

Microsoft®

Updated for Windows Server 2008 R2

Understanding
Microsoft®

2
SECOND
EDITION

Virtualization Solutions

From the Desktop to the Datacenter



Mitch Tulloch with the
Microsoft Virtualization Teams

PUBLISHED BY

Microsoft Press

A Division of Microsoft Corporation

One Microsoft Way

Redmond, Washington 98052-6399

Copyright © 2010 by Microsoft Corporation

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Control Number: 2010920178

A CIP catalogue record for this book is available from the British Library.

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office or contact Microsoft Press International directly at fax (425) 936-7329. Visit our Web site at www.microsoft.com/mspress. Send comments to mspinput@microsoft.com.

Microsoft, Microsoft Press, Access, Active Directory, ActiveX, Aero, Authenticode, BitLocker, Excel, Hyper-V, Internet Explorer, MS, MSDN, MS-DOS, Outlook, SharePoint, Silverlight, SQL Server, Visual Basic, Visual C++, Windows, Windows Media, Windows Mobile, Windows NT, Windows PowerShell, Windows Server and Windows Vista are either registered trademarks or trademarks of the Microsoft group of companies. Other product and company names mentioned herein may be the trademarks of their respective owners.

The example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred.

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the authors, Microsoft Corporation, nor its resellers, or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

Acquisitions Editors: Ken Jones

Developmental Editor: Devon Musgrave

Project Editor: Valerie Woolley

Editorial Production: Waypoint Press, www.waypointpress.com

Technical Reviewer: Bob Hogan; Technical Review services provided by Content Master, a member of CM Group, Ltd.

Cover: Tom Draper Design

Table of Contents

Acknowledgments	ix
Introductions	xi
1. Why Virtualization?	1
Understanding Dynamic IT	2
Microsoft's Infrastructure Optimization Model	2
Virtualization and the Infrastructure Optimization Model	4
Benefits of Virtualization	7
How Virtualization Enables Dynamic IT	7
Achieving the Benefits of Datacenter Virtualization	9
Achieving the Benefits of Client Virtualization	10
Achieving the Benefits of Cloud Virtualization	11
Windows Optimized Desktop Scenarios	12
Mobile Worker Scenario	12
Office Worker Scenario	13
Task Worker Scenario	13
Contract/Offshore Worker Scenario	14
Anywhere-Access Scenario	14
Microsoft's Integrated Virtualization Solution	15
Microsoft's Commitment to Virtualization	16
Additional Resources	18
General	18
Microsoft's IT Infrastructure Optimization Model	18
Microsoft's Dynamic IT	18
Microsoft Virtualization Technologies and Solutions	19
Windows Optimized Desktop Scenarios	19

 **What do you think of this book? We want to hear from you!**

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

www.microsoft.com/learning/booksurvey/

2. Server Virtualization	21
Understanding Server Virtualization	21
Understanding Virtual Machines	21
Understanding Hypervisors	23
Understanding the Hyper-V Architecture	27
Understanding the Parent Partition	28
Understanding Child Partitions	33
Key Features of Hyper-V	36
New Features in Hyper-V R2	37
Comparing Hyper-V and Virtual Server	40
Key Benefits of Using Hyper-V	41
Hyper-V Usage Scenarios	42
Server Consolidation	43
Business Continuity and Disaster Recovery	43
Testing and Development	43
The Dynamic Datacenter	43
Working with Hyper-V	46
Hyper-V Role vs. Microsoft Hyper-V Server	46
System Requirements for Using Hyper-V R2	48
Supported Guest Operating Systems	50
Functionality Provided by Integration Services	52
Planning for Hyper-V Deployment	53
Installing the Hyper-V Role	54
Using the Hyper-V Management Snap-in	57
Using the Virtual Machine Connection Tool	70
Creating a Virtual Machine	74
Working with Virtual Machines	79
Working with Live Migration	88
Tools for Managing Hyper-V and Virtual Machines	97
Additional Resources	104
General	104
Planning for Hyper-V	105
Deploying Hyper-V	105
Managing and Maintaining Hyper-V	106
Securing Hyper-V	107
Resources for Hyper-V Developers	107

Hyper-V Bloggers at Microsoft	107
Other Hyper-V Bloggers.....	108
Hyper-V Forum on TechNet.....	108
3 Local Desktop Virtualization	109
Examining the Benefits of Each Technology	110
Key Benefits of Windows Virtual PC and the Windows XP Mode Environment	110
Key Benefits of MED-V	111
Key Benefits of App-V	111
Examining Usage Scenarios for Each Technology	112
Usage Scenarios for Windows Virtual PC and the Windows XP Mode Environment	113
Usage Scenarios for MED-V	113
Usage Scenarios for App-V	113
Availability of Each Technology	114
Availability of Windows Virtual PC and the Windows XP Mode Environment	114
Availability of MED-V	114
Availability of App-V	115
Understanding Windows Virtual PC and the Windows XP Mode Environment	115
Understanding Windows Virtual PC.....	115
Understanding Virtual Applications	119
Understanding Windows XP Mode	121
Requirements for Windows Virtual PC.....	123
Installing Windows Virtual PC	124
Requirements for Windows XP Mode	125
Installing the Windows XP Mode Environment	126
Configuring Virtual Machine Settings	129
Using Windows XP Mode.....	138
Understanding MED-V	147
Introducing Microsoft Enterprise Desktop Virtualization	147
How MED-V Works	150
Understanding App-V.....	162
App-V Terminology.....	167
How App-V Works.....	169
App-V Components.....	176
App-V Architecture	182

Working with App-V	185
App-V Deployment Scenarios	185
Using the Management Console	192
Using the Sequencer	201
Working with App-V Clients	214
Additional Resources	219
Resources for Windows Virtual PC and Windows XP Mode	219
Resources for MED-V	220
Resources for App-V	221
4 Remote Desktop Virtualization	223
Examining the Benefits of Remote Desktop Virtualization	224
Examining Usage Scenarios for Remote Desktop Virtualization	224
Usage Scenarios for Remote Desktop Services	224
Usage Scenarios for App-V for RDS	226
Usage Scenarios for Microsoft VDI	227
Availability of Remote Desktop Virtualization Technologies	227
Availability of Remote Desktop Services	227
Availability of App-V for RDS	227
Availability of Microsoft VDI	228
Understanding Remote Desktop Services	228
Understanding Remote Desktop Connection Client Experience Improvements	230
Understanding the Remote Desktop Session Host	231
Understanding Remote Desktop Web Access	256
Understanding RemoteApp and Desktop Connections	263
Understanding Remote Desktop Connection Broker	272
Understanding Remote Desktop Gateway	276
Understanding Remote Desktop Licensing	281
Understanding Remote Desktop Virtualization Host	284
Deploying Remote Desktop Services	299
Understanding Microsoft Application Virtualization for Remote Desktop Services	301
Understanding Microsoft Virtual Desktop Infrastructure	303
Understanding Microsoft's VDI Architecture	304
How Microsoft VDI Works	307

Additional Resources	309
Additional Resources on Remote Desktop Services.	309
Additional Resources for App-V for RDS	310
Additional Resources for Microsoft VDI.	311
5 Virtualization Management.	313
Understanding Virtual Machine Manager	313
Terminology	313
VMM Components	315
VMM Architecture	316
Key Features of VMM	330
Features and Improvements Introduced in VMM 2008	330
New Features and Enhancements in VMM 2008 R2	333
Key Benefits of VMM.	338
Usage Scenarios for VMM	340
Server Consolidation	340
Provisioning of Virtualized Resources	340
Business Continuity	341
Working with VMM 2008 R2	342
Planning for Deploying VMM 2008 R2	342
System and Infrastructure Requirements.	343
Installing VMM 2008 R2	348
Using the VMM Administrator Console	356
Working with Managed Hosts.	361
Working with the Library	378
Working with Virtual Machines.	384
Performing P2V Conversions.	400
Performing V2V Conversions.	411
Configuring User Roles	413
Using the Self-Service Portal	421
Microsoft System Center Solutions	425
System Center Server Management Suite Enterprise	426
System Center Essentials.	426
Other System Center Products	427
Benefits of System Center for Virtualization	427

Additional Resources	430
General	430
Administering VMM	430
System Center Blog	430
VMM Forums on TechNet	430
6 Cloud Computing	431
What Is Cloud Computing?	431
Private vs. Public Cloud	432
Examining the Benefits of Cloud Computing	433
Benefits of Using a Private Cloud vs. a Public Cloud	433
Increasing Use of IT Resources	434
Examining Cloud-Computing Usage Scenarios	435
Understanding Microsoft's Cloud-Computing Platform	435
Understanding Different Cloud Services	435
Implementing Cloud Services	437
Understanding the Dynamic Data Center Toolkit	438
Comparing the Toolkits	440
Understanding the Private-Cloud Architecture	441
Implementing a Private-Cloud Solution	443
Windows Azure	444
The Dynamic Data Center Alliance	446
Availability of Microsoft's Cloud-Computing Platform	446
Additional Resources	447
Additional Resources for Microsoft's Cloud-Computing Initiative	447
Additional Resources for Windows Azure	447
Index	449

 **What do you think of this book? We want to hear from you!**

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

www.microsoft.com/learning/booksurvey/

Acknowledgments

This book would not have been possible without the support and assistance of numerous individuals. First, I would like to especially thank Michael Cooper, Senior Product Manager and Field Technical Community Lead for System Center and Virtualization Business Group; Aurora Santiago, Product Manager for System Center and Virtualization Technical Readiness; and Kenon Owens, Technical Product Marketing Manager for Integrated Virtualization, all of whom helped drive this project forward and provided liaison with other members of the virtualization team at Microsoft. Next, my sincere thanks to each of following experts at Microsoft who provided key technical insights, peer-reviewed chapter content, contributed *Direct from the Source* sidebars, and assisted me in many other ways with this project:

Aaron Holzer, Alex Balcanquall, Arun Jayendran, Balagopan Nikhil, Chuck Timon, Fei Lu, Isaac Roybal, Karri Alexio-Tiernan, Max Herrmann, Megan Kidd, Mohit Srivastava, Peter Ballantyne, Prashant Ketkar, Ran Kohavi, Ran Oelgiesser, Rick Kingslan, Vipul Shah and Wole Moses

Thanks also to Brett Polen of Xtreme Consulting Group, Rex Backman of Simplicity Consulting, and Nicole Pargoff of iSoftStone who assisted me with certain topics covered in this book. Special thanks to Bill Noonan, Mark Kitris, and the CTS Global Technical Readiness (GTR) team at Microsoft for contributing their expertise to this project.

I'd also like to thank again others at Microsoft together with several Microsoft Most Valuable Professionals (MVPs) who contributed their expertise to the previous edition of this book, namely:

Anshul Rawat, Baldwin Ng, David Greschler, Edwin Yuen, Falko Gräfe, James O'Neill, Jason Leznec, Jeff Woolsey, Kalle Saunamäki Kyle Beck, Michelle Foley, Ming Zhu, Peter Larsen, Sean Donahue and Tim Mangan

Next, special thanks Devon Musgrave, the development editor for this project, and Valerie Woolley, the project editor for this project, both of whom work at Microsoft Press. I've enjoyed working with on this book and hope to do so again on another one in the near future. Thanks also to Steve Sagman of Waypoint Press who managed the editing and production for this book, to Bob Hogan the technical editor for this project, and to the ever-insightful Roger LeBlanc who copy edited the manuscript. Thanks also to Ken Jones who was project planner for this title.

As always, heartfelt thanks to my friend and agent, Neil Salkind of the Salkind Agency, which is part of Studio B Productions, Inc.

And last but never least, thanks to my wife, Ingrid, for her encouragement and support during this project.

—Mitch Tulloch

Introduction

Welcome to *Understanding Microsoft Virtualization Solutions, From the Desktop to the Datacenter, 2nd Edition*. This is the book for IT professionals who want to learn more about the latest Microsoft virtualization technologies, including Hyper-V and Remote Desktop Services in Windows Server 2008 R2, Microsoft Virtual Desktop Infrastructure, Microsoft Application Virtualization 4.5, Microsoft Enterprise Desktop Virtualization, Windows Virtual PC and Windows XP Mode, System Center Virtual Machine Manager 2008, and Microsoft's private and public cloud computing platforms including Windows Azure.

Who Is This Book For?

The primary target audience for this book is IT administrators, implementers, and decision makers of large and mid-sized organizations who want to learn about the benefits of the latest virtualization technologies and how to plan, implement, and manage virtual infrastructure solutions based on these technologies. The book assumes that you are familiar with core Windows Server technologies and how to implement an Active Directory Domain Services infrastructure. The book also assumes you have experience working with the latest client and server versions of Windows, namely Windows 7 and Windows Server 2008 R2. Finally, the book assumes you are already familiar with earlier Microsoft virtualization products such as Microsoft Virtual Server 2005 and Microsoft Virtual PC 2007.

How This Book Is Organized

The book is intended to be read from cover to cover and will give you a good understanding of the capabilities, features, and operation of Microsoft virtualization technologies. You can also read individual chapters to gain an understanding of a particular product or technology.

The topics covered by the various chapters are as follows:

- **Chapter 1: Why Virtualization?** This chapter provides an overview of Microsoft's integrated virtualization solution and how it plays a key role in Dynamic IT, Microsoft's strategy for enabling agile business. The chapter also describes the benefits businesses can achieve through virtualization and how Microsoft's virtualization platforms, products and technologies can help these businesses move their IT infrastructures toward the goal of Dynamic IT.

- **Chapter 2: Server Virtualization** This chapter covers the Hyper-V role of Windows Server 2008 R2 and Microsoft Hyper-V Server 2008 R2 and how these platforms can be used to manage virtualization server workloads in the datacenter. The chapter explores features of Hyper-V including the new Live Migration feature of Windows Server 2008 R2. It also describes the benefits of deploying Hyper-V, and various usage scenarios.
- **Chapter 3: Local Desktop Virtualization** This chapter describes various Microsoft virtualization technologies that enable client operating systems and applications to run within a virtualized environment hosted on the user's computer. The platforms and products covered in this chapter include Windows Virtual PC and the Windows XP Mode environment, Microsoft Enterprise Desktop Virtualization (MED-V), and Microsoft Application Virtualization (App-V).
- **Chapter 4: Remote Desktop Virtualization** This chapter describes various Microsoft virtualization technologies that enable client operating systems and applications to run within a virtualized environment hosted on a server located in the datacenter. The platforms and products covered in this chapter include Remote Desktop Services in Windows Server 2008 R2, Microsoft Virtual Desktop Infrastructure (VDI), and App-V for Remote Desktop Services.
- **Chapter 5: Virtualization Management** This chapter describes how System Center Virtual Machine Manager (VMM) 2008 can be used to centrally manage all aspects of a virtualized IT infrastructure. The chapter explains how VMM works and explores how to use the platform to manage virtual machines running on Windows Server 2008 R2 Hyper-V servers. The chapter also describes the benefits of the other members of the System Center family of products.
- **Chapter 6: Cloud Computing** This chapter examines Microsoft's emerging cloud computing platform, how it works, and what benefits businesses can obtain from it. The chapter examines both private and public cloud solutions including Windows Azure, and describes how Microsoft's Dynamic Data Center Toolkit can be used to integrate cloud computing as a part of your virtualized IT infrastructure.

Conventions Used in This Book

The following elements have been used in this book to help keep the text clear and easy to follow:

- **Note** Provides additional detail or a sidelight on the topic under discussion
- **Tip** Gives you some cool pointers that you'll probably want to know because they will make your job easier
- **Caution** Informs you of things to be aware of so that you can avoid potential pitfalls

An important feature of the book is sidebars written by Microsoft product groups and experts in the field. Sidebars written by individuals or teams inside Microsoft are titled *Direct from the Source*; sidebars written by other experts such as Microsoft Most Valuable Professionals (MVPs) are titled *Direct from the Field*.

Other Virtualization Resources

While this book is intended as a broad introduction to the technical aspects and benefits of Microsoft virtualization technologies, Microsoft provides many other useful resources from which you can learn more about these technologies. These resources include Microsoft TechNet, Microsoft bloggers, and other online resources. To help you find the most relevant resources, each chapter concludes with an “Additional Resources” section that provides descriptions and URLs for these resources.

Contact the Author

Feel free to contact me at virtual@mtit.com if you have comments, questions, or suggestions regarding anything in this book. Although I respond to all queries from readers and will do my best to answer your question to your satisfaction, I cannot provide readers with technical support.

Support

Every effort has been made to ensure the accuracy of this book. As corrections or changes are collected, they will be added to a Microsoft Knowledge Base article accessible via the Microsoft Help and Support site. Microsoft Press provides support for books, including instructions for finding Knowledge Base articles, at the following Web site:

<http://www.microsoft.com/learning/support/books>

If you have questions regarding the book that are not answered by visiting the site above or viewing a Knowledge Base article, send them to Microsoft Press via e-mail to

mspinput@microsoft.com

Please note that Microsoft software product support is not offered through these addresses.

We Want to Hear from You

We welcome your feedback about this book. Please share your comments and ideas via the following short survey:

<http://www.microsoft.com/learning/booksurvey>

Your participation will help Microsoft Press create books that better meet your needs and your standards.



Note We hope that you will give us detailed feedback via our survey. If you have questions about our publishing program, upcoming titles, or Microsoft Press in general, we encourage you to interact with us via Twitter at *<http://twitter.com/MicrosoftPress>*. For support issues, use only the email address shown above.

Chapter 1

Why Virtualization?

I'm writing this chapter as the first decade of the twenty-first century draws to a close, and the fundamental point I want to convey in my book is this: businesses of all sizes need virtualization more than ever before.

Why? Because times have changed. With the global economy in the doldrums, the key concern for many businesses is survival—how to keep the lights turned on. And virtualization—with its potential gains in efficiency and ability to lower costs—is viewed by many businesses as a key strategy to enable their future survival in the marketplace.

Survival is not their only concern, however. Based on surveys conducted by Microsoft, some specific concerns of businesses at this time in history include

- How to control capital expenditures (CapEx) and make operational expenditures (OpEx) more predictable
- How to create more business value by providing services faster yet cheaper
- How to make their business more "Green" while avoiding higher costs
- How to ensure security, safeguard privacy, and meet compliance regulations and standards

These are some of the key concerns that are driving the spending decisions of today's businesses, and virtualization can address each one of these concerns. A recent study that surveyed Chief Information Officers (CIOs) indicated that more than one third of respondents identified server, storage, and cloud virtualization as drivers of their spending decisions for 2009 and 2010, and almost one quarter of respondents also identified desktop virtualization as similar drivers. The study even indicated that virtualization would influence their spending decisions more than issues such as labor optimization, wireless computing, Green computing, or security concerns.

As a result of these concerns, more businesses than ever are investigating how Microsoft virtualization platforms, products, and solutions can help them address their concerns. In other words, more businesses than ever are aligning themselves with Dynamic IT, Microsoft's strategic vision for implementing IT infrastructures that can automatically adjust to changing business conditions by aligning computing resources with business objectives.

Understanding Dynamic IT

The goal of Dynamic IT is to lower cost and improve business agility by establishing a highly optimized IT infrastructure that can respond to changing needs in an automated fashion. A dynamic IT infrastructure is one that is

- Logically based instead of physically based
- Managed by policy
- Services based
- Federated and connected
- State aware and self-healing
- Highly available
- Secure

Dynamic IT infrastructures have efficient development-to-IT operations and can include any combination of on-premise, off-premise, and hosted IT services.

Most businesses today have not achieved the ideal of Dynamic IT, but many are steadily progressing toward this vision. Microsoft's Infrastructure Optimization Model can help you evaluate where your own organization is located on the continuum that ranges from the traditional enterprise to the dynamic datacenter.

Microsoft's Infrastructure Optimization Model

Without an IT environment that is efficient, reliable, and easily managed—at the lowest possible cost—businesses today won't survive. This observation is true particularly in the midmarket sector, where the pressure to grow your business is greatest, where competition for IT talent is most intense, and where budgetary constraints are often felt the hardest.

As a response to the needs of this segment, Microsoft developed the IT Infrastructure Optimization Model, a framework that helps you understand and improve your organization's IT infrastructure by using specific, concrete actionable items. The framework outlines the steps a business can take to determine where it is today with its IT infrastructure, where it needs to go, and exactly how to get there. The application of this framework can help a business create an IT environment that is easy to manage and makes the most efficient use of IT resources—including people, hardware, and software—as might be possible.

Microsoft's Infrastructure Optimization Model is particularly helpful for midsized businesses because they generally don't have the luxury of having a large IT staff. Yet they do have a critical need for an IT infrastructure that provides the level of service their workers

need when operating in today's business environments. Workers today need quick access to corporate resources, an ability to easily communicate and collaborate online, and the most up-to-date business tools in order to perform their jobs. Microsoft's Infrastructure Optimization Model—together with virtualization as a key enabler of this framework—can help make this happen for your company.

As shown in Figure 1-1, Microsoft's Infrastructure Optimization Model defines your existing IT infrastructure as being in one of four possible categories: Basic, Standardized, Rationalized, and Dynamic. These categories range from least optimized (Basic) to most optimized (Dynamic).




FIGURE 1-1 Microsoft's Infrastructure Optimization Model.

Basic IT Infrastructure

A typical Basic IT infrastructure is one characterized by manual, localized business processes; minimal central control of resources; and nonexistent or unenforced IT policies and standards for security, backup, deployment, compliance, and other common IT practices. In a Basic IT infrastructure, the health of your applications and services is generally unknown because of a lack of suitable tools and resources for gauging this. In addition, patch management, software deployment, and desktop services are provisioned and maintained manually.

Standardized IT Infrastructure

A typical Standardized infrastructure builds on the Basic one by introducing controls through implementing standards and policies for managing desktops and servers, controlling provisioning and deployment of computers onto the network, and using Active Directory Domain Services (AD DS) to centralize management of network resources, security policies, and access control. Patch management, software deployments, and desktop services have been

partially automated using light-touch technologies. Efforts are made for inventorying hardware and software and for managing licenses. Security at the perimeter of the network has been enhanced by the use of a firewall and malware filtering, but security inside the network is not yet a primary focus.

Rationalized IT Infrastructure

A typical Rationalized infrastructure is one in which the costs involved in managing desktops and servers has been significantly lowered and the processes and policies that support your business have been optimized. The approach to security is now proactive both at the perimeter and within the network, and threat response is methodical in its approach. Zero-touch deployment technologies simplify software deployment and minimize cost. Hardware and software are carefully inventoried, and the business purchases only the licenses it needs.

Dynamic IT Infrastructure

A typical Dynamic infrastructure is one in which the business is fully aware of the strategic value of its IT infrastructure, and this awareness enables it to run its business efficiently and remain ahead of competitors. IT costs are now fully controlled, and there is tight integration between users, data, desktops, and servers. Collaboration between users is pervasive, and mobile users have nearly the same level of access as desktop users. IT processes have been fully automated, which facilitates managing IT according to the needs of the business. Any additional technology investments made by your IT department tend to yield specific, measurable benefits for the operation of the business. Both manageability and security have been greatly enhanced by the use of self-provisioning software and quarantine-like systems, and these systems help ensure compliance with established security policies to improve reliability, lower costs, and increase service levels.

Virtualization and the Infrastructure Optimization Model

Where does your IT infrastructure fit in the Infrastructure Optimization Model? Are you at the Basic or Standardized stage, or are you moving toward having a Rationalized or Dynamic infrastructure? And how can implementing virtualization technologies help move your infrastructure further along toward the goal of an efficient, reliable Dynamic IT infrastructure? Figure 1-2 summarizes some of the many benefits of virtualization technologies and how these benefits align with the Infrastructure Optimization Model.




FIGURE 1-2 How the business benefits of virtualization align with Microsoft's IT Infrastructure Optimization Model.

From Basic to Standardized

As the figure illustrates, from a business perspective, virtualization can help move your IT infrastructure from the Basic stage to the Standardized stage by reducing your total cost of ownership (TCO). Virtualization does this in the following ways:

- **Reducing power and space requirements** By reducing the number of physical computers you need on your network to host your services and applications, virtualization can help you reduce your power requirements and thus the cost of supplying power for your IT infrastructure. And reducing the number of physical computers you need also means less space is required in your server room or hosting area. The cost savings involved can be especially significant in large data centers, where the amount of electricity needed to run thousands of computers, the amount of floor space you need to lease to locate them, and the expensive cooling equipment needed to maintain them can cause costs to reach six figures or higher each year.

- **Accelerating server provisioning and consolidation** By enabling you to consolidate multiple servers onto fewer physical computers, virtualization makes the utilization of your IT assets more efficient. And by providing you with tools for quickly and easily provisioning servers, your infrastructure becomes more flexible and adaptive to change.
- **Eliminating application incompatibility issues** By allowing you to run older applications in their own virtual environments while running other applications simultaneously on the latest physical hardware, application compatibility issues are minimized or even eliminated. In addition, you can retire the inefficient, old hardware your older applications used to run on.
- **Rapid return on investment (ROI)** All the aforementioned benefits of using virtualization to move your IT infrastructure toward the Standardized stage can result in a significant ROI for your business, even in as short a time frame as six months.

From Standardized to Rationalized

From a business perspective, virtualization can also help move your infrastructure from the Standardized stage to the Rationalized stage by increasing the availability of business-critical applications and services. Virtualization does this in the following ways:

- **Simplifying backup and recovery** By enabling you to back up and recover entire virtual machines easily, virtualization helps rationalize your backup and restore processes, making them simpler to use and more reliable. Recovering from backup becomes a rapid process that minimizes service interruptions for workers and customers, thus increasing availability of business-critical network services.
- **Enhancing business continuity** By allowing you to capture point-in-time snapshots of running virtual machines, you can save an image of the machine that you can then return to at any later stage if needed. From a business-continuity perspective, this means you can recover your business more quickly after a disaster. For example, snapshots of virtual machines can quickly be restored onto hardware located at a standby location, allowing the business to resume operations with minimal interruption. This enhanced capability for business continuity is built right into the virtualization business model. You can also gain enhanced disaster recovery capabilities by replicating your storage area network (SAN) from one location to another and by using stretch clustering.
- **Enabling dynamic storage provisioning** Virtualization also enables the addition or removal of virtualized storage resources as needed. The flexibility provided by dynamic storage provisioning not only reduces costs by preventing underutilization of storage resources, it also prevents storage capacity from running out and causing application crashes. For example, virtual machines can be automatically unmapped from a lower priority use, reconfigured as might be needed, and quickly brought up for some new use to meet evolving demand.

From Rationalized to Dynamic

From a business perspective, virtualization can also help move your infrastructure from the Rationalized stage to the Dynamic stage by increasing the agility (flexibility and responsiveness to change) of your infrastructure. Virtualization does this in the following ways:

- **Providing a logical IT infrastructure** By enabling you to view and manage your IT infrastructure as a series of logical layers instead of a collection of physical hardware, virtualization simplifies the provisioning, management, and troubleshooting of systems and applications. These benefits can be felt throughout your infrastructure—from the data center to the desktop.
- **Facilitating self-managing dynamic systems** The holy grail of business computing is an agile IT infrastructure that enhances the dynamic capabilities of people, processes, and technology. Microsoft's Dynamic IT, formerly known as Dynamic Systems Initiative (DSI), is designed to provide technology and solutions that enable businesses to be as agile as possible, and agility is a key to success in the fast-moving world of the Internet economy. And as described next, virtualization is one of the key enablers of Dynamic IT.

Benefits of Virtualization

Virtualization is a key driver that can help you move your IT infrastructure to the Dynamic stage of Microsoft's Infrastructure Optimization Model. To accomplish this, Microsoft provides a broad selection of virtualization technologies, platforms and products that make it easier than ever to achieve the goal of Dynamic IT.

How Virtualization Enables Dynamic IT

Three major business benefits of virtualization are lower total cost of ownership (TCO), increased availability, and improved business agility. Your business can achieve these benefits by implementing a virtualization solution that takes advantage of the different capabilities of Microsoft virtualization technologies, platforms, and products that can help you realize the goal of Dynamic IT.

Cost is usually the most important consideration for a business upgrading its IT infrastructure, and implementing a Microsoft virtualization solution allows you to lower your TCO by

- Increasing hardware utilization to ensure more efficient use of resources through server consolidation
- Reducing power consumption and use of space in the datacenter
- Reducing your licensing costs and other up-front costs
- Simplifying application and desktop life-cycle management
- Lowering operational costs for maintenance and training

Another important consideration for today's dynamic businesses is to increase server availability to ensure that services are provided to customers without delay or interruption. Microsoft's virtualization technologies can help you in this area by enabling you to do the following:

- Increase your service levels, and minimize disruption to services
- Reduce application and desktop deployment time by using virtual applications and virtual desktops
- Enhance desktop business continuity across your organization

Another key goal of Dynamic IT is to improve business agility. In today's highly competitive global marketplace, being able to respond to market changes by being flexible and efficient is a critical key to ongoing success. Microsoft virtualization products help you increase your business agility by

- Integrating your physical, virtual, and application management
- Enabling flexible desktop and application deployments
- Resolving application compatibility issues that can block desktop migration to the latest version of Windows
- Providing capacity on demand that accelerates responses to changing business needs

These three drivers of lowering TCO, increasing availability, and making business more agile can be applied to three different areas of how IT services are delivered:

- **Datacenter virtualization** The virtual datacenter forms the foundation of a dynamic IT infrastructure that can rapidly change in response to changing workloads and demands. The Hyper-V role in Windows Server 2008 R2 and Microsoft Hyper-V Server R2 form the foundation of Microsoft's solution for virtualizing the datacenter. These platforms are described in detail in Chapter 2, "Server Virtualization."
- **Client virtualization** Virtualizing desktop operating systems and applications allows users to have access to the resources they need from anywhere in order to get their work done. This optimized desktop scenario is enabled by a number of Microsoft technologies and products, including Remote Desktop Services in Windows Server 2008 R2, Microsoft Virtual Desktop Infrastructure (VDI), Microsoft Enterprise Desktop Virtualization, Microsoft Application Virtualization, and Windows Virtual PC and the Windows XP Mode environment. These client virtualization technologies are described in detail in Chapter 3, "Local Desktop Virtualization" and Chapter 4, "Remote Desktop Virtualization."
- **Cloud virtualization** Cloud computing makes your business more agile by enabling you to expand or contract your IT infrastructure on demand to provide your users with the resources they need to consume when they need them. Microsoft's private and public cloud computing platforms are described in Chapter 6, "Cloud Computing."

While these three pillars of datacenter, client, and cloud virtualization are necessary to enable Dynamic IT, they aren't sufficient. The final piece needed for this is a fourth pillar: a unified management platform that can manage all aspects of the physical, virtual, and cloud computing resources, including hardware, servers, desktops, applications, and user settings and data. (See Figure 1-3.). Without such an integrated management platform, your virtualized IT infrastructure is just so many disconnected pieces and cannot fulfill the potential of the Dynamic IT vision. Microsoft's unified management platform for doing this is the Microsoft System Center family of products, particularly a key product from this family—System Center Virtual Machine Manager 2008—is described in detail in Chapter 5, "Virtualization Management."




FIGURE 1-3 The four pillars of Dynamic IT.

Achieving the Benefits of Datacenter Virtualization

Server consolidation—running multiple virtual servers on a single physical host—is clearly one technique you can use to save your organization money by enabling you to better use your IT resources. Hyper-V in Windows Server 2008 R2 and Microsoft Hyper-V Server R2 provide you with the tools you need to consolidate servers in your datacenter.

Server consolidation by itself is not enough, however, because it results in a series of silos—separate physical hosts with virtual machines running on them. This structure creates a situation in which when a physical host goes down, the virtual machines on that host also go down, and that leads to service interruptions that can cost your organization money.

A more dynamic solution is to create a clustered pool of hosts in which virtual servers can be easily moved from one host to another. When a host needs to be taken down for maintenance, the virtual servers running on the host are temporarily moved to another host so that services are uninterrupted. The result is higher availability of services and happier

customers. The new Live Migration feature of Hyper-V in Windows Server 2008 R2 and Microsoft Hyper-V Server R2, enabled through the enhanced Failover Clustering feature of Windows Server 2008 R2, makes it possible to create such clustered pooled resources in your datacenter.

Another benefit of being able to create such clustered pools is the ability to increase the agility of your IT infrastructure by allowing you to move virtual servers between hosts to better balance the use of your resources. By adding Microsoft System Center Virtual Machine Manager (VMM) to the mix, not only can you easily perform a Live Migration, but you can also easily add virtual servers to hosts or remove virtual servers from hosts to meet changing demand.

From the perspective of the users who access these servers to perform their work, it doesn't matter which host a particular virtual server is running on. All that matters is that the virtual server is always available.

Achieving the Benefits of Client Virtualization

Business agility can be significantly increased through being able to work anytime, anywhere, using any device. Today's business users need to be able to access applications and services on a variety of platforms, including desktop computers, laptop computers, netbooks, and mobile phones. Whether the user is sitting at her own desk or at a branch office, at a partner site, at home, in an airport lounge, in a coffee shop, or in her car—if she can't run her applications or access corpnet services, she can't do her work. Microsoft desktop virtualization technologies, such as Remote Desktop Services in Windows Server 2008 R2 and Microsoft VDI, can help make such "anytime, anywhere, any device" access a reality for your business.

Significant costs can be saved by making application deployment and maintenance faster and easier. Traditional application deployment ties the client operating system and installed applications to the hardware. The result is that the more different types of hardware devices the user needs to access, the more work it is for your IT department to deploy and maintain these devices. For example, if a user has a desktop PC, laptop, and mobile phone that are all managed devices on your network, the operating systems for each device must be separately installed and maintained, and the applications for each device must also be separately installed and maintained.

By using virtualization to separate desktop operating systems and applications from the hardware they run on, however, the deployment and maintenance of such operating systems

and applications becomes easier. This can be achieved in two different ways: by virtualizing desktop operating systems or applications as they run on the user's device (an approach called *local desktop virtualization*) or by centralizing the execution of virtual desktop operating systems or applications on servers located in the datacenter (an approach called *remote desktop virtualization*). Microsoft Enterprise Desktop Virtualization, Microsoft Application Virtualization, and Windows Virtual PC and the Windows XP Mode environment are three Microsoft virtualization solutions that can be used to implement local desktop virtualization for your business. And to implement remote desktop virtualization, you can use Remote Desktop Services in Windows Server 2008 R2 alone or as part of a Microsoft VDI solution.

Desktop virtualization not only reduces the effort and cost of deploying and maintaining desktops and applications; it also helps resolve application compatibility issues that can make traditional deployment difficult and time-consuming. By using Microsoft Application Virtualization, for example, you can reduce application-to-application conflicts that can arise when a user needs to run two different versions of the same application but is unable to install both versions locally on the same computer. And for resolving application-to-operating system compatibility issues, you can use a managed solution such as Microsoft Enterprise Desktop Virtualization when many users face such issues. Or if only a few users have such issues, you can use Windows XP Mode and Windows Virtual PC.

Achieving the Benefits of Cloud Virtualization

Cloud computing can significantly increase the agility of your business by enabling you to have capacity on demand. As your business needs and resource requirements change, you can quickly consume either more or less IT infrastructure resources as needed. Such virtualized infrastructure resources can be located either on your premises (by implementing a private cloud solution) or off-premises (by using a public cloud solution implemented by a hosting service provider).

Another benefit from cloud virtualization is that it reduces the management complexity of your IT infrastructure. Your IT department can simply purchase the cloud services it needs, in the amount it needs, whenever the need arises. The result is increased business agility and lower cost. By combining Microsoft server, desktop, and application technologies with the System Center management platform and the automation capabilities of Microsoft's Dynamic Data Center Toolkit, both private enterprises and hosting service providers can implement cloud computing solutions that deliver the benefits of Dynamic IT.

Windows Optimized Desktop Scenarios

Users can benefit in many ways from implementing virtualization technologies within your IT infrastructure. Here are five client computing scenarios common to today's businesses that can directly benefit from virtualization:

- The mobile worker
- The office worker
- The task worker
- The contract or offshore worker
- The worker who needs to access to his applications or data from anywhere

As we briefly examine each of these different client computing scenarios, known as Windows Optimized Desktop Scenarios, you will see how users in your organization can benefit from the implementation of Microsoft's virtualization technologies, platforms, and products.

Mobile Worker Scenario

Businesses are increasingly relying on a mobile workforce to meet the demands of the evolving marketplace, and Microsoft Application Virtualization (App-V) together with Windows roaming desktop technologies (Roaming User Profiles, Folder Redirection and Offline Files) can bring big benefits to both the mobile users themselves and the IT departments that manage them. Specifically, App-V and Windows roaming desktop technologies provide the following end-user benefits for mobile workers:

- A rich user experience that enables users to run multiple applications simultaneously and that, from the user's perspective, feels the same as having these applications installed locally and having data accessible locally on his computer
- Flexible configurations for managing the roaming user data and settings in different ways
- The ability to access applications and data when not connected to the company network

And adding BitLocker Drive Encryption to the formula ensures that the environment for the mobile workforce is a secure, reliable, and efficient option for enterprises to move toward.

From the IT side, implementing App-V and Windows roaming desktop technologies provides the following benefits when managing mobile workers and their computers:

- Safeguarding user data through centralized profile storage, redirection of folders to network file servers, or both

- The ability to migrate user data and settings from previous versions of Windows to Windows 7 by using the User State Migration Tool (USMT)
- The ability to share user data between v1 and v2 user profiles by using Folder Redirection

Office Worker Scenario

In a traditional office worker environment where users have desktop computers, three virtualization technologies (App-V, Remote Desktop Services, and Windows roaming desktop technologies) can provide significant benefits for both users and the IT department.

From the perspective of the office worker, the benefits are similar to two of those for mobile users described in the previous section—namely, a rich user experience and flexible configurations for managing roaming user data and settings. The ability to access applications and data when not connected to the network is not an issue, however, because uninterrupted network connectivity is a defining aspect of this scenario.

From the IT department perspective, implementing these virtualization technologies provides the following direct benefits:

- Simplifies desktop maintenance and application upgrades, and reduces the need to replace PCs
- Simplifies the task of moving users to different desktop computers when they are transferred between departments or locations
- Helps to ensure compliance by enabling sensitive applications to be executed centrally on servers instead of on less secure desktop computers

Task Worker Scenario

Task workers generally need access to only a few task-specific applications to be able to perform their job. Examples of such workers include bank tellers, customer service personnel, shipping/receiving personnel. Remote Desktop Services, Microsoft's presentation virtualization technology, is ideal in this scenario because it can provide users with the specific remote applications they need by using Remote Desktop Services RemoteApps or it can provide them with an entire remote desktop if this is required. Combining Remote Desktop Services with Group Policy in an Active Directory environment allows administrators to lock down functionality presented to users and provide them with a limited, task-oriented user interface that enables them to do their job and nothing else.

IT departments also benefit significantly in this scenario because Remote Desktop Services enables centralized management and enhanced security at a lower cost than other client-computing scenarios. In addition, businesses that want to leverage older hardware can also use Windows Fundamentals for Legacy PCs (WinFLP), which gives organizations the opportunity to extend the life of their older PCs and provides task workers with a cost-efficient, no-frills client device that, when used with Remote Desktop Services, supplies them with all the business application functionality they require.

Contract/Offshore Worker Scenario

Businesses that need to hire outside contractors or offshore developers often cope with having unmanaged, noncorporate PCs connected to their network. Connecting unmanaged PCs to your corporate network can expose your network to possible threats from malware-infected computers. And if access is not carefully controlled, your sensitive business information might also be exposed. Microsoft's recommended virtualization solution in this case is to use Windows Virtual Enterprise Centralized Desktop (VECD). VECD is a unique licensing option of Microsoft's Virtualized Desktop Infrastructure that—when implemented with Hyper-V, App-V, WinFLP, and Vista Roaming Desktop technologies—can provide access to the right applications and data, and nothing else.

IT departments also benefit in this scenario by having centralized management of applications and data. Security and compliance are also facilitated by these technologies. And although this scenario might be a little more complex to implement than, say, setting up a Remote Desktop Session Host server or two, the benefit of having your business secrets safe while allowing lower-cost contract or offshore workers to access your network are clear and compelling for most organizations.

Anywhere-Access Scenario

Sometimes a user needs to access her applications or data but cannot get into the office—for example, during a snowstorm, while away on vacation, or simply when at home during the evening. If your users need “anywhere access” like this from computers not owned by the company, you can provide it for them by using Remote Desktop Gateway, which allows users to access their individual applications or their entire desktop through Remote Desktop Services over the Internet from any computer running Windows 7, Windows Vista with Service Pack 1, or Windows XP with Service Pack 3 or by using Internet Explorer.

The benefits for users in this scenario are clear: secure access to their desktop, applications, and data from anywhere, at anytime. The benefits for IT departments that implement these solutions are also compelling and include centralized management, security, and compliance.

Microsoft's Integrated Virtualization Solution

What really distinguishes Microsoft's vision and strategy for virtualization from that of its competitors is this: instead of providing products that implement only one or two types of virtualization technologies, Microsoft offers businesses a comprehensive and integrated set of virtualization platforms and products that range from the datacenter to the desktop and allows all your IT assets—including both physical and virtual assets—to be easily managed from a single integrated management platform.

Figure 1-4 summarizes how Microsoft's various virtualization platforms, products, and technologies provide a comprehensive, integrated solution to organizations seeking to implement the benefits of virtualization in their IT infrastructures.




FIGURE 1-4 Microsoft's integrated virtualization solution.

The remaining chapters of this book will dig deeper into each of the virtualization technologies, platforms, and products listed in the Figure 1-4. But for now, let's conclude this introductory chapter with a simple question: Why should I use Microsoft's virtualization technologies, platforms, and products instead of those of a competitor? Four answers to this question immediately come to mind.

First, Microsoft is the platform you know and has created the tools you are familiar with. Windows Server 2008 R2 is the latest version of Microsoft's Windows Server operating system, and in addition to features such as Hyper-V, it also includes enhanced versions of features present in previous Windows Server operating systems, such as Remote Desktop Services and Failover Clustering. And because Microsoft virtualization technologies leverage existing product features, such as Active Directory Domain Services (AD DS), the path to learning is clear and training costs are reduced for both IT personnel and end users.

Second, in addition to the integrated suite of virtualization technologies, platforms, and products offered by Microsoft, there also exists a large and thriving partner ecosystem surrounding these technologies, platforms, and products. This clearly gives you another compelling reason to implement Microsoft virtualization solutions instead of those from competing vendors.

Third, Microsoft not only provides you with the technologies to virtualize your servers, desktops, and applications, it also provides you with the tools to manage these technologies, regardless of which components of your infrastructure are physical and which are logical or virtual. Microsoft's management solutions also support interoperability with third-party virtualization solutions from vendors such as VMware, allowing you to manage a cross-hypervisor virtual environment.

Finally, Microsoft virtualization solutions can provide better TCO than any vendor's virtualization solution and the fastest return on investment. With both lower up-front cost and lower ongoing cost in many scenarios, with a comprehensive and integrated set of virtualization products, and with tools that let you manage both virtual and physical computing resources from a single management platform, Microsoft virtualization products and technologies can clearly help you solve the critical technological and business issues facing your business.

Microsoft's Commitment to Virtualization

In conclusion, I remind you of what Steve Ballmer, CEO of Microsoft, said at the Microsoft Management Summit back in 2005 (from "Microsoft CEO Steve Ballmer Affirms Commitment to Dynamic Systems Initiative" on Microsoft PressPass, April 5, 2005):

"We've heard from our enterprise IT customers loud and clear that they need their systems to be more automated and flexible. That's why we're investing in the Dynamic Systems Initiative and areas like virtualization, more secure network access and interoperability—we're committed to helping IT deliver greater efficiency and value."

This year (2009) saw the release of the following Microsoft virtualization products:

- Windows Server 2008 R2 with its enhanced Remote Desktop Services and Hyper-V roles and its support for Cluster Shared Volumes that enables Live Migration

- Microsoft Hyper-V Server 2008 R2
- Windows Virtual PC and Windows XP Mode
- Microsoft Desktop Optimization Pack 2009 R2 with Windows 7 support for Microsoft Enterprise Desktop Virtualization 1.0 and Microsoft Application Virtualization 4.5 SP1

With the release of these products, Microsoft continues to fulfill its commitment to providing an integrated virtualization solution that reaches all the way from the datacenter to the desktop, down to individual applications, and up to the cloud.

Direct from the Source: The Compelling Argument for Virtualization

There's no doubt that virtualization is one of the most compelling technologies, and for good reason. From server consolidation to business continuity to accelerated desktop deployments to Green IT, virtualization is a key enabler for making IT more dynamic, efficient, and agile. However, today a number of barriers are blocking wide-scale adoption of virtualization, including deployment complexities, training, and high costs.

To overcome those barriers, Microsoft is offering virtualization products that cover everything from the datacenter to the desktop, including comprehensive management solutions. The Microsoft approach is to make it as seamless as possible to integrate virtualization into your existing IT environment. On the server, Microsoft Hyper-V and Remote Desktop Services are key features of Windows Server 2008 R2. So if you know Microsoft Windows Server, you'll know virtualization. On the desktop, Microsoft Application Virtualization integrates with Active Directory, making it easy to add to your existing desktop environments while accelerating migrations to Windows 7. And *all* Microsoft virtualization products can be managed with System Center, so you can manage your virtual infrastructure the same way you manage your physical infrastructure.

But it doesn't end here. Microsoft sees virtualization as a strategic investment that will provide a foundation for the next generation of IT innovation, so it's continuing to develop its virtualization products and work closely with its partners so that Microsoft customers have all they need to make virtualization ubiquitous across their entire IT infrastructure. Virtualization is just the beginning of something much bigger.

—David Greschler, Director, Integrated Virtualization Strategy

Additional Resources

The resources listed in this section were either referred to in the chapter or provide additional information concerning concepts and products discussed in this chapter.

General

For a brief history of IBM virtual machines, see <http://www.cap-lore.com/Software/CP.html>.

Bob Muglia's executive e-mail on virtualization technologies can be found at <http://www.microsoft.com/mscorp/execmail/2008/01-21virtualization.mspix>.

Steve Ballmer's commitment to Microsoft's Dynamic Systems Initiative (now called Dynamic IT), which includes virtualization, can be found at <http://www.microsoft.com/presspass/press/2005/Apr05/04-20VirtualizationInvestmentsPR.mspix>.

Microsoft's IT Infrastructure Optimization Model

You can find a description of Microsoft's IT Infrastructure Optimization Model, together with specific recommendations for moving your infrastructure from the Basic stage toward the Dynamic stage, at the Infrastructure Optimization TechCenter on Microsoft TechNet at <http://technet.microsoft.com/en-ca/infrastructure/default.aspx>.

Another good source of information is the Infrastructure Optimization blog on TechNet at <http://blogs.technet.com/io/default.aspx>.

Microsoft's Dynamic IT

Bob Muglia, senior vice president of the Server and Tools Business at Microsoft, outlined Microsoft's strategic vision for Dynamic IT at TechEd 2007. You can read about it on Microsoft PressPass at <http://www.microsoft.com/presspass/press/2007/jun07/06-04TechED07PR.mspix>.

You can also watch videos of Bob Muglia and other senior leaders at Microsoft explaining the benefits of Dynamic IT at <http://www.microsoft.com/presspass/press/2007/jun07/06-04DynamicITVideos.mspix>.

Learn more about how virtualization technologies provide one of the architectural underpinnings of dynamic systems by downloading the virtualization white paper found at <http://www.microsoft.com/business/dsi/virtualization.mspix>.

You can download a white paper in PDF format titled "Optimize and Secure Your Core Infrastructure" from the Microsoft Download Center at <http://download.microsoft.com/download/8/d/8/8d8fd1f8-9af5-4ae9-898d-ffbe130d1ca9/Whitepaper-OSCI-all-up-CoreIO-FY08.pdf>.

Microsoft Virtualization Technologies and Solutions

Get the big picture concerning Microsoft's integrated virtualization vision and learn about Microsoft's various virtualization products and technologies at <http://www.microsoft.com/virtualization/default.aspx>.

Read technical details concerning virtualization and partitioning on the Windows Hardware Developer Central Web site at <http://www.microsoft.com/whdc/system/platform/virtual/default.aspx>.

Stay up to date on the latest virtualization developments from Microsoft by subscribing to the newsfeed of the Microsoft Windows Virtualization Product Group Team Blog at <http://blogs.technet.com/virtualization/default.aspx>.

Windows Optimized Desktop Scenarios

For more information about how Microsoft virtualization technologies can benefit users in different client computing scenarios, see <http://www.microsoft.com/windows/products/windowsvista/enterprise/scenarios.aspx>.

For additional examples of virtualization in action for different types of workers, see <http://www.microsoft.com/virtualization/action.aspx>.

Chapter 2

Server Virtualization

A key benefit of virtualization technologies is server virtualization—the ability to virtualize server workloads. Server virtualization can save businesses money and simplify management overhead by allowing them to reduce the number of physical servers they need through server consolidation. Microsoft offers three business-level server virtualization products. The first is Microsoft Hyper-V, a hardware-assisted virtualization server role that is available in Microsoft Windows Server 2008 R2. The second is Microsoft Hyper-V Server 2008 R2, a free standalone hypervisor-based server virtualization product that lets businesses virtualize workloads onto a single physical server. And the third is Microsoft Virtual Server 2005 R2 SP1, a server virtualization product that runs on Windows Server 2003 or later and can use, but does not require, hardware-assisted virtualization.

This chapter focuses primarily on Hyper-V and describes how Hyper-V works, its key features, and the new features and enhancements found in the R2 release of both the Hyper-V role of Windows Server 2008 and the standalone Hyper-V Server product. The chapter then describes how to deploy and manage Hyper-V and how to use it to create and work with virtual machines. The chapter concludes by summarizing the key benefits of using Hyper-V and looking at common Hyper-V usage scenarios.

Understanding Server Virtualization

Hyper-V is an example of server (or machine) virtualization technology. What this means is that Hyper-V allows you to virtualize entire computers by running multiple operating systems (usually server operating systems) on a single physical computer (typically server-class hardware). Each guest operating system thinks (if operating systems could think) that it owns the computer and has exclusive use of the computer's hardware resources (or to whatever subset of the total machine resources that have been allocated to the virtual machine). Each operating system is therefore said to be running in a separate virtual machine, with these multiple virtual machines running on the same physical computer. In a typical nonvirtualized environment, only one operating system can run on a computer—it's Hyper-V that makes running multiple virtual machines possible. Clearly, we need to dig deeper into the concept of virtual machines before we can understand how Hyper-V works.

Understanding Virtual Machines

A *virtual machine* is a computing environment that is implemented in software and that abstracts the hardware resources of the physical computer so that multiple operating systems can run simultaneously on a single computer. Each operating system runs in its

own virtual machine and is allocated logical instances of the computer's processors, hard disks, network cards, and other hardware resources. An operating system that is running in a virtual machine is unaware that it is executing in a virtual environment and behaves as if it exclusively controls the underlying physical computer's hardware.

Realizing virtual machines as described in the preceding paragraph means that server virtualization must be implemented in a way that meets the following requirements:

- **Management interfaces** Server virtualization requires management interfaces so that administrators can create, configure, and monitor virtual machines running on the computer. These interfaces should also support programmatic administration, and they must be able to work over the network so that virtual machines can be managed remotely.
- **Memory management** Server virtualization requires a memory manager, which ensures that each virtual machine receives its allocation of memory resources and that those memory resources are isolated between each virtual machine.
- **Scheduler** Server virtualization requires a scheduler to manage access to physical resources by different virtual machines. The scheduler must be configurable by the administrator so that different virtual machines can be given different priority to hardware as might be needed.
- **State machine** Server virtualization requires a state machine that can track information concerning the current state of all virtual machines on the computer. State information for a virtual machine includes its CPU, memory, devices, and whether the virtual machine is running or stopped. The state machine must also be designed to manage transitions between different states.
- **Storage and networking** Server virtualization requires functionality that can abstract storage and networking resources on the computer so that each virtual machine is presented with the view that it owns its own exclusive hard disks and network interfaces. In addition, machine virtualization must be able to multiplex access to physical devices in a way that is consistent, isolated, and secure.
- **Virtualized devices** Server virtualization requires virtualized devices that can provide operating systems running in virtual machines with logical representations of devices that behave in a similar manner as their physical counterparts. In other words, when an operating system running in a virtual machine needs to access a physical device on the computer, it does so by accessing a corresponding virtualized device, and this virtualized device is accessed in the same manner as a physical device would be accessed.
- **Virtual device drivers** Server virtualization requires that virtual device drivers be installed on operating systems running in virtual machines. These virtual device drivers enable applications to access the virtual representations of hardware and I/O connections in the same manner that they would access hardware and I/O connections on the underlying physical hardware.

As we'll see in a moment, Microsoft designed Hyper-V, its server virtualization solution, to meet all the above requirements. But first let's examine the key software component that makes server virtualization possible: the hypervisor.

Understanding Hypervisors

A *hypervisor* is a virtualization platform that enables you to run multiple operating systems on a single physical computer called the *host computer*. The main function of the hypervisor is to provide isolated execution environments for each virtual machine and to manage access between the *guest operating systems* running in virtual machines and the underlying hardware resources on the physical computer.

The term “hypervisor” goes way back to 1972 when IBM updated the control program of its System/370 mainframe computing platform to support virtualization. The creation of the hypervisor was a milestone in the evolution of computing because it provided a way to overcome the architectural limitations and high cost of using mainframe computers.

Hypervisors come in several different flavors. They can be categorized, for example, by type—that is, by whether they run directly on the physical hardware or within (hosted by) an operating system environment. Hypervisors can also be categorized by design—that is, whether they are monolithic or microkernel.

Type 1 Hypervisor

Type 1 hypervisors run directly on the underlying physical hardware of the host computers and function as a control program. In other words, they run on bare-metal systems. Guest operating systems then run within multiple virtual machines positioned above the hypervisor layer as shown in Figure 2-1.




FIGURE 2-1 Type 1 hypervisors run directly on bare metal.

Because Type 1 hypervisors run directly on bare metal instead of within an operating system environment, they can generally provide the best performance, availability, and security of

any form of hypervisor. Some examples of server virtualization products that implement Type 1 hypervisors include these:

- Microsoft Hyper-V
- Citrix XenServer
- VMware ESX Server

Type 2 Hypervisor

Type 2 hypervisors run within an operating system environment running on the host computer. Guest operating systems then run within virtual machines above the hypervisor as shown in Figure 2-2. This type of virtualization is typically referred to as *hosted virtualization*.




FIGURE 2-2 Type 2 hypervisors run within an operating system environment.

As you can see by comparing Figure 2-2 with Figure 2-1, guest operating systems running in virtual machines on Type 2 hypervisor platforms are one level further separated from underlying physical hardware than guest operating systems on Type 1 hypervisor platforms. This extra level of separation between the virtual machines and the hardware results in a performance hit being incurred on Type 2 hypervisor platforms, and the effect of this added overhead limits the number of virtual machines you can realistically run on Type 2 platforms.

Examples of server virtualization products that use Type 2 hypervisors include these:

- Microsoft Virtual Server
- VMware Server

The desktop machine virtualization product Microsoft Virtual PC also uses a Type 2 hypervisor architecture.

Monolithic Hypervisor

Monolithic hypervisor design involves using hypervisor-aware device drivers that are hosted within and managed by the hypervisor as shown in Figure 2-3.




FIGURE 2-3 Monolithic hypervisor platforms require hypervisor-aware device drivers.

The monolithic design choice results in some benefits but also some drawbacks. For example, monolithic hypervisors do not need a controlling, or parent, operating system because all guest operating systems interact directly with the underlying physical hardware of the host computer by using hypervisor-aware device drivers. This is an example of the benefit of the monolithic design.

On the other hand, the fact that device drivers must be specifically developed for the hypervisor creates significant challenges because there are so many different types of motherboards, storage controllers, network adapters, and other types of hardware devices on the market. The result is that vendors of monolithic hypervisor platforms have to work closely with manufacturers of hardware devices to ensure these manufacturers develop hypervisor-aware versions of device drivers for their hardware. It also means that vendors of monolithic hypervisor platforms are dependent on manufacturers of hardware devices to supply such drivers for their products. The result is that the number of devices that can be used in virtualized operating system environments running on monolithic hypervisor platforms can be more limited than when those same operating system environments are run directly on physical computers.

One important point is that in this design you're ignoring one of the most important security tenets: defense in depth. With defense in depth, you provide multiple layers of defense to prevent against attacks. In this model, there is no defense in depth because everything is running in the most privileged part of the system.

An example of a server virtualization product that uses a monolithic hypervisor design is VMware ESX Server.

Microkernel Hypervisors

Microkernel hypervisors do not require hypervisor-aware device drivers because they have an operating system acting as the root, or parent, partition. This parent partition then provides the execution environment needed for device drivers to access the underlying physical hardware of the host computer. We'll talk more about partitions in a moment, but for now, simply think of the term "partition" as being equivalent to the previously introduced concept of a virtual machine.

On microkernel hypervisor platforms, you need to install device drivers only for physical devices in the operating system running in the parent partition. You do not need to install these drivers in guest operating systems running in child partitions because when these guest operating systems need to access physical hardware on the host computer, they simply do so by communicating with the parent partition. In other words, in the microkernel design, the guest operating systems do not have direct access to the underlying hardware. They can access physical devices only by communicating with the parent partition. Figure 2-4 shows the microkernel hypervisor design.




FIGURE 2-4 Microkernel hypervisor platforms require that guest operating systems that want to access hardware must do so via the parent partition.

The microkernel hypervisor design has several advantages over the monolithic design. First, because microkernel hypervisors do not need hypervisor-aware drivers, they can immediately use the wide range of existing drivers that are available from device manufacturers. Second, because device drivers are not part of the hypervisor, the hypervisor has less overhead, which means it's smaller and might therefore be more reliable. Third and perhaps more importantly, the attack surface is minimized because foreign code is not loaded in the hypervisor. (Device drivers are manufactured by third parties and are therefore considered to be foreign code from the standpoint of the hypervisor vendor.) After all, the last thing you want to have happen is for malware to infect your hypervisor and thus take control of all the virtual operating systems running on your computer!

The only downside of the microkernel design is that a special partition, the parent partition, is required. This adds measurable (but usually minimal) overhead to your system because of the communication between parent and child partitions that is required to allow the child partitions to access the hardware through the parent.

One great benefit to Hyper-V microkernelized architecture is the use of the defense in depth. Hyper-V has been architected to run as little as possible in its hypervisor and push more functionality up into the stack, such as its state machine and management interfaces, which reside up the stack in user mode.

So what's an example of a server virtualization platform that implements the microkernel design? You guessed it—it's Microsoft Hyper-V, which runs Windows Server 2008 or later in the parent partition.

Understanding the Hyper-V Architecture

Figure 2-5 shows the big picture concerning the Hyper-V architecture.




FIGURE 2-5 Overview of Hyper-V architecture.

As you can see from Figure 2-5, Hyper-V consists of the Microsoft Hypervisor component running on top of bare metal, which demonstrates that Hyper-V is a Type 1 hypervisor platform. Running on top of the hypervisor are one parent partition and one or more child partitions. In virtualization terminology, a *partition* is simply a unit of isolation within the hypervisor that is allocated physical memory address space and virtual processors. There are two types of partitions:

- The *parent partition* is the controlling partition in which the virtualization stack runs. The parent partition is also the partition that owns the hardware devices and manages resources for the child partitions
- A *child partition* is any partition that has been created by the parent partition. Guest operating systems and their applications run in child partitions.

In the Microsoft implementation of the Type 1 hypervisor model—that is, in Hyper-V—the parent partition runs either a Full or Server Core installation of the Standard, Enterprise, or Datacenter edition of Windows Server 2008 or later as its operating system. For more information on what kind of computer you need for running Hyper-V, see the section titled “System Requirements for Using Hyper-V R2” later in this chapter.

Partitions communicate with the hypervisor layer by using *hypercalls*, which are application programming interfaces (APIs) that partitioned operating systems can use to leverage the optimizations that the hypervisor provides. Developers who are interested in learning how to develop applications that use hypercalls can learn more about them in the MSDN Library at <http://msdn.microsoft.com/en-us/library/bb969694.aspx>.

Understanding the Parent Partition

In the Hyper-V implementation of server virtualization, the parent partition includes a number of special components not present in child partitions. Figure 2-6 shows the various components of the parent partition in more detail, including both user-mode (ring 3) and kernel-mode (ring 0) processes.

The parent partition is the first partition created on the system when the hypervisor is started. The parent partition is created for the Windows Server 2008 R2 operating system instance that hosts the Hyper-V server role. The parent partition in Hyper-V serves the following purposes:

- The parent partition is used for creating and managing other (child) partitions on the system and includes the WMI provider, which provides an interface for remote administration.
- The parent partition manages and assigns hardware devices, except for processor scheduling and physical memory allocation, which are handled by the hypervisor.

- The hardware resources of the parent partition are shared or allocated for use by child partitions.
- The parent partition handles power management, plug and play operations, and logging of any hardware failure events when they occur.




FIGURE 2-6 Detailed view of components of parent partition in Hyper-V.

The Virtualization Stack

The virtualization components hosted in the parent partition are referred to collectively as the *virtualization stack*. The virtualization stack runs in the parent partition and has direct access to the hardware of the underlying host computer through mechanisms described later. In Microsoft's Hyper-V implementation of the Type 1 hypervisor model, the virtualization stack consists of the following components:

- Virtual Machine Management Service
- Virtual Machine Worker Process
- Virtual Devices
- Virtualization Infrastructure Driver
- Windows Hypervisor Interface Library

Other components of the parent partition include the following:

- Virtualization Service Providers
- Virtual Machine Bus

The sections that follow examine each of the remaining components of the parent partition in detail.

Virtual Machine Management Service

The Virtual Machine Management Service (VMM Service or VMMS) is responsible for managing the state of all virtual machines in child partitions. This includes managing stopped or offline virtual machines, handling the creation of snapshots, and managing the addition or removal of devices. When a virtual machine in a child partition is started, the VMMS spawns a new Virtual Machine worker process, which is used to perform the management tasks for that virtual machine.

The VMMS also controls which operations can be performed on a virtual machine in a given state. For example, when you are deleting a snapshot of a virtual machine, the VMMS prevents you from applying the snapshot. (See the section titled “Working with Snapshots” later in this chapter for more information concerning snapshots.) Specifically, the VMMS manages the following virtual machine states:

- Starting
- Active
- Not Active
- Taking Snapshot
- Applying Snapshot
- Deleting Snapshot
- Merging Disk

Online virtual machine operations—such as Pause, Save, and Power Off—are not managed by the VMMS. Instead, they are managed by the Virtual Machine worker process that the VMMS spins up for the virtual machine being managed.

The VMMS is implemented in both user mode and kernel mode as a system service (VMMS.exe) and has dependencies on the Remote Procedure Call (RPC) and Windows Management Instrumentation (WMI) services. The VMMS comprises a number of components, one of which is a WMI Provider that exposes a set of WMI-based APIs for managing and controlling virtual machines. These Hyper-V WMI APIs can be used together with Visual Basic Scripting Edition (VBScript) or Windows PowerShell to manage most aspects of a Hyper-V environment either from the command line or by using scripts. The Hyper-V WMI APIs also allow Microsoft System Center products to manage Hyper-V servers. We’ll talk more about how System Center leverages Hyper-V in Chapter 5, “Virtualization Management.”

Virtual Machine Worker Processes

A Virtual Machine worker process (vmwp.exe) is a user-mode process that provides virtual machine management services from the Windows Server 2008 R2 instance in the parent partition to the guest operating systems in the child partitions. The VMMS spawns a separate VM worker process for each running virtual machine to isolate one virtual machine from another. That way, if one VM worker process fails, only the virtual machine associated with that VM worker process is affected. For enhanced security, VM worker processes run under the Network Service built-in identity.

The VM worker process manages the following aspects of its associated virtual machine:

- Creation, configuration, and running of the virtual machine
- Pausing and resuming the virtual machine
- Saving and restoring the virtual machine
- Taking snapshots of the virtual machine

In addition, the VM worker processes contains the Virtual Motherboard (VMB). The VMB exposes guest memory, IRQ generation, and memory-mapped and port-mapped I/O to the virtual machine as separate devices. The VMB is also responsible for the management of virtual devices, which are described next.



Tip You can view the globally unique identifier (GUID) for the virtual machine associated with a particular VM worker process by opening Task Manager, selecting the Processes tab, and adding the Command Line column to the Processes view. This displays each running instance of vmwp.exe on the computer along with its GUID.

Virtual Devices

Virtual Devices (VDevs) are software modules that provide device configuration and control for child partitions. The VMB includes a basic set of VDevs, including a PCI bus and the chipset-level devices found on the Intel 440BX motherboard. VDevs come in two types:

- **Core VDevs** These virtual devices model existing hardware devices and are available to each virtual machine. They are typically used in situations where compatibility is important so that existing software such as the BIOS or device drivers can work properly without needing modifications. Core VDevs can be either of the following:
 - **Emulated devices** These virtual devices emulate a specific hardware device, such as a VESA video card. Most Core VDevs are emulated devices like this, and examples include BIOS, DMA, APIC, ISA Bus, PCI Bus, PIC Device, PIT Device, Power Mgmt device, RTC device, Serial Controller, Speaker device, 8042 PS/2 keyboard/mouse controller, Emulated Ethernet (DEC/Intel 21140), Floppy controller, IDE Controller, and VGA/VESA video.

- ❑ **Synthetic devices** These virtual devices do not model specific hardware devices. Examples of synthetic devices include a synthetic video controller, synthetic Human Interface Device (HID) controller, a synthetic network interface card (synthetic NIC), a synthetic storage devices, synthetic interrupt controller, and memory service routines. These synthetic devices are available only to guest operating systems that support Integration Services, which are discussed later in this section.
- **Plug-in VDevs** These virtual devices do not model existing hardware devices and are used to instantiate, configure, and manage Virtualization Service Providers running in the parent partition, which is the partition that controls the hardware. Plug-in VDevs enable direct communication between the parent and child partitions through the VMBus.

Virtualization Infrastructure Driver

The Virtualization Infrastructure Driver (Vid.sys) is the kernel-mode component of the virtualization stack and provides partition management services, virtual processor management services, and memory management services for all child partitions. The Vid.sys also enables user-mode components of the virtualization stack to communicate with the hypervisor.

Windows Hypervisor Interface Library

The Windows Hypervisor Interface Library (WinHv.sys) is a kernel-mode dynamic-link library (DLL) that loads within the Windows Server 2008 R2 instance running in the parent partition, and within the guest operating system in any child partition where the guest is Hyper-V-aware. WinHv.sys abstracts the hypercall implementation details and enables the operating system's drivers to call the hypervisor by using standard Windows calling conventions.

Virtualization Service Providers

Virtualization Service Providers (VSPs) are hosted in the parent partition and provide a way of publishing device services to child partitions by providing I/O-related resources to Virtualization Service Clients (VSCs) running in child partitions. VSPs are the server endpoint and VSCs are the client endpoint for client/server communications for device functionality. All communications between VSPs and VSCs take place over the VMBus.

Virtual Machine Bus

The Virtual Machine Bus (VMBus) is a logical, channel-based, interpartition communication mechanism between the parent partition and child partitions. The purpose of the VMBus is to provide a high-speed, highly optimized communications mechanism between virtualized partitions rather than other techniques that are slower because of the higher overhead that emulation imposes.

Guest operating systems that do not support Integration Services are not hypervisor-aware and must use emulation. That means that the hypervisor must intervene to intercept calls to the physical hardware from these guests and route them to the emulated device, which runs in the VM worker process in the parent partition. Emulation requires much more overhead for processing than communication using the VMBus, which is why it is recommended that users install the Hyper-V Integration Services after the guest operating system is installed.

The way it works is that the parent partition hosts VSPs, which communicate over the VMBus to handle device access requests from child partitions. Child partitions host VSCs, which redirect device requests to VSPs in the parent partition via the VMBus. The communication process between parent and child partitions over the VMBus is transparent to the guest operating system.

Understanding Child Partitions

As shown in Figure 2-7, the Hyper-V implementation of the Type 1 hypervisor model supports three types of child partitions:

- Child partitions hosting Hyper-V-aware Windows operating systems
- Child partitions hosting Hyper-V-aware non-Windows operating systems
- Child partitions hosting non-Hyper-V-aware operating systems, either Windows or other types




FIGURE 2-7 Different types of child partitions supported by Hyper-V.

Child Partitions Hosting a Hyper-V-Aware Windows Operating System

As shown on the left side of Figure 2-7, child partitions running Windows operating systems that are Hyper-V aware include the following kernel-mode virtualization components:

- **Virtualization Service Clients** VSCs are synthetic devices residing in the child partition that use hardware resources provided by the VSPs in the parent partition by communicating over the VMBus. VSCs are automatically made available for installation when Integration Services are installed in the child partition, which enables the child partition to use synthetic devices. Without Integration Services installed, a child partition can only use emulated devices as shown on the right side of Figure 2-7.
- **Enlightenments** This term refers to modifications made to operating system code to make the operating system hypervisor aware so that it runs more efficiently when it detects that it is running as a guest within a hypervisor environment. Hyper-V supports enlightenment of the following resources: storage, networking, graphics, and input subsystems. An *enlightened guest* is an operating system whose kernel can detect whether or not it is running in a virtualized environment. Windows Server 2008 R2 is an enlightened guest and is therefore a fully enlightened operating system. Windows Vista is an operating system that can reach a degree of enlightenment by installing Integration Services onto it.



Note The VMBus and WinHv.sys components were discussed previously in the section titled “Understanding the Parent Partition” earlier in this chapter.

Child Partitions Hosting a Hyper-V-Aware Non-Windows Operating System

As shown in the middle portion of Figure 2-7, child partitions running non-Windows operating systems that are Hyper-V aware use third-party VSCs to communicate over the VMBus with VSPs in the parent partition in order to access hardware. These VSCs are provided to the child partition by installing Integration Services in the partition.

Integration Services are primarily used to address usability issues that occur because of the isolated environment that is inherent to virtual machines. Integration services also provide the components that allow child partitions to communicate with other partitions and the hypervisor. In previous Microsoft virtualization platforms, such as Microsoft Virtual Server and Microsoft Virtual PC, integration services were referred to as Virtual Machine Additions.

Integration Services also provides the following functionality to the child partition:

- **Heartbeat** Used to verify that the child partition is responding to requests from the parent partition.
- **Key\Value Pair Exchange** Registry key pairs exchanged between child and parent partitions (used in management tools).
- **Time Synchronization** Synchronizes the child partition time with the parent partition.
- **Shutdown** Allows the child partition to respond to shutdown requests from the parent partition.
- **Volume Shadow Copy Service** Works with the VSS component in the parent partition to facilitate data-consistent backups.

Hyper-V includes Integration Services for both x86 and x64 versions of the following Windows guest operating systems:

- Windows XP with Service Pack 3 (SP3)
- Windows Vista with Service Pack 1 (SP1) or later
- Windows 7
- Windows Server 2003 SP2
- Windows Server 2008
- Windows Server 2008 SP2
- Windows Server 2008 R2 (x64 only)

Microsoft has also developed Linux Integration Components For Hyper-V, which can be obtained by searching for "Linux Integration Components for Windows Server 2008 Hyper-V R2" at the Microsoft Download Center at <http://www.microsoft.com/downloads>. These components have been released by Microsoft under the GPL v2 license, and the code is being integrated into the Linux kernel tree via the Linux Driver Project found at <http://www.linuxdriverproject.org>.

Child Partitions Hosting a Non-Hyper-V-Aware Operating System

As shown on the left side of Figure 2-7, child partitions running non-Hyper-V-aware operating systems (whether versions of Microsoft Windows or some third-party operating system) cannot have Integration Services installed on them. This means that these guest operating systems must use emulated devices instead of synthetic devices and suffer the performance hit that is incurred by the use of such emulated devices.

Key Features of Hyper-V

The following are some of the key features of the original release of Microsoft's Hyper-V platform:

- **Broad operating system support** Hyper-V includes broad support for simultaneously running different types of operating systems, including both 32-bit and 64-bit operating systems across different server platforms, such as Windows, Linux, and others.
- **Extensibility** Hyper-V has standards-based Windows Management Instrumentation (WMI) interfaces and application programming interfaces (APIs) to enable independent software vendors (ISVs) and developers to quickly build custom tools, utilities, and enhancements for the virtualization platform.
- **Network Load Balancing** Hyper-V includes virtual switch capabilities that provide the ability to use the Windows Network Load Balancing (NLB) service to load-balance across virtual machines running on different servers.
- **Microkernelized architecture** Hyper-V has a 64-bit microkernelized hypervisor architecture that enables the platform to provide a broad array of device support methods and enhanced performance and security.
- **Hardware-assisted virtualization** Hyper-V requires and uses either Intel-VT or AMD-V hardware-assisted virtualization technologies.
- **Hardware-sharing architecture** Hyper-V includes a Virtualization Service Provider (VSP) and Virtualization Service Client (VSC) architecture, which provides enhanced access and use of hardware resources, such as disks, networking, and video.
- **Quick migration** Hyper-V provides the ability to migrate a running virtual machine from one physical host computer to another with minimal downtime, leveraging the high-availability capabilities of Windows Server 2008 and System Center management tools.
- **Scalability** Hyper-V includes support for multiple processors and cores at the host level and improved memory access within virtual machines. This support enables virtualization environments to be scaled to support a large number of virtual machines on a given host, while continuing to leverage quick migration for scalability across multiple hosts.
- **Symmetric multiprocessor (SMP) support** Hyper-V includes support for up to four processors in a virtual machine environment in order to take advantage of multi-threaded applications running in a virtual machine.
- **Virtual machine snapshots** Hyper-V provides the ability to take snapshots of a running virtual machine to enable you to easily revert to a previous state, thus improving backup and recoverability solutions.

Many of these features are discussed elsewhere in this chapter, but what's really exciting is what has been added to Hyper-V in the R2 release, which is discussed next.

New Features in Hyper-V R2

New features have been added to the Hyper-V role in Windows Server 2008 R2. These new features improve the flexibility, performance and scalability of Hyper-V. Let's explore them now.

Enhanced Flexibility

Hyper-V R2 includes the following new features that can provide enhanced flexibility to how you deploy and maintain your server virtualization infrastructure:

- **Live Migration** Hyper-V R2 introduces Live Migration, which allows you to move a virtual machine from one Hyper-V server to another without dropping the network connection and without any user-perceived downtime or service interruption other than performance slowing for a few seconds. Live Migration can provide high availability to servers and applications running on clustered Hyper-V servers within a virtualized datacenter environment. Live Migration can also simplify the process of upgrading and maintaining host hardware, and it can enable new scenarios such as load balancing virtual machines for maximum power efficiency or optimal processor usage. Live Migration is described in detail in the section titled "Working with Live Migration" later in this chapter.
- **Cluster Shared Volumes** Cluster Shared Volumes is a new feature of Failover Clustering available in Windows Server 2008 R2 that provides a single consistent file namespace so that all nodes in the cluster see the same storage. Cluster Shared Volumes are highly recommended for Live Migration scenarios and are described further in the section titled "Working with Live Migration" later in this chapter.
- **Support for hot adding and hot removal of storage** In the R2 release of Hyper-V, you can now add or remove virtual hard drives (VHDs) and passthrough disks to a running virtual machine without the need of shutting down and restarting the virtual machine. This enables you to reconfigure the disk storage capacity used by a virtual machine in response to changing workloads without the need of any downtime, and it can enable new scenarios for Microsoft SQL Server, Microsoft Exchange Server, and datacenter backup. To use this feature, the VHDs and passthrough disks must be attached to the virtual machine using the virtual SCSI controller. For more information about adding SCSI controllers to virtual machines, see the section "Managing Virtual Machines" later in this chapter.
- **Processor compatibility mode** The new processor compatibility mode available in Hyper-V R2 allows you to migrate a VM between host machines having the same

processor architecture (either AMD or Intel). This makes it easier for you to upgrade your Hyper-V host infrastructure by simplifying the migration of VMs from host machines running on older hardware to host machines running on newer hardware. It also provides more flexibility for the migration of virtual machines between the nodes of a cluster. For example, you can use processor compatibility mode to migrate virtual machines from an Intel Core 2 host to an Intel Pentium 4 host or from an AMD Opteron host to an AMD Athlon host. Note that processor compatibility mode allows you to migrate VMs only between host machines having the same processor architecture. (In other words, it supports both AMD-to-AMD and Intel-to-Intel migrations.) It does not allow you to migrate VMs between host machines running different processor architectures. (In other words, it does not support either AMD-to-Intel or Intel-to-AMD migrations.) To learn more about processor compatibility mode and how to configure it, see the sidebar titled “How It Works: Processor Compatibility Mode” later in this chapter.

Improved Performance

Hyper-V R2 includes the following new features that can improve the performance of your server virtualization infrastructure:

- **Support for up to 384 concurrently running virtual machines with up to 512 virtual processors per server**

If your server hardware is sufficient, you can use Hyper-V R2 to achieve greater levels of server consolidation than ever before. For example, on a single Hyper-V host machine you can run

- ❑ 384 single virtual processor VMs (which is well beneath the 512 virtual processor limit)
- ❑ 256 dual virtual processor VMs (totaling 512 virtual processors)
- ❑ 128 quad virtual processor VMs (again totaling 512 virtual processors)

Or you can run any combination of single, dual, and quad virtual processor VMs as long as your total number of VMs is less than or equal to 384 and the total number of virtual processors assigned to your VMs is less than or equal to 512. This enhanced capability allows Hyper-V R2 to achieve the highest virtual machine density that is currently available in the market at the time of this writing. By comparison, the previous version of Hyper-V in Windows Server 2008 SP2 supported only up to 24 logical processors and up to 192 running virtual machines. Note that when Failover Clustering is used, Hyper-V R2 only supports running up to 64 VMs per cluster node.

- **Support for second-level address translation (SLAT)** In Hyper-V R2, the processor handles address translations across virtual machines instead of the Hyper-V code doing page table remapping in software. This means that SLAT adds a second level of paging below the architectural x86/x64 paging tables found in x86/x64 processors by providing an indirection layer from virtual machine memory access to physical memory access.

With the right processor, such as an Intel processor with Extended Page Tables (EPT, first introduced with i7) or an AMD processor with Nested Page Tables (NPT, which most current AMD processors have), Hyper-V R2 can provide significant performance gains in many scenarios. These gains are a result of the improved memory management and reduction in memory copies needed when these processor features are used, and the gains are especially significant with large working sets (for example, with Microsoft SQL Server). In fact, the memory usage for the Microsoft Hypervisor can shrink from 5 percent to 1 percent of the total physical memory. This means that more memory can be available for child partitions, which in turn can mean higher consolidation ratios.

- **VM Chimney** This feature enables TCP/IP traffic for a virtual machine to be offloaded to a physical network adapter on the host computer. The physical network adapter and operating system must, of course, support TCP Chimney Offload for this to benefit the performance of the virtual machine by reducing the CPU burden on the logical processors. (Support for TCP Chimney Offload was first added to Microsoft Windows in Windows Vista and Windows Server 2008.) VM Chimney is disabled by default and must be enabled on both the physical network adapter and the host operating system. (See <http://support.microsoft.com/kb/951037> for instructions.)

Note that not all applications can benefit from this feature. In particular, applications that use pre-posted buffers and those that use long-lived connections with large data transfers will benefit most from enabling this feature. Note also that physical network adapters that have TCP Chimney Offload-capable hardware support only a fixed number of offloaded connections, and these are shared between all virtual machines running on the host.

- **Support for Virtual Machine Queue (VMQ)** Hyper-V R2 includes support for Virtual Machine Device Queues (VMDq), a feature of Intel Virtualization Technology For Connectivity. VMQ offloads the sorting burden of virtual machine data traffic from the Virtual Machine Manager to the network controller. This enables a physical network adapter to appear as multiple network adapters (queues) on the host, which improves CPU utilization and helps accelerate network throughput, providing better traffic management capabilities to the virtual machine's data traffic. The host machine no longer has device Direct Memory Access (DMA) data in its own buffer because the network adapter can use DMA to direct packets into the virtual machine's memory. The resulting shorter path length for input/output leads to the performance gain. For more information about VMDq, see the Intel Web site at http://www.intel.com/network/connectivity/vtc_vmdq.htm.
- **Support for jumbo frames** Jumbo frames are Ethernet frames that contain more than 1500 bytes of payload. Jumbo frames have been available previously in nonvirtual environments; Hyper-V R2 extends this capability to virtual machines, supporting jumbo frame sizes up to 9014 bytes if the underlying physical network supports this.

The result is improved network throughput and reduced CPU utilization when large file transfers are performed by virtual machines.

Greater Scalability

Hyper-V R2 includes the following new features that can make your server virtualization infrastructure more scalable:

- **Support for up to 64 logical processors in the host processor pool** The number of logical processors supported by Hyper-V in this release has been increased fourfold over the previous, original version of Hyper-V. This allows businesses to take advantage of the latest large, scale-up server systems to maximize their benefits from consolidating their existing server workloads at a higher server density than has previously been possible. In addition, using such server systems lets you more easily provide multiple processors per virtual machine. (Hyper-V supports up to four logical virtual processors per virtual machine.)
- **Support for core parking** Core parking allows Windows and Hyper-V to consolidate processing onto the fewest number of possible processor cores. This is done by suspending inactive processor cores by putting them in deep C state (that is, by “parking” them or putting them to “sleep”). This lets you schedule virtual machines on a single host machine for higher density instead of dispersing them onto multiple hosts. The benefit of doing this is to significantly enhance green IT objectives by allowing you to reduce the power required by CPUs in the hosts in your datacenter.

Comparing Hyper-V and Virtual Server

Because of its advanced capabilities, Hyper-V has already replaced Microsoft Virtual Server for many enterprises that have been using Virtual Server for server consolidation, business continuity, and testing and development. However, Virtual Server can still have its place in the virtualization infrastructure of today’s enterprises. Table 2-1 compares some of the features and specifications of Hyper-V and Virtual Server.

TABLE 2-1 Comparison of Features and Specifications of Virtual Server 2005 R2 SP1 and Hyper-V R2

Feature or Specification	Virtual Server 2005 R2 SP1	Hyper-V R2
Architecture		
Type of virtualization	Hosted	Hypervisor based
Performance/Scalability		
32-bit VMs	Yes	Yes
64-bit VMs	No	Yes

Feature or Specification	Virtual Server 2005 R2 SP1	Hyper-V R2
32-bit hosts	Yes	No
64-bit hosts	No	Yes
Multiprocessor VMs	No	Yes
Maximum guest RAM per VM	3.6 GB	64 GB
Maximum guest CPUs per VM	1	4
Maximum host RAM	256 GB	1 TB
Maximum number of running VMs	64	384
Resource management	Yes	Yes
Availability		
Guest-to-guest failover	Yes	Yes
Host-to-host failover	Yes	Yes
Host migration	Yes	Yes
VM snapshots	No	Yes
Management		
Scriptable/extensible	Yes, COM	Yes, WMI
User interface	Web Interface	MMC 3.0 Interface
SCVMM integration	SCVMM 2007	SCVMM 2008



More Info For more information about the capabilities of Virtual Server and to download it, see <http://www.microsoft.com/windowsserversystem/virtualserver/downloads.aspx>. To learn how to migrate virtual machines from Virtual Server to Hyper-V, see the topic “Virtual Machine Migration Guide: How To Migrate from Virtual Server to Hyper-V” in the TechNet Library at <http://technet.microsoft.com/en-us/library/dd296684.aspx>.

Key Benefits of Using Hyper-V

The benefits of using Hyper-V in business environments of all sizes can be numerous. The following is a quick summary of three key benefits:

- Hyper-V allows you to easily consolidate systems, workloads, and operating environments. For example:
 - You can use Hyper-V to combine multiple workloads and operating systems onto one physical server, thus reducing the costs of hardware and operations.

- ❑ You can use Hyper-V to test versions of software on the hardware that they will later use in production mode, without affecting your production workloads.
- ❑ You can use Hyper-V virtual systems as low-cost test systems without jeopardizing your production workloads.
- ❑ You can run multiple operating system types and releases on a single physical computer, with each virtual system running the operating system that best matches its application and user requirements.
- Hyper-V allows you to optimize use of your computing resources. For example:
 - ❑ Using Hyper-V can allow you to achieve high resource usage by assigning virtual resources such as processors and memory to physical resources through mechanisms such as dispatching and paging. The virtual resources that can be provided in this manner can exceed the physical system resources in both quantity and functionality.
 - ❑ Using Hyper-V can allow you to dynamically share physical resources and resource pools. The result is higher resource usage, particularly for variable workloads where the average needs are much less than those that might be supplied by using an entire dedicated resource.
 - ❑ Because different workloads tend to show peak resource usage at different times of the day and week, implementing multiple workloads in the same physical server using Hyper-V can help you improve system use, cost, and performance.
- Hyper-V can improve the flexibility and responsiveness of your IT infrastructure and thus bring the same kinds of benefits to your business. For example:
 - ❑ Hyper-V can allow service providers to create one virtual system or clone many virtual systems on demand, thus facilitating dynamic resource provisioning.
 - ❑ Hyper-V allows you to implement virtual systems with variable resources to enable the manual or automated management of workload resources.

Hyper-V Usage Scenarios

Four common usage scenarios involving Hyper-V are

- Server Consolidation
- Business Continuity and Disaster Recovery
- Testing and Development
- The Dynamic Datacenter

Server Consolidation

A key use of server or machine virtualization is to help consolidate many servers onto a single system while maintaining isolation between the servers. One of the main benefits of using Hyper-V for this purpose is the lower total cost of ownership (TCO), which is achieved not only by lowering hardware requirements but also by lowering the costs of power, cooling, physical hosting space, network hardware and cabling costs, and hardware maintenance fees. Another benefit of using Hyper-V for this purpose is its ability to integrate 32-bit and 64-bit workloads in the same environment.

Business Continuity and Disaster Recovery

Business continuity is the ability to minimize both scheduled and unscheduled downtime. Hyper-V includes powerful business continuity features, such as live backup and quick migration, that allow businesses to meet stringent uptime and response metrics. Disaster recovery is also a key component of business continuity, and by leveraging the Failover Clustering feature of Windows Server 2008 R2, Hyper-V provides support for disaster recovery within IT environments and across datacenters, even for geographically dispersed clusters.

Testing and Development

Testing and development are important business functions that can leverage the use of virtualization technologies such as Hyper-V. By using virtual machines in place of physical systems, developers can create and test a wide variety of scenarios in an isolated, self-contained environment that closely resembles the behavior of physical systems. With its extensive guest operating system support and checkpoint features, Hyper-V also helps maximize the use of test hardware, which can help reduce development costs, improve software life-cycle management, and improve test coverage.

The Dynamic Datacenter

When integrated with Microsoft System Center, Hyper-V can help you realize the promise of the dynamic datacenter—the vision of self-managing dynamic systems and operational agility. Because Hyper-V includes features such as automated virtual machine reconfiguration, flexible resource control, and quick migration, datacenter administrators can create a dynamic IT environment that employs virtualization not only for responding to problems but also to anticipate increased demands. For more information about Microsoft System Center products and how they can be used with Hyper-V to provide an integrated virtualization solution, see Chapter 5.

Direct from the Source: Virtualizing Server Applications

More and more companies are using virtualization to deploy server applications such as Microsoft SQL Server, Microsoft Exchange Server, and Microsoft Office SharePoint Server, as well as packaged and custom line-of-business (LOB) applications. These customers can recognize significant benefits, including increased resource utilization, enhanced business continuity, and a more efficient management solution.

Microsoft Virtualization provides the best choice for virtualizing Microsoft server applications with recommended deployment scenarios that increase deployment options with one-stop support and that offer a complete management solution at a competitively low cost.

Key Benefits of Virtualizing Server Applications

The benefits your business can achieve by using Microsoft server virtualization solutions include the following:

- **Microsoft Server Applications Built for Windows** The built-in architecture of Windows Server 2008 R2 with Hyper-V eliminates the need for an additional hypervisor purchase, and you also benefit from one-stop support. Features from server applications and Windows Server together help provide increased deployment options. Additionally, the broad network of experienced Microsoft partners can provide support by rapidly responding to your diverse business needs.
- **Complete Management Solution** Microsoft Virtualization offers you a complete, end-to-end management solution. It enables you to optimize your assets and centrally monitor and manage all your physical and virtual resources across multiple hypervisors, from the hardware to the application level. With Microsoft System Center, you will have the ability to patch and deploy your applications with System Center Configuration Manager and access best practices and deep knowledge through System Center Operations Manager management packs. Additionally, System Center delivers a backup and recovery solution with System Center Data Protection Manager and dynamic management of virtualized infrastructure through System Center Virtual Machine Manager and Performance and Resource Optimization (PRO) packs.
- **Low-Cost, Complete Solution** Because a comparable virtualization solution can cost up to six times more than Microsoft Virtualization, organizations worldwide are using Microsoft Virtualization to help reduce costs and deliver greater value. You can simplify your IT infrastructure and processes, resulting in lower ongoing costs by saving on space, power, and resources. Microsoft also offers virtualization-friendly licensing specifically designed for your virtualization needs, allowing you to move virtualized applications to a different server when you need to, without licensing restrictions.

