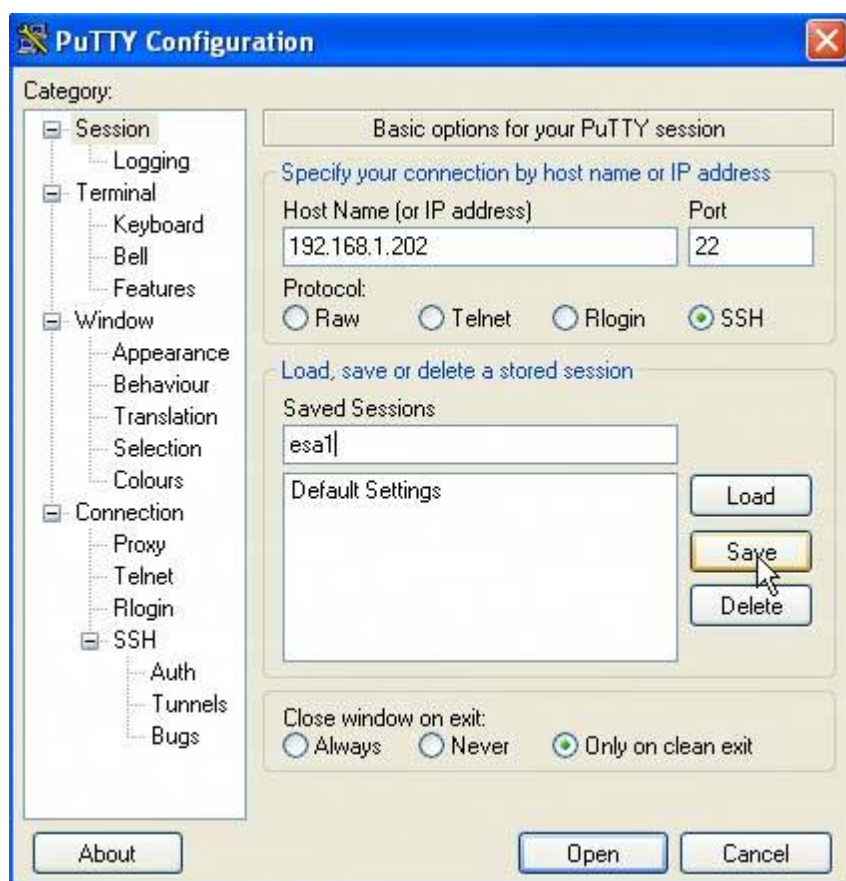
So, what is SSH? Well, SSH allows you to connect to remote computers, such as your ESX Server, using an encrypted link. This helps prevent network packet sniffing, and the discovery of unencrypted passwords and text can be read by unauthorized individuals. Once you've connected to your ESX Server using PuTTY, which uses SSH for the encrypted connection, you can run commands, scripts, and perform all of the functionality that comes native to ESX Server.

You can download PuTTY at the following site:

[www.chiark.greenend.org.uk/~sgtatham/putty/download.html](http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html).

Once you've downloaded and installed this, PuTTY is quite simple to use. Double-click the **PuTTY icon** to start the executable and bring up the window shown in [Figure 8.11](#)

**Figure 8-11. The PuTTY Window**



You can type in the IP address of your ESX Server or if you have created an A record in your DNS server, you can input the host name of your ESX Server. In the Saved Sessions area, type in a name for this session, which in this example is simply the name of our ESX Server. Once you've done that you can click the **Save** button to add the settings, host name, or IP

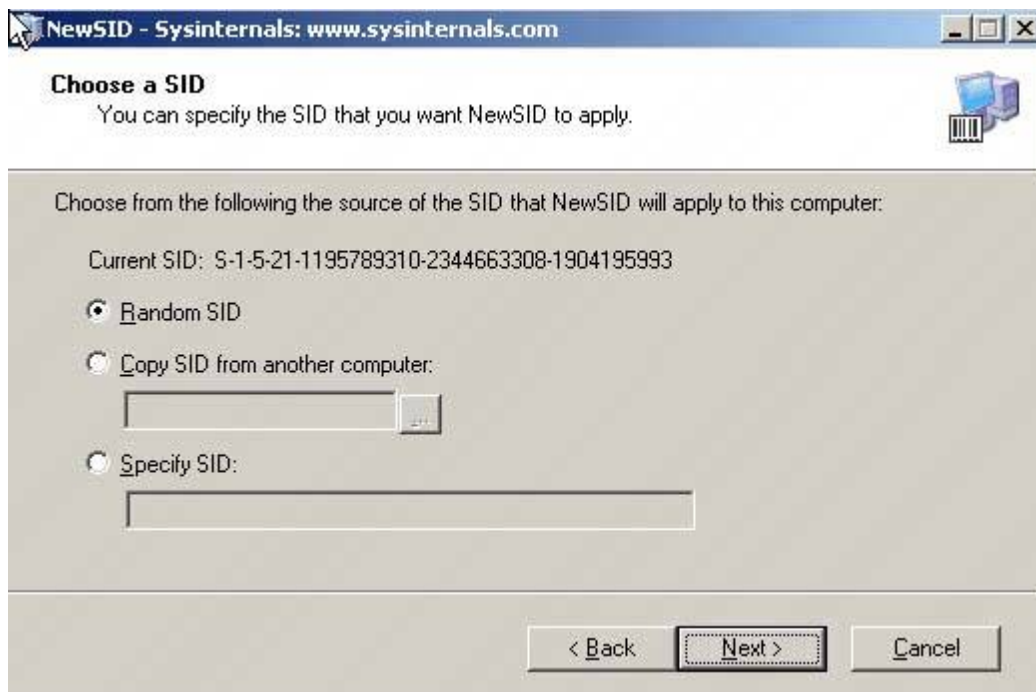




## 8.5. NewSID

NewSID is a free tool you can download from [www.sysinternals.com](http://www.sysinternals.com). You need to run this tool if you clone template virtual machines. When you run it you can select a random SID or specify an SID yourself (see [Figure 8.13](#)). We recommend the former.

**Figure 8-13. Choosing an SID**



Click **Next**. You should now be able to rename your virtual machine (see [Figure 8.14](#)).

**Figure 8-14. Renaming the Computer**



Click **Next** and you'll be given the option to reboot the computer after the new SID is applied. The SID will then be applied to the virtual machine (see [Figure 8.15](#)).

**Figure 8-15. Applying the SID**





## 8.6. Platespin's PowerP2V and PowerRecon

Platespin produces two tools you can buy to manage your virtual infrastructure. We discussed the physical-to-virtual migration capabilities of PowerP2V in [Chapter 4](#). This tool allows for high-level server profiling combined with the ability to make virtual machines out of physical ones. The tool supports both ESX Server as well as Microsoft's Virtual Server, and provides the ability to move virtual machines between both virtual platforms.

Platespin also produces PowerRecon. This tool allows for ESX Server sizing by gathering the detailed utilization and trending information of your physical server. This data can then be used to provision your ESX Server accordingly.

Of course, other products exist for this type of measuring, such as Nagios.

Platespin and PowerRecon are referenced only so you can understand the data gathering tool and the P2V tool requirements you may need for your virtual infrastructure.

## 8.7. VirtualCenterVmotion

VirtualCenter is a single console from which you can manage all your virtual machines on ESX Server. This tool is discussed in greater detail in [Chapter 9](#), "Administering a Virtual Infrastructure."

This section's topic, though, is Vmotion, which is probably the coolest tool in the virtual infrastructure. What Vmotion does is allow you to move a running virtual machine from one physical ESX Server to another. The authors have been working with Vmotion since it was first available and when we move a running server from one physical box to another without any service interruption, we're still amazed. This tool theoretically allows for the 100-percent uptime of your virtual machines and the services they provide. If you plan to have an ESX Server environment with more than one ESX Server, this tool is an absolute must. Of course, the coolest tools are never free and neither is this one. Nevertheless, the price you pay is well worth it. Make sure VirtualCenter and Vmotion are included in your budget.





## 8.8. vmkfstools and vmware-cmd

Two powerful tools VMware provides natively for creating and managing both your ESX Server and your virtual machines are the commands:

- *vmkfstools*
- *vmware-cmd*

The commands and their associated options can be run from the command line, or in scripts that you can schedule manually or run as cron (backup) jobs.

### 8.8.1. vmkfstools

Use PuTTY to log in to your ESX Server and type **vmkfstools help**.

The various options for this command will appear, and there are many (see [Figure 8.16](#)).

**Figure 8-16. vmkfstools Options**

```
root@esx1:~
Last login: Mon May 30 10:00:13 2005 from 192.168.1.100
[root@esx1 root]# man vmkfstools
[root@esx1 root]# vmkfstools --help
vmkfstools: unrecognized option '--help'
vmkfstools -b --blocksize #[mMkK]
            -m --commit
            -F --config [public|shared|writable]
            -C --createfs [vmfs1|vmfs2]
            -Z --extendfs extension-SCSIDevice
            -P --querypartitions
            -c --createfile #[gGmMkK]
            -r --maprawdisk raw-SCSIDevice
            -e --exportfile dstFile
            -X --extendfile #[gGmMkK]
            -i --importfile srcFile
            -g --geometry srcFile
            -l --list
            -h --human-readable
            -M --verbose mappings
            -L --lock [reserve|release|reset|lunreset]
            -n --numfiles #
            -R --recover
            -s --scan
            -S --setfsname fsName
            -z --nozero
            -k --createswapfile #[gGmMkK]
            -w --activateswapfile
            -y --deactivateswapfile [swapFileID]
            -T --tovmfs2
            scsiDevice[:file]
[root@esx1 root]#
```

Many of this tool's options can be used in the MUI as well. The first option to discuss here is one the author used many times in fact, more than any other the `extendfile` option. What this option allows you to do is increase the size of a virtual disk. If you recall, in [Chapter 3](#) we discussed creating gold images or templates of virtual machines. These templates' virtual disks are exactly the same size. When you clone them, however, you must be able to adjust the size of the virtual disk. This is where the `vmkfstoolsX` command comes in handy.

Let's say the template you're cloning has a virtual disk of 5 gigabytes, but the virtual server you need to build requires 10 gigabytes. Of course, you could build a number of clones of various sizes and use the one you need accordingly, or you could use *vmkfstools* to expand the virtual disk to the size required. The following is an example. First, we'll clone the virtual disk of your template with a simple `cp` command and rename it appropriately (or you could go into the MUI and select **Manage Files** and copy the file there). We named our new `.vmdk`





## 8.9. VMware Tools

We first discussed loading VMware Tools in [Chapter 3](#) when we began building a virtual machine. VMware Tools loads drivers that allow your virtual machines to run more optimally and allow you to manage them more efficiently. VMware Tools also provides a number of additional functionalities, which we'll describe next.

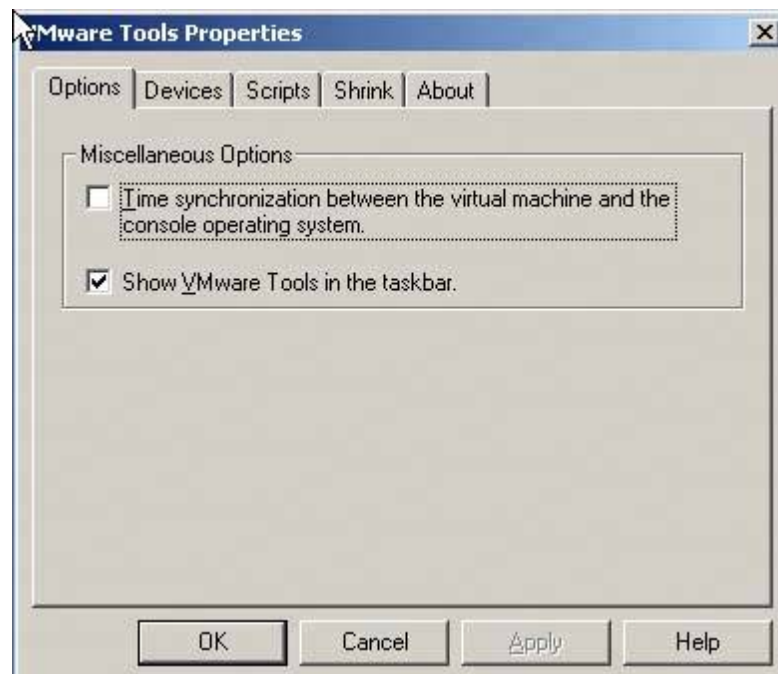
After you've loaded VMware Tools, you'll notice in your Windows virtual machines a small icon in the System Tray (see [Figure 8.23](#)).

**Figure 8-23. The System Tray Icon**



If you double-click the icon, a window appears with a number of tabs (see [Figure 8.24](#)).

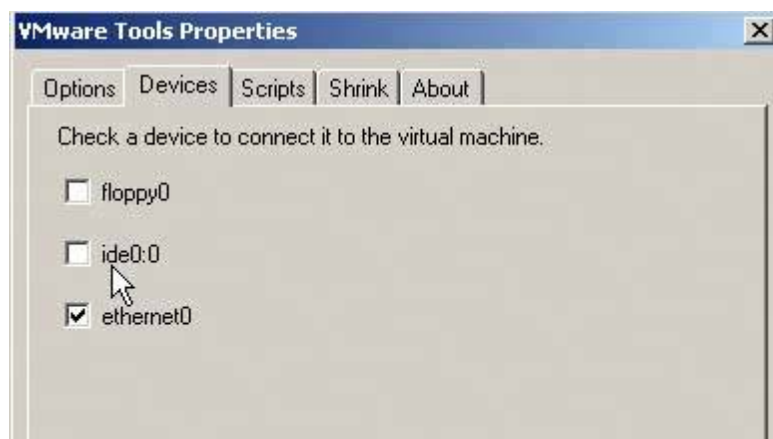
**Figure 8-24. The VMware Tools Properties Dialog**



From the **Options** tab, you can choose to synchronize your virtual machine with your ESX Server. If your ESX Server is not relying on its BIOS Clock, but has been configured to receive its time from an NTP server, you select this. Or, you can leave this option unchecked and set the virtual machine to receive its time from the server.

From the **Devices** tab, you can select from all the removable devices such as the CD-ROM and floppy drive that the virtual machine has access to (see [Figure 8.25](#)).

**Figure 8-25. Selecting a Removable Device**





## 8.10. Summary

There are many cool tools that you'll find indispensable in your virtual infrastructure. Some of the basic have been covered in this chapter. From free tools to costly ones, using the right tool for the job will make your virtual evolution progress go more smoothly. Many independent software vendors (ISVs) are currently working on incorporating support for virtual environments and you can expect the list of tools for managing and working within a virtual environment to grow significantly.



## 8.11. Solutions Fast Track

### 8.11.1.

#### 8.11.1.1. ISO Images and WinISO

- ☒ ISO images give you access to your CD-ROM library without having to physically slap a CD into the drive of your ESX Server.
- ☒ WinISO is one of a number of tools that allows you to create ISO images that can be used within your virtual infrastructure.
- ☒ Creating a library of ISO images greatly reduces the need for you to ever leave your chair...which is the goal of any good admin.

#### 8.11.1.2. WinSCP

- ☒ WinSCP allows you to move files from your Windows desktop to your ESX Server.
- ☒ Use this tool to transport your ISO images onto your ESX Server securely.
- ☒ You can use this tool to move files off your ESX Server as well.

#### 8.11.1.3. PuTTY

- ☒ PuTTY allows you secure access to your ESX Server console utilizing SSH from your Windows desktop.
- ☒ This tool lets you run a number of native commands in ESX Server that you can't perform from the MUI.

#### 8.11.1.4. NewSID

- ☒ NewSID allows you to reSID your cloned templates.
- ☒ It's easier to use then SysPrep.

#### 8.11.1.5. Platespin's PowerP2V and PowerRecon

- ☒ Platespin PowerP2V lets you easily migrate your physical server into your virtual infrastructure.
- ☒ PowerRecon allows you to gather the utilization data required to size your ESX Server and provision your virtual machines accordingly.

#### 8.11.1.6. VirtualCenterVmotion

- ☒ The king of all vm tools, Vmotion is a central management console that allows you to manage your ESX Servers and the virtual machines they run.
- ☒ Vmotion provides you with the ability to move running virtual machines to other ESX Servers without powering them down or experiencing any noticeable drop in service.

#### 8.11.1.6.1. vmkfstools and vmware-cmd

- ☒ *vmkfstools* allows you to manipulate many aspects of your ESX Server and your virtual machines.
- ☒ *vmware-cmd* lets you manage aspects of your virtual machines, some of which can't be performed from the MUI.

#### 8.11.1.7. VMware Tools

- ☒ VMware Tools loads the specific drivers needed for more advanced virtual machine functionality.
- ☒ It's highly recommended that you load VMware tools on all of your virtual machines for optimal performance and stability and to give you added functionality per virtual machine.





## 8.12. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** What is the VMlibrary?

**A:** The VMlibrary (/vmlib) is simply a directory on your ESX Server that lets you organize the tools and files you need to manage and maintain your virtual infrastructure,

**Q:** Can ISOs be shared between ESX Servers?

**A:** If you place your VMlibrary on a LUN that's shared amongst your ESX Servers, you can share your ISO images or anything else you place in your VMlibrary.

