

**Essential Microsoft Operations Manager**

By Chris Fox

.....

Publisher: O'Reilly  
Pub Date: January 2006  
Print ISBN-10: 0-596-00953-4  
Print ISBN-13: 978-0-59-600953-3  
Pages: 380

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

## Overview

For system administrators, ensuring that all Windows servers are performing optimally is a tall order. The larger the enterprise, the greater the chance for irritating, time-consuming configuration problems. Sometimes, you can determine the root cause of the problem yourself-but that's only if you're lucky.


With Microsoft Operations Manager (MOM), the diagnosis is done for you. MOM monitors server operations and automatically notifies you of problems by sending an immediate alert to your console, email address, or pager. To help you better understand how MOM works, O'Reilly presents *Essential Microsoft Operations Manager*.

The goal of this comprehensive tutorial is to give first-time MOM administrators a solid foundation for planning, implementing, and administering MOM 2005. Author Chris Fox, a renowned MOM expert, offers the type of practical, real-world advice that you need to improve the performance of your IT infrastructure.

After taking you through the entire process of setting up MOM on the network, the book moves on to more advanced administration issues. It carefully instructs you how to program and automate MOM and the agents that reside on the servers themselves. You'll also learn how to manage the scripts that determine which server agents are relevant to report.

By capturing system data, intelligently analyzing it, and then notifying you with a suggested course of action, MOM makes extinguishing fires a breeze. And now, thanks to *Essential Microsoft Operations Manager*, learning how to use MOM is a breeze, too.

PREV



Essential Microsoft Operations Manager

By Chris Fox

.....

Publisher: O'Reilly  
Pub Date: January 2006  
Print ISBN-10: 0-596-00953-4  
Print ISBN-13: 978-0-59-600953-3  
Pages: 380

Table of Contents | Index

- Copyright
- Dedication
- Preface
  - Who Should Read This Book
  - What's in This Book
  - Conventions Used in This Book
  - Using Code Examples
  - Safari® Enabled
  - We'd Like Your Feedback
  - Acknowledgments
- Part I: Introducing Operations Management and MOM 2005
  - Chapter 1. Introduction to MOM 2005
    - Section 1.1. Using MOM 2005
    - Section 1.2. Basic MOM Structure
    - Section 1.3. The Life of a MOM 2005 Alert
    - Section 1.4. MOM 2005 Components
    - Section 1.5. Additional Components
    - Section 1.6. Summary
  - Chapter 2. Designing, Planning, and Implementing MOM 2005
    - Section 2.1. Requirement Gathering
    - Section 2.2. Design Decisions
    - Section 2.3. Pre-Installation Configuration Decisions
    - Section 2.4. Testing and Piloting
    - Section 2.5. Implementation
    - Section 2.6. Installation Specifics
    - Section 2.7. Summary
- Part II: Managing and Using MOM on a Daily Basis
  - Chapter 3. Managing Agents
    - Section 3.1. Agent Functions
    - Section 3.2. Preparing to Deploy Agents
    - Section 3.3. Deploying and Managing Agents in a Trusted LAN
    - Section 3.4. Deploying and Managing Agents Across a Firewall and a Slow WAN Link

- [Section 3.5. Deploying and Managing Agents Into an Untrusted Environment](#)
- [Section 3.6. Deploying and Managing Agents from Multiple Management Groups](#)
- [Section 3.7. Troubleshooting](#)
- [Section 3.8. Tools](#)
- [Section 3.9. Summary](#)
- [Chapter 4. Administering Management Packs](#)
  - [Section 4.1. Management Pack Life Cycle](#)
  - [Section 4.2. Importing Management Packs into Preproduction](#)
  - [Section 4.3. Transfer the Management Pack to Production](#)
  - [Section 4.4. Creating Simple Management Packs](#)
  - [Section 4.5. Summary](#)
- [Chapter 5. Administering Global Settings](#)
  - [Section 5.1. Alerts](#)
  - [Section 5.2. Connections](#)
  - [Section 5.3. Maintenance](#)
  - [Section 5.4. Summary](#)
- [Chapter 6. Operator Console](#)
  - [Section 6.1. Console Scopes](#)
  - [Section 6.2. Creating a Custom Console Scope](#)
  - [Section 6.3. Creating a Custom Computer Group](#)
  - [Section 6.4. Creating the Console Scope](#)
  - [Section 6.5. Using the Console](#)
  - [Section 6.6. Building a Filter in the Operator Console](#)
  - [Section 6.7. Views](#)
  - [Section 6.8. Customizing the Operator Console](#)
  - [Section 6.9. Summary](#)
- [Chapter 7. MOM 2005 Database Fundamentals](#)
  - [Section 7.1. SQL Server Enterprise Manager](#)
  - [Section 7.2. Data Transformation Service](#)
  - [Section 7.3. MOM 2005 Reporting Databases](#)
  - [Section 7.4. Summary](#)
- [Chapter 8. MOM 2005 Reporting](#)
  - [Section 8.1. Installation](#)
  - [Section 8.2. Administering MOM 2005 Reporting](#)
  - [Section 8.3. Summary](#)
- [Part III: MOM 2005 Enterprise Integration](#)
  - [Chapter 9. Connecting MOM 2005](#)
    - [Section 9.1. Partitioning](#)
    - [Section 9.2. Connecting MOM to MOM](#)
    - [Section 9.3. Summary](#)
  - [Chapter 10. Extending Monitoring](#)
    - [Section 10.1. Understanding SNMP](#)
    - [Section 10.2. Windows and MOM Implementation of SNMP](#)
    - [Section 10.3. Configuring Windows and MOM for SNMP](#)
    - [Section 10.4. Syslog](#)
    - [Section 10.5. Summary](#)

[← About the Author](#)

[← Colophon](#)

[← Index](#)

[← PREV](#)





Essential Microsoft Operations Manager™

by Chris Fox

Copyright © 2006 O'Reilly Media, Inc. All rights reserved.

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles ([safari.oreilly.com](http://safari.oreilly.com)). For more information, contact our corporate/institutional sales department: (800) 998-9938 or [corporate@oreilly.com](mailto:corporate@oreilly.com).

Editor: John Osborn and Robbie Allen

Editor: John Osborn and Robbie Allen

Developmental Editor: Brian MacDonald

Production Editor: Laurel R.T. Ruma

Copyeditor: Laurel R.T. Ruma

Proofreader: Sada Preisch

Indexer: John Bickelhaupt

Cover Designer: Hanna Dyer

Interior Designer: David Futato

Cover Illustrator: Karen Montgomery

Illustrators: Robert Romano, Jessamyn Read, and Lesley Borash

Printing History:

February 2006:	First Edition.
----------------	----------------

Nutshell Handbook, the Nutshell Handbook logo, and the O'Reilly logo are registered trademarks of O'Reilly Media, Inc. *Essential Microsoft Operations Manager*, the image of a beaver, and related trade dress are trademarks of O'Reilly Media, Inc.

Microsoft, MSDN, the .NET logo, Visual Basic, Visual C++, Visual Studio, and Windows are registered trademarks of Microsoft Corporation.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and O'Reilly Media, Inc. was aware of

a trademark claim, the designations have been printed in caps or initial caps.

While every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

ISBN: 0-596-00953-4

[M]



 [PREV](#)

# Dedication

*To my wife Dora and our children, Mariana, Anastasia, Gabrielle, and Catalina. This book only exists because of your love, encouragement, support, and sacrifice.*

Chris Fox

 [PREV](#)

# Preface

Microsoft Operations Manager (MOM) 2005 can deliver a huge amount of value to organizations of any size. It automates burdensome and lengthy diagnostic tasks so that you are notified of an event in your environment almost as soon as it happens. Armed with this information, and a set of integrated tools engineered to help you fix whatever is wrong, a solid MOM implementation helps any IT infrastructure reduce outages and simply run better. And if your machines and applications are healthier, chances are your life will be just a bit easier.

System administrators of any network must perform operations management duties. If your network is small and relatively uncomplicated, you probably handle these tasks without the assistance of a tool like MOM. Instead, you rely on your end users to tell you about an outage or an incident with a system. More than likely, you rely on built-in tools such as event logs, performance counters, Dr. Watson logs, and application-specific logs to provide the data you need to diagnose issues with your systems. In addition, you perform tasks such as pinging the IP address of a device to see if it is up on the network. You run diagnostic tools to get more in-depth diagnostic information, such as DCDIAG and REPLMON for domain controller issues. You rely on your experience, your knowledge of the systems involved, and external support coming from online research or a phone call to a support engineer to determine a course of action to fix issues and restore service.

Sometimes, issues resolve themselves, or they arise intermittently evaporating before you can capture the data needed to diagnose them. If you are lucky and can determine the root cause of the incident, you will need a good deal of self-discipline to record the facts of the issue and the steps you took to resolve it for future reference. Then, during the budgetary cycle when you have to justify why new server hardware is needed, or why a different backup solution is appropriate, you will probably scramble to find that supporting documentation.

Ultimately, you would probably describe your workday as being interruption driven you spend most of your time putting out fires. You never have time for systems design and new implementations. You are always in a reactive mode.

An operations management system can help you with many of these duties, not by eliminating them no product can ever do that but by capturing system data, intelligently analyzing it, correlating it with data from other systems, and notifying you about an issue with actionable information. The operations management system can then help you perform further diagnostic tests and suggest courses of action based on vendor knowledge. In addition, it can give you a place to capture the resolution steps that are associated with the issue, so that the next time it happens your troubleshooting information is already there for quick resolution. It can help you track system performance and events over a long period of time for reporting purposes. Probably the greatest help that an operations management system can give you is that it, unlike you, can be in many different places doing many different things at the same time. And through the implementation of such a system, your organization can gain operational awareness, which is knowledge about the state of your IT systems at any given point in time. Once an organization has this level of insight into the health of its distributed systems, it can start preventing issues before they become outages. This allows an IT organization to shift from being primarily in a reactive mode, to being primarily in a proactive mode.



MOM 2005 takes a unique approach, developing your operational awareness by providing you with information that focuses on the *state* of an application or server, based on the condition of its components. The current state of a monitored computer or application is judged against a definition of its health *as defined by the vendor that produced it*. Health definitions, otherwise known as *management packs*, are developed by the application product teams and contain their distilled diagnostic and troubleshooting knowledge. Using MOM gives you the benefit of their experience. This combined with MOM's ability to monitor all the servers that play a role in delivering an application service (such as email or Active Directory (AD)), can let you know how healthy your whole environment is at a single glance.





 PREY

# Who Should Read This Book

When you transition into a new job or to new duties in your current job, you naturally first want to learn only what you *must know* to do whatever it is you have to do. How you find out what you *must know* is an experience who's pain lies somewhere between that of a root canal and sore muscles from a hard workout. The goal of this book is to give you what you *must know* to have a solid foundation for planning, implementing, and administering MOM 2005, which is a pretty good workout.

I wrote this book with the first-time MOM administrator in minde. To get the most from this book, it is helpful to have some background as a Windows system administrator in a multiple-server environment. If you are a complete novice to the Windows OS and are not even sure what event logs or performance counters are, then you are going to struggle a bit with this material until you get some basic Windows OS administration under your belt.

 PREY

# What's in This Book

This book is divided into three parts.

[Part I](#), "[Introducing Operations Management and MOM 2005](#)," brings you up to speed on the basic concepts of operations management and how to get MOM 2005 up and running in your environment.

- [Chapter 1](#), *Introduction to MOM 2005*, introduces the core components of MOM 2005, describes how an alert is generated, and explains how to use the information in an alert to troubleshoot an issue. The chapter introduces the Operator console and some of the tools and views that are used during the troubleshooting process. It also provides an overview of all the components in a MOM management group and what they do.
- [Chapter 2](#), *Designing, Planning, and Implementing MOM 2005* addresses the critical design and implementation decisions that must be made when starting a MOM 2005 rollout. Like [Chapter 1](#), which introduces the components of MOM via the processing of an alert, this chapter takes the reader through the business decision points of a rollout via the design and implementation efforts of a fictitious company Leaky Faucet.

[Part II](#), "[Managing and Using MOM on a Daily Basis](#)," takes you through the components and tasks that you'll be working with most often as an administrator.

- [Chapter 3](#), *Managing Agents*, provides detailed information on how agents work, agent components, and types of agents. It covers how to target computers for agent installation, and the installation and uninstallation process. Basic troubleshooting tips and tools are discussed as well.
- [Chapter 4](#), *Administering Management Packs*, lays out a framework for controlling the management pack life cycle. It covers importation into preproduction, tuning, importation into production, and evolution and synchronization of preproduction and production. It introduces rule types, overrides, notifications, and responses, as well as backup and restore of management packs. [Chapter 4](#) closes by showing you how to create a rudimentary management pack using the management pack creation wizard.
- [Chapter 5](#), *Administering Global Settings*, explains the configurations made in the global settings node of the Administrator console, plus their purpose and effect on the behavior of the whole management group.
- [Chapter 6](#), *Operator Console*, provides practical instruction on how to navigate the Operator console, what its components are, and how and when to use them. It also covers how to create new views and console scopes.
- [Chapter 7](#), *MOM 2005 Database Fundamentals*, describes the operations and reporting databases and their inter-relationship. The basic use of the SQL Enterprise Manager and SQL Query Analyzer tools is covered, paying particular attention to database grooming, backup, and restore.

- [Chapter 8](#), *MOM 2005 Reporting*, takes you through the setup of MOM 2005 Reporting, including the installation of SQL Server 2000 Reporting Services. It explains what a Data Transformation Service (DTS) package is and how the operational data is extracted from the production database and moved to the reporting database. It provides details on how to administer Reporting, such as creating custom reports and configuring subscriptions and backup procedures.

[Part III](#), "[MOM 2005 Enterprise Integration](#)," discusses the specific issues you'll encounter when using MOM in a larger environment, with a variety of platforms.

- [Chapter 9](#), *Connecting MOM 2005*, focuses on those situations when it is necessary to implement multiple management groups and explains how to configure inter-management group communication (MOM-to-MOM Product Connector) so that operational information from multiple management groups is made available for viewing and resolution in a single Operator console.
- [Chapter 10](#), *Extending Monitoring*, shows how you can use MOM and the Windows OS to collect and analyze data from non-Windows platforms such as Unix machines (via syslogs) and network switches and routers (via SNMP).



# Conventions Used in This Book

The following typographical conventions are used in this book:

`Constant width`

Indicates command-line elements, computer output, and code examples.

*Constant width italic*

Indicates placeholders (for which you substitute an actual value) in examples and in registry keys.

**Constant width bold**

Indicates user input.

*Italic*

Introduces new terms and URLs, commands, file extensions, filenames, directory or folder names, and UNC pathnames.

Indicates a tip, suggestion, or general note. For example, we'll tell you if you need to use a particular version or if an operation requires certain privileges.

Indicates a warning or caution. For example, we'll tell you if MOM does not behave as you'd expect or if a particular operation has a negative impact on performance.



## Using Code Examples

This book is here to help you get your job done. In general, you may use the code in this book in your programs and documentation. You do not need to contact us for permission unless you're reproducing a significant portion of the code. For example, writing a program that uses several chunks of code from this book does not require permission. Selling or distributing a CD-ROM of examples from O'Reilly books *does* require permission. Answering a question by citing this book and quoting example code does not require permission. Incorporating a significant amount of example code from this book into your product's documentation *does* require permission.

We appreciate, but do not require, attribution. An attribution usually includes the title, author, publisher, and ISBN. For example: "*Essential Microsoft Operations Manager*, by Chris Fox. Copyright 2006 O'Reilly Media, Inc., 0-596-00953-4."

If you feel your use of code examples falls outside fair use or the permission given above, feel free to contact us at [permissions@oreilly.com](mailto:permissions@oreilly.com).





# Safari® Enabled



When you see a Safari® Enabled icon on the cover of your favorite technology book, that means the book is available online through the O'Reilly Network Safari Bookshelf.

Safari offers a solution that's better than e-books. It's a virtual library that lets you easily search thousands of top tech books, cut and paste code samples, download chapters, and find quick answer: when you need the most accurate, current information. Try it for free at <http://safari.oreilly.com>.





## We'd Like Your Feedback

The information in this book has been tested and verified to the best of our ability, but mistakes and oversights do occur. Please let us know about errors you may find, as well as your suggestions for future editions, by writing to:

O'Reilly Media, Inc.  
1005 Gravenstein Highway North  
Sebastopol, CA 95472  
(800) 998-9938 (in the U.S. or Canada)  
(707) 829-0515 (international or local)  
(707) 829-0104 (fax)

You can also email us. To be put on our mailing list or to request a catalog, send email to:

[info@oreilly.com](mailto:info@oreilly.com)

To ask technical questions or comment on the book, send email to:

[bookquestions@oreilly.com](mailto:bookquestions@oreilly.com)

For corrections and amplifications to this book, check out O'Reilly Media's online catalog at:

<http://www.oreilly.com/catalog/microsoftopmgr>



# Acknowledgments

I didn't write this book alone, although I claim all responsibility for any inaccuracies or outright mistakes. It started as an idea that I discussed with John Osborn, an executive acquisitions editor at O'Reilly, at TechEd 2004 in San Diego. John was immediately receptive, and his encouragement and involvement throughout the process is greatly appreciated. My thanks to Robbie Allen for taking the book proposal, getting it approved, providing guidance, and being the technical editor for this book.

Brian MacDonald, what can I say? You naturally speak the way I wish I could write. I can't thank you enough for helping me to gain a small amount of skill in translating disjointed thoughts into coherent sentences, then into paragraphs and eventually into chapters. You showed me the path.

The idea for this book started at a Microsoft Management Summit (2003, I think) where I was able to meet and talk to some incredibly talented people from the MOM product team at Microsoft. Thanks to Ashvin Sanghvi, Shawn Bice, Vlad Joanovic, and Brian Wren. A little later on, I ran into Baelson Duque, MOM Program Manager, who has to be the hardest-working guy at Microsoft. Baelson volunteered to be the sole technical contact on the MOM team. Why he signed on for battle in the trenches, I'll never know, but he was available for endless questions via all modes of communication that have been created, and the book is immeasurably better for his involvement. Thank you, Baelson!

Thanks to the technical reviewers, Stuart Renes, John Hann, James Hedrick, and Doug Bradley, for their contributions, comments, and corrections. Special thanks to Jamie Peebles at Microsoft, who took on the technical reviewing tasks and went through the whole book in about two weeks. I still think my wizard would own your dwarf in Dungeon Siege. You have been a constant source of support and trusted advice.

My deepest gratitude to the many other truly wonderful people at Microsoft, including Randy Young at Microsoft Consulting Services, who guided me in my earliest use of MOM; Don Bryner and Scott Brown, for ongoing moral support; Mike Kellogg, for answering questions and trying (many times) to get permission; and Suzanne Sylliaasen, Paul Cholak, and Chris Furlin, for discussions. Special appreciation to Bill Anderson, who helped set me on this road years ago when he guided me into Rainier Technology and who always gives trusted advice.

Thank you to my parents, Brian and Marjorie, for their support and encouragement and for making sure I was well educated. Thank you to Jesse Liberty, who doesn't know how much he helped me when I needed inspiration on how to turn a phrase.

My gratitude to Jim Minatel, who gave me my first entry into the world of published writing and Elise Peterson, who made introductions. Thanks to Jon Shrier, who taught me the ropes of operations management, and to all of my friends and colleagues who have offered best wishes, encouragement, and advice along the way.

Thank you to my family for putting up with my spending so much time in the basement while working on this project. Mariana, you read the proposal and the earliest chapters when I needed an English major's eye. Ana, you were my early-morning companion. Cate, you would always let me know when it was time for bed (and when you needed help putting the lid on your cup). Gabby, I would have

been late for many dinners and events if you had not invented the Vent-O-Phone. The vacation is coming. Thank you to Max, Sammy, and Molly, companions and guardians through endless hours.

*Mi esposa, te amo para siempre* The only reason that I was able to do this is because of you. I thank God for you and the girls and all the blessings that He has given me. I thank Him for putting me here, now with all the support in the world.

Now, it is time to fix those light fixtures and that darn Leaky Faucet.

 **PREV**

 PREY

# Part I: Introducing Operations Management and MOM 2005

[Chapter 1, Introduction to MOM 2005](#)

[Chapter 2, Designing, Planning, and Implementing MOM 2005](#)

 PREY



# Chapter 1. Introduction to MOM 2005

As you manage your Microsoft-based infrastructure, your goal is to develop a higher level of operational awareness concerning your unique IT environment. This awareness can be summed up in a simple statement: "I know what is going on in my IT environment right now." When you can say this with a high degree of confidence, you have arrived at your goal.

You can reach this goal by performing *operations management*. Operations management is not system administration. System administration, and the skill sets that go along with it, are used in operations management, but system administration is narrower in focus in that it applies to a single system or platform. Operations management has a broader scope. It looks at how multiple systems work together to provide IT services to a company. It involves troubleshooting, managing, and reporting on all those systems as a whole.

You can perform operations management manually by examining Windows event logs, gathering performance monitor data, and depending on your users to tell you that something has gone awry. This can be very time-consuming, even if you have a small number of machines; ultimately, this approach leaves you reacting to events rather than preventing them. Microsoft Operations Manager 2005 (MOM 2005 or just MOM) performs many of these tasks for you and generates an alert when it detects a malfunction in the monitored applications or a condition that can lead to a malfunction. In the alert, MOM 2005 tells you what is wrong, what the most common causes of the problem are, and the likely initial steps to fix the issue. Whether these issues have a large impact or are unnoticeable to the end user, resolving them quickly has two effects. If the impact is large, say an Exchange server mailbox store is down, then quick resolution is the obvious goal, with obvious benefits. If the issue is small, for example a missed Active Directory (AD) replication cycle, then fixing it now often prevents a small issue from causing larger issues that will continue to snowball until you have a major outage on your hands. Being able to fix small issues promptly because you know what is going on in your environment lets you be proactive. This is the biggest benefit of successful operations management and is why developing your operational awareness is critical to that success.

The goal of this book is to teach you how to use MOM 2005 correctly so that you can raise your level of operational awareness for your environment. Every IT environment will have unique monitoring, alerting, and reporting requirements; there is no cookie-cutter implementation for MOM. The key to using MOM correctly for your environment is to understand how MOM works and how to use it. Throughout this book, I will show you ways to implement and use MOM in different environments. You can then plan and use MOM to the greatest effect in your environment.

## 1.1. Using MOM 2005

Although the design and implementation of MOM 2005 will be unique to every environment, the steps you follow to make use of MOM will be the same:

1. Plan and design. Essentially, these tasks involve taking inventory of your environment and determining what your business and technical needs are. This has to be the starting point of your MOM deployment, and it is where the uniqueness of your environment is baked into your MOM implementation. Here you create a design that will meet your business needs.
2. Install MOM 2005. Every installation of MOM 2005 has the same basic components, just as every car has an engine, tires, a seat for the driver, and controls that the driver uses. But the placement and configuration of these parts varies from car to car. So, too, will the location and configuration of MOM parts vary from installation to installation. And they have to vary to meet the unique business needs of the environments they are being installed into.
3. Deploy agents . This is the first task that you will do once you have MOM 2005 installed. Agents are deployed to machines that MOM 2005 goes out to and discovers. Discovery is performed according to rules that you configure. The most common rule tells MOM to discover all machines in a domain. Deploying agents onto machines that host applications you want to monitor lets MOM 2005 do one thing that you cannot be in many places at the same time. Agents monitor applications on servers, as well as the servers themselves, and compare the collected data to sets of customizable health rules defined by the application vendors. When an exception is found, an alert is generated. The deployment and management of agents is covered in detail in [Chapter 3](#).
4. Monitor your environment and fix what is wrong. MOM 2005 will tell you things about your environment that you are completely unaware of. Some of these things will be informational in nature, and others will require immediate action for resolution.
5. Tweak MOM. Not every health rule that MOM 2005 provides out-of-the-box will be useful or even appropriate for your environment. If MOM is alerting you to issues that you don't want to be alerted about, turn the rule off or adjust the thresholds in the rule so that MOM is giving you relevant, actionable information.

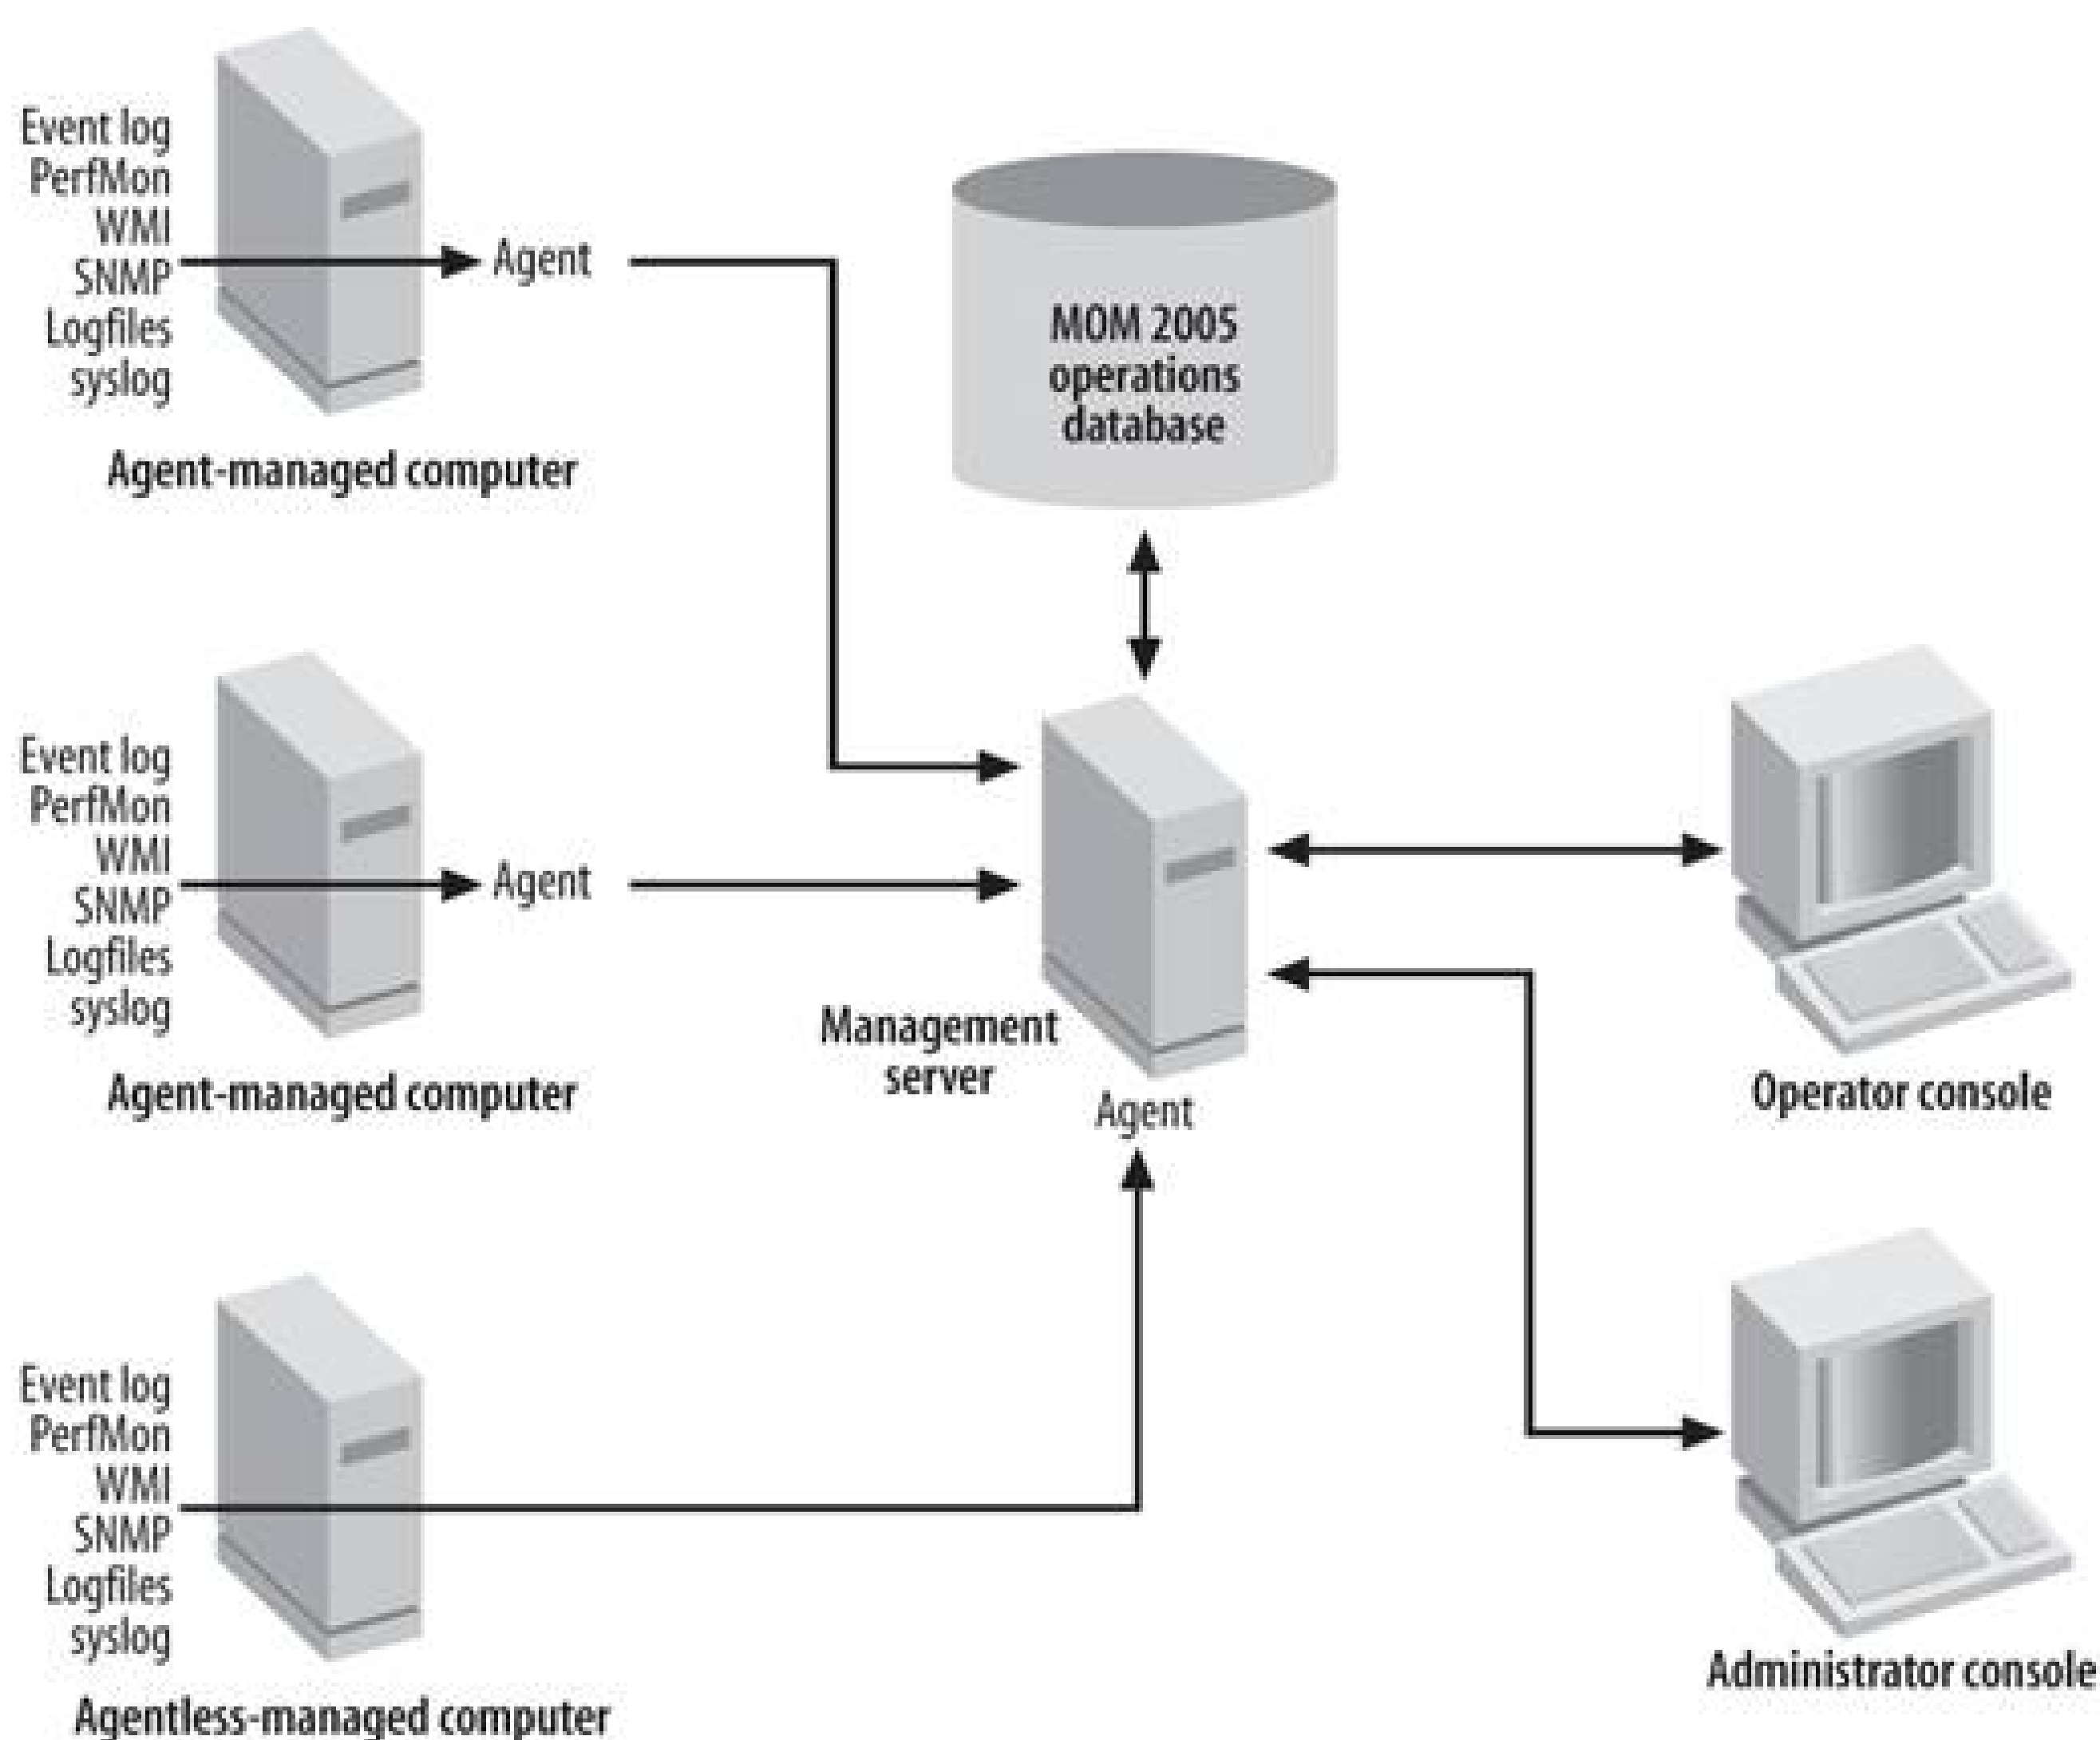
Of course these are not all the things you can do with MOM, but you will do at least these things, no matter what. In the last section of this chapter, additional services of MOM 2005, such as its reporting, will be covered.



## 1.2. Basic MOM Structure

The functional unit of a MOM 2005 implementation is a management group. All management groups consist of an operations database, one or more management servers, one or more managed computers, the Administrator console, and the Operator console. [Figure 1-1](#) shows the minimum components of a MOM 2005 management group.

Figure 1-1. A basic MOM 2005 management group



The operations database maintains all the settings that define the configuration of the management group. It also stores the live or operational data that has been collected by the agents until it is groomed out by automated SQL jobs when the data is no longer useful or to free-up database space. The database can be Microsoft SQL Server 2000 with SP3A or MSDE the desktop version of the SQL database for smaller test installations.

The management server performs all the centralized tasks for the management group. It deploys and manages agents, and it proxies communication between the database, the agents, and the user consoles. In addition, all configuration tasks for the management group are performed on the management server, including rule management.

Managed computers are machines that are monitored by the management group and most of them will have agents installed. But MOM 2005 can also monitor a limited number of computers in an *agentless* fashion, which means that the monitoring of a remote computer is performed by the agent on the management server. This means an increase in the load on the management server and on the network between the two machines. As such, Microsoft has limited the number of agentless-managed machines to 60 per management group.

You will not need to make modifications to these components on a daily basis. In fact, with proper planning these components are set up during installation and not modified again until you dismantle the management group.

The specific health rules and thresholds will be adjusted during the tweaking period of your MOM 2005 deployment; but again, once they are producing the type of information that you want, you should leave them alone.

 **PREV**

## 1.3. The Life of a MOM 2005 Alert

Now for the life cycle of a MOM 2005 alert. This explanation shows how MOM 2005 processes information, introduces you to the Operator console, and discusses how MOM 2005 is used on a daily basis to manage your environment.

A MOM 2005 alert tells you when something significant happened somewhere in one of your systems. Not all alerts are created equal; they come in different levels of severity, from the benign informational and success alerts to urgent service unavailable and critical error alerts.

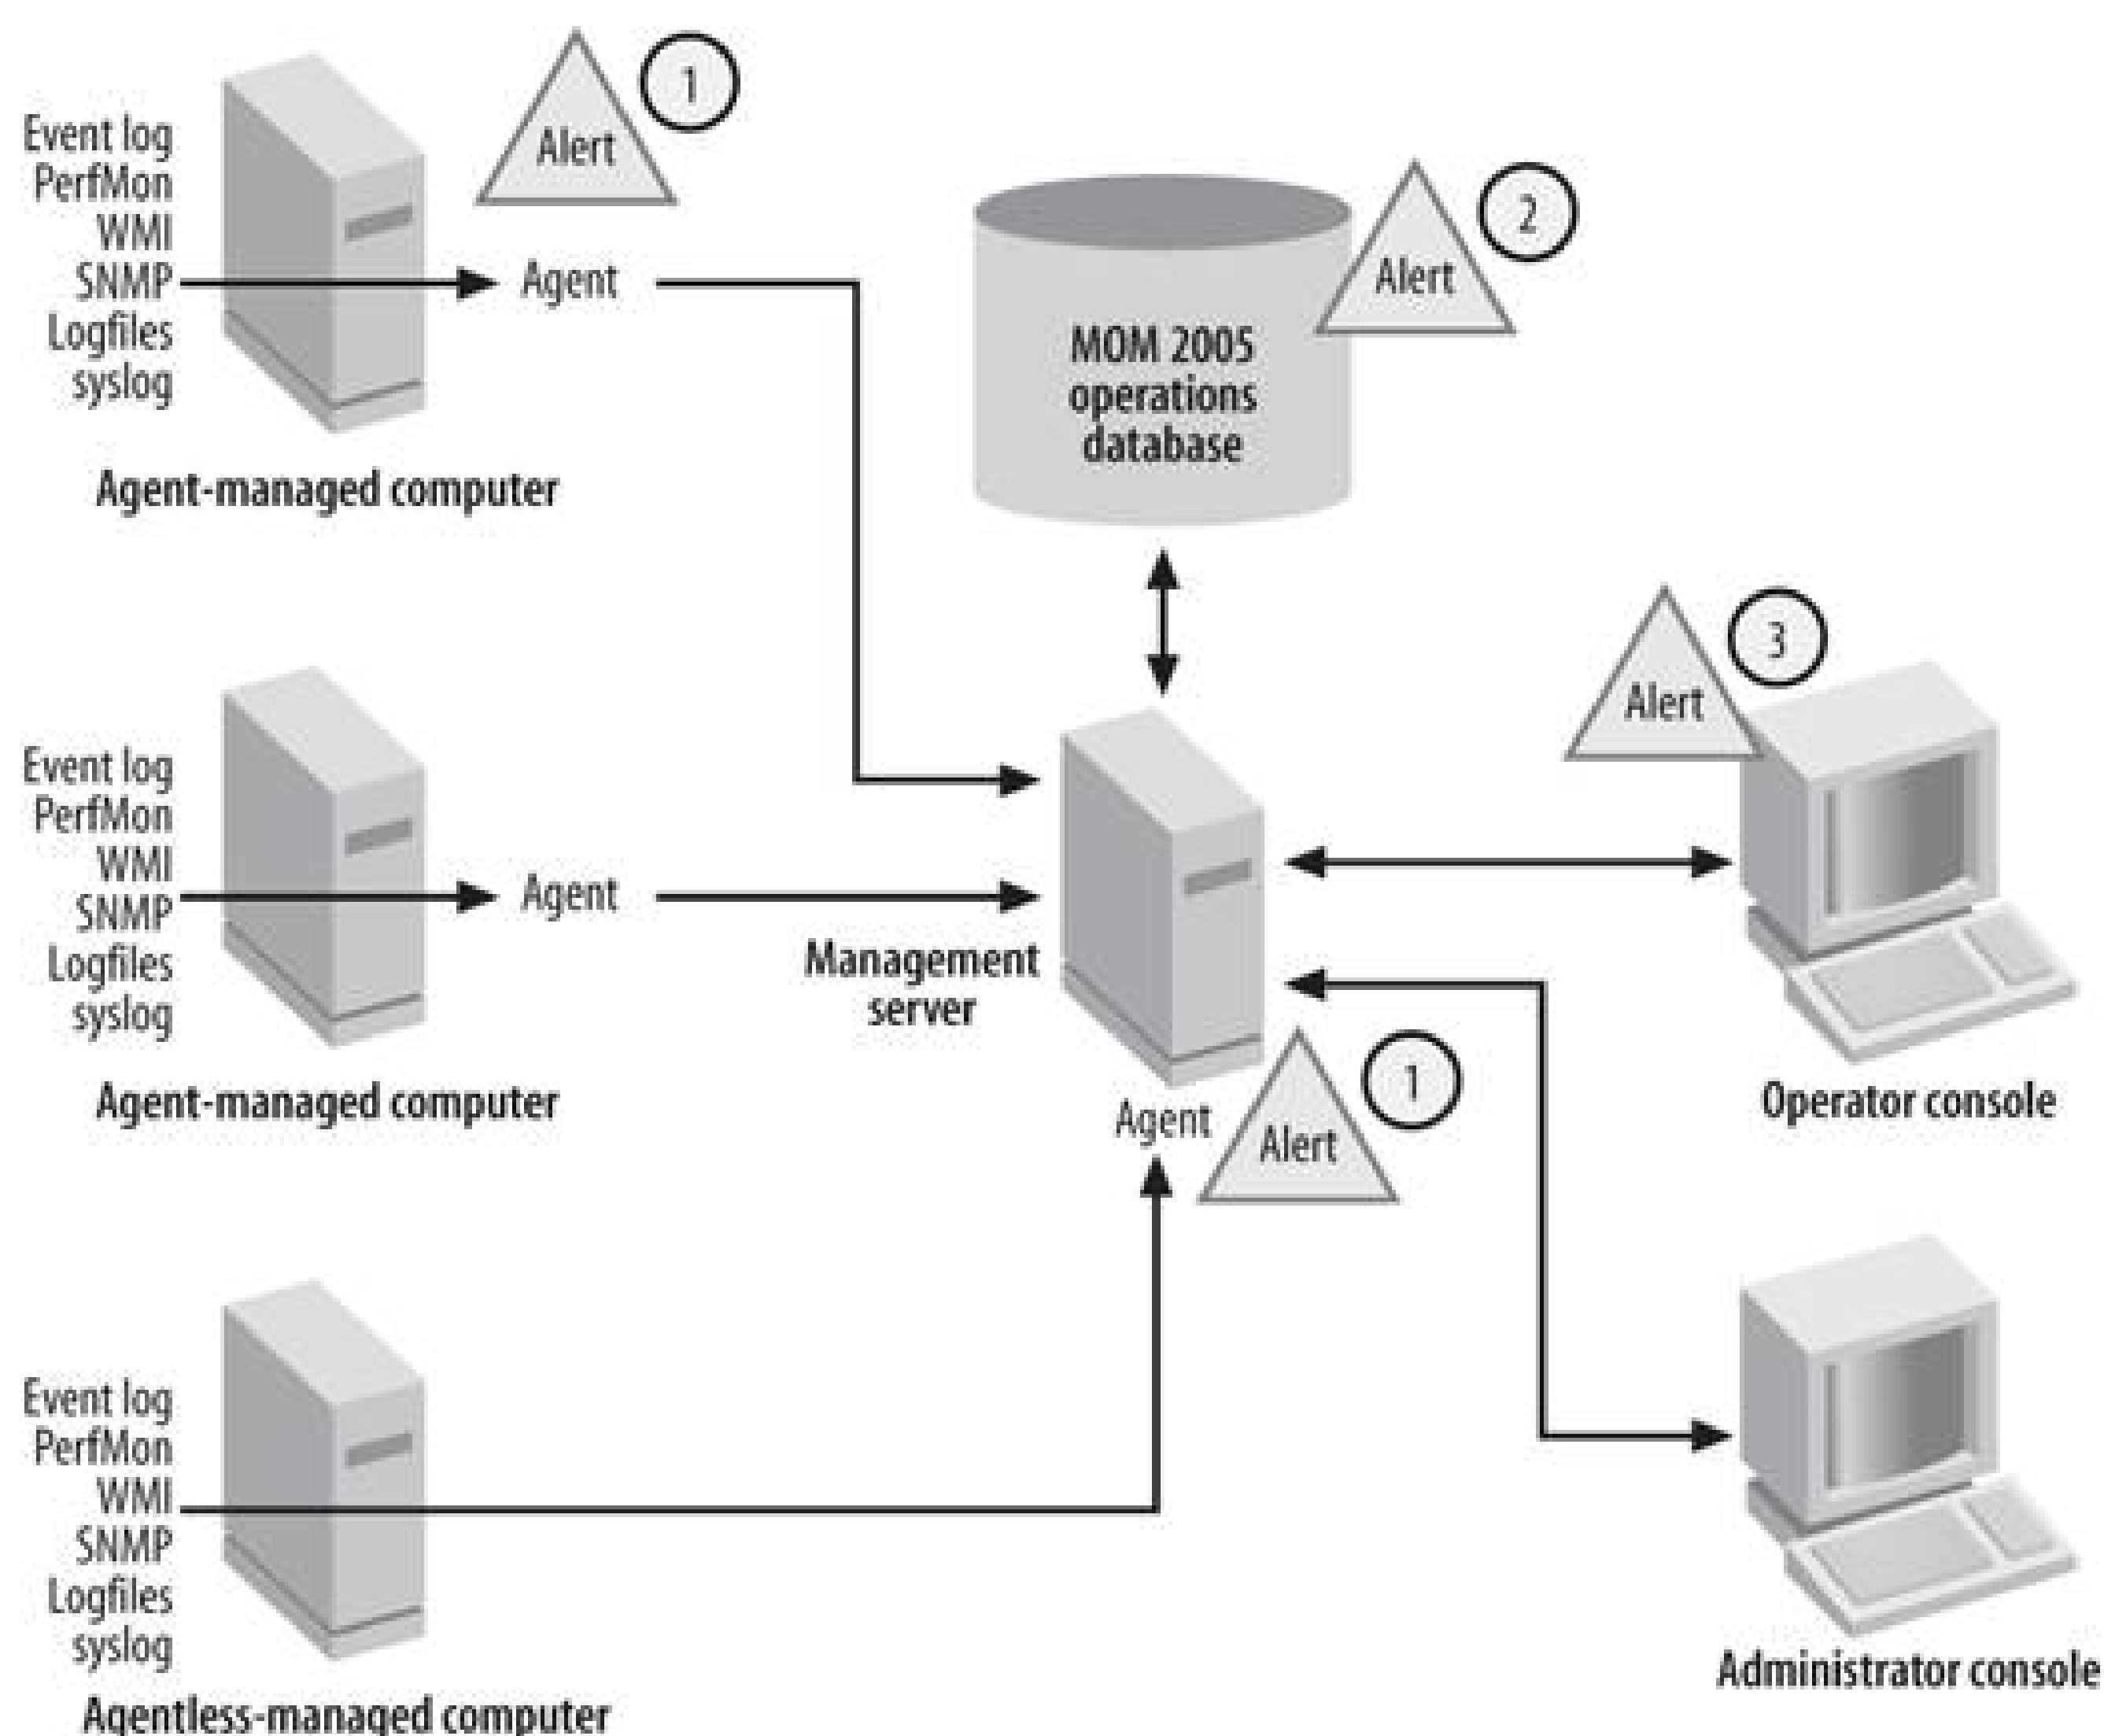
Note that a MOM 2005 alert is not the same as a Windows Event log event. A Windows event is written to the event log on the server that the event occurred on and goes no further. It is specific to a service or component of that server and is, essentially, restricted to that computer.

Sometimes, a single event provides you with enough information to take action, but most of the time it doesn't. Events do play a role in the generation of alerts, but they themselves are not alerts. In the world of MOM, think of an event as the indication of a symptom that something is wrong, not a diagnosis in and of itself. For example, if someone is sick and has a 103-degree fever, the measurement of the 103-degree fever would be like an event, as would a cough and stiff neck. This person would then go to a doctor who would consider all the symptoms before making a diagnosis (an alert) of the flu and prescribing bed rest, plenty of fluids, and pain reliever.

All alerts are generated by MOM 2005 agents (see point 1 in [Figure 1-2](#)), whether that agent is running on a monitored computer or on the management server. In the normal course of operations, agents function independently of administrator intervention. They receive instructions on what work to perform and how to perform it from the management packs (the health rules). At the same time, agents collect data from their host machines by monitoring things such as the event logs, performance monitor counters, and executing scripts that use the Windows Management Interface (WMI) API against monitored computers. Agents then compare the collected data to values that have been predefined by Microsoft (for Microsoft-written applications) to describe health for that application. When the collected data meets the criteria in the health rules, the agent generates an alert. The alert is sent to the management server (see point 2 in [Figure 1-2](#)), recorded in the operations database, cross-correlated with other alerts to determine if a summary alert or other alert should be generated, and then surfaces in the Operator console (see point 3 in [Figure 1-2](#)). After a period of time, most alerts are groomed (deleted) out of the operations database.

Figure 1-2. Alert flow through MOM

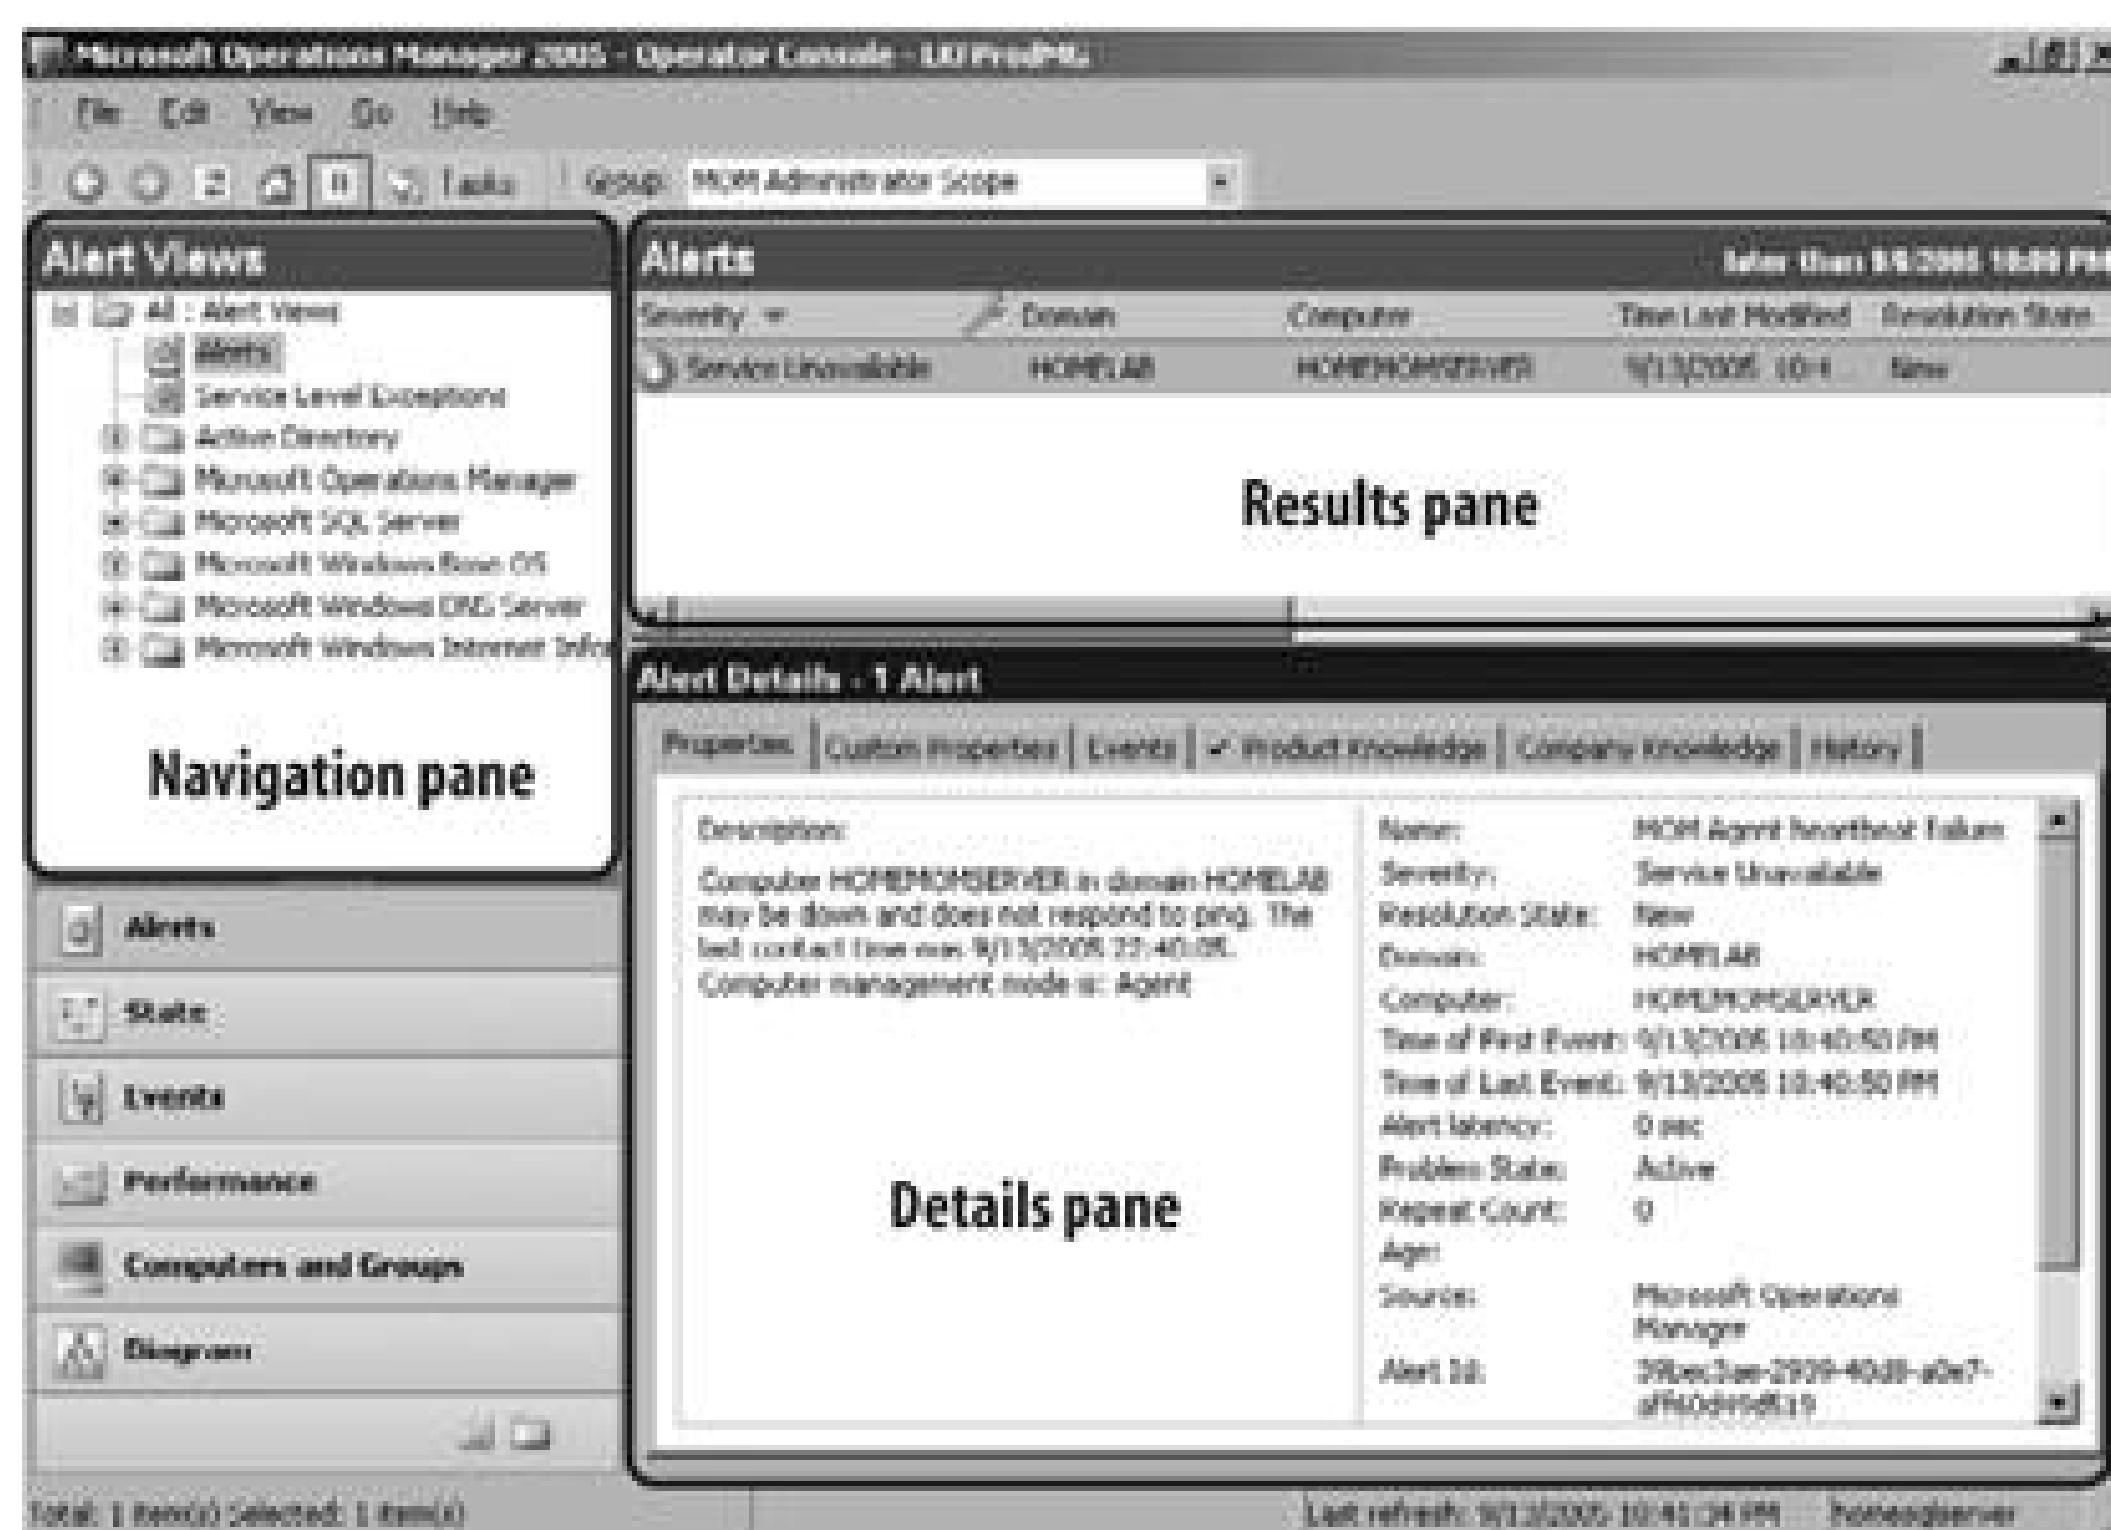




You will manage alerts, and almost all the other information that MOM gives you, in the Operator console . Once the alert surfaces in the Operator console, you can act on it, drill down into it for more information, and modify it by adding your own information on how it was fixed.

[Figure 1-3](#) shows the Operator console with an alert that was created by unplugging the network cable from the computer *homemomserver*. The Alerts view shown here, sorted by the Time Last Modified field, is the default view for the Operator console. The whole console is patterned after Outlook 2003, so there are some panes intentionally hidden here for the sake of simplicity. The Operator console is discussed in detail in [Chapter 6](#).

Figure 1-3. The heartbeat failure alert means that homemomserver is not communicating with the management server



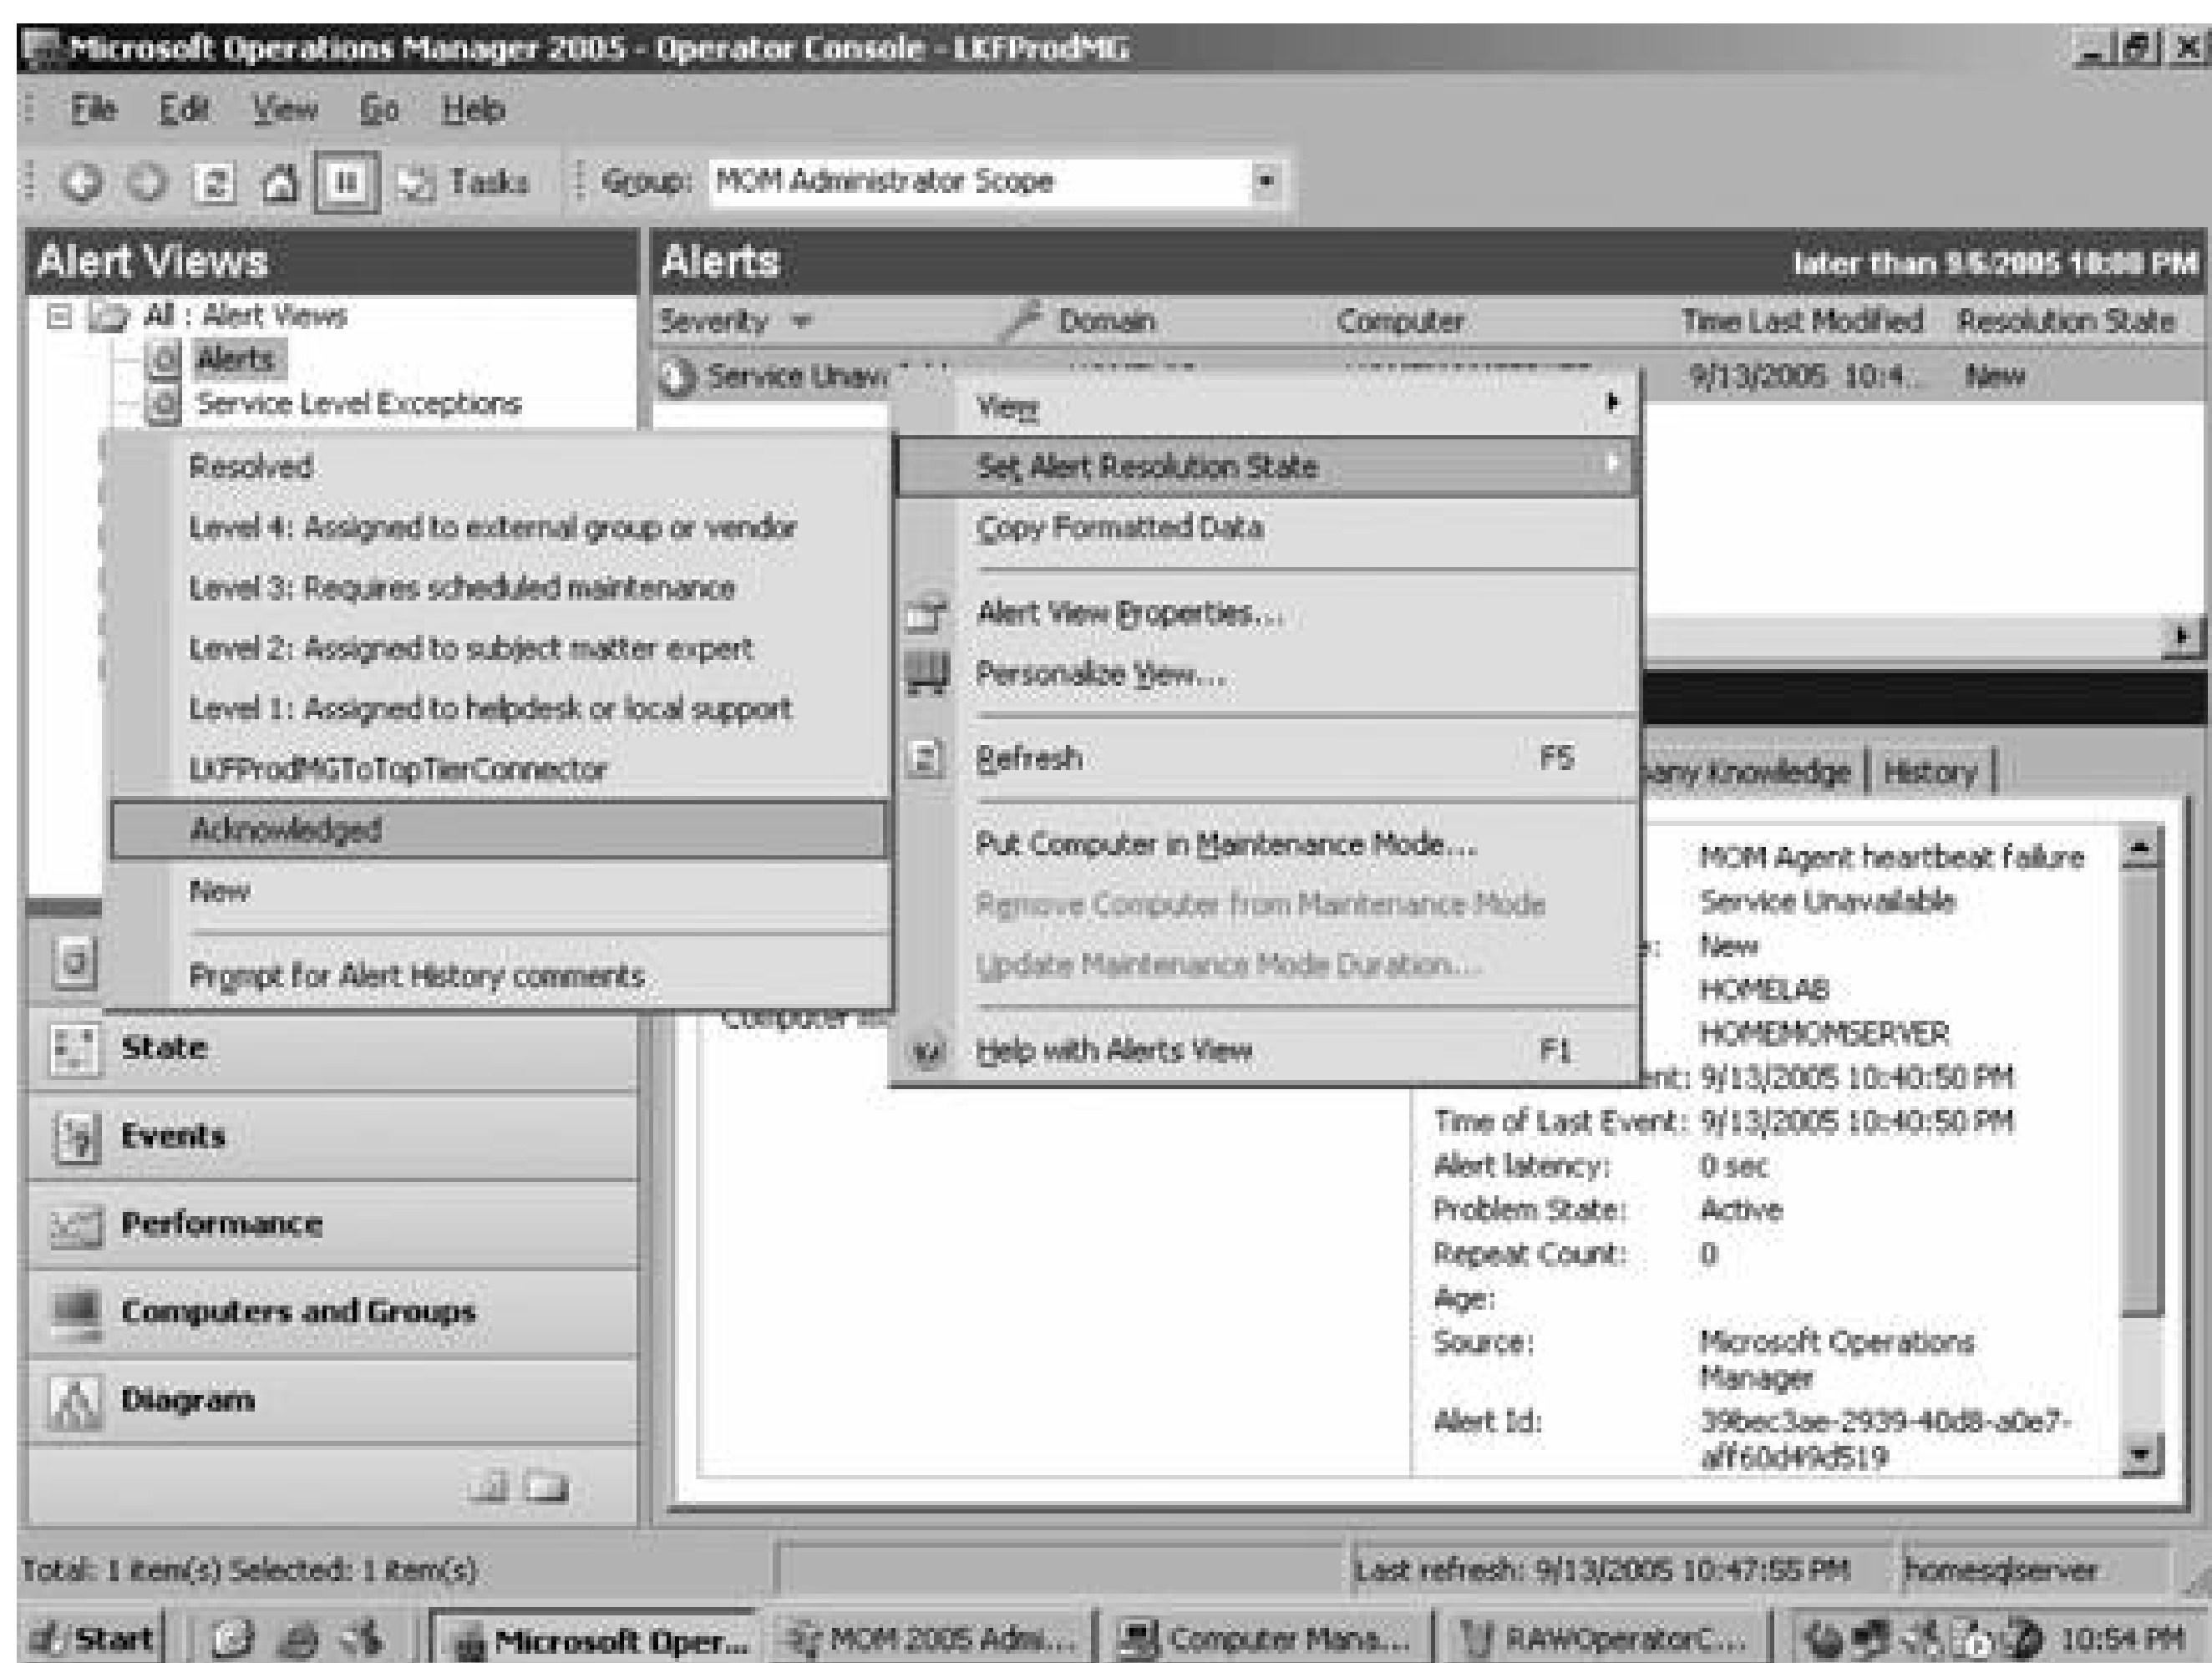
In the Alert Details pane , MOM automatically displays the Properties tab of the alert that has been selected in the Results pane. What is displayed in the Results pane is controlled by your selections in the Navigation pane.

Depending on what is going on in your environment, and how often you clear the alerts by resolving them, the Results pane of the Alerts view may be full of alerts or only have a few. If there are multiple alerts present there, sort them by severity in descending order by clicking on the Severity column header this places the most severe alerts at the top. With most alerts, the first thing you want to do is read the information that it contains and get more information if you are not familiar with the alert. You do this by going through the tabs of the alert in the Details pane.

The Properties tab gives the description of the alert; in this case, "Computer HOMEMOMSERVER in domain HOMELAB may be down and does not respond to ping. The last contact time was 9/13/2005 22:40:05. Computer management mode is: Agent." Along with the name of the alert, "MOM Agent heartbeat failure," this tells you a few things. First off, the agent on the computer missed its regularly scheduled heartbeat its last good heartbeat was at 10:40 p.m. which could be due to any number of reasons, but the computer is also not responding to a TCP/IP ping.

Two other things to note on the Properties tab are the Resolution State and the Repeat Count fields. An alert can exist in one of several resolution states. When it first surfaces in the Operator console, it will always be in a resolution state of New unless you manually configure the rule that generated the alert with a different default state. Since this alert is now being examined and the issue that caused it is being resolved, the resolution state is changed to Acknowledged. This lets anyone else who is viewing the Operator console know that the alert has been seen and is being acted on (see [Figure 1-4](#)).

Figure 1-4. Setting the alert resolution state to Acknowledged



If additional help is needed to fix the issue, the resolution state can be updated again to one of the other values (Levels 1 through 4), which assigns the disposition of an alert to a different group. This will help you keep track of what is going on with the alert and who owns the alert. MOM 2005 also tracks how long an alert is in any of the resolution states. When the time spent in a resolution state exceeds a configurable limit, the alert is flagged and it will appear in the Service Level Exceptions Alert view . This lets you know that the person the alert was assigned to has not updated or resolved the alert in the allotted time. It will also show if the company's service-level agreements are being met on problem response and resolution time.

Out-of-the-box, a change in the resolution state does not kick off a workflow, like a help desk trouble ticketing application. However, since resolution state is a property of an alert, you could script a response that fires when an alert is placed into a state, such as "Level 1: Assigned to help desk or local support." This would notify a predefined group of users (Level 1 operators) via email, pager, or another mechanism that they have been assigned an alert and are responsible for resolving it. See the "[Alert notification](#)" section in [Chapter 4](#).

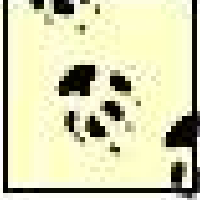
The Repeat Count field indicates how many times this specific alert has been raised in the Operator console. MOM 2005 will increment the value in this field every time a duplicate alert is generated rather than placing a new alert into the Alerts view, which could possibly generate a page or email. This saves you from being flooded with duplicate alerts. For an alert to be considered a duplicate, there must be an existing alert in the Operator console with a resolution state of anything except Resolved, and it must have been generated by the exact same rule on the exact same machine. The criteria can be further refined in the rule under Alert Suppression. This is performed in the Administrator console and is covered in [Chapter 5](#).

The next place to look for more information is on the Events tab in the Details pane. This tab lists all



the Windows events that are associated with the "MOM Agent heartbeat failure alert" (see [Figure 1-5](#)).

This event actually occurred on the management server *homesq/server* and has the severity of Warning. MOM 2005 uses a small lightning bolt in a yellow triangle to indicate the event associated with an alert. When you start examining the different views in the Operator console, you will see the Events view. There, you can examine all Windows events that have been collected from managed computers (no more reviewing the event logs one machine at a time) and you will notice that some events are not associated with alerts and some are.



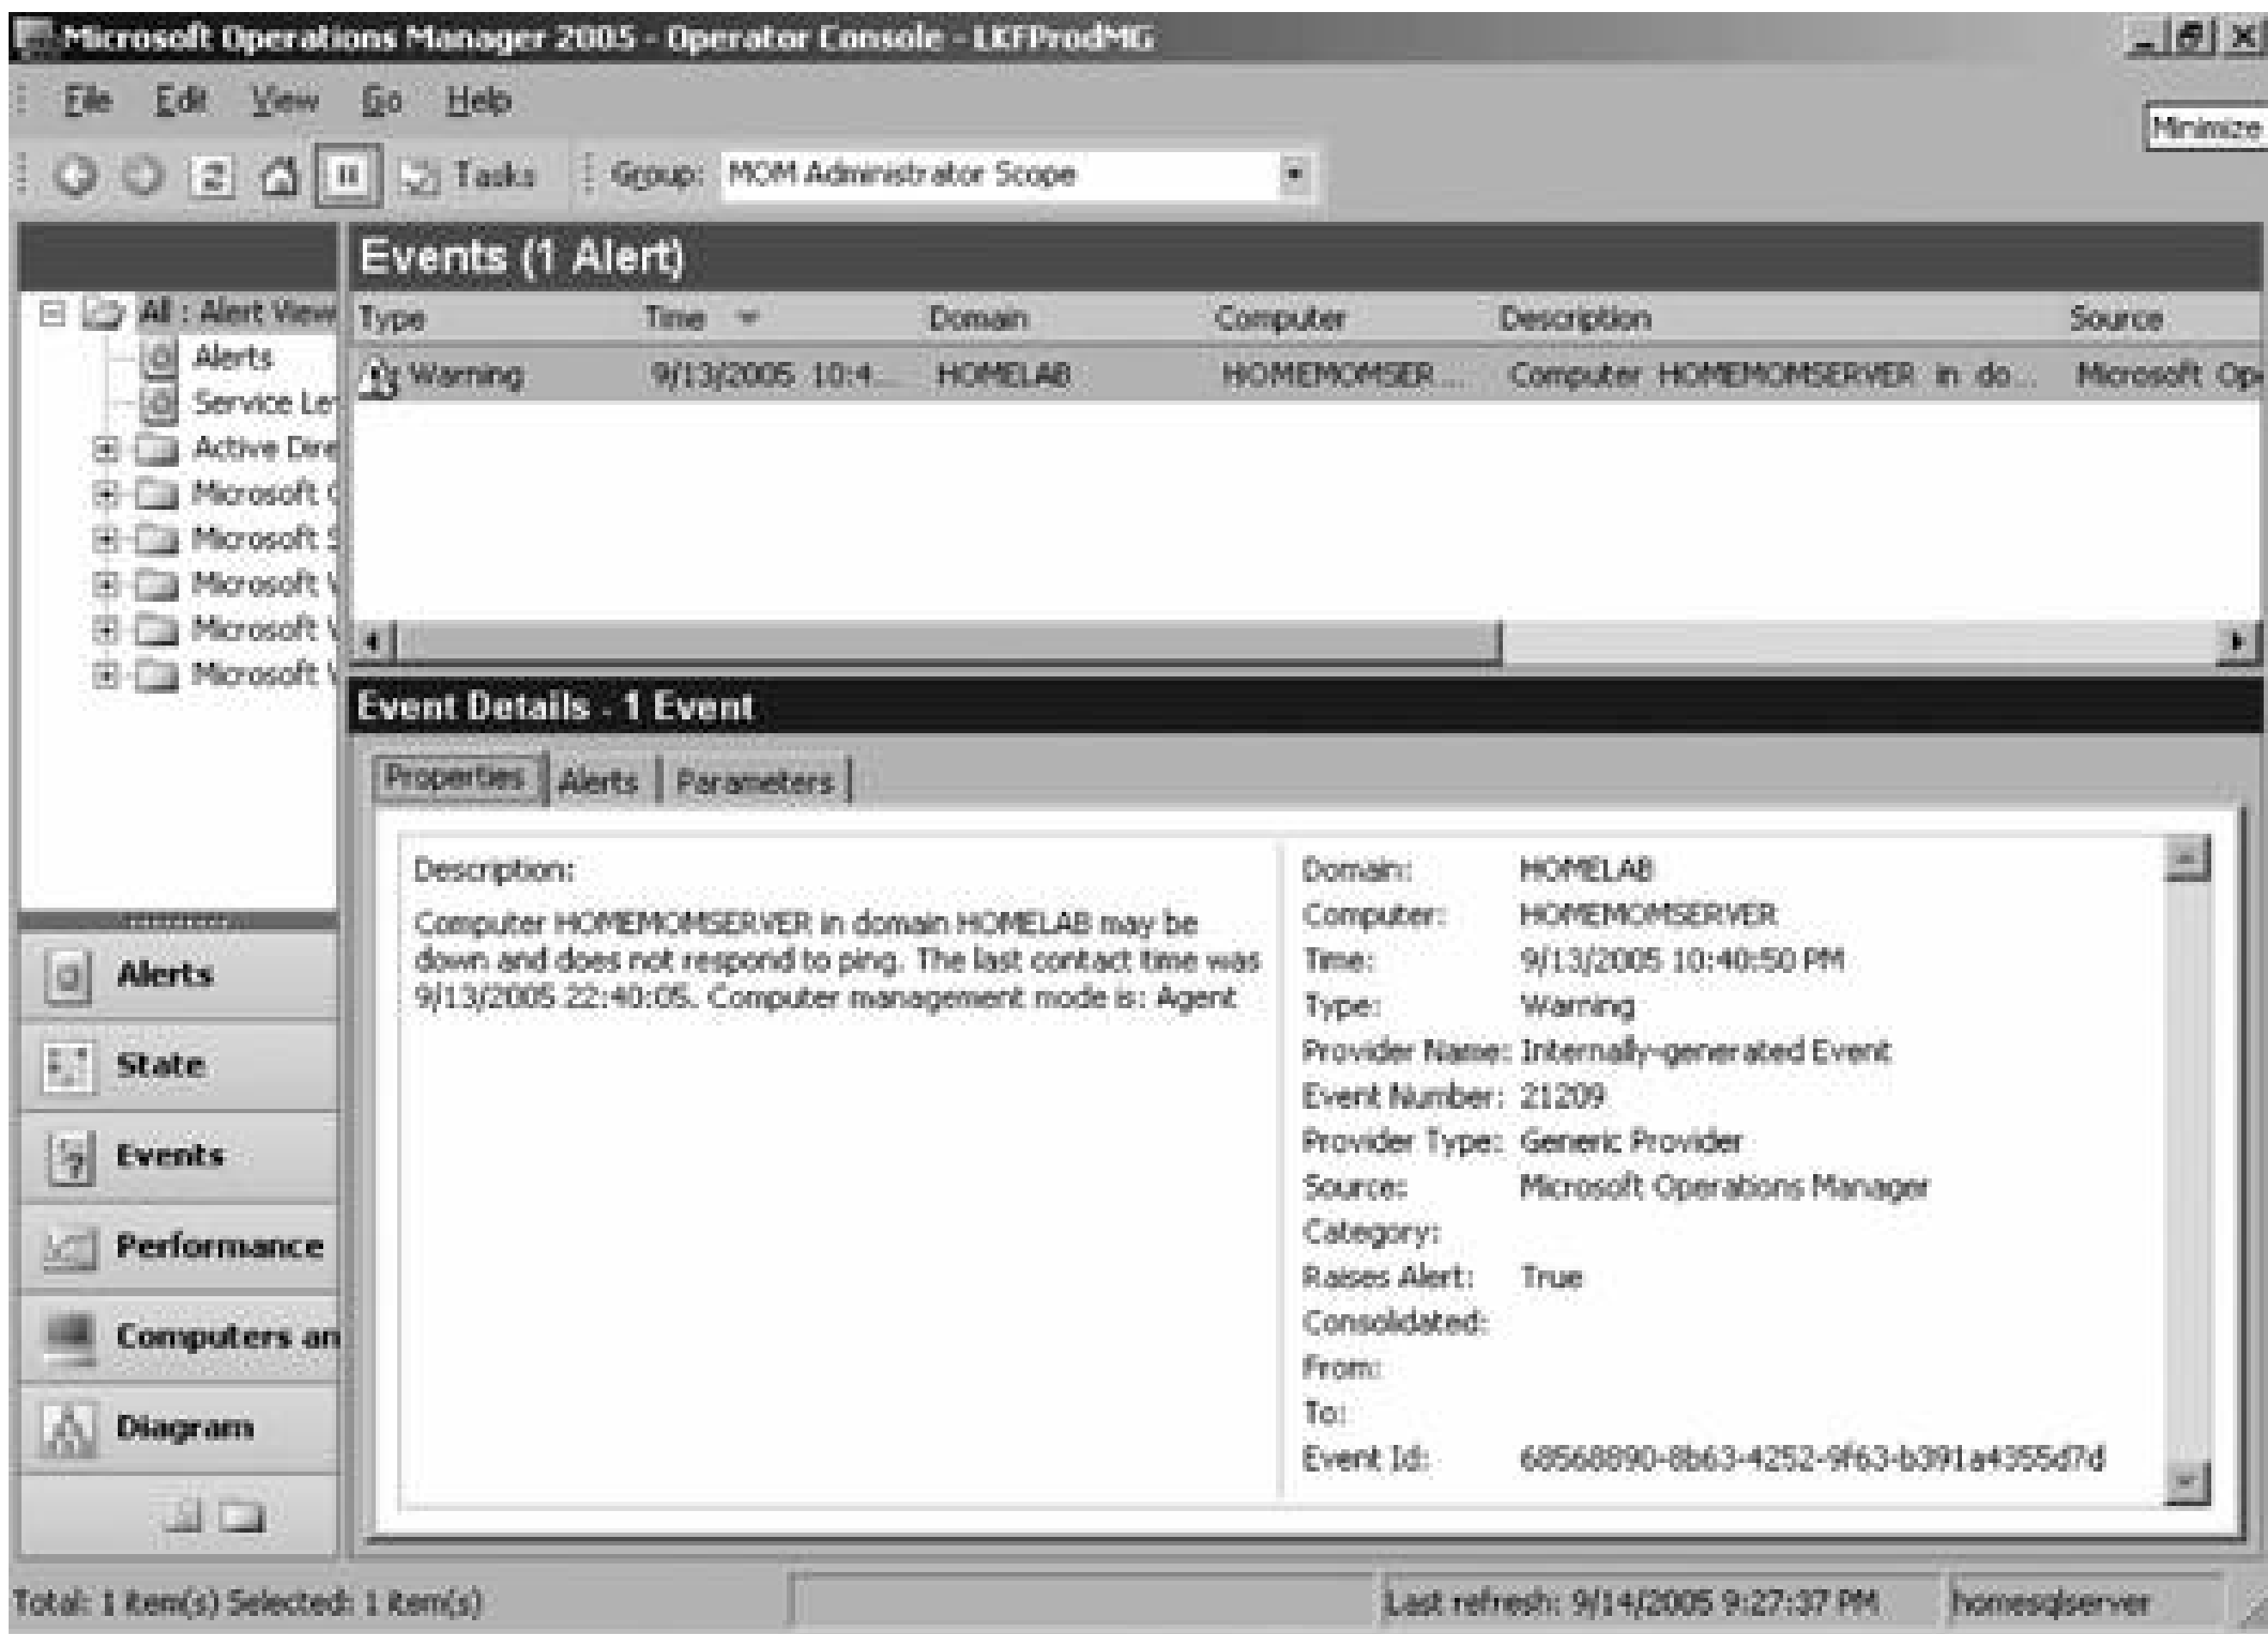
You should be asking why this event was generated on a machine other than the one that is not communicating (*homemomserver*). This is because machines cannot perform heartbeat checking on themselves. If they did, then when they went down or became unavailable, the MOM agent would not be able to communicate to the outside world to notify the management server that the machine was down.

Next, drill down into the Event Details view by double-clicking to see what the specific event number is (see [Figure 1-6](#)).

Figure 1-5. The Windows event that triggered the "MOM Agent heartbeat failure" alert



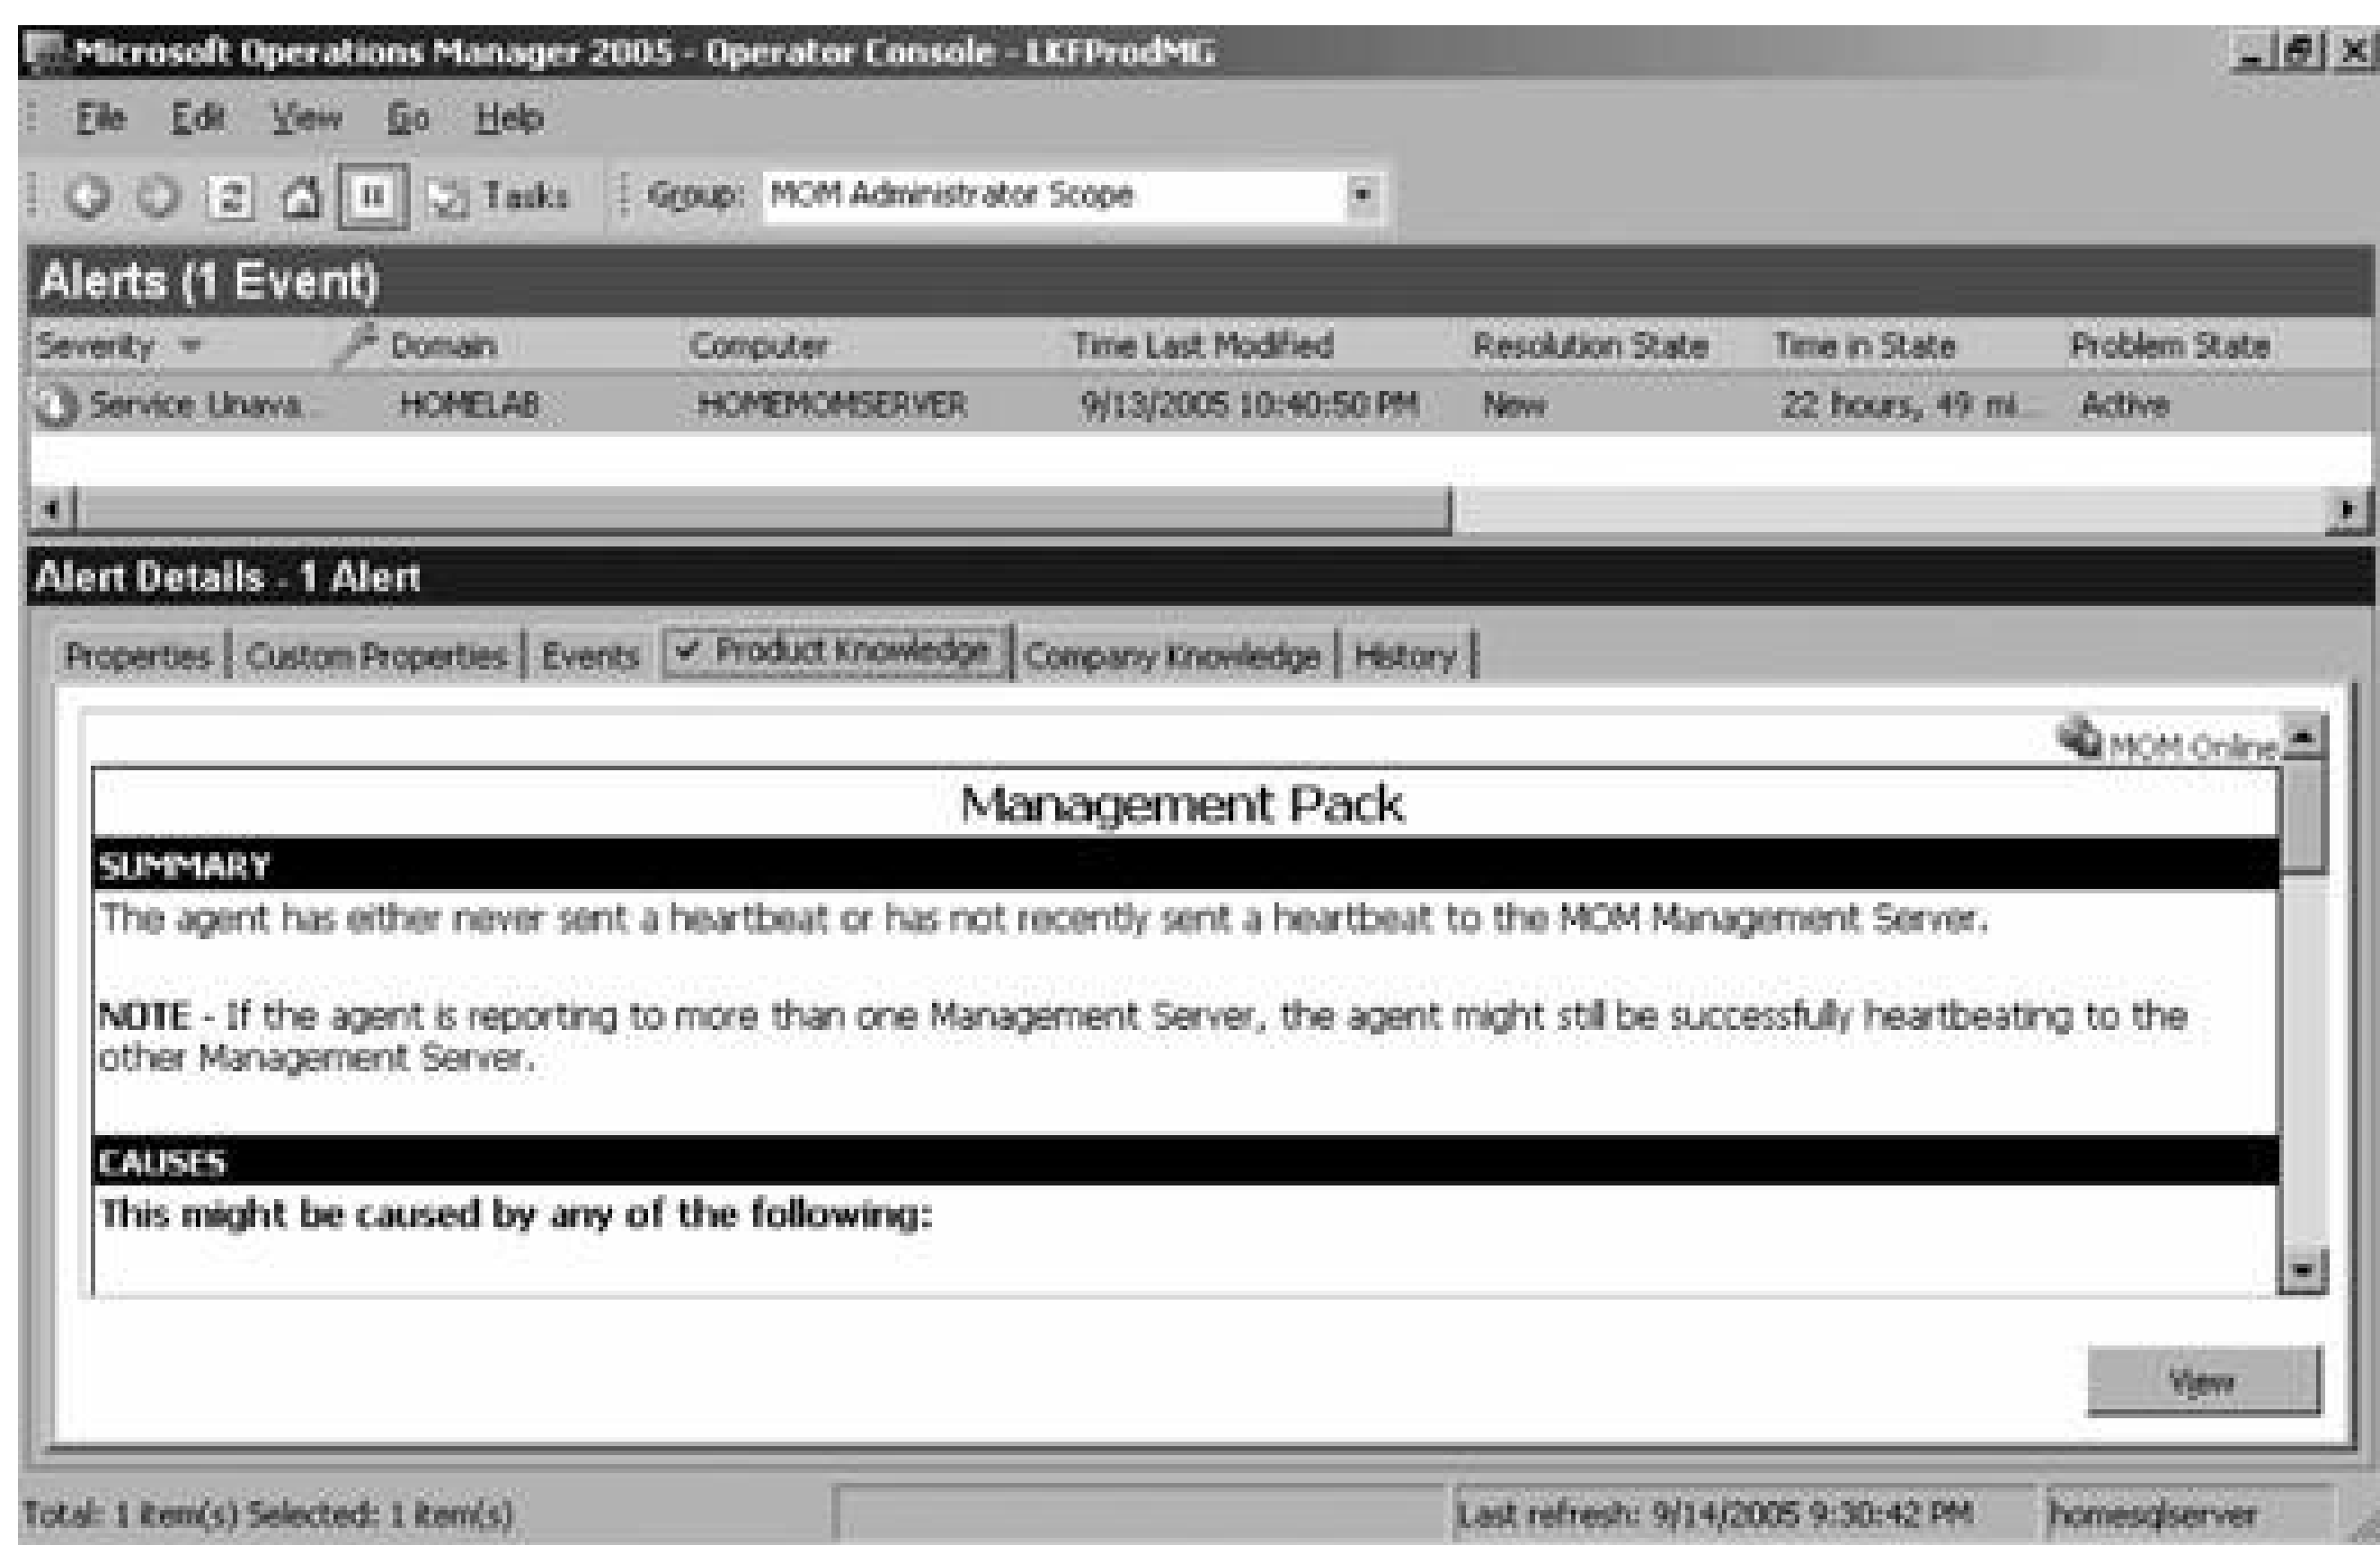
Figure 1-6. The details of a Windows event



This event number is 21209, and you can use this information for further research if necessary. The Alerts tab displays the alerts that this event is associated with, which you already know about.