Microsoft Virtualization for SQL Server

Virtualizing Microsoft SQL Server helps you reduce hardware and maintenance costs with a flexible server consolidation solution. You can centralize data services on fewer physical servers and leverage virtualization for high availability. SQL Server can be deployed in a pure virtual or a mixed physical and virtual environment, depending on the flexibility and service levels you require. You can consolidate database servers and improve business continuity by using a combination of Windows Server Hyper-V and SQL Server technologies, such as Live Migration, guest clustering, and database mirroring. You can perform rapid provisioning and scale out Business Intelligence (BI) infrastructure components such as OLAP cubes and reporting services. And you can streamline development, testing, and staging environments for database applications.

Microsoft Virtualization for SharePoint

Virtualizing Microsoft Office SharePoint Server allows your IT department to build and manage scalable Office SharePoint Server farms to provide for an infrastructure that supports collaboration and content management. SharePoint architects can create a deployment model that is reliable and scalable without introducing unnecessary costs or overarchitecting the environment. All SharePoint roles are supported for virtualization; however, certain roles make better virtualization candidates than others. SharePoint Web, Query, and Application roles are the best candidates for virtualization, whereas Index and Database roles can be deployed as physical or virtual roles, depending on resource needs.

Microsoft Virtualization for Exchange Server

For most businesses today, e-mail is the mission-critical communication tool. Virtualization offers IT professionals new options to provide deployment flexibility and realize the benefits in terms of reduced hardware and energy costs. Exchange Client-Access Server (CAS), Hub Transport, and Mailbox roles, which are underutilized, are best candidates for virtualization. Additionally, the Edge Transport role along with other security gateways on the edge server can be considered for virtualization to maximize hardware utilization.

Learn More

To learn more about how you can virtualize Microsoft server applications, see <http://www.microsoft.com/virtualization/solutions/business-critical-applications>. Additional Virtualization resources can be found at <http://www.microsoft.com/virtualization/resources>.

—Vipul Shah, Senior Product Manager, Virtualization

Working with Hyper-V

This section provides a brief overview of how you can deploy, configure, manage, and use Microsoft's Hyper-V platform. Topics are treated in different levels of detail and are intended as an overview only rather than as a comprehensive Hyper-V operations guide. For detailed technical information on implementing and maintaining Hyper-V, see the Hyper-V product and feature information available from the Virtualization TechCenter on Microsoft Technet at <http://technet.microsoft.com/en-us/virtualization/default.aspx>.

Hyper-V Role vs. Microsoft Hyper-V Server

The new "R2" version of Hyper-V is available in two forms:

- The Hyper-V server role, which you can install on systems running Windows Server 2008 R2 Standard, Enterprise, or Datacenter edition.
- Microsoft Hyper-V Server 2008 R2, a standalone, hypervisor-based server virtualization product that lets you virtualize workloads onto a single physical server.

Both products are based on the same hypervisor technology and Windows Server 2008 R2 operating system. Microsoft Hyper-V Server, however, is

- Totally free.
- Has no graphical user interface (think Server Core—you'll have to manage it remotely).
- Has no guest virtualization rights, which means you'll need licenses for any Windows operating systems you run on it.
- Does not support any other server roles. In other words, you can't use it as a DNS server, DHCP server, and so on. You can of course use it to host a virtual machine that runs these roles, but you'll need a license for the guest operating system installed on the virtual machine.

On the face of it, it sounds like Microsoft Hyper-V Server is a bit limited in functionality. However, the R2 version of Microsoft Hyper-V Server also includes support for

- Installing the Failover Clustering feature, which means you can use Microsoft Hyper-V Server systems as nodes in a cluster
- Live Migration, which is another great reason for using Microsoft Hyper-V Server systems as nodes in a cluster

- Up to 8 physical processor sockets with up to 32 logical cores, and up to 1 terabyte of memory on the host system
- Almost every other feature of the Hyper-V role in Windows Server 2008 R2, including virtual machine snapshots, multiple virtual processors per virtual machine, Live Backup support through Volume Shadow Services, and so on

You can even manage Microsoft Hyper-V Server systems using System Center Virtual Machine Manager 2008 R2. And to help you get Microsoft Hyper-V Server up and running, it includes the Server Configuration Utility (SConfig), which is the same initial configuration command-line interface available on the Server Core installation of Windows Server 2008 R2. (See Figure 2-8.)




FIGURE 2-8 The Server Configuration Utility for Microsoft Hyper-V Server 2008 R2.

Note that you cannot upgrade from a previous version of Windows Server to Microsoft Hyper-V Server because Microsoft Hyper-V Server is a separate, standalone product. You also cannot upgrade from Microsoft Hyper-V Server to Windows Server 2008, again because they are different products. However, you can migrate from a Microsoft Hyper-V Server environment to a Windows Server 2008 Hyper-V environment by exporting the virtual machines from Microsoft Hyper-V Server, installing Windows Server 2008, enabling the Hyper-V role, and importing the virtual machines.

Table 2-2 compares the features of Hyper-V Server 2008 R2 with the Hyper-V role of Windows Server 2008 R2 Standard, Enterprise, and Datacenter editions.

TABLE 2-2 Feature Comparison Between Hyper-V Server and the Hyper-V Role

Feature	Hyper-V Server 2008 R2	Windows Server 2008 R2 Standard	Windows Server 2008 R2 Enterprise and Datacenter
Host clustering	✓		✓
Live Migration	✓		✓
Large Memory support (> 32 GB) on the host system	✓		✓
Support for > 4 processors on the host system	✓		✓
Local graphical user interface		✓	✓
Support for adding additional server roles		✓	✓
Guest virtualization rights included in the host server license		✓	✓
Application failover			✓

Hyper-V Server 2008 R2 can be used for almost any server virtualization scenario, including server consolidation, branch server consolidation, testing and development, and mixed operating system (Windows and Linux) virtualization—even as part of a Virtual Desktop Infrastructure (VDI). The only place this product is not really suitable for is the dynamic datacenter, where the more powerful capabilities of Windows Server 2008 R2 Enterprise and Datacenter editions are more suitable.



More Info For more information about Microsoft Hyper-V Server and to download a free copy, see <http://www.microsoft.com/hyper-v-server/en/us/r2.aspx>.

System Requirements for Using Hyper-V R2

The Hyper-V server role is available only in the Standard, Enterprise, and Datacenter editions of Windows Server 2008 R2. In addition to having the correct operating system, the Hyper-V role also has certain requirements concerning the host machine on which it is installed. These requirements include the following:

- Support for *hardware-assisted virtualization*, which is included in the Intel VT and AMD-V line of processors. Processors that support hardware-assisted virtualization include extensions to provide the ability to load a hypervisor virtualization platform in between the computer hardware and the main, or host, operating system.

- Support for hardware-based *Data Execution Prevention (DEP)*, a security feature that prevents a process from executing code from a nonexecutable memory region. Although an implementation of DEP can be hardware based, software based, or a combination of the two, support for hardware-based DEP is required in order to use Hyper-V. Hardware-based DEP requires processors that can mark memory pages as nonexecutable. Examples of processors that support hardware-based DEP include the Intel XD (Execute Disable) and AMD NX (No-Execute) lines of processors.
- Sufficient physical memory (RAM) to allow you to run the virtualized workloads you plan to run on the system. The Standard edition of Windows Server 2008 R2 supports up to 32 GB of RAM and up to four x64 processor sockets. The Enterprise edition supports up to 2 TB of RAM and up to eight x64 processor sockets. The Datacenter edition supports up to 2 TB of RAM and up to 64 x64 processor sockets. In addition, each virtual machine on the Enterprise or Datacenter edition can address up to 64 GB and the sum of the memory assigned to virtual machines cannot exceed the system's physical RAM minus 1 GB allocated for the parent partition.



Note Hyper-V is not supported on Itanium versions of Windows Server 2008 R2 and Windows Server 2008.

Direct from the Source: Hardware Assisted Virtualization

Since the introduction of the 80286 CPU, operating system architectures have supported four modes of execution called *rings* (for example, ring 0 – 3). Ring 0 is the most privileged mode, and components running in ring 0 have direct access to the underlying hardware. Ring 3 is the least-privileged mode, and operations to modify the hardware are generally not allowed in this ring. Windows historically has used only ring 0 (for kernel-mode components) and ring 3 (for user-mode components).

Virtual Server uses ring compression, or ring depriving, so the Virtual Machine Manager (VMM) can control the execution of a guest operating system in a virtual machine. With this design, kernel-mode operations in a virtual machine are performed in ring 1. Most privileged operations issued by the kernel in a guest operating system result in a transition to the VMM to interact with the underlying hardware. Because it is unlikely that the guest operating system running in a virtual machine is aware of the VMM, virtual machine additions are implemented to facilitate the VMM transitions for these operating systems.

These VMM transitions are expensive in terms of CPU cycles, and therefore, they affect system performance. To overcome this limitation, Intel and AMD have implemented extensions to the classical four-ring architecture to provide an additional level, often

called ring -1, for a VMM to execute. This allows virtual guest operating systems' kernels to run at ring 0 and invoke the VMM in ring -1 for critical operations with much less overhead. The hypervisor is synonymous with the VMM in this context. The basic concept is that components running in the new ring -1 can control components running in ring 0. These extensions also implement extended page tables and tagged Translation Lookaside Buffers (TLBs) to support the isolation of virtual machines.

The Microsoft implementation of a hypervisor requires these processor virtualization extensions. The extensions are currently implemented in the Intel VT and AMD-V lines of processors.

—CSS Global Technical Readiness (GTR) team

Supported Guest Operating Systems

Below is a list of operating systems that are supported at the time of this writing for use as guest operating systems running in virtual machines on Hyper-V.

The following editions of Windows Server 2008 and Windows Server 2008 R2 can be used as a supported guest operating system on a virtual machine configured with from one to four virtual processors:

- Windows Server 2008 R2 Standard, Windows Server 2008 Standard, and Windows Server 2008 Standard without Hyper-V
- Windows Server 2008 R2 Enterprise, Windows Server 2008 Enterprise, and Windows Server 2008 Enterprise without Hyper-V
- Windows Server 2008 R2 Datacenter, Windows Server 2008 Datacenter, and Windows Server 2008 Datacenter without Hyper-V
- Windows Web Server 2008 R2 and Windows Web Server 2008
- Windows High Performance Computing (HPC) Server 2008 edition

The following editions of Windows Server 2003 can be used as a supported guest operating system on a virtual machine configured with either one or two virtual processors:

- Windows Server 2003 R2 Standard edition with Service Pack 2
- Windows Server 2003 R2 Enterprise edition with Service Pack 2
- Windows Server 2003 R2 Datacenter edition with Service Pack 2
- Windows Server 2003 Standard edition with Service Pack 2
- Windows Server 2003 Enterprise edition with Service Pack 2
- Windows Server 2003 Datacenter edition with Service Pack 2

- Windows Server 2003 Web edition with Service Pack 2
- Windows Server 2003 R2 Standard x64 edition with Service Pack 2
- Windows Server 2003 R2 Enterprise x64 edition with Service Pack 2
- Windows Server 2003 R2 Datacenter x64 edition with Service Pack 2
- Windows Server 2003 Standard x64 edition with Service Pack 2
- Windows Server 2003 Enterprise x64 edition with Service Pack 2
- Windows Server 2003 Datacenter x64 edition with Service Pack 2

The following versions of Windows 2000 can be run on a virtual machine configured with one virtual processor:

- Windows 2000 Server with Service Pack 4
- Windows 2000 Advanced Server with Service Pack 4

The following versions of Windows 7 can be used on a virtual machine configured with either one or two virtual processors:

- Windows 7 Professional
- Windows 7 Enterprise
- Windows 7 Ultimate

The following versions of Windows Vista can be used on a virtual machine configured with either one or two virtual processors:

- Windows Vista Business with Service Pack 1 or later
- Windows Vista Enterprise with Service Pack 1 or later
- Windows Vista Ultimate with Service Pack 1 or later

The following versions of Windows XP can be run on a virtual machine as specified:

- Windows XP Professional with Service Pack 3 (configured with one or two virtual processors)
- Windows XP Professional with Service Pack 2 (configured with one virtual processor)
- Windows XP Professional x64 edition with Service Pack 2 (configured with one or two virtual processors)

The following Linux distributions can be run on a virtual machine configured with one virtual processor:

- Suse Linux Enterprise Server 10 with Service Pack 2 (x86 or x64 edition)
- Suse Linux Enterprise Server 10 with Service Pack 1 (x86 or x64 edition)

- Suse Linux Enterprise Server 11 (x86 or x64 edition)
- Red Hat Enterprise Linux (RHEL) 5.2 and 5.3 (x86 Edition or x64 Edition) (Emulated devices only)



Note You can run both 32-bit and 64-bit guest operating systems at the same time on a single server running Hyper-V.

Functionality Provided by Integration Services

Although Hyper-V comes with Integration Services for all supported guest operating systems (with the exception of Linux Integration Components For Windows Server 2008 Hyper-V, which is provided out-of-band as a download), not all guests receive the same usability and performance enhancements from these Integration Services. Table 2-3 describes the enhancements provided by Integration Services for each supported guest.

TABLE 2-3 Usability and Performance Enhancements Provided by Integration Services for Different Guest Operating Systems

Guest Operating System	Device and Service Support
Windows Server 2008 R2	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchroni- zation, data exchange, heartbeat, and online backup
Windows Server 2008 (x64 editions)	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchroni- zation, data exchange, heartbeat, and online backup
Windows Server 2008 (x86 editions)	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchroni- zation, data exchange, heartbeat, and online backup
Windows Server 2003 (x64 editions) with Service Pack 2	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchroni- zation, data exchange, heartbeat, and online backup
Windows Server 2003 (x86 editions) with Service Pack 2	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchroni- zation, data exchange, heartbeat, and online backup
Windows 2000 Server with Service Pack 4	Drivers: IDE, networking, video, and mouse Services: Operating system shutdown, time synchroni- zation, data exchange, heartbeat, and online backup
Windows 2000 Advanced Server with Service Pack 4	Drivers: IDE, networking, video, and mouse Services: Operating system shutdown, time synchroni- zation, data exchange, heartbeat, and online backup

Guest Operating System	Device and Service Support
Windows 7 (x64 editions)	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchronization, data exchange, heartbeat, and online backup
Windows 7 (x86 editions)	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchronization, data exchange, heartbeat, and online backup
Windows Vista (x64 editions) with Service Pack 1 or later	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchronization, data exchange, heartbeat, and online backup
Windows Vista (x86 editions) with Service Pack 1 or later	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchronization, data exchange, heartbeat, and online backup
Windows XP Professional (x86 editions) with Service Pack 2 or 3	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchronization, data exchange, and heartbeat
Windows XP Professional x64 edition with Service Pack 2	Drivers: IDE, SCSI, networking, video, and mouse Services: Operating system shutdown, time synchronization, data exchange, and heartbeat
Suse Linux Enterprise Server 10 (x64 edition) with Service Pack 1 or 2	Drivers only: IDE, SCSI, networking, and mouse
Suse Linux Enterprise Server 10 (x86 edition) with Service Pack 1 or 2	Drivers only: IDE, SCSI, networking, and mouse

Planning for Hyper-V Deployment

The Windows Server Virtualization Guide can help you design and plan for your Hyper-V deployment to ensure its success. This guide is part of Microsoft's Infrastructure Planning and Design (IPD) Guides for Virtualization, which help streamline and clarify the process for designing a virtualization infrastructure by describing the critical architectural decisions that need to be addressed and the available options for these decisions. The guides also provide you with tools for validating your design decisions to ensure the solutions you are planning to implement meet both your business requirements and the requirements of your organization's IT stakeholders.

Version 2.0 of the Windows Server Virtualization Guide explains the critical infrastructure design elements necessary for a successful implementation of a server virtualization solution using Windows Server 2008 R2. The guide leads you through the following nine steps of the server virtualization design process:

Step 1: Determine the Virtualization Scope

Step 2: Create the List of Workloads

Step 3: Select the Backup and Fault-Tolerance Approaches for Each Workload

Step 4: Summarize and Analyze the Workload Requirements

Step 5: Design and Place Virtualization Host Hardware

Step 6: Map Workloads to Hosts

Step 7: Design Backup and Fault Tolerance

Step 8: Design the Storage Infrastructure

Step 9: Design the Network Infrastructure



More Information For more information about the Windows Server Virtualization Guide and to download this guide, go to <http://technet.microsoft.com/en-ca/library/bb897507.aspx>.

Installing the Hyper-V Role

Hyper-V is implemented as a server role on both Full and Server Core installations of Windows Server 2008 R2 Standard, Enterprise, and Datacenter editions. You can install Hyper-V on a Full installation of Windows Server 2008 R2 using the following methods:

- By launching the Add Roles Wizard from the Initial Configuration Tasks (ICT) interface
- By launching the Add Roles Wizard from the Server Manager MMC snap-in
- By using the ServerManagerCmd.exe command-line tool

To install Hyper-V on a Server Core installation of Windows Server 2008 R2, you must use the Ocsetup.exe utility by typing **start /w ocsetup Microsoft-Hyper-V** at the command prompt.

After you install Hyper-V, you must restart your computer before the role can take effect. If you discover that Hyper-V does not start properly after performing one of the listed procedures, try shutting down your computer completely and performing another cold boot of the system. If Hyper-V still fails to start, follow these steps to troubleshoot the problem:

- Verify with the manufacturer that the processors in your system support both hardware-assisted virtualization and hardware Data Execution Prevention.
- Make sure that hardware-assisted virtualization is enabled in the BIOS. If it isn't, enable it and then shut down your computer before rebooting it so that the BIOS change can take effect.
- Check the manufacturer's Web site to see whether an updated version of the BIOS is available for your computer; install the update if one is available.

- Verify that the BCD store is configured properly by typing **bcdedit /enum** and verifying that HypervisorLaunchType is set to AUTO.

For more information about deploying Hyper-V, see the “Hyper-V Planning and Deployment Guide,” which is available from the TechNet Library at <http://technet.microsoft.com/en-us/library/cc794762.aspx>.

Direct from the Source: Considerations for Physical Servers Hosting the Hyper-V Role

Before setting up a physical server to host the Hyper-V role, download, read, and understand information included in the white paper “Performance Tuning Guidelines for Windows Server 2008” available at http://www.microsoft.com/whdc/system/sysperf/Perf_tun_srv.msp. Three sections in this white paper that can have a significant impact on the performance of the physical server discuss tuning the server hardware and setting up the networking and storage subsystems. These are especially critical for Hyper-V because the hypervisor itself sits on top of the hardware layer as described earlier and controls all hardware in Windows Server 2008. The operating system itself essentially runs in a virtual machine, better known as the Parent Partition.

Best practices for physical servers hosting the Hyper-V role are described in the following sections of this sidebar.

Avoid Overloading the Server

Determining the number of virtual machines that will be hosted on the Hyper-V server and the workloads they will be handling is critical. The version of the operating system that will be installed on the physical server can help in this regard, so the first “best practice” is to consider using Windows Server 2008 Datacenter x64 with Hyper-V. The Datacenter x64 edition supports up to 64 processors, 2 terabytes of physical memory, and 16 failover cluster nodes for Quick Migration scenarios and allows unlimited virtual machines to be run in Hyper-V. Selecting a Server Core installation provides added benefits, including enhanced security and lower maintenance.

Ensure High-Speed Access to Storage

For storage, consider using a storage area network (SAN) that is configured with high-speed (10,000 rpms or greater) drives (SATA or SAS) that support queued I/O and Raid 0 +1 configurations. You can use either Fibre Channel or iSCSI SAN hardware.

Install Multiple Network Interface Cards

For networking, be sure to have more than one network card installed on the physical server and dedicate one network interface to Hyper-V server administration. This means no virtual networks in Hyper-V will be configured to use this NIC. For high-workload virtual machines, you might want to dedicate a physical network adapter on

the server to the virtual network the virtual machine is using. Ensure virtual machines that share a physical adapter do not oversubscribe to the physical network. Use the Reliability And Performance Monitor to establish a performance baseline for the load and then adjust NIC configurations and loads accordingly.

If you have only a single NIC in the machine that you are configuring the Hyper-V role on and you are doing the configuration remotely—say, in an RDP session—if you choose to bind the Virtual Switch Protocol to the single NIC in the machine, you will be disconnected from your session and a reconnection might not be possible until the newly created virtual network adapter has been properly configured.

Avoid Mixing Virtual Machines That Can Use Integration Services with Those That Cannot

Do not mix on the same physical server virtual machines that can take advantage of Hyper-V Integration Services with those that cannot. Virtual machines that cannot use Integration Services must use legacy network adapters to gain access to the physical network. To accommodate legacy network adapters, you might need to disable some high-end features on the network interface, which can unnecessarily limit the functionality of the synthetic devices. Additionally, using emulated devices places an extra workload on the Hyper-V server.

Configure Antivirus Software to Bypass Hyper-V Processes and Directories

If you are running antivirus software on the physical server, you might want to consider excluding the Vmms.exe and Vmswp.exe processes. Also, exclude the directories that contain the virtual machine configuration files and virtual hard disks from active scanning. An added benefit of using passthrough disks in your virtual machines is that you can use the antivirus software running on the physical server to protect that virtual machine.

Avoid Storing System Files on Drives Used for Hyper-V Storage

Do not store any system files (Pagefile.sys) on drives dedicated to storing virtual machine data.

Monitor Performance to Optimize and Manage Server Loading

When running multiple high-workload virtual machines on a Hyper-V server, ensure a proper aggregate performance baseline is obtained over a specified period of time—say, five days during normal working hours—to ensure the hardware configuration for the physical server is optimal to support the load being placed on it by the virtual machines. If adding more memory, processors, or higher performing storage is not possible, you might need to migrate the virtual machines to other Hyper-V servers.

—CSS Global Technical Readiness (GTR) team

Using the Hyper-V Management Snap-in

When you install the Hyper-V role on a Full installation of Windows Server 2008 R2, the Hyper-V Manager snap-in is also installed and is available from both the Administrative Tools menu as a separate MMC console (Virtmgmt.msc) and from within the Server Manager console. You can also install the Hyper-V Manager console on the following systems:

- A computer running a Full installation of Windows Server 2008 R2 that does not have the Hyper-V role installed. Install Hyper-V Tools from the Remote Administration Tools section of the Remote Server Administration Tools (RSAT) feature by using the Add Features Wizard.
- A computer running Windows 7, Windows Server 2008, Windows Server 2003, or Windows 7 Professional, Enterprise, or Ultimate edition on which you have downloaded and installed the Remote Server Administration Tools (RSAT) for Windows 7. RSAT for Windows 7 lets you manage roles and features that are installed on remote computers that are running Windows Server 2008 R2 (and, for some roles and features, Windows Server 2008 or Windows Server 2003) from a remote computer that is running Windows 7. RSAT for Windows 7 includes support for remote management of computers that are running either the Full or Server Core installation options of Windows Server 2008 R2, and for some roles and features, Windows Server 2008. Some roles and features on Windows Server 2003 can also be managed remotely by using RSAT for Windows 7, although the Server Core installation option is not available with the Windows Server 2003 operating system. You can obtain RSAT for Windows 7 from the Microsoft Download Center.

Figure 2-9 shows the Hyper-V Manager snap-in remotely connected to a Hyper-V server named SEA-SCV running Server Core, which can be used to manage both the Hyper-V server (parent partition) itself and the virtual machines that are running in child partitions on the server. Note that no virtual machines (child partitions) have been created yet on this server. You can connect to more Hyper-V servers by right-clicking on the root node in the left pane to bring up the Select Computer dialog box. This enables you to use the Hyper-V Manager snap-in running from a single computer running Windows Server 2008 R2 (or from a Windows 7 computer that has RSAT for Windows 7 installed) and use it to manage multiple Hyper-V servers remotely.



Tip If you plan to use a Server Core installation of Windows Server 2008 as your Hyper-V platform, you must manage the Hyper-V role remotely because the Hyper-V Manager snap-in is not available on Server Core. However, the R2 version of Server Core does include a new command-line interface tool, named SCONFIG, that simplifies the initial configuration of a Server Core installation, making it easier to use Server Core as a Hyper-V host. For more information on SCONFIG, see “Windows Server 2008 R2 Core: Introducing SCONFIG” on the Microsoft Virtualization Team Blog at <http://blogs.technet.com/virtualization/archive/2009/07/07/windows-server-2008-r2-core-introducing-sconfig.aspx>.



FIGURE 2-9 The Hyper-V Manager snap-in.

Configuring Server Settings

As shown in Figure 2-9, when a Hyper-V server is selected in the left pane, a series of configuration actions is displayed in the Actions pane on the right side. These configuration actions can also be selected from a context menu by right-clicking on the server node in the left pane. Some of these server-level management actions are self-explanatory (such as Stop Service or Remove Server). Others require some degree of explanation as follows:

- **New | Virtual Machine** Selecting this action starts the New Virtual Machine Wizard, which steps you through creating a new virtual machine (child partition) on the host computer. The steps for creating a new virtual machine are as follows:
 1. Name the virtual machine.
 2. Select a location to store the virtual machine configuration file.
 3. Assign memory.
 4. Configure networking.

5. Configure storage.
 6. Install an operating system.
 7. Start the virtual machine (optional).
- **New | Hard Disk** Selecting this action starts the New Virtual Hard Disk Wizard, which steps you through creating a new virtual hard disk. The steps involved are as follows:
 1. Choose the type of disk (Dynamically Expanding, Fixed Size, Differencing).
 2. Name and choose a storage location for the disk.
 3. Configure the disk (including specifying the size of disk or copying the contents of an existing disk).
 - **New | Floppy Disk** Selecting this action creates a 1.4-MB virtual floppy disk (.vfd) file in the location specified.
 - **Import Virtual Machine** Selecting this action enables you to import previously exported virtual machines by pointing to the appropriate virtual hard disk (.vhd) file. Selecting this option allows you to move virtual machines between different Hyper-V servers.



Note The Import Virtual Machine action does *not* let you import virtual machines created in Virtual Server 2005.

- **Hyper-V Settings** Selecting this action allows you to configure both Server and User settings as shown in Figure 2-10.

There are two Server settings you can configure:

- **Default Location for Virtual Hard Disks** Selecting this option enables you to specify the default location for virtual hard disk (.vhd) files to be used by the virtual machines running on the server. The default location is the C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks directory.
- **Default Location for Virtual Machine Configuration Files** Selecting this option enables you to specify the default location for the virtual machine configuration (.vmc) files on the server. The default location is the C:\ProgramData\Microsoft\Windows\Hyper-V directory.



Tip For the best performance, move the location for virtual hard disks to a non-system drive and move the configuration files to the same location.




FIGURE 2-10 Configuring Server and User settings for a Hyper-V server.

There are five User settings that apply to users managing the Hyper-V server, the virtual machines running on the server, or both. These settings are as follows:

- ❑ **Keyboard** This option allows you to specify the default setting for how keyboard combinations such as Alt+Tab will be used.
- ❑ **Mouse Release Key** This option allows you to specify the setting for the key-stroke sequence used to release the mouse from inside a virtual machine. The default is Ctrl+Alt+Left Arrow, and you can change this to Ctrl+Alt+Right Arrow, Ctrl+Alt+Space, or Ctrl+Alt+Shift as desired. Note that this configuration option does not apply to guests that have had Integration Services installed on them.
- ❑ **User Credentials** This option allows you to specify the credentials that will be used to connect to a virtual machine. The default is to use the same credentials you used to run the Hyper-V Manager snap-in; otherwise, you will be prompted when connecting.
- ❑ **Delete Saved Credentials** This option allows you to delete any saved user credentials on the Hyper-V server for enhanced security.

- ❑ **Reset Check Boxes** This option allows you to reset all Hyper-V confirmation messages and wizard pages to their defaults when the Hyper-V role was installed.
- **Virtual Network Manager** Selecting this action enables you to configure virtual networking settings for the virtual machines running on the server. As illustrated by Figure 2-11, there are three types of virtual networks you can configure:
 - ❑ **External virtual networks** This type of virtual network binds to a physical network adapter on the host computer. An external network is required in order to access the Internet or to connect to organizational resources that do not reside in the parent partition. You can bind only one external network per physical adapter or port, so if multiple external networks are needed, additional physical adapters or ports will have to be installed in the Hyper-V server. VLAN access is also supported on external virtual networks if the physical network the parent partition is connected to is properly configured. You should use the external type of virtual network when you want to allow communication between different virtual machines running on the same host computer, between virtual machines and the parent partition, and between virtual machines and externally located (physical or virtual) servers.
 - ❑ **Internal virtual networks** This type of virtual network is an external virtual network that is not bound to a physical network adapter. You should use this type of virtual network when you want to allow communication between different virtual machines running on the same host computer and between virtual machines and the parent partition, but not between virtual machines and externally located servers. You can also isolate virtual machines on an internal virtual network by selecting the Enable Virtual LAN Identification For Parent Partition check box for a particular adapter or port. A common use for internal virtual networks is to build test environments where you need to connect to the virtual machines from the parent partition.
 - ❑ **Private virtual networks** This type of virtual network is an internal virtual network without a virtual network adapter in the parent partition. You should use this type of virtual network when you want to allow communication only between different virtual machines running on the same host computer. A typical use for private virtual networks is when you want to isolate virtual machines from network traffic in the parent partition and in the external networks.




FIGURE 2-11 Configuring a virtual network.

New in the R2 release of Hyper-V is the setting Allow Management Operating System To Share This Network, which is shown in Figure 2-11. This new configuration option can be enabled when the virtual network type is External, and it determines whether the specified physical network adapter can be used to access the management (host) operating system that runs the Hyper-V role. You can use this option to isolate the management operating system from communications between virtual machines and other computers on your physical network. Note that if this option is not enabled, you will not be able to connect to the management operating system remotely through this physical network adapter.

Also new in Hyper-V R2 is the MAC Address Range option found under Global Network Settings in the Virtual Network Manager settings page. (See Figure 2-12.) This new configuration option has been added to help prevent MAC address conflicts when multiple Hyper-V servers are hosting virtual machines on the same physical network subnet. By default, the MAC address range option is dynamically configured, but you can also configure this setting manually if desired.




FIGURE 2-12 Configuring the MAC address range for virtual machines hosted by the Hyper-V server.

Table 2-4 summarizes the different types of virtual networks you can configure and the connectivity allowed by each type. For more information about how networking works in Hyper-V, see the sidebar titled “Direct from the Source: The Hyper-V Networking Model” later in this section.

TABLE 2-4 Connectivity Allowed by Different Types of Virtual Networks

Type of Virtual Network	Between VMs on the Host Computer	Between VMs and the Parent Partition	Between VMs and External Servers
External	✓	✓	✓
Internal	✓	✓	
Private	✓		




Note Microsoft recommends that you have at least two physical network adapters on the host computer running Hyper-V: one network adapter dedicated to the physical machine and used for remote management, and one or more network adapters dedicated to the virtual machines.

- **Edit Disk** Selecting this action starts the Edit Virtual Hard Disk Wizard, which enables you to make changes to existing virtual hard disks. The actions you can perform include the following:
 - ❑ **Compact** This option allows you to shrink a disk by removing blank space that remains when data is deleted from the disk.
 - ❑ **Convert** This option allows you to convert a dynamic virtual hard disk to a fixed hard disk by creating a new fixed-size virtual disk having a different name and then copying the contents of the dynamic disk to the new fixed disk. The new fixed disk can then be associated with the virtual machine, and the virtual machine can be started, after which the old dynamic disk can be deleted.
 - ❑ **Expand** This option allows you to expand the capacity of a virtual hard disk.
 - ❑ **Inspect Disk** Selecting this option displays information about a virtual hard disk.

Direct from the Source: The Hyper-V Networking Model


One of the more challenging concepts to grasp in Hyper-V is how networking is implemented both from the perspective of the Hyper-V server, which is essentially a virtual machine itself called the *parent partition*, and also from the perspective of the child partitions created and supported by the parent partition.

In general, when the Hyper-V role is installed and prior to creating any virtual networks, the parent partition is functioning as seen in this diagram:



In the configuration just shown, all protocols are bound to the physical network card, which provides direct connectivity for the Hyper-V server to the physical network.

After a virtual network is created and configured to be used by virtual machines (child partitions) running on the Hyper-V server, the networking model shifts as seen the diagram that follows.



—CSS Global Technical Readiness (GTR) team

Managing Virtual Machines

In addition to using the Hyper-V Manager snap-in to manage the Hyper-V server, you can also use this snap-in to manage various aspects of virtual machines running on the server. Figure 2-13 shows a Hyper-V server with a new virtual machine that has just been created on it. This virtual machine has not been started and has no operating system installed on it.

As you can see from Figure 2-13, the Actions pane now has two sections: an upper section named after the Hyper-V server (SEA-SCV) that displays the server-level settings that we have already described, and a lower section named after the new virtual machine (SEA-SRV1-V), which provides options you can select for managing various aspects of the virtual machine. If your Hyper-V Manager console is connected to additional Hyper-V servers and/or the servers have additional virtual machines, more sections are displayed in the Actions pane

accordingly. Some of these VM-level management actions are self-explanatory (such as Start or Delete). Others require more explanation, as follows:

- **Connect** Let's you connect to and manage an individual virtual machine using the Virtual Machine Connection interface. There are two ways to connect to a virtual machine whether it is running or not:
 - By selecting the virtual machine in the center pane and then selecting the Connect action.
 - By double-clicking on the virtual machine in the center pane.




FIGURE 2-13 Hyper-V server with a new virtual machine created.



More Info For more information about connecting to virtual machines, see the section titled “Using the Virtual Machine Connection Tool” later in this chapter.

- **Settings** Selecting this action allows you to configure or modify different settings that apply to the selected virtual machine. These virtual machine settings fall into two categories: Hardware and Management settings.

The hardware settings you can configure for each virtual machine include the following:

- ❑ **Add Hardware** Allows SCSI, network adapters, and legacy network adapters to be added to a virtual machine. (See Figure 2-14.) Each virtual machine can be configured with up to 12 virtual network adapters where eight of these can be of the “network adapter” type and four can be of the “legacy network adapter” type. The “network adapter” type provides better performance it’s a synthetic device that takes advantage of the new high-speed Hyper-V I/O architecture. For more information about network adapters and legacy network adapters, see the sidebar titled “Direct from the Source: Network Adapters in Hyper-V” later in this section.
- ❑ **BIOS** Allows changes to be made to the virtual machine’s BIOS, such as changing the Numlock status or the boot order.
- ❑ **Memory** Allows changes to be made to the amount of physical RAM allocated to the virtual machine.
- ❑ **Processor** Allows changes to be made to the number of virtual processors allocated to the virtual machine and the host physical processor resources allocated to virtual machines.
- ❑ **IDE Controller** Allows you to add hard drives and DVD drives to the virtual machine. There are two IDE controllers, and two drive types can be supported per controller. Note that Hyper-V virtual machines can boot only from IDE drives, not from SCSI drives.
- ❑ **SCSI Controller** Allows you to add up to four SCSI controllers with up to 64 disks each, for a total of 256 drives, per virtual machine.
- ❑ **Network Adapter** Displays the configuration for the virtual network the virtual machine is connected to. You can also configure network adapter MAC addresses and enable VLAN IDs.
- ❑ **COM 1\2** Allows for communication with the parent partition via a named pipe connection. This can be useful for debugging a virtual machine.
- ❑ **Diskette Drive** Provides a connection to a 1.4-MB floppy disk created as a .vfd file.




FIGURE 2-14 Configuring the Add Hardware settings for a virtual machine.

In addition to configuring the hardware settings just described, you can also use the Hyper-V Manager snap-in to configure different management settings for each virtual machine as follows:

- ❑ **Name** Allows the name of the virtual machine to be changed and for descriptive notes to be added.
- ❑ **Integration Services** Displays what Integration Services have been installed and are currently in effect. There are five services available: Operating System Shutdown, Time Synchronization, Data Exchange, Heartbeat, and Backup (volume snapshot).
- ❑ **Snapshot File Location** Displays the location of snapshots that have been taken of the selected virtual machine. Note that if there is an active snapshot of the virtual machine, the file location cannot be changed.
- ❑ **Automatic Start Action** Specifies the action the virtual machine executes when the parent partition starts. The default action is Automatically Start If It Was Running When The Service Stopped. A start delay can also be configured so that virtual machines starting up do not contend with the parent partition for resources on the host computer.

- Automatic Stop Action** Specifies the action the virtual machine executes when the parent partition or the Virtual Machine Management Service stops. The default action is to save the state of the virtual machine.
- Snapshot** Selecting this action takes a point-in-time snapshot of a virtual machine. The virtual machine can be running, saved, or stopped when the snapshot is taken. As shown in Figure 2-15, taking consecutive snapshots of a virtual machine builds a snapshot tree in the Snapshots pane with a green arrow followed by Now, indicating which snapshot is active in the virtual machine. Snapshots can also be annotated by adding descriptive notes to them. For more information about snapshots, see the section titled “Working with Snapshots” later in this chapter.



FIGURE 2-15 Snapshots of a virtual machine.

- Rename** Selecting this action allows you to rename a virtual machine. This change is reflected only in the Hyper-V Manager interface—the names of the virtual machine files or the machine name (as known to the guest operating system) are not changed.

Direct from the Source: Network Adapters in Hyper-V

Hyper-V contains two types of network adapters that can be used by guests: a legacy network adapter and a network adapter.

The *legacy network adapter* is an emulated adapter (Intel 21140 PCI) that is available to guests who either cannot take advantage of Integration Services or must have connectivity to the physical network to download and install prerequisites before they can take advantage of Integration Services (for example, Windows XP Professional x86 must download and install Service Pack 3).

A *network adapter* is a synthetic device that can be used only after Integration Services are installed in nonenlightened guests. Enlightened guests already have the necessary components installed in the operating system to begin taking advantage of this type of network adapter.

The default is to configure a network adapter when creating a new virtual machine. If a legacy network adapter needs to be added, it must be done after the virtual machine is created. This is accomplished by selecting the virtual machine in the Hyper-V Manager console and modifying the settings by using the Add Hardware process. If this is not done, there will be no virtual NIC present in the guest after it boots.

After Integration Services are installed on a guest, the legacy network adapter can be removed and replaced with a network adapter (synthetic NIC).



Note If a guest operating system is going to be installed using PXE boot to download an image, a legacy network adapter must be used and the boot order must be modified in the virtual machine settings.

—CSS Global Technical Readiness (GTR) team

Using the Virtual Machine Connection Tool

You can use the Virtual Machine Connection tool to connect to and manage an individual virtual machine running on a Hyper-V server. The Virtual Machine Connection tool uses the same Remote Desktop Protocol (RDP) technology used for remotely connecting to Windows desktops.

VMConnect links RDP to a simulated virtualized display of the VM. It's sometimes easier to think of it as a virtual keyboard video mouse (KVM) device in the parent partition—on one side, it's RDP; on the other, it's the keyboard and screen of the VMs.

Figure 2-16 shows the Virtual Machine Connection tool connected to a virtual machine named SEA-SRV1-V running on a Hyper-V server named SEA-SCV. The Virtual Machine Connection tool is currently running in windowed mode and shows that a guest operating system—a Full installation of Windows Server 2008 x64 Standard edition—is in the process of being installed in the virtual machine.




FIGURE 2-16 The Virtual Machine Connection tool.

The Virtual Machine Connection tool's menu bar provides the following options to select from:

- **File** This menu option allows you to access the Settings dialog box for the virtual machine (as shown earlier in Figure 2-13) or to exit from (close) the Virtual Machine Connection tool.
- **Action** This menu option allows you to perform any of the following actions for the virtual machine:
 - ❑ Display the Windows Logon screen (the screen that would result if you pressed Ctrl+Alt+Del in the virtual machine).
 - ❑ Turn Off, Shut Down, or Save the state of the virtual machine.
 - ❑ Pause or Reset the virtual machine.
 - ❑ Take a Snapshot of the virtual machine or Revert to a previous snapshot.

- ❑ Choose the Insert Integration Services Setup Disk option, which launches the process of installing Integration Services on the virtual machine.
- **Media** This menu option includes a DVD Drive option that allows you to insert a DVD/CD/ISO file into the virtual machine's DVD drive or eject a disk from it, or to capture the physical DVD/CD drive on the host computer. It also provides a Diskette Drive option that allows you to insert or eject a virtual floppy disk drive (.vfd file).
- **Clipboard** This menu option includes a Type Clipboard Text option that allows you to transfer text from the parent partition into the child partition, and a Capture Screen option that allows you to capture a screen shot in the child partition so that you can paste it into an imaging program, such as Microsoft Paint, that is running in the parent partition.
- **View** This menu option includes a Full Screen Mode option that allows you to expand the Virtual Machine Connection interface so that it fills the screen, and a Toolbar option that allows you to return the Virtual Machine Connection tool to windowed mode.



Tip The Virtual Machine Connection tool also displays a toolbar when running in windowed mode. This toolbar is shown in Figure 2-16 immediately below the menu bar. You can use it to send Ctrl+Alt+Del to the virtual machine, start or stop the virtual machine, pause the machine, take a snapshot of the virtual machine's state, and perform other actions.

Installing the Virtual Machine Connection Tool

The Virtual Machine Connection tool is installed by default when you add the Hyper-V role to a Full installation of Windows Server 2008 R2. When you use the Virtual Machine Connection tool to connect to and manage virtual machines running on your local Hyper-V server, the title bar of each Virtual Machine Connection tool displays *localhost* as the name of the Hyper-V server on which each virtual machine runs. If you use the Virtual Machine Connection tool to connect to and manage virtual machines running on a different Hyper-V server, the title bar of each Virtual Machine Connection tool displays the name of the Hyper-V server on which the virtual machines are running.

You can also install the Virtual Machine Connection tool onto a Windows Server 2008 R2 computer that is not running the Hyper-V role, and then use the Virtual Machine Connection tool on this computer to manage virtual machines running on your Hyper-V servers. To do this, install Hyper-V Tools from the Remote Administration Tools section of the Remote Server Administration Tools (RSAT) feature using the Add Features Wizard.

You can also install the Virtual Machine Connection tool onto a computer running Windows 7 Professional, Enterprise, or Ultimate edition and use this computer to manage

virtual machines running on your Hyper-V servers. To do this, download and install RSAT for Windows 7 from <http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=7d2f6ad7-656b-4313-a005-4e344e43997d> and follow the instructions on that page for enabling the tools you need.



Note If your Hyper-V server is running Server Core, you will not be able to install either the Hyper-V Manager snap-in or the Virtual Machine Connection tool directly on your server. Instead, you must use these tools remotely from a different computer running either Windows 7 with RSAT or a Full installation of Windows Server 2008 R2.

Connecting to a Virtual Machine

To use the Hyper-V Manager console to connect to and manage a virtual machine by using the Virtual Machine Connection tool, do one of the following:

- Right-click on a virtual machine listed in the top center Virtual Machines pane of the Hyper-V Manager console and select Connect.
- Click on a virtual machine listed in the top center Virtual Machines pane of the Hyper-V Manager console to select the virtual machine, and then click Connect in the Actions pane or select Connect from the Action menu.
- Double-click on a virtual machine listed either in the top center Virtual Machines pane or on a thumbnail of a virtual machine displayed in the lower center pane.

You can also launch the Virtual Machine Connection tool from the command prompt. For example, to connect to a virtual machine named SEA-SRV1-V running on the local Hyper-V server, type **"C:\Program Files\Hyper-V\vmconnect.exe" localhost SEA-SRV1-V** at the command prompt. If you have two or more virtual machines with the same display name, you can connect to the desired virtual machine by specifying the GUID of the virtual machine instead of its display name as shown in the preceding command example. To specify a GUID, you must use the **-G** switch like this: **"C:\Program Files\Hyper-V\vmconnect.exe" localhost -G <GUID>**.



Note Although the Hyper-V Manager snap-in always runs in an elevated state, the Virtual Machine Connection tool does not. To avoid authentication issues when running the Virtual Machine Connection tool on a computer running Windows 7, you must either run the tool with elevated privileges or add your user account to the Hyper-V Administrators role by using the Authorization Management tool.

Using Windows Keyboard Accelerators

Standard Windows keyboard accelerators must be replaced by their equivalents when managing a virtual machine using the Virtual Machine Connection tool. These differences are summarized in Table 2-5.

TABLE 2-5 Changes to Standard Keyboard Accelerators When Managing a Virtual Machine Using the Virtual Machine Connection Tool

Standard Windows Keyboard Accelerator	Virtual Machine Connection Equivalent	Description
CTRL+ALT+DEL	CTRL+ALT+END	Displays the Windows Security dialog box
ALT+TAB	ALT+PAGE UP	Switches between programs from left to right
ALT+SHIFT+TAB	ALT+PAGE DOWN	Switches between programs from right to left
ALT+ESC	ALT+INSERT	Cycles through the programs in the order they were launched
CTRL+ESC	ALT+HOME	Displays the Windows Start menu
N/A	CTRL+ALT+PAUSE	Changes the Virtual Machine Connection window to or from full-screen mode
N/A	CTRL+ALT+LEFT ARROW	Releases the mouse and keyboard focus from the Virtual Machine Connection window when Integration Services is not installed



Note By default, standard Windows keyboard accelerators are not sent to the virtual machine unless you are working in full-screen mode. You can modify this, however, so that standard Windows keyboard accelerators are always sent to the virtual machine when the Virtual Machine Connection tool has the focus. You do this by opening the Hyper-V Manager console, selecting Hyper-V Settings, choosing Keyboard, and selecting the Use On The Virtual Machine option. Note, however, that Ctrl+Alt+Del will always go to the host (physical) computer, so you must use Ctrl+Alt+End regardless of the setting you have selected here.

Creating a Virtual Machine

Once your Hyper-V server has been deployed and configured and you are familiar with how to use the Hyper-V Manager console and the Virtual Machine Connection tool, you can create new virtual machines (child partitions) on your server and install a guest operating system onto each virtual machine.

Based on how you configure disk storage, there are two types of virtual machines you can create:

- Virtual machines that have their guest operating system installed on a virtual hard disk, which is implemented as a file on the hard drive of the host computer
- Virtual machines that have their guest operating system installed on a separate hard drive on the host computer, a configuration that is known as a *passthrough disk*

The following sections outline the general procedures for creating virtual machines using each of these storage configurations. For additional information concerning how Hyper-V storage works and how it can be configured, see the sidebar titled “Direct from the Source: Understanding Hyper-V Storage” later in this section.



Tip Before you create a new virtual machine, make sure you have sufficient disk storage space on your host computer for the operating system and applications you will be installing in the virtual machine. Also make sure that you have configured a virtual network so that guests will be able to access the physical network if needed.

Direct from the Source: Understanding Hyper-V Storage

Hyper-V supports several different storage options, including Direct Attached Storage (DAS)—for example, SATA or SAS—and SAN Storage—for example, FC or iSCSI. After the Hyper-V server has been connected to storage, it can be made available to guests in many different ways.

After storage has been exposed to the Hyper-V server, there are two choices available for hosting the guest operating system.

- Creating a virtual hard disk (VHD) on one of the volumes on the Hyper-V server. The virtual hard disk is simply a file that is stored on one of the storage volumes on the Hyper-V server. There are two types of virtual hard disks: dynamic and fixed. The maximum size of a VHD file is 2040 gigabytes (just short of 2 terabytes).
- Using a passthrough disk, which allows the virtual machine to access the disk directly. The raw disk (no size limit) can be a disk local to the Hyper-V server or a logical disk (logical unit number [LUN]) on a SAN. Before configuring a guest with a passthrough disk, the disk must be placed in an offline state so that there is no contention between the virtual machine and the Hyper-V server. This is accomplished in the Windows Disk Management snap-in or by using the Diskpart.exe command-line interface (CLI).

Connecting Storage to the Guest

There are three methods available to connect storage to a virtual machine:

- **IDE** The Hyper-V IDE controller allows for disks up to 2048 gigabytes. Additionally, the new filter driver used for IDE in Hyper-V essentially bypasses the emulation path for IDE, providing much higher performance that is almost on par with SCSI. The IDE controller can support either virtual hard disks or passthrough disks. There can be up to four IDE disks configured on a guest (2 controllers with 2 disks each). One important note is that Hyper-V virtual machines can boot only from IDE. Booting from virtual SCSI is not supported. This is mainly because a SCSI controller is a synthetic device and must be added only after Integrated Services have been installed on the guest.
- **SCSI** The Hyper-V SCSI controller is a synthetic device and therefore cannot be added to a guest configuration until after Integrated Services have been installed. There can be up to four SCSI controllers configured per guest. Each controller can support 64 disks each, for a total of 256 disks per virtual machine. SCSI disks backed with VHD are limited to 2040 GB. A guest cannot be configured to boot from a SCSI controller.
- **iSCSI** Guests connected to a physical network can take advantage of iSCSI storage. Guests can connect directly to iSCSI storage over an iSCSI network, completely bypassing the Hyper-V server itself. All that is required is the proper configuration of an iSCSI client in the guest and an iSCSI target running somewhere on the network that is accessible by the guest. There is no limit to the number of iSCSI disks that can be supported on the guest. A guest cannot boot from an iSCSI disk.

—CSS Global Technical Readiness (GTR) team



Tip You can bypass the 2048-GB size limitation for IDE and SCSI virtual disks by using passthrough disks.

Creating a Virtual Machine Using a Virtual Hard Disk

The general procedures for creating a new virtual machine that uses a virtual hard disk (.vhd file) and installing a guest operating system are as follows:

1. Launch the New Virtual Machine Wizard from the Hyper-V Manager console. Follow the steps of the wizard to assign memory, configure a network, and perform other required steps. If the operating system you plan on installing in your virtual machine is an unenlightened guest, do not configure a network at this point because you need

to either install Integration Services first or configure a legacy network adapter for connectivity.



Note If your virtual machine will be running a 64-bit version of either Windows XP or Windows Server 2003, there is no driver for the legacy network adapter included in Hyper-V. This means the synthetic network adapter must be used, which requires Integration Services to be installed first.

2. When you get to the Connect Virtual Hard Disk page of the wizard, choose the Create A Virtual Hard Disk option and verify the location for the disk.
3. When you finish the wizard, the virtual machine will start and you can install the guest operating system from CD or DVD media. Connect to the virtual machine using the Virtual Machine Connection tool so that you can respond to any prompts displayed during the install process.
4. After you install any guest operating system, the next thing you should do is install the Integration Services components on your virtual machine.
5. At this point, shut down your virtual machine to configure additional storage controllers, additional hard disks, and additional processors as needed. Then boot your virtual machine and install roles and features, install applications, join the domain, and perform other initial configuration tasks as needed.

Creating a Virtual Machine Using a Passthrough Disk

The general procedures for creating a new virtual machine that uses a passthrough disk and installing a guest operating system are as follows:

1. Ensure that you have at least one dedicated disk volume of sufficient size on your host computer to use as the system/boot volume for your new virtual machine.
2. Ensure that you have a separate location available for storing the virtual machine configuration (.xml) files for your new virtual machine. This is required because a virtual machine configured with a passthrough disk uses the entire passthrough disk volume for its operating system. The .xml files for the virtual machine must be stored on a different volume. The location you choose for storing the .xml files can be a different hard disk volume on your Hyper-V server, or it can even be a shared folder on a network file server. For more information, see the sidebar titled “Direct from The Source: Relocating Virtual Machine Configuration Files” later in this section.
3. Launch the New Virtual Machine Wizard from the Hyper-V Manager console. Follow the steps of the wizard to assign memory, configure a network, and perform other required steps. If the operating system you plan on installing in your virtual machine is an unenlightened guest, do not configure a network at this point because you need

to either install Integration Services first or configure a legacy network adapter for connectivity.

4. When you get to the Connect Virtual Hard Disk page of the wizard, choose the Attach A Virtual Hard Disk Later option and continue through the wizard.
5. This time, when you finish the wizard, the virtual machine will not start because no storage has been configured for it to use.
6. Select your virtual machine in the Hyper-V Manager console, and click Settings in the Actions pane.
7. Select IDE Controller 0, click Add, select the Physical Hard Disk option, and then choose the correct physical disk volume from the drop-down list.
8. Select IDE Controller 1, and select the physical CD/DVD drive, or mount an .iso image file that has your operating system files.
9. Apply the settings you have configured, start your virtual machine, and install the guest operating system. Connect to the virtual machine by using the Virtual Machine Connection tool so that you can respond to any prompts displayed during the install process. Continue as described in the previous section.

Direct from the Source: Relocating Virtual Machine Configuration Files

The typical configuration of a virtual machine (guest) includes storing the virtual machine configuration files in the default location on the system drive under `\ProgramData\Microsoft\Windows\Hyper-V` in a folder corresponding to the name given to the virtual machine in the New Virtual Machine Wizard. This location can be changed by manipulating the settings for the Hyper-V server using Hyper-V Settings in the Actions pane in the Hyper-V Manager console.

There are scenarios where storing virtual machine configuration files on a remote server that is not running Hyper-V is a very real possibility. Storing Hyper-V configuration files in an alternate location is required when using passthrough disks. This is because the entire disk is used for the operating system files and there is no room for the configuration files. This is true whether this configuration is used in a standalone Hyper-V server or when making a virtual machine highly available in a failover cluster.

As an example, follow these steps to store configuration files in a remote location (for example, File Server):

1. First, configure a folder on a remote machine that will be shared and used to store the virtual machine configuration files. In this example, the shared folder VMCONFIG is created on a remote server that can be accessed from the Hyper-V server.

2. When configuring the virtual machine, use the network share for the location of the virtual machine configuration files (UNC path).
3. If the default permissions on the share are not modified, an error will be encountered when you are trying to complete the New Virtual Machine Wizard.
4. Permissions need to be modified on the share such that the user running the New Virtual Machine Wizard and the Hyper-V server computer account both have Write permissions to the share. After that has been accomplished, the wizard will complete and the configuration files will be placed on the file share.

—CSS Global Technical Readiness (GTR) team

Working with Virtual Machines

You can use the Hyper-V Manager console to perform a number of management tasks involving virtual machines. This section briefly examines three of these tasks:

- Exporting and importing virtual machines
- Working with snapshots
- Working with Live Migration (new in Hyper-V R2)

Exporting and Importing Virtual Machines

You can use the Hyper-V Manager console to export a virtual machine from one Hyper-V server so that you can import it onto a different Hyper-V server. This import/export functionality allows you to migrate a virtual machine from one host computer to another using a process called Quick Migration.

The procedure for exporting a virtual machine from one Hyper-V server and importing it into another involves two steps:

1. Export the virtual machine from the first Hyper-V server as a collection of exported files and folders.
2. Import the exported files and folders onto your destination Hyper-V server.

The following is an outline of the steps involved for exporting a virtual machine:

1. Begin by shutting down the virtual machine you want to move. To shut down a virtual machine, select the virtual machine in the Hyper-V Manager console and click Shut Down in the Actions pane.

2. Decide on the location to which you will export your virtual machine. Your export location could be any one of the following:
 - ❑ A temporary folder on an external hard drive to transport the exported virtual machine files from the first Hyper-V server to the destination server.
 - ❑ A shared folder on a network file server used to temporarily store the virtual machine files until they are moved to the destination server.
 - ❑ A shared folder on your destination server that represents the final location to which your virtual machine is being migrated.
3. Select the virtual machine you want to export, and click Export in the Actions pane. When the Export Virtual Machine dialog box is displayed, type or browse to the path of the export location. If the destination folder is a shared folder on the network, specify the path to the folder as a UNC path.
4. Click the Export button to initiate the export process.

When the export process is finished, the following files and folders will be present in the export location:

- **Config.xml** An XML file that contains information about the original locations of all virtual hard disks configured for the exported virtual machine.
- **Virtual Machines** A folder that contains an export file whose name is of the form <GUID>.exp. This export file contains configuration information for the exported virtual machine and is converted during the import process into an XML configuration file.
- **Virtual Hard Disks** A folder that contains the virtual hard disks (.vhd files) for the exported virtual machine.
- **Snapshots** A folder that contains information about any snapshots taken of the virtual machine, including the snapshot differencing disks (.avhd files) and state information files (.vsv and .bin files) for those snapshots.

After you have exported your virtual machine and copied the export files and folders to their final locations on your destination server, you are ready to import these files and folders so that you can re-create your virtual machine on the destination server.

But there are two things you need to know about first concerning this import process. First, you can import only virtual machines that were exported from another Hyper-V server. You cannot import virtual machines that were imported from either Virtual Server 2005 or Virtual PC. This is because even though all three server virtualization products (Hyper-V, Virtual Server, and Virtual PC) use the same virtual hard drive (.vhd) file format, they store virtual machine configuration information differently and also have additional incompatibilities in terms of the features they each support.

Second, you can perform the import process only once per exported virtual machine. This is because, during the import process, the export (.exp) files are converted into XML configuration (xml) files. What this also means is that if the import process fails or is performed incorrectly—for example, if you import the exported files and folder into the wrong location—the only way to recover is to delete the virtual machine, relocate the .vhd files to the correct location, and then re-create the virtual machine by recalling the settings that were used.

The following is an outline of the steps involved for importing your exported virtual machine files and folders:

1. Make sure your exported files and folders are in their correct locations on your destination server.
2. Connect to your destination server using the Hyper-V Manager console, and click Import Virtual Machine in the Actions pane.
3. In the Import Virtual Machine dialog box, type or browse to the location of the exported files and folders.
4. Click the Import button to initiate the import process.

After the virtual machine has been imported, try starting it and make sure it is functioning properly.

Working with Snapshots

A *snapshot* is a point-in-time picture of the state and settings of a virtual machine. Hyper-V allows you to capture snapshots of virtual machines and revert back to those snapshots. For example, you could install a guest operating system on a virtual machine, take a snapshot, make some configuration changes to the guest, and then revert back to your snapshot to undo your configuration changes.

Snapshots can be taken when a virtual machine is running, saved, or shut down. Snapshots cannot be taken, however, when a virtual machine is paused. You can take multiple snapshots of a virtual machine to create a snapshot tree, which is a sequence of snapshots taken at different times. You can manage this tree of snapshots by deleting individual snapshots or an entire subtree of snapshots. And you can revert to any particular snapshots in a tree by applying that snapshot to your virtual machine.

Snapshots can be particularly useful during the test and development stages of a product development cycle. For example, you can install an application you are developing on a virtual machine, take a snapshot, and then try working with the application. If the application crashes, you can revert back to your snapshot and try to reproduce the steps that led to the crash, which can help you troubleshoot the cause of the crash.



Note Snapshots should generally not be used in production environments because they are not intended as replacements for proper backup and recovery processes. For example, although running domain controllers in virtual machines is supported on Hyper-V, taking snapshots of domain controllers and then reverting to them later can cause replication problems and should therefore not be done in a production environment.

You can use the Hyper-V Manager console to take a snapshot of a virtual machine. To do this, select the virtual machine in the Virtual Machines pane and click Snapshot in the Actions pane for the selected virtual machine. As Figure 2-17 shows, when a new snapshot is created, an icon for the snapshot is displayed in the Snapshots pane in the center of the console. The new snapshot is given a descriptive name that includes the name of the virtual machine from which the snapshot was made and the date and time when the snapshot occurred.



FIGURE 2-17 Snapshot of a virtual machine.

As shown in Figure 2-17, when you take a snapshot of a virtual machine, a green arrow labeled Now is also displayed in the Snapshots pane. This Now arrow represents the current running configuration of the virtual machine.

You can also take snapshots of a virtual machine using the Virtual Machine Connection tool. You can do this in two ways:

- By selecting Snapshot from the Action menu bar item.
- By clicking the Snapshot icon on the toolbar.

When you take a snapshot using the Virtual Machine Connection tool, a Snapshot Name dialog box is displayed, prompting you to provide a descriptive name for the new snapshot.

By default, all snapshot files are stored in the following folder on your Hyper-V server:

`%SystemRoot%\ProgramData\Microsoft\Windows\HyperV\Snapshots`

You can change this location on a per-VM basis by configuring the settings for each VM.

Taking a snapshot of a virtual machine creates the following types of snapshot files:

- Virtual machine configuration (.xml) file
- Virtual machine saved state (.vsv) files
- Virtual machine memory contents (.bin) files
- Snapshot differencing disk (.avhd) files

As shown in Figure 2-17 when you select a snapshot in the Snapshots pane, the Actions pane displays various actions you can perform with that snapshot. These include the following:

- **Apply** Selecting this action allows you to copy the complete virtual machine state from the selected snapshot to the active virtual machine. This allows you to revert your virtual machine to the state contained in the selected snapshot. When you select this action, any unsaved data in your currently active virtual machine will be lost. Because of this, you are prompted to choose whether you want to create a new snapshot of your current virtual machine state before the state contained in the selected snapshot is applied.
- **Rename** Selecting this action allows you to modify the descriptive name of the selected snapshot.
- **Delete Snapshot** Selecting this action allows you to remove only the files associated with the selected snapshot. (Files for other snapshots will not be affected.) After you delete a snapshot, you will be unable to revert to the state contained in that snapshot. The current state of the active virtual machine is not affected by this action.
- **Delete Snapshot Tree** Selecting this action allows you to delete the selected snapshot and any snapshots hierarchically beneath it. The current state of the active virtual machine is not affected by this action.



Note Snapshots are read-only. The only settings you can configure for a snapshot are its name and any attached descriptive notes.

If you select the virtual machine in the Virtual Machines pane, the tasks displayed in the Actions pane change to the following:

- **Snapshot** Selecting this action allows you to take another snapshot of your virtual machine.
- **Revert** Selecting this action allows you to apply the previous snapshot (the snapshot directly above the green Now arrow in the Snapshots pane).



Tip When you delete an entire snapshot tree, the result will be the last snapshot applied to the running virtual machine. If your intention, instead, is to have the result be the pristine installation of your virtual machine, your first snapshot should be taken after your virtual machine is configured and before you make any alterations for testing your configuration. That way, you can apply your first snapshot (the root snapshot) before deleting the snapshot tree, and the result is that your virtual machine's configuration will return to where you started before you made your alterations.

Direct from the Source: Best Practices for Configuring Virtual Machines

Virtual machine performance is affected not only by how the physical server is configured but also by the selections made when configuring the virtual machine itself. The following sections discuss best practices that should be considered when configuring virtual machines in Hyper-V.

Change Default Locations for Virtual Hard Disk and Machine Configuration Files

Change the default locations for storing the virtual hard disks and the virtual machine configuration files. By default, they are stored on the drive where the operating system is installed. For better performance, move the location to another disk on a SAN, if possible. If no SAN storage is configured, use another internal, fault-tolerant drive or drives that can be dedicated to storing virtual machine data and are not supporting the operating system.

Install Integration Services

The first, and probably most important, best practice for virtual machines is to install Integration Services, which comes with Hyper-V, as soon as possible if the operating system running in the virtual machine is supported. Then update Integration Services as needed.

Uninstall VM Additions and Compact VHDs

When migrating virtual machines from Virtual PC or Virtual Server 2005 R2, uninstall the VM Additions and compact the virtual hard disk before moving the disk to the Hyper-V server.

Set Display for Best Performance

For the best display in a virtual machine, ensure the display interface is set for Best Performance. This ensures the hardware acceleration is set to Full.

Configure Fixed-Size VHDs

Choose to configure fixed-size virtual hard disks rather than dynamically expanding disks. Performance is faster, the file system is less likely to fragment, and managing space on the physical disk is easier. Always defragment a physical disk before creating a virtual hard disk.


Use SCSI Virtual Adapters for Data Drives

Hyper-V requires the virtual machine to boot from a virtual IDE controller; however, SCSI virtual adapters can be used after that for mounting additional virtual hard disks. Although performance differences between a virtual IDE controller and a virtual SCSI controller in Hyper-V is negligible (with Integration Services installed), more and larger capacity virtual hard disks can be attached to a virtual SCSI controller (4 controllers with 64 virtual disks each, for a total of 256). So, if you need more than four virtual hard disks attached to a virtual machine, use a virtual SCSI controller.

Allocate CPU Resources Based on Anticipated Usage

It is also important to determine virtual machine performance to ensure CPU resource allocation on the physical server is adequate to support the workload inside the virtual machine. The default in Hyper-V server is to treat all virtual machines equally. In reality, this might not be a practical or wise business decision. When allocating physical machine CPU resources to a virtual machine, it is important not to over-subscribe—that is, trying to allocate more physical machine resources than are really available. The next version of System Center Virtual Machine Manager (SCVMM 2008) will play a key role in monitoring virtual machine performance.

To help with this process, the following figure shows the Processor configuration setting for a virtual machine:



The interpretation of the Processor configuration settings shown in the preceding figure are as follows:

- **Virtual Machine Reserve** Percent of the logical CPU that is set aside for the running virtual machine. As each VM is started, the available capacity on the Hyper-V server itself is reduced.
- **Virtual Machine Limit** Percentage of logical CPU that a running virtual machine is not allowed to exceed.
- **Relative Weight** Determines how CPU is distributed when there is contention among all running virtual machines. The higher the number, the more processing power allocated to the VM. Relative weight can range from 1 to 10,000.
- **Migrate To A Physical Computer With A Different Processor Version** This configuration option is new in Hyper-V R2 and enables or disables processor compatibility mode. For more information on this feature, see the sidebar titled “How It Works: Processor Compatibility Mode” later in this chapter.
- **Run An Older Operating System, Such As Windows NT** Reduces vulnerability of some operating systems to high central processing unit identification

(CPUID) values. Unexpected, high CPUID values can cause a crash. This option was called Limit Processor Functionality in the original release version of Hyper-V.

Consider Using Passthrough Disks

When creating a virtual machine, it is a best practice to use virtual hard disks; however, circumstances might dictate using passthrough disks. Performance using passthrough disks is slightly better than performance achieved using a virtual hard disk (VHD), you can conserve drive letters, and you can configure disks larger than two terabytes (if the external storage supports that). However, when using passthrough disks, the virtual machine configuration files need to be relocated to either another hard disk or a file share. Additionally, you lose snapshot functionality when using passthrough disks, and they are not portable like a file (VHD).

Ensure File Share High Availability

If a file share is being used to store virtual machine configuration data, it is a best practice to ensure the file share is highly available (for example, a file share being hosted in a failover cluster). You also need to modify the security on the file share to allow the Hyper-V server (all nodes of it if it's in a failover cluster) write access to the share.

Configure Domain Controllers to Optimize Performance

Domain controllers are supported in Hyper-V. The following best practices are recommended for these configurations:

- Never save state in a domain controller because this might cause synchronization issues in the domain.
- Never pause a domain controller virtual machine for long periods of time because this might adversely affect replication.
- Always shut down a domain controller.
- Do not take snapshots of a domain controller.
- Make a determination regarding time synchronization. The decision is either to use the Hyper-V Integration Service For Time Synchronization or not. If the decision is to treat the virtualized domain controllers like hardware-based domain controllers, disable the Time Synchronization capability in the settings for each virtual machine and point the PDC Emulator to an external time source and allow all the other domain controllers to synchronize with the PDC Emulator. If the decision is to synchronize with the parent partition, enable only the Time Synchronization capability for the domain controller holding the PDC Emulator FSMO role.

—CSS Global Technical Readiness (GTR) team

Working with Live Migration

Live Migration is a new feature of Hyper-V in Windows Server 2008 R2 that makes running virtual machines highly available by allowing them to be transparently moved between the nodes of a failover cluster without perceived downtime or dropped network connections. Live Migration relies on the new Cluster Shared Volumes feature of the Failover Clustering feature of Windows Server 2008 R2, which is described in the next section. Live Migration can also take advantage of the new Processor Compatibility Mode feature of Hyper-V R2 to transparently migrate running virtual machines between host machines that have the same processor architecture (AMD or Intel) but different processor features. For more information on processor compatibility mode, see the sidebar titled “How It Works: Processor Compatibility Mode” later in this chapter.

Understanding Cluster Shared Volumes

Cluster Shared Volumes are a new type of storage volume supported by the Failover Clustering feature of Windows Server 2008 R2. The feature allows multiple cluster nodes to concurrently read from and write to a single shared volume. Cluster Shared Volumes is intended only for use by Hyper-V and is not supported for other uses unless specified by Microsoft.

When Cluster Shared Volumes is implemented on a failover cluster of Hyper-V servers, the clustered virtual machines (the virtual machines residing on the different nodes of the cluster) can all access their virtual hard disks at the same time. These virtual hard disks all reside on the cluster shared volume, which is a single logical unit number (LUN) in the cluster storage array. This arrangement enables the clustered virtual machines to fail over independently of each other. At any given time, the state of each running virtual machine is managed by only one of the cluster nodes.

Without the Cluster Shared Volumes feature, making clustered virtual machines highly available requires additional complexity. This is because each LUN in the storage cluster array can be accessed only by one cluster node at a time. Each clustered virtual machine, therefore, requires its own separate LUN, which makes implementing and managing clustered virtual machines and LUNs more complex. The new Cluster Shared Volumes feature helps remove this complexity, making highly available virtual machines easier to implement.

Implementing Clustered Hyper-V

The simplest form of highly available clustered Hyper-V is a two-node failover cluster. The cluster typically consists of the following:

- Two servers running Windows Server 2008 R2 Enterprise or Datacenter edition and having the Hyper-V role and Failover Clustering feature installed. The same version of Windows Server 2008 R2 must be used for each node—you cannot have one node

running Enterprise edition and the other running Datacenter. The same architecture must also be used—you cannot have one be an Itanium server and the other x64. The same installation option must also be used for each node—you cannot have one node running a Full installation of Windows Server 2008 R2 and the other running a Server Core installation. The hardware for these servers should also generally be as identical as possible, and the servers should have the same service packs and software updates installed. Finally, both servers should be member servers of the same Active Directory Domain Services (AD DS) domain.

- A storage array that consists of two LUNs configured at the hardware level. The storage array can use Serial Attached SCSI (SAS), iSCSI, or Fibre Channel. The storage array should include
 - One LUN configured as the witness disk for the cluster. The witness disk holds a copy of the cluster configuration disk. The quorum configuration for the cluster will be Node And Disk Majority, which is the default for a two-node cluster and which stores the cluster configuration on the nodes and the witness disk.
 - One LUN configured as the shared storage on which the virtual machines and their virtual hard disks reside. Cluster shared volumes must be basic disks, not dynamic. They can be either master boot record (MBR) or GUID partition table (GPT) disks and should be formatted using NTFS.

A network infrastructure that connects the cluster nodes with each other and with the storage array. This can be implemented in different ways, but you should do it in a way that avoids single points of failure. In particular, make sure that the network connection used by the cluster shared volume is fault tolerant.


All storage and networking hardware in the cluster, including the servers themselves, must be “Certified for Windows Server 2008 R2.” You should use the Validate A Configuration Wizard to verify that your hardware and configuration will support failover clustering before you implement your clustering solution.

Setting Up for Live Migration

The steps for setting up for Live Migration using the Hyper-V and Failover Clustering roles of Windows Server 2008 R2 are as follows:


1. Set up the server, storage, and network architecture described in the previous section to support failover clustering and management of your virtual machines.
2. Install Windows Server 2008 R2 on each server, and add the Hyper-V role.
3. Install the Failover Clustering feature on each server, and configure the servers as nodes in a failover cluster. Validate your cluster configuration by running the Validate A

Configuration Wizard. You can do this by opening the Failover Cluster Manager console and clicking Validate A Configuration in the Management central pane:



4. Using Failover Cluster Manager, right-click on your failover cluster and select Enable Cluster Shared Volumes. After you have done this, all nodes in the cluster will be able to use shared volumes. You can enable Cluster Shared Volumes only once per failover cluster. A node named Clustered Shared Volumes will be displayed in the console tree after you complete this step.
5. Select the cluster you want to manage, and click Add Storage in the Actions pane. In the Add Storage dialog box, select a disk from the list of available disks and click OK to add the disk to the cluster shared volume. The cluster shared volume storage location will be displayed as SystemRoot\ClusterStorage for all nodes of the failover cluster, and each disk you add will be displayed as a folder underneath this.
6. Now create your clustered virtual machines, choosing a cluster shared volume as the location to store both the virtual machines and their virtual hard disks. For the cluster shared volume to be available to the virtual machines, you must create the virtual machines on the cluster node that owns the cluster shared volume.

7. Now use Hyper-V Manager to open the settings for each virtual machine, select Automatic Start Action in the left pane, and select Nothing to specify what you want the virtual machine to do when the physical computer starts. Doing this enables the management of the virtual machine state to be controlled through the Cluster service.
8. Now use Failover Cluster Manager to make your clustered virtual machines highly available. To do this, select Services And Applications in the console tree and click Configure A Service Or Application in the Actions pane to launch the High Availability Wizard. Click Virtual Machine on the Select Service Or Application page; then on the Select Virtual Machine page, select the check box beside the virtual machine you want to make highly available and finish the wizard. The virtual machine will be displayed under the Services And Applications node, and you should start the virtual machine by right-clicking on it and selecting Bring This Service Or Application Online. Then install the guest operating system and applications on the virtual machine. Repeat the process with any other virtual machines that you want to make highly available.
9. Now in Failover Cluster Manager, under Services And Applications, select the clustered virtual machine for which you want to configure the cluster network for Live Migration. In the center pane, right-click on the virtual machine resource and select Properties, click the Network For Live Migration tab, and select one or more cluster networks to use for Live Migration.



You can use the Up and Down buttons to re-order the networks, with the network at the top being the preferred one for use during a live migration. Generally, you will want to ensure that a private cluster network is the preferred network for a live migration,


while a network used only for cluster traffic and by cluster shared volumes should be last on your list. Configuring this setting for one clustered virtual machine automatically makes it apply globally to all virtual machines on the cluster.

At this point, you have set up a failover cluster of Hyper-V servers with highly available clustered virtual machines running on them that support Live Migration. The next section shows how to perform a live migration.

Performing a Live Migration


To perform a live migration and move a virtual machine from one cluster node to another, do the following:

1. Open Failover Cluster Manager, select the cluster, and expand Services And Applications.
2. Right-click the virtual machine resource in the center pane, and click Live Migrate Virtual Machine To Another Node; then select the node you want to move the virtual machine to.




In the preceding screen shot, the two virtual machines are initially owned by R2-NODE1, and the virtual machine CONTOSO-FS5 is being migrated to R2-NODE2.

3. The progress of the live migration will be displayed in the center pane.



The time it takes to perform the live migration depends on the amount of RAM configured for the virtual machine, the load on the source and destination cluster nodes, and the network connection speed and bandwidth available for moving the virtual machine.

4. After the migration is complete, the ownership of virtual machine CONTOSO-FS5 has been transferred from R2-NODE1 to R2-NODE2.



In the R2 release of Hyper-V and Failover Clustering, you can also use Windows PowerShell to perform a live migration. The command for doing this is as follows:

```
Get-Cluster "<Cluster Name>" | Move-ClusterVirtualMachineRole -Name "<VM group name>" -Node "<Destination node name>"
```

In the preceding command, <Cluster Name> is the name of the cluster on which the virtual machine resides, <VM group name> is the name of the virtual machine resource group, and <Destination node name> is the name of the destination cluster node to which you want to move the virtual machine using Live Migration.



More Info To learn more about how to configure Failover Clustering for highly available virtual machines, see the topic "Hyper-V: Using Hyper-V and Failover Clustering" in the TechNet Library at <http://technet.microsoft.com/en-us/library/cc732181.aspx>. To learn more about how to perform a live migration, see the topic "Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2" at <http://technet.microsoft.com/en-us/library/dd446679.aspx>.

How It Works: Processor Compatibility Mode

Processor compatibility mode is a new feature of Hyper-V in Windows Server 2008 R2 that lets you migrate a virtual machine between host machines having the same processor architecture (either AMD or Intel). When you start a virtual machine on a Hyper-V host, the hypervisor exposes the supported processor features that are available on the host's hardware. The exposed processor features are known as guest visible processor features, and these processor features are available to the virtual machine until it is restarted.


If you now enable processor compatibility mode on the virtual machine, Hyper-V normalizes the processor feature set. What this means is that only processor features that are available on all processors having the same architecture are exposed as guest visible processor features. In other words, the guest operating system sees either a "standard" Intel processor or a "standard" AMD processor, regardless of the actual physical processor on the host machine. By hiding the real processor features and exposing only a standard processor that has limited features, the ability to migrate the virtual machine between processors having the same architecture is enabled. The actual way this hiding of processor features occurred was by having the hypervisor intercept the virtual machine's CPUID instruction in order to clear the returned bits that correspond to the processor features being hidden.

You do not have to make any changes in your system BIOS to enable processor compatibility mode. In addition, your processor does not need to support advanced features such as Extended Migration or Flex Migration for processor compatibility mode to function. For AMD processors, the processor features you lose when you enable processor compatibility mode include SSSE3, SSE4.1, SSE4.A, SSE5, POPCNT, LZCNT, AMD 3DNow!, Extended AMD 3DNow!, and Misaligned SSE. For Intel processors, the processor features you lose include SSSE3, SSE4.1, SSE4.2, POPCNT, Misaligned SSE, XSAVE, and AVX.

Note that if a virtual machine has processor compatibility mode enabled and is running a third-party application that can make use of one or more of these hidden processor features, it is up to the application to determine how to behave in the absence of these features. For this reason, processor compatibility mode is disabled by default on Hyper-V virtual machines—just in case having it enabled might cause problems with a third-party application running in the virtual machine. Be sure to test all applications installed on your virtual machine before enabling processor compatibility mode on the machine. Multimedia applications in particular might fail to perform as well on virtual machines that have processor compatibility mode enabled on them. This is because multimedia applications often use specialized processor features to optimize how they perform.

Note also that enabling processor compatibility mode is not required if you are moving a virtual machine to a new host machine that has a superset of the processor features of the old host machine. You have to enable this mode only if you are moving the virtual machine to a new host machine that has a subset of the processor features of the old host machine.

Processor mode is enabled on a per-virtual-machine basis. To enable processor compatibility mode for a virtual machine, begin by shutting down the virtual machine using Hyper-V Manager. Open the properties sheet for the virtual machine, and click Processor in the left pane. In the right pane, select the check box labeled *Migrate To A Physical Computer With A Different Processor Version* as shown in the following screen shot:



Enabling processor compatibility mode on a virtual machine allows you to perform either a live migration or a quick migration of the virtual machine between Hyper-V hosts having the same processor architecture. Processor compatibility mode also makes saved states and snapshots compatible as well. For more information on using processor compatibility mode, see the topic "Configure Memory and Processors" in the TechNet Library at <http://technet.microsoft.com/en-us/library/cc742470.aspx>.

Tools for Managing Hyper-V and Virtual Machines

We've already described in detail two tools you can use for managing Hyper-V servers and virtual machines: the Hyper-V Manager snap-in and the Virtual Machine Connection tool, which can be installed on any Windows Server 2008 R2 computer and can be used on a Windows 7 computer by installing RSAT for Windows 7 on the computer. There are several other ways, however, that you can remotely manage Hyper-V servers and virtual machines running on them, including the following:

- Remote Desktop Connection
- RemoteApp
- Windows Management Instrumentation (WMI)
- Windows PowerShell
- System Center Virtual Machine Manager 2008 R2

System Center Virtual Machine Manager 2008 R2 is covered in detail in Chapter 5, so the sections that follow deal with the remaining management tools.

Managing Hyper-V Using Remote Desktop Connection

Instead of connecting to a remote Hyper-V server using the Hyper-V Manager snap-in, you can use Remote Desktop Connection (Mstsc.exe) to connect to the desktop of the remote server and then run the Hyper-V Manager console locally on that server. To do this, you simply have to enable Remote Desktop on the remote Hyper-V server.

There are some downsides to this approach for managing Hyper-V, however:

- If your Hyper-V server is running on a Server Core installation of Windows Server 2008 R2, the Hyper-V Manager snap-in and Virtual Machine Connection tools are not available locally on the server, so this management approach won't work in this case.
- If you are connected to the remote desktop of a Hyper-V server and you open the Virtual Machine Connection tool on the server, you might have to contend with the following issues:
 - You will not have any mouse control inside the virtual machine unless Integration Services has been installed. The workaround is of course to make sure Integration Services is installed on your virtual machines.
 - Certain keyboard accelerators such as Ctrl+Alt+Del might have unexpected results when the Virtual Machine Connection tool is running within a Remote Desktop session. This is because Remote Desktop Connection intercepts these keyboard accelerators before the Virtual Machine Connection can see them. To resolve this, you have to modify your Hyper-V Server settings to allow Windows

keyboard accelerators to go to the virtual machine—for example, by changing the release key combination to something other than Ctrl+Alt+Left Arrow and by using the toolbar button or Action menu of Virtual Machine Connection to send a Ctrl_Alt+Del signal to the virtual machine.

Managing Hyper-V Using RemoteApp

If you want to manage Hyper-V servers from a computer running an earlier version of Microsoft Windows, such as Windows XP Professional, you can do so by using RemoteApp to publish the Hyper-V Manager application on the Hyper-V server using Remote Desktop Services. In brief, the procedure for doing this is as follows:

1. Install the Remote Desktop Services role on a server running a Full installation of Windows Server 2008 R2. Be sure to include the Remote Desktop Web Access role service in your Remote Desktop Services role installation.
2. Install the Hyper-V role or Hyper-V role management tools on the server.
3. Configure user/group membership as needed for the Remote Desktop Users and Remote Desktop Web Access Computers security groups. Also, configure RDP and security settings as needed.
4. Launch the RemoteApp Wizard from RemoteApp Manager, and add the Hyper-V Manager console (Virtmgmt.msc) to the list of published applications on the server.

Now, from the computer running the earlier Windows operating system, connect to the server using Remote Desktop Web Access, select the remotely published application (Hyper-V Manager) to launch the connection screen, and authenticate with the server. At this point, the Hyper-V Manager console will be running on your computer—it will look and work just as if the console was installed locally on your computer, with the exception of the word *Remote* in the title bar indicating that it is a RemoteApp and not a local program.

Managing Hyper-V Using Windows Management Instrumentation

Hyper-V also includes a Windows Management Instrumentation (WMI) provider that enables developers and scripters to build custom tools, utilities, and scripts for most aspects of a Hyper-V platform. This WMI provider exposes WMI classes for the following types of functionality:

- BIOS
- Input
- Integration components
- Memory
- Networking

- Processor
- Profile registration
- Resource management
- Serial devices
- Storage
- Video
- Virtual system
- Virtual system management

For example, the BIOS classes include the *Msvm_BIOSElement* class, which represents virtual BIOS software that is loaded into memory to configure and start the system, and the *Msvm_SystemBIOS* class used to associate a virtual system with its BIOS. The *Msvm_BIOSElement* class exposes properties such as *BaseBoardSerialNumber*, *BIOSGUID*, *BIOSNumLock*, *BootOrder*, and so on. Some of these properties are read-only, while others are read/write.



More Info You can find more information concerning the Hyper-V WMI provider in the MSDN Library at <http://msdn.microsoft.com/en-us/library/cc136992.aspx>.

Managing Hyper-V Using Windows PowerShell

You can also use Windows PowerShell to manage most aspects of a Hyper-V platform. Windows PowerShell is a command-line shell and task-based scripting technology that helps administrators control and automate system administration tasks. Windows PowerShell includes numerous system administration utilities, has consistent syntax and naming conventions, and enables improved navigation of common management data, such as the Windows registry, the certificate store, and WMI namespaces. Windows PowerShell also includes an intuitive scripting language specifically designed for Windows administration.

Version 2.0 of Windows PowerShell, which is included in Windows Server 2008 R2 and Windows 7, has many additional features and enhancements including the following:

- **Support for remoting** Lets you run remote commands on one or many computers using a single Windows PowerShell command.
- **Support for background jobs** Makes use of remoting to allow you to run local and remote commands in the background while using the Windows PowerShell console for other tasks. You can then retrieve the results of your jobs at your convenience.
- **Support for ScriptCmdlets** Lets you write new Windows PowerShell cmdlets in Windows PowerShell language instead of C#.

- **Integrated Scripting Environment (ISE)** A new graphical user interface for Windows PowerShell that includes a scripting interface with syntax coloring and selective execution.
- **Windows PowerShell debugger cmdlets** You can use these new cmdlets to set breakpoints, step through a script or function, and display the contents of the call stack.
- **Script internationalization** Lets you display user messages and help text for a script in the user's local language.
- **Data sections** A special section of a script or function where you can isolate your data from the script logic instead of mixing data with logic. Data sections also let you keep resource file strings together.
- **Out-GridView** A cmdlet that lets you create interactive tables from command output so that you can manipulate data using the mouse or keyboard.
- **New and improved WMI cmdlets** These include `Remove-WmiObject`, `Set-WmiInstance`, and `Invoke-WmiMethod`. In addition, all WMI cmdlets now also support the *EnableAllPrivileges* parameter as well as *Impersonation*, *Authentication*, and *Authority* parameters.
- **Improved Help files** Many new examples have been added to the Help files for various cmdlets.



More Info For more information about Windows PowerShell, see the Windows PowerShell Blog at <http://blogs.msdn.com/PowerShell/>. You can also find a collection of useful resources for learning Windows PowerShell at the Script Center on Microsoft TechNet at <http://www.microsoft.com/technet/scriptcenter/hubs/msh.mspx>.

A growing Windows PowerShell management library for Hyper-V can be found on the CodePlex Project at <http://www.codeplex.com/PSHyperv>. CodePlex is Microsoft's open-source project-hosting Web site that allows you to start a new project, join an existing one, or download software created by the CodePlex community. Note that the CodePlex site is provided by Microsoft to the developer community solely as a Web storage site and service—Microsoft does not control, review, revise, endorse, or distribute any third-party projects hosted on this site. For more information about CodePlex, see its Terms Of Use at <http://www.codeplex.com/Legal/Terms.aspx>.

At the time of this writing, the CodePlex Windows PowerShell management library for Hyper-V includes 80 Windows PowerShell functions that can be used to perform Hyper-V management tasks, such as finding a VM, connecting to a VM, discovering and manipulating the state of a VM, backing up a VM, exporting a VM, taking a snapshot of a VM, and many other common administrative tasks. Table 2-6 lists the various cmdlets currently available for different categories of Hyper-V management tasks.

TABLE 2-6 Windows PowerShell Available from CodePlex

Type of management task	Available cmdlets
Finding a VM	Get-VM Choose-VM Get-VMHost
Connecting to a VM	New-VMConnectSession
Discovering and manipulating virtual machine states	Get-VMState Set-VMState Convert-VMState Ping-VM Test-VMHeartBeat Shutdown-VM Start-VM Stop-VM Suspend-VM Get-VMKVP Add-KVP Remove-KVP Get-VMJPEG
Backing up, exporting, and taking snapshots of VMs	Export-VM Import-VM Get-VMSnapshot Choose-VMSnapshot Apply-VMSnapshot New-VMSnapshot Remove-VMSnapshot Rename-VMSnapshot Update-VMSnapshot Get-VMSnapshotTree Get-VMBackupScript
Adding and removing VMs, and configuring motherboard settings	New-VM Remove-VM Set-VM Get-VMCPUCount Set-VMCPUCount Get-VMMemory Set-VMMemory Set-VMSerialPort

Type of management task	Available cmdlets
Manipulating disk controllers, drives, and disk images	Get-VMDiskController Add-VMSCSIController Remove-VMSCSIController Get-VMDriveByController Add-VMDRIVE Remove-VMdrive Get-VMDiskByDrive Add-VMDisk Set-VMDisk Get-VMDisk Get-VMFloppyDisk Add-VMFloppyDisk Add-VMNewHardDisk
Manipulating network interface cards	Get-VMNic List-VMNic Choose-VMNIC Add-VMNIC Remove-VMNIC Set-VMNICAddress Set-VMNICConnection Get-VMNicPort Get-VMnicSwitch Choose-VMSwitch New-VMSwitchPort Get-VMByMACAddress Choose-VMExternalEthernet, New-VMExternalSwitch New-VMInternalSwitch New-VMPrivateSwitch
Working with VHD files	Get-VHDDefaultPath Get-VHDInfo New-VHD Compact-VHD Test-VHD Convert-VHD Merge-VHD Mount-VHD Unmount-VHD



Tip Because the Server Core installation option of Windows Server 2008 R2 now includes Windows PowerShell as an optional feature, you can use these Windows PowerShell cmdlets to remotely manage virtual machines running on Server Core Hyper-V servers.

