

"The perfect way to understand both the technology and benefits of virtualization" Simon Crosby, CTO, XenSource

Virtualization

FOR DUMMIES®

**A Reference
for the
Rest of Us!®**

FREE eTips at dummies.com®

Bernard Golden, MBA

CEO of Navica, Inc. & expert blogger
for CIO.com

**Boost server
utilization to
maximize your
IT ROI**



Todos los derechos reservados por Direitos de Autor

Virtualization FOR **DUMMIES®**

by Bernard Golden



Wiley Publishing, Inc.

Virtualization For Dummies®

Published by
Wiley Publishing, Inc.
111 River Street
Hoboken, NJ 07030-5774

www.wiley.com

Copyright © 2008 by Wiley Publishing, Inc., Indianapolis, Indiana

Published by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Legal Department, Wiley Publishing, Inc., 10475 Crosspoint Blvd., Indianapolis, IN 46256, (317) 572-3447, fax (317) 572-4355, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, the Wiley Publishing logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies Daily, The Fun and Easy Way, Dummies.com, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, please contact our Customer Care Department within the U.S. at 800-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002.

For technical support, please visit www.wiley.com/techsupport.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Control Number: 2007940109

ISBN: 978-0-470-14831-0

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1



Contents at a Glance

<i>Foreword.....</i>	<i>xvii</i>
<i>Introduction</i>	<i>1</i>
<i>Part I: Getting Started with a Virtualization Project</i>	<i>7</i>
Chapter 1: Wrapping Your Head around Virtualization	9
Chapter 2: Making a Business Case for Virtualization	33
Chapter 3: Understanding Virtualization: Technologies and Applications	49
Chapter 4: Peeking at the Future of Virtualization	83
<i>Part II: Server Virtualization.....</i>	<i>99</i>
Chapter 5: Deciding Whether Server Virtualization Is Right for You.....	101
Chapter 6: Performing a Server Virtualization Cost-Benefit Analysis	111
Chapter 7: Managing a Virtualization Project	129
Chapter 8: Choosing Hardware for Your Server Virtualization Project.....	159
<i>Part III: Server Virtualization Software Options.....</i>	<i>187</i>
Chapter 9: Migrating to Your New Virtualized Environment	189
Chapter 10: Managing Your Virtualized Environment	209
Chapter 11: Creating a Virtualized Storage Environment	221
<i>Part IV: Implementing Virtualization.....</i>	<i>243</i>
Chapter 12: Implementing VMware Server	245
Chapter 13: Implementing Fedora Virtualization	267
Chapter 14: Implementing XenExpress.....	291
<i>Part V: The Part of Tens</i>	<i>319</i>
Chapter 15: Ten Steps to Your First Virtualization Project	321
Chapter 16: Ten Virtualization Pitfalls to Avoid	329
Chapter 17: Ten Great Resources on Virtualization	335
<i>Index</i>	<i>341</i>

Table of Contents

<i>Foreword</i>	<i>xvii</i>
-----------------------	-------------

<i>Introduction</i>	<i>1</i>
---------------------------	----------

Why Buy This Book?	2
Foolish Assumptions	2
How This Book Is Organized	3
Part I: Getting Started with a Virtualization Project	3
Part II: Server Virtualization	4
Part III: Server Virtualization Software Options	4
Part IV: Implementing Virtualization	5
Part V: The Part of Tens	5
Icons Used in This Book	5
Where to Go from Here	6

<i>Part I: Getting Started with a Virtualization Project</i>	<i>7</i>
--	----------

Chapter 1: Wrapping Your Head around Virtualization	9
--	----------

Virtualization: A Definition	10
Why Virtualization Is Hot, Hot, Hot —	
The Four Drivers of Virtualization	11
Trend #1: Hardware is underutilized	11
Trend #2: Data centers run out of space	13
Trend #3: Energy costs go through the roof	14
Trend #4: System administration costs mount	15
Four trends mean virtualization is hot	16
Sorting Out the Types of Virtualization	16
Client virtualization	17
Server virtualization	20
Storage virtualization	28
Creating the Virtualized Enterprise	31

Chapter 2: Making a Business Case for Virtualization	33
---	-----------

Virtualization Lowers Hardware Costs	34
When servers just coast along	35
A handy (if anonymous) real-world example	36

Virtualization Increases IT Operational Flexibility	37
Virtualization's promise	38
A handy (named) example	39
Virtualization Reduces IT Operations Costs	41
Virtualization makes hardware maintenance and upgrades easier and cheaper	41
Virtualization means fewer servers and lower IT costs	43
Virtualization Lowers Energy Costs.....	44
Software Licensing Costs: A Challenge for Virtualization.....	45
General software licensing practices	46
Software licensing meets virtualization.....	47

Chapter 3: Understanding Virtualization: Technologies and Applications49

Virtualization Technologies	50
Operating system virtualization	51
Hardware emulation.....	54
Paravirtualization	57
Assessing the implications of virtualization technology	61
Virtualization Applications	62
Development and testing	62
Training.....	65
Server consolidation.....	67
Failover/high availability/load balancing.....	69
Products for failover/high availability/clustering/ load balancing: The usual suspects	75
Future directions in high availability	75
Server pooling.....	76
Disaster recovery	77

Chapter 4: Peeking at the Future of Virtualization83

Virtualization Gets Integrated into Operating Systems	84
Virtualized Software: Delivered to Your Door Preinstalled.....	85
The software installation headache.....	86
The virtual appliance: Oh, what a relief it is!	87
Virtualization Diffusing into the Internet	90
Software as a Service (SaaS)	90
Ever more virtualization = ever less hardware.....	91
The Changing Skill Set of IT Personnel.....	93
Software Pricing: How Will It Respond to Virtualization?	95

Part 11: Server Virtualization99**Chapter 5: Deciding Whether Server Virtualization
Is Right for You101**

How to Decide Whether You Should Use Server Virtualization.....	102
Figuring out whether virtualization is right for your organization.....	102
Looking for clear financial benefits.....	103
Checking to see whether important prerequisite conditions are in place.....	103
When Not to Use Virtualization.....	106
When your computing environment is static	106
When your software providers refuse support within a virtualized infrastructure	107
When your applications don't lend themselves to virtualization.....	107
When your organization is unwilling to invest to improve operations	110

**Chapter 6: Performing a Server Virtualization
Cost-Benefit Analysis111**

Getting Your Cost-Benefit Ducks in a Row	111
Defining a solution	112
Looking at current costs.....	112
Identifying virtualization costs.....	116
Identifying the financial benefits of virtualization	119
Creating your virtualization cost-benefit spreadsheet.....	122
The Cost-Benefit Bottom Line	127

Chapter 7: Managing a Virtualization Project129

Understanding the Virtualization Life Cycle.....	130
Planning your virtualization journey	131
Implementing your virtualization solution	132
Operating your virtualization solution.....	133
Creating Your Virtualization Plan.....	134
Identify your use cases	134
Evaluate your organizational structure	137
Select your architecture	139
Implementing Your Virtualization Solution	142
Confirming your planning assumptions and conclusions.....	143
Choosing your virtualization software product(s)	144
Choosing your virtualization hardware.....	145
Performing a pilot implementation.....	148
Migrating to your production virtualization environment.....	152
Administering your new virtualized infrastructure	154

Chapter 8: Choosing Hardware for Your Server Virtualization Project159

Taking Hardware Seriously	160
The multiple-eggs-in-one-basket syndrome	163
The Big Four resources of servers	164
The big four resources — the story so far	174
Choosing Servers	175
Reusing your existing 32-bit hardware	176
Using 64-bit servers.....	178
Designed for virtualization: The new generation of servers	180
Making the Hard Hardware Choices	181
But Wait, There's More: Future Virtualization Hardware Development ..	183
Virtualization-enabled chips	183
More efficient memory	184
Faster networking for virtualized machines	184
Better device support in virtualized systems	186

Part III: Server Virtualization Software Options 187**Chapter 9: Migrating to Your New Virtualized Environment189**

Moving from Physical to Virtual: An Overview	190
Getting Ready to Move: Preparing the Virtualized Environment.....	191
Preparing existing systems for migration	192
Preparing the new virtualization servers.....	193
Install necessary infrastructure (hardware and software)	197
Migrating Your Physical Servers	200
Manual migration: The hands-on approach.....	202
Automated migration: Let software do the work	203
Moving to Production	207

Chapter 10: Managing Your Virtualized Environment209

Managing Virtualization: The Next Challenge	211
Managing Free Virtualization	212
Virtualization Management: The Two Philosophies	213
Virtualization as a specialized resource	214
Virtualization as an equal member of the data center	215
Making Sense of Virtualization Management.....	217
Where you are in the virtualization journey	217
Your current system management approach	218
What your current hardware mix looks like	219
Deciding on Your Virtualization Management Approach	220

Chapter 11: Creating a Virtualized Storage Environment223

Storage Overview	223
What the heck is DAS?	224
Shared storage	226
Why use shared storage?	228
LUNS and RAID	228
Data redundancy	230
Cost implications of storage	231
Choosing Storage for Virtualization.....	231
DAS and virtualization	231
Shared storage and virtualization	234
Storage and the Different Types of Virtualization.....	236
Operating system virtualization(containers)	236
Hardware emulation virtualization	237
Paravirtualization	238
Storage and the Virtualization Journey	239

Part IV: Implementing Virtualization243**Chapter 12: Implementing VMware Server245**

Understanding VMware Server Architecture: Pros and Cons	246
Getting Your (Free) Copy of VMware Server	247
Acquiring VMware Server	248
Installing VMware Server	248
Creating a Guest Virtual Machine	252
Starting the VMware Server Console	252
Installing a new virtual machine.....	253
Installing an Operating System.....	258
Can I Skip the Boring OS Installation Process?.....	261
Can I Skip the Boring Application Installation Process?	263

Chapter 13: Implementing Fedora Virtualization267

Obtaining Fedora 7.....	268
Installing Fedora 7.....	269
Creating a Guest Virtual Machine	278
Installing a Guest Operating System.....	286

Chapter 14: Implementing XenExpress291

What Is XenSource, Anyway?.....	292
Obtaining XenSource XenExpress	292
Installing XenExpress	293
Installing XenConsole	296
Working with XenConsole	297

Creating a Guest Virtual Machine.....	300
Installing Paravirtualized Drivers.....	311
A bit of background, or QEMU who?.....	311
Baby, you can drive my virtualization environment.....	312
Accessing a Windows Guest VMwith an RDP Client.....	314

Part V: The Part of Tens319

Chapter 15: Ten Steps to Your First Virtualization Project321

Recite After Me: Virtualization Is a Journey, Not a Product	321
Evaluate Your Use Cases	322
Review Your Operations Organizational Structure.....	323
Define Your Virtualization Architecture.....	324
Select Your Virtualization Product(s).....	325
Select Your Virtualization Hardware.....	325
Perform a Pilot Implementation	326
Implement Your Production Environment.....	326
Migrate Your Physical Servers	326
Manage Your Virtualized Infrastructure.....	327

Chapter 16: Ten Virtualization Pitfalls to Avoid329

Don't Wait for All the Kinksto Be Worked Out.....	329
Don't Skimp on Training.....	330
Don't Apply Virtualization in AreasThat Are Not Appropriate	330
Don't Imagine That VirtualizationIs Static	331
Don't Skip the "Boring" Stuff.....	332
Don't Overlook a Business Case.....	332
Don't Overlook the Importanceof Organization.....	332
Don't Forget to Research Your Software Vendor Support Policies	333
Don't Overlook the Importanceof Hardware	333
Don't Forget to Have a Project Party	334

Chapter 17: Ten Great Resourceson Virtualization335

Get Free Virtualization Software	336
Get Great Content about Virtualization.....	336
Get the Latest News about Virtualization	336
Read Blogs about Virtualization.....	337
Keep Up with Hardware Developments Relating to Virtualization.....	337
Find Out More about Virtualization	338
Attend Virtualization Events.....	338
Take Advantage of Vendor Information.....	339
Keep Up with Storage Virtualization	339
Get the Latest and Last Word on Virtualization	339

Index.....341

Foreword

When Bernard invited me to write an introduction to this book, I found myself reminded of a frequently repeated conversation with my father, who is a retired engineer. Typically, it goes like this: “Simon, what does virtualization do?” — followed by a lengthy reply from me and then a long pause from my father — “And why is that useful?” Now, I certainly don’t think that my father really has much use for server virtualization, but a lot more people do need it — and need to understand it — than currently use it.

Although virtualization is all the rage in the tech industry press, and savvy market watchers have observed the exciting IPO of VMware, and Citrix’s acquisition of my own company, XenSource, the market for virtualization software is largely unaddressed. Depending on whose research you read, only 7 percent or so of x86 servers are virtualized, and only a tiny fraction of desktop or mobile PCs are virtualized. But the virtualization market is white hot, and every day new announcements in storage, server, and network virtualization make the picture more complex and harder to understand.

Virtualization For Dummies is the perfect way to develop a complete understanding of both the technology and the benefits of virtualization. Arguably, virtualization is simply a consequence of Moore’s Law — the guideline developed by Intel founder Gordon Moore that predicts a doubling in the number of transistors per unit area on a CPU every couple of years. With PCs and servers becoming so incredibly powerful, the typical software suites that most users would install on a single physical server a few years ago now consume only a few percent of the resources of a modern machine. Virtualization is simply a consequence of the obvious waste of resources — allowing a machine to run multiple virtualized servers or client operating systems simultaneously. But if that were all that were needed, there wouldn’t be such a fuss about virtualization. Instead, virtualization is having a profound impact on data center architectures and growth, on software lifecycle management, on security and manageability of software, and the agility of IT departments to meet with new challenges. And it is these opportunities and challenges that urgently need to be articulated to technologists and business leaders alike in an accessible and understandable way.

