

MasteringTM

System Center Operations Manager 2007

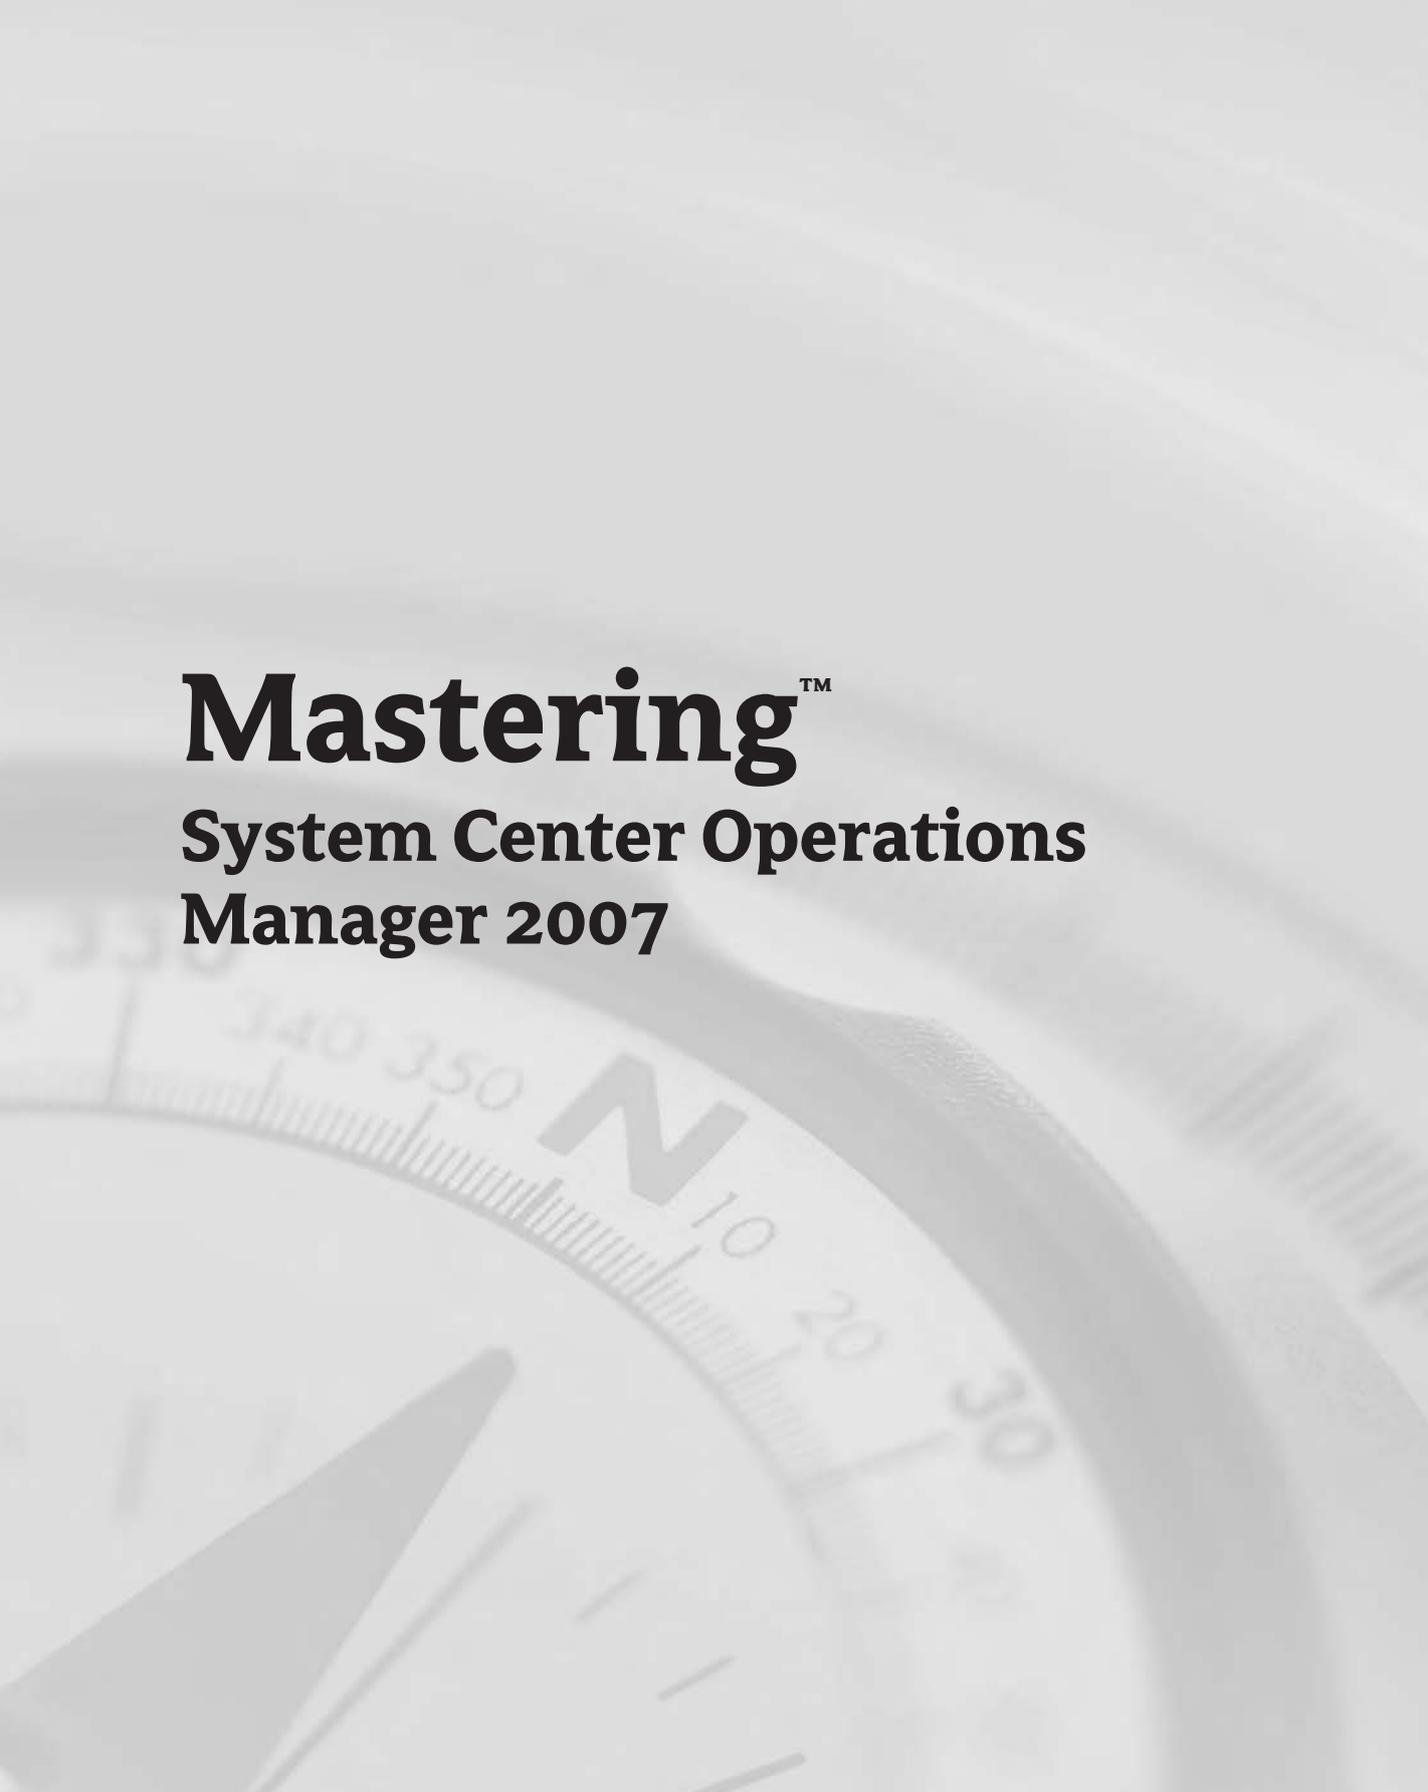
Brad Price

John Paul Mueller

Scott Fenstermacher



Wiley Publishing, Inc.



MasteringTM

**System Center Operations
Manager 2007**

MasteringTM

System Center Operations Manager 2007

Brad Price

John Paul Mueller

Scott Fenstermacher



Wiley Publishing, Inc.

Acquisitions Editor: Tom Cirtin
Development Editor: David Clark
Technical Editor: John Hann
Production Editors: Elizabeth Britten and Rachel Gunn
Copy Editor: Liz Welch
Production Manager: Tim Tate
Vice President and Executive Group Publisher: Richard Swadley
Vice President and Executive Publisher: Joseph B. Wikert
Vice President and Publisher: Neil Edde
Book Designer: Maureen Forys
Compositor: Craig Johnson, Happenstance Type-O-Rama
Proofreader: Candace English
Indexer: Ted Laux
Anniversary Logo Design: Richard Pacifico
Cover Designer: Ryan Sneed
Cover Image: © Pete Gardner / Digital Vision / GettyImages

Copyright © 2007 by Wiley Publishing, Inc., Indianapolis, Indiana

Published by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-0-470-11930-3

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

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Website is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Website may provide or recommendations it may make. Further, readers should be aware that Internet Websites listed in this work may have changed or disappeared between when this work was written and when it is read.

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

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

Library of Congress Cataloging-in-Publication Data is available from the publisher.

TRADEMARKS: Wiley, the Wiley logo, and the Sybex logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

10 9 8 7 6 5 4 3 2 1

Dear Reader

Thank you for choosing *Mastering System Center Operations Manager 2007*. This book is part of a family of premium-quality Sybex books, all written by outstanding authors who combine practical experience with a gift for teaching.

Sybex was founded in 1976. More than 30 years later, we're still committed to producing consistently exceptional books. With each of our titles we're working hard to set a new standard for the industry. From the paper we print on to the authors we work with, our goal is to bring you the best books available.

I hope you see all that reflected in these pages. I'd be very interested to hear your comments and get your feedback on how we're doing. Feel free to let me know what you think about this or any other Sybex book by sending me an email at nedde@wiley.com, or if you think you've found a technical error in this book, please visit <http://sybex.custhelp.com>. Customer feedback is critical to our efforts at Sybex.

Best regards,

A handwritten signature in black ink, appearing to read 'Neil Edde', written in a cursive style.

Neil Edde
Vice President & Publisher
Sybex, an Imprint of Wiley

In memory of Senior Airman Daniel Miller Jr., Staff Sergeant Kyle B. Wehrly, Private First Class Caleb A. Lufkin, Petty Officer First Class Gary T. Rovinski, Petty Officer First Class Jennifer A. Valdivia, Captain Joshua Steele, and all of the other men and women who have sacrificed so much in the line of duty.

—Brad

I dedicate this book to my devoted wife, Rebecca, who has been my best friend now for 27 years.

—John

Acknowledgments

As always, I have to start out by thanking my family. DeAnn—you have been there beside me for over 20 years. Each year brings new challenges and changes, and you are the constant that keeps my life sane. I love you and all that you do for me. Thank you for being my partner in life. Jami—you have made me proud as you weathered all of the changes in your life this year. As each year goes by, I marvel as I watch you grow and mature. And I promise I will put more pictures on the fridge. Becca—you have some major changes of your own as you graduate and make your way out into the world. I know you will make me proud. Keep hold of that sense of humor and your compassionate nature; they will take you far. To both of you smart, beautiful, considerate, and caring girls, I love you very much!

When it comes to parents, I was blessed with the best. And although you keep saying that I had to earn everything that I have because we did not have much as I grew up, you gave me the things that truly matter and cannot be bought—love, honesty, trust, and a great work ethic. Thank you for everything you did for me. And to the rest of my insane family—thank you for putting up with me. I know I am not always available to help out when needed, but I will try to make up for it in the near future. Of course I could always milk this for all it's worth. . . .

And once again, my brother John was there on the sidelines to help out. He wrote the first chapter and was my sounding board as I worked through this book. He is an invaluable asset professionally and personally.

As for my Sybex family, I have to start by thanking Tom Cirtin for all he has done for me. He is not only a great developmental editor; he has become a good friend. David Clark, Rachel Gunn, John Hahn, Liz Welch, Liz Britten, Candace English—thank you for all of your hard work getting this crazy book to print. I have had nothing but good experiences with the Sybex crew over the years, and each of you helped make this book a success.

—Brad Price

Thanks to my wife, Rebecca, for working with me to get this book completed. I really don't know what I would have done without her help in researching and compiling some of the information that appears in the chapters that I wrote. She also did a fine job of proofreading my rough draft and page-proofing the result.

Matt Wagner, my agent, deserves credit for helping me get the contract in the first place and taking care of all the details that most authors don't really consider. I always appreciate his assistance. It's good to know that someone wants to help.

Finally, I would like to thank Tom Cirtin, David Clark, John Hahn, Rachel Gunn, Liz Britten, Liz Welch, Candace English, and the rest of the editorial and production staff at Sybex for their assistance in bringing this book to print. It's always nice to work with such a great group of professionals and I very much appreciate the friendship we've built over the few years.

—John Paul Mueller

First, I would like to thank Brad Price for asking me once again to share the book-writing experience with him. I thought you had learned your lesson after the first book. This one was definitely a bigger challenge. Rest assured, one day I'll get even with you. And you can't write a book without the publisher and people like Tom Cirtin and David Clark that make it tick, so a big thanks goes out to the entire Wiley staff for their work on getting these words onto paper and into your hands. It's a task with details beyond my comprehension, so I'll just stick to the simple things like writing programs and figuring out Microsoft beta software.

Not said nearly often enough, special thanks go out to my wife Lori. Keeping the family going is a tough enough job, and allowing me the time to work on this project can't make it any easier. I'm sure jewelry is probably in order after this one. And I can't forget the young'uns, Jaina and Shelby. You might not have contributed words to the book or work around the house, but every day you manage to put a smile on my face without even trying. Working on projects like this means sometimes missing the little moments, but I know you have many big moments ahead and I wouldn't miss them for world.

Of course, I wouldn't be who I am today without the molding at the hands of my parents. To Carol, my mother—the words escape me sometimes, but always know that I am proud to call you Mom. To Craig, my father, I miss you. You taught me lessons about life that I'm still just realizing. Your time with us was too short, but you made every moment count. Happy hunting.

And last but not least, the friends who probably think I'm just naturally not right in the head (usually Linux is involved in those conversations), if for nothing else than simply putting up with me. To Jeff and Ken, for putting up with my sleep-deprived notions. To Michelle and Penny, for keeping me entertained with the soap opera. To "Fun Aunt Susan," for always being the voice of encouragement. To Phil, for the constant stream of jokes to my inbox. And finally to Darrin, Jay, Chris, Tim, Paul, Doug, and Mark, just for being a great bunch of guys to work with.

—*Scott Fenstermacher*

About the Authors

Brad Price has held several roles within the industry over the past 17 years. From programming barcode-scanning equipment and databases on midrange computers, to designing Active Directory, Exchange, SMS, and Operations Manager implementations, he has seen the IT industry grow and mature around him. He has taught classes at all levels of technology over the past 15 years and continually enjoys the interaction he has with students from around the globe. When not spending time working with new technologies, writing, and technical-editing, he enjoys spending time with his wife and two daughters, woodworking, listening to music, and snatching the occasional moment of sleep.

John Mueller is a freelance author and technical editor. He has writing in his blood, having produced 76 books and over 300 articles to date. The topics range from networking to artificial intelligence and from database management to heads-down programming. Some of his current books include a Windows command-line reference book, books on VBA and Visio 2007, and programmers' guides for web development with Visual Studio 2007 and Visual Web Developer. His technical editing skills have helped over 52 authors refine the content of their manuscripts. John has provided technical editing services to both *Data Based Advisor* and *Coast Compute* magazines. He's also contributed articles to magazines like *DevSource*, *InformIT*, *SQL Server Professional*, *Visual C++ Developer*, *Hard Core Visual Basic*, *asp.netPRO*, *Software Test and Performance*, and *Visual Basic Developer*. Be sure to read John's blog at <http://www.amazon.com/gp/blog/id/AQQA2QP4X1YWP>.

When John isn't working at the computer, you can find him in his workshop. He's an avid woodworker and candle maker. On any given afternoon, you can find him working at a lathe or putting the finishing touches on a bookcase. He also likes making glycerin soap and candles, which comes in handy for gift baskets. You can reach John on the Internet at JMue11er@mwt.net. John is also setting up a website at <http://www.mwt.net/~jmue11er/>. Feel free to look and make suggestions on how he can improve it. One of his current projects is creating book FAQ sheets that should help you find the book information you need much faster.

Scott Fenstermacher is currently a network engineer for a top 200 software company. He has a degree in computer science and holds several industry certifications from Microsoft and other vendors, including MSCE, MCSD and MCDBA.

Contents at a Glance

<i>Introduction</i>	<i>xxi</i>
Part 1 • Implementation and Administration	1
Chapter 1 • Overview of Operations Management	3
Chapter 2 • Installing System Center Operations Manager 2007	19
Chapter 3 • Management Group Settings	63
Chapter 4 • Installing and Configuring Agents	89
Chapter 5 • Managing Management Packs	109
Chapter 6 • Authoring and the Management Pack Life Cycle	155
Chapter 7 • Monitoring with the Console	173
Chapter 8 • Audit Collection Services and Agentless Exception Monitoring	205
Chapter 9 • Reporting	239
Chapter 10 • Working with Complex Environments	275
Part 2 • Maintenance and Troubleshooting	301
Chapter 11 • Optimizing Your Environment	303
Chapter 12 • Backup, Restore, and Disaster Recovery	331
Chapter 13 • Troubleshooting	345
Chapter 14 • Scripting Operations Manager 2007	361
<i>Index</i>	<i>411</i>

Contents

<i>Introduction</i>	<i>xxi</i>
---------------------------	------------

Part 1 • Implementation and Administration **1**

Chapter 1 • Overview of Operations Management **3**

Understanding IT Service Management	4
Exploring the IT Infrastructure Library (ITIL)	4
Exploring the Microsoft Operations Framework (MOF)	8
Tying It All Together	14
Defining Operations Management	16
Systems Management	16
Operations Management	16
The Bottom Line	16

Chapter 2 • Installing System Center Operations Manager 2007 **19**

Identifying Server Components	19
Operations Manager 2007 Server	20
DataStore (SQL Server 2005)	21
Verifying Prerequisites	22
Identifying Prerequisites	22
Running the Prerequisite Check Utility	26
Installation Criteria	28
Creating Accounts	28
Configuring Active Directory	30
Installing the Database	32
Installing Operations Manager 2007	39
Installing the Root Management Server on Clustered Hardware	42
Verifying the Installation	54
The Bottom Line	55

Chapter 3 • Management Group Settings **63**

Configuring Management Group Global Settings	63
Alerts	64
Database Grooming	66
Notification Channels	67
Privacy Settings	71
Reporting Settings	74
Web Addresses	75
Configuring Global Server Settings	76
Heartbeat	76
Security	76
Exploring Global Agent Settings	77
Understanding Individual Server Options	78
Agent Management	78

Heartbeat	84
Security	84
Proxy Settings	84
The Bottom Line	86
Chapter 4 • Installing and Configuring Agents	89
Exploring Agent Installation Prerequisites	89
Installing by Using the Computer and Device Management Wizard	90
Installing Manually	97
Installing from CD	97
Other Manual Installation Methods	99
Multihomed Agents	101
Adjusting Individual Agent Settings	102
Properties	103
Open	105
Change Primary Management Server	105
Repair	105
Uninstall	106
Delete	106
The Bottom Line	106
Chapter 5 • Managing Management Packs	109
Understanding Monitoring Capabilities	110
Identifying Management Pack Requirements	110
Current Management Packs Available	111
Requirements for Management Pack Monitoring	111
Exploring Management Packs	113
Discovery	114
Groups	116
Monitors	120
Using Rules	141
Reporting	148
Overrides	148
Converting Legacy Management Packs	149
Removing Management Packs	152
The Bottom Line	152
Chapter 6 • Authoring and the Management Pack Life Cycle	155
Authoring Management Packs	155
Creating New Management Packs	156
Exploring a Typical Life Cycle	156
Creating Management Pack Life-Cycle Controls	158
Release Management	158
Versioning	158
Management Packs That You Created	159
Management Packs Provided by Microsoft	161
Documenting Changes	162
Using MP Studio 2007	162

Analysis and Documentation	163
Manage Overrides	163
Management Pack Store	164
Testing	164
Comparison	164
Management Pack Life Cycle	164
Management Pack Development	165
Management Pack Implementation	166
Management Pack Maintenance	167
Management Pack Considerations	169
The Bottom Line	169
Chapter 7 • Monitoring with the Console	173
Console Overview	173
Workspaces	174
Views	175
Tasks	182
Operators	186
Customizing the Console	190
Limiting the Data in the Details Pane	191
Creating Views As an Operator	196
Configuring Notification Channels	197
Maintenance Mode	198
The Bottom Line	201
Chapter 8 • Audit Collection Services and Agentless Exception Monitoring ..	205
Exploring Audit Collection Services	205
Audit Collection Services Components	206
Deployment Options	209
Optimizing the ACS Collector	219
Security	220
Performance Monitoring	224
Admin Tools	225
Uninstalling ACS	227
Using Agentless Exception Monitoring	228
Configuring Client Monitoring	229
Editing Group Policy	232
Agentless Exception Monitoring Views	235
The Bottom Line	235
Chapter 9 • Reporting	239
Installing Reporting Services	239
Installing SQL Reporting Services	240
Configuring SQL Report Server	245
Installing Operations Manager Reporting	253
Using the Reporting Workspace	258
Predefined Reports	260
Creating New Reports	262
The Bottom Line	273

Chapter 10 • Working with Complex Environments	275
Working with Multihoming Agents	275
Multihoming Overview	275
Configuring Management Groups	278
Agent Setup	279
Agent Settings	286
Configuring Connected Management Groups	286
Overview of Connected Management Groups	287
Configuring Connectors	287
Providing Notification Channels	289
Forwarding and Responding to Alerts	290
Interoperating with Third-Party Monitoring Solutions	293
Installing Connectors	293
Configuring Connectors	294
Editing and Deleting Subscriptions	299
The Bottom Line	299

Part 2 • Maintenance and Troubleshooting

Chapter 11 • Optimizing Your Environment	303
Optimizing Management Packs	303
Identifying Useful Settings	304
Customizing Management Packs	306
Optimizing Rules	309
Modifying Existing Rules	310
Creating New Rules	316
Self-Tuning Thresholds	319
Optimized Performance Counters	328
The Bottom Line	328
Chapter 12 • Backup, Restore, and Disaster Recovery	331
Backup Strategy	331
Backing Up Secure Storage	332
Backing Up the Registry	333
Backing Up the Operations Manager 2007 Database	335
Backing Up Management Packs	337
Restoring Your Environment	338
Performing the Initial Installation	339
Restoring the Management Packs	339
Restoring Secure Storage	339
Restoring the Registry	340
Restoring the Database	340
Disaster Recovery	342
Using OpsManager 2007 to Monitor Your Disaster Recovery Site	343
Replicating Operations Manager Data to Your Disaster Recovery Site	343
The Bottom Line	344

Chapter 13 • Troubleshooting	345
Operations Manager 2007 Troubleshooting Tools	345
Resource Kit Tools	346
Management Pack Catalogs	346
Third-Party Tools	348
Other Interesting Tools	349
Troubleshooting Agents	350
Making the Agent Available to the Server	350
Operations Manager Won't Remove an Agent Entry in the Pending Management Folder	352
Operations Manager Can't Push Content or Manage the Client Remotely	352
Operations Manager Can't Install the Agent Remotely	354
Agent Warnings and Errors in Event Viewer	354
Troubleshooting Management Packs	355
The Entity Health for One or More Management Servers Displays Critical	355
A Management Pack Won't Work After Conversion	355
Operations Manager Won't Let You Save a Recorded Web Monitor	356
Management Pack Warnings and Errors in Event Viewer	356
Troubleshooting Security Issues	356
Overcoming Issues with the W32Time Service	357
Understanding Why the Discovery Wizard Never Completes	358
Security Failure Audits, Warnings, and Errors in Event Viewer	359
The Bottom Line	360
Chapter 14 • Scripting Operations Manager 2007	361
Introducing PowerShell	361
Useful PowerShell Cmdlets	363
PowerShell and Operations Manager	373
Operations Manager Cmdlets	375
Computer and Device Cmdlets	376
Discovery Cmdlets	381
Management Pack Cmdlets	384
Rule Cmdlets	386
Task Cmdlets	389
Action Cmdlets	391
Notification Cmdlets	392
Alert Cmdlets	394
Miscellaneous Cmdlets	397
Management Server Cmdlets	400
Maintenance Window Cmdlets	403
Monitoring Cmdlets	405
Performance Counter Cmdlets	407
The Bottom Line	409
<i>Index</i>	411

Introduction

Welcome to the first book released on a very exciting topic. System Center Operations Manager 2007 is as much different from its predecessor, Microsoft Operations Manager 2005, as Windows Vista is from Windows 95. So much has changed and so much more has been added. This is the product that administrators have wanted since they first witnessed the monitoring capabilities of Microsoft Operations Manager 2000.

Due to all the changes and new features, you are probably looking for a resource that will help you navigate. And you are not wrong for wanting to have a little guidance as you move forward with the implementation of Operations Manager 2007. Just as with most of the new Microsoft products, this one has become very complex and difficult to implement without some insight. That is where this book comes in. You can read each of the chapters in order to glean some insight into how Operations Manager 2007 works and how it integrates and affects your current organization. You can also use it as a reference when you are not sure how something works or is supposed to work.

Who Should Read This Book?

The answer to this question: *everyone*. Well, maybe anyone who wants to monitor their network using Operations Manager 2007. Included between the covers of this book is a comprehensive look at deploying, managing, troubleshooting, and working with the new command set for PowerShell.

As complex as products are becoming, no one can be an expert on all of them. If you are like most administrators, you have time to learn only enough about a product so that you can manage it effectively. However, there is probably a lot more that you could be doing with any one product. This book is meant to get you up to speed quickly, and then help you through some of the more arcane topics.

Not every administrator will have the same type of infrastructure to work with. What works well in a large corporation does not always work for the small mom-and-pop companies. What works well for small companies may not scale well for large organizations. Microsoft has attempted to address the differences among companies and deliver a product that can be implemented quickly for a small company, yet will still scale well for large organizations. No matter which scenario fits you, you will want to learn how this product will work for you.

But most of all, any administrator who wants to try to get to the “proactive” side of managing their infrastructure should consider looking through these pages to see how they can start monitoring their systems effectively. Being on the other side—the “reactive” management side—means that you are constantly having to respond to emergencies and continually

“putting out fires.” If you are a reactive administrator, you probably already understand how hard it is to try to make your infrastructure more efficient when you don’t have enough time to work on anything else but emergencies.

The Contents at a Glance

There are two parts to this book. Part 1 covers implementation and administration. Part 2 covers maintenance and troubleshooting. As you read through each section, you will find that the material flows from one subject to another, building as you go. By the time you finish Part 1, you should have a good understanding of what goes into deploying and managing your management group. Throughout Part 2, you will learn some of the tips and tricks to keep your management group running smoothly.

Part 1: Implementation and Administration

Chapter 1: Overview of Operations Management This chapter explains what operations management is and how Microsoft has implemented it with Operations Manager 2007. The topics covered include Microsoft Operations Framework, Information Technology Infrastructure Library, and the role played by Operations Manager 2007.

Chapter 2: Installing System Center Operations Manager 2007 This chapter identifies the prerequisites and examines the installation options in detail.

Chapter 3: Management Group Settings The settings that control the management servers within the management group are discussed in this chapter. You will see how to manage the global settings applied to all management servers, and then how to override the settings on a server-by-server basis.

Chapter 4: Installing and Configuring Agents Each of the monitored systems needs to have an agent installed in order for the monitored system to send detailed information to the management server and to be monitored and managed. The options for installing and configuring the client are discussed in this chapter.

Chapter 5: Managing Management Packs Management packs define what is monitored on each of the managed systems. They also contain the rules for creating reports and tasks within the Operations Console. This chapter covers importing management packs and working with all of the objects contained within them.

Chapter 6: Authoring and the Management Pack Life Cycle Management packs are dynamic in nature. After you create a management pack, it will need to be fine-tuned and additional object modifications will need to be made. In this chapter you will be introduced to life cycle management for your management packs.

Chapter 7: Monitoring with the Console After the management group has been created and the management packs and agents deployed, you will need to start monitoring. The Operations Console is the primary tool for working with all of the objects for the management group. Here you will find out how to work with the Operations Console effectively.

Chapter 8: Audit Collection Services and Agentless Exception Monitoring In the past, auditing was always difficult because the audit logs were spread out all over your organization. Now with Audit Collection Services, you can centralize the log entries. Agentless Exception Monitoring allows you to take the Dr. Watson logs and control where they are distributed.

Chapter 9: Reporting Reporting Services allows you to create some impressive reports very quickly, thus having quick access to the inner workings of the management group. You can also control access to the reports, thus allowing other users to print out their own.

Chapter 10: Working with Complex Environments You may find yourself having to work with multiple management groups or allow your management group to interface with another management product. This chapter explains how connectors can be used and how you can create a hierarchy of management groups.

Part 2: Maintenance and Troubleshooting

Chapter 11: Optimizing Your Environment Not everyone can use the management packs right out of the box. You will need to test them within your organization to see how they work. This chapter will cover some of the options that you have to make your management group work a little more efficiently.

Chapter 12: Backup, Restore, and Disaster Recovery You do not want to start over from scratch, and you definitely don't want to lose all of the data that you have collected. This chapter covers your options for making sure that you have your management group protected in case a disaster strikes.

Chapter 13: Troubleshooting Even though we would like to think that everything works perfectly all of the time, we know that is not the case. Some troubleshooting tips are detailed in this chapter.

Chapter 14: Scripting Operations Manager 2007 PowerShell is a very powerful tool that will help you manage your management group. This chapter is meant to act as a primer on PowerShell as it is used in Operations Manager 2007.

Final Comments

Make sure you take the time to become familiar with Operations Manager 2007. The more comfortable you are with it, the more you will be able to do with it. Some of the chapters include a Master It! section at the very end. Each Master It! is a small lab that helps reinforce the topics in the chapters. Instructions have been included that allow you to create a virtual environment using Microsoft Virtual Server, which is a free download. Building a virtual test environment can come in handy when you are trying to work through a new topic or troubleshoot a problem.

Most of all, have fun as you are going through the topics contained herein. Once you find out how much power this product has in store for you, you will be amazed at some of the things you can do. Just looking at the surface, being able to monitor servers and services, may be impressive enough, but the additional features can be equally impressive, such as monitoring the health of an application from the time a user submits a request until the response is returned from your servers, reporting on the availability of servers and services, and creating scripts that will automatically alter the way the service functions when a problem arises.

Part I

Implementation and Administration

The first part of this book is going to present you with insight on how to install, configure, and manage System Center Operations Manager 2007. Microsoft had shipped earlier versions of Operations Manager, Microsoft Operations Manager 2000 and 2005, but Systems Center Operations Manager 2007 is a completely redesigned product. You will find only a few similarities between the products.

As you work through this first section you will read about topics that you may visit only once in the lifetime of your management group, and others that you will see on a day-to-day basis. With the complexity of this product, you may well find yourself returning to this part to assist you with duties you may not perform regularly, or you may want to see if there is another way of performing recurring tasks.

We open this part with an overview of Microsoft's MOF and DSI initiatives. These are the backbone of System Center Operations Manager. Once we have introduced you to those concepts, we move into installing the management servers and creating the management group. After the management group is created, you will need to fine-tune it by creating the objects that are used to monitor servers and services within the management group. All of this will be presented to you in an orderly fashion so you can follow along and build your environment to meet your needs.

Chapter 1

Overview of Operations Management

Before we delve into the System Center Operations Manager product, we must define what operations management is, what it defines, and why we need it. Every day we talk to businesses about the Microsoft System Center line of products. Only a small percentage of those clients actually understand the concept and reasoning behind proactive IT service management.

In a meeting with a client, we were approached by an IT manager and asked, “Why should I care about the business? I just run the IT department.” IT service management is very much ingrained in the business as a whole. IT service management is defined as a way to organize IT services and provide constant improvement to the quality of those services. This is what caught the manager’s attention: IT service management should be tailored to delivering and supporting IT services that directly affect the organization’s business requirements.

He was correct in a way. As an IT manager, you are not responsible for certain key business activities. When those activities are being processed on your servers, however, you become a critical piece of the puzzle in overall IT systems management. You may control the SQL servers, but they house information that is critical to day-to-day operation of the billing department, for example. Suddenly, you start to see how everything ties together. A missing or damaged link in the chain or an unplanned removal of the chain may cause much more damage than you originally thought.

This is just one of the many reasons Microsoft created the Microsoft Operations Framework (MOF), based on the IT Infrastructure Library (ITIL). The idea behind MOF and ITIL is to create a complete team structure with the ultimate goal of service excellence. Numerous groups fall under the IT Department tag, but we often see many of them acting as separate departments rather than as one cohesive unit. Desktop support, application developers, server support, storage administrators, and so forth are all members of IT, but they are not always as tight as they should be.

Operations Manager is so much more than just a centralized console view of the events and processes in your network. Operations Manager was built with MOF in mind. We would like to start the book with a background of both MOF and ITIL.

In this chapter you will:

- ◆ Define IT service management
- ◆ Learn how ITIL is the foundation of IT service management
- ◆ Understand how MOF expands ITIL
- ◆ Explore the Dynamic Systems Initiative
- ◆ Learn about the Microsoft System Center products
- ◆ Learn how to use System Center Operations Manager

Understanding IT Service Management

ITIL and MOF were introduced as a way to deliver consistent IT Service Management (ITSM). Some of the key objectives of ITSM are:

- ◆ To align IT services with current and future needs of the business and its customers
- ◆ To improve the quality of IT services delivered
- ◆ To reduce the long-term cost of service provisioning

Think of ITSM as a conduit between the business and the technology that helps run the business. Without a proper conduit in place, one cannot function properly without the other.

Exploring the IT Infrastructure Library (ITIL)

Before we dig into the guts of ITIL, it is important for the ITIL beginner to understand that ITIL and its counterpart, Microsoft Operations Framework (MOF), are not based on technology. Both ITIL and MOF are based on IT processes. This is important to understand before proceeding. For those readers interested in IT processes and procedures, as well as how the Microsoft System Center line of products fits into these processes, you may find the rest of this chapter very interesting. For those of you who yawned and rolled your eyes, we'll meet you at Chapter 2.

If you start researching ITIL, you will find that it is a series of books that describe an approach to IT service management. Originally created in the United Kingdom to address strict operations management standards, it has become the accepted standard in IT service management. The library is owned by the UK government's Office of Government Commerce (OGC). If you want to get real cozy with ITIL, be prepared to spend a lot of time reading. In its original form, the ITIL volumes were at a count of 60 books. These books were created by industry leaders of the time and described best practices for IT processes.

There is much more to ITIL than just the books, however. ITIL as a whole includes the books, certification, ITIL consultants and services, and ITIL-based training and user groups. ITIL is mainly updated by its own user group, known as the IT Service Management Forum (itSMF). The last piece of the puzzle, ITIL certification, is administered by the Netherlands Examination Institute for IT (EXIN) and the Information Systems Examination Board (ISEB).

ITIL can be divided into two categories: service support and service delivery. Numerous processes are divided up into the two different categories. Service support is described as a user-focused interface point, whereas service delivery is considered a customer-focused interface point. The reasoning behind this is to differentiate what is considered a user of the system and what is considered an actual customer of the system. Now you may be thinking, "I run an internal network. Everyone on my network is a user; we don't have any customers who connect into the network."

In all actuality, every admin has both users and customers on their network, and often the same user can be both a user and a customer. For example, HallieM is a *user* of the network when she interacts with the service desk. HallieM is also a *customer* of the network when she obtains certain services from another department, such as services that she must pay for or services that have availability management in place. Table 1.1 shows the breakdown of the difference between service support and service delivery.

SERVICE DESK

We will first look at the service desk, as it is unique from the other items in Table 1.1. The service desk is a *function* as opposed to a *process*, as are the other items listed. All incident reporting and service requests go through the service desk. It is the function that ties the service providers with the users, keeping users informed of service events and actions that may impact their day-to-day activities. The service desk becomes a single point of contact for customers and users to interact with the IT department. This approach helps expedite the call process by managing it in a timely and satisfactory way.

TABLE 1.1: ITIL Service Support and Service Delivery Differences

CATEGORY	FOCUS	TYPE	AREAS
Service support	user-focused	Operational	Service desk Incident management Problem management Configuration management Release management Change management
Service delivery	customer-focused	Tactical	Service-level management Financial management Capacity management IT service continuity management Availability management

INCIDENT MANAGEMENT

The incident-management process is mainly concerned with restoring normal service operations as soon as possible. This will help minimize any adverse effects on business operations and will ensure high levels of service quality and availability. Service-Level Agreements (SLAs) will determine what a “normal” service operation is. Information is collected about the incident to allow changes or enhancements in the environment to prevent future incidents. This information can also be used to compare against SLA compliance metrics and service quality information.

PROBLEM MANAGEMENT

The problem-management process is mainly concerned with minimizing the impact of incidents and problems. The goal is to reduce incident-resolution times by providing circumventions for known errors and removing the underlying causes. This strategy improves IT service quality by helping the service desk resolve incidents at the time of logging. If an incident can be resolved at the time of logging, business impact is reduced, business efficiency is improved, and IT efficiency is improved.

The problem-management process should not only be considered a “reactive” approach, however. When dealing with incident management, problem control, or error control, it is very reactive. The problem-management process can be viewed as “proactive” when you consider how it is used for problem prevention.

Problem investigation and diagnosis is used when known errors are created. During this investigation and diagnosis time, circumvention details, or “work-arounds” of the known errors are documented and distributed to the service desk and other support personnel until a fix for the problem is found. This approach helps with the staffing of the incident-management process, thus ensuring there aren’t too many IT staff members duplicating work while trying to fix the same issue.

CONFIGURATION MANAGEMENT

The configuration-management process is responsible for keeping an accurate and up-to-date model of the entire IT infrastructure. It uses this information to help support a number of areas:

- ◆ Allows for assessment of change- or problem-management functions
- ◆ Allows financial information to be gathered to help determine lease, rental, maintenance, and support costs for IT infrastructure components
- ◆ Supplies information about component performance and reliability to support capacity and availability management
- ◆ Improves security by identifying the location and details of assets, making it difficult for unauthorized changes to be carried out or undetected
- ◆ Helps with legal obligations by identifying the location of unauthorized software, determined by enabling authenticity checks on software and making sure current, correct versions of the software are being used

Configuration management uses this information to identify relationships between items that are going to be changed and any other components of the infrastructure that an item is tied to. Such a strategy enables the owners of the other components to be notified and involved in the impact-assessment process.

CHANGE MANAGEMENT

The change-management process is used to ensure standard methods are used when implementing change, and for developing and documenting reusable processes. You can reduce the possibility that a change in the environment could cause a failure, thus resulting in an incident, by having proven methods in place.

The IT infrastructure is changing all the time. Patches, service packs, updates, bios updates, and so forth are released on an almost daily basis. Having a safe and repeatable process in place is vital to service management.

RELEASE MANAGEMENT

Changes in the environment often result in the need for new versions of software, new hardware, new documentation, and so forth. The release-management process works closely with change management and configuration management to produce a secure and managed rollout of the new item. Consequently, physical changes to the environment are taken into account and the transition to live operation is successful—including both hardware and software releases.

The quality of a new version of software is tested in this process, along with tests to determine whether patches and updates are going to affect a piece of “approved” software. In this way, the process guarantees that only the authorized versions of software releases are being installed.

SERVICE-LEVEL MANAGEMENT

The Service-Level Management (SLM) process is responsible for creating SLAs and making sure Operation-Level Agreements (OLAs) are met at all times. During this process, changes to the environment are assessed to determine the effect on SLAs.

SLAs play an important role in SLM. They help set expectations for IT by determining what the customer's service-level requirements are, and they help customers by having a measurable understanding of what "good" service is. Both sides can agree on time lines for deliverables for everything from service upgrades, to updates, to incident resolution. SLAs also provide a clear understanding of what value customers are receiving from IT and can be used as a basis for charging for IT services. This brings us to the financial-management process.

FINANCIAL MANAGEMENT

The financial-management process is responsible for determining the costs of IT services as well as figuring the return on IT service investments. It is also a key in the role of recovering costs from customers if you charge for your services. As mentioned earlier, having SLAs in place to manage expectations is very important.

Budgeting can become much more accurate as well, because financial management is responsible for tracking costs of IT assets and resources. Financial management allows you to break down the money spent on IT services so you can clearly view where IT budget money went. Since budgeting is so much more precise, it helps support future business decisions on IT investments.

If you are considering charging for IT services, a fair recovery system is determined by data gathered through the financial-management process. Charging for internal services has its advantages and disadvantages. One advantage to charging for IT services is that it helps customers and users see the value of IT. Customers and users may also behave differently if they are faced with a "charging" model. Such a model helps the customers decide whether the services they are receiving are cost-justified. Using a model could lower the demands on the IT department.

One of the disadvantages of charging for services is that the customer has the ability to take their business elsewhere, which could have a severe effect on budgeting. Also, charging systems are often expensive, and the cost of such a model could offset the money that is generated by the system.

CAPACITY MANAGEMENT

The capacity-management process involves determining the required service delivery, determining the current service delivery for the IT infrastructure, and ensuring that all current and future capacity and performance requirements from the business are met. Capacity management also needs to take into account changes in new technology and the increase in performance that new technology brings to the table. Basically, this process is responsible for identifying the current service delivery as well as the service delivery potential at any given time.

Capacity management is responsible for making sure business requirements for system capacity are met at all times. Again, taking off the technical cap for a second, this does not directly relate to a technical capacity. It is related to the business requirements for the system, not necessarily the performance of the system.

IT SERVICE CONTINUITY MANAGEMENT

The IT service continuity management process ensures that an organization can continue to function with predetermined and agreed-on levels of IT services to support the minimum business requirements following an interruption to the business. The idea behind this process is that the organization will always have a base level of IT services that are required at all times.

Each IT service is examined to determine the minimum level it can function at to meet the business requirements. A plan is then put in place to guarantee that this level of service can be reached at all times under any circumstances.

AVAILABILITY MANAGEMENT

The availability-management process deals with the design, implementation, and management of IT services to guarantee that certain business requirements for availability are obtained. This requires information from both incident management and problem management to determine why an IT service failed and the time it took to resume service. This process can help IT departments meet SLAs that define availability levels. These SLAs cannot be met without a thorough understanding of the availability and reliability of IT components.

Availability management is a very high-profile process. Take an accounting server offline during a month-end run and see what kind of attention it gets. Because of this high-profile status, it is beneficial to have a single process owner for all availability issues to ensure consistent and comprehensive measures are taken for managing and improving availability of IT systems.

Exploring the Microsoft Operations Framework (MOF)

As stated earlier, MOF is the basis of Operations Manager. MOF was developed by Microsoft and a group of partners to expand on the best practices developed by ITIL. MOF includes a plethora of resources that are available to help you achieve mission-critical system reliability, manageability, supportability, and availability with Microsoft products and technologies. These resources are in the form of white papers, operations guides, assessment tools, best practices, case studies, templates, support tools, courseware, and services. All of these resources are available on the official MOF website at <http://www.microsoft.com/mof>.

HOW MOF EXPANDS ITIL

While ITIL is based on IT operations as a whole, MOF has taken the route of providing a service solution as its core. MOF focuses on the release and life cycle of a service solution, such as an application or infrastructure deployment.

Since ITIL was based on a philosophy of “adopt and adapt,” Microsoft decided to use it as its basis for MOF. Although Microsoft supports ITIL from a process perspective, Microsoft decided to make a few changes and add a few things when they built MOF. One of these changes and additions includes moving to a “prescriptive” process model. Microsoft defines the ITIL process model as “descriptive.” It has more of a “why” approach, whereas MOF has more of a “prescriptive,” or “how,” approach.

MOF also introduced the concept of Service-Management Functions (SMFs). As Table 1.2 illustrates, there are now 21 SMFs that describe the series of management functions performed in an IT environment. All of these SMFs map to an ITIL-based best practice for performing each function.

TABLE 1.2: MOF Quadrants Breakdown

QUADRANT	SMF	OMR (AT END OF QUADRANT)
Optimizing	Service-Level Management Financial Management Service Continuity Management Availability Management Capacity Management Workforce Management Security Management Infrastructure Management	Change Installation Review
Changing	Change Management Configuration Management Release Management	Release Readiness Review
Operating	System Administration Security Administration Service Monitoring and Control Job Scheduling Network Administration Directory Services Administration Storage Administration	Operations Review
Supporting	Service Desk Incident Management Problem Management	SLA Review

MOF also introduced the Team model. This gives MOF two core models; the other is the Process model. The Team model was added to fill a gap in ITIL. ITIL identifies roles for process owners of each operation process, whereas MOF creates seven distinct role clusters that describe the functional role or team:

Service Primary responsibility is to make sure all IT services are at a satisfactory level to customers and users. This is done by creating SLAs and ensuring that they are being met on a regular basis.

Infrastructure Responsible for ensuring plans are in place to keep networking, telecommunications, hardware, and software running in order to satisfy business requirements.

Support Maps to the service desk, incident management, and problem management functions in ITIL.

Operations Responsible for making sure that the day-to-day tasks of running the IT systems are met, according to SLAs.

Partner This is more of a “virtual” team in the IT department, usually made up of outsource vendors, IT partners, resellers, service providers, consultants, and so forth.

Security Responsible for data confidentiality, data integrity, and data availability.

Release Transitions a release between development or test environments into production. A release could be a new software package, an update, a patch, and so forth. The Release role also has the responsibility of maintaining accurate inventory management and asset management.

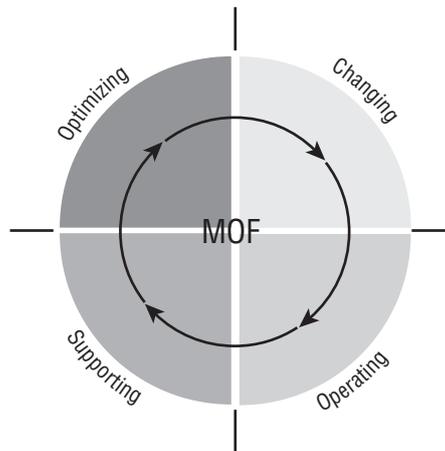
The Risk Management discipline was added to provide the management of risk to its own management discipline. ITIL provides only theory discussion, not detailed steps, about handling of risk for each IT operations process.

Explicit management review checkpoints are also built into MOF to guarantee that there is involvement by management at each key step in the process. The ITIL books do not include these checkpoints. This is another value-add that Microsoft provides with MOF.

MICROSOFT OPERATIONS FRAMEWORK PROCESS MODEL

The MOF Process model breaks down a complex environment into an easy-to-manage and easy-to-understand set of functions. This is accomplished by the numerous SMFs there were added by Microsoft when they created MOF. SMFs are just a portion of the overall Release Cycle that MOF employs.

Microsoft defines a release as any change, or set of changes, that is incorporated into a managed environment. A release includes not only changes in applications or operating system updates, but also changes in operations processes or changes in the physical environment. These releases have a defined life cycle. The life cycle is defined by quadrants, Operations Management Reviews (OMRs), and SMFs. The four quadrants are divided by the different SMFs that relate to each quadrant. SMFs are groups of best practices. These SMFs are broken into four categories that explain the activities of an operations environment. Graphic 1.1 shows the four MOF quadrants in the MOF life cycle.



Changing Quadrant

The Changing quadrant is a group of SMFs that define the proper introduction of approved changes into a managed IT environment. This can include changes in applications, hardware, and systems, as well as changes in policies and procedures. The Changing quadrant maps to the ITIL discipline

of service support. The three SMFs that reside in the Changing quadrant are Change Management, Configuration Management, and Release Management.



Operating Quadrant

The Operating quadrant is a group of SMFs that are used to monitor, control, manage, and administer service solutions to achieve and maintain service levels. All of the SMFs in the Operating quadrant are items that Microsoft has specifically added to expand ITIL.



System Administration The day-to-day administration of services and systems in an IT infrastructure. This could include user and group account administration; administration of file, print, database, and applications servers; low-level monitoring; and troubleshooting of the systems in the IT infrastructure.

Security Administration The administration of security in an IT infrastructure. This includes monitoring the environment in both a reactive and a proactive way, thus ensuring that the environment is safe from attack. This is done through many facets, including identification and authorization control, access control, and auditing, to name a few.

Service Monitoring and Control The administration and monitoring of the health of an IT service. This SMF ensures that SLAs are in place and that business requirements for IT services are being met.

Job Scheduling The administration and scheduling of jobs and processes so that an efficient sequence is utilized. This could include scheduling batch jobs to maximize system throughput and utilization and to meet SLAs.

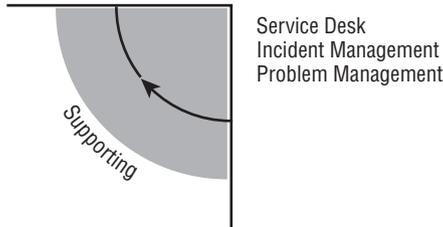
Network Administration Administration of the network to ensure that the network operates at an efficient level at all times. This includes the administration of people, processes and procedures, vendors and service providers, as well as the administration of the network hardware.

Directory Services Administration The administration of resources in Active Directory, such as users, applications, servers, printers, and so forth. The goal of this SMF is not only to make sure that directory access is always available, but also to ensure that information from the directory is available via a simple and centralized process.

Storage Management The administration and control of data, both electronic and physical, for the purposes of restoration and historical archiving. This includes both onsite and offsite storage. Storage Management was put into place to help guarantee the physical security of backups and archives.

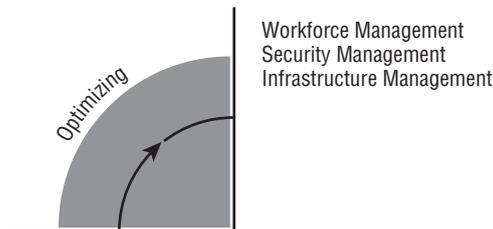
Supporting Quadrant

The Supporting quadrant is a group of SMFs that identify, assign, diagnose, track, and resolve incidents and problems in a timely manner within SLAs. The Supporting quadrant maps to the ITIL discipline of service support. The three SMFs that reside in the Supporting quadrant are Service Desk, Incident Management, and Problem Management.



Optimizing Quadrant

The Optimizing quadrant is a group of SMFs that help maintain business and IT alignment by attempting to decrease IT costs while maintaining or improving service levels. The Optimizing quadrant introduces three new SMFs to help expand the base ITIL disciplines.



Workforce Management This function was added to specifically address staffing issues in the IT infrastructure team. It helps with the process of attracting, developing, and retaining a properly trained and prepared IT staff. It also ensures that the work environment is safe and efficient.

Security Management This function was created to help an IT infrastructure define and communicate the business’s security plans and policies, based on the guidelines and regulations that apply to that business.

Infrastructure Engineering Think of this function as the “project manager” of MOF. The processes and tasks in the Infrastructure Engineering SMF could be linked to any other SMF to help coordinate engineering policies and standards.

OMRs are either event-based or time-based. The Change Initiation and Release Readiness reviews are event-based and occur at the initiation and final installation of a review into the target environment.

Change Initiation Review The Change Initiation Review is triggered when approval has been requested for a proposed change to the environment. This begins the process for actually implementing the release. Investments in money, time, equipment, and staff will now begin to work on the process and get it ready for release.

Release Readiness Review The Release Readiness Review determines when a release is confirmed as ready for production. The proposed release is checked to ensure standards, policies, and quality metrics are in place to support the release.

The Operations Review and Service Level Agreement Review occur at regular intervals to assess the internal operations as well as performance against customer service levels.

Operations Review The Operations Review is a regularly scheduled review to assess and improve IT operations based on business need and SLAs. Operations reviews use information from operations guides, company policies and procedures, and operating-level agreements to measure and evaluate the performance of the operations staff.

Service Level Agreement Review The Service Level Agreement Review is a regularly scheduled review to assess and improve the alignment of business needs with IT service delivery defined in SLAs. During this review, the operations staff and service-level management take current information and measure that against published SLAs to determine whether the service has met its service-level requirements.

Inside these four quadrants is a collection of 21 SMFs. Each quadrant consists of a group of SMFs that break down the quadrant into logical procedures and tasks. Each SMF is assigned to a home quadrant, but SMFs are by nature cross-functional and cross-quadrant. If you look back at Table 1.2, you can see how these processes define an SMF, each of which is a series of actions or operations that are designed to achieve a goal. Each process is then broken down into procedures, which allows for coordination between departments. Each procedure has a series of tasks that must be performed to complete the procedure. The task is the lowest level of effort on a project.



Real World Scenario

MIXING IT UP

So you say you don't have enough coworkers to fill all of the positions for your project? According to MOF recommendations, you don't need to have one person for every role. Looking at any of the MOF guides for services, you will find role clusters that define the responsibilities for each of the roles. The separation of responsibilities is meant to define who is allowed to perform each action, which also grants a level of accountability. When one administrator needs to make a change, there are others who can validate that the change is required, and still others who can verify that the change was made appropriately. Not every organization can have enough personnel to fulfill every one of the roles. If your company has a small staff, you will need to assign multiple roles to each of your administrators. The best option is to look at each of the role clusters and have one administrator take on all of the actions that are defined for a cluster. Not enough personnel to assign one person per role cluster? Look at each role cluster and determine which role clusters do not pose a possible conflict of interest. You want to maintain some accountability, so assign responsibilities so that there are checks and balances in place.

Tying It All Together

Microsoft set out to provide you with tools that help you manage your IT systems. They achieved this aim by integrating systems-management tools and knowledge of the systems to help you with day-to-day operations of the environment, as well as ease your troubleshooting efforts and time and improve planning capabilities. Microsoft took ITIL and MOF, and created a “family” of products—known as the System Center—that helps you in your quest to align with the practices set forth in those documents.

DYNAMIC SYSTEM INITIATIVE

The System Center suite of products helps IT organizations capture and use information to design more manageable systems and automate IT operations. As software becomes more and more complex, thus introducing new components and systems to the infrastructure, the infrastructure will in turn become much more complex. For example, an inventory application moves from being client server-based, to multitier, to a web service-based application. As the application grows and more users start using it, the decision is made to install a hardware load balancer in front of it. Then the data is moved to a Storage Area Network (SAN) to give the IT department better control over backup and recovery options.

All of these changes result in many different “teams” in the IT department being involved with this “application.” You quickly see how a change to the application can affect more than just the application developers. You now have to coordinate changes with the Web Server team, the database administrators, the Networking team, and the Storage team.

Whether these teams are made up of one person (you), or they are made up of dozens of people on each team, you quickly realize how complex the infrastructure can become, and why there is a need for management of these distributed systems. The Dynamic Systems Initiative (DSI) is a plan to build software that incorporates ITSM capabilities and MOF best practices with the software (System Center) in order to match IT capabilities and operations with business needs.

DSI will help IT organizations deliver end-to-end offerings that will:

- ◆ Increase productivity and reduce costs across the entire IT organization
- ◆ Reduce time and effort required to troubleshoot and maintain systems
- ◆ Improve system compliance with business and IT policies
- ◆ Increase responsiveness to changing business demands

Currently, the goals of DSI are accomplished through management packs in Operations Manager, which we will discuss in greater detail in Chapter 5. In the future, this strategy will be achieved through a process called the System Definition Model (SDM).

The SDM is at the heart of DSI components and products. It is an XML-based schema that is used to create “models” of IT systems. The SDM can then manipulate these models using software tools. The SDM will be able to capture data that is vital to the operations of the system. This approach will give IT departments end-to-end solutions that are integrated across applications, operating systems, hardware, and management tools and will provide reduced costs, improved reliability, and increased responsiveness in the entire IT life cycle.

SYSTEM CENTER

Before we look at the System Center line of products, let's first look at the different System Center management disciplines that were introduced by Microsoft to help define IT service management:

- ◆ Operations Management
- ◆ Change Management
- ◆ Configuration Management
- ◆ Release Management
- ◆ Asset Management
- ◆ Data Protection Management
- ◆ Problem Management
- ◆ Capacity Management
- ◆ Incident Management

Looking at this list and considering Microsoft's plan for the System Center family of products, you can quickly see how they have embraced this framework. Through internal products, close work with partners, and acquisitions of software from other companies, Microsoft has addressed each one of these disciplines with a product in the System Center suite.

System Center Operations Manager Formerly Microsoft Operations Manager (MOM). This product provides you with tools to help you proactively monitor your network as well as reduce the complexity associated with managing an IT infrastructure.

System Center Configuration Manager Formerly Systems Management Server (SMS). This product provides a comprehensive solution for the Change Management and Configuration Management disciplines. It touches Release Management and Asset Management as well since it is often used to roll out patches and updates and has a widely used inventory feature.

System Center Data Protection Manager Data Protection Manager (DPM) is a backup and recovery product from Microsoft that gives the end user some recovery options to help take some of the burden off the IT staff. DPM is a centralized backup solution that captures changed files to disk, providing rapid and reliable recovery.

System Center Capacity Planner Capacity Planner is a tool to be used in a predeployment scenario when planning a new deployment, infrastructure changes, or upgrades. It provides best-practice guidance and hardware-specific knowledge to help planning around Exchange 2003 and 2007, as well as MOM 2005 and Operations Manager.

System Center Reporting Manager Reporting Manager is a data warehouse and reporting program that gives you reporting capabilities and management data from products in the System Center line, such as SMS 2003 and Configuration Manager, MOM 2005 and Operations Manager, and other data sources, such as Active Directory. Reporting Manager gives IT managers the ability to support decisions that will affect the corporation.

Service Desk This is another product that touches more than one management discipline. Service Desk is an incident management tool (think help desk) that gives the end user a direct interface to the IT department to provide information about IT infrastructure issues. Service Desk can also be used as a repository for tracking information about IT assets and processes, allowing it to also touch the Change Management and Asset Management disciplines.

Defining Operations Management

There is often some confusion when it comes to the actual definition of *operations management*. Microsoft's System Center family of products comprise several products that span a wide range of "management" ground. The most confusing portions of this area are between systems management and operations management. We will look at the difference between the two.

Systems Management

Systems management is typically defined as software that is used to centrally manage large groups of computer systems. This software contains the tools to control and measure the configuration of both hardware and software in the environment.

Microsoft's entry into this arena is with a product called System Center Configuration Manager. Configuration Manager provides remote control, patch management, software distribution, hardware and software inventory, user activity monitoring, capacity monitoring, and much more.

Operations Management

Now that you have an understanding of what falls under the category of systems management, we will focus on operations management. *Operations management* is mainly focused on ensuring that business operations are efficient and effective through processes that are aimed at improving the reliability and availability of IT systems and services. You accomplish this by gathering information from your current systems, having the proper people in place to decipher that data, and having proper procedures in place to carry out any tasks that may arise if there is a current or potential problem in your environment.

The System Center product that addresses this need is System Center Operations Manager. This is the product that will be the focus of this book. Operations Manager is based on MOF, which in turn is based on ITIL. System Center Operations Manager is a product that allows centralized monitoring of numerous computers on a network. Many hardware and software products can be monitored by Operations Manager, such as Active Directory, SQL Server, and Exchange Server. Operations Manager provides you with the information you need to help reduce time and effort in managing your IT infrastructure by automating tasks and giving you a proactive approach at determining possible problems.

The Bottom Line

Define IT service management. Start thinking in terms of the customer's and user's perspective when it comes to IT systems and how they operate. Along with the technical approach to IT, consider implementing a process-based approach to complement the technical side. This will help minimize downtime and help your business meet requirements set forth in SLAs.

Learn how ITIL is the foundation of IT service management. ITIL was created by the United Kingdom's Office of Government Commerce to create a framework of best-practice procedures to help support business. This increases quality and value of both IT systems and services through the creation of a set of policies for service support and service delivery.

Understand how MOF expands ITIL. Using ITIL as its foundation, Microsoft set out to customize this set of best practices and tune them to fit the Microsoft philosophy. MOF takes the ITIL processes in service delivery and service support and breaks them out into SMFs. The SMFs are located in four quadrants: Changing, Operating, Supporting, and Optimizing.

Explore the Dynamic Systems Initiative. To help support MOF, Microsoft has started to build systems that will ease the administrative burden of managing the growing complexity of the IT infrastructure. This is known as the Dynamic Systems Initiative (DSI). DSI will help IT administrators tie together the MOF best practices with the software and systems that they manage.

Learn about the Microsoft System Center products. Microsoft has aligned a suite of products under the title of System Center to help an IT organization meet the best practices set forth in MOF. All of the current and future System Center products align perfectly with the MOF quadrants and SMFs.

Learn how to use System Center Operations Manager. Focusing on the Operating quadrant of MOF, this product is a management tool that helps administrators gather information from the Windows servers and software in your environment. This information can then be viewed from a central console. This information can help IT administrators make decisions for the management, tuning, and security of the servers and software that they manage.

Chapter 2

Installing System Center Operations Manager 2007

In the previous chapter we discussed the framework that Microsoft has adopted for the System Center family. In this chapter we are going to take a look at what you should do when deploying Operations Manager 2007. Most administrators consider this the most rewarding part of the job since they get to apply hands-on what they have planned. Building the servers and configuring them so they work as expected can be very time-consuming, but time tends to pass quickly when you are setting up the systems to work for you. This chapter covers the installation requirements and steps when configuring the Operations Manager systems within your infrastructure. Since there are different ways to implement Operations Manager components, we will tackle each of the ways.

In this chapter you will:

- ◆ Identify the system requirements
- ◆ Install the database on a remote SQL 2005 server
- ◆ Install the database locally on the Operations Manager system
- ◆ Install Operations Manager 2007
- ◆ Verify the installation of Operations Manager 2007

Identifying Server Components

At the heart of the Operations Manager infrastructure are the server and database components. The Operations Manager *server* is responsible for notifying the agents how they are supposed to function, collecting alert and event data as it is sent by the agents, and responding to those events that require it. The *database* stores all the data that is collected by the agents and sent to the Operations Manager server. These are two components that you cannot live without in your Operations Manager infrastructure. You will have at least one of each, possibly more.

Management servers come in a couple of flavors. To start with, the first Operations Manager 2007 system you build within a management group becomes the Root Management Server (RMS). The RMS is the coordinator for health monitoring. As other management servers receive data from monitors within the management group, the RMS is responsible for validating all of the health criteria for distributed applications and services. This way, you can have multiple management servers receiving data from all of the monitored systems within your network and still monitor end-to-end services.

The RMS also acts as a coordinator for changes to rules, monitors, and alerts. Whenever an operator opens the Operations Console and uses the Authoring or Administration views, they are connecting to the RMS. After the modifications are made, the RMS updates the database and relays the changes to the other management servers.

Management Servers are responsible for communications with the agents that are loaded on the monitored systems. As new rule, alert, and monitor configurations are made, the management server will send the changes to the agent. As the agent collects data based on the rules, alerts, and monitors, the agent will deliver the data to the management server, which is then responsible for updating the database.

Gateway servers are new in Operations Manager 2007. To monitor systems that reside in untrusted domains, the gateway server is configured to accept data from the agents within the untrusted domain and pass the data to management servers within the trusted domain. Doing so allows you to monitor systems from partner organizations, Demilitarized Zones (DMZs), and isolated domains.

The final role is known as the Audit Collection Server (ACS). The ACS is a management server that has been configured to accept forwarded security log events from a specially configured agent. The ACS is not installed by default. You have to install the collector service on a management server and enable the forwarder function on the agent. A database used specifically for audit collection is also required. Once configured, any event recorded in the security log on a monitored system will be saved in the ACS database.

Operations Manager 2007 Server

The primary component of the System Center Operations Manager 2007 infrastructure is the Operations Manager 2007 management server. When thinking of the management server, we think of it as the reporting point for alerts and events. It is the component that is responsible for making sure that the criteria we set to monitor is actually monitored. When we talk about the Operations Manager 2007 server from here on out, we are going to refer to it as the management server. When we are specifically talking about functionality that is found on the Root Management Server, we will identify it by name. We will refer to the entire Operations Manager 2007 infrastructure as Operations Manager. That way, we can differentiate between the server and the supporting services.

Although the management server is responsible for monitoring systems, it also provides additional functionality. If monitoring were the only thing that we were interested in, we have other tools at our disposal that can provide that functionality. But we want to have a system that will also help us support our organization. Two terms are bandied about when administrators talk about service management and monitoring: proactive and reactive. Reactive management is not a situation most administrators want to find themselves in. As the term implies, you are in a position of reacting to events. As problems occur, you are going to have to scramble to right them. If you are constantly in a state of “putting out fires,” you will not have the ability, or time, to manage other parts of your organization.

Proactive management, on the other hand, is the ideal situation for administrators. A proactive administrator will monitor systems, constantly watching for signs of impending problems. Many problems that occur have warning signs that appear before something tragic happens. Drive failures are usually preceded by read or write errors. Services failing will usually emit warnings prior to an actual failure. By monitoring all the systems, you can see patterns and issues that may forewarn you of disaster. That is not to say that you won't ever have to react to a problem. There will be those issues that appear out of nowhere. Disasters strike and we have to deal with that. But being proactive will help you alleviate many problems.

You may be asking yourself, “How do I stay proactive with all of the systems I have?” It is true that the more servers you add to your organization to alleviate workload and automate or complement your business processes, the more points of failure and problems you present. And that is where Operations Manager can help. Not only will Operations Manager become the centralized monitoring solution, capturing alerts and events as they appear, but it can also provide support by reacting to those alerts and events.

In Chapter 5, “Managing Management Packs,” we will discuss how you can configure Operations Manager to react to events and alerts either by having the management server provide a response or by configuring the agent running on the managed server to respond. Another function that helps administrators monitor their systems is the consolidation of events. That way, if several events appear on one system, you can present them to the management server as a single event notification, thus keeping the size of the database smaller. You can also monitor multiple systems for the appearance of the same event or alert, possibly identifying a distributed attack or multiple system failure.

You are not limited to a single management server in your organization. You can create a management group and then populate it with multiple management servers, distributing the load across them to make monitoring more efficient. You can also set up failover criteria; in case one of the management servers fails, you can have a managed system automatically start reporting to another management server. This strategy reduces the possibility of missing alerts or events.

Also new to Operations Manager 2007 is the ability to install the Root Management Server on clustered hardware. This gives you the peace of mind of knowing that the primary management server will be highly available. You have always been able to install the SQL database on clustered equipment, but now you have another layer of availability that you did not have in the past.

The management server itself becomes the central location for reporting alerts and events, but by itself it does not provide much control over the agents that reside on the client systems. To be able to control them and to specify how they are going to work, you need to add in or create management packs. Management packs contain the rules that let the agents know what they are supposed to be looking for and how to respond. In Chapter 5 we will look at management packs and see how they affect the agents and servers.

DataStore (SQL Server 2005)

Management servers need to have a storage location for the data that they collect. For Operations Manager 2007, the only supported database is SQL Server 2005. Those installations that have not yet implemented SQL Server 2005 will find that they are going to have to come up to speed to take advantage of Operations Manager 2007. That is not a bad thing, considering SQL Server 2005 is much more efficient and robust than its predecessor. However, some companies have not budgeted for the time and training necessary to implement SQL Server 2005. But SQL Server 2005 is the database server of choice according to Microsoft. After all, the previous version of SQL had been out for five years before SQL Server 2005 came on the scene, and there have been a lot of changes in the database world in that amount of time.

The database itself can coexist with Operations Manager on the same server if you have a small organization. If that is the case, you can take solace in the fact that the two services will play nice with each other, although they will struggle for resources. That is why you don’t want to place the database on the same server as Operations Manager if you have a large organization. So what is a “large” organization, you ask? That depends on how much monitoring you are doing. The rule of thumb is, if you have more than 10 servers that you need to monitor, you should separate the database from the Operations Manager server. And while this may hold true for many companies, you should run tests against your systems. Some companies decide they want to collect more data

than others, and the more management packs you add and the more rules that you enable, the harder your systems will have to work to process the information.

Although the management servers within a management group will use only a single database server, you can cluster the server to make sure that you have some fault tolerance built in. Setting up a clustered SQL server is quite easy once you have the hardware in place. And although it is more expensive to implement clustering, it may be worth the added initial monetary investment when it comes to making sure your database is available. You can take comfort in the fact that, if the database server goes offline, the management server can hold the data it collects. That way, even if the database server cannot be contacted, you will not lose monitoring information that you want to retain in the database for reporting or trend analysis purposes.

Verifying Prerequisites

Every piece of software has its bare minimum requirements for installation. It seems like every time a list of requirements and prerequisites is published, everyone groans and takes it for granted that they could install it on a platform that meets the requirements but they wouldn't dare to run a system that way. So with that in mind, we present the list of minimum requirements for Operations Manager 2007.

Identifying Prerequisites

Starting off with the supporting services, you need to make sure that you have Active Directory in place within your environment if you want to take advantage of all the new features. Of course, that means that your Domain Name System (DNS) infrastructure needs to be sound and stable. Active Directory will allow you to take advantage of creating and using the Operations Manager container, which makes configuring and controlling managed systems easier. Instead of making you pull up the documentation for Operations Manager 2007, we will present Tables 2.1 through 2.8, which detail the prerequisites you will need to meet for each of the server types.

TABLE 2.1: Prerequisites for the Database Server

REQUIREMENT	MINIMUM
Processor	x86 1.8 GHz
	x64 1.8 GHz
Memory	1 GB
Free space	10 GB database size
Operating system versions	Windows Server 2003 SP1 Standard or Enterprise
	Windows Server 2003 R2 Standard or Enterprise
Database	SQL Server 2005 SP1 Standard or Enterprise

TABLE 2.2: Prerequisites for the Management Server (Root/Non-Root)

REQUIREMENT	MINIMUM
Processor	x86 1.8 GHz
	x64 1.8 GHz
Memory	1 GB
Free space	5 GB
Operating system versions	Windows Server 2003 SP1 Standard or Enterprise
	Windows Server 2003 R2 Standard or Enterprise
.NET Framework	Version 2.0
	Version 3.0
Microsoft Core XML Services	Version 6.0

TABLE 2.3: Prerequisites for Installing the Operations Console

REQUIREMENT	MINIMUM
Processor	x86 1 GHz
	x64 1 GHz
Memory	1 GB
Free space	5 GB
Operating system versions	Windows Server 2003 SP1 Standard or Enterprise
	Windows Server 2003 R2 Standard or Enterprise
	Windows XP Professional SP2
.NET Framework	Version 2.0
	Version 3.0
Office Word	Office Word 2003 with .NET Programmability Support
PowerShell	Version 1.0
Visual Studio Tools for Office System	Needed for knowledge authoring

TABLE 2.4: Prerequisites for Installing the Agent

REQUIREMENT	MINIMUM
Processor	x86 minimum OS requirements
	x64 minimum OS requirements
	IA-64 minimum OS requirements
Memory	Minimum OS requirements
Free space	30 MB
Operating system versions	Windows 2000 SP4 Professional, Standard, Advanced or Datacenter editions
	Windows Server 2003 SP1 Standard, Enterprise, Datacenter, Small Business, and Storage Server editions
	Windows Server 2003 R2 Standard, Enterprise, Datacenter, and Small Business editions
	Window XP Professional SP2
	Windows Vista Business, Enterprise, and Ultimate editions
Microsoft Core XML Services	Version 6.0
Windows Installer	Version 3.1

TABLE 2.5: Prerequisites for the Reporting Database Server

REQUIREMENT	MINIMUM
Processor	x86 1.8 GHz
	x64 1.8 GHz
Memory	1 GB
Free space	5 GB
Operating system versions	Windows Server 2003 SP1 Standard or Enterprise
	Windows Server 2003 R2 Standard or Enterprise
Operating system components	Internet Information Server 6
Database	SQL Server 2005 SP1 Standard or Enterprise

TABLE 2.6: Prerequisites for the Gateway Server

REQUIREMENT	MINIMUM
Processor	x86 1.8 GHz
	x64 1.8 GHz
Memory	1 GB
Free space	5 GB
Operating system versions	Windows Server 2003 SP1 Standard or Enterprise
	Windows Server 2003 R2 Standard or Enterprise
.NET Framework	Version 2.0
	Version 3.0
Microsoft Core XML Services	Version 6.0

TABLE 2.7: Prerequisites for the Audit Database

REQUIREMENT	MINIMUM
Processor	x86 1.8 GHz
	x64 1.8 GHz
Memory	1 GB
Free space	20 GB database size
Operating system versions	Windows Server 2003 SP1 Standard or Enterprise
	Windows Server 2003 R2 Standard or Enterprise
Database	SQL Server 2005 Enterprise

TABLE 2.8: Prerequisites for the Management Server with Audit Collector or Agentless Exception Monitoring File Share

REQUIREMENT	MINIMUM
Processor	x86 1.8 GHz
	x64 1.8 GHz
Memory	1 GB

TABLE 2.8: Prerequisites for the Management Server with Audit Collector or Agentless Exception Monitoring File Share (*CONTINUED*)

REQUIREMENT	MINIMUM
Free space	10 GB
Operating system versions	Windows Server 2003 SP1 Standard or Enterprise Windows Server 2003 R2 Standard or Enterprise
.NET Framework	Version 2.0 Version 3.0
Microsoft Core XML Services	Version 6.0

Even though you will be able to install and run Operations Manager with the aforementioned minimums, you will find that it will not run efficiently. Tests have shown that management servers need to have 2 GB of RAM to be effective. As with most servers, the more resources you throw at it, the better.

Running the Prerequisite Check Utility

Built into the Operations Manager installer web page is a nice tool that will check your system to make sure it meets the minimum requirements for the components that you are going to install. When you drop the Operations Manager CD into the CD tray of your system, or when you initiate the `setup.hta` file from the installation media or a network share, the setup options page appears. As you can see in Figure 2.1, a handy link to the prerequisites viewer is available when you click Check Prerequisites.

FIGURE 2.1
Operations Manager
splash screen



On the prerequisites screen, you can choose to check the prerequisites for each of the components that are used in your Operations Manager implementation. When you are ready to configure the database server, you can select the database server option and leave the rest of the options deselected. If the database server is also going to act as your reporting server, you can select both options to make sure that all the required settings are available.

If your system does not meet the requirements, you will see a large Failed listed, as shown in Figure 2.2. The design team has done a rather commendable job of making sure you have enough information about any error condition that arises. If you click the More button beside any of the error conditions, you will be presented with a detailed view of how you can resolve the problem, as shown in Figure 2.3.

FIGURE 2.2
Prerequisite Checker showing failed items

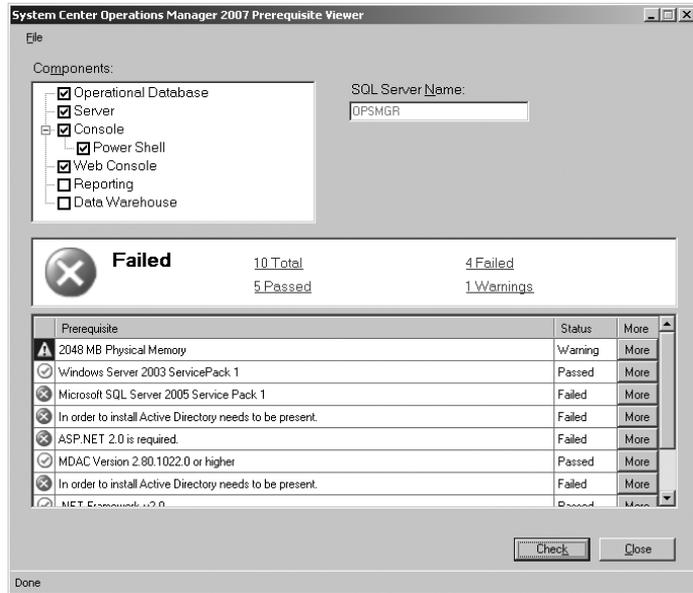
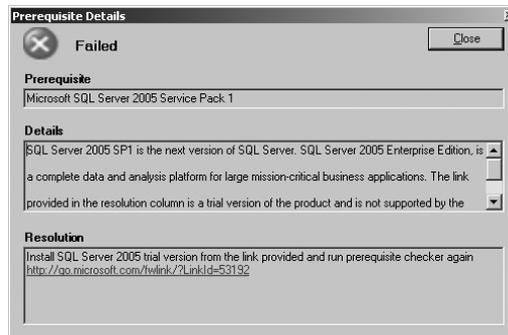


FIGURE 2.3
Prerequisite Checker showing detailed information about the failed item



After you have corrected any of the issues that appeared in the Prerequisite Checker, run it again to make sure that all criteria show up with a green check mark next to them. If that is the case, you are ready to start installing the appropriate components for your Operations Manager management group.

Installation Criteria

As you are configuring your environment to support Operations Manager, you will need to make sure not only that the servers have all of the prerequisites in place, but also that you prepare the services and accounts you will need. That means going back to your design and planning documents to make sure you have everything you need. Depending on your current business infrastructure, you may need to create more accounts, or you may need to make sure that all your components run on separate servers. As we go through this next section, we will take each scenario into account.

Creating Accounts

Operations Manager relies on accounts to perform certain actions in the management group. One of the key functions is installing the agent. Without the agent, Operations Manager will not be able to collect all the pertinent data from the managed systems. Of course, you could use the agentless management approach, but you still need to make sure that you have some of these accounts created and enabled if you want Operations Manager to function correctly.

AGENT ACTION ACCOUNT

The Agent Action account is used by the agent on the managed system. This account is used to allow the agent to access components that will be monitored. Most of the components that are monitored are accessible only to services that have sufficient privileges. To keep from assigning accounts to groups that have a high level of privileges in the domain, you can create regular user accounts and then grant them specific rights on the managed systems.

At a bare minimum, you need to make sure that the Agent Action account has the following privileges:

- Member of the local Users group
- Member of the local Performance Users group
- Allow Log On Locally

Once you grant these privileges, the agent will be able to run as if it were a user account logging onto the system. Once logged on, the agent can start monitoring the system. In most organizations, however, a user account is not allowed to have much access to the operating system. Limiting what a user account can access allows an organization to control their user environments. At the same time, limiting accounts in this way severely limits service-based accounts. Some of the actions that the Action Account may need to perform could require more rights and permissions than what you have granted here. You will need to make sure that the Action Account has all the rights required to perform any action that the management pack calls for. You should always read the documentation that ships with the management pack to see if you have to elevate any rights or permissions.

For example, when you start collecting performance data from the monitored systems, limited accounts may not be able to retrieve the performance counter data from the operating system and the Windows Management Instrumentation (WMI) database. You will need to make a couple of

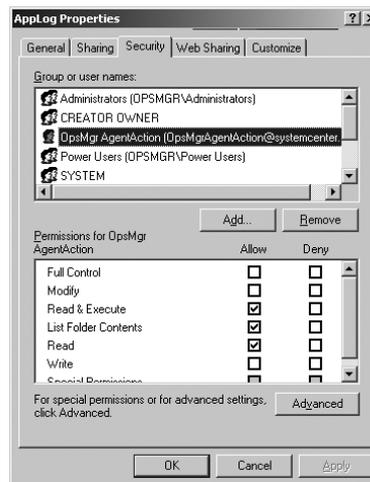
changes to ensure that the account responsible for collecting this information, the Agent Action account, has the appropriate rights. If you make the account a member of the performance users group, the account will have the appropriate access to performance data from the system.

Most of the management packs that will be added to the management group need to access even more information on a managed system. Most management packs will require access to event log and application log data. To ensure that the agent can perform its duties, you should grant the Agent Action account the Manage Auditing And Security Log privilege. Doing so will allow the agent to have access to event log data on a managed system, but you will have to take additional steps for allowing access to log files that applications create and use.

Applications create their own log files. If a management pack monitors the data that is captured in an application log, the agent requires access to the log files. In this case, make sure that the Agent Action account has read permissions to the directory where the application logs are written.

Allowing read access to an application's log directory is as simple as adding the account used for the Agent Action account to the access control list with an Allow permission for Read. If you look at Figure 2.4, you will see the access control list (ACL) for the log directory showing the Agent Action account set for Read.

FIGURE 2.4
Log directory with appropriate permissions for the Agent Action Account



SDK AND CONFIG SERVICE ACCOUNT

Each management server needs to “talk” to the database. To do so, an account needs to be created and assigned so that the Operations Manager system can authenticate. As with the Agent Action account, you have the option of creating an account within the domain, or you can use the Local System account. We recommend the Local System account since it is a rather secure account. Note, however, that if a system is compromised, any permissions that are assigned to the Local System account may be taken advantage of by the attacker.

To reduce the exposure to resources, consider using a different account for the SDK and Config Service account than you used for the Agent Action account. Doing so will limit the amount of access an attacker will have if they are able to compromise the service. Of course, if you plan on using the Local System account for both of these accounts, you won't have the ability to separate the access.

OPERATIONS MANAGER ADMINISTRATORS

This is also a good time to create the group that will be used for the Operations Manager administrators. The members of this group will be allowed to work with all aspects of your management groups. Creating this group will allow you to identify the group as the administrative account when you configure Active Directory and when you install your root management server. This way, you will not be restricted to using a single user account as the administrator for the management group.



Real World Scenario

MANAGING OPERATIONS MANAGER ACCOUNTS

One way to keep track of the accounts that you are using is to create an Organizational Unit (OU) within Active Directory for the sole purpose of organizing Operations Manager accounts. For large organizations that have several management groups controlled by different administrative teams, you can create an OU for organizing Operations Manager accounts, and then create lower-level OUs based on the management groups.

Configuring Active Directory

System Center Operations Manager 2007 is the first Operations Manager product to take advantage of Active Directory. Other Microsoft products—Exchange Server, Systems Management Server, and Internet Security and Acceleration Server, to name a few—have all used Active Directory as a repository for configuration data. And for good reason—taking advantage of Active Directory’s replication model makes for an efficient way to guarantee that each of the managed systems have the configuration data close by.

At this point in the game, we are assuming that organizations willing to implement a monitoring solution such as Operations Manager 2007 have moved away from Windows NT 4 and have an Active Directory implementation in place. If your organization has not, you should think about doing so. For one thing, NT 4 has moved past its support lifetime and can be costly to maintain. Even though you will need to purchase new software and licenses for Windows Server 2003, the gains you receive from an administrative standpoint are staggering. But this is not the place to start going into the administrative and management cost savings. Instead, we need to talk about configuring Active Directory to support Operations Manager 2007.

Like its sibling product within the System Center line, Configuration Manager 2007, Operations Manager enters configuration settings in Active Directory so that an agent can look to the nearest domain controller and determine how it is supposed to function. With Active Directory supplying the configuration settings to the agent, the client system does not need to know which management server it needs to “talk to” in order to start sending responses. Instead, it can look at the domain controller, determine how it is supposed to work in the management group, and then start collecting data.

The first step in configuring Active Directory is making sure that you have the appropriate rights. The only accounts that are allowed to create the Operations Manager container are those that are members of the Domain Admins group for that domain, or members of the Enterprise Admins group within the forest. If you are a member of one of these groups, you can continue. If you haven’t been granted the appropriate credentials, you must ask the Active Directory team to create the container for you. More often than not, this is the scenario that administrators find themselves in. They work in environments where the responsibilities are divided up.

The next thing you need to verify is that your domain is at the Windows 2000 Native or Windows Server 2003 functional level. If not, you will receive an error as you attempt to run the command that creates that container and assigns permissions.

The tool that you will need to use is `MOMADAdmin.exe`. After you install a management server, you will find this tool in the Operations Manager installation directory. For most administrators, the next step is to copy this utility to a USB drive or a network location and then hand it over to someone on the Active Directory team. No matter who is allowed to run this tool, you will find that it is pretty uneventful when it runs.

The command-line syntax for the `MOMADAdmin` utility is as follows:

```
MOMADAdmin.exe <ManagementGroupName>
➤ <OpsMgrAdministratorsSecurityGroup>
➤ <RootHealthServiceComputerName> <DomainName>
```

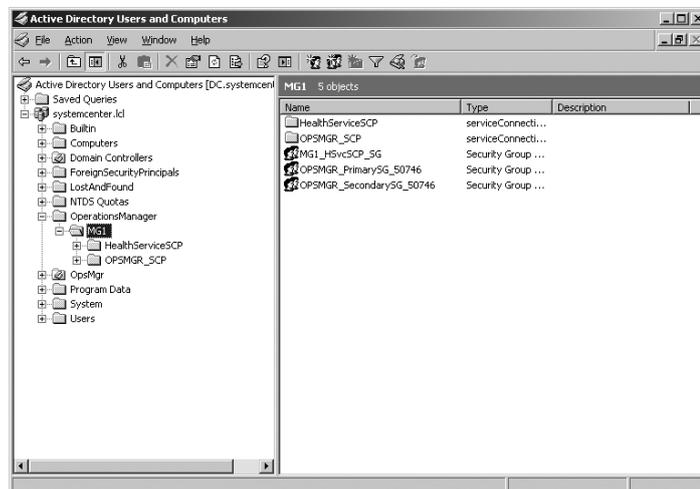
To run this tool, make sure that the system you run it from is running Windows 2000 Service Pack 4 or later. The tool has to run from a command prompt, and if any of the options entered on the command line contain spaces, you must enclose them in quotes.

If you want to run the `MOMADAdmin` utility to create a container named Exchange in the `zygort.1c1` domain and assign the ExchangeOpsMgr computer as the primary Operations Manager system and the Exchange Administrators as the administrative group, you would enter the following command at the command line:

```
MOMADAdmin.exe Exchange "Exchange Administrators" ExchangeOpsMgr zygort.1c1
```

After you have created the container, you can then view it in Active Directory Users And Computers or use ADSIEdit. Using ADSIEdit will allow you to view more information about the container and the objects that are contained within it. Figure 2.5 shows the OperationsManager container that you will find in Active Directory Users And Computers. When you first open the console, you will not see this container. That is because there is no reason for an operator or administrator to manipulate the objects within. However, for troubleshooting purposes, you may need to view the objects so that you can determine that they were created correctly or are configured correctly. If you select **View** ➤ **Advanced View**, you will be able to see the objects shown here.

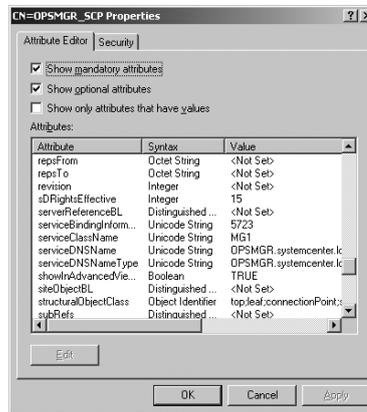
FIGURE 2.5
Viewing the
OperationsManager
container in Active
Directory using Active
Directory Users And
Computers



The objects that you see in the OperationsManager container represent the service connection point objects that are used by the management servers and the agents. The management points register themselves with Active Directory using the service connection points. The agents then look at the settings of the service connection points to determine how they need to function in the management group. The information they retrieve allows them to determine which management point they are supposed to communicate with, and in the event of a management point failure, which management point they will fail over to.

ADSIEdit will show you a little more information about each of the objects. Using this tool, you can view the properties of the objects and find out which management servers have successfully registered service connection point settings. Figure 2.6 show the properties of the OpsMgr_SCP container and the serviceDNSName attribute with the OpsMgr server name listed as a value.

FIGURE 2.6
ADSIEdit property view of a service connection point



To populate the groups that you find in this container, you must run the Agent Assignment and Failover Wizard. For more information about this wizard and how to use it, see Chapter 4.

Installing the Database

There are two ways that you can run the Operations Manager database: locally or remotely. Having the database on the same server as Operations Manager can be easier for some, since they know where all the services are running, but in the long run, the server will undergo a performance hit. Having both services vying for resources can be detrimental to your Operations Manager efficiency.

If your company is like many others, you probably have an administrative group that is responsible for managing the SQL servers. If this is the case, make sure that you interface with them during deployment. They will be responsible for configuring the SQL server and creating the appropriate databases. You will probably encounter several questions along the way. Have your design criteria prepared so that you can answer any questions that may arise. Remember—people tend to become territorial with their systems. If you can effectively present your case to them, they will probably work well with you.

The following installation steps are not meant to be an exhaustive discussion of SQL Server 2005. Instead, we are going to give you enough information to properly install it for use with Operations

Manager. It is understandable that you may not have time to become an expert on all of the Microsoft technologies, but you should seriously consider becoming familiar with SQL Server 2005 if you are going to support Operations Manager 2007 or Configuration Manager 2007. The more comfortable you are with all of the services that support your environment, the easier troubleshooting will become. There are several books on the subject, such as *Professional SQL Server 2005 Administration* (Wrox Press).

When you start the installer, the first thing you are presented with is an End User License Agreement (EULA). Accept the terms and move on to the prerequisites installer. The installer will scan your system to identify which of the prerequisites need to be installed and will ask you to install them. The prerequisites you see here are components that are required for SQL Server 2005's setup program. The two components that will probably appear are the Microsoft SQL Native Client and the Microsoft SQL Server 2005 Setup Support Files. Click the Install button to proceed.

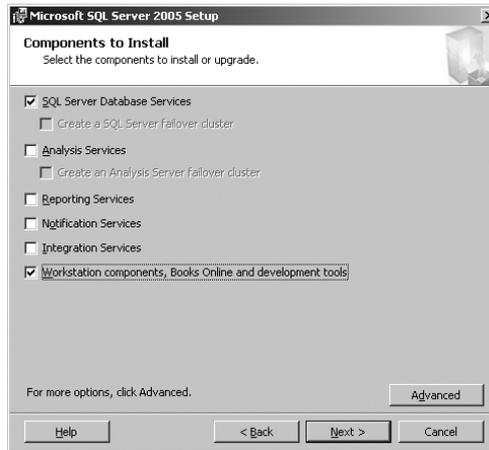
Once the prerequisite components are installed, the installer will run a prerequisite check again to make sure the components have installed as well as the current system configuration. This should not take long, and when it completes you will find yourself looking at the installation wizard page. Clicking the Next button will take you to the System Configuration Check. This is where the wizard checks to make sure that all the hardware and software requirements are met. As it runs, it will let you know whether your system meets the requirements, and it will also warn you about items that are not required but are highly recommended. View the messages to make sure that there is nothing that will affect your Operations Manager configuration. For the most part, if the installation requirements are not met for SQL Server 2005 itself, you will not want to install until you have fixed the issues. Operations Manager does not present any additional requirements for the database server other than the server requirements.

After making sure the system requirements are met, you are presented with the wizard pages that collect the details about your installation. The first page, Registration Information, is the typical information page that you will find in nearly every Microsoft installer. Here you can enter your name and company information so that your software can be personalized. This is also the page that contains the Product Key text boxes. Enter the product key that you have been given and click Next to move on.

On the next several screens are database options. The selections you make should be based on what you need for your Operations Manager database only. If you already have a SQL Server 2005 server in place, you may not be able to change some of the existing settings. For example, if the databases that are already in use cannot use the Local System account for the service account, you will have to use a domain-based account for your database services. Also, if you already allow SQL Authentication, you will probably not be able to change the authentication method to Windows Authentication only. Talk with your database administrators to find out how your servers are currently configured so that you know how you are going to use the database server for Operations Manager.

As shown in Figure 2.7, the Components To Install page presents you with several options to choose from. You do not have to install all the components on this page to support Operations Manager. At a bare minimum, you should install the SQL Server Database Services as well as the Workstation Components, Books Online And Development Tools option. When you are selecting your choices, keep in mind that there may be several subcomponents that you do not want to install listed on a page. Although it may be easy to select the two options here, you may want to control the components and subcomponents a little more efficiently.

FIGURE 2.7
Component selection



Click the Advanced button, and you will see a page that allows you to have more granular control over the components and subcomponents that you are going to install. As Figure 2.8 shows, you will see the components that are installed when you select the options on the previous page. Chances are you will not need replication services or some of the database design tools that are included with these options. If you want to streamline your database server, you can expand the components and choose not to install some of the features. If your server is only supporting your Operations Manager infrastructure, you will not have to install the Replication subcomponent under Database Services. You can also safely eliminate the Business Intelligence Development and Software Development Kit subcomponent under Client Components. Take note that if the SQL server is also going to host Reporting Services, you should keep the Business Intelligence Development selected so that you can create reports.

Notice that the Feature Selection page that appeared when you clicked the Advanced button is the only way to specify an alternate location for the installation path. If you would like to use a location other than the default of `C:\Program Files\Microsoft SQL Server`, choose it now, then click Next to move on to the Instance Name page. If this is the first time you have installed SQL Server 2005 on this server, you must leave the Default Instance selected and move on. If this is a new instance of SQL Server 2005, specify a name for your server.

Next up is the Service Account settings. On this page you have the opportunity to specify the account that will be used by the SQL services. At this point you should consult your domain administrators to make sure that they have provided you with the appropriate accounts to use for the services. As you can see in Figure 2.9, you can select the Customize For Each Service Account check box and you will be able to provide a different account for each of the services. You then have the option of using the Local System account or providing a domain account. You should remember that the accounts you use will have unimpeded access to the services. If the account becomes compromised, the attacker will have the same level of permissions and rights as the account. For this reason alone, you should use a different account for each of the services. That way, if one account becomes compromised, the other services will not be affected.

The next screen that you will see is Authentication Mode. For Operations Manager you will want to make sure that you have selected Windows Authentication Mode. This allows the management

servers to use their domain-based authentication in order to access the SQL Server. When you click Next, the Collation Settings page appears. You can click Next on this page, accepting the defaults.

At this point, you will find yourself making choices about how you should answer the error reporting options. As with all of the latest Microsoft products, whenever the product has a problem and an error condition arises, a debugger takes over and collects information about the error condition. These error condition results can be sent to Microsoft so that they can use the data to improve the stability and functionality of the program. You have two options at this point, as shown in Figure 2.10. If you choose the first option, you can have the error reports automatically delivered. The second option allows you to send basic usage reports that detail how SQL is used and how it is functioning.

FIGURE 2.8
Selected components

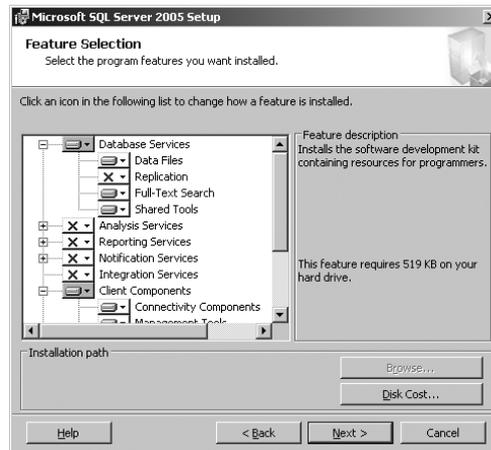


FIGURE 2.9
Service accounts

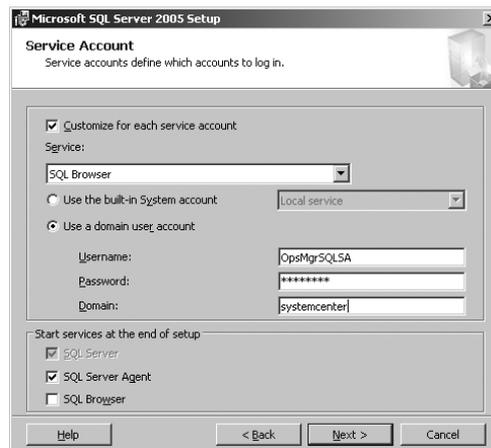
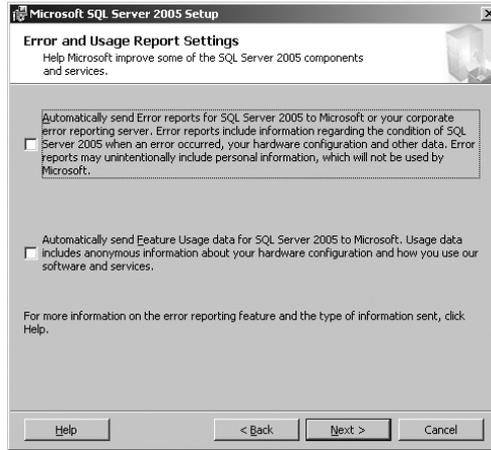


FIGURE 2.10
Error and usage
reporting options



Before you bypass these settings, you should take something into consideration. Operations Manager 2007 has introduced Agentless Exception Monitoring. By setting a group policy option, you can have servers within your organization send these reports to your management servers instead of Microsoft. You can then control how the reports are disseminated, either having them available to your staff or forwarded on to Microsoft. We will discuss Agentless Exception Monitoring in Chapter 8.

You are now ready to install the database server. The final page of the wizard will detail the options you have selected and give you a summary of what will occur once you click the Install button. This is your last chance to make changes before the server is installed. Of course, you can always run the installation wizard again to alter your server, but you should try to get it right the first time. After installing the database server, you are ready to install the Operations Manager database.

INSTALLING THE DATABASE ON A REMOTE SQL SERVER

When you start the installation wizard, you are presented with the usual EULA, so make sure you read through all of the legalese. Then select the I Accept The Licensing Terms And Conditions check box and click Next to move on. The Product Registration page appears next, prompting you to enter your username, your organization name, and the product key. Fill these in and then click Next.

Since we are adding the Operations Manager database to a remote SQL server, we want to make sure that none of the other Operations Manager components are installed. When the Custom Setup page appears, make sure that you pull down the installation menu options next to the Management Server and User Interfaces components and select This Component Will Not Be Available, as shown in Figure 2.11. That will leave only the Database component set to install. Notice that you can also choose the location where the Operations Manager files will be installed.

On the Management Group Configuration page, you must enter the name of the management group that the database server will support. The other configuration setting on this page allows you to choose the group that will be considered Operations Manager administrators. The members of the security group that is set here will have the ability to open the Operations Manager console and have unrestricted access to all the Operations Manager functions in the management group. You must supply this information here so that the data can be entered into the database. You will also enter this information when you install the Operations Manager server. Clicking the Browse button on this page, shown in Figure 2.12, will bring up an Active Directory selection page, where you can select the group you want to make the Operations Manager administrator group.

When you install the Operations Manager database, the setup program will automatically determine the database instances that are available on the database server. As you can see in Figure 2.13, the SQL Server Database Instance page allows you to use the pull-down menu to choose the instance you plan to use.

FIGURE 2.11
Selecting only the MOM Database component

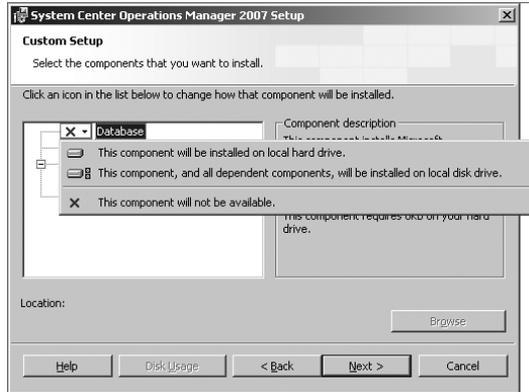


FIGURE 2.12
Selecting the Operations Manager administrators

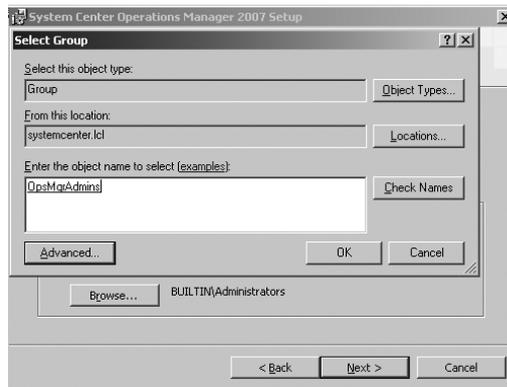


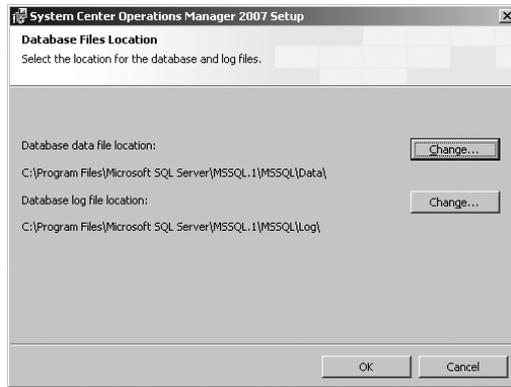
FIGURE 2.13
Selecting the database instance



When you reach the Database And Log File Options page, you will need to consult your design documentation. The first option to consider here is the Database Size entry. The default size for the database is 1 GB. Remember, this is the initial size of the database, as well as the criteria for log-file retention size. The log file size will be 20 percent of the initial database size. The maximum log-file size, if you keep the database size at 1 GB, will be 20 MB.

The default name of the database is OperationsManager. The default location for the log files is C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Log. The database data files default location is C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Data. However, as you can tell from the Database And Log File Options page, you are not allowed to make changes to these components. If you click the Advanced button, you will be presented with the Database Files Location page, shown in Figure 2.14. Here you can edit the locations and names of the database components. Change them as necessary for your environment.

FIGURE 2.14
Database Files
Location page



When you are setting the location of the database data files and log files, remember that you will receive a performance boost on your SQL server if you move the log files to a different physical drive than the drive where the data files are located. If you move them to a separate partition on the same drive, you will not see any performance gains. If the drive that you place the log files on is dedicated to just serving the log files, you will receive an even better performance boost. The other advantage of separating the database and log files onto separate drives is fault tolerance. If you were to lose the drive that contained the database, you could restore the database onto another drive; then the log files could be used to bring the database back to the point it was in at the time of failure. In that scenario, you should not have to worry about losing data.

On the SC Error Reporting page, specify whether you want to send Operations Manager error reports to Microsoft. If you do, you have the option to automatically send the errors, or you can queue the errors and then approve which errors will ultimately be sent. Once you have selected your reporting choice, you are ready to perform the installation of the database. Click Install to continue.

INSTALLING A LOCAL DATABASE

In all fairness, there isn't much difference between installing a database locally and installing it remotely. Step one is still installing SQL Server 2005 on a server, but this time, you need to make sure it is installed on the same server that will host Operations Manager. Once SQL Server 2005 is installed, you will then start the Operations Manager setup and select to install the database at the same time you install the rest of the service. The options that were specified in the previous section still apply; they are just presented during the setup of Operations Manager.

INSTALLING A CLUSTERED DATABASE

Installing the database on clustered hardware is not overly difficult. You do have to perform a few additional actions in order to install everything correctly, though. You will have to start by installing SQL Server 2005 SP1 or later in the clustered environment. Once you do, you can perform the installation of the Operations Manager database.

The account that is used when installing the database not only has to have administrator privileges on the SQL server instance; it will also need to be a local administrator for the cluster node. Once you're logged on, the actual installation steps are similar to installing the database on a remote SQL server, but you will need to make sure that you install the database to the SQL cluster virtual server instead of the node name. Once you've defined the server name, you can proceed with the installation as normal.

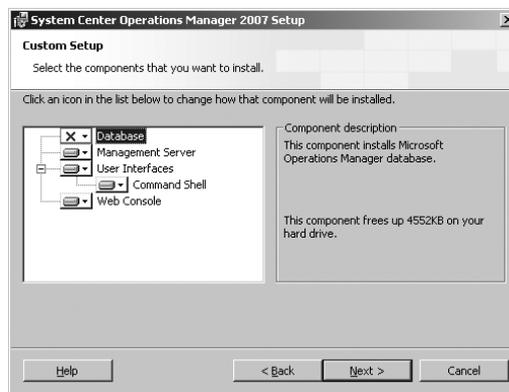
Installing Operations Manager 2007

Now that everything else has been configured, it is time to turn our attention to the Operations Manager management server. Make sure you have your planning documentation handy, as you will need to make sure you answer the prompts during the wizard. Some of the settings that you configure cannot be modified once you have finished the install.

As with all of the Microsoft products, once you have clicked Next to go past the initial splash page that welcomes you to the product's setup wizard, you will have to agree to the EULA to go on. After you do so, you will be prompted to enter your username, organization name, and product key. After you have entered all of that information, click the Next button to begin configuring the individual components of Operations Manager.

On the Custom Setup page, as shown in Figure 2.15, you have the option to select each component for Operations Manager. In the previous section, we detailed the steps to install the Operations Manager database on a remote SQL server. If you have already performed that step, choose the pull-down menu next to Database and select This Component Will Not Be Installed. If the database server and the management server are one and the same, you can leave the option at its default setting. The other options should be left selected, depending on your needs. You need to make sure the Management Server option remains selected, but you are not required to leave the other two components.

FIGURE 2.15
Selecting components



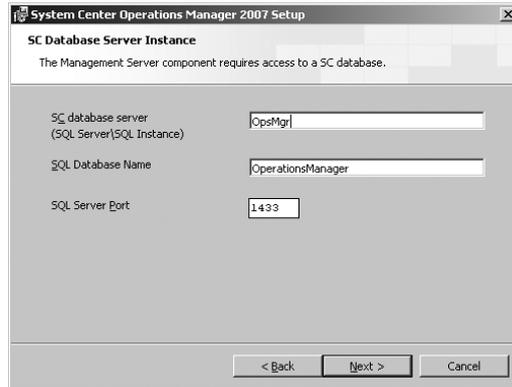
Both of the remaining components are administration tools. The User Interfaces selection installs the console on the server. The component beneath the User Interfaces option is the Power-Shell Command Shell. To tap into the full potential of the Operations Manager server, you should

install the PowerShell component. Microsoft has invested a lot of time and energy in their new script-based management tool—and for good reason. It works well for most administrators and it doesn't have a large learning curve as some other scripting tools do. For more information on using PowerShell to manage Operations Manager, see Chapter 14.

Before we move on, we'd like to mention that if you decide to use a local database, you will need to make sure you select the option for the database. Of course, this is assuming that SQL Server 2005 has already been installed on the server prior to you running the Operations Manager setup. After making your installation choices, click Next and review the EULA. Select the I Accept The Terms In The License Agreement option and click Next to move on to the next setup option.

If you are not installing a local database, the next option you will see is the SC Database Server Instance, shown in Figure 2.16. The settings here allow you to specify the SQL server that will host the Operations Manager database, as well as the name of the database and the port used to communicate. By default, the database name is OperationsManager. If you specified a different name during the installation of the database on the SQL server, make sure you enter the name correctly. SQL communicates across port 1433. If you decided to use a different port, make sure you change this entry or you will not be able to communicate with the database.

FIGURE 2.16
Selecting the database instance to use



As you can see in Figure 2.17, the next page in the wizard allows you to specify the account that will be used as the management server action account. As we mentioned in the “Creating Accounts” section, you have the option of creating a domain account to use as the management server action account, or you can use the Local System account. Once you have entered the account information, click Next to move to the SDK And Config Service Account information. Just as you did with the management server action account, specify the credentials that will be used and click Next. If you are sure all the options you entered are correct, you can click the Install button to install Operations Manager.

After the installation is complete, you have one more important setting to configure. If you had created the service connection point containers in Active Directory, make sure that the management point will register itself. By default, Active Directory integration is not enabled, although the root management server will be listed simply because the MOMADAdmin utility created the appropriate settings. On all your management servers, you should configure the Registry entry that controls integration.

After you open the Registry Editor, navigate to HKEY_Local_Machine\SYSTEM\CurrentControlSet\Services\HealthService\Parameters\ConnectorManager and double-click the EnableADIntegration key. Change the Value to 1 and click OK.

FIGURE 2.17
Choosing the
management server
action account



Installing the Root Management Server on Clustered Hardware

As we mentioned at the beginning of the chapter, the RMS is the only management server that hosts the SDK and Config services. Without these services, the management group cannot function. Having a single point of failure is a problem for any enterprise-class service. When the server in question is responsible for monitoring your organization, allowing it to be a single point of failure is completely unacceptable. To alleviate this situation, the RMS can be installed on clustered hardware.

As you may have surmised, installing the RMS on clustered hardware is not as straightforward as installing on a single server. Several steps are involved in performing the installation, which includes not only installing the software on each node, but also copying the encryption key to the additional nodes and then adding the additional nodes to the clustered solution.

CLUSTER RESOURCES

If you are unfamiliar with the clustering service in Windows Server 2003 Enterprise Edition, you can find a plethora of good information on the Microsoft site. If you navigate to <http://www.microsoft.com/windowsserver2003/technologies/clustering/resources.mspx>, you will find white papers and how-tos that will get you started. Remember that a clustered solution will increase the total cost of a project, but the security in knowing that you have a higher level of availability for your RMS may well outweigh the additional costs involved.

The first step to installing on clustered hardware is to prepare the cluster. We are going to assume that you already know how to configure your clustered hardware, since going into the details of setting up a clustered system is beyond the scope of this book. And since there are different methods of setting up the shared storage and network connections, we are not going to assume that we know how you've built your physical systems that make up the clustered nodes. Instead, we'll start at the point of preparing the cluster for installation.

Because any of the nodes within the cluster can act as the RMS, you must make sure that each node has the appropriate accounts added to the local administrators group. To do so, go into computer management on each one of the nodes and add the domain account that you use for Operations Manager into the local administrators group. You also need to add the cluster service account into the Operations Manager administrators security group. Doing so allows the cluster service to work

with the appropriate services in Operations Manager to guarantee that the RMS will fail over properly and function correctly in the cluster. The service account that you are going to use for the SDK and Config services also has to be added to the local administrators group. Once you have added the appropriate accounts in the local administrators group, you are ready to create the resources that will be used in the cluster.

You have to create the resources that you'll use for your RMS cluster before you can proceed with the installation. These groups and resources are used in the cluster to control the failover of any of the RMS services. Unlike with Exchange Server or SQL Server, none of the groups or resources are created during the installation of the RMS. Instead, you have to create all these objects manually.

CREATING THE CLUSTER GROUP AND INITIAL RESOURCES

To create the group that will be used for failover, open Cluster Administrator and right-click on Groups. Then select New ➤ Group from the context menu. You can name this group anything you like, but you should probably choose something that fits the standards for your organization. For our example, let's name ours **RMS cluster group**.

After you have created the group, the next step is to create the resources that will be used for the RMS cluster. You'll create three resources: an IP address, a network name, and a physical disk. You should already know the physical disk that you'll use to install the RMS software on. To create each of the resources, follow these steps.

First, to create the Physical Disk resource:

1. Right-click the group and select New ➤ Resource.
2. Name the new resource **RMS Physical Disk**, as shown in Figure 2.18.
3. From the pull-down menu, choose Physical Disk and click Next.
4. Choose all of the cluster nodes that will act as RMS failover nodes in the Available Nodes list and click the Add button to move them into the Possible Owners list, as shown in Figure 2.19.
5. The Physical Disk resource does not rely on any dependencies, so you can click Next on the Dependencies page.
6. On the Disk Parameters page, shown in Figure 2.20, use the Disk drop-down to specify the physical disk letter that you wish to use, and then click Finish.

FIGURE 2.18
RMS Physical Disk

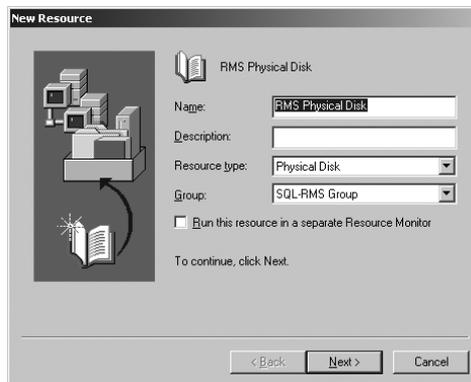


FIGURE 2.19
Selecting available
nodes

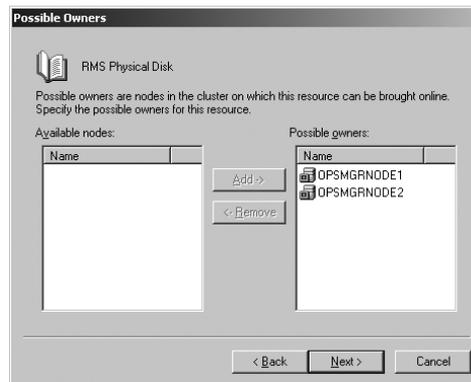
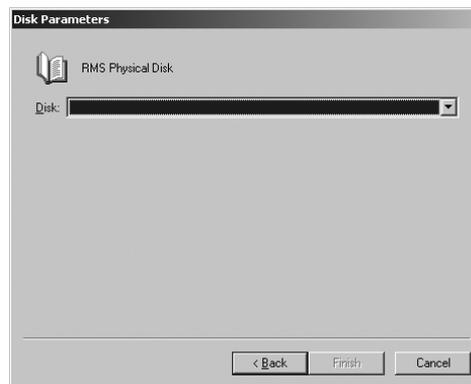


FIGURE 2.20
Choosing the drive



Next, to create the IP Address resource:

1. Right-click the group and select **New > Resource**.
2. Name the new resource **RMS IP Address**, as shown in Figure 2.21.
3. From the Resource Type drop-down list, choose IP Address and click Next.
4. Make sure all of the cluster nodes are selected as possible owners, and then click Next.
5. The IP Address resource does not rely on any dependencies, so you can click Next on the Dependencies page.
6. On the TCP/IP Address Parameters page, shown in Figure 2.22, enter an IP address and a subnet mask that is available on the public network.
7. From the Network drop-down list, choose your public network and click Finish.

FIGURE 2.21
Naming the IP Address
resource

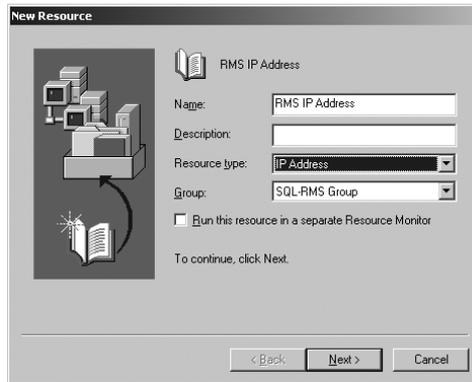
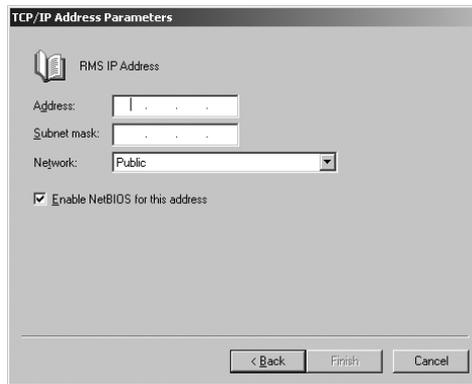


FIGURE 2.22
Entering the IP
Address parameters



Finally, to create the Network Name resource:

1. Right-click the group and select **New > Resource**.
2. Name the new resource **RMS Network Name**.
3. From the Resource Type drop-down list, choose **Network Name**, as shown in Figure 2.23, and click **Next**.
4. On the Possible Owners page, make sure all of the cluster nodes are selected as possible owners.
5. On the Dependencies page, shown in Figure 2.24, select the RMS IP Address resource from the Available Resources list and click the **Add** button to move it into the Resource Dependencies list. Then click **Next**.
6. On the Network Name Parameters page, shown in Figure 2.25, the Name field within the parameters should include the name of the resource. For our example, let's call ours **OpsMgrRMS**.
7. Make sure the **Enable Kerberos Authentication** check box is selected, and click **Finish**.

FIGURE 2.23
 Naming the Network
 Name resource

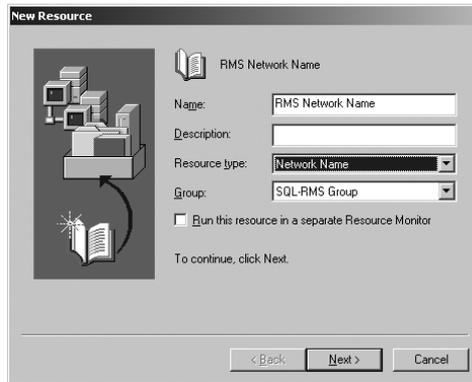


FIGURE 2.24
 Defining the
 dependencies

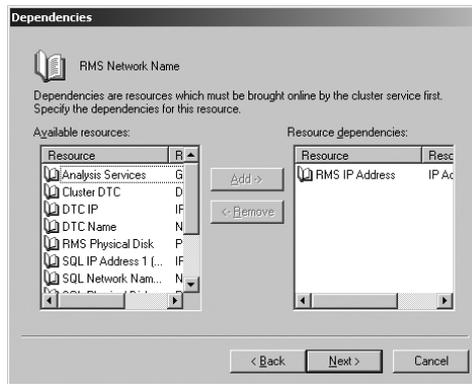
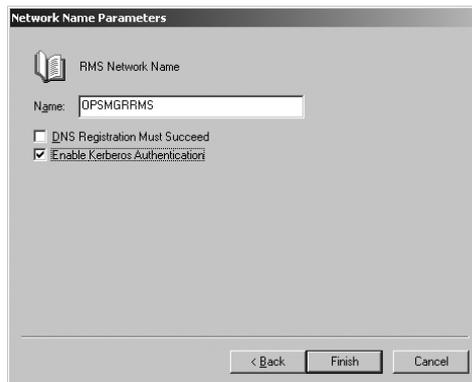


FIGURE 2.25
 Entering the name for
 the resource

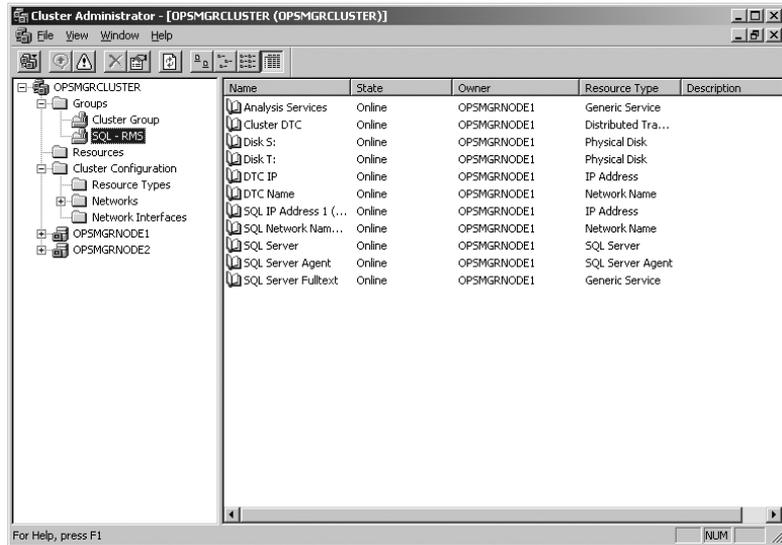


After you have created all the resources, right-click the group and select Bring Online.

INSTALLING THE RMS

At this point, you are ready to install the RMS. Just as with a standard installation, you must verify the prerequisites. You must run the prerequisite viewer on the node that owns the group you are planning to use for the RMS. You can find out which node owns the resource by opening the Cluster Administrator and looking at the group information. Figure 2.26 shows that the group is owned by OpsMgrNode1. If you want to run the prerequisite check on each node, move the group to each node and perform the prerequisite check again.

FIGURE 2.26
Determining node ownership



Once you have completed the prerequisite check, make sure that you are logged on to the node that owns the group. The installation program must have access to all the resources, especially the physical disk. Note that you will be installing Operations Manager on each of the nodes that will participate in the cluster. So your first step is to bring the group online. Launch Cluster Administrator and open Groups. Right-click the RMS cluster group that you created and select Bring Online. You are now ready to start the installation.

The steps to install Operations Manager at this point are not all that different from a standard install. However, if you plan on installing the SQL Server on the cluster, it must be installed separately. The setup program isn't configured to install SQL in this manner, so it is up to you. Given that, after you start up SetupOM.exe, click the Install Operations Manager 2007 link; the only two components that you should select at this time are the Server and Console items.

The steps for the install are as follows:

1. Run SetupOM.exe from the CD or network location where the installation files are located.
2. Click the Install Operations Manager 2007 link.
3. When the wizard starts, click Next.

4. Agree to the EULA and click Next.
5. Enter your name, organization, and product key and then click Next.
6. When the Custom Setup page appears, select This Component, And All Dependent Components, Will Be Installed On The Local Disk Drive for both the Management Server and User Interfaces options.
7. Configure the Database, Command Shell, and Web Console options to This Component Will Not Be Available, as shown in Figure 2.27, and then click Next.
8. Leave the default installation location, then click Next.
9. When the SC Database Server Instance page appears, as shown in Figure 2.28, enter the name of the SQL Server that you are going to use as your database server. If the database is the default instance, you can simply enter the database name. If it is not the default instance, you must enter the server / instance name. If your database is clustered, enter the name of the virtual server.
10. You can name the database to fit your organization's standards or leave the default name OperationsManager.
11. Make sure the port number is 1433 and click Next.
12. Enter the username and password that you created for the Management Server Action account, and click Next.
13. Enter the username and password for the SDK and Config Service account and click Next.
14. Specify your preferences on the Customer Experience Improvement Program page and click Next.
15. Specify your preference on the Microsoft Update page and click Next.
16. When the Install The Program page appears, click Install.
17. After installation has completed, click Finish to end the wizard and start the Operations Console.

FIGURE 2.27
Configuring the
components to install

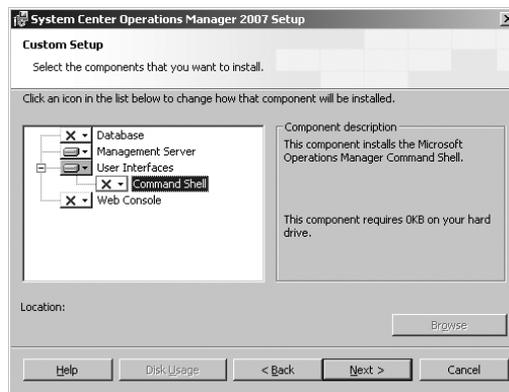
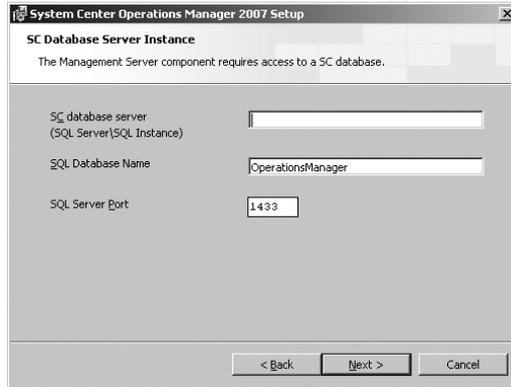


FIGURE 2.28
SC Database Server
Instance page



INSTALLING THE SECONDARY MANAGEMENT SERVERS

Now that you have installed the RMS, you must install Operations Manager on the other nodes of the cluster. As you do so, you will not find anything out of the ordinary for the setup. None of the settings will identify that you are installing onto a cluster. As a matter of fact, the additional nodes will be considered secondary management servers until we specify that they are to be failover nodes for the RMS. As you install each node, install only the same options that you selected for the RMS. After each node's installation is complete, the interesting part begins.

CREATING THE CLUSTER RESOURCES FOR OPERATIONS MANAGER SERVICES

We created the initial clustered resources before installing the management servers, but there are still a few resources left to install. The Health Service and the SDK and Config Service both need cluster resources created for them so that the cluster service can monitor their availability. Both of these resources are Generic resource types, which means that there is only rudimentary monitoring available for the service.

To create the new resources:

1. Open Cluster Administrator.
2. Right-click your RMS group and select New ► Resource.
3. Name the resource **RMS Health Service**.
4. Select Generic Service from the Resource Type drop-down list, as shown in Figure 2.29, and click Next.
5. Select all nodes as possible owners and click Next.
6. Select the Physical Disk and Network Name resources for your RMS group and add them to the Resource Dependencies list, as shown in Figure 2.30; then click Next.
7. Enter **HealthService** as the Service Name, as shown in Figure 2.31.
8. Select the Use Network Name For Computer Name check box.