Direct from the Source: Managing Hyper-V with Windows PowerShell

Hyper-V provides an MMC snap-in that can be used to manage it remotely, but no command-line tools. Fortunately, the MMC invokes a set of WMI interfaces that are all documented on MSDN, which allows for development of tools using any language capable of supporting WMI.

Windows PowerShell has built-in support for WMI with a `Get-WmiObject` cmdlet. After the cmdlet retrieves the object, its properties and methods are available to a Windows PowerShell script. Because you need to script Hyper-V operations, and because Windows PowerShell is positioned as Microsoft's scripting tool for the future, I set about developing a library of Windows PowerShell functions that are available from Microsoft's CodePlex Open Source repository <http://www.codePlex.com/PSHyperV>. These functions range in complexity from a simple `Get-VM`—which takes a name and returns the matching *MSVM_VirtualMachine* WMI object (in its raw state)—to complex tools for modifying the configuration of VMs and their virtual hard disks. The same handful of Windows PowerShell techniques are used again and again. For example, here is the basic form of *GET-VM*:

```
Function Get-VM
{
    Param ($machineName, $Server=".")

    $WQL="Select * From MSVM_ComputerSystem Where ElementName Like '$machineName' AND
    Caption Like 'Virtual%' "

    Get-WmiObject -computername $Server -Namespace "root\virtualization" -Query $WQL
}
```

As you can see, the function has two parameters. The *Server* parameter defaults to ".", the local machine. Later, I decided to default the machine name to a wildcard, which is the % sign in the WMI query language. Because I kept using *Name**, I added code to replace * with % in the name, and I added switches that modify the WMI Query Language (WQL) to return only running machines. So now I use the command like this:

```
Get-VM -Server "HV-Core"-Running | suspend-VM
```

The *MSVM_ComputerSystem* object has a method called `Request State Change`, so *suspend-VM* calls this asking for a change in state to `Suspended`. Changing the settings for a VM is a task that needs a bit more work. This involves calling the `Modify Virtual System Resources` method of the *MSVM_VirtualSystemManagementService* WMI object.

This is given a machine name, a block of XML, and a *Null* variable to contain the result. Building the XML usually means getting a WMI object, which describes the object being changed, and then calling its *getText* method and asking for XML text, like this:

```
Function Set-VMemory
{Param ($VM , $memory, $server=".")

  if ($VM -is [String]) {$VM=(Get-VM -Machinename $VM -Server $Server) }

  $memSettingData=Get-WmiObject -computerName $vm.__server Namespace
    "root\virtualization" Query "select * from Msvm_MemorySettingData
    where instanceId Like 'Microsoft:$($vm.name)%' "

  $memsettingData.Limit          =$Memory / 1MB
  $memsettingData.Reservation    =$Memory / 1MB
  $memsettingData.VirtualQuantity =$Memory / 1MB

  $arguments=@($VM.__Path, @($memsettingData.psbase.GetText([System.Management.
  TextFormat]::WmiDtd20)) , $null)

  $VSMgtSvc = (Get-WmiObject -computerName $vm.__server -Namespace
    "root\virtualization" -Class "Msvm_VirtualSystemManagementService")

  $result=$VSMgtSvc.psbase.invokeMethod("ModifyVirtualSystemResources",$arguments)

  if ($result -eq 0) {"Set memory for '$($vm.elementName)' to $memory."} else
    {"Failed to set memory for '$($vm.elementName)', result code: $result."} }
}
```

With these two techniques under your belt and documentation of the WMI provider from MSDN (not forgetting the examples from CodePlex), you can put together your own scripts to do pretty much anything.

—James O'Neill, IT Pro Evangelist, Microsoft UK

Additional Resources

The resources listed in this section were either referred to in the chapter or provide additional information concerning concepts and products discussed in this chapter.

General

A general overview of the features and benefits of Hyper-V can be found on the Windows Server 2008 R2 product information page at <http://www.microsoft.com/windowsserver2008/en/us/hyperv-main.aspx>.

For detailed technical information concerning implementing and maintaining Hyper-V, the best place to start is the Virtualization TechCenter on Microsoft Technet at <http://technet.microsoft.com/en-us/virtualization/default.aspx>. On this page in the "Products and Technologies" section, you will find a Hyper-V link that takes you to another page that displays a growing list of technical resources for IT administrators who want to learn more

about deploying and managing Hyper-V in their organizations. The following two TechNet resources are especially helpful:

- On the page “Browse Windows Server Technologies” in the TechNet Library at <http://technet.microsoft.com/en-us/library/dd283012.aspx>, scroll down and click the Hyper-V link to find useful information on how to plan, install, configure, migrate, evaluate, and get started working with the Hyper-V role in Windows Server 2008 R2. Be sure to explore each of these sections—there’s a lot you can learn in them.
- The installed Help file for Hyper-V Manager can be found online at <http://technet.microsoft.com/en-us/library/cc730764.aspx>. Be sure to familiarize yourself with the material contained in this Help file.

If you want to learn more about the hypervisor in Hyper-V, there’s no better source of information than the “Hypervisor Functional Specification v2.0: For Windows Server 2008 R2,” which is available from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=0c34932c-1bac-41a6-9b69-bc03d63ad739>.

Planning for Hyper-V

Be sure to download the latest version (4.0) of the Microsoft Assessment and Planning (MAP) Toolkit if you are planning on deploying Hyper-V in your environment. MAP 4.0 can take an inventory of your current server environment and determine which servers are underutilized. MAP can then generate recommendations for server placements and virtualization candidate assessments for Hyper-V implementation within your environment. MAP also includes a Power Savings Calculator you can use to calculate potential power cost savings with Hyper-V prior to deployment. Finally, MAP includes a VMware discovery feature you can use to identify already-virtualized servers running on VMware servers that you can manage using System Center Virtual Machine Manager or migrate to Hyper-V servers. For more information, see <http://technet.microsoft.com/en-us/solutionaccelerators/dd537566.aspx>.

The Windows Server Virtualization Guide, which is part of Microsoft’s Infrastructure Planning and Design (IPD) Guides for Virtualization, can also be used to help you design and plan for your Hyper-V deployment to ensure its success. For more information concerning the Windows Server Virtualization Guide and to download this guide, go to <http://technet.microsoft.com/en-ca/library/bb897507.aspx>.

Deploying Hyper-V

For a quick introduction on installing and using Hyper-V, see the “Step-by-Step Guide to Getting Started with Hyper-V” white paper available from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?FamilyID=bcaa9707-0228-4860-b088-dd261ca0c80d&DisplayLang=en>.

The most up-to-date resource for deploying Hyper-V is the "Hyper-V Planning and Deployment Guide," which is available from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?FamilyID=5da4058e-72cc-4b8d-bbb1-5e16a136ef42&displaylang=en>.

Managing and Maintaining Hyper-V

A useful resource to use after you've got Hyper-V installed is the "Microsoft Hyper-V Server 2008 Configuration Guide," which is available from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?familyid=E1E111C9-FA69-4B4D-8963-1DD87804C04F&displaylang=en>.

To remotely manage Hyper-V servers from a Windows 7 administrative workstation, use the Remote Server Administration Tools for Windows 7, which can be downloaded from <http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=7d2f6ad7-656b-4313-a005-4e344e43997d>.

Detailed information concerning the Hyper-V WMI provider can be found in the MSDN Library at [http://msdn.microsoft.com/en-us/library/cc136992\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/cc136992(VS.85).aspx).

A growing Windows PowerShell management library for Hyper-V can be found on the CodePlex Project at <http://www.codeplex.com/PSHyperv>.

To learn more about networking in Hyper-V, download the guide "Understanding Networking with Hyper-V" from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=3fac6d40-d6b5-4658-bc54-62b925ed7eea>.

The white paper "Windows Server 2008 R2 & Microsoft Hyper-V Server 2008 R2 - Hyper-V Live Migration Overview & Architecture" describes the Live Migration feature of Windows Server 2008 R2 Hyper-V in detail, including how Live Migration moves running virtual machines and the requirements for implementing Live Migration. You can obtain this white paper from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=fdd083c6-3fc7-470b-8569-7e6a19fb0fdf>.

The white paper "Windows Server 2008 Hyper-V and BitLocker Drive Encryption," available from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?FamilyID=2c3c0615-baf4-4a9c-b613-3fda14e84545&DisplayLang=en>, explains how to use Hyper-V and BitLocker Drive Encryption together for enhanced security.

The white paper "Performance Tuning Guidelines for Windows Server 2008" contains a section on tuning performance for the Hyper-V role and can be downloaded from Windows Hardware Developer Central (WHDC) at http://www.microsoft.com/whdc/system/sysperf/Perf_tun_srv.mspx.

Securing Hyper-V

The Hyper-V Security Guide is a Solution Accelerator that provides instructions and recommendations to help strengthen the security of computers running the Hyper-V role on Windows Server 2008. You can obtain this guide from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?familyid=2220624B-A562-4E79-AA69-A7B3DFFDD090&displaylang=en>.

Ensuring that your virtual machines are fully up to date with software patches is also important to ensure the security of these virtual machines. To help you do this, you can use the Offline Virtual Machine Servicing Tool 2.1, which provides best practices and automated tools to help you keep your offline virtualized machines updated to prevent vulnerabilities from being introduced into your IT infrastructure. This tool combines the Windows Workflow programming model with the Windows PowerShell interface to allow you to automatically bring groups of virtual machines online, service them by applying the latest security updates, and return them to an offline state. For more information about Solution Accelerator and to download it, go to http://technet.microsoft.com/en-ca/library/cc501231.aspx?SA_CE=OVMST21-Release-VIRTPROD-2009-12-07.

Note that Hyper-V is certified at the Common Criteria level EAL4 augmented by ALC_FLR.3 (also known as EAL4+). Common Criteria certification is vital, especially to government agencies, as it provides them with reassurance that Hyper-V has gone through a rigorous and internationally accepted security review. For more information about Hyper-V certification, see https://www.bsi.bund.de/cln_136/SharedDocs/Downloads/DE/BSI/Zertifizierung/Reporte05/0570b_pdf.html.

Resources for Hyper-V Developers

Developers who are interested in learning how to develop applications that use hypercalls can learn more about them in the MSDN Library at <http://msdn.microsoft.com/en-us/library/bb969694.aspx>.

Hyper-V Bloggers at Microsoft

A good place to begin traversing the blogosphere in search of the latest Hyper-V information is the Microsoft Virtualization Team Blog found at <http://blogs.technet.com/virtualization>.

John Howard, Senior Program Manager, Hyper-V team, Windows Core Operating System Division, has many helpful posts concerning Hyper-V on his blog at <http://blogs.technet.com/jhoward/archive/tags/Hyper-V/default.aspx>. Mike Kolitz, Software Design Engineer in Test on the Virtualization Technologies team at Microsoft, has lots of helpful "how-to" posts concerning Hyper-V on his blog at <http://blogs.msdn.com/mikekol/archive/tags/Microsoft+Hyper-V/default.aspx>.

James O'Neill, a Windows Server Evangelist at Microsoft UK, has tons of helpful information about managing Hyper-V using Windows PowerShell on his blog at <http://blogs.technet.com/jamesone/archive/tags/Virtualization/default.aspx>.

Another Microsoft blogger who often posts about managing Hyper-V using Windows PowerShell is the Virtual PC Guy, Ben Armstrong, Program Manager on the core virtualization team at Microsoft. You can find the Hyper-V posts on Ben's blog at http://blogs.msdn.com/virtual_pc_guy/archive/tags/Operating+Systems+under+Virtual+PC+_2F00_+Virtual+Server+_2F00_+Hyper-V/Hyper-V/default.aspx.

Matthijs ten Seldam, a Principal Consultant with Microsoft Consulting Services who focuses on virtualization, has been developing a "Virtual Server to Hyper-V import tool" that can simplify the task of taking virtual machines created on Virtual Server/Virtual PC and using them on Hyper-V. You can read more about this tool in his blog at <http://blogs.technet.com/matthts/default.aspx>.

Taylor Brown, Test Lead for Windows Core OS Division on the Hyper-V Team, has a helpful post titled "Hyper-V WMI: Creating/Applying/Deleting Virtual Machine Snapshots" on his blog at <http://blogs.msdn.com/taylorb/archive/2008/06/16/hyper-v-wmi-creating-applying-deleting-virtual-machine-snapshots.aspx>. Other useful posts about Hyper-V on Taylor's blog can be found at <http://blogs.msdn.com/taylorb/archive/tags/Hyper-V/default.aspx>.

Tony Voellm, the lead of the Hyper-V Performance Team, has a blog called All Topics Fundamental that has many useful posts about Hyper-V performance monitoring at <http://blogs.msdn.com/tvoellm/archive/2008/06/06/hyper-v-performance-faq.aspx>.

Clive Watson, Virtualization Architectural Product Technical Specialist at Microsoft UK, often posts about Hyper-V on his blog at http://blogs.technet.com/clive_watson.

The Microsoft Enterprise Support Windows Server Core Team has some helpful posts about Hyper-V on their blog ASKCORE found at <http://blogs.technet.com/askcore/archive/tags/Hyper-V/default.aspx>.

Other Hyper-V Bloggers

Mark Wilson, a Senior Customer Solution Architect for a leading IT services company, has compiled a useful list of videos from the Hyper-V product team on his blog at <http://www.markwilson.co.uk/blog/2008/07/how-hyper-v-works-product-team-videos.htm>.

Hyper-V Forum on TechNet

To obtain help with your questions and problems concerning Hyper-V, and to help others, use the Hyper-V forum on Microsoft TechNet at <http://forums.technet.microsoft.com/en-US/winserverhyperv/threads>.

Chapter 3

Local Desktop Virtualization

The term *desktop virtualization* refers to any type of technology that creates an additional virtualized operating system environment (the guest operating system installed in a virtual machine) on a physical desktop computer (the host). This additional virtualized environment can be either visible or invisible to the user. If the environment is visible to the user, the user sees a window that displays the desktop of the guest operating system within its virtual machine. (The virtual machine can also be displayed in full-screen mode, hiding the underlying host computer's desktop.) If the environment is invisible, the virtualized applications are displayed on the host computer's desktop even though these applications are actually running within the guest operating system. Such virtual applications look and behave just like native applications (applications that are locally installed on the host computer), and the user might not even know that virtualization is being used. This second scenario, where the virtual machine runs invisibly on the client computer, is sometimes known as *application virtualization*.

Desktop virtualization can be either local or remote. In *local desktop virtualization*, the virtual environment is running on the user's computer (the host), while in *remote desktop virtualization* the virtual environment runs on a server, typically a Hyper-V server or a server with the Remote Desktop Service role installed. This chapter deals with local desktop virtualization and covers three technologies currently available from Microsoft:

- **Windows Virtual PC and the Windows XP Mode Environment** Windows Virtual PC is an optional component of the Windows 7 operating system that lets users run more than one operating system at the same time on a single computer. Windows XP Mode is a preconfigured virtual machine with Windows XP Service Pack 3 (SP3) pre-installed.
- **Microsoft Enterprise Desktop Virtualization (MED-V)** MED-V is an enterprise solution for desktop virtualization that allows administrators to create, deliver, and manage corporate Virtual PC images on any Windows-based desktop.
- **Microsoft Application Virtualization (App-V)** App-V lets administrators transform applications into centrally managed virtual services to reduce the cost of application deployment, eliminate application conflicts and reboots, simplify your base image footprint to expedite PC provisioning, and increase user productivity.

We'll dig deeper into how each of these technologies works later in this chapter. But first let's examine how these technologies can benefit your organization.



Note A related virtualization technology called *user state virtualization* allows application and desktop users to virtualize their user settings and data by storing them on the network. Three Microsoft technologies make user state virtualization possible: roaming user profiles, Folder Redirection, and Offline Files. For information on how to implement these technologies in a Windows 7 and Windows Server 2008 R2 environment, refer to Chapter 14 of the Windows 7 Resource Kit from Microsoft Press. More information about this title can be found at <http://www.microsoft.com/learning/en/us/Books.aspx?Id=13811&locale=en-us>.

Examining the Benefits of Each Technology

Which of these local desktop virtualization technologies you deploy in your organization depends on what benefits they can provide to your business. The sections that follow highlight the key benefits of each technology.

Key Benefits of Windows Virtual PC and the Windows XP Mode Environment

If your business has only a small number of users and you want to migrate their computers to Windows 7 but are concerned about whether your older applications will continue to work on the new platform, Windows Virtual PC with Windows XP Mode is the answer. With Windows Virtual PC with Windows XP Mode, each user can have a separate virtual instance of Windows XP SP3 running on their computers together with older applications installed on this virtual machine. Then, when a user wants to run one of these older applications, she launches it from her Start menu just as if the application was installed locally on her computer. (The application is actually installed in and runs within the virtual machine.) In other words, the key benefit of Windows Virtual PC with Windows XP Mode is that you can use it to mitigate application compatibility issues that are blocking your desktop migration to the newest version of Windows. With Windows Virtual PC with Windows XP Mode, you can migrate your desktop computers to Windows 7 and take advantage of all the new and exciting features of Windows 7 while still being able to run your older applications.

The limitation with Windows Virtual PC with Windows XP Mode is the lack of centralized management functionality—each virtual machine running on each desktop computer must be managed locally on that desktop computer. This means that Windows Virtual PC with Windows XP Mode is mainly suitable for small networks of computers. If you are contemplating migrating a larger network to Windows 7 but are concerned about application compatibility, MED-V provides the answer as described next.

Key Benefits of MED-V

The ability for MED-V to allow you to deploy Virtual PC images onto Windows desktops and to manage them while maintaining a seamless end-user experience can provide businesses with many advantages. The key benefit from deploying MED-V is being able to maintain support for running legacy applications when upgrading desktop operating systems. MED-V allows you to run legacy applications in a virtual machine running an older version of Microsoft Windows, thus accelerating the deployment of the latest version of Windows by resolving application incompatibility issues. MED-V also lets you test your migration plans using virtual machines instead of physical computers, and it reduces user training costs by making virtual desktops invisible and seamlessly integrating legacy applications into the user's local desktop.

Other benefits of deploying MED-V can include

- **Accelerating application development** MED-V reduces the time and work involved in deploying and reconfiguring applications on users' desktops, and it increases quality assurance by allowing you to test and document application functionality on multiple operating systems using virtual machines.
- **Centralizing desktop management and deployment** MED-V lets you use policies to lock down corporate virtual machines and to easily deploy managed virtualized applications to any desktop computer, including less controlled assets such as employee PCs, contractor PCs, and desktop computers in partner subsidiaries, branch offices, and offshore operations. And because virtual machines deployed using MED-V live locally on the user's computer, users have access to their virtual desktop and applications even when their computers are disconnected from the corporate network. This makes MED-V an ideal solution for enterprises that have a large proportion of mobile users with laptops.
- **Driving business continuity** MED-V allows you to rapidly reconstitute corporate-managed virtual desktops on any Windows computer independent of the underlying hardware. MED-V also lets you test and deploy software on different versions of Windows more easily because one failed application or operating system won't affect others.

Key Benefits of App-V

Some of the key benefits of deploying an App-V infrastructure within your organization include

- Centralized management of the entire application life cycle
- Faster application deployment
- Simplified application versioning

- Fewer side-by-side application compatibility issues
- Reduced need for regression testing
- On-demand application delivery
- Integrates with existing Terminal Services or electronic software distribution (ESD) infrastructures

See also the sidebar titled “MED-V and App-V: Comparing the Benefits” later in this chapter.

MED-V and App-V: Comparing the Benefits

The key benefit of MED-V is that it helps enterprises deal with incompatibility between applications and the operating system. For instance, if a user needs to run an early version of Internet Explorer and that version of Internet Explorer is not supported on Windows Vista, the administrator can use MED-V 1.0 to deploy this early version of Internet Explorer to the user as part of a Windows XP virtual image. (And when MED-V 1.0 SP1 becomes available in Q1 of 2010, the user will be able to do the same thing on computers running Windows 7.) The user can then have two copies of Internet Explorer running simultaneously on his desktop—the most recent version (running on the host computer) and the earlier version (running in the MED-V workspace). From the user’s perspective, both copies of Internet Explorer appear as if they were running on the local computer. MED-V does this by allowing users to run legacy applications within a virtual machine that has an earlier version of Microsoft Windows installed. The user can then access these applications either from a virtual desktop (as with Virtual PC 2007 running natively on a system) or by using application windows that are seamlessly integrated into the local desktop of the user’s computer (similar to RemoteApp in Remote Desktop Services).

Microsoft Application Virtualization (App-V) also helps enterprises handle application compatibility issues, but it addresses challenges differently than MED-V does. Specifically, App-V lets you resolve conflicts that arise between different applications or different versions of the same application; MED-V, on the other hand, allows users to run older versions of Microsoft Windows concurrently with the local desktop of their computers, which can help with issues where legacy applications are unable to run natively on the most recent version of Windows installed on the user’s computer.

Examining Usage Scenarios for Each Technology

Each of these local desktop virtualization technologies can be implemented in various ways to bring benefit to an organization. The sections that follow describe some of these scenarios.

Usage Scenarios for Windows Virtual PC and the Windows XP Mode Environment

The key usage scenario for Windows Virtual PC and Windows XP Mode is to resolve application-to-operating system incompatibility on a machine-by-machine basis. A typical scenario is a small business that has computers running Windows XP and that has older applications installed that don't work properly when installed on Windows 7. Typically, such a situation would be a showstopper that would prevent the business from migrating its computers to Windows 7 because small businesses typically don't have the resources to shim applications using the Application Compatibility Toolkit, and a small business also might not have the budget to purchase newer versions of these applications (if such newer versions exist). In such a scenario, Windows Virtual PC and Windows XP Mode can enable the business to upgrade its desktop computers while still enabling older applications to run properly—all at no extra cost because Windows Virtual PC and Windows XP Mode are free downloads from Microsoft.

Usage Scenarios for MED-V

The key usage scenario for MED-V is resolving application-to-operating system incompatibility to accelerate the upgrade path to a new operating system. Businesses that need to continue to run legacy line-of-business applications on users' desktop computers can do so by using Virtual PC. Incompatibility between legacy applications and newer versions of Microsoft Windows can often be a primary blocking issue preventing an enterprise from upgrading to the latest version of Windows, such as Windows Vista, to take advantage of the many new features and enhancements offered by this version. By delivering those applications in a Virtual PC that runs a previous version of the operating system (for example, Windows XP or Windows 2000), MED-V allows administrators to break the tight dependency between a computer's underlying hardware and the operating system, and it can help remove such blocking issues so that your users can benefit from having the latest version of Windows deployed on their desktop computers. From the user's perspective, with MED-V, these applications are accessible from the Start menu and appear side by side with regular applications—so there is minimal change to the user experience.

Usage Scenarios for App-V

App-V supports a wide range of different usage scenarios, ranging from a full application virtualization infrastructure to a lightweight infrastructure to standalone deployment. Specifically, App-V supports the following usage scenarios:

- **Full Infrastructure** This scenario uses the App-V Management Server, which provides full streaming capabilities, Desktop Configuration Service, active/package upgrade, and basic licensing and metering. This infrastructure requires Active Directory and SQL

Server and is an update to the existing SoftGrid Virtual Application Server that version 4.2 customers are familiar with using.

- **Lightweight Infrastructure** This scenario uses the App-V Streaming Server, which includes streaming capabilities such as active/package upgrade without the Active Directory or SQL Server requirements. However, it does not have a Desktop Configuration Service or licensing or metering capabilities. This service relies on the manual or scripted addition of a manifest file for virtual application configuration. The Desktop Configuration Service of the App-V Management Server can also be used in conjunction with the App-V Streaming Server such that the Management Server configures the application but the Streaming Server delivers it.
- **Standalone mode** The App-V Sequencer has an option to create an .msi file that automates the addition of the virtual application. The .msi contains metadata so that an ESD system can recognize it and control the virtualized applications. Standalone mode requires the App-V Client to go into Standalone mode, which allows only .msi-based updates of the virtual applications. (Streaming is not allowed while in Standalone mode.) This mode is meant for rarely connected users that need the power of virtualized applications but do not have access to a server.

For more information about various App-V usage scenarios, see the section titled “App-V Deployment Scenarios” later in this chapter.

Availability of Each Technology

Now that we’ve examined the benefits and usage scenarios for each technology, let’s look at how you can obtain them from Microsoft.

Availability of Windows Virtual PC and the Windows XP Mode Environment

Windows Virtual PC and Windows XP Mode are available for Windows 7 as free downloads from Microsoft. To download the appropriate version of these products, go to <http://www.microsoft.com/windows/virtual-pc> and click Get Windows XP Mode And Windows Virtual PC Now. This takes you to a page where you can specify your edition and language of Windows 7 so that you can download the correct installation packages for each product.

Availability of MED-V

MED-V 1.0 is available to volume-licensed customers as part of the Microsoft Desktop Optimization Pack (MDOP) for Software Assurance, an add-on subscription license available

to Software Assurance (SA) customers. The current version as of this writing (MDOP 2009) includes five other key technologies that help enterprises manage their desktops more easily:

- Microsoft Application Virtualization
- Microsoft Asset Inventory Service
- Microsoft Advanced Group Policy Management
- Microsoft Diagnostics and Recovery Toolset
- Microsoft System Center Desktop Error Monitoring

For more information about MDOP, see <http://www.microsoft.com/windows/enterprise/products/mdop/default.aspx>.

For more information about Microsoft's Software Assurance (SA) licensing program, see <http://www.microsoft.com/licensing/sa/default.mspx>.

Availability of App-V

Microsoft Application Virtualization is also available as part of the Microsoft Desktop Optimization Pack (MDOP) for Software Assurance. For more information about MDOP, see <http://technet.microsoft.com/en-us/windows/bb899442.aspx>.

Understanding Windows Virtual PC and the Windows XP Mode Environment

As described earlier in this chapter, Windows Virtual PC and the Windows XP Mode Environment complement each other to provide businesses with a simple way of mitigating application compatibility issues that can block deployment of the latest versions of the desktop Windows operating system. This section examines how Windows Virtual PC and the Windows XP Mode Environment work and then demonstrates how to install, configure, and use them together.

Understanding Windows Virtual PC

Figure 3-1 shows a high-level view of the architecture of Windows Virtual PC. The left side of the diagram shows the components of the host operating system, while the right side shows the components of the guest.




FIGURE 3-1 Architecture of Windows Virtual PC.

The architecture of the host side of Windows Virtual PC can be further broken down into user interface (UI) components, user-mode engine components, and kernel-mode engine components.

Host-Side Components

The UI components of the Windows Virtual PC host include the following:

- **VPC Settings dialog** Lets you modify configuration options such as networking, memory, integration features, and virtual hard disks for each virtual machine.
- **VPC Wizard** Walks you through the steps of creating new virtual machines.
- **VM Window (VMWindow.exe)** When you start a virtual machine (VM), an instance of VMWindow.exe is launched to manage the display window that you use to interact with that VM. VMWindow.exe also loads MSTSCAX.dll, which functions as a Remote Desktop Protocol (RDP) client and is essentially the same ActiveX control that is used to run RemoteApps and Remote Desktops from a Remote Desktop Web Access server in Windows Server 2008 R2. One unique instance of VMWindow.exe is launched for each

running VM, except for VMs running virtual applications, which is discussed in the next bullet.

- **VM SAL (VMSAL.exe)** When you launch a virtual application from the host, an instance of the Virtual Machine Seamless Application Launcher is launched to initiate, monitor, and control the application. As with VMWindow.exe, the in-process ActiveX control MSTSCAX.dll acts as the RDP client.

The user-mode engine components of the Windows Virtual PC Host include the following:

- **VPC (VPC.exe)** The core Virtual PC engine that manages virtual machines and provides services for them. VPC.exe includes the following subcomponents that provide specialized services for virtual machines: RDP Encoder Technology, device emulators, COM servers, Network Address Translation (NAT), and Integration Components (ICs). VPC.exe also provides a set of COM APIs you can use to develop custom applications for performing tasks such as creating and managing virtual machines, creating and managing virtual hard drive (VHD) images, and modifying the configuration settings of VMs.
- **RDP ET (RDP Encoder Technology)** A group of components that uses RDP to provide the console experience for accessing a virtual machine and converts keyboard, mouse, and video actions between the RDP format and the format used by the VM device emulators.
- **Devices** Device emulators for devices such as virtual hard drives, COM ports, and network interfaces.
- **COM port redirector** Provides access for the virtual machine to remote serial devices such as modems.
- **NAT** Allows a virtual machine to use the physical network adapter for network connectivity.
- **Integration Components (ICs)** Provides advanced features such as video resizing and audio redirection within virtual machines.

The kernel-mode engine components of the Windows Virtual PC Host include the following:

- **Virtualization Server Provider (VSP)** Provides I/O device-related resources to Virtualization Service Clients (VSCs) running in virtual machines.
- **VPCBus.sys** A kernel-mode bus driver used by the VSP to communicate between the host and guests.
- **VMM.sys** The Virtual Machine Monitor, which virtualizes the physical processing resources across the host and virtual machines and provides resource management, including memory and interrupts.

- **USB Connector (vpcusb.sys)** Provides USB virtualization to the guest operating systems, and manages the virtual root hubs for connected USB devices. Each virtual machine has one virtual hub that can be assigned between zero and eight devices.
- **USB Stub Driver (vpcuxd.sys)** A stub driver that is loaded by the operating system in lieu of the normal USB client driver.

Guest-Side Components

The architecture of the guest side of Windows Virtual PC can be further broken down into Integration Components, RAIL (Remote Applications Installed Locally)/RDP components, and kernel-mode components.

The Integration Components of a Windows Virtual PC guest include the following two services, which provide Integration Component services to the guest:

- Virtual PC Integration Components Services Application service (VMSvc.exe)
- Virtual Machine User Services (VMUSvc.exe)

The RAIL/RDP components of a Windows Virtual PC guest include the following:

- **RDP Server service** Listens for RDP connections from the RDP clients running in a virtual machine window or application.
- **RDP Shell (RDPShell.exe)** A shell designed to present virtual applications as if they are running locally on the host and to make the seamless running of virtual applications possible.



Note For more information about RAIL, see the sidebar titled “Direct from the Source: RAIL and RemoteApp” later in this chapter.

The kernel mode of a Windows Virtual PC guest includes the following:

- **VSC** Consumes resources provided to it by the VSP running on the host.
- **VMX/SVM Root Kernel** Built upon the Virtual Machine Extensions (VMX) of Intel Virtualization Technology (Intel VT) technology. It includes the Virtual Machine Monitor (VMM) runtime layer, which provides support for virtual machine execution, memory management, intercept and exception handling, and routing of interrupts raised by virtual machines. For more information, see the sidebar titled “Direct from the Source: Windows Virtual PC vs. Hyper-V” later in this section.

Direct from the Source: Windows Virtual PC vs. Hyper-V

Windows Virtual PC is not built on hypervisor technology in the way that Hyper-V server is. Instead, Windows Virtual PC uses the VMX kernel to provide support similar to that provided by the hypervisor.

In Virtual PC and Virtual Server, device support was done primarily through emulation of hardware; this is also done in Windows Virtual PC. In Windows Virtual PC, the disk, network, and display subsystems manifest themselves as physical devices that are detected by the guest operating system at boot and are indistinguishable (to the guest) from real hardware. The drivers for these corresponding devices get loaded by the guest operating system and execute I/O commands as they would in a real environment. These I/O commands are intercepted by the VMM runtime, which is the VMX/SVM kernel that triggers callbacks of device emulators running within the user mode process VPC.exe. Windows Virtual PC uses VPCBus-based devices coexisting with the current device framework.

—CSS Global Technical Readiness (GTR) Team

Understanding Virtual Applications

Virtual applications are applications installed on virtual machines that run on the desktop of the host. From the user's perspective, a virtual application is launched the same way as a local application (an application installed on the host)—that is, by clicking on the application's shortcut in the Start menu and similar methods. Virtual applications are a key feature of Windows Virtual PC. They let you transparently run applications that are not fully compatible with the host operating system by running them in a guest operating system.

Two Windows Virtual PC components play a key role in allowing applications installed on the guest to run transparently on the host computer's desktop:

- **VMSAL.exe** When you start a virtual application from the host, an instance of VMSAL.exe is launched to initiate, monitor, and control the virtual application. An in-process ActiveX control called MSTSCAX.dll then performs the role of the RDP client and creates a named pipe to connect to the appropriate RDP server service on the guest using the TCP port on which the service is listening, which by default is 3389. The difference between opening a virtual application and opening the virtual machine is that in the first case the management process on the host is VMSAL.exe, while in the second case it is VMWindow.exe.
- **RDPSHELL.exe** To display a virtual application as if it is running locally (instead of displaying a full virtual desktop), a unique virtual application shell called RDPSHELL.exe is started when the RDP session is initiated.

Launching a virtual application works like this:

1. Starting the application initiates a Remote Desktop session on the server.
2. The normal logon process (WinLogon.exe) calls the user initialization process (UserInit.exe) to process group policies, run logon scripts, and perform similar tasks.
3. If the logon was initiated by starting a virtual application, UserInit.exe then loads RDPInit.exe, which is an initialization process specific to virtual applications.
4. RDPInit.exe then loads RDPShell.exe instead of loading the standard desktop shell (Explorer.exe) so that the remote application is presented to the user as if it is running locally.

Enabling Support for Virtual Applications

The Windows XP Mode virtual machine is preconfigured to support virtual applications, making it easy to use this environment to mitigate application compatibility issues that might be blocking your desktop refresh cycle. However, any virtual machine installed on Windows Virtual PC can run virtual applications if the guest operating system has RemoteApp support enabled. Windows 7 supports RemoteApp by default, but for earlier operating systems you must download a software update to provide RemoteApp support. Specifically,

- If your guest operating system is running Windows XP SP3 and you want to install virtual applications on the guest, you must install the KB961742-v3.exe update, which is available from <http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=e5433d88-685f-4036-b435-570ff53598cd>. For more information concerning this update, see <http://support.microsoft.com/kb/961742>.
- If your guest operating system is running Windows Vista SP1 or later and you want to install virtual applications on the guest, you must install the Windows6.0-KB961741-x86.msu update, which is available from <http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=26a2de17-8355-4e8d-8f33-9211e48651fb>. For more information concerning this update, see <http://support.microsoft.com/kb/961741>.

After you have installed the appropriate update in your guest operating system, you must also perform the following additional steps to enable virtual application support on the virtual machine:

- You must install the Integration Components on the virtual machine.
- You must enable the Auto Publish setting on the virtual machine.

For more information on performing these tasks, see the section titled “Configuring Virtual Machine Settings” later in this chapter. For more information on how RemoteApp is used by Windows Virtual PC, see the sidebar titled “Direct from the Source: RAIL and RemoteApp” later in this section.

Direct from the Source: RAIL and RemoteApp

A key technology that enables virtual application support in Virtual PC is commonly referred to as RAIL (Remote Applications Installed Locally). This is the same technology that is used in Remote Desktop Services in Windows Server 2008 and later to enable RemoteApps (Remote Applications). RemoteApps are applications that are installed and executed on a Terminal Server or Remote Desktop Session Host server, but they are integrated with the Remote Desktop client machine in a way that makes them appear to be running locally on the client. Although a Remote Desktop Session has been established to a server, the user does not see the desktop of that session; he just interacts with the application as if it was installed locally. These applications are referred to as *virtual applications*, but when the terms *RAIL* or *RemoteApp* are used in the context of Virtual PC, they should be considered synonymous with “virtual applications.” For more information, see the section titled “Understanding Virtual Applications” earlier in this chapter.

—CSS Global Technical Readiness (GTR) Team

Installing Virtual Applications

After support for virtual applications has been enabled on a virtual machine, you can publish applications installed on the virtual machine to the host simply by installing the application on the guest operating system.



Tip When you install an application on a guest, if an option is displayed that allows you to choose whether the application should be made available for either the current user or for all users, you must select the All Users option for the application to be published. If the installation routine does not install an application shortcut to the All Users Start menu, you must copy or move the application’s program groups and shortcuts to the %SystemDrive%\Documents and Settings\All Users\Start Menu\Programs folder for the application to be published.

Understanding Windows XP Mode

Windows XP Mode is a virtual machine and a virtual hard disk (VHD) that has Windows XP SP3 x86 pre-installed and pre-activated on it. The virtual machine has also been preconfigured for publishing virtual applications to provide businesses with a rapid way of mitigating application compatibility issues that can block desktop deployment of the latest version of Windows. When older applications are installed into this virtual machine, users can transparently run such applications on the host even if the applications cannot be installed locally on the host.

Installing and using a Windows XP Mode virtual machine requires that Windows Virtual PC has already been installed on the host computer. You cannot use Windows XP Mode with older Microsoft desktop virtualization platforms such as Virtual PC 2007 or with Microsoft Virtual Server or Hyper-V.

When Windows XP Mode is installed, two virtual hard disks are created on the host computer:

- A parent virtual hard disk named Windows XP Mode base.vhd located in the %SystemDrive%\Program Files\Windows XP Mode folder. This parent disk is write-protected and approximately 1.2 GB in size.
- A differencing virtual hard disk named *VM_name.vhd*, where *VM_Name* is the name of the virtual machine. This differencing disk varies with size (it grows as needed) and is located in the hidden %SystemDrive%\Users\username\AppData\Local\Microsoft\Windows Virtual PC\Virtual Machines folder, where *username* is the user's profile folder. The virtual machine configuration file (.vmc file) for the virtual machine is also located in this folder.



Tip You should back up the parent disk in case it becomes corrupted, because the differencing disk won't work without the parent.

Windows XP Mode is designed for running older business productivity applications that cannot be installed on Windows 7 because of compatibility issues. Windows XP Mode is not designed to run graphic-intensive applications such as games or AutoCAD.



Tip You should install antivirus and antimalware software on your Windows XP Mode virtual machine. You should also make sure that Automatic Updates is enabled on the virtual machine.

Direct from the Source: Comparing XP Mode to a Custom XP Virtual Machine

Windows XP Mode does not differ much from a virtual machine that a user can manually create and then install Windows XP SP3 into. The primary differences are

- The Windows XP Mode installation is pre-activated.
- The Windows XP Mode installation is preconfigured.

The main steps used to install and configure the Windows XP Mode virtual machine are

1. Install Windows XP SP3.
2. Install the latest Windows Virtual PC Integration Components.

3. Install the Windows XP RemoteApp update required for virtual applications.
4. Apply all of the latest mandatory Windows XP updates (as of the date the XP Mode VM was built).

None of the features of Windows Virtual PC, such as virtual application support, are limited to Windows XP Mode VMs. A user can use the steps just listed to create her own XP SP3 virtual machine (with modifications if desired) and use the same feature set as she can with a Windows XP Mode VM.

—CSS Global Technical Readiness (GTR) Team

Requirements for Windows Virtual PC

This section outlines the requirements for installing Windows Virtual PC.

Host Operating System

Windows Virtual PC requires that Windows 7 be installed as the host operating system. Specifically, Windows Virtual PC can be installed on host computers that have any of the following Windows 7 editions installed:

- Windows 7 Starter (32-bit only)
- Windows 7 Home Premium (32-bit or 64-bit)
- Windows 7 Professional (32-bit or 64-bit)
- Windows 7 Ultimate (32-bit or 64-bit)
- Windows 7 Enterprise (32-bit or 64-bit)

Guest Operating System

The supported guest operating systems (32-bit only) for Windows Virtual PC are as follows:

- Windows XP Professional SP3
- Windows Vista Business SP1 or later
- Windows Vista Ultimate SP1 or later
- Windows Vista Enterprise SP1 or later
- Windows 7 Professional
- Windows 7 Ultimate
- Windows 7 Enterprise

Note that the virtual application feature of Windows Virtual PC is not supported for the following guest operating systems:

- Windows Vista Business SP1 or later
- Windows 7 Professional

CPU

Windows Virtual PC requires a processor capable of hardware virtualization using either AMD-V or Intel VT technology. In addition, hardware virtualization must be turned on in the BIOS of the host computer.



Tip You can use the Microsoft Hardware-Assisted Virtualization Detection Tool to determine whether your computer supports AMD-V or Intel VT hardware virtualization. You can obtain this tool from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?FamilyID=0ee2a17f-8538-4619-8d1c-05d27e11adb2&displaylang=en>.

Memory

Microsoft recommends that your host computer have at least 2 GB of memory to use Windows Virtual PC.

Installing Windows Virtual PC

To install Windows Virtual PC, begin by downloading the appropriate version (32-bit or 64-bit) from the Microsoft Download Center at <http://www.microsoft.com/windows/virtual-pc/>. Then double-click on the downloaded .msu file and follow the prompts to install Windows Virtual PC.

After Windows Virtual PC has been installed on your computer, a Windows Virtual PC program group and shortcuts are added to your Start menu. (See Figure 3-2.)




FIGURE 3-2 The Windows Virtual PC program group on the Start menu.

Requirements for Windows XP Mode

This section outlines the requirements for installing Windows XP Mode.

Host Operating System

Windows XP Mode requires that one of the following Windows 7 editions be installed on your host computer:

- Windows 7 Professional (32-bit or 64-bit)
- Windows 7 Ultimate (32-bit or 64-bit)
- Windows 7 Enterprise (32-bit or 64-bit)

In addition, your host computer must have Windows Virtual PC installed and must meet all of the requirements of Windows Virtual PC.

Disk Space

The disk space requirements for Windows XP Mode are as follows:

- At least 2 GB for installing Windows XP Mode
- An additional 15 GB for the virtual Windows environment

Installing the Windows XP Mode Environment

To install Windows XP Mode on a host computer that has Windows Virtual PC installed, begin by downloading the appropriate language version from the Microsoft Download Center. For example, the English language version of Windows XP Mode (WindowsXPMode_en-us.exe) can be downloaded from <http://www.microsoft.com/windows/virtual-pc/download.aspx>. Note that the Windows XP Mode download is almost 500 MB in size.



Tip If you need a localized version of Windows Virtual PC and Windows XP Mode, you can download it from <http://www.microsoft.com/windows/virtual-pc/download.aspx>.

After you have downloaded the installation file, double-click on it to begin installing Windows XP Mode. During installation, you have the option of specifying where you want the VHD file installed. (See Figure 3-3.)




FIGURE 3-3 Specify where you want to install the Windows XP Mode VHD file.

When Setup is finished, leave the Launch Windows XP Mode check box selected so that you can begin configuring Windows XP Mode. (See Figure 3-4.)

Read and accept the licensing agreement, and click Next. Then on the Installation Folder And Credentials page of the wizard, type a password for the XPMUser account. (See Figure 3-5.) This account is created by the wizard and has administrator privileges on the guest operating system. You can leave the Remember Credentials check box selected so that you won't have to enter the password each time you start the virtual machine.




FIGURE 3-4 Leave the check box selected to begin configuring Windows XP Mode on your computer.




FIGURE 3-5 Configuring credentials for running virtual applications.

On the next wizard page, select the option to enable automatic updates to keep the virtual machine up to date with the latest software updates. When you finish the wizard, a screen is displayed indicating that Windows XP Mode is being set up on your computer. (This might take several minutes.)

Once Windows XP Mode has been installed, the virtual machine starts and a window opens displaying the desktop of the virtual machine (as shown in Figure 3-6).




FIGURE 3-6 The Windows XP Mode virtual machine console window open on the desktop of the host computer.



Tip You can resize this virtual machine console window by dragging its corners.

To view the installed virtual machines on your computer, click Start, click All Programs, click Windows Virtual PC, and click Windows Virtual PC. This opens an Explorer window showing the installed virtual machines (.vmcx files) and their status, memory used, VHD file, and configuration file. (See Figure 3-7.)

The toolbar of this Explorer window lets you create new virtual machines on the host computer. In addition, when the virtual machine is selected a Settings button is displayed that lets you configure virtual machine settings, which are described in the next section.



Tip You can also right-click on a virtual machine displayed in the Virtual Machines folder and select Settings from the context menu to modify the virtual machine's settings.




FIGURE 3-7 Viewing the installed virtual machines on the host computer.

Configuring Virtual Machine Settings

Windows Virtual PC lets you easily modify the settings for a virtual machine. To do this, open the Virtual Machines folder as shown in the previous figure, select the virtual machine you want to modify, and click **Settings** on the folder toolbar. Doing this opens the Settings dialog for the selected virtual machine as shown in Figure 3-8.




FIGURE 3-8 Configuring the settings for a virtual machine.

The sections that follow examine some of the settings you can configure for a virtual machine.

Name

You can use the Name settings page to assign a descriptive name to the virtual machine. This descriptive name is displayed in the Virtual Machines folder, and it is not the same as the NetBIOS name you assign in the guest operating system. For example, the descriptive name for a Windows XP Mode virtual machine is Windows XP Mode.vmcx, while the NetBIOS name assigned by default to the guest operating system is VirtualPC-#####, where ##### is a randomly assigned numeric string. Note that you must shut down the virtual machine before you can modify this setting.

Memory

You can use the Memory settings page to assign some of your host computer's physical memory to the virtual machine. Note that virtual machines consume such memory only when they are running. Be sure to assign sufficient memory so that the virtual machine runs well, but not so much as to starve the host computer of the memory it needs. Note that you must shut down the virtual machine before you can modify this setting.

Hard Disk

You can use the Hard Disk settings pages to do the following:

- Create a new virtual hard disk, and attach it to the virtual machine. You can create any of the following types of virtual hard disks:
 - Dynamically expanding
 - Fixed-size
 - Differencing
- Attach an existing virtual hard disk to the virtual machine.
- Compact, convert, or merge existing virtual hard disks connected to the virtual machine.

Note that performing any of these actions requires that you first shut down the virtual machine.

Note also that for Windows XP Mode, the Hard Disk 1 page displays the settings for the differencing disk and you can add up to three additional virtual hard disks to the virtual machine.

Undo Disks

You can use the Undo Disks settings page to configure an undo disk for your virtual machine. Undo disks enable changes to a virtual disk to be saved to a separate undo disk file (.vud file) while keeping the original disk unmodified. When Enable Undo Disks is selected, the undo disk setting configured here applies to all virtual hard disks attached to the virtual machine. Undo disks are typically used in testing environments, so this feature is disabled by default for all virtual machines. Note that you must shut down the virtual machine before you can modify this setting.



Note Changes made using undo disks can be applied or discarded when the virtual machine is turned off. Changes cannot be selectively applied or discarded; it is all or nothing.

DVD Drive

You can use the DVD Drive settings page to specify a physical or virtualized DVD drive for use by the virtual machine (as shown in Figure 3-9). When a physical DVD device has been attached to the guest, the device is unavailable on the host while the virtual machine is running. You can also select an ISO image file on the host computer and attach it to the DVD drive of the virtual machine, which makes it easy to install applications in the virtual machine because you do not need to burn DVD media first.




FIGURE 3-9 Configuring DVD Drive settings for a virtual machine.

COM

You can use the COM port settings pages to specify a physical COM port on the host machine, a named pipe, or a text file. Possible uses for these settings include the following:

- Select the Physical Serial Port option to bind the virtual machine to the specified physical serial port. Optionally, select the Wait For Modem Command To Open Port check box if you want the serial port to be captured when the virtual machine attempts to access the COM port. (Selecting this check box causes the virtual machine to wait for the AT modem command to be sent to the port.)
- Select the Named Pipe option to allow the virtual machine to connect to a Windows named pipe on the host or across the network. Named pipes can be used to create virtual null modem cables between two virtual machines to allow for kernel or user-mode debugging of the virtual machine.
- Select the Text File option to send COM port output from the virtual machine to a text file for troubleshooting purposes. You can configure the output location of the text file when you choose this option.

Networking

You can use the Networking settings page (shown in Figure 3-10) to attach up to four virtual network adapters to a virtual machine, and the network connectivity for each network adapter can be configured as any of the following:

- **Not Connected** The virtual machine has no network connectivity through this network adapter.
- **Internet Network** Select this option to enable the virtual machine to communicate only with other virtual machines running on the same host. The virtual machine will be isolated from the host's network, but you will still be able to access the virtual machine using RDP and run virtual applications on it. Choose this approach if you will not be running any network-aware applications on your virtual machine and you want to completely isolate the virtual machine from the host computer's physical network for increased security.
- **Shared Networking (NAT)** Select this option to enable the virtual machine to share the host computer's physical network adapter. The virtual machine will have an IP address of the form 192.168.131.x automatically assigned using a built-in Dynamic Host Configuration Protocol (DHCP) server that has an IP address of 192.168.131.254. The virtual machine is then connected to the host computer's physical network through a NAT gateway that uses the same address as the DHCP server. The result is that the virtual machine can access the host computer's physical network but the virtual machine is not displayed on this network as a separate computer. Choose this approach if you

need to move your host computer between different networks—for example, from a local area network (LAN) connection to a virtual private network (VPN) connection.

- **NIC_name** This option displays the name of the physical network interface card (NIC) on the host. If you have more than one physical NIC, each will be displayed as a separate option here. This option, sometimes called *Bridge mode*, bridges the virtual network adapter to the selected physical NIC and allows the virtual network adapter to acquire an IP address from a DHCP server running on the host computer's physical network. The result is that the virtual machine is displayed on the host computer's physical network as a separate computer from the host computer.




FIGURE 3-10 Configuring Networking settings for a virtual machine.

You can change the network connectivity of a virtual network adapter without needing to shut down the virtual machine. To add a new virtual network adapter, however, you must first shut down the virtual machine.



Note Bridge mode has poorer performance than Shared Networking (NAT) but is more secure because the guest operation system is hidden behind a NAT.

Integration Features

You can use the Integration Features page to selectively enable or disable Integration Features for the virtual machine provided that Integration Components (ICs) have been installed on the virtual machine. (ICs are installed by default on a Windows XP Mode virtual

machine). If access to drives on the host is enabled, you can also selectively enable or disable access by the guest to drives on the host. (See Figure 3-11.)




FIGURE 3-11 Configuring Integration Features settings for a virtual machine.

For more information on Integration Features and Integration Components in Windows Virtual PC, see the sidebar titled “Direct from the Source: Understanding Integration Features” later in this chapter.

Direct from the Source: Understanding Integration Features

Integration of the host operating system with a guest virtual machine can be described as “none,” “basic,” or “enhanced.” Host-guest integration features are not available until the Integration Components (ICs) are installed on a guest VM. Basic integration is activated by *installing* the ICs on the guest. Enhanced integration is activated by *enabling* the Integration Features.

No Integration

When the ICs are not installed on a guest VM, the only built-in interaction between the host and the guest is through an emulated “native” console (VMWindow.exe). The benefit of this method is that the user can interact with the virtual machine from the beginning of the boot process before the guest operating system has been initialized, and it even allows them to enter the setup for the BIOS. The main disadvantage of this method is that the user does not get the benefit of some Integration Features, such as the redirection of host devices to the guest machine. For example, mouse movement will not be seamless across host and guest. The mouse will be stuck inside the guest

window, or it will be outside of the guest window. The user will have to use the host key combination, Ctrl+Alt+Left Arrow, to leave the guest window.

Basic Integration Mode

Basic integration mode is made available by installing the ICs on a guest VM. Basic mode Integration Features include

- **Time synchronization between guest and host** This feature is responsible for synchronizing the date and time on the guest operating system with the host. For example, if the time (System clock) is out of sync because the system is sleeping or hibernating, the time synchronization feature is used to force an adjustment immediately after the machine is awakened. The time synchronization feature also periodically adjusts the guest VM time if it is out of sync because of clock drift (lost clock interrupts). Additionally, the time settings of the guest are synchronized with the host during a boot operation. Note that the synchronization operation does not affect the time zone or Daylight Savings Time settings on the guest. The date and time are adjusted based on the GMT date and time derived from the host machine.
- **Guest heartbeat monitoring** The heartbeat provides a mechanism to detect an unresponsive (crashed, blue-screened, or hung) guest VM. Heartbeat requests are sent every 60 seconds by the host. If the guest is unresponsive, three retries are made for getting the heartbeat response. If the correct heartbeat response is not received within the allowed heartbeat latency, the guest is declared dead. The heartbeat information flow consists of a heartbeat request sent from the host to the guest followed by a heartbeat response sent from the guest to the host.
- **Host-initiated shutdown** This feature allows shutdown of the guest operating system with a request message generated by the host. COM APIs provide the messaging channel. The options available through the virtual machine window are Sleep, Restart, Hibernate, Shutdown, and Turn Off. The API additionally provides the option to log off the current user.
- **Guest video resize (arbitrary guest resolution)** This is a feature of Virtual PC (VPC) by which the guest resolution is changed in basic mode when the guest window is resized such that the guest desktop fits in the host side VM window.
- **Mouse movement across host and guest** The same “native” console emulation available when Integration Components are not installed is available in basic mode. But because the mouse integration component is activated in basic mode, users will be able to move their mouse cursor inside and outside of the virtual machine window without using the release key.

Enhanced Mode

Installing the Integration Components and enabling the Integration Features on a guest VM activates the enhanced host-guest integration mode. When the Integration Features are enabled, a console experience that more resembles a “classic” Remote Desktop Connection is enabled. This experience is very similar to the one a user would observe when establishing an RDP connection from a Remote Desktop Connection client (Mstsc.exe) to a Terminal Server or Remote Desktop Session Host server. This experience includes the ability to redirect the local devices in the following list. Enhanced mode includes the ability to share the following devices and features located on the host with the guest:

- **Audio** This setting controls whether audio input and output for the virtual machine is redirected to audio devices in the host or is managed by an emulated audio device. To improve audio performance, clear the check box for a virtual machine running Windows XP, and select the check box for a virtual machine running Windows Vista or Windows 7. Note that for Vista and later guest operating systems, a virtual audio device is used for redirection of sound from the guest to the host. For earlier versions of Windows, a SoundBlaster 16 audio adapter is emulated.
- **Clipboard** You can copy and paste data and files between the host operating system and the guest operating system. For example, you can copy a URL from the browser in a guest operating system and paste it to a browser in the host operating system.
- **Printer** This IC gives the user the ability to redirect printing from the guest machine to printers on the host machine. To share printers when Windows XP is the guest operating system, you must also install the printer drivers. In enhanced mode, the drivers need to be present on the host and on the guest. However, basic mode does not require host side drivers to be present.
- **Smart Cards** This IC provides the ability to support smart cards connected to the host inside the guest operating system. The smart-card redirection works in a similar way as printer redirection.
- **Drives** This feature shares the drives you select on the host with the virtual machine so that you can easily access host data from within the virtual machine. This feature also makes it possible to access the host desktop and Documents folder from virtual applications when you select those resources to share. Note that host hard drives are listed in the guest operating system by using the computer name of the host operating system. For example, on a host computer named WindowsTest, the C drive would be listed in the guest operating system as “C on WindowsTest”.

—CSS Global Technical Readiness (GTR) Team

Keyboard

You can use the Keyboard settings page to configure whether Windows key combinations such as Alt+Tab are sent

- In Full Screen View Only (the default)
- To Virtual Machine
- To This Computer

Windows Virtual PC includes a number of keyboard shortcuts. Different keyboard shortcuts are available for Windowed mode and for Running mode. Interaction with a virtual machine window is in Windowed mode when the focus of the mouse and keyboard are on the border, title bar, or menu bar of the VM window, and it is in Running mode when the focus of the mouse and keyboard are on the desktop of the guest operating system. To switch from Windowed mode to Running mode, click anywhere on the desktop of the guest. To switch from Running mode to Windowed mode, press the “release” key combination of Ctrl+Alt+Left Arrow. Table 3-1 lists the keyboard shortcuts available in Windowed mode, while Table 3-2 lists those available in Running mode.

TABLE 3-1 Keyboard Shortcuts Available in Windowed Mode

Operation	Keyboard Shortcut
Ctrl+Alt+Del	Ctrl+Alt+End
Alt+Tab	Alt+Page Up
Alt+Shift+Tab	Alt+Page Down
Windows key/Ctrl+Esc	Alt+Home
Print Scrn	Ctrl+Alt+Plus Sign (+)
Alt+Print Scrn	Ctrl+Alt+Minus (–)
Release keyboard and mouse	Ctrl+Alt+Left Arrow

TABLE 3-2 Keyboard Shortcuts Available in Running Mode

Operation	Keyboard Shortcut
Pause	Alt+P
Resume/Wakeup	Alt+U
Settings	Alt+E
USB Menu	Alt+B
Full screen toggle	F11/Ctrl+Alt+Break
Help	F1

Logon Credentials

You can use the Logon Credentials settings page to delete the saved credentials of the user account used to log on to the virtual machine. For Windows XP Mode, this user account is VirtualPC-#####\XPMUser, where ##### is a numeric string that is randomly assigned to the virtual machine when Windows XP Mode is installed. When you create a new virtual machine, you have the option of storing the user name and password so that you can be logged on automatically to the virtual machine when it is booted.

Auto Publish

By selecting the Automatically Publish Virtual Applications check box on the Auto Publish settings page, applications you have installed on the guest operating system can be launched using the Start menu of the host operating system. There are additional requirements before autopublishing will work, however; see the section titled “Enabling Support for Virtual Applications” earlier in this chapter for more information.

Close

You can use the Close settings page to configure the action that Windows Virtual PC performs when you click the Close button at the top right of the virtual machine window. You can either configure Windows Virtual PC to prompt you to choose an action to perform or configure it to perform one of the following actions: Hibernate, Shut Down, Turn Off, or Turn Off And Discard Options. (The last of these options is available only when the Undo Disks feature is enabled for the virtual machine.)

Using Windows XP Mode

Although you can use Windows Virtual PC to create your own custom virtual machines, the most likely scenario in business environments is to use Windows XP Mode to mitigate problems running older applications. The following sections describe some of the things you can do using Windows XP Mode.

Working with Windows XP Mode

You can open the Windows XP Mode virtual machine in the Windows Virtual PC console window to interact with the guest operating system installed on the virtual machine. You might do this, for example, to configure the guest operating system by installing software updates or attaching USB devices. You would also do this to install applications on the guest.

As shown in Figure 3-12, the Windows Virtual PC console window has four menu items you can choose from: Action, USB, Tools, and Ctrl+Alt+Del. The Action menu has four options:

- **View Full Screen** Displays the guest desktop in full-screen mode with a Remote Desktop Services bar at the top
- **Sleep** Puts the guest into sleep mode without closing the window
- **Restart** Reboots the guest operating system
- **Close** Hibernates the guest operating system




FIGURE 3-12 The Windows Virtual PC console window showing Action menu items.

The USB menu displays USB devices attached to the host computer and allows you to attach them to the virtual machine using USB redirection. To do this, click the USB menu and then click Attach for the USB device. (See Figure 3-13.)




FIGURE 3-13 A USB device on the host is available for redirection to the virtual machine.

The host computer then installs a stub driver that enables redirection of the USB device to the virtual machine, while the guest operating system uses Plug and Play to install drivers for the device. (See Figure 3-14.)




FIGURE 3-14 Drivers are installed on the guest and the host.

After the USB device has been attached (redirected) to the virtual machine, it can be configured and used using the appropriate programs. (See Figure 3-15.)




FIGURE 3-15 The USB device is available to the guest operating system.

While attached to the guest, the USB device is exclusively available to the guest and cannot be used from the host unless you use the USB menu again to release it from the guest. (See Figure 3-16.)




FIGURE 3-16 You must release a redirected USB device before you can use it again on the host.

Some USB devices—including printers, storage devices, and smart card readers—can be shared by both the guest and the host. When you connect such devices to the host computer, they are displayed in the Action menu as Shared (as shown in Figure 3-17). Other USB devices cannot be shared and must be redirected instead, which enables them to be used exclusively by the guest until they are released. Examples of USB devices that must be redirected include digital cameras, MP3 players, cell phones, PDAs, CD/DVD readers/writers, wireless network adapters, and webcams.




FIGURE 3-17 Shared USB devices can be accessed simultaneously from the guest and the host.

The Tools menu lets you choose from the following actions to perform:

- **Disable Integration Features** Selecting this option disables all Integration Features except basic features and locks the desktop on the guest. Select the option again to re-enable Integration Features on the guest.
- **Settings** Selecting this option displays the Windows Virtual PC Settings dialog for the virtual machine.

Finally, the Ctrl+Alt+Del menu lets you display the Windows Security dialog on the guest.

Installing and Using a Virtual Application

This overview of Windows Virtual PC and Windows XP Mode concludes with examining how to install and use a virtual application on Windows XP Mode. To illustrate how to do this, we will install Microsoft Office XP Professional on a Windows XP Mode virtual machine and show how to access Microsoft Word XP from the Start menu of the host.

Begin by configuring the DVD Drive settings page of the virtual machine settings to access the physical drive on the host, and then insert the Office XP DVD media in the DVD drive on the host computer. Alternatively, you can configure the DVD Drive settings page to open an .iso image, copy the .iso image for Windows XP to the host, and double-click on the DVD

drive icon in My Computer. Either method will launch the Setup program for Office XP on the host. (See Figure 3-18.)




FIGURE 3-18 Installing Office XP on the Windows XP Mode guest.

After Office XP has been installed on the guest, you can launch Word XP from the Start menu of the guest, which opens a document window within the virtual machine console window. (See Figure 3-19.)




FIGURE 3-19 Running Word XP within the Windows XP Mode console window.

Close the console window, which puts the virtual machine into hibernation. Now, on the host computer, click Start, All Programs, Windows Virtual PC, and then Windows XP Mode Applications. This displays a list of published applications installed on the guest (as shown in Figure 3-20).




FIGURE 3-20 The Office XP suite of applications have been published on this virtual machine.

Now select Microsoft Word (Windows XP Mode) from the list of virtual applications available on the host computer. A progress bar is displayed indicating that Windows Virtual PC is trying to launch the application, as shown in Figure 3-21.




FIGURE 3-21 The virtual machine is being started in order to run the virtual application.

A dialog box is then displayed indicating that the XPMUser account is still logged on to the Windows XP Mode console window and must be logged off before the virtual application can be launched. (See Figure 3-22.)



FIGURE 3-22 You must log off from the Windows XP Mode console window before the virtual application can start.

At this point, the progress bar indicates that the virtual application is being started. (See Figure 3-23.)




FIGURE 3-23 The virtual application is being started.

Moments later, a Word XP document window is displayed on the host computer. (See Figure 3-24.)




FIGURE 3-24 The published application is now running on the host.

When you select Save As from the File menu in this application, the open document is saved by default to the user's Documents folder on the host computer, not the guest. (See Figure 3-25.) In other words, you can work with this virtual instance of Word XP as if you had the program installed locally on your host computer instead of on the guest.




FIGURE 3-25 The default save location for a virtual application is the Documents folder on the host.

Finally, if you close the virtual application and then re-open it from the Start menu on the host computer, this time the progress dialogs are not displayed. That's because closing a running virtual application leaves the virtual machine running in the background for a period of time (5 minutes), during which the virtual application can be immediately restarted without having to restart the virtual machine.



More Info For more information about Windows Virtual PC and Windows XP Mode, see the resources listed in the "Additional Resources" section at the end of this chapter.

Understanding MED-V

Microsoft Enterprise Desktop Virtualization (MED-V) is a desktop virtualization technology that builds upon the highly popular and easy-to-use Microsoft Virtual PC 2007, a first-generation desktop virtualization product. The coverage of MED-V in this chapter is limited to the current version, MED-V 1.0, which is available to Software Assurance customers as part of MDOP 2009 R2 and currently works on desktop computers running Windows Vista. The next version, MED-V 1.0 SP1, which is scheduled for release in the first quarter of 2010, will include the following new features:

- Support for 32-bit and 64-bit hosts running Windows 7.
- Support for virtual machines (guests) running Windows XP SP3 (recommended), Windows XP SP2, and Windows 2000 SP4. (No support is planned for Windows 7 or Windows Vista guests in the immediate future.)

Note that MED-V 1.0 SP1 will still use Microsoft Virtual PC 2007 as its foundation, not Windows Virtual PC. The reason for this is because Windows Virtual PC requires that the computer support hardware virtualization while Virtual PC 2007 does not require such support.

Introducing Microsoft Enterprise Desktop Virtualization

Microsoft Enterprise Desktop Virtualization enhances deployment, management, and user experience for Virtual PC images to streamline operating system upgrades and to increase IT control and user flexibility in enterprise environments. With MED-V, application-to-operating system compatibility issues are minimized; operating system migrations are accelerated; and delivery and reconstitution of corporate desktops are made easy, simplifying support tasks, business continuity, and incorporation of heterogeneous IT environments.

Based on the technology acquired from Kidaro, MED-V adds four additional features on top of Virtual PC 2007. These mechanisms are designed to facilitate the creation, storage, delivery, and management of Virtual PC images to desktop computers. At a high level, the four additional technologies provided by MED-V are as follows:

- **Virtual images repository and delivery** Simplifies the process of creating, testing, delivering, and updating virtual images
- **Centralized management and monitoring** Manages the life cycle of virtual images, provisions virtual images to authenticated users according to Microsoft Active Directory users and groups, and aggregates client events for monitoring and reporting purposes.

- **User policy and data transfer control** An endpoint agent enforces usage policies and data transfer permissions on the virtual machine.
- **Seamless end-user experience** The user remains unaware of the virtualization running in the background and keeps one desktop environment.

Virtual Image Repository and Delivery

MED-V provides mechanisms for storing and delivering standard Virtual PC images (also called virtual images) onto user's desktop computers. These mechanisms simplify the process of creating, testing, deploying, and maintaining virtual images from a central location.

The virtual image repository and delivery capabilities provided by MED-V include

- A central repository for storing, versioning, and delivering virtual images you create
- An administrator console for virtual image creation and testing
- A format for packaging the MED-V client and virtual images for automatic deployment over the network, over the Web, or via removable media such as DVD media or USB key drives
- A client component that uses standard MSI installation and allows users to retrieve virtual images using a standard Web infrastructure, together with an automated process for keeping users' computers updated with the most recent image build without interrupting their work
- An efficient, bandwidth-conserving Trim Transfer mechanism for delivering virtual images over both high-speed local area network (LAN) and slow wide area network (WAN) connections
- An auto-installation package that allows self-deployment of the client component and the virtual images using removable media such as DVD or from a Web site
- Support for image delivery using standard enterprise content distribution systems

For a closer look at some of these technologies, see the section titled "How MED-V Works" later in this chapter.

Centralized Management and Monitoring

MED-V provides the means for managing the entire life cycle of virtual machines deployed on desktop computers throughout an enterprise. The centralized management and monitoring capabilities provided by MED-V include

- A central management server that can be used to control virtual machines that have been deployed onto desktop computers

- Integration with Microsoft Active Directory Domain Services to enable provisioning of virtual images based on group membership or user identity
- The requirement that users must authenticate using a valid account before being granted access to the virtual desktop, regardless of whether the virtual machine is online or offline
- A mechanism for automating the first-time setup of virtual machines—for example, by specifying a unique computer name, performing initial network setup, joining a domain, and performing other needed deployment steps
- The ability to remotely assign the amount of RAM allocated and the network settings used by Virtual PC on endpoint computers, which facilitates deployment of virtual images across diverse computers in heterogeneous environments
- A central database of all client activity and events, making it easy for helpdesk personnel to remotely monitor for problem conditions and to facilitate troubleshooting
- The ability to revert a virtual machine back to its base image, making it easy for helpdesk personnel to support the virtual desktops

Usage Policy and Data Transfer Control

MED-V includes an agent on the endpoint (client) that can be used to enforce the application of corporate usage policies and permissions to virtual machines for specified users, groups, or both. The usage policy and data-transfer control capabilities provided by MED-V include

- The ability to protect a virtual image from unauthorized execution
- The ability to configure an expiration date for a virtual machine and specify a time limit for offline use of the virtual machine.
- The ability to allow or block inbound and outbound data transfer between the virtual machine and the endpoint, regardless of whether data transfer is performed using copy/paste, file transfer, or printing.
- The ability to automatically redirect specified Web sites (such as the corporate intranet or sites that require an older version of the browser) from the endpoint browser to the virtual machine so that they run automatically when the browser installed on the virtual machine is started.

Seamless End-User Experience

MED-V provides a seamless experience for end-users, making users unaware of the virtual machines running in the background. Overall, it reduces the training required for deploying Virtual PC images by making the deployment process transparent to the end user and

simplifying the work process with virtual machines. The end-user experience provided to users by MED-V includes

- The user does not have to learn how virtualization works or how to work with a separate desktop as is required when running Microsoft Virtual PC 2007 without MED-V. Instead, the virtual machine is “invisible” to the user and the user interacts with it only through an icon in the system notification area of the taskbar on the user’s computer.
- Applications that are published to endpoints using MED-V and are installed in Virtual PC are made available to the user via the Start menu, desktop shortcuts, or both on the user’s computer. These applications run within Virtual PC and are seamlessly integrated into the desktop of the user’s computer in a way that makes the applications work as if they were locally installed on the user’s computer.
- Advanced users and administrators have access to a power user mode that lets them view the virtual machine loading processes and can even display the desktop of the virtual machine if needed for troubleshooting purposes.
- The user has access to any applications published to him via MED-V even when the user’s computer is disconnected from the corporate network (offline) or when network connectivity is limited.

How MED-V Works

The following description of how MED-V works is based on a prerelease version of the product and is subject to change. The sections below cover the following topics:

- MED-V terminology
- System requirements
- Supported applications
- Deploying packages
- Workspace initialization
- Using the workspace
- Configuring usage policies
- Managing and maintaining virtual images
- Domain managed vs. self-cleaning virtual machines
- Single desktop vs. full desktop modes

MED-V Terminology

The following terminology is used for describing different components of a MED-V environment:

- **Guest** The operating system installed on the virtual machine.
- **Host** The end user's physical computer, typically a desktop or laptop computer. It's also called an endpoint computer.
- **MED-V Client** Software that runs on the host that can download and run Virtual PC images seamlessly on the host, according to MED-V usage policies.
- **MED-V Image Repository** An IIS Web server that stores and distributes virtual images to endpoints.
- **MED-V Management Server** A MED-V server that authenticates, provisions, and controls all users of the system. Client-server communication is based on HTTP or HTTPS.
- **MED-V Server** The server that holds the main image repository and is the management server.
- **MED-V Package** A mechanism for installing Virtual PC, the MED-V client, and optionally a virtual image on a host.
- **Virtual image** A file representing the file system of a virtual machine. This file can be delivered to multiple endpoints, independent of their hardware or software.
- **Workspace** The Virtual PC image that the MED-V client runs on the host. It's also called a virtual machine or guest.

System Requirements

The system requirements for version 1.0 of MED-V can be broken down into the following categories:

- System Requirements for MED-V Clients
- System Requirements for the MED-V Server
- Active Directory Domain Services Requirements
- Database Server Requirements

System Requirements for MED-V Clients The system requirements on the client side (the user's host computer) for MED-V are shown in Table 3-3.

TABLE 3-3 System Requirements for MED-V 1.0 on the Client Side

Memory	Minimum: 1 GB Recommended: 2 GB
Operating system	Windows XP SP2/3 (Professional, Home), Vista SP1 (Enterprise, Home Basic, Home Premium, Business, Ultimate) 32-bit
Web browser	Microsoft Internet Explorer 6 SP2, 7.0
File system	NTFS
Supported locales	English, French, German, Italian, Portuguese (Brazil), Spanish

Note The system requirements for MED-V 1.0 SP1 are not available at the time of this writing. However, it will support both 32-bit and 64-bit versions of Windows 7 on the client side.

System Requirements for the MED-V Server The system requirements for deploying the MED-V server are shown in Table 3-4.

TABLE 3-4 System Requirements for MED-V Server Version 1.0

Memory	4 GB RAM or greater
Processor	Dual Processor (2.8 GHz)
Operating system	Windows Server 2008 Standard/Enterprise x86 & 64-bit
Database	Microsoft SQL Server 2005 Enterprise SP2 Microsoft SQL Server 2008 Express/Standard/Enterprise editions
Supported locales	English, French, German, Italian, Portuguese (Brazil), Spanish

Active Directory Domain Services Requirements The MED-V server can be installed either on a member server belonging to an Active Directory Domain Services domain or on a standalone server when using a workgroup scenario. If a domain environment is being used, authentication and authorization of the user is performed against Active Directory Domain Services. (If users are not part of the same domain the server belongs to, a trust must be created between the domains.) If a workgroup environment is being used, authentication and authorization is performed using user and group accounts stored locally on the server.

Database Server Requirements Microsoft SQL Server is needed for hosting the MED-V report database where MED-V workspace logs are stored. This log database is then used for generating MED-V reports. SQL Server can be installed either on the same server as the MED-V server or on a remote server. As Table 3-4 indicates, the following versions of SQL Server are supported for hosting the MED-V database:

- Microsoft SQL Server 2005 Enterprise SP2
- Microsoft SQL Server 2008 Express/Standard/Enterprise editions

Supported Applications

Because MED-V does not add any additional layer of virtualization functionality on top of Virtual PC 2007, any application that has vendor support for running on Virtual PC 2007 will typically run in MED-V without encountering any issues. There are no special restrictions on the type of applications that can run in MED-V workspaces.

Deploying Packages

MED-V packages are used to deploy full Virtual PC images to host computers as MED-V workspaces. These virtual images include a guest operating system, applications, and user data. Guest operating systems can be either domain members or standalone systems, and if the host computer is a domain member, the virtual machine can even belong to a different domain if needed. MED-V workspaces are fully functional virtual desktop computers and can be associated with any Active Directory Domain Services environment. For example, a workspace that is a domain member can be locked down using Group Policy.

MED-V packages are typically deployed onto host computers that already have Virtual PC installed on them. If a host computer does not have Virtual PC installed, a MED-V package can be constructed that will deploy Virtual PC together with the Virtual PC image being deployed to the host.



Note Local Administrator privileges on the host system are required in order to install a MED-V package. Alternatively, MED-V and Virtual PC can be installed using a software distribution mechanism such as Microsoft System Center Configuration Manager (SCCM).

When a MED-V package is installed on a host system and the workspace is active, a prespecified amount of RAM on the host system is allocated for running the Virtual PC image. In other words, when you create a MED-V package to deploy a virtual image to a host system, you can specify how much RAM the host system will use to run the virtual image using Virtual PC.

MED-V packages can be deployed over any network connection, including high-speed LAN connections, slow WAN links, or virtual private network (VPN) connections. You can also deploy virtual images to users on removable media, such as USB key drives or CD/DVD media, when network connectivity is not available for the user. Packages that are deployed via removable media can be encrypted so that only authorized users will be able to install and use the virtual images.



Note To enforce MED-V usage policies, Virtual PC images that have been deployed to host computers via MED-V packages will not run natively in Virtual PC on systems where the MED-V client has not been installed. In other words, users cannot bypass MED-V by running MED-V–deployed virtual images directly in Virtual PC. The reason this is not possible is because MED-V creates encrypted virtual images that can be accessed only by the MED-V client.

Workspace Initialization

After a MED-V package has been deployed to a host computer, the host has the following software installed:

- MED-V client software
- Virtual PC 2007
- A virtual machine image

The next step is to initialize the MED-V workspace. Figure 3-26 illustrates the steps involved in the process of initializing the workspace on a host computer to which a MED-V package has been deployed.




FIGURE 3-26 Connection sequence for initializing the workspace on a host.

The steps involved in this initialization sequence are as follows:

1. When the MED-V client is installed and started on the host, the client tries to establish a communications channel with the MED-V server using a preconfigured port.
2. After this communications channel has been established, the user's credentials are authenticated against the MED-V server.
3. If user authentication is successful, the MED-V server provides the host with an encryption key that can be used to decrypt the virtual machine (VM) image that the MED-V server has previously deployed to the host. After the VM has been decrypted, the MED-V client running on the host uses Virtual PC to launch the VM, which initializes the MED-V workspace on the host. The user can then access either of the following:
 - ❑ The full virtual desktop of the active virtual machine. In this case, the virtual desktop is displayed as a separate window on the user's local desktop just like when using Virtual PC natively.
 - ❑ Individual applications that the administrator has published from the virtual machine. Any application window that appears on the virtual desktop is fully

integrated into the host desktop; this includes context menus, system tray icons, and system tray notifications. Applications from the system tray of the workspace are displayed as “shadow icons” in the system tray on the host. The way this works is that the virtual machine starts in the background when the MED-V client is launched, and the MED-V client can display applications running in the virtual machine as windows on the local desktop of the host. If the user closes all applications running in his workspace, the state of the virtual machine is saved. To restart the virtual machine, the user clicks the MED-V client icon in the system tray.



Note Any number of Virtual PC images can be installed on a host system. However, only one Virtual PC image can be running on the host at any given time. When starting the MED-V client, the user will be asked which workspace he would like to run if more than one is configured for him.

Using the Workspace

The MED-V workspace and the user’s local desktop are two separate and distinct entities, and by default there is no direct interaction between them. For example, if the user launches a local copy of Microsoft Office Word, creates a document, and tries to save the document, the document is saved by default in the Documents (in Windows Vista and later) or My Documents (in Windows XP) folder on her local computer. If, however, the user launches a copy of Word from within the workspace, creates a second document, and tries to save the document, this second document is saved by default in the Documents (or My Documents) folder in the user’s workspace—that is, in the virtual machine running in Virtual PC on the user’s computer. This means that by default the user will have two Documents (or My Documents) folders to deal with—the one on the local computer and the one in the user’s workspace.

Although this situation can be confusing to some users, administrators can give users permissions to use the MED-V file transfer tool (accessible from the MED-V tray menu) to copy files or directories between the host and the guest. Alternatively, they can work around this issue by implementing Folder Redirection on the host (the user’s physical computer), the workspace (the user’s virtual computer), or both computers. For example, an administrator could use Group Policy to redirect both Documents (or My Documents) folders—the physical one and the virtual one—to the same shared folder on a network file server.

In fact, because a virtual machine running Windows looks and behaves just like any physical computer running Windows, the full power of Group Policy can be used within an Active Directory Domain Services environment to manage user data and application settings within the workspace. For example, in addition to configuring Folder Redirection for MED-V virtual machines, you can also implement Roaming User Profiles and Offline Files if these are needed. For more information on using these different corporate roaming technologies, see Chapter 15 of the *Windows 7 Resource Kit* from Microsoft Press.

Virtually anything a user can do on his physical host computer, he can also do in his MED-V workspace, provided the guest operating system and applications in the Virtual PC image are installed and configured appropriately. For example, if the virtual image includes third-party VPN client software, the user will be able to launch this software from within his workspace and use it to connect to a remote network to which he is authorized to connect.

It's important to realize, however, that MED-V provides users with two separate computers—their physical host computer and a virtual workspace—and that these two computers are completely separate and do not normally interact with one another. For example, the user cannot use Microsoft Outlook running in the workspace to open a PDF file attachment using Adobe Acrobat Reader running on the local desktop. Instead, the user needs a copy of Adobe Acrobat Reader installed in the virtual machine image. In other words, local applications and virtual applications cannot interact with one another unless they can do so as applications running on separate computers residing on the same network. Simple data exchange between the host and guest is possible, however, using a shared clipboard. The process for doing this is explained in the next section.

Configuring Usage Policies

MED-V usage policies allow administrators to configure the following kinds of behaviors:

- Set expiration dates for the virtual machine and time limits for offline work.
- Control inbound/outbound data transfer control (for example, copy/paste, file transfer, printing) between the virtual machine and the endpoint.
- Enable automatic redirection of predefined Web sites (for example, corporate intranet) from the endpoint browser to the virtual machine.
- Configure published applications. (Applications that are installed in the virtual machine become available through the user's Start menu.)

The MED-V server is used to assign usage policies and data transfer control permissions, which are then enforced by an endpoint agent. The following are additional details concerning two of the policies you can configure.

Configuring Clipboard Behavior MED-V usage policies allow administrators to configure the behavior of the clipboard for copying/pasting between the workspace and the host computer. Options for configuring the sharing of clipboard data between the host and the workspace include

- Disabling shared clipboard functionality entirely
- Enabling one-way copy/paste functionality—for example, from the workspace to the host or from the host to the workspace
- Enabling two-way copy/paste functionality

Configuring Virtual Images to Expire MED-V allows administrators to configure Virtual PC images so that they will expire under certain conditions, including

- If a new connection to the MED-V server is not made within a specified interval of time
- At a date and time you specify

The way MED-V configures the expiration of virtual images is similar to the way App-V configures the expiration of applications. This is a security feature that ensures virtual desktops and applications will not be used outside of the purposes for which they are intended. This capability of configuring the expiration of virtual images is especially important because MED-V enables users to continue to use their virtual desktop and applications even when their computer becomes disconnected from the network—for example, by undocking a laptop computer from its docking station to take the computer home or on the road.

Creating, Managing, and Maintaining Virtual Images

Management of a MED-V environment is performed by using the MED-V Management console. Using this console, administrators can manage the inventory and versioning of Virtual PC images, create packages, and perform other related management tasks. In future versions of MED-V, such management functionality might be integrated into the Microsoft System Center family of products.

Virtual images can be retrieved by the MED-V client from the MED-V server using Microsoft Background Intelligent Transfer Service (BITS) version 2, which streams data over HyperText Transfer Protocol (HTTP) or Secure HTTP (HTTPS) sessions. The first time a virtual image is delivered to the host, the entire image must be streamed before the virtual image can start working. After this is done, however, if the image is later updated on the MED-V server, only the deltas between the old and new images must be streamed. This is possible because the MED-V client employs advanced de-duplication techniques to detect blocks in the virtual machine image that already exist on the host, and then it downloads only the missing blocks that might be needed. The result of this implementation is to significantly reduce the time it takes to launch applications that have been updated in virtual images stored in the repository on the server. This implementation also enables MED-V to work well in low-bandwidth scenarios—for example, over a slow WAN link or modem connection—because it significantly reduces the amount of data that needs to be downloaded when an image is distributed or updated over a low-bandwidth connection.

Administrators need to maintain Virtual PC images stored in the repository on the MED-V server. For standardized virtual images that do not contain any user data, administrators can choose to create a master virtual image and then create a new version of this image in the repository. They can then update the new image by making configuration changes, installing or upgrading applications, applying service packs or software updates, and performing similar maintenance tasks on the image. Then the next time the client tries to run an application

using the deployed image, the MED-V client on the user's computer will download the deltas from the server, thus updating the local copy of the image on the user's computer.

For Virtual PC images that have already been deployed and that contain user data—for example, data stored within user profiles in the guest operating system—administrators can choose to update the deployed virtual image using traditional software maintenance methods. For example, they can deploy or upgrade applications on the guest by using Group Policy Software Installation. And they can ensure the guest is fully up to date with the latest critical security updates by using Group Policy to configure the guest to download software updates from Windows Update or from a server running Windows Server Update Services (WSUS). In this scenario, the virtual machine is being maintained in the same way that a physical system would be maintained using standard tools.

The following list of steps provide a high-level overview of the tasks involved in creating, deploying, and maintaining virtual images using MED-V:

1. Create the virtual image within Microsoft Virtual PC.
2. Define a MED-V workspace by doing the following:
 - ❑ Create a list of applications installed in the virtual image that are to be made available to end users through their Start menu.
 - ❑ Specify Web sites that should be viewed inside or outside the virtual machine browser and which should be redirected to the appropriate location by the MED-V client.
 - ❑ Provision the MED-V workspace to users and groups in Active Directory.
 - ❑ Configure a usage policy (such as expiration, permission to work offline, and so on) and data transfer permissions (such as file transfer, copy and paste, and printing) for these users and groups.
3. Test the virtual image using the MED-V management console, and load it into the MED-V Image Repository.
4. Deploy the MED-V client to users' computers by using one of these methods:
 - ❑ Use an enterprise software distribution tool to deploy the MED-V client and Virtual PC software as standard Windows Installer files.
 - ❑ Create a MED-V self-install package that includes the MED-V client and Virtual PC software, and deliver the package via a self-service Web site or on removable media such as a CD or DVD.
5. Deliver the virtual image to the users' computers by using one of these methods:
 - ❑ Over the network using HTTP or HTTPS
 - ❑ Using an enterprise software distribution tool
 - ❑ Using removable media such as a CD or DVD.

At this point, users can authenticate against the MED-V management server and are ready to work within the virtual image deployed to their computers. After the first online authentication, offline work is also supported if this has been enabled by the administrator.

Ongoing maintenance involves managing and updating the MED-V workspace. The MED-V management console lets administrators update usage policies, provision MED-V workspaces to additional users, deprovision existing users, and update virtual images. Updates are automatically distributed to appropriate users when they work online. The MED-V management console also lets administrators monitor clients by generating a report that contains detailed information concerning client events. This report can help the administrator understand the source of the problem remotely when an error occurs and instruct the user on how to resolve the problem.

After a virtual image has been deployed, it must be managed in the same way that any desktop operating system must be managed within a corporate IT environment. This means delivering new applications to the virtual machine, patching the guest operating system of the virtual machine, updating security definitions and policies, and so on. To manage deployed virtual machines, simply join them to an Active Directory Domain Services domain and manage them just like physical desktop computers on your network using standard administration tools and services to patch, update, deliver applications, and apply policies on the virtual machine.

Single Desktop vs. Full Desktop Modes

The MED-V user experience can take two forms:

- **Single desktop mode** In this approach, which is recommended, the administrator publishes the applications installed on the virtual machine to make these applications available on the Start menu on the users' computers. Users then launch these virtual applications from their Start menu or by double-clicking on desktop shortcuts, and the running virtual applications are displayed side by side, together with native applications that are locally installed on the user's computer. (The virtual applications can optionally be differentiated by enclosing them in a colored frame.) In this approach, the virtual machine desktop is not visible to the user, which simplifies the user experience and makes it easier for the user to use the virtual applications to perform his work.
- **Full desktop mode** In this approach, the user sees and works with the whole virtual machine, either in a console window or full screen. The user has to manually toggle between his physical desktop and the virtual machine desktop to switch between using native applications and virtual applications.

Figure 3-27 shows the Start menu on a user's computer and displays a list of virtual applications that are delivered to the user's computer using MED-V.




FIGURE 3-27 Launching a MED-V application from the Start menu.

Figure 3-28 shows several virtual applications running on the user's computer in single desktop mode. Note that these applications look just like they are locally installed on the user's computer.

Figure 3-29 shows the MED-V system tray icon and the menu that can be displayed by right-clicking on this icon.

These screen shots show that the MED-V user experience is similar to the Windows XP Mode user experience described earlier in this chapter. The difference between these technologies is that Windows XP Mode cannot be centrally managed the way MED-V does it, which makes MED-V suitable for large environments with hundreds or thousands of computers to deliver virtual applications to.




FIGURE 3-28 MED-V applications running on a host computer's desktop.




FIGURE 3-29 Using the MED-V system tray icon.

Finally, Figure 3-30 shows the MED-V Management Console with the Policy section selected and the focus on the Applications tab, which lists the published applications and menus on the clients. You can verify that Microsoft Office 2003 applications have been published to the Start menu on the clients by referring back to Figure 3-27.




FIGURE 3-30 The MED-V Management Console showing which applications and menus have been published to the clients.



More Info For more information about MED-V, see the MED-V resources listed in the “Additional Resources” section at the end of this chapter.

Understanding App-V

Another important component of Microsoft’s integrated virtualization strategy is application virtualization, which refers to any technology that lets you decouple applications from desktop operating systems to dynamically deliver applications on demand to your users. When you run applications centrally instead of installing them on each user’s computer, software update management is simplified, application-to-application conflicts are reduced, and application compatibility regression testing is made easier.

Microsoft’s primary platform for delivering application virtualization solutions is Microsoft Application Virtualization (App-V) 4.5, formerly known as SoftGrid Application Virtualization.

App-V improves upon the earlier SoftGrid Application Virtualization platform in four broad areas:

- **Dynamic Suite Composition** App-V includes Dynamic Suite Composition (DSC), which provides a method for administrators to control which virtual applications will be combined to create a unified, virtual working environment for an application set. DSC also provides a way for administrators to specify mandatory or optional dependencies between virtual applications. This means that when a virtual application is run on the client, it will also launch the dependent virtual application's environment, allowing for the combination of both virtual environments.

DSC also enables a one-to-many scenario for middleware applications. For example, let's say that you have some applications that require the Java Runtime Environment (JRE). You first sequence the JRE into its own virtual application. Then you reset the Sequencer, install the JRE locally, and sequence the dependent application; then you edit the OSD file to manually create a dependency between the single virtual JRE package and the different virtual dependent applications that enables multiple virtual applications to share the same virtual JRE package. DSC also reduces the sequencing overhead because only one JRE needs to be sequenced instead of having to resequence the JRE into each individual package. Package updates are also simplified because you need to update only the single JRE package instead of multiple packages.