Having spent many enjoyable hours with Bernard Golden, a recognized open source guru, President and CEO of Navica, and self-taught virtualization expert, I cannot think of a better-qualified author for a book whose objective is to cut through the hype and clearly and succinctly deal with virtualization and its effects on IT and users alike. I always look forward to reading

Bernard's frequently published commentaries on Xen, VMware, and Linux, which combine his hands-on experience with those products and a rare depth of insight into industry dynamics. I know firsthand that Bernard is a master of the subject of virtualization because he is one of the most persistent and demanding beta testers of XenEnterprise, XenSource's server virtualization product, where his feedback has provided us with terrific guidance on how to improve the product overall. This, together with Bernard's incisive, clear, and articulate style, makes this book a pleasure to read and a terrific contribution to the virtualization industry — a concise categorization of virtualization that will further the understanding of the technology and its benefits, driving uptake of virtualization generally. It is with great pleasure that I strongly recommend that you read this book.

Simon Crosby

CTO, XenSource, Inc.

Introduction

If you work in tech, there's no way you haven't heard the term *virtualization*. Even if you don't work in tech, you might have been exposed to virtualization. In August 2007, virtualization's leading company, VMware, went public with the year's most highly anticipated IPO. Even people who confuse virtualization with visualization sit up and pay attention when a blockbuster IPO comes to market. To show how hot the sector is, VMware was bought by the storage company EMC for \$625 million in 2004, but it has, as of this writing, a market capitalization of \$25.6 *billion*.

The excitement and big dollars illustrate a fundamental reality about virtualization: It's transforming the way computing works. Virtualization is going to fundamentally change the way you implement and manage data centers, the way you obtain and install software, and the way you think about the speed with which you can respond to changing business conditions. The changes that virtualization will cause in your work environment will be so profound that, in ten years time, you'll look back on the traditional ways of managing hardware and software the way your grandparents looked back on operator-assisted telephone dialing after the introduction of direct dialing.

I wrote this book because I'm convinced that the world is on the cusp of an enormous change in the use of information technology, also known as IT. In the past, IT was expensive, so it was limited to must-have applications such as accounting and order tracking. In the past, IT was complex, so it had to be managed by a group of wizards with their own special language and incantations. That's all changing.

In the future, IT will be cheap, so applications will be ubiquitous, and low-priority applications will finally get their day in the sun. In the future, implementing IT will be simple, so groups outside of IT will shun the wizards' robes and arcane language and implement their own applications, which will, of course, make central IT's role even more important because it will have to create a robust yet malleable infrastructure.

Instead of IT being this special thing that supports only certain aspects of a business, it will become pervasive, suffusing throughout every business operation and every business interaction. It's an incredibly exciting time for IT; I compare it to the rise of mass production made possible by Henry Ford. Because of Henry Ford, automobiles went from playthings for the wealthy to everyday belongings of the masses, and society was transformed by mobility and speed. Virtualization is the mass production of IT.

Just as the automobile industry underwent rapid transformation after Ford invented mass production in 1913, the virtualization marketplace is transforming the IT industry today. One of the biggest challenges for this book is to present a coherent and unified view of the topic even though virtualization is evolving at an incredible pace. At times, I felt that writing this book was like trying to nail Jell-O to the wall. During just one week of the writing of this book, the IPO of VMware went from an event no one had even considered to the technology financial event of the year; in the same week, XenSource, the commercial sponsor of the open source Xen virtualization project, was purchased by Citrix for \$500 million. Furthermore, myriads of virtualization technology and product announcements occurred, making me, at times, wish I could push a Pause button on the market so that I could have a hope of completing an up-to-date book. Alas, virtualization's fevered evolution shows no sign of diminishing — good for the virtualization user, challenging for the virtualization writer.

Why Buy This Book?

Even though virtualization is changing the face of technology, it is, unfortunately, still riddled with the complexities and — especially — the arcane language of tech. Two seconds into a conversation about virtualization, you'll start hearing terms like *hypervisor* and *bare metal*, which sound, respectively, like something from *Star Wars* and an auto shop class.

It's unfortunate that virtualization can be difficult to approach because of this specialized terminology. It's especially unfortunate because understanding and applying virtualization will be, in the near future, a fundamental skill for everyone in IT — and for many people working in other disciplines like marketing and finance. Consequently, having a strong grounding in virtualization is critical for people wanting to participate in the IT world of the future.

This book is designed to provide a thorough introduction to the subject. It assumes that you have no knowledge of virtualization and will leave you (I hope) with a good grasp of the topic. My objective is for you to be completely comfortable with virtualization and its many elements and to be able to participate and contribute to your organization's virtualization initiatives. The book also serves as a jumping-off point for deeper research and work in virtualization.

Foolish Assumptions

This book doesn't assume you know much about virtualization beyond having heard the term. You don't have to know any of the technical details of the topic, and you certainly don't need to have done hands-on work with

virtualization. (The book provides the opportunity to do hands-on work with virtualization, with three chapters devoted to installing and implementing different virtualization products.)

I define every virtualization term you encounter. I also make it a point to thoroughly explain complex topics so that you can understand the connections between different virtualization elements.

The book *does* assume that you have a basic understanding of computers, operating systems, and applications and how they work together to enable computers to do useful work. Because virtualization shuffles the placement and interaction of existing system software and hardware layers, it's important to have a grasp of how things are traditionally done. However, if you've worked with computers, used an operating system, and installed applications, you should have the knowledge base to make use of the book's content.

How This Book Is Organized

As is the case with other *For Dummies* books, this book doesn't assume that you'll begin on page one and read straight through to the end. Each chapter is written to stand alone, with enough contextual information provided so that you can understand the chapter's content. Cross-references are provided to other chapters that go into more detail about topics lightly touched on in a given chapter.

You'll soon notice, though, that individual chapters are grouped together in a somewhat-less-than-random order. The organizing principle here is the *part*, and this book has five of them.

Part I: Getting Started with a Virtualization Project

Getting a good grounding in a subject is critical to understanding it and, more important, to recognizing how you can best take advantage of it. Part I provides a whirlwind tour of the world of virtualization — from where it is today to beyond where it will be tomorrow.

Chapter 1 is where you get an overview of virtualization, including an introduction to why it's such a hot topic. Chapter 1 also discusses the basic philosophy of virtualization — the abstraction of computer functionality from physical resources. Chapter 2 describes the business reasons that are driving virtualization's explosive growth, and it discusses how you can make a business case for your virtualization project. If you want a deeper understanding of the different technologies that make up virtualization as well as the different

ways virtualization is applied in everyday use, Chapter 3 is for you. Finally, if you want to get a sense of where virtualization is heading, Chapter 4 provides a glimpse of the exciting initiatives that are being made possible by virtualization.

Part II: Server Virtualization

Server virtualization is where the hottest action is in today's virtualization world. The most obvious use cases and the most immediate payoffs are available with server virtualization, and this part covers it all.

Chapter 5 gives you information for a litmus test to inform you whether server virtualization makes sense for you. The chapter lets you do a bit of self-testing to see whether your organization is ready to implement virtualization. Just as important, the chapter gives you the tools you'll need to find out whether it doesn't make sense for you to implement virtualization. Chapter 6 provides in-depth information on how to make a financial assessment of your virtualization project, including how to create a spreadsheet to calculate all the costs and benefits of a potential virtualization project. Chapter 7 discusses the all-important topic of how to manage a virtualization project; there's far more involved here than just installing a virtualization hypervisor. Finally, Chapter 8 discusses a very important topic — the hardware you'll use to run your virtualization software. There are many exciting developments in hardware with significant influence on the operational and financial benefits of virtualization.

Part III: Server Virtualization Software Options

Sometimes I see a movie with a happy ending and wonder "Yeah, but how did the rest of their lives turn out?" Virtualization can be something like that. If you listen to vendors, you just install their software and — presto! — instant virtualized infrastructure. Just like real life isn't like the movies, real infrastructure isn't like that, either, and Part III helps you have a true happy ending.

Chapter 9 deals with the critical issue of how to migrate an existing physical infrastructure to a virtualized one. (*Hint:* It's more complex than the vendors claim.) Chapter 10 addresses managing a virtualized infrastructure; there are a plethora of options, and this chapter provides help in deciding which option is a good choice for you. Chapter 11 addresses a topic that's often an afterthought in virtualization: storage. For many organizations, virtualization provides the impetus to move to shared (also known as *virtualized*) storage. It's important to know what your storage options are and how to select one.

Part IV: Implementing Virtualization

If you're like me, theoretical understanding goes just so far. Then I want to roll up my sleeves and get my hands dirty. If you're like that as well, rejoice! Part IV can feed your hands-on hunger. In this part, I present three different examples of how to install and use virtualization. Best of all, each of the products used for the examples is available at no cost, making it possible for you to work along with them with no financial commitment.

Chapter 12 illustrates how to implement VMware Server as well as how to install a guest virtual machine. Chapter 13 works through Xen virtualization via the open source Linux distribution Fedora. Chapter 14 also illustrates a Xen-based virtualization, but the chapter uses the free XenExpress product from XenSource to share a different way of applying Xen virtualization.

Part V: The Part of Tens

Every *For Dummies* book concludes with a few chapters that provide a final burst of valuable information delivered in a sleek, stripped-down format — the time-honored ten-point list.

In Chapter 15, you get a list of the ten must-do steps for your first virtualization project. Chapter 16 shares ten no-no's to avoid in a virtualization project. And Chapter 17 gives you ten great virtualization resources for you to use after you finish this book.

Icons Used in This Book



This icon flags useful, helpful tips, or shortcuts.



This icon marks something that might be good to store away for future reference.



Pay attention. The bother you save might be your own.



This icon highlights tidbits for the more technically inclined that I hope augment their understanding — but I won't be offended if less-technically inclined readers hurry through with eyes averted.

Where to Go from Here


Pick a page and start reading! You can use the Table of Contents as a guide, or my parts description in this introduction, or you can leaf through the index for a particular topic. If you're an accounting type, you might jump right into Chapter 6 — the chapter with all the lovely spreadsheets. If you're a hardcore techie type, you might want to check out Chapter 8 — yes, yes, the hardware chapter. Wherever you start, you'll soon find yourself immersed in one of the more exciting stories to come down the tech pipe in a long time.

Part I

Getting Started with a Virtualization Project

The 5th Wave

By Rich Tennant



"Why, of course. I'd be very interested in seeing this new milestone in our server virtualization project."

In this part . . .

Virtualization is a hot topic. But what if you don't know enough to even fake it around the water cooler? Not to worry. This part gives even the virtualization-challenged a thorough introduction to the subject of virtualization.

It begins with an overview of virtualization — what it is and what all the different types of virtualization are. You didn't know there are different types? Well, this part is perfect for you.

For good measure, this part also includes an in-depth look at the technologies and applications of virtualization so that you can form an opinion of how it might be usefully applied to your own environment.

I conclude this part with a look at the future of virtualization. The area is evolving rapidly, and looking forward to future developments is exciting and, perhaps, sobering. Certainly I expect that virtualization will transform the way hardware and software is sold, so a peek into the future is well worth your time.

Chapter 1

Wrapping Your Head around Virtualization

In This Chapter

- ▶ Finding out what virtualization is
 - ▶ Figuring out what has everyone so excited
 - ▶ Seeing how virtualization is applied
 - ▶ Working through the challenges of virtualization
-

It seems like everywhere you go these days, someone is talking about virtualization. Technical magazines trumpet the technology on their covers. Virtualization sessions are featured prominently at technology conferences. And, predictably enough, technology vendors are describing how *their* product is the latest word in virtualization.

If you have the feeling that everyone else in the world understands virtualization perfectly while you're still trying to understand just what it is and how you might take advantage of it, take heart. Virtualization *is* a new technology. (Actually, it's a pretty well-established technology, but a confluence of conditions happening just now has brought it into new prominence — more on that later in this chapter.) Virtualization *is* a technology being widely applied today with excellent operational and financial results, but it's by no means universally used or understood. That's the purpose of this book: to provide you with an introduction to the subject so that you can understand its promise and perils and create an action plan to decide whether virtualization is right for you, as well as move forward with implementing it should you decide it *is* right for you.

Sadly, not even this book can protect you from the overblown claims of vendors; there is no vaccine strong enough for that disease. This book helps you sort out the hope from the hype and gives you tools to feel confident in making your virtualization decisions.

Virtualization: A Definition

Virtualization refers to a concept in which access to a single underlying piece of hardware, like a server, is coordinated so that multiple guest operating systems can share that single piece of hardware, with no guest operating system being aware that it is actually sharing anything at all. (A *guest operating system* is an operating system that's hosted by the underlying virtualization software layer, which is often, you guessed it, called the *host system*.) A guest operating system appears to the applications running on it as a complete operating system (OS), and the guest OS itself is completely unaware that it's running on top of a layer of virtualization software rather than directly on the physical hardware.

Actually, you've had experience with something like this when you used a computer. When you interact with a particular application, the operating system "virtualizes" access to the underlying hardware so that only the application you're using has access to it — only your program is able to manipulate the files it accesses, write to the screen, and so on. Although this description oversimplifies the reality of how operating systems work, it captures a central reality: The operating system takes care of controlling how applications access the hardware so that each application can do its work without worrying about the state of the hardware. The operating system *encapsulates* the hardware, allowing multiple applications to use it.

In *server* virtualization — the most common type of virtualization — you can think of virtualization as inserting another layer of encapsulation so that multiple operating systems can operate on a single piece of hardware. In this scenario, each operating system believes it has sole control of the underlying hardware, but in reality, the virtualization software controls access to it in such a way that a number of operating systems can work without colliding with one another. The genius of virtualization is that it provides new capability without imposing the need for significant product or process change.

Actually, that last statement is a bit overbroad. A type of virtualization called paravirtualization does require some modification to the software that uses it. However, the resulting excellent performance can make up for the fact that it's a little less convenient to use. Get used to this exception business; the subject of virtualization is riddled with general truths that have specific exceptions. Although you have to take account of those exceptions in your particular project plans, don't let these exceptions deter you from the overarching circumstances. The big picture is what you need to focus on to understand how virtualization can help *you*.



