

Virtualization

From the Desktop to the Enterprise

CHRIS WOLF AND ERICK M. HALTER

Virtualization: From the Desktop to the Enterprise
Copyright © 2005 by Chris Wolf and Erick M. Halter

All rights reserved. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage or retrieval system, without the prior written permission of the copyright owner and the publisher.

ISBN: 1-59059-495-9

Printed and bound in the United States of America 9 8 7 6 5 4 3 2 1

Trademarked names may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, we use the names only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

Lead Editor: Jim Sumser

Technical Reviewer: Harley Stagner

Editorial Board: Steve Anglin, Dan Appleman, Ewan Buckingham, Gary Cornell, Tony Davis,

Jason Gilmore, Jonathan Hassell, Matthew Moodie, Chris Mills, Dominic Shakeshaft, Jim Sumser

Assistant Publisher: Grace Wong

Project Manager: Kylie Johnston

Copy Manager: Nicole LeClerc

Copy Editor: Kim Wimpsett

Production Manager: Kari Brooks-Copony

Production Editor: Kelly Winkquist

Compositor: Van Winkle Design Group

Proofreader: April Eddy

Indexer: Carol Burbo

Artist: Diana Van Winkle, Van Winkle Design Group

Interior Designer: Diana Van Winkle, Van Winkle Design Group

Cover Designer: Kurt Krames

Manufacturing Manager: Tom Debolski

Distributed to the book trade in the United States by Springer-Verlag New York, Inc., 233 Spring Street, 6th Floor, New York, NY 10013, and outside the United States by Springer-Verlag GmbH & Co. KG, Tiergartenstr. 17, 69112 Heidelberg, Germany.

In the United States: phone 1-800-SPRINGER, fax 201-348-4505, e-mail orders@springer-ny.com, or visit <http://www.springer-ny.com>. Outside the United States: fax +49 6221 345229, e-mail orders@springer.de, or visit <http://www.springer.de>.

For information on translations, please contact Apress directly at 2560 Ninth Street, Suite 219, Berkeley, CA 94710. Phone 510-549-5930, fax 510-549-5939, e-mail info@apress.com, or visit <http://www.apress.com>.

The information in this book is distributed on an “as is” basis, without warranty. Although every precaution has been taken in the preparation of this work, neither the author(s) nor Apress shall have any liability to any person or entity with respect to any loss or damage caused or alleged to be caused directly or indirectly by the information contained in this work.

This book is dedicated to my wonderful wife, Melissa, and son, Andrew.

*True success is not measured by professional accomplishments
but rather by the love and respect of one's family.*

*As George Moore says, "A man travels the world over
in search of what he needs and returns home to find it."*

—Chris Wolf

*This book is dedicated to my family: Holly, Zack, Ella, and Gates,
and to the teachers who taught me to write and think...and Elvis too!*

—Erick M. Halter

Contents at a Glance

About the Authors	xvii
About the Technical Reviewer	xix
Acknowledgments	xxi
Introduction	xxiii
CHAPTER 1 Examining the Anatomy of a Virtual Machine	1
CHAPTER 2 Preparing a Virtual Machine Host.	35
CHAPTER 3 Installing VM Applications on Desktops	69
CHAPTER 4 Deploying and Managing VMs on the Desktop	85
CHAPTER 5 Installing and Deploying VMs on Enterprise Servers.	139
CHAPTER 6 Deploying and Managing Production VMs on Enterprise Servers.	173
CHAPTER 7 Backing Up and Recovering Virtual Machines.	223
CHAPTER 8 Using Virtual File Systems	263
CHAPTER 9 Implementing Failover Clusters	309
CHAPTER 10 Creating Load-Balanced Clusters.	347
CHAPTER 11 Building Virtual Machine Clusters	379
CHAPTER 12 Introducing Storage Networking	413
CHAPTER 13 Virtualizing Storage.	435
CHAPTER 14 Putting It All Together: The Virtualized Information System	471
APPENDIX A Virtualization Product Roundup	495
INDEX	527

Contents

About the Authors	xvii
About the Technical Reviewer	xix
Acknowledgments	xxi
Introduction	xxiii
CHAPTER 1 Examining the Anatomy of a Virtual Machine	1
Introducing VM Types	2
Hardware Emulators	2
Application Virtual Machines	3
Mainframe Virtual Machine	3
Operating System Virtual Machines	3
Parallel Virtual Machines	3
Deploying VMs	5
Choosing VM Hardware	5
Introducing Computer Components	8
CPU	9
RAM	9
Hard Drive	10
Introducing Virtual Disk Types: Microsoft and VMware	11
Virtual Hard Disk and Virtual Disks	12
Dynamically Expanding and Dynamic Disks	12
Fixed and Preallocated Disks	13
Linked and Physical Disks	13
Undo and Undoable Disks	13
Differencing Disks	14
Persistent and Nonpersistent Independent Disks	14
Append Disks	15
Resizing Disks	15
Introducing Networking	16
VM Networking Protocols	17
TCP/IP	17
DHCP	17
NAT	18
Introducing Networking VMs	20

Introducing Hardware	21
Network Interface Card	22
Switches	22
BIOS	24
Generic SCSI	25
I/O Devices	25
Introducing VM Products	26
Virtual PC	26
VMware Workstation	28
Microsoft Virtual Server 2005	29
VMware GSX Server	31
VMware ESX Server	32
Virtual Infrastructure	33
VMware VirtualCenter and VMotion	33
VMware P2V Assistant	33
Migrating Between VMs	34
VMware ACE	34
Summary	34

■ CHAPTER 2 **Preparing a Virtual Machine Host**

Implementing Best Practices	35
Evaluating Host Requirements	36
Selecting a Motherboard	38
CPU Speed and Quantity	40
Controller Chipset	40
Memory Requirements	41
Bus Types	42
Integrated Devices	43
Board Form Factor	44
Overall Quality	44
Considering Your Network	45
Public or Private VMs	45
Availability and Performance	47
Simplicity	48
Mesh Networks	48
Teaming and Load Balancing	49
Network Adapter Teaming	49
VM Networking Configurations	51
Supporting Generic SCSI	53
Windows Guests	53