Back in the Alert Details pane, review the Product Knowledge tab. This tab is prepopulated with information from the application vendor, in this case Microsoft (see [Figure 1-7](#)).

Figure 1-7. The information on the Product Knowledge tab suggests likely causes and the resolution steps



You can skip the Summary section since you already know what this alert is about, but the Causes section reads like this:

- The computer is unavailable.
- The computer's domain is unavailable.
- The MOM Service on the computer is unavailable.
- The agent was uninstalled from the computer.

The Resolutions section gives you various suggestions depending on the error, and reads like this:

The agent has never sent a heartbeat to the MOM Management Server.

- Make sure the computer is available by using the Ping task in the Task pane of the MOM Operator console.
- If the ping fails, make sure the computer still exists and that it is available on the network.
- Make sure the MOM Service is running on the computer. You can start the MOM Service by using the Start MOM 2005 Service task in the Task pane of the MOM Operator console.
- Try to update the agent settings by using the Update Agent Settings dialog on the Management Server.
- Reinstall the agent.

The agent has not recently sent a heartbeat to the MOM Management Server.

- Make sure the computer is available by using the ping command.
- If the ping fails, make sure the computer still exists and that it is available on the network.

The agent failed to send a heartbeat within the allotted time.

- Make sure the computer is available by using the ping command.
- If the ping fails, make sure the computer still exists and that it is available on the network.
- Make sure the MOM Service is running on the computer.

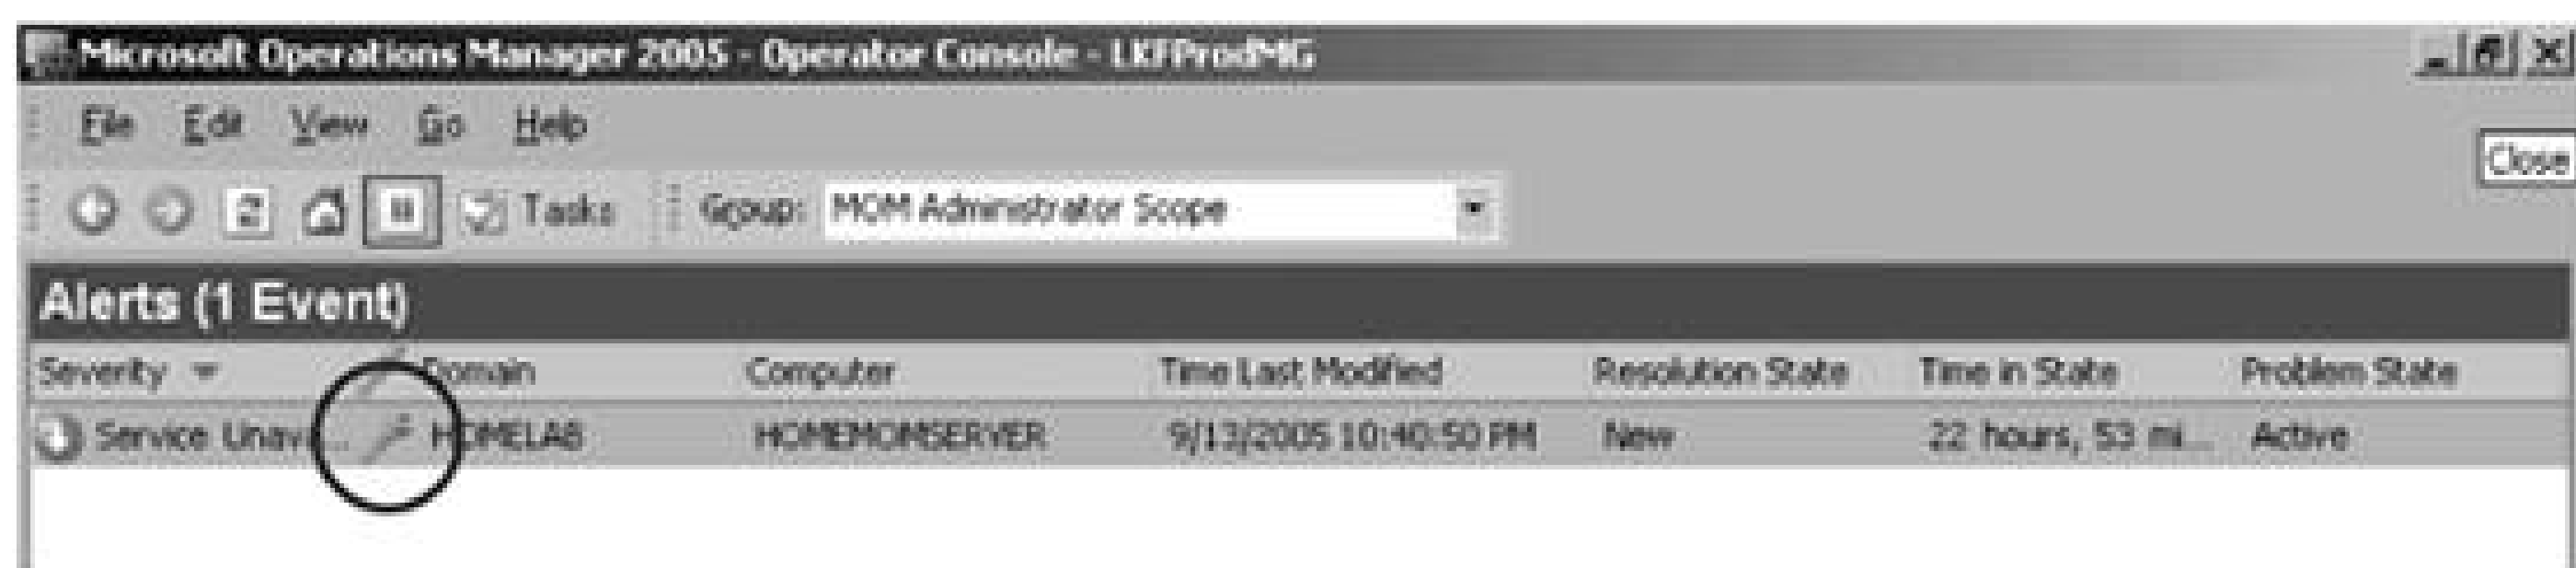
Going through the Causes section, you can immediately eliminate the unavailable domain and the uninstalled agent. You know that the domain is available because all of the other computers in the domain have not lost communication with the domain controller. You also know that you did not uninstall the agent, either automatically, through the Administrator console, or manually. That leaves the unavailable computer and the unavailable MOM Service causes to further diagnose. By looking back at the alert, you can see that MOM pinged *homemomserver* and it did not respond, so the unavailable MOM service can be eliminated as a cause for the missing heartbeat. So, now you know that at the time of the alert being generated the server was either down or disconnected from the network.

Moving to the Resolutions section, you can also go straight to the description that most closely matches the current situation in this case, "The agent has not recently sent a heartbeat to the MOM Management Server" and follow the steps listed there. The first step is to attempt to ping the server again. But since the server is down, there may be other alerts being generated that are associated with *homemomserver*. These alerts could be flooding the console, but you can't deal with these until the server is up. To help manage a flood of alerts from a downed machine, you can put the machine in maintenance mode. To do this, you right-click on the computer name in the Results pane and select Put Computer in Maintenance Mode.

When you place a managed computer into maintenance mode, all alerts generated for that specific computer are automatically resolved by the management server they won't surface in the Alerts view and you don't need to deal with them. You can still examine them later because they have been captured in the operations database. When you place the server in maintenance mode, you configure the maintenance mode duration (it automatically expires on a configurable date and time) and the reason why the server is being placed in maintenance mode. This mode is especially useful when there is work planned on a server that will generate alerts, say the reboot of a server after the installation of a service pack. A small wrench icon in the column between Severity and Domain (see [Figure 1-8](#)) indicates that a computer is in maintenance mode.

Figure 1-8. *homemomserver* in maintenance mode





Now that MOM is protected from an alert flood from this computer, you can continue with the first step suggested in the Resolution section, pinging *homemomserver*.

To do this, you'll need to open the Tasks pane in the Operator console, either by selecting Tasks pane from the View menu, pressing Ctrl-T, or clicking the toggle Tasks pane icon on the toolbar. The Tasks pane contains folders and leaf objects (see [Figure 1-9](#)). The leaf objects are preconfigured to execute commands against whatever machine has the focus in the Results pane. One of these functions is the ping command. This will produce the same result as opening a command prompt and running ping against *homemomserver*. Microsoft has included hooks into these tools in the Operator console to give you an integrated environment in which to perform your monitoring and troubleshooting tasks. You can leave it and make use of the tools manually, but why do things the hard way?

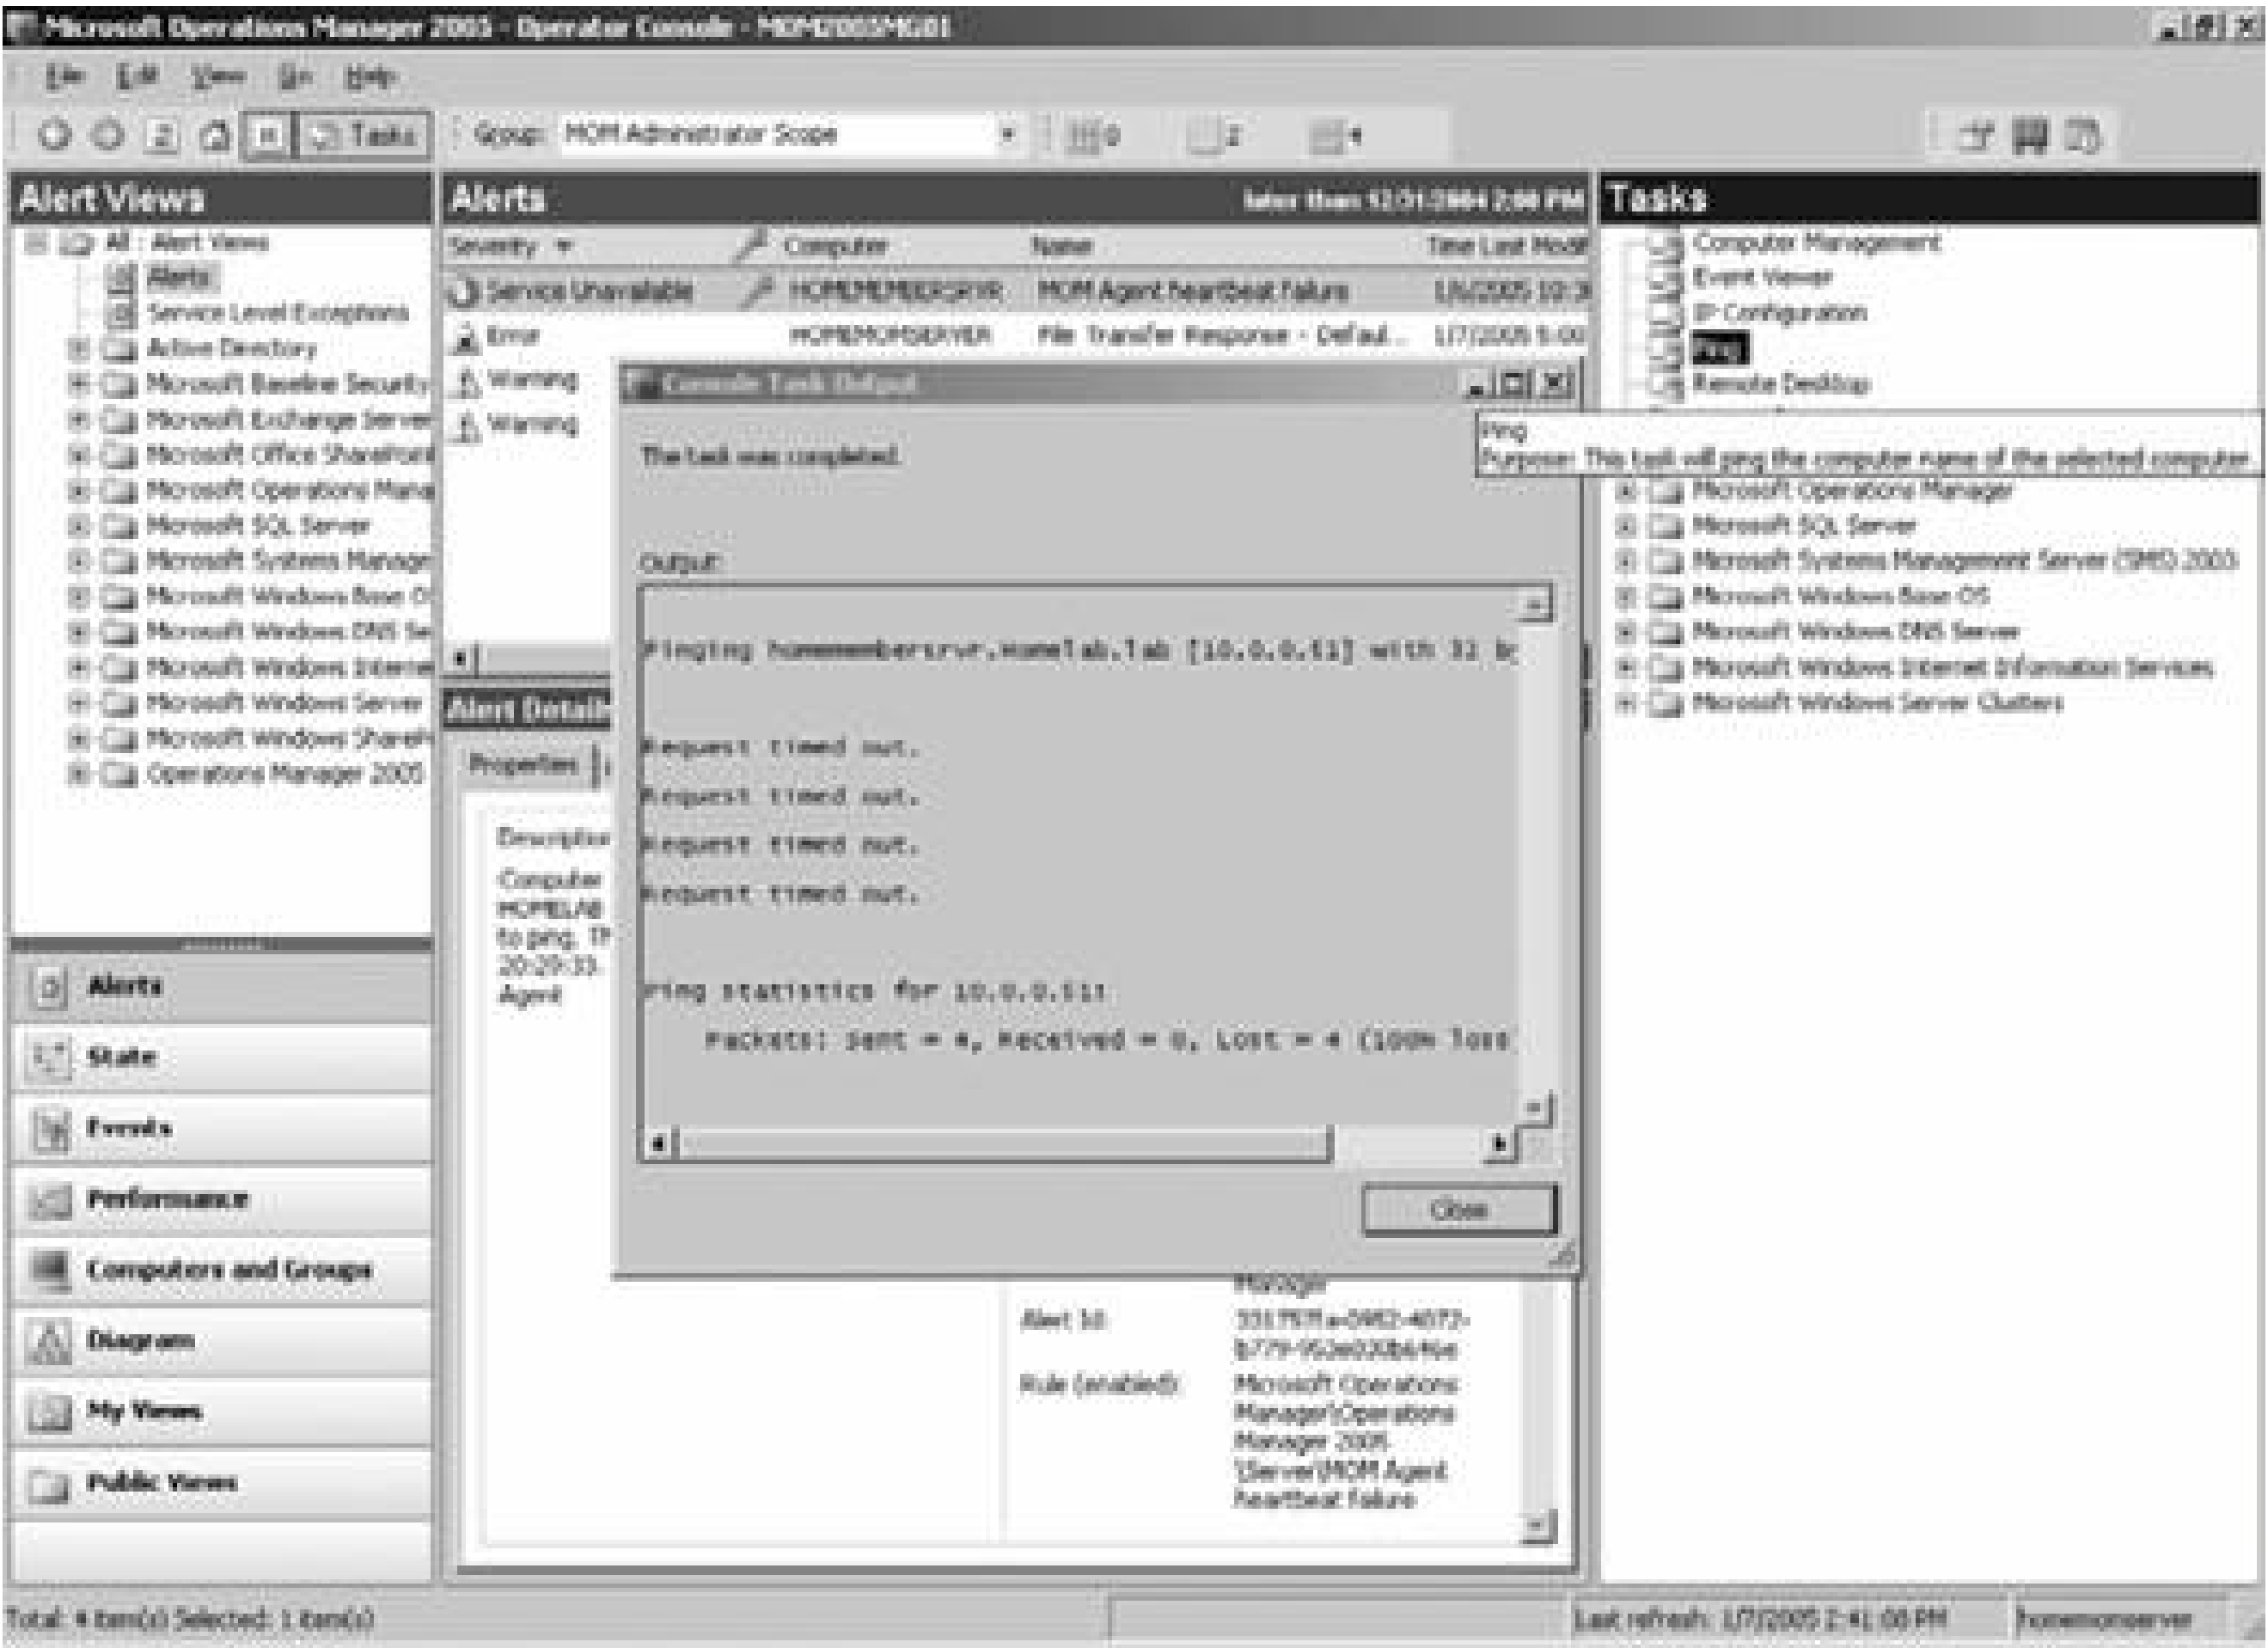
When you invoke a Task object, you will be given the opportunity to enter custom command-line parameters to be passed to the tool when it executes. So, you have to know what the tool does and how to use it before clicking on it in the Operator console. Otherwise, you won't get the desired result and you won't know what the output means.

The folders in the Tasks pane correspond to the application management packs (the health rules) that have been imported into MOM 2005. Each management pack is preconfigured with tasks that provide you with access to the tools that are most commonly used in troubleshooting that application. Additional tasks can be created as needed in your environment through the Administrator console.

Since you can't ping the server, there is no point in trying to make use of some of the other tasks listed, such as Remote Desktop or Computer Management, in this remote diagnosis process. These do, however, invoke the same tools that you are

Figure 1-9. Ping results continue to indicate that homemembersrvr is not accessible from the network





already familiar with Remote Desktop is the Windows 2003 version of terminal services in administration mode, and Computer Management is the same interface you get when you right-click the My Computer icon on your desktop and select the Manage option.

At this point, you have to physically check the computer and plug the disconnected network cable back in. Then you can run the ping task again and you will get a successful response. Some other alerts pop up, telling you that the network connection had been disconnected.

Now that you know what to do to resolve this issue, you should capture that information into the alert so that the next time an alert of this type occurs, information specific to your environment is available along with the product knowledge from the vendor. You will do this on the Company Knowledge tab in the Details pane (see [Figure 1-10](#)). Don't record historical data here, only solution-specific data. After all, for future alerts it doesn't matter that someone kicked the network cable out of the server on the night of the Christmas party. Future troubleshooters will only want to know how to fix the problem.

You can now change the state of the "MOM Agent heartbeat failure" alert to Resolved and remove *homemomserver* from maintenance mode. These actions will remove the alert from the Alert view of the Operator console and allow other heartbeat failure alerts for *homemomserver* to surface as new alerts in the console.

Figure 1-10. Enter solution-specific information in the Company Knowledge tab



← PREV

## 1.4. MOM 2005 Components

[Figure 1-11](#) shows the core components of MOM 2005 that are involved with producing the end product of MOMan alert and the tools you can use to view and act on the alert. To introduce these components, let's work backward through the system, tracing the path of the alert through the core components to its origin.

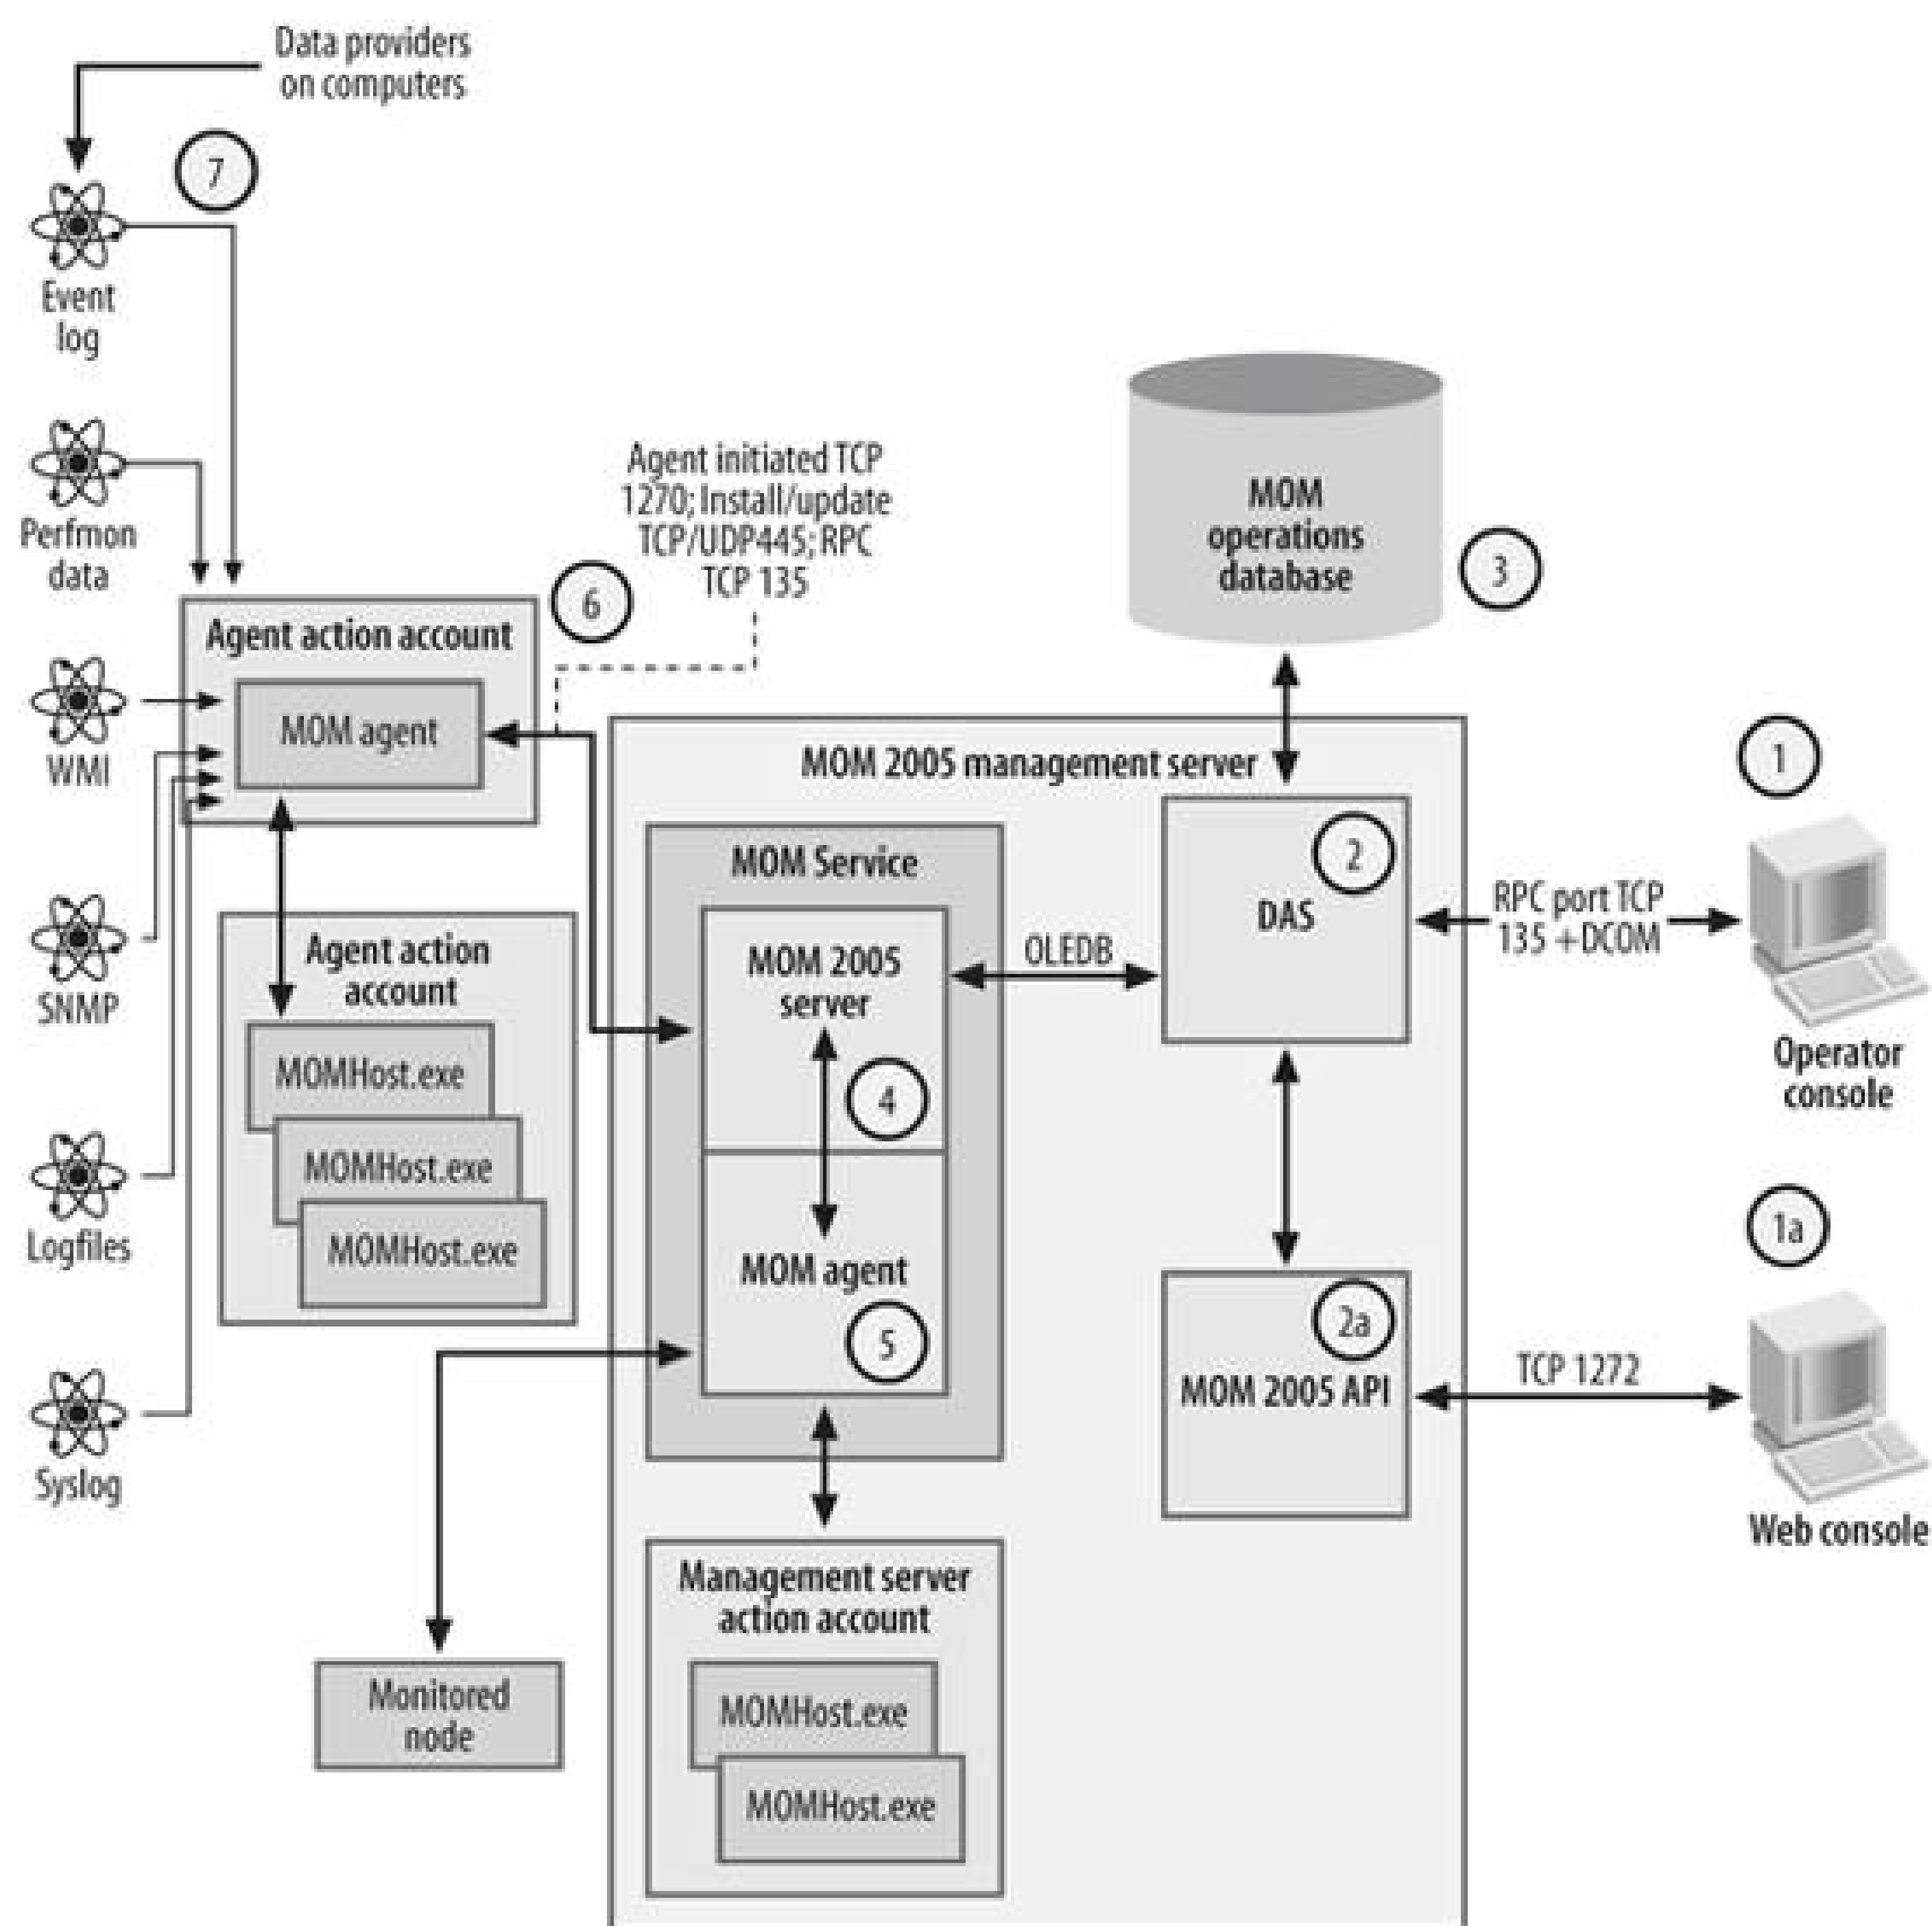
### 1.4.1. The Operator Console

MOM 2005 gives you access to all the information it collects through the Operator console (point 1 in [Figure 1-11](#)). The Operator console is also where you will manage alerts and perform troubleshooting steps. You can also use the Web console (point 1a in [Figure 1-11](#)) for accessing the same information remotely, although you don't get the same level of functionality.

The Operator console is based on the console that Microsoft's internal IT group, the Operations and Technology Group (OTG), developed for its own use in working with MOM 2000 SP1.

In [Figure 1-12](#), there are four panes, three of which you can display or hide at your discretion. On the lefthand side is the Alert Views pane. What you select here controls what you see in the middle two panes. In our example, All: Alert Views is selected and the resulting Alerts and Alert Details are shown in the middle two panes. On the far right side is the Tasks pane. When you select an object in the Tasks pane, you can execute that operation against the computer that has the focus in the middle two panes. For example, if you select the Ping object in the Tasks pane, the Ping command will execute against the computer named *homemomserver3*. You can think of the objects in the Tasks pane as buttons that cause an action to occur.

Figure 1-11. MOM 2005 core components



Alerts and Alert Details panes display all the details that MOM 2005 embeds in an alert. Initially, you make the most use of the information contained on the Properties, Events, Product Knowledge, History, and Company Knowledge tabs.

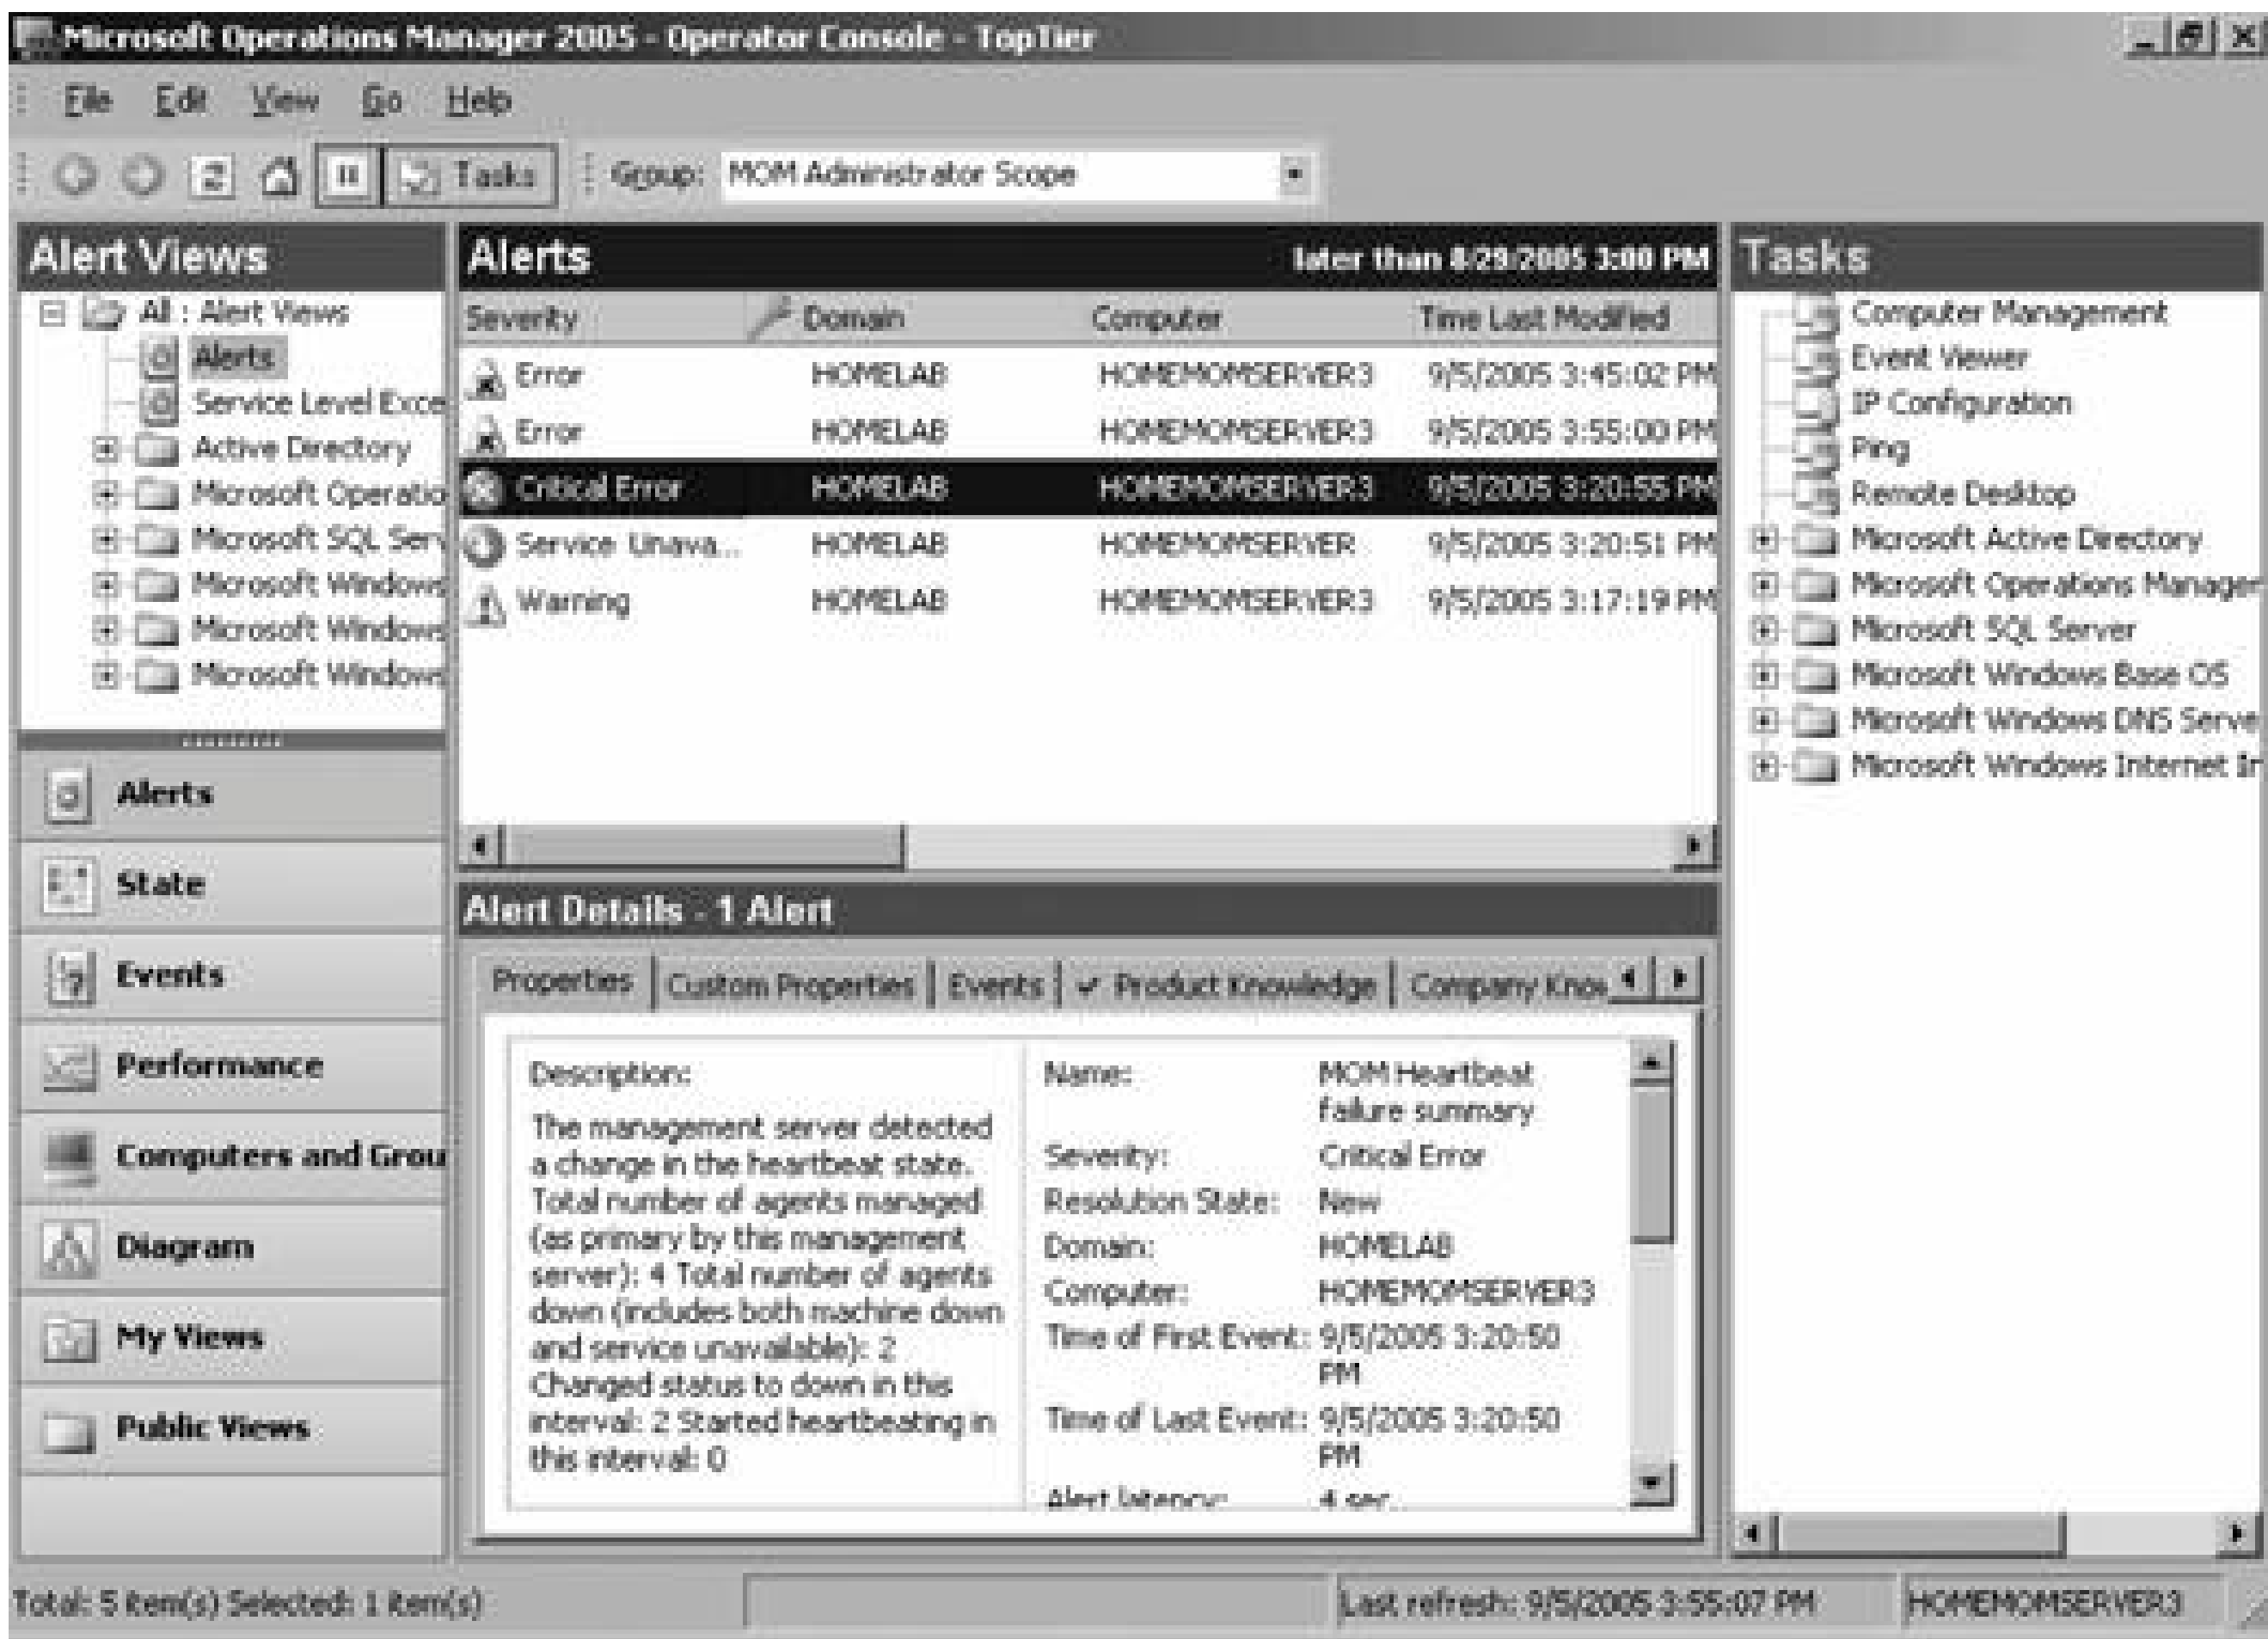
The Properties tab contains detailed information about the alert, including the description, the name of the rule that created the alert, and the time it first and most recently occurred (in the case of a repeating alert).

The Events tab contains links to all the Windows Event log events (or events from other sources such as scripts or logfiles) that caused the alert to be triggered. This is invaluable when analyzing the causes of an alert.

The Product Knowledge tab contains a summary of the issue that caused the alert, possible technical causes of the issue, and proposed resolution steps. The information on this tab represents what the product team identified as the most likely causes



Figure 1-12. The Operator console



of the issue and the most likely resolutions. If you saw an alert about an Exchange server, the Product Knowledge tab would have the summary, causes, and resolution steps developed by the Microsoft Exchange team themselves.

The History tab keeps track of the life history of an alert from its creation in the Operator console to its resolution.

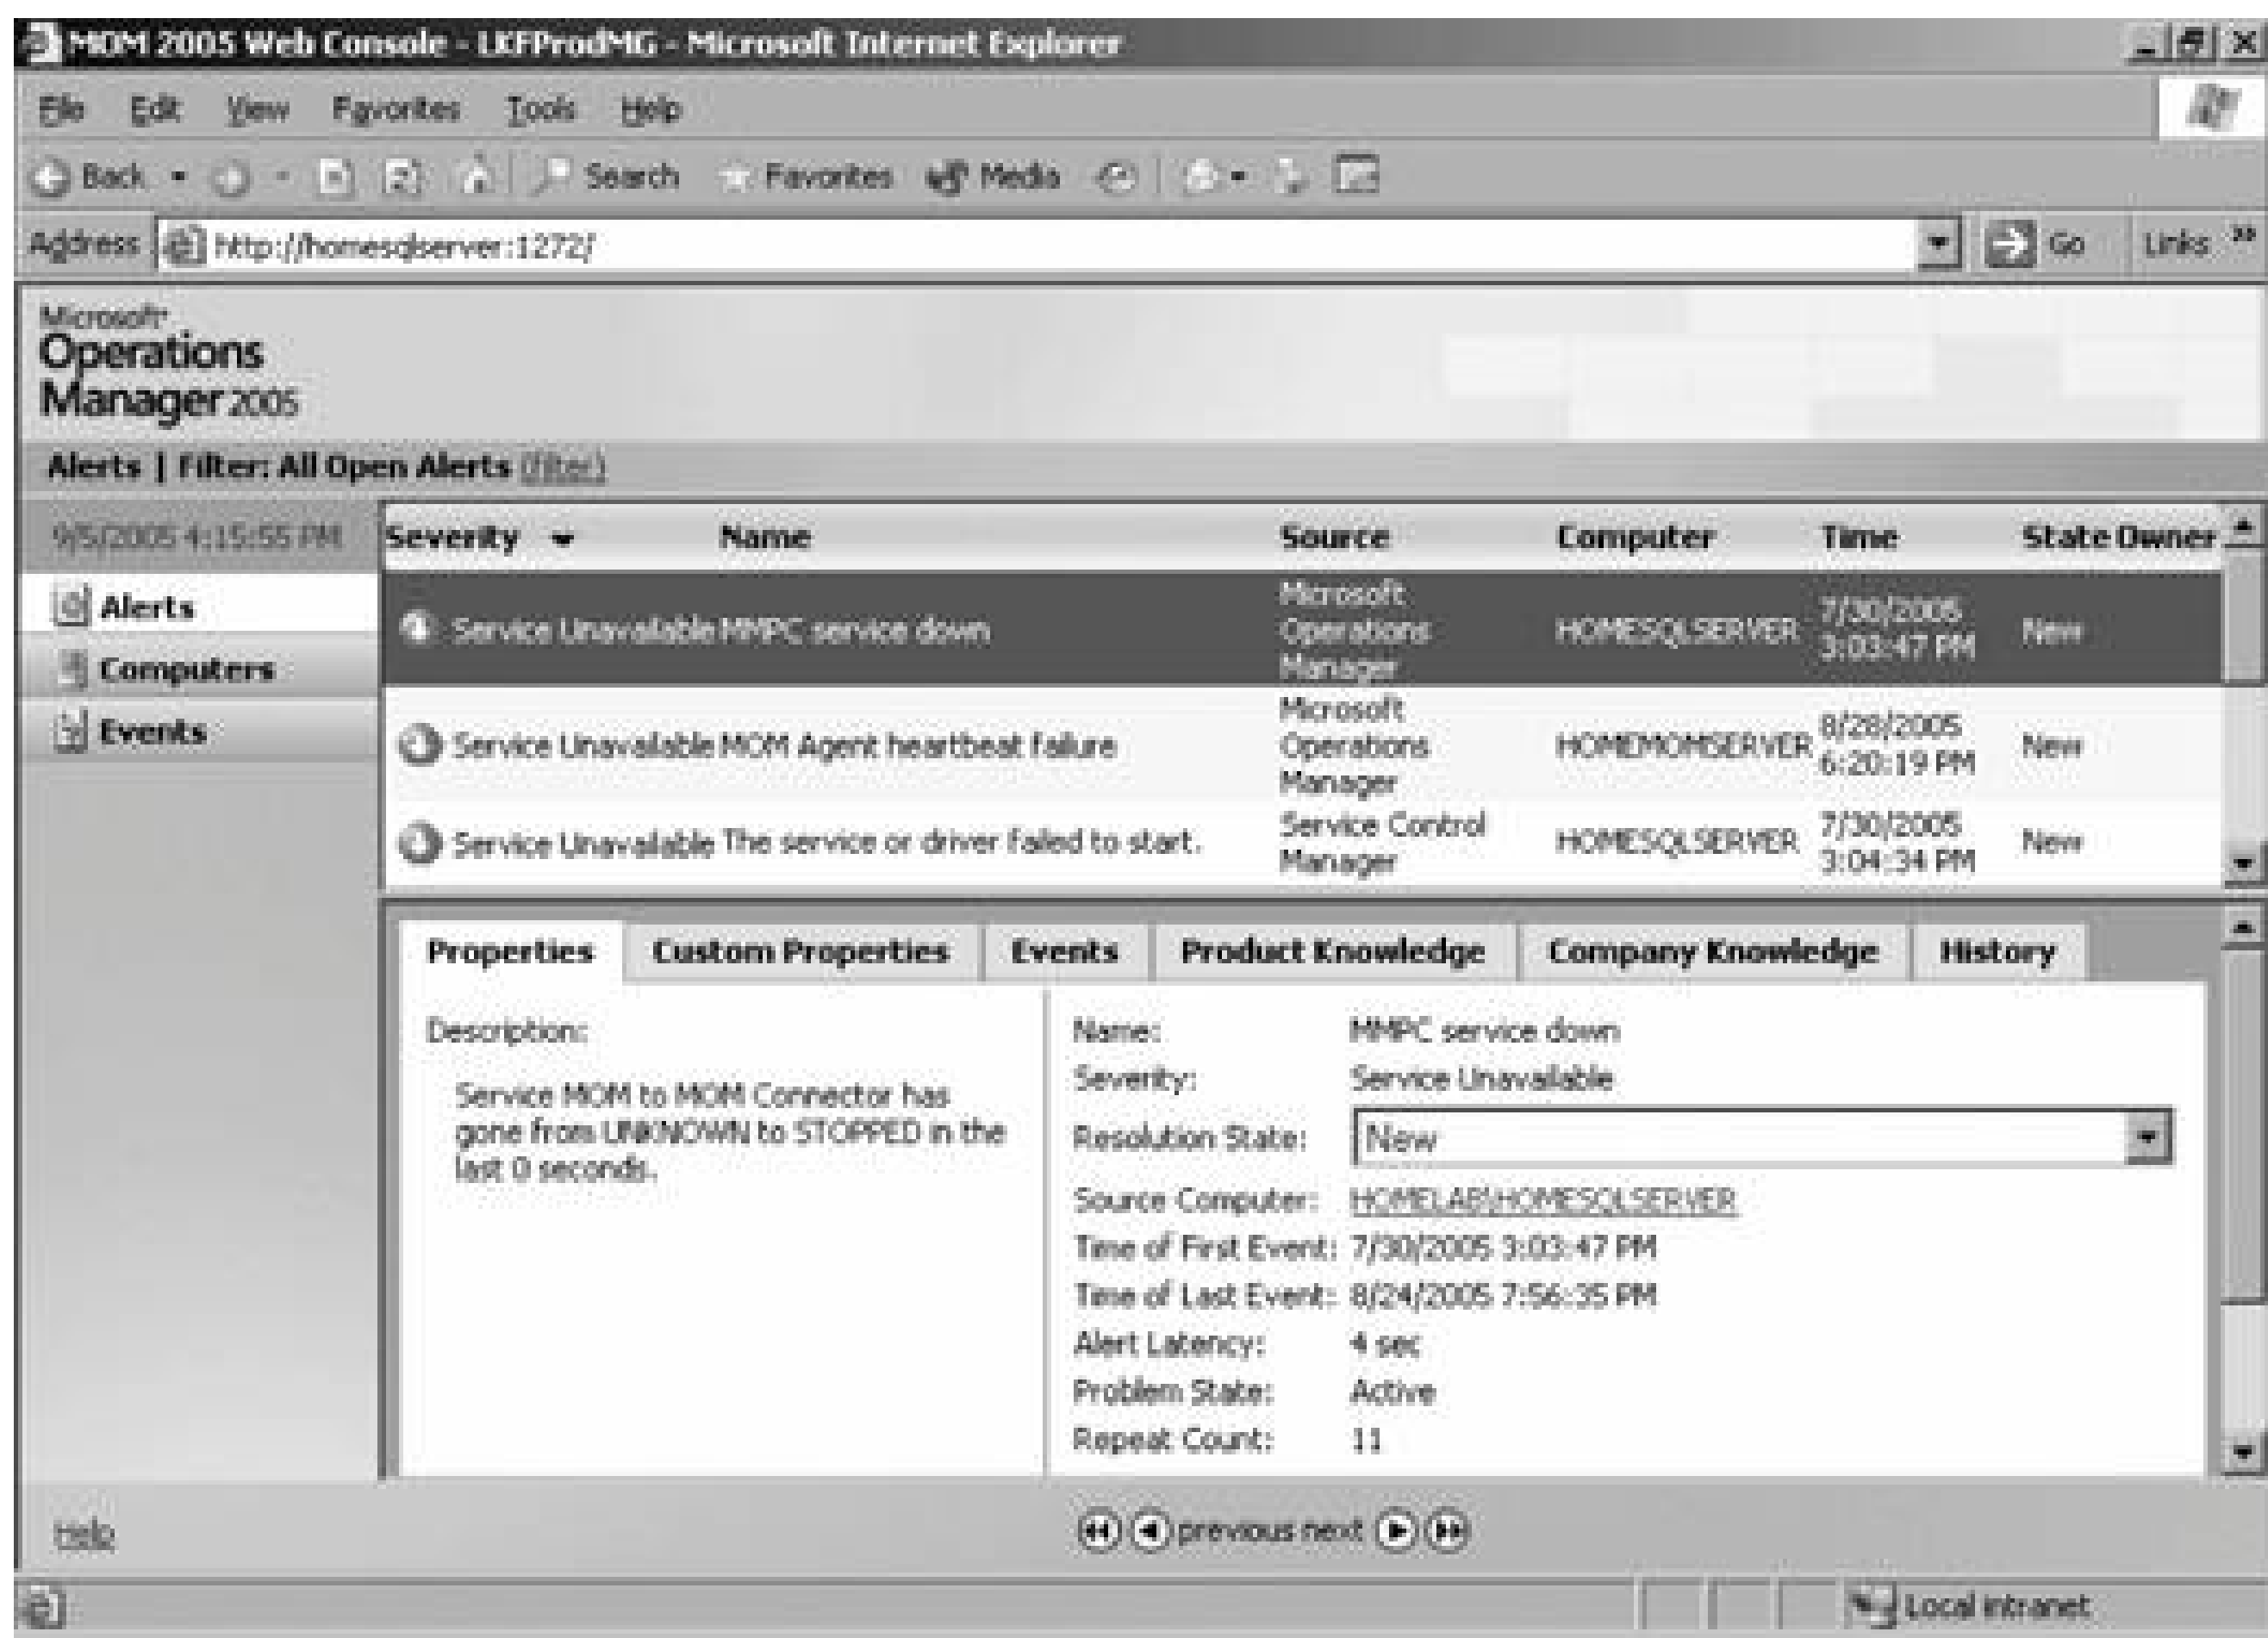
The Company Knowledge tab for each alert type starts out blank. Every time you resolve an alert, you will be prompted to record information about that alert. This is a useful way to capture the specific steps that you took to resolve the issue that caused the alert in the context of your environment. Start using this tab right away. Don't record historical information here because when this alert comes up again, you will want to know the steps for resolutions, such as "stop IIS and reboot server when this problem occurs," not "this happened because a contractor knocked over the equipment rack on the server floor that holds the Internet-facing router."

### 1.4.2. The Web Console

The Web console (point 1a in [Figure 1-11](#) and shown in [Figure 1-13](#)) displays only the information essential to managing your environment, alerts, computers in the managed environment, and Windows Events. You will be using this console when the Operator console is not available due to firewall issues or when working remotely. This will let you get enough information to determine if

further action is necessary.

Figure 1-13. MOM 2005 Web console



In this console, there are three panes: a views pane, a summary pane, and a details pane. The Web console communicates with the Data Access Service through the MOM 2005 Application Programming Interface (API) (point 2a in [Figure 1-11](#)). The MOM 2005 API is fully documented in the MOM 2005 SDK, available at <http://www.microsoft.com/mom/downloads/sdk/default.mspx>. The components of the API include:

- MOM Connector Framework Version 2
- MOM Managed Code Library
- MOM Runtime Library

### 1.4.3. The Data Access Server (DAS)

The DAS (point 2 in [Figure 1-11](#)) is the next component in the path back to the origin of an alert. The DAS is a Component Object Model (COM+) application that manages all read/write access to the operations database. The actions of the DAS enable the Operator console and the Web console to

render information to the user. Because the DAS manages all communications between the user interfaces, the database, and the MOM service on the management server, it is responsible for enforcing permissions inside of MOM. This is done via COM+ roles and impersonation. When you install MOM 2005, you will have to specify a Windows account that the DAS will use to access the operations database. This account must be assigned the db\_owner role in the MOM OnePoint database in SQL and have permit server access as a SQL server security login. Through the DAS and the DAS account, MOM 2005 manages the database by executing the stored procedure used for grooming the database tables, as well as data insert functions. The DAS communicates with the MOM 2005 service (*MOMService.exe*) over OLEDB.

## 1.4.4. The MOM Operations Database

MOM 2005 keeps two types of information in the operations database:

### *Configuration data*

All configuration data for a management group is stored in the operations database. Because the data is stored centrally, it can be accessed by all the management servers in the management group. Information stored here includes management pack rules and their thresholds, agent configuration settings, and the global settings for the management group, such as which email server to route SMTP traffic through, and security settings.

### *Operations data*

This is the live data gathered from the agents and includes all events, performance monitor data, and alerts.

Before an alert appears in the Operator or Web consoles, it is written to the operations database. If you look in SQL Enterprise Manager, you will not find a MOM database. The actual name of the operations database is OnePoint, which is a leftover from the original versions of MOM when it was developed and marketed by NetIQ.

The operations database is a SQL 2000 database (point 3 in [Figure 1-11](#)) that is best kept small for performance reasons. The largest operations database supported by MOM 2005 is 30 GB, of which you should keep 40 percent as free space. This free space is required for successful execution of the stored procedures, namely the reindexing job.

The DAS communicates with the operations database over TCP/UDP port 1433. This is where the data that composes an alert actually lives. All modifications to that data, either through the user interfaces or through the MOM service, are persisted here until they are groomed out.

As of the Release To Manufacturing (RTM) version of MOM 2005, the MOM product team does not have a clear plan regarding the use of SQL Server 2005 for either the Operations database or the Archiving database.



## 1.4.5. The MOM 2005 Service

On the management server, the MOM service plays two roles: it is the MOM server (point 4 in [Figure 1-11](#)) and the MOM agent (point 5 in [Figure 1-11](#)). Both of these processes run under the *MOMService.exe* service on the management server and under the security context of NT Authority/Network Service, Local System on Windows Server 2003, or Local System only on Windows 2000 servers. Running the *MOMService.exe* process under any other security context is not supported. When the agent on the management server needs to take some action, it spawns an instance of *MOMHost.exe* that runs under the management server action account. In MOM 2005, the agent can launch multiple instances of *MOMHost.exe*, which is responsible for executing scripts and running managed code responses. If one of the *MOMHost.exe* instances hangs, the others and the *MOMService.exe* are unaffected.

## 1.4.6. The MOM 2005 Server

The MOM 2005 server sits between the agents on the managed nodes, the agent on the management server, and the DAS. The MOM server communicates with the DAS using OLEDB calls, and with the agents on the managed nodes over TCP port 1270, RPC 135, and TCP/UDP 445.

The primary responsibility of the MOM server is to manage communications with the agents on the managed nodes. It sends configuration information down to the agents and receives operational data from them. Working with the management server agent, it consolidates the data from the agents on managed nodes and passes it to the DAS for insertion into the operations database. It is also responsible for computer discovery and pushing agents to computers. Alerts travel through the MOM server and can be modified there based on actions taken by the management server agent.

## 1.4.7. The Management Server Agent

The management server is itself a managed node, and the management server agent is responsible for collecting all the event and performance data from the MOM management server. It then compares this data (event log, performance monitor, WMI data, etc.) to a set of criteria to determine if there is a match. If there is a match, the agent can execute a response. When an agent needs to execute a response, it spawns an instance of *MOMHost.exe* running under the management server's action account credentials. Responses include generating an alert, or running a script against the management server itself or against a managed node. The agent can generate an email, transfer a file, or execute managed code as well.

What is unique about this agent is that it also uses the data stream coming from the managed nodes as one of its data providers. This gives it the ability to correlate alerts and events coming from multiple managed nodes and generate new alerts that reflect a significant event across a wider set of machines. This is precisely the case for the example "MOM Agent heartbeat failure" alert. The management server agent is expecting a heartbeat message from every remote agent (which run over UDP) at 15-second intervals. When it does not receive one, it generates the example alert.

The other responsibility that the management server agent has is the administration of remote agents and the monitoring of agentless-managed computers. In this capacity, the management server agent uses a *MOMHost.exe* instance to install agents remotely on discovered computers, run



discovery tasks on remote computers, and update settings on remote agents. In the case of agentless-managed computers, this agent performs discovery of the computer's role, remotely collects information from the computer, applies a set of criteria against the incoming data, and initiates responses based on matches to the criteria. Because it is performing these tasks remotely, there is a performance hit on the management server that needs to be planned for. Microsoft does not recommend performing agentless management against more than 60 machines per management group.

## 1.4.8. MOM Agent

The MOM agent (point 5 in [Figure 1-11](#)) represents a remote agent. This is an agent that resides on a computer other than the management server. This agent receives all its configuration information from the MOM 2005 management server and sends its processed and filtered data to the MOM server on the management server. Just like the MOM agent on the management server, it creates *MOMHost.exe* processes for collecting data from the data providers. It then applies the criteria to this data and executes the appropriate response if a match is found, such as generating an alert. Below this are the data providers that you would manually examine for clues to the cause of an issue, if you don't have an operations management system in place. Agents perform the bulk of the work in MOM 2005 and hopefully send only actionable information up to the management server.

## 1.4.9. Management Packs and Processing Rule Groups

So, at this point you are asking yourself, how does an agent know what criteria to apply to a set of data and what response to take? Rules applied by agents are imported into the management group in files called *management packs*.

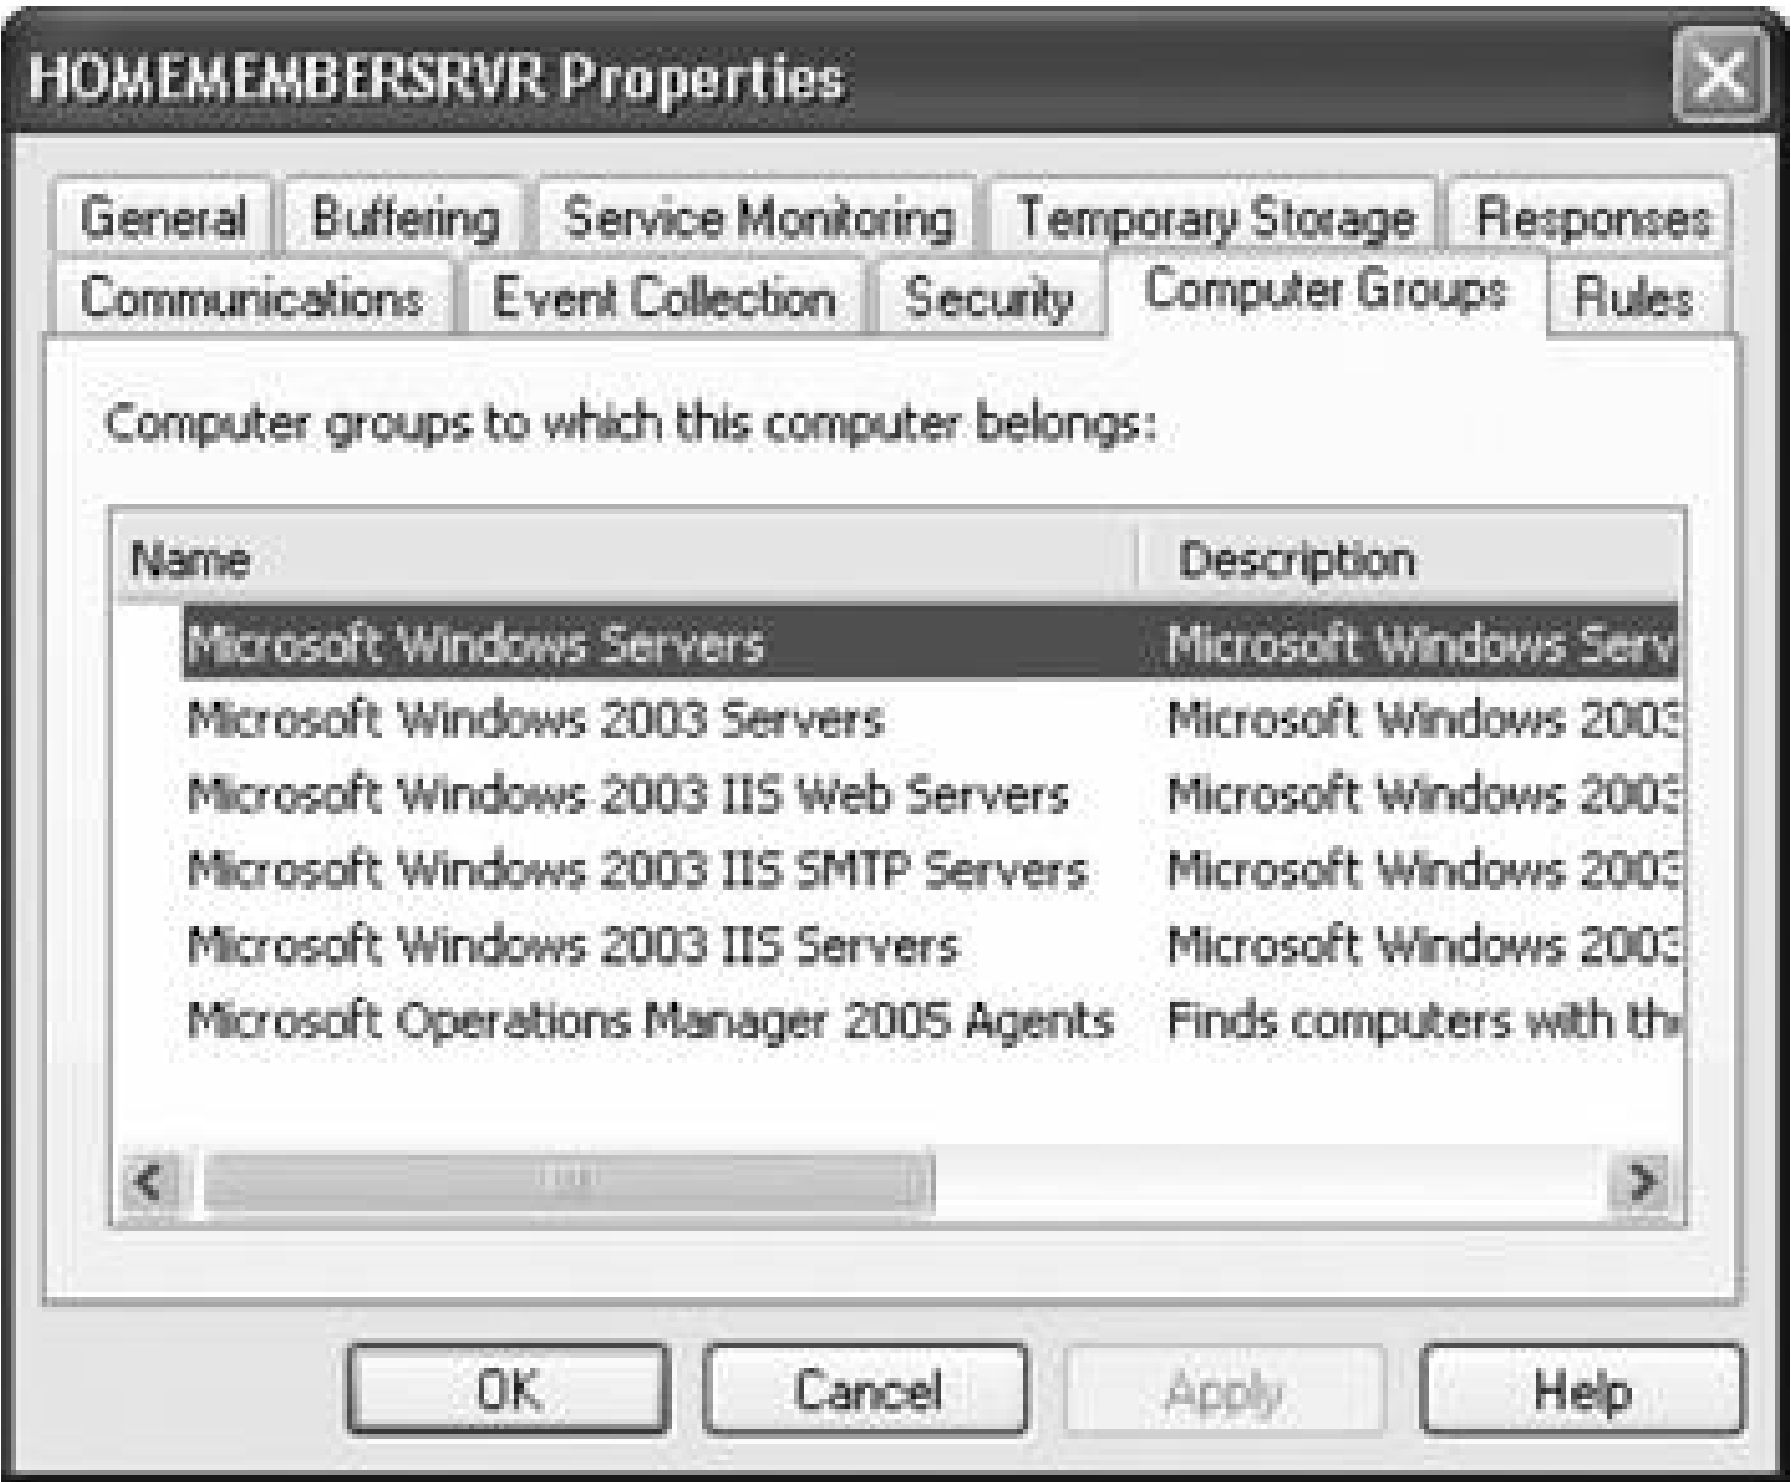
Every server-based application that Microsoft releases has a management pack. There are management packs for operating systems, Exchange 2000 and 2003, MOM 2005, and SMS 2.0 and 2003, to name a few. At the time of this writing, Microsoft is shipping 55 different management packs, with more being added.

In each management pack are application definitions that the agents use to identify the role of a computer, such as it being a domain controller or a MOM server or an Internet Information Server (IIS). The agents determine this based on the discovery process that they execute. The discovery process looks at a number of computer attributes such as registry values and the existence of certain files, directories, and registered services. Once a computer role has been identified, that computer is placed in a computer group inside of MOM.

A computer group in MOM is not a computer security group in Active Directory and is used only inside of MOM. Membership in MOM computer groups is dynamic, based on the discovered role of a computer.

Computers always belong to more than one MOM computer groupthis is normal. For example, *homemembersrvr* in [Figure 1-14](#) belongs to six computer groups.

Figure 1-14. MOM computer group membership for a Windows 2003 member server running IIS



In each management pack there are also three groupings of rules:

### Event rules

These rules tell an agent to collect information from various event logs, to filter out certain types of alerts, to look for missing events, to consolidate multiple alerts with similar characteristics, and to suppress duplicate alerts. [Figure 1-15](#) shows an event rule.

### Performance rules

These rules instruct an agent to either sample and report specific performance monitor data or to generate an alert when the performance monitor data cross over a threshold defined in the rule.

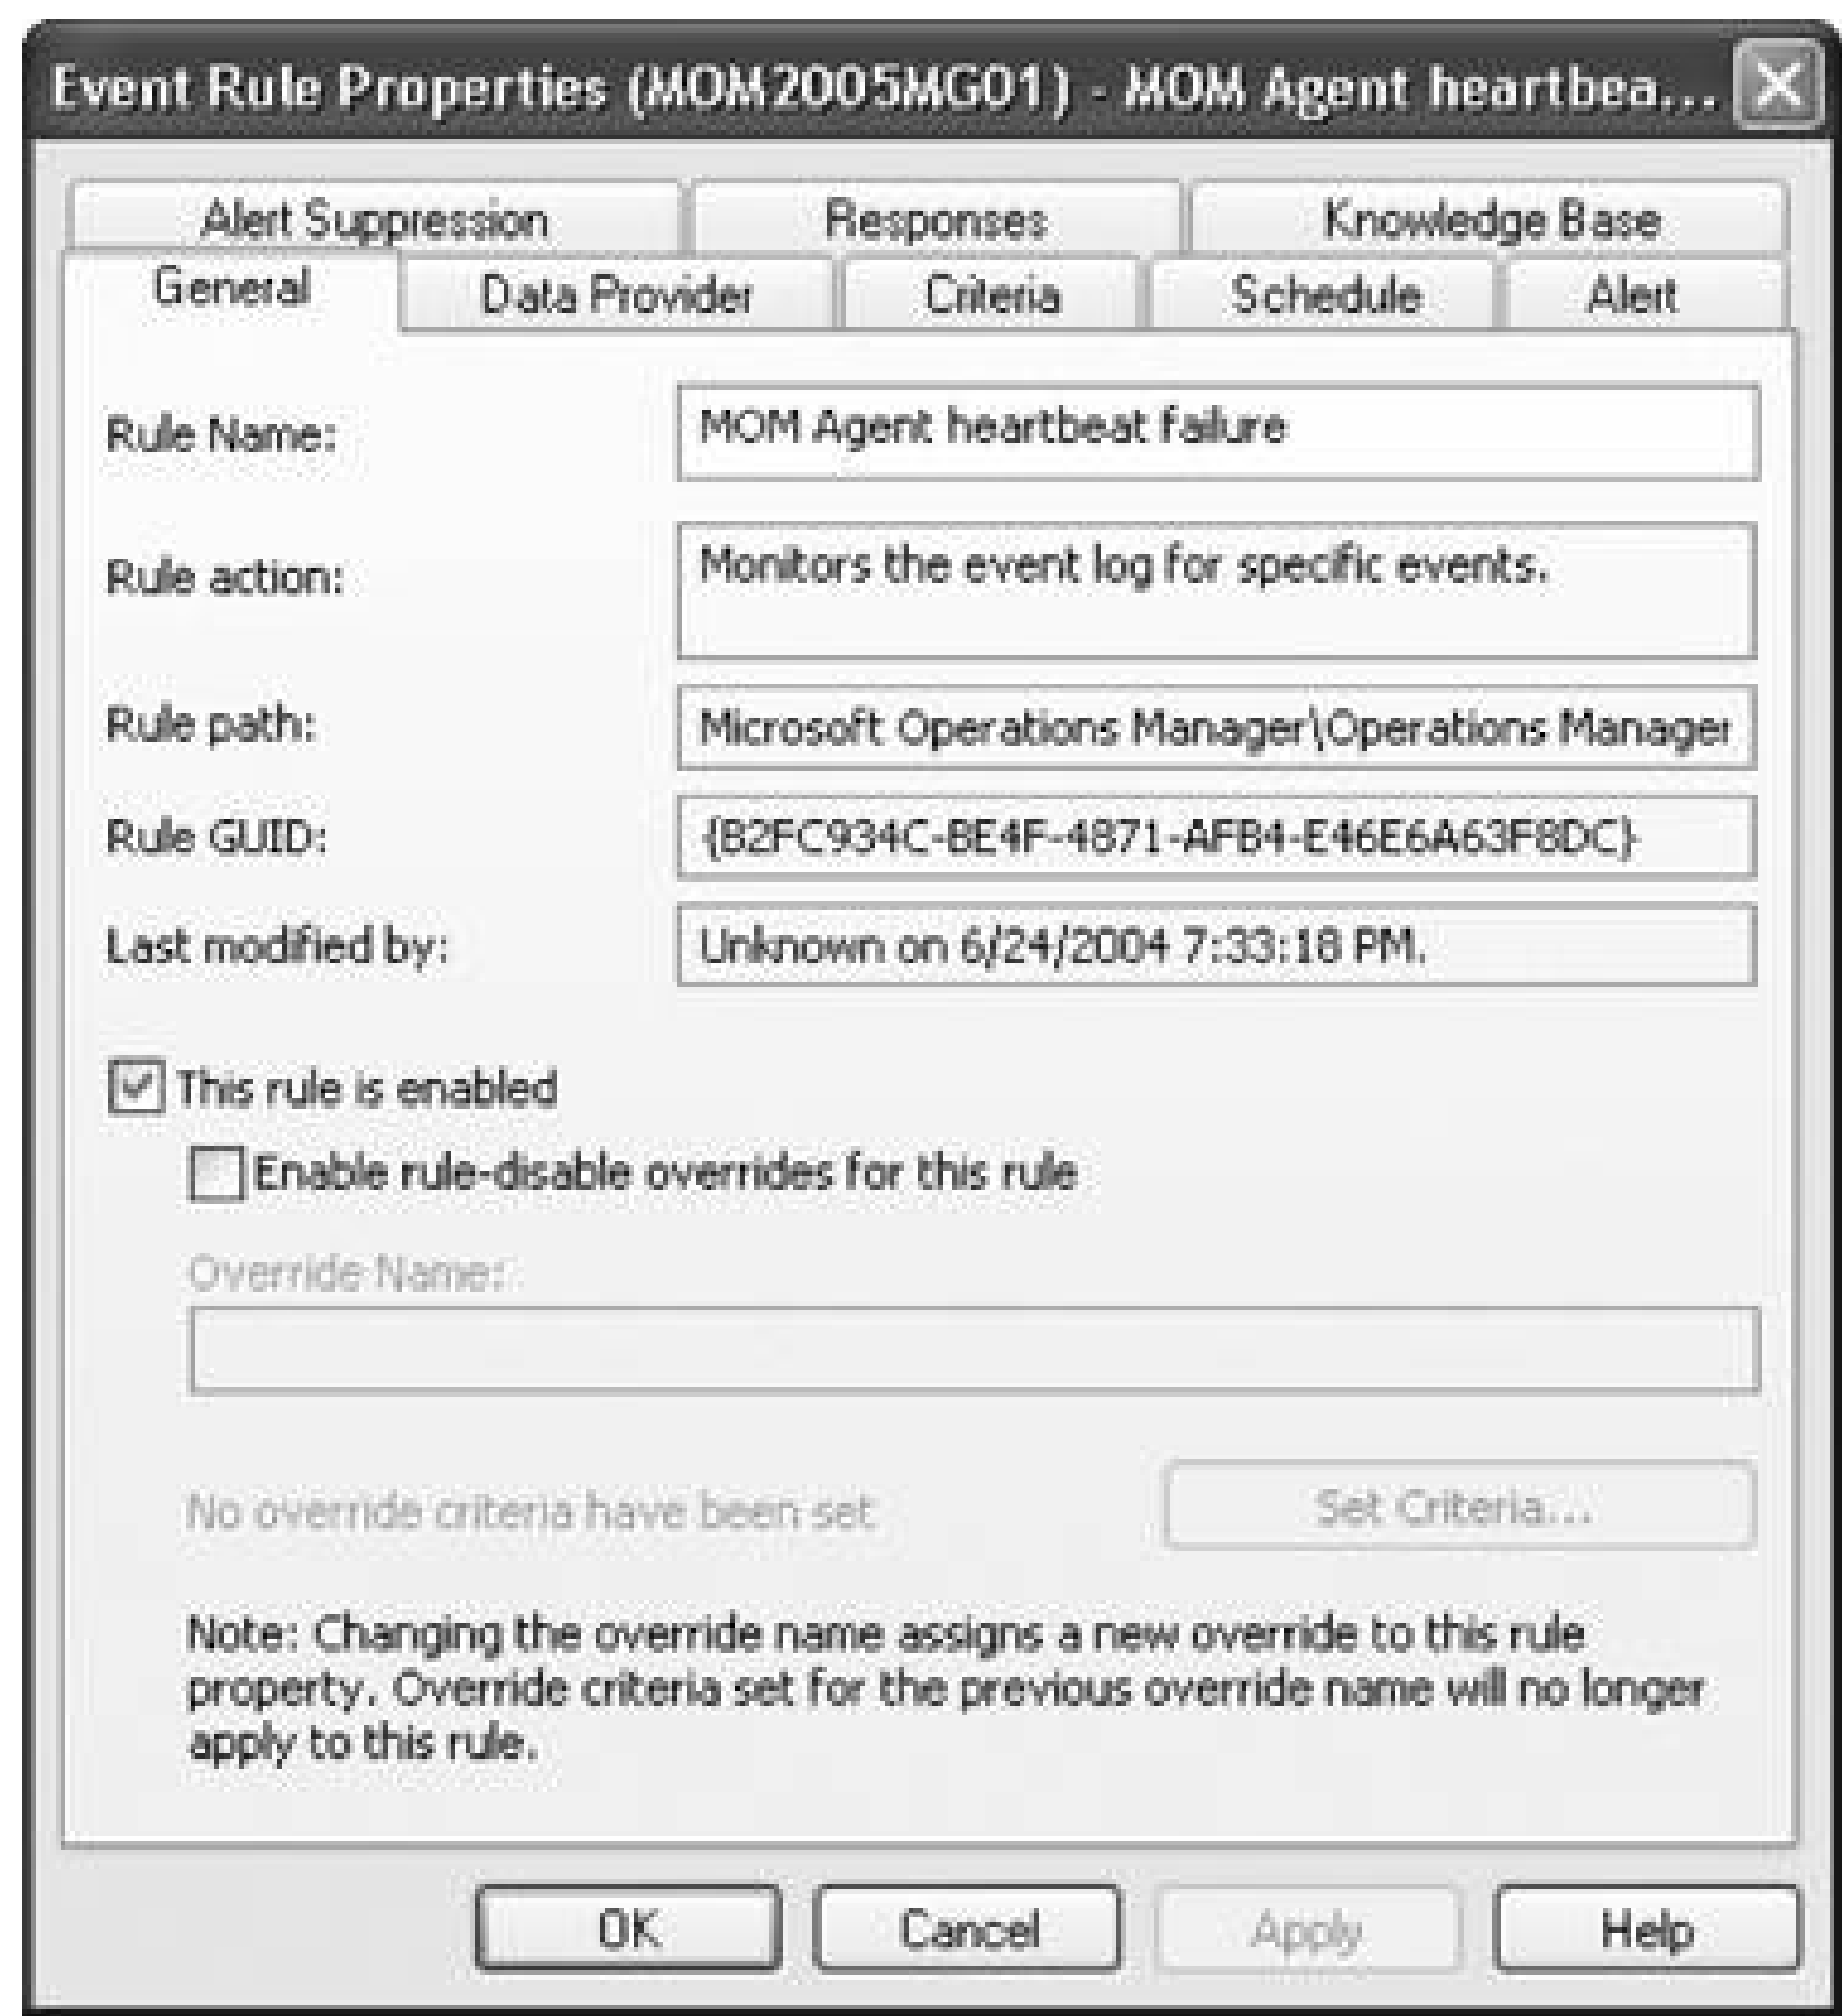
### Alert rules

Rather than having each event or performance rule generate its own alert, a single alert rule can instruct an agent to generate a response for an entire group of rules.

Each rule identifies the provider from which data will be examined, the criteria to look for, the response to take when a match between the raw data and the criteria is found, and the vendor and company knowledge. This is where the text on the Product Knowledge tab of an alert in the Operator console comes from. The rules are defined by the product teams themselves and as such represent their definition of health for the application. [Figure 1-15](#) shows the General tab of the "Agent

heartbeat failure" rule.

Figure 1-15. The General tab of the "MOM Agent heartbeat failure" event processing rule



Each computer group can have one or more rule groups associated with it. It is the responsibility of the MOM server to assign the correct set of rules to each agent it manages, based on the role of the computer. Only those rules that are associated with a computer group to which a computer belongs are sent to the agent on that computer. This is how MOM ensures that an agent is only processing the necessary rules for the computer it monitors.

The last items included in a management pack are definitions for the views in the Operator console and predefined reports that will be viewed in the Reporting console.

### 1.4.10. Data Providers

You are already familiar with the data providerson a Windows server (point 7 in [Figure 1-11](#)). These are the tools you examine manually now. They include the event logs and performance monitor counters and data collected from WMI and application logfiles. This is the raw data used to generate an alert, but you won't find alerts here.

Data provider objects are defined in management packs. An individual provider can be used by many rules. For example, event logs are defined as individual providers and are used as the data source or provider by many event rules.





## 1.5. Additional Components

The MOM 2005 components introduced so far are directly involved with the creation of an alert, the management of an alert, and the storage of an alert. But these are not all the components available. This section introduces other components and gives an overview of where they fit and their purpose. [Figure 1-16](#) shows these additional components; however, note that none of these are involved with the generation of an alert.

### 1.5.1. Administrator Console

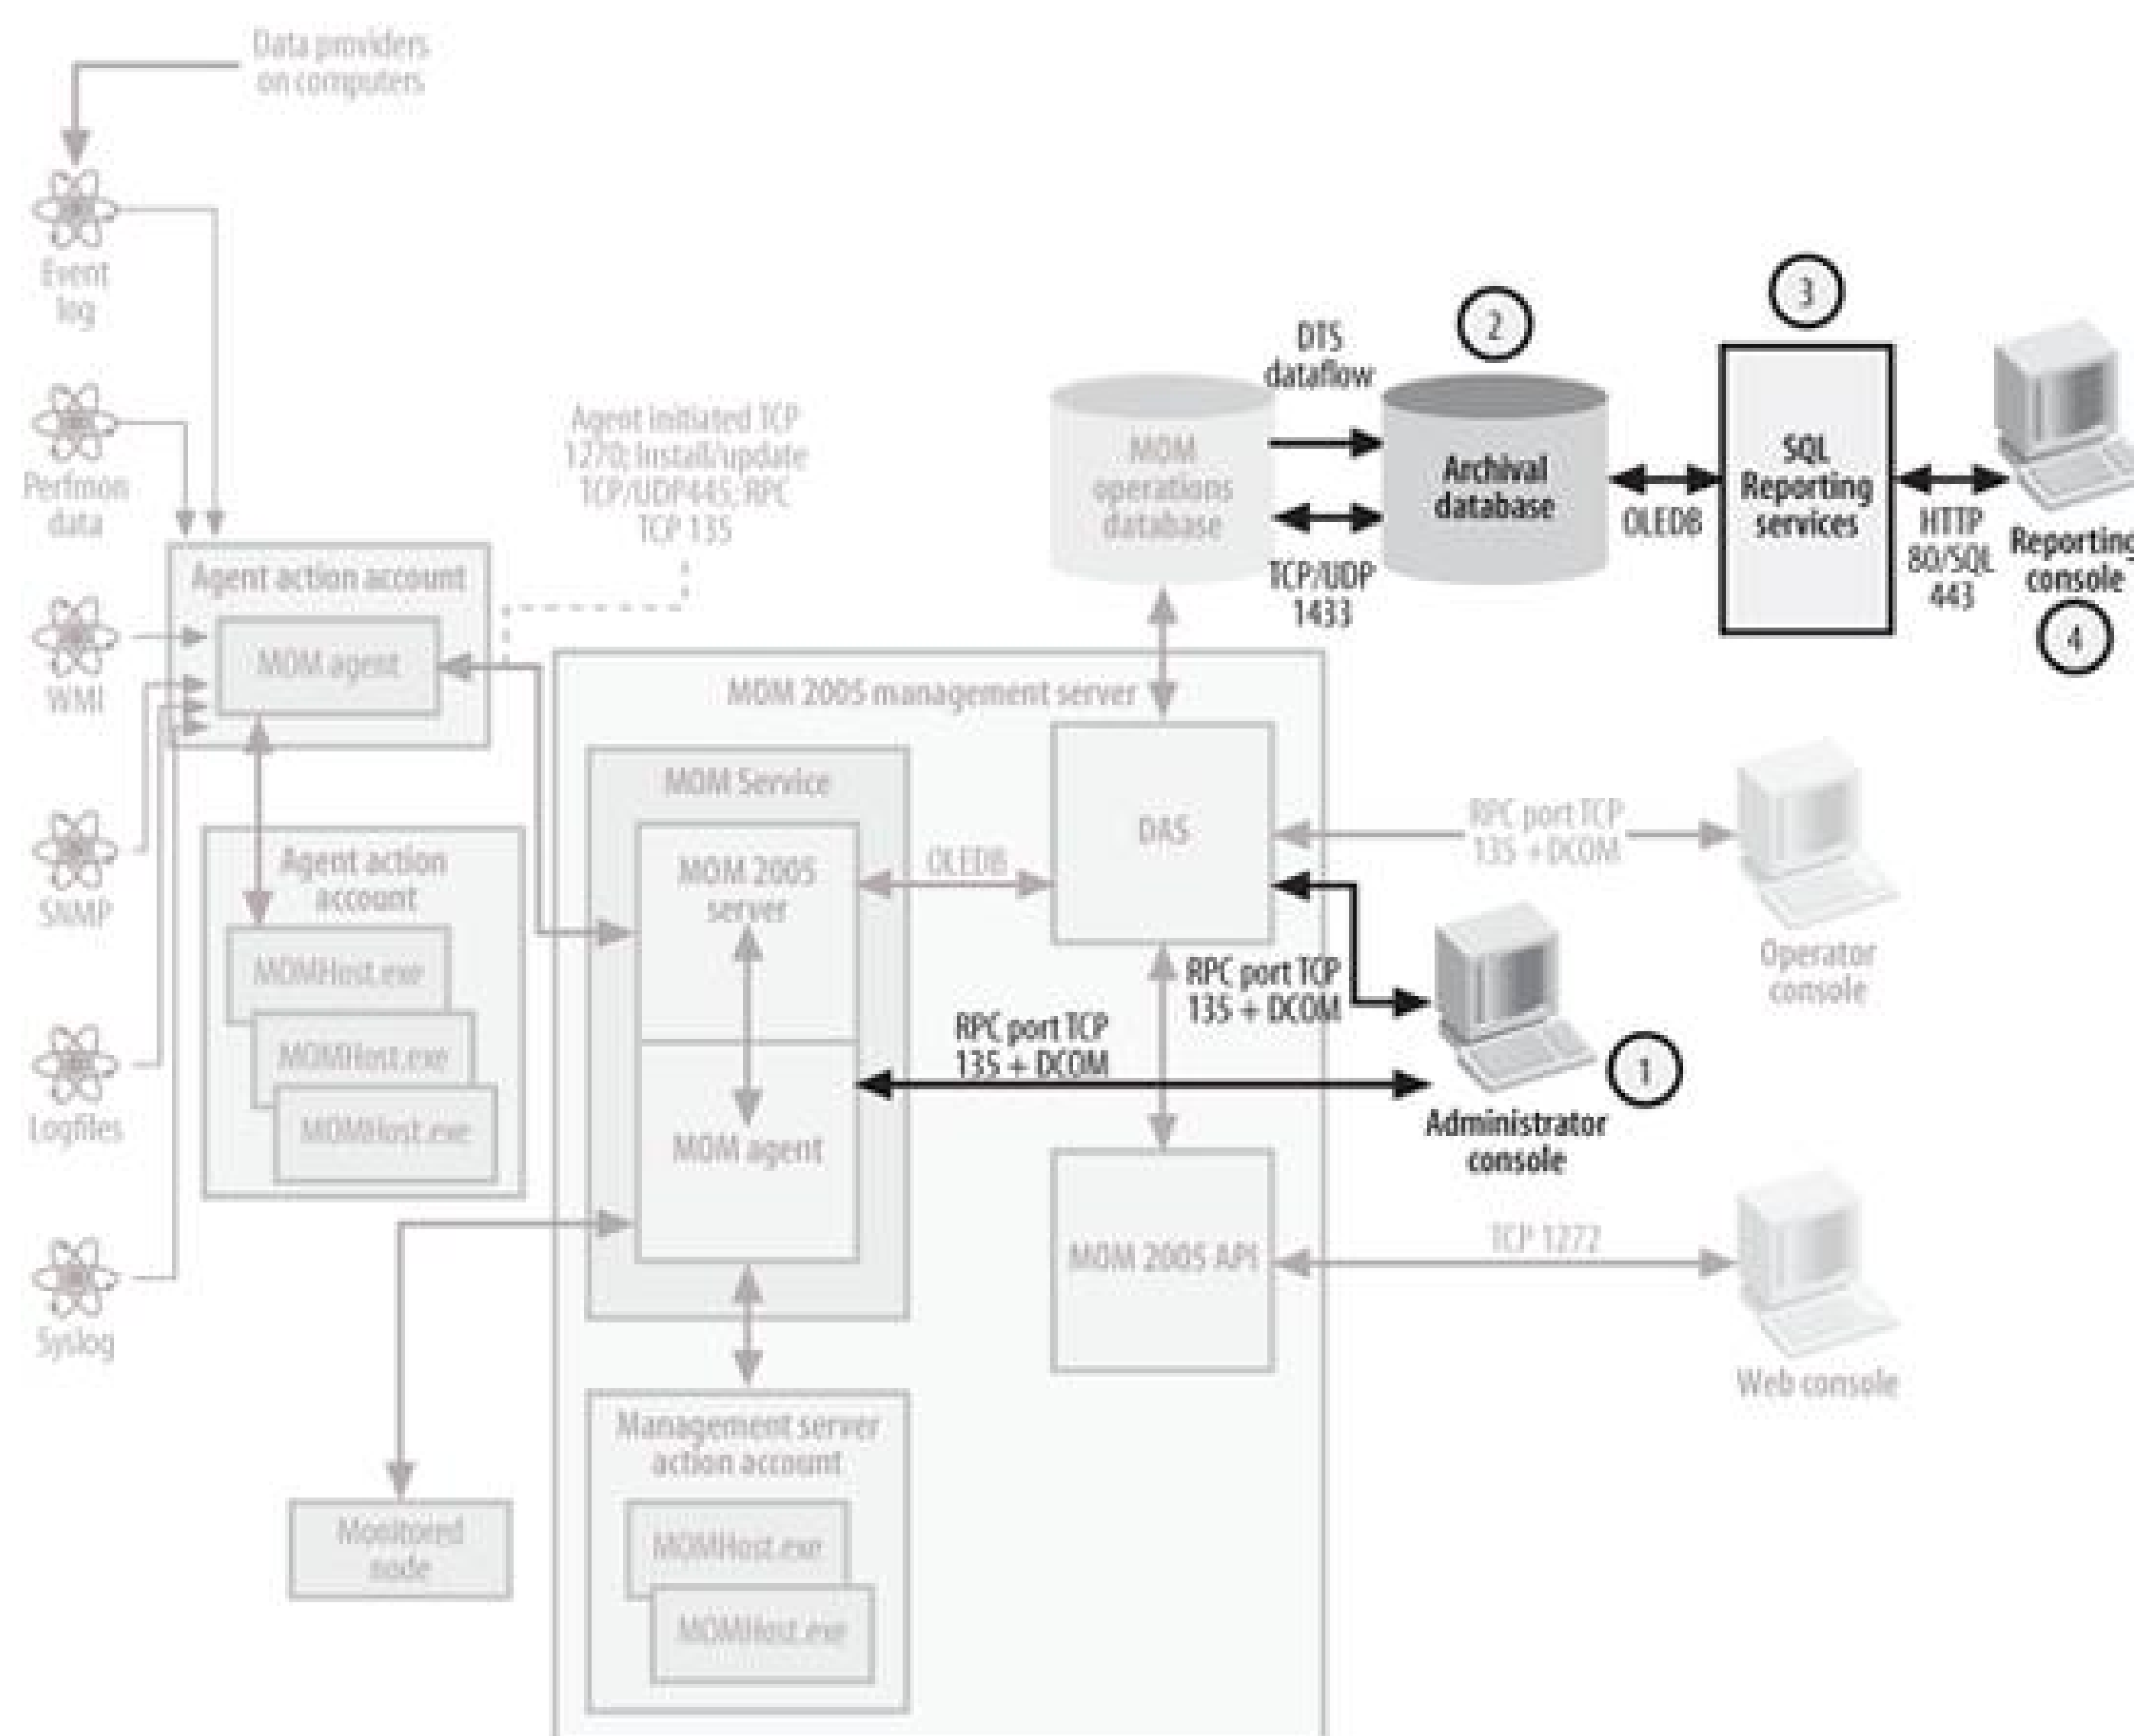
While the Operator console is used to consume and manage the alert, event, performance, and status data that MOM 2005 produces, the Administrator console (point 1 in [Figure 1-16](#)) is used to manage the configuration of MOM 2005. The Administrator console performs most of its work through the DAS on the MOM 2005 management server, but for some specific tasks, it works directly with *MOMService.exe*. Like the Operator console, the Administrator console is a Microsoft Management Console (MMC) snap-in.