Virtualization is actually a simple concept made complex by all the exceptions that arise in particular circumstances. It can be frustrating to find yourself stymied by what seems to be a niggling detail, but unfortunately, that's the reality of virtualization. If you stop to think about it, the complexity makes sense — you're moving multiple operating systems and applications onto a

new piece of software called a *hypervisor*, which in turn talks to underlying hardware. Of course it's complex! But don't worry, if you hang in there, it usually comes out right in the end. Chapters 12 through 14 offer real examples of how to install several flavors of virtualization software and successfully put guest OSes onto the software. Work through the examples, and you'll be an expert in no time!

So, if you take nothing more away from this section than the fact that virtualization enables you to share a hardware resource among a number of other software systems, that's enough for you to understand the next topic — what's making virtualization so important *now*.

Why Virtualization Is Hot, Hot, Hot — The Four Drivers of Virtualization

Despite all the recent buzz about it, virtualization is by no means a new technology. Mainframe computers have offered the ability to host multiple operating systems for over 30 years. In fact, if you begin to discuss it, you might suffer the misfortune of having someone begin to regale you with tales of how *he* did virtualization in the old days.

The truth is that your old gaffer is right. Yeah, virtualization as a technology is nothing new, and yeah, it's been around for many years, but it was confined to “big iron” (that is, mainframes). Four trends have come together in just the past couple of years that have moved virtualization from the dusty mainframe backroom to a front-and-center position in today's computing environment.



When you take a look at these trends, you can immediately recognize why virtualization is much more than the latest technology fad from an industry that has brought you more fads than the fashion industry.

Trend #1: Hardware is underutilized

I recently had an opportunity to visit the Intel Museum located at the company's headquarters in Santa Clara, California. The museum contains a treasure trove of computing — from a giant design-it-yourself chip game to a set of sterile clothes (called a *bunny suit*) you can put on while viewing a live camera feed from a chip-manufacturing plant. It's well worth a visit. But tucked near the back of the museum, in a rather undistinguished case, is ensconced one of the critical documents of computing. This document, despite its humble presentation, contains an idea that has been key to the development of computing in our time.

I refer to the article by Intel cofounder Gordon Moore in the April 1965 issue of *Electronics Magazine*, in which he first offered his observation about the compounding power of processor computing power, which has come to be known as “Moore’s Law.”

In describing the increasing power of computing power, Moore stated: “The complexity for minimum component costs has increased at a rate of roughly a factor of two per year.” Clearly, Moore wasn’t in charge of marketing at Intel, but if you translate this into something the average human can understand, he means that each year (actually, most people estimate the timeframe at around 18 months), for a given size processor, twice as many individual components can be squeezed onto a similarly sized piece of silicon. Put another way, every new generation of chip delivers twice as much processing power as the previous generation — at the same price.

This rapid doubling of processing power has had a tremendous impact on daily life, to an extent that’s difficult to comprehend for most people. Just over ten years ago, I ran the engineering group of a large enterprise software vendor. Everybody in the group knew that a major part of the hassles involved in getting a new release out the door involved trying to get acceptable performance out of the product on the then-latest generation of hardware. The hardware just wasn’t that capable. Today’s hardware, based upon the inexorable march of Moore’s Law, is around 100,000 times as powerful — in ten short years!

What you need to keep in mind to understand Moore’s Law is that the numbers that are continuing to double are themselves getting larger. So, if you take year one as a base, with, say, processing power of 100 million instructions per second (MIPS) available, then in year two, there will be 200; in year three, 400; and so on. Impressive, eh? When you get out to year seven or eight, the increase is from something like 6,400 to 12,800 in one generation. It has grown by 6,400. And the next year, it will grow by 12,800. It’s mind boggling, really.



Moore’s Law demonstrates increasing returns — the amount of improvement itself grows over time because there’s an exponential increase in capacity for every generation of processor improvement. It’s that exponential increase that’s responsible for the mind-boggling improvements in computing — and the increasing need for virtualization.

And that brings me to the meaning of this trend. Unlike ten years ago, when folks had to sweat to get software to run on the puny hardware that was available at that time, today the hardware is so powerful that software typically uses only a small portion of the available processing power. And this causes a different type of problem.

Today, many data centers have machines running at only 10 or 15 percent of total processing capacity. In other words, 85 or 90 percent of the machine’s power is unused. In a way, Moore’s Law is no longer relevant to most companies because they aren’t able to take advantage of the increased power

available to them. After all, if you're hauling a 50-pound bag of cement, having a truck come out that can carry 20,000 pounds instead of this year's model that can only carry 10,000 pounds is pretty irrelevant for your purposes. However, a lightly loaded machine still takes up room and draws electricity, so the cost of today's underutilized machine is nearly the same as if it was running at full capacity.

It doesn't take a rocket scientist to recognize that this situation is a waste of computing resources. And, guess what? With the steady march of Moore's Law, next year's machine will have twice as much spare capacity as this year's — and so on, for the foreseeable future. Obviously, there ought to be a better way to match computing capacity with load. And that's what virtualization does — by enabling a single piece of hardware to seamlessly support multiple systems. By applying virtualization, organizations can raise their hardware use rates from 10 or 15 percent to 70 or 80 percent, thereby making much more efficient use of corporate capital.



Moore's Law not only enables virtualization, but effectively makes it mandatory. Otherwise, increasing amounts of computing power will go to waste each year.

So, the first trend that's causing virtualization to be a mainstream concern is the unending growth of computing power brought to you by the friendly folks of the chip industry. By the way, the same trend that's described in chips by Moore's Law can be observed in the data storage and networking arenas as well. They just don't have a fancy name for their exponential growth — so maybe Gordon Moore was a marketing genius after all. However, the rapid improvement in these other technology areas means that virtualization is being explored for them as well — and because I like being complete, I work in coverage of these areas later on in this book.

Trend #2: Data centers run out of space

The business world has undergone an enormous transformation over the past 20 years. In 1985, the vast majority of business processes were paper based. Computerized systems were confined to so-called backroom automation: payroll, accounting, and the like.

That has all changed, thanks to the steady march of Moore's Law. Business process after business process has been captured in software and automated, moving from paper to computers.

The rise of the Internet has exponentially increased this transformation. Companies want to communicate with customers and partners in real time, using the worldwide connectivity of the Internet. Naturally, this has accelerated the move to computerized business processes.

To offer a dramatic example, Boeing's latest airliner, the 787 Dreamliner, is being designed and built in a radically new way. Boeing and each of its suppliers use Computer-Aided Design (CAD) software to design their respective parts of the plane. All communication about the project uses these CAD designs as the basis for discussion. Use of CAD software enables testing to be done in computer models rather than the traditional method of building physical prototypes, thereby speeding completion of the plane by a year or more.

As you might imagine, the Dreamliner project generates enormous amounts of data. Just one piece of the project — a data warehouse containing project plans — runs to 19 terabytes of data.

Boeing's experience is common across all companies and all industries. How big is the explosion of data? In 2003, the world's computer users created and stored 5 exabytes (each exabyte is 1 million terabytes) of new data. A recent study by the Enterprise Strategy Group predicted that governments and corporations will store over 25 exabytes of data by the year 2010. Certainly, the trend of data growth within organizations is accelerating. The growth of data can be easily seen in one key statistic: In 2006, the storage industry shipped as much storage in one month as it did in the entire year of 2000. The research firm IDC estimates that total storage shipped will increase 50 percent per year for the next five years.



The net effect of all this is that huge numbers of servers have been put into use over the past decade, which is causing a real-estate problem for companies: They're running out of space in their data centers. And, by the way, that explosion of data calls for new methods of data storage, which I also address in this book. These methods go by the common moniker of *storage virtualization*, which, as you might predict, encapsulates storage and abstracts it from underlying network storage devices.

Virtualization, by offering the ability to host multiple guest systems on a single physical server, helps organizations to reclaim data center territory, thereby avoiding the expense of building out more data center space. This is an enormous benefit of virtualization because data centers cost in the tens of millions of dollars to construct. You can find out more about this in Chapter 4, where I discuss this trend, which is usually referred to as *consolidation* and is one of the major drivers for organizations to turn to virtualization.



Take a look at your data center to understand any capacity constraints you're operating with. If you're near capacity, you need virtualization — stat!

Trend #3: Energy costs go through the roof

In most companies' strategic thinking, budgeting power costs used to rank somewhere below deciding what brand of soda to keep in the vending machines. Companies could assume that electrical power was cheap and endlessly available.

Several events over the past few years have changed that mindset dramatically:

- ✓ The increasing march of computerization discussed in Trend #2, earlier in this chapter, means that every company is using more power as their computing processes expand.
- ✓ The assumption regarding availability of reliable power was challenged during the California power scares of a few years ago. Although later evidence caused some reevaluation about whether there was a true power shortage (can you say “Enron?”), the events caused companies to consider whether they should look for ways to be less power dependent.
- ✓ As a result of the power scares and Pacific Gas & Electric’s resulting bankruptcy, power costs in California, home to Silicon Valley, have skyrocketed, making power a more significant part of every company’s budget. In fact, for many companies, electricity now ranks as one of the top five costs in their operating budgets.

The cost of running computers, coupled with the fact that many of the machines filling up data centers are running at low utilization rates, means that virtualization’s ability to reduce the total number of physical servers can significantly reduce the overall cost of energy for companies.



Data center power is such an issue that energy companies are putting virtualization programs into place to address it. See Chapter 5 to find out about an innovative virtualization rebate program Pacific Gas & Electric has put into place.

Trend #4: System administration costs mount

Computers don’t operate on their own. Every server requires care and feeding by system administrators who, as part of the operations group, ensure that the server runs properly. Common system administration tasks include monitoring hardware status; replacing defective hardware components; installing operating system (OS) and application software; installing OS and application patches; monitoring critical server resources such as memory and disk use; and backing up server data to other storage mediums for security and redundancy purposes.

As you might imagine, this job is pretty labor intensive. System administrators don’t come cheap. And, unlike programmers, who can be located in less expensive offshore locales, system administrators are usually located with the servers due to their need to access the physical hardware.

The steady increase in server numbers has meant that the job market for system administrators has been good — very good.

As part of an effort to rein in operations cost increases, virtualization offers the opportunity to reduce overall system administration costs by reducing the overall number of machines that need to be taken care of. Although many of the tasks associated with system administration (OS and application patching, doing backups) continue even in a virtualized environment, some of them disappear as physical servers are migrated to virtual instances. Overall, virtualization can reduce system administration requirements by 30 to 50 percent per virtualized server, making virtualization an excellent option to address the increasing cost of operations personnel.



Virtualization reduces the amount of system administration work necessary for hardware, but it doesn't reduce the amount of system administration required for guest OSes. Therefore, virtualization improves system administration, but doesn't make it vanish.

Four trends mean virtualization is hot

Looking at these four trends, you can see why virtualization is a technology whose time has come. The exponential power growth of computers, the substitution of automated processes for manual work, the increasing cost to power the multitude of computers, and the high personnel cost to manage that multitude all cry out for a less expensive way to run data centers. In fact, a newer, more efficient method of running data centers is critical because, given the four trends, the traditional methods of delivering computing are becoming cost prohibitive. Virtualization is the solution to the problems caused by the four trends I outline here.

Sorting Out the Types of Virtualization

If you've made it this far in this chapter, you (hopefully) have a rough idea of virtualization and why it's an important development. Your next step involves determining what your options are when it comes to virtualization. In other words, what are some common applications of the technology?

Virtualization has a number of common uses, all centered around the concept that virtualization represents an abstraction from physical resources. In fact, enough kinds of virtualization exist to make it a bit confusing to sort out how you might apply it in your organization.

I do what I can to sort out the virtualization mare's nest. If you're okay with gross generalizations, I can tell you that there are three main types of virtualization: client, server, and storage. Within each main type are different approaches or *flavors*, each of which has its benefits and drawbacks. The next few sections give brief descriptions of each of the three types of virtualization, along with examples of common implementations of them.

Client virtualization

Client virtualization refers to virtualization capabilities residing on a *client* (a desktop or laptop PC). Given that much of the earlier discussion of the driving forces behind virtualization focuses on the problems of the data center, you might wonder why virtualization is necessary for client machines at all.



The primary reason organizations are interested in pursuing client virtualization solutions has to do with the challenges they face in managing large numbers of computers controlled by end users. Although machines located in data centers typically have strict procedures about what software is loaded on them and when they're updated with new software releases, end user machines are a whole different story.

Because loading software is as easy as sticking a disc into the machine's CD drive (or a thumb drive into a USB slot), client machines can have endless amounts of non-IT-approved software installed. Each application can potentially cause problems with the machine's operating system as well as other approved applications. Beyond that, other nefarious software can get onto client machines in endless ways: via e-mail viruses, accidental spyware downloads, and so on. And, the hard truth is that Microsoft Windows, the dominant client operating system, is notorious for attracting attacks in the form of malware applications.

Added to the end user-caused problems are the problems inherent to client machines in general: keeping approved software applications up to date, ensuring the latest operating system patches are installed, and getting recent virus definitions downloaded to the machine's antivirus software.

Mixed together, this stew is a miserable recipe for IT. Anything that makes the management of client machines easier and more secure is of definite interest to IT. Client virtualization offers the potential to accomplish this.

Three main types — or flavors, if you will — of client virtualization exist: application packaging, application streaming, and hardware emulation.

Application packaging

Although the specifics of how application packaging is accomplished vary from one vendor to another, all the methods share a common approach: isolating an application that runs on a client machine from the underlying operating system. By isolating the application from the operating system, the application is unable to modify underlying critical operating system resources, making it much less likely that the OS will end up compromised by malware or viruses.

You can accomplish this application-packaging approach by executing the application on top of a software product that gives each application its own virtual set of system resources — stuff like files and registry entries. Another

way to accomplish application packaging is by bundling the application and the virtualization software into a single executable program that is downloaded or installed; when the executable program is run, the application and the virtualization software cooperate and run in an isolated (or *sandboxed*) fashion, thereby separating the application from the underlying operating system.