9. Click Next and then click Finish.
10. Right-click your RMS group and select New ➤ Resource.
11. Name the resource **RMS Config Service**.
12. Select Generic Service from the Resource Type drop-down list and click Next.
13. Select all nodes as possible owners and click Next.
14. Select the Physical Disk and Network Name resources for your RMS group and add them to the Resource Dependencies list; then click Next.
15. Enter **OMCFG** as the Service Name, as shown in Figure 2.32.
16. Select the Use Network Name For Computer Name check box.
17. Click Next and then click Finish.
18. Right-click your RMS group and select New ➤ Resource.
19. Name the resource **RMS SDK Service**.
20. Select Generic Service from the Resource Type drop-down list and click Next.
21. Select all nodes as possible owners and click Next.
22. Select the Physical Disk and Network Name resources for your RMS group and add them to the Resource Dependencies list; then click Next.
23. Enter **OMSDK** as the Service Name.
24. This time, do not select the Use Network Name For Computer Name check box.
25. Click Next and then click Finish.

There is still some configuration required in order to complete the clustered setup, so make sure you do not bring the group online yet. Even though the resources are now available, we still need to make sure that all of the nodes can access the database. To do so, we must make sure that every node has the appropriate encryption keys and are registered as the RMS.

FIGURE 2.29
Creating the generic
service

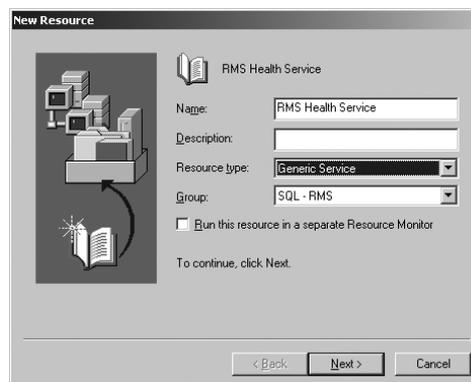


FIGURE 2.30
Identifying dependencies

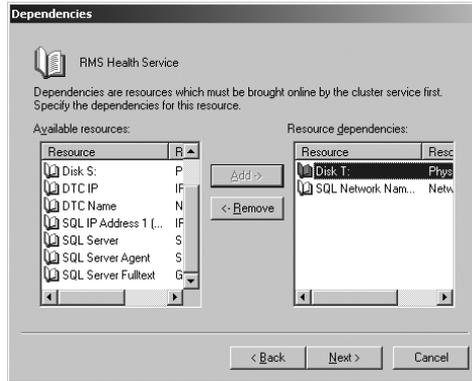


FIGURE 2.31
Naming the Health service

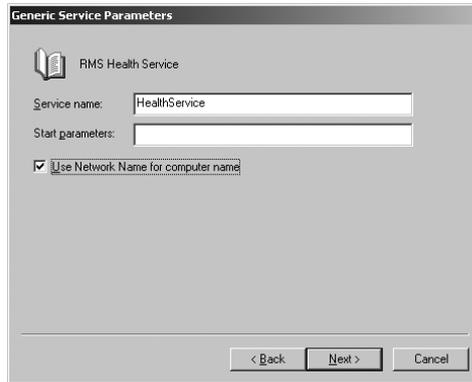
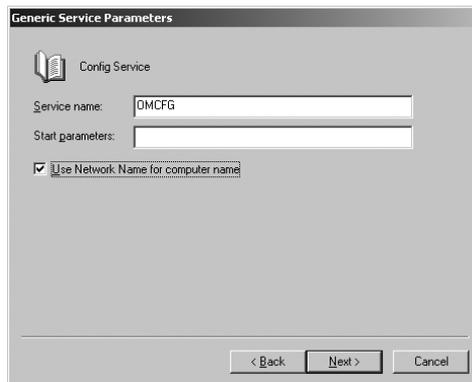


FIGURE 2.32
Naming the Config service



CREATING THE RMS VIRTUAL SERVER

We use the term *virtual server* in conjunction with the cluster service. Each node in the cluster requires the ability to become the owner of the resources at one point or another. If you were to use the name of the node itself, clients would become confused when failover occurred and the name of the server changed. Instead, when a node becomes the owner of a group, the network name, or virtual server, is the name that is used by the node.

This introduces a level of complexity with Operations Manager due to the fact that the encryption key for the RMS is tied to the first node that you installed. All of the nodes will need access to this same key so that they can access the data when required. As you prepare to configure each node for the encryption key, make sure that you have a shared folder available that each of the nodes can access. This is where you will place the key that you are going to export. If you have not already done so, add the cluster service account as a member of the Operations Manager Administrators group.

Before you create the virtual server, perform these actions:

1. Create a shared storage location that is available to all nodes and secure it so that only the Operations Manager Administrators group has access.
2. Add the Cluster Administrator account to the Operations Manager Administrators group, as shown in Figure 2.33.
3. Copy the `SecureStorageBackup.exe` and `ManagementServerConfigTool.exe` programs from the Support Tools folder on the Operations Manager installation disk to the shared storage location.

To export the RMS key:

1. Log on to the RMS node as the cluster service account.
2. Connect to the shared storage location and copy `SecureStorageBackup.exe` and `ManagementServerConfigTool.exe` to the installation location for Operations Manager, typically found in `Program Files\System Center Operations Manager 2007`.
3. Log on to each node as the cluster service account and copy `SecureStorageBackup.exe` and `ManagementServerConfigTool.exe` to the installation location for Operations Manager, typically found in `Program Files\System Center Operations Manager 2007`.
4. On the RMS node, open a command prompt and navigate to the Operations Manager installation directory.
5. Enter the following command, substituting the **servername** and **sharename** for your shared storage location: `securestoragebackup.exe Backup \\servername\sharename\OpsMgrClusterKey.bin`.
6. Enter a password that is at least eight characters long.
7. On each remaining node, open a command prompt and navigate to the Operations Manager installation directory.
8. Enter the following command, substituting the **servername** and **sharename** for your shared storage location: `securestoragebackup.exe Restore \\servername\sharename\OpsMgrClusterKey.bin`.

9. Enter the password for the file in order to import the key.

To create the virtual server:

1. Prior to creating the virtual server, you need to back up the SQL database in case of corruption. Do not skip this step!
2. On the RMS node, open a command prompt and navigate to the Operations Manager installation directory.
3. Enter the following command: **ManagementServerConfigTool.exe InstallCluster /vs:clusternetworkname /Disk:physicaldisk.**

To configure the remaining nodes:

1. Log on to the node with the cluster service account.
2. Open Cluster Administrator and move the group to the node you are logged on to.
3. Open Administrative Tools ➤ Services and set the OpsMgr SDK service to start automatically.
4. Start the OpsMgr SDK service.
5. Open a command prompt and navigate to the Operations Manager installation folder.
6. Enter the following command: **ManagementServerConfigTool.exe AddRMSNode /vs:clusternetworkname /Disk:physicaldisk.**
7. When all of the nodes have been configured, open Cluster Administrator and bring the RMS group online.

At this point, the RMS is configured as a highly available service and you should see all services started in the Cluster Administrator, as shown in Figure 2.34. If there is a problem with a node that is hosting the RMS, failover will occur and a remaining node will take over. If everything is working as it should, you will see the virtual server name in the Administration workspace appearing in a Healthy state. Each of the nodes within the cluster should appear in a Healthy state in the Agentless Managed node.

FIGURE 2.33
Adding the Cluster Administrator account to the Operations Manager Administrators group

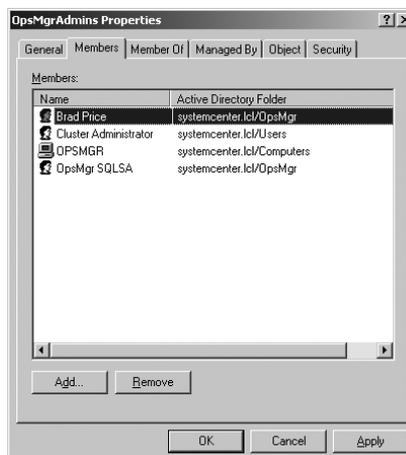
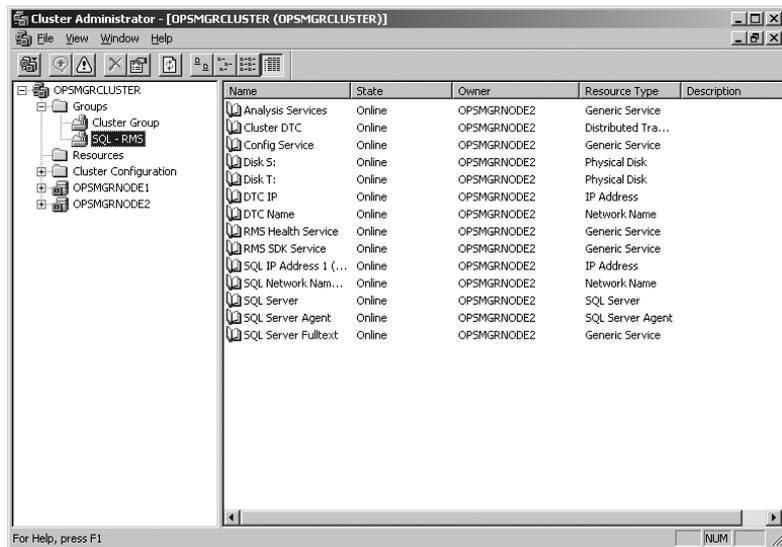


FIGURE 2.34
All services running in
Cluster Administrator



Verifying the Installation

After you have installed the database and the management server, you can run a couple of tests to make sure that everything is working correctly. The easiest test you can run is to start the console. As a matter of fact, after the installation completes, you have the option to start the console when you click Finish. If the console appears and you can open each of the views, then chances are good that everything installed correctly. There are a few other things you should check to verify that the install completed successfully.

EVENT VIEWER

Part of the installation consists of creating a new event log type that appears within the Event Viewer on the management server. This new event log, named Operations Manager, will appear in the list and detail events as they are recorded from the Operations Manager functions. Look through the log entries to see if there were any complications during setup. You should find, at the very least, two entries that specify the OpsMgr Config Service, Event ID 29000, and OpsMgr SDK Service, Event ID 26361, started successfully. Other entries within this log will help you determine the overall health of the service as well as identify additional steps you may need to take to ensure that everything is working correctly.

INSTALLATION LOG FILES

During installation, Operations Manager writes three log files containing the installation steps. You can find these installation logs in the directory where you chose to install Operations Manager, which by default is %Program Files%\System Center Operations Manager 2007. The first of these logs records the installation of the Configuration Service, called Microsoft.MOM.ConfigService.InstallLog. The second log file records the database configuration, called Microsoft.MOM.DatabaseWriteModules.InstallLog. The final log file, Microsoft.MOM.Sdk.ServiceDataLayer.InstallLog, records the changes made to the system as the console and supporting tools are added.

INSTALLATION DIRECTORIES

Beneath the installation directory you will find the supporting directories for Operations Manager.

AgentManagement Includes the agents installation programs for each supported operating system platform

Config Service State Includes directories that hold configuration data that is used to control how the management server and agents are configured

Health Service State Includes directories that control how the agents and management servers interact when events and alerts are raised

SDK Service State Used to host configuration files that control how the management consoles are configured

Service Configurations Contains default configuration files for each service that is running on the management server

Helper Objects Contains installer files for additional tools that can be used when monitoring your systems

Help Contains the help documentation

One other directory exists within this list, named after the two-digit code for the language version that is installed. The English version will be named `en`, the Spanish version will be named `es`, and so on. In this directory you will find the dynamic link libraries (DLLs) that are responsible for controlling the display on the consoles, database entries, and reports.

The Bottom Line

Identify the system requirements. Although it may not seem that the system requirements for installing Operations Manager represent the hardware you would want to run your Operations Manager system on, you should be aware of the hardware and software requirements so that you know which systems in your environment can support it.

Install the database on a remote SQL 2005 server. Installing the database on a remote SQL 2005 server requires that you install SQL 2005 and then run the Operations Manager setup program on that server. When you run the setup program, you will only choose the MOM Database option so that the database can be created and used by the server that will run Operations Manager.

Install the database locally on the Operations Manager system. In some cases the Operations Manager server will host the database. Although this will impact the server's performance, some smaller installations may find this the best method to keep their total number of servers to a minimum.

Install Operations Manager 2007. The steps to install Operations Manager are straightforward, but you need to make sure that you have identified the accounts that will be used by the services, and you will have to know where the database is located within your infrastructure.

Verify the installation of Operations Manager 2007. Log files are generated during the installation of Operations Manager. They will tell you what occurred during the setup. You can also check the installation directory to make sure that all the required files and directories are available.

Master It! Microsoft has made their Virtual Server 2005 R2 product available to download and use for free. As we go through the chapters in this book, we are going to assume that you are using Virtual Server 2005 R2 to build your own virtual test environment. You could very well be using Virtual PC 2004 or VMWare Server, all available for free, but we are going to provide steps for only the Virtual Server platform.

With that said, what follows are steps to set up and configure three systems: one providing Active Directory and DNS functionality, a second server that will become your database server, and a third that will run Operations Manager 2007. Another assumption is going to be made at this point regarding memory. For best results, the physical system that you are running should have 4 GB of memory. We realize that not everyone can afford a system that has that much memory, so you should allocate as much memory as you can to your management server, since it will be the system doing most of the work. Later we will be setting up an additional machine that will become our managed system. If you have two or more servers or workstations that you can use to run Virtual Server, you may want to divide up the virtual machines instead of running them all on one system. That way, you can allocate more memory to each one.

If you do not have access to Windows Server 2003 R2, SQL Server 2005, or System Center Operations Manager 2007, you can contact Microsoft for evaluation editions of them. Although they will expire, you should be able to perform all of the exercises in this book during the amount of time that they give you.

Creating the Virtual Machine Base Drive:

1. Open Virtual Server Administration Website.
2. Create a new virtual disk by clicking Virtual Disks > Create > Dynamically Expanding Virtual Hard Disk.
3. Select the path to your new virtual disk in the Location drop-down box and then type the name **OpsMgrBase** in the Virtual Hard Disk File Name text box.
4. Enter **32** in the Size text box and leave the Units as GB. Click Create.
5. Create the new virtual machine by clicking Virtual Machines > Create.
6. In the Virtual Machine Name text box, type the full path to the location where you created the virtual hard disk and append the name **OpsMgrBase**.
7. In the Memory text box, enter **750**, choose Use An Existing Virtual Hard Disk, and select the **OpsMgrBase.vhd** file.
8. Select External Network from the Virtual Network Adapter options. Leave all the other options at their defaults and click Create.
9. If you are not at the OpsMgrBase configuration screen, click Virtual Machines > Configure > OpsMgrBase.
10. Click the Hard Disks link on the OpsMgrBase Configuration section and make sure the Enable Undo Disks option is selected; then click OK.
11. Insert the Windows Server 2003 CD into the host machines CD drive.
12. In the OpsMgrBase Configuration section, click CD.DVD.

13. Make sure the Physical CD/DVD Drive option is selected and that the CD drive for the host machine is selected. Click OK.
14. Start the OpsMgrBase virtual machine by clicking the arrow next to OpsMgrBase in the OpsMgrBase Status section and choosing Turn On.
15. Open the Virtual Machine Remote Control Client and select the OpsMgrBase virtual machine.
16. Press Enter when the Welcome To Setup screen appears, then press the F8 key when you see the license agreement.
17. To create the installation partition, press Enter to use the full size of the drive; then select the option Format The Partition Using The NTFS File System (Quick).
18. After the format completes, the system files will be copied to the virtual drive, the system will reboot, and the Windows Server 2003 setup will begin.
19. On the setup screens, provide the following information:
 - ◆ Enter your name and organization.
 - ◆ Enter your product key on the Product Key page.
 - ◆ Select Per Device Or Per User on the Licensing Modes page.
 - ◆ Enter **OpsMgrBase** in the Computer Name field and enter a password for the administrator account.
 - ◆ Specify your Date and Time settings.
 - ◆ Choose Typical Settings when the Network Settings appear.
 - ◆ Leave the Workgroup Or Computer Domain options at their defaults.
20. Once the system starts up, log on as administrator. If you are asked to install the additional tools for Windows Server 2003 R2, or to update the server, do so now.
21. Switch back to the OpsMgrBase Configuration web page and click the CD/DVD link.
22. Click the Known Image Files radio button and choose the VMAdditions.iso image; then click OK.
23. Switch back to the Remote Control client and follow the wizard prompts, selecting the defaults to install the Virtual Machine Additions. Then reboot the virtual machine.
24. Make sure that the Windows Server 2003 CD is in the host computer's CD drive; then switch to the configuration web page and select the Physical CD/DVD Drive option. Make sure the physical drive with the Windows Server 2003 CD is selected.
25. On the remote OpsMgrBase remote control screen, open a Windows Explorer window and copy the I386 directory from the CD to the C: drive's root.
26. Open REGEDIT from Start ► Run. Navigate to HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Setup and change the SourcePath value to C:\. Do the same for the ServicePackSource value.
27. In Windows Explorer, navigate to the Support\Tools directory on the Windows Server 2003 CD and double-click deploy.cab. Extract the contents to C:\deploy.

28. In the C:\deploy directory, double-click `sysprep.exe`. Click OK. When the System Preparation Tool 2.0 screen appears, click Reseal and then click OK.
29. After the virtual machine shuts down, switch to the OpsMgrBase configuration web page. Click the OpsMgrBase drop-down arrow in the status section and select Merge Undo Disks.
30. After the drive has been committed, click the OpsMgrBase drop-down arrow again and select Remove; then click OK.

Creating the Domain Controller:

1. On the Virtual Server Administration Website, create a new virtual drive by clicking Virtual Disks > Create > Differencing Virtual Hard Disk.
2. Specify the location where you have the `OpsMgrBase.vhd` file and name the new virtual disk `ZygortAD`. Select the `OpsMgrBase.vhd` file in the Known Virtual Hard Disks drop-down list and click Create.
3. Create the new virtual machine by clicking Virtual Machines > Create.
4. In the Virtual Machine Name text box, type the full path to the location where you created the virtual hard disk and append the name **ZygortAD**.
5. In the Memory text box enter **750**, choose Use An Existing Virtual Hard Disk, and select the `ZygortAD.vhd` file.
6. Select External Network from the Virtual Network Adapter options. Leave all the other options at their defaults and click Create.
7. If you are not at the `ZygortAD` configuration screen, click Virtual Machines > Configure > `ZygortAD`.
8. Click the Hard Disks link in the `ZygortAD` Configuration section and make sure the Enable Undo Disks option is selected. Then click OK.
9. On the `ZygortAD` Status screen, select the `ZygortAD` drop-down arrow and click Turn On.
10. Switch to the remote control client and select the `ZygortAD` virtual machine.
11. When the setup wizard starts, accept the defaults, except:
 - ◆ Enter your name and company on the Personalize Your Software page.
 - ◆ Enter your product key on the Your Product Key page.
 - ◆ Select Per Device Or Per User on the Licensing Modes page.
 - ◆ Enter **ZygortAD** in the Computer Name field and enter a password for the administrator account.
 - ◆ Specify your Date and Time settings.
 - ◆ Choose Custom Settings and enter a static IP address, default gateway, and DNS server settings.
12. After the `ZygortAD` virtual machine reboots, log on as administrator.
13. Go to Start > Run, type **dcpromo**, and click OK.

14. In the Active Directory installation wizard, use the defaults on all screens except:
 - ◆ Specify **Zygart.1c1** for the Full DNS Name For New Domain.
 - ◆ Specify a password for the administrator account on the Directory Services Restore Mode Administrator Password screen.
15. Reboot the virtual machine after the Active Directory installation completes.
16. Log on and open Active Directory Users And Computers.
17. Right-click **zygart.1c1** and select New ➤ Organizational Unit. Enter **OpsMgr Accounts** and click OK.
18. Right-click **OpsMgr Accounts** and select New ➤ User.
19. In User Logon Name, type **OpsMgrAction**. In Full Name, enter **Operations Manager Action Account**. Click Next.
20. Create a password for the account, deselect User Must Change Password At Next Logon, and select Password Never Expires. Then click Next and click Finish.
21. Perform steps 18–20 for the following accounts: Operations Manager SDK Account, OpsMgrSDK, SQL Server Service Account, and SQLSA.
22. Create a security group with the name **OpsMgrAdmins** and then shut down ZygartAD.
23. After the virtual machine shuts down, switch to the ZygartAD configuration web page and select the ZygartAD drop-down arrow in the status section. Select Merge Undo Disks.
24. After the drive has been committed, click the Memory link and enter **128** in the Virtual Machine Memory (In MB) text box and click OK.
25. Select the ZygartAD drop-down arrow again and select Turn On.

Creating the SQL Server:

1. On the Virtual Server Administration Website, create a new virtual drive by clicking Virtual Disks ➤ Create ➤ Differencing Virtual Hard Disk.
2. Specify the location where you have the **OpsMgrBase.vhd** file and name the new virtual disk **ZygartSQL**. Select the **OpsMgrBase.vhd** file in the Known Virtual Hard Disks drop-down list and click Create.
3. On the Virtual Server Administration Website, create a new virtual drive by clicking Virtual Disks ➤ Create ➤ Dynamically Expanding Virtual Hard Disk.
4. Choose the location where you save your virtual disks and name the new drive **ZygartSQLDatabase**. Specify the size as 127 GB and click Create.
5. On the Virtual Server Administration Website, create a new virtual drive by clicking Virtual Disks ➤ Create ➤ Dynamically Expanding Virtual Hard Disk.
6. Choose the location where you save your virtual disks and name the new drive **ZygartSQLLogs**. Specify the size as 30 GB and click Create.
7. Create the new virtual machine by clicking Virtual Machines Create.

8. In the Virtual Machine Name text box, type the full path to the location where you created the virtual hard disk and append the name **ZygortSQL**.
9. In the Memory text box, enter **750**; choose Use An Existing Virtual Hard Disk; and select the ZygortSQL.vhd file.
10. Select External Network from the Virtual Network Adapter options. Leave all of the other options at their defaults and click Create.
11. If you are not at the ZygortSQL configuration screen, click Virtual Machines Configure > ZygortSQL.
12. Click the Hard Disks link in the ZygortSQL Configuration section and make sure the Enable Undo Disks option is selected.
13. Click Add Disk and select the Primary Channel (1) Attachment option and the ZygortSQLDatabase.vhd file for Known Virtual Hard Disks.
14. Click Add Disk and select the Secondary Channel (1) Attachment option and the ZygortSQLLogs.vhd file for Known Virtual Hard Disks.
15. On the ZygortSQL Status screen, select the ZygortSQL drop-down arrow and click Turn On.
16. Switch to the remote control client and select the ZygortSQL virtual machine.
17. When the setup wizard starts, accept the defaults except:
 - ◆ Enter your name and company on the Personalize Your Software page.
 - ◆ Enter your product key on the Your Product Key page.
 - ◆ Select Per Device Or Per User on the Licensing Modes page.
 - ◆ Enter **ZygortSQL** in the Computer Name field and enter a password for the administrator account.
 - ◆ Specify your Date and Time settings.
 - ◆ Choose Custom Settings and enter a static IP address, default gateway, and the IP address of the domain controller for the DNS server settings.
 - ◆ Choose the Domain option and add the server to the Zygort domain.
18. After the ZygortSQL virtual machine reboots, log on to the zygort domain as administrator.
19. Open Add/Remove Programs and select Add/Remove Windows Components.
20. Select Application Server and click Details. Select Internet Information Services (IIS) and click Details.
21. Click the check box for World Wide Web Service and click OK twice, then click Next.
22. Shut down ZygortSQL.
23. After the virtual machine shuts down, switch to the ZygortSQL configuration web page and click the ZygortSQL drop-down arrow in the status section. Select Merge Undo Disks.

Creating the Operations Manager System:

1. On the Virtual Server Administration Website, create a new virtual drive by clicking Virtual Disks ➤ Create ➤ Differencing Virtual Hard Disk.
2. Specify the location where you have the `OpsMgrBase.vhd` file and name the new virtual disk **OpsMgr1**. Select the `OpsMgrBase.vhd` file in the Known Virtual Hard Disks drop-down list and click Create.
3. Create the new virtual machine by clicking Virtual Machines Create.
4. In the Virtual Machine Name text box, type the full path to the location where you created the virtual hard disk and append the name **OpsMgr1**.
5. In the Memory text box, enter **750**; choose Use An Existing Virtual Hard Disk; and select the `OpsMgr1.vhd` file.
6. Select External Network from the Virtual Network Adapter options. Leave all of the other options at their defaults and click Create.
7. If you are not at the `OpsMgr1` configuration screen, click Virtual Machines Configure ➤ `OpsMgr1`.
8. Click the Hard Disks link in the `OpsMgr1` Configuration section and make sure the Enable Undo Disks option is selected; then click OK.
9. On the `OpsMgr1` Status screen, select the `OpsMgr1` drop-down arrow and click Turn On.
10. Switch to the remote control client and select the `OpsMgr1` virtual machine.
11. When the setup wizard starts, accept the defaults except:
 - ◆ Enter your name and company on the Personalize Your Software page.
 - ◆ Enter your product key on the Your Product Key page.
 - ◆ Select Per Device Or Per User on the Licensing Modes page.
 - ◆ Enter **OpsMgr1** in the Computer Name field and enter a password for the administrator account.
 - ◆ Specify your Date and Time settings.
 - ◆ Choose Custom Settings and enter a static IP address and default gateway and the IP address of the domain controller for the DNS server settings.
12. After the `OpsMgr1` virtual machine reboots, log on as administrator.
13. Open Add/Remove Programs and select Add/Remove Windows Components.
14. Select Application Server and click Details. Select Internet Information Services (IIS) and click Details.
15. Click the check box for World Wide Web Service and click OK twice; then click Next.
16. Shut down `OpsMgr1`.
17. After the virtual machine shuts down, switch to the `OpsMgr1` configuration web page and click the `OpsMgr1` drop-down arrow in the status section. Select Merge Undo Disks.

Installing SQL Server 2005:

1. Turn on ZygortSQL and log on as administrator.
2. On the ZygortSQL configuration page, make sure the CD/DVD setting is configured for the physical CD drive and put the SQL Server 2005 CD in the drive.
3. Select the Server Components, Tools, Books Online And Samples link beneath the Install option.
4. Accept the EULA and install the prerequisites.
5. During the System Configuration Check, you will receive a warning due to the amount of memory that is allocated to the system. Make sure there are no errors and click Next.
6. Enter your name and company and the product key for SQL Server 2005; then click Next.
7. Select the SQL Server Database Services and Workstation Components, Books Online And Development Tools check box, and then click Advanced.
8. Expand Database Services and change the Replication option to Entire Feature Will Be Unavailable. Do the same for the Software Development Kit under Client Components. Click Next.
9. Click Next on Default Instance.
10. For the Service Account, enter **SQLSA** in the Username field, type **Zygort** in the Domain field, and enter the password you used when you created the account in the Password field. Select the SQL Server Agent and SQL Browser check boxes, and then click Next.
11. Click Next on the Authentication Mode, Collation Settings and Error And Usage Report Settings screens, and then click Install.
12. Click Next and Finish to complete the installation.

Installing the Operations Manager Database:

1. Make sure that you have installed all the prerequisites. If you are unsure, the prerequisite check will tell you what you need to install. Make sure you download and install all the prerequisite software.
2. Click the Install Operations Manager 2007 link after you insert the CD into the physical CD drive.
3. Click Next, accept the license agreement, and click Next again.
4. Enter your name, organization, and product key information and click Next.
5. Choose This Component Will Not Be Available for all of the selections except MOM Database and click Next.
6. Enter **OpsMgrGroup1** in the Management Group Name text box. Click the Browse button and select OpsMgrAdmins as the administrators; then click Next.
7. Make sure ZygortSQL is selected for the SQL Server Database Instance and click Next.
8. Click Advanced on the Database And Log File Options page. Change the database location to the E: drive and the logs to the F: drive. If the drives do not appear, go to Disk Management and create and format the partitions. Click OK and then click Next.

9. Click Next on the SC Error Reporting page, and then click Install.

Installing Operations Manager 2007:

1. Make sure that you have installed all the prerequisites. If you are unsure, the prerequisite check will tell you what you need to install. Make sure you download and install all of the prerequisite software.
2. Click the Install Operations Manager 2007 link after you insert the CD into the physical CD drive.
3. Click Next, accept the license agreement, and click Next again.
4. Enter your name, organization, and product key information and click Next.
5. Choose This Component Will Not Be Available for the MOM Database and click Next.
6. Enter **ZygortSQL** and click Next.
7. Enter **OpsMgrAction** in the User Account field and the password you used in the Password field; then click Next.
8. Enter **OpsMgrSDK** in the User Account field and the password you used in the Password field; then click Next.
9. Click Install to finish the installation.
10. Once the installation has completed, open a command prompt.
11. At the command prompt, enter **MOMADAdmin OpsMgrGroup1 OpsMgrAdmins OpsMgr1 zygort**.
12. After you receive confirmation that the MOMADAdmin utility ran successfully, close the command prompt.
13. Close each of your virtual machines and save the changes.

Chapter 3

Management Group Settings

Now that we have installed our first management group and have the management server in place, let's concentrate on configuring the management group so that it can perform its duties. We can start working with three setting types: the management group, the management server, and the agent. When you configure the management group, you are setting the default options for all of the management servers within that group. Then, if the need arises, you can override the global settings at the individual management server or agent to suit your needs.

In this chapter you will learn how to:

- ◆ Configure the global settings that apply to all management servers
- ◆ Configure the global agent properties
- ◆ Override the global settings on individual management servers
- ◆ Configure Operations Manager settings within Active Directory

Configuring Management Group Global Settings

When you set up a management group, the management servers in that group will more than likely all need to have the same settings applied to them. Each of the servers will be performing basically the same function in the management group. We can apply several settings to the management group so that the servers will take on the same functions. Known as the *global settings*, all of the management servers will abide by these settings unless overridden on a per-server basis.

Clicking on the Settings node brings up the screen shown in Figure 3.1. In the details pane you will notice the three types of settings: Agent, General, and Server. The settings in the General list control the management group and all of the systems within it. The settings that you can configure in this list are:

Alerts To control the behavior of the alerts, you can configure settings that specify the alert resolution state. These property pages allow you to create custom field properties for the alert criteria that are not covered by the default option. You can also configure auto-resolution parameters here.

Database Grooming To maintain an efficient database, you need to clean out the old data. The settings for database grooming allow you to control how often the data is removed from the database.

Notification Operations Manager supports four notification channels (otherwise known as methods of transmitting notifications to operators): E-mail, Instant Messaging, Short Message Service (SMS), and Command. In this area you can configure the notification channel that suits

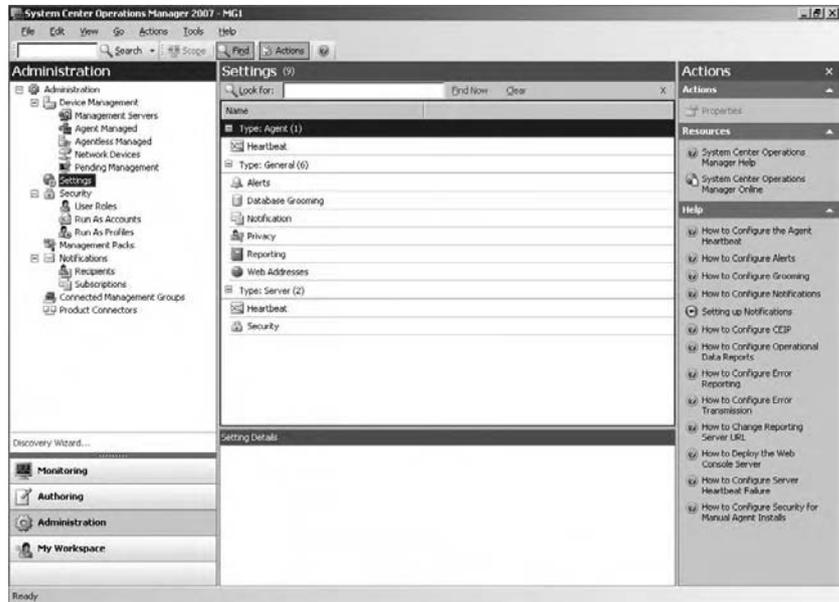
your needs, and then you can configure the recipients who will receive the notifications in the Notifications node.

Privacy As exceptions occur within the management group, Operations Manager will capture the exceptions. Based on the setting configured in this section, you can control whether the exceptions are sent to Microsoft.

Reporting To direct the management servers to the correct database server and database used for the reporting server, the Reporting settings are used. These settings include the Data Warehouse Server Name, Data Warehouse Database Name, and Reporting Server URL.

Web Addresses Two web addresses are used within the management group: one for the Web Console and the other for your own company's Online Product Knowledge.

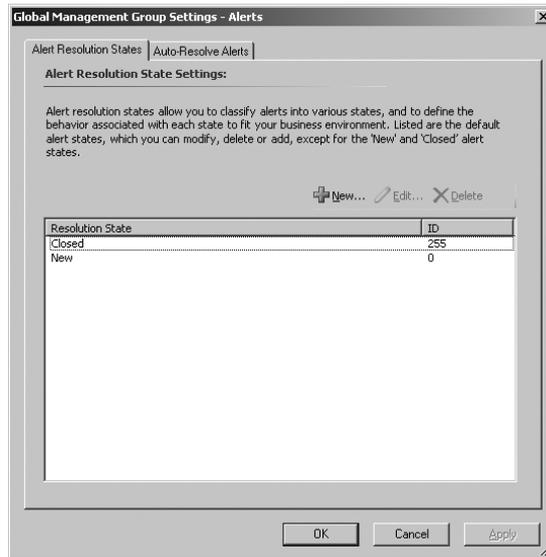
FIGURE 3.1
Settings options



Alerts

In the Alerts settings you will find three property pages: Alert Resolution States, Custom Alert Fields, and Auto-Resolve Alerts. Each of these settings helps you control an aspect of the alerts that arise within your organization. The first property page, Alert Resolution States, specifies the resolution levels that you use in your organization. You can assign each of the resolution states a unique identifier that you can use when searching the database for alert resolution status. As you can see in Figure 3.2, two states exist when you install: New and Closed. New is granted the ID of 0, and Closed is set to 255. All other values are available for you to use.

FIGURE 3.2
Alert resolution states



When configuring the resolution states, you should refer to the design documentation. During the design phase, the design team should have identified the alert resolution states that will be used by the operations monitoring staff. From the list they have provided, you can create the necessary resolution states by clicking the New button and entering the name for the resolution state and its ID. Figure 3.3 shows the alert resolution state of First Level Support being added with an ID of 10. After you use an ID, it will no longer be available from the list, so you don't have to worry about accidentally entering the same value for two IDs.

FIGURE 3.3
Adding a new alert resolution state



The third Alerts property page is Auto-Resolve Alerts. Here you can specify the amount of time that is allowed to pass on an open alert before it is automatically set to the resolved state. Two settings, shown in Figure 3.4), can be applied here. The first, Resolve All Active Alerts In The New Resolution State After, is set to 30 days by default. If an alert has remained in the New state for that long, it is assumed that it has been ignored because it is unimportant. The second option is Resolve All Active Alerts When The Alert Source Is Healthy After. This option “auto-resolves” all alerts that have not been confirmed as resolved after the source of the alert has been healthy for the amount of time specified. The default setting is one week. You can alter these settings as you deem necessary.

FIGURE 3.4
Auto-Resolve Alert
settings



Database Grooming

The primary purpose of auto-resolving alerts is to maintain the size of the database. After alerts have been set to the resolved state, you can remove them from the database in an action known as *grooming*. Database grooming is nothing new; it has been available for quite some time with most databases. The Database Grooming settings will help you maintain the size of your database.

As you can see in Figure 3.5, only one property page is available for database grooming. All of the grooming options that apply to the Operations Manager database can be found here. To change a grooming setting, select it and click the Edit button. The grooming options are:

Resolved Alerts These are alerts that have had their alert resolution status set to Resolved. The default setting is 7 days.

Event Data Events that are collected by the management servers populate the database so that you can make sure that specific events do occur, or that several events of the same type are not occurring on multiple systems at the same time (which could indicate a distributed attack). The default setting is 7 days.

Performance Data Performance data can be collected so that you can monitor how your managed systems are functioning, and possibly raise alerts if the performance data falls outside constraints you have set. The default setting is 7 days.

Task History From the console, you can run tasks against the managed systems. These tasks could include anything from running a quick diagnostic command to performing actions on the managed server. The details from the tasks are retained in the database until groomed. The default is 7 days.

Monitoring Job Data Job data on the managed systems can be collected and monitored. The default setting is 7 days.

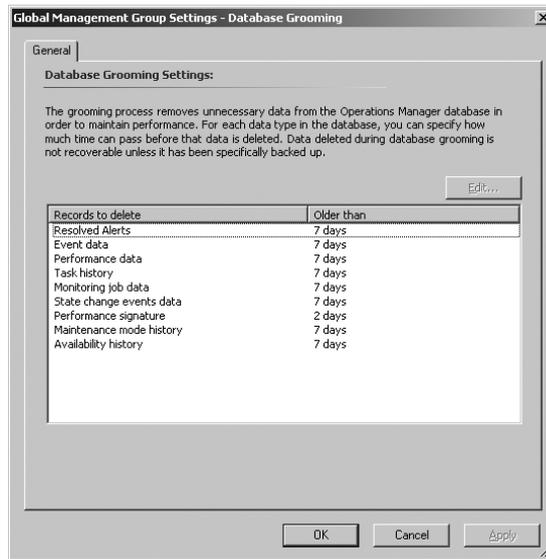
State Change Events Data The database keeps track of the state of alerts: when they were changed and who made the change. This data is groomed by default after 7 days.

Performance Signature The default setting is 2 days.

Maintenance Mode History Managed systems can be put in maintenance mode whenever updates, repairs, or changes are made to the systems configuration. Once in maintenance mode, the agent will not report any “out of spec” events or alerts. This history information is groomed after 7 days by default.

Availability History Most organizations are concerned about Service-Level Agreements (SLAs). System availability is a key ingredient to their SLAs. Availability information is collected from each of the managed systems and populated in the database. The default grooming setting is 7 days.

FIGURE 3.5
Database Grooming
Settings



Notification Channels

Of all the settings that you can modify for the management group, the most important are the Notification parameters. Without a notification channel defined, you will not be able to send any notification messages to your personnel, and you will have to rely on someone to discover the alert in the console.

E-MAIL CHANNEL

Looking at the options that appear in Figure 3.6, note that you have several settings that can be configured for the email notification channel. The E-mail settings are probably some of the most important. However, if the email server that you are relying on fails, the alert may never reach the intended recipient. To combat a failed email server scenario, Operations Manager allows you to configure multiple email servers and specify the order in which they will be used. If you click the Add button next to SMTP Servers, you will be presented with a dialog box like the one shown in Figure 3.7. Here you can enter the name of the SMTP server, the port that you will use when you connect to the SMTP server, and the authentication method that the SMTP server requires, either anonymous or Windows-integrated.

FIGURE 3.6
E-mail settings

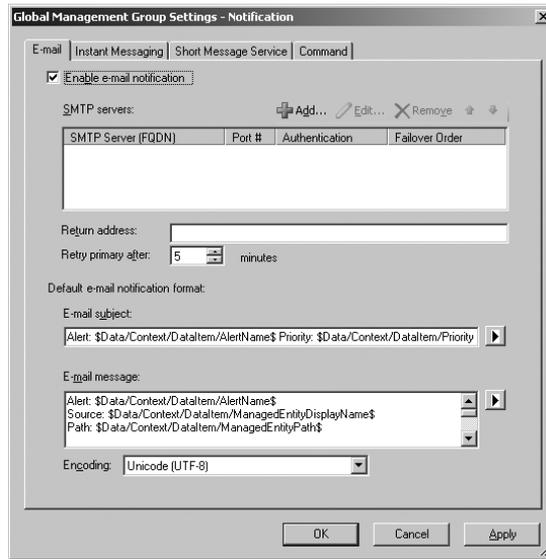


FIGURE 3.7
The Add SMTP Server
dialog box



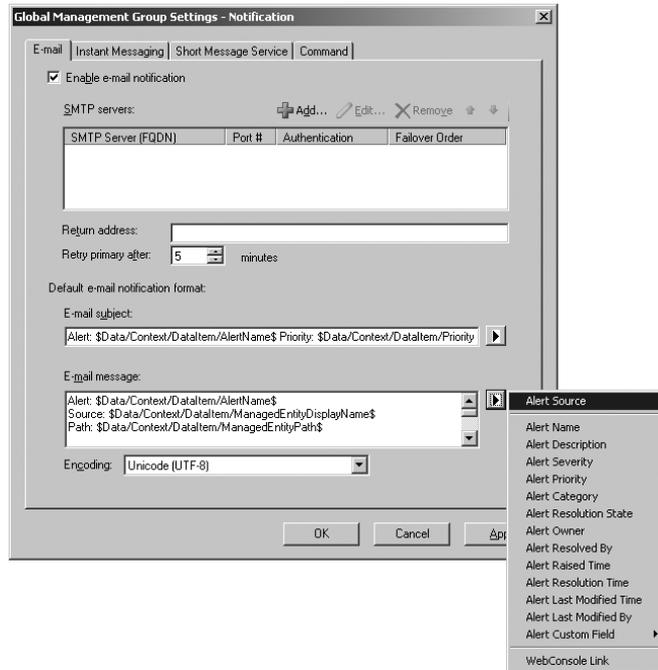
If you are not planning on using anonymous authentication for your SMTP server, you will need to make sure that the account used for the Agent Action account is granted submit permissions to the SMTP server. Otherwise, the management server will not be able to send the notification e-mail to the appropriate recipient.

When you have entered the SMTP servers that you will use for e-mail, you can select one and change the failover order by clicking the up and down arrows above the SMTP server list. You can also select a server and edit the properties by clicking the Edit button, or you can completely remove the SMTP server from use by the management servers by clicking the Remove button.

Directly beneath the SMTP Servers list you will find the Return Address and Retry Primary After settings. The return address should be the address of an administrator who will receive any email reply. Some spam filters will mark an email as unsolicited commercial email if a sender address is not configured. For the second option, if the SMTP server that is configured as the primary server fails, the time frame that you enter in this box will control how long the management server will wait until the server is retried.

At the bottom of the E-mail property page you will find the format of the email that is to be sent to the recipients. Here you can set the subject, message, and encoding used within the email. Messages can include any text that you want to enter as well as variables that will be used to notify the recipient of the system status. As you can see in Figure 3.8, when you click the Placeholder button, you can select the variables that you want to include in the subject or message body.

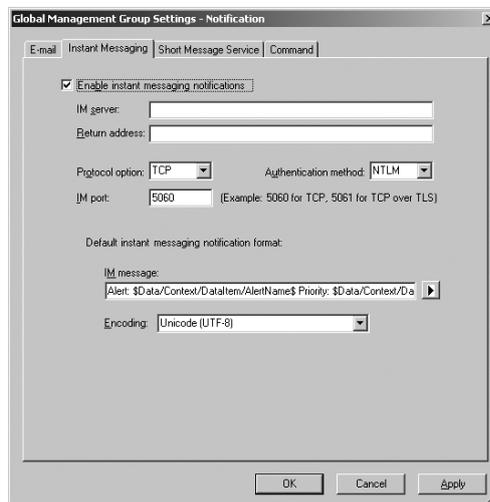
FIGURE 3.8
Email variables



INSTANT MESSAGING CHANNEL

In Figure 3.9, the Instant Messaging options are shown. If you are using an instant messaging service such as Live Communications Server, you can enter the name of the IM server as well as the protocol, port number, and authentication method to use (NTLM or Kerberos). The Return Address entry allows you to configure the address that will be seen by the recipient when the instant message appears within their messaging window.

FIGURE 3.9
Instant Messaging settings



As with the E-mail channel, you can configure the message that is sent to the recipient. Here you have only the option to set the message itself since there are no subject-line options for instant messaging. The same placeholders are available as those available for email. Note that most instant messaging services have restrictions on the total number of characters that you can send in a single message. Test to make sure that your message will not be truncated and that the relevant data is at the beginning of the message.

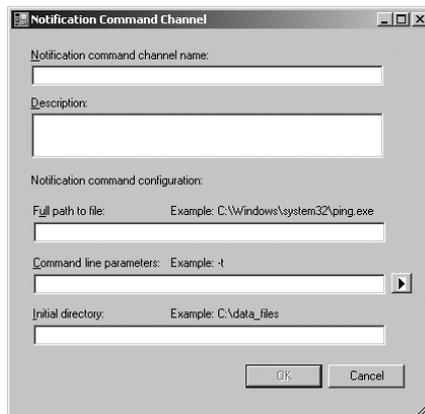
SHORT MESSAGE SERVICE CHANNEL

The third channel uses Short Message Service (SMS) messaging to send the notification. SMS is used by several cell phone providers to send messages from phone to phone. On this property page, you have the option of configuring the SMS message that will be sent to the recipient when an alert condition arises. As with the Instant Message channel, you may be restricted to the maximum number of characters that you can send in one message, so test to make sure that the message placeholder you are using is configured to not pass too much data.

COMMAND CHANNEL

When you first select the Command channel property page, there doesn't seem to be much to work with. The commands that you can use for notification purposes are hidden behind the scenes. The label that appears on this page is simply the name of the channel that you are configuring. When you click the Add button, or select a command channel name and click the Edit button, you are presented with a window that looks like the one shown in Figure 3.10. In the Notification Command Channel Name field, enter a descriptive name that will be meaningful for all the Operations Manager administrators. You can optionally use the Description field to add a meaningful explanation of what the channel provides.

FIGURE 3.10
Notification
Command Channel
settings



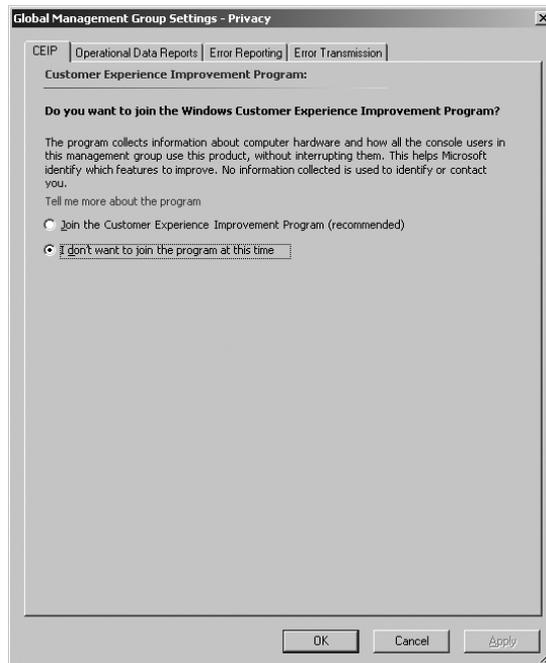
At the bottom of the Notification Command Channel page is the heart of the channel. You can define the command that will be used as well, as the parameters for the command. In the Full Path To File field, enter the command line. Then in the Command Line Parameters field, enter the parameters that will be used with the command.

Privacy Settings

Two camps of thought exists when it comes to the optional reporting feature that is becoming part of all Microsoft products. Some folks do not trust Microsoft, and believe that they are probably collecting more information than they are willing to admit. On the other side, you have folks who want to assist Microsoft in error collection so that hot fixes, service packs, and future products will correct problems that they encounter.

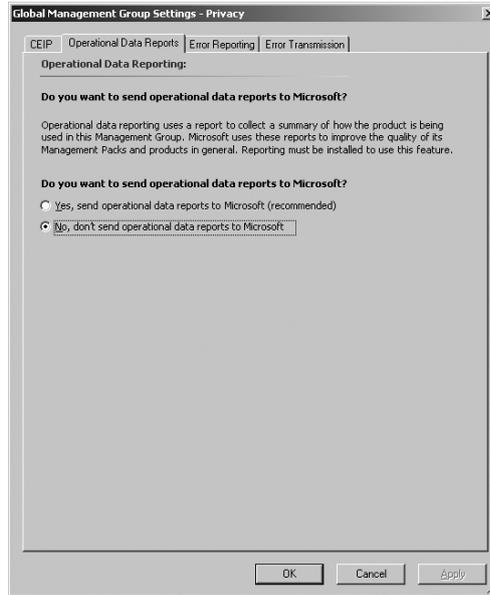
In the Privacy settings dialog box, you will find the CEIP tab, shown in Figure 3.11. Notice that you have the power to control whether data is sent to Microsoft. The first of the four tabs allows you to opt into the Customer Experience Improvement Program. If you choose the option to join, data will be collected on the systems in your management group. This information will include the hardware and software that you have installed on the management servers and the systems that the operators use to monitor. As you can see from the description of the program, Microsoft wants to make sure that they are building their products to be the most efficient, so collecting this information helps them target the types of systems and the monitoring methods that are used by operators.

FIGURE 3.11
Privacy settings
dialog box



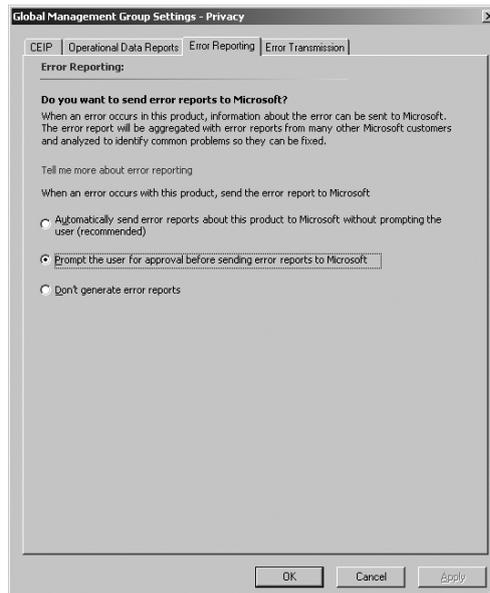
The second property page, shown in Figure 3.12, is named Operational Data Reports. The data reporting feature takes advantage of the reporting functionality of Operations Manager. Therefore, if you do not have reporting enabled, you will not be able to take part in this program. Once you have enabled reporting, one of the report types, Operational Data Reports, details how the management group and all of the managed systems are functioning. This data can be used to see whether managed systems are falling out of the specifications you have set, and how the management group is functioning as a whole. These reports can be automatically sent to Microsoft so that they can also have an idea of how systems are functioning and how the management packs are working for those products.

FIGURE 3.12
Operational Data
Reports settings



The third property page is titled Error Reporting. Whenever an error condition arises that causes Operations Manager to fail, the operating system captures the memory information that pertains to the management server. Three settings exist for this feature, two of which will collect error data and prepare it to be sent to Microsoft. The first option shown in Figure 3.13, Automatically Send Error Reports About This Product To Microsoft Without Prompting The User (Recommended), does exactly what it says: it collects the data and immediately sends it to Microsoft if a connection can be made.

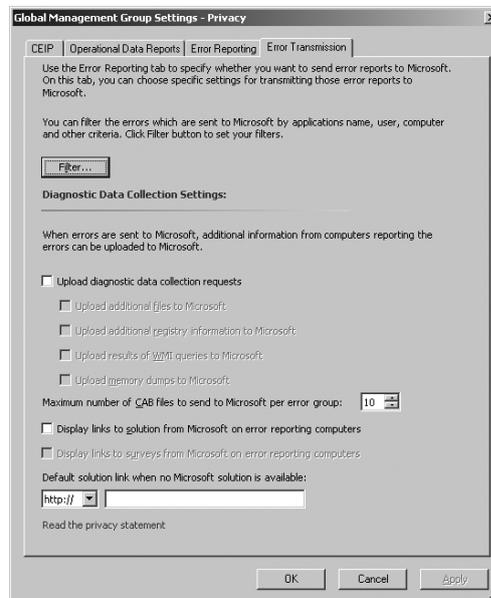
FIGURE 3.13
Error Reporting
settings



The second option, Prompt the User For Approval Before Sending Error Reports To Microsoft, gives you a chance to approve or deny which error reports will be sent to Microsoft. Once you select this option, when you log on you will be presented with a prompt that will detail the error reports that have been queued. You can choose which will be sent and which will not.

The final property page in the Privacy settings dialog box is Error Transmission. As shown in Figure 3.14, you can configure what data is sent to Microsoft. When you click the Filter button on this page, you are presented with a new page that is reminiscent of the Rules wizard from Outlook. You can select the filter type you would like to use and then enter the criteria that you would like to filter on. For example, if you would like to make sure that errors from specific users are not delivered to Microsoft, select the That Come From Specific Users check box and enter the account names of those users, as shown in Figure 3.15.

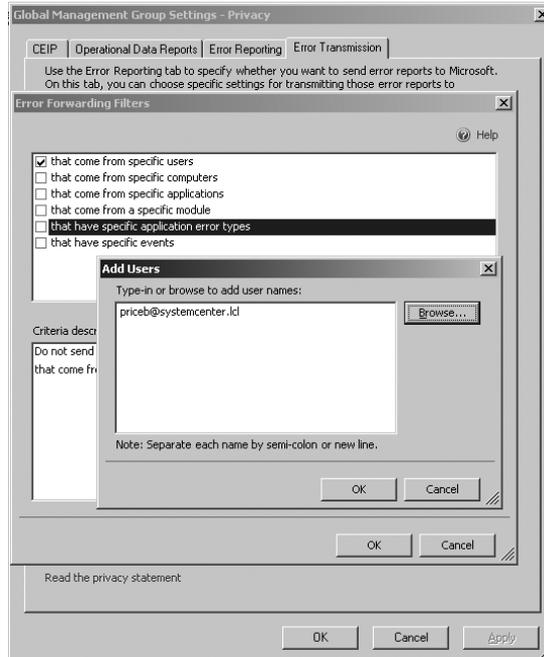
FIGURE 3.14
Error Transmission
settings



The lower portion of the Error Transmission page allows you to control the amount of information that is sent. When you select the Upload Diagnostic Data Collection Requests check box, four other selections become available: Upload Additional Files To Microsoft, Upload Additional Registry Information to Microsoft, Upload Results Of WMI Queries To Microsoft, and Upload Memory Dumps To Microsoft. When you provide more data to Microsoft, you make it easier for them to troubleshoot the problem and possibly provide resolutions for your staff and other organizations that are having the same problems. The data is then encapsulated in a cabinet file and delivered to Microsoft.

The last two check boxes control how data is displayed on the computers that are reporting the error conditions. If you select the first option, Display Links To Solutions From Microsoft On Error Reporting Computers, the web page that provides a resolution to the error can be provided to the computers. The second option, Display Links To Surveys From Microsoft On Error Reporting Computers, will present survey websites to the computers.

FIGURE 3.15
Controlling accounts
for error transmission



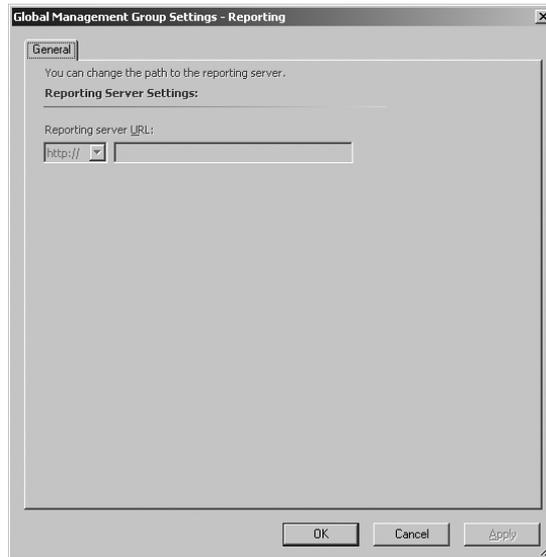
If you have a default support web page, you can present that to the system that generated the error by entering the URL of the support site in the Default Solution Link When No Microsoft Solution Is Available text box. With this URL, you could display some default troubleshooting steps that could be performed, or you could simply display details about the reporting steps that have just been taken. Some companies refer to this page as the “chill out” page because they are trying to get the users to calm down whenever an error occurs.

Reporting Settings

To take advantage of the reporting features of Operations Manager, you must make sure that the management servers can reach the database server that holds the data warehouse database. When you right-click on the Reporting feature and choose Properties, a property page titled General appears, as shown in Figure 3.16. Here you can specify the Reporting Server Settings.

In Chapter 9, “Reporting,” we are going to step through the configuration settings necessary to enable reporting services. A key ingredient to reporting is making sure you have an IIS web server available to generate the reports. Once everything has been configured for reporting, you can enter the URL in the Reporting Server URL text box.

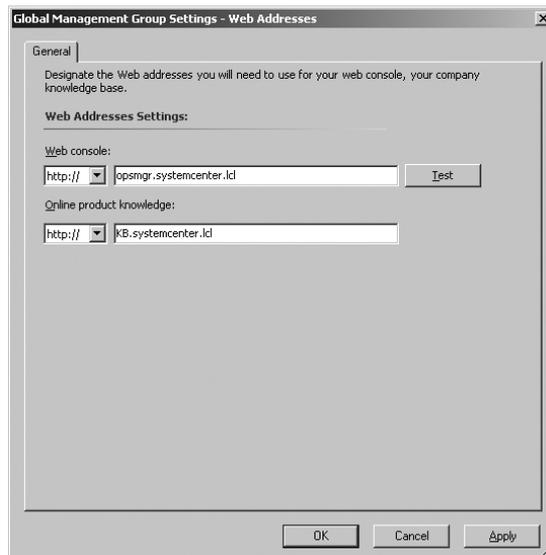
FIGURE 3.16
Reporting Server
Settings



Web Addresses

The final global setting option for the management group is the Web Addresses properties. Right-clicking this option and choosing Properties presents you with a page that looks like Figure 3.17. In the Web Console text box, you can enter the URL that defines the path to the web console on your IIS server supporting your management group. Notice that you have a Test button next to the text box that will allow you to make sure that the console is configured correctly.

FIGURE 3.17
Web Addresses
settings



If your organization has generated its own knowledgebase, whether it is for supporting the systems in your organization or for disseminating technical data on applications developed by your staff, you can specify a path to the intranet server that hosts the data.

Configuring Global Server Settings

If you are lucky, you will be able to set the global settings for your management servers and not have to worry about overriding those settings on an individual management server basis. By default, these settings affect all the management servers in the management group. The settings you find here can come in handy when you want to add a new management server to the group. It will fall under the global settings until you have to make an adjustment to how it functions. Two settings are applied for management servers: the agent heartbeat detection and the manual agent installation behavior.

Heartbeat

Each agent sends a small “heartbeat” message to its assigned management server on a periodic basis. To keep the management server from panicking if a single heartbeat fails, you can configure the number of missing heartbeats that will be allowed before the management server will attempt sending a ping request to the managed server. The default is to allow three heartbeats to go missing before the ping is sent. So if the agent heartbeat interval is set to 30 seconds and 90 seconds have passed since the last heartbeat was detected, the management server will send a ping request to the managed server. By sending a ping request, the management server can determine whether the managed system is still responding on the network. If the server replies to the ping request, the management server will assume there is a problem with the agent on the managed system. If the ping request fails, the management server can assume the managed system has failed.

Security

As you can see in Figure 3.18, the Security settings control how the management server handles agents that have been manually installed. An agent is assumed to be manually installed if it has been installed on the managed system by any means other than by having a management server push the client out during a discovery phase. So if you install the agent from the installation CD or use a software distribution method such as SMS 2003 or Configuration Manager 2007, the management server considers the agent manually installed.

The default setting is to reject any manually installed agents. This protects the management group from having a rogue agent installed on a system. Rogue agents could possibly inject inaccurate information into the database, causing incorrect information to be relayed through reports. A denial of service could occur if that agent fills the database.

If you do want to allow manually installed agents to work within your management group, select the option Review New Manual Agent Installations In Pending Management View. When this option is selected, manually installed agents appear in the Pending folder and you will have to right-click the entries and approve them. Once approved, they function as any other agent in the management group.

FIGURE 3.18
Manually installed
agent settings



Be careful with the check box that appears once you allow manually installed agents to function in the management group. By selecting the check box for Auto-Approve New Manually Installed Agents, you will be allowing all agents to start functioning within the management group without your express approval. While this setting may work well in a test environment where you want to bring agents online quickly, you should consider the ramifications in a production environment. Remember: data that is collected within the management group could affect how your SLAs are met, and if a rogue agent is allowed to function, you may find yourself trying to explain why your data is skewed.

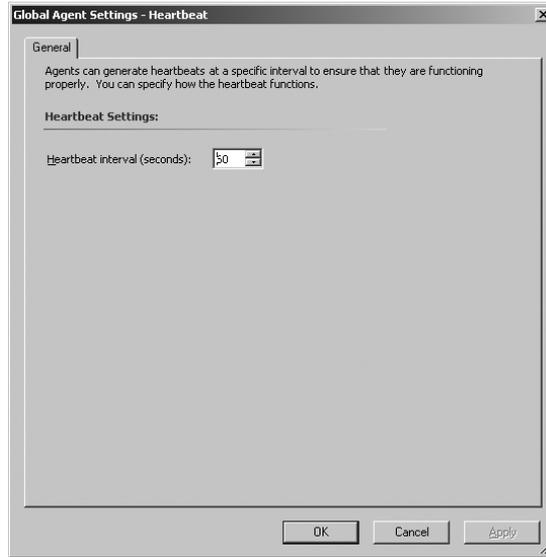
Exploring Global Agent Settings

You won't find very many global setting for the agents. As you will notice in Figure 3.19, the only option that is available under the Agent settings is Heartbeat. When you right-click Heartbeat and select Properties from the context menu, Enable Agent Heartbeat and Heartbeat Interval (Seconds) are the only two configurable options.

The default settings for these options are to have heartbeats enabled, transmitting every 60 seconds. For most organizations, these settings will suffice. You want to make sure that the agents are active and able to respond in case there is an alert or event that needs to be sent back to a management server. As long as the heartbeat is going out as planned, the management server knows that the agent is still performing its job.

If you want to extend the amount of time that passes before the heartbeat is sent, use the scroll arrows to increase the time. But be aware that if you do, the management server will not be notified as often that the agent is active and functional. This could extend the amount of time that it takes for a management server to recognize that a system is no longer available. On the flip side, reducing the time frame will increase the amount of network traffic that you will have. Make sure that you know what the ramifications are in your environment before you make changes to this setting.

FIGURE 3.19
Global Agent
Settings – Heartbeat



Understanding Individual Server Options

Whereas the global management group and global server settings apply to all the management servers in the management group, the individual server options define the overrides on a per-server basis. For instance, you may have managed servers in the management group that require you to manually install the agent. You could define one management server to allow manual agent installations, and then configure all those systems to that one management server. It will become responsible for monitoring those servers while the other management servers will not.

Agent Management

In Chapter 2, “Installing System Center Operations Manager 2007,” we discussed how to prepare Active Directory for use with Operations Manager. After setting up the required objects in Active Directory, agents can look to a domain controller to determine which management server they will report to. To populate the correct settings in Active Directory, you need to run the Agent Assignment And Failover Wizard. To do so, right-click on the management server and select Properties. As shown in Figure 3.20, the first property page is Auto Agent Assignment. When you click Add on this page, the wizard appears. The introduction page will tell you what the wizard is used for. Click the Do Not Show This Page Again check box to keep from having to view it again the next time you run the wizard.

After clicking the Next button, you are presented with the Domain page shown in Figure 3.21. Here you can specify whether the agents that will be managed by this server reside within a trusted domain or within a domain where a trust relationship has not been defined.

FIGURE 3.20
Auto Agent
Assignment
property page

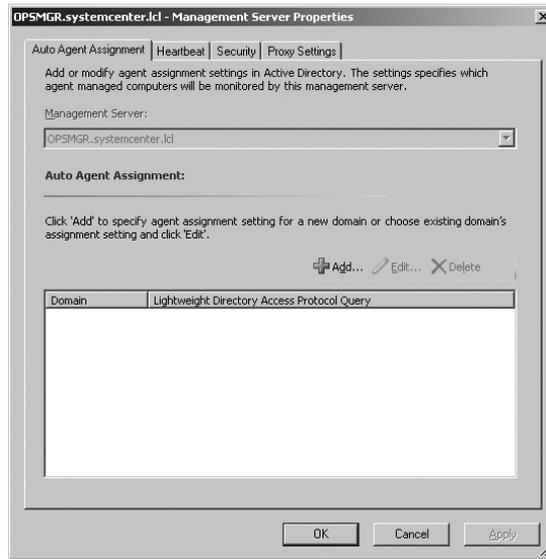
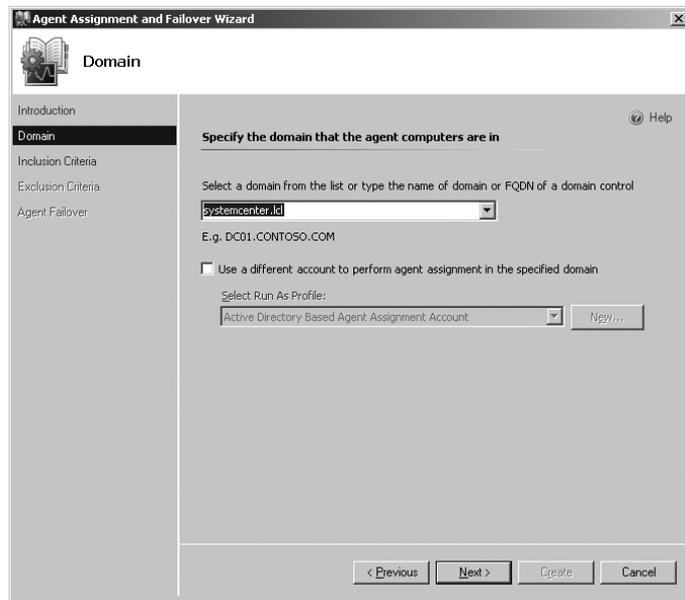


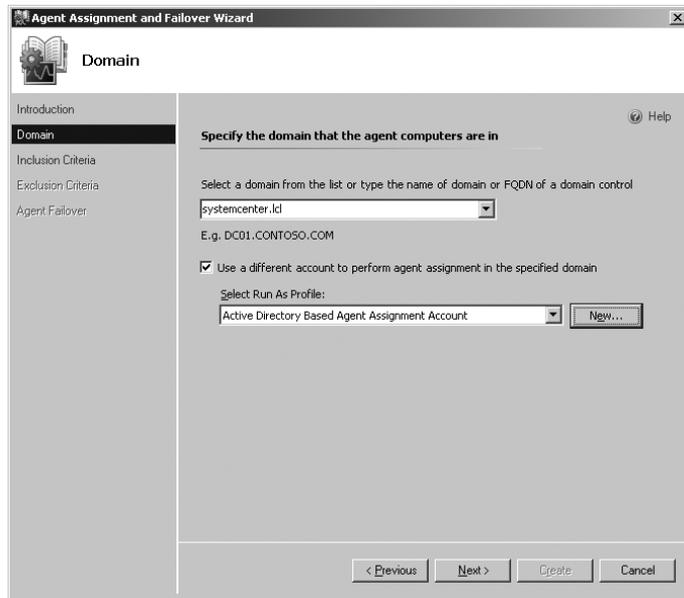
FIGURE 3.21
Domain page



When you click Next, you are presented with the Inclusion Criteria page, which allows you to define a Lightweight Directory Access Protocol (LDAP) query that will be used to locate the computers in the domain. At its most basic, the inclusion criteria can include a list of computer names.

You can also make the LDAP query as intricate as you need by searching for values in specific attribute fields. In Figure 3.22, notice that the option to use a different account to install the agent has been selected. By default, the Agent Action account is used to send the LDAP query to the domain controller. If you do not want to use that account, or you are using the Local System account of the computer for the Agent Action account and need to use an account that has permissions to read entries from Active Directory, you can select the appropriate Run As account from this list. The options available from the pull-down list include all of the Run As profiles that have been created. If you want to create your own, you can click the New button, which will take you to the Create Run As Profile Wizard. For more information on the Run As Profile Wizard and how to create new Run As accounts, see Chapter 7, “Monitoring with the Console.”

FIGURE 3.22
Domain page
properties



After clicking Next, you are presented with the second Inclusion Criteria page. Here is where you can create the LDAP query that is used to define which systems will be monitored by the management server. Clicking the Configure button brings a query builder into focus. The two tabs control which computers are included within the query. On the Computers tab, enter in the Computer Name field the name of the computer you wish to monitor. You can wildcard the name so that you are not restricted to setting up failover for one system. For example, you can enter `DC*` to select all the computers whose names start with `DC`.

On this tab you can also specify the owner of the computer. This field should be named Managed By because that is the Active Directory attribute that is used. If you have entered information in the Managed By property for each of the computer accounts in Active Directory, you can use this field to specify which computers should be monitored by the management server.

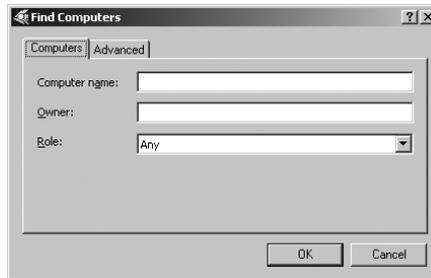
The final field on the Computers tab, shown in Figure 3.23, is the Role option. Here you can select which computer roles are assigned to the management server. The options found here are Any, Workstations And Servers, and Domain Controllers. Since you usually do not want to monitor all of the different roles on a management server, you can control whether domain controllers will be included in the monitoring, or specify that they are the only role type you will monitor.

Real World Scenario

USING THE MANAGED BY FIELD

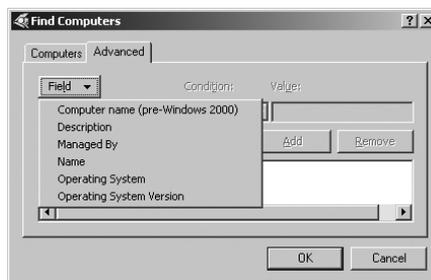
You cannot enter a computer account in the Managed By field of an Active Directory object, but you can specify user or group accounts. An organization wanting to streamline the assignment of monitored systems created different Operations Manager Admins groups for each of the management groups in its organization. They then assigned one administrator as the primary administrator for each management server. Once the primary administrator was identified, they created scripts that would populate the Managed By field for all their systems. The agent assignment could then be configured using the Owner option when they configured the inclusion rule for the management server.

FIGURE 3.23
Computers
tab options



If you select the Advanced tab, you will find that you have far more granular control over how you identify the systems that will be assigned to the management server. Clicking the Field button allows you to identify the attributes in Active Directory that will be used for the assignment. Notice that Computer Name and Managed By are both available as well, as shown in Figure 3.24. As a matter of fact, you can enter the same criteria on the Advanced tab as on the Computers tab and both entries will be populated.

FIGURE 3.24
Advanced
tab options



As you can see in Figure 3.25, there are six conditions that you can use to build the criteria:

Starts With This option will wildcard the text that you enter in the Value field so that anything you type will be appended with the * wildcard character when the LDAP query is built.

Ends With This option will wildcard the text that you enter in the Value field so that anything you type will be prefixed with the * wildcard character when the LDAP query is built.

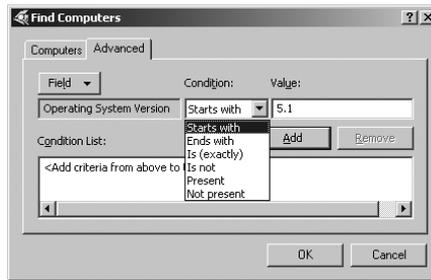
Is (Exactly) This option will use the literal value that you enter in the Value field.

Is Not This option will create a query option that will search for anything that does not include the literal value that you enter in the Value field.

Present This option will create a query option that will include any computer that has the chosen field populated with a value.

Not Present This option will create a query option that will include any computer that does not have a value entered in the chosen field.

FIGURE 3.25
Conditions you can use to build criteria



Sometimes we need a query that is a little more complex than what the query builder will generate for us. Using the query builder, we used two options: Computer Name (Pre-Windows 2000) with a value that starts with *wkstn* and Computer Name (Pre-Windows 2000) with a value that starts with *dc*. These two options created the query shown in Figure 3.26. The initial configuration would have created a query that looked for a computer with a name that started with *wkstn* and *dc*. Since that is a query that would not return any results, it was instead used as a basis for creating the custom query seen in Figure 3.27.

FIGURE 3.26
Query created by the query builder

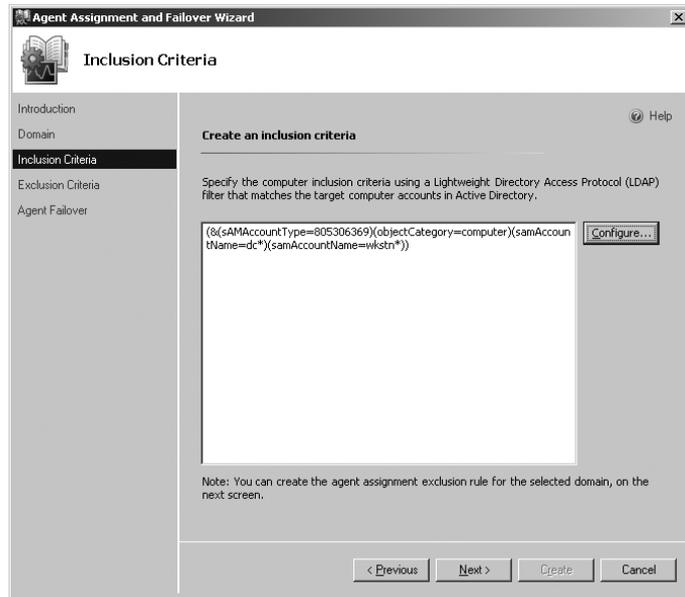
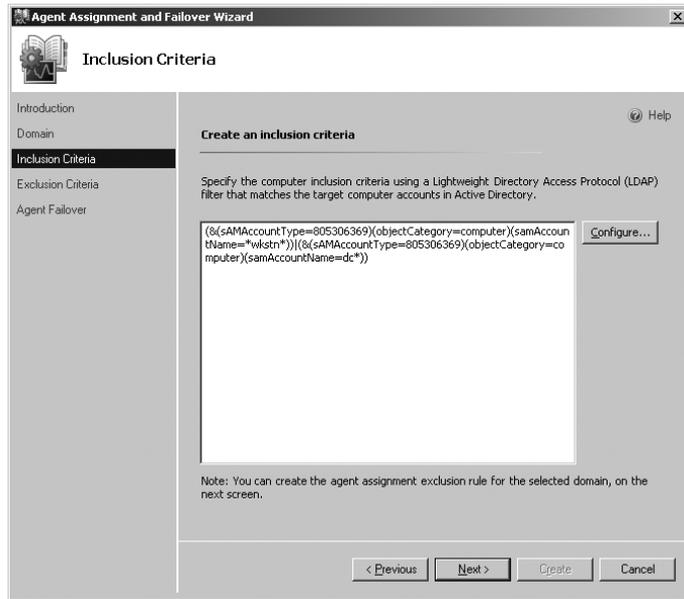
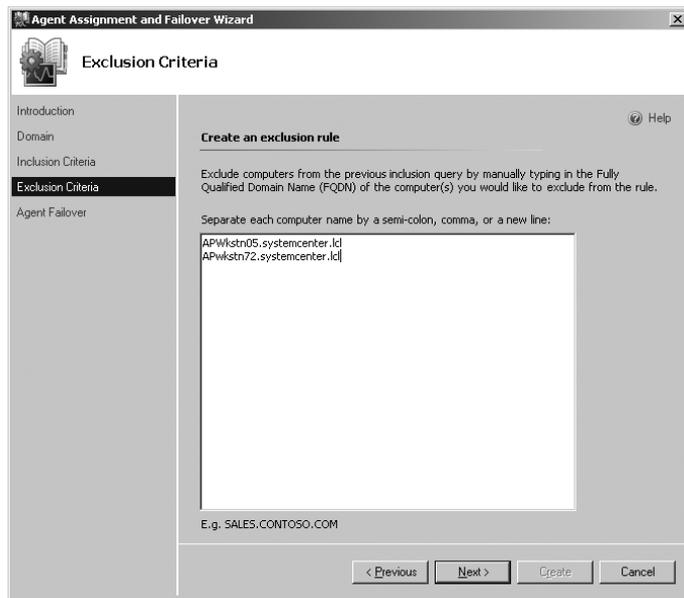


FIGURE 3.27
 Custom query to find computers whose names start with *DC* or contain *wkstn*



There may be times that you will create a query that includes all the computers you want to monitor, but also includes one or more that you do not wish the selected management server to monitor. You can include those exceptions on the Exclusion Criteria page. Unlike on the Inclusion Criteria page, you will have to enter the computers that should not be monitored one at a time using their fully qualified domain names, as shown in Figure 3.28.

FIGURE 3.28
 Computers placed on the Exclusion list



The final page of the wizard allows you to define how the failover will occur. If you select the first option, *Automatically Manage Failover*, the monitored systems' agents will automatically fail over to any other management server in the management group. This may not be the most optimal solution for you, however. You may have a management server that is already close to capacity, or you may have a dedicated management server that supports a specific set of computers. In either case, or if you have any reason to not allow agents to fail over to a specific system, select the *Manually Configure Failover* option and deselect any management server that should not be used in a failover scenario.

Once you click *Next*, the LDAP query will be saved. Then when agents are installed, the computer accounts for the managed systems will be added to the management computer's service connection point group in Active Directory.

Heartbeat

The global heartbeat setting has already been defined in the global management group properties. If you want to override the global setting on a per-server basis, select the check box *Override Global Server Settings* and then change the number of heartbeats that can be missed before the management server raises attempts to ping the agent managed system.

Typically, this is not a setting that is overridden, but you may have reason to. For instance, if you have a group of managed servers that reside in a subnet that has intermittent communications issues, you may want to extend the number of heartbeats that the management server will allow to be skipped. This will keep the management server from having to send out too many ping requests when you know that the heartbeats will not arrive from time to time.

Security

In the global management group settings, the Security settings allowed you to configure whether manually installed agents were permitted to function in the management group. The same settings are available on a per-server basis. If you would like to block manual installations on all your management servers in the management group except for one management server, you could set the global properties to reject the manual installations and then allow an override on one management server. Remember that systems residing behind firewalls that do not allow the proper ports for agent installation will have to utilize some form of manual installation. Since that is the case, you will need to allow at least one of your management servers to approve the manually installed agent. Failure to do so will leave some of your systems unable to be managed.

As you can see in Figure 3.29, there is another option on the Security property page that was not part of the global settings: *Allow This Server To Act As A Proxy And Discover Managed Objects On Other Computers*. Selecting this option will allow the management server to collect information from other managed server types. If you have SNMP in your environment, you could allow your management server to collect the SNMP traffic from managed devices and use the information in the Operations Manager database.

Proxy Settings

The settings found on the Proxy Settings page allow you to specify the proxy server that your management server will use when sending error reports to Microsoft. Often servers are not allowed to communicate across the Internet for security reasons. If you specify a proxy server to use, as shown in Figure 3.30, you can allow the Operations Manager server to pass data through a proxy server and not have to worry about trying to find another means of sending the data.

FIGURE 3.29
Proxy settings
on the Security
property page

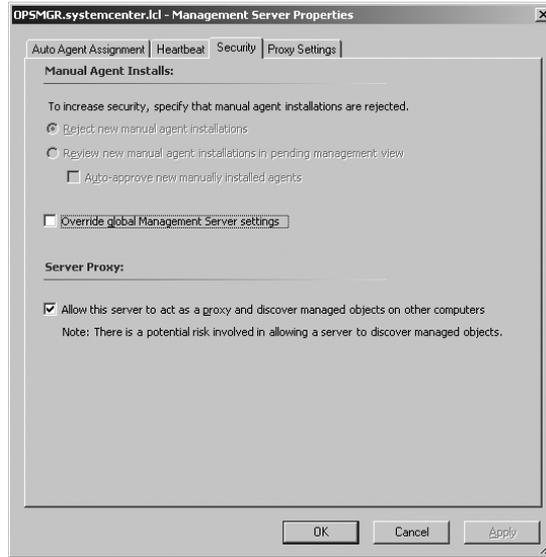
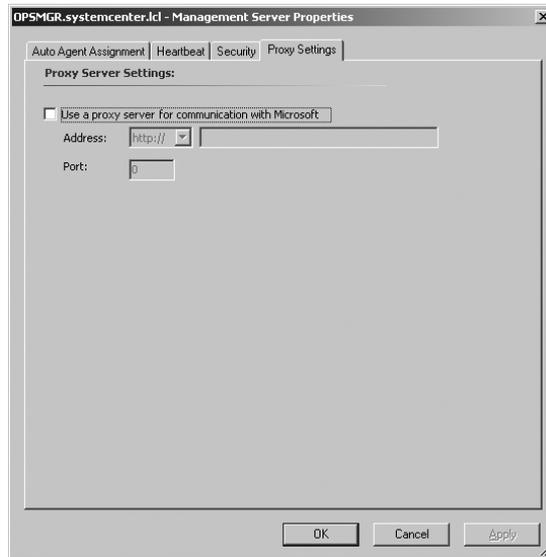


FIGURE 3.30
Proxy server settings



The Bottom Line

Configure the global settings that apply to all management servers. The management group settings can be configured so that all of the management servers behave the same way. These global settings are applied by each management server by default. The management group settings that you can configure globally are:

Alerts To control the behavior of the alerts, you can configure settings that specify the alert resolution state. These property pages allow you to create custom field properties for the alert criteria that are not covered by the default option. You can also configure auto-resolution parameters here.

Database Grooming To maintain an efficient database, you need to clean out the old data. The settings for database grooming allow you to control how often the data is removed from the database.

Notification Operations Manager supports four notification channels: E-mail, Instant Messaging, Short Message Service (SMS), and Command. In the Notification area, you can configure the notification channel that suits your needs, and then you can configure the recipients who will receive the notifications in the Notifications node.

Privacy As exceptions occur in the management group, Operations Manager will capture the exceptions. Based on the setting configured in this section, you can control whether the exceptions are sent to Microsoft.

Reporting To direct the management servers to the correct database server and database used for the reporting server, the Reporting settings are used. These settings include the Data Warehouse Server Name, the Data Warehouse Database Name, and the Reporting Server URL.

Web Addresses Two web addresses are used within the management group: one for the Web Console and the other for your company's Online Product Knowledge.

Configure the global agent properties. The agent heartbeat can be configured as a global property so that each agent in the management group sends its heartbeat at the same frequency. This allows the management server to know how often the agents should be reporting and allows you to control the network traffic for all the agents.

Override the global settings on individual management servers. Each management server will use the settings applied in the global properties unless you override the settings on a per-server basis. Once you override the settings, the management server has to be managed as a unique entity instead of using the global options.

Configure Operations Manager settings within Active Directory. In the management server properties, you can configure the managed systems that will be monitored by each management server in the management group.

Master It! At this point we are ready to configure the management group and management servers. We are going to start by configuring the global settings, and then we will override some of the settings on the management server. For this exercise, you will need to start all three of the virtual machines that were created in Chapter 2. Remember to start the domain controller first, followed by the SQL server, and then finally the management server.

Setting Global Properties and Overrides on Management Server:

1. Open the Operations Console by selecting Start > All Programs > System Center Operations Manager 2007 > Operations Console.
2. If not already selected, select the Administration view.
3. Click the Settings node.
4. Right-click the Alerts option under General in the Details pane and select Properties.
5. Click New on the Alert Resolution States page.
6. In the Alert Resolution State dialog box, enter **Support Level 1** in the Resolution State text box and set the Unique ID to 10. Click OK.
7. Perform steps 5 and 6 for the following:
 - Resolution State: Unique ID
 - Support Level 2: 20
 - Support Level 3: 30
 - Vendor Support: 35
 - Problem Management: 50
8. Right-click Security under Server and select Properties.
9. Review the settings and make sure Reject New Manual Agent Installations is selected.
10. Expand Device Management, select Management Servers, right-click your management server, and select Properties.
11. On the Security tab, select the Override Global Server Settings check box.
12. Under Manual Agent Installs, select Review New Manual Agent Installations In Pending Management View.
13. Make sure you leave the check box for Auto-Approve New Manually Installed Agents deselected and click OK.
14. Make sure you keep all undo disk changes if you shut down or save the virtual machines.

Chapter 4

Installing and Configuring Agents

As mentioned in previous chapters, you can set up as many servers as you think you may need, but unless you have the Operations Manager agent loaded on systems in your environment, you won't have an effective monitoring solution. The agent becomes the conduit that connects the management servers with the systems they monitor and maintain. In the simplest sense, you can think of the agent as a telephone and the caller as the monitored system, with the phone company acting as the management server.

This chapter concentrates on getting the agent onto the managed systems. Since there are different methods you can use to install the agent, you will need to decide which method works best for your environment. We'll look at the most popular methods and discuss the pros and cons of each one. As you review each installation method, envision your own environment and try to determine which methods would benefit you.

In this chapter you will learn how to:

- ◆ Automatically install the agent by using a discovery
- ◆ Perform a manual installation from CD and command line
- ◆ Automate the manual install of the agent by using command-line options
- ◆ Override the default management group settings on an individual client

Exploring Agent Installation Prerequisites

As with any piece of software, there are requirements that must be met in order for the software to run. Even though the Operations Manager client does not consume a lot of resources, you must make sure that a few things are in place before you try to run it. And depending on how you perform the installation of the agent, you should make sure that the management systems allow agent communication.

By default, the management systems are an elitist group. They won't talk to an agent that they did not install. If the agent has been installed in any fashion other than by the discovery / autoinstall method, the management servers will simply ignore the agent's request to communicate. Even though we don't usually care for this type of snobbish attitude in our lives, we commend the management servers for their actions.

In this day and age of security consciousness, we want to make sure that we allow communication to only those devices that are approved. If we were to allow our management servers to talk to any machine that had an agent installed, we could allow them to inject "bad" data into our management database. The data that is collected at the management server is used to ensure that our critical services are available to clients and that our network continues to provide mission-critical applications. If we were to permit unplanned data to enter the system, we could allow attacks such as denial of service and incorrect reporting.

In Chapter 3, we detailed the global settings for the management group as well as the individual server override options. One of these settings, Accept New Manual Agent Installations, allows you to manually install the agent on managed systems. Otherwise, the manually installed agents will be ignored. The other option you may want to consider is Auto-Approve New Manually Installed Agents. Be forewarned, however: any system that has an agent will be automatically approved and added to the Agents View in the console.

AUTO-APPROVAL

As mentioned in Chapter 3, automatically approving an agent-based system can introduce the risk of having an attacker attempt to inject false data into your environment. False data can cause problems—anything from simply generating false alerts to providing incorrect reporting to causing a denial of service of the management server. Before selecting the Auto-Approve New Manually Installed Agents setting, consider the ramifications of your action.

Installing by Using the Computer and Device Management Wizard

Up to this point, we have gone through the options that we need to consider when configuring the management servers and the management groups. We are now ready to take on the client deployment. The two methods, as with many installations, are the unattended installation and the manual installation. With Operations Manager 2007, there is a mechanism built into the management group that will allow you to choose the systems where the agent will install, as well as allowing that installation to proceed without any user intervention. Of course, you could always walk around to each system and install the agent yourself, but what fun is that? We like to have control over the systems that we work on, and like the mad scientist, we don't want to have to interact with our subjects—we just want to control them.

In Chapter 3 we configured most of the settings that were required for defining our management group and the agents that will function within the management group. Now we need to find the systems that we are going to manage. To do so, we will perform what is termed a *discovery*. Based on the rules that we apply, systems will be found that meet the rule criteria. So just what criteria does the rule use, you ask? Actually, there are very few options. You have the option to discover server-based systems as well as supported client operating systems.

This is where you need to go back and take a look at the design options you chose. If you are interested in monitoring your client workstations, then select that option. And, as common sense will tell you, select the option to monitor servers. Select them both if you want a full monitoring solution. There are trade-offs, however. The more you monitor, the more data you will have to retain in your database. If you collect everything, the database will grow extremely quickly and possibly out of control. Remember, step back and decide what is most important to you, and then expand upon it as you get a feel for what you are collecting and working with.

Bear with us for a few pages as we take a look at the agent deployment. Some of what we will go through is the step-by-step tasks. But unlike many of the books you pick up, this one will let you know just why you are selecting these options, and give you the important screen shots that go with them.

Start off by opening the Operator Console. You should be presented with the Administration view. If you are not, select the Administration view from the View menu at the bottom of the Navigation pane. You can see the discovery options by clicking the Computer and Device Management Wizard link in the Navigation pane. You will also see the Computer and Device Management Wizard option when you right-click on any of the nodes.

SUPPORT LIMITATIONS

When planning the deployment of your agents, make sure you know what is and is not supported for agent monitoring. You are able to monitor systems in a cluster, but you have to install the agent on each of the physical nodes in the cluster; you cannot install the agent on the virtual server. This could mean that you will encounter duplicate events on the two systems, so you may need to plan how you are going to monitor and replicate information between the nodes.

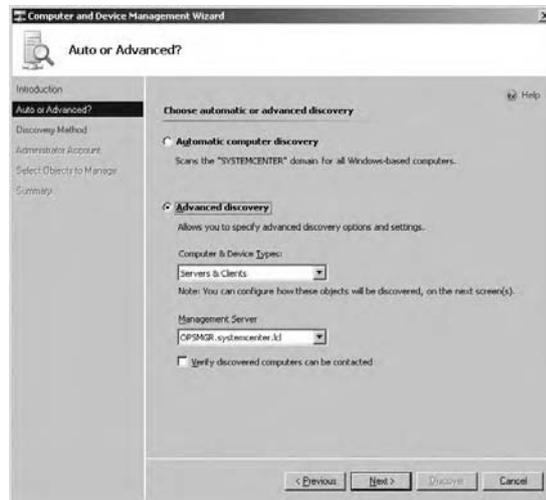
Windows NT 4 is not agent supported. If there are any remaining Windows NT 4 systems in your environment, you will need to use agentless monitoring on them. Or, if possible, you can upgrade the operating system to a supported platform.

Before you will be able to install the client using the Computer and Device Management Wizard, make sure that all the required ports are opened between the server and the client. The management server responsible for deploying the client has to have both the Server Message Block (SMB) port, which is TCP/UDP port 445, and the Remote Procedure Call (RPC) port, TCP port 135, opened on the server as well as the client. If either system does not have these ports available, the deployment will fail. On top of that, if there is a firewall in between the systems, the deployment will probably fail due to requirements with the RPC port range. Most administrators will not open all the ports that can possibly be used by RPC, so a manual installation will be required in that case.

When you start the Computer and Device Management Wizard, you are presented with an introductory splash page. As with many of the wizards, you are given information as to how you should proceed. Luckily, you are also given the option to prevent this page from showing up again. Click the Do Not Show This Page Again check box at the bottom of the page and you will not have to look at it again.

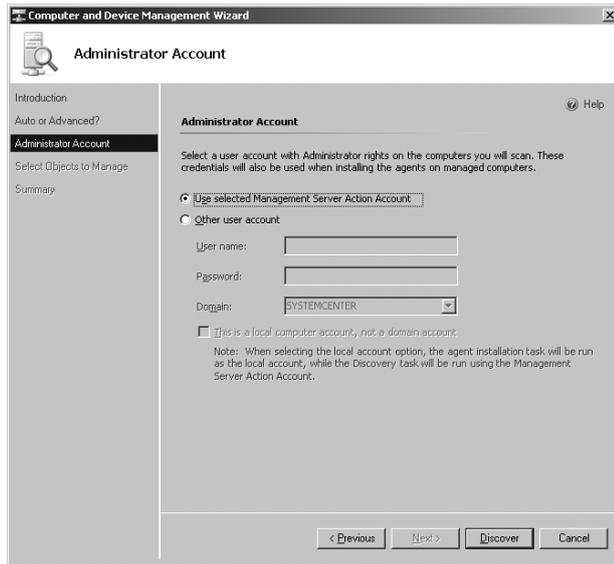
As shown in Figure 4.1, the first of the configuration options in the Computer and Device Management Wizard asks you if you would like to perform an automatic or advanced discovery. The default option at this point is to perform an advanced discovery. If you select the pull-down menu, you have the option to discover Servers & Clients, Servers Only, Clients Only, and Network Devices. Choosing one of these options and clicking Next gives you a greater level of control over your agent rollout than using the automatic method. If you choose the automatic method, all the systems within your domain will be discovered.

FIGURE 4.1
Computer and Device
Management Wizard
options



In Figure 4.2, we've selected the automatic option, and the screen that opens is the Administrator Account page. Here you get to choose the account that you will use to install the agent on the managed computers. Usually, you will choose the Action Account option. When you created the Agent Action account earlier, you granted it the ability to install the agent on the managed computers. Note that if the systems you are managing are not in a trusted domain, you will have to provide credentials for an account that does have the appropriate rights to install software. Once you have selected the Administrator account, click Discover to let the management servers start discovering systems to manage.

FIGURE 4.2
Automatic discovery options



If you choose the advanced discovery option, you are presented with the Discovery Method page, as shown in Figure 4.3. Here you can choose whether you want to scan Active Directory for accounts or enter computer names. If you choose to enter the computer names, make sure that the names are separated by semicolons. This gives you a lot of control over exactly which systems will be targeted for agent install. If you decide to test the agent on a small group of systems, this is an efficient and easy method to use.

When scanning Active Directory to find computers, you can find all of the systems in the domain by selecting the domain name and leaving the object types option empty. But at this point we assume that you did not want to do that; otherwise, you would have chosen the automatic discovery option. Clicking the Configure button allows you to specify exactly the systems that you would like to discover in the domain. As you can see in Figure 4.4, there are two configuration tabs: Computers and Advanced. If you have a naming convention that you use for your systems, type in the first part of the computer name and the wizard will automatically apply a wildcard at the end of the name for you. You can use the * symbol to denote a wildcard at the beginning and end of the naming convention. Let's say you want to look for all systems that contain the characters *CHI* within the computer name. You would enter **CHI** at the prompt, as shown in Figure 4.5. For our example, we have entered **DC**, and the resulting query appears in Figure 4.6.

FIGURE 4.3
Advanced discovery options

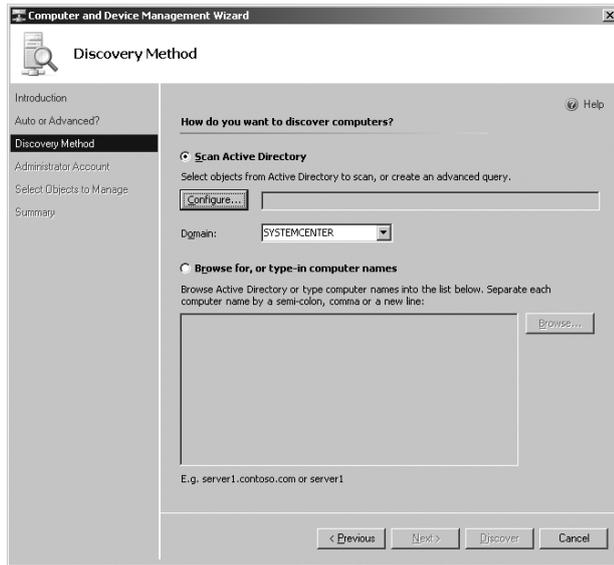


FIGURE 4.4
Determining systems to discover

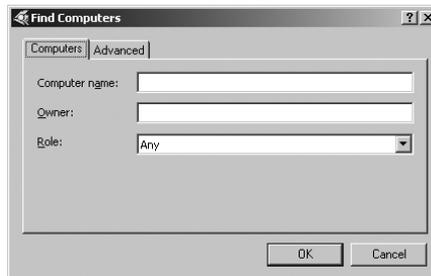


FIGURE 4.5
Searching for systems that include CHI in the computer name

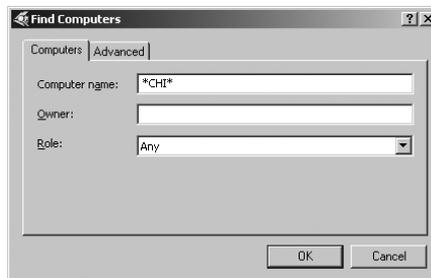
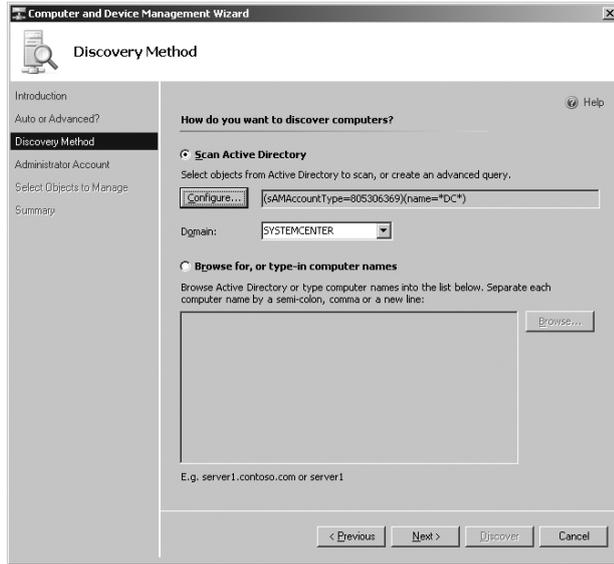


FIGURE 4.6
Query used to discover systems with names that start with DC



The Advanced tab allows you to specify the fields in the computer account that you would like to use for discovery, as well as the details that you would like to use when creating the discovery query. As you can see in Figure 4.7, the fields that are available are Computer Name, Description, Managed By, Name, Operating System, and Operating System Version. In Figure 4.8, you can see the conditions that are available to build the query criteria: Starts With, Ends With, Is (Exactly), Is Not, Present, and Not Present.

FIGURE 4.7
Advanced search options

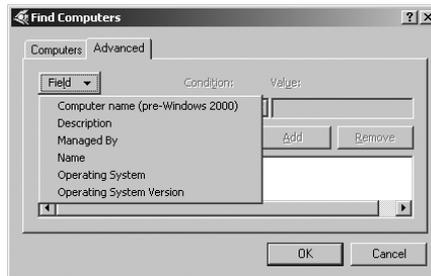
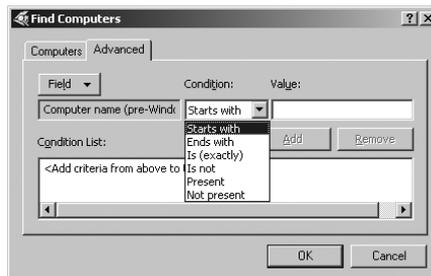


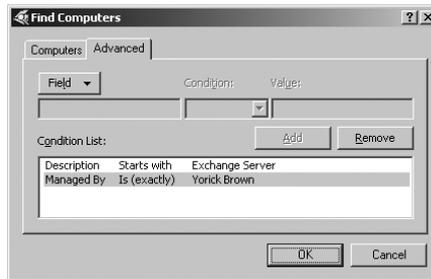
FIGURE 4.8
Conditions available to define the search



Using these options, you can build a fairly sophisticated query. Let's assume that you want to discover all of the servers that contain *Exchange Server* in the Description field and that are managed by Yorick Brown. You could create a query that looks like the one shown in Figure 4.9. Of course, the results are only as good as the data that is entered in the Active Directory object for the systems you are discovering. If anyone has been lax in entering the required information, you may receive incomplete results, or worse, incorrect results.

FIGURE 4.9

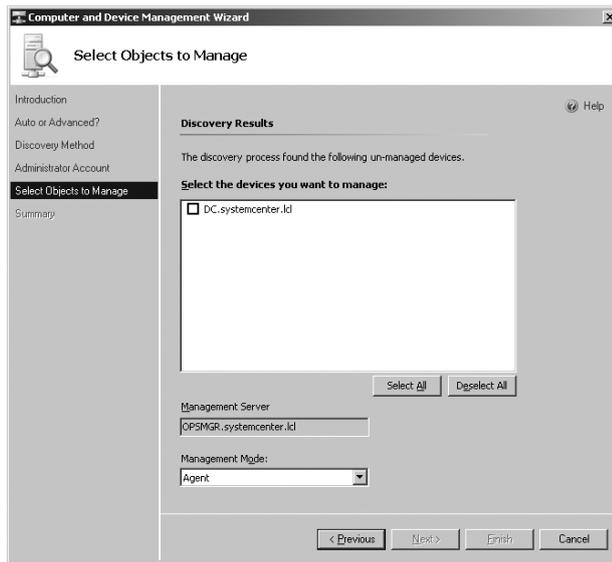
Query used to discover Exchange servers managed by a specific user



After choosing your discovery options, you are presented with the same Administrator Account page that we looked at when discussing the automatic discovery option. The final page that appears is the Select Objects To Manage page, shown in Figure 4.10. Here you can select not only the systems you wish to manage, but also which management server will be responsible for the selected systems and the initial management mode. Select the check boxes beside the appropriate systems, pull down the list of management servers, and select the one to you want to manage the selected systems. When choosing the management mode, remember that you cannot collect detailed information about a system if it is monitored using the agentless method. To get a full array of data, you will have to install an agent on your systems.

FIGURE 4.10

Select Objects To Manage page



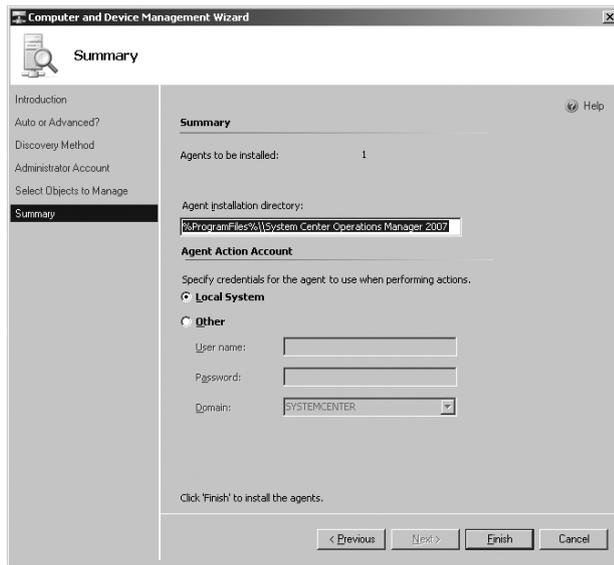
You need to keep something in mind when selecting the management server from the list of available management systems: when a management group contains more than one management server, the other management servers that are not selected become the failover systems for the agent. That way, if the assigned management system fails or cannot be contacted, one of the other systems can start accepting the management information passed from the agents. Once you have selected the appropriate options, click Next, and you will find yourself at the final wizard page, the Summary page.

The Summary page may look a little deceiving to those of you who are familiar with the typical wizard summary page that simply details the options you have selected as you've gone along. Instead, the Summary page has a few final options, as shown in Figure 4.11. The options of interest here are the Agent Installation Directory and Agent Action Account. For the Agent Installation Directory, enter a path that is available on all the systems where the agent will be installed. The default path, `%Program Files%\System Center Operations Manager 2007`, is usually a viable option. The `%Program Files%` variable will be evaluated for each system and the path that is specified in the environment variables will be used to install the agent. You can specify other paths explicitly by entering the exact path, or you can use other environment variables such as `%windir%` or `%systemroot%`.

The Agent Action Account is an account that will have administrative privileges on the managed system. Usually you can select the Local System option and not worry about it. Local System is the only account that is allowed to use the system's computer account for authentication purposes. By using the Local System account, you don't have to worry about anyone trying to log in with the account, since it cannot be used as an interactive account. This gives you a higher level of security than if you were to use a domain account.

Once you click Finish, the agent installation will occur. You will notice that the status will go into a Queued state, and then the installation will be performed. If all goes well, you will receive a Succeeded status.

FIGURE 4.11
Options on the
Summary page



Installing Manually

As you read through the previous section, you may have wondered why you would want to perform a manual installation. For most instances, you won't want to. However, there are situations that will force you to make that decision. For instance, if you have a firewall in place, hardware- or software-based, and you cannot open the ports to allow the automated install, you will be forced to find another installation method. Of course, you could use a software distribution method such as Microsoft's Systems Management Server to push the install to the client instead of walking to each of the systems to install the agent. No matter how you end up installing the agent, if you do not allow a management server to install the agent, Operations Manager will consider it a manual installation.

Before you perform a manual installation, make sure that you allow the management servers to allow manually installed agents to register with them. As you saw in the previous chapter, you can select the option to allow manually installed agents on a per-management-server basis.

After setting up the management server to allow manual installations, you then need to go about installing the agent. As we mentioned earlier, you can walk around with the CD and run the setup files, you can run Setup from a network share, or you can use a software distribution method such as Group Policy or Systems Management Server.

Installing from CD

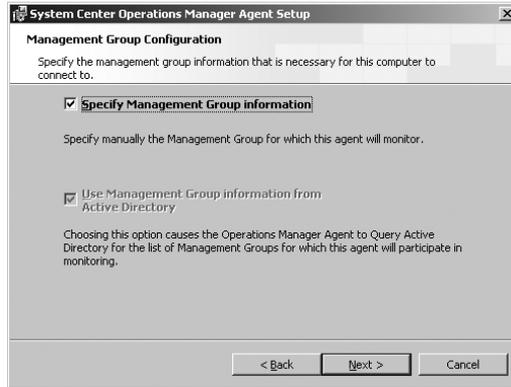
Whenever you perform a manual installation, the setup needs to be performed by an account that has local administrator privileges. Normally, you will not want to give your users the ability to install software, let alone be a member of any administrative group. So that means it is up to you, or someone on your administrative staff, to go around and install the software.

After dropping the CD in the drive, you will be presented with a wizard splash screen that gives you some options. At this point we want to install the agent on the client machine, so select the option Install Operations Manager 2007 Agent. This will start up the agent installation wizard, where you can go on past the typical Welcome page and start configuring the settings. First up is the Destination Folder setting. If you do not want to choose the default location, `C:\Program Files\System Center Operations Manager 2007\`, then click the Change button and select the location where you want to place the files.

Figure 4.12 shows the Management Group Configuration page that you see when you click Next on the Destination Folder page. Two options appear here: Specify Management Group Information (the default) and Use Management Group Information From Active Directory. For those organizations that have not extended Active Directory to include the Operations Manager schema extensions, leaving the Specify Management Group Information check box selected is their only option. If you have extended the Active Directory schema, you should deselect that check box, which enables the Use Management Group Information From Active Directory option.

Using Active Directory as the repository for management group information only makes sense. As with many other server products, Systems Management Server, Configuration Manager, and Exchange, the configuration information that is used by these products is replicated throughout the organization's infrastructure and stored on domain controllers. With proper planning, domain controllers can be placed strategically so that they are near the systems they support, not just the users who use them for authentication. Having one replication model to be concerned about makes troubleshooting much easier.

FIGURE 4.12
Management Group
Configuration page



Leaving the Specify Management Group Information check box selected means that you will have to provide more information to the wizard instead of being able to automatically configure the agent using Active Directory. Thus, you will see two additional wizard screens that would not appear if you had deselected the check box. The first of these pages is named Management Group Configuration, just as the prior page was. However, this page, shown in Figure 4.13, requests the Management Group Name, Management Server, and the Management Server Port entries.

FIGURE 4.13
Second agent-
configuration page



The Agent Action account that you enter on the following page is an account that has the rights required to perform actions on the managed system. As mentioned in Chapter 3, you need to carefully consider how you will configure the Agent Action account because this account has a lot of rights on the managed servers. Your best option is to use the Local System account because it cannot be used as an interactive logon, thus protecting you from an attacker trying to use the account. If you do select the Domain Or Local Computer Account option, you will be presented with text boxes where you can enter the appropriate credentials, as shown in Figure 4.14.

At this point, when you click the Next button, you will arrive at the Ready To Install page, which is nothing more than a summary screen of the actions you chose during the wizard. If you decide to change anything before moving on, you still have the chance to click the Back button. Once you click Install, you will commit the changes and install the agent on the server.

FIGURE 4.14
Entering credentials



Other Manual Installation Methods

Some administrators like to control their software installations by using a software distribution method. These tools are handy for most software installation purposes, but you should try to avoid them when installing the Operations Manager agent. The main reason we say that is because you don't have to open a potential security hole by allowing manual installations. Microsoft has made it easy for an administrator to push the agent out to the clients that will be monitored, so you should take advantage of that functionality.

If you would like to use another distribution method, note that no automated installation scripts or unattended installation files come with the agent. You will have to create your own "wrapper" or program to perform a silent install, or you will have to use the installation switches that are available with the agent install file `MOMAgent.msi`. Otherwise, you will have to give your users the ability to interact with the setup wizard as it runs, and instruct them on the proper entries to use when it runs. Most administrators, myself included, have a hard time trusting users with a vital component of a monitoring solution. More often than not they carelessly click a button or "fat-finger" an entry, leaving more work when trying to troubleshoot an issue.

Several switches are available for use when installing the agent; in fact, every option that you have when using the wizard is available from the installation file. Here is a quick rundown of the options and possible values:

ACTIONS_USE_COMPUTER_ACCOUNT This switch allows you to configure the account that will be used as the Agent Action account. The two values that are available are 0 and 1. Setting the value to 0 will tell the installer to use a domain-based account for the Agent Action account. Setting the value to 1 uses the Local System account. The default value is 1.

Using a value of 0 requires setting properties for the `ACTIONSUSER`, `ACTIONSDOMAIN`, and `ACTIONSPASSWORD` switches.

ACTIONSUSER This switch is used in conjunction with the `ACTIONS_USE_COMPUTER_ACCOUNT` switch to specify the domain-based user account required by the Agent Action account. The switch is not required when the `ACTIONS_USE_COMPUTER_ACCOUNT` value is set to 1.

ACTIONS_DOMAIN Use this switch in conjunction with the `ACTIONS_USE_COMPUTER_ACCOUNT` switch to specify the domain in which the Agent Action account resides. The switch is not required when the `ACTIONS_USE_COMPUTER_ACCOUNT` value is set to 1.

ACTIONS_PASSWORD This switch is used in conjunction with the `ACTIONS_USE_COMPUTER_ACCOUNT` switch to specify the password required by the Agent Action account when the service authenticates. The switch is not required when the `ACTIONS_USE_COMPUTER_ACCOUNT` value is set to 1.

USE_SETTINGS_FROM_AD This switch has possible values of 0 and 1. The default is 0, which means the management group settings will have to be entered via the switches `MANAGEMENT_GROUP` and `MANAGEMENT_SERVER_DNS`. If the DNS name and Active Directory name of the management server differ, the `MANAGEMENT_SERVER_AD_NAME` switch will have to be supplied also.

If this switch is set to a value of 1, the agent will query a domain controller for the configuration settings.

MANAGEMENT_GROUP When the `USE_SETTINGS_FROM_AD` switch is set to 0, the value that is supplied by this switch will identify the management group name that the agent will be a member of.

MANAGEMENT_SERVER_DNS When the `USE_SETTINGS_FROM_AD` switch is set to 0, the value that is supplied by this switch will specify the hostname of the primary management server of the management group.

MANAGEMENT_SERVER_AD_NAME When the `USE_SETTINGS_FROM_AD` switch is set to 0, the value that is supplied by this switch will specify the Active Directory name of the primary management server for the management group. This switch is used only if the DNS hostname and the Active Directory object name differ.

SECURE_PORT The `Secure_Port` switch is used to specify the port that is used for communication between the agents and the management servers. The default value is 6270.

If you choose to install the agent from the command line and you are going to use Active Directory, the command line is not too difficult to work out. Let's say that you want to install the agent onto a system and you have stored the installer file on a server named `ZygortFS` in the `OpsMr 2007` directory. You also want to use the Local System account for the Agent Action account. The command-line syntax would look like this:

```
%WinDir%\System32\msiexec.exe /i
➤ \\ZygortFS\OpsMgr 2007\momagent.msi /qn
➤ USE_SETTINGS_FROM_AD=1 ACTIONS_USE_COMPUTER_ACCOUNT=1
```

If you were to change the previous example to use the domain user account `BloomM` for the Agent Action account, the command line would change to look like the following:

```
%WinDir%\System32\msiexec.exe /i \\ZygortFS\OpsMgr
➤ 2007\momagent.msi /qn USE_SETTINGS_FROM_AD=1
➤ ACTIONS_USE_COMPUTER_ACCOUNT=0 ACTIONUSER=BloomM
➤ ACTIONS_DOMAIN=zygort ACTIONS_PASSWORD=P@ssw0rd
```

For an example of an installation that does not take advantage of using Active Directory, let's assume that you are going to use the same accounts and settings, but you want to manually specify

that the agent is going to be installed into the Exchange management group and the ExchOM1 server is going to be the primary management point. Here is what the command line would look like:

```
%WinDir%\System32\msiexec.exe /i \\ZygortFS\OpsMgr 2007\momagent.msi /qn
➤USE_SETTINGS_FROM_AD=0 MANAGEMENT_GROUP=Exchange
➤MANAGEMENT_SERVER_DNS=zygort\ExchOM1 ACTIONS_USE_COMPUTER_ACCOUNT=0 ACTIONUSER=BloomM
➤ACTIONSDOMAIN=zygort ACTIONSPASSWORD=P@ssw0rd
```

The command-line options can also be used for creating a program for Systems Management Server 2003, System Center Configuration Manager 2007, or even Group Policy software installs. But once again, it must be stressed, any type of installation other than the push installation that is performed using the Computer and Device Management Wizard is considered a manual installation and will be rejected by the management servers unless you explicitly allow manual installations.

A popular method of deploying software is to preload the software on an operating system image. This way, you can install the software once, and then when you put the image on a new system, all of the software is already in place and ready to go. For your agents, this is not a good idea. Again, this is considered to be a manual installation and the management servers are configured to reject by default. Think hard about the security implications you may be introducing into your monitoring solution before you put the agent into an image.

Multihomed Agents

We have touched on how to make your management servers work together to provide failover in case the primary management server fails. You can also configure an agent to communicate with multiple management groups at the same time. This allows you to have separate management groups for specific monitoring purposes and have a managed system report to each one.

Take, for instance, an Exchange server. Exchange relies on Internet Information Server (IIS) for the underlying protocol, which is used to send and retrieve mail as well as generate web pages for Outlook Web Access. Many companies have different groups responsible for Exchange and IIS. If this is the case, each one of the teams would have its own Operations Manager management group. The Exchange server would then have an agent installed and that agent would be configured to send Exchange-related events and alerts to the Exchange group, while the IIS group would receive IIS-related events and alerts from the agent.



Real World Scenario

TOO MANY MANAGEMENT GROUPS

The example that we give here is an actual case. A large client that we work with has divided up the administration and monitoring responsibilities to a great extent. They have one team responsible for Active Directory, another for DNS, another for Exchange, another for IIS—you get the point. Exchange became an interesting problem for them: Exchange uses IIS and they had two groups, one for each of those products, and they wanted to monitor them separately.

On top of that, they also had an antivirus program that the security team wanted to monitor, the server manufacturer had provided a hardware-monitoring management pack that the infrastructure team wanted to monitor, and the Windows operating system team wanted to monitor the base operating system.

Since Operations Manager only allows for an agent to work with four management groups, something had to give. It was determined that since they were not taking advantage of Outlook Web Access for their clients, they would allow the Exchange team's management group to collect the IIS data. A separate program was then written that copied the data from the SQL database in the Exchange management group to the IIS management group. Although the information was not provided in real time as it was with the other IIS servers, the IIS group could still report on all the IIS implementations.

You could take this approach one step further and have another group that is responsible for monitoring and maintaining the Windows Server operating system and another responsible for the antivirus solution. With Operations Manager 2007, you can have an agent interacting with up to four management groups.

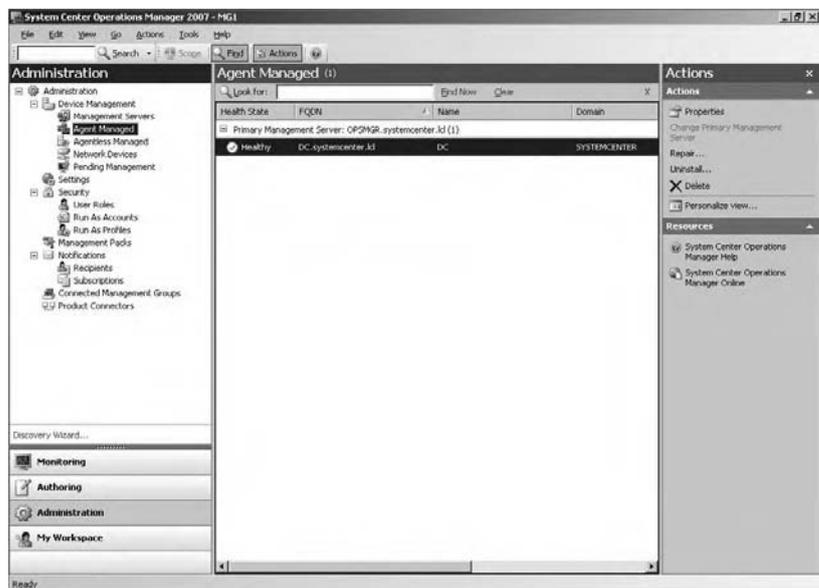
To do so, you do not have to install up to four different agents; instead, you can install one agent, and when you perform a discovery with another management group, the agent is configured to communicate with that management group as well. If you are using Active Directory to host the Operations Manager configuration data, you can configure Active Directory with the necessary information to tell the agent which management groups to start working with.

Adjusting Individual Agent Settings

Since the interface for Operations Manager has changed from the familiar MOM 2000/2005 Administrator Console/Operator Console that we have grown to know and love, we will have to get used to the new Operations Manager Console. As you have seen already, the configuration, administration, and monitoring is all performed from this one console.

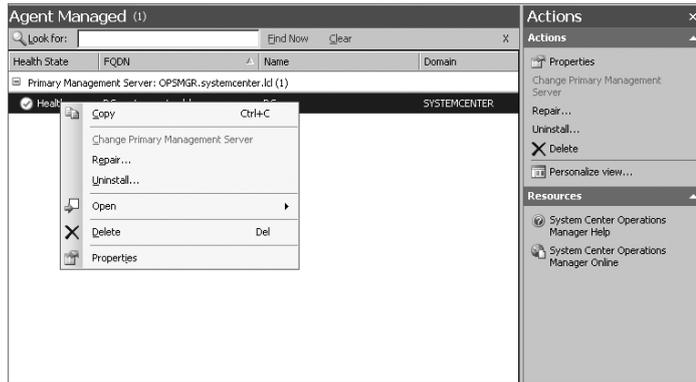
When you are ready to configure the agent settings, open the Operations Manager Console and select the Administration view. In the Navigation pane, expand Device Management and select the Agent Managed option. You should be presented with a view that looks like Figure 4.15. Each of the agent-managed systems will appear in this view. If you think you have installed the agent on systems that do not appear here, open Pending Actions from the Navigation pane to make sure the client has been allowed to interact with your management group.

FIGURE 4.15
Agent-managed
systems



Notice the Actions pane on the right side of the screen. Here you can work with the agent that is installed on the agent-managed system that is selected in the Details pane. These same options can also be found on the context menu when you right-click on any of the agent-managed systems, as shown in Figure 4.16. Using either of these two methods you find Properties, Change Primary Management Server, Repair, Uninstall, and Delete.

FIGURE 4.16
Context menu options



Properties

When you select the Properties option, the configuration items that are shown (or that are missing) might surprise you. Two tabs appear on the Properties page: Heartbeat and Security. The Heartbeat tab, shown in Figure 4.17, shows how often the agent will generate a heartbeat, which is sent back to the management server. As you can see, you can change the Heartbeat Interval (Seconds) setting to the time value that suits your needs. The interesting thing about this page is that even though you can change the time value, the agent won't actually change its heartbeat rate until you select the option to override the default setting. In Chapter 3 we configured the default settings for the management group. By default, all the managed systems will use those values until you override them on a per-agent basis by selecting the Override Global Agent Settings check box.

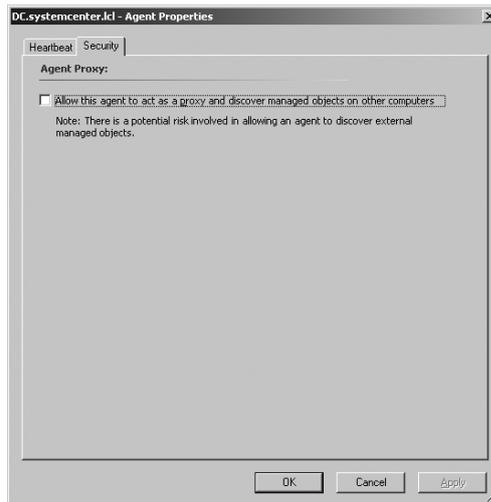
FIGURE 4.17
Heartbeat settings



Once you select the Override Global Agent Settings check box, you can then change the heartbeat rate to any value that you would like. You can even completely stop the heartbeat by deselecting the Enable Agent Heartbeat check box. If you decide to change either of these settings, make sure you understand the ramifications of your choice. Operations Manager is going to look for a heartbeat from each of the agents that it knows about. If the agent has not sent a heartbeat when Operations Manager thinks it should, Operations Manager will perform a ping attempt against the agent-managed system. So if you decide to change the heartbeat to once every 6,000 seconds, or 10 minutes, but the management server is configured to check every 180 seconds (the default), you may cause Operations Manager to perform additional work and generate unnecessary events and alerts.

The Security tab, as shown in Figure 4.18, presents you with a single setting, and this is an important setting depending on the computer you are managing. The Allow This Agent To Act As A Proxy And Discover Managed Objects On Other Computers setting enables the agent to send information to the management server on behalf on unmanaged systems. While this could potentially be a security risk—someone could take advantage of the proxy—there are times when you may need to use this setting. One case in point is when you are monitoring Exchange servers. In order to discover the Exchange server topology, you need to allow the agent to discover Exchange services that can then be monitored, thus allowing for a full end-to-end monitoring solution.

FIGURE 4.18
Security tab controlling proxy and discovery of externally managed objects



Open

Choosing the Open menu option allows you to quickly change to another view. The views that are available from this option are Event, Alert, Performance, Diagram, and State. If you are looking at the list of agent-managed systems, you can quickly check out the current state of one of the computers by simply choosing the state view.

Change Primary Management Server

If you are running more than one management server in your management group, each of the agents will be assigned a primary management server. As long as the primary management server is online, the agent will send any events or alerts that it has captured to the primary management server. This is also the management server that is responsible for notifying the agent about the configuration details. The agent also sends heartbeats to this management server.

If the management server goes offline or the agent loses contact with it, the agent will start communicating with a backup management server, also known as the failover system. This redundancy allows the managed system to continually stay in touch with a management server, keeping the monitoring available. When the agent can finally start communicating with the primary management server again, it will automatically return to communicating with the primary.

There may be occasions when you would want to change the primary management server. There are times when you may want to decommission a server. Or you may add in another management server to your management group. In either of these cases, make sure that the clients are communicating efficiently. In the case of decommissioning a server, change the primary management server for those agents that used that server. That keeps the clients from trying to reestablish communications with a server that no longer exists. When you add an additional server into your management group for load-balancing purposes, redistribute the load so that the new server has some work to do and the original servers are relieved of some of their responsibilities.

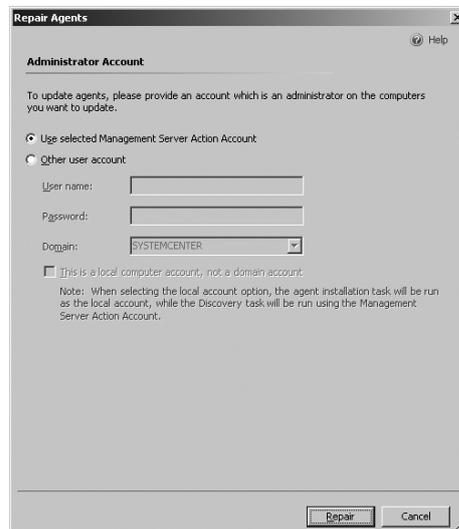
Repair

This menu option does exactly what it sounds like it does: it repairs the agent on the agent-managed server. So if you have an agent that is not responding, or the files for the agent have become corrupted and the agent will not run, you can reinstall the agent using this option.

There is another function for this option, however. If at any time you make a change to the account that the agent runs under, you can use the Repair option to reset the account without having to walk to each machine and make changes in the services settings. Of course, this assumes that you are using a domain-based user account for the Agent Action account instead of the Local System account. Additionally, you can change from using a domain-based user account to the higher-level security option of using the Local System account using this option.

Figure 4.19 shows the Repair Agents page, which allows you to select an account that has permissions to install software on the managed server. Just as when you installed the agent, you have the option to use the Agent Action account or you can specify a domain-based user account. Just remember that any domain-based account you use must have the ability to install software on the managed system. If you are using advanced clients such as Windows XP or Windows Vista, most of your user accounts will probably not have that capability.