- **Enhanced scalability** App-V provides numerous scalability improvements. For example, App-V has flexible deployment modes and is interoperable with Microsoft Systems Management Server (SMS), Microsoft System Center Configuration Manager (SCCM), and third-party ESD systems. App-V also supports a standalone mode for virtual application delivery and has increased supportability over the earlier SoftGrid version. These scalability enhancements can benefit enterprises of all sizes, especially those with branch offices and those that have existing ESD systems in place.

Other scalability enhancements in version 4.5 of App-V include

- **Background streaming with autoload options** The client stream policy can be set so that the entire virtual application will be delivered on first launch or on login. The application can also be used while it is still streaming in the background.
- **Offline availability** The virtual application can be accessed when offline as well as online.
- **More applications** The DSC capabilities described earlier, together with the enhanced support for side-by-side applications, expands the number of applications that can be serviced by the platform.

- ❑ **Application Source Root** The Application Source Root (ASR) provides the capability to specify a server name through Group Policy or a script on the target client. When the manifest or domain controller refresh occurs, the .osd file in its current state is delivered to the client. However, at runtime, the ASR value replaces the DNS server name in the RSTP URL of the .osd file. This will route the stream request to the ASR server name. This server is typically a server that is local to the client.
- ❑ **Windows Server 2008 Terminal Services support** Application Virtualization for Terminal Services supports Microsoft Windows Server 2008 Terminal Services (32-bit only).
- ❑ **Enhanced data metering** The App-V Windows Management Instrumentation (WMI) provider collects application usage information and provides a simplified way to pull the data into your organization's reporting store. Data metering can even occur while users are offline.
- ❑ **Microsoft Update** App-V includes Microsoft Update support for virtualized applications at sequencing time (but this is not available at run time).
- ❑ **Backup and recovery** App-V includes Volume Shadow Copy Service (VSS) Writer support.
- ❑ **Enhanced management** You can manage your App-V system using tools such as the System Center Operations Manager 2007 Management Pack, ADM template, and Best Practice Analyzer.
- ❑ **Improved diagnostics** App-V has Watson integration and event log support on both the client and the server.
- ❑ **MSI creation capability** App-V allows for Microsoft Installer (MSI) creation for standalone use. App-V also supports streaming MSIs, where the MSI sets up the virtual application and settings but the application is streamed from the server when the user clicks on the application.
- ❑ **Enhanced command-line interface** Additional capabilities for batch operations are provided by the SFTMIME command. Batch MSI creation is also supported.
- ❑ **Differential SFTs** App-V allows for the creation of package content files (SFTs) with only sequenced differences (updates) for use with Standalone mode only (merged by client action).
- **Globalization** The globalization and localization features of App-V support localized applications and operating systems. Specific enhancements include
 - ❑ Installation on any Windows language version supported by the earlier SoftGrid version.

- ❑ Installation in mixed-language environments (server/client). For example, a Japanese employee traveling to her German branch office could use the Japanese App-V Client with the German App-V Server.
- ❑ Autodetecting the system and user locale, and autoloading the appropriate resource files.
- ❑ Respecting all user locale and regional settings.
- ❑ Sequencing non-English or localized applications.
- ❑ Support for foreign language applications with special characters.
- ❑ Foreign language Active Directory and server support.
- ❑ Run-time locale detection.
- ❑ Localization in 12 Languages: Portuguese (Brazil), Chinese (Simplified), Chinese (Traditional), Dutch (client only), French, German, Italian, Japanese, Korean, Russian, and Spanish.
- **Enhanced security** The enhanced security features in this version of App-V include
 - ❑ Support for Internet-facing scenarios where users run virtual applications over the Internet without needing to use a VPN connection. App-V can securely provide virtual applications to users when both the App-V Server and the App-V Client are on untrusted networks.
 - ❑ A secure-by-default configuration out of the box that includes locked-down client privileges, Transport Layer Security (TLS) turned on by default, Kerberos support, and certificate-based server authentication.

App-V provides the ability to deliver applications to end users without actually installing the applications on their client computers. This section is designed to help you understand App-V and covers the following topics:

- App-V terminology
- How App-V works
- App-V components
- App-V architecture

App-V 4.5 SP1 and App-V 4.6

This sidebar provides information on the latest and soon-to-be-released updates to App-V.

App-V 4.5 SP1

At the time of this writing, App-V 4.5 SP1 has been released as part of MDOP 2009 for Software Assurance customers. This latest release of App-V includes the following new features:

- Full support for Windows 7, including the ability to pin virtual applications to the taskbar and leverage jump lists for navigation to maintain user productivity.
- Increased IT control using AppLocker integration, which helps enforce compliance of virtual applications to provide consistent policy management for all application types.
- Support for BranchCache, a feature of Windows Server 2008 R2 that enables content from file and Web servers on a WAN to be cached on computers at a local branch office. This means that virtual applications only need to traverse the WAN once and will thus be available to users more quickly, eliminating the need for an Internet Information Services (IIS) server in every branch location.
- The ability to secure applications on removable devices with BitLockerToGo and stream virtual applications from USB, thus allowing only authorized users, including remote disconnected users, to access the virtual applications.
- Integration with third-party Lightweight Directory Access Protocol (LDAP) directories to reduce administrative overhead when user accounts are maintained in these directories.
- Several other additional improvements, including instant access to or removal of applications assigned to users.

App-V 4.6

Still in beta at the time of writing this chapter, App-V 4.6 is scheduled to become available in the first half of 2010 and promises the following improvements to the sequencing experience and more:

- A convenient Welcome page that helps you get started with commonly used tasks, such as package creation, editing, and upgrading.

- A redesigned Monitoring page that helps guide you through the various steps of monitoring, including starting the virtual environment, installing applications, and stopping to collect system changes.
- Support for both x64 and x86 Windows platforms, including the ability to sequence true 64-bit applications.

Other improvements in App-V 4.6 haven't been announced yet—stay tuned to the App-V Team Blog at <http://blogs.technet.com/softgrid/default.aspx> for more information as it becomes available.

App-V Terminology

The following are some of the key concepts and terms you need to understand when working with App-V:

- **Active Upgrade** A feature of App-V that provides for automatically upgrading an application on all end-user computers at their next publishing refresh cycle. To gain the benefit of this feature, you must have either an App-V Management Server or an App-V Streaming Server in your environment.
- **Content Folder** A directory, named Content by default, where the virtual application package contents (.sft files) are stored and streamed from. This directory can be in a shared folder on your App-V Management Server, in a highly available Distributed File System (DFS) share, or on a storage area network (SAN) or network-attached storage (NAS) device.
- **Desktop Client** An application that resides on a Microsoft Windows-based computer desktop and that communicates and authenticates with the Microsoft System Center Virtual Application Server to receive the application code and allow a sequenced application to be run locally.
- **Dynamic Suite Composition** A feature that enables a virtual application package to allow dependent plug-ins or middleware packages to use the primary package's registry settings so that the packages behave and interact with one another in the same way as if they were installed locally on a computer. This feature allows applications to be sequenced in separate virtual environments yet selectively communicate with each other.
- **Installation directory** The directory where the installer for the application virtualization sequencer places its files.

- **Management Console** A Microsoft Management Console (MMC) snap-in that is used to administer a specific deployment of the App-V platform that includes all of the components that are managed by a single data store.
- **Management Server** One of two App-V server types (the other being Streaming Server) from which a sequenced application package can be streamed. The Management Server also offers other services, such as publishing, management, reporting, and so on.
- **Microsoft Application Virtualization for Terminal Services** Refers to the client component of App-V running in a Terminal Services environment.
- **Publishing an application** This makes an application available to authorized users whose computers have the App-V Client installed. Publishing delivers the icons (.ico file), package definition information, and content source location (.osd file) to each computer where the App-V Client has been installed.
- **Q: drive** The default virtual application client drive from which sequenced applications are “run”. For more information, see the sidebar titled “Direct from the Source: The Q: Drive” later in this chapter.
- **Sequenced application** An application that has been monitored by the Sequencer, broken up into primary and secondary feature blocks, streamed to a computer running the App-V Terminal Services Client or the App-V Desktop Client, and that can run inside of its own virtual environment. A sequenced application is an application that has been transformed from a traditional installed application to one that runs inside an App-V virtual environment.
- **Sequenced application package** The files that make up a virtual application and allow the virtual application to run. These files are created after sequencing and include .osd, .sft, .sprj, and .ico files.
- **Sequencer** A utility that monitors and records the installation and setup process for applications so that an application can be sequenced and run in the virtual environment. **Sequencing** The process of creating an application package by using the Application Virtualization Sequencer. In this process, an application is monitored, its shortcuts are configured, and a sequenced application package is created containing the .osd, .sft, .sprj, and .ico files. Sequencing is performed by using the Sequencing Wizard, which walks you through sequencing an application, including configuring a package, installing the application or applications to be sequenced, and sequencing the application package for streaming.
- **Sequencing computer** The computer used to perform sequencing and create a sequenced application package.

- **Streaming** The process of obtaining content from a sequenced application package (.sft file) starting with the primary feature block (feature block 1) and then obtaining additional blocks as needed.
- **Streaming Server** One of two App-V server types (the other being Management Server) from which a sequenced application package can be streamed. The Streaming Server only streams applications to the client machines and does not offer other services, such as publishing, management, reporting, and so on.
- **Terminal Services Client** An application that resides on a terminal server and that communicates and authenticates with the App-V Server to receive the application code and allow a sequenced application to be run locally.
- **Virtual application** An application packaged by the Sequencer to run in a self-contained, virtual environment that contains the information necessary to run the application on the client without installing the application locally.

How App-V Works

App-V lets you create virtual applications, which are applications that have been packaged so that they can run within a self-contained virtual environment or “sandbox” on client computers. This virtual environment contains all the information needed to be able to run the virtual application on the client computer and runs within the App-V Client software on the client computer. After you have sequenced applications and deployed the App-V Client software to client computers, you can deliver these applications to the client computers in various ways that resolve many of the issues associated with the traditional application deployment life cycle. The sections that follow examine these processes in more detail.

App-V Virtual Environment

The App-V virtual environment is a run-time container that defines the resources available to application processes launched from a sequenced application package. The resources that are defined by the virtual environment include

- **Virtual COM** A subsystem that manages COM objects created by application processes running in the virtual environment and prevents conflict with the same objects created outside the virtual environment.
- **Virtual directory** An opaque directory where only files and subdirectories defined in the virtual application package or created through interaction with an application in a virtual environment are visible. Any files that are in an identically named local directory are not visible to the virtual application.

- **Virtual file** A file name within the virtual environment that is mapped to an alternate target location. A virtual file appears alongside other files in the containing directory, regardless of whether that directory is virtual or local.
- **Virtual file system** A subsystem that intercepts and redirects file system requests from application processes running in a virtual environment. These requests are processed based on the virtual files and directories defined in the application package and created or modified through interaction with a virtual application.
- **Virtual registry** A subsystem that intercepts and redirects registry requests for keys and values from application processes running in a virtual environment. The redirection is based on the registry information defined in the application package and created or modified through interaction with a virtual application.
- **Virtual services** A subsystem that acts as the Service Control Manager for services running in a virtual environment.

This virtual environment is created by the App-V Client software, which runs on the client computer and enables the end user to interact with virtualized applications after they have been delivered to the client computer. For more information concerning App-V Client software, see the section titled “App-V Clients” later in this chapter.

Sequencing Applications

Before you can use App-V to deliver applications to users on client computers, you first need to package the applications for delivery. The process of packaging an application to enable it to run within its own self-contained virtual environment on a client computer is called *sequencing the application*. Sequenced applications are virtualized and are completely isolated from one another, which eliminates any application conflicts that might occur between two applications.

A sequenced application package contains four types of files that make up a virtual application and allow the virtual application to run. These files are created after sequencing and include the following types of files:

- **.ico file** This is the type of file for the icon on the client’s desktop used to launch a sequenced application.
- **.osd file** This is an XML-based Open Software Descriptor file that instructs the client on how to retrieve the sequenced application from the App-V Management Server or Streaming Server and how to run the sequenced application in its virtual environment.
- **.sft file** This type of file contains one or more sequenced applications that the Sequencer has packaged into streaming blocks, as well as the associated delivery information. An .sft file is stored on each server that must stream the packaged applications to a client.

- **.sprj file** This is an XML-based Sequencer Project file in which the Sequencer stores its Exclusion Items and Parse Items information. An .sprj file is used in the creation of application records and when upgrading a package.

In addition, a sequenced application package can also contain a Microsoft Windows Installer (.msi) file that can be used for standalone distribution of virtual applications, for publishing application packages using an electronic software distribution (ESD) system such as Microsoft System Center Configuration Manager 2007, or for both purposes.

For more information concerning sequencing applications, see the sections titled “App-V Sequencer” and “Using the Sequencer” later in this chapter.

Publishing Applications

After an application has been sequenced to create a virtual application package consisting of the aforementioned files, the application must be published on the App-V Management Server. Publishing an application delivers the icons, package definition information, and content source location to each client that has the App-V Client installed. There are three publishing delivery methods supported by App-V:

- Using the App-V Management Server
- Using an ESD system such as System Center Configuration Manager 2007
- Standalone delivery

For organizations that already have an existing ESD system in place, using this publishing delivery mechanism provides the benefits of reducing the cost of acquiring and deploying additional hardware, operating systems, and database licenses. Leveraging your existing ESD infrastructure can also help your organization avoid the support issues associated with needing to maintain two infrastructures.

If you use ESD as your publishing delivery mechanism, you can choose from the following three approaches for publishing the application to the clients:

- **MSI files** Uses Microsoft Windows Installer (.msi) files
- **MSI Manifest** Uses the MSI Manifest contained in the .msi file.
- **SFTMIME commands** Uses a command-line window and SFTMIME commands for adding the applications and loading the .sft file.

Streaming Packages

After an application has been published and its .ico and .osd files have been streamed to the client, the virtual application package content file (.sft file) must be delivered to the client. App-V supports various ways of doing this, including using the App-V Management Server,

an Internet Information Services (IIS) Web server, a file server, standalone delivery, or a distribution point running IIS within a System Center Configuration Manager 2007 environment.

The first time a user double-clicks on an application icon that has been placed on a computer via the publishing process, the App-V Client first performs authorization and license checking. The client then begins streaming the virtual application package content (.sft file) from the configured streaming source location. The way this works is that the .sft file is mounted in RAM on the streaming server, which then delivers the application in blocks of 32 KB size by default over the wire to the client. The streaming source location is typically a server that is “local” to (accessible over a well-connected network) the user’s computer, but some electronic distribution systems such as System Center Configuration Manager 2007 can distribute .sft files to a folder on the user’s computer and then stream the package from that local folder. A streaming source location for virtual application packages can even be set up on a computer that is not a server—that is, on a workstation. This type of solution can be especially useful in a small branch office location that has no server.

Virtual Application Management

App-V greatly simplifies application deployment by helping you resolve several key issues that often arise during the traditional application management life cycle:

- App-V can help resolve the kind of problems that can arise when you install two applications that are incompatible with one another onto the same computer. Because each virtual application deployed using App-V runs within its own isolated virtual environment, registry and file conflicts are significantly reduced between different virtual applications running on the same client computer.
- App-V can help reduce or eliminate the time-consuming regression testing that is needed before deploying applications onto client computers to ensure that application-compatibility issues are detected before the applications are installed.
- App-V can help reduce the headache of maintaining applications installed on client computers by applying service packs, security fixes, and other types of software updates.
- App-V helps prevent the mess that can result when applications that are no longer needed by users are uninstalled from their computers but leave remnant files and registry settings that can create conflicts later on when other applications are installed.

For more information on how App-V helps resolve these problems and other kinds of issues associated with the traditional application management life cycle, see the sidebar titled “Direct from the Source: App-V and the Application Management Life Cycle” in this chapter.

Direct from the Source: App-V and the Application Management Life Cycle

Every organization faces the challenges of deploying, updating (that is, installing patches, service packs, and upgrades), supporting (such as troubleshooting, license compliance, and training), and terminating all the applications in the enterprise in an ongoing cycle. There are many diverse solutions to this problem, but most of these solutions target only one or two areas of this life cycle.

Deployment is the initial process that organizations take to install their applications onto the client machines. This can be done using any of several traditional methods, ranging from having a support engineer touch every client PC with the install media to remotely using an electronic software distribution method to having the client access a terminal server (or Citrix XenApp server) remotely.

Updates are a natural and necessary process in the application management life cycle. As applications gain maturity in the market place, they will undoubtedly be revised and updated through service packs or hot fixes. It is the burden of the support engineer in an enterprise to update every client PC that has a particular application version installed on it.

Providing support for the entire library of applications in an enterprise environment can be a daunting task. Issues from the most basic elements of application conflict to users inadvertently damaging their own application installations by deleting critical files can result in considerable overhead to an organization's support team.

Termination of installed applications is the last phase of the application management life cycle. Applications are eventually replaced or retired and, as such, need to be removed from the client environment. In a traditional management model, this might necessitate the support engineer visiting those clients PCs and uninstalling that application. In doing so, it is not 100 percent guaranteed that the process will remove all of the files and registry entries, meaning some are left to be orphaned and possibly to cause an issue later on.

One of the most basic challenges that enterprises encounter with the standard application management life cycle is the possibility that two or more applications will conflict with each other. When an application is allowed to install onto a host or client computer, its programmed behavior is to add or modify files and registry settings on that client's operating system. If that application did not add or modify settings in the registry and was contained exclusively in its own directory on the file system, conflicts would never occur. However, because thousands of applications have been made available since 1995, the application will almost always place its own files (.dll, .vxd, .sys, and

so on) in numerous directories throughout the file system. It will also populate the registry of that client with its own values or modify existing values. One of the advantages of deploying applications this way is that they will take advantage of the local resources on the device onto which they were installed.

Having incompatible applications on a system would not be so frequent an occurrence if enterprises ran only one or a very select group of applications. Because most enterprises need to run countless applications, ending up with applications that conflict with one another is almost guaranteed. Until App-V Application Virtualization, the solution for many organizations was to separate these applications by placing them on different computers. In essence, a user would have two or more computers on her desk, one to run Application A and another separate computer to run Application B, although one of these computers could be a virtual machine running on the user's physical computer. Although this is a possible solution, it is very impractical from an expense and support perspective.

Applications that are App-V enabled are never allowed to install or modify the local file system or local registry. When an application is App-V enabled, it is made to run inside its own virtual environment. (See Figure 3-31.) Contained inside of this virtual environment are all the files, registry information, fonts, COM, embedded services, and environment variables that the application normally would have installed and been expected to use on the client PC. Instead, with all of these assets residing inside the virtual environment, the application leverages them from this virtual environment and remains isolated from other applications, which are also running inside of their own separate virtual environment. The process of creating the virtual environment is known as sequencing.




FIGURE 3-31 The App-V virtual environment.

The Sequencer is the component in the App-V system that is vital in creating the SystemGuard for an application or suite of applications. App-V-enabled applications will be able to use local and network drives, CPU, RAM, the local Windows Installer

Service, and other local resources on the App-V Client to which they are streamed to, cached, and run.

Regression testing has always been a top priority inside enterprises when the enterprise is deploying any new or updated application. In enterprises with formalized processes for doing this, every application is tested against every known configuration that could exist within that organization. This could, and often does, exceed 40 hours of effort for a single application. With an App-V-enabled application, however, this is reduced to only the time required to sequence that application. When the application is deployed, it is isolated from any other applications that were sequenced or that are still locally installed on the client, guaranteeing a conflict-free environment.

Updates

It is a natural part of an application's life cycle to have updates in the form of service packs or hot fixes become available. These updates need to be applied to that application, and this is often done by having the support engineer visit every client PC or terminal server machine and manually apply the update in question. This process can be very time consuming, and it can also increase the likelihood of causing an application conflict because the update modifies files and registry settings on the client. With an App-V-enabled application, updates are performed centrally and occur at only one time. The sequence engineer takes the original App-V-enabled application's package back to a clean sequencer workstation and performs a package upgrade, appending the original package with the updates. This updated package is then used to replace the original package on the App-V Server, and the App-V Client receives the updated files seamlessly.

Support

In addition to eliminating application conflicts, the App-V platform can solve many other support-related issues. When you run each application inside of its own protected SystemGuard environment, App-V-enabled applications remain immune to users inadvertently or intentionally deleting critical files needed by that application to run. Because the App-V-enabled applications are running inside of their own SystemGuard environment, users and local system administrators never see any of the application's files or registry entries if they look at those local resources. This can effectively reduce the number of help desk calls an organization requires. Another issue facing support personnel inside an enterprise is the concern surrounding licensing. App-V enables organizations to control the number of users who can gain access to App-V-enabled applications concurrently. This licensing feature is administered centrally from the sole administrative utility for Microsoft App-V Application Virtualization, the App-V Management Console.

Termination

At the end of an application's life cycle, it is time to retire or terminate that application. In a traditional method, someone is required to visit every client PC or terminal server machine and uninstall that application. This approach has the potential to leave some files and registry settings orphaned and create conflicts later on. With App-V-enabled applications, the organization simply needs to deactivate or remove that retired application centrally from the App-V Management Console. By doing so, the users subsequently have the application removed from their desktop, and all previously cached data blocks of the application are removed as needed. Because applications are no longer truly installed when App-V is used, there is never a need to remove the application from the client's computer.

*—Sean Donahue, Senior Program Manager,
System Center Alliance, Microsoft Corporation*

App-V Components

The App-V environment consists of the following components:

- App-V Management Server
- App-V Management Web Service
- App-V Data Store
- App-V Streaming Server
- App-V Management Console
- App-V Sequencer
- App-V Client

In addition, you can publish your virtual application packages using your existing electronic software distribution system such as Microsoft System Center Configuration Manager 2007 instead of using the App-V Management Server. You might also need one or more file servers, Web servers, or both, depending on how you want to use App-V.



Note Not all of these components need to be installed in your environment; the components that need to be installed depend on how you plan on using App-V to deploy virtual applications to users. For more information on different App-V deployment scenarios, see the section titled “App-V Deployment Scenarios” later in this chapter.

The sections that follow provide more information concerning each component of the App-V environment.

App-V Management Server

The App-V Management Server is used for streaming the virtual application package content and for publishing virtual application shortcuts and file type associations to the App-V Client. Because the Management Server streams virtual applications to end users on demand, these servers are ideally suited for environments that have reliable, high-bandwidth local area networks (LANs), such as head office environments. The Management Server also supports Active Upgrade; the Publishing Service, which is used by the client to retrieve the applications that the logged-in user has access to; and licensing and metering capabilities.

The Management Server should be installed on a dedicated server computer and needs access to a Microsoft SQL Server database that can either be installed on the same server or on a different server on your network. Microsoft SQL Server is used to manage the database and data store for the App-V environment. You can deploy a single Management Server or use many of them. In a typical App-V environment, multiple Management Servers share a common data store for configuration and package information. For more information concerning the App-V Data Store, see the section titled “App-V Data Store” later in the chapter.

The Management Server also needs access to the Content folder, which is a repository for the virtual application packages you want to publish and stream to the client computers on your network. The Content folder is where the SFT files are loaded and stored, and it can be located on either the Management Server itself, on a separate file server on your network, on a Distributed File System (DFS) share, or on a SAN. For more information concerning the Content folder, see the sidebar titled “Direct from the Source: Using the Content Folder” in this chapter.

The Management Server handles user requests for application data and then streams this data on demand to authorized users. This streaming of application data takes place using one of the following protocols:

- Real-Time Streaming Protocol (RTSP)
- Real-Time Streaming Protocol Secure (RSTPS), which is RTSP over Transport Layer Security (TLS)
- Hyper-Text Transfer Protocol (HTTP)
- Hyper-Text Transfer Protocol Secure (HTTPS), which is HTTP over Transport Layer Security (TLS)

You configure and manage the Management Server by using the Application Virtualization Management Console, which is described in the section titled “App-V Management Console” later in the chapter.

App-V Management Web Service

The App-V Management Web Service is the component responsible for communicating read/write requests to the App-V Data Store. The App-V Web Service functions as an intermediary between the Management Console and the Data Store.

Note that even though the administrator makes his changes in the GUI of the App-V Management Console, those changes do not get written to the Data Store by this MMC console. Instead, the Management Console makes a .NET Remoting connection to the Management Web Service. This service then makes an OLE DB connection to the SQL Data Store and performs the actual read/write operations.

The App-V Management Web Service can be installed either on the Management Server itself or on a separate server that has IIS 6.0 or higher installed. In addition, Microsoft Data Access Components (MDAC) 2.7 or higher and the .NET Framework 2.0 must be installed on the server running the App-V Management Web Service in order to allow connectivity with the data store.

App-V Data Store

The App-V Data Store is a required component when you deploy an App-V Management Server. The data store is responsible for storing all information related to the App-V infrastructure, including the following:

- App-V Management Server configuration information
- App-V Management Server reporting information
- Application records
- Application assignments
- Application licensing information
- Logging information

The Data Store consists of a SQL Server database that can be installed on either Microsoft SQL Server 2005 or Microsoft SQL Server 2008.

When a user tries to launch an application that has been virtualized using App-V, the Management Server that receives the user's request contacts Active Directory Domain Services for authorization and the data store for application licensing information.

App-V Streaming Server

The App-V Streaming Server is responsible for hosting and streaming virtual application packages to App-V clients. You can think of the Streaming Server as a lightweight version of the Management Server that includes only streaming functionality, doesn't include the

App-V Management Web Service or the Management Console, and doesn't require using a Microsoft SQL Server database. Instead, the Streaming Server uses access control lists (ACLs) for granting user access to the package files. The Streaming Server also supports Active Upgrade, but it doesn't have a Publishing Service, licensing or metering capabilities.

Like the Management Server, the Streaming Server also needs access to the Content folder, which is the repository for your virtual application packages. The Content folder can be located either on the Streaming Server itself, on a separate file server on your network, or on a SAN.

The Streaming Server can be used in environments that have an existing ESD, such as System Center Configuration Manager 2007. The Streaming Server can be used together with the Management Server. For example, the Streaming Server can be used at a branch office while the Management Server is deployed at the head office, with a slow wide area network (WAN) link between the two locations. Alternatively, the Streaming Server can be used alone without the Management Server in environments that don't have the infrastructure to support the Management Server. For more information on different App-V deployment scenarios, see the section titled "App-V Deployment Scenarios" later in this chapter.

App-V Management Console

The App-V Management Console is an MMC snap-in you can use to manage your App-V environment. Using the Management Console, an administrator can do the following:

- Import applications
- Manage file type associations for applications
- Manage application licenses
- Create and manage server groups
- View and configure server settings
- Create provider policies
- Generate reports

The Management Console can be installed locally on the Management Server, and it can also be installed on any workstation that has MMC 3.0 and .NET Framework 2.0 installed to allow remote management of the App-V environment.

Figure 3-32 shows the layout of the Management Console and displays a list of the applications that have been sequenced on the local Management Server.




FIGURE 3-32 The Application Management Console showing a list of virtual application packages on the local App-V Management Server.

App-V Sequencer

The App-V Sequencer is a wizard-based tool that can be used to monitor and capture the installation of an application to create a virtual application package you can publish and stream to client computers. After an application has been sequenced, the resulting App-V-enabled application package can be delivered to users on demand to run within an isolated virtual environment on the user's computer.

The output of the sequencing process includes an application's icon (.ico) files, an .osd file containing the package definition information, a package manifest file (manifest.xml), and the .sft file that contains the application program's content assets. The sequencing process is performed once for each application or suite of applications you want to virtualize, and the process protects the application's integrity by not making any modifications to the source code of the application. After an application has been sequenced, its files must be copied to the Content folder before they can be streamed or published to the App-V Client. Alternatively, the .ico and .osd files can be hosted on a Web server and delivered to the App-V Client using HTTP or HTTPS.

The Sequencer component typically must be installed on a separate computer from the other App-V components. This separate computer is called the *sequencing computer*. This sequencing computer needs to be a clean image that can be restored back to its virgin state at the end of every successful sequencing operation.

During the sequencing process, the Sequencer is first placed in monitor mode. The application to be sequenced is then installed on the sequencing computer. The sequenced application is then started, and common tasks are performed with the application so that the monitoring process can configure the primary feature block, which contains the minimum application package content that is needed for the virtualized application to run properly. When all of these steps are finished, monitoring mode is stopped and the sequenced application is saved. The sequenced application should then be thoroughly tested to ensure that it works properly when virtualized.



Tip Some applications cannot be sequenced, including Internet Explorer, device drivers, applications that start services at boot time, and some other parts of the Windows operating system.

For more information about sequencing applications, see the section titled “Using the Sequencer” later in this chapter.

App-V Client

The App-V Client is the software component that resides on the client computer and provides the virtual environment for running virtual applications. The App-V Client also handles the streaming of the application content from a Management Server and also from a Streaming Server if one has been deployed. The streaming process structures the application content so that the initial user interaction is streamed to the client computer first. This is done so that the user can launch the application immediately without needing to wait for the entire application content to be streamed to the client. Users can launch virtual applications by clicking on icons on their desktop or Start menu, or by double-clicking on file types associated with the application.

There are two kinds of App-V Client software:

- **App-V Desktop Client** This client is used on standard desktop computing environments. The App-V Desktop Client is included in the Microsoft Desktop Optimization Pack (MDOP) for Software Assurance (SA). The App-V Desktop Client is installed on end-user workstations to C:\Program Files\Microsoft Application Virtualization Client and is responsible for caching and launching virtualized applications. The App-V Desktop Client turns desktop applications into services to be deployed on demand without installation and without administrators having to be concerned about conflicts with any existing applications. The App-V Desktop Client also allows applications to be centrally managed with real-time license compliance.
- **App-V Terminal Services Client** This client is used in Terminal Services environments. The App-V Terminal Services Client behaves much like the App-V Desktop Client except that it provides for installation on a terminal server, which hosts the virtualized application instead of having the virtualized application run directly on the client computer.

The App-V Terminal Services Client allows administrators to deliver any application to any Terminal Services or Citrix XenApp server without having to perform the installation, without being concerned about conflicts or testing, and without disruption of service.

The App-V Client must be configured at installation time using the Client Management Console to specify the name or IP Address of the Publishing Server that it contacts at login to retrieve application icons and .osd files that the user has access to. If you are going to use the App-V Desktop Client, you must also deploy this software to your client computers. This is typically done using ESD system such as Microsoft System Center Configuration Manager 2007, but it can also be done using other methods, such as Group Policy Software Installation, scripting, or even manual installation.

App-V Architecture

Figure 3-33 illustrates the App-V architecture, showing the different components of the App-V platform and the protocols and other transport mechanisms used for communications between these components.




FIGURE 3-33 The App-V architecture.

The section titled “App-V Components” that preceded this section summarized the function of each of the different App-V components. The following is a summary of the protocols and other transport mechanisms used for communications between these components:

- **.NET Remoting** A component of the .NET Framework that enables client applications to use objects in other processes on the same computer or on any other computer available on its network. This used by the Management Web Service to connect to the SQL Data Store.
- **ADSI** Active Directory Service Interfaces, a set of COM interfaces used to access the features of directory services from different network providers. ADSI is used in a distributed computing environment to present a single set of directory service interfaces for managing network resources. This is used by the Management Server to retrieve user group associations from Active Directory.
- **HTTP** Hypertext Transfer Protocol, an application-level protocol for distributed, collaborative, hypermedia information systems, such as text, graphic images, sound, video, and other multimedia files on the World Wide Web. This can be used to stream the SFT file to the App-V clients.
- **HTTPS** Hypertext Transfer Protocol over Secure Sockets Layer, an extension of HTTP that securely encrypts and decrypts Web page requests using Secure Sockets Layer (SSL). SSL is a security protocol that supports confidentiality and integrity of messages in client and server applications communicating over open networks. This can be used to stream the SFT file to the App-V clients with added security.
- **OLE DB** A set of COM-based interfaces that expose data from a variety of sources. OLE DB interfaces provide applications with uniform access to data stored in diverse information sources or data stores. This is used by the Management Web Service to connect to the SQL Data Store.
- **ODBC** Open Database Connectivity, a universal data access interface that enables applications to concurrently access, view, and modify data from multiple, diverse databases. This is also used by the Management Web Service to connect to the SQL Data Store.
- **RTSP** Real-Time Streaming Protocol, an application-level protocol that controls the transport of multimedia content, session announcements, and tear-downs. When the App-V client communicates with the App-V Management Server, the client uses RTSP over port 554 to establish the initial connection with the server. After the initial connection has been made, however, the client continues to send and receive blocks of the streamed application package content using two other protocols, the Real-Time Transport Protocol (RTP) and the Real-Time Control Protocol (RTCP). These two protocols open connections with the clients, starting with ports 49152 up to 65535, concurrently for send/receive.

- **RTSPS** Real-Time Streaming Protocol Secure, which is RTSP over Transport Layer Security (TLS). When the App-V client communicates with an App-V Management Server that has a certificate assigned to it, the client uses RTSPS over port 322 to establish the initial connection with the server and then uses RTP and RTCP for streaming of blocks of application package content. If there is no certificate assigned to the server, the communication uses RTSP over port 554 if the option to allow nonsecure connections is selected.
- **SMB** Server Message Block, a protocol used to request file and print services from server systems over a network. Standard ports are used when App-V is deployed in a trusted environment, such as a corporate LAN, while restricted ports are used when App-V is delivering virtual applications to untrusted clients—for example, over the Web. Restricted reports require that a server certificate be installed on the Management Server during installation of the server, and also on any file, Web servers, or both that are used to stream application package content to clients.

Table 3-5 lists the various ports that must be open for App-V components to communicate with one other.

TABLE 3-5 App-V Communications Ports

Communications Function	Standard Port	Protocol	Restricted Port	Secure Protocol
Between the Management Console and the Management Web Service	80	HTTP	443	HTTPS
Between the Data Store and the Management Web Service	1433	ODBC	1433 (IPsec)	ODBC
Between the Data Store and the Management Server	1433	ODBC	1433 (IPsec)	ODBC
Between App-V clients and the Management Server	554	RTSP	322	RTSPS
Used by RTSP and RTSPS to manage communications after initial communication has been established between App-V clients and the Management Server	49152-65535	RTP RTCP		



Note For more information on the protocols and other transport mechanisms used by App-V, see article KB 932017 in the Microsoft Knowledge Base on Microsoft TechNet at <http://support.microsoft.com/kb/932017>.

Working with App-V

Microsoft App-V provides organizations with powerful and flexible solutions for delivering virtualized applications to end users. This section covers some of the basics of working with App-V and includes the following topics:

- App-V deployment scenarios
- Obtaining App-V
- Using the Management Console
- Using the Sequencer
- Working with App-V clients



Note A full treatment of how to deploy, configure, use, and maintain an App-V environment is beyond the scope of this chapter. For detailed information on these topics, see the “Planning and Deployment Guide for the Application Virtualization System” and “Operations Guide for the Application Virtualization System” on Microsoft TechNet at <http://technet.microsoft.com/en-us/library/cc843770.aspx>.

App-V Deployment Scenarios

Microsoft App-V has a great deal of flexibility in how it can be deployed within an organization. For example, you can choose from the following three publishing delivery methods for transferring the .osd and .ico files from their designated location (typically, \Content) to the App-V Client software running on client computers:

- Using the App-V Management Server alone as the core of your virtual application deployment platform to transfer the .osd and .ico files from the Content folder on your Management Server. You can also use an App-V Streaming Server, such as an IIS Web server or file server, to host your Content folder if you already have such infrastructure in place and want to use it.
- Using an electronic software distribution (ESD) system, such as System Center Configuration Manager 2007, to move .osd and .ico files to clients via standalone Microsoft Windows Installer (.msi) files, the MSI manifest contained within .msi files, or SFTMIME commands.
- By standalone (locally-installed) delivery of .msi files to users via network shares, removable media, or some other method.

You also have several choices as to which package delivery method you use to stream the virtual application packages or .sft files from your Management Server to the App-V Client software running on client computers:

- An IIS Web server, which streams package content using HTTP or HTTPS
- A file server, which streams package content using SMB
- An App-V Streaming Server, which streams package content using RTSP, RTSPS, or HTTP/S (if IIS is installed)

You can also deliver package content via standalone delivery or by using an ESD system, but the most typical ways of streaming packages are by using a Streaming Server, IIS server, or file server. Table 3-6 lists some of the advantages and disadvantages of each of these package delivery methods.

TABLE 3-6 Advantages and Disadvantages of Different Package Delivery Methods

Package Delivery Method	Advantages	Disadvantages
App-V Streaming Server	Supports Active Upgrade and RTSPS for enhanced security; needs only one firewall port open	Requires supporting a dual infrastructure; requires additional server administration
IIS server	Supports streaming to remote clients via the Internet and HTTPS for enhanced security; highly scalable; needs only one firewall port open	Adds the overhead of managing IIS servers
File server	Supports using existing file servers to provide a simple low-cost solution	Does not support Active Upgrade

For more information on the power and flexibility of the various ways you can deploy App-V in your organization, see the sidebar titled “Direct from the Source: App-V—A Scalable Solution for Application Virtualization” in this chapter.

Direct from the Source: App-V—A Scalable Solution for Application Virtualization

One of the features in the new version of App-V that I am most excited about is the number and variety of deployment methods for the virtual application packages. In SoftGrid versions 4.2 and earlier, you were limited to deploying your virtual applications with a full SoftGrid back-end infrastructure only. This could prove to be quite limiting for a branch office environment. In short, you had to deploy a SQL data store at each of your branch offices along with a SoftGrid Streaming Server. App-V, however, is much like choosing from a cafeteria line of selections for how you deploy your virtual applications. You pick and choose the options you want—and those you don’t want, you leave under the warming lights for later.

The way I look at App-V, today is by looking at the virtual application package first. Because this package can be deployed to multiple clients in various ways, the package is the constant. For example, I can take the same package and deploy it to a collection of Windows Vista clients at Corporate HQ using what is considered to be a traditional method. That is, I use the App-V Management Console to publish the application to a group of Active Directory users. This is written in the SQL data store, and the package files (.SFT, .ICO, .OSD, .SPRJ) are stored in the Content folder of the Streaming Server. The client logs in and contacts his Publishing Refresh Server, which is most often the same server as the App-V Streaming Server, and gets a list of his applications from the SQL data store. The icons and OSD files are transferred to the client, and upon initial launch of the icon the SFT file starts streaming, using RTSPS, to the client. Nothing new here.

However, if I'm in a branch office instead, I could use what is referred to internally as the Lightweight Streaming Server. This is an App-V server whose only purpose and function is to stream App-V SFT files using RTSPS (by default). As an administrator, I publish the App-V applications to my user groups in Active Directory as I did in the traditional model. I then copy the SFT file to a server in the local branch office that had the lightweight streaming server (LWS) installed on it. The user logs on to his computer, authenticates, and then contacts the Publishing Refresh Server to receive a list of applications he has permissions to. The icons and OSD files transfer to the client as they do with the method used at Corporate HQ. However, when the user launches the application shortcut, the application streams from the local LWS instead of over the WAN from Corporate HQ. This happens because as the administrator, I have set the Application Source Root in the client's registry that told the clients in the branch to override whatever the HREF line in the OSD file said and use the local branch server instead. Also, I can set the Icon Source Root and the OSD Source Root in the same way and have all traffic, except the refresh, occur from the local LWS. But wait! There's more!

What if I told you that in addition to this you could now deploy your App-V packages to remote users who did not have regular access to an office connection? During the sequencing process you could select the check box that generates an MSI file in addition to the standard App-V package's files. You could then deploy the MSI and SFT files to a location accessible to this remote user—for example, to a DVD or local share. The user double-clicks the MSI file, and it uses the Windows Installer Service to "install" the virtual application. Rest assured that nothing is actually installed. Instead, what really happens is the installer service calls one of the App-V client executables called Sftmime.exe. If you open the MSI in Orca or another MSI edit utility, you see a bunch of SFTMime commands that basically add the application to the client, publish the shortcuts, add the OSD, and load the SFT file into the local file system cache, sftfs.fsd. This is combined with a registry setting that sets the RequireAuthorizationIfCached option to 0. I use this all the time when testing applications from independent software

vendors (ISVs) that I've sequenced. I simply copy the MSI and SFT to a client and launch it. No back-end whatsoever is needed.

As if that weren't enough, you also have the ability to use an ESD, such as System Center Configuration Manager R2. With this option, the administrator advertises the package in System Center Configuration Manager as he would with a physically installed application. However it is really a Virtual Application instead. By using System Center Configuration Manager, you would still need the App-V client on the desktop in addition to the System Center Configuration Manager Advanced Client. System Center Configuration Manager uses a new file added to the App-V packages called the manifest file (`_manifest.xml`) which stores information about all of the applications in the package. This file is used to populate several of the fields in the new applications added to the System Center Configuration Manager console. System Center Configuration Manager also allows distribution points, which are most likely already established in the organization, to act as the streaming points for the App-V packages. When the System Center Configuration Manager Advanced Client does its policy refresh it will pick up the App-V applications just like it does the physically installed Apps. The beauty here is that this requires no special App-V infrastructure, it uses all of the existing System Center Configuration Manager configuration. It also allows the publishing of App-V packages to physical collections and not just users.

Sometimes I feel like a proud Papa bragging on and on about how great their prodigy child is. But this last feature on scalability that I will call out here in this sidebar is the introduction of HTTPS streaming instead of RTSPS. Now to some this might seem an anticlimactic way of ending. "Isn't this an obvious evolution?" One might say. As obvious as it might seem it introduces a whole new world to App-V. Imagine being able to stream virtual applications that never change or alter the foot print on a client devices, over the Internet. Yes Virginia, there is a Santa Claus. This means that a company, or an ISV, could host their applications on a Web server and deliver those virtual applications to their clients anywhere in the world. The client would still need the App-V client and a Web server would need to be configured with a content location under the root. But the question is, "Who doesn't have a Web server?"

Back in my day we didn't have these fancy delivery methods. We had to have a Streaming Server and SQL data store in every branch office. And we liked it! Nowadays you kids have it easy. You take your App-V package and pick and choose your delivery method. "I think I'll stream this to these clients with RTSPS. But deliver it to these clients with an LWS, and then burn to DVD and mail it to these remote users. Oh but for this segment of machines I'll use System Center Configuration Manager to advertise to a collection and then stream it over the Web using HTTPS to these customers." What's next? A better network topology than ARCNet?

*—Sean Donahue, Senior Program Manager,
System Center Alliance, Microsoft Corporation*

Deploying App-V at a Single Site

If your organization is located at a single site and has a fast, reliable LAN throughout, you can deploy App-V using the traditional or classic approach familiar to earlier SoftGrid 4.2 administrators. (See Figure 3-34.) Here are the App-V components you need to deploy for this scenario:

- App-V Management Server
- App-V Management Web Service
- App-V Data Store
- Content Folder location
- App-V Management Console
- App-V Client software

For smaller sites, all of the components just listed can be installed on a single server with the Content folder located on any of the following:

- A share on the server itself
- A highly available DFS share
- A highly available SAN/NAS device




FIGURE 3-34 Deploying App-V at a single site.

After you've deployed the App-V components needed for this scenario, you can use the Sequencer to create virtual application packages and copy the package folder to the Content

share to create a subfolder under Content for each application. Then the administrator can publish each application to groups of users in Active Directory Domain Services so that when the user logs on to his computer he sees shortcuts on his Start menu and desktop to launch these applications. When the user double-clicks on a shortcut, the App-V Client on the user's computer streams the .sft file for the application package from the Management Server and then launches the application for the user to use. The application package is also cached locally on the user's computer so that the application can be launched more quickly next time the user needs to run the application.

Deploying App-V at Branch Offices

Larger organizations that include branch offices at remote sites can add another App-V component, the App-V Streaming Server, to enable users to efficiently use virtual applications that are provisioned from the head office site over slower WAN links. For this scenario, you can deploy a Streaming Server at each branch office and your remaining App-V components at the central head office location. (See Figure 3-35.) In this scenario, virtual applications are published to client computers at the branch office over the WAN link while application package content is streamed to these clients over the branch office LAN. In this branch office scenario, it is possible to either have the .ico and .osd files delivered to the client over the WAN or to modify the client's registry so that these files are also delivered from the local branch office's server.




FIGURE 3-35 Deploying App-V for branch offices.

Deploying App-V Using an Existing ESD System

Large enterprises that already have an electronic software distribution (ESD) system in place can leverage their existing infrastructure to provide application virtualization to users in a

number of different ways. For example, Figure 3-36 shows an enterprise that has two sites (head and branch offices) plus external users who need to access virtual applications over the Internet. In this particular implementation, client computers at the head office stream virtual application package content from an App-V Streaming Server. This gives these clients the benefits of Active Upgrade, a feature of App-V that requires App-V servers to be deployed and that enables automatic upgrading of virtual applications on end-user computers at their next publishing refresh cycle. By contrast, client computers at the branch office stream their virtual application package content directly from the ESD distribution point running on a file server on their local network. This scenario is a simple, low-cost solution that uses SMB to stream the package content from an existing file server that hosts the Content Folder for the clients, but it doesn't support Active Upgrade. Finally, an IIS server on the perimeter network at the head office is used to stream virtual application package content over the Internet using HTTPS to external client computers such as mobile users with laptops, which is another scenario that is simple to implement but doesn't support Active Upgrade.




FIGURE 3-36 Sample App-V deployment leveraging an existing ESD system.

Standalone Deployment of Virtual Applications

The App-V Sequencer also has an option to create an .msi file that automates the “installation” of a virtual application. This .msi file contains additional metadata that enables an ESD system to recognize and control virtualized applications. Standalone mode requires the App-V Client to be configured for standalone mode, which allows only .msi-based updates of virtual applications. (Streaming is not allowed while in standalone mode.) This mode is meant to be used by rarely connected users that need the power of virtualized applications but do not have access to a server. So, in this scenario, you need to provide the .msi file to the user directly—for example, on CD or DVD media.

It is important to note that the package’s content file (.sft) is not included inside of the .msi file. The .sft file must be available in tandem with the .msi. The .ico and .osd files, however, are in the .msi file. When the user double-clicks the .msi file for a virtual application, a series of commands using the App-V client’s built in SFTMIME command are used to add the application and load the .sft file into the local cache. In this standalone deployment scenario, third-party products such as FullArmor’s GPAnywhere can be used to apply Group Policy settings to the virtual applications.

Deploying App-V with Terminal Services

You can also deploy App-V with Terminal Services so that users can run virtual applications on a terminal server instead of on their local computers. In this scenario, the user employs the App-V Terminal Services Client instead of the App-V Desktop Client. For more information on the benefits of deploying App-V together with Terminal Services, see the blog post titled “SoftGrid and Terminal Services: Better Together” on the Microsoft Application Virtualization team blog at <http://blogs.technet.com/softgrid/archive/2008/04/10/softgrid-and-terminal-services-better-together.aspx>.

Using the Management Console

The App-V Management Console is your central location for performing all App-V-related management tasks. The following sections describe common administrative tasks you can perform when using this console.

Managing Applications

The Applications node of the Management Console lets you perform the following tasks:

- **Import an application** This action makes an application available for streaming from an App-V server. To import an application, you need to have either its .osd or .sprj file available on the server. Importing an application automatically creates a package for the application during the import process.

- **Manually add an application** This action requires that you manually specify all the information that is normally determined automatically by the Import Applications Wizard. Adding an application manually also requires that you manually add a package for the application. For information about adding packages, see the next section.
- **Grant or deny access to an application** These actions let you specify which groups of users will be allowed to access the application.

Other actions you can perform using this node include renaming, deleting, and moving an application and changing an application icon.

As an example of how you can manage applications using this console, the following procedure demonstrates how you can import an application. Figure 3-37 shows the Applications node before any applications have been imported.




FIGURE 3-37 No applications have been imported yet.

Right-clicking on the Applications node brings up a shortcut menu. From this menu, select Import Applications. In the Open dialog box that appears, browse to locate the .osd or .sprj file for the application. (See Figure 3-38.) By using the .osd file, the administrator would need to add each application in a suite individually. By importing from the .sprj file, however, all the applications in the suite will be imported at one time.




FIGURE 3-38 Selecting an application to import.

After you've selected the application you want to import, click Open. This launches the New Application Wizard. The first page of this wizard is automatically populated with the name, version, OSD path, icon path, application license group, and server group of the application. (See Figure 3-39.) If you want to stream the application to clients, make sure that the Enabled check box is selected as shown. You can also add a description for the application if desired. Verify that the remaining information displayed is correct before proceeding with the wizard. The OSD path and Icon path should reference either a UNC path or a URL because this will tell the user where those files can be copied from during the client's refresh.




FIGURE 3-39 The New Application Wizard.

The next page of the wizard lets you specify the locations where you would like application shortcuts to appear on client computers. (See Figure 3-40.) By default, these will mimic what the application's installation would have populated as captured during sequencing.




FIGURE 3-40 Specifying the application shortcuts that should appear on client computers.

The next page of the wizard displays the file associations that are currently configured for the application. (See Figure 3-41.) This screen also allows you to add new file associations for the application by clicking the Add button and to edit or remove existing file associations by clicking the corresponding buttons for these actions. By default, these will mimic what the application's installation would have populated as captured during sequencing.




FIGURE 3-41 Viewing and modifying file associations for the application.

The next wizard page lets you assign which groups of users will be granted permission to use the application. (See Figure 3-42.)




FIGURE 3-42 Granting access to the application.

The final page of wizard ask you to confirm the selections you have made. (See Figure 3-43.) If any conflicts were detected (such as File Type association) with an already existing application, you will receive a warning in this dialog box.




FIGURE 3-43 Summary page of the wizard.

Clicking Finish imports the application you have selected. The result of the import process is shown in Figure 3-44.




FIGURE 3-44 The application has been imported.

After the application has been imported, it is ready to be streamed to clients.

You can also use the Applications node to organize your applications into groups for easier management. For example, you might create application groups for specific departments, divisions, or sites in your organization. You can also create application groups for specific types of applications such as enterprise resource planning (ERP) applications, customer resource management (CRM) applications, and so on. Using application groups also makes it easier to grant permissions to applications and manage application licenses.

Managing Packages

Virtual application packages can be managed using the Packages node in the Management Console. (See Figure 3-45.) Packages let you control virtual application versions on your App-V Management Servers. Using the Packages node, you can

- Manually add a package by specifying the package name and the path to the application's .sft file.
- Add a new version of a package. (You can leave the previous version in place for compatibility reasons if needed.)
- Delete all versions of a package or only the specified version.
- Upgrade a package. (An automatic upgrade occurs when you perform an Open For Package Upgrade action of an existing package in the Sequencer.)




FIGURE 3-45 Managing packages.

Managing Application Licenses

You can use the Application Licenses node to add, remove, configure, and control application license groups. Depending on the type of license group you create, you will be able to control which users have access to your applications and how many users will be allowed to access applications at a given time. App-V helps administrators ensure license compliance: if there is a license available when a user tries to launch an application, the user is allowed to launch the application; if there is no available license, the client’s system tray will report Launch Failed and an error message will be displayed indicating that there is no available license.



Note License Groups are not application specific. This means that one license group can be applied to multiple applications, although license groups are typically created with specific applications in mind.

Three types of licenses can be created using the Management Console:

- **Concurrent License** This type of license allows a limited number of users to have simultaneous access to the applications that have the license groups assigned to them. Concurrent License groups are the most common type of licensing used for virtual applications and can reduce licensing costs by limiting the number of copies of an application that can be run concurrently.
- **Unlimited License** This type of license allows any number of users to have simultaneous access to the applications that have the license groups assigned to them. Unlimited License groups are useful for evaluating the number of licenses that will be required for an application, and when used in conjunction with Reporting they can assist in your purchasing decisions for applications.

- **Named License** This type of license allows only the specified users to have access to the application associated with the license. Named License groups are typically used for applications whose use must be restricted to certain groups of users, such as administrators, members of the management team, and specially trained users.

Managing Servers

You can use the Server Groups node of the Management Console to manage the App-V servers in your environment, including your Management Servers and Streaming Servers, provided these servers share the same SQL Server database. Specifically, you can

- Create or remove server groups to organize your servers for easier management.
- Add a server to a server group or remove it from a group.
- Adjust the maximum memory allocation for the server cache of a server, specify the maximum block size to be used when streaming content from the server, or modify the port number used for RTSP or RTSPS streaming from the server.

By right-clicking on a server group and selecting Properties, you can manipulate the settings of all App-V servers in the selected server group.

The Default Server Group is created when you set up your App-V environment. (See Figure 3-46.) Small and mid-sized organizations might be able to get by with using only the Default Server Group.




FIGURE 3-46 The Default Server Group.

Managing Reports

The Reports node of the Management Console can be used to generate different kinds of reports that contain information about your App-V system. These reports are created using a SQL Reports run-time agent, which is installed by default with the Management Console. Reports can be viewed, printed, or exported to PDF format to gather information about how your App-V system has been used over the course of daily, weekly, or monthly periods.

Before running a report, metering and logging to a database must be enabled in your App-V system.

The following types of reports can be generated using this node:

- **Application Utilization** This report graphs the total daily and concurrent sessions over time during the reporting period for the specified application. The report uses a simple line graph with an independent y-axis for each metric. The report also lists all users who used the application, as well as the number of sessions, total session duration, average session duration for each user, and a summary of total usage for all users.
- **Software Audit** This report lists usage information during the reporting period for all applications defined in the database. For each application, this report lists the top N users who used the application together with the number of sessions, total session duration, average session duration for each user, and a summary of total usage for all users.
- **System Utilization** This report graphs the total daily and concurrent usage over time during the reporting period for the specified server, server group, or entire enterprise. This report uses a simple line graph with an independent y-axis for each metric. The report also graphs usage by day of week and by hour of day.
- **System Error** This report graphs the number of fatal errors, errors, and warnings logged by the specified server, server group, or entire enterprise during the reporting period. This report uses a simple stacked bar graph with an independent y-axis for each metric. The report also lists each of the fatal errors, errors, and warnings logged in ascending order by time.

File Type Associations

The File Type Associations node lets you display and manage all the file type associations for all the applications you have imported. (See Figure 3-47.)




FIGURE 3-47 Displaying file type associations for all imported applications.

Provider Policies

The Provider Policies node lets you specify a set of rules that are applied to users making connections to virtual applications. As connections come into the server group (provider), the server appends several rules (provider policy) to the connection. If the user does not specify a custom provider policy, the rules of the Default Provider are applied. (See Figure 3-48.)




FIGURE 3-48 The Provider Policies node.

Administrators

The Administrators node lets you view the group specified during installation as responsible for the administration of your App-V system. You can also add new groups to this node or remove groups you no longer need.

For more information on using the App-V Management Console to manage an App-V system, see the "Operations Guide for the Application Virtualization System" in the TechCenter Library on Microsoft TechNet at <http://technet.microsoft.com/en-us/library/cc843770.aspx>.

Using the Sequencer

The Sequencer is used to create a virtual application package for an application. The sequencer does this by monitoring and recording the installation and setup processes for an application. The result of sequencing an application is a set of files (.ico, .osd, .sft, .sprj, _manifest.xml, and optionally .msi) that contain all the necessary information for running the application within a virtual environment on the client.

To sequence an application, log on to your sequencing computer and select Microsoft Application Virtualization Sequencer from under Microsoft Application Virtualization in the Programs section of your Start menu. This launches the App-V Sequencer Console as shown in Figure 3-49.




FIGURE 3-49 The Application Virtualization Sequencer Console.

From the File menu of your Sequencer Console, select New Package. This launches the Sequencing Wizard and displays the first page of the wizard, which lets you specify a name and add an optional comment for your new package. (See Figure 3-50.)




FIGURE 3-50 Specifying a package name, and adding an optional comment.

Select the Show Advanced Monitoring Options check box if you want to display the Advanced Options page of the wizard. The Advanced Options page can be used to specify

the block size for your virtual application, which determines how the .sft file will be divided up when the package is streamed to client computers.

The third page of the wizard is called Monitor Installation. (See Figure 3-51.) Click the Begin Monitoring button to start monitoring the installation of your application.




FIGURE 3-51 Monitoring the installation of an application.

After the virtual environment has been loaded, begin installing your application. (See Figure 3-52.)



FIGURE 3-52 Launch setup for your application.

Choose the Custom installation option for your application. (See Figure 3-53.)




FIGURE 3-53 Select the Custom installation option.

Specify a location on your App-V virtual drive (Q: drive) where the application will be installed. (See Figure 3-54.) For more information about the Q: drive, see the sidebar titled "Direct from the Source: The Q: Drive" later in the chapter.



FIGURE 3-54 Install the application in a folder on the Q: drive.

After your application has been installed, click the Stop Monitoring button on the Sequencing Wizard.

The next page of the wizard lets you specify additional files to be added to the virtual file system. (See Figure 3-55.) You can also click the Reset button to clear any existing files from the virtual file system.




FIGURE 3-55 Adding files to the virtual file system.

On the next page of the Sequencing Wizard, you can configure shortcuts and file associations for the virtual application if needed. (See Figure 3-56.)




FIGURE 3-56 Configuring shortcuts and file associations for the package.

The next page of the wizard is named Launch Applications. (See Figure 3-57.) Select your application, and click the Launch All button to start the application to ensure that the virtual application package is properly optimized for streaming. Doing this is useful for several reasons:

- It allows you to configure how the application initially runs on client computers.
- It allows you to accept any license agreement for the application prior to making the application available to users.




FIGURE 3-57 The Launch Applications page of the wizard.

When you have finished with this step, click Next to sequence the application. The final page of the wizard, named Sequence Package, appears and displays the progress of the sequencing process. (See Figure 3-58.)




FIGURE 3-58 Sequencing progress.

When the sequencing process is complete, click Finish to close the wizard and return to the App-V Sequencer Console. The console now displays the results of the sequencing operation. The Properties tab shows basic information for the package, such as creation date, maximum block size, launch size, compression algorithm, and block size. (See Figure 3-59.)




FIGURE 3-59 Properties tab of the new package.

The Deployment tab lets you configure the streaming protocol, Streaming Server, port, application path, supported operating systems, and other options. (See Figure 3-60.) You can also select the Generate Microsoft Windows Installer (MSI) Package check box on this tab to generate an .msi file for the package.

The Change History tab displays the version history of the virtual application package.




FIGURE 3-60 Deployment tab of the package.

The Files tab displays the files the application copied, modified, or created and where those files reside. (See Figure 3-61.)

The Virtual Registry tab displays every registry setting that was created or modified and lets you view or change the settings and manually create or delete keys and values. (See Figure 3-62.)




FIGURE 3-61 Files tab of the package.




FIGURE 3-62 Virtual Registry tab of the package.

The Virtual File System tab displays the hierarchical directory of the files that comprise the package in common system folders.

The Virtual Services tab displays information about any Windows services that were detected and that are included as part of the sequence. (See Figure 3-63.)




FIGURE 3-63 Virtual Services tab of the package.

Finally, the OSD tab displays a hierarchical representation of the .xml descriptor file contents and lets you modify these values if needed to suit the application.

After you’ve reviewed your package settings and modified them as needed, select Save from the File menu and specify the name and location where your package should be saved. (See Figure 3-64.)




FIGURE 3-64 Saving the package.

Direct from the Source: The Q: Drive

As you may well be aware, one of the best practices for sequencing an application is to have the application install itself during the monitoring phase to an 8.3 directory on the Q: drive. What came up during launch of App-V is the need for deeper clarification on the reasons and even an extension to this particular best practice.

During sequencing, you enter the information in the package configuration phase that is common for all .osd files in the suite, and then you enter what is called the “installation phase.” It is during this phase that you Install, Test, and Configure the application or applications being virtualized. Also during this phase, the Sequencer monitors the actual installation of the application so that it can capture all of the files, registry settings, COM settings, embedded services, fonts, and so on that the application would normally lay down on a client during its native install.

The best practice states, “You should install the application to an 8.3 root directory on the Q: drive of the Sequencer. Each installer in the package should then have its own subdirectory under this 8.3 root. For example, if sequencing Microsoft Office 2007, during the installation phase the installer of Office will ask you where it should install into:

- Default: C:\Program Files\Microsoft Office\
- Best Practice: Q:\Off2k7.v1\Microsoft Office\

In the preceding example, we followed the best practice with an 8.3 root and each piece of the suite having its own subdirectory. If we were to add Microsoft Communicator to this SoftGrid suite, we would put it during the monitoring phase into the Q:\Off2k7.v1\Microsoft Office Communicator\ folder.

But why?

Good question. Let’s break this into pieces, shall we?

Question: “Why do I need to install it to Q:\?”

Answer: The App-V System has been built to try to remove any restrictions on where you install applications to during sequencing and which drive is used for the App-V File System on your deployed clients. This is done by searching for paths in the sequenced application that point to the installed path and replacing them with the variable *%SFT_MNT%*. However, some applications might have paths hard-coded in nonstandard configuration files that are not found by App-V. When that happens, the application will stream to the client, launch into the virtual environment and look to Q:\, yet the Mount Point of the client might actually be B:\ and the application will not work properly.

Essentially, the application is looking for what it was always told to look for, yet when it gets there its ideal does not exist.

Question: "Why do I need to install it to an 8.3?"

Answer: Most, if not all, applications will generate a backward-compatible 8.3 directory name even when they install into a long folder name. If you do not even remember the days when we were limited by our directory and file names to an 8.3 convention please, do me a favor, and just skip ahead. You're too young. Now an application such as Microsoft Office 2000 will install into a long folder of "Microsoft Office." When it auto-generates the 8.3, it follows the algorithm of first 6 characters, a tilde (~) and a number (1). So Office 2000 would be Micros~1.

Following proper sequencing practices, you would revert the Sequencer back to its clean state at the end of every successful sequence and start over fresh. If you were then to sequence Office 2003, it would install to a long folder of "Microsoft Office" again. And again, because the Sequencer is clean, it would autocreate its 8.3 as Micros~1.

So if you stream these two packages to a single client, there is a short name collision. By specifying an 8.3 name, you avoid the autogenerated short name and generate packages that won't have the short name collision.

Question: "Why should I put each component of the suite in its own subdirectory of the 8.3 root?"

Answer: Even if it is its own single application and no other applications will coexist in the suite, it should get its subdirectory under the 8.3 root. In the path of the application, it might have been "told" to always look under the relative path to that directory for its components. For example, Microsoft Office Help is coded to always look to a path relative to "Microsoft Office." In the preceding example, the Q:\Off2k7 takes the place of the C:\Program Files and the \Microsoft Office is still in its expected relative path.

Question: "When I click Stop Monitoring, the sequencer prompts me to select the directory that the application was installed to. Why do I have to choose the 8.3 root as the directory the application was installed to? Why wouldn't I select the actual subdirectory of the 8.3 root?"

Answer: Primarily because you did install the application to the 8.3 root, albeit you put it into a subdirectory of that root. What if you had installed both Office and Communicator to the same 8.3 root during the same sequence? By selecting the 8.3 root, you avoid the short path generation we mentioned previously.

Also, by selecting the 8.3 root here you are essentially saying, "OK. I installed the bulk of the assets under this root. But as with almost all applications, some files went to common folders such as C:\Windows\System32."

The Sequencer caught those common file locations. And it is at this point, as a result of you selecting the 8.3 root as the install folder, that the Sequencer will then create the Virtual File System (VFS) structure. As you may well know, the VFS folder structure and the Virtual Environment file get placed in whatever directory is selected at the end of monitoring. These are two components that are shared by all applications in this suite.

What you will end up with is the following:

Q:\0ff2k7

 \Microsoft Office

 \Microsoft Office Communicator

 \VFS

 \Osguard.cp

When teaching the App-V class, I used to use this analogy:

- Q:\ is the town in which you live.
- The 8.3 root is your house.
- The First Subdirectory (Microsoft Office) is the boys' bedroom.
- The second subdirectory (Microsoft Office Communicator) is the girls' bedroom.
- The VFS directory is the common dining room.
- The Osguard.cp file is the common rumpus room.
- Each child (application) gets his or her own bedroom where the bulk of that child's assets live, but they all share the common areas of the dining room and the rumpus room.

I hope this better clarifies the what and the why behind the 8.3 root.

—Sean Donahue, Senior Program Manager,
System Center Alliance. Microsoft Corporation



More Info For more information on using the App-V Sequencer to sequence applications, see “Operations Guide for the Application Virtualization System” in the Virtualization TechCenter Library on Microsoft TechNet at <http://technet.microsoft.com/en-us/library/cc843770.aspx>.

Working with App-V Clients

App-V Clients are small programs residing on desktop computers or terminal servers that communicate and authenticate with the App-V Server, receive the streamed application code, and enable the application to be executed for the user to use it. Administrators can configure the App-V Desktop Client and App-V Terminal Services Client and manage applications by using the Application Virtualization Client console.

Applications

The Applications node in this console can be used to manually manage virtual applications. (See Figure 3-65.) By selecting this node and then right-clicking on a virtual application, you can perform various tasks, such as the following:

- Load or unload an application from the cache.
- Clear an application from the console, which also removes the application's settings, shortcuts, and file type associations.
- Repair an application to remove any customizations and restore the application's default settings.
- Import an application into the cache.
- Lock or unlock an application. (A locked application cannot be removed from the cache to make room for new applications.)
- Delete an application, which means that the application will no longer be available to any users on that client. This operation also removes any shortcuts and file type associations for the application. The operation also removes the application from the cache unless another application refers to the selected application's file system cache data.
- Change the icon associated with an application.
- Manually add an application to the client by selecting New Application from the shortcut menu.
- Publish shortcuts to an application on the desktop, Quick Launch toolbar, Send To menu, Programs section of the Start menu, or some other location.




FIGURE 3-65 Managing applications using the Application Virtualization Client console.

You can also right-click on the Applications node itself to perform the following actions:

- Change the cache size and drive letter designation for the client.
- Change the log reporting level for the client.
- Modify user access permissions for the client.
- Configure the import search path where the client looks for .sft files when you try to import them.

File Type Associations

The File Type Associations node of the Application Virtualization Client console lets you add or delete a file type association for the application. When you add a new file association, you specify the file name extension, whether the new file type association should be global for all users, and which existing file type the new extension should be associated with.

Publishing Servers

The Publishing Servers node lets you set up new publishing servers and perform related tasks on the client. (See Figure 3-66.)




FIGURE 3-66 Managing publishing servers using the Application Virtualization Client console.

For example, to set up a new publishing server, first add the server by right-clicking on the Publishing Servers node and selecting New Server. Follow the steps of the wizard to specify a display name and select a server type. The supported server types are these:

- Application Virtualization Server—Uses RTSP as its streaming protocol
- Enhanced Security Application Virtualization Server—Uses RTSPS as its streaming protocol
- Standard HTTP Server—Uses HTTP as its streaming protocol
- Enhanced Security HTTP Server—Uses HTTPS as its streaming protocol

After you have added the publishing server, right-click on it and select Properties to display its Properties dialog box. You can use the tabs on this dialog box to configure the following:

- Server name and type
- Host name and port
- Whether to refresh publishing on user login
- Publishing refresh rate

Managing App-V Clients from the Command Line

You can also manage App-V Clients from the command line by using the SFTMIME command. For example, to add a virtual application package for all users of the computer, type **SFTMIME ADD PACKAGE:<name> /MANIFEST <path> /GLOBAL** at the command prompt.

For more information on using the SFTMIME command to manage App-V Clients, see the SFTMIME Command Reference on Microsoft TechNet at <http://technet.microsoft.com/en-us/library/cc817090.aspx>.

Direct from the Source: App-V Troubleshooting

The following are some tips for troubleshooting different aspects of an App-V system.

Troubleshooting Publishing and Streaming Issues

A common problem among those who are new to App-V is understanding how application publishing works with App-V and what paths, directories, and protocols are involved when using classic streaming-based publishing when a Management Server is involved. It's somewhat unfortunate that the publishing process does not take full advantage of the existing delivery channel—RTSP or its secure version RTSPS—that already exists between the server and client components, but that's just the way it is as of now.

Sometimes confusion also stems from the purpose of the various file formats used by App-V not being understood thoroughly. Each App-V virtual application package consists of a package data file (SFT) and a bunch of help files (an OSD file for describing one individual application from the package data file, an ICO file for creating the short-cut, and so on), some of which are delivered using different methods than the others.

When troubleshooting application publishing and streaming problems, you should double-check all the components involved. Following are a few things to check for that affect the delivery and visibility of the applications.

One thing to check is the content directory path, which is the file system path you need to select during the installation of every App-V server. This directory represents a logical starting point for the service that listens on the network for the clients' requests. The content directory path, as defined during installation (for example, C:\content\), is used solely by the App-V service on the server when it tries to find physical SFT files for streaming. If you ever see or define this path in any of the management screens, that's an error because the server itself handles SFT delivery through RTSP or RTSPS channels. When using the classic application delivery with App-V, SFT files cannot be published through UNC paths or HTTP URLs (except, of course, if you are using that new-fangled HTTP-streaming introduced in 4.5, but that's a whole different scenario).

Also check the package relative path. This path is the one you see under the Packages node in the Management Console for each version of each package (for example, myapp\package.sft). This path usually resolves correctly automatically when you have imported App-V packages, but sometimes it might be incorrect if the person who packaged the virtual application wasn't careful. The import procedure reads this relative path from the first OSD file for each package it imports. This path is the second piece of information that the App-V Management Server uses to find a package a client requests when it tries to stream something in. Just as the content directory path was the logical starting point, this relative path is the logical conclusion of that path.

Together, they form a full, valid file system path to SFT file (for example, C:\content\ + myapp\package.sft = C:\content\myapp\package.sft).

Completely unrelated path settings from the ones just described are the OSD and ICO paths. These paths, as defined in the application publishing records (the screen you get when running Import in the Management Console), are delivered by the Management Server to the App-V Client when it issues a refresh operation against the server. Refresh, as the name implies, refreshes the client's list of applications that are known to the user under which the refresh is running.

Part of this procedure, in addition to getting the list of applications available for the user, is to get paths to the OSD and ICO files. The content of those files are not delivered as part of the refresh data itself, but rather as a separate reference. So you have a choice of which delivery channel you want to use: a file system path (effectively a UNC path) or an HTTP URL. And here's what causes so much confusion: because you can actually use a file system path as OSD and ICO paths, you are tricked into thinking that those paths can refer to content directory using server's local path (for example, C:\content\myapp\myapp.osd) or in the way relative package paths are used for SFT files (for example, myapp\myapp.osd). The issue, however, is that those paths are not interpreted by the server but by the client when it gets the paths as part of a refresh operation and the client doesn't have C:\content\ to get OSDs and ICOs from! To make matters worse, the Management Console defaults to using whatever path you used to browse to SPRJ (the project file that rules all other files) as OSD and ICO paths, unless you happened to set something called Default Content Path before the import.

Troubleshooting Virtual Service Issues

Not many packages contain virtual services because services are normally associated more with server-side software. However, sometimes virtual services do exist in an App-V package. Virtual services are very straightforward (the App-V Sequencer picks up newly created services during packaging) and people tend not to touch them. There are some catches to using them, and being aware of those oddities might help in the longer run.

The first issue with the virtual services is what the startup type of Automatic causes to happen on the client side. Because virtual services do not exist on the client system's real service list, nothing starts them automatically when you start your machine. As a result, the App-V Client must start virtual services just before the main application starts from the package. Sometimes these services will start fast enough, but there are some type of services and environmental conditions that cause them to take some time. And this additional time, much to the dismay of the user starting the application, results in delaying the startup of the actual application. And if there are many such

services in the package, the delay can be even longer, causing the user to wonder if the start of the application failed for some reason.

The solution? Either set virtual services to start manually if possible or disable virtual services totally if they are not needed. In the case where such services were present in the package, the application launch time exceeded one minute solely because of slow-starting services, and changing to a manual start caused the launch time to drop to a reasonable level.

Another issue that might arise with virtual services is a conflict with the locally present identical service. This is not uncommon with some software licensing components that install themselves as services. If you happen to use multiple products that use that same licensing service, each one might contain a virtual copy if they are sequenced separately. Typically, what happens is that the copy that starts first (either the one on the virtual machine or the one on the local machine) will run perfectly, but any subsequent copies will not. The service will either terminate on start or consume lots of CPU time when trying to do something that overlaps with the existing instance.

The solution? There's no easy way out of this problem, but something worth trying is to disable all but one copy of the service. Usually, the locally installed service is preferred, as the ordering of which virtual application (and thus the virtual service in it) will launch first is unknown.

—Kalle Saunamäki, MVP

Additional Resources

The resources listed in this section were either referred to in the chapter or provide additional information about concepts and products discussed in this chapter.

Resources for Windows Virtual PC and Windows XP Mode

General information about Windows Virtual PC and Windows XP Mode, including how to obtain them, can be found at <http://www.microsoft.com/windows/virtual-pc/>.

The Windows Virtual PC Evaluation Guide, which shows you how to set up Windows Virtual PC and Windows XP Mode, can be found in the TechNet Library at <http://technet.microsoft.com/en-us/library/dd744684.aspx>.

Release notes for Windows Virtual PC and Windows XP Mode can be found in the TechNet Library at <http://technet.microsoft.com/en-us/library/ee681620.aspx>.

The Windows Virtual PC and Windows XP Mode Setup and Installation Guide can be found in the TechNet Library at <http://technet.microsoft.com/en-us/library/ee681616.aspx>.

Windows Virtual PC Help can be found in the TechNet Library at <http://technet.microsoft.com/en-us/library/ee449411.aspx>.

Ben Armstrong's blog, called the "Virtual PC Guy's Blog," has lots of useful tips and tricks on deploying, managing, and using Windows Virtual PC and Windows XP Mode. You can find it at http://blogs.msdn.com/virtual_pc_guy/archive/tags/Windows+Virtual+PC/default.aspx and http://blogs.msdn.com/virtual_pc_guy/archive/tags/Windows+XP+Mode/default.aspx.

To ask questions about Windows Virtual PC or Windows XP Mode, or to help others with their questions, use the Windows 7 Virtualization forum on Microsoft TechNet at <http://social.technet.microsoft.com/Forums/en-US/w7itprovirt/threads>.

Resources for MED-V

For a description of what's available in MDOP 2009, see <http://www.microsoft.com/windows/enterprise/products/mdop/default.aspx>. For additional information, see the Official MDOP Blog at <http://blogs.technet.com/mdop/default.aspx>.

For a general description of MED-V, see <http://www.microsoft.com/windows/enterprise/products/mdop/med-v.aspx>.

The "MED-V 1.0 Architecture Overview" white paper is available from the Microsoft Download Center at <http://download.microsoft.com/download/A/A/F/AAF7988A-94F0-483A-9610-E0E6AB51DA79/MEDV%20Architecture%20June09.pdf>.

Release notes for MED-V 1.0 can be found on TechNet at <http://technet.microsoft.com/en-us/library/ee348918.aspx>.

"MED-V Planning, Operations and Deployment Guide" can be found at <http://technet.microsoft.com/en-us/library/ee348978.aspx>.

The MED-V Team Blog can be found at <http://blogs.technet.com/medv/default.aspx>.

At the time of writing this chapter, there is no TechNet Forum dedicated to discussing MED-V.

Resources for App-V

For a general description of App-V, see <http://www.microsoft.com/windows/enterprise/products/mdop/app-v.aspx>.

The launching page for all App-V product documentation and related information is <http://www.microsoft.com/systemcenter/appv/default.mspx>.

For detailed technical information about App-V, see the Application Virtualization TechCenter on Microsoft TechNet at <http://technet.microsoft.com/en-us/appvirtualization/default.aspx>.

"Planning and Deployment Guide for the Application Virtualization System" can be found at <http://technet.microsoft.com/en-us/library/cc843778.aspx>.

Also, be sure to review Microsoft Application Virtualization Management System Release Notes at <http://technet.microsoft.com/en-us/library/cc817171.aspx> prior to deploying App-V within your organization.

If you plan on upgrading your existing SoftGrid 4.2 system to App-V, be sure to review the "Upgrading to Microsoft Application Virtualization 4.5 Frequently Asked Questions" article found at <http://technet.microsoft.com/en-us/appvirtualization/cc664494.aspx>.

"Operations Guide for the Application Virtualization System" can be found at <http://technet.microsoft.com/en-us/library/cc843770.aspx>.

Also, be sure to read the various white papers available on the Application Virtualization 4.5 Documentation section of the Application Virtualization TechCenter at <http://technet.microsoft.com/en-us/appvirtualization/cc843994.aspx>.

To keep abreast of the latest developments and find tips about App-V, subscribe to the RSS feed for the App-V Team Blog at <http://blogs.technet.com/softgrid/default.aspx>.

To ask questions about App-V or SoftGrid, or to help others with their questions, use the Microsoft Application Virtualization forums on Microsoft TechNet at <http://social.technet.microsoft.com/forums/en-US/category/appvirtualization/>.

Chapter 4

Remote Desktop Virtualization

As described in the previous chapter, desktop virtualization can be either local or remote. In local desktop virtualization, the virtual environment is running on the user's computer (the host). Microsoft provides several virtualization technologies for local desktop virtualization, including Windows Virtual PC and the Windows XP Mode environment, Microsoft Enterprise Desktop Virtualization (MED-V), and Microsoft Application Virtualization (App-V).

Remote desktop virtualization on the other hand is different. In this approach, the virtual environment runs on a server, typically on a Windows Server host. This chapter deals with remote desktop virtualization and covers the following Microsoft technologies and solutions:

- **Remote Desktop Services** With the release of Microsoft Windows Server 2008 R2, the Terminal Services role has now been renamed Remote Desktop Services and has been enhanced with the capability of delivering virtual machine desktops to users using Remote Desktop Protocol (RDP). Remote Desktop Services is a key component of all other Microsoft remote desktop virtualization solutions.
- **Microsoft Application Virtualization for Remote Desktop Services (App-V for RDS)** Formerly known as Application Virtualization for Terminal Services (App-V for TS), App-V for RDS lets you transform applications into centrally managed virtual services and deliver them to users using RDP. This can help you reduce the cost of application deployment, eliminate application conflicts and reboots, simplify your base image footprint to expedite PC provisioning, and increase user productivity.
- **Microsoft Virtual Desktop Infrastructure (VDI)** Microsoft VDI is an architectural model composed of a number of components, including Hyper-V, Remote Desktop Services, Microsoft Desktop Optimization Pack (MDOP), and Microsoft System Center products. Microsoft VDI enables entire desktop operating systems such as Windows 7 Enterprise to run on a hypervisor server located in a datacenter and be delivered to users as virtual desktops using RDP. By deploying a Microsoft VDI infrastructure, users can access either their own personal virtual desktop, which they can customize as desired, or a shared pool of identically configured virtual desktops.

This chapter begins by examining the benefits, usage scenarios, and availability of each remote desktop virtualization technology or solution. The chapter then describes each technology or solution in detail by describing how it works and what its capabilities are.



Note At the time of this writing, the final name for App-V for RDS has not yet been finalized.

Examining the Benefits of Remote Desktop Virtualization

Which remote desktop virtualization technologies and solutions you deploy in your organization depends on what benefits they can provide to your business. Table 4-1 compares the advantages and disadvantages of deploying the following remote desktop virtualization technologies and solutions:

- Remote Desktop Services alone
- Remote Desktop Services together with App-V for RDS
- Microsoft VDI using shared virtual desktop pools
- Microsoft VDI using personal virtual desktops

TABLE 4-1 Comparing the Benefits of Different Remote Desktop Virtualization Scenarios

Comparison of benefits	Remote Desktop Services	Remote Desktop Services with App-V for RDS	Microsoft VDI using shared virtual desktop pools	Microsoft VDI using personal virtual desktops
Cost of image management	Low	Low	Medium	High
Ease of administration	Least complex	More complex	More complex	Most complex
Amount of resources needed	Least	More	More	Most
Compatibility with legacy applications	Some	Better	Better	Best

Examining Usage Scenarios for Remote Desktop Virtualization

Each remote desktop virtualization technology or solution can be implemented in various ways to bring benefit to an organization. The sections that follow describe some of the scenarios in which you might deploy these different technologies and solutions advantageously for your business.

Usage Scenarios for Remote Desktop Services

Common usage scenarios for Remote Desktop Services include

- **Branch Office** Deploying line-of-business software applications for users at multiple branch offices can be time-consuming and costly. Remote Desktop Services can help

reduce this cost by allowing you to run such software on a Remote Desktop Session Host server located at a central headquarters. Employees at branch offices can then access these applications on an as-needed basis via remote desktops or as RemoteApp programs, even over low-bandwidth connections.

- **Controlled Partner Access or Outsourcing** The burden and complexity of deploying and maintaining line-of-business software applications on computers belonging to business partners or outsourcing firms can be significantly reduced by using Remote Desktop Services. Partners and outsourced workers are able to access the applications they need to do their job without having to obtain and install these applications on their own computers. This can make for a smoother business relationship between your organization and your business partners or outsourcing firms, plus it provides added security by limiting the access these businesses will need to resources on your corporate network.
- **Easing the Burden of Regulatory Compliance** By running applications and storing application data on centrally located servers, Remote Desktop Services helps reduce the risk of accidental data loss caused, for example, by the accidental loss of a laptop. The zero-application footprint and data delivery model employed by Remote Desktop Services also helps to ensure that as little data as possible resides on the client computing device. And if you use Remote Desktop Gateway together with RemoteApp, your employees, partners, and customers no longer require full access to your corporate network and computers. Instead, you can limit the applications they can access, even to using a single application, if needed.
- **Merger Integration** During a corporate merger, companies typically need to use consistent line-of-business (LOB) applications on a variety of Windows operating system versions and configurations. Rather than going through the effort and incurring the high cost of deploying all your LOB applications to all the computers in the merged company, these applications can simply be installed on a Remote Desktop Session Host server and made available to those who need it via RemoteApp. This can be especially useful when an application is difficult to maintain, cannot be deployed easily, or has other management issues.
- **Mobile Workers** Organizations that support employees who work from home or work while traveling can implement a Remote Desktop Services solution to help enable employee productivity anytime and anywhere. Remote Desktop Services can also increase effective collaboration between users without compromising security, and it can offer secure access to applications over low-bandwidth connections without requiring those applications to be installed on client computers.
- **Task Workers** Organizations that have structured task workers—such as call center employees, factory floor workers, or both—can use Remote Desktop Services to provide such employees with a more productive user experience. Typically, task workers like these do not need to access many applications to complete the tasks they have

been assigned, and RemoteApp together with Remote Desktop Web Access provide an easy way for them to access the applications they need, when they need them. A similar user experience can be provided even if the user's computer is an older desktop computer running an earlier version of Windows, a non-PC desktop computer, or a mobile computing device. Deploying applications for task workers in this way can help extend the reach of Windows-based applications within an enterprise and is a valuable, cost-effective way to deliver the right business tools to the people who need them.

Usage Scenarios for App-V for RDS

Because App-V for RDS is implemented within a Remote Desktop Services environment, the usage scenarios for this technology are similar to those just described. However, App-V for RDS also provides a number of additional benefits to these scenarios, including being able to

- **Consolidate servers and end server siloing** App-V for RDS eliminates the need for server silos, and thus it significantly improves server utilization and increases server farm return on investment (ROI) for organizations. This is because App-V for RDS allows any application to run alongside any other application, including applications that normally conflict, multiple versions of the same application, and many applications that cannot run RD Session Host servers without App-V.
- **Accelerate application deployment** Because applications deployed using App-V for RDS need to be packaged only a single time for deployment to both desktops and RD Session Host servers, application deployment is streamlined.
- **End application conflicts and regression testing** Because App-V for RDS eliminates the need to permanently install applications on RD Session Host servers, the need to perform lengthy regression testing of applications is significantly reduced, which also helps to speed the deployment of applications within your Remote Desktop Services environment.
- **Reduce application deployment risk** Because using App-V for RDS applications can be deployed and updated on demand to users without the need to reboot servers or log off users from their sessions, installing applications on RD Session Host servers is easier and less can go wrong.
- **Simplify profile management** Because App-V for RDS allows application settings and data to be stored in a single network location, each user's application settings can be available regardless of which RD Session Host server the user connects to—without the need to implement roaming profiles. App-V for RDS also makes mandatory profiles viable within a Remote Desktop Services environment because operating system settings can remain locked within the mandatory profile while per-application settings can still be modified by users.

Usage Scenarios for Microsoft VDI

Usage scenarios for Microsoft VDI also have many similarities with those for Remote Desktop Services. This is because Remote Desktop Services is one of the key components of a Microsoft VDI solution. However, Microsoft VDI also adds capabilities that enable new kinds of usage scenarios that include the ability for users to

- Centralize all desktop deployment, updating, and management in the datacenter
- Provision virtual desktops dynamically on demand from gold images
- Work from home or at offsite contractor locations
- Personalize their own virtual desktop by configuring it and installing applications
- Access shared pools of identically configured virtual desktops
- Hot-desk between different desktop PCs

Availability of Remote Desktop Virtualization Technologies

Now that we've examined the benefits and usage scenarios for each remote desktop virtualization technology, let's look at how you can obtain them from Microsoft.

Availability of Remote Desktop Services

Remote Desktop Services is included as an installable server role in the Standard, Enterprise, and Datacenter editions of Windows Server 2008 R2. No additional download is needed.

Availability of App-V for RDS

Obtaining App-V for RDS requires that you have

- A valid licensed copy of Microsoft Windows Server 2008 or Windows Server 2008 R2
- A valid Windows Server 2008 Terminal Services (TS) client access license (CAL) or Windows Server 2008 Remote Desktops Services (RDS) CAL

If you meet these conditions, you can obtain App-V for TS/RDS from

- The Microsoft Volume License Services (MVLS) Web site at <http://www.microsoft.com/licensing/Default.aspx>. Note that after December 6, 2009, the MVLS site will automatically redirect to the new Volume Licensing Service Center (VLSC).
- By searching for "Microsoft Application Virtualization for Terminal Services / Remote Desktop Services" on the Microsoft Download Center, or by downloading directly from

<http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=0890d6cd-0d3b-4c9d-b208-231c65d3e55a>. The App-V for RDS client is a free download, but the site requires that you enter a valid 20-digit Windows Server 2008 Remote Desktop Session Host Product Identification Key or Windows Server 2008 R2 Remote Desktop Services Product Identification Key.

Note that at the time of writing this chapter, the latest version of this product is called Application Virtualization 4.5 for Remote Desktop Services Service Pack 1. (The corresponding App-V desktop client is called Application Virtualization Hosting for Desktops Version 4.5 Service Pack 1.) It is currently available only for the x86 architecture, which means that App-V for RDS SP1 can be installed only on Windows Server 2008 x86 terminal servers, and not on Windows Server 2008 R2 RD Session Hosts (because R2 is x64 only). However, an x64 version of App-V for RDS SP1 is expected to be released in the first half of 2010 when App-V 4.6 is released.

Availability of Microsoft VDI

Microsoft VDI suites are available only to volume licensed customers. You can purchase the different components of the server and management infrastructure required to deploy a Microsoft VDI infrastructure through a single licensing vehicle. This vehicle has two SKUs: VDI Standard suite and VDI Premium suite. In addition, to use Windows 7, Windows Vista, or Windows XP within a Microsoft VDI infrastructure, you need to purchase a Windows Virtual Enterprise Centralized Desktop (VECD) license. These licenses are available to volume licensed customers under Enterprise Agreement, Select, Open Value, and Campus arrangements and to Software Assurance (SA) customers.

Understanding Remote Desktop Services

Remote Desktop Services provides a server-based execution environment that enables users to run Windows-based programs as if they were locally installed on the users' computers, when in fact the programs are centrally installed on a server. Remote Desktop Services also enables users to access entire session-based desktops running on an RD Session Host server or virtual machine desktops running on an RD Virtualization Host server.

Remote Desktop Services works by transmitting the key presses and mouse clicks on a user's computer over the network to a server (either an RD Session Host or RD Virtualization Host), where they are accepted as input for actions performed using the remote program, remote desktop, or remote virtual machine being accessed. The server then transmits the result over the network to the user's computer, where it is displayed on the user's monitor. Keyboard and mouse activity and display information is transmitted between the terminal server and the client by using Remote Desktop Protocol 7.0 (RDP 7.0), which runs over a Transmission Control Protocol/Internet Protocol (TCP/IP) network via TCP port 3389 as shown

in Figure 4-1. The software running on the client side that makes such communications possible is called Remote Desktop Connection (RDC).




FIGURE 4-1 How Remote Desktop Services works.