Application packaging is a great way to isolate programs from one another and reduce virus transmission, but it doesn't solve the problem of end users installing nonpackaged software on client machines.

One thing to keep in mind with this approach is that it causes additional work as the IT folks prepare the application packages that are needed and then distribute them to client machines. And, of course, this approach does nothing to solve the problem of end users installing other software on the machine that bypasses the application packaging approach altogether. If you're loading a game onto your business laptop, you're hardly likely to go to IT and request that someone create a new application package so that you can run your game securely, are you?

Products that provide application packaging include SVS from Altiris, Thinstall's Virtualization Suite, and Microsoft's SoftGrid.

Application streaming

Application streaming solves the problem of how to keep client machines loaded with up-to-date software in a completely different fashion than application packaging. Because it's so difficult to keep the proper versions of applications installed on client machines, this approach avoids installing them altogether. Instead, it stores the proper versions of applications on servers in the data center, and when an end user wants to use a particular application, it's downloaded on the fly to the end user's machine, whereupon he or she uses it as though it were natively installed on the machine.

This approach to client virtualization can reduce the amount of IT work necessary to keep machines updated. Furthermore, it happens transparently to the end user because the updated application is automatically delivered to the end user, without any physical software installation on the client. It also has the virtue of possibly allowing client machines less capability to be deployed because less disk space is required to permanently store applications on the client hard drive. Furthermore, if this approach is taken to its logical conclusion and the client machine has no hard drive, it is possible that less memory is required because only the official IT applications can be executed on the machine. This result is because the end user can't execute any programs other than the ones available from the central server.

Although at first glance, this approach might seem like a useful form of virtualization, it is really appropriate only in certain circumstances — primarily situations in which end users have constant connectivity to enable application

downloads when required. Examples of these situations include call centers and office environments where workers rarely leave the premises to perform work duties. In today's increasingly mobile workforce world, these circumstances apply to a small percentage of the total workforce. Perhaps the best way to think about this form of virtualization is as one that can be very useful in a restricted number of work environments.

This type of virtualization is offered by AppStream's Virtual Image Distribution, Softricity's Softgrid for Desktops, and Citrix's Presentation Server. Softricity has recently been acquired by Microsoft, and its SoftGrid product will soon be available as part of the Windows Server platform. SoftGrid will offer the capability of streaming applications to remote desktops.



Application streaming is best suited for static work environments where people don't move around much, such as call centers and form-processing centers, although some organizations are exploring using it for remote employees who have consistent network connectivity to ensure that applications can be streamed as necessary.

Hardware emulation

Hardware emulation is a very well-established form of virtualization in which the virtualization software presents a software representation of the underlying hardware that an operating system would typically interact with. (I discuss hardware emulation in more detail in the "Server virtualization" section, later in this chapter.) This is a very common type of virtualization used in data centers as part of a strategy to get higher utilization from the expensive servers that reside in them.

Because of the spread of *commodity hardware* (that's to say, hardware based on Intel's x86 chip architecture; these chips power everything from basic desktop machines to huge servers), the same hardware emulation type of virtualization that can be used in data centers can also be used on client machines. (The term *commodity* refers to the fact that the huge volumes of x86 processors sold make them so ubiquitous and inexpensive that they're almost like any other mass-produced, unspecialized product — almost as common as the canned goods you can get in any grocery store.

In this form of client virtualization, the virtualization software is loaded onto a client machine that has a base operating system already loaded — typically Windows, but client hardware emulation virtualization is also available for systems running Mac and Linux operating systems.

After the hardware emulation software is loaded onto the machine, it's ready to support guest operating systems. Guest OSes are installed via the virtualization software; that is, rather than just sticking a CD into the machine's drive and rebooting it to install the operating system directly onto the hardware, you use the virtualization software's control panel to indicate your

desire to install a guest OS (which can be either Windows or Linux). It sets up the container (often called the *virtual machine*, or VM for short) for the guest operating system and then directs you to put the CD in the drive, whereupon the normal installation procedure occurs.

After the installation completes, you control the virtual machine (which is a normal Windows or Linux system) through the virtualization software's control panel. You can start, stop, suspend, and destroy a VM from the control panel.



Interacting with the VM guest OS is just like interacting with it if it were the only OS on the machine. A guest OS displays graphics on the screen, the VM responds to keyboard and mouse commands, and so on. That's why it's called virtualization!

Products offering this type of virtualization are VMware's VMware Server and Microsoft's Virtual Server. On the Macintosh, SWsoft's Parallels product provides hardware emulation virtualization.

Server virtualization

When discussing trends driving virtualization, you'll soon discover that most of the examples that come up are focused on issues of the data center — the server farms that contain vast arrays of machines dedicated to running enterprise applications, databases, and Web sites.

That's not an accident. Most of the action in the virtualization world right now focuses on server virtualization — no surprise, then, if you see me spending most of my time in this book on precisely that topic.

IT organizations are avidly pursuing virtualization to gain more control of their sprawling server farms. Although client virtualization is interesting to them, server virtualization is critical because many IT organizations are running out of room in their data centers. Their inability to add more machines means they can't respond to important business initiatives, meaning they can't deliver necessary resources so that the other parts of the business can implement the company's strategy. Obviously, this inability to provide IT resources is unacceptable, and many, many IT organizations are turning to server virtualization to solve this problem.

Three main types of server virtualization exist:

- ✓ **Operating system virtualization:** Often referred to as *containers*
- ✓ **Hardware emulation:** Similar to the same type of virtualization described in the client virtualization section, earlier in the chapter

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

of them raises the probability that you won't have a good outcome for your virtualization project. Fortunately, I address these topics in this book and provide tools for you to better find virtualization happiness.

In this book, I introduce the concept of the *virtualization life cycle*, the overall process of successfully introducing virtualization to your organization. Many people feel that all they need to do to be successful is buy some virtualization software and install it. That might get them started on the virtualization life cycle, but doesn't complete it. You'll be more successful if you address all the steps in the life cycle.

Okay, the whirlwind tour of virtualization has concluded. Time to move on to the real thing.

*image
not
available*

*image
not
available*

*image
not
available*

better economic sense. Approach this server utilization evaluation with tact in order to avoid conflict. See Chapters 6 and 7 for more information about how to manage the interpersonal aspects of a virtualization project.

Virtualization enables IT organizations to retire many of their low-performing servers and then migrate the applications to new hardware much better suited to supporting virtualization. (This is often referred to as *server consolidation* and is typically the initial step on the virtualization journey for IT organizations.) Although it's possible to use existing hardware to support virtualization, the ratio of retired to continuing servers is typically no more than four or five to one. Using the new generation of hardware, the ratios can be more like 10 or even 20 to one; in fact, the limiting factor of how many servers to host on a single physical server is likely to be your comfort level of having so many applications residing on a single box, rather than the pure processing capacity of the machine.

Chapter 8 goes into much more detail about selecting virtualization hardware and how to mitigate the single point of failure (SPOF) risk inherent in supporting so many virtual servers on a single piece of hardware.

A handy (if anonymous) real-world example

What kind of financial benefits are available through server consolidation? Here's one example, which sadly must remain anonymous as to the true identity of the company.

A large financial institution faced an unpalatable decision. To its current complement of branch-based servers, it needed to add encryption software to each one in order to comply with regulatory requirements regarding customer data privacy. To implement the software (an expensive item in its own right) at each branch, an additional server would need to be installed as well. Total cost: \$35 million.

Even for a large financial institution, \$35 million is more than chump change. Consequently, it decided to explore other options. It reviewed virtualization, and after due consideration, decided to migrate all the existing servers onto virtualized servers running in its data center, replacing the old servers with new, more powerful, virtualization-ready machines. It also decided to virtualize the data storage at the same time.

Because it was able to centralize its servers (and because they were so much more powerful), the financial institution no longer needed to purchase multiple copies of the encryption software, not to mention no longer needing to buy additional machines to run it.

*image
not
available*

*image
not
available*

*image
not
available*

bubble, if you prefer) of the Internet, has gone through tough times and emerged as a success. In fact, it's fulfilling its original mission of making art available to a new audience that previously never participated in the fine art market.

art.com's business has changed substantially as it has perfected its approach to the market. Part of its strategy has been to acquire other complementary companies in order to broaden the offerings of art.com.

And here is the IT operational flexibility challenge: As you might expect, changing your business offerings and merging companies means you face doing an IT infrastructure mashup. (Having lived through a couple of these situations, I assure you that it can be extremely challenging.)

The folks at art.com wanted to consolidate the acquired company's data center into art.com's facility. Ordinarily, that would be a protracted exercise in juggling hardware, doing backups, reinstalling systems, and so on. Instead, art.com took a different tack. First, the IT folks installed hardware at the art.com data center to support virtualized servers. Then, with that task out of the way, they used virtualization technology from VMware to migrate the systems from the acquired company's data center to art.com's new (and improved) data center.

The company now has taken its virtualization efforts a step further. Using VMware's Virtual Infrastructure 3 (VI3) product, art.com operates a pool of virtualized servers. When business initiatives require a new server instance, here's what happens:

1. IT instantiates an instance based on an existing image template.
In other words, the IT staff has preconfigured virtualized servers ready for immediate use.
2. IT installs the necessary application into the virtual server.
3. IT brings up the new virtual server by making it available through the network.

This three-step setup dramatically reduces the time necessary to get a new server up and running. Think about it: Instead of ordering a new server, waiting for it to arrive, installing it into a rack, getting it connected to the data center network, and then finally installing and configuring an application — a process that can take weeks — art.com can get a server up in just a few hours. (I heard art.com's CTO claim that his group could bring up a new server in minutes.) And here's the really striking thing: art.com doesn't keep track of where the new virtualized server goes. VI3 decides which physical server is best suited to support the new virtual machine, and VI3 decides where to put it. Instead of IT worrying about physical installation, it can focus on monitoring and tuning the running infrastructure.

*image
not
available*

*image
not
available*

*image
not
available*

By using storage virtualization, data backup is simplified. Most storage virtualization solutions come with a backup capability, enabling the same storage management system to handle the entire life cycle of data, from allocation to offline storage. By collapsing the number of tasks related to managing storage, storage virtualization reduces IT operations costs.

Virtualization Lowers Energy Costs

In Chapter 1, I discuss the cost of energy as being one of the four main motivations causing companies to implement virtualization.

Concern with the cost of power is a recent development for companies. The move to digitized business processes, both internal and Internet based, has been so dramatic that organizations often have used historic assumptions about energy use in their business planning. However, companies have implemented so many more systems recently that those historic assumptions are sadly out of date.

The lack of insight into just how much energy was costing companies has been exacerbated by the traditional split in how data centers are managed. Many companies assign the responsibility (and costs) for computer hardware like servers and network switches to IT, but assign the responsibility (and costs) for data center infrastructure services like power and air conditioning to . . . facilities. After all, a data center is a building, right? So facilities should manage it.

Unfortunately, that split in responsibility has meant that no single organization really kept track of how the growth in the number of servers was affecting other costs like power and air conditioning. As server sprawl has occurred, the consequence of running thousands of servers hasn't been clear.

That's all changed now. Skyrocketing energy costs have raised the issue of IT power costs well beyond the backwaters of facilities. Depending on how much of a company's business is digitized and runs across the Internet, energy might be one of the top-line items in terms of its overall cost structure. For many "Web 2.0" companies (Internet-powered companies like MySpace and Facebook), power is the second- or third-highest cost for the company overall.

Clearly, reducing overall energy consumption is important for companies, and the amount of energy use in running data centers is an obvious target. Beyond the bottom-line issue of energy cost, the rising concern with global warming is likely to make itself felt in increasing pressure upon companies to reduce energy consumption.

Virtualization is ideally suited for companies wanting to reduce energy use. Rather than powering thousands of machines, most of which are running at

*image
not
available*

*image
not
available*

*image
not
available*

quite reasonably, feel that position is unreasonable and that they should pay for only a two-processor license. In fact, Oracle's license pricing is even more complex because it charges a premium for customers running the software on multicore processors.

At the time of this writing, a spat was going on between VMware and Microsoft about the latter's virtualization licensing policies. VMware states that Microsoft's licensing inhibits customer flexibility in deploying Microsoft-based virtual machines. I expect this bickering to continue for the foreseeable future, because Microsoft is gearing up to compete with VMware when Microsoft Server virtualization is released. The details of the conflict keep changing as Microsoft adjusts its position, but the impact is clear: End users are confused about their licensing rights and responsibilities in a virtualized environment, and the situation appears to be driven more by vendor business needs rather than by customer satisfaction.

Whether or not you agree with these vendors' approaches to software licensing in a virtualized world, one thing is clear: When it comes to virtualization, software licensing is a mess.

What this means for you, as you embark on your virtualization journey, is that you need to carefully check the conditions of your software licenses if you decide to migrate physical systems to virtualized ones.



Take a look at the total number of licenses you have from your software vendors. If you're having trouble with a vendor about moving some of your licenses to a virtualized environment, it's amazing how much more supportive they end up being about your virtualization plans once you inform them how many licenses you actually have from them.

What might you expect to find as you interact with vendors regarding their virtualization licensing policies? You'll probably find that vendors who are leaders in their category will be the most unyielding with regard to pricing. After all, if you're the lead dog, you have the most to lose with any change in the market. Vendors that have lower market shares will be the most flexible about working with you in your virtualization efforts. Again, that makes sense: They're much less threatened by market changes and might even see embracing virtualization as a strategy to improve their market position.

Here's one other thing to remain aware of as you consider vendor policies with regard to virtualization: Although every vendor is aware of the momentum of virtualization, some of them are less enthusiastic about supporting their products running in a virtualized environment. Chapter 5 contains some information about evaluating your vendors' support policies, and it presents a few strategies for mitigating a lack of enthusiasm on their part.