	Linux Guests	55
	Considering Storage Options	56
	Physical Hard Drive Specifications	57
	RAID	59
	Host Disk Sizing	61
	Guest Disk Sizing	62
	Storage Area Networks	63
	Summary	67
CHAPTER 3	Installing VM Applications on Desktops	69
	Deploying VMs with Microsoft Virtual PC	69
	Installing VMware Workstation for Windows	71
	Installing VMware Workstation for Linux	74
	Installing the RPM	75
	Installing the TAR	76
	VM Host Tuning Tips	78
	Summary	84
CHAPTER 4	Deploying and Managing VMs on the Desktop	85
	Deploying VMs with VMware Workstation	85
	Installing VM Tools	92
	VMware Tools for Windows	93
	VMware Tools for Linux	95
	VMware Virtual Hardware Options for Windows and Linux	95
	Microsoft Virtual PC: Building a Windows VM	101
	Microsoft Virtual PC: Building a Linux VM	105
	Virtual PC Virtual Hardware Options	106
	Installing Virtual Machine Additions	108
	Managing VMs	109
	Backing Up and Modifying VM Configurations	109
	VMware *.vmx Configuration Files	115
	Virtual PC *.vmc Configuration Files	117
	Copying and Moving VMware Workstation Guest VMs	119
	VMware Universally Unique Identifiers	121
	Copying and Moving Virtual PC VMs to Other Hosts	121
	Running VMs As Services	123
	Introducing VM CLI Administration and Keyboard Shortcuts	128
	VMware Workstation CLI	129
	Virtual PC CLI Administration	131

Monitoring and Configuring VM Performance	134
VMware Performance Counters	135
Virtual PC Performance Options	137
Summary	138
CHAPTER 5 Installing and Deploying VMs on Enterprise Servers	139
Installing Microsoft Virtual Server	140
Installing VMware GSX Server for Windows	143
Installing VMware GSX Server for Linux	146
Installing the RPM	147
Installing the TAR	149
Installing the VMware Management Interface	150
Working with the VMware Virtual Machine Console	152
Changing GSX Server's Remote Console Port Number	153
Installing VMware ESX Server	154
Verifying ESX Server Configuration Information	161
Viewing Configuration Files	161
Using Linux Survival Commands	163
Working with the Management Interface	165
Understanding MUI and SSL	165
Configuring the ESX Server Installation: Part One	166
License Agreement	167
Startup Profile	167
Storage Configuration	167
Swap File	168
Network Configuration	168
ESX Security	169
Configuring the ESX Server Installation: Part Two	170
Summary	171
CHAPTER 6 Deploying and Managing Production VMs on Enterprise Servers	173
Deploying VMs with VMware GSX Server and ESX Server	173
Building VMware GSX Server VMs	173
Building VMware ESX Server VMs	176
Mounting ISO Images	177
Installing VM Tools for GSX Server and ESX Server VMs	178
Using VMware Tools for Windows	178
Using VMware Tools for Linux	179

Configuring VMware GSX Server and ESX Server Virtual Hardware Options	180
Hard Disk	181
DVD/CD-ROM Drive	182
Floppy Drive	182
Ethernet Adapter	183
Sound Adapter	183
Configuring Legacy Devices	183
Configuring Generic SCSI Devices	184
Configuring a USB Controller	185
Scripting ESX Server USB Connectivity	191
Building Microsoft Virtual Server VMs	193
General Properties	194
Virtual Machine Additions	194
Memory	195
Hard Disks	195
CD/DVD	195
SCSI Adapters	196
Network Adapters	196
Scripts	196
Floppy Drives	196
COM Ports	197
LPT Ports	197
Managing Server-Class VMs	197
Modifying VM Configurations: Renaming and Moving	197
Using VMware Universally Unique Identifiers (UUIDs)	200
Importing Workstation and GSX Server VMs into ESX Server	202
Working with VMware GSX Server and ESX Server *.vmx Configuration Files	207
Working with Virtual Server *.vmc Configuration Files	210
Performing Command-Line Management	211
VMware	211
Microsoft	213
Using the Windows System Preparation Tool	214
Monitoring VM Performance	215
Monitoring ESX Server Performance	215
Monitoring VMware GSX Server Performance	219
Monitoring Virtual Server Performance	220
Performing Fault Monitoring and Fault Tolerance	221
Summary	222

CHAPTER 7	Backing Up and Recovering Virtual Machines	223
	Performing Traditional Agent-Based Backups	224
	Running Backup Agents on VMs	224
	Running Backup Agents on the Host	230
	Performing Non-Agent-Based Backups	230
	Using Windows Backup	231
	Backing Up Linux File Systems	234
	Performing Flat-File Backups	239
	Running VMware Workstation Flat-File Backups	241
	Running VMware GSX Server Flat-File Backups	247
	Running Virtual PC 2004 Flat-File Backups	251
	Running Virtual Server 2005 Flat-File Backups	252
	Taking Online Snapshots	257
	Performing a Full System Recovery	259
	Restoring Online VM Backups	260
	Restoring Flat-File VM Backups	261
	Summary	261
CHAPTER 8	Using Virtual File Systems	263
	Introducing DFS	263
	Implementing Windows DFS	265
	Implementing Linux DFS	271
	Using Samba with Kerberos Authentication	284
	Adding Samba to Active Directory	285
	Setting Up Samba DFS Shares	286
	Introducing AFS	288
	Implementing AFS	292
	Installing AFS for Linux Systems	302
	Summary	307
CHAPTER 9	Implementing Failover Clusters	309
	Introducing Failover Clustering	310
	Defining Essential Terms	311
	Introducing Cluster Architecture	312
	Introducing N-tier Clustering	313
	Working with Failover Cluster Products	314
	Planning for Failover Clusters	315
	Choosing the Right Model	316
	Configuring Cluster Hardware	319