FIGURE 4.19
Repair Agents



Uninstall

When you select this option, you are going to remove the agent from the agent-managed system. It's as simple as that. Well, actually, you first have to go through a single wizard page that asks you for an account to use for the uninstall program. Typically you would use the same account you used to install the agent. The wizard page that appears looks almost exactly like Figure 4.19, but you will find an Uninstall button to click.

Delete

One question may pop into your mind as you look at the two options, Uninstall and Delete: "What is the difference between these two options?" Whereas Uninstall will remove the agent from the agent-managed system, Delete will remove the agent-managed system from the Operations Manager database.

The Bottom Line

Automatically install the agent by using a discovery. The best method for installing the client is to use a discovery from the Administration view in the Operations Manager console. A discovery will allow you to configure which systems will receive the agent as well as automatically install the agent and configure it to start monitoring in the management group.

Perform a manual installation from CD and command line. There are times when you cannot use the automated method to install the agent. Especially if there are firewalls in place, you will have to perform some type of manual installation. However, when you do so, you will have to make sure the management server allows for manual installations. The default setting is to reject manual installations.

Automate the manual install of the agent by using command-line options. The Windows installer file that is used to install the agent allows for several command-line options to help silence the install. These options come in especially handy when you are using a software distribution method such as System Center Configuration Manager.

Override the default management group settings on an individual client. Not every system is created equal, so you may need to alter the settings on some of the agents so that they do not match the management group default settings.

Master It! After deploying three Operations Manager 2007 servers in your management group, you need to assign agent-managed systems to each one. You have installed the Active Directory container for Operations Manager and want to use Active Directory to make the assignment of each agent easy.

Configuring Active Directory:

1. Open the Operations Console by selecting Start > System Center Operations Manager 2007 > Operations Console.
2. In the Operations Console, select the Administration view.
3. In the Administration view, expand Administration, then expand Device Management.
4. Select the Management Servers container.
5. Right-click OpsMgr1 and select Properties from the context menu.

6. On the Agent Management property page, click Add.
7. When the wizard starts, click Next and then select Assign Agents To This Management Server In A Trusted Domain.
8. Enter **Zygort** in the Domain Name text box and click Next.
9. On the Inclusion Criteria page, click Configure.
10. On the Find Computers page, click OK, then click Next.
11. On the Create An Exclusion Rule page, click Next.
12. On the Agent Failover page, click Create.

Creating the discovery rule:

1. In the Operations Console, click Discovery Wizard.
2. On the Introduction page, click Next.
3. On the Auto Or Advanced page, make sure Advanced Discovery is selected and verify OpsMgr1.zygort.lc1 is entered as the management server, and then click Next.
4. On the Discovery Method page, click Configure; then click OK and Next.
5. On the Administrator Account page, click Discover.
6. Once the discovery is finished, select the check boxes by the two discovered systems, leave the management mode as Agent, and click Next.
7. Leave the settings on the Summary page at their defaults and click Finish.
8. Once you get a Success message, click Close.
9. To verify that the systems are now considered agent-managed, click on Agent Managed in the Operations Console and verify that both systems appear.

Chapter 5

Managing Management Packs

Up to this point we have put in place the management server, database, and agents. We have configured them so that they are ready to monitor each of our servers. Global settings are in place, and specialized settings for those unique systems have been configured. Now we need to tell the agents what they are supposed to start looking for and how they are to react when certain criteria arise.

In the past, monitoring meant that you would check a service periodically to see if it was running. When it failed, you could be notified by the operating system, or you could have it automatically restart. Windows 2000 and Windows Server 2003 have the ability to restart services or even force the operating system to reboot if a critical service is not running. That type of monitoring gives you limited control over your systems.

Watching for a service to be in a running or failed state is rudimentary and is considered reactive. To move toward proactive management, you must watch how the services and systems are running, looking at the health of each to determine whether problems are starting to arise. Operations Manager takes monitoring to the proactive side by looking at each of the services and the components that make up that service, watching to see how they are running. When services can be monitored for their health, you have the ability to recognize which systems are starting to have problems and you can rectify the situation before you have a serious issue.

Management packs give us the ability to fine-tune our monitoring. If you take a look at a management pack, you will find several rules that can be used to start collecting information about a managed system. However, most of the settings are turned off so that we don't have too much information being collected. "Noisy" rules can quickly fill up our database with unnecessary information and consume network bandwidth as well as system resources.

If you are familiar with Microsoft Operations Manager 2000 or 2005, you will notice several changes with management packs. First, they are now formatted using XML. You can use any text editor to modify them, although you will need to know what you are changing. Another change is the introduction of sealed management packs. You cannot make changes to sealed management packs. That means that you cannot directly modify which settings are enforced and which ones are ignored. To specify whether a monitor or rule is turned on or off, you must configure overrides. MOM 2005 had the ability to use overrides, but now they are vital when it comes to fine-tuning your monitoring solution.

In this chapter you will learn how to:

- ◆ Identify management pack formats
- ◆ Work with discovery rules
- ◆ Create groups based upon discovery
- ◆ Create monitors and rules

Understanding Monitoring Capabilities

Operations Manager 2007 provides two monitoring options: agentless monitoring and agent-managed monitoring. In an agentless scenario, the management server is responsible for checking on a server to see how it is performing. Using this method, you can watch some of the processes and services, but you are limited by what you can collect. Also, without having an agent on a system, you cannot monitor systems that reside on the far side of a firewall. To do so would mean that you would have to open up too many ports in your firewall to allow RPC communication to slide through. Doing so would negate the firewall's effectiveness.

Agent-managed monitoring is the most effective method of monitoring. Agents running on each system have the ability to look at more of the system. Since the agent runs as a service with system-level rights, it has access to parts of your system that you could only hope to have one day. The agent is able to parse through your event log files, view application log files, dig into the Windows Management Instrumentation (WMI) database, and monitor services at their lowest level. Using the agent, you can monitor nearly every aspect of your servers and workstations in your organization.

Now you are probably wondering why you would ever want to perform agentless management. Honestly, in most scenarios you won't. But there are times when you may have to. Some companies are still supporting servers running Windows NT 4. Even though the support life cycle has officially ended for NT 4, due to reasons beyond their control, some administrators are still stuck supporting those servers. Sometimes it may be a budgeting issue; other times it may be a critical business application that is not supported on another platform. And many times, getting rid of the NT 4 systems is easier said than done.

Another instance in which you may need to perform agentless management is when you cannot install the agent due to software conflicts or hardware limitations. You should always test installing the agent on a nonproduction system to make sure there are no conflicts when it runs in conjunction with other software. If the problem is not software compatibility but hardware limitations, you can perform agentless monitoring until you improve your hardware to support running the agent.

No matter which type of monitoring you end up using, the monitoring criteria come from the management packs installed on the management server. If the management server is the heart of our monitoring solution, then the management pack is the lifeblood. Without the management pack, the management server cannot do its job. It can still discover systems and install the agent on them, but then it will sit there, idly waiting for something to do.

Identifying Management Pack Requirements

If you were to simply install Operations Manager and configure the settings necessary to install the agent on the systems you would like to monitor and then do nothing else, you would find that Operations Manager would start monitoring the operating systems on those systems. By default, a core set of management packs is installed and made available as soon as the software is installed. You will notice that you do not get every management pack available. You must choose exactly which management packs will be used in your environment, going so far as deciding which ones will be installed in each of your management groups. It only makes sense that you do not want to import the Exchange management packs into the SQL management group, where they would not be used.

At the same time, you want to make sure that you install only the management packs that are necessary, because the discovery rules that are configured for each management pack will still reach out and try to determine whether systems are available for it to monitor. If you install the DNS management pack into the Exchange management group, and there are any DNS servers for

it to discover, it will find them and start monitoring them, which could introduce more data into your database as well as present alerts to service desk personnel who are not responsible for monitoring those systems.

When you are importing the management packs, make sure that you know all of the management packs that are required, as well as the parent management packs that they rely on. Usually you will have documentation that ships with the management pack that will inform you of the requirements. Make sure you read the documentation and follow any instructions that are included. Some management packs even come with a setup and configuration wizard that aids in the rollout of complex management packs.

Current Management Packs Available

Management packs come in several flavors. You can pick the ones you like the best and ignore the rest. In other words, you can install the management packs that you need to use in order to support the applications, servers, and hardware in your organization—you don't have to install every management pack that is available. In fact, you really don't want to. The more management packs you install, the more space you take up in the database and the more cluttered your reporting solution becomes. Determine which applications, servers, and hardware you would like to monitor and then obtain the management packs to support them.

Most vendors are supplying management packs for their products. Microsoft's philosophy is that they will not ship a product until the management pack is available. That doesn't include things like the Xbox game console or the Zune portable audio player, but their business and enterprise products are included. Just think if they did support some of those products and we could monitor when our kids were using their game consoles, which games they were playing—and, well, never mind; let's move on.

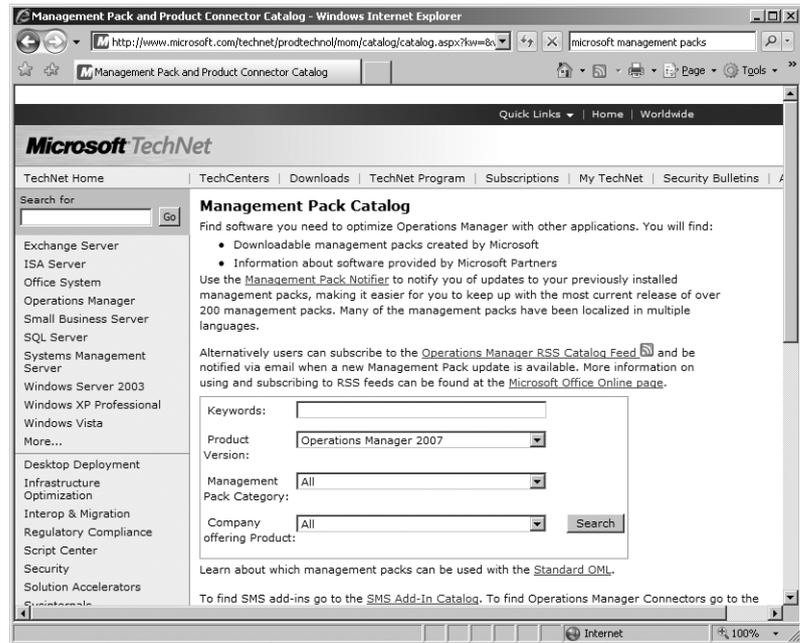
As we mentioned, Microsoft has determined that all of their server products, such as Exchange, SQL, and SMS, as well as their workstation operating systems, will have a management pack available when they are released. Exchange Server 2007, System Center Configuration Manager 2007, SQL Server 2005, and many others are supported. Other vendors are also supplying management packs for their products so that you can monitor them using a single solution.

You can find the management packs that are currently available by going to <http://www.microsoft.com/technet/prodtechnol/mom/catalog/catalog.aspx?kw=&vs=2007&ca=&co=AT>, as shown in Figure 5.1. Here you are able to search for management packs or view a complete listing. Using this website, you can determine whether the application that you are planning on implementing is currently supported by Operations Manager 2007. Many vendors realize that organizations may be making decisions based on an end-to-end solution that includes support for monitoring using their monitoring solution. You can expect this list to expand greatly over the next few years.

Requirements for Management Pack Monitoring

Implementing a management pack involves importing the management pack and all of the parent management packs. Management packs within Operations Manager 2007 take advantage of inheriting settings from other management packs so that there is not much duplication. Another advantage is that the management packs can be smaller, more compact units, or they can be created to alter the initial settings of a parent management pack. The base management packs that have been made available for Operations Manager 2007 have several elements. In some cases, the elements are not configured to be used in the Operations Console; instead, the elements are defined within the base management pack, only to be used by other management packs that inherit the settings.

FIGURE 5.1
Microsoft's
management pack
search site

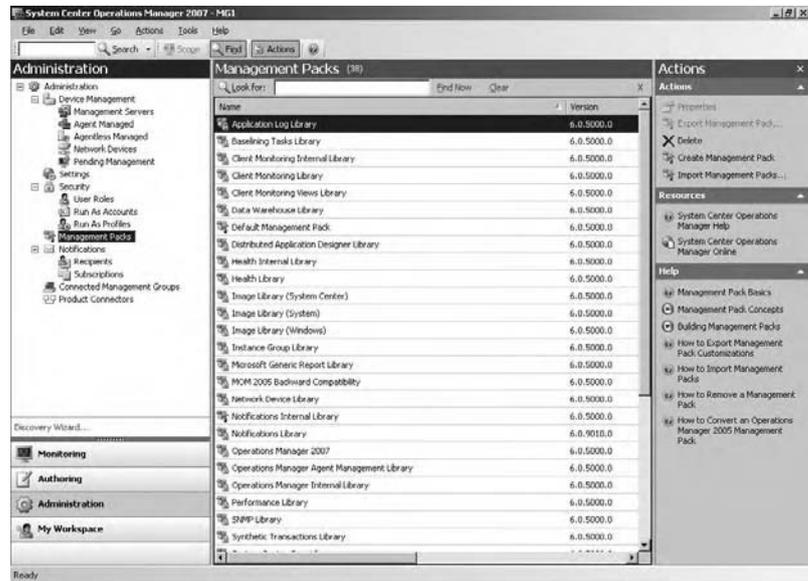


One case in point is the SQL Server management packs. When you decide to import the SQL Server 2005 management pack, you will find that you need to import not only the Microsoft SQL Server 2005 Discovery and Microsoft SQL Server 2005 Monitoring management packs, but also the SQL Server Library management pack. Common elements are defined within the library that can be used with SQL Server 2000 as well as SQL Server 2005 management packs. Unique configuration details are then defined within the monitoring management packs for the two separate products.

In Figure 5.2, all of the management packs that are installed in the management group are shown. The management packs that are installed by default are there to primarily monitor Operations Manager or to support other management packs. Most of these are called the libraries because they contain the initial building blocks for most of the other management packs that you will use. As with the SQL example in the preceding paragraph, you will find that the basic Windows operating systems libraries are included in this collection of management packs. Once defined, the object setting within these libraries will not have to be redefined in other management packs. Instead, the settings can simply be referenced from these libraries and the monitoring criteria can be set.

There is one management pack that you should take note of since it comes in to play quite often: the Default Management Pack. This is the management pack that is used by default when you create a new event, monitor, or rule. It is also the management pack that is used to configure any overrides that you may assign against sealed management pack settings. Of course, you can always create a new management pack to use instead of the default, but you will need to use one for overrides. For more information on the on the Default Management Pack, overrides, and saving settings to other management packs, see Chapter 11.

FIGURE 5.2
Management packs included with the installation



Exploring Management Packs

Management packs have matured since their introduction in Microsoft Operations Manager 2000. Originally, management packs were a set of instructions that controlled how the agent would perform on a client. In Operations Manager 2007, management packs are now XML-based files. This makes the management pack easy to work with, as well as standardized in the same format that Microsoft is using for many of its applications and servers. If you have worked with other applications that take advantage of XML, you should have no problem bringing that knowledge to Operations Manager.

When using XML files, you need to have a definition, or schema, for the configuration files. The schema defines how the system uses the entries within the XML files. For Operations Manager 2007, the schema file is named `MPSchema.xsd`. It should go without saying that you shouldn't modify the `MPSchema.xsd` file. Doing so could render the entire Operations Manager operation useless. Every management pack is written based on the schema definitions included in this file.

For those of you who are familiar with management packs as they existed within MOM 2000 or MOM 2005, you will find that they look nothing like they did in those products. You may want to grab a book such as *Beginning XML, 4th Edition* (Wrox Press, 2007), seek out information on the Web, or attend a class to get you up to speed. Typically you do not need to worry about directly editing a management pack, so it is not vitally important that you learn how to work with XML files.

We have alluded to the fact that developers are creating management packs for their own products. Usually they will take the time to optimize their product on systems of their own. They then build the management pack so that the settings contained within are already configured to what they feel are the optimal settings. They will then take it one step further and seal the management pack so that you cannot directly change the settings. If you try to "open" a sealed management

pack, you will find that it does not look anything like an XML file. Sealing the management pack in this manner protects the intellectual property of the developers. But don't worry; you can make changes to the management group so that the management servers actually work the way you want them to.

Discovery

Discovery is the process of finding systems, services, and objects to monitor. Think of a discovery as a means of determining what will be monitored. Every management pack has discovery rules defined that, once imported, will seek out the object types the discovery method identifies. Back in Chapter 4, "Installing and Configuring Agents," we performed one form of discovery to find systems on which we installed agents. When we started the Discovery Wizard, we provided information on which systems we wanted to locate in order to install the agent. Once the entries were complete, systems were located based on the criteria we provided. This is not the only type of discovery that we have available to us, though. We can find many other objects that we may want to monitor.

Prepackaged management packs will already have discovery methods defined. The authors of the management packs know what type of objects they are monitoring in your network. Of course, they won't know exactly what objects exist in your network, but they do know what they would like to monitor. The discovery rules that they have defined will query and probe the network in search of the objects. Once discovered, the objects are automatically added to predefined groups. These groups are then used to organize the objects so that monitoring methods can be assigned to them more easily. For more information on groups, see the next section.

When a management pack author creates a discovery rule, the target for the rule needs to be chosen. The author needs to know what type of object they are going to monitor and how it exposes itself. This exposure is commonly referred to as an *attribute*. Operations Manager is not selective about the methods that can be used to discover object attributes. You could parse a systems registry to find a specific key, or you could perform an Lightweight Directory Access Protocol (LDAP) query against a directory service to locate a system. You can use any of the following methods when creating a discovery rule:

- ◆ LDAP query
- ◆ OLE DB
- ◆ Registry
- ◆ Scripts
- ◆ WMI query
- ◆ Custom managed code

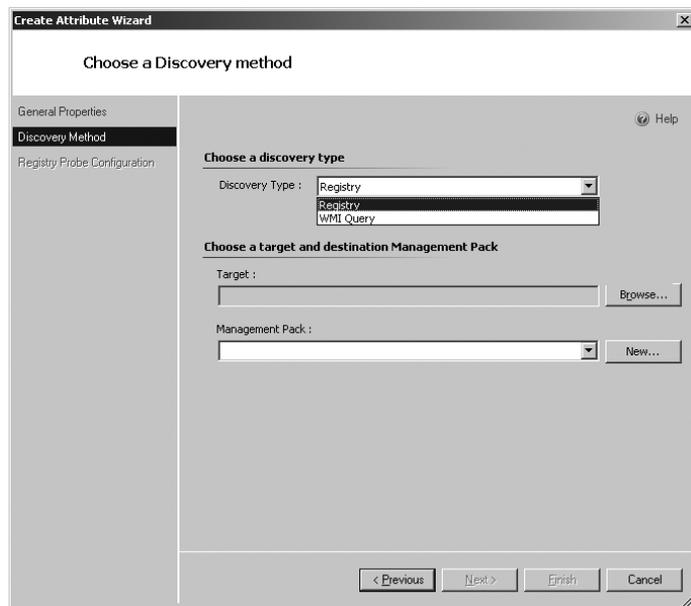
The authors of the management packs know the best way to discover the objects, but you may run into a scenario where you discover objects that you do not want to monitor. When a management pack is sealed, you will not be able to directly modify the discovery rule so that you can limit the discovered objects. You can always create an override that will limit the objects that are discovered. This way, you still have control over the monitored objects in your network. For more information, see the section "Overrides" later in this chapter.

If you feel the need to create a new group based on an attribute that is not defined within a management pack, you can create your own rule by using the Create Attribute Wizard. This wizard is available in the Operations Console; select the Authoring view and right-click Attributes in the Monitoring pane. Then select Create A New Attribute from the context menu.

As with most wizards, you need to supply a name for the item you are creating—in this case, the attribute. Most administrators simply supply the name here and move on. You have the option to supply a description as well. Take the time to further identify what the attribute is and why you are discovering it. At a later time, when you or another administrator is reviewing the attributes, the few moments you take to further identify the information could make this process easier.

After you click Next, you are presented with the Choose A Discovery Method page, which allows you to choose the type of discovery that you will use. As you can see in Figure 5.3, you have the options available to you that were discussed a few paragraphs back. After you select the type of discovery you are planning to use, choose the management pack that will host the discovery method for this attribute. If you are working with a management pack that you created, or if the attribute you are discovering is for a management pack that is unsealed, you will be able to select the management pack from the list and move on. However, if the management pack is sealed, you will not be able to modify it, which means you will not be able to save the new attribute discovery in the management pack.

FIGURE 5.3
Selecting the
discovery type
from the Choose
A Discovery
Method page



Sealed management packs introduce another challenge, then. You will have to identify an unsealed management pack where you will save the attribute discovery, or you could create a new management pack just for this occasion. Of course, just as with anything else, you should plan your attack. You don't want to have excess management packs lying around just because you wanted to create a new attribute discovery or an override. Think about and plan what you would like to do in your management group. If you decide to create a new management pack, instead of selecting a management pack from the pull-down menu, you can click the New button.

After you have decided which discovery type and management pack you are going to use, define how the discovery is going to work. This is where the wizard deviates depending on the discovery method you use; if you are searching the registry, you must specify the registry location and data that you are going to use. If you are performing a search with WMI, you must supply the WMI query and data that you are searching for.

Once you fill in your criteria and click Finish, you will have an attribute ready for discovery. That is only part of the equation, however. Sure, you know what attribute you are trying to find, and the management servers will be able to discover objects that match the attribute, but you do not have a way of organizing those objects so that you can start monitoring them. That is where groups come in.

Groups

In Operations Manager 2007, you can organize discovered objects into groups. If you have just imported a management pack that is used to monitor a storage area network (SAN), the discovery methods included in the management pack will probably be configured to monitor several objects, including physical drives and logical drives. After the discovery process is complete, the objects will be added to the defined groups according to their discovered attributes.

The criteria that are used to populate groups can be anything from a list of names to identify the objects, to information that is unique to the object that is written into the WMI. As discoveries find objects, groups are populated based on their configurations. As with discoveries, developers of the management packs have to work in generalities because they do not know the inner workings of every company. They create the groups that they believe will be useful for a majority of the companies that will use their management pack. For more specific groups—those that are tailored to your own environment—you will have to create your own groups to use.

Once configured, groups can be quite useful. Since they identify unique object types, you can specify who will be able to monitor and manipulate the object in the group. You can also control the views in the administrator console by using groups. As such, you can take control of your monitoring solution and assign access at a granular level. Take, for instance, a scenario where you want to allow a user to be able to view and monitor specific databases in your SQL server infrastructure. You can create a group that identifies the databases by name, and then assign permissions for the user to be able to view and monitor them. You can then create a specialized view that only displays information pertaining to the objects in the group. This way, you can isolate the data that you want the operator to view, keeping them out of data they do not need to view.

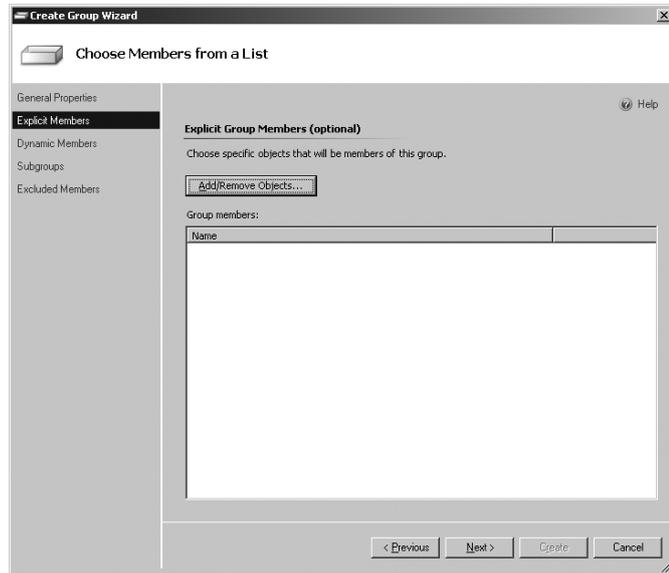
CREATING GROUPS

If the groups included in a management pack do not meet your needs, you will have to create your own. The process itself is not difficult; a wizard is available to step you through the process. You must decide what the group will be used for. As mentioned, a group can be created to organize systems or objects. Once you decide what you want to organize and monitor, start the Create Group Wizard: in the Operations Console, select the Authoring view and right-click Groups in the Authoring pane. Then select Create A New Group.

First, enter the group name and a description that will help you remember what the group is used for. After clicking Next, you are presented with the Management Pack page. Select the unsigned management pack that you are going to use. Notice that the Default Management Pack is selected, but you can choose any existing management pack or click the New button to create one.

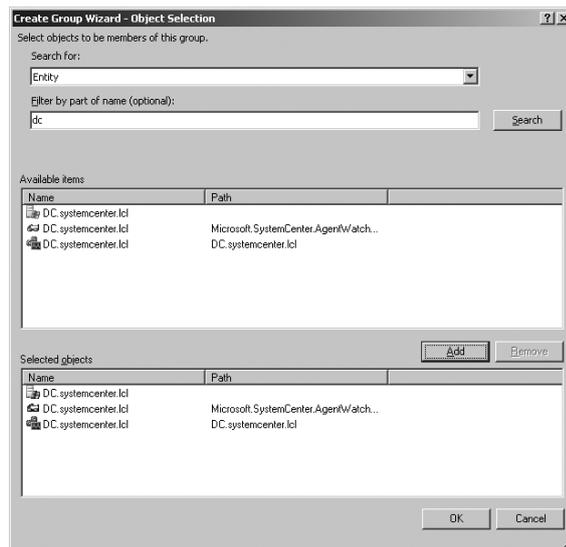
When the Choose Members From A List page appears, you need to make a choice. The page that you are looking at, as shown in Figure 5.4, is used if you want to explicitly add objects to the group. One of the most popular uses for this page is to add computers by name. You could just as easily add databases by name. If you have decided to add objects to the group in this manner, click the Add/Remove Objects button to specify the objects.

FIGURE 5.4
Explicitly adding
objects



The Object Selection dialog box that appears presents you with search criteria that you will use to find the objects. The Search For list contains object types that you can use to identify the objects. If you are planning on adding computers by their name, select the Windows Computer object type, as shown in Figure 5.5. At this point, you could go on and view a list of all the computers in your domain, but you may have too many to go through. If you have a naming scheme that you use for your systems, enter a portion of the computer name that is used for the group that you would like to add. Enter the text in the Filter By Part Of The Name text box and then click Search.

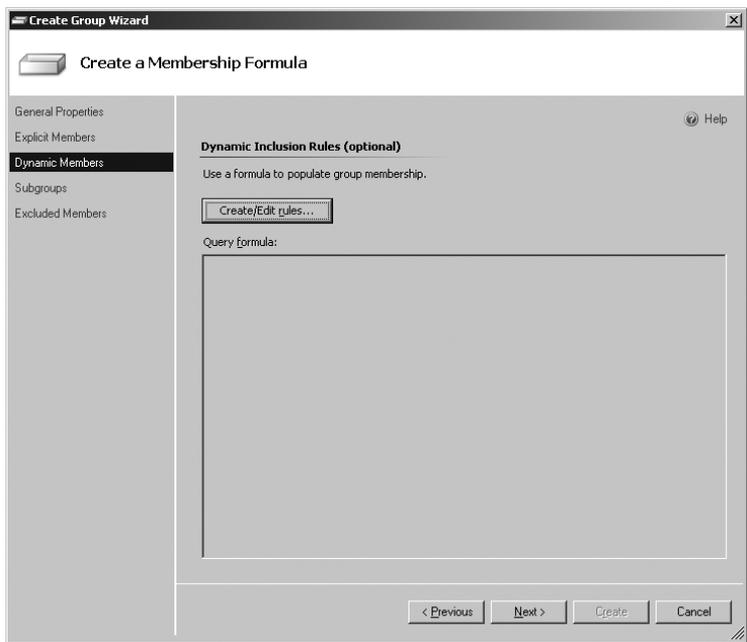
FIGURE 5.5
Adding objects
explicitly



After the query has run, you are presented with a results screen that allows you to select the appropriate objects. Select those that meet your needs, click the Add button, and then click Next. The only problem with this method is that the group membership is static in nature. If you want to add or remove objects from the group, you have to manually add or remove them. Allowing the system to dynamically control the group membership is usually the more efficient method of controlling the group. To do so, you have to create the group based on attributes.

If you plan to create a group based on an attribute, you have to move past the Choose Members From A List page and work with the Create A Membership Formula page, shown in Figure 5.6. When you click the Create/Edit Rules button, you are presented with a query builder. The query builder is handy if you are not proficient at writing your own queries. The attributes that you can use are included in the list of objects type on this page. When you created an attribute discovery method in the Discovery area, you chose an object type that the attribute was related to. Now you are targeting the attribute at a specific object type. You can select an attribute from the list, or if you have created your own attribute, click Add and select your attribute from the list of attributes.

FIGURE 5.6
Create A Membership
Formula page



The operators that are available from the query builder let you fine-tune your query. You can use the standard Is Equal and Does Not Equal operators as well as some less common operators such as Does Not Match Wildcard. Using the Insert button, you can add additional expressions, as well as create AND and OR groups to control the order of precedence for query evaluation. When you create the groups, you are creating a list of expressions that are evaluated together. For instance, if you enter three expressions in an AND group, all of the expressions have to evaluate to true for it to be processed.

As Figure 5.7 shows, as you are building your query you will see just a few of the classes that are available for you to choose to populate the group. When you select a class and click Add, an AND group is created in the query window. When you go back to the classes and add another,

an OR group is created so that the two AND groups are evaluated. When either of the AND groups evaluate to true, the expression becomes true. In Figure 5.8, we have selected the Virtual Servers class and chosen the attributes for Virtual Machine and Organizational Unit. The settings we have specified allow us to create a group that includes all the DNS servers that are running as virtual servers.

FIGURE 5.7
Choosing classes for
the group query

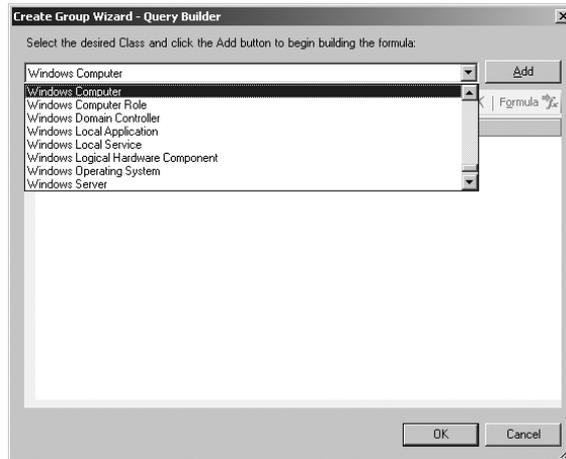
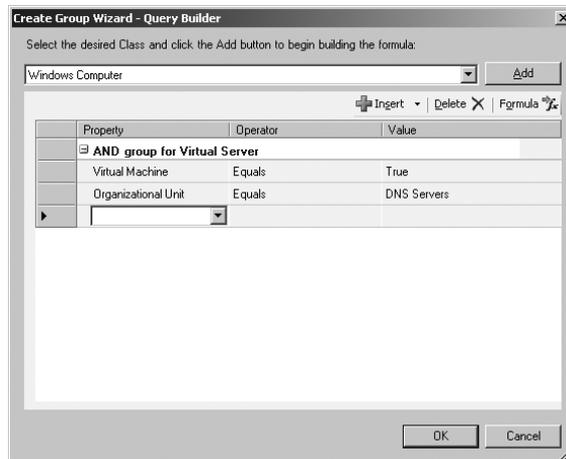
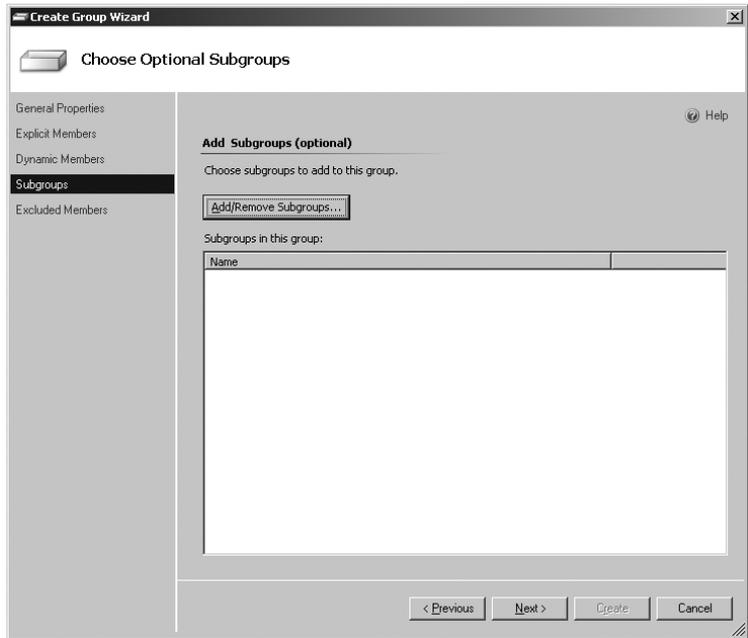


FIGURE 5.8
Creating the query



Once your query is completed, click the Formula button to view the completed query statement. Click Next and you will be presented with the Add Subgroups page. Here you can create groups that have their own distinct membership but are related to the parent group. For instance, you could have a group of operators who are responsible for monitoring all the Exchange public folders, and then create subgroups and assign other operators the ability to monitor a subset of those public folders. The only caveat here is that the group that you will configure as a subgroup must already exist. You can search for the group names, as shown in Figure 5.9.

FIGURE 5.9
Adding subgroups



The final page of the wizard allows you to exclude specific objects from the group membership. This is helpful when you are using dynamic group membership and there are objects that you want to explicitly exclude from monitoring. The method for adding excluded objects to the list is identical to the one for adding subgroups. Once you have them added, click Finish to complete the wizard and your group will be created and populated.

Monitors

Monitors are used to determine the health of objects that you plan on monitoring. Several monitor types are available, each providing you with a means of extracting information about an object and determining how that object is functioning. Depending on what type of monitor you are working with, you could determine the health of an object, which can alter the health view of other monitors. You can also create alerts when a monitor determines that an object is in a warning or error state.

ROLLUP MONITORS

The fundamental change in monitoring is what sets Operations Manager 2007 apart from its counterparts. As we have mentioned several times throughout this book, moving from server monitoring to service monitoring allows administrators to determine how a service and its supporting services are functioning. Two of the most important tools in Operations Manager's arsenal are the aggregate rollup monitor and the dependency rollup monitor. These monitors work with all the monitor types to create a monitoring *chain*.

Aggregate Rollup Monitors

Supporting services act as dependencies to related services in the chain. When a dependency service starts having problems, the monitor that is watching the service is changed to the appropriate

state. The change does not stop there, however. All the monitors further up the chain also take on the state of the affected service if so configured. In this way, you can see that there is a problem with the service due to issues with another service. Of course, you do have the option of not aggregating the state of the monitor, in which case the chain will be affected only to the degree that you specify.



Real World Scenario

FORMATTING OUTPUT

It is interesting to see how data can be viewed within a management group. As data is collected from managed systems, operators can view the data and determine what needs to be done. In the text we mentioned that you can create a rollup policy based upon the state that operators need to see. Just as in the text, there were instances when the data that was collected was viewed by two different operators, each one needing to monitor a different part of the organization. Each operator was concerned with the services that they were responsible for, and needed to make sure that the supporting services were working correctly.

In this organization, there were several Exchange servers for the organization's email and calendaring needs. To route the email outside of the organization, there were SMTP servers within the perimeter network. The administrators of the SMTP servers were interested in making sure that all of the SMTP servers were running. If a single server were to fail, they wanted to know. At the same time, the Exchange administrators were not concerned about an individual SMTP server failing, but they did want to know if there weren't any SMTP servers available.

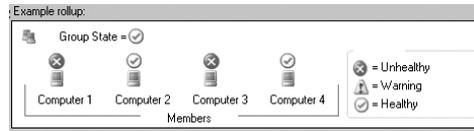
Two aggregate rollup monitors were created. One for the SMTP administrators, which had a worst-case rollup policy configured so that they could be notified whenever any of the SMTP servers were down. The second one was configured as a best-case rollup policy so that the Exchange administrators knew that the SMTP server were available. This way both groups were happy because the Exchange administrators didn't get worried when a single SMTP server failed, and the SMTP administrators were notified as soon as something happened to one of their servers, and they could start fixing the problem as soon as it was reported.

In the past, when a server was being monitored, only services and applications residing on that server would be monitored. If the server did not have a problem but a dependency service on another server was having an issue, the person responsible for the unaffected server would not see anything wrong. Rare is the case anymore when a server can work as a stand-alone unit. Email servers rely on networking services and underlying infrastructure to run correctly. Database servers use Web-based front ends for users to access data. Domain controllers rely on DNS servers. For instance, if you are running Exchange, you must make sure that DNS is available. You also need to have domain controllers configured as Global Catalogs. In previous versions of Operations Manager, you could have a DNS server fail but the technicians responsible for managing the Exchange servers might not have realized there was an issue.

Aggregate rollup monitors can be configured so that they show a worst- or best-case state of the service being monitored. When you configure a monitor as best-case, you are essentially saying that you want the monitor to appear in the state of the service or system that is functioning the best, which Microsoft shows in their graphic seen in the Health Rollup Policy configuration page (Figure 5.10). You would typically use a best-case state when you want to make sure you have at

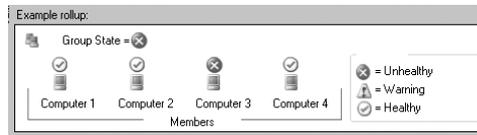
least one service still providing functionality. Such is the case when you have three SMTP servers in a network load-balanced scenario. You could configure a monitor that shows an operator that the service is still available even though one of the SMTP servers may have gone offline.

FIGURE 5.10
Best-case
rollup policy



As you can imagine, the worst-case state is the polar opposite of the best-case state. When you configure a monitor to show worst-case state, whenever a single child monitor goes into a warning or error state, that state is propagated up the hierarchy, which Microsoft shows in their graphic seen in the Health Rollup Policy configuration page (see Figure 5.11). In the case of the SMTP service that we mentioned earlier, if you are using the worst-case state, the operator would see a problem as soon as one of the SMTP servers failed or had a problem. For some monitors, you can decide whether you are a glass-half-full or glass-half-empty kind of person. Although most people want to see if there is any problem at all, you may want to devise different viewing strategies for different groups in your organization. Some may need to see that there is a problem as soon as a child monitor detects one, while others don't have to see that there is a problem until the service is completely down.

FIGURE 5.11
Worst-case
rollup policy

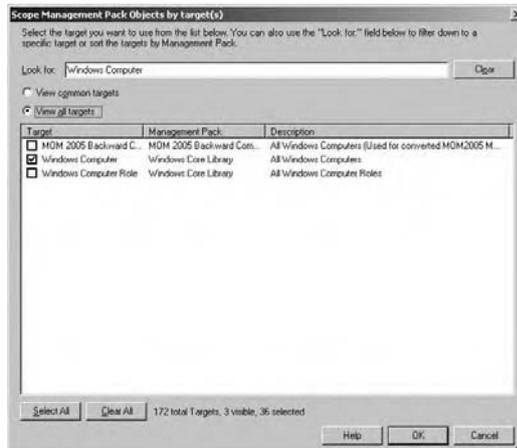


If you are working under an SLA that specifies the availability of a service, you may find that showing a best-case state is the way to go. Most high-availability solutions provide for some type of redundancy or failover. Consider a network load-balanced web service that provides access to an e-commerce site. The web servers that provide access to the website will all provide the same functionality to the customers who are visiting the site. If one of the web servers fails, the customer requests would then simply be redirected to the other web servers in the network load-balanced cluster. Access to the website would not be affected.

A positive state for the service would appear in the operator's console as well as in all the reports that could be generated from the database. As long as one web server was still providing access to the website, the SLA would be intact. Of course, there are other criteria usually included in an SLA, but you get the picture.

To create an aggregate rollup monitor, begin by opening the operator's console and choosing the Authoring view. In this view, navigate through the Authoring pane until you reach the monitors. To do so, select Authoring > Management Pack Objects > Monitors. Since several object types can be monitored, you will have to generalize the view. At the top of the operator's console, click the Scope button. When the Scope Management Pack Objects By Target(s) dialog box appears, as shown in Figure 5.12, type **Windows Computer** in the Look For text box. Select the Windows Computer check box and then click OK.

FIGURE 5.12
Scoping by Windows
Computer



Now the view should show only the Windows Computer object type. Once you expand the Windows Computer object type, follow these steps:

1. Expand Entity Health in order to see the Performance Monitor level.
2. Right-click the Performance node, point to Create A Monitor, select Aggregate Rollup Monitor, and then select an unsigned management pack or create a new one to hold the monitor. This will start another wizard: the Create A New Aggregate Rollup Monitor Wizard.
3. On the General Properties page, shown in Figure 5.13, name and describe the monitor.
4. The final option on this page is the Parent Monitor list. Define the monitor that will be affected when this monitor changes states, and click Next.
5. The next page, Health Rollup Policy, allows you to choose the way in which the parent is affected when the child monitors change state. The two options are Worst State Of Any Member and Best State Of Any Member. Choose whether you want to see the monitor showing problems or whether the service is running on at least one system; then click Next.
6. The Configure Alerts page, shown in Figure 5.14, allows you to configure the alerts that are raised by this monitor. Make sure that the Generate Alerts For This Monitor option is selected.
7. In the Generate An Alert When pull-down, select the level that you want to present an alert: The Monitor Is In A Critical Health State or The Monitor Is In A Critical Or Warning Health State.
8. The Automatically Resolve The Alert When The Monitor Returns To A Healthy State check box should be deselected only if you want to have an operator be responsible for resolving the alert.
9. Select the alert Priority: Low, Medium, or High.
10. Select a Severity: Information, Warning, or Critical.

FIGURE 5.13
Aggregate rollup
General Properties

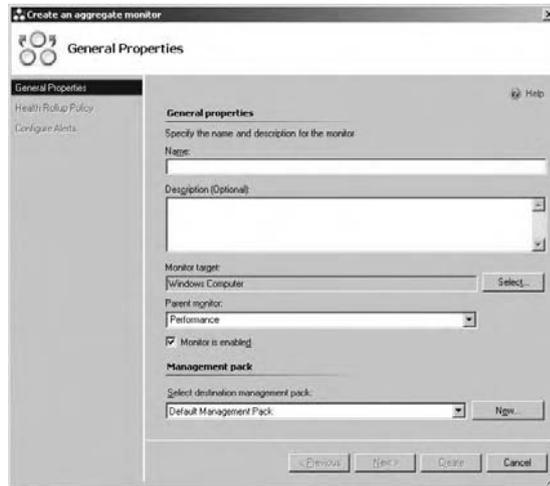
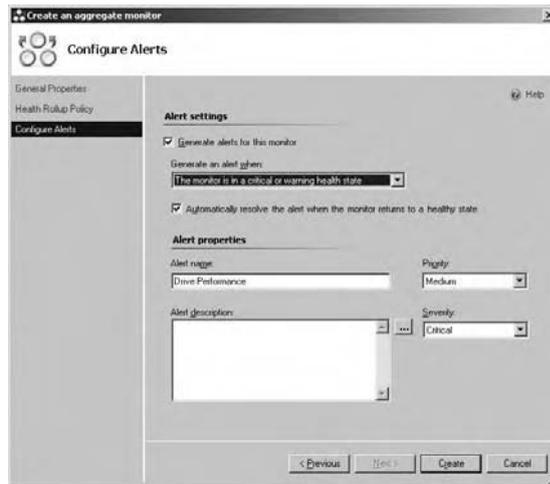


FIGURE 5.14
Configuring alerts for
the aggregate rollup
monitors



Dependency Rollup Monitor

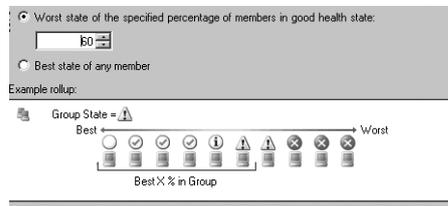
At first glance, you might think the dependency rollup monitor is the same thing as an aggregate rollup monitor, but there is a difference between the two. Whereas the aggregate rollup monitor relies on child monitors for services in a solution, a dependency rollup monitor can change the health state for a dependency that would otherwise not be included in an aggregate rollup monitor chain. Exchange is a good example of this. In an Exchange solution, you have several services that need to run in order for the Exchange solution to appear healthy. You have network services that need to run in order for Exchange to function. For instance, DNS has to be functioning for the Exchange servers to locate their Global Catalog servers as well as each other. SMTP servers are responsible for sending email messages between each of the Exchange servers as well as foreign mail systems on the Internet.

If one of the SMTP servers running on an Exchange server fails, the aggregate rollup monitors propagate the health state up the chain. However, to show the Windows server in the same state, you have to use a dependency monitor. Without the dependency monitor, you would see each of

the servers in a healthy state and only the Exchange services having a problem. To identify which server is having the problem, the dependency rollup monitor should be used.

As with the aggregate rollup monitor, you can change the health state of the dependency rollup monitor so that it reflect the health state of the dependency according to your needs. You have the same two options that were available from the aggregate rollup monitor: Worst State of Any Monitor and Best State of Any Monitor. In addition, a third option is available: Worst State Of The Specified Percentage Of Members In Good Health State (see Figure 5.15). This third state allows you to fine-tune the monitor's health; for example, you could have six SMTP servers and the system could be shown as healthy as long as three are still functioning properly. In this case, you could specify that 50 percent of the systems have to show in a healthy state for the monitor to show as healthy. As soon as the fourth system in a group of 10 has a problem, the monitor would show in a warning or error state.

FIGURE 5.15
Worst state of
percentage
of members



In addition to the health state rollup, the dependency rollup monitor allows you to change the state of the monitor whenever monitoring becomes unavailable or if the system is placed in maintenance mode so that you can perform scheduled maintenance. For either of these settings, you can configure the health to Error, Warning, or Ignore. Of course, Ignore leaves the state as healthy.

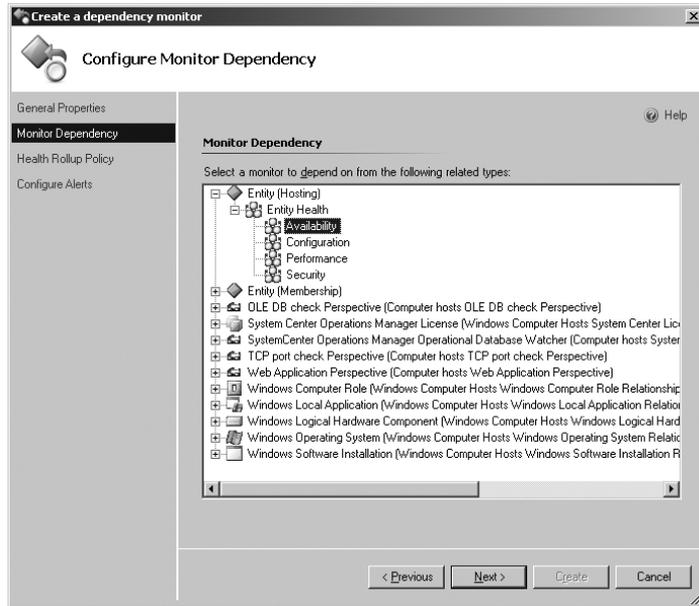
To create a dependency rollup monitor, open the operator's console and choose the Authoring view. In this view, navigate through the Authoring pane until you reach the monitors. To do so, select Authoring > Management Pack Objects > Monitors. Since several object types can be monitored, you must generalize the view. At the top of the operator's console, click the Scope button. When the Scope Management Pack Objects By Target(s) dialog box appears, as shown in Figure 5.12, type **Windows Computer** in the Look For text box. Select the Windows Computer check box and then click OK.

Now the view should show only the Windows Computer object type. Once you expand the Windows Computer object type, follow these steps:

1. Expand Entity Health in order to see the Performance monitor level.
2. Right-click the Performance node, point to Create A Monitor, select Dependency Rollup Monitor, and then select an unsigned management pack or create a new one to hold the monitor. This will start another wizard: the Create A New Dependency Rollup Monitor Wizard.
3. On the General Properties page, name and describe the monitor.
4. The final option on this page is the Parent Monitor list. Define the monitor that will be affected when this monitor changes states and click Next.
5. On the Configure Monitor Dependency page, shown in Figure 5.16, expand Windows Logical Hardware Component > Entity Health and click Availability. Click Next.
6. On the next page, Health Rollup Policy, choose the way in which the parent is affected when the child monitors change state. The three options are Worst State Of Any Member, Best State Of Any Member, and Worst State Of The Specified Percentage Of Members In Good Health State. Choose whether you want to see the monitor showing problems or whether the service is running on at least one system.

7. In the Monitoring Unavailable And Maintenance Mode area, select from the Monitoring Unavailable list how you want the health of the monitor to appear when the monitoring is not available.
8. In the Monitoring Unavailable And Maintenance Mode area, select from the Maintenance Mode list how you want the health of the monitor to appear when the system is in maintenance mode and then click Next.
9. The Configure Alerts page allows you to configure the alerts that are raised by this monitor. Make sure that the Generate Alerts For This Monitor option is selected.
10. In the Generate An Alert When pull-down menu, select the level that you want to present an alert: The Monitor Is In A Critical Health State or The Monitor Is In A Critical Or Warning Health State.
11. The Automatically Resolve Alert When The Monitor Returns To A Healthy State check box should be deselected only if you want to have an operator be responsible for resolving the alert.

FIGURE 5.16
Choosing the monitor
dependency



12. Select the alert Priority: Low, Medium, or High.
13. Select a Severity: Information, Warning, or Critical.

UNIT MONITORS

Unit monitors are meant to monitor discrete events or services so that you can detect when you have an issue arising. Unlike performance-based alerts, these unit monitors probe the system by gathering data from event logs or monitor services to determine whether they are running. Even though unit monitors are sometimes deemed as reactive, you should still consider their use since they will give you a well-rounded view of your monitoring solution. There are also applications and

services that do not expose themselves to the agent; therefore, you cannot gather health statistics on them. If that is the case, you will have to create a unit monitor to keep track of events they report.

As you are creating your unit monitors, you can test them using the `eventcreate.exe` utility that is included with Windows Server 2003 or Windows XP. After you create a monitor, you can raise an event with the event ID that you are watching for. If you are monitoring the System event log for event ID 1012, you could create an event by using the command:

```
eventcreate /L System /ID1012 /T Error
```

The syntax for the `eventcreate` command allows you to specify a computer that you want to target; assign credentials for the command to run under; and specify which event log you are targeting, the type of event, and a description. So if you were looking for a specific event that had a specific text string, you could enter it.

NOTE If you want more information about `eventcreate`, check out the Microsoft website at <http://www.microsoft.com/resources/documentation/windows/xp/all/proddocs/en-us/eventcreate.mspx?mfr=true>.

Simple Windows Event Unit Monitor

A simple Windows event unit monitor keeps an eye on event logs, watching for an event that matches a query you have built. If you decide that you want to be notified whenever event ID 9036 appears in the system log, you can create it here. There are three types of simple Windows event unit monitors:

Manual Reset Manual reset unit monitors will not automatically reset themselves. It is up to an operator to reset the monitor, thus restoring the health.

Timer Reset Timer reset unit monitors are configured to reset their health status after a specified amount of time has elapsed.

Windows Event Reset Windows event reset relies on two events to complete. When the initial event occurs, the status of the monitor is set accordingly and will not be reset until the second defined event occurs.

To create a simple Windows event unit monitor, open the operator's console and select the Authoring view. In this view, navigate through the Authoring pane until you reach the monitors. To do so, expand Authoring > Management Pack Objects > Monitors. Since several object types can be monitored, you must generalize the view. At the top of the operator's console, click the Scope button. When the Scope Management Pack Objects By Target(s) dialog box appears, as seen in Figure 5.12, type **Windows Computer** in the Look For text box. Select the Windows Computer check box and then click OK.

Now the view should show only the Windows Computer object type. Once you expand the Windows Computer object type, follow these steps:

1. Expand Entity Health in order to see the Availability monitor level.
2. Right-click the Availability node, point to Create A Monitor, and select Unit Monitor. This starts another wizard: the Create Monitor Wizard.
3. On the Select A Monitor Type page, expand Windows Events > Simple Event Detection, as shown in Figure 5.17, and select Windows Event Reset. Select an unsigned management pack that will host the monitor; then click Next.

4. On the General Properties page, name and describe the monitor. Then choose the Monitor Target by clicking Select and choosing the target from the list, as shown in Figure 5.18.
5. The final option on this page is the Parent Monitor list. Define the monitor that will be affected when this monitor changes states and click Next.
6. On the Event Log Name page, start defining the unhealthy condition by clicking the ellipsis button.
7. On the Select Event Log page, click the ellipsis under Computer and select the name of a computer that hosts the appropriate event log.
8. Choose the appropriate event log from the list and click OK.
9. Click Next.
10. The Build Event Expression page, shown in Figure 5.19, appears, allowing you to define the event that you want to watch for that indicates an unhealthy condition.
11. On the Event Log Name page, start defining the healthy condition that will reset the monitor by clicking the ellipsis button.
12. On the Select Event Log page, click the ellipsis under Computer and select the name of a computer that hosts the appropriate event log.
13. Choose the appropriate event log from the list and click OK.
14. Click Next.
15. The Build Event Expression page appears, allowing you to define the event that you want to watch for that indicates a healthy condition.
16. Click Next.
17. The Configure Health page, shown in Figure 5.20, allows you to configure the event states for the events that you just identified in the previous steps. In the First Event Raised option, select the name in the Operational State column and enter a new name for the event. In the Health State column, choose Critical or Warning.
18. In the Second Event Raised row, select the name in the Operations States column and enter a new name for the event. In the Health State column, choose Healthy.
19. Click Next.
20. The Configure Alerts page allows you to configure the alerts that are raised by this monitor. Make sure that the Generate Alerts For This Monitor option is selected.
21. In the Generate An Alert When pull-down, select the level that you want to present an alert: The Monitor Is In A Critical Health State or The Monitor Is In A Critical Or Warning Health State.
22. The Automatically Resolve Alert When The Monitor Returns To A Healthy State check box should be deselected only if you want to have an operator be responsible for resolving the alert.
23. Select the alert Priority: Low, Medium, or High.
24. Select a Severity: Information, Warning, or Critical.

FIGURE 5.17
Selecting the Win-
dows Event Reset

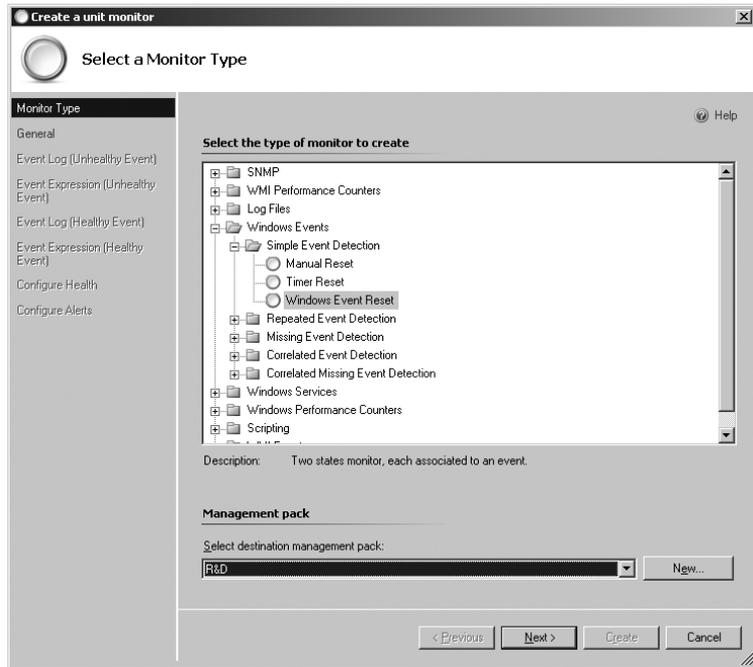


FIGURE 5.18
Choosing the monitor
target

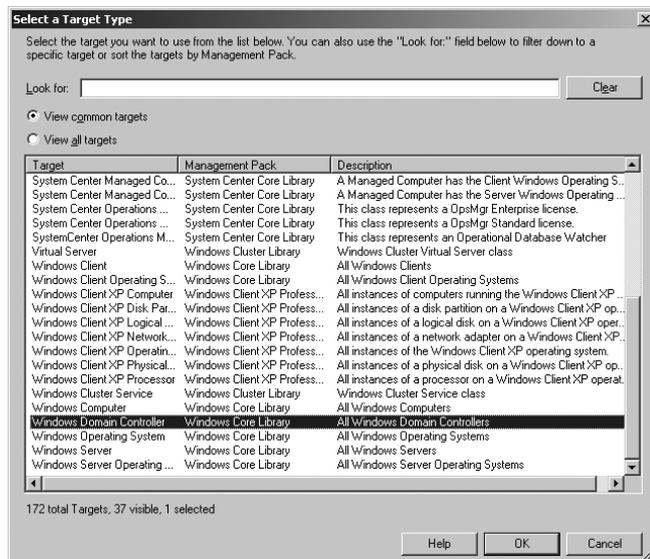


FIGURE 5.19
Building the expression

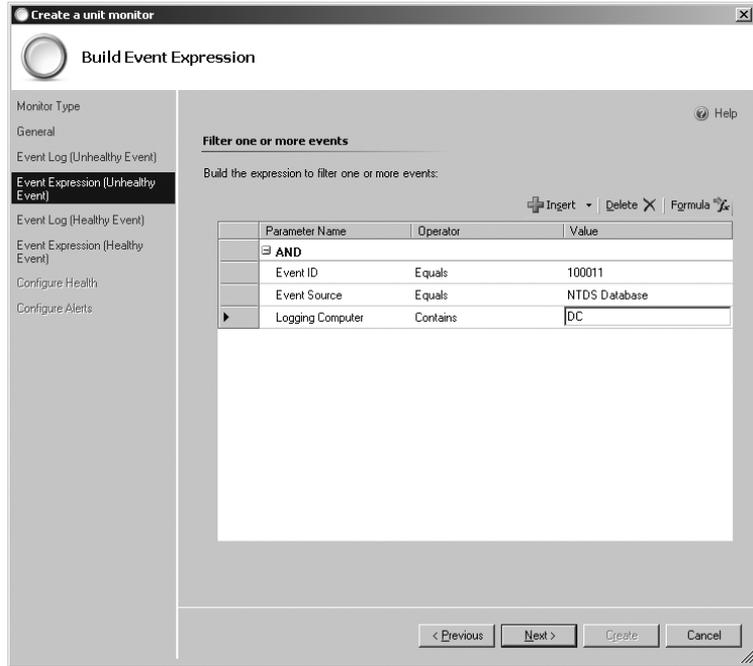
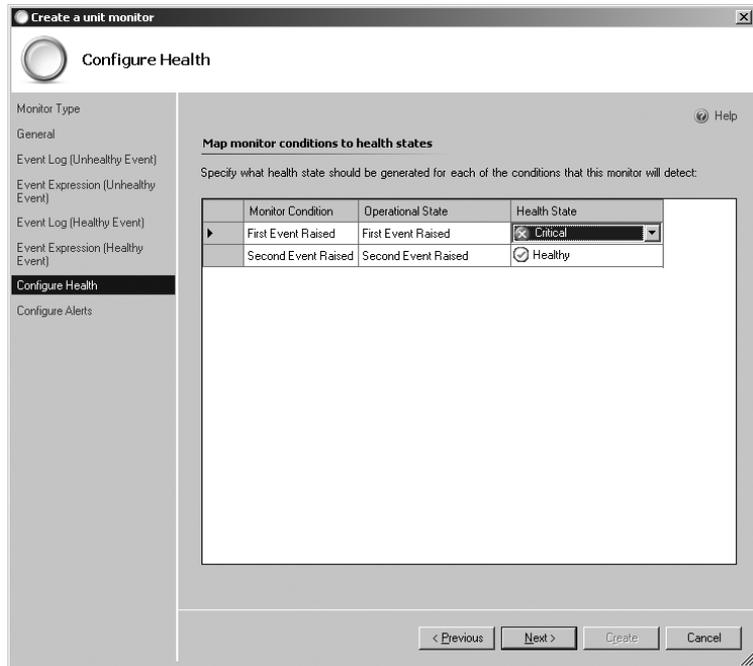


FIGURE 5.20
Configuring the health states



Correlated Windows Event Unit Monitor

Correlated Windows event unit monitors are an interesting monitor type. Whereas simple Windows event unit monitors wait for an event to occur, the correlated Windows event unit monitor will not change the health state unless a second event also occurs. This is because there are times when an event is not a critical event unless a second related event also appears. When you create this type of unit monitor, you need to actually configure three events: one to make the monitor take notice, a second to change the state, and a third to reset the state back to healthy.

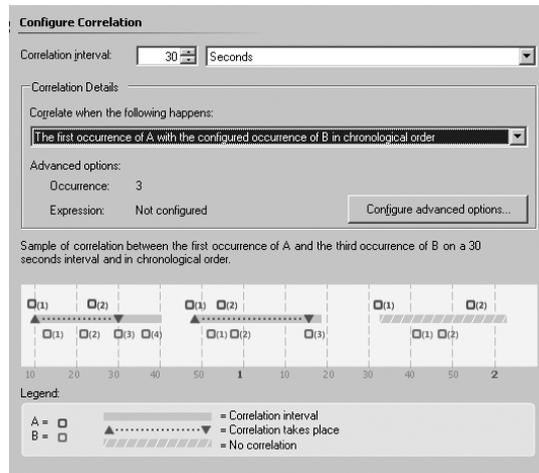
The first event is configured as a simple monitor that initiates a timer. The timer settings determine how long the monitor will wait for the second event. The second event has to occur during the timer settings. Make sure that you know the time it takes for the second event to occur. If you set the timer too short and the second event occurs outside the parameters, you will not see a state change. If the timer is set too long, you could possibly encounter phantom state changes that don't reflect a true problem.

Just like the simple Windows event unit monitor that is configured as a Windows event reset type, the third event in the correlated monitor watches for another event that reports the system is running nominally. Here are the correlations you can define:

The First Occurrence Of A With The Configured Occurrence Of B In Chronological Order

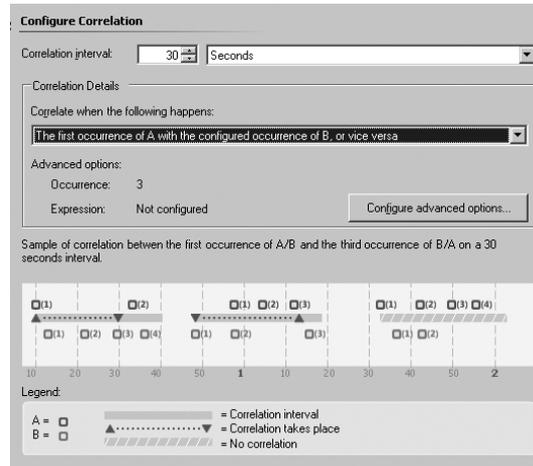
This monitors for the event defined in Log A to occur, and then will monitor for the specified number of events defined in Log B to occur. When this number is reached in the time frame specified, the monitor health will change to the defined state. Figure 5.21 shows the monitor configured to watch for three events from Log B to appear within 30 seconds after the event from Log A appears.

FIGURE 5.21
Monitoring for the occurrence of A and then three B events in a specified time frame



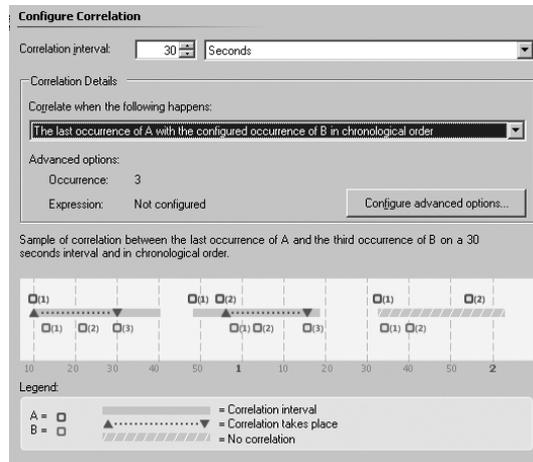
The First Occurrence Of A With The Configured Occurrence Of B, Or Vice Versa This monitors for the event defined in either log to occur, and then will monitor for the specified number of events defined from the other log to occur. When this number is reached in the time frame specified, the monitor health will change to the defined state. Figure 5.22 shows the monitor configured to watch for three events from Log A or Log B to appear within 30 seconds after the event from the other log file appears.

FIGURE 5.22
Monitoring for the occurrence of three A or B events in a specified time frame



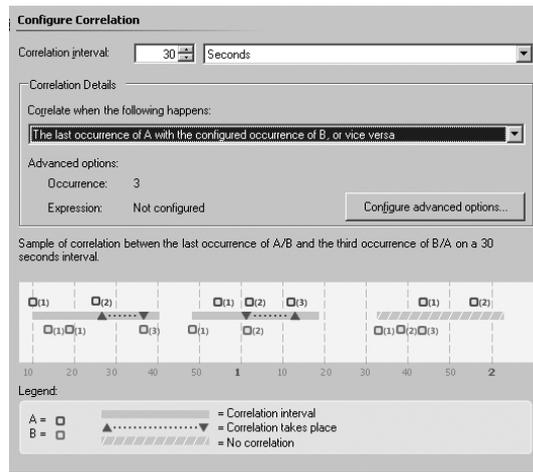
The Last Occurrence Of A With The Configured Occurrence Of B In Chronological Order
This monitors for the occurrence of the specified number of events from Log B since the last occurrence of the defined event from Log A, in the specified amount of time. Figure 5.23 shows the monitor configured to watch for the three defined events from Log B to appear since the last appearance of the defined event from Log A.

FIGURE 5.23
Monitoring for three occurrences from Log B since the last occurrence from Log A



The Last Occurrence Of A With The Configured Occurrence Of B, Or Vice Versa
This monitors for the occurrence of the specified number of events from Log A or Log B in the last occurrence of the defined event from the other log file, in the specified amount of time. Figure 5.24 shows the monitor configured to watch for the three defined events from Log A or Log B to appear since the last appearance of the defined event from the other log.

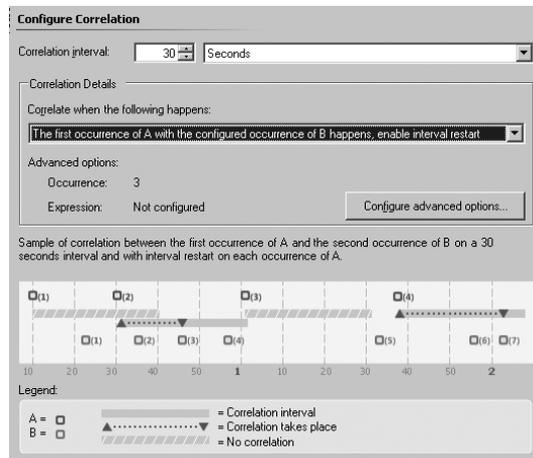
FIGURE 5.24
Monitoring for three occurrences from Log A or Log B since the last occurrence from the other log



The First Occurrence Of A With The Configured Occurrence Of B Happens, Enable Interval Restart This monitors for the occurrence of the specified number of events from Log B to appear after the event from Log A has appeared. On each successive appearance of the event from Log A, the timer is reset. Figure 5.25 shows the correlations between the two event logs when you are monitoring for two events from Log B to appear within 30 seconds after the event from Log A appears.

To create a correlated Windows event unit monitor, open the operator’s console and select the Authoring view. In this view, navigate through the Authoring pane until you reach the monitors. To do so, expand Authoring > Management Pack Objects > Monitors. Since several object types can be monitored, you must generalize the view. At the top of the operator’s console, click the Scope button. When the Scope Management Pack Objects By Target(s) dialog box appears, as shown in Figure 5.12, type **Windows Computer** in the Look For text box. Select the Windows Computer check box and then click OK.

FIGURE 5.25
Monitoring for two events to occur from Log B and then resetting the timer whenever an event from Log A appears

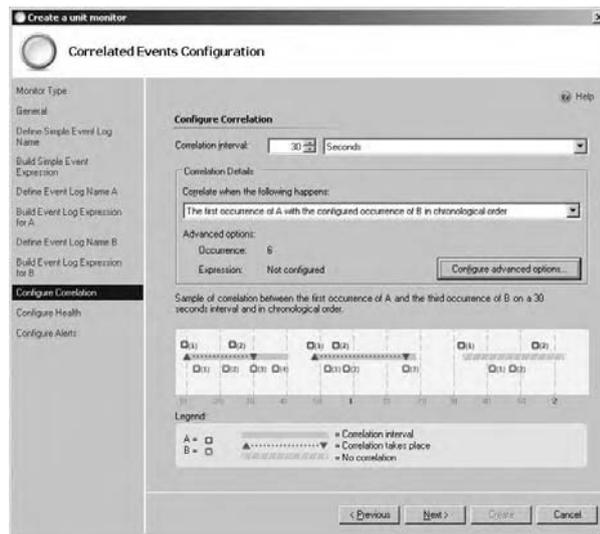


Now the view should show only the Windows Computer object type. Once you expand the Windows Computer object type, follow these steps:

1. Expand Entity Health in order to see the Availability monitor level.
2. Right-click the Availability node, point to Create A Monitor, and select Unit Monitor. This will start another wizard: the Create Monitor Wizard.
3. On the Select A Monitor Type page, expand Windows Events > Correlated Event Detection, select Windows Event Reset, select an unsigned management pack that will host the monitor, and click Next.
4. On the General Properties page, name and describe the monitor.
5. The final option on this page is the Parent Monitor list. Define the monitor that will be affected when this monitor changes states and click Next.
6. On the Event Log Name page, start defining the first event you are watching for by clicking the ellipsis button.
7. On the Select Event Log page, click the ellipsis under Computer and select the name of a computer that hosts the appropriate event log.
8. Choose the appropriate event log from the list and click OK.
9. Click Next.
10. The Build Event Expression page appears, allowing you to define the event that you want to watch for that indicates the first event. Click Next.
11. On the Event Log Name page, start defining the correlated event that will set the monitor to an unhealthy condition by clicking the ellipsis button.
12. On the Select Event Log page, click the ellipsis under Computer and select the name of a computer that hosts the appropriate event log.
13. Choose the appropriate event log from the list and click OK.
14. Click Next.
15. The Build Event Expression page appears, allowing you to define the event that you want to watch for that indicates an unhealthy condition.
16. Click Next.
17. On the Event Log Name page, start defining the correlated event that will reset the monitor to a healthy condition by clicking the ellipsis button.
18. On the Select Event Log page, click the ellipsis under Computer and select the name of a computer that hosts the appropriate event log.
19. Choose the appropriate event log from the list and click OK.
20. Click Next.
21. The Build Event Expression page appears, allowing you to define the event that you want to watch for that indicates a healthy condition.
22. Click Next.

23. The Correlated Events Configuration page, shown in Figure 5.26, allows you to identify the relationship between the correlated events and the time interval between the time the first event is raised and the correlated event occurs. Enter the time frame in the Correlation Interval box. This value can be anywhere from 1 second to 2,174,483,647 seconds (68 years!).
24. In the Correlation Details, select the entry that specifies the relationship between the two correlated events in the Correlate When The Following Happens list and click Next.
25. The Configure Health page allows you to configure the event states for the events that you just identified in the previous steps. In the Correlated Event Raised option, select the name in the Operations States column and enter a new name for the event. In the Health State column, choose Critical or Warning.
26. In the Event Raised row, select the name in the Operations States column and enter a new name for the event. In the Health State column, choose Healthy.
27. Click Next.
28. The Configure Alerts page allows you to configure the alerts that are raised by this monitor. Make sure that the Generate Alerts For This Monitor option is selected.
29. In the Generate An Alert When pull-down menu, select the level that you want to present an alert: The Monitor Is In A Critical Health State or The Monitor Is In A Critical Or Warning Health State.
30. The Automatically Resolve Alert When The Monitor Returns To A Healthy State check box should be deselected only if you want to have an operator be responsible for resolving the alert.
31. Select the alert Priority: Low, Medium, or High.
32. Select a Severity: Information, Warning, or Critical.

FIGURE 5.26
Configuring the
correlated events



Windows Services Unit Monitor

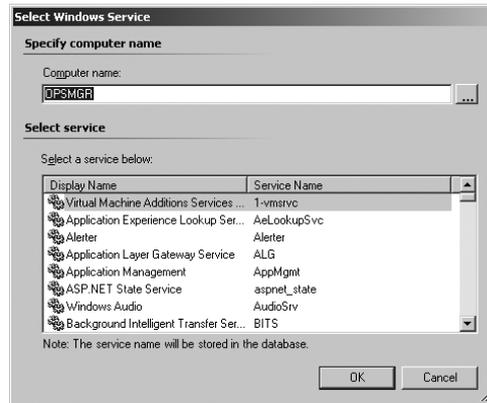
Windows services unit monitors constantly watch a service running on a computer and change the state of the object when the service is found to be not running. Although many of the services running on Windows systems can report some health information to monitors, it can be helpful to watch for a failed service in this manner.

To create a Windows services unit monitor, open the operator's console and select the Authoring view. In this view, navigate through the Authoring pane until you reach the monitors. To do so, expand Authoring > Management Pack Objects > Monitors. Since several object types can be monitored, you must generalize the view. At the top of the operator's console, click the Scope button. When the Scope Management Pack Objects By Target(s) dialog box appears, as shown in Figure 5.12, type **Windows Computer** in the Look For text box. Select the Windows Computer check box and then click OK.

Now the view should show only the Windows Computer object type. Once you expand the Windows Computer object type, follow these steps:

1. Expand Entity Health in order to see the Availability monitor level.
2. Right-click the Availability node, point to Create A Monitor, and select Unit Monitor. This will start another wizard: the Create Monitor Wizard.
3. On the Select A Monitor Type page, expand Windows Events > Simple Event Detection, select Windows Event Reset, and select an unsigned management pack that will host the monitor. You can also choose to create a new management pack at this point by clicking the New button. Click Next.
4. On the General Properties page, name and describe the monitor.
5. The final option on this page is the Parent Monitor list. It is here that you need to define the monitor that will be affected when this monitor changes states. Click Next.
6. On the Service Name page, start defining the service you are going to monitor by selecting the ellipsis button next to Service Name.
7. On the Select Windows Service page shown in Figure 5.27, select the ellipsis under Computer Name and select the name of a computer that hosts the appropriate event log.
8. Choose the appropriate service from the list and click OK.
9. Click Next.
10. The Configure Health page allows you to specify the state you are looking for. Typically you can leave the setting here as configured.
11. Click Next.
12. The Configure Alerts page allows you to configure the alerts that are raised by this monitor. Make sure that the Generate Alerts For This Monitor option is selected.
13. In the Generate An Alert When pull-down menu, select the level that you want to present an alert: The Monitor Is In A Critical Health State or The Monitor Is In A Critical Or Warning Health State.
14. The Automatically Resolve Alert When The Monitor Returns To A Healthy State check box should be deselected only if you want to have an operator be responsible for resolving the alert.

FIGURE 5.27
Choosing the
appropriate service
to monitor



15. Select the alert Priority: Low, Medium, or High.
16. Select a Severity: Information, Warning, or Critical.

SNMP Probe-Based Unit Monitor

Simple Network Management Protocol (SNMP) is used in many organizations to monitor and control devices, which includes network devices such as routers and switches as well as computers. SNMP has been around for a long time, providing an efficient monitoring solution that many administrators trust. Part of its efficiency is due to the fact that SNMP does not have many commands that it relies on. SNMP clients use simple rules to watch for a monitored object to fall outside of the parameters set by the administrator, and then will send in a “trap.”

SNMP is based on a set of standards and has a governing body that controls the identifiers for all of the monitored objects. Each of the objects is assigned an identifier, known as an object identifier (OID). A collection of OIDs make up the Managed Information Base (MIB) for monitored solutions. MIBs exist for many of the Windows servers and services as well as managed network devices. Each company that wants to have their solution monitored using SNMP submits a request to have their OIDs ratified and accepted as a standard.

NOTE If you would like to read a more detailed description of SNMP, check out <http://msdn2.microsoft.com/en-us/library/aa379100.aspx>.

That quick summary of SNMP was meant to give you a rudimentary understanding of how SNMP works. As you create and work with SNMP probe-based unit monitors, you will find that you will have to identify the OID of the object you want to monitor. This unit monitor queries the object and receives responses, which you can then use to set the health of the object. With the SNMP probe-based monitor, you have the opportunity to set two different probes. The first probe monitors for the object to go into an unhealthy state, at which time it sets the health to warning or error. The second probe you define is used to set the health back to the success state.

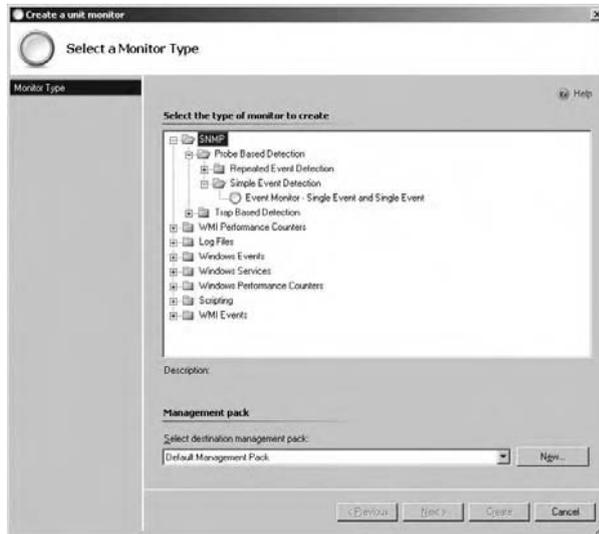
To create an SNMP probe-based unit monitor, open the operator’s console and select the Authoring view. In this view, navigate through the Authoring pane until you reach the monitors. To do so, expand Authoring > Management Pack Objects > Monitors. Since several object types can be monitored, you must generalize the view. At the top of the operator’s console, click the Scope button. When the Scope Management Pack Objects By Target(s) dialog box appears, as

shown in Figure 5.12, click the View All Targets radio button and type **SNMP Network Device** in the Look For text box. Select the SNMP Network Device check box and then click OK.

Now the view should show only the SNMP Network Device object type. Once you expand the SNMP Network Device object type, follow these steps:

1. Expand Entity Health in order to see the Availability monitor level.
2. Right-click the Availability node, point to Create A Monitor, and select Unit Monitor. This will start another wizard: the Create Monitor Wizard.
3. On the Select A Monitor Type page, expand SNMP > Trap Based Detection > Simple Trap Detection; select Event Monitor – Single Event And Single Event, as shown in Figure 5.28; select the unsigned management pack you want to use, and then click Next.
4. On the General Properties page, name and describe the monitor.
5. The final option on this page is the Parent Monitor list. Define the monitor that will be affected when this monitor changes states, and click Next.
6. On the First SNMP Probe page, enter the community name in the Community String text box.
7. In the Frequency box, enter the frequency that you want to collect. Event collection frequency can be anywhere from 30 seconds to 604,800 seconds (7 days).

FIGURE 5.28
Selecting the SNMP
Probe event monitor



8. In Object Identifier Properties, enter the OIDs of the objects you are monitoring.
9. Click Next.
10. The Build Event Expression page appears, allowing you to define the event that you want to watch for that indicates an unhealthy condition.
11. Click Next.

12. On the Second SNMP Probe page, enter the community name in the Community String text box.
13. In the Frequency box, enter the frequency that you want to collect. Event collection frequency can be anywhere from 30 seconds to 604,800 seconds (7 days).
14. In Object Identifier Properties, enter the OIDs of the objects you are monitoring.
15. Click Next.
16. The Build Event Expression page appears, allowing you to define the event that you want to watch for that indicates an healthy condition.
17. Click Next.
18. The Configure Health page allows you to configure the event states for the events that you just identified in the previous steps. In the First Event Raised option, select the name in the Operations States column and enter a new name for the event. In the Health State column, choose Critical or Warning.
19. In the Second Event Raised row, select the name in the Operations States column and enter a new name for the event. In the Health State column, choose Healthy.
20. Click Next.
21. The Configure Alerts page allows you to configure the alerts that are raised by this monitor. Make sure that the Generate Alerts For This Monitor option is selected.
22. In the Generate An Alert When pull-down, select the level that you want to present an alert: The Monitor Is In A Critical Health State or The Monitor Is In A Critical Or Warning Health State.
23. The Automatically Resolve Alert When The Monitor Returns To A Healthy State check box should be deselected only if you want to have an operator be responsible for resolving the alert.
24. Select the alert Priority: Low, Medium, or High.
25. Select a Severity: Information, Warning, or Critical.

THRESHOLD MONITORS

Threshold monitors are used to keep an eye on the performance of an object. Anything that you could add into a performance log can be monitored using a threshold monitor. That means that any counter can be proactively monitored by Operations Manager, and an action can be taken if the counter falls out of the criteria you have selected. If you want to know when processor utilization peaks over 90 percent, or when excessive page faults occur, threshold monitors can let you know. Operations Manager 2007 offers two types of threshold monitors: static and self-tuning.

Static Threshold Monitor

Static threshold monitors are used to make sure that an object does not fall outside of safe parameters. Of course, you have to know which limits you are watching for. Some are well documented; others you will have to glean from a baseline of your system. Once you have determined what the

peak value is, you can create a static threshold monitor. During the schedule that you set, the agent will watch for the system to exceed the threshold, and an alert or health state can be modified.

Single Threshold Monitor When creating a single threshold monitor, you only get to define one threshold limit. A good use for this monitor is to monitor free space on a drive. You can determine what would be a good notification level and set the monitor to alert you when the free space falls beneath that level. Network utilization can be watched to make sure it does not exceed a specific percentage of total availability.

Double Threshold Monitor The double threshold monitor allows you to control the monitor state at different threshold levels. You can monitor an object and set a warning level and an error level. You set the first value so that you can detect when a problem may be approaching. For instance, you can monitor the network utilization for a SQL server and set the monitor to a warning state when the network utilization reaches 60 percent. Then set the second value to warn you if the counter continues to climb or recedes past a second level you define. For our SQL server, we can set the monitor value to 75 percent so that the warning state will appear once the total utilization reaches or exceeds that level.

Thresholds come in four flavors. When deciding which one you would like to use, determine what type of monitoring you are trying to accomplish. Each threshold type has its place in your monitored solutions; use the one that will provide the best method of monitoring.

Simple Threshold The simple threshold is the easiest to comprehend and use: you simply decide a value for the maximum or minimum level you want your object to achieve. If you want to make sure that a storage group in your Exchange server does not exceed 40 GB of drive space, you set the threshold value to 40 GB, and if that level is achieved, the monitor will go into a warning or error state. As long as the amount of space consumed is less than 40 GB, the state of the monitor will show as Success.

Consecutive Samples Over Threshold This type of monitor alleviates some of the false positives that you may obtain when using a simple threshold. Simple thresholds only monitor when an object exceeds a limit, but do not take into account those objects that may spike periodically. Processor utilization and network utilization are good examples of this. There may be short bursts of network activity, or the processor might be momentarily taxed, but they quickly return to normal activity. The consecutive samples over threshold will not change its state until several samples have been taken that indicate the object is still over the threshold value. For this type of threshold, you must specify the warning or error level that you are watching for, as well as the number of samples that must be taken that show the threshold being exceeded.

Average Threshold Average thresholds provide the ability to average the values that are obtained during a span of time. With this threshold type, you can specify that you want to make sure that the object has been over the threshold value based on the average value of the sampled data. You must specify the value that you consider the average threshold, and also specify the number of samples you want to average. If, for instance, you want to make sure that the number of page faults have been averaging less than 2 over the past 10 samples, you can set the average threshold to obtain a sample once every minute, and then average the past 10 minutes' activity.

Delta Threshold Sometimes you do not need to configure a set threshold value but instead want to determine if there has been a great deal of change over time. This threshold is comparable to the average threshold, but it takes into account the amount of change that an object has. A good example of this is an SMTP queue for your email solution. As email enters and exits your corporate email servers, the queue will grow and fall depending on the email that is flowing at that moment. You may want to monitor to make sure that your queues are not filling up due to a problem communicating with outside SMTP servers, or that you do not have an attack occurring that is causing your SMTP queues to grow abnormally. With this threshold, you must

configure the number of samples that you want to take as well as the threshold value that you want to monitor.

Self-Tuning Threshold Monitor

Chapter 11 has a complete section devoted to self-tuning threshold monitors. To have a better look at how these monitors function, check out that chapter. We are only going to provide a brief description of these monitors here because they round out this topic.

Self-tuning threshold monitors are new to Operations Manager 2007 and do not have a comparable function in any of the previous versions. These advanced monitors take samples of object counters and learn the normal operating parameters of your systems. This means that you do not have to perform the same baseline analysis of your systems as you do with other threshold monitors. The self-tuning threshold monitor performs its own baseline analysis and then starts monitoring for behavior discrepancies. The baseline that is created includes the standard operating behavior of your organization.

Since most systems take on a different load level at different times of the day, a standard threshold monitor will not be as effective. You would have to perform a baseline, and then decide the out-of-range values that you want to use for the threshold values. The value that you use has to be high or low enough that peak activity will not trip the monitor. At the same time, it has to be set at a level that abnormal activity will still be detected.

The self-tuning threshold monitor detects abnormal behavior during each of the phases of your operation. The database of samples that are collected determines what the typical operating parameters of the counter are at specific times of the day. The operating parameters at 8:00 AM are probably different than those at 12:00 PM. The database will contain the baseline for the time period, and if the current sample falls outside that baseline, the monitor's state will change.

When you create a self-tuning threshold, you have the option of configuring a two- or three-state monitor. Both these types of monitors have three areas that you can define: the area in the baseline, the area above the baseline, and the area below the baseline. These areas are then used to define where the healthy, warning, and error states are. In a two-state monitor, you define whether you want the error or warning state to be defined above or below the baseline. In a three-state monitor, you define whether the error state is above or below the baseline, just as with the two-state, but since the baseline is considered healthy, the remaining area is considered the warning level.

Note that you will have to create a baseline for each instance for the counter you are monitoring. For example, if you have a dual-processor system, each processor is considered an instance. You must configure the monitor to watch each processor individually. You cannot select the All Instances option to have both processors monitored as a single entity. This is because the monitor creates baselines on each instance and cannot alert on a combined view.

Using Rules

We have been discussing monitors up to this point, and for good reason—monitors allow you to view the health of objects to make sure that those objects are working as desired. Monitors can raise alerts based on the health of the object, and you can use them to verify the health of the entire service path. Rules can also be used to raise alerts based on how an object is functioning, but they do not allow you to monitor the health of the object.

If you have an object for which you do not need to view health information but you still want to monitor the status of that object, rules work very well. You must decide whether the object is part of an end-to-end service monitoring solution. Monitors can be used as either parent or child monitors in an end-to-end monitoring solution, but rules cannot. If the object you are planning to monitor is a stand-alone object, you can use a rule to monitor it.

There are several rules types that you can use when monitoring objects: alert-generating, collection-based, performance, and SNMP rules. For some of these, you can set criteria that you want to use for the rules.

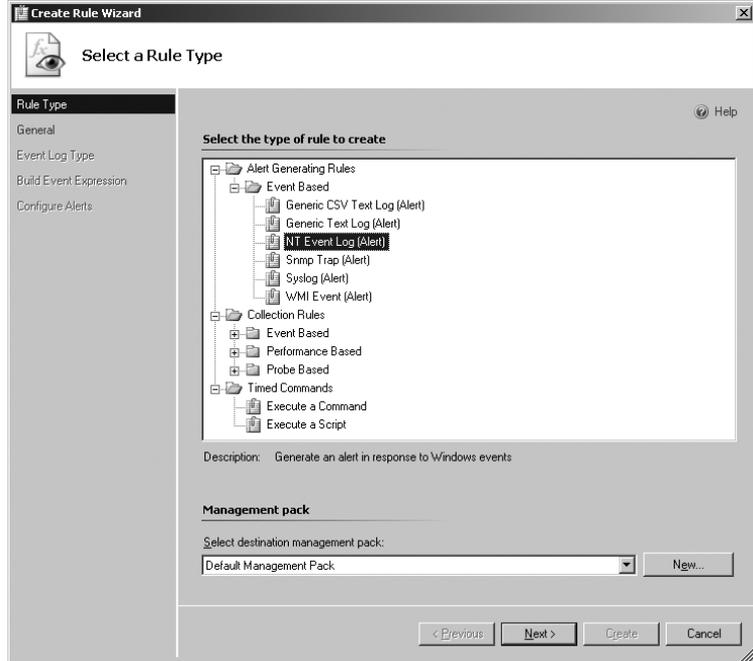
Each of the rule types can be created in the Operations Console by selecting the Authoring view. After expanding the Management Pack Objects node and right-clicking Rules, select Create A New Rule from the context menu. The Create Rule Wizard starts and presents the Select A Rule Type page. Here, you will start creating each of the various rules.

NT EVENT LOG–BASED RULE

An NT event log–based rule is used to parse an event log and discover an event based on the criteria that you enter in the wizard. Using this rule, you can detect when a specific event occurs on a system and take immediate action on it. Since many of the events are well documented and there are solutions already devised, you can provide your operations staff with exact instructions to follow to alleviate the problem.

1. On the Select A Rule Type page, select Alert Generating Rules > Event Based > NT Event Log (Alert), as shown in Figure 5.29.
2. Choose the management pack that you would like to use, or click New to create a new management pack to use. Click Next.
3. On the Rule Name And Description page, enter a name in the Rule Name text box and a description in the Description text box.

FIGURE 5.29
Choosing the NT event
log rule



4. Select a target by clicking the Select button and clicking the target you want to monitor. The target list can be scoped or filtered, just as the Details pane can be.
5. Rules are enabled by default, but you can disable the rule by clearing the Rule Is Enabled check box. Click Next.
6. Choose the log file that you want to monitor by selecting its name from the list on the Event Log Name page, and then click Next.
7. On the Build Event Expression page, select the appropriate event information and click Next.
8. When the Configure Alerts page appears, enter a name for the alert that will be generated.
9. You can enter a description to provide further information about the alert. Notice that you can click the ellipsis and select variables that are included in the Description text field.
10. Select a priority and severity from the lists provided.
11. The Custom Alerts field allows you to define how the alert is going to report the alert to you. You have the option of entering the message you would like to send, or you can use the ellipsis to input variables that will be used.
12. You also have the option to suppress alerts by clicking the Alert Suppression button and then providing the information about the fields you would like to base the suppression on.
13. Click the Create button to finish the wizard.

NT EVENT LOG EVENT COLLECTION RULE

The event log collection rule gathers the events that match the settings you specify in the rule. Once collected, they appear in the Event View for the agent-managed system.

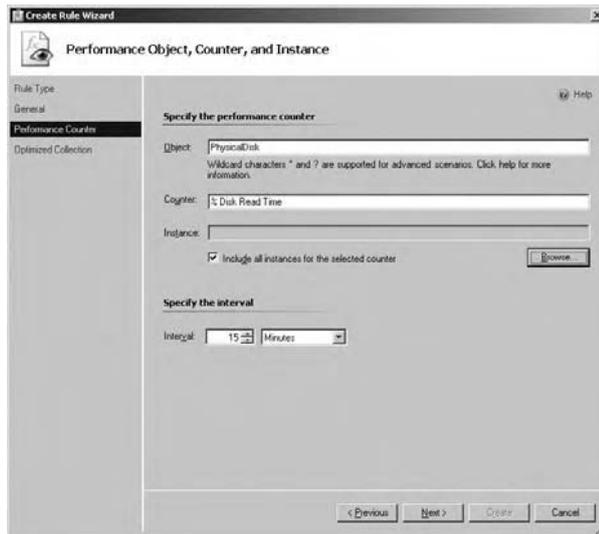
1. On the Select A Rule Type page, select Collection Rules > Event Based > NT Event Log (Alert).
2. Choose the management pack that you would like to use, or click New to create a new management pack to use. Click Next.
3. On the Rule Name And Description page; enter a name in the Rule Name text box and a description in the Description text box.
4. Select a target by clicking the Select button and clicking the target you want to monitor.
5. Rules are enabled by default, but you can disable the rule by clearing the Rule Is Enabled check box. Click Next.
6. Choose the log file that you want to monitor by selecting its name from the list on the Event Log Name page and click Next.
7. On the Build Event Expression page, select the appropriate event information and click Next.
8. Click the Create button to finish the wizard.

WINDOWS PERFORMANCE COLLECTION RULE

As with the event collection rules, the performance collection rule gathers the specified object performance data and stores it in the database. Once stored, the performance data can be viewed from reports or the Performance view.

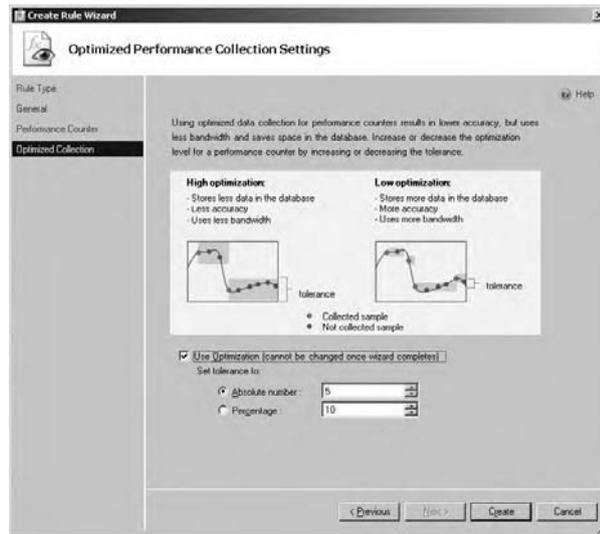
1. On the Select A Rule Type page, select Collection Rules > Performance Based > Windows Performance.
2. Choose the management pack that you would like to use, or click New to create a new management pack to use. Click Next.
3. On the Rule Name And Description page, enter a name in the Rule Name text box and a description in the Description text box.
4. Select a target by clicking the Select button and clicking the target you want to monitor.
5. Rules are enabled by default, but you can disable the rule by clearing the Rule Is Enabled check box. Click Next.
6. When the Performance Object, Counter, And Instance page, shown in Figure 5.30, appears, click the Browse button and enter the name of a computer that has the counter you want to monitor.
7. From the Object list, select the performance object that you want to use in order to identify the counter.
8. When the counters appear, choose the counter you want to monitor, or click the All Instances button. Then click OK.
9. In the Interval page, enter the time interval that you want to use when collecting performance data and click Next.

FIGURE 5.30
Performance Object,
Counter, and
Instance page



10. On the final page, shown in Figure 5.31, you can select Use Optimization to enter a value that is used to optimize the data sent by the performance collector. This way, you do not send too much data across the network. You can select Absolute Number and enter a digit, or you can select Percentage and enter a percentage of change that has to occur before the counter data is forwarded.
11. Click the Create button to finish the wizard.

FIGURE 5.31
Optimization settings



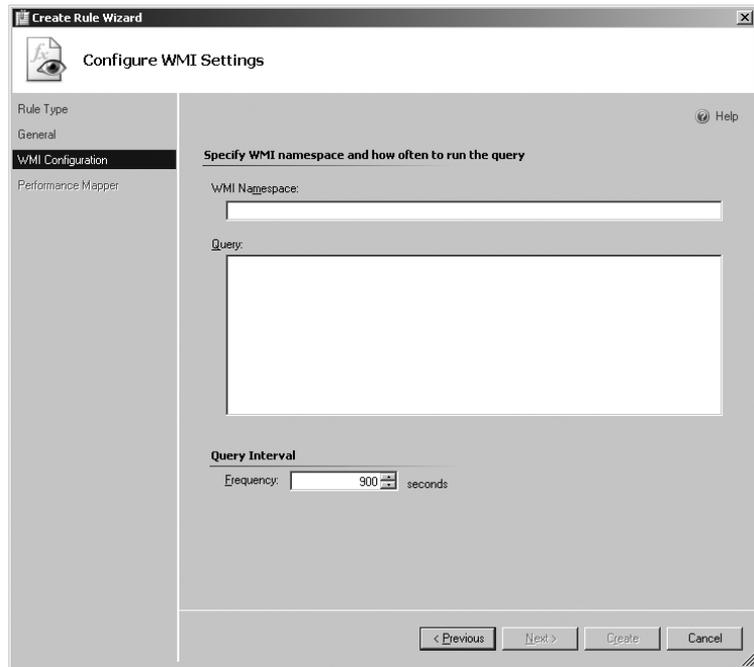
WMI PERFORMANCE COLLECTION RULE

Whereas the Windows performance rule collects performance counter data as it is occurring, the WMI performance collection rule gathers the performance data that is collected by WMI.

1. On the Select A Rule Type page, select Collection Rules > Performance Based > WMI Performance.
2. Choose the management pack that you would like to use, or click New to create a new management pack to use. Click Next.
3. On the Rule Name And Description page, enter a name in the Rule Name text box and a description in the Description text box.
4. Select a target by clicking the Select button and clicking the target you want to monitor.
5. Rules are enabled by default, but you can disable the rule by clearing the Rule Is Enabled check box. Click Next.
6. On the Configure WMI Settings page, shown in Figure 5.32, type the namespace, typically root\cimv2.
7. Type your query in the Query box.
8. Enter the frequency at which the counter should be sampled and click Next.

9. The Performance Mapper page appears; enter the object name in the Object text box.
10. Enter the counter name you want to collect in the Counter text box.
11. Enter the instance of the counter that you want to collect information from.
12. Enter the name of the property for which you are monitoring by clicking the Value button and entering it in the dialog box that appears after you click Data ➤ Property.
13. Click the Create button to finish the wizard.

FIGURE 5.32
Configuring the WMI namespace



PROBE-BASED EVENT COLLECTION RULE

When no other rule will do, the probe-based collection rule will handle the job. The other rule types are already defined as to what the scope of the rule is, but the probe-based rule allows *you* to define how the rule will collect data and access the data. You'll notice that as you start working with this rule type it is designed to collect information only; it should not be used to make any changes to the object that you are collecting from. If you did use the rule to alter the current state of the object, you could alter the way the object functions.

1. On the Select A Rule Type page, select Collection Rules ➤ Probe Based ➤ Script (Event).
2. Choose the management pack that you would like to use, or click New to create a new management pack to use. Click Next.
3. On the Rule Name And Description page, enter a name in the Rule Name text box and a description in the Description text box.

4. Select a target by clicking the Select button and clicking the target you want to monitor.
5. Rules are enabled by default, but you can disable the rule by clearing the Rule Is Enabled check box. Click Next.
6. The Schedule page allows you to configure the interval at which the rule will run. Enter the interval value, and then select Synchronize At and specify a start time for the rule.
7. On the Scripts page, name the script and then enter the script lines that will collect the performance data.
8. Enter the maximum run time in the Timeout box and click Next.
9. On the second Script page, use the Computer, Event Source, EventID, Category, and Level fields to map the script parameters to the event.
10. Click the Create button to finish the wizard.

PROBE-BASED PERFORMANCE COLLECTION RULE

Whereas the probe-based event collection rule will collect data from nearly any type of source file that you need, the probe-based performance collection rule can gather performance data from nearly any type of performance counter that is stored in the WMI. So if you have a product or service that stores any type of performance information in the WMI, you can gather that data using this rule.

1. On the Select A Rule Type page, select Collection Rules > Performance Based > WMI Performance.
2. Choose the management pack that you would like to use, or click New to create a new management pack to use. Click Next.
3. On the Rule Name And Description page, enter a name in the Rule Name text box and a description in the Description text box.
4. Select a target by clicking the Select button and clicking the target you want to monitor.
5. Rules are enabled by default, but you can disable the rule by clearing the Rule Is Enabled check box. Click Next.
6. The Schedule page allows you to configure the interval at which the rule will run. Enter the interval value, and then select Synchronize At and specify a start time for the rule.
7. On the Scripts page, name the script and then enter the script lines that will collect the performance data.
8. Enter the maximum run time in the Timeout box and click Next.
9. The Performance Mapper page appears; enter the object name in the Object text box.
10. Enter the counter name you want to collect in the Counter text box.
11. Enter the instance of the counter that you want to collect information from.
12. Enter the name of the property for which you are monitoring by clicking the Value button and entering it in the dialog box that appears after you click Data > Property.
13. Click the Create button to finish the wizard.

Reporting

Most management packs contain report configuration settings that are imported into the Reporting Server if you have it set up. Having the reports already built for you allows you to start generating reports immediately. While you may find some of the included reports handy, you should not assume that all the reports that you need will be included. You will find yourself building some reports of your own to fulfill your own needs.

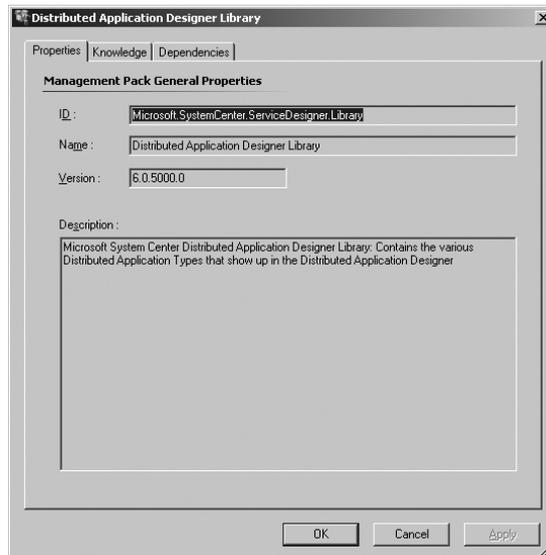
In Chapter 9, we are going to review reports in greater detail. At that point we will look at some of the predefined reports available when you import a management pack, as well as work with Report Builder to create a new report.

Overrides

Overrides have been available throughout the different iterations of Operations Manager. In previous versions, overrides were not widely utilized because an administrator could open a management pack and make changes to the settings. Many organizations that had implemented a strict change management philosophy placed restrictions on the changes that could be made to a management pack's settings without first going through a complete change management review. To implement a quick fix to silence a noisy management pack setting and restore a little order for the service desk, overrides could be used. The override would effectively change the way the setting was working, but it wouldn't change the original structure of the underlying management pack. Instead, the override would target a specific rule and disable it from acting on specific agent-managed systems.

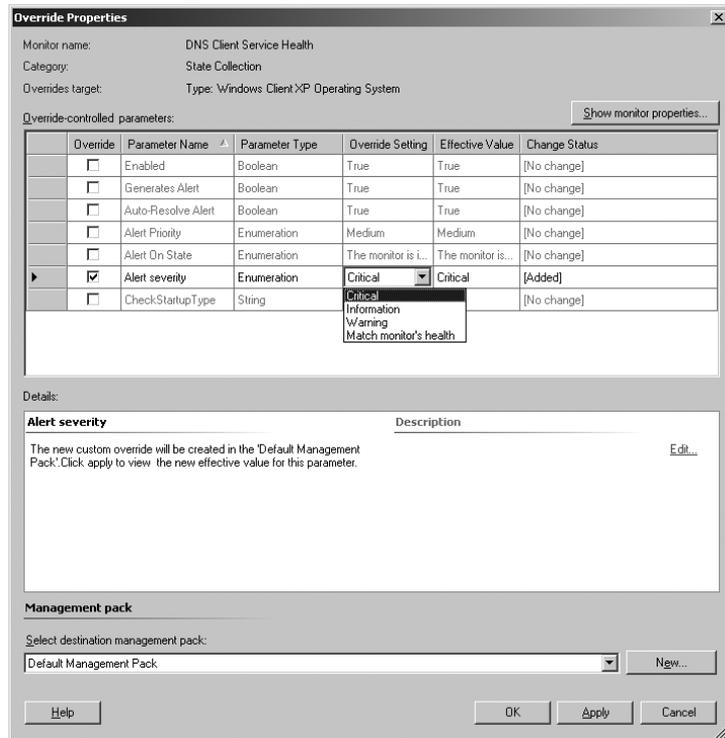
Sealed management packs help protect a vendor's intellectual property. Settings and configuration details that they don't want other companies to have access to can be hidden away in the sealed management pack. Of course, that means administrators no longer have the option to modify the settings in the management pack—not directly at least. If you right-click on a sealed management pack setting in the Administration view, then click Properties to open the Properties page, as shown in Figure 5.33, you will notice that there are no configurations options that directly affect the way the setting functions. To allow an administrator to change the setting, an override has to be used.

FIGURE 5.33
The Properties page
for a sealed manage-
ment pack



In essence, an override is a means of making sure that a setting in the management pack works effectively in the network. Depending on what type of setting you are attempting to override, you will find options in the override criteria that allow you to modify the settings parameters. In Figure 5.34, an override is being configured for a setting in the Windows XP Professional management pack. Notice that the management pack settings that can be changed through the override are exposed in the override window.

FIGURE 5.34
Override options



The settings that are exposed are controlled by the provider of the management pack. If they don't want you to see something, they don't have to make it available. Usually the settings that are available are those that directly affect the functioning of the management group. In the DNS setting shown in Figure 5.34, notice that the only settings available are the options to turn the alert on or off. Using the pull-down menu, you can specify that the alert not report in the management pack. Thus, you can suppress some of the rules that you do not find pertinent. In previous incarnations, you could open the management pack and make the change directly. Now, as long as the management pack is sealed, you have to use an override.

Since overrides are so important to optimizing the way your Operations Manager infrastructure runs, we have provided more information about them in Chapter 11.

Converting Legacy Management Packs

Although most administrators will not create management packs, those who already have Microsoft Operations Manager 2000 or 2005 in their environments may be required to convert existing management packs. Not all of the management packs that are in use will have to be converted.

Microsoft supplies native management packs for its services and systems, but third-party providers may not have a management pack available when you are ready to start running Operations Manager 2007.

For this reason, Microsoft has supplied a management pack conversion utility. Your first step is to identify the management packs that need conversion. If you have a homemade management pack for a service that your company provides, or a homegrown application that you monitor, you can add them immediately to your list. For services and applications from other vendors, make sure that the vendors have not released an updated management pack. Check their website and/or contact them directly to determine what their plans are. You may find that although they do not have something currently available, they may have a release coming out in a time frame that you can live with.

Also, make sure you know which customized settings you have configured in the existing management packs. While the conversion may keep those settings intact, you may find that some of them don't convert the way you want them to. If this is the case, make sure that your documentation is up-to-date so that you know which settings to further modify.

Step one when converting your existing management packs is to convert them to an XML file. The native AKM format is not supported by the Operations Manager 2005 to Operations Manager 2007 conversion program. The program that you will need to use is included in the Microsoft Operations Manager 2005 Resource Kit: it's called MP2XML. Step two is to convert the new XML file that you created into an Operations Manager 2007 management pack. The program you will use for this purpose is included on the Operations Manager 2007 CD: it's called MPCConvert.

The syntax for MP2XML is very straightforward. All you need to do is specify the name of the management pack that you want to convert and then enter the name of the resulting file. However, you can run this command only on a MOM 2005 server, so don't attempt it from your workstation. The easiest way to run the program is to copy it to the directory where the original management pack exists and then define the resulting XML filename. However, if you want to place the XML file in another location, or the original is not in the same directory as MP2XML, you can always use the fully qualified paths. The syntax is as follows:

```
MP2XML.exe SourceDirectory\SourceManagementPackName.akm
➔DestinationDirectory\DestinationManagementPackName.xml
```

After you have a management pack in XML format, you can convert it to be used within Operations Manager 2007. Unlike MP2XML, this conversion can be performed on any Windows Server 2003 or Windows XP system as long as .NET Framework 2.0 is installed. Just like MP2XML, the conversion command is quite easy to use. All you need to know is the location and name of the XML file you created and the location and name of the new Operations Manager 2007 management pack. The syntax for the MPCConvert command is as follows:

```
MPCConvert.exe SourceDirectory\OpsMgr05XMLFileName
➔DestinationDirectory\OpsMgr07XMLFileName
```

Once that step is completed, you will have a management pack that is ready to import into your Operations Manager 2007 management group. Before you jump right into importing the management pack, however, verify that the conversion worked and that there is nothing syntactically incorrect with the management pack. To help you with this, Microsoft has included the MPVerify utility. It also resides on the Operations Manager 2007 CD. Copy it to the system where you have created the new management pack and then run it using the management pack name in the command line.

So at this point you know that we can import our newly converted management pack into Operations Manager 2007, but you also know that there are new features in Operations Manager 2007 that were not available in MOM 2005. Additionally, you know there has been some deprecated functionality from MOM 2005. Moving from the old to the new will require you to know the differences between the management packs so that you can identify the correct functionality. Let's take a look at what has changed and the corresponding functionality between the two versions.

Alert Severity/Health State Operations Manager 2005 used alert severity to define the state of the object that was being monitored. The severity levels were Success, Warning, and Error. The same three levels exist in Operations Manager 2007, but they are now referred to as health states.

Rule/Monitor Rules that change the state of the object from healthy to unhealthy, or vice versa, are converted to a monitor. Monitors can change the state value as well as generate alerts.

Rule/Rule For those rules that do not affect the state but do generate alerts, a rule is created.

Rule/Collection Rule Performance-based rules are converted to a collection rule type.

Computer Group/Computer Group Class Operations Manager 2005 was used for monitoring servers. Once a computer was discovered, scripts from the management packs that were installed on management servers would discover the applications on the server. In Operations Manager 2007, since the primary responsibility of the management group is to monitor services, the discovery methods used to discover and populate the computer groups are converted to computer groups.

Computer Group/Installation Class Operations Manager 2005 was used for monitoring servers. Once a computer was discovered, scripts from the management packs that were installed on management servers would discover the applications on the server. In Operations Manager 2007, since the primary responsibility of the management group is to monitor services, the discovery methods that are used to discover applications and services are converted to installation classes.

Script/Module Type All of your scripts that have been created in Operations Manager 2005 still function in Operations Manager 2007, even though Operations Manager 2007 uses a different scripting engine. As you familiarize yourself with the new scripting capabilities, go back and rewrite the scripts to make them more efficient.

Task/Task Only those tasks that are configured to run through the agent are converted. Other types of tasks, such as those set to run on the management server or via the Operations Console, are not converted.

View/View Administrator console views are converted to Operations Console views.

Notification Group/Notification Rule Notification groups are not converted because there is no corresponding notification item in Operations Manager 2007. The notification settings are converted to notification rules.

Knowledge/Knowledge Article Company knowledge is converted to knowledge articles.

Topology/View Topologies are converted to views so that you can still view the diagrams created by the management pack.

No longer used in Operations Manager 2007 are the Operator, Report, Notification Group, Filter Rules, and Console Scope objects. These objects are not converted when the MPConvert utility is run against the Operations Manager 2005 management pack XML file. You may need to configure analogous objects in Operations Manager 2007 to make sure you are providing the same level of monitoring.

Removing Management Packs

As your organization changes, you may find yourself no longer having to monitor certain services. Or you may no longer be responsible for monitoring something as other divisions become accountable for the service. If that is the case, you can remove a management pack from your management group. Removing a management pack removes the management pack and all associated views from the Operations Console; it also allows the data that had been collected and stored in the database to be deleted.

Removing the management pack is as easy as selecting the Delete option from the context menu. Keep in mind that most management packs rely on parent management packs to obtain some of their settings. Do some investigating to make sure you understand the relationship between the management packs you are using. Especially make sure that a management pack you are removing is not required by any other management pack.

To remove a management pack:

1. Open the Operations Console.
2. From the Administration view, click the Management Packs node.
3. Select the management pack that you no longer need.
4. Right-click on the obsolete management pack and select Delete from the context menu.
5. Click Yes in the warning dialog box that appears.

The Bottom Line

Identify management pack formats. Management packs come in two flavors: sealed and unsealed. Sealed management packs can be imported into your management group, but you cannot modify the settings. Unsealed management packs are XML files that you can alter to meet your needs.

Work with discovery rules. Along with discovering computers so that we can install agents, discovery rules also determine which monitored objects exist in the management group. Using discovery rules, you can target the monitoring rules and responses to the correct objects.

Create groups based on discovery. After discovering the objects in the management group, groups can be used to organize the objects so that they can be monitored easily.

Create monitors and rules. Monitors examine the objects in your management group and will not only alert an operator if necessary but also display the health status of the object. Rules are used when you only need to alert, not display health status.

Master It! Create a new group to monitor five specific domain controllers. These domain controllers are used by the research and development department, and you want to make sure they are not overloaded.

1. Open the Operations Console.
2. Select the Authoring workspace.
3. In the Authoring pane, select Groups.
4. Right-click Groups and select Create A New Group.

5. In the Name field, enter **R&D Domain Controllers**.
6. Click the New button to create a new management pack.
7. In the Name field, enter **R&D Domain Controllers** and click Next.
8. Click Create on the Knowledge Article page.
9. Click Next.
10. On the Select Members From A List page, click the Add/Remove Objects button.
11. On the Object Selection page, click Search.
12. When the list of entities appears, select the domain controller and click Add, then OK.
13. Click Next.
14. On the Create A Membership Formula Page, click Next.
15. On the Choose Optional Subgroups page, click Next.
16. On the Specify Exclude List page, click Create.

Chapter 6

Authoring and the Management Pack Life Cycle

Every management pack that you import into a management group goes through the same changes. Administrators need to know how to control the changes that will be made to management packs as well as the effect those changes will have on their organization. Adversely affecting your existing system will cause you to have a very long day, if not a long week. Continual problems could even brand you as a troublemaker or worse. To gain the trust of the rest of the staff, as well as the entire employee population, you need to make sure that you can introduce and maintain the management packs in an efficient and controlled manner.

Throughout this chapter we are going to discuss some of the methods you can use to ensure that the changes you introduce will not cause undue distress among the rank and file. We'll show you a method that has worked out well for us in the past, and then you can take this knowledge and customize it to fit your needs.

In this chapter you will learn how to:

- ◆ Document management pack changes
- ◆ Prepare the management pack for release
- ◆ Control versioning of the management pack

Authoring Management Packs

If you are familiar with the legacy management packs, the changes that have been made to the management pack structure since the introduction of Operations Manager will take some getting used to. Most management packs are not stand-alone units any longer. They build on the configuration settings of other management packs to provide additional management capabilities. With this in mind, you will find that creating new management packs has become an art unto itself.

A prerequisite for creating management packs is making sure you know how other management packs are built. Knowing what has already been configured in other management packs will keep you from duplicating settings. If you take a look at the management packs that are provided from Microsoft, you will see that there are only a couple of base management packs, and the others inherit those settings, further defining what is to be monitored.

Most administrators will probably never need to create their own product management packs. Instead, they will import ready-made management packs provided by the developer of the product they are monitoring. For those that are providing the management packs for their products, however, they will need to understand how the new management packs are structured. Operations Manager administrators, on the other hand, will create management packs that will be used for special monitors or rules that they need to define, as well as to host overrides.

Creating New Management Packs

As mentioned, creating a management pack is something that will probably be taken on only by vendors, not by administrators. But vendors will need to have access to some tools that will allow them to build their management packs. Microsoft has issued a management pack authoring tool as well as documents on how to create management packs. Of course, all that is truly needed is a text editor. A simple text editor can be used to create the XML files that will be imported into Operations Manager. This may be too simplistic of a solution. Being able to work with an XML file in a text editor such as Notepad is nice, but with all of the settings that make up a management pack, it is probably out of the question to use it. Trying to troubleshoot problems can be very difficult.

When you are creating new rules or monitors, or if you are creating an override for an existing management pack object, you can create a new management pack by clicking the New button next to the list of available management packs. As you create the new management pack objects, they are added to the management pack without you having to know how the XML format should look.

Exploring a Typical Life Cycle

As you will probably find out, rare is the product that you can install and use as it comes out of the box. Typically, you need to configure and customize the solution so that it fits your needs. Making sure that an application or service meets your needs is always the start, and then you will tweak and optimize it so that it works the way you want it to. Once you have it working perfectly, you can sit back and relax, allowing it to do its job. Well, maybe that is a bit of a stretch.

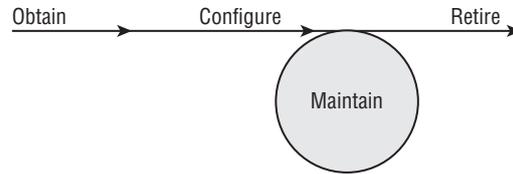
Just as it is rare that you will find a product that completely meets your needs, having a solution that does not need updating and periodic maintenance is nearly impossible. Several factors come into play that may force you to change the way a solution is configured. Such factors include security holes, changes in the infrastructure, and updates to operating systems and other applications.

Management packs are no exception to these examples. While you might think that once you have the management pack in place you will be able to let it run unimpeded, you would be mistaken. Outside influences will force you to make changes from time to time. You may introduce new systems into your infrastructure that start causing some of your rules to become a little noisier than you had originally seen. Your management team may decide they need to start receiving additional information about the managed systems, and you may have to make changes to suit their needs.

No matter what, your management solution will periodically be in a state of flux. To maintain the changes that you have to introduce without causing additional headaches, introduce management pack life-cycle controls. This approach allows you to make changes, track the effects of those changes, and revert back to the original settings if necessary.

This chapter provides instructions and discussions on versioning control of your management packs. Note that the methods described here have proven themselves, but they are not the only way to provide versioning control. Other methods have been employed by companies that have proven just as effective as what you will find here. This chapter is meant to be a guideline for those who have not implemented their own versioning controls. We won't take any offense if you consider your own methods to be better than what we describe here.

Each management pack will go through the same basic life cycle. Some management packs may end up going through many more configuration changes than others, and some may have a much longer lifetime than others. But throughout their lifetime, they will see the same phases, shown here:



Obtain This can be the easiest of the phases for most management packs. Vendors are supplying management packs for their products on an ever-increasing basis. They are realizing that Microsoft is making a sincere investment in their System Center line, and Operations Manager is one of the pinnacle servers in that line. If the vendor of the application or service provides a management pack for their product, this part of the process is easy. If a management pack does not exist, you may have to build a management pack. Either way, once you have the management pack ready to install into your environment, you are ready to move on to the next phase: configure.

Configure Few management packs are going to meet your needs 100 percent of the time. Vendors try to judge which settings are needed by most of their customers. They interview their customers and work with them to determine what they want to monitor. They also review the application data that is generated while it is running and make some educated guesses about what a customer would want to be alerted about or see in a report. But not all the settings that they think should be turned on are appropriate for your organization. The same goes for disabled settings. You may want to monitor something that the vendor didn't find critical. During the configure phase, you will need to evaluate the alerts, events, and performance data collected by the management pack and decide how they will affect your organization.

Maintain This is the phase that you will find most of your management packs in. Once the initial configuration is finished, you have to support the management pack in your environment. As your network grows and changes, the management packs you use will have to be tweaked so that they run efficiently in your organization. You will also run into vendor updates to the management packs, which will force you to modify your existing management packs.

Retire Nothing lasts forever, and so it is with management packs. You may find that the application or service is no longer needed and you are ready to remove it from duty, or you may move to another vendor's product and decommission the original product. Either way, the management pack for the old product will no longer be needed. After removing the application, you will need to remove the management pack from the management group.

Removing the management pack will not be the last step, however. Make sure that you archive the management pack and all its associated files and documentation. If you have kept previous versions of the management pack in an archive and retained the documentation, make sure that all those files are safely backed up and stored somewhere that you can get to easily in case you need to reinstate the management pack. Yes, there could be a reincarnation of the management pack, which would start this entire process over again.

Creating Management Pack Life-Cycle Controls

As with any set of rules or policies, making sure that they are followed is the hardest part of the job. Many network and systems administrators like to “shoot from the hip,” making changes as they see fit, and if a certain change doesn’t work out the way they want, they make another change to see what happens. As you can imagine, in most organizations this strategy doesn’t work out too well. For one thing, making changes without any controls in place may result in unforeseen problems. Also, poorly documented experimentation could lead to more work as the changes are forgotten and the administrator has to go back and try to figure out what they have done.

The four phases that a management pack goes through during its life cycle—obtain, configure, maintain, and retire—should be documented and tracked so that you know which management pack version is in place and which versions have been archived. All the settings applied to each version should also be documented so that you can find the correct version when you are looking for it. In addition, make sure that you have change management controls in place. Every organization performs these steps differently, but the results should be the same: getting a management pack into production with the fewest effects to your infrastructure.

Release Management

A good change management plan will take you from a test environment to production. In Chapter 1, “Overview of Operations Management,” we discussed the ITIL and MOF. Both set standards for change and configuration management. And while everything that we discussed in Chapter 1 is important, one thing that should not be overlooked is release management. Unfortunately, a majority of organizations do ignore this aspect. For some, they need to make a change quickly and will act rashly, hoping that what they do will not adversely impact the rest of the organization.

In smaller companies, this approach may work most of the time. Larger organizations that have complex heterogeneous networks will not be able to make quick untested changes. One error could impact several other systems and have dire consequences. Trying to fix a widespread problem could be more difficult than performing extensive testing to begin with. But performing the testing assumes that you have the proper test environment that mimics the production environment. Sometimes that is just not the case.

Testing must go on, however, and you must make sure your tests are as thorough as possible and are documented. The documentation should include the interactions with other systems and the effect of the change. As you are performing the tests, choose the scientific route. Make a single change and test the results of the change. As you perform each alteration to the system, document the effects and do not move to the next change until you have reached your goal. This way, if you do find that there is a problem with one of your changes, you will know which configuration change you just made.

Once you are ready to move the new management pack into production, take your changes to the release review board. The review board should consist of individuals who have a stake in each of the systems in the organization. They will know their own systems and can tell if you have skimped on your testing, so make sure you are thorough before bringing the change to them. If you have documented all your changes and detailed the effects of those changes, you should be able to get the release approved easily.

Versioning

You now have your management pack configured just the way you want it. It has been moved into production. All of the alerts and events that should be making their way to the database are doing

so perfectly. Performance information is being collected and monitored so that you can see exactly how your systems are performing. All this activity is not causing too much additional traffic on your network. You are identifying problems with your systems before they become serious issues, and you are responding to alerts and rectifying problems faster than ever. Life is good.

Then something happens. You start getting a spike in traffic after adding a few new servers to the mix. Or you find out that a new management pack has been released for something that you are monitoring. Maybe your boss has decided that she wants to view performance criteria that you have not been capturing. Suddenly something has to change.

From the time that you obtain the management pack until the time comes that you retire it, you must make sure you are documenting all the changes. As you create new versions of management packs and alter existing management pack settings, ensure that you are tracking each change and that you keep track of which changes work and which ones fail. Failure to do so may cause you to perform a lot of additional work as you try to re-create the problems or figure out what worked.

Version tracking has been available for developers for quite some time. Developers found out a long time ago that if they did not maintain a record of the changes they made to their programs, they would often lose track of what they had done and would re-create the same thing over and over. Also, they would not have any documentation to help them figure out which changes worked and which ones didn't, so they would have a hard time trying to find a version of their program that was stable before the changes were introduced.

Management Packs That You Created

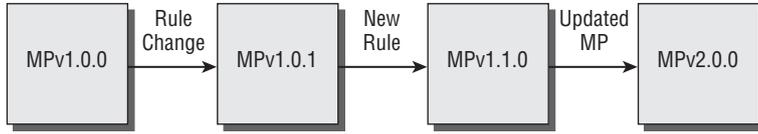
Successful administrators document every change they introduce into their management packs. If they disable or enable a rule, they save a copy of the original management pack, make the change to the rule, and then test the management pack with the new settings. Once the test is complete, they document whether or not the change they made was successful. If it was successful, they introduce the change to their production environment. If it failed, they can roll back the changes by removing the management pack that failed and import the last good management pack back into the management group.

Here are some things you should think about as you are planning your versioning solution. First, decide where you are going to store the old management pack and settle on a naming convention to use. The naming convention is entirely up to you. Choose something that is going to be descriptive enough that you will understand which management pack is which. For most companies, a naming convention based on the management pack name, the version number, and the date of the change works well. You can always get the date from the properties of the file, but having it in the name of the file is handy—especially if you are looking at the filename in a dialog box or pull-down menu.

One way of versioning is to use the full name of the management pack and append a version number. Once you have the version number in place, append the date to the filename. For instance, after you have created a management pack, save the new management pack as `management_pack_name_1.0.0_date.xml`. If this management pack is meant to monitor access to entries in an Active Directory Application Mode (ADAM) instance that you use for federated access (controlling access into your Active Directory forest using Active Directory Federation Services), and you are configuring changes on January 9, 2008, you can save the first version that you are working with as `Federated-ADAM_Management_Pack_1.0.0_010907.xml`.