*image
not
available*

*image
not
available*

*image
not
available*

- ✓ **System libraries** are the system resources that applications call upon to interact with the operating system, and through the operating system, to the underlying hardware.

So, OS virtualization provides a software emulation of a complete operating system environment. You can load applications into a virtual environment and run them as though they were on a completely different machine from another application that's running in a separate virtual environment on the same physical server. And each virtual container can be assigned its own IP address and have that mapped to a domain name (such as the `amazon.com` in `www.amazon.com`), which means that each virtual environment can be treated as a separate system.

OS virtualization does clever things to reduce the amount of processing necessary to provide containerization. It maps the virtual file system to the underlying operating system's file system very cleanly, so little processing overhead is necessary for applications to read and write to physical drives. Files that are used by multiple virtual environments are held in a single copy to reduce unnecessary duplication.

One benefit of OS virtualization is that it imposes the least overhead of any virtualization solution. Consequently, it achieves the highest performance of all the virtualization solutions (although other solutions, particularly paravirtualization, achieve near-native performance as well). Also, because OS virtualization imposes the least overhead of any virtualization approach, it can support the highest *density* (that is, the highest number) of virtual environments for a given piece of hardware. Many OS virtualization environments support dozens and even hundreds of containers on a single machine.

Of course, every story has two sides, and OS virtualization has some drawbacks as well. Although containers are great for isolating applications, each container reflects the configuration of the underlying base operating system. In particular, that means that each container must be the same type, version, and patch level of the base operating system. Furthermore, each container is limited to the device drivers present in the base OS, which limits hardware flexibility, because hardware can't be put in place for just one container that requires it. And, of course, the entire collection of containers is dependent on the base OS: If it crashes, all of the containers will go down. (Please note that operating system virtualization providers are working hard to address the version and patch level limitations and may be able to address this issue in the future; for today, however, the limitations are present and inflexible.)

Consequently, OS virtualization supports only one operating system environment, which means that you can't mix guest operating systems. Depending on your desired virtualization architecture, this limitation might present a problem. On the other hand, many organizations find that they end up segregating their virtualized operating systems: Some virtualization servers run only Windows guests, and others run only Linux guests. So the single OS limitation of OS virtualization isn't a fatal drawback for these organizations.

*image
not
available*

*image
not
available*

*image
not
available*

and coordinates access to the underlying hardware directly. As you might also suspect, this approach provides much higher performance than the version of hardware emulation described in the previous paragraph.



Hardware emulation is the type of virtualization implemented by VMware, the category leader in x86 virtualization. (Microsoft's Virtual Server also implements hardware emulation.)

VMware provides products that implement both approaches to hardware emulation. VMware's VMware Server is installed atop an existing operating system (which can be Windows or one of several flavors of Linux), whereas its ESX Server implements bare-metal hardware emulation. Although VMware Server's approach to hardware emulation — the Install-it-on-top-of-the-OS path — imposes a significant performance hit, it can be usefully applied where performance isn't critical. For example, lightly loaded servers can usefully apply this approach to hardware emulation. It's also useful for virtualizing client machines like a desktop or notebook PC.

Hardware emulation does have some drawbacks:

- ✓ **It imposes a performance penalty, even for the bare-metal approach.** This problem occurs because translation is still going on, eating up some machine cycles as the hypervisor executes the translation. VMware does clever optimizations on this translation, caching data that is repeatedly accessed, and so on, but some performance penalty still exists. On the other hand, if you're moving from a Pentium III-based server that's running at only 15 percent utilization to a virtualized dual-core machine that offers 100 times the performance of the old machine and can thereby support six virtual machines, you might not be too worried that you're not getting 100 percent of the potential performance of the new server. You'd be happy just taking advantage of what's available from the new machine.
- ✓ **The idealized image of the hardware provided offers idealized device drivers.** The hypervisor contains real device drivers that talk to the underlying hardware. In other words, the hypervisor drives the underlying hardware, and as an end user, you must depend on the hypervisor to include the appropriate drivers for your hardware. Consequently, you're somewhat limited in your choice of hardware because you need the hypervisor to include drivers to interact with your hardware. Although VMware does a very good job of trying to keep up with new hardware releases and their accompanying device drivers, you might find your hardware is unsupported. So checking the hardware compatibility list for ESX Server support is critical if you go down this path. Note that VMware Server, while imposing a performance penalty, actually suffers less from the device driver issue because it is able to leverage the device drivers present in the underlying OS.

*image
not
available*

*image
not
available*

*image
not
available*

Table 3-1 Virtualization Technologies		
<i>Virtualization Type</i>	<i>Comments</i>	<i>Products</i>
Operating System Virtualization	Provides an abstracted view of OS resources (root file system, process table, and so on); useful for homogeneous operating environments where consistent OS version is desirable (such as Web hosting); very scalable.	SWsoft OpenVZ (open source) and Virtuozzo (commercial); Sun OpenSolaris (open source) and Solaris (commercial)
Hardware Emulation	Provides an abstracted view of underlying hardware (that is, the emulated motherboard); based on binary translation of runtime guest OSes; offers support for guest OSes that might be heterogeneous; market-leading approach to virtualization	Microsoft Virtual Server (commercial); VMware VM server (commercial, but free); VMware ESX Server (commercial, leading virtualization product in market)
Paravirtualization	Light-weight virtualization approach multiplexes access to underlying hardware resources; high performance; requires guest OS modification prior to deployment, therefore somewhat more complex than hardware emulation; native guest OS support on latest generation of processor chips; available bundled with commercial operating systems	Xen (open source) available or free download; Xen also bundled with Linux distributions from Red Hat and Novell as well as other community distributions; XenSource (sponsor of Xen open source product) offers a commercial version of Xen with additional functionality; upcoming Microsoft Server product will include paravirtualization

This table will help you decide which virtualization technology is best suited for your environment; of course, your ultimate decision will be affected by the use you'll make of virtualization. If you look at the next chapter, you'll find a discussion of all the different ways virtualization can be applied to provide a more efficient and cost-effective computing infrastructure.

*image
not
available*

*image
not
available*

*image
not
available*



Even a developer's laptop is plenty powerful enough to support several guest machines. By using virtualization, a developer or tester can replicate a distributed environment containing several machines on a single piece of hardware. This setup negates the need to have a bunch of servers sitting around for the occasional use of developers or testers. It also avoids the inevitable conflict that occurs when an organization attempts to have its members share machines for distributed use. Just when one person wants to test some distributed functionality, someone else is already using the machines. Even worse, the machines are incorrectly configured, so the first thing a new user has to do is remove all the software on the machine and reinstall new software to get the right configuration for his or her purposes. As you can imagine, this repetitive teardown and reinstallation is a real drag on productivity.

Virtualization is also useful in testing and development in another way. One of the side effects of exercising software is that early versions often crash and damage not only the application, but also the underlying operating system as well as other applications in the software stack. To recover, it's necessary to reinstall all the software. Again, this is a real drag on productivity.

One of the great things about virtualization is that guest machines can be captured in a file. That is, an image of the complete virtualized machine — operating system, software stack, and applications — can be saved to a file. This makes sense when you consider that the virtualization software is running software virtual machines, so it has to have a standard format for the machines. Virtualization software can save these images onto a disk file and then load the file at a later time, ready to run.

This capability is a boon to developers and testers. Rather than having to repeatedly rebuild test instances, they can just save a complete virtual machine image and load it each time they trash a virtual machine instance. In terms of recovery time, the consequences of a machine crash go from hours to minutes (or even seconds).

This ability to save and restore virtual machine instances is an incredibly useful aspect of virtualization. The usefulness of virtualization for development was brought home to me in a recent conversation with a friend who works in development for a chip company. He noted that he was using virtualization on his laptop and running development applications in a virtual machine. The inevitable happened (that's why it's inevitable, right?), and the base operating system got corrupted, necessitating a complete OS reinstall. Dreading the extended reinstallation of numerous applications, he decided to take a chance on reloading the virtual machine image from a backup: two minutes later he was up and running with his entire tool set and data files available. He estimated he saved at least six hours of boring software loading. This is why virtualization is being widely adopted in development and testing.

*image
not
available*

*image
not
available*

*image
not
available*

The outcome is much more dramatic than this simple list would seem to indicate. Companies implementing server consolidation often move from running 150 physical servers to running 150 virtual machines on only 15 servers, with an understandable reduction in power, hardware investment, and employee time.

Naturally, the process of migrating systems requires many smaller tasks that aren't listed here. For example, each virtual machine must have a network address assigned and storage allotted (or have the existing virtualized storage the physical server is using assigned to the new virtual machine). The backup process for the new setup must be modified when it is no longer a question of just sticking a tape into a lot of separate servers. And, of course, the virtual machines must still be managed after the server consolidation project is complete.

Conceptually, however, the list captures the steps. And, in the overall scheme of things, migrating 150 servers to new virtual machines is much, much easier than migrating 150 systems to new physical servers — now that's a lot of work!



Many organizations are surprised at how pain-free the implementation of a server consolidation project can actually be. You can find more information about migration and virtualization management in Chapters 9 and 10, as well as more about the virtualization life cycle and all the elements of a virtualization project.

If you're in the market for a virtualization product for a server consolidation project, check out the following (commercial and open source) options:

- ✓ **SWsoft Virtuozzo:** This is a commercial operating system virtualization (also called *container*) product that is very well suited to support large numbers of virtualized systems. It's especially suited for consolidating environments that have large numbers of identical systems. For example, if an organization wants to centralize an environment in which a large number of branch offices all have their own on-premises servers, each of which is identically configured, then container virtualization is ideal. Unfortunately, this kind of homogeneous environment is pretty rare. SWsoft also sponsors the free open source container virtualization product OpenVZ, which you can use for the same server consolidation purposes.
- ✓ **Sun Solaris containers:** Similar to SWsoft's Virtuozzo, you can use Sun Solaris containers for server consolidation projects with a significant degree of homogeneity. Solaris is a commercial product that is also available in open source form as OpenSolaris. Sun also sells physical servers that are architected to support containers with better isolation between individual container instances.
- ✓ **VMware ESX Server:** The most widely deployed commercial server consolidation virtualization product today, VMware ESX Server is very scalable and therefore able to support high numbers of virtual machines per

*image
not
available*

*image
not
available*

*image
not
available*

- ✓ The hardware representation of the virtual machine, including what resources are available to it. These include networking and storage adapters (that is, cards).
- ✓ External connections that the virtual machine uses, including IP addresses and storage locations and identifiers.
- ✓ The amount and settings of all memory the virtual machine is using. (Remember, this memory is used by the virtual machine but is actually controlled by the hypervisor.)

Every virtual machine, even one that has just been started and, as yet, has not done any work on behalf of a user, has a state. As applications execute on behalf of users, the state changes constantly. As you can see, failover, while simple conceptually, can be quite difficult to execute properly.

Failover: The simple case

The role of the hypervisor is to start, run, suspend, restore, and delete virtual machines. Part of doing that work is to constantly monitor the state of each virtual machine and react to requests by the virtual machine for hardware resources.

Because the hypervisor is constantly monitoring each virtual machine's status, it's relatively straightforward to configure the hypervisor to start a new instance of a virtual machine should it notice a previously running virtual machine is no longer present.

Because all the hypervisor has to do is start a new virtual machine based on the machine's image, the outage duration of a machine might be mere seconds. Obviously, this is a huge improvement over the minutes-to-days durations typical of nonvirtualized system restores.

However, this approach has a shortcoming. Although simple failover can restart a virtual machine, it can't restore the state of the original virtual machine's memory at the time of the crash. Instead, it brings up a new VM, ready to begin work. Any work in process at the time of the VM crash is lost.

Although losing memory state might seem like a fatal flaw, don't sell simple failover short. Just being able to recover systems in seconds instead of hours is an enormous improvement for most IT organizations because very few of them have any kind of failover mechanism in place today. For them, simple failover represents a huge improvement over the status quo.

A more serious issue is that simple failover deals with a virtual machine crash but does nothing to protect against hardware failure. In other words, if the virtual machine continues to operate but the underlying virtualization hardware fails, suddenly you have no hardware for the simple failover to execute upon. You can, however, address this hardware issue, as I discuss in the following section.

*image
not
available*

*image
not
available*

*image
not
available*

Because the hypervisor needs to keep track of every virtual machine's memory, it would seem possible for the memory state of each machine to be automatically transferred to somewhere off the physical server the hypervisor resides on, ready to be re-created in a new virtual machine should the original go down.

There are some developments in that direction. Today the effort is mostly focused on Xen, because its open source characteristics make it possible to add additional software to track and communicate virtual machine memory state. I expect that there will be much more of this in the future, as virtualization is extended to enable the highest possible availability. To put it another way: It's a future development, but expect it in the *near* future.

Server pooling

If you've read through all the other descriptions of how virtualization can be applied, you're probably thinking "Wow, virtualization sure can be applied in a lot of useful ways, but there seems to be a lot of installing and configuring in those descriptions. Wouldn't it be great if the virtualization software was arranged for the installation and configuration so that I automatically got failover and load balancing?"

And in fact, that functionality exists. It's called *server pooling*, and it's a great application of virtualization.

With server pooling, your virtualization software manages a group (or pool) of virtualized servers. Instead of installing a virtual machine on a particular server, you merely point the virtualization software at the virtual machine image, and the software figures out which physical server is best suited to run the machine.

The server-pooling software also keeps track of every virtual machine and server to determine how resources are being allocated. If a virtual machine needs to be relocated to better use the available resources, the virtualization software automatically migrates it to a better-suited server.

You manage the pool through a management console, and if you notice that the overall pool of servers is getting near to your defined maximum utilization rate, you can transparently add another server into the pool. The virtualization software then rebalances the loads to make the most effective use of all the server resources.



Because you have no way of knowing which physical server will be running a virtual machine, your storage must be virtualized so that a VM on any server can get access to its data.