Setting Up Microsoft Server Clusters	320
Looking Under the Hood	320
Planning Resource and Group Configuration	323
Installing the Windows Server 2003 Cluster Service	326
Using the Cluster Administrator	328
Setting Up Linux Failover Clusters	331
Setting Up the Red Hat Cluster Suite	331
Using Linux-HA Clusters	333
Summary	345
CHAPTER 10 Creating Load-Balanced Clusters	347
Round-Robin DNS: The Beginning	349
Planning for Load-Balanced Clusters	351
Selecting Applications	351
Verifying Licensing	351
Analyzing Risks	352
Estimating Server Capacity	352
Building Windows Network Load-Balanced (NLB) Clusters	353
Enabling the NLB Service	353
Understanding Unicast and Multicast	354
Understanding Convergence	356
Setting Priority	357
Setting Port Rules	358
Understanding Remote Control	358
Using the Network Load Balancing Manager	358
Implementing Best Practices for NLB Cluster Implementations	359
Configuring and Managing Windows NLB Clusters	360
Building Linux Virtual Server (LVS) Clusters	374
Understanding LVS Architecture	375
Implementing LVS Implementation	377
Summary	377
CHAPTER 11 Building Virtual Machine Clusters	379
Building Microsoft VM Clusters	379
Setting Up Windows Server Clusters	380
Setting Up iSCSI Windows Server Clusters	391
Installing the Windows Server 2003 Cluster Service	404
Setting Up Windows NLB Clusters	406
Building Linux VM Clusters	409
Summary	411

CHAPTER 12	Introducing Storage Networking	413
	Introducing SCSI	414
	Speaking SCSI	414
	ID vs. LUN	415
	Using SCSI Buses	415
	Understanding Termination	418
	Introducing Fibre Channel	421
	Introducing Fibre Channel Cables	423
	Introducing Fibre Channel Hardware Devices	424
	Understanding Zoning	428
	Configuring Fibre Channel Hardware	429
	Extending the SAN with FCIP and iFCP	430
	Introducing iSCSI	430
	Understanding iSCSI Architecture	430
	Securing iSCSI	431
	Using SAN Backup and Recovery Techniques	432
	Performing LAN-Free Backups	432
	Performing Server-Free Backups	433
	Performing Serverless Backups	434
	Summary	434
CHAPTER 13	Virtualizing Storage	435
	RAID: The Root of Storage Virtualization	435
	Introducing Common RAID Levels	435
	Implementing RAID	442
	Introducing the SNIA Shared Storage Model	443
	Why a Model for Shared Storage?	444
	Benefits of the Model	445
	A Note on the Graphical Conventions Used in the Model	445
	The Classic Storage Model	446
	The SNIA Shared Storage Model	447
	Applying the SNIA Shared Storage Model	461
	Understanding Host-Based Architecture	461
	Understanding Storage-Based Architecture	463
	Understanding Network-Based Architecture	463
	Adding Fault Tolerance to the SAN	466
	Performing Backups	466
	Introducing Hierarchical Storage Management	466

Using Virtual Tape Libraries	467
Dividing Physical Libraries	468
Writing to Magnetic Disk	468
Summary	469
CHAPTER 14 Putting It All Together: The Virtualized Information System	471
Reviewing the Elements of the Virtual IS	471
Failover Cluster	472
Load-Balanced Cluster	473
Virtual Machine Host	474
Storage Area Network	476
Distributed File System	479
Maintaining a Standby VM Server	480
Setting Up the VM	482
Maintaining a Standby Server with Scheduled Backups and Restores	483
Maintaining a Standby Server with Shared Storage	483
Maintaining a Standby Server with Disk Mirroring	485
Automating Standby Server Startup	489
Summary	494
APPENDIX A Virtualization Product Roundup	495
Global Namespace:	
The New Paradigm in Distributed Data Management	495
The Problem: Network Data Management and Movement	496
The Solution: Global Namespace	497
StorageX Uses Global Namespace to Deliver a Complete Network Data Management Platform	502
What Is Unique About the StorageX Global Namespace?	502
Conclusion	504
Server Consolidation and Beyond:	
Enhancing Virtual Machine Infrastructure Through Automation	505
Evolution of Virtualization in the Data Center	506
Enhancing Data Center Flexibility Through PlateSpin PowerP2V™	506
Comparing Other Methods of Converting Between Physical and Virtual Machine Infrastructure with PlateSpin PowerP2V	507

Replicating Entire Test Lab Environments Using Virtual Machines . . .	510
Using Virtual Machines As Hot Backup Servers	
for Planned and Unplanned Downtime	511
Moving a Virtual Machine from One VM Host to Another (V2V)	512
Production Server Virtualization.	513
Server Migrations Across Geographic Regions.	514
The Need for Continuous Resource Analysis and Rebalancing . . .	514
Dynamic Virtual Machine Portability:	
Using Virtual Machines to Prevent SLA Violations	515
How Dynamic Virtual Machine Portability	
Can Enhance Business Service Management	516
Conclusion.	519
Rocket Division Software	519
iSCSI Target.	519
StarPort iSCSI Initiator, RAM Disk, and Virtual DVD Emulator. . . .	521
Mission and Technologies.	522
Market Focus	523
Network Instruments' Observer	523
Too Much Traffic in the Kitchen	523
Two Unique Networks Working Together	523
Maximizing the Network	524
Complete Network Control	524
Capacity Analysis	524
Proactive Management	525
About Jack in the Box, Inc.	525
About Real-Time Expert	525
About Advanced Single Probes, Advanced Multi-Probes, and Advanced Expert Probes	525
About Network Trending	526
About "What-If" Analysis	526
About Network Instruments, LLC.	526
INDEX	527

About the Authors



■ **CHRIS WOLF** has worked in the IT trenches for more than a decade, specializing in virtualization, enterprise storage, and network infrastructure planning and troubleshooting. He has written four books and frequently contributes to *Redmond* magazine and *Windows IT Pro* magazine. Chris has a master's of science degree in information technology from the Rochester Institute of Technology and a bachelor's of science degree in information systems from the State University of New York–Empire State College.