In Windows Server 2008 and earlier, this server role was called Terminal Services; with the release of Windows Server 2008 R2, however, this server role is now called Remote Desktop Services and new capabilities have been added. Table 4-2 summarizes the changes to the names of the role services, while Table 4-3 summarizes the changes to the names of the tools for managing this server role. The sections that follow then describe each of the role services included in Remote Desktop Services and examines the new and enhanced capabilities that have been added to the role service.

TABLE 4-2 How Terminal Services Role Services in Windows Server 2008 Have Been Renamed in Windows Server 2008 R2

Terminal Services role services in Windows Server 2008	Remote Desktop Services role services in Windows Server 2008 R2
Terminal Server	Remote Desktop Session Host (RD Session Host)
Terminal Services Gateway (TS Gateway)	Remote Desktop Gateway (RD Gateway)
Terminal Services Licensing (TS Licensing)	Remote Desktop Licensing (RD Licensing)
Terminal Services Session Broker (TS Session Broker)	Remote Desktop Connection Broker (RD Connection Broker)
Terminal Services Web Access (TS Web Access)	Remote Desktop Web Access (RD Web Access)
N/A	Remote Desktop Virtualization Host (RD Virtualization Host)

TABLE 4-3 How Terminal Services Management Tools in Windows Server 2008 Have Been Renamed in Windows Server 2008 R2

Terminal Services management tools in Windows Server 2008	Remote Desktop Services management tools in Windows Server 2008 R2
Terminal Services Manager	Remote Desktop Services Manager
Terminal Services Configuration	Remote Desktop Session Host Configuration
TS Gateway Manager	Remote Desktop Gateway Manager
TS Licensing Manager	Remote Desktop Licensing Manager
TS RemoteApp Manager	RemoteApp Manager
N/A	Remote Desktop Services PowerShell provider

Understanding Remote Desktop Connection Client Experience Improvements

The Remote Desktop Connection (RDC) client has been enhanced in Windows 7 and Windows Server 2008 R2 to make your experience of using a remote desktop closer to the experience of using the local desktop on your computer. The enhancements found in the RDC 7.0 client include the following:

- **Audio and video playback redirection** Audio and video content played back using Windows Media Player can now be redirected from the RD Session Host server to the client computer in its original format and rendered using the client computer's resources. Other types of multimedia content, such as Microsoft Silverlight content, are still rendered on the server.
- **Audio recording redirection** Audio recording devices, such as microphones, can now be redirected from the client computer to the remote desktop session—an improvement that can be useful for organizations using voice chat or Windows Speech Recognition.
- **Desktop composition** Windows Aero is now supported within an RD Session Host session when using the RDC 7.0 client.
- **Language bar redirection** You can now use the language bar on the client computer to control the language settings within your RemoteApp programs.
- **Multiple-monitor support** Remote desktop sessions can now support up to 16 monitors using any monitor configuration supported on the client computer. Note that desktop composition is not supported on an RD Session Host session when using multiple monitors.

In addition to the preceding client experience improvements, RDC 7.0 also includes support for other new features:

- Web Single Sign-On (SSO) and Web forms–based authentication
- Access to personal virtual desktops by using RD Connection Broker
- Access to virtual desktop pools by using RD Connection Broker
- Status and disconnect system tray icon
- RD Gateway–based device redirection enforcement
- RD Gateway system and logon messages
- RD Gateway background authorization and authentication

- RD Gateway idle and session timeouts
- Network Access Protection (NAP) remediation with RD Gateway

Many of these additional enhancements are described in more detail in later sections of this chapter.

RDC 7.0 is also available as a free download for Windows Vista SP1 or later and for Windows XP SP3. For more information and to obtain RDC 7.0 for these earlier versions of Windows, see Microsoft Knowledge Base article KB969084 at <http://support.microsoft.com/kb/969084>. However, the following features of RDC 7.0 are available only when connecting from Windows 7 to Windows Server 2008 R2:

- Support for Aero Glass and language-bar docking
- Support for starting applications and desktops by using RemoteApp and Desktop Connection
- Using the remote application task scheduler to automatically start remote applications on the Remote Desktop client that might be required by the user

Understanding the Remote Desktop Session Host

The Remote Desktop Session Host (RD Session Host) role service of Windows Server 2008 R2 was formerly called Terminal Server in Windows Server 2008. Installing the RD Session Host role service enables the server to deliver Windows-based programs (called RemoteApp programs) or entire Windows desktops to users over the network. New features of the RD Session Host role service introduced in Windows Server 2008 R2 include the following:

- **Configure Client Experience page** Adds a new wizard page to the Add Roles Wizard when installing the RD Session Host role service of the Remote Desktop Services role. This new wizard page lets you enable the following advanced experiences for RD Session Host session users:
 - ❑ **Audio and video playback redirection** Lets users redirect audio and video output from their computer to an RD Session Host session.
 - ❑ **Audio recording redirection** Lets users redirect the output of an audio recording device, such as a microphone, from their computer to an RD Session Host session.
 - ❑ **Desktop composition** Provides Windows Aero user interface elements within an RD Session Host session.
- **Per-user RemoteApp filtering** Lets you filter the list of RemoteApp programs available to a user account when logged on using RD Web Access.

- **Fair-share CPU scheduling** Dynamically distributes processor time across RD Session Host sessions based on the number of active sessions and the load on those sessions using the kernel-level scheduling mechanism of Windows Server 2008 R2. The result is that one user of an RD Session Host server will not affect the performance of another user's session even when the RD Session Host server is under heavy load.
- **Windows Installer RDS compatibility** Allows per-user application installations to be queued by the RD Session Host server and then handled by the Windows Installer. This enables you to install a program on the RD Session Host server in the same way you would install the program on a local desktop, provided you install the program for all users of the computer and provided all components of the application are installed locally on the RD Session Host server.
- **Roaming user profile cache management** Lets you limit the size of the overall profile cache for users of your RD Session Host server. Then if the size of the profile cache grows larger than the configured size, the least recently used profiles are deleted until the cache size goes below the quota. The size of the cache can be configured using the following Group Policy setting:

Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Profiles\Limit the size of the entire roaming user profile cache
- **Remote Desktop IP Virtualization** Lets IP addresses be assigned to Remote Desktop connections on either a per-session or per-program basis. Assigning IP addresses for multiple programs causes them to share a session IP address. To configure this feature, use the new RD IP Virtualization tab in the Remote Desktop Session Host Configuration snap-in. If your RD Session Host server has more than one network adapter, you must choose one network adapter for Remote Desktop IP Virtualization.

Installing an RD Session Host Server


You can install the RD Session Host role service on the Standard, Enterprise, or Datacenter edition of Windows Server 2008 R2, with the Standard edition limited to 250 Remote Desktop Services connections. You can use any of the following methods to install this role service:

- Launching the Add Roles Wizard from either Server Manager (ServerManager.msc) or the Initial Configuration Tasks window (Oobe.exe)
- Using ServerManager.cmd from the command line


- During the installation of Windows Server 2008 R2 by using Microsoft Deployment Toolkit 2010 (MDT 2010), System Center Configuration Manager 2007 SP2 (SCCM 2007 SP2), or both
- Using the Deployment Image Servicing and Management (DISM.exe) tool included in the Windows Automated Installation Kit 2.0
- Using Windows PowerShell

Follow these steps to add the RD Session Host role service to you server by using the Add Roles Wizard:


1. Launch the wizard from either Server Manager or the Out-Of-Box Experience (OOBE) window. (Skip the Before You Begin page if it is displayed.)
2. Select the Remote Desktop Services role:




3. Select Remote Desktop Session Host from the list of available role services:



4. The next page of the wizard warns you that any user applications already installed on the server might need to be uninstalled before they can work properly in a multi-user environment (for more information, see the sidebar titled "Direct from the Source: Best Practices for Installing Applications on RD Session Host Servers" later in this chapter):



5. The next few pages of the wizard allow you to configure security (discussed later in the section titled “Configuring RD Session Host Security”) and licensing (discussed later in the section titled “Understanding Remote Desktop Licensing”) for your RD Session Host server.
6. The next-to-last page of the wizard allows you to configure the client experience. For example, by selecting Desktop Composition you can enable session-based desktops to be delivered to users from this RD Session Host server to display Windows Aero features:



Note that enabling these client experience features requires additional network bandwidth and system resources on the server.

7. The final page of the wizard summarizes the selections you have made so that you can review them before installing the role service.

Configuring RD Session Host Security

Security can be configured for your RD Session Host server either during the installation of the role or afterward. You must take the following additional steps to provide secure access to your RD Session Host server:

- Specify an authentication method for your RD Session Host server.
- Specify which groups of users are allowed access to your RD Session Host server.

The sections that follow provide more details concerning these steps. In addition, the sidebar titled “Direct from the Source: Using SSL for Remote Desktop Services Connections” provides additional information on how you can protect your RD Session Host server.

Configuring RD Session Host Authentication Support for Network Level Authentication (NLA) was first introduced in the Terminal Services role of Windows Server 2008. NLA is an authentication method that completes the user authentication process before an RD Session Host session is established and the logon screen is displayed. NLA provides greater security and requires fewer initial resources during the authentication process. As Figure 4-2 shows, the option to require NLA can be configured during the installation of the RD Session Host role service.




FIGURE 4-2 Requiring Network Level Authentication for remote connections to an RD Session Host server.

To use NLA for authenticating remote connection attempts to an RD Session Host server, the computer must be using RDC 6.0 or later client software. This means that the computer must be running Windows 7, Windows Vista SP1 or later, Windows XP SP3 or later, Windows Server 2008, or Windows Server 2008 R2.

Configuring RD Session Host Access Before any users can establish RD Session Host sessions with your RD Session Host server, you must add their appropriate security groups to the Remote Desktop Users group on your RD Session Host server. You can do this either during the RD Session Host role service installation process or afterward. By default, the local

Administrators group on your RD Session Host server has such access, which by default also provides access to members of the Domain Admins group. To allow domain users the ability to establish RD Session Host sessions with your RD Session Host server, add the Domain Users global group as shown in Figure 4-3.




FIGURE 4-3 Allowing users access to an RD Session Host server.

Direct from the Source: Using SSL for Remote Desktop Services Connections

By default, RD Session Host sessions use native RDP encryption. However, RDP does not provide authentication to verify the identity of an RD Session Host server. You can enhance the security of RD Session Host sessions by using Secure Sockets Layer (SSL) Transport Layer Security (TLS 1.0) for server authentication and to encrypt RD Session Host communications. The RD Session Host server and the client computer must be correctly configured for TLS to provide enhanced security. The three available security layers are shown in Table 4-4.

TABLE 4-4 RDP Security Layers

Security Layer	Description
SSL (TLS 1.0)	SSL (TLS 1.0) will be used for server authentication and for encrypting all data transferred between the server and the client.
Negotiate	The most secure layer that is supported by the client will be used. If supported, SSL (TLS 1.0) will be used. If the client does not support SSL (TLS 1.0), the RDP Security Layer will be used. This is the default setting.
RDP Security Layer	Communication between the server and the client will use native RDP encryption. If you select RDP Security Layer, you cannot use Network Level Authentication.

A certificate is needed to authenticate an RD Session Host server when SSL (TLS 1.0) is used to secure communication between a client and an RD Session Host server during RDP connections. You can select a certificate that you have already installed on the RD Session Host server, or you can use the default self-signed certificate. Figure 4-4 shows how to enable SSL for Remote Desktop connections using the RDP-Tcp Properties dialog box, which is accessed from the Remote Desktop Session Host Configuration snap-in.




FIGURE 4-4 Configuring the security layer settings for your RD Session Host server.

For Remote Desktop connections, data encryption protects data by encrypting it on the communications link between the client and the server. Encryption protects against the risk of interception of the client/server communication.

By default, Remote Desktop connections are encrypted at the highest level of security available (128-bit). However, some older versions of the Remote Desktop Connection client application do not support this high level of encryption. If a high level of encryption is needed to support legacy clients, the encryption level of the connection can be configured to send and receive data at the highest encryption level supported by the client. There are four levels of encryption available, as shown in Table 4-5.

TABLE 4-5 RDP Encryption Levels

Level of Encryption	Description
Low	Data sent from the client to the server is encrypted using 56-bit encryption. Data sent from the server to the client is not encrypted.
Client Compatible	Encrypts client/server communication at the maximum key strength supported by the client. Use this level when the terminal server is running in an environment containing mixed or legacy clients. This is the default encryption level.
High	Encrypts client/server communication using 128-bit encryption. Use this level when the clients accessing the terminal server also support 128-bit encryption. When encryption is set at this level, clients that do not support this level of encryption will not be able to connect.
FIPS Compliant	All client/server communication is encrypted and decrypted with the Federal Information Processing Standards (FIPS) encryption algorithms. FIPS 140-1 (1994) and its successor, FIPS 140-2 (2001), describe U.S. government requirements for encryption.

Figure 4-5 shows how to configure the encryption level using the RDP-Tcp Properties dialog box, which is accessed from the Remote Desktop Session Host Configuration snap-in.

RD Session Host authentication and encryption settings also can be configured by applying the following Group Policy settings:

- Set Client Connection Encryption Level
- Require Use Of Specific Security Layer For Remote (RDP) Connections
- Server Authentication Certificate Template
- Require User Authentication For Remote Connections By Using Network Level Authentication




FIGURE 4-5 Configuring the encryption layer settings for your RD Session Host server.

These Group Policy settings are located in the following container:

Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Security

FIPS can be specified as the encryption level by applying the System Cryptography: Use FIPS Compliant Algorithms For Encryption, Hashing And Signing Group Policy setting located in the following container:

Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options

—CSS Global Technical Readiness (GTR) team

Installing Applications on an RD Session Host Server

Applications developed for end users of your system must be installed properly on an RD Session Host server so that they can be accessed remotely. Before you install an application on an RD Session Host server, make sure you understand any compatibility issues associated with the application when it is running in a Remote Desktop Services environment.

Make sure you install the RD Session Host role service on your server before you install any applications that users will need to run within their remote sessions or as RemoteApp programs. If you install the RD Session Host role service after you have installed your applications, the applications might not function correctly in a multi-user environment.

To install an end-user application on an RD Session Host server, the RD Session Host server must first be switched into a special install mode called RD-Install to ensure that the application will be able to run in a multi-user environment. After your applications have been installed on your RD Session Host server, you must switch the server back into execution mode (RD-Execute) before users can remotely connect to your server. You can switch between the install and execute modes from the command line using these commands:

```
change user /install
```

```
change user /execute
```

To determine the current install mode of your RD Session Host server, use this command:

```
change user /query
```

You can also install applications on your RD Session Host server by using the Programs portion of Control Panel (shown in Figure 4-6). If you do this on an RD Session Host server, an additional option, Install Application On Remote Desktop Server, is displayed that is not available on a Windows Server 2008 R2 computer that does not have the RD Session Host role service installed. Clicking this option launches the Install Program From Floppy Disk Or CD-ROM Wizard, which walks you through the installation process by automatically switching the server to RD-Install mode, installs the program, and switches the server back to RD-Execute mode.




FIGURE 4-6 Installing an application on an RD Session Host server.

For additional tips on installing applications on terminal servers, see the sidebar titled “Direct from the Source: Best Practices for Installing Applications on RD Session Host Servers” in this chapter.

Direct from the Source: Best Practices for Installing Applications on RD Session Host Servers

Install related applications, or applications that have dependencies on other local applications, on the same RD Session Host server. For example, you should install Microsoft Office as a suite on the same RD Session Host server instead of installing individual Office programs on separate RD Session Host servers. This reduces the installation, management, and servicing overhead required for the programs.

You should consider installing individual applications on separate RD Session Host servers in the following circumstances:

- The application has compatibility issues that might affect other programs.
- The number of application users might exceed server capacity.
- The applications are resource (CPU, memory, and so forth) intensive and negatively affect performance when multiple instances are running at the same time.

It is also a good practice always to consult with the application vendor to ensure that the application being installed will function correctly for multiple users in an RD Session Host server environment. Application vendors sometimes provide fixes or compatibility scripts for applications to ensure that they function correctly in a multi-user Remote Desktop Services environment.

—CSS Global Technical Readiness (GTR) team

Managing RD Session Host Servers

After the RD Session Host role service has been installed and configured on your servers, you are ready to begin administering your RD Session Host servers, licenses, and users. Windows Server 2008 R2 includes a number of tools for administering RD Session Host servers, including the following:

- Microsoft Management Console (MMC) snap-ins
- Command-line tools
- Group Policy
- Windows Management Instrumentation (WMI)
- Windows PowerShell

Managing RD Session Host Servers Using MMC Snap-ins Installing the RD Session Host role service installs four MMC snap-ins in the Remote Desktop Services program group under Administrative Tools on your Start menu. These snap-ins include:

- **Remote Desktop Services Manager** Lets you manage users, sessions, and applications running on a local or remote RD Session Host server. For example, using this snap-in you can connect to and disconnect from sessions; display information about servers, sessions, users, and processes; log off users; monitor sessions; remotely control a user's session; reset sessions; send messages to users; and terminate processes. You can also use the My Group feature to organize your RD Session Host servers into groups to manage them more easily, and you can even import a list of RD Session Host servers from an RD Connection Broker farm. Figure 4-7 shows the Remote Desktop Services Manager console with user Karen Berg (CONTOSO\kberg) remotely connected to an RD Session Host server named SEA-RDS4. The Action Pane on the right side shows some of the actions you can perform for this connected user.




FIGURE 4-7 The Remote Desktop Services Manager console.

- **Remote Desktop Session Host Configuration** Lets you configure various properties of an RD Session Host server, including listeners, temporary folders, security, and licensing. For example, using this snap-in you can enable or disable automatic logons; enable or disable logons through the connection; enable or disable session remote control; rename a connection; override user profile settings for wallpaper; set client device mapping and connection parameters; set connection timeouts; set permissions on the connection; set the level of encryption; set the maximum number of sessions allowed; set whether to disconnect broken connections; specify a connection transport and transport properties; specify a connection type; and specify a program to run when a user logs on.

- **Remote Desktops** Lets you manage connections to RD Session Host servers and to client computers that have the Remote Desktop feature enabled on them. For example, using this snap-in, you can remotely administer multiple computers from a central location by connecting directly to the remote desktop of each computer. Note that by default, when you add a connection to the Remote Desktops console, the console connects you to an administrative session on the remote computer. This is equivalent to launching the Remote Desktop Connection client using the **/admin** option, which was first introduced in Windows Server 2008 and Windows Vista SP1. For more information about the **/admin** option, see the article titled “Changes to Remote Administration in Windows Server 2008” in the Microsoft Knowledge Base at <http://support.microsoft.com/kb/947723>.
- **RemoteApp Manager** Lets you make programs installed on an RD Session Host server available to users as RemoteApp programs so that users can launch these programs from the user’s local Start menu. After they are launched, RemoteApp programs actually run on an RD Session Host server, but to the user they appear as if they are running on the user’s local computer. For more information about this feature, see the section titled “Understanding RemoteApp” later in this chapter.

In addition to the snap-ins just mentioned, you can also use the Active Directory Users and Computers snap-in to manage certain per-user aspects of an RD Session Host environment. Specifically, you can use the settings on the following tabs of the properties sheet for a user account:

- **Remote Control** Used to enable or disable remote control, specify the desired level of remote control (view or interact), and require the user’s permission to observe or control sessions.
- **Remote Desktop Services Profile** Used to set the path to the Remote Desktop Services user profile for each user, set the path to the user’s home folder, and enable or disable RD Session Host logons for the user.
- **Environment** Used to specify a program to run when a user logs on, specify that client drives and printers connect at logon, and select to default to the client’s main printer.
- **Sessions** Used to set the maximum duration on sessions, set the maximum idle time for a session, set the maximum time a disconnected session remains active, specify whether to disconnect or reset a broken connection, and modify settings for reconnecting disconnected sessions.



Note Many of these settings can also be configured using the Remote Desktop Services Configuration snap-in or by configuring Group Policy settings. When settings configured using these different methods conflict, they will be overridden according to the following order of precedence:

1. Active Directory Users and Computers snap-in
2. Remote Desktop Services Configuration snap-in
3. Group Policy

In other words, Group Policy settings trump any other method for configuring Remote Desktop Services settings.

Managing RD Session Host Servers Using Command-Line Tools Many aspects of RD Session Host servers can be administered from the command line. Table 4-6 lists the command-line tools for doing this that are available in Windows Server 2008 and Windows Server 2008 R2, while Table 4-7 lists some command-line tools that were deprecated beginning with Windows Server 2008.

TABLE 4-6 Remote Desktop Services Command-Line Tools Available in Windows Server 2008 and Later

Command	Description
Change	Changes RD Session Host settings for logons, Component Object Model (COM) port mappings, and install mode
Change logon	Enables or disables logons from client sessions on an RD Session Host server, or displays current logon status
Change port	Lists or changes the COM port mappings to be compatible with MS-DOS applications
Change user	Changes the install mode for the RD Session Host server
Chglogon	Enables or disables logons from client sessions on an RD Session Host server, or displays current logon status
Chgport	Lists or changes the COM port mappings to be compatible with MS-DOS applications
Chguser	Changes the install mode for the RD Session Host server
Flattemp	Enables or disables flat temporary folders
Logoff	Logs off a user from a session on an RD Session Host server, and deletes the session from the server
Msg	Sends a message to a user on an RD Session Host server
Mstsc	Creates connections to an RD Session Host server or other remote computers
Qappsrv	Displays a list of all RD Session Host servers on the network
Qprocess	Displays information about processes that are running on an RD Session Host server

Command	Description
Query	Displays information about processes, sessions, and RD Session Host servers
Query process	Displays information about processes that are running on an RD Session Host server
Query session	Displays information about sessions on an RD Session Host server
Query termserver	Displays a list of all RD Session Host servers on the network
Query user	Displays information about user sessions on an RD Session Host server
Quser	Displays information about user sessions on an RD Session Host server
Qwinsta	Displays information about sessions on an RD Session Host server
Reset session	Enables you to reset (delete) a session on an RD Session Host server
Rwinsta	Enables you to reset (delete) a session on an RD Session Host server
Shadow	Enables you to remotely control an active session of another user on an RD Session Host server
Tscon	Connects to another session on an RD Session Host server
Tsdiscon	Disconnects a session from an RD Session Host server
Tskill	Ends a process running in a session on an RD Session Host server
Tsprof	Copies the Remote Desktop Services user configuration information from one user to another

TABLE 4-7 Remote Desktop Services Command-Line Tools That Have Been Deprecated in Windows Server 2008 and Later

Command	Function
Tsshutdn	Shuts down an RD Session Host server
Register	Registers a program so that it has special execution characteristics
Cprofile	Removes user-specific file associations from a user's profile

Managing RD Session Host Servers Using Group Policy Many aspects of RD Session Host servers can be administered using Group Policy. Table 4-8 summarizes the various policy areas for managing RD Session Host servers. Of particular interest are the policy settings for configuring application compatibility for RD Session Host servers, which are found here:

Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\ Remote Desktop Session Host\Connections

These policy settings are new in Windows Server 2008 R2, and there are four of them:

- **Do Not Use Remote Desktop Session Host Server IP Address When Virtual IP Address Is Not Available** This policy setting specifies whether a session uses the IP address of the Remote Desktop Session Host server if a virtual IP address is not

available. If you enable this policy setting, the IP address of the RD Session Host server is not used if a virtual IP address is not available. The session will not have network connectivity. If you disable or do not configure this policy setting, the IP address of the RD Session Host server is used if a virtual IP address is not available.

- **Select The Network Adapter To Be Used For Remote Desktop IP Virtualization** This policy setting specifies the IP address and network mask that corresponds to the network adapter used for virtual IP addresses. The IP address and network mask should be entered in Classless Inter-Domain Routing notation—for example, 192.0.2.96/24. If you enable this policy setting, the specified IP address and network mask are used to select the network adapter used for the virtual IP addresses. If you disable or do not configure this policy setting, Remote Desktop IP Virtualization is turned off. A network adapter must be configured for Remote Desktop IP Virtualization to work.
- **Turn Off Windows Installer RDS Compatibility** This policy setting specifies whether Windows Installer RDS Compatibility runs on a per-user basis for fully installed applications. Windows Installer allows one instance of the msixexec process to run at a time. By default, Windows Installer RDS Compatibility is turned on. If you enable this policy setting, Windows Installer RDS Compatibility is turned off, and only one instance of the msixexec process can run at a time. If you disable or do not configure this policy setting, Windows Installer RDS Compatibility is turned on, and multiple per-user application installation requests are queued and handled by the msixexec process in the order in which they are received.
- **Turn On Remote Desktop IP Virtualization** This policy setting specifies whether Remote Desktop IP Virtualization is turned on. By default, Remote Desktop IP Virtualization is turned off. If you enable this policy setting, Remote Desktop IP Virtualization is turned on. You can select the mode in which this setting is applied. If you are using Per Program mode, you must enter a list of programs to use virtual IP addresses. List each program on a separate line—for example, explorer.exe\mstsc.exe. (Do not enter any blank lines between programs.) If you disable or do not configure this policy setting, Remote Desktop IP Virtualization is turned off.



Tip For detailed information concerning what policy settings are available in each policy area and what each policy does, see “Group Policy Settings for Remote Desktop Services in Windows Server 2008 R2,” found at <http://technet.microsoft.com/en-us/library/ee791928.aspx>. For additional information, see “Group Policy Settings Reference for Windows and Windows Server,” available from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=18c90c80-8b0a-4906-a4f5-ff24cc2030fb>.

TABLE 4-8 Group Policy Settings for Administering RD Session Host Servers

Type of Policy Settings	Path in Group Policy Management Editor
Computer Configuration Policy Settings	
Policy settings for configuring application compatibility for RD Session Host servers (new in Windows Server 2008 R2)	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Connections
Policy settings for configuring Connections settings for RD Session Host servers	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Connections
Policy settings for configuring Device And Resource Redirection settings for RD Session Host servers	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Device and Resource Redirection
Policy settings for configuring Licensing settings for RD Session Host servers	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Licensing
Policy settings for configuring Printer Redirection settings for RD Session Host servers	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Printer Redirection
Policy settings for configuring Profiles settings for RD Session Host servers	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Profiles
Policy settings for configuring RD Connection Broker settings for RD Session Host servers	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Remote Session Environment
Policy Settings for configuring Remote Session Environment settings for RD Session Host servers	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Remote Session Environment
Policy settings for configuring Security settings for RD Session Host servers	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Security
Policy settings for configuring Session Time Limits settings for RD Session Host servers	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Session Time Limits
Policy settings for Temporary Folders settings for RD Session Host servers	Computer Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Temporary folders

Type of Policy Settings	Path in Group Policy Management Editor
User Configuration Policy Settings	
Policy settings for configuring Connections settings for RD Session Host servers	User Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Connections
Policy settings for configuring Device And Resource Redirection settings for RD Session Host servers	User Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Device and Resource Redirection
Policy settings for configuring Printer Redirection settings for RD Session Host servers	User Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Printer Redirection
Policy settings for configuring Remote Session Environment settings for RD Session Host servers	User Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Remote Session Environment
Policy settings for configuring Session Time Limits settings for RD Session Host servers	User Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\Remote Desktop Session Host\Session Time Limits

Managing RD Session Host Servers Using WMI There is also a WMI provider to enable administration of RD Session Host servers and also other Remote Desktop Services role services by using WMI interfaces. Specifically, there are five categories of WMI classes available for managing Remote Desktop Services in Windows Server 2008 R2:

- Remote Desktop Services Configuration Classes
- Remote Desktop Gateway Classes
- Remote Desktop License Server Classes
- RemoteApp Classes
- Remote Desktop Connection Broker Classes

For more information on these WMI classes, properties, and methods supported by Windows Server 2008 Terminal Services, see “Remote Desktop Services WMI Provider Reference” in the MSDN Library at <http://msdn.microsoft.com/en-us/library/aa383515.aspx>.

Managing RD Session Host Servers Using Windows PowerShell Beginning with Windows Server 2008 R2, you can use Windows PowerShell to manage many aspects of RD Session Host servers and also aspects of other Remote Desktop Services role services. This allows administrative tasks to be scripted, allowing you to automate complex and recurring administrative tasks. You can also change settings and perform tasks directly from the Windows PowerShell command line, without having to write, save, or run a script. Examples

of administrative tasks you can perform by using the Remote Desktop Services Provider for Windows PowerShell include the following:

- View configuration settings for a remote desktop server.
- Edit configuration settings for a remote desktop server.
- Create and configure a remote desktop server connection.
- Publish or remove a RemoteApp program.
- Create and configure a remote desktop server farm.
- Configure RemoteApp and Desktop Connection for virtual desktops and RemoteApp programs.
- Manage a Remote Desktop Licensing server.
- Manage a Remote Desktop Gateway server

.For more information about using Windows PowerShell with Remote Desktop Services, see "Remote Desktop Services Management" found at <http://technet.microsoft.com/en-us/library/dd939782.aspx>.

Understanding RemoteApp

Terminal Services on Windows Server 2003 and earlier could provide users only with entire session-based desktops that included Terminal Services-enabled applications. This was sometimes confusing to users because it meant they had to contend with having two desktops—their local computer's desktop and the remote desktop—presented to them via Terminal Services.

For example, at any given time a user might have several applications running on her local desktop plus additional applications running on her remote desktop. This presented users with interesting challenges. For instance, if a user wanted to quickly switch between applications on her local desktop, she could use the Alt+Tab or Alt+Esc keyboard accelerators to do this. But if the user wanted to do the same with applications on her remote desktop, she needed to use Alt+Page Up, Alt+Page Down, or Alt+Insert instead. And if she wanted to switch from a local application to a remote one, the easiest way was probably just to use the mouse! Such confusion and occasional frustration introduced by having two desktops created inefficiencies that resulted in loss of worker productivity.

Beginning with Windows Server 2008, however, Terminal Services introduced the ability for terminal servers to provide users with access to individual applications running on a terminal server. These remote applications, known as RemoteApp programs, can be launched from the user's Start menu or desktop shortcuts, and when they are open they look and feel the same as locally installed programs. This look and feel extends to resizing, maximizing,

minimizing, and cascading program windows; cut and paste operations between program windows; drag and drop support between multiple monitors; and notification icons displayed in the notification area. The result is that a user running both local and RemoteApp programs on his computer's local desktop might be unaware of the difference between a program installed locally and one running on a terminal server. RemoteApp thus integrates Terminal Services applications into the user's own desktop instead of presenting the user with a second, separate desktop.

In Windows Server 2008, this functionality was known as Terminal Services RemoteApp (TS RemoteApp). Beginning with Windows Server 2008 R2, however, this functionality is now simply known as RemoteApp.

How RemoteApp works RemoteApp functionality is implemented as part of the RD Session Host role service and requires Windows Server 2008 or later on the server side and RDC 6.0 or later on the client side. This means that the user's computer must be running either Windows 7, Windows Vista SP1 or later, or Windows XP SP3.

The process by which a RemoteApp program is launched on an RD Session Host server is as follows:

1. When a remote user tries to launch a RemoteApp program, a new instance of Rpdinit.exe, the RemoteApp Logon Application, is started in the session space on the RD Session Host server.
2. Rpdinit.exe then launches Rdpshell.exe, which provides the functionality for running the RemoteApp program's process. Rdpshell.exe replaces the usual Explorer.exe shell, which does not support RemoteApp functionality, and provides event hooks and APIs for monitoring the state of the user's taskbar, the position of windows, notification icons, and so on. Rdpshell.exe also opens a virtual channel that allows RemoteApp-specific commands to be transmitted from the client to the server.
3. Rpdinit.exe then monitors the RemoteApp program's process during the lifetime of the process. For example, if the process terminates abnormally, Rpdinit.exe will restart it.
4. If the user then launches additional RemoteApp programs on the same RD Session Host server, these programs all share the same Remote Desktop Services session with the first program launched above.



Tip When you are using a RemoteApp program and save a document or other file you are working on with the program, the document or other file is saved within your user profile on the RD Session Host server, not on your local workstation. The user profiles for Remote Desktop Services users are stored on RD Session Host servers under the %SystemDrive%\Users directory. If you need to save a document or other file on your local computer instead of on the RD Session Host server, save it to a redirected drive by browsing to \\tsclient in the Save file dialog box and selecting a redirected drive.

When a user terminates a RemoteApp program, the RemoteApp session between the user and the RD Session Host server is placed into a disconnected state. The RD Session Host server then applies heuristics to determine whether the RemoteApp session should remain disconnected or whether the user should be logged off from the RD Session Host server to end the session. The termination logic for RemoteApp programs is configurable using Group Policy and basically works like this:

1. The user closes a RemoteApp program window on his desktop.
2. The RD Session Host server checks the user's session to see whether there are any remaining active RemoteApp windows still open on the user's desktop. If the answer to this is Yes, the user's session remains connected.
3. If there are no remaining active RemoteApp program windows on the user's desktop, the RD Session Host server next checks whether there are any RemoteApp program notification icons being displayed in the system tray on the user's desktop. If the answer to this is Yes, the user's session continues to remain connected.
4. If there are no RemoteApp program windows or notifications remaining on the user's desktop, the RD Session Host server waits 20 seconds before terminating the user's session, just in case the user decides to launch another RemoteApp program immediately. If no RemoteApp program is launched during this time interval, the RD Session Host server disconnects the user's session and the RDC client process exits. If another RemoteApp program is launched within the time interval, the user's session remains connected and the new process runs within the existing session.
5. After the user's session has been disconnected, the session remains disconnected for a configurable period of time, after which the user is logged off from the disconnected session. The time interval between disconnection and logoff can be configured using the Group Policy setting Set Time Limit For Logoff Of RemoteApp Sessions. The reason for providing for the configurability of this time interval is because it is much faster to connect to a disconnected session than to start a new session.

Adding RemoteApp Programs You must add RemoteApp programs to your RD Session Host server before they can be made available to users on your network. To add RemoteApp programs to your RD Session Host server, you use the RemoteApp Manager snap-in, which is found in the Remote Desktop Services program group under Administrative Tools on the Start menu. Figure 4-8 shows the RemoteApp Manager snap-in with no RemoteApp programs yet configured.

In this example, Microsoft 2007 Office System has been installed on an RD Session Host server running Windows Server 2008 R2 and named SEA-RDS4 in the contoso.com domain. To add Word 2007 to the list of RemoteApp programs on the RD Session Host server, open RemoteApp Manager and click Add RemoteApp Programs in the Actions pane to launch the RemoteApp Wizard. Then on the Welcome screen, click Next and select the check box for Word 2007. (See Figure 4-9.)




FIGURE 4-8 The RemoteApp Manager snap-in.




FIGURE 4-9 Adding Word 2007 as a RemoteApp program.

When you finish the wizard, Word is listed as a RemoteApp program in the RemoteApp Manager snap-in (as shown in Figure 4-10).




FIGURE 4-10 Word is now available as a RemoteApp program.

Configuring and Deploying RemoteApp Programs To configure a RemoteApp program, select the program in the RemoteApp Manager snap-in and then select the action you want to perform in the Actions pane on the right side.

For users to make use of RemoteApp programs, these programs must first be deployed to their client computers. This can be done by

- Packaging the RemoteApp programs as Windows Installer (.msi) files, and distributing them to users using the Software Installation feature of Group Policy
- Packaging the RemoteApp programs as either Remote Desktop Protocol (.rdp) files or Windows Installer (.msi) files, and making them available to users by copying them to a network share
- Publishing the RemoteApp programs on the intranet or over the Internet using RD Web Access and, optionally, RD Gateway when greater security is required

The Create .Rdp File and Create Windows Installer Package actions that are available in the Actions pane when a RemoteApp program is selected allow you to package the RemoteApp program for deployment in enterprise environments. For smaller environments, RD Web Access provides an easy way of deploying RemoteApp programs to users. By default, adding a RemoteApp program also makes it available to users through RD Web Access if this additional role service has been installed. For more information about using RD Web Access

to deploy RemoteApp programs to users, see the sections titled “Understanding Remote Desktop Web Access” and “Understanding RemoteApp and Desktop Connections” later in this chapter.

To configure the properties of the selected RemoteApp program, click Properties in the Action pane. Beginning with Windows Server 2008 R2, the properties sheet for RemoteApp programs includes an additional tab named User Assignment. (See Figure 4-11.) This new tab lets you specify which domain users and domain groups will be able to see the icon for the RemoteApp program when they access the RD Web Access Web page. By default, all authenticated domain users can see the icon for a RemoteApp program on the RD Web Access Web site unless otherwise configured. Note that to assign users to a RemoteApp program, the RD Session Host server on which the RemoteApp program is configured must be a member of an Active Directory domain. In addition, to run the RemoteApp program the user must be a member of the Remote Desktop Users group on the RD Session Host server.




FIGURE 4-11 Configuring which users are allowed to access the RemoteApp program using the RD Web Access Web page.

To use this new feature to assign users, groups, or both to a RemoteApp program, follow these steps:

1. On the RD Session Host server, open RemoteApp Manager. To open RemoteApp Manager, click Start, point to Administrative Tools, point to Remote Desktop Services, and then click RemoteApp Manager.

2. In the RemoteApp Programs list, click the program to which you want to assign domain users and domain groups.
3. In the Actions pane for the program, click Properties, and then click the User Assignment tab.
4. Click the specified domain users and domain groups, and then click Add.
5. Use the Select Users Or Groups dialog box to select the domain users and domain groups to assign to the RemoteApp program.
6. Click OK to close the Select Users Or Groups dialog box.
7. Click OK to close the RemoteApp Properties dialog box.

Understanding Remote Desktop Web Access

The Remote Desktop Web Access (RD Web Access) role service of Windows Server 2008 R2 was formerly called Terminal Services Web Access (TS Web Access) in Windows Server 2008. Installing the Remote Desktop Web Access role service lets you use Internet Information Services (IIS) to simplify the deployment of RemoteApp programs, session-based desktops, and virtual desktops to users on your network. New features of the RD Web Access role service introduced in Windows Server 2008 R2 include the following:

- **Per user RemoteApp program filtering using RemoteApp and Desktop Connections** RD Web Access is now integrated with RemoteApp and Desktop Connections, a new feature of Windows 7 and Windows Server 2008 R2 that provides a personalized view of RemoteApp programs, session-based desktops, and virtual desktops to users. This enables RD Web Access to filter the view on a per-user basis so that each user logging on to RD Web Access sees only the programs that the administrator has configured for them to see. For more information on this feature, see the section titled “Understanding RemoteApp and Desktop Connections” later in this chapter.
- **Single sign-on between RD Session Host and RD Web Access** This enhancement allows users to enter their user name and password only once when connecting to a RemoteApp program by using RD Web Access.
- **Public and private computer option** There are now two ways for users to access the RD Web Access Web page: public and private mode. When a user selects public mode, her user name is not remembered in the Web browser and RD Web Access cookies storing her user name time out in 20 minutes. When the user selects private mode, cookies storing her user name remain available for four hours. In either mode, passwords are not stored.
- **Forms-based authentication** This enables applications to provide their own logon page and perform their own credentials verification, and it uses ASP.NET to authenticate users, redirect unauthenticated users to the logon page, and perform all the necessary cookie management.

RD Web Access enables administrators to deploy RemoteApp programs via a Web browser to users on a corporate intranet. Using RD Web Access, a user who needs to run a RemoteApp program simply connects to a special Web site and clicks on the RemoteApp program's icon, and the program starts on the user's desktop. Although deploying RemoteApp programs using Group Policy is effective for a large enterprise, RD Web Access can provide a simple means of deploying RemoteApp programs for small and medium-sized businesses because of its ease of use. In addition, by implementing RD Web Access together with RD Gateway, administrators can securely deploy RemoteApp programs to users over an unsecure Internet connection without the need to configure a virtual private network (VPN) for those users.

The experience of running RemoteApp programs deployed using RD Web Access is identical to the experience of launching them from .rdp files or from .msi files deployed using Group Policy. Whichever way a user launches a RemoteApp program, the user can interact with the program just as if it was locally installed on her computer.

RD Web Access also includes a customizable Web Part that provides flexibility in how you display RemoteApp programs you want to deploy. By using this customizable Web Part, administrators can also create their own customized Web page or Windows SharePoint Services site for deploying RemoteApps to users.

RD Web Access also provides users with the option of connecting to the remote desktop of any computer on which they have logon privileges. When implemented together with RD Gateway, this enables a user to remotely access the desktop of his corporate desktop computer over the Internet. For more information concerning RD Gateway, see the section titled "Understanding Remote Desktop Gateway" later in this chapter.

How RD Web Access Works

RD Web Access is implemented as a separate role service of the Remote Desktop Services role of Windows Server 2008 R2. Installing the RD Web Access role service on a server also installs the Web Server (IIS) role along with some of its components, which is needed to host the Web site that users connect to using their Web browsers to launch RemoteApp programs.

Both the RD Web Access and RD Session Host role services must be present for RD Web Access to work. The simplest configuration is to install both the RD Web Access and RD Session Host role services on a single server. Users can then connect to the RD Web Access Web page using their Web browser, log on using their domain credentials, and launch RemoteApp programs, which then run on the RD Session Host server.

The RD Web Access and RD Session Host role services can also be installed on separate servers if needed. If this is done, however, you must add the computer account of the RD Web Access server to the RD Web Access Computers security group on your RD Session Host.

For larger deployments, you might install RD Web Access on a front-end Web server to service multiple RD Session Host servers on the back end. You can then configure RD Web Access to populate its list of RemoteApp programs from all your RD Session Host servers, including servers that belong to an RD Session Host farm.

To use RD Web Access and fully use the new features available in Windows Server 2008 R2, your client computers need Remote Desktop Connection 7.0 (RDC 7.0), which is included by default in Windows 7. If your client computers are running Windows Vista SP1 or SP2, or Windows XP SP3, you can download RDC 7.0 from <http://support.microsoft.com/kb/969084>. To connect to the RD Web Access server, a user opens a Web browser such as Internet Explorer and types **https://<server_name>/rdweb** in the address bar as described in the next section.

Installing, Configuring, and Using RD Web Access

You can install the RD Web Access role service on the Standard, Enterprise, or Datacenter edition of Windows Server 2008 R2 using any of the methods described earlier in the section titled “Installing an RD Session Host Server.” After you have installed the role service, you need to log on to the RD Web Access Web site to finish configuring your server.

To access the RD Web Access Web page (shown in Figure 4-12), use one of these methods:

- On the RD Web Access server, click Start, point to Administrative Tools, point to Remote Desktop Services, and then click Remote Desktop Web Access Configuration.
- Open the URL https://<server_name>/rdweb in Internet Explorer, where <server_name> is the fully qualified domain name (FQDN) of the RD Web Access server.

If you connect to the RD Web Access Web site using a public computer—for example, a kiosk computer at a coffee shop or airport, or using a computer that you share with others—select the This Is A Public Or Shared Computer option. Selecting this option requires that you provide both your user name and password each time you sign in to the RD Web Access Web site. However, if you connect to the RD Web Access Web site using your own work computer (which you don’t share with other users), select the This Is A Private Computer option. Selecting this option means that your user name is remembered so that you only need to enter your password each time you sign in to the RD Web Access Web site.

To configure your RD Web Access server, log on to the RD Web Access Web site using either the local Administrator account or an account that belongs to the RD Web Access Administrators group on the RD Web Access server. After you’ve logged on, you must specify the source that provides the RemoteApp programs and desktops to users. To do this, select the Configuration tab of the RD Web Access Web page as shown in Figure 4-13.




FIGURE 4-12 The login page for the RD Web Access Web site.



FIGURE 4-13 Specify the source from which the RD Web Access server will pull RemoteApp lists.


Formerly in Windows Server 2008, the TS Web Access server could pull the RemoteApps configuration only from a single source, which was either a single terminal server or a terminal server farm. Beginning with Windows Server 2008 R2, however, you can configure the RD Web Access server so that users can pull RemoteApp lists from multiple RD Session Host and RD Virtualization Host servers while routing the connection request to the proper server. To provide users with access to RemoteApp programs, session-based desktops, and virtual desktops, you must first configure your RD Web Access to specify the source that hosts those resources. That source can be either of the following:

- An RD Connection Broker server
- One or more RD Session Host servers

An RD Connection Broker server can provide users with access to RemoteApp programs, session-based desktops hosted on RD Session Host servers, and virtual desktops hosted on RD Virtualization Host servers. RD Session Host servers provide users with access only to RemoteApp programs and session-based desktops.

After the source has been configured, the RemoteApp Programs tab of the RD Web Access Web page displays the available RemoteApp programs and desktops. Users can then launch RemoteApp programs on their client computers by following these steps:

1. Open Internet Explorer, and type the URL **https://<server_name>/rdweb**, where <server_name> is the FQDN of the RD Web Access server.
2. On the RD Web Access logon page, type your credentials in the form <domain>\<username> along with your password and click Sign In.
3. On the RD Web Access Web page, select the RemoteApp Programs tab to display the available RemoteApp programs and desktops you can access:



4. Double-click on a RemoteApp program or desktop to launch the item. (You might need to supply additional credentials.)




Note Ordinary users do not see the Configuration tab that administrators see when they access the RD Web Access Web site.

In addition to being able to access RemoteApp programs and desktops, users can also use the RD Web Access Web site to connect to the desktop of any computer on the network, a feature known as Remote Desktop Web Connection. For this to work, however, the remote computer must have Remote Desktop enabled, and the user must be a member of the Remote Desktop Users group on the remote computer. This can be especially useful, for example, when the user at home needs to connect to the desktop of his computer on the corporate network. (To do this in a secure fashion, you must implement RD Web Access together with RD Gateway.)

For example, let's say that you are a user named Karen Berg who wants to connect to a computer named SEA-CLI-888 by using Remote Desktop Web Connection from another computer. You could do this as follows:


1. Open Internet Explorer, log on to the RD Web Access Web site, and select the Remote Desktop tab.
2. Type the name of the computer you want to connect to, and select other options as desired:



3. Specify your credentials in the Windows Security dialog:



4. The desktop of the remote computer is now displayed using Remote Desktop Connection:



5. You can now run programs and open files on the remote computer as if you were sitting in front of it and using its keyboard.

Although you can use the preceding approach to make RemoteApp programs and desktops available to users using a Web browser, Windows 7 and Windows Server 2008 R2 also includes a new feature called RemoteApp and Desktop Connection that gives you even greater control over deploying RemoteApp programs and desktops for your users. This is described in the next section.

Understanding RemoteApp and Desktop Connections

RemoteApp and Desktop Connection is a new feature of Windows 7, and Windows Server 2008 R2 provides users with a personalized view of RemoteApp programs, session-based desktops, and virtual desktops. Specifically,

- When you start a RemoteApp program or a session-based desktop, a Remote Desktop Services session is started on the RD Session Host server that hosts the remote desktop or RemoteApp program.
- When you connect to a virtual desktop, a Remote Desktop connection is made to a virtual machine running on an RD Virtualization Host server.

By using RemoteApp and Desktop Connection, you can create an aggregated, customized view of RemoteApp programs and desktops and assign them to users. After you have assigned them, users can then subscribe to a Web feed that will seamlessly integrate the RemoteApp programs and desktops into the Start menu on their Windows 7 computers and automatically update the list as the published RemoteApp programs and desktops change. Specifically, RemoteApp and Desktop Connections provides the following benefits:

- Published RemoteApp programs can be launched from the user's Start menu just like locally installed applications.
- Published Remote Desktop Connections are included alongside RemoteApp programs on the user's Start menu.
- Any changes to a published connection—for example, additional published RemoteApp programs—automatically appear on the user's Start menu without any action needed on the user's part. Users have to log on only once when they create the connection, after which updates happen automatically with no prompt for user credentials.
- Users can use Start menu search to find RemoteApp programs in the same way they search for locally installed programs.
- RemoteApp and Desktop Connections supports per-user application filtering of RemoteApp programs.
- Client computers do not have to be domain-joined to use RemoteApp and Desktop Connections.

In addition, RemoteApp and Desktop Connections is built on standard technologies such as XML and HTTPS. This means that developers can use standard tools to build solutions around RemoteApp and Desktop Connections.

How RemoteApp and Desktop Connections Works

RemoteApp and Desktop Connection functionality is not provided by a single component. Instead, different aspects of RemoteApp and Desktop Connection functionality are provided by Remote Desktop Services role services and by other components including the following:


- Remote Desktop Session Host servers
- Remote Desktop Virtualization Host servers
- Remote Desktop Connection Broker
- Remote Desktop Web Access
- Remote Desktop Gateway
- Remote Desktop Connection 7.0
- Windows 7 client computers

Configuring RemoteApp and Desktop Connections

You can use RemoteApp and Desktop Connections to configure which RemoteApp programs, session-based desktops, and virtual desktops should be available for users on your network. For example, to make RemoteApp programs available via a Web browser and using an RD Connection Broker server, follow these steps:


1. Install the Remote Desktop Connection Broker role service on a server running Windows Server 2008 R2. For more information about this role service, see the section titled “Understanding Remote Desktop Connection Broker” later in this chapter.
2. On your RD Session Host server, use RemoteApp Manager to add applications to the RemoteApp programs list. Make sure these programs are enabled for RD Web Access. For information on how to do this, see the section titled “Adding RemoteApp Programs” earlier in this chapter.
3. Install the RD Web Access role service on a server running Windows Server 2008 R2. This can be the same server as your RD Session Host server, but it should not be the same as your RD Connection Broker server.
4. If your RD Web Access and RD Session Host servers are separate servers, add the computer account of the RD Web Access server to the RD Web Access Computers security group on your RD Session Host server.

5. Open the RD Web Access Web site using one of the methods described in the section titled "Installing, Configuring, and Using RD Web Access" earlier in this chapter. Log on to the site using either the local Administrator account on the RD Web Access server or an account that is a member of the RD Web Access Administrators group on the RD Web Access server.
6. Select the Configuration tab on the RD Web Access Web site, choose An RD Connection Broker Server, and type the NetBIOS name or FQDN of your RD Connection Broker server:




7. On your RD Connection Broker server, open the Remote Desktop Connection Manager snap-in, right-click on the root node labeled Remote Desktop Connection Manager: `<server_name>`, and select Properties to display the RemoteApp and Desktop Connection Properties dialog.

8. On the Connection Settings tab, specify a display name and connection ID or accept the defaults:



The display name is used to identify RemoteApp and Desktop Connection resources (RemoteApp programs, session-based desktops, and virtual desktops) that are being provided to users by the RD Connection Broker server. Specifically, the display name appears in the user's Start menu under All Programs in a folder named RemoteApp And Desktop Connections on a computer running Windows 7 that has been configured to use RemoteApp and Desktop Connections provided by the RD Connection Broker server. The connection ID is displayed when the user runs a RemoteApp program from an RD Web Access Web site that has been configured to provide RemoteApp and Desktop Connections from the RD Connection Broker server.

9. On the RD Web Access tab, make sure your RD Web Access server's name is displayed:



At this point, you have published the RemoteApp programs so that they can be available to users through RemoteApp and Desktop Connections. To access these published programs, the client side of RemoteApp and Desktop Connections must be configured on the user's Windows 7 computer. There are several ways this can be done:

- The user can manually set up a new connection using the RemoteApp and Desktop Connections Control Panel item together with a special URL provided to the user by the administrator. For more information on how to manually configure the client side of RemoteApp and Desktop Connections, see the section titled "Configuring RemoteApp and Desktop Connections" earlier in the chapter.
- The administrator can create a Workspace Configuration (.wcx) file using an RD Connection Broker server and distribute it to Windows 7 users so that RemoteApp and Desktop Connection can be configured without the need of having the user manually configure the RemoteApp and Desktop Connections Control Panel item. For more information on how to create a .wcx file, see the section titled "Understanding Remote Desktop Connection Broker" later in this chapter.
- The administrator can create a .wcx file and use Group Policy to silently run a script on Windows 7 computers so that RemoteApp and Desktop Connection is set up automatically when users log on to their computers.

After the client side of RemoteApp and Desktop Connections has been configured, Windows 7 users will see a new RemoteApp and Desktop Connections program group on their Start menu, which they can use to launch RemoteApp programs, session-based desktops, and virtual desktops that have been published for them to use. (See Figure 4-14.)




FIGURE 4-14 RemoteApp and Desktop Connections program group on the Start menu.

After the user launches an available RemoteApp program, session-based desktop, or virtual desktop from the user's Start menu, a RemoteApp And Desktop Connections icon is displayed in the notification area of the taskbar. (See Figure 4-15.) The presence of this icon indicates that the user has successfully connected to RemoteApp and Desktop Connections. The user can also use this icon to perform the following tasks:

- List all the connections the user is currently connected to.
- Disconnect one or all connections that are currently connected.
- Open the RemoteApp Desktop Connections in Control Panel to view the properties of the connection, update any or all of the connections, or both.




FIGURE 4-15 RemoteApp And Desktop Connections icon in the taskbar notification area.

After the user is no longer connected to any connections, the icon disappears from the taskbar notification area.

If you also want to allow users to access the RD Web Access server from the Internet, you can deploy an RD Gateway server on your perimeter network and configure it for doing this. For more information, see “Understanding Remote Desktop Gateway” later in this chapter.

In smaller environments, users can manually set up RemoteApp and Desktop Connections to pull a feed of available RemoteApp programs, session-based desktops, and virtual desktops from your RD Web Access server. To do this, the user needs to follow these steps:


1. Open the RemoteApp and Desktop Connections Control Panel item. Initially, there will be no connections available on the user’s computer:




2. Next, the user needs to type the URL for the RD Web Access Web site, which is always in the following form:

`https://<server_name>/RDWeb/Feed/webfeed.aspx`


where `<server_name>` is the FQDN of the RD Web Access server. In this example, it is SEA-RDS4.contoso.com:




3. After the user's computer has established a connection with the specified RD Web Access server, the user's computer is ready to pull a feed of the RemoteApp programs and desktops available to the user:




4. After the feed has been pulled from the RD Web Access server, RemoteApp and Desktop Connections is set up. The following example shows that two RemoteApp programs are available to the user (these programs are Word and Excel as you can see by referring back to Figure 4-14):



5. Returning to the main page of the RemoteApp and Desktop Connections Control Panel item displays additional information concerning the available connections:



6. The list of available connections is automatically updated on the user's computer on a periodic basis. However, the user can also force the list of connections to update by clicking Properties. This opens the View Properties Or Update This Connection page, and clicking Update Now forces the user's computer to pull a new feed from the RD Web Access server immediately:



For more information concerning RemoteApp and Desktop Connections, see the "Additional Resources" section at the end of the chapter.

Understanding Remote Desktop Connection Broker

The Remote Desktop Connection Broker (RD Connection Broker) role service in Windows Server 2008 R2 was formerly called Terminal Services Session Broker (TS Session Broker) in Windows Server 2008. This service was first introduced in Windows Server 2003 as the Terminal Server Session Directory service, and it allowed disconnected Terminal Services users to reconnect to a previously established terminal server session. The feature was renamed TS Session Broker in Windows Server 2008 and was enhanced with the addition of TS Session Broker Load Balancing, which enables load-balancing of terminal servers in a terminal server farm and allows disconnected users to reconnect to an existing session in a load-balanced

terminal server farm. Previously, such load-balancing support was available only through third-party software and through Windows Network Load Balancing (NLB).

RD Connection Broker further extends the capabilities of this feature by providing a unified administrative experience for traditional session-based remote desktops and for newer virtual machine–based remote desktops, which can be implemented using the new Remote Desktop Virtualization Host role service of Windows Server 2008 R2. RD Connection Broker also provides additional support for RemoteApp Centralized Publishing through RemoteApp and Desktop Connections (described in the previous section) and for Remote Desktop Connection Manager (described later).

How RD Connection Broker Works

Similar to the TS Session Broker role service of Windows Server 2008, the RD Connection Broker role service allows a user to reconnect to an existing session of a load-balanced RD Session Host server farm. RD Connection Broker does this by storing session state information, including session IDs and their associated user names, and the name of the RD Session Host server where each session resides. When a user having an existing session connects to an RD Session Host server in a load-balanced farm, the RD Connection Broker server redirects the user to the RD Session Host server where the user's session resides, preventing the user from being connected to a different server in the farm and thus having to start a new session.

And if you enable the RD Connection Broker Load Balancing feature, the RD Connection Broker server can also do the following:

- Evenly distribute the session load between RD Session Host servers in a load-balanced RD Session Host server farm.
- Track the number of user sessions on each RD Session Host server in the farm.
- Redirect users who don't have an existing session to the RD Session Host server that has the fewest sessions.

Installing and Configuring RD Connection Broker

You can install the RD Connection Broker role service on the Standard, Enterprise, or Datacenter edition of Windows Server 2008 R2 using any of the methods described in the section titled "Installing an RD Session Host Server" earlier in this chapter. After you have installed the role service, you can configure it using the Remote Desktop Connection Manager snap-in. (See Figure 4-16.)




FIGURE 4-16 The Remote Desktop Connection Manager snap-in is used to configure and manage the RD Connection Broker server.


Using the Remote Desktop Connection Manager snap-in, you can perform various tasks on your RD Connection Broker server, such as adding RemoteApp sources (RD Session Host servers), adding RD Virtualization Host servers, creating virtual desktop pools, adding virtual machines to a pool, configuring virtual desktops, assigning personal desktops to users, and performing other tasks. Some of these tasks are described in more detail in the section titled “Understanding Remote Desktop Virtualization Host” later in this chapter.

An example of a configuration task you can perform using this snap-in is to configure the RemoteApp and Desktop Connection Web feed by creating a Workspace Configuration (.wcx) file you can distribute to users of Windows 7 computers on your network. Doing this allows the client side of RemoteApp and Desktop Connection to be configured on users’ computers without the need for users to manually configure the RemoteApp and Desktop Connection Control Panel item. Performing this task requires that you be a member of the local Administrators group on the RD Connection Broker server that you plan to configure.

To create the configuration file, follow these steps:

1. Open the Remote Desktop Connection Manager snap-in on your RD Connection Broker server by clicking Start, Administrative Tools, Remote Desktop Services, Remote Desktop Connection Manager.

2. In the console tree pane on the left, click Remote Desktop Connection Manager: `<server_name>`, where `<server_name>` is the name of your RD Connection Broker server.
3. In the Actions pane on the right, click Create Configuration File.
4. In the Create Configuration File dialog, enter the RemoteApp and Desktop Connection URL that specifies the RD Web Access server that provides the RemoteApp and Desktop Connection resources to the targeted users. Be sure to specify the FQDN of the RD Web Access server:



5. Click Save, specify a file name, select a folder location, and click Save.
6. Distribute the configuration file to the targeted users using an appropriate method such as Group Policy as described in the section titled "Configuring RemoteApp and Desktop Connections" earlier in this chapter.

Another configuration task involving RD Connection Broker servers is configuring security groups. Specifically,

- For an RD Web Access server to provide RemoteApp and Desktop Connection information from an RD Connection Broker server, you must add the computer account for the RD Web Access server to the RD Web Access Computers security group on the RD Connection Broker server. You must be a member of the local Administrators group on the RD Connection Broker server to do this.
- For an RD Session Host server to provide redirection to virtual desktops, you must add the computer account for the RD Session Host server to the Session Broker Computers security group on the RD Connection Broker server. And if you have deployed a load-balanced RD Session Host server farm to provide RemoteApp programs to users through RemoteApp and Desktop Connection, you must add the computer account for each RD Session Host server in the farm to the Session Broker Computers security group.

Finally, to configure the RD Connection Broker Load Balancing feature of an RD Connection Broker server, you must use Remote Desktop Session Host Configuration snap-in described earlier in this chapter.

Understanding Remote Desktop Gateway

The Remote Desktop Gateway (RD Gateway) role service in Windows Server 2008 R2 was formerly called Terminal Services Gateway (TS Gateway) in Windows Server 2008. Installing the RD Gateway role service enables users to securely connect over the Internet to RD Session Host servers or RD Virtualization Host behind the corporate firewall to access session-based desktops or virtual desktops, run RemoteApp programs, and access client computers that have Remote Desktop enabled. RD Gateway provides an alternative to virtual private network (VPN) connections as a means for secure access to remote desktops and applications on an organization's internal network.

RD Gateway enables authorized remote users to connect to resources on an internal corporate network over the Internet using the Remote Desktop Protocol (RDP). The resources that external users can connect to via RD Gateway include

- RD Session Host servers running RemoteApp programs and session-based desktops
- RD Virtualization Host servers running virtual desktops
- Client computers that have Remote Desktop enabled

By simplifying the task of enabling secure access to an organization's internal network from over the Internet, RD Gateway provides an alternative to deploying VPN servers and configuring VPN client software on users' computers. RD Gateway also allows organizations that block VPN connections at the firewall to still allow users to remotely connect to the corporate network from home or when traveling.

RD Gateway enhances security by allowing you to place RD Session Host and RD Virtualization Host servers needed by external users for access inside the corporate network instead of on the perimeter network. With Windows Server 2003 Terminal Services, allowing external users RDP connectivity with terminal servers required that the terminal servers reside on the perimeter network. This configuration potentially exposed these terminal servers to attack from outside the corporate network, an unsatisfactory situation for most organizations. With RD Gateway, however, you can safely place your RD Session Host and RD Virtualization Host servers inside the corporate network; only the RD Gateway server itself needs to reside on a screened subnet of the perimeter network. This means that only the RD Gateway server is directly exposed to outside attack. And the attack surface of the RD Gateway server is lower than that of an RD Session Host and RD Virtualization Host server placed in a similar location because the only external port that needs to be open on the RD Gateway server is TCP port 443.

RD Gateway also enhances security by providing a point-to-point RDP connection between the remote client and the internal terminal server or a computer having Remote Desktop enabled. RD Gateway thus allows remote users to access anything on the corporate network that allows RDP access (provided they have appropriate privileges for doing so), including networks that are hidden behind firewalls or must be accessed across a Network Address Translation (NAT) device.

RD Gateway can also be implemented together with RemoteApp and RD Web Access to allow authorized users to connect over the Internet to an RD Session Host and RD Virtualization Host server on the corporate network and run RemoteApp programs or access session-based desktops or virtual desktops. A remote user at home or on the road can also use RD Gateway together with Remote Desktop Web Connection, a feature of RD Web Access, to securely connect to the desktop of her computer on the corporate network as if she is sitting at her desk in the office.

New features of RD Gateway in Windows Server 2008 R2 include the following:

- **Configurable idle and session timeouts** Allows you to configure idle and session timeouts on the RD Gateway server to reclaim resources used by inactive user sessions and to periodically enforce new policies on active user connections. Disconnected users can reestablish their sessions by using Remote Desktop Connection.
- **Background session authentication and authorization** Allows a remote session to either be disconnected or silently re-authenticated and reauthorized when the session timeout has been reached. Selecting the option Silently Re-authenticate And Reauthorize After Session, enables session authentication and authorization to happen in the background, and sessions for users whose property information has not changed are not affected.
- **Device redirection enforcement** Allows you to configure remote desktop clients so that they can connect only to RD Session Host servers that enforce device redirection. This feature requires RDC 7.0 and an RD Session Host server running Windows Server 2008 R2.
- **Network Access Protection (NAP) remediation** Allows you to update client computers that are not in compliance with your network's health policy to ensure managed clients are in compliance by having the latest software updates. You can also configure Connection Authorization Policy (CAP) policies so that unmanaged clients do not receive updates; instead, feedback is provided to these users, allowing them to manually update their systems.
- **Pluggable authentication and authorization** Provides APIs that can be used to write authentication and authorization plug-ins that integrate with RD Gateway.

- **System and logon messages** Allows you to configure and display system messages that inform users about server maintenance issues such as shutdown and restarts, and logon messages such as logon notices when users try to access remote resources. Remote desktop clients must be running RDC 7.0 to use this feature.

How RD Gateway Works

In Windows Server 2003 Terminal Services, RDP connectivity from outside the corporate network required that TCP port 3389 be left open in the organization's perimeter firewall. For security reasons, however, many businesses choose to keep this port closed to help safeguard computers on the internal network from unauthorized access. RD Gateway overcomes this limitation by transmitting RDP traffic over TCP port 443 using an encrypted HTTPS—Secure Sockets Layer/Transport Layer Security (SSL/TLS) over HTTP—tunnel. RD Gateway thus requires only that TCP port 443, the SSL/TLS port, be left open on perimeter firewalls, which is usually the case in most corporate environments.

Figure 4-17 shows a basic RD Gateway scenario where home office workers use RD Gateway to remotely connect to terminal servers and their company computers over the Internet.




FIGURE 4-17 How RD Gateway works for the home office usage scenario.

Installing and Configuring RD Gateway

At a high level, installing and configuring an RD Gateway server involves the following steps:

1. Install the RD Gateway role service on a Windows Server 2008 R2 server using any of the methods described in the section titled “Installing an RD Session Host Server” earlier in this chapter.
2. Obtain an SSL-compatible X.509 certificate, and install it on your new RD Gateway server. This certificate can be issued by a trusted public certification authority (CA) that participates in the Microsoft Root Certificate Program Members program, or you can use a certificate issued by Microsoft Certificate Services running on a server in your corporate network or even a self-signed certificate that has been exported and imported into the computer certificate store of the client computers that will need to access the RD Gateway server from outside the corporate network.
3. Create a Remote Desktop Connection Authorization Policy (RD CAP) on your RD Gateway server. This RD CAP allows administrators to specify the connection criteria that must be met to connect to the RD Gateway server. For example, a simple RD CAP might require a connecting user’s account to be a member of a specific security group or it might require the user’s computer to have a computer account that belongs to a specific group. The RD CAP could also require that the user use a smart card to connect through the RD Gateway. Users and computers are granted access to the RD Gateway server if they meet the conditions specified in the RD CAP. RD CAPs simplify administration and enhance security by providing administrators with greater control over which users and computers are allowed to access resources on the internal network. RD CAPs can also be used to control how device redirection is handled when users remotely connect to terminal servers. And you can create multiple RD CAPs to define different connectivity conditions for different groups of users and computers.
4. Create a Remote Desktop Resource Authorization Policy (RD RAP) on your RD Gateway server. This RD RAP allows administrators to specify the internal resources (RD Session Host and RD Virtualization Host servers and client computers that have Remote Desktop enabled on them) that remote users will be allowed to connect to through the RD Gateway server. When you create an RD RAP, you can create a security group containing computer accounts and associate this group with the RD RAP. For example, an RD RAP might specify that users who are members of the HR Users security group are allowed to connect only to computers that are members of the HR Computers group. Then you could create a second RD RAP that specifies that users who are members of the Finance Users group are allowed to connect only to computers that are members of the Finance Computers group. Remote users using RD Gateway to connect to the internal corporate network are then granted access to computers on the network only if they meet the conditions specified in at least one RD CAP and one RD RAP.

5. After you've installed the RD Gateway role service, obtained and installed an SSL certificate on your server, and configured at least one RD CAP and one RD RAP, the last step you need to perform is to configure the RDC client software on users' computers to use your RD Gateway server. For small-scale deployments, client configuration can be done manually through the Settings button on the Advanced tab of the Remote Desktop Connection dialog box. Clicking this button displays the RD Gateway Server Settings dialog box:



For larger deployments, configuration can be done using Group Policy with the settings found at the following location:

User Configuration\Policies\Administrative Templates\Windows Components\Remote Desktop Services\RD Gateway

6. After your RDC clients have been properly configured, remote users can begin to use RD Gateway to access resources on your organization's internal network. For example, users will be able to access internal servers from over the Internet to run RemoteApp programs installed and published on RD Session Host servers, access session-based desktops running on RD Session Host servers, or access virtual desktops running on RD Virtualization Host servers. Users will also be able to access the desktop of their own computers on the internal corporate network if Remote Desktop has been enabled on these computers.

7. At this point, you can use the RD Gateway Manager snap-in to further configure and manage your RD Gateway server. For example, you can limit the maximum number of simultaneous connections allowed through your RD Gateway server to optimize server performance or ensure compliance with your organization's security policies. And you can use RD Gateway Manager to view information about active connections, including the domain and user ID of the user logged on to the client, IP address of the client, name of the target computer to which the client is connected, target port through which the client is connected, date and time when the connection was initiated, length of time that the connection is idle, connection duration, and amount of data (in kilobytes) that was sent and received by the client through the RD Gateway server. You can also specify which types of events to monitor for the RD Gateway server, such as successful or unsuccessful connection attempts. These events will be displayed in Event Viewer under Application and Services Logs\Microsoft\Windows\TerminalServices-Gateway.

Understanding Remote Desktop Licensing

The Remote Desktop Licensing (RD Licensing) role service in Windows Server 2008 R2 was formerly called Terminal Services Licensing (TS Licensing) in Windows Server 2008. RD Licensing allows you to manage Remote Desktop Services client access licenses (RDS CALs) for users and devices that connect to your terminal servers. The RD Licensing role service supports managing RDS CALs for Remote Desktop Session Host servers running Windows Server 2008, Windows Server 2003, and Windows 2000 Server.

Microsoft licensing requirements require that each user or device that connects to the Remote Desktop Session Host server have its own RDS CAL to establish a valid connection with the server. These Remote Desktop Services licenses are separate from and in addition to the valid Windows Server 2008 R2 licenses and Windows Server CALs required on a Windows-based network.

By installing the RD Licensing role service on a Windows Server 2008 R2 server, administrators can easily deploy, manage, and revoke RDS CALs in a Remote Desktop Services environment. RD Licensing in Windows Server 2008 R2 includes support for per-user tracking and reporting, manual revocation of licenses, improved diagnostics, and a Windows Management Instrumentation (WMI) provider.

RD licensing can be configured either during the installation of the role or afterward. Figure 4-18 shows the licensing configuration options in the Add Roles Wizard.




FIGURE 4-18 Installing and configuring the RD Licensing role service.

Remote Desktop Services can be licensed either per device or per user as desired. In addition, if you select the Configure Later option during role installation, you have a 120-day grace period before an RD license server (a server running the RD Licensing role service) must be deployed on your network. This grace period allows you to configure and test your Remote Desktop Services deployment before rolling it out to your production environment.

For tips on planning licensing for your Remote Desktop Services deployment, see the sidebar titled “Direct from the Source: Tips for Planning Remote Desktop Services Licensing” later in this chapter.

New features of RD Licensing in Windows Server 2008 R2 include the following:

- **Removal of support for automatic license server discovery for RD Session Host servers** For Windows Server 2008 R2, you must now use Remote Desktop Session Host Configuration to specify the name of the license server an RD Session Host or RD Virtualization Host server should use. For earlier versions of Windows Server—including Windows Server 2008, Windows Server 2003, and Windows 2000 Server—you could specify a discovery scope when you installed the RD Licensing role service, where the

discovery scope determined how the license servers were automatically discoverable by terminal servers running these earlier operating systems.

- **Changes to the Licensing tab in Remote Desktop Session Host Configuration** In Remote Desktop Session Host Configuration in Windows Server 2008 R2, you must specify a license server for the RD Session Host server to use. You can either choose from a list of known license servers or manually enter the name. License servers that are registered as a service connection point (SCP) in Active Directory Domain Services (AD DS) will appear in the list of known license servers in Remote Desktop Session Host Configuration. You can add more than one license server for the RD Session Host server to use. If more than one license server is added, the RD Session Host server contacts the license servers in the order in which they appear in the Specified License Servers area on the Licensing tab in Remote Desktop Session Host Configuration.
- **Manage RDS CALs Wizard** Windows Server 2008 R2 now includes a new Manage RDS CALs Wizard in Remote Desktop Licensing Manager that allows you to easily migrate RDS CALs from one license server to another. Migrating RDS CALs can be useful when you need to replace a license server on your network. If you migrate RDS CALs from a license server running a version of Windows Server earlier than Windows Server 2008 R2, you need to manually delete the RDS CALs from the original license server after you have migrated the RDS CALs to the new server. The Manage RDS CALs Wizard also allows you to rebuild the RD Licensing database on your licensing server, but be aware that doing this deletes any RDS CALs currently installed on the license server. This means you need to reinstall the RDS CALs on the license server after the licensing database has been rebuilt.
- **Service Connection Point registration** When you install the RD Licensing role service using Server Manager, the license server tries to register itself as an SCP in AD DS. This causes the license server to appear in the list of known license servers in Remote Desktop Session Host Configuration. If AD DS is not available when you install the RD Licensing role service, you can still manually register the license server using Remote Desktop Licensing Manager.
- **Single RDS CAL pack** In versions of Windows Server earlier than Windows Server 2008 R2, you could only purchase RDS CALs in packs of 5 or 20. Beginning with Windows Server 2008 R2, however, you can now purchase single RDS CALs.

Direct from the Source: Tips for Planning Remote Desktop Services Licensing

For environments with a single Remote Desktop Services license server, if the license server becomes unavailable, previously licensed clients will still be able to connect to the RD Session Host servers if the clients' licenses have not expired. However, unlicensed clients will not be able to connect to the RD Session Host servers because they cannot obtain a license. To provide a window to bring the license server back online without disrupting users, a redundant license server can be used.

RD licensing server redundancy can be achieved using either of the following methods:

- Install all RDS CALs on the primary license server and none on the secondary license server. This approach provides a period of redundancy in the event of a failure of the primary license server. If you add a second license server with no licenses, unlicensed clients will be able to connect and obtain a temporary license from this server if the primary license server is not available. Previously licensed clients can also still connect if their license has not expired.
- Install RDS CALs on both the primary license server and the secondary license server. Ideally, each RD licensing server should contain 50 percent of the available RDS CALs. If the primary licensing server is not available or does not have valid RDS CALs, clients are redirected to the secondary licensing server for license issuance.

Because licensed clients attempt to renew their license seven days prior to expiration, these redundant configurations help to ensure uninterrupted RD Session Host server connectivity for up to seven days if one of the license servers fails.

It is also recommended that you regularly back up Remote Desktop Services licensing data to protect the data from accidental loss because of hardware or storage failure. In the event that the original data on the hard disk is inaccessible for any reason, the licensing data can be restored from the archived copy. For additional security, a redundant license server can be used.

—CSS Global Technical Readiness (GTR) team

Understanding Remote Desktop Virtualization Host

Remote Desktop Virtualization Host (RD Virtualization Host) is a completely new role service of the Remote Desktop Services role in Windows Server 2008 R2. Similar to how the RD

Session Host role service allows users to run RemoteApp programs and access session-based desktops, the RD Virtualization Host service allows them to access virtual desktops (the desktops of virtual machines) through Remote Desktop Services. These virtual machines that users access using the RD Virtualization Host role service must be running on a server running Windows Server 2008 R2 that has the Hyper-V server role installed.

RD Virtualization Host is an important component of Microsoft's Virtual Desktop Infrastructure (VDI), which is described in the section titled "Understanding Microsoft Virtual Desktop Infrastructure" later in this chapter.

How RD Virtualization Host Works

The RD Virtualization Host role service integrates with the Hyper-V server role of Windows Server 2008 R2 and must be installed on a server that has the Hyper-V server role installed. This means that all the requirements of running Hyper-V, such as support for hardware virtualization, must be met by servers you want to install the RD Virtualization Host role service on. For more information about the requirements for installing the Hyper-V server role, see Chapter 2, "Server Virtualization."

You can install the RD Virtualization Host role service on a Windows Server 2008 R2 server using any of the methods described in the section titled "Installing an RD Session Host Server" earlier in this chapter. After you've installed the role service, you can use your RD Virtualization Host server to make virtual desktops available to your users in two forms:

- **Personal virtual desktops** In this scenario, you assign a single virtual machine to a single domain user account. Only one virtual machine can be assigned to each user, and only that particular user can remotely access that specific virtual machine. Each time the user establishes a Remote Desktop connection to the RD Virtualization Host server, the user is connected to the same virtual machine. Any customizations the user makes to her virtual desktop during a session are saved so that they can be available the next time the user accesses her virtual machine.
- **Virtual desktop pools** In this scenario, you begin by creating a pool of identically configured virtual machines. These virtual machines must have the same operating system and applications installed, the same service packs and updates applied, the same configuration settings, and so on. These virtual machines must also not have already been assigned to users as personal virtual desktops. When the user establishes a Remote Desktop Connection to the RD Virtualization Host server, he is connected to any of the virtual machines in the pool. Because all virtual machines in the pool are configured identically, the user experience is the same regardless of which virtual desktop he connects to. By default in this scenario, any customizations the user makes to

his virtual desktop during a session are not saved and are discarded when the user logs off of his Remote Desktop session. However, by combining roaming user profiles with Folder Redirection and storing the roaming profiles and redirected folders on a separate server, any changes that the user makes to his virtual desktop can be saved so that they will be available the next time the user accesses a virtual desktop from the pool.

With personal virtual desktops, each user has exclusive use of a single, dedicated virtual machine. The user can also have Administrator access to her virtual machine and is thus able to configure any aspect of her virtual machine. With virtual desktop pools on the other hand, users share access to a collection of identically configured virtual machines, and they do not have Administrator access to these virtual machines.

Personal virtual desktops and virtual desktop pools can also be provisioned to users in one of two possible ways:

- By using RemoteApp and Desktop Connection as described in the section titled “Configuring RemoteApp and Desktop Connections” found earlier in this chapter.
- By using RD Web Access as described in the section titled “Installing, Configuring, and Using RD Web Access” found earlier in this chapter.

For additional information on how the RD Virtualization Host role service works, see the sidebar titled “Direct from the Source: How RD Virtualization Host Works.”

The section following this one demonstrates how to provision personal virtual desktops to users by using RD Web Access. For full walkthroughs of each of these four possible scenarios, see the following Step-by-Step guides in the TechNet Library:

- Deploying Personal Virtual Desktops by Using RemoteApp and Desktop Connection Step-by-Step Guide, found at <http://go.microsoft.com/fwlink/?LinkId=154801>.
- Deploying Virtual Desktop Pools by Using RemoteApp and Desktop Connection Step-by-Step Guide, found at <http://go.microsoft.com/fwlink/?LinkId=154802>.
- Deploying Personal Virtual Desktops by Using Remote Desktop Web Access Step-by-Step Guide, found at <http://go.microsoft.com/fwlink/?LinkId=147909>.
- Deploying Virtual Desktop Pools by Using Remote Desktop Web Access Step-by-Step Guide, found at <http://go.microsoft.com/fwlink/?LinkId=147906>.

Direct from the Source: How RD Virtualization Host Works

The components needed for remote desktop virtualization include the following:

- **RD Virtualization Host server** The virtual machine (VM) images are hosted here. A service, the VM Host Agent, runs on this machine. This service is controlled by the RD Connection Broker server and can perform certain actions such as spinning up a VM image.
- **RD Connection Broker server** Given an authenticated user and a request for a particular application or desktop (for example, Microsoft Office Excel or “my desktop”), the RD Connection Broker server determines which RD Virtualization Host server or VM image can best satisfy the RDS client’s request.
- **Redirector** This component is an RD Session Host server whose purpose is to query the RD Connection Broker server on behalf of the RDS client. After querying the RD Connection Broker server, the redirector sends an RDP redirection packet back to the RDS client.
- **RD Web Access server** This component lets an end user view a Web page that shows the user all the applications and desktops he can access.
- **RDS Assignment Database** This is a representation of the Active Directory Schema extensions that provide user mappings to a particular RD Virtualization Host server image.

Role of the RD Connection Broker Server


Remote desktop virtualization uses an RD Connection Broker to determine where the user is redirected. If a user is assigned and requests a personal virtual desktop, the RD Connection Broker server redirects the user to this virtual machine. If the virtual machine is not turned on, the RD Virtualization Host server turns on the virtual machine and then connects the user. If the user is connecting to a shared virtual machine pool, the RD Connection Broker server first checks to see if the user has a disconnected session in the pool. If the user has a disconnected session, she is reconnected to that virtual machine. If the user does not have a disconnected session, a virtual machine in that pool is dynamically assigned to the user, if one is available.

—CSS Global Technical Readiness (GTR) team

Provisioning Personal Virtual Desktops Using RD Web Access

To demonstrate some of the capabilities of the new RD Virtualization Host role service, let's walk through the steps of assigning a personal virtual desktop to a user and enabling the user to access this virtual desktop using an RD Web Access Web site on the corporate network.

1. Begin by installing the Remote Desktop Virtualization Host role service on a server running Windows Server 2008 R2 that already has the Hyper-V role installed on it. You can install the RD Virtualization Host role service on the Standard, Enterprise, or Datacenter edition of Windows Server 2008 R2 using any of the methods described earlier in the section titled "Installing an RD Session Host Server." For example, you can launch the Add Roles Wizard from Server Manager and use this wizard to add the Remote Desktop Virtualization Host role service to your server:



2. Make sure you also have an RD Session Host server, RD Web Access server, and RD Connection Broker server deployed on your network. These three servers perform the following functions:
 - The RD Session Host server is used to provide redirection for virtual desktops. By default, an RD Session Host server uses IP address redirection, which means that a client that queries the RD Connection Broker is redirected to its existing session


by using the IP address of the server where its session resides. When you install the RD Session Host role service on a server for this purpose, make sure that none of the check boxes are selected on the Configure Client Experience page of the Add Roles Wizard. Also, don't install any RemoteApp programs on the server because users will not be able to access any RemoteApp programs on this server or remotely connect session-based desktops on the server. In other words, the RD Session Host server you use to provide redirection for virtual desktops should be used only for this purpose and not to host RemoteApp programs or session-based desktops.

- The RD Web Access server is used to make virtual desktops running on the RD Virtualization Host server available to users on your network.
- The RD Connection Broker server is used for coordinating the availability of virtual desktops for users and making sure a virtual machine is running before a user connection is established. Note that you need an RD Connection Broker server only if you are using RD Web Access to provision virtual desktops to users as you are doing here; if instead you are using RemoteApp and Desktop Connections to provision virtual desktops, you don't require an RD Connection Broker server.

All three of these servers, plus the RD Virtualization Host server you deployed in the previous step, must be joined to your AD DS domain.

3. On your RD Virtualization Host server, which in this example is named HVR2.contoso.com, create a new virtual machine using the Hyper-V Manager snap-in. For information on how to do this, see the section titled "Creating a Virtual Machine" in Chapter 2. In this walkthrough, you assign the new virtual machine the name SEA-CLI-789.
4. Install a Windows client operating system such as Windows 7 Enterprise edition in your new virtual machine. Note that you cannot provision Windows Server operating systems such as Windows Server 2008 R2 by using RD Virtualization Host—the role service is designed to provision only Windows client operating systems to users.
5. After the OOBE appears after the final reboot during the install, assign the same name to the guest operating system as the virtual machine itself has for its name, which in this example is SEA-CLI-789.
6. Now log on to the guest operating system of the virtual machine using a user account that is a member of the local Administrators group on the machine. Then join the virtual machine to the domain.

7. Next, log on again using an account that has administrative privileges and enable Remote Desktop in the guest operating system of the virtual machine:



8. Click Select Users, and add the domain user who will be assigned this personal virtual desktop to the local Remote Desktop Users group on the machine. In this example, the user who will be assigned this personal virtual desktop is user Karen Berg (CONTOSO\kberg).
9. The next thing you need to do is enable Remote RPC for Remote Desktop Services. To configure this, open Registry Editor and navigate to the following key:
HKLM\System\CurrentControlSet\Control\Terminal Server
Under this registry key, change the value of the AllowRemoteRPC registry entry from 0 to 1, and then close Registry Editor.
10. Open Windows Firewall from Control Panel, and allow the Remote Service Management firewall exception for the domain firewall profile.

11. Next, open a command prompt, with the Run As Administrator option, and type the following series of commands (modify these commands if your RD Virtualization Host server has a name other than HVR2 and if your domain is not CONTOSO):

```
wmic /node:localhost RDPERMISSIONS where TerminalName="RDP-Tcp" CALL  
AddAccount "contoso\HVR2$",1
```

```
wmic /node:localhost RDACCOUNT where "(TerminalName='RDP-Tcp' or  
TerminalName='Console') and AccountName='contoso\\HVR2$'" CALL  
ModifyPermissions 0,1
```

```
wmic /node:localhost RDACCOUNT where "(TerminalName='RDP-Tcp' or  
TerminalName='Console') and AccountName='contoso\\HVR2$'" CALL  
ModifyPermissions 2,1
```

```
wmic /node:localhost RDACCOUNT where "(TerminalName='RDP-Tcp' or  
TerminalName='Console') and AccountName='contoso\\HVR2$'" CALL  
ModifyPermissions 9,1
```

The new WMI provider for Remote Desktop Services included in Windows Server 2008 R2 provides full read/write capabilities for configuring nearly all Remote Desktop Services server settings, and the commands previously listed use the Windows Management Instrumentation Command-line (WMIC) interface to add RDP protocol permissions to the virtual machine. You can find more information about the RDS WMI provider on MSDN at <http://msdn.microsoft.com/en-us/library/aa383515.aspx>.


12. After you've add RDP protocol permissions to the virtual machine, restart Remote Desktop Services to make the changes take effect. You can restart Remote Desktop Services from the command line by using these commands:

```
net stop termervice
```


```
net start termervice
```

13. Log off of the guest operating system on the virtual machine. If you like, you can also shut down or hibernate the virtual machine because after the virtual machine has been provisioned to Karen she will be able to start it using RPC when she tries to access it.
14. At this point, the virtual machine is ready to be provisioned to Karen as her personal virtual desktop. Before you can do this, however, you need to configure your RD Web Access server to work together with your RD Connection Broker server. To do this, begin by adding the computer account for your RD Web Access server to the RD Web Access Computers security group on your RD Connection Broker server.


15. Next, log on to the RD Web Access Web site using an account that has administrative privileges for the domain.
16. On the Configuration tab of the RD Web Access Web site, select An RD Connection Broker Server as the source your RD Web Access server will use and then type the name of your RD Connection Broker server:




17. Now you are ready to configure the personal virtual desktop on the RD Connection Broker server and assign the personal virtual desktop to user Karen Berg. To do this, begin by logging on to your RD Connection Broker server using an account that has administrative privileges for your domain. After you have logged on, click Start, Administrative Tools, Remote Desktop Services, Remote Desktop Connection Manager. Then with the Remote Desktop Connection Manager snap-in open, click Configure Virtual Desktops Wizard in the Actions pane to launch the Configure Virtual Desktops Wizard:




18. On the Specify An RD Virtualization Host Server page, type the NetBIOS name or FQDN of your RD Virtualization Host server (which is HVR2.contoso.com in this example) and click Add:




19. On the Configure Redirection Settings page, type the NetBIOS name or FQDN of the RD Session Host server that will be used to provide redirection to the virtual desktops:



20. On the Specify An RD Web Access Server page, your RD Web Access server should already be specified because you previously used the RD Web Access Web site to specify your RD Connection Broker server as the source for your RD Web Access server:




21. The final page of the wizard summarizes the changes that will be made:




Leave the Assign Personal Virtual Desktop check box selected as shown in the preceding screen shot, as doing this will automatically launch the Assign Personal Virtual Desktop Wizard after you've clicked Finish.


22. When the Assign Personal Virtual Desktop Wizard opens, select user Karen Berg from AD DS and then select the virtual desktop you want to assign to her from the Virtual Machine list box:




23. Alternatively, if your AD DS domain functional level has already been raised to Windows Server 2008 R2, you will also be able to assign personal virtual desktops to users by using the new Personal Virtual Desktop tab on the properties sheet for the user's account:




24. If you want to assign additional personal virtual desktops to other users in your organization, you can use the Remote Desktop Connection Manager snap-in on your RD Connection Broker server to do this:




25. Now user Karen Berg can at last connect to her personal virtual desktop. To do this, she logs on to her physical computer, opens Internet Explorer, types **https://SEA-RDS4.contoso.com/RDWeb** in the address bar, and presses Enter. This connects her to the RD Web Access server (which is SEA-RDS4 in this example) and displays the logon page, where she now enters her credentials:



26. After Karen has logged on to the RD Web Access Web site, she sees a My Desktop item on the RemoteApp Programs tab of the page:



27. Karen now double-clicks on the My Desktop icon and the virtual machine associated with her personal virtual desktop spins up. After a short time—the length of which depends on whether the virtual machine was stopped, hibernated, or running—a Remote Desktop Services connection is established with the virtual machine and the virtual desktop is displayed, either in full-screen mode or as shown in the following screen shot as a separate window on her physical computer's desktop.



For more information on the RD Virtualization Host role service, see the “Additional Resources” section at the end of this chapter.

Deploying Remote Desktop Services

Depending on your organization’s needs, you might decide to deploy one, some, or all Remote Desktop Services role services on one or more Windows Server 2008 R2 servers in your environment. Figure 4-19 illustrates a simple Remote Desktop Services deployment where all of the role services are installed on individual servers in an AD DS domain.




FIGURE 4-19 An example of a Remote Desktop Services deployment.