*image
not
available*

*image
not
available*

*image
not
available*

Table 3-2 (continued)		
Application	Product	Comment
Failover	XenSource	Commercial paravirtualization product provides good scalability and focus on supporting Windows guests.
	Virtuallron	Commercial virtualization product using Xen technology; very scalable.
	SWsoft	Commercial Virtuozzo product provides failover functionality.
	Sun	Solaris operating system virtualization can be configured to provide failover.
	VMware	ESX Server can be configured to automatically restart crashed virtual machine.
	Xen	Open source paravirtualization product contained in many Linux distributions can be configured to automatically restart crashed guest operating systems.
	XenSource	Commercial paravirtualization products can be configured to automatically restart crashed guest operating systems.
	Virtuallron	Commercial virtualization product can be configured to automatically restart crashed guest operating systems.
	VMware ESX Server	Virtual Infrastructure 3 (VI3) provides sophisticated management services for collections of virtual machines.
	Virtuallron	Commercial product incorporating Xen hypervisor provides high availability functionality.
Load Balancing	VMware ESX Server	VI3 provides load-balancing capability.
	Virtuallron	Product provides load-balancing capability.

*image
not
available*

*image
not
available*

*image
not
available*

Virtualization Gets Integrated into Operating Systems

Really, you don't have to be much of a fortune teller to predict that virtualization will become a regular feature of operating systems, just like hundreds of other features that come pre-integrated into the core operating system.

In fact, you don't have to be a fortune teller at all. Both Red Hat and Novell include Xen virtualization as part of their core operating systems. And Microsoft is scheduled to deliver virtualization as part of its upcoming Microsoft Server 2008 release.

The inclusion of virtualization in operating systems reflects the OS providers' recognition that virtualization is so critical a technology that it needs to be a fundamental part of the operating system. Of course, you might take a more cynical view: The OS vendors recognize that the virtualization market is too important to cede to competitors like VMware. More importantly, OS vendors don't want to lose their primary role in the data center. If end users focus their attention on how well their applications integrate with software other than the underlying OS, that means OS vendors no longer are the key supplier to those end users. This is bad from their perspective.



So, you can expect to see increasing virtualization integration from OS providers in their flagship products. In the future, look for this integration to be extended even further into areas like virtualization management, storage, and server pooling.

Of course, such integration will present a challenge to independent vendors of virtualization products, particularly the leading vendor, VMware. Currently, it has significantly better functionality compared with the virtualization functionality present in operating systems, but that advantage will undoubtedly be reduced as OS vendors improve their products.

The cliché in the virtualization world is that basic virtualization functionality will become a commodity — just another assembly-line product — and that future differentiation in virtualization products will be at the management level; in other words, what will distinguish virtualization products will be how easy they are to run on a day-to-day basis. VMware offers significant capabilities in this area and is continually improving this dimension of its products. It remains to be seen, however, how successfully it can compete against virtualization capabilities that are present as a default (and free) feature in the underlying OS.

Don't count VMware out by any means, though. Although many OS features started out as separate products (for example, TCP/IP stacks were once a separate product but now are included by default in every OS), other

*image
not
available*

*image
not
available*

*image
not
available*



Unfortunately, virtual appliances aren't quite as simple as that last sentence. You'll still need to perform some minimal configuration in terms of assigning an IP address and making the rest of your machines aware of the new system. Overall, though, this kind of configuration is like a walk in the park compared to the traditional drudgery of installing new software.

Software vendors love this approach, because it means

- ✓ **Fewer support calls:** Even with the best of intentions, users make mistakes during installation and configuration, resulting in support calls to the vendor. Cutting installation and configuration out of the picture means fewer user mistakes and fewer support calls, which translates into lower support costs.
- ✓ **Easier problem debugging:** If a support call comes in, the vendor can dispense with much of the usual back-and-forth required to confirm that the product is installed properly (though, of course, they'll still need to confirm the classic reason a system isn't working properly — it isn't plugged in). This will enable easier and quicker problem resolution, and thereby lower support costs.
- ✓ **Happier customers:** Encountering problems during application installation is frustrating, particularly when the installation is performed under a deadline. Making application installation simpler and faster makes for more satisfied customers, which results in quicker application adoption and additional sales.

End users will love virtual appliances as well, for the following reasons:

- ✓ **Virtual appliances align with the move to virtualization:** If you have a virtualized infrastructure, it's natural to run new applications in virtual machines.
- ✓ **Virtual appliances reduce the manual effort involved in bringing up new applications to a minimum:** The vendor, rather than IT staff, does all the configuration, shifting the work from user to vendor. Users are bound to like that equation.
- ✓ **Virtual appliances reduce the chance for error:** There are so many steps required to provision a new application that one or more things are bound to go wrong with the process, particularly when the work is, inevitably, being done in a hectic, interruption-filled, work environment. Avoiding the need for this work means far fewer problems.
- ✓ **Virtual appliances get applications up and running much faster:** The speed of provisioning a virtual appliance versus installing and configuring applications by hand means that applications are available for use much more quickly. Because the whole point of the exercise is to make applications available for real work, using virtual appliances speeds availability.

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

image
not
available

image
not
available

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

*image
not
available*

installation (*continued*)

- of virtual machine on Fedora 7, 278–287
- of virtual machine on VMware Server, 253–258
- virtualization implementation options, 61
- of virtualization software on server, automated, 194–196, 203–207, 327
- of virtualization software on server, manual, 153, 194, 202–203, 327
- of VMware Server, 248–251

Integrated Drive Electronics (IDE), 225

Intel

- multicore processors, 46
- Preboot Execution Environment (PXE), 196
- VT (virtualization-enabled) chip, 28, 59, 118, 166, 183–184

Internet

- SaaS (Software as a Service), 53, 90–91
- storage requirements increased by use of, 13
- virtualization diffusion into, 90–93

Internet resources

- blogs about virtualization, 337
- DeepBurner image burner, 269
- Fedora 7 download, 268
- Fedora image files, 262
- Google's Docs and Spreadsheets, 53
- InfoWorld virtualization newsletter, 336–337
- Invirtus, 204
- Leostream, 204
- PG&E rebate program, 109, 127
- PlateSpin, 147, 204
- P2V vendors, 204
- SearchServerVirtualization site, 336
- storage virtualization information, 339
- vendor information, 339
- virtualization information, 336–339
- virtualized pricing information, 97
- VMware Capacity Planner, 147
- VMware Server appliances, 89, 264
- VMware Server download, 248
- VMware user group listings, 338
- XenExpress download, 292

Intrepid management product (Levanta), 196

Invirtus migration tools, 204

IP addresses

- for remote data centers, 93
- for virtual containers, 52
- VMware Server options, 255–256
- for XenServer, 299

iSCSI

- with ESX Server, 170, 171
- network traffic increased by, 31
- overview, 171, 228
- preparing for migration, 198
- SCSI used for data communication, 228
- standard corporate network used by, 30
- standard IP Protocols used for, 171
- subnetting for, 171

IT. *See also* system administration

- future of, 1–2, 93–95
- higher-level management with virtualization, 94–95
- job dissatisfaction in, 42
- operational flexibility increased by virtualization, 37–41
- operations costs reduced by virtualization, 41–44
- proprietary attitude toward servers by, 35–36, 137–138
- system administration costs as virtualization driver, 15–16, 34, 101
- work reduced by application streaming, 18

• J •

JBOD (just a bunch of drives), 169, 229

• K •

Krug, Kris (*BitTorrent For Dummies*), 262

• L •

Leostream migration tools, 204

Levanta's Intrepid management product, 196

*image
not
available*

managing the virtualized infrastructure

(continued)

- creating your own solution, 156–157
- current approach's impact on, 218–219, 220–221
- current hardware mix issues, 219–220
- in departmental implementation, 217–218, 220–221
- as equal member of data center, 213, 215–216, 221
- expecting change, 331–332
- in experimentation phase, 217, 220, 221
- with free virtualization, 212–213
- importance of, 154, 327
- incompatibility issues, 216
- life cycle phase, 133
- multiple tools for, 156
- open source software for, 157
- philosophies of, 213
- in pilot implementation, 217–218, 220–221
- planning for, 133
- postponing virtualization, 156
- in production environment, 218, 221
- recommendations, 220–221
- solutions for, 155–157
- as specialized resource, 213, 214–215, 220–221
- storage administration, 229–230
- table summarizing recommendations, 220–221
- toolset choices for, 136, 154–155
- virtualization journey steps, 213, 217–218, 220–221
- virtualization software tools for, [84](#), 155

manual migration, 153, 194, 202–203, 327

memory

- amount per virtual machine, 173
- in blade servers, 179
- for Fedora [7](#) virtual machine, 280
- future technology, 184
- as key server resource, 172–174, 175
- requirements for virtualization, 172–174
- 32-bit server issues, 177
- for XenExpress virtual machine, 302

Microsoft. *See also* Microsoft Server virtualization; Microsoft Windows

- blog site on virtualization, 337

dispute with VMware, [48](#)

hardware emulation offered by, [20](#), 26

pricing model for virtualization, 96

SoftGrid, [18](#), [19](#)

Softricity acquired by, [19](#)

virtual appliances from, 89

Virtual Server, [20](#), 26, 65, 66

Windows Server, [19](#)

Microsoft Server virtualization

- child partition as guest OS, 57

- memory requirements, 173

- paravirtualization approach with, 28, 57, 61

- pricing model for, 96

- root partition as privileged guest, 57

- Virtual Server superceded by, 26

- virtualization integration into OS, [84](#)

Microsoft Windows

- file structure, 51

- installing XP SP2 as guest OS on XenExpress, 300–311

- malware attacks on, [17](#)

- paravirtualization for, 28

- workgroups in organization structure, 138–139

migrating to virtualization. *See also*

- managing a virtualization project

- automated software installation for, 194–196, 203–207

- backing up applications and data, 193

- challenges of, 31–32, 189–190, 200–201

- cleaning up the source system, 192–193

- cloning for, avoiding, 201

- custom management products for, 195–196

- database options, 153

- file-based data options, 153

- free tools for, 205

- hardware inconsistencies, 201

- installing infrastructure, 197–200

- with Linux systems, 202–203

- manual software installation for, 153, 194, 202–203, 327

- migrating physical servers (P2V), 191, 200–207

- moving to production (cutover), 191, 207–208
 - options for, 152–153
 - phases overview, 190–191
 - planning for, 133
 - planning needed for, 152, 153
 - preparing the virtualized environment, 190–200
 - preparing virtualization servers, 193–196
 - production environment, 133, 152–153, 326
 - removing hardware-dependent software, 192
 - removing unused and/or obsolete software, 192
 - security issues, 198–199
 - storage choices impacting, 240
 - system management tools for, 195, 203
 - testing the infrastructure, 200
 - third-party tools for, 152–153, 203
 - time constraints and pressures during, 152, 153
 - updating current OS and applications, 193
 - vendor-specific tools for, 152, 203
 - in virtualization life cycle, 133
 - migration. *See also* migrating to virtualization
 - between hypervisors, 23
 - of pilot systems, 149
 - between servers, 42
 - software for applications, 38
 - mirrored systems, 71
 - mission-critical applications, shared storage for, 235
 - Moore, Gordon (Intel cofounder), 12, 13
 - Moore's Law, 12–13, 35, 164
 - multicore processors, 46, 48, 165–166
 - multimachine architectures, 109
 - multiplexing with paravirtualization, 26, 57
 - multisystem configurations with VMware Server, 247
- **N** •
- NAS (Network-Attached Storage). *See also* storage virtualization
 - cost-effectiveness of, 30
 - DAS versus, 169–170
 - defined, 224
 - drawbacks of, 30
 - Host Bus Adapter (HBA) for, 170
 - identifying virtualization project costs, 118
 - migration between servers with, 42
 - network traffic increased by, 30
 - overview, 29–30, 226–227
 - scalability and reliability of, 227–228
 - as shared storage, 226–227
 - types of implementations, 169–170
 - NAT (Network Address Translation), 256
 - net present value (NPV), 123
 - NetSuite's Enterprise Resource Planning (ERP) application, 53
 - Network Time Protocol (NTP), 295
 - networking
 - Fedora 7 options, 284
 - infrastructure preparation for smigration, 198
 - iSCSI-based SAN increasing traffic over, 31
 - NAS increasing traffic over, 30
 - operating system use of, 51
 - shared storage on networks, 226
 - trend like Moore's Law for, 13
 - VMware Server options, 255–256
 - XenExpress setup for, 295
 - New Energy data center, 181
 - NFS (Network File System), 226
 - NICs (Network Interface Cards)
 - defining for Fedora 7, 270
 - for ESX Server implementations, 168
 - future technology, 184, 186
 - as key server resources, 167–168, 174–175
 - not used by SANs, 30
 - performance considerations, 167
 - preinstalled on servers, 167
 - redundant, advantages of, 168, 174–175, 230
 - 32-bit server issues, 176
 - Novell Linux. *See* Linux
 - NPV (net present value), 123
 - NTP (Network Time Protocol), 295