Currently, Chris is a full-time member of the faculty at ECPI Technical College in Richmond, Virginia. When not teaching, Chris stays very involved in consulting projects for midsize to enterprise-class organizations and is a regular speaker at computer conferences across the nation. Chris is a two-time Microsoft MVP award recipient and currently has the following IT certifications: MCSE, MCT, CCNA, Network+, and A+.



■ **ERICK M. HALTER**, an award-winning IT/networking and security management educator for more than three years, is now the senior security administrator for a technology-based law firm where he's virtualizing the network and optimizing system processes for the Web. Erick has earned several industry certifications, including CCNP, MCSE: Security, Master CIW Administrator, SCNP, Security+, Linux+, and Net+, and he's an industry consultant. Erick completed his undergraduate studies in English and is currently earning a graduate degree in IT. He has more than a decade of practical experience in troubleshooting Microsoft, Cisco, and Linux networking technologies. He resides in Richmond, Virginia, with his wife and three dogs.

About the Technical Reviewer

■ **HARLEY STAGNER** has been an IT professional for seven years. He has a wide range of knowledge in many areas of the IT field, including network design and administration, scripting, and troubleshooting. He currently is the IT systems specialist for WWBT Channel 12, a local NBC affiliate television station in Richmond, Virginia. He is a lifelong learner and has a particular interest in storage networking and virtualization technology.

Harley has a bachelor's of science degree in management information systems from ECPI Technical College in Richmond, Virginia. He also has the following IT certifications: MCSE, CCNA, Network+, and A+.

Acknowledgments

Writing a book of this magnitude has certainly been a monumental task, and to that end I owe thanks to many. First, I'd like to thank you. Without a reading audience, this book wouldn't exist. Thank you for your support of this and of other books I've written to date.

Next, I'd like to thank my wonderful wife, Melissa. Melissa has been by my side throughout many of my writing adventures and is always willing to make the extra cup of coffee or do whatever it takes to lend support. I must also thank my mother, Diana Wolf, and my father, the late William Wolf. Thank you for encouraging me to always chase my dreams.

At this time, I must also thank my coauthor, Erick Halter. My vision for this book may not have been realized if not for Erick's hard work and persistence.

Several contributors at Apress were also extremely dedicated to this book's success. First, I must thank my editor, Jim Sumser, who shared in my vision for this book. Next, I must thank Kylie Johnston, the project manager. After having worked with Kylie before, I already knew that I'd be working with one of the industry's best. However, I also realize how easy it is to take everything Kylie does for granted. Kylie, you're a true professional and a gem in the mine of IT book publishing. Next, I must thank Kim Wimpsett, whose keen eye for detail truly made this book enjoyable to read. A technical book's true worth is most measured by its accuracy. With that in mind, I must also thank our technical editor, Harley Stagner, who diligently tested every procedure and script presented in this book.

I must also thank my agent, Laura Lewin, and the great team of professionals at Studio B. For one's ideas and vision to have meaning, they must be heard. Studio B, you're my virtual megaphone.

Finally, I must thank several technical associates who also added to the content of this book with their own tips and war stories. Please appreciate the contributions and efforts of the following IT warriors: Mike Dahlmeier, Jonathan Cragle, David Conrad, Scott Adams, Jim Knight, John Willard, Iantha Finley, Dan Vasconcellos, Harvey Lubar, Keith Hennett, Walt Merchant, Joe Blow, Matt Keadle, Jimmy Brooks, Altaf Virani, and Rene Fourhman.

—Chris Wolf

I aggravated and neglected a lot of people during this writing project: thank you for being patient and not giving up on me. Moreover, I am indebted to Chris, the crew at Apress, and the folks at Studio B for providing me with this opportunity to write about virtualization. Thank you. Thank you. Thank you.

—Erick M. Halter

Introduction

Virtualization is a concept that has evolved from what many first recognized as a niche technology to one that's driving many mainstream networks. Evidence of virtualization exists in nearly all aspects of information technology today. You can see virtualization in sales, education, testing, and demonstration labs, and you can see it even driving network servers.

What's virtualization? Well, to keep it simple, consider *virtualization* to be the act of abstracting the physical boundaries of a technology. Physical abstraction is now occurring in several ways, with many of these methods illustrated in Figure 1. For example, workstations and servers no longer need dedicated physical hardware such as a CPU or motherboard in order to run as independent entities. Instead, they can run inside a virtual machine (VM). In running as a virtual machine, a computer's hardware is emulated and presented to an operating system as if the hardware truly existed. With this technology, you have the ability to remove the traditional dependence that all operating systems had with hardware. In being able to emulate hardware, a virtual machine can essentially run on any x86-class host system, regardless of hardware makeup. Furthermore, you can run multiple VMs running different operating systems on the same system at the same time!




Figure 1. A virtualized information system

Virtualization extends beyond the virtual machine to other virtualization technologies such as clustering. *Clustering* allows several physical machines to collectively host one or more virtual servers. Clusters generally provide two distinct roles, which are to provide for continuous data access, even if a failure with a system or network device occurs, and to load balance a high volume of clients across several physical hosts. With clustering, clients won't connect to a physical computer but instead connect to a logical virtual server running on top of one or more physical computers. Clustering differs from virtual machine applications in that it allows for automated failover between physical hosts participating in the cluster. You could view failover as the movement of a virtual server from one physical host to another.