**Q:** Why is Vmotion so cool?

**A:** Try it...and then imagine the possibilities and applications. You'll be a convert soon enough.

**Q:** You mentioned that additional ISVs were creating management capabilities for virtual infrastructures in their products. Which ones?

**A:** Look at HP Insight Manager, Dell OpenManage, IBM Director, BMC's Patrol product line, Computer Associates' Unicenter, and many others.

# Chapter 9. Administering a Virtual Infrastructure

Solutions in this chapter:

- VMware Tools for Managing the Virtual Infrastructure
- Understanding VirtualCenter Terms and Concepts
- Installation and Configuration of VirtualCenter 1.2
- Using VirtualCenter to Administer Your Virtual Infrastructure
- Using the Management Interface and Service Console to Administer the Virtual Infrastructure
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions

## 9.1. Introduction

The focus of this chapter will be the administration of a virtual infrastructure. Changing from the legacy mind-set of a single OS per physical server to a virtualized infrastructure is a paradigm shift that will require a change in thought process. You no longer need to be as concerned with rack space because your focus will now be on managing a pool of virtual resources and how to partition physical servers and administer virtual machines. Your server resources (VMs) are no longer tied to a physical servers, and you have the ability to move a running VM from one physical server to another without missing a beat (If you haven't seen VMotion in action, all I can say is sweet!). This chapter delves into VirtualCenter, which is one of the coolest products on the market for managing a virtual infrastructure. For those organizations that are strapped for cash we will also explore using the ESX Server Management Interface (MUI) for administering the virtual infrastructure. Topics such as systems requirements, VirtualCenter terms and concepts, performance management, as well as security, will be covered. You will learn how to deploy templates and perform migrations of VMs with VMotion. VMware has created three product categories that are as follows:

- **vManage** Includes VMware VirtualCenter with VMotion
- **vPlatform** Includes VMware ESX Server, Virtual SMP, GSX Server, Workstation and ACE
- **vTools** Currently includes VMware P2V assistant

VirtualCenter is the focus of this chapter, but in reality you'll probably use a hybrid approach to manage your virtual infrastructure because depending on your task, the MUI may work just fine, whereas other times VirtualCenter will be your best option. So once you get over the learning curve associated with this paradigm shift you'll start grasping the most efficient way to manage your virtual infrastructure. We hope this chapter points you in the right direction and increases your knowledge of the products.



## 9.2. VMware Tools for Managing the Virtual Infrastructure

The following is a brief introduction to the main tools you'll use to manage your virtual infrastructure. Depending on the size and scope of your VM deployment, you may not require VirtualCenter at this point and time. Without VirtualCenter, however, you will need to become very familiar with the MUI, the Service Console, and VMware Remote Console. It would be in your best interest to become familiar with those interfaces regardless of whether you deploy VirtualCenter or not. The three interfaces you'll utilize to manage your environment are:

- VMware VirtualCenter
- VMware Management Interface (MUI) and the Service Console
- VMware Remote Console

Now we'll give a brief description of each of the management interfaces you'll be using.

### 9.2.1. ESX Server Management Interface and the Service Console

The ESX Server Management Interface is the out-of-the-box administration option available to you to administer your virtual infrastructure. When you're working with a small VM deployment such as a test lab, these tools will be adequate and cost-effective. When you start consolidating and migrating more physical servers into a production virtualized environment, using the out-of-the-box tools to administer your environment will be difficult to manage in the long run. Security becomes more complex, and the ability to understand the placement and function of all your VMs will grow cumbersome. The MUI and Service Console give you the following capabilities:

- Review system status information for VMs hosted on ESX Server
- Start, stop, suspend and reset VMs
- View memory usage and performance statistics for ESX Server and VMs
- Launch remote console sessions for VMs
- Configure start-up options for VMs and ESX Server
- Modify virtual hardware
- Create new VMs from scratch or from a template
- Modify your ESX Server configuration
- Manage VMFS volumes and Service Console files and directories
- Change hardware allocations

Probably the biggest drawback with native tools is the fact that you can only view the resources hosted on a single ESX Server, and you don't have the ability to migrate running virtual machines from one host to another without a third-party tool. With the introduction of other third-party management agents and their management console, you can get a more detailed view of your environment without having to deploy VirtualCenter, but none of those options give you the same capabilities and flexibility that VirtualCenter does at this time. That may change in the very near future. This chapter goes into more detail on using the out-of-the-box tools, and we'll cover topics like integrating ESX Server authentication with Active Directory, setting security on VMs to allow different levels of access, and monitoring performance. For more information on creating and deploying templates, please refer to [Chapter 3](#). We'll cover templates with VirtualCenter in this chapter.

### 9.2.2. VMware VirtualCenter

VirtualCenter is a Windows-based application with a database back end used to centrally manage ESX Servers, GSX Servers, and the virtual machines they run. Both ESX and GSX







## 9.3. Understanding VirtualCenter Terms and Concepts

[Chapter 1](#) introduced terms and concepts related to server virtualization, VMware-specific terms and concepts. This chapter either adds to or reinforces previously mentioned terms and concepts with the goal of this section being terms and concepts that are unique to VirtualCenter. We hope this will have us all using the same vocabulary.

### 9.3.1. Managed Components

VirtualCenter would have no reason to exist if it had nothing to manage. The following components are managed by VirtualCenter:

- **Virtualization Platform** The VMware products used to create virtual machines include ESX Server, GSX Server, and VMware Workstation.
- **Host (Managed Host)** These physical servers run either ESX or GSX server and actually run virtual machines.
- **Virtual Machines** This guest operating system, as well as its associated applications that run on virtualized hardware, runs on the ESX or GSX Server (Host).
- **Datastores** These are the various storage locations for associated virtual machine files. They include VMFS, SAN, and Local Storage.

### 9.3.2. Software Components

The software components of VirtualCenter are the underlying technologies that work together to provide the services that make VirtualCenter so good. VirtualCenter Application consists of the following software components:

- **VirtualCenter Server** The server service that actively communicates with the managed host and writes data to the VirtualCenter Database. This service coordinates actions with VMs and managed hosts.
- **VirtualCenter Web service** This optional service can be installed with VirtualCenter Server. This is a required component for third-party applications that use the VMware SDK Application programming interface.
- **VirtualCenter Database** This database is used to store persistent information on the status of managed hosts, VMs, and VirtualCenter users. It can be installed locally or on a remote database server.
- **VirtualCenter Client** Windows management interface that can be installed on a Windows desktop or server and used to manage the VirtualCenter Application.
- **Virtual Center Agent** This software component is installed on each managed host. It collects, communicates, and performs actions received from the VirtualCenter Server.
- **VMotion** A cool feature that allows running virtual machines to move from one managed host to another without a loss in service.

### 9.3.3. VirtualCenter Organizational Hierarchy

[Table 9.1](#) covers the organizational hierarchy and objects that exists within VirtualCenter.

Table 9-1. VirtualCenter Containers and Objects	
Container Type	What Is It For?
Server Farms	There is only one "Server Farms" group. It can contain other farm groups or farms.
Farm Group	Can be used to further organize or secure your virtual





## 9.4. Installation and Configuration of VirtualCenter 1.2

This section covers the system requirements of each component that makes up the VirtualCenter deployment. The components that make up a VirtualCenter deployment include the VirtualCenter Server, VirtualCenter database, managed host and the VirtualCenter client. This section also covers the database setup and installation of VirtualCenter 1.2. In addition, this section covers the installation and how to add licenses to the VirtualCenter Server. First, the following concepts will be discussed:

- VirtualCenter 1.2 System Requirements
- Setting up Microsoft SQL Server for VirtualCenter 1.2
- Installing VirtualCenter

### 9.4.1. VirtualCenter Server 1.2 System Requirements

This section covers the requirements for the VirtualCenter Server, clients and managed host. You will need to ensure these requirements are met so that VirtualCenter will function properly.

#### 9.4.1.1. VirtualCenter Server

The following section covers the requirements for the VirtualCenter Server, including hardware, supported OS versions, and supported databases and licenses.

##### 9.4.1.1.1. Recommended Hardware

VMware recommends the following hardware for your VirtualCenter Server, and it provides additional recommendations depending on the number of managed hosts and VMs you will manage in your virtual infrastructure:

- At least 2GB of memory if fewer than 50 hosts will be managed. If more than 50 hosts will be managed, it is best to have at least 3GB of memory. For environments with more than 100 managed hosts and 2,000 VMs, your server should have at least 4GB of memory.
- At least a 2.0-GHz Pentium IV processor. Two processors are recommended if more than 25 hosts will be managed.
- At least one 10/100 NIC (Gigabit NIC recommended)

##### 9.4.1.1.2. Supported OS Versions

VirtualCenter Server 1.2 supports Windows 2003 Standard, Enterprise and Web Editions; Windows 2000 Standard and Advanced Editions; and Windows XP Professional.

##### 9.4.1.1.3. Databases Support

VirtualCenter can reside on a remote system if SQL Server or Oracle is to be used. Microsoft Access (default) is not recommended for production use; Access is best for trial and demonstration purposes only. VirtualCenter can be installed on the following database platforms:

- Microsoft SQL 2000 or SQL 7.0
- Oracle 8i or 9i

##### 9.4.1.1.4. Licensing Requirements

You'll need one license for each VirtualCenter Server.

#### 9.4.1.2. VirtualCenter Client

The VirtualCenter Client is the user interface used to manage the VirtualCenter Server. It can be installed on the VirtualCenter Server or on a Windows desktop. The following requirements should be met before installing the VirtualCenter Client:

- Minimum of 256MB of memory (512MB recommended)
- Windows NT (SP6a), Windows 2000 (all versions), Windows XP (all versions), and Windows 2003 (all versions)






## 9.5. Using VirtualCenter to Administer Your Virtual Infrastructure

VirtualCenter is the most versatile and essential administrative tool for managing a virtual infrastructure available today. A single VirtualCenter server deployment enables you to manage hundreds of managed hosts and thousands of virtual machines. Even in the largest deployments, you can easily scale VirtualCenter servers to meet all the administrative needs of your virtual infrastructure. This section outlines the most common administrative tasks you'll need to perform within VirtualCenter to maintain your virtual machines and manage new deployments of virtual machines from templates to cloned VMs. The installation and security of VirtualCenter have already been discussed. This section focuses on provisioning new servers (VMs), customizing templates, managing performance and alarms, and migrating virtual machines between hosts using cold migration and VMotion (hot migrations).

### 9.5.1. Creating Virtual Machine Groups

We'd like to quickly cover the creation of virtual machine groups. This task is very similar to creating the other organizational components in VirtualCenter. VM Groups can be created within farms or other VM groups. To create a new VM group, simply right-click on a farm or virtual machine group and select **New Virtual Machine Group**, as shown in [Figure 9.48](#)

**Figure 9-48. A Popup Menu to Create New Virtual Machine Group**



This will create a new virtual machine group that you can name based on the types of VMs it will contain or the function of the VMs in the group. Now would be a good time to set the appropriate permissions for the VM group.

### 9.5.2. Managing Virtual Machine Groups and Virtual Machines

Once you have created your organizational hierarchy of farm groups, farms, and virtual machine groups, you can now focus on deploying new VMs from templates and clones or building VMs from scratch. You also have the ability to migrate machines with VMotion or move them between physical hosts to satisfy performance needs or to perform maintenance on a managed host. This section covers creating new VMs, cloning VMs, and creating templates. First, let's cover the difference between the three. When creating a new VM, you'll need to specify all the virtual hardware devices and install a guest OS into the VM, unless you build a VM and attach to an existing virtual disk (which will be the example we use in this chapter). Cloning lets you copy an existing virtual machine and customize it. It does require the VM you are cloning to be powered down. Cloning is a great option if you have a standard VM build that meets your needs and you need to make a copy of it for a new VM. You can use the cloning option as long as you can afford your source VM to be powered off temporarily. When you create a template VM, you have the ability to store the template and build a repository of templates that you can use to provision standard server builds without impacting a running VM. This is the downside of cloning.

For example, let's say you just created a Web server VM that is now in production. You installed virus scanning, backup software, IIS, OS patches, and so on. You now decide to clone this virtual machine and use it to create your next Web server.

To clone the VM you will need to power it off. If you had a template of this server, you could create a new Web server without having to power off a production machine. Both options achieve the same goal, but do it in a different way. If you know you'll have to provision a standard Web server on an ongoing basis, it would make sense to create a template of the VM. If you just need to make a copy of a known good VM to create a new server, then cloning is a good option.







## 9.6. Using the Management Interface and Service Console to Administer the Virtual Infrastructure

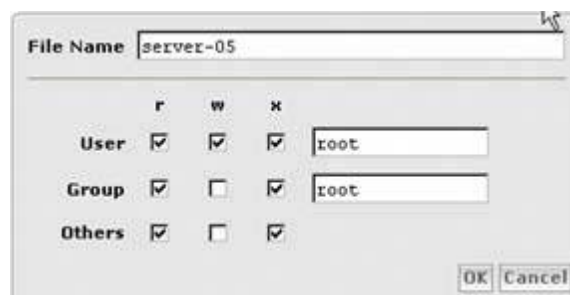
This section focuses on the day-to-day administration of your ESX server using the out-of-the-box administrative tools. Previous chapters of this book showed how to install ESX Server, create new VMs, and configure networking using the Management Interface. This chapter focuses on maintaining an existing ESX Server installation by modifying permissions on a virtual machine, integrating ESX Server with Active Directory, and monitoring performance. The specific tools we'll be using are the MUI and Service Console. Even if you have VirtualCenter deployed in your environment, some tasks can be accomplished just as easily with the MUI, Service Console, or Remote Console. You'll also find that certain tasks can be performed only from the MUI or Service Console. Examples include:

- Setting advanced options like masking LUNs
- Creating and formatting new VMFS volumes
- Creating and activating SWAP files
- Viewing system logs files
- Modifying the system startup profile

### 9.6.1. Virtual Machine Security in ESX Server

The ESX Server security model is much different from the VirtualCenter Security model. For one, security is not inherited, so it must be explicitly defined. Also, the permissions are applied differently. On ESX Server a directory/file has an owner (User), an affiliated group, and Others (everyone else). An example of the permissions is shown in [Figure 9.142](#). In this example the user who owns the file is root, and the affiliated group is root. The user has read, write, and execute permissions, and the group and others have read and execute permissions.

**Figure 9-142. ESX Server Permissions**



[Table 9.9](#) provides a breakdown of each set of permissions and what they will allow you to perform on a virtual machine.

Table 9-9. ESX Server permissions	
Permissions on VM files	Rights
R	View the state of VM in the MUI
W	Modify VM hardware
X	Perform power operations on a VM and use the Remote Console



## 9.7. Summary

This chapter focused on administering the virtual infrastructure with VMware VirtualCenter. A whole book could easily be written to cover the VirtualCenter application, so there are still a lot of areas that were not covered in this chapter. We did cover the installation and configuration of Virtual Center and the different security roles that exist within the application. We walked through several options for creating Virtual Machines, such as creating new VMs, cloning, creating templates, and deploying new VMs from templates. We also covered alarms, triggers, and actions. We also discussed how VirtualCenter allows you to migrate a running VM using VMotion without having to power the VM down.. You can also perform a cold migration of a VM, which allows you to relocate VMs to another managed host within the same farm.

This chapter also covered integrating ESX Server with Active Directory by configuring the Pluggable Authentication Module and using ESX Server's built-in native performance-monitoring tools. We also covered the concepts of proportional shares and minimum and maximum resource values.



## 9.8. Solutions Fast Track

### 9.8.1.

#### 9.8.1.1. Tools for Administering the Virtual Infrastructure

- ☒ The VirtualCenter Client is the interface used to manage the virtual infrastructure through the VirtualCenter Server.
- ☒ The Management Interface (MUI) and Service Console are used to administer ESX Server with or without VirtualCenter in the environment.
- ☒ The VMware Remote Console is used to remotely connect to and manage VMs. It's similar to a KVM switch in your data center, but is implemented as software.

#### 9.8.1.2. Understanding VirtualCenter Terms and Concepts

- ☒ Managed components include virtualization platforms, managed hosts, virtual machines, and datastores.
- ☒ Software components include the VirtualCenter Server, VirtualCenter Web service, VirtualCenter Client, VirtualCenter database, VirtualCenter agents, and Vmotion.
- ☒ Organizational components include server farms, farm groups, farms, hosts, virtual machine groups, and virtual machines
- ☒ Functional components include inventory, scheduled tasks, templates, alarms, and events.

#### 9.8.1.3. Installation and Configuration of VirtualCenter 1.2

- ☒ VirtualCenter supports Microsoft Access, SQL Server, and Oracle database platforms.
- ☒ Installing VirtualCenter requires an ODBC connection to the VirtualCenter database.
- ☒ VirtualCenter is a Windows-based application that requires the .NET Framework.
- ☒ VirtualCenter requires licenses for the VirtualCenter Server, VMotion, and VirtualCenter Agents.