• 0 •

- “one application, one server” approach
 - as dominant data center reality, 38
 - for isolation of applications, 162
 - limitations of, 37–38
 - processors underutilized with, 165
 - reasons for historical practice, 160–163
 - resource contention avoided by, 161
 - simplicity of, 160
 - territoriality with, 137, 161–162
- open source software. *See also specific products*
 - for development and testing environments, 65
 - free virtualization products, 61
 - for server consolidation, 68
 - for training environments, 67
 - for virtualization management, 157
- OpenSolaris (Sun), 61, 65, 67, 68
- OpenView software manager (HP), 195
- OpenVZ (SWsoft), 23, 61, 67
- operating system crashes. *See also* outage recovery
 - in development and testing environments, 63, 64
 - OS virtualization dangers for, 52
- operating system virtualization (containers). *See also* guest operating systems; migrating to virtualization
 - applications contained in virtual OS, 21
 - benefits of, 22, 52
 - density of containers with, 52
 - device driver issues, 52
 - ease of installation, 194
 - environments suitable for, 23, 53–54
 - free products, 61
 - hardware emulation versus, 21, 22, 51, 60
 - illustrated, 22
 - IP addresses for containers, 52
 - licensing requirements less with, 22
 - limitations of, 23, 52–53
 - multisystem configurations with VMware Server, 247
 - OS choice limited with, 23, 52–53
 - overview, 21–23, 51–54, 129
 - paravirtualization versus, 51, 60
 - patches propagated to all containers, 53
 - performance benefits, 22, 52
 - products providing, 23, 53, 60, 61
 - storage options, 236–237
 - terminology, 51–52
 - uses for, 21–22, 53–54
 - virtual OSES with, 21
 - VMware products, 209
- operating systems. *See also* guest operating systems
 - choices limited with OS virtualization, 23, 52–53
 - described, 50
 - integration of virtualization with, 84–85, 93–94, 210–211
 - virtualization functionality with, 61
- Oracle, licensing approach of, 47–48
- organizational structure
 - evaluating skill sets and experience levels, 139
 - need for understanding, 137, 332–333
 - “one application, one server” approach, 137, 161–162
 - reviewing before technical decisions, 323–324
 - territoriality (hands off “my” server attitude), 137–138
 - Windows versus Linux workgroups, 138–139
- “orphan” applications in data centers, 35
- outage recovery. *See also* DR (disaster recovery); SPOF (single point of failure)
 - failover for, 39, 69–72
 - hardware requirements for redundancy, 141
 - hardware risk exposure with virtualization, 163–164
 - high availability for, 73
 - hypervisor role in, 69–70, 72
 - load balancing for, 74
 - products providing, 75
 - server pooling, 39
 - usual steps for, 38–39
 - with virtualization, 39
- overtime costs, 41–42

• **p** •

- Pacific Gas & Electric (PG&E), 15, 108–109, 127
- Parallels (SWsoft), 20
- paravirtualization. *See also* migrating to virtualization
 - benefits of, 27, 58
 - described, 10, 21, 26–27, 129
 - drawback of, 28, 59
 - with Fedora 7, 282
 - free products, 61
 - full virtualization versus, 282, 311–312
 - guest OS modification needed for, 28, 57, 59
 - hardware emulation versus, 27, 60, 129
 - illustrated, 27
 - multiplexing by, 26, 57
 - native device drivers with, 27, 58
 - OS virtualization versus, 51, 60
 - overview, 26–28, 57–59
 - “para” in, 27
 - performance benefits, 27, 52, 58
 - privileged guest with, 57
 - products providing, 28, 57, 60, 61
 - shared memory used by, 27, 58
 - storage options, 238–239
 - stub drivers with, 58
 - virtualization-enabled chips for, 28
 - with XenExpress, 311–314
- patches. *See* maintenance and upgrades
- performance
 - application requirements and suitability for virtualization, 107–108
 - capacity planning issues, 165
 - DAS issues, 169
 - hardware emulation drawbacks, 25, 56
 - memory’s impact on, 172–174, 175
 - NIC considerations, 167
 - OS virtualization benefits, 22, 52
 - paravirtualization benefits, 27, 52, 58
 - of pilot implementation, assessing, 150–151
 - QEMU penalty, 312
 - VMware Server issues, 246, 247
 - PG&E (Pacific Gas & Electric), 15, 108–109, 127
 - physical to virtual. *See* migrating to virtualization; P2V
 - pilot implementation
 - agile systems development versus, 148
 - assessing performance, 150–151
 - bounded timeframe for, 150
 - costs saved by, 148
 - DAS in, 232
 - debriefing session after, 151
 - defined, 148
 - gaining management support for, 148
 - importance of, 326
 - installing the mini-production environment, 149
 - managing, 149–150, 217–218, 220–221
 - migrating systems, 149
 - moving into production, avoiding, 150
 - performing, 149–150
 - selecting subset of production environment for, 149
 - trying more than one virtualization product, 145
 - planning a virtualization project
 - architecture selection, 134, 139–142, 324, 332
 - confirming assumptions and conclusions, 143
 - costs of poor planning, 134–135
 - expecting change, 331–332
 - high-level tasks, 132, 134
 - importance of, 321–322
 - life cycle phase, 131–132
 - lobbying and politics, 132
 - organizational structure evaluation, 132, 134, 137–139, 323–324
 - reviewing the plan, 143, 332
 - use cases identification, 132, 134–137, 322–323
 - PlateSpin
 - migration-analysis tool, 152–153, 204
 - PowerRecon planning tool, 147
 - Web site, 204
 - power costs. *See* energy costs
 - PowerRecon planning tool (PlateSpin), 147

- Preboot Execution Environment (PXE), 196
- Presentation Server (Citrix), 19
- price elasticity, 212
- process tables, defined, 51
- processors. *See also* x86 chips
 - assigning to XenExpress virtual machine, 302
 - in blade servers, 179
 - commodity hardware with x86 chips, 19, 50, 159
 - hardware selection issues, 166
 - as key server resources, 165–166, 174
 - licensing practices related to, 46, 47–48
 - Moore’s Law for, 12–13, 35, 164
 - multicore, 46, 48, 165–166
 - multiple, for avoiding SPOF, 166, 174
 - price/performance curve, 35
 - 64-bit servers, 175, 178–180, 181
 - software emulation of x86 chips, 24
 - 32-bit servers, 175, 176–178, 181
 - underutilization as virtualization driver, 12–13, 33–37
 - underutilized in “one application, one server” approach, 165
 - virtualization-enabled, 28, 59, 118, 166, 183–184
- production environment
 - implementing, 326
 - migrating to virtualization, 133, 152–153
 - selecting subset for pilot implementation, 149
 - virtualization management in, 218, 221
- products. *See* virtualization software products
- project management. *See* managing a virtualization project
- proprietary attitude toward servers, 35–36, 137–138
- prototyping, DAS for, 232
- P2V (physical to virtual). *See also* migrating to virtualization
 - automated software installation for, 194–196, 203–207, 327
 - challenges of, 200–201
 - cloning for, avoiding, 201
 - custom management products for, 195–196
 - database options, 153
 - defined, 190
 - file-based data options, 153
 - hardware inconsistencies, 201
 - with Linux systems, 202–203
 - manual software installation for, 153, 194, 202–203, 327
 - options for, 152–153
 - planning needed for, 152, 153
 - source system, 192
 - system management tools for, 195, 203
 - target system, 192
 - third-party tools for, 152–153, 203
 - time constraints and pressures during, 152, 153
 - vendor-specific tools for, 152, 203
- P2V wizard (XenSource), 204
- PXE (Preboot Execution Environment), 196

• Q •

- QEMU technology for virtualization-enabled chips, 311–312
- quality assurance, hardware emulation for, 25

• R •

- RAID (redundant array of inexpensive disks), 229, 230
- RAM. *See* memory
- RDP (Remote Desktop Protocol), 314–317
- Red Hat Enterprise Linux (RHEL), 267, 291
- Red Hat Fedora Linux. *See* Fedora virtualization; Linux
- redundancy. *See also* capacity planning
 - for avoiding SPOF risk, 166, 168, 174
 - hardware requirements for, 141
 - multiple processors, 166, 174
 - of NICs, 168, 174–175, 230
 - with RAID, 229, 230
 - in storage infrastructure, 230

redundant array of inexpensive disks
(RAID), 229, 230
Remember icon, [5](#)
Remote Desktop Protocol (RDP), 314–317
resource contention, 161, 246
RHEL (Red Hat Enterprise Linux), 267, 291
root file system, defined, 51

• S •

SaaS (Software as a Service), 53, 90–91, 92–93
Salesforce.com's CRM application, 53, 90
SAN (Storage Area Network). *See also* storage virtualization
costs of, 31
DAS versus, 169–170
defined, 224
Fibre Channel versus iSCSI, 31
Host Bus Adapter (HBA) for, 30, 170, 227
identifying virtualization project
costs, 118
migration between servers with, 42
network protocols for, 30
overview, 30–31, 227–228
as shared storage, 227–228
switch fabric, 227
types of implementations, 169–170
sandboxing, [18](#)
SATA (Serial Advanced Technology Attachment), 225
scalability, 175, 227–228, 233
SCSI (Small Computer System Interface), 225, 227, 228, 229
SearchServerVirtualization site, 336
seasonal demand variations, 37, 39, 106
security, 198–199, 275
SELinux, 275
Serial Advanced Technology Attachment (SATA), 225
server consolidation
backup process with, [68](#)
consolidation defined, 67
costs reduced by, 36–37, 43
defined, 31, 67, 130

as first step in virtualization, 31, [36](#), 39, 67
hardware emulation for, 25
hardware risk exposure with, 163–164
overview, 67–69
products providing, 68–69, 79–80
proprietary attitude toward servers, 35–36
ratio of retired to continuing servers, [36](#)
real-world example, 36–37
typical project tasks, 67–68
underutilization rate for servers, 35
server pooling. *See also* load balancing
Art.com example, 39–41
defined, 39, 130
integration into OSes, [84](#)
outage recovery with, 39
overview, 76–77
products providing, 77, 80–81
seasonal demand variations eased by, 39
shared storage for, 236
SPOF avoided by, 164
storage virtualization required for, [76](#)
in times of company restructuring, 39
server virtualization. *See also* hardware emulation; migrating to virtualization; operating system virtualization (containers); paravirtualization
applications not suitable for, 107–109
benefits, 101
checking prerequisite conditions, 103–106
company characteristics important for, 105–106
deciding about using, 101–110
described, [10](#)
determining financial benefits, 103
installing software manually, 153, 194, 202–203, 327
installing software using automation, 194–196, 203–207, 327
investment in virtualization, 104–105, 110
as main virtualization arena, [20](#), 31
organizing physical servers, 140–141
overview, 20–28
power company support for, 108–109
resistance by employees, 105

- server virtualization (*continued*)
 - situations not suitable for, 106–110, 330–331
 - storage type impacting hardware choices, 171–172
 - table comparing types of, [60](#)
 - training requirements, 104–105, 330
 - types of, 21–28, 49, 59–60
 - vendor support issues, 104, 107, 333
 - x86-based, 50–51
- servers. *See also* server virtualization; *specific products and applications*
 - backing up before migration, 193
 - blade servers, 179
 - choices impacted by storage, 171–172
 - choosing, 175–182
 - designed for virtualization, 45, 175, 180–181, 182
 - existing, cleaning up before migration, 192–193
 - key resources, 164–175
 - migration between, benefits of virtualization for, 42
 - new, preparing for migration, 193–196
 - NICs preinstalled on, 167
 - “one application, one server” approach, 37–38, 137, 160–163, 165
 - proprietary attitude toward, 35–36, 137–138
 - reusing existing 32-bit hardware, 175, 176–178, 181
 - 64-bit, 175, 178–180, 181
 - table summarizing choices, 181–182
 - total data per server limitation with DAS, 43
 - underutilization rate for, 35
- shared memory, 27, 58
- shared storage. *See* NAS (Network-Attached Storage); SAN (Storage Area Network); storage virtualization
- Simple Storage Service (S3), 92
- single point of failure. *See* SPOF
- 64-bit servers, 175, 178–180, 181–182
- SLES (SUSE Linux Enterprise Server), 214. *See also* Linux
- small business virtualization, 337
- Small Computer System Interface (SCSI), 225, 227, 228, 229
- soft costs. *See also* cost-benefit analysis
 - backup, 116
 - for current infrastructure, 112–114, 115–116
 - in data center, 115–116
 - defined, 113
 - difficulties establishing, 113
 - estimating, 122
 - identifying reductions with virtualization, 122
 - keeping separate from hard costs, 113
 - preponderance of, 112, 113–114
 - in spreadsheets for cost-benefit analysis, 123–124
 - system administration, 115–116, 122
- Softgrid for Desktops (Softricity), [19](#)
- SoftGrid (Microsoft), [18](#), [19](#)
- Softricity’s Softgrid for Desktops, [19](#)
- software applications
 - backing up before migration, 193
 - contained in virtual OS, 21
 - data generation by, 225
 - end user installation issues, [17](#), [18](#)
 - evaluating support for, 144
 - installation headaches, 86–87
 - migration software for, 38
 - “one application, one server” approach, 37–38, 137, 160–163, 165
 - “orphans” in data centers, 35
 - SaaS environments, 53
 - suitability for virtualization, 107–109
 - user-level, defined, 50
- Software as a Service (SaaS), 53, 90–91, 92–93
- software development. *See* development and testing environments
- software licensing. *See* licensing
- software management products, 195, 203
- software products. *See* software applications; virtualization software products
- Solaris OS (Sun)
 - for development and testing environments, 65