Aside from virtual machines and clustering, you'll also see the reach of virtualization extend to network file systems and storage. In reaching network file systems, technologies such as Distributed File System (DFS) allow users to access network resources without knowing their exact physical location. With storage virtualization, administrations can perform restores of backed-up data without having to know the location of the physical media where the backup resides.

Now, if your head is already spinning, don't worry, because you're probably not alone. With such a vast array of virtualization technologies available, it can be difficult to first tell one from another and also come to an understanding as to which technologies are right for you. That's why we decided to piece together a reference that explains each available virtualization technology, whether you're interested in running virtual machines on your desktop or are planning to add virtualization layers to an enterprise network environment. In this book, we'll guide you through all aspects of virtualization and also discuss how to fit any and all of these complex technologies into your IT life. Let's start by looking at the format of each chapter in this book.

Chapter 1: Examining the Anatomy of a Virtual Machine

Two major software vendors, EMC (Legato) and Microsoft, are leading the virtual machine software charge. In spite of their products' architectural differences, the terminology and the theory that drives them are similar.

In Chapter 1, we'll start by explaining the buzzwords and the theory associated with virtual machines. We'll address such topics as virtual networks, virtual hard disks, and CPU emulation. We'll also give you an overview of the major virtual machine products, outlining the differences between EMC's VMware Workstation, GSX Server, and ESX Server products, as well as Microsoft's Virtual PC and Virtual Server 2005.

Chapter 2: Preparing a Virtual Machine Host

With an understanding of the ins and outs of virtual machines, the next logical step in VM deployment is to prepare a host system. Several factors determine a host's preparedness to run a VM application, including the following:

- Physical RAM
- CPU
- Hard disk space
- Networking

When selecting a host, you'll need to ensure that the host meets the VM application's minimum hardware requirements and also that enough resources are available for the number of VMs you plan to run simultaneously on the host. Properly preparing a host prior to running virtual machines on it will almost certainly result in better stability, scalability, and long-term performance for your virtual machines. After finishing Chapter 2, you will not only be fully aware of the techniques for preparing a host system for virtual machines but will also be aware of the many gotchas and common pitfalls that often go unrecognized until it's unfortunately too late.

Chapter 3: Installing VM Applications on Desktops

When deciding to run virtual machines on workstations, you have two choices of workstation application: VMware Workstation (shown in Figure 2) and Virtual PC (shown in Figure 3). VMware Workstation is supported on Windows NT 4.0 (with SP6a) or higher Windows operating systems and can also run on Linux (Red Hat, Mandrake, or SuSE). Microsoft Virtual PC is supported on Windows 2000 Professional, Windows XP Professional, and Windows XP Tablet PC Edition.