The next time you modify the management pack, you should change the versioning numbers according to the modification that you are making. Notice that in our earlier example, the versioning scheme uses three version numbers. The first number is known as the management pack version

number. The second is the major revision number, and the third is the minor revision number. The following shows a simple life cycle:



Management Pack Version When you first create a management pack, it should become your 1.0.0 version. If you create a fully revamped version of your management pack that has changed how the management pack functions, you should then change the management pack version. The updated management pack will become version 2.0.0 in your management group.

Major Revision Number Major revision numbers are updated when you make changes to the original structure of the management pack. For instance, if you add a new rule to the management pack, when you save the management pack you would increase the major revision number to 1. So if you have just obtained the management pack and have added in a new rule, the management pack becomes `management_pack_name_1.1.0_date`.

Minor Revision Number Minor revision numbers reflect changes to the settings in the management pack. These changes are not as drastic as adding or removing rules or other components; instead, they only alter how the management pack functions. Let’s say you have finished adding your new rule to the management pack and now you want to alter a few of the settings. You could enable a rule or disable an alert. After you have configured the settings, you should save your changes and name the new management pack as `management_pack_name_1.1.1_date`.

In Table 6.1, the versions of the management pack are shown as they progress through the versioning process. When the original management pack is obtained, version 1.0.0 is introduced into the test environment. The original vendor-supplied or developer-created management pack will probably not be introduced into production. Instead, it will be evaluated and modified to fit your environment. Chances are, the management pack will go through several versions before it is considered stable and is introduced into production. Also notice that we don’t change the version number when it is introduced into production. The test version and the production version are the same management pack, so we simplify the documentation process by not changing the version number.

TABLE 6.1: Version Numbers for Management Pack Revisions

OBTAIN	TEST ENVIRONMENT	PRODUCTION ENVIRONMENT
v 1.0.0	v 1.0.0	
	v 1.0.1	
	v 1.0.2	
	v 1.0.3	v 1.0.3
	v 1.1.0	
	v 1.1.1	

TABLE 6.1: Version Numbers for Management Pack Revisions (*CONTINUED*)

OBTAIN	TEST ENVIRONMENT	PRODUCTION ENVIRONMENT
	v.1.1.2	v 1.1.2
	v.2.0.0	
	v.2.0.1	

As you look at Table 6.1, note that the first set of modifications consists of configuration changes that are recorded as minor revision changes. As soon as the first rule addition or deletion or other modification to the structure of the management pack is made, the version changes to 1.1.0. As the management pack is deemed stable—in other words, it meets the criteria that the organization is looking for—it is released into production. Versions 1.0.3 and 1.1.2 reflect the introduction of the management pack into production. Version 2.0.0 is a completely new upgrade to the management pack that comes from either the vendor or the developer of the product as an unsealed management pack or is a complete revision of one of your own management packs.

When an updated version of the management pack is made available, you will need to perform additional testing. Usually, the new versions contain settings that had not appeared in the previous incarnation. You may have also introduced new settings to the previous version that may not be reflected by the new one. As you can probably imagine, you will have to test all the new settings that were introduced as well as determine if your custom settings will still work the same as they did under the previous management pack.

As we maintain our management packs, the version numbers will continue to increase, and we will be able to see the evolution of the management pack as we make changes to the structure, and then configure settings. Make sure that you are testing all of the changes, whether major changes, minor changes, or upgrades to new versions. Having virtualization tools such as Microsoft Virtual Server, Microsoft Virtual PC, VMware, and Xen available make testing much easier. You can set up your test environment and examine the effects of changing your management packs.

Having tested the effects of the management pack in a closed, isolated environment, it is time to take your management pack out into the real world. Before you release it and walk away, you should make sure that it is going to play nice with the rest of your systems. We don't always have all of the data we need when we start working in test environments. No matter how much we would like to think that we know everything that goes on, unknowns can rear their ugly heads at the most inopportune moments. Because of this, you should make sure the initial deployment is a limited deployment and that you control the results.

One of the nice things about Operations Manager is that you can control who will receive notifications. During your preproduction testing, test the management pack effects against your servers, but make sure that the results are not being sent to the service desk. Instead, initially configure the notification settings to report to your test group so that you do not adversely affect the service desk. This also allows you to identify whether the management pack causes any problems, or is too “noisy,” before directing the notifications to the service desk.

Management Packs Provided by Microsoft

If you take a look at management packs that Microsoft provides, you will see that they use a versioning scheme that uses the major.minor.build.revision numbering method. Due to their versioning

method, you cannot version changes to their management packs in the same manner. Of course, most of the management packs that they provide are sealed so you cannot make direct changes to them. One of the exceptions to the rule is the Default Management Pack. You can export this management pack and view the settings in it. As you make changes to the Default Management Pack, save it and increment the revision number. You can keep track of the versions by incrementing the numbers by one each time you make a change.

If Microsoft releases a new version of the management pack, they include a new, higher build number that supersedes the number you have provided as the revision. It will be up to you to make sure that the settings you have overridden are still performing the way you intended. You may find that you have to make some changes to get the overrides to work again. Be sure to test before putting the management pack into production.

For example, let's assume that you have just imported the Active Directory management pack. You have tested the management pack and decided that you want to disable some of the settings. After overriding the default settings and saving the changes in the Default Management Pack, you discover that Microsoft has issued a new version of the Active Directory Management Pack. Ideally, your first step would be to view the documentation that accompanies the management pack, and then compare the original management pack with the new. Sealed management packs provide a challenge, however, since you cannot easily view the information that is provided in them.

In such a case, you can compare your override criteria with the options in the new management pack to make sure the GUIDs are the same. If they are, chances are the override will still be valid. You must still test to make sure, but you can go into the test feeling a little better about how the management pack is going to work. Some reports are available when you install the reporting functionality. We discuss these in Chapter 9, "Reporting."

Documenting Changes

None of the versioning that you have been doing is worthwhile if you do not document your changes. Having good documentation is critical if you are planning to restore your management group back to a previous setting. You can always look at the version number that you have applied to the management pack, and maybe check the date to see when the change occurred, but without good documentation, you cannot be absolutely sure which management pack has the setting you would like to use.

This is where you need to set up standards for documentation and let everyone know what you expect to see in the change log. Some change logs are as simple as a binder with pages used to keep track of the changes. Once your change has been tested and saved, you pull out the binder and enter the changes that were enacted, the date, and the person responsible for the change.

If you want to get a little fancier with your change log and automate some the functions, which would include the ability to search for information, you can create your own database. Having a database store the changes that are made to your management pack can save you a lot of time when you are trying to determine when a change was made, which settings were modified, and which management pack you should use when returning to a previous version.

Using MP Studio 2007

Some applications are available to assist you with your management pack changes. One of the premier tools is Silect Software's MP Studio. This is a powerful tool that combines a change management feature with the ability to test and compare management packs. MP Studio is a vehicle that lets you incorporate the capabilities of Information Technology Service Management (ITSM) and Microsoft

Operations Framework (MOF) best practices and take charge of the life cycle of management packs. With its Microsoft Management Console (MMC)-style interface, the learning curve to use this tool is minimal. Under the hood of this tool is a sophisticated engine. You can make changes to a management pack and test to see what the effect will be on managed systems, as well as the amount of traffic that will be generated and the database activity that will occur. MP Studio comes in three versions tailored to fit any organization's needs:

MP Studio Enterprise Enables full life cycle support for Operations Manager 2007 management packs, including development, creation, testing, deployment, and maintenance to further help customers maximize the return on their Operations Manager investments.

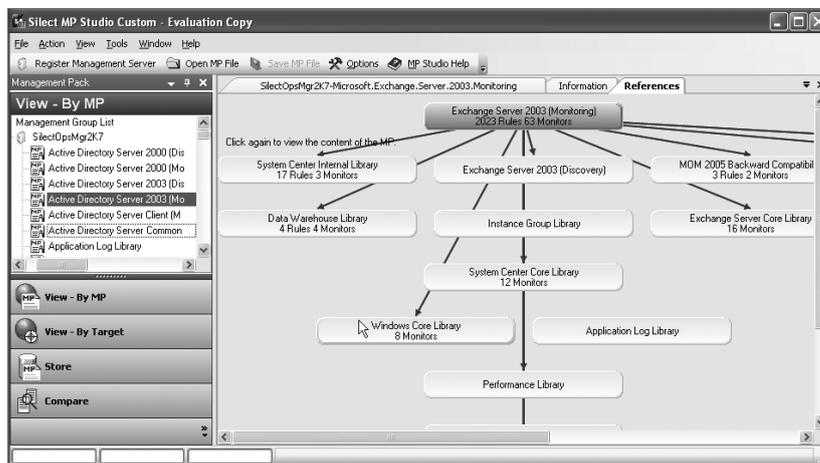
MP Studio Professional Provides management pack analysis, testing, reporting, editing, and tuning capabilities.

MP Studio Express Provides efficient reporting and analysis of Operations Manager 2007 management packs.

Analysis and Documentation

MP Studio gives you everything you need to fully analyze a management pack. This includes a graphical reference tree report showing the relationships between management packs, as shown in Figure 6.1. MP Studio's customizable reporting capabilities ensure that you can find and report on the information you need within a management pack, including the ability to share this information with team members. Reports can be in one of two formats: a report of the entire management pack in HTML or sections such as rules or overrides only, exported to Excel.

FIGURE 6.1
MP Studio 2007
console—
reference tree



Manage Overrides

Effective creation and deployment of overrides is critical to controlling alert volume and efficiently monitoring application service levels. MP Studio provides valuable override authoring, editing, and management capabilities in a highly controlled and accountable way. These capabilities include the ability to tune management packs prior to implementation, ensuring manageable alert volumes when deploying new management packs. MP Studio also gives you the ability to track who has

made changes to which overrides, thus providing a complete history of changes to your operational environment.

Management Pack Store

The management pack life cycle is a repeatable, structured process that helps organizations exploit the power of management packs. Whether you are a large organization with multiple management groups or an individual managing a smaller environment, having a process and the necessary tools to manage and track changes to management packs is critical to ensuring consistent service monitoring. MP Studio's MP Store feature provides critical change-management capabilities, including backup and version control. Full versions of management packs can be backed up to the store on a regular basis, and changes that have been made to the packs are automatically recorded and stored in an audit history report. You can sort by person, date, or other fields to get the information you need on who changed what and when.

Testing

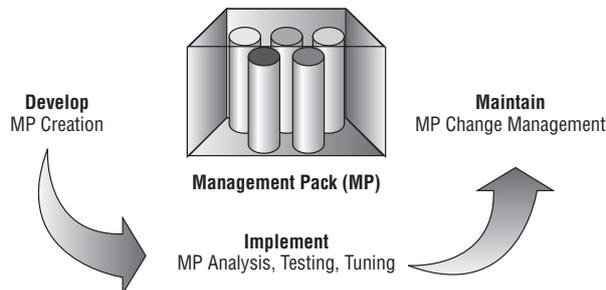
MP Studio lets you test the key features of management packs against live application servers or historical data such as event and performance logs without requiring you to install the management pack in your environment. MP Studio generates a complete report showing the Resultant Set of Alerts that would have been raised by the management pack.

Comparison

Ensuring consistency of service monitoring across different management groups involves detailed comparisons of management packs. This can be a tedious, time-consuming activity that is prone to errors. Alternatively, you may need to compare a tuned management pack against the original management pack to see what changes have been made. In either of these cases, MP Studio makes comparing management packs easy. MP Studio will show you all the details of the differences between the packs. Furthermore, comparisons can be scheduled to take place automatically and can include comparing any number of packs (installed in production or management pack files located in the management packs store or on the file system) to a standard management pack configuration.

Management Pack Life Cycle

Organizations using Operations Manager to deliver a service-monitoring function will want to utilize a common series of steps or processes when working with management packs. The aim is to provide a repeatable, efficient, and effective way of leveraging management packs, from initial planning and development to implementation and maintenance. The typical life cycle of a management pack broadly follows the stages shown here:



From initial creation to analysis and override management to ongoing change management, MP Studio works as the foundation to help you manage management packs through these stages. MP Studio delivers a set of capabilities that reduce the time and the effort required as a management pack passes through each stage of the life cycle. By leveraging these capabilities, IT groups can deliver a higher-quality management pack to the production environment in less time, resulting in better service monitoring and less alert noise.

A key enabling technology in MP Studio and the foundation for managing change throughout all stages of the management pack life cycle is the MP Store. The MP Store is a version-control system for management packs. It provides a place to store complete versioned copies of all management packs and provides a single point of reference for a management pack regardless of where it is in the life cycle. Version control, check-in/check-out, and a complete audit history are all available via the MP Store.

Management Pack Development

MP Studio's version-control system is a great place to start the management pack development process. The version-control environment provides a central shared location to store management packs, including all versions as they are being developed. This ensures team members are using the most current version as they add their expertise in the development process. If desired, users can use the check-out capability to allow only a single user to make changes to a management pack at any given time. When the management pack is checked back in, changes are automatically detected and the user is given an opportunity to comment on the changes that were made. This change history is maintained within the MP Store alongside the management pack itself, as shown in Figure 6.2, thus ensuring that the reasons for each change will never be lost.

As development of the management pack progresses and changes are made, MP Studio's testing capability allows developers to test their changes in an integrated test environment that does not require the management pack to be installed in the environment. This "sandbox" allows features to be tested as they are developed without having to wait for results while interacting with the production environment. Testing gathers data based on rules and monitors contained within the management pack. The user is presented with the results of testing, including detailed performance and event information as well as the Resultant Set of Alerts generated based on the data collected, as shown in Figure 6.3. The management pack can then be edited or overrides tuned using these results. After making changes to the management pack or overrides, you can quickly test it against the same test data to confirm that the alert volume has been reduced to an expected level.

FIGURE 6.2
Version control using
MP Studio's MP Store

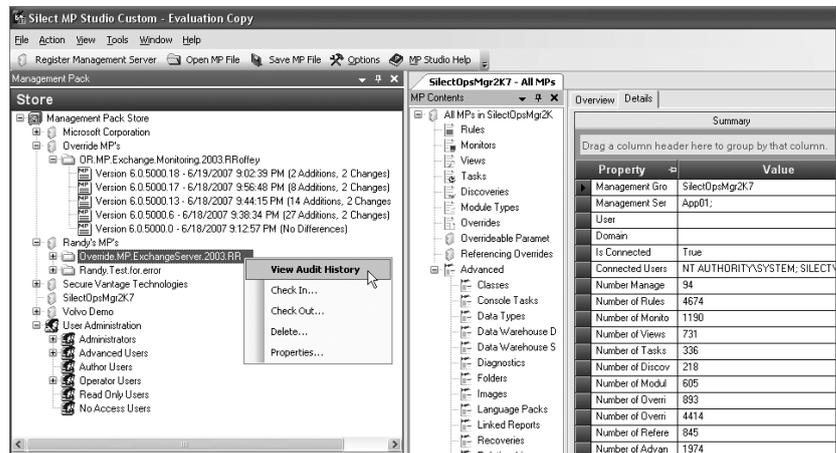
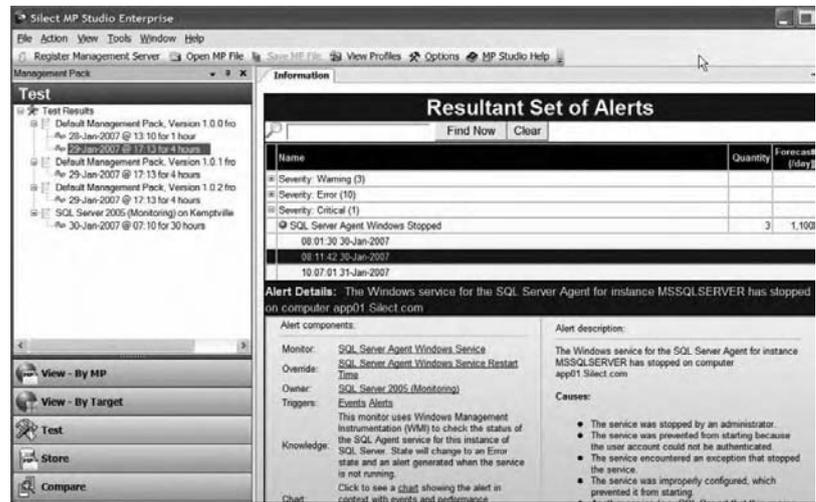


FIGURE 6.3
MP Studio: testing
and the Resultant Set
of Alerts



MP Studio's extensive documentation capabilities are also useful at this stage to help capture and share knowledge among team members as the management pack development process continues.

Management Pack Implementation

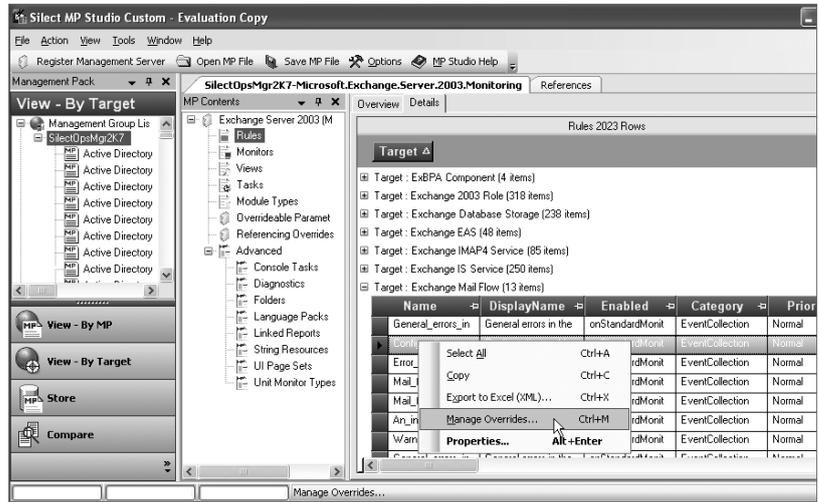
Prior to implementation, MP Studio's MP Store should be used to store the original released management pack. This can be used as a reference point for all future work with the management pack. Any referenced management packs that are sealed will automatically be imported into a special project folder, such as Microsoft or Secure Vantage.

Next, MP Studio's reporting and documentation features can be used to perform a detailed analysis of the management pack prior to implementation. A management pack file can be opened and viewed using MP Studio's customizable reporting user interface. You can group, sort, hide, and filter columns so that only the information that is important to you is displayed. Information from this customized view can be exported to XML or to Excel for further sharing, analysis, or review. MP Studio also includes a one-touch documentation feature that generates a complete management pack report within seconds. This report can have the header/footer and content customized, printed, and/or saved to a file for future reference.

MP Studio's testing feature allows effective testing of a management pack in an isolated environment. Testing gathers data based on rules and monitors contained within the management pack. The user is presented with the testing results, including detailed performance and event information as well as the Resultant Set of Alerts generated based on the data collected. A management pack can be quickly tested against any server in an environment, including a production server, regardless of whether or not it has an agent. Testing production servers will result in a much higher-quality management pack than if tuning took place in a lab environment. The Resultant Set of Alerts is useful for tuning the management pack before further lab or preproduction testing takes place.

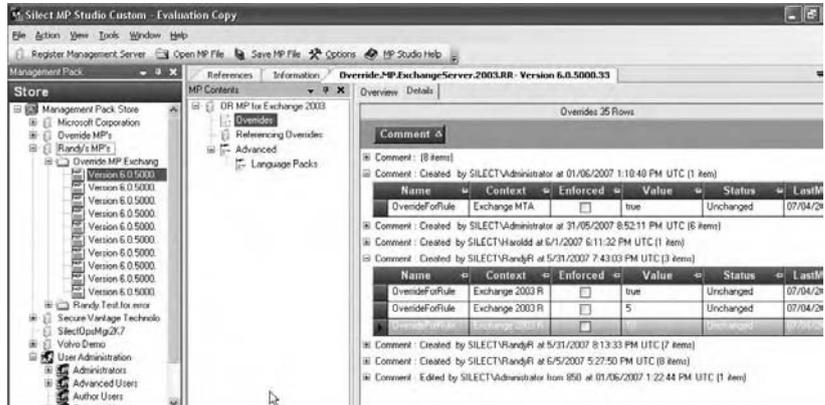
Tuning of the Operations Manager environment is easy using MP Studio. According to Microsoft best practices, override management packs should be created for each key management pack, such as Exchange, SQL, and Active Directory. You accomplish this by using the Manage Overrides option in MP Studio, as shown in Figure 6.4. Overrides can be created, edited, and then saved to the override management pack previously created.

FIGURE 6.4
MP Studio: Manage Overrides option



To address the all-important change management requirements, any override edits made using MP Studio can be tracked, thus providing a complete audit history of who changed what and when. This metadata will be saved with the management pack and can be viewed in MP Studio only using the Audit History report of the MP Store feature, shown in Figure 6.5, or by viewing the contents of the management pack in its file format using MP Studio.

FIGURE 6.5
MP Studio: auditing override editing



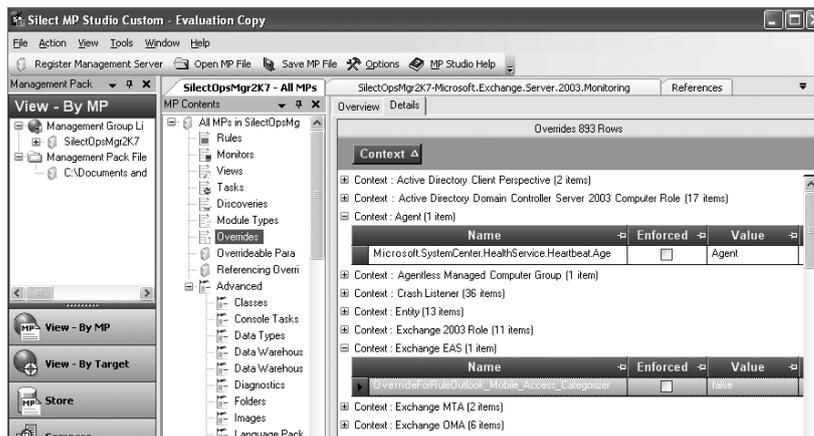
Management Pack Maintenance

MP Studio’s MP Store provides key maintenance capabilities, including management pack backup and version control. Full versions of management packs can be backed up to the store on a regular basis and changes that have been made are automatically recorded and stored in an audit history report.

MP Studio is also invaluable when upgrading management packs. MP Studio’s Compare feature can show all the differences between an existing pack and a newly released one, as shown in Figure 6.6. This makes it easy to drill down and quickly pinpoint deletions, additions, and changes within the new management pack. The Compare capability is flexible, and you can generate a comparison report in several ways. The report can include any combination of file-based or installed

MP Studio provides the ability to view the contents of the entire management group. The resulting view, shown in Figure 6.8, displays the resulting set of rules, monitors, views, tasks, parameters that can be overridden, and overrides for the group. This single rolled-up view allows focused analysis of the management group and individual components of the management packs.

FIGURE 6.8
Resultant Set of
Overrides for a
management group



Management Pack Considerations

You cannot modify sealed management packs. The only way to control how the management pack will run in your environment is to create overrides. As mentioned in Chapter 5, “Managing Management Packs,” overrides allow you to nullify a rule so that the agent on a single computer or group of computers will no longer process the rule. When this is the case, you will find that you are not actually making changes to the management pack where the rule resides; instead, you are setting an override property in the default management pack.

Taking that into consideration, you probably realize that you don’t have to worry about versioning sealed management packs. Well, that’s not completely true. Vendors will release updates to their management packs, and you will have to make sure that you document the Operations Manager systems that have had the new management packs imported. You will also need to test to determine whether the override is still functioning when the newly updated management pack is put into place.

Good documentation will help you determine which overrides have been put into place, so you should be able to quickly modify or reactivate and override. Exporting the default management pack and maintaining versioning controls should make it easier for you to troubleshoot issues that occur when the new management pack just doesn’t work the way you thought it should.

The Bottom Line

Document management pack changes. For every change that is made to a management pack, make sure you are keeping track of the change as well as the effect the change has. Some changes may simply inject data into the database, while others may automatically call upon a script to alter how the managed server functions.

Prepare the management pack for release. Keeping track of all the changes and their effects allows you to show the Release Board what the proposed changes will do when introduced into production. If you have performed all your testing and have documented everything well, you should have no problem getting approval for the release.

Control versioning of the management pack. While you are making changes to the management pack, you will need to have version controls in place so that you know what the last stable version of the product was. This allows you to go back to an earlier version of the management pack that you know worked the way you wanted it to in your environment. While there is no perfect method, any method that defines a change so that it can easily be identified will work.

Master It! Let's make some changes to a management pack and configure the version number as we make those changes.

1. Log on to OpsMgr1 as Administrator.
2. Open Windows Explorer and navigate to the C: drive.
3. Create a new folder named Management Packs and open the folder.
4. Create a folder named R&D Domain Controllers Management Pack and open the folder.
5. Create a folder named Base and open the folder.
6. Open the Operations Console and open the Administration view.
7. In the Administration view, select Management Packs and then right-click RD Domain Controllers Management Pack in the results pane.
8. In the Actions pane, click Export Management Pack.
9. In the Browse For Folder windows, click Local Disk > (C:) > Management Packs > R&D Domain Controllers Management Pack > Base.
10. In Windows Explorer, navigate to the Base directory under Default Management Pack.
11. Right-click the RD.Domain.Controllers.xml management pack file and select Rename.
12. Type **RD.Domain.Controllers.1.0.0.[date].xml** where *[date]* is the current date.
13. In the Operations Console, select the Authoring workspace and click Groups in the Authoring pane.
14. Right-click RD Domain Controllers and select Properties.
15. Select the Explicit Members tab and click the Add/Remove Objects button.
16. Click Search and select the other two domain controller objects from the list; then click Add and OK.
17. Click OK.
18. Select the Administration workspace and click Management Packs in the Administration pane.
19. Right-click RD Domain Controllers and select Properties.
20. In the Version text box, enter **1.0.1**.

21. In the Description box, enter **Added additional domain controller objects** and click OK.
22. Right-click RD Domain Controllers Management Pack in the results pane.
23. In the Actions pane, click Export Management Pack.
24. In the Browse For Folder window, click Local Disk > (C:) > Management Packs > R&D Domain Controllers Management Pack.
25. Click the New Folder button, name the folder **1.0.1**, and click OK.
26. In Windows Explorer, navigate to the Base directory under Default Management Pack.
27. Right-click the `RD.Domain.Controllers.xml` management pack file and select Rename.
28. Type **RD.Domain.Controllers.1.0.1.[date].xml**, where *[date]* is the current date.

Chapter 7

Monitoring with the Console

The Operations Console is where your operators interact with all the monitored systems. It is also the tool administrators and authors use to design the objects that the operators use when monitoring the organization's services and servers. Several options are available that allow you to fine-tune access to components in your organization. You may want to define the users who can monitor individual components, which means you will need to create additional Operations Manager components to control access to only the views they need to work with.

You may also want to create your own customized views of the organization's objects instead of using only those that are made available from the imported management packs. Within the views, you may also want to create specialized tasks that run against the objects.

In this chapter you will:

- ◆ Identify the workspaces and views that are available from the console
- ◆ Learn how to customize a view
- ◆ Create views and limit the operator's view
- ◆ Learn what tasks are and how to create and use them

Console Overview

If you have worked with the previous versions of Microsoft's Operations Manager, you know that there were two consoles: the administrator's console and the operator's console. The two consoles allowed for a separation of responsibilities. You could allow the MOM administrative staff within your organization to have access to the administrator's console so that they could configure the settings necessary for monitoring your environment. Anyone responsible for monitoring systems could be given access to the operator's console so that they could see what was occurring on those systems.

One of the nice things about the operator's console was the ability to present different views of your monitoring solution. Not every operator needed to see all of the monitored systems in the organization. More often than not, they needed only a small subset of the data that was being received by MOM. New views could be created that would present the data in a concise manner, and then permissions could be applied so that the operators had access to only their own views.

Operations Manager 2007 moves away from the two-console scenario and exposes everything in a single console. Now called the Operations Console, this single repository allows administrators to view all aspects of the monitoring solution from a single console, and at the same time gives them granular control over what operators are allowed to see.

Workspaces

When you open the Operations Console, you will immediately notice the four workspaces that appear in the lower-left side of the console, as shown in Figure 7.1. If you have installed Reporting, you will see a fifth workspace, appropriately named Reporting. Each of these workspaces is protected by access control lists that allow you to create roles for users and then assign those roles so that users have access to specific workspaces. Each of the workspaces allows you to work with different aspects of Operations Manager:

Monitoring The Monitoring workspace allows you to view the alerts and health status of services and servers that are under the control of the management servers.

Authoring Using the Authoring workspace, you can create new rules, alerts, and monitors to complement the management packs used by the management servers.

Reporting The Reporting workspace exposes reports that are available to the operator and allows reports to be customized. Until you install the Reporting component, this workspace does not appear in the list of workspaces.

Administration The Administration workspace is used to manage the configuration settings for the entire management group. This is where you configure the settings that control the management servers and the agents, as well as import and control the management packs.

My Workspace You use My Workspace to create a customized console view of the monitored services, servers, and alerts that you are responsible for.

FIGURE 7.1

The workspaces available within the Operations Console

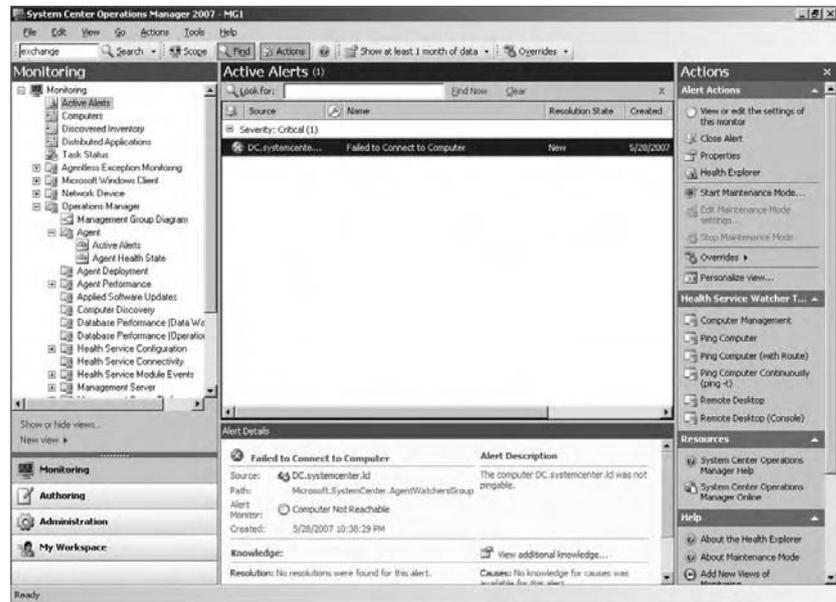


When you click one of the workspace buttons, the Navigation pane above it changes to display the views that are available in the workspace. The views can be controlled by the scope of views that are available to the operator. Using a scope is simply modifying the data returned by the query so that only the information you want to view is visible. Using scopes allows you to display only the information you want to see, or to hide information that another operator does not need to see. If an operator does not have permissions to a view, it simply does not appear in the Navigation pane.

To the right of the Navigation pane is the Display pane. The Display pane contents reflect the view you have selected in the Navigation pane. For instance, if you select the Active Alerts view, all of the alerts that have not been resolved appear, as shown in Figure 7.2. The Display pane is divided into two sections: the top section shows the alerts and the lower presents details of an alert when you select it from the upper section.

Finally, the extreme-right section, the Actions pane, contains actions that can be performed by an operator. The actions that appear in this pane are context-sensitive, so what is selected in the upper section of the display pane dictates what is available from the Actions pane. In the next few sections, we discuss what is available from the different views.

FIGURE 7.2
Unresolved alerts
from the Active
Alerts view



Views

In reality, a *view* is nothing more than a response to a query. When you define a view, you are creating a query that is sent to the database, and then the results of the query are presented to the operator. Several views are already defined in management packs, but you can create your own views, either for your own use or to give to operators who will perform a specific function. Typically, most operators do not need to see all the information available through a default view.

As you import management packs, the views that have already been created by the author of the management pack become available in the Monitoring workspace. These are views that the developer of the management pack feels an operator would need to use to monitor the alerts and health of their product. Most of the management packs that you import will be sealed management packs and impossible to modify. If you want to show different views than those provided by the author, you have to create a view in a different management pack.

In Figure 7.3, you see the views that are available by default when you install Operations Manager 2007. These views, especially the Active Alerts view, allow you to start monitoring right out of the box.

FIGURE 7.3
Default views
from Operations
Manager 2007



Active Alerts The Active Alerts view presents an operator with all of the current alerts that have yet to be resolved. This becomes a one-stop shopping place to see the problems that need to be addressed. Since all of the alerts appear in this view, it could become difficult to determine exactly what is happening in your environment and which systems are being affected. However, if you want to see what problems still need to be resolved, and how long each alert has been active, this is one view that becomes indispensable.

Computers In the Computers view you will find all the systems that are members of the management group. Using this view, you can determine which systems are monitored by the management servers in the management group and the state each system is in. This view also helps you determine which systems are not managed, giving you a place to start identifying which systems may not have the agent installed correctly. You can also use this view as a starting point for monitoring each system since you can access other views from the context menu or Actions pane.

Discovered Inventory All of the objects that are discovered by the discovery rules contained in each imported management pack can be found in the Discovered Inventory view. It is here that you will find all the individual monitored objects and their current state. As with the Computers view, you can select an object and then switch to another view by selecting the view on the context menu or the Actions pane.

Distributed Applications Because Operations Manager monitors on an end-to-end basis, you can create distributed application views that allow you to monitor all the components that tie an application together. Here you can see an overview of the health of the distributed application. You can then navigate to other views via the context menu or Actions pane.

Task Status Not all of the tasks that you initiate will return a response directly to the console screen, and for those that do, you don't always want to keep the results onscreen—you will probably want to move to another view to continue with your troubleshooting. The Task Status view allows you to revisit the results of tasks that you have run in the past, or view those that do not return results directly to the screen.

Although most manufacturers are cognizant of the monitoring that is required for their products, they are not aware of the other products you have in your organization. The views that are delivered through their management packs represent only their products, but your operators may need to monitor a wide range of products and services. Having the ability to create views that represent their monitoring needs comes in very handy.

CREATING YOUR OWN VIEWS

Different types of views grant you different aspects of the collected information and operator actions. Understanding what is available from each view helps you when it comes time to create a customized view of the information.

Alert View Alert views display the alerts that are generated based on the rules and monitors that are defined in the management packs. The Active Alerts view, mentioned in the default views, is an alert view that aggregates all unresolved alerts in the management group. You can create additional alert views that meet your needs. You can scope the alert view to display alerts specific to a service, severity, resolution state, or any combination thereof.

Diagram View The diagram view presents a graphical representation of managed objects and their relationships to one another. Using a diagram view, you can monitor solutions, such as a

group of SMS servers and the supported clients. The diagram view that appears shows the high-level state of each system.

Event View Events that are collected can be viewed from an event view. As with other views, you can display all of the events, which can be quite overwhelming, or you can scope the view to display specific types of events so that you can target your troubleshooting efforts.

Performance View Performance views can be used to display performance counter data that has been collected from the monitored systems. The view can be controlled so that only performance data from a specific system is displayed, or you can display the same performance object and counter for the entire set of monitored objects in the management group.

State View The state view presents the current status of a monitored object. Because objects can include a group of servers and underlying services, the state view shows the service's state. If an object goes into a warning or error state, you can select the object and choose another view, which can help you troubleshoot the problem.

Task Status View The task status view displays the results of tasks that have been run so that you don't have to rerun any tasks or document the results while you are troubleshooting problems.

Web Page View The web page view can be used to launch a new web browser window in the Operations Console in order to display the information you are viewing.

Dashboard View This view does not specify the type of data that will appear in the view; instead, it is a placeholder for other views so that you can look at multiple views simultaneously.

Also, knowing what types of views are available and their settings can come in handy. Sealed management packs may include views that you can use when monitoring, but these views cannot be modified to meet your specific needs. If you are going to modify any view from a sealed management pack, you will have to save the customized view in your My Workspace workspace using the Favorite Views node. It is up to you whether you will use the Default Management Pack or a custom management pack.

Creating an Alert View

Follow these steps to create an alert view:

1. In the Monitoring workspace, right-click an unsealed management pack and choose New ➤ Alert View.
2. In the Name field, type a name to identify the view.
3. On the Criteria tab, define the alert view:
 - A. On the Show Data Related To box, click the ellipsis and select the alert data type to view.
 - B. In the Show Data Contained In A Specific Group box, select the group that you want to filter the data.
 - C. In the Rule Entries area, select the check boxes for the alert criteria that you wish to show in the view.
 - D. In the Criteria Description box, click the underlined words to define the alert criteria that you wish to show in the view.

4. On the Display tab, define how the data is displayed:
 - A. Choose the columns to display by selecting the check box for each column.
 - B. Move the column into the appropriate position in the view by using the up and down arrows.
 - C. Choose whether to sort the columns by Age, Created, Maintenance Mode, Name, Resolution State, Severity, or Source by selecting the appropriate option from the Sort Columns By pull-down.
 - D. Select whether the sorting should be Ascending or Descending.
 - E. Select the grouping of the alert items by using the Group Items By pull-downs to create three levels of reporting.
 - F. Select whether each grouping should be Ascending or Descending.

Creating a Diagram View

Follow these steps to create a diagram view:

1. In the Monitoring workspace, right-click an unsealed management pack and New ➤ Diagram View.
2. In the Name field, type a name to identify the view.
3. Click Browse next to Choose Target and select the object type that you want to add to the diagram view.
4. If you have already defined a template, you can select the template from the pull-down under Choose From A Template.
5. If you want to create a new template, select Create Your Own Template and then define the template:
 - A. In the Diagram Properties tab, select how many levels deep the diagram will show when initially opened and the layout direction of the objects in the diagram.
 - B. On the Object Properties tab, select whether you want the diagram to be contained in a box or not, then specify how many nodes will occupy each line of the box.
 - C. On the Line Properties tab, select the line color for each type of object, then specify the style and thickness of the line.

Creating an Event View

Follow these steps to create an event view:

1. In the Monitoring workspace, right-click an unsealed management pack and choose New ➤ Event View.
2. In the Name field, type a name to identify the view.

3. On the Criteria tab, define the event view:
 - A. On the Show Data Related To box, click the ellipsis and select the event data type to view.
 - B. In the Show Data Contained In A Specific Group box, select the group that you want to filter the data.
 - C. In the Rule Entries area, select the check boxes for the event criteria that you wish to show in the view.
 - D. In the Criteria Description box, click the underlined words to define the event criteria that you wish to show in the view.
4. On the Display tab, define how the data is displayed:
 - A. Choose the columns to display by selecting the check box for each column.
 - B. Move the column into the appropriate position in the view by using the up and down arrows.
 - C. Choose whether to sort the columns by Date And Time, Event Number, Level, Name, or Source by selecting the appropriate option from the Sort Columns By pull-down.
 - D. Select whether the sorting should be Ascending or Descending.
 - E. Select the grouping of the event items by using the Group Items By pull-downs to create three levels of reporting.
 - F. Select whether each grouping should be Ascending or Descending.

Creating a Performance View

Follow these steps to create a performance view:

1. In the Monitoring workspace, right-click an unsealed management pack and choose New > Performance View.
2. In the Name field, type a name to identify the view.
3. On the Criteria tab, define the performance view:
 - A. On the Show Data Related To box, click the ellipsis and select the performance data type to view.
 - B. In the Show Data Contained In A Specific Group box, select the group that you want to filter the data.
 - C. In the Rule Entries area, select the check boxes for the performance criteria that you wish to show in the view.
 - D. In the Criteria Description box, click the underlined words to define the performance criteria that you wish to show in the view.

4. On the Display tab, define how the data is displayed:
 - A. In the Date And Time area, select whether to show all data, or only the data for a specific time frame.
 - B. In the Chart area, select to format the data as a line or spline graph, and then provide further formatting by selecting the Enable 3D and Point Labels check boxes if appropriate for your view.
 - C. The X Axis and Y Axis areas allow you to further define how the view will be formatted by allowing you to show the axis details as well as gridlines and colors to define the gradations.
 - D. Click the Change buttons on each of the axis areas to define the color for each axis.

Creating a State View

Follow these steps to create a state view:

1. In the Monitoring workspace, right-click an unsealed management pack and choose New ➤ State View.
2. In the Name field, type a name to identify the view.
3. On the Criteria tab, define the state view:
 - A. On the Show Data Related To box, click the ellipsis and select the state data type to view.
 - B. In the Show Data Contained In A Specific Group box, select the group that you want to filter the data.
4. On the Display tab, define how the data is displayed:
 - A. Choose the columns to display by selecting the check box for each column.
 - B. Move the column into the appropriate position in the view by using the up and down arrows.
 - C. Choose whether to sort the columns by Maintenance Mode, Name, Path, or State by selecting the appropriate option from the Sort Columns By pull-down.
 - D. Select whether the sorting should be Ascending or Descending.

Creating a Task Status View

Follow these steps to create a task status view:

1. In the Monitoring workspace, right-click an unsealed management pack and choose New ➤ Task Status View.
2. In the Name field, type a name to identify the view.

3. On the Criteria tab, define the task status view:
 - A. On the Show Data Related To box, click the ellipsis and select the task status type to view.
 - B. In the Show Data Contained In A Specific Group box, select the group that you want to filter the data.
 - C. In the Rule Entries area, select the check boxes for the task status criteria that you wish to show in the view.
 - D. In the Criteria Description box, click the underlined words to define the task status criteria that you wish to show in the view.
4. On the Display tab, define how the data is displayed:
 - A. Choose the columns to display by selecting the check box for each column.
 - B. Move the column into the appropriate position in the view by using the up and down arrows.
 - C. Choose whether to sort the columns by Run Location or Schedule Time by selecting the appropriate option from the Sort Columns By pull-down.
 - D. Select whether the sorting should be Ascending or Descending.
 - E. Select the grouping of the alert items by using the Group Items By pull-downs to create up to two levels of reporting.
 - F. Select whether each grouping should be Ascending or Descending.

Creating a Web Page View

Follow these steps to create a web page view:

1. In the Monitoring workspace, right-click an unsealed management pack and choose New ► Web Page View.
2. In the Name field, type a name to identify the view.
3. Enter the website address in the Target Website text box.

Creating a Dashboard View

Follow these steps to create a dashboard view:

1. In the Monitoring workspace, right-click an unsealed management pack and choose New ► Dashboard View.
2. In the Name field, type a name to identify the view.
3. Select the number of views to show in the dashboard from the pull-down.
4. Select the layout of the dashboard from the Select The Layout Template area and click OK.
5. When the dashboard appears, click each of the Click To Add A View links to select the view that will appear in the dashboard.

Tasks

As with views, there are tasks that are available when Operations Manager is installed, tasks that are imported along with management packs, and tasks that you can create on your own. Since there are so many tasks available from the default and imported management packs, you should verify whether a task that can perform the action you need already exists. Of course, there are those tasks that are part of sealed management packs that you do not have access to outside of the management pack. If this is the case, you may very well have to duplicate the functionality of the task.

Tasks come in two flavors: command-line tasks and scripts. Command-line tasks allow you to create a single command that is run on either the management server or the agent. Scripts can be written and then targeted to the management server or agent in order to run on them. With either of these approaches, the operator can select a system to target, and then run the task manually from the Operations Console. These manually run tasks enable the operators to run diagnostic tasks to help determine what is wrong, or run recovery tasks to help solve the problem causing the alert.

Commands can also be run automatically when certain events or alerts arise. Since developers of applications and services know how their products function, they also know how to repair some of the problems that arise from time to time as the product is functioning. Some recovery solutions that they have identified can be incorporated into their management packs. When an event or alert occurs, the command line can run automatically to alleviate the problem or perform an action that will further help the operator identify the problem.

Tasks can be as simple as issuing the `ipconfig` command against a monitored system in order to find out what its IP settings are, or as involved as running a complex script that allows you to modify the settings in a service-level agreement database and then notifying an operator that problems exist with a service via pager or cell phone text message. Task capabilities are limited only by the functionality of the services and systems in your organization.

If the task that you want to use is part of an unsealed management pack, you can make changes to the task to fit your environment. It probably goes without saying that tasks in sealed management packs are off limits. To create new tasks, you use the Create Task Wizard. To access this wizard, select the Authoring workspace and navigate to the Management Pack Objects node. Once the Management Pack Objects node is open, you can manage tasks by selecting the Tasks node. When you right-click Tasks, select Create A New Task to start the Create Task Wizard.

Two types of tasks exist: command line or script-based. You can configure the command-line tasks to be initiated at the Operations Console and run on the client, or you can configure them to run by the agent whenever a specific alert or event occurs. As you can see in Figure 7.4, the two options, Agent Tasks and Console Tasks, are defined and displayed so that when you are creating the task, you can figure out where the command or script will be run. Also take note of the Management Pack option at the bottom of the page. You can use the pull-down to select an existing management pack to store the task, or you can create a new management pack by clicking New and entering the identifying information for the new management pack.

AGENT TASKS

Agent tasks are initiated by the agent that is installed on the management server or managed system. When these tasks are exposed in the Operations Console, running them will cause the agent to run the command or script. As the task is executing, you will not have control over the task. Instead, to make sure that the task does not run out of control, you have the ability to specify how long the task is allowed to run.

When you run the Create Task Wizard, you are presented with pages and prompts that allow you to define exactly what the task will do. The first page you'll see is the General Properties page.

The information that can be entered on this page helps you identify the task when you are getting ready to run or modify the task. The Task Name and Description (Optional) text boxes are self-explanatory, but at the bottom of the General Properties page for Command Line or Run A Script tasks there is a Task Target option, as shown in Figure 7.5.

FIGURE 7.4
Tasks types

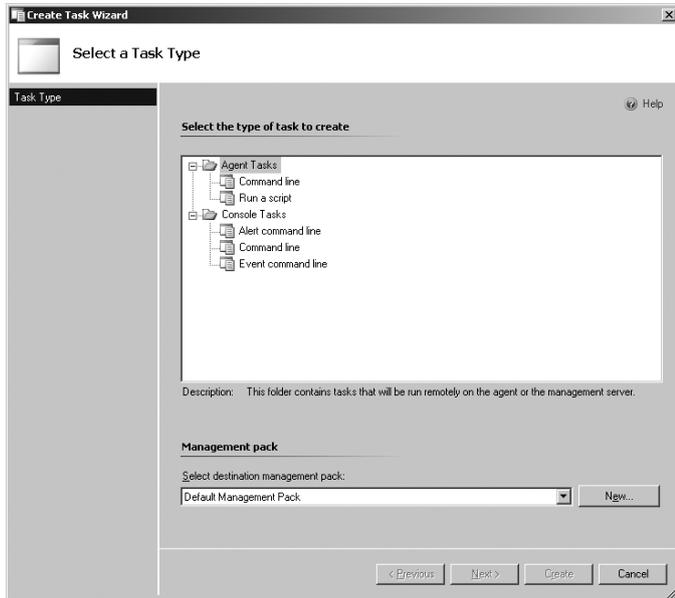
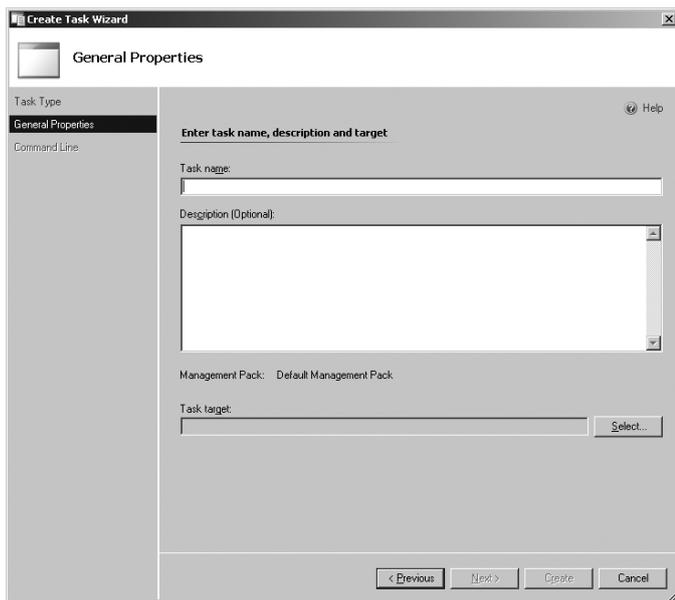
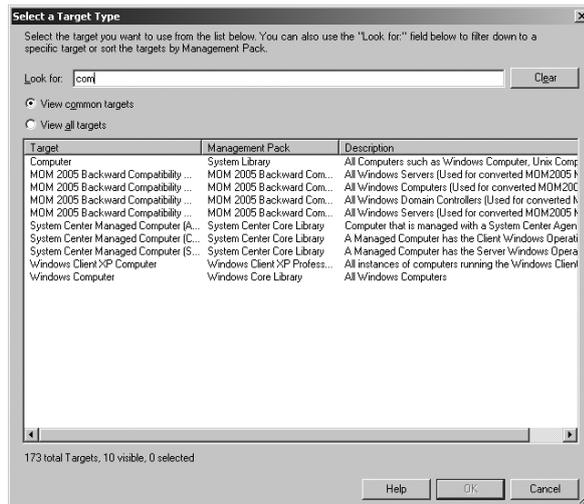


FIGURE 7.5
General Properties of
a task



When you click the Select button, you are presented with a list of target types that exist for the management group. This list will be different based on the management packs that exist in the management group. As you are creating a task, you should know what the task is going to affect, so selecting the target should not be a hard choice; however, the list can seem daunting at times. You can filter the targets that appear in the list by typing in a word or part of a word. Those entries that don't include the typed information in the Look For text box will be removed from the list. Figure 7.6 shows the list after we entered **com** in the Look For text box.

FIGURE 7.6
Filtering targets

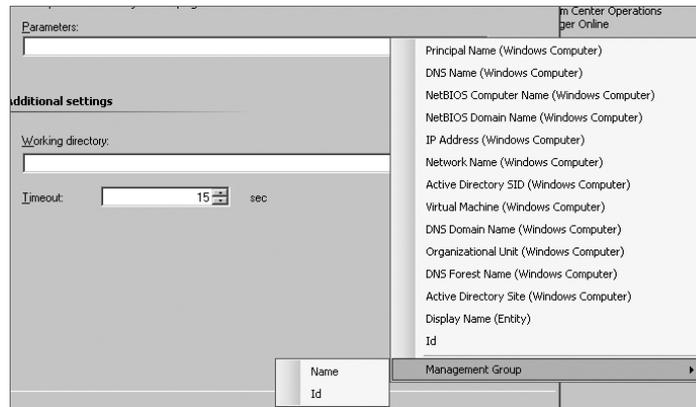


After you click Next, you are presented with the Configure Command Line Execution Settings page if you are creating a command-line task, or the Script page if you are creating a script task. This is where you define the actual command line or script that will be executed as well as how it will be executed. For command-line tasks, the first option, the Full Path To File text box, allows you to enter the path to the file. You have two options when entering the path: you can use an absolute value, such as `c:\winnt\system32\ping.exe`, or you can use environment variables such as `%systemroot%\system32\ping.exe`. Using variables is the preferred method because there may be monitored systems that do not use the standards set forth by your company. If one system has the system root defined as `c:\winnt`, another as `c:\windows`, and a third as `d:\windows`, using an absolute value will not execute on two of the three systems.

The second text box, Parameters, allows you to define the parameters that will be passed to the file that you defined earlier. Again, this can be an absolute value that you type into the text box, or you can supply data from the parameters selection. To do so, click the arrow beside the Parameters text box and select the parameter from the list, as shown in Figure 7.7. For our ping example, we are selecting DNS Name (Windows Computer). This will allow us to check name resolution when performing a ping against a selected computer in the Operations Console. We could also create a second task that pings the IP address of the computer by selecting the IP Address (Windows Computer) option.

At the bottom of the Configure Command Line Execution Settings page is the Working Directory and Timeout options. If there is a directory that you want to execute the task, you can enter it in the Working Directory text box. As for the Timeout option, you can define how long a task is allowed to run before it will be forcibly ended. Entering a value other than 0 will define how many seconds the task can run, effectively giving you a way to end a program or script that is hung or is executing longer than allowed.

FIGURE 7.7
Parameters area
of a task



If you are creating a script task, the Script page allows you to name the script, set a timeout value, and enter the script. Notice that the script task is meant for a script that is written specifically for the task. You don't have the ability to specify a path to an existing script. If you do have an existing script, you could use the command-line task to call it. When entering the script, if you don't like to use the small text box that is available from the Script page, you can click the handy Edit In Full Screen button to give yourself some more room to type. When you have finished with the script text page, simply click Save And Close to return to the Script page.

CONSOLE TASKS

Moving down the tasks types, we enter the realm of console tasks. Although these are very similar to the agent tasks, console tasks are run against objects in the Operations Console and are initiated by the system that is running the console. Whenever you select an object from the Details pane that meets the target requirements for a task, the task becomes visible in the Actions pane.

There are not many differences between the agent and console task settings, so we will detail just the differences between the two. In Figure 7.8, you see that there is an option to specify whether the output is displayed when the task is run instead of a timeout option. If you decide that you do not need to have immediate feedback from a task, you can deselect this check box and then view the results from the Task Status node.

ANNOYANCES

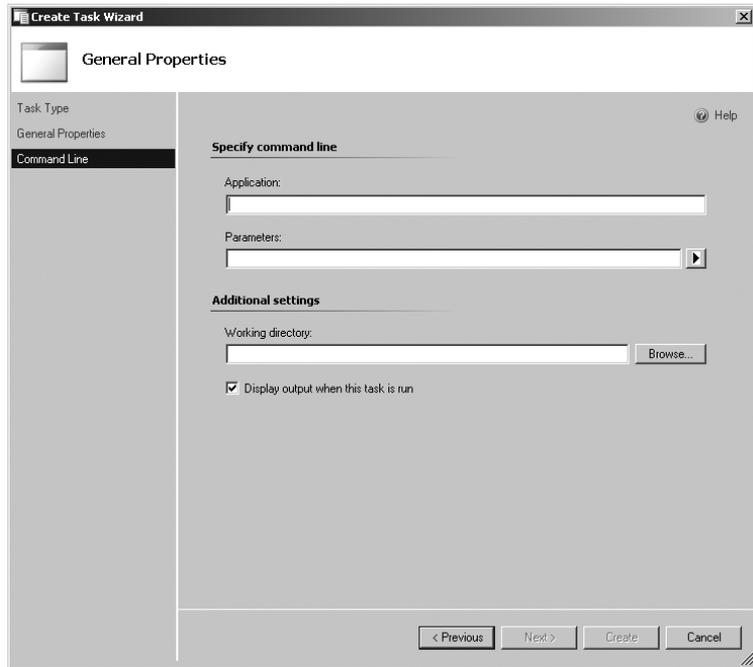
There are a few annoyances that you may find in Operations Manager 2007. For instance, if you look at Figure 7.8, you will see that the Command Line page of the Create Task Wizard is labeled as General Properties. This is apparent only if you are creating console tasks—the agent tasks specify the correct labeling.

Another annoyance is the grammar found on a few of the pages. If you look at the Description field when you are selecting a target, you will find entries that should have probably been caught by a grammar check.

If you are creating a script task, you will find that the documentation specifies that you can identify whether the script is VBScript or JScript, but that functionality is not exposed in the product. Instead, you can enter only VBScript into the Script text box.

While none of these annoyances affect the functionality of the product, it leads one to think that Operations Manager is not a complete or mature product. These issues will probably be rectified when the first service pack is issued, but having them exposed in a final product is a little disconcerting.

FIGURE 7.8
Console task
Command Line page



Operators

Not every operator in the management group will need the ability to work with all the objects that are available. Not only will you probably want to divide up the responsibilities of your staff, but you will also want to limit the information they can view. For this reason, you can create operator types, or user roles, and then add accounts that are members of these roles.

Five user roles are available when you install operations manager: administrators, advanced operator, authors, operators, and read-only operators. These roles cannot be modified or deleted.

Operations Manager Administrators This role has full access to all aspects of the management group.

Operations Manager Advanced Operators This role has the ability to override rules and monitors in the management group.

Operations Manager Authors This role has the ability to create tasks, rules, monitors, and views in unsealed management packs in the management group.

Operations Manager Operators This role has the ability to interact and manage alerts, tasks, and views in the management group.

Operations Manager Read-Only Operators This role has the ability to view alerts and views in the management group.

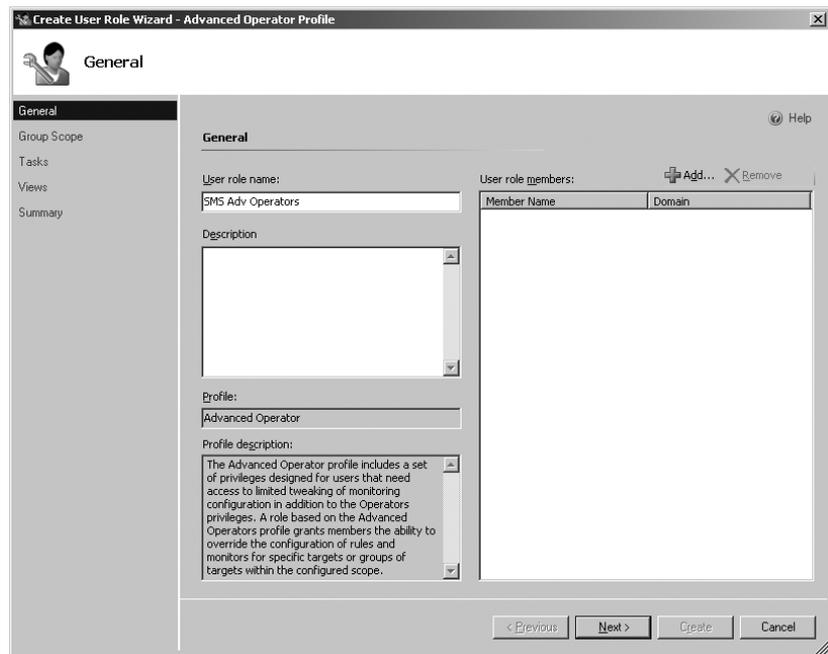
You can create your own user role by right-clicking User Roles in the Administration workspace and selecting New User Role, then selecting the role type. The only role type that is not visible is

administrators. You can select any of the other role types and then limit the access the group has to information. If you use the default roles, you are granting access to every group that resides in the management group. In most cases this is far too much power to hand out; instead, you should create a new role and limit its scope of power.

Each of the roles has abilities and options available to it that ensure it the right to perform actions on objects in the management group. For instance, the author role has the right to author new objects, but the operator role does not. The operator role is allowed to manipulate objects, but the read-only operator is limited to viewing the data. As you create a new role, use the appropriate role type so that you do not grant too much power to a member of the role.

When the Create User Role Wizard starts, you are presented with the General page, as shown in Figure 7.9. Here you can name the role and add members, but it is also a good place to double-check that the role is appropriate for the level of control you want to grant the members. The Profile Description section gives you an idea of what the role can do, and the page links on the left side tell you what rights are granted. The example in Figure 7.9 is an advanced operator role; note that the Author option is not available. If this were an author role, you would see the Author page listed on the left side of the page.

FIGURE 7.9
Create User Role
Wizard, General page



The Author Scope page allows you to define the object types that can be authored by this role. The default is to allow the role to author any and all targets, including those that will be imported at a later time. Selecting the Only Targets Explicitly Added To The Approved Targets Grid Are Approved check box will allow you to limit the objects that can be authored. In Figure 7.10, you see a sample of the targets that are available. Not all of the targets are listed here; to see the entire list, you need to select the View All Targets radio button. You can filter the list by entering part of the target name in the Look For text box.

FIGURE 7.10
 Authoring scope targets

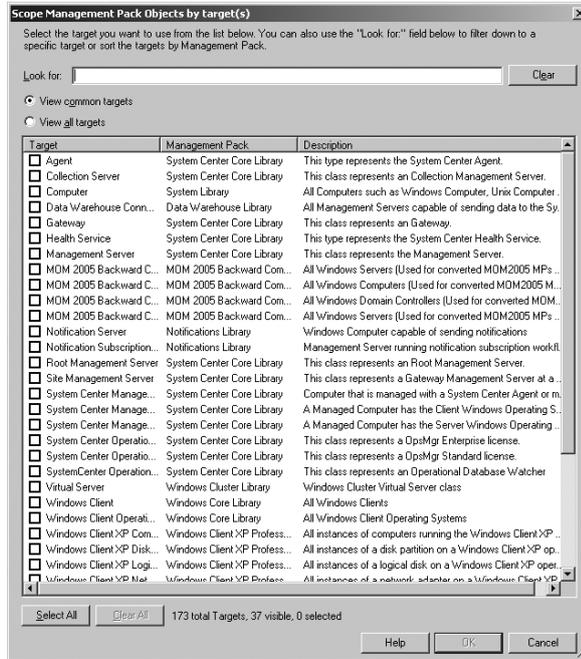
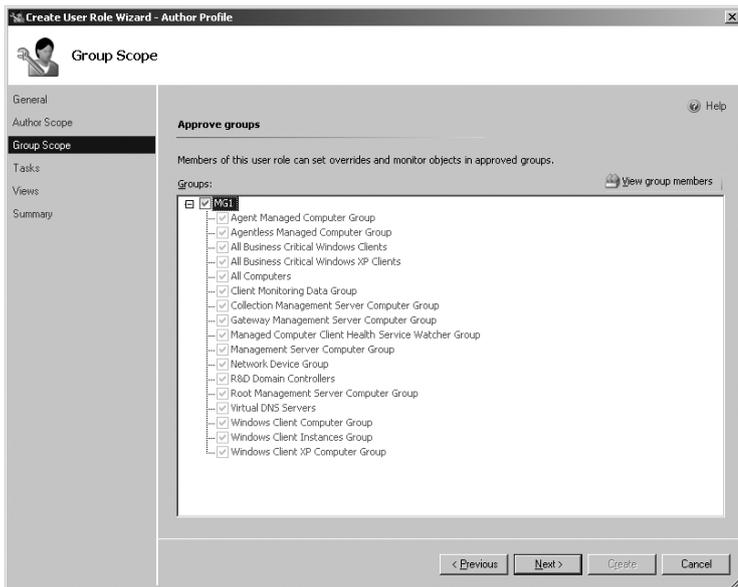


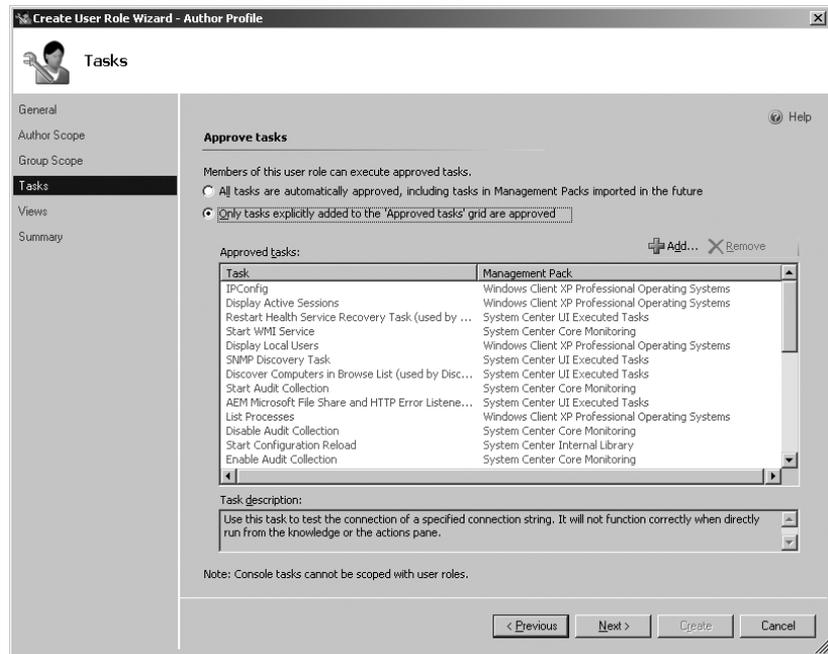
Figure 7.11 shows the groups that are available for the role to manage. By default, all of the groups are available. If you want to limit the groups that the role has access to, deselect the check box by the management group name and select only the groups that you want to allow the role to access.

FIGURE 7.11
 Group scope targets



There are several tasks that you cannot limit a user to. As you can see in Figure 7.12, the tasks in the grid are grayed out and cannot be removed. You can add additional tasks to the grid by selecting Add and then clicking the check box next to each task you want to add. Once added, these tasks can be removed from the list.

FIGURE 7.12
Approving tasks for the operator



The Views page, shown in Figure 7.13, lets you define the views that the user will have access to. By default, all views are available. If you choose the Only The Views Selected Below Are Approved option, the views that are required due to the Author and Group scopes you chose are selected. If you grant access to any of the other views, you will be granting read-only access to those views.

As we mentioned earlier, each of the user roles can perform specific actions. Naturally, some of the actions that can be performed by each of the roles will overlap with those from another role. Simply choose which role best fits the accounts that will need to work with objects in your management group.

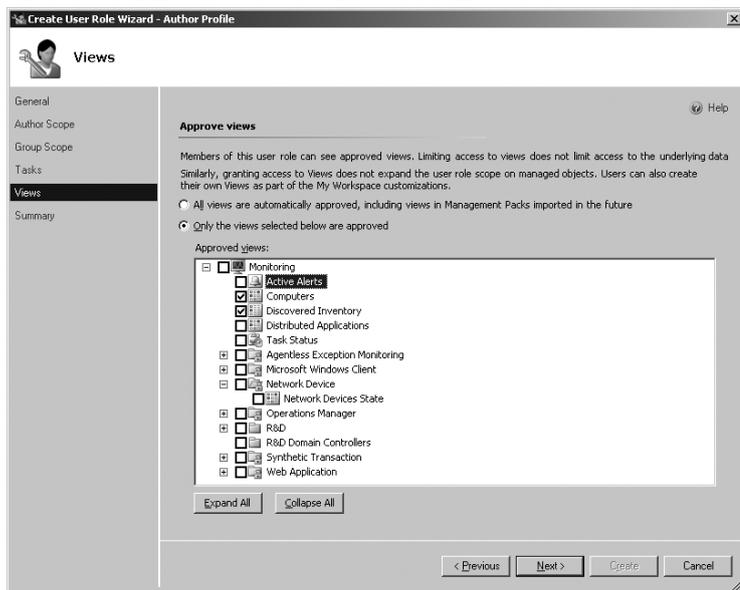
Members of the administrators user role have the ability to manage all aspects of the management group, which includes deploying agents, creating objects, importing management packs, configuring and managing connectors, and managing alerts and health states. They can do it all, which means you need to monitor the membership of this user role. Add only those individuals you trust and that require a high level of control in the management group.

All the other user roles can be scoped to allow the members access to only the groups and views that are necessary. The four built-in roles that exist when the management group was created should be used sparingly. Unless your organization is small, you have probably already divided up responsibilities among your administrative staff. Take the time to review those responsibilities and then create views that allow access to the appropriate objects for the user groups. Once you have the views defined, you can create user roles and assign them to have access to the views. Create a user role that uses the read-only operators role and add the users who will only need to view the

alerts and monitor states. Since read-only operators cannot run any tasks, you will then need to create a user role that uses the operator role. These users can respond to alerts and monitor states as well as run tasks. Finally, you should create a user role that has the advanced operator role only if you have users who are responsible for creating and maintaining overrides. If you allow only administrators to create overrides, you can omit creating the advanced operator role.

As you start creating roles and adding user accounts as members of the role, you may find that you have users who need differing levels of control over various aspects of the management group. When you have a user account that is a member of multiple roles, the user will be granted all of the rights and permissions that are given to every role. This means that a user could be an author for one group but only an operator for other groups. The roles are evaluated independently for each role, so the author permissions are not granted to the other groups. If you do want a user to have the ability to create alerts, monitors, and rules as well as to be able to monitor and manage objects, you can add the user to both the author and operator roles for the same group, and the user will be granted both sets of permissions for the group.

FIGURE 7.13
Defining views



Customizing the Console

So far in this chapter we have explored configuring and manipulating the views that are available to Operators. However, our discussion has been mainly from an administrator's viewpoint. If you are an Operator who does not have the permissions required to create a view and save it in a management pack, you can still create your own view of the data that is available from your console. Some of the steps are similar to what you see when you create a view as an Author, but you will not see the Author workspace available. Only the Monitoring and My Workspace workspaces are exposed to operators.

In each workspace, you will find the views that you have permissions to see. As you navigate through the list, either in the My Workspace or Monitoring workspaces, you may find yourself a little

overwhelmed by the amount of data presented to you. There are methods to limit the total amount of data that appears in the Details pane.

Limiting the Data in the Details Pane

Let's take a look at the ways you can manipulate the information that is presented to you within the Operations Console.

SCOPE

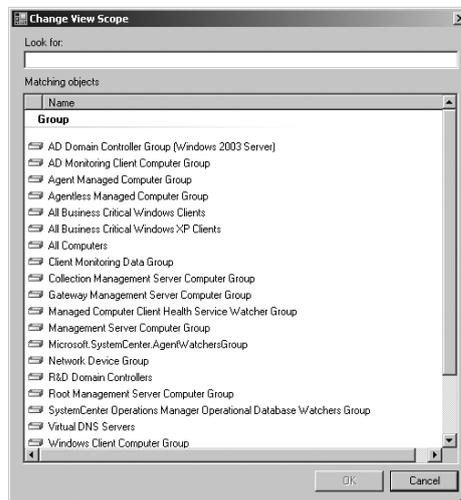
When you initially install Operations Manager, you probably won't have very many monitored systems to view. As you extend your monitoring to other servers and client systems, you may find that you have several hundred or even thousands of systems appearing in the console. You may also see far too many alerts appearing at the same time, making it difficult to determine which alerts you should work with.

To filter the data that appears in the console, you can scope the view. For instance, you can create a scope that only displays Exchange edge servers, and then give the operators who are responsible for monitoring the edge servers permissions to manage those servers.

You also have the ability to temporarily modify a view by choosing to scope the data and filter out anything that you do not want to see. This way, you can eliminate the extraneous data that appears in the Operations Console and make decisions on the remaining information. If you have the ability to open up the operations console and view aspects of the monitored infrastructure, you can modify the details that are shown.

When you open the Monitoring workspace, you are presented with the Monitoring Overview, which provides a quick summary of what is happening in your monitored environment as well as a summary of actions you can perform. Notice that the Scope button is available directly above the Monitoring workspace. If you click the Scope button, you are presented with the Change View Scope dialog box (see Figure 7.14). This is the same dialog box that will appear no matter which view you are focused on. The options will change as you add and remove management packs in your management group. So when you add in the Exchange Server 2007 management pack, the groups that are imported are displayed. Any group that you manually create will also appear in this list.

FIGURE 7.14
Change View Scope
dialog box



Once you select a group and click OK, only the data that meets the scope criteria, which is any system in the group, will be displayed in the view's Details pane. Notice that you can quickly see what the current scope is by looking at the information bar directly above the Details pane, shown in Figure 7.15. The information bar also allows you to immediately revert to a nonscoped view by clicking the X in the top-right corner. You can also cancel out the scope view by clicking the Scope button once more.

FIGURE 7.15
Scope information bar



One option that comes in handy is the ability to change scopes quickly by selecting the Change Scope link from the information bar. When you click this link, the Change View Scope dialog box appears, which allows you to select another scope to work with. To do the same thing with the Scope button, you would have to click the button to turn off the existing scope, and then click it once again to set a new scope.

FIND

After you have defined the scope that you want to view, you may still have far too many entries in the Details pane. Having so much to sift through is often frustrating as you are looking for one specific entry. That is where the Find button comes in handy. When the Find button is selected, the Look For search box is displayed. From this search box, you can enter the search criteria that you want to use to parse through the data.

When you enter your search criteria and click the Find Now button, the Details pane will change to display the entries that meet the search criteria. Filtering out all the other data makes it easier to discover exactly what you were looking for.

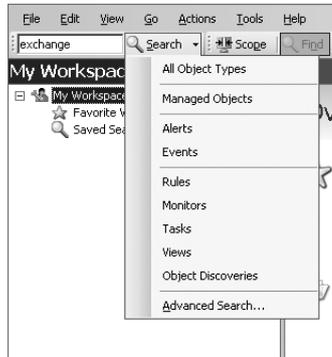
SEARCH

Whereas Find will filter the data in a scoped view, Search allows you to parse through the data and will return information from each of the object types that meet the search criteria. This comes in especially handy if you are unsure of the scope that you need to work with when trying to find pertinent information. Entering the text string in the text field at the top of the Monitoring pane and clicking Search will start a search through all of the data. Once you have the data returned to you, you will find yourself with a new window, where you can scroll and locate the information you desire.

If you know the type of data you are looking for, you can select the object type from the pull-down list, as shown in Figure 7.16. The default option from this list is All Object Types, but if you want to filter on any of the object types, you can do so by entering the search criteria and then using this pull-down.

More often than not, when you decide to search for something, you end up needing to search for the exact same thing again at a later time. Having to continually re-create the search every time can be time-consuming. To cut down on the number of times you have to re-create the search, you can create a saved search and then simply rerun it as necessary. Saved searches become available from the My Favorites view.

FIGURE 7.16
Search object selection



To create a saved search, click the down arrow next to the search text box and select Advanced Search. Once you do, you will see the Search For Specific Object Types drop-down list, which lists the objects types you can limit your search to. As you can see in Figure 7.17, there are several conditions that you can use when creating your search. The example here displays the search criteria that exist when you select the Events object type. If you select a different object type, the list will be different. Figure 7.18 shows the criteria that exist for managed objects. If you are interested in the object types that are available for you to search, Figure 7.19 displays the available options from the pull-down.

FIGURE 7.17
Advanced Search criteria

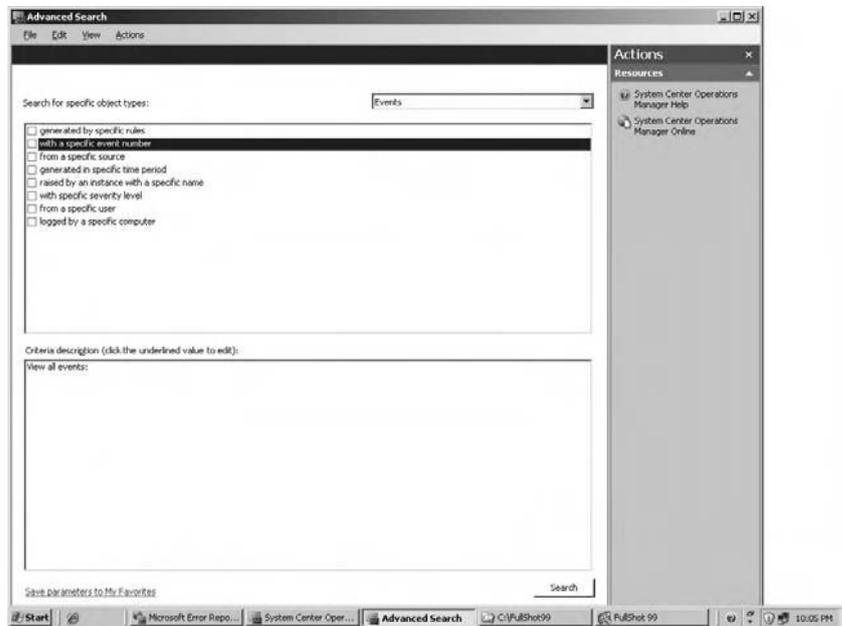


FIGURE 7.18
Search criteria for managed objects

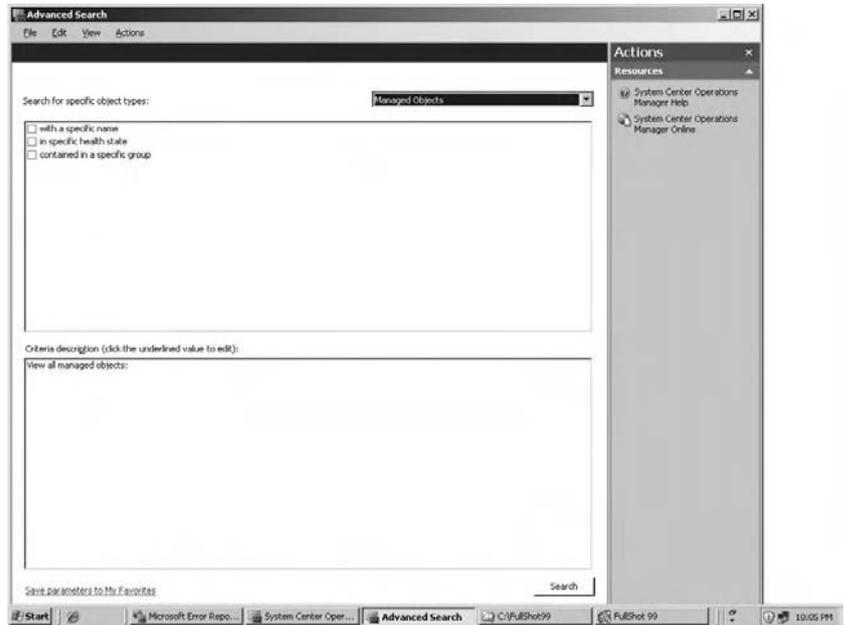
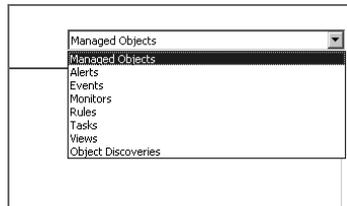


FIGURE 7.19
Search object types



As you build your search criteria, you can select as many of the search options as you need to define your query and they are added together as an And condition. Figure 7.20 shows a query being built that searches for an aggregate monitor that has Entity Health in the description. Just as you can with programs such as Outlook, you can build the query by clicking the links in the Criteria Description. Once you have entered the appropriate information, you can test the query by clicking the Search button. If you like what you see returned from the search, you can click the Show Parameters link at the top of the Advanced Search results page, which will bring you back to the Advanced Search builder.

After you enter the search criteria, click Save Parameters To My Favorites to save the search. You will be prompted to name your search, as shown in Figure 7.21, and then your search will become available from your My Workspace. When you open My Workspace and click on the Saved Searches view, you will see your search, as in Figure 7.22.

FIGURE 7.20
Advanced Search
criteria builder

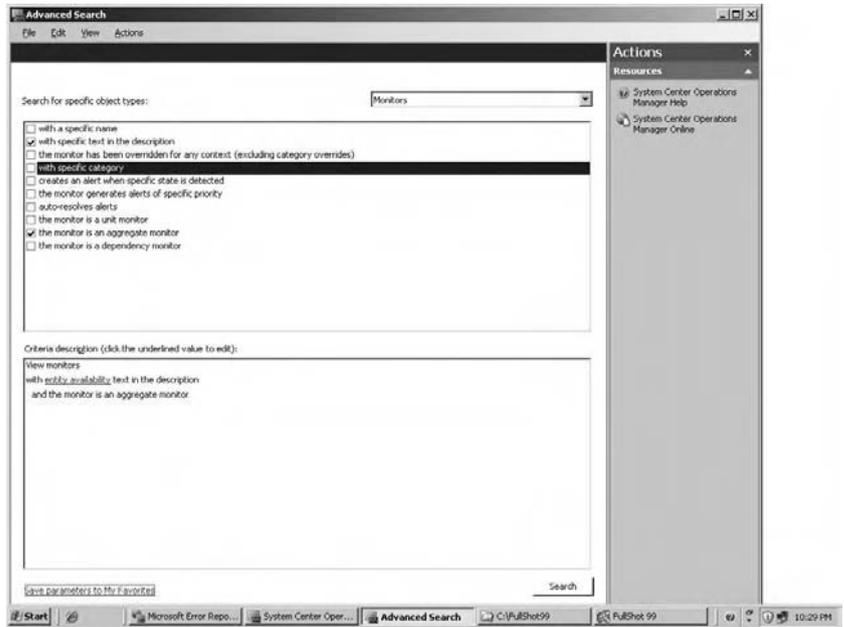


FIGURE 7.21
Naming the search

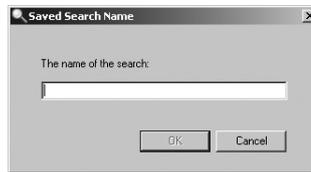
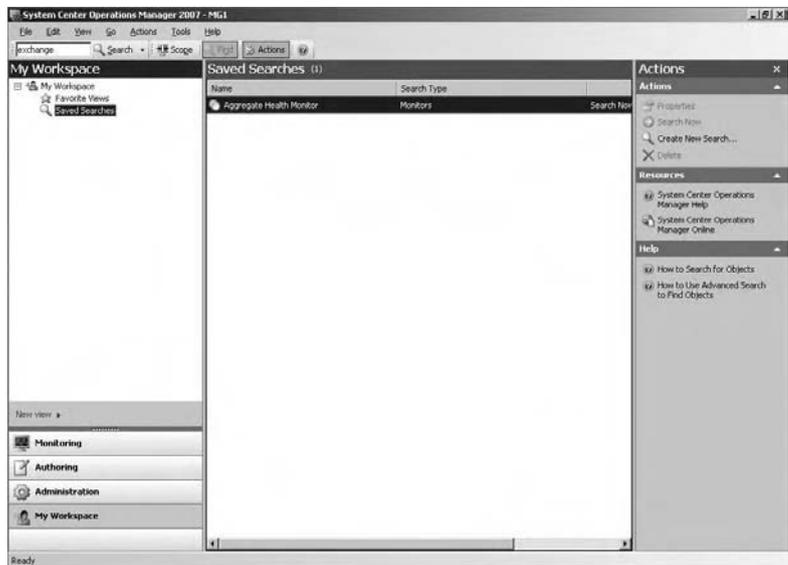


FIGURE 7.22
Saved searches in
My Workspace



Creating Views As an Operator

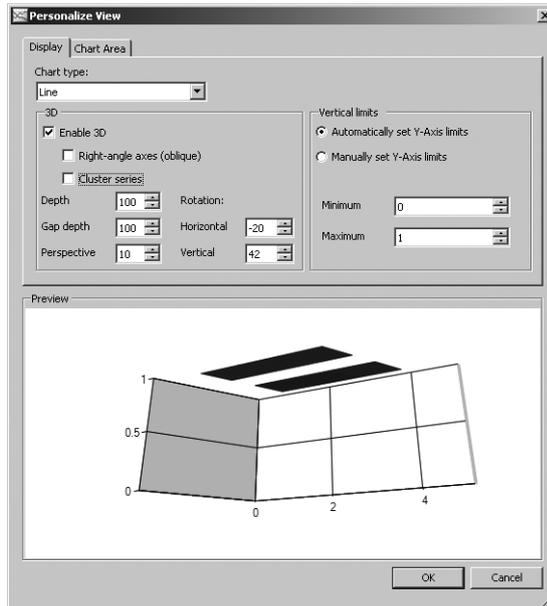
If you are an operator who does not have permissions to create views that will be included in management packs, you may feel as if you are limited to only using the views that have been created for you. However, you do have the ability to create views that are personalized for your own use. You can't include them in a management pack, but you can save them in your My Favorites workspace. Once a view is saved, you can go back to the view as often as necessary, and you have the ability to modify the view whenever you deem necessary.

Creating the views is the same whether you are an operator or an author; the only difference is that as an operator you will not have the option to choose a management pack to store the view. To create the new view, select My Workspace. Right-click Favorite Views and select New. All of the view types are available as well as a Folder option that allows you to organize your views. Following the directions from the "Creating Your Own Views" section earlier in this chapter, you can create a view for any of the groups to which you have access.

The easiest way to create a view in My Workspace is to find one that you like in the Monitoring workspace, right-click it, and select Add To My Workspace. This will create a copy of the view in your My Favorites using the name that you supply. You can even specify the folder under My Favorites where it will be added. From there, you can modify the view to meet your needs, and then it is always easy for you to find in My Favorites.

You can also modify the views that are available from the Monitoring workspace, but you can't save the modification in any of the management packs unless you are an author. Instead, you can manipulate the view as much as you like and the Operations Console will remember your view settings. When you right-click a view, you are presented with a context menu that includes the option Personalize View at the bottom of the list. Clicking this option brings up a formatting page for the view you have selected. Manipulate it to your specifications and then click OK to see if it is formatted as you had planned. Figure 7.23 shows the Personalize View page for a performance view.

FIGURE 7.23
Personalizing a
performance view



Configuring Notification Channels

In Chapter 3, “Management Group Settings,” we examined the settings that are available for configuring notification channels. At that point, we discussed the high-level options. As you learned, there are a lot of options for you to work with when you start configuring notification channels. Because you want to make sure that your operators are aware of alerts that are raised, you will have to make sure the channels are configured—then you can take advantage of the channel when something happens.

Take, for instance, an alert being raised because the drive space on a critical system is running critically low. Instead of letting the alert only appear on the Operations Console for the operators, you may want to immediately notify the person responsible for maintaining that system. Of course, the notification channel that you use is entirely up to you, and will be based on the technology that you have implemented in your organization.

The four channels available—E-mail, Instant Messaging, Short Message Service, and Command—allow you to control how the messages are delivered to operators. You can use a single option for all of your notifications, or you can mix and match according to your needs. Configuring the channels is not the end of the process, however. You must define the operators who will be receiving the notifications. To do so, you will have to open the Notifications node of the Administration workspace.

The first node beneath Notifications is Recipients. Right-click the Recipients node and select New Notification Recipient from the context menu to open the Notification Recipient Properties dialog box, shown in Figure 7.24. You can specify the user or group that will receive the notification. As a best practice, you should always define a group in this text box. More often than not, as soon as you define an individual, another person will need to receive the notification. It is much easier to add a user to a group than it is to create another notification recipient.

FIGURE 7.24
Notification Recipient Properties

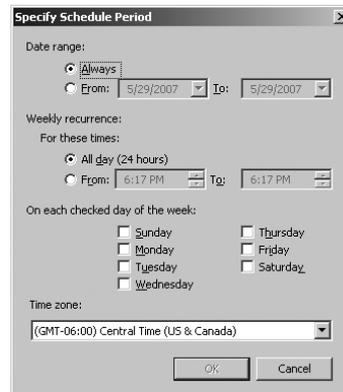


You can send notifications to this recipient at any time of the day, any day of the week, by leaving the Always Send Notifications radio button selected. Most companies attempt to divide up the responsibilities for most services and as such send notifications to different operators based on their shift. You can create a schedule for an operator by selecting the Only Send Notification During The

Specified Times button and defining the schedule using the Schedules To Send and Exclude Schedules boxes.

Click the Add button on either of the schedule boxes and you will see the dialog box shown in Figure 7.25. This dialog box lets you create a schedule that meets just about any rotation. If you have a specific date range, select the first option, Date Range, and enter the starting and ending dates, or use the Always option. If you have a standard time rotation, such as three shifts, you can enter the shift's time of operation in the For These Times section and use the From schedule, or you can use the All Day (24 Hours) option if you do not have shifts. Finally, if you have different operators working on alternating days, you can select the days that the notification will be sent by selecting them under the One Each Checked Day Of The Week section.

FIGURE 7.25
Scheduling
notifications



For example, if you have three shifts that run Monday through Friday and then on-call personnel Saturday and Sunday, you can create four notification recipients, one for each shift and one for the weekend. Each of the recipients would be a group that included the appropriate accounts that work each of the shifts. Then as you hire or dismiss users, or move them from one shift to another, you can place them in the appropriate group for their shift.

Maintenance Mode

The Monitoring workspace brings everything together into one location. Every operator will have access to the Monitoring workspace when they open it up. The details that they have access to differ based on the permissions they have been granted, but this is the area that shows an operator what is happening on the network. As the authors create the accounts and views, operators are limited to only the information they require.

As the alerts and events occur, or health status information is reported from the client systems, the appropriate operators can view the data and perform actions against them. As an operator watching the changes as they occur, you will see many different things appear on the screen. If everything has been configured correctly, a majority of alerts and events will be automatically resolved by command and scripts, so you may not have to perform any actions. For those alerts, events, and health status messages that are not automatically resolved, you must determine the best way to fix the problem.

There are other actions an operator can take besides trying to rectify problems in the organization. Occasionally a server needs to be taken offline, such as when service packs or hotfixes have to be installed or if hardware needs to be changed out. Many service packs, hotfixes, and updates

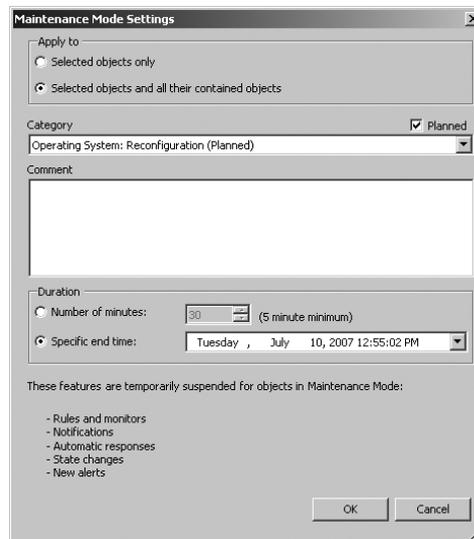
require a reboot for the installation to complete. There are also times when you need to perform periodic maintenance on a system and you will have to stop specific services or take the system offline. When you do any of these actions, you will not want Operations Manager to report that you have problems occurring on the monitored system. At the very least, when a monitored system goes offline, the heartbeat will cease and the management server will start generating alerts and possibly notifications that the server is no longer available.

With Maintenance Mode, the management servers are notified that a monitored system is no longer available for monitoring and that the management servers should ignore it until told differently. All alerts, monitors, notifications, rules, and state changes are suppressed at the agent and the server stops monitoring for a heartbeat.

Operators, advanced operators, and administrators have the ability to place any managed system, in any view in which they have access, into Maintenance Mode. From either the My Workspace or Monitoring workspace, select a view that includes the system you want to place into Maintenance Mode. When you right-click the system, you will see a Maintenance Mode ➤ Start Maintenance Mode menu item. You can also select Start Maintenance Mode from the Actions pane.

The dialog box that appears, shown in Figure 7.26, allows you to configure exactly how the Maintenance Mode will function. In the Apply To area, you can choose the Selected Objects Only option, which will only stop the monitoring of the object or objects you chose in the view. Or you can choose the Selected Objects And All Their Contained Objects option to place every monitored entity on the monitored system in Maintenance Mode. This comes in very handy if the monitored system supports monitored services or parts of a distributed application.

FIGURE 7.26
Maintenance Mode settings



The Category pull-down allows you to define why Maintenance Mode was being initiated. This is similar to the Shutdown Event Tracker that is part of Windows Server 2003. You can select the Planned check box if you have planned your downtime, or you can leave the check box deselected to choose unplanned outage reasons. The pull-down includes:

- ◆ Hardware: Maintenance – Planned or Unplanned
- ◆ Hardware: Installation – Planned or Unplanned

- ◆ Operating System: Reconfiguration – Planned or Unplanned
- ◆ Application: Maintenance – Planned or Unplanned
- ◆ Application Installation – Planned
- ◆ Application: Unresponsive – Unplanned
- ◆ Application: Unstable – Unplanned
- ◆ Loss of Network Connectivity – Unplanned
- ◆ Security Issue – Planned
- ◆ Other – Planned or Unplanned

After selecting the reason from the Category pull-down, you should leave a detailed reason in the Comment text box. At a later time, or when you have to document gaps in the monitoring, the comment will help you determine why you had stopped monitoring. You can make this as simple or as complex as you wish. You should come up with a standard that includes the exact reason Maintenance Mode was initiated, as well as the operator who performed the action and who authorized the actions and the fix that was put into place.

In the Duration area, specify how long Maintenance Mode will be enabled. Use the options here to define a specific time frame by selecting the Number Of Minutes option and entering the number of minutes. Notice that 5 minutes is the minimum and 1,051,200 is the maximum. This effectively gives you two years to rectify the problem you may have. The other Duration option, Specific End Time, allows you to define the exact date and time that Maintenance Mode will end. Again, you can select any day and time in a two-year period.

After you have performed your maintenance on the system, you will want to reenable monitoring so that you can see how the system is performing. The method to stop Maintenance Mode is the same as starting it. Right-click on the system that is in Maintenance Mode and select the Stop Maintenance Mode option.

In Figure 7.27, you can see that the server that was placed in Maintenance Mode now has a Maintenance Mode icon next to it. The icon, which looks like a wrench, allows you to tell at a glance which systems are in Maintenance Mode. You can then select any object that is in Maintenance Mode and the Actions pane will change to grant you access to two other options: Edit Maintenance Mode Settings and Stop Maintenance Mode.

FIGURE 7.27
Maintenance
Mode icon

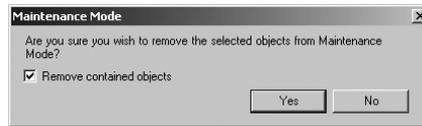
State	Name	Agent	Management Server
Healthy	DC.systemcente...	Healthy	
Healthy	OPSMGR syste...	Healthy	
Warning	XPWkstn1.syste...	Warning	

Choosing Edit Maintenance Mode Settings allows you to make some quick changes to the settings that you had specified when you initially put the system into Maintenance Mode. So if you decide that the initial duration was not long enough, you can extend the amount of time. You should always place more information into the comments if you make any further changes. That will allow you to go back to the comments at a later date and view documentation on what occurred.

The other action, Stop Maintenance Mode, allows you to take an object out of Maintenance Mode before the duration time is reached. When you click the action, or select it from the context menu when you right-click the object, you are presented with a dialog box, as shown in Figure 7.28.

You can remove the object and all other objects it contains by leaving the check box selected, or you can remove only the individual object, which means you will have to manually stop Maintenance Mode on all the other objects.

FIGURE 7.28
Stopping
Maintenance Mode



Real World Scenario

ADDING ACCOUNTS TO ROLES

Throughout this chapter we have discussed how membership in the author and operator user roles can affect what the user account can do in the Operations Console. As users are added to roles, they can perform the actions those roles provide to the views that have been configured for the role. Whenever you are adding accounts to the User Role Members text box, take a moment to think about the ramifications of adding those accounts.

Of course, not only are we referring to the abilities the users will have once they have been added to the role; we are also taking into consideration the ease of administrative management. When you click the Add button to choose the accounts that will be made members of the role, you have the opportunity to add user accounts as well as groups and computer accounts.

When you start working with user accounts, you increase the amount of administrative overhead. When a user is identified as a new operator, you need to open each of the user roles that the user should be a member of, and then add the user to the role. If the user must be added to five user roles, you will have to perform this action on each of the five user roles. If instead you create a group that is used as the member and then add the user accounts to the group, any new user who needs to be added in will only have to be added to the group. Once added to the group, the user will become a member of all the user roles where the group is a member. This same theory is the basis for allowing access to resources in your organization and the assignment of rights to users. Assign rights and permissions to one object, and then add other objects to the initial object so that they are granted the rights and permissions they require.

One other thing to think about: if you add a computer account to any of the user roles, you are essentially granting the computer account, any services on that computer account, and any user who logs on to that computer account the same level of permissions that are granted by the user role. This could potentially be a security problem, so you will want to make sure you do not add computer accounts to these roles.

The Bottom Line

Identify the workspaces and views that are available from the console. Once you install Operations Manager 2007, workspaces are created for you that allow users who have been granted access the ability to perform certain functions in the management group. The workspaces Monitoring, Authoring, Administration, and My Workspace can be made available to users, or you can limit what they can see.

Views can be created that filter the data from the database. Each management pack that is imported includes views that you can use to display only the data that meets the requirements of the view. Other views can be created by administrators and authors.

Learn how to customize a view. Administrators and authors can create new views that are made available to anyone in the management group. These views are stored in a management pack, which can be exported and used in other management groups if necessary. Operators can also personalize their own views of the data by selecting the Personalize View option and then changing the display properties of the view. My Favorites and Saved Searches in My Workspace can be used to create completely new views and filtered searches.

Create views and limit the operator's view. The default user roles allow access to all the groups and views in the management group. You can create your own user roles and then select the groups and views that the user role can access.

Learn what tasks are and how to create and use them. Tasks are commands or scripts that you use to perform actions or run tests on monitored systems. You can create two types of tasks: console and agent. Agent tasks are executed by the agent on either the monitored system or the management server. Console tasks are initiated by the system that is running the Operations Console. The tasks associated with the object selected in the Operations Console become available for use. Administrators, advanced operators, and operators are allowed to run the tasks.

Master It! In this exercise you'll create a new view in a new management pack. You'll then create a task to go along with the view. A new operator user role will be created, and a group will be added to the role. Finally, you'll log on as a member of the new role, create a new view in the My Workspace Favorite Views, and set up a saved search.

Creating Accounts:

Before creating the views, management packs, tasks, or roles, create a few accounts to be used:

1. Log on to the domain controller and open Active Directory Users and Computers.
2. If you do not have an OpsMgr OU, create one and then select it in the console.
3. Right-click the OpsMgr OU and select New ➤ User.
4. In the New Object – User dialog box, create a user that will be used as an author user role account.
5. Right-click the OpsMgr OU and select New ➤ User.
6. In the New Object – User dialog box, create a user that will be used as an operator user role account.
7. Right-click the OpsMgr OU and select New Group.
8. In the New Object – Group dialog box, create a global security group named **OpsMgr Author**.
9. Right-click the OpsMgr OU and select New Group.
10. In the New Object – Group dialog box, create a global security group named **OpsMgr Operator**.
11. In the properties for the OpsMgr Author group, add to the Members tab the Author user account that you created.

12. In the properties for the OpsMgr Operators group, add to the Members tab the Operator user account that you created.
13. Log on to the Operations Manager server as the Operations Manager administrator.
14. Open the Operations Console and select the Administration workspace.
15. Select the User Roles node.
16. Double-click the Operations Manager author role.
17. On the General page, click Add.
18. Enter or find the OpsMgr Authors group you created and click OK.
19. Right-click the Management Packs node and select Create Management Pack.
20. In the Name field, enter **Research and Development**; then click Next and click Create.
21. Log off.

Creating a new view:

1. Log on to the Operations Manager system with the Author account you created.
2. Open the Operations Console and select the Monitoring workspace.
3. Right-click the R & D Domain Controllers folder and select New ➤ Alert View.
4. In the Name field, enter **R&D Domain Controller Alerts**.
5. Next to the Show Data Related To field, click the ellipsis (...) button.
6. In the Look For field, type **r&d**.
7. Select the R&D Domain Controllers target and click OK.
8. Click the Of A Specific Severity check box.
9. In the Criteria Description area, click the Specific link.
10. Select the Critical check box and click OK.
11. Click OK to save the view.

Creating a new user role:

1. Select the Administration workspace.
2. Under the Security node, right-click User Roles and select New User Role ➤ Operator.
3. In the User Role Name field, enter **R&D Operators**.
4. Add in any user who you would like to be a member of this role and click Next.
5. Deselect the management group check box to make the individual groups available.
6. Select RD Domain Controllers from the list of groups and click Next.
7. Click the button Only Tasks Explicitly Added To The 'Approved Tasks' Grid Are Approved.
8. In the Look For field, type **active**.

9. In the Tasks grid, select the check boxes for DCDIAG, DCDIAG Verbose, GP Update, List Top Processes On DC, NETDIAG, NETDOM Query FSMO, Start Netlogon Service, and Stop Netlogon Service. Then click OK.
10. Click Next.
11. Select the button Only The Views Selected Below Are Approved.
12. Click the R & D Domain Controller Alerts check box and click Next.
13. Review the summary information and click Create.

Creating views as an operator:

1. Log on as a user who is a member of the R&D operators role.
2. Open the Operations Console.
3. Open My Workspace.
4. Right-click Favorite Views and select New ➤ State View.
5. Name the new view **DCs in Maintenance Mode**.
6. Click the ellipsis (...) next to the Show Data Related To field and select the R&D Domain Controllers group.
7. Click the Is In Maintenance Mode check box and click OK.

Chapter 8

Audit Collection Services and Agentless Exception Monitoring

Yes, the title of this chapter is rather long. We are presenting two topics here that are new to Operations Manager, yet have been on people's radar for a while now. Audit Collection Services (ACS) has been through a couple of trial runs and is now being included as part of Operations Manager. Although many administrators were looking for it to become a stand-alone product, it makes sense that it will be rolled into Microsoft's monitoring solution, since one of Operations Manager's functions is to parse log files.

You probably know one of the primary components of Agentless Exception Monitoring by a different name: Dr. Watson. Yes, that service that runs whenever an application has a problem is still being used to collect application exception events, but we are now directing the results to the Operations Manager systems. Now you can review the exception data that has been collected as well as compare events across several machines to see if there is a correlation between the application problems.

In this chapter you will learn how to:

- ◆ Install the components for Audit Collection Services
- ◆ Configure the ACS Forwarder
- ◆ Configure the ACS Collector
- ◆ Configure Agentless Exception Monitoring

Exploring Audit Collection Services

Let's face it; it seems like auditors don't trust anyone. They come into your organization looking for weaknesses and breaches. They want to find out where your security procedures are lacking. As you prepare for an auditor to enter your organization and start reviewing policies, procedures, and settings, you will find that you have a lot of work to do to make sure you can provide security-auditing logs. And then if you do have them readily available, some auditors will look at them as if they are not trustworthy. This is because they know that anyone with administrator-level credentials has the ability to manipulate the audit settings or logs.

Even if you are not worrying about an auditor coming in to review your environment, you probably trust your audit settings to let you know when there are security breaches. You configure the audit settings based upon the role of the computer, and then review the settings periodically to make sure everything is working as planned. Entries in the security logs on each system help identify whether you are having problems, either inadvertent or intentional.

The only problem with using the built-in auditing scheme is that the auditing is decentralized. You do not have one central repository for security log information. To review the logs, you have to access them one by one. In large environments, this can be a very time-intensive approach. Due

to this fact alone, many administrative teams have identified just those systems that they deem a high priority. Using this method, you could potentially have data that is accessed, or possibly damaged, by someone. You also run the risk of falling out of compliance. In the United States, for instance, new legislative restrictions that have been introduced, such as the Sarbanes–Oxley Act and the Health Insurance Portability and Accountability Act (HIPAA) that may determine what level of auditing you need to employ.

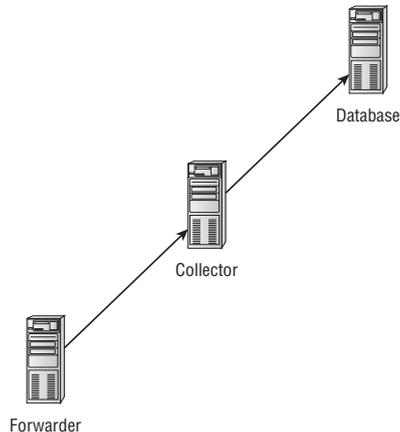
Audit Collection Services (ACS) may be the answer to many of the limitations that you may have encountered when trying to configure your auditing solution. Using this tool, you can centralize the audit database and alleviate some of the possible risks involved in decentralized auditing. As events occur, the ACS database is notified of the event. You can then monitor for specific events much more easily than before, as well as report on the state of your infrastructure. You can also completely remove the local administrator accounts on each monitored server from the auditing equation. Since the events are forwarded to the database almost as soon as they occur, the local administrator account cannot adversely affect what is collected. Used in conjunction with a strict group policy infrastructure, you can start protecting and monitoring your security settings.

You can also fine-tune what you are collecting so that you don't populate the database with too much information. Whereas you know that you can control the audit settings in your group policies, and then control the level of auditing based on the role of the system, you can also control which systems participate in the ACS reporting. Not every system in your organization has to start forwarding security log events to the database. You can choose which systems will forward data, thus monitoring only the sensitive systems.

When you decide that you want to move to a centralized collection solution, you will need to turn on some additional functionality as well as install and configure some additional components. ACS is not installed and functional when you turn on Operations Manager 2007. Microsoft realizes that not everyone is going to be interested in auditing, but for those who are, it is relatively easy to get ACS up and running.

Audit Collection Services Components

When you start investigating ACS, you will quickly find that only three components make up the service: the forwarder, the collector, and the database. These three components working together allow you to move entries from a computer's security log to a SQL-based database. As an added bonus, when you install the components, queries and reports are made available so that you can view the data that is collected.



ACS COLLECTOR

The collector is an aptly named component responsible for taking the data that is delivered by the forwarder, parsing it, and sending it to the database in a normalized form. Raw data that is forwarded usually has far too much information, and most of it is redundant. The collector decides what should be kept and what will be disposed of. For instance, the event description is discarded because that information already resides in the database. The data that is passed—time, data, event ID, source, and so on—is normalized so that the database does not include megabytes of redundant data.

When you start planning your rollout of ACS, decide which of your management servers will take on the role of the collector. Since the service will consume resources on your management server as it parses and normalizes the data forwarded to it, you may want to allocate a management server that is not going to be consumed with monitoring many other systems. Your best bet is to dedicate a management server to the ACS task, but that may not always be a practical solution. Determine which server you would like to act as the collector, and then limit the number of agents that report to it. Then you can configure the forwarders to send their security-related events to the collector.

You will find that the systems that usually forward the most events are domain controllers. Because they are the central authentication and authorization systems, most of the security events are generated by them. Microsoft claims that a collector can handle approximately 100 domain controllers. (This figure is based on a system that is acting as a dedicated collector running on a system with a 2.8 GHz processor and 2 GB of RAM.) But as they say, your mileage may vary. Test the system to determine the load it can handle before committing too many forwarders to it.

One limitation with this initial release of ACS on Operations Manager 2007 is the single collector/database relationship. You cannot allocate multiple collectors with a single database. If you find that you have too many forwarders for your collector and you want to direct the remaining forwarders to another collector, you will have to create a new database for the second collector. This means that the data collected by one collector will not populate the same database as another, thus making reporting and querying more of a challenge.

There are a couple of things you can do to improve this situation. First, if you are well-versed in SQL Server 2005, you could create queries and reports that use both databases as sources. Of course, this makes creating the queries and reports more difficult, but ultimately more useful than individual reports from each database. Second, you could create multiple management groups based on the audit collection requirements. While this approach may not be conducive to your initial monitoring plan, in some cases it will work out nicely.

ACS DATABASE

The database is the central repository for the collected data. As the forwarder forwards the events and after the collector normalizes the data, the data is stored in the database. Once the data has been loaded into the database, you can access it to determine what is going on with those systems under an audit policy. You could have scheduled queries running that are monitoring for specific events. You could have created report subscriptions that automatically generate a report and send it to your security staff. The possibilities are limited only by your imagination.

As with the Operations Manager 2007 database, the ACS database must be on a SQL Server 2005 database server. For best results, you may want to think about placing the ACS database on a dedicated server. Due to the processing requirements, the database is extremely resource-intensive. If you were to use the same database server that hosts the Operations Manager database, you could introduce performance issues for one or both of the services.

Which brings us to another point to consider: you have the choice of using either the standard or enterprise version of SQL Server 2005. As you may know, the enterprise version will cost you more than the standard version. That may not be the most important criterion that you take into consideration when you are purchasing it, however. Daily maintenance tasks are handled differently between the two versions. The standard version stops accepting data during the maintenance cycle, whereas the enterprise version does not. This introduces a couple of interesting scenarios. First, if you plan on using the standard version, when the maintenance cycle starts the collector will no longer be able to send data to the database, thus causing the collector to queue the data it receives from the forwarders. If the collector's queue fills to capacity, the forwarders will be disconnected and will not be allowed to forward any new events.

To alleviate this problem, make sure that you allocate as much queue space to the collectors as you believe will be required while the maintenance is proceeding. You should also make sure that the security log on the servers you are monitoring have enough space allocated to them to hold the data that cannot be forwarded to the collector. For more instructions on how to increase the queue size on the collector, see the upcoming section "Deployment Options."

The second problem that arises from the use of the standard edition is a possible security hole in your auditing solution. It is assumed that local administrators cannot be trusted and you need to deliver events to the database as soon as possible. When the daily maintenance cycle starts, the event data that you are trying to collect becomes vulnerable to attack. If a malevolent administrator knows that the data is queued during a specific time frame on a daily basis, that administrator could attack a system and then poison the cache of unprocessed data, thus eliminating their tracks. While this scenario would take a knowledgeable individual, it would not be outside the realm of possibility.

The enterprise edition continues to allow connections as the daily maintenance is proceeding. Although the maintenance cycle is resource-intensive and the performance of the system is slowed, you do not run as large a risk of losing data. The collector queue will not have to hold as much data and an attacker will have far less time to try to cover their tracks. You should plan for the daily maintenance to run when the processing requirements are lightest and performance is not such an issue.

ACS FORWARDER

The ACS forwarder is one component that you usually don't have to worry about installing—at least if you have the Operations Manager 2007 agent already installed on a system. The forwarder is a subcomponent of the agent; it is simply not activated by default. That way, it is not consuming additional resources when it is not needed. In order for you to start using the forwarder, you must configure the other two services, the collector and the database, and then use the Operations Console to enable it.

After you have installed the collector and created the database, you need to identify the systems that will report to the collector. Not every system will have data to report. The best place to start when trying to decide which forwarders to enable is the audit policies. If you are not responsible for creating the policies that enable auditing, talk with your security staff or the group policy staff to determine which systems fall under an audit policy. In some organizations, every system is monitored, whereas others only audit sensitive systems.

Once you have determined which systems fall under an audit policy, you can determine which ones you want to collect events from. Again, you should consult with your security staff on this one. They may not be concerned with events coming from workstations but will want to see anything that comes from the servers, especially domain controllers.

Deployment Options

To deploy ACS, perform the following four tasks:

- ◆ Make sure you have a SQL Server 2005 system on which you will create the database.
- ◆ Choose the management server that will be the collector.
- ◆ Create an audit policy for your systems.
- ◆ Configure the forwarders.

DEPLOYING THE DATABASE

If you are not responsible for your SQL Servers, make sure that you interface with the SQL administrators in order to create your database. If you determine that you need to have a dedicated system, you may run into resistance not only from the SQL administrators, who may think their existing systems will not be able to handle the load, but also from management, who may not want the extra expense in the budget. Run some tests and develop a baseline of data that you expect to see. Let them know what is required of the server. Having a test bed to start with allows you to have representative data for them to review.

If you are responsible for your SQL servers, be sure to test the system with a representative set of data. You want to build a system that can handle the load you are going to be placing on it, as well as plan for future growth. Rare is the organization that reduces the total amount of data it collects. It is more likely that your data-collection requirements are going to increase, so take growth into consideration.

You can always limit the data that is collected by controlling the audit settings or filtering data that is parsed by the collector. Both solutions have their merits, and it depends on how you want to control the situation. If you can, you should change the audit policy so that you are delivering fewer events from the forwarders to the collector. This strategy has the added benefit of reducing the network traffic generated by the forwarders. However, the drawback is that you are reducing the audit details included in the security log.

There are times when you do not have the ability to change the audit policy. You may not have permissions to alter the group policies that enforce the audit policy because it may be under control of another group in your organization. If that is the case, and you cannot convince them to make changes to reduce the events, you should filter the events that the collector sends to the database. This method has the unfortunate drawback that you are sending data across the network that never makes it into the database, but at least you are controlling the size of the database.



Real World Scenario

MODIFYING THE AUDIT POLICY SETTINGS

We have assumed that you have a good understanding of how an audit policy is configured and managed. The audit policy becomes a central tool within ACS in that it is responsible for controlling the data collected from the monitored systems. The audit policy is not a component of ACS, however. ACS allows you to centralize the results of the audit collection from each of the monitored systems and store it in a single database.

Because of the sheer number of events that can be collected from auditing, you should audit only for those events that you need to watch. While it may appear to be a good idea to enable all of the audit policy settings, doing so will inundate you with more information than you may be able to parse. And thinking that ACS will make it easier for you to find security breaches is sheer folly.

Keep in mind that the more you start auditing, the more the database will need to hold. Sit back and think about the settings you have to monitor. For instance, enabling the auditing of successful logon events for domain controllers will start filling the security event log with all the logons in the domain. This includes all of the users as well as the computer accounts as they access data from other systems. More than likely, you will want to limit your logon event auditing to report only failed attempts. This approach should cut down the number of events that are generated and still give you an idea as to how many attempts were made to access a system.

If you are creating filters, do some research to find out which events you can safely discard. Although it may seem safe to discard many of the success events, someone may be watching for specific events to appear. Even if you think the omission of an event makes sense, talk with the other groups in your organization before you make the decision to excise it. For more information on how to create a filter, see the command syntax later in the “Admin Tools” section.

DEPLOYING THE COLLECTOR

The collector is a busy little fellow. Responsible for sending data to the database, the collector becomes the traffic cop that validates the data sent from the forwarder and ensures that the database sees only the information that it should.

Earlier in this section we discussed the impact that daily maintenance makes on the ACS infrastructure. Performance can also be affected by the database and log file options that you choose. When you are running the installation wizard, you will encounter settings that allow you to specify the database and log file locations. Best practices state that database and log files should be located on separate drives. Not only does this improve performance, but also it is important from a disaster-recovery viewpoint. The log files include data that has not been committed to the database. If the server were to fail due to a drive failure, the data could be recovered from a combination of a restored backup and the log files.

To aid performance, consider placing the log files on a RAID-1 (mirror) set or a RAID-01 (mirrored stripe) set. The database should reside on a RAID-5 (stripe set with parity) drive set. This allows the database to reside on a fault-tolerant drive array that can be recovered if a single drive fails. Since the log files are mirrored copies of the data that is processed in the system’s memory, they should be written to the drive as fast as possible. Having a drive dedicated to this purpose reduces the latency involved when trying to write the log data. Having them mirrored allows you to have a fault-tolerant solution as well.

There are two authentication methods you can use when you are dealing with the database. Windows authentication is usually the method you will apply to authenticate an account against the database, since it uses the domain Kerberos authentication method. The other method is SQL authentication. You should use SQL authentication when you have collectors from foreign domains connecting to the database. If you choose SQL authentication, make sure you have an account created that can then be used to authenticate the collector from the remote domain.

Both the database and the collector can be installed at the same time. The setup program allows you to configure a management server that has SQL Server 2005 installed on it to act as the collector

and host the database. Since this is not the recommended configuration, you can also install the database on a remote SQL server when you install the collector locally. The installation files are included with the Operations Manager 2007 CD, but ACS is not automatically loaded with Operations Manager. Follow these steps to load it:

1. From the Operations Manager 2007 CD or the network location where you have copied the installation files, double-click `SetupOM.exe`.
2. On the Setup page, click Install Audit Collection Server.
3. On the Welcome page of the Audit Collection Services Collector Setup wizard, click Next.
4. Accept the license agreement and click Next.
5. Click Create A New Database on the Database Installation Options page, seen in Figure 8.1, and then click Next.
6. Enter a name for your ODBC data source name (DSN) on the Data Source page, seen in Figure 8.2, and then click Next. The default name of OpsMgrAC is entered automatically.
7. If the database is hosted on a remote server, click Remote Database Server and enter the name of the SQL server as seen in Figure 8.3. If the database is hosted on the local management server, click Database Server Running Locally.
8. Enter the name you would like to use for your database. You have the option of leaving this field blank in order to use the default database name of OperationsManagerAC. Click Next.
9. If the collector and database server are in the same domain, you can choose to use Windows authentication on the Database Authentication page, seen in Figure 8.4. If the two services are not in the same domain, select SQL authentication. Click Next.
10. As seen in Figure 8.5, the Database Creation Options page, select the Specify Directories option and enter the path where you would like to store the database and log files; preferably place the log files on a RAID-01 set and the database on a RAID-5 set.
11. The Event Retention Schedule page seen in Figure 8.6 allows you to configure the daily maintenance schedule and the retention limit for data in the database. Under Local Hour Of Day To Perform Daily Maintenance, enter a time of day when you expect to have the least number of security-related events.
12. Under Number Of Days An Event Is Retained in Database, enter the total number of days that you wish to keep events stored in the database. The default is 14 days, but you can alter this setting to reflect your auditing requirements. Click Next.
13. If you want to have all of the data reflecting the local time for your time zone, select Local from the ACS Stored Timestamp Format. If you want to show the time according to Greenwich Mean Time, select Universal Coordinated Time (UCT), shown in Figure 8.7. Click Next.
14. Review the settings on the Summary page, as seen in Figure 8.8, and click Next to install the database and collector.
15. When the installation is complete, click Finish.

FIGURE 8.1
Database installation
options

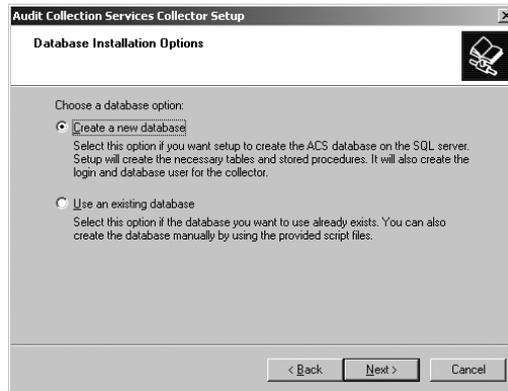


FIGURE 8.2
Data Source page

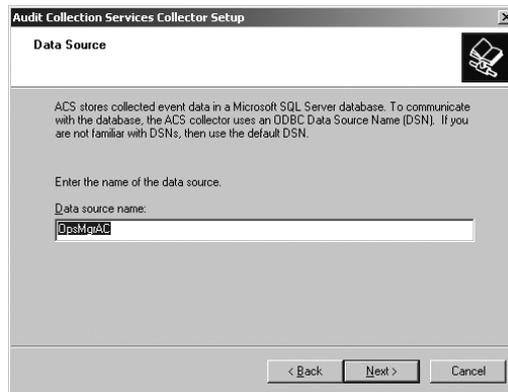


FIGURE 8.3
Database name

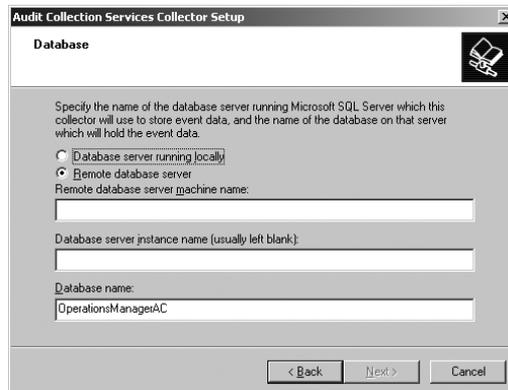


FIGURE 8.4
Authentication type

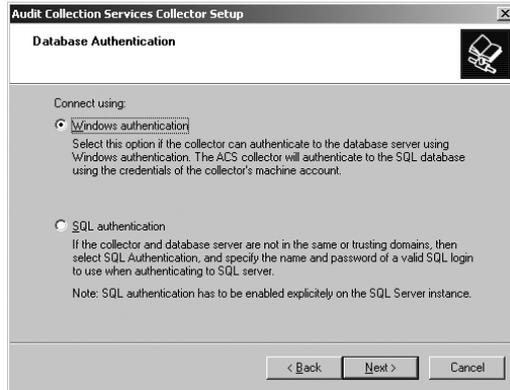


FIGURE 8.5
Database and log file locations

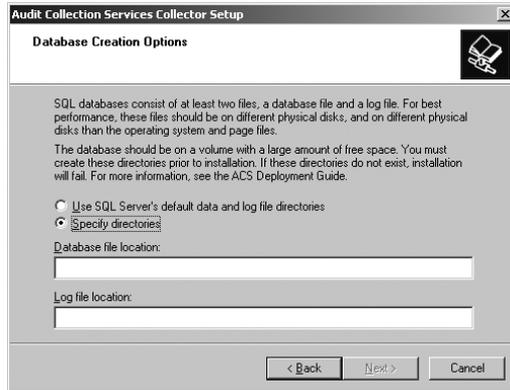


FIGURE 8.6
Cleanup and retention

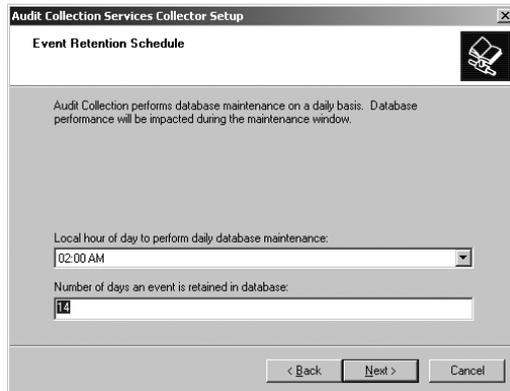


FIGURE 8.7
Time stamp

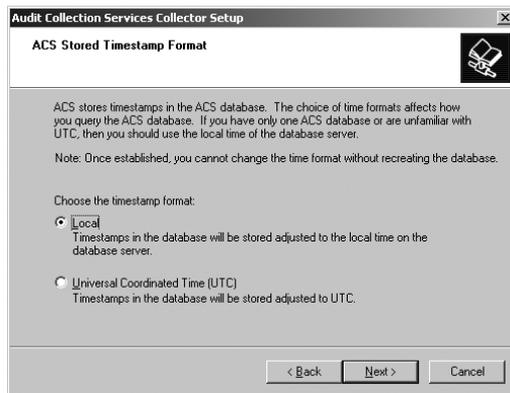
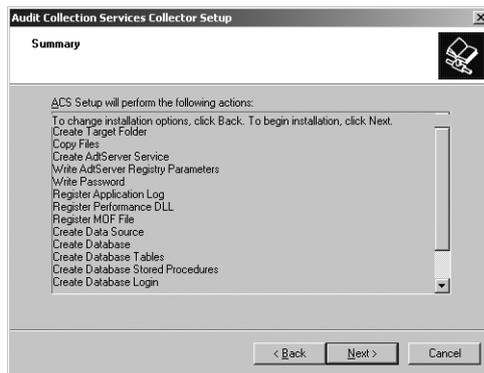


FIGURE 8.8
Summary page



CREATING THE AUDIT POLICY

Many companies already have an audit policy in place while others have avoided implementing any type of auditing. If you have any need for monitoring access to resources, auditing provides you with a good overview of what is happening. This book is not meant to cover audit policies and group policies in great detail; there are many other sources of information on those topics. A good place to start is the Microsoft website. From there you will be directed to a large number of pages that will help you determine your audit policy needs.

If you are responsible for your company's audit policy, you should evaluate exactly what it is you are trying to monitor. Many audit policies were created when a security template was imported. Some administrators are auditing without even knowing what they are collecting. A review of the security log will let you know what is being collected and will give you a good indication of the data you will be injecting into the database.

ENABLING THE FORWARDERS

Enabling a forwarder is not as simple and straightforward a process as installing the collector. There are two procedures you can perform, but neither is intuitive. If you look through the Operations

Manager documentation, you will find that you can enable audit collection through the Agent Health State View. When you open this view, note that the Details pane is divided into two windows: one for the Agent State From Health Service Watcher and another for the Agent State. Selecting a computer in the Agent State pane brings up the Health Service Tasks action pane. When you are performing either of these procedures, you must use an account that has local administrator privileges on the system for which you are enabling the forwarder. The steps to enable the forwarder are as follows:

1. Open the Operations Console.
2. Select the Monitoring view.
3. In the Monitoring pane, expand Monitoring Operations > Manager > Agent.
4. Select the Agent Health State view.
5. In the Agent State Details pane, select the systems you want to enable as forwarders.
6. Under Health Service Tasks in the Actions pane, click Enable Audit Collection. Alternately, you can right-click the system and select Enable Audit Collection as shown in Figure 8.9
7. If the account you are using is not a member of the local administrators group on the computers, fill in alternate credentials in the Run Task – Enable Audit Collection page (shown in Figure 8.10).
8. Click Run.
9. After the task is complete, click Close.