- for failover/high
 - availability/clustering/load balancing, 75
 - OS virtualization with, 23, 53
 - for server consolidation, 68
 - for training environments, 67
- specialized resource, managing
 - virtualization as, 213, 214–215, 220–221
- SPOF (single point of failure). *See also* failover
 - blade servers for avoiding, 179
 - defined, 70
 - failover for avoiding, 70, 74, 164
 - load balancing as protection against, 74
 - multiple processors for avoiding, 166, 174
 - processor issues, 166
 - RAID for avoiding, 229
 - redundant NICs for avoiding, 168, 174
 - risk exposure with virtualization, 163–164
 - server pooling for avoiding, 39, 164
- spreadsheets for cost-benefit analysis
 - assumptions for sample project, 125
 - benefits shown by, 125–127
 - Current Cost Structure, 123, 124
 - hard cost and soft cost sections, 123–124
 - improving the project case, 126–127
 - justifying financial projections, 126
 - NPV (net present value) in, 123
 - Project Cost Summary, 123, 125, 126
 - reexamining assumptions, 127
 - sample set, 123–125
 - showing financial implications across time, 123
 - Total Yearly Hard Costs in, 126
 - usefulness of, 122, 125
 - Virtualized Cost Structure, 123, 124
- S3 (Simple Storage Service), 92
- storage. *See also specific types*
 - administration, 229–230
 - approaches to, 29–31
 - assigning to XenExpress virtual machine, 302
 - backup problems, 29
 - in blade servers, 179
 - bottlenecks, 29
 - costs of implementing, 231, 241
 - data redundancy, 230
 - defined, 224
 - hardware emulation options, 237–238
 - importance for virtualization, 223
 - increasing capacities for, 224–226
 - JBOD (just a bunch of drives), 169, 229
 - as key server resource, 168–172, 175
 - LUNs (Logical Unit Numbers), 229, 230
 - network, 226
 - OS virtualization options, 236–237
 - paravirtualization options, 238–239
 - RAID (redundant array of inexpensive disks), 229, 230
 - requirements for virtualization, 168–172
 - server hardware choices impacted by, 171–172
 - shared, types of, 226–228
 - space issues as virtualization driver, 13–14, 28, 33
 - 32-bit server issues, 176
 - trend like Moore's Law for, 13
- Storage Area Network. *See* SAN
- storage virtualization. *See also* storage; *specific types of storage*
 - applications suitable for shared storage, 234–236
 - approaches to storage, 29–31
 - clustering software, 230
 - cost effectiveness of, 241
 - DAS limitations for, 43, 169–170, 198, 228, 232–233
 - data backup simplified with, 44
 - external storage protocols, 170–171
 - for failover, 235–236
 - with hardware emulation, 237–238
 - for high availability, 235–236
 - horizontal scaling, 234–235
 - infrastructure setup before migration, 198
 - integration into OSes, 84
 - IT costs lowered by, 43–44
 - for large numbers of VMs, 235
 - for load balancing, 74
 - migration between servers with, 42
 - need for, 28–29
 - with OS virtualization, 236–237
 - overview, 28–31

- storage virtualization (*continued*)
 - with paravirtualization, 238–239
 - ramifications for networking and system administration, 43
 - for server pooling, 76, 236
 - shared storage for, 234–236
 - system cloning, 234
 - types of implementations, 169–170
 - virtualization journey steps, 239–240
 - Web site for, 339
- stub drivers, 58
- SugarCRM installation on VMware Server, 264–266
- Sun. *See also* Solaris OS (Sun)
 - Fire 4600 server, 180
 - Fire X4200 server, 181
 - OpenSolaris, 61, 65, 67, 68
 - OS virtualization offered by, 23, 53
- support issues
 - evaluating when choosing products, 144, 333
 - for software installation, 86–87
 - for virtualized software, 104, 107
- SUSE Linux Enterprise Server (SLES), 214. *See also* Linux
- SVS (Altiris), 18
- switch fabric for SANs, 227
- SWsoft. *See also* Virtuozzo (SWsoft)
 - OpenVZ, 23, 61, 67
 - OS virtualization offered by, 23, 53
 - Parallels, 20
 - VZP2V migration tool, 204
- Symantec's Veritas Provisioning Manager, 196
- system administration. *See also* IT; maintenance and upgrades
 - in cost-benefit analysis, 115–116, 122
 - costs as virtualization driver, 15–16, 34, 101
 - soft costs in the data center, 115–116
- system cloning, shared storage for, 234
- System Insight Manager (HP), 204
- system libraries, defined, 52
- System x3550 server (IBM), 180

• T •

- Technical Stuff icon, 5
- Technology Without An Interesting Name (TWAIN), 229
- territoriality
 - proprietary attitude toward servers, 35–36, 137–138
 - Windows versus Linux workgroups, 138–139
- testing applications. *See* development and testing environments
- thin client information, 337
- Thinstall's Virtualization Suite, 18
- 32-bit servers, 175, 176–178, 181–182
- Tip icon, 5
- Tivoli software manager (IBM), 195
- torrents, 262
- training environments
 - as common virtualization application, 65–66
 - OS virtualization suitability for, 53–54
 - typical scenario, 66
 - virtualization products for, 66–67
- training requirements for virtualization
 - avoiding skimping on, 330
 - in deciding about virtualization, 104–105
 - identifying costs of, 119
- transaction count licensing method, 47
- TWAIN (Technology Without An Interesting Name), 229

• U •

- Ubuntu Linux as guest OS, 288–290
- Unicenter software manager (CA), 195
- Unix file structure, 51
- upgrades. *See* maintenance and upgrades
- use cases
 - for architecture selection, 140
 - asking people for, 137
 - business problems addressed by, 135
 - defined, 323
 - for desired management toolset, 136

- for desired virtual machine density, 136
- identifying, defined, 134
- importance of, 134–135, 322–323, 332
- items to evaluate when identifying, 135–136
- need for, 132
- for overall business requirements, 136
- for potential future requirements, 136
- user-level applications. *See* software applications

• U •

- vendor support issues. *See* support issues
- Veritas Provisioning Manager (Symantec), 196
- virtual appliances
 - benefits for end users, [88](#), 210
 - benefits for vendors, [88](#), 210
 - defined, 87
 - as future of virtualization, 89
 - impact on IT personnel, 93–94
 - installing on VMware Server, 263–266
 - Microsoft offerings, 89
 - need for, 87
 - VMware offerings, 89
- Virtual Image Distribution (AppStream), [19](#)
- Virtual Infrastructure [3](#) (VI3) by VMware
 - Art.com server pooling using, [40](#)
 - bare-metal approach of, 210
 - DR capability of, 78
 - DRS (Distributed Resource Scheduler), 77
 - for failover/high
 - availability/clustering/load balancing, 75
- Virtual Iron, 69, 75, 77, 78
- Virtual Machine Manager (Red Hat), 214
- virtual machine monitor (VMM), 23, 25
- virtual machines. *See* VMs
- virtual OSES. *See* guest operating systems; operating system virtualization (containers)
- Virtual PC Guy blog site, 337
- Virtual Server (Microsoft), [20](#), 26, 65, 66
- virtualization events, 338–339

- virtualization life cycle
 - defined, [32](#)
 - graphical representation, 131
 - infrastructure, 327
 - implementation phase, 132–133
 - management choices affected by, 213, 217–218, 220–221
 - need for understanding, 130–131
 - operating phase, 133
 - planning phase, 131–132
 - storage methods in, 239–240
- virtualization pilot. *See* pilot implementation
- virtualization software products. *See also* free virtualization products; *specific products and companies*
 - for application packaging, [18](#)
 - for application streaming, [19](#)
 - application-level, 291
 - for automated software installation, 195–196, 204–207
 - custom management products, 195–196
 - for development and testing environments, 65, 79
 - for disaster recovery, 78, 81
 - evaluating application software support, 144
 - evaluating costs, 145
 - for failover/high
 - availability/clustering/load balancing, 75, 80
 - hardware choices affecting, 145, 174, 181–182, 333–334
 - for hardware emulation, [20](#), 26, [56](#), [60](#)
 - management choices affected by, 213
 - management systems with, [84](#), 155
 - for OS virtualization, 23, 53, [60](#), 61
 - for paravirtualization, 28, 57, [60](#), 61, 65, 67
 - pilot program for testing, 145
 - selecting for virtualization
 - implementation, 132, 144–145, 325
 - for server consolidation, 68–69, 79–80
 - for server pooling, 77, 80–81
 - software management products, 195, 203, 216, 219
 - table summarizing, 79–81
 - for training environments, 66–67

- Virtualization Suite (Thinstall), [18](#)
- virtualization technology types, [60](#)
- virtualization-enabled chips
 - Fedora virtualization with, [282](#)
 - full virtualization with, [282](#), [311–312](#)
 - hardware selection issues, [166](#)
 - Linux QEMU technology for, [311–312](#)
 - need for, [118](#)
 - overview, [28](#), [59](#), [183–184](#)
 - for Xen-based environment, [166](#)
 - with XenExpress, [311–312](#)
- virtualized storage. *See* storage
- virtualization
 - virtualization
- Virtuozzo (SWsoft)
 - for failover/high
 - availability/clustering/load balancing, [75](#)
 - OS virtualization with, [23](#), [53](#), [61](#)
 - for server consolidation, [68](#)
 - for training environments, [67](#)
- VI3. *See* Virtual Infrastructure 3 (VI3) by VMware
- VMM (virtual machine monitor), [23](#), [25](#)
- VMs (virtual machines)
 - accessing Windows guest VM with RDP client in XenExpress, [314–317](#)
 - determining desired density, [136](#)
 - in development and testing environments, [64](#)
 - in hardware emulation for client virtualization, [20](#)
 - in hardware emulation for server virtualization, [23](#)
 - for high availability, [73](#)
 - installing on Fedora [7](#), [278–287](#)
 - installing on VMware Server, [253–258](#)
 - installing on XenExpress, [300–311](#)
 - IT management of, [94](#)
 - managing during pilot implementation, [150](#)
 - memory requirements, [173](#)
 - mixed workload issues, [138–139](#)
 - price elasticity of, [212](#)
 - security issues, [199](#)
 - shared storage for large numbers, [235](#)
 - state of, [71–72](#)
- VMware. *See also* ESX Server (VMware); Virtual Infrastructure 3 (VI3) by VMware; VMware Server
 - bare-metal products, [210](#)
 - blog site on virtualization, [337](#)
 - Capacity Planner, [147](#)
 - challenges from virtualization integration into OSes, [84–85](#), [210–211](#)
 - Converter migration tool, [152](#), [204](#), [205](#)
 - dispute with Microsoft, [48](#)
 - DRS (Distributed Resource Scheduler), [77](#)
 - EMC acquisition of, [1](#)
 - hardware compatibility list, [118](#)
 - hardware emulation offered by, [26](#), [56](#)
 - IPO of, [1](#)
 - as leader in virtualization deployment, [61](#)
 - managing virtualization as specialized resource, [214](#)
 - OS virtualization products, [209](#)
 - user group listings, [338](#)
 - virtual appliance offerings, [89](#)
 - VMware Workstation, [209](#)
- VMware Server
 - acquiring your free product, [248](#)
 - for development and testing environments, [65](#)
 - ESX Server versus, [246](#), [291](#)
 - flexibility of host OS with, [247](#)
 - hardware emulation provided by, [20](#), [26](#), [56](#)
 - installing a guest OS, [258–261](#)
 - installing a virtual machine, [253–258](#)
 - installing Fedora image files, [261–263](#)
 - installing virtual appliances, [263–266](#)
 - installing VMware Server, [248–251](#)
 - memory requirements, [173](#)
 - multisystem configuration capabilities, [247](#)
 - performance issues, [246](#), [247](#)
 - pros and cons of, [246–247](#)
 - redundant NICs for, [175](#)
 - resource contention with, [246](#)
 - serial number for, [248](#)
 - starting the Server Console, [252–253](#)
 - for training environments, [66](#)

virtual appliances, 89, 264
 virtualization approach of, 209, 246
 Xen versus, 291
 VMware Workstation, 209
 VMworld event, 339
 VZP2V migration tool (SWsoft), 204

• W •

Warning! icon, 5
 Web resources. *See* Internet resources
 Web-hosting companies, OS virtualization
 for, 21–22
 Windows. *See* Microsoft Windows
 Windows Server (Microsoft), SoftGrid to be
 incorporated with, 19

• X •

x86 chips. *See also* processors
 blogs about virtualization, 337
 commodity hardware based on, 19, 50, 159
 limitations of, 50, 54
 software emulation of, 24
 virtualization architected for, 50–51
 x86-based systems, 50
 Xen (XenSource). *See also* Fedora
 virtualization
 AWS's use of, 92
 bare-metal approach of, 291
 blog site on virtualization, 337
 described, 291
 for development and testing
 environments, 65
 Domain0 as privileged guest, 57
 DomainU as guest OS, 57
 for failover/high
 availability/clustering/load
 balancing, 75
 as free virtualization product, 61
 functionality extensions, 61
 hardware requirements, 117–118
 inclusion in Linux distributions, 28, 84
 managing virtualization as specialized
 resource, 214

memory requirements, 173
 as paravirtualization example, 28, 57
 PXE for installing, 196
 for server consolidation, 69
 for training environments, 66
 virtualization-enabled chips for, 166, 183
 VMware Server versus, 291
 XenEnterprise (XenSource), 69, 75, 205
 XenExpress (XenSource)
 accessing a Windows guest VM with an
 RDP client, 314–317
 dual-boot mode not possible for, 294
 installing paravirtualized (PV) drivers,
 311–314
 installing Windows XP SP2 as guest
 virtual machine, 300–311
 installing XenConsole, 296–297
 installing XenExpress, 293–295
 obtaining, 292–293
 paravirtualization versus full
 virtualization, 311–312
 for server consolidation, 69
 working with XenConsole, 297–299
 XenServer (XenSource), 69, 75
 XenSource. *See also specific products*
 described, 292
 paravirtualization products, 28, 57,
 61, 65, 67
 P2V wizard, 204
 server consolidation products, 69

• Z •

zealotry, avoiding, 117



consumption for a
greener world

Understand virtualization and choose the best options for your situation

Virtualization is the hottest topic in the data center today, but how do you separate the hope from the hype? This book gives you the tools you need to make virtualization work for you. From selling the idea to your boss to choosing the right hardware and migrating to a virtualized environment, here's virtually all you need to know.

David Golden, MBA, is CEO of Navica, Inc., an IT consulting firm. He has experience in all levels of the IT industry from software engineer to executive and is an expert blogger for CIO.com.

**DUMMIES
WAY**

Explanations in plain English
"Get in, get out" information
Icons and other navigational aids
Tear-out cheat sheet
Top ten lists
A dash of humor and fun

Discover how to

*Understand the different
types of virtualization*

*Weigh the pros and cons
of virtualization*

*Administer a virtualized
infrastructure*

*Choose hardware for
your project*

*Integrate a storage
virtualization strategy*

**Get smart
@www.dummies.com**

- ✓ Find listings of all our books
- ✓ Choose from many different subject categories
- ✓ Sign up for eTips at etips.dummies.com

ISBN 978-0-470-14831-0



5 2 9 9 9

Obras protegidas por Direitos de Autor