#### 9.8.1.4. Using VirtualCenter to Administer Your Virtual Infrastructure

- ☒ VirtualCenter can be used to manage hundreds of managed hosts and thousands of virtual machines.
- ☒ Clones can be created from a VM in the VirtualCenter inventory or from a template.
- ☒ Managed hosts have alarms for CPU, memory, and state.
- ☒ The following roles exist within VirtualCenter: VirtualCenter administrator, virtual machine administrator, virtual machine user, and Read-Only user.
- ☒ VMs have alarms for CPU, memory, state, and heartbeat.
- ☒ You can schedule the following tasks within VirtualCenter:
  - ☒ Deploy a virtual machine from a template
  - ☒ Clone an existing virtual machine
  - ☒ Change the power status of a virtual machine
  - ☒ Migrate a virtual machine
  - ☒ Migrate a virtual machine with VMotion
  - ☒ Change resource settings of a virtual machine

#### 9.8.1.5. Using the Management Interface and Service Console to Administer the Virtual Infrastructure

- ☒ ESX Server can be integrated with Active Directory by configuring the pluggable authentication module.
- ☒ The esxtop and vmkusage utilities are used to monitor the performance of VMs on your ESX Server.
- ☒ You use CPU, memory, and disk shares to allocate resources to virtual machines.
- ☒ You can specify minimum CPU and memory requirements for VMs that are sensitive



## 9.9. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** I can't add a new virtual hard drive to my virtual machine. What should I do?

**A:** You may not have the appropriate permissions to add virtual hardware to the VM, or the VM may be powered on and won't allow the addition of virtual hardware while powered on.

**Q:** I just created a new virtual machine and attached to an existing disk. Every time I power on the VM, I get a blue screen. What could the problem be?

**A:** Make sure the OS type that you selected when you created the VM matches the OS type installed on the virtual disk. For example, if you selected Windows 2003 Standard for the VM when you created it, but the existing virtual disk has Windows 2000 Standard installed, you'll probably encounter issues.

**Q:** I followed the directions for Active Directory authentication, and I still can't log in using an account and password in AD. What should I do?

**A:** Check the time on your ESX Server and Active Directory and make sure they're synced up. Kerberos is very sensitive to being out of sync and could reject credentials if the time is not within specific limits. You can also review the System event logs for clues.



# Chapter 10. Best Practices, Gotchas, and Common Problems

Solutions in this chapter:

- Best Practices
- Gotchas
- Common Problems
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions

## 10.1. Introduction

In this chapter, we discuss some of the best practices mentioned in previous chapters as well as others not yet covered. Also highlighted are some gotchas and problems that are common to those new to the virtual world. We'll show how to easily avoid them.



## 10.2. Best Practices

In this section, the best practices in the following areas will be explored:

- ESX Server hardware
- ESX Server software
- Virtual machine builds
- Administrative tasks

### 10.2.1. ESX Server Hardware

In this section, we'll discuss the best practices related to ESX Server hardware, including configuration, storage management, and documentation.

#### 10.2.1.1. Purchasing Supported Hardware

When considering hardware for your virtual infrastructure, the first thing you want to consider is the hardware supported by VMware. Most major computer hardware vendors, including Dell, HP, and IBM, all have platforms that are supported by VMware. To ensure you're purchasing supported hardware, read the *Hardware Compatibility Guide* which you can obtain from the following URL: [www.vmware.com/pdf/esx\\_systems\\_guide.pdf](http://www.vmware.com/pdf/esx_systems_guide.pdf).

You should also familiarize yourself with the minimum and maximum system requirements outlined in the *ESX Installation Guide*, which is available at [www.vmware.com/pdf/esx25\\_install.pdf](http://www.vmware.com/pdf/esx25_install.pdf).

You can avoid many problems and save loads of time by purchasing supported hardware. Nothing is more frustrating than troubleshooting hardware issues only to find out from VMware Support that the piece of hardware you're having an issue with isn't supported.

After you receive your hardware, burn it in for a while. There are tools you can use to benchmark or simply install an operating system onto it and thereby keep it active for a few days.

#### 10.2.1.2. Multiple NICs

Although it isn't absolutely mandatory, it's highly recommended you get multiple NICs for your ESX Server. It's possible to have the Service Console NIC use a VMnic, which is great to test with and so on. But if this is a production ESX Server, or for that matter a serious test ESX Server, use multiple NICs. NICs are inexpensive and for network-intensive virtual machines, you may want to allocate a separate NIC to achieve greater serviceability. Placing the NIC for the Service Console on a different network (subnet) than your virtual machines is also a recommended best practice and another reason to have them on separate physical NICs.

#### 10.2.1.3. Memory

ESX Server uses memory efficiently and allows you to overallocate the physical memory you have to your virtual machines. That said, if you were to over-spec any aspect of your ESX Server, we'd recommend you choose to have more memory than, say, the fastest CPUs. By having enough memory for both the Service Console and your virtual machines, you can prevent excessive paging (thrashing) which causes performance degradation.

If you plan on loading systems management software on your ESX Server itself, you may want to increase the amount of memory you give the Service Console. This is especially true if you plan to run a large number of virtual machines on your ESX Server. Although not stated in any documentation we've read, we believe a best practice for memory in high-end production ESX Servers (especially these days since memory prices are low) is to allocate a maximum amount of memory to the Service Console, the maximum being 800MB.

#### 10.2.1.4. Documentation

Document the hardware configuration of your ESX Server. Keep that document up-to-date if you make any changes or upgrades to the ESX Server hardware. The document can be as elaborate or basic as you want, but it is to make it accurate. If you have a document versioning product such as Microsoft's SharePoint Server, keep the document in there and allow the versioning feature to track any changes you or others make to the document.

#### 10.2.1.5. Server Sizing

If you're doing a server consolidation project, you need to gather the historical utilization trend information from the physical servers you're going to virtualize. This data allows you to better gauge the performance requirements of your applications and helps you size your ESX Server and virtual machines appropriately.





## 10.3. Gotchas

The key to a smooth migration from a physical to a virtual environment is avoiding the pitfalls, or gotchas. In this section, we'll discuss best practices that will ensure that your migration to the virtual world goes smoothly. We focus on the biggest gotcha of all: Microsoft's lack of support for VMware.

### 10.3.1.

#### 10.3.1.1. Virtual Evolution

Virtual evolution is probably the most difficult aspect of moving into a virtual environment. It's the transition from the physical to the virtual world. The evolution for some is easy since it makes so much sense. For others, trappings of the physical world slow the process of thinking virtually, but even for these folks, once they embrace the technology, they never want to go back to installing operating systems on physical hardware and then layering applications, drivers, firmware, and so on... and then only achieve 10- to 15-percent utilization from the server.

#### 10.3.1.2. Microsoft Support

Undoubtedly, the largest gotcha out there right now is Microsoft's support of VMware, or should we say lack of support. Until very recently, Microsoft's stance was that if you were experiencing issues with a Microsoft application or operating system that had been virtualized on a VMware virtualization platform, you had to re-create the same issue on a physical server. We've been asked many times by clients about this and how they could possibly proceed with rollouts of production ESX Servers given Microsoft's unwillingness to support their products. Our answer has been that Dell, HP, and IBM have all agreed to support operating system issues as well as some application issues if you've purchased the hardware and software through them and have active support contracts. These hardware manufacturers have seen the benefits of VMware and both heard the needs of their customers and responded appropriately.

Microsoft's perspective was somewhat understandable since they too have virtualization products: Microsoft Virtual Server and Virtual PC. Although these products are decent, neither comes close to rivaling the power and flexibility of ESX Server. Microsoft was a little late in the game on this offering, but I wouldn't take my eye off Microsoft if I were you. Anyway, as we already intimated in this chapter, there has been a shift in Microsoft's support position for VMware. If you have a Premier Support contract intact with Microsoft, the company will support you to "commercially reasonable efforts" before having you move the problematic application or operating system on either their virtualization product or onto a physical machine. We've included the current support options verbatim, taken from the following URL on VMware's Web site: [www.vmware.com/pdf/ms\\_support\\_statement.pdf](http://www.vmware.com/pdf/ms_support_statement.pdf). Be able to answer this question about support if ever queried.

### **Designing & Planning... Support for Microsoft Software in VMware Virtual Machines**

The rapid increase in production deployments of VMware virtual infrastructure solutions by enterprises worldwide often leads to inquiries from customers on the support status of Microsoft products (operating systems and applications) that are used within VMware virtual machines. This sidebar summarizes the range of support options available to VMware customers intending to run licensed Microsoft operating systems and applications within VMware virtual machines. The support options vary, depending on how customers purchase the VMware and Microsoft software.

1. For customers who purchase OEM VMware products with Dell hardware and Gold Support, the vendor provides end-to-end support, including the VMware software and certified OEM licensed Microsoft operating systems that are run within virtual machines.
2. For customers who purchase VMware products with HP hardware, the vendor provides end-to-end support, including the VMware software and







## 10.4. Common Problems

Even if your initial physical-to-virtual migration is glitch-free, you are bound to encounter some common problems. In this section we tell you what to do when you cannot boot from your operating system ISO or when your virtual machine's performance is glacial.

### 10.4.1.

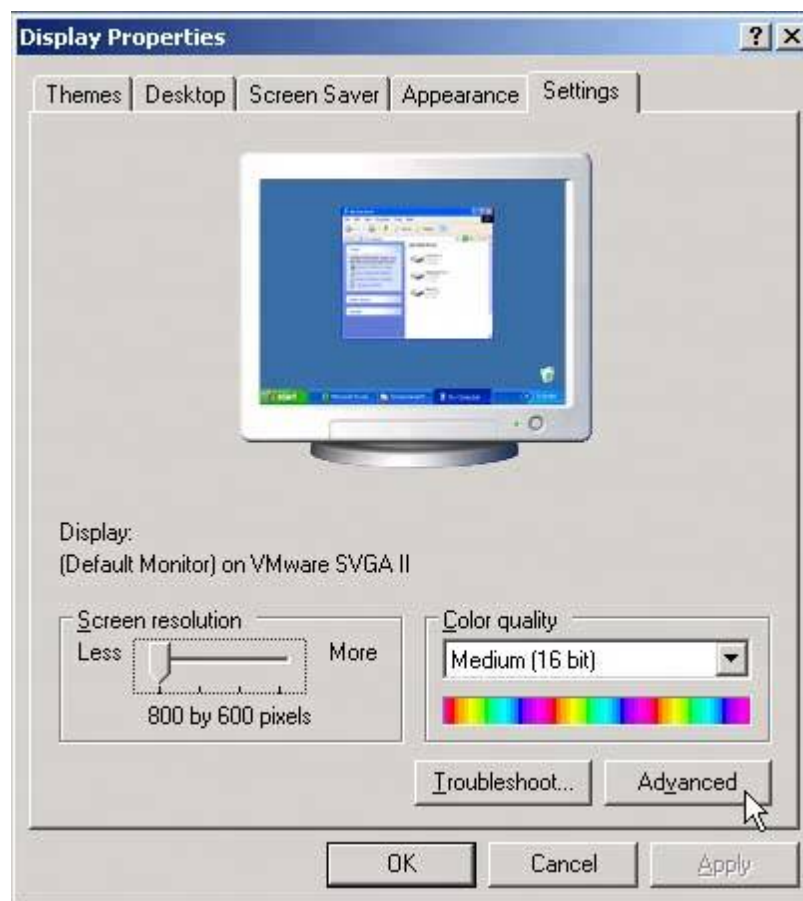
#### 10.4.1.1. Being Unable to Boot from Your Operating System ISO

If you're unable to boot from your operating system ISO, it's most likely because the ISO you created isn't bootable. Make sure you follow the appropriate instructions for creating **bootable ISO files** that are given in the application you use to create ISO images.

#### 10.4.1.2. Sluggish Windows

Let's say you just built a Windows virtual machine and you're receiving very sluggish performance: the mouse barely moves, applications are slow to open, and so on. A fix for this is cranking the Hardware Acceleration on the virtual machine itself. To do this right-click your desktop and select **Properties**. Go to the **Settings** tab and click the **Advanced** button (see [Figure 10.2](#)).

**Figure 10-2. Virtual Machine Display Properties**



Select the **Troubleshoot** tab and adjust the **Hardware Acceleration** slider to **Full**. Click **Apply** and your virtual machine should start responding faster (see [Figure 10.3](#)).

**Figure 10-3. Adjusting the Hardware Acceleration**





## 10.5. Summary

Best practices are an ever-evolving list of do's and don'ts. As VMware's ESX Server gets deployed evermore ubiquitously, this list will lengthen. We have presented here some of the things to do, avoid doing, and consider when deploying ESX Server and virtual machines. You'll find this a great place to start, but keep in mind that within your environment this list will expand. Incorporate best practices where you can and you'll make your transition into a virtual infrastructure more painless and easier to manage while reaping the uptime, scalability, and flexibility that ESX Server provides.

## 10.6. Solutions Fast Track

### 10.6.1.

#### 10.6.1.1. Best Practices

- ☑ Buy your supported ESX Server hardware from a major vendor such as Dell, HP, or IBM. Ensure you spec out your hardware according to recommendations from both the *ESX Installation Guide* and the *Systems Compatibility Guide*.
- ☑ Practice installing, configuring, breaking, and rebuilding your ESX software. You want a solid ESX Server when it's deployed into production. The way to get this type of solid build is to practice, practice, practice. Once you've got it down, document it.
- ☑ Use clones of your template virtual machines as often as you can. Your templates should have all the unnecessary services turned off which provides greater scalability for your ESX Server.

#### 10.6.1.2. Gotchas

- ☑ Learning to think virtually takes time, practice, and education.
- ☑ Know and follow support issues surrounding Microsoft and follow the issue as it evolves, as we're sure it will.

#### 10.6.1.3. Common Problems

- ☑ Follow the procedures from the manufacturer of the tool you use to create bootable ISOs. Don't forget, ESX Server is Linux, by placing a CD-ROM into the drive and then from the Service Console typing **cat /dev/cdrom > cdname.iso**.
- ☑ To remedy a sluggish Windows VM, always ensure you set the Hardware Acceleration slider to Full.
- ☑ Don't forget to restart your virtual machines after each reboot of your ESX Server. You can specify this to happen automatically and in any order you prefer.

## 10.7. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** What is the most important aspect of deploying ESX Server and virtual machines?

**A:** The ability to provide service at least as good as that which you had with physical servers. With adherence to best practices, you should easily be able to provide this and exceed it by a phenomenal degree.

**Q:** ISOs have been mentioned in several of the chapters. Why are they so important?

**A:** Good administrators strive to never leave their chairs, cubes, or offices. If they do, they may just run into a user (we're kidding, of course). Having a library of ISOs can help you in this endeavor. In addition, they run a lot faster than regular CDs.

**Q:** What will happen if I deploy systems management software on the ESX Server itself?

**A:** If you're going to do this, make sure you allocate enough memory so that the application doesn't impact the performance of your ESX Server and thus your virtual machines. Also, use a system management package that is supported and test your installation on your ESX Server thoroughly. If you notice a degradation of performance, contact the systems management software vendor. The major vendors have instructions and best practices for deploying their products onto ESX Server.

**Q:** What is the most important best practice to follow?

**A:** The one that ensures your smooth evolution into a virtual infrastructure education. Learn every aspect of a virtual environment to the best of your ability. Be passionate about it and you'll reap the rewards. See the recommended reading list, troll the VMware Web site regularly, as well as other Web sites dedicated to virtualization such as p2v.net, vmguru.com, and virtual-strategy.com.

# Chapter 11. VMware Under the Covers

Solutions in this chapter:

- Understanding the ESX Server File and Directory Structure
- The ESX Server Boot Process
- Interaction between the Service Console, VMkernel, and Virtual Machines
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions

## 11.1. Introduction

In [Chapter 2](#), this book covered the installation of VMware ESX Server. This chapter goes into further detail regarding the directory structure, configuration files, boot process, and the interaction between the Service Console, VMkernel, and virtual machines. VMware ESX Server uses a modified version of Red Hat Linux 7.2 for the boot process and to run the Service Console. Each component serves a specific purpose, and each interacts with the others to provide the services and processes necessary to run and manage virtual machines. This chapter explains the system architecture of ESX Server to ensure you have a thorough understanding of the underlying components and their dependencies. Another goal of this portion of the book is to help fill in the blanks for individuals without a Linux background so you gain a better understanding of ESX Server and better manage your virtual infrastructure.