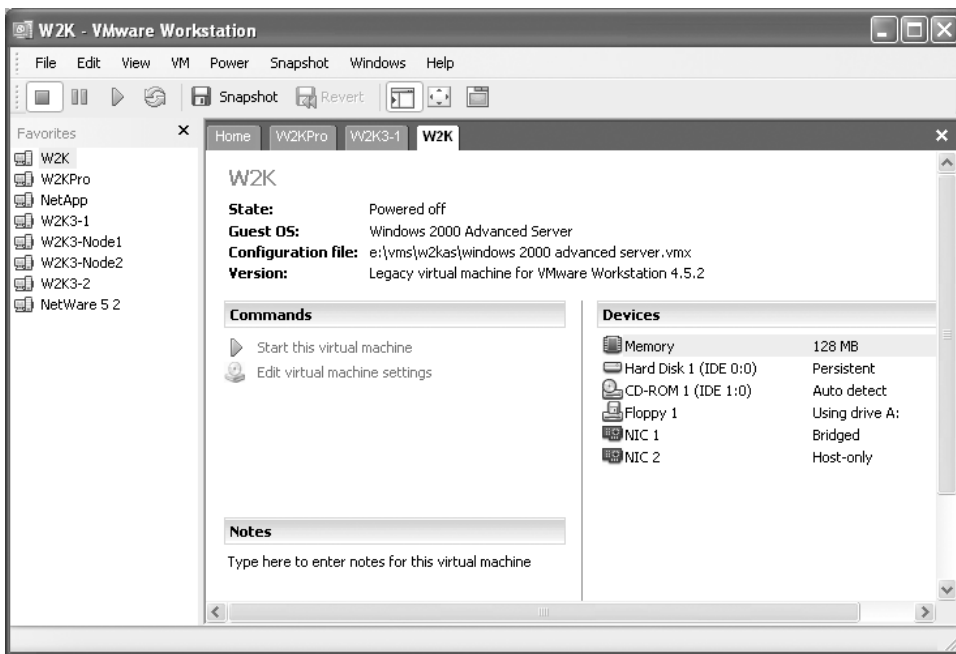


Figure 2. VMware Workstation UI

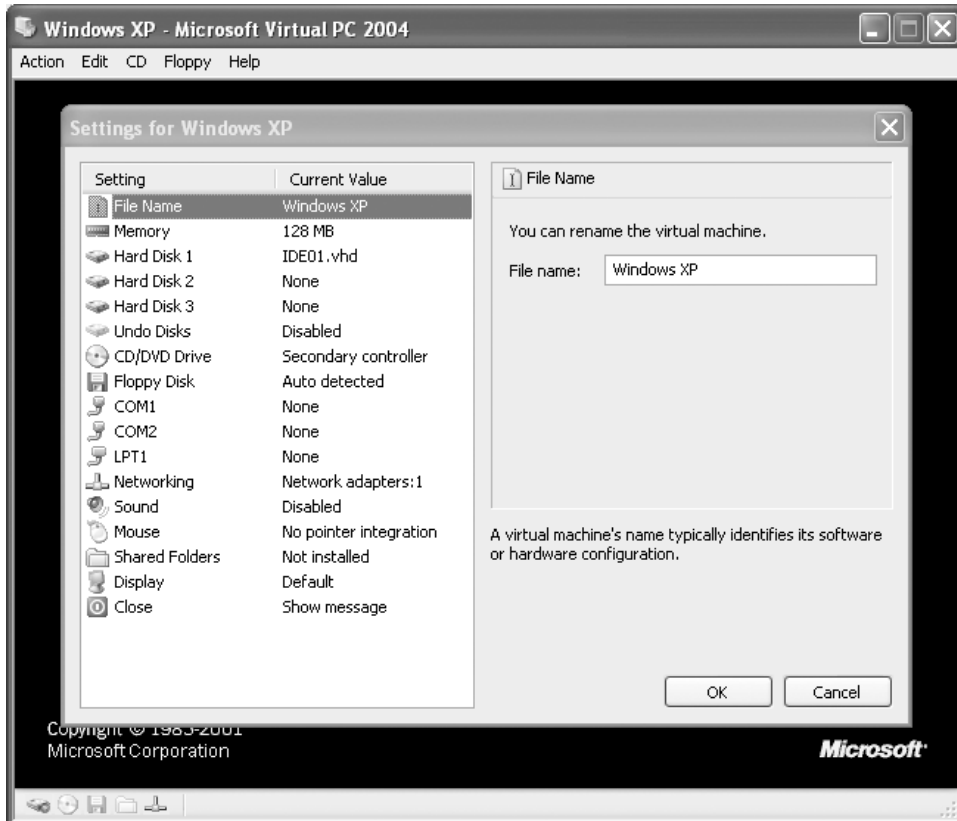


Figure 3. *Virtual PC UI*

As you can see, your current operating system may decide the VM application you choose. Chapter 1 will also help with the VM application decision, as it will outline all of the differences between each program. Once you've decided on a VM application, you'll be able to use this chapter to get you through any final preparations and the installation of VMware Workstation and Virtual PC.

Chapter 4: Deploying and Managing VMs on the Desktop

Chapter 4 provides guidance on deploying specific VM operating systems on your workstation system, including examples of both Windows and Linux VM deployments. Once your VM is up and running, you can perform many tasks to optimize and monitor their performance and to also ease future VM deployments. Topics in this chapter include the following:

- Monitoring the performance of VMs
- Staging and deploying preconfigured VMs

- Running VMs as services
- Configuring VMs to not save any information
- Administering through the command line and scripts

As you can see, this chapter is loaded with information on VM management. This is the result of years of experience, so you'll find the tips and techniques presented in this chapter to be as valuable as a microwave oven. Although they may not heat your food any faster, they definitely will save you plenty of time managing your virtualized network.

Simply installing and running virtual machines is really just the tip of the iceberg. Once your VMs are configured, you can perform many tasks to make them run better and run in ways you never thought possible. In Chapter 4, you'll see all of this and more.

Chapter 5: Installing and Deploying VMs on Enterprise Servers

Chapters 5 and 6 are similar in format to Chapters 3 and 4. In Chapter 5, we'll walk you through the installation and deployment process of VM applications on server systems, with the focus on using the VMs in a production role.

As with Chapter 3, in this chapter we devote time to both VMware GSX Server deployments on Linux and Windows operating systems and Microsoft Virtual Server 2005 deployments on Windows operating systems. With the decision to run VMs in production comes a new list of responsibilities. Whether you're looking to run domain controllers, file servers, or even database servers as VMs, you must consider several performance factors for each scenario before setting up the VMs. Many of us have learned the sizing game the hard way, and we don't want you to have to suffer as well. An undersized server can sometimes be the kiss of death for an administrator, as it may be difficult to get additional funds to "upgrade" a server that's less than two months old, for example. Sizing a server right the first time will not only make VM deployment easier but it may also help to win over a few of the virtualization naysayers in your organization. When deploying cutting-edge technology, you'll always have pressure to get it to run right the first time, so pay close attention to the advice offered in this chapter. We've made plenty of the common deployment mistakes already, so you shouldn't have to do the same!

Chapter 6: Deploying and Managing Production VMs on Enterprise Servers

Running VMs in production may involve managing fewer physical resources, but the fact that multiple OSs may depend on a common set of physical hardware can cause other problems. For example, rebooting a VM host server may affect four servers (a host plus three VMs) instead of one. Problems such as this put more pressure on you as an administrator to consider how an action on one server can impact several other servers.

With VMs in production, keeping the host system happy will likely result in well-running virtual machines. Again, with many systems depending on the hardware of one system, a hung CPU, for example, could have devastating consequences. This is why monitoring and mainte-

nance is still crucial after deployment. To help ease your mind, in this chapter we'll give you guidance on all the most common VM server administrative tasks, including the following:

- Automating the startup and shutdown of virtual machines
- Monitoring VM and host performance and alerting when performance thresholds are passed
- Using Perl and Visual Basic scripts for management, monitoring, and alerting
- Spotting host system bottlenecks and taking the necessary corrective action

As with its predecessors, you'll find Chapter 6 to have valuable information for managing production virtual machines.

Chapter 7: Backing Up and Recovering Virtual Machines

Now that you have your VMs up and running, you can't forget about protecting them. In this chapter, we'll cover all the methods for backing up and recovering virtual machines. We'll discuss the following:

- Backing up VMs with backup agent software
- Backing up VMs as "flat files"
- Backing up the VM host
- Using the available scripted backup solutions

As you can see, you'll have plenty of alternatives when deciding on how to best protect your virtual machines. In this chapter, we'll not only show you each VM protection methodology but we'll also outline the common pitfalls that exist with certain choices.

Chapter 8: Using Virtual File Systems

While the first half of the book is devoted to the design, deployment, and management of virtual machine solutions, the second half deals with the remaining virtualization technologies currently available. Chapter 8 leads off by explaining virtual file systems. With virtual file systems, you can manage files and file shares transparent to their physical location. For users, this means they won't need to know where a file is physically located in order to access it. For administrators, this means that if a file server needs to be brought down for maintenance, you can move its data temporarily to another server so that it remains available to users, without the users ever noticing a difference. Also, it's possible to configure replication with your virtual file system solution, which can allow for both load balancing and fault tolerance of file server data.