There are four top-level objects used to break down the functions in the Administrator console. They are the Information Center object; the Operations object used for launching the Operators, Reporting, and the Web consoles; the Management Packs object; and the Administration object. [Figure 1-17](#) shows the Administrator console.

If you have administrator permissions to MOM 2005, you will have full access to all objects in the Administrator console. A MOM 2005 administrator controls the configurations that govern the overall behavior of MOM 2005, including:

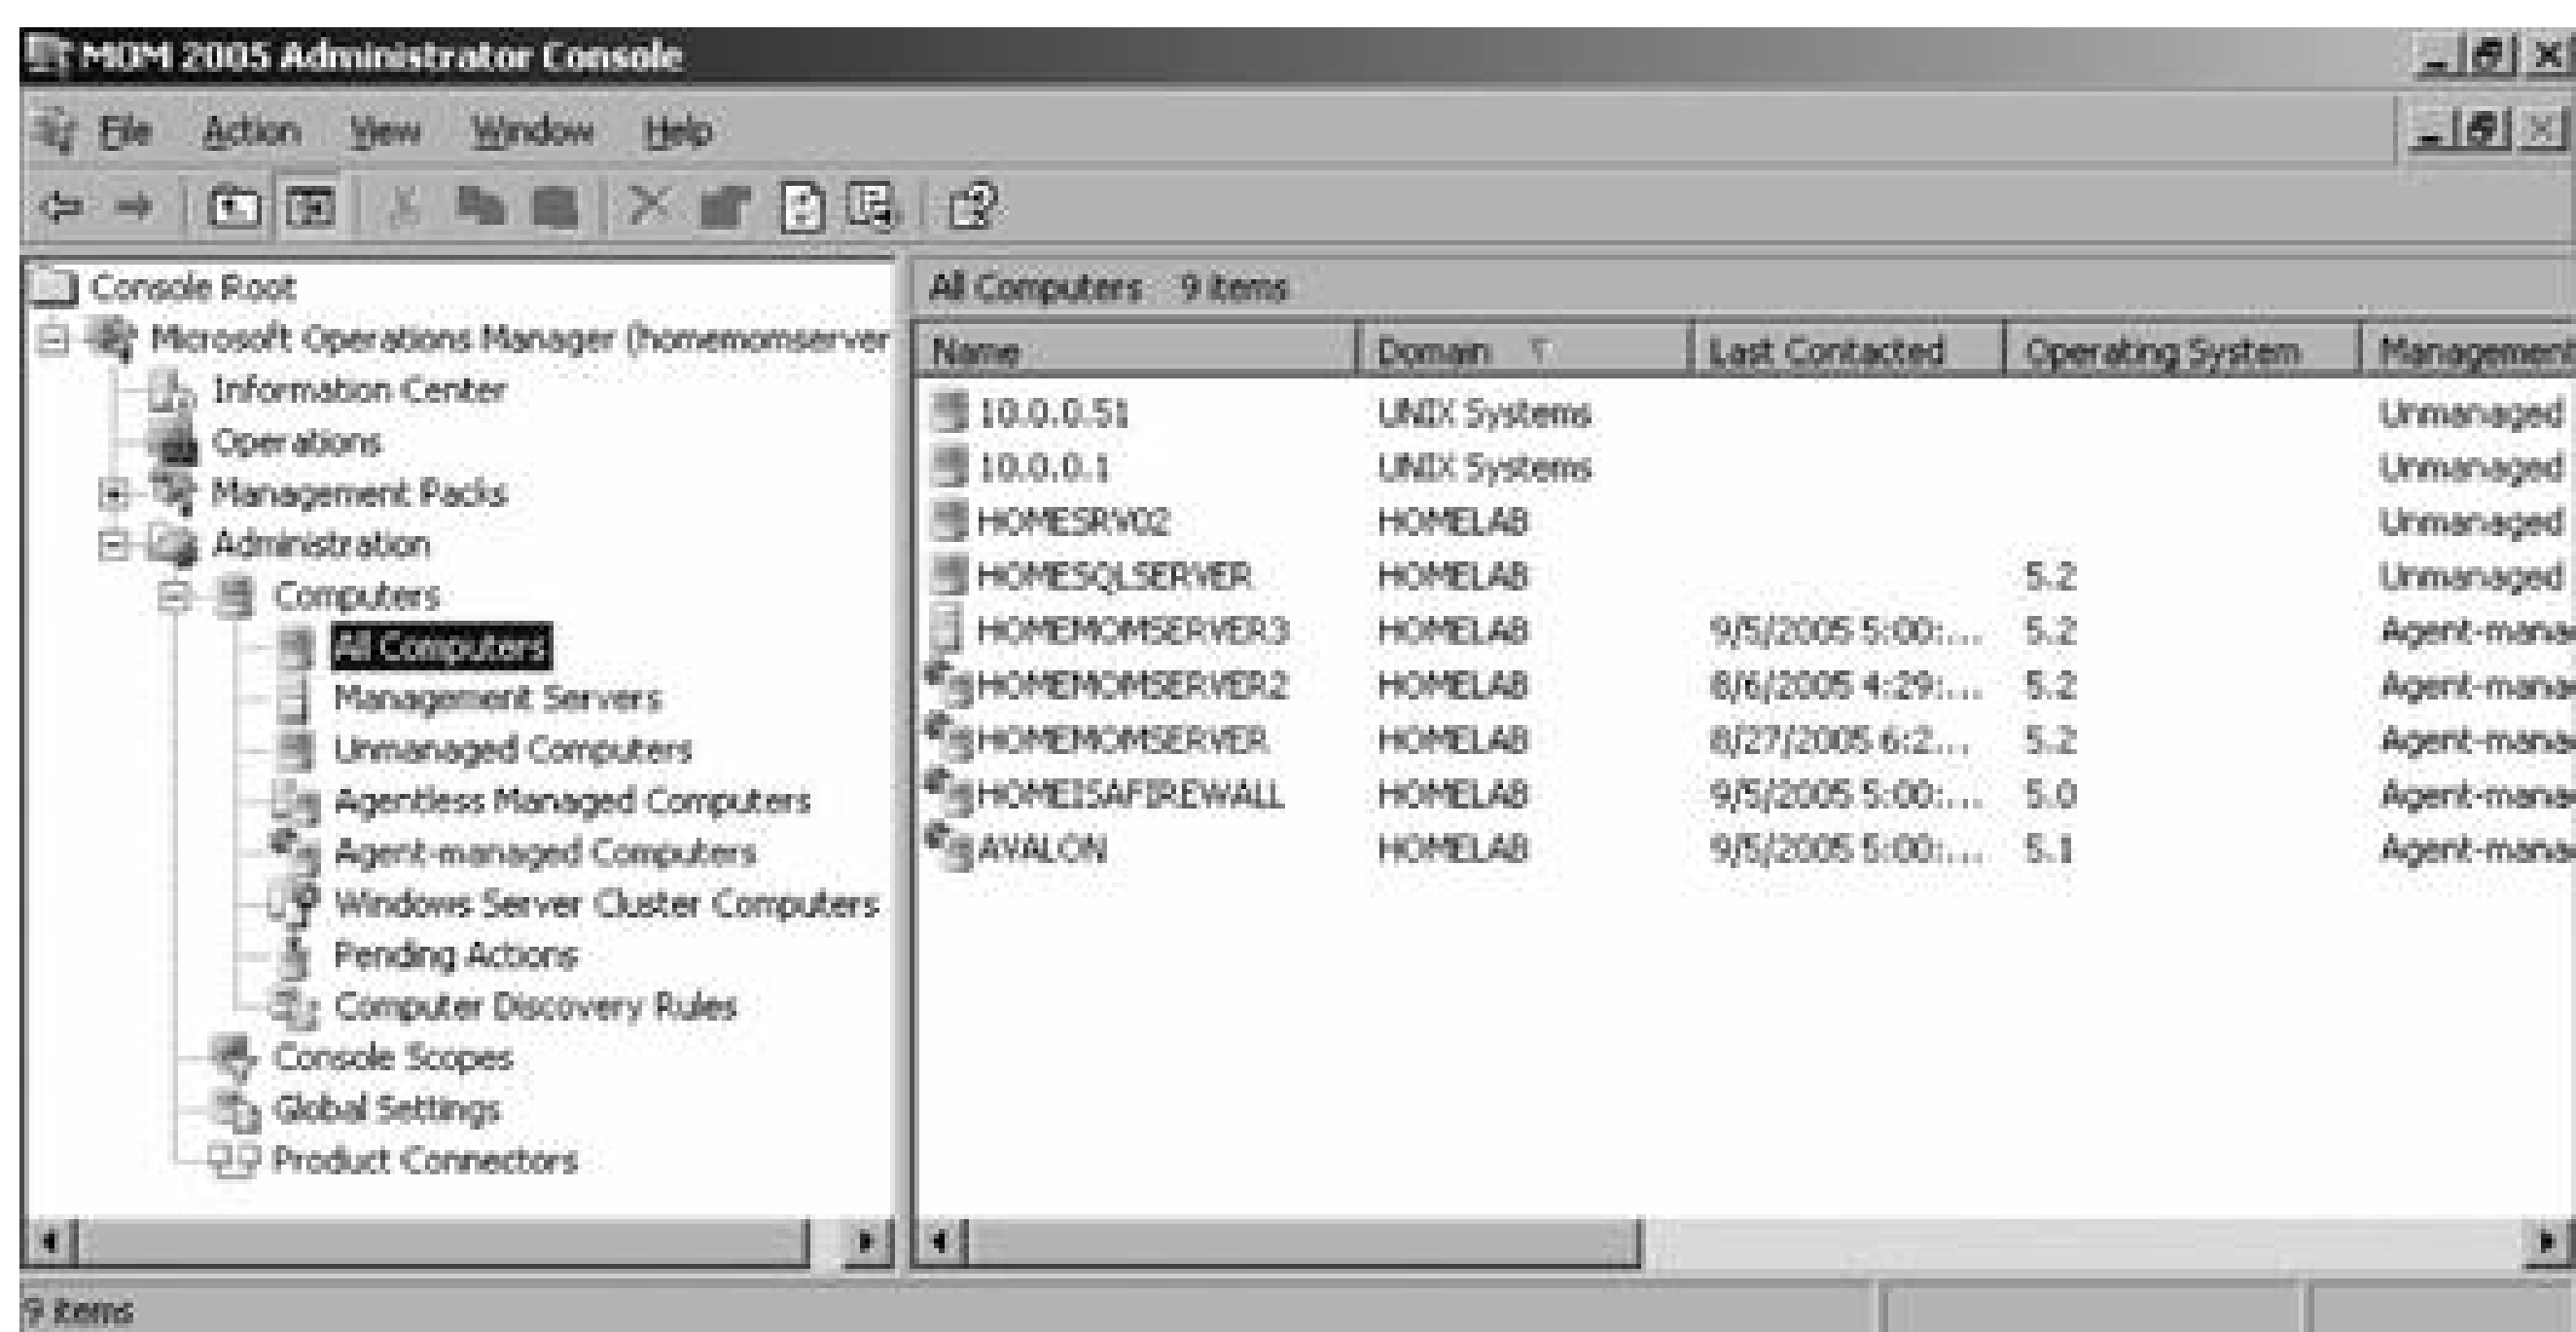
- Deciding which computers the management server will manage via a set of discovery rules, e.g., all computers in the homelab domain or only those computers whose names start with XYZ.
- Installing and uninstalling agents.

Figure 1-16. Additional MOM 2005 components



- Deciding the type of management to be applied to a managed node: agent-managed, agentless-managed, or unmanaged.
- Creating and controlling console scopes that filter information in the Operator console based on the user's Active Directory ID.
- Controlling the global settings for all of the management servers that operate together. This includes database grooming, managing security settings, determining which email server to send notifications through, defining alert resolution states, defining custom fields for alerts, and deciding the communications ports to be used.
- Creating and managing product connectors that are objects built using the MOM Connector Framework (MCF) to allow bidirectional communication between groups of MOM servers and other operations management products.

Figure 1-17. Top level of the Administrator console



Only individuals who have been granted administrative permissions to MOM 2005 can access these objects.

The MOM 2005 Administrator console is also used to manage the MOM 2005 management packs. People who are responsible for maintaining the role-based computer groups and the sets of rules that are pushed down to the agents will use this node. These individuals are called MOM authors. Administrators also have full access to this node.

MOM 2005 administrators must have in-depth knowledge of how MOM 2005 works. This does not mean that they must be experts on all the applications and server types that MOM can monitor. Usually, each application that MOM monitors has an associated group of experts at your company who are responsible for administering that application. These are the folks who receive the notifications from MOM when something goes wrong with that application and are responsible for fixing it. Give these application experts full authority to manage their applications rules by making them MOM authors. Otherwise, the MOM administrator will be an unnecessary middleman when it's the application administrators who really need to tune the monitoring of their specific management packs.

In the management packs node, MOM authors will manage:

- Import/export and creation of management packs
- Role-based computer group membership
- Event, performance, and alert rules for their applications
- Performance and event rule thresholds and criteria to override default settings for selected computer groups
- Tasks available for a computer role type in the Operator console
- Who gets notified and how they are notified of an application-specific alert



- Scripts that are used to gather information and execute responses on a managed node
- Computer attributes that are used to classify the role of a discovered computer
- Data Provider definitions, for example, the definitions for performance counter objects and counters that are used by performance rules

MOM authors also have the necessary rights to open the Operations object to launch the Operator, Reporting, or Web consoles. MOM authors cannot access the Administration object.

There is no need to grant administrative or authoring permissions to individuals who will only resolve alerts. These are typically the help desk and datacenter employees who are the recipients of MOM 2005 notifications. For them, access to the Administrator console is restricted to the operations object. From there, they can only launch the Operator, Web, or Reporting consoles, according to the rights they have been granted.

When you select the Information Center object, the Details pane brings up links to MOM 2005-specific web sites. From here, you can access the MOM 2005 homepage, MOM 2005 updates, the product documentation, community web sites, technical support, licensing, and security information.

## 1.5.2. Warehouse Database

While the operations database supports the monitoring, alerting, and configuration functions of MOM 2005, the data warehouse database (point 2 in [Figure 1-16](#)) supports the reporting function. MOM 2005 reports are based on historical data and give you a longitudinal view of the performance of your monitored applications and the servers they run on. A scheduled task runs once a day by default, which transfers the "live" data from the operations database to the data warehouse database. The dataflow in the transfer is unidirectional, going from the operations database to the data warehouse database only. The transfer is accomplished by a SQL Server DTS package and is coordinated with the grooming jobs on the operations database to ensure that no data is groomed out of operations before it has been transferred.

## 1.5.3. SQL Server Reporting Services

MOM 2005 reporting makes use of SQL Server Reporting Services. SQL Server Reporting Services (point 3 in [Figure 1-16](#)) is a separate product from MOM 2005. SQL Server Reporting Services separates the presentation of the data from the retrieving of the data, thereby allowing you to present the same data in different formats, including HTML, XML (for Office 2003 applications), and PDF. SQL Server Reporting Services does not include a printing engine, so you must export a report to one of these other formats before printing.

The reporting services in MOM 2000 SP1 were based on a Microsoft Access frontend. They did not scale well and did not reliably deliver scheduled reports. Because the reporting engine for MOM 2005 is server-based, it has the flexibility and robustness that its predecessor lacked. This is one of the most improved features in MOM 2005.

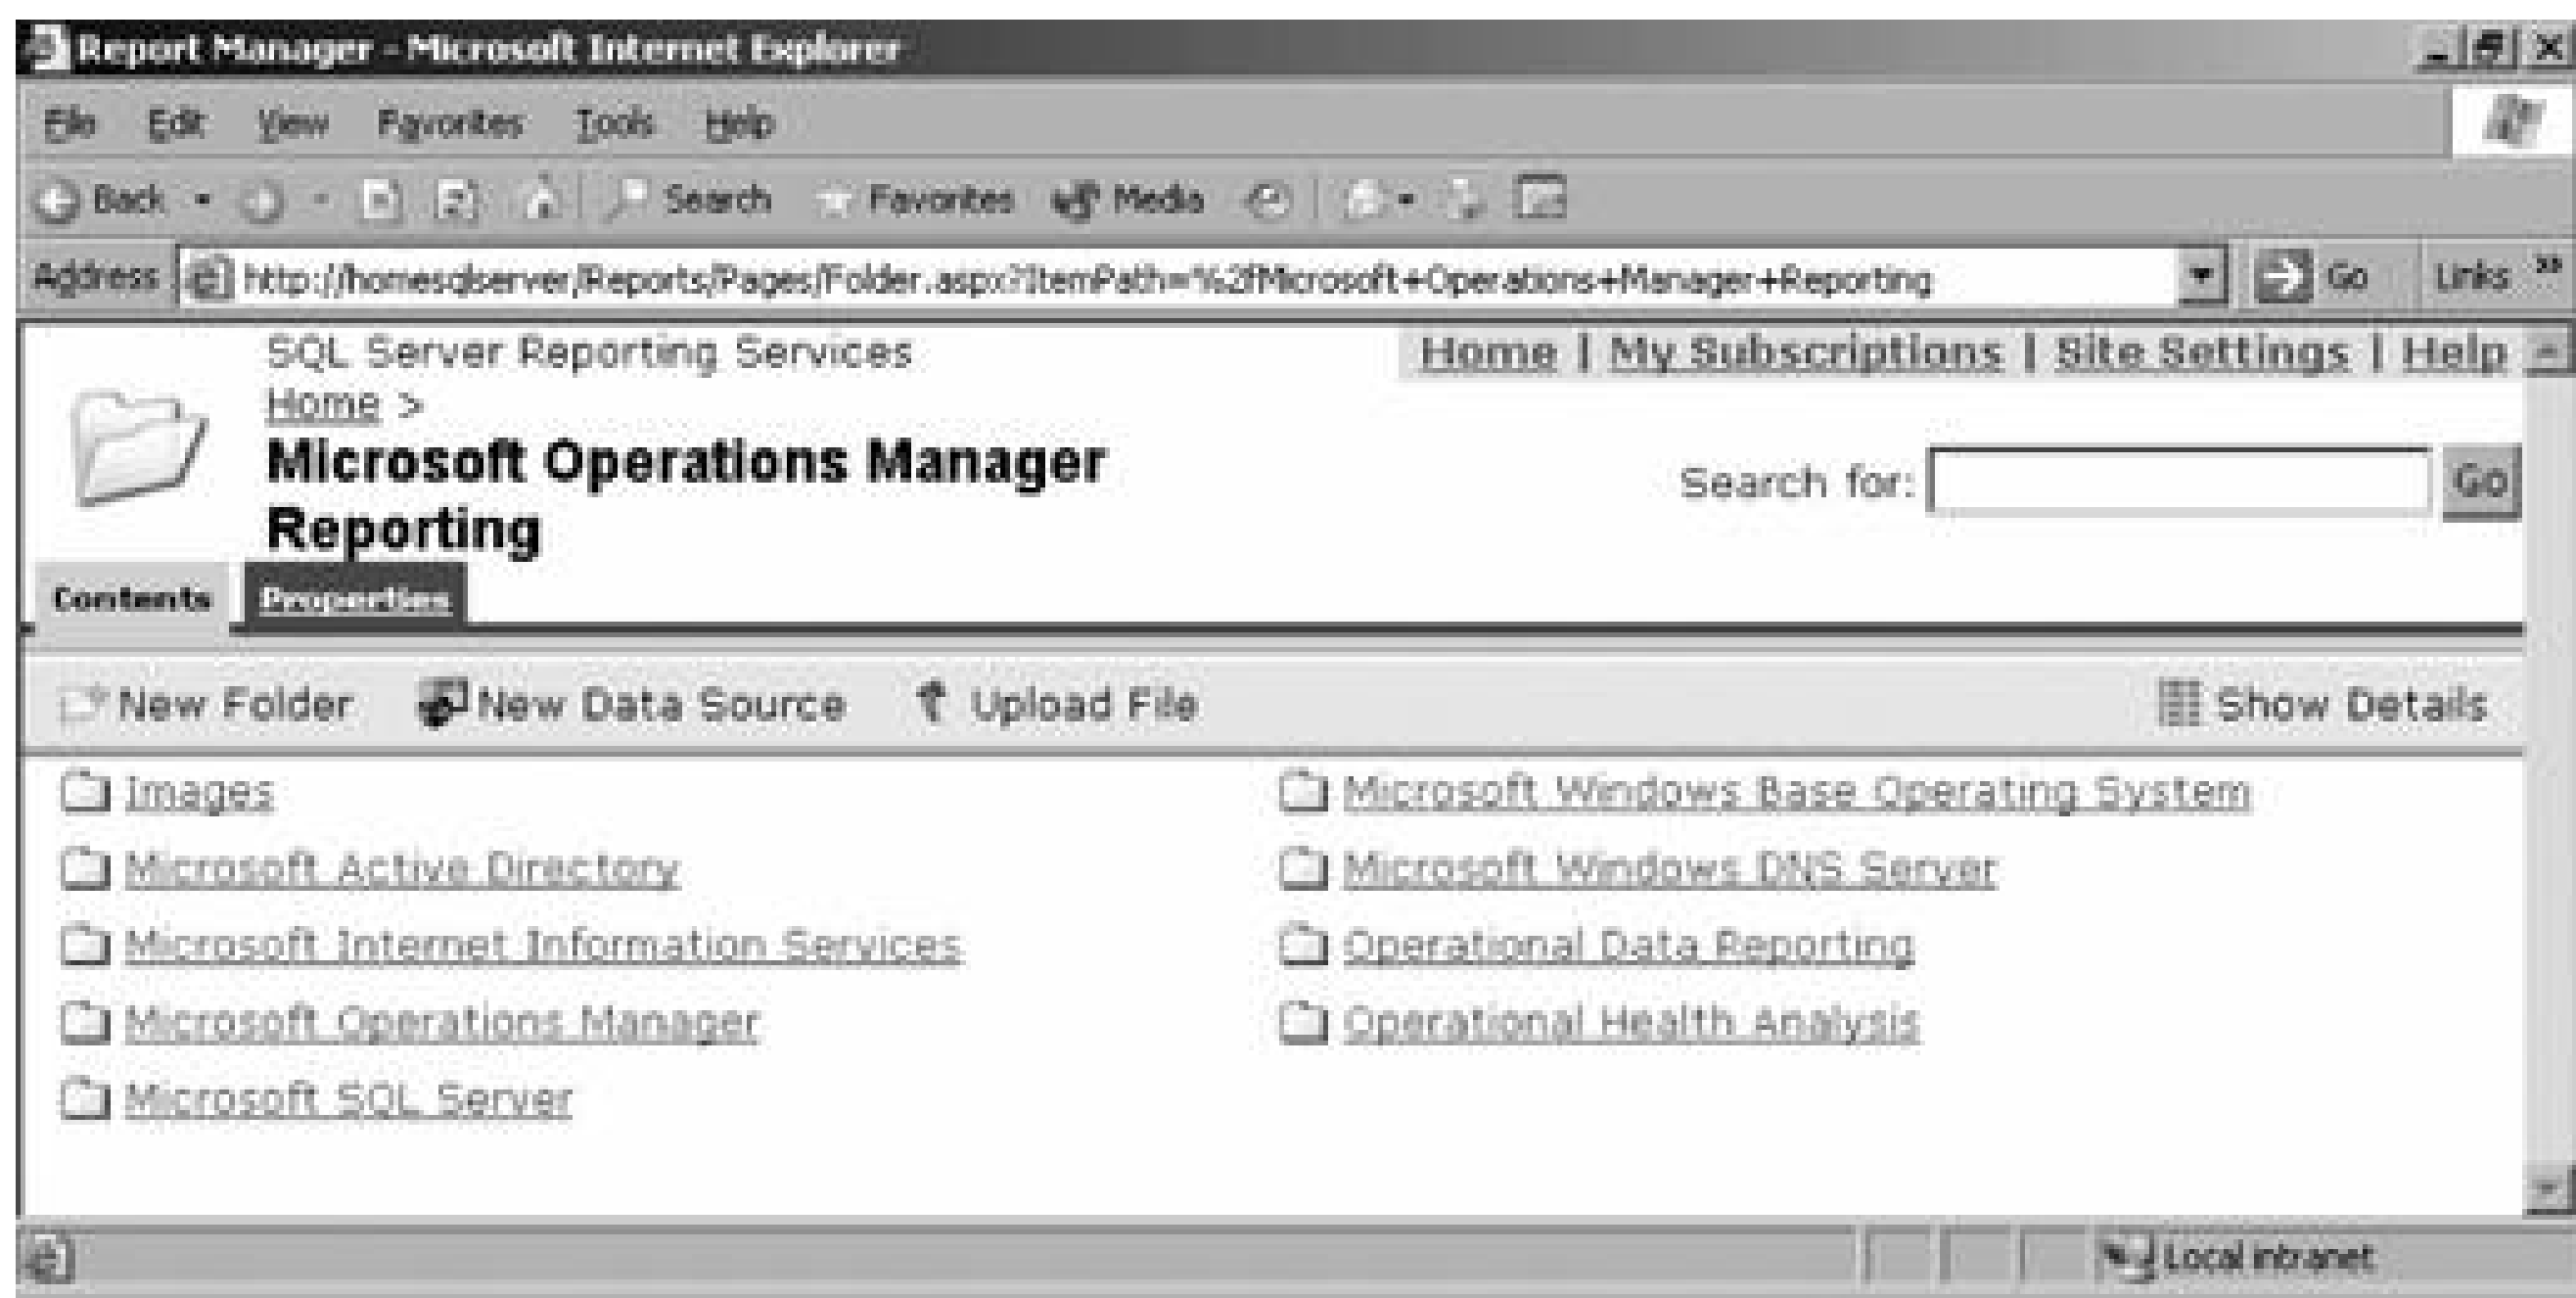
Security can be applied on a per-report basis. By borrowing some functionality from Microsoft Windows Sharepoint Services, reports can be subscribed to.



# 1.5.4. Reporting Console

MOM 2005 takes the SQL Server Reporting Services Report Manager and presents it as the Reporting console (point 4 in [Figure 1-16](#) and shown in [Figure 1-18](#)). All MOM 2005 reports are accessed and managed through the Reporting console.

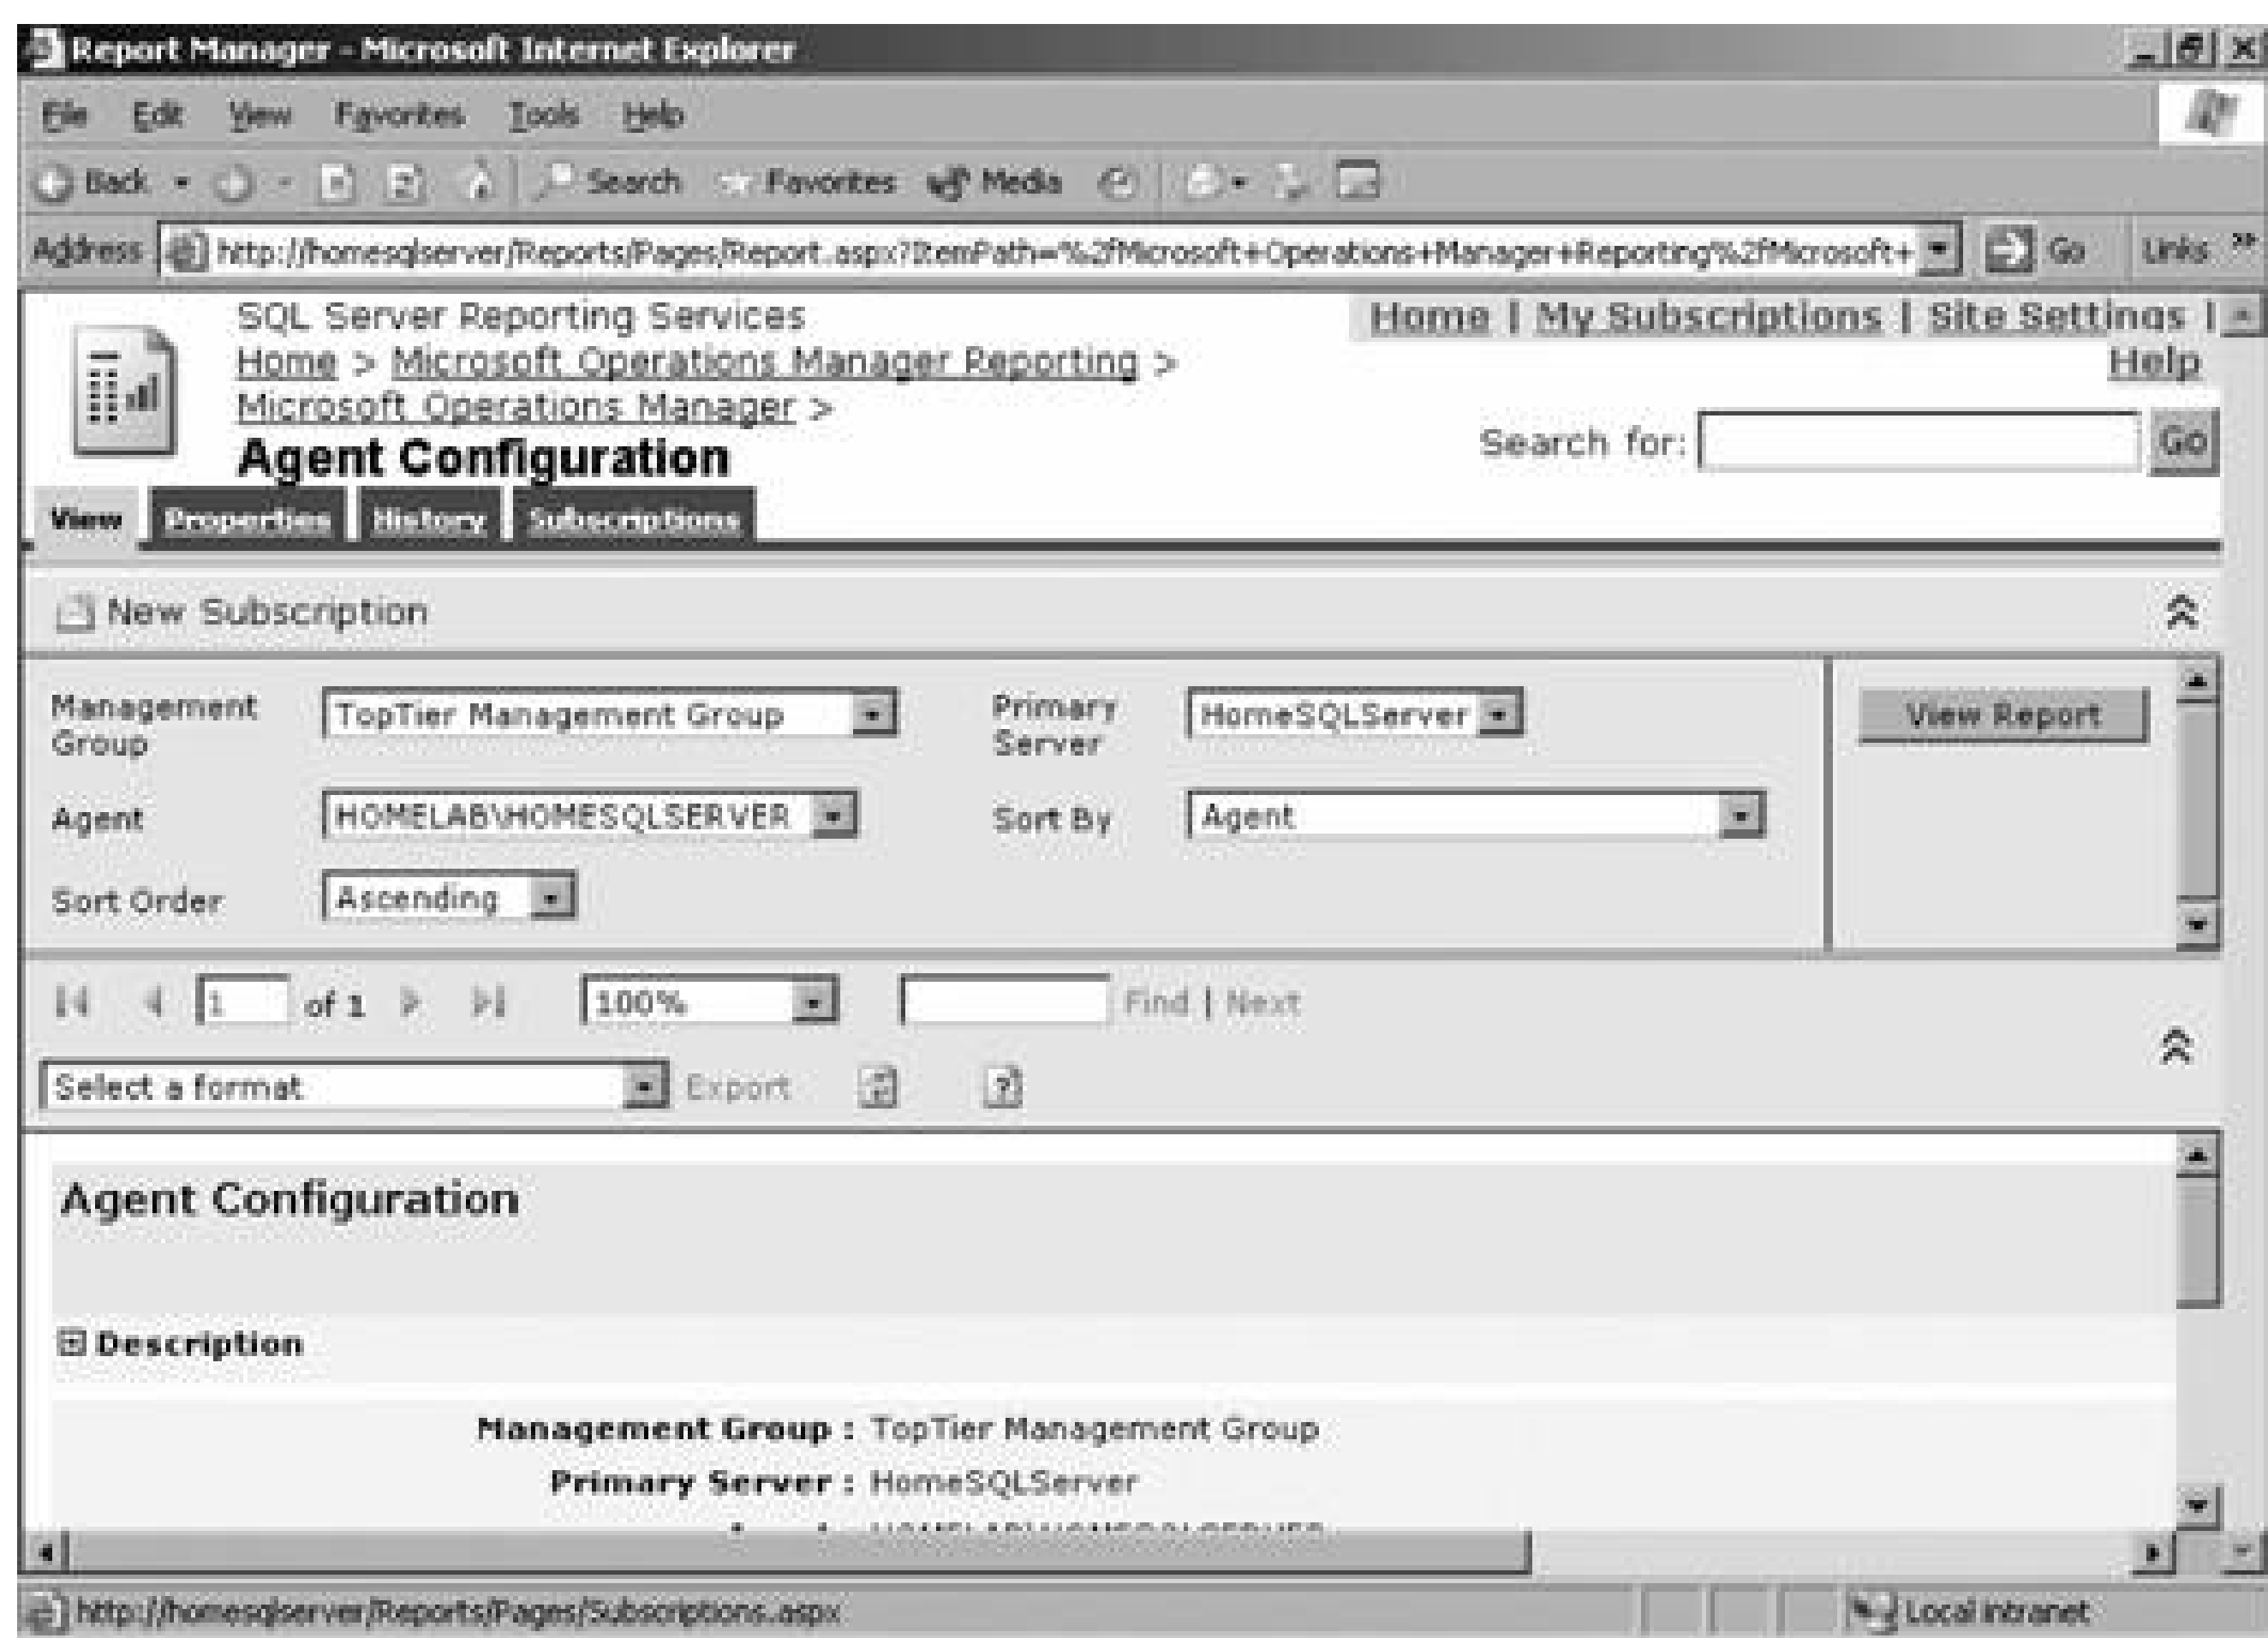
Figure 1-18. MOM 2005 Reporting console



Inside individual reports, users can enter parameters for filtering the information that they want to see. For example, in [Figure 1-19](#), the Agent Configuration report, the user can specify the management group, the agent, the sort order, and the primary management server for this report.

For each report, a user with the appropriate permissions can also access the Properties, History, and Subscriptions tabs that are located across the top of the report. Through these tabs, you can set the default parameters for a report, change the data source, and set credentials for accessing the data source. You can create cached reports that have a Time To Live (TTL) value and create snapshots of reports that capture a report at a fixed point in time.

Figure 1-19. The MOM 2005 Agent Configuration report



The Microsoft-defined reports are included in the application management packs. As of the MOM 200! release, Microsoft is shipping in excess of 200 reports. Custom reports are now designed using the Report Designer in Microsoft Visual Studio .NET 2003 as a Report Project. Through this tool set, you can access the objects defined in the new Report Definition Language (RDL) that was created for SQL Server Reporting Services.

### 1.5.5. System Center Data Warehouse (SCDW)

The data warehouse database, SQL Server Reporting Services, and the Reporting console compose the System Center Data Warehouse. Microsoft came up with this name because it plans on combining the data from future versions of MOM and Systems Management Server (SMS) into a single database. Further down the road, the two products will become more tightly integrated in their user interfaces and eventually merge into one product, called System Center Suite.

## 1.6. Summary

To keep your IT environment running smoothly, the greatest advantage you can give yourself is to know what is going on in your environment right now. When you have this operational awareness, you can resolve major issues faster and prevent larger outages by fixing smaller issues before they snowball. MOM 2005 assists you by actively monitoring as many machines and applications in your environment as you wish and sending you an alert when anything significant happens.

MOM 2005 does this by deploying intelligent agents from a central management server to all managed computers. These agents compare the state of the monitored machine to a vendor-defined health state, generate alerts when health state rule criteria are met, and take certain actions on the managed computer. The agent sends the alert to the central management server, which captures the information in a database and then makes the alert object available to you in the Operator console.

You can use the Operator console to consume all the information that is contained in an alert and to manipulate the state of the alert. To help you troubleshoot issues, the Operator console gives you certain predefined actions that can be executed against the managed computers, along with specific troubleshooting recommendations from the application vendor. After you have completed your troubleshooting steps, you can capture that knowledge immediately into MOM through the Operator console so that it is available for future troubleshooting efforts.

MOM 2005 has a core set of components that are involved in generating and managing alerts. These are agents that read data from providers, the MOM server, the DAS, and the Operator console. Other components include the Administrator, Web and Reporting consoles, the MOM 2005 APIs, and the reporting database.

Now that you have had a brief look at MOM 2005, I'll guide you through the necessary steps to plan for and implement a basic MOM 2005 installation in [Chapter 2](#). [Chapters 3](#) through [8](#), which are the bulk of the book, cover the most common administrative tasks that you will perform after you complete the setup wizard and use MOM to monitor and troubleshoot your environment.

[Chapters 9](#) and [10](#) look at the facilities MOM 2005 has for extending its functionality and connecting to other operations management frameworks.



# Chapter 2. Designing, Planning, and Implementing MOM 2005

MOM 2005 is easy to install, but you need to make certain design decisions before starting the installation process. If you attempt to implement MOM 2005 without answers to these questions, you'll have a lot of work to redo down the road. These design decisions are based on the management group structure that you need, which is based on your requirements.

The design process is also known as *topology planning*. The design of MOM focuses on the management groups and management servers in those groups. It involves the relationships between those groups and the required features and additional servers installed to support those features. The design process includes decisions on redundancy and overall performance.

This chapter goes through the development of a MOM 2005 design and provides sufficient configuration to get you up and running. I'll provide an example of the MOM planning and deployment efforts at a fictitious company called Leaky Faucet. Leaky Faucet is a plumbing supply manufacturer with 40 to 60 Windows servers and centralized administration, but has distributed offices and services.





# 2.1. Requirement Gathering

The design process starts with finding out what your company needs MOM to do and stating those requirements plainly for all interested parties to review. As you develop the list of requirements, talk to people in IT management and technical roles: business process owners, executive management, and information security. It is important to solicit different points of view to ensure a best-fit design for your company. Your questions should be open-ended and broad. Start simply with "What do you want MOM to do?" Answers will range from the very functional, such as, "MOM needs to tell me when something has gone wrong in my environment," to something that is more of a consequence of implementing MOM, such as, "Once implemented, MOM must reduce our downtime by enabling us to be proactive." Hopefully, the answers will be consistent with each other. You should be able to answer, "Yes, MOM will tell you when something has gone wrong in the environment" and "yes, implementing MOM will help us reduce downtime." An inconsistent answer means that someone expects MOM to do something that it doesn't do. You must then educate end users about what MOM can do so expectations are brought in line.

The question to ask next is, "What does our company need MOM 2005 to be optimized for?" A few options include performance, availability, redundancy, cost to implement, integration with existing operations management products, and network impact. These discussions provide two benefits: you will have the input needed to build the most beneficial operations management solution for your organization and expectations can be set with key stakeholders.

The following example is how a Windows administratorMax at Leaky Faucethandled the requirement gathering process. Max must improve system uptime. He is on a small team of 3 people that provide services to 11 sales offices and a central financial and IT office. The team supports about 50 servers. They have been plagued with intermittent loss of communication between the workstations and the file and print servers. The financial and order-taking systems have not been affected. So, although these outages are disruptive, they have not yet hindered the company's ability to service its customers or negatively affected the revenue stream.

Max is bringing MOM 2005 in-house to help monitor these systems and is starting the design process. He has completed interviews with the director of IT (his manager), the director of the business units most affected by the outages, the IT support staff at the remote sites, and the CFO. Here are their responses:

## *Director of IT*

MOM 2005 must provide real-time monitoring and alerting. Since the IT department is not staffed to provide around-the-clock coverage, MOM must have minimal outages and notify the IT staff when they are out of the office. The director of IT views MOM 2005 as an autopilot for his IT operations, but it is no replacement for hands-on human monitoring. For his needs, MOM must provide solid reliability and availability.

## *Business unit director*

The business unit director is the IT department's primary customer and is more than a little annoyed when her staff's daily work is interrupted. She expects IT to provide a higher quality of service, especially since she receives a monthly charge-back for the disk space her staff uses and other IT overhead charges. She wants IT to give her reports on service uptime and resource consumption. Her business unit is also purchasing a new document management application and IT is to report on that as well. She has no need to interact directly with MOM and, in fact, doesn't really want to know that it exists; she just wants the IT services to work.

### *Remote site support staff*

The remote site support staff have job responsibilities in addition to local desktop, printer, and backup/recovery support. They want insight into the computers that are their responsibility and are not interested in anything else. The remote site support staff must be able to get the information that MOM produces, but are not interested in MOM configurations or rules groups. If MOM is not giving them the information needed, then the IT staff at headquarters will be called for support. The remote site support staff have limited administrative abilities in Active Directory over their organizational unit (OU).

### *CFO*

The CFO is responsible for ensuring that governmental audit requirements on the financial systems are met. Although the Microsoft space is not significantly impacted by these requirements yet, she wants similar controls here. Primarily, she wants accountability for changes introduced into the IT systems and to know who accesses certain restricted data. The auditing reports from MOM must be available on a weekly basis and, for reasons of confidentiality, only she and the rest of the executive staff can have access to them. They also need to get the reports without relying on IT. Because IT operations are under her responsibility, she is the executive sponsor for all IT efforts and represents IT to the rest of the executive staff and the board of directors. She expects that MOM will be implemented in the most cost-effective way possible.

The Leaky Faucet physical environment consists of:

- 128 Kbps and 256 Kbps wide area network (WAN) connections that have fairly heavy use, with a firewall between the remote sites and the central site and an DMZ configuration between the central site and the Internet. All Internet connectivity goes through the central site connection.
- They have upgraded to Windows Server 2003 for all domain controllers and member servers in Leaky Faucet's single production AD domain. In this environment, the company runs Active Directory 2003, Exchange 2003, IIS 6.0, SQL 2000 ISA 2004, SharePoint Portal Server 2003 and Windows SharePoint Services, and BizTalk.





# 2.2. Design Decisions

The business requirements are as follows:

- MOM 2005 must monitor all Windows servers (about 50) and the Microsoft applications that they run. The volume of data that MOM will have to process can be estimated from this requirement and goes directly into the capacity planning calculations (covered later in the chapter). MOM must also support the monitoring of third-party applications in a reliable fashion
- MOM must be available with minimal outages. These are the availability and redundancy requirements .
- Because MOM 2005 must provide reports that are used in auditing and charge-backs, the data must be secured while going from the agent to the database and then to the presentation layer. This will direct the security and report planning.
- The remote site staff only wants to see their information and will not tweak MOM.
- Some agents will be located across firewalls and slow WAN links.
- Access to MOM information is not required from the Internet.

## 2.2.1. Versions

Based on this information, the first design decision to be made is which version of MOM 2005 to implement. There are two versions : the MOM 2005 edition, which includes all features and scales from the smallest installations to the largest, and the MOM 2005 Workgroup edition.

The Workgroup edition has the same core infrastructure as the full edition, but it only supports up to 10 managed computers (agents). It does not include reporting or connectivity features. MOM 2005 Workgroup edition requires that all components, except for the consoles, are installed on a single piece of server hardware. Other features and restrictions of the MOM 2005 Workgroup edition include:

- It can only be installed on Windows Server 2003.
- It does not support agent failover.
- It uses the same management packs as MOM 2005, so it provides the same event and performance monitoring and product knowledge as the full edition.
- It pre-imports the most commonly used management packs into Workgroup: Baseline Security Analyzer, Exchange 2000 and 2003, SMS 2003, AD, Base OS, DNS, IIS, MOM, and Clustering.
- It has the same Operator, Administrator, and Web consoles as the full edition.



The MOM 2005 Workgroup Edition is intended, and priced, for very small environments that require limited functionality. Because the Leaky Faucet environment has more than 10 servers to monitor, and reporting and redundancy are required, the full MOM 2005 version must be used.

### 2.2.1.1. Management groups

The next step is to decide how many management groups are required and determine the number of management servers in each management group. As mentioned in [Chapter 1](#), all management groups have 1 to 10 management servers (maximum), an operations database, managed computers and Administrator and Operator consoles. All MOM infrastructures start with a single management group, which can be added to if required. This is similar to AD planning, which starts with a single domain and grows from there.

A single management server can manage up to 2,000 agents, and a management group can administer up to 4,000 agents. If you think the math doesn't add up, you're right. The operations database, of which there can only be one per management group, is the limiting factor of the management group, not the management servers. There is more to capacity planning than this (see the "[Pre-Installation Configuration Decisions](#)" section later in this chapter). For each category, the maximum amount is as follows:

- Management servers in a management group: 10
- Agents per management server: 2,000
- Agents per management group: 4,000
- Agentless-managed computers per management group: 60
- Management servers an agent can report to: 4

Since Leaky Faucet only has 50 servers to manage, a single management group will suffice for the production environment. But what about testing? The Leaky Faucet Windows team knows new applications and servers will be introduced into their environment. So, they have to provide a stable production environment, but testing new applications and servers in production is unacceptable. To solve this issue, Leaky Faucet creates a preproduction management group.

A preproduction management group is used for two tasks:

1. To have a place where a newly built server can be monitored for the amount and types of alerts it will generate without having any impact on the production MOM environment and network. An agent from the preproduction management group can be placed on the new server and data can be collected. If something on the server is misconfigured or it misbehaves in some other way, then these issues are discovered and dealt with before the server goes into production.
2. The preproduction management group can create new management packs and rules and modify existing ones. Management packs, which contain rules, tell MOM how to identify an application and what to monitor. They define what information is collected and determine alert generation. Without an application-specific management pack, MOM won't recognize the application and,

therefore, cannot monitor it.

Tweaking MOM 2005 (see the "[Using MOM 2005](#)" section in [Chapter 1](#)) is actually tweaking the management packs. During the tweaking stage, you can enable or disable processing rules and adjust the alert thresholds to get usable information. Managed code (.NET programs) can be written to govern how an agent will respond to an event on a managed computer. You may also need to develop management packs, with all of their components, from scratch.

A misconfigured management pack will produce unpredictable and potentially disruptive results. For successful MOM 2005 operations, it is essential to test management pack changes before you introduce them into your production environment. This is explained in [Chapter 4](#).

Management packs are very complextweaking, testing, and developing them in production is a bad idea.

There are other reasons to have multiple management groups, such as for security boundaries between groups. For example, if a business has multiple independent units, each with its own IT infrastructure and administration, then it would be appropriate to separate management groups along the administrative boundaries. Another reason to have multiple management groups is to separate certain types of data, such as the security event logs, from the rest of the alert, performance, and event data. You may want to do this for capacity or administrative reasons. Scalability is another reason for multiple management groups. Say a company has more than a few hundred managed computers at multiple remote sites, but wants to manage alerts at a central datacenter location. That company can place management groups at each remote site and configure a MOM-to-MOM Product Connector (MMPC) to forward alerting and other information to the central datacenter site for administration. This creates a tiered MOM infrastructurethe lower tier is the source management group and the top tier is the destination management group. For more information on integrating multiple management groups into a single reporting structure and integrating MOM to other operations management products, see [Chapter 9](#).

Leaky Faucet has only one Windows Active Directory domain in the production environment but will perform all management at the central headquarters location. Leaky Faucet creates two separate management groups, one for production and one for preproduction, with no connectivity between them.

### **2.2.1.2. Management servers and the operations database**

Since Leaky Faucet only has 50 servers to manage, even if you add the servers that will host MOM 2005 and allow for 100 percent growth, the number of servers is comfortably within the capacity limits of a single management server (2,000 agents). But since the director of IT wants this solution to be highly available, every effort must be made to eliminate single points of failure. Therefore, a second management server will be added to the management group. All agents that are owned by that management group are automatically made aware of the additional management servers. The agents can dynamically failover to a secondary management server if there is a loss of communications with the primary management server.

When planning for failover with a large number of agents, the sum of the number of failed-over agents and the number of agents already managed by a single management server cannot exceed 2,000. For example, if you have a management group that manages 3,000 agents and you split the agents evenly between two management servers, there is no capacity left for failover. To allow for



continued operations in the event of a failure, you would need to have at least 3 management servers with 1,000 agents each.

Leaky Faucet implemented two management servers in the production management group and one in the preproduction management group since high availability is not a requirement. Without the redundancy requirements for preproduction, the cost of an extra server with software licensing cannot be justified.

So then, how will Leaky Faucet respond to a single point of failure in the single operations database? High availability for this component can be addressed in three ways.

1. Install a Windows Server 2003 Enterprise Edition and create a Microsoft Cluster Service cluster for the SQL 2000 Server instance. This mitigates, as best as possible, hardware failure and allows for rolling upgrades to SQL and the cluster servers. For more information on Windows Server 2003 clusters, see *Learning Windows Server 2003* (O'Reilly).
2. Implement a solid backup and restore plan, as discussed in [Chapter 7](#).
3. Configure the drive array configuration for at least RAID 5 (striping with parity) but preferably RAID 10 (striping across multiple mirrored sets) and separate the database transaction logs and database files onto different spindles. If you have the luxury of configuring your cluster-shared drives on a storage area network (SAN), then the RAID choice is already made.

However, a clustered solution requires that one cluster node is idle while the other node hosts the clustered application. If the first node fails, the second node takes ownership of the application and continues to run it. Leaky Faucet cannot justify the expense of a computer sitting idle in addition to the cost of the external drive enclosure that the two cluster nodes would share. Max, with the support of the director of IT, chooses not to cluster the management group database server, but the database will be on a RAID 10 disk configuration.

### 2.2.1.3. Additional feature and services

Leaky Faucet needs to provide services that are not included in the basic MOM 2005 feature set. Both the business unit manager and the CFO want usage and tracking reports for auditing purposes and to have a record of IT service. The CFO wants secure reports and access to them without going through the IT staff.

The remote site support staff will access MOM 2005 data across a slow WAN link and through a firewall. They will not perform configuration duties.

Max recognizes that a MOM 2005 Reporting Server component must be included in Leaky Faucet's deployment of MOM 2005.

Where the operations database supports the monitoring, alerting, and configuration functions of MOM 2005, the Reporting Server database supports the reporting function. MOM 2005 reports are based on historical data and provide a longitudinal view of monitored application performance and the servers the applications run on. A scheduled task runs once a day by default, which transfers the live data from the operations database to the Reporting Server database. The dataflow in the transfer is unidirectional, going from the operations database to the Reporting Server database only. The



transfer is accomplished by a SQL Server DTS package. This scheduled task is coordinated with the grooming jobs that delete old information from the operations database. This ensures that no data is removed before it has been transferred over to the Reporting Server database.

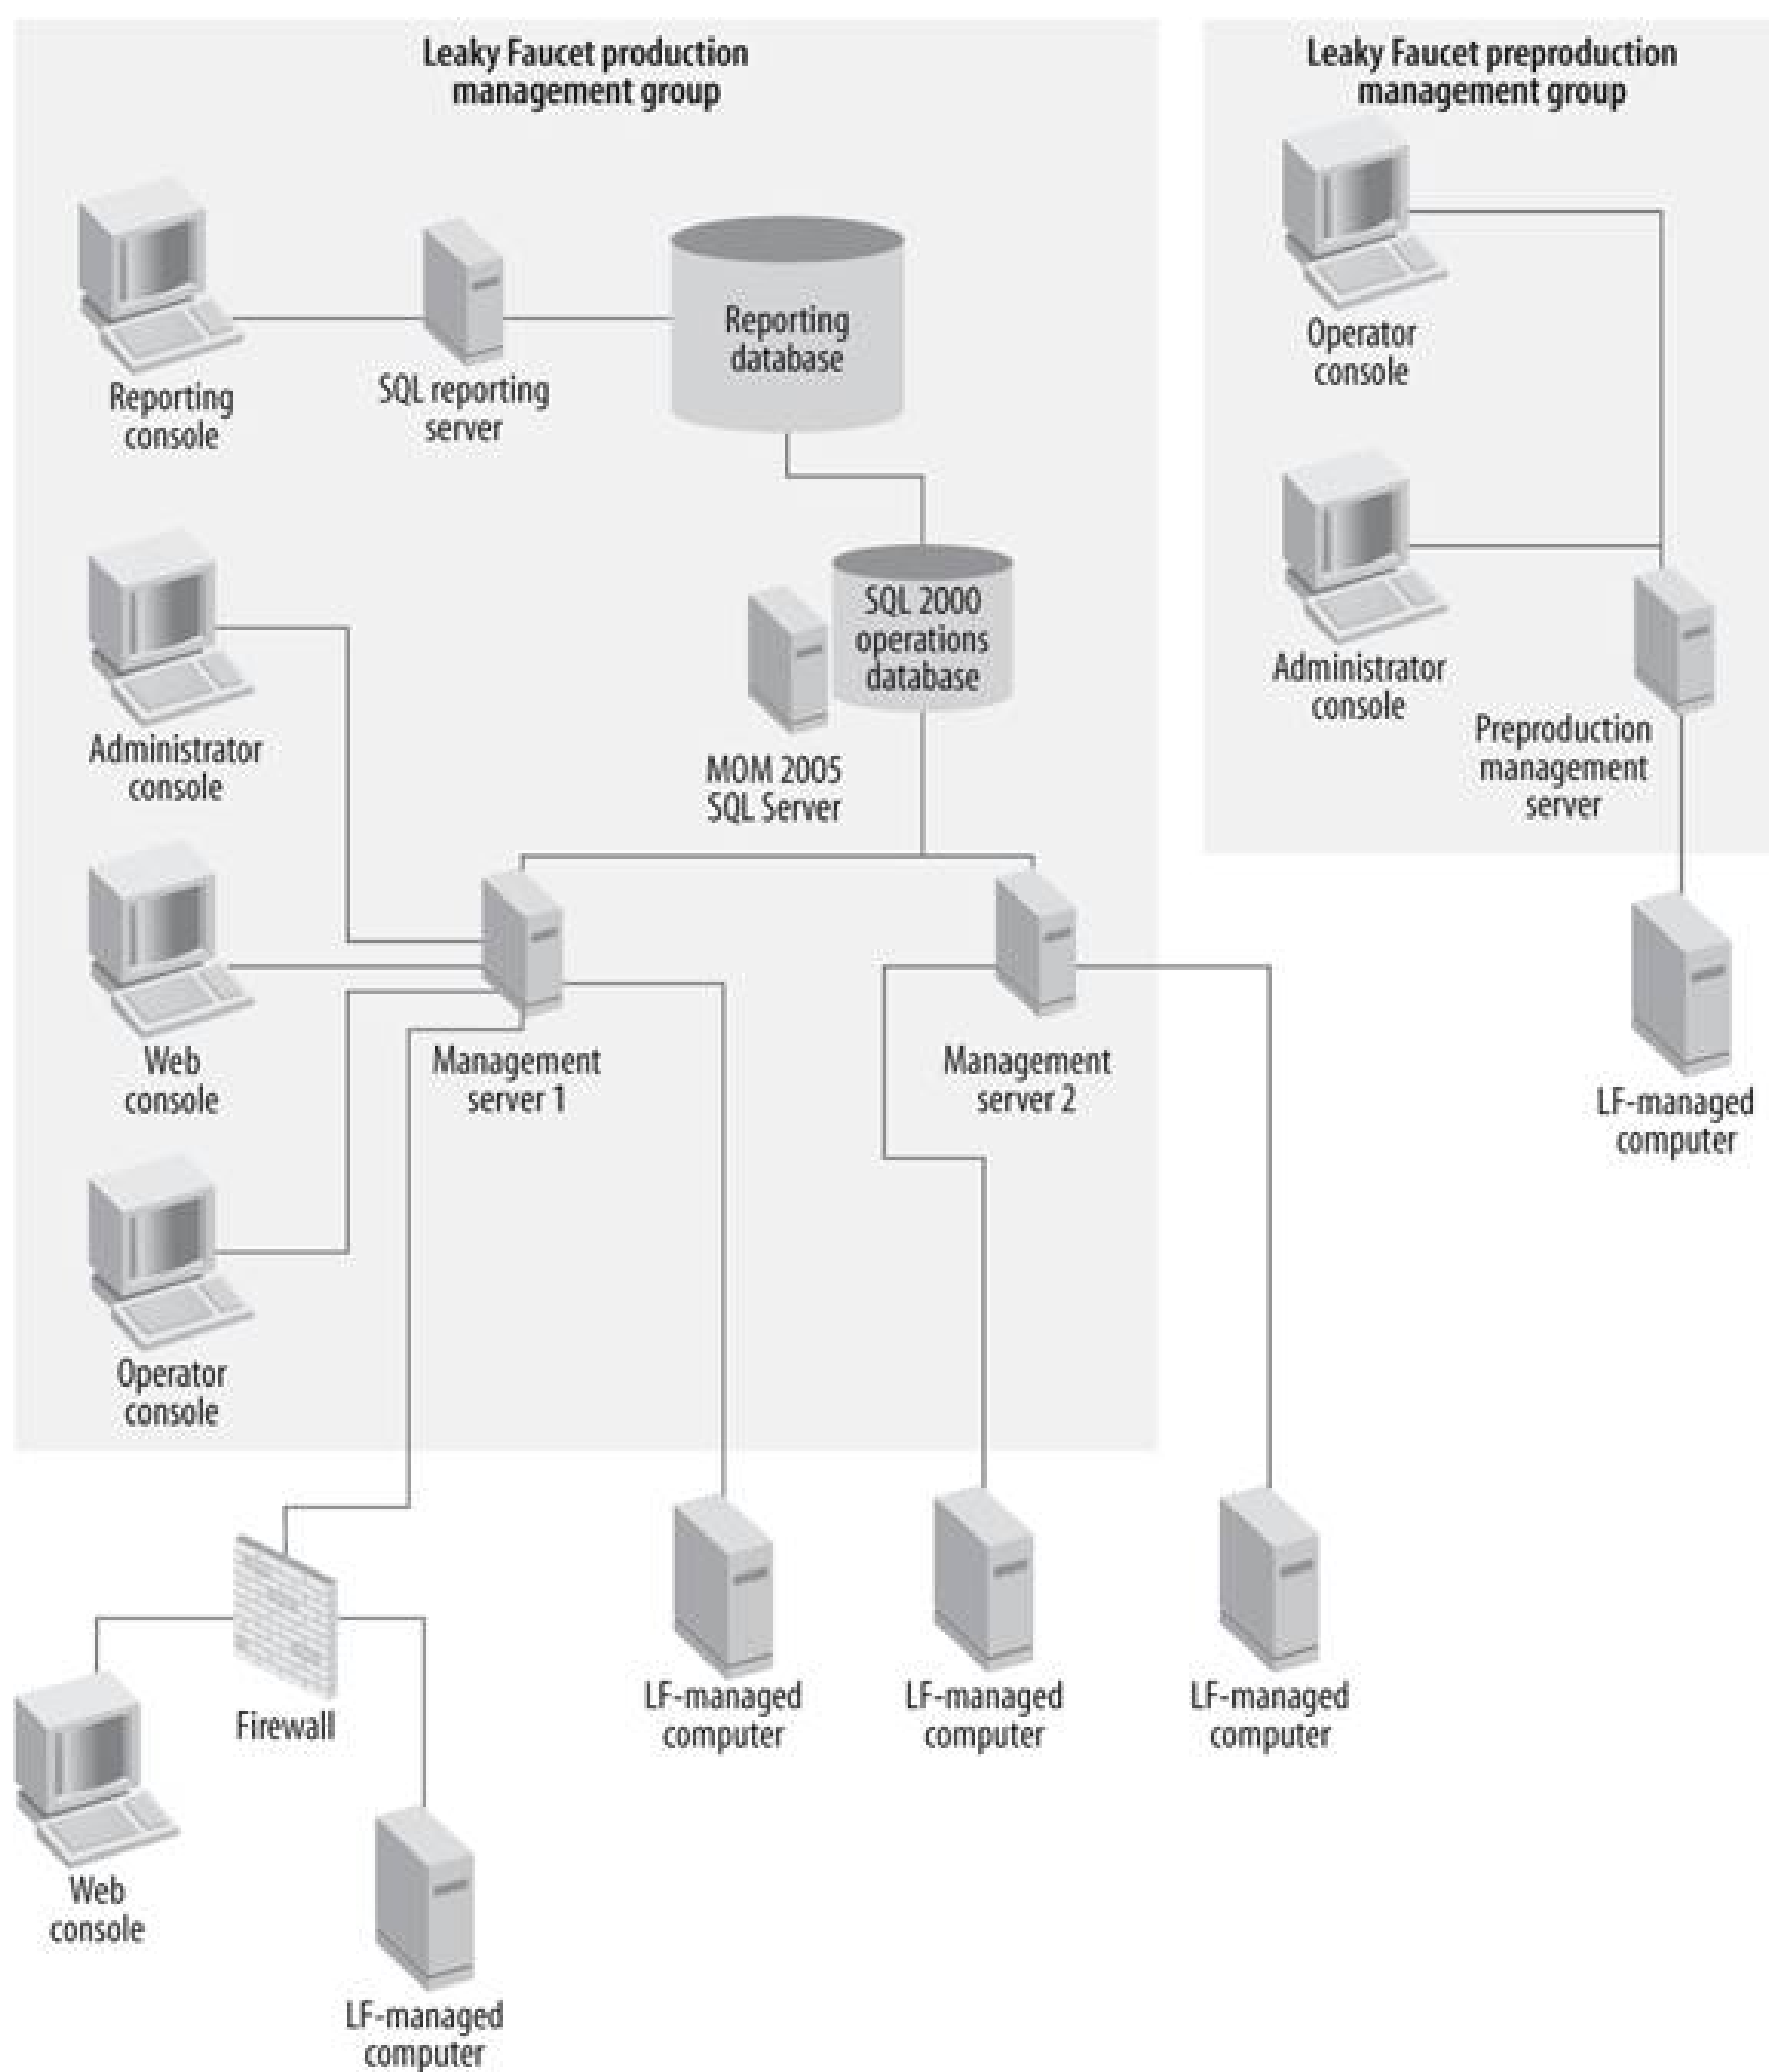
MOM 2005 reporting uses SQL Server Reporting Services. SQL Server Reporting Services is a separate product from MOM 2005 and requires separate installation. Security can be applied to individual reports or their folders, much like NTFS permissions applied in a file structure hierarchy. This will make the CFO happy and the reporting web site can be SSL secured. For the business unit manager, MOM 2005 reports include service uptime and outage information. These can be mailed directly to her Exchange mailbox.

Unlike the full-featured Operator console, the MOM 2005 Web console displays only the information essential to managing your environment. It presents the operational data by alert view, computer view, and event view. The Web console can be used by the remote site support staff or they can use the Operator console. What the remote site support staff sees can be controlled via *console scopes*. A console scope is an association created by a MOM administrator between a user account and one or more computer groups. It allows users to only see data from computer groups that are associated with their user accounts. The information they have access to will be filtered by the scope assigned to the logged-on user. The Web console is best when there are firewall issues or when working remotely. It is lightweight and all communication occurs over port 1272, so arranging an open port in the firewall should be easy.

## 2.2.2. Initial MOM 2005 Design

By this time, Max has made enough design decisions to develop a design diagram, as shown in [Figure 2-1](#). All management servers will be centrally located at headquarters. The Web and Operator consoles will be available for remote access, and the Operator and Administrator consoles will be available locally. Note that when you run MOM 2005 setup for the MOM user interfaces (UIs), both the Operator and Administrator consoles are loaded, so they are always on the same machine.

Figure 2-1. High-level diagram of the MOM 2005 infrastructure at Leaky Faucet



Based on this design, Leaky Faucet will need to purchase the following:

- Three servers for the management servers (two in production, one for the all-in-one preproduction)
- One server for SQL Server 2000 (production). The preproduction management group can use the Microsoft SQL Desktop Edition (MSDE) for its database or a full version of SQL Server. To

keep the two environments as similar as possible, Leaky Faucet uses SQL Server 2000 in both environments.

- One server for the Reporting Server. This server requires SQL 2000, IIS, and SQL 2000 Reporting Services to support the report manager interface and the Reporting Server database.
- Five Windows Server 2003 Standard Licenses.
- One MOM 2005 Operations Management License (OML) for each management server and database server, as well as agent licenses for each computer to be monitored.

This basic infrastructure should meet Leaky Faucet's needs for some time and can be easily extended. If a third management group is needed, then a similar planning exercise should be undertaken to develop redundancy and capacity requirements. Connectivity requirements will also have to be developed, as discussed in [Chapter 9](#).

The choice of hardware for a MOM 2005 setup should be influenced by three factors: desired performance, desired redundancy, and available funds. Your hardware solution will be unique to your situation but here are some useful guidelines:

- As an application, MOM 2005 is most influenced by the performance of the operations database server. MOM imposes a 30 GB size limit on the operations database to keep it relatively small and performant. Performance of the operations database server is most influenced by its disk performance, so don't skimp on the hardware. Don't optimize the operations database for space (no more than 18 GB out of a 30 GB database should be used), optimize it for speed.
- Take what you determine as the hardware need for a management server and double it for the operations database server. For example, if you need a single 3-GHz processor and 1 GB RAM in a management server, the operations database server should be a dual 3-GHz with at least 2 GB RAM. If you need a dual processor in the management server, make the database a quad.
- Size your hardware to accommodate reasonable load growth for three years. The last thing you want is a hardware upgrade a year after MOM 2005 is installed.

Microsoft has well-published minimum hardware standards for each possible type of server in a management group. [Figure 2-2](#) is a screenshot of the MOM 2005 Sizer.xls tool that gives you an idea of the load you can expect and the minimum-size hardware needed. This tool is included in the MOM 2005 Resource Kit and available at <http://www.microsoft.com/downloads/details.aspx?FamilyId=93930640-FA0F-48B3-8EB0-86836A1808DF&displaylang=en>. However, it is no substitute for testing in your environment.

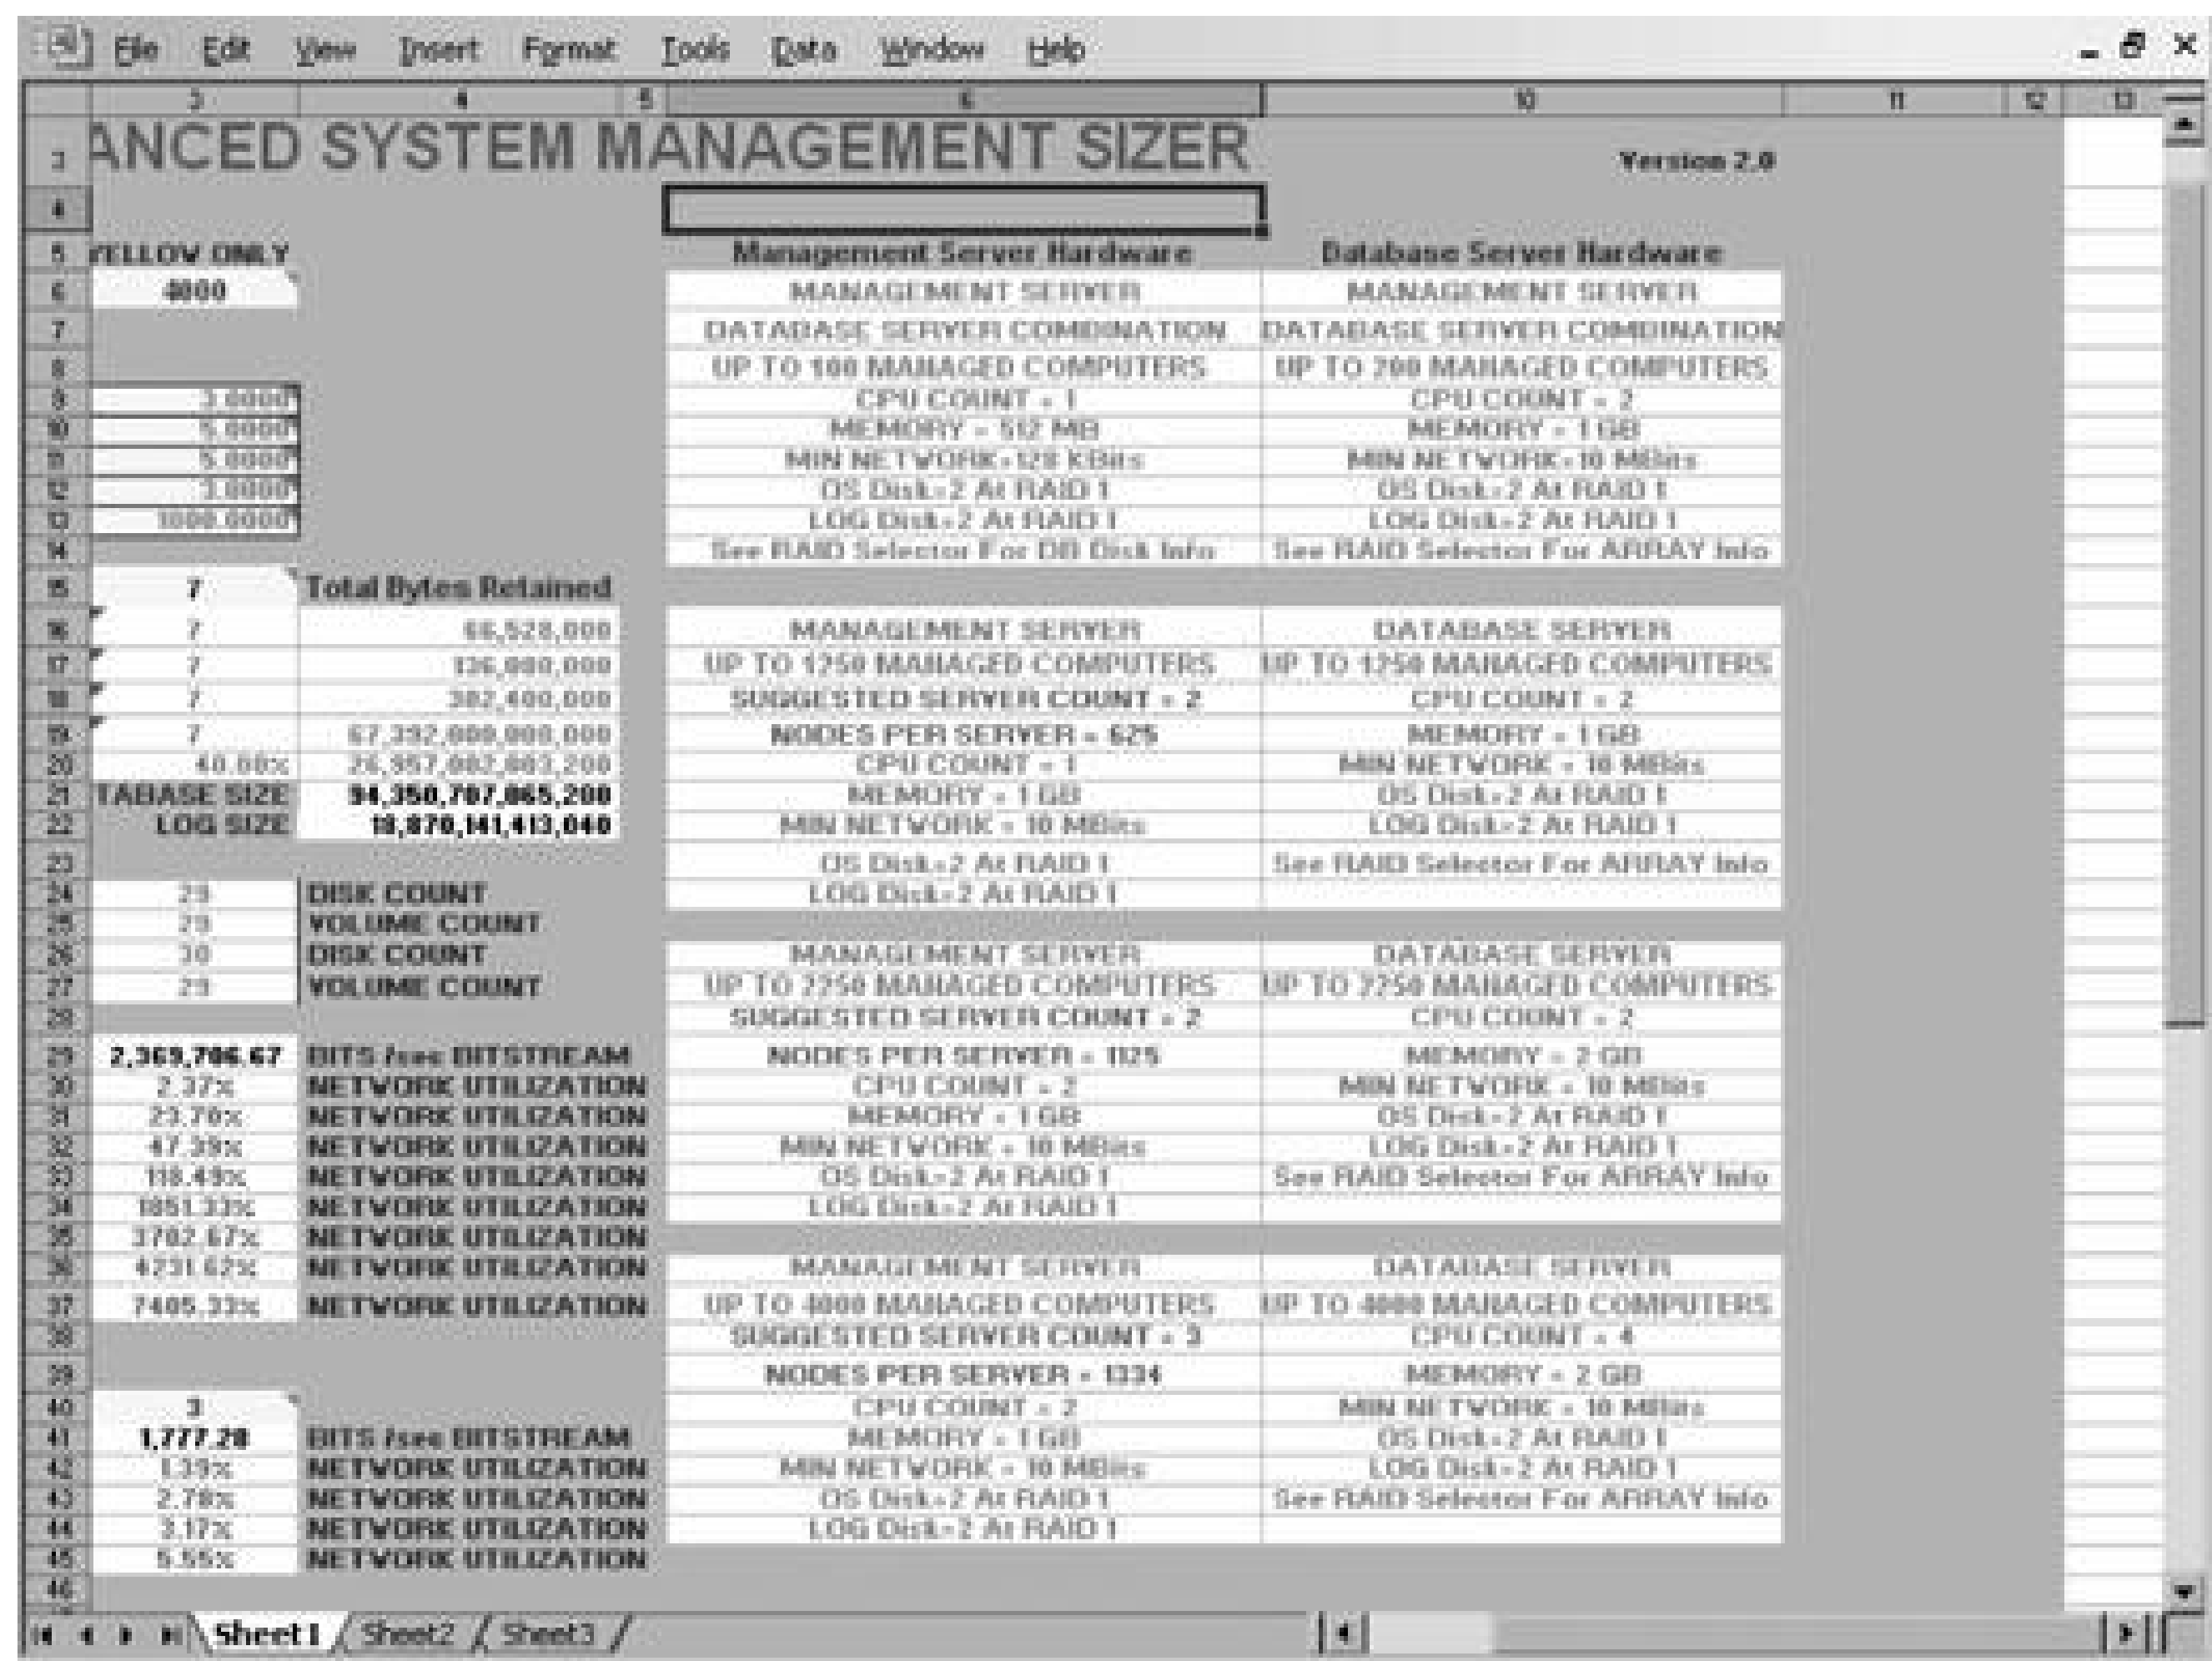


## 2.3. Pre-Installation Configuration Decisions

Setting the design and getting the hardware up and configured are just the first steps. More preparation must be done before installing MOM 2005:

- Planning out the user accounts, groups, and roles that will be used

Figure 2-2. The MOM 2005 Sizer tool gives you estimates for database size, network load, and minimum hardware suggestions



- Calculating the starting sizes for the operations and reporting databases
- Preparing the environment and security

During the installation, you will be prompted for specific accounts and starting sizes for the databases. You must have this information before you start. You will not be prompted for other special security requirements, but it will be to your advantage to configure certain security options before installing some components, such as the Web and Reporting consoles. Otherwise, you'll have

to redo your work.

## 2.3.1. User Accounts, Groups, and Roles

The MOM 2005 installation prompts you for two accounts, the management server action account and the database access service (DAS) account. These accounts are not the typical service accounts. You have no choice regarding what user account the MOMService (named *momservice.exe* in the computer management (MMC) services node) logs on as. For Windows Server 2003, the MOMService runs as the network service account. If you install MOM 2005 on a Windows 2000 server, it will run as the local system account. Since Leaky Faucet has upgraded all of its servers to Windows Server 2003 Standard Edition, only the network service account is covered here. The network service account is a new account in Windows Server 2003. It is a built-in account with limited rights on the local machine and can interact with other systems across the network using the local computer account credentials. For more information on the network service account, see *Learning Windows Server 2003* (O'Reilly).

### 2.3.1.1. Management server action account

In MOM 2005, an agent interacts with the computer that it is on, collecting information from event logs, the registry, and Windows Management Interface (WMI) objects. An agent also runs responses and takes actions on the managed computer based on instructions given to it by management packs. These actions can include running scripts or managed code (e.g., C# or VB.NET).

Any agent may be running more than one action at any given time. To prevent a misbehaving script from impacting the execution of other scripts or the MOMService (the same name is used for this on both management servers and managed computers), the MOMService spawns a new process called *MOMHost.exe*. This process is responsible for executing performance counter collection, script responses, managed code responses, and batch file responses. If you were to watch the executing processes on an agent-managed computer, you would see multiple instances of *MOMHost.exe* running from time to time. *MOMHost.exe* runs in the security context of the computer's configured action account. In the case of a management server, this is called the *management server action account* and for all other managed computers, it is called the *agent action account*.

Because agents perform actions on managed computers in the context of the configured action account, this account must have sufficient rights on the local computer to carry out whatever actions are defined in the processing rules that apply to that computer. For example, if the server is a SQL 2000 Server, the agent action account would need to have corresponding rights to interact with SQL objects.

The management server action account has some additional duties. It is used to manage agents on managed computers and can also be used to install and uninstall those agents automatically. In the case of agentless-managed computers, all monitoring is performed by MOMService using the management server action account.

Because the management server action account must interact with machines across the network, it is easiest to use a *domain account*. Microsoft strongly recommends not using an account with domain administrator rights for the management server action account. A standard domain user account with



local administrative rights on the management servers will suffice. If you are going to use fully automated agent deployment, then the management server action account must have full administrative rights on all of the computers that are to be managed. Fortunately, this is not the only way to get an agent installed; you can also provide credentials with administrative rights on the target machines at the time of agent installation. These credentials are disposed of after the agent has been successfully installed. See [Chapter 3](#) for more details.

### 2.3.1.2. DAS account

The DAS account proxies all communication with the operations database for MOM. The *MOMService.exe* and the Operator, Administrator, and Web consoles interact with it to query, insert, modify, and delete data in the database. The DAS is a COM+ application and, as such, requires an identity to provide a security context when communicating with the database. This is the role of the DAS account. It is assigned the db\_owner role on the operations database and is a SQL Server login with Permit server access. The Reporting Server can use the DAS account for accessing the operations database as well. SQL-stored procedures are also executed using the DAS account.

The DAS account has no duties outside of the management group, so it does not need any elevated rights in the domain. You can either have a low privilege domain account or the network service account with a DAS account. The low privilege domain account is preferable when the management group contains multiple management servers. This gives you one account to manage on the database server and in SQL, rather than the multiple computer accounts that are required with the network service account.

### 2.3.1.3. Groups and roles

Access to various levels of functionality in MOM 2005 is controlled by membership in local security groups on the management server and on the operations database server. These local groups are created automatically by the setup process, so you won't be prompted for any information to create them. To use MOM 2005 as quickly as possible, prepare your domain-level security groups so they can be added to the local groups as soon as the setup process is complete. This is in line with the best practice of putting users into global groups, putting global groups into local groups, and assigning local groups to permissions on servers. These local groups are:

#### *MOMService group*

This group is for accounts that perform internal MOM functions. This group is empty by default on new installations. Remember that the *MOMService.exe* process runs in the security context of the local system account for Windows 2000 and the network service account for Windows 2003 servers.

#### *MOM administrators group*

Members of this group have full access to all functionality in the Administrator, Operator, and Web consoles.



### *MOM authors group*

This group is for individuals who will be working with management packs in the Administrator console. Users in the MOM authors group can also use all the features and functions in the Operator console.

### *MOM users group*

The MOM users group is for first responders to alerts. These people are responsible for acting on the alerts and other information in the Operator console. They view the Operator console through scopes that have been defined by a MOM administrator.

### *Systems Center Data Warehouse (SCDW) readers group*

This group is created on the server where MOM Reporting Services are installed and will be prepopulated with the account designated for use in communication with the operations database . Members of the SCDW readers group view and run reports in the Reporting console.

Create domain-level groups for the MOM administrators, authors, users, and SCDW readers groups and add your users to them. No accounts or groups should be added to the local MOMService group.

Max created a domain user account for both the management server action account and the DAS account. Because the Leaky Faucet AD domain has a domain-wide password expiration policy set, an additional preparatory step must be taken. To keep the passwords on these accounts from expiring, they should be placed in an organizational unit that is configured not to inherit group policy from the domain. Max works with his team to plan domain group membership for the MOM 2005 groups. His team will be placed in the MOM 2005 administrators domain group and the SCDW readers group. The remote site support staff will be in the MOM 2005 Users domain group and the SCDW readers group. The CFO, business unit director, and the rest of the executive staff will be assigned to the SCDW readers group. The MOM 2005 authors group remains empty. When the business unit director launches a new document management application, her staff could have input into the usage and uptime processing rules, so a domain group is created and left empty in anticipation of this need.

## **2.3.2. MOM 2005 Operations and Reporting Database Planning**

Planning the configuration of the operations and reporting databases completely depends on the amount of data that will be moved. A certain predictable volume will flow into the operations database (in SQL the operations database is called OnePoint; it is covered in [Chapter 7](#)) from the agents. After a configurable period of time, the data is copied to the reporting database. After another configurable interval, the data is groomed out of the operations database. The copying and grooming schedules must be coordinated so that multiple copy cycles occur before a grooming cycle deletes the data.

So, the two databases handle the exact same data, just at different times and for different purposes. The focus of the operations database is current information, the "What is going on in my environment right now" information. The reporting database takes the same information and keeps it because it is

no longer relevant to what is going on right now. It provides the answers to the "What happened when..." type of questions.

The operations database needs to be relatively small to quickly move current information. The database can be between 300 MB and 30 GB, but you cannot use all of that 40 percent must be kept free. The free space is used by the SQL stored procedures that ship with MOM 2005 (grooming and reindexing) as working space. This leaves you with actual usable space of 180 MB to 18 GB. With this database, you are prompted for a maximum size during setup and the database is configured not to allow growth after that. You can change the setting to allow growth, but it can very quickly get away from you and you'll be in an unsupportable configuration with very poor performance.

Microsoft does not support reducing the database after it has been created although it is possible.

When the maximum operations database size is set, choose a size that can accept all of the of data that will flow into it without requiring growth, and hold it until it is groomed out. You don't want to see your database to be too large because then you will have slower performance.

The Reporting Server database is a different beast. This database is expected to grow and it can grow quite large. As of this writing, Microsoft has tested the Reporting Server database up to 500 GB. However, you should be concerned with how big it will become in your environment so you can plan for disk space appropriately. You also want the jobs that copy that data from operations to reporting, which are DTS packages, to get all of the data in one copy. If the reporting database size is initially set too small and the "Automatically grow file" increment is not large enough, the reporting database will spend unnecessary time growing itself, which can impact the DTS transfer. By default, the DTS transfer occurs every day at 1:00 a.m. It transfers all of the alert, performance, and event data that is more than five minutes old and is new in the operations database since the last transfer.

The sizing of both databases, the grooming and DTS transfer schedules, and the automatic growth settings all depend on how much data will be flowing into the operations database on a daily basis.

Three types of information flow in and are groomed out of the operations database: alerts, events, and performance data.

- An individual alert is about 6,000 bytes in size, but you will see few of these.
- Events are about 2,500 bytes for each collected event. They occur moderately frequently.
- A single sample of collected performance monitor data is the smallest piece of data, at about 200 bytes, but they occur very frequently.

Leaky Faucet will monitor 50 machines and will install about a dozen management packs. So, the last piece for the calculation is how many of each type of data can be expected each day from each monitored machine. Given that each managed computer will only be running the management packs that are appropriate to it and will be generating different amounts of data, an average of the amount of generated data sent across the machines can be developed.

Microsoft publishes guidelines on how many of each data type is generated per agent-managed computer per day. These guidelines can be found in the MOM 2005 Deployment Planning Guide.

## *Alerts*



Approximately four alerts per managed computer, per day

*Events*

Approximately 200 events per managed computer, per day

*Performance*

Approximately seven performance counters per minute

The guidelines, data, and number of machines are the base to create a rough estimate of the amount of data written to the operations database per day. This is summarized in [Table 2-1](#).

Table 2-1. Amount of data collected per day and written to the operations database

Type	Size (bytes)	Number per day	Number of machines	Bytes per day	MB per day
Alerts	6,000	4 per machine	50	1,200,000	1.1
Events	2,500	200 per machine	50	25,000,000	24
Performance	200	10,000 per machine	50	100,000,000	95

About 120 MB of data per day will be written to the operations database. Now, since the operations database must have 40 percent of free space for the SQL reindex to complete successfully, multiply 120 MB by 1.4 to yield the minimum size of the operations database that keeps collected data for only one day, which is 168 MB.

Since the minimum size for the operations database is 300 MB, Leaky Faucet can keep about two days' worth of data online just by accepting the minimums. Because so little data is coming in, Leaky Faucet can easily increase data retention to a total of 12 days and create a 2-GB operations database. The operations database transaction logs are automatically created to be 20 percent of the database size, so in this case the transaction logs are about 200 MB.

The reporting database stores the data differently than the operations databaseit includes more indexes, which allows faster searching during reporting operations. So, to store the same data in the reporting database, about twice as much space is consumed than in the operations database. The 120 MB per day into the operations database turns in about 240 MB per day in the reporting database. At this rate, it will take about 100 days to accumulate about 24 GB of data in the reporting database. This gives Leaky Faucet the amount needed to size the disk for the reporting database.

In addition, the DTS transfer uses the tempdb database on the SQL Server, which requires about 200 MB space. The DTS transfer also involves the reporting database transaction logfiles, which requires space equal to about five to six times the amount of data that is being transferredroughly 600 MB (120x5, plus some) for the transaction logs. During the reporting setup, the reporting database size



is increased to 2 GB and the transaction log size is set to 500 MB. After setup is complete, the reporting database is configured (its name in SQL Server is SystemsCenterReporting) to grow in 200 MB increments.

Max chooses to make the operations database 2,048 MB. Setup automatically sets the transaction log size to 410 MB or 20 percent of the operations database size. After MOM 2005 is set up and running, Max will change the default grooming values to 12 days. See [Chapter 7](#) for more details on database configuration.

## 2.3.3. Security and Other Configuration Considerations

With the planning finished, there are a few more details to confirm before moving on to the installation. Specifically, decisions about security and remote operations need to be made.

### 2.3.3.1. Secure reports

The Leaky Faucet CFO requires secure access for the executive-level reports. This is a two-fold request: the reports cannot be accessible to the public and they must be viewed over a secure channel. Because the reporting console is strictly web-based, the MOM 2005 Reporting web site can be SSL-secured and accessed over port 443 using the HTTPS protocol. For ease of installation, the IIS server should be SSL secured prior to the installation of MOM 2005 Reporting. This is done by obtaining and applying an SSL certificate to the IIS server.

Access to reports is restricted by setting the appropriate permissions on the folders that contain the reports or on the reports themselves in the Reporting console. To meet the second part of the CFO's requirement, a new folder will be created on the MOM 2005 Reporting web site to hold the sensitive reports. Group permissions will be applied to the folder, thereby restricting access (see [Chapter 8](#)).

### 2.3.3.2. Remote Web console

At Leaky Faucet, the remote site support administrators require easy access to the information on the machines that they help manage. There aren't enough managed computers at these locations to warrant a remote management server or additional management groups (see [Figure 2-1](#)). Leaky Faucet has chosen to make the MOM 2005 Web console available to the remote site support administrators. The IIS server that hosts the web site will be located at the central headquarters, with a firewall between that location and all remote sites.

The Web console uses port 1272, so the Leaky Faucet administrator, Max, must ensure that port is open on the Internet Security and Acceleration Server (ISA) firewall. If further security was needed, Max could make the Web console web site available via application publishing in ISA.

To filter the machines for the remote site support administrators, Operator console scopes are configured in the Administrator console (see the '[Console Scopes](#)' section in [Chapter 6](#)). This is the association of AD user accounts groups with MOM 2005 computer groups. This scope is applied in both the Operator console and the Web console.

### 2.3.3.3. Remote agent installation and agent management

Leaky Faucet must manage two types of agents: agents inside the remote site and internet firewalls, and the agents outside of the firewalls. The steps for installing both types of agents are covered in [Chapter 3](#).

Installing and managing an agent across a firewall is only slightly more complicated and requires two additional preparatory steps in the Leaky Faucet environment:

1. Ensure port 1272 is open for the Web console users, and open port 1270, TCP, and UDP on both sets of firewalls. All agent-to-management server communication occurs over port 1270 and is encrypted.
2. Develop a plan to manually install the agents on the remote machines and protect the management servers from rogue agent installations.

In addition to these two steps, there are some management group-wide security and agent management settings that must be planned for. There is no prompt to configure these settings during setup. Most of these settings can be overridden at the individual agent or management server level later on, but some cannot. The following settings cannot be overridden but are relevant to Leaky Faucet's requirements:

#### *Mutual authentication*

In a MOM 2005 management group, agents and management servers can be required to authenticate each other before data and configuration information is exchanged. This is a Kerberos v5 authentication and serves the same purpose in MOM as it does in ADit blocks against man-in-the-middle attacks. This provides additional security for the operations data, which is in line with the CFO's interests.

#### *Reject new manual agent installations*

The MOM 2005 agent is a Windows installer-based *.msi* package. As such, the MOM 2005 agent can be manually installed on any computer simply by launching the *.msi* package. During the installation process, the installer prompts for a management server or group to connect to. If no controls exist at the management group level, anyone with access to the agent *.msi* could add any computer to any management group. When enabled, this setting automatically rejects all new requests from manually installed agents to join the management group.

#### *Communications ports*

By default, the management server and agents communicate over port 1270. This communication is encrypted for security. This port can be changed, but to allow managed computer failover between management servers, all management servers and agents must use the same port.





## 2.4. Testing and Piloting

The implementation of MOM 2005 at Leaky Faucet is a significant event that will touch every server. To reduce the risk of something going wrong during the implementation, the MOM administrators have gathered the business and technical requirements, planned out their design, estimated how much data will be collected per day, planned out database sizes, and prepared the environment for inter-component communication. The next step is to test MOM 2005 and figure out what changes are needed. This is very hard to do until there is something to work with.

The director of IT requires that the following tasks be completed before the production rollout is authorized:

1. Take every reasonable measure to prevent outages and mistakes during the rollout.
2. Adapt the current support procedures to take advantage of new features. Prepare the support staff to use this new tool.
3. Ensure that the new solution meets the needs of all the key stakeholders before production implementation.

The Windows team creates a test environment that mirrors the production environment on critical points of configuration, but not scale. It consists of a single management server, a SQL Server that will house both the operations database and the reporting database and web site, and one of each type of server to be managed. All of the consoles will be installed on the management server as well as on desktops. The remote site support administrators will help test the overall security and Web console functionality.

Next, the test plan is reviewed to make sure the deliverables are covered. This plan starts after MOM 2005 has been successfully installed and configured. During piloting, the following must be accomplished:

1. Deploy agents to one of each type of server. This will provide a sample of the type of information that MOM will produce. Agents must be deployed to servers that are outside of the firewalls. Procedures for uninstalling agents and testing agent failover between management servers must be developed.
2. Install the management packs for all of the applications that will be monitored. Import the management packs one at a time so the impact of each on the monitored system and on MOM can be seen.
3. Configure notifications via email and pager.
4. Work on the process to resolve alerts so all team members use the same process every time.

This will ease the transition of alerts from one team member to another or between teams.

5. Practice backing up and restoring MOM 2005 and management packs.
6. Configure security on the reports and teach the CFO how to navigate the Reporting console.
7. Populate domain-level MOM groups with key stakeholders that are participating in the pilot to ensure appropriate access to the correct objects.
8. Send all team members to formal MOM 2005 administration training.

Leaky Faucet sets a start and end date for the pilot to avoid scope creep and prevent the pilot from rolling straight into the production implementation. The pilot duration is estimated to take 21 days (training aside). At the end of this time, Leaky Faucet will take all of the feedback from the key stakeholders, as well as all that they learned, and modify the design if necessary. The rollout plan will be finalized in documented form. This primary deliverable from the pilot will help reassure everyone.