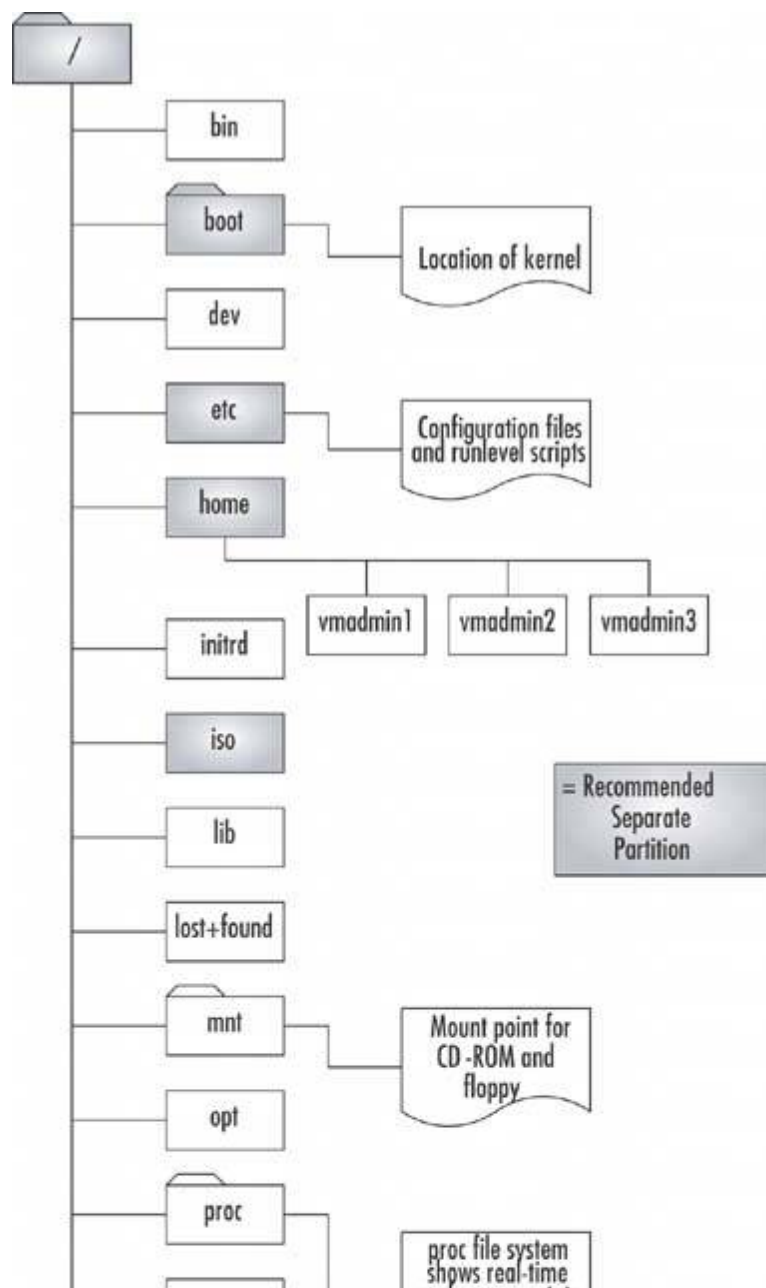
## 11.2. Understanding the ESX Server File and Directory Structure

Whether you are working at the physical ESX Server Service Console or connecting to the Service Console remotely, you need to know several things about navigating the directory structure and accessing specific files that are critical to the function of ESX Server. This section covers the directory structure of the Service Console. Understanding the location of directories and configuration files will help you in configuring, troubleshooting, and maintaining ESX Server.

### 11.2.1. The Service Console Directory Structure

The directory structure of the Service Console contains many of the usual Linux directories, plus a few additional directories that are recommended for, or created by, the ESX Server installation. A typical directory structure for your ESX Server Service Console is shown in [Figure 11.1](#). VMware recommends creating separate partitions for `/`, `/boot`, `/home`, `/tmp`, `/var`, and `/vmimages`. This helps isolate problems in the event that any of these partitions become full, and keeps them from overwriting files in the root file system. The purpose of several of these directories is explained in this section. Your directory structure may look a little different from the one shown in [Figure 11.1](#), depending on which partitions or directories you created for storing ISO images, templates, and virtual machine configuration files.

**Figure 11-1. The Service Console Top-Level Directories**







## 11.3. The ESX Server Boot Process

Several boot loaders are used on Linux systems, such as the Grand Unified bootloader (GRUB) and the Linux Loader (LILO). ESX uses LILO as the boot loader and has system components that expect the presence of LILO as the boot loader, so don't replace LILO with another boot loader, or your server may experience problems. The configuration parameters for the bootloader are contained in **/etc/lilo.conf** in a human-readable format, but the actual bootloader is stored in a binary format on the boot sector of the default boot disk. This section explains the boot process of ESX Server, as well as how to load the VMkernel and configuration files.

### 11.3.1. High-Level Boot Process for ESX Server

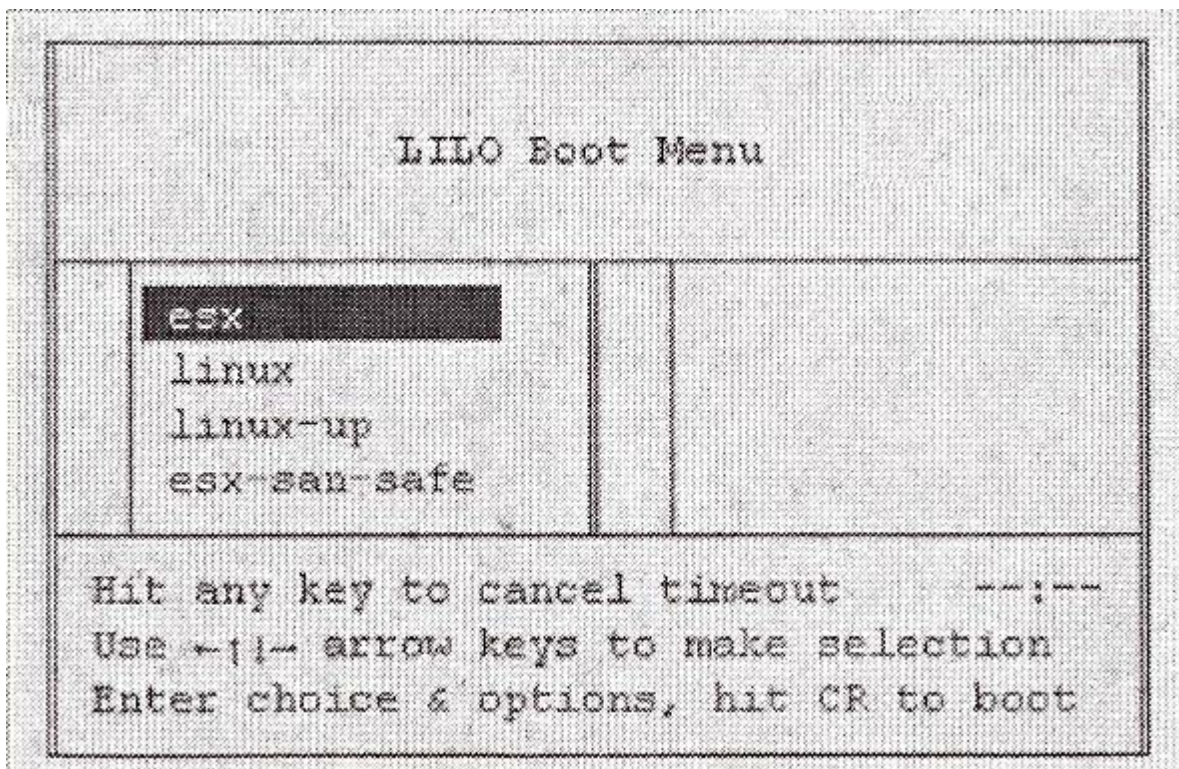
1. BIOS is executed on the server.
2. BIOS launches LILO from the default boot drive.
3. LILO loads Linux Kernel for the Service Console.
4. The Service Console launches VMkernel.
5. MUI Server is started.
6. Virtual machines can then be launched by VMkernel and managed through MUI.

### 11.3.2. Detailed Boot Process

As you can see in [Figure 11.3](#), esx is the default boot image that loads automatically after the timeout period. This is actually configured in the `/etc/lilo.conf` file shown in [Figure 11.4](#) on the line **default=esx**. The Linux kernel for the Service Console is loaded in the lowest part of memory when it is started and occupies the amount of memory specified during the installation of ESX Server. If you look at the line in the `/etc/lilo.conf` file shown in [Figure 11.4](#) that reads **append="mem=272M cpci=0;\*;1:\*;2:\*;3:\*;6:\*;"**. This shows that the Service Console occupies the first 272MB of memory on the server. [Figure 11.5](#) shows a screen shot from the MUI where the Reserved Memory is set in the **Options|Startup Profile** for the server.

If the Reserved Memory option is changed in the MUI, it will update the `/etc/lilo.conf` file and write those changes to the Master Boot Record/Boot Sector in binary format.

**Figure 11-3. The LILO Boot Menu**







## 11.4. Interaction between the Service Console, VMkernel, and Virtual Machines

This section covers the interactions and functions of the Service Console, VMkernel, and virtual machines. The previous section explained the process of booting ESX Server and loading all the necessary components to create the virtualization layer for running virtual machines and managing ESX Server. Once the VMkernel is loaded and running the Service Console, each one provides services to virtual machines. The following sections cover more of the specifics of these interconnected relationships.

### 11.4.1. Service Console Services

The Service Console is a modified version of Red Hat Linux 7.2 that is actually managed by the VMkernel much like a virtual machine. The Service Console may also be referred to as the Console OS or VMnix. It is responsible for managing all the processes related to each virtual machine. Specifically, the Service Console is responsible for mouse, keyboard, and screen (MKS) processes for virtual machines, as well as access to CD-ROM and floppy drives. The Service Console also provides command-line management of ESX Server and allows for the integration of thirdparty utilities for systems management, disaster recovery, and high availability.


#### 11.4.1.1. Virtual Machines Processes

To see how the Console OS manages the VM processes for ESX Server, take a look at [Figure 11.13](#). The first column shows the owner of the process, the second column shows the process ID, the third column shows the parent process, and the last column shows the command for that process. To determine the primary process for a specific VM, you can open the properties of the VM in the MUI and locate the PID shown in [Figure 11.14](#). Once you know the PID for the VM, you can then determine all the processes related to that virtual machine in the Service Console by using the `ps ef |grep vmware` command. In [Figure 11.14](#), the PID for the VM is 1908. If you look in [Figure 11.13](#) and find *process id 1908* in the second column with a parent process of *1* in the third column that is running `/usr/lib/vmware/bin/vmware-vmx`, then, bingo, that is your virtual machine. You can also trace additional processes related to your VM by looking at the parent child processes. Again, in [Figure 11.13](#) if you look at the *process id 1909*, which is a child process of *1908*, it is running `vmware-mks` (mouse, keyboard, and screen) for that VM. The main point of this is to show how the Service Console manages the processes for the virtual machines.

**Figure 11-13. Output of the `ps ef |grep VMware` Command to Show VMware Processes Managed by the Console OS**

```
[root@esx-server01 root]# ps -ef |grep vmware
root      990      1  0 20:08 ?        00:00:12 /usr/sbin/vmware-serverd -s -d
root     1908      1  0 20:23 ?        00:00:00 /usr/lib/vmware/bin/vmware-vmx -
root     1909    1908  0 20:23 ?        00:00:00 vmware-mks -A 11 -D 13 -S -L /tm
root     1914    1908  0 20:23 ?        00:00:00 vmware [Floppy]
root     1915    1908  0 20:23 ?        00:00:00 vmware [ide0:0]
root     1916    1908  0 20:23 ?        00:00:00 /usr/lib/vmware/bin/vmware-vmx -
root     1917    1916  0 20:23 ?        00:00:03 /usr/lib/vmware/bin/vmware-vmx -
root     1918    1916  0 20:23 ?        00:00:00 /usr/lib/vmware/bin/vmware-vmx -
```

**Figure 11-14. Partial Screenshot for VM from the MUI Identifying Process ID**



### 11.4.2. VMkernel Services

VMkernel manages the virtualization layer for virtual machines, which includes memory, CPU, disk,





## 11.5. Summary

Having a detailed understanding of the directory and file structure, the boot process, important processes, and interaction between VMkernel and Service Console components can greatly aid in the troubleshooting process in the event your ESX Server encounters problems. By understanding the directory structure, you will know where to locate files critical to the function of ESX Server and understand the general purpose of those directories. Knowing where configuration files are and which ones to edit if you need to reconfigure your server is a very important component for managing both your virtual server and virtual infrastructure in an efficient manner. Such information allows you to reconfigure the following settings if and when they're needed:

- Service Console IP
- Default gateway
- DNS Server
- Hostname
- Default boot timeout

Understanding the relationship between the Service Console processes and VMkernel services will help you locate your running VM's processes and view realtime performance data for specific VMs on the Service Console. Understanding the interaction between all of these components is critical to troubleshooting your virtual environment.

## 11.6. Solutions Fast Track

### 11.6.1.

#### 11.6.1.1. Understanding the ESX Server File and Directory Structure

- ☑ The directory structure is very similar to a standard Linux installation.
- ☑ Virtual machine configuration files should be stored in the home directory of the user who will manage them, or a logical directory structure should be created for storing the configuration files. This will help to organize files so that security can be set appropriately for users and groups who will manage the Virtual Machines.
- ☑ Understanding the directory structure can help with the troubleshooting process because you will know the location of critical files and directories.

#### 11.6.1.2. The ESX Server Boot Process

- ☑ ESX Server uses a modified version of Red Hat Linux 7.2 for booting and running the service console.
- ☑ ESX Server uses LILO as the boot loader. Configuration settings for LILO are stored in the `/etc/lilo.conf` file.
- ☑ The default runlevel for ESX Server is `runlevel3` and should not be changed.
- ☑ Understanding the boot process is the key to having a detailed understanding of your ESX Server architecture.

#### 11.6.1.3. Interaction between the Service Console, VMkernel, and Virtual Machines

- ☑ The Service Console controls VM's virtual console mouse, keyboard, and screen (MKS) processes.
- ☑ The Service Console manages the initialization and power state of VMs and access to the CD-ROM and floppy drives.
- ☑ The Service Console manages all processes related to virtual machines (vmware-vmx).
- ☑ The Service Console provides management functions for ESX Server, including HTTP and SNMP. It also offers authentication and a control API, which allows management of virtual machines and resource allocations.
- ☑ The VMkernel creates a world for the virtual machine and assigns a world ID.
- ☑ The VMkernel manages the virtualization layer for virtual machines and manages access to physical memory, the chipset, networks, and the hard disk.
- ☑ The VMkernel is the ultimate authority on scheduling what runs and when.
- ☑ Virtual machines depend on the Service Console and VMkernel for resources.

## 11.7. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** Which version of VMware ESX Server supports Boot from SAN?

**A:** ESX Server 2.5 supports Boot from SAN but has the following limitations: ESX server had to be installed in boot from SAN mode, the HBA can only be used by the Service Console, LUN masking should be used to restrict other ESX servers from the boot LUN, the HBA of the boot LUN must be a QLogic HBA, and the boot LUN must be the lowest numbered LUN controlled by the storage processor.

**Q:** Where are log files for VMware ESX Server written to?

**A:** `/var/logs/vmware`

**Q:** What do you do if you forget the root password of the Service Console?

**A:** You will need to boot into single-user mode from the Service Console by selecting linux from the LILO boot menu and appending `-s` to your boot choice. This will boot the console into single-user mode and will allow you to use the `passwd` command to change the root user password.

**Q:** Why do I still see processes for my virtual machine when running the `ps` command on the Service Console even though my virtual machine is powered down?

**A:** If there is still a virtual console session running for your VM, you will still see processes associated with it for mouse, keyboard, and screen (MKS) even though it is powered off.

**Q:** I accidentally unmounted the VMFS volume on my ESX Server. How can I re-mount the volume without re-booting

**A:** You can type `mount -t vmfs vmfs /vmfs`

# Chapter 12. Backup Strategies for ESX Server and Virtual Machines

Solutions in this chapter:

- Backing Up the Service Console
- Backing Up Virtual Machines
- Best Practices
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions

## 12.1. Introduction

This chapter focuses on backup strategies and the different approaches available to ensure you can recover your Service Console and VMs after a disaster. VMware provides greater flexibility in backing up virtual servers compared with traditional backup methods; however, that flexibility comes at the cost of a greater degree of complexity. VMware offers two general approaches for backing up your server environments, each with its own benefits and limitations. It's important to thoroughly understand the application model of VMware in order to select the approach best suited to your environment. This model touches a wide variety of factors, including:

- Host operating systems
- Guest operating systems
- SAN/NAS hosting
- Network bandwidth
- Resource utilization
- Current backup equipment and capabilities

### **Designing & Planning... Disaster Recovery and Business Continuity Considerations**

Backups are the cornerstone of any business/disaster recovery or business continuity plan without backups, you have nothing to recover. Having the ability to recovery quickly using a current and reliable backup enables a business to continue operating, if perhaps at a reduced level of service. So make certain that your operations strategy considers carefully the short-term and long-term manageability of your backup strategy.