FIGURE 8.9
Enabling audit collection

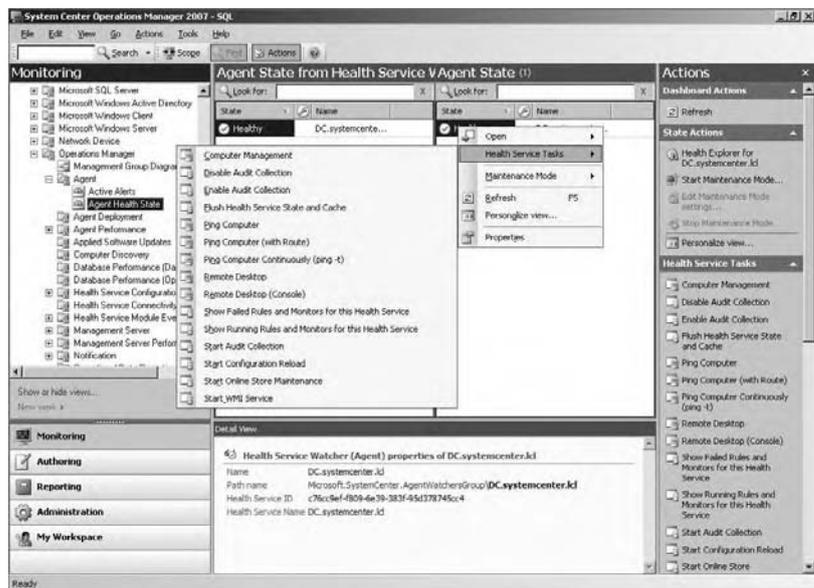
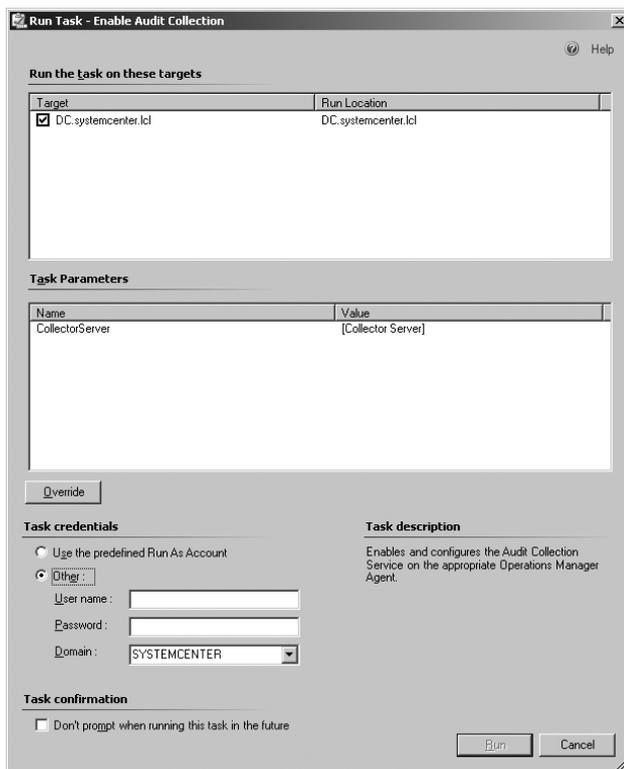


FIGURE 8.10
Providing alternate
credentials



Another method you can use that will allow you to more easily identify the forwarders is to create a new state view named after Audit Collection Services. You will need to create a new view, and then select the appropriate systems in the Operations Console and enable them. The steps are as follows:

1. Open the Operations Console.
2. Select the Monitoring view.
3. Right-click the Monitoring node, select New, and then click State View as seen in Figure 8.11.
4. On the View Properties page, name the view something descriptive, such as **Audit Collection** or **ACS Forwarders**, as seen in Figure 8.12.
5. On the Criteria tab in the Show Data Related To field, click the ellipsis (...) and select Agent, as shown in Figure 8.13, and click OK twice.
6. Select the computers you want to enable.
7. In the Actions pane, select Enable Audit Collection under Health Service Tasks.

8. If the account you are using is not a member of the local administrators group on the computers, fill in alternate credentials on the Run Task – Enable Audit Collection page.
9. Click Run Task.
10. After the task is complete you will see the successful task completion summary as seen in Figure 8.14, and then you can click Close.

FIGURE 8.11
Creating a state view

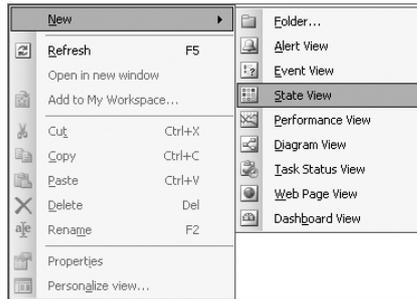


FIGURE 8.12
Configuring the state view

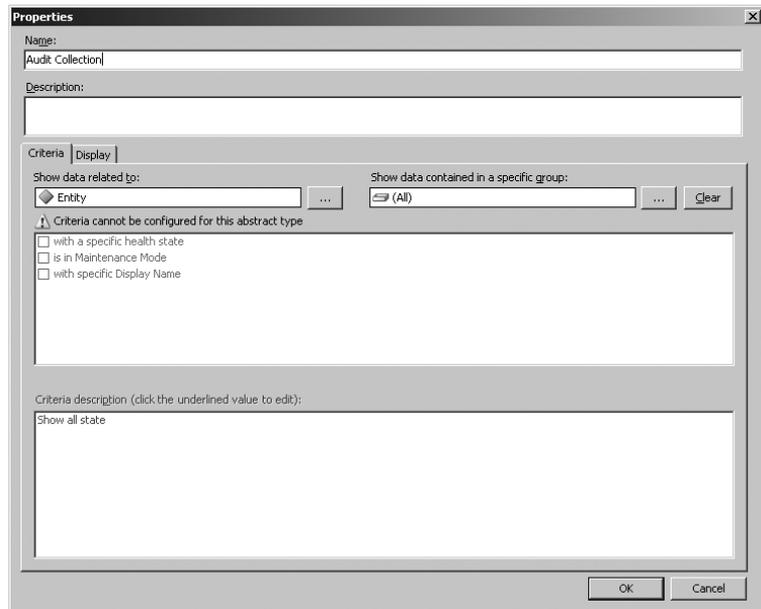


FIGURE 8.13
Choosing Target
for View

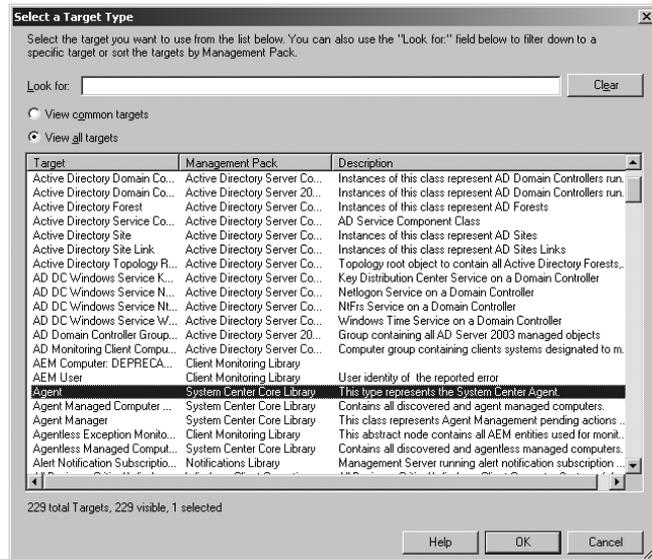
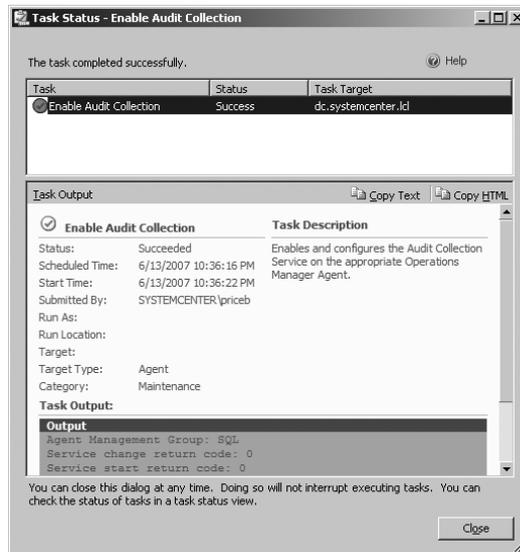


FIGURE 8.14
Successful completion
of the Enable Audit
Collection task



After the forwarder has been enabled, it will start collecting the events in the security log and forwarding them to the collector. Each forwarder can be configured to be part of a forwarder group and can be assigned a priority. If you create a forwarder group, you can configure all of the forwarders in the group to be identical to one another. Using this construct, a set of forwarders that share the same capabilities can be easily configured at the same time.

The priority for your forwarder can be used to determine which forwarders are disconnected first when the collector's queue starts to reach capacity. The forwarders with the lowest priority number are disconnected before those with a higher priority. This way, you can make sure that

sensitive systems remain connected, passing their audit data to the collector for as long as possible. For more information on how to configure the group and priority values, see the section “Admin Tools” later in this chapter.

Optimizing the ACS Collector

Security events are usually deemed high-priority events, especially in environments where sensitive information is found. Even if you do not think of your data as sensitive, all information that you use to keep your company up and running is vital to your company’s well-being. If any of that data is deleted or destroyed, intentionally or not, you could risk crippling your organization, and quite possibly lose your job in the process.

The ACS collector is a vital component in Audit Collection Services. Forwarders are programmed to send the events that appear in the security log to the collector they are configured to communicate with. They do this as quickly as possible after the event appears in the log. The collector accepts the data and processes it, evaluating whether the event should be included within the database, and if so, what information from the event should be included. Usually the data is sent to the database immediately, but there are times when the data is slowed down.

When the collector becomes overburdened during times of heavy traffic, the events that arrive are placed in a queue, where they wait to be processed. The queue is also used when the SQL server becomes busy, or when it is not able to accept incoming requests. An example is when the daily maintenance cycle hits. The collector has to queue up the information that it is preparing to send to the database until the database is able to accept.

By default, the collector is configured to queue 262,144 events. While this may seem like a rather large number of events, when you have several systems generating security events, you could consume this number of entries in a short period of time. This is another reason you may want to limit the types and amount of auditing that you enforce. If you find that you need to increase the queue size, you must edit the Registry settings that control the queue. You can also edit the Registry settings that control how the collector handles accepting events from forwarders. All of these settings are found under HKEY_Local_Machine\System\CurrentControlSet\Services\AdtServer\Parameters:

MaximumQueueLength The setting that controls the queue size is `MaximumQueueLength`. As mentioned, the default value is set as a hexadecimal value of 40,000, which has a decimal equivalent of 262,144. As you increase this value, you are essentially configuring the amount of space that will be allocated to the queue. Although it is assumed that each of the events that are forwarded to the collector consumes 512 bytes—which means the default value allocates a possible 128 megabytes—some events may be larger than others. Using this average, however, will give you an idea of how much memory you will need in order to allow the collector to queue events.

BackOffThreshold The `BackOffThreshold` value determines how full the queue can grow before new connections from forwarders are denied. Existing connections will be maintained, but any new forwarder connection will be denied, thus making sure that you do not drop any events from forwarders that may already have events being forwarded and queued. The default here is 75 percent. If you deem it necessary, you can increase this value. For instance, if a collector is only gathering events from domain controllers, you may want to forward entries from all of them until the queue becomes 85 percent full. If you do increase this value, make sure that you increase the `DisconnectThreshold` value.

DisconnectThreshold The `DisconnectThreshold` value determines when the collector will start disconnecting forwarders. Since forwarders can be configured to have a priority, the collector will disconnect those that have a lower priority in order to squelch the incoming requests while allowing the high-priority forwarders to still report their events. The default value assigned

to this entry is 90 percent. Make sure the `DisconnectThreshold` value is higher than the `BackOffThreshold` value so that you stop new connections from being made to the collector before disconnecting existing forwarders based on priority.

MaxPurgingQueueLength Before the database purge cycle can run, the `MaxPurgingQueueLength` value is used to determine whether the purge cycle is allowed to run. If there are more events in the queue than the value you have set here, the queue will be processed before the purge will run. The default is 500 events, but you can alter this setting to meet your needs. For instance, you could increase this value if you are monitoring client systems, but you may want to leave this value set if the collector is used to collect events from domain controllers or sensitive servers.

Security

Auditing should be part of your complete security plan. You need to make sure that users are able to access only that which they need, and you should give them only the level of access they require. You don't want to have someone inadvertently deleting files that they are not supposed to be able to delete, or viewing data that should be restricted from them. This holds true for your auditors as well. When you think about what level of access the auditors need to your data, you will probably determine that they need only to read the information and not manipulate it. If they are allowed to manipulate the data, they will not be able to show a true representation of how your security plan works.

Making sure that the auditors do not have too much power over the data should be a prime concern to the administrators of the Operations Manager Audit Collection solution. You need to maintain a level of autonomy over the database, allowing only the collectors to add data. To do so, assign the auditors the ability to read the data from the database, but not to manipulate it, by adding them to the `db_datareader` role for the database. Of course, the most efficient method of doing this is to create a group, and then assign the group to the role. Once the group has permissions to the database, it is as simple as adding an auditor's account to the group.

As we mentioned before, according to auditors, the local administrator accounts at each server being monitored cannot be trusted. If you are an administrator, you could manipulate the system so that it would not send events to the security log. Since the administrator account is essentially un-auditable, meaning you cannot tell who has logged on with the administrator account to make changes, you leave a big hole in your audit solution. To alleviate this problem, the forwarder will send events immediately to the collector. Any action taken on the system is recorded in the database, and an alert can be raised when something out of the ordinary occurs.

The forwarder cannot be modified by a local administrator, either. There are no command-line or scriptable solutions that allow someone to affect the operation of the forwarder. The only way a forwarder can be modified is for the collector to assign it new settings. In other words, you would have to be an Operations Manager administrator to be able to modify how the forwarders function.

I know what you are probably thinking at this point. If someone stopped the agent from running, could they perform whatever action they wanted while the forwarder is disabled? To kill off the agent, they could go into Task Manager and end the `AdtAgent.exe` process. Once the agent is down, they could perform their nefarious tasks. Of course, the agent availability should go into a critical state, which would make any good operator take notice and restart the service. If the operator is paying attention, the service should be up and running rather quickly, reporting on what is going on.

You can take this one step further and create a unit monitor that will restart the service when it is stopped. Note that this is not an immediate restart, however. After the AdtAgent service is started, the Operations Manager system will recognize that the service is not running, and then will send a start command. This could take anywhere from 30 to 75 seconds, during which time someone could perform an action on the system that will not be audited. Therefore, you may want to check the log files on the local machine once the service has been compromised.

To create a unit monitor that will restart the service, follow these steps:

1. Select the Authoring workspace.
2. In the Authoring pane, select Monitors from the Management Pack Objects list.
3. In the details pane, expand Agent > Entity Health > Availability as in Figure 8.15.
4. Right-click Availability and select Create A Monitor > Unit Monitor.
5. When the Create A Unit Monitor wizard starts, expand Windows Services and select Basic Service Monitor, as seen in Figure 8.16.
6. Under Management Pack, click New.
7. Type **ACS Overrides** in the Name field and click Next.
8. Click Create.
9. Click Next.
10. In the Name field, type **AdtAgent Availability**.
11. Make sure the Monitor Target is set to Agent, the Parent Monitor is set to Availability, and the Monitor Is Enabled check box is selected, as seen in Figure 8.17; then click Next.
12. In the Service Name field, enter **adtagent** and click Next.
13. Verify that the Service Is Running Health State option is set to Healthy and that the Service Is Not Running Health State option is set to Critical; then click Next.
14. Click to enable the Generate Alerts For This Monitor check box.
15. Make sure the Generate An Alert When pull-down selection is set to The Monitor Is In A Critical Health State.
16. If you want to automatically resolve the alert when the agent restarts, make sure the Automatically Resolve The Alert When The Monitor Returns To A Healthy State check box, as seen in Figure 8.18, is selected.
17. Change the Priority to High so that this alert will be processed quickly, and then click Create.

FIGURE 8.15
Creating the monitor

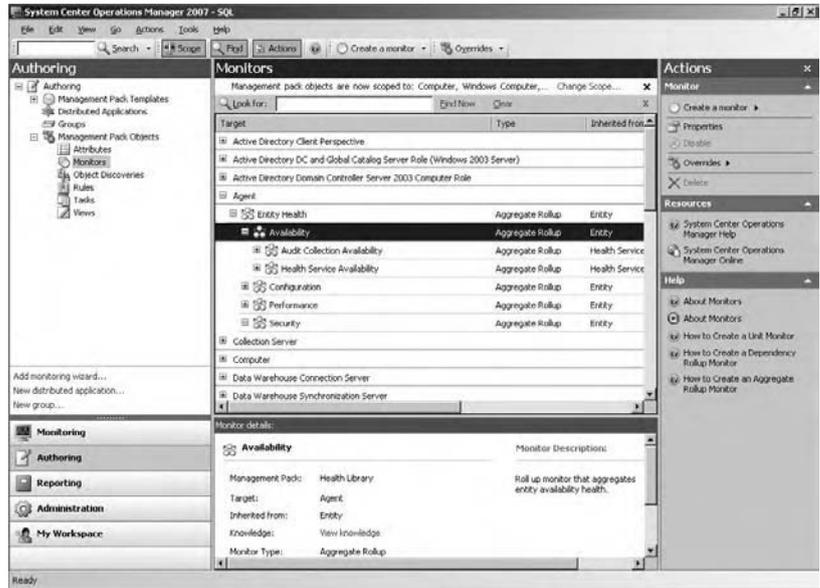


FIGURE 8.16
Choosing the Basic Service Monitor type

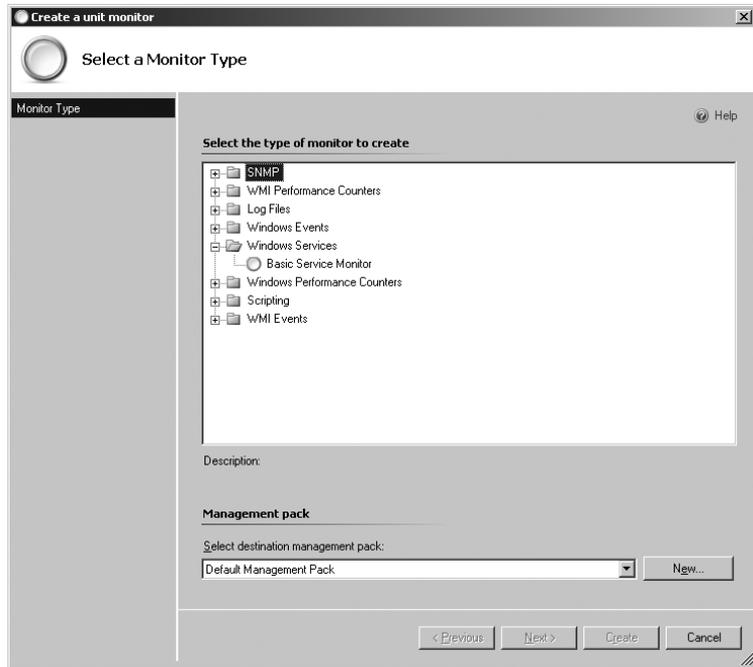


FIGURE 8.17
Setting the monitor parameters

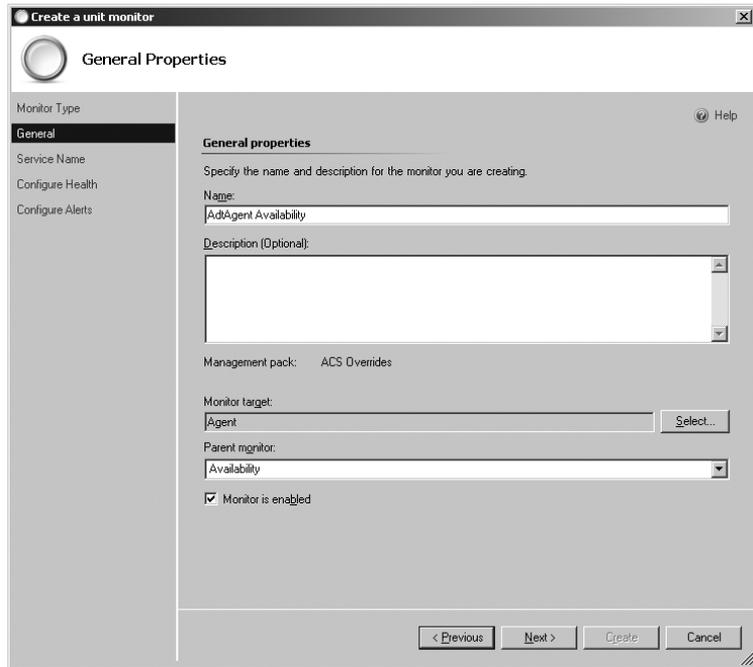
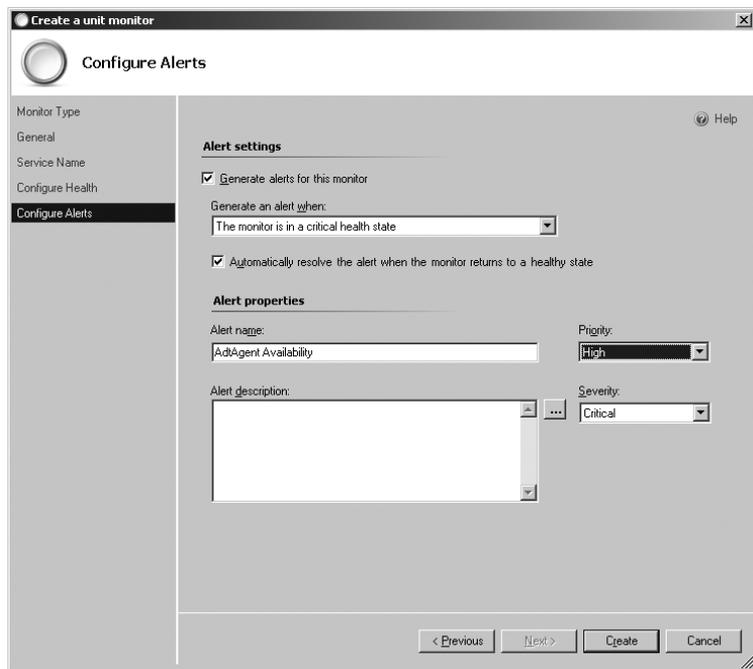


FIGURE 8.18
Configuring the monitor alert settings



If you decide to have the monitor automatically reset the alert, as we did in the steps earlier, the alert will appear only for a short period of time. The Agent Availability management pack also monitors for the agent to stop and will raise an alert. This alert is not automatically reset when the service is restarted. If you don't want to see two alerts when the service stops, disable the monitor by using an override. To do so, follow these steps:

1. Select the Authoring workspace.
2. In the Authoring pane, select Monitors from the Management Pack Objects list.
3. In the Details pane, expand Agent > Entity Health > Availability > Audit Collection Availability.
4. Right-click Audit Collection Forwarder – Windows Service and select Overrides > Disable The Monitor > For All Objects Of Type: Agent.
5. Click Yes at the prompt.

Performance Monitoring

As with most of the services and applications that you can install on a Windows platform, performance counters are supplied that allow you to monitor how the service or application is performing. When you install ACS, two sets of counters are made available for you to track how the ACS infrastructure is running. One set is used for the collector and another for the forwarders. Between the two sets of counters, you can monitor how well the ACS solution is performing. In Chapter 11, "Optimizing Your Environment," we will cover how these counters can be used to monitor the efficiency of your systems.

COLLECTOR OBJECT COUNTERS

The following 14 counters can be used to monitor the collector performance:

Connected Clients Tracks the number of forwarders that are currently connected to the collector.

Database Queue % Full Displays the percentage of events that remain in the queue. This percentage is calculated based on the total number of remaining events divided by the maximum number of events that had resided in the queue.

Database Queue Length Shows the total number of events in the queue.

DB Loader Events Inserts/Sec Shows the average number of events per second that are added to the dtEvent database table.

DB Loader Principal Inserts/Sec Shows the average number of events per second that are added to the dtPrincipal database table.

DB Loader String Inserts/Sec Shows the average number of events per second that are added to the dtString database table.

DB Principal Cache Hit % Shows the percentage of the handling requests that the principal cache has serviced instead of being handled by the dtPrincipal database table.

DB Request Queue Length Shows the number of requests that are in the request queue, waiting to be sent to the database.

DB String Cache Hit % Shows the percentage of the handling requests that the string cache has serviced instead of being handled by the dtString database table.

Event Time In Collector In Milliseconds Tracks the time it takes for an event to be sent to the database after it has arrived at the collector.

Incoming Events/Sec Shows the total number of events per second that arrive at the collector from all the forwarders.

Interface Audit Insertions/Sec Shows the total number of records per second that are sent to WMI when an audit query is requested.

Interface Queue Length Shows the total number of WMI requests pending.

Registered Queries Shows the total number of requests that WMI has received for audit queries since the collector has been running.

COLLECTOR CLIENT OBJECT COUNTERS

To check out the efficiency of the forwarders in your environment, you can use the two counters that are included in this object. It should be noted that the instances of the counters that appear when you are trying to add the counters show all the forwarders that you have currently enabled unless your environment has grown to include more than 100 forwarders. In a case such as this, when you have a large environment, the forwarders that have the highest priority will appear as instances.

By default, the priorities for forwarders are assigned by the collector. Forwarders that generate the most events are assigned higher priority values. As the priority of the forwarder is assigned, the higher the priority, the more important the connection is deemed, and the collector will accept connections from them over the lower-priority forwarders. This allows forwarders that have a lot of events to transfer to send those events before systems that do not generate a lot of traffic. If you need to change the priority of a forwarder so that it is considered a high-priority system, you can use the `adtadmin.exe` utility discussed in the next section, “Admin Tools.”

The two collector client object counters are:

Average Time Between Event Generation Shows the average number of milliseconds it takes for an event to reach the collector after it has been raised

Incoming Audits/Sec Shows the total number of events that the forwarder has sent to a collector

Admin Tools

One of the things that many administrators are going to dislike about ACS is the interface—or lack thereof. The only method of working with the collectors and forwarders is to use a command line. Although this may seem like a limited method of managing ACS, it can also be seen in a positive light. There are many people who simply do not like using command-line tools. For years those people have enjoyed the ease of use of the graphical user interface, and using a command line is somewhat foreign to them.

No matter how you feel about using the command line, it is your only choice when tweaking ACS. The command itself is `adtadmin.exe`. With `adtadmin.exe`, you can control the collector as well as the forwarders connected to it. The syntax is pretty straightforward, even though it does have several parameter options to work with.

```
adtadmin.exe /parameter /options
```

Depending on what you are trying to accomplish, you will always specify one parameter, but you may have none, one, or multiple options. The 12 parameters are:

- /AddGroup** Creates a new group on the collector
- /DelGroup** Deletes a group on the collector
- /ListGroup** Shows the groups that reside on the collector
- /UpdGroup** Used to rename a group
- /ListForwarders** Used to show the forwarders that are connected to the collector
- /UpdForwarder** Used to change the priority value, group, and name of a forwarder
- /Disconnect** Used to disconnect a forwarder or a group of forwarders from the collector
- /Stats** Used to show statistical information about the forwarders
- /GetDBAuth** Used to show the authentication method used by the collector in order to connect to the database
- /SetDBAuth** Used to configure either Windows or SQL authentication between the collector and the database
- /GetQuery** Used to show the queries currently in used to filter the events before they are saved to the database
- /SetQuery** Used to configure a query that will be used to filter events before they are saved to the database

Ten options are used to further target the command that you are issuing. Most of the options are used to define values.

- /Collector:Collectorname** Used to target the command at the collector specified
- /Group:GroupName** Used to target the command at a group of forwarders
- /GroupID:GroupID** Used to target or change a group's identifier in the database
- /Forwarder:ForwarderName** Used to target a forwarder
- /ForwarderID:ForwarderID** Used to target or change a forwarder's identifier in the database
- /ForwarderSID:ForwarderSID** Used to target the forwarder based on the forwarder's Security Identifier
- /Name:NewName** Used to specify a new name for a forwarder or group when used in conjunction with /Group or /Forwarder
- /Value:PriorityValue** Used to assign a new priority to a forwarder or group
- /GroupValue** Used to specify that the group's priority should be used for the forwarder's priority instead of the forwarder's defined value
- /pwd** Prompts for a password

CREATING A GROUP

To create a group, all you have to do is specify the name of the group. For instance, if you are going to create a group by the name of *DomainControllers*, you would issue the command

```
Adtadmin /addgroup /group:DomainControllers
```

ASSIGNING A FORWARDER TO A GROUP

Forwarders can be assigned to groups by issuing the command that specifies the forwarder you want to modify and the group it will be a member of. If you were to add the forwarder for a machine named DC1 to a group named DomainControllers, you would issue this command:

```
Adtadmin /updforwarder /forwarder:DC1 /group:DomainControllers
```

CONFIGURING FORWARDER PRIORITY

Forwarders are assigned a priority value of 1 by default when the forwarder is enabled on a monitored system. This may not be right for your environment, because you may have specific systems that you want to keep monitoring even if the database is busy or shut down. While the events won't make it to the database until it becomes available again, the events will still be passed from the forwarder to the collector, where they will be held in the queue.

To configure a system to continue monitoring while others are being disconnected, you need to configure a higher priority. When the collector's queue fills to the `DisconnectThreshold` value, systems with forwarders that have lower priorities are disconnected from the collector before those with a higher priority. If you have a management group where you are monitoring domain controllers as well as client systems, you could essentially create two groups, one for each system role, and then assign the domain controllers a higher priority.

If you decide to change the priority for the forwarder, you essentially have two options: you can change the priority for the individual forwarder, setting its value higher or lower, or you can configure the forwarder to use the value assigned to the group of which it is a member. The value range you can use for an individual forwarder is 1–99, with 99 the highest priority, trumping all the other values. If you assign the forwarder to use the value of the group, you need to assign the forwarder a value of -1.

There is one other value you can set: 0. Using a value of 0 essentially disables the forwarder. While the forwarder will still attempt to send events to the collector that it detects in the security log, the collector will ignore all the data forwarded from the forwarder. Since this method sends data across the network, you should consider disabling the forwarder on that system instead of ignoring it.

The command that you would use to set the priority of a forwarder named CHIExch1 to 50 is as follows:

```
Adtadmin /updforwarder /forwarder:CHIExch1 /value:50
```

To set the forwarder to use the value of the group, you would issue this command:

```
Adtadmin /updforwarder /forwarder:CHIExch1 /groupvalue
```

Uninstalling ACS

If you have reason to uninstall ACS, be sure to uninstall and disable all the services. Because ACS consists of separate components that run on systems throughout your organization, each component requires a different method to stop it. The database must be deleted, the forwarder needs to be disabled, and the collector must be uninstalled from the management server. Before getting rid of the collector and the database, you should disable the forwarders that send data to the collector. Once data is no longer being forwarded to the collector, you can uninstall the collector. Finally, you can open the SQL Server 2005 Manager and delete the database.

DISABLING FORWARDERS

To disable your forwarders, follow these steps:

1. Open the Operations Console with an account that is an administrator operator, and then select the Monitoring view.
2. Click the ACS state view that you created when you enabled the forwarders.
3. Select all the systems, and then click Disable Audit Collection in the Action pane.
4. When the Run Task – Enable Audit Collection page appears, enter alternate credentials if the account you are logged in with is not a member of the local administrators group on the monitored systems; then click Run Task.
5. After the task completes successfully, click Close.

REMOVING THE COLLECTOR

To remove the collector, follow these steps:

1. Run the SetupOM.exe program from the product CD or network share where the product installation files have been copied.
2. Click Next on the Welcome page.
3. Select Remove The ACS Collector on the ACS Collector Maintenance page.
4. Click Next on the Summary page.
5. After the uninstall process completes, click Finish.

DELETING THE DATABASE

To delete the database, follow these steps:

1. Open SQL Server Management Studio.
2. Expand Databases and right-click on the name of your ACS database, which by default is named OperationsManagerAC.
3. From the context menu, select Delete.

Using Agentless Exception Monitoring

If you are like most Windows users, you have been annoyed with the service known as Dr. Watson. Whenever an application crash occurs, Dr. Watson is responsible for creating an error exception report. Once the report has been generated, you end up looking at a dialog box that asks you whether or not you would like to send the report to Microsoft. If you are like most people, you click the Don't Send button and restart the application. The only problem with doing so is that you are not informing Microsoft about the problem. Microsoft uses information it gathers from customers to determine what problems they have with their software and can then determine what course of action they want to take to correct issues.

In a corporate environment, even though you may not want to send the error reports to Microsoft, you can gather those reports for your own troubleshooting purposes. Even if you do decide to send

the reports to Microsoft in an attempt to help them discover bugs in their products, you may want to collect the data before sending it on.

Corporate Error Reporting was the first tool that Microsoft released that would allow a company to collect error reports before deciding to send the reports on to Microsoft. This stand-alone solution allowed the organization to capture the error reports and view them to determine if there was something that needed to be done in the organization. As an added benefit, each machine did not have to be configured to send the reports to Microsoft, so in a strict security environment, the systems would not have to be configured to have Internet access.

Agentless Exception Monitoring (AEM) with Operations Manager 2007 is the next incarnation of the Corporate Error Reporting (CER) tool. Although the CER tool is available as a stand-alone product, the decision to include it with Operations Manager does make sense. Operations Manager already collects and reports on data collected by systems in the organization. By integrating the two products, you have a centralized common database of monitored information.

Configuring Client Monitoring

Instead of having a separate tool to run as you do with Client Error Reporting, AEM is embedded into Operations Manager. There are two components that you need to configure to enable it, though. The first component is in the Operations Console, and the second is part of your Active Directory infrastructure or stand-alone computers.

To start configuring the client-monitoring component, run the Client Monitoring Configuration Wizard to specify the configuration options and enable client monitoring. As you are running this wizard, you are going to be specifying a shared folder where the error reports will be sent, as well as an administrative template that will be used to direct the client systems where they should deliver the error reports. The shared folder should reside on an NTFS-formatted partition. This is because the folder will need to be protected using NTFS permissions. Also make sure that the partition has enough space to hold all the report data that is forwarded. At a bare minimum, you need to have at least 2 GB of drive space.

To run the Client Monitoring Configuration Wizard:

1. Open the Operations Console.
2. Open the Administration workspace by clicking the Administration button.
3. In the Administration pane, expand Administration ➤ Device Management and select Management Servers.
4. Right-click the management server that will be responsible for collecting the error reports and select Configure Client Monitoring from the context menu.
5. When the Client Monitoring Configuration Wizard starts, click Next on the splash page.
6. On the Customer Experience Improvement Program page, select the option that fits your scenario. If you want to participate in the program select, Yes; otherwise click No.
7. If you have installed an SSL certificate on the management server, you can leave the default selection for Use Secure Socket Layer (SSL) Protocol seen in Figure 8.19. If you have not installed an SSL certificate, deselect the option.
8. The default port that is used by AEM is 51907. If you have changed the settings for the port, identify the port you are using and then click Next.

9. On the Error Collection page, seen in Figure 8.20, enter the path to the folder that will be used as the shared folder for the forwarders to send the error reports.
10. If you are collecting reports from Vista clients, select the Collect Application Error From Windows Vista-Based Or Later Clients check box.
11. If you selected the Vista client option, you will have the opportunity to choose the port the Vista clients will use, 51906 by default, as well as SSL and authentication settings.
12. The final option on this page is Organization Name. If you do not want the default option of Microsoft, enter your own organization name, and then click Next.
13. Choose Automatically Forward All Collected Errors To Microsoft if you want to send all the reports that are forwarded. Otherwise, leave the option cleared to refrain from sending them.
14. If you require more verbose error reporting, select Detailed. If you only want to collect summary information, select Basic.
15. Click Next.
16. On the Create File Share page, seen in Figure 8.21, choose an account that has permissions to create a file share. The path that was entered earlier will be shared. Click Next.
17. After the share is created, click Next.
18. On the Deploy Configuration Settings page, enter the path where you would like to save the ADM file that will be imported into a group policy, and then click Finish.

FIGURE 8.19
Configuring AEM

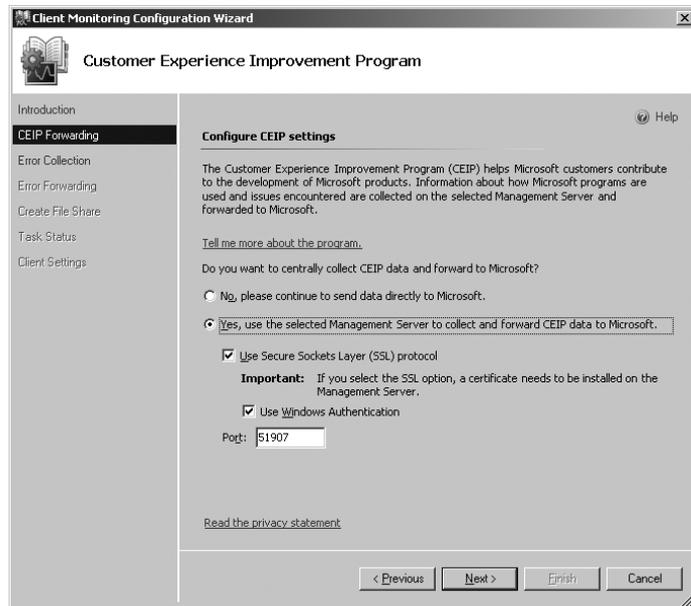


FIGURE 8.20
Error collection settings

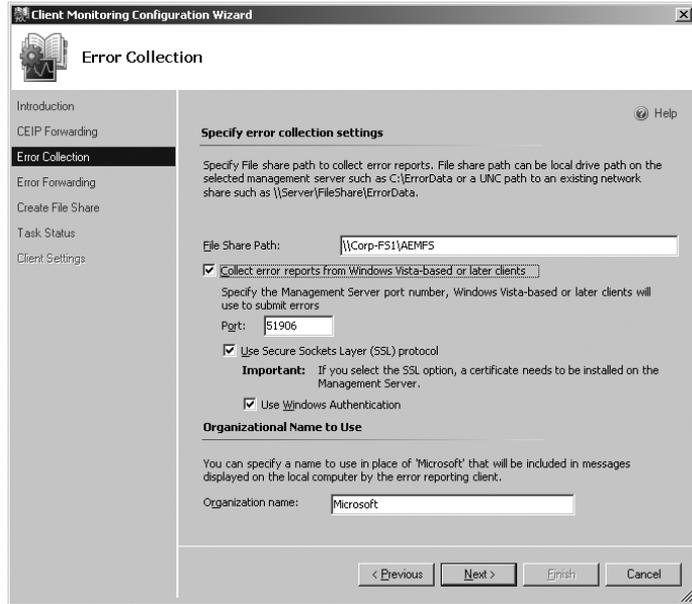
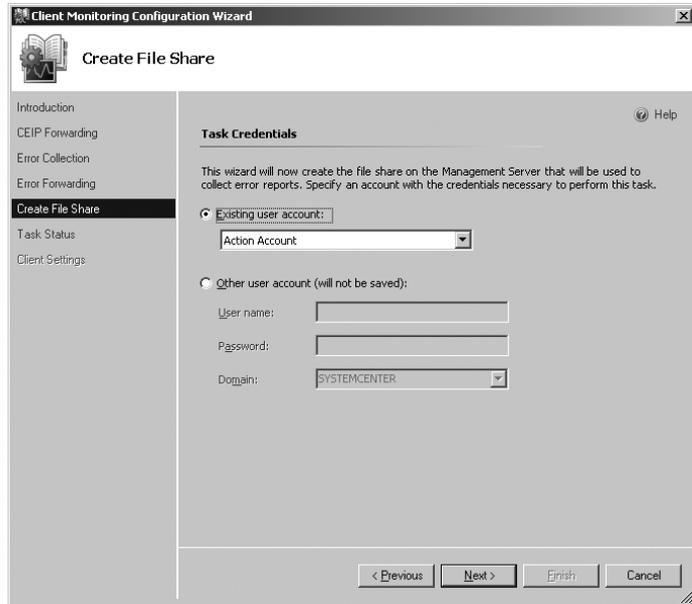


FIGURE 8.21
Account used to create the AEM share



As you work your way through the Client Monitoring Configuration Wizard, the settings that you provide are used to populate the defaults in the ADM file that you create in step 18. This ADM file will then be imported into a group policy that you will apply to all the clients that have forwarders enabled on them. The settings include the management server that you are configuring for client monitoring from step 4, the shared folder that was specified in steps 9 and 16, and the ports that are used in client communication, configured in steps 8 and 11.

Editing Group Policy

After configuring the server to accept the exceptions that are generated, make sure the client systems know that they are supposed to send the exception reports to the server. You want to collect these reports that are generated by the Dr. Watson process, and you don't want to have the users control whether the reports are delivered. Also ensure that you are collecting them on your server and not automatically sending them to Microsoft.

First, turn off the error-reporting functionality of the client systems, which seems a little counterproductive at first. Why would you want to disable error reporting? Dr. Watson is configured to notify the user that there was a problem with an application, and then ask whether the user would like to send the collected fault information to Microsoft. For stand-alone systems and home users, this might be a good solution so that they can send in the fault information and Microsoft can use the data to make their applications more reliable. In a corporate environment, you will probably want to collect this information so that you can review it before sending it on to Microsoft. This way, you can monitor to identify trends. To change the reporting behavior, turn off the default error reporting and configure your client systems to send the reports to your own collection point.

When you ran the Client Monitoring Wizard in the previous section, one piece of output was an ADM file. This file can be used in conjunction with your group policy solution to redirect the error reports to a network share that you specify. The error reports themselves are packaged into CAB files and sent to the shared directory you created. Once there, the management server that you configured as your client-monitoring system will retrieve the CAB file, process it, and store the results so that you can review them and send the file on to Microsoft if necessary.

Before you import the ADM file into your organization and start monitoring all your systems, sit back and identify which systems you really want to monitor. Then identify a test group to work with. In most organizations, there is a test bed that can be used to validate new configurations. After the test systems prove that the settings work, pilot the ADM file to a select group to make sure the test will work in your organization. One way to do this is to create a Group Policy Object (GPO) that will be used to test the pilot group. Then you can filter the permissions to apply it against the system you have identified as the pilot group.

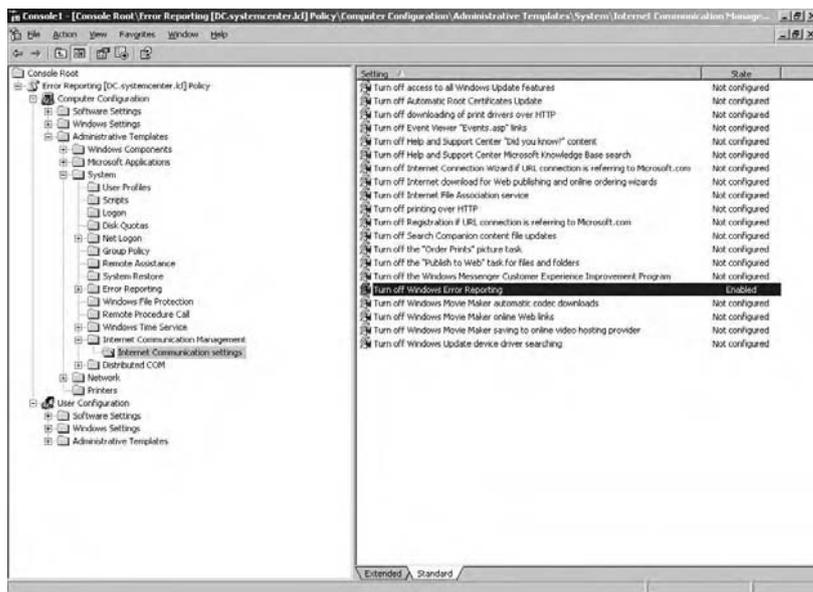
To import the ADM file and configure the group policy for Agentless Exception Monitoring, follow these steps:

1. Open Group Policy Management Console.
2. Expand *Your_Forest* > *Your_Domain* (steps to get to policy).
3. Right-click the group policy that you want to use and select Edit.
4. When the Group Policy Object editor opens, navigate to Computer Configuration > Administrative Templates > System > Internet Communication Management > Internet Communication.
5. Double-click the Turn Off Windows Error Reporting policy (shown in Figure 8.22).
6. Click the Enable button and click OK.

7. Navigate to Computer Configuration > Administrative Templates.
8. Right-click Administrative Templates and select Add/Remove Templates from the context menu.
9. When the Add/Remove Templates dialog box appears, click Add.
10. Navigate to the location where you had saved the ADM file, select the ADM file, and click Open.
11. After the ADM template appears in the Add/Remove Templates dialog box, click Close.
12. Navigate to Computer Configuration > Administrative Templates > Microsoft Applications > System Center Operations Manager (SCOM).
13. Enable the settings that are required in your environment, as described in the next section, "SCOM Client Monitoring."

The ADM template that you created and imported will control how the client machines will report errors. For the group policy to be effective, you have to enable the settings. The template includes three primary folders, and under SCOM Client Monitoring there are not only settings to configure, but also another folder called Advanced Error Reporting Settings. When you enable the settings, you will see that the details that you had input during the wizard appear as the defaults.

FIGURE 8.22
Turning off Windows Error Reporting



SCOM CLIENT MONITORING

Three settings and a subfolder appear in the SCOM Client Monitoring folder.

Configure Error Notification This setting allows you to specify that error reporting will be enabled or disabled. If you enable error reporting, two options appear that allow you to control whether the user is notified of the error. If you want to silently capture the error reports, you can leave the ShowUI check box blank. Select the check box to present the user with the error when

it arises. The DoNotDebugErrors check box allows you to control whether the built-in debuggers for applications will run when an error arises. If you want to turn off debuggers, select the check box.

Configure Error Reporting For Windows Operating Systems Older Than Windows Vista

This setting allows you to identify the shared folder on the collector where the error reports will be forwarded. In the options pane, enter the UNC path to the error reporting folder in the Corporate Upload File Path setting. The default of identifying Microsoft as the error reporting recipient can be changed to your own corporate name in the Replace Instances Of The Word 'Microsoft' With text box.

Configure Error Reporting For Windows Vista And Later Operating Systems Vista relies on a different error reporting mechanism than previous versions of Windows. When you enable this setting, you can identify the collector server in the Error_Listener text box. You can then specify the port that the listener uses, as well as whether SSL communication and Windows authentication are used. The default port specified in the Error_ListenerPort is 51906.

Advanced Error Reporting Settings The settings in the Advanced Error Reporting Settings folder control which type of error reports are forwarded.

Application Reporting Settings (All Or None) When enabled, three settings control whether error reports are sent: Report All Application Errors, which forwards the application errors from any vendor; Report All Errors In Microsoft Applications, which reports errors that arise from Microsoft applications; and Report All Errors In Windows Components, which forwards all errors in the built-in operating system application components.

Report Operating System Errors This setting allows you to specify whether errors that occur in the operating system itself will be sent to the collector.

Report Unplanned Shutdown Events When unplanned shutdown events occur, if you have this setting selected the error report generated by the operating system will be forwarded to the collector.

SCOM CLIENT MONITORING FOR OFFICE 10.0 APPLICATIONS

Two settings appear in the SCOM Client Monitoring for Office 10.0 Applications folder, which controls how Office applications report errors.

Configure Error Notification This setting allows you to specify that error reporting will be enabled or disabled. If you enable error reporting, an option appears that allows you to control whether the user is notified of the error. If you want to silently capture the error reports, leave the ShowUI check box blank. Select the check box to present the user with the error when it arises.

Configure Error Reporting For Windows Operating Systems Older Than Windows Vista

This setting allows you to identify the shared folder on the collector where the error reports will be forwarded. In the options pane, enter the UNC path to the error reporting folder in the Corporate Upload File Path setting. The default of identifying Microsoft as the error reporting recipient can be changed to your own corporate name in the Replace Instances Of The Word 'Microsoft' With text box.

SCOM CLIENT MONITORING FOR WINDOWS MEDIA PLAYER

Two settings appear in the SCOM Client Monitoring for Windows Media Player folder, which controls how Windows Media Player reports errors.

Configure Error Notification This setting allows you to specify that error reporting will be enabled or disabled. If you enable error reporting, an option appears that allows you to control whether the user is notified of the error. If you want to silently capture the error reports, you can leave the ShowUI check box blank. Select the check box to present the user with the error when it arises.

Configure Error Reporting For Windows Operating Systems Older Than Windows Vista

This setting allows you to identify the shared folder on the collector where the error reports will be forwarded. Within the options pane, the Corporate Upload File Path should be configured with the UNC path to the error reporting folder. The default of identifying Microsoft as the error reporting recipient can be changed to your own corporate name within the Replace Instances Of The Word 'Microsoft' With text box.

Agentless Exception Monitoring Views

In the Monitoring workspace you will find the views used to monitor AEM. By default, five views are created that allow you to see which systems have forwarded error reports, as well as let you check on the components of AEM to ensure everything is working as it should. Depending on the settings you enabled in the ADM file, the views will display the data that was placed in the database. For example, you won't find information in the System Crash View if you have not enabled the Report Unplanned Shutdown Events setting.

Application Error Events This state view displays any of the events that are forwarded. Since this view could become cluttered, the other views can be used to filter the events that you want to view.

Application View This state view displays the application failures that have been forwarded.

Crash Listener View This state view displays all the computers configured as listeners in the management group.

Error Group View This state view displays all error groups configured for application errors.

System Crash View This state view displays all the operating system failures that have occurred on monitored systems.

The Bottom Line

Install the components for Audit Collection Services. Audit Collection Services is made up of three components. The ACS database holds the events from the clients that you are monitoring. The ACS collector is responsible for accepting the data sent to it and delivering it to the database. The ACS forwarder resides on the client systems and forwards all security events to the collector. When installing the service, the installation wizard creates the database and installs the collector service. After the service is available, you can enable the forwarder.

Configure the ACS Forwarder. The forwarder is already installed as part of the Operations Manager client, but it is not enabled by default. You can use the `AdtAdmin.exe` utility to configure how the forwarder operates.

Configure the ACS Collector. The collector captures the data that is sent from the forwarder. After installing the collector, you can configure it using the `AdtAdmin.exe` utility. Not every event is required to be captured and sent to the database, so the collector can be configured to filter certain events. The `AdtAdmin.exe` utility can be used for this purpose also.

Configure Agentless Exception Monitoring. Agentless Exception Monitoring allows a company to centralize the application and operating system error reporting in the framework of the Operations Manager infrastructure. Instead of having a distributed error reporting solution controlled by the end user, AEM allows administrators to control how the error reports are collected and which ones are distributed to Microsoft.

Master It! Your internal auditors have informed you that you need to start saving your security events due to new government compliance regulations. Instead of saving each security log from each server, you decide to implement Audit Collection Services.

1. From the Operations Manager 2007 CD, double-click SetupOM.exe.
2. On the Setup page, click Install Audit Collection Server.
3. On the Welcome page of the Audit Collection Services Collector Setup wizard, click Next.
4. Accept the license agreement and click Next.
5. Click Create A New Database on the Database Installation Options page, and then click Next.
6. Enter **OpsMgrACS** for your ODBC data source name (DSN) on the Data Source page and then click Next.
7. Click Database Server Running Locally.
8. Leave the field blank in order to use the default database name of OperationsManagerAC. Click Next.
9. Select to use Windows authentication on the Database Authentication page.
10. On the Database Creation Options page, select the Specify Directories option and enter **c:\OpsMgrACS** as the path where you would like to store the database and log files.
11. Under Local Hour Of The Day To Perform Daily Maintenance, enter **1:00 AM**.
12. Under Number Of Days To Retain Events, leave the default of 14 days.
13. If you want to have all the data reflecting the local time for your time zone, select Local from the ACS Stored Timestamp Format list. If you want to show the time according to Greenwich Mean time, select Universal Coordinated Time (UCT). Click Next.
14. Review the setting on the Summary page and click Next to install the database and collector.
15. When the installation is complete, click Finish.

Create a monitor to restart the forwarder service if it stops:

1. Open the Operations Console.
2. Select the Authoring workspace.
3. In the Authoring pane, select Monitors from the Management Pack Objects list.
4. In the Details pane, expand Agent > Entity Health > Availability.
5. Right-click Availability and select Create Monitor > Unit Monitor.
6. When the Create A Unit Monitor wizard starts, expand Windows Services and select Basic Service Monitor.

7. Under Management Pack, click New.
8. Type **ACS Override Management Pack** in the Name field and click Next.
9. Click Create.
10. Click Next.
11. In the Name field, type **Forwarder Availability**.
12. Make sure that Monitor Target is set to Agent and that Parent Monitor is set to Availability, select the Monitor Is Enabled check box, and then click Next.
13. In the Service Name field, enter **adtagent** and click Next.
14. Verify that the Service Is Running Health State option is set to Healthy and that Service Is Not Running Health State is set to Critical. Then click Next.
15. Click to enable the Generate Alerts For This Monitor check box.
16. Make sure the Generate An Alert When pull-down selection is set to The Monitor Is In A Critical Health State.
17. If you want to automatically resolve the alert when the agent restarts, make sure the Automatically Resolve The Alert When The Monitor Returns To A Healthy State check box is selected.
18. Change the Priority option to High so that this alert will be processed quickly, and then click Create.

Configure the forwarder group and assign the domain controller forwarder to be a member:

1. Open a command prompt.
2. Type the command **Adtadmin /addgroup /group:DomainControllers**.
3. Enter the command **Adtadmin /updforwarder /forwarder:DC1 /group:DomainControllers**.

Chapter 9

Reporting

So what good is all of the data that we have been collecting if we cannot review it to see what has been captured? As events and alerts arrive at the management server, they are passed through the Operations Manager infrastructure, where they will ultimately be stored in the Operations Manager database. As you resolve alerts and fix issues that are causing health problems, the data is marked as resolved and then purged from the database as the grooming process runs.

Operations Manager Reporting allows you to run reports that you can use to determine how your managed systems are working, view the uptime of services, check out the efficiency of the organization, and create your own custom reports. Since it is based on SQL Reporting Services, you don't have to worry about learning another reporting system. Instead, you can create your reports using the same tools you use to create the reports for SQL, or you can use Report Builder, which is included with Operations Manager.

The database that you are currently using for your management group will not be the same one that you use for reporting. Instead, during the installation of reporting, you will have a new database created for this purpose. Known as the Reporting Data Warehouse, this database is used to house all the data that you will need to use for reports—which means you will not have to worry about holding all the additional data in your Operations Manager database. Using a separate database allows you to keep the Operations Manager database groomed and efficient.

In this chapter you will learn how to:

- ◆ Install and configure SQL Reporting Services
- ◆ Install Operations Manager Reporting
- ◆ Run predefined reports
- ◆ Build a Report Manager model
- ◆ Create custom reports with Report Manager

Installing Reporting Services

When you install Operations Manager, Reporting is not installed along with the other services. This is because Microsoft cannot be assured that you have installed SQL Reporting Services, and even if you have, they cannot assume which server hosts SQL Reporting Services in your organization. Consequently, they leave this service out of the initial installation.

You have probably guessed at this point that the primary prerequisite is having SQL Reporting Services installed. If you have never installed Reporting Services before, you may want to consult the Microsoft SQL website to take a look at the requirements. We will assume that you have already installed Reporting Services when you installed the SQL Server for Operations Manager, but if you have not done so, you should probably take care of that now. The actual installation is not difficult;

it is the service configuration that might take you a few minutes. Before you jump right into installing Reporting Services, make sure you have taken the time to plan your installation. As you install, you will be making changes to existing SQL Reporting Services, so make sure that you do not adversely affect the existing reporting services that you are providing for other departments. When you install Operations Manager Reporting, the account information that you use when accessing the database is used to modify the access for SQL Reporting Services. If you had previously configured SQL Reporting Services to access the database, you may find that the original settings you had applied are no longer viable.

One other thing to consider as you move forward: if there is a failure during the installation of Operations Manager Reporting, you could potentially put the entire Reporting Services in an unusable state. Your best bet is to install SQL Reporting Services for sole use by Operations Manager Reporting. Although in some organizations that may not be a viable solution, you will at least know that you will not be adversely affecting other services.

Installing SQL Reporting Services

When we installed the SQL Server that supports our Operations Manager infrastructure, we installed the database, analysis services, and client components but did not install Reporting Services. To use Operations Manager Reporting, SQL Reporting Services must be installed as it is the underlying component used to generate the reports.

Before starting up the install, make sure you have the accounts that you will use planned out. Just as we planned the accounts back during the installation of Operations Manager and its SQL Server, we must make sure that the accounts have just the level of rights and permissions required to perform their duties—but not too much power in case they are compromised.

You will also need to locate your installation media, or navigate to the location on the network where you have stored the installation files. Once you have the installation files, run the `setup.exe` program located in the Servers directory. The initialization screen appears; then the Installing Prerequisites page, as shown in Figure 9.1, appears. Once you've finished, click Next to see the splash screen in Figure 9.2. Clicking Next once again brings you to the System Configuration Check page, shown in Figure 9.3.

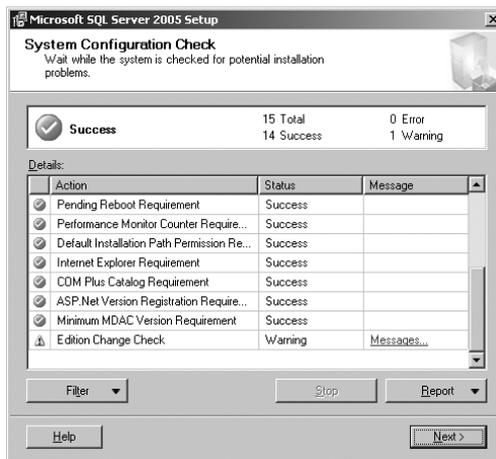
FIGURE 9.1
Installing prerequisites



FIGURE 9.2
Splash screen



FIGURE 9.3
System Configuration
Check page



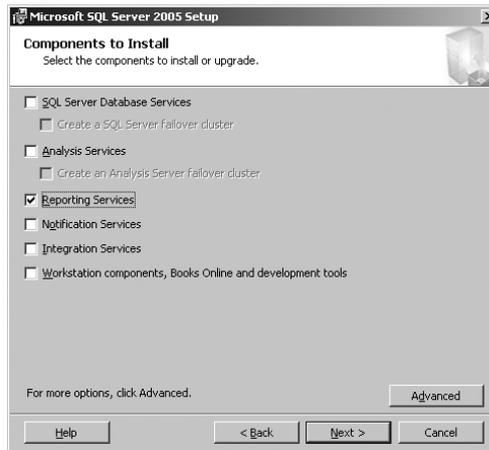
If you installed SQL Server 2005 and then installed the service pack and hotfixes, you will see the same thing that appears in Figure 9.3: the Edition Change Check warning. After installing SQL Reporting Services, you should install the appropriate service pack and all of the required hotfixes to bring Reporting Services up to the same build level as the rest of the SQL Server services. Click Next, and the installer prepares the installation files and then presents the Registration Information page shown in Figure 9.4. Enter your identifying information and click Next to continue.

When you are presented with the Components To Install page, select only the Reporting Services option, as shown in Figure 9.5. Many of the Microsoft installation wizards require that you also select the options that are already installed; otherwise it is assumed that you are going to uninstall those options. With SQL Server 2005, any modifications to existing services are made through the Add/Remove Software application in Control Panel. Any new component added to the installation has to be installed through the setup wizard. Thus, selecting only the Reporting Services option notifies the setup program to add only this one service to the existing installation.

FIGURE 9.4
Registration
Information page



FIGURE 9.5
Choosing to install
Reporting Services



Unless you have installed an additional instance of SQL Server to support your Operations Manager environment, leave the default selection of Default Instance selected, as shown in Figure 9.6. Otherwise, select Named Instance and enter the instance that identifies the installation of SQL Server in the `server\instance_name` format. Alternatively, you can click the Installed Instances button to discover the installed instances and select the appropriate instance. Click Next to move on.

In Figure 9.7, notice that you can define the service account that SQL Reporting Services will run under when it starts. You can use the same account used by the other SQL services, or you can specify a new domain account here—assuming you want to use a domain account for the service account. You also have the option to use a built-in system account, which you can select from the Use The Built-in System Account drop-down.

FIGURE 9.6
Defining the instance
to use

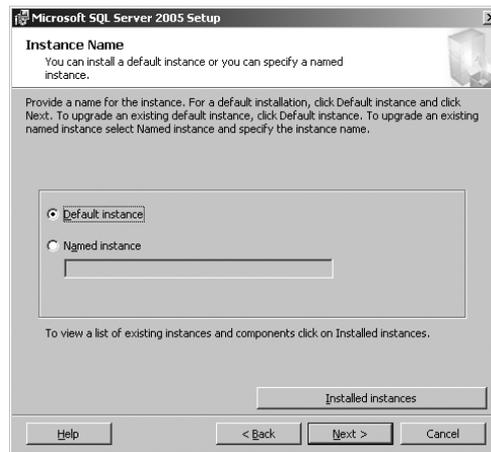
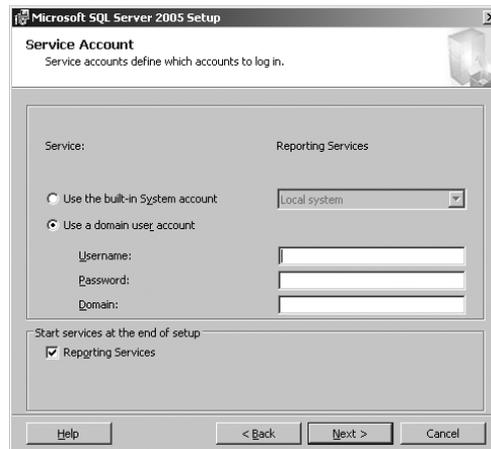


FIGURE 9.7
Selecting the service
account



When you move to the next page, Report Server Installation Options, as shown in Figure 9.8, you will not have the ability to automatically install and configure Report Server unless you use the same service account for Report Server as you did for the SQL Server service account. Click Next and you are presented with the Ready To Install summary page, shown in Figure 9.9. Notice that the service pack warning appears here, letting you know that the prerequisite check had determined that the SQL installation is at a different service pack level than the Report Server installation. If everything looks right, click the Install button to start the Report Server installation. Once the installation is complete, you will see the final Setup Progress page, shown in Figure 9.10, letting you know that the installation either succeeded or failed.

FIGURE 9.8
Report Server
Installation Options
page

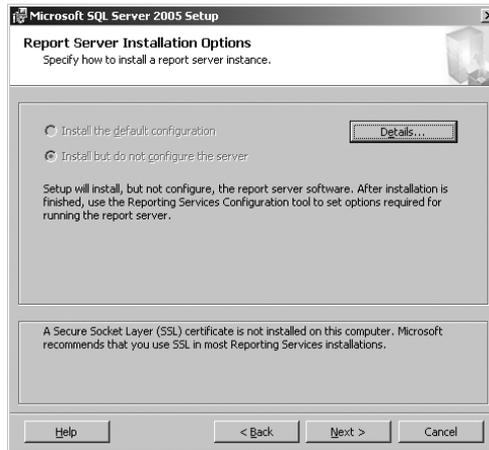


FIGURE 9.9
Ready To Install
summary page

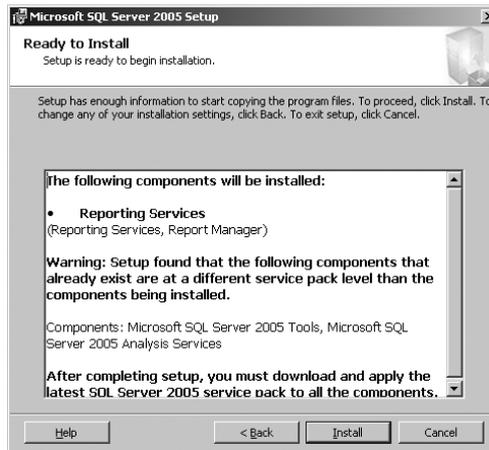
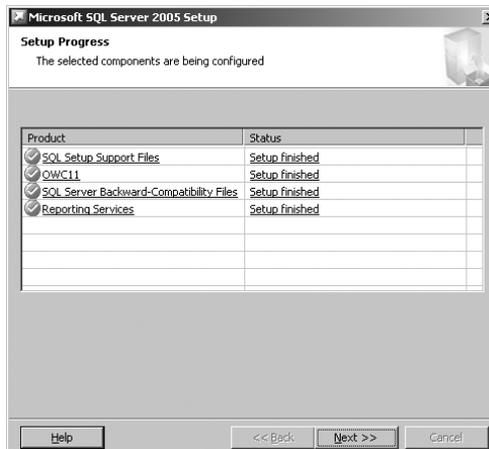


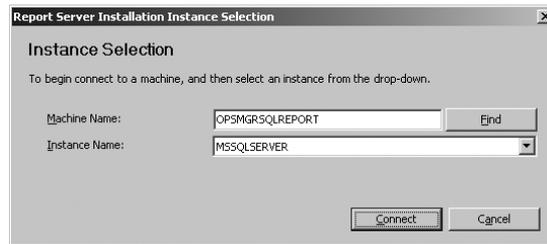
FIGURE 9.10
Setup Progress page



Configuring SQL Report Server

After installing Reporting Services and applying the appropriate service pack and hotfixes, you will need to configure the service so that it can be used to generate the Operations Manager reports. You will find the configuration utility in the Start Menu group for SQL Server 2005. Select Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > Reporting Services Configuration to open the Report Server Installation Instance Selection dialog box, shown in Figure 9.11, where you can select the server and instance that host Reporting Services.

FIGURE 9.11
Selecting the Reporting Services instance to configure



When the Configure Report Server page appears, notice that some of the configuration options are not set, as shown in Figure 9.12. Anything that has a red X beside it has to be configured before Reporting Services will function. Those items with a green check mark are already configured and working, while a yellow triangle with an exclamation mark denotes a recommended component that should be configured but will not adversely affect the installation. Finally, any component with a blue circle that contains an exclamation mark is an optional component. The only one that is considered optional by the configuration page is Encryption Keys. In most organizations this will not be considered an optional setting, but Microsoft does not force you to consider working with this option.

FIGURE 9.12
Configure Report Server page



Selecting the Report Server Virtual Directory option allows you to configure the virtual directory that will be used in Internet Information Server (IIS) for SQL Server reporting to use when generating reports. The configuration page, shown in Figure 9.13, allows you to define the virtual directory name and the website that will host the virtual directory. Clicking the New button opens a dialog box that you can use to define these settings, as shown in Figure 9.14. Use the drop-down to select the website name and then type the virtual directory name in the text box. Then click OK.

Figure 9.15 shows the settings that can be affected when you click the Report Manager Virtual Settings link. It is here that you can define the virtual directory that will be used when a user attempts to connect and run a report. The steps to configure this virtual directory are the same as those for creating the Reporting Services virtual directory in the previous paragraph. Make sure that you specify a different directory than you chose for Reporting Service.

Clicking the Windows Service Identity link presents you with the page you see in Figure 9.16. The account that appears in the Service Account text box should be the account that you selected when you installed Report Server. You can change the account here by entering a new domain account or using the Built-in Account drop-down to specify a system account. Once you click Apply, the account will be changed.

The Web Service Identity page allows you to define the identity that Reporting Services will run under in IIS. There are two options to work with from this page, as shown in Figure 9.17. The first is the IIS 5.0 service account. If you have IIS 5.0 as the web server, you will have to use the account that ASP.NET uses as its service account. The second option is IIS 6.0 and later. IIS 6.0 allows you to specify different application pools for each of the web applications. To define the application pool used for Report Server and Report Manager, either select the application pool you want to use from the appropriate drop-down, or click New and create a new application pool.

FIGURE 9.13
Report Server
Virtual Directory
Settings page

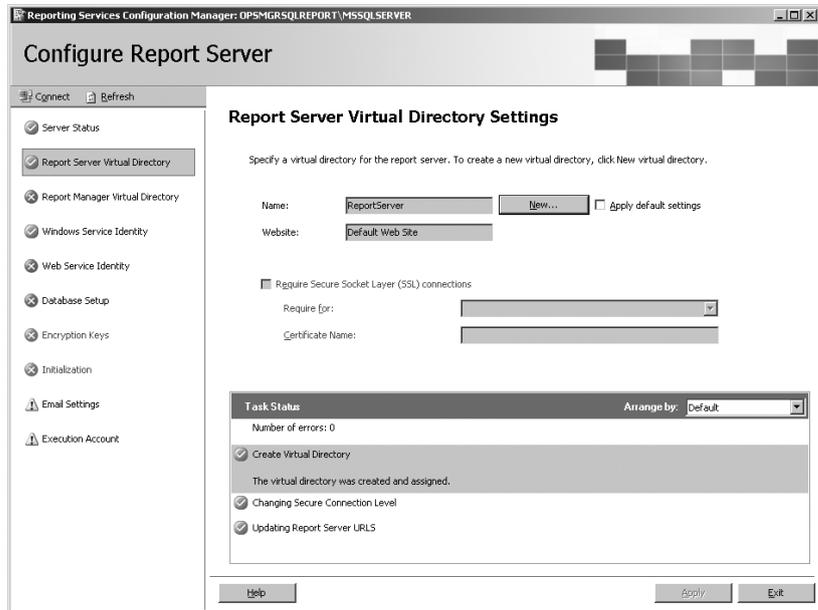


FIGURE 9.14
Creating the virtual directory for Reporting Services



FIGURE 9.15
Configuring the report directory

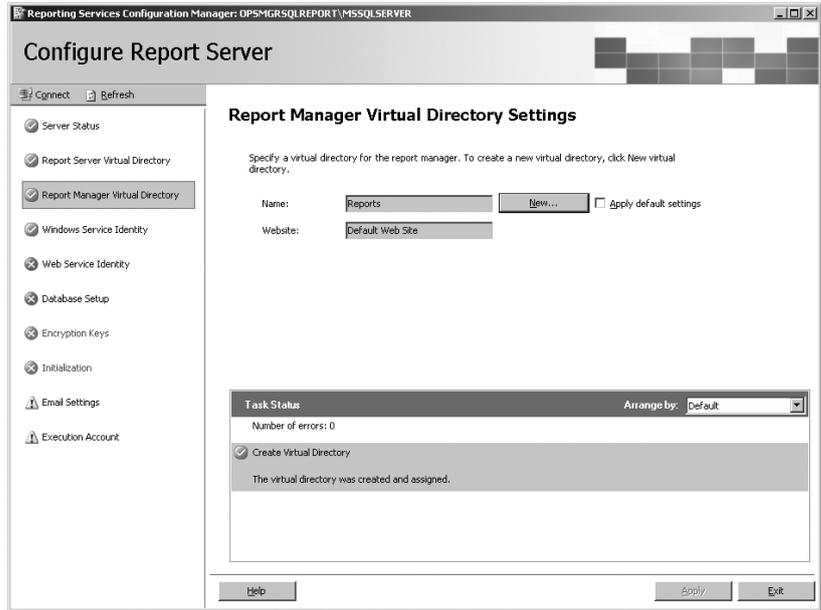


FIGURE 9.16
Windows Service Identity page

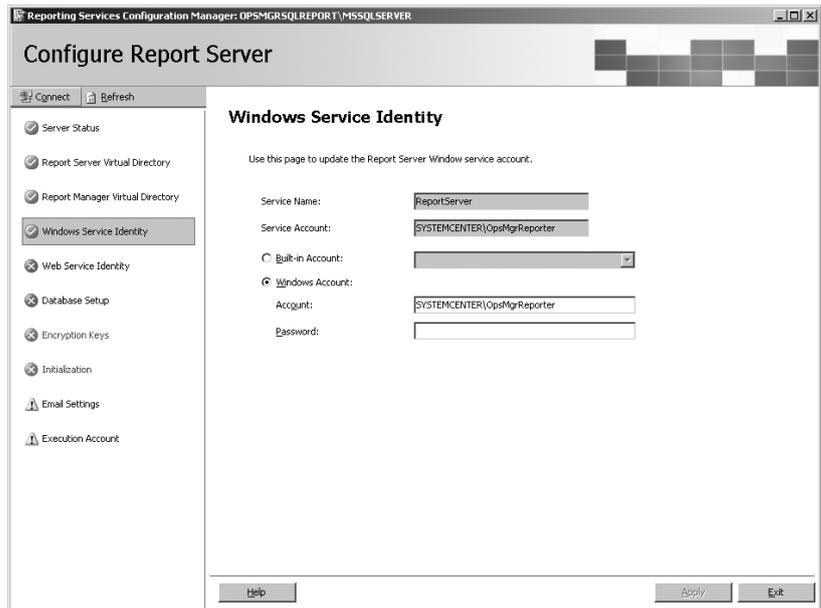


FIGURE 9.17
Configuring the
web identity



When creating a new application pool, you are presented with the dialog box shown in Figure 9.18. Here you enter a name for the application pool in the Application Pool Name text box and then identify the account that will be used to authenticate the application pool when the application pool starts. As with the other identities we have been discussing, the account can either be a domain account or a built-in system account. Once the application pool has been identified, click OK, and then on the Web Service Identity page, click the Apply button to have the application pools created and the identities associated. When the process has completed, you will see the Tasks Status at the bottom of the page, as shown in Figure 9.19.

FIGURE 9.18
Creating a new
application pool



FIGURE 9.19
Results of creating the application pools and setting the identities



The Database Connection page is used to either make a connection to an existing database or create a new database. To start, make a connection to the server running SQL Server by clicking the Connect button and entering the server name and an account that has administrative permissions to the SQL Server in the SQL Server Connection Dialog, as shown in Figure 9.20. Once connected, you can click the New button to create a new database, or you can use the drop-down to select the database you want to use. If you click the New button, you will see the dialog box shown in Figure 9.21. Here you specify the credentials you need to use to create a database, as well as the name of the Report Server database and the language to be used. Once you click OK, you must click the Apply button on the Database Connection page. Once the database is created, you will see the results in the Task Status pane, as shown in Figure 9.22.

FIGURE 9.20
Connecting to the SQL Server

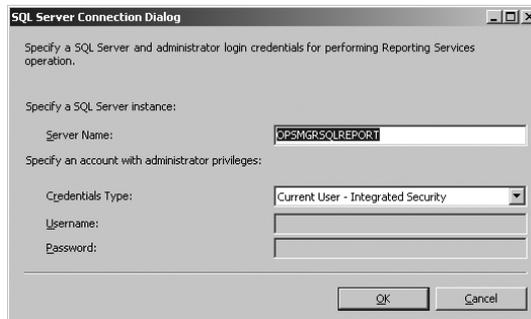


FIGURE 9.21
Creating a new database

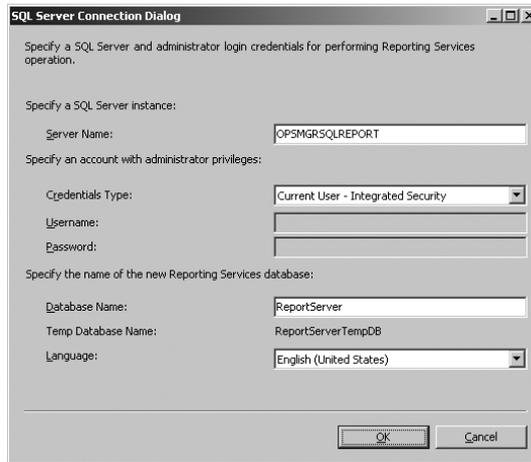
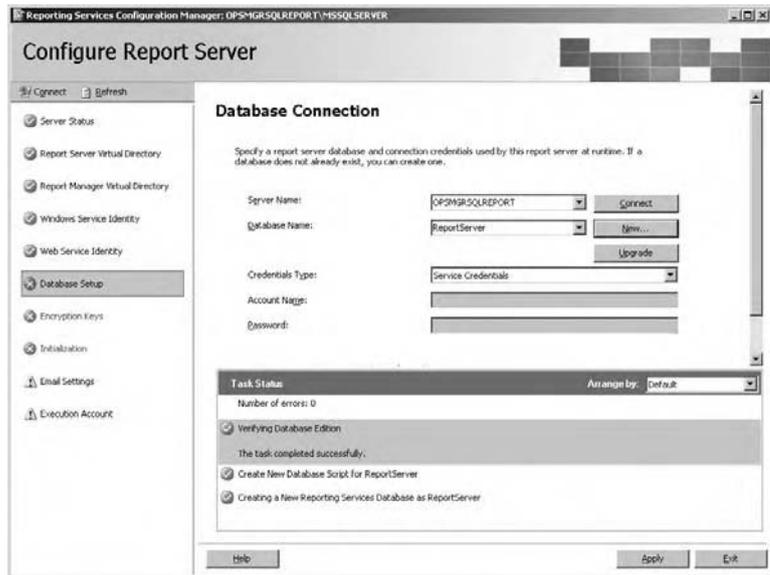
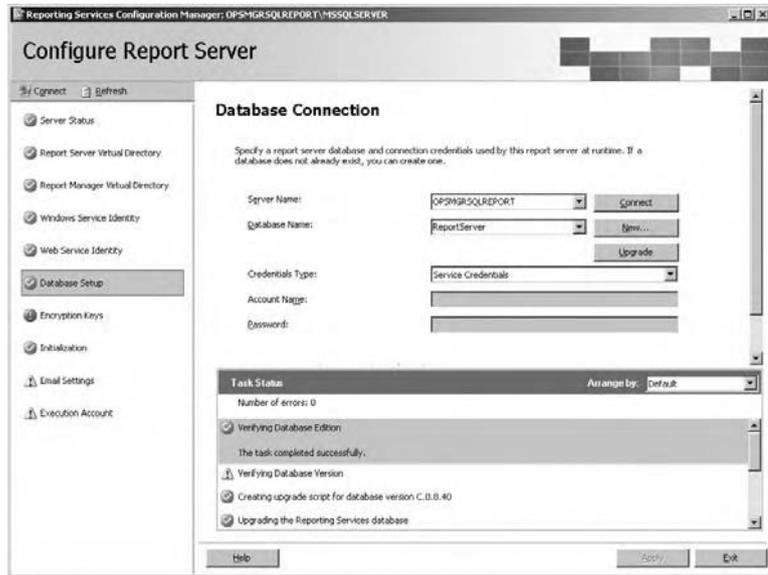


FIGURE 9.22
Results of creating the database



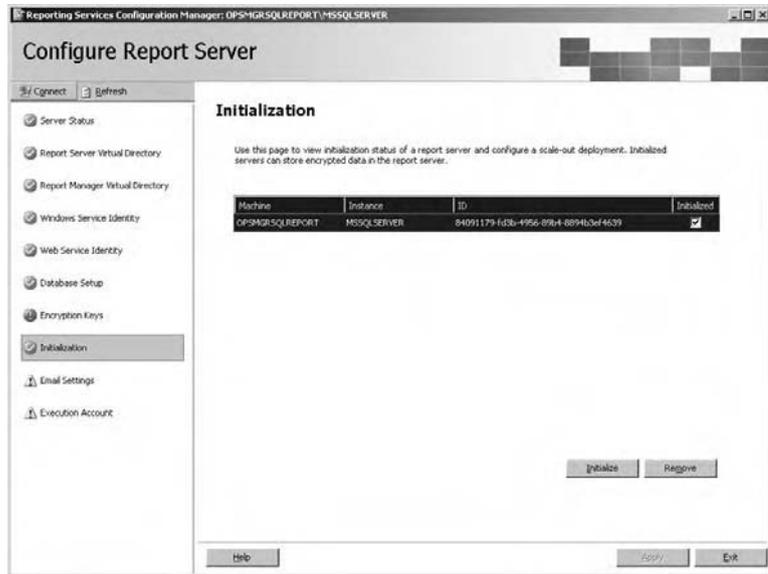
If you did select an existing database, you will be prompted as to whether you want to upgrade the database. If you click Yes, you will then associate that account used on this page with the database. The Task Status pane, shown in Figure 9.23, reflects the upgrade. Notice the credentials on this page. You can either associate the service account used for Reporting Services to have access to the database, or you can select a new account by entering it on this page and clicking Apply.

FIGURE 9.23
 Upgrading an existing
 database



Moving down to the Initialization page, shown in Figure 9.24, you can specify which Report Servers will work in conjunction with one another in a server farm. Your server should be shown as initialized. If not, click the Initialize button.

FIGURE 9.24
 Initializing the server



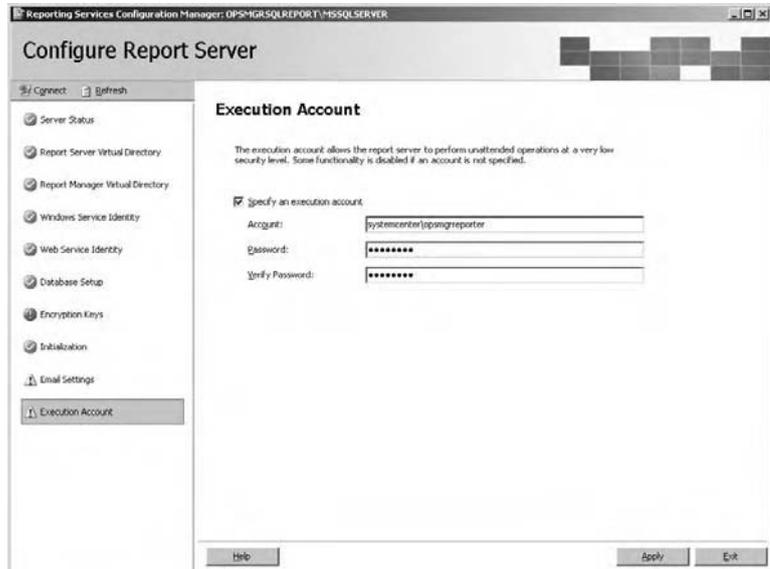
The E-mail Settings page, shown in Figure 9.25, allows you to configure the email server that will be used to send subscription emails. Subscriptions can be created by users of Reporting Services so that they can receive their reports via email. You do not have to configure this option, but some users may appreciate the capabilities.

FIGURE 9.25
E-mail Settings page



The last of the configuration pages that we will look at, with the exception of the optional page for encryption keys, is the Execution Account page, shown in Figure 9.26. Here you configure the account that will be used to run tasks, such as the subscription execution. This account should be a low-level access account and typically will not need to have any additional rights.

FIGURE 9.26
Execution Account settings



The one optional page that appears, Encryption Key, shown in Figure 9.27, is used to back up, restore, and change the symmetric encryption key that is used when accessing the database. After you have configured everything else, navigate to this page and click the Backup button. The dialog box that appears, shown in Figure 9.28, is used to define a password will be needed to restore the keys if necessary. You also have a text box and a browse button that can be used to specify the directory where the keys will be saved. If you back them up to a directory on a local hard drive, make sure that you copy them to a safe location. You will need to restore the keys if you ever have to rebuild the Report Server.

FIGURE 9.27
Encryption Keys page



FIGURE 9.28
Choosing the
password and backup
directory



Installing Operations Manager Reporting

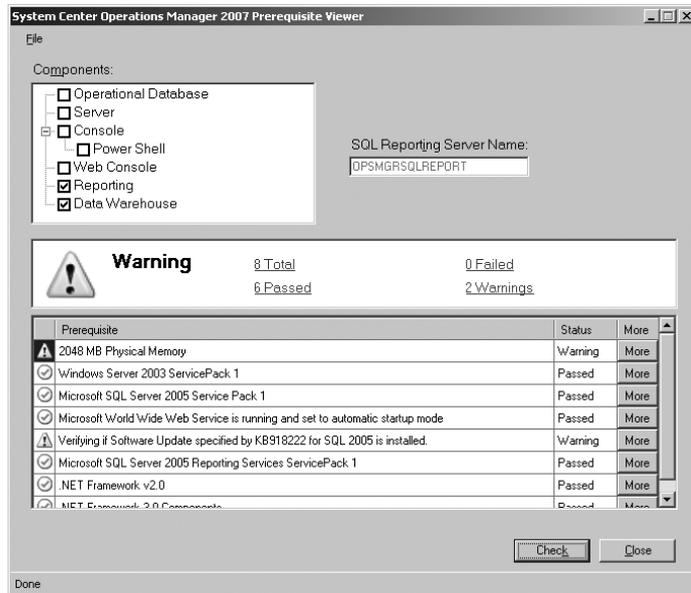
Up to this point in the chapter, we have discussed installing the dependent SQL service for Operations Manager Reporting. Once you have the dependencies installed, you will be ready to perform the rest of the install. Like SQL Report Server, installing Operations Manager Reporting is not difficult but does require some time spent configuring the options.

Make sure you have configured the account used by the SDK and Config services to authenticate as a domain account. If you have configured the account as a system account, you could potentially cause problems with the installation. You may discover that the SDK and Config services will not properly access the database during installation, and you will not find any reports available after you have completed the installation.

Make sure that you have installed the hotfixes as mentioned in Knowledge Base article 918222. There are three hotfixes that you will need to download and apply: one for the SQL Server, one for Analysis Services, and one for Reporting Services. Note that when you install the hotfixes for SQL Server and Analysis Services, they recognize if you have those services installed on clustered hardware and will patch all nodes, but the Reporting Services hotfix has to be applied to each node individually.

Once everything is in place, you are ready to start installing. Begin by locating the Operations Manager installation source that you used when installing the Operations Manager database and management servers. When you start SetupOM.exe, you are presented with the same splash page that you saw when installing the other Operations Manager services. Click Next, and you are presented with the EULA and the Product Registration pages in turn. After you enter your product key, a prerequisite check is performed. If you have missed any of the prerequisites, you will see a warning such as the one in Figure 9.29.

FIGURE 9.29
Warning that you are missing the 918222 hotfixes



As you take a look at the Custom Setup page shown in Figure 9.30, you will see two options that you can work with: Data Warehouse and Reporting Server. If the SQL Server that hosts the data warehouse is going to be different than the Reporting Services server, you can install the two options independently of one another. After you choose the options to install, click Next.

The Connect To The Root Management Server page appears, as shown in Figure 9.31, requesting the name of the RMS. Type the name of the RMS in the text box and click Next to move to the SQL Server Database Instance page, shown in Figure 9.32. Choose the database instance from the drop-down. Once you have identified the database instance that you want to use, click Next to define the database name.

FIGURE 9.30
 Selecting the components to install

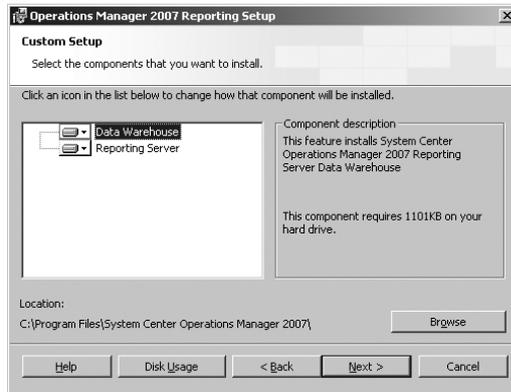


FIGURE 9.31
 Identifying the Root Management Server

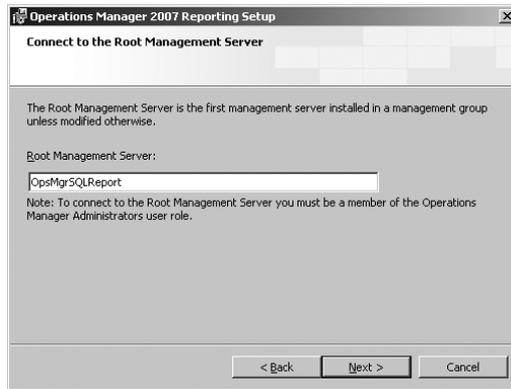
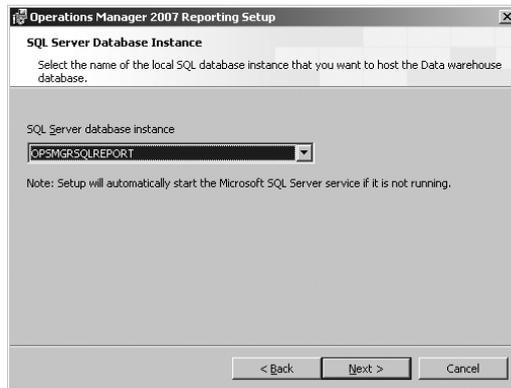


FIGURE 9.32
 Choosing the database instance



The default database name for the data warehouse is `OperationsManagerDW`, although you can enter whatever name you need to use in the Database Name field of the Database and Log File Options. This page, shown in Figure 9.33, is the same page that we discussed when installing the Operations Manager database. You must specify the initial database size and location of the database and log files in order to make the system as efficient as possible. Click the Advanced button to change the path to the files and click Next when finished.

After identifying the database settings, you are asked to identify the SQL Server Reporting Services server by selecting it from the drop-down shown in Figure 9.34. Clicking Next will take you to the Data Warehouse Write Account page, where you define the domain account that will be used to write data to the data warehouse, as shown in Figure 9.35. Clicking Next goes to the Data Reader Account page, which looks very similar, but this page is used to specify the account used to read data from the data warehouse. The account that you enter on this page, shown in Figure 9.36, is used whenever a query is sent to the data warehouse, and is also used by the application pool to connect to the Root Management Server. Then clicking Next will present the completion page, where you can click Finish to install.

FIGURE 9.33
Database and log file
locations

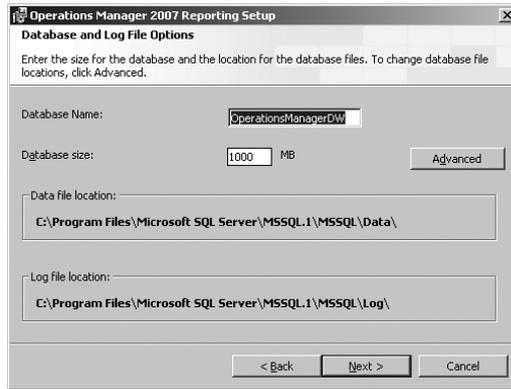


FIGURE 9.34
Choosing the Report-
ing Services server

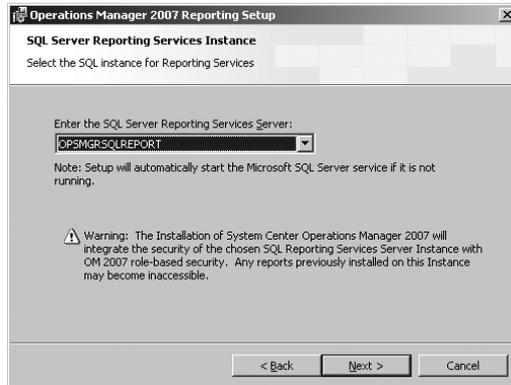


FIGURE 9.35
Defining the data
warehouse write
account

Operations Manager 2007 Reporting Setup
Data Warehouse Write Account

The account provided will be assigned write permissions on the data warehouse and read permissions on the Operations database.

User Account:
 Password:
 Domain: SYSTEMCENTER

< Back Next > Cancel

FIGURE 9.36
Specifying the data
warehouse reader
account

Operations Manager 2007 Reporting Setup
Data Reader Account

The Data Reader Account will be used to define what user SQL Reporting Services uses to execute queries against the Operations Manager reporting data warehouse. This account is also used for the SQL Reporting Services IIS Application Pool account to connect to the Principal Management Server.

Setup will add this account to the Report Administrator User Profile.

User Account:
 Password:
 Domain: SYSTEMCENTER

< Back Next > Cancel



Real World Scenario

FRUSTRATING BEHAVIOR

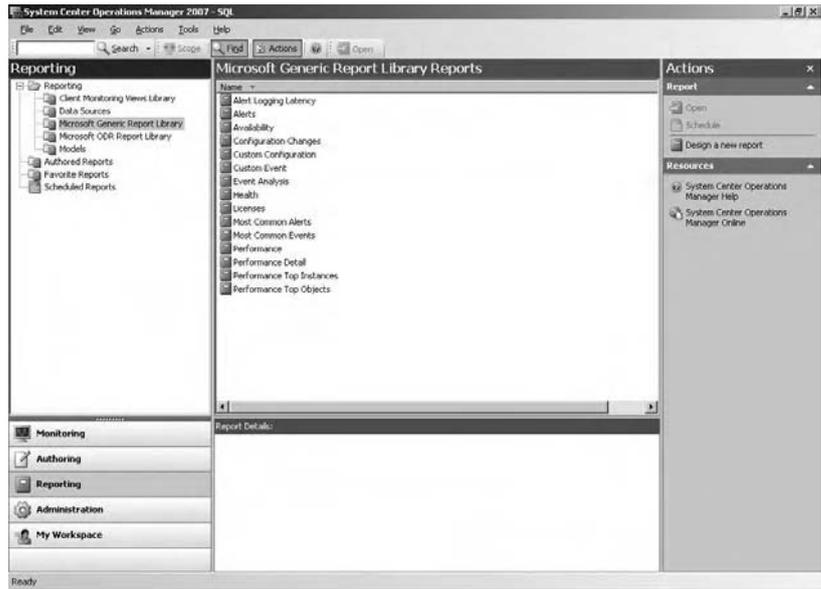
Once you install the Reporting component and open the Reporting workspace, you may find that you do not have any reports to work with. As with any installation, check the setup log files and the event logs to see if there is anything out of the ordinary. Chances are you will not have to worry about not having anything to view yet. Operations Manager has to process the reporting data from the management packs. As the reporting data is processed, the data model used by SQL Reporting Services has to be built before it can be used. If you wait, you will find reports available from the Reporting workspace. Our experience has shown that this can take as few as 5 minutes and as long as a couple of hours.

The same thing holds true for any management packs that are imported. You will not see the reports become available immediately, but if you are patient, the reports will be ready in a short amount of time.

Using the Reporting Workspace

In Figure 9.37, you will see that the Operations Console has been opened to the Reporting workspace and the Microsoft Generic Report Library is opened. As you can see, there are already some useful reports available for you to use. These reports can be run manually from this workspace, or you can open a web browser and connect to the SQL Reporting Services website and run them from there.

FIGURE 9.37
Reporting workspace



Once you have decided which report you need to run, double-click it to launch another window. Each report that you launch either will generate a page with the results of a query or will require information. The reports that require data to be entered will display an intermediary page with the available options for the report. As you can see in Figure 9.38, the Management Pack ODR Report simply displays results without prompting for input. Figure 9.39 shows the page that appears when you choose to run the Most Common Alerts report.

The prompts that appear on this page, shown in Figure 9.40, allow you to configure the query that is sent to the data warehouse. For this report, you have the ability to configure the time frame of the report by selecting the date range in the From and To drop-downs, and specifying the time by using the scroll wheels. You can change the time zone as well by choosing a different time zone from the drop-down. To the right of these options is a list of management packs that you can select to query upon. If you only want to look at the most common alerts from a specific management pack, deselect all others from this list and leave only the one you are curious about selected. It is completely up to you as to how you select the management packs from the list. Once you have chosen all of the parameters, click the Run button to generate the report.

As you can see from those two examples, each of the reports can have different parameters to work with. Take the time to look over the reports that are available when you add management packs.

FIGURE 9.38
Results of running
Management Pack
ODR Report

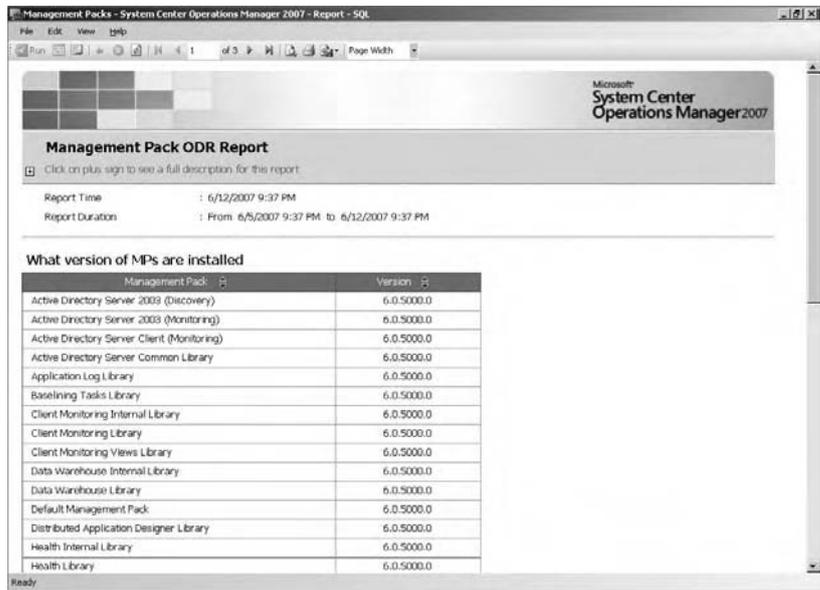


FIGURE 9.39
Most Common Alerts
parameters page

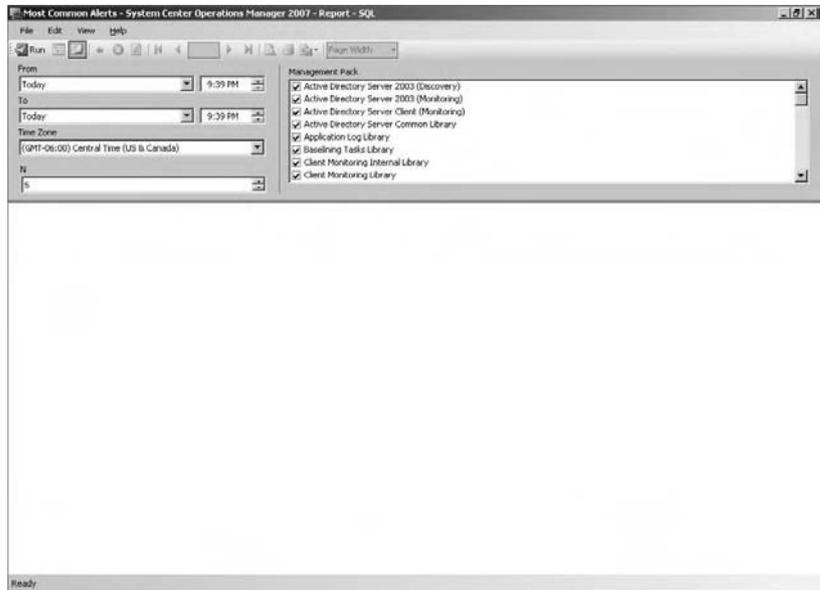
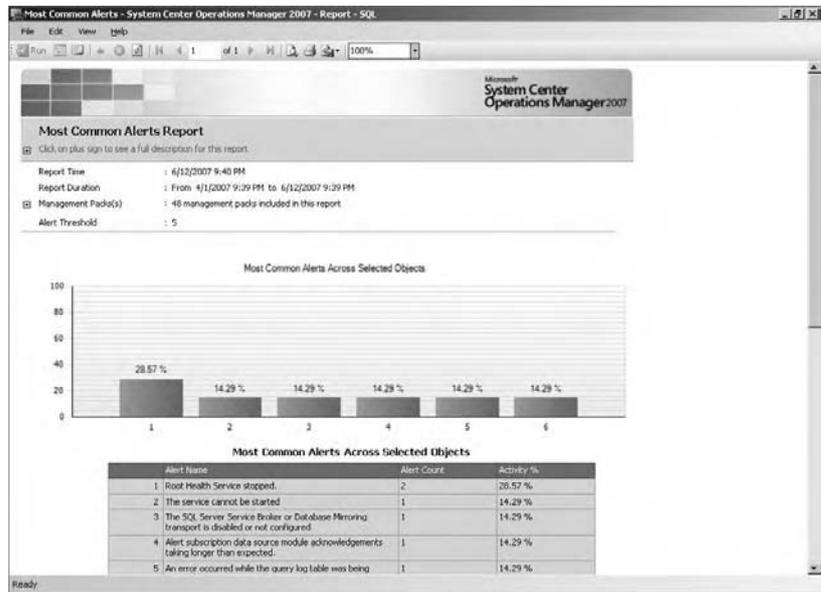


FIGURE 9.40
Parameters for the
Most Common Alerts
report



Predefined Reports

Most management packs that you import will include reports that will appear in the Reporting workspace. These are the reports that have been designed and built by the providers of the management packs, based on what those vendors deem important to administrators. Some of these reports will detail the availability of the servers or services while others will show downtime and alert resolution.

Not readily evident are some of the options you have when running reports. Of course, you know by now that you can open the Reporting workspace and access any of the reports. You can also access certain reports through the Actions pane. In Figure 9.41, the Operations Console is open to the Monitoring workspace. With a computer name selected, reports appear in the Windows Computer Reporting windows in the Actions pane. If you click one of the reports, the computer name that is selected will be passed to the report, as shown in Figure 9.42. You still have to supply some other parameters, such as the time frame you wish to report on, but at least some of the information is already supplied for you.

Reports that appear in the Actions pane are considered *targeted* reports because they already have an object or group that can be associated with the report. The reports that you will find in the Reporting workspace are considered *generic* reports because they can be associated with any object or group. Of course, some of the objects that you can add on a report parameter page may not actually return any data, so you will have to test to see which reports work with each group or object.

FIGURE 9.41
 Selecting a computer
 and the reports that
 appear

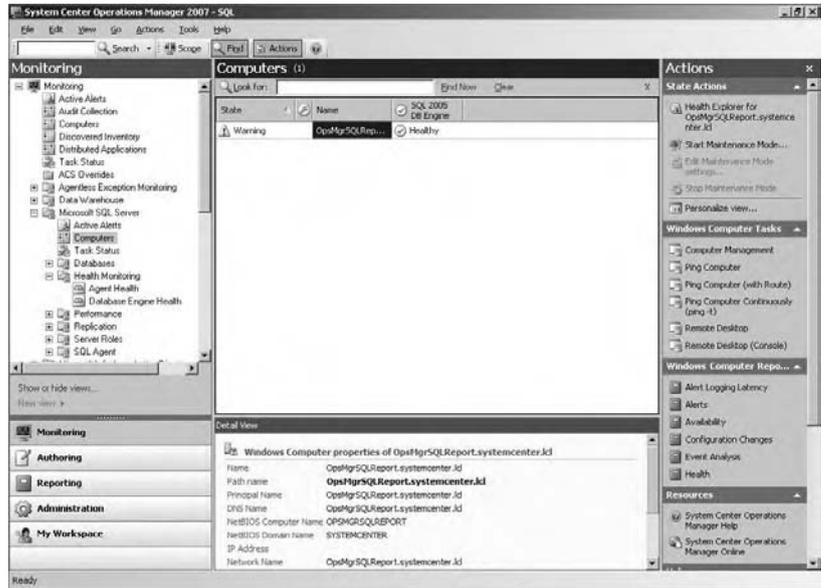
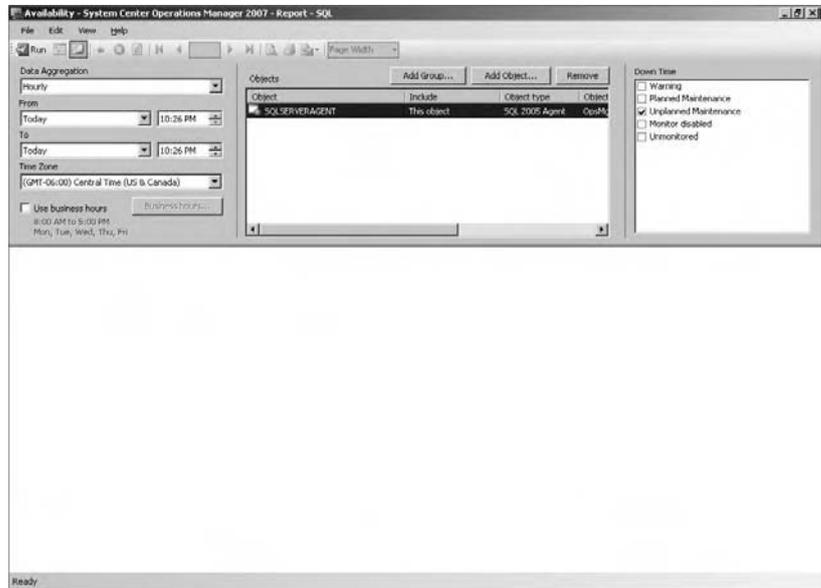


FIGURE 9.42
 The report parameters
 page with the selected
 computername already
 supplied



IMPORTING AUDIT COLLECTION SERVICES REPORTS

At the end of this chapter, there is a section named “Master It!” that has a step-by-step set of instructions that you can use when importing the Audit Collection Services reports. For the sake of simplicity, and to reduce the number of virtual machines that we have running at any one time, we have chosen to import the ACS reports into an existing Reporting Services server that already hosts the Operations Manager Reporting. While you can do this, it is not a recommended practice. Your auditing data should be kept secure and unavailable to everyone except for the auditors who need to review the information. When you import them into the existing database, you may be allowing users who do not need to see the auditing data the ability to view it. Take the time to build another separate server to provide ACS reporting.

Creating New Reports

It is not very likely that you will want to settle for using only the predefined reports. Although they will give you some valid information, you will probably determine rather quickly that you want to know more than they give you. However, learning how to create a report using the traditional Reporting Services tools, which includes Visual Studio tool, can seem like a daunting task. Luckily, SQL Server 2005 Reporting Services includes a tool called Report Builder.

Report Builder is a simplified interface that allows you to create rudimentary reports that are customized to your needs. With its drag-and-drop interface and easy editing tools, you can create a report in a matter of minutes. The hardest part of Report Builder is determining the objects that you need to add to the report.

Before you jump into creating reports, you must perform some preliminary actions so that Report Builder understands the database and objects that you will be using in the reports you create. Report Builder allows you to create reports based on a model that it builds from the database. Once the model has been created, you can then open Report Builder and start generating your own report definitions.

CONFIGURING A REPORT MODEL FOR REPORT BUILDER

To create a model, you use the SQL Server Business Intelligence Development Studio. Although creating reporting models is not the only thing the Development Studio will do, it makes the task of creating a model for building reports appear easy.

Start off by selecting Start > All Programs > Microsoft SQL Server 2005 > SQL Server Business Intelligence Development Studio. You will see the tool open and present you with the workspace, as shown in Figure 9.43. On this screen you start building a new project that will be used to create the model for your reports. Click New > Project to get started.

The New Project dialog box appears, where you define the project that you are going to work on. For our reporting model we want to select Report Model Project from the list of Visual Studio Installed Templates. Once you have selected the project type, name and define the location of the files at the bottom of the page, as shown in Figure 9.44. Make sure that you name the solution something that is going to be understood by anyone who needs to work with this in the future. It is also a good idea to create a separate directory for the solution so that you do not interfere with any other solutions that may already exist.

After the project files have been created, you will see the initial options for your solution in Solution Explorer, as shown in Figure 9.45. Here you define how to access the database and which database will be used when building the model. Right-click on the Data Sources folder and select Add New Data Source. You will be presented with a splash page welcoming you to the wizard. Click Next to move to the Select How To Define The Connection page, as shown in Figure 9.46.

FIGURE 9.43
SQL Server 2005
Business Intelligence
Development Studio

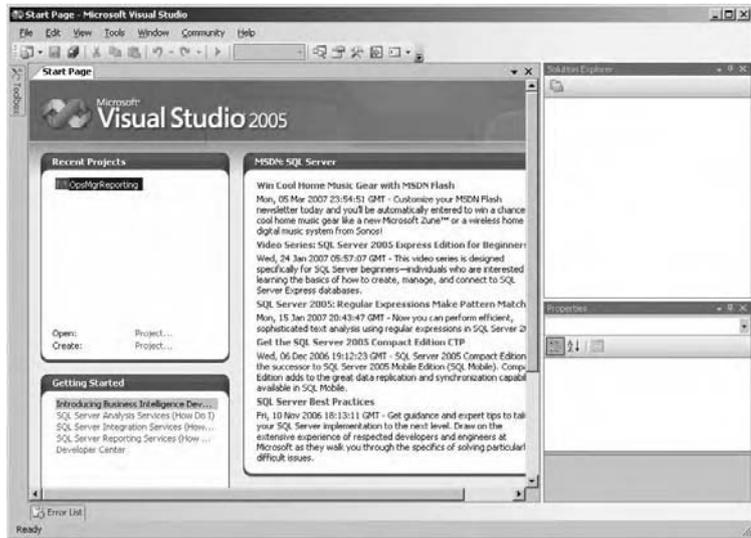


FIGURE 9.44
Creating a report
model project

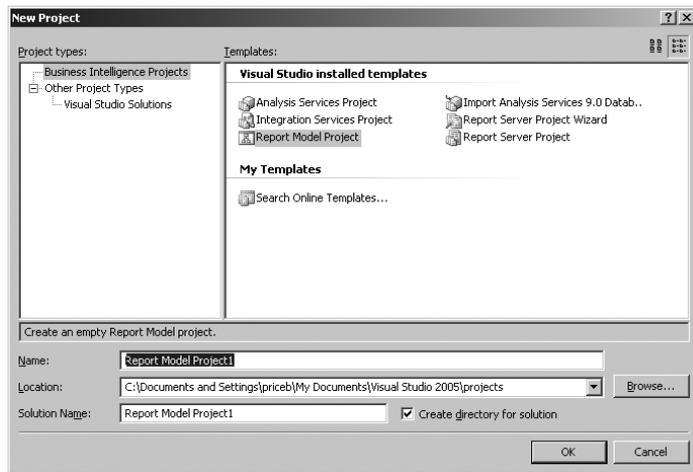
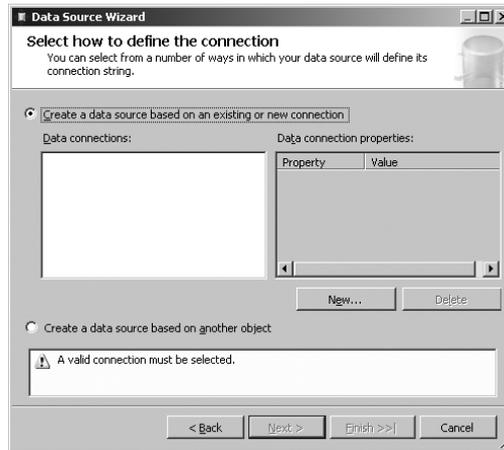


FIGURE 9.45
Solution Explorer



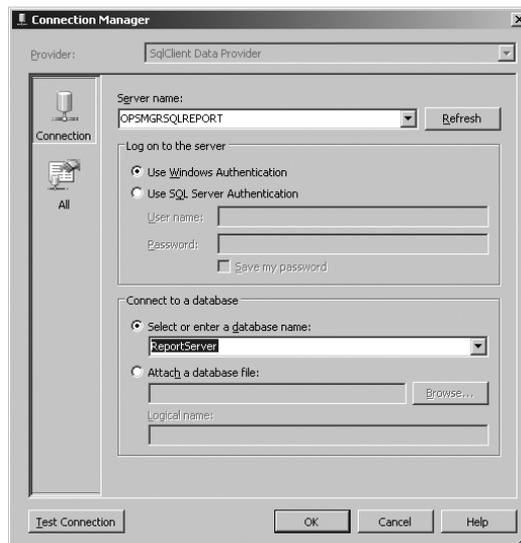
FIGURE 9.46
Defining the connection using the Data Source Wizard



Since this is the first data source we are configuring, click the New button and define the connection using the Connection Manager shown in Figure 9.47. Using this page, you select the server that is hosting the database, the database instance, and the authentication that is used. If you do not see your server listed when you click the Server Name drop-down, enter the name manually and then click Refresh. You should then see the databases that are available on the server. For our model we want to use the data warehouse database, which is named OperationsManagerDW by default. If you changed the name when you installed Operations Manager Reporting, make sure you select the right database.

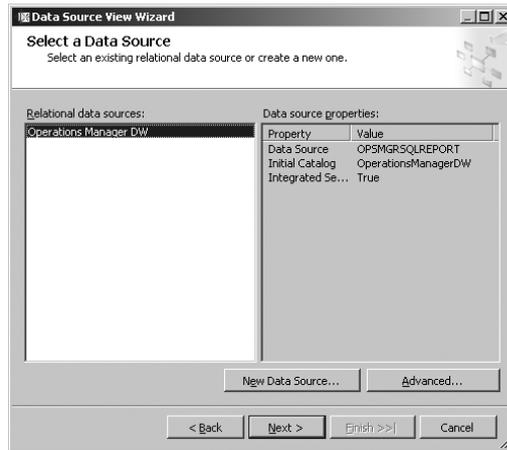
Notice that at the bottom of the Select How To Define The Connection page you have the option of clicking Next or Finish. At this point either button will do the same thing: create the data source. So the choice is yours. Once you click the button, the data source connection will be created and will appear in the Data Sources folder in Solution Explorer using a .ds extension.

FIGURE 9.47
Defining the connection to the database



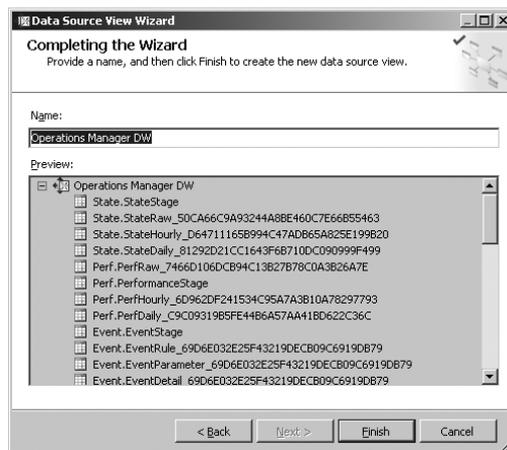
Next up is creating the data source view. As you did with the Data Sources folder, right-click Data Source Views and select New Data Source View. Move past the welcome screen and click Next when the database name appears in the Relational Data Sources field, shown in Figure 9.48. The next page that appears will allow you to specify the objects that will be used in your report model. Look over the objects and determine which you will need for your reporting purposes. If you are completely unsure, a good rule of thumb is to include any of the objects that do not include a prefix of dbo. Once the objects are selected, click the right arrow to move them to the Included Objects list, and then click Next.

FIGURE 9.48
Including objects for
the model



On the Completing The Wizard page, shown in Figure 9.49, enter a name for the data source view and click Finish. This page will also show you all the objects that will be used in the model.

FIGURE 9.49
Completing the
wizard



The final step to creating the model is defining the model. Right-click on the Report Models folder and select Create Report Model. Click Next on the welcome screen and then select your data source view you just created, as shown in Figure 9.50. Once you click Next, you will be taken to the Select Report Model Generation Rules page, which allows you to control how the metadata for

the model is generated. Accept the defaults as shown in Figure 9.51 and click Next; then select how you would like the model statistics to be generated. Select the Update Model Statistics Before Generating option; while this will increase the build time, it will also guarantee that you will not be using stale statistics from an earlier build.

FIGURE 9.50
Data source selection

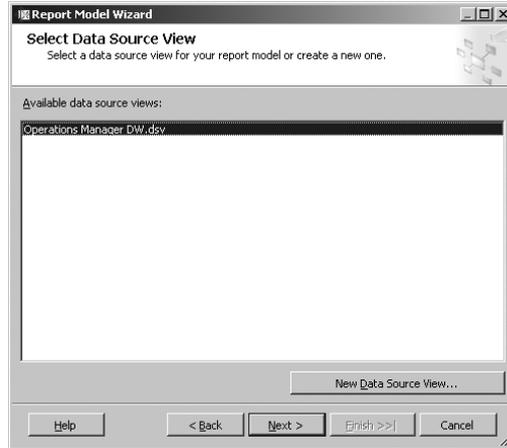
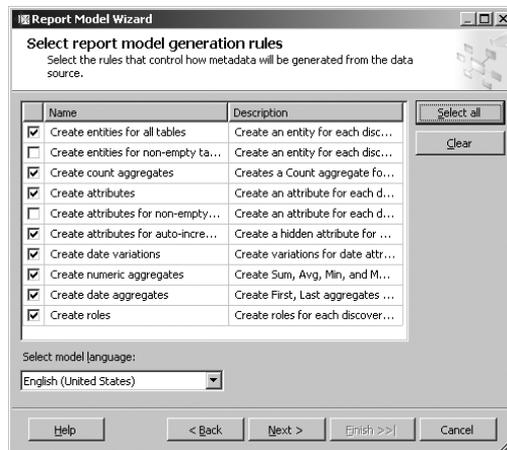


FIGURE 9.51
Selecting the
metadata for the
model



After specifying the data source view options, you have the opportunity to name the model. On the Completing The Wizard page, shown in Figure 9.52, you will also see the connection string that will be used. Click the Run button and the model will be created. As it is building, you will see the model generation, as shown in Figure 9.53. When the build completes, click the Finish button and you will see the model displayed in the Development Studio window.

Now that you have completed the build, deploy the model so that it can be used with Report Builder. Before the actual deployment, though, you need to define where the model will be stored. Right-click the solution name in the Solution Explorer window and select Properties. This brings up the window shown in Figure 9.54. Make sure the URL for your Report Server is correct and that the Configuration drop-down is set to Active(Production).

FIGURE 9.52
Naming the model

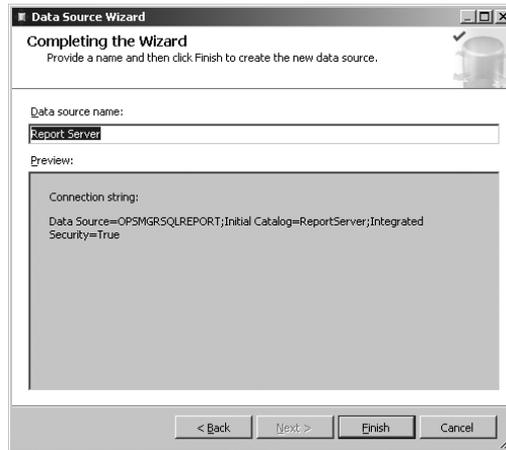


FIGURE 9.53
Building the model

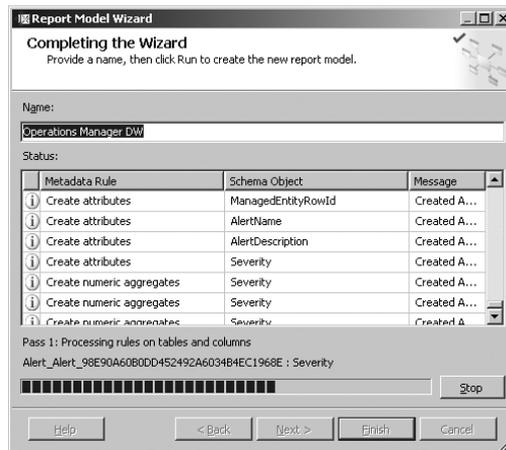
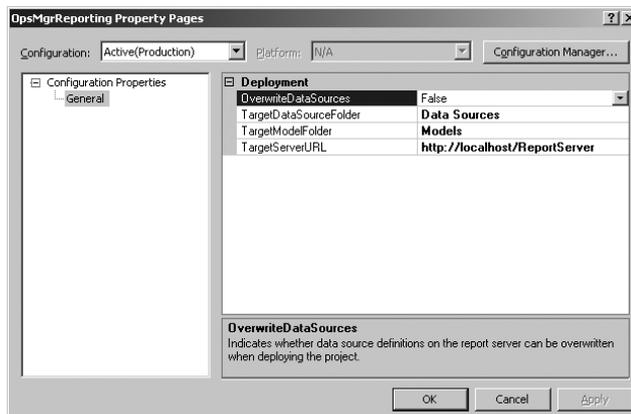
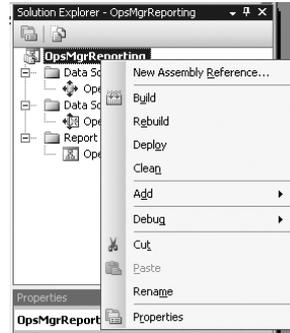


FIGURE 9.54
Setting the model publishing properties



In the Solution Explorer window, right-click your solution title and select **Deploy**, as shown in Figure 9.55. This will take all of the data that you had defined for the project and build the model, then save it in the folder you specified earlier when defining the project.

FIGURE 9.55
Deploying the model



USING REPORT BUILDER

Once you have the model created, you can start using Report Builder to create your own custom reports. When you are in the Reporting workspace, you will see a **Design A New Report** link in the Actions pane. The first time you click this link, you are presented with a prompt asking if you want to run the Report Builder application, as shown in Figure 9.56. Click **Run** and the appropriate files will be installed; then the Report Builder interface will appear.

FIGURE 9.56
Installing Report Builder



At first it does not appear that there is much you can do with Report Builder. If you look at the interface when it first starts, as shown in Figure 9.57, you won't see too much to work with. You have probably noticed that your model that you just created appears in the window on the right side of the interface. It is located just above the Report Layout options. Select the layout format that you want to use—Table, Matrix or Chart—then click **OK**. Alternatively, you could open an existing report and modify it by clicking one of the links at the bottom of the page under **Open**.

Once you select the format, you will be presented with the designer interface, shown in Figure 9.58. Here you can select objects from the Entities and Fields windows and drag them onto Report Designer. You can also add a title and text to the report to help identify and describe the data that is being gathered. In Figure 9.59, we are designing a simple report that will show the alert name and

severity of the alert. We dragged the Alert Name and Total Severity fields to Report Designer and added the title **Alert Totals**. Once you have your report designed, click the Run Report button to see the results. Our final report appears in Figure 9.60.

FIGURE 9.57
The Report Builder interface

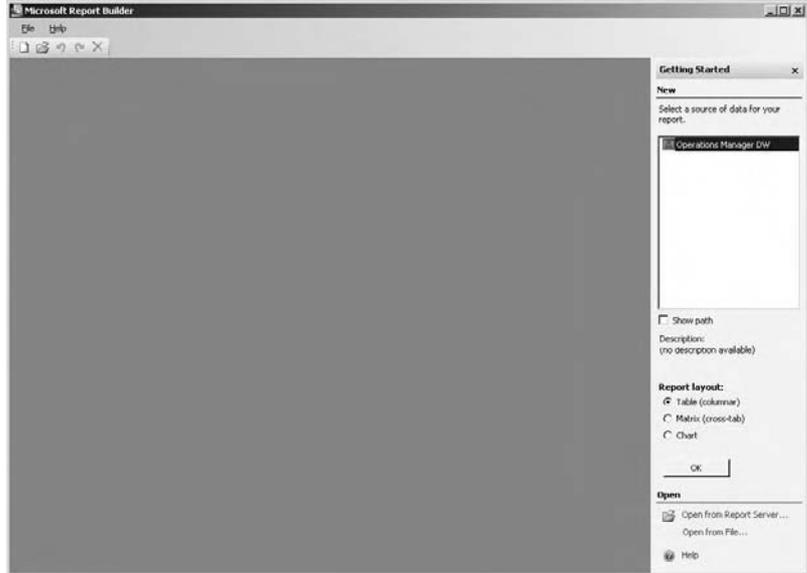


FIGURE 9.58
Report Builder designer interface

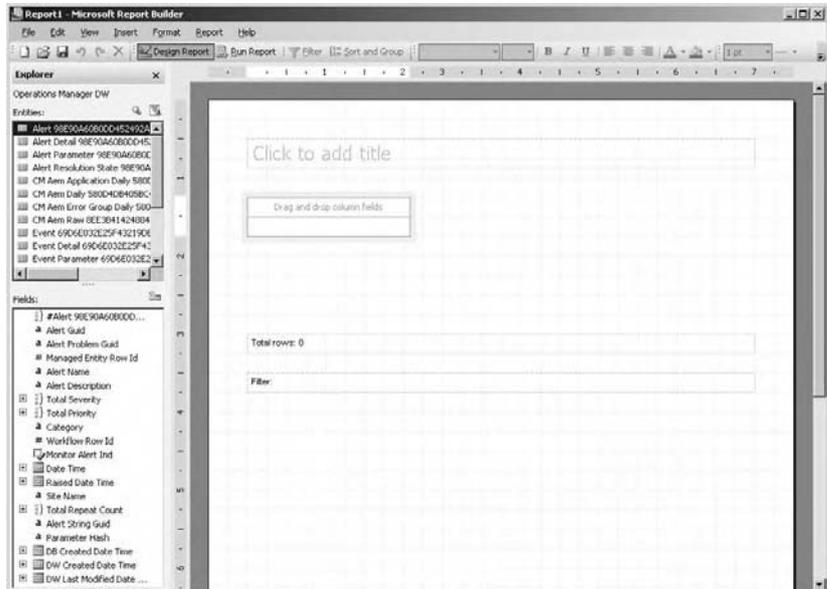


FIGURE 9.59
Choosing the report parameters

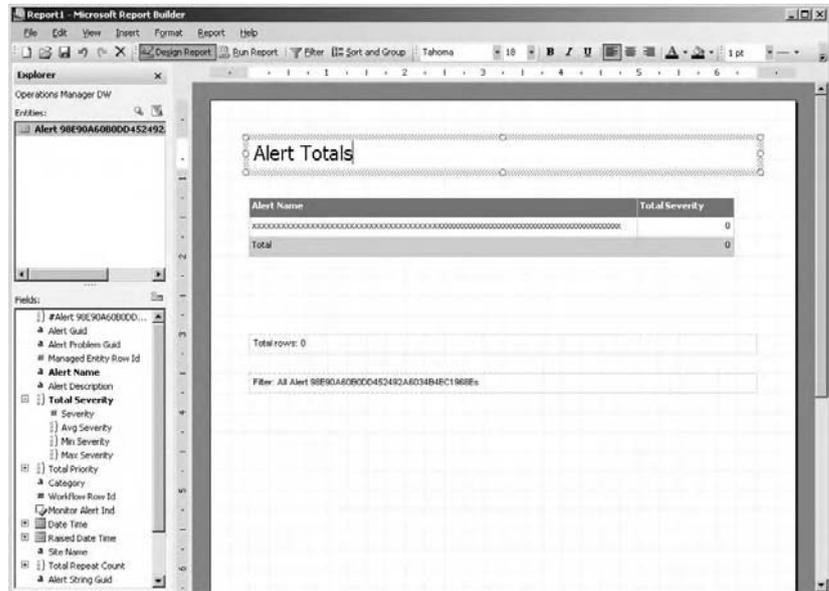
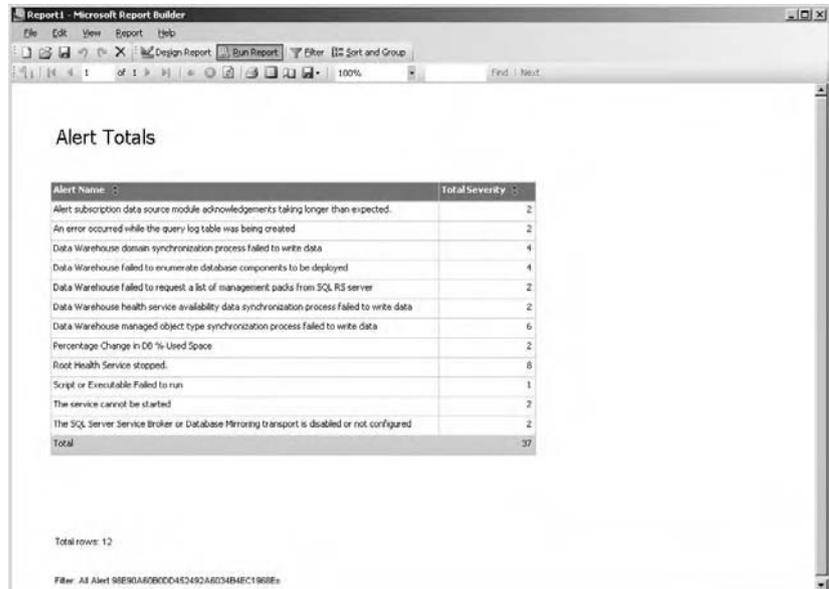


FIGURE 9.60
Finished report



The reports that you create will be saved in the Authored Reports node. You can now go back and rerun the report at any time to see what the updated totals are.