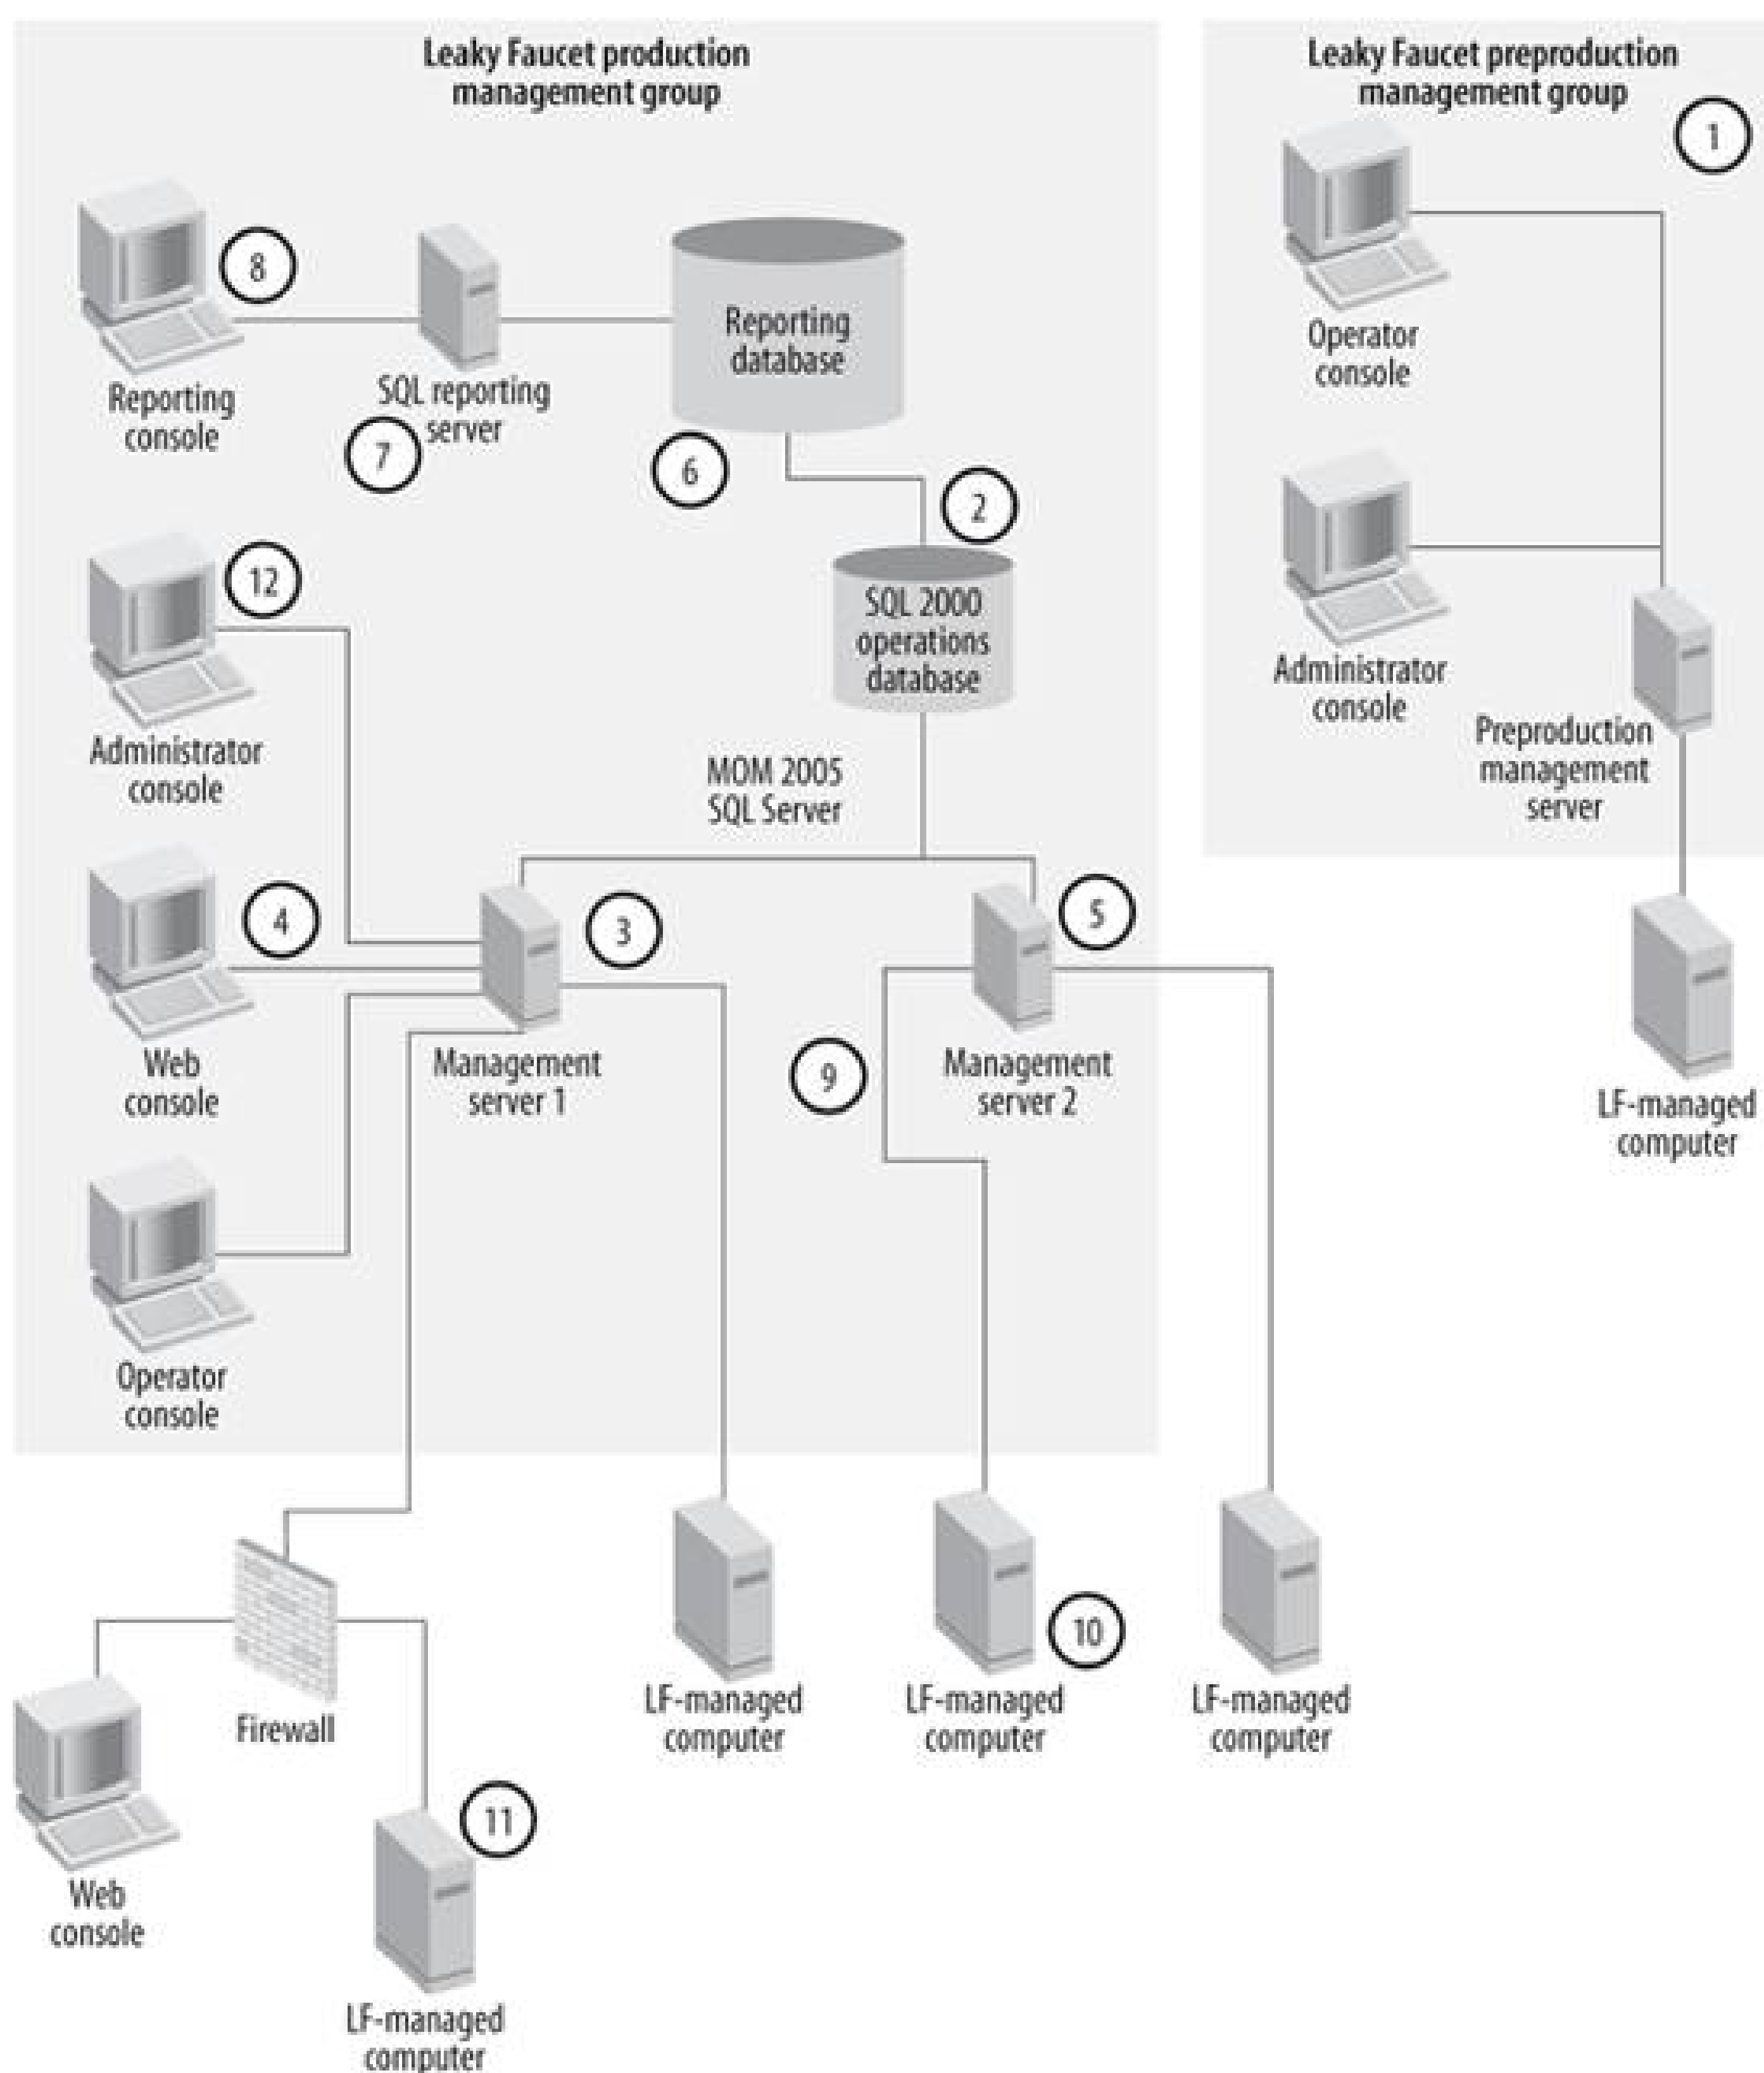
## 2.5. Implementation

The goal of all the planning, testing, and preparation is to make this last step, the implementation, as boring as possible. The best thing that can be said after any technology and process implementation is that it was no big deal and everything just worked. As mentioned at the beginning of the chapter, this is an easy task if you have done your planning. Microsoft documented the steps that the setup wizard goes through however, there are critical database and management server installation steps that should be examined. The planning is already done, so you have the necessary responses for the setup of the individual components.

The order of the component setup is just as important as the setup of the individual components themselves. [Figure 2-3](#) is the build process for Leaky Faucet's MOM 2005 environment.

Figure 2-3. The MOM 2005 build order at Leaky Faucet





Leaky Faucet's MOM 2005 build consists of two separate management groups, the preproduction and the production groups LKFPreProd and LKFProd.

Before step 1 of the build, the management server action account and the DAS account are created, as well as the desired domain groups. Because there is a domain-wide AD group policy for password expiration, a special OU is created for service accounts. This account does not have the password expiration policy applied to it since the "do not inherit group policy" feature is enabled.

The 12-step build process is as follows:

1. Build the LKFPreProd management group, which has all features installed on one computer. The installation occurs in this order:
  - a. Install Windows Server 2003 and required hotfixes and service packs.
  - b. Install IIS 6.0.
  - c. Install SQL 2000 SP3A standard edition.
  - d. Install MOM 2005 Edition. Here they are prompted for the management group name and the management server action account.
  - e. Deploy agents to desired servers and monitor. When prompted for the credentials to install the agents, the domain admin account is given. The agent action account is set to be the local system account.
2. Install and configure the MOM 2005 operations database on the SQL Server. Be sure to install SQL Server SP3A; do not install SQL Server SP4, because it will cause the prerequisite checker to fail at the first install of MOM 2005. SQL SP4 can be installed on the MOM database server after MOM 2005 SP1 has been installed. The server has been named LKFProdOpsDB. Using the MOM 2005 Edition, only the operations database component is installed. During this installation, Leaky Faucet is prompted for the size and location of the operations database, its transaction logs, the management group name LKFProd and the DAS account. Mutual authentication is enabled.
3. Install and configure the first management server. Ensure IIS 6.0 is installed. The first management server has been named LKFProdMS1 and the management server is selected, as well as the the Administrator, Operator, and Web console components. This is the server that the remote site support administrators will connect to when accessing the Web console. During installation Leaky Faucet is prompted for the operations database instance, the management server action account, and the DAS account.
4. Load the Operator and Administrator consoles onto the team's desktops. In the MOM 2005 setup, select user interfaces only. Access each of the consoles to test connectivity and to examine any alerts that may have been raised by the agents that are automatically installed on the management server.
5. Install the second management server, LKFProdMS2, by following the same procedures as for LKFProdMS1 (step 3). The team opts not to install the Web console on this server for security reasons. If LKFProdMS1 goes down, the Web console can be quickly loaded on LKFProdMS2.
6. Install and configure SQL 2000 on the Reporting Server:
  - a. Install SQL Server Standard edition SP3A.
  - b. Install IIS 6.0 and an SSL certificate.
  - c. Install SQL Reporting Services with SP1.

7. Install and configure MOM 2005 Reporting. During this installation, Max is prompted for the name of the server that SQL Reporting Services are installed on (LKFPProdReporting), the SQL Server instance that holds the operations database (LKFPProdOpsDB), the local SQL Server instance to install the reporting database to, and the database and transaction log sizes. They are also prompted for the account to connect to the operations database and the reporting database; in both cases they have chosen to use the DAS account.
8. Connect to the reporting console. Max connects to LKFPProdReporting over HTTPS to ensure access is encrypted. Later on, after the agents have been deployed and the management packs are imported, they will start to configure the secure folder for the CFO's reports.
9. Configure the computer discovery rules. At this point, all the installations from the MOM 2005 CD have been completed. The rest of the deployment is agent- and management pack-oriented. The Leaky Faucet team wants a very simple discovery rule that includes all computers in the Leaky Faucet AD domain.
10. Install agents to the local computers. This process is detailed in [Chapter 3](#), but at a high level, it involves:
  - a. Selecting the computers that have been discovered and running the agent installation wizard.
  - b. Choosing the credentials to be used for the installation. These credentials must have full administrative rights on the target computers. Because Leaky Faucet has chosen to use a regular domain account for the management server action account, that account cannot be used. Instead, they provide the domain administrator's account credentials to be used only for the installation. These credentials are stored securely and disposed of after the agent installation is complete.
  - c. Choosing the account to use as the agent action account on the managed computers. The local system default account and the default installation location for the agent files are selected.
11. Install agents on remote site computers. Since port 1270 on the firewalls is open and the "automatically reject manually installed agents setting" has been disabled, the team runs the agent installation setup from the installer `.msi/package` or runs the setup program from the MOM 2005 CD. During the installation, they are prompted for the name of the management group, the installation directory, the management server name, the communication port, and the agent action account, among other things. Back in the Administrator console, the agents are manually approved and the computers are now managed.
12. Import all desired management packs. The team downloads the desired management packs from the web. The management packs, along with their reports, are imported one at a time into MOM 2005. An interval is left between the imports to allow the service discovery process to occur and push the rules out to the agents.

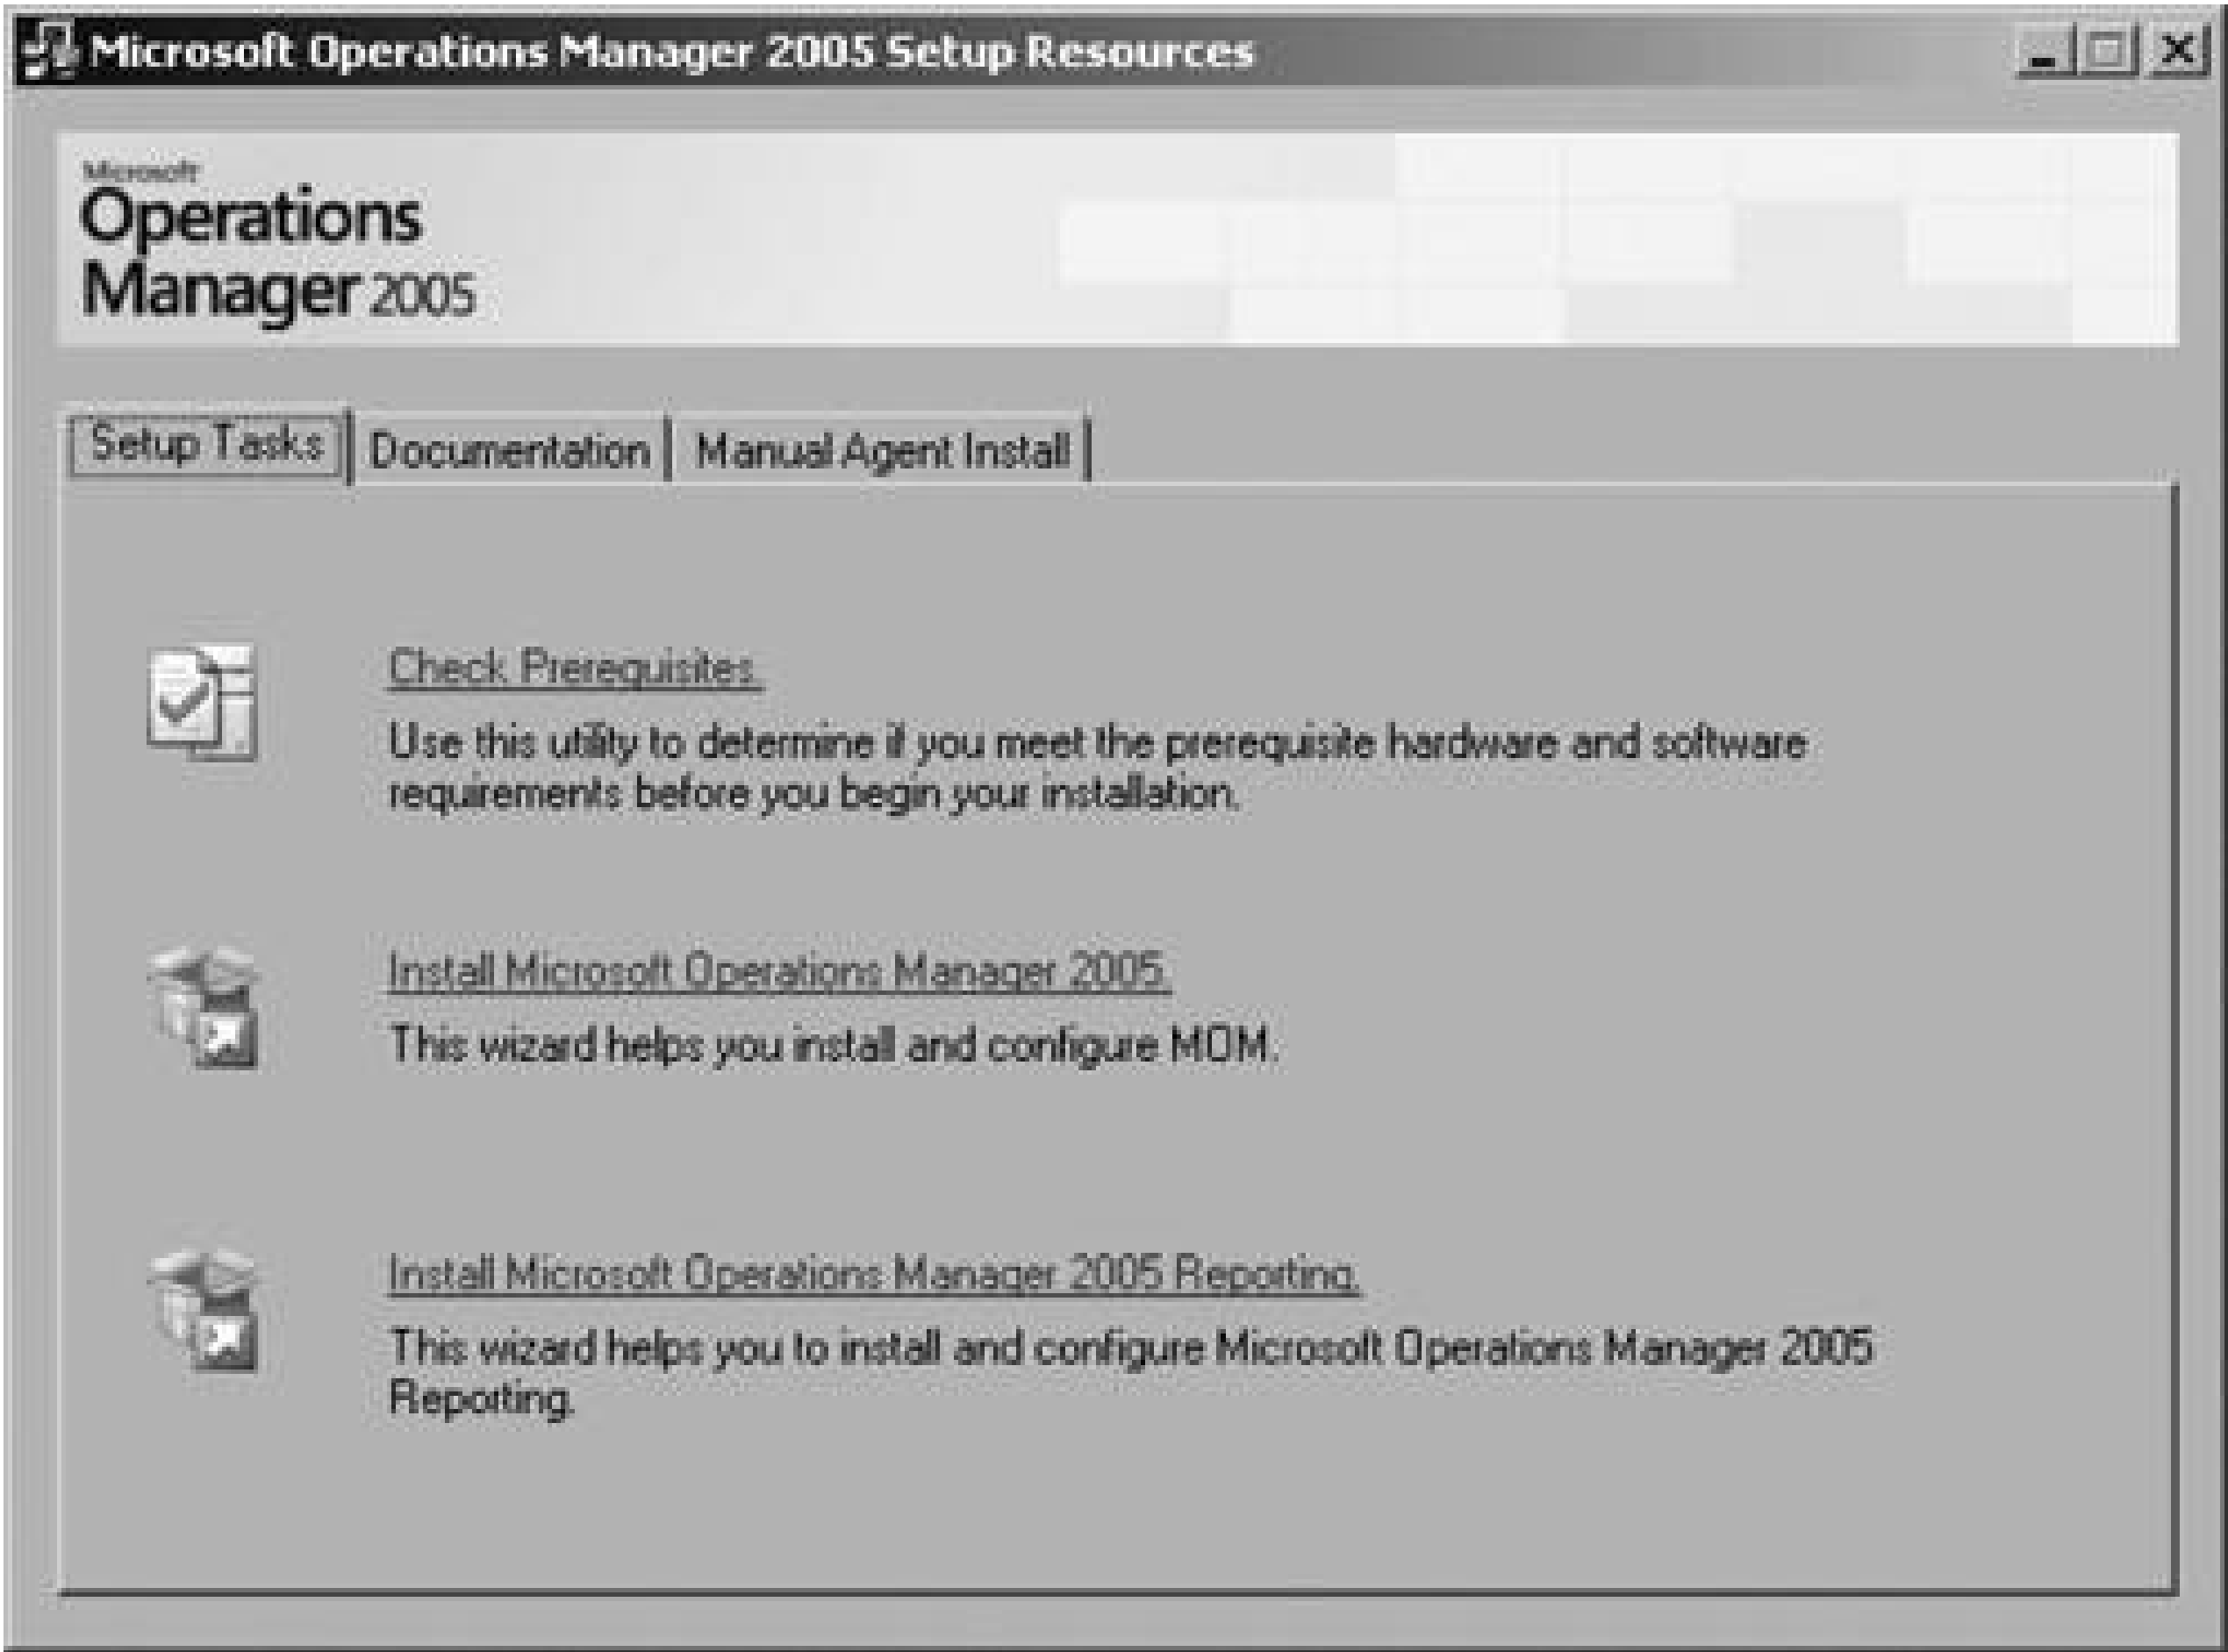


# 2.6. Installation Specifics

The 12-step build list is very broad and covers more of the life cycle than is necessary to simply run and install MOM. The installation of the MOM 2005 operations database and the first management server starts here.

The installation of the production management group starts on the production SQL Server LKFProdOpsDB (step 2 in the 12-step build list). SQL 2000 has already been installed and SP3A applied. Log on with an account that has local administrator rights to the server and launch the MOM 2005 setup from the MOM CD. This brings up the Setup Resources pageshown in [Figure 2-4](#).

Figure 2-4. The MOM 2005 setup start page



This is where all MOM 2005 setup tasks can be initiated, as well as accessing the online documentation and the prerequisite checker. Select the Install Microsoft Operations Manager 2005 option. The next three pages are the welcome page, the end-user license agreement (to which you must agree in order to continue), and the product registrationpage. On the product registration page, enter the 25-digit CD key (if that field is not already populated). Proceed through these three

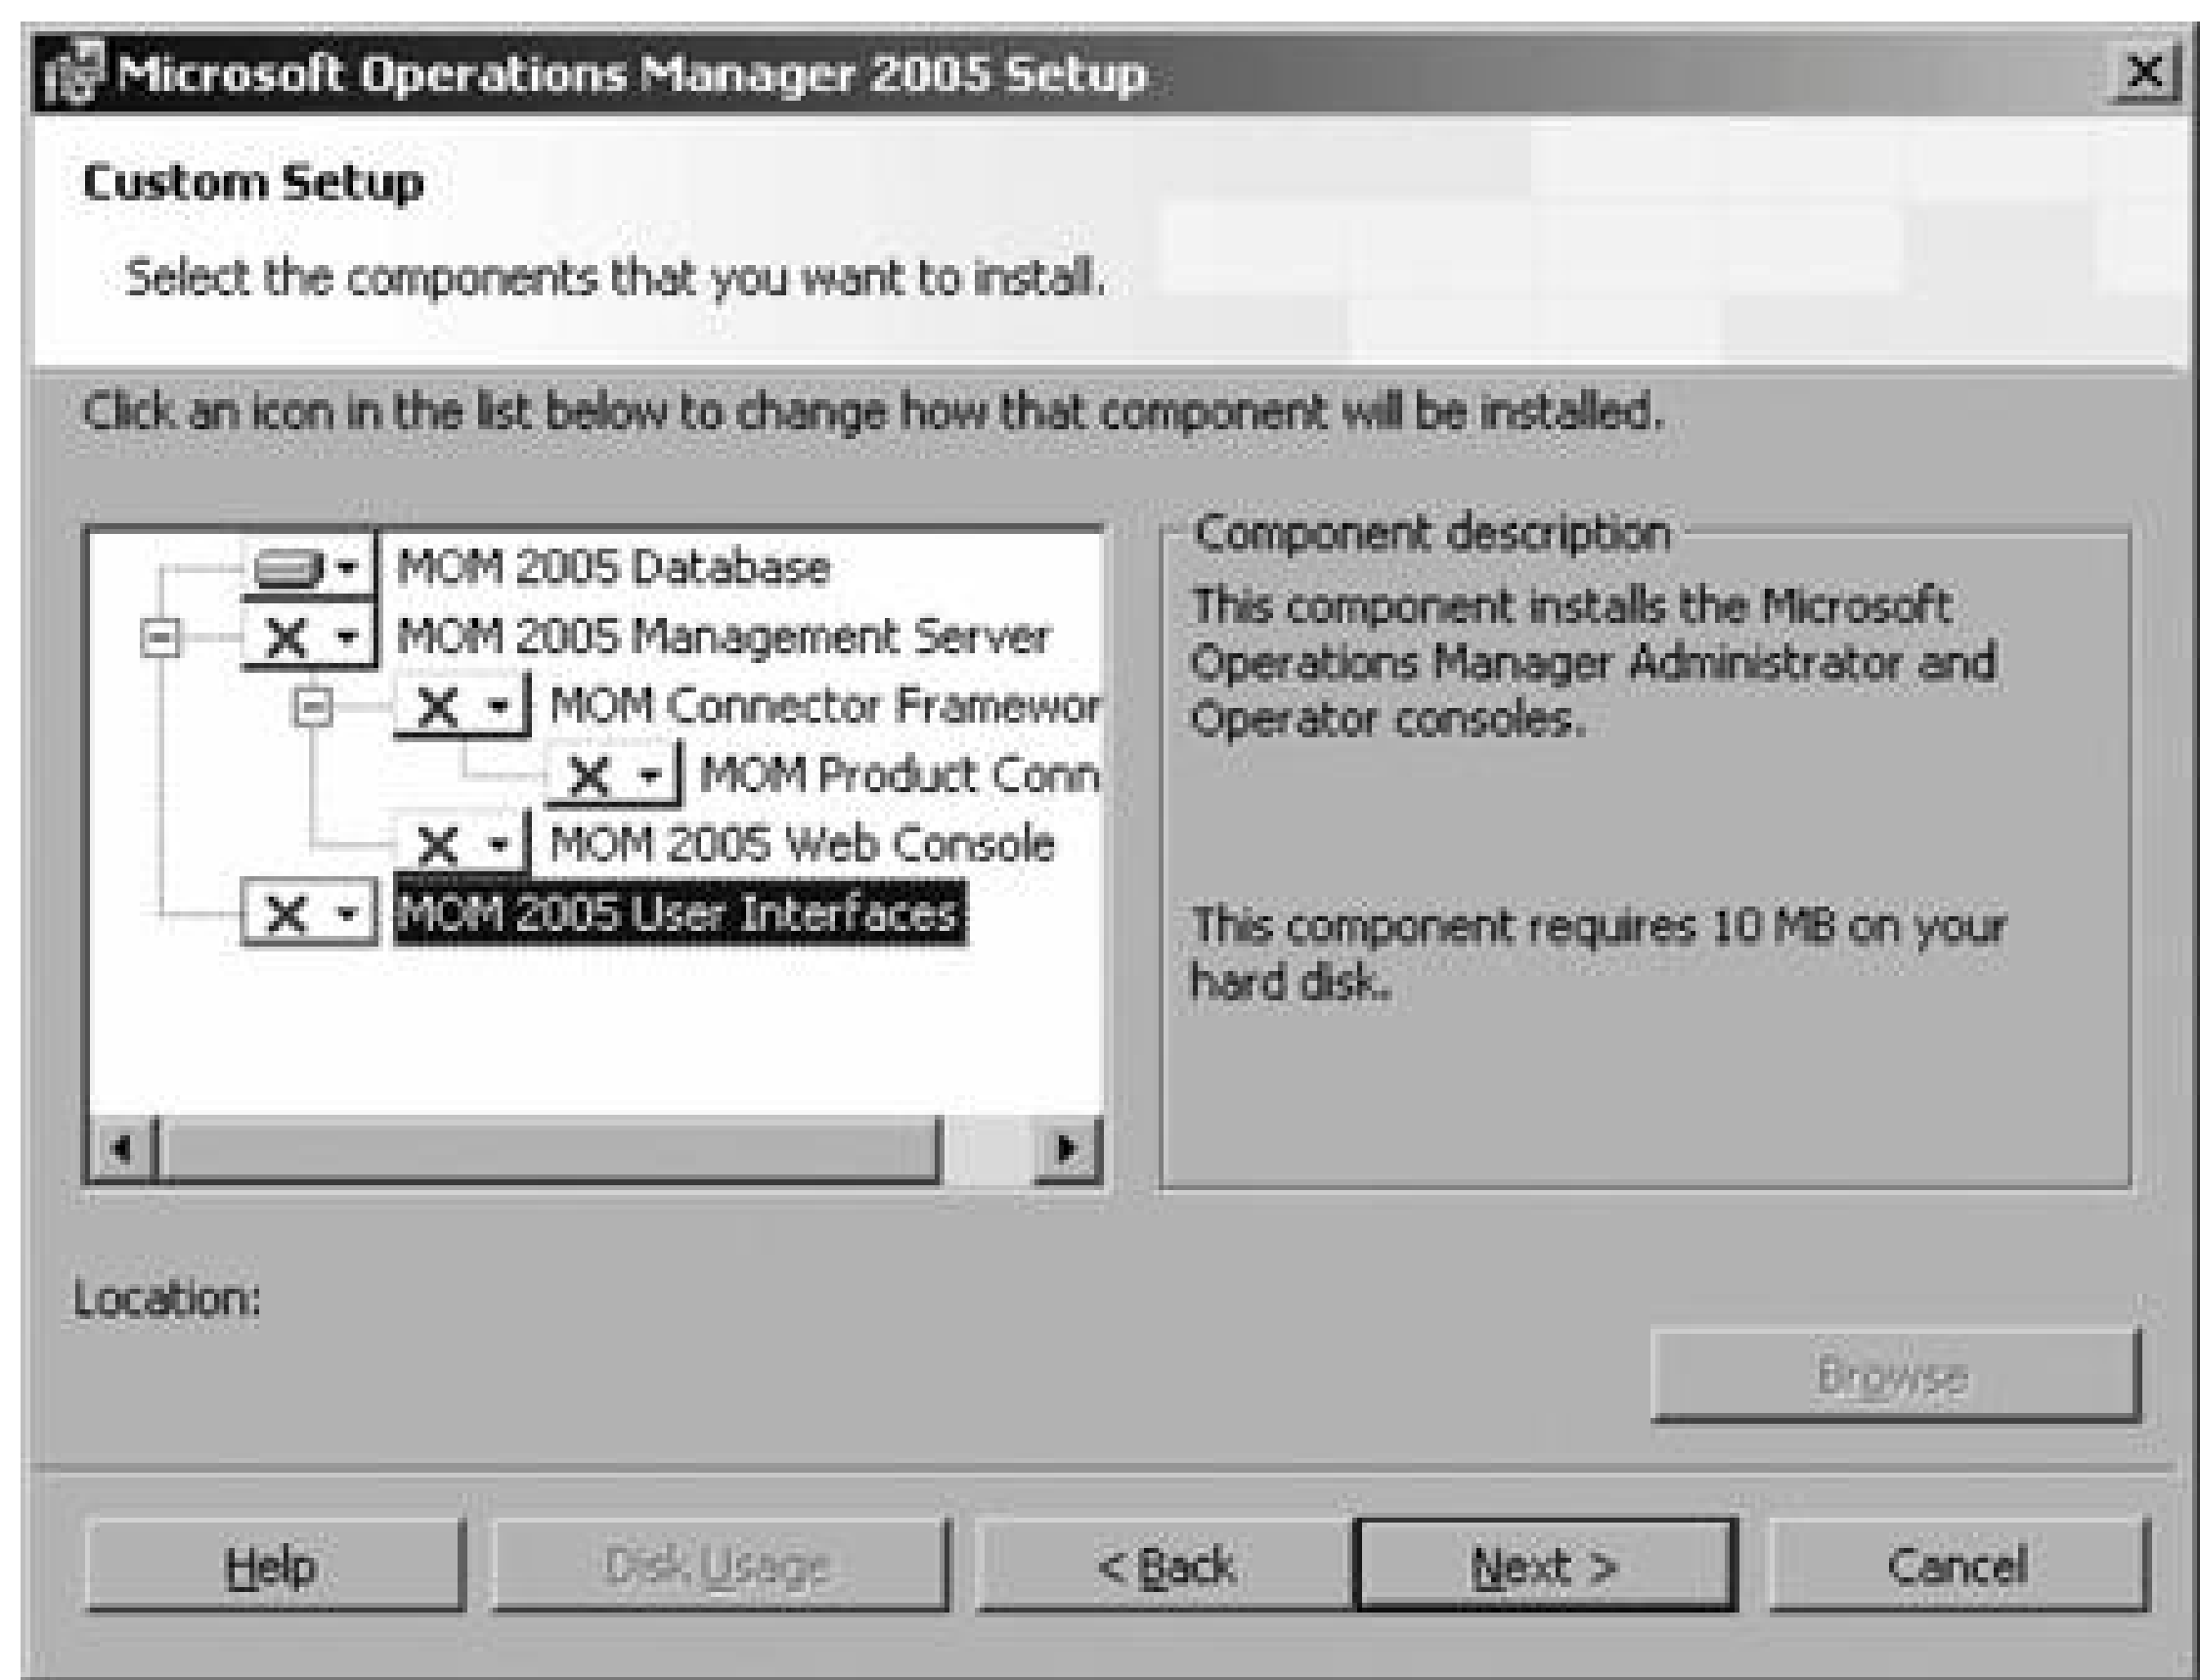
pages to the Installation Options page (see [Figure 2-5](#)). Select the Custom option.

To perform anything but an all-in-one installation (all components on the same server), you must choose the Custom option. This option allows you to install and uninstall individual components on any given server and is the only way to create a management group with distributed components. Clicking Next brings up the Custom Setup page (see [Figure 2-6](#).)

Figure 2-5. Selecting the type of MOM installation to perform



Figure 2-6. The Custom Setup page selections for installing the database only

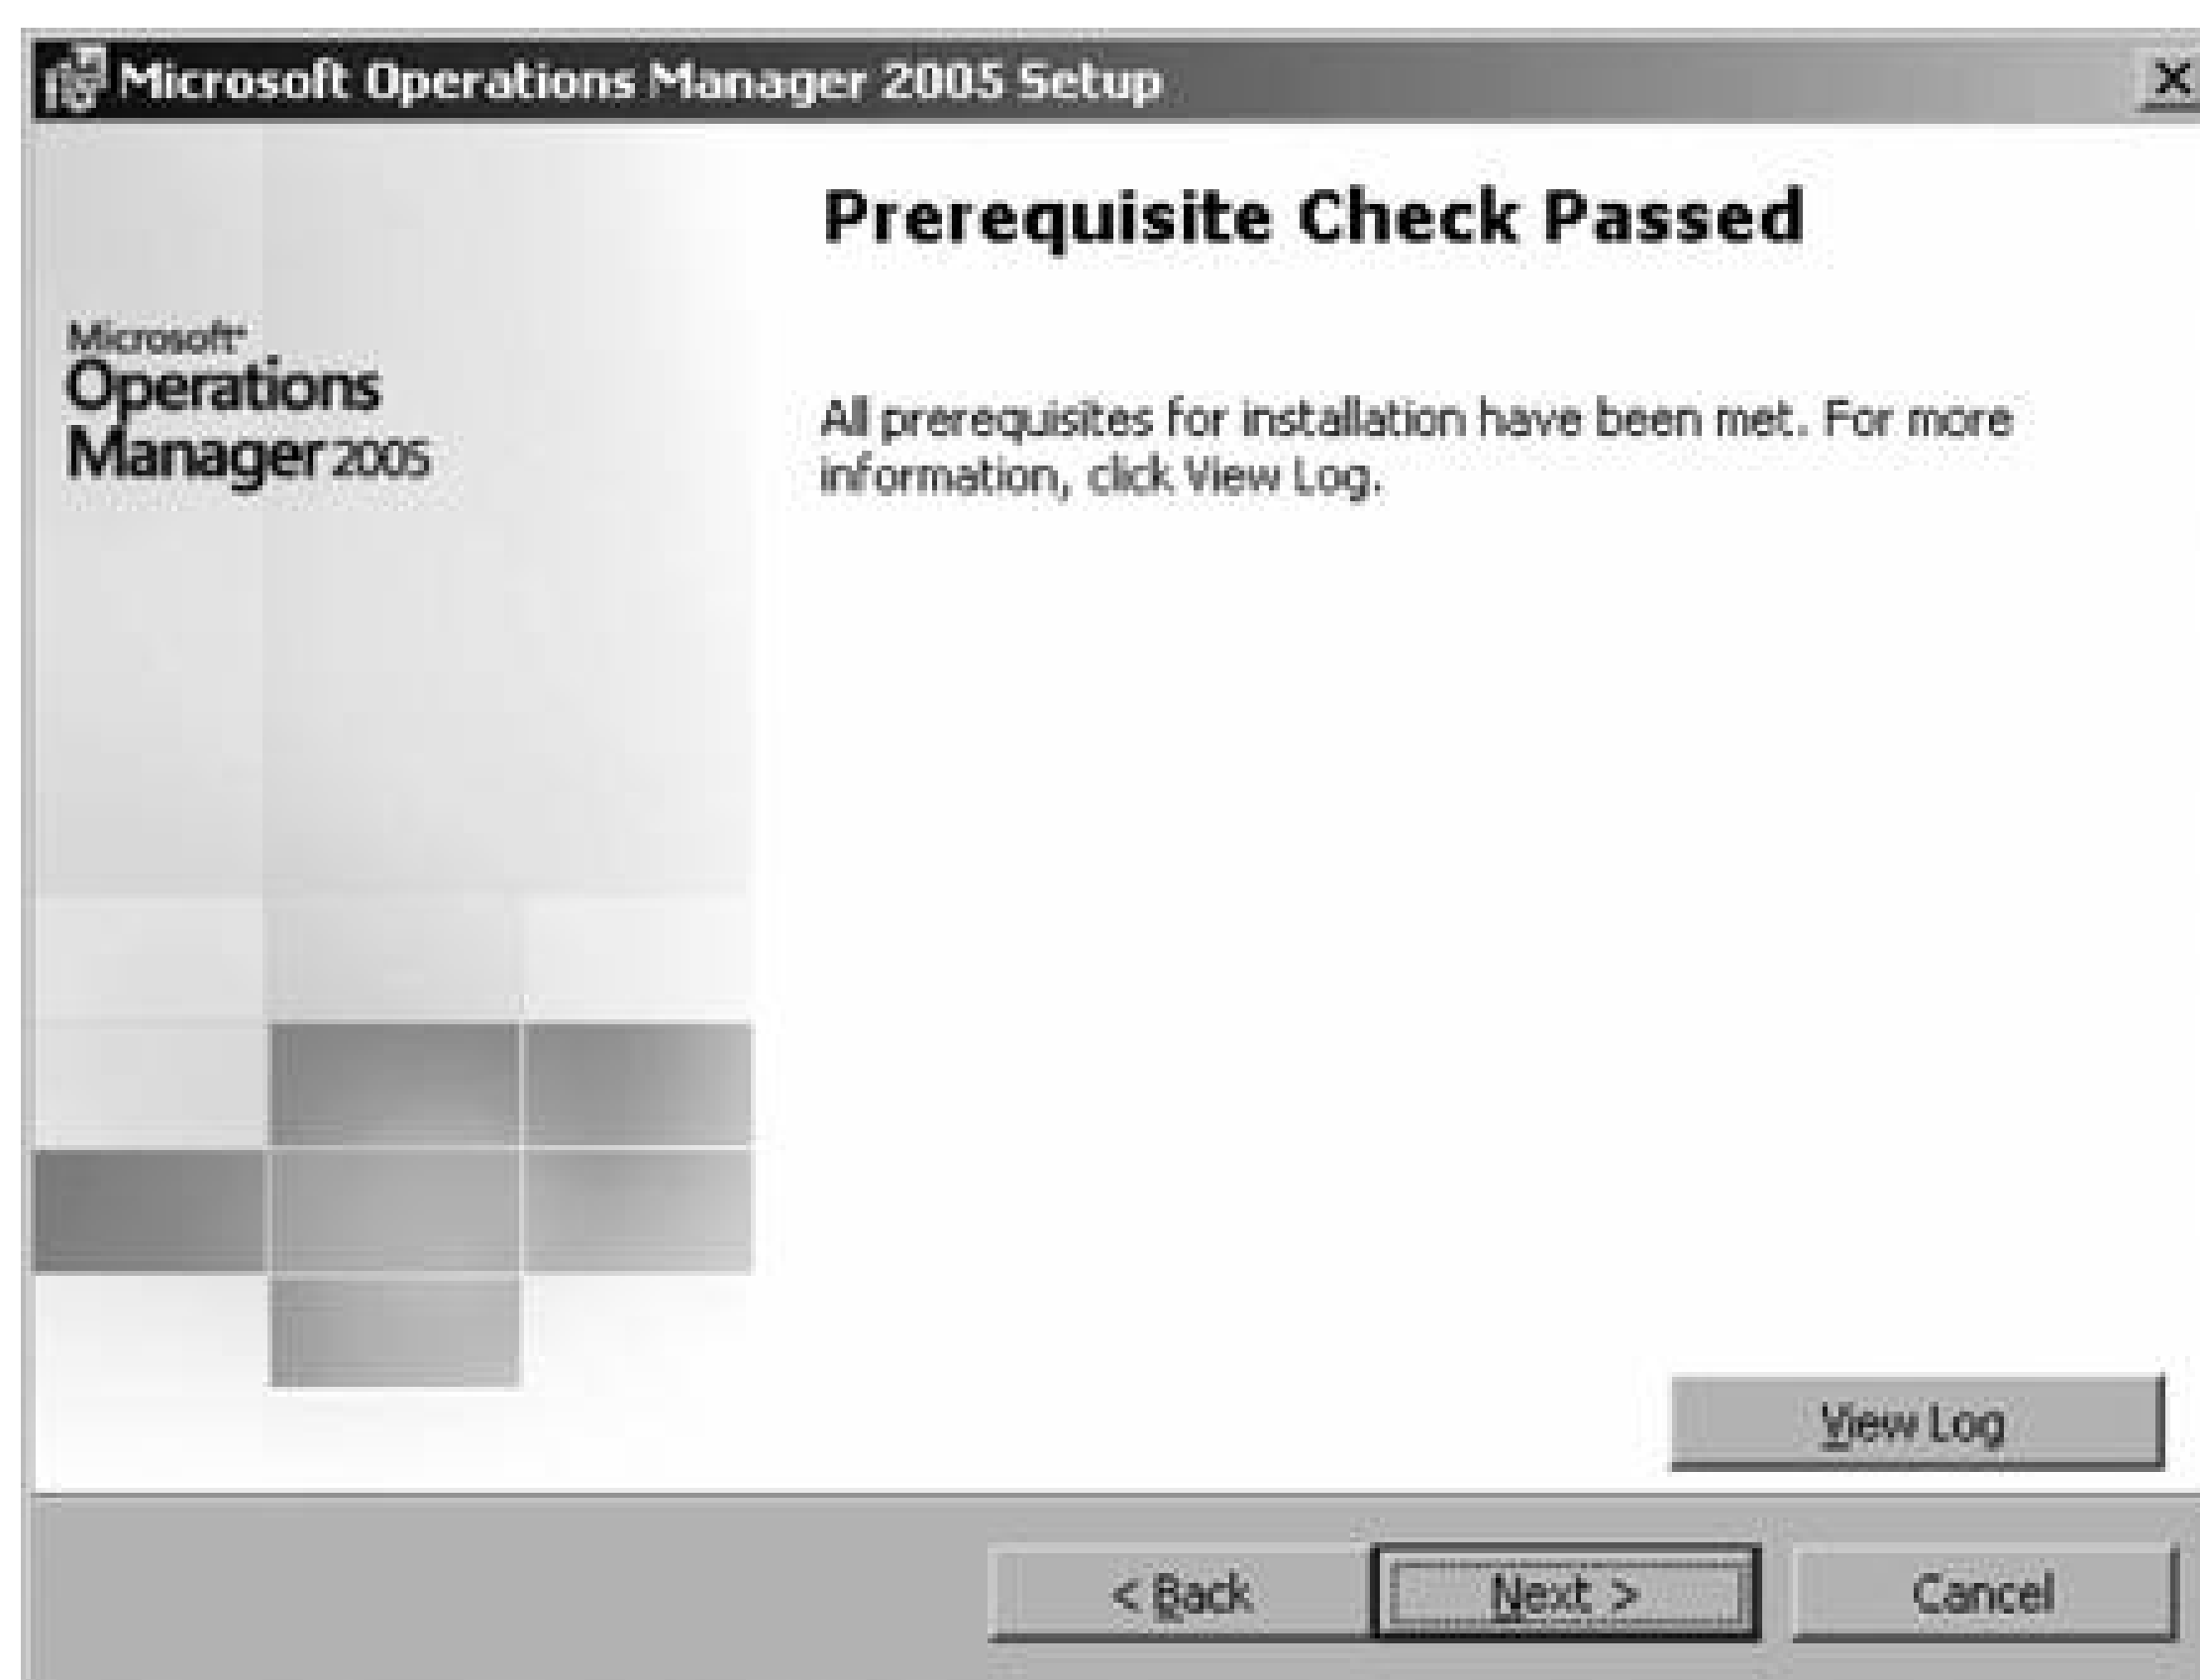


Since the database is being installed on a dedicated machine, select the "This component will not be installed on this server" option for the MOM 2005 Management Server components and the MOM 2005 User Interfaces components. Do this by selecting the down arrow in the box to the left of the component's name.

Components marked with an "X" will not be installed. [Figure 2-6](#) shows the configuration for installing the database only. Proceeding through this page starts the prerequisite checker in the background and the results page appears (see [Figure 2-7](#)).

Figure 2-7. Prerequisites result page





The prerequisite checker can also return a failed or a passed message with warning status depending on the state of the machine. The prerequisite checker performs different evaluations depending on the components that are being installed. For example, with a database-only installation, it checks for the presence of SQL Server 2000 with SP3A, Windows 2000 SP4 or Windows 2003, and confirms that the server is running on a 32-bit architecture. You can view the results log, which is rendered through Internet Explorer by selecting the View Log button. Since this server passed the prerequisite checker, proceed through by clicking Next. This brings up [Figure 2-8](#), the SQL Server Database Instance page.

SQL Server 2000 can run multiple instances on a single physical server. Since this server is dedicated to this role, there are no other instances of SQL Server 2000 running. Select the default value, which is the default instance on the local server; for this example it is *homemomserver3*. If multiple instance of SQL Server 2000 are running on *homemomserver3*, then use the drop-down menu to select the server. Click next to bring up the Database and Log File Options page (see [Figure 2-9](#)).

As calculated earlier in this chapter, the production management group will have a 2 GB database. If your server is configured with a separate drive for applications, install the databases there instead of accepting the default, which places them on the C: drive. The Advanced button lets you change the installation paths. Completing this page takes you to the Management Group Name page (see [Figure 2-10](#)).

Figure 2-8. Select the instance of SQL Server that the operations database will be installed to



Figure 2-9. Set the database size limit and installation paths for the operations database and its transaction log

Carefully select the name of your management group, since it cannot be changed later. Fortunately, the management group name is only evident to administrators and others that have access to the Administrator console and the title bar of the Operator console MMC. In all situations where you might need to use the management group name, such as changing the focus of a console from one management group to another, attach it to the management server rather than the management group. Proceed through this page to the Data Access Server Account page (see [Figure 2-11](#)).

Complete the fields on this page using the domain account created for the DAS account. This then

takes you to the MOM Error Reporting page (see [Figure 2-12](#)).

Figure 2-10. Enter the name that has been chosen for the management group



Figure 2-11. Designate the account to be used for the DAS Account



The choice made in [Figure 2-12](#) has absolutely no effect on your MOM 2005 installation. If you choose not to send the reports, MOM will generate alerts periodically to remind you that error reporting is not enabled. These alerts can be safely resolved and ignored. Click Next to bring up the Active Directory Configuration page (see [Figure 2-13](#)).

Clicking Next brings up the Ready to Install page. From there, the actual installation is started. Once it has started you can only cancel the installation, no parameters can be changed. The database install finishes with the usual installation completed page.

Figure 2-12. Select to enable sending error reports to Microsoft

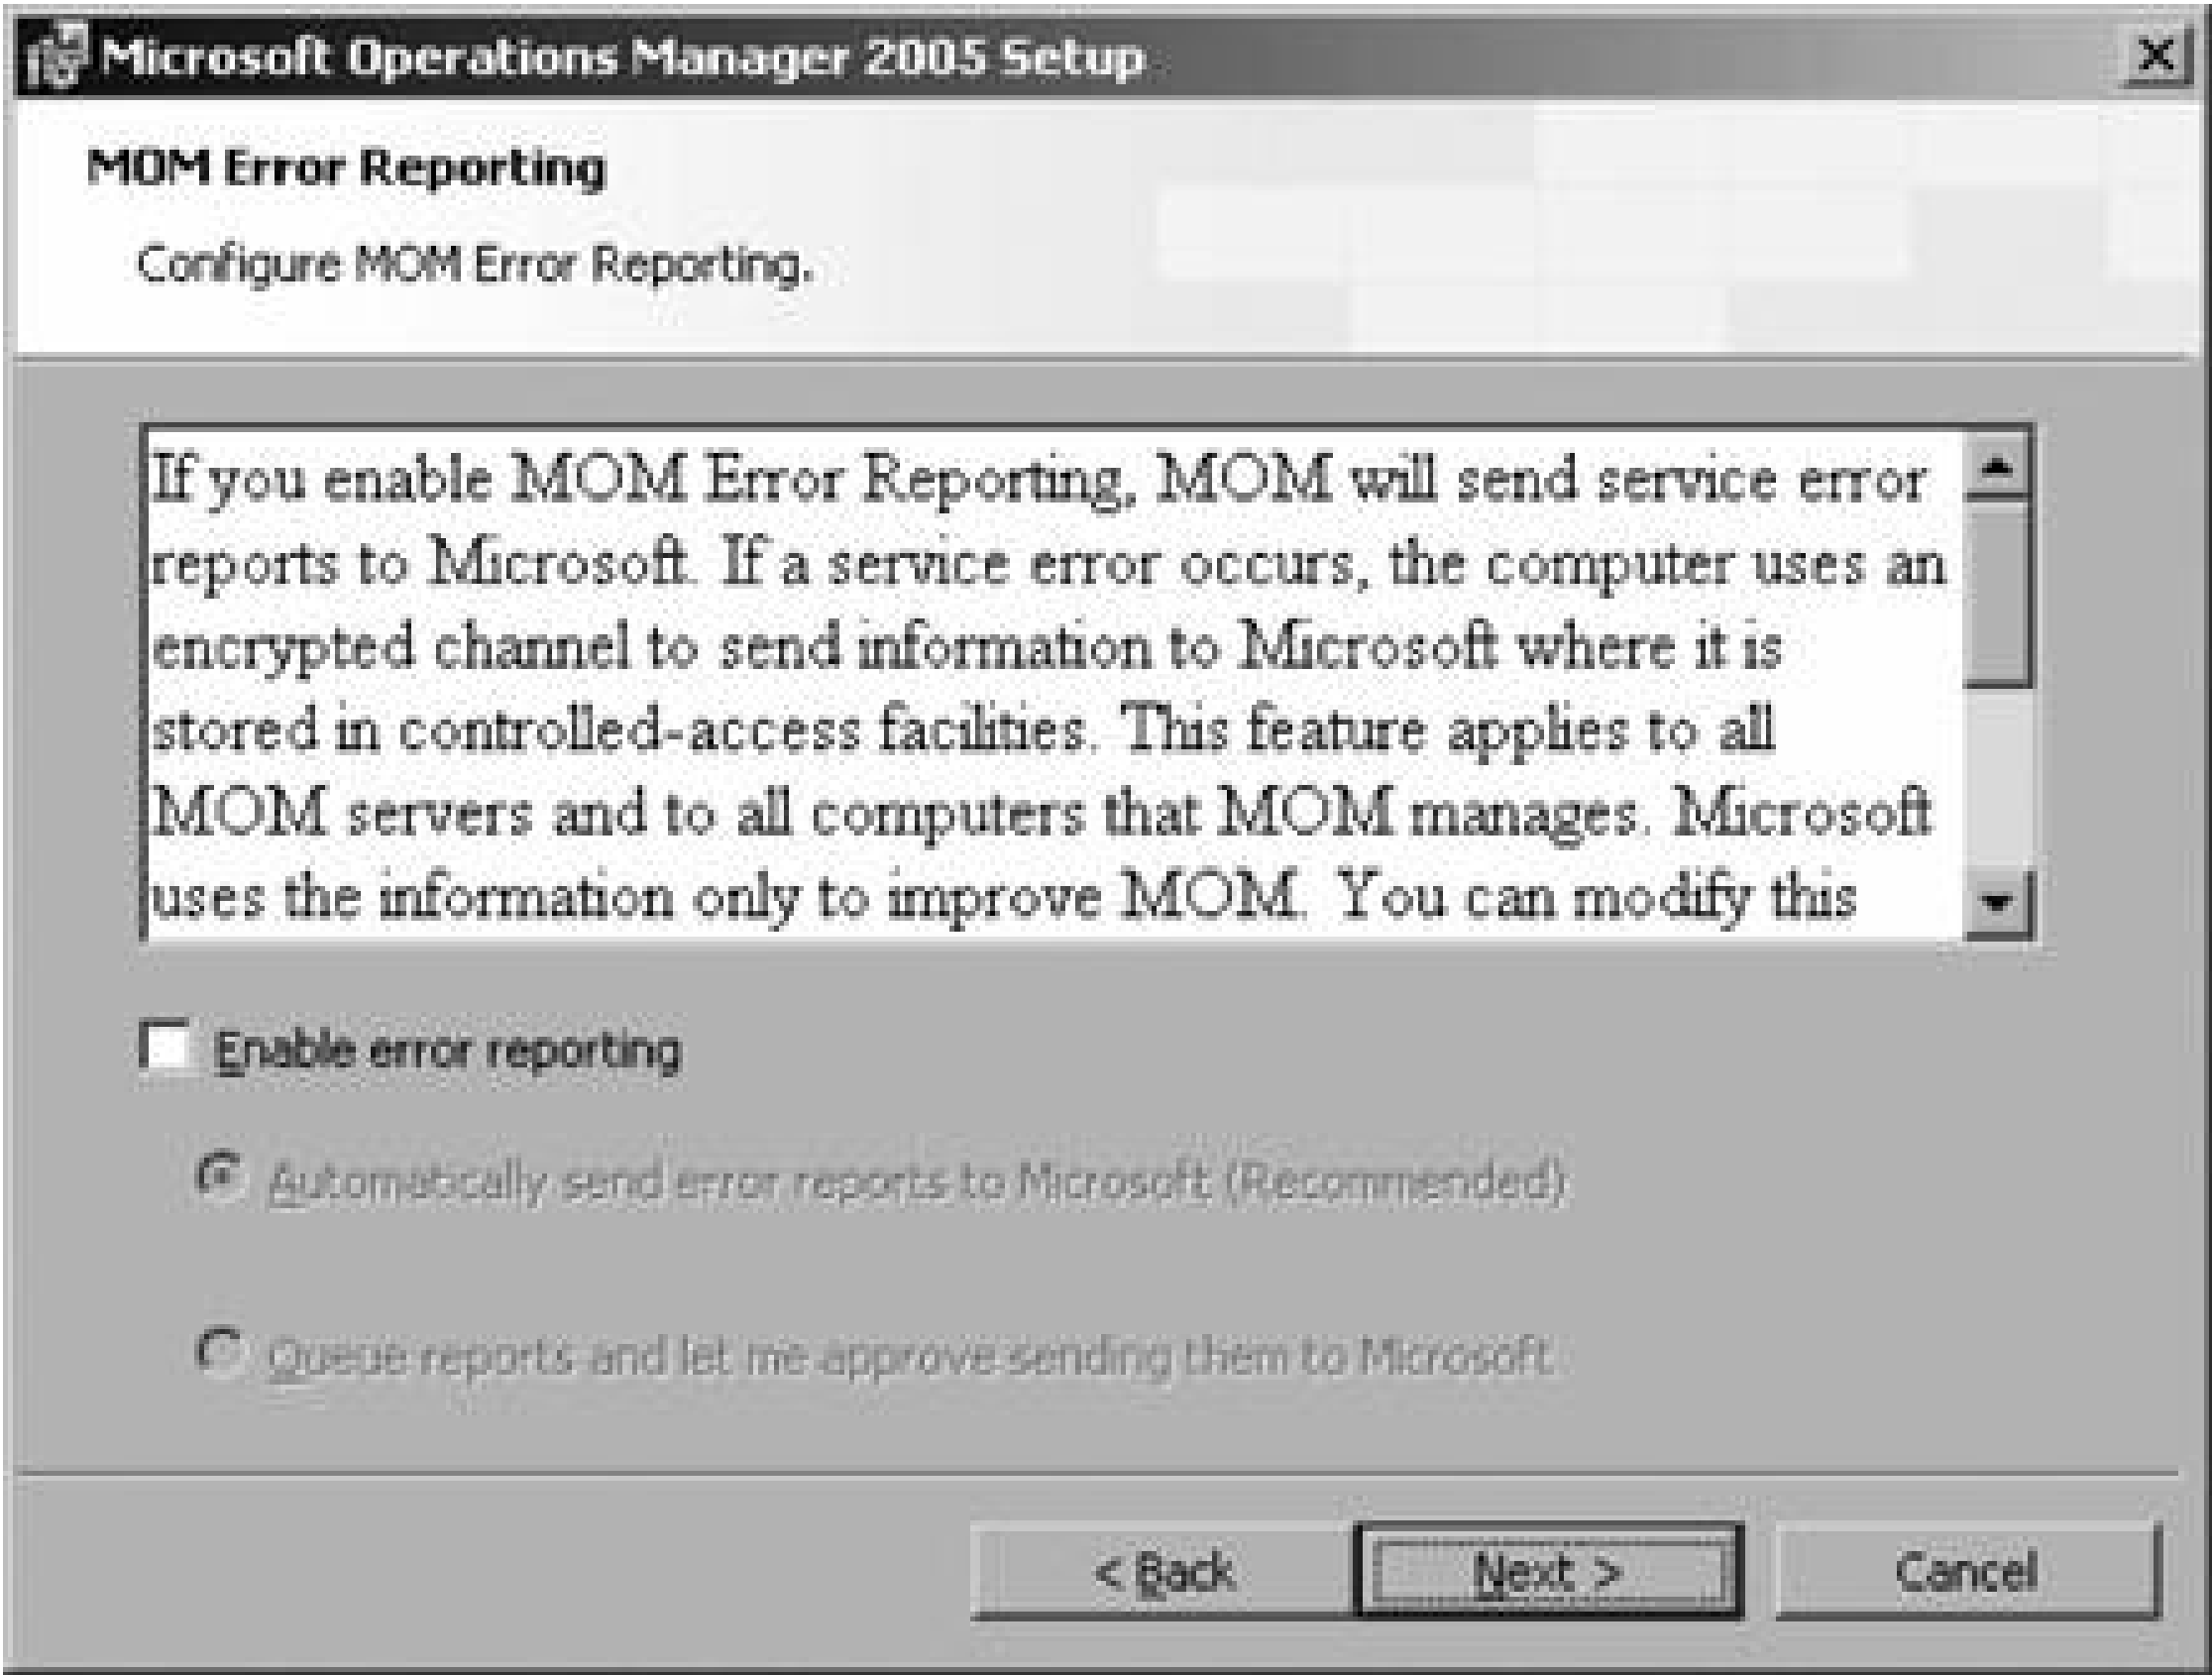


Figure 2-13. By selecting Yes on this page, mutual authentication between agents and management servers is enabled

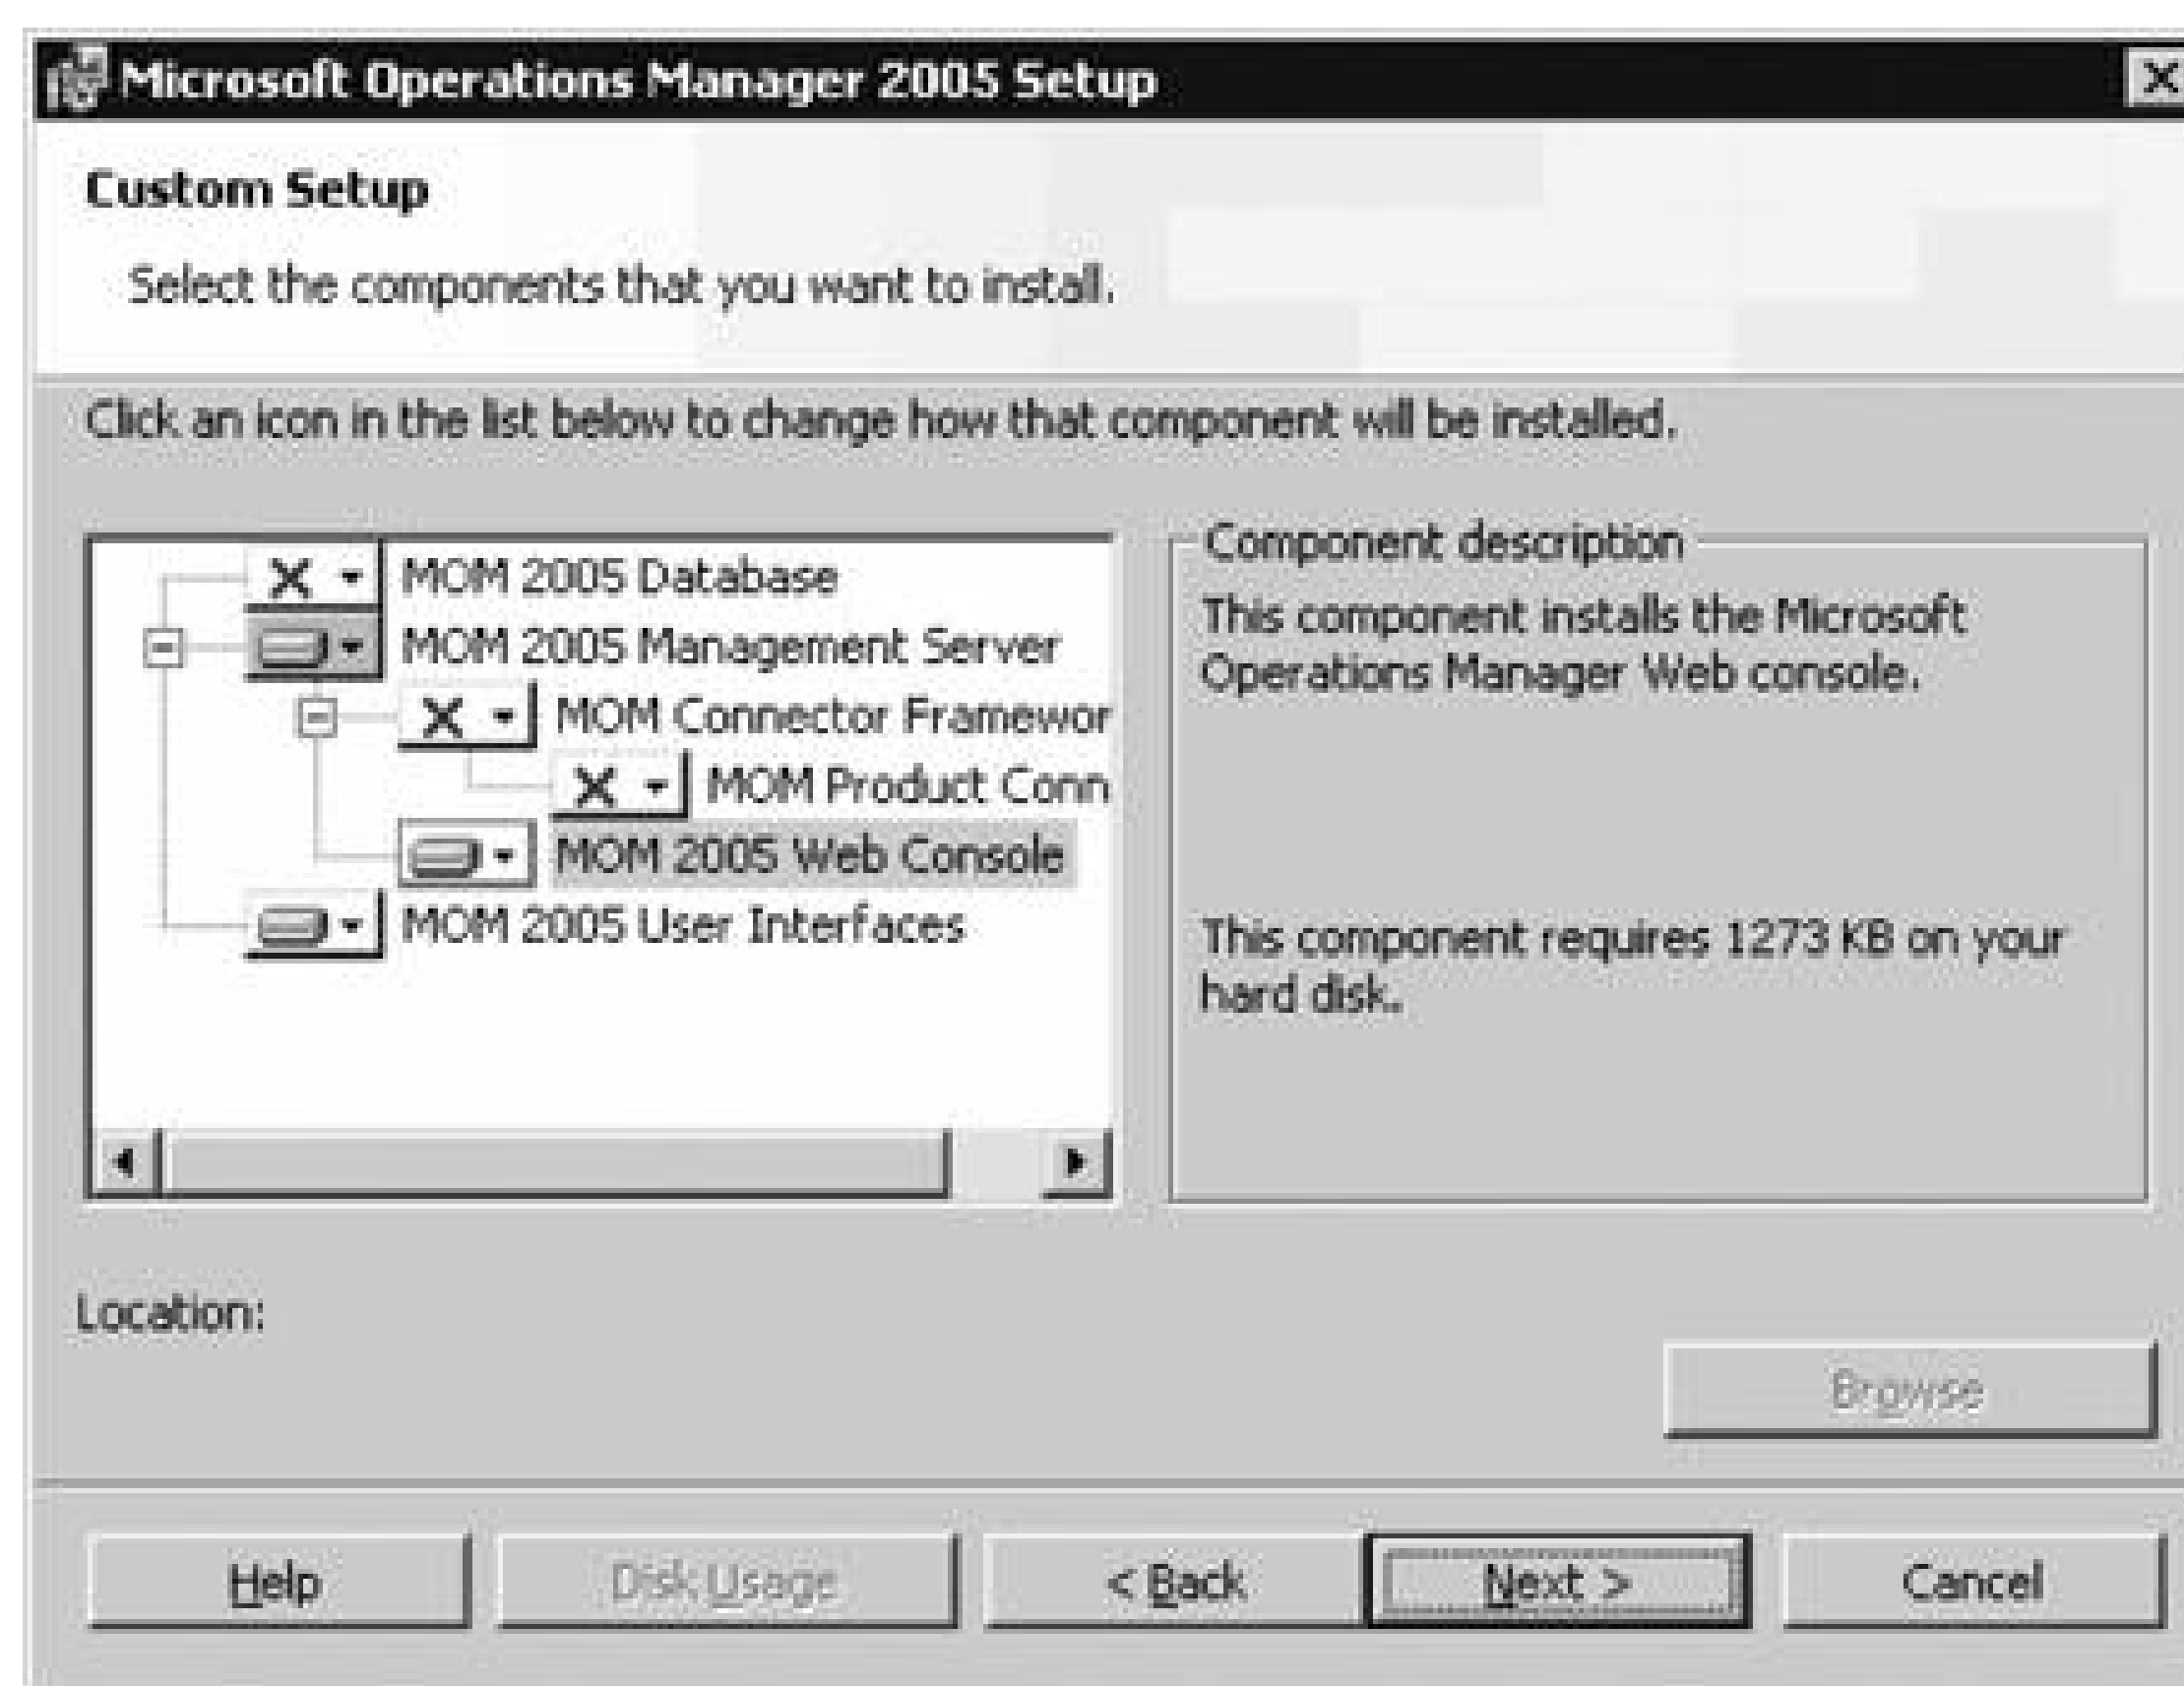


After installing the operations database, install the first management server in the management group. Log on as an administrator to the server that you are installing the management server component on and start the MOM 2005 setup. Because the same setup program is used, many of the same pages are presented.

The following is detailed information on the pages that are different from the database server setup and the management server setup:

- Setup Resources pageSelect the Install Microsoft Operations Manager 2005 option.
- Welcome pageProceed through.
- End-User License AgreementAccept the agreement; proceed through.
- Product RegistrationEnter the appropriate username, organization name, and 25-digit CD key; proceed through.
- Installation OptionsSelect the Custom option; proceed through.
- Custom Setup[Figure 2-14](#) shows the selections for installing the management server, Web console, and UI components. Here the MOM 2005 Database has been deselected, as well as the MOM Connector Framework and the MOM Product Connector.


Figure 2-14. Component configuration for management server installations



- Prerequisite checkerIf this fails to return, view the log and it will tell you exactly what needs to be installed or configured to continue with the installation.
- MOM Database Server InstanceOn this page, enter the name of the server where the MOM 2005 Operations database is installed (see [Figure 2-15](#)). The default port that is used for communications with the database server can be changed here if you have configured SQL to use a port other than 1433. Unless you have specifically changed the port that SQL is listening on, accept the default setting of port 1433.

Figure 2-15. Enter the MOM database server name and accept the default communication port





- Management Server Action Account Provide the credentials for the account to be used as the management server action account; proceed through.
- Data Access Server Account Provide the credentials of the same DAS account that you entered during the operations database install; proceed through.
- Ready to Install Proceed through.
- Completing Installation Accept the default selection of "Start the MOM Administrator console," and click finish.



## 2.7. Summary

Planning and implementing MOM 2005 starts with gathering the business and technical requirements that the operations management solution is expected to fulfill. These requirements come from the key stakeholders in your company. Be sure to include the technical staff, business management, executive management, and anyone else who has a vested interest.

Once you have these requirements, start the engineering process to map the MOM 2005 features to the requirements. This process results in a MOM 2005 design document. The next step is to build a proof-of-concept installation for MOM 2005 to verify your engineering design, gain valuable hands-on experience with the product in your environment, and validate your build and operations procedures. The whole point of the test environment is to mitigate risk.

Complete your MOM 2005 project with the production implementation. This can be a long process, but all the preparation and testing has led to this. Once this is complete, your company moves into monitoring and maintaining activities this is where MOM 2005 will earn its keep.

# Part II: Managing and Using MOM on a Daily Basis

[Chapter 3, Managing Agents](#)

[Chapter 4, Administering Management Packs](#)

[Chapter 5, Administering Global Settings](#)

[Chapter 6, Operator Console](#)

[Chapter 7, MOM 2005 Database Fundamentals](#)

[Chapter 8, MOM 2005 Reporting](#)



# Chapter 3. Managing Agents

[Chapter 2](#) explained planning and deployment of a MOM 2005 infrastructure, but gave light treatment to the deployment and management of agents. This chapter provides the details on how agents function and provides guidance on common tasks that will manage your agents.

Agent management consists of installing and uninstalling tasks, troubleshooting tasks, and updating tasks. Out of these, you will do mostly installing and uninstalling tasks. Like the deployment of the management group components, once an agent is installed on a managed computer there should be little hands-on interaction with it. If the agent is not responding or you need to get other information from it, you may have to do some troubleshooting.

MOM 2005 is very flexible with managing agents and monitoring computers. It has the ability to monitor computers in a wide variety of network environments and under various security requirements. For example, you can deploy agents across domain boundaries, into a workgroup, or across a firewall or a slow WAN link. MOM 2005 provides this flexibility through groups of settings that can be applied to individual management servers or across the management group as a whole.

Since understanding the different combinations of these settings can be difficult, agent management will be discussed in the context of the Leaky Faucet environment where it is useful. Through these examples, you'll learn the considerations for deploying and managing agents across firewalls, into untrusted networks, and how to multi-home an agent.

## 3.1. Agent Functions

The agent is the engine of execution in MOM 2005. Fundamentally, it does two things: it collects data from the data providers on the managed computer and it executes the management pack instructions given by the agent on the management server. When there is a match between the collected data and the parameters described in the rules, the agent takes whatever action has been defined in the rule. This could be to generate an alert, pass the data to the management server, or to run a script or piece of managed code as a response.

A MOM agent runs as the *MOMService.exe* process using the local system account as its security context. The MOMService process is responsible for communicating with the management server, reading and writing to local event logs, reading WMI data, and performing file transfer sends and receives. When an action needs to be taken on the local computer, the *MOMService.exe* spawns a new process named *MOMHost.exe*. The *MOMHost.exe* runs in the security context of the agent action account. *MOMHost.exe* collects performance counter collection and runs scripts, managed code, and batch responses. Other actions are executed by the *MOMService.exe* impersonating the agent action account. The maximum number of responses that can run simultaneously can be configured, but the default is five. The isolation provided by running responses under separate process adds greatly to the stability and reliability of agents. [Figure 3-1](#) shows where an agent sits and the tasks it performs relative to the local computer and the management server.

Figure 3-1. MOM 2005 agent details

The agent always initiates communications with the management server over TCP port 1270 and UDP port 1270 for heartbeats. Agents are designed to run independently of their owning management server in case there is a communications failure or if the management server is otherwise unavailable. By default, agents cache up to 3 MB of data on disk and this can be adjusted to suit your environment. For example, if there are frequent communications failures between the agent and the management server, you could increase this value so that valuable data is not lost. How long 3 MB lasts depends on how many management packs the agent is executing and how active the server is. On average, 3 MB will actually last a long period of time. For example, by using the sizing calculation from the "[MOM 2005 Operations and Reporting Database Planning](#)" section in [Chapter 2](#)--four 6-KB alerts per machine per day, 200 2.5-KB events per machine per day, and 10,000 200-byte performance monitor counter data a single machine generates about 2.75 MB of data per day. So, the default 3 MB will last about a day, but it will vary.

[Table 3-1](#) shows the default values for agents that can be set at a global level to apply across the management group. Some of these settings can be overridden at the management-server level, allowing agent configuration flexibility. These settings control the behavior of agents after they have been installed and act at the agent level.

Table 3-1. Agent settings that can be configured globally

Global agent setting	Default value	Management server override
Request configuration information (rule updates and configuration changes)	1 minute	No
Send heartbeat signal (lets the management server know the agent is alive and working)	10 seconds	No
Number of simultaneous responses allowed	5	Yes
Amount of temporary storage	3,000 KB	Yes
Service status check interval (monitors for changes in status of services)	On; 20 seconds	Yes
Event and performance data buffer duration	30,000 ms	Yes
Alert and response buffer duration	1,000 ms	Yes
Packet size	50 KB	Yes
Maximum amount of data to send per second	1,000 KB	Yes
Collect event binary data (information in the data frame of the properties of a Windows event log event)	Disallowed	Yes

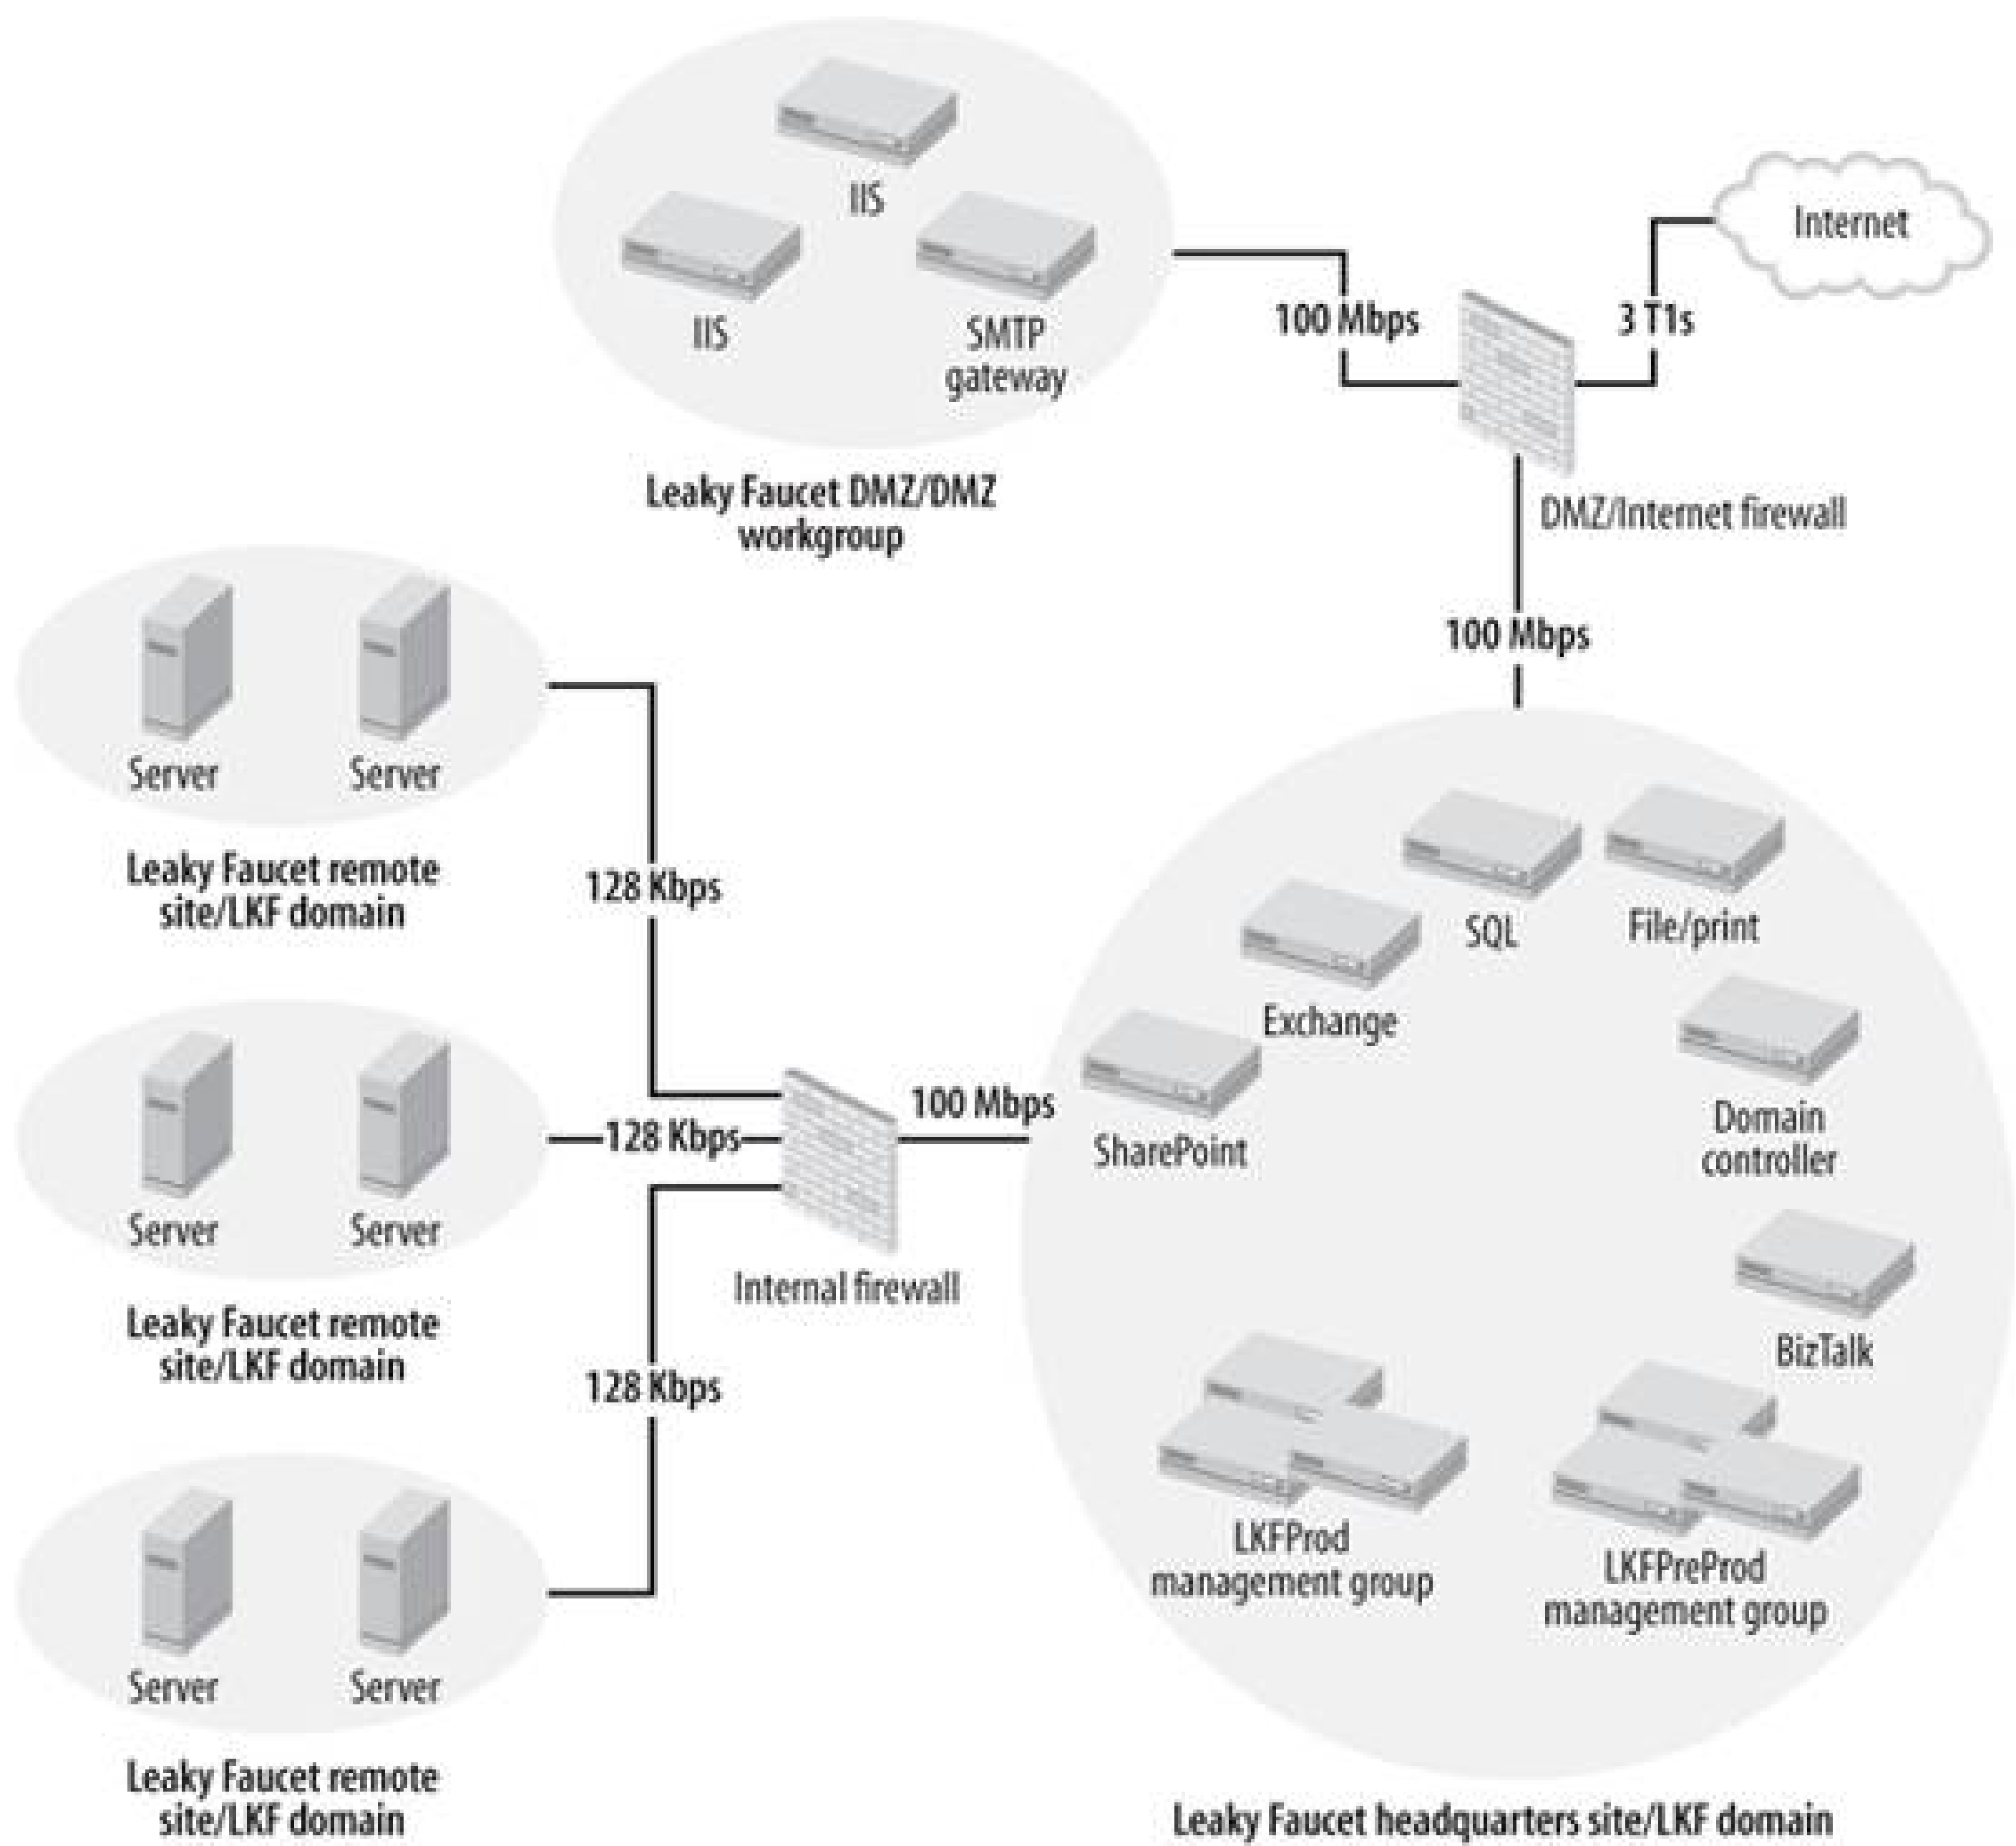


## 3.2. Preparing to Deploy Agents

Before deploying agents , identify the machines you are going to deploy to. Depending where the machines are in your network, you will use different methods to install and manage the agents.

Leaky Faucet has to deploy agents into three different environments: the local LAN, across slow WAN links and firewalls but still in the same domain, and into a workgroup in their DMZ. The Leaky Faucet network is depicted in [Figure 3-2](#).

Figure 3-2. Leaky Faucet's network



Although the method of agent deployment and management will vary based on the target computer and the environment it is in, you will go through the same deployment process and management steps.

1. Identify the machines that you want to monitor, including network location and security context
2. Prepare the machines that will be monitored, ensuring that they have sufficient disk space for the agent to use and that the Windows event logs and application logs are of sufficient size.
3. Set the desired configurations at the management-group level. For example, you can configure the management group to automatically install agents on every computer it becomes aware of, or require administrative approval before installing an agent.
4. Configure MOM 2005 to include or exclude computers in its management scope through *computer discovery rules*. Computer discovery is the process in which MOM actively scans the network for computers that satisfy some preconfigured criteria. For example, you can configure the computer discovery rules to search for all computers in a domain using the domain name and the any character(s) wildcard (\*), or only those with a certain string of characters in the name. Regular expressions and Boolean regular expressions are available to help craft your search query. Once MOM finds these computers, it adds an entry for them to the operations database and displays that information in the Administrator console.
5. Deploy the agent, or start agentless management on a computer.
6. Ensure the agent is working by checking the Operator console for entries and data that are specific to the monitored computer.
7. Import management packs into the management group if this is the first time agents are being deployed. The agents will then classify the computer into roles based on discovery information and apply application-specific rules.
8. Manage and troubleshoot the agent through its life cycle, updating it manually if necessary.
9. Uninstall the agent when the monitored computer is decommissioned.

## 3.3. Deploying and Managing Agents in a Trusted LAN

The trusted LAN network and security environment is the easiest to work with. The management groups and target computers are all in the same domain and there is high-speed, reliable connectivity between them with no port restrictions. [Figure 3-3](#) shows the Leaky Faucet headquarters site in the LKF domain. You will probably deploy the bulk of your agents into this type of environment.