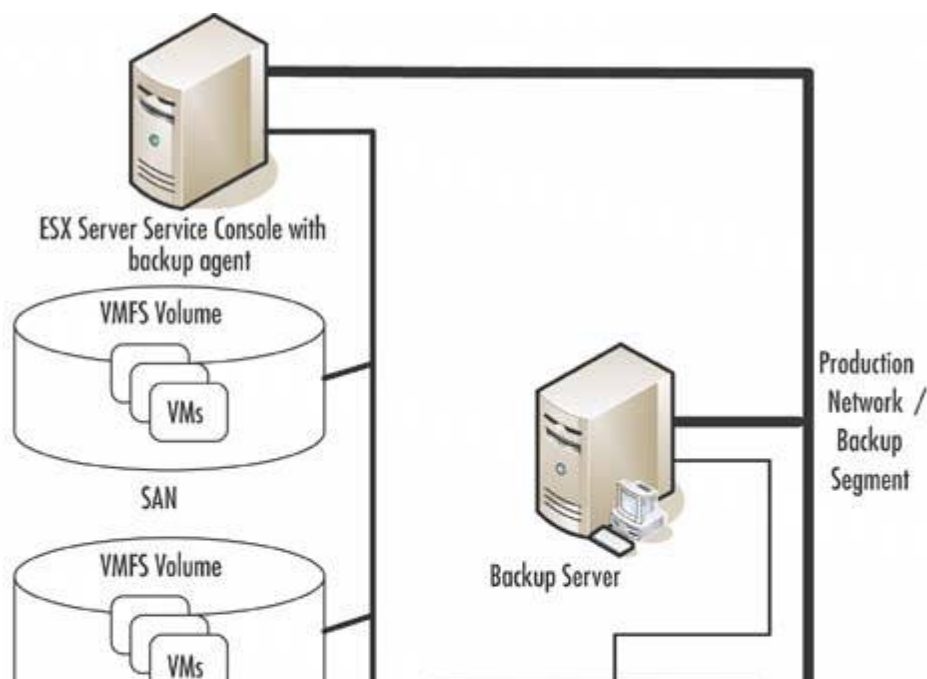


## 12.2. Backing Up the Service Console

The Service Console is a modified version of Linux that can run agents and backup software created for the Linux platform. To back up the Service Console, you have the option to install backup agents on the console and perform the backup over the network from a backup server, or you can install backup server software on the Service Console and back up to a local tape drive. (This method is not recommended because it utilizes Service Console resources that are critical to the function of ESX Server.) The state of your Service Console should not change very often and therefore will probably not require full backups as frequently as your VMs. [Table 12.1](#) lists the backup options for the Service Console. [Figure 12.1](#) shows an example of backing up the Service Console over a network.

Table 12-1. Different Methods for Backing Up the Service Console	
Backup Method	Comments
Install backup server software on Service Console with local tape drive	Utilizes Service Console resources which is not recommended
Install local tape drive but assign the SCSI tape device to virtual machines. Install backup server software on a VM to back up Service Console locally (uses the virtual network internal to ESX Server).	The backup server is a VM that requires the Service Console to be up and running in order to function. If the Service Console needed to be restored, this option would require importing the backup tape to another backup server and restoring the Service Console over the network before the backup server VM could be brought online. Final analysis: not recommended.
Install backup agents on Service Console and back up over the network to an external backup server.	This is a standard backup strategy for most enterprises. This is the recommended backup approach for the Service Console.

**Figure 12-1. Backing Up the Service Console over a Network**









## 12.3. Backing Up Virtual Machines

For the purpose of backup, virtual machines can be considered the same as physical machines. The backup software that operates on physical machines works just as well on virtual machines. Generally speaking, whatever tools and processes you currently have in place for your physical machines can be used to back up your virtual machines.

However, in addition to backing up the operating system, application software and data, you can also back up a server's virtual disk files as well. Moreover, you can back up an entire virtual machine (hardware, operating system, applications, and data) as a single entity.

The following subsections detail these two approaches: the physical server approach and the virtual machine approach.

### 12.3.1. Backing Up Virtual Machines Like Physical Servers

A virtual machine can be backed up using most commercially available backup agents. Similar to backing up physical servers, this backup approach can specify files, directories, shares, and so on. A virtual machine can back up itself and/or other virtual machines using standard backup technologies and approaches. A virtual machine can be backed up locally or across the network. Each virtual machine must be licensed for the appropriate agents used in this approach, however.

This method has several well-served points that follow conventional backup wisdom and requirements. The foremost lies in the ability to perform file-level backups and restores; you back up what you want, when you want, and how you want. This approach is generally considered to be the foundation of any backup environment because you have very finite control points in the backup and it allows for continuous operation of the server/services during the backup.

Always remember that a virtual machine should be considered as another device in your production environment. It may lack in physical occupation, but provides many advantages that would encumber non-virtual machines. Remember that physical devices have physical constraints in terms of power, environmental occupation, hardware management, and tape media resources. The case for virtual machines becomes very attractive when one considers the opportunities to leverage existing technologies for backup media, including SAN, NAS, tape, and DVD approaches that are not limited to one physical datacenter facility or even a facility owned by the client (in terms of outsourcing the backups to an ISP type IDC). Based on the opportunities presented, virtual disks and files can reside on any number of servers, SANs, and NASes at the local or distributed datacenter level. All virtual disks can be recovered to any ESX Server, regardless of the physical hardware or base architecture of the recovery facility or resources. In the following examples, a general comparison is made and the pro and cons of each strategy are listed.

#### 12.3.1.1. Pros

The pros to backing up virtual machines like physical servers include the following:

- General backup approaches apply; allows incremental, differential backup.
- The maintenance windows are extended to include any time allowed by the operational schedule. This approach may be well founded in your daily operations and strategy for intra-day backups as prescribed by regulated or unregulated legal requirements. This approach permits the archiving of applications without virtual machine shutdown.
- Leverages existing backup techniques
- This approach is generally deemed best for application data. Remember that regulations under HIPAA and any credit card data environments may not consider the abilities of virtual technologies in their current state.

#### 12.3.1.2. Cons

The cons to backing up virtual machines like physical servers include the following:





## 12.4. Best Practices

Best practices are developed from the review of many installations and implementations. They are not construed as the only approach to any situation, but rather a culmination of general approaches for most situations.

### 12.4.1. General Recommendations for Backups <sup>[9]</sup>

© VMware Education Services, Virtual Infrastructure with ESX Server and VirtualCenter, Student Manual, ESX Server 2.5 and VirtualCenter 1.2

The following approaches are based on the foundational applications of the VMware toolset. Some of the recommendations summarized in this list are based on requirements of VMware, and others are based on experience within the backup technology and planning methodology. Care should be taken to construct your own deployment and use of these recommendations based on the particular and specific needs and requirements of your operating environment.

- Separate a virtual machine's system data from its application data, control the location of frequently changing data on a separate file system of media (such as SAN, separate directory structure, and so on). Maintain a backup environment of these data sets using your normal backup or SAN's repository strategy.
- Back up the virtual machine's systems disks via the Service Console while in a suspended or powered down mode of operation.
- Back up the virtual machine's data disks via backup agents in those virtual machines. This approach is accomplished by using media backup technology or a combination of SANs replication technology.
- A snapshot approach can be administered but should be done only after considering the risk exposure introduced by this approach. Use snapshot backups with care. Remember that these snapshots impose risk due to the open file conditions. Snapshots require that redo logs be maintained at all times. Based on the type of snapshot dynamics, multiple redo logs may be implemented. Restoration of a snapshot requires that the redo logs be committed in a chronological order. Although the advantages of snapshots allow for minimal downtime opportunities, a long-term plan that evaluates the risk versus the reward should be considered before deploying this approach. Applications and services should be quiesced, suspended, or shut down if possible during the snapshot events. There should always be a cold backup performed and maintained on a regular basis in order to support a recovery in case of unplanned events.
- Store data on physically separate disks. This approach not only makes backups more flexible but also allows for more flexible deployment from templates. You can configure the applications in your templates to keep their data on separate disk drives, and then provide a new blank disk for data whenever the virtual machine is deployed.
- Store application data on virtual disks that are separate from those containing system images.
- Use backup agents inside guest operating systems for application data.
- Use Service Consolebased backups for system images, or plan to redeploy from a template.

VMware has compiled the results of many approaches to backing up virtual machines, as well as the associated pros and cons attached to each particular process. The preceding listing basically consolidates all approaches into a general best practice. Additionally, as a particular production environment changes, we may recognize the need, or acquired ability to change the backup environment based on evolving conditions. The previous recommendation appears to hold merit regardless of the current technology evolution stage. "If the data is changing, economics dictate to back it up relative to the dynamics of the frequency and volume of change." This adage also applies to the virtual machine construction. As the applications or operating systems change, make backups to reflect the dynamics and frequency of those changes. Keep in mind that someday you'll be called upon to restore those recovery images and, as always, time and technology will be your enemy. In order to reduce the effects of



## 12.5. Summary

To have a bulletproof implementation of VMware, you need to have a very complete set of backup approaches that meet your environmental and regulatory requirements in terms of SOX, Graham Leach Bliley, HIPAA, and other regulated issues. The best approach currently is to use two methods for backups, the Service Console and a data backup, either direct via SANs or indirect via tape and a licensed agent.

The SANs approach offers the fastest and most versatile approach. This method is more attractive if your environment supports offsite SAN-to-SAN replication. As described in earlier chapters, build a master set of virtual disks with templates. Store the masters with an easy to understand naming convention that describes the contents. Back up the virtual disks via the Service Console onto your local and remote SANs or other storage devices that can support VMFS file structure. Back up the data and applications separately, based on the size and time windows of your backup environment. Document the naming conventions used for backups/rollbacks and redos. Validate that all backups, direct and indirect, maintain their integrity. Perform some form of regular recovery tests for process and stability. Validate the procedures for master backup recovery and the preparation of the master for recovery (SID preparation, DNS modifications, network addressing, and host-naming).

Be cognizant of the issues associated with backing up a virtual machine in an active production environment. Practice recovering with redo files and document your internal processes before introducing the recovery processes into a production environment. Use the facilities of redo files to assist you in developing multiple deployments of the same template-based build. Establish well-documented practices when using recovery processes in the ESX environment. ESX offers power and flexibility in deploying and managing multiple recovery and production environments, as long as a set of documented processes are in place and managed according to your Business Continuity and Disaster Recovery plans.



## 12.6. Solutions Fast Track

### 12.6.1.

#### 12.6.1.1. Backing Up the Service Console

- ☑ The Service Console supports agents and backup server software that runs on the Linux platform.
- ☑ The Service Console can be backed up across the network or local tape drive.
- ☑ The Service Console can run backup server software and back up to a local tape drive, but this is not recommended because it requires higher resource utilization on the Service Console.

#### 12.6.1.2. Backing Up Virtual Machines

- ☑ Virtual machines can be backed up just like physical servers.
- ☑ You can back up virtual disk files to create a complete backup of the VM.
- ☑ Certain Linux and third-party applications can't handle files larger than 2GB.
- ☑ You should separate the VMs' system data from their application data, control the location of frequently changing data on a separate file system of media (for example, SAN, separate directory structure, and so on). Maintain a backup environment of these data sets using your normal backup or SAN's repository strategy.
- ☑ Back up VMs' systems disks via the Service Console in a suspended or powered-down mode of operation, a snapshot approach can be administered but should be done only after considering the risk exposure induced by this approach.
- ☑ Back up VMs' data disks via backup agents in those VMs. This approach works best if using media backup technology or a combination of SANs replication technology.
- ☑ Use snapshot backups with care. Remember that these snapshots impose risk owing to open file conditions. Snapshots require that redo logs be maintained at all times. Based on the type of snapshot dynamics, multiple redo logs may be implemented. Restoration of a snapshot requires that the redo logs be committed in a chronological order. Although the advantages of snapshots allow for minimal downtime opportunities, a long-term plan that evaluates the risk versus the reward should be considered before deploying this approach. Applications and services should be quiesced, suspended, or shut down if possible during the snapshot events. There should always be a cold backup performed and maintained on a regular basis in order to support a recovery in case of unplanned events.

#### 12.6.1.3. Best Practices

- ☑ Separate system data from application data in your backups.
- ☑ Develop a detailed backup schedule for your Service Console and VMs.
- ☑ Validate and test the backup and restore process to ensure it will work before you actually need it.

## 12.7. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** What software is qualified for use with VMware 2.5.1?

**A:** Please see [http://vmware.com/pdf/esx\\_backup\\_guide.pdf](http://vmware.com/pdf/esx_backup_guide.pdf) for an up-to-date list of software that's compatible with ESX.

**Q:** What SANs are compatible with ESX Server?

**A:** Please see [http://vmware.com/pdf/esx\\_SAN\\_guide.pdf](http://vmware.com/pdf/esx_SAN_guide.pdf) for details on ESX-to-SAN server compatibility.

**Q:** What are the certified backup tools for ESX?

**A:** Please see [http://vmware.com/support/esx25/doc/backup\\_tools\\_links.html](http://vmware.com/support/esx25/doc/backup_tools_links.html).

**Q:** Can I back up my entire virtual machine from the Service Console?

**A:** Yes, but it is not advised. The console services should be left alone to manage the entire virtual machine infrastructure. It's best that backups be performed by separating the applications and data from the operating systems because backups can become quite large very quickly. You should consider a backup of the environments that change frequently (data and applications) with an agent specifically designed to perform this function. A backup of those virtual disks that change infrequently and need VMFS formats should be backed up from the Service Console. These console-based backups should be performed in a powered-down or suspended approach (preferred), or a suspended environment using redo logs as an alternative. All backups performed from a systems console mode must be restored in an all-or-nothing approach. Recovery for a single file or directory can only be accomplished via a backup agent or from a backup to a SAN environment using a variety of different recovery/restoral techniques.

# Chapter 13. Common and Necessary Command-Line Interface (CLI) Commands

Solutions in this chapter:

- Using the Command-Line Interface
- Configuring and Managing ESX Server Service Console from the Command Line
- Managing Virtual Machines from the Command Line
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions

## 13.1. Introduction

This chapter focuses on the commands that are most likely to be used to manage your ESX Server from the Service Console. As mentioned in other chapters, the Service Console is a modified version of Red Hat Linux which is used to boot and load the VMkernel for VMware ESX Server and provide services such as the Apache web server for the MUI and command-line management which, of course, is the focus of this chapter. The target audience for this chapter is mainly Windows administrators who are not as familiar with the Linux command-line environment. It is not meant to cover every available command for Linux or to explain Linux in depth. Hopefully, this chapter can be used to provide a quick reference for the commands that will most likely be used at the Service Console command line. This chapter should also provide some insight to administrators on how commands can be used to effectively manage their ESX Server deployments. The Service Console has several commands that are unique to ESX Server such as *vmware-cmd*, *vmkfstools*, *vmkpcidivvy*, *esxstop*, and custom scripts that can be used to manage your virtual environment. These commands, as well as the standard Linux commands, are covered in this chapter and examples are provided.



## 13.2. Using the Command-Line Interface

In order to fully and efficiently manage your ESX Server, you need to become very familiar with the command-line interface of the ESX Server Service Console. Although the MUI can be used to perform the majority of the administrative tasks for your ESX Server, several tasks can be carried out only from the command line. This section walks you through the most common task you must perform at the Service Console. Specifically, you will learn the following:

- How to connect to the Service Console
- Command-line shortcuts
- Basic commands for files and directories
- How to move and copy files
- How to find files and search their content
- How to edit files
- How to manage users and groups
- How to manage file and directory ownership and permissions
- How to archive and compress files and directories
- How to mount file systems in the Service Console (ISO images, external file systems, and Virtual Disks)

### 13.2.1. Connecting to the Service Console

In order to use the command-line interface, you must connect to the Service Console with either an SSH client or work at the physical Service Console and use the virtual terminals that are available by using the **Alt + F2** through **Alt + F6** keys.

#### 13.2.1.1. Virtual Terminal

Working from the physical console, you can access the virtual terminals by pressing **Alt + F2** through **Alt + F6**. Pressing **Alt + F1** brings you back to the Service Console Status screen. Thus, you have five virtual terminals to work with at the physical console.

#### 13.2.1.2. Remote Access Card

It's highly recommended that you include a remote access card in your ESX Server, which will give you the ability to perform remote management. Some of the most common vendor solutions allow for remote management of the server with virtual power, virtual CD-ROM, and virtual floppy capabilities. Some of these solutions are listed in the next paragraph. Once connected, you can use the **Alt + F2** through **Alt + F6** commands to access virtual terminals as if you were at the physical server. Most will run through a web browser. The main reason this is recommended is because you'll have the ability to control virtual power and still maintain access to the console even if someone configured the Service Console incorrectly.

Major vendor remote access solutions include the following:

- HP Insight Lights Out (ILO)
- DELL Dell Remote Access Card (DRAC)
- IBM Remote Supervisor

#### 13.2.1.3. SSH Client

If you're working from a Linux workstation, you can connect to the Service Console remotely using an SSH client. For Windows-based clients, you can download the PuTTY ssh client for free.