SCHEDULING REPORTS

Chances are you do not want to return to the Reporting workspace every time you need to run a report. There may be several reports that you want to view on a daily or weekly basis. Continually running and rerunning the same report can be tedious. So it's scheduling to the rescue. You can create a schedule for any of the existing reports and have the report automatically generated and stored for you to view when you arrive at the office.

Right-click on a report and choose Schedule, as shown in Figure 9.61. The Subscribe To A Request wizard appears. On the first page, Delivery Settings (see Figure 9.62), enter a description for the subscription and then choose a delivery method that you want to use after the report has been generated. Depending on the delivery method, you will be presented with more options to define, as shown in Figure 9.63. We have chosen to store the report in a file share.

FIGURE 9.61
Starting a schedule

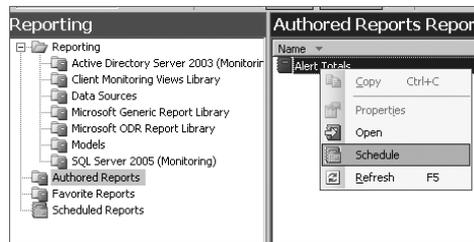


FIGURE 9.62
Initial Delivery Settings page

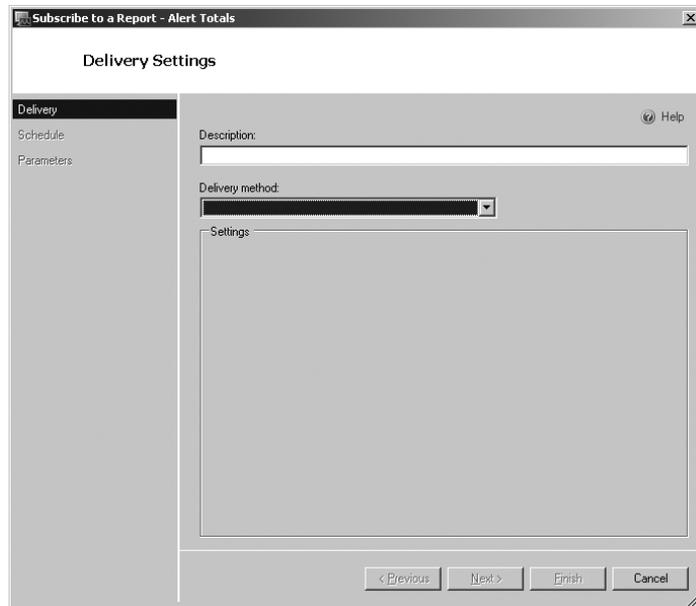


FIGURE 9.63
Additional delivery
options

If you take a look at the options that appear here, you will see text fields where you can define the network path to store the report as well as the credentials used to access the network share. There are also settings you can use to define the format for the report and the storage option. Figure 9.64 shows the format options that are available for you to store the report in. Choose the format that best fits your needs. If this is going to be a web page that will be viewed from your intranet, choose the HTM format that best fits. If it is going to be used in a spreadsheet or document, choose the appropriate format for the application.

FIGURE 9.64
Format options

Figure 9.65 shows the options for saving the report file. If you want to keep earlier versions of the report, you should use the Autoincrement option, whereas Overwrite will only store the newest report, deleting the previous version. Once all of the settings are entered, click Next to move to the Subscription Schedule page.

FIGURE 9.65
Save options

The Subscription Schedule page allows you to specify when the report will be generated and delivered. You can set nearly any schedule you can think of on the page, as shown in Figure 9.66. Once, hourly, daily, weekly, and monthly schedules are available in this configuration. Choose the

radio button beside the configuration options you want, and then define the schedule parameters. Notice that at the bottom of the page, you can specify not only when the report generation should begin but also when it should expire.

FIGURE 9.66
Configuring the
schedule

The screenshot shows a dialog box titled "Subscribe to a Report - Alert Totals" with a sub-tab "Subscription Schedule". On the left, there are tabs for "Delivery", "Schedule", and "Parameters", with "Schedule" selected. The main area is titled "Specify the schedule to use with this subscription" and contains a "Generate the report" section. Under "Generate the report", there are radio buttons for "Once", "Hourly", "Daily", "Weekly", and "Monthly", with "Monthly" selected. Below these are checkboxes for each month of the year, all of which are checked. There are also options for "On the following calendar days" (set to "1, 3-5") and "Every [2nd] week of the month on the following days: [Mon] [Tue] [Wed] [Thu] [Fri] [Sat] [Sun]". The "The subscription is effective beginning" section has a date of "Saturday, June 16, 2007" and a time of "12:48 AM". The "The subscription expires on" checkbox is unchecked, and the date is "Saturday, June 16, 2007". At the bottom, there are buttons for "< Previous", "Next >", "Finish", and "Cancel".

In Figure 9.66 we have selected the Monthly option and set the report generation options so that the report will be generated on the second Monday of every month. After you've defined your settings, click Next and then click Finish to set the schedule. Once the schedule has been created, you will find the schedule in the Scheduled Reports node in the Reporting workspace.

The Bottom Line

Install and configure SQL Reporting Services. Operations Manager Reporting relies on the underlying Reporting Services that is part of SQL Server 2005. You will need to install and configure Reporting Services on your SQL Server so that the reports that are defined in the management pack will operate.

Install Operations Manager Reporting. After installing SQL Server 2005 Reporting Services, you must install the Operations Manager reporting components. When you do so, you define a new reporting data warehouse. This new database contains the data that will be used to create the reports

Run predefined reports. Management packs contain definitions for reports. Once you install the reporting functionality, reports will appear in the Reporting workspace. Every management pack that is introduced in the management group generates new reports to run. When you run a report, you will typically be prompted for report criteria.

Build a Report Manager model. Before Report Manager can be used to create custom reports, you must create a model. Models define how the pieces of data in the database relate to one another and how the reports can use the data.

Create custom reports with Report Manager. Using Report Manager, define the criteria that is used in the report by dragging the objects to the appropriate location on the report page and then entering the text that will remain static on the page. Once created, the report may be run the same way as any other report.

Master It! Perform the installation of Reporting Services and Operations Manager Reporting as detailed within this chapter. Once everything is up and running, follow these steps to install the Audit Collection Services reporting.

1. Open Windows Explorer.
2. Create a directory named `OpsMgr Reporting Temp`.
3. Navigate to the location where you have stored your installation files for Operations Manager.
4. Open `ReportModels\acs` and copy both directories and the file to the `OpsMgr Reporting Temp` directory.
5. In the installation files for Operations Manager, navigate to `SupportTools`.
6. Copy `ReportingConfig.exe` to the `OpsMgr Reporting Temp` directory.
7. Open a command prompt and change the directory to the `OpsMgr Reporting Temp` directory.
8. Enter the command to upload the reports using the format `uploadauditreports.cmd ACSDBServer\Instance ReportingServicesURL TemporaryFolderName`.
9. Open Internet Explorer and navigate to `http://ReportingServicesServer/Reports`.
10. Click the Audit Reports link.
11. Click Show Details.
12. Click the Db-Audit data source link.
13. Select Windows Integrated Security in the Connect Using section and then click Apply.
14. Open the Operations Console and then open the Reporting workspace.
15. Right-click Reporting and select Refresh.
16. When the Audit Reports folder appears, run one of the reports.

Chapter 10

Working with Complex Environments

Management tasks often involve more than one site and could even involve multiple companies (or companies that have merged into a single entity). As you add different management scenarios, the management environment becomes more complex. You no longer have a single, discrete system to consider. In fact, the environment could include third-party solutions.

In this chapter, you'll learn how to:

- ◆ Determine whether you really need multihoming
- ◆ Create an administrator hierarchy
- ◆ Perform the multihomed agent installation
- ◆ Create connected management groups
- ◆ Add third-party solutions

Working with Multihoming Agents

As mentioned in the “Multihomed Agents” section of Chapter 4, the most common reason to use multihomed agents is to allow one machine to communicate with more than one management group. This communication strategy lets each group check the machine for specific issues. For example, one group may be interested only in how well Office is working, while another is concerned about Exchange. However, multihoming has significantly more utility than just this scenario. The following sections provide more information about multihoming as a useful feature, show how to configure the agent for multihoming, and describe how to configure multiple management groups.

Multihoming Overview

The main use for multihoming in Operations Manager is as a means for distributing the workload. The agent reports to several management groups on one or more servers. Each group has one or more support persons who check the client for problems in a particular area, such as Internet Explorer or IIS. In fact, this is the only use of multihoming that Microsoft supports, and many experts view this as the only practical use of multihoming as part of Operations Manager.

A few people have come up with other uses for multihoming, and you might want to think about them. For example, it's possible to use multihoming as a reliability measure. You could have an agent on a system communicate the same information to two completely different servers in physically different locations. If the connection to one server fails, the second server can continue



Real World Scenario

CONSIDERING THE MANAGEMENT-APPROACH CHANGES IN OPERATIONS MANAGER 2007

When working with Microsoft Operations Manager (MOM) 2005 and earlier products, the administrator used a server-centric management approach. The reason for this approach is that MOM 2005 provides separation between management groups. A server might report to multiple groups based on application requirements. For example, the Exchange server would report to the Exchange management group, while Internet Information Server (IIS) might report to the IIS management group. Of course, the problem with this approach is that the administrator must now consolidate the information manually in order to obtain a view of the enterprise as a whole, which often means that the administrator misses vital clues about the health of the computers as a whole.

System Center Operations Manager (SCOM) 2007 changes the management approach from server-centric to service-oriented. Consequently, it's possible to obtain a view of the enterprise as a whole and simply drill down into the data you need. The various data views let you ignore the data you don't need in a particular case. Because of this change in management approach, it's very likely that you'll only need one management group per server and possibly only one management group even when you have multiple servers working together to maintain the status of system health. The ease of managing scope in System Center Operations Manager 2007 means that you can restrict data views to just what a particular administrator needs—an IIS administrator won't see Exchange data, even though the data is available as part of the management group.

A change in strategy always incurs some penalty for the enterprise. In this case, you may find that you need to reconfigure your system to obtain maximum benefit from Operations Manager. The administrators working with you may require some additional training on the new management approach as well. However, the result of these changes is worth all of the time you spend because the new management approach results in an overall reduction in system complexity. You may very well not need multihoming in your organization because of the new management approach, and you should consider this change before setting up a complex environment that System Center Operations Manager 2007 really doesn't require.

to monitor the client. Obviously, you'd need to use this approach only for a system that has a very high priority and requires the extra monitoring.

TIP Another way to tackle the secondary site problem is to replicate the data from the primary SQL Server to the secondary site. In this case, the secondary site still receives the information, but not in real time. In addition, this technique doesn't necessarily protect you from the loss of the primary connection, but it does let you get the secondary site up and running quite quickly.

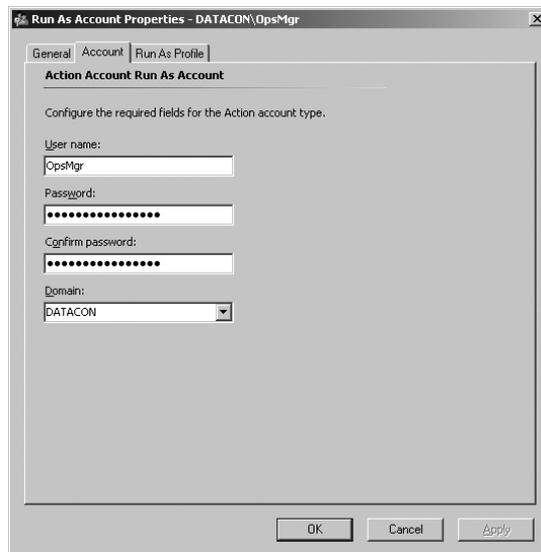
Some people will attempt to use multihoming for a number of unsuitable tasks. For example, one person wanted to use multihoming to set up agents for three servers. The theory was that each server would monitor the other two for failure. Although this idea has a certain appeal, the Operations Manager software doesn't support it very well. An alternative to this management technique is to set up a single server to monitor the health of all of the servers on the network. Using this approach, the system might warn you of some impending server issues. Obviously, the administrator

is the first one to know when a server goes completely offline (the cries of disgruntled users is hard to ignore).

It's important to consider the security issues regarding multihomed agents. You want to be sure that the security configuration will guarantee a consistent level of access across systems so that the agent doesn't have more access than is needed to accomplish the tasks for a particular management group. Giving the agent too much access opens a potential security hole in your setup. Normally, the service runs under the LocalSystem account unless you configure it to run under another account. When it needs to perform a task, the agent opens a separate process that runs under the Action account you configure in the Administration\Security\Run As Accounts folder of the Operations Manager. Simply double-click the account you want to change and you'll see the Run As Account Properties dialog box. Select the Account tab and you'll see the configuration information shown in Figure 10.1.

FIGURE 10.1

Configure the Run As account to give the agent only the access it actually requires to perform a task.



When there are multiple administrators working with a particular machine (assuming one per management group), it's important to choose a lead administrator. The lead administrator will address concerns about the multihoming installation based on the input of the other administrators. Some of the issues you should consider include the following:

Agent Installation Define how the group will install the agent on the client machine. Generally, you'll have better success if you choose to install all of the management groups using either remote installation or manual installation, but not both. It's also important to coordinate the installation effort to ensure the client machine isn't attempting to install more than one management group at a time.

Agent Version and Updates Make sure that everyone is using the same agent version with the same updates. You can't install an older version of an agent over a newer version. Any attempt to do so will result in errors, most of which won't tell you anything about trying to install the wrong version. This is another reason to coordinate the overall agent installation to ensure that one person is performing the required work.

Resource Requirements Verify that the client has sufficient disk space to hold the additional management groups. In addition, you don't want to use multihoming on a system that's already running slow or lacks sufficient RAM. In some cases, it might be better to continue in a single-homed mode and share the required data, rather than bog down the client machine until it becomes unusable.

Operations Manager Log File Settings The 3 MB space estimate for the management group doesn't include the log file requirements. You must add this space to the space for the additional management group installation. All of the management groups use a single log, Operations Manager, which the agent automatically sets at 15 MB during installation. If the log space isn't sufficient to meet the combined management group's requirements, make sure you increase the log size and account for this change as part of your disk space calculations.

Finger-Pointing Resolution One of the worst problems with multihoming is that administrators tend to start pointing fingers at each other when it comes time to diagnose a problem. To avoid this problem, choose an administrator to coordinate problem resolution.

Configuring Management Groups

Before you can do anything else, you must configure the management groups you want to use for multihoming. Remember that you don't need to use management groups strictly as a means of keeping data separate—that's the task of the management and scoping features of Operations Manager. What you really need to do is set up management groups based on a service-oriented approach. For example, if the administrators who are managing IIS use a server that's physically located in another country and these administrators never communicate with those who work with Exchange, it's possible that you'll be best served by creating a separate management group. You can create multiple management groups using the same techniques you use to create a single management group.

The starting point is to ensure you have the proper Active Directory setup. Use the MOMADAdmin utility to create the proper entries in Active Directory before you do anything else. The entries include the management group name, the Operations Manager security administration group, the name of the server, and the domain hosting the management group. Figure 10.2 shows a typical example of what you'll type at the command line, along with the standard response for a successful setup.

FIGURE 10.2
Create the required Active Directory entries using the MOMADAdmin utility.

```

C:\WINDOWS\system32\cmd.exe
C:\>MOMADAdmin DBMgmt SCOMAdmin WinServer DataCon
Microsoft System Center Operations Manager 2007 -- MOM AD Configuration Tool
(C) Copyright 2000-2006 Microsoft Corp.

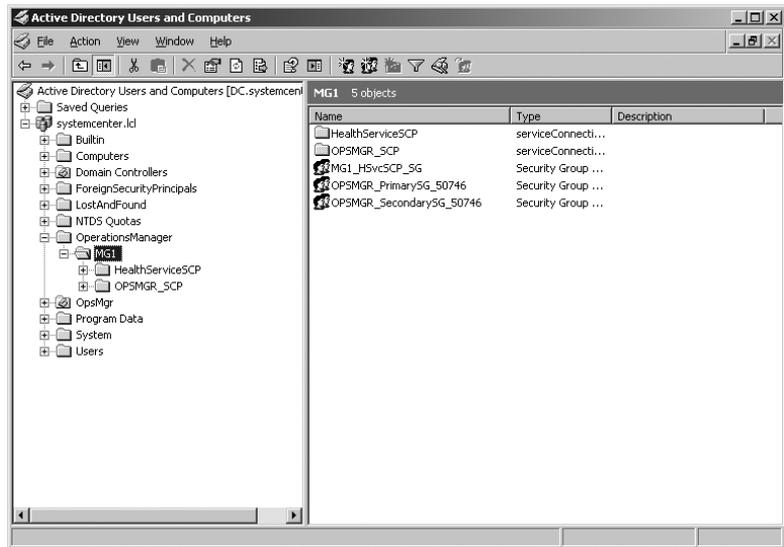
Successfully Created a container for ManagementGroup DBMgmt
Successfully added WINSERVER5 to SCOMAdmin security group

C:\>_

```

It pays to verify your setup after you perform the configuration, even if you obtain the required success message. Open the Active Directory Users And Computers console. Check the View > Users, Groups, And Computers As Containers and View > Advanced Features options. You should see each management group that you configure on a particular server along with the required subfolders, as shown in Figure 10.3. If you don't see these entries, there's a problem with your setup and you need to check it out before you proceed. Chapter 13 ("Troubleshooting") describes some of the troubleshooting issues you'll experience as you work with Operations Manager.

FIGURE 10.3
Make certain that the server setup is correct before you begin adding multihomed agents.



Agent Setup

The first task in configuring a multihomed agent is to identify the computers that actually require this functionality. It's tempting to think that you want to have all computers communicating with all management groups on your system. However, the reality is that attempting to create such a setup will:

- ◆ Bog down the client systems
- ◆ Chew up a ton of network bandwidth
- ◆ Probably result in information overload for the administrators

Normally, you'll want to choose clients based on the applications they have installed and the actual requirement to monitor them. For example, you really don't need to monitor the IIS setup on a developer's machine. The developer will know when the IIS setup isn't working, and may very well have caused the problem.

As with the initial agent installation, you have the choice of performing a remote or a manual installation. Each additional agent installation requires approximately 3 MB on the client machine. You'll want to be sure that the client system has the required memory (it shouldn't be a problem with the huge hard drives supported by today's systems). In addition, you can use Active Directory

Domain Services to assign computers to a management group. The technique for performing these tasks is slightly different from the initial agent installation.

NOTE All clients start out as *single-homed*, which means that they begin with a single connection to a single management group on the server. A client becomes *multihomed* when it has multiple connections to the server. Each of these connections is to a different management group. In short, the client starts out with a single home (a single management server to report to) and ends up with multiple homes (more than one management server to report to).

One thing does change from working with a single-homed setup: you don't have a choice about agent setups. If you want to create a multihomed agent setup, the client must use agent management rather than agentless management. The following sections describe all three kinds of multihoming agent assignments.

PERFORMING A REMOTE INSTALLATION

Most of the rules for working with a remote installation of a single-homed setup also apply to a multihomed setup. Consequently, most of the material in Chapter 4 also applies to a multihomed setup. However, you need to consider where to begin the installation. When you want the client to use a particular management group, you must perform the installation from the server that supports that management group. You perform the same steps that you use for the Discovery Agent (see Chapter 4 for details).

PERFORMING A MANUAL INSTALLATION

Manual installations are truly manual for a multihomed installation. You can theoretically create a wrapper for the process using the techniques described in the "Other Manual Installation Methods" section of Chapter 4, but these techniques won't work very well and are error-prone. If you must perform a manual installation, rather than pushing the agent to the client using the Discovery Agent, use the following manual steps:

1. Place the CD in the drive. If the autorun feature works as it should, you'll see the Start screen shown in Figure 10.4. If you don't see the Start screen, double-click SetupOM.EXE on the CD to start the program.

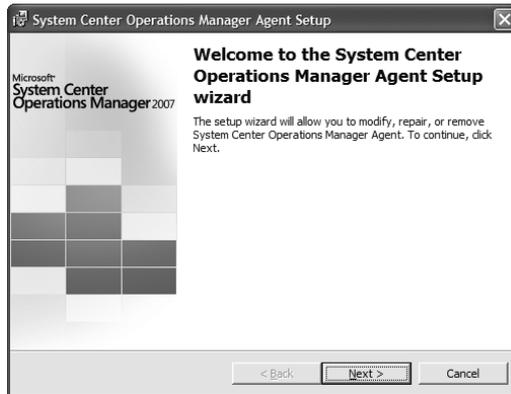
FIGURE 10.4
Use the CD, if desired, to begin the installation.



2. Click Install Operations Manager 2007 Agent. You'll see an initial dialog box like the one shown in Figure 10.5. You can also display this dialog box by double-clicking one of the MOMAgent.MSI files found in the \Program Files\System Center Operations Manager 2007\AgentManagement folder of the Operations Manager server. This folder contains three sub-folders: x86 (for 32-bit machines), AMD64 (for machines equipped with a 64-bit AMD processor), and IA64 (for machines equipped with a 64-bit Intel processor). Choose the copy of MOMAgent.MSI that matches the client machine's setup. You can only use the 64-bit options when working with a 64-bit version of Windows.

FIGURE 10.5

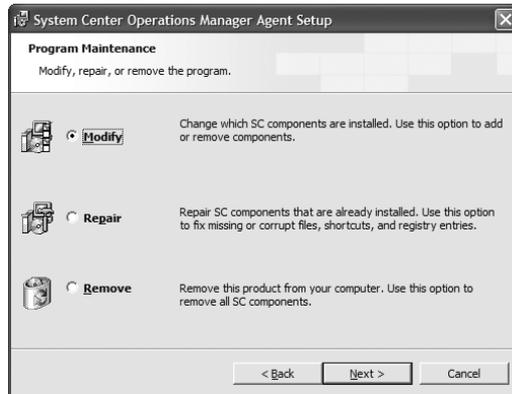
It's also possible to use the files on the Operations Manager server to start the agent installation.



3. Click Next. Instead of the display you normally see for an initial installation, you'll see the selections shown in Figure 10.6. You use these options as follows:
 - ◆ **Modify**—Lets you perform a number of tasks, which includes removing, modifying, or installing management groups. Even though the option name says Modify, it really does much more. You'll use the Modify option for most tasks.
 - ◆ **Repair**—Repairs or updates the software for an existing installation.
 - ◆ **Remove**—Removes all of the software. Use this option to remove all of the management groups, not just a single group.

NOTE You'll very likely see a warning message for the computer as you perform the configuration process in Operations Manager for any management groups currently associated with the client. The system can't perform the required client heartbeat check during the configuration process, so the warning is perfectly normal. The warning appears in several places in the Operations Manager, including the Monitoring\Computers folder and the Administration\Device Management\Agent Managed folder. The warning messages may also appear in the event logs of both the client and the server.

FIGURE 10.6
The multihomed installation displays different options than a standard installation.



4. Select **Modify** and then click **Next**. You'll see the modification options in Figure 10.7. You use these options as follows:
 - ◆ **Add Management Group**—Lets you add a new management group to the list of groups that the client already supports. This option helps you turn a single-homed client into a multihomed client. You also use it as the means for changing management groups for the client. Simply remove the existing group and add a new group.
 - ◆ **Remove Management Group**—Lets you delete a management group from the list of groups that the client supports. This option helps you turn a multihomed client into a single-homed client. Simply remove management groups one at a time until only one management group appears in the list to turn the client into a single-homed client.
 - ◆ **Modify Management Group**—Changes the server associated with a particular management group. You can either request that the agent obtain the information directly from Active Directory or manually provide the information.
5. Choose **Add Management Group** and click **Next**. You'll see the dialog box shown in Figure 10.8. Because the client is normally part of a domain, you'll seldom (if ever) see the **Use Management Group Information From Active Directory** option enabled. Normally, this option appears grayed out, as shown in Figure 10.8, because the client is part of a domain.
6. Check **Specify Management Group Information**. If you don't check this option, the installation program will merely reinstall the existing management group and not create a multihomed setup. Click **Next**. You'll see the dialog box shown in Figure 10.9.

FIGURE 10.7

Select a maintenance option depending on the task you want to perform.

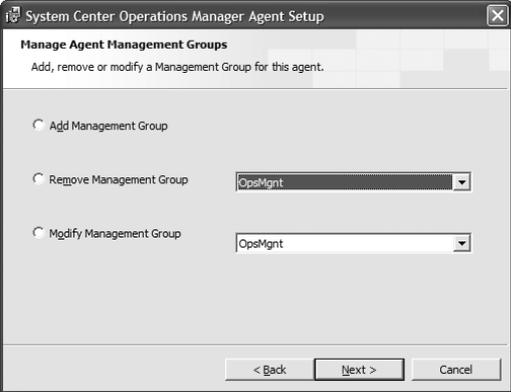


FIGURE 10.8

Choose the options for adding the management group.



FIGURE 10.9

Provide the information for the new management group for this client.



7. Provide the information required for the new management group. The fields will always contain the information for the first management group in the list for this client. The following list describes each of the field entries:
 - ◆ Management Group Name—Type the name precisely as it appears in Active Directory. If you forget the precise name, you can always look it up using the Active Directory Users And Computers console shown in Figure 10.3. See the “Configuring Management Groups” section earlier in this chapter for additional details.
 - ◆ Management Server—Type the Fully Qualified Name (FQN) of the server that supports the management group. The FQN always contains complete domain information using the normal dot syntax.

WARNING One of the most common multihoming errors is to change the Management Group Name field without changing the Management Server field. If you make this mistake, you can easily go back later and change the server information using the Modify Management Group option shown in Figure 10.7.

- ◆ Management Server Port—Defines the port used for TCP/IP communication on the network. Unless you specifically change this port on the server, don’t change it in the agent configuration. The messages between the client and server must have the correct port or the server won’t see the client messages. The wizard always contains the default port for Operations Manager.
8. Click Next. You’ll see the dialog box shown in Figure 10.10. The Agent Action account affects the local system, not the server. Consequently, using the Local System account won’t normally affect the client’s ability to perform tasks. In addition, using the Local System account is safer from a security perspective because it limits the activities of the agent to the local system. However, you may find that, in some cases, you need to use a domain account. The domain account allows a broader range of activity, but also opens potential security holes because the client now has greater flexibility. Use the Local System account whenever possible to favor security. See Chapter 13 for details on potential client problems and their fixes.

FIGURE 10.10
Choose an Agent
Action account for the
local client.



- Choose an Agent Action account. If you choose Domain Or Local User Account, the dialog box changes as shown in Figure 10.11. Provide the correct account information.

FIGURE 10.11

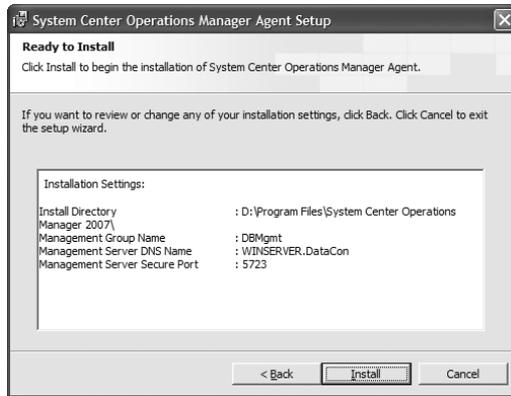
Using a domain or local user account means that you must provide credentials.



- Click Next. You'll see a summary dialog box like the one shown in Figure 10.12. Make sure you compare every entry with your server setup. Pay special attention to the Management Server DNS Name field.
- Click Install when you're certain that all of the entries in the summary are correct.

FIGURE 10.12

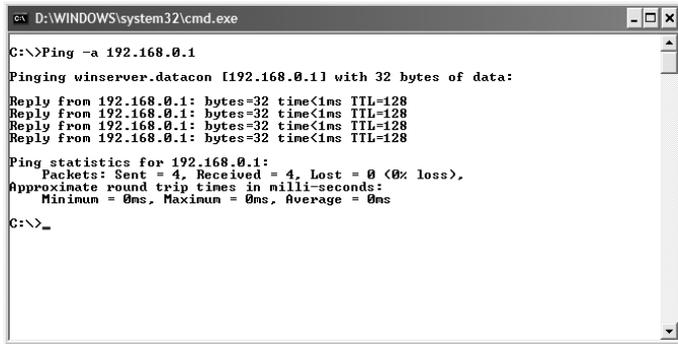
Verify that all of the information you provided for the new management group is correct.



Some administrators have a hard time figuring out how the client sees the server. You can always determine the server's DNS information by right-clicking Network Neighborhood or My Network Places on the server and choosing Properties from the context menu. Double-click the network connection. Select the Support tab and you'll see the IP address for the server. At this point, you can use the Ping utility with the `-a` command-line switch as shown in Figure 10.13 to verify the server's FQN. The FQN appears as part of the Pinging entry.

You can perform this check from either client or server. If you perform it from the client, you can also verify that the client actually sees the server and can contact the server using the FQN. If client can't see the server, your system has a DNS configuration issue (in most cases) and you'll need to fix it before you can get the agent to work properly.

FIGURE 10.13
Check the server's
FQN using the Ping
utility.



```

D:\WINDOWS\system32\cmd.exe
C:\>Ping -a 192.168.0.1
Pinging winserver.datacon [192.168.0.1] with 32 bytes of data:

Reply from 192.168.0.1: bytes=32 time<1ms TTL=128

Ping statistics for 192.168.0.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\>_

```

USING ACTIVE DIRECTORY DOMAIN SERVICES

Working with Active Directory Domain Services lets you automate the process of installing an agent. The server pushes the agent out to clients that meet the criteria that you specify. All this happens automatically, so you don't need to worry about performing the work yourself. A downside to this approach is that automation is useful only when it works as expected. Unless you're careful in specifying the criteria for agent installation, you could suddenly find that you have network bandwidth or other problems that result when too many systems suddenly have multihoming. Installing multiple agents using this feature is the same as working with a single-homed setup. See the "Understanding Individual Server Options" section of Chapter 3 for more details.

Agent Settings

The agent settings for a multihomed setup are very much like those for a single-homed setup. The only major difference is that you must change any settings for the agent at the server that supports it, which means making multiple changes for a single multihomed setup. You can discover all of the changes for an agent in the "Adjusting Individual Agent Settings" section of Chapter 4.

Configuring Connected Management Groups

The essential purpose of the Connected Management Groups node is to let you connect a local management group to other management groups in your organization. You need this feature if you want to provide centralized management of your network while allowing various groups to set up a separate management group. For example, you might want to create a separate management group for all the Exchange administrators. The sections that follow describe connected management groups in detail.

NOTE If you worked with management group hierarchies in MOM 2003, you'll have a good idea of what connected management groups do in Operations Manager. You use this feature to provide connectivity with management groups on other servers. Once you create a connection, you can perform tasks such as forwarding and responding to alerts that normally appear on the other server.

Overview of Connected Management Groups

Most organizations need centralized management of network resources, yet want separate management for special groups. The only way to achieve both goals is to create a connected management group scheme. When you create a connected management group scheme, you let one group know about alerts that occur on another group. In addition, the local group can send tasks to other management groups. The goals you achieve include:

- ◆ Monitoring a large number of management objects from one location
- ◆ Isolating particular management objects to a particular group

The group that has the connected management group is a centralized group—the one responsible for everyone on the network. Think of this group as the center with multiple satellites hanging off it. Although the central group doesn't maintain a close watch on other group activities on a daily basis, they can help in an emergency. In addition, the centralized group can act as a repository of information for everything that occurs on the network.

NOTE It's important to realize that a connected management group interacts only with the local group. The connection doesn't affect other management groups. The local group always receives alerts from the connected management group—the connection isn't two-way. Likewise, the local management group always sends task requests to the connected management group, not the other way around. The whole purpose is to provide a central management point for everything in the case of an emergency.

After you create the required connections to other management groups, you can use these connections to interact with the other groups automatically. This means setting up the system to forward alerts and respond to them in an intelligent manner. Consequently, setting up a connected management group is a four-step process:

1. Create the connectors in the Connected Management Groups node.
2. Set up the required notification channels in the Settings node.
3. Define recipients in the Recipients node.
4. Create the required subscriptions in the Subscriptions node.

You've already performed two of these tasks, Steps 2 through 4, although you might need to augment your setup to address the connected management groups. Even so, it's important to perform all four steps to ensure you have a complete setup.

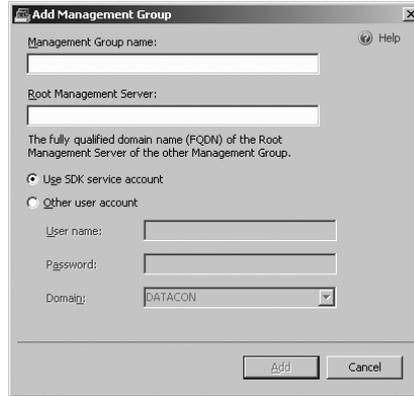
Configuring Connectors

The first step in the process of creating a connected management group is to create the connections. You perform this task to tell Operations Manager where to locate the other management groups. In order to complete this task, you must know the name of the other group and the server on which it resides. The following steps describe how to perform the required configuration.

1. Choose Start > Programs > System Center Operations Manager 2007 > Operations Console to open the Operations Console.
2. Provide any required login information. Make sure you use an account that has full privileges in Operations Manager.

3. Click Administration. You'll see the Administration folder.
4. Select the Administration\Connected Management Groups entry.
5. Choose Add Management Group in the Actions pane and you'll see the Add Management Group dialog box shown in Figure 10.14.

FIGURE 10.14
Define the new
management group
connection.



6. Type the management group name in the Management Group Name field.
7. Type the name of the server that hosts the management group (where the management group is the local group).

WARNING Don't type the name of a server that has a connection to the management group: the connection won't work. Always use the name of the server that actually hosts the management group.

8. Select an account to access the management group. The account must work on the remote machine, not on the local machine. In most cases, this means you'll have to provide an Other User Account entry, rather than rely on the Use SDK Service Account option.
9. Click Add. Operations Manager adds the new connection to the server.

WARNING Operations Manager doesn't verify the management group for you. It does verify that the server exists and that the account you provide has access to the server. However, this check doesn't guarantee that the connection will actually work. If you find that the connected management group doesn't work as expected, make sure you check the connection information for errors.

Once you create a connection, you can't modify it. Highlighting the connection in the Connected Management Group node and clicking Properties only shows the connection information. If you find that you made an error in entering the information, highlight the entry and click Delete. Operations Manager will remove the errant entry. Follow steps 4 through 9 to create a new entry with the correct information.

UNDERSTANDING MANAGEMENT TIERS

An Operations Manager setup can provide multiple management tiers. The lower level begins with a local server that supports a single management group for a particular need, such as Exchange. This and other local groups can connect to a central local server using the Connected Management Group node. This central server manages all the data for a particular region. You might divide it by company location, state, or other physical criteria.

If your company has more than one location, you might need to create another tier for regional reporting. Again, you'd create connections in the Connected Management Groups node, but this time the entries would be for the local servers, not for individual management groups.

You can create as many tiers as necessary to accomplish the task. The regional servers might connect to a country central server and finally to a server for the entire network. The number of tiers you need depends solely on the size and complexity of your network.

Of course, your tiers can also include product connectors. Some parts of the organization may rely on other kinds of management servers. Given the right product connectors, you can include information from these other local servers into your management scheme. See the "Interoperating with Third-Party Monitoring Solutions" section later in this chapter for additional details on this topic.

Providing Notification Channels

You'll find complete information about configuring notification channels in the "Notification Channels" section of Chapter 3. However, you must consider a few additional issues when working with connected management groups. Because the central server, the one with the connections, receives input from the entire network, you must now reconsider how the notification channels work. If you configure the notification channels incorrectly, you could end up with a confusing mess as administrators rush to fix a problem, or worse yet, fail to respond because they think someone else will address the concern. Here are the issues you should consider when configuring the notifications on the central server:

Notification Location Configure all the central management notifications on the central server. This means setting up the Notification properties on that server as well. It's important that you consider how the server is used when performing the setup. If you need to perform other kinds of notifications, then use the local server, the one on which the management group is installed.

Notification Recipient It's important to remember the local administrators will very likely receive notifications from the server that supports their management group. Consequently, you don't want to support notifications for the local managers on the central machine unless you want them to receive notifications for groups other than their own. For example, you might want to have some administrators assigned to nighttime or weekend duties who do need to receive notifications for all groups in an emergency.

Notification Type When the central server normally acts as a means of providing emergency services, rather than first-line services, you need to consider the kind of notification it provides. Email isn't going to be particularly useful, in this case, unless the administrator monitors it constantly at home. Consider using instant messaging (IM) or Short Message Service (SMS) instead. In fact, if you can't be certain that the administrator will receive the IM or SMS promptly, rely

on a custom solution to page the administrator or send a text message to a cell phone. In this case, you might actually need to rely on the command channel instead.

Redundancy An error that's easy to make is to use the same servers to provide support for the central server as you do the local servers. After all, it seems like a good idea to use the same server for all IM messages. Unfortunately, the server that provides the IM service might be the one experiencing problems. If you choose to use it to send the IM from the central server as well, the message will never get out. Always have a second, unrelated server provide services for the central server so that you can ensure there are at least two ways for the administrator to receive alerts.

Alert Type Don't send low-priority alerts through the central server unless they really are important to the parties concerned. The central server is providing a management function and provides a central locus for all information. If you try to forward every message through this central locus, the number of alerts could quickly overwhelm the server, not to mention the administrators. Think about the messages you provide here in the same way that you provide messages to management. In both cases, it's important that the person receiving the message get just an overview so that it's possible to keep an eye on the big picture, rather than become overwhelmed with the details of a particular area.

Message Content The message you send from the central server should differ from the one sent by local servers. For one thing, you should make it clear that the central server is issuing the message. Otherwise, administrators might be tempted to disregard the message because it doesn't necessarily fall into their normal area of expertise. It's important to make it clear that something more significant than a localized error has occurred.

Forwarding and Responding to Alerts

Configuring notifications for connected management groups works very much like the same setup for a simple system. The only difference is in thinking about the management tiers when you perform the configuration. For example, you might configure a recipient to receive all the messages from a server for a particular group but only part of the messages from a local central server. The central server could also send messages during a particular time frame, as shown in Figure 10.15. In this case, the recipient only receives messages from the central server on Saturdays. The rest of the time, someone else is responsible for handling emergencies monitored by the central server.

Operations Manager provides a great deal of flexibility in configuring the setup, and you can provide special exclusions as needed. For example, you can provide an exclusion for the recipient's vacation each year.

Make sure you set up the entries on the Notification Devices tab correctly as well. For example, you might expect the user to track events on a home system tied into the network during specific hours. Theoretically, you can rely on email during these hours to send alerts, backed up with IM or other secondary means. However, once the administrator logs off for the day, the notification may have to rely on SMS. In this case, create multiple Notification Devices entries and provide a schedule for each. As with the notifications themselves, you can configure a schedule for using the notification channel on the Schedule page of the wizard, shown in Figure 10.16.

FIGURE 10.15
Create notification-recipient setups based on the server tier.

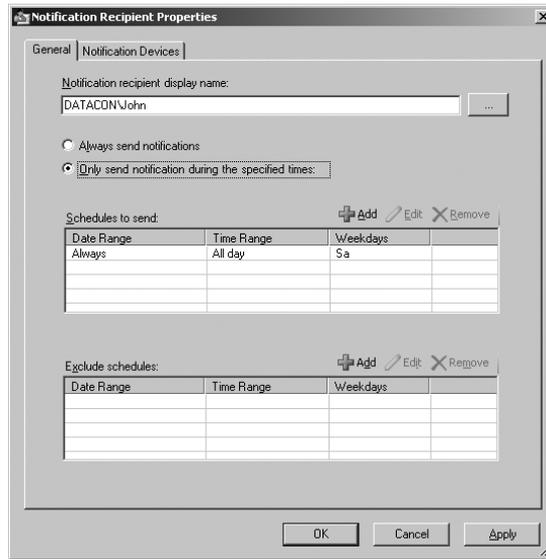
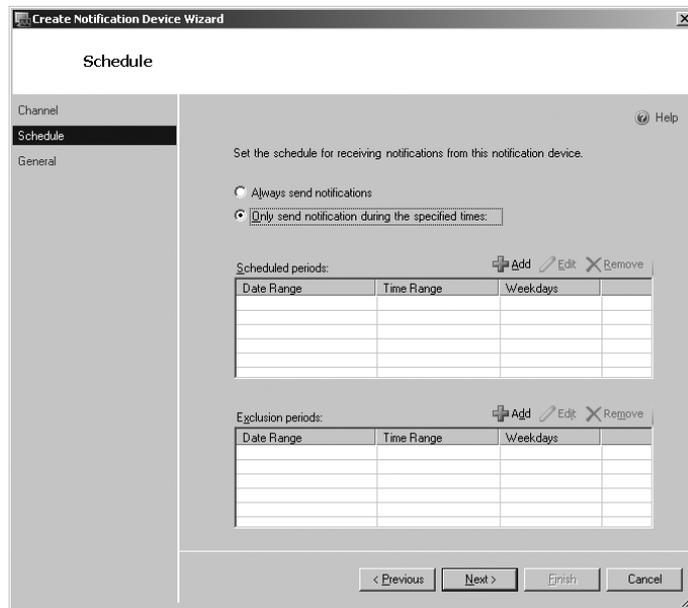


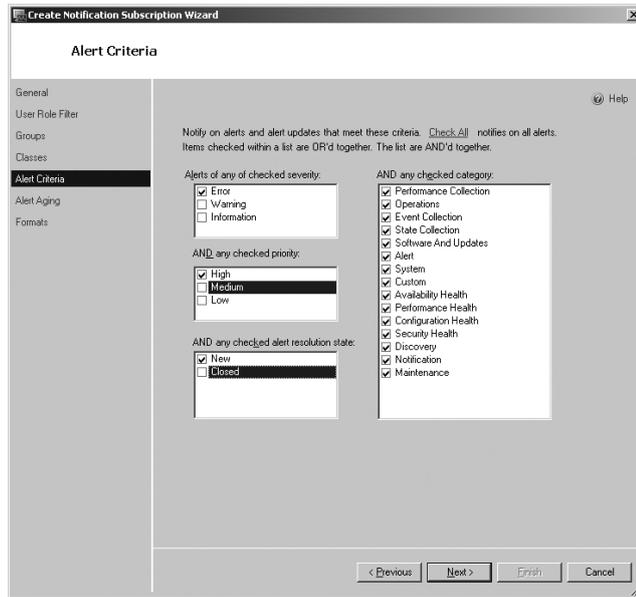
FIGURE 10.16
Define times to use each notification channel based on the needs of the administrator.



Once you have a recipient in place, you must define subscriptions. Because of the way that you've set up the recipient, you already know that the recipient will only receive emergency messages from the server during a given time. It's important to remember that the subscription won't show the connected groups—only the group supported by the particular server. You'll also want to set the Alert Criteria page to reflect the emergency nature of the alert, as shown in Figure 10.17.

FIGURE 10.17

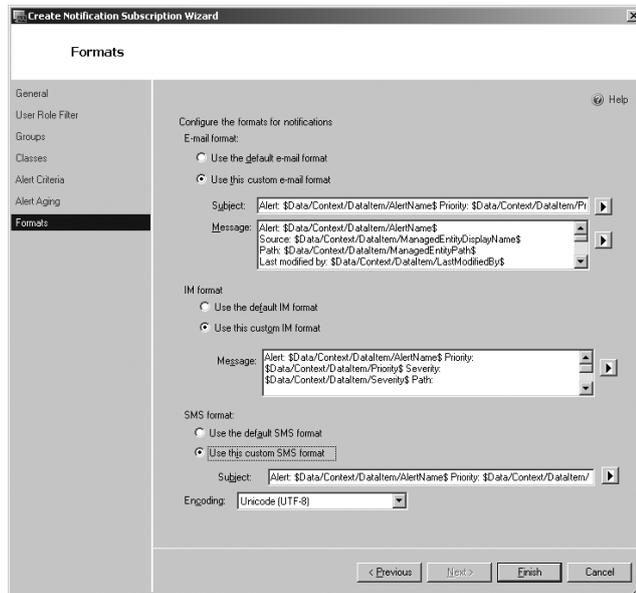
Configure the emergency alerts to reflect a special need.



An important part of configuring the alert, in this case, is the Formats page shown in Figure 10.18. The page will default to the standard messages for the server. Given that this is an emergency alert, you'll want to provide the administrator with a different message, one that actually tells the administrator that this is something different than the standard message. The message should contain the same alert items, such as the source, path, and description, but it should also include additional text that specifies that the alert is due to some extraordinary condition. The alert may simply occur on a weekend, but it might also be due to a loss of the normal alerting channel. The administrator needs to understand the difference.

FIGURE 10.18

Modify the Formats tab to provide message content that emphasizes the alert.



Interoperating with Third-Party Monitoring Solutions

Your company may have multiple monitoring solutions and use them for a variety of environments. Not every company uses Microsoft products exclusively; in fact, most don't today. Mergers, corporate environment, performance requirements, and other needs all make it unlikely that your organization will have a single solution in place. Consequently, you need some means of making all the monitoring solutions that you do have work together. In general, that means installing a third-party monitoring solution such as Ancoris Extension Framework for MOM 2005/SCOM 2007 (<http://www.ancoris.com/s/mon/ancmgmtfrm.shtml>), iWave Integrator (<http://www.iwaveintegrator.com/>), Engyro Product Connector Suite (<http://www.engyro.com/products/integration/>), or another solution.

Because all these products have differing capabilities, installation requirements, and other needs, the following sections can't provide you with a full discussion of every aspect of working with third-party monitoring solutions. For example, the following sections only provide a quick overview of installation because each vendor has a different process for accomplishing the task. You'll also find that your product will probably vary some in capability and flexibility from the solutions mentioned in this section. However, the results are the same in all cases. You end up with a product connector that sends alerts from Operations Manager to the other monitoring systems (and, hopefully, vice versa). The following sections describe third-party monitoring solutions.

Installing Connectors

Before you can do anything with a connector, you must install it. Every connector vendor will provide specific instructions to perform this task with their product connector. All vendors seem to use terminology that differs at least a little from the Microsoft terminology, so make sure you understand how the third-party product relates to Operations Manager. In addition, you'll need to know about the connectivity requirements between the machines. For example, you should know which products the connectivity vendor supports and the version of products that the vendor supports. A vendor might support HP OpenView but not Computer Associates Unicenter. Some of the most common management frameworks include:

- ◆ Aprisma (Cabletron) Spectrum
- ◆ BMC Magic Service Desk
- ◆ BMC Service Impact Manager
- ◆ Computer Associates Unicenter TNG
- ◆ FrontRange HEAT
- ◆ HP OpenView Network Node Manager
- ◆ HP OpenView Operations (U/W)
- ◆ HP Service Desk
- ◆ Micromuse Netcool/OMNIBus
- ◆ Remedy AR System
- ◆ Tivoli Enterprise
- ◆ Tivoli TEC

After you install the product, make sure that it's working properly with the management framework before you configure Operations Manager to use it. Otherwise, you might find that the setup

UNDERSTANDING OPERATIONS MANAGER RELATIONS WITH OTHER MANAGERS

Not all product connectors are created equal. The product connector you receive from a third party can be unidirectional or bidirectional. When the product connector is unidirectional, Operations Manager can only provide information to or receive information from the third-party product. In this case, Operations Manager can't really provide a complete monitoring solution, at least not at the top-tier management solution. Only when the product connector is bidirectional can Operations Manager provide full management capabilities.

If you use bidirectional product connectors and Operations Manager is the top-tier management solution, you can use it to provide full management functionality. Operations Manager can provide a centralized data store for all issues on the network, regardless of their source.

In addition, because Operations Manager provides only alerts to other management systems, and not informational messages, you can use it to create a trouble ticketing system. Whenever an alert occurs, a trouble ticketing system can automatically generate a trouble ticket for the problem, saving administrators considerable time and effort.

isn't working and not know where to begin troubleshooting it. In most cases, the time spent checking each step of the installation will make it significantly easier to locate problems when they do occur (and they will occur at some point in a complex setup like this).

Configuring Connectors

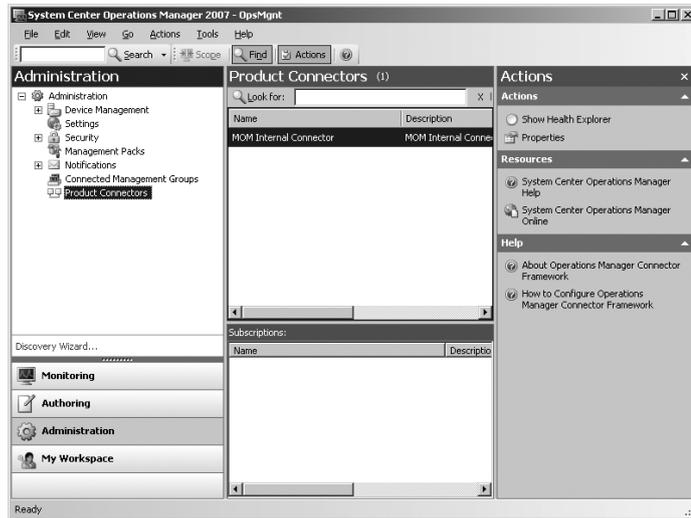
After you've installed a product connector and tested it so that you know it works, you can configure it for use in Operations Manager. The following steps help you configure the product connector for use. These steps should work about the same with any product connector that you use, but be sure to check the vendor documentation for any differences.

1. Choose Start > Programs > System Center Operations Manager 2007 > Operations Console to open the Operations Console.
2. Provide any required login information. Make sure you use an account that has full privileges in Operations Manager.
3. Click Administration. You'll see the Administration folder.
4. Select the Administration\Product Connections entry. After a few moments, you'll see one or more product connectors in the Product Connectors pane. Figure 10.19 shows the default connector, the MOM Internal Connector.

NOTE Wait at least a minute before you give up on seeing your product connector. Depending on the capabilities of your system, it may require some time for you to see anything at all. You should at least see the default MOM Internal Connector. If you aren't seeing this connector as a minimum, then you know you haven't waited long enough or that there's something wrong with your setup.

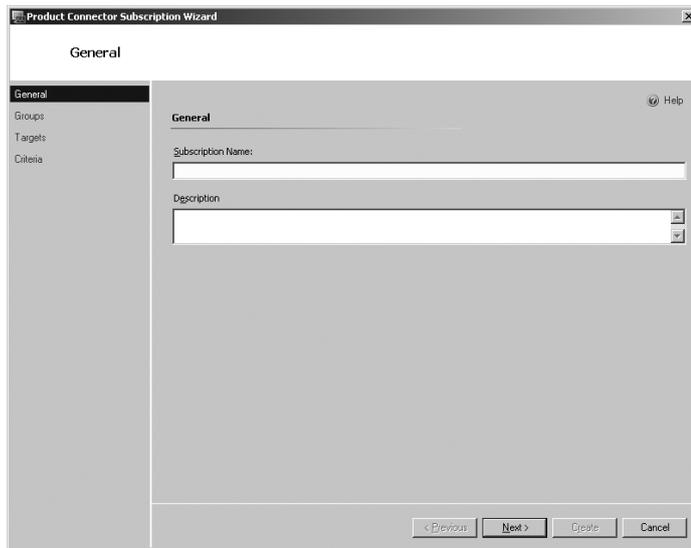
WARNING Never try to modify the settings of the MOM Internal Connector. Doing so can cause problems with your Operations Manager setup. The only connectors you should modify are those supplied by third-party vendors.

FIGURE 10.19
The Product Connectors pane should include at least one connector in it.



5. Right-click the connector entry and choose Properties from the context menu. You'll see the connector's Properties page. This page includes a General tab, as a minimum, where you see the connector name, description, and a list of subscriptions (there aren't any at the outset).
6. Click Add in the Subscriptions areas. After a few seconds, you see the Product Connector Subscription Wizard shown in Figure 10.20.

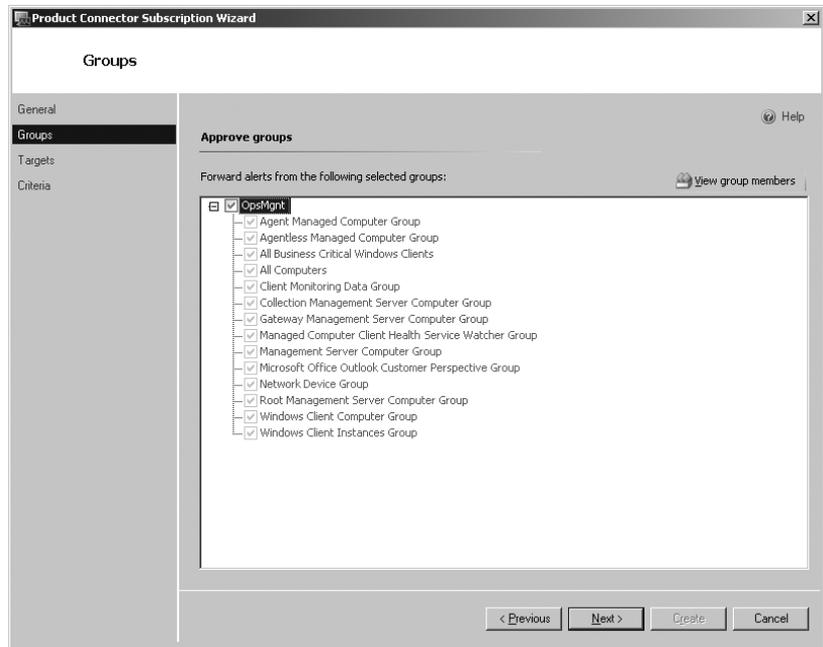
FIGURE 10.20
Add subscriptions to the product connector as part of the configuration process.



7. Type a name for the subscription in the Subscription Name field and a description in the Description field. Make sure you use a descriptive name and provide a full explanation for the subscription so that you can understand the configuration later.
8. Click Next. You'll see the Groups screen in Figure 10.21, where you can choose one or more management groups. If your system has only one management group, the wizard automatically selects it for you.

TIP If you only want to select a few of the child groups, clear the check mark next to the management group and then check the child groups you want to work with. For example, in Figure 10.21, you'd clear the OpsMgmt option and then choose one of the child groups, such as All Computers.

FIGURE 10.21
Choose one or more management groups for the subscription.

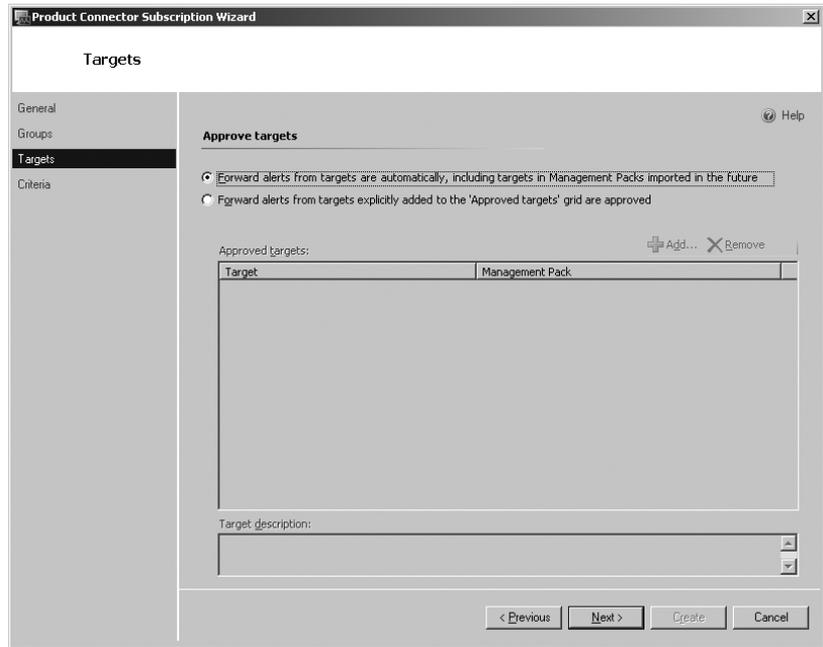


9. Choose one or more management groups and/or child groups for the subscription.
10. Click Next. You'll see the Targets screen shown in Figure 10.22. This screen provides you with two choices. You can send the alerts to everyone, or you can choose to filter the recipients.

The default option is to send the alerts to everyone, as shown Figure 10.22. Use the following steps when you want to configure the subscription to filter the recipients.

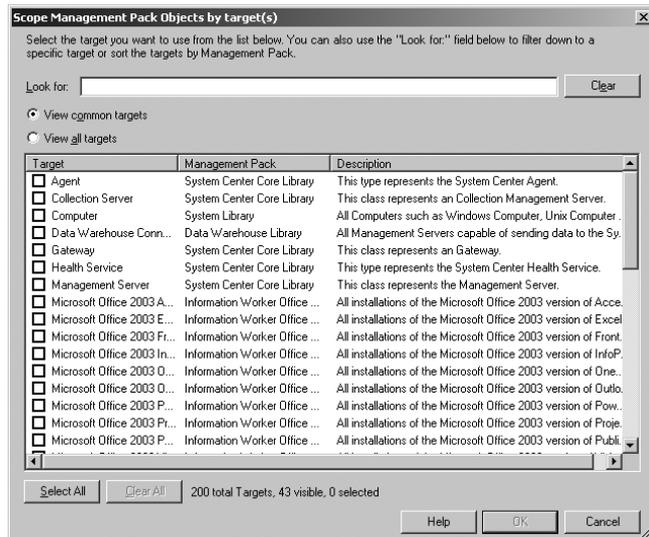
1. Select the Forward Alerts From Targets Explicitly Added To The 'Approved Targets' Grid Are Approved option.
2. Click Add. You'll see the Scope Management Pack Objects By Target(s) dialog box, shown in Figure 10.23.

FIGURE 10.22
Select the targets you want to use for the subscription.



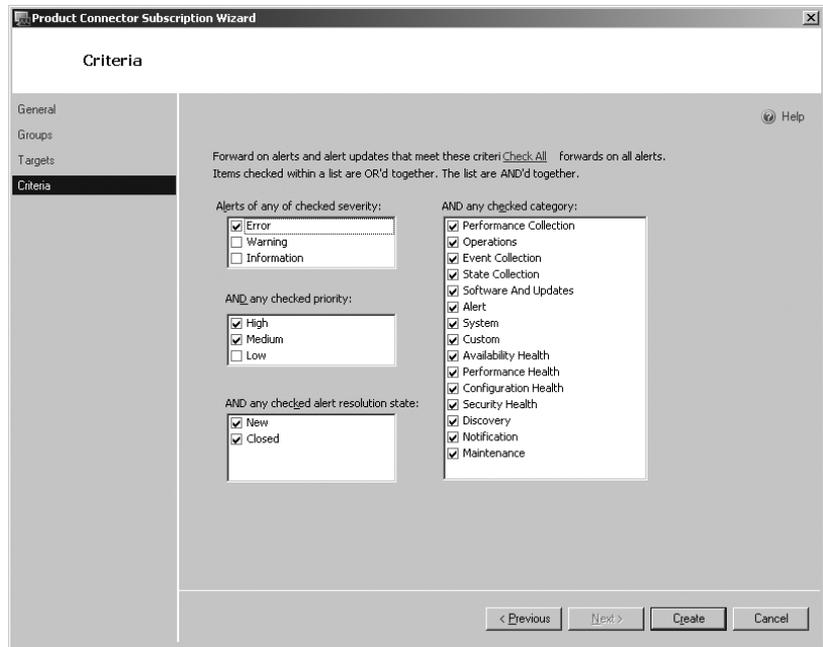
3. Check one or more target options and click OK. The agents you select will appear in the Approved Targets list in the Targets dialog box shown in Figure 10.22.

FIGURE 10.23
Choose one or more targets from the list supplied.



- Click Next. You'll see the Criteria dialog box, shown in Figure 10.24.

FIGURE 10.24
Add criteria to define what the subscription forwards using the product connector.



- Check all of the options you need to define the criteria for the subscription. The criteria determine what the subscription sends through the product connector to the other system. The information falls into four categories:

Alerts Of Any Checked Severity This group defines the severity of the information sent to the other system. The default setting sends only errors. However, you can choose to send warning and information alerts as well.

And Any Checked Priority This group defines the priority of the alerts. Normally, the subscription sends only high- and medium-level errors. However, you can choose to send low-priority-level alerts of any type.

And Any Checked Alert Resolution State This group defines the resolution state of the alert, whether it's new or closed. The default settings send both to the remote system. However, you can probably focus on new alerts since those are the alerts that the administrators need to worry about in managing the system.

And Any Checked Category This group defines what kind of alerts the subscription sends. The default setting sends every category. However, you might choose to focus on just the categories that would interest the administrator using the other product. For example, you might not want to forward alerts in the Custom category since they likely require a knowledge of precisely what the Custom category monitors.

- Click Create. The Product Connector Subscription Wizard creates the new subscription for you. Be patient; this process can require a minute or so to complete. Eventually, you'll return to the product connector's Properties dialog box, where you'll see the subscription in the Subscriptions area.

14. Perform steps 6 through 13 for any additional subscriptions you want to create.
15. Click OK to close the connector's Properties dialog box.
16. View the subscription in the Subscriptions pane (Operations Manager may require a few moments to update the Details pane—choosing View > Refresh won't make the process any faster). Make sure you have all of the subscriptions that the third-party product connector requires to work properly. You can add, edit, or delete subscriptions as needed. See the next section for details on modifying existing subscriptions.

Editing and Deleting Subscriptions

You won't create a perfect subscription every time, and sometimes the conditions for using the subscription change. This section discusses how to edit and delete existing subscriptions. In both cases, you'll begin in the Operations Manager, just as you do when creating the subscription. Use the following steps to open the subscription for modification:

1. Click Administration. You'll see the Administration folder.
2. Open the Administration\Product Connection folder.
3. Select the product connector you want to modify and verify the subscription appears in the Subscriptions list. If it doesn't appear as expected, make sure you're working with the correct server and that another administrator hasn't already performed the task (and forgotten to tell anyone).
4. Right-click the product connector and choose Properties from the context menu. You'll see the product connector's Properties dialog box.
5. Select the subscription you want to modify. At this point, you can perform one of two tasks:
 - ◆ Click Edit to modify the existing entry. Follow steps 7 through 13 of the procedure in the "Configuring Connectors" section of the chapter to modify the existing entry.
 - ◆ Click Delete to remove the subscription from the product connection.

WARNING Operations Manager won't ask you before it completes the subscription deletion. Microsoft assumes that you know what you're doing and don't want to be asked about deleting the subscription. Once you click Delete (even if you didn't mean to do so), the subscription is gone.

6. Click OK to close the product connector's Properties dialog box. You see any modifications to the subscription in the Subscriptions pane.

The Bottom Line

Determine whether you really need multihoming. A principle change in System Center Operations Manager 2007 management from previous Microsoft products is the use of a service-oriented approach. Previous Microsoft products relied on a server-centric approach that made use of multiple management groups (and, therefore, multihoming) a requirement. System Center Operations Manager 2007 does away with this requirement, in most cases, by providing you with better scoping. An administrator's role can easily determine what the person sees despite the availability of a single view of the enterprise through one management group.

Consider creating a single management group for your organization and then rely on scoping to limit data views based on an administrator's needs. Using this approach will greatly reduce the complexity of your setup.

Create an administrator hierarchy. You don't necessarily have to define one administrator as the leader and all others as subordinate, but you should have lead administrators for specific tasks. If you don't create some type of hierarchy, the resulting mess of confusing directions taken by each individual will make management of your network nearly impossible. The following list describes administrator management concerns, and you can choose a single lead for each of these requirements.

Multihoming Requirements This leader helps coordinate multihoming requirements, such as determining which computers must report to multiple management groups. The other administrators will need to provide input to this person, who then coordinates the efforts of the group as a whole.

Agent Installation Approach This leader decides how to manage the agent installation. Management means weighing the need for information against the cost of obtaining it. If every system on the network is multihomed, the system as a whole will probably fail to provide the desired results.

Perform the multihomed agent installation. After you decide how to perform the multihomed setup, it's time to perform the actual installation. Make sure that you set up and configure the management groups first. It's important to verify that the proper security settings appear in Active Directory before you actually begin the installation. You have three essential choices for creating a multihomed setup: remote installation, manual installation, and using Active Directory Domain Services.

Create connected management groups. If you decide to use multiple management groups and you do need to create connections between these groups, then you'll have to set up connected management groups. Generally speaking, these connections help you route information between groups. However, you want to use them with care to ensure you don't use up network bandwidth and system resources needlessly. Always keep the number of groups small—if you're creating too many of these connections, you might want to reconsider your setup.

Add third-party solutions. Once you have a stable and usable SCOM setup, you might want to add some third-party alternatives to it. If your organization uses several management products, create connections between them to create a whole-enterprise view of everything. Of course, Microsoft would prefer that you used SCOM for everything, but the reality is that very few organizations use a single solution today. All it takes is one merger or a dispute with another vendor to create a situation where you're using two or more management products. When this situation occurs, the only choice is to make sure everything works together to create a single view of your organization.

Part 2

Maintenance and Troubleshooting

Up to this point we have been discussing how to monitor your organization's services and servers using Operations Manager 2007. We have covered how you can use Operations Manager to proactively maintain and monitor nearly every aspect of the network and the supporting services.

But how do you monitor the monitoring systems? What can you do to make sure that the Operations Manager systems are working correctly? What are some of the common troubleshooting options available? Although Operations Manager has built-in, self-monitoring available from the management packs that are installed by default, there are still other steps you can take to make sure that the management servers are running as well as possible. And you should always take into consideration that one day you will encounter a disaster. We always hope that it doesn't happen to us, but we know that forces beyond our control can have devastating effects.

This part of the book explores the options that are available to help minimize the downtime that may occur when problems arise. From making sure that your management group is running as efficiently as possible to having a disaster recovery plan, the chapters in this section will point you in the right direction.

Chapter 11

Optimizing Your Environment

Nothing runs well if it isn't tuned. This fact holds true for anything you can think of, from cars to air conditioners. Operations Manager requires tuning as well, but we'll use the fancier term *optimizing* in this chapter. The goal of optimization is to ensure Operations Manager runs reliably, without any security problems, and at the highest performance possible. Those three goals might seem unattainable, but you can achieve all three to an extent. In some respects, optimization is also a balancing act—you must consider the needs of everyone working with Operations Manager when optimizing the system to ensure it meets all of the requirements but still interacts properly with the administrators who use it.

Part of optimizing a setup is maintenance. Just as you don't tune your car or electronics just once, you can't tune Operations Manager just once and expect it to continue performing well. Maintenance tells you about the health of Operations Manager and helps you know when you need to perform additional tuning. In short, optimization is a mix of tuning and maintenance applied in such a way that it places minimum stress on the system as a whole. Much of the maintenance you perform is during off-peak times to reduce the impact of the maintenance on the system.

In this chapter, you'll learn how to:

- ◆ View management pack optimization as an opportunity
- ◆ Optimize rules to improve reliability, security, and performance
- ◆ Configure self-tuning thresholds to make monitoring more efficient

Optimizing Management Packs

The management packs you receive contain settings that are more on the order of a guideline than an absolute requirement. In fact, the management packs are tuned to meet the needs of the average system, an elusive beast that doesn't really exist. The first step in optimizing a management pack is locating settings that don't quite match your system setup and changing them. You don't have to change every setting—just those that don't meet your needs. Of course, that means spending some time looking at the settings, which isn't always easy to do given the time constraints of most administrators. Even so, it's important to at least try to look for targets of opportunity in the tuning scheme.

Once you've optimized your management pack, you can't save the settings under the old name. The developer who created the management pack probably sealed it, making it impossible to change. Consequently, you'll have to save your optimized settings as a new management pack. That's where creating your own management pack comes into play. The customized management pack provides the means to store the settings for later use on another system or as a way to back up the changes you've made. The following sections describe both requirements.

CONVERTING YOUR EXISTING MANAGEMENT PACKS

Operations Manager 2007 won't use the AKM file format management packs from Operations Manager 2005. The basic reason is that Operations Manager 2007 is architecturally different from Operations Manager 2005—the two use different strategies for managing your network, so naturally they have different management pack requirements. In addition, the AKM file is a binary format and Operations Manager 2007 relies on an XML format. Fortunately, your investment in creating management packs for Operations Manager 2005 isn't lost. You can convert these older management packs into a form that Operations Manager 2007 will accept. The process requires two steps:

1. Convert the existing AKM file into an XML file in Operations Manager 2005 format using the MP2XML utility.
2. Convert the Operations Manager 2005 XML file into an Operations Manager 2007 XML file using the MPConvert utility.

The first step requires that you download and install the MOM 2005 Resource Kit found at <http://www.microsoft.com/technet/opsmgr/2005/downloads/tools/reskit.mspx>. This resource kit requires that you also have the .NET Framework 1.1 installed on your system. You can download the .NET Framework 1.1 at <http://www.microsoft.com/downloads/details.aspx?FamilyId=262D25E3-F589-4842-8157-034D1E7CF3A3>. Once you have the MOM 2005 Resource Kit installed, you'll find the MP2XML utility in the `\Program Files\Microsoft Operations Manager Resource Kit\Tools\Convert Management Packs to XML` folder. To use the MP2XML utility, you must provide an input AKM file and an output XML file like this:

```
MP2XML <InputAKMFile> <OutputXMLFile>
```

Now that you have an XML file in Operations Manager 2005 format, you can use the MPConvert utility located in the `\Program Files\System Center Operations Manager 2007` folder of your system. To use the MPConvert utility, you must provide an input XML file in Operations Manager 2005 format and an output filename. You can also include the `/trace` command-line switch, which provides verbose output of the conversion process, and the `/version` command-line switch, which provides versioning support, like this:

```
MPConvert.exe [/trace] [/version <OutputVersion>] <Mom2005MP.xml>
<OutputFileName.xml>
```

Identifying Useful Settings

Management packs include hundreds, sometimes thousands, of settings. Going through all of these settings one at a time probably isn't the best way to find the settings you need to change. A useful setting is one that can help the management pack perform better. For example, when a management pack samples the status of workstations on your network too often, it wastes server resources and slows everything down. You can use the Performance console to monitor the various Operations Manager features and locate items that are consuming many server resources. It's also possible to use the OpsMgr Connector object on the workstation to determine how the setup is affecting the workstation and therefore the server as a whole. After viewing how Operations Manager is working on the systems, you can decide which areas are consuming too many resources and then look for settings to make Operations Manager work more efficiently.

Once you have a basic area of concern to consider, open the Monitoring tab. Look in the Operations Manager folder for areas of interest. For example, you might want to look in the Monitoring\Operations Manager\Agent Performance\Performance Data folder shown in Figure 11.1 to see how the agent is performing when you have a problem in that area. You might find that the Health Service object is consuming a lot of processing cycles and want to create a rule override to control it.

Performance monitoring isn't your only tool in gauging the usefulness of the management packs you have installed. For example, you should use the Event Viewer console to check the various event logs—not just the one associated with Operations Manager. Innocuous warnings might point to a need for additional monitoring. For example, you might find that systems with a particular CD drive tend to fail regularly, so adding a rule to monitor the CD drive activity is a good idea. Of course, you won't know about this issue until you check the event log or the CD actually fails (CDs often don't fail immediately—the event log will simply say it couldn't read the CD at a given time and then the CD will work again for a while).

The event logs don't always have all the answers you need, however. You should also include your maintenance logs and discussions you encounter online as resource information for finding settings that need a tweak in Operations Manager. It's important to know what the management pack does as part of this research. For example, you might have installed the Office 2003 management pack. To understand the information that the management pack is providing, you must look in the Monitoring\Microsoft Windows Client\Health Monitoring\Information Worker Application Health Monitoring\Office Application Health folder shown in Figure 11.2. Notice that you can detect the state of Office 2003, along with any alerts that have occurred, making it easier to determine whether you need additional monitoring of this application.

FIGURE 11.1
Use monitoring tools to check the performance of the management packs.

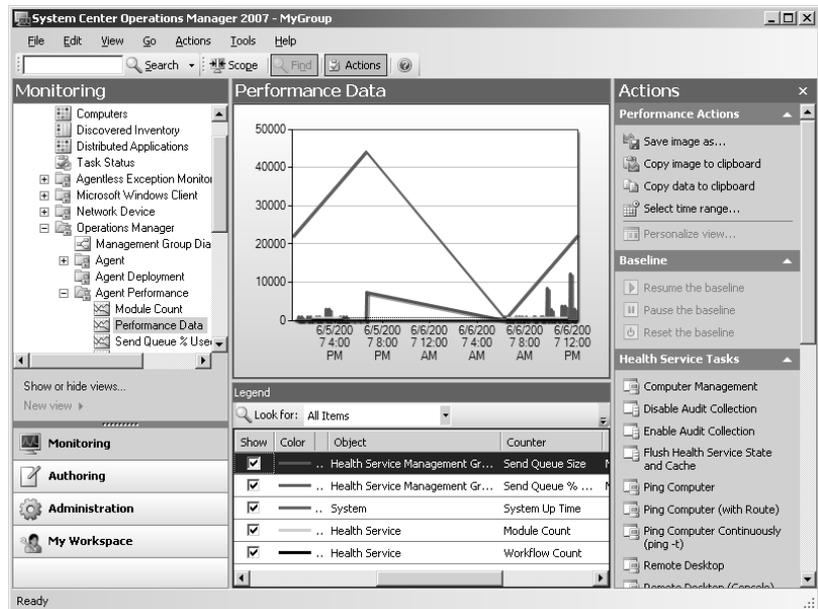
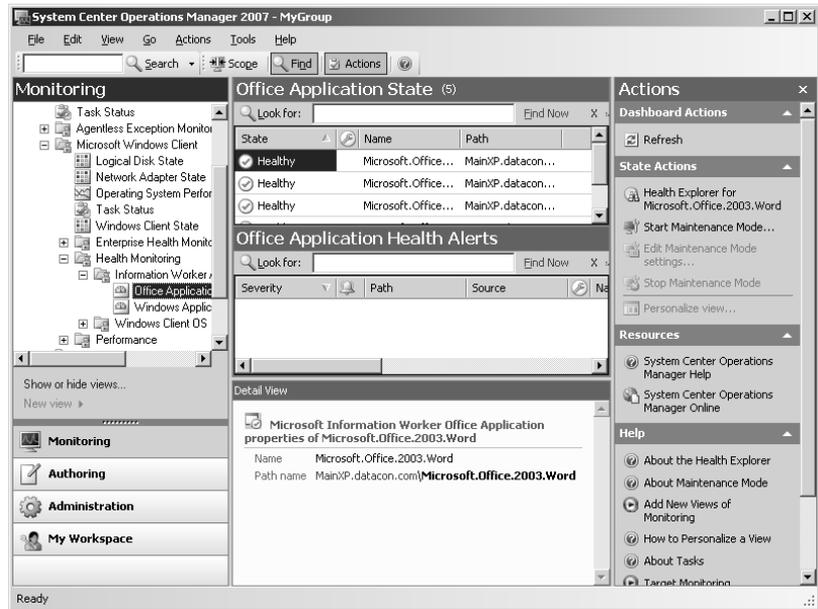


FIGURE 11.2
Monitoring the activity of management packs can give you clues about which settings to change.



Customizing Management Packs

Once you have a list of issues for Operations Manager, you'll want to implement them in some way. The management packs you receive from third-party vendors are normally sealed. A sealed management pack contains a host of settings that you can't change, which would seem to defeat the whole purpose of looking for settings to change. However, you can use overrides to customize sealed management packs. For example, you can disable a rule found in the sealed management pack and create your own version of the same rule (see the "Optimizing Rules" section later in this chapter for additional details on creating rule overrides). You can create overrides for:

- ◆ Monitors
- ◆ Object discoveries
- ◆ Rules

You can place the overrides in the same management pack as the original setting. However, this can cause a number of problems. First, using the original management pack can become confusing if you encumber it with many new rules. Of course, the original entries do receive the overrides, so you'll at least need to make this change in the original, but having the new rules in a custom management pack places all of them in one location where you can easily find and modify them.

Second, you can't export a sealed management pack. You'll very likely want to export the new rules and import them on servers, rather than repeat your efforts. Using a custom management pack makes it considerably easier to move the custom monitors, object discoveries, and rules around. It also makes it easy to create a backup of your custom settings (see the "Backing Up Management

Packs” section of Chapter 12, “Backup, Restore, and Disaster Recovery,” for details). Consequently, creating a custom management pack not only makes the changes clearer, but also helps you preserve the work you’ve done when disaster strikes. The following steps describe how to create a custom management pack:

1. Open the Administration tab.
2. Click Create Management Pack in the Actions pane. You’ll see the Create A Management Pack wizard, shown in Figure 11.3.
3. Type a name for your management pack in the Name field. Notice that Operations Manager automatically adds the name field information to the ID field. The name should reflect the purpose of the management pack accurately. For example, if the management pack modifies the Health Library, then a name such as Health Library Additions is a good description and will place the custom management pack in a position for easy viewing in the management packs list.
4. Type a version number for your management pack. Good version information is important. You should change the version number to reflect the significance of the change. Here are the entry meanings, beginning at the left and moving toward the right.

Major Use major numbers to show significant changes. For example, if you choose to add to the functionality of a management pack, you’d change the major number.

Minor Use minor numbers to show significant additions that don’t affect the overall purpose of the management pack. For example, if you add a new rule to refine the functionality of a management pack further, you’d update the minor number.

FIGURE 11.3
Define the identifying information for your management pack.

The screenshot shows a wizard window titled "Create a Management Pack" with a "General Properties" tab selected. The "Management Pack General Properties" section includes the following fields:

- ID:** An empty text box.
- Name:** An empty text box.
- Version:** A text box containing "1.0.0.0" with a note "e.g. 1.0.0.0" to its right.
- Description:** A large empty text area.

At the bottom of the window are four buttons: "< Previous", "Next >", "Create", and "Cancel".

Build Use build numbers to show significant changes in existing features. For example, if you update a rule to provide better performance characteristics, you'd update the build number.

Revision Use revision numbers to show minor changes in existing features. For example, if you change the documentation for a rule but don't change the operation of the rule, you'd update the revision number.

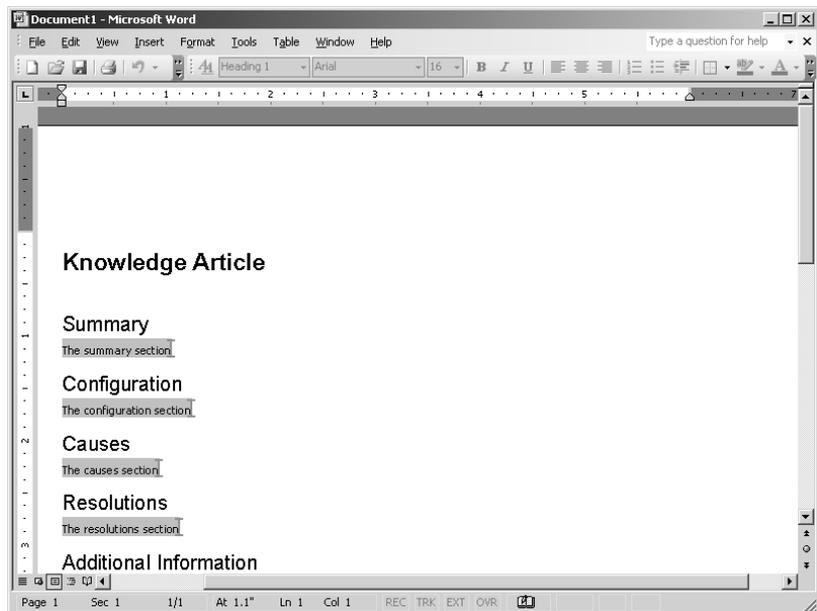
NOTE It's never a good idea to view revision number updates as cast in concrete. For example, an accumulation of revisions may warrant an update to the build number. You should create a policy that defines how to update the revision numbers in your company. Nothing's worse than revision numbers that don't reflect the reality of a change.

5. Type a description of the management pack in the Description field. Make sure you include the reason for the creating the management pack. The description should provide a good overview. You'll provide detailed information later in the creation process.

NOTE Before you can perform some rule-editing tasks, such as changing the content of the Company Knowledge tab, you need to download and install the Microsoft Visual Studio 2005 Tools for Office Second Edition Runtime (<http://www.microsoft.com/downloads/details.aspx?FamilyID=f5539a90-dc41-4792-8ef8-f4de62ff1e81>). In addition, you need to install a copy of Word on your server.

6. Click Next. You'll see a Knowledge Article page. Click Edit. If you have the proper support installed, Operations Manager will start a copy of Microsoft Word on your machine. You'll see a blank knowledge article, as shown in Figure 11.4.

FIGURE 11.4
Operations Manager uses Word to display the knowledge article entry for the Company Knowledge tab.



7. Click Save in Word to save the changes. Close Word. At this point, you'll see the knowledge article you created appear as part of the new management pack.
8. Click Create. Operations Manager creates the new management pack.
9. Add new monitors, object discoveries, and rules as needed to create the custom management pack. Make sure you enable the rules for the management packs you want to override.



Real World Scenario

THE SOFTWARE DECISION

More and more organizations are taking a look at open source options when it comes to software. Open source software is intriguing when you take a look at the savings in licensing costs that an organization can reap. At the same time, there are several companies that are limiting where software can be installed. This again is due to the cost of licensing. Many companies make their employees prove that they need to have a software package before it will be installed on their systems.

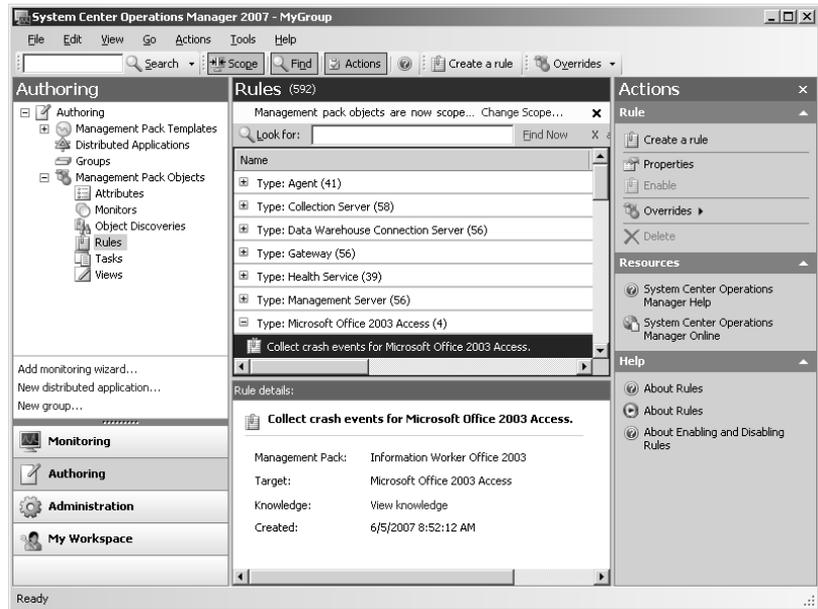
Such is the case of one company that had determined that the system administrators did not need to have Word installed on their workstations. The administrators had all of their management tools installed but did not see a need to install Word when they could use WordPad or Notepad to create their documents. That is, until Operations Manager came along. To create the Company Knowledge content, the administrators were suddenly faced with either having to install Word on their systems, or use another workstation when writing up company knowledge base articles. After a short period of time, they relented and installed Word. Faced with the inconvenience of moving to another workstation whenever they needed to document a fix quickly, installing Word won out over the price of a few more copies of Word.

Optimizing Rules

Every management pack you install creates rules. These rules control how the management pack affects the Operations Manager setup—how Operations Manager scans the systems on your network for adverse changes. The rules appear in the `Authoring\Management Pack Objects\Rules` folder, as shown in Figure 11.5. The Rules pane contains entries for each management package. When you open a particular management pack, you see the individual rules, such as those shown for Microsoft Office 2003 Access in Figure 11.5. Double-click a rule so that you can open it for optimization purposes.

The following sections describe some of the ways in which you can optimize rules (see the “Exploring Management Packs” section in Chapter 5, “Managing Management Packs,” for additional information on working with rules and rule groups). You’ll discover techniques for modifying existing rules and creating new rules. In addition, you’ll discover the nuances of self-tuning thresholds and optimizing performance counters.

FIGURE 11.5
Management pack rules change the way Operations Manager interacts with systems on the network.



Modifying Existing Rules

You'll find two kinds of management packs on your system. The first type is sealed, which means that you can't modify most of it. Vendors seal management packs to ensure that no one can modify them and then pass the rule off as coming from the vendor. Someone with ill intent could possibly damage your system if there weren't any way to seal the management from outsiders. Sealed management packs have an .mp file extension and you can't delete rules found in them. You can override the rules, which is part of the discussion in this section. The second type is an unsealed or custom rule that you create. You use custom rules to provide overrides for sealed management packs. A custom rule can account for differences between your system and the target system that the developer used to create the management pack. Consequently, even though you can't delete a rule, you can override it with a new rule.

SEALED RULES

Sealed rules normally come from vendors. The only element you can change with a sealed rule is the Company Knowledge tab. You can certainly see the other tabs, but you can't modify the information on them. Figure 11.6 shows the Company Knowledge tab with an entry already included.

To add an entry to this tab, click Edit. If you have the proper support installed, Operations Manager will start a copy of Word on your machine. You'll see a blank knowledge article, as shown in Figure 11.4. When you finish making your comments, click Save in Word and then close the Word document. At this point, you'll see the knowledge article you created in the Company Knowledge tab. Click Save to save the article.

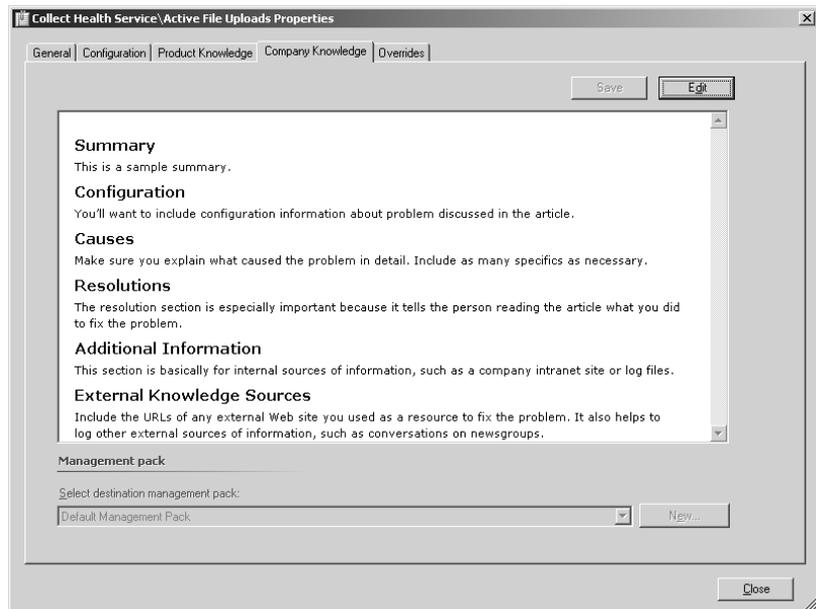
You'll eventually find that some of these sealed rules don't work as anticipated. For example, you might find that a vendor samples data from your network too often or hasn't taken a configuration issue on your system into account. Consequently, the inflexibility of a sealed rule can be a problem. An override provides a means of telling Operations Manager not to use a particular rule with an object. You can also override parameters to make the rule work better with your system. When you override a rule, you can create a new rule to take its place that works with your system. Place this new rule in the type, such as Agent, that provides the required support for the rule or create a custom management pack to hold it.

The Overrides tab, shown in Figure 11.7, provides the means of overriding certain aspects of the rule. You can choose to disable the rule completely for some objects or you can choose to override a parameter. Disabling the rule means that Operations Manager won't use it at all. Overriding a parameter means that Operations Manager uses the value you supply, rather than the default value. When you finish making changes to the overrides, you can click View Summary to see the results, as shown in Figure 11.8. You can choose to view the overrides for the rule as a whole or just for the currently selected object. The Overrides Summary dialog box is important because you can use it to edit and delete existing overrides—the Disable and Override buttons only let you add new overrides.

When you disable a rule or override a parameter, you have a choice of what to disable or override. Clicking Disable or Override displays a menu of options that includes:

- ◆ The current object as a whole. (For example, when you choose the Computer object, Operations Manager disables or overrides the rule for all computers.)
- ◆ A particular group.
- ◆ A specific object of the current type (for example, a specific computer on the network).
- ◆ All objects of another type (such as Agent).

FIGURE 11.6
The Company Knowledge tab contains information your company has acquired about a particular rule.



After you make a selection, Operations Manager displays a list of appropriate objects. You select the object you want to modify and click OK. At this point, you'll see the Override Properties dialog box shown in Figure 11.9 when working with a parameter override. Check the parameter you want to override, choose the override value, and then click OK. When disabling a rule, Operations Manager simply asks whether you're sure you want to disable it. Click Yes to complete the process.

After you complete the configuration process, you can click OK to make the changes permanent. Make sure you test the new rule to ensure it works as anticipated. For example, if you're working with a performance rule and you change the sampling interval to reduce the amount of work Operations Manager must perform, make sure that you're still sampling the object often enough to obtain valid statistics for it. Otherwise, Operations Manager could see an error condition long after it's time to fix it.

FIGURE 11.7
Define overrides
for your rule using
the options on the
Overrides tab.



FIGURE 11.8
Use the Overrides
Summary dialog
box to edit and
delete overrides you
created earlier.

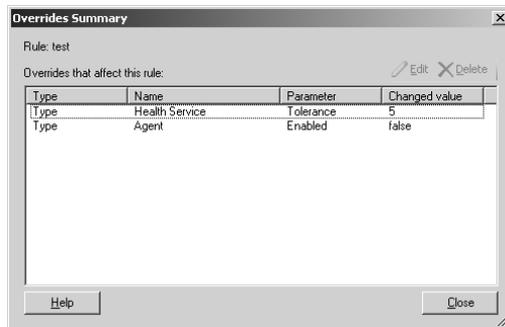
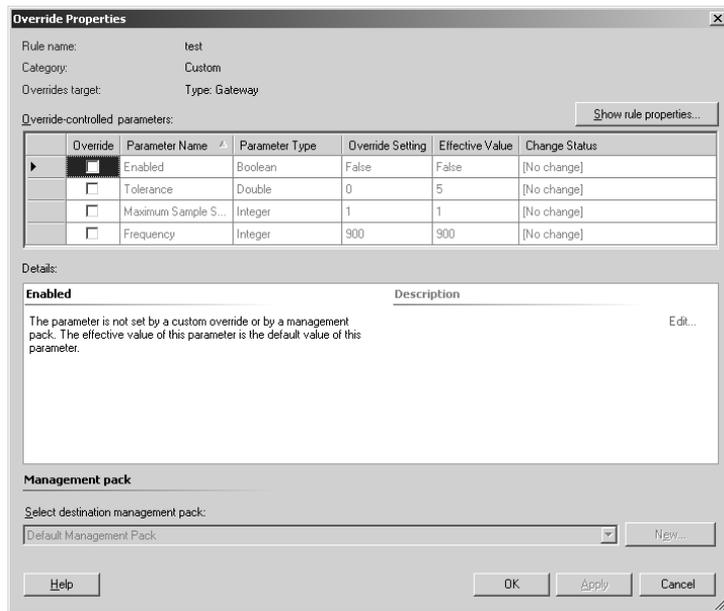


FIGURE 11.9
Choose the parameter
you want to override
and choose a new
value for it.



CUSTOMIZED RULES

When you work with a rule that isn't sealed, such as one that you've created yourself, you'll see only four tabs instead of the five shown in Figure 11.6. The dialog box won't include the Company Knowledge tab because you can modify the rule information using the Product Knowledge tab instead. The process for performing the edit is precisely the same as you use for the Company Knowledge tab.

The General tab contains the name of the rule and its description. Although Microsoft tries to tell you that the description is optional, providing a detailed description is always helpful because it makes the purpose of the rule clearer. Even though the person using the rule can get details on the Product Knowledge tab, sometimes all the person really needs is a good overview. The Description field can provide this overview for the user. The final entry on the General tab is the Rule Is Enabled option. Check this option when you want to make the rule active, and clear it when you want to disable the rule.

The Configuration tab, shown in Figure 11.10, helps you configure the rule. The dialog box contains two sections. The upper section, Data Sources, controls the rule input. When you click Edit, you'll see the data source you provided when creating the rule, such as the performance counter information shown in Figure 11.11. The content of this dialog box changes according to the kind of data source you choose when creating the rule (see the "Creating New Rules" section later in this chapter for details). You can change the actual data collection information in this dialog box easily. However, you can't change the optimization information shown on the Optimized Collection tab. In order to change the optimization information, you must create a new rule.

The lower section contains information on how the rule reacts to new input. You can add new responses, edit existing responses to match the requirements of your setup, and remove responses you don't need. When you click add, Operations Manager displays a little menu that asks whether you want to create a command or a script entry. In both cases, you'll see a beginning dialog box where you type the name of the new response. It's important to provide a descriptive response name so that others don't have to figure out what you intend by the new entry. When you click OK, you'll see either a

Script (Figure 11.12) or a Configure Command Line Execution Settings (Figure 11.13) dialog box where you can provide the particulars for the new script or command. When working with a script, you provide a script name, timeout value, the actual script content, and any parameters the script requires. When working with a command, you simply provide the path, arguments, and working directory for the external executable program. Editing an existing script or command is similar to creating a new response, except you don't have to provide a script or command name. When you want to remove a script or command you no longer need, highlight the entry, click Remove, and click Yes when Operations Manager asks if you're sure you want to remove the entry.

FIGURE 11.10

The Configuration tab makes it possible for you to change how the rule works.

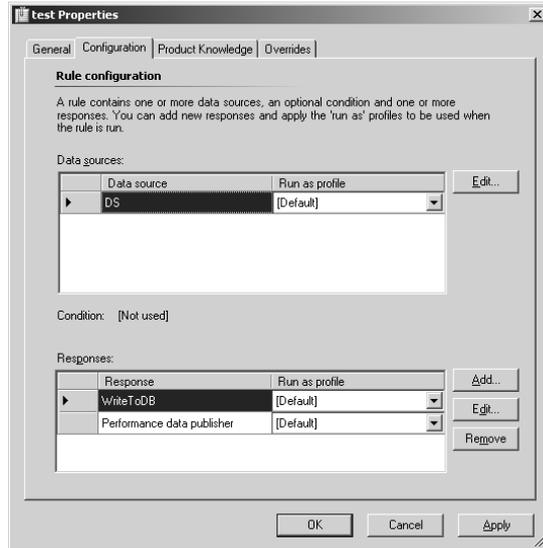


FIGURE 11.11

A data source dialog box changes according to the kind of data source the rule uses.

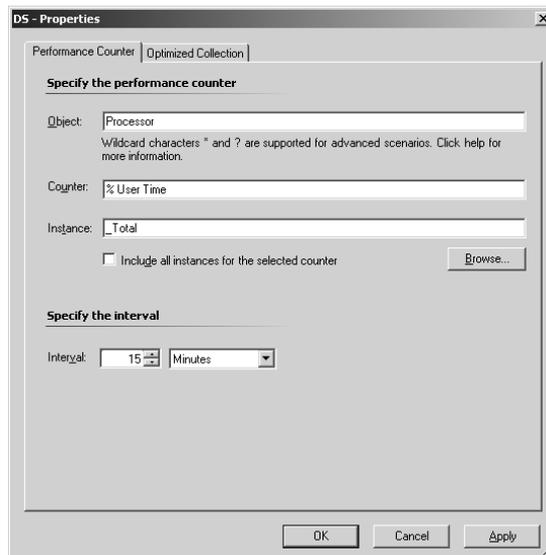


FIGURE 11.12

Create a new script to perform scripted tasks within Operations Manager.

The screenshot shows the 'Create Response Wizard' dialog box with the 'Script' tab selected. The window title is 'Create Response Wizard'. The main title is 'Script'. Below the title bar is a 'Specify Script Details' section with a 'Help' icon. The 'Enter script information' section contains a 'File Name' text box with the example 'MyScript.vbs' and an 'Timeout' spinner set to '1' with a 'Minutes' dropdown. Below this is a large 'Script' text area with an 'Edit in full screen...' button. At the bottom, there is a 'Parameters...' button and a note: 'Click the "Parameters..." button to specify the script parameters.' Navigation buttons at the bottom are '< Previous', 'Next >', 'Finish', and 'Cancel'.

FIGURE 11.13

Create new command-line settings to tell Operations Manager which external application to use.

The screenshot shows the 'Create Response Wizard' dialog box with the 'Command Line' tab selected. The window title is 'Create Response Wizard'. The main title is 'Configure Command Line Execution Settings'. Below the title bar is a 'Command Line' section with a 'Help' icon. The 'Specify command line execution settings' section contains a 'Full path to file' text box with the example 'c:\windows\system32\ping.exe' and a 'Parameters' text box with a right-pointing arrow button. Below this is an 'Additional settings' section with a 'Working directory' text box and a 'Timeout' spinner set to '15' with a 'sec' label. Navigation buttons at the bottom are '< Previous', 'Next >', 'Finish', and 'Cancel'.

Creating New Rules

The settings you get with a management pack are unlikely to fulfill every need. For example, you might have a problem segment on a network and want to check that segment regularly to ensure it's still providing good throughput for the users. A management pack developer won't know to include that rule because it's particular to your installation. As you work with Operations Manager, you might find yourself wondering whether a particular management feature provides enough information. You can always add more information by defining your own rules. Creating new rules provides the means for overriding existing management pack functionality and makes the management pack work more as you expect it to.

Once you decide to create a new rule, you have to consider how to create it. Of course, one of the considerations is the kind of rule you want to create. The "Using Rules" section of Chapter 5 provides a complete description of the various rule types. However, here's a quick list you can use for the example in this chapter:

- ◆ Event Based
 - ◆ Generic CSV Test Log
 - ◆ Generic Test Log
 - ◆ NT Event Log
 - ◆ SNMP Event
 - ◆ SNMP Trap
 - ◆ Syslog
 - ◆ WMI Event
- ◆ Performance Based
 - ◆ SNMP Performance
 - ◆ WMI Performance
 - ◆ Windows Performance
- ◆ Probe Based
 - ◆ Script (Event)
 - ◆ Script (Performance)
- ◆ Timed Commands
 - ◆ Execute a Command
 - ◆ Execute a Script

When you create a rule for optimization purposes, it's important to consider the rule that it replaces. Make sure you record any information about the old rule before you begin creating the new one. The following steps describe how to create a rule to override an existing one:

1. Open the Authoring tab.
2. Click Create A Rule in the Actions pane. You'll see the Create Rule wizard, shown in Figure 11.14.

3. Choose a rule type from the list. When replacing an existing rule, make sure you choose the same rule type. Otherwise, the rule won't work as expected.
4. Select a management pack. The "Customizing Management Packs" section earlier in this chapter tells how to create a new management pack. Click New if you want to create a new management pack for the rule that you're creating.
5. Click Next. You'll see the Rule Name And Description dialog box shown in Figure 11.15.
6. Type a rule name in the Rule Name field. When overriding an existing rule, give the new rule a similar name, but not precisely the same name as the original. For example, if the original rule is named Collect Health Service/ Active File Uploads, use a name such as Collect Health Service/ Active File Uploads (Performance Enhanced). The information in the parentheses tells the casual viewer what has changed in the new version of the rule.
7. Click Select. You'll see the Select A Target Type dialog box, shown in Figure 11.16.
8. Highlight the desired target and click OK.
9. Check Rule Is Enabled if you want to enable the rule immediately after you create it.
10. Click Next. At this point, you'll see various dialog boxes depending on the kind of rule you're creating. Refer to Chapter 5 for details on working with specific rule types.

FIGURE 11.14
Choose the type of rule you want to create for optimization purposes.

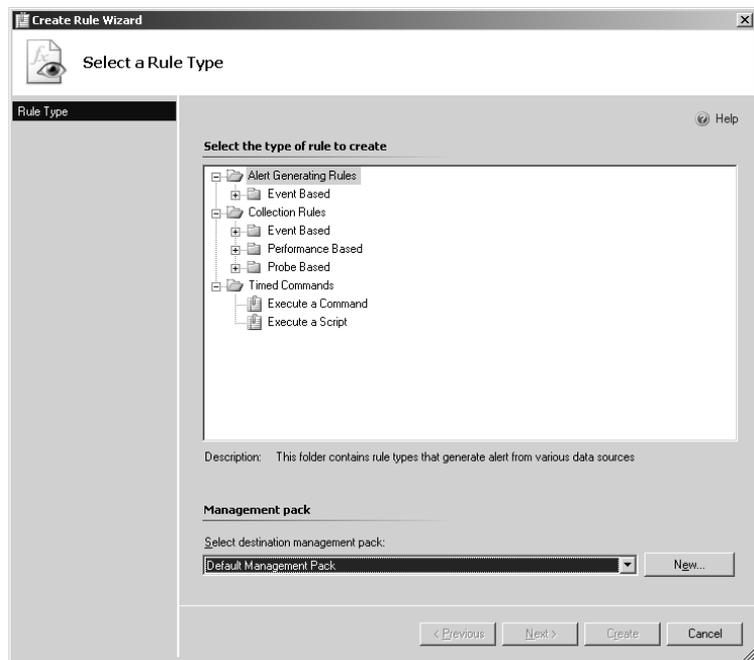


FIGURE 11.15
Provide a name, description, and a rule target for the new rule.

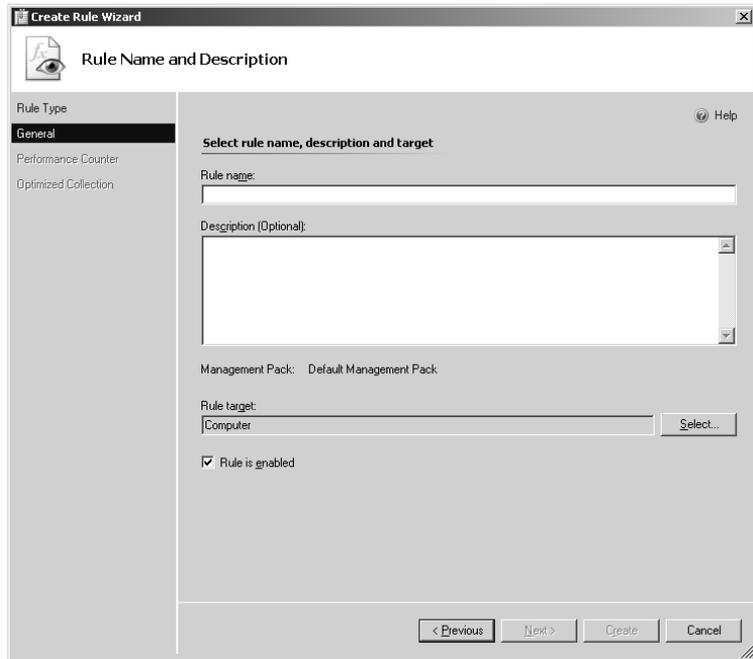
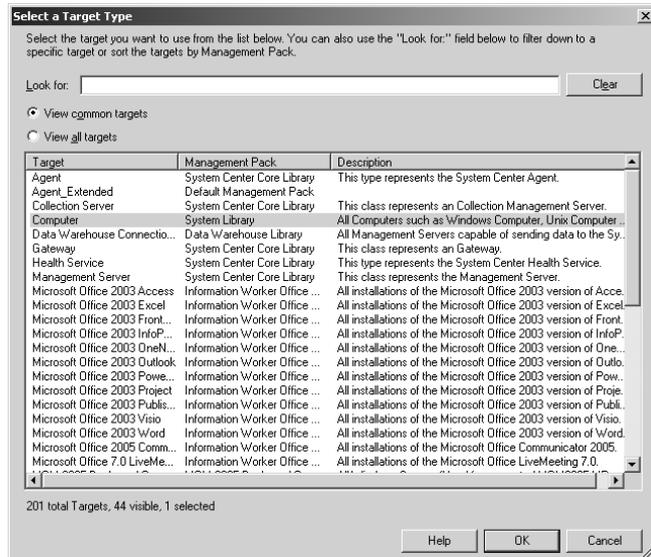
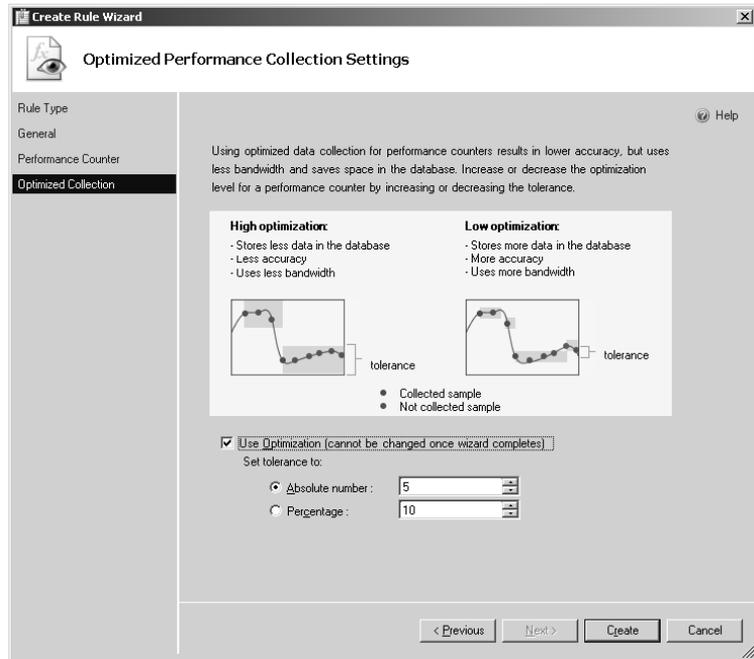


FIGURE 11.16
Choose a target type for the rule.



11. Define the rule specifics for your rule. Keep clicking Next until you see the Optimized Performance Collection Settings dialog box, shown in Figure 11.17. This is the only opportunity you have to select an optimization strategy. If you decide to change the strategy later, you must re-create the rule.

FIGURE 11.17
Choose an optimization setting for your rule to ensure it works as efficiently as possible.



TIP Providing an optimization strategy for the rules you create is one way to ensure you get better performance from your system. Because a third-party vendor can't guess about your system setup, most third-party rules don't include any optimization strategy. Adding a strategy to your rule will help you enhance system performance significantly.

12. Click Create. Operations Manager creates the new rule for you.

Self-Tuning Thresholds

What is normal? Many people ask this question, yet no one can seem to answer it because there isn't a correct answer. The question of normal depends on a significant number of factors. One system may have different normal operating characteristics even if it uses precisely the same hardware as another system. In the past, the lack of a definition for normal caused significant problems for administrators because a monitor set up for one machine probably wouldn't work for another machine even when both machines were in their normal state.

UNDERSTANDING SELF-TUNING THRESHOLDS

Self-tuning thresholds provide a way around the whole question of normal. The monitor you create goes through a learning process where it discovers what normal means in a specific situation. As the monitor continues tuning itself, the number of alerts the administrator sees decreases until alerts appear only when something is truly wrong with the system. Of course, you have to begin with a baseline threshold or the learning process would seem quite long indeed. Operations Manager automatically establishes a baseline for new monitors based on usage patterns and other data it acquires as it monitors the system.

A threshold can occur in a number of ways and the method of reacting to them varies according to the way you design the monitor. Monitors have to check for three conditions:

- ◆ The area above the baseline (a)
- ◆ The area within the baseline (w)
- ◆ The area below the baseline (b)

These three areas are important. For example, a CPU monitor may see the area above the baseline as anything over 90 percent. The area within the baseline may include everything from 10 percent to 89 percent. The area below the baseline may include everything 9 percent or less. The combination of these three areas defines the entire range of possible values, whether or not the device can achieve a particular value. Because monitoring situations vary, Operations Manager provides two different categories of self-tuning thresholds:

Two-State Monitor A two-state monitor monitors only two of the three areas. You can define an alert between any of the two areas: a-w, w-b, or a-b. This form of monitor is perfect for overflow or underflow conditions. For example, when monitoring the CPU, you don't care that the CPU is idle, but you do care if it goes beyond the normal 89 percent, so you set an alert for this condition.

Three-State Monitor A three-state monitor checks all three areas. It defines the healthy state as the area within the baseline. You choose one of the other two areas as the error state and the other as the warning state. A three-state monitor is useful in situations where you must maintain a specific level. For example, when monitoring the power settings for a UPS, you don't want the UPS to encounter an over-voltage condition (the error state) or an under-voltage condition (the warning state).

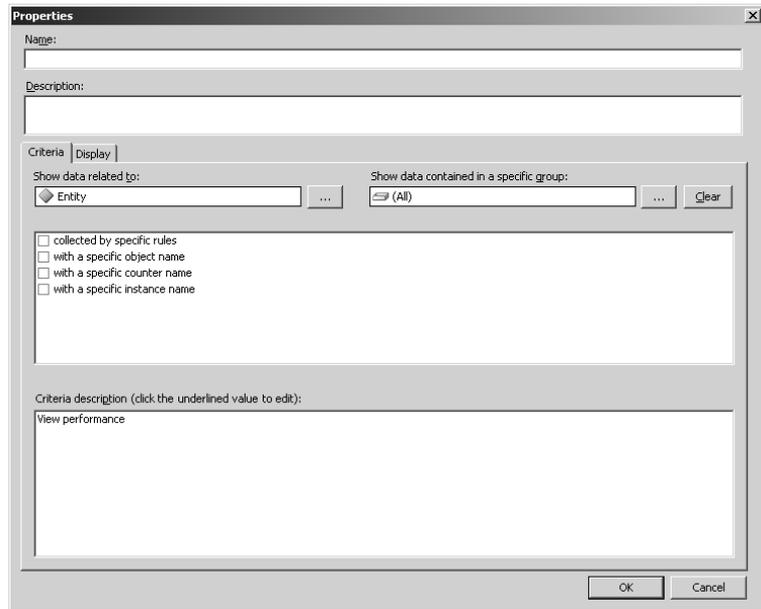
VIEWING A PERFORMANCE BASELINE

To work with self-tuning thresholds, you need to know about performance baselines. You can create a performance baseline to monitor any object that Operations Manager supports. For example, you can track the amount of memory that the server uses on average. It's often helpful to view the performance baseline of a particular object before you create a self-tuning threshold to monitor it. The following steps describe how to view a performance baseline:

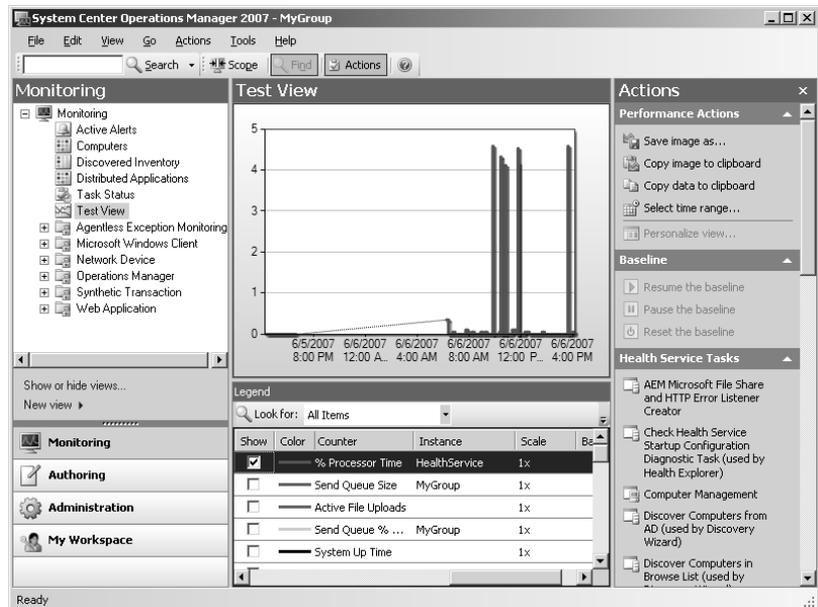
1. Open the Monitoring tab.
2. Right-click the Monitoring entry and choose New > Performance View from the context menu. You'll see the Properties dialog box, shown in Figure 11.18.
3. Type a name for the view in the Name field and then click OK. Operations Manager creates the new view for you.
4. Select the view you just created. You'll see a list of common entries in the Legend.
5. Select one of the entries, such as Process\% Processor Time, for the server. This entry shows the Process object, the % Processor Time counter, and the HealthService instance, as shown in Figure 11.19.
6. Select other performance baseline entries in turn to see how they appear.

FIGURE 11.18

Define a new performance view to see a performance baseline.

**FIGURE 11.19**

See the performance baseline for Process\% Processor Time.



DEFINING A SELF-TUNING THRESHOLD

You can define a self-tuning threshold for any monitoring need that doesn't have a specific (static) range. It's important to keep the learning nature of this kind of monitor in mind because the initial monitor may raise alerts when there really isn't a problem. As the monitor learns the range of the object that it monitors, you'll see fewer false alarms and considerably more true alerts. The following steps describe how to create a self-tuning threshold:

1. Open the Authoring tab.
2. Choose the Authoring\Management Pack Objects\Monitors entry.
3. Click Scope to display the Scope Management Pack Objects By Target(s) dialog box shown in Figure 11.20. You may actually have to click the Scope button twice to display the dialog box when you already have the Scope button selected.
4. Type the name of the object you want to use, such as Windows Computer, in the Look For text box. Operations Manager shows you the object entry.
5. Check the object you want to use to create the monitor and then click OK. Operations Manager chooses just the object you selected in the Target area, as shown in Figure 11.21.
6. Right-click one of the entries under Entity Health (see Figure 11.21), such as Performance, and choose Create a Monitor > Unit Monitor from the context menu. You'll see the Select A Monitor Type page of the Create A Unit Monitor wizard, shown in Figure 11.22.
7. Open the Windows Performance Counters\Self-Tuning Thresholds folder and select one of the self-tuning thresholds in the list. (The "Understanding Self-Tuning Thresholds" section earlier in this chapter describes the meanings behind these entries.)

FIGURE 11.20
Choose a scope for
the new monitor.

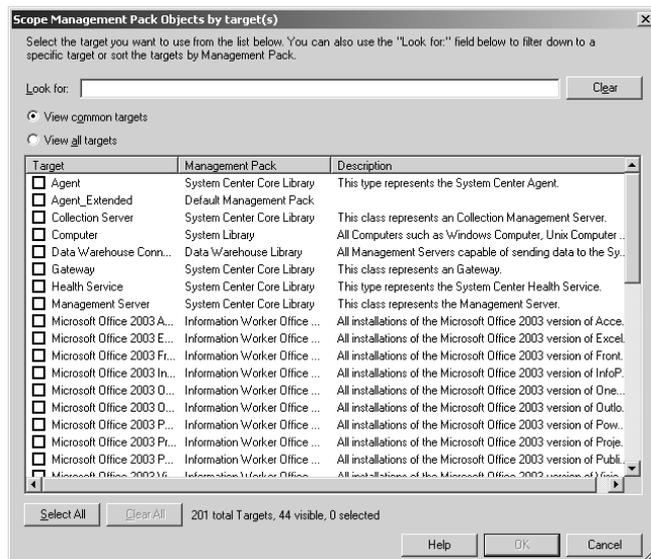
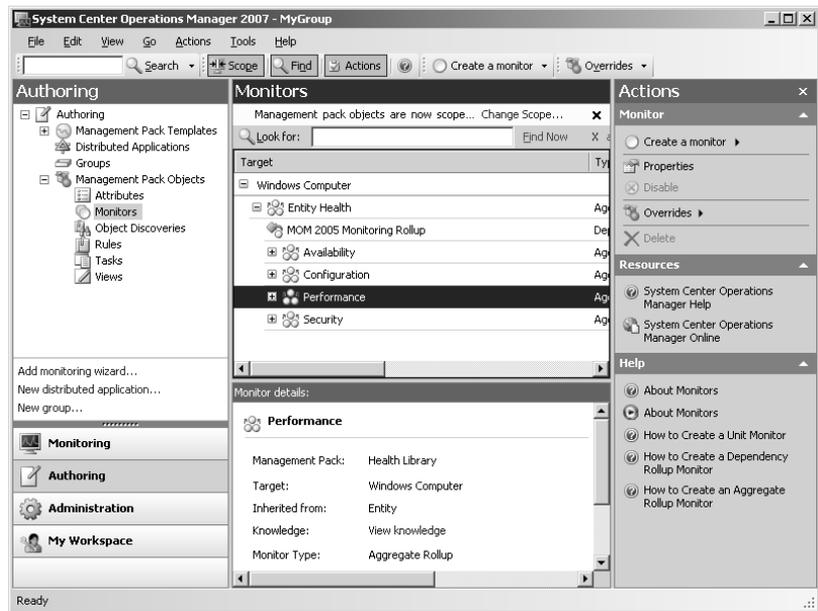
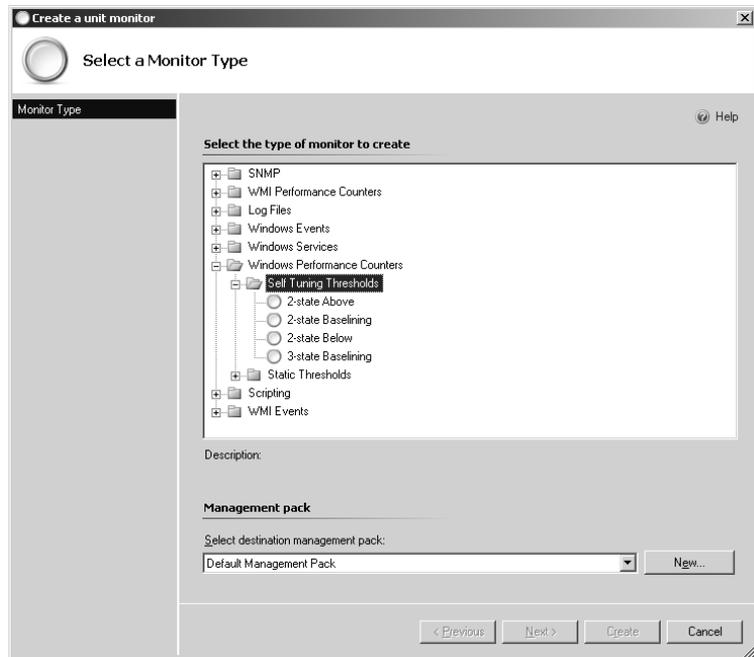


FIGURE 11.21

Verify that Operations Manager has selected the object you want to work with.

**FIGURE 11.22**

Choose the kind of monitor you want to create.



8. Click Next. You'll see the General Properties dialog box, shown in Figure 11.23.
9. Type a name and description for the self-tuning threshold. Make sure you provide descriptive information so others know how to use the monitor. In most cases, you won't need to change the Monitor Target, Parent Monitor, or Monitor Is Enabled option. If necessary, change the Monitor Target field to a custom target you've created.
10. Click Next. You'll see the Performance Object, Counter, and Instance dialog box, shown in Figure 11.24.
11. Click Browse. You'll see the Select Performance Counter dialog box, shown in Figure 11.25.
12. Click OK. Operations Manager adds all of the counter information to the Performance Object, Counter, and Instance dialog box.

NOTE Some counters don't include any instances. In this case, the Instance field appears blank. Seeing a blank Instance field doesn't always mean that the input information is incorrect.

13. Click Next. You'll see the Configure The Baseline Values Used To Determine The Thresholds dialog box, shown in Figure 11.26. This is where you include the information that Operations Manager will use to create an initial baseline for the self-tuning threshold. The default values generally provide a good starting point, but you can use other values when experience shows that the defaults won't work.

FIGURE 11.23
Provide a description of the self-tuning threshold that you want to create.

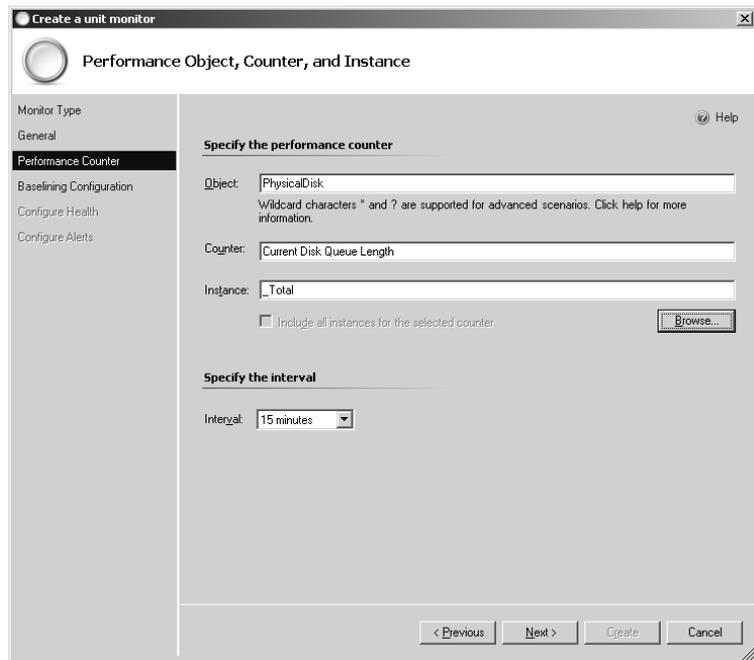
The screenshot shows the 'Create a unit monitor' dialog box with the 'General Properties' tab active. The 'Monitor Type' is 'General'. The 'General properties' section contains the following fields and options:

- Name:** A text input field.
- Description (Optional):** A multi-line text area.
- Management pack:** Set to 'Default Management Pack'.
- Monitor target:** Set to 'Windows Computer' with a 'Select...' button.
- Parent monitor:** Set to 'Performance'.
- Monitor is enabled:** A checked checkbox.