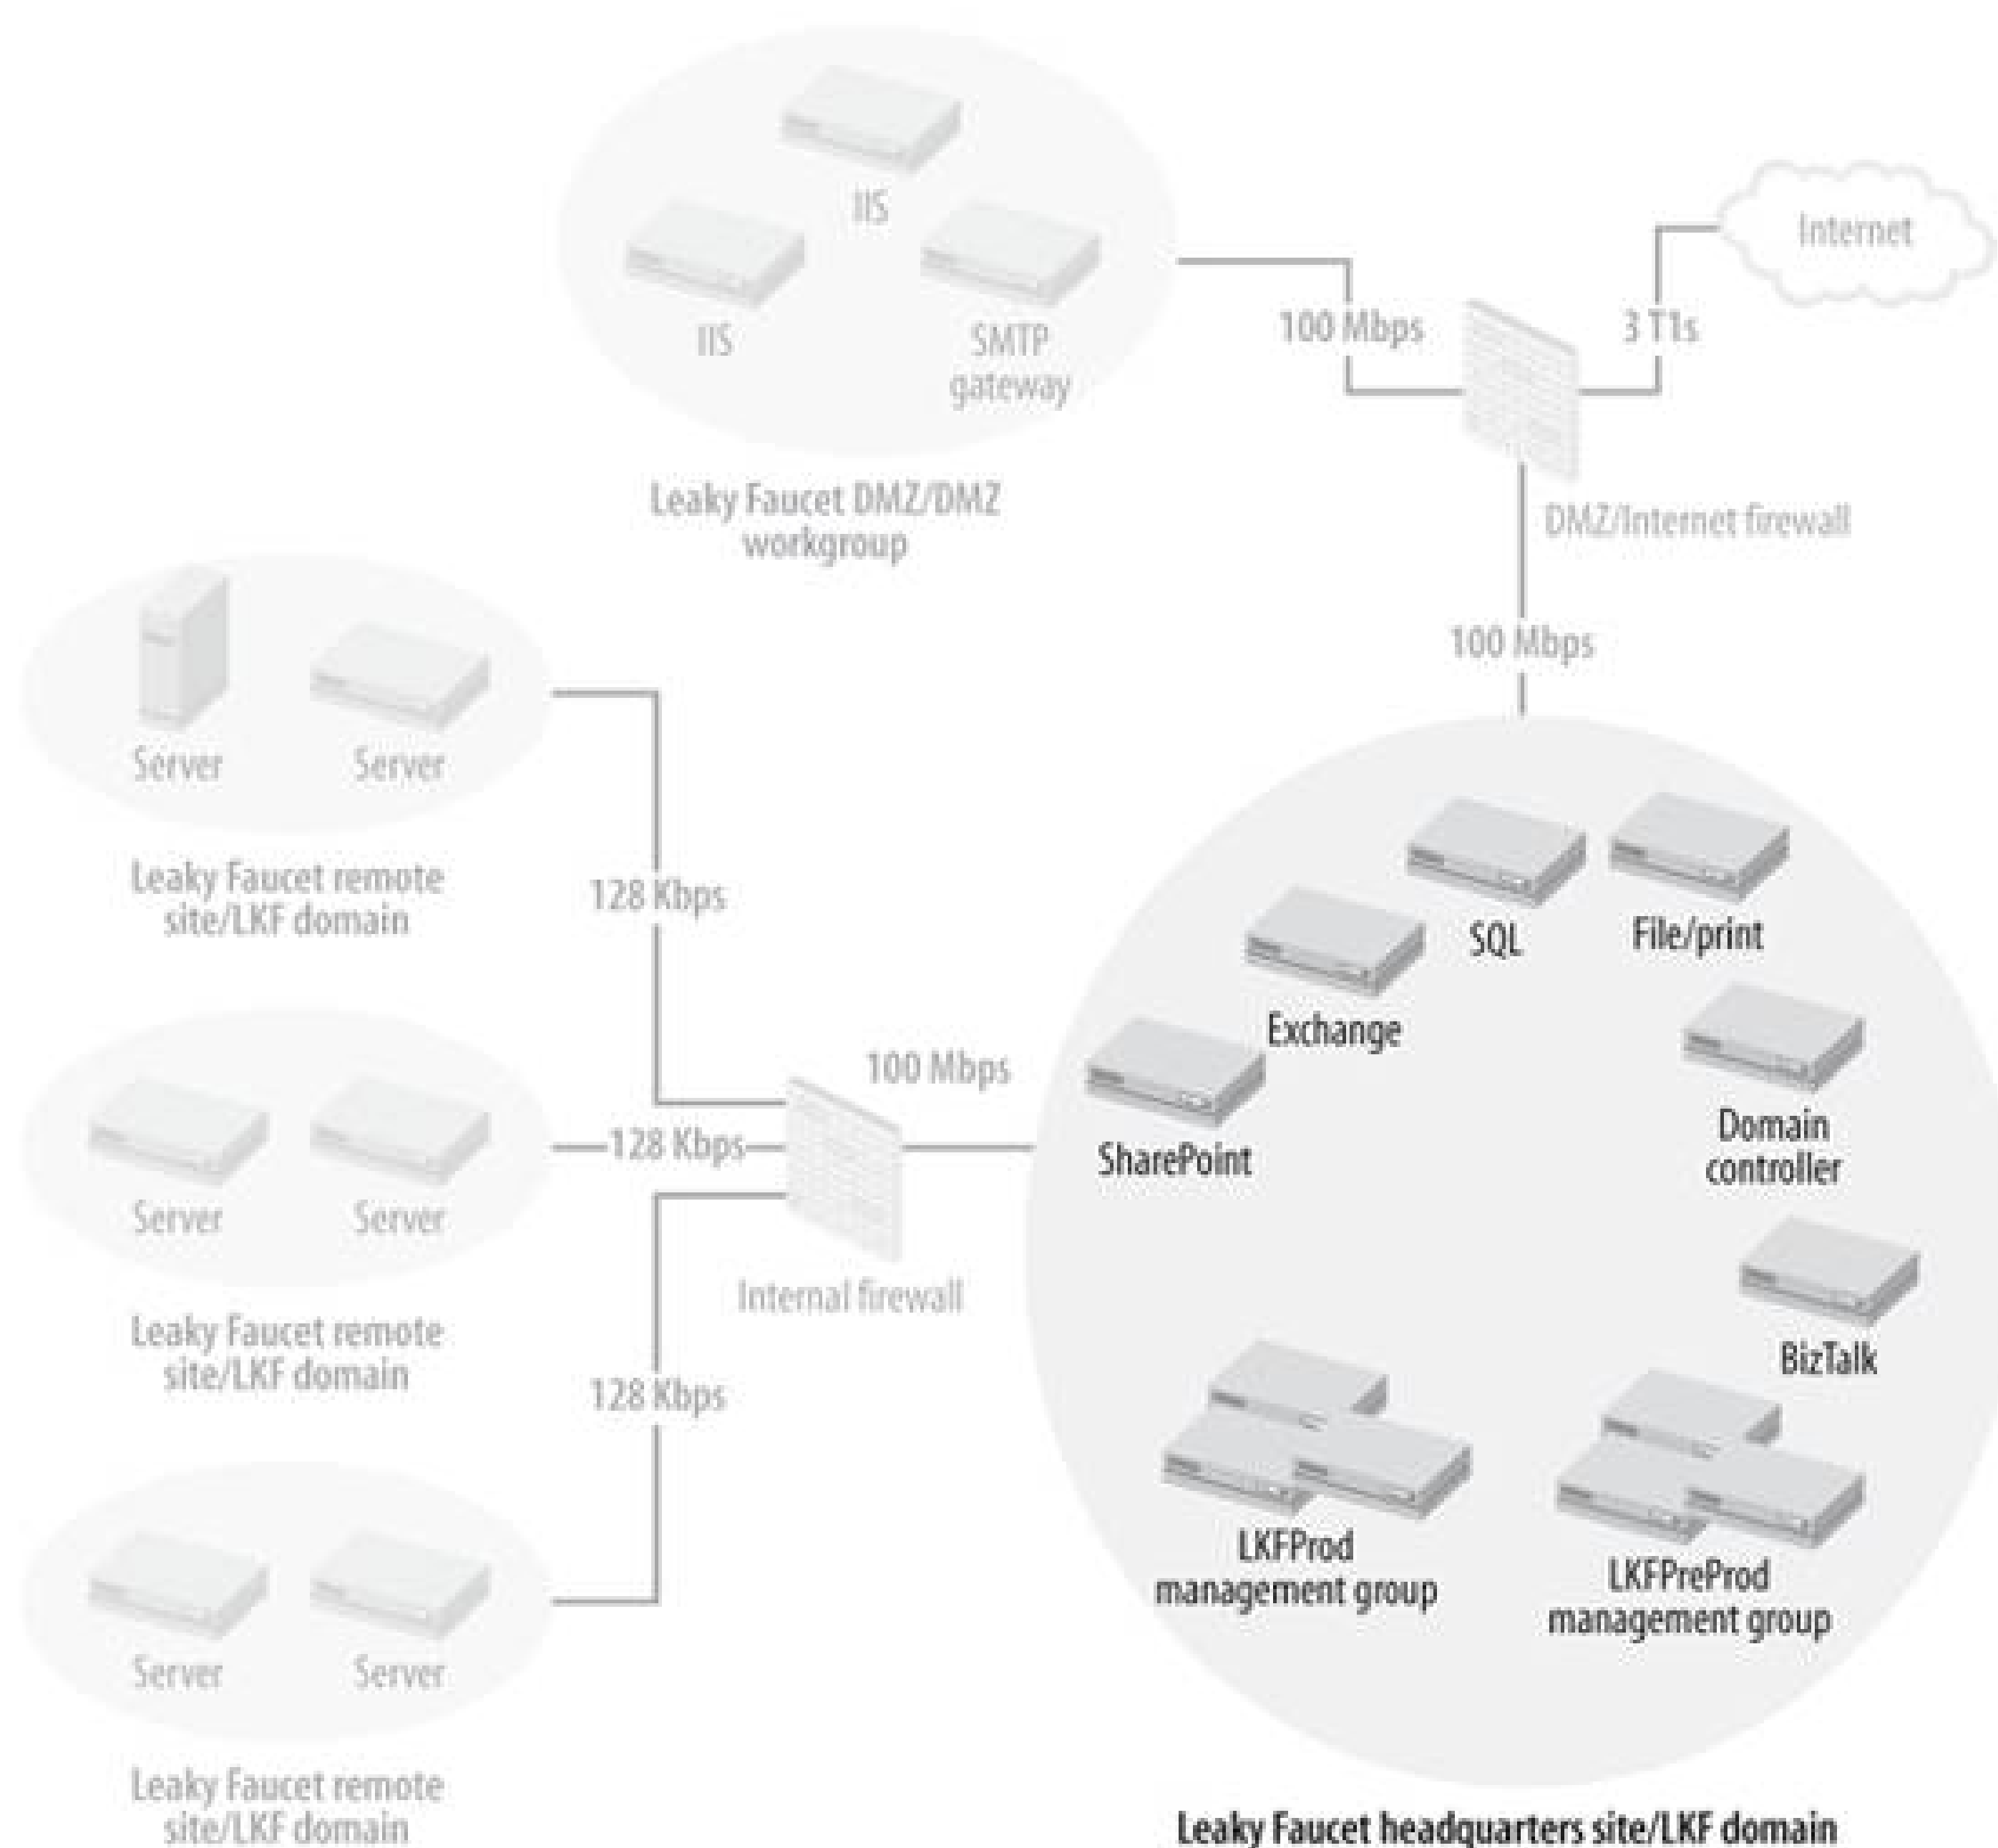
### 3.3.1. Preparing the Managed Computers

Preparing for agent deployment starts with preparing the computers that you want to deploy to. Leaky Faucet already knows it must deploy agents to all of its Windows-based servers, so the preparation for the deployment process will be the same for all machines. For each machine that is agent-managed, inventory is taken of the OS and patch level, the amount of free disk space on the system partition, and the size and configuration of the event logs.

MOM 2005 can install agents to computers that are running Windows 2000, Professional, or Server with SP4 or higher. Computers that are running an OS and patch level less than this can be monitored in MOM 2005 via agentless management. Because it is noninvasive, agentless management is attractive, but there are trade-offs in functionality.

Figure 3-3. Leaky Faucet headquarters site in the LKF domain






All of Leaky Faucet's servers are running Windows Server 2003 Standard edition. They are not at a consistent patch level across the board, but that is acceptable here. In fact, the Microsoft Baseline Security Analyzer (MBSA) management pack in MOM will tell them where the servers are deficient.

The MOM agent is relatively small, requiring only about 3.5 MB when installed with no additional management packs. That being said, there should be at least 100 MB of free space on the drive that is targeted for the agent install. This allows plenty of room for additional management packs and buffers. If the agent is reporting multiple management groups, then add 3 MB more additional space for each additional management group. By default, agents are installed to *C:\Program Files\Microsoft Operations Manager 2005*; this is configurable based on your needs and can be changed globally, at the management-server level, or when the agent is installed on a per-agent basis.

Max, the Leaky Faucet administrator, accepts the default installation location after he inventories the free space on all of the server system partitions via script.



Microsoft Systems Management Server 2003 is a server-based tool that provides hardware and software inventory, remote desktop control, software deployment, and software metering. In its hardware inventory capacity, it would be ideal for collecting the free disk space on all drives on all servers with SMS agents.

It is important to ensure that the Windows event logs are configured correctly. MOM 2005 gathers a great deal of information from the Windows event logs. If the server stops logging events because the event logs are full, MOM won't be able to report on the state of the managed computer accurately.

All logs should be configured to "Overwrite events as needed" with a maximum logfile size of at least 25 MB. If you leave the event logs configured to "Overwrite events older than X days" or "Do not overwrite events (clear log manually)," then you will likely miss events while you are manually clearing the logfiles. This is especially true of the security event log, which, on a domain controller, which can generate enough events to wrap on itself in less than a day. [Figure 3-4](#) shows a correctly configured Windows Server 2003 security event log for a domain controller in a small environment.

Figure 3-4. A Windows 2003 Security event log configured for MOM 2005 monitoring

One way to ensure the logfiles are configured correctly is by using AD group policy. AD group policy i

a mechanism that enforces configuration settings on computers. Group policies can be applied at the AD site, domain, organizational unit, and local computer levels. Group policies are applied in a specific order, so a policy that is higher in the order can be overridden by one that is lower or inheritance can be disallowed. To access these settings, open Active Directory Users and Computers and select the domain and open its properties. From there, click the Group Policy tab and edit the default domain policy. The settings are under Computer Configuration → Window Settings → Security Settings → Event Log - Settings for Event Logs. The settings are:

- Maximum application log size
- Maximum security log size
- Maximum system log size
- Retention method for application log
- Retention method for security log
- Retention method for system log
- Shut down the computer when the security audit log is full

You should also be aware of the effective Security Options Audit Group Policy setting that specifies whether the system should be shut down immediately if it is unable to log audit events to the security log. It may be appropriate to have the machine shut down automatically if it can no longer write to the security event log, depending on the security nature of its role. However, if the machine is not in a role that demands a high level of auditing, make sure this setting is disabled.

Configuration and application of domain-wide group policy is beyond the scope of this book, so for a detailed treatment of AD group policy, see *Learning Windows Server 2003* and *Active Directory*, both published by O'Reilly.

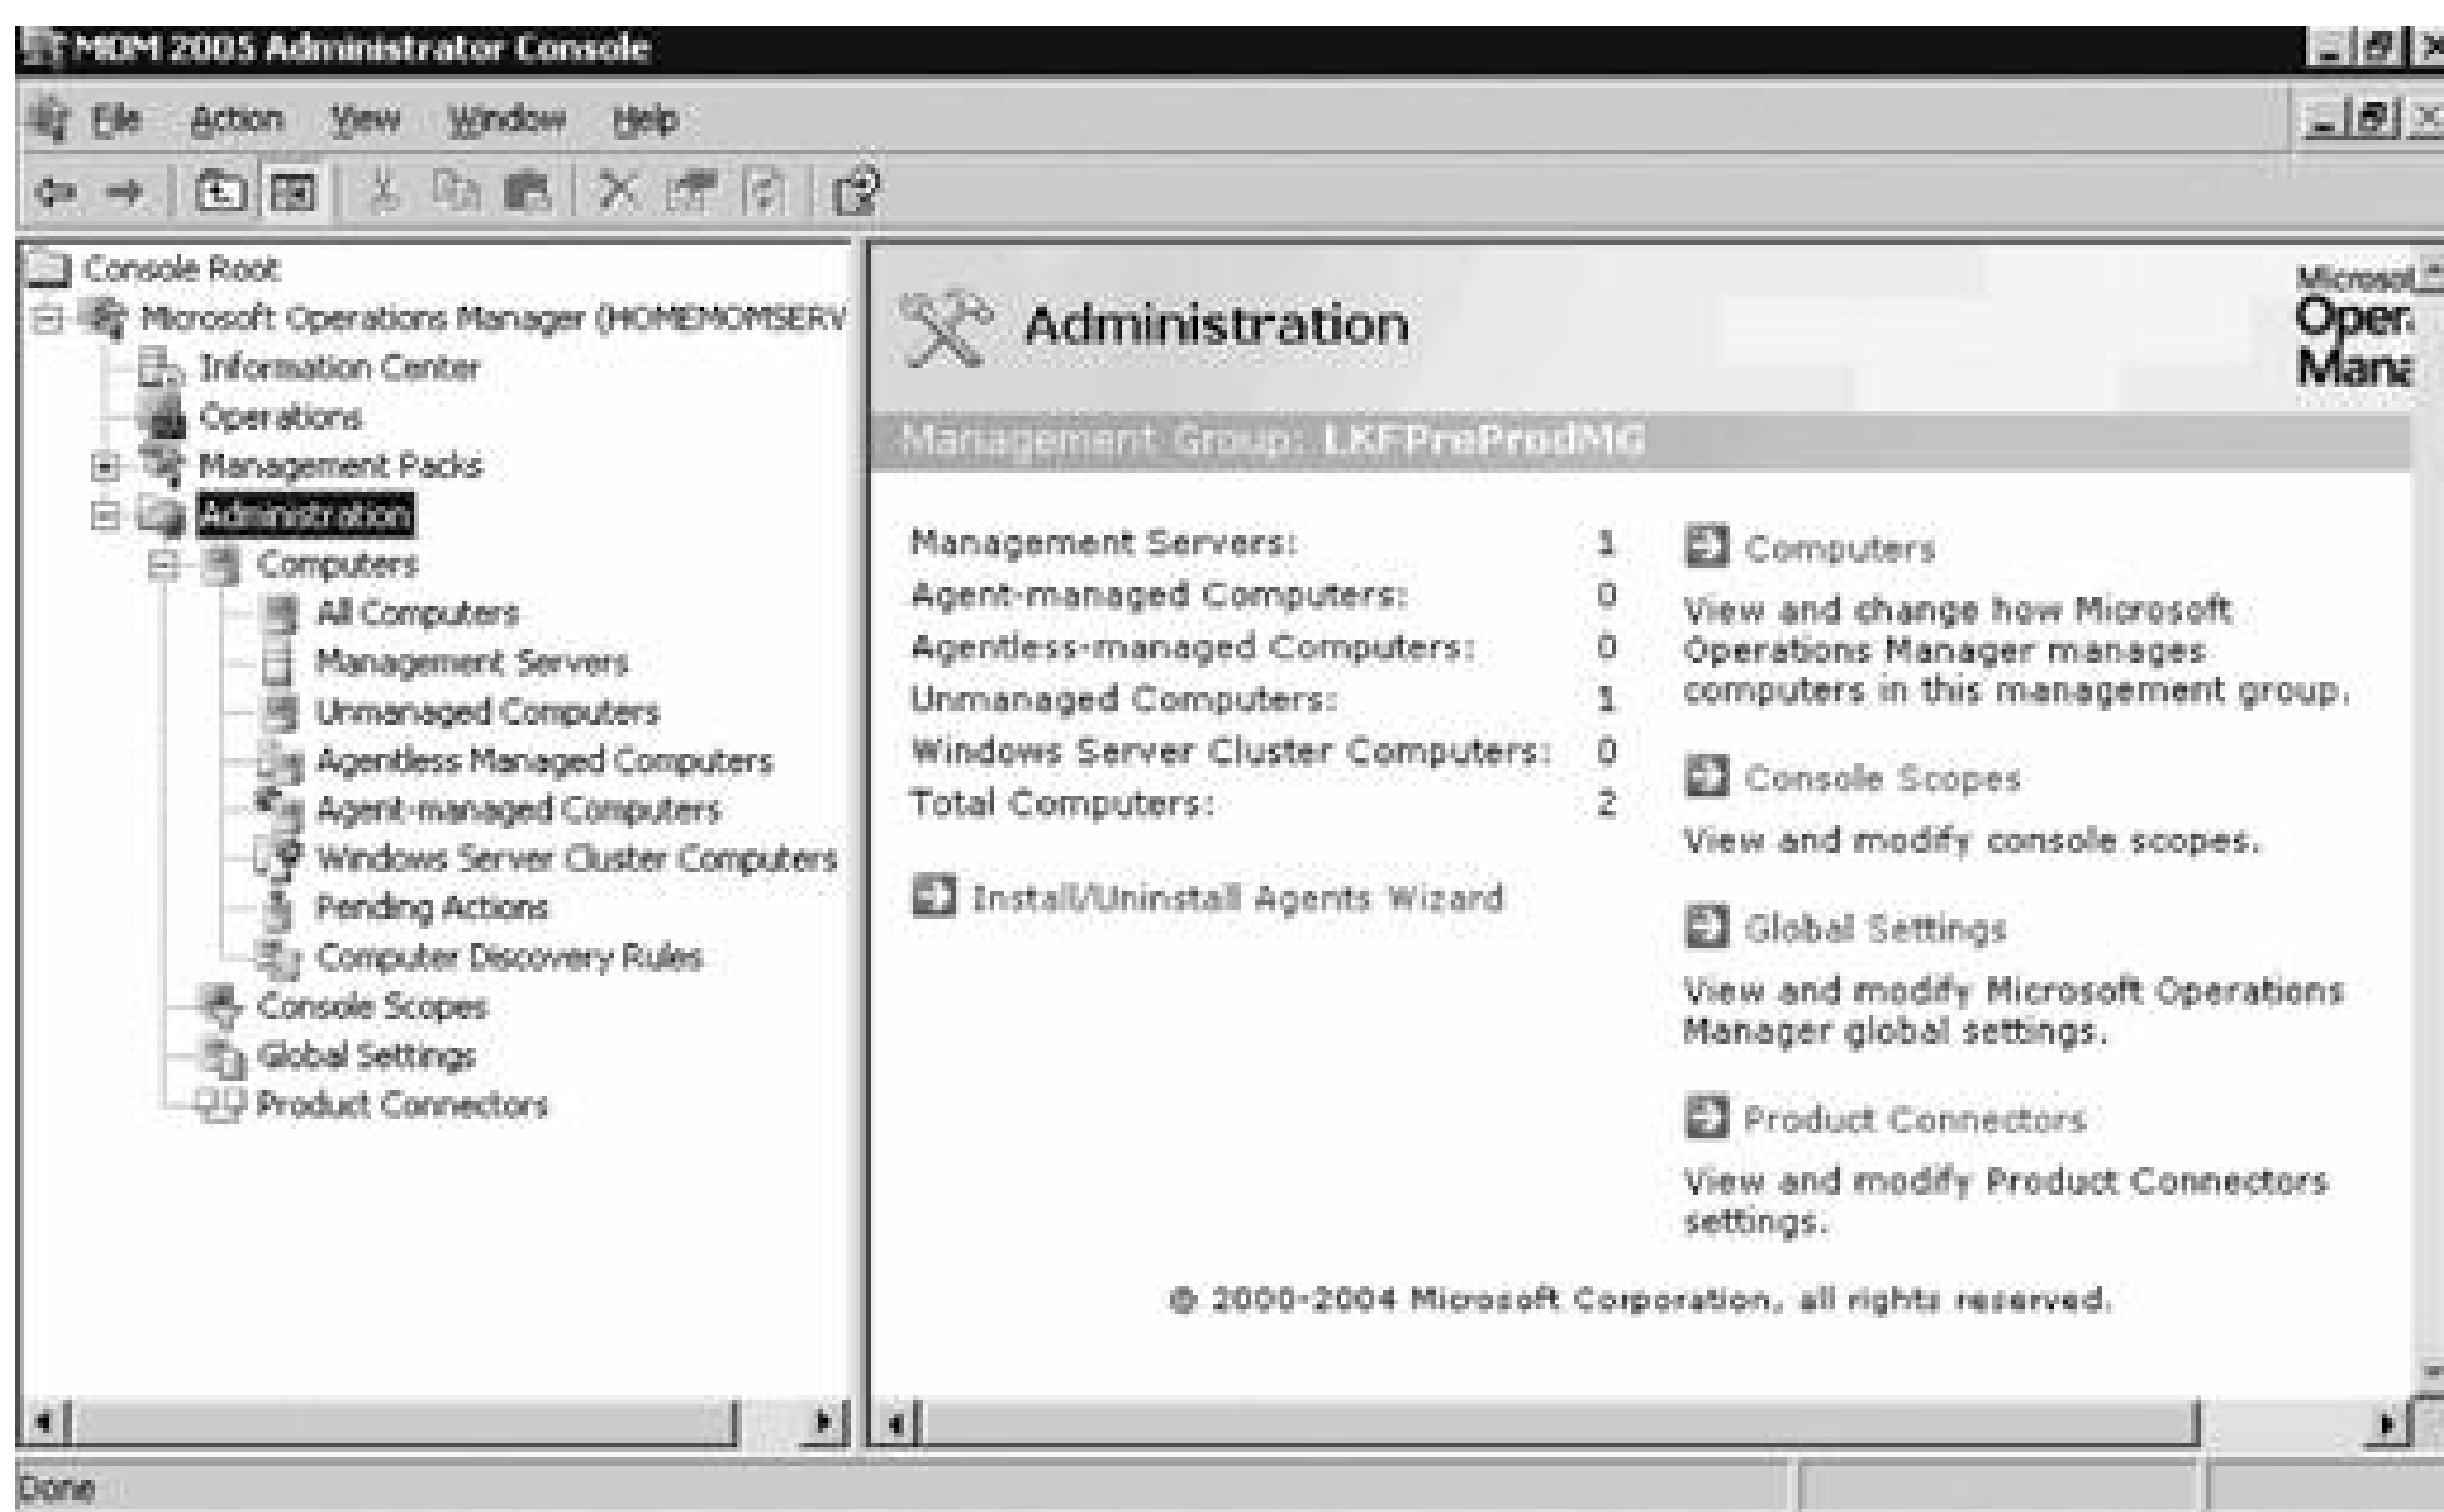
### 3.3.2. Preparing the Management Group

All management group configuration tasks will be performed in the Administration node of the Administrator console. This includes setting the default values for the agents, which was covered in the "[Agent Functions](#)" section earlier in this chapter, individual management servers, and the management group globally. To access this node, you must use an account that is a member of the MOM Administrators group. [Figure 3-5](#) shows the default view of the Administrator console with the Administration node selected in the details pane.

You can access this Microsoft management console (MMC) tool on the management server and on any computer where the MOM 2005 User Interfaces are installed. Click on Start → All Programs → Microsoft Operations Manager 2005 → Administrator console.

Figure 3-5. Administration node of the MOM 2005 Administrator console





Notice that the details pane shows an inventory and classification of the computers that MOM is aware of and a list of hyperlinks that will take you to the four nodes that are below the Administration node in the hierarchy. To prepare the management group for agent deployment, navigate to the Global Settings node. You can do this either by clicking on the Global Settings hyperlink in the results pane or clicking on the Global Settings node in the navigation pane. Settings configured in the Global Settings node become the defaults for the entire management group. Some defaults can be overridden lower down in the hierarchy and some cannot (see [Table 3-1](#)).

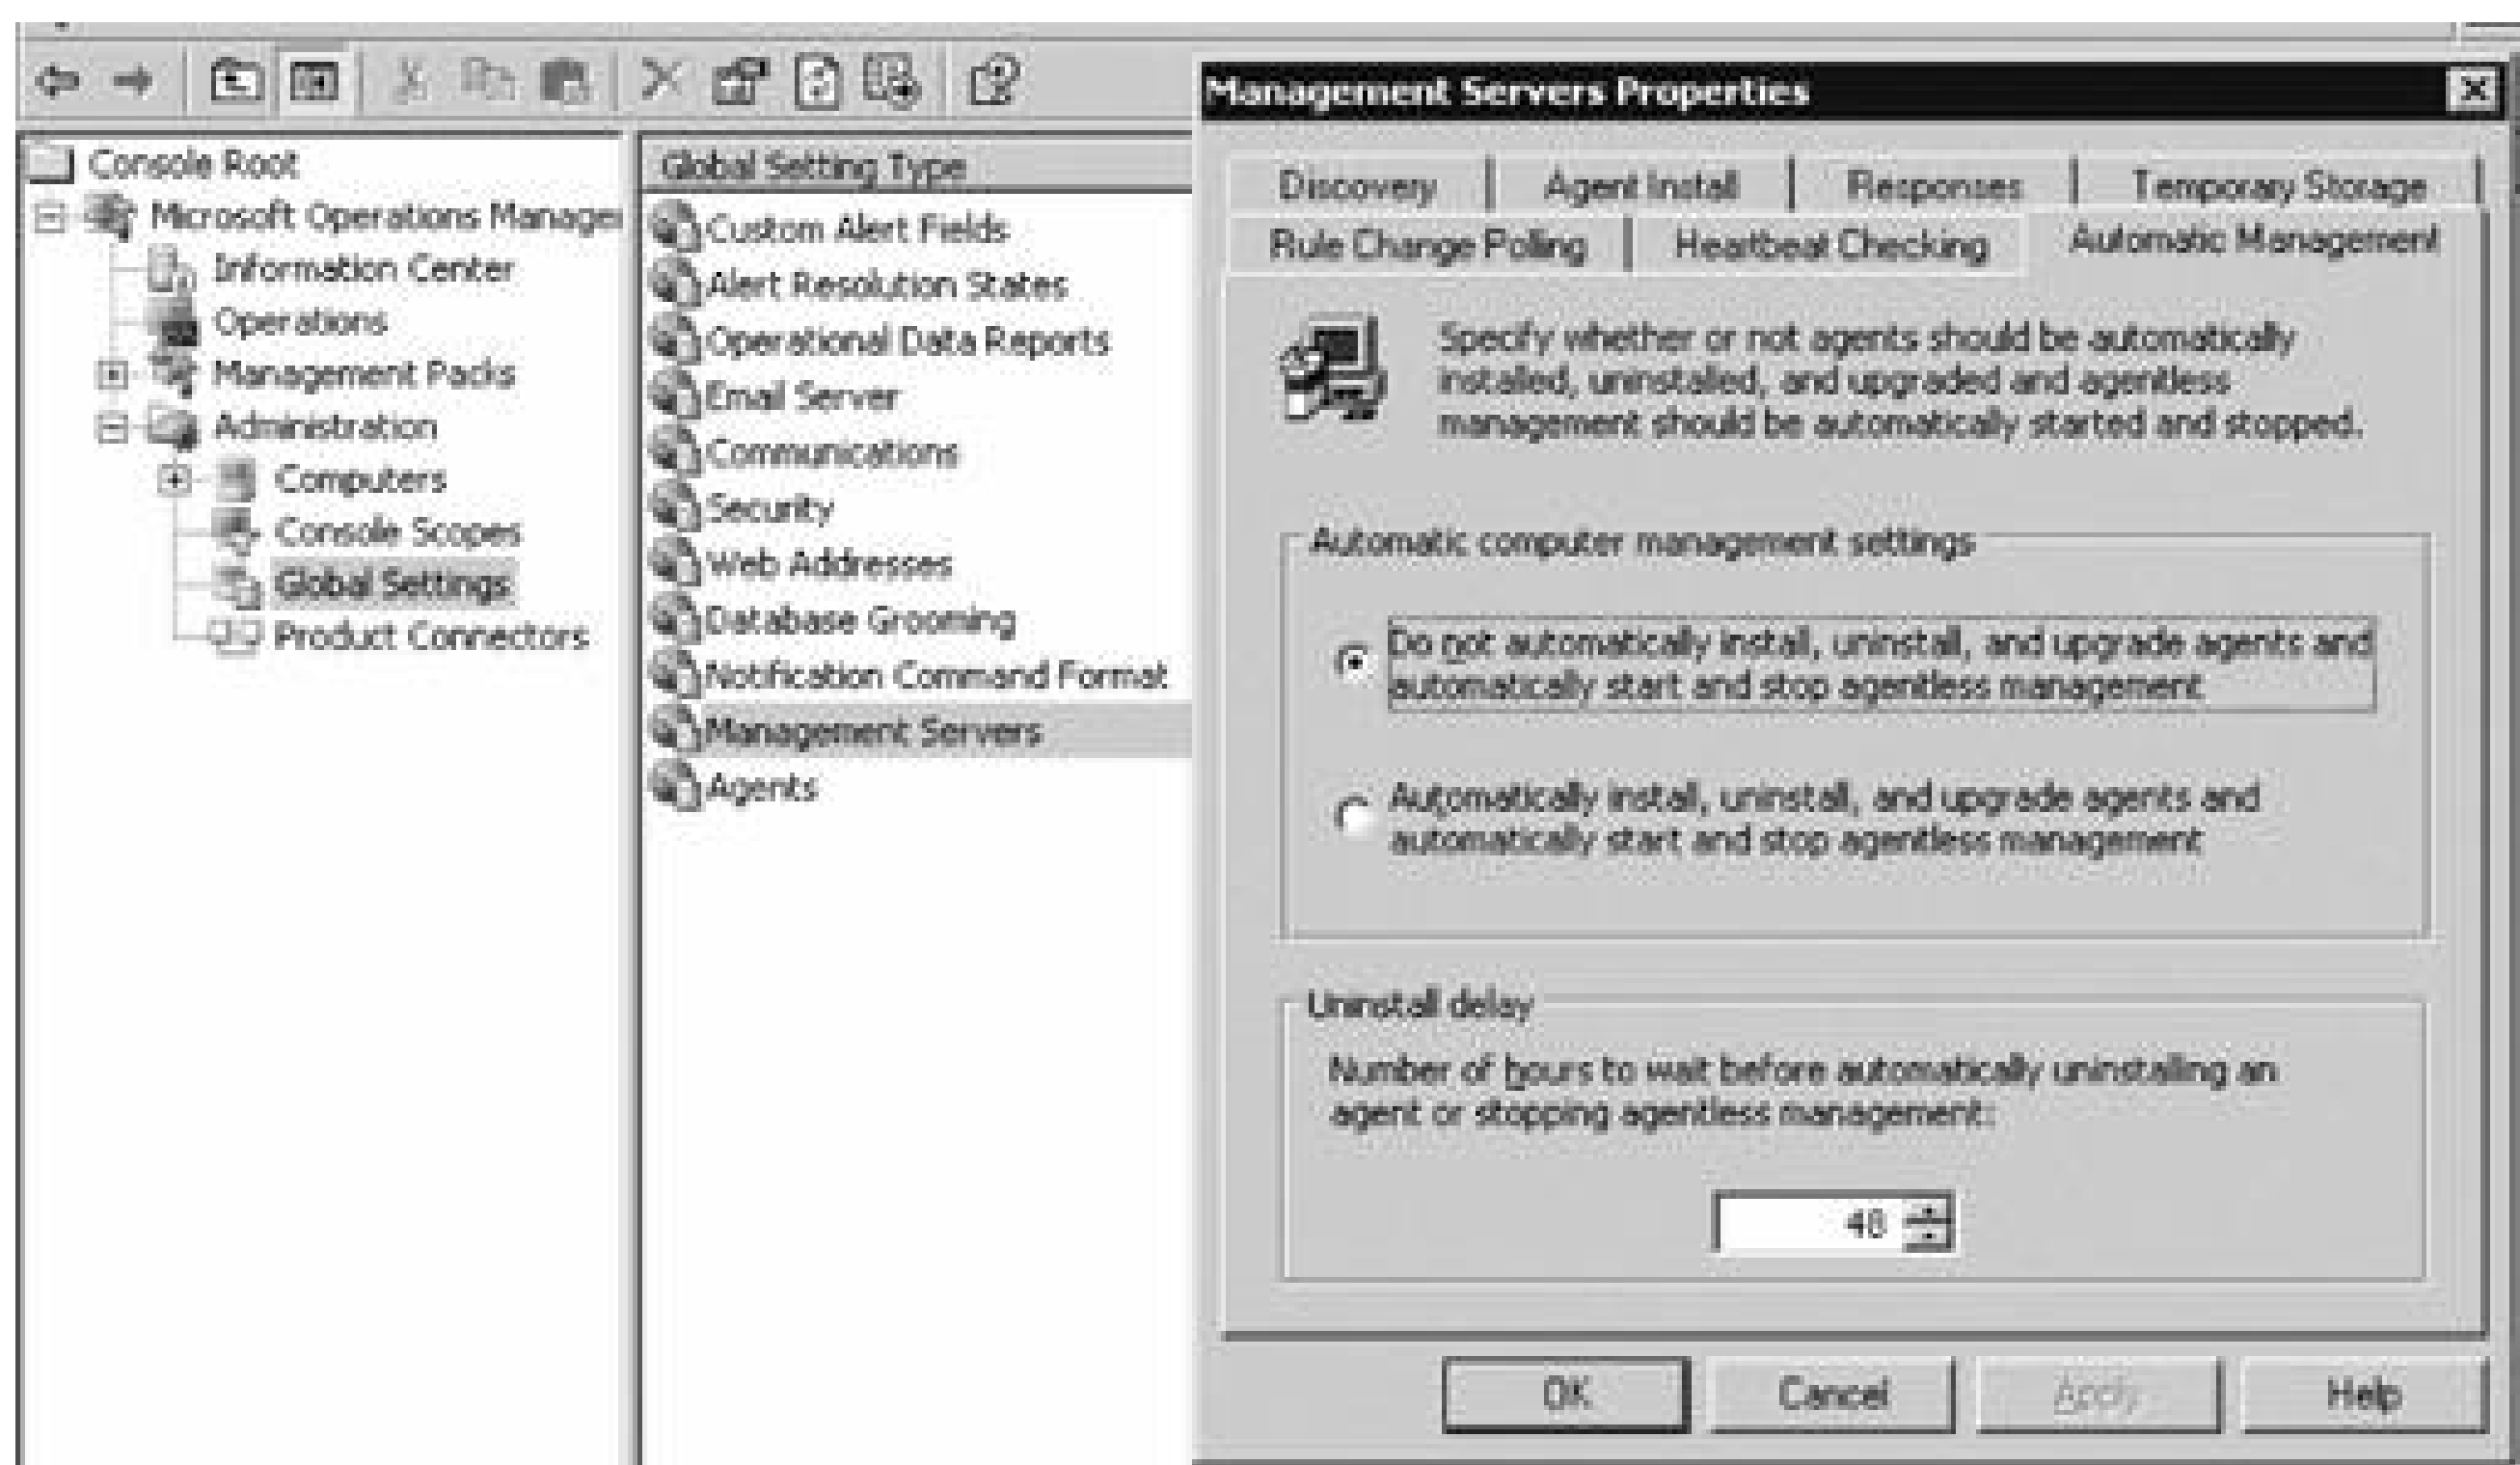
Familiarizing yourself to the Administrator console, and for that matter the Operator console, can take a bit of time, so be patient. There are multiple entry points into almost every function in both consoles. For the sake of clarity, this book will avoid using hyperlinks and instead show you how to navigate to that point manually. I have found that knowing where an object is located in the hierarchy helps in understanding its function and scope relative to the rest of the objects in the hierarchy. Use the hyperlinks once you know how to navigate the hierarchy and know the purpose of each node.

### 3.3.2.1. Configure automatic agent installation

MOM 2005 can automatically install agents to computers discovered through the computer discovery process, which runs once a day, and uninstall them as well. You may or may not want to this to happen. If you don't care what computers are added to MOM and then start generating alerts, set the discovery process to allow automatic installation, uninstallation, and upgrades. However, if you want more administrative control, accept the default configuration to disallow automatic installation.

[Figure 3-6](#) shows the default Automatic Management tab of the Management Servers Properties dialog box.

Figure 3-6. The Automatic Management tab controls what MOM 2005 does with computers after each computer discovery cycle is completed



If you allow automatic installation, MOM will attempt to execute whatever install, uninstall, or upgrade action specified on this tab using the management server action account. If the account does not have full administrative permissions on the target computer, these operations will fail. You would receive failure notifications in the Operator console.

The "Uninstall delay" setting is useful for allowing a period of time between the creation of a computer discovery rule that will cause the removal of an agent and the actual removal of the agent. This interval can be used to ensure that all critical functions are removed from the node before monitoring is stopped. It is not necessary to modify the default setting for "Uninstall delay" in preparation for agent deployment.

By leaving this setting at the default, which is "Do not automatically install" MOM 2005 will place the target computers into the Pending Actions node with an entry indicating what the desired action is after each computer discovery cycle. You can then approve the action and it will be carried out during the next computer discovery cycle, or you can execute the action immediately. This setting can be overridden in the Properties section of each management server.

Since Leaky Faucet's administrator, Max, does not want unknown computers using MOM to generate alerts and he doesn't want unexpected configuration changes, he elects to retain the default setting.

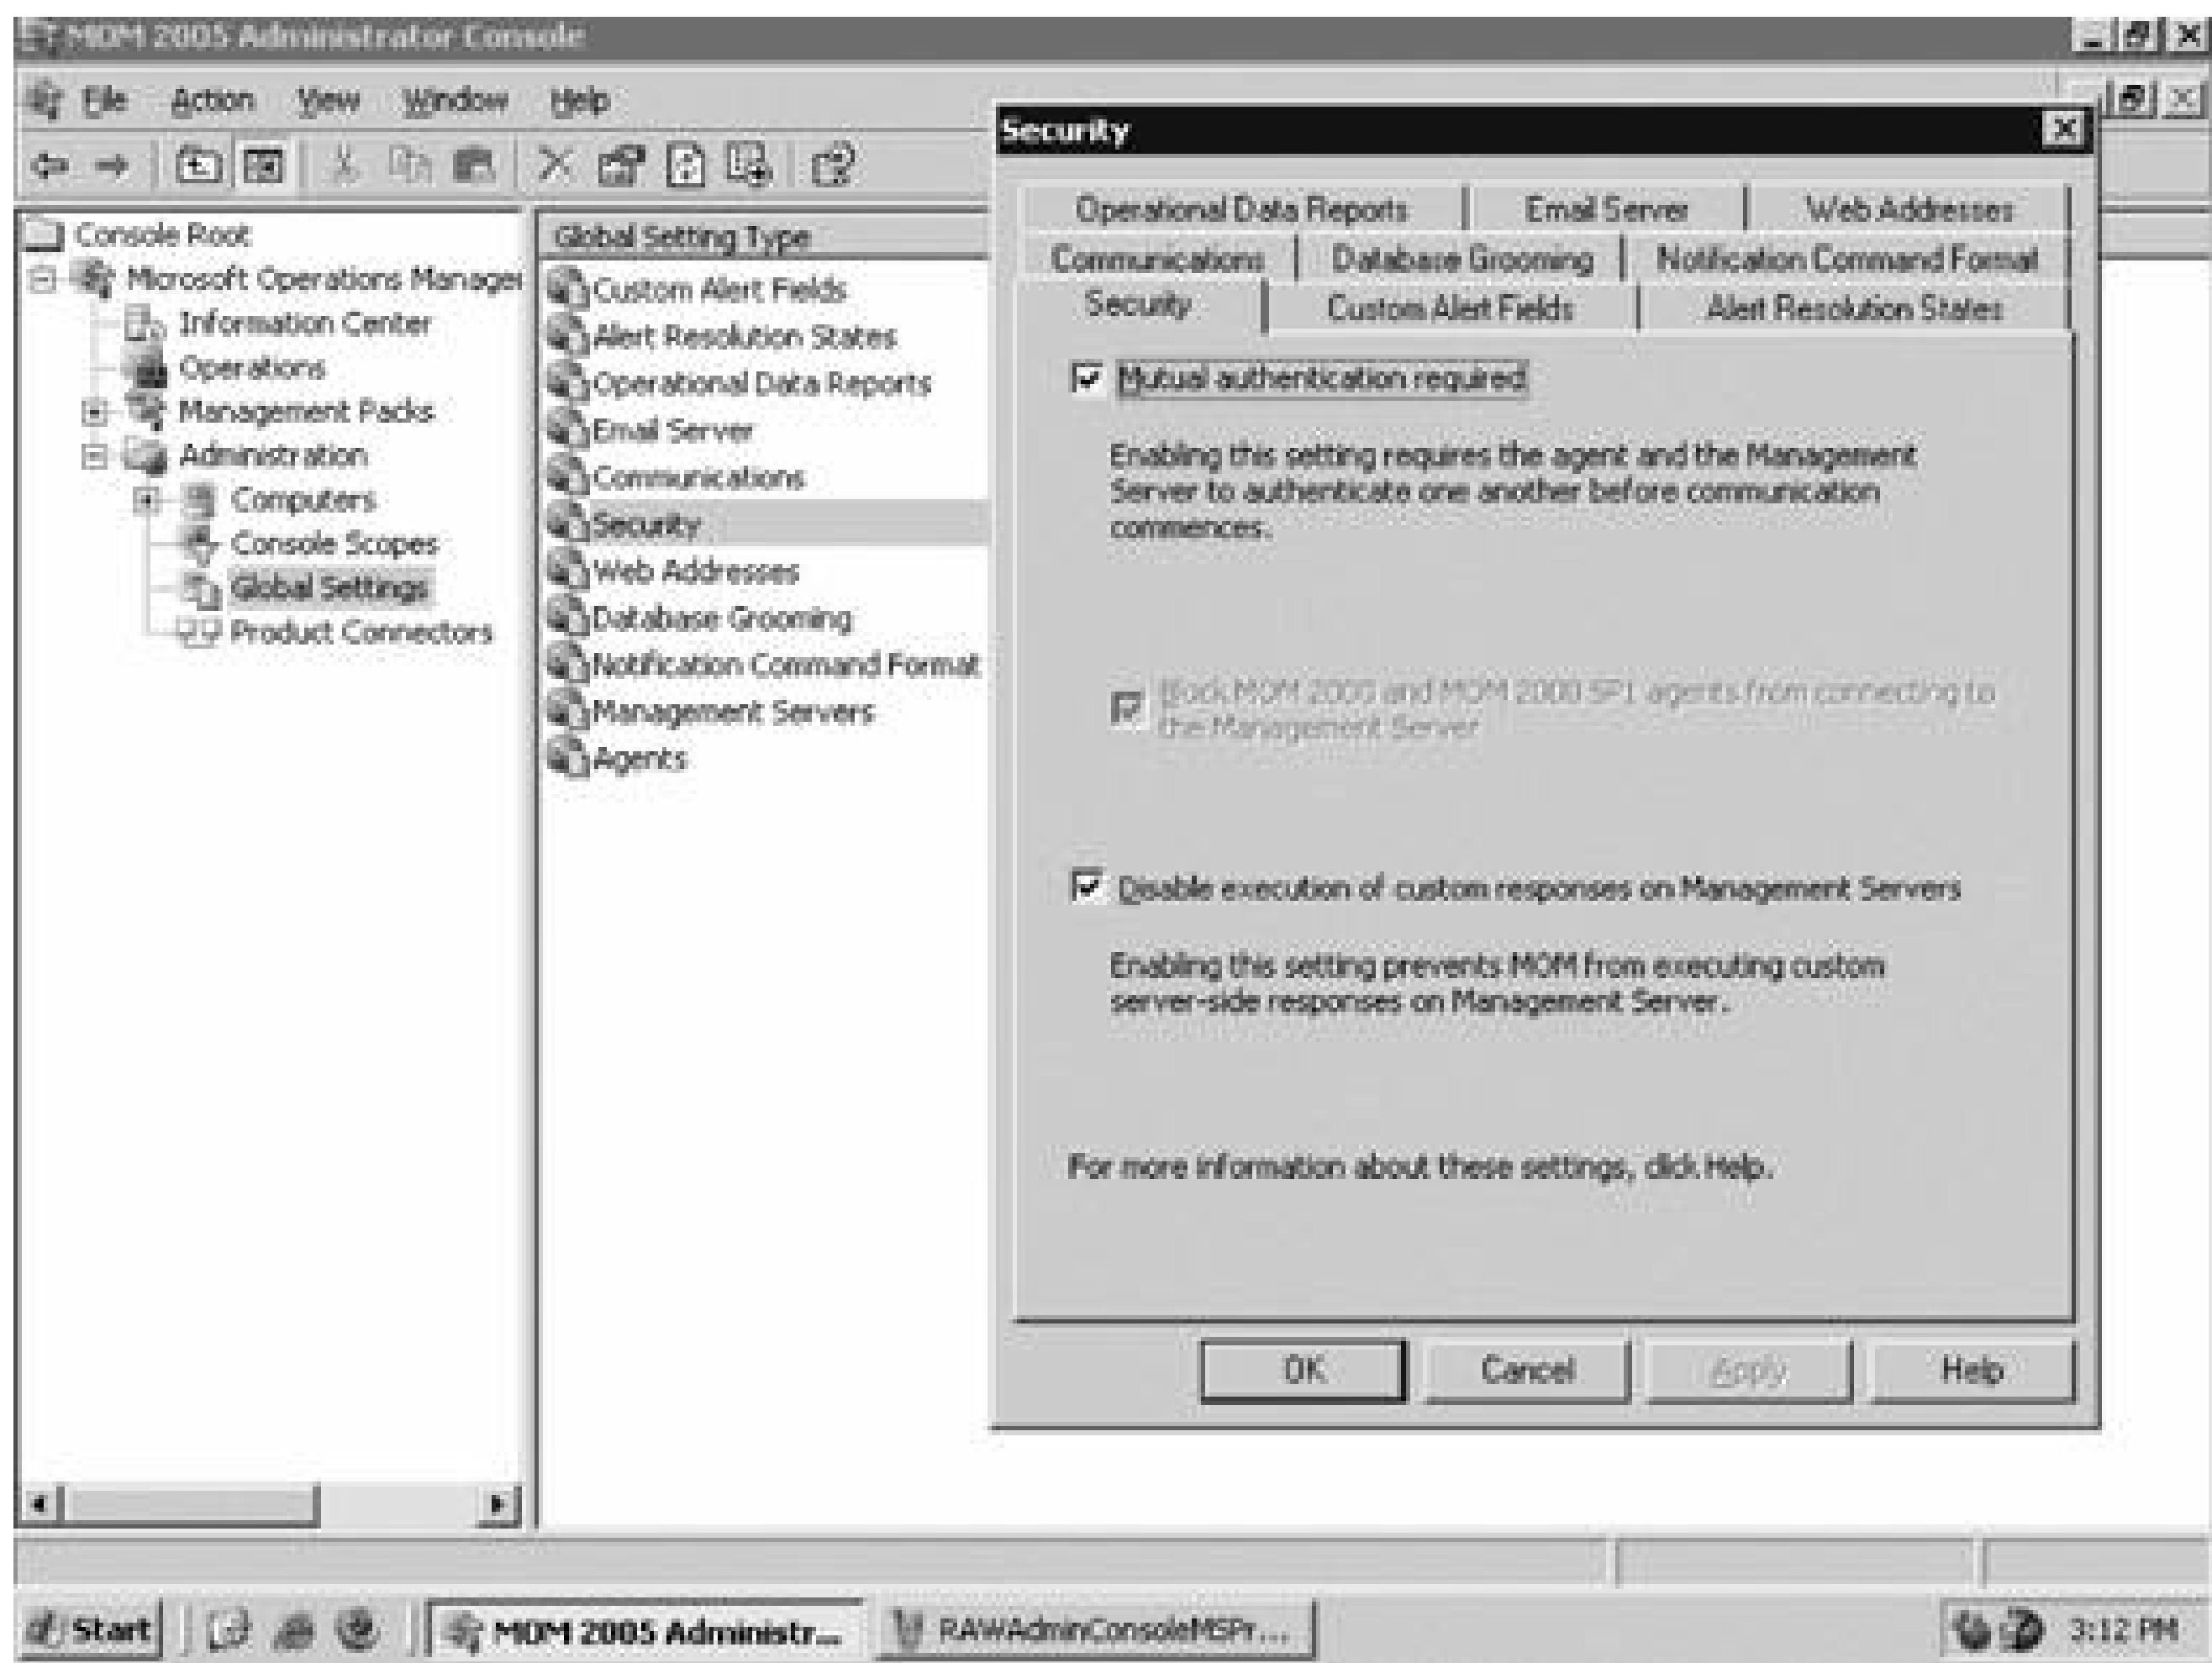
### 3.3.2.2. Configure mutual authentication

When you install MOM 2005 into an AD directory services environment, MOM can take advantage of the native Kerberos v5 authentication present in AD. Both clients (managed computers) and servers (management servers) can be required to authenticate each other prior to sending operational or configuration data. AD is required because there must be a valid, secure communications channel

between the two machines. Mutual authentication can be used between machines in the same domain and between machines in the same AD tree (AD domains that have a common namespace), but not across domain namespaces (between AD forests) or into untrusted spaces, such as workgroups. For example, if your management groups and agents are in AD domains that share a DNS namespace (are in the same AD tree), such as [leakyfaucet.com](http://leakyfaucet.com) and [headquarters.leakyfaucet.com](http://headquarters.leakyfaucet.com), you can use mutual authentication. If they are not in domains that share a namespace, such as [leakyfaucet.com](http://leakyfaucet.com) and [drippysink.org](http://drippysink.org), then you can't. For an in-depth treatment of Kerberos, see *Kerberos: The Definitive Guide*(O'Reilly), and for more details on Microsoft's Active Directory see *Learning Windows Server 2003*(O'Reilly) and *Active Directory* (O'Reilly).

Mutual authentication can only be configured globally, so there is only one setting for the entire management group. To configure mutual authentication, navigate to the Administrator console Administration node → Global Settings → Security, as shown in [Figure 3-7](#).

Figure 3-7. The Global Settings Security tab with mutual authentication enabled



Leaky Faucet has implemented a single AD domain, so mutual authentication was enabled during installation. They choose to accept this configuration because it enhances the overall security of their monitoring solution.



### 3.3.2.3. Configure rejection of manually installed agents

MOM 2005 agents can be installed on a discovered computer remotely from the management server, or locally at the discovered computer either by manually running the setup from the MOM 2005 installation media, which calls *momagent.msi*, or by copying the *momagent.msi* file to the local computer and running it there. The critical difference here is that the remote installation is initiated at the management server and is, therefore, intentional on the part of a MOM administrator; manual installation is not. Manual installation can be initiated by anyone with sufficient rights to the machine. When a manual installation is performed, you are prompted for the management group, the management server, and the port to use for communicating with the management group, among other things. The agent connects after it has been given this information. If MOM 2005 couldn't block this type of connection, the management group could be exposed to malicious actions. For manual installation to be enabled, mutual authentication must be enabled as well. Point 1 in [Figure 3-8](#) shows the default configuration at the global level that enables automatic rejection of manually installed agents.

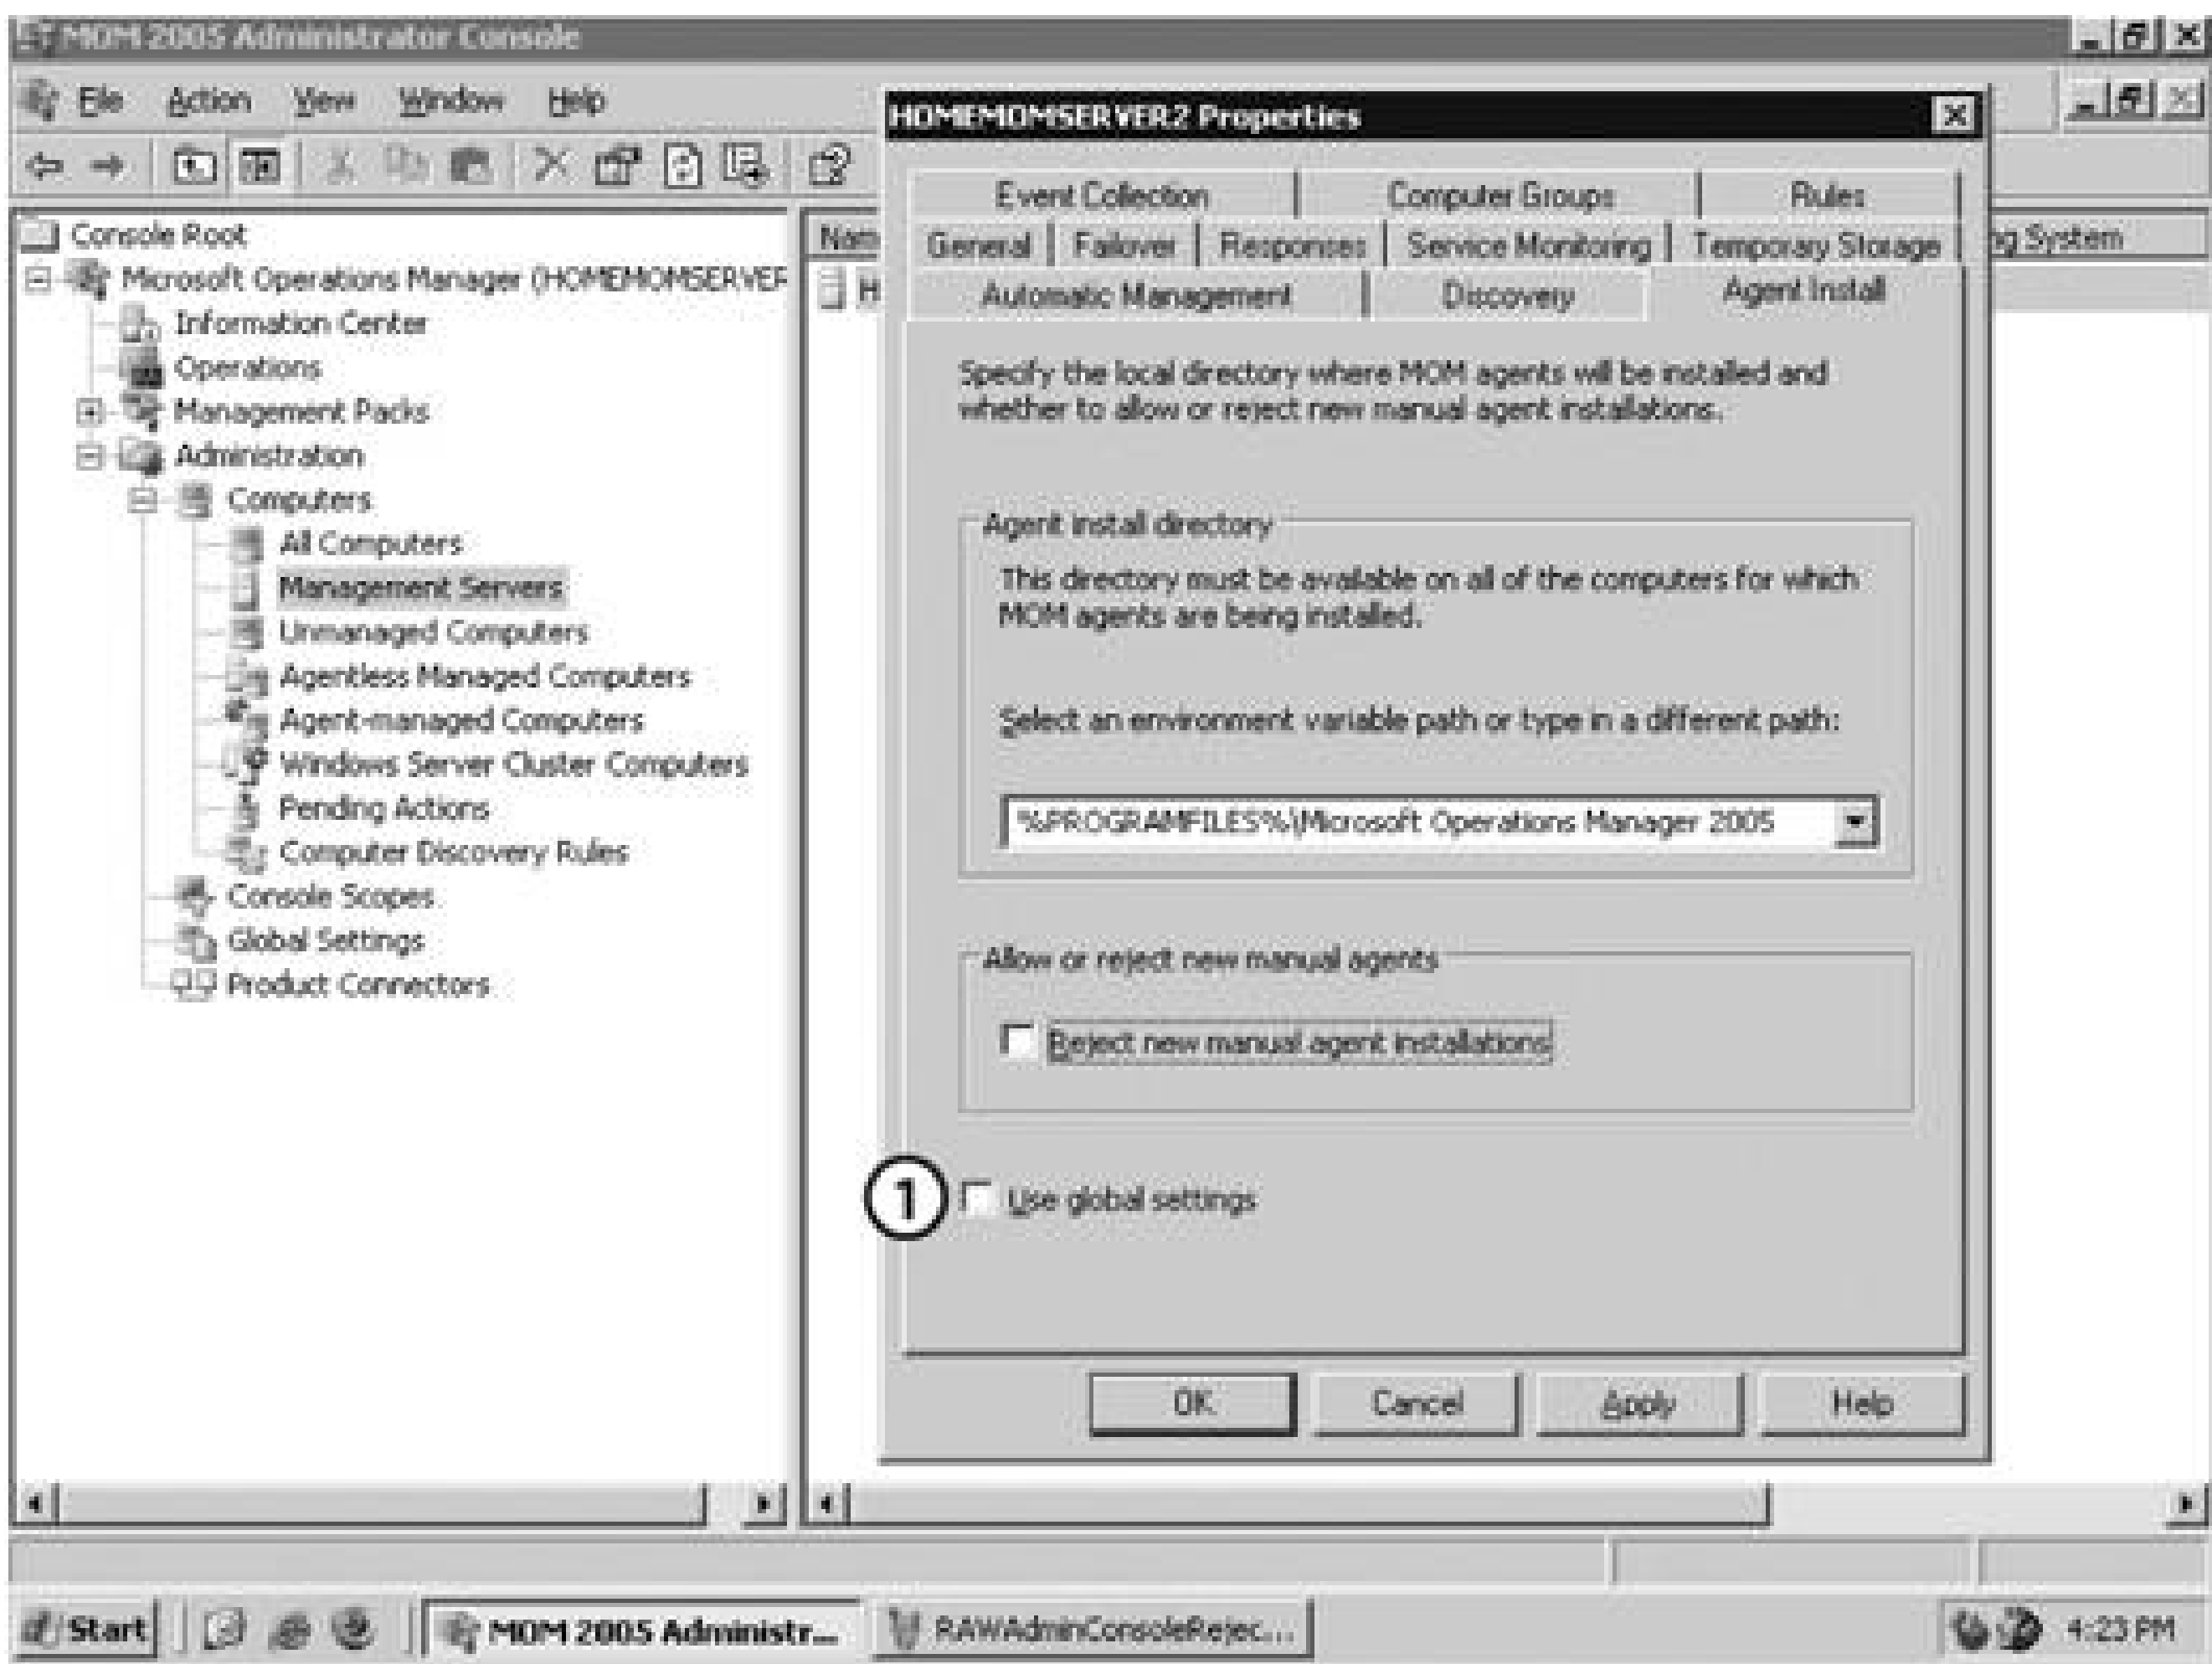
Figure 3-8. The default Global Settings rejects manually installed agents

This setting is configurable globally, but can be overridden at the management server level. To install agents manually, which is the only method of installing agents on computers that are separated from the management server by a firewall, you should override this setting on a single management server. When you change this setting, you must also right-click the Management Packs node, select Commit Configuration Change, and then restart the MOM service on all the management servers in the management group for it to take effect.

Opening a single management server to manually install agents, rather than all of the management servers in a management group, gives you a higher degree of control and reduces your exposure. You can configure failover to other management servers during manual agent installation without opening up the secondary management server during manual agent installation. After the manual agent installation is complete and the agents have been added to the management group, re-enable this setting until next time.

Because Leaky Faucet is first installing agents into a trusted network with no port restrictions between targeted computers and the management servers, Max will use the remote push-type installation of agents. Max overrides the default setting, see [Figure 3-9](#).

Figure 3-9. Override the globally set "Reject new manual agent installations" on a specific management server



### 3.3.3. Configuring and Running Computer Discovery

So, the servers that you want to manage are now prepared and you've configured the pertinent parameters in the management group to allow the deployment of agents in this environment. The remaining task to do before the agents are deployed and the machines are monitored is to make the management servers aware of the target machines. You do this by creating *computer discovery rules*, which are really just filters that look for matches to search specified criteria. When a match is found between the scanned computers and the search criteria, a mapping to the managed computer is



created on the management server. This mapping makes the management server responsible for agent administration and monitoring on that machine.

MOM 2005 has three ways to create the mapping between the scanned computers and the search criteria. Each method has its advantages and drawbacks:

- Create discovery rules in the Computer Discovery Rules node, and then either run an on-demand computer discovery cycle, or wait for the daily automated cycle to run. The default is every day at 2:05 a.m., but you can change this by going to the Administrator console, and in the Administration node choose Global Settings Management Servers Discovery tab. This cycle will use all of the computer discovery rules that have been configured. Agents are automatically installed or uninstalled based on the Automatic Agent Discovery setting. If you've prohibited automatic agent installation, the computer is placed in the Pending Actions container for further disposition.
- Use the install/uninstall agents wizard. This tool helps create a specific computer discovery rule, performs a limited discovery using only the rule created during the current session of the wizard, prompts you for necessary information, and then executes the action against the target machine. During the agent installation (or uninstallation) process, the wizard provides feedback on the status of the process. The install/uninstall process occurs regardless of the current automatic agent installation setting.
- Create a text file that contains the names of the servers that you want to manage. Call this file *ManualMC.txt* and place it in the `\%SystemRoot%\Program Files\Microsoft Operations Manager 2005` folder. If automatic agent installation is allowed, then when the daily discovery cycle runs these computers will be added to MOM and agents will be deployed.

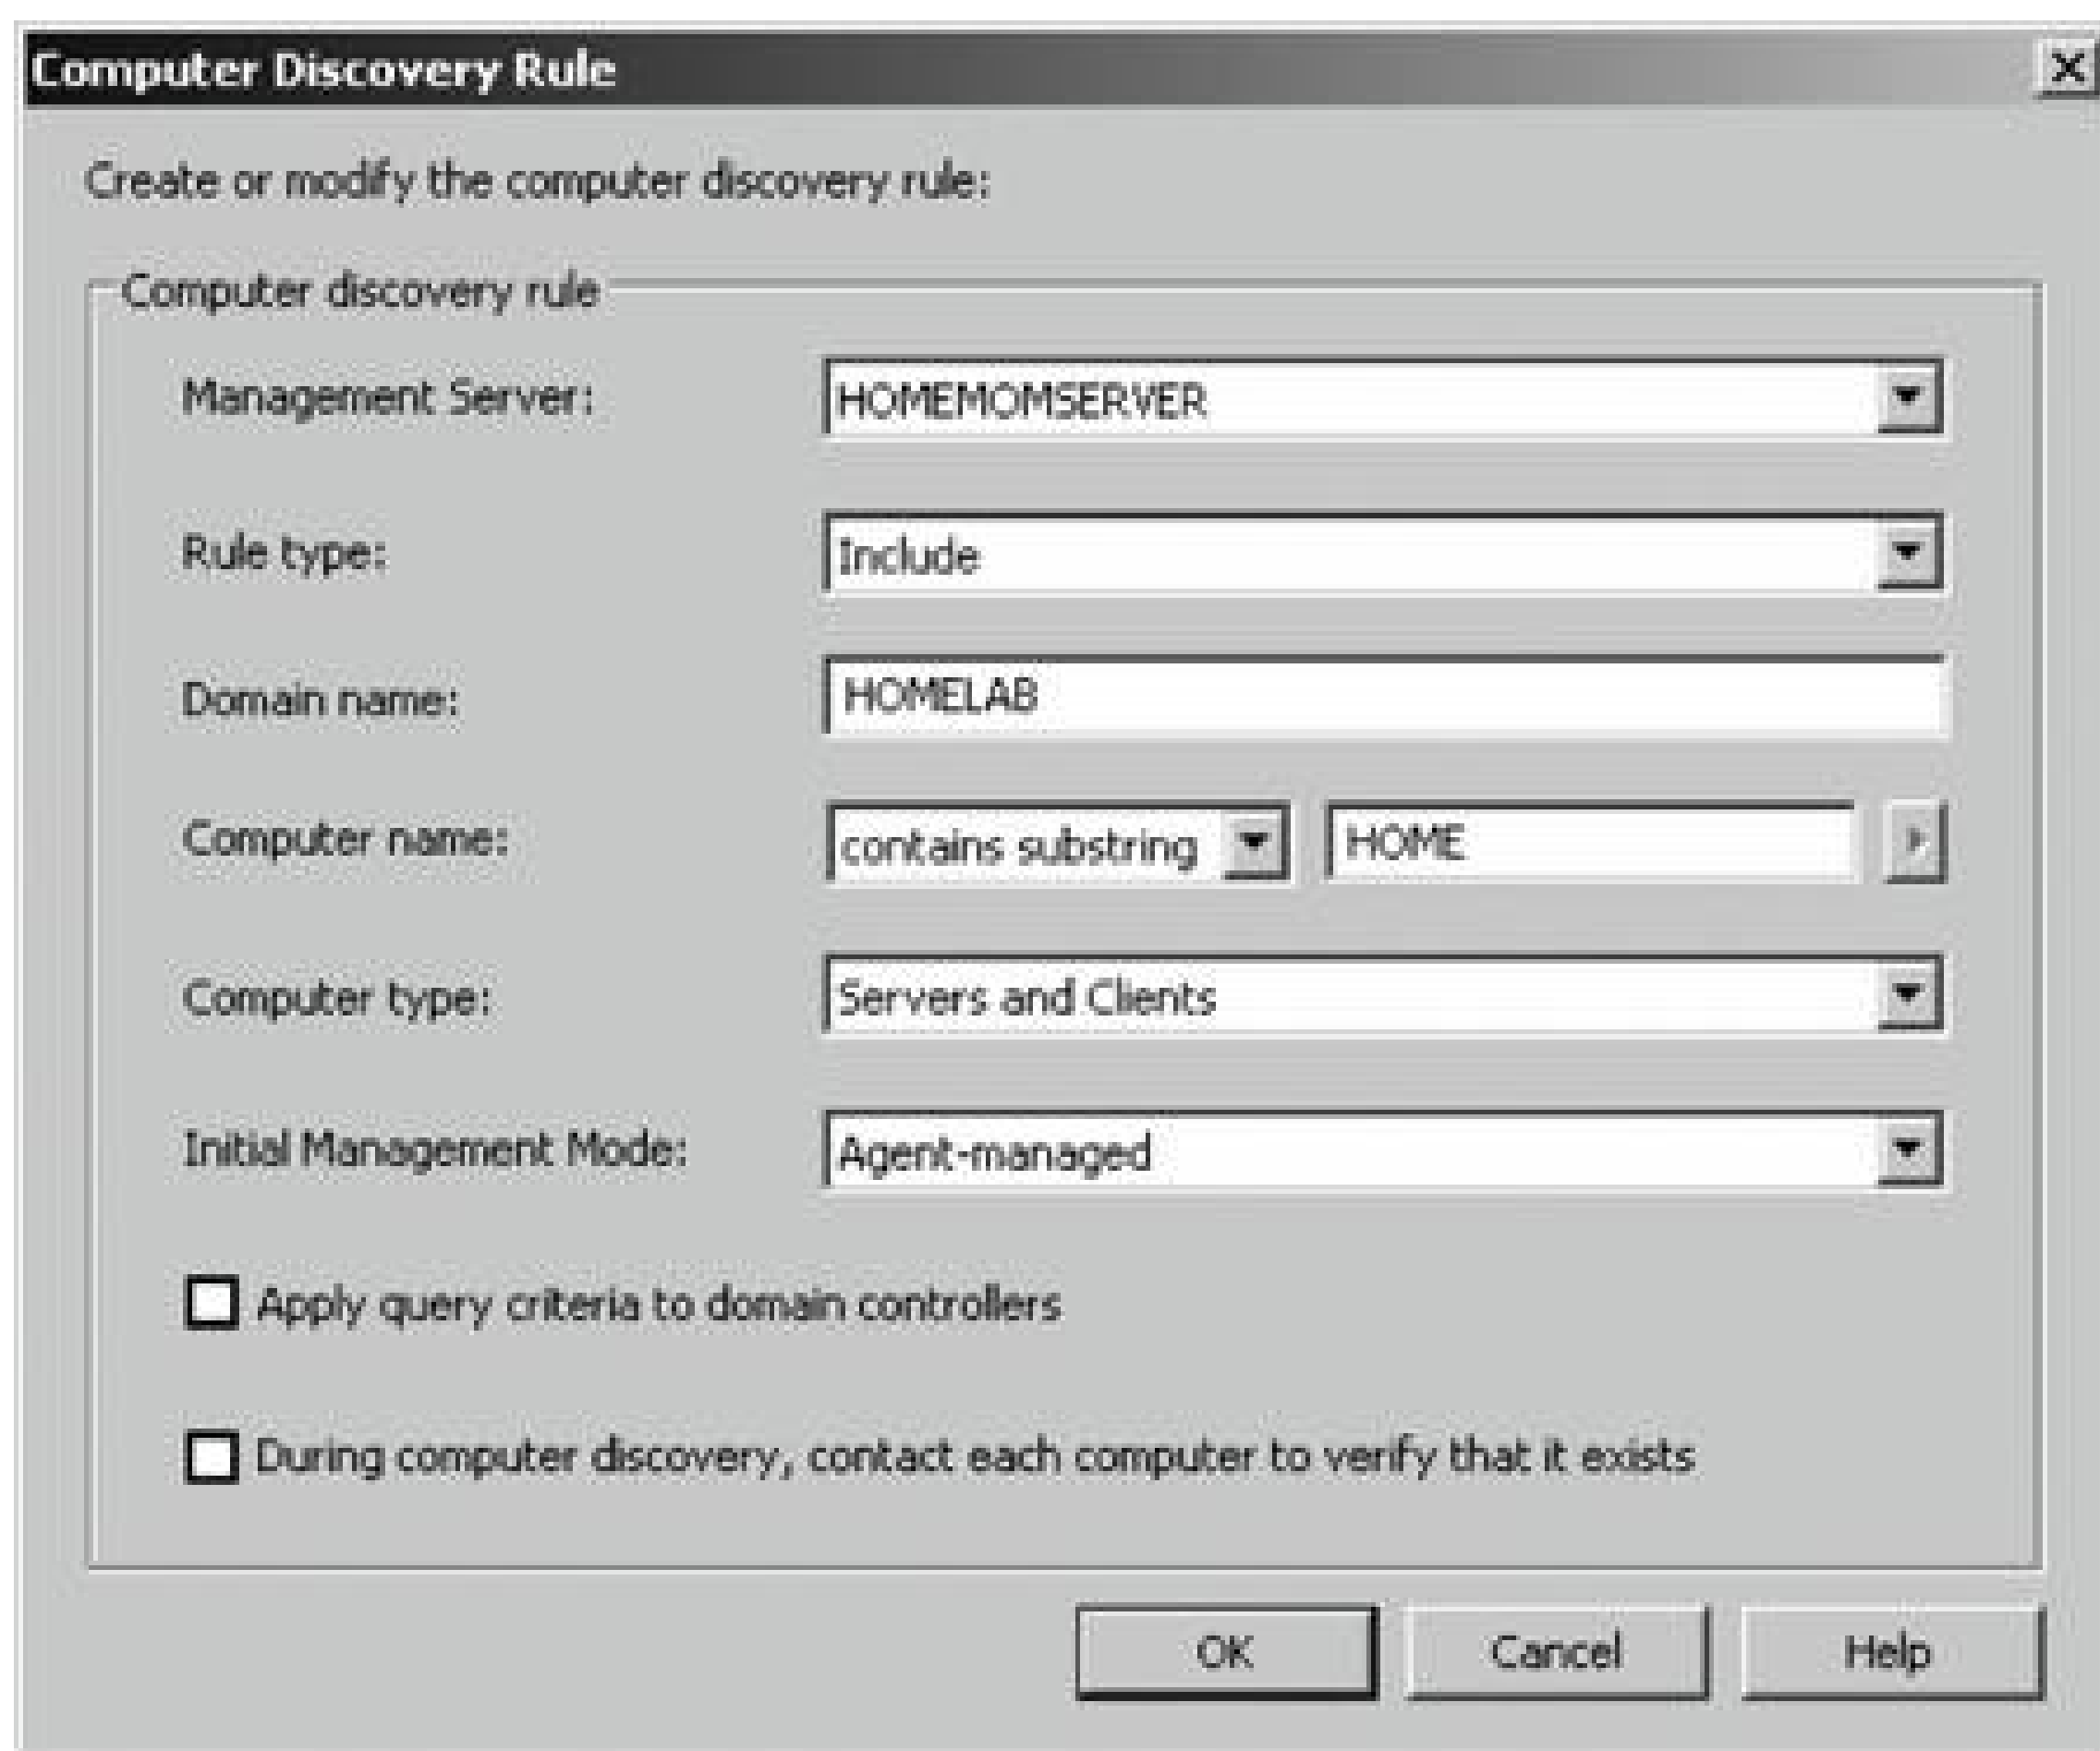
Do your initial discovery by creating your rules in the Computer Discovery Rules node. This method gives you the greatest flexibility and describes the configurable parameters. Skip the wizard to start with, because it does other things at the same time and there are certain restrictions associated with it. The *ManualMC.txt* method is a holdover from the previous version of MOM. So, you either have to maintain it manually on each management server or develop scripts to do so; either way is prone to human error.

### 3.3.4. Creating Rules in the Computer Discovery Rules Node

To start creating rules, navigate to the Computer Discovery Rules node in the Administrator console, then right-click and select Create Computer Discovery Rule. This brings up the Computer Discovery Rule options dialog box (see [Figure 3-10](#)). This box holds all the fields that can be configured for this filter. In this management group there are two management servers, *homemomserver* and *homemomserver2*, and a separate database server that holds the operations and reporting databases. By default, a fresh installation of MOM 2005 will only install agents on the management servers. So, if the operations database is not on a management server, you will need to configure a discovery rule to map it to a management server and install an agent.

Figure 3-10. Computer Discovery Rule options dialog box





The computer discovery rule options are as follows:

### *Management Server*

When you create a computer discovery rule, it is linked to a specific management server. When the discovery cycle runs, computers that match this filter become "owned" by that management server for the purposes of agent administration. If there are multiple management servers in a management group, this drop-down list will allow you to select which management server you want to bind this rule to. *homemomserver* has been selected for this example.

### *Rule type*

Use the "Rule type" field to specify whether the rule will be used to include computers for management by the management server or exclude them from being managed. Exclude rules always override include rules. Since the only two computers that are currently being managed by this management group are the management servers themselves, an Include rule for *homesqlserver* that houses the MOM operations database needs to be created.

### *Domain name*

In the "Domain name" field, enter the NetBIOS name of the domain that you want searched, the Fully Qualified Domain Name (FQDN), or no value at all. The more the search scope is

narrowed by providing distinguishing criteria, the quicker the search will run. If the rule searches for computers that are not in the same domain as the management server you are configuring the rule on, use the FQDN, because it gives the rule a more exact focus. *homesqlserver* is located on homelab, so that is entered here.

### *Computer name*

The "Computer name" field is where you specify most of the distinguishing criteria. The first drop-down menu allows you to select from the following options:

Equals. Use this to match an exact string. When selected, wildcard characters are not available in the next field.

- a. Contains substring. This is useful if your server naming convention includes characters that specify a server role and there are multiple servers in that role, such as EXCH for Exchange servers. In [Figure 3-9](#), the string HOME is specified because all servers in the homelab domain have that substring in them, but workstations do not.
- b. Matches wild card. This allows you to insert ? (any character), \* (any number of characters), or # (any digit) values into the target text in the next field.
- c. Matches regular expression and matches Boolean regular expression. To build more complex filters, select these options from the "Computer name" field. This will make the matching types of operators available for insertion into the search string.

This rule can discover all of the servers in the homelab domain. By specifying contains substring HOME, *homesqlserver* and all of the other servers in the domain will return as a result. You can then choose to install agents or start agentless management without having to configure another rule.

### *Computer type*

Unless you have very specific needs, such as a management server that will be dedicated to managing either servers or clients, leave this at the default setting of Servers and Clients. If you change it to one or the other, the discovery process must query each computer that has matched the filter for its role in the domain. It is left at the default setting here.

### *Initial Management Mode*

The default Initial Management Mode is Agent-managed, but you can also select Unmanaged or Agentless-managed. This field doesn't play a role in the filter, but labels the discovered computers with the type of management you want to be performed.

- The Agent-managed option is self-explanatory.
- The Agentless-managed mode provides some of the functionality as the full agent, but it cannot monitor across firewalls (it uses the RPC port range), and not all data providers are supported (specifically, application event logs). Windows event log descriptions do not



come through in alerts and you must carefully test each management pack to see how well the agentless management mode supports it. In addition, you can only monitor up to 10 computers per management server and 60 computers per management group in agentless management mode. Because this type of management is performed by the management server, the management server action account must have full administrative rights to the target machine.

- Unmanaged is a management mode that indicates that MOM is aware of the computer but is not currently monitoring it. Before you can delete a computer from MOM, the agent must be removed or agentless management must be stopped. This places the computer into the unmanaged category.

Finally, the two checkboxes, "Apply query criteria to domain controllers" and "During computer discovery, contact each computer to verify that it exists," change how the query runs. By default, queries are not applied to domain controllers, nor do they verify the existence of a server. So, if you have stale computer entries in AD, the "Verify existence" setting can help weed them out at the cost of increasing the overhead and time required to complete the search.

If a computer is accidentally included for management on different management servers in the same management group, it will switch primary management servers every time the discovery cycle is run. This is something you want to avoid, so be careful with your filter construction. If a computer is included for management on management servers that are in different management groups, then it becomes a multi-homed computer and will accept configurations from and send data to both management groups.

### 3.3.4.1. Running computer discovery

Now that the discovery rule has been configured, you can either initiate a full computer discovery cycle or wait overnight for MOM to kick discovery off automatically. To initiate discovery immediately right-click the Computer Discovery Rules node or the Management Servers node and select Run Computer Discovery Now. A pop-up box will appear, confirming that a computer discovery task has been submitted to the management servers. The management server then queries the domain that was specified in the rule and applies the search criteria. If you have configured the rule to verify existence or to determine the computer role of the matches it finds, it will perform those actions now as well.

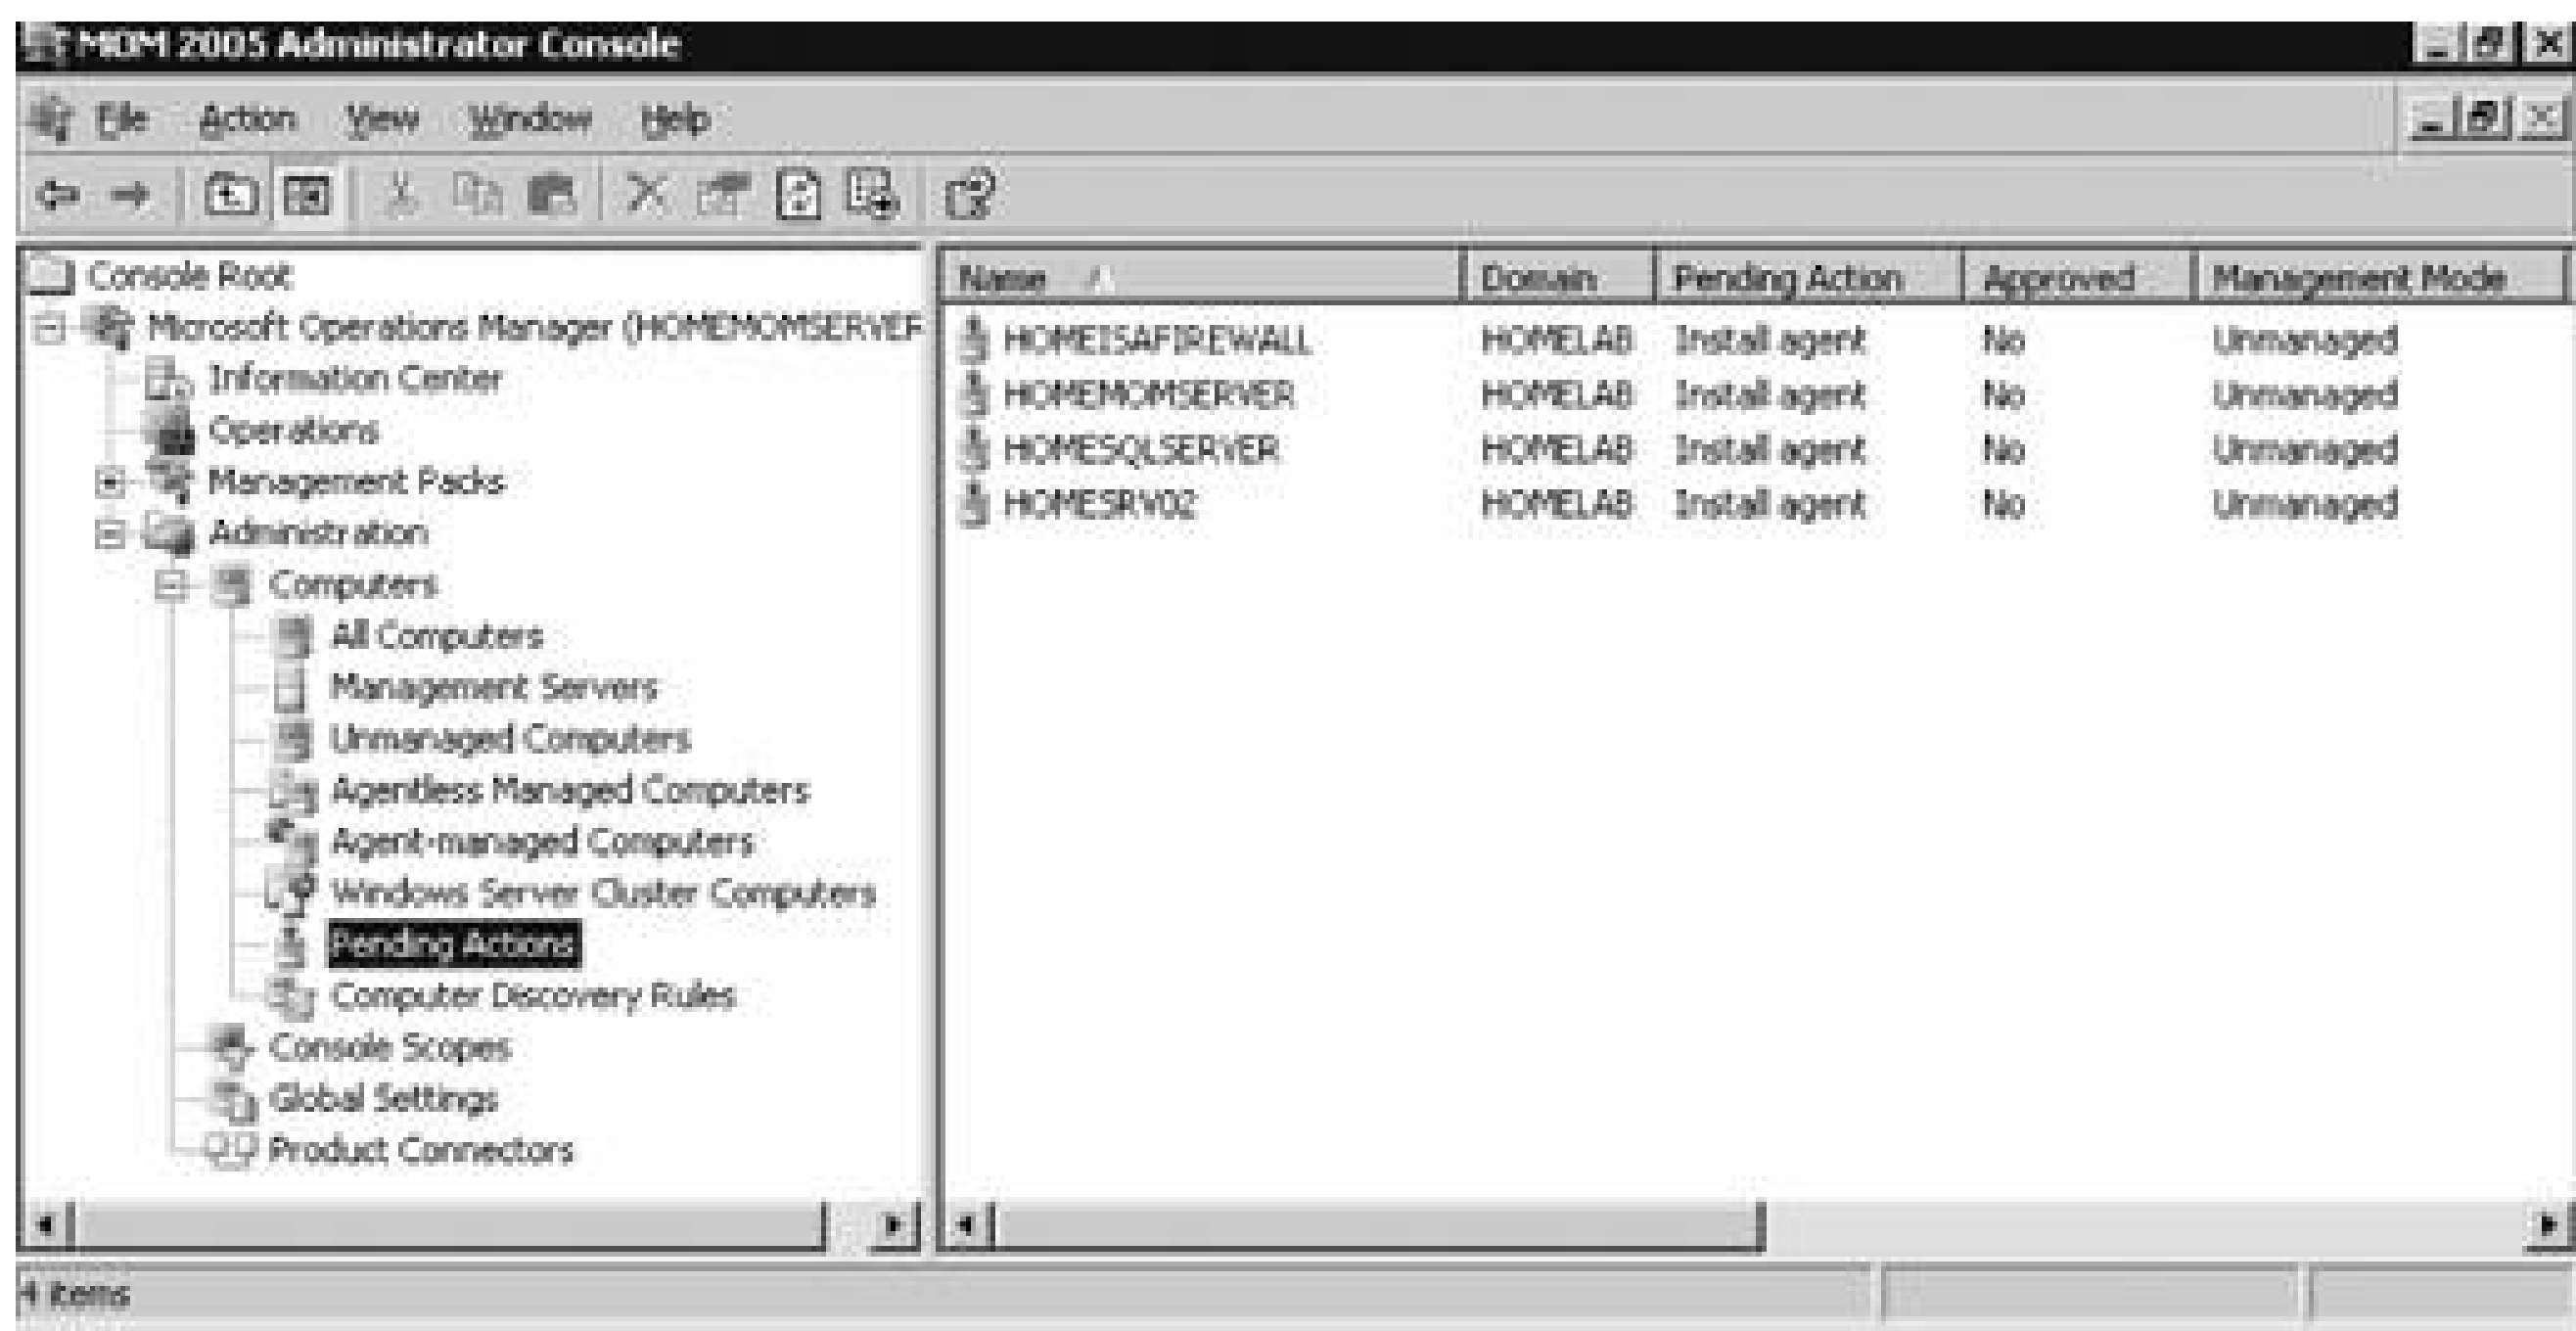
What happens next depends on how you configured automatic agent installation. If you allowed it, the management server will attempt to execute whatever action the rule specifies. For example, if you have created an Include rule that specifies the initial management mode as agent-managed, the management server will attempt to push an agent to the newly discovered computer using the management server action account credentials. If you have created an Exclude rule that has all the other settings configured identically, it will schedule an uninstall action for the targeted agents that will wait until the uninstall delay interval has passed and then uninstall the agent.

If automatic agent installation has been allowed, make sure that the management server action account has local administrator rights on the target machines, or the install/uninstall action will fail.



If you have disallowed automatic agent installation, then the management server will place an entry for the discovered computers in the Pending Actions container, as shown in [Figure 3-11](#).