#### 13.2.1.4. Telnet

If you have set your ESX Server security to Medium or Custom to allow telnet, you'll also be able to use a







## 13.3. Configuring and Managing ESX Server Service Console from the Command Line

The previous section of this chapter focused on some of the more common commands used to navigate the directory structure, to manage files and directories and manage users and groups. This section covers commands which are most applicable for managing your processes, devices, and the configuration of your ESX Server from the Service Console. The command line can be used to manage your ESX server VMFS volumes, device allocations, network configuration, running processes, and virtual machines. This section discusses the following:

- Managing processes and viewing performance statistics
- Managing physical and virtual disks
- Installing and managing software packages
- Managing devices and changing hardware and network settings

### 13.3.1. Managing Processes, Restarting Services, and Viewing Performance Statistics

From time to time you may find it necessary to view running processes for troubleshooting purposes. You may also need to kill a running process that is not running correctly. The following sections show the command-line tools you can use to manage processes, restart services, and view the performance data of your Service Console.

#### 13.3.1.1. Managing Processes

You can use the *ps* and *kill* commands to view your running processes.

To list all processes running on the Service Console, type

*ps ef*

Each process will be listed with its process ID (PID). If you want to list all processes shown with its parent process, type

*ps efH*

[Figure 13.27](#) shows a partial output of the command. Notice the PPID, which is circled. This is the process ID of the parent controlling process.

**Figure 13-27. Partial Output of ps Command**

```
[root@esx-server01 ~]# ps -efH
UID          PID    PPID    C  STIME TTY          TIME CMD
root          7        0  0  08:53 ?        00:00:00 [kupdated]
root          6        0  0  08:53 ?        00:00:00 [bdflush]
root          5        0  0  08:53 ?        00:00:00 [kreclaimd]
root          4        0  0  08:53 ?        00:00:00 [kswapd]
root          3        0  0  08:53 ?        00:00:02 [ksoftirqd_CPU0]
root          1        0  0  08:53 ?        00:00:03 init [3]
root          2        1  0  08:53 ?        00:00:00 [keventd]
root         10        1  0  08:53 ?        00:00:00 [kjournald]
root         77        1  0  08:53 ?        00:00:00 [khubd]
root        244        1  0  08:53 ?        00:00:00 [kjournald]
root        245        1  0  08:53 ?        00:00:00 [kjournald]
root        246        1  0  08:53 ?        00:00:00 [kjournald]
root        247        1  0  08:53 ?        00:00:00 [kjournald]
root        248        1  0  08:53 ?        00:00:01 [kjournald]
root        249        1  0  08:53 ?        00:00:00 [kjournald]
root        607        1  0  08:54 ?        00:00:00 [vmfs_flush]
root        918        1  0  08:55 ?        00:00:00 /usr/sbin/vmklogger
root        991        1  1  08:55 ?        00:09:01 /usr/sbin/vmware-serverd -s -d
root       1016        1  0  08:55 ?        00:00:00 syslogd -m 0
root       1024        1  0  08:55 ?        00:00:00 klogd -2
root       1116        1  0  08:55 ?        00:00:00 /usr/sbin/sshd
```





## 13.4. Managing Virtual Machines from the Command Line

Some actions on your virtual machines need to be performed from the Service Console. To give you an example, one time the authors had a virtual machine get stuck during a power-up operation in Virtual Center and we couldn't power it down from the Virtual Center client or the ESX Server MUI. So, we had to perform a hard power off from the command line to stop the VM. To manage virtual machines from the Service Console, use the *vmware-cmd* command. There are also scripts available to perform actions like backing up and restoring virtual disk files. In this section, we'll cover the following:

- Using *vmware-cmd* to manage virtual machines
- Viewing performance data for virtual machines with *esxtop*
- Adding new virtual blank disks and extending existing virtual disks
- Performing backups and restores of virtual machine virtual disks

### 13.4.1. Using *vmware-cmd* to Manage Virtual Machines

The *vmware-cmd* command allows you to perform many different tasks related to virtual machines. [Figure 13.33](#) shows the server operations that can be performed on the Service Console using *vmware-cmd*. As you can see from the list, you can register and unregister virtual machines and get/set resource variables.

**Figure 13-33. A List of Server Operations that Can Be Performed with *vmware-cmd***

```
Server Operations:
/usr/bin/vmware-cmd -l
/usr/bin/vmware-cmd -s register <config_file_path>
/usr/bin/vmware-cmd -s unregister <config_file_path>
/usr/bin/vmware-cmd -s getresource <variable>
/usr/bin/vmware-cmd -s setresource <variable> <value>
```

[Figure 13.34](#) shows the VM operations that can be performed from the Service Console using the *vmware-cmd* utility.

**Figure 13-34. A List of VM Operations that Can Be Performed with *vmware-cmd***

```
VM Operations:
/usr/bin/vmware-cmd <cfg> getconnectedusers
/usr/bin/vmware-cmd <cfg> getstate
/usr/bin/vmware-cmd <cfg> start <powerop_mode>
/usr/bin/vmware-cmd <cfg> stop <powerop_mode>
/usr/bin/vmware-cmd <cfg> reset <powerop_mode>
/usr/bin/vmware-cmd <cfg> suspend <powerop_mode>
/usr/bin/vmware-cmd <cfg> setconfig <variable> <value>
/usr/bin/vmware-cmd <cfg> getconfig <variable>
/usr/bin/vmware-cmd <cfg> setguestinfo <variable> <value>
/usr/bin/vmware-cmd <cfg> getguestinfo <variable>
/usr/bin/vmware-cmd <cfg> getid
/usr/bin/vmware-cmd <cfg> getpid
/usr/bin/vmware-cmd <cfg> getproductinfo <prodinfo>
/usr/bin/vmware-cmd <cfg> connectdevice <device_name>
/usr/bin/vmware-cmd <cfg> disconnectdevice <device_name>
/usr/bin/vmware-cmd <cfg> getconfigfile
/usr/bin/vmware-cmd <cfg> getheartbeat
/usr/bin/vmware-cmd <cfg> getuptime
/usr/bin/vmware-cmd <cfg> getremoteconnections
/usr/bin/vmware-cmd <cfg> gettoolslastactive
/usr/bin/vmware-cmd <cfg> getresource <variable>
/usr/bin/vmware-cmd <cfg> setresource <variable> <value>
/usr/bin/vmware-cmd <cfg> addredo <disk_device_name>
/usr/bin/vmware-cmd <cfg> commit <disk_device_name> <level> <freeze> <wait>
/usr/bin/vmware-cmd <cfg> answer
```



## 13.5. Summary

Understanding how to utilize the command line on the Service Console is a critical component of managing your virtual infrastructure. You can use the command line to manage the Service Console configuration, the VMkernel, and virtual machines. Although you can perform most tasks from the ESX Server MUI, you won't be able to manage every aspect of your server through the MUI. Specifically, the MUI does not give you the ability to search for files, search the contents of files for strings, or edit text files. The command line gives you the ability to compress and archive files as well as view and manage processes on your server. If you need to add additional software packages such as backup software or management agents this needs to be done from the Service Console command line as well.

The following is a list of tasks that can only be performed from the Service Console command line:

- Editing configuration files
- Starting and stopping services such as httpd and network
- Killing processes
- Searching the directory structure for files
- Searching files based on string values
- Installing third-party software
- Configuring Pluggable Authentication Modules
- Changing the Service Console IP address, subnet mask, default gateway, and host name
- Import vmdk files from VMware Workstation or GSX Server disk files
- Mounting CD-ROM and floppy driver

## 13.6. Solutions Fast Track

### 13.6.1.

#### 13.6.1.1. Using the Command-Line Interface

- ☑ You can connect to the Service Console via the physical console, SSH client, Remote Management Card, or Telnet (if security is set to Medium or Custom).
- ☑ You can use the Tab key to help auto-complete commands and file names to speed up typing.
- ☑ Using the Up- and Down-arrow keys allows you to cycle through previously typed commands.
- ☑ Using the command line, you can search for files, as well as for strings and patterns in files.
- ☑ The command line gives you the ability to archive and compress files.
- ☑ You can use either vi or nano to edit text files.

#### 13.6.1.2. Configuring and Managing ESX Server Service Console from the Command Line

- ☑ You can use the *vmkpcidivvy* command to manage device allocation between the Service Console and the VMkernel.
- ☑ The command line allows you to view the list of device drivers loaded in the VMkernel using the *vmkload\_modlist* command.
- ☑ You can use the *vmware-mount.pl* script to mount a virtual disk in the service console.
- ☑ You can use *vmkfstools* to import vmdks from VMware Workstation and GSX Server into ESX Server.
- ☑ Use the *top* command to view real-time performance data for the Service Console.
- ☑ You can manage the Service Console and VM processes from the command line, which allows you to kill processes if necessary.
- ☑ Use the Service Console to install software packages in RPM format.

#### 13.6.1.3. Managing Virtual Machines from the Command Line

- ☑ The *esxstop* utility can be used to view performance statistics of your virtual machines.
- ☑ You can use *vmware-cmd* to start, stop, reset, and suspend virtual machines from the command line.
- ☑ The *vmsnap.pl* script lets you back up a virtual machine's virtual disk while it is powered on.
- ☑ With the *vmres.pl* script, you can restore a virtual machine that was backed up with *vmsnap*.

## 13.7. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** I'm not able to connect to the Service Console over the network. What could the issue be?

**A:** You may have allocated the Service Console NIC to the VMkernel. Use *vmkpcidivyi* to reassign the NIC to the Service Console.

**Q:** I have a virtual machine that did not start up correctly, but now I can't power it down from the MUI or Virtual Center. How can I get this VM to shut down?

**A:** You can use the *vmware-cmd* utility to force a hard power down. The following syntax should work:

```
vmware-cmd /path-to-vm/vm-directory/vm.vmx stop hard
```

**Q:** I find using commands to be very difficult. Why can't I use X Windows on the Service Console?

**A:** You actually could run the X Windows system on the Service Console, but it will eat up valuable resources that are needed by the system to manage all the processes related to Virtualization. VMware specifically says not to run X Windows on the Service Console. So, it's best to just buck up and deal with it.

**Q:** Is there a way to mount the vmfs volumes if they accidentally get unmounted without having to reboot?

**A:** Yes. You can run *mountt vmfs vmfs /vmfs*.

**Q:** How do I check the speed and duplex setting of the Service Console NIC? Also, how do I change it if needed?

**A:** You'll need to cat out the eth0.info file for your type of adapter. This file can be found at */proc/net/type-of-nic/eth0.info*. To give you an example, our server has an Intel Pro 100 Nic for the Service Console, so for us to find the speed and duplex information we would type: *cat /proc/net/PRO\_LAN\_Adapters/eth0.info*.



# Chapter 14. A Peek at the Future: It's Totally Cool

Solutions in this chapter:

- Evolved Virtually You Have Become
- Bigger and Badder Brains
- ESX Server 3.0...Better the Beta
- Faultless Fault Tolerance
- The Rise of Utility Computing
- Your Moment of Xen
- Microsoft Virtual Server...Watch Your Six
- Virtual Storage Is the Place to Live
  - ☒ Summary
  - ☒ Solutions Fast Track
  - ☒ Frequently Asked Questions

## 14.1. Introduction

In this chapter, we'll take a look at what the future holds. We'll discuss topics ranging from the evolution of your physical environment into a virtual one to new versions of ESX Server and other virtualization products. We'll also take a glimpse at the storage on which your virtual machines may live, and see if the future of your virtual environment is as exciting as it is today. Getting a head start on these elements of virtualization is a step you don't want to miss.

## 14.2. Evolved Virtually You Have Become

By this chapter, we hope that you have at least one ESX Server humming efficiently away with a large number of virtual machines utilizing the resources of their host optimally. Maybe you even have VirtualCenter up and running and you have a library of templates from which you can provision servers more quickly than your SLAs demand. If you have more than one ESX Server and VirtualCenter, you may have witnessed the wonder that is VMotion. Perhaps you have dabbled with the power of `vmware-cmd` and `vmkfstools` and even created some basic scripts to manage and control your virtual machines. If so, you have taken the first steps toward virtual nirvana.

The further you move down the path of virtualization, the greater the rewards. Buzzwords like ROI and TCO will actually become meaningful as the tangible costs of running your server environment decrease while performance, stability, flexibility, and availability all increase.

## 14.3. Bigger and Badder Brains

Back in April 2004, VMware announced that its virtualization platforms would support 64-bit platforms within 18 months. Good as the company's word, VMware Workstation 4.5 and GSX 3.1 began experimental support on 64-bit processor hosts a couple of months after the announcement. VMware Workstation 5.0, recently released, now fully supports a 64-bit host, either AMD or Intel. As we've mentioned a number of times in this book, virtual machines perform better with more memory. The 64-bit architecture allows for huge increases in the amount of physical and virtual (swap) memory. Although ESX Server supports the AMD Opteron 64-bit chip, at the time of this writing, you could load ESX on a 64-bit host, in 32-bit mode only. With ESX 3.0 in a tightly controlled beta, we expect support for 64-bit platform hosts to take the possibilities of virtual infrastructures to a whole new level; thus, staging ESX on a 64-bit host in 32-bit mode may be a good idea.

VMware is definitely setting its sights on 64-bit computing. From a memory perspective alone, it's a technology that will greatly enhance consolidation ratios, thus making VMware's product have a higher ROI. Theoretically, 64-bit processors can address 16 Exabytes of memory, which is two orders of magnitude larger than a terabyte we're talking serious memory, here. Imagine having a virtual machine(s) running completely in physical (fault-tolerant) memory and writing to disk only during idle cycles...someday in the not-too-distant future this will happen. Now, combine 64-bit processing with dual-core (or more) technology and the power and brains of your ESX Server host increase significantly.

For virtual machines themselves, whether a 64-bit processor(s) is fully supported or not on the host (it will be), in the future four-way virtual SMP will be available. This, of course, will allow more robust applications such as databases to be virtualized. Some may say that if you're allowing a database access to all four of your host's CPUs, and your host is a four-way, then you shouldn't stay on a physical. Perhaps. But don't forget about the dual-core processors that are coming out. Soon your four-ways will be eight-ways, or more. Now, some vendors like Oracle have decided to treat dual core chips as two processors (for licensing purposes), while others like Microsoft have stated that they believe a dual core chip is one processor keep in mind some of the physical-to-virtual migration logic. If your current database server (say, Oracle) is running on a legacy four-way box (e.g., four P3 1.0GHz processors, which was smokin' for its day), and you move into a virtual environment running P4 3.2GHz or higher, you'll see a performance increase.

Speaking of the ESX Server 3.0 Beta, let's tackle that in the next section.

## 14.4. ESX Server 3.0...Better the Beta

ESX Server 3.0 is just around the corner. At the time of this writing, VMware slated its release for the second half of 2005. Since the VMWorld Conference is near the end of the year, this may prove to be the venue for the big release. We've heard rumors that the new version will be based on the 3.0 Linux kernel. We've also heard (amongst other things) that it will come with metadata enhancements, new operating system support, 64-bit support, four-way virtual-SMP, improvements in SAN storage, and iSCSI support.

Unfortunately, this is still all rumor. VMware has restricted the beta release to so few people that even the authors of this book, who do large scale virtualization consulting and are VMware VIP partners haven't been allowed to participate in the beta program. So although we've heard the aforementioned gossip from VMware system engineers in Palo Alto, we can't tell you exactly what the future holds for this product.

## 14.5. Faultless Fault Tolerance

With each release of ESX Server, support for clustered virtual servers has increased. We imagine this trend will continue with the next release. A recently published paper from VMware examines the possibilities of geographically dispersed clusters (see [Chapter 4](#)) utilizing EMC hardware and software. VMware in the past tried to show nearly equal support between the SAN manufacturers, so you can bet that new and interesting fault-tolerant solutions will become available, ensuring a greater degree of uptime and service availability. Support for iSCSI, of course, would create interesting possibilities, and in a report featuring EMC on iSCSI, EMC said that it "estimates cost differentials between iSCSI and FC (Fiber Channel) to be between 18 percent and 40 percent depending on the hardware used and number of servers connected." A cost-effective transport for storage supported by VMware would definitely open up more possibilities for virtual infrastructures.

## 14.6. The Rise of Utility Computing

Imagine if you will, the following: From a Web site, an application owner selects from a menu of options the specification for a server that she needs for a new application. She selects a hard-drive size, the amount of memory, the number of processors, and the operating system. She clicks the **Submit** button. From a small workflow process, approvals are gathered, licensing is paid for, additional questions are asked and answered and in a very short time a new server is provisioned. The application owner is sent an e-mail with a link that when she clicks it opens her new virtual machine. Instead of days, weeks, or months, you've cut server provisioning down to minutes or hours. When the application owner is done with her virtual machine, the resources are reabsorbed for future demand using a de-provisioning process.

With VirtualCenter, all of your ESX Servers are basically one pool of resources, CPUs, disks, memory, and so on. Utilization of each of these resources is easily monitored and availability is obtained. It's a small step to provisioning.

With the use of VMware's SDK, the possibilities of utility computing become closer and closer. Some time in the not-too-distant future, the ease of virtual machine provisioning will exist ubiquitously throughout IT departments.

Virtualization allows for the natural extension or evolution of this process. Some interesting technologies exist in VMware's ACE product, such as expirations and policy-based virtual machines that if applied to some future version of ESX Server would allow for tremendous utility computing and server provisioning capabilities.