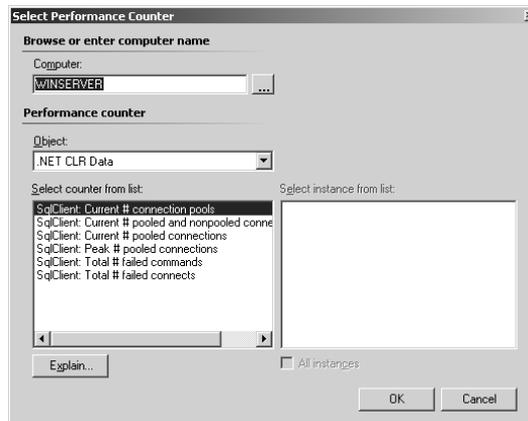
At the bottom of the dialog, there are four buttons: '< Previous', 'Next >', 'Create', and 'Cancel'.

FIGURE 11.24

Choose the performance object, counter, and instance you want to monitor.

**FIGURE 11.25**

Use this dialog box to make performance information selection easier.



14. Choose the threshold settings you want to use for the monitor.
15. Optionally, click Advanced when you need to set either the Learning Rate or Time Sensitivity value. When you're satisfied with the values, click OK to close the Baselining – Advanced dialog box.
16. Click Next. You'll see the Configure Health dialog box shown in Figure 11.27.

FIGURE 11.26
Set the information used to create a baseline configuration for the monitor.

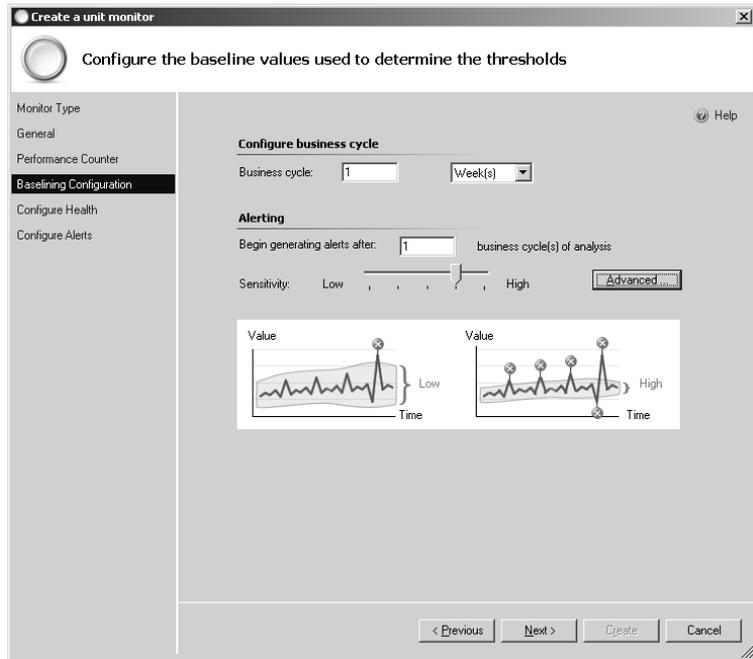
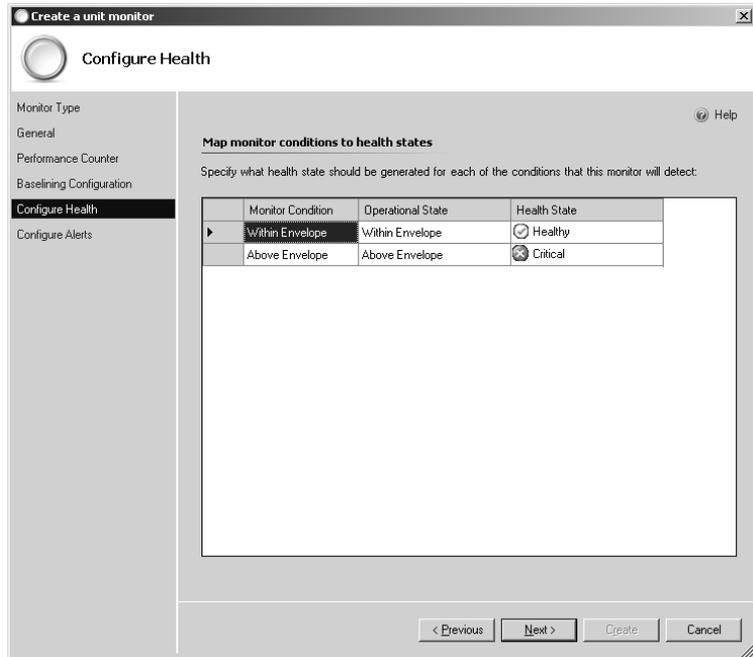


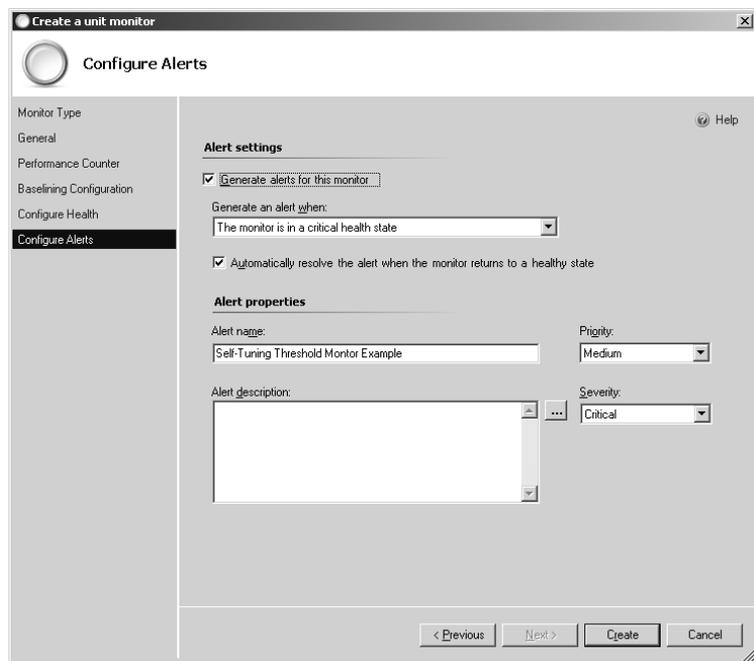
FIGURE 11.27
Set the limits that define the health of the monitor.



17. Type a value for the Operational State field for each of the entries. This value should tell the viewer the operational state of the self-tuning threshold. For example, you could simply type **Healthy** in the Within Envelope row if a within envelope state is indeed healthy.
18. Choose a Health State option for each of the entries. The Health State can be one of three values:
 - ◆ Critical
 - ◆ Warning
 - ◆ Healthy
19. Click Next. You'll see the Configure Alerts dialog box, shown in Figure 11.28. It's not always necessary to configure an alert for the self-tuning threshold. You may simply want to monitor a condition, rather than create an alert for it. When you don't want to create an alert, click Create at this point to create the new monitor.
20. Check Generate Alerts For This Monitor.
21. Choose one of the options in the Generate An Alert When field. If you choose The Monitor Is In A Warning Health State option, you'll receive alerts for both the warning and critical states.
22. Check Automatically Resolve The Alert When The Monitor Returns To A Healthy State option when you want the alerts turned off automatically. This is a good setting when the monitor doesn't reflect an actual failure, such as a temporary out-of-memory condition. However, you'll want to clear this option when working with a failure condition, such as a hard drive error rate that's too high.

FIGURE 11.28

Create an alert for the self-tuning threshold when you need to know about critical conditions.



23. Type a name for the alert in the Alert Name field and a description in the Alert Description field. Remember that this information will guide the viewer to fix a problem described by the monitor you create. Consequently, you need to make this information as descriptive as possible. Be sure to include information such as the associated object, counter, and instance as part of the description.
24. Set a priority and severity for the alert. Make sure you use values that actually reflect the alert status. A full hard drive is probably high priority and critical, while a processor that's working too hard might only rate a medium priority and a warning severity.
25. Click Create. Operations Manager creates the new self-tuning threshold for you.

Optimized Performance Counters

Performance counters are simple in theory but can become quite a resource sink in practice. There isn't anything magic about a performance counter; it simply counts something. The focus of a performance counter is a particular object, such as the processor. The particular counter determines what the counter is counting. For example, it may count the number of user tasks performed per second. In many cases, you can also tell a counter to focus on one particular object only, rather than all the objects of a given type. For example, when working with a multiprocessor system, you can choose to count only the user tasks for the second processor, rather than all the processors in a system.

You can't really change anything about a performance counter—a developer builds it into the code. However, you can change how you use the performance counter with Operations Manager. The two criteria for performance counter optimization include:

Resource Usage When you choose to monitor all the processors in a system, rather than the single processor you actually need to monitor, Operations Manager wastes the resources for the other processors. Even in a two-processor system, this means that the system will see 50 percent waste. Consequently, you can optimize performance counter use by selecting only the instances that you need to monitor, even if that means creating multiple monitors to do it.

Information Value It doesn't take long to realize that you can quickly go into information overload mode with Operations Manager. It provides so much information that you could find yourself looking in several directions at once while trying to figure out what to fix first. Adding more monitors than you need can significantly add to the information overload. After all, what you want to do is monitor conditions that you know or at least suspect will occur. When you monitor everything, you eventually find that you can't monitor anything very well.

The Bottom Line

View management pack optimization as an opportunity. Many people view optimization tasks as drudgework they must accomplish at some time. The optimization process is actually an opportunity in disguise because it helps you become more familiar with the product. As you become more familiar with the product, you discover ways to make it more efficient, see potential problem areas well in advance, and discover new ways to diagnose problems. In addition, optimizing forces an organization to make choices about how to use the management packs effectively. This action results in a setup that's easier to support and use. In short, optimization helps you create a better working environment.

Optimize rules to improve reliability, security, and performance. Someone one said that rules are made to be broken. In most cases, breaking the rules has serious consequences, but you have to break the rules in some cases when working with Operations Manager. The rules supplied with a management pack reflect the standard or default condition that the developer who created the management pack expected. Seldom does the real world actually reflect the default—just about every installation has quirks that make modification of the rules necessary. Of course, just because you have a license to break the rules doesn't mean you should do so arbitrarily. Make sure you have a good reason for creating a new rule, document the reason, and perform thoughtful analysis of the change you must make before you create the new rule. Always make sure you document every aspect of the new rule so that people who follow behind don't have to guess about your motivations.

Configure self-tuning thresholds to make monitoring more efficient Self-tuning thresholds use algorithms that allow them to determine what the operational level of an object is. As the self-tuning threshold runs, it gathers information about the object and determines if the threshold values need to be altered because of changes in the system. As the self-tuning threshold learns about how the objects perform, the threshold levels are adjusted, which reduces the number of false alerts and incorrect health status updates that are sent to the database.

Chapter 12

Backup, Restore, and Disaster Recovery

Disaster always strikes. The only problem is that you don't know when it will strike. Of course, you could sit huddled in the center of your office day and night, but few people would consider that a good strategy. After all, what happens if you're asleep when the disaster strikes? For that matter, what can you do about some disasters when they strike except ride them out? That's why preparedness is so important. Begin prepared for a disaster includes four stages: devising a strategy, preparing a response, compensating for the disaster, and recovering from the disaster. Unless you provide for all four stages, the results can be overwhelming failure of your entire setup.

This chapter helps you prepare for disaster by preparing a backup of your data (the preparing a response stage), restoring backups as needed (compensating for the disaster), and creating a recovery strategy. Throughout the chapter, you'll see a lot of planning information because nothing works without a plan. You must devise a strategy for surviving a disaster because everything else falls apart when you don't.

In this chapter you'll learn how to:

- ◆ Define a backup strategy
- ◆ Perform various kinds of backup
- ◆ Restore previously created backups
- ◆ Create a disaster recovery strategy

Backup Strategy

The most expensive and irreplaceable asset in your organization is the data it owns. You can replace the hardware and, in some cases, it's not too hard to replace the people, but the data is irreplaceable. In fact, as organizations make a move (finally) to less paper backup, the irreplaceable nature of the data increases. The loss of data can be so devastating that some companies go out of business after a data loss. Consequently, the one thing that you must protect is your organization's data.

Data doesn't just include critical business information any longer. Most people know that you have to back up the company database. All those invoices, contacts, contracts, and other information would be impossible to replace later should a disaster occur. However, most organizations don't take proper care of other databases that mean just as much. For example, many organizations don't properly back up the Windows Active Directory data, and later find that reconfiguring their network after a failure is a time-consuming process. Yes, you can do it, but you won't enjoy it. Creating a backup is a better plan. Likewise, you need to back up Operations Manager. When a disaster occurs, you want to be able to move your setup wherever you need it in order to continue working and you can't do that without a backup.

Microsoft doesn't really provide a backup strategy as part of the Operations Manager help file. You also won't find much in the way of help online. The following sections represent a homemade strategy that produces a usable server installation. The result is not quite a carbon copy of the installation you had earlier, but it's very close—close enough that you can get work done almost immediately using this strategy. Microsoft will likely provide better backup tools in the future, so it's important that you look for updates and additional information online since backup is such an important topic.



Real World Scenario

BACKUPS ALONE DON'T WORK

You've probably heard so many stories about people not performing backups that it seems odd that anyone would make the mistake of leaving their data to chance any longer. Most organizations do perform backups today, if for no other reason that the insurance company requires it. However, backups alone don't work, as witnessed by a company that performed daily backups only to find that their tape drive was faulty. The tape head had stopped working some time earlier, so the tapes were effectively blank. The operator made a daily backup of nothing. Obviously, testing should be a part of your backup strategy. Unless you test the backup, you have no idea of whether it'll actually work.

Another organization had faithfully made backups over several years and had even tested them. The same person performed both the backup and the test. Unfortunately, the person was only testing to see that the data was actually stored on the tape and not that the data was complete or even current. The day of disaster arrived and the person responsible for making the backup proudly presented the backup for restoration. The backup was incomplete, it turned out, and it included old data that the company had abandoned years before and maintained only for archival purposes. Consequently, the good backup turned out to be less than helpful. The lesson learned, in this case, was to have one person perform the backup and another restore it to a blank machine for full testing. Otherwise, you can't be sure that the backup is actually useful.

Operations Manager backups require that you have a consistent backup strategy in place—one in which backups do occur often enough to let you get up and running quickly in case of disaster. You need some type of distributed storage plan to ensure that the backups are actually available—a backup vault in the same city might not provide sufficient protection in case of a natural disaster such as an earthquake. The backups you create must contain valid data and you have to test them regularly by installing the data on a fresh machine. This might seem like a lot of work, but the alternative is spending a lot more time when you can least afford it reconfiguring systems because you didn't have a reliable backup.

Backing Up Secure Storage

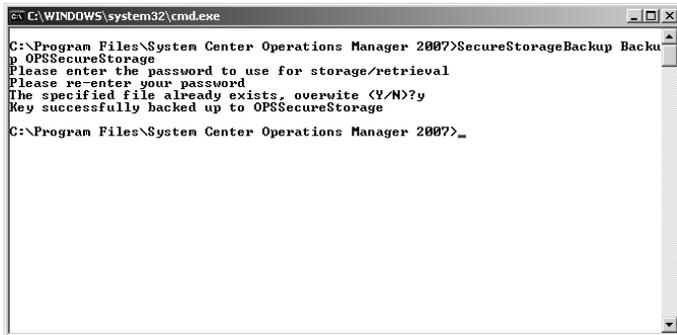
This is the only part of the backup process that Microsoft documents and even it isn't ready for prime time when you start out with Operations Manager. The first step you must perform is to copy the `SecureStorageBackup.exe` file from the `\SupportTools` folder on the Operations Manager media to the `\Program Files\System Center Operations Manager 2007` folder on your system. This tool helps you create a backup of the keys required to make Operations Manager work. Without these keys, you can't hope to make a restore work—it's like having a car without the keys to drive it.

NOTE Don't attempt to use the `SecureStorageBackup.exe` file from the support media. You'll see an error message when you attempt to use the utility this way because the utility won't be able to find the files it needs. The only location where this utility works is the `\Program Files\System Center Operations Manager 2007` folder.

Once you copy the `SecureStorageBackup.exe` file to the proper location, you can make the backup. Open a command line and change directories to the `\Program Files\System Center Operations Manager 2007` folder. Type **SecureStorageBackup Backup OPSSecureStorage** at the command prompt and press Enter. The `SecureStorageBackup` utility will ask you a series of questions, as shown in Figure 12.1 (you'll see only the overwrite question at the end of the sequence when the backup file already exists).

FIGURE 12.1

Use the `SecureStorageBackup` utility to create a backup of the keys to your Operations Manager installation.



```

C:\WINDOWS\system32\cmd.exe
C:\Program Files\System Center Operations Manager 2007>SecureStorageBackup Backup
OPSSecureStorage
Please enter the password to use for storage/retrieval
Please re-enter your password
The specified file already exists, overwrite (Y/N)?y
Key successfully backed up to OPSSecureStorage
C:\Program Files\System Center Operations Manager 2007>_

```

Answer each question in turn and you'll have a backup of secure storage. You must include the `OPSSecureStorage` file as part of your backup or you won't have the key when you need it. See the "Restoring Secure Storage" section later in this chapter for details on restoring this file.

WARNING You must save your secure storage password in a safe location. Microsoft recommends a secure storage password at least eight characters long. You can override this default when the `SecureStorageBackup` utility asks, but using a short password (or no password at all) isn't recommended if you want to keep your Operations Manager setup secure. Anyone who gains access to your backup (including this secure storage backup) could easily view the contents of the database and find ways to compromise your system.

Backing Up the Registry

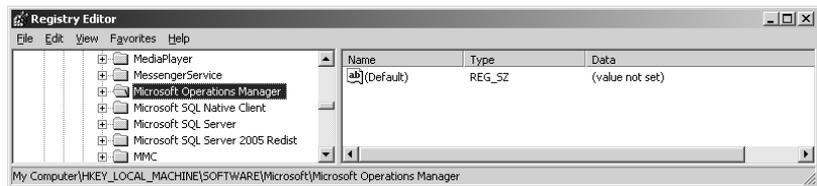
Backing up secure storage saves the keys to your setup, but not the actual settings. To create a backup of your settings, you must resort to using the `RegEdit` tool to export the required keys. This technique works only when you want to create a duplicate of your original setup. All of the particulars must remain the same, including the machine name and the names of the local administrator accounts. The following steps describe how to perform the global setting backup:

1. Select Start ➤ Run. In the Open field, type **RegEdit** and click OK. You'll see the Registry Editor utility. The backup takes place in two phases. You always perform the first phase, which is to back up the global setups. The second phase is to back up individual administrator settings.

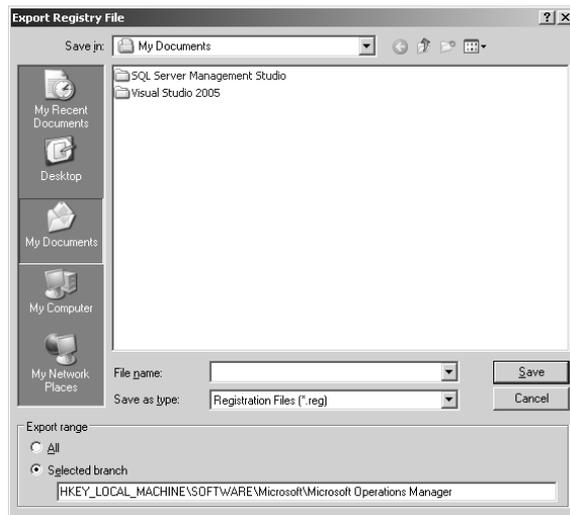
2. Locate the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft Operations Manager key, shown in Figure 12.2.
3. Choose File ➤ Export. You'll see the Export Registry File dialog box, shown in Figure 12.3. The Registry Editor should show the correct key in the Export Range area, as shown in Figure 12.3. If not, make sure you've selected the key mentioned in step 2.
4. Type a name for the backup file, such as **OpsMgrGloba1**, in the File Name field.
5. Click Save. The Registry Editor saves a copy of the Operations Manager global settings for you. At this point, you can exit the backup procedure, but you won't save any individual administrator settings. You should at least create a backup of the main administrator's settings.
6. Locate the HKEY_CURRENT_USER\Software\Microsoft\Microsoft Operations Manager key.
7. Choose File ➤ Export. You'll see the Export Registry File dialog box, shown in Figure 12.3.
8. Type a name for the backup file, such as **OpsMgrAdmin**, in the File Name field and click OK to save the settings.
9. Perform steps 6 through 8 for each of the other administrators. You can also use these steps to back up administrator settings on remote systems.

FIGURE 12.2

Use the Registry Editor to create a backup of the settings keys for Operations Manager.

**FIGURE 12.3**

The Registry Editor should provide all the settings you need except a filename.



Backing Up the Operations Manager 2007 Database

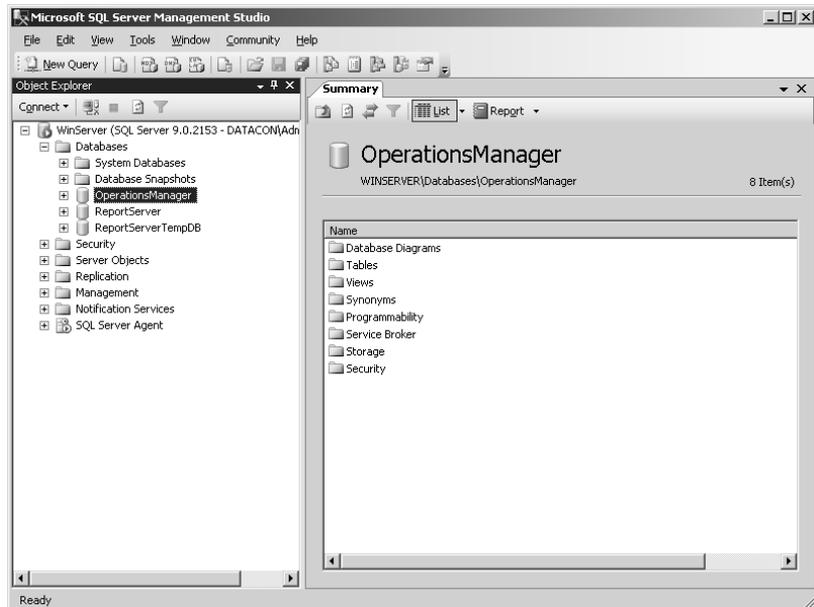
The Operations Manager database is managed by SQL Server, which means you must work with SQL Server to back it up. The person performing the backup must have rights to work not only with Operations Manager but with SQL Server as well—at least to perform backups of the OperationsManager database. The following steps tell you how to perform this part of the process:

1. Select Start > Programs > Microsoft SQL Server 2005 > SQL Server Management Studio. You'll see the login dialog box shown in Figure 12.4 for the SQL Server Management Studio.
2. Provide the proper credentials to log into the Database Engine. The SQL Server Management Studio opens.
3. Open the Databases\OperationsManager folder, as shown in Figure 12.5.

FIGURE 12.4
Provide the credentials required to log into SQL Server Management Studio.



FIGURE 12.5
Select the database you want to back up.



4. Right-click OperationsManager and choose Tools ➤ Back Up from the context menu. You'll see the Back Up Database dialog box, shown in Figure 12.6. Most of the settings are acceptable, but you need to verify and change a few of the settings.
5. Verify that the Backup Type field has Full selected. Otherwise, you can't count on your backup providing a complete restore later.
6. Type a name in the Name field or accept the default. The important issue is that you choose a descriptive backup name so that everyone knows what the file contains.
7. Verify that the Backup Set Will Expire After option is selected and set to 0 days so the backup doesn't expire.
8. Choose a backup location for the database or accept the default option.
9. Select the Options page. You'll see a list of options such as the ones shown in Figure 12.7. These settings control how SQL Server makes the backup, and some of the defaults won't work for an Operations Manager backup.
10. Choose the Overwrite All Existing Backup Sets option to ensure you have just the latest backup available for an Operations Manager restore.
11. Check the Verify Backup When Finished option to ensure you have a good backup. Otherwise, you could create a backup that won't work.
12. Click OK. SQL Server will begin creating the backup file. When the backup is complete, you'll see a success message. Click OK to clear the success message.

FIGURE 12.6
Select the database
you want to back up.

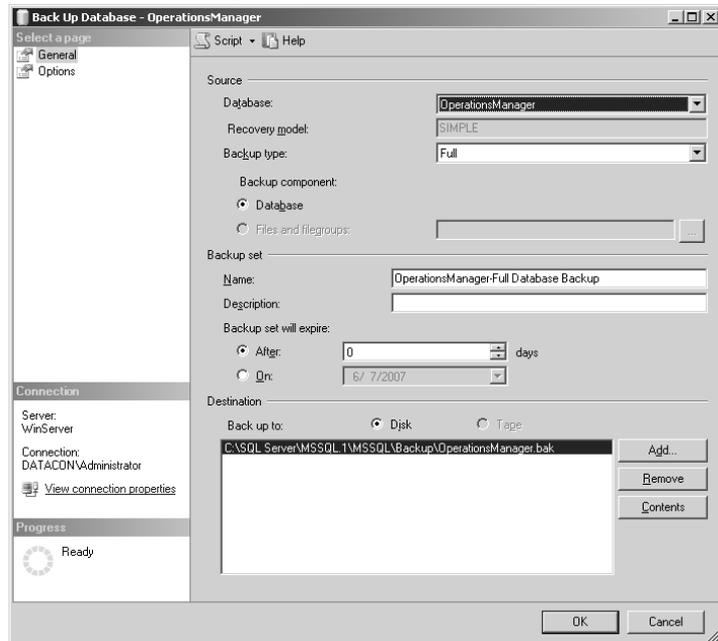
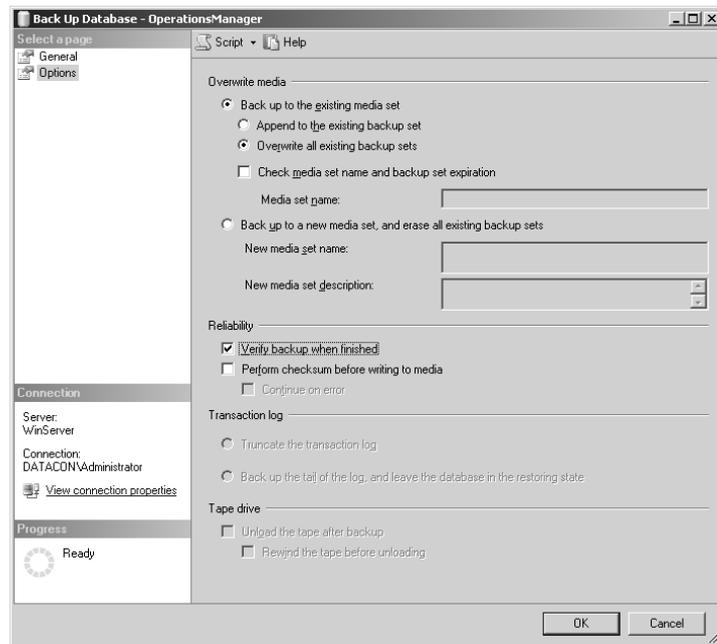


FIGURE 12.7
Set the backup options
to ensure the backup
works correctly.



Backing Up Management Packs

You'll need to perform a backup of more than just the management packs. The other sections have created a number of files that you need to back up as well. To perform this task, you need a good backup program and some type of portable media, such as CDs, DVDs, or tape drive. Make sure the backup program includes a full verify feature because you need to verify the backup when it's complete against the original files. Follow the backup instructions that come with the backup program you purchased to configure the backup program and media. You need to back up the following files and folders:

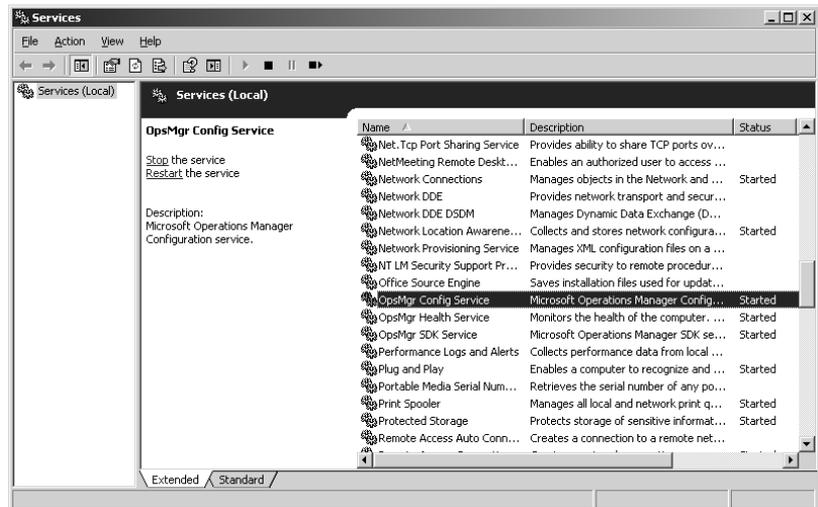
- ◆ \Program Files\System Center Operations Manager 2007\OPSSecureStorage (the secure storage file)
- ◆ The registry files you created, such as OpsMgrGlobal.reg
- ◆ \Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Backup\OperationsManager.bak (the Operations Manager database)
- ◆ The \Program Files\System Center Management Packs folder
- ◆ The \Program Files\System Center Operations Manager 2007\Config Service State folder
- ◆ The \Program Files\System Center Operations Manager 2007\Health Service State folder

- ◆ The \Program Files\System Center Operations Manager 2007\Health Service State\Health Service Store folder
- ◆ The \Program Files\System Center Operations Manager 2007\Health Service State\Management Packs folder

Before you can perform the backup, you must stop the Operations Manager. Otherwise, the files you need to back up are going to be in use. The following steps tell how to perform this task:

1. Close Operations Manager if you have it open.
2. Open the Services console found in the Administrative Tools folder of Control Panel.
3. Locate the OpsMgr services, as shown in Figure 12.8.
4. Highlight each OpsMgr service in turn and click the Stop link. You must stop all of the OpsMgr services before you can perform a backup.
5. Perform the backup using the instructions that come with your backup program.
6. Highlight each OpsMgr service in turn in the Services console and click Start.

FIGURE 12.8
You must stop the OpsMgr services before you perform a restore.



Restoring Your Environment

As previously mentioned, a good backup of your system, one that you've tested as far as possible, is akin to buying insurance. At some point, you'll need to break out your insurance policy and use it to restore your system. Although it's hard to categorize disasters, the restoration process normally follows two courses:

- ◆ An existing system has had a failure that causes damage to the setup.
- ◆ A new system requires configuration after the existing system fails completely.

In both cases, you need to replace files that are somehow not usable. Operations Manager may start, but it won't function sufficiently well to perform any management tasks. The only difference between the two scenarios is that when you install a new system, you have to perform the installation process first. The following sections describe the process you use to restore Operations Manager functionality after a disaster.

Performing the Initial Installation

When your old server dies, you need new hardware, which means you don't have the Operations Manager software installed on your system. The backup you create doesn't provide a complete installation—all you saved were the settings you require for Operations Manager. Use the procedures found in Chapter 2, "Installing System Center Operations Manager 2007," to create a basic Operations Manager setup. Perform the rest of the restoration before you start Operations Manager. Otherwise, you may find that the restoration doesn't work as well as anticipated and it could fail to work at all. You won't need to create accounts or perform other basic setup tasks because the backup restores this information for you—all you need to perform is the initial Operations Manager installation.

Restoring the Management Packs

Restoring the management packs also means restoring the rest of the files you need to complete the restoration process. You'll need the same backup program you use to create the backup and the removable media containing the backup files. Follow the procedures your backup software provides for performing a restore. This chapter assumes that you're using the same configuration as your original server, so you might end up making a few changes to the procedures when you use a different setup. Use the following steps to ensure the restore process works as intended:

1. Close Operations Manager if you have it open.
2. Open the Services console found in the Administrative Tools folder of Control Panel.
3. Locate the OpsMgr services, as shown in Figure 12.8.
4. Highlight each OpsMgr service in turn and click the Stop link. You must stop all of the OpsMgr services before you can perform a restore.
5. Perform the restore using the instructions that come with your backup program.
6. Highlight each OpsMgr service in turn in the Services console and click the Start link.

Restoring Secure Storage

After you reinstall Operations Manager, you'll want to restore secure storage. If you haven't done so already, copy the `SecureStorageBackup.exe` file from the `\SupportTools` folder on the Operations Manager media to the `\Program Files\System Center Operations Manager 2007` folder on your system. Now that you have the `SecureStorageBackup.exe` file to the proper location, you can make the backup. Open a command line and change directories to the `\Program Files\System Center Operations Manager 2007` folder. Type **SecureStorageBackup Restore OPSSecureStorage** at the command prompt and press Enter. The `SecureStorageBackup` utility will ask you a series of questions, as shown in Figure 12.9. Answer each question in turn and you'll restore secure storage. (Depending on your setup, you may not see the question about overwriting secure storage.)

FIGURE 12.9

Use the SecureStorage-Backup utility to create a backup of the keys to your Operations Manager installation.

```

C:\WINDOWS\system32\cmd.exe
C:\Program Files\System Center Operations Manager 2007>SecureStorageBackup Restore OPSSecureStorage
Please enter the password to use for storage/retrieval
The Secure Storage Keys already exists, overwrite (Y/N)?y
Key successfully restored from OPSSecureStorage
C:\Program Files\System Center Operations Manager 2007>_

```

Restoring the Registry

The backup of the registry created one or more REG files. All you need to do to restore the registry is double-click the REG files. When the Registry Editor asks whether you're sure you want to add the information in the REG file to the registry, click Yes. You need to perform this task once for each of the REG files you created as part of the backup.

Restoring the Database

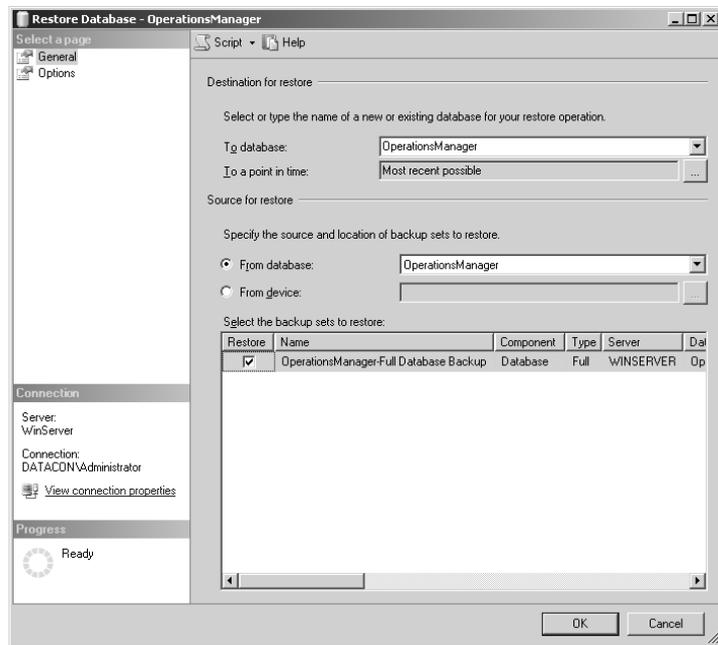
When you install Operations Manager, it creates a default OperationsManager database for you. However, this database doesn't contain any of your settings. You need to restore the original database before you can begin performing any useful work with Operations Manager. To restore the database, you must use the SQL Server Management Studio, which means having rights to work in SQL Server. Unlike in the backup process, you must have rights to create databases and perform other administrator-level tasks within SQL Server to perform a restore. The following steps tell you how:

1. Copy the `OperationsManager.bak` file to the `\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Backup` folder. (You may have your data located in another place on the server, so make sure you choose the correct backup folder.)
2. Close Operations Manager if you have it open.
3. Open the Services console found in the Administrative Tools folder of Control Panel.
4. Locate the OpsMgr services, as shown in Figure 12.8.
5. Highlight each OpsMgr service in turn and click the Stop link. You must stop all of the OpsMgr services before you can perform a restore.
6. Select `Start > Programs > Microsoft SQL Server 2005 > SQL Server Management Studio`. You'll see the login dialog box shown in Figure 12.4 for the SQL Server Management Studio.
7. Provide the proper credentials to log into the Database Engine. The SQL Server Management Studio opens.
8. Open the `Databases\OperationsManager` folder as shown in Figure 12.5.
9. Right-click `OperationsManager` and choose `Tools > Restore > Database` from the context menu. You'll see the Restore Database dialog box, shown in Figure 12.10. You'll need to modify at least the source of the backup so that SQL Server uses the file you created earlier.

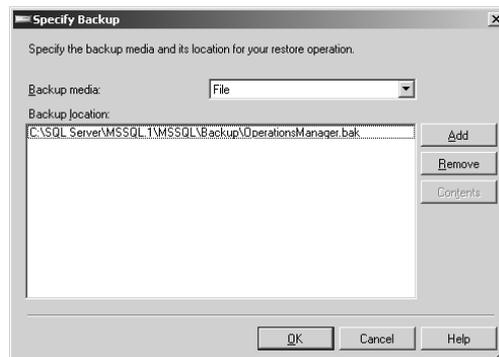
10. Select From Device. Click the ellipsis next to this option. You'll see the Specify Backup dialog box, shown in Figure 12.11.
11. Click Add. You'll see a Locate Backup File dialog box, where you can choose the backup file. The default location to look is the \Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Backup folder. You should see the OperationsManager.bak file in the list.
12. Select the OperationsManager.bak file and click OK. Click OK in the Specify Backup dialog box. The Restore Database dialog box should now show the location of the file you selected on the hard drive. In addition, the Select The Backup Sets To Retrieve list should now contain a single entry for the backup set you created.
13. Check the backup set to restore in the Select The Backup Sets To Retrieve list.

FIGURE 12.10

Set up SQL Server to restore the Operations Manager database.

**FIGURE 12.11**

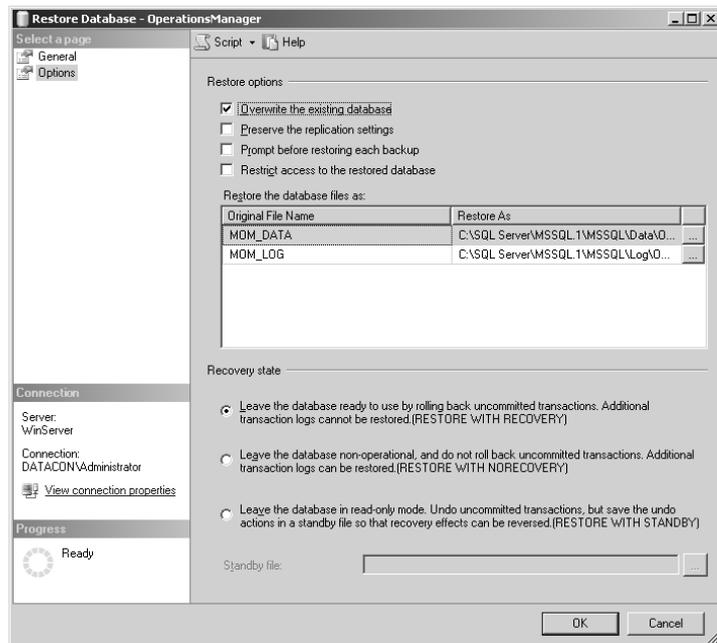
Tell SQL Server where to find the backup file for Operations Manager.



14. Select the Options page. You'll see the list of options shown in Figure 12.12.
15. Check the Overwrite the Existing Database option. If you don't check this option, the restore won't work properly. It will appear to work, but Operations Manager will continue using the default database installed during the installation process.
16. Click OK. SQL Server will begin the restoration process. When the restore is complete, you'll see a success message. Click OK to clear the success message.
17. Highlight each OpsMgr service in turn in the Services console and click Start.

FIGURE 12.12

Define how SQL Server should perform the Operations Manager database restore.



Disaster Recovery

More and more companies are creating setups where they have two or more management centers in physically separate portions of the country or even in the world. The company designates a main site and several satellite locations. The satellite locations look like the main site, but they use data replicated from the main site in most cases. In other words, the main site still does all the work, but any of the satellite locations can immediately become a main site when disaster occurs. This kind of setup is extremely complex to set up, but using it means that you have an instant replacement for any site that becomes nonfunctional due to any disaster. Even a natural disaster won't likely affect your system when you have the satellite locations in different areas. The following sections provide some basic concepts for disaster recovery using satellites.

NOTE Because of the complexities of creating this kind of setup, there isn't any way this chapter can provide a precise procedure for performing this task. Your organization's network and satellite setup are likely to have unique configuration issues that a single procedure can't cover, and it isn't possible to discuss all the potential permutations in a book. However, this chapter does provide good general advice on creating a satellite setup so you know how to get started and what you need to provide for a successful installation.

Using OpsManager 2007 to Monitor Your Disaster Recovery Site

It's possible that your Operations Manager setup can still fail even if you provide several satellite setups in other locations. Disasters could overcome more than one site at once, or you might not have planned the setup as well as you thought you did. The one thing you don't want to happen is have several servers fail at one time. You also want to check the health of all the applications on the various servers and ensure that everything's configured as it should be. The servers should all have the same service packs installed, and you'll want to be sure that the other administrators know the setup as well as the administrators for the main site do. All of these concerns reflect a need for monitoring, and you'll want to use Operations Manager to do it.

When setting up the monitoring strategy, you don't want to configure every site to monitor every other site because that would consume a large amount of system resources. What you really want to do is set up the main site to monitor all the satellites. The main site administrators then know that all the satellites are ready to use in case of a disaster. When a disaster does strike, the designated satellite should take over and begin monitoring all the remaining satellites as the new main site.

Replicating Operations Manager Data to Your Disaster Recovery Site

When creating a disaster recovery site, you need to replicate the setup at the main site. You can use the procedures found in this chapter to create a backup of the main site and then restore it to the satellite location. The only problem is that restoring the backup only provides you with a snapshot of the main site when you made the backup. Obviously, this strategy won't work; you need to replicate the OperationsManager database on the other sites.

To perform the level of replication required for Operations Manager, you must have the Enterprise Edition of SQL Server installed on every site. You create a publication for the database on the main site. The publication must include everything in the database. The reason you need everything is that if the main site goes down, you must have a complete copy of the database to keep Operations Manager running. This is also the reason that you need the Enterprise Edition of SQL Server—other editions don't provide the support required to replicate everything in a database to another location fully.

Once you have a publication that the satellites can subscribe to, you must create subscriptions at each of the satellite sites to the publication that you've created on the main site. The subscription will let the satellites obtain all the information they require (essentially the entire database) from the main site. When the main site goes down, every satellite will have a reasonably current copy of the entire main site database and will be able to continue providing Operations Manager support with few, if any, problems.

The Bottom Line

Define a backup strategy. Most companies at least know that they need to perform backups, but they do so in a haphazard fashion. The backups aren't consistent, tested thoroughly, or used to reconstruct a setup until an actual emergency occurs. When the emergency occurs, it's too late to test the backup or do anything about the deficiencies in the backup strategy. The only way to ensure that your backup strategy for Operations Manager actually works is to create the backup and restore it to a new machine—one that doesn't have any other software installed. Make sure the environment on the new machine isn't compromised by performing a complete reformat of the drive and installing a fresh copy of the operating system before you begin. If you can restore Operations Manager to a working state from your backup, then you have a backup strategy that works. Otherwise, you need to keep working with the backup strategy until you find one that does the job.

Perform various kinds of backup. Backing up the data on your drive isn't sufficient. An Operations Manager backup has to provide more than just the data in order to work because Operations Manager requires proper access to the database and must have the correct settings in place. In addition to your data, you need to back up the registry keys, secure storage, and your management packs regularly. It's also a good idea to include any tools or other software you need to restore Operations Manager with the backup. For example, you'll want to include service packs with the backup as Microsoft issues them because your current setup will have those service packs. The backup you create must be able to provide a completely usable Operations Manager installation on a new machine—you can't assume the new system will have anything installed on it at all.

Restore previously created backups. Restore your backups regularly to ensure they actually do contain the data you think they provide. Test your setup at specific intervals to ensure that you maintain the skills required to perform a restore in an emergency. Think about it this way: when the emergency occurs, you won't have time to learn how to perform the restore—you have to know how to perform the process. You should be able to go through the steps to perform a restore without much help.

Create helpful restoration aids. No one remembers every step precisely relying on memory alone. That's why airline pilots rely on checklists. Make sure you that create a checklist for restoring Operations Manager for your company. The time you take to create a checklist today will make restoration considerably easier in an emergency. Yes, you'll already have a good idea of what to do, but the checklist will serve to prod your memory and help you get the job done faster and with fewer errors.

Have a human backup. You're going to need a backup for yourself as well. What happens if you get stuck on the road leading out of town in a natural disaster? The company can't wait until you arrive several hours late to get the secondary site set up and ready for use. Always have a backup operator trained to perform the restoration should the need arise. This person should work side by side with you during testing and should be able to use the checklist you create to fill in any gaps.

Create a disaster recovery strategy. Disasters happen. You're going to encounter a disaster at some point—the only certainty you'll have when working in IT. The only problem is that the disaster can take many forms and it always happens when you least expect it. Backup, disaster recovery plans, recovery training, and the like are all forms of insurance. You invest in these items today hoping that you won't need them, but knowing that you probably will and will be happy about the investment when the time comes.

Chapter 13

Troubleshooting

Finding problems with your setup is never fun. In fact, it's probably one of the worst parts about administering a network. However, when the problem isn't with the network and it's with your main diagnostic tool, things become even worse because now your monitoring tool isn't working. This chapter discusses some troubleshooting tools for Operations Manager. You'll begin by discovering that there are tools to fix Operations Manager. Although these tools can't guarantee absolute success, they can get you started in the right direction.

Troubleshooting falls into three major areas when working with Operations Manager: agents, management packs, and security. Of the three, security issues can be the most troublesome because you aren't just working with the local system. A security issue can pop up on any machine on the network, which means that you can invest a lot of time into tracking down the precise source of an error. Agents are usually the easiest of the three to fix because you get all kinds of good visual cues about the problem. For example, visiting the Event Viewer usually yields useful clues into the agent problems. The management packs fall into a middle ground of difficulty. The level of difficulty depends largely on the complexity of the management pack and the skill of the person who authored it in creating a good design.

In this chapter you'll learn how to:

- ◆ Locate and use troubleshooting tools
- ◆ Troubleshoot agents
- ◆ Troubleshoot management packs
- ◆ Troubleshoot security issues

Operations Manager 2007 Troubleshooting Tools

Administrators can become gadget freaks. Any gadget that comes along attracts their attention and if it does anything even remotely interesting, it ends up in their virtual toolbox. Being a gadget freak isn't necessarily a bad thing as long as you keep track of the gadgets and know which one to use in a crisis. The right gadget at the right time can save you considerable time and effort. In some cases, having a gadget is the only way to accomplish a task. For example, if you want to convert your existing management packs to something that Operations Manager 2007 can use, you must have the MP2XML utility found only in the MOM 2005 Resource Kit. Not everyone needs to perform this task, so not everyone needs this particular tool; but the MP2XML utility is irreplaceable when you do need it. The following sections discuss some other irreplaceable tools you should consider adding to your toolbox.

Resource Kit Tools

As of this writing, Microsoft hasn't provided a specific resource kit for Operations Manager 2007. However, the older resource kit for MOM 2005 still has some interesting offerings in it. You can download this older resource kit at <http://www.microsoft.com/technet/opsmgr/2005/downloads/tools/reskit.msp>. Most of the tools definitely won't work with Operations Manager 2007. For example, don't expect to use most of the Management Pack Toolkit. However, the Clean-up MOM tool does work. The MOM Inventory tool is also exceptionally useful when you upgrade your current installation.

Although some of these tools are useful, many aren't. In general, you'll want to wait until Microsoft releases a resource kit for Operations Manager 2007 to obtain anything of significant value. Look at this older resource kit as a source of tools to help you overcome installation and file conversion problems.

WHERE DO YOU GO FROM HERE?

Most people need somewhere to go to discover new techniques, learn about fixes for problems, and discuss issues with other people. Websites provide static information that's helpful, in a lot of cases, but you have to take it as is because there isn't any way to interact with anyone. Newsgroups (list servers) provide a significant level of interaction, but sometimes it's hard to find someone with any more experience than you have and the information isn't presented very professionally. If you're truly interested in finding out about the latest Operations Manager tools from Microsoft, don't waste too much time browsing the Microsoft websites or newsgroups; try the various blogs online instead. Blogs provide professionally presented information and a means to discuss the blog content, which many professionals find quite helpful. Here's a list of blogs you should try when working with Operations Manager:

Clive Eastwood's Blog <http://blogs.technet.com/cliveeastwood/default.aspx>

Pete's Management Blog <http://www.it-jedi.net/>

Operations Manager Product Team Blog <http://blogs.technet.com/momteam/default.aspx>

Mother (Microsoft Operations Manager 2005 and System Center Operations Manager 2007) <http://blogs.msdn.com/incarnato/>

System Center Operations Manager Command Shell <http://blogs.msdn.com/scshell/>

Jakub@Work (Programming with System Center Operations Manager) <http://blogs.msdn.com/jakuboleksy/>

Advisec Blog <http://advisec.wordpress.com/>

Ops Manager – John Hann <http://msmvps.com/blogs/jfhann/default.aspx>

Management Pack Catalogs

You may not have found the particular management pack you wanted mentioned in this book. In fact, Microsoft produces new management packs all the time. Because this information is constantly changing, Microsoft or a third-party vendor may answer a need you have today with a new product tomorrow.

NOTE Microsoft provides other catalogs on their website, such as the Microsoft Operations Manager Product Connectors at <http://www.microsoft.com/mom/partners/momprodconnectors.aspx>. Unfortunately, these catalogs are still set up for MOM 2005. Since the Operations Manager 2007 structure is so different from MOM 2005, you can't use these old solutions. Microsoft will likely update the other catalogs as some point.

The Management Pack Catalog appears at <http://www.microsoft.com/technet/prodtechnol/mom/catalog/catalog.aspx>. You can use it to search for management packs by keyword, by management pack category, or by company, as shown in Figure 13.1.

After you fill in the search criteria and click Search, you'll see a list of possible offerings such as the one shown in Figure 13.2. Notice that each of the entries tells you the product name, provides a good description of what the management pack will do for you, includes a company name, and tells you the date of the last management pack update. When you want to download one of the management packs, simply click the product name link. At the time of this writing, clicking the Microsoft link takes you to the MOM website at <http://www.microsoft.com/mom/default.aspx>. Clicking the other vendor links takes you to a "link not found" error page.

After you become familiar with the offerings that Microsoft provides, you don't even need to use the catalog search tool unless you need something special. Simply subscribe to the Really Simple Syndication (RSS) feed (<http://www.microsoft.com/windowsserver2003/evaluation/rss/momrss.aspx>) provided as part of the Management Pack Catalog web page. You can find the instructions for subscribing to an RSS feed in Outlook at <http://office.microsoft.com/en-us/outlook/HA101595391033.aspx>.

FIGURE 13.1
Use the Management Pack Catalog to search for the management pack you need.

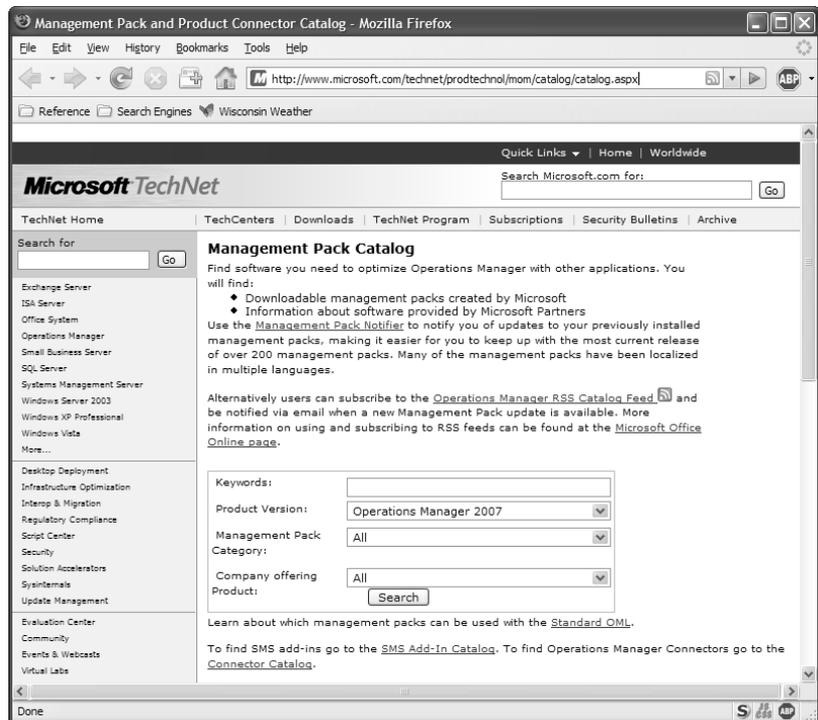


FIGURE 13.2
The list of management pack offerings for Operations Manager grows daily.

The screenshot shows a web browser window with the address bar containing the URL: <http://www.microsoft.com/technet/prodtechnol/mom/catalog/catalog.aspx?kw=8vs=2t>. The page title is "Management Pack and Product Connector Catalog - Mozilla Firefox". The main content area displays a search result summary: "Your query returned 41 result(s)". Below this is a table with the following data:

Product	Description	Company	Last Updated
Office SharePoint Server 2007	This Management Pack quickly brings any failures or configuration problems to your attention, which increases the availability and performance of Office SharePoint Server 2007. This Management Pack also provides the knowledge and expertise you need to leverage Operations Manager 2007 and get an immediate return. Requires OpsMgr 2007 MOM 2005 Backward Compatibility MP Update 6.0.5000.1 (or higher).	Microsoft	06/04/2007
Windows DHCP Server 2000/2003	Collects events generated by the Microsoft Dynamic Host Configuration Protocol Service (DHCP) in Windows 2000 and Windows Server 2003. Ensures this critical part of your Windows infrastructure is working correctly, and helps prevent costly service outages. Requires OpsMgr 2007 MOM 2005 Backward Compatibility MP Update 6.0.5000.1 (or higher).	Microsoft	06/04/2007
Windows File Replication Service 2000/2003 (FRS)	Monitors the performance and availability of Windows File Replication Service, ultrasound, and replica sets. Requires OpsMgr 2007 MOM 2005 Backward Compatibility MP Update	Microsoft	06/04/2007

NOTE Make sure you use the RSS feed if you want to know about new Operations Manager 2007 management packs. The Microsoft Management Pack Notifier for Microsoft Operations Manager 2005 at <http://www.microsoft.com/downloads/details.aspx?FamilyID=a24cea3a-1920-4b18-8cf2-8bf78c94c917> won't work with Operations Manager 2007 (even if you convert it).

Third-Party Tools

A search on the Internet won't yield too many hits for Operations Manager 2007 tools. Fortunately, Microsoft is doing most of the work by gathering a list of third-party tools for you. You can find this list of tools at <http://www.microsoft.com/systemcenter/opsmgr/partners/default.msp>. Here's a description of some of the more interesting third-party tools from Microsoft partners available as of the time of this writing:

AVIcode (<http://www.avicode.com/>) This company produces a special monitor that shows the status of all .NET applications on your system. You can use it to discover both application failures and performance bottlenecks.

F5 (<http://www.f5.com/>) This specialty monitor helps you assess the health of all F5 devices on your network.

Quest Software (<http://www.quest.com/>) The Quest offering is a connector that helps you monitor Unix and Linux systems. In addition, it helps you monitor Oracle databases.

Tidal Software (<http://www.tidalsoftware.com/>) The Tidal Software offering is a connector that helps you manage SAP solutions.

eXc Software (<http://www.excsoftware.com/>) The eXc Software offering is a connection that helps you monitor Linux systems. In addition, you can use it to monitor Oracle databases. This product actually provides connectivity for a number of other platforms, including VMWare, so you'll want to check the vendor website for details.

nworks (<http://www.nworks.com/>) The nworks offering is a connector that helps you monitor VMWare and Blackberry Enterprise Server.

Secure Vantage (<http://www.securevantage.com/ProductsACS.html>) This product extends the Audit Collection Service to provide advanced compliance and security monitoring. The tool provides support for CoBITs, FISMA, HIPAA, ISO, PCI, SOX, and others, so you'll want to check the vendor website for details.

Silect Software (<http://www.silect.com/>) Silect's MP Studio product helps you control the management pack life cycle, including management pack authoring, analysis, testing, and version control.

You'll also find a few companies that aren't Microsoft partners but that do provide interesting tools. The Ancoris Extension Framework for MOM 2005 and SCOM 2007 extends Operations Manager using technologies such as NetIQ's Extended Management Packs, Jalasoft Network Management, Tidal Software's SAP monitoring, and Skywire Software's connectors for frameworks and help desks. Using this product, you'll be able to manage a wider variety of operating systems, databases, and platforms using Operations Manager. You can find Ancoris at <http://www.ancoris.com/s/mon/ancmgmtfrm.shtml>.

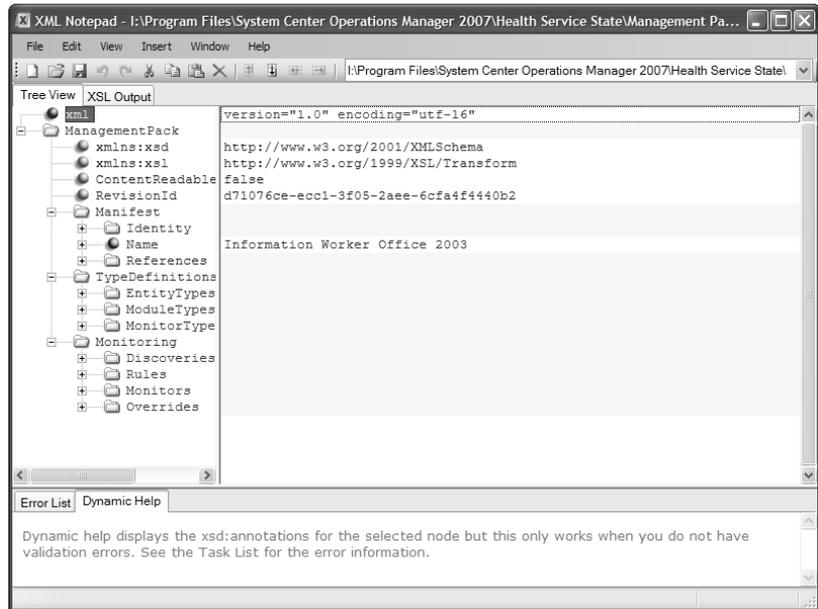
Some other companies are working pretty hard on products and they may actually have something by the time you read this chapter. Two of these companies are Engyro (<http://www.engyro.com/products/integration/>) and iWave integrator (<http://www.iwaveintegrator.com/>). Both companies are working on connectors today, but expect to find them working on other tools as well in the future.

Other Interesting Tools

You won't find many tools just yet for Operations Manager 2007. However, beside the tools that Microsoft is promoting, there are a few interesting tools that they aren't. For example, Microsoft Product Support's Reporting Tools (<http://www.microsoft.com/downloads/details.aspx?FamilyId=CEBF3C7C-7CA5-408F-88B7-F9C79B7306C0>) can provide you with a complete inventory of any system. Once you have this inventory, you can use it as a means of validating the Operations Manager 2007 content. By comparing the two results, you can find missing entries and determine whether Operations Manager is reporting the true state of the systems that you're trying to manage.

It may take a while to get used to the idea that Operations Manager 2007 uses many XML files because MOM 2005 uses mostly binary files. However, once you get used to the idea, you may become curious about the content of those XML files. Trying to read them in Notepad or even Word will prove frustrating, and an accidental change can prove disastrous. You could invest in an expensive third-party solution, but XML Notepad 2007 (<http://www.microsoft.com/downloads/details.aspx?familyid=72d6aa49-787d-4118-ba5f-4f30fe913628>) is free and it does a great job. For example, Figure 13.3 shows an example of the Microsoft.InformationWorker.Office.2003.{D71076CE-ECC1-3F05-2AEE-6CFA4F4440B2}.{7B96BA62-AE3C-FA0C-749B-27681DD96E4C}.xml file found in the \Program Files\System Center Operations Manager 2007\Health Service State\Management Packs folder.

FIGURE 13.3
View the actual content of unsealed management packs using XML Notepad 2007.



Troubleshooting Agents

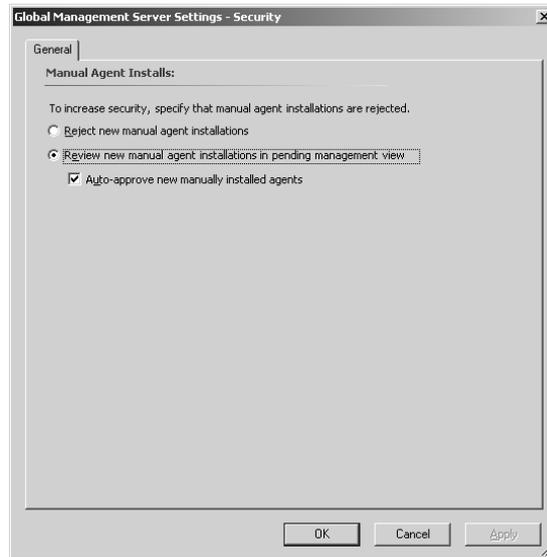
Agents provide the communication that Operations Manager requires between the client and the server. Needless to say, agents can become a major source of problems because just about everything interferes with their operation. You won't find too many system health issues that don't affect agents in some way. Everything from faulty networks to misconfigured firewalls can prevent proper communication. The following sections describe some of the more common issues that affect agents.

Making the Agent Available to the Server

You may run into a problem where the agent you manually install on a client doesn't appear to work. In fact, the server might not even seem to know that the client exists. In most cases, the cause of this problem is simple to fix. Microsoft assumes that everyone's going to use push technology to install agents. Consequently, the option to accept manually installed agents isn't even available. Use the following steps to resolve this problem:

1. Open the Administration\Settings folder in Operations Manager.
2. Highlight the Security entry in the Type Server (2) group and click Properties. You'll see the Global Management Server Settings – Security dialog box shown in Figure 13.4.
3. Select the Review New Manual Agent Installations In Pending Management View option.
4. Check the Auto-Approve New Manually Installed Agents option to allow these new agents to start immediately.

FIGURE 13.4
Set the policy
for manual agent
installations.



5. Click OK. The server is now open to working with agents. However, the server won't automatically process any of the agents that you manually installed previously.
6. Perform the manual installation tasks for your clients.



Real World Scenario

FINDING 'LOST' SYSTEMS

Joel had finally convinced his boss that they should automatically allow agents to be approved for two remote offices. The wide area network connection between the remote offices and the management servers in the corporate office was far too slow to try and push the agent to the clients. Approximately 50 clients and 4 servers were hosted at the remote locations, and one of the junior technicians had been chosen to install the agents on all of the systems.

After Joel had turned on the Auto-Approve New Manually Installed Agents option, he thought that everything was good. However, as he started checking the managed systems within the Operations Console, he noticed that he was missing two of the servers from one remote site. He approached the technician about the servers and was told that all of the systems had agents installed. After not finding the servers in the Agent Managed container, he decided to double-check the technician's work. He started a remote desktop session into one of the servers and found that the agent had indeed been installed.

Not sure why the two servers were not in the Agent Managed container, Joel posted a question within an Operations Manager forum, only to receive a quick response asking him if he had checked the Pending Management container. Sure enough, when he checked, both systems were in there, along with three workstations that Joel had overlooked. As it turned out, the technician had made it to the remote facility and installed five agents before Joel had a chance to turn on the automatic approval option.

Operations Manager Won't Remove an Agent Entry in the Pending Management Folder

Sometimes you'll see an agent entry in the `Administration\Device Management\Pending Management` folder. You know that you want to accept the agent, so you highlight it and click `Approve` in the Actions pane. Nothing happens immediately, but you feel that the system will eventually remove the entry. A little while later, you come back and the entry is still there. The entry also appears in the Agent Managed folder. Repairing the agent entry doesn't fix the problem. Normally, this entry signifies that Operations Manager experienced an error with the agent. In fact, you may see an event log entry that tells you that the client was trying to perform an illegal action, such as connecting to a management group that doesn't exist on the current system.

Before you can do anything else, you'll need to clean up the error on the client and monitor the event log. Use these steps to fix the problem when you see the client is no longer generating the errant entries in the Event Log:

1. Approve the entry again in the Pending Management folder.
2. Repair the agent in the Agent Managed folder.
3. Highlight the entry in the Pending Management folder and choose `Reject`. After a few seconds, choose `View > Refresh`. The errant entry should go away.
4. Verify that the agent entry still appears in the Agent Managed folder. Reinstall the agent, if necessary.

Operations Manager Can't Push Content or Manage the Client Remotely

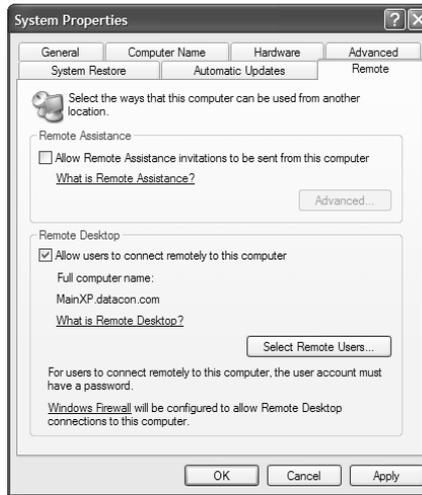
In some cases, an incorrect configuration can actually hide as a security benefit. For example, many websites recommend that users disallow Remote Desktop connectivity, and they provide precise instructions to turn this feature off. Unfortunately, you need this feature to provide proper connectivity for Operations Manager. If you suddenly find that some features are working (the agent still has a heartbeat) but other features aren't (you can't create a remote connection), this is a good place to look. The following instructions tell you how to turn the Remote Desktop feature back on. Start at the client machine.

NOTE The instructions in this section are for a Windows XP client system. The same principles apply to other versions of Windows, but you may need to modify the instructions slightly to make the instructions work in these other environments.

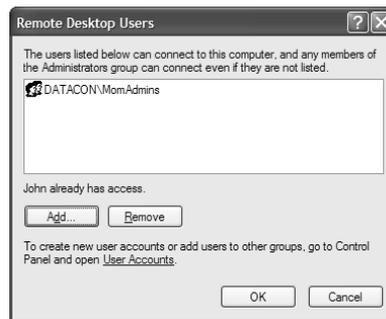
1. Right-click `My Computer` and choose `Properties` from the context menu. You'll see the System Properties dialog box.
2. Select the `Remote` tab, shown in Figure 13.5.
3. Verify that the `Remote Desktop` option is selected.
4. Click `Select Remote Users`. You'll see the `Remote Desktop Users` dialog box, shown in Figure 13.6.
5. Verify the administrators have proper Remote Desktop access to the system. If the administrators do have proper access, click `OK` twice to close the dialog boxes and exit this procedure.
6. Click `Add`. You'll see a `Select Users` dialog box.

FIGURE 13.5

Check the Remote Desktop setting and Remote Desktop Users security settings on the client system.

**FIGURE 13.6**

Verify that the administrators have proper access to the system.



7. Click Locations and provide any required security information. You'll see the Locations dialog box.
8. Choose the domain controller so you can access the domain accounts. Click OK to show the domain controller as the search location.
9. Click Advanced. You'll see an extended form of the Select Users Or Groups dialog box.
10. Click Object Types. You'll see the Object Types dialog box.
11. Check the Groups option and then click OK.
12. Click Find Now. Windows displays a list of domain accounts that you can use to access the client.
13. Choose the group or users that will require remote access to the client system. Click OK twice. The new users or groups will appear in the Remote Desktop Users dialog box.
14. Click OK twice to apply the changes.

Operations Manager Can't Install the Agent Remotely

You may encounter a problem where Operations Manager can't push the agent to the client. If you keep seeing an 80070643 error code, then the cause might be a disabled Automatic Updates service on the client. Microsoft is currently working on this issue. In the meantime, you need to set Automatic Updates to start automatically on the client system or you can perform the agent installation manually.

Agent Warnings and Errors in Event Viewer

More often than not, the agent will tell you about a problem it has encountered. The only problem is that many administrators aren't listening. The warnings and errors appear in the Event Viewer console located in the Administrative Tools folder of Control Panel. Open Event Viewer and search through the Operations Manager event log first. Most warnings and all errors in this folder are a cause for concern (every time you clear the event log Operations Manager outputs a host of warning messages that you can safely ignore).

After you check the Operations Manager event log, look for errors in the Application and System event logs as well. These logs can contain error messages that point to other potential sources of agent communication problems.

Of course, the natural assumption is that most of the information you need will appear on the server. However, the client can also provide an invaluable source of information. Again, begin your search with the Operations Manager event log on the client machine, and then move on to the Application and System event logs.

Unfortunately, some of the event log entries are difficult to read. In fact, some people feel they're downright impossible to read because they're coated in jargon and dunked in Microsoft-speak. For example, the event log entry in Figure 13.7 falls into the nearly-impossible-to-understand category.

Click the link at the bottom of the event log entry, and you'll see an Event Viewer dialog box like the one shown in Figure 13.8. When you click Yes, Windows will open your browser or Help and Support Center and take you directly to a web page with additional information. It would be great to say that this web page is always helpful or that it contains any information at all, but that's not the case. Sometimes you'll see a "We're Sorry" error message.

FIGURE 13.7

Event log entries can prove difficult or impossible to understand and often don't point to the correct error.

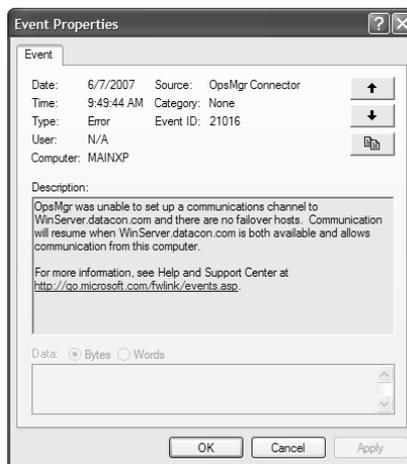
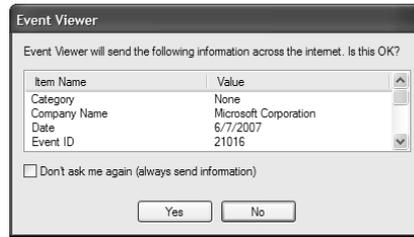


FIGURE 13.8
Sending the error information to Microsoft can help you get additional information.



In this case, you have another recourse that actually works more often than the Microsoft solution does. Copy the beginning of the message shown in Figure 13.7 (in this case, you'd copy "OpsMgr was unable to set up a communications channel") and open the Google Advanced Search web page (http://www.google.com/advanced_search?hl=en). Copy this information to the Find Results With The Exact Phrase Field and click Google Search. In this case, the search yielded four results, one of which cured the problem.

Troubleshooting Management Packs

Management packs normally don't experience problems of the sort that cause Operations Manager to fail completely or even cause problems with an agent unless they have significant configuration problems. However, management packs can contain errors that cause the output to contain erroneous information. For example, if you create a self-tuning threshold monitor and place it within a management page, the monitor can produce incorrect readings and therefore generate unwanted alerts. The result is more annoying than disabling, but it's problematic anyway. These configuration settings represent the problems that you'll find most often.

Sometimes a management pack configuration issue can cause problems that are more significant. The following sections describe some management pack issues that you'll want to review and fix immediately.

The Entity Health for One or More Management Servers Displays Critical

It's important to remember that some of your best sources of information for problem resolution aren't on the Microsoft website but on third-party sites. In this case, someone has found a management pack error that causes the servers to show critical. The problem is one where a threshold is incorrectly set so that the Health Service Handle Count Threshold or Health Service Private Bytes Threshold monitors show a critical level. It turns out that the threshold is set too low and you need to create an override to fix it. You can read more about this particular problem and its fix at <http://blogs.technet.com/cliveeastwood/archive/2007/06/06/entity-health-for-one-or-more-management-servers-displays-critical-in-the-operations-manager-console.aspx>.

A Management Pack Won't Work After Conversion

Many people will want to use their MOM 2005 management packs in Operations Manager 2007. However, after conversion, the following error appears onscreen:

```
Invalid Management Pack FileName.xml.:
➔XSD verification failed for management pack. [Line 1, Position 18]
```

This error always occurs when you fail to convert the management pack correctly. The management pack conversion process requires two steps. First, you must convert the AKM file to MOM 2005 XML format, and then you must convert the MOM 2005 XML file to an Operations Manager 2007 XML file. See the “Converting Legacy Management Packs” section in Chapter 5, “Managing Management Packs,” for additional information.

Unfortunately, correct conversion doesn’t always mean that the management pack will work. For example, you might convert the management pack and discover that it requires Notification Services to work correctly. The lack of MOM 2005 Notification Services support may mean that the management pack won’t work in Operations Manager 2007. Consequently, it’s important to know the specifics of the management pack and determine if there are dependencies that Operations Manager 2007 won’t support.

Operations Manager Won’t Let You Save a Recorded Web Monitor

Sometimes a problem is a matter of having input that Operations Manager can’t work with properly. For example, you may record a web monitor and find that you can’t save it. You see the following error message:

```
Array index out of bounds.  
↳Cannot fetch element with index=0 (Collection size=0)
```

The problem is that URLs within the recorded session may contain characters that Operations Manager can’t understand. In this case, the URL contains curly braces ({}). You can replace the curly braces with the correct escape sequences to fix the problem. Use %7b for the opening curly brace ({}) and %7d for the closing curly brace ({}). You can find a complete list of URL escape codes at <http://www.december.com/html/spec/esccodes.html>.

Management Pack Warnings and Errors in Event Viewer

As with agents, you’ll find a considerable number of management pack errors described in the event logs. However, unlike when working with agents, you won’t normally find the errors described on the client—they usually appear exclusively on the server unless a management pack is looking for a performance characteristic or registry entry on the client that doesn’t exist. See the “Agent Warnings and Errors in Event Viewer” section earlier in this chapter for details on using Event Viewer to help you locate problems in the event logs.

Troubleshooting Security Issues

Because Operations Manager affects every area of your network, you can encounter a number of security issues, some of which are hard to find. Most administrators know to check the user settings when an administrator can’t gain access to the Operations Manager console. In addition, the chapters in this book have spent a good deal of time telling you how to configure Operations Manager so you won’t encounter security errors. However, in the grand scheme of things, these errors are usually easy to locate. Sometimes an error is so obscure that finding it proves elusive. The following sections describe some of the issues you need to consider when troubleshooting security issues for Operations Manager.

Overcoming Issues with the W32Time Service

You may run into a problem where Kerberos is reporting seemingly unrelated security issues in your system. The problem might affect SQL Server in a way that doesn't seem to have anything to do with Operations Manager, yet you're finding that Operations Manager doesn't seem to work as it should. It appears to forget things, and some client operations never take place. The problem is normally intermittent. The distinguishing characteristic of this problem is that you also see a W32Time service warning in the system event log that says you haven't synchronized W32Time to an external time source. The issue, in this case, is one of timing. Kerberos is quite fussy about the time stamp provided on messages and a matter of even a few milliseconds can cause problems in a marginal setup.

You can fix this problem by synchronizing the W32Time service to an external source. The following steps describe the process to use the `tock.usno.navy.mil` time server. This time server isn't the only one available on the Internet; the site at <http://ntp.isc.org/bin/view/Servers/StratumOneTimeServers> contains a number of time servers you can use.

WARNING Working with the registry isn't for the faint of heart. This procedure assumes that you know how to work with the registry and have made a recent backup of your system. Always make changes to the registry with extreme care.

1. Open the Registry Editor by choosing Start ► Run, typing **RegEdit** in the Open field of the Run dialog box, and clicking OK.
2. Locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Parameters\Type value, shown in Figure 13.9. This entry tells the server where to obtain time synchronization information. The default entry tells the server to look for a time server on your network, but, in this case, you don't have one.
3. Change the Type value from NT5DS to NTP. This change lets the server look through the list found in the NtpServer value for a time server.
4. Locate the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config\AnnounceFlags value, shown in Figure 13.10. This entry lets the server announce itself as a reliable time source. The server uses this entry only when you set the Type value in step 3 to NTP.

FIGURE 13.9
Locate the entry that tells the server where to look for time information.

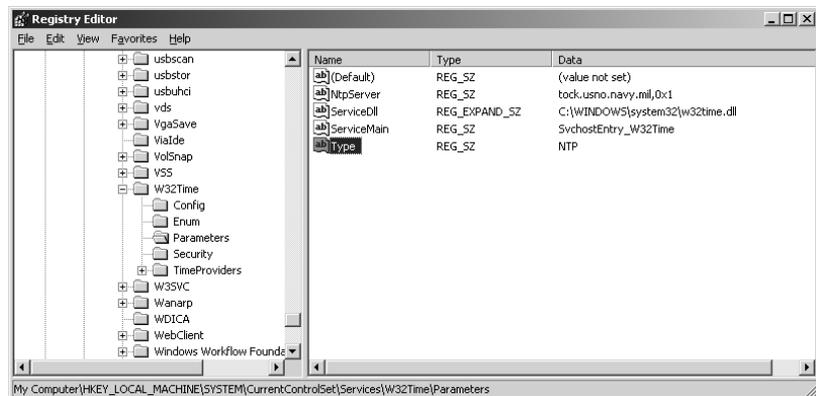
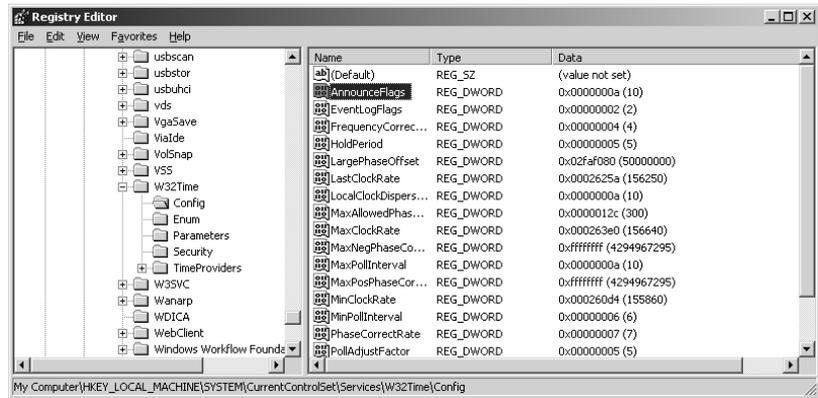


FIGURE 13.10

Locate the entry that tells the server to announce that it is a time server.



- Change the `AnnounceFlags` value from 5 to 10 (decimal). The server will now announce itself as a reliable time source.
- Locate the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Parameters\NtpServer` value, shown in Figure 13.9. This entry controls which time servers the server uses for synchronization. The default value is `time.windows.com`, which will probably work just fine. However, many administrators prefer non-Microsoft sources for reliability reasons, one of which is `tock.usno.navy.mil`. You can include as many time servers as desired.
- Open the `NtpServer` value. Type each time server value, followed by a comma and `0x1`. For example, to use the `tock.usno.navy.mil` server, type `tock.usno.navy.mil,0x1`. If you want to add a second server, press the spacebar and type its URL, following by the comma and `0x1`.
- At the command line, type `net stop w32time` and press Enter. When the `W32Time` service has stopped, type `net start w32time` and press Enter to restart the `W32Time` service.
- Type `w32tm /resync /rediscover` at the command line and press Enter. This action ensures that the server quickly synchronizes with the new time server.

The `W32Time` service supports a wealth of other configuration options that are out of the scope of this book. You can learn more about some of these options in the Knowledge Base article at <http://support.microsoft.com/kb/816042>.

Understanding Why the Discovery Wizard Never Completes

You can encounter errors that look like they are application errors in nature, but no one seems to know about an application problem that could cause the error. After many hours of searching, you don't find anything that can help you—certainly not a patch, registry edit, or other common form of fix. In at least some cases, what appears to be an application error is actually a symptom of a security error. This problem can occur when the Discovery Wizard won't complete. It keeps telling you that it's working on the discovery, but that's all it will do. When you cancel the discovery, you go to the task status and see that the task completed successfully. This looks like an application problem of the most grievous sort, but it's really a security error.

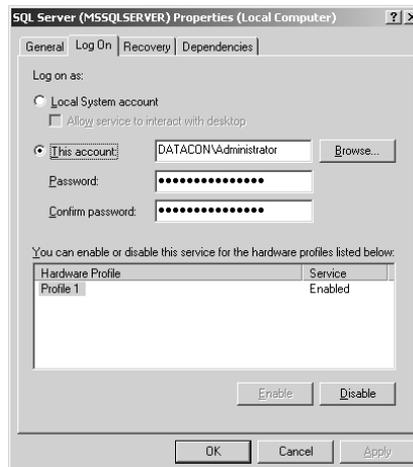
If you encounter this problem, look at the Application log in Event Viewer and see if you have an error with an Event ID value of 28005. The actual error message reads, "An exception occurred while

enqueueing a message in the target queue. Error: 15404, State: 19. Could not obtain information about Windows NT group/user 'MyDomain\AccountName, error code 0x5." The event log entry doesn't really tell you very much and it doesn't look like a security error either, but it is.

This problem can occur when the SQL Service Broker doesn't have proper access to Active Directory. It could have partial access, but it doesn't have the full access that Operations Manager requires it to have in order to complete the Discovery Wizard. In fact, this particular security problem can cause a number of Operations Manager features to act in a strange manner and not complete their tasks. You'll normally encounter this problem when you configure SQL Server to use something other than a domain account. For example, the Local System Account setting won't work. You must configure SQL Server to use a domain account that has sufficient access to ensure it can perform any Active Directory task. This account need not be the Administrator account (for demonstration purposes I did use the Administrator account, but using this account could open your server to security holes). The following steps tell how to fix this problem:

1. Open the Services console found in the Administrative Tools folder of Control Panel.
2. Locate the SQL Server service.
3. Right-click the service and choose Properties from the context menu. Select the Log On tab. You'll see the SQL Server Properties dialog box, shown in Figure 13.11.
4. Provide a domain account that has full Active Directory access privileges.
5. Click OK to complete the process.
6. Restart the service to ensure it logs on using the new account.
7. Perform steps 2 through 6 for the other SQL Server services in the Services console.

FIGURE 13.11
Choose a domain account with full Active Directory access to solve Discovery Wizard errors.



Security Failure Audits, Warnings, and Errors in Event Viewer

As with agents, you'll find a considerable number of management pack errors described in the event logs. Unlike with an agent, the first place you'll want to check for security errors is the Security log. Look for Failure Audit entries that tell you about an access that failed. Interestingly enough, sometimes Failure Audit entries also appear in the Application log. After you look for the Failure Audit entries, begin checking for warning and error message that point to a lack of access, such as a log entry

that tells you the client couldn't create a connection with the server (or vice versa). See the "Agent Warnings and Errors in Event Viewer" section earlier in this chapter for details on using Event Viewer to help you locate problems in the event logs.

The Bottom Line

Locate and use troubleshooting tools. Good tools can be hard to find unless you know where to look. For example, Microsoft often makes tools available for particular needs but then doesn't advertise them. Using a Google site search can help you locate precisely what you need in far less time than Microsoft's search engine. Go to the Google Advanced Search at http://www.google.com/advanced_search?hl=en, type **SCOM** as a search term, and type **download.microsoft.com** in the Domain field. The search results will tell you about every Operations Manager-related entry on the download site. After you try the download site, try msdn.microsoft.com, msdn2.microsoft.com, and www.microsoft.com. In a few searches, you'll have found everything that Microsoft has to offer. Locating what you need quickly and easily through a search is just one example of how you can find good tools for your Operations Manager toolkit. Make sure you check out newsgroups, newsletters, and other potential sources for tools as well. The next good tool you find could be lurking anywhere.

Troubleshoot agents. Agents provide the communication path from the server to the client. Without good communication, you can't check on the health of a client. Of course, the first thing that administrators tend to blame when an agent isn't working is the network the communication link must be damaged in some way. However, a good deal of other problems can cause the agent to stop working. A simple check box in the System Properties dialog box can cause problems. A new application can interfere with the agent. Security settings can also cause problems, even when the administrator doesn't remember making a change. Firewalls can cause communication problems. In fact, the number of issues that can cause a communication problem is nothing short of amazing, so it pays to look at the network for a while and then move on to other areas.

Troubleshoot management packs. Management packs help Operations Manager discover what to search for and, once discovered, how to interact with it. For example, when you install the Office 2003 management pack on your system, it discovers all of the systems that have Office 2003 installed on them. Now that it knows about those systems, the management pack helps Operations Manager monitor the status of those installations so that you can quickly determine which installations are healthy and which require service. Of course, each management pack has a specialty, so you need many of them to discover and monitor system health. When the management pack contains flaws, it no longer helps you discover or monitor anything; instead it can give you a false sense of security about the health of your system. When you see a critical condition in Operations Manager, you have to know that the condition is the result of a client error and not a flaw in the management pack.

Troubleshoot security issues. Security is the focus of many IT efforts today. Every week presents a list of new terrifying threats accompanied by an equal list of fixes that may or may not actually fix anything. Considering the amount of information you store in Operations Manager, it's essential that the security it provides be beyond reproach because a single break-in can expose your entire organization to unwarranted risk. That's why you want to examine every security issue closely and test it thoroughly before you implement a security solution. Make sure you pay attention to those seemingly minor event log entries because they can come back to haunt you. Check for updates, proper settings, good passwords, and the like as well. Of course, the biggest security hole is the human behind the screen—it pays to scrutinize the people you hire, too.

Chapter 14

Scripting Operations Manager 2007

In this chapter, we'll introduce you to the PowerShell scripting environment and explain how it relates to management of Operations Manager 2007. PowerShell is a new command-line interface and scripting environment introduced in conjunction with Windows Vista and Longhorn Server. You can think of it as an upgrade for the command window, which we can trace back to the good old days of DOS.

PowerShell is powerful. Hugely powerful. Mastering it takes more room than we have available in this book, so we're going to focus on the functionality of PowerShell as it pertains to Operations Manager 2007 and cover only enough basic PowerShell methodology to get you up and running. With that in mind, it's time to jump in and get your feet wet.

In this chapter, you will learn how to:

- ◆ Begin using the PowerShell environment
- ◆ Use PowerShell's built-in commands and utilities
- ◆ Pipe together individual components named cmdlets to produce powerful results
- ◆ Use additional cmdlets provided with Operations Manager
- ◆ Leverage PowerShell to create scripts for Operations Manager

We'll cover the bases of managing Microsoft Operations Manager through PowerShell. We'll start with the basics to bring you up to speed, introducing you to the environment and the syntax of PowerShell. You'll see how individual components named cmdlets, while performing simple functions themselves, can be piped together to produce powerful results. We'll walk through the basic cmdlets available with PowerShell, and finish the chapter by working through the additional cmdlets provided with Operations Manager that are designed specifically for its management.

Introducing PowerShell

We'll start our crash course in PowerShell in the standard PowerShell command shell. Begin by clicking Start > Programs > Windows PowerShell 1.0 > Windows PowerShell. Starting the command shell usually takes a few seconds, so don't be alarmed if nothing happens right away.

First, we need to get a little terminology straight. PowerShell does not have commands—it has *cmdlets* (pronounced *command-lets*). The cmdlets are really just small Microsoft .NET programs designed for a specific task, such as copying files or changing your current location in the file system (or the registry). PowerShell provides a very far-reaching environment for system management.

To provide this powerful management environment, PowerShell has a comprehensive set of cmdlets. You can obtain a list of the native PowerShell cmdlets by issuing one of the following commands (do not type the `>`; that is simply part of our command-line prompt):

```
> get-command
> get-command | select-object Name
```

Either of these will give you a long listing of cmdlets to browse through. One important thing to note is that these cmdlets do not, as a general rule, return simple strings of text like the DOS commands you are probably familiar with. Instead, these cmdlets usually return an object of some type. That object may represent anything from a file to a folder to (in our ultimate purpose for this chapter) an Operations Manager task, action, and so on.

In the second command, you probably noticed the familiar pipe symbol (`|`). The pipe symbol is used to connect the output of one program to the input of another, much like connecting two pieces of plumbing pipe. This ability to plumb cmdlets together is what puts the “Power” in PowerShell. Each cmdlet is small, quick, and focused on its particular job. While it may seem as if we have a lot more commands (sorry, cmdlets) to deal with, this approach is more practical in the long run.

We could spend all day, and about 14 more chapters, talking about writing PowerShell scripts, but our time is short and we have much ground to cover in regard to using PowerShell with Operations Manager. Therefore, we’ll finish this general PowerShell discussion by examining some tools that help you learn the basic PowerShell cmdlets and hit some of the highlights.

Let’s start by showing you how to help yourself. It wouldn’t make any sense for Microsoft to give you this awesome management environment without any documentation. It also doesn’t make much sense to make you leave that environment to read the documentation, so it’s available right from your command prompt. Enter the following command:

```
> get-help get-command
```

The Get-Help cmdlet will give you the basics of a specified cmdlet (in this case, the Get-Command cmdlet). You’ll get a synopsis, some basic information about the cmdlet syntax, a description, and some related links. Would you like some more information? OK; enter this command then:

```
> get-help get-command -detailed
```

Now you have more details about the parameters and some examples as well. Still not enough information for you? Then try this command:

```
> get-help get-command -full
```

That’s all the helpful information you’re going to get out of the command-line environment, and it should be plenty to get you going. When you find yourself lost about how to use a cmdlet, just use the Get-Help cmdlet to find your way.

Moving right along, let’s check out the format of PowerShell cmdlets. You’ll notice that cmdlets follow a certain grammatical syntax, specifically a *verb-noun* format. This format is designed to help you remember which cmdlets are used for a particular purpose. Think about our earlier examples; we were “getting help” or “getting commands.” If you want to stop a service, your best guess would probably be Stop-Service (and you’d be right!).

The verb-noun format is useful in remembering the cmdlet names, but it definitely doesn’t make them any shorter, does it? And longer names usually mean more typos, thus resulting in syntax

errors. To make life a little simpler on that front, many cmdlets also have aliases assigned to them. For example, to change your location in the file system, you would execute a command like this:

```
> set-location C:\
```

Rather than typing the 12 characters for the Set-Location cmdlet, wouldn't you rather just type two? Here's an example:

```
> cd c:\
```

Hey, that looks familiar, doesn't it? Since almost all the DOS commands you are already accustomed to using have a PowerShell cmdlet equivalent, it only makes sense to assign an alias to the cmdlet that matches the DOS command name.

Of course, there are also many new PowerShell cmdlets. Not all cmdlets have aliases assigned to them, but you can view the alias list with one of the two following commands:

```
> get-alias
> get-alias | select-object Name
```

Excellent. You've put your toes in the water, and now it's time to start getting those feet even more wet. In the next section, we're going to talk about some of the basic PowerShell cmdlets that you'll want to know about to get the most out of our upcoming Operations Manager scripts.



Real World Scenario

POWERSHELL

You are working in an IT shop responsible for the care and feeding of several servers, well over 100. To maintain the necessary availability of the servers, Microsoft Operations Manager has been installed and configured to create order out of the chaos.

Over time, not only have new servers come in to replace the old, but newly acquired divisions in the company have brought an entirely new fleet of servers into your realm. Managing both the existing and new systems has required a redesign in the Operations Manager landscape.

To implement the new design, you have two choices. The first is to manually configure each new server by hand; the second is to automate the process. Manually reconfiguring every system is not only time-consuming, but also introduces the possibility of inconsistent changes and outright errors. Neither of these are acceptable to the corporate leadership.

But Operations Manager 2007 has considerably eased the automation situation through the adoption of PowerShell. By providing a comprehensive set of cmdlets targeted at Operations Manager management, you are able to create a script that quickly and accurately implements the new design policies on each managed server under your care.

Useful PowerShell Cmdlets

Now that you've seen the cmdlets a few times, let's take a look at some of the ones we're going to be taking advantage of in this chapter. Again, we don't have room to hit them all, but the goal is to familiarize you with them enough to make sense of the upcoming Operations Manager scripts.

SET-EXECUTIONPOLICY

This is one of the first and most important cmdlets we'll work with. It sets the security regarding PowerShell scripts. The default installation of PowerShell does not allow the execution of PowerShell scripts. If you've ever dealt with a malicious script, you can understand the importance of maintaining security in your scripted management environment.

However, this is going to hinder us down the line. Using the Set-ExecutionPolicy cmdlet, we can relax the security standards a bit. There are four execution policies we can select, as shown in Table 14.1.

TABLE 14.1: PowerShell Execution Policies

POLICY	DESCRIPTION
-Restricted	Does not run scripts. This is the default policy.
-AllSigned	Requires that all scripts be signed by a trusted publisher, including local scripts.
-RemoteSigned	Requires that only remote scripts be signed.
-Unrestricted	Runs all scripts and configuration files.

As you can see, security depends on where the script is stored and whether or not the script has been signed with a code-signing certificate issued by a trusted certificate publisher. Specifically, the publisher must be trusted by the computer executing the script. For example, to set our script execution policy to require signed scripts when executing from a remote source, we would run the following command:

```
> set-executionpolicy -remotesigned
```

You must determine which security policy applies to your situation. At a minimum, we recommend using the RemoteSigned policy on production servers. In a testing environment without a Certificate Authority server installed, you will have to use the Unrestricted policy. Remember to sign those scripts before moving them up to your production servers, though!

FOREACH-OBJECT CMDLET

This cmdlet comes in handy for working objects, especially multiple objects. For example, think back to the Get-Command cmdlet. It returned a collection of cmdlet objects. Each cmdlet object has the ability to display its own help information, which is used by the Get-Help cmdlet.

Now let's say you want to view the help information for several cmdlets. You could type each Get-Help cmdlet by itself, but that's just not much fun. Instead, give this command a try (and sit back—it might take a minute to finish scrolling by):

```
> get-command | foreach-object {get-help $_.Name}
```

There are a few syntax elements in there that will take a little explaining, but first let's look at the big picture. First, a collection of cmdlet objects was returned by the Get-Command cmdlet. That collection was passed to the ForEach-Object cmdlet. The ForEach-Object cmdlet took the name of each cmdlet object, one at a time, and passed it to the Get-Help cmdlet. And the Get-Help cmdlet did what it does best: display the help text for each cmdlet.

That command wasn't too tough, but we can make it a little easier on the eyes and fingers. Think back to our alias list and you might remember seeing one assigned to the `ForEach-Object` cmdlet. That alias is called, simply enough, `foreach`. Substitute the alias and we get this command:

```
> get-command | foreach {get-help $_.Name}
```

It's not a big savings in terms of typing, but every little bit helps. We'll be employing the `ForEach-Object` cmdlet several times in the future, so using the alias now is a good habit to get into.

Now let's clear up those syntax elements. First, the curly braces (`{}`) are used to mark the beginning and ending of a script block. Remember that the `ForEach-Object` executes some script code with the collection of objects that it is passed. The script code is placed between those curly braces. In our previous example, we used only one statement. But we can also separate multiple statements by placing a semicolon between the individual statements, as shown here:

```
> get-command | foreach {write "=====" ; get-help $_.Name}
```

This will place a nice (albeit short) double line between each set of help text. You can probably guess that our sample commands are going to start getting longer and appearing on multiple lines, and that's OK.

Next, let's talk about the `$_`.Name that appears in the curly braces. To understand it, we need to break it into two pieces. The first piece is the `$_`, not including the period. This is a pointer used by the `ForEach-Object` cmdlet. As the `ForEach-Object` cmdlet cycles through each individual object, the `$_` symbol always refers to the current object. In other words, it points to the first object, then the second, and so on until the collection of objects has been cycled to the end.

The next piece is after the period, specifically the `Name`. This is the specific property of the method of the current object that we want to reference. In our previous example, it referred to the `Name` property, but it could just as easily refer to any other visible property of the object. As we get into the Operations Manager cmdlets, you'll see syntax such as `$_`.Id or `$_`.ManagementGroup. These will be the globally unique IDs and management groups, respectively, of different Operations Manager objects.

WHERE-OBJECT

One thing you'll quickly learn about the `Get-*` cmdlets is that they can return a lot of unwanted objects. Did you realize that the `Get-Command` cmdlet will return over 120 cmdlets? Most of the time, we don't need the full range of objects; we only want a subset that matches a given criterion. To that end, we can use the `Where-Object` cmdlet (whose alias is `Where`) to narrow down our results. Take a look at this example:

```
> get-command | where {$_.Name -like "Get-*"}
```

Much as we can with a SQL query statement, we can use the `Where-Object` cmdlet to specify exactly the objects that we want. Those objects can then be passed off to another cmdlet for further action or processing. This example narrows down the list of cmdlets to only those that begin with the verb `Get`.

Before moving on, we need to look at comparison operators in PowerShell. You might have noticed in the previous example that we didn't use the equals sign (`=`) when making the comparison. That is because the equals sign is specifically for making an assignment in PowerShell, like assigning a value to a variable. It doesn't work in comparisons.

PowerShell has a number of comparison operators to use, as shown in Table 14.2. Note that they all begin with a single hyphen. Don't forget the hyphen, because PowerShell won't like it!

TABLE 14.2: PowerShell Comparison Operators

OPERATOR	DESCRIPTION
-like	Looks for match using wildcards
-notlike	Looks for non-matches using wildcards
-match	Looks for matches using regular expressions
-notmatch	Looks for non-matches using regular expressions
-eq	Equal to (=)
-ne	Not equal to (<>)
-lt	Less than (<)
-gt	Greater than (>)
-lte	Less than or equal to (<=)
-gte	Greater than or equal to (>=)

These operators will also be useful in our scripts. They can be used to make comparisons between values and objects using an If-Then statement (more on that later).

SELECT-OBJECT

As you execute different cmdlets, you're eventually going to see extraneous information, data that we just don't need. Let's illustrate the point with an example that we know is going to spew out a lot of text. You'll recall this little piece of code:

```
> get-command | foreach {get-help $_.Name}
```

One thing you might not notice as all those cmdlets are scrolling by is that they are all displaying similar information, such as Name, Synopsis, Syntax, Detailed Description, Related Links, and Remarks. Let's say that we really only want to list the commands and a synopsis of the purpose of each. Adding one more cmdlet to our command will do the trick:

```
> get-command | foreach {get-help $_.Name}
| select Name, Synopsis
```

The Select-Object cmdlet (whose alias is Select) lets us extract only the information from the object that we want. It can take a comma-separated list of object property names (in this example, those would be the Name and Synopsis properties) to determine what to display. If you have a property name that has embedded spaces (such as Related Links), you can enclose the property name in single quotes.

OK, WE'RE CHEATING A LITTLE

When you see the syntax `Select Name, Synopsis`, we're cheating a bit. The proper syntax for the `Select-Object` cmdlet would be `Select -property Name, Synopsis`. The parameter `-property` is the first positional property of the cmdlet, so we're taking advantage of that fact to shorten the length of our commands. We'll take advantage of this in other cmdlets as well. Strictly speaking, it's not the correct syntax, and if we were using more than one parameter, we would specify them all to be sure of our results. For these simple examples, let's save the extra typing.

We can also use the `Select-Object` cmdlet to limit the number of returned objects. The `-first` parameter, followed by the number of objects we want, will return the number of objects we want from the top of the list. Likewise, we can use the `-last` parameter (also followed by a number) to return a number of objects from the bottom of the list. The following two commands show you each in action:

```
> get-command | select -first 1
> get-command | select -last 1
```

These two options become extremely useful when used with the `Sort-Object` cmdlet, which we'll talk about next. After sorting a list of objects by the appropriate criteria, you can come up with all kinds of useful information, such as which processes on your computer are using the greatest amount of memory and which are using the least amount.

SORT-OBJECT

As you wade through the output of various formats, you'll find it might or might not appear in some sorted order. Even when it is sorted, it might not be in the order you want. To that end, the `Sort-Object` cmdlet (whose alias is `Sort`) is here to help!

Going back to our list of aliases that we examined earlier, you might have noticed that they appear in a relatively random order. This makes it tough to find the alias you are looking for. Let's look at employing the `Sort-Object` cmdlet to fix that:

```
> get-alias | sort Name
> get-alias | sort Definition
```

There you go—two sorted lists of aliases. Now you can look them up by either the alias or the definition (which refers to the cmdlet). Sorting by the definition, now you can see that some cmdlets have two or more aliases assigned to them.

You are not limited to one sort criterion. The `Sort-Object` cmdlet will accept a comma-separated list of property names to sort by. Be sure to test this to make sure you are getting the sort results you are looking for.

By default, the `Sort-Object` cmdlet will sort a collection of objects in ascending order. You can reverse the order by using the `-descending` parameter, as shown here:

```
> get-process | sort VM -descending
```

The `Get-Process` cmdlet lists the running processes on the local computer and displays various bits of information about each process, including the amount of virtual memory used. With this command, you can now see which of those processes is chewing up RAM in your system.

Before we move on, let's jump back to the `Select-Object` cmdlet. You saw how it allowed us to limit our lists to a certain number of objects, either from the top of the list or the bottom of the list. Let's employ that here to find the top ten memory-consuming processes:

```
> get-process | sort VM -descending | select -first 10
```

FORMAT-LIST

As you can probably guess by the name, the `Format-List` cmdlet (whose alias is `FL`) is used to format the output of a command operation as a list of properties. Each retrieved property will appear on a new line. Take a look at the following example:

```
> get-command | foreach {get-help $_.Name}
| select Name, Synopsis | format-list
```

Rather than cram the synopsis into its own column, the `Format-List` cmdlet gives it an entire line (or more if necessary). If you don't want to include the `Select-Object` cmdlet, the `Format-List` cmdlet also lets you specify which properties you want to see, as shown here:

```
> get-command | foreach {get-help $_.Name}
| format-list -property Name, Synopsis
```

The `-property` parameter takes a comma-separated list of property names to display. Another useful parameter is `-groupBy`. This lets you group the output based on a shared property or value. Let's look at some code that will show this in action:

```
> get-command | select Name, Verb | format-list
> get-command | select Name, Verb
| format-list -groupby Verb
```

As you can see, we used the verb as the grouping criteria. All the cmdlets with the same verb were displayed in their own group, and each group had a short prefix that noted the shared value. This can be handy in differentiating Operations Manager objects.

When using the `Format-List` cmdlet, keep in mind that there are usually only 80 characters available on a standard console screen. This can be stretched out to upwards of 120 characters, but you're still bound to run into some size and formatting constraints, so experiment to see what you're going to get.

GET-DATE

We know what you're thinking: how hard can it be to get a date (the calendar type)? It's really not, but PowerShell does it a little differently. Remember that PowerShell returns objects, and this cmdlet returns a `DateTime` object. Let's just execute the cmdlet and see what happens:

```
> get-date
```

You'll get a fully formed date, including day of the week, full month name, and time right down to the second. Nice, but maybe not the format you are looking for. Maybe you just want the date and not the time, or the time and not the date. Some cmdlets also require a starting and ending date, for which you will use a `DateTime` object provided by the `Get-Date` cmdlet.

Date manipulation is a very important subject for Operations Manager administrators. We deal with dates all the time, such as reviewing the last month's worth of event logs, setting the date for a new maintenance window, reviewing alert histories, and so on. It would be in your best interest to understand how to take advantage of the Get-Date cmdlet.

First, let's look at methods to change the format of the current date and time. The easiest way is to use the `-format` parameter, which will accept .NET-format specifiers. You can look these up on MSDN, or check out the short list of useful specifiers in Table 14.3 (note that the format specifiers are case-sensitive).

TABLE 14.3: Get-Date Format Specifiers

SPECIFIER	DESCRIPTION	RESULT
d	Short date	12/1/2006
D	Long date	Friday, December 01, 2006
t	Short time	6:30 PM
T	Long time	6:30:45 PM
f	Full date-time	Friday, December 01, 2006 6:30:45 PM
g	General date-time	12/1/2006 6:30 PM
G	General date-time	12/1/2006 6:30:45 PM
u	Universal	2006-12-01 18:30:45Z
s	Sortable	2006-12-01T18:30:45
y, Y	Month-Year	December, 2006

This isn't the entire list, but I think these are the most useful. Among them you might see some results you aren't used to seeing. The Universal and Sortable formats are useful if you plan to do some sorting by dates later in a PowerShell script. The Month-Year format is handy when you are appending results to a monthly text file and would like some kind of label prefixing the data.

To see these in action, try out some of the following examples:

```
> get-date -format d
> get-date -format D
> get-date -format f
> get-date -format s
```

You can also use the Get-Date cmdlet for creating a `DateTime` object. This object can be used with other cmdlets when a parameter calls it. These dates might be in the past (such as in reviewing an alert history) or in the future (if you are setting up a maintenance window). So let's explore the process of creating a new `DateTime` object.

First, we can manually specify the date we want. The `Get-Date` cmdlet has six parameters we can use to specify the new date, right down to the second. Those parameters are `-year`, `-month`, `-day`, `-hour` (this is in 24-hour format, so 3 p.m. would be 15), `-minute`, and `-second`. Putting those parameters to work, here are some commands you might use:

```
> get-date -year 2007 -month 7 -day 4
> get-date -hour 15 -minute 0
```

Alternatively, we can let `Get-Date` do the work for us. Remember, the `Get-Date` cmdlet returns a `DateTime` object. This object has some built-in methods for date manipulation that are extremely useful. Let's look at a few examples to start down that path:

```
> (get-date).AddDays(30)
> (get-date).AddDays(-30)
> (get-date).AddDays(30).Date
> (get-date).AddHours(-6)
> (get-date).AddMinutes(-90)
```

The first example will return the date that is 30 days in the future; the second will return the date 30 days in the past. The third also returns the future date, but it strips out the time value (leaving it at midnight). Also available are methods to add years (`AddYears`), months (`AddMonths`), hours (`AddHours`), minutes (`AddMinutes`), and even second (`AddSeconds`).

There is only one trick here: enclosing the `Get-Date` cmdlet in parentheses. Simply put, this allows the `Get-Date` function to run and complete, giving us that `DateTime` object we need. We can then access the object's properties and methods. In the previous examples, we used the `AddDays` method (be sure to prefix the property or method name with a dot—this is the proper syntax when it comes to objects).

ACCESSING THE CLASS DIRECTLY

This is a little more advanced, but I have no doubt you'll use it soon enough. The method of calling the `Get-Date` cmdlet is fine if you only have one date-manipulation operation you want to perform. However, it becomes inefficient if you have several operations you want to perform. Instead, you can access the class directly and use the static methods of the class (such as `AddHours`) directly.

To reference a class in the .NET Framework, you enclose the class name in square brackets (such as `[DateTime]`). Next, you use a pair of colons to access the static methods of the class. Check out this example:

```
> [DateTime]::Now.AddDays(30).Date
```

This does the same thing as the previous examples, except it accesses the class directly without running the `Get-Date` cmdlet. The `Now` function returns the current date and time, to which we can add our 30 days.

There are other useful date- and time-manipulation methods available with the `DateTime` object. Not only can you manipulate the date and time, but you can also extract portions of the date and time without messy string parsing. Take a look at Table 14.4 for a listing of the more useful properties and methods of the `DateTime` object.

TABLE 14.4: Useful *DateTime* Properties and Methods

METHOD	DESCRIPTION
AddMinutes	Adds a specified number of minutes to the current value of the <code>DateTime</code> object
AddHours	Adds a specified number of hours to the current value of the <code>DateTime</code> object
AddDays	Adds a specified number of days to the current value of the <code>DateTime</code> object
AddYears	Adds a specified number of years to the current value of the <code>DateTime</code> object
Hour	Returns the current hour of the <code>DateTime</code> object
Minute	Returns the current minute of the <code>DateTime</code> object
Day	Returns the current day of the <code>DateTime</code> object
Month	Returns the current month of the <code>DateTime</code> object
Year	Returns the current year of the <code>DateTime</code> object

You can make the use of these properties and methods as simple or as complex as you want. As you'll see from the following examples, you can come up with some interesting results by not only using the properties and methods by themselves, but also by combining them into longer commands:

```
> (get-date).AddMinutes(45)
> (get-date).AddHours(2)
> (get-date).AddHours(2).AddMinutes(45)
> (get-date).Hour
> (get-date).Day
> (get-date).Month
> (get-date).Year
> (get-date).AddHours(2).Hour
> (get-date).AddDays(45).Month
```

As you can see, there are several usage scenarios for the `DateTime` object. Once you have gotten the hang of using the properties and methods of the `DateTime` class, you are ready to start applying the same technique to other object classes in the .NET Framework.

OUT-PRINTER

I don't know about you, but I still like paper. It's like having a signed confession in your hand about what the system is doing (besides, we have too many things to remember already). PowerShell gives us a convenient cmdlet, named `Out-Printer`, to redirect output to a printer. Let's look at a quick example to get started:

```
> get-process | sort VM | select -first 10
   | out-printer
```

This example will send a list of the top ten virtual memory processes to the default printer. We can send the output to a different printer with a slight modification to the command:

```
> get-process | sort VM | select -first 10
  | out-printer -name Laserjet5
```

This command redirects the output to a printer named Laserjet5. If the printer has a name that contains spaces, you will have to enclose the name in quotes. Personally, we think it's easier to use short printer names.

OUT-FILE

As you are monitoring Operations Manager through PowerShell, at some point you might decide that it's a good idea to save some of your monitoring results. The PowerShell environment does support redirection of output using the `>` and `>>` symbols (the first creates a new file; the second appends to an existing file or creates it if the file doesn't exist). In case you haven't seen these in action before, here are a couple of command examples:

```
> get-process | sort VM | select -first 10 > C:\Top10VM.txt
> get-process | sort VM | select -first 10 `
  >> C:\Top10VM.txt
```

This first example gets the top ten processes consuming virtual memory and sends the results to a file named `C:\Top10VM.txt`. The problem is that if you run this command a second time, the first results will be lost. The single `>` symbol will overwrite an existing file. To keep a running log file (of sorts), we need to use the `>>` symbol, as shown in the second command. The double `>>` symbol appends the new command output to the end of the existing file. File sizes will obviously grow over time with this technique, so you will want to keep an eye on it.

There is also a PowerShell cmdlet to redirect output to a file, and it can give us a few more options. The cmdlet name is `Out-File` (it has no alias assigned to it by default). Let's put our earlier examples into a command using the `Out-File` cmdlet:

```
> get-process | sort VM | select -first 10
  | out-file C:\Top10VM.txt
> get-process | sort VM | select -first 10
  | out-file C:\Top10VM.txt -append
> get-process | sort VM | select -first 10
  | out-file C:\Top10VM.txt -noclobber
```

So here we have three examples of the `Out-File` cmdlet in action. The first duplicates the action of the `>` redirection symbol (remember, this will overwrite an existing file). The second uses the parameter `-append` to add the new output to the end of the existing file. The third example uses the parameter `-noclobber`. This parameter is like a safety net; it will not let you overwrite an existing file. Instead, you will get an error message telling you that the file already exists. Pay particular attention to this one—it can save you from wiping out a file at the command line. On the downside, using it in a script can cause you debugging problems.

Now, with your introductory lesson to PowerShell out of the way, it's time to start digging into PowerShell as it relates to Operations Manager.

PowerShell and Operations Manager

Now that we've covered the basics of the PowerShell world, it's time to see how it applies to Operations Manager. To start, we aren't going to be in the standard PowerShell environment, so if you still have your PowerShell window open, go ahead and close it.

To use PowerShell with Operations Manager, we need to tweak the environment a little. Microsoft has created the tweaks for us, and made them available in a menu shortcut named Command Shell under the Operations Manager 2007 program menu. Go ahead and start the command shell, and you'll see the first hint of the customizations that are taking place.

The customizations take place through a custom startup script and console file assigned to the command shell shortcut. These files perform some useful functions, chief among which is loading a PowerShell snap-in that makes the Operations Manager cmdlets available to us. Without this snap-in, none of the cmdlets in the rest of this section would be accessible. They also automatically connect our PowerShell session to the default management server assigned to the computer and modify the prompt to make the command line a bit less cluttered.

After starting the command shell, you will see a line that tells you which management server you are currently connected to (in our case, it is a server named SCOM.zygort.lcl). You can add a connection to another server later if the need arises, but for this discussion we will stay right where we are. The prompt is set up to remind us where we are in the management server at all times, so let's look at an example:

```
PS Monitoring:\SCOM.zygort.lcl
```

Here's the quick rundown: the PS stands for PowerShell, to remind that we aren't in DOS anymore, Toto. `Monitoring` tells us that we are using the monitoring provider (a.k.a. Operations Manager). The rest tells us our current location, specifically the root of SCOM server.

This would be a good time to increase the size of the command shell window. If you've forgotten how to do this in a DOS-esque window, just click the icon in the upper-left corner and select Properties. Select the Layout tab and increase the width and height properties in the Window Size frame. We prefer at least 120 characters wide and 40 characters deep, if possible.

Now, here's what you have to get used to, and it might take awhile. You have to realize that this is just like navigating a DOS environment. We can use commands like `cd` (Change Directory) to move about and `dir` (Directory) to list the contents of a "folder." You have to remember that these aren't really folders at all but rather hierarchical groupings of Operations Manager objects. At the top level is the server, which is where we are now. Let's execute the following command and see what we get:

```
> dir
```

Yikes—that wasn't quite what you expected, was it? All that text that just scrolled across your screen represented a list of the Operations Manager objects available at the top level of our current management server. Let's try that command again, but this time we'll give it a little PowerShell-addition to try to clean that list up a bit:

```
> dir | Select PathName, DisplayName
```

Well, it's more readable (because we eliminated much of the extraneous information), but it's still a little confusing. Some of the objects you see are system-specific objects that you won't ever

have to deal with. Others relate to specific Operations Manager management groups. In fact, you should see several groups listed under the `DisplayName` field.

In a small test installation, many of these groups will be empty. We can test this out with the `Dir` command on some of the pathnames you can now read. Go ahead and try some to see what you can come up with.

PowerShell uses the `Tab` key for character completion. Start typing the beginning portion of a pathname, press the `Tab` key, and the first match will be returned. You can keep pressing the `Tab` key until the correct pathname shows up.

```
> dir Microsoft.Exchange
> dir Microsoft.Windows.Client.ComputerGroup
> dir Microsoft.Windows.Server.ComputerGroup
> dir Microsoft.SystemCenter.AgentManagedComputerGroup
```

Hopefully you found some objects underneath a couple of those group listings. In fact, you might have found some that look like duplicates, such as these on our system (we've shortened the output to prevent boredom):

```
> dir Microsoft.Windows.Server.ComputerGroup
Id                : 64353e34-ba2f-2f94-7989-fd23b56cc986
PathName          : SCOMAD.zygort.lcl

> dir Microsoft.SystemCenter.AgentManagedComputerGroup
Id                : 64353e34-ba2f-2f94-7989-fd23b56cc986
PathName          : SCOMAD.zygort.lcl
```

How can that be, you ask? Simple; they are the same object (a monitoring object, to be exact, something we are keenly interested in managing) that belongs to two different groups (at least). You've seen this all the time with Active Directory users belonging to multiple security and distribution groups—we just happen to be looking at it through the lens of a folder structure.

Now the question becomes, "Why would you represent monitoring objects this way?" To answer your question, consider two of the groups you've seen listed:

```
Microsoft.SystemCenter.AgentManagedComputerGroup
Microsoft.SystemCenter.AgentLessManagedComputerGroup
```

From earlier chapters, you can deduce that systems with an Operations Manager agent would be placed in the first group and systems that are remotely managed would be placed in the second group. Now let's say we want to apply some type of change to the agent-managed systems, such as changing to a new management server. This change would not apply at all to remotely managed systems because they use a proxy agent.

So, we would go to the "folder" that represents the group of computers we want to change. Just for practice, let's move to that folder and list its contents:

```
> cd Microsoft.SystemCenter.AgentManagedComputerGroup
> Dir | Select PathName, DisplayName
```

You should see a list of objects contained in that group. Earlier we mentioned that these were monitoring objects, but let's make sure first. We'll just check out the first object in the group:

```
> Dir | Select -first 1 | Get-Member
```

Wow, more information scrolling across the screen! Now scroll back up to the top and you'll see a tag labeled `TypeName`, which should be set to `Microsoft.EnterpriseManagement.Monitoring.MonitoringObject`. All that text simply means that we are looking at an object of the class `MonitoringObject` that is defined in the .NET Framework namespace `Microsoft.EnterpriseManagement.Monitoring`. All we care about right now is that we have a monitoring object in our sights.

The rest of that text that scrolled by will give you an idea of what that monitoring object is capable of. We are interested in the `Property` and `Method` values listed under the `MemberType` property. Properties tell us what kind of information we can get and possibly set for that object. Methods describe a function that the object can perform for us. For example, in the list you'll see a method called `ExecuteMonitoringTask`. With this, the monitoring object can execute an associated monitoring task as necessary.

Now that we know we're dealing with monitoring objects, we can plan how we want to deal with the task at hand. In our example, this task was setting the management server for the agents. To accomplish this, we would use the Operations Manager cmdlet `Set-ManagementServer` (which is described in more detail later in this chapter). The cmdlet needs to know which monitoring objects to change the management server for, but it can only change one monitoring object at a time. Guess what—the `Dir` command passes one object at a time through a pipe, which makes it perfect for funneling the monitoring objects to the cmdlet. When all is said and done, the command would look something like this:

```
> dir | Set-ManagementServer `
    -ManagementServer (Get-ManagementServer -Root)
```

At this point, our commands are getting longer than a single line (at least on a book page). In order to continue a PowerShell command on multiple lines, we have to use a line-continuation character. To do this, we use a grave symbol (```) at the end of the line to be continued. This character is typically located below the tilde (`~`) on the keyboard. You will see this symbol used frequently throughout the remainder of this chapter for readability.

Before moving on to the Operations Manager cmdlets, you should take some time explore the “file system” as provided by the monitoring provider. Find the location of different objects and relate them to what you are used to seeing in the Operations Console. As you find objects, use the `Get-Member` cmdlet to discover what class the object belongs to (this is listed by the `TypeName` tag) and the functionality of that class.

Operations Manager Cmdlets

As we mentioned earlier, PowerShell does not natively support Operations Manager functions. We have to add those functions in through something called a *snap-in*.

When you installed Operations Manager 2007 and specified the PowerShell functionality, a shortcut was created in the Operations Manager menu called `Command Shell` that starts the PowerShell environment and automatically invokes the Operations Manager snap-in. Without this snap-in, you're stuck with the standard set of PowerShell cmdlets. You might want to copy this shortcut to the Desktop or even the Quick Launch bar to make it easier to run.

When starting the Operations Manager PowerShell environment, a connection is made to a management server. This is important because you need to know which management server you are currently connected to. As we execute cmdlets, their operations are going to be based on the setup of that management server. For example, if we use PowerShell to install an agent on a system to be monitored, then that agent is going to be assigned to the current management server.

Depending on the situation, you may want to change the target Management Server. Most cmdlets provide a parameter that allows you to specify a different target management server if this is the case. This is probably the most reliable method, as you will be specifying the target in each case. Avoid assumptions whenever possible—be specific.

With the Operations Manager snap-in, we have access to all the additional Operations Manager cmdlets. Before we dive into scripting with them, though, let's take some time to examine the bounty of cmdlets we've just been given access to. In the following section, we are going to take a brief look at all the available cmdlets. I'll describe their purpose, what type of objects they return (if any), and what kind of information they need to do their job. Finally I'll give you a few snippets of code to get you started.

As you've seen with our brief introduction to PowerShell, cmdlets use the form verb-noun when describing their purpose. This can make locating the correct cmdlets a little tough sometimes. We're going to categorize the cmdlets to make referencing them a little easier on you.

Computer and Device Cmdlets

The cmdlets that follow will deal with the systems targeted for management by Operations Manager. We'll start with a set of cmdlets for installing, removing, and working with the Operations Manager agent software. Agent objects are going to play a big role in this section, so be sure you understand what agents are and how they function.

GET-AGENT

This cmdlet is used to retrieve the agent objects that are associated with a management server. By default, the current management server will be referenced or you can specify a different management server when executing the cmdlet.

Using this cmdlet, we can get references in our script to specific agent objects. That reference allows us to invoke functionality specific to the agent, such as starting a task or reading a current property value.

Parameters of this cmdlet include:

- ManagementServer** Specifies the management server from which to retrieve agents. This object can be retrieved using the `Get-ManagementServer` cmdlet.
- Path** Specifies the path to the desired management groups. This parameter is used only if the management server is not specified with the `-ManagementServer` parameter.

Here are some examples:

```
> Get-Agent
> Get-Agent | Where {$_.ComputerName -eq "SCOMAD"}
> $Server = Get-ManagementServer -Root
> Get-Agent -ManagementServer $Server
```

INSTALL-AGENT

This cmdlet is used to install the Operations Manager agent on a specified computer system. By default, the current management server will host the agent operations.

Before we can install an agent, we have to obtain a reference to a target computer (or computers). To do this, we need to execute the `Start-Discovery` cmdlet with an appropriate discovery object.

Parameters for this cmdlet include:

- ManagementServer** Specifies the management server to be used by the agent. This object can be retrieved by using the `Get-ManagementServer` cmdlet.
- AgentManagedComputer** Specifies the computer(s) to install the Operations Manager agent on. This should be a set of monitoring objects retrieved from the `Start-Discovery` cmdlet.
- InstallAgentConfiguration** Specifies agent configuration information. This type of object can be created only through the `New-Object` cmdlet and using the .NET Framework classes.
- FailoverServers** Specifies a list of servers to be used by the agent as failover servers. Multiple servers should be separated by commas.
- Confirm** Prompts for confirmation before uninstalling the agent.

Here are some examples:

```
> $Results = Start-Discovery `
    -ManagementServer (get-ManagementServer) `
    -WindowsDiscoveryConfiguration $DiscoverConfig
> Install-Agent -Managementserver (Get-ManagementServer) `
    -AgentManagedComputer $Results.CustomMonitoringObjects
```

UNINSTALL-AGENT

This cmdlet is used to remove the Operations Manager agent from managed systems. We begin by creating a list of agent objects we wish to remove, and then we pass that list to the `Uninstall-Agent` cmdlet.

Parameters for this cmdlet include:

- AgentManagedComputer** Specifies the computer(s) to remove the Operations Manager agent from. Unlike the `Install-Agent` cmdlet, this parameter requires agent objects (which we can get with the cmdlet `Get-Agent`).
- AgentConfiguration** Specifies agent configuration information. This type of object can be created only through the `New-Object` cmdlet and using the .NET Framework classes.
- Confirm** Prompts for confirmation before uninstalling the agent.

Here are some examples:

```
> $AgentList = Get-Agent `
    | Where {$_ .Name -eq 'srv01.zygort.com'}
> Uninstall-Agent `
    -AgentManagedComputer $AgentList `
    -Confirm
```

ADD-REMOVEDMANAGEDCOMPUTER

This cmdlet is used to add a remotely managed (agentless) computer to an Operations Manager agent. To use this cmdlet, you must create a list of computers to manage using the `Start-Discovery` cmdlet, similar to preparing to use the `Install-Agent` cmdlet. You also need to create a reference to the agent object you wish to use as the proxy agent for the remote computer(s).

Parameters for this cmdlet include:

-Computer Specifies the remote computer(s) to manage. This parameter expects a monitoring object, which we can retrieve through the Start-Discovery cmdlet.

-ProxyAgent Specifies a proxy agent object to use. This object can be retrieved by using the Get-Agent cmdlet.

Here are some examples:

```
> $LDAPQuery = New-LdapQueryDiscoveryCriteria `
    -Domain zygort
    -Ldapquery "(cn=RemoteServer01)"
> $DiscoverConfig = New-WindowsDiscoveryConfiguration `
    -ComputerType Server
    -LdapQueryDiscoveryCriteria $LDAPQuery
> $Results = Start-Discovery `
    -ManagementServer (get-ManagementServer) `
    -WindowsDiscoveryConfiguration $DiscoverConfig
> $ProxyAgent = Get-Agent `
    | Where {$_.Name -eq "ProxyAgent"}
> Add-RemotelyManagedComputer `
    -Computer $Results.CustomMonitoringObjects `
    -ProxyAgent $ProxyAgent
```

ADD-REMOTELYMANAGEDDEVICE

This cmdlet is used to add a remotely managed (agentless) device to an Operations Manager agent. To use this cmdlet, you must create a list of devices to manage using the Start-Discovery cmdlet. Rather than using the New-WindowsDiscoveryConfiguration cmdlet, though, you will use the New-DeviceDiscoveryConfiguration cmdlet. You also need to create a reference to the agent object you wish to use as the proxy agent for the remote computer(s).