Figure 3-11. The contents of the Pending Actions container after the computer discovery task has run



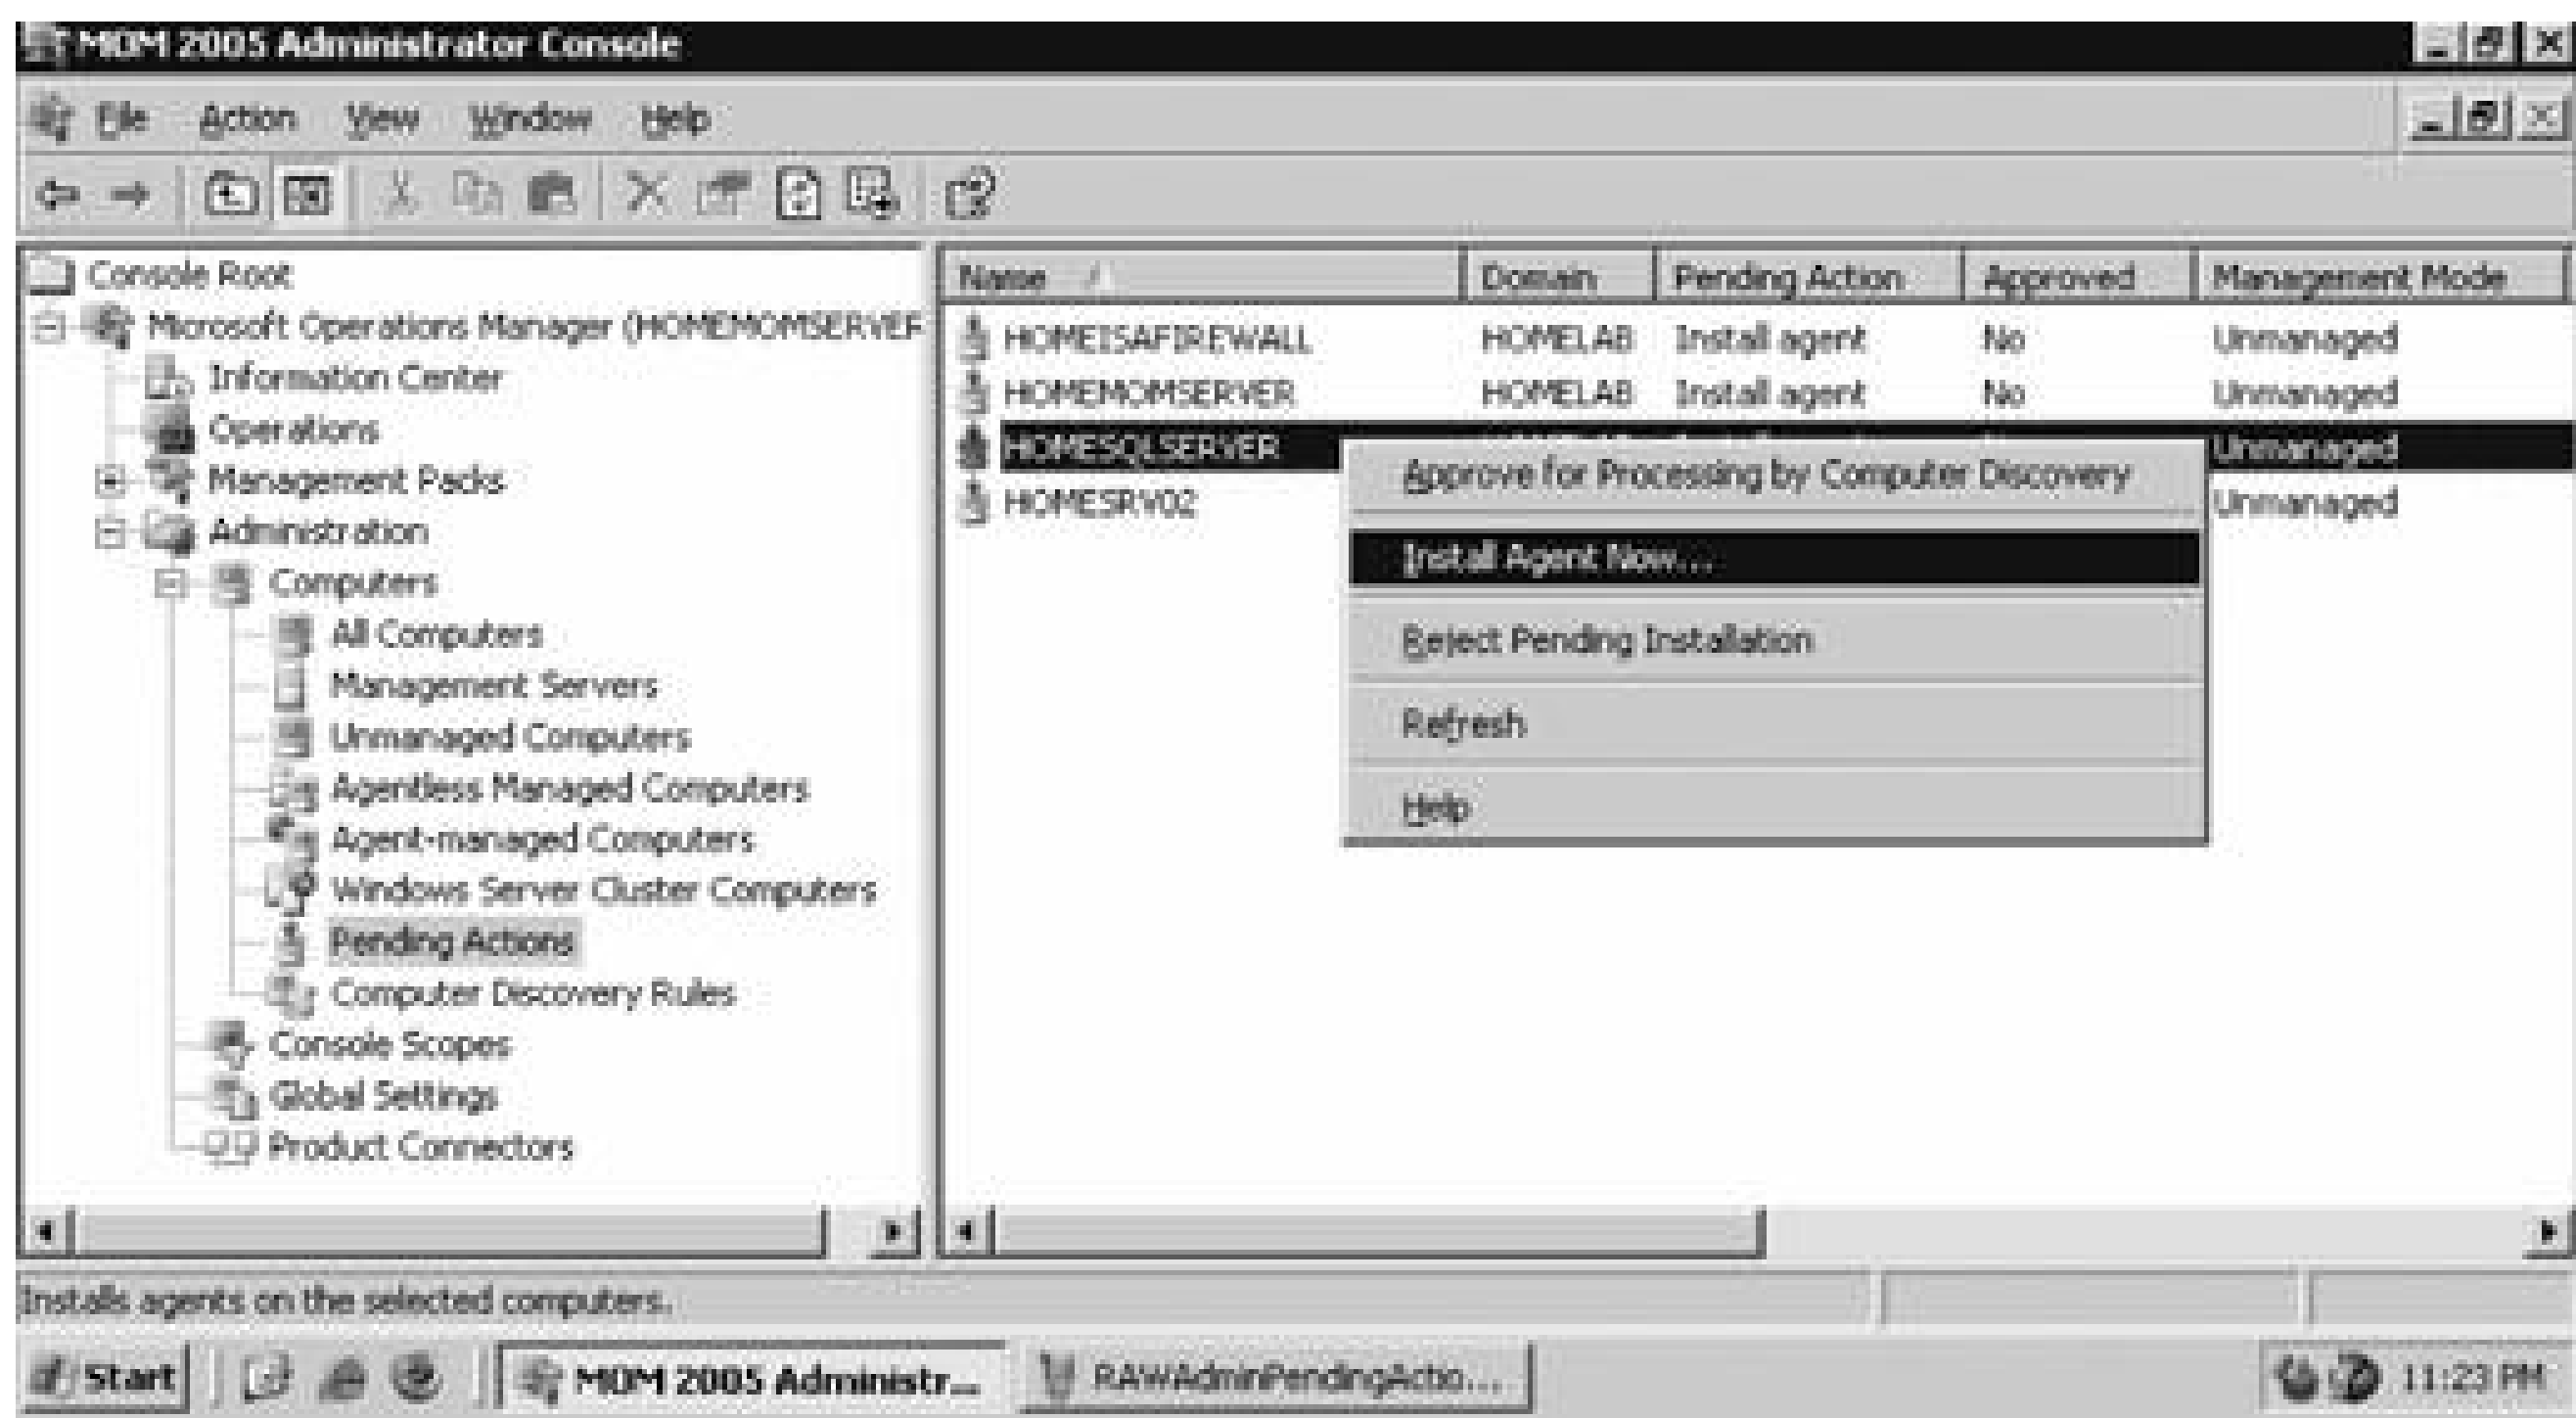
[Figure 3-11](#) shows that four computers have been found that match the search criteria. The domain controller *homesrv02* is listed because the search criteria were applied to domain controllers. *homesqlserver* has been discovered, its management mode is currently Unmanaged, and the desired Management Mode is Agent-Managed. MOM 2005 will hold the discovered computers in this state until an administrator approves or rejects the desired action, or initiates the action immediately. You make this choice by right-clicking the computer or group of computers and selecting the desired action. Again, if you choose to approve the action, MOM will attempt to execute it during the next discovery cycle using the management server action account credentials.

Using the Computer Discovery Rules node to create discovery rules, run discovery cycles, and install/uninstall agents gives you the most flexibility of the three methods. It also gives you access to all options and incorporates the best features of the other two methods.

### 3.3.5. Installing Agents via Computer Discovery Rules

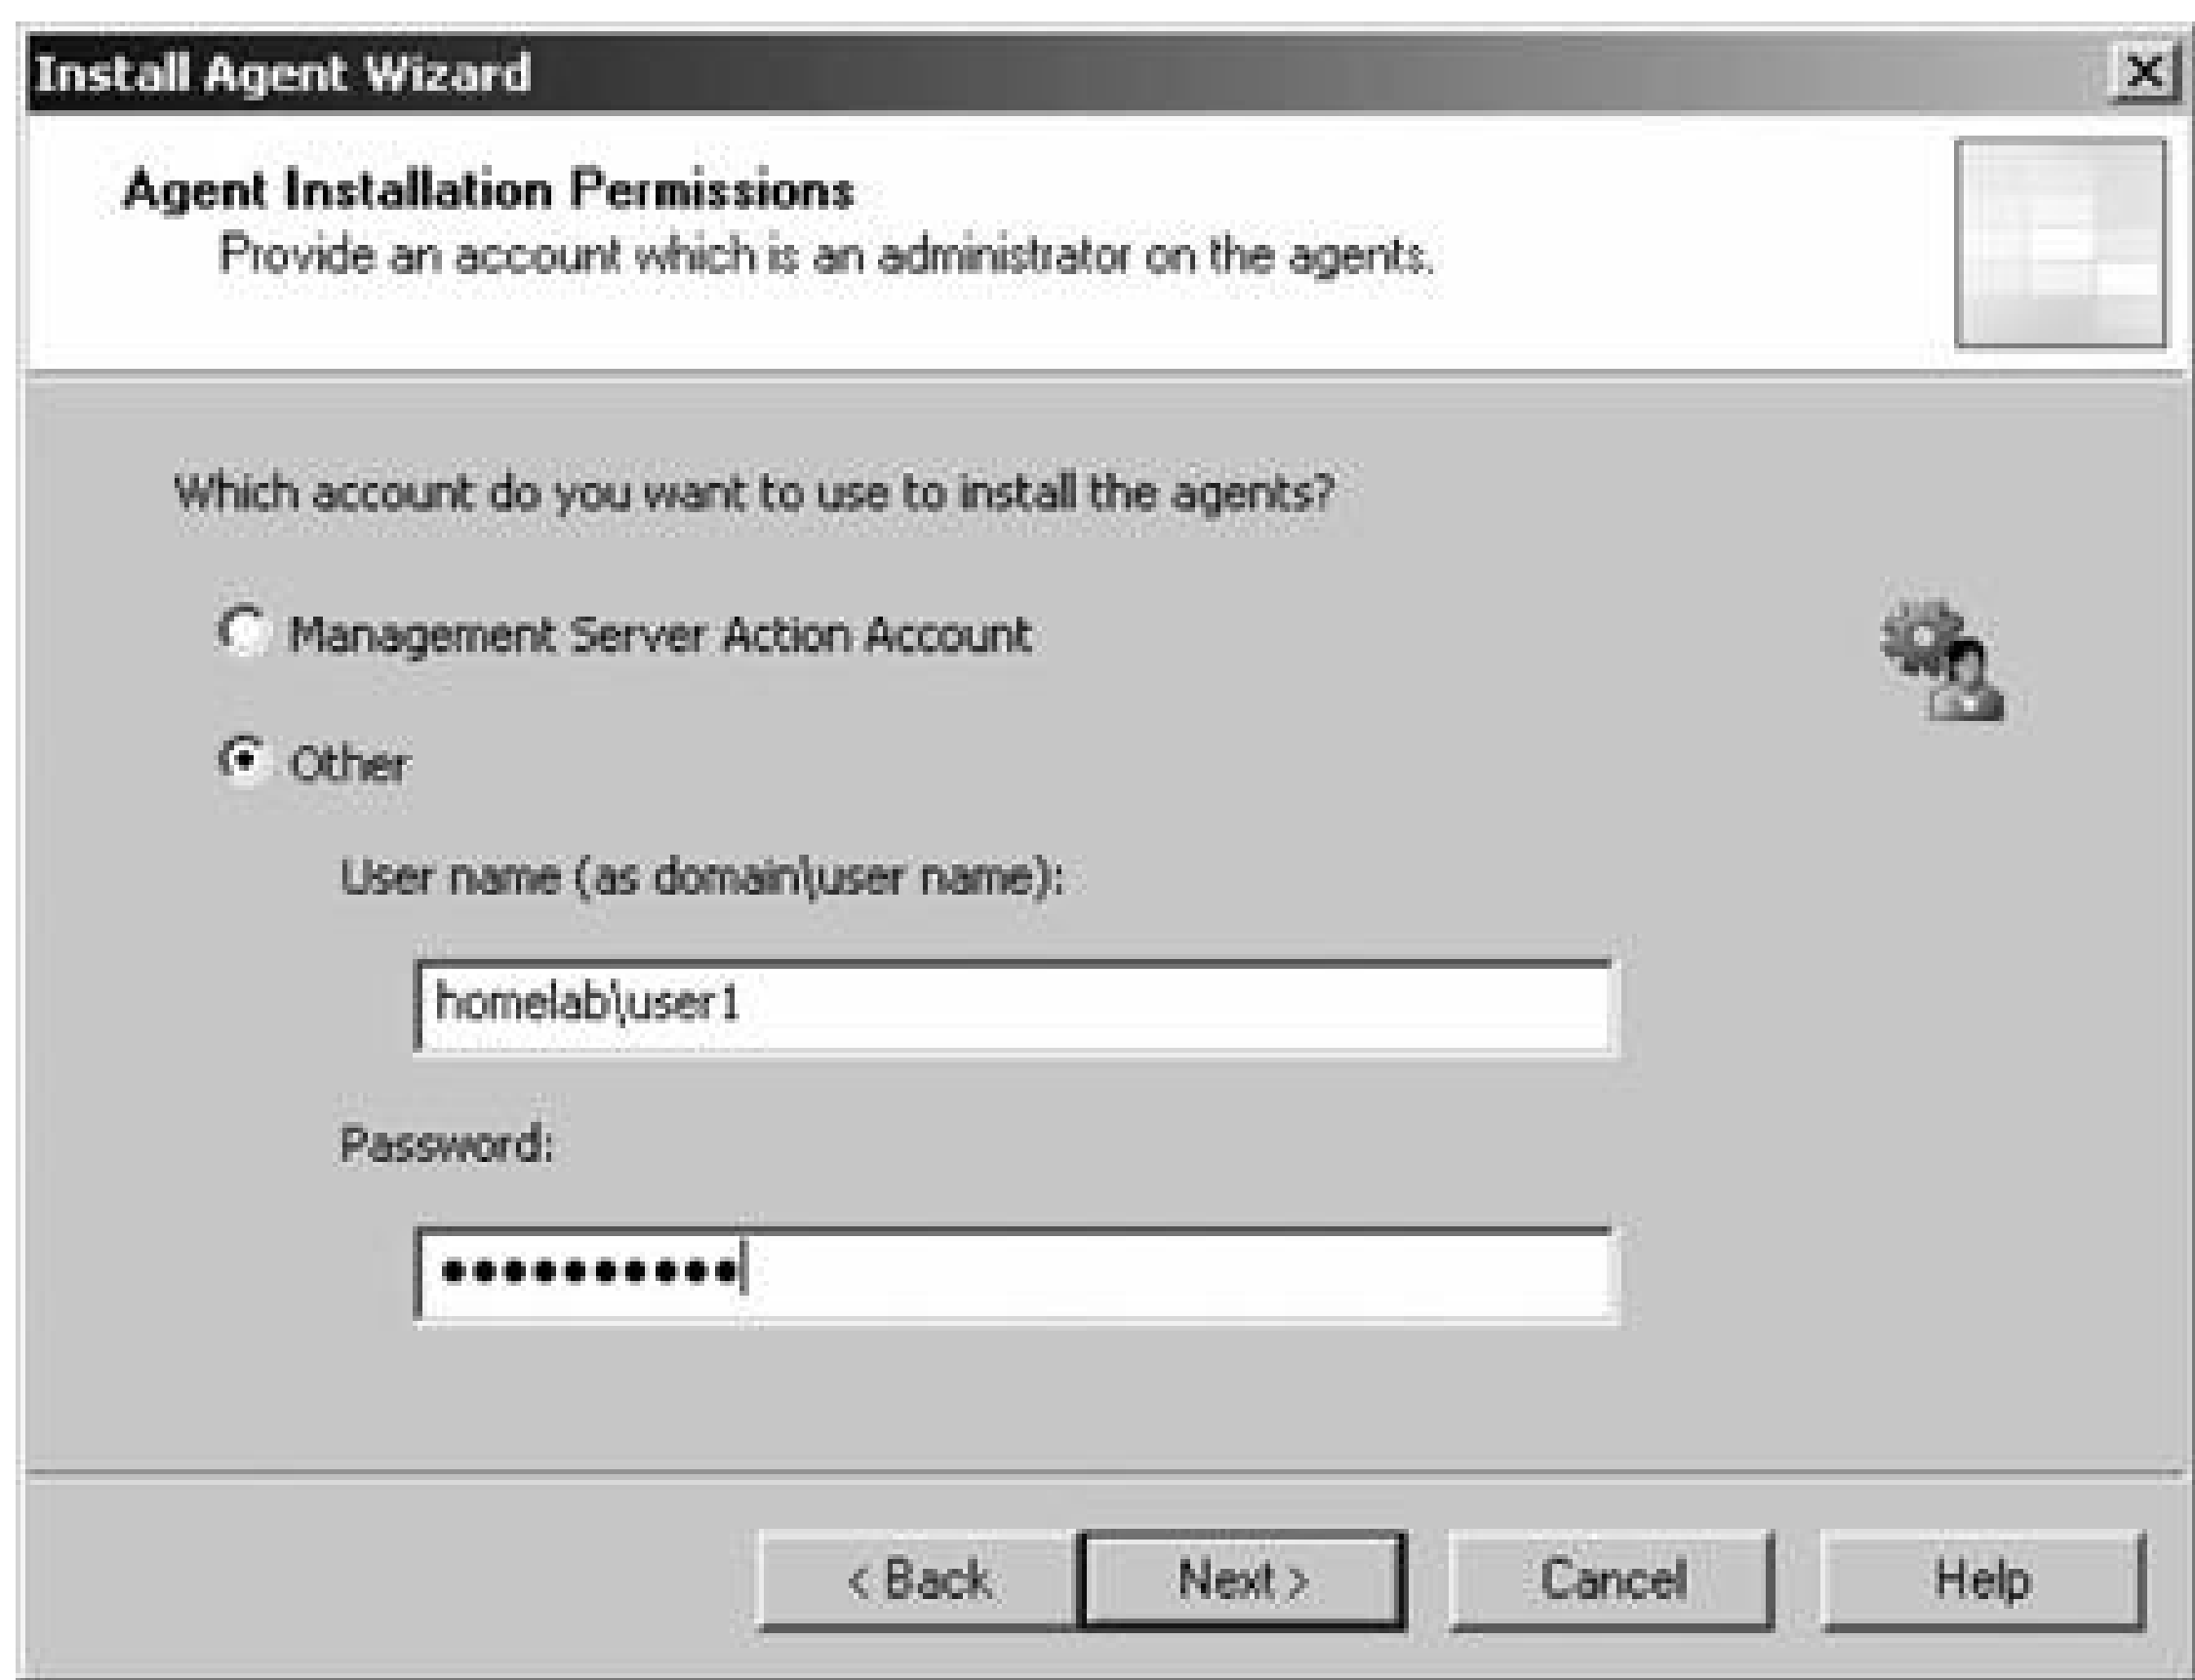
Because Leaky Faucet (and homelab) chose a management server action account that is a domain user account with no administrative rights on any computers except the management servers, approving the desired action for the discovered computers will always fail. The next step is to choose the Install Agent Now option, as shown in [Figure 3-12](#).

Figure 3-12. Selecting the Install Agent Now option



This starts the Install Agent Wizard, which prompts you for a set of credentials to install the agent or the target computer, as shown in [Figure 3-13](#). These credentials are stored securely during the instal process and are then disposed of. Choosing to install agents in this manner does not impose any restrictions. This technique is used to build security into your procedures. For example, if the installation account is the domain administrator account or equivalent, or the target machine's local administrator account, then an administrator would enter the credentials at this time. If this person isn't the original MOM administrator, then an additional individual will have to be added to the agent deployment or removal process, which introduces a human checkpoint and bottlenecksecurity is always a trade-off.

Figure 3-13. Provide the necessary credentials at agent install time



Next, you are prompted to identify and provide the credentials for the Agent Action Account, as shown in [Figure 3-14](#). The local system account is chosen and since this account has full rights to the local machine and the password is automatically managed, you are relieved of that responsibility.

The next prompt is for the agent installation directory. Remember this defaults to the *%Program Files%\Microsoft Operations Manager 2005* folder, which is on the C: drive. Since this is the drive that was inventoried for space, accept the default.

This takes you to the Completing the Install Agent Wizard page shown in [Figure 3-15](#), which provides you with a summary of the actions and allows you to choose to view the task progress. Click Finish to start the installation of the agent.

By selecting the "Show task progress" checkbox, the progress of the agent installation is tracked and displayed, as shown in [Figure 3-16](#). Note that if you selected multiple machines, this single display window would have tracked all of the agent installations.

Once finished, the tracking window will display the final status of the action. If a successful status is returned, this is the first indicator that the agent installation has completed and the target server is now being monitored. If the installation process encounters an error and fails, the tracking window will return a failure status along with initial diagnostic information for the failure (see [Figure 3-17](#)). The failure was caused by not providing the correct agent installation credentials to the agent installation wizard.

Figure 3-14. Selecting the agent action account



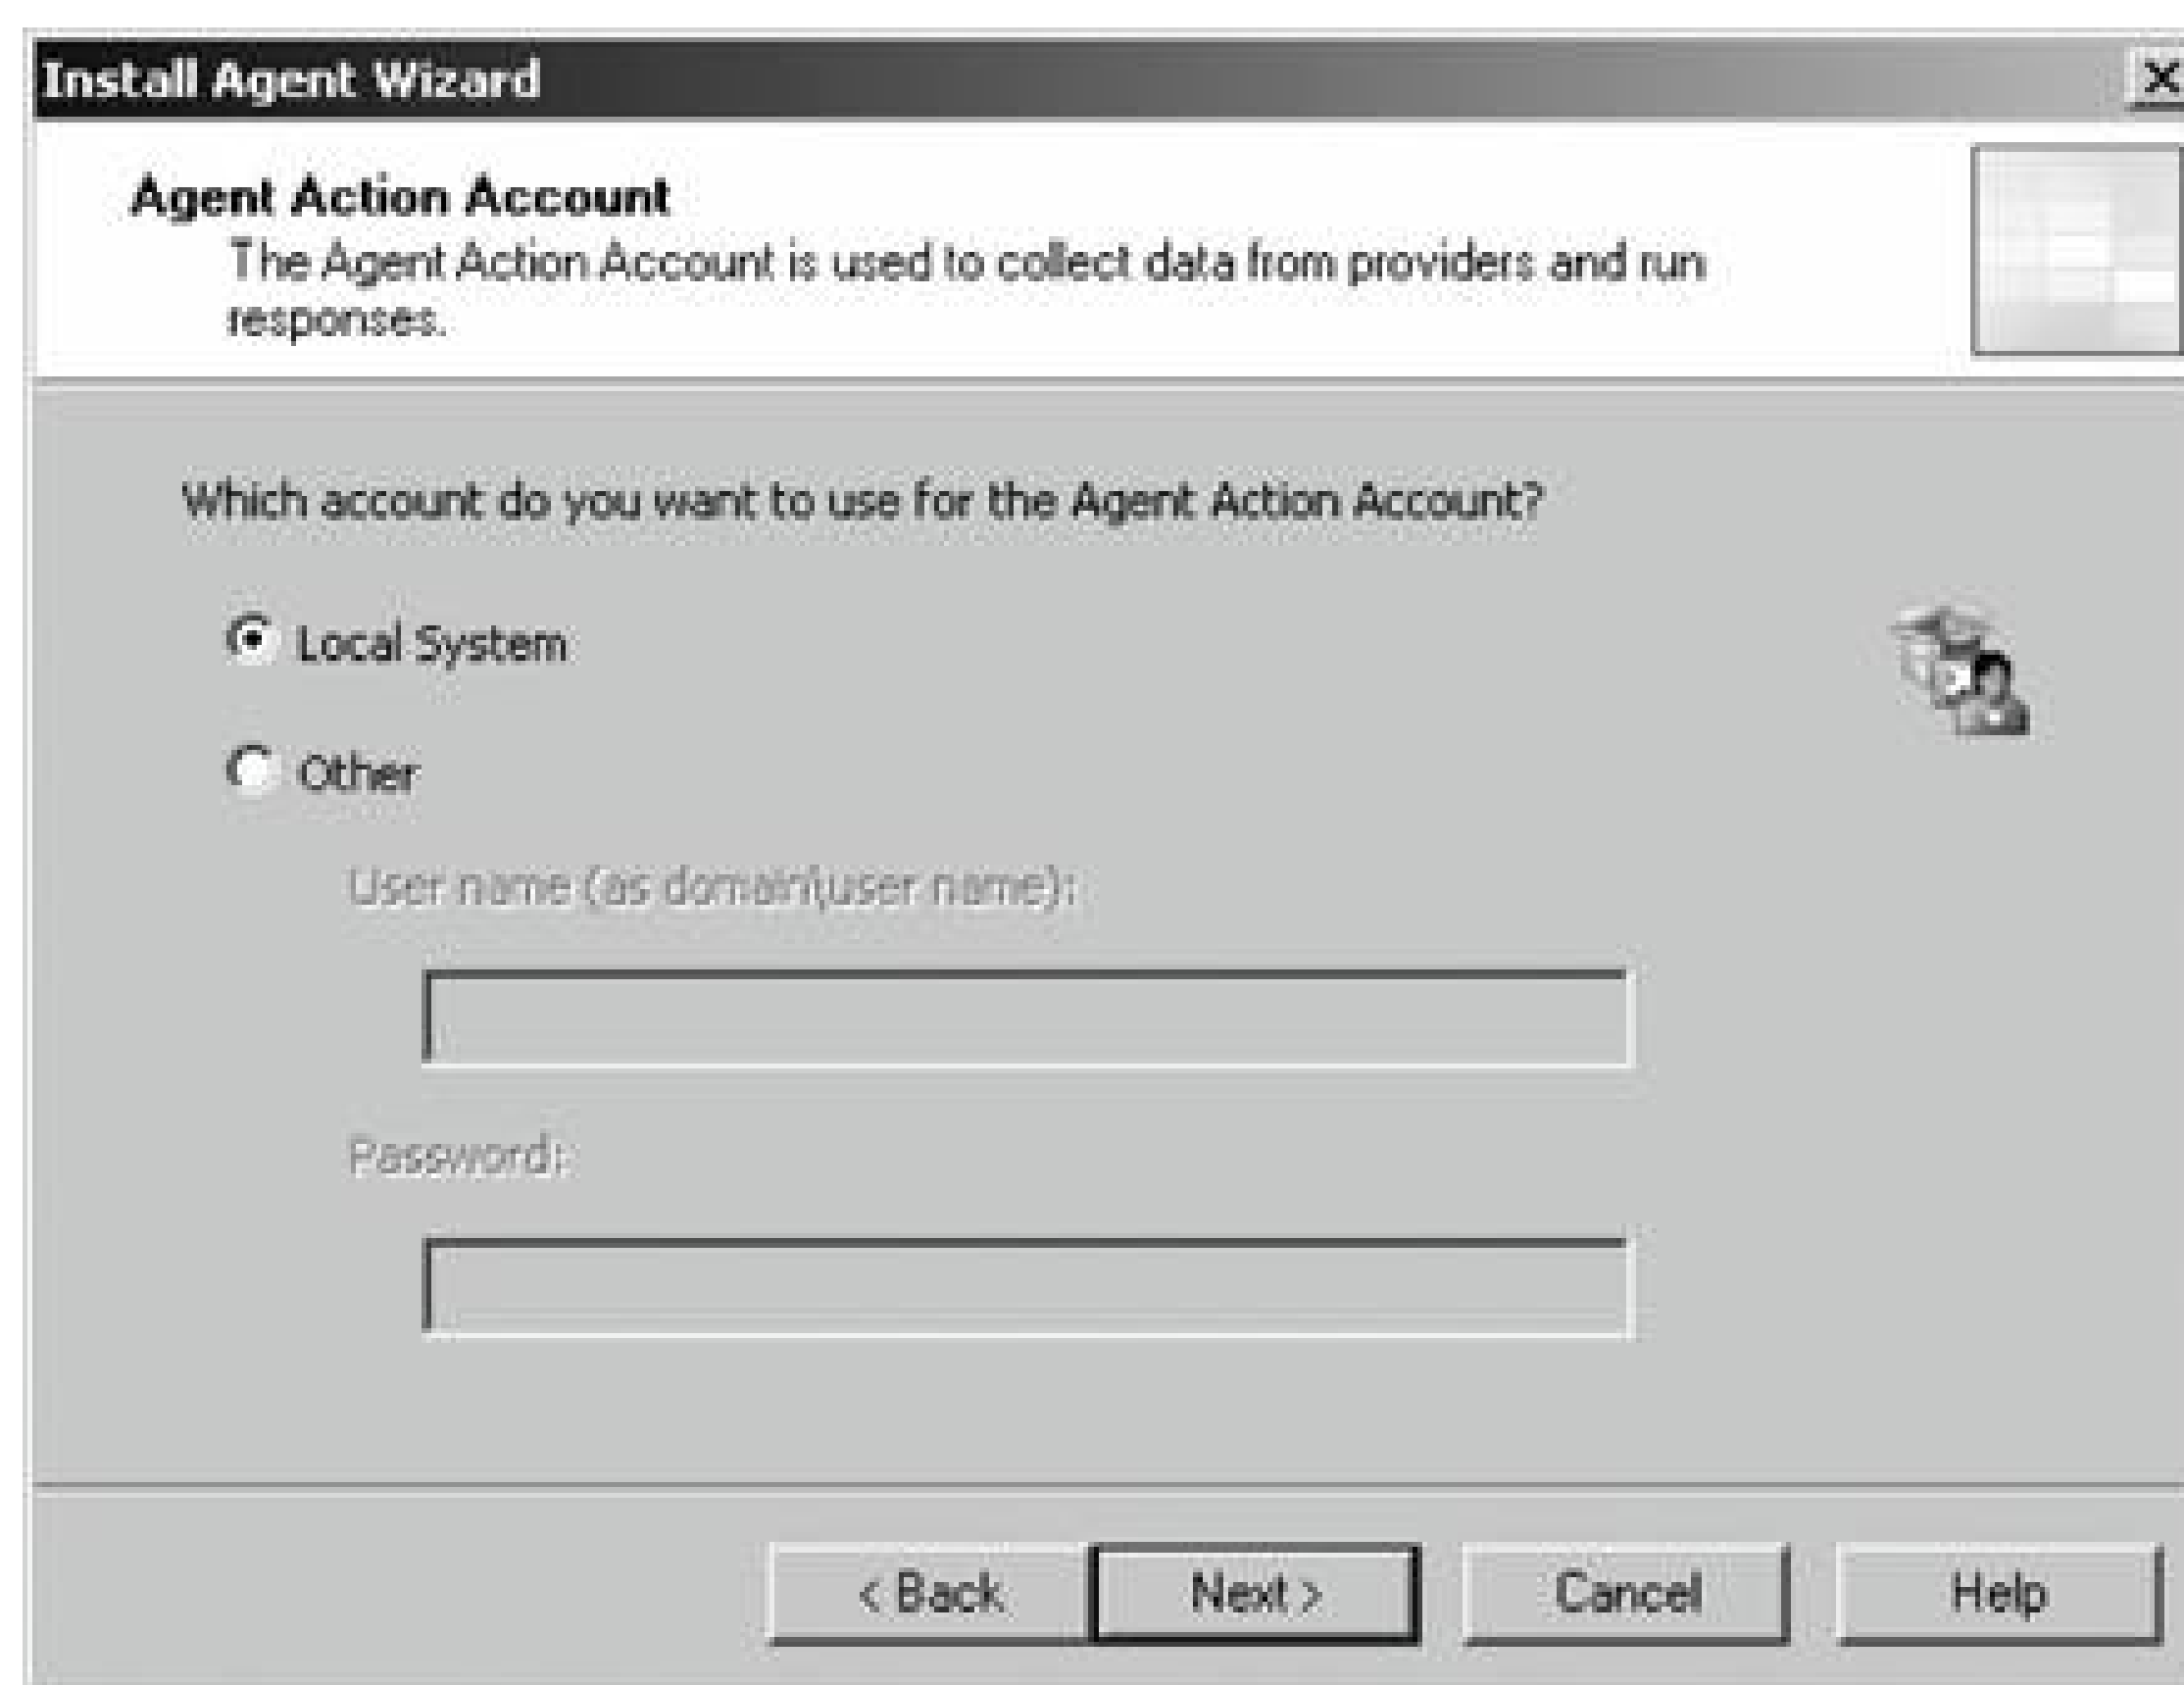


Figure 3-15. The summary of the Agent Installation Wizard

At this point, the server *homesql/server* has been removed from the Pending Actions container and placed into the Agent-managed Computers container. The computer is being monitored and will now

appear in the operator console as well. The computer that the installation failed on will remain in the Pending Actions container.

Figure 3-16. Tracking the progress of the agent installation

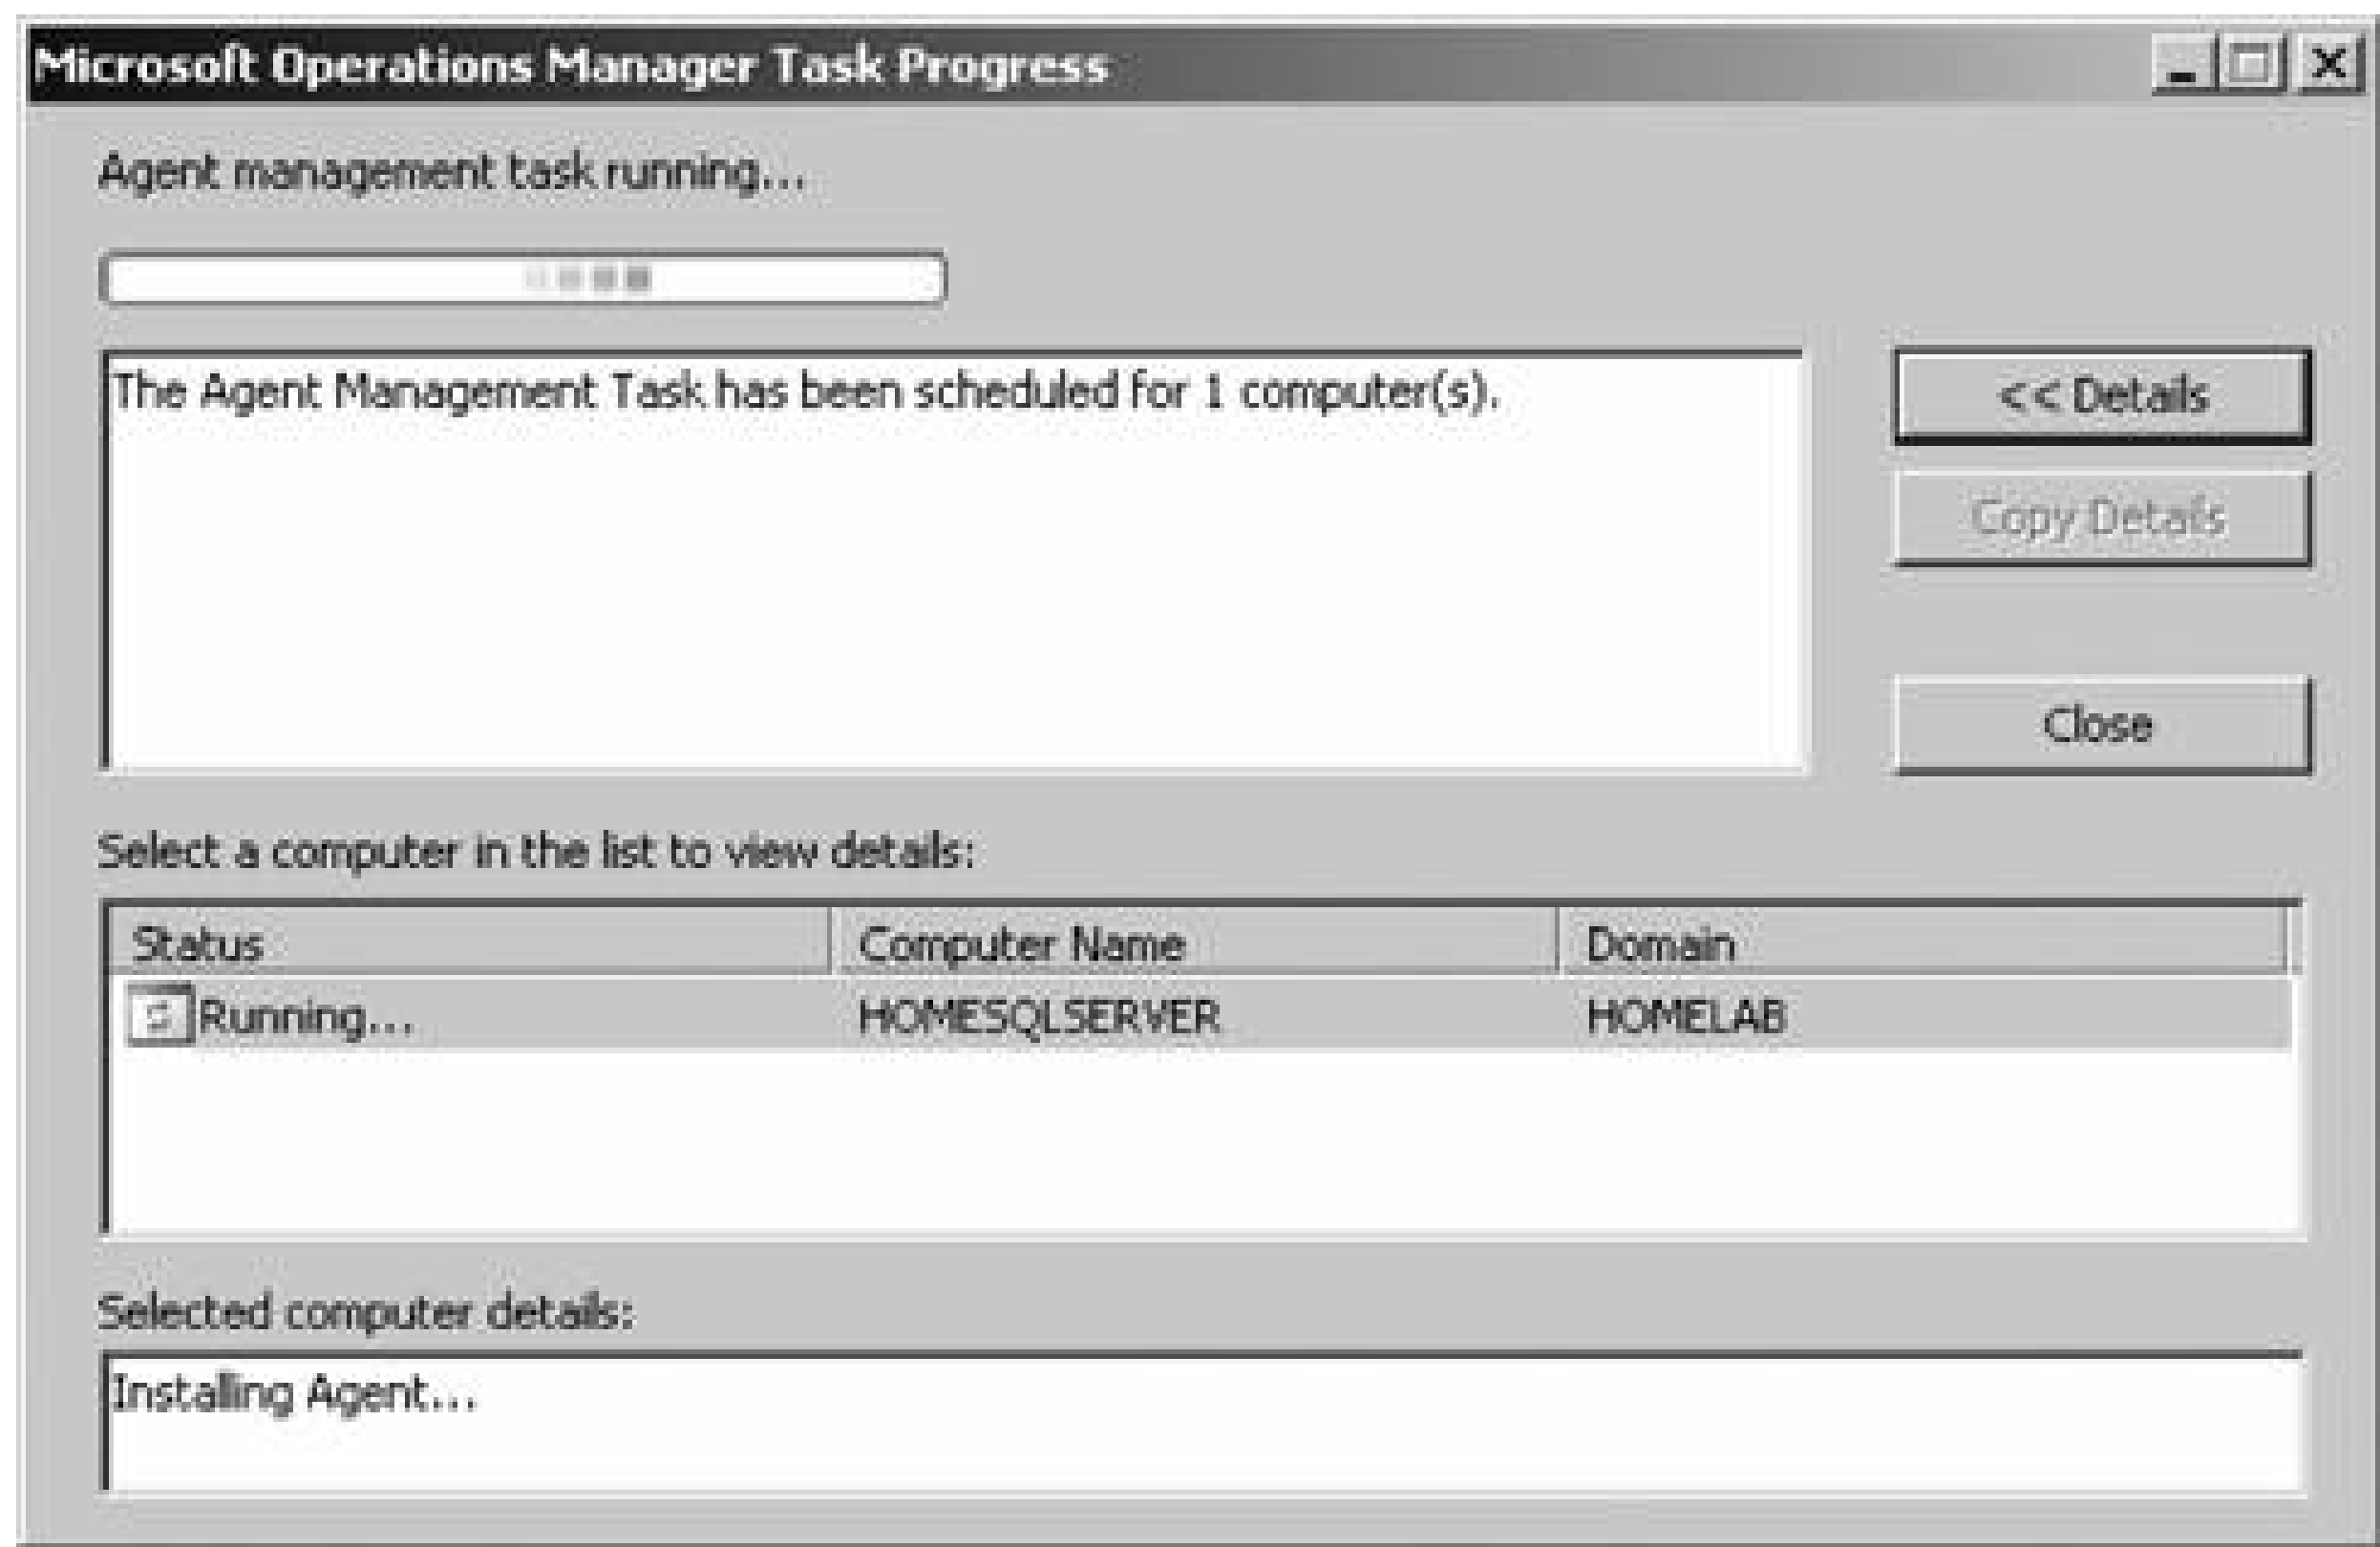
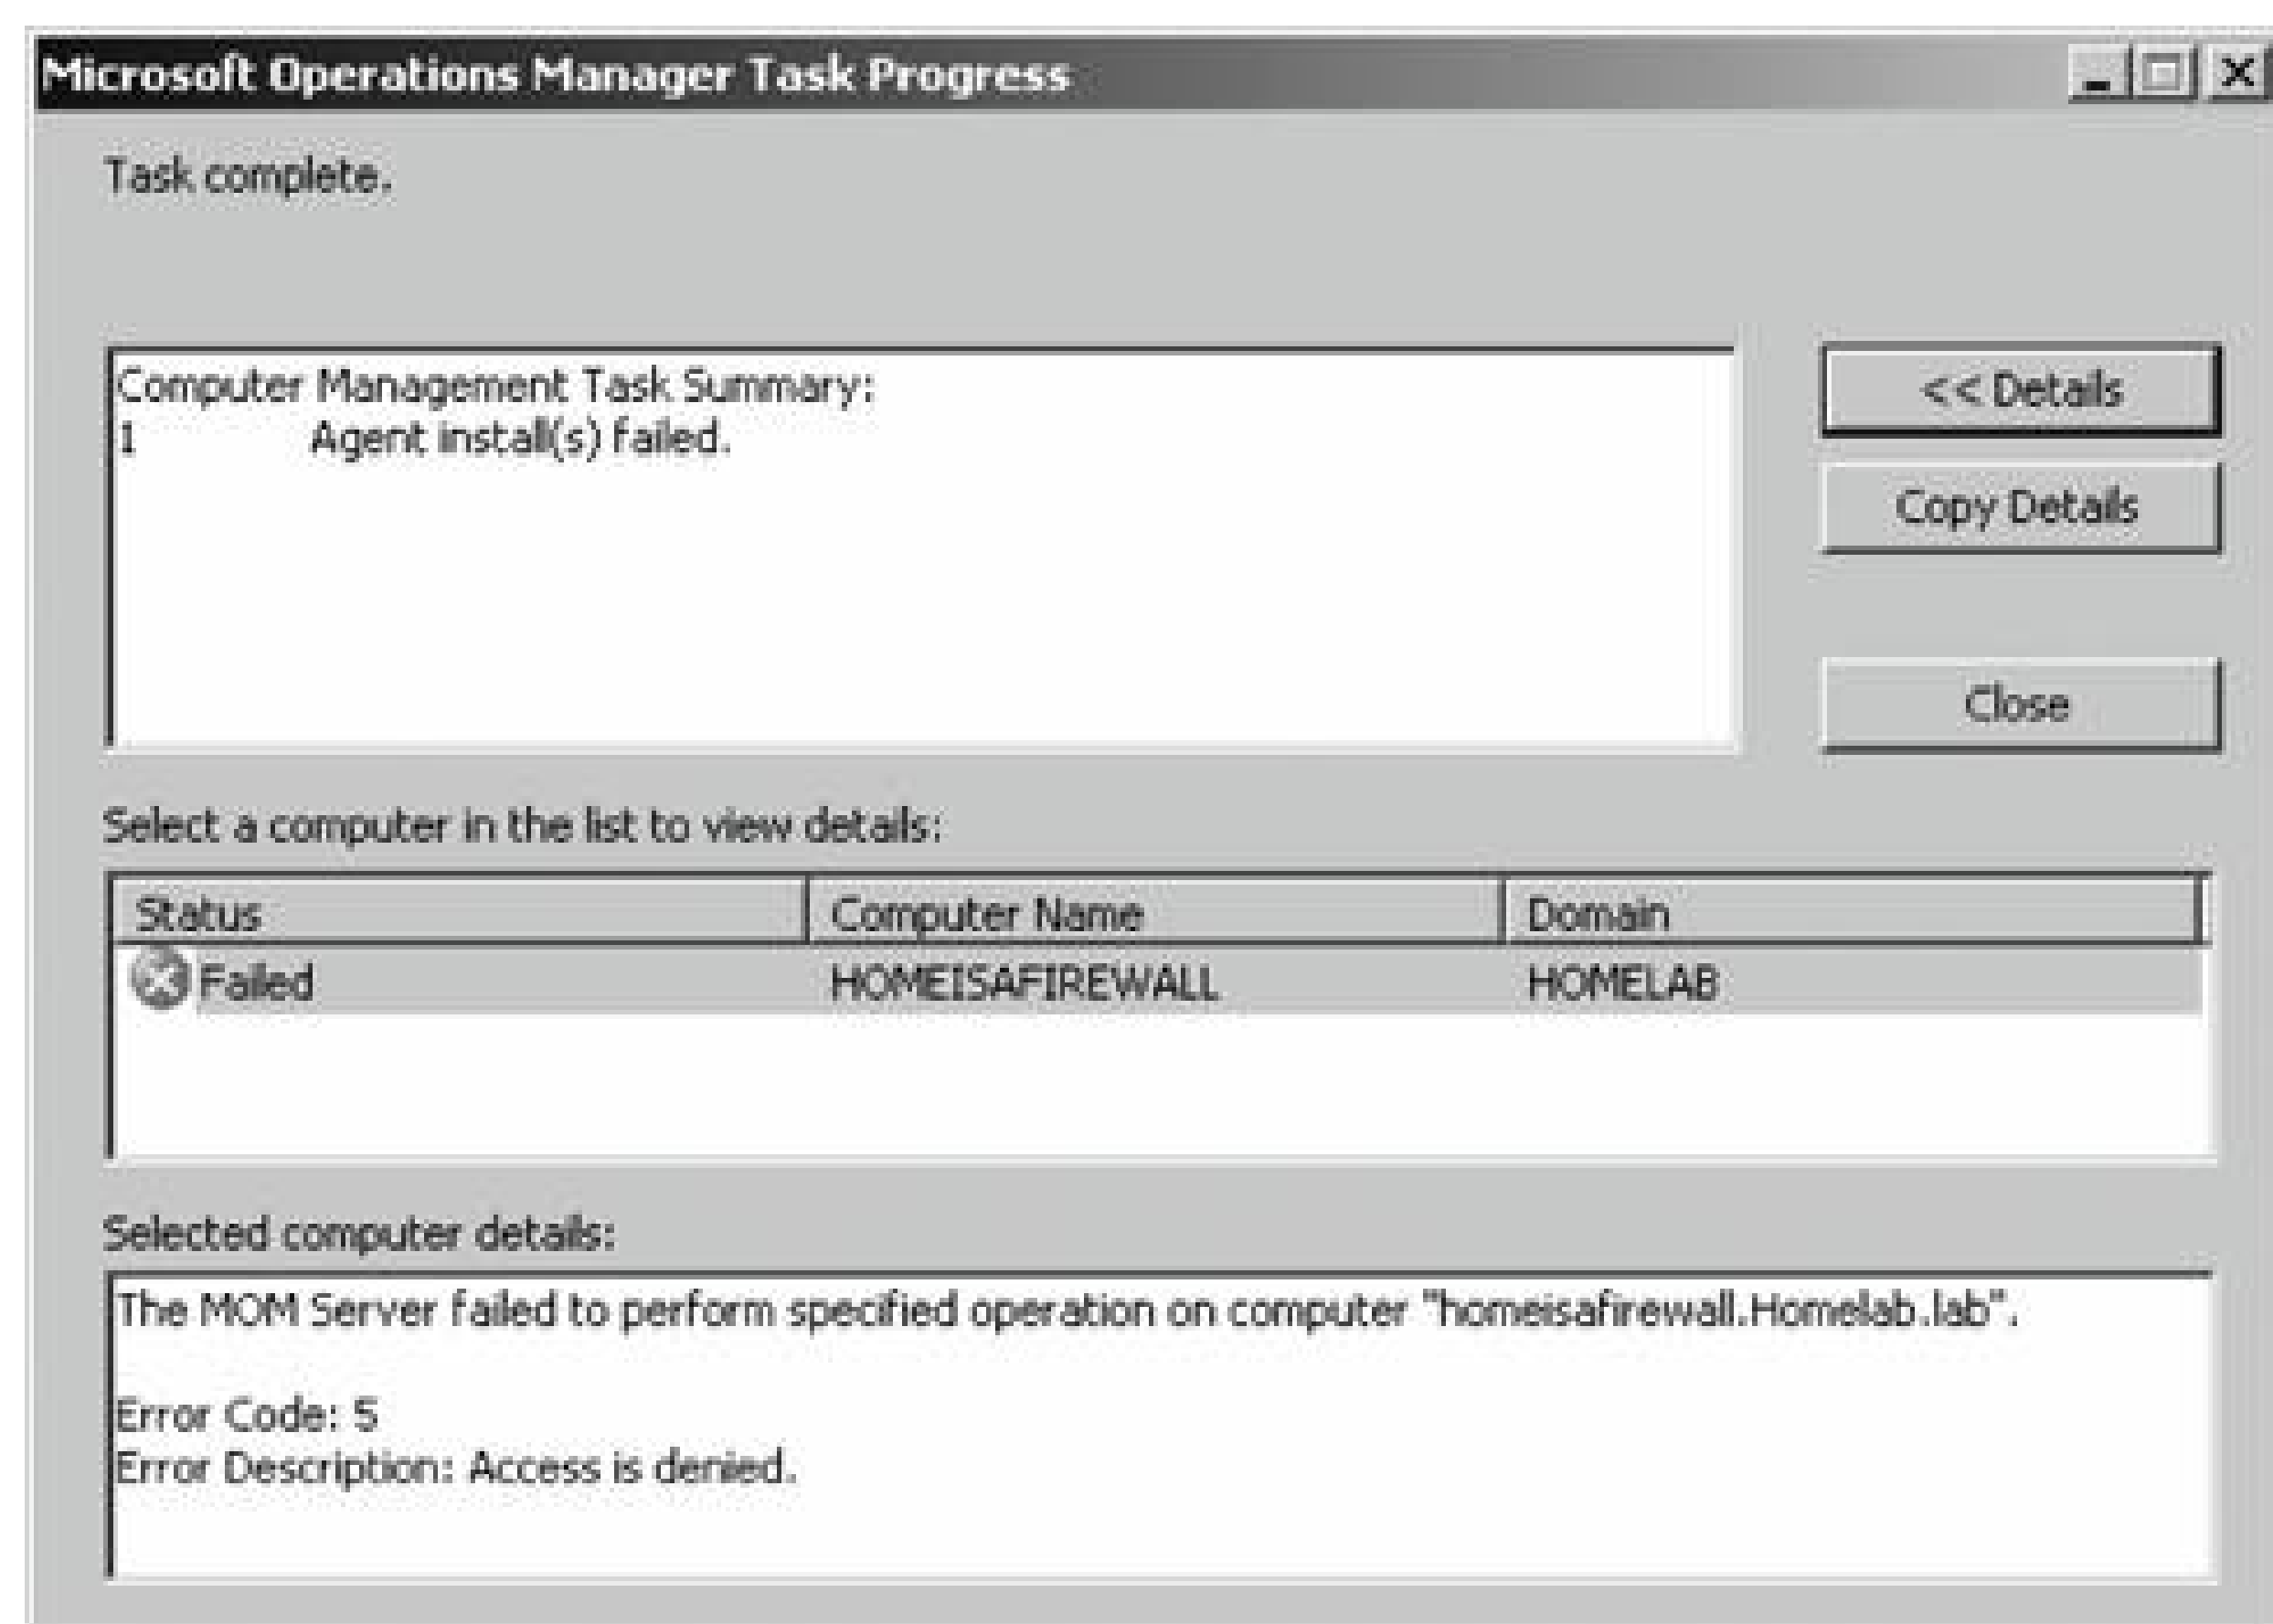


Figure 3-17. Returned failure status

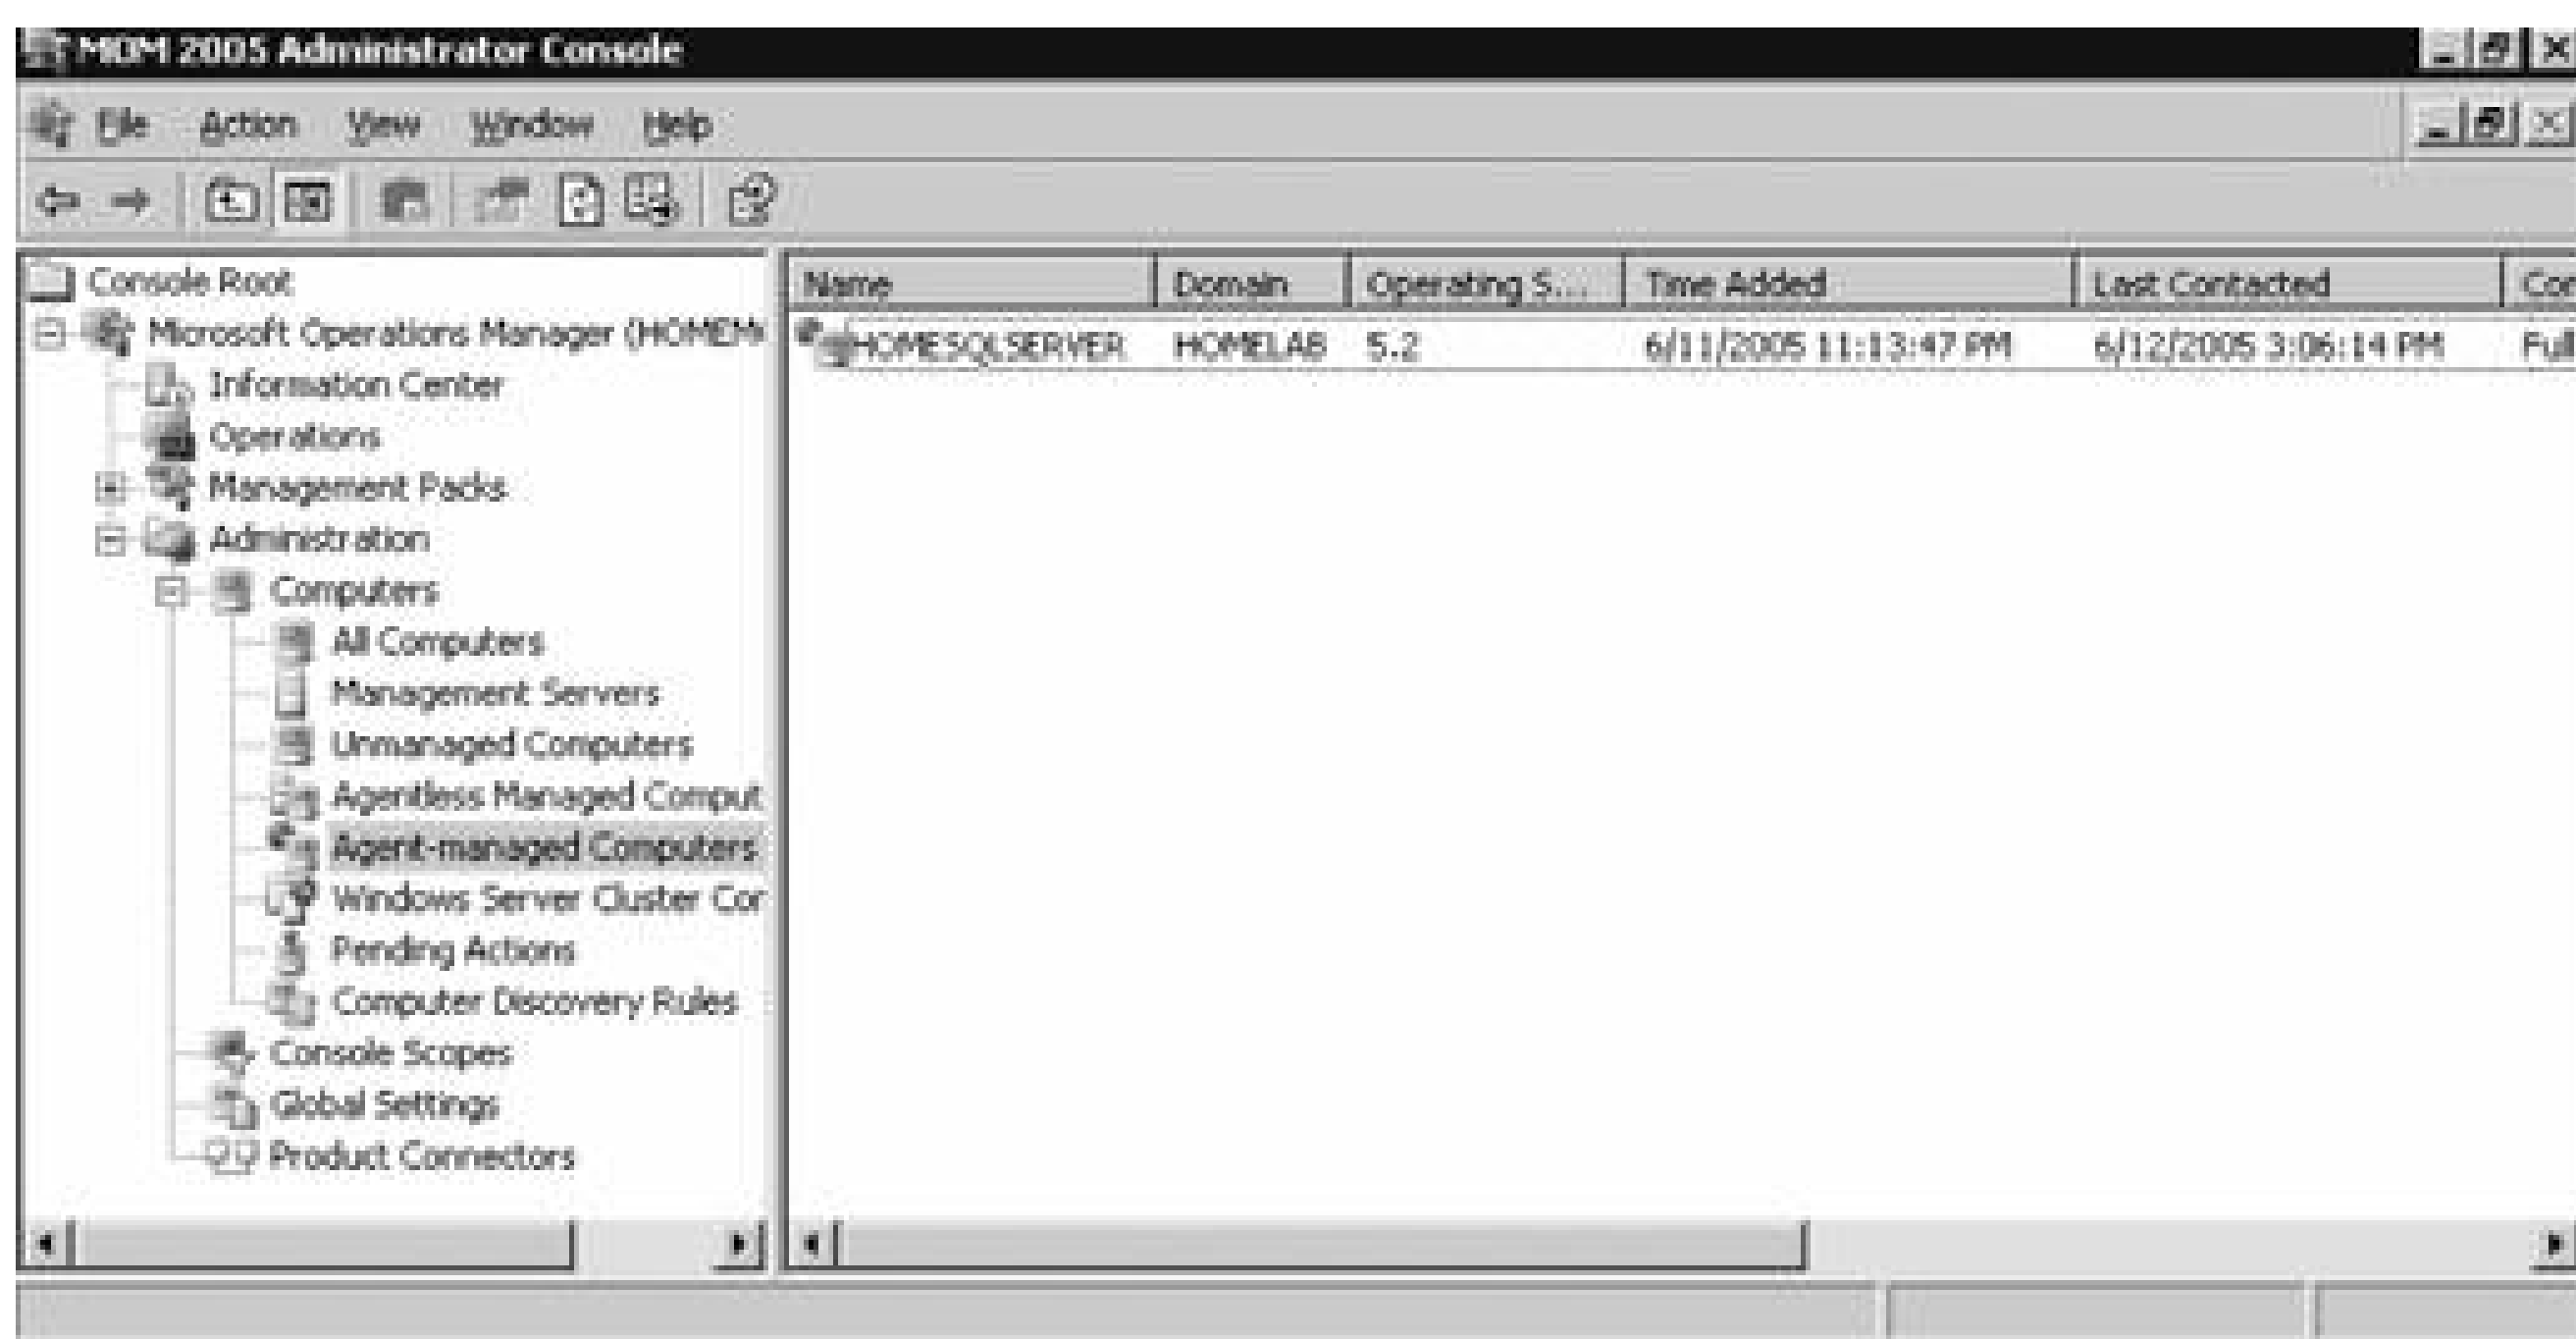


### 3.3.6. Confirming Agent Functionality

The last task in successful agent deployment is to confirm that the agent has installed successfully and is returning data to MOM. The first indication of a successful deployment is the success status returned at the end of the agent installation wizard. The second indication is that you can examine the computer data in the results pane of Agent-managed Computers container and access the properties of the agent on the target computer, in this case *homesql/server* ([Figure 3-18](#)).

Figure 3-18. Agent-managed Computers data in the Administrator console



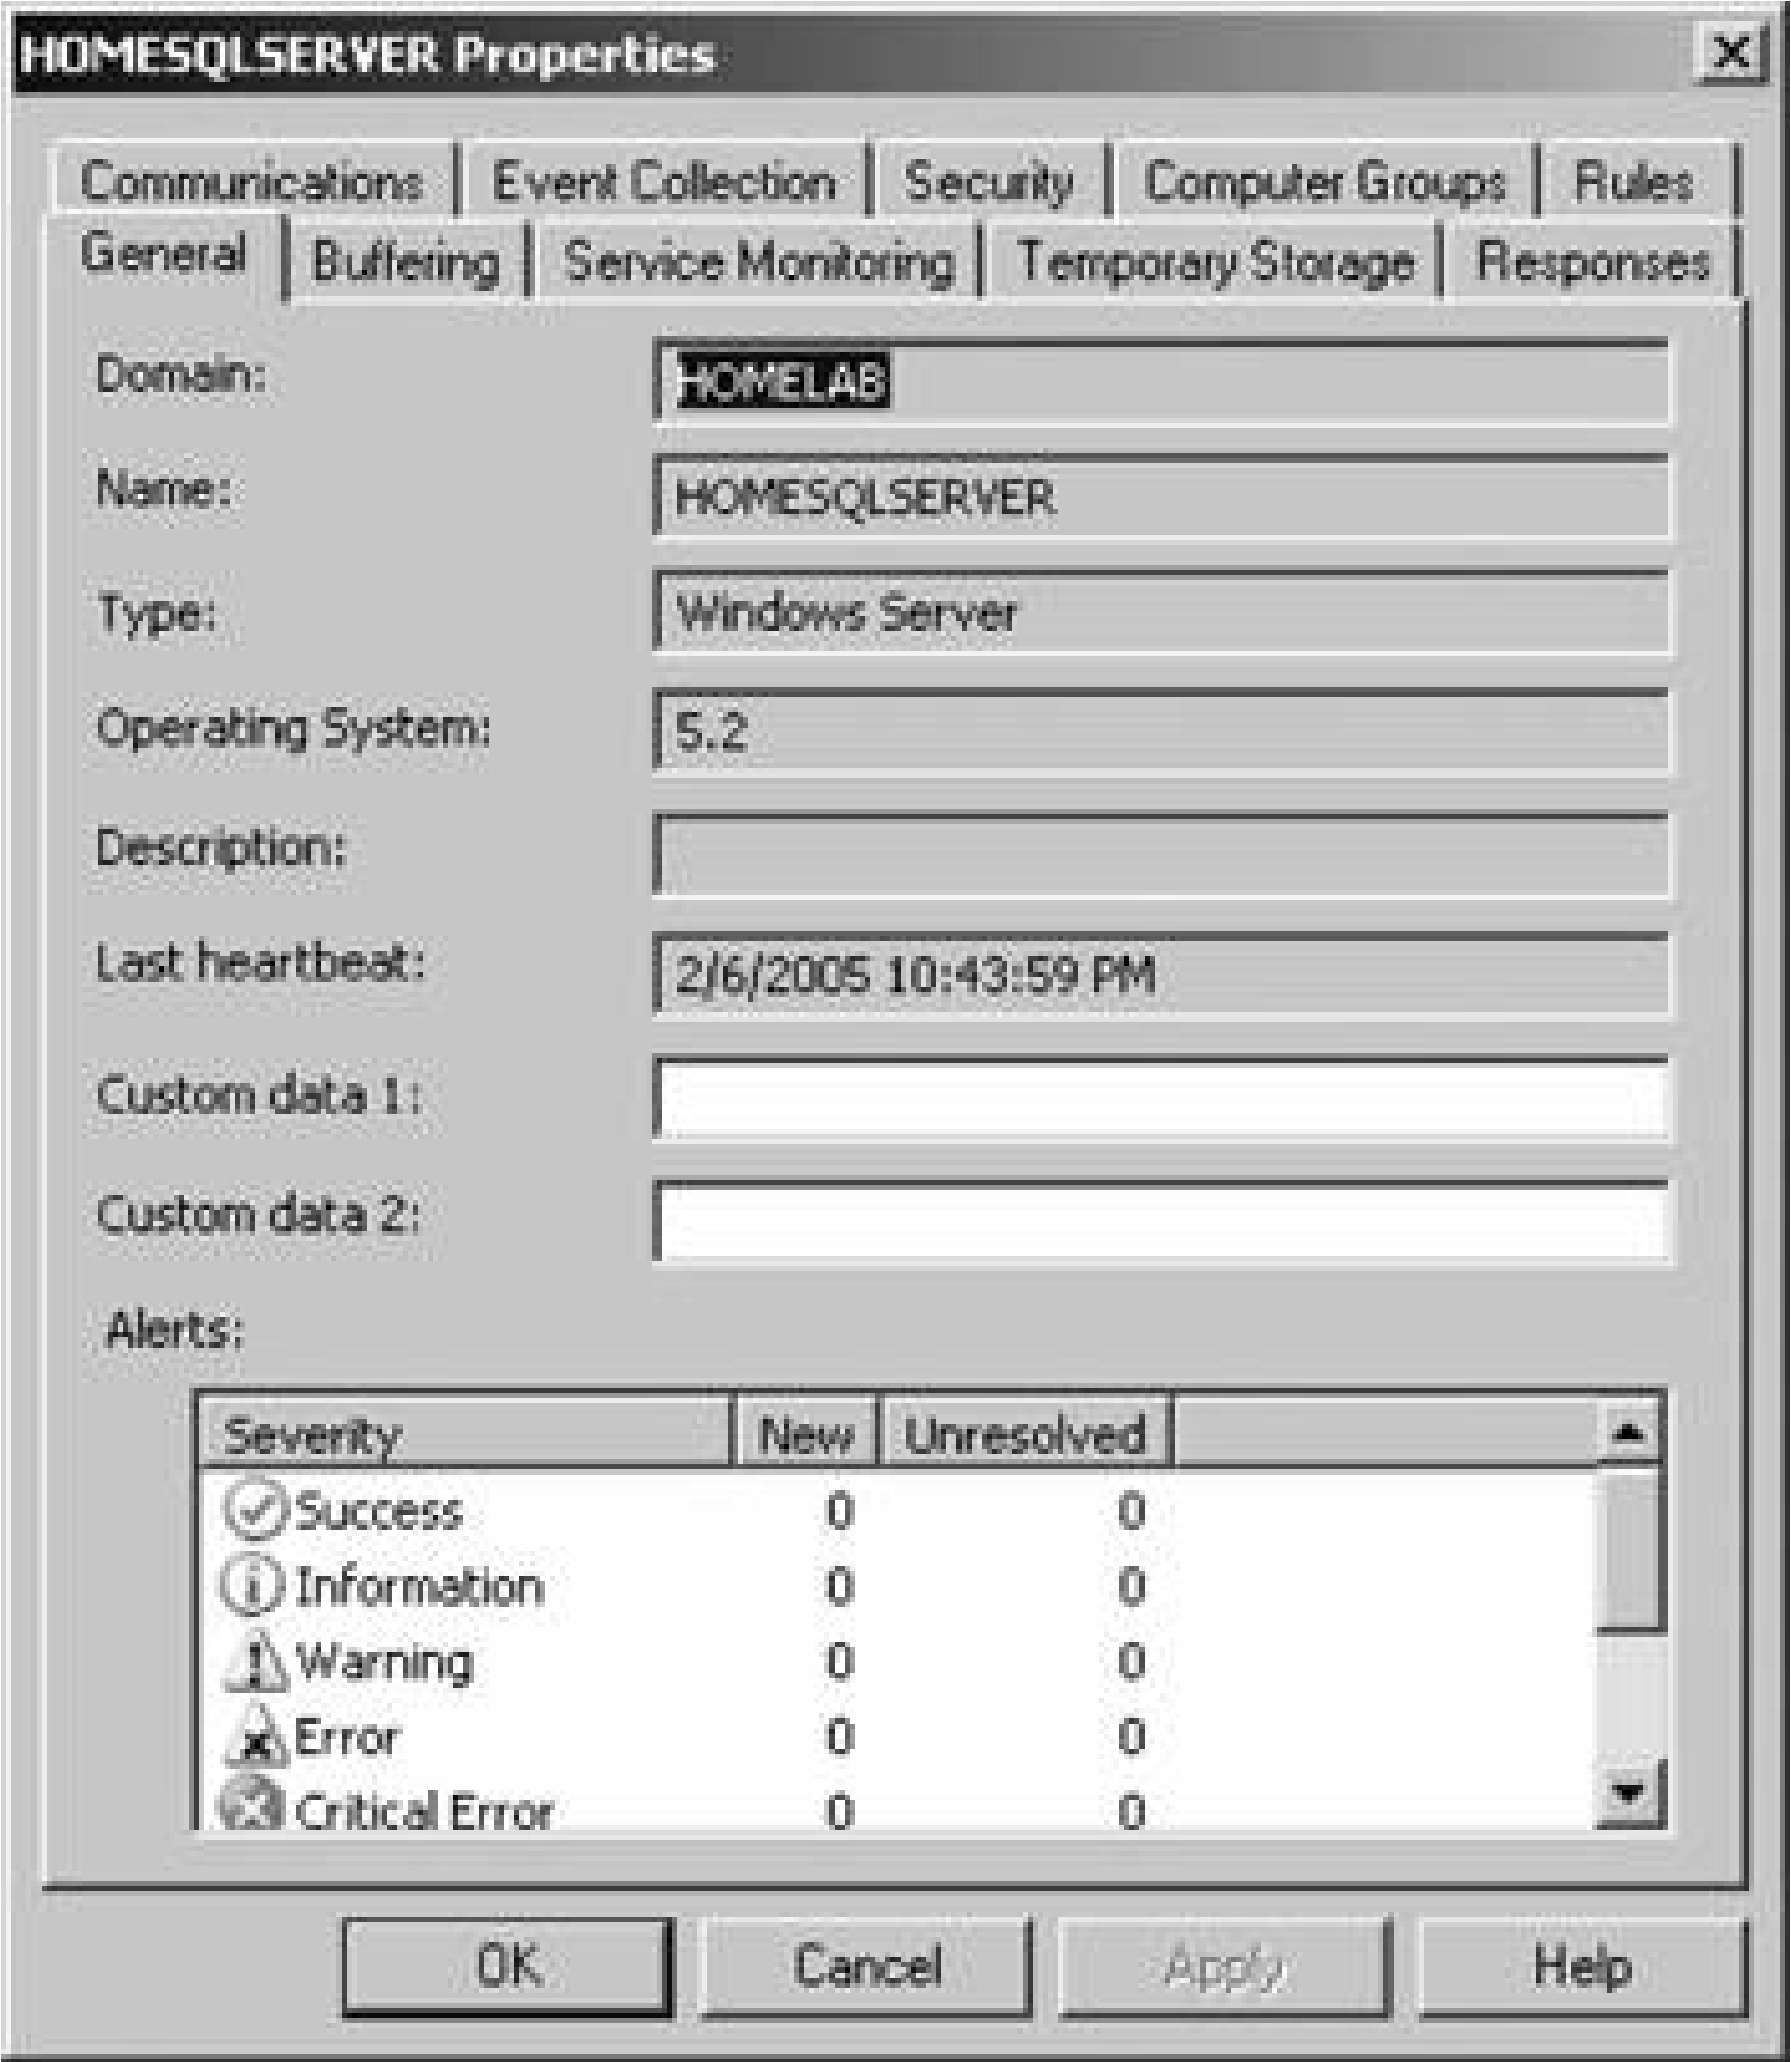


[Figure 3-18](#) shows that the agent is returning data, the time the server was last contacted, the time and date that it was successfully added to MOM, and the primary management server. If you right-click the managed computer and bring up its properties, as shown in [Figure 3-19](#), you can see additional information such as the time of the last good heartbeat and any alerts generated for this computer. These are all definitive indicators that the agent has installed successfully.

The last step is to confirm that end-to-end monitoring is occurring successfully. This is done in the Operator console. Navigate to the State view, select the target computer, and launch the Test End to End Monitoring task against it. This should produce a 9898 informational event that is viewable in the Public Views Task Status object indicating the following:

```
Type: Information
Domain: HOMELAB
Computer: HOMESQLSERVER
Description: The task 'Microsoft Operations Manager\Test End to End Monitoring' has
successfully executed against 'Computer:HOMELAB\HOMESQLSERVER'.
Launched By: HOMELAB\user1
The following output has been generated:
Source: Microsoft Operations Manager
Event Number: 9898
Provider Type: Generic Provider
Provider Name: Internally-generated Event
Source Domain: HOMELAB
Source Computer: HOMESQLSERVER
```

Figure 3-19. The properties of a newly installed agent as shown in the Administrator console



This shows that the agent was successfully deployed and is working.[Chapter 4](#) discusses how processing rules are passed to the agent and what the agent does with them. As far as agent administration, there are only a few other tasks that you can perform: uninstalling the agent, forcing it into unmanaged mode, and updating the agent's settings.

### 3.3.7. Creating Computer Discovery Rules and Installing Agents via the Install/Uninstall Agents Wizard

One alternate method of installing agents combines the creation of the computer discovery rule with a limited discovery cycle and agent install/uninstall. This is probably the quickest method for installin or uninstalling agents, but you lose the ability to configure some of the deployment options. For example, you can only create Include computer discovery rules. But you can browse for target computers in addition to creating the search criteria.

To start the Install/Uninstall Agents Wizard, click on the hyperlink by that same name in the results pane, as previously shown in [Figure 3-5](#). Or right-click on the Computers node, the All Computers node, or the Agent-managed Computers node in the navigation pane of the Administrator console.

After the welcome page, select the management server to apply the computer discovery rule (see [Figure 3-20](#)). In this case, *homemomserver2* is selected.

Because *homemomserver2* currently has no agents reporting to it, the wizard automatically proceeds

through the steps to install or uninstall agents. If you selected a management server that was already managing agents, the wizard would prompt you to proceed with the install or uninstall process.

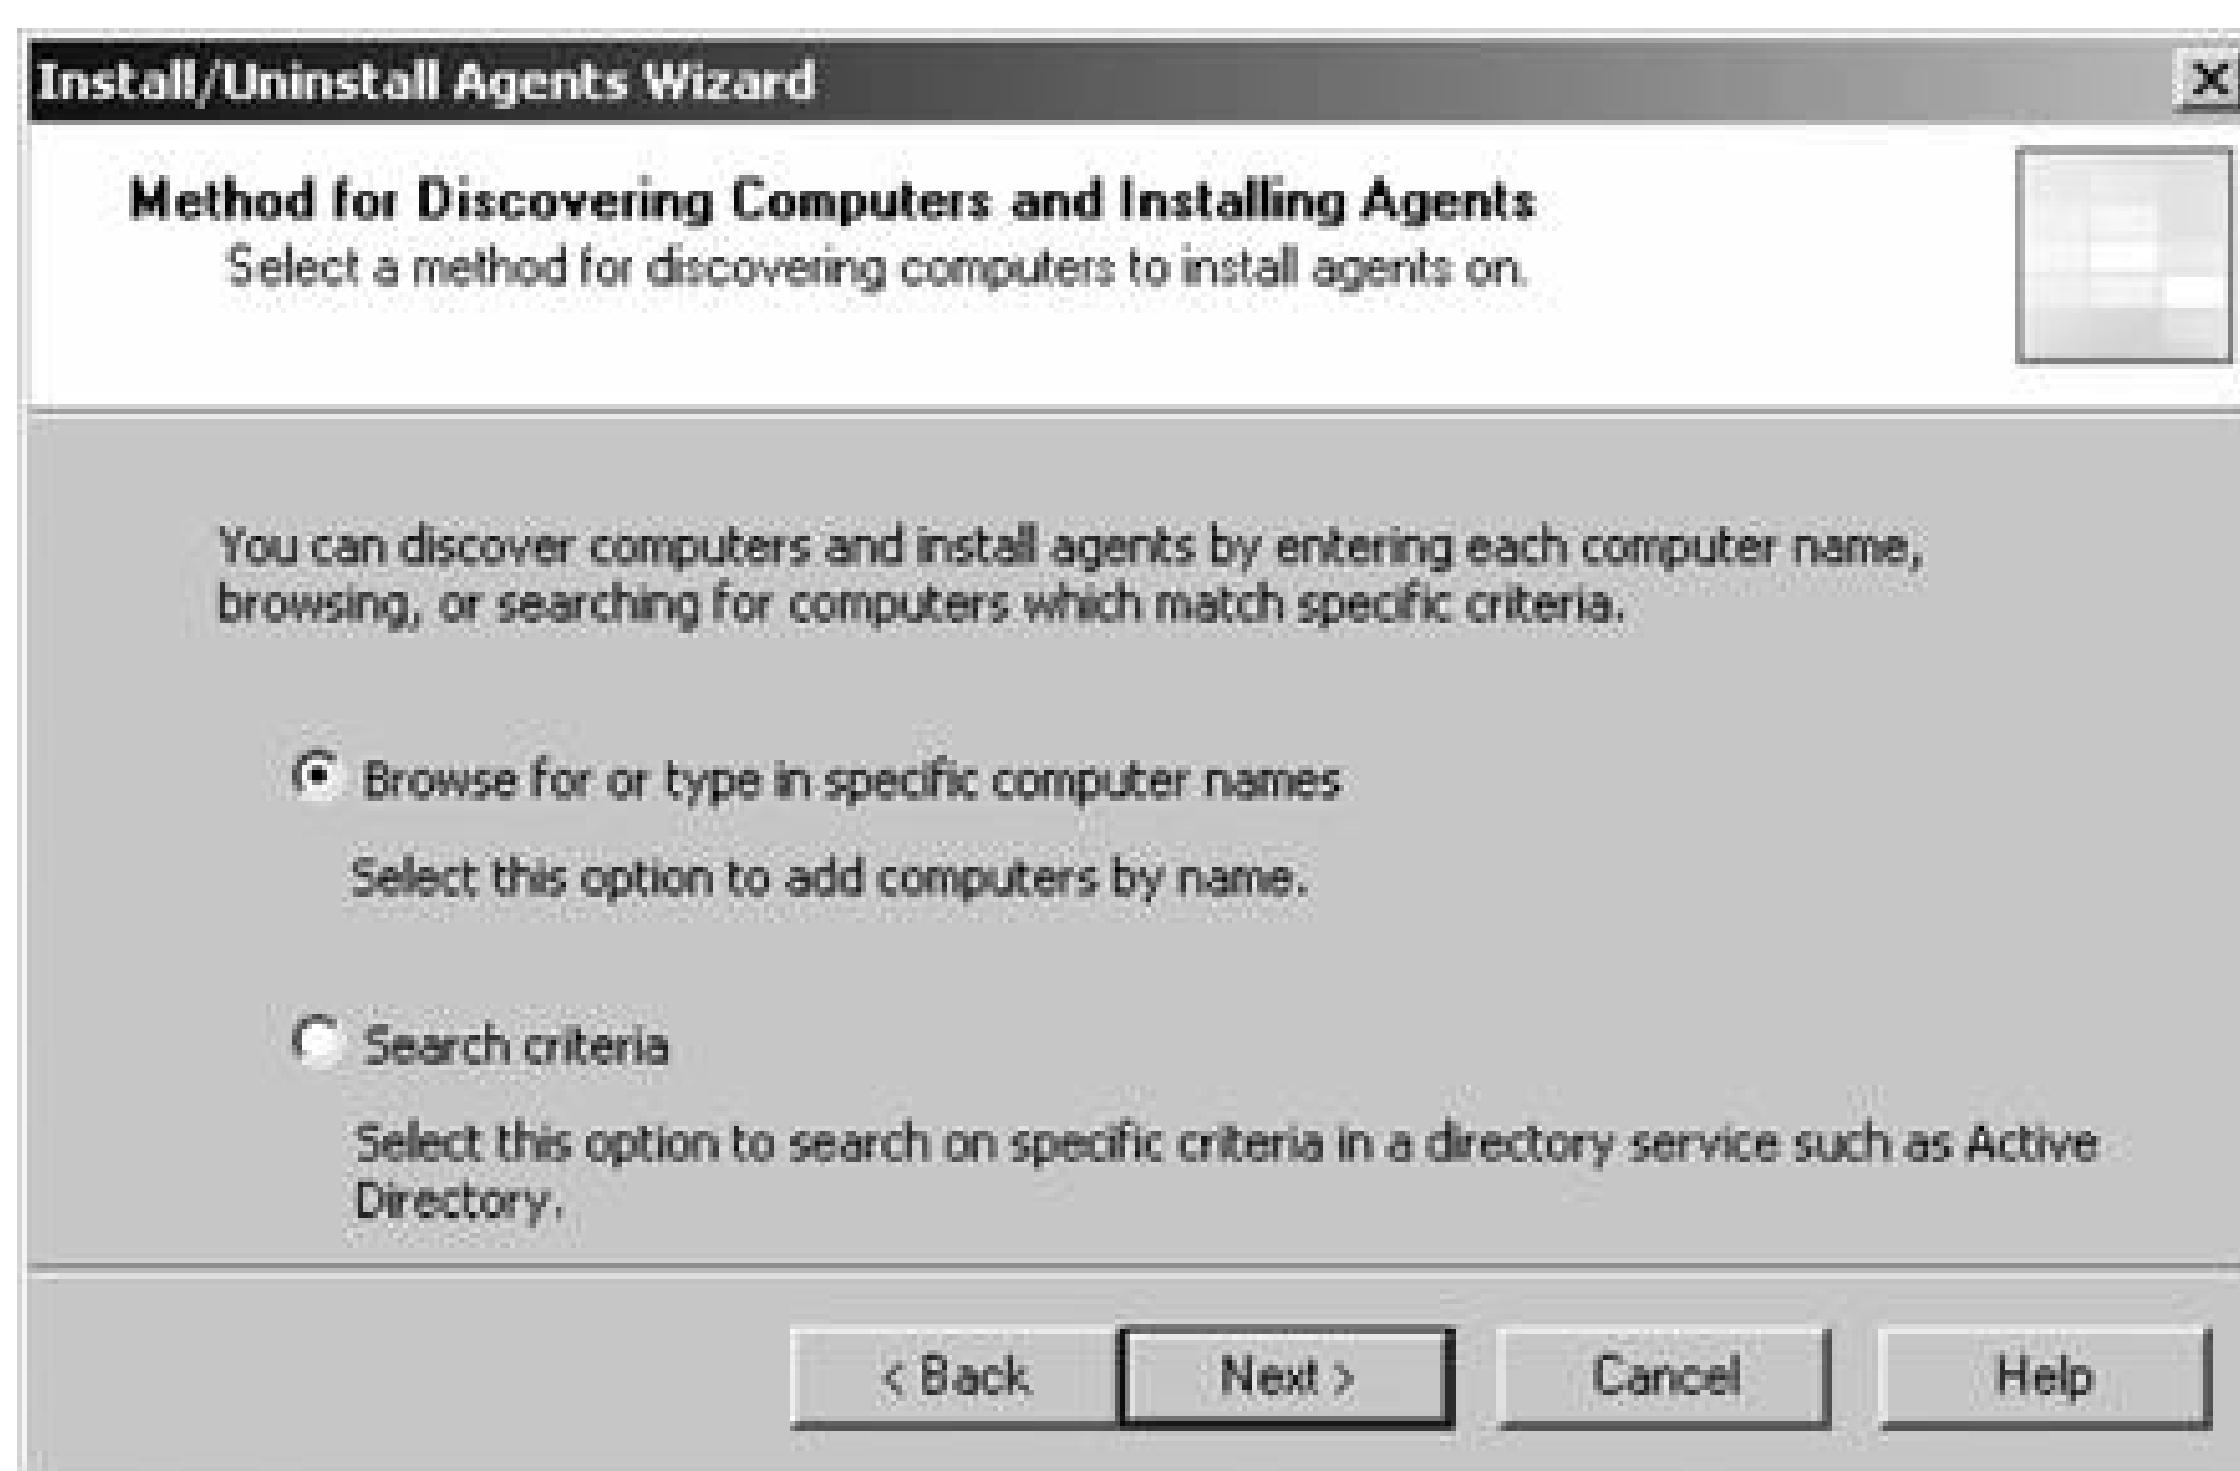
Figure 3-20. Selecting a management server to create the Computer Discovery Rule on



Next, you can browse for a computer or group of computers, or create the search criteria using the same process as the Computer Discovery Rules method (see [Figure 3-21](#)). This is described in the previous section, "[Confirming Agent Functionality](#)."

Figure 3-21. Selecting the method for computer discovery and agent installation





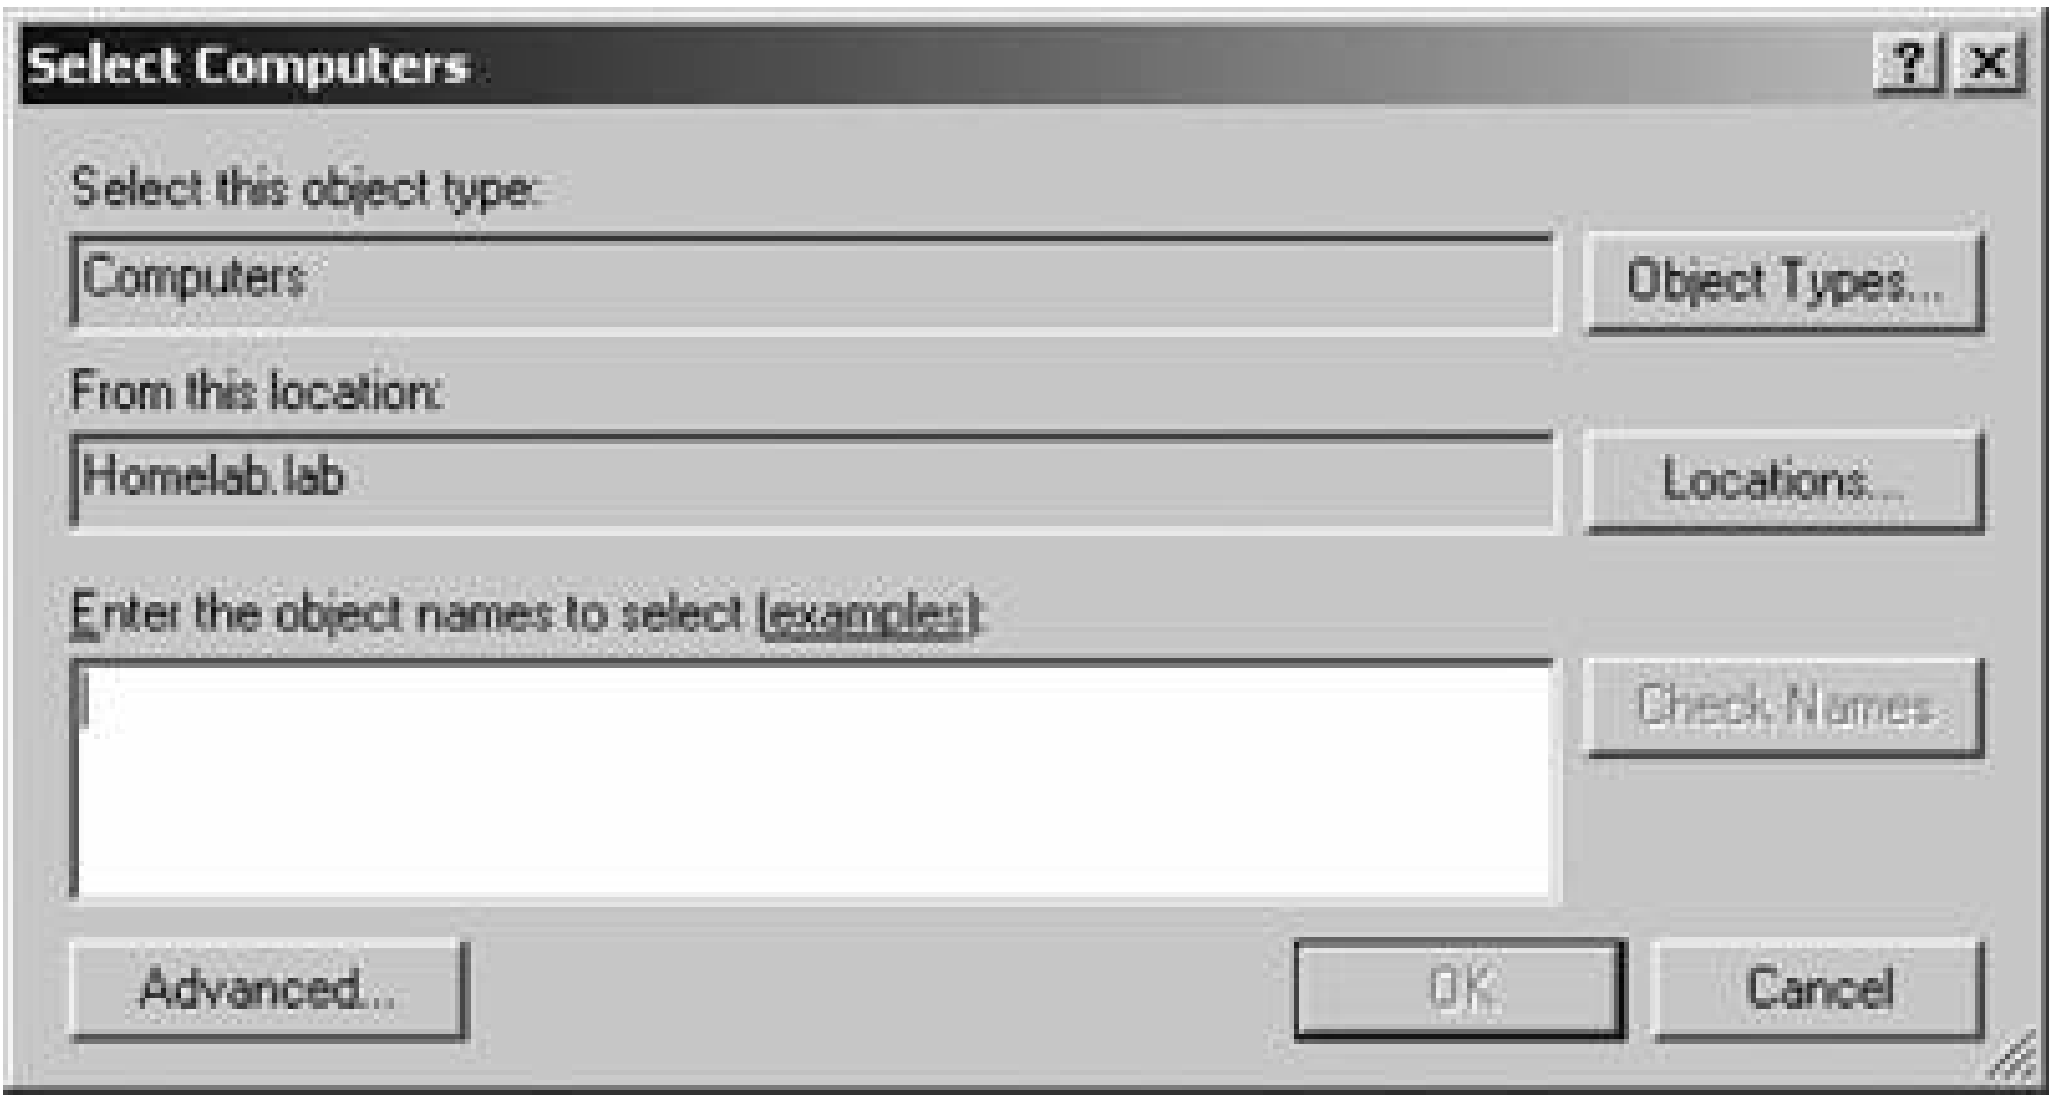
If you select the "Browse for or type in specific computer names" option here, you can enter the specific computer names in either FQDN or domain\NetBIOS name format, as shown in [Figure 3-22](#).

Figure 3-22. Enter the specific computer names for the Computer Discovery Rule to search for

The browse button presents you with the familiar Select Computers (object picker) query configuration box that you use to find objects in the Active Directory Users and Computers tool, as shown in [Figure 3-23](#).