## 14.7. Your Moment of Xen

Xen is a developing virtual platform for x86 platforms. The virtual machine monitor (VMM) or hypervisor was developed in Cambridge, England, and has morphed into a company called XenSource. According to the XenSource site, currently Xen only supports some Linux and does not support any Windows virtual machines. But in the future the company says it will. Xen is an open source product licensed under the GNU General Public License terms, so it's free. Xen is gaining a lot of support from the industry, so it'd be smart to increase your understanding of this product. The recent release of SUSE 9.3 comes with the most current release of Xen built into it.

For more information on Xen, go to [www.xensource.com](http://www.xensource.com).



## 14.8. Microsoft Virtual Server...Watch Your Six

The early code for Microsoft's Virtual Server product was purchased from Connectix in 2003. The product is a quasi-equivalent to VMware's GSX Server. Some of the major similarities are that both products, GSX and Virtual Server, can run on Windows as the host operating system (although GSX can run on Red Hat Linux as well). Both support only one virtual CPU per virtual guest. Microsoft Virtual Server and VMware's GSX, as well as the virtual machines, run as processes within the host operating system.

Microsoft Virtual Server runs well enough, and we've had very good stability with the virtual servers we created on the Microsoft Virtual Server platform. Recent announcements made by top brass in Microsoft suggest that they are serious about competing in the virtual space. In the future, Microsoft plans to have its hypervisor built into the operating system. By coincidence, we ran into a Microsoft TAM in the Oakland Airport. We chatted, and the conversation turned towards Virtual Server. The TAM suggested that Microsoft would soon be catching up to VMware, which we don't doubt. How and in what time frame, however, we're still uncertain about, and the fellow didn't offer up much in the way of details. It'd be a good idea to watch the evolution of this technology...someday it may rule the world.

For more details on Microsoft Virtual Server, go to [www.microsoft.com/win-dowsserversystem/virtualserver/default.mspx](http://www.microsoft.com/win-dowsserversystem/virtualserver/default.mspx).

## 14.9. Virtual Storage Is the Place to Live

For now, a SAN (virtual storage) is where you want your virtual infrastructure to live, if possible. It is a requirement for VMotion, and quality SANs provide stellar redundancy and availability. If you have one already, you're set; if not, try to get it in the budget (keep in mind, it'll be a major line item since the costs can be high). If you can budget for one but can't go whole hog, ensure you purchase one that in the future can be scaled without significant costs, like for an entire new enclosure. However, with some clever scripting, you can get some of the benefits of a SAN such as snaps and BCV-like restore possibilities even though you still don't have VMotion. VMware is owned by the SAN manufacturer EMC, so you can bet as SAN technology continually develops with more virtual capabilities, VMware and ESX Server will leverage those capabilities, allowing for greater amounts of flexibility in your virtual infrastructure.

## 14.10. Summary

The virtual evolution is happening in data centers and IT shops everywhere. The use of VMware's virtualization products has now moved out of the test and development shops into production. ESX Server is the virtualization platform on which departmental and data center servers and services should be run. The future of virtualization is exciting since growing support from many areas of the tech industry will spawn new products that will assist you in achieving a 100-percent uptime environment while reducing costs and complexities. ESX Server is at the core of your virtual infrastructure and provides you with the immediate tools you need to jump-start your virtual evolution. Many years from today, other virtualization products may equal ESX Server, but for now and for the foreseeable future, ESX Server is by far the leader in the marketplace and will remain so for quite some time.



## 14.11. Solutions Fast Track

### 14.11.1.

#### 14.11.1.1. Evolved Virtually You Have Become

- ☒ The road ahead involves constant learning, something which you should embrace.
- ☒ You've already taken the first steps to a virtual infrastructure. As you continue forward, you'll find ever-increasing options and possibilities, and you'll certainly want to virtualize what you can in your environment.
- ☒ The further you move down the path of virtualization, the greater the rewards. Buzzwords like ROI and TCO will actually become meaningful.

#### 14.11.1.2. Bigger and Badder Brains

- ☒ Virtual machines perform better with more memory.
- ☒ Sixty-four-bit processing and the increased capacity for memory inherent with this technology will elevate virtual possibilities far above today's levels.
- ☒ From a memory perspective alone, 64-bit computing is a technology that will greatly enhance consolidation ratios.

#### 14.11.1.3. ESX Server 3.0...Better the Beta

- ☒ VMware needs to expand its beta program and be more inclusive instead of exclusive.
- ☒ We've heard rumors that ESX Server 3.0 may be based on the 3.0 Linux kernel.
- ☒ We've also heard that it will come with metadata enhancements, new operating system support, 64-bit support, four-way virtual-SMP, improvements in SAN storage, and iSCSI support.

#### 14.11.1.4. Faultless Fault Tolerance

- ☒ With each release of ESX Server, support for clustered virtual servers has increased.
- ☒ In the past VMware tried to show nearly equal support between the SAN manufacturers, so you can bet that new and interesting fault-tolerant solutions will become available
- ☒ A cost-effective transport for storage supported by VMware would definitely open up more possibilities for virtual infrastructures.

#### 14.11.1.5. The Rise of Utility Computing

- ☒ With VirtualCenter, all of your ESX Servers are basically one pool of resources, CPUs, disks, memory, and so on.
- ☒ With the use of VMware's SDK, the possibilities of utility computing become closer and closer.
- ☒ Some interesting technologies exist in VMware's ACE product, such as expirations and policy-based virtual machines that if applied to some future version of ESX Server would allow for tremendous utility computing and server provisioning capabilities.

#### 14.11.1.6. Your Moment of Xen

- ☒ Xen is a growing virtualization platform that is gaining industry support. Keep your eye on it.
- ☒ According to the XenSource site, currently Xen only supports some Linux and does not support any Windows virtual machines.
- ☒ The recent release of SUSE 9.3 comes with the most current release of Xen built into it.

#### 14.11.1.7. Microsoft Virtual Server...Watch Your Six

- ☒ Microsoft's virtual platform is a decent one on which to run Microsoft operating systems. Although the product and its support isn't that old, Microsoft has always found a way of catching up. So definitely keep your eye on it.



## 14.12. Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the "**Ask the Author**" form. You will also gain access to thousands of other FAQs at ITFAQnet.com.

**Q:** How long has Xen been around?

**A:** Since 2004, and they're located in Palo Alto, California. Hmm...what other virtualization company is in Palo Alto?

**Q:** What is a hypervisor?

**A:** In VMware parlance, it's the virtualization layer.

**Q:** Why did VMware limit its beta of ESX Server 3.0 to so few?

**A:** Good question...We don't know, and we hope that VMware's future beta programs are opened up to a larger audience.

**Q:** If I can't get a SAN, will local storage with a RAID device be sufficient?

**A:** Absolutely. You won't get some of the cooler tools like VMotion, but being virtual on local storage is better than remaining physical.

# Appendix A. Recommended Reading

[Section A.1. Guides](#)

[Section A.2. Books](#)

[Section A.3. White Papers](#)

[Section A.4. Web Sites](#)



## A.1. Guides

*VMware ESX Server 2.5 Administrative Guide*

*VMware ESX Server 2.5 Installation Guide*

## A.2. Books

Dalheimer, Matthias Kalle, Terry Dawson, Lar Kaufman, and Matt Welsh.

*Running Linux*, Fourth Edition. O'Reilly, 2002

## A.3. White Papers

*Building Virtual Infrastructure with VMware*

[www.vmware.com/pdf/vi\\_wp.pdf](http://www.vmware.com/pdf/vi_wp.pdf)

*Using Raw Device Mapping*

[www.vmware.com/pdf/esx/esx25\\_rawdevicemapping.pdf](http://www.vmware.com/pdf/esx/esx25_rawdevicemapping.pdf)

*NIC Teaming*

[www.vmware.com/pdf/esx2\\_NIC\\_Teaming.pdf](http://www.vmware.com/pdf/esx2_NIC_Teaming.pdf)

*SAN Compatibility Guide for ESX Server 2.X*

[www.vmware.com/pdf/esx\\_SAN\\_guide.pdf](http://www.vmware.com/pdf/esx_SAN_guide.pdf)

*Best Practices Using VMware Virtual SMP*

[www.vmware.com/pdf/vsmp\\_best\\_practices.pdf](http://www.vmware.com/pdf/vsmp_best_practices.pdf)

*Best Practices for ESX Server 2*

[www.vmware.com/pdf/esx2\\_best\\_practices.pdf](http://www.vmware.com/pdf/esx2_best_practices.pdf)

## A.4. Web Sites

Read VMware's discussion forums for issues, questions, and solutions. Visit this public forum at [www.vmware.com/community/index.jspa](http://www.vmware.com/community/index.jspa).

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

['-C option \(vmkfstools command\)](#)

['-extendfile option \(vmkfstools command\)](#)

['-k command](#)

[.vmdk files](#)

[importing from VMware Workstation of GSX Server 2nd](#)

[naming](#)

[64-bit platforms 2nd](#)

[\ backslash](#)

# Index

[SYMBOL] [A] [B] [C] [D] [E] [F] [G] [H] [I] [J] [K] [L] [M] [N] [O] [P] [R] [S] [T] [U] [V] [W] [X] [Z]

- [Active Directory](#)
- [Active Directory authentication](#)
- [adduser command](#)
- [administrative tasks](#)
- [alarms 2nd](#)
  - [configuring resource allocations and 2nd 3rd 4 h](#)
- [algorithm for MAC address generation](#)
- [application virtualization](#)
- [applications](#)
  - [centralized repository/database of modeling](#)
- [archiving files/directories](#)
- [autostart](#)

# Index

[SYMBOL] [A] [B] [C] [D] [E] [F] [G] [H] [I] [J] [K] [L] [M] [N] [O] [P] [R] [S] [T] [U] [V] [W] [X] [Z]

- [backslash \(\\)](#)
- [backups 2nd 3rd 4th 5th 6th 7th 8th 9th 10th 11 h 12th 13th 14th](#)
  - [best practices for 2nd 3rd](#)
  - [for Service Console 2nd 3rd 4th](#)
  - [for VMs 2nd 3rd 4th 5th 6th 7th 8th](#)
- [bandwidth allocation](#)
- [BartPE 2nd 3rd 4 h 5 h 6 h 7 h](#)
  - [legal issues and](#)
- [best practices 2nd 3rd](#)
  - [server consolidation methodology and](#)
- [bg command](#)
- [/bin directory](#)
- [BIOS](#)
- [bonded NICs 2nd 3rd](#)
- [/boot directory 2nd](#)
- [boot process 2nd 3rd 4th 5th 6th 7th 8th 9th](#)
- [boot timeout 2nd](#)
- [build document 2nd 3rd](#)
- [bunzip2 tool](#)
- [bus sharing 2nd](#)
- [business case 2nd 3rd 4th 5th 6th 7th 8th 9th](#)



# Index

[SYMBOL] [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[capacity planning](#)  
[case sensitivity](#)  
     [configuration files and](#)  
[cat command](#) [2nd](#)  
[cat p-to-v process](#) [2nd](#) [3rd](#) [4 h](#) [5 h](#) [6 h](#)  
[cd /vm command](#)  
[cd command](#)  
[CD-ROMs](#)  
[chgrp command](#)  
[chown command](#)  
[Clone VM wizard](#) [2nd](#) [3rd](#) [4th](#) [5th](#)  
[clones](#) [2nd](#)  
[cluster heartbeat](#) [2nd](#)  
[clustering](#) [2nd](#) [3rd](#)  
     [adding cluster nodes and](#)  
     [future trends and](#)  
     [naming clusters and](#)  
     [types of](#)  
[command & command](#)  
[command-line management](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#)  
[command-line utilities](#) [2nd](#) [3rd](#)  
[commands](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#) [10th](#) [11th](#) [12 h](#) [13th](#) [14th](#)  
     [auto-completion for](#) [2nd](#)  
         [for files/directories](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#) [10th](#) [11th](#) [12 h](#) [13th](#)  
[compliance](#)  
[compressing files/directories](#)  
[configuration files](#) [2nd](#)  
[configuring for Active Directory](#) [2nd](#) [3rd](#) [4 h](#)  
[configuring virtual machines](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#) [10th](#) [11th](#) [12 h](#) [13th](#) [14th](#) [15th](#) [16 h](#) [17th](#) [18th](#) [19th](#) [20 h](#) [21st](#) [22nd](#) [23rd](#) [24th](#) [25th](#) [26 h](#) [27th](#)  
[28th](#) [29 h](#) [30th](#) [31st](#) [32nd](#) [33rd](#) [34th](#) [35 h](#) [36th](#) [37th](#) [38th](#) [39 h](#) [40th](#) [41st](#) [42nd](#)  
[cost estimates](#)  
[cost performance](#)  
[cp command](#)  
[CPU](#)  
     [virtualization and](#)  
[CPU affinity](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

- [data for 2nd 3rd](#)
- [data gathering 2nd 3rd 4th 5th 6th 7th 8th](#)
- [default gateway 2nd](#)
- [Dell Remote Access Card \(DRAC\)](#)
- [/dev directory](#)
- [device drivers](#)
- [devices](#)
- [directories](#)
  - [managing](#)
- [ownership/permissions for](#)
  - [managing](#)
- [directory structure 2nd 3rd 4 h 5 h](#)
- [disaster recovery](#)
  - [backups for](#)
  - [scripts for 2nd](#)
- [disk bandwidth 2nd](#)
- [disk I/O 2nd](#)
- [disk images of ISO 9660](#)
- [disk mode options](#)
- [disk space](#)
- [disks](#)
  - [virtualization and](#)
- [DNS server](#)
  - [IP address](#)
  - [settings](#)
- [drivers](#)
- [du command](#)
- [dual-boot systems](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[egrep command](#)

[encapsulation](#)

[EST \(External Switch Tagging\) 2nd](#)

[ESX Server 2.5](#)

[Discussion Forum](#)

[future trends in 2nd 3rd 4th 5th 6th 7th](#)

[host name for](#)

[installing 2nd 3rd 4th 5th 6th 7th](#)

[step by step 2nd 3rd 4th 5th 6th 7th 8th 9th 10 h 11th 12th 13th 14 h 15th 16th](#)

[integrating with Active Directory 2nd 3rd 4 h](#)

[managing from Service Console 2nd 3rd 4th 5th 6th 7th](#)

[raw device mapping and 2nd](#)

[specifications for 2nd 3rd 4th](#)

[ESX Server 3.0 Beta 2nd](#)

[esxstop command 2nd 3rd](#)

[/etc directory](#)

[/etc/... configuration files](#)

[/etc/passwd file](#)

[exit command](#)

[extensions to ISO 9660](#)

[External Switch Tagging \(EST\) 2nd](#)

# Index

[SYMBOL] [A] [B] [C] [D] [E] **[F]** [G] [H] [I] [J] [K] [L] [M] [N] [O] [P] [R] [S] [T] [U] [V] [W] [X] [Z]

- [fault tolerance 2nd](#)
- [file names 2nd](#)
- [files](#)
  - [managing](#)
- [find command 2nd](#)
- [findnic command 2nd](#)

# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[J](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)] [[X](#)] [[Z](#)]

[gateway \(Service Console\)](#)

[geo-clustering](#)

[get/set resource variables](#)

[Ghost 2nd 3rd 4 h 5 h](#)

[P2V Assistant 2 and](#)

[Gold Master](#)

[moving to ESX Servers](#)

[Grand Unified bootloader \(GRUB\)](#)

[grep command 2nd 3rd](#)

[groups](#)

[from command line](#)

[via MUI](#)

[guest operating system memory requirements](#)

[gunzip tool](#)

[gzip tool](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[Hal file 2nd](#)  
    [BartPE and](#)  
    [PowerP2V and](#)  
[hard drives](#)  
[head virtual\\_disk.vmkd | file command](#)  
[Heartbeat vmxnet](#)  
[history command](#)  
[/home directory 2nd](#)  
[host](#)  
[hostname 2nd 3rd 4 h](#)  
[HP Insight Lights Out \(ILO\)](#)

# Index

[SYMBOL] [A] [B] [C] [D] [E] [F] [G] [H] [I] [J] [K] [L] [M] [N] [O] [P] [R] [S] [T] [U] [V] [W] [X] [Z]

[IBM Remote Supervisor](#)

[id command](#)

[id username command](#)

[idealized hardware](#)

[ifconfig command](#) 2nd 3rd

[image process](#)

[source servers](#) 2nd 3rd 4th 5th

[target servers](#)

[implementation plan](#)

[independent software vendors \(ISVs\)](#)

[infrastructure](#)

[/initrd directory](#)

[installing](#)

[ESX Server 2.5](#) 2nd 3rd 4th 5th 6th 7th 8th 9th 10th 11 h 12th 13th 14th 15 h

[methods of](#) 2nd

[operating systems](#)

[Remote Console](#)

[VMware Tools](#)

[IP address \(Service Console\)](#)

[IP addresses](#)

[binding to VMNIC](#) 2nd

[changing](#)

[reviewing information about](#)

[ISO 9660](#)