To tell the complete virtual file system story, we'll explain the most widely employed solutions available today, including both DFS and Andrew File System (AFS).

Chapter 9: Implementing Failover Clusters

The general concept of clustering is to allow multiple physical computers to act as one or more logical computers. With server or failover clusters, two or more physical servers will host one or more virtual servers (logical computers), with a primary purpose of preventing a single point of failure from interrupting data access. A single point of failure can be any hardware device or even software whose failure would prevent access to critical data or services. With the server cluster, a virtual server will be hosted by a single node in the cluster at a time. If anything on the host node fails, then the virtual server will be moved by the cluster service to another host node. This allows the virtual server running on the cluster to be resilient to failures on either its host system or on the network. Figure 4 shows a typical server cluster.

One aspect of the server or failover cluster that's unique is that all physical computers, or nodes, can share one or more common storage devices. With two nodes, this may be in the form of an external Small Computer System Interface (SCSI) storage array. For clusters larger than two nodes, the shared external storage can connect to the cluster nodes via either a Fibre Channel bus or an iSCSI bus.




Figure 4. Two-node mail server cluster

In Chapter 9, we'll fully explain server clustering, outlining its deployment options and common management issues. To tell the complete clustering story, we'll cover the deployment and management of both Windows and Linux clustering solutions.

Chapter 10: Creating Load-Balanced Clusters

Load-balanced clusters give you the ability to relieve some of the load on an overtaxed server. With load balancing on Microsoft servers, you can configure up to 32 servers to share requests from clients. On Linux, you can even go beyond the 32-node limit imposed by Microsoft server operating systems.

In short, load balancing allows you to configure multiple servers to act as a single logical server for the purpose of sharing a high load of activity imposed by network clients. In a load-balanced cluster, two or more physical computers will act as a single logical computer, as

shown in Figure 5. Client requests are evenly distributed to each node in the cluster. Since all clients attempt to access a single logical (or virtual) server, they aren't aware of the physical aspects of the network server they're accessing. This means a client won't be aware of which physical node it's in communication with. Configuring load-balanced clusters will give you a great deal of flexibility in a network environment that requires a high level of performance and reliability. Because of its natural transparency to clients, you can scale the cluster as its load increases, starting with two nodes and adding nodes to the cluster as the demand requires.




Figure 5. *Four-node load-balanced Web server cluster*

Since load-balanced clusters don't share a common data source (as with server or failover clusters), they're typically used in situations that require fault tolerance and load balancing of read-only data. Without shared storage, writing updates to a load-balanced cluster would be difficult to manage, since each node in the cluster maintains its own local copy of storage. This means it's up to you to make sure the data on each cluster node is completely synchronized. Because of this limitation, load-balanced clusters are most commonly a means to provide better access to Web and FTP services.

In Chapter 10, we'll take you through the complete design and implementation process for load-balanced clusters. We'll show you examples of when to use them and also detail how to deploy load-balanced clusters on both Windows and Linux operating systems.

Chapter 11: Building Virtual Machine Clusters

Although many organizations have or plan to run clusters in production, few have resources to test cluster configurations. That's where building virtual machine clusters can help. With virtual machine clusters, you can build and test working cluster configurations on a single system. Having this ability gives you several advantages, including the following:

- The ability to have others train on nonproduction equipment
- The ability to perform practice restores of production clusters to virtual machines
- The ability to perform live demonstrations of cluster configurations using a single system

Before virtual machines, most administrators had to learn clustering on production systems, if at all. Few organizations had the resources to run clusters in lab environments. When it comes time to test disaster recovery procedures, many organizations can't practice recovery for clusters, again because of limited resources. By being able to run a working cluster inside of virtual machines, organizations now have a means to test their backups of server clusters and in turn prepare recovery procedures for the production cluster. Applications and resources that run on server clusters are often the most critical to an organization. Oftentimes, when a disaster occurs, it's the clusters that must be restored first. Without having the ability to test recovery of the cluster and thus practice for such an event, recovering a cluster in the midst of a crisis can be all the more nerve-racking.

After reading Chapter 11, you'll understand the methodologies needed to configure nearly any cluster configuration using virtual machine technology.

Chapter 12: Introducing Storage Networking

The development of storage virtualization has been fueled by the rapid growth of storage networking. In short, storage networking allows you to network storage resources together for the purpose of sharing them, similar to a TCP/IP network of workstations and servers.

To understand the methodologies and benefits of virtualized storage resources, you must first be comfortable with storage networking technologies. This chapter lays the technical foundation for Chapter 13 by fully dissecting storage networking. Topics covered in Chapter 12 include storage area networks (SANs), network-attached storage (NAS), and direct-attached storage (DAS). Several modern storage networking protocols will be discussed, including the following:

- Fibre Channel Protocol (FCP)
- Internet Fibre Channel Protocol (iFCP)
- Fibre Channel over Internet Protocol (FCIP)
- Internet SCSI (iSCSI)

In addition to covering all the relevant storage networking protocols, we'll also dive into the hardware devices that drive storage networks. The following are some of the most common storage networking hardware devices that will be examined in this chapter:

- Fibre Channel switches
- Fibre Channel bridges and routers
- Fibre Channel host bus adapters (HBAs)
- Gigabit interface converters (GBICs)

Many in IT don't find storage to be the most thrilling topic, but nearly all realize its importance. Because of the inherently dry nature of storage, you may find that this chapter serves dual purposes: it lays the foundation for understanding storage virtualization, and it may substitute as an excellent bedtime story for your children, guaranteed to have them sleeping within minutes!