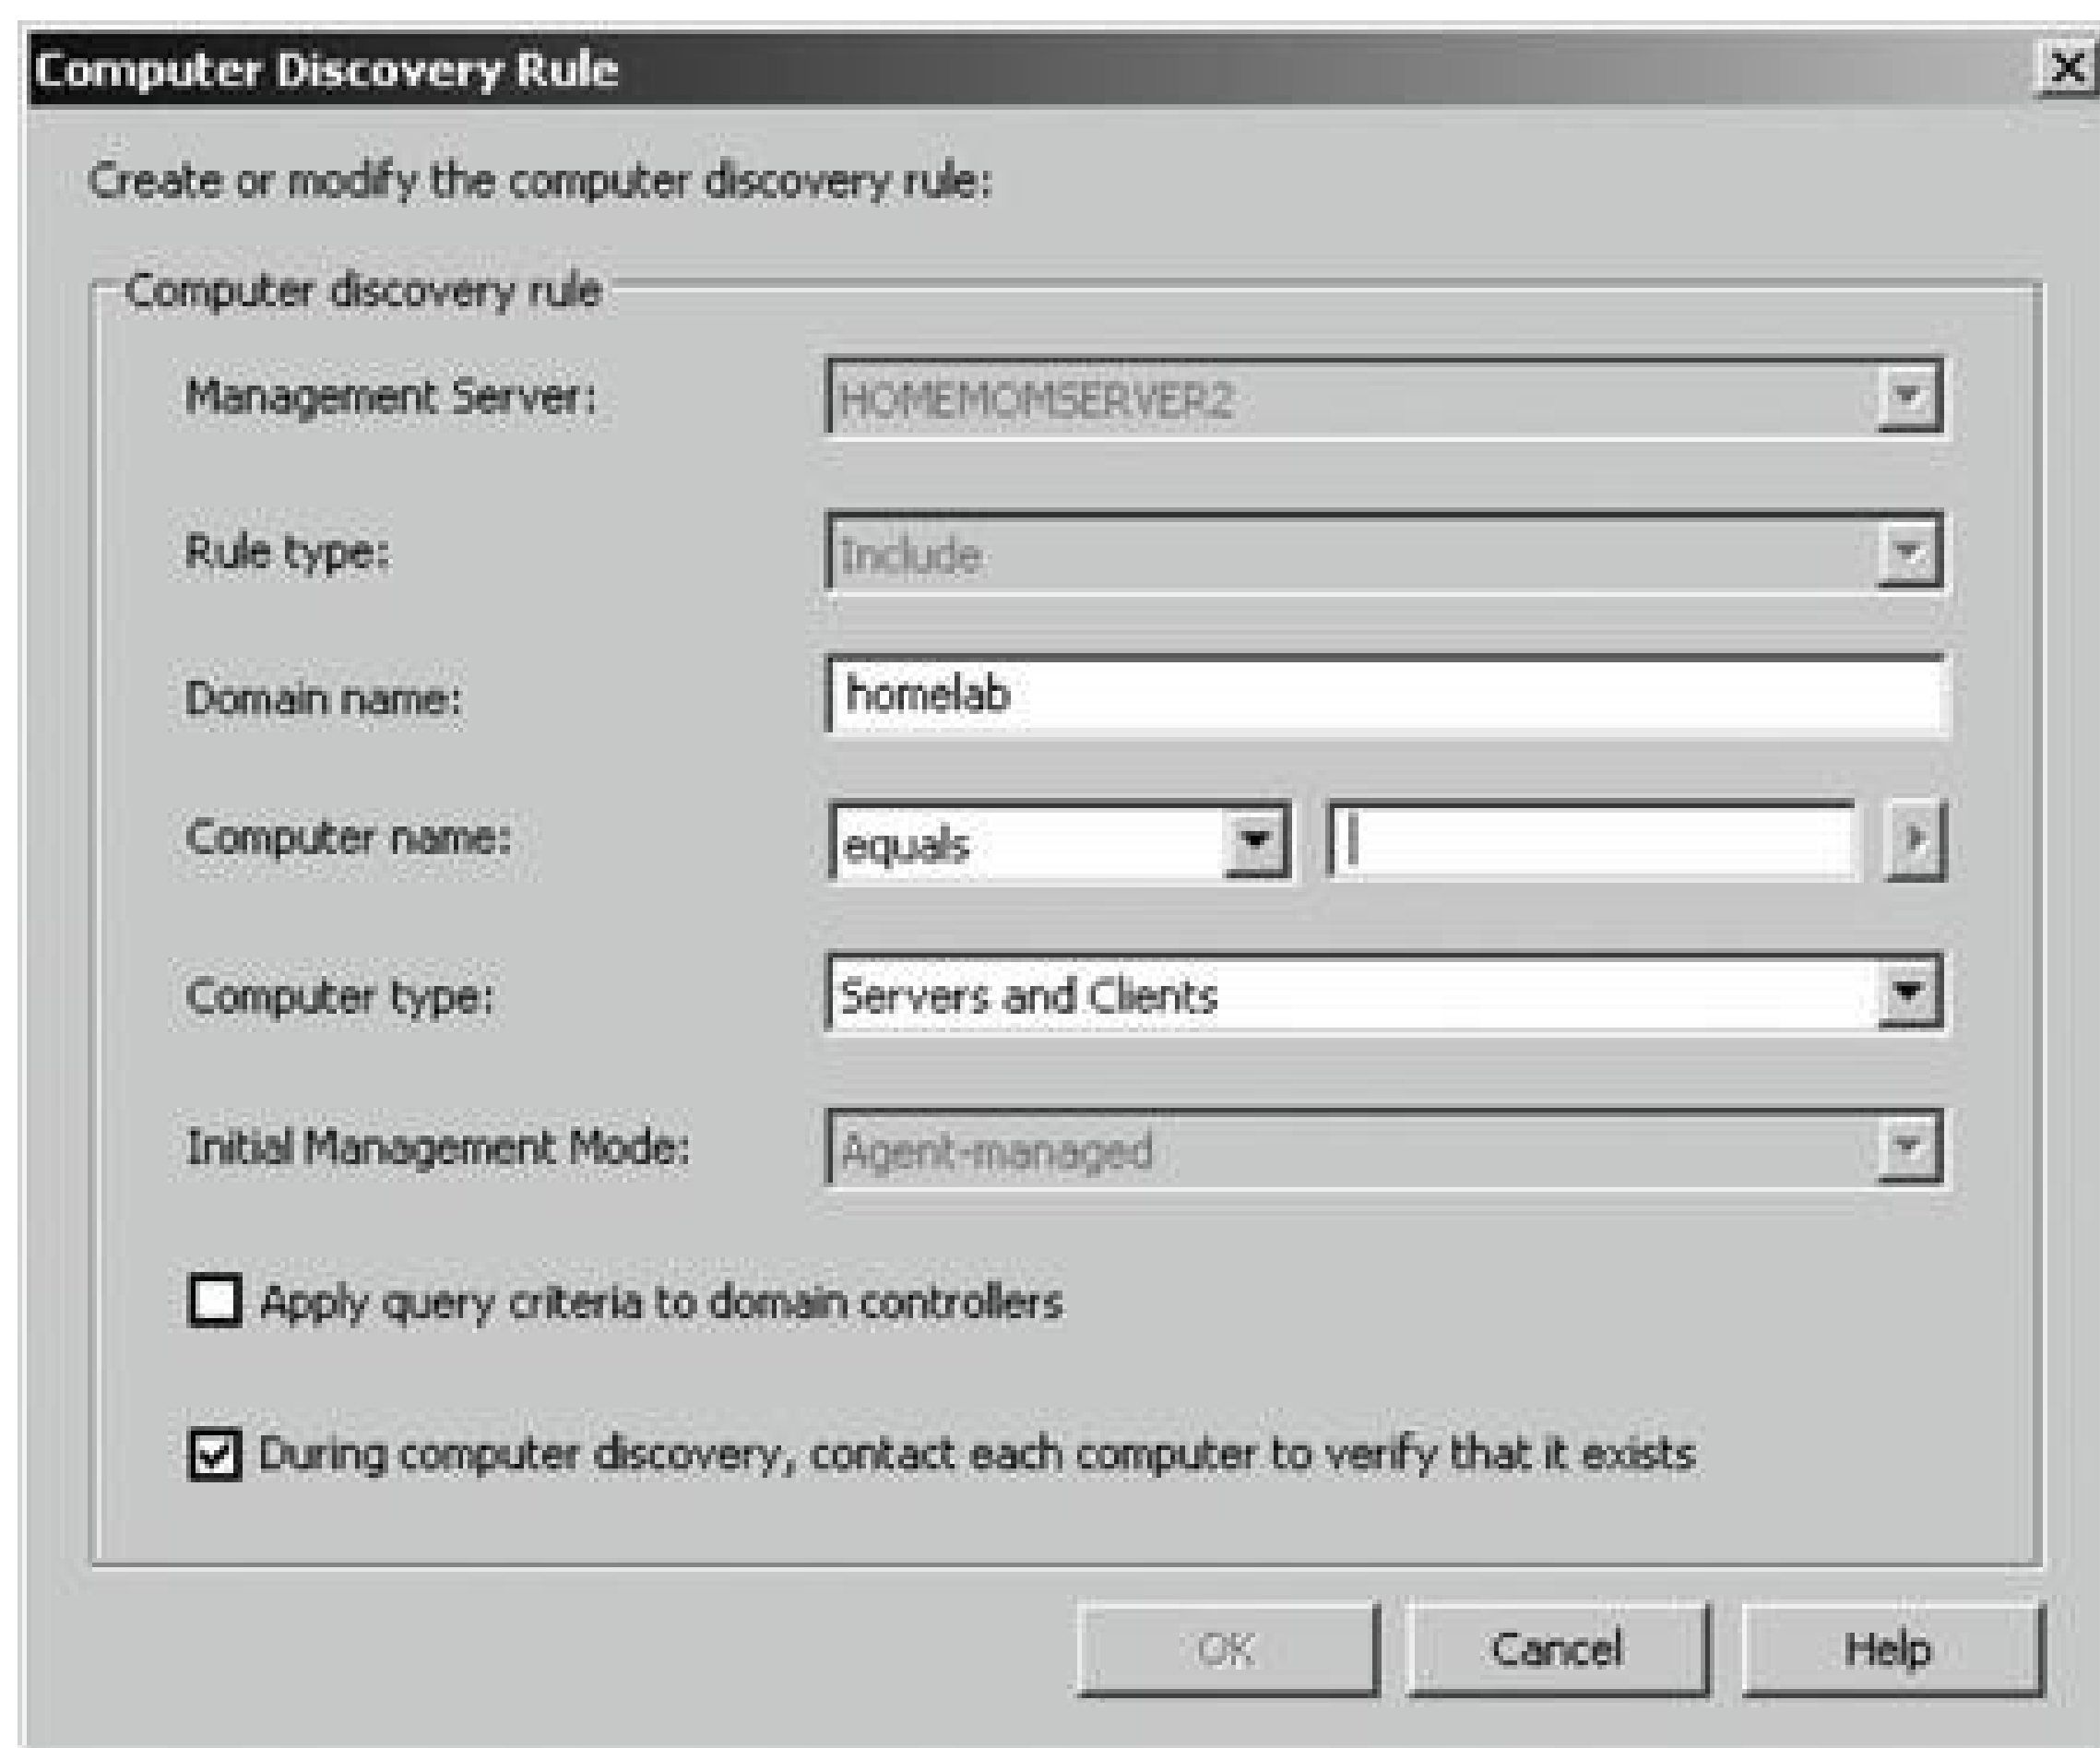
Figure 3-23. The Select Computers (object picker) query configuration

tool



If you select the "Search criteria" option in the Method for Discovering Computers and Installing Agents page (see [Figure 3-21](#)), you are presented with the same Computer Discovery Rule dialog box as discussed in the previous section. However, your options for rule configuration are not available, namely the management server, the rule type, and the initial management mode. By default, the wizard also selects to verify the existence of the target computers ([Figure 3-24](#)). If you need to control the rule type or management mode options, you should not use the install/uninstall agents wizard.

Figure 3-24. Using the agent install/uninstall wizard limits your options for configuring discovery rules



Once you have identified the target computers through one of these methods, the install/uninstall agents wizard behaves identically to the agent install wizard. It prompts you for the credentials to install the agent, for the agent action account, and the directory to install the agent. When you finish the wizard, the same task progress feedback appears.

The install/uninstall agents wizard then performs a limited computer discovery cycle that only looks for computers according to the rule/filter that was created in the wizard and executes the install/uninstall action against it. This is unlike the daily full discovery cycle, which will search the network and apply all discovery rules that have been configured, execute the actions specified in those rules, execute any actions that have been approved for computers in the Pending Actions container, and place discovered computer names into the operations database.

The difference between the previous method and this one is that in this method, the action will be executed regardless of the automatic agent installation setting. The target computers are never placed into the Pending Actions container, so there is no opportunity to reconsider or to check your work once you click Finish in the wizard. Although this method of installing agents may seem easier, the functionality that is lost is not offset by the ease of use. I don't used this wizard because it adds little value in my eyes.

At this point, if you used the wizard to perform an agent install, you should go through the steps to confirm agent functionality.

### 3.3.8. Discovering and Installing Agents via ManualMC.txt



To use this method of installing agents, simply create a text file named *Manual/MC.txt* on the management server that will own the agents and populate it with the names of the desired target computers. You can use FQDN, NetBIOS name, or domain\NetBIOS format. Here is an example of the contents of a *Manual/MC.txt* file using the FQDN format:

```
avalon.homelab.lab
homesrv02.homelab.lab
homee2k3server.homelab.lab
homememberserver.homelab.lab
```

Note that there is only one entry per line and that there are no spaces between the lines. Also, for the sake of clarity, you should use a single format for the computer name for all entries. If you need to add computers that are not members of a domain, use the workgroup\NetBIOS name format. Place this file in the *I%SystemRoot%\Program Files\Microsoft Operations Manager 2005* or the directory you installed MOM into.

When the next full discovery cycle runs, the management server will either place these computers into the Pending Actions container for further action or execute the installation based on the automatic agent installation settings. To uninstall an agent that was installed using this method, simply delete its entry from the *Manual/MC.txt* file and either wait for the daily discovery cycle to run or initiate a full discovery cycle yourself. Remember, if you want the management server to perform the install/uninstall tasks automatically, then the management server action account must have full administrative rights on the target computers.

This text file method for discovering and installing agents is the easiest to configure but also gives you the fewest options and no immediate feedback except by monitoring the Administrator and Operator consoles for the expected changes. With this method, as with the wizard, you can only identify computers that will be agent managed, and you cannot use the *Manual/MC.txt* file to exclude a computer from being managed by a management server. Using this method overrides all other Include rules but not Exclude rules for that management server. This method is error-prone because it must be manually maintained on a daily basis.

### 3.3.9. Miscellaneous Agent Management Tasks

Uninstalling the agent is a very simple task. Right-click the managed computer in the Agent-managed Computers container in the Administrator console, select Uninstall Agent, and provide the appropriate credentials. With this method there isn't a convenient toolbar. Forcing the computer into unmanaged mode is the same except there is no prompt for credentials.

You'll need to update an agent to change things like the storage space allocation on the agent, the communication port, the "Collect event binary data" settings, and the agent control level if it is set to None. Most significantly, if you need to change the agent action account or if the mutual authentication setting is changed, these modifications will need to be pushed down to the agent by the Update Agent Settings feature in the Administrator console. You can access this feature in the Agent-managed Computers container by right-clicking on the managed computer and selecting Update Agent Settings. When invoked, this prompts for the necessary administrative credentials to update the agent settings, as well as the account to be used as the agent action account, which you

do not need to provide.

Pushing agent updates to managed computers is a different process than pushing management pack updates to managed computers. For example, if you make a change to a threshold where a processing rule generates an alert, or if you disable a rule, these changes are pushed down automatically once a minute. Updating agent settings is a manual process and is only needed when reconfiguring the agent.



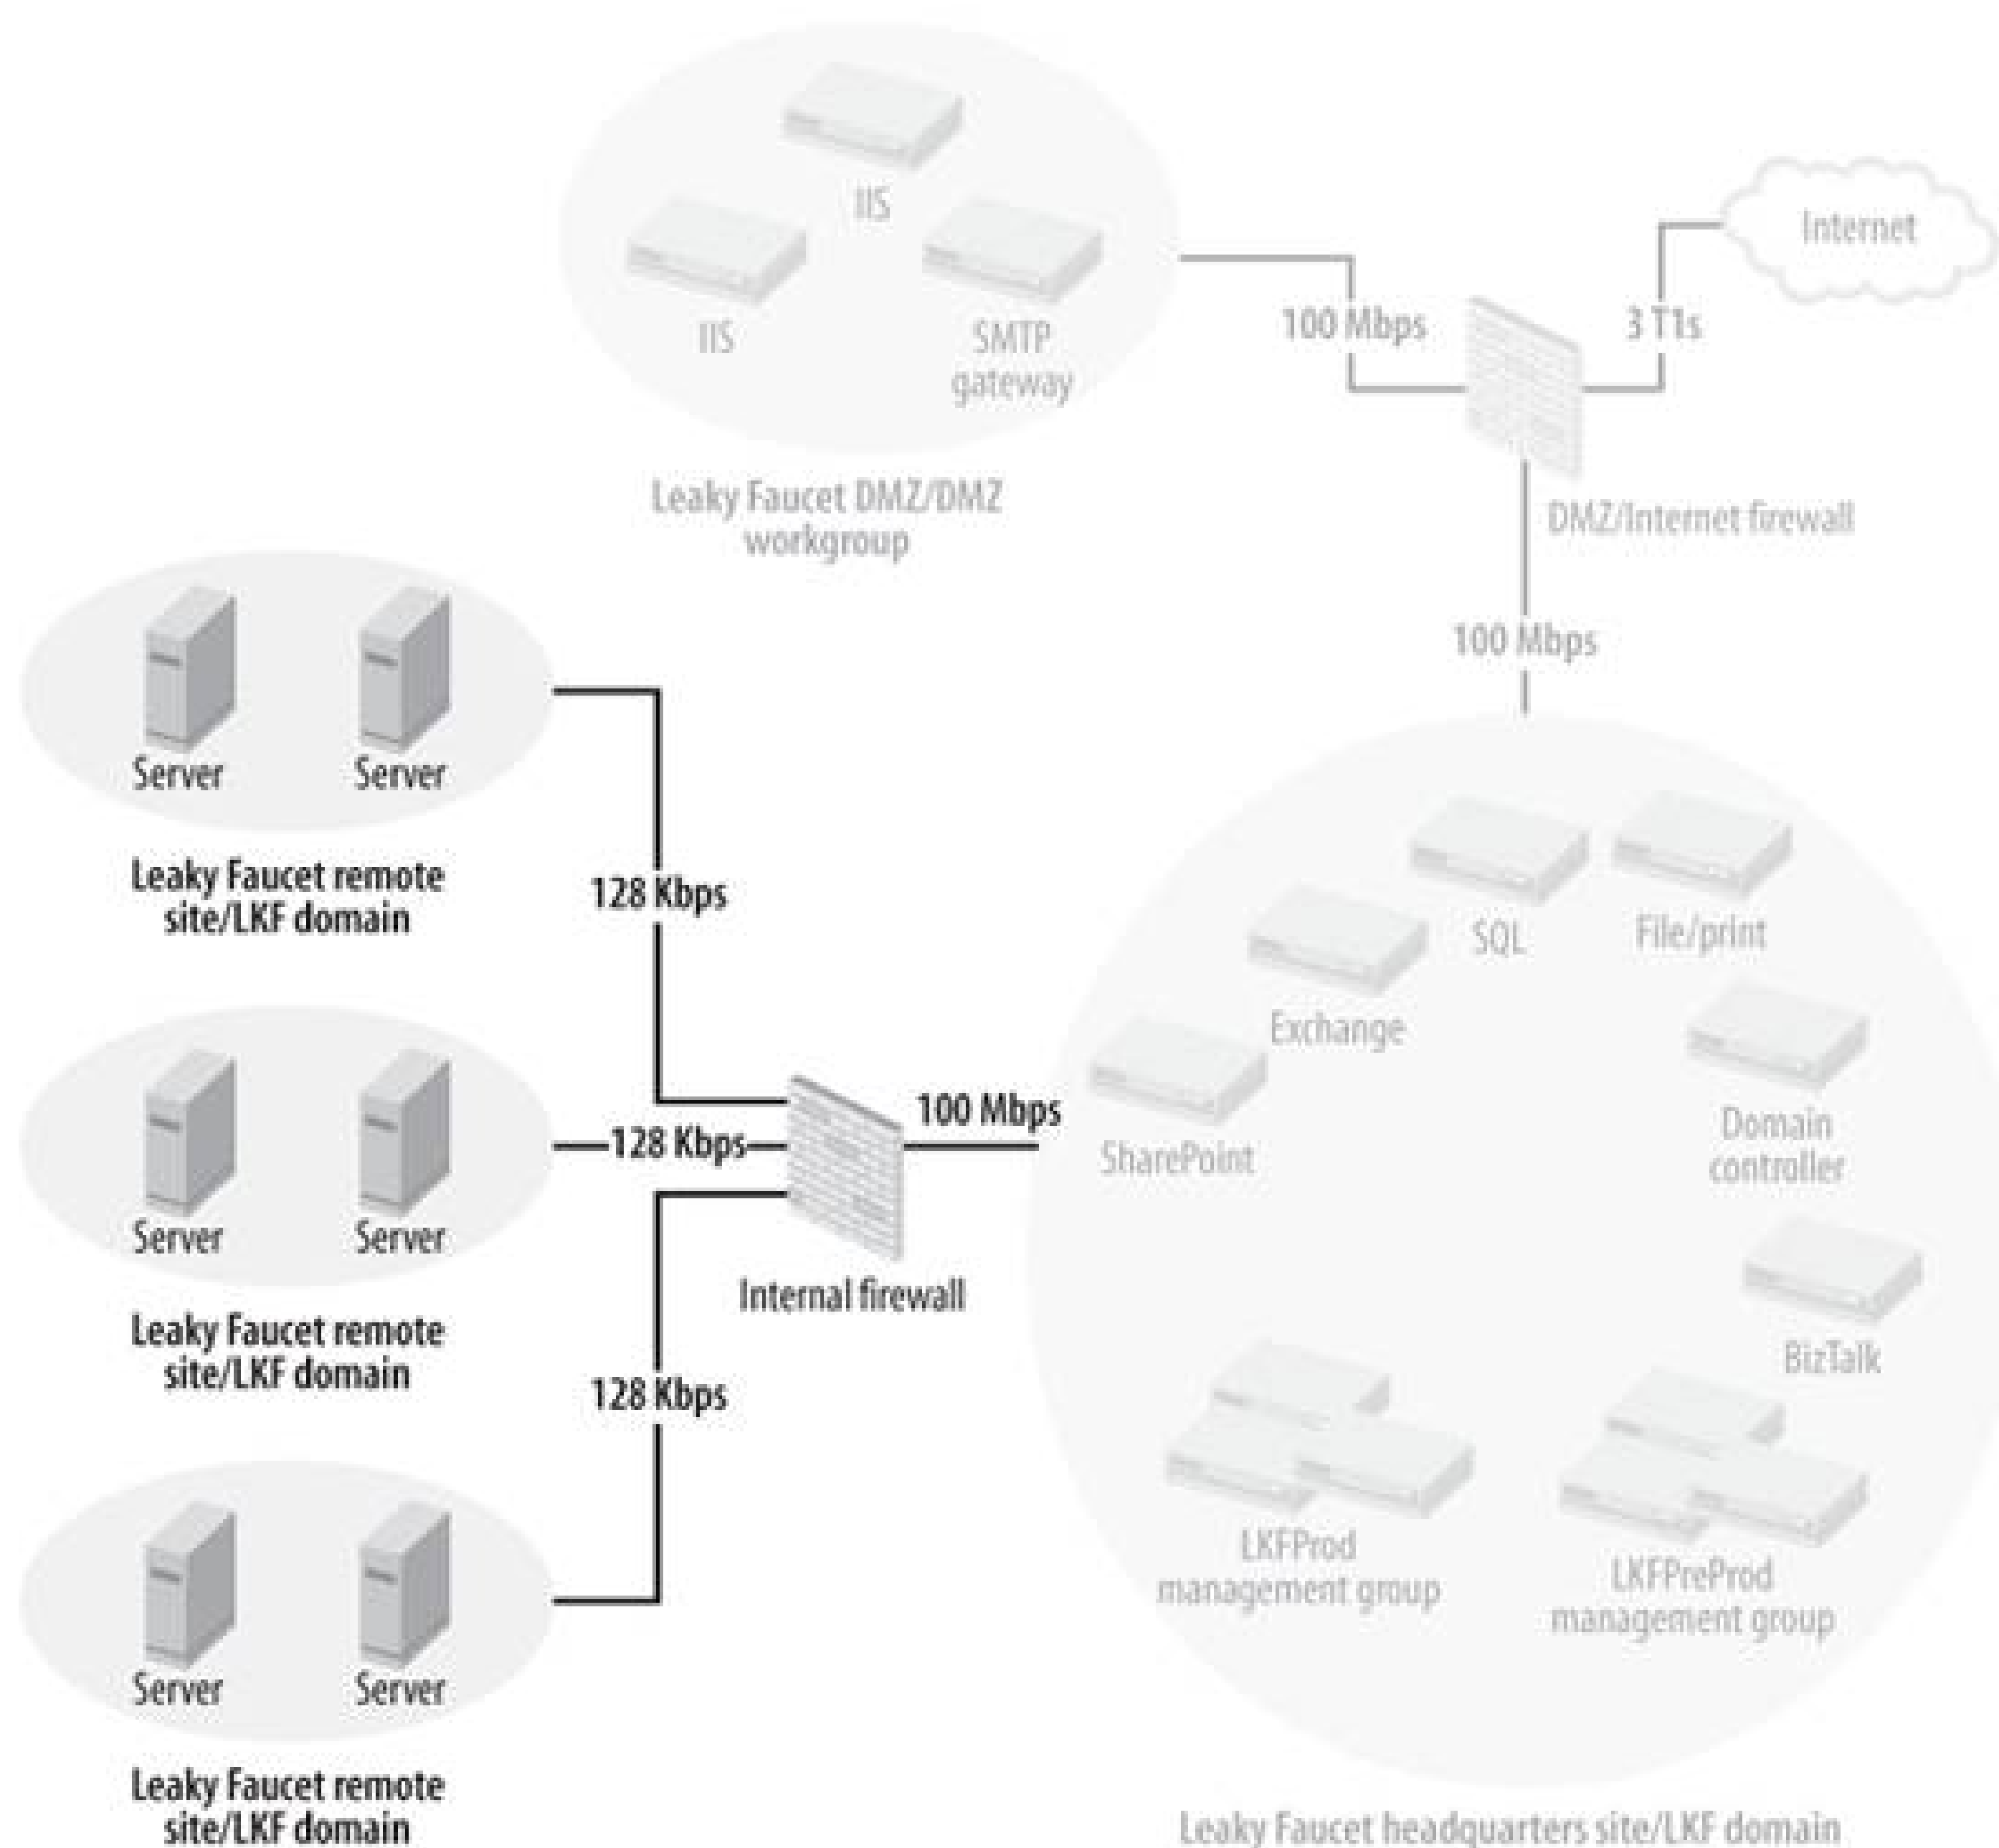
## 3.4. Deploying and Managing Agents Across a Firewall and a Slow WAN Link

The major considerations for deploying and managing servers in this type of environment are the presence of a firewall between the target computers and the management servers, and the limited bandwidth. Of the two, the limited bandwidth is more likely to cause problems. If the environment outside the firewall is not in the same AD domain, then there are additional considerations. All of the servers at the Leaky Faucet remote sites are members of the LKF domain, so they don't have to change the underlying security structure. [Figure 3-25](#) shows the area of interest in the Leaky Faucet network.

For operational data, heartbeats, and management pack update communication, MOM agents communicate with their management server over TCP port 1270. It is easy to keep this port is open, so the presence of the firewall poses no problem. Agent-to-management-server communication resists being passed across a proxying firewall. If you try to, the management server will not recognize the agent, because the network traffic is coming from the proxying firewall, not the server on which the agent is installed. The workaround is to turn off proxying for port 1270 on the intervening firewall and pass the port 1270 traffic straight through using stateful inspection. However, agent installation is a different matter. For a management server to install an agent remotely and update the agent configuration, a variety of protocols and ports are used; for example, TCP 135 (RPC End Point Mapper) and the RPC port range must be open, as well as UDP 445, DCOM, and SMB. Because these ports are usually closed on a firewall, you cannot remotely install agents from the management servers. You must manually install and update the agents on the target machines.

Figure 3-25. In this environment the target computers are separated from the management servers by slow WAN links and a firewall





The volume of traffic over port 1270 can be estimated with the MOM 2005 Sizer.xls tool. This tool is available at <http://www.microsoft.com/downloads/details.aspx?FamilyId=93930640-FA0F-48B3-8EB0-86836A1808DF&displaylang=en>. You can also use this tool for estimating operations database sizes, but you may get mixed results. It is a very straightforward tool, consisting of a single-page spreadsheet on which the computer counts or other appropriate information is entered in the yellow highlighted cells. Leaky Faucet's servers are split about evenly between the 12 remote sites and headquarters. The amount of traffic that can be expected cumulatively from the 24 remote site servers is only about 3,000 bps or 2.5 percent of a 128-Kbps line ([Figure 3-26](#)).

Figure 3-26. Estimation of the network traffic for 24 servers

AGENT TO MANAGEMENT SERVER NETWORK SIZING UTILITY		
ENTER MANAGED COMPUTER COUNT ACROSS SMALL LAN		
	24	
	3,274.49	BITS /sec BITSTREAM
NETWORK SIZE = 128 Kbits/sec, 16 KBytes/sec	2.56%	NETWORK UTILIZATION
NETWORK SIZE = 64 Kbits/sec, 8 KBytes/sec	5.12%	NETWORK UTILIZATION
NETWORK SIZE = 56 Kbits/sec, 7 KBytes/sec	5.85%	NETWORK UTILIZATION
NETWORK SIZE = 32 Kbits/sec, 4 KBytes/sec	10.23%	NETWORK UTILIZATION

This number is a rough estimate, since the tool performs its calculations based on a network connection without intervening port restrictions, so it is including configuration update data as well. Even so, the number is useful because it gives an order-of-magnitude estimate for the traffic that can be expected. In Leaky Faucet's case, there are only one or two servers at the end of each 128-Kbps line, so the actual volume of traffic will be less than the estimate.

Certain steps can be taken to accommodate successful operations across the slow link. You can reduce the amount of data that is sent, increase the communications interval so that data transmission occurs less frequently, and increase the buffer size on the agents to hold the extra data between those communication sessions. See the ["Accommodating Low Bandwidth"](#) section later in this chapter.

### 3.4.1. Manual Agent Deployment

Before you manually install the agent, you have to identify the primary management server for the agents, then change the "Reject new manual agent installations" setting. Remember that once you override this at the management-server level, you must right-click the Management Packs node, select Commit Configuration Change, and restart the MOM service on each management server in the management group. For this example, *homemomserver* is the designated management server.

To start the manual agent installation, log on to the target machine locally with administrative rights and either launch *setup.exe* from the root of the MOM 2005 install media and select the Manual Agent Install tab, or locate the *momagent.msi* file in the i386 directory and launch that. Clicking through the Welcome panel brings you to the select install location page [Figure 3-27](#)).

Accept the default location, as was planned for, and click Next. If you planned to install the agent to a different location, click on the Change button and enter the path to the desired directory. On the next page of the setup process, shown in [Figure 3-28](#), you are prompted for all the necessary parameters required to install the agent.

Enter the appropriate information in the management Group Name, Management Server, and Management Server Port fields. I recommend that you accept the default

Figure 3-27. Manual agent install destination folder selection



Figure 3-28. Manual Agent Installation Configuration page



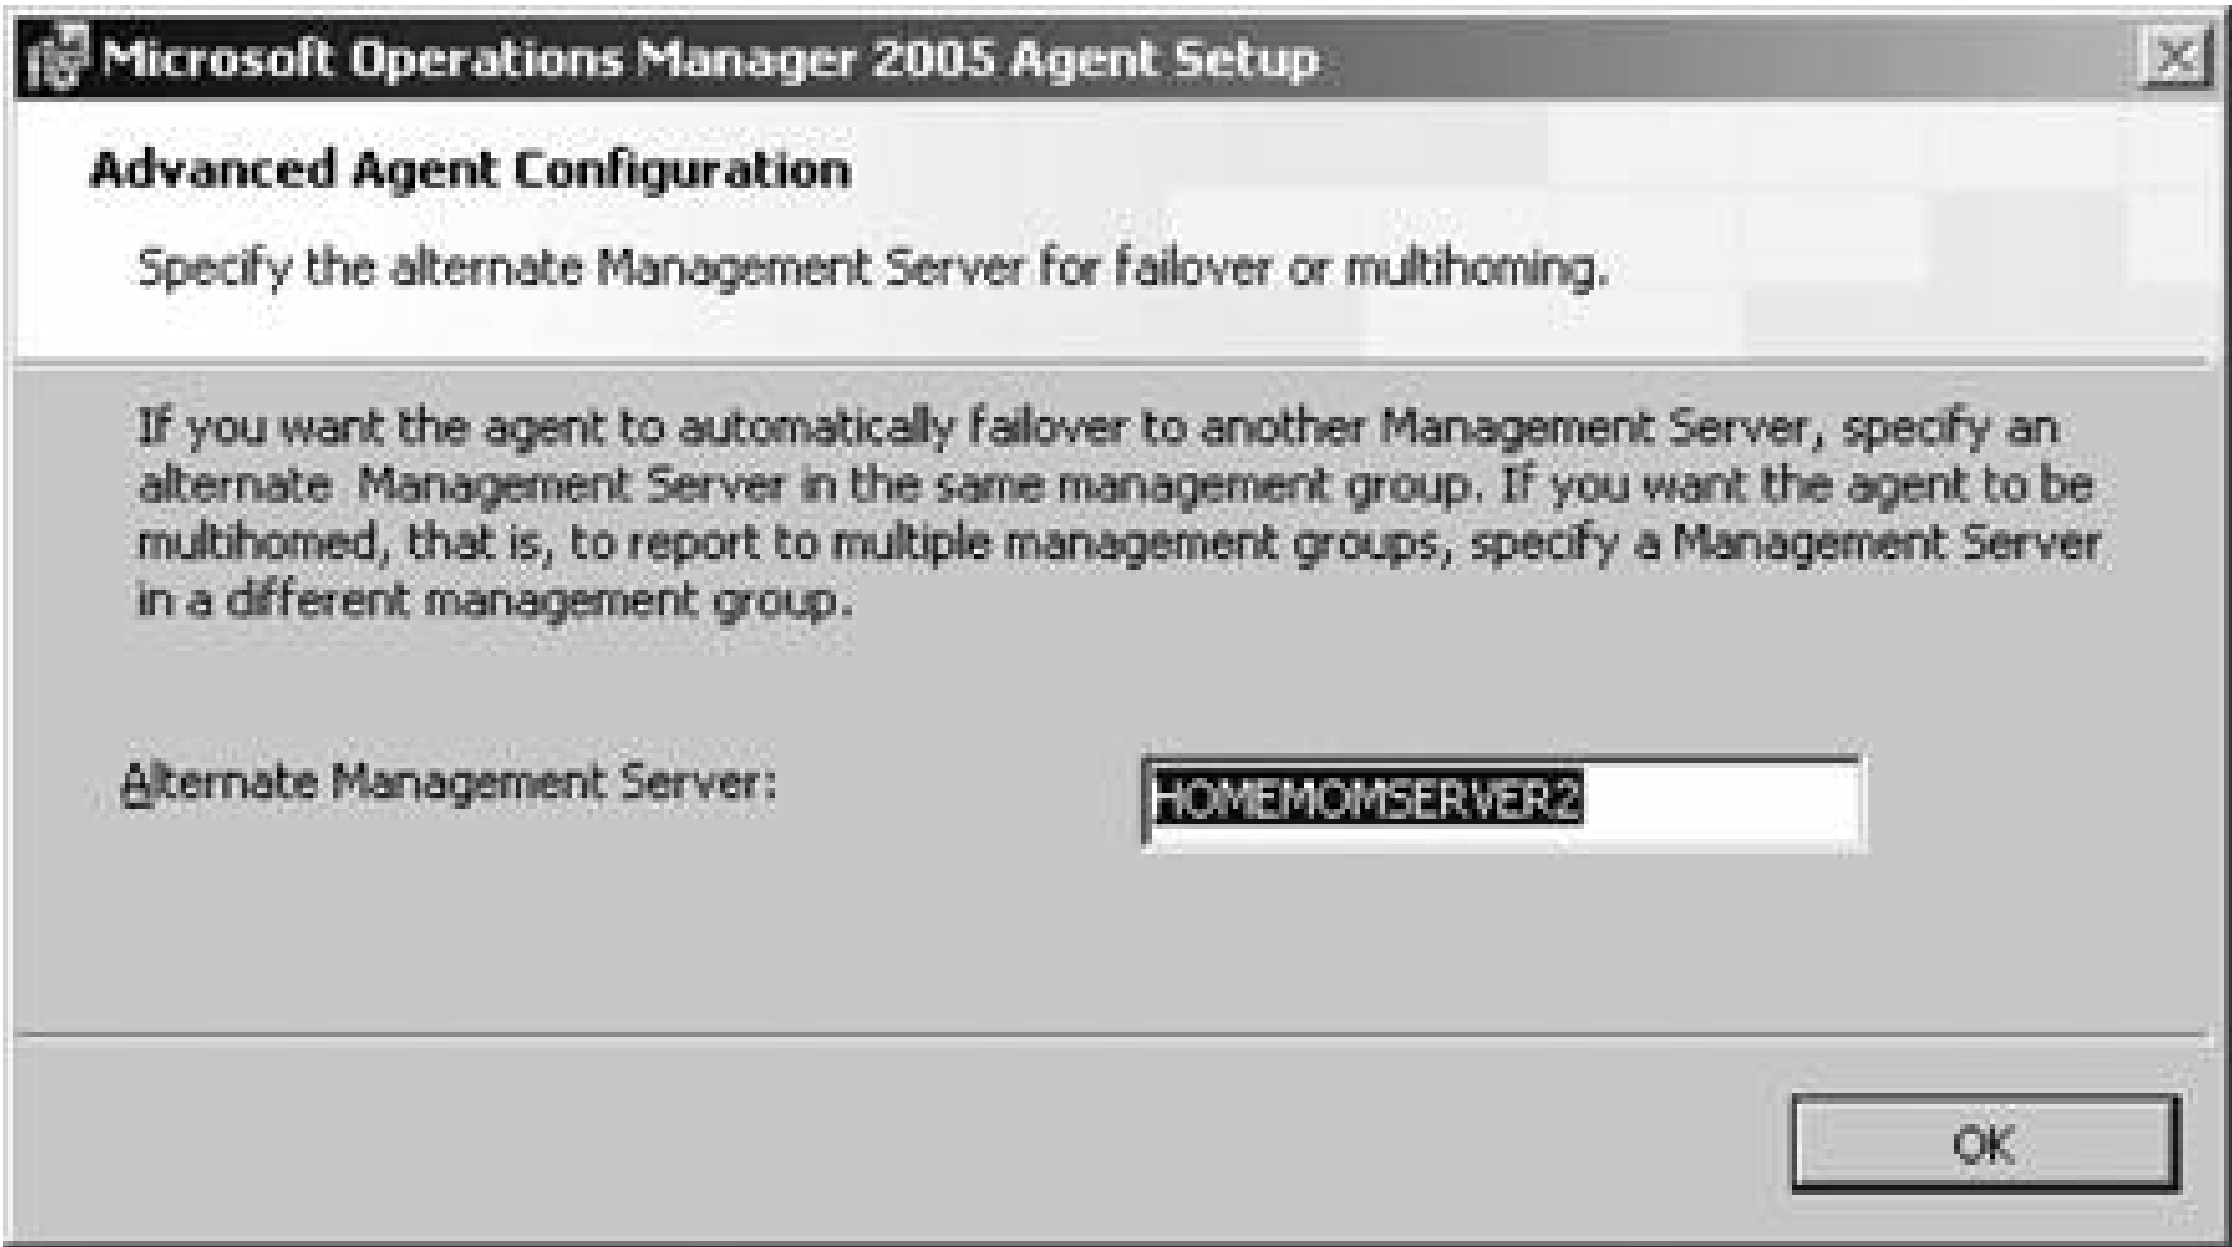
port (1270). If you decide to use another port, you will have to configure this on the management servers as well. This is addressed in [Chapter 5](#), but briefly: navigate to the Administration node in the Administrator console → Global Settings node → Security object on the Communications tab and change the port setting there. You will have to make the change on the management server before a manually installed agent can communicate with it. To allow agent failover, you have to change the communication port on all of the management servers.

Unlike the remote agent installation from a management server process, during a manual install you can configure the failover server. To do this, click on the Advanced button and enter the desired management server name ([Figure 3-29](#)). If the agent is being multi-homed, this is where you can



enter the management server name from a different management group.

Figure 3-29. Configuring the failover server for a manually installed agent



### 3.4.2. Agent Failover

In management groups with multiple management servers, you can configure which management servers agents will failover if communication is lost with an agent's primary management server. This is done in the properties of the management server, on the Failover tab. Here you can choose to allow or disallow the agents owned by this management server to failover to other management servers in the management group. The Failover tab also indicates how many agents are currently assigned to the other management server, as shown in [Figure 3-30](#). If you don't want agents to failover to a specific management server, simply clear the checkbox. Remember that agent failover is configured automatically in multi-management server management groups; you do not need to take any additional steps to enable it, only to modify it.

Next, you must set the level of control that the management server will have on this agent. Agents that are fully controlled will have all updates, patches, and configuration changes pushed to them by the management server and can be uninstalled by the management server. Basically, with full control, all management functions can be performed remotely by the management server. Full control requires all remote procedure call (RPC) ports to be available, as well as a number of others. Typical firewall configurations do not permit RPC ports to be opened. Agents whose control level is set to None must have patches, upgrades, and configuration changes manually applied, and must also be uninstalled manually. In Leaky Faucet's case, because the firewall is blocking the ports used for these management tasks, the Agent Control Level is set to None, as shown in [Figure 3-28](#).

Figure 3-30. Configuring agent Failover



Manual install prompts you for the agent action account and as before, you can choose to use the local system account or some other domain or local account. The next panel prompts you for the AD configuration ([Figure 3-31](#)). This is necessary so mutual authentication can be configured correctly or the agent. Leaky Faucet is running MOM 2005 in AD and mutual authentication is required between the agents and the management servers.

Figure 3-31. This setting tells the manually installed agent if it is in an AD domain or not



The installation wizard will next present you with the familiar Ready to Install summary page where you can review the configuration choices for accuracy and then start the installation.

After the installation is complete, MOM places the computer into the Pending Actions container in the Administrator console, with an "Approve manual agent installation" status in the Pending Actions column (see [Figure 3-32](#)).

Figure 3-32. Manually installed agents are placed into the Pending Actions container

To allow the management server to start managing this agent, right-click the computer and select to approve the install. The computer will then be placed into the Agent-managed Computers container and the process is complete.

If the agent does not appear in the Pending Actions container, after the manual agent installation process, you probably did not correctly configure the management server to accept new manually installed agents. Don't panic you don't have to uninstall and reinstall the agent. Simply confirm your override settings on the management server, commit the configuration change in the Administrator console, and restart the MOMService on all management servers in the management group. The agent should then appear in the Pending Actions container. When you are done with manual agent installation, re-enable the Reject New Manually Installed Agents configuration.



### 3.4.3. Accommodating Low Bandwidth

One configuration step to minimize the bandwidth used by the agent/management server is to set the agent control level to None. By doing this, the traffic associated with automatic agent configuration updates (not processing rule updates) is eliminated. Other steps can include modifying the heartbeat interval and the request configuration interval on the Global Settings Agents Agent Heartbeat tab. This is practical only if you have a management group that is dedicated to agents across slow links. Otherwise, you are forcing all agents to behave as if they were across a slow link even if they are not. Even if you increase the request configuration interval, I don't recommend increasing the heartbeat interval. Heartbeat communications are very small, so you won't recover much bandwidth by reducing them but you will sacrifice a good deal of functionality.

On the Properties tab for individual agents (Agent-managed Computers) you can override the global settings and increase the Service Monitoring Status check interval so that it reports less frequently. When a service on the target computer changes state, say from started to stopped, it will take longer for the status change to be reflected in the State view in the Operator console. Also, in the agent properties, you can change the event, performance, and alert buffer settings so they report less frequently.

The most effective way to reduce agent-to-server operations traffic is to keep the applied management packs to a minimum. If it is possible, create separate versions of the required management packs for the low-bandwidth agents and only enable rules that are required. This is especially applicable to rules that collect performance- counter data. See the [Management Pack Tuning](#)" section in [Chapter 4](#).

## 3.5. Deploying and Managing Agents Into an Untrusted Environment

The untrusted environment at Leaky Faucet consists of IIS servers in a workgroup configuration in the DMZ. These servers are separated from the management servers by a firewall, do not share a security account structure with the management server nor do they have an AD trust between them. If mutual authentication is left enabled it will prevent agents from installing, so it must be disabled. However, Leaky Faucet has a high-bandwidth and a reliable connection to the management servers. [Figure 3-33](#) shows this portion of the Leaky Faucet network.

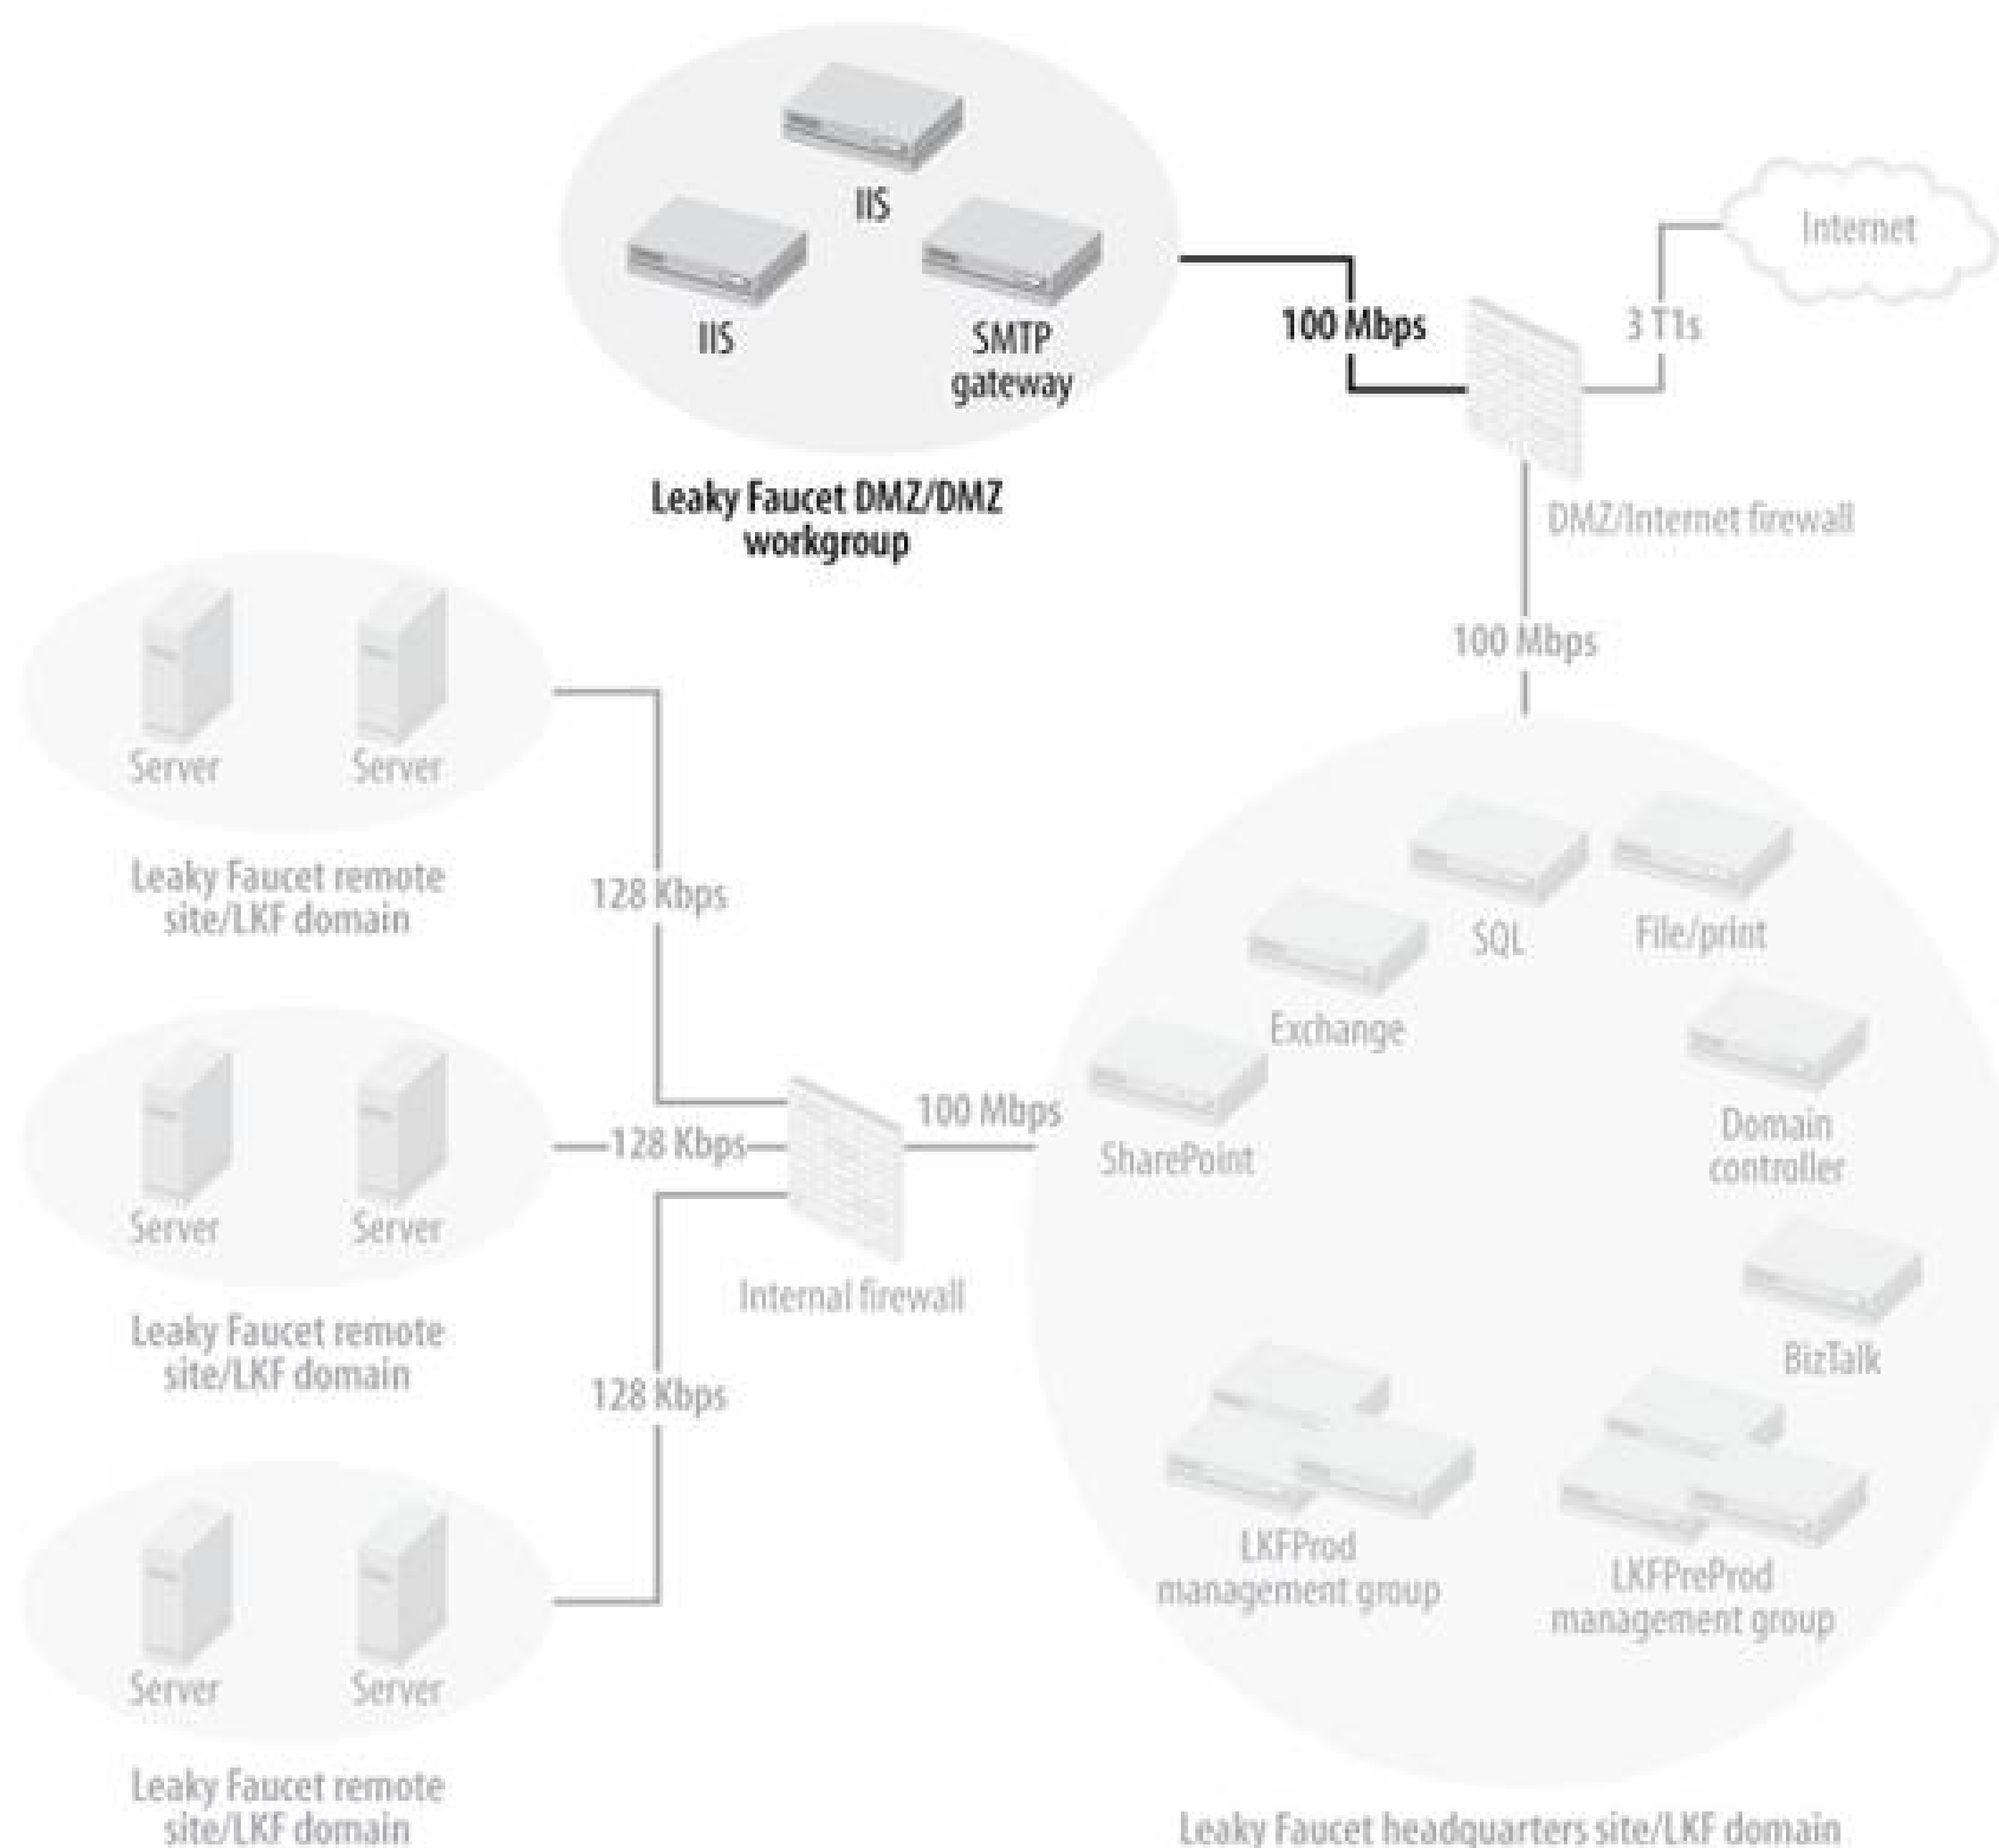
Fortunately, these servers do share name resolution services (WINS) and that means that the management servers can browse to the DMZ servers and vice versa. NetBIOS must be enabled on the management servers and the appropriate ports must be open on the firewall. If there was no firewall between them MOM would support remote installation into a workgroup via the Install/Uninstall Agents Wizard or the Create computer discovery rule methods of agent deployment. During the installation process, you must substitute the *domain name\computer name* with *workgroupname\computer name*. Then provide credentials that have local administrator rights to the workgroup machine for the "install as" credentials. Also, you will have to use the local system account or another local account as the agent action account. This is because there is no trust to the AD domain that the management servers are in.

In this case, because there is an intervening firewall, agents will have to be manually installed and port 1270 opened between the two sets of servers. Also, the agent control level must be set to None

## 3.6. Deploying and Managing Agents from Multiple Management Groups

*Multi-homing* is when an agent reports operational data to and accepts configuration and processing rules from more than one management group. A multi-homed computer only has one agent on it, but the agent interacts with each management group as if that management group were the only one it is reporting to.

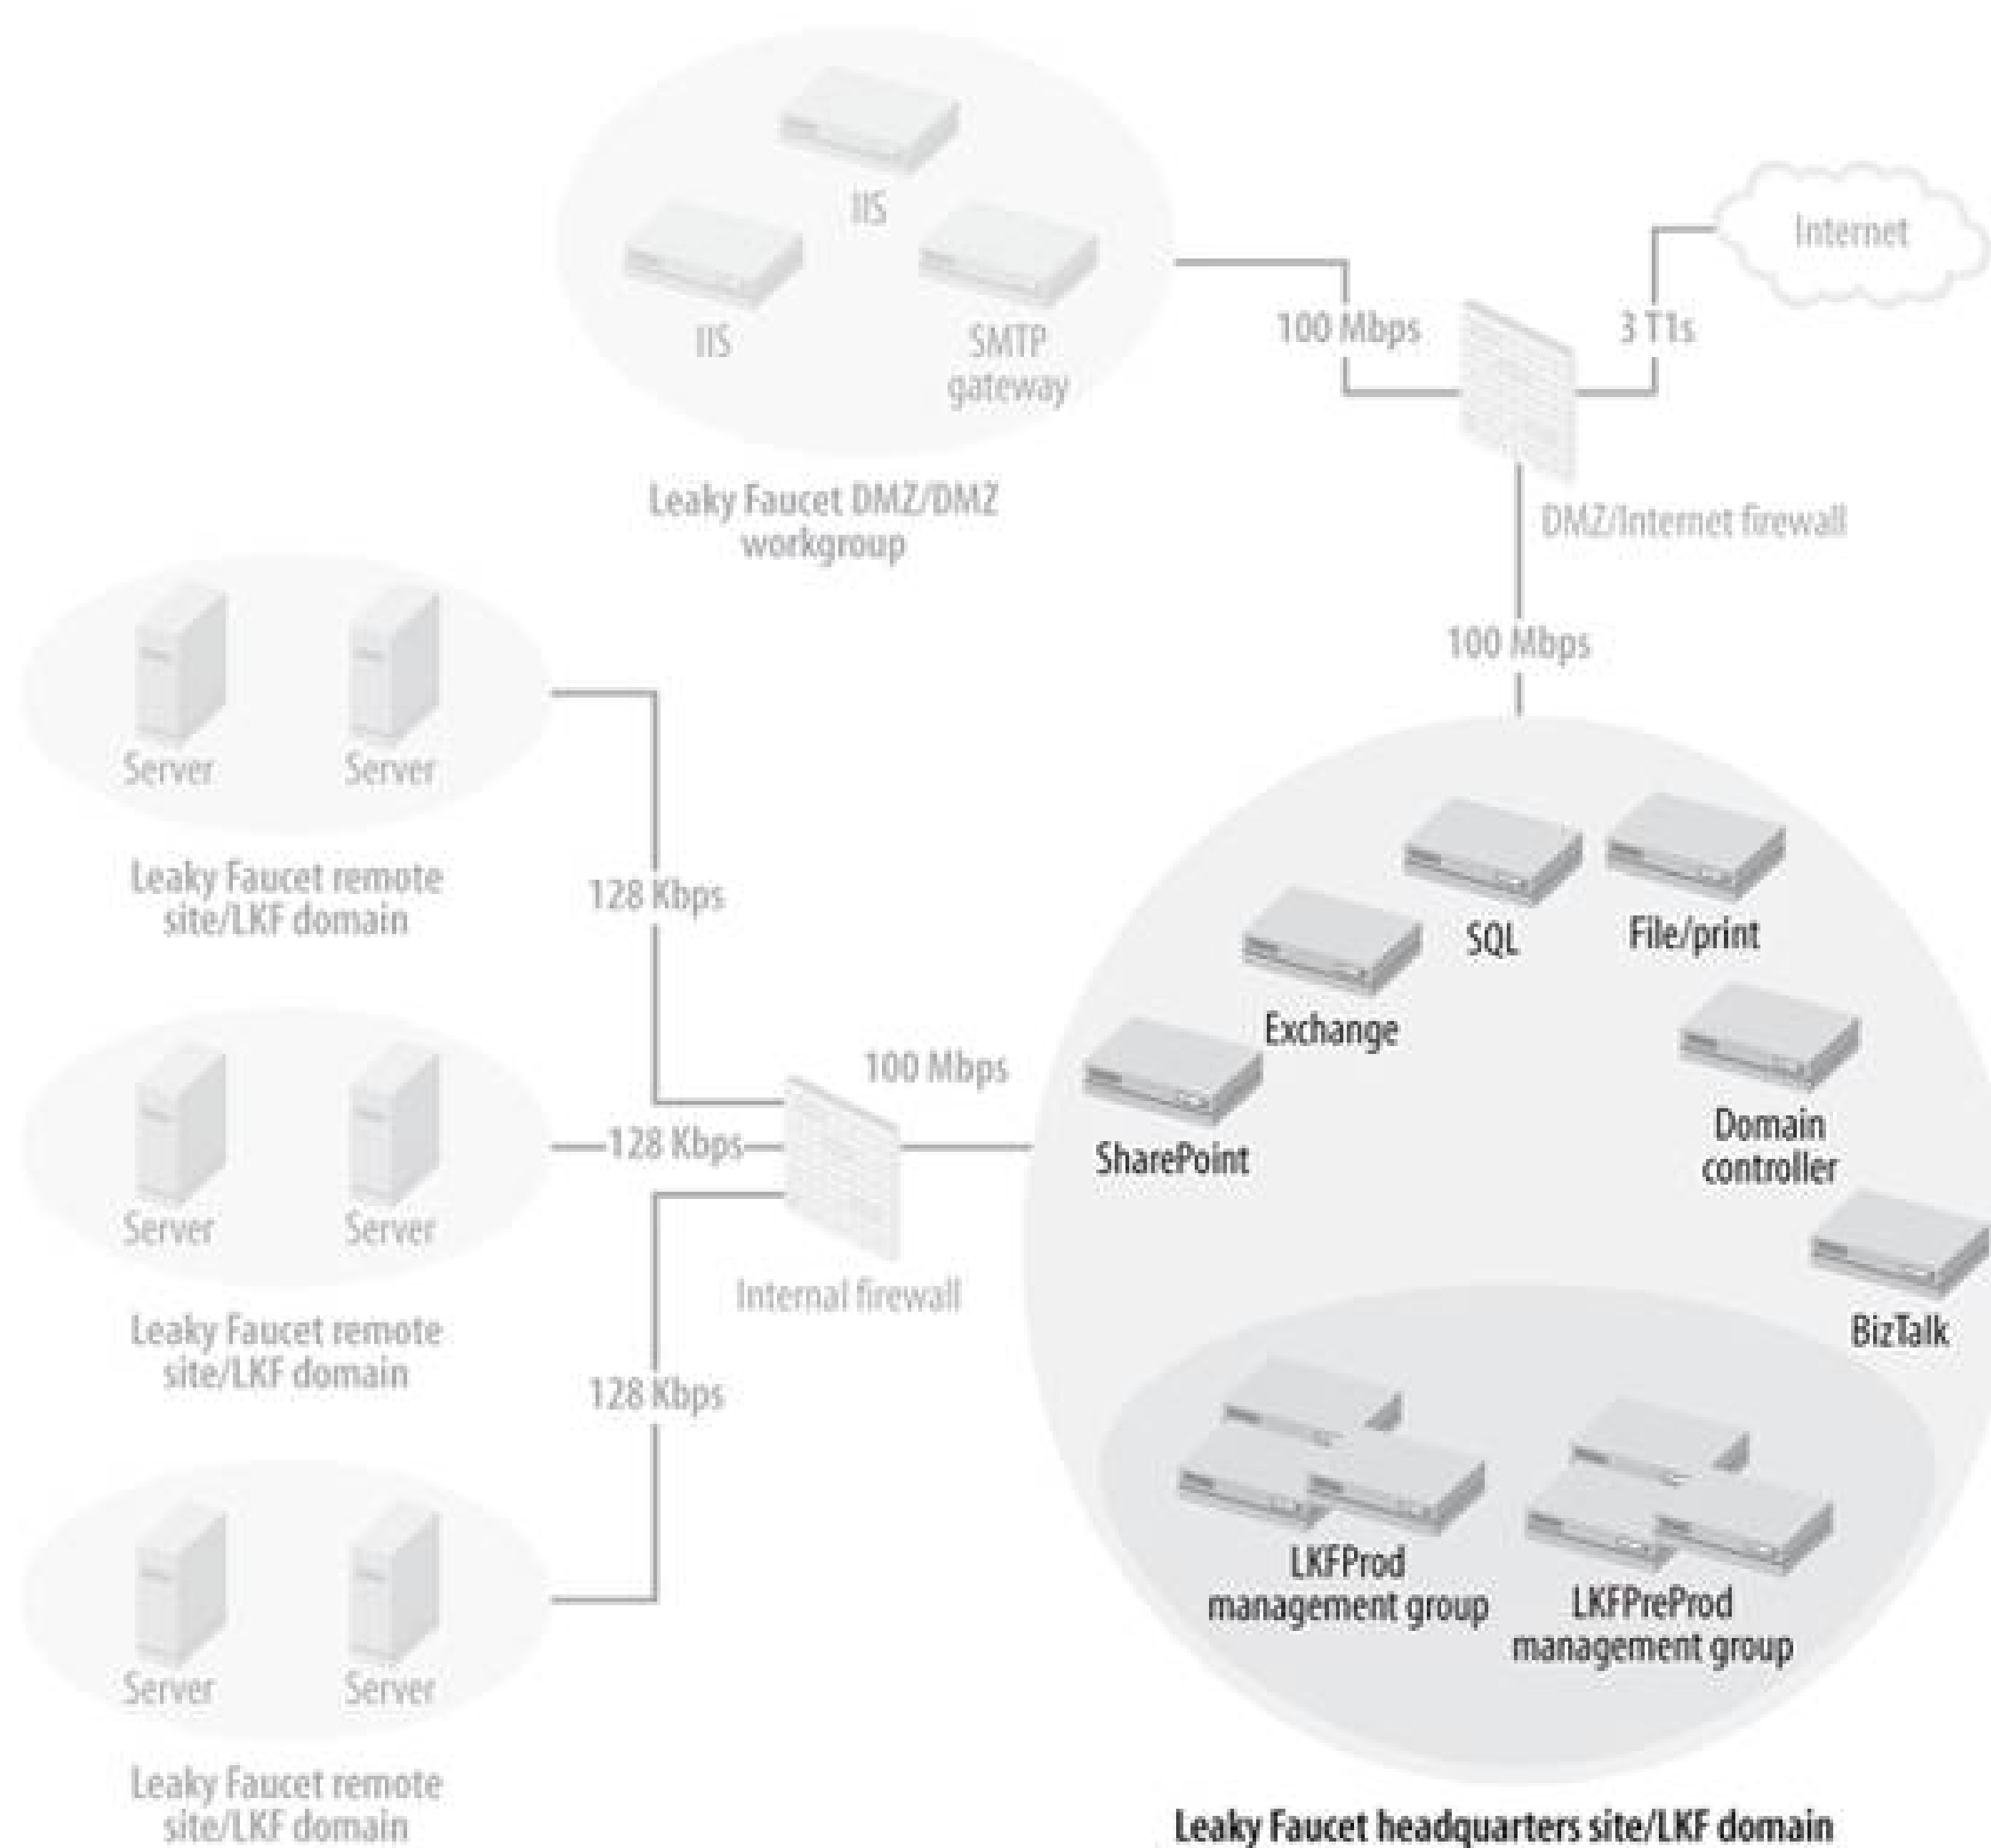
Figure 3-33. These servers don't share security or a trust relationship with the management servers





In the Leaky Faucet MOM 2005 infrastructure there are two management groups, as shown in [Figure 3-34](#). The production management group, LKFProd, is used for real-time monitoring and alerting, and feeds its data to the reporting services server. The preproduction management group, LKFPreProd, is used for testing configuration changes at the global, management server, and agent levels; testing new management packs; and developing new management packs and modifying existing ones. In addition to management pack and processing rule testing, the LKFPreProd management group can be used to initially monitor new servers and applications before they are introduced into the production environment. After a new server is monitored through a "burn-in" period, it is joined to the production network and its MOM agent is multi-homed to the LKFProd management group. To finish the migration from preproduction to production, the server is removed from the computer discovery rules that included it in the LKFPreProd group, now it is single-homed again.

Figure 3-34. The two Leaky Faucet management groups



There are other reasons for multi-homing agents. For example, if you expect a particular management pack to generate a large volume of data, say security auditing data, you may want to separate that data stream from the application and OS operational data. In this case, you would bring up a second management group and create computer discovery rules in each management group that includes the target computers. This technique is also useful if there are incompatible security or administration needs, such as the need to use different communications ports (1270 for one group and something else for the other group).

In a multi-homed environment, only the agent is aware of the multiple management groups. As far as the management groups are concerned, they each have ownership of all of their agents. The agent code is only installed on the target computer once by the first management group to include the target computer or when a manual installation is performed. This means that there is only one *MOMService.exe* installed on the agent managed computer. When a computer discovery rule is created in another management group that also includes the target computer, an entry for the additional management group is created in the agent so that the agent knows about the new management group and management server. The agent is restarted on the managed computer (the *MOMService.exe* is restarted). When the agent comes up again, it will start reporting data to and accepting configuration information from the second management group. An agent can be managed by no more than four management groups.

To prepare the agent for multi-homing, make sure that there is an additional 3 MB of disk space available for each management group that the agent will belong to. If you are going to be applying a large management pack to the agent, review its disk space needs and allow for that as well. You should also analyze the event log size requirements of each management pack and set the event log size to accommodate the largest requirement.

To multi-home an agent, you can use any of the deployment methods presented in this chapter; all are supported. To multi-home an agent that is beyond a firewall, go through the manual agent installation process once for each management group that you want the agent to report to.



# 3.7. Troubleshooting

Not surprisingly, the MOM 2005 management pack includes 95 event processing rules and 17 performance processing rules that specifically target agents, as well as hundreds of other rules that monitor the other MOM components. MOM is well equipped for self-diagnosis. You will rely on MOM's ability to report on itself in almost all troubleshooting scenarios, and almost all troubleshooting is done in the Operator console.

The main concerns of troubleshooting agents are:

- Is the agent up or down?
- If it is up, is the agent providing a heartbeat?
- If the agent is not providing a heartbeat, when was the last successful contact with the agent?
- If the agent is up, is it successfully sending event, alert, and performance data to the management server?
- Is the agent successfully receiving updates from the management server?

These concerns are listed from most critical to least, but an agent has to be doing all of these things successfully to be fully functional. The first place to check the current status of any service is in the State view in the Operator console (see [Figure 3-35](#)).

Here, the focus is on *homesqlserver* in the results pane. The current state of each machine is reported in the leftmost column and is the worst state of any of the monitored components for that machine. In the MOM agent column, you can see that the

Figure 3-35. The current agent condition is reported in the State view



state is reported as successful. In the Details pane, the heartbeat component of the overall MOM agent state is also good. If the agent service on *homesql/server* was stopped to simulate an agent failure, within a minute the State view would update to reflect the change in the condition of the agent on *homesql/server* (see [Figure 3-36](#)).

Figure 3-36. State change indicates a missing heartbeat and possibly a failed agent



The agent failure also generates a MOM Agent heartbeat failure alert in the Alerts view. This is the same alert examined in "[The Life of a MOM 2005 Alert](#)" section in [Chapter 1](#). From the information in the alert, you know that MOM pinged *homesql/server* and got a successful response, but the agent is not responding. On the product knowledge tab, one of the possible solutions is to ensure that the agent service is running on the target computer. To further troubleshoot this, open the Tasks pane while in the State view to keep the focus on *homesql/server* and run the Start MOM 2005 Service task ([Figure 3-37](#)).

The output of the task is returned in the console task output box and the MOM agent and heartbeat status are both returned successfully. The State view also shows the time and date of the last heartbeat, which is the time that the computer was last contacted, when the target computer's name in the computer column is selected (see [Figure 3-38](#)).

Figure 3-37. Running the Start MOM 2005 Service task in the tasks pane

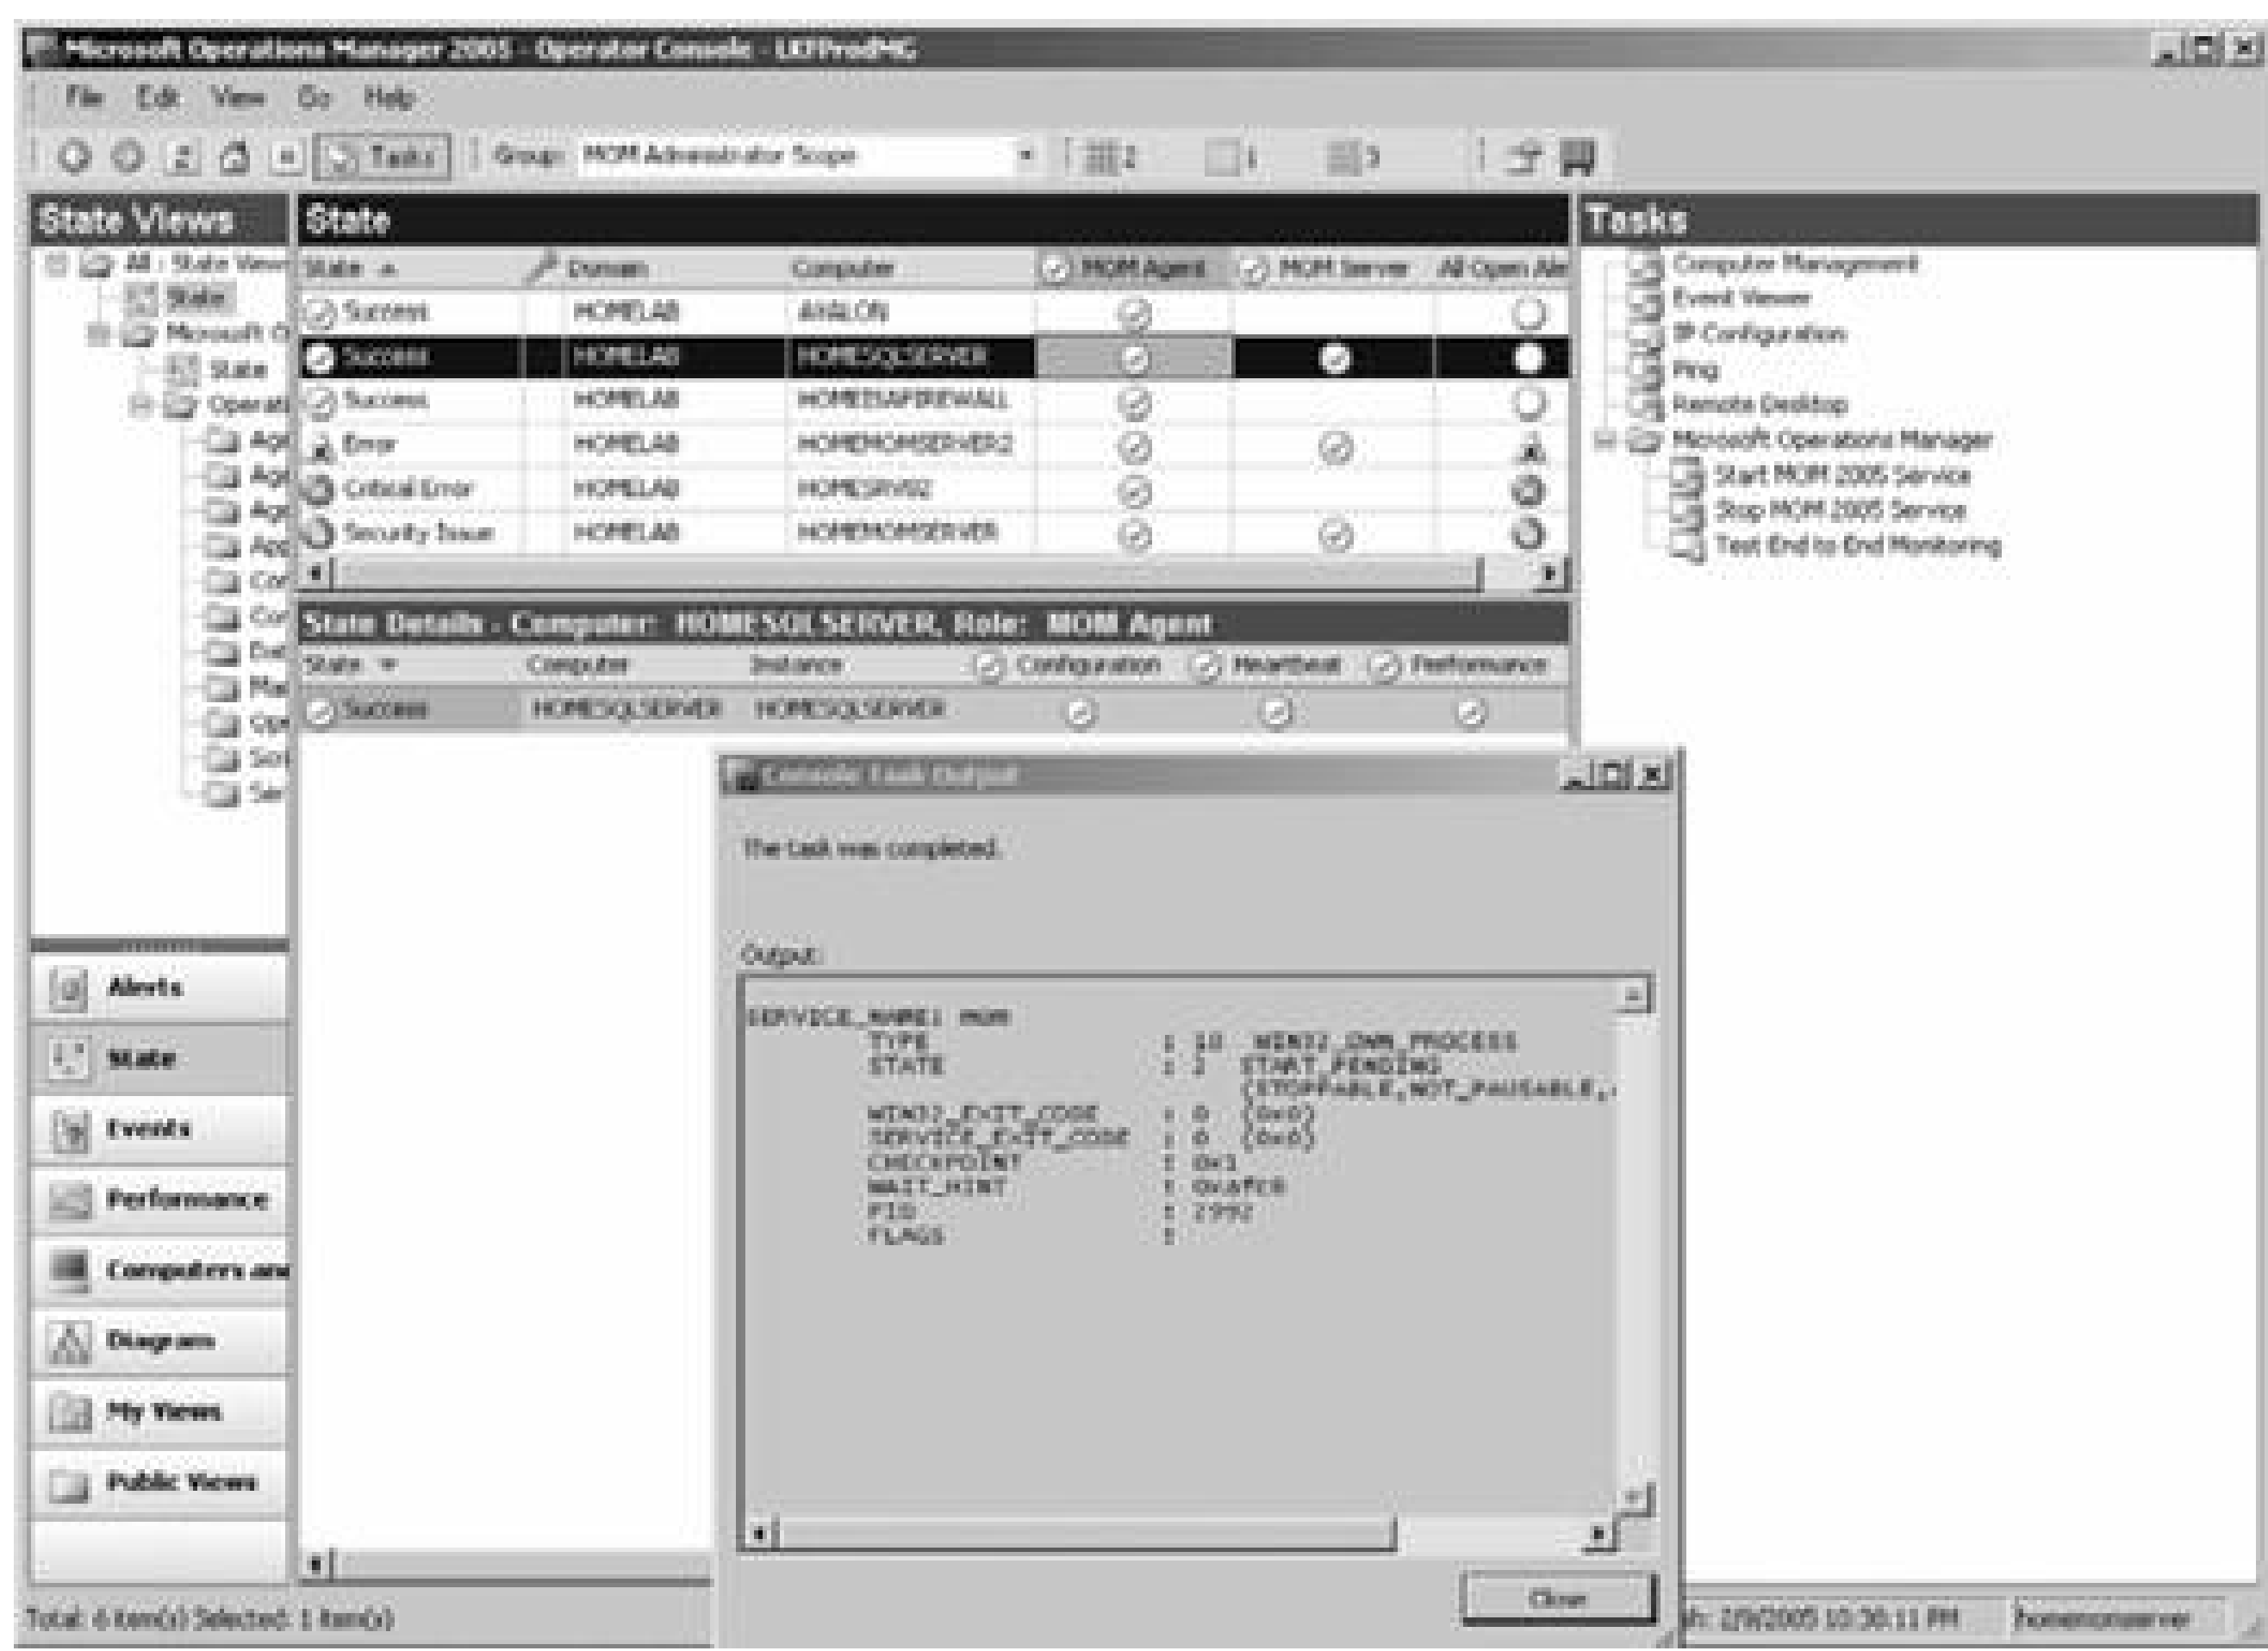


Figure 3-38. Date and time of the last heartbeat

Next, to ensure that the agent is running correctly, run the Test End to End Monitoring task in the Tasks pane while the focus is on *homemomserver*. This task causes the agent to place a specific event in the target server's Application event log (event ID 22078), which then generates an informational alert that reports back to MOM (see [Figure 3-39](#)).

To track the status of the task, switch to the Public Views/Task Status folder shown in [Figure 3-40](#). The Test End to End Monitoring task generates two events in this console: a 9897, which states that the Test End to End Monitoring task has been scheduled, and a 9898, which states that the task completed successfully (see [Figure 3-40](#)).