The interactions between the various role services in this example include the following:

- A remote client connecting over the public Internet is using the RD Gateway server to access an RD Session Host and an RD Virtualization Host located inside the corporate network.
- The RD Connection Broker server connects clients to sessions (RemoteApp programs and session-based desktops) on the RD Session Host server and to virtual machines (virtual desktops) on the RD Virtualization Host server.
- The various Remote Desktop Services servers are validated by the RD Licensing server.

If needed, more complex Remote Desktop Services infrastructures can be deployed that include load-balanced farms of RD Session Hosts and RD Virtualization Hosts. For more information, see Remote Desktop Services online help at <http://technet.microsoft.com/en-us/library/cc770412.aspx>.

Understanding Microsoft Application Virtualization for Remote Desktop Services

In Chapter 3, “Local Desktop Virtualization,” you learned about Microsoft Application Virtualization (App-V), a technology that allows you to transform applications into centrally managed virtual services by sequencing (virtualizing) them for delivery across the network. This allows you to reduce the cost of application deployment, eliminate application conflicts, simplify the base image footprint to expedite provisioning of client computers, and generally increase user productivity. By using App-V 4.5 SP1 (the version available at the time of writing this chapter), you can deploy sequenced applications to users in two different ways:

- **Directly to their client computers** By installing the App-V desktop client software onto client computers running 32-bit versions of Windows 7, Windows Vista, or Windows XP, you can deliver applications to users of those computers without having to install the applications locally on their computers. The App-V 4.5 SP1 desktop client is available as part of the Microsoft Desktop Optimization Pack (MDOP) for Software Assurance (SA) for volume-licensed customers.
- **Via RD Session Hosts/Terminal Servers** By installing the Microsoft Application Virtualization for Remote Desktop Services (App-V for RDS) client software on terminal server computers running 32-bit versions of Windows Server 2008 or Windows Server 2003, you can deliver applications to Terminal Services users without having to install the applications locally on their computers. As indicated previously in the section titled “Availability of App-V for RDS” in this chapter, 64-bit App-V for RDS client software that can be installed on Windows Server 2008 R2 is expected to be released in the first half of 2010 as part of App-V 4.6, the next version of the product.

The first client software, the App-V desktop client, lets you realize the benefits of App-V in a standard AD DS networking environment and is typically deployed to resolve side-by-side application compatibility issues that have been blocking your desktop computer migration plans. The second client software, App-V for RDS, lets you realize the benefits of App-V for session-based desktops running on RD Session Hosts/Terminal Servers and helps you reduce server sprawl and eliminate RD Session Host/Terminal Server silos. Many organizations deploy more RD Session Hosts/Terminal Servers than necessary because of application conflicts that require different versions of an application to be separated and run on different servers. The result is often RD Session Hosts/Terminal Servers running at only 25 percent of their capacity or lower. By using App-V for RDS, however, you do not need to separate incompatible applications onto different servers because the applications are no longer installed on the RD Session Host/Terminal Server; instead, they are packaged into managed virtual services that can be streamed to users who need them via Remote Desktop Services sessions.

Figure 4-20 illustrates a basic App-V for RDS infrastructure in which the App-V Sequencer is first used to virtualize Windows applications. An App-V Streaming Server is then used to stream the virtualized applications to an RD Session Host/Terminal Server that has the App-V

for RDS client software installed. Users who establish sessions with the RD Session Host/Terminal Server can then run the applications within Remote Desktop Services sessions—the applications run as if they were locally installed on the users' computers.




FIGURE 4-20 How App-V for RDS works.



More info For more information on implementing App-V for RDS, see the white paper titled "Application Virtualization 4.5 for Terminal Services" available from the Microsoft Download Center at <http://download.microsoft.com/download/6/9/0/69095D7C-649D-4A0E-AF0B-17B26E-ACCF67/App-V%20Terminal%20Services.docx>. See also the section titled "Additional Resources for App-V for RDS" at the end of this chapter.

Understanding Microsoft Virtual Desktop Infrastructure

Whereas Remote Desktop Services is an installable feature of Windows Server platforms and App-V for RDS is a downloadable client application, Microsoft Virtual Desktop Infrastructure (VDI) is much more than a feature or application. Microsoft VDI, in fact, is an architectural model that enables entire desktop operating systems, such as Windows 7 Enterprise edition, to run on a server located in a datacenter.

In a typical VDI deployment, hundreds or even thousands of desktop virtual machines run on a small number of centralized servers using shared storage such as a storage area network (SAN). Users then access these virtual machines over the network (and even over the Internet) using a remote desktop protocol such as Microsoft's RDP or Citrix's Independent Computing Architecture (ICA). A device called a *connection broker* manages the connection between the user's computer and the servers. Users can either be assigned their own virtual desktops or access a shared pool of virtual desktops, in which case they simply use the next virtual desktop that is available when one is needed. With VDI, administrators can dynamically provision virtual desktops as needed, move them across different hardware and storage platforms, back up running and stored virtual machines, reassign user rights to another device in case of endpoint failure, and manage other aspects of the virtualization environment.

The result of implementing VDI is that the organization's entire desktop infrastructure becomes centralized in the datacenter. Desktop operating systems and applications are stored and executed on the server and thus use the computing resources of the server. In other words, the workload is on the server side while the presentation is on the client side. At first glance, this sounds a lot like Terminal Services, a technology that has already been available on Windows Server platforms for over a decade and which is now called *session virtualization* in Windows Server 2008 R2 and is based upon Remote Desktop Services. However, there are some significant differences between VDI and session virtualization.

For instance, in session virtualization multiple users simultaneously access a single desktop (specifically, a Windows Server desktop) or application (a RemoteApp program) and therefore must share the operating system or application. This means that session virtualization users cannot have Administrator access, because this would allow them to perform certain tasks (such as rebooting the operating system or uninstalling the application) that would cause problems for other users connected to the RD Session Host server. In VDI, however, each user has access to an individual desktop (specifically, a Windows client desktop), and users can even be granted exclusive access to their own personal virtual desktops. The result is that VDI users can have (but are not required to have) Administrator access if desired to an entire virtual desktop, which can give them greater flexibility and power. And although session virtualization offers a highly scalable solution for desktop and application virtualization, VDI is ideal for environments where virtual machine-based isolation is an important consideration, because the virtual desktop of one VDI user is completely isolated from that of another user.

You already learned about the concept of personal virtual desktops and virtual desktop pools in the section titled “Understanding Remote Desktop Virtualization Host” earlier in this chapter. And Windows Server 2008 R2 Remote Desktop Services is indeed one of the key components of Microsoft’s VDI architectural model. But there are other components in this model, so let’s examine these now.

Understanding Microsoft’s VDI Architecture

At a high level, Microsoft’s VDI architecture consists of three components:

- **Hardware layer** This layer includes one or more datacenter servers that support hardware virtualization and shared storage such as a SAN, where the virtual machines can be stored. A good place to start for planning the hardware layer of your VDI infrastructure is the Windows Server Marketplace found at <http://www.windowsservercatalog.com/marketplace/>.
- **Client access points** This component includes client computing devices, which can be either rich clients (Windows PCs) or thin clients (Windows terminal devices) connected to the datacenter over an internal private network or even over the Internet.
- **Licensing** There are two types of licensing requirements for implementing a Microsoft VDI solution:
 - **VDI suite licensing** The use rights for technology developed by Microsoft that provides virtualization, management, desktop-delivery, and application-delivery capabilities you can use to deploy a VDI infrastructure within your organization.
 - **Additional licensing** In addition to the use rights for the server and management infrastructure included in the VDI suite, you also need to purchase licenses to run virtual copies of Windows client operating systems on your servers so that your users can legally access the virtual desktops. These licenses are known as Windows Virtual Enterprise Centralized Desktop (Windows VECD).

In addition, Microsoft’s VDI offering can be enhanced by integrating it with third-party products such as Citrix XenDesktop. Before we examine the various VDI suites currently offered by Microsoft, let’s examine the licensing requirements for VDI clients.

Understanding Windows Virtual Enterprise Centralized Desktop

Windows Virtual Enterprise Centralized Desktop (Windows VECD) is a type of license offered by Microsoft that lets you license virtual copies of Windows client operating systems (such as Windows 7) running on servers (such as Windows Server 2008 R2 servers with the Hyper-V server role installed). Windows VECD is a device-based annual subscription, which means the total number of licenses must equal the total number of devices (thick or thin clients) accessing the virtual environment.

Microsoft offers two versions of Windows VECD:

- **Windows VECD for SA** This version is intended for PCs that are already covered under Software Assurance (SA) volume licensing. You can use this type of licensing if your VDI deployment includes blade PCs residing in the datacenter, if you need to be able to remote boot virtual machines from a network storage device, or if you need to deploy virtual machines on portable media.
- **Windows VECD** This version is intended for all other devices, such as PCs not covered under SA, thin clients, and anything else. This version comes bundled with Software Assurance, however, so that you can get more out of your license if needed. You can use this type of licensing if your VDI deployment includes any of the requirements just described and you also need to deliver virtual machines to work-from-home employees who use their own PCs or to contractors or offshore workers who use their own PCs.



Tip For a good explanation of why Windows 7 is the best client operating system for Microsoft VDI solutions, see “Windows 7 with RDP7: Best OS for VDI” on the Remote Desktop Services (Terminal Services) Team Blog at <http://blogs.msdn.com/rds/archive/2009/11/02/windows-7-with-rdp7-best-os-for-vdi.aspx>.

Understanding Microsoft VDI Suites

Microsoft offers two packaged VDI solutions that include server and management platforms and tools together with licensing (as shown in Figure 4-21):

- **VDI Standard suite** This VDI offering is designed to help organizations deploy the basic infrastructure for VDI and includes the following components:
 - Hyper-V Server 2008 R2 as the virtualization host (or you can separately license Windows Server 2008 R2 and use the Hyper-V server role for this purpose)
 - An integrated management suite consisting of System Center Virtual Machine Manager 2008 R2, System Center Operations Manager 2007 R2, and System Center Configuration Manager 2007 R2
 - Microsoft Desktop Optimization Pack (MDOP), which, among other technologies, includes Microsoft Application Virtualization (App-V)
 - Connection Brokering capability through Windows Server 2008 R2 Remote Desktop Services

Note that the use rights for Remote Desktop Services in the Standard VDI suite are restricted. This means that delivery of session-based desktops or RemoteApp programs using RD Session Host servers is not allowed for this VDI offering—RD Session Host servers can be used only for VDI-specific purposes—that is, they can run only in re-direction mode. The Standard VDI suite allows VDI desktop virtualization only where

there is a separate virtual machine image of the operating system for each accessing device.

- **VDI Premium suite** This VDI offering is designed for customers who want additional flexibility in their VDI deployment and includes everything the Standard Suite includes plus the following:
 - ❑ Full (unrestricted) Remote Desktop Services capability, including the option to deploy session-based desktops in addition to VDI desktops.
 - ❑ Microsoft Application Virtualization for Remote Desktop Services (App-V for RDS)

	VDI Standard Suite	VDI Premium Suite
Application Delivery	Microsoft Desktop Optimization Pack for Software Assurance	Microsoft Desktop Optimization Pack for Software Assurance Microsoft RemoteApp
Desktop Delivery	Windows Server 2008 Remote Desktop Services (Restricted to VDI scenario only)	Windows Server 2008 Remote Desktop Services (Unrestricted use rights)
Management Use rights for System Center components restricted to VDI scenario	Microsoft System Center Operations Manager 2007 R2 Microsoft System Center Virtual Machine Manager 2007 R2 Microsoft System Center Configuration Manager 2007 R2	Microsoft System Center Operations Manager 2007 R2 Microsoft System Center Virtual Machine Manager 2007 R2 Microsoft System Center Configuration Manager 2007 R2
Virtualization Platform	Microsoft Hyper-V Server 2008 R2	Microsoft Hyper-V Server 2008 R2

FIGURE 4-21 Comparing the Standard and Premium VDI suites.

Both of Microsoft’s VDI suites are licensed using the same device-based annual subscription model used for Windows VECD, which makes it easier for your organization to budget the implementation of a Microsoft VDI solution. Licensing of System Center applications are also limited in both suites to accessing devices to which the VDI suite license is assigned.

As Figure 4-22 shows, both of these VDI suites can also be deployed within your organization in two different ways:

- **Using an RD Connection Broker server** In this scenario, the connection broker for the VDI infrastructure is a Windows Server 2008 R2 server that has the Remote Desktop

Connection Broker role service installed. This all-Microsoft VDI scenario mainly targets low-complexity deployments as indicated in Table 4-9.

- Using a third-party connection broker such as Citrix XenDesktop This Microsoft+Partner scenario targets high-complexity deployments as indicated in Table 4-9.




FIGURE 4-22 Microsoft VDI solutions can be deployed using either a Microsoft or third-party connection broker.

TABLE 4-9 Examples of Low-Complexity and High-Complexity VDI Deployments

Comparison	Low-complexity environment (Microsoft inbox solution)	High-complexity environment (Microsoft+Partner solution)
Management	Single site Static image placement Basic user profile management	Multiple sites Dynamic image placement Advanced profile management
Client experience	Local area network (LAN) only USB support limited to PnP devices	Local area network/wide area network (LAN/WAN) Broad USB device support

How Microsoft VDI Works

Because Microsoft VDI solutions can be deployed in so many different ways to meet the needs of individual customers, understanding how VDI works can be difficult to grasp at first. To get you started, Figure 4-23 shows a simple Microsoft inbox VDI environment and walks you through the steps of a launching a virtual desktop.



Note In this example, each Remote Desktop Services role service is running on a separate server. This was done only to simplify your understanding of the VDI connection process. In fact, if you don't have any users connecting remotely from over the Internet, you can conceivably run all Remote Desktop Services role services on a single server, though this would be suitable only for very small VDI deployments.




FIGURE 4-23 An example of how a user accesses his virtual desktop in a simple VDI environment.

The steps shown in this diagram are as follows:

1. The user opens Internet Explorer on his computer and logs on to the RD Web Access Web page. In this example, the user works from home and has to use his Internet connection to access the RD Web Access Web server, which is behind (or on) the corporate firewall.
2. The RD Web Access server queries the RD Connection Broker server to retrieve a list of virtual desktops that the user is allowed to use, and this list is then displayed to the user on the RD Web Access Web page. In this example, assume that the RD Virtualization Host has been configured for personal virtual desktops; this means the RD Web Access Web site will publish only one virtual desktop to the user—namely, the virtual desktop that has been assigned to the user.
3. The user clicks on a virtual desktop icon on the RD Web Access server to try and launch the virtual desktop. This causes the Remote Desktop Connection (RDC) client running on the user's computer to establish a Remote Desktop Protocol (RDP) connection with

the RD Gateway server, which is also located behind (or on) the corporate firewall. (All other Remote Desktop Services servers are safely hidden behind the corporate firewall.)

4. After the RD Gateway server has authorized the user's connection attempt, the RD Gateway server forwards the user's connection to the RD Session Host. This RD Session Host is running in redirection mode, which means its sole purpose is to redirect users to their virtual machines.
5. The RD Session Host requests the RD Connection Broker server to orchestrate the requested virtual machine for the user.
6. The RD Connection Broker server queries AD DS to find the virtual machine assigned to the particular user.
7. The RD Connection Broker requests the correct RD Virtualization Host server to orchestrate the delivery of the user's virtual machine. The RD Connection Broker then returns the IP address of the user's virtual machine via a redirection packet to the RD Session Host server, which forwards the packet via the RD Gateway server to the RDC client running on the user's computer.
8. The RDC client now establishes a connection to the virtual machine running on the RD Virtualization Host server via the RD Gateway server, and the user's personal virtual desktop is displayed on the user's computer.

If you find any of the steps confusing, refer back to the section in this chapter that deals with the particular Remote Desktop Services role service or services mentioned in the step.

Additional Resources

The resources listed in this section were either referred to in the chapter or provide additional information concerning concepts and products discussed in this chapter.

Additional Resources on Remote Desktop Services

If you're already familiar with Terminal Services in Windows Server 2008, a good place to start learning about the new features and enhancements found in Windows Server 2008 R2 is "What's New in Remote Desktop Services" in the TechNet Library at <http://technet.microsoft.com/en-us/library/dd560658.aspx>.

Another good place to start is the Remote Desktop Services features page at <http://www.microsoft.com/windowsserver2008/en/us/rds-product-home.aspx>. Here you will find information about the benefits of this technology, product details, case studies, and links to technical resources, communities, and partners. This site is useful to both IT pros and business decision makers.

The Remote Desktop Services page on the Windows Server TechCenter is the starting point for IT pros who want to dig deeper into all aspects of Remote Desktop Services. This page can be found at <http://technet.microsoft.com/en-us/windowsserver/ee236407.aspx>.

You can use the Infrastructure Planning and Design (IPD) guide for Windows Server 2008 R2 Remote Desktop Services to help you plan Remote Desktop Services deployment for your environment. You can download this guide from <http://go.microsoft.com/fwlink/?LinkId=177881>.

The walkthroughs found in this chapter for various Remote Desktop Services role services are only the beginning of demonstrations of RDS functionality. For more detailed walkthroughs you can set up and test in your lab environment, see the Remote Desktop Services Step-by-Step Guides available in the TechNet Library at <http://technet.microsoft.com/en-us/library/dd736539.aspx>.

The installed Help content for each Remote Desktop Services role is an important learning resource. This content can be found in the TechNet Library at <http://technet.microsoft.com/en-us/library/cc770412.aspx>.

To learn more about the Windows PowerShell scripting capabilities of Remote Desktop Services, see the Remote Desktop Services page on the Script Center at <http://technet.microsoft.com/en-us/scriptcenter/ee364707.aspx>.

If you have any technical questions about Remote Desktop Services, try posting them to one of these TechNet forums:

- Hyper-V forum at <http://social.technet.microsoft.com/Forums/en-US/windowsserver2008r2virtualization/threads> (preferred forum)
- Terminal Services forum at <http://social.technet.microsoft.com/Forums/en-US/winserverTS/threads>

Finally, to keep abreast of the latest news about Remote Desktop Services and for walkthroughs, tips, and tricks, see the Remote Desktop Services (Terminal Services) Team Blog at <http://blogs.msdn.com/rds/>.

Additional Resources for App-V for RDS

For information about licensing App-V for RDS for your environment, see "Remote Desktop Services Licensing" at <http://www.microsoft.com/windowsserver2008/en/us/rds-product-licensing.aspx>.

The starting place for IT pros who want to learn more about App-V is Application Virtualization TechCenter, found at <http://technet.microsoft.com/en-us/appvirtualization/default.aspx>.

You can find a link where you can download the white paper "Application Virtualization 4.5 for Terminal Services" at <http://technet.microsoft.com/en-us/appvirtualization/cc843994.aspx>.

To learn the latest news about App-V 4.6 and App-V for RDS, see the App-V Team Blog at <http://blogs.technet.com/softgrid/>.

Additional Resources for Microsoft VDI

You can find a brief overview of the benefits and available offerings of Microsoft VDI on the Windows "Enterprise Solutions: Virtualization" page at <http://www.microsoft.com/windows/enterprise/solutions/virtualization/improve-flexibility.aspx>.

Information about the components and licensing of VDI suites can be found on the Remote Desktop Services Licensing page at <http://www.microsoft.com/windowsserver2008/en/us/rds-vdi.aspx>.

For information on why Windows 7 is the best desktop operating system for VDI, see the following post on the Remote Desktop Services Team Blog: <http://blogs.msdn.com/rds/archive/2009/11/02/windows-7-with-rdp7-best-os-for-vdi.aspx>.

Matt McSpirit, a Partner Technology Specialist at Microsoft, has a three-part video series that walks you through the steps of setting up a VDI infrastructure. The series is available at the following locations:

- <http://blogs.technet.com/mattemcspirit/archive/2009/09/24/virtualboytv-com-microsoft-vdi-part-i-server-side-configuration.aspx>
- <http://blogs.technet.com/mattemcspirit/archive/2009/09/24/virtualboytv-com-microsoft-vdi-part-ii-virtual-desktop-configuration.aspx>
- <http://blogs.technet.com/mattemcspirit/archive/2009/09/24/virtualboytv-com-microsoft-vdi-part-iii-client-side-experiences.aspx>

More information about Microsoft VDI is also available at <http://www.microsoft.com/vdi/>.

Chapter 5

Virtualization Management

Windows Server 2008 R2 Hyper-V is a foundational part of Microsoft's integrated virtualization solution for the enterprise. While Hyper-V alone might be sufficient for some organizations, others might have virtualized workloads running on the earlier Microsoft Virtual Server 2005 R2 SP1 platform or on VMware ESX 3.x Server computers within a VMware VI3 environment. These organizations can benefit from implementing Microsoft System Center Virtual Machine Manager 2008 R2, which has the capability of managing virtualized workloads running on all three host platforms—Hyper-V, Virtual Server, and VMware ESX Server—from a single, centralized platform. This chapter delves into the workings of Virtual Machine Manager 2008 R2, and explains how it works, its key features and benefits, key usage scenarios, and how to use it.



Note Unless otherwise indicated, the information in this chapter applies to both Virtual Machine Manager 2008 and Virtual Machine Manager 2008 R2. An earlier product called Virtual Machine Manager 2007, which was built upon Microsoft Virtual Server 2005 and could be used to manage only Virtual Server hosts, is discussed only briefly in this chapter when the key features of VMM 2008 and VMM 2008 R2 are summarized.

Understanding Virtual Machine Manager

VMM 2008 R2 consists of a number of components that work together at different layers to facilitate the provisioning and management of virtualized workloads across an enterprise. This section describes the terminology, different components, and underlying architecture of VMM 2008 R2 to explain how the product works.

Terminology

The following are some of the key concepts and terms you need to understand when working with VMM 2008 R2:

- **Guest operating system profile** A saved collection of settings that provide customization of the guest operating system. This profile is analogous to a setup answer file and contains information about the system settings, administrator account, and domain. Specific guest operating system profiles can be saved in the library and then used to quickly apply the settings to new virtual machines that are created from templates.

- **Hardware profile** A saved collection of settings that define the hardware characteristics of a virtual machine. These settings include items such as processors, memory, network, and DVD drives. Specific hardware profiles can be saved in the library and then used to quickly apply the settings to new virtual machines and templates.
- **Host** Also known as a virtual machine host, a physical computer that can host one or more virtual machines. Examples of hosts that can be managed using VMM 2008 R2 include servers running Microsoft Windows Server 2008 with the Hyper-V role installed, servers running Microsoft Hyper-V 2008 Server and Microsoft Hyper-V Server 2008 R2, servers running Microsoft Virtual Server 2005 R2, and servers running VMware ESX. Hosts are added by using the Add Hosts Wizard in the VMM Administrator Console, and until you add a host you cannot use VMM to create virtual machines.
- **Library server** The component of VMM 2008 that holds stored virtual machines, virtual hard disks, .iso files (CD/DVD software images), post-deployment customizations scripts, hardware configurations, and templates. The library provides a single interface for all of these virtualization building blocks.
- **Managed host** A host that has been added to a VMM library. VMM 2008 R2 allows for controlling multiple hosts by adding them to a central library and managing these hosts from one centralized location. After being added to the library, a host then becomes a managed host that can be managed by only one VMM Server at a time. Although multiple VMM Servers can exist on a network, a host can be managed by only one of these at a time. Should a different VMM Server want to add a host to its library and manage that host, it would have to take the host away from whatever other VMM Server had it.
- **Performance and Resource Optimization (PRO)** A feature first introduced in VMM 2008 that leverages the monitoring and alerting capabilities of Microsoft System Center Operations Manager (OpsMgr) 2007 to surface tips or recommendations within VMM 2008 R2 that can help administrators ensure high performance and an efficient virtualized environment.
- **Physical-to-Virtual (P2V)** A process in VMM 2008 R2 that converts a physical machine into a virtual machine.
- **Stored virtual machine** A managed virtual machine whose .vhd files and other properties are stored in the VMM library. A new .vmc file is created for a new virtual machine created from a stored virtual machine.
- **Template** A combination of a guest operating system profile, hardware profile, and one or more .vhd files. The .vhd file containing the operating system files has computer identity information removed using Sysprep. Templates are used to create new virtual machines.

- **Virtual machine host** Also known simply as a host, a virtual machine host is a physical computer that can host one or more virtual machines. Examples of hosts that can be managed using VMM 2008 R2 include servers running Windows Server 2008 or Windows Server 2008 R2 with the Hyper-V role installed, servers running Microsoft Hyper-V Server 2008 or Microsoft Hyper-V Server 2008 R2, servers running Microsoft Virtual Server 2005 R2 SP1, and servers running VMware ESX 3.x.
- **Virtual-to-Virtual (V2V)** A process in VMM 2008 R2 that converts a virtual machine running in a VMware environment (specifically, virtual machines running in an ESX environment using the .vmdk format) into a virtual machine running in a Windows Hyper-V environment.

VMM Components

VMM 2008 R2 consist of several core components, as described in the following sections. These components can be either installed together on a single server or distributed across multiple servers (or even workstations in some instances).



More Info For more information on installing VMM components, see the section titled “Installing VMM 2008 R2” later in this chapter.

Virtual Machine Manager Server

The VMM Server is the core component of a VMM 2008 R2 infrastructure—all other VMM components interact and communicate through the VMM Server. The Virtual Machine Management Service runs on the VMM Server and enables the running of commands and transferring of files throughout your VMM infrastructure. The VMM Server also controls all communications with other VMM components and with virtual machine hosts and VMM Library Servers via VMM Agents installed on these other computers. The VMM Server also connects to the Microsoft SQL Server 2005 or 2008 database used to store all VMM configuration information. By default, the VMM Server is also the default VMM Library Server. You configure and manage the VMM Server using the VMM Administrator Console.

Virtual Machine Manager Library Server

The VMM Library Server maintains the VMM library, a catalog of resources that can be used to create and configure virtual machines within a VMM infrastructure. The library contains files stored on library shares and can contain file-based resources, including .iso images, scripts, virtual hard disks, virtual floppy disks, virtual machine templates, guest operating system profiles, and hardware profiles. The library can also contain stored virtual machines that are not in use.

Virtual Machine Manager Agent

The VMM Agent manages virtual machines on virtual machine hosts and enables both and Library Servers to communicate with the VMM Server. Using the VMM Administrator Console to add a host or a Library Server belonging to a trusted domain automatically installs the VMM Agent on that managed computer. If the host is not joined to a domain, belongs to an untrusted domain, or resides on the perimeter network, the VMM Agent must be installed locally on the host before you can add the host using the VMM Administrator Console.

Virtual Machine Manager Administrator Console

The VMM Administrator Console is an MMC console you can use to manage global configuration settings; manage and monitor hosts and Library Servers; and create, deploy, and manage virtual machines. You can install the VMM Administrator Console on the same computer as the VMM Server or on a different computer. Installing the VMM Administrator Console also installs the Windows PowerShell console, which provides Virtual Machine Manager cmdlets you can use to perform all tasks that you can do using the VMM Administrator Console from the command line.

Virtual Machine Manager Self-Service Portal

The VMM Self-Service Portal is an optional, Web-based component you can use to allow end users to create and manage their own virtual machines. You do this by configuring self-service policies that control which templates self-service users can use to create their virtual machines, how many virtual machines they can create, which hosts their virtual machines can run on, and which actions the users can perform on their virtual machines.



Tip You must install the VMM Administrator Console on the same computer as your OpsMgr server if you plan on using the Reporting feature of VMM. This is because the Windows PowerShell console, used by Virtual Machine Manager, is needed by the System Center Operations Manager 2007 R2 administrator in order to perform tasks from within the Virtualization Management Pack.

VMM Architecture

VMM 2008 R2 was designed using a modular architecture to provide the greatest flexibility for managing the entire virtualization infrastructure of an enterprise. Using VMM 2008 R2, you can manage virtualized workloads on hosts running Microsoft's Hyper-V or Virtual Server 2005 R2 platforms, or on VMware's ESX Server platform.

As shown in Figure 5-1, the modular architecture of VMM 2008 R2 consists of three layers that communicate with one another using well-known documented interfaces:

- Client Layer
- Engine Layer
- Managed Computer Layer




FIGURE 5-1 Modular architecture of VMM 2008.

Interlayer Communications

Communications between the different layers of the VMM architecture involve the following transport mechanisms:

- **Windows Communication Foundation** WCF is a Microsoft messaging platform for building service-oriented applications that is part of the .NET Framework 3.0 and later. In VMM 2008 R2, WCF is used for all communication between the Client Layer applications and the Engine Layer.
- **Windows Remote Management** WinRM is the Microsoft implementation of the WS-Management Protocol, a standard Simple Object Access Protocol (SOAP)-based, firewall-friendly protocol that enables hardware and operating systems from different vendors to interoperate. WinRM uses a client/server architecture that includes a Hypertext Transfer Protocol (HTTP) listener on the WinRM server that awaits calls from WinRM clients. In VMM 2008 R2, WinRM is used for communication between the Engine Layer and Hyper-V hosts, Virtual Server hosts, and Library Servers.

- **Distributed Component Object Model** DCOM is the Microsoft Component Object Model (COM) specification that defines how components communicate over Windows-based networks. In VMM 2008 R2, the Engine Layer communicates with Physical-to-Virtual (P2V) source servers in the Managed Computer Layer using Windows Management Instrumentation (WMI) via DCOM.
- **Hypertext Transfer Protocol over Secure Sockets Layer** HTTPS is a combination of Hypertext Transfer Protocol (HTTP) over Transport Layer Security (TLS). In VMM 2008 R2, HTTPS is used by the Engine Layer for calling virtual infrastructure (VI) Web Services application programming interfaces (APIs) on VMware ESX hosts in the Managed Computer Layer.

In addition to the preceding transport mechanisms, VMM also uses ADO.NET for communications between the VMM Server and the Microsoft SQL Server database within the Engine Layer. ADO.NET is a suite of data access technologies included in the .NET Framework class libraries that provide access to relational data and XML.

Communications Ports

Communication between Client Layer applications, the Engine Layer, and managed hosts and Library Servers in the Managed Computer Layer of VMM take place using specific well-known TCP ports. This means that if the components of VMM are distributed across multiple computers, the host firewalls on these computers must be configured to enable communications over these ports. If Windows Firewall is enabled on these computers when VMM components are installed on them, the necessary exceptions are automatically opened in Windows Firewall.

Table 5-1 summarizes the transport mechanisms and default ports used for communications by VMM 2008 R2.

TABLE 5-1 Transport Mechanisms and Default Ports Used for Communications Between VMM Architectural Components

Communications Endpoints	Transport	Port
VMM Server to VMM Agent (control)	WinRM	80
VMM Server to VMM Agent (data)	BITS	443
VMM Server to remote database	HTTP	1433
VMM Server to P2V source	WinRM	135
VMM Administrator Console to VMM Server	WCF	8100
VMM Administrator Console to VMM Self-Service Portal	WCF	80
VMM Library to virtual machine hosts	BITS	80

Communications Endpoints	Transport	Port
Virtual machine host to virtual machine host	BITS	80
VMM Self-Service Portal user Internet Explorer session to Virtual Server host	VMRC	5900
VMM Self-Service Portal user Internet Explorer session to Hyper-V host	RDP	3389
VMM Administrator Console reporting view to System Center Operations Manager (OpsMgr) reporting server	HTTP	80
VMConnect on Hyper-V hosts for single-class console view	—	2179

Table 5-2 summarizes the default ports used by each component of VMM 2008 R2.

TABLE 5-2 Default Ports Used by Each Component of VMM 2008 R2

Component	Ports Needed
VMM Server	80 (HTTP, WS-MAN) 443 (HTTPS, BITS) 8100 (WCF connections to Windows PowerShell or Administrator Console)
VMM Library Server	80 (HTTP, WS-MAN) 443 (HTTPS, BITS) 3389 (RDP) 2179 (VMConnect on Hyper-V hosts for single-class console view) 5900 (VMRC on Virtual Server hosts)
Virtual machine hosts	80 (HTTP, WS-MAN) 443 (HTTPS, BITS) 3389 (RDP) 2179 (VMConnect on Hyper-V hosts for single-class console view) 5900 (VMRC on Virtual Server hosts)
SQL Server	1433 (Remote SQL instance connection) 1434 (SQL browser service—only needed for initial setup)
VMware Virtual Center server	443 (HTTPS for calling VI Web Services APIs)
VMware ESX hosts	443 (HTTPS for calling VI Web Services APIs) 22 (SSH for sFTP files to/from ESX hosts—not needed for ESX version 3.5i)

Client Layer

The Client Layer of the VMM architecture represents the applications and interfaces you use to interact with VMM 2008 R2. These applications include

- VMM Administrator Console
- VMM Self-Service Portal
- System Center Operations Manager
- Windows PowerShell Command-Line Interface

Using the applications in the Client Layer, you can perform actions such as creating and deploying a new virtual machine, starting or stopping virtual machines, monitoring virtual machines, and other tasks.

When you perform an action using one of the applications in the Client Layer, these actions are transformed into a Windows PowerShell script, which is run by the Engine Layer of VMM to perform the specified action.



More Info For more information concerning Windows PowerShell integration with VMM 2008 R2, see the sidebar titled “Direct from the Source: Windows PowerShell Integration in VMM 2008 R2” in this chapter.

Direct from the Source: Windows PowerShell Integration in VMM 2008 R2

Most user actions performed using the Client Layer applications result in the creation of a Windows PowerShell script that specifies the tasks that will be performed by the Engine Layer of VMM 2008 R2. The computer on which you install the VMM components must have either version 1.0 or 2.0 of Windows PowerShell installed. Windows PowerShell 1.0 is included in Windows Server 2008 as an optional feature that can be installed using the Add Features Wizard. Windows PowerShell 2.0 is included in Windows Server 2008 R2 and is installed by default. Windows PowerShell is a Windows command-line shell that provides an interactive command prompt and scripting environment built on the .NET Framework. Windows PowerShell introduces the concept of a cmdlet (pronounced “command-let”), a simple, single-function command-line tool built into the shell. Windows PowerShell 2.0 includes more than 200 core cmdlets that can be used alone or in conjunction with one another to perform complex tasks.

Windows PowerShell is also extensible in that additional cmdlets can be created and loaded as a snap-in dynamic-link library (DLL). Windows PowerShell is

extended to include VMM 2008 R2 cmdlets using the Microsoft.SystemCenter.VirtualMachineManager.dll snap-in DLL, located in the following folder:

`%ProgramFiles%\System Center Virtual Machine Manager 2008\Bin`

Launching the Windows PowerShell Virtual Machine Manager Command-Line Interface

The VMM 2008 R2 Windows PowerShell snap-in DLL is loaded when you start Windows PowerShell using the Windows PowerShell Virtual Machine Manager shortcut on the Start menu or when clicking the Windows PowerShell button in the VMM Administrator Console. The Start menu shortcut uses the following command to start Windows PowerShell and load the snap-in DLL:

**`C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe
-PSConsoleFile "C:\Program Files\Microsoft System Center Virtual Machine
Manager 2008\bin\cli.psc1" -NoExit`**

The command starts Windows PowerShell and specifies the console file cli.psc1. The cli.psc1 Windows PowerShell console file is an XML file containing the following elements, including the reference to Microsoft.SystemCenter.VirtualMachineManager.dll:

```
<?xml version="1.0" encoding="utf-8"?>  
<PSConsoleFile ConsoleSchemaVersion="1.0">  
  <PSVersion>1.0</PSVersion>  
  <PSSnapIns>  
    <PSSnapIn Name="Microsoft.SystemCenter.VirtualMachineManager" />  
  </PSSnapIns>  
</PSConsoleFile>
```

Note that the Microsoft.SystemCenter.VirtualMachineManager.dll snap-in is *not* loaded in Windows PowerShell if Windows PowerShell is started by clicking the Windows PowerShell icon pinned to the taskbar in Windows Server 2008 R2.

Configuring Windows PowerShell for Use with VMM 2008

The following four items are needed to configure Windows PowerShell for use with VMM 2008 R2:

- An application-specific profile
- An application-specific snap-in
- A security context with sufficient permissions to run custom scripts
- An application context

The application-specific profile and application-specific snap-in are configured when the cli.psc1 Windows PowerShell console file is loaded. The security context must be set by the user.

Windows PowerShell 2.0 operates using one of the following security contexts:

- **Restricted** This is the default PowerShell execution policy. When this policy is in effect, script execution is disabled; however, Windows PowerShell can still be used as an interactive command interpreter.
- **AllSigned** Specifies that only Authenticode-signed scripts can be executed. When running a signed script, users are asked if they want to trust the signer of the script.
- **RemoteSigned** Specifies that all scripts that are downloaded from a remote location must be Authenticode-signed before they can be executed.
- **Unrestricted** Windows PowerShell will execute any script. It will still prompt the user when it encounters a script that has been downloaded from a remote location. This is the least-secure setting.
- **Bypass** Nothing is blocked, and there are no warnings or prompts. This execution policy is designed for configurations in which a Windows PowerShell script is built in to a larger application or for configurations in which Windows PowerShell is the foundation for a program that has its own security model.

Two Windows PowerShell cmdlets, *Get-ExecutionPolicy* and *Set-ExecutionPolicy*, can be used to view and change the current execution policy. To set the execution policy, open a Windows PowerShell console with administrative rights and type the following command:

Set-ExecutionPolicy -ExecutionPolicy <executionpolicy>

where <executionpolicy> is either *Restricted*, *AllSigned*, *RemoteSigned*, *Unrestricted*, or *Bypass*.

The current execution policy can also be viewed by typing the following command:

Get-ExecutionPolicy

The application context sets the specific VMM Server to which all subsequent commands apply. By default, no application (server) context, even the local machine that is running Windows PowerShell, is assumed. Before executing VMM Windows PowerShell commands, the application context must be set. *Application context* refers to the specific instance of the VMM Server against which the requested script is to be run. To get the application context, type the following command:

get-VMMServer -ComputerName <FQDN computer name>

Note that configuring these Windows PowerShell settings is not needed when tasks are executed using the VMM Administrator Console user interface. This is because these tasks are executed internally within the VMM engine without using Powershell.exe.

—CSS Global Technical Readiness (GTR) team

VMM Administrator Console The VMM Administrator Console provides a central location for administering virtual machines and virtual machine hosts throughout an enterprise. The Administrator Console does this by creating and executing jobs, which are collections of tasks that perform virtual machine-related tasks, such as

- Creating and deploying virtual machines
- Converting physical machines to virtual machines (P2V)
- Converting from one type of virtual machine to another type, such as converting a VMware virtual machine to a Hyper-V virtual machine (V2V)
- Managing virtual machines running on local or remote virtual machine hosts
- Refreshing the VMM Administrator Console display
- Creating and managing virtual machine templates
- Managing the library resources needed to build virtual machines
- Delegating management tasks through Role-Based Authorization profiles and the Self-Service Portal user roles

The main executable for the Administrator Console is Vmmadmin.exe, which is found in the following directory:

%ProgramFiles%\System Center Virtual Machine Manager 2008\Bin

Vmmadmin.exe loads the Virtual Machine Remote Control ActiveX Control (Vmrcactivexclient.dll), which provides functionality for viewing (but not interacting with) virtual machines within the Administrator Console.

The Administrator Console also includes the Virtual Machine Viewer process (Virtualmachineviewer.exe), which provides remote access to virtual machines running on Hyper-V-based managed hosts. Virtualmachineviewer.exe is located in the following directory:

%ProgramFiles%\Microsoft System Center Virtual Machine Manager 2008\Bin

Virtualmachineviewer.exe is automatically started by Vmmadmin.exe when a connection is made from the Administrator Console to a Hyper-V-based virtual machine. Virtualmachineviewer.exe uses the Terminal Services ActiveX control, Mstscax.dll, to make a Remote Desktop Protocol (RDP) connection to the remote connection port specified when the host was added to the Administrator Console as a managed host. By default, TCP port 2179 is used for the RDP communication to Hyper-V-based hosts because this is the default port that the Hyper-V Virtual Machine Management Service (VMMS) listens on for incoming management connections.

Virtualmachineviewer.exe can also be launched separately from outside the Administrator Console interface by using a command prompt as follows:

virtualmachineviewer.exe /host *host* /computername *compname* [/port *portnumber*] [/vmid *vmid*] [/vmname *name*] [/vgsinstalled *true* | *false*]

Table 5-3 lists the various Virtualmachineviewer.exe command-line options.

TABLE 5-3 Command-Line Options for Virtualmachineviewer.exe

Option	Description
/host	The name of the host to connect to. When a host is provided, the connection feed is routed through the host so that the virtual machine is visible even when rebooting or disconnected from the network. This option cannot be used with the <i>computername</i> parameter.
/computername	The machine name of the virtual machine to connect directly to through remote desktop. This option cannot be used with the <i>host</i> parameter.
/port	Sets the port on which to communicate with the host. This option is ignored unless used with the <i>host</i> parameter. By default, this is set to 2179.
/vmid	The ID of the virtual machine to connect to. This must be provided to connect to the virtual machine when the <i>host</i> parameter is specified.
/vmname	A display name for the virtual machine used in the title bar of the viewer. This is required when the <i>host</i> parameter is specified.
/vgsinstalled	A Boolean value of true or false that indicates whether the Virtual Guest Services are installed on the virtual machine. By default, this is set to false.

Vmmadmin.exe also starts the Virtual Machine Remote Control process (Vmrc.exe) when a connection is made from the Administrator Console to a Virtual Server 2005 SP1–based virtual machine. Vmrc.exe provides remote access to virtual machines running on Virtual Server 2005 SP1–based managed hosts and is found in the following folder:

%ProgramFiles%\Microsoft System Center Virtual Machine Manager 2008\Bin

Vmrc.exe uses a modified Virtual Network Computing (VNC) protocol to communicate with Virtual Server 2005 SP1–based virtual machines using the remote connection port that was specified when the host was added to SCVMM as a managed host. Vmrc.exe can also be launched directly from outside the Administrator Console by using a command prompt as follows:

vmrc.exe [-fullscreen] [-viewonly] [-reducecolors] [vmrc server URL]

Table 5-4 lists the various Vmrc.exe command-line options.

TABLE 5-4 Command-Line Options for VMRC.exe

Option	Description
-fullscreen	Opens the VMRC console in full-screen mode
-viewonly	Disables mouse and keyboard input
-reducecolors	Enables reduced colors mode to increase performance
vmrc server URL	URL to a virtual server

Vmmadmin.exe also starts the VMware viewer process (Vmwareviewer.exe) when a VMware-based virtual machine is selected in the Administrator Console. Vmwareviewer.exe provides the initial connection to virtual machines running on VMware ESX Server–based managed hosts and loads the VMware WebCenter Remote MKS Plug-in control to enable communication with the VMware VirtualCenter Server console.

VMM Self-Service Portal The VMM Self-Service Portal is a Web-based ASP.NET application that provides an interface to allow designated users to perform specific virtual machine management tasks on managed hosts. The options that are available to Self-Service Portal users depend on the rights assigned to them by an administrator when their specific user role is created.

The Self-Service Portal for VMM 2008 R2 can be installed on Internet Information Services (IIS) 7.0 on Windows Server 2008 or on IIS 7.5 on Windows Server 2008 R2. During installation, you specify which port will be used for the Self-Service Portal. By default, port 80 is used unless another Web site, such as the Default Web Site, is already using port 80.

The Self-Service Portal is accessed by using Internet Explorer and requires the installation of ActiveX controls to provide remote access to virtual machines on managed hosts. When a user uses the Self-Service Portal to view or connect to any of her virtual machines, she is prompted to install the appropriate ActiveX control in Internet Explorer on her local machine.

Direct from the Source: System Center Operations Manager and VMM 2008 R2

System Center Operations Manager (OpsMgr), formerly Microsoft Operations Manager (MOM), is a performance-monitoring and event-monitoring application that can be used to monitor the health of an enterprise's IT infrastructure, including devices, operating systems, and applications. An OpsMgr agent is installed on monitored computers to monitor objects based on a predefined health model. The agents monitor several sources for events or other alerts and forward the information to a central OpsMgr server that maintains a database that includes a history of alerts. The OpsMgr server applies filtering rules to alerts as they arrive and takes appropriate action if needed. A filtering rule can trigger a notification, such as an e-mail or a pager message; generate

a network support ticket; or trigger some other workflow intended to correct the cause of the alert.

OpsMgr uses the term *management pack* to refer to a set of filtering rules that are specific to a monitored application. Management packs are available for VMM 2008 as well as several other Microsoft products, such as SQL Server and Exchange Server. OpsMgr also provides facilities for authoring custom management packs. Although an administrator role is needed to install agents, configure monitored computers, and create or install management packs, rights to monitor alerts can be given to any valid user account.

OpsMgr agents can be installed on the VMM Server, Windows-based virtualization hosts, and Windows-based virtual machines, including Windows-based virtual machines running on VMware ESX Server. OpsMgr communicates with the agents using the WCF over TCP port 5723.

OpsMgr integration with VMware is agentless and is accomplished using third-party management packs, such as nWorks, communicating directly with the VMware Web service using public VMware APIs. Users also get limited OpsMgr integration from VMM itself.

OpsMgr Integration with VMM

System Center Operations Manager 2007 R2 provides support for Performance and Resource Optimization (PRO) for VMM 2008 R2. PRO relies on OpsMgr to monitor and collect performance data from hosts and virtual machines within an environment. It is designed to make recommendations or take actions that take advantage of the capabilities provided by a virtualized environment. PRO collects performance and configuration data, which is used to generate tips to help users place, migrate, or reconfigure virtual machines to ensure workload uptime. PRO tips can be managed from the VMM Windows PowerShell console or from the VMM Administrator Console.

PRO leverages Intelligent Placement to help determine the best location for any VMs that are migrated as a result of tips that are generated. PRO provides workload-aware and application-aware resource optimization within host clusters that are managed jointly by VMM 2008 and OpsMgr. PRO is designed to help you evaluate whether physical resources provided by clustered host servers to the virtual machines deployed on those hosts are used efficiently.

PRO tips suggest how to remedy the issues raised in the alerts generated by OpsMgr. For example, a PRO tip might recommend increasing performance for a virtual machine by moving it to a new host with more resources or by adding an additional CPU to the virtual machine itself.

Capabilities of PRO

PRO in VMM 2008 R2 has the following capabilities:

- **Intelligent Placement** Provides intelligent placement of multimachine configurations across multiple hosts.
- **Clustering** PRO works in a clustered environment.
- **Health-based decisions** PRO uses health to determine when remediation is necessary.
- **In-Guest aware** PRO uses data collected from within the guest to suggest remediation.
- **Virtual machine right-sizing** PRO can make recommendations to alter the configuration of VMs to improve performance.
- **Host-level load balancing** PRO can provide a solution that ensures that the host load is balanced.
- **Automatic Remediation** PRO can be set up to automatically implement suggested remediation.

As noted above, you can configure VMM 2008 R2 to take corrective action automatically based on PRO tips or you can choose to respond to PRO tips manually. The following paragraphs describe what happens if you specify that VMM 2008 R2 will implement PRO tips automatically.

For Hyper-V hosts configured in a host cluster, VMM 2008 R2 can monitor and report at both the guest and host level and can use the Hyper-V Quick Migration feature to move virtual machines transparently between nodes in the cluster.

For VMware hosts configured in a host cluster, VMM 2008 R2 can also monitor and report at both the guest and host level and can use the VMware Live Migration feature (VMotion) to move virtual machines transparently between nodes in the cluster.

PRO Windows PowerShell Cmdlets

The following Windows PowerShell cmdlets are provided for PRO in VMM 2008 R2:

- Get-PROTip
- Dismiss-PROTip
- Set-PROTip
- Invoke-PROTip
- Set-VMHostCluster

The Set-VMMServer SCVMM cmdlet also supports parameters to enable PRO support in VMM 2008 R2.

—CSS Global Technical Readiness (GTR) team

Engine Layer

The VMM Engine Layer performs the tasks communicated to it by the Client Layer applications. The Engine Layer also controls the resources used by the Managed Computer Layer components. The Engine Layer is comprised of the following:

- Virtual Machine Management Service
- Microsoft SQL Server
- Other components

Virtual Machine Management Service The Virtual Machine Management Service (VmmService) is a system service implemented in VmmService.exe that provides the WinRM, WMI, HTTP, and WCF interfaces used for communication with the Client Layer applications and the Managed Computer Layer components. The Virtual Machine Management Service is also responsible for executing the individual tasks that make up a job. The Virtual Machine Management Service is dependent on the SQL Server (Microsoft\$VMM\$) system service.

SQL Server VMM 2008 R2 employs a Microsoft SQL Server database to maintain the database used to store library resource objects and certain configuration options, such as the VMware port number to use for communication with a VMware VirtualCenter Server. The version of SQL Server that you use with VMM depends on the needs of your environment. For example, if there is no previously installed version of SQL Server in your environment, installing VMM 2008 will install Microsoft SQL Server Express SP3. If you don't plan on using the PRO functionality of VMM 2008, SQL Server Express can provide the database functionality needed by VMM. Using the PRO feature of VMM 2008 R2, however, requires using OpsMgr 2007 R2, which requires a 32-bit or 64-bit version of either the Standard or Enterprise edition of Microsoft SQL Server 2008 or SQL Server 2005.



Note The structure of the database and the manner in which VMM 2008 R2 interacts with the database is the same regardless of which version of SQL Server you are using.

VMM 2008 R2 employs ADO.NET to communicate with SQL Server. ADO.NET is a .NET Framework-based data-access technology that enables applications to connect to data stores and manipulate the data contained in them. ADO.NET implements a disconnected database access model whereby a connection to the database is opened to serve an application request and is then closed after the request has been completed. This mechanism conserves system resources, has less impact on system performance, and enhances the security of databases. ADO.NET employs XML when interacting with a database and converts all of the data into XML format for database-related operations.



Note If a remote SQL Server database is used, TCP port 1433 must be open for the communication.

Other Engine Layer Components Some of the additional components within the Engine Layer of VMM 2008 include

- Authorization Manager
- A backup and restore engine
- Virtual Disk Service Providers

Managed Computer Layer

The Managed Computer Layer consists of the various types of virtualization hosts that VMM 2008 manages, including the Library Server and P2V Source servers. The Library Server maintains two kinds of resources:

- Resources not currently in use, such as library resources and P2V Source servers
- In-use resources, such as Hyper-V hosts, Virtual Server 2005 R2 SP1 hosts, VMware ESX hosts, and V2V Source hosts

VMM Agents With the exception of VMware-based hosts, communication between VMM 2008 R2 and the hosts it manages is accomplished using an agent service installed on the hosts. There are two types of VMM Agents:

- Host agent
- P2V agent

A host agent is automatically installed on a managed host when the host is added to VMM as a managed entity. The host agent is implemented as the VMMAgent service (Vmmagent.exe) and is used for managing Hyper-V and Virtual Server hosts. In addition to installing the VMMAgent service on remote managed hosts, the VMMAgent service is also installed by default on the VMM Server because it is needed for managing the library that is created on the server. The VMMAgent service depends on the Background Intelligent Transfer Service (BITS), Windows Management Instrumentation (WMI), and Windows Remote Management (WinRM) services.

A P2V agent is installed during a P2V conversion and is implemented using Vmmp2vagent.exe.

Library Servers and Resources The VMM library is a collection of resources that can be used to create and configure virtual machines. Library resources include the following:

- Virtual Hard Drive (.vhd) and Virtual Machine Configuration (.vmc) files
- Virtual machine templates and guest operating system profiles
- Scripts and Sysprep answer files
- ISO image files
- Virtual floppy disks

The VMM Library Server is a file server that has one or more shares. A VMM Agent runs on the Library Server to enable communication with a VMM Server and to identify the Library Server to the VMM Server as a Library Server. Only one VMM Agent can run on each Library Server, which means that a Library Server can connect only to one VMM Server. A VMM Server can connect to more than one Library Server, however.

Library Servers basically have two functions:

- To store the objects that can be used to create and configure virtual machines
- To transfer these objects to the hosts where they will be instantiated as running virtual machines



Tip Because the images that are stored in the library can be very large, a significant amount of network traffic can occur when these images are transferred to a virtual machine host during the virtual machine creation process. You should therefore locate your Library Servers on the same subnet as the hosts that they will be servicing.

The VMM library is more than just a collection of files in a share folder. Instead, library resources are organized into groups of objects that exist physically in the library share along with other objects that exist only within the SQL database. An example of a type of library object that does not exist physically in the library share but is still tracked as a library object is virtual machines, which run on managed hosts while being catalogued in the SQL database as part of the library infrastructure.

Key Features of VMM

This section lists some of the key features of the original release of System Center Virtual Machine Manager 2008, followed by a summary of the new features and enhancements found in Virtual Machine Manager 2008 R2.

Features and Improvements Introduced in VMM 2008

The original release of VMM 2008 improves upon the earlier VMM 2007 product in the following areas:

- The types of virtualization hosts that can be managed
- Support for highly available virtual machines through Failover Clustering
- Delegated administration based on role-based authorization
- Tuning, alerting, and reporting through integration with OpsMgr 2007

The sections that follow briefly describe these improvements made in the original release of VMM 2008 and which are all also found in VMM 2008 R2.

Windows Server 2008 Hyper-V Management

VMM 2008 introduced support for managing hosts running Microsoft Hyper-V. Specifically

- VMM 2008 is designed to fully use the foundational features and services of Windows Server 2008 Hyper-V, including Hyper-V's 64-bit architecture and attack-hardened security model.
- VMM 2008 supports installing the Hyper-V role remotely from the VMM 2008 console.
- VMM 2008 integrates with the Failover Clustering feature of Windows Server 2008 to allow you to create fault-tolerant and cluster-aware virtual machines.
- VMM 2008 supports Hyper-V functionality while providing additional VMM-specific functions such as Intelligent Placement, the VMM Self-Service Portal, and the integrated VMM library.

VMware (VI3) Management

VMM 2008 introduced support for managing VMware ESX Server hosts running within a VMware VI3 environment. Specifically

- VMM 2008 integrates multihypervisor management into one tool with its support for virtual machines running on VMware ESX infrastructure.
- VMM 2008 provides comprehensive support for VMware VI3, including moving virtual machines among virtual hosts with no downtime by using VMotion through integration with VMware's VirtualCenter.
- VMM 2008—specific features—such as Intelligent Placement, consolidation candidate recommendations, and others—can be run against virtualized infrastructure on any supported platform, including VMware VI3.
- Windows PowerShell scripts for customization or automation are supported across Hyper-V, VMware ESX, and Virtual Server implementations.

Windows Server 2008 Failover Clustering Integration

VMM 2008 included expanded support for failover clusters to provide enhanced high-availability capabilities for managing mission-critical virtual machines. Specifically

- VMM 2008 is fully cluster aware and can detect and manage Hyper-V host clusters as a single unit.

- VMM 2008 has automatic detection of virtual hosts that are added or removed from the cluster, which eases the burden on the administrator to manage this function.
- VMM 2008 creates a highly available virtual machine in a one-step process: the administrator selects a check box that designates a virtual machine as highly available. Behind the scenes, VMM 2008 orchestrates the creation of that highly available virtual machine by instructing the Intelligent Placement feature of VMM 2008 to recommend only hosts that are part of a host cluster.
- VMM 2008 includes improved highly available virtual machine management features, such as the Failover Cluster Management console for performing cluster-related tasks such as designation and management of cluster reserves, letter-less disk drives, guest clusters, and so on.
- VMM 2008 also supports VMware host clusters in which the nodes of the cluster are VMware ESX Servers.

Delegated Administration Based on Role-Based Authorization

VMM 2008 includes support for delegated administration based on role-based authorization. Specifically

- VMM 2008 allows administrators to create new user roles of the following types:
 - Administrator role
 - Delegated Administrator role
 - Self-Service User role
- VMM 2008 lets you define a scope for the Delegated Administrator or Self-Service User role to define which objects (such as host groups and Library Servers) the user can take actions on.
- VMM 2008 also lets you define permissions for a Self-Service User role to define what actions a user can perform on his virtual machines.

Performance and Resource Optimization

VMM 2008 introduced a feature called Performance and Resource Optimization (PRO), which can dynamically respond to failure scenarios or poorly configured components that are identified in hardware, operating systems, or applications. Specifically, PRO provides these capabilities:

- Working via PRO-enabled management packs and using System Center Operations Manager 2007's monitoring capabilities, PRO can either alert an administrator of an unhealthy system or application state and recommend corrective action, or it can

respond automatically to its recommendations to provide a self-healing virtualization infrastructure.

- Because of the highly granular level of monitoring available to PRO, a wide range of hardware, operating system, or application variables can trigger PRO to take corrective action.
- PRO's capabilities are also available to VMware ESX or Virtual Server hosts, allowing administrators to manage their entire virtualized environment regardless of the virtualization platform they are using.

New Features and Enhancements in VMM 2008 R2

VMM 2008 R2 introduces a number of new features and enhancements that make VMM even more powerful, more flexible, and easier to use than before. First, VMM 2008 R2 supports all of the improvements to Hyper-V found in Windows Server 2008 R2, including Live Migration, hot addition/removal of virtual hard disk (VHD) or iSCSI passthrough disk storage, and support for optimized networking technologies such as Virtual Machine Queue (VMQ) and TCP Chimney. VMM 2008 R2 also supports the new Cluster Shared Volumes (CSV) feature of Failover Clustering in Windows Server 2008 R2 that makes Live Migration possible. These improvements in Failover Clustering and Hyper-V in Windows Server 2008 R2 were described earlier in this book in Chapter 2, "Server Virtualization," so they won't be discussed further here. Instead, let's examine what improvements are specifically found in VMM 2008 R2.

Improved SAN Transfers

Enterprise environments that use storage area networks (SANs) will benefit from significantly improved SAN transfers in VMM 2008 R2 in several ways. Specifically

- **SAN migration into and out of clustered hosts** In VMM 2008 R2, you can now use SAN transfers to migrate highly available virtual machines into and out of a cluster. When you migrate a virtual machine into a cluster using a SAN transfer, VMM checks all nodes in the cluster to ensure each node can see the logical unit number (LUN). VMM also automatically creates a cluster disk resource for the LUN. Note that even though VMM automatically configures the cluster disk resource, it does not validate the resource. You must therefore still use the Validate A Configuration Wizard in Failover Cluster Management to validate the newly created cluster disk resource. Finally, to migrate a virtual machine out of a cluster, the virtual machine must reside on a dedicated LUN that is not using cluster shared volumes (CSV).
- **Expanded support for iSCSI SANs** VMM 2008 R2 now supports SAN transfers of virtual machines that use initiator-based iSCSI target connections. For this to work, there must be one iSCSI target for each LUN. VMM 2008 R2 also adds support for LUN

masking, which means you can now have multiple LUNs for each iSCSI target. Finally, VMM 2008 R2 includes expanded support for iSCSI SAN vendors.

Improved Support for Shared Storage

In addition to the support for CSVs that is provided by Failover Clustering in Windows Server 2008 R2, VMM 2008 R2 also provides the following additional shared storage improvements:

- **Sanbolic Clustered File System** VMM 2008 R2 now supports the Sanbolic Clustered File System (CFS), a third-party shared volume solution you can use for quick migration on hosts running Windows Server 2008 with Hyper-V and for live migration on hosts running Windows Server 2008 R2 with Hyper-V.
- **Veritas Storage Foundation for Windows** VMM 2008 R2 now supports Veritas Storage Foundation 5.1 for Windows (SFW), which is an online storage management solution you can use for creating virtual storage devices from physical disks and arrays. Any volumes created as part of a cluster resource group using SFW are automatically detected by VMM 2008 R2. These volumes can then be selected during virtual machine placement or migration. Note that an SFW volume is limited to a single virtual machine.

Improved Windows PowerShell Support

In addition to continuing to support Windows PowerShell 1.0, VMM 2008 R2 now also supports Windows PowerShell 2.0. This means you can now take advantage of the new features found in Windows PowerShell 2.0 such as remote management, background jobs, the debugger, and Integrated Scripting Environment (ISE). For more information on these new features of Windows PowerShell 2.0, see "What's New in Windows PowerShell" in the TechNet Library at <http://technet.microsoft.com/en-us/library/dd378784.aspx>.

VMM 2008 R2 also includes two new Windows PowerShell cmdlets:

- **Disable-VMHost** Places a VMM host into maintenance mode
- **Enable-VMHost** Removes a VMM host from maintenance mode, and returns it to service

VMM 2008 R2 also includes the following new parameters for existing VMM 2008 cmdlets:

- ***AllowUnencryptedTransfers*** Lets you specify that network files transfer into or out of a library or into, out of, or within a host group. It does not require encryption to improve performance when neither the source computer nor the destination computer requires encryption. You can use this parameter with the Set-LibraryServer and Set-VMHostGroup cmdlets.

- **BlockLMIfHostBusy** When used with the Move-VM cmdlet, this parameter enables you to stop attempting to restart a live migration that could not previously start because the source host or destination host was already performing a live migration.
- **LimitCPUForMigration** Limits the processor features for a virtual machine so that the virtual machine can be migrated to a different physical computer having a different version of the same processor as the source computer. Note that migrating virtual machines between physical computers with processors from different manufacturers is not supported. You can use this parameter with the New-HardwareProfile, Set-HardwareProfile, New-VM, Set-VM, and Set-Template cmdlets.
- **MoveWithinCluster** When used with the Disable-VMHost cmdlet to place a host that is a member of a host cluster into maintenance mode, all virtual machines currently deployed on the host will be migrated to another host in the same host cluster.
- **RemoveLibraryStoreSharePath** When used with the Set-VMMUserRole cmdlet, this parameter lets you clear the path to the specified library share.
- **RetainDeletedObjects** When used with the Get-VMMServer cmdlet, this parameter preserves objects in the cache that are marked for deletion.
- **RetainObjectCache** When this parameter is used with the Get-VMMServer cmdlet, the objects in the cache remain in memory and are not reclaimed by garbage collection.
- **UseCluster** When you move a virtual machine using the Move-VM cmdlet, this parameter lets you force the transfer of a virtual machine using Windows Server 2008 Cluster Migration even when Hyper-V live migration is available.
- **UseLocalVirtualHardDisks** When used with the New-VM cmdlet, this parameter lets you specify that the VHD file for the new virtual machine being created is at a specified location on the destination host, and that no VHD files will be copied from the library.
- **VMNetworkOptimizationEnabled** When used with the New-VirtualNetworkAdapter, Set-VirtualNetworkAdapter, New-P2V, and New-V2V cmdlets, this parameter enables virtual machine network optimization detection to allow you to improve network performance for virtual machines using network adapters that support Virtual Machine Queue (VMQ) or TCP Chimney Offload.
- **VMWarePortGroup** Lets you specify a VMware port group.

These new cmdlets and parameters in VMM 2008 R2 enable new scenarios as described in the sections that follow.

Maintenance Mode

If you need to perform maintenance on a VMM host—for example, by applying software updates or replacing a hardware component—you can use the new maintenance mode

feature of VMM 2008 R2 to easily prepare your host for maintenance. When you start maintenance mode on a host, either from the Administrator Console or by using the new `Disable-VMhost` cmdlet just described, VMM 2008 R2 does the following:

- Places all running virtual machines on the host into a saved state
- Prevents the creation of any new virtual machines on the host
- Displays the host status as In Maintenance Mode in the Host view of the VMM Administrator Console
- Excludes the host from the host ratings during any placement activities

In addition, if you start maintenance mode on a host that belongs to a Windows Server 2008 R2 cluster with highly available virtual machines, you also have the option of using live migration to move all the virtual machines to another host in the cluster so that connected users will not experience loss of service.

After you've finished performing maintenance on your host, you should stop maintenance mode on the host either from the Administrator Console or by using the new `Enable-VMhost` cmdlet. When you do this, the host status changes from In Maintenance Mode to OK in the Host view of the VMM Administrator Console. Note that when you stop maintenance mode, any virtual machines saved as a result of initiating maintenance mode are not automatically restarted—you must restart these virtual machines manually or by using scripts. Furthermore, if you used live migration to move highly available virtual machines to another host in a cluster when you initiated maintenance mode, stopping maintenance mode does not automatically migrate the virtual machines back to their original host—you must do this manually or by using scripts.

Support for Rapid Provisioning of Virtual Machines

In VMM 2008 R2, you can now quickly create virtual machines from a template using the new `UseLocalVirtualHardDisks` parameter of the `New-VM` cmdlet. The `UseLocalVirtualHardDisks` parameter specifies that New-VM should use an existing virtual hard disk file stored locally on the destination host instead of copying a VHD file from the library by using BITS. This allows you to provision virtual machines more rapidly when they are needed. For more information about how to use this parameter, type **Get-Help New-VM -detailed** at the VMM Windows PowerShell prompt.

In the TechNet Library at <http://technet.microsoft.com/en-us/library/cc967317.aspx>, you can find the sample scripts *RapidProvisionVM.ps1* and *RapidProvisionVMwithAnswerFile.ps1*, which demonstrate how to use rapid provisioning to create a new virtual machine. There is also the sample script *RecoverVMUsingRapidProvisioning.ps1*, which demonstrates how to use rapid provisioning to recover a virtual machine.

Support for Quick Storage Migration

VMM 2008 R2 also supports using the `Move-VirtualHardDisk` cmdlet with the `Move-VM` cmdlet to migrate a running virtual machine and its files (either Windows-based .vhd files or VMware-based .vmdk files) to a different storage location on the same host or on a different host with minimal or no downtime. This is called Quick Storage Migration, and it is supported for both Windows Server 2008 R2 hosts and VMware Storage VMotion-capable hosts. For more information on how to do this, type **Get-Help Move-VM –examples** at the Windows PowerShell prompt.

Support for VMware Port Groups for Virtual Switches

VMM 2008 R2 uses the network location and tag specified in the hardware configuration for the virtual network adapter to determine the network availability of a virtual machine on a host. In addition, VMM 2008 R2 now lets you select from the VMware port groups that are available for virtual switches when you are deploying the virtual machine to a VMware ESX Server host.

Enhanced Administrator Console

Although the basic layout and use of the VMM Administrator Console remains the same in this release, the design of the Administrator Console has been updated in VMM 2008 R2 to make it friendlier and easier to use.

Updated User Role Processing

In the original release of VMM 2008, access to virtual machines, hosts, and resources was determined using only the rights and permissions associated with VMM user roles. VMM 2008 did not actually make any changes to Hyper-V role definitions and role memberships. Instead, it just ignored the Hyper-V authorization store for the hosts and virtual machines under its management. In VMM 2008 R2, however, any changes you make to role definitions or role memberships are preserved in the root scope of the Hyper-V authorization store. All changes to any other scopes are overwritten every half hour by the VMM User Role Refresher.

Shared ISO Images for Self-Service Users

In the original release of VMM 2008, virtual machines could use a shared ISO image file. Shared ISO images enable virtual machines to share ISO image files stored in the VMM library in a Hyper-V environment managed by VMM. Unfortunately, in VMM 2008 this did not work for self-service users. Instead, VMM 2008 attaches a copy of the ISO image file to the new virtual machine. VMM 2008 R2, however, now supports shared ISO images for self-service users.

Key Benefits of VMM

The key benefits of using System Center Virtual Machine Manager 2008 R2 to manage your entire virtualized environment include the following:

- **Designed for virtual machines running on Windows Server 2008 R2**
Hyper-V Windows Server 2008 R2 Hyper-V is a hypervisor-based virtualization platform from Microsoft that is designed to offer high performance, enhanced security, high availability, scalability, and many other improvements. VMM 2008 R2 is designed to take full advantage of these foundational benefits through a powerful yet easy-to-use console that streamlines many of the tasks necessary to manage virtualized infrastructure. Administrators can also manage their traditional physical servers right alongside their virtual resources through one unified console.
- **Support for Microsoft Virtual Server and VMware ESX Server** VMM 2008 R2 can manage a VMware ESX virtualized infrastructure in conjunction with VMware VirtualCenter Server. This means that administrators running multiple virtualization platforms can rely upon a single tool to manage virtually everything within their infrastructure. With its compatibility with VMware VI3 through VirtualCenter Server, VMM 2008 R2 supports features such as VMware's VMotion and can even provide VMM-specific features such as Intelligent Placement to VMware ESX Servers.
- **Performance and Resource Optimization (PRO)** The PRO feature of VMM 2008 R2 enables dynamic management of virtual resources through PRO-enabled management packs for System Center Operations Manager 2007. PRO lets administrators establish remedial actions for VMM to execute if poor performance or pending hardware failures are identified. As an open and extensible platform, PRO also encourages partners to design custom management packs that promote compatibility of their products and solutions with PRO's powerful management capabilities.
- **Maximizing datacenter resources through consolidation** A typical physical server in the datacenter operates at only 5 to 15 percent CPU capacity. VMM 2008 R2 can assess and then consolidate suitable server workloads onto a virtual machine host infrastructure, thus freeing up physical resources for repurposing or retirement. Through physical server consolidation, continued datacenter growth is less constrained by space, electrical, and cooling requirements.
- **Easier machine conversions** Converting a physical machine to a virtual one used to be a daunting undertaking that was slow and problematic, and that typically required you to halt the physical server. Thanks to the enhanced Physical-to-Virtual (P2V) conversion in VMM, however, P2V conversions can become routine. In addition, VMM 2008 R2 also provides a straightforward wizard for converting VMware virtual machines to VHDs via a quick and easy Virtual-to-Virtual (V2V) transfer process.

- **Fast provisioning of new machines** In response to new server requests, an agile IT department can use VMM 2008 to deliver new servers to its business clients anywhere in the network infrastructure with a very quick turnaround. VMM 2008 enables this agility by providing IT administrators with the ability to deploy virtual machines in a fraction of the time it would take to deploy a physical server. Through one console, VMM 2008 allows administrators to manage and monitor virtual machines and hosts across an organization to ensure they are meeting the needs of business groups within that organization.
- **Minimizing virtual machine guesswork in deployment via Intelligent Placement** VMM 2008 does extensive data analysis of a number of factors before recommending which physical server should host a given virtual workload. This is especially critical when administrators are determining how to place several virtual workloads on the same host machine. And with access to historical data provided by System Center Operations Manager 2007, the Intelligent Placement process is able to factor in past performance characteristics to ensure the best possible match between the virtual machine and its host hardware.
- **Delegated virtual machine management** Virtual infrastructures are commonly used in test and development environments, where there is constant provisioning and tear down of virtual machines for testing purposes. This new version of VMM has a thoroughly reworked and improved Web-based Self-Service Portal through which administrators can delegate this provisioning role to authorized users while maintaining precise control over the management of virtual machines.
- **Organizing virtual machine components** To keep a datacenter's virtual house in order, VMM 2008 provides a centralized library to store various virtual machine building blocks, such as offline machines, templates, virtual hard disks, and other virtualization components. With the library's easy-to-use, structured format, IT administrators can quickly find and reuse specific components and thus remain highly productive and responsive to new server requests and modifications. Additionally, multiple Library Servers can be deployed throughout the organization if needed for increased scalability.
- **A rich management and scripting environment that uses Windows PowerShell** The entire VMM 2008 R2 application is built on the command line and scripting environment provided by Windows PowerShell 2.0. This version of VMM adds more Windows PowerShell cmdlets and View Script controls that allow administrators to explore customizing or automating operations at an unprecedented level.

Usage Scenarios for VMM

The following are some common usage scenarios for System Center Virtual Machine Manager 2008 R2:

- Server consolidation
- Provisioning of virtualized resources
- Business continuity
- Performance and resource optimization

Server Consolidation

One of the key scenarios that can take maximum advantage of the features found in VMM 2008 R2 is server consolidation. By consolidating physical servers, organizations can realize significant business benefits, including power savings and increased asset utilization. Organizations can take one of two approaches to server consolidation using VMM 2008 R2:

- A conservative, incremental approach whereby new applications are added to the virtual infrastructure while old applications remain on dedicated physical assets until retired
- A more aggressive approach that takes the form of a server consolidation project where the IT group identifies candidate applications for virtualization and then migrates the identified workloads to appropriate physical resources

VMM 2008 facilitates an aggressive approach to server consolidation by providing the following capabilities:

- Integration with System Center Operations Manager (OpsMgr) 2007 to enable VMM 2008 to identify potential consolidation candidates within the OpsMgr database
- A straightforward P2V conversion process that uses BITS to move data to the virtual machine and that can be automated using Windows PowerShell scripts
- A straightforward V2V conversion process that allows migration of a virtual machine running on VMware ESX Server to either Virtual Server 2005 R2 or Hyper-V
- An Intelligent Placement feature that helps identify the best host for running a virtualized workload

Provisioning of Virtualized Resources

Another key scenario for VMM 2008 R2 is facilitating the provisioning of virtualized resources to users and business groups who need them. By using virtualization, IT administrators

no longer have to procure and configure physical servers for new applications, a task that sometimes takes weeks or months. Instead, IT administrators can provision new virtual machines in a matter of minutes using VMM 2008 R2. And administrators can also delegate this provisioning role to authorized users by creating delegated administrators for groups and departments while maintaining precise control over the management of virtual machines. Authorized users can also access a Web page that enables them to provision their own virtual machines within predefined restraints laid down by administrators.

VMM 2008 enables such provisioning of virtualized resources by providing the following capabilities:

- Support for creating Administrator, Delegated Administrator, and Self-Service User roles and defining the scope of management action and virtual machine permissions for members of these roles
- A Web-based Self-Service Portal that an administrator can use to provide ordinary users with the capability to easily create and manage their own virtualized resources

Business Continuity

Business Continuity Planning refers to the creation and validation of a practiced logistical plan for how an organization will recover and restore partially or completely interrupted critical functions within a predetermined time after a disaster or extended disruption. VMM 2008 R2 can help ensure an organization's business continuity by enabling the management virtualized resources running on Windows Server 2008 R2 to be integrated with the following:

- An enterprise-level backup and recovery platform such as System Center Data Protection Manager 2007 R2
- Quick migration of virtualized workloads between nodes of a Windows Server 2008 failover clustering to ensure high availability of business-critical resources

Performance and Resource Optimization

Performance and Resource Optimization (PRO) is a feature of VMM 2008 R2 that helps administrators ensure that virtual machine hosts and their virtual machine guests are operating in the most efficient possible manner. PRO leverages OpsMgr 2007 R2 to monitor a complete end-to-end IT infrastructure, including hardware, host and guest operating systems, and applications. PRO also enables an administrator to create operational policies and automatically take actions based on those policies. For instance, when an event occurs that triggers a policy, PRO can be configured to present the issues to the administrator along with recommended resolutions. PRO can also be configured to automatically implement the preconfigured corrective actions for lights-out operations.

Working with VMM 2008 R2

System Center Virtual Machine Manager 2008 R2 is a powerful and flexible platform for managing different virtualized resources, including Hyper-V, Virtual Server, and VMware ESX Server virtual machines. This section examines how to install and use VMM 2008 R2 and covers the following topics:

- Planning for Deploying VMM 2008 R2
- System and Infrastructure Requirements
- Installing VMM 2008 R2
- Using the VMM Administrator Console
- Working with Managed Hosts
- Working with the Library
- Working with Virtual Machines
- Performing P2V Conversions
- Performing V2V Conversions
- Configuring User Roles
- Using the Self-Service Portal

Planning for Deploying VMM 2008 R2

Before you jump in and install VMM 2008 R2 in your organization, you need to do some initial planning. One of the things you should consider is the scope of your planned virtualization infrastructure. Will you be deploying hundreds of hosts or only a handful? This question is important to consider because VMM 2008 R2 uses a component-based architecture that can be installed in one of two ways:

- **Installing all VMM components on a single computer** If you are going to be using VMM to manage only a dozen or two hosts, you might consider installing all the VMM components—the VMM Server, VMM Administrator Console, VMM Self-Service Portal, VMM Library Server, and even your VMM Database—on a single server running Windows Server 2008 R2 with the Hyper-V role installed.
- **Installing each VMM component on a separate computer** If you are planning on using VMM to manage dozens or even hundreds of hosts, you should install the individual VMM components on dedicated computers to distribute the workload of these component for better performance.

Other questions to consider when planning your VMM 2008 R2 deployment include the following:

- Will you be implementing high availability using Failover Clustering?
- Will you be using a storage area network (SAN) environment for storing your virtual machine files?
- Will you deploy (or have you already deployed) other System Center products such as System Center Operations Manager 2007 R2 in your organization?
- Do you have an existing VMM 2008 infrastructure that you want to upgrade to VMM 2008 R2, or is this a new virtualization infrastructure you want to deploy?
- Will your hosts be running only Windows Server 2008 R2 Hyper-V, or will you also have Virtual Server 2005 hosts, VMware hosts, or both?

System and Infrastructure Requirements

After you've finished your initial planning and before you actually deploy VMM 2008 R2 in your production environment, you need to ensure that all prerequisites have been met. The prerequisites for installing VMM 2008 R2 include

- Hardware requirements
- Software requirements
- Infrastructure requirements

Hardware Requirements

The hardware requirements for installing VMM 2008 R2 on a computer depend on whether you are installing all or some VMM components on a particular computer. For example, if you are going to be managing only a handful of hosts, you can install all VMM components on a single computer having the following recommended hardware:

- A dual-processor or dual-core processor system running at 2 GHz or higher
- At least 2 GB of RAM
- At least 40 GB of hard disk space (more if you will be running a local instance of SQL Server)

If, however, you plan on using VMM to manage around a hundred hosts, you will want to install the various VMM components on different computers. In this case, the recommended hardware for your VMM Server is

- A dual-processor or dual-core-processor system running at 2.8 GHz or higher

- At least 4 GB of RAM
- At least 40 GB of hard disk space (more if you will be running a local instance of SQL Server)

And the recommended hardware for your VMM Library Server is

- A dual-processor or dual-core-processor system running at 3.2 GHz or higher
- At least 2 GB of RAM
- As much hard disk space as you will need

Of course, using more cores, faster processors, and more RAM will usually boost performance. For detailed information on VMM hardware requirements, see “VMM System Requirements” at <http://technet.microsoft.com/en-us/library/cc764328.aspx>.

Software Requirements

The software requirements for installing VMM 2008 R2 on a computer depend on whether you are installing all or some VMM components on a particular computer. For example, if you install all VMM components on a single computer, the computer must be running the Standard, Enterprise, or Datacenter edition of either Windows Server 2008 SP2 with Hyper-V or Windows Server 2008 R2 with Hyper-V. Before you install all the VMM components on a single server, you must install the Web Server (IIS) role, including the following role services:

- IIS 6 Metabase Compatibility
- IIS 6 WMI Compatibility
- Static Content
- Default Document
- Directory Browsing
- HTTP Errors
- ASP.NET
- .NET Extensibility
- ISAPI Extensions
- ISAPI Filters
- Request Filtering

And if you are installing all VMM components on a single server, you will also want to install the IIS Management Console so that you can delete the Default Web Site from the server. (See Figure 5-2.) You will want to do this because otherwise you'll need to create a host header for the VMM Self-Service Portal so that this portal and the Default Web Site can co-exist using the same IP address and port number.




FIGURE 5-2 Delete the Default Web Site before installing all VMM 2008 R2 components on a single computer.

Additional software requirements for VMM 2008 R2 include

- Windows PowerShell 1.0 or 2.0
- Windows Remote Management (WinRM) 1.1 or 2.0
- Microsoft .NET Framework 3.0 or higher

Because Windows Server 2008 R2 has Windows PowerShell 2.0, WinRM 2.0, and the .NET Framework 3.0 SP1 already installed, it's the preferred platform for installing VMM 2008 R2.

Finally, you need a supported version of Microsoft SQL Server, which includes SQL Server 2008, SQL Server 2005, or SQL Server 2005 Express Edition. If you will be managing many hosts, you should use a full version of either SQL Server 2008 or 2005 and install your VMM Database on a separate server. If you are managing few hosts, you can install the VMM Database server on your VMM Server computer and choose to have SQL Server 2005 Express Edition SP3 installed on the computer as part of the VMM installation process.



Tip For the best performance, especially when your VMM Server will be managing a large number of hosts, you should use a separate server for your VMM database.

Some VMM 2008 components can even be installed on computers running operating systems other than Windows Server 2008 (x64). Table 5-5 lists the supported operating systems for installing the VMM 2008 R2 components and for Windows-based virtual machine hosts. The table also shows which operating systems you use when installing all VMM components on a single computer. Note that VMM 2008 R2 requires Windows Server 2008 with Hyper-V 64-bit with Service Pack 2 for Hyper-V hosts.

TABLE 5-5 Supported Operating Systems for Installing the VMM 2008 R2 Components and Windows-Based Virtual Machine Hosts

Operating System	Single Computer	VMM Server	Hyper-V Hosts	Virtual Server Hosts
Windows Server 2008 R2 with Hyper-V 64-bit, Standard, Enterprise, and Datacenter editions	Yes	Yes	Yes	No
Windows Server 2008 R2 without Hyper-V 64-bit, Standard, Enterprise, and Datacenter editions	No	Yes	No	No
Windows Server 2008 R2 with Hyper-V 64-bit Server Core Installation, Standard, Enterprise, and Datacenter editions	No	No	Yes	No
Windows Server 2008 with Hyper-V 64-bit, Standard, Enterprise, and Datacenter editions with Service Pack 2	Yes	Yes	Yes	No
Windows Server 2008 with Hyper-V 64-bit, Standard, Enterprise, and Datacenter editions with Service Pack 1 or earlier	Yes	Yes	Yes (required for VMM 2008 R2)	No
Windows Server 2008 without Hyper-V 64-bit, Standard, Enterprise, and Datacenter editions	No	Yes	Yes (VMM 2008 only)	No
Windows Server 2008 without Hyper-V 32-bit, Standard, Enterprise, and Datacenter editions	No	No	No	Yes
Windows Server 2008 — Server Core installation, Standard, Enterprise, and Datacenter editions with Service Pack 2	No	No	Yes (required for VMM 2008 R2)	No
Windows Server 2008 — Server Core installation, Standard, Enterprise, and Datacenter editions with Service Pack 1 or earlier	No	No	Yes (VMM 2008 only)	No

Operating System	Single Computer	VMM Server	Hyper-V Hosts	Virtual Server Hosts
Windows Server 2003 with Service Pack 2, Standard, Enterprise, and Datacenter editions	No	No	No	Yes
Windows Server 2003 R2 with Service Pack 2	No	No	No	Yes
Windows Server 2003 x64 edition with Service Pack 2	No	No	No	Yes
Windows Server 2003 R2 x64 edition with Service Pack 2	No	No	No	Yes

Table 5-6 lists the supported operating systems for the VMM Administrator Console, Self-Service Portal, and Library Server components.

TABLE 5-6 Supported Operating Systems for the VMM Administrator Console, Self-Service Portal, and Library Server Components

Operating System	VMM Administrator Console	VMM Self-Service Portal	VMM Library Server
Windows Server 2008 R2 with Hyper-V 64-bit, Standard, Enterprise, and Datacenter editions	Yes	Yes	Yes
Windows Server 2008 R2 without Hyper-V 64-bit, Standard, Enterprise, and Datacenter editions	Yes	Yes	Yes
Windows Server 2008 R2 — Server Core installation, Standard, Enterprise, and Datacenter editions	No	No	Yes
Windows Web Server 2008 R2	No	Yes	No
Windows Server 2008 with Hyper-V 64-bit, Standard, Enterprise, and Datacenter editions	Yes	Yes	Yes
Windows Server 2008 without Hyper-V 64-bit, Standard, Enterprise, and Datacenter editions	Yes	Yes	Yes
Windows Server 2008 without Hyper-V 32-bit, Standard, Enterprise, and Datacenter editions	Yes	Yes	Yes
Windows Server 2008 - Server Core installation, Standard, Enterprise, and Datacenter editions	No	No	Yes
Windows Web Server 2008	No	Yes	No

Operating System	VMM Administrator Console	VMM Self-Service Portal	VMM Library Server
Windows Server 2003 with Service Pack 2, Standard, Enterprise, and Datacenter editions	Yes	Yes	Yes
Windows Server 2003 R2 with Service Pack 2	Yes	Yes	Yes
Windows Server 2003 x64 with Service Pack 2	Yes	Yes	Yes
Windows Server 2003 R2 x64 with Service Pack 2	Yes	Yes	Yes
Windows 7	Yes	No	No
Windows Vista with Service Pack 1	Yes	No	No
Windows XP Professional with Service Pack 2 or Service Pack 3	Yes	No	No
Windows XP Professional x64 with Service Pack 2	Yes	No	No

For more information on VMM software requirements, see “VMM System Requirements” at <http://technet.microsoft.com/en-us/library/cc764328.aspx>.

Infrastructure Requirements

VMM 2008 must be deployed within an Active Directory Domain Services environment. Specifically, the server or servers on which you will be installing the VMM 2008 components must be joined to a domain. In addition, any hosts that your VMM Server will manage must also be a domain member, either of the same domain where your VMM Server resides or in a trusted domain.

Although a Fast Ethernet (100 Mbps) network infrastructure is usually sufficient connectivity for the servers in your VMM 2008 deployment, you should use Gigabit Ethernet (1000 Mbps) for best performance.

Installing VMM 2008 R2

After you’ve verified that your environment meets the system and infrastructure requirements for deploying VMM 2008 R2, you can begin the installation process.

Using the Virtual Machine Manager Configuration Analyzer

The first step of the installation process involves running the Virtual Machine Manager Configuration Analyzer (VMMCA), a diagnostic tool that verifies the configuration settings

for computers to verify that they can run VMM 2008. Using the VMMCA, you can scan your computers to verify whether they are suitable to function as a VMM Server, run the Administrator Console, function as a Self-Service Portal, or be a managed virtual machine host.

Begin by downloading the 64-bit version of Microsoft Baseline Configuration Analyzer (MBCA) from the Microsoft Download Center at <http://go.microsoft.com/fwlink/?LinkID=97952>. You must install the MBCA before you can use the VMMCA. Be sure to install the MBCASetup64.msi file, which is the 64-bit version of MBCA. After you've installed the MBCA, download and install the VMMCA from the Microsoft Download Center at <http://www.microsoft.com/downloads/details.aspx?FamilyID=02d83950-c03d-454e-803b-96d1c1d5be24&displaylang=en>.

After you've installed both tools, open the VMMCA from the Start menu and specify the names of the computers on which you want to install each VMM role. (See Figure 5-3.)




FIGURE 5-3 Use the Virtual Machine Manager Configuration Analyzer to evaluate important configuration settings for computers that might serve VMM roles or other VMM functions.

Now click Scan and wait for VMMCA to scan each computer you specified. After the scan is complete, Internet Explorer opens and displays the results of the scan. (See Figure 5-4.) If any issues are displayed, be sure to resolve them now before you proceed with deploying VMM 2008 R2.




FIGURE 5-4 Results of a VMMCA scan show the computer is ready for installing VMM 2008 R2.

Installing the VMM Server

After you've run the VMMCA and verified that your computer or computers will support VMM 2008 R2, you're ready to begin installing the various components of the product. As mentioned earlier, these different components can be installed either on a single server or on multiple servers on your network. For purposes of illustration, the following walkthrough installs all VMM components on a single server running Windows Server 2008 R2 with the Hyper-V role installed.

To begin the installation, insert the product media into the server's DVD drive. The splash screen is displayed, showing the different VMM components you can install and other options. (See Figure 5-5.)

The first component you should install should be the VMM Server. When you install this component, the domain account you are currently logged in with is automatically added to the VMM Administrator user role. You can add other user accounts to the VMM Administrator user role afterwards—see the “Managing User Roles” topic in VMM Help for information on how to do this.

When you install the VMM Server component, the VMM setup program performs a check to ensure that you have met the necessary hardware and software prerequisites for installing the component. (See Figure 5-6.) Note that this check is in addition to and complementary to the check you performed earlier using the VMMCA. Do not rely only on the check performed by the VMM setup program—be sure to use the VMMCA first as well.




FIGURE 5-5 Splash screen for installing different VMM components.




FIGURE 5-6 The VMM setup program performs its own check to ensure the computer meets the necessary hardware and software prerequisites for installing different VMM components.

The VMM setup program also prompts you to configure SQL Server settings. VMM uses a Microsoft SQL Server database to store the information you see in the VMM Administrator Console, such as virtual machines, virtual machine hosts, virtual machine library servers, and

so on. You can either install SQL Server 2005 Express Edition SP3 locally on your VMM Server or specify an existing remote instance of SQL Server 2005 or 2008. If you are installing all VMM components on a single computer, you will likely want to use SQL Server 2005 Express Edition. (See Figure 5-7.)




FIGURE 5-7 Configure SQL Server settings for locating the VMM Database component.