Chapter 13: Virtualizing Storage

Storage virtualization stays within the general context of virtualization by giving you the ability to view and manage storage resources logically. Logical management of storage is a significant leap from traditional storage management. For many backup administrators, having to restore a file often meant knowing the day the file was backed up and also knowing the exact piece of backup media on which the file was located. With storage virtualization, many backup products now abstract the physical storage resources from the administrator. This allows you as an administrator to simply tell the tool what you want, and it will find the file for you.

With data continuing to grow at a near exponential rate, it can be easy to become overwhelmed by the task of managing and recovering data on a network. As your data grows, so do the number of storage resources you're required to track. Having the right tools to give you a logical view of physical storage is key to surviving storage growth without having to seek mental counseling. Okay, maybe that's a stretch, but you'll certainly see how much easier your life as an administrator can become after you finish reading this chapter.

Chapter 14: Putting It All Together: The Virtualized Information System

Following a theme common to most technical references, Chapters 1–13 cover virtualization technologies one at a time, making it easy for you to find specific information on a particular topic. However, although it's nice to understand and appreciate each technology, it's also crucial to understand their interrelationships. In this chapter, you'll see examples of networks running several combinations of virtualization technologies simultaneously.

Many find relating virtualization technologies to their organization's networks to be challenging. Some common questions we run into quite frequently include the following:

- How can I justify an investment in virtualization to upper management?
- What are the best uses for server-class virtual machine products?

- What are the best uses for workstation-class virtual machine products?
- How can I optimize data backup and recovery between VMs and my production storage area network?
- What situations are best suited for clustering solutions?
- What questions should I ask when sizing up suitable hardware and software vendors?
- What precautions must be observed when integrating different virtualization technologies?

In addition to answering the most common questions surrounding running virtualization technologies in production and test environments, we'll also provide examples of the methods other organizations are using to make best use of their virtualization investments. This chapter wraps up with a detailed look at a process for maintaining standby virtual machines that can be automatically brought online if a production VM fails.

Appendix

Although virtualization product vendors such as EMC and Microsoft will go far to aiding you with support utilities, several other vendors also offer products to add virtualization layers to your existing network and to also aid in managing virtualized network resources.

In this appendix, we'll shower you with examples of some of the latest and greatest virtualization products on the market today. Several of these vendors have allowed us to include evaluation versions of their products on the book's companion CD. With so many virtualization software and hardware vendors contributing to the virtualized information system, you'll probably find managing your virtual resources to be much simpler than you ever imagined.

Summary

Virtualization is no longer an umbrella for disconnected niche technologies but is rather what's seen by many as a necessity for increasingly complex information systems. Virtual machine technology has broad appeal to many different levels of IT professionals. Sales associates can run software demonstrations on virtual machines. Instructors now have a tremendous amount of flexibility when teaching technical classes. They can now demonstrate several operating systems in real time, and their students no longer have to team up with a partner in order to run client-server networking labs. Now students can run client and server operating systems on a single computer in the classroom.

VM technology today has reached several other IT professionals as well. Network architects and administrators can now test software deployments and design solutions on virtual machines, prior to introducing new technologies to a production environment. For our own testing, we've found that our wives are much happier since we no longer each need half a dozen computers in the home office. Now a single computer with several virtual machines works just fine. However, both of us no longer receive a Christmas card from the electric company thanking us for our high level of business each year.

Although the savings on electricity might be a bit of a stretch, the expanded roles of virtualization certainly aren't. When approaching this book, start with Chapter 1 if your first interest is

virtual machines. This will set a solid foundation for Chapters 2–7. Following Chapter 7, you’ll find that the remaining chapters serve as independent references on the various technologies that drive virtualization. Since Chapter 14 focuses on making all virtualization technologies seamlessly operate together in the same information system, you’ll want to read it after Chapters 1–13.

To safely integrate virtualization technologies into your IT life, you first need to have a good understanding of what’s available, when to use each technology, and also how to use them. That being said, let’s not waste any more time discussing how great virtualization technology is. Instead, turn to Chapter 1 to get started with the anatomy of a virtual machine.



Examining the Anatomy of a Virtual Machine

In the preface, you learned how to quickly slice and dice this book to get good, quick results. Now you'll look at what's going on under the hood of virtual machines (VMs). At the component level of VMs and physical hardware, you'll revisit some fairly basic concepts, which you might take for granted, that contribute to the successful virtualization of a physical computer. By taking a fresh look at your knowledge base, you can quickly tie the concept of virtualized hardware to physical hardware.

Wrapping your mind around the term *virtual machine* can be daunting and is often confusing. This stems from the varying definitions of the term *virtual machine*. Even the term *virtual* brings unpredictability to your understanding if you don't understand *virtual* to mean "essence of" or "properties of." In short, if a list of specifications is equal to a machine (think: personal computer), and if software can create the same properties of a machine, you have a VM. If you further reduce a computer to its vital component, electricity, it's easier to understand a VM: the entirety of a computer deals with "flipping" electricity on or off and storing an electrical charge. If software exists such that it can provide the same functionality, a hardware emulator, then a VM exists.

Hardware virtualization offers several benefits, including consolidation of the infrastructure, ease of replication and relocation, normalization of systems, and isolation of resources. In short, VMs give you the ability to run multiple virtual computers on the same physical computer at the same time and store them on almost any media type. VMs are just another computer file offering the same ease of use and portability you've grown accustomed to in a drag-and-drop environment.

Virtualization software protects and partitions the host's resources, central processing units (CPUs), memory, disks, and peripherals by creating a virtualization layer within the host's operating system (OS) or directly on the hardware. "Running on the metal" refers to virtualization software that runs directly on the hardware: no host operating system is required to run the software. Every virtual machine can run its own set of applications on its own operating system. The partitioning process prevents data leaks and keeps virtual machines isolated from each other. Like physical computers, virtual machines require a physical network, a virtual network, or a combination of both network types to communicate.