Parameters for this cmdlet include:

-Device Specifies the device to manage. This parameter expects a monitoring object, which we can retrieve through the Start-Discovery cmdlet.

-ProxyAgent Specifies a proxy agent object to use. This object can be retrieved by using the Get-Agent cmdlet.

Here are some examples:

```
> $MonitoringClass = Get-MonitoringClass `
    -Name "System.NetworkDevice"
> $DiscoveryConfig = New-DeviceDiscoveryConfiguration `
    -MonitoringClass $MonitoringClass `
    -FromIPAddress 192.168.1.200
    -ToIPAddress 192.168.1.254
> $Results = Start-Discovery `
    -ManagementServer (Get-ManagementServer) `
    -DeviceDiscoveryConfiguration $DiscoveryConfig
```

```
> $ProxyAgent = Get-Agent `
    | Where {$_Name -eq "ProxyAgent"}
> Add-RemotelyManagedDevice `
    -Device $Results.CustomMonitoringObjects `
    -ProxyAgent $ProxyAgent
```

GET-REMOTELYMANAGEDCOMPUTER

This cmdlet is used to retrieve the specified agentless managed computers. A monitoring object is returned, which can be operated on by other cmdlets.

The only parameter for this cmdlet is `-Path`, which specifies the management path (or paths, separated by commas) to the remotely managed computer. The `Where-Object` cmdlet can also be used to specify the remotely managed computer.

Here's an example:

```
> Get-RemotelyManagedComputer | `
    Where {$_Name -eq 'RemoteServer01'}
```

GET-REMOTELYMANAGEDDEVICE

This cmdlet is used to retrieve the specified agentless managed devices. A monitoring object is returned, which can be operated on by other cmdlets.

The only parameter for this cmdlet is `-Path`, which specifies the management path (or paths, separated by commas) to the remotely managed device. The `Where-Object` cmdlet can also be used to specify the remotely managed computer.

Here's an example:

```
> Get-RemotelyManagedDevice `
    | Where {$_Name -eq "RemoteSwitch01"}
```

REMOVE-REMOTELYMANAGEDCOMPUTER

This cmdlet removes a remotely managed (agentless) computer from an Operations Manager proxy agent. Parameters for this cmdlet include:

-Computer Specifies the computer to remove from the proxy agent. This object can be retrieved with the `Get-RemotelyManagedComputer` cmdlet.

-Confirm Prompts to confirm removal of the remote computer.

Here are some examples:

```
> $Remote = Get-RemotelyManagedComputer `
    | Where {$_Name -eq "RemoteServer01"}
> Remove-RemotelyManagedComputer `
    -Computer $Remote `
    -Confirm
```

REMOVE-REMOTELYMANAGEDDEVICE

This cmdlet removes a remotely managed (agentless) device from an Operations Manager proxy agent. Parameters for this cmdlet include:

- Device** Specifies the device to remove from the proxy agent. This object can be retrieved with the `Get-RemotelyManagedDevice` cmdlet.
- Confirm** Prompts to confirm removal of the remote device.

Here are some examples:

```
> $Remote = Get-RemotelyManagedDevice
    | Where {$_.Name -eq 'RemoteSwitch01'}
> Remove-RemotelyManagedDevice `
    -Device $Remote `
    -Confirm
```

SET-PROXYAGENT

This cmdlet is used to set the proxy agent for a set of remotely managed computers or devices. To use this cmdlet, you must create a list of remotely managed computers and/or devices to reset the proxy agent on. The ability to create a large list makes it easy to decommission one proxy agent and move its clients over to a new agent. You also need to create a reference to the desired agent object you wish to use as the proxy agent for the remote computers and/or devices.

Parameters for this cmdlet include:

- Computer** Specifies the remote computer(s) on which to change the proxy agent setting. This object can be retrieved with the `Get-RemotelyManagedComputer` cmdlet.
- Device** Specifies the remote device(s) on which to change the proxy agent setting. This object can be retrieved with the `Get-RemotelyManagedDevice` cmdlet.
- ProxyAgent** Specifies the new proxy agent for the remote computer(s) or device(s). This object can be retrieved with the `Get-Agent` cmdlet.
- Confirm** Prompts for confirmation before assigning the proxy agent.

Here are some examples:

```
> $Remote = Get-RemotelyManagedComputer `
    | Where {$_.Name -eq 'RemoteServer01'}
> $Proxy = Get-Agent `
    | Where {$_.Name -eq 'ProxyAgent02'}
> Set-ProxyAgent `
    -Computer $Remote
    -ProxyAgent $Proxy
```

GET-MANAGEMENTSERVER

This cmdlet is used to retrieve instances of the management servers for all currently connected management groups. This is useful for cmdlets that have a management server parameter as it can be passed on the command line without the need to create a separate variable.

Parameters of this cmdlet include:

- Root** Specifying this parameter causes the cmdlet to return only the root management server.
- Path** Specifies the path (or paths, separated by commas) to the management servers to retrieve.

Here are some examples:

```
> $MgmtSvr = Get-ManagementServer `
    | Where (_.PrincipalName -eq 'OpsMgr')
> $Results = Start-Discovery `
    -ManagementServer $MgmtServer `
    -WindowsDiscoveryConfiguration $DiscoverConfig
```

SET-MANAGEMENTSERVER

This cmdlet is used to set the default Management Server for an Operations Manager agent. Parameters for this cmdlet include:

- ManagementServer** Specifies the management server to be used by the agent. This object can be retrieved with the Get-ManagementServer cmdlet.
- AgentManagedComputer** Specifies the agent managed computer(s) on which to change the assigned management server. This object can be retrieved with the Get-Agent cmdlet.
- FailoverServer** Specifies the failover server(s) for the agent(s).
- GatewayManagementServer** Specifies the gateway management server(s) for the agent(s).

Here are some examples:

```
> $Agent = Get-Agent `
    | Where {$_.Name -eq 'Server01.zygort.com'}
> $MgmtSvr = Get-ManagementServer `
    | Where (_.PrincipalName -eq 'OpsMgr')
> Set-ManagementServer `
    -AgentManagedComputer $Agent `
    -ManagementServer $MgmtSvr
```

Discovery Cmdlets

The cmdlets in this section are used when configuring Operations Manager discovery options. The cmdlet Start-Discovery is used to return a set of monitoring objects that can be used for operations such as installing agents. But before you can use the Start-Discovery cmdlet, you must configure it with criteria objects created by the other cmdlets listed in this section.

NEW-DEVICEDISCOVERYCONFIGURATION

This cmdlet is used to create a discovery criteria object based on a specified monitoring class and TCP/IP address range. This is useful when attempting to assign agents to monitor devices other than Windows client or server systems.

The monitoring class and address range are required parameters of this cmdlet, as shown in the example that follows. To view a list of the available monitoring classes, just execute the cmdlet Get-MonitoringClass.

Parameters for this cmdlet include:

- MonitoringClass** Specifies the Operations Manager monitoring class to scan for.
- FromIPAddress** Specifies the beginning of the TCP/IP address range to scan.
- ToIPAddress** Specifies the ending of the TCP/IP address range to scan.

Here are some examples:

```
> $NetworkDevice = Get-MonitoringClass `
    -Name 'System.NetworkDevice'
> $Criteria = New-DeviceDiscoveryConfiguration `
    -MonitoringClass $NetworkDevice
    -FromIPAddress 192.168.0.1 `
    -ToIPAddress 192.168.0.100
```

NEW-LDAPQUERYDISCOVERYCRITERIA

This cmdlet is used to create a discovery criteria object based on a Lightweight Directory Access Protocol (LDAP) query. Simply put, we're creating an object that contains an LDAP query to discover computers or network devices that are registered in a directory service. For example, it can be used to find a set of systems that we want to install Operations Manager agents on.

In order to use the cmdlet, you will want to either be familiar with the format of LDAP queries or have a tool to generate them. One of the most readily available tools is Active Directory Users And Computers. First, create a custom query that matches the results you are looking for, then you can copy and paste the LDAP query string into this cmdlet.

Parameters for this cmdlet include:

- Domain** Specifies the name of the domain to execute the LDAP query against.
- ldapQuery** Specifies the LDAP query to use. Improper syntax of an LDAP query can cause this cmdlet to fail, so it's a good idea to verify your syntax in another application first and copy it to this parameter.

Here is an example:

```
> $Criteria = new-LdapQueryDiscoveryCriteria `
    -Domain zygort
    -ldapquery "(cn=*)"
```

NEW-WINDOWSDISCOVERYCONFIGURATION

This cmdlet is used to create a discovery criteria object based on a set of specified computer parameters. This criteria object is necessary when using the Start-Discovery cmdlet. In order to create this object, you must first create an LDAPQueryDiscoveryCriteria object using the cmdlet new-ldapdiscoverycriteria.

Parameters for this cmdlet include:

- ComputerName** Allows you to specify which specific computers returned from the LDAP query you would like to include in the discovery results.
- ComputerType** Allows you to specify what computer type you would like to include in the discovery results. Values include Workstation, Server, and Both.

-ldapQuery Allows you to pass an LDAP criteria object (created through the New-LDAPQueryDiscoveryCriteria cmdlet) to the cmdlet. This is used to narrow the result set returned during the discovery process.

-PerformVerification Specifies whether to perform a verification of the discovered system. This parameter does not need a value; it only needs to be present on the command line.

-ActionAccountCredential Specifies the account credentials to use in the discovery. This object can be retrieved using the Get-Credential cmdlet.

Here are some examples:

```
> $DiscoverConfig = new-WindowsDiscoveryConfiguration `
    -ComputerName "SCOMAD.zygort.com", "Srv1.zygort.com"
> LDAPQuery = $Criteria = new-LdapQueryDiscoveryCriteria `
    -Domain zygort
    -ldapquery "($(objectCategory=computer)(cn=Svr01))"
> $DiscoverConfig = new-WindowsDiscoveryConfiguration `
    -ComputerType Server
    -ldapQueryDiscoveryCriteria $LDAPQuery
```

START-DISCOVERY

This cmdlet starts a task to discover computers and devices. Before running this cmdlet, you must decide which management server you want to handle the newly discovered systems. You must also create a search criteria object through either the New-WindowsDiscoveryConfiguration or New-DeviceDiscoveryConfiguration cmdlet.

Parameters for this cmdlet include:

-ManagementServer Required; specifies the management server that will be responsible for the computers or devices discovered. This can easily be retrieved by using the cmdlet Get-ManagementServer.

-WindowsDiscoveryConfiguration Specifies Windows discovery configuration objects (created by New-WindowsDiscoveryConfiguration) to use in describing the computers to discover. Use commas to separate multiple Windows discovery configuration objects.

-DeviceDiscoveryConfiguration Specifies device discovery configuration objects (created by New-DeviceDiscoveryConfiguration) to use in describing the computers to discover. Use commas to separate multiple device discovery configuration objects.

Here are some examples:

```
> $LDAPQuery = New-LdapQueryDiscoveryCriteria `
    -Domain zygort `
    -ldapquery "(cn=*)"
> $DiscoverConfig = new-WindowsDiscoveryConfiguration `
    -ComputerType Server `
    -ldapQueryDiscoveryCriteria $LDAPQuery
> $Results = Start-Discovery `
    -ManagementServer (get-ManagementServer) `
    -WindowsDiscoveryConfiguration $DiscoverConfig
> $Results | Select -expand customMonitoringObjects
```

Management Pack Cmdlets

The cmdlets that follow are designed for administration of Operations Manager management packs. Through these cmdlets, you will see how to list, add, remove, and export management packs from your system. At the end, there is a script that will let you list the rules contained in each installed management pack.

GET-MANAGEMENTPACK

This cmdlet is used to retrieve Operations Manager management pack(s). A `ManagementPack` object is returned, which can be used to uninstall or export the specified management pack with other cmdlets.

Parameters for this cmdlet include:

- Id** Specifies the GUID of the management pack to retrieve
- Name** Specifies the name of the management pack to retrieve

Here are some examples:

```
> Get-ManagementPack -Name 'System.Health.Library'
> Get-ManagementPack | Select Name | Sort Name
```

When looking at the rule cmdlets, you will note that some cmdlets will not work on sealed management packs. By using a `Where-Object` cmdlet, we can distinguish between sealed and unsealed management packs, as shown here:

```
> Get-ManagementPack | Where {$_.Sealed -eq $true}
> Get-ManagementPack | Where {$_.Sealed -eq $false}
```

INSTALL-MANAGEMENTPACK

This cmdlet is used to install an Operations Manager management pack from a specified XML file. Parameters for this cmdlet include:

- FilePath** Specifies the file path to the XML file for the management pack to install. Only one management pack at a time can be specified with this parameter.
- Path** Specifies the path (or paths, separated by commas) to the management groups into which management packs should be installed.
- Confirm** Invokes a prompt to confirm the installation of a management pack.

Here is an example:

```
> Install-ManagementPack -FilePath "C:\NewMP.xml"
```

UNINSTALL-MANAGEMENTPACK

This cmdlet is used to remove installed management packs from the current management server. The cmdlet is expecting a set of management pack objects to remove, which we can create by using the cmdlet `Get-ManagementPacks`.

Parameters for this cmdlet include:

- ManagementPack** Specifies the name of the management pack to remove.

-Path Specifies the path (or paths, separated by commas) to the management groups from which management packs should be uninstalled.

-Confirm Invokes a prompt to confirm the removal of the management pack.

Here are some examples:

```
> $MP = get-managementpack
    -name "Microsoft.Windows.Server.2000"
> uninstall-managementpack -managementPack $MP
```

EXPORT-MANAGEMENTPACK

This cmdlet will export a single management pack to an XML file. The cmdlet is expecting a management pack object to export, which we can retrieve by using the cmdlet `Get-ManagementPacks`. We must also supply a path for the new XML file to create. The newly created file will bear the name of the exported management pack.

Note that this cmdlet will export only one management pack at a time. To export multiple management packs, we can create a set of management pack objects and use a `ForEach` loop to export each pack individually. The second example that follows will illustrate this approach.

Parameters for this cmdlet include:

-ManagementPack Specifies the management pack(s) to export.

-Path Specifies the file system path to export the management pack(s) to. This is not the same as the `Path` parameters you'll see in other cmdlets, which refer to management paths.

Here are some examples:

```
> $MP = Get-ManagementPack -Name "MyCustomMP"
> Export-ManagementPack `
    -ManagementPack $MP `
    -FilePath "C:\ExportedManagementPacks"

> $MPList = get-managementpack `
    | where {$_.sealed -eq $false}
> foreach ($MgmtPack in $MPList) {
    Export-ManagementPack `
        -ManagementPack $MP `
        -FilePath "C:\ExportedManagementPacks"
}
```

SCRIPT: LISTING RULES IN A MANAGEMENT PACK

This script is designed to allow you to list the rules of a specified management pack on the current management server. Given that management pack names are sometimes difficult to recall off the top of one's head, this script uses the `-match` criteria when locating the specified management pack. This will allow the user to actually search multiple management packs with similar names.

```
Write-Host "List Management Pack Rules";
$MPName = Read-Host "Enter the Management Pack"
$MPList = Get-ManagementPack | Where {$_.Name -match $MPName};
If ($MPList -eq $null) {
```

```

        Write-Host "No matching Management Pack found!";
    }
    Else {
        $MPList | ForEach {
            Write-Host "Rules in Management Pack" $_.Name;
            Write-Host " -----";
            $RuleList = Get-Rule -ManagementPack:$_ | Sort Name;
            If ($RuleList -eq $null) {
                Write-Host "No rules found.";
            }
            Else {
                $RuleList | ForEach {
                    Write-Host $_.Name;
                }
            }
            Write-Host " ====="
        }
    }
}

```

SCRIPT: CREATING RULE FILES FOR EACH MANAGEMENT PACK

You can easily create a list of the rules contained in each management pack in a file format. Note that in the first portion of the Where-Object cmdlet, a variable called \$FN, is created. This is the file-name that will be used by the Out-File cmdlet, and is created by prefixing a path onto the name of the current management pack and adding the extension .txt.

```

Write-Host "Create Management Pack Rule Files"
$FileLocation = Read-Host "Enter a file location"
Get-ManagementPack | `
    ForEach { `
        $FilePath = $FileLocation + $_.Name + ".txt"; `
        Get-Rule -ManagementPack:$_ | Select Name | `
        Out-File -FilePath:$FilePath -Encoding:ASCII
    }

```

Rule Cmdlets

The cmdlets in this section are designed for managing rules in Operations Manager. Our options include listing, enabling, and disabling rules. At the end of the section are some additional scripts for working with rules.

GET-RULE

This cmdlet is used to retrieve Operations Manager rule objects. After obtaining a reference to a set of rules, you can check their state and enable or disable them as you wish.

Parameters for this cmdlet include:

-Id Used to specify the GUID of a specific rule to work with.

-MonitoringObject When specified, only rules associated with the specified monitoring object are returned.

-ManagementPack When specified, only rules associated with the specified management pack are returned. You can use the `Get-ManagementPack` cmdlet to get an appropriate argument.

Here are some examples:

```
> Get-Rule | Select Name | Sort Name
> Get-Rule | Where {$_.Name -match "Exchange"}
> $MP = Get-ManagementPack `
    -name Microsoft.SystemCenter.Internal
> Get-rule -managementpack $MP
```

DISABLE-RULE

This cmdlet disables Operations Manager monitoring rule(s). Rules that belong to a sealed management pack cannot be disabled. Rule objects can be retrieved by using the `Get-Rule` cmdlet. You can disable either one or many rules.

Parameters of this cmdlet include:

- Rule** Specifies the monitoring rule to disable. This object can be retrieved with the `Get-Rule` cmdlet. This parameter accepts only a single monitoring rule object.
- Confirm** Prompts for confirmation before disabling the monitoring rule.

Here is one example:

```
> $MP = Get-ManagementPack -Name:MyCustomMP
> $Rule = get-rule -ManagementPack:$MP `
    | Where {$_.Name -match "Application.Alert"}
> Disable-Rule $Rule
```

Another example, a bit more streamlined:

```
> Get-Rule -ManagementPack:$MP `
    | Where {$_.Name -match "Application.Alert"} `
    | Disable-Rule
```

ENABLE-RULE

This cmdlet enables Operations Manager monitoring rule(s). Rules that belong to a sealed management pack cannot be enabled. Rule objects can be retrieved by using the `Get-Rule` cmdlet. You can enable either one or many rules, depending on the scripting technique used.

Parameters of this cmdlet include:

- Rule** Specifies the monitoring rule to enable. This object can be retrieved with the `Get-Rule` cmdlet. This parameter only accepts a single monitoring rule object.
- Confirm** Prompts for confirmation before enabling the monitoring rule.

Here are some examples:

```
> $MP = get-managementpack -Name:MyCustomMP
> $Rule = Get-Rule -ManagementPack:$MP |
    Where {$_.Name -match "Application.Alert"}
```

```
> Enable-Rule $Rule
```

Another example, a bit more streamlined:

```
> Get-Rule -ManagementPack:$MP `
  | Where {$_.Name -match "Application.Alert"} `
  | Enable-Rule
```

SCRIPT: LIST ALL RULES

The PowerShell script will list all the rules available on the current management server. We could run the `Get-Rule` cmdlet, but that would tend to scatter rules about in an unorganized fashion. Instead, we'll list each set of rules according to their respective management pack.

```
# ListAllRules.ps1
Write-Host "List All Rules";
$MPList = Get-ManagementPack;
$MPList | ForEach {
    Write-Host "Rules in Management Pack" $_.Name;
    Write-Host " -----";
    $RuleList = Get-Rule -ManagementPack:$_ | Sort Name;
    If ($RuleList -eq $null) {
        Write-Host "No rules found.";
    }
    Else {
        $RuleList | ForEach {
            Write-Host $_.Name;
        }
    }
    Write-Host " ====="
}
}
```

SCRIPT: SEARCHING FOR SPECIFIC RULES

Next I have two scripts for your consideration. This first script prompts the user for the name of a rule to search for. The search uses the `-match` criteria, so the entire rule name does not have to be entered at the prompt. For example, you can simply search for `EventLogFull` rather than `Microsoft.Windows.Server.2003. OperatingSystem.EventLogFull.Alert`.

```
# Searches the entire scope of rules
# SearchAllMPRules.ps1
Write-Host "Find Management Pack Rule";
$RuleName = Read-Host "Enter the Rule name";
$RuleList = Get-Rule | Where {$_.Name -match $RuleName};
If ($RuleList -eq $null) {
    Write-Host "No matching rules found! "
}
Else {
    Write-Host $RuleList.Count "rule(s) found."
    $RuleList | ForEach {
        Write-Host $_.Name
    }
}
}
```

We can narrow the search down to specific management packs by asking the user for a little more information at the start of the script:

```
# Narrows the search down to Management Pack
# SearchMPRules.ps1
Write-Host "Find Management Pack Rule";
$MPName = Read-Host "Enter the Management Pack to search"
$RuleName = Read-Host "Enter the Rule name";
$RuleList = Get-Rule | Where {$_.Name -match $RuleName};
If ($RuleList -eq $null) {
    Write-Host "No matching rules found!"
}
Else {
    Write-Host $RuleList.Count "rule(s) found."
    $RuleList | ForEach {
        Write-Host $_.Name
    }
}
}
```

Task Cmdlets

The cmdlets in this section are designed for managing tasks in Operations Manager. Our options include listing, enabling, and disabling tasks. At the end of the section are some additional scripts for working with tasks.

GET-TASK

This cmdlet is used to retrieve the specified Operations Manager monitoring task objects. With a reference to a specific set of tasks, you can manually run the task using the Start-Task cmdlet. Results of running the task can be viewed with the Get-TaskResult cmdlet.

Parameters of this cmdlet include:

- Path** Specifies the path (or paths, separated by commas) to the monitoring objects from which to retrieve tasks.
- MonitoringObject** When specified, the tasks associated with the specified monitoring object are retrieved.
- Id** Specifies the GUID of the monitoring object task to retrieve.

Here are some examples:

```
> Get-Task | Select Name, Category
> Get-Task | Where {$_.Name -match "CheckDCs"}
> Get-Task | Where {$_.Enabled -eq $true} | Select Name
> Get-Task | Group Category
> Get-Task | Where {$_.Category -match "Maintenance"}
```

START-TASK

Starts a specified Operations Manager task. Task objects can be retrieved by using the Get-Tasks cmdlet, as shown earlier. Be aware that this cmdlet will block for the duration of the task (in other

words, you won't get a command prompt back until it's finished). You can run the task in the background by specifying the `-asynchronous` parameter.

Parameters for this cmdlet include:

-Task Specifies the monitoring task to start. This is a `MonitoringTask` object, which can be retrieved with the `Get-Task` cmdlet.

-Overrides Specifies any task overrides. Override objects for a given task can be found with the `Get-Override` cmdlet. To see if a task has any overrides, you can use the following example. Replace `CheckDCs` with the task you are interested in:

```
> Get-Task `
    | Where {$_.Name -match "CheckDCs"} `
    | Get-Override
```

-Credential Specifies the credentials under which to run the monitoring task. Credentials can be obtained using the cmdlet `Get-Credential`.

-Asynchronous When specified, the monitoring task runs asynchronously. In other words, the command prompt returns right away while the task runs in the background. If this parameter is not specified, the command will wait until the task completes before returning control.

-TargetMonitoringObject Optionally specifies the monitoring object to run the task against. Monitoring objects can be retrieved with the `Get-MonitoringObject` cmdlet. This parameter accepts only a single monitoring object.

-Confirm Prompts for confirmation before running the task.

Here are some examples:

```
> $Task = Get-Task `
    | Where {$_.Name -match "CheckDCs"}
> $Overrides = Get-Override -Task $Task
> $Monitor = Get-MonitoringObject `
    | Where {$_.PathName -match "DomainController"}
> Start-Task -Task $Task -Asynchronous `
    -Override $Overrides -MonitoringObject $Monitor
```

GET-TASKRESULT

This cmdlet is used to retrieve the results of a task. The easiest way to take advantage of this is to assign the task started with the `Start-Task` cmdlet to a variable. You can then pass the ID contained in the task variable to the `Get-TaskResult` cmdlet. Task results are stored in the `Output` property in an XML format.

When run by itself, `Get-TaskResult` will list all the accumulated task results known on the management server. Over time, this will grow to quite a list, so you will want to try narrowing it down with either a `Where-Object` cmdlet or the `-criteria` parameter.

Parameters of this cmdlet include:

-Path Specifies the path (or paths, separated by commas) of the management group connections for the desired tasks. Task results are returned for tasks run in their respective management groups.

-Id Specifies the GUID of the task result to obtain.

-BatchID Specifies the GUID of the batch of the tasks to retrieve the results of.

Here are some examples:

```
> Get-TaskResult | Where {$_ .Status -eq "Succeeded"}
> $Start = (get-date).AddDays(-1)
> Get-TaskResult | Where {$_ .TimeStarted -lt $Start}\
> $Task = Get-Task `
    | Where {$_ .DisplayName -match "Run Chkdk"}
> $RunTask = Start-Task -task $Task -asynchronous
> Get-TaskResult $RunTask.ID | Select Output
```

Action Cmdlets

The cmdlets in this section are designed for managing agent actions in Operations Manager. Our options include listing the pending agent actions, and approving or denying a pending agent action.

GET-AGENTPENDINGACTION

This cmdlet is used to retrieve pending Operations Manager actions. The only parameter used by this cmdlet is **-Path**, which specifies the path to the management group (or groups, separated by commas) from which to retrieve the list of pending agent actions.

Here are some examples:

```
> Get-AgentPendingAction
> Get-AgentPendingAction `
    | Where {$_ .AgentName -match "SCOMAD"}
```

APPROVE-AGENTPENDINGACTION

This cmdlet is used to approve a pending agent action. Pending actions can be retrieved by using the `Get-AgentPendingAction` cmdlet.

Parameters of this cmdlet include:

-AgentPendingAction Specifies the pending actions to approve. This parameter expects an `AgentPendingAction` object, which is retrieved with the `Get-AgentPendingAction` cmdlet. You can approve multiple actions by separating the `AgentPendingAction` objects with commas.

-Confirm Prompts for confirmation before approving the pending actions. This can be useful when processing many pending actions.

Here are some examples:

```
> $PendingActions = Get-AgentPendingAction
> Approve-AgentPendingAction `
    -AgentPendingAction $PendingActions

> Get-AgentPendingAction `
    | Approve-AgentPendingAction -Confirm
```

REJECT-AGENTPENDINGACTION

This cmdlet is used to reject a pending agent action. Pending actions can be retrieved by using the `Get-AgentPendingAction` cmdlet.

Parameters of this cmdlet include:

-AgentPendingAction Specifies the pending actions to reject. This parameter expects an `AgentPendingAction` object, which is retrieved with the `Get-AgentPendingAction` cmdlet. You can approve multiple actions by separating the `AgentPendingAction` objects with commas.

-Confirm Prompts for confirmation before rejecting the pending actions. This can be useful when processing many pending actions.

Here are some examples:

```
> $PendingActions = Get-AgentPendingAction
> Reject-AgentPendingAction `
    -AgentPendingAction $PendingActions

> Get-AgentPendingAction `
    | Reject-AgentPendingAction -Confirm
```

Notification Cmdlets

The cmdlets in this section are designed for managing notification operations in Operations Manager. Our options include listing the current subscription settings, and enabling and disabling rules. Notifications must be configured in the Operations Console prior to using these cmdlets. At this time, the ability to create or configure notification subscriptions and recipients is not supported in Operations Manager PowerShell.

GET-NOTIFICATIONSUBSCRIPTION

This cmdlet is used to retrieve notification subscriptions associated with a particular notification subscription ID. You can retrieve all the related subscriptions or one particular subscription.

Parameters of this cmdlet include:

-Path Specifies the path (or paths, separated by commas) of the management groups from which to retrieve the notification subscriptions.

-Id Specifies the GUID of the notification subscription to retrieve.

Here are some examples:

```
> Get-NotificationSubscription `
    | Select DisplayName, ToRecipients, Enabled
> Get-NotificationSubscription `
    | Where { $_.Disabled -eq $false }
> Get-NotificationSubscription `
    | Where { $_.DisplayName -eq "SQL Subscription" }
```

ENABLE-NOTIFICATIONSUBSCRIPTION

This cmdlet is used to enable a notification subscription. Parameters of this cmdlet include:

- NotificationSubscription** Specifies the subscription to enable. This object can be retrieved with the Get-NotificationSubscription cmdlet.
- Confirm** Prompts for confirmation before disabling the subscription.

Here are some examples:

```
> $Subscription = Get-NotificationSubscription `
    | Where {$_.DisplayName -eq "SQL Subscription" }
> Enable-NotificationSubscription `
    -NotificationSubscription $Subscription `
    -Confirm
```

DISABLE-NOTIFICATIONSUBSCRIPTION

This cmdlet is used to disable a notification subscription. Parameters of this cmdlet include:

- NotificationSubscription** Specifies the subscription to disable. This object can be retrieved with the Get-NotificationSubscription cmdlet.
- Confirm** Prompts for confirmation before disabling the subscription.

Here are some examples:

```
> $Subscription = Get-NotificationSubscription `
    | Where {$_.DisplayName -eq "SQL Subscription" }
> Disable-NotificationSubscription `
    -NotificationSubscription $Subscription `
    -Confirm
```

GET-NOTIFICATIONACTION

This cmdlet is used to retrieve notification actions associated with a specific notification action ID. You can retrieve all the related notification actions or one particular action.

Parameters of this cmdlet include:

- Path** Specifies the path (or paths, separated by commas) of the management groups from which to retrieve the notification actions.
- Id** Specifies the GUID of the notification action to retrieve.

Here is an example:

```
> Get-NotificationAction `
    | Select Endpoint, Name, Body
```

GET-NOTIFICATIONENDPOINT

This cmdlet is used to retrieve notifications endpoints associated with a specific notification endpoint ID. You can retrieve all the related notification endpoints or one particular endpoint.

Parameters of this cmdlet include:

- Path** Specifies the path (or paths, separated by commas) of the management groups from which to retrieve the notification endpoints.
- Id** Specifies the GUID of the notification endpoint to retrieve.

Here are some examples:

```
> Get-NotificationEndpoint `
    | Select Name, Description
> Get-NotificationEndpoint `
    | Where {$_ .Name -match "SmtPEndpoint"}
> Get-NotificationEndpoint `
    | Where {$_ .Name -match "ImEndpoint"}
> Get-NotificationEndpoint `
    | Where {$_ .Name -match "SmsEndpoint"}
```

GET-NOTIFICATIONRECIPIENT

This cmdlet is used to retrieve notification recipients associated with a specific notification recipient ID. You can retrieve all the related notification recipients or one particular recipient.

Parameters of this cmdlet include:

- Path** Specifies the path (or paths, separated by commas) of the management groups from which to retrieve the notification recipients.
- Id** Specifies the GUID of the notification recipient to retrieve.

Here are some examples:

```
> Get-NotificationRecipient `
    | Select Name, Devices
> Get-NotificationRecipient `
    | Where {$_ .Name -match "SCOMAdmin"} `
    | Select Name, Devices
```

Alert Cmdlets

With alert objects, you can view and manage pending alerts on your Operations Manager server. You can also wade through the alert history as recorded in the Operations Manager database.

GET-ALERT

This cmdlet is used to retrieve Operations Manager alerts. This allows you to view pending alerts and either resolve an alert or manage other aspects of the alert. The information contained in an alert object can be very extensive, so take advantage of the Format-List cmdlet (whose alias is FL) to make the output more readable (especially if you decide to check out the `Description` property).

The Get-Alert cmdlet will return all known alerts in the Operations Manager database, regardless of resolution state. In new alerts, the `ResolutionState` property will have a value of zero.

In closed alerts, the property will have a value of 255. You might also decide to prioritize your alert management by the value of the `Severity` property, which will be either `Error` (these will appear as `Critical` in the Operations Manager console), `Warning`, or `Information`.

Parameters of this cmdlet include:

-Criteria Used to specify a specific set of alerts to retrieve. A `Where-Object` cmdlet can also be used to narrow down the result set. You might find it easier to use the `Where-Object` cmdlet instead.

-Recurse By default, only alerts for the current monitoring object are retrieved. When the `-Recurse` parameter is used, alerts for the child objects of the current monitoring object are also retrieved.

-Id Retrieves the alert with the specified GUID.

-MonitoringObject Retrieves alerts for the specified monitoring object.

Here are some examples:

```
> Get-Alert | Where {$_.ResolutionState -eq 0} | `
    Select PrincipalName, Name, Severity, Description | fl
> Get-Alert | `
    Where {$_.ResolutionState -eq 0 -and `
        $_.Severity -eq "Error"} | `
    Select PrincipalName, Name, Description | fl
> Get-Alert | Group Severity | `
    Sort Name | Select Count, Name
```

RESOLVE-ALERT

This cmdlet is used to resolve a set of active Operations Manager alerts. The alert objects can be retrieved by using the `Get-Alert` cmdlet, as shown earlier. This cmdlet can accept only one alert object at a time, so care must be taken when scripting this cmdlet.

Parameters of this cmdlet include:

-Alert Specifies the alert to resolve.

-Comment An optional comment to be included with the history of the alert object.

-Confirm Causes a prompt to confirm resolution of the alert. This is a good parameter to take advantage of when sifting through a set of alerts.

Here are some examples:

```
> Get-Alert | Where {$_.ResolutionState -eq 0} | `
    Resolve-Alert -confirm
> $InformationAlerts = Get-Alert | `
    Where {$_.ResolutionState -eq 0 -And `
        $_.Severity -match "Information"}
> ForEach ($A in $InformationAlerts) { `
    Resolve-Alert -alert $A `
        -comment "Resolving Information Alerts"}
```

GET-ALERTDESTINATION

This cmdlet is used to retrieve the destination connection for a specified set of Operations Manager alerts. This cmdlet can accept only one alert object at a time, so care must be taken when scripting this cmdlet.

This cmdlet has only one parameter, `-Alert`, which is used to specify the alert to retrieve the destination connector for. This parameter can accept only one alert object at a time. We can either pipe the output from the `Get-Alert` cmdlet into the `Get-AlertDestination` cmdlet, or we can assign the set of alerts to a variable and use a `ForEach-Object` cmdlet to loop through each single alert.

Here are some examples:

```
> Get-Alert | Where {$_.Name -match "Exchange"} | `
    Get-AlertDestination
> $Alerts = Get-Alert | Where {$_.Name -match "Exchange"}
> ForEach ($A in $Alerts) {Get-AlertDestination -alert $A}
```

SET-ALERTDESTINATION

This cmdlet is used to set the destination connector for a specified set of Operations Manager alerts. This cmdlet can accept only one alert object at a time, so care must be taken when scripting this cmdlet. Before using this cmdlet, it is recommended that you verify the current destination of alerts through the `Get-AlertDestination` cmdlet.

Parameters for this cmdlet include:

- Alert** Specifies the alert to modify the destination connector for.
- Connector** Specifies the Operations Manager connector to set as the new alert destination. This can be retrieved by using the `Get-Connector` cmdlet.
- Comment** An optional comment about the new alert destination setting.
- Confirm** Causes a prompt to verify the change in destination connector.

Here are some examples:

```
> $NewDestination = Get-Connector | `
    Where {$_.Name -match "OpsMgr2"}
> Get-Alert | Where {$_.Name -match "Exchange"} | `
    Set-AlertDestination -connector $NewDestination
```

GET-ALERTHISTORY

This cmdlet allows you to retrieve the history for a set of specified Operations Manager alerts. This cmdlet can accept only one alert object at a time, so care has to be taken when using this cmdlet. You will see this in action in the examples that follow.

This cmdlet has only one parameter, `-Alert`, which specifies the alert to retrieve the history of. Alerts retrieved from the `Get-Alert` cmdlet can be used in this parameter (one at a time, such as in a `ForEach-Object` loop), or the `Get-Alert` cmdlet can be piped directly into the `Get-AlertHistory` cmdlet.

Here are some examples:

```
> Get-Alert | Where {$_.ResolutionState -ne 0} | `
    Get-AlertHistory
```

```
> $Alerts = Get-Alert | Where {$_ .ResolutionState -ne 0}
> ForEach ($A in $Alerts) { Get-AlertHistory -alert $A }
> ForEach ($A in $Alerts) { `
    Write-host $A.PrincipalName $A.Name;
    Get-AlertHistory -alert $A | Sort TimeModified | `
        Select TimeModified, Comments
    }
}
```

Miscellaneous Cmdlets

These are the cmdlets that don't fit neatly into any specific category but that you'll still find essential from time to time. They also are not strictly Operations Manager cmdlets—you will find a sprinkling of standard PowerShell cmdlets in this section as well that you will want to take advantage of.

GET-CREDENTIAL

This cmdlet is used to create a PowerShell `Credential` object (also known as a `PSCredential` object). This type of object is used when a cmdlet needs a set of user credentials in order to execute, usually with a parameter named `-Credential`. `PSCredential` objects are used not only in Operations Manager cmdlets, but also with other PowerShell cmdlets such as `Get-WMIObject`.

The only parameter used with the cmdlet is the `-Credential` parameter. It is used to specify a username to be used in the `PSCredential` object. The username can be accepted in the form `MyUser` or `MyDomain\MyUser` as needed. If this parameter is omitted, a dialog box will be displayed to retrieve a username and password.

You might have noticed there is no `-Password` parameter. This is for security reasons—to prevent a password from being plainly visible in a stored PowerShell script. As a result, whenever a cmdlet is executed that uses the `Get-Credential` cmdlet, a dialog box will be displayed to prompt the user for a password.

Here are some examples:

```
> $Cred = Get-Credential -Credential zygort\SCOM_Admin
> Get-WMIObject Win32_Product `
    -ComputerName SCOMAD `
    -Credential (Get-Credential)
```

GET-EVENT

This cmdlet is used to retrieve Operations Manager events. Events are generated according to rules on the management server. The GUID of the rule is returned with the event, which can be passed into the `Get-Rule` cmdlet to retrieve the rule definition.

Parameters of this cmdlet include:

- Path** Specifies the path (or paths, separated by commas) to monitoring objects from which events should be retrieved.
- MonitoringObject** Specifies the monitoring object from which events should be retrieved.
- Criteria** Specifies the criteria that events must match to be returned.
- Recurse** Determines whether to retrieve events for only the specified monitoring object(s) or to include child monitoring object events.
- Id** Specifies the GUID of the event to retrieve.

Here are some examples:

```
> Get-Event -Path "SCOM.zygort.com"
> Get-Event | Where {$_.LoggingComputer -match "SCOMAD"}
> Get-Event `
    | Select LoggingComputer, Description
> Get-Event `
    | Where {$_.LoggingComputer -match "SCOMAD"}
    | Select MonitoringRuleID `
    | ForEach {Get-Rule -Id $_.MonitoringRuleID} `
```

GET-USERROLE

This cmdlet retrieves the specified Operations Manager user role objects. You can use this cmdlet to either list the available user roles on a management server or modify the user membership in a specified role.

Parameters of this cmdlet include:

- Path** Specifies the path (or paths, separated by commas) to the management groups from which to retrieve user roles.
- Id** Specifies the GUID of the user role to retrieve.

Here are some examples:

```
> Get-UserRole
> Get-UserRole | Select Name, Description | Format-List
```

You can also use this cmdlet to search the Users collection of each specified user role or to find which user roles a specific user belongs to:

```
> Get-UserRole | ForEach {Write-Host $_.name "-" $_.Users}
> Get-UserRole `
    | Where {$_.Users -match "Administrators"} `
    | Select Name
```

SCRIPT: LISTING USER ROLE MEMBERSHIP

This basic script can be used for a quick and dirty listing of user roles and their group memberships:

```
Get-UserRole `
    | ForEach {Write-Host $_.Name " - " $_.Users}
```

To make the output easier on the eyes, we can use a couple of ForEach loops to separate the user role and individual members onto separate lines:

```
Write-Host "User Role Membership"
$Roles = Get-UserRole
$Role | ForEach {
    Write-Host $_.Name;
    $_.Users | ForEach {
        Write-Host "---- " $_
    }
}
```

We can take user role listing one step further and send the output to a file for later sorting or other analysis. The script that follows will display the same listing as the earlier script, but first it will prompt for a file system location in which to create the user role listing file:

```
Write-Host "User Role Memberships"
$Roles = Get-UserRole
$Role | ForEach {
    Write-Host $_.Name;
    $_.Users | ForEach {
        Write-Host "---- " $_
    }
}
```

SCRIPT: SEARCH FOR USER ROLE MEMBERSHIP

The following script will prompt for a name (either username or group name will work) and search the available user roles. When a user role is found whom the specified name belongs to, the script displays the name of the user.

```
Write-Host "Search for User Role membership"
$Member = Read-Host "Enter name to search for"
Get-UserRole | Where {$_ .Users -match $Member}
    | ForEach {write-host $_.Name}
```

SCRIPT: ADD MEMBER TO USER ROLE

With this script, you can interactively add a new member to a specified user role. The script will first prompt for the user role to add a new member to. It will then try to find that user role to verify it actually exists and stop with a warning if it is not found on the current management server.

If the user role is successfully found, then a loop is started that allows you to add multiple users to the user role. Typing anything other than "Y" or "y" will end the loop and return to the command prompt.

```
Write-Host "Add Member to User Role"
$UserRoleName = Read-Host "Enter User Role"
$UserRole = Get-UserRole
    | Where {$_ .Name -match $UserRoleName}
If ($UserRole -eq $null) {
    Write-Host "User Role not found!"
}
Else {
    Do {
        Write-host "Found User Role " $UserRole.Name
        $Member = Read-Host "Enter new member name: "
        $UserRole.Users.Add($Member)
        $UserRole.Update()
    } While ((Read-Host "Add another(Y/N)?" -eq "Y")
}
```

SCRIPT: REMOVE MEMBER FROM USER ROLE

Almost identical to the Add Member script, this script will remove members from a specified user role:

```
Write-Host "Remove Member from User Role"
$UserRoleName = Read-Host "Enter User Role"
$UserRole = Get-UserRole
    | Where {$_.Name -match $UserRoleName}
If ($UserRole -eq $null) {
    Write-Host "User Role not found!"
}
Else {
    Do {
        Write-host "Found User Role " $UserRole.Name
        $Member = Read-Host "Enter member name: "
        $UserRole.Users.Remove ($Member)
        $UserRole.Update()
    } While ((Read-Host "Remove another(Y/N)?" ) -eq "Y")
}
```

The last two scripts to add and remove members of user roles take advantage of the `-match` criteria when searching for the user role. This can be good or bad. On the good side, you can type a shortened version of the name, such as **AdvancedOperators**, rather than the full role name, **OperationsManagerAdvancedOperators**. On the bad side, if you use a short name that is common to several user roles, you might wind up connecting the user to the wrong role!

Management Server Cmdlets

These cmdlets are used to manage connections to Operations Manager management servers and administer the default settings. When the Operations Manager command shell is opened, a connection is always attempted to the default management server of the system. When performing administrative tasks on other servers, you must first connect to that management server.

GET-MANAGEMENTSERVER

This cmdlet retrieves the specified management server(s) from the current management group. We can either retrieve information about the current management server or specify a different management server on which to execute more management tasks.

This cmdlet has one parameter:

-Root When specified, the cmdlet retrieves only the root management server of the group. This parameter does not take a value; it only needs to appear in the command to be considered true. By default, this value is false, which will return all the management servers in the current management group.

Here are some examples:

```
> Get-ManagementServer
> Get-ManagementServer -Root
> Get-ManagementServer | Where {$_.Name -match "SCOM"}
> Get-ManagementServer | Select Name, ManagementGroup
```

GET-DEFAULTSETTING

This cmdlet is used to retrieve the default configuration settings for the current management server. This cmdlet accepts one parameter, `-Path`, which specifies the path (or paths, separated by commas) to the management groups from which to retrieve default settings. In order to retrieve a specific default setting, you must use the `Where-Object` cmdlet, as shown in these examples:

```
> Get-DefaultSetting
> Get-DefaultSetting | Where {$_.Name -match "Heartbeat" }
```

The last command will usually yield several results. We can narrow it down by being more specific about the default setting name:

```
> Get-DefaultSetting `
  | Where {$_.Name -match "Agent\Heartbeats\Interval" }
```

SET-DEFAULTSETTING

This cmdlet is used to set one or more of the default configuration settings for an Operations Manager management server. `DefaultSetting` objects can be retrieved by using the `Get-DefaultSetting` cmdlet or by specifying the full name of the object to modify.

Parameters of this cmdlet include:

-Name The name of the default server setting you want to modify. This parameter is case-sensitive.

-Value The new value for the default server setting.

-Path Specifies the path (or paths, separated by commas) to the management groups to the new default settings should be applied.

-Confirm Prompts for confirmation before applying the new default settings to the management groups.

Here is an example:

```
> Set-DefaultSetting `
  -Name "Agent\Heartbeats\Interval" `
  -Value 30
```

GET-CONNECTOR

This cmdlet is used to retrieve the connectors for a management group. These connector objects can be used with other cmdlets, such as `Set-AlertDestination`.

Parameters of this cmdlet include:

-Path Specifies the path (or paths, separated by commas) to the management groups to retrieve connectors for.

-MonitoringAlert Specifies the alert for which to retrieve connectors. The `Get-Alert` cmdlet can be used to retrieve the alert object for this parameter.

-Id Specifies the GUID of the connector to retrieve.

Here are some examples:

```
> Get-Connector | Where {$_.Name -Match "SQLConnector" }
```

```
> Get-Connector -Path Microsoft.SQLServer.ComputerGroup
> $Alert = Get-Alert | Where {$_.Name -match "SQL"}
> Get-Connector -MonitoringAlert $Alert
```

GET-MANAGEMENTGROUPCONNECTION

This cmdlet is used to find all open connections to Operations Manager management groups. This cmdlet does not take any parameters.

Here are some examples:

```
> Get-ManagementGroupConnection
> Get-ManagementGroupConnection `
  | Where {$_.ManagementGroup -match "zygort"}
```

NEW-MANAGEMENTGROUPCONNECTION

This cmdlet creates and opens a new connection to a server in a management group. After successfully executing this cmdlet, any future operations will be directed to that management group.

Parameters of this cmdlet include:

-ConnectionString Specifies the connection string to the desired management group. This is usually the DNS name of the primary management server in the management group.

-Credential Specifies the `Credential` object to use when connecting to the new management group. The object can be retrieved using the `Get-Credential` cmdlet.

Here are some examples:

```
> New-ManagementGroupConnection `
  -ConnectionString:"scm2.zygort.lcl"
> New-ManagementGroupConnection "scm2.zygort.lcl"
```

REMOVE-MANAGEMENTGROUPCONNECTION

This cmdlet is used to remove a connection to an Operations Manager management group. `ManagementGroupConnection` objects can be retrieved by using the `Get-ManagementGroupConnection` cmdlet.

Parameters of this cmdlet include:

-Path Specifies the path (or paths, separated by commas) to the management groups from which to remove the connection.

-Connection Specifies the management connection to remove. This object can be retrieved with the `Get-Connector` cmdlet.

-Confirm Prompts for confirmation before removing the management group connection.

Here is an example:

```
> Get-ManagementGroupConnection `
  | Where {$_.ManagementGroup -match "zygort2" } `
  | Remove-ManagementGroupConnection
```

Maintenance Window Cmdlets

These cmdlets are used to create, modify, and retrieve information about monitoring object maintenance windows. Maintenance windows are used to suspend monitoring of a system while work is being performed on the system.

When creating or modifying a maintenance window, a reason can be assigned to the window. The list of reason codes accepted by the `New-MaintenanceWindow` and `Set-MaintenanceWindow` cmdlets include:

- ◆ `PlannedOther`
- ◆ `UnplannedOther`
- ◆ `PlannedHardwareMaintenance`
- ◆ `UnplannedHardwareMaintenance`
- ◆ `PlannedHardwareInstallation`
- ◆ `UnplannedHardwareInstallation`
- ◆ `PlannedOperatingSystemReconfiguration`
- ◆ `UnplannedOperatingSystemReconfiguration`
- ◆ `PlannedApplicationMaintenance`
- ◆ `UnplannedApplicationMaintenance`
- ◆ `ApplicationInstallation`
- ◆ `ApplicationUnresponsive`
- ◆ `ApplicationUnstable`
- ◆ `SecurityIssue`
- ◆ `LossOfNetworkConnectivity`

GET-MAINTENANCEWINDOW

This cmdlet is used to retrieve the maintenance window information for a specified Operations Manager monitoring object. The monitoring object must be in maintenance mode.

Parameters for this cmdlet include:

- History** Specifies that the history of the maintenance window should be retrieved. This parameter does not take a value; it only needs to be present in the parameter list.
- Path** Specifies the path (or paths, separated by commas) to the monitoring objects on which to retrieve the maintenance window.
- MonitoringObject** Specifies the monitoring object to retrieve the maintenance window. This parameter accepts only a single monitoring object, which can be retrieved with the `Get-MonitoringObject` cmdlet.

Here are some examples:

```
> Get-MaintenanceWindow
```

```

> $Monitor = Get-MonitoringObject | `
    Where {$_.DisplayName -match "SQL 2005 Computers"}
> Get-MaintenanceWindow -MonitoringObject $Monitor

> Get-MaintenanceWindow `
    -Path Microsoft.SqlServer.ComputerGroup `
    -History

```

NEW-MAINTENANCEWINDOW

This cmdlet is used to create a new Operations Manager maintenance window for a monitoring object. Maintenance windows are used to suspend monitoring while maintenance is being performed on the monitoring object system.

Parameters for this cmdlet include:

- StartTime** Specifies the beginning time of the maintenance window. This parameter will accept either a string value or a `DateTime` object.
- EndTime** Specifies the ending time of the maintenance window. This parameter will accept either a string value or a `DateTime` object.
- Path** Specifies the path (or paths, separated by commas) to the monitoring objects on which to set the maintenance window.
- MonitoringObject** Specifies the monitoring object that will be placed in maintenance mode for the duration of the window. The `Get-MonitoringObject` cmdlet can be used to supply this object.
- Reason** Optionally assigns a reason to the maintenance window. Any code from the list at the beginning of this section is acceptable.
- Comment** An optional comment about the maintenance window.
- Confirm** Causes the cmdlet to prompt the user to confirm the creation of the new maintenance window.

Here are some examples:

```

> New-MaintenanceWindow `
    -Path Microsoft.SqlServer.ComputerGroup
    -StartTime "12/01/2007" -EndTime "12/02/2007" `
    -Reason PlannedHardwareInstallation `
    -Comment "Additional memory install"

```

SET-MAINTENANCEWINDOW

This cmdlet is used to update the properties of a specified maintenance window. The monitoring object must be in maintenance mode. `MaintenanceWindow` objects can be retrieved by using the `Get-MaintenanceWindow` cmdlet.

Parameters for this cmdlet include:

- Path** Specifies the path (or paths, separated by commas) to the monitoring objects on which to set the maintenance window.

- MonitoringObject** Specifies the monitoring object on which to set the maintenance window.
- Endtime** Specifies the time the maintenance window will end. The value can be either a string format or a `DateTime` object.
- Reason** This parameter specifies the reason for the maintenance window. Accepted codes from the list at the beginning of this section are acceptable.
- Comment** An optional comment about the maintenance window.

Here are some examples:

```
> Set-MaintenanceWindow `
    -Path Microsoft.SQLServer.ComputerGroup `
    -Reason PlannedApplicationsMaintenance
    -Comment "System Updates" `
    -EndTime (Get-Date).AddHours(6)

> $Monitor = Get-MonitoringObject | `
    Where {$_.DisplayName -match "SQL 2005 Computers"}
> Set-MaintenanceWindow `
    -MonitoringObject $Monitor `
    -Reason PlannedApplicationsMaintenance
    -Comment "System Updates" `
    -EndTime (Get-Date).AddHours(6)
```

Monitoring Cmdlets

These cmdlets are used to work with monitoring objects in Operations Manager. We can retrieve monitoring objects (which can be used as parameters in other cmdlets), search the monitoring classes for specific monitoring objects, and retrieve overrides assigned to the monitoring objects. We can also examine the properties of a specified monitoring object and retrieve the management path to the monitoring object.

GET-MONITORINGOBJECT

This cmdlet is used to retrieve a set of specified Operations Manager monitoring objects. Monitoring objects are used in many cmdlet operations, as you will see in some of the included examples.

Parameters for this cmdlet include:

- Path** Specifies the path (or paths, separated by commas) to the monitoring objects to retrieve.
- Id** Specifies the GUID of the monitoring object to retrieve.
- MonitoringClass** Specifies that only monitoring objects of a specified monitoring class should be returned. Monitoring class objects can be retrieved through the `Get-MonitoringClass` cmdlet.
- Criteria** Used to narrow down the returned set of monitoring objects. A `Where-Object` cmdlet can also be used for this function.

Here are some examples:

```
> Get-MonitoringObject `
    | Select Displayname, HealthState
```

```
> Get-MonitoringObject `
  | Where {$_HealthState -eq "Error"}
```

GET-MONITORINGOBJECTPATH

This cmdlet is used to retrieve the path to a specified Operations Manager monitoring object. This path can be used as a parameter for other cmdlets where a `-Path` parameter is available.

Parameters for this cmdlet include:

- MonitoringObject** Specifies the monitoring object to retrieve the management path for. This object can be retrieved using the `Get-Monitoring` cmdlet.
- Id** Specifies the GUID of the monitoring object to retrieve.
- DriveQualified** Specifies that a drive-qualified path should be returned for the monitoring object.
- Path** Specifies the path (or paths, separated by commas) to the monitoring objects to retrieve.

Here is an example:

```
> Get-MonitoringObject `
  | Where {$_HealthState -eq "Error"} `
  | Get-MonitoringObjectPath
```

GET-MONITORINGOBJECTPROPERTY

This cmdlet is used to retrieve the properties of a specified Operations Manager monitoring object. The only parameter for this cmdlet is `-MonitoringObject`, which specifies the monitoring object to retrieve the properties of. This object can be retrieved using the `Get-Monitoring` cmdlet.

Here is an example:

```
> Get-MonitoringObject `
  | Where {$_HealthState -eq "Error"} `
  | Get-MonitoringObjectProperty `
  | Format-List ParentElement, Name, Value
```

GET-MONITORINGCLASS

This cmdlet is used to retrieve a set of specified Operations Manager monitoring classes. Monitoring classes are used to group or categorize various Operations Manager monitoring objects. We can use these classes to narrow down requests for specific monitoring objects to be used with other Operations Manager cmdlets.

Parameters for this cmdlet include:

- Path** Specifies the path (or paths, separated by commas) to the monitoring objects from which to retrieve monitoring classes.
- Id** Specifies the GUID of the monitoring class to retrieve.
- Name** Specifies the name of the monitoring class to retrieve.

- MonitoringObject** Specifies a particular monitoring object to return the monitoring class.
- ManagementPack** Specifies a management pack from which to return a list of monitoring classes.

Here are some examples:

```
> Get-MonitoringClass
> Get-MonitoringClass | `
    Select Name, Displayname, Description
> $Class = Get-ManagementClass `
    -name "Microsoft.SystemCenter.AllComputerGroup"
```

GET-OVERRIDE

This cmdlet is used to retrieve the overrides provided by either an Operations Manager management pack or a monitoring task.

Parameters for this cmdlet include:

- ManagementPack** Specifies the management pack from which to list available overrides. This parameter can accept only a single management pack object.
- Task** Specifies the monitoring task from which to list available overrides. This parameter can accept only a single monitoring task object.
- Criteria** Used to narrow the result set of overrides. A Where-Object cmdlet can also be used for this function.
- Path** Specifies the path (or paths, separated by commas) to the monitoring objects from which to retrieve overrides.

Here are some examples:

```
> Get-ManagementPack | Get-Override | `
    Where {$_.Enforced -eq $true}
> Get-Task | Get-Override
```

Performance Counter Cmdlets

Performance counter cmdlets are used to reveal performance data on Operations Manager clients. The process usually begins by retrieving a list of the performance counter objects supported by a particular monitoring object. With that information, we can retrieve the performance counter values recorded by the Operations Manager agent for the monitoring object.

GET-PERFORMANCECOUNTER

This cmdlet is used to retrieve the specified Operations Manager performance data items available on the current management server. The data items may be specified by their ID or by matching a specified criterion. The first example that follows will list the available performance counter objects on the current management server.

Parameters used by this cmdlet include:

- Path** Specifies the path (or paths, separated by commas) to the monitoring objects from which to retrieve performance counters.
- Id** Specifies the GUID of the performance counter to retrieve.

-Criteria Used to specify a specific set of performance counter objects to retrieve. A Where-Object cmdlet can also be used to narrow down the result set.

Here are some examples:

```
> get-performancecounter
    | select ObjectName, CounterName, InstanceName
    | sort ObjectName,CounterName

> get-performancecounter
    | where {$_.ObjectName -match "Memory"}
    | select ObjectName, CounterName, InstanceName
```

GET-PERFORMANCECOUNTERVALUE

This cmdlet is used to retrieve the values for a Operations Manager performance data item during the specified time interval. This cmdlet can be used in a standalone command or in conjunction with the Get-PerformanceCounter cmdlet.

The parameters **-StartTime** and **-EndTime** are used with this cmdlet to narrow the range of performance data retrieved from a specified performance counter. While these parameters accept a `DateTime` object, we can use a string value and let the framework do the conversion for us.

Parameters used by this cmdlet include:

-PerformanceCounter Specifies which performance counter object to view data from. This parameter accepts only one performance counter object at a time.

-StartTime Specifies the beginning of the date and time range of the desired performance counter data. This can be either a string value or a `DateTime` object.

-EndTime Specifies the ending of the date and time range of the desired performance counter data. This can be either a string or a `DateTime` object.

Here are some examples:

```
> get-performancecounter
    | where {$_.ObjectName -match "Memory"}
    | get-performancecountervalue
      -StartTime: "12/15/06" -EndTime: "12/16/06"
```

Static dates are well and good, but we can take advantage of the built-in date-handling abilities of PowerShell to do the work for us, as shown in these next examples:

```
> get-performancecounter
    | where {$_.ObjectName -match "Memory"}
    | get-performancecounterdata
      -StartTime: (get-date).AddMinutes(-90)
      -EndTime: (get-date)

> get-performancecounter
    | where {$_.ObjectName -match "Memory"}
    | get-performancecountervalue
      -StartTime: (get-date).AddHours(-6)
      -EndTime: (get-date).AddMinutes(-30)
```

The Bottom Line

In this long and technical chapter, we've covered a lot of PowerShell ground. You've learned how to:

Begin using the PowerShell environment We started with the basics, starting the PowerShell environment and navigating through the system. We discussed how PowerShell provides different providers for accessing not only the file system, but also the registry, WMI information, and, through the monitoring provider, Operations Manager configuration data. Through PowerShell snap-ins, the PowerShell environment can be expanded to include the functionality for managing Operations Manager.

Use PowerShell's built-in commands and utilities Out of the box, PowerShell provides a number of useful cmdlets. You learned that cmdlets follow the verb-noun naming convention to make them more intuitive and easier to remember. Understanding these basic cmdlets and their functionality is key to making the best use of the additional cmdlets provided with the Operations Manager PowerShell snap-in.

Pipe together individual components named cmdlets to produce powerful results PowerShell cmdlets do not actually return strings of text, but rather objects. By using the pipe symbol, the output of one cmdlet can be passed as the input for another cmdlet. You saw how the ability to clumb multiple cmdlets together can create a solution greater than the sum of its individual parts.

Use additional cmdlets provided with Operations Manager We looked at the additional cmdlets that become available when the Operations Manager PowerShell snap-in is applied. Though each cmdlet is tasked with a single job, we saw how the output from one cmdlet could be used to pass necessary information to another cmdlet to perform its specified task.

Leverage PowerShell to create scripts for Operations Manager The ability to pipe together different cmdlets is very useful, but sometimes it's not enough. Through PowerShell scripts, we can build a more complex solution than a single command line would give us. By using variables and looping structures, we have the ability to manage several sets of computers, groups, or other management structures available in Operations Manager.

Index

Note to the Reader: Throughout this index **boldface** page numbers indicate primary discussions of a topic. *Italicized* page numbers indicate illustrations.

A

Accept New Manual Agent Installations option, 90

Account tab, 277, 277

accounts

creating, **28–30**, 29, **202–203**

multihomed agents, 277, 277

roles, **201**

ACS. *See* Audit Collection Services (ACS)

actions and Actions pane

agents, 103, 103

approving, **391**

console, 174

rejecting, **392**

Reporting workspace, 260

retrieving, **391**, **393**

ACTIONS_USE_COMPUTER_ACCOUNT
option, 99

ACTIONSDOMAIN option, 100

ACTIONSPASSWORD option, 100

ACTIONSUSER option, 99

Active Alerts view, 176

Active Directory

with agent, 97–98

configuring, **30–32**, 31–32, **106–107**

MOMADAdmin for, 278–279, 278

Active Directory Application Mode (ADAM), 159

Active Directory Domain Services, **286**

Active Directory installation wizard, 58

Active Directory selection page, 36

ADAM (Active Directory Application Mode), 159

Add Management Group dialog box, 282, 283, 288,
288

Add New Data Source option, 262

Add-RemotelyManagedComputer cmdlet,
377–378

Add-RemotelyManagedDevice cmdlet, **378–379**

Add/Remove Templates dialog box, 233

Add SMTP Server dialog box, 67, 68

AddDays method, 369–370

AddHours method, 371

AddMinutes method, 369–370

AddMonths method, 369–370

AddSeconds method, 369–370

AddYears method, 369–370

Administration page, 307

Administrator Account page, 92, 92, 95, 107

administrator tools in ACS, **225–227**

Administrator workspace, 174

Administrators role, 186–187

ADSIEdit utility, 31–32, 32

adtadmin.exe tool, **225–227**

AdtAgent.exe tool, 220

Advanced Discovery option, 107

Advanced Error Reporting Settings folder, **233–234**

Advanced Operators role, 186–187

Advanced page

agents, 92–94, 93

LDAP queries, 81, 81

Advisec Blog, 346

AEM (Agentless Exception Monitoring), 36,
228–229

client monitoring, **229–232**, 230–231

group policies, **232–235**, 233

views, **235**

Agent Action Account option, 96

Agent Action accounts, **28–29**, 29, 96, 98

Agent Assignment And Failover Wizard, 78

Agent Failover page, 107

Agent Installation Directory option, 96

Agent Managed folder, 352

agent-managed monitoring, 110

Agent Managed option, 102

Agent Management property page, 107

Agent State Details pane, 215

Agentless Exception Monitoring (AEM), 36,
228–229

client monitoring, **229–232**, 230–231

group policies, **232–235**, 233

views, **235**

agentless monitoring, 110

AgentManagement directory, 54

AgentManager folder, 281

agents

- availability, **350–351**, 351
 - heartbeat settings
 - global, **76–77**
 - individual, 103, 103
 - installing, 277, **376–377**
 - Computer and Device Management Wizard, **90–96**, 91–96
 - manual, **97–101**, 98–99
 - prerequisites, **89–90**
 - remote, **354**
 - multihomed. *See* multihomed agents
 - Pending Management folder entries, **352**
 - prerequisites, 24
 - retrieving, **376**
 - settings
 - global, **76–77**, 77–78
 - individual, **102–106**, 102–105
 - management groups, **78–84**, 79–83
 - tasks, **182–185**, 183–185
 - troubleshooting, **350–355**, 351, 353–355
 - uninstalling, **377**
 - versions and updates, 277
 - warnings and errors, **354–355**, 354–355
- aggregate rollup monitors, **120–123**, 122–124
- Alert Criteria page, 291, 292
- Alert Resolution State dialog box, 64, 65, 87
- alerts
- cmdlets, **394–395**
 - destinations, 396
 - forwarding and responding to, **290–292**, 291–292
 - history, **396–397**
 - management group settings, **64–65**, 65–66
- monitors
- aggregate rollup, 123, 124
 - dependency rollup, 126, 126
 - SNMP probe-based unit, 139
 - Windows event unit, 128, 135
 - Windows services unit, 136
- resolving, **395**
- retrieving, **394–395**
- rules. *See* rules
- views, **177–178**
- Alerts Of Any Checked Severity group, 298
- aliases, cmdlets, 363, 365

- Allow This Agent To Act As A Proxy And Discover Managed Objects On Other Computers option, 104
- Allow This Server To Act As A Proxy And Discover Managed Objects On Other Computers option, 84
- Always Send Notifications option, 197
- AMD64 folder, 281
- Ancoris Extension Framework, 293, 349
- And Any Checked Alert Resolution State group, 298
- And Any Checked Category group, 298
- And Any Checked Priority group, 298
- AND groups, 118–119
- AnnounceFlags value, 358
- Application Error Events view, 235
- Application log, 358
- application pools, 248, 248
- Application Reporting Settings (All Or None) option, 234
- Application view, 235
- approving pending actions, 391
- Assign Agents To This Management Server In A Trusted Domain option, 107
- attributes, management packs, 114
- Audit Collection Services (ACS), 20, **205–206**
 - administrator tools, **225–227**
 - audit policy settings, **209**, **214**
 - collectors, **207**
 - deploying, **210–211**, 212–214
 - object counters, **224–225**
 - optimizing, **219–220**
 - removing, **228**
 - components, **206–208**
 - databases, **207–208**
 - deleting, **228**
 - deploying, 209–210
 - forwarders, 208
 - assigning to groups, 227
 - disabling, **228**
 - enabling, **214–219**, 215–218
 - priority, **227**
 - groups, **226–227**
 - implementing, **236**
 - monitors, 221, 222–223, **236–237**
 - security, **220–224**, 222–223
 - uninstalling, **227–228**
- Audit Collection Services Collector Setup wizard, 211, 236

audits

- ACS policy settings, 209, 214
- database prerequisites, 25
- failure, 359–360

authentication

- ACS, 210–211
- cluster groups, 44
- Authentication Mode setting, 34
- Author Scope page, 187
- authoring and Authoring page
 - management packs, 155–156
 - rules, 316
 - self-tuning thresholds, 322
- Authoring workspace, 174
- Authors role, 186, 188
- Auto Agent Assignment page, 78, 79
- Auto-Approve New Manually Installed Agents
 - option, 77, 87, 90, 350
- Auto Or Advanced page, 107
- Auto-Resolve Alerts page, 64–65, 66
- Autoincrement report option, 272
- Automatically Forward All Collected Errors To
 - Microsoft option, 230
- Automatically Manage Failover option, 84
- Automatically Resolve Alert When The Monitor
 - Returns To A Healthy State option
 - ACS, 221, 237
 - aggregate rollup monitors, 123
 - dependency rollup monitors, 126
 - self-tuning thresholds, 327
 - SNMP probe-based unit monitors, 139
 - Windows event unit monitors, 128, 135
 - Windows services unit monitors, 136
- Automatically Send Error Reports About This
 - Product To Microsoft Without Prompting
 - (Recommended) option, 72
- Availability History option, 67
- availability-management process, 8
- average thresholds, 140
- Average Time Between Event Generation counter, 225
- AVIcode tool, 348

B

- Back Up Database dialog box, 336, 336
- BackOffThreshold setting, 219
- Backup Set Will Expire After option, 336

backups, 331–332

- management packs, 337–338, 338
- Operations Manager database, 335–336, 335–337
- Registry, 333–334
- secure storage, 332–333, 333
- testing, 332
- baselines, 320, 321
- Baselining - Advanced dialog box, 325
- best-case rollup policy, 121–122, 122
- Best State Of Any Member option
 - aggregate rollup monitors, 123
 - dependency rollup monitors, 125
- bidirectional product connectors, 293
- blogs, 346
- budgeting, 7
- Build An Expression page, 128, 130, 134
- Build Event Expression page
 - NT event log-based rules, 143
 - NT event log event collection rules, 143
 - SNMP probe-based unit monitors, 138–139
- build numbers for management packs, 308
- Business Intelligence Development
 - subcomponent, 34

C

- CAB files, 232
- capacity-management process, 7
- Capacity Planners, 15
- catalogs, management pack, 346–348, 347–348
- cd command, 374
- CEIP (Custom Experience Improvement
 - Program), 71
 - AEM, 229, 230
 - root management server, 47
- CEIP page, 71, 71
- Change Initiation Review, 13
- change-management process in ITIL, 6
- Change View Scope dialog box, 191–192, 191
- Changing quadrant in MOF, 10–11, 11
- Choose A Discovery Method page, 115, 115
- Choose Members From A List page, 116, 117, 118
- classes
 - accessing, 370
 - monitoring, 406–407
- Clean-up MOM tool, 346
- client monitoring, 229–232, 230–231

- Client Monitoring Configuration Wizard, 229–232, 230–231
- Clive Eastwood blog, 346
- Cluster Administrator, 52, 52–53
- cluster resources
 - databases, 39
 - groups, 42–44, 42–45
 - for operations manager services, 48–49, 49–50
 - root management server installation on, 41–52, 42–50, 52–53
- cmdlets, 361–363
 - actions, 391–392
 - alerts, 394–395
 - computer and device, 376–381
 - discovery, 381–383
 - maintenance window, 403–405
 - management packs, 384–386
 - management servers, 400–402
 - miscellaneous, 397–400
 - monitoring, 405–407
 - notifications, 392–394
 - Operations Manager, 375–376
 - performance counters, 407–408
 - rules, 386–389
 - tasks, 389–391
 - useful, 363–372
- Collation Settings page, 34
- Collect Application Error From Windows Vista Or Later Computers option, 230
- collection rules, 143–147, 144–146
- collectors, ACS, 207
 - deploying, 210–211, 212–214
 - object counters, 224–225
 - optimizing, 219–220
 - removing, 228
- Command channel settings, 70, 70
- command-line tasks, 182–183
- Company Knowledge page, 310, 311, 313
- comparisons
 - cmdlet operators, 365–366
 - management packs, 168, 168
 - MP Studio 2007 for, 164
- Components To Install page
 - database installation, 33, 34
 - SQL Reporting Services, 241
- Computer and Device Management Wizard, 90–96, 91–96
- computer cmdlets, 376–381
- Computer Name option, 82, 94
- Computers page
 - agents, 92
 - LDAP queries, 80, 81
- computers view, 176
- conditions
 - agent discovery process, 94–95
 - LDAP queries, 81–82
- Config Service account, 29
- Config Service State directory, 54
- ConfigService.InstallLog file, 53
- configuration-management process, 6
- Configuration Managers, 15
- Configuration page, 313, 314
- Configure Alerts page
 - aggregate rollout monitors, 123, 124
 - dependency rollout monitors, 126
 - NT event log-based rules, 143
 - self-tuning thresholds, 327
 - SNMP probe-based unit monitors, 139
 - Windows event unit monitors, 128, 135
 - Windows services unit monitors, 136
- Configure Command Line Execution Settings page, 184, 315
- Configure Error Notification setting, 233–234
- Configure Error Reporting For Windows Operating Systems Older Than Windows setting, 234
- Configure Error Reporting For Windows Vista And Later Operating Systems setting, 234
- Configure Health page
 - SNMP probe-based unit monitors, 139
 - thresholds, 325, 326
 - Windows event unit monitors, 128, 130, 135
 - Windows services unit monitors, 136
- Configure phase in management pack life cycle, 157, 157
- Configure Report Server page, 245–253, 245–253
- Configure The Baseline Values Used To Determine The Thresholds dialog box, 324, 326
- Configure WMI Settings page, 145
- Connect To The Root Management Server page, 254, 255
- Connected Clients counter, 224
- connected management groups, 286
 - alerts, 290–292, 291–292
 - configuring, 287–288, 288
 - notification channels, 289–290
 - overview, 287
- connection cmdlets, 402

- Connection Manager, 264, 264
 - connectors
 - retrieving, 401–402
 - third-party monitoring solutions
 - configuring, 294–299, 295–298
 - installing, 293–294
 - consecutive samples over threshold monitors, 140
 - console
 - Details pane, 191–194, 191–195
 - Maintenance Mode, 198–201, 199
 - notification channels, 197–198, 197
 - operators, 186–190, 187–189, 191
 - overview, 173
 - prerequisites, 23
 - tasks, 182–185, 183–186
 - views, 175–182, 175, 196, 197, 204
 - workspaces, 174, 175
 - containers, 31, 31
 - converting management packs, 149–152, 304, 355–356
 - Corporate Error Reporting tool, 229
 - Correlate When The Following Happens list, 135
 - Correlated Event Raised option, 135
 - Correlated Events Configuration page, 135, 135
 - correlated Windows event unit monitors, 131–135, 131–133
 - counters
 - ACS, 224–225
 - optimizing, 328
 - Windows performance collection rules, 144
 - Crash Listener view, 235
 - Create A Management Pack wizard, 307, 307
 - Create A Membership Formula page, 118, 118
 - Create a Monitor option
 - aggregate rollup monitors, 123
 - dependency rollup monitors, 125
 - SNMP probe-based unit monitors, 138
 - Windows event unit monitors, 127, 134
 - Windows services unit monitors, 136
 - Create a New Aggregate Rollup Monitor Wizard, 123
 - Create A New Attribute option, 114
 - Create a New Database option, 211
 - Create A New Dependency Rollup Monitor Wizard, 125
 - Create a New Group option, 116
 - Create A Rule wizard, 316–317, 317
 - Create A Unit Monitor wizard, 221, 236, 322, 323
 - Create An Exclusion Rule page, 107
 - Create Attribute Wizard, 114–115, 115
 - Create File Share Page, 230, 231
 - Create Group Wizard, 116
 - Create Monitor Wizard
 - SNMP probe-based unit monitors, 138
 - Windows event unit monitors, 127, 134
 - Windows services unit monitors, 136
 - Create Rule Wizard, 142
 - Create Run As Profile Wizard, 80
 - Create Task Wizard, 182
 - Create User Role Wizard, 187, 187
 - Credential objects, 397
 - credentials, retrieving, 397
 - criteria and Criteria page
 - ACS forwarders, 216
 - alert views, 177–178
 - connector subscriptions, 298, 298
 - console searches, 193–195, 193–194
 - discovery, 381–382
 - event views, 179
 - LDAP queries, 79–84, 79–83, 382
 - performance views, 179
 - state views, 180
 - task status views, 181
 - curly braces ({} in cmdlets, 365
 - Custom Alert Fields page, 64
 - Custom Experience Improvement Program (CEIP), 71
 - Custom Setup page
 - database installation, 36, 37
 - Operations Manager 2007, 39, 39
 - Operations Manager Reporting, 254, 255
 - root management server, 47, 47–48
 - Customer Experience Improvement Program page
 - AEM, 229, 230
 - root management server, 47
 - Customize For Each Service Account option, 34
- ## D
- dashboard views, 177, 181
 - Data Grooming settings, 66–67, 67
 - Data Protection Managers, 15
 - Data Reader Account page, 256, 257
 - Data Source page, 211, 212, 236
 - Data Warehouse option, 254
 - Data Warehouse Write Account page, 256, 257
 - Database And Log File Options page
 - Operations Manager Database, 38, 62
 - Operations Manager Reporting, 256, 256

- Database Authentication page, 211, 213, 236
- Database Connection page, 249, 250–251
- Database Creation Options page, 211, 213, 236
- Database Files Location page, 38, 38
- Database Installation Options page, 211, 212, 236
- Database Queue % Full counter, 224
- Database Queue Length counter, 224
- Database Size entry, 38
- databases
 - ACS, 207–208
 - deleting, 228
 - deploying, 209–210
 - backups, 335–336, 335–337
 - installing, 32–39, 61–62
 - prerequisites, 22
 - requirements, 22
 - restoring, 340–342, 341–342
- DatabaseWriteModules.InstallLog file, 53
- DataStore, 21–22
- DateTime objects, 368–371
- Day method, 371
- DB Loader Events Inserts/Sec counter, 224
- DB Loader Principal Inserts/Sec counter, 224
- DB Loader String Inserts/Sec counter, 224
- DB Principal Cache Hit % counter, 224
- DB Request Queue Length counter, 224
- DB String Cache Hit % counter, 225
- Default Management Pack, 112
- default management server settings, 401
- Default Solution Link When No Microsoft Solution
 - Is Available option, 74
- defining self-tuning thresholds, 322–328, 322–327
- Delete option, 106
- deleting
 - ACS databases, 228
 - agents, 106
 - connector subscriptions, 299
- Delivery Settings page, 271, 271–272
- delta thresholds, 140–141
- dependencies
 - cluster groups, 42–45, 45
 - secondary management servers, 49, 50
- Dependencies page, 42–45, 45
- dependency rollout monitors, 124–126, 125–126
- Deploy Configuration Settings page, 230
- Destination Folder page, 97
- destinations, alert, 396–397
- Details pane
 - agents, 103
 - console, 191–194, 191–195
- developing management packs, 165–166, 165–166
- device cmdlets, 376–381
- Device Management option, 102
- Diagram Properties page, 178
- diagram views, 176–178
- dir command, 373–374
- directories
 - installation, 54
 - services administration, 12
- Disable-NotificationSubscription cmdlet, 393
- Disable-Rule cmdlet, 387
- disabling
 - forwarders, 228
 - notification subscriptions, 393
 - rules, 387
- disaster recovery, 342–343
- DisconnectThreshold setting, 219–220, 227
- discovered inventory views, 176
- Discovery Method page, 92, 93–94, 107
- discovery process
 - agents, 90–96, 91–96, 107
 - cmdlets, 381–383
 - criteria, 381–382
 - management packs, 114–116, 115
 - starting, 383
- Discovery Wizard
 - agents, 107
 - troubleshooting, 358–359
- Disk Parameters page, 42, 43
- Display Links To Solutions From Microsoft On
 - Error Reporting Computers option, 73
- Display Links To Surveys From Microsoft On
 - Error Reporting Computers option, 73
- Display or Override option, 311
- Display page
 - alert views, 178
 - event views, 179
 - performance views, 180
 - state views, 180
 - task status views, 181
- Display pane, console, 174
- distributed applications views, 176
- documentation
 - management pack changes, 162
 - MP Studio 2007, 163, 163

Does Not Equal operator, 118
 Does Not Match Wildcard operator, 118
 dollar signs (\$) in cmdlets, 365
 Domain Admins group, 30
 domain controllers, 57–58
 Domain Or Local Computer Account option, 98
 Domain page, 78–80, 79–80
 DOS commands, 363
 double threshold monitors, 140
 DR sites, 343
 Duration area, console, 200
 Dynamic System Initiative (DSI), 14
 Dynamically Expanding Virtual Hard Disk option, 58

E

E-mail channel settings, 67
 E-mail property page, 68
 E-mail Settings page, 251, 252
 Edit Maintenance Mode Settings option, 200
 Enable Agent Heartbeat option, 77, 104
 Enable Audit Collection option, 215, 215
 Enable Kerberos Authentication option, 44
 Enable-NotificationSubscription cmdlet, 393
 Enable-Rule cmdlet, 387–388
 Enable Undo Disks option, 57, 59–60
 enabling

- ACS forwarding, 214–219, 215–218
- notification subscriptions, 393
- rules, 387–388

 Encryption Key page, 253, 253
 Encryption Keys component, 245
 End User License Agreement (EULA)

- database installation, 33
- Operations Manager, 39–40
- root management server, 47

 endpoints, notification, 393–394
 Ends With condition

- agent discovery process, 94
- LDAP queries, 81

 Engyro company, 293, 349
 Enterprise Admins group, 30
 entity health, 355
 -eq operator, 366
 Error Collection page, 230, 231
 Error Group view, 235

Error Reporting page, 72, 72
 Error Transmission page, 73
 errors

- agents, 354–355, 354–355
- database installation reports, 35, 36
- Event Viewer, 359–360
- management packs, 356

 EULA (End User License Agreement)

- database installation, 33
- Operations Manager, 39–40, 254
- root management server, 47

 EULA page, 254
 Event Data option, 66
 Event Log, 352
 Event Log Name page

- NT event log-based rules, 143
- Windows event unit monitors, 128, 134

 event logs, 352, 354, 354
 Event Retention Schedule page, ACS, 211, 213
 Event Time In Collector In Milliseconds counter, 225
 Event Viewer

- agent warnings and errors, 354–355, 354–355
- Discovery Wizard completion, 358
- failure audits, 359–360
- management pack warnings and errors, 356
- verifying installation, 53

 event views, 177–179
 events, retrieving, 397–398
 eXc Software tools, 349
 Exclusion Criteria page, 83, 83
 ExecuteMonitoringTask method, 375
 Execution Account page, 252, 252
 Export-ManagementPack cmdlet, 385
 Export Registry File dialog box, 334, 334
 exporting

- management packs, 384
- RMS keys, 51

F

F5 monitor, 348
 failover system agents, 105
 failure audits, 359–360
 Feature Selection page, 34
 financial-management process in ITIL, 7
 Find Computers page, 107
 Find feature in console, 192

- finger-pointing resolution, 278
- First Occurrence Of A With The Configured Occurrence Of B, Or Vice Versa correlations, 131, 132
- First Occurrence Of A With The Configured Occurrences Of B In Chronological Order correlations, 131, 131
- First SNMP Probe page, 138
- foreach-object cmdlet, 364–365
- format-list cmdlet, 368
- Format page, 292, 292
- Format The Partition Using The NTFS File System (Quick) option, 56
- formats
 - date and time, 369–370
 - notifications, 292, 292
 - reports, 272, 272
- Forward Alerts From Targets Explicitly Added To The 'Approved Targets' Grid Are Approved option, 296
- forwarders, ACS, 208
 - assigning to groups, 227
 - disabling, 228
 - enabling, 214–219, 215–218
 - priority, 227
- forwarding alerts, 290–292, 291–292
- free space requirements, 22–26
- Frequency setting, 138–139

G

- gateway server prerequisites, 25
- general date-time format, 369
- General properties
 - aggregate rollup monitors, 123
 - connectors, 295
 - dependency rollup monitors, 125
 - management groups, 63–64, 64
 - reporting server, 74, 75
 - roles, 187
 - rules, 313
 - self-tuning thresholds, 324, 324
 - SNMP probe-based unit monitors, 138
 - tasks, 182–183, 183
 - Windows event unit monitors, 128, 134
 - Windows services unit monitors, 136
- Generate Alerts For This Monitor option
 - ACS, 221, 237
 - dependency rollup monitors, 126
 - Windows services unit monitors, 136
- generic reports, 260
- Get-Agent cmdlet, 376
- Get-AgentPendingAction cmdlet, 391
- Get-Alert cmdlet, 394–395
- Get-AlertDestination cmdlet, 396
- Get-AlertHistory cmdlet, 396–397
- get-alias cmdlet, 363
- get-command cmdlet, 362, 364–365
- Get-Connector cmdlet, 401–402
- Get-Credential cmdlet, 397
- get-date cmdlet, 368–371
- Get-DefaultSetting cmdlet, 401
- Get-Event cmdlet, 397–398
- get-help cmdlet, 362, 364
- Get-MaintenanceWindow cmdlet, 403–404
- Get-ManagementGroupConnection cmdlet, 402
- Get-ManagementPack cmdlet, 384
- Get-ManagementServer cmdlet, 377, 380–381, 400
- get-member cmdlet, 375
- Get-MonitoringClass cmdlet, 381, 406
- Get-MonitoringObject cmdlet, 405
- Get-MonitoringObjectProperty cmdlet, 406
- Get-MonitoringObjectPath cmdlet, 406
- Get-NotificationAction cmdlet, 393
- Get-NotificationEndpoint cmdlet, 393–394
- Get-NotificationRecipient cmdlet, 394
- Get-NotificationSubscription cmdlet, 392
- Get-Override cmdlet, 407
- Get-PerformanceCounter cmdlet, 407–408
- Get-PerformanceCounterValue cmdlet, 408
- get-process cmdlet, 367–368, 371–372
- Get-RemotelyManagedComputer cmdlet, 379
- Get-RemotelyManagedDevice cmdlet, 379
- Get-Rule cmdlet, 386–387
- Get-Task cmdlet, 389
- Get-TaskResult cmdlet, 390–391
- Get-UserRole cmdlet, 398
- Global Management Server Settings - Security dialog box, 350, 351
- global settings for management groups, 63–64, 64
 - agents, 77, 78
 - alerts, 64–65, 65–66
 - Data Grooming, 66–67, 67
 - notification channels, 67–70, 68–70
 - privacy, 71–74, 71–74
 - reporting, 74, 75
 - server, 76–77, 77
 - web addresses, 75–76, 75
- GPOs (Group Policy Objects), 232

- grave symbol (`) in PowerShell, 375
- grooming settings, 66–67, 67
- group policies, 232–235, 233
- Group Policy Objects (GPOs), 232
- groups
 - ACS, 226–227
 - creating, 116–120, 119–120, 152–153
 - management. *See* management groups
 - management packs for, 116–120, 119–120
 - multihomed agents, 278–279, 278–279
- Groups screen, 296, 296
- gt operator, 366
- gte operator, 366

H

- Health Rollup Policy page
 - aggregate rollup monitors, 121–123, 122
 - dependency rollup monitors, 125
- Health Service, cluster resources for, 48, 50
- Health Service Handle Count Threshold monitor, 355
- Health Service Private Bytes Threshold monitor, 355
- Health Service State directory, 54
- Health State option, 327
- Heartbeat Interval (Seconds) setting, 77, 103
- Heartbeat page, 103, 103
- heartbeat settings
 - global, 76–77
 - individual, 84, 103, 103
- Help directory, 54
- Helper Objects directory, 54
- history, alerts, 396–397
- hotfixes, 254
- Hour method, 371

I

- IA64 folder, 281
- identifying
 - prerequisites, 22–28, 26–27
 - server components, 19–22
- implementing management packs, 166–169, 167–169
- incident-management process in ITIL, 5
- Inclusion Criteria page, 79–80, 82, 83, 107
- Incoming Audits/Sec counter, 225
- Incoming Events/Sec counter, 225

- individual settings for management groups
 - agents, 78–84, 79–83, 102–106, 102–105
 - heartbeat, 84
 - proxy, 84, 85
 - security, 84
- Information Value criteria, 328
- Infrastructure cluster, 9
- Infrastructure engineering, 12
- Initialization page, 251, 251
- Install-Agent cmdlet, 376–377
- Install-ManagementPack cmdlet, 384
- Install The Program page, 47
- installation
 - agents, 277, 376–377
 - Computer and Device Management Wizard, 90–96, 91–96
 - manual, 97–101, 98–99
 - multihomed, 280–286, 280–286
 - prerequisites, 89–90
 - remote, 354
 - connectors, 293–294
 - management packs, 384
 - Operations Manager, 19
 - accounts, 28–30, 29
 - Active Directory configuration, 30–32, 31–32
 - database, 32–39, 61–62
 - domain controllers, 57–58
 - Operations Manager System, 60
 - prerequisites, 22–28, 26–27
 - process, 39–40, 39–41, 62
 - root management server on clustered hardware, 41–52, 42–50, 52–53
 - server component identification, 19–22
 - SQL Server, 58–59, 61
 - verifying, 53–54
 - Virtual Machine Base Drive, 55–57
 - Operations Manager Reporting, 253–257, 254–257
 - Reporting Services, 239–244, 240–244, 274
- Installing Prerequisites page, 240, 240
- Instance selection dialog box, 245, 245
- Instant Messaging options, 69–70, 69
- Interface Audit Insertions/Sec counter, 225
- Interface Queue Length counter, 225
- interoperating with third-party monitoring solutions, 293
 - configuring connectors, 294–299, 295–298
 - installing connectors, 293–294

Interval page, 144
 Introduction page, 107
 IP addresses for cluster groups, 43, 44
 ipconfig command, 182
 Is Equals condition, 118
 Is (Exactly) condition
 agent discovery process, 94
 LDAP queries, 81
 Is Not condition
 agent discovery process, 94
 LDAP queries, 82
 IT Infrastructure Library (ITIL), 4–8
 IT Service Management (ITSM)
 MOF, 8–13, 10–12
 service desk, 4–8
 iWave Integrator, 293, 349

J

Jakub@Work blog, 346
 job scheduling, 11

K

Kerberos authentication
 ACS, 210
 cluster groups, 44

L

Last Occurrence Of A With The Configured
 Occurrence Of B, Or Vice Versa correlations, 132,
 133
 Last Occurrence Of A With The Configured
 Occurrence Of B Happens, Enable Interval
 correlations, 133, 133
 Last Occurrence Of A With The Configured
 Occurrence Of B In Chronological Order
 correlations, 132, 132
 LDAP (Lightweight Directory Access Protocol)
 queries
 criteria, 79–84, 79–83, 382
 management packs, 114
 LDAPQueryDiscoveryCriteria objects, 382
 legacy management pack conversions, 149–152,
 304
 license agreements and Licensing Modes page
 domain controllers, 57
 Operations Manager 2007, 62

Operations Manager Database, 61
 Operations Manager System, 60
 SQL Server, 59, 61
 Virtual Machine Base Drive, 56
 life cycle of management packs, 156–157, 157
 MP Studio 2007 for, 164–165, 164
 release management, 158
 versioning, 158–162, 170–171
 Lightweight Directory Access Protocol (LDAP)
 queries
 criteria, 79–84, 79–83, 382
 management packs, 114
 -like operator, 366
 line-continuation characters, 375
 Line Properties page, 178
 listing
 rules, 385–386, 388
 user role membership, 398–399
 local database installation, 38
 Local System Account, 284
 Local System option, 96
 Locate Backup File dialog box, 341
 Locations dialog box, 353
 log-based rules, 142–143, 142
 log files, 29, 29
 agent warnings and errors in, 354, 354
 installation, 53
 multihomed agents, 278
 Log On page, 359, 359
 long date format, 369
 long time format, 369
 -lt operator, 366
 -lte operator, 366

M

maintaining management packs
 in life cycle, 157, 157
 MP Studio 2007 for, 167–169, 168–169
 Maintenance Mode, 198–201, 199
 Maintenance Mode History option, 67
 maintenance windows
 cmdlets, 403–405
 new, 404
 properties, 404–405
 retrieving, 403–404
 major revision numbers for management packs,
 160, 307

- Managed By field, 80–81
- Managed Information Base (MIB), 137
- Management Group Configuration page, 36, 97–98, 98
- Management Group Name setting, 284
- MANAGEMENT_GROUP option, 100
- management groups, 63
 - connected, 286
 - alerts, 290–292, 291–292
 - configuring, 287–288, 288
 - notification channels, 289–290
 - overview, 287
 - connections, 402
 - exercise, 86–87
 - global settings, 63–64, 64
 - agents, 77, 78
 - alerts, 64–65, 65–66
 - Data Grooming, 66–67, 67
 - notification channels, 67–70, 68–70
 - privacy, 71–74, 71–74
 - reporting, 74, 75
 - server, 76–77, 77
 - web addresses, 75–76, 75
 - individual settings
 - agents, 78–84, 79–83
 - heartbeat, 84
 - proxy, 84, 85
 - security, 84
 - multihomed agents, 278–279, 278–279
- management packs, 109
 - authoring, 155–156
 - available, 111
 - backups, 337–338, 338
 - catalogs, 346–348, 347–348
 - cmdlets, 384–386
 - considerations, 169
 - converting, 149–152, 304, 355–356
 - creating, 156
 - customizing, 306–309
 - discovery process, 114–116, 115
 - documenting changes, 162
 - exploring, 113–114
 - exporting, 384
 - for groups, 116–120, 119–120, 296
 - installing, 384
 - life cycle, 156–157, 157
 - release management, 158
 - versioning, 158–162, 170–171
 - listing rules, 385–386
 - monitors. *See* monitors and monitoring
 - MP Studio 2007 for. *See* MP Studio 2007 for management
 - packs
 - optimizing, 303
 - overrides, 148–149, 148–149
 - removing, 152
 - report configuration settings, 148
 - requirements, 110–112, 112–113
 - restoring, 339
 - retrieving, 384
 - rules. *See* rules
 - settings, 304–305, 305–306
 - troubleshooting, 355–356
 - uninstalling, 384–385
 - warnings and errors, 356
- MANAGEMENT_SERVER_AD_NAME option, 100
- MANAGEMENT_SERVER_DNS option, 100
- Management Server setting, 284
- Management Server Port setting, 284
- management servers
 - agents, 104–105
 - cmdlets, 400–402
 - default settings, 401
 - multihomed agents, 284
 - prerequisites, 23, 25–26
 - retrieving, 380–381, 400
 - setting, 381
- management tiers, 289
- ManagementGroupConnection objects, 402
- ManagementPack objects, 384
- ManagementServerConfigTool.exe program, 51
- manual multihomed agent installations, 280–286, 280–286
- manual reset unit monitors, 127
- Manually Configure Failover option, 84
- match operator, 366
- MaximumQueueLength setting, 219
- MaxPurgingQueueLength setting, 220
- Members tab, 202
- MemberType property, 375
- memory requirements, 22–25
- Merge Undo Disks option, 58
- message content for notifications, 290
- methods, PowerShell, 375
- MIB (Managed Information Base), 137

- Microsoft management packs, **161–162**
 - Microsoft.MOM.ConfigService.InstallLog file, 53
 - Microsoft.MOM.DatabaseWriteModules.InstallLog file, 53
 - Microsoft.MOM.Sdk.ServiceDataLayer.InstallLog file, 53
 - Microsoft Operations Framework (MOF), **8–13**, 10–12
 - Microsoft Operations Manager 2005 and System Center Operations Manager blog, 346
 - Microsoft Update page, 47
 - minor revision numbers for management packs, 160, 307
 - Minute method, 371
 - MOF (Microsoft Operations Framework), **8–13**, 10–12
 - MOMADAdmin utility, 31, 278–279, 278
 - MOMAgent.msi file, 99, 281
 - Monitor Dependency page, 125, 126
 - Monitor Is Enabled option, 237
 - ACS, 221
 - self-tuning thresholds, 324
 - Monitor Is In A Critical Health State option
 - ACS, 237
 - aggregate rollup monitors, 123
 - dependency rollup monitors, 126
 - SNMP probe-based unit monitors, 139
 - Windows event unit monitors, 128, 135
 - Windows services unit monitors, 136
 - Monitor Is In A Critical Or Warning Health State option
 - aggregate rollup monitors, 123
 - dependency rollup monitors, 126
 - SNMP probe-based unit monitors, 139
 - Windows event unit monitors, 128, 135
 - Windows services unit monitors, 136
 - Monitor Is In A Warning Health State option, 327
 - Monitor Target setting, 324
 - monitoring classes, **406–407**
 - Monitoring Job Data option
 - management groups, 66
 - management packs, 305, 305
 - self-tuning thresholds, 320
 - Monitoring Unavailable And Maintenance Mode area, 126
 - Monitoring workspace, 174, 191
 - monitors and monitoring
 - ACS, 221, 222–223, **236–237**
 - AEM, 36, **228–229**
 - clients, **229–232**, 230–231
 - group policies, **232–235**, 233
 - views, **235**
 - cmdlets, **405–407**
 - with console. *See* console
 - DR site, **343**
 - management packs
 - capabilities, **110**
 - requirements, **111–112**, 112–113
 - objects, **405–406**
 - rollup, **120**
 - aggregate, **120–123**, 122–124
 - dependency, **124–126**, 125–126
 - saving, **356**
 - third-party, **293**
 - configuring connectors, **294–299**, 295–298
 - installing connectors, **293–294**
 - threshold, **139–141**
 - unit, **126–127**
 - SNMP probe-based, **137–139**, 138
 - Windows event, **127–135**, 129–133, 135
 - Windows services, **136–137**, 137
- Month method, 371
 - month-year date format, 369
 - Most Common Alerts report, 258, 259–260
 - Mother blog, 346
 - MP Store feature, **164**
 - MP Studio 2007 for management packs, **162–163**
 - analysis and documentation, **163**, 163
 - comparisons, **164**
 - development, **165–166**
 - implementation, **166–167**, 167
 - life cycle, **164–165**, 164
 - maintenance, **167–169**, 168–169
 - overrides, **163–164**
 - stores, **164**
 - testing, **164**
 - MP2XML utility, 150, 304, 345
 - MPConvert command, 150, 304
 - MPSchema.xsd file, 113
 - MPVerify utility, 150
 - multihomed agents, **101–102**, **275**
 - management groups, **278–279**, 278–279
 - manual installations, **280–286**, 280–286
 - overview, **275–278**, 277
 - remote installations, **280**
 - settings, 286
 - setup, **279–280**
 - My Workspace, 174

N

Navigation pane, 174
 -ne operator, 366
 .NET Framework requirements, 23, 25–26
 network administration, 11
 Network Name Parameters page, 44, 45
 New Application Pool dialog box, 248, 248
 New Data Source View option, 265
 New-DeviceDiscoveryConfiguration cmdlet, 378, 383
 New-LdapQueryDiscoveryCriteria cmdlet, **382**
 New-MaintenanceWindow cmdlet, **404**
 New-ManagementGroupConnection cmdlet, **402**
 New Object - Group dialog box, 202
 New Object - User dialog box, 202
 New Project dialog box, 262, 263
 New-WindowsDiscoveryConfiguration cmdlet, **382–383**
 Not Present condition
 agent discovery process, 94
 LDAP queries, 82
 Notification Command Channel page, 70, 70
 Notification Devices page, 290
 Notification Recipient Properties dialog box, 197, 197
 notifications
 actions, **393**
 alerts, **290–292**, 291–292
 channels
 connected management groups, **289–290**
 settings, **67–70**, 68–70, **197–198**, 197
 endpoints, **393–394**
 subscriptions, **392–393**
 -notlike operator, 366
 -notmatch operator, 366
 Now function, 370
 NT event log-based rules, **142–143**, 142
 NT event log event collection rules, **143**
 NtpServer value, 358
 nware tools, 349

O

object identifiers (OIDs), 137, 139
 Object Properties page, 178
 Object Selection dialog box, 117, 117
 Object Types dialog box, 353
 Obtain phase in management pack life cycle, **157**, 157
 ODR Report, 258, 259
 Office Application Health folder, 305, 306
 Office applications, SCOM for, **234**
 OIDs (object identifiers), 137, 139
 OLAs (Operation-Level Agreements), 7
 OMRs (Operations Management Reviews), 10
 Only Send Notification During The Specified Times option, 197–198
 Only Targets Explicitly Added To The Approved Targets Grid Are Approved option, 187
 Only The Views Selected Below Are Approved option, 189
 Open menu option for agents, 104
 Operating quadrant, **11–12**, 11
 operating system versions requirements, 22–26
 Operation-Level Agreements (OLAs), 7
 Operational Data Reports page, 71, 72
 Operations cluster, MOF, 9
 Operations Console. *See* console
 operations management, **3**
 Dynamic System Initiative (DSI), **14**
 IT Service Management (ITSM)
 MOF, **8–13**, 10–12
 service desk, **4–8**
 System Center, **15**
 systems management, **16**
 Operations Management Reviews (OMRs), 10
 Operations Manager administrator groups, 30
 Operations Manager Product Team Blog, 346
 Operations Manager Reporting, **253–257**, 254–257, **274**
 Operations Review in Optimizing quadrant, 13
 operators
 console, **186–190**, 187–189, 191
 view creation by, **196**, 197, **204**
 Ops Manager - John Hann blog, 346
 OpsMgr1 Status screen, 60
 OpsMgrBase configuration screen, 55–56
 optimization
 ACS collectors, **219–220**
 management packs, **303**
 customizing, **306–309**
 settings, **304–305**, 305–306
 performance counters, **328**
 rules, **309**, 310
 creating, **316–319**, 317–319

- modifying, **310–314**, 311–315
 - self-tuning thresholds, **319–328**, 321–327
 - Windows performance collection rules, 145, 145
 - Optimized Performance Collection Settings dialog box, 318, 319
 - Optimizing quadrant, **12–13**, 13
 - Options page
 - database backups, 336
 - database restoring, 342
 - OR group, 118–119
 - out-file cmdlet, **372**
 - out-printer cmdlet, **371–372**
 - Override Global Agent Settings option, 103–104
 - Override Global Server Settings option, 84, 87
 - Override Properties dialog box, 312, 313
 - overrides
 - agent settings, 103–104
 - management packs, **148–149**, 148–149
 - MP Studio 2007, **163–164**
 - retrieving, **407**
 - rules, 311–312, 312–313
 - server settings, 84, 87
 - Overrides page, 311, 312
 - Overrides Summary dialog box, 311, 312
 - Overwrite All Existing Backup Sets option, 336
 - Overwrite the Existing Database option, 342
- P**
- parameters, tasks, 184, 185
 - Parent Monitor setting, 324
 - parentheses () in date and time formats, 369
 - Partner cluster, 9
 - Password Never Expires option, 58
 - passwords
 - domain controllers, 58
 - Operations Manager 2007, 62
 - SQL Report Server, 253
 - paths in monitoring objects, **406**
 - pending actions, **391–392**
 - Pending Management folder, **352**
 - performance
 - ACS, **224–225**
 - management groups, 66–67
 - performance baselines, **320**, 321
 - performance counters
 - optimizing, **328**
 - retrieving, **407–408**
 - Performance data folder, 305, 305
 - Performance Data option, 66
 - Performance Mapper page
 - probe-based performance collection rules, 147
 - WMI performance collection rules, 146
 - Performance Object, Counter, And Instance page
 - self-tuning thresholds, 324
 - Windows performance collection rules, 144, 144
 - Performance Signature option, 67
 - performance views, 177, **179–180**
 - Personalize View page, 196, 197
 - Personalize Your Software page
 - domain controllers, 57
 - Operations Manager System, 60
 - SQL Server, 59
 - Pete’s Management BLOG, 346
 - Ping utility, 285, 286
 - pipe symbol (|) in cmdlets, 362
 - pointers in cmdlets, 365
 - Possible Owners page, 44
 - PowerShell, **361–362**
 - cmdlets. *See* cmdlets
 - installing, 39–40
 - requirements, 23
 - working with, **373–375**
 - predefined reports, **260**, 261
 - Prerequisite Check Utility, **26–28**, 27
 - prerequisites
 - agents, **89–90**
 - database installation, 33
 - Operations Manager, **22–28**, 26–27
 - Present condition
 - agent discovery process, 94
 - LDAP queries, 82
 - primary management servers, **104–105**
 - printer output, **371–372**
 - priorities and Priority setting
 - ACS, 221, **227**, 237
 - aggregate rollup monitors, 123
 - alerts, 298
 - dependency rollup monitors, 126
 - NT event log-based rules, 143
 - SNMP probe-based unit monitors, 139
 - Windows event unit monitors, 128, 135
 - Windows services unit monitors, 137
 - Privacy settings for management groups, **71–74**, 71–74

- privileges for Agent Action account, 28
- proactive management, 20
- probe-based collection rules, **146–147**
- probe-based unit monitors, **137–139**, 138
- problem-management process, **5–6**
- Process model, **9–13**, 10–12
- processor requirements, 22–25
- Product Connector Subscription Wizard, 295
- Product Key page, 56
- Product Knowledge page, 313
- Product Registration page
 - database installation, 36
 - Operations Manager Reporting, 254
- Programming with System Center Operations Manager blog, 346
- Prompt the User For Approval Before Sending Error Reports To Microsoft option, 73
- properties
 - agents, 103–104, 103–104
 - connectors, 295
 - maintenance windows, **404–405**
 - management pack overrides, 148, 148
 - monitoring objects, **406**
 - performance baselines, 320, 321
 - PowerShell, 375
- proxy agents, **380**
- proxy servers, **84**, 85
- Proxy Settings page, **84**, 85
- PSCredential objects, 397

Q

- quadrants, MOF, **9–13**, 10–12
- queries
 - agent discovery process, **92–96**, 93–96
 - groups, **117–119**, 117
 - LDAP, **79–84**, 79–83, **114**, **382**
 - management packs, **114**
- Quest Software tools, 348

R

- reactive management, 20
- Read-Only Operators role, 186, 189–190
- Ready To Install page
 - agents, 98
 - SQL Reporting Services, 243, 244
- Really Simple Syndication (RSS) feeds, 347–348

- recorded web monitor problems, **356**
- redirecting printer output, **371–372**
- redundancy for notifications, 290
- Registered Queries counter, 225
- Registration Information page
 - database installation, 33
 - SQL Reporting Services, 241, 242
- Registry and Registry Editor, 333–334, 334
 - ACS collectors, 219
 - backups, **333–334**
 - Operations Manager installation, 40
 - restoring, **340**
 - Virtual Machine Base Drive, 56
 - W32Time service, 357
- Reject-AgentPendingAction cmdlet, **392**
- rejecting pending actions, **392**
- Release cluster, 10
- release management
 - ITIL, **6**
 - management packs, **158**
- Release Readiness Review, 13
- remote clients, **352–353**, 353
- Remote Desktop Users dialog box, 352–353
- remote installations for agents, **280**, **354**
- Remote page, 352, 353
- Remote Procedure Call (RPC) port, 91
- remote SQL Server, database installation on, **36–38**, 37–38
- remotely managed computers
 - adding, **377–378**
 - removing, **379**
 - retrieving, **379**
- remotely managed devices
 - adding, **378–379**
 - removing, **380**
 - retrieving, **379**
- Remove Management Group option, 282
- Remove-ManagementGroupConnection cmdlet, **402**
- Remove-RemotelyManagedComputer cmdlet, **379**
- Remove-RemotelyManagedDevice cmdlet, **380**
- removing
 - collectors, **228**
 - management group connections, **402**
 - management packs, **152**
 - members from user roles, **400**
 - remotely managed computers, **379**
 - remotely managed devices, **380**

- Repair Agents page, 105, 105
- replicating data to DR sites, 343
- Report Builder, 262
 - configuring, 262–268, 263–268
 - working with, 268–270, 268–270
- Report Operating System Errors option, 234
- Report Server Installation Instance Selection dialog box, 245, 245
- Report Server Installation Options page, 243, 244
- Report Server Virtual Directory Settings page, 246, 246
- Report Unplanned Shutdown Events option, 234
- reporting, 239
 - Operations Manager Reporting installation, 253–257, 254–257
 - prerequisites, 24
 - Reporting workspace, 174, 258, 258–260
 - creating reports, 262–273, 263–273
 - predefined reports, 260, 261
 - scheduling, 271–273, 271–273
 - settings, 74, 75, 148
 - SQL Reporting Services
 - configuring, 245–253, 245–253
 - installing, 239–244, 240–244
- Reporting Managers, 15
- Reporting Tools, 349
- resolution levels for alerts, 64, 65
- Resolution State setting, 87
- Resolve-Alert cmdlet, 395
- Resolve All Active Alerts In The New Resolution State After setting, 65
- Resolve All Active Alerts When The Alert Source Is Healthy After setting, 65
- Resolved Alerts option, 66
- resolving alerts, 65–66, 395
- resource kit tools, 346
- Resource Usage criteria, 328
- resources for multihomed agents, 278
- responding to alerts, 290–292, 291–292
- Restore Database dialog box, 340–341, 341
- restoring, 338–339
 - database, 340–342, 341–342
 - management packs, 339
 - registry, 340
 - secure storage, 339, 340
 - server installation for, 339
- Retire phase in management pack life cycle, 157, 157
- Retry Primary After setting, 68
- Return Address setting, 68–69
- Review New Manual Agent Installations In Pending Management View option, 76, 87, 350
- revision numbers for management packs, 160, 307–308
- risk management, 10
- rolesRMS (Root Management Server), 19, 41–52, 42–50, 52–53
 - accounts for, 201
 - console, 186–190, 187–189, 191
 - creating, 203–204
 - listing, 398–399
- rollup monitors, 120
 - aggregate, 120–123, 122–124
 - dependency, 124–126, 125–126
- Root Management Server (RMS), 19, 41–52, 42–50, 52–53
- RPC (Remote Procedure Call) port, 91
- RSS (Really Simple Syndication) feeds, 347–348
- Rule Is Enabled option, 317
 - NT event log event collection rules, 143
 - probe-based event collection rules, 146
 - probe-based performance collection rules, 147
 - Windows performance collection rules, 144
 - WMI performance collection rules, 145
- Rule Name And Description page, 317, 318
 - NT event log-based rules, 142
 - NT event log event collection rules, 143
 - probe-based event collection rules, 146
 - probe-based performance collection rules, 147
 - Windows performance collection rules, 144
 - WMI performance collection rules, 145
- rules, 141–142
 - cmdlets, 386–389
 - creating, 316–319, 317–319, 386
 - customized, 313–314, 314–315
 - disabling, 387
 - enabling, 387–388
 - listing, 385–386, 388
 - managing, 386–387
 - modifying, 310–314, 311–315
 - NT event log-based, 142–143, 142
 - NT event log event collection, 143
 - optimizing, 309, 310
 - probe-based event collection, 146–147
 - probe-based performance collection, 147
 - searching for, 388–389

- self-tuning thresholds, **319**
 - defining, **322–328**, 322–327
 - monitors, **141**
 - overview, **319–320**
 - performance baselines, **320**, 321
- Windows performance collection, **144–145**, 144–145
- WMI performance collection, **145–146**, 146

Rules folder, 309

Run As Account Properties dialog box, 277, 277

Run As Profile Wizard, 80

Run Task - Enable Audit Collection page, 216, 216, 218, 228

S

Saved Searches view, 194, 195

saving recorded web monitors, **356**

SC Database Server Instance page, 40, 40, 47

SC Error Reporting page, 38, 62

schedules

- notifications, 197–198, 198, 290, 291
- probe-based event collection rules, 147
- reports, **271–273**, 271–273

SCOM (System Center Operations Manager), **233–235**, 346, 349

SCOM Client Monitoring folder, 233–235

scope, console, **191–192**, 191–192

Scope Management Pack Objects By Target(s) dialog box

- aggregate rollup monitors, 122, 123
- connectors, 296
- dependency rollup monitors, 125
- self-tuning thresholds, 322, 322
- SNMP probe-based unit monitors, 137
- Windows event unit monitors, 127, 133
- Windows services unit monitors, 136

Script page, 184–185, 314, 315

scripts, 182–183, **361–362**. *See also* cmdlets

Scripts page, 147

SDK

- account, **29**
- cluster resources for, 48, 50

SDK Service State directory, 54

Sdk.ServiceDataLayer.InstallLog file, 53

SDM (System Definition Model), 14

sealed rules, **310–312**, 311–313

searches

- agent discovery process, **92–96**, 93–96
- console, **192–194**, 193–195
- rules, **388–389**
- user role membership, **399**

Second SNMP Probe page, 139

secondary management servers, **48**

SECURE_PORT option, 100

secure storage

- backing up, **332–333**, 333
- restoring, **339**, 340

Secure Vantage tools, 349

SecureStorageBackup utility, 51, 332–333, 333, 339, 340

security and security settings

- ACS, 208, **220–224**, 222–223
- agents, 104, 104
- global, **76–77**, 77
- individual, **84**, 85
- in Operating quadrant, 11–12
- set-executionpolicy cmdlet, **364**
- troubleshooting, **356–360**, 357–359

Security cluster, 10

Security page, 104, 104

Select a Data Source page, 265, 265

Select A Monitor Type page

- self-tuning thresholds, 322
- SNMP probe-based unit monitors, 138, 138
- Windows event unit monitors, 127, 129, 134
- Windows services unit monitors, 136

Select A Rule Type page

- NT event log-based rules, 142, 142
- NT event log event collection rules, 143
- probe-based event collection rules, 146
- probe-based performance collection rules, 147
- Windows performance collection rules, 144
- WMI performance collection rules, 145

Select A Target Type dialog box, 317, 318

Select Data Source View page, 265, 266

Select Event Log page, 128, 134

Select How To Define The Connection page, 262, 264, 264

select-object cmdlet, **366–367**

Select Objects To Manage page, 95, 95

Select Performance Counter dialog box, 324

Select Report Model Generation Rules page, 265, 266

Select Users dialog box, 352

Select Users Or Groups dialog box, 353

- Select Windows Service page, 136
- Selected Objects And All Their Contained Objects option, 199
- Selected Objects Only option, 199
- self-tuning thresholds, **319**
 - defining, **322–328**, 322–327
 - monitors, **141**
 - overview, **319–320**
 - performance baselines, **320**, 321
- Server Message Block (SMB) port, 91
- servers
 - component identification, **19–22**
 - management group settings, **76–77**, 77
 - for restoring environment, **339**
 - virtual, **51–52**, 52–53
- Service Account settings
 - database installation, 34
 - SQL Reporting Services, 242, 243
- Service cluster, 9
- Service Configurations directory, 54
- service connection points, 32, 32
- Service delivery category, 5
- service desk, **5–8**, 15
- Service Is Not Running Health State option, 221, 237
- Service Is Running Health State option, 221, 237
- Service Level Agreement Review, 13
- Service-Level Management (SLM) process, 7
- Service-Management Functions (SMFs), 8
- service monitoring and control, 11
- Service Name page, 136
- Service support category, ITIL, 5
- Set-AlertDestination cmdlet, **396**
- Set-DefaultSetting cmdlet, **401**
- set-executionpolicy cmdlet, **364**
- set-location cmdlet, 363
- Set-MaintenanceWindow cmdlet, **404–405**
- Set-Managementserver cmdlet, 375, **381**
- Set-ProxyAgent cmdlet, **380**
- Setup Progress page, 243, 244
- SetupOM.exe utility, 254, 280
- Severity setting
 - aggregate rollup monitors, 123
 - dependency rollup monitors, 126
 - SNMP probe-based unit monitors, 139
 - Windows event unit monitors, 128, 135
 - Windows services unit monitors, 137
- short date format, 369
- Short Message Service (SMS) messaging settings, **70**
- short time format, 369
- Silect Software tools, 349
- Simple Network Management Protocol (SNMP)
 - probe-based unit monitors, **137–139**, 138
- simple thresholds, **140**
- single-homed agents, 280
- single threshold monitors, **140**
- SMB (Server Message Block) port, 91
- SMFs (Service-Management Functions), 8
- SMS (Short Message Service) messaging settings, **70**
- SNMP (Simple Network Management Protocol)
 - probe-based unit monitors, **137–139**, 138
- Software Development Kit subcomponent, 34
- Solution Explorer windows, 262, 263, 268
- sort-object cmdlet, **367–368**
- sortable date-time format, 369
- Specify Backup dialog box, 341, 341
- Specify Management Group Information option, 97–98, 282
- SQL authentication, 210–211
- SQL Reporting Services
 - configuring, **245–253**, 245–253
 - installing, **240–244**, 240–244
- SQL Server Business Intelligence Development Studio, 262, 263
- SQL Server Connection Dialog, 249
- SQL Server creation, **58–59**, 61
- SQL Server Database Instance page, 37, 254, 255
- SQL Server Management Studio
 - database backups, **335–336**, 335–337
 - database restores, **340–342**, 341–342
- SQL Server Properties dialog box, 359, 359
- Start-Discovery cmdlet, 376–377, **383**
- Start-Task cmdlet, **389–390**
- starting tasks, **389–390**
- Starts With condition
 - agent discovery process, 94
 - LDAP queries, 81
- State Change Events Data option, 66
- state views, 177, **180**
- static threshold monitors, **139–140**
- Stop Maintenance Mode option, 200, 201
- storage
 - MP Studio 2007, **164**
 - in Operating quadrant, 12

- secure
 - backing up, **332–333**, 333
 - restoring, **339**, 340
- subgroups, **119–120**, 120
- Subscribe To A Request wizard, 271, 271
- Subscription Schedule page, 272–273, 273
- subscriptions
 - connectors, **295–299**, 295, 297–298
 - reports, 272–273, 273
- Summary page
 - ACS, 236
 - agent discovery process, 96, 96
 - agents, 107
- Support cluster, 9
- Support page, 285
- supporting quadrant, **12**, 12
- system administration in Operating quadrant, 11
- System Center, **15**
- System Center Operations Manager (SCOM), **233–235**, 346, 349
- System Configuration Check page, 33, 240, 241
- System Crash view, 235
- System Definition Model (SDM), 14
- System Preparation Tool 2.0 screen, 57
- System Properties dialog box, **352–353**, 353
- systems management, **16**

T

- Tab key in PowerShell, 375
- targeted reports, 260
- Targets screen, 296–297, 297
- Task History option, 66
- task status views, 176–177, **180–181**
- tasks, **182**
 - agent, **182–185**, 183–185
 - console, **185**, 186
 - management groups, 66
 - results, **390–391**
 - retrieving, **389**
 - starting, **389–390**
- TCP/IP Address Parameters page, 43, 44
- TCP port, 91
- Team model, 9
- templates, 233
- testing
 - backups, **332**
 - MP Studio 2007 for, **164**

- That Come From Specific Users option, 73
- third-party tools
 - monitoring solutions, **293**
 - configuring connectors, **294–299**, 295–298
 - installing connectors, **293–294**
 - for troubleshooting, **348–349**
- This Component Will Not Be Available option
 - database, 36, 61
 - Operations Manager, 62
 - root management server installation, 47
- This Component Will Not Be Installed option, 39
- three-state monitors, 320
- thresholds
 - entity health, 355
 - monitors, **139–141**
 - self-tuning, **319**
 - defining, **322–328**, 322–327
 - monitors, **141**
 - overview, **319–320**
 - performance baselines, **320**, 321
- Tidal Software tools, 348
- tiers, management, 289
- time service, **357–358**, 357–358
- Timeout option for tasks, 184
- timer reset unit monitors, 127
- timestamps, 211, 214
- troubleshooting, **345**
 - agents, **350–355**, 351, 353–355
 - Discovery Wizard, **358–359**
 - interesting tools, **349**, 350
 - management pack catalogs, **346–348**, 347–348
 - management packs, **355–356**
 - resource kit tools, **346**
 - security issues, **356–360**, 357–359
 - third-party tools, **348–349**
- Turn Off Windows Error Reporting policy, 232, 233
- two-state monitors, 320

U

- unidirectional product connectors, 293
- Uninstall-Agent cmdlet, **377**
- Uninstall-ManagementPack cmdlet, **384–385**
- Uninstall option, **106**
- uninstalling
 - ACS, 227–228
 - agents, 106, 377
 - management packs, 384–385

- unit monitors, **126–127**
 - SNMP probe-based, **137–139**, 138
 - Windows event, **127–135**, 129–133, 135
 - Windows services, **136–137**, 137
- universal date-time format, 369
- updates, agent, 277
- Upload Additional Files To Microsoft option, 73
- Upload Additional Registry Information to Microsoft option, 73
- Upload Diagnostic Data Collection Requests option, 73
- Upload Memory Dumps To Microsoft option, 73
- Upload Results Of WMI Queries To Microsoft option, 73
- Use An Existing Virtual Hard Disk option
 - domain controllers, 57
 - Operations Manager System, 60
 - SQL Server, 59
- Use Management Group Information From Active Directory option, 97
- Use Network Name For The Computer Name option, 48–49
- Use Optimization option, 145
- Use Secure Socket Layer (SSL) Protection option, 229
- USE_SETTINGS_FROM_AD option, 100
- User Must Change Password At Next Logon option, 58
- user roles
 - adding members to, **399**
 - listing members, **398–399**
 - removing members from, **400**
 - retrieving, **398**
 - searching for members, **399**

V

- variables
 - email, 68, 69
 - environment, 96, 184
 - tasks, 184
- verb-noun cmdlet format, 362–363
- Verify Backup When Finished option, 336
- verifying installation, **53–54**
- versions
 - agent, 277
 - management packs, **158–162**, **170–171**, 307
- View Properties page, 216

- views
 - AEM, **235**
 - console, **175–182**, 175
 - operators for, **196**, 197, **204**
 - tasks, 189, 190
- Views page, 189, 190
- Virtual Machine Base Drive, **55–57**
- Virtual Machine Memory setting, 58
- Virtual Network Adapter options
 - domain controllers, 57
 - Operations Manager System, 60
 - SQL Server, 59
- virtual servers, **51–52**, 52–53
- Visual Studio Tools for Office Systems, 23

W

- W32Time service, **357–358**, 357–358
- warnings
 - agents, **354–355**, 354–355
 - Event Viewer, **359–360**
 - management packs, **356**
- Web Addresses settings, **75–76**, 75
- web page views, 177, **181**
- Web Service Identity page, 246–248, 248–249
- Where-Object cmdlet, **365–366**
- Windows event unit monitors, **127–135**, 129–133, 135
- Windows Installer requirements, 24
- Windows Media Player, **234–235**
- Windows performance collection rules, **144–145**, 144–145
- Windows Service Identity page, 246, 247
- Windows services unit monitors, **136–137**, 137
- WMI performance collection rules, **145–146**, 146
- Word program requirements, 23
- Workforce management, 12
- Working Directory tasks, 184
- workspaces
 - console, **174**, 175
 - Reporting, **258**, 258–260
- worst-case rollup policy, 121–122, 122
- Worst State Of Any Member option
 - aggregate rollup monitors, 123
 - dependency rollup monitors, 125
- Worst State Of The Specified Percentage Of Members In Good Health State option, 125

X

XML Notepad 2007, 349

XML services and files

management packs, 113, 150, 304

requirements, 23–26

for troubleshooting, 349

Y

Year method, 371

Your Product Key page

domain controllers, 57

Operations Manager System, 60

SQL Server, 59

Z

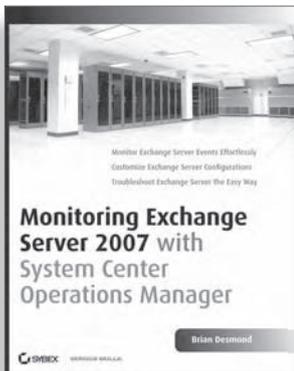
ZygortAD configuration screen, 57–58

ZygortAD Status screen, 57

ZygortSQL configuration screen, 59, 61

ZygortSQLDatabase.vhd file, 59

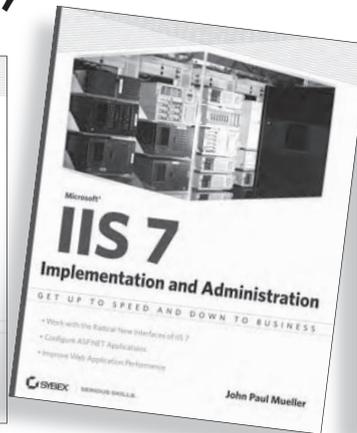
Get the job done.



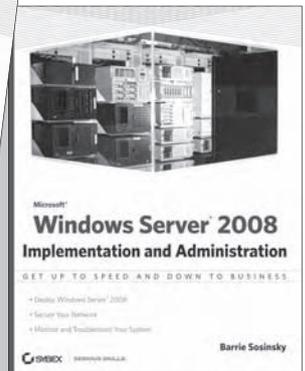
978-0-470-14895-2



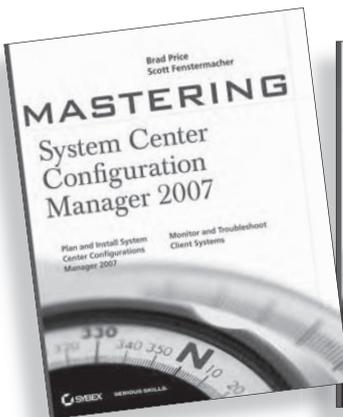
978-0-470-18741-8



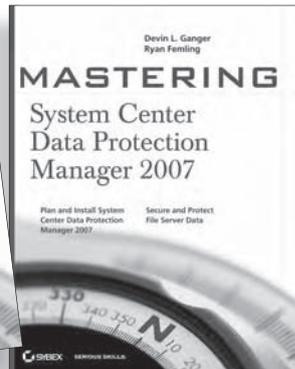
978-0-470-17893-5



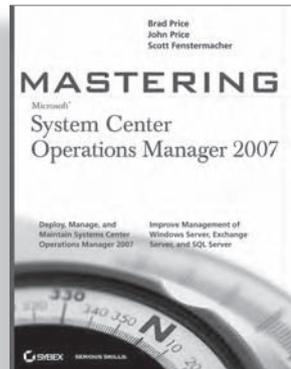
978-0-470-17459-3



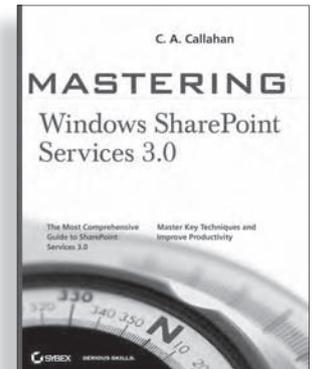
978-0-470-17367-1



978-0-470-18152-2



978-0-470-11930-3



978-0-470-12728-5

Visit www.sybex.com.

Wiley, the Wiley logo, and the Sybex logo are registered trademarks of John Wiley & Sons, Inc. and/or its affiliates. All other trademarks are the property of their respective owners.