During installation of the VMM Server, you are also prompted to create a new library share on your server or specify a preconfigured share on another server. The VMM library serves as a central, secure store for the resources used to create virtual machines in your VMM 2008 environment and helps enable the re-use of approved images and configurations. The library share is created as part of the VMM Server Setup process. The default is to create a share on the VMM Server on the system drive at %SystemRoot%\ProgramData\Virtual Machine Manager Library Files using a share name of MSSCVMMLibrary. (See Figure 5-8.) Additional Library Servers and shares can then be added afterwards by using the VMM Administrator Console. Note that you cannot delete or relocate the default library server or its library share, so carefully consider the location of your default library share before installing the VMM Server.

A best practice is to make the Library Server by using a highly available file server running on a Windows Server 2008 or 2008 R2 Failover cluster. If a Failover cluster is not available, locate the share on a nonsystem, high-speed (minimum 10,000 rpm) hard disk drive for best performance. Make your selection carefully because you cannot remove or relocate the default Library Server or its library share after installing your VMM Server.




FIGURE 5-8 Specify a default share for the VMM library.

Next, accept the default port assignments and service account for the VMM Server component or change these to ports and an account you specify. (See Figure 5-9.) If you specify a domain user account for the VMM service account, the account should be a dedicated account not used for any other purpose. In particular, it should not be the same account used for VMM communications with SQL Server. The service account should also belong to the local Administrators group on the VMM Server, and if you plan on using PRO the account should also be a member of the Administrator role in OpsMgr 2007 R2.




FIGURE 5-9 Specify the ports and service account for the VMM Server component.

You must specify a domain account instead of Local System on this Setup page if any of the following are true:

- Your AD DS environment has the Restricted Groups group policy implemented, because this policy prevents machine accounts from being members of the local Administrators group.
- VMM will be used to manage hosts in a disjoint namespace environment, where the fully qualified domain name (FQDN) of a Windows Server–based host in AD DS does not match the FQDN of the host in Domain Name System (DNS).
- You want to use shared ISO images with your Hyper-V virtual machines.

As for port assignments, you might want to consider modifying the default port assignments for HTTP and HTTPS to harden your server, although this can introduce some added complexity in managing your virtualization infrastructure.

Now finish the installation to install the VMM Server component and other required features, such as the Windows Automated Installation Kit 1.1, SQL Server, and SQL Server Tools. (See Figure 5-10.)




FIGURE 5-10 Completing installation of the VMM Server component.

Installing the VMM Administrator Console

After you finish installing your VMM Server, return to the splash screen and install the VMM Administrator Console on the local server. Installing the Administrator Console also installs

the Windows PowerShell Virtual Machine Manager console on the computer. If you modified the default port assignments when installing your VMM Server, be sure to take this into account when installing the VMM Administrator Console. (See Figure 5-11.)




FIGURE 5-11 Be sure to specify the correct port so that your VMM Administrator Console can connect to your VMM Server.

You might also want to install the Administrator Console on additional computers later to remotely access and manage your VMM Server, but it's helpful to install this console locally on your VMM Server in case you need it for troubleshooting purposes. In addition, if you plan on using the reporting feature of VMM 2008 R2, you must install the Administrator Console on the same computer as the VMM Server component. This is because the reporting feature of VMM 2008 R2 relies on System Center Operations Manager (OpsMgr) 2007 R2, and it relies on OpsMgr administrators to perform tasks on hosts and virtual machines from within the Server Virtualization Management Pack, which requires that the Windows PowerShell Virtual Machine Manager console be installed on the VMM Server.

Installing the VMM Self-Service Portal

An optional deployment step is to install the VMM Self-Service Portal, a Web-based component that lets users create and manage their own virtual machines within a controlled environment. If you choose to install the Self-Service Portal, you can install it on a separate computer from your VMM Server for best performance, but if you're managing only a handful of hosts you can install the Self-Service Portal on your VMM Server computer. Note that installation of the Self-Service Portal on a domain controller is not supported.

When you install the Self-Service Portal, you are prompted to specify the FQDN of your VMM Server and the TCP port for connecting to the server. You are also prompted to specify a TCP port for self-service users to use the Self-Service Portal, which by default is port 80. (See Figure 5-12.) Because the Default Web Site in IIS also uses port 80, if your VMM Server has only one IP address assigned to the network adapter used by the Self-Service Portal, you need to either delete the Default Web Site (as indicated previously) or specify a host header name so that self-service users will be able to connect to and use the Self-Service Portal.




FIGURE 5-12 Make sure you have the correct port assignments for the VMM Self-Service Portal.

For more information on installing VMM 2008 R2 components, see "New Installation of VMM" at <http://technet.microsoft.com/en-us/library/cc793149.aspx>.

Using the VMM Administrator Console

The VMM Administrator Console can be used to manage all aspects of a virtualized environment, including virtual machines on managed hosts running Microsoft Hyper-V, Microsoft Virtual Server 2005 R2, and VMware ESX Server 3.x within a VMware VirtualCenter 2.5 or 2.0.1 environment. As mentioned previously, the Administrator Console can be installed both locally on your VMM Server and also on additional computers in your environment for remote management purposes. By default, any user account that has local administrator privileges on your VMM Server can use the Administrator Console.

When you first open the Administrator Console, the Connect To Server dialog box opens. (See Figure 5-13.) If you are opening the Administrator Console on your VMM Server, you

should specify `localhost:port` in the Server Name field of the dialog box, where *port* is the TCP port you assigned when you installed the VMM Server. (The default is TCP port 8100.) If you are opening the Administrator Console on a different computer, you can specify the server using its name, FQDN, or IP address.




FIGURE 5-13 Connecting to a VMM Server using the VMM Administrator Console.

The first time you open the Administrator Console after installing the VMM components in your environment, no hosts are displayed in the console—even if you installed all the VMM components on a Hyper-V server that has virtual machines on it. As you can see from Figure 5-14, the Administrator Console includes the following user interface items:

- **Command menu** Lets you connect to other VMM Servers, change the current view, hide or display the navigation pane, and perform actions relating to the view selected.
- **VMM taskbar** Lets you hide or display the Actions pane, display additional columns of information in the Results pane, open the Jobs window, use PRO tips, display the network configuration, and open the Windows PowerShell VMM console.
- **Navigation pane** Lets you browse hosts and host groups, browse library servers and profiles, and display jobs and other administrative information.
- **Filters pane** Lets you filter by status, owner, operating system, and other criteria depending on the current view selected.
- **View buttons** Lets you change the current view to perform related tasks for managing hosts, virtual machines, library servers, and so on.
- **Results pane** Displays more information about whatever is currently selected in the navigation pane.
- **Details pane** Displays more information about whatever is currently selected in the results pane.
- **Actions pane** Lets you perform management tasks relating to whatever is currently selected in the navigation and results panes.




FIGURE 5-14 Layout of the VMM Administrator Console.



Tip The current view persists between the closing and opening of the Administrator Console.

Understanding Views

Depending on how you have deployed VMM 2008 R2, the Administrator Console can show up to seven possible views: Hosts, Virtual Machines, Library, Jobs, Administration, Reporting, and Diagram. It is possible that not all views are available for your installation. For example, the Reporting and Diagram views are not available unless VMM 2008 R2 has been configured to work with System Center Operations Manager 2007 R2.

The sections that follow list the seven views available in the Administrator Console and describe the different administrative tasks you can perform using each view. Additional information about using some of these views can be found in the sections that follow.

- **Hosts view:**
 - Add, remove, and monitor the status of virtual machine hosts.

- ❑ Configure virtual networks, placement options, virtual machine paths, and custom properties on a host.
- ❑ Enable remote connections to virtual machines. Register existing virtual machines on a host.
- ❑ Create and delete host groups for ease of monitoring and managing hosts.
- ❑ Configure host reserves and self-service policies for host groups.
- Virtual Machines view:
 - ❑ Create, deploy, migrate, operate, connect to, clone, repair, store, and remove virtual machines.
 - ❑ Create checkpoints so that you can restore virtual machines to a previous state.
- Library view:
 - ❑ Add file-based resources to the Virtual Machine Manager library for use in creating virtual machines.
 - ❑ Add Library Servers and library shares.
 - ❑ Refresh a library share to immediately index its files in Virtual Machine Manager. (By default, all library shares are refreshed every hour.)
 - ❑ Configure guest operating system profiles, hardware profiles, and virtual machine templates for use in virtual machine creation.
- Jobs view:
 - ❑ Monitor, cancel, restart, search, sort, filter, and group jobs.
 - ❑ View the changes that a job made to objects.
- Administration view:
 - ❑ Overview—View graphical summary information of the environment (hosts, virtual machines, recent jobs, and library resources).
 - ❑ General—Configure global VMM settings, such as Remote Control and PRO.
 - ❑ Managed Computers—Manage Virtual Machine Manager agents on managed hosts and Library Servers, update the agent, remove agent roles, and re-associate agents with the current VMM Server.
 - ❑ Networking—View the MAC address range used by VMM across all managed hosts.
 - ❑ User Roles—View all existing user roles grouped by profile type.
 - ❑ System Center—View reports generated by System Center Operations Manager.
 - ❑ Virtualization Managers—View all available virtualization managers, such as Virtual Machine Manager and VMware VirtualCenter Server.

- Reporting view:
 - View and open reports. Reporting view is available only if you have configured Operations Manager.
- Diagram view:
 - Displays the health of the Virtual Machine Manager Server, database server, Library Servers, hosts, virtual machines, and VMware VirtualCenter Server servers. Diagram view is available only if you have configured Operations Manager.



Note In addition to the seven views just described, there is also a special Networking view that can be used to display a graphical representation of the current network configuration. The Networking view does not have a view button and is invoked differently than the other views. See the section titled “Using Networking View” later in this chapter for more information.

Using Filters

The information displayed in each view can be filtered by using the Filters pane of the Administrator Console. The type of filtering you can perform depends on which view is currently selected. For example, Figure 5-15 shows that the Virtual Machines view is selected and the Filters pane of the console is being used to display only virtual machines whose status is Stopped. The effect of applying this filter is that the results pane shows only virtual machines that are stopped.




FIGURE 5-15 The filters available in the Virtual Machines view.

Working with Managed Hosts

You can use the Administrator Console to manage virtual machine hosts across a virtualization infrastructure. For example, you can create host groups to organize your hosts, add hosts to a host group, and configure and manage different kinds of hosts, including Hyper-V, Virtual Server, and VMware ESX hosts. The following sections provide further details concerning some of the tasks you can perform for managing hosts.

Creating and Using Host Groups

A *host group* is a logical container for organizing managed hosts within the Administrator Console. Creating host groups and adding hosts to them is a simple way of making it easier for you to manage hosts across your virtualization infrastructure.

You create and work with host groups within the Hosts view of the Administrator Console. By default, all managed hosts belong to the All Hosts group. When you create a new host group, the new group is added beneath the All Hosts group. You can also nest host groups to create three, four, or more levels of host groups if desired.

After you've created your hierarchical structure of host groups with the All Hosts group at the top, you can then add managed hosts and host clusters to each group as desired.

Figure 5-16 shows an example of how to use host groups to organize managed hosts according to geographical location. In this figure, three host groups—New York, Hong Kong, and Redmond—were created within the All Hosts group. The New York hosts group is currently selected.

The Virtual Machine Manager section allows you to perform actions on the selected managed host in the top middle Results pane. In this figure, a managed host named HVR2.contoso.com is currently selected in the Results pane, and the status and other information concerning this virtual machine is displayed on the Summary tab in the bottom middle Details pane. The actions you can perform on this managed host include creating a new virtual machine, converting a physical server, converting a virtual machine, and other tasks. You can perform these actions either by clicking the appropriate link in the Actions pane or by right-clicking on a managed host in the Results pane and using the shortcut menu options.

The Hosts section of the Actions pane allows you to perform other actions on managed hosts, including moving the selected host to a different host group, removing a host from the console, displaying and configuring the properties of the host, and so on.




FIGURE 5-16 Working with host groups.

The Actions pane on the right side in Figure 5-16 has three sections: Virtual Machine Manager, Host, and Host Group.

The Host Group section of the Actions pane lets you create new host groups for organizing your managed hosts for easier administration. You can also use this section of the Actions pane to view and configure the properties of a host group. For example, you can configure the properties of a host group to specify how much reserved computing resources (percentage of CPU usage, MBs of memory, MBs of disk space, and so on) you want to allocate to the host. You would do this to ensure that the host does not become starved for resources by running too many virtual machines on it.

Adding a Managed Host You add managed hosts to a host group by using the Add Hosts wizard. (See Figure 5-17.) This wizard can be used to add hosts belonging to an Active Directory domain or residing on a perimeter network. You can also add VMware ESX Servers to host groups as managed hosts provided these servers are managed by VMware VirtualCenter Server. When you add a host, you can select the host group you want to add the host to.




FIGURE 5-17 Adding a host in VMM using the Add Hosts wizard.

You can also add managed hosts that are part of a managed host cluster. If you try and add a host that belongs to a cluster, all the nodes in the cluster are added as hosts.

Removing a Managed Host If you no longer want to manage a host with VMM, you can remove the host using the Administrator Console. After you remove a host, it is no longer managed by VMM and the VMM Agent is automatically removed from the host. When you remove a host from VMM, the host and all of its associated virtual machines are removed from the VMM database and from the Administrator Console views. However, both the host and its virtual machines remain available and can be added as a host again to VMM if desired, or you can manage the host and its virtual machines outside of VMM using standard Hyper-V, Virtual Server, or VMware administrative tools.

To remove a managed host when the VMM Server can no longer communicate with that host or when you do not have the credentials for that host, use the *Remove-VMHost* Windows PowerShell cmdlet together with the *Force* parameter. When you specify the *Force* parameter, VMM does not prompt for or check credentials, and VMM does not attempt to connect to the host and uninstall the VMM agent. Using the *Force* parameter is recommended only when you need to remove stale host records from the VMM database.



Note If a single computer is serving as both a host and a library server and you remove the host role from the computer, the VMM Agent remains on the computer until the Library Server role is also removed.

Direct from the Source: Managed Host Agent Management

This sidebar deals with installing, reassociating, and updating managed host agents on managed hosts.

Installing the Agent Automatically

When you add a virtual machine host or Library Server, VMM 2008 R2 remotely installs a VMM Agent on the managed computer. The VMM Agent deployment process uses both the Server Message Block (SMB) ports and the Remote Procedure Call (RPC) port (TCP 135) and the DCOM port range. You can use either SMB packet signing or IPSec to help secure the agent deployment process. You can also install VMM Agents locally on hosts, discover them in the VMM Administrator Console, and then control the host using only the WinRM port (default port 80) and BITS port (default port 443).

Installing the Agent Manually

To install the VMM Agent manually, log on to the intended host computer and run Setup.exe from your VMM 2008 R2 product media. When the Setup splash screen is displayed, select Local Agent from under the Setup options.

Installing the Agent on Server Core

To manually install the VMM Agent on a Server Core installation of Windows Server 2008 R2, navigate to the `..\Prerequisites\VCRedist\amd64` directory on your VMM 2008 R2 product media and install the Visual C++ 2005 Redistributable Package by running `Vcredist_x64.exe`. After this finishes, navigate to the `..\amd64\msi\Agent` directory and execute the following `Msiexec.exe` command line to install the VMM Agent:

```
msiexec /I vmmagent.msi
```

After the agent has been installed, the Add Host Wizard can be initiated on the VMM Server to add the host to a host group or to the library.

Reassociating an Agent

Occasionally, you might want to move one or more hosts from one VMM Server to another. For example, you might want to consolidate hosts that are being managed by two or more VMM Servers onto a single VMM Server.

Alternatively, you might want to move one or more hosts back to a VMM Server on which they had been previously managed but from which they have not yet been removed. In this situation, the hosts still appear in the Managed Computers pane of Administration view on the original VMM Server. However, because the agents on the hosts have been associated with a different VMM Server, the host status on the original VMM Server is reported as Needs Attention in Hosts view and Access Denied in the Managed Computers pane in Administration view. Before you can manage the hosts again on the original VMM Server, you must reassociate the hosts with the original VMM Server.

To associate a host with a VMM Server, select the Reassociate Host With This Virtual Machine Manager Server check box in the Add Hosts Wizard or choose the host from the Hosts view and select Reassociate from the Actions pane. Note that the Reassociate command is enabled only when the status of an agent is reported as Access Denied because it is no longer associated with the current VMM Server.

Updating an Agent

After upgrading VMM 2008 R2 from its previous version, you must update the agents on all your managed hosts. For the agent update to work properly, the version of the agents on managed computers must match the version of the VMM Server. To update a VMM Agent on a managed host, select either the Hosts view or the Administration view in Managed Computers, and select the managed hosts on which you want to update the agent. In the Actions pane, click Update Agent.

—CSS Global Technical Readiness (GTR) team

Managing Hosts

When you select a managed host in the Administrator Console and click Properties in the Actions pane, the properties dialog box for the managed host is displayed. You can use the different tabs to configure various aspects of the host. The sections that follow describe the different configuration options available for managed hosts.

Summary tab This tab provides descriptive information about the host, including some limited system information such as CPUs, memory, storage capacity, operating system, and other information. (See Figure 5-18.) The version of the installed VMM Agent is also displayed on this tab.




FIGURE 5-18 The Summary tab of the properties of a managed host.

Status tab This tab provides overall status information for the host, as well as specific information about certain configuration settings on the host. (See Figure 5-19.)




FIGURE 5-19 The Status tab of the properties of a managed host.

The types of status information and their possible values that the Status tab can display include the following:

- Overall status
 - ❑ OK—No issues exist with any host status.
 - ❑ OK (Limited)—The VMware ESX Server host requires credentials and other security configurations before it can be fully managed by VMM.
 - ❑ Needs Attention—A problem exists with the status of one or more hosts.
- Connection status
 - ❑ Responding—The VMM Server is able to communicate with the agent on the host.
 - ❑ Not Responding—The VMM Server is unable to communicate with the agent on the host.
 - ❑ Access Denied—The VMM Agent is no longer associated with the VMM Server.
- Agent version
 - ❑ Up-to-date—The version of the VMM Agent is up to date.
 - ❑ Upgrade Available—The version of the VMM Agent must be upgraded to match the version of VMM Server.
 - ❑ Unsupported—The version of the VMM Agent is not supported for any VMM functions.
- Virtualization service status
 - ❑ Running—The virtualization service is started.
 - ❑ Stopped—The virtualization service is stopped.
- Virtualization service version
 - ❑ Up-to-date—The version of the virtualization software is up to date.
 - ❑ Upgrade Available—The version of the virtualization software must be upgraded to a version supported by VMM.
 - ❑ Unsupported—The version of the virtualization software is not supported for any VMM functions.

Additional status information will be displayed if the host belongs to a cluster. The host status values do not change in the Administrator Console until the VMM Server performs a host refresh, which by default runs automatically every 30 minutes. You can also perform a refresh on demand by right-clicking the host and clicking Refresh.

VMs This tab provides basic information concerning the virtual machines that are hosted on the selected managed host. In Figure 5-20, for example, there are four virtual machines on the host, and the selected virtual machine displays its status, up time, and virtual hardware resources on the tab.




FIGURE 5-20 The VMs tab of the properties of a managed host.

There is also an additional setting on this tab, Register The Following Virtual Machines, that enables you to register virtual machines that have files located on the host but have not been added as a virtual machine. Registering a virtual machine is done automatically when you create, deploy, or migrate a virtual machine to a host. However, you might have files for a virtual machine that have not been added on a host or in VMM as a managed virtual machine. Registering these virtual machines adds them to the host and allows for management using VMM.

Reserves This tab lets you to configure resource parameters (CPU percentage, memory, disk space, maximum disk I/O, and network capacity) that will be used to determine whether a virtual machine can be hosted on the host. (See Figure 5-21.). If the host does not meet all these requirements, virtual machines cannot be placed on that host.




FIGURE 5-21 The Reserves tab of the properties of a managed host.

Hardware This tab displays information concerning the hardware configuration of the host. (See Figure 5-22.) You can use this tab to specify whether or not a specific volume is available for placement, to override the discovered network location for a specific network adapter, and for other purposes.




FIGURE 5-22 The Hardware tab of the properties of a managed host.

Networking This tab displays information concerning the virtual networks that are configured on the host to support virtual machines running on the host. Virtual Networks can be added, modified, or removed by using this tab. The type of virtual networks displayed here depends on the type of virtualization running on the host. For example, Figure 5-23 shows the virtual networks configured on a Hyper-V host.




FIGURE 5-23 The Networking tab of the properties of a managed host.



More Info For more information concerning virtual networks, see Chapter 2 in this book.

Placement This tab displays information concerning the default paths available on the host for placement of virtual machines. (See Figure 5-24.) If a path was not specified when the host was added to VMM, all known paths will be displayed.




FIGURE 5-24 The Placement tab of the properties of a managed host.

Remote This tab displays the port used by VMRC for remote connection. (See Figure 5-25.) The default port for VMRC in Hyper-V is port 2179. (In Virtual Server 2005 R2 SP1, this was port 5900.)




FIGURE 5-25 The Remote tab of the properties of a managed host.

Custom This tab lets you add custom properties to the host for informational purposes. (See Figure 5-26.) To add more columns, including custom properties, to the Results pane of the Hosts view, select Columns from the toolbar to open the Select Columns dialog box. You can also right-click on the column header to add or remove columns.




FIGURE 5-26 The Custom tab of the properties of a managed host.

Other settings for managed hosts In addition to the settings available on the Properties page of a managed host, more information concerning the host is displayed in the Details pane at the bottom middle of the Administrator Console. To view this information, choose the Host view to display a list of hosts in the Results pane and then select one of these hosts.

The Details pane displays three tabs of information: Summary, Networking And Storage, and Latest Job. (See Figure 5-27.) The Summary tab provides general information concerning the host including which virtual machines are currently being hosted. The Networking And Storage tab shows which networks are connected and what storage is available, plus some usage statistics. The Latest Job tab displays information concerning the last job that was run against the host.




FIGURE 5-27 Additional information concerning managed hosts is displayed in the Details pane.

Using Networking View

There is also a Networking view that is available for managed hosts. The Networking view shows what networks the host is connected to and what networks its hosted virtual machines are connected to. Figure 5-28 shows how you can specify the scope of the Networking view by selecting the hosts, host groups, or both that you want to display using this view.




FIGURE 5-28 Specifying the scope to be used for the Networking view.

After you've specified the scope you want to use, the Networking view is displayed as a separate window as shown in Figure 5-29.




FIGURE 5-29 Displaying the Networking view window.

Managing Host Clusters

You can manage host clusters in a similar way to how you manage individual hosts as previously described. The main difference is that in a host cluster the hosts are grouped together. Plus there is some additional management functionality available. A benefit of using host clusters in VMM 2008 is that you can take advantage of the integration between VMM 2008 and OpsMgr 2007 using the Performance and Resource Optimization (PRO) feature introduced in VMM 2008. PRO provides workload-aware and application-aware resource optimization within host clusters and host groups. PRO is implemented through PRO tips, which can recommend or automatically implement actions such as virtual machine migration or virtual machine right-sizing based on policies implemented using OpsMgr 2007. For example, if you configure a host cluster for automatic implementation of PRO tips, VMM 2008 can migrate virtual machines automatically between nodes in a host cluster.

Direct from the Source: Managing a VMware Infrastructure 3 Environment

VMM 2008 manages VMware through the VMware Infrastructure API, so to start managing VMware, you must add a VMware VirtualCenter Server to VMM. When you add a VMware VirtualCenter Server to Virtual Machine Manager, all existing ESX Server hosts managed by the VirtualCenter Server are also added to VMM.

Requirements for Managing VMware

Virtual Machine Manager supports the following VMware releases:

- Virtualization managers:
 - VMware VirtualCenter 2.5
 - VMware VirtualCenter 2.0.1
- Virtual machine hosts:
 - VMware ESX Server 3.5
 - VMware ESX Server 3.0.2
 - VMware ESX Server 3i

Adding ESX Server Hosts

Your first step in configuring VMM to manage a VI3 environment is to add the VirtualCenter Server so that VMM can use the VMware Infrastructure API to manage the ESX Server hosts and virtual machines. When you add the VirtualCenter Server, you specify whether or not to communicate with the ESX Server hosts in Secure Mode. If you choose Secure Mode, you must provide a certificate and a public key in addition to credentials for each ESX Server host.

To Add a VMware VirtualCenter Server

1. In any view of the VMM Administrator Console, click Add VMware VirtualCenter Server.
2. Supply the computer name, port (default of 443) and administrative credentials of the VMware VirtualCenter Server to add.
3. Under Security, specify whether or not to communicate with the ESX Server hosts in secure mode.

In secure mode, a certificate and public key are required for each ESX Server host. Clear this option if you want to trust communications and require only credentials for the hosts. (See Figure 5-30.)

4. If the VirtualCenter Server has a self-signed certificate, to verify the server's identity, you must import the server's security certificate to the local machine certificate store.

After the server is added to VMM, the ESX Server hosts are added to VMM in their own host group hierarchy. This operation might take several minutes although VMM adds the ESX Server hosts and then refreshes virtual machine data on the new hosts.




FIGURE 5-30 Adding a VMware VirtualCenter Server

Configuring Security for Individual VMware ESX Server Hosts

Newly discovered ESX Server hosts initially have OK (Limited) status in VMM. To fully manage the virtual machines on the hosts, you must enter credentials in the host properties. If you are managing the VirtualCenter Server in Secure Mode, you also must provide a certificate and public key. For information about the restrictions that apply to virtual machines on a host that has OK (Limited) status, see Table 5-7.

TABLE 5-7 Supported Virtual Machine Actions for VMware ESX Server Hosts by Host Status

Virtual Machine Action	OK(Limited)	OK
Start	Yes	Yes
Stop	Yes	Yes
Pause	Yes	Yes
Modify properties	Yes	Yes
Create a new checkpoint	Yes	Yes
Manage checkpoints	Yes	Yes
Remove	Yes	Yes
Migrate with VMotion	Yes	Yes
Migrate across VirtualCenter Server	No	Yes
Save state	No	Yes

Virtual Machine Action	OK(Limited)	OK
Discard a saved state	No	Yes
Store in the VMM library	No	Yes
Clone within the same VirtualCenter	No	Yes
Clone on the same ESX Server host	No	Yes
Perform V2V conversion	No	Yes
Create from a VMM template	No	Yes
Create from a blank disk	No	Yes

To change the status of the ESX Server host to OK so that you can fully manage virtual machines, you must provide credentials and, if security mode is enabled on the host, upload a certificate and public key from the host to VMM.

Setting Credentials for Host Communication

To communicate with a virtual machine host that is on a perimeter network in a domain that does not have a two-way trust with the VMM Server's domain, VMM uses credentials for an account that has administrative privileges on the host. To communicate with a VMware ESX Server host, VMM uses credentials for an account that has administrative privileges on the host and, if the host is in secure mode, a certificate and public key.

To update the credentials for communicating with a host, follow these steps:

1. In the Hosts view, select the host on which you want to update the credentials.
2. In the Host area of the Actions pane, click Properties.
3. Click the Security tab and then do one of the following:
 - ❑ For a host on a perimeter network, type the credentials for the local agent service account and then click OK.
 - ❑ For a host on a nontrusted domain, type the credentials for account on the host that has administrative privileges and then click OK.
 - ❑ For a VMware ESX Server host, in the Credentials For This Host area, type the user name and password for an account that has administrative privileges on the host, and then type the password again to confirm it. If the host is in secure mode, in the Certificate And Public Key area, click Retrieve to upload the certificate and public key from the host select the Accept Both The Certificate And Public Key For This Host check box, and then click OK. Note that after adding a VirtualCenter Server to VMM, you can use the Add Hosts Wizard to manually add any new ESX Server hosts to VMM. VMM does not discover the new hosts automatically. Note also that removing an ESX Server host removes it immediately from VMM and also from the VirtualCenter Server.

Working with the Library

The VMM library is basically a catalog that provides access to file-based resources such as ISO images, virtual hard disks, scripts, and other types of files that are stored on VMM Library Servers. The library also provides access to virtual machine templates, guest operating system profiles, and hardware profiles you create and that reside in the VMM database. You can also store virtual machines in the library when they are not in use.

When you install your VMM Server, a Library Server is automatically installed on the server. The default location for the Library Server share on a VMM Server is `%SystemRoot%\ProgramData\Virtual Machine Manager Library Files`, and the share name is `MSSCVMMLibrary`. By default, when you install a Library Server, the only directory that is created is the one that holds two built-in .vhd files: Blank Disk Small and Blank Disk Large. You can then expand the folder structure under the library share to include folders for storing operating system installation images (ISO files), Windows PowerShell scripts, stored virtual machines, and other kinds of resources.

Library Server Requirements

Library Servers must meet the following requirements:

- Must be running either Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003 SP1 or later. For highly available file servers, Windows Server 2008 Failover Clustering must be configured.
- Must be in an Active Directory domain that has a two-way trust relationship with the VMM Server's domain.



Note VMM does not support file servers configured with the case-sensitive option for Windows Services for UNIX. (The NFS Case Control is set to Ignore.)

Adding Library Servers

In a distributed environment with branch offices, it's a good idea to have additional Library Servers available to support remote locations, especially when users are using a Self-Service Portal to create virtual machines for their own use. This is because having all the necessary files available locally to support a virtualized environment can enhance performance and reduce network traffic. To accomplish this, you can add more Library Servers to VMM and then arrange for the replication of the necessary files to these new servers.

To add a Library Server, open the Administrator Console and select any view. In the Actions pane on the right side, select Add Library Server to launch the Add Library Server Wizard and then follow the prompts. If you have multiple Library Servers, you can organize them

into Library Server groups in a similar way to how managed hosts can be organized into host groups. Figure 5-31 shows the Library view in the Administrator Console with a single Library Server and its resources visible. After you've deployed a Library Server, you can create additional library shares on it and add files to your library.




FIGURE 5-31 Library view in the Administrator Console.

Adding Virtual Machine Templates

Virtual machine templates can be used to create new virtual machines repeatedly using standardized hardware and software settings. Self-service users must use assigned templates to create their own virtual machines.

A virtual machine template is a library resource consisting of the following components:

- **Virtual hard disk** You can use a generalized virtual hard disk from the library, or you can create a virtual hard disk from an existing virtual machine. If the source virtual machine for your template has multiple virtual hard disks, select the disk that contains the operating system. To simplify the generalization process, include Virtualization Guest Services (such as Virtual Machine Additions or Integration Services) in your template.
- **Hardware profile** To define a standard set of hardware settings, you can create a hardware profile and associate it with a template. When you create a new template or create a virtual machine from a template, you can specify the virtual hardware settings

or reuse an existing hardware profile from the library. For more information about hardware profiles, see the section titled “Adding Hardware Profiles” that follows.

- **Guest operating system profile** To maintain guest operating system standardization in deployed VMs using templates, you can attach a guest operating system profile from the library to the template. When you create a new template or create a virtual machine from a template, you can specify the settings manually or use an operating system profile associated with your answer files. For more information about guest operating system profiles, see the section “Adding Guest Operating System Profiles” that follows.

Virtual machine templates can be created in the Library view in the Administrator Console. In the Actions pane on the right side, select New Template to launch the New Template Wizard and then follow the prompts.

Adding Hardware Profiles

Hardware profiles are library resources that contain hardware specifications that can be applied to a new virtual machine or to a virtual machine template. For example, hardware profiles can contain specifications for CPU, memory, network adapters, a DVD drive, a floppy drive, COM ports, and the priority given to a virtual machine when allocating resources on the virtual machine’s host.

You can create hardware profiles that import a standard hardware configuration into a template or into a virtual machine. The options are the same whether you update the hardware configuration of a virtual machine, hardware profile, or virtual machine template. After specifying a hardware profile for a virtual machine, you can change the settings that were imported. The changes do not affect the hardware profile. No association is maintained with the hardware profile after the virtual machine is created.

Hardware profiles are managed in the Library view of the Administrator Console. To create a hardware profile, click the New Hardware Profile action in the Library view. You can also save a new hardware profile based on the hardware configuration of an existing virtual machine or virtual machine template.

Adding Guest Operating System Profiles

A *guest operating system* is the operating system that runs on a virtual machine. By contrast, a *host operating system* is the operating system that runs on the physical computer (the virtual machine host) on which one or more virtual machines are deployed. In VMM 2008, a guest operating system profile is a collection of operating system settings that can be imported into a virtual machine template to provide a consistent operating system configuration for virtual machines that are created from that template.

You create guest operating system profiles to provide standard settings for the operating systems on your virtual machines. Guest operating system profiles are database objects that do not have any physical files associated with them. Guest operating system profiles are configured in the Library view, where they are displayed in a special VMs And Templates folder as shown in Figure 5-32.



FIGURE 5-32 The VMs And Templates folder, which contains guest operating system profiles.

To create a guest operating system, use the New Guest OS Wizard in the Library view. Alternatively, you can save a guest operating system based on your current settings while you are creating a template. After the template is created, no association is maintained between the template and the guest operating system profile that was used with it. Any changes that are made to a guest operating system profile affect only new templates that are created after the changes are made.

Removing Library Resources

Library resources can be removed as well as added. In addition, specific files stored in the library can be disabled, library shares can be removed, and Library Servers can be removed from VMM.

When you no longer need a file in the VMM library, you can remove the file using the Administrator Console. If you remove the file from a library share outside of VMM, any template, guest operating system profile, hardware profile, or virtual machine that uses the file must be repaired to remove the reference to the deleted file. If you use the Remove action in the Library view to remove the file, VMM lists any virtual machines, templates, or guest operating system profiles that reference the file, and, if you choose to proceed, VMM removes references to the deleted file from the dependent resources.

Table 5-8 summarizes the resource dependencies that exist for different resource types. These dependencies involve only physical files. When a guest operating system profile, hardware profile, or template is used to create a virtual machine or a template, the settings are imported without maintaining any link to the source configuration.

TABLE 5-8 Resource Dependencies for Different Resource Types

Resource Type	Dependent Resources
Templates	ISO images, scripts, answer files, virtual hard disks, and virtual floppy disks.
Guest Operating System Profiles	Scripts and answer files.
Hardware Profiles	ISO images, virtual floppy disks.
Virtual Machines	The virtual hard disks and virtual floppy disks attached to a virtual machine are indexed in Virtual Machine Manager, but they are not displayed in the library because the files are not available for use with other resources.

Direct from the Source: Backing Up and Restoring the VMM 2008 R2 Database

The VMM 2008 R2 library comprises both flat files in the library share and the VMM database. The VMM database is a Microsoft SQL Server database that contains not only Library information but all VMM configuration information, including managed hosts, host configuration settings, VM settings, and other information. It is therefore important to back up the VMM database regularly as part of a comprehensive backup plan for protecting all VMM data, including data on hosts, virtual machines, and Library Servers. In addition to using the tools provided in VMM, you can also use SQL Server Management Studio to back up and restore the VMM database.

To Back Up the VMM Database

1. In the Administration view, click General, and in the Actions pane, click Back Up Virtual Machine Manager.
2. In the Virtual Machine Manager Backup dialog box, type the path for a destination folder for the backup file.

Note The folder must not be a root directory and must be accessible to the SQL Server.

To Restore the VMM Database on the Same Computer

1. To restore the VMM database, run the Scvmmrecover.exe tool from the command line. The Scvmmrecover.exe tool is located on the product CD at the following path: %cddrive%\i386\bin for a 32-bit computer, or %cddrive%\amd64\bin for a 64-bit computer.
2. On the VMM database computer, open a command-prompt window with elevated privileges, and then run the Scvmmrecover.exe tool using the following syntax:

SCVMMRecover [-Path <location>] [-Confirm]

where <location> is the Location of the Virtual Machine Manager database backup.

3. If any hosts has been added or removed since the database backup and the physical computer that you are restoring the VMM database on has the same System Identification (SID) number as the computer it was on before, then in the VMM Administrator Console, in the Hosts view, you must remove any hosts that might have been removed from VMM since the last backup was created. If a host has been removed from VMM after the last backup was created, it will have a status of Needs Attention in the Hosts view, and any virtual machines on that host will have a status of Host Not Responding in the Virtual Machines view. Then you must add any hosts that might have been added since the last backup.

To Restore the VMM Database on a Different Computer

1. To restore the VMM database, run the Scvmmrecover.exe tool from the command line. The Scvmmrecover.exe tool is located in %ProgramFiles%\Microsoft System Center Virtual Machine Manager 2008\bin.
2. On the VMM database computer, open a command-prompt window with elevated privileges, and then run the Scvmmrecover.exe tool using the following syntax:

SCVMMRecover [-Path <location>] [-Confirm]

where <location> is the location of the Virtual Machine Manager database backup.

3. Because the VMM Server is a different computer and the existing hosts are now not associated with this new VMM Server computer, you must perform the following steps to reassociate the hosts with the VMM Server:
 - a. In the VMM Administrator Console, in the Administration view, click Managed Computers, and in the Results pane, identify any managed computers with a status of Access Denied. Then click a managed computer with a status of Access Denied, and in the Actions pane, click Reassociate.
 - b. If any hosts have been added or removed since the database backup, in the VMM Administrator Console, in the Hosts view, remove any hosts that might have been removed from VMM since the last backup was created. If a host has been removed from VMM after the last backup was created, it will have a status of Needs Attention in the Hosts view and any virtual machines on that host will have a status of Host Not Responding in the Virtual Machines view. Then add any hosts that might have been added since the last backup.
 - c. If any VMs have been removed since the database backup, in the VMM Administrator Console, in the Virtual Machines view, remove any of those virtual machines. If a host is present but has a virtual machine that was removed since the last backup, the virtual machine will have a status of Missing in the Virtual Machines view.

Working with Virtual Machines

You can use VMM 2008 R2 to create and manage virtual machines hosted on Hyper-V, Virtual Server, or VMware ESX Server managed hosts. You can also use VMM to perform P2V and V2V conversions and perform other tasks relating to virtual machines.

Creating New Virtual Machines

You can use VMM 2008 to create a new virtual machine from a template, from an existing virtual machine, or using a blank virtual hard disk. You can also use VMM to clone a virtual machine. The following walkthrough illustrates how to create a new virtual machine from a template.

Figure 5-33 shows a host named HVR2 running Windows Server 2008 R2 with Hyper-V. The Results pane shows that the cluster currently has four virtual machines on it: NYC-SRV-001, NYC-SRV-005, NYC-SRV-008, and NYC-CLI-233.




FIGURE 5-33 A managed host running three virtual machines.

Let's create a new virtual machine and place it on the host. To do this, click New Virtual Machine in the Actions pane. This launches the New Virtual Machine Wizard. (See Figure 5-34.)




FIGURE 5-34 New Virtual Machine Wizard.

The New Virtual Machine Wizard gives you two options:

- You can create a new virtual machine from a template, an existing virtual machine, or an existing virtual hard disk.
- You can create a new virtual machine using a blank virtual hard disk. When you install your VMM Server, two blank virtual hard disks are added by default to the Library Server. These blank VHDs are named Blank Disk – Small and Blank Disk – Large. Both of these disks are dynamically expanding virtual hard disks (VHDs) that differ only in how large they can expand to. Specifically, the small VHD can expand to a size of 16 GB while the large VHD can expand to 60 GB.

To create a new virtual machine from a template, select the first option and click Browse. This opens the Select Virtual Machine Source dialog box, which displays a list of virtual machine templates and available virtual hard disks to choose from. We'll select the template named NYC-SRV-W2K8R2ENT, which is configured to use Windows Server 2008 R2 Enterprise Edition as the guest operating system for the new virtual machine. (See Figure 5-35.)




FIGURE 5-35 Selecting a virtual machine template.

Selecting the desired template and clicking OK returns you to the New Virtual Machine Wizard. (See Figure 5-36.)




FIGURE 5-36 Creating a new virtual machine from a template.

Continuing through the wizard, the next thing you must do is give your new virtual machine a name. (See Figure 5-37.)




FIGURE 5-37 Naming the new virtual machine.

On the next page of the wizard, you configure the virtual hardware settings for your new virtual machine. Virtual hardware settings include BIOS, processor, memory, IDE controllers, network adapters, and other virtual devices that collectively represent the hardware profile of the virtual machine. Figure 5-38 shows the virtual machine being configured to use two logical processors.




FIGURE 5-38 Configuring virtual hardware settings for the virtual machine.

On the next wizard page, you specify a password for the local Administrator account of the guest operating system for the virtual machine. (See Figure 5-39.)




FIGURE 5-39 Specifying a password for the Administrator account.

The next page lets you choose whether to place your new virtual machine on a host or store it in the library. (See Figure 5-40.)




FIGURE 5-40 Choosing whether to deploy or store the virtual machine.

Selecting the option to deploy the virtual machine on a host brings up a list of possible hosts you can deploy the machine on. Figure 5-41 shows that we have chosen to deploy the new virtual machine onto HVR2, which is the only available host in the New York hosts group. The Details portion of this wizard page shows that this cluster node currently has four other virtual machines on it.




FIGURE 5-41 Choosing to place the new virtual machine on a node in a host cluster.

The next wizard page lets us specify the path on the host to where you will save the file associated with the virtual machine. (See Figure 5-42.)




FIGURE 5-42 Specifying the path for storing the virtual machine files.

In the next page of the wizard, you select the virtual network that the new virtual machine will use. (See Figure 5-43.)




FIGURE 5-43 Selecting a virtual network for the new virtual machine.

The Additional Properties page of the wizard lets you specify what action should be performed on the virtual machine when the physical server stops or starts. (See Figure 5-44.)




FIGURE 5-44 Specifying what happens to the virtual machine when the host stops or starts.

The final page of the wizard lets you review the settings you have chosen for your new virtual machine. (See Figure 5-45.) After you complete the wizard, the new virtual machine is created and deployed to the host cluster you specified.




FIGURE 5-45 Summary page of New Virtual Machine Wizard.

Clicking Create now creates the new virtual machine using the template you specified. The Jobs window opens to display the progress of this task. (See Figure 5-46.)




FIGURE 5-46 The progress of virtual machine creation is displayed in the Jobs window.

Managing Virtual Machines

You can use the Administrator Console to manage virtual machines. (See Figure 5-47.) The following are some of the management actions that can be performed on virtual machines:

- **Start** Starts a virtual machine that is stopped, paused, or in a saved state.
- **Stop** Stops a virtual machine and does not save any state information. This action has the same effect on the virtual machine as does pulling the plug on a physical computer.
- **Pause** Suspends execution of a virtual machine, and keeps all virtual machine state in memory.
- **Save State** Suspends execution of a virtual machine, and saves the current virtual machine state to disk to release memory and CPU resources for other virtual machines. When the virtual machine is restored from the saved state, it returns to the condition that it was in when its state was saved.

- **Discard Saved State** Discards the state that was saved for a virtual machine that is in a saved state, and turns off the virtual machine.
- **Shut Down** Shuts down the guest operating system on the virtual machine.
- **Connect To Virtual Machine** Connects to a virtual machine by using Remote Desktop Protocol (RDP). VMM attempts to connect to a running virtual machine that you select in the Results pane and adds a thumbnail of the connection—the desktop of the virtual machine—to the virtual machine details. To open a larger connection window so that you can log on to the virtual machine, you can either double-click the thumbnail or click Connect To Virtual Machine in the Actions pane.
- **Repair** Repairs a failed virtual machine by retrying the action that caused the failure, restoring the virtual machine to its state before the action caused it to fail, or refreshing the data for the virtual machine in VMM after mitigating the issue outside VMM.
- **Install Virtual Guest Services** Installs Virtual Guest Services, such as Integration Services (Hyper-V) or Virtual Machine Additions (Virtual Server) on the virtual machine.
- **Disable** Disables a virtual machine that is stored in the library to temporarily prevent use of the virtual machine. A disabled virtual machine remains in the library but cannot be deployed or repaired. (This option is available in Library view only.)
- **Properties** Modifies the properties of a virtual machine.




FIGURE 5-47 Managing virtual machines using the Administrator Console.

You can also use the Administrator Console to view and configure the settings for virtual machines. The settings for a virtual machine can be accessed either in the Virtual Machines view or the Library view, depending on whether or not the virtual machine is deployed to a host or is stored in the library. The settings that can be modified also depend on the current state of the virtual machine.

Configuring Virtual Machine Settings Virtual machine settings can be accessed by right-clicking on the virtual machine and selecting Properties or by clicking Properties in the Actions pane when in the Virtual Machine view. As shown in Figure 5-48, the properties dialog box for a virtual machine has six tabs with configurable settings as follows:

- **General** This tab is used to view general information concerning the virtual machine. The only setting that is not available for modification if the virtual machine is running, saved, or paused is the name of the virtual machine. Significant entries here are Owner and Tag. For example, if the virtual machine will be used as part of the Self-Service Portal, the Self-Service Portal user or users should be designated as the owner or owners of the virtual machine or they will be denied access. Additionally, tag entries can be used as part of the deployment or placement process, where virtual machines can be deployed only to managed hosts that meet specific requirements based on tags that are set on templates in the library that Self-Service Portal users have access to.
- **Hardware Configuration** This tab documents the hardware configuration of the virtual machine. The settings here can be changed only if the virtual machine is not running and its status shows as Stopped.
- **Checkpoints** This tab shows the checkpoints (snapshots in Hyper-V terminology) that have been taken for the virtual machine. You can also create a new checkpoint, remove a checkpoint, or restore a configuration to a virtual machine.
- **Custom Properties** This tab allows for configuration of up to 10 custom fields for the virtual machine. You can use custom fields to identify, track, and sort virtual machines by any property, including department, geographic area, or function.
- **Settings** This tab allows for the configuration of Self-Service Quota Points and Physical Resource Optimization (PRO) settings. By default, all virtual machines are allocated 1 quota point.
- **Actions** This contains settings that determine what happens when the physical host server starts and stops.

The settings available for when the physical server starts are

- ☐ Never Automatically Turn On The Virtual Machine
- ☐ Always Automatically Turn On The Virtual Machine
- ☐ Automatically Turn On The Virtual Machine If It Was Running When Virtual Server Stopped

The settings available for when the physical server stops are

- ☐ Save State
- ☐ Turn Off Virtual Machine
- ☐ Shut Down Virtual Machine




FIGURE 5-48 Configuring the properties of a virtual machine.

Connecting to a Virtual Machine To connect to a running virtual machine that is visible in the Administrator Console, do one of the following:

- Select the virtual machine in the Results pane, and click Connect To Virtual Machine in the Actions pane.
- Right-click on the virtual machine in the results pane, and select Connect To Virtual Machine.
- Double-click on the thumbnail image of the virtual machine in the Details pane.

When you perform one of these actions, the VMM Virtual Machine Viewer (VirtualMachineViewer.exe) opens a separate window that allows you to interact with the running virtual machine by using the keyboard and mouse. (See Figure 5-49.)




FIGURE 5-49 Interacting with a running virtual machine using the Virtual Machine Viewer.

Direct from the Source: Connecting to Virtual Machines

Remote connections to virtual machines running on managed hosts can be made in the VMM Administrator Console. VMM uses Virtual Machine Remote Control (VMRC), which is a feature of Virtual Server 2005 that lets you enable, disable, and configure virtual machines from within Virtual Machine Manager. Remote control options and the connection ports for virtual machines on virtual machine hosts vary depending on the virtualization software running on the host and the operating system running on the computer on which the VMM Administrator Console is installed. The settings for the VMRC are configured under the General settings in the Administration view.

To connect to a virtual machine, the following requirements must be met:

- Credentials must be provided to log on to the virtual machine. Local administrator credentials can be configured during virtual machine creation (if using a template).
- For Microsoft Virtual Server, you must have specified the port on the VMM Server to use for remote control. RDP must be enabled on the virtual machine host.

You can also use remote control to access the guest operating system of a virtual machine much like you can access Windows by using Remote Desktop. However, unlike Remote Desktop, which allows you to connect only if the operating system is running, remote control in VMM allows you to access a virtual machine prior to the guest operating system startup.

Controlling Hyper-V Hosts

To control virtual machines on a Hyper-V host by remote control, VMM uses either VMConnect or RDP, depending on the operating system that is running on the computer that the VMM Administrator Console is installed on.

For all supported versions of Windows Server 2008, Windows Server 2008 R2, Windows 7, and Windows Vista with SP1 or later, VMM uses VMConnect with the default port of 2179. You can change the default port that is used for connecting to virtual machines on new Hyper-V hosts. For all other supported operating systems, VMM uses RDP with the default port of 3389.

Controlling Virtual Server Hosts

To control virtual machines on a Virtual Server host by remote control, VMM uses VMRC.

The VMRC client connects to an instance of Virtual Server on a host and allows you to access its virtual machines. Through VMRC, you can use a virtual machine as if you were using it through the Virtual Server Administration Web site. However, VMRC does not provide the administrative capabilities available in the Administration Web site, such as creating a new virtual machine or changing a virtual machine configuration. You can perform those functions in VMM.

You can enable and configure VMRC settings when you add a host or after you have added the host.

You can set up VMRC access accounts to give administrators access to all virtual machines on all managed hosts. To gain access to a virtual machine through VMRC, an administrator must be logged on under an account that has administrative credentials on the local host computer.

By default, when you add a host to VMM, VMRC is enabled and uses the following settings:

- The connection port is set to the global default port setting for Virtual Server hosts as specified in General settings in the Administration view.
- No connection time-out is enabled.

- Only one user at a time is allowed to connect to a virtual machine.
- The VMRC connection is not encrypted. (You can modify this setting only after a host has been added.)



Note It is recommended that you implement Secure Sockets Layer (SSL) security for VMRC connections, particularly if you use Basic authentication, which transmits passwords in plain text.

If you change the VMRC port, the port setting you assign for the hosts must identically match the port settings that are assigned in Virtual Server.

You can allow multiple users to connect to the same virtual machine. However, each user can access the guest operating system without the knowledge of the other users. This is by design for training and lab scenarios, where one user wants to demonstrate a task to other users and have them connect to and view the same remote session. Because VMRC connections do not use sessions, allowing more than one user to connect can result in collisions.

You can use SSL to encrypt communications over the VMRC connection by uploading a certificate from an appropriate internal or third-party certification authority.

Controlling ESX Server Hosts

To control virtual machines on VMware ESX Server hosts by remote control, VMM uses the VMware mouse keyboard and screen (MKS) Client with default port 902. You cannot change the remote control settings or default port for an ESX Server host in VMM.

Before connecting to a VM on an ESX Host, you must install the VMware ActiveX control. To install the ActiveX control, in the Virtual Machines view, from the Results pane, select an ESX host. In the Summary tab of the details pane, the expected thumbnail view of the VM will have the following message: "VMware ActiveX control not installed. To view a thumbnail for this virtual machine, you must install the VMware ActiveX component." Select Install ActiveX Control, and follow the instructions. Note that thumbnails are not visible for 64-bit VMware clients.

Connecting to a Virtual Machine

To connect to a virtual machine, locate the host it is running on and either double-click the thumbnail view in the middle pane or select the Connect To Virtual Machine action. This will open the Virtual Machine Viewer (Hyper-V), the VMRC Viewer (Virtual Server 2005), or the VMware Viewer (VMware ESX), accordingly.

—CSS Global Technical Readiness (GTR) team

Deploying Virtual Machines

You can use VMM 2008 R2 to deploy virtual machines stored in the library to managed hosts. Deploying a virtual machine stored in the library removes the virtual machine from the library and places it on the target host. Alternatively, you can use the New Virtual Machine Wizard by selecting to use an existing virtual machine or VHD in the library if you want to keep the original virtual machine in the library as a source for additional future deployments.

To deploy a virtual machine, select the virtual machine from the appropriate Library Server and click Deploy. This launches the Deploy Virtual Machine Wizard. Then you follow the prompts.

Migrating Virtual Machines

You can use VMM 2008 R2 to migrate virtual machines between hosts that are running the same virtualization platform (Hyper-V, Virtual Server, or VMware ESX Server). You can also migrate virtual machines from Virtual Server to Hyper-V.

The following methods are available to migrate a deployed virtual machine to a different host when in the Virtual Machines view:

- Click the Migrate action to launch the Migrate Virtual Machine Wizard, which enables you to select a suitable host, specify the path that will store the virtual machine, and select the type of transfer to be performed.
- Drag and drop the virtual machine onto a host.
- Drag and drop the virtual machine onto a host group. Through automatic placement, the virtual machine is placed on the most suitable host that is available in the host group based on the virtual machine's requirements and the group or individual host rating metrics.

The transfer methods available to migrate Hyper-V and Virtual Server virtual machine include SAN transfers (when a properly configured SAN is available) and Network (LAN) transfers. Note that if you migrate a virtual machine that is connected to SAN storage, the virtual machine will not be able to reconnect to the SAN unless the destination host also has access to that SAN. VMM is not able to detect whether a virtual machine is connected to a SAN or whether the destination host is connected to the same SAN, and therefore it cannot provide a warning. You must ensure that the new host is configured to allow the virtual machine to reconnect to the SAN before you migrate the virtual machine.

Cloning Virtual Machines

You can use VMM 2008 R2 to “clone” (make an exact copy) of a virtual machine. The virtual machine you want to clone can be one that already exists in the library or one that is already deployed to a host (in which case, it must be turned off). The cloning process can also be used to back up virtual machines.



Note A cloned virtual machine has the same security identifier (SID) as the original, so they cannot be running simultaneously on the same network.

When you clone a virtual machine, you cannot make changes to the operating system settings, but you can make changes to the hardware profile.

To begin the cloning process, first make sure that the virtual machine is shut down on the host. Then select the virtual machine in the Administrator Console, click the Clone action to launch the New Virtual Machine Wizard, and then follow the prompts.

Performing P2V Conversions

You can use VMM 2008 R2 to convert a physical computer to a virtual machine. The process of doing this is known as Physical-to-Virtual (P2V) conversion. You might typically use the P2V conversion process as part of a server consolidation project where underutilized physical servers are converted into virtual machines so that these physical servers can be repurposed or decommissioned.



Tip You can use System Center Operations Manager 2007 R2 to generate a Virtualization Candidates report that can help identify underutilized physical servers that could be good candidates for the P2V conversion process.

When identifying your best candidates for P2V conversion, consider converting the following types of servers, which are listed in order of preference:

- Non-business-critical underutilized servers. By starting with the least utilized servers that are not business critical, you can learn how the P2V process works with relatively low risk. Web servers can often make good candidates here.
- Servers with low utilization that are hosting less critical in-house applications.
- Servers with higher utilization that are hosting less critical applications.
- Other underutilized servers.

The requirements for a physical source computer depend on whether you are going to perform an online or offline P2V conversion:

- **Online P2V conversion** In this scenario, VMM uses BITS to copy data while the source computer continues to service user requests. The source computer is not restarted during the conversion, and the Volume Shadow Copy Service (VSS) is used to ensure data consistency. To perform an online P2V conversion, the source computer must have at least 512 MB of RAM.
- **Offline P2V conversion** In this scenario, the physical source computer restarts into the Windows Preinstallation Environment (Windows PE) before VMM converts the physical disks to VHDs. The Windows PE environment is used by the conversion process to capture the disk and therefore requires network and storage controller drivers for Windows PE. By default, VMM scans for these devices during the P2V conversion process and injects the required drivers into the Windows PE environment automatically. If third-party network interface card (NIC) or storage controller drivers are needed, however, these drivers must be provided by the administrator.

In either case, VMM temporarily installs the VMM Agent on the physical source computer that you want to convert.

The following operating systems are supported for P2V conversions:

- Windows Server 2008 R2
- Windows Server 2008 (32-bit)
- Windows Server 2008 (64-bit)
- Windows Server 2003 (32-bit)
- Windows Server 2003 (64-bit)
- Windows 2000 Server SP4 or later (Offline P2V only)
- Windows 2000 Advanced Server SP4 or later (Offline P2V only)
- Windows XP Professional (32-bit) SP2 or later
- Windows XP Professional (64-bit) SP2 or later
- Windows Vista SP 1 (32-bit)
- Windows Vista SP1 or later (64-bit)
- Windows 7 (32-bit)
- Windows 7 (64-bit)

Although P2V conversion of source computers running Windows NT Server 4.0 is not supported by VMM 2008 R2, you can use the Microsoft Virtual Server 2005 Migration Toolkit (VSMT) or third-party solutions for converting computers running Windows NT Server 4.0 into virtual machines.



Note When performing a P2V conversion of Windows 2000 Server or Advanced Server, the following components must be installed on the source computer:

- Service Pack 4
- Windows Installer 3.1 Redistributable (v2)
- Microsoft Visual C++ 2005 Redistributable Package
- .NET Framework 2.0
- BITS 2.0 Update KB842773

Let's walk through an example of performing a P2V conversion to illustrate how it works. Figure 5-50 is a screen shot taken on a physical server named NYC-SRV-B in the contoso.com domain showing the File Services role installed. Our goal here is to use VMM 2008 R2 to perform an online conversion of this physical server into a virtual machine.




FIGURE 5-50 Screen shot of the physical server that will be converted to a virtual machine.

To start the conversion process, click Convert Physical Server in the Actions pane of VMM to launch the Convert Physical Server (P2V) Wizard. On the Select Source page of this wizard, type or browse to specify the name or IP address of the physical server you want to convert to a virtual machine and specify credentials that have local administrator privileges on the physical server. (See Figure 5-51.)




FIGURE 5-51 Specify the name of the physical server and local administrator credentials on the server.

The next wizard page tells us that the virtual machine you are going to create will have the same name as your physical server. (See Figure 5-52.) You can change this identity if desired, but let's leave it configured as it is.




FIGURE 5-52 Specify the name of the virtual machine to be created.

On the next wizard page, click Scan System to install a VMM agent on the physical server to gather information about its hardware. (See Figure 5-53.)




FIGURE 5-53 Click Scan System to install a VMM agent on the physical server and scan its hardware.

When the scan is complete, the results are displayed as shown in Figure 5-54.




FIGURE 5-54 Results of hardware scan of the physical server.

The next wizard page displays the disk volumes on the physical server that you will be duplicating on the virtual machine. The default VHD type for these volumes is *dynamic*, but you can change this to *fixed* if desired. Clicking Conversion Options displays options for choosing between an online or offline conversion. For this walkthrough, select the Online Conversion option. (See Figure 5-55.)




FIGURE 5-55 Specifying disk options and conversion method.

The next wizard page lets you specify the number of processors and amount of RAM to be assigned to the virtual machine being created. By default, the values displayed here are those of the physical server's hardware. (See Figure 5-56.)




FIGURE 5-56 Specify the number of processors and amount or RAM to be used by the virtual machine.

On the next wizard page, you select a host for the virtual machine you are going to create. In Figure 5-57, you first select the New York hosts group, which displays HVR2 as the only available host in this host group.




FIGURE 5-57 Choose a host to place the new virtual machine on.

The next figure lets you select a storage location on the host for the virtual machine files that will be created. (See Figure 5-58.)




FIGURE 5-58 Specify a location for storing the virtual machine files.

On the next wizard page, you can select a virtual network for the new virtual machine. (See Figure 5-59.)




FIGURE 5-59 Specify a virtual network for the virtual machine.

The next wizard page lets you configure the automatic start actions for the virtual machine when the physical server starts or stops. (See Figure 5-60.)




FIGURE 5-60 Configuring automatic start options.

The next-to-last wizard page does a final check to determine whether there are any issues that might block the conversion process from happening or cause the process to fail. (See Figure 5-61.) If any issues are detected at this point, you can try resolving them on the physical server and then click Check Again to verify.




FIGURE 5-61 Verifying that there will be no issues with the planned conversion.

The final screen of the wizard displays summary information of how you have configured the conversion process. (See Figure 5-62.)




FIGURE 5-62 Summary page of the Convert Physical Machine (P2V) Wizard.

Now click Create to begin the P2V conversion process. The Jobs window opens at this point and displays the progress of the conversion. (See Figure 5-63.)




FIGURE 5-63 P2V conversion is underway.

The conversion process might take some time to complete. After it is finished, the virtual machine will be displayed in the Administrator Console and you can start it up to get it running. (See Figure 5-64.)




FIGURE 5-64 The new virtual machine in the VMM Administrator Console.

Double-clicking on the thumbnail image of the virtual machine opens it in the Virtual Machine Viewer. As you can see by comparing Figure 5-65 with the earlier screen shot captured on the physical server (shown in Figure 5-50 earlier), the P2V conversion process has indeed worked as expected.




FIGURE 5-65 The new virtual machine in the Virtual Machine Viewer.

Performing V2V Conversions

You can use VMM 2008 R2 to convert a VMware ESX Server–based virtual machine to a Hyper-V or Virtual Server–based virtual machine. This process is called virtual-to-virtual (V2V) conversion and is different from migration, which is the process of moving a virtual machine from one host to another and was discussed earlier in this chapter. Table 5-9 summarizes the difference between virtual machine migration and V2V conversion.

TABLE 5-9 Options for Virtual Machine Migration and V2V Conversion

From	To	Method
Hyper-V	Hyper-V	Migration
Virtual Server	Virtual Server	Migration
Virtual Server	Hyper-V	Migration
VMWare ESX Server	VMWare ESX Server	Migration
VMWare ESX Server	Hyper-V	V2V
VMWare ESX Server	Virtual Server	V2V

The following guest operating systems are supported for virtual-to-virtual (V2V) conversion:

- Windows Server 2008 R2
- Windows Server 2008 (32-bit)
- Windows Server 2008 (64-bit)
- Windows Server 2003 (32-bit)
- Windows Server 2003 (64-bit)
- Windows 2000 Server SP4 or later (Offline P2V only)
- Windows 2000 Advanced Server SP4 or later (Offline P2V only)
- Windows XP Professional (32-bit) SP2 or later
- Windows XP Professional (64-bit) SP2 or later
- Windows Vista SP 1 (32-bit)
- Windows Vista SP1 or later (64-bit)
- Windows 7 (32-bit)
- Windows 7 (64-bit)

V2V conversion can be performed by using the Convert Virtual Machine Wizard, which converts the .vmdk files to .vhd files and makes the guest operating system on the virtual machine compatible with Microsoft virtualization technologies. The virtual machine created by the wizard matches the VMware virtual machine properties, including its name, description, memory, disk-to-bus assignment, CD and DVD settings, network adapter settings, and other parameters.

Direct from the Source: Using P2V to Convert a Running VM from Another Hypervisor

One task that administrators ask for is the ability to do a Virtual-to-Virtual (V2V) conversion of a running VMware virtual machine to a Hyper-V server. The current V2V process in VMM is an offline process, but you can do an easy online conversion of a VMware VM. We have to remember that virtual machines are machines first and virtual second, and that VMs present themselves to users and administrators just like physical machines. Thus, administrators can do a Physical-to-Virtual conversion of the VMware VM. Even though it is a virtual machine, the VM can look like a physical machine to VMM, and if the operating system is a supported P2V operating system, VMM can do a live migration of that VMware VM to a Hyper-V VM.

—Edwin Yuen, Senior Product Manager, Integrated Virtualization Strategy

Configuring User Roles

VMM 2008 R2 employs user roles to determine the level of access users and groups can have to different kinds of resources. You can create three types of user roles using the Administrator Console:

- **Administrator role** Able to perform all actions using the Administrator Console, including creating new Delegated Administrator and Self-Service User roles and adding members to these roles. The default membership in the Administrator role includes members of the local Administrator group on the VMM Server, the domain account of the user who installed the VMM Server, and the computer account of the VMM Server.
- **Delegated Administrator role** Able to perform most actions using the Administrator Console as restricted by the scope defined for the role. The scope defines which host groups and Library Servers the Delegated Administrator is permitted to manage. Members of this role can also create new Delegated Administrator and Self-Service User roles and add members to these roles.
- **Self-Service User role** Able to use the Self-Service Portal to perform tasks on virtual machines as defined by the scope and permissions for the role. Members of this role cannot create any new user roles.

You can create user roles by using the New User Role Wizard as demonstrated in the following sections.

Creating a Delegated Administrator Role

If you belong to either the Administrator or Delegated Administrator role, you can use the Create User Role Wizard to create a new Delegated Administrator role. To do this, begin by selecting User Roles under Administration in the Administration pane. (See Figure 5-66.)

Now click New User Role under User Role in the Actions pane. This launches the Create User Role Wizard. Type a name for the new role you will create and an optional description. Then select Delegated Administrator as the kind of role you will create. (See Figure 5-67.)




FIGURE 5-66 Managing user roles.




FIGURE 5-67 Creating a new Delegated Administrator role.

On the next page of the wizard, browse Active Directory to select the user accounts you want to add as members of your new role. (See Figure 5-68.)




FIGURE 5-68 Adding members to the role.

On the next wizard page, specify the scope of your new role. (See Figure 5-69.)




FIGURE 5-69 Specify the scope for the Delegated Administrator role.

The final page of the wizard presents a summary of the selections you have made on the previous pages. After you're satisfied you are configuring the new role appropriately, click Create to create the new Delegated Administrator role. See Figure 5-70 shows the result of creating the new Delegated Administrator role.



Tip You can click View Script on the final wizard page to view the Windows PowerShell script that will be used to perform the actions you specified in the wizard. This can be useful for automation purposes. For example, you could copy the script from this role and modify it to create multiple roles in one batch operation.




FIGURE 5-70 Summary page of Create User Role Wizard.

Creating a Self-Service User Role

If you belong to either the Administrator or Delegated Administrator role, you can also use the Create User Role Wizard to create a new Self-Service User role. To do this, begin again by selecting User Roles under Administration in the Administration pane as shown in Figure 5-70. Click New User Role under User Role in the Actions pane to launch the Create User Role Wizard. Type a name for the new role you will create and an optional description. Now select Self-Service User as the kind of role you will create. (See Figure 5-71.)




FIGURE 5-71 Creating a new Self-Service User role.

Browse Active Directory to select the user accounts you want to add as members of your new role, and then add these accounts to the role. Figure 5-72 shows user Jacky Chen assigned to the Self-Service User role.




FIGURE 5-72 Adding members to the role.

Next, specify the scope of your new role. (See Figure 5-73.)




FIGURE 5-73 Specify the scope for the Self-Service User role.

Next you specify the permissions that members of this role will have for managing virtual machines running on hosts that belong to host groups that are within the scope of the role. You can either specify that role members can perform any action on the virtual machines, or you can select specific actions to allow them to perform. (See Figure 5-74.)




FIGURE 5-74 Specify the virtual machine permissions for the Self-Service User role.

On the next page of the wizard, you specify whether members of this role are allowed to create new virtual machines on hosts that belong to host groups that are within the scope of the role. You can also specify which virtual machine templates the role members can use when creating their new virtual machines. (See Figure 5-75.)




FIGURE 5-75 Specify the templates that role members can use for creating new virtual machines.

On this wizard page, you can also set a quota for deployed virtual machines for members of the Self-Service User role. A virtual machine quota in a Self-Service User role is a limit to the number of virtual machines that members of the role can deploy. Quota points are assigned to the templates that self-service users use to create their virtual machines, and they apply only to virtual machines on a virtual machine host. If a self-service user is allowed to store virtual machines, the quota does not apply to virtual machines stored in the library. When the self-service user's quota is reached, the user cannot create any new virtual machines until an existing virtual machine is removed or stored in the library.

On the next page of the wizard, you can specify whether the user is allowed to store her virtual machines in the library. (See Figure 5-76.)




FIGURE 5-76 Specify whether role members can store their virtual machines in the library.

The final page of the Create User Role Wizard displays a summary of the selections you have made. Figure 5-77 shows the result of creating the new Self-Service User role.




FIGURE 5-77 Summary page of the Create User Role Wizard.

Using the Self-Service Portal

The VMM Self-Service Portal provides Web-based functionality that allows users to create, manage, and operate their own virtual machines independently within a controlled environment. Setting up the Self-Service Portal involves the following steps:

1. Installing the VMM Self-Service Portal on a Web server running IIS. You can configure the Self-Service Portal on your VMM Server or on a dedicated Web server in your domain.
2. Creating a host group to use for virtual machine self-service, and moving hosts into the host group. This configuration is useful for limiting self-service users to a specific group of managed hosts.
3. Creating virtual machine templates that your self-service users can use to create their own virtual machines.
4. Creating Self-Service User roles to grant users permissions to administer their virtual machines and to allow them to create their virtual machines from templates.
5. Specifying the e-mail address of the administrator who will support the self-service users.



Tip If self-service users will use virtual machines that you have created, create the virtual machines, configure them for virtual machine self-service, and deploy the virtual machines on a host in the host group that you use for self-service.

After the Self-Service Portal has been configured, users who are members of a Self-Service User role can use a Web browser such as Internet Explorer to open the home page of the portal and log on using their domain credentials. (See Figure 5-78.)




FIGURE 5-78 Logging on to the Self-Service Portal.

After members of the Self-Service User role log on, they can manage their virtual machines using the portal by clicking the Computers tab on the Portal.aspx page. By clicking List View, you can view the name, status, owner, assigned memory, disk space used, date deployed, and quota points for each virtual machine you have access to. (See Figure 5-79.) And by selecting a virtual machine, you can perform the actions you have permissions to perform as determined by the Self-Service Portal role to which you belong.

If you select a running virtual machine in List View and click Connect To VM in the Actions pane on the right, a separate Internet Explorer window opens displaying a remote control session to the virtual machine. (See Figure 5-80.) When you do this for the first time on the Self-Service Portal, you are prompted to install the ActiveX control that makes this connection possible.




FIGURE 5-79 Viewing a list of the user's virtual machines.




FIGURE 5-80 Connecting to a running virtual machine using a remote control session.

If you return to the Self-Service Portal page and click Thumbnail View, you can view a thumbnail image of your running virtual machines. (See Figure 5-81.) And by selecting one of the thumbnails, you can perform the actions you have permissions to perform as determined by the role. If for some reason a thumbnail image isn't displayed for a particular running virtual machine, click the green circular arrow icon to reconnect to the virtual machine.




FIGURE 5-81 Viewing thumbnails of the user's virtual machines.

Finally, by clicking New Computer in the Create pane on the right side of the portal page, you can create a new virtual machine from a template based on the way the Self-Service User role was configured. (See Figure 5-82.)




FIGURE 5-82 Creating a new virtual machine.

Microsoft System Center Solutions

System Center Virtual Machine Manager 2008 R2, which we've dived into in this chapter, is only one part of the Microsoft System Center family of products. System Center solutions can help you manage an organization's physical and virtual infrastructure from the desktop to the datacenter. System Center solutions play a central role in Microsoft's vision for enabling IT organizations to benefit from self-managing, dynamic systems.

For the dynamic datacenter, System Center solutions enable

- Configuration management
- Server compliance
- End-to-end monitoring
- Data protection and recovery

For the dynamic desktop, System Center solutions facilitate

- Adaptive application delivery
- Simplified Windows 7 deployment
- Endpoint security management
- Configuration compliance
- Client infrastructure monitoring
- Remote PC diagnostics and repair

Dynamic IT, developed according to Microsoft's strategy for providing critical technologies that enable IT to become more strategic, represents Microsoft's vision of what an agile business looks like—where IT works closely with a business to meet the demands of a rapidly changing and adaptable environment.

System Center Server Management Suite Enterprise

System Center Server Management Suite Enterprise encompasses four core System Center products to provide a complete set of server management capabilities for managing applications and platforms that span both physical and virtual environments. System Center Server Management Suite Enterprise also provides the right to manage an unlimited number of operating system environments on a single virtual machine host. The four core System Center products within this suite include

- **System Center Configuration Manager 2007 R2** A solution that comprehensively assesses, deploys, and updates servers, client computers, and devices across physical, virtual, distributed, and mobile environments. System Center Configuration Manager is optimized for Windows and is extensible so that it can control virtually any IT system.
- **System Center Operations Manager 2007 R2** An end-to-end service-management product that works seamlessly with Microsoft software and applications, helping organizations increase efficiency while enabling greater control of their physical and virtual infrastructure.
- **System Center Data Protection Manager 2007 SP1** A solution for Windows backup and recovery that delivers continuous data protection for Microsoft application and file servers using seamlessly integrated disk and tape media.
- **System Center Virtual Machine Manager 2008 R2** A product that allows unified management of physical and virtual machines, consolidation of underused physical servers, and rapid provisioning of new virtual machines.

System Center Essentials

System Center Essentials 2007 SP1 is specifically designed for midsize businesses with up to 500 client computers and 30 servers. It provides a unified management solution that enables midsize organizations to proactively manage their IT environment with increased efficiency. System Center Essentials provides monitoring and alert resolution for servers, clients, applications, hardware, and network devices; software distribution; update management; and software and hardware inventory.

Other System Center Products

Other System Center products include the following:

- **System Center Capacity Planner 2007** A predeployment capacity-planning and postdeployment change-analysis solution for Microsoft server applications, such as Microsoft Exchange Server 2007, Windows SharePoint Services 3.0, and Microsoft Office SharePoint Server 2007. System Center Capacity Planner provides tools and guidance to deploy these servers efficiently, while also helping you plan for the future by enabling “what-if” analyses.
- **System Center Mobile Device Manager 2008 SP1** An end-to-end product for single-point access of line-of-business (LOB) applications and corporate data on devices running Windows Mobile 6.1. System Center Mobile Device Manager provides secure access to sensitive corporate data on such devices in a seamless manner.
- **System Center Service Manager** A new Microsoft System Center product designed to meet the needs of the modern IT help desk by providing capabilities for incident, problem, asset, and change management.

For more information about System Center solutions, their benefits, how to purchase or try them out, and for partner and technical information, see the Microsoft System Center portal at <http://www.microsoft.com/systemcenter/en/us/default.aspx>.

Benefits of System Center for Virtualization

Implementing System Center solutions can have numerous direct benefits for your organization's virtualization infrastructure. For example,

- System Center Configuration Manager 2007 R2 facilitates operating-system and application configuration management, patch management and deployment, and software upgrades for physical and virtual machines from the datacenter to the desktop.
- System Center Operations Manager 2007 R2 provides end-to-end service management, server and application health monitoring and management, and performance reporting and analysis for physical and virtual machines.
- System Center Virtual Machine Manager 2008 R2 allows for management of virtual machines, server consolidation and resource utilization optimization, and Physical-to-Virtual (P2V) and Virtual-to-Virtual (V2V) conversions for managing your virtualized infrastructure and making more efficient use of your physical hardware resources.
- System Center Data Protection Manager 2007 SP1 provides live, host-level virtual machine backup, in-guest consistency, and rapid recovery for Hyper-V hosts to eliminate downtime for both your physical infrastructure and virtual infrastructure.

Direct from the Source: Choosing a DPM 2007 SP1 Backup Solution

With Microsoft System Center Data Protection Manager (DPM) 2007 SP1, you can use disk-based storage, tape-based storage, or both.

Tape-Based Backup and Archive

Traditional data archiving methods have relied on backing up data to tape media. This method is referred to as disk-to-tape (D2T). Tape is a popular medium for offsite storage, because magnetic tape and similar storage media offer an inexpensive and portable form of data protection that is particularly useful for long-term storage. A thorough disaster recovery plan based on tape media includes offsite storage of critical information to allow data recovery in situations where a facility is damaged or destroyed.

Using DPM, you can back up data from a server directly to tape. Data can be backed up to tape as frequently as daily for short-term protection, and it can be maintained as long as 99 years for long-term protection.

For long-term, tape-based protection, DPM backs up data from the replica in the storage pool to tape so that there is no impact on the protected server. If a file was open when the replica was synchronized last, the backup of that file from the replica will be in a crash-consistent state. A crash-consistent state of the file will contain all file data that was preserved to disk at the time of last synchronization. This applies only to file system backups. Application backups will always be consistent with the application state.

Disk-Based Protection and Recovery

With the advent of large and relatively inexpensive disk subsystems, organizations have begun to implement disk-based backup solutions, referred to as disk-to-disk (D2D). This method of backup is used to store data from one computer on the hard disk of another computer. One advantage of disk-based data protection is the potential to save time. Disk-based data protection eliminates the need to locate a specific tape required for a recovery job, load the tape, and position the tape to the correct starting point. The potential time savings and ease of use inherent in disk backup solutions encourages sending incremental data more frequently, which reduces the impact to the server being protected and on network resources. Additional benefits of this approach include

- **Increased reliability** The reliability of data recovery with disk-based data protection is also better than that of tape-based systems. Disk drives typically have much greater mean time between failure (MTBF) ratings than tape systems.
- **Faster recovery** Recovery of data from disk is quicker and easier than recovery from tape. Recovering data from disk is a simple matter of browsing through previous versions of the data on the DPM server and copying

selected versions directly to the protected file server. In comparison, a typical file recovery from tape takes hours, which can be costly, and administrators in a medium-size datacenter can usually expect to perform 10 to 20 or more of these recoveries each month.

Using DPM and disk-based data protection, data can be synchronized as frequently as every 15 minutes and maintained as long as 64 days with 1 recovery point per day.

Combining Disk and Tape Backup Methods

You can also back up data from a disk-based replica. Often referred to as D2D2T or “disk-to-disk-to-tape,” this method combines the ease of use inherent to disk-based backup with the long-term and offsite storage capabilities of tape media by replicating data to disk for subsequent backup to tape. The primary advantage to this method is that the tape backup operation can occur at any time with no impact on the server that is being protected.

DPM allows administrators to schedule disk-based replication and backup to tape, providing the flexibility to create focused, detailed backup strategies that result in efficient and economic data protection. When you need to restore a single file or an entire server, recovery is typically fast and simple, because once you identify the data to be restored, DPM locates the data and retrieves it from the disk-based replica. Should a disk replica fail, tape backups can be used to restore the replica and resume protection.

Determining Which Storage Method to Use

To determine which storage method to use, you must consider the relative importance of your organization’s protection requirements.

- **How much data your organization can afford to lose** Realistically, not all data is equally valuable. Organizations must weigh the impact of loss against the costs of protection.
- **How quickly recovered data must be made available** Recovery of data that is critical to ongoing operations will typically be more urgent than routine data. On the other hand, organizations should identify servers providing essential services during working hours that must not be disrupted by recovery operations.
- **How long your organization must maintain data** Long-term storage might be necessary for business operations, depending on the type and contents of the data. An organization might also be subject to legal requirements for data retention.
- **How much your organization can spend on data protection** When considering how much to invest in data protection, organizations must include not only the cost of hardware and media, but also the personnel costs for administration, management, and support.

—CSS Global Technical Readiness Team (GTR)

Additional Resources

The resources listed in this section were either referred to in the chapter or provide additional information concerning concepts and products discussed in this chapter.

General

The main SCVMM portal on microsoft.com can be found at <http://www.microsoft.com/systemcenter/virtualmachinemanager/en-us/default.aspx>. You should begin there if you are looking for general information about VMM 2008 R2.

Administering VMM

For detailed and authoritative technical guidance about how to deploy, configure, use, maintain, and troubleshoot VMM 2008 R2, see the SCVMM TechCenter on Microsoft TechNet at <http://technet.microsoft.com/en-us/scvmm/default.aspx>.

To dive in quickly into the product, check out the following two sections of the TechNet Library:

- “System Center Virtual Machine Manager 2008 and Virtual Machine Manager 2008 R2” found at <http://technet.microsoft.com/en-us/library/cc917964.aspx>.
- “Online Help for Virtual Machine Manager” found at <http://technet.microsoft.com/en-ca/library/bb740738.aspx>.

System Center Blog

Be sure to subscribe to the RSS feed of the System Center Virtual Machine Manager Team Blog at <http://blogs.technet.com/scvmm/default.aspx> if you want to keep up with all the latest news, tips and best practices concerning VMM.

VMM Forums on TechNet

To obtain help with your questions and problems concerning VMM, and to help others, use the VMM forums on Microsoft TechNet at <http://social.technet.microsoft.com/forums/en-US/category/virtualmachinemanager>.

Chapter 6

Cloud Computing

The previous chapters of this book have dealt with specific virtualization products and technologies that are already available from Microsoft. This final chapter looks to the future and describes how these various products and technologies make possible a new type of virtualization called *cloud computing*. This chapter examines the benefits of cloud computing, possible usage scenarios, the availability of Microsoft's offerings in this area, and how the Microsoft cloud-computing platform works. Because Microsoft's cloud-computing platform is still in early beta at the time of this writing, the technical details described in this chapter are subject to change.

What Is Cloud Computing?

Cloud computing is an emerging solution whereby virtualized computing resources running "in the cloud" can be provisioned and delivered as services to those who need them, an approach often called *IT as a service*. These computing resources can include servers, storage, and networking resources running in virtualized environments using Microsoft Hyper-V, and the virtualized environment itself resides in a datacenter that can be either privately or publicly owned.

A cloud-computing solution basically includes the following elements:

- A datacenter of Hyper-V servers and storage resources that forms the underlying infrastructure of the cloud
- A scalable pool of virtual servers, storage, and networking resources *hosted* or running in the cloud
- Tools for managing the datacenter servers and storage resources as a single fabric
- Software for packaging and delivering virtualized computing resources as services
- Automation software for authorizing the provisioning of resources to users upon their request
- A reliable, high-bandwidth network for delivering these resources to users

In a cloud-computing scenario, the user isn't concerned about where the virtual servers, storage, or networking resource are actually located; all he cares about is that he can quickly access a resource when he needs it. For example, when he needs a new virtual server for his

department, he doesn't care whether the virtual server is running on a Hyper-V server down the hall, in a nearby building, or in a datacenter halfway across the continent. He also doesn't care whether the Hyper-V server is owned and maintained by the company he works for, a partner organization, or a public hosting provider. All he wants is to be able to submit a request for a new virtual server and then gain access to this server as quickly as possible.

Private vs. Public Cloud

Microsoft's vision for cloud computing centers around two types of solutions:

- **Private cloud** This approach uses your organization's existing virtualization infrastructure, and it provides your organization with complete control over what virtualized computing resources can be made available to your users and how these resources can be provisioned and delivered. In a private-cloud scenario, you use software and tools provided by Microsoft to pool the resources of your existing virtualized infrastructure by wrapping up these resources and offering them as a service to your users. A private cloud is therefore not a separate product offering from Microsoft, but rather it's a computing infrastructure deployment option that involves virtualized infrastructure with fabric-integrated automated management that results in an internal, service-oriented environment for an enterprise.
- **Public cloud** This approach uses a hosting provider's infrastructure to provide virtualized computing resources on demand to organizations that need them. The hosting provider in this case is located outside your corporate firewall and can either be Microsoft itself (at a Microsoft datacenter) or a third-party hosting provider using software and tools available from Microsoft. A public cloud is therefore not a separate product offering from Microsoft, but rather it's a deployment option for enterprises whereby virtualized infrastructure services are provided by a hosting partner.

Large organizations with an existing virtualization infrastructure might choose to implement a private-cloud solution because it provides them with complete control over the solution, making it more secure because the solution is completely implemented behind the organization's corporate firewall. Small and mid-sized organizations without the infrastructure or in-house expertise to deploy their own private-cloud solution can use the public-cloud approach through a pay-as-you-go arrangement with a hosting provider that offers the cloud computing services they require. Some organizations might even choose to follow a hybrid cloud approach whereby they combine their own in-house, private-cloud solution with the offerings of a public-cloud service provider. The ability for organizations to blend together custom solutions based on a combination of private and public cloud offerings from Microsoft is called the *cloud continuum*.

Examining the Benefits of Cloud Computing

The fundamental benefit of the cloud-computing approach is business agility. In fact, cloud computing is the natural and next evolutionary stage of Microsoft's Dynamic IT initiative, in which the datacenter becomes capable of dynamically responding to the changing needs of businesses. Cloud computing provides this benefit by making your IT infrastructure more able to quickly respond to rapid changes in your business environment. For example, in a traditional enterprise computing environment, a department that needs a new server might need to wait weeks or even months for the provisioning process to work its way from approval to supplier to delivery and deployment. Even in a virtualized IT infrastructure, it might still take several days or even weeks from requesting a new virtual server to having this server configured, tested, and provisioned. With a private-cloud infrastructure, however, the approval and provisioning of virtual servers can be predefined using policies so that the entire process from submitting a request to gaining access to the virtual server can be almost immediate. And with a public-cloud solution, the provisioning process can be just as rapid, and it can happen without any changes being made to your organization's existing infrastructure.

Benefits of Using a Private Cloud vs. a Public Cloud

Specific benefits of Microsoft's private-cloud offering include the ability to

- Quickly and cost-effectively scale up as needed
- Deliver scalable applications and workloads to users who need them
- Manage the datacenter fabric as a single pool of virtual computing resources
- Focus on the management of the datacenter service and its dependencies
- Enable federated services across the full cloud continuum
- Eliminate unnecessary hardware purchases by ensuring that existing hardware resources are used more effectively
- Ensure trustworthiness by providing you with complete control of your cloud

The benefits of Microsoft's public-cloud solution includes the ability to

- Lower business costs by renting virtualized computing resources "in the cloud" on a pay-as-you-go basis
- Lower IT support costs by always having the latest versions of software available without the need for internal IT support
- Mitigate risk by transferring responsibility for data security to the hosting company
- Increase business efficiency by delivering computing resources over the Internet to anywhere a user might need it

Increasing Use of IT Resources

The bottom line is that cloud computing increases the agility of your business, reduces your IT operational costs, and ensures that your IT infrastructure is being efficiently used. This last point concerning efficient use of IT resources is worth a closer look. (See Figure 6-1.)




FIGURE 6-1 How the evolution of computing infrastructure leads to greater efficiency through virtualization.

The traditional (nonvirtualized) IT infrastructure is usually well-known, stable, and secure. However, servers in such a traditional environment are typically underused and, on average, might be running at a utilization of under 15 percent. This means that although your organization might benefit from having a secure and stable traditional IT infrastructure, you are not using your IT resources efficiently.

Adding virtualization to the traditional IT infrastructure using Hyper-V maintains the existing benefits of stability and security while also increasing the average utilization of your servers to 50 percent or more. In addition, management tools such as Microsoft System Center can further reduce costs by making it easier to deploy and maintain virtual servers than physical ones.

Implementing Microsoft's private-cloud solution on top of your existing virtualized infrastructure allows you to provision IT resources to departments and users as a service. Although this doesn't necessarily increase average server utilization, it does decrease management costs significantly, making your IT infrastructure more efficient.

Finally, using the public-cloud approach adds further efficiencies to your business by allowing you to quickly and cost-effectively increase the capacity of your IT infrastructure on demand, anywhere in the world it might be needed.

Examining Cloud-Computing Usage Scenarios

Cloud computing offers capacity on demand so that consumers of this capacity can scale up or down as needed. Some examples of scenarios where this can be useful include the following:

- Susan is a developer who needs several virtual servers in order to emulate the production environment to which she will later be deploying the application she is developing. Susan uses her organization's private cloud to procure the servers she needs within hours of her request instead of having to wait weeks for new hardware to arrive.
- Bob is a branch office administrator who needs a new server and additional storage to meet the needs of the new employees who have just arrived at his remote site. He uses the public-cloud services provided by his organization's hosting provider to deploy these servers by the end of the day, ensuring that the new employees can begin their work immediately instead of having to wait around for resources to become available for them.
- George works for a system integrator as a developer of Web-based applications for a vertical industry segment. Instead of developing a custom solution for each customer, he uses the Microsoft Windows Azure services platform to develop and host these applications on datacenters owned and operated by Microsoft. For more information about this service, see the section titled "Windows Azure" later in this chapter. This speeds development time and allows customers to access these applications over the Internet from anywhere in the world.

Understanding Microsoft's Cloud-Computing Platform

To understand how Microsoft's cloud-computing platform works, you first need to understand its various components and how these components can be tied together and used to deliver virtualized computing resources to those who need them. There are several ways of looking at this matter, and these are explored in the following sections.

Understanding Different Cloud Services

To begin with, the concept of "IT as a service" can be understood in several ways, depending on which types of cloud services are being provided. Specifically, cloud computing can deliver

- **Software as a Service** SaaS refers to business applications being hosted in and delivered from a private or public-cloud infrastructure. Users might use installed desktop applications or Web browsers for interacting with these hosted applications.

SaaS providers can even offer headless (without a UI) Web services that enable enterprises to integrate data and business processes with SaaS applications.

- **Platform as a Service** PaaS refers to providing operating system and application services to users from a cloud in order to build applications that are hosted in and run in the cloud. To enable this, PaaS provides system resource management functions for allocating memory space, scheduling processing time, and ensuring system and application integrity within a multi-user environment.
- **Infrastructure as a Service** IaaS refers to a providing processing and storage resources to users as dynamically scalable, virtualized services. IaaS reduces an organization's need to invest in low-level hardware such as servers and storage devices.

In Microsoft's cloud-computing approach, each of these different types of cloud services relies on one or more Microsoft products, technologies, or platforms for hosting them. Table 6-1 lists some of the different products and platforms Microsoft offers that enable the private and public-cloud versions of SaaS, PaaS, and IaaS.