[ISO images](#) 2nd 3rd 4 h 5 h 6 h 7 h 8 h

[mounting on Service Console](#)

[WinSCP and](#) 2nd 3rd

[isolation](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[jobs command](#)



# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[K scripts](#)  
[kernel modification](#)  
[kill command 2nd](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[/lib directory](#)

[LILO \(Linux Loader\) 2nd](#)

[Linux Loader \(LILO\)](#)

[load balancing 2nd](#)

[/lost+found directory](#)

[ls -l command 2nd](#)

[ls command](#)

[lsmod command](#)

# Index

[SYMBOL] [A] [B] [C] [D] [E] [F] [G] [H] [I] [J] [K] [L] **M** [N] [O] [P] [R] [S] [T] [U] [V] [W] [X] [Z]

- [MAC addresses 2nd 3rd 4th 5th](#)
- [makewhatis command](#)
- [man pages](#)
- managed hosts
  - [adding to VirtualCenter 2nd 3rd](#)
  - [performance data for 2nd](#)
- managing
  - virtual machines (VMs)
    - [from command line 2nd 3rd 4th 5th](#)
- [mapped drive setup 2nd 3rd](#)
- [memory affinity 2nd](#)
- [methodology 2nd 3rd](#)
  - [best practices and phases of \(list\)](#)
- Microsoft
  - [Network Load Balancing](#)
  - [SysPrep 2nd](#)
  - [virtual machine support and](#)
  - [Virtual Server](#)
  - [VMware support and 2nd 3rd](#)
- Microsoft Cluster Service
  - [account name/password for](#)
  - [raw device mapping and](#)
- [mii-tool 2nd](#)
- [mkdir command](#)
- [MKS processes \(mouse\)](#)
- [/mnt directory](#)
- [monitoring resources 2nd 3rd 4th 5th](#)
- [mouse](#)
- [MUI 2nd 3rd](#)
- [MUI Status Monitor](#)
- [MUI Web server](#)
- [mv command](#)
- [mware-authd \(Remote Console port\)](#)

# Index

[SYMBOL] [A] [B] [C] [D] [E] [F] [G] [H] [I] [J] [K] [L] [M] [N] [O] [P] [R] [S] [T] [U] [V] [W] [X] [Z]

[nano editor](#)

[network card setup](#) 2nd 3rd

[network labels](#)

[Network Load Balancing](#)

[network traffic shaping](#) 2nd

[network virtualization](#)

[New Template wizard](#) 2nd

[NewSID tool](#) 2nd 3rd 4th

[nfshaper module](#)

[NICs \(network interface cards\)](#) 2nd 3rd 4 h

[adding for cluster heartbeat](#) 2nd

[changing speed of](#)

[drivers and](#)

[mii-tool and](#) 2nd

[Non-Uniform Memory Access nodes \(NUMA nodes\)](#) 2nd

[Ntoskrnl file](#)

[NTP server](#) 2nd

[ntpd daemon](#)

[NUMA nodes \(Non-Uniform Memory Access nodes\)](#) 2nd

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

ODBC connection  
    [creating on VirtualCenter Server 2nd 3rd 4 h 5 h](#)  
[/opt directory](#)  
[ownership of files/directories 2nd 3rd 4th](#)  
ownership/permissions for  
    directories  
        [viewing/editing with MUI](#)

# Index

[SYMBOL] [A] [B] [C] [D] [E] [F] [G] [H] [I] [J] [K] [L] [M] [N] [O] [P] [R] [S] [T] [U] [V] [W] [X] [Z]

[p-to-v \(physical to virtual\) migration](#) 2nd 3rd 4th 5th 6th 7th 8th 9th 10th 11th 12 h 13th  
    [advanced configurations for](#) 2nd 3rd 4 h  
    [migration strategy planning and](#)  
    [table summarizing methods for](#) 2nd  
[P2V Assistant](#) 2nd  
[PAM \(Pluggable Authentication Module\)](#)  
    [parallelization clustering](#)  
    [partitions](#) 2nd  
[performance data](#)  
    [viewing](#)  
[permissions](#)  
    [for ESX Server](#) 2nd  
    [for files/directories](#) 2nd 3rd  
[physical disks](#)  
[physical extents](#)  
[physical NICs](#) 2nd 3rd 4th 5th 6th 7th  
[PI \(process ID\)](#)  
[piloting](#)  
[PlateSpin](#)  
    [PowerP2V](#) 2nd 3rd 4 h 5 h  
    [PowerRecon](#)  
[Pluggable Authentication Module \(PAM\)](#)  
[port groups](#) 2nd 3rd 4th 5th 6th 7th 8th 9th  
[portability](#)  
[PowerP2V](#) 2nd 3rd 4th 5th  
[PowerRecon](#) 2nd  
[/proc directory](#)  
[/proc/vmware directory](#)  
[process ID \(PID\)](#)  
[processes](#)  
    [managing](#) 2nd  
    [related to VMs](#)  
[project methodology](#) 2nd 3rd 4 h 5 h 6 h 7 h 8 h 9 h 10th 11th 12th 13 h 14th 15th 16th 17 h 18th 19th 20th 21st 22nd  
[ps -ef |grep vmware command](#)  
[ps command](#)  
[public mode](#)  
[PuTTY](#) 2nd 3rd 4th  
[PuTTY SSH client](#)  
[pwd command](#)

# Index

[SYMBOL] [A] [B] [C] [D] [E] [F] [G] [H] [I] [J] [K] [L] [M] [N] [O] [P] [R] [S] [T] [U] [V] [W] [X] [Z]

- [rationalization](#)
- [raw device mapping](#) 2nd
- [rc script](#)
- [readiness](#) 2nd
  - [migration strategies and](#)
- [Red Hat Linux 7.2](#)
  - [Service Console and](#)
- [Red Hat Package Manager \(RPM\)](#) 2nd
- [redo files](#) 2nd
  - [backups and](#) 2nd 3rd 4th
- [remote access cards](#)
- [Remote Console](#)
  - [changing port address for](#)
  - [loading](#)
- [remote UnixLinux computers](#) 2nd
- [resources](#)
  - [monitoring](#) 2nd
- [resources for further reading](#)
  - [backup strategies](#)
  - [device support](#)
  - [Microsoft support options for VMware](#)
  - [raw device mapping](#)
  - [Security Response](#)
  - [systems compatibility](#)
  - [VLAN tags](#)
  - [VLANs](#)
  - [VMware recommendations](#)
- [return on investment \(ROI\)](#)
- [risk reduction](#)
- [rm command](#)
- [ROI \(return on investment\)](#)
- [root directory](#) 2nd
- [route print command](#)
- [RPM \(Red Hat Package Manager\)](#) 2nd
- [rpm command](#)
- [run-level scripts](#)
- [runlevels](#) 2nd

# Index

[SYMBOL] [A] [B] [C] [D] [E] [F] [G] [H] [I] [J] [K] [L] [M] [N] [O] [P] [R] [S] [T] [U] [V] [W] [X] [Z]

- [S scripts](#)
- [S10network script](#)
- [S11vmware script](#)
- [S12syslog](#)
- [S55sshd](#)
- [S56xinetd script](#)
- [S91httpd.vmware script](#)
- [SANs \(storage area networks\)](#)
  - [backups and](#)
  - [ESX Server compatibility and](#)
- [screenshots](#)
- [scripts 2nd 3rd](#)
  - [for p-to-v migration](#)
- [SCSI drivers](#)
- [Secure Shell \(SSH\)](#)
- [secure shell daemon](#)
- [security](#)
  - [in ESX Server 2nd 3rd](#)
  - [in VirtualCenter 2nd 3rd](#)
- [server consolidation 2nd 3rd 4th 5th 6th 7th 8th 9th 10 h 11th 12th 13th 14 h 15th 16th 17th 18 h 19th 20th 21st 22nd 23rd 24 h 25th 26th 27th 28 h 29th 30 h 31st 32nd 33rd 34th 35th](#)
  - [defined 2nd 3rd](#)
  - [methodology for 2nd](#)
  - [people/process components of 2nd](#)
  - [planning for 2nd 3rd 4th 5th 6th 7th 8th 9th](#)
- [server consolidation database 2nd](#)
- [server sizing](#)
- [server virtualization 2nd 3rd 4th](#)
- [service agility](#)
- [service command](#)
- [Service Console 2nd 3rd 4 h](#)
  - [backing up 2nd 3rd 4th](#)
  - [connecting to 2nd 3rd 4th 5th](#)
  - [CPU management and 2nd](#)
  - [directory structure of](#)
  - [installing/managing software on 2nd 3rd](#)
  - [viewing performance data of](#)
- [VMs](#)
- [Service Console NIC 2nd](#)
  - [troubleshooting](#)
- [Service Console swap file](#)
- [service network restart command 2nd](#)
- [service servicename restart command](#)
- [services](#)
  - [restarting](#)
- [shared mode](#)
- [source servers](#)
- [SQL Server \(Microsoft\) 2nd 3rd 4 h 5 h 6 h 7 h 8 h 9 h 10th](#)
- [SSH \(Secure Shell\)](#)
- [SSH client](#)
- [start-up process](#)
- [Startup and Shutdown settings 2nd 3rd 4th](#)
- [storage virtualization](#)
- [su command](#)
- [subnet mask \(Service Console\)](#)
- [support services](#)
- [swap space](#)
- [swing servers](#)
- [symbolic links](#)
- [syslogd daemon](#)
- [SysPrep 2nd 3rd](#)
- [system architecture 2nd 3rd 4th 5th](#)
- [system requirements 2nd 3rd](#)
  - [for VirtualCenter 2nd](#)





# Index

[SYMBOL] [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[tar command](#) [2nd](#)  
[target servers](#)  
[Ta ham](#)  
[technical requirements](#) [2nd](#) [3rd](#) [4th](#) [5th](#)  
[technology design](#) [2nd](#) [3rd](#) [4th](#)  
[technology refresh](#) [2nd](#)  
[telnet client](#)  
[templates](#)  
    [creating](#)  
    [VMs](#) [2nd](#) [3rd](#)  
[terminology](#)  
    [VirtualCenter and](#)  
[testdir directory](#)  
[testing](#)  
[text editors](#) [2nd](#) [3rd](#) [4th](#) [5th](#)  
[time](#)  
[/tmp directory](#)  
[top command](#)  
[trending](#)  
[troubleshooting](#)  
    [directory structure and](#) [2nd](#)  
    [server consolidation challenges and](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[Undoable Mode](#)

[utilities](#)

[ISO images](#) [2nd](#)

[NewSID](#) [2nd](#) [3rd](#) [4 h](#)

[P2V Assistant](#) [2nd](#)

[PuTTY](#) [2nd](#) [3rd](#) [4th](#)

[vmkfstools](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#)

[vmware-cmd](#) [2nd](#) [3rd](#) [4 h](#) [5 h](#)

[WinISO](#) [2nd](#)

[WinSCP](#) [2nd](#)

[utility computing](#)



# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[/var directory](#)

[vdf -h command](#) [2nd](#)

[VGT \(Virtual Guest Tagging\)](#) [2nd](#) [3rd](#)

[vi editor](#)

[virtual disk files](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#)

[virtual disks](#)

[backing up](#) [2nd](#)

[extending size of](#)

[options for](#)

[selecting existing](#)

[shrinking](#)

[Virtual Guest Tagging \(VGT\)](#) [2nd](#) [3rd](#)

[virtual infrastructure](#) [2nd](#) [3rd](#) [4 h](#) [5 h](#) [6 h](#) [7 h](#) [8 h](#) [9 h](#) [10th](#) [11th](#) [12th](#) [13 h](#) [14th](#) [15th](#) [16th](#) [17 h](#) [18th](#) [19th](#) [20th](#) [21st](#) [22nd](#) [23rd](#) [24th](#) [25th](#) [26th](#) [27 h](#) [28th](#)

[29th](#) [30 h](#) [31st](#) [32nd](#) [33rd](#) [34th](#) [35th](#) [36 h](#) [37th](#) [38th](#) [39th](#) [40 h](#) [41st](#) [42nd](#) [43rd](#) [44th](#) [45th](#) [46 h](#) [47th](#) [48th](#) [49th](#) [50 h](#) [51st](#) [52nd](#) [53rd](#) [54th](#) [55th](#) [56 h](#) [57th](#)

[58th](#) [59 h](#) [60th](#) [61st](#) [62nd](#) [63rd](#) [64th](#) [65 h](#) [66th](#) [67th](#) [68th](#) [69 h](#) [70th](#) [71st](#) [72nd](#) [73rd](#) [74th](#) [75 h](#) [76th](#) [77th](#) [78th](#) [79 h](#) [80th](#) [81st](#)

[tools for](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#) [10th](#) [11th](#) [12 h](#) [13th](#) [14th](#) [15th](#) [16 h](#) [17th](#) [18th](#) [19th](#) [20 h](#) [21st](#) [22nd](#) [23rd](#)

[via VirtualCenter](#) [2nd](#) [3rd](#) [4 h](#) [5 h](#) [6 h](#) [7 h](#) [8 h](#) [9 h](#) [10th](#) [11th](#) [12th](#) [13 h](#) [14th](#)

[virtual local area networks \(VLANs\)](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#)

[virtual machine configuration files](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#)

[editing](#) [2nd](#) [3rd](#)

[finding information in](#)

[moving](#)

[virtual machine groups \(VM groups\)](#)

[managing](#)

[virtual machine monitor \(VMM\)](#) [2nd](#)

[virtual machines \(VMs\)](#) [2nd](#)

[backing up](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#)

[configuring](#) [2nd](#)

[creating](#) [2nd](#) [3rd](#)

[determining processes related to](#)

[synchronizing with ESX Server via VMware Tools](#)

[templates of](#) [2nd](#)

[vendor support for](#)

[virtual networking](#) [2nd](#) [3rd](#) [4 h](#) [5 h](#) [6 h](#) [7 h](#) [8 h](#) [9 h](#) [10th](#) [11th](#) [12th](#) [13 h](#) [14th](#) [15th](#) [16th](#) [17 h](#) [18th](#) [19th](#) [20th](#) [21st](#) [22nd](#) [23rd](#) [24th](#) [25th](#) [26th](#) [27 h](#) [28th](#) [29th](#)

[30th](#)

[virtual NICs](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#)

[Virtual Server \(Microsoft\)](#) [2nd](#)

[Virtual Switch Tagging \(VST\)](#) [2nd](#) [3rd](#)

[virtual switches](#) [2nd](#) [3rd](#) [4 h](#) [5 h](#) [6 h](#) [7 h](#) [8 h](#) [9 h](#) [10th](#) [11th](#)

[adding/removing physical NICs and](#)

[assigning to new VM](#)

[creating](#) [2nd](#)

[virtual terminals](#)

[VirtualCenter](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#) [7th](#) [8th](#) [9th](#) [10th](#) [11th](#) [12 h](#) [13th](#) [14th](#) [15th](#) [16 h](#) [17th](#) [18th](#) [19th](#) [20 h](#) [21st](#) [22nd](#) [23rd](#) [24th](#) [25th](#) [26 h](#) [27th](#) [28th](#) [29th](#) [30 h](#)

[features of](#)

[installing](#)

[logging into](#)

[terminology/components of](#) [2nd](#)

[using](#) [2nd](#) [3rd](#) [4 h](#) [5 h](#) [6 h](#) [7 h](#) [8 h](#)

[VirtualCenter Agent](#)

[VirtualCenter Client](#)

[VirtualCenter database](#)

[creating on SQL Server](#) [2nd](#) [3rd](#) [4 h](#)

[VirtualCenter Server](#) [2nd](#) [3rd](#) [4 h](#)

[virtualization](#) [2nd](#) [3rd](#) [4th](#)

[of hardware](#)

[of servers](#) [2nd](#) [3rd](#)

[virtualization layer](#)

[viance driver](#) [2nd](#)

[VLANs \(virtual local area networks\)](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[VM alarms](#) [2nd](#) [3rd](#)

[VM console](#)

[.vmdk files](#) [2nd](#) [3rd](#) [4th](#) [5th](#)

[vmfs partitions](#)

[VMFS volumes](#) [2nd](#) [3rd](#)

[backup best practices and](#)

[viewing contents of](#)

[VMID number](#) [2nd](#)

[/vmimages directory](#)

[VMkernel](#) [2nd](#) [3rd](#)

[vmkfstools](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#)

[vmkfstools command](#) [2nd](#)

[creating/extending virtual disks and](#)

[importing/exporting virtual disks and](#)

[vmkload \\_mod command](#) [2nd](#)

[vmkpcidiv command](#) [2nd](#) [3rd](#)

[VMlibrary](#) [2nd](#)

[VMnets](#)

[VMotion](#) [2nd](#) [3rd](#)



# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

web sites

[BartPE](#)

[VMware 2nd](#)

[Xen](#)

Windows 2003 Server

[Enterprise Edition 2nd](#)

[Standard Edition 2nd](#)

[WinISO 2nd 3rd](#)

[WinSCP 2nd 3rd](#)

[world ids](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[x86 platforms](#)

[Xen 2nd](#)

[xinetd daemon](#)



# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[zip files](#)