TABLE 6-1 How Microsoft Products and Platforms Enable SaaS, PaaS, and IaaS for Private and Public Cloud Solutions

Type of Cloud Service	Microsoft Products and Platforms that Enable These Services to Be Delivered	
	Through a private cloud	Through a public cloud
Software as a Service (SaaS)	Microsoft SharePoint Services Microsoft Exchange Microsoft Dynamics	Microsoft SharePoint Services Microsoft Online Services Microsoft Office Live
Platform as a Service (PaaS)	Microsoft SQL Services Microsoft .NET	SQL Services Microsoft .NET Services Live Services Windows Azure
Infrastructure as a Service (IaaS)	Windows Server with Active Directory, Hyper-V, and Failover Clustering Microsoft System Center Microsoft Dynamic Data Center Toolkit for Enterprises	Windows Server with Active Directory, Hyper-V, and Failover Clustering Microsoft System Center Microsoft Dynamic Data Center Toolkit for Hosters

Some of these platforms, such as Windows Server with Hyper-V and Microsoft System Center, have been discussed in detail in earlier chapters of this book. Other products should be generally familiar to IT professionals who work with Microsoft products and solutions. A new and key component of any Microsoft cloud-computing solution, however, is the last item in each column—namely, the Microsoft Dynamic Data Center Toolkit, which will be described shortly.

Implementing Cloud Services

The categories of SaaS, PaaS, and IaaS are really just ways of thinking about the different components that make up a Microsoft cloud solution. Implementing a real-world cloud solution involves using at least one product or platform from each of these three categories of cloud services and layering them together to make your cloud-computing infrastructure. For example, Figure 6-2 shows what a full private-cloud infrastructure might look like when layered upon your organization's existing virtualization-capable IT infrastructure.




FIGURE 6-2 Example of a private-cloud infrastructure

Layered on top of your server and storage hardware is the IaaS layer, which includes Microsoft System Center, Windows Server with Hyper-V, and Microsoft's Dynamic Data Center Toolkit for Enterprises. These components are essential to the functioning of any Microsoft private-cloud solution because they provide the necessary "plumbing" for making the provisioning of virtualized computing resources work, a platform for hosting these resources, and tools for managing the virtualized infrastructure and resources. Every private-cloud infrastructure must include these three components.

On top of IaaS comes the PaaS layer, which includes server applications (SQL Server) and technologies (the .NET Framework) for building custom solutions that support the provisioning and delivery of IT services. Finally, on top of PaaS comes the SaaS layer, which lets you provision workspaces, business applications, and collaboration services to users who need them.

Because the IaaS layer is the key to making the cloud-computing infrastructure possible, it's worth examining this layer in more detail. (See Figure 6-3.)




FIGURE 6-3 Examining the IaaS layer of the private-cloud infrastructure.

The IaaS layer of the cloud-computing infrastructure is based on three Microsoft platforms:

- **Windows Server 2008 R2** IaaS relies on a number of key Windows Server roles and features:
 - ❑ Active Directory Domain Services (AD DS) and Domain Name System (DNS) for enabling provisioning of virtualized computing resources to users who need them
 - ❑ Hyper-V for hosting virtual servers and managing virtual networks and storage
 - ❑ Failover Clustering with Live Migration for ensuring that virtual servers are highly available
- **System Center** IaaS relies on System Center Virtual Machine Manager for managing the fabric of your virtualized infrastructure, and on System Center Operations Manager for delivering cloud-computing services to users
- **Dynamic Data Center Toolkit for Enterprises** Enables you to deliver virtualized computing resources to users on demand as consumable services

Understanding the Dynamic Data Center Toolkit

While most of the components of Microsoft's cloud computing platform are existing products or services offered by Microsoft, the glue that ties them all together and makes cloud computing a reality is the Dynamic Data Center Toolkit. This toolkit is a collection of documentation, tools and code you can use to rapidly build and launch a scalable virtualized infrastructure that can deliver managed IT services.

The Dynamic Data Center Toolkit comes in two different forms:

- **Dynamic Data Center Toolkit for Hosters** This version can be used by hosting service providers to enable them to build a public cloud they can use to provide virtualized IT infrastructure and managed services to their customers. The Dynamic Data Center Toolkit for Hosters includes
 - ❑ Step-by-step instructions
 - ❑ Best practices
 - ❑ Sample code
 - ❑ Demos you can view
 - ❑ Marketing collateral
 - ❑ White papers and other documentation
- **Dynamic Data Center Toolkit for Enterprises** This version can be used by businesses to build a private cloud in their datacenters that they can use to dynamically pool, allocate, and manage virtualized computing resources within their organization. The Dynamic Data Center Toolkit for Enterprises includes
 - ❑ Architectural roadmap
 - ❑ Deployment guidance
 - ❑ Best practices
 - ❑ Familiar tools
 - ❑ Lightweight provisioning engine
 - ❑ Prebuilt user and admin portals
 - ❑ White papers and other documentation

The basic structure of these toolkits aligns closely with the Infrastructure Planning and Design (IPD) series of guides that provide architectural guidance for planning and implementing Microsoft infrastructure products. These IPD guides provide concise planning guidance that helps clarify and streamline the design process for deploying Microsoft infrastructure technologies and walks you through the critical decisions you will need to address concerning each available options. The guides also help you validate your design decisions to ensure that the solution meets the requirements of your business and infrastructure stakeholders. Each IPD guide in the series addresses a particular infrastructure technology or scenario. For more information on available IPD guides and to download them, go to <http://technet.microsoft.com/en-us/solutionaccelerators/ee382254.aspx>.

Comparing the Toolkits

The Dynamic Data Center Toolkit for Hosters is not an out-of-the-box solution you can just install and use within your environment. This is because hosting service providers typically have complex, heterogeneous environments that often include provisioning, control, and management tools developed in-house by the hoster for its specific needs. However, the Toolkit for Hosters does include a sample portal that can be used to provide the hoster's customers with an integrated control and view of the services provided by the hoster.

By contrast, the Dynamic Data Center Toolkit for Enterprises is closer to the ideal of a turnkey cloud-computing solution because it assumes an existing Microsoft-centric datacenter infrastructure based on Hyper-V. The Toolkit for Enterprises includes a lightweight provisioning engine with two prebuilt portals (shown in Figure 6-4), including

- An admin portal for allocating virtualized computing resources to users who request them. Using this portal, a datacenter administrator can provision infrastructure once in response to a user's request and then delegate the further operation of the infrastructure to the user.
- A self-service portal for users that provides them with a list of actions they are authorized to perform, such as creating new virtual servers. The user is typically the IT person or team of a smaller business unit within the organization, such as a departmental administrator.




FIGURE 6-4 Microsoft's private-cloud solution includes two prebuilt portals.

The following is an example of how these two portals can be used:

1. Al, the datacenter administrator, has implemented a private-cloud solution for his organization using the Dynamic Data Center Toolkit for Enterprises. He has done this by “laying down the plumbing” of implementing the toolkit on top of the existing virtualization-capable infrastructure in the datacenter. This enables Al to use the admin portal to manage the procurement of virtualized computing resources and for managing and monitoring the physical and logical fabric of the datacenter.
2. Bob, a departmental administrator, needs a focused, service-level agreement (SLA)-driven solution for meeting a business need for his department. He determines that this solution will require additional servers, so he contacts Al, who tells Bob to sign up using the self-service portal for the private-cloud infrastructure.
3. Bob signs up using self-service portal by providing his organizational details, and he outlines the justification for requesting these new servers. Bob’s request includes specific details about the computing, networking, and storage requirements for the servers he is requesting. For example, Bob might request an infrastructure that includes 20 virtual CPUs, 500 GB of storage, and two network subnets.
4. Al uses the admin portal to validate the availability of the resources Bob has requested, configures these resources, and allocates them to Bob. At this point, Al’s work is finished—he has provisioned the infrastructure for Bob and delegated to Bob the ability to implement it.
5. Bob now uses the self-service portal to create and use as many virtual servers as he needs within the “sandbox infrastructure” that Al created for him. For example, Bob could create five new dual-CPU virtual servers having 50 GB of storage each. If he later finds that he needs additional capacity to scale up his solution, he can add up to five more dual-CPU virtual servers based on the infrastructure Al has provisioned to him. Bob can thus add or reduce capacity as needed to tailor his solution to his needs and available budget, without the need of asking Al to intervene in any way.
6. For monitoring purposes, Al can access reports that let him monitor how Bob has been using the infrastructure that he provisioned to him.

Understanding the Private-Cloud Architecture

Figure 6-5 illustrates the logical architecture of the private-cloud solution based on layering the Dynamic Data Center Toolkit for Enterprises on top of a virtualization-capable infrastructure that uses Microsoft products and technologies. (This diagram is based on prerelease information and is subject to change.)




FIGURE 6-5 Architecture of Microsoft's private-cloud solution.

As the diagram indicates, the toolkit is built upon a foundation of Windows PowerShell, which is a Microsoft .NET Framework-connected environment designed for administrative automation and is integrated into the Microsoft System Center family of infrastructure management products. On top of Windows PowerShell resides code that uses the Windows Workflow Foundation (WF), which provides a programming model, in-process workflow engine, and rehostable designer to implement long-running processes as workflows within .NET applications. On top of this sits processing logic and a queuing agent. The admin and self-service portals are then implemented as Web services using Windows Communication Foundation (WCF), Microsoft's unified programming model for building service-oriented applications.

The Store (shown in the top right of Figure 6-5) uses a SQL Server database to keep track of interactions between the toolkit and its users through the two portals. The Action/Tasks Definition XML is a key component that defines what actions can be performed using the toolkit's portals. For example, the action "Create a VM" is mapped by the XML to a collection of Windows PowerShell scripts that perform a series of actions to create a new virtual machine for the user.

Implementing a Private-Cloud Solution

The Dynamic Data Center Toolkit for Enterprises enables businesses to implement a private-cloud infrastructure on top of their existing virtualization-capable datacenter. Accomplishing this will typically involve performing a series of steps to plan, design, deploy, and configure various Microsoft products, technologies, and services together with infrastructure hardware (servers, routers, switches, SANs, and so on) that support the IaaS, PaaS, and SaaS components of Microsoft's private-cloud computing platform.

As an example of how this will work, the high-level steps for implementing the IaaS portion of a private-cloud infrastructure for an enterprise can be summarized as follows (this guidance is based on prerelease information and is subject to change):

1. Determine the scope of your private cloud. This step involves determining the proposed workloads that the private cloud will host and the technical requirements for sustaining the workloads.
2. Design your physical fabric infrastructure. This step involves determining the physical fabric resources required to support the planned workloads, including
 - ❑ Types of virtualization hosts used
 - ❑ Host storage requirements and fabric
 - ❑ Cloud storage requirements and fabric
 - ❑ Physical network and topology requirements
 - ❑ Load balancer requirements
 - ❑ Types of firewalls needed and their topology
3. Design your core infrastructure. This step involves determining how you will use Hyper-V virtualization to provide hardware abstraction for the workloads and provide additional services needed, such as
 - ❑ Infrastructure services using AD DS and DNS
 - ❑ High availability and Live Migration using Failover Clustering
 - ❑ Business continuity using System Center Data Protection Manager
 - ❑ Hosting virtual desktops if needed using Remote Desktop Services
4. Design your management infrastructure. This step involves determining what services will be needed to facilitate deploying workloads, patching and monitoring hosts, and managing the life cycle of physical fabric components and workloads. Tasks for this step include designing the following:
 - ❑ Hypervisor and workload management solution using System Center Virtual Machine Manager, System Center Configuration Manager, and System Center Operations Manager

- ❑ Hardware management solution
- ❑ Physical server and life-cycle management solution
- ❑ On-board and provisioning engine solution

The Dynamic Data Center Toolkit for Enterprises will also provide step-by-step guidance for planning, designing, deploying, and configuring the hardware and software that is needed to support the PaaS and SaaS components of Microsoft's private cloud-computing platform.

Windows Azure

Complementing Microsoft's public-cloud offering is Windows Azure, a highly scalable service platform that provides users with on-demand computing and storage resources that can be used to create cloud applications and services that can be hosted in Microsoft's own data-centers. At the present time, Microsoft's offerings in this area provide two complementary types of services:

- **Windows Azure** You can think of this as the "operating system in the cloud."
- **Microsoft SQL Azure** You can think of this as a "fully relational database in the cloud."

Windows Azure runs your organization's Web applications without the need of hosting them yourself. Windows Azure, therefore, complements hosting service providers that have implemented Microsoft's public-cloud solution. (See Figure 6-6.)




FIGURE 6-6 Windows Azure and Microsoft's public-cloud solution complement each other.

Windows Azure provides a reliable, scalable, and fault-tolerant cloud platform suitable for running mission-critical, line-of-business (LoB) enterprise applications. By hosting your organization's Web applications using Windows Azure, you no longer need to worry about the following:

- Procuring and setting up hardware
- Installing and configuring operating systems
- Maintaining operating systems by patching them
- Configuring and managing network devices and settings
- Troubleshooting hardware issues, and resolving hardware failures
- Adding more memory or storage to your hardware
- Upgrading and patching applications without downtime

Instead of having you be concerned with all of these issues, Windows Azure requires only that you deploy your code plus some additional metadata onto the Windows Azure fabric, which is hosted and maintained by Microsoft at its datacenters. This metadata is called "the Service Model," and it specifies how many instances you require to run your applications and what roles these instances will play in your deployment. Examples of such roles might be Web Role, Worker Role, and so on. The Windows Azure fabric then automatically provisions these roles onto the instances, spins up the necessary network infrastructure to support those roles, and begins monitoring these roles to maintain their health. If you require more instances in order to scale up, the built-in elasticity of the Windows Azure fabric makes this as easy to do as turning a dial. Later, if you no longer need the added capacity, you can easily scale the number of running instances down with the same degree of ease to reduce costs. The Windows Azure fabric does everything else that is needed to ensure that these instances are ready and handling your application.

Windows Azure leverages the same familiar tools and technologies used for developing applications you host on your own Windows servers. Although some architectural changes are needed to take full advantage of hosting your applications in the cloud using Windows Azure, these changes can be implemented easily because Windows Azure also provides you with a local development fabric that emulates the same services provided by the cloud. After you've tested your application using the local fabric, you can push it out to the cloud so that it can be accessed from anywhere using the Internet. For more information about Windows Azure, see <http://www.microsoft.com/windowsazure>.

The Dynamic Data Center Alliance

The Dynamic Data Center Alliance is a program designed to create an ecosystem of partners that will work with Microsoft to offer customers cloud-computing solutions and support. The Alliance consists of a number of types of groups, including

- Hosting service providers
- System integrators
- Hardware manufacturers
- Software vendors

At the time of this writing, the following nine hosting service providers have already joined the Alliance and are now delivering managed hosting, cloud services, and on-demand virtualized computing resources:

- Applied Innovations (<http://www.appliedi.net>)
- Cloudmore (<http://www.cloudmore.com>)
- HostBasket (<http://www.hostbasket.com>)
- Hostway (<http://www.hostway.com>)
- MaximumASP (<http://www.maximumasp.com>)
- Peak 10 (<http://www.peak10.com>)
- PoundHost (<http://www.poundhost.com>)
- StarUK (<http://www.star.co.uk>)
- Reliance Data Center (<http://www.relianceidc.com>)

Availability of Microsoft's Cloud-Computing Platform

Microsoft's Dynamic Data Center Toolkit for Hosters is already available for free today and is updated as Microsoft's public cloud-computing platform continues to evolve.

Microsoft's Dynamic Data Center Toolkit for Enterprises will be available for free, and its release date is currently scheduled for the first half of 2010.

At the time of this writing, Windows Azure is currently in Community Technology Preview (CTP) and the services are scheduled to remain free for evaluation purposes until February 1, 2010, after which time Microsoft will begin charging customers for use.

Additional Resources

The resources listed in this section were either referred to in the chapter or provide additional information about concepts and products discussed in this chapter.

Additional Resources for Microsoft's Cloud-Computing Initiative

For general information about Microsoft's public and private cloud offerings, including links to white papers, presentations, videos, and other material, see <http://www.microsoft.com/privatecloud>.

Information about Microsoft's Dynamic Data Center Toolkit for Hosters can be found at <http://www.microsoft.com/dynamicdatacenter>.

For the latest news about Microsoft's public-cloud offering, see the Dynamic Data Center Alliance Blog at <http://blogs.technet.com/ddcalliance/default.aspx>.

Additional Resources for Windows Azure

For general information about Windows Azure, including links to white papers, presentations, videos, and other material, see <http://www.microsoft.com/windowsazure>.

For the latest news about the Windows Azure platform, see the Windows Azure Team Blog at <http://blogs.msdn.com/windowsazure>.

Index

A

access

- anytime, anywhere, any device, 10
- to RD Session Host servers, 236

- Action menu (Virtual Machine Connection tool), 71

- Active Directory Domain Services.
See AD DS (Active Directory Domain Services)

- Active Directory Service Interfaces (ADSI), App-V use, 183

- Active Directory Users and Computers console, RD Session Host management, 244

- Active Upgrade feature (App-V), 167

- Add Hardware settings (Hyper-V Manager Action pane), 67, 68

- Add Hosts Wizard, 363

- address translation, 38

- Add Roles Wizard (Server Manager), 54, 232, 233–236

- Configure Client Experience page, 231

- licensing options, 281

- AD DS (Active Directory Domain Services), 149, 155

- in cloud-computing solutions, 438
- for management of resources, 3
- requirements for, 152

- VMM and, 348

- Add Storage dialog box, 90

- Administration view (VMM

- Administrator Console), 359

- Administrator roles (VMM), 341, 350, 413

- agent versions, managed host, 367

agility

- achieving, 7
- cloud computing and, 433, 434
- improving, 7, 8

- Allow Management Operating System To Share This Network setting (Virtual Network Manager), 62

- AllowUnencryptedTransfers parameter, 334

- AMD processors, processor

- compatibility mode and, 95

- antimalware software, for Windows XP Mode virtual machines, 122

- antivirus software

- on servers, 56
- for Windows XP Mode virtual machines, 122

- anytime, anywhere, any device access, 10

- optimizing, 14

- application compatibility issues, 112.

- See also* virtual applications

- Med-V and, 111

- resolving, 6, 8, 11, 113

- Windows Virtual PC and Windows XP Mode and, 110, 113

- application deployment

- accelerating, 226

- optimizing, 10

- risks of, 226

- application failover, 48

- application life-cycle management, 7

- application packages. *See* virtual

- application packages

- applications. *See* virtual applications

- Application Source Root (ASR), 164

- application virtualization, 109,

- 162. *See also* App-V; virtual application packages; virtual applications

- Application Virtualization.

- See* App-V (Application Virtualization)

- Application Virtualization 4.5 for Remote Desktop Services Service Pack 1, 228

- Application Virtualization for Terminal Services (App-V for TS), 164, 168, 223. *See also* App-V for RDS

- Applied Innovations, 446

- Apply option (Snapshots pane), 83

- App-V, 8, 11

- for mobile worker scenario optimization, 12

- App-V (Application Virtualization), 8, 11, 109, 115, 162–219

- administrators, 201

- applications, managing, 172–176

- applications, publishing, 171

- applications, sequencing, 170–171

- architecture of, 182–184

- availability of, 115

- benefits of, 111–112, 112

- branch office deployment, 190

- client stream policy, 163

- communication ports, 184

- components of, 176–182

- deployment scenarios, 185–192

- Dynamic Suite Composition, 163

- ESD system deployment, 190–191

- for mobile worker scenario

- optimization, 12

- for office worker scenario

- optimization, 13

- globalization features, 164–165

- implementation benefits, 12–13

- management of, 164

- packages, streaming, 171–172

- reports on, 199–200

- resources on, 221–222

- scalability features, 163, 186–188

- security features, 165

- single-site deployment, 189–190

- standalone deployment, 192

- Terminal Services deployment, 192

- terminology, 167–169

- troubleshooting, 217–219

- usage scenarios, 113–114

- version 4.5 SP1, 166, 301

- version 4.6, 166–167

- virtual environments, 169–170, 174

- App-V Client console, 182, 214

- Applications node, 214–215

- File Type Associations node, 215

- Publishing Servers node, 215–216

- App-V Clients, 181–182, 214–216
- managing from command line, 216

- Standalone mode, 114

- App-V Data Store, 178

- App-V Desktop Client, 181, 301

- App-V for RDS, 223, 301–302

- availability of, 227–228

- benefits of, 224

- resources on, 310–311

- usage scenarios, 226

- App-V Management Console, 168, 177, 179–180, 192–201

- Administrators node, 201

- Application Licenses node, 198–199

- Applications node, 192–197

- File Type Associations node, 200

- Packages node, 197

- Provider Policies node, 201

- Reports node, 199–200

- Server Groups node, 199

- App-V Management Server, 113, 168, 177

- for application deployment, 185

- App-V Management Web Service, 178

- App-V Sequencer, 168, 174, 180–181, 201–213. *See also* Sequencing Wizard

App-V Sequencer (*continued*)
 automated installation feature, 192
 .msi file creation, 114
 App-V Sequencer Console, 201, 207
 Change History tab, 208–209
 Deployment tab, 208
 Files tab, 208, 209
 OSD tab, 210
 Properties tab, 207
 Virtual File System tab, 209
 Virtual Registry tab, 208, 209
 Virtual Services tab, 209, 210
 App-V servers, managing, 199
 App-V Streaming Server, 178
 App-V Terminal Services Client, 181
 archives, tape-based, 428
 Armstrong, Ben, 108
 Assign Personal Virtual Desktop Wizard, 295
 Attach A Virtual Hard Disk Later option (New Virtual Machine Wizard), 78
 audio integration component (virtual machines), 136
 authentication
 forms-based, 256
 NLA, 236
 Automatic Start Action setting (Hyper-V Manager console), 68, 91
 Automatic Stop Action setting (Hyper-V Manager console), 69
 auto publish settings (virtual machines), 120, 138
 availability
 of App-V, 115
 increasing, 7
 of MED-V, 114–115
 of Windows Virtual PC, 114
 of Windows XP Mode, 114
 .avhd (differencing disk) files, 83

B

background jobs, Windows PowerShell support for, 99
 backup and recovery, simplifying, 6
 Backup service, 68
 Ballmer, Steve, 16, 18
 Basic IT infrastructures, 3
 description of, 3
 moving to Standardized, 5–6
 .bin (memory contents) files, 83
 BIOS setting (Hyper-V Manager Action pane), 67
 BitLocker Drive Encryption, 12
 BlockLmifHostBusy parameter, 335
 blogs on Hyper-V, 107–108

branch office scenarios
 cloud computing solutions for, 435
 Remote Desktop Services for, 224–225
 Bring This Service Or Application Online option (Hyper-V Manager Actions pane), 91
 Brown, Taylor, 108
 business agility. *See* agility
 business continuity
 enhancing, 6, 8
 Hyper-V features for, 43
 Med-V capabilities for, 111
 business continuity planning with VMM, 341–342
 businesses
 concerns of, 1
 environmental friendliness of, 1
 spending decisions, drivers of, 1
 survival of, 1
 business partners, Remote Desktop Services and, 225
 business value, creating, 1

C

capacity on demand, 8, 11. *See also* cloud virtualization
 capital expenditures, controlling, 1
 centralized execution, 11
 central processing unit identification (CPUID) values, 86
 certificates for RD Session Host authentication, 238
 child partitions, 33–37. *See also* parent partitions; partitions
 communication with parent, 33
 creating and managing, 28
 device configuration and control for, 31–32
 running non-Hyper-V-aware operating systems, 35
 running non-Windows operating systems, 34–35
 running Windows operating systems, 34
 types of, 28, 33
 Citrix XenApp servers, 173
 Citrix XenServer, 24
 client computing scenarios, optimizing with virtualization, 12–14
 client virtualization, 8
 benefits of, 10–19
 clipboard, configuring behavior, 156
 Clipboard integration component (virtual machines), 136
 Clipboard menu (Virtual Machine Connection tool), 72
 cli.psc1 console file, 321
 cloning virtual machines, 400
 close settings (virtual machines), 138
 cloud computing, 8
 benefits of, 433
 core infrastructure, 443
 description of, 431–432
 Dynamic Data Center Alliance and, 446
 elements of, 431
 implementing solutions, 437–439
 management infrastructure, 443–444
 physical fabric infrastructure, 443
 private-cloud solutions, 432. *See also* private-cloud solutions
 public-cloud solutions, 432. *See also* public-cloud solutions
 resources on, 447
 usage scenarios, 435
 Windows Azure and, 444–445
 Cloudmore, 446
 cloud virtualization, 1, 8. *See also* cloud computing
 benefits of, 11
 clustered virtual machines
 creating, 90
 high availability of, 88, 91
 cluster networks, Live Migration configuration, 91
 Cluster service, VM state management, 91
 Cluster Shared Volumes (CSV)
 feature, 37, 88
 enabling, 90
 storage, adding, 90
 VMM support for, 333
 cmdlets for Managing Hyper-V, 101–103
 CodePlex Open Source repository, 103
 CodePlex Project, 100, 106
 collaboration between users, 4
 COM 1\2 setting (Hyper-V Manager Action pane), 67
 command line, App-V Client management from, 216
 command-line tools, RD Session Host management, 245
 Common Criteria certification, 107
 communication ports for VMM, 318–319
 compatibility issues, resolving, 6, 8, 11
 compliance
 meeting, 1
 self-provisioning software and quarantine systems for, 4

- COM port redirector (Virtual PC Host), 117
- COM port settings (virtual machines), 132–133
- computing resources, optimizing use of, 42
- configuration files (virtual machines)
 - best practices for, 84–85
 - default location, 84
 - relocating, 78–79
- Configure A Service Or Application setting (Hyper-V Manager Actions pane), 91
- Configure Virtual Desktops Wizard, 292–295
- Confix.xml file, 80
- Connect action (Hyper-V Manager console), 66
- Connection Authorization Policy (CAP) policies, 277
- connection status, managed hosts, 367
- Connect To Server dialog box (VMM Administrator Console), 356
- Content directory (App-V), 167
- contract/offshore worker scenarios, 14
- Control Panel, installing RD Session Host from, 241
- conversions, machine
 - vs. migration, 411
 - P2V, 400–411
 - V2V, 411–412
 - with VMM, 338
- Convert Physical Server (P2V) Wizard, 402–409
- Convert Virtual Machine Wizard, 412
- core parking, 40
- Core VDevs, 31–32
- corporate mergers, Remote Desktop Services and, 225
- corporate networks, external user access, 276
- cost efficiency of virtualization, 1
- costs of Microsoft virtualization solutions, 16
- CPU resource allocation, 85, 86
- Create A Virtual Hard Disk option (New Virtual Machine Wizard), 77
- Create User Role Wizard, 416–420
- credentials
 - for connecting to virtual machines, 60
 - deleting, 60
- Custom tab (managed host properties), 372

D

- datacenters, 9. *See also* server consolidation
- datacenter virtualization, 8
 - benefits of, 9–19
- data, centralized management of, 14
- Data Exchange service, 68
- Data Execution Prevention (DEP), 49
- data sections, Windows PowerShell support for, 100
- data transfer control, 148
 - with MED-V, 149
- debugger cmdlets (Windows PowerShell), 100
- deep C state, 40
- Default Server Group (App-V), 199
- Default Web Site (IIS), deleting, 344–345, 356
- defense in depth
 - microkernel hypervisors and, 27
 - monolithic hypervisors and, 25
- delegated administration, 332–333
- Delegated Administrator roles (VMM), 341, 413
 - creating, 413–415
- delegated management, 339
- Delete Snapshot option (Snapshots pane), 83
- Delete Snapshot Tree option (Snapshots pane), 83
- Deploy Virtual Machine Wizard, 399
- desktop applications, virtualizing, 8
- desktop clients, 167
 - migrating to Windows 7, 110
 - virtualized environments on, 109
- desktop life-cycle management, simplifying, 7
- desktop management, centralizing, 111
- desktop operating systems, virtualizing, 8, 11
- Desktop Optimization Pack 2009 R2, 17
- desktop services
 - management, manual, 3
 - management, optimized, 4
 - partially automated, 3
 - standards and policies for, 3
- desktop virtualization, 1, 12–19
 - definition of, 109
 - local, 109
 - remote, 109
- Details pane (VMM Administrator Console), 372
- development
 - Hyper-V and, 43
 - snapshots and, 81
- device drivers
 - hypervisor-aware, 25
 - security considerations, 26
- device redirection, 277
- devices, adding and removing, 30
- Devices component (Virtual PC Host), 117
- Diagram view (VMM Administrator Console), 360
- differencing disks, 130
- differential SFTs, 164
- disaster recovery, 6, 43
- disk-based storage, 428–429
- disk controllers, managing with Windows PowerShell, 102
- disk drives, managing with Windows PowerShell, 102
- Diskette Drive setting (Hyper-V Manager Action pane), 67
- disk images, managing with Windows PowerShell, 102
- disruption to services, minimizing, 8
- Distributed Component Object Model (DCOM), 318
- domain controllers, 87
- Donahue, Sean, 176, 188, 213
- Do Not Use Remote Desktop Session Host Server IP Address When Virtual IP Address Is Not Available policy setting, 246
- drives integration component (virtual machines), 136
- DSI (Dynamic Systems Initiative). *See also* Dynamic IT
- dual virtual processors, 38
- DVD drives, adding to VMs, 67
- DVD drive settings (virtual machines), 131, 142
- Dynamic Data Center Alliance, 446
- Dynamic Data Center Alliance Blog, 447
- dynamic datacenters, 43
- Dynamic Data Center Toolkit, 436, 438–439
- Dynamic Data Center Toolkit for Enterprises, 439
 - availability of, 446
 - in cloud-computing solutions, 437, 438
 - vs. Dynamic Data Center Toolkit for Hosters, 440
 - private-cloud infrastructure implementation, 443
- Dynamic Data Center Toolkit for Hosters, 439
 - availability of, 446
 - vs. Dynamic Data Center for Enterprises, 440–441
 - resources on, 447

- Dynamic IT, 2–7
 - pillars of, 9
 - strategic vision for, 18
- Dynamic IT infrastructures, 2–3
 - description of, 4
- dynamic storage provisioning, 6
- Dynamic Suite Composition (DSC), 163, 167
- Dynamic Systems Initiative (DSI), 7

E

- Edit Virtual Hard Disk Wizard, 64
- efficiency gains from virtualization, 1
- electronic software distribution (ESD) systems, 112
 - for application deployment, 185, 190
 - publishing virtual applications with, 171
- emulated devices, 31
 - for non-Hyper-V-aware operating systems, 35
- emulation, 33, 119
- Enable Cluster Shared Volumes
 - option (Failover Cluster Manager), 90
- encryption, 238–239
- endpoint computers, 151
- end-user experience, 148
- enlightenments, 34
- Enterprise Desktop Virtualization, 8, 11
- environmentally friendly businesses, 1
- ESX Server hosts
 - configuring, 376
 - controlling, 398
 - credentials for communication, 377–378
 - managing, 374
- ESX Servers, 411–412
- execution environments, isolating, 23
- execution policy, changing and viewing, 322
- .exp (export) files, 81
- exporting virtual machines, 79–81
- Export option (Hyper-V Manager console), 80
- Export Virtual Machine dialog box, 80
- Extended Page Tables (EPT), 39
- external users, internal network access, 276

F

- failover clustering
 - in cloud-computing solutions, 438
 - VMM support for, 331, 333
- Failover Clustering feature, 88
 - installing, 89
 - leveraging, 43
 - Microsoft Hyper-V Server and, 46
 - VMs per cluster, 38
- Failover Cluster Manager, 90
- failover clusters
 - Cluster Shared Volumes feature, 88
 - two-node, 88–89
 - witness disk, 89
- File menu (Virtual Machine Connection tool), 71
- file shares for VM configuration files, 87
- file transfer tool (MED-V), 155
- file type associations, 200, 215
- filters (VMM Administrator Console), 360
- firewalls, 4
- floppy disks, connecting to, 67
- Folder Redirection, 155
 - for mobile worker scenario optimization, 12–19
 - for sharing user data, 13
 - for user state virtualization, 110
- forms-based authentication, 256
- full application virtualization infrastructure, 113
- functionality, locking down, 13

G

- globalization, 164–165
- globally unique identifier (GUID), 7
- graphical user interface in Hyper-V, 48
- Green businesses, 1
- green IT objectives, 40
- Greschler, David, 17
- Group Policy
 - RD Gateway configuration with, 280
 - for RD Session Host management, 246–250
 - RD Session Host authentication and encryption settings, 239, 240
- Group Policy Software Installation, 158
- guest operating system profiles, 313
 - in virtual machine templates, 380
 - VMM library, adding to, 380–381
- guest operating systems, 21, 151

- applications on, 144
- desktop display, 139
- hardware resources, access to, 23
- hibernating, 139
- Hyper-V support for, 50–52
- installing, 76–77
- installing applications on, 121
- Integration Services functionality, 52–53
- integration with host, 134–136
- interaction with hardware, direct, 25
- resolution of desktop, 135
- restarting, 139
- separation from hardware, 24, 26
- sleep mode, 139
- 32-bit and 64-bit, 52
- for Windows Virtual PC, 123
- guest virtualization rights, 48
- GUID (globally unique identifier), 31

H

- hard disk settings (virtual machines), 130
- hard drives
 - adding to VMs, 67
 - managing with Windows PowerShell, 102
- hardware-assisted virtualization, 36
 - support for, 48
 - technologies for, 36
- hardware failure events, 29
- hardware, inventorying, 4
- hardware profiles, 314
 - in virtual machine templates, 379
 - VMM library, adding to, 380
- hardware resources
 - adding to virtual machines, 67
 - allocating, 28
 - failover clustering support, 89
 - guest operating system access to, 23
 - managing and assigning, 28
 - sharing, 29, 36
 - virtual machine control of, 22
- Hardware tab (managed host properties), 369
- hardware utilization, increasing, 7
- heartbeat monitoring, 135
- heartbeats, 35
- Heartbeat service, 68
- Help files in Windows PowerShell, 100
- high availability of clustered virtual machines, 88, 91
- High Availability Wizard, 91
- home office settings, 278
- host agents, 329

- HostBasket, 446
 - host clustering, 48, 374
 - hosted virtualization, 24
 - host groups, 361–365
 - hosting service providers, 446
 - host operating systems, integration with guests, 134–136
 - host processor pools, 40
 - hosts, 151. *See also* managed hosts
 - clustered pools of, 9, 10
 - communication ports, 319
 - definition of, 314
 - Intelligent Placement process and, 339
 - local desktop virtualization on, 109
 - maintenance on, 335
 - managing with VMM Administrator Console, 361
 - moving virtual servers between, 10
 - virtualized environments on, 109
 - Hosts view (VMM Administrator Console), 358
 - Hostway, 446
 - Howard, John, 107
 - HTTP, and App-V, 183
 - HTTPS, 318
 - App-V use, 183
 - hypercalls
 - definition of, 28
 - resources on, 107
 - Hyper-V, 8, 9, 21
 - antivirus software and, 56
 - architecture, 27–35
 - benefits of, 41–42
 - child partition support, 33
 - in cloud-computing solutions, 438
 - Cluster Shared Volumes feature, 37, 88
 - Common Criteria certification, 107
 - deployment planning, 53–54
 - deployment resources, 105–106
 - domain controller support, 87
 - extensibility, 36
 - features, 36–37
 - Full or Server Core installations, 28
 - hardware-assisted virtualization, 36
 - hardware-sharing architecture, 36
 - hot adding and removal of storage, 37
 - Hyper-V Manager snap-in, 57–70
 - installing, 54–55
 - Integration Services, 52–53. *See also* Integration Services
 - key features, 36–41
 - Linux Integration Components For Hyper-V, 35
 - Live Migration feature, 10, 37. *See also* Live Migration
 - management and maintenance resources, 106
 - managing with RemoteApp, 98
 - managing with Remote Desktop Connection, 97
 - managing with Windows PowerShell, 99–103, 103–104
 - managing with WMI, 98–99
 - microkernelized architecture, 27, 36
 - versus Microsoft Hyper-V Server, 46
 - network adapters in, 70–71
 - networking model, 64
 - new R2 features, 37–40
 - operating system consolidation, 41
 - operating system support, 36
 - parent partition components, 29
 - performance features, 38–40
 - planning resources, 105
 - processor compatibility mode, 37
 - RD Virtualization Host role and, 285
 - for resource optimization, 42
 - resources on, 104
 - scalability, 36, 40
 - Server settings, 59
 - SMP support, 36
 - system consolidation, 41
 - system requirements, 48–50
 - troubleshooting startup problems, 54
 - type 1 hypervisors, 24
 - usage scenarios, 42–45
 - User settings, 60–61
 - versus Virtual Server, 40–42
 - VMM management of, 331, 338
 - Windows PowerShell cmdlets, 101–103
 - vs. Windows Virtual PC, 119–120
 - WMI APIs, 30
 - working with, 46–104
 - workloads, consolidating, 41
 - Hyper-V Administrators role, 73
 - Hyper-V forum on Microsoft TechNet, 108
 - Hyper-V hosts, 397
 - Hyper-V I/O architecture, 67
 - hypervisors, 23–28
 - microkernel, 26–27
 - monolithic designs, 25
 - overhead of, 26
 - security of, 25, 26–29
 - Type 1, 23–24
 - Type 2, 24
 - Hyper-V Manager console, 57–70
 - Actions pane, 65–67
 - Allow Management Operating System To Share This Network option, 62
 - Edit Disk option, 64–65
 - exporting and importing virtual machines, 79–81
 - Hyper-V Settings option, 59
 - Import Virtual Machine option, 59
 - Keyboard settings, 74
 - MAC Address Range option, 62
 - New | Floppy Disk option, 59
 - New | Hard Disk option, 59
 - New | Virtual Machine option, 58
 - for remote server management, 57–58
 - server settings, configuring, 58–59
 - snapshots actions, 82–83
 - virtual machine hardware settings, 65–70
 - virtual machine management settings, 68–69, 79–87
 - Virtual Network Manager option, 61
 - “Hyper-V Planning and Deployment Guide”, 55
 - Hyper-V Security Guide Solution Accelerator, 107
 - Hyper-V Server
 - versus Hyper-V server role, 47–48
 - migrating to Windows Server 2008, 47
 - Hyper-V servers
 - moving virtual machines, 59
 - remote desktop virtualization on, 109
 - Hyper-V WMI provider, 98, 99
- I**

 - IBM, 23
 - IBM virtual machines, 18
 - .ico files, 170
 - IDE controllers, 78
 - IDE Controller setting (Hyper-V Manager Action pane), 67
 - importing virtual machines, 79–81
 - compatibility issues, 80
 - Import Virtual Machine dialog box, 81
 - Import Virtual Machine option (Hyper-V Manager console), 81
 - Infrastructure as a Service (IaaS), 436–438
 - implementing, 443–444
 - Infrastructure Optimization blog, 18
 - Infrastructure Optimization Model, 2–4
 - moving between stages, 5–7
 - virtualization and, 4–7

Infrastructure Optimization TechCenter, 18

Infrastructure Planning and Design (IPD) guides, 439

Infrastructure Planning and Design (IPD) Guides for Virtualization, 53

Install Application On Remote Desktop Server option (Control Panel), 241

installation directory (App-V), 167

Install Program From Floppy Disk Or CD-ROM Wizard, 241

integrated Microsoft virtualization solutions, 15–17

benefits of, 16

Integrated Scripting Environment (ISE), 100

Integration Components (ICs) (Virtual PC), 117, 118

basic integration mode, 135–136

enhanced integration mode, 136–137

integration features (virtual machines), 133–136

disabling, 142

Integration Services, 34

displaying, 68

functionality provided by, 52–53

installing, 77, 85

network adapters and, 70

for unenlightened guests, 77

for Windows guest operating systems, 35

Integration Services setting (Hyper-V Manager console), 68

Intelligent Placement process, 339

Intel processors, processor compatibility mode and, 95

Intel Virtualization Technology For Connectivity, 39

internal networks, external user access, 276

IP addresses, assigning to Remote Desktop connections, 232

iSCSI SANs, 333

IT as a service, 431, 435

IT costs, 4

IT infrastructures

- efficient use of, 434
- flexibility of, improving, 42–43
- management complexity of, 11
- responsiveness of, improving, 42–43
- strategic value of, 4
- types of, 3
- virtualized, 11

IT processes, fully automated, 4

IT resources, efficient use of, 434

J

Jobs view (VMM Administrator Console), 359

jumbo frames, 39

K

keyboard accelerators

- Remote Desktop sessions and, 97
- for use with virtual machines, 74

keyboard settings, 60, 137

keyboard shortcuts in Virtual PC, 137

key/value pairs, exchange of, 35

Kidaro, 147

L

large memory support, 48

legacy applications, running in virtual machines, 110, 111

legacy network adapters, 70–71

library servers. *See also* VMM library; VMM Library Servers

definition of, 314

Library view (VMM Administrator Console), 359

licensing

- costs, 7
- managing, 4
- VDI, 304, 306
- for virtual applications, 198
- Windows VECD for SA version, 305
- Windows VECD version, 305

light-touch technologies, 4

lightweight application

- virtualization infrastructure, 114

LimitCPUForMigration parameter, 335

Linux, as guest operating system, 51

Linux Integration Components For Hyper-V, 35

Live Migrate Virtual Machine To Another Node (Failover Cluster Manager), 92

Live Migration, 10, 37, 88–96

- in cloud-computing solutions, 438
- Hyper-V support for, 48
- Microsoft Hyper-V Server and, 46
- setting up for, 89–92

live migrations, 92–94

- time to perform, 93
- Windows PowerShell for, 94

load balancing terminal servers, 272

local desktop virtualization, 11, 109

localization, 164–165

logical IT infrastructures, 7

logical processors, 40

logical unit numbers (LUNs), 88

login credentials settings (virtual machines), 138

M

MAC address conflicts, 62

MAC Address Range option (Virtual Network Manager), 62–63

machine conversions.

- See* conversions, machine

maintenance

- costs, 7
- optimizing, 10

malware filtering, 4

managed host agents

- installing, 364
- reassociating, 364

managed hosts

- adding, 362
- definition of, 314
- managing, 365–373
- Networking view, 373–374
- removing, 363
- status of, 367–368
- working with, 361–377

management interfaces for server virtualization, 22

management operating systems

- isolating, 62
- remote connection to, 62

management platforms, unified, 9

management settings for virtual machines, 68–69

Manage RDS CALs Wizard, 283

MaximumASP, 446

Media menu (Virtual Machine Connection tool), 72

MED-V, 8, 11, 109, 147–162

- availability of, 114–115
- benefits of, 111–112
- centralized management capabilities, 148–149
- centralized monitoring capabilities, 148–149
- data transfer control, 149
- description of, 150
- desktop modes, 159–162
- end-user experience, 149–150
- file transfer tool, 155
- full-desktop mode, 159
- Management console, 157, 161, 162
- package deployment, 153
- resources on, 220
- single-desktop mode, 159, 160
- supported applications, 153

MED-V (*Continued*)
 system requirements, 151–152
 technologies provided by, 147–148
 terminology, 151
 usage policies, configuring, 156–157
 usage policies, enforcing, 149
 usage scenarios, 113
 virtual image management, 148, 157–159
 workspace, initializing, 154–155
 workspace, using, 155–156

MED-V clients, 151
 deploying, 158
 system requirements for, 151

MED-V image repository, 151

MED-V management servers, 151
 user authentication, 159

MED-V packages, 151. *See also* virtual application packages

MED-V Server, 151
 system requirements for, 152

memory management for server virtualization, 22

Memory setting (Hyper-V Manager Action pane), 67

memory settings (virtual machines), 130

mergers, Remote Desktop Services and, 225

microkernel hypervisors, 26–27, 36

Microsoft
 virtualization, commitment to, 16–17
 virtualization products and technologies resource, 19

Microsoft Advanced Group Policy Management, 115

Microsoft Application Virtualization (App-V). *See* App-V (Application Virtualization)

Microsoft Application Virtualization for Remote Desktop Services. *See* App-V for RDS

Microsoft Assessment and Planning (MAP) Toolkit, 105

Microsoft Asset Inventory Service, 115

Microsoft Background Intelligent Transfer Service (BITS) version 2, 157

Microsoft Baseline Configuration Analyzer (MBCA), 349

Microsoft cloud-computing platform, 435–446. *See also* cloud computing availability of, 446

Microsoft Desktop Optimization Pack (MDOP) 2009 R2, 17

Microsoft Desktop Optimization Pack (MDOP) for Software Assurance, 181

App-V availability in, 115

Med-V availability in, 114

Microsoft Diagnostics and Recovery Toolset, 115

Microsoft Dynamic Data Center Toolkit. *See* Dynamic Data Center Toolkit

Microsoft Dynamic IT initiative cloud computing and, 433

Microsoft Enterprise Desktop Virtualization (MED-V). *See* MED-V

Microsoft Enterprise Support Windows Server Core Team, 108

Microsoft Exchange Server, virtualizing, 45

Microsoft Hardware-Assisted Virtualization Detection Tool, 124

Microsoft Hyper-V Server 2008 R2, 8, 9, 17, 21
 versus Hyper-V server role, 46–48
 management of, 47

Microsoft Installer (MSI), 164

Microsoft Management Console (MMC) snap-ins, 243

Microsoft partners, 16

Microsoft Silverlight, 230

Microsoft SQL Azure, 444–445

Microsoft System Center, 9

Microsoft System Center Desktop Error Monitoring, 115

Microsoft System Center Virtual Machine Manager. *See* VMM (Virtual Machine Manager)

Microsoft Update, App-V support for, 164

Microsoft Virtual Desktop Infrastructure. *See* VDI (Virtual Desktop Infrastructure)

Microsoft Virtualization, 44–46

Microsoft Virtualization Team Blog, 107

Microsoft Virtual PC. *See* Virtual PC

Microsoft Virtual Server
 versus Hyper-V, 40–42
 online resources for, 41
 Type 2 hypervisors, 24
 VMM support for, 338

Microsoft Virtual Server 2005
 exporting virtual machines from, 80

R2 SP1, 21

Microsoft Volume License Services (MVLS) Web site, 227

Microsoft Windows Azure. *See* Windows Azure

Microsoft Windows Virtualization Product Group Team Blog, 19

middleware applications, one-to-many arrangement, 163

midsized businesses
 business survival, 2
 Infrastructure Optimization Model for, 2

Migrate To A Physical Computer With A Different Processor Version setting (Processor configuration), 86, 96

migration, 36
 vs. conversion, 411
 Quick Storage Migration, 337
 with VMM, 399

migration to Windows 7
 legacy applications and, 110, 111
 testing plans on virtual machines, 111

mobile workers
 access levels, 4
 optimizing scenarios, 12–13
 profile cache management, 232
 Remote Desktop Services and, 225
 VDI and, 227

monolithic hypervisors, 25

motherboard settings, 101

Mount-VHD cmdlet (Windows PowerShell), 102

mouse movement, 135
 Remote Desktop sessions and, 97

mouse release key, 60

MoveWithinCluster parameter, 335

.msi files, standalone delivery of, 185

MSTSCAX.dll, 119

Muglia, Bob, 18

N

named pipes, 132

Name setting (Hyper-V Manager console), 68

name settings (virtual machines), 130

NAT component (Virtual PC Host), 117

Nested Page Tables (NPT), 39

.NET Remoting, 183

Network Access Protection (NAP), 277

network adapters
 adding, 67

network adapters (*continued*)
 configuring, 67
 in Hyper-V, 70
 physical, 63
 synthetic, 77
 Network Adapter setting (Hyper-V Manager Action pane), 67
 network infrastructure for failover clusters, 89
 networking
 Hyper-V model, 64–65
 resources for server virtualization, 22
 networking settings (virtual machines), 132
 Networking tab (managed host properties), 370
 Networking view (VMM Administrator Console), 360, 373–374
 network interface cards
 manipulating with Windows PowerShell, 102
 multiple, 55
 Network Level Authentication (NLA), 236
 network load balancing, 36
 network resources, centralized management of, 3
 Network Service built-in identity, 31
 New Application Wizard (App-V), 194–197
 New Guest OS Wizard, 381
 New User Role Wizard, 413–420
 New Virtual Hard Disk Wizard, 59
 New Virtual Machine Wizard, 58, 76, 77, 384–391

O

Ocsetup.exe utility, 54
 office worker scenarios, optimizing, 13
 Offline Files
 for mobile worker scenario optimization, 12–19
 for user state virtualization, 110
 Offline Virtual Machine Servicing Tool 2.1, 107
 OLE DB, App-V use, 183
 O'Neill, James, 104, 108
 Oobe.exe, 232
 Open Database Connectivity (ODBC), App-V use, 183
 operating systems
 consolidating, 41
 enlightened, 34
 Hyper-V support for, 36
 incompatibilities, 113

multiple, 21
 P2V conversion support, 401
 for V2V conversions, 412
 for VMM, 346–348
 Operating System Shutdown service, 68
 operational expenditures, predictability of, 1
 “Optimize and Secure Your Core Infrastructure” white paper, 18
 .osd files, 164, 170
 outsourcing firms, 225
 overall status, managed hosts, 367

P

P2V agents, 329
 P2V (physical-to-virtual) conversions, 314, 338, 400–411
 requirements for, 401
 VMware to Hyper-V, 412
 paging tables, 38
 parent partitions, 28–33
 communication with children, 33
 components of, 29
 definition of, 28
 microkernel hypervisors and, 26
 overhead of, 27–28
 purposes of, 28
 virtualization stack, 29
 VM communication with, 67
 partitions, 19
 definition of, 28
 parent partitions, 26, 28–33
 types of, 28
 passthrough disks
 configuration files location, 78
 considerations for use, 87
 performance and, 87
 virtual machines, creating with, 77–78
 patch management
 manual, 3
 partially automated, 3
 Peak 10, 446
 performance
 Hyper-V features for, 38–40
 monitoring, 56
 passthrough disks and, 87
 of Type 2 hypervisors, 24
 Performance and Resource Optimization. *See* PRO (Performance and Resource Optimization)
 “Performance Tuning Guidelines for Windows Server 2008” white paper, 55
 personal virtual desktops, 285–286
 connecting to, 297–298
 provisioning, 288–299
 physical desktop computers, virtualized, 109
 physical devices. *See also* hardware resources
 accessing, 22
 Physical Hard Disk option (Hyper-V Manager snap-in), 78
 physical network adapters, 63
 physical serial ports, 132
 physical servers. *See also* hardware resources; servers
 hosting Hyper-V role, 55–57
 physical-to-virtual conversions. *See* P2V (physical-to-virtual) conversions
 Placement tab (managed host properties), 370
 Platform as a Service (PaaS), 436, 437
 plug and play operations, 29
 plug-in VDevs, 32
 port groups, 337
 PoundHost, 446
 power management, 29
 power requirements, reducing, 5, 7
 printer integration component (virtual machines), 136
 privacy, 1
 private-cloud solutions, 11, 432
 architecture of, 441–442
 benefits of, 433
 implementing, 443–444
 infrastructure of, 437
 Microsoft products for, 436
 portals, 440–441
 processor compatibility mode, 37
 disabling, 86
 enabling, 86, 96
 Processor Compatibility Mode feature, 88, 95–97
 Processor configuration settings (virtual machines), 86
 Processor setting (Hyper-V Manager Actions pane), 67
 processors, multiple, 48
 processor virtualization extensions, 50
 profile cache management, 232
 profile management, 226
 PRO (Performance and Resource Optimization), 314, 326–327, 338, 341–342, 374
 capabilities of, 332
 Windows PowerShell cmdlets, 327
 public-cloud solutions, 11, 432
 benefits of, 433
 Microsoft products for, 436
 publishing servers, 215, 216
 publishing virtual applications, 168

Q

Q: drive, 168
 installing applications on, 211–213
 quad virtual processors, 38
 quarantine-like systems, 4
 Quick Migration process, 79–81
 Quick Storage Migration, 337

R

RAM
 for Hyper-V server role, 49
 for virtual images, 153
 for virtual machines, 67
RapidProvisionVM.ps1 script, 336
RapidProvisionVMwithAnswerFile.ps1 script, 336
Rationalized IT infrastructures, 3
 description of, 4
 moving to Dynamic, 7
RD Connection Broker, 272–276
 functionality, 273
 installing, 273
 Load Balancing feature, 273, 276
RD Connection Broker servers, 287
 configuring, 274–276
 function of, 289
 as RemoteApp program source, 260
 security groups, configuring, 275–276
RD Connection clients, 230–231
 enhancements to, 230–231
 features of, 231–232
RD Connection (RDC), 97–98, 229
RD connections
 encryption of, 238–239
 IP addresses, assigning to, 232
RD Gateway, 14, 276–281
 authentication and authorization plug-ins, 277
 configuring, 280–281
 functionality, 278
 home office usage scenario, 278
 installing, 279
 new features, 277
 security and, 276–277
RD Gateway Manager console, 281
RD Gateway servers, 277
RD Licensing, 281–284
 installing, 282
 new features, 282–283
RDP Encoder Technology (Virtual PC Host), 117
RDP encryption, 237
RDP Server service, 118
RDPShell.exe, 118, 119
RDP-Tcp Properties dialog box
 (Remote Desktop Session Host Configuration console), 239
RDS CALs
 managing, 281, 283
 single versus packs, 283
RD Session Host, 231–256
 access, configuring, 236–237
 applications, installing on, 240–242
 authentication settings, 236–237
 client experience, 235
 Group Policy settings, 246, 248–250
 installing, 232–235
 managing, 242–250
 new features, 231–232
 RD-Execute mode, 241
 RD Install mode, 241
 security, configuring, 235–240
 SSL and, 237–240
RD Session Host servers
 redirection for virtual desktops, 288–289
 as RemoteApp program source, 260
RD Session Host sessions, 232
RD Session Hosts/Terminal Servers, 301, 302
RD Virtualization Host, 284–299
 functionality, 285–287
 installing, 285
 personal virtual desktops, provisioning, 288–299
RD Web Access, 256–262
 functionality, 257–258
 installing, 258
 new features, 256
 public and private modes, 256
 using, 260–262
 Web Part, 257
RD Web Access servers
 configuring, 258–260
 function of, 289
RD Web Access Web site
 accessing, 258–259
 Configuration tab, 259, 261
 logon page, 259
 My Desktop item, 298
 RemoteApp Programs tab, 260
Real-Time Streaming Protocol (RTSP), 183
Real-Time Streaming Protocol Secure (RTSPS), 184
 recovery, 6
 regression testing
 reducing need for, 112
 virtual applications, 175
 regulatory compliance, 225

Relative Weight setting (Processor configuration), 86
Reliance Data Center, 446
RemoteApp, 121, 250–256
 functionality, 251
 managing Hyper-V with, 98
 support for, 120
RemoteApp and Desktop Connections, 262–272
 benefits of, 263
 client configuration, 267
 configuring, 264–273
 functionality, 264
 Start menu program group, 268
 Taskbar icon, 268–269
 Web feed, 269–272, 274–275
Remote Applications Installed Locally (RAIL), 121
RemoteApp Manager console, 244, 253
RemoteApp programs, 231
 adding, 252, 253
 configuring, 254–255
 displaying, 257
 filtering per user, 231, 256
 launching, 251, 260
 look and feel, 250
 packaging, 254
 Start menu access, 263
 terminating, 252
 user access, 255
 Web browser access, configuring, 264–273
RemoteApp Wizard, 98
 remote control, managing, 244
Remote Desktop Connection Authorization Policy (RD CAP), 279
Remote Desktop Connection Broker, 272–275
Remote Desktop Connection Manager console
 RD Connection Broker settings, 274
 virtual desktop assignment, 296
Remote Desktop Connection (RDC), 6, 97–98, 229
Remote Desktop Gateway. *See* **RD Gateway**
Remote Desktop Licensing. *See* **RD Licensing**
Remote Desktop Resource Authorization Policy (RD RAP), 279
 remote desktops. *See also* **virtual desktops**
 accessing, 257
Remote Desktops console, 244

Remote Desktop Services, 8, 10, 11, 109, 121, 223, 228–229
 availability of, 227
 benefits of, 224
 centralized management, 14
 command-line tools, 245–246
 deploying, 299–300
 for office worker scenario
 optimization, 13
 for task worker scenario
 optimization, 13
 functionality, 228
 interactions between role
 services, 300
 licensing, 281, 284
 RD Connection Broker, 272–276
 RD Gateway, 276–281
 RD Licensing, 281–284
 RD Session Host, 231–256
 RD Virtualization Host, 284–299
 RD Web Access, 256–262
 RemoteApps, 13
 resources on, 309–310
 security, 14
 usage scenarios, 224–226
 WMI provider for, 291

Remote Desktop Services
 Configuration console, 245

Remote Desktop Services Manager
 console, 243

Remote Desktop Services user
 profile, 244

Remote Desktop Session Host.
See RD Session Host

Remote Desktop Session Host
 Configuration console, 239, 243
 Licensing tab, 283

Remote Desktop sessions,
 encrypting, 238–239

Remote Desktop Users group, 98

remote desktop virtualization, 11,
 109, 223. *See also* Remote
 Desktop Services
 availability of, 227–228
 benefits of, 224
 usage scenarios, 224–227

Remote Desktop Virtualization Host.
See RD Virtualization Host

Remote Desktop Web Access.
See RD Web Access

Remote Desktop Web Access
 Computers group, 98

Remote Desktop Web Connection,
 261, 262

Remote Server Administration Tools
 for Windows 7, 106

remote sessions
 authentication and authorization
 in background, 277

system messages, 278

Remote tab (managed host
 properties), 371

remote users, internal network
 access, 280

remoting, 99

RemoveLibraryStoreSharePath
 parameter, 335

Rename option (Hyper-V Manager
 console), 69

Rename option (Snapshots pane), 83

Reporting view (VMM
 Administrator Console), 360

reports on App-V environment,
 199–200

Reserves tab (managed host
 properties), 368

RetainDeletedObjects parameter,
 335

RetainObjectCache parameter, 335

return on investment. *See* ROI
 (return on investment)

ring architecture, 49–50

roaming desktop technologies,
 12–13

roaming users, 12–19, 110, 232

ROI (return on investment)
 of Microsoft's virtualization
 solution, 16
 on virtualization, 6

Run An Older Operating System,
 Such As Windows NT setting
 (Processor configuration), 86

S

Sanbolic Clustered File System, 334

schedulers for server virtualization,
 22

SCONFIG tool, 57

ScriptCmdlets, 99

script internationalization, 100

SCSI controllers
 adding to VMs, 67
 using, 85

SCSI Controller setting (Hyper-V
 Manager Action pane), 67

second-level address translation
 (SLAT), 7

Secure Sockets Layer (SSL), 237–240

security
 App-V features, 165
 BitLocker Drive Encryption, 12
 Data Execution Prevention, 49
 device driver considerations, 26
 ensuring, 1
 of Hyper-V server, 107
 isolation of virtual machines, 132
 microkernel hypervisors and, 27
 monolithic hypervisors and, 25
 proactive approach, 4
 RD Gateway and, 276–277
 for RD Session Host, 235–240
 virtual image expiration, 157
 VM worker processes and, 31
 security groups, 275–276
 Seldam, Matthijs ten, 108

Select The Network Adapter To Be
 Used For Remote Desktop IP
 Virtualization policy setting,
 247

self-managing dynamic systems, 7

self-provisioning software, 4

Self-Service Portal. *See* VMM Self-
 Service Portal

Self-Service User roles (VMM), 341,
 413
 creating, 416–430

sequenced application packages,
 168, 170–171. *See also* virtual
 application packages

sequenced applications, 168. *See
 also* virtual applications

Sequencer. *See* App-V Sequencer

sequencing computers, 168

Sequencing Wizard, 202–203. *See
 also* App-V Sequencer

Add Files To Virtual File System
 page, 204

Advanced Options page, 202

Configure Applications page,
 205–206

Launch Applications page,
 206–207

Monitor Installation page, 203

Sequence Package page, 206

server availability, 8

Server Configuration Utility
 (SConfig), 47

server consolidation, 7, 9, 21
 accelerating, 6
 Hyper-V for, 43
 VMM and, 338, 340

Server Core installations, 103

Hyper-V role management and,
 57

ServerManagerCmd.exe command-
 line tool, 54

ServerManager.msc, 232

Server Message Block (SMB), 184

server provisioning, 6

server roles, adding, 48

servers
 antivirus software on, 56
 consolidating with App-V for RDS,
 226
 CPU resource allocation, 85

servers (*continued*)

- hosting Hyper-V role, 55–57. *See also* Hyper-V servers
- overloading, avoiding, 55
- remote desktop virtualization on, 109
- Server Core Hyper-V servers, 103
- utilization rates of, 434

Server settings (Hyper-V Manager), 59

server virtualization, 44–46

- benefits of, 44–45
- design of, 53
- management interfaces for, 22
- memory management for, 22
- networking resources for, 22
- requirements of, 22
- scheduler for, 22
- state machine for, 22
- storage for, 22
- understanding, 21–27
- virtual device drivers for, 22
- virtualized devices for, 22

service levels, 8

Service Model (Windows Azure), 445

services, faster and cheaper, 1

Settings action (Hyper-V Manager console), 66

.sft files, 169, 170, 203

- streaming, 171

SFTMIME command, 216

Shah, Vipul, 45

shared ISO image files, 337

shared storage, 334

shared volumes, 89

shutdown, host-initiated, 135

Shut Down option (Hyper-V Manager console), 79

shutdown requests, 35

single sign-on, 256

single virtual processors, 38

SLAT (second-level address translation), 38

smart cards integration component (virtual machines), 136

SMP (symmetric multiprocessor), 36

Snapshot File Location setting (Hyper-V Manager console), 68

Snapshot Name dialog box, 83

Snapshot option (Hyper-V Manager console), 82

snapshots, 81–84

- applying state from, 83
- configuration options, 84
- creating, 30
- deleting, 83
- Hyper-V support for, 36
- location of, 68

- managing with Windows PowerShell, 101
- naming conventions, 82–83
- in production environments, 82
- renaming, 83
- reverting to previous, 84
- snapshot trees, 81, 83, 84
- taking, 69, 84
- of virtual machines, 6

Snapshot setting (virtual machines), 69

Snapshots folder, 80, 83

Snapshots pane, 82

SoftGrid Application Virtualization, 162

software

- inventorying, 4
- self-provisioning, 4

Software as a Service (SaaS), 435

software audits (App-V), 200

software deployment

- controlling, 3
- flexibility in, 8
- manual, 3
- partially automated, 3
- reducing time for, 8
- zero-touch, 4

space requirements, 5, 7

spending decisions, 1

.sprj files, 171

SQL Server

- in cloud-computing solutions, 437
- communication ports, 319
- as MED-V report database, 152
- virtualizing, 45
- with VMM, 328, 345
- VMM database, 351

SSL certificates, 279

SSL for Remote Desktop Services connections, 237–240

Standalone virtualization mode, 114

Standardized IT infrastructures, 3

- description of, 3
- moving to Rationalized, 6

StarUK, 446

state machines for server virtualization, 22

Status tab (managed host properties), 366, 367

storage

- disk-based, 428–429
- high-speed access to, 55
- hot adding and removal, 37
- for server virtualization, 22
- of system files, 56

storage area networks (SANs), 333

storage arrays, 89

stored virtual machines, 314

streaming servers, 169

Summary tab (managed host properties), 365

symmetric multiprocessors (SMPs), 36

synthetic devices, 32–33

synthetic network adapters, 77

system/boot volumes, 77

System Center Capacity Planner, 427

System Center Configuration Manager, 171, 426, 427

System Center Data Protection Manager, 426–429

System Center Essentials, 426

System Center Mobile Device Manager, 427

System Center Operations Manager, 426, 427

- Virtualization Candidates report, 400
- VMM and, 325–327

System Center Server Management Suite Enterprise, 426

System Center Service Manager, 427

System Center solutions, 9, 425–429

- benefits of, 427
- in cloud-computing solutions, 438

System Center Virtual Machine Manager, 9, 47, 85, 426, 427

System Center Virtual Machine Manager Team Blog, 430

system error reports (App-V), 200

SystemGuard environment, 174, 175

system utilization reports (App-V), 200

T

tape-based backup, 428

task worker scenarios

- optimizing, 13–14
- Remote Desktop Services and, 225–226

TCO (total cost of ownership) of Microsoft's virtualization solution, 16

- reducing, 5–7
- of servers, 43

TCP Chimney Offload, 39

TCP ports

- for SQL Server communications, 328
- for VMM communication, 318–319

templates, 314

Terminal Server Session Directory service, 272

terminal servers, load balancing, 272

Terminal Services, 112, 229, 250. *See also* Remote Desktop Services

- App-V client for, 181
- App-V support for, 164

Terminal Services Client, 169

Terminal Services Gateway, 276

Terminal Services Licensing, 281

Terminal Services RemoteApp, 251.
See also RemoteApp

Terminal Services Session Broker, 272

Terminal Services Web Access, 256

test hardware, 43

testing

- Hyper-V and, 43
- snapshots and, 81

time synchronization, 35, 87, 135

Time Synchronization service, 68

total cost of ownership. *See* TCO (total cost of ownership)

training costs, 7

Transport Layer Security (TLS), 237–238

troubleshooting

- App-V, 217–219
- COM port output, 132
- Hyper-V startup problems, 54
- license compliance, 173
- virtual applications, 149
- virtual machine desktops, 150

Turn Off Windows Installer RDS

- Compatibility policy setting, 247

Turn On Remote Desktop IP

- Virtualization policy setting, 247

Type 1 hypervisors, 23–24

- Hyper-V, 28. *See also* Hyper-V

Type 2 hypervisors, 24

- performance of, 24

U

undo disks settings (virtual machines), 131

usage policies

- configuring, 156–157
- enforcement of, 148, 149, 153

usage scenarios

- for App-V, 113–114
- for MED-V, 113
- for Windows Virtual PC, 113
- for Windows XP Mode, 113

USB Connector (vpcusb.sys) (Virtual PC Host), 118

USB devices

- configuring, 141
- redirecting, 139, 140
- sharing, 142

USB Stub Driver (vpcuxd.sys) (Virtual PC Host), 118

UseCluster parameter, 335

UseLocalVirtualHardDisks parameter, 335, 336

Use On The Virtual Machine option (Hyper-V Manager console), 74

user data and settings

- migrating from previous versions, 13
- virtualizing, 110

user role processing, 337–338

user roles, configuring with VMM, 413–420

users

- collaboration between, 4
- moving, 13

User settings (Hyper-V Manager console), 60–61

- resetting to defaults, 61

User State Migration Tool (USMT), 13

user state virtualization, 110

V

V2V (virtual-to-virtual) conversions, 315, 411–412

- operating systems supported for, 412

Validate A Configuration Wizard, 89

vDevs (Virtual Devices), 29, 31–32

VDI (Virtual Desktop Infrastructure), 8, 10, 223, 304–307

- architecture of, 304
- availability of, 228
- benefits of, 224
- deploying, 306–307
- deployment scenarios, 303, 307
- desktop infrastructure, centralizing with, 303
- functionality of, 307–309
- licensing, 306
- Premium suite, 306
- RD Virtualization Host component, 285
- resources on, 311
- Standard suite, 305
- usage scenarios, 227

Veritas Storage Foundation for Windows, 334

.vfd (virtual floppy disk) files, 67

- creating new, 59

.vhd (virtual hard disk) files

- default location, 59
- managing with Windows PowerShell, 102

video resize, guest, 135

Vid.sys component, 32

View menu (Virtual Machine Connection tool), 72

Virtmgmt.msc. *See* Hyper-V Manager console

virtual adapters, 85–86

virtual application packages

- configuring, 207–211
- creating, 201
- delivery methods, 186
- deploying, 186–188
- managing, 197
- publishing, 176
- saving, 210
- sequenced, 168, 170–171
- streaming, 171–172, 178

virtual applications, 119–120

- access permissions, 196
- access to, 193
- adding, 193
- application data, streaming, 177
- block size, 203
- centralized management of, 14
- closing, 146
- compatibility of, 173
- default save locations, 146, 155
- definition of, 169
- delivery, 112, 185–186
- deploying, 111, 173, 301
- development, 111
- file type associations, 195, 200, 215
- importing, 192, 193
- installing, 121, 142–147, 203–204
- installing on RD Session Host server, 240–242
- installing per user, 232
- launching, 120, 143, 144
- license management, 198–199
- life cycle management, 111
- look and feel of, 109
- managing, 172–176, 192–197, 214–215
- offline availability, 163
- packaging, 170–171
- provider policies, 201
- publishing, 121, 168, 171
- Q: drive, installing on, 211–213
- regression testing, 175
- sequenced applications, 168
- shortcuts, 195
- supporting, 120–121, 173, 175
- terminating, 173, 176
- updating, 173, 175
- usage data, 164
- utilization reports, 200
- versioning, 111
- virtual environments for, 174
- virtualizing, 8, 11

virtual COM, 169

- Virtual Desktop Infrastructure.
 - See VDI (Virtual Desktop Infrastructure)
- virtual desktops
 - accessing, 308
 - centralizing execution of, 11
 - connecting to, 297–298
 - personal, 285, 286
 - pools of, 285–286
 - provisioning, 227, 288–299
 - redirection for, 288–289
 - sharing, 227
- virtual device drivers, 22
- Virtual Devices. *See* VDevs (Virtual Devices)
- virtual directories, 169
- virtual environments, 169–170, 174
- virtual files, 170
- virtual file systems, 170, 204
- virtual hard disks
 - on cluster shared volumes, 88
 - compacting, 64, 85
 - configuring, 130
 - converting, 64
 - creating new, 59
 - default location of, 84
 - editing, 64–65
 - expanding, 64
 - fixed size, 85
 - inspecting, 64
 - undo disks, 131
 - in virtual machine templates, 379
 - VM_name.vhd, 122
 - Windows XP Mode, 121. *See also* Windows XP Mode
 - Windows XP Mode base.vhd, 122
- Virtual Hard Disks folder, 80
- virtual images, 151
 - centralized management of, 147
 - creating, 158
 - delivering, 147, 148, 158
 - expiration of, 157
 - maintaining, 157, 159
 - managing, 157–159
 - RAM for, 153
 - repository for, 147, 148
 - retrieving, 157
 - updating, 158
- virtualization. *See also* Hyper-V
 - barriers to, 17
 - benefits of, 5, 7, 427–428
 - developments in, 19
 - efficiency of, 434
 - hardware-assisted, 36, 49
 - Infrastructure Optimization Model and, 4–7
 - integrated Microsoft solution, 15–17
 - need for, 1
 - technical details resource, 19
- Virtualization Candidates report, 400
- Virtualization Infrastructure Driver (Vid.sys), 29, 32
- virtualization infrastructures
 - full, 113
 - lightweight, 114
 - Standalone mode, 114
- Virtualization Server Provider (VSP) (Virtual PC Host), 117
- Virtualization Service Clients (VSCs), 34, 118
- Virtualization Service Providers (VSPs), 29, 32
- virtualization service status,
 - managed hosts, 367
- virtualization service version,
 - managed hosts, 367
- virtualization stack, 29
- Virtualization TechCenter, 104
- virtualization technologies, 16
- virtualized computing resources, 431. *See also* cloud computing
- virtualized devices, 22
- virtualized environments, visible
 - and invisible, 109
- virtualized resources
 - managing with VMM
 - Administrator Console, 356
 - provisioning with VMM, 340–341
- virtualized storage resources, 6
- Virtual Machine Additions, 34
- Virtual Machine Bus (VMBus), 29, 32–33
- Virtual Machine Connection tool, 66, 70–74
 - authentication issues, 73
 - connecting to virtual machines, 73
 - installing, 72–73
 - keyboard accelerators, 74
 - launching, 73
 - menu options, 71–72
 - remote use, 73
 - snapshot options, 83
 - title bar, 72
 - toolbar, 72
- Virtual Machine Device Queues (VMDq), 39
- virtual machine hosts, 315
- Virtual Machine Limit setting (Processor configuration), 86
- Virtual Machine Management Service. *See* VMMS (Virtual Machine Management Service)
- Virtual Machine Manager. *See* VMM (Virtual Machine Manager)
- Virtual Machine Manager
 - Configuration Analyzer (VMMCA), 348–349
- Virtual Machine Queue (VMQ), 39
- Virtual Machine Reserve setting (Processor configuration), 86
- virtual machines
 - adding and removing, 101
 - auto publish settings, 120, 138
 - cloning with VMM, 400
 - close settings, 138
 - clustering, 88. *See also* clustered virtual machines
 - components library, 339
 - COM port settings, 132
 - concurrently running, 38
 - configuration best practices, 84–85
 - configuration files, 77–79
 - configuring, 128–138
 - configuring with VMM
 - Administrator Console, 394
 - connecting to, 66–67, 73, 101, 395–398
 - connecting to running machines, 422–423
 - connection credentials, 60
 - console windows
 - converting physical computers to, 338, 400–411
 - CPU resource allocation for, 85
 - creating, 58, 76–79, 128, 384–392, 424
 - delegated management of, 339
 - deploying with VMM, 399
 - display interface settings, 85
 - domain controllers, 87
 - dual virtual processor, 38
 - DVD drive settings, 131
 - exporting, 79–81
 - fast provisioning, 339
 - finding, 101
 - full-screen mode, 109
 - guest operating systems on, 50–52. *See also* guest operating systems
 - GUIDs of, 31
 - hard disk settings, 130
 - hardware, adding, 67
 - hardware settings, 67–68
 - hibernation, 144
 - hot adding and removal of storage, 37
 - importing, 59, 79–81
 - integration features, 133–134
 - Integration Services use, 56
 - Intelligent Placement process, 339
 - isolated execution environments for, 23

virtual machines (*continued*)

- jumbo frames support, 39
- keyboard accelerators for, 74
- keyboard settings, 137–138
- legacy applications, running in, 110, 111
- local management of, 110
- logon credentials, 138
- management settings, 68–69
- managing, 31, 65, 148–149, 384–400
- memory settings, 130
- migrating between host machines, 37
- migrating with VMM, 337, 399
- mouse control in, 97
- naming, 68, 130
- network adapter configuration, 70
- networking settings, 132
- output settings, 132
- ownership of, 93
- passthrough disks for, 87. *See also* passthrough disks
- performance, 85
- processor compatibility mode and, 95–97
- Processor configuration settings, 86
- properties dialog box, 394
- provisioning, 336
- quad virtual processor, 38
- remote management of, 103
- renaming, 69
- security resources, 107
- shared ISO image files, 337
- shutting down, 79
- single virtual processor, 38
- snapshot management, 101
- snapshots of, 6, 81–84. *See also* snapshots
- start actions, 68
- stop actions, 69
- system/boot files for, 77
- templates, creating from, 385–392
- understanding, 21–23
- undo disk settings, 131
- USB devices, connecting, 139–141
- USB devices, sharing, 142
- V2V conversions, 411–412
- viewing, 128, 129
- virtual networking settings, 61
- Windows XP Mode, 121
- Virtual Machines folder, 80
- virtual machine states, 22
 - applying, 83
 - discovering and manipulating, cmdlet for, 101
 - managing, 30
- Virtual Machines view (VMM Administrator Console), 359
- virtual machine templates
 - adding, 379–380
 - virtual machines, creating from, 385–392
- Virtual Machine User Services (VMUSrv.exe), 118
- Virtual Machine Viewer, 396
- Virtualmachineviewer.exe, 323, 395
 - command-line options, 324–325
- Virtual Machine worker processes, 29, 31
 - online operations management, 30
 - spawning, 30
- virtual mages, 158
- Virtual Motherboard (VMB), 31
- virtual network adapters, 132, 133
- Virtual Network Manager, 61
- virtual networks
 - connectivity allowed by, 63
 - types of, 61
- Virtual PC
 - exporting virtual machines from, 80
 - Type 2 hypervisor, 24
- Virtual PC images. *See* virtual images
- Virtual PC Integration Components Services Application service (VMSrv.exe), 118
- virtual processors
 - allocating to VMs, 67
 - types of, 38
- virtual registries, 170
- Virtual Server 2005 R2 SP1 platform, 313
- Virtual Server hosts, 397
- virtual servers
 - availability of, 10
 - moving between hosts, 10
 - provisioning, 433
 - utilization rates of, 434
- virtual services, 170
- virtual switches, 337
- Virtual Switch Protocol, 56
- virtual-to-virtual conversions. *See* V2V (virtual-to-virtual) conversions
- VM Additions, 85
- VMB (Virtual Motherboard), 31
- .vmc (virtual machine configuration) files, 59
- VM Chimney, 39
- VMConnect. *See* Virtual Machine Connection tool
- .vmcx files, 128
- Vmmadmin.exe, 323
- VMM Administrator Console, 316, 322–324
- Actions pane, 362
- configuring virtual machines, 394
- connecting to virtual machines, 395–398
- Connect To Server dialog box, 356
- filters, 360–361
- host management, 361–377
- Hosts section, 361
- installing, 354
- Library view, 379
- managing virtual machines with, 392–394
- supported operating systems, 347–348
- updates to, 337
- user interface, 357–358
- using, 356–360
- views, 358–360
- Virtual Machine Manager section, 361
- Virtual Machines view, 360
- VMM Agents, 316, 329
 - installing, 364–365
- VMM database, 345, 352
 - backing up and restoring, 382–383
- VMM library, 329, 330, 352, 378–383
 - connecting to virtual machines, 396–398
 - deploying virtual machines from, 399
 - guest operating system profiles, adding, 380–381
 - hardware profiles, adding, 380
 - resource dependencies, 381
 - resources, removing, 381
 - virtual machine templates, adding, 379–380
- VMM Library Servers, 315, 330, 378
 - adding, 378–379
 - communication ports, 319
 - installing, 352
 - requirements for, 378–379
- VMM Self-Service Portal, 316, 325, 339, 341, 421–425
 - installing, 355–356
 - logging on, 422
 - setting up, 421
 - supported operating systems, 347–348
- Thumbnail View, 424
- VMM Server, 315
 - communication ports, 319
 - connecting to, 357
 - domain account for, 354
 - installing, 350–354
 - port assignments, 353, 354

VMM Server (*continued*)

- Self-Service Portal installation, 355
- VMM Administrator Console installation, 355
- VMM Server Library, 347–348
- VMMS.exe, 30
- VMMS (Virtual Machine Management Service), 29, 30, 328
 - RPC dependencies, 30
 - user mode and kernel mode, 30
 - WMI dependencies, 30
- VMM.sys (Virtual PC Host), 117
- VMM (Virtual Machine Manager), 10, 313
 - architecture, 316–330
 - benefits of, 338–339
 - Client Layer, 320–321
 - cloning virtual machines, 400
 - communication ports, 318–319, 353
 - components of, 315–316
 - creating new virtual machines, 384–392
 - delegated administration support, 332–333
 - Engine Layer, 317, 320, 328–329
 - failover cluster support, 331
 - guest operating system control, 49–50
 - hardware requirements, 343–344
 - host clusters, 374
 - Hyper-V management, 331
 - infrastructure requirements, 348
 - installing components, 348–356
 - Intelligent Placement process, 339
 - interlayer communications, 317–318
 - latest improvements, 330–333
 - library, 339
 - maintenance mode feature, 335
 - Managed Computer Layer, 329–330
 - migrating virtual machines, 399
 - new features, 333–337
 - P2V conversions, 400–411
 - Performance and Resource Optimization, 326–327
 - planning for deployment, 342–343
 - PRO capabilities, 332, 338, 341–342, 374
 - resources on, 430
 - SAN transfers, 333
 - Self-Service Portal. *See* VMM Self-Service Portal
 - for server consolidation, 338, 340
 - shared storage, 334
 - software requirements, 344–348

- SQL Server database for, 351
- System Center Operations Manager and, 325–327
- system requirements, 343–348
- terminology, 313–315
- usage scenarios, 340–341
- user role configuration, 413–420
- user role processing, 337–338
- V2V conversions, 411–412
- virtual machine management capabilities, 384–400
- VMware ESX Server host management, 331
- VMware Infrastructure 3 environment management, 374–377
- Windows PowerShell and, 320–322, 334–336, 339
- VM_name.vhd, 122
- VMNetworkOptimizationEnabled parameter, 335
- Vmractivexclient.dll, 323
- Vmrc.exe, 324
- VMSAL.exe, 117, 119
- VMSrv.exe service, 118
- VMs tab (managed host properties), 368
- VMUSrv.exe service, 118
- VMware ESX hosts, 319
- VMware ESX Server monolithic hypervisor, 25
 - type 1 hypervisors, 24
 - VMM support for, 338
- VMware Infrastructure 3 environment, 374–377
- VMWarePortGroup parameter, 335
- VMware port groups, 337
- VMware Server, 24
- VMware VI3 environment, 313
 - VMM host management, 331
- VMwareviewer.exe, 325
- VMware Virtual Center servers adding, 375
- communication ports, 319
- VMWindow.exe, 116
- vmwp.exe, 31
 - viewing instances of, 31
- VMX/SVM Root Kernel, 118
- Voellm, Tony, 108
- Volume Shadow Copy Service (VSS), 35
 - App-V support for, 164
- VPCBus.sys, 117
- VPC.exe, 117
- VPC Settings dialog box (Virtual PC Host), 116
- VPC Wizard (Virtual PC Host), 116
- .vsv (saved state) files, 83

W

- Watson, 164
- Watson, Clive, 108
- .wcx (Workspace Configuration) file, 274
- Wilson, Mark, 108
- Windows 7
 - as guest operating system, 51
 - host operating system for Windows Virtual PC, 123
 - migrating to, 110
 - P2V conversion support, 401
 - RemoteApp support, 120
 - V2V conversions support, 412
 - VDI and, 305
 - VMM support, 348
 - Windows Virtual PC availability, 114
 - Windows XP Mode and, 114, 125
- Windows 2000, as guest operating system, 51
- Windows 2000 Server, 402
- Windows Aero, 230
- Windows Azure, 435, 444–445
 - availability of, 446
 - resources on, 447
- Windows Communication Foundation (WCF), 317
- Windows Fundamentals for Legacy PCs, 14
- Windows Hardware Developer Central Web site, 19
- Windows Hypervisor Interface Library (WinHv.sys), 29, 32
- Windows Installer, 232
- Windows keyboard accelerators, 74
- Windows Management Instrumentation (WMI)
 - managing Hyper-V with, 98–99
 - RD Session Host management with, 249
- Windows Network Load Balancing (NLB) service, 36
- Windows operating systems
 - early versions, 98
 - RemoteApp support, 120
 - running with Med-V, 112
- Windows Optimized Desktop Scenarios, 12–14
 - anywhere-access scenario optimization, 14
 - contract/offshore worker scenario optimization, 14
 - mobile worker scenario, 12–13
 - office worker scenario optimization, 13
 - resources on, 19
 - task worker scenario optimization, 13–14

- Windows PowerShell
 - in cloud-computing solutions, 442
 - cmdlets for managing Hyper-V, 101–103
 - features and enhancements, 99
 - for live migrations, 94
 - management library for Hyper-V, 100
 - managing Hyper-V with, 99–104
 - new parameters for cmdlets, 334–335
 - PRO cmdlets, 327
 - RD Session Host management with, 249–250
 - security contexts, 322–323
 - VMM cmdlets, 334, 336
 - VMM command-line interface, 321–322
 - VMM, integration with, 320–322, 334–335, 339
- Windows PowerShell Virtual Machine Manager console, 355
- Windows Remote Management, 317
- Windows roaming, 12–13
- Windows Server 2003
 - as guest operating system, 35, 50, 52
 - with Hyper-V, 57
 - legacy network adapter driver, 77
 - VMM support, 347, 348
- Windows Server 2008
 - in cloud-computing solutions, 438
 - Group Policy settings, 246–247
 - P2V conversion support, 401
 - RD Session Host role service, 231–256
 - Remote Desktop Services role services, 229
 - Terminal Services, 229, 250
 - V2V conversion support, 412
 - VMM support, 346, 347
- Windows Server 2008 R2, 16
 - as guest operating system, 34, 35, 50
 - Datacenter x64 with Hyper-V, 55
 - Failover Clustering feature, 37, 43
 - Hyper-V role, 8, 9, 21, 27, 28, 44, 47–48
 - Itanium versions, 49
 - parent partition for, 28
 - release of, 16
 - Remote Desktop Services, 8
 - Server Core installation option, 103
- Windows Server Update Services (WSUS), 158
- Windows Server Virtualization Guide, 53, 54, 105
- Windows Speech Recognition, 230
- Windows VECD for SA licensing, 305
- Windows VECD licensing, 305
- Windows VECD (Virtual Enterprise Centralized Desktop), 14, 304–305
- Windows Virtual PC, 8, 11, 17, 109, 115–147
 - Actions menu, 139
 - architecture of, 115–116
 - availability of, 114
 - benefits of, 110
 - console window, 139–140
 - CPU requirements, 124
 - guest operating systems, 123
 - guest-side components, 118
 - host operating system, 123
 - host-side components, 116–118
 - vs. Hyper-V, 119–120
 - installing, 124
 - keyboard shortcuts, 137
 - localized versions, 126
 - memory requirements, 124
 - requirements for, 123–124
 - resources on, 219–220
 - Settings dialog box, 142
 - Tools menu, 142
 - usage scenarios, 113
 - USB menu, 139
 - virtual application feature, 124
 - virtual machines, configuring, 129–138
 - Windows Security dialog box, 142
- Windows Virtual PC Host
 - kernel-mode engine components, 117
 - UI components, 116
 - user-mode engine components, 117
- Windows Vista
 - enlightenment of, 34
 - as guest operating system, 35, 51, 53
 - P2V conversion support, 401
 - TCP Chimney Offload support, 39
 - V2V conversion support, 412
 - VMM support, 348
 - Windows6.0-KB961741-x86.msu update, 120
- Windows XP, 121
 - as guest operating system, 35, 51–53
 - KB961742-v3.exe update, 120
 - legacy network adapter driver, 77
 - managing Hyper-V servers from, 98
 - P2V conversion support, 401
 - V2V conversion support, 412
 - VMM support, 348
- Windows XP Mode, 8, 11, 17, 109, 115–147
 - availability of, 114
 - benefits of, 110
 - vs. custom Windows XP virtual machines, 122
 - disk space requirements, 125
 - host operating system, 125–126
 - installing, 126–129
 - localized versions, 126
 - logon credentials, 138
 - requirements for, 125
 - resources on, 219–220
 - usage scenarios, 113
 - using, 138–147
 - virtual application support, 120
 - working with, 138–142
- Windows XP Mode base.vhd, 122
- WinFLP (Windows Fundamentals for Legacy PCs), 14
- witness disks, 89
- WMI cmdlets, 100
- WMI providers, 30
 - in parent partition, 28
 - for Remote Desktop Services, 291
- workloads
 - consolidating, 41
 - 32-bit and 64-bit, consolidating, 43
- workspaces (MED-V), 151
 - defining, 158
 - initializing, 154–155
 - using, 155–156

X

.xml (configuration) files, 83

Y

Yuen, Edwin, 412

About the Author



Mitch Tulloch, lead author for the *Windows 7 Resource Kit*, is a widely recognized expert on Windows administration, networking, and security and has been repeatedly awarded Most Valuable Professional (MVP) status by Microsoft for his outstanding contributions in supporting users who deploy Microsoft platforms, products, and solutions. Mitch has written or contributed to two dozen books including the *Windows Vista Resource Kit, Second Edition* and *Understanding Microsoft Virtualization Solutions: from the Desktop to the Datacenter*, both published by Microsoft Press.

Mitch has published hundreds of articles on sites such as WindowsNetworking.com and in leading industry magazines such as *BizTech Magazine* and *FedTech Magazine*. He has also developed e-learning courses on Windows 7 for Microsoft Learning and graduate-level courses on Information Security Management (ISM) for the Masters of Business Administration (MBA) program of Jones International University.

Mitch resides in Winnipeg, Canada, where he runs an IT content development business with his wife Ingrid. Prior to starting his business in 1998, Mitch worked as a Microsoft Certified Trainer (MCT) for Productivity Point International. For more information about Mitch, see <http://www.mtit.com>.

What do you think of this book?

We want to hear from you!

To participate in a brief online survey, please visit:

microsoft.com/learning/booksurvey

Tell us how well this book meets your needs—what works effectively, and what we can do better. Your feedback will help us continually improve our books and learning resources for you.

Thank you in advance for your input!

Microsoft
Press

Stay in touch!

To subscribe to the *Microsoft Press® Book Connection Newsletter*—for news on upcoming books, events, and special offers—please visit:

microsoft.com/learning/books/newsletter