Figure 3-39. Launching the Test End to End Monitoring task

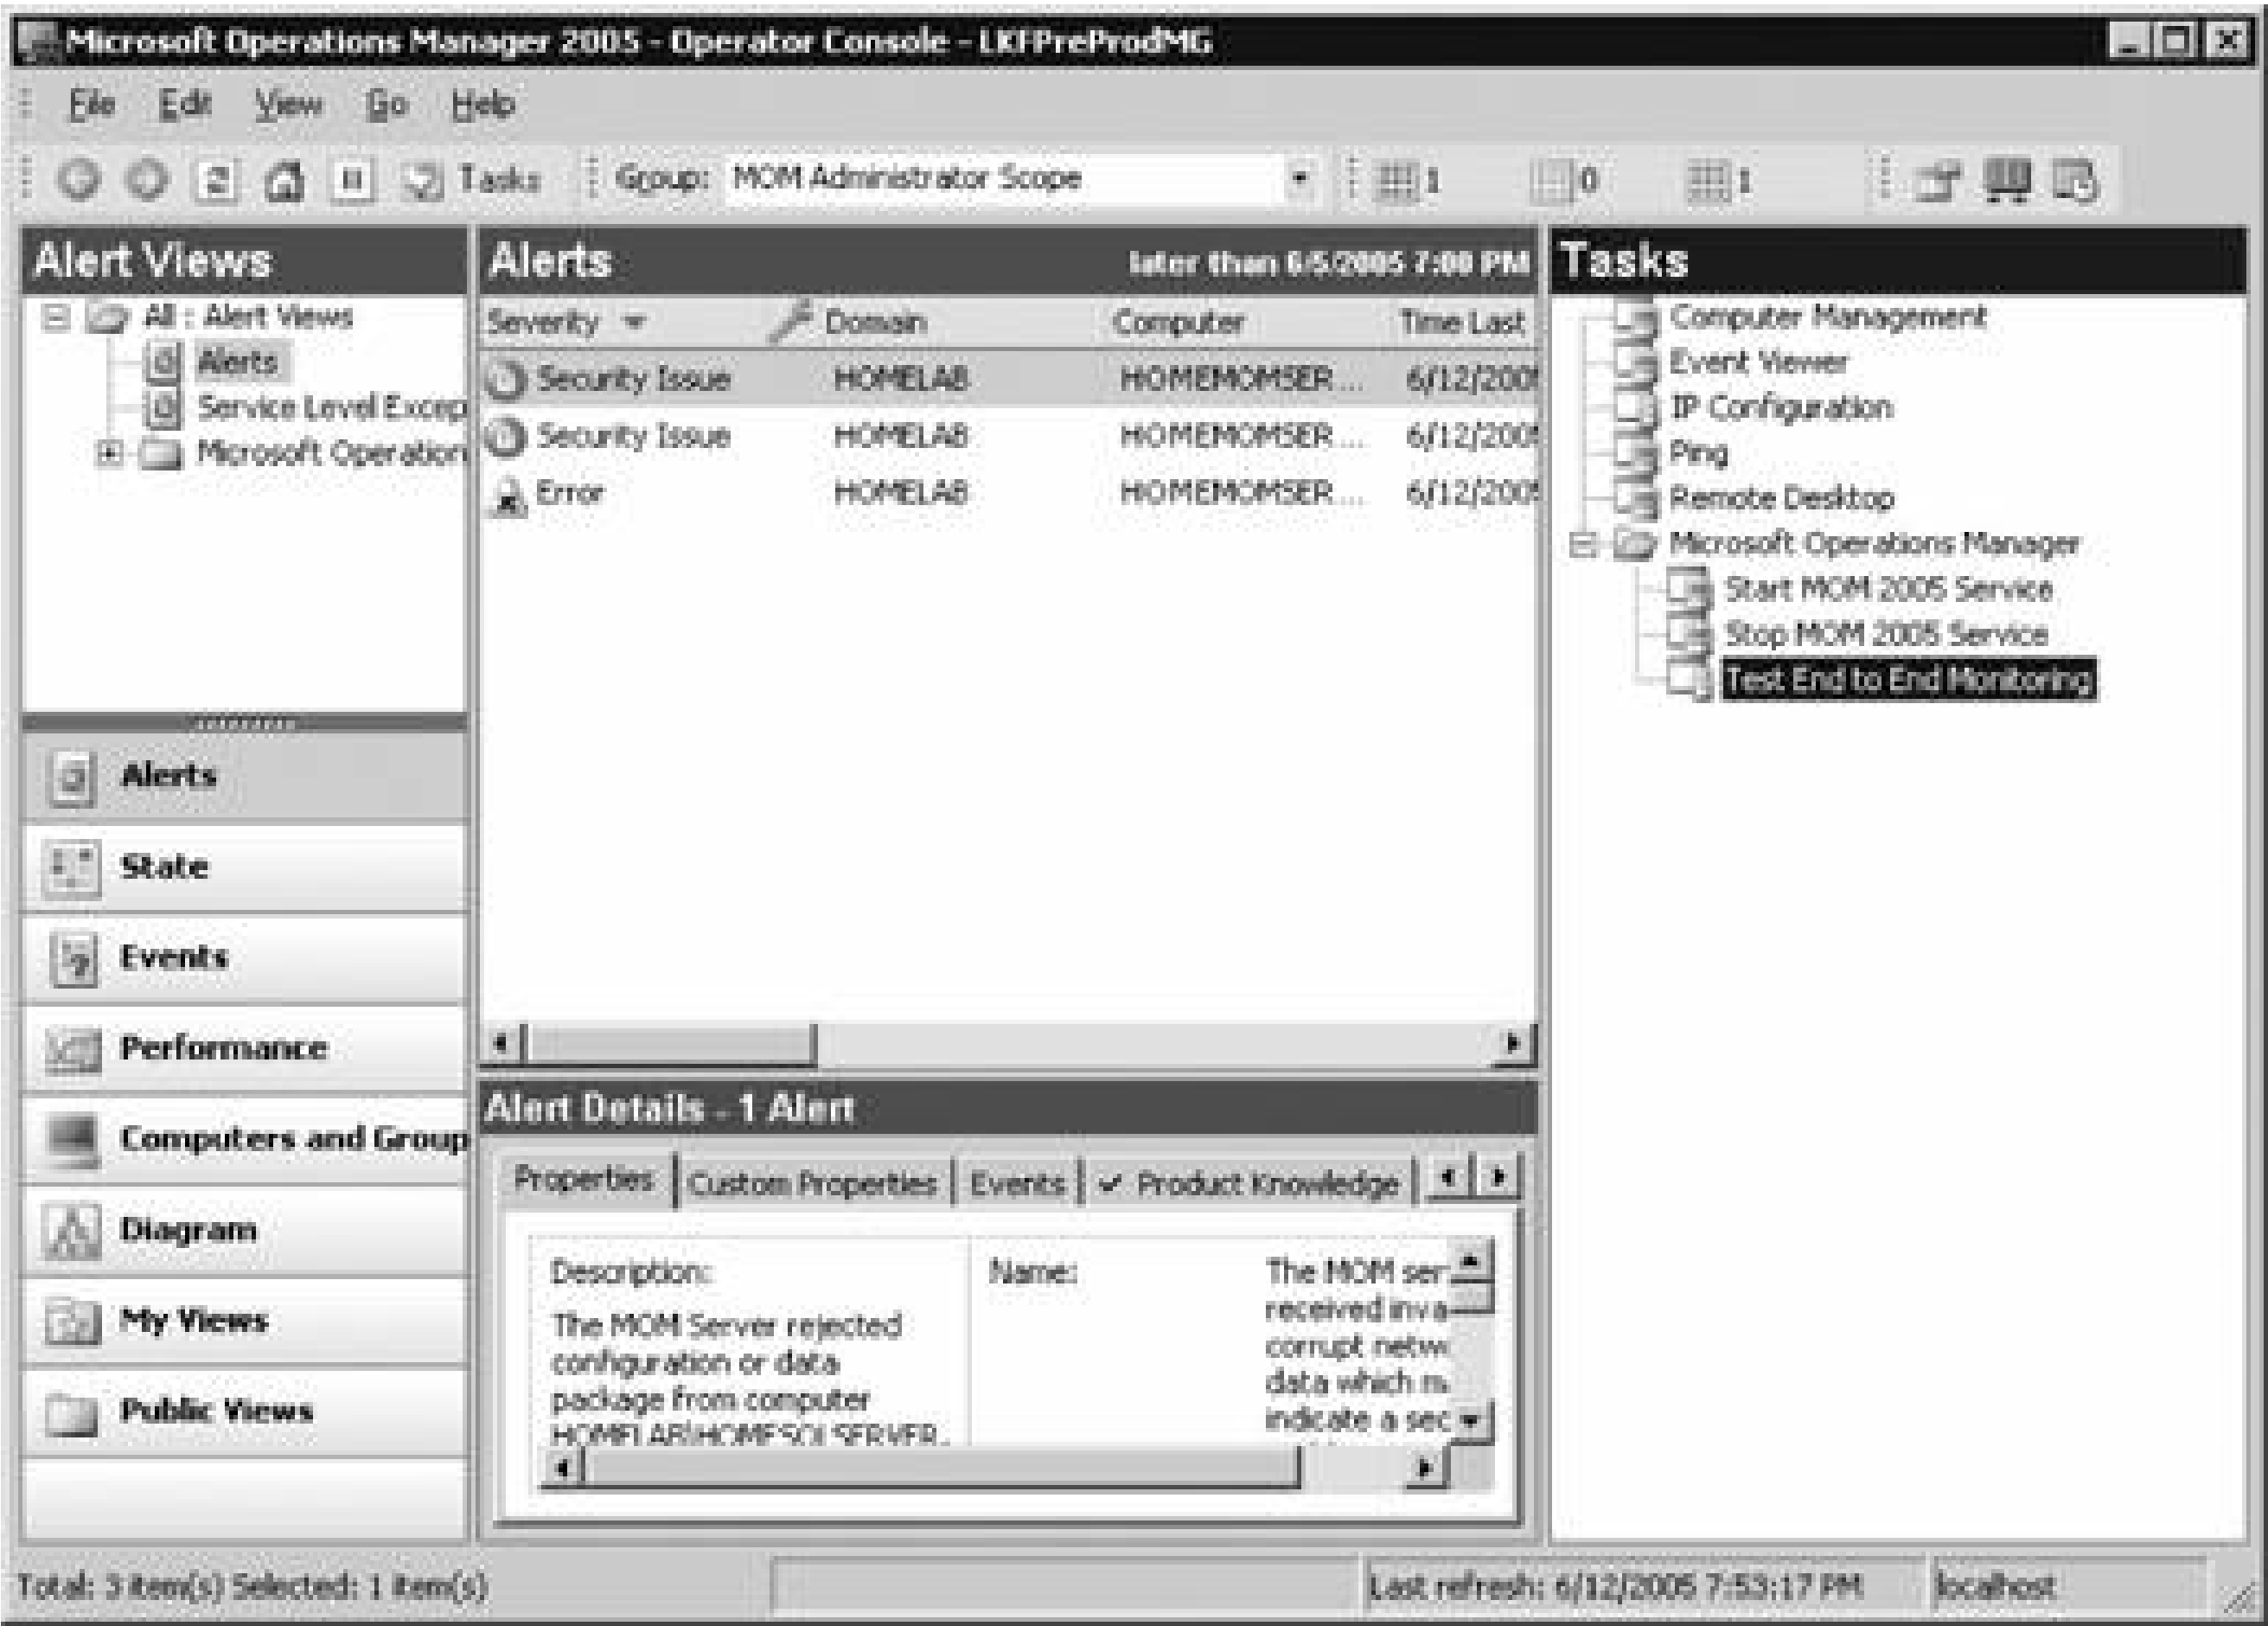
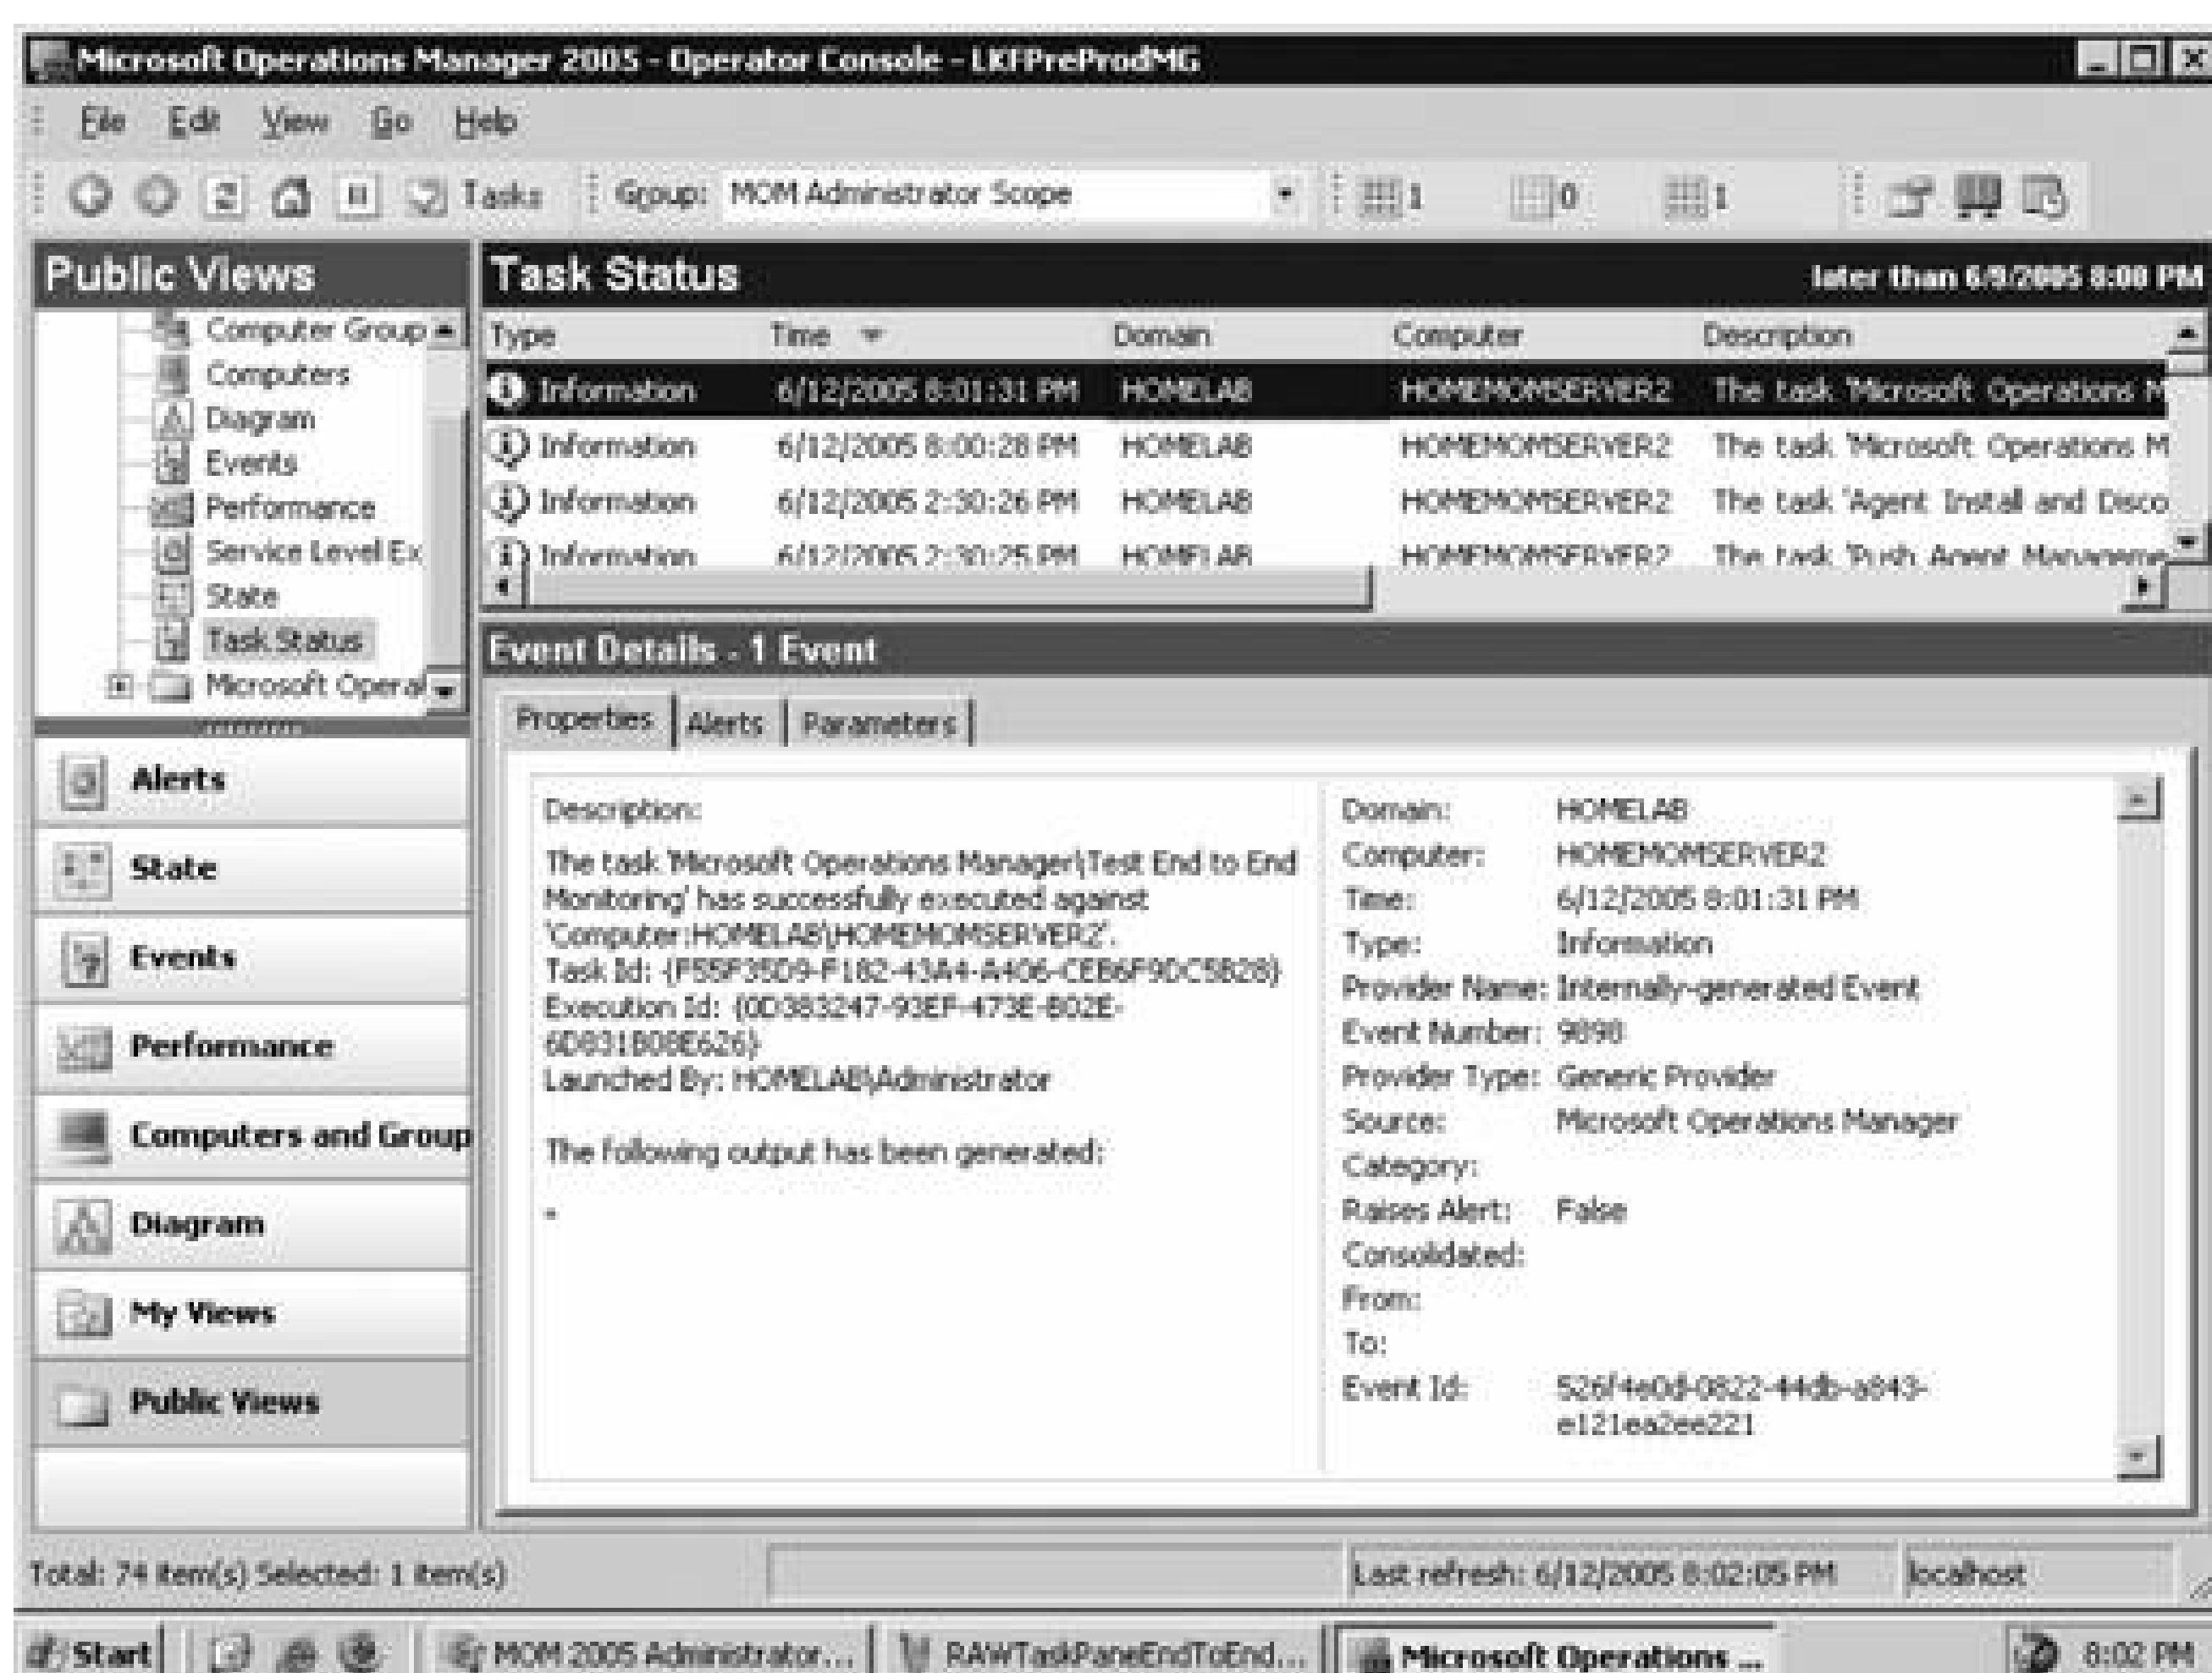


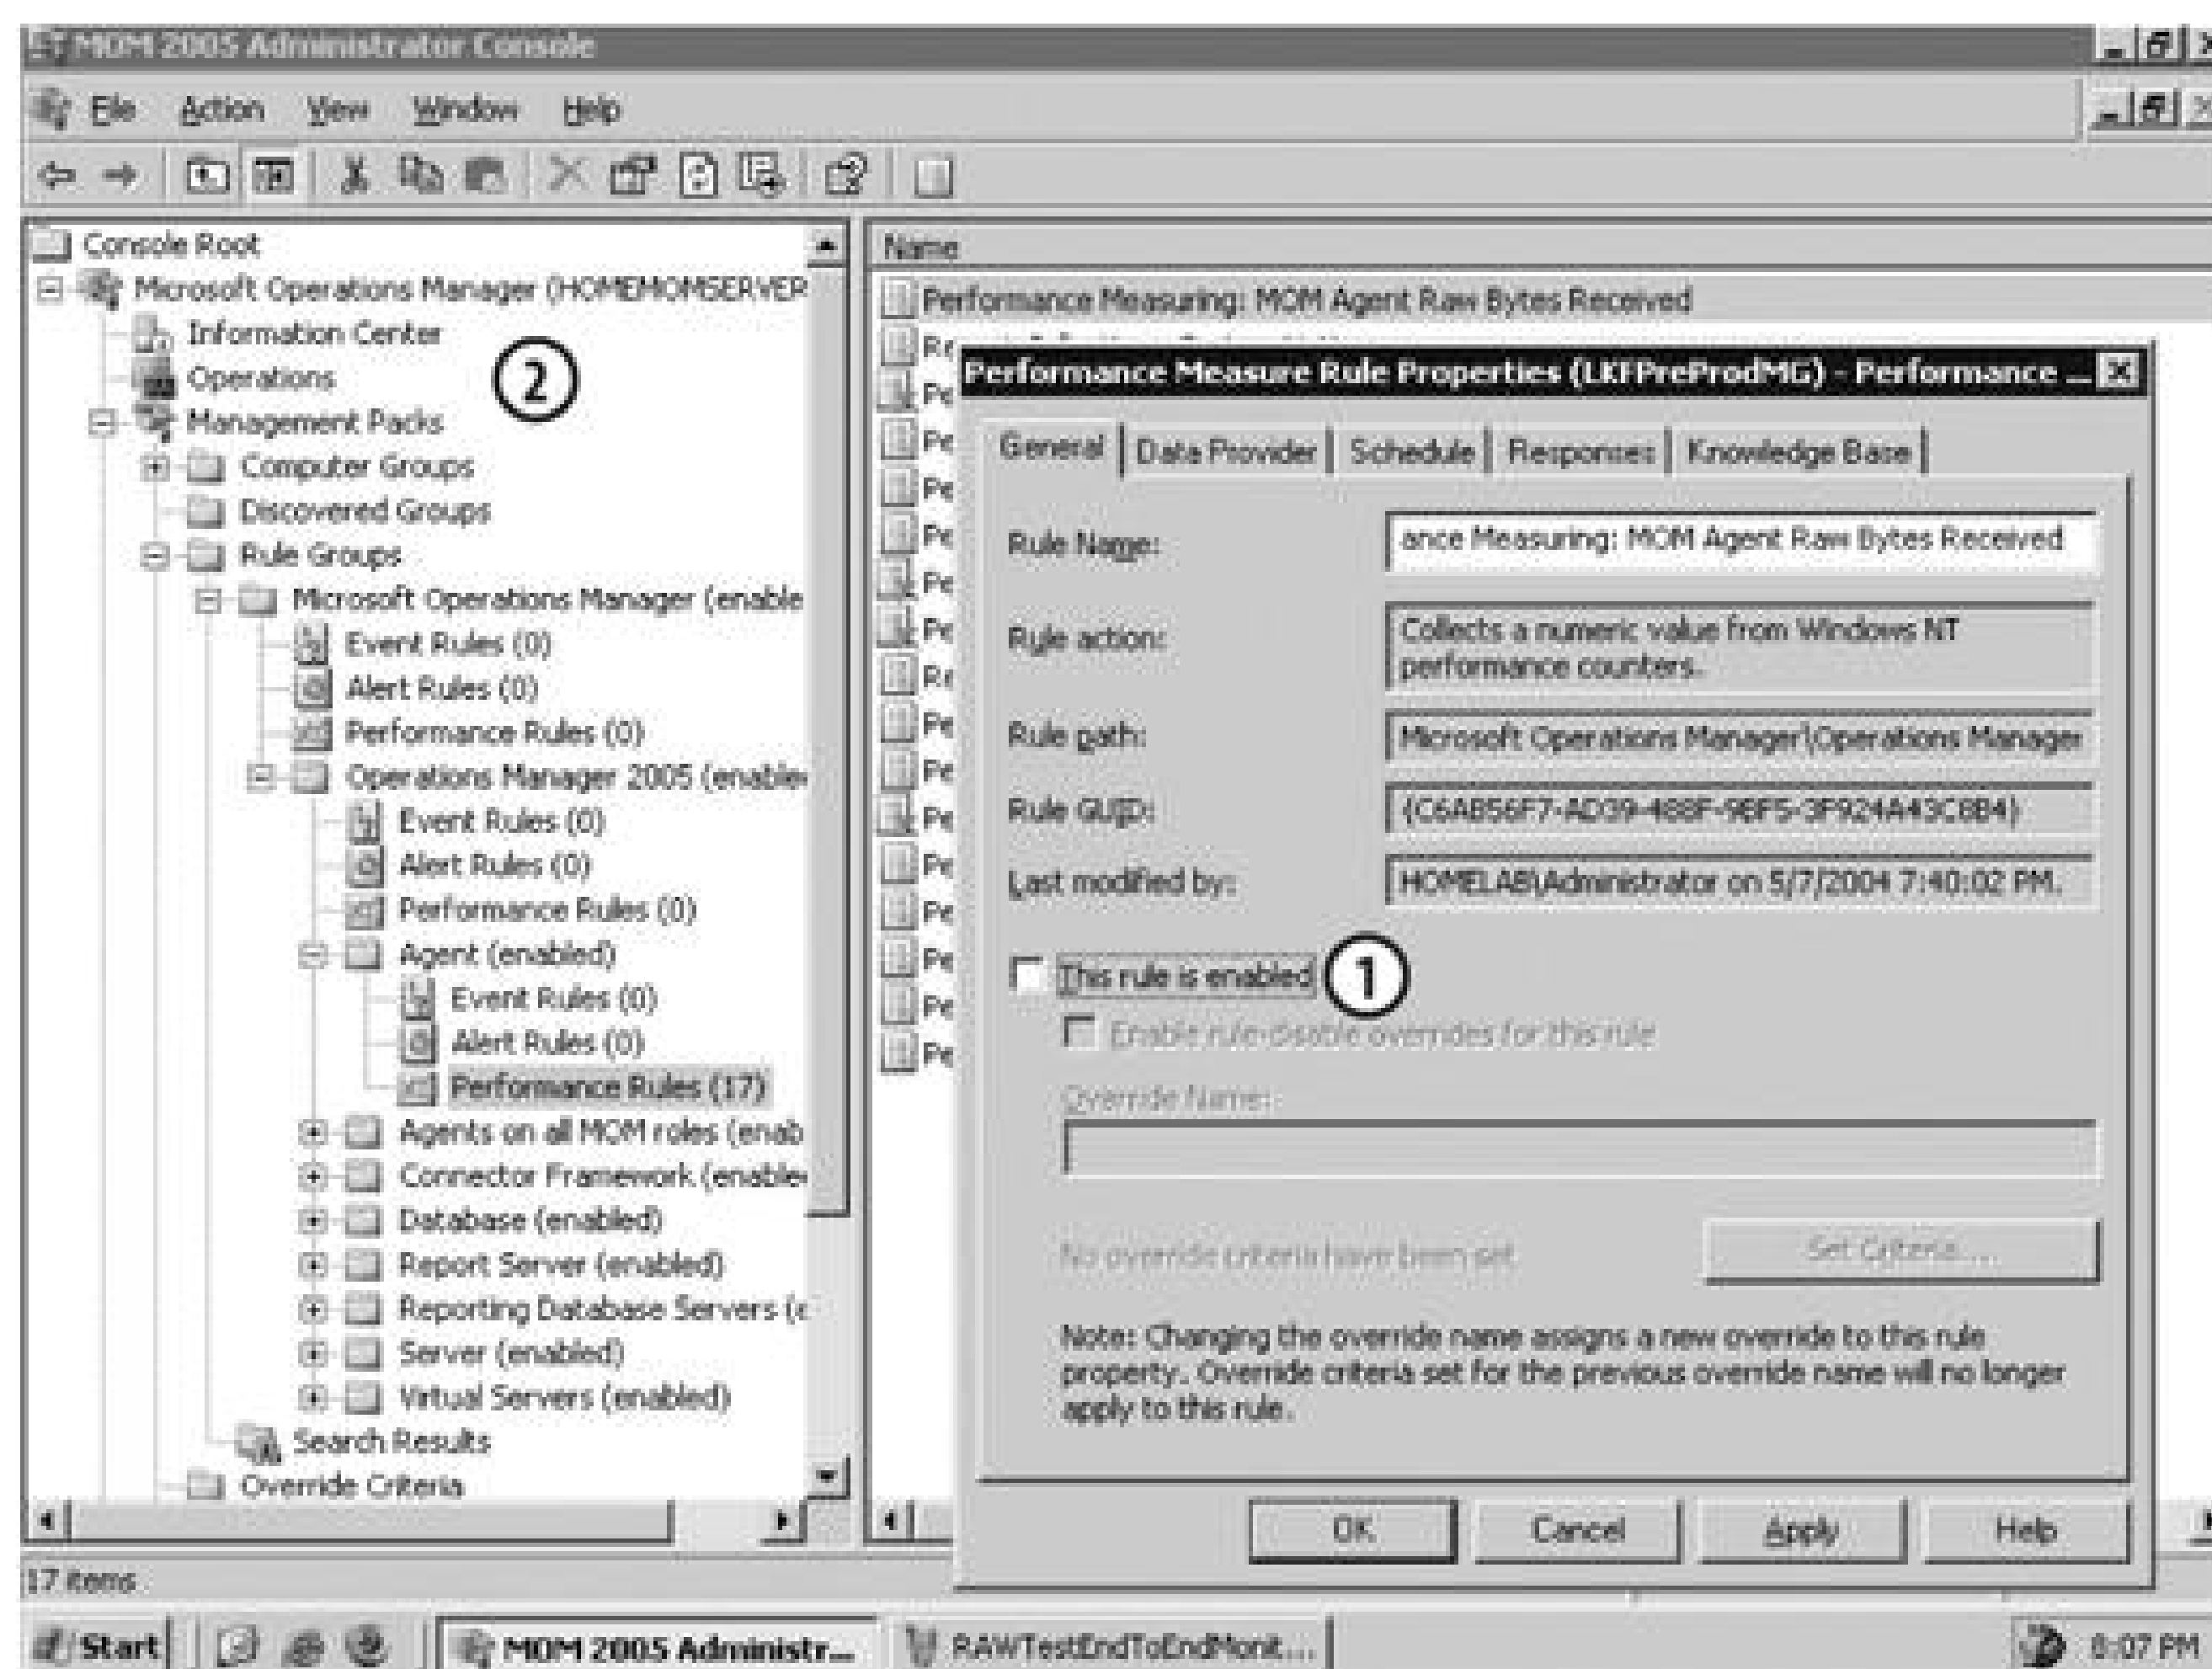
Figure 3-40. Task status tracking





The first four troubleshooting concerns have been addressed, but you still need to know if the agent is receiving configuration update information correctly. To do this, you make a minor change to a processing rule in the Administrator console (point 1 in [Figure 3-41](#)) and then Commit Configuration Change (see point 2 in [Figure 3-41](#)). In this case, disable a rule in the Management Packs Rule Groups Microsoft Operations Manager Operations Manager 2005 Agent Performance Rules rule group. MOM 2005 will submit a task to update the rules on the affected agents.

Figure 3-41. Disable a rule in the Administrator console



When this happens, an event ID 21240 is generated in the Application Event Log on the target server. You then can either watch the Application Event Log on the target computer or switch to the Public Views/Events container to see that event be returned from the target machine (see [Figure 3-42](#)).

## 3.8. Tools

The MOM 2005 Resource Kit has two tools that are particularly helpful in dealing with agents: the Resultant Set of Rules (RSOR) tool and the Agent Helper tool. These tools are located in the Tools directory of the Resource Kit, which is downloadable from <http://www.microsoft.com/mom>. These tools will help you obtain special detailed information about a required agent when troubleshooting agent issues.

Figure 3-42. The 21240 event indicates that the agent successfully received a configuration update from its management server



Because these tools are Resource Kit utilities, Microsoft does not support them via the regular support channels.

### 3.8.1. Resultant Set of Rules Tool

The RSOR queries the MOM 2005 operational database for all of the rules in a management group that are applied to an agent. Although this is a command-line tool, it is similar to the Resultant Set of Policy query in the AD group policy management console. It lets you know exactly which processing rules an agent is executing. You can run this tool from any machine that can access the operations database server. It accesses the data via the ActiveX Data Object (ADO) model. To start the tool, use a command similar to this:



```
C:\Rsor.exe HOMESQLSERVER Avalon
```

This will query *homesqlserver* for all the rules that apply to the agent on computer avalon. The RSOR tool creates a directory on the root of the drive that it is run from, called *ResultantSetOfRules*. On this drive the output file is automatically named, in this case *AVALON RSOR 2005-2-10 001220.log*. The file contents look like the following (many lines were deleted because they weren't necessary, but the file ends up with a total count of the rules applied):

```
Microsoft Operations Manager - Resultant Set of Rules 1.1
Exporting rules applied to managed node AVALON
Microsoft Operations Manager 2005 Agents (Enabled)
-----
(Level1) Agent (Enabled)
-----
1. PMC Measure - Performance Measuring MOM Agent Alert Incoming Rate (Enabled)
2. PMC Measure - Performance Measuring MOM Agent Data Incoming Rate (Enabled)
3. PMC Measure - Performance Measuring MOM Agent Raw Bytes Received (Enabled)
4. PMC Measure - Performance Measuring MOM Agent Raw Bytes Transmitted (Enabled)
5. PMC Measure - Performance Measuring MOM Host - %Processor Time (Enabled)
6. PMC Measure - Performance Measuring MOM Host - Handle Count (Disabled)
7. PMC Measure - Performance Measuring MOM Host - Private Bytes (Disabled)

(Level1) Agents on all MOM roles (Enabled)
-----
1. Event - A management pack script was unable to complete successfully (Enabled)
2. Event - A script failed (Enabled)
3. Event - A script failed to get or set a varset (Enabled)
7. Event - Agent communication failure troubleshooting events (Disabled)

The total number of rules exported was: 112
```

This tool is useful to see if a rule was applied but isn't showing up, or if the wrong rule was applied. You would then compare the output of the RSOR tool to the rule groups in that computer by virtue of computer group membership and look for discrepancies. You would do this in a situation where an override to a processing rule has been configured for a particular computer and you want see if the override has been applied properly. See the "[Overrides](#)" section in [Chapter 4](#).

### 3.8.2. Agent Helper

The Agent Helper is actually a management pack; for more about management packs see [Chapter 4](#). When an agent fails to return a heartbeat, this tool will execute a managed code response (.NET code) against the managed computer. It does this from the management server, so it is called a *server-side response*. To use this tool, you must import the management pack in the Administrator

console and copy the *AgentHelper.dll* into the installation directory, which is usually *C:\Program Files\Microsoft Operations Manager 2005* on the management servers. Note that while you only need to import the management pack once into the management group, you need to copy the *AgentHelper.dll* into the installation directory on each management server in the management group because the responses could run from any of them.

When a failed heartbeat is detected, the responses will be one of three actions depending on the situation. The rule could attempt to restart the agent, reinstall the agent if it is missing, update the agent's settings and thereby force a re-establishment of communications between the agent and the management server.



## 3.9. Summary

In MOM 2005, agents are the intelligent data miners that perform most of the real work. By performing most of the pattern matching between the processing rules and the collected data locally, rather than just pushing raw data up to the management server for processing, agents increase the overall performance of MOM 2005. Besides collecting data, agents also perform actions on the managed computer when a processing rule response instructs them to do so. The agent runs the responses under a different process (*MOMHost.exe*) and security context (agent action account) to protect the agent service (*MOMService.exe*) from hung scripts or other poorly executed actions that could impede its function.

In MOM 2005, when agent code is installed locally on a computer, that computer is said to be "agent-managed." MOM 2005 can also perform monitoring on a computer via the management server agent which is called "agentless monitoring." MOM 2005 agents can be deployed to and manage computers that are in a wide range of network and security environments, including across firewalls and slow WAN links and into trusted domains or untrusted domains and workgroups. This can be done remotely from the management server or locally on the target computer via a manual installation process.

MOM 2005 agents can execute the monitoring requirements of multiple management groups simultaneously and will failover automatically between management servers that are in the same management group in case of agent-to-management server communication failures.

Troubleshooting MOM 2005 agents primarily involves making sure the agent is running, monitoring heartbeats, testing round-trip monitoring, and ensuring configuration updates are being received from its management group.

So far, the agents have only executed the MOM 2005 management pack. There are many more that need to be implemented so agents have the rules needed to monitor all the applications. But there is more to implementing management packs than just importing them. Management packs undergo a life cycle and evolve to more closely fit your environment the longer they are used; therefore, you must be careful to administer them correctly. [Chapter 4](#) provides a solid foundation and methodology for doing so.



# Chapter 4. Administering Management Packs

The MOM 2005 components discussed so far provide the functional infrastructure of MOM. Each component plays a specific role in the management group. The agents interact with the managed computers, the management servers communicate with the database and manage the agents, the databases store collected data and configuration information, and the consoles allow interaction with MOM. But the picture is not yet complete. At this point, all the components of a management group are like the members of an orchestra. They are sitting where they are supposed to be in relation to one another, they have a conductor that coordinates the actions of the various instruments, they have tuned up, and the audience is waiting for them to perform. But no meaningful sounds are being made because they have no sheet music—the management packs. Management packs allow the management group to identify the applications on a computer, to group computers by function, and to determine the health of the applications on those servers. They then tell management group components how to notify the outside world of pertinent information, format the information for consumption, and offer tools for troubleshooting the application.

Install reporting services before you import management packs into the management group. By doing this, the report definitions can be imported at the same time.

Microsoft has publicly stated that every server that is a member of the Windows Server System grouping will have an associated management pack developed by the application product team. They have met this goal for the most part, although sometimes the release of the management pack lags well behind the release of the application.

In MOM 2005, management packs have grown up. Earlier versions of management packs were *alert-oriented*, but the MOM 2005 management packs are *state-oriented*. These management packs provide insight into the current state of an application or a server and each of its components. They do so in an easy-to-interpret red-, yellow-, and green-light fashion. You have already seen this in the State view in the Operator console in [Chapters 1](#) and [3](#). The alert helps determine the health state of an application or server. Alerts need to be handled individually, but they are not the highest level of data that MOM 2005 produces.

In the future, under Microsoft's Dynamic Systems Initiative, management packs may evolve into or be replaced by objects called system definition models (SDMs). These models will logically represent entire computer systems (e.g., hardware, OS, and applications) to developers and monitoring applications.

This chapter discusses the most common tasks involved with administering management packs. You will see how Leaky Faucet administers management packs, and some beneficial tools will be introduced. More extensive treatment will be given to the rules that the agents execute, as well as the components of management packs. From this you will learn how to modify existing management

packs and rules for your environment and how to create a simple management pack from scratch.

How well MOM produces actionable, pertinent information and how much it raises your operational awareness depends on how the management packs are administered and how the rules in those packs are adjusted.



## 4.1. Management Pack Life Cycle

There are two types of management packs: those developed by parties outside your company (either by Microsoft or a third party, such as NetIQ or eXc Software); and those developed internally because there are no third-party management packs for the application. An application may not have a vendor-developed management pack because it is obscure and too small to warrant one, because it doesn't have a service that can be monitored and doesn't write to the event log, or because the software vendor just isn't interested in developing one. Most likely, a management pack does not exist for an application if it was homegrown.

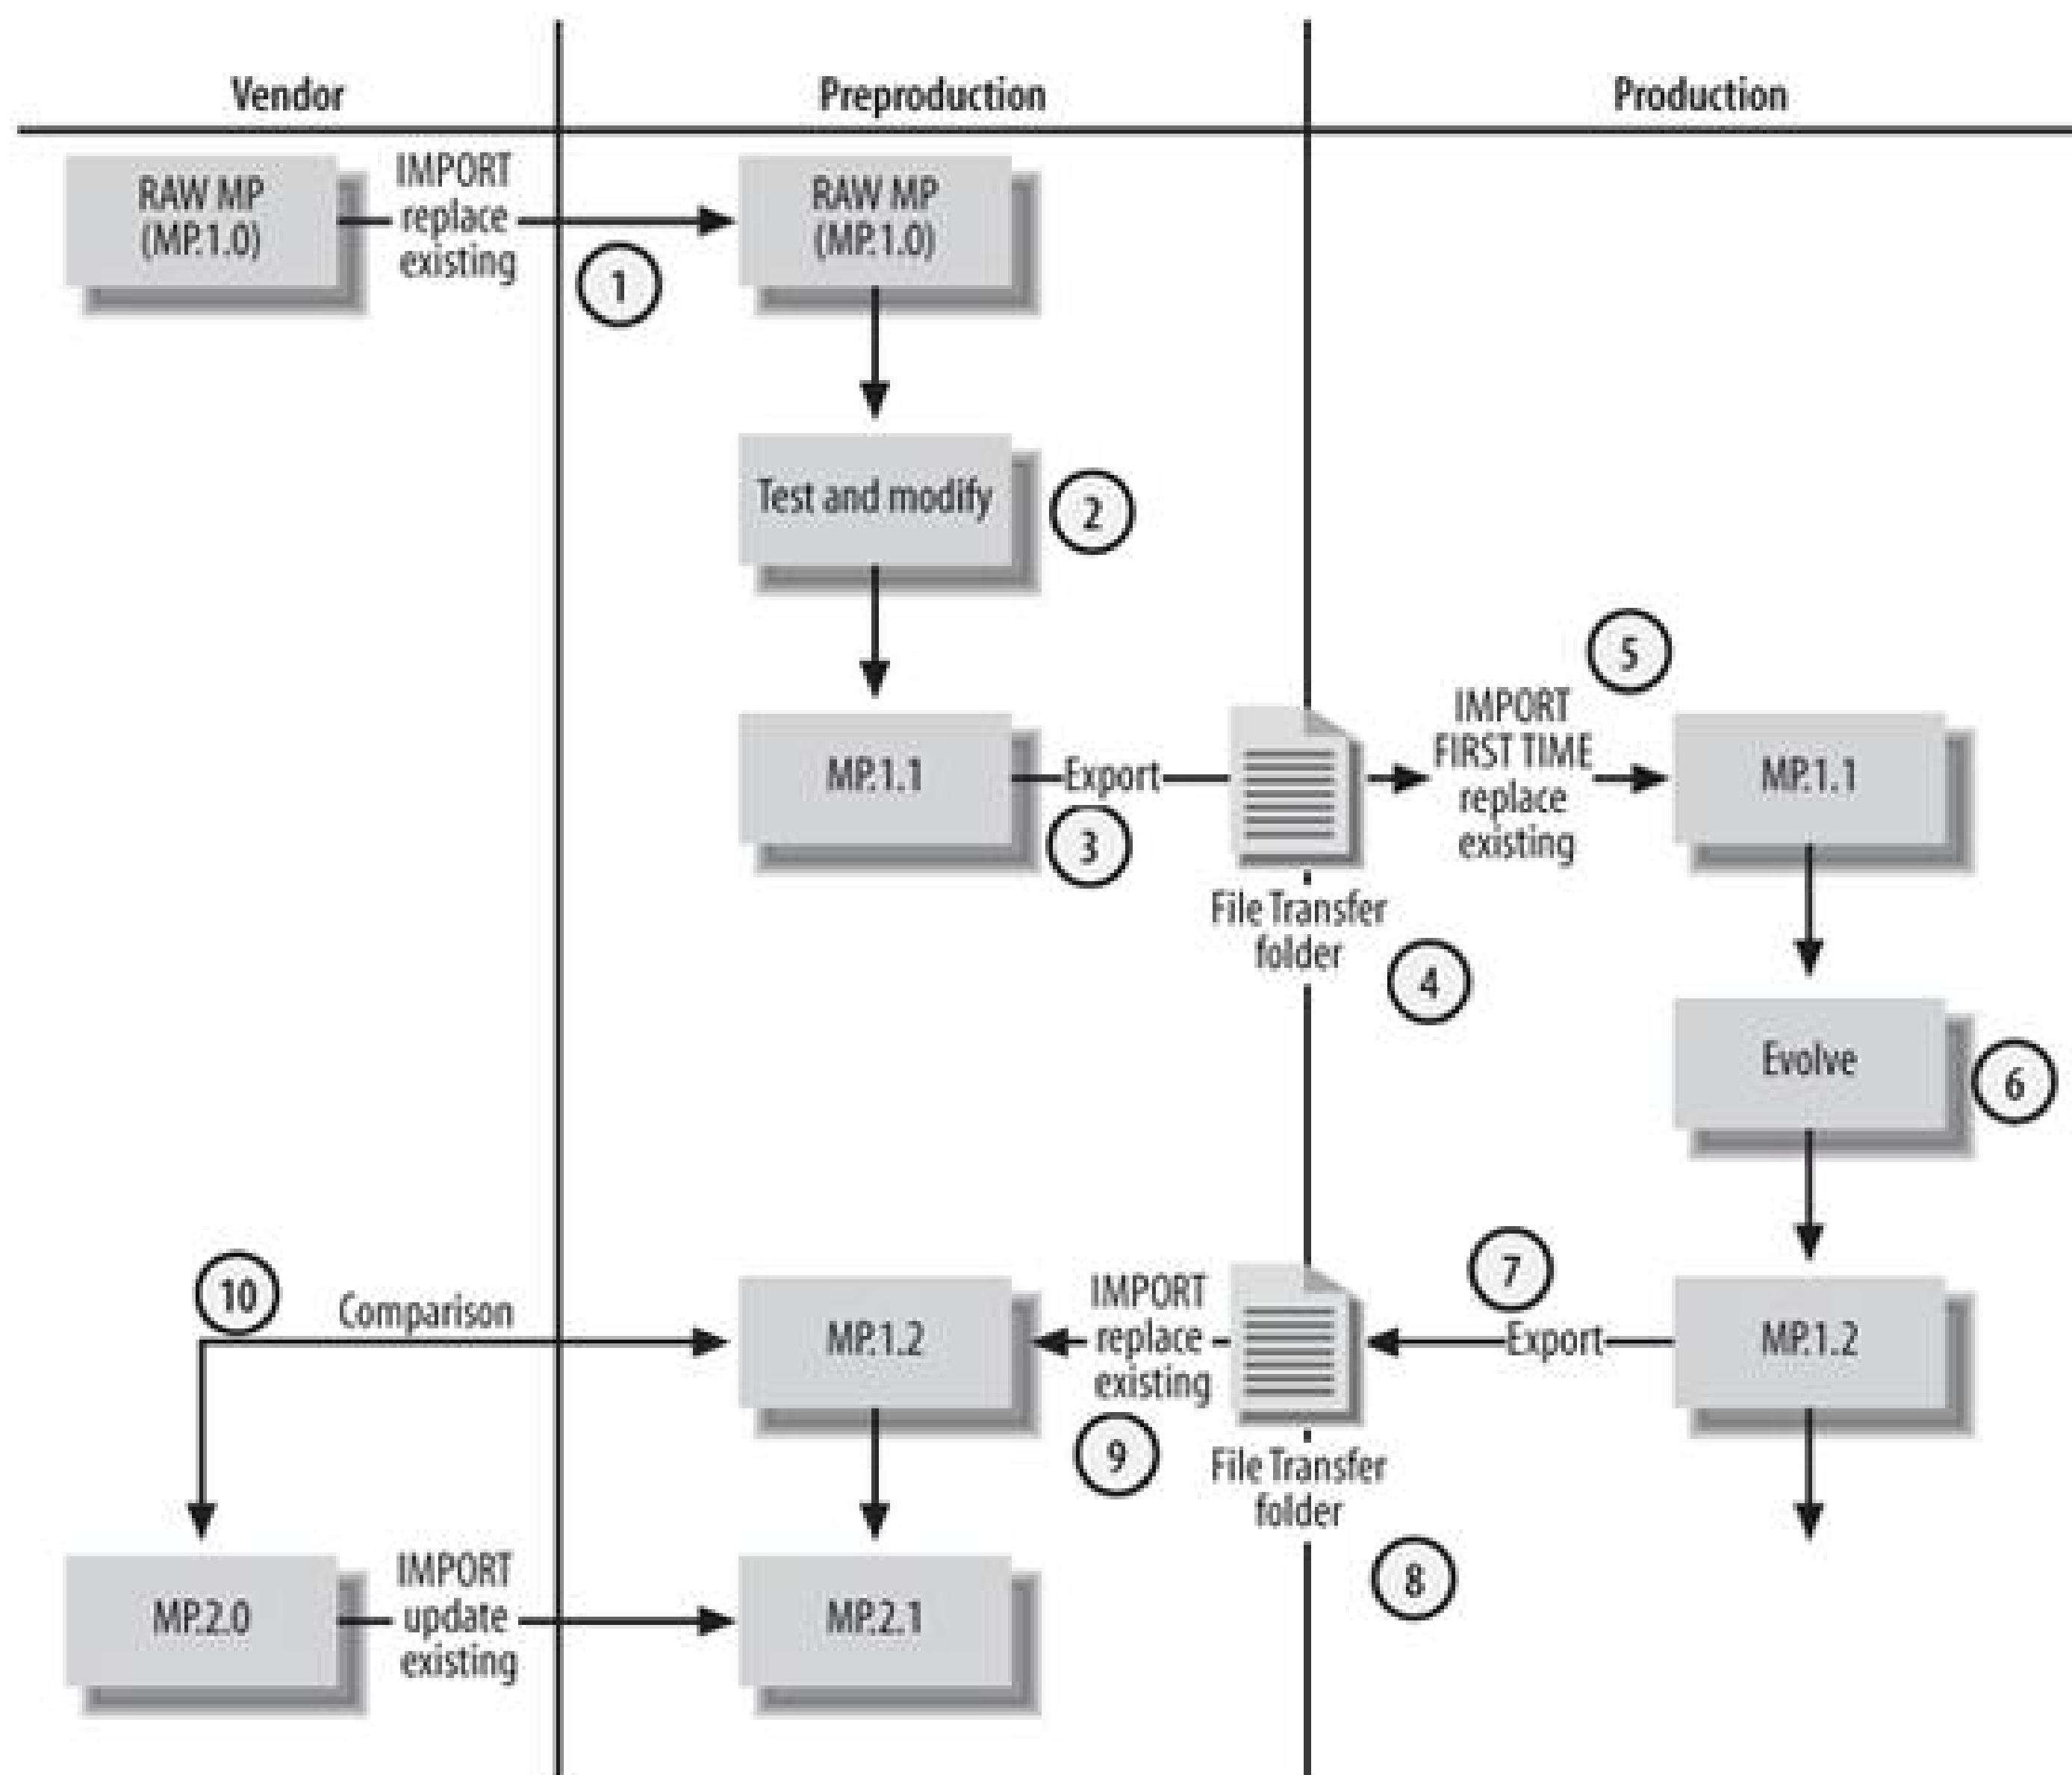
Once imported into MOM, both management packs follow the same life cycle. Managing the life cycle of management packs is what the rest of this chapter is about. *Life cycle* describes the process that the management pack goes through as you take it from its original vendor-created state and tailor it to your environment, first by testing it and modifying rules and then by using it in your production and preproduction environments. As the management pack is used and you troubleshoot alerts, you will capture those troubleshooting steps in the management pack. You will also likely change the thresholds for rules that generate alerts and the severity of those alerts. All of these management pack modifications mean that the management pack you are using today will not be the same as it was the day before, or as it will be tomorrow.

Most companies get into trouble with their management packs because of poorly implemented version control . Because companies don't have a plan for managing change, they end up with multiple versions of the same management pack and don't know which one has the most up-to-date rule sets and there is no record as to why the rules were changed. When a vendor update needs to be applied, it is unclear which existing version of the pack it should be merged with, so the existing version is overwritten and all customization is lost. In this regard, managing management packs is somewhat like maintaining code versions. There must be a solid management pack version control methodology and all the changes to the management pack rules must be documented.

[Figure 4-1](#) shows a process that can be used for moving, modifying, and tracking versions of any given management pack. This workflow is the foundation for most of this chapter, and it will be referred to often. The workflow brings a management pack into the preproduction environment, tunes it, and deploys it into production. From there, it covers the evolution of the management pack in production and how to synchronize those daily changes back to the preproduction environment. Finally, the workflow shows how to integrate changes into existing versions of a management pack when the vendor provides an update.

Figure 4-1. Life cycle workflow of a management pack





The workflow focuses on maintaining control through the preproduction and production environments. The vendor column is included as another entry point for change.

## 4.2. Importing Management Packs into Preproduction

[Chapter 3](#) left off with agents deployed and their health confirmed. Some agent heartbeat data was returned from the agents to the management server. This is because the MOM 2005 management pack is automatically imported when the management server is installed and distributed to the agents. This management pack was expanded in [Figure 3-40](#).

Microsoft supplies management packs for most of its server products; see <http://www.microsoft.com/management/mma/catalog.aspx> for the current listing. They also ship 11 of the most commonly used management packs on the MOM 2005 product CD. If you are working with the MOM 2005 Workgroup edition, all of these are imported during setup.

For the MOM 2005 full edition of the product, you must import all management packs and reports manually. To start importing a management pack, first make sure that you have downloaded the most current version. Create a network share to store these management packs. In [Figure 4-1](#), it is called the File transfer folder. In this folder, create three subfolders: *CurrentMP*, *OldMP*, and *VendorSupplied*. Place all the downloaded management packs in the *VendorSupplied* folder. Each of the management pack executables creates a folder as it extracts. In the contents of the extracted management pack folder, each management pack file ends with the *.akm* extension, and reports have the same name but end with *Reports.xml*.

On one of the management servers, open the Administrator console and right-click the Management Packs node to bring up the context menu. Select Import/Export Management Packs. You can also select the Import/Export Management Packs hyperlink in the Results pane, but then you won't know how to manually navigate the consoles. This starts the Import/Export Management Packs wizard (point 1 in [Figure 4-1](#)). After the Welcome page, choose either "Import Management Packs and/or reports" or Export Management Packs (see [Figure 4-2](#)). You cannot choose to export report definitions through this tool. It can be done with the Report Utility (*rptutil.exe*) tool found in the root of the MOM 2005 installation directory (usually *C:\Program Files\Microsoft Operations Manager 2005*).

On the next page of the wizard, browse to the folder that contains the management pack file that you want to import. In this case, *\\MPTransferFolder\VendorSupplied* has been mapped to the Y: drive, as shown in [Figure 4-3](#). You are prompted for the type of import to perform: management packs only, reports only, or management packs and reports. Notice that the options for importing reports are disabled in this

Figure 4-2. Select to import a management pack

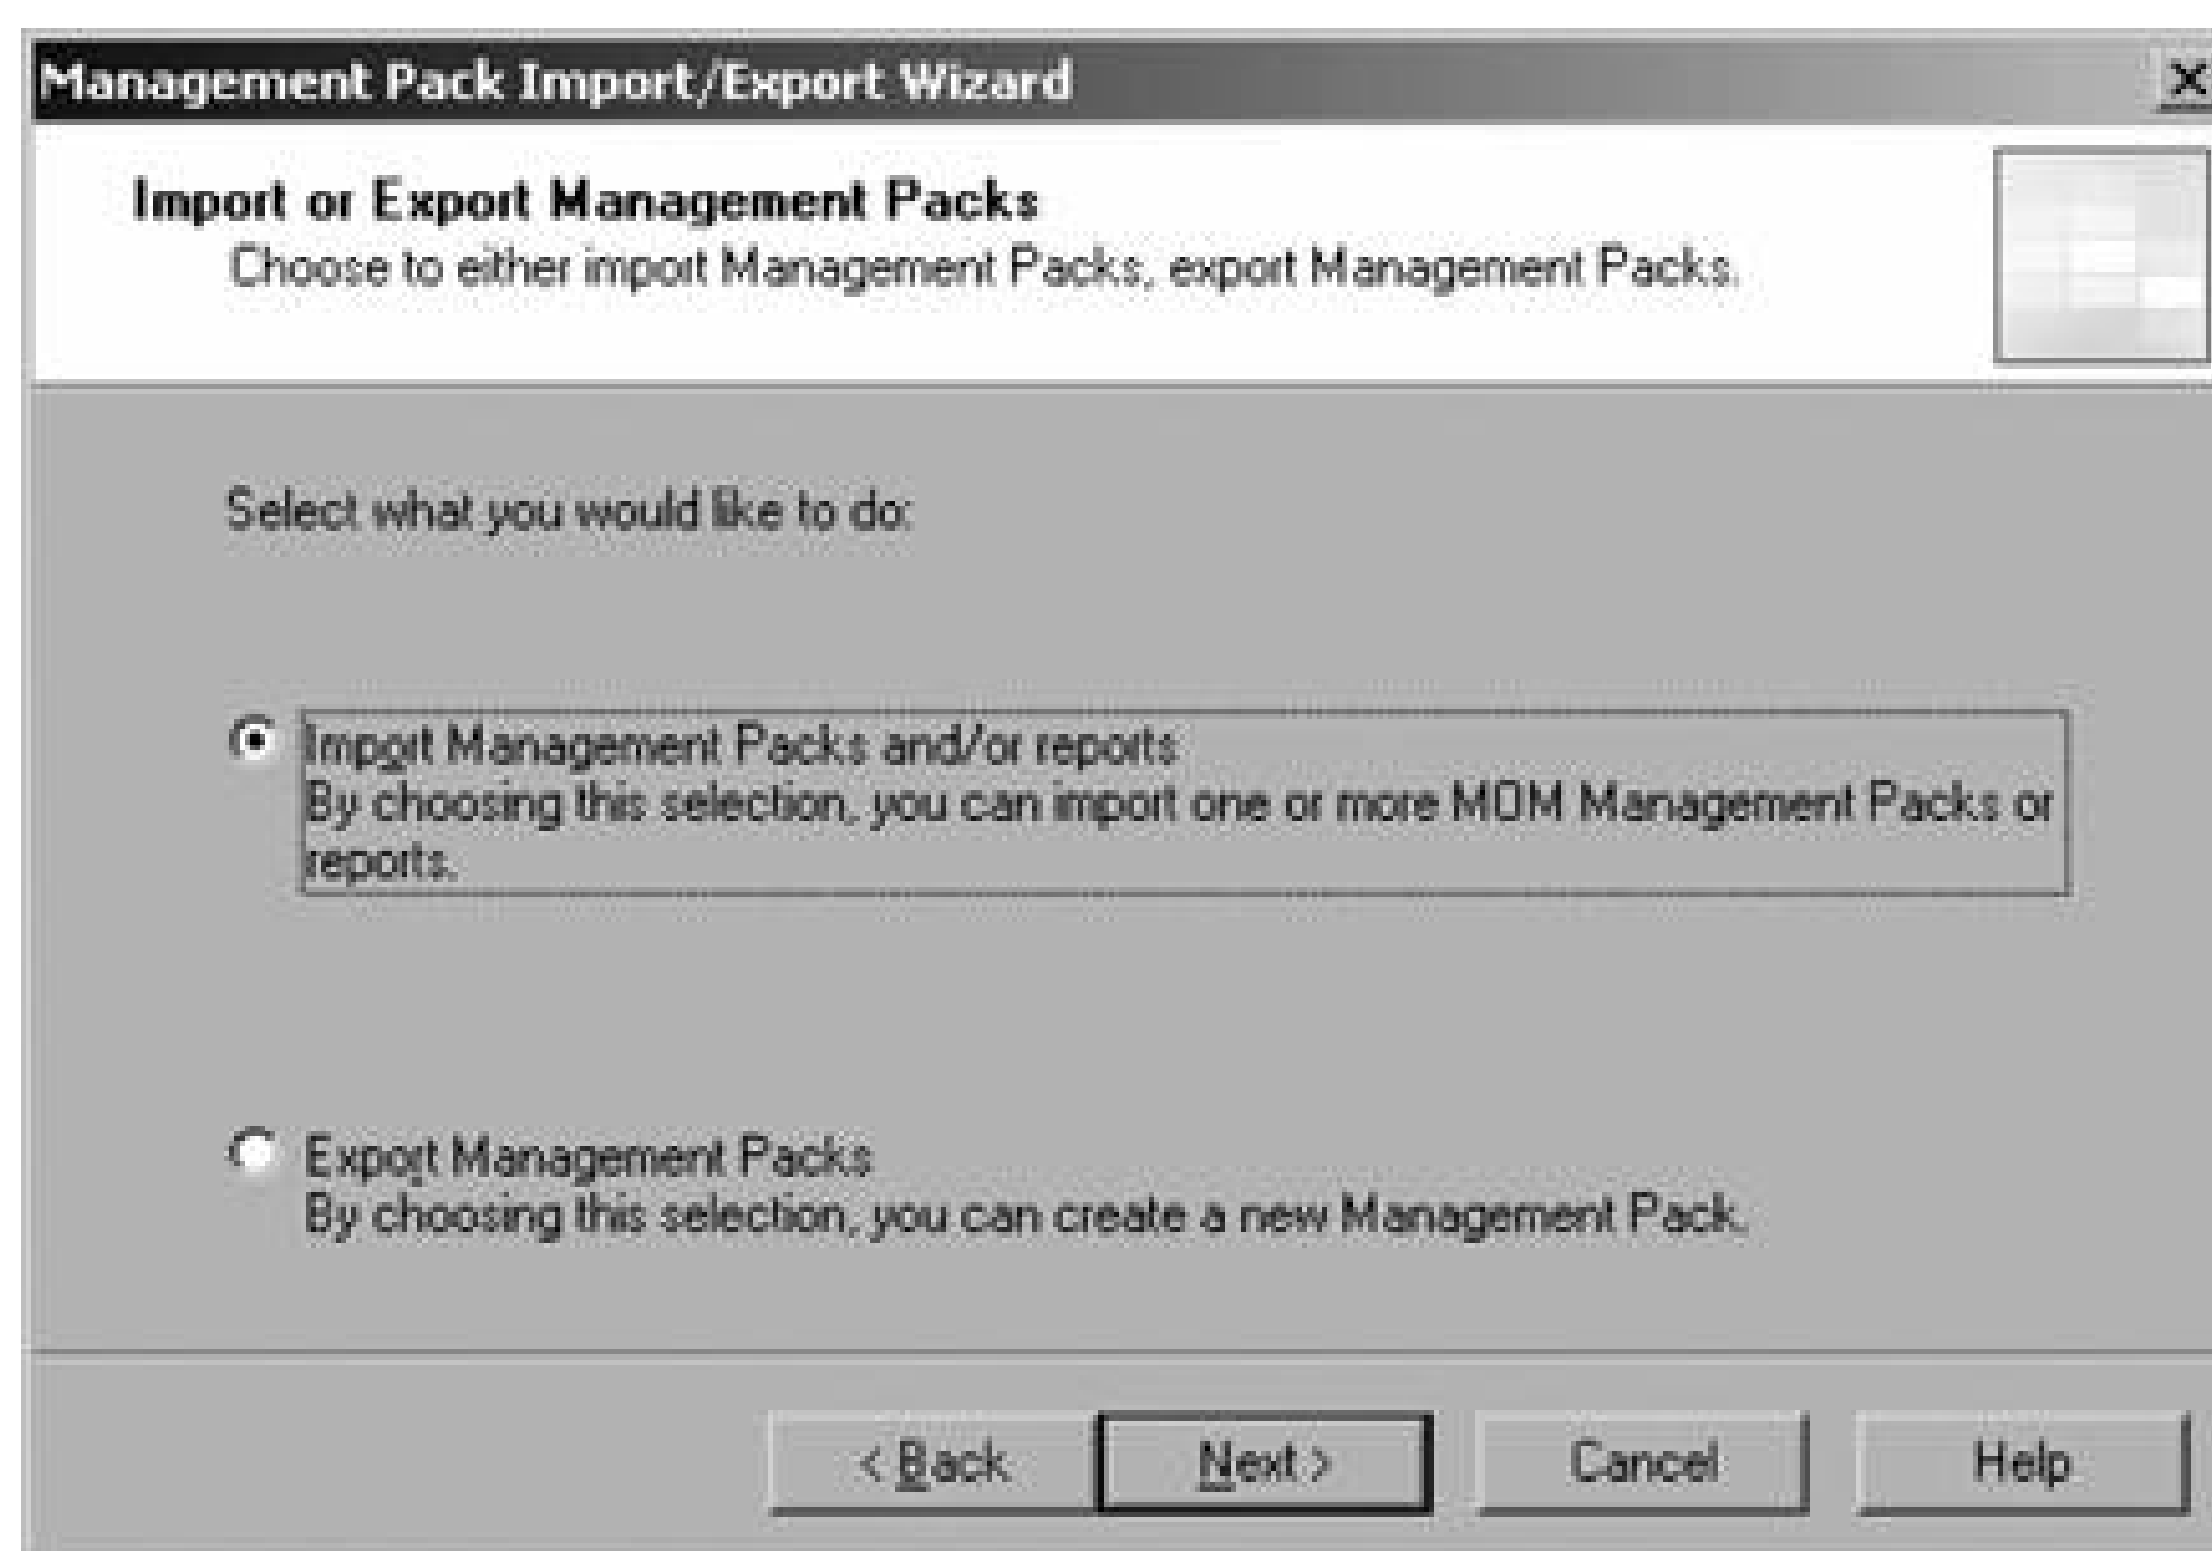


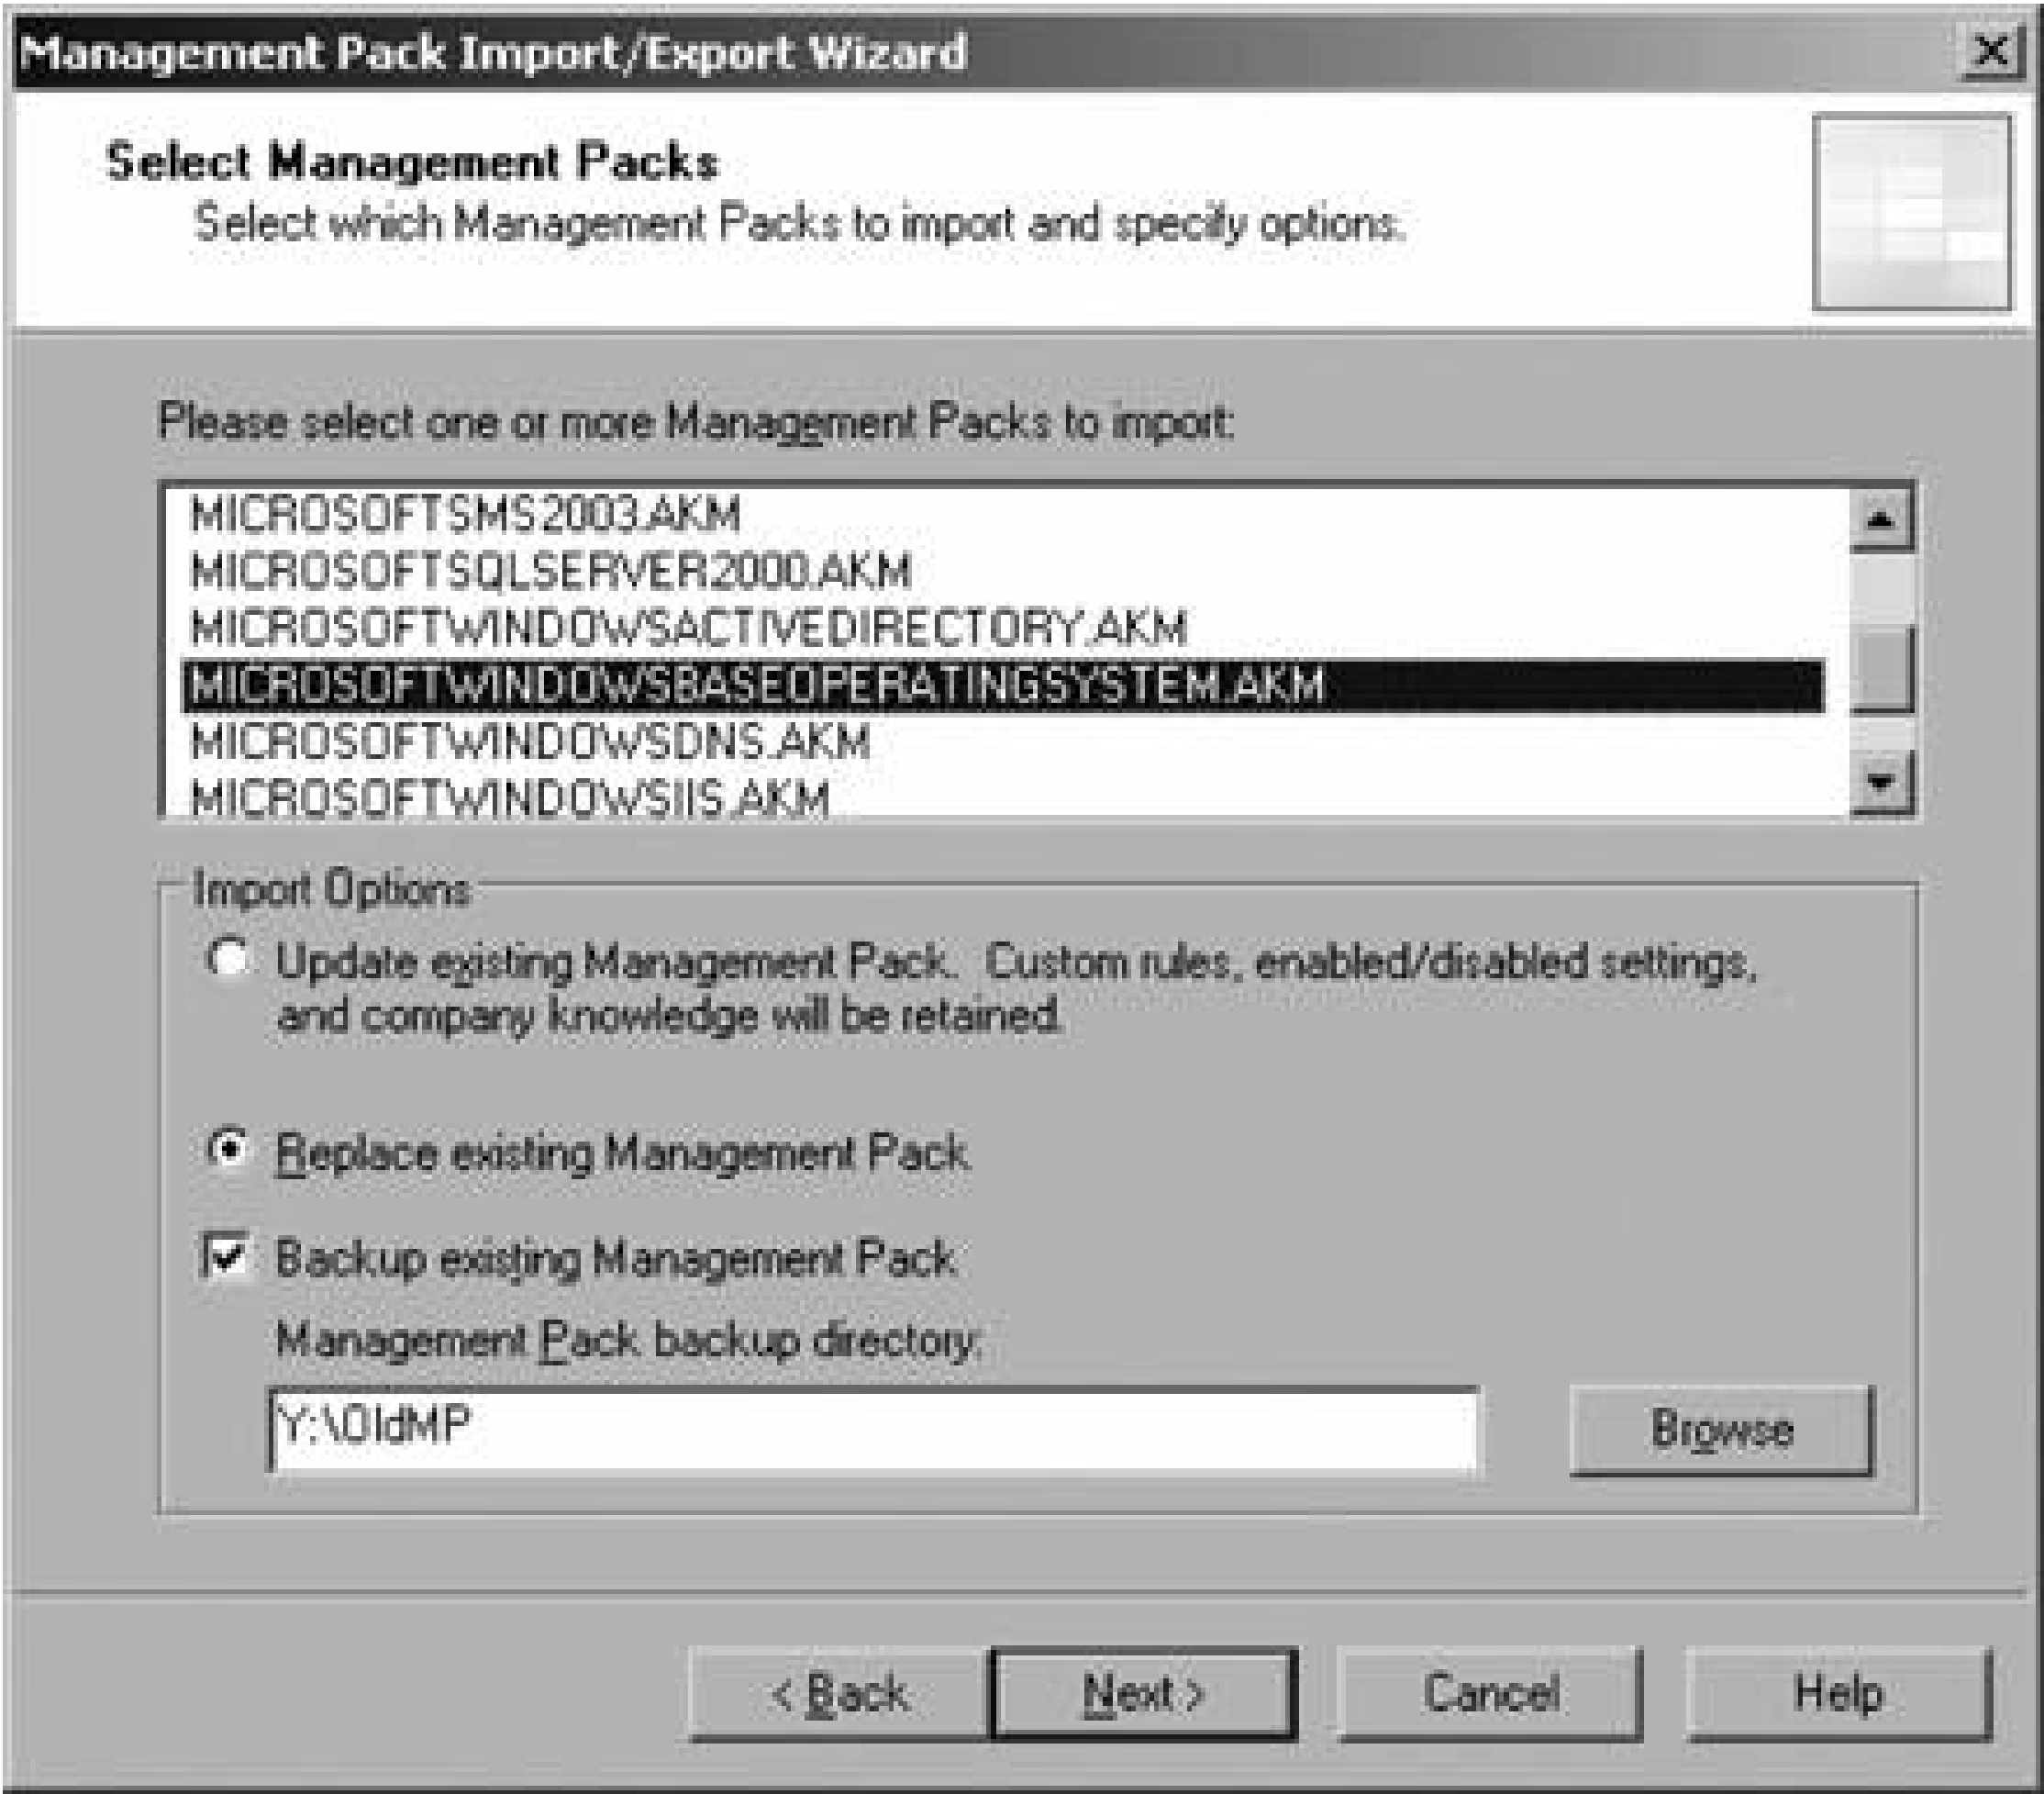
Figure 4-3. Because Reporting Services are not installed in this management group, only management packs can be imported



figure because the management group does not have MOM 2005 Reporting Services installed. If you install Reporting Services later, you will need to re-run the Import/Export wizard and select "Import reports only."

This next page of the wizard is critical to version control of management packs [Figure 4-4](#)). The top box lists all the management packs in the folder. [Figure 4-4](#) shows the base OS management pack selected. When the Import/Export wizard executes, it will search the operations database for the presence of this management pack by its globally unique ID (GUID), not by its name. Therefore, you can have two different management packs with identical names active in the same management group at the same time. Similarly, each rule in a management pack is identified by its GUID, not its name. This is a feature that can be made to work in your favor when you need to alter a vendor-produced management pack rule to create a custom rule.

Figure 4-4. The choices you make in this window are critical in the version control process

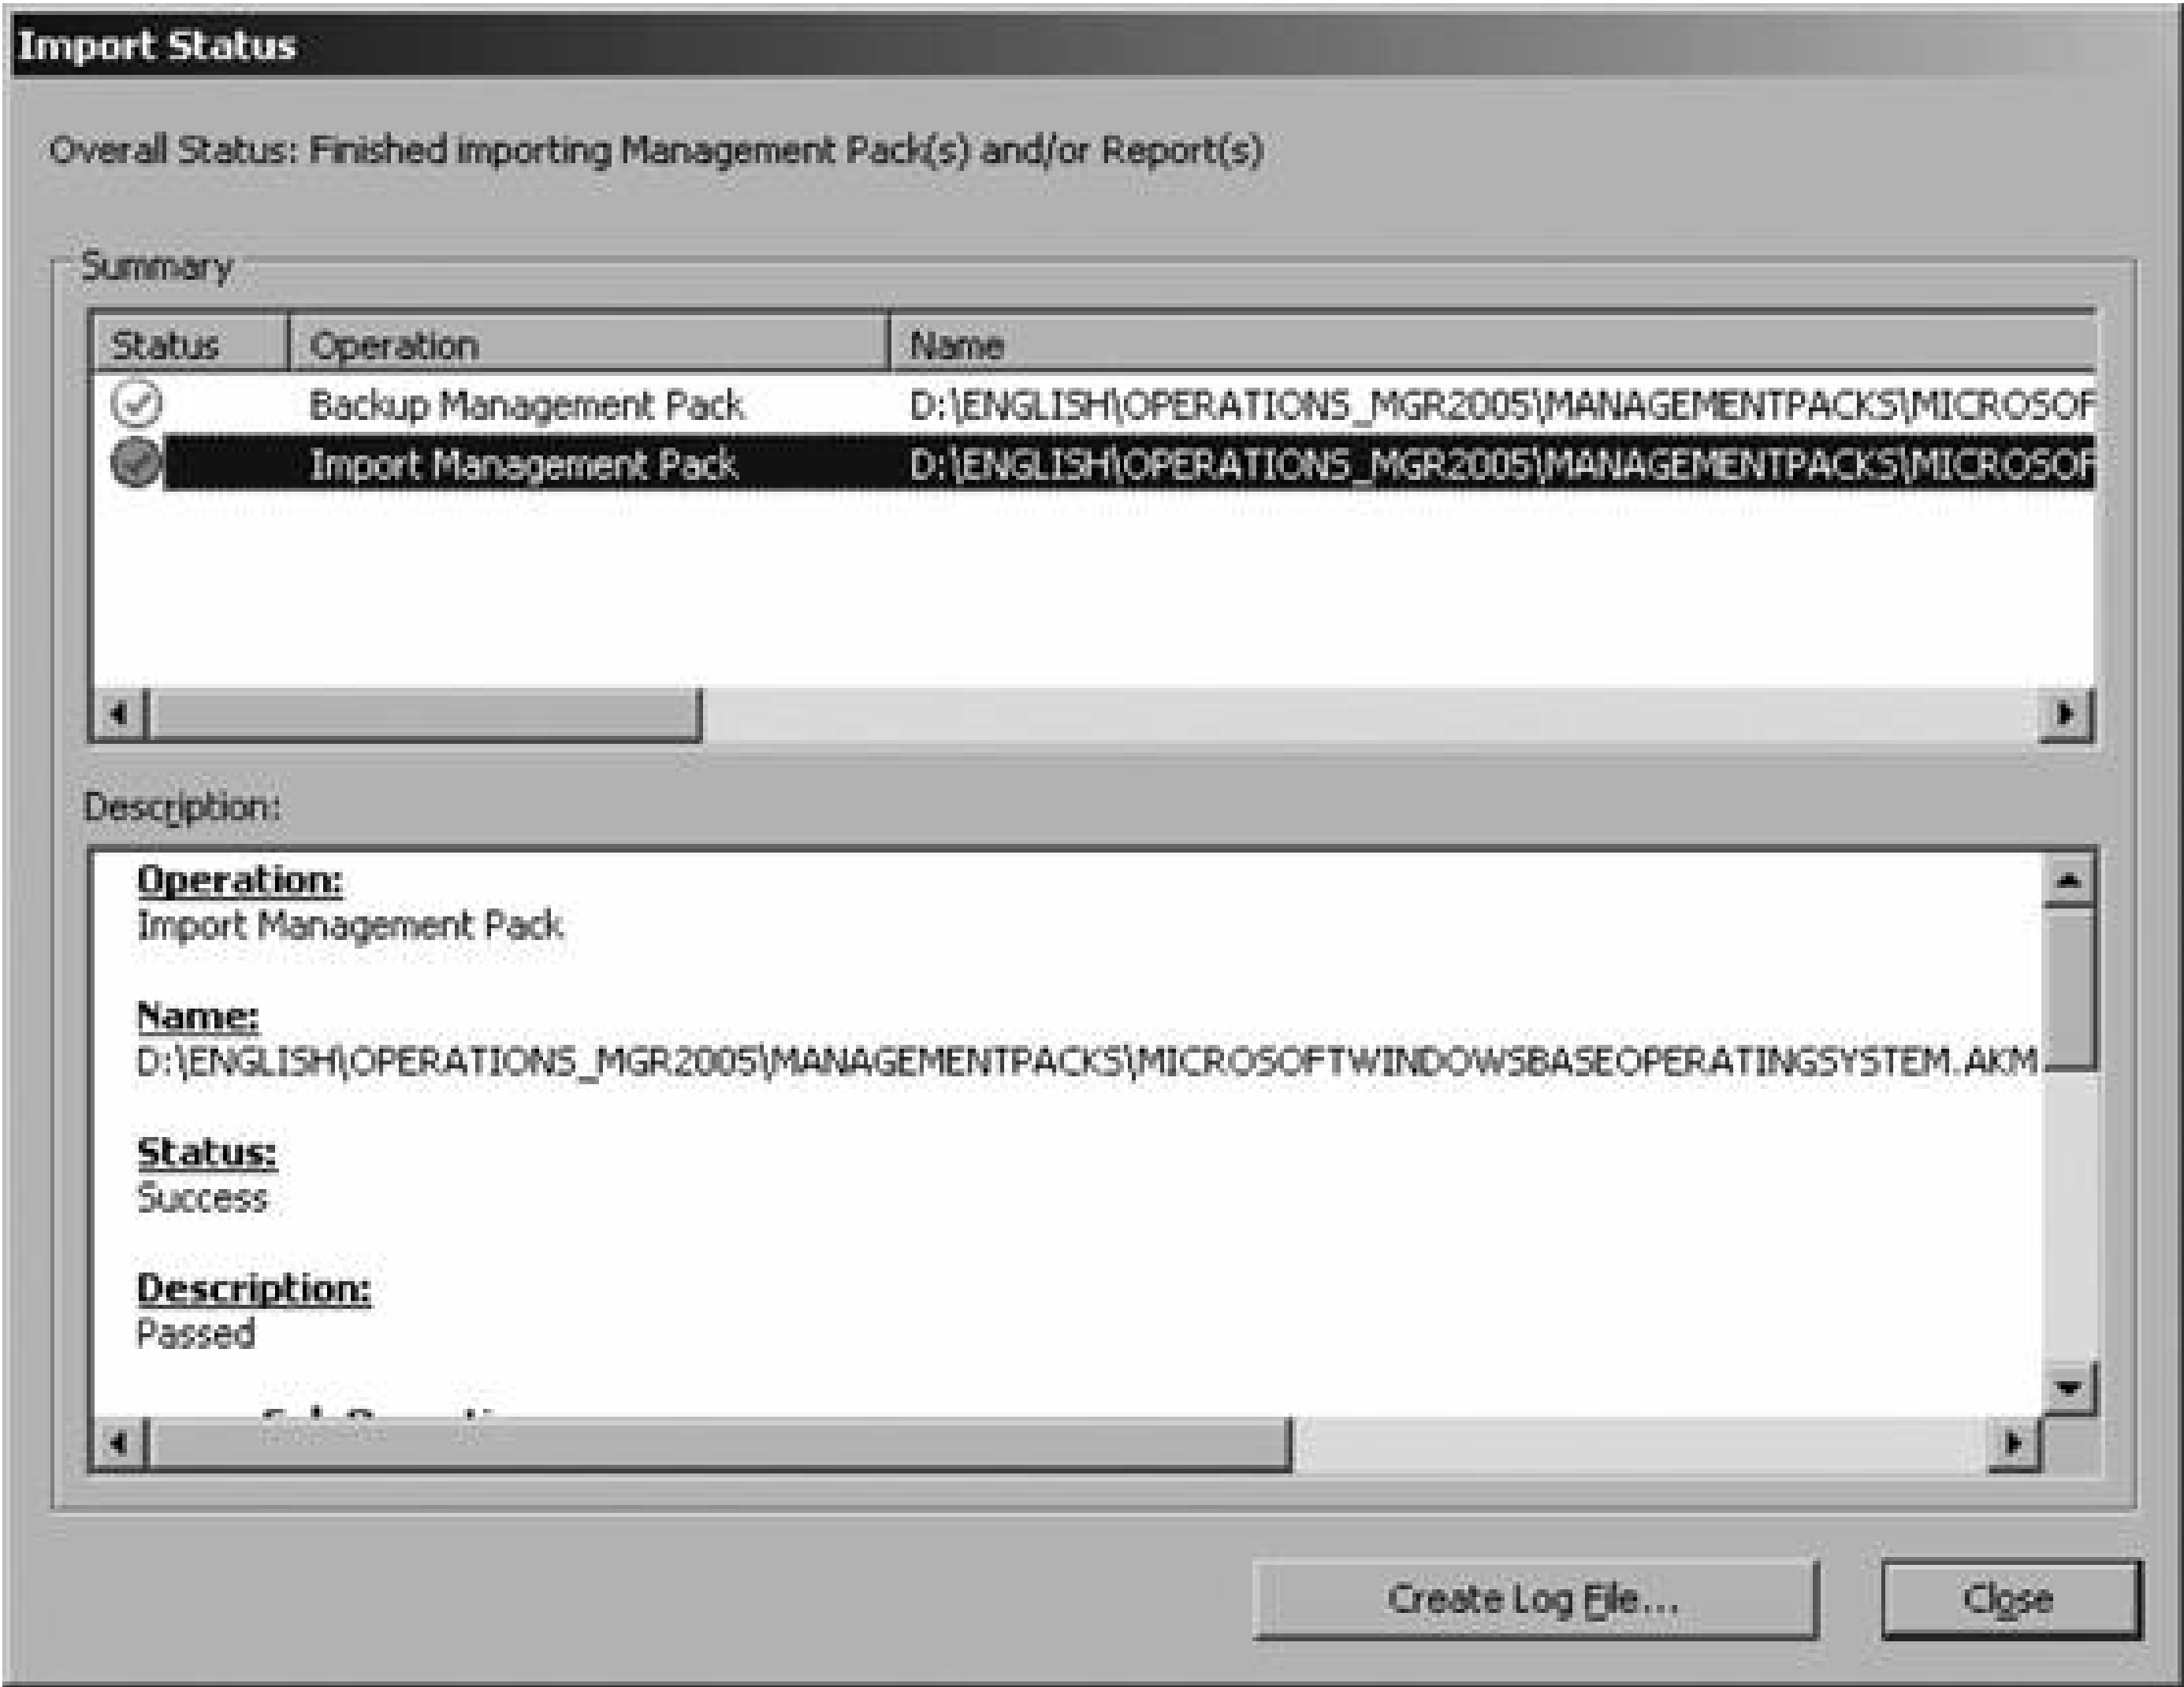


In the Import Options box, you have two choices: update the existing management pack or replace it, or back up the existing management pack. Your first choice depends whether you need to preserve existing data in a management pack or overwrite it completely. Unless the management pack is completely useless, you should always select to back it up. Save your backup to the *\\MPTransferFolder\OldMP* directory. The Import/Export tool will create the backup .*akm* file with a timestamp appended, allowing multiple versions of the .*akm* file to exist in the same folder at the same time. For example, this particular run of the Import/Export tool created a backup file named *MICROSOFTWINDOWSBASEOPERATINGSYSTEM\_02.23.05 21.27.52.akm*. This management pack is

being imported into a preproduction management group, so there is no existing data to preserve. To ensure that the management pack is in a known state after the import has completed, choose the Replace option.

After reviewing the Summary page, click through it. The Import/Export wizard presents you with an Import Status window for immediate feedback on the success or failure of the import process (see [Figure 4-5](#)).

Figure 4-5. Management pack Import Status window



If you wish, you can export the entire text of the import process using the Create Log File option. If you do, give it the same name and timestamp (*ManagementPackName\_MM.DD.YY.HH.MM.SS.akm*) as the backup file and place it in the *OldMP* directory.

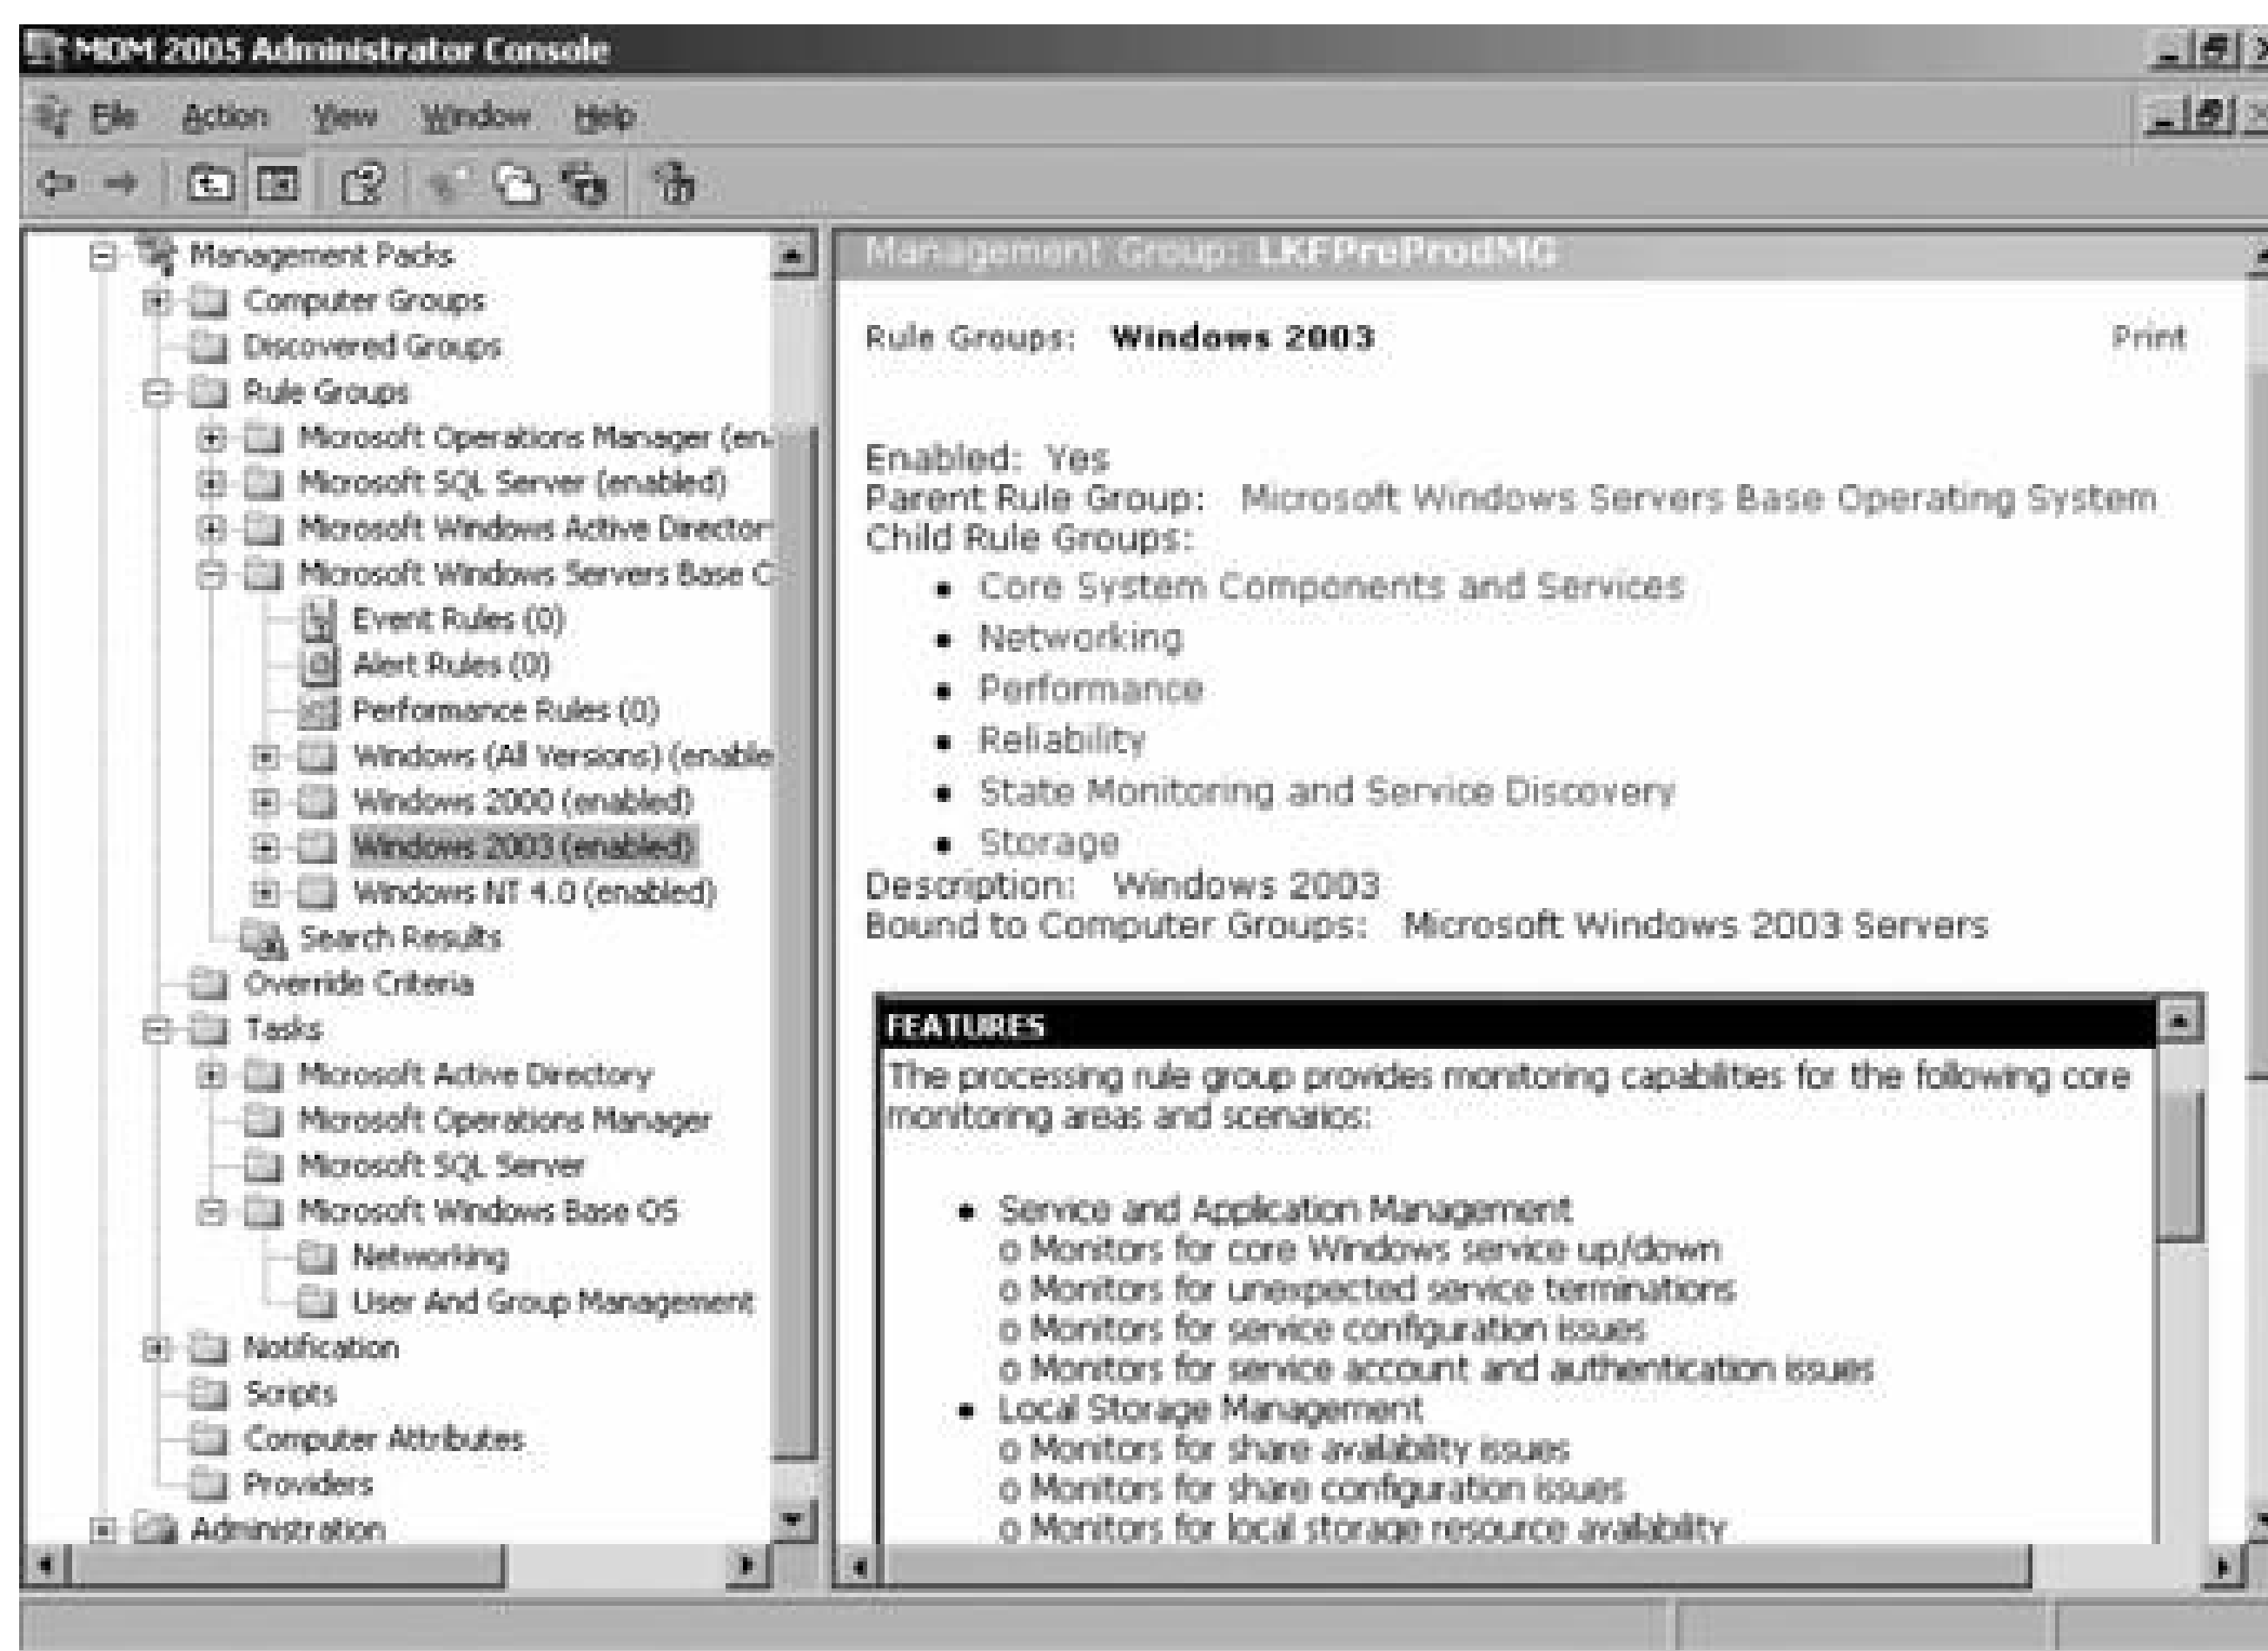
Referring back to [Figure 4-1](#), the management pack is now in the preproduction environment and is represented by the RAW MP (MP.1.0) object.

### 4.2.1. Management Pack Processing

In the Administrator console, several new objects now exist in the Management Packs node, as shown in [Figure 4-6](#):

- In the Computer Groups node, groups now exist that will be used to gather computers with common attributes together, such as the Microsoft Windows 2003 Servers group. Membership in a computer group is calculated based on a formula containing attributes that define the computer group on the discovered servers. When a match is found, the discovered computer is added to the computer group. Computer groups are used to target the application of processing rules.

Figure 4-6. Objects that are created in the Administrator console when a management pack is imported



- A rule group bearing the name of the management pack that was just imported. Rule groups contain three top-level rule types: event rules that monitor the event logs or that run on a timed basis, alert rules that can execute a specific notification for all the alerts raised in the rule group, and performance rules that can be used to collect performance monitor statistics and generate alerts if a threshold is crossed.
- Application-specific task definitions for use in the Operator console.
- Scripts that are called by event processing rules.
- Computer attribute entries that define data that is used in the computer group membership calculation formula. For example, the Microsoft Windows Base OS management pack creates an



attribute called Windows Current Version. This attribute defines the Windows registry path search to collect the running version of the OS on a managed computer.

- Providers, which are predefined sources of data, such as a specific performance monitor counter or a timed event, that are used by the rules.

When a new management pack is imported, the agents request configuration data from the management server. The management server responds with the agent settings (see [Chapter 3](#)), the list of additional management servers in the management group, and the registry key-based attributes to collect. The agent then returns the registry key information to the management server and the computer group membership is calculated.

For example, in the LKFPredProdMG, after the Windows Base OS management pack is imported, a computer attribute definition named Microsoft Windows Current Version returns the value found at *HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\WindowsNT\CurrentVersion\CurrentVersion*. When the membership for the computer group Microsoft Windows 2003 Servers is calculated, a computer is placed into that group if this formula is evaluated as true for that computer:

```
AttributeValue(Microsoft Windows Current Version)="5.2"
```

Each computer group is associated with one or more sets of rule groups. At the next agent configuration request interval (default is 1 minute), the management server responds to the agent configuration request with all the processing rules that apply to the computer group members.

Computer groups do not need to be created only for applications. They can be created based on geographic location; administrative responsibilities; or on a line of business application, such as SAP or PeopleSoft, that include multiple components. MOM carries out this registry-based computer attribute discovery on all monitored computers, which allows you to classify machines into broad groupings. Computer groups can be nested and individual computers can belong to more than one group. For example, all agent-managed Windows 2003 domain controllers are also in the Windows 2003 Servers, the Windows Server, and the MOM 2005 Agents computer groups.

Some of the rules that are pushed out to the agents are *service discovery rules*. These rules gather additional data about the service components, the relationships between computers, and the roles of computers. The data that service discovery gathers is used to monitor the health of components that support a service as well as other things.

For example, for an IIS server to be healthy, the WWW, SMTP, and FTP services need to be healthy. MOM must be aware of the relationship between a frontend Outlook Web Access (OWA) server and its backend mailbox server.

Service discovery data enables MOM 2005 to render diagrams that show the relationships between computers, to dynamically add and remove computers from computer groups, and to identify computers that can be targets of certain types of tasks. Much of the state roll-up process draws on the data that is provided by service discovery.

Service discovery is performed by scripts that are run by event rules at timed intervals. The timing for running the service discovery rules is predetermined by the management pack author and it is not something that should be changed. A timed event provider is a timer that can be used to trigger a rule, task, or response to occur at predefined time intervals.

So at this point, Leaky Faucet has imported management packs into the LKFPredProd management group. MOM has performed both the registry key-based attribute discovery and the more in-depth script-based service discovery, and the computer group membership has been calculated. All appropriate event, alert, and performance rules are being executed by the agents. Leaky Faucet is now at point 2 in the management pack life cycle workflow (see [Figure 4-1](#))--and they are ready to tweak the management pack processing rules to meet their needs.

## 4.2.2. Management Pack Tuning

The default configurations of the rules in a management pack represent what the management pack author defines as the health of that application. Microsoft puts considerable effort into authoring management packs that yield just the right amount of information, by tuning and tweaking the rules as best as possible for the application. Based on this, you may ask yourself, "Who am I to believe that I know more about what a healthy Exchange server or domain controller looks like than the vendor?" Why should you change any of the rules at all if that is how Microsoft says it should be?

The answer is simple: you must examine the rules from the perspective of applying them in *your* environment. You may choose to accept the default settings for some or even most of the rules, but keep in mind that Microsoft authored and tuned the management packs for a generic environment, not yours. To get the most out of any management pack, you will need to go through the rules and decide if they are:

### *Necessary for your environment*

Some management packs contain child rule groups for versions of the software that are not in your environment. For example, the Base OS management pack contains rule groups for Windows NT, 2000, and 2003. If you only have 2003 in your environment, then the NT and 2000 rules will get in the way.

### *Raising alerts of the appropriate severity*

Most event and performance threshold rules generate alerts and the alerts are created with a default severity. By default, it is the rule that monitors if a user cannot map a printer generate an error alert, but in your environment that may only be an informational alert.

### *Being seen by the right people*

Alerts appear in the Operator console, but they can also be forwarded to mailboxes, public folders, pagers, and any other communication channel that can be called from a command prompt or script.

### *Applied identically to all computers in the associated computer groups*

By default, all rules in a rule group are applied the same way to the computer groups that are associated with that rule group. However, you might have a subset of those computers to



which you want the rule applied in a different manner. Perhaps you don't want the rule applied at all, or perhaps there should be different matching criteria. For example, say you want MOM to generate an alert when the percent processor utilization on any server in the domain exceeds 75 percent for more than 10 minutes. But you don't want to be alerted when this condition occurs on a server that is being load tested, or in that case, you want a threshold of 90 percent. Overrides allow you to make that kind of change. This type of tuning is very difficult to perform in the preproduction environment because you don't have the same load on the managed computers. Extensive use of overrides should wait until the management pack is in production. For more information see the ["Overrides"](#) section later in this chapter.

Review the management pack guide for the management pack you are working with. Management pack guides contain information on the rules, scripts, and the overall configuration of the management pack. However, not all management packs have a guide, only the more complicated ones do.

Management packs will need to be tuned many times, especially when you are new to MOM. At this point, you are rough-tuning the rules and management packs. Once you have more experience with the data these rules produce in your environment, you can go through them again and refine them further. In this first pass, only consider if a rule or rule group should be enabled or disabled, if the alert is of the correct severity, and if the correct notifications are being sent out. The goal of this is to ensure that only the necessary rules are processed and that the data produced by the rules is what you want.

### 4.2.3. Types of Rules

There are three main rule categories: event rules, alert rules, and performance rules. There are also different flavors of event and performance rules.

#### *Event rule*

Monitors Windows event logs for the presence or absence of events. When a match is found, the event rule takes whatever action is defined in the responses tab for that rule.

#### *Alert/response rule*

This is the most generic type of event rule. It monitors a provider data source for a match, or it may use a timed event provider. A timed event provider is a timer that causes tasks or responses to occur at predefined intervals.

#### *Filter rule*

Filter event rules do one of three things:

##### *A pre-filter rule*



Watches for an event and deletes it; the event never makes it to the operations database and no alert or action can be taken based on that event.

### *A database filter*

Keeps the event data for use in correlating other data. The sum of this event data may raise an alert or cause other action to be taken, but the filter does not insert the individual event into the operations database.

### *A conditional filter*

Causes an event to be disregarded unless another rule would process that event.

### *Detect missing event rule*

This type of rule alerts you when something did not happen in the time frame that it was supposed to. For example, say backups generate a 123456 event in the application event log when they complete successfully and your backups should finish between 3 a.m. and 6 a.m. You can configure this rule to look for the 123456 event in the application log between the hours of 3 a.m. and 6 a.m. and raise an alert if the event is not there.

### *Consolidate similar events rule*

This event rule looks for multiple occurrences of similar events (you define what similar is) occurring within a given time frame on one computer and generates a summary in a single event.

### *Collection event rule*

Use this rule to collect events from a certain source and store them in the operations database for later examination.

### *Performance rule*

This rule always watches performance monitor counter providers for data.

### *Sample performance data*

This type of rule is much like the event collection rule. It collects performance monitor data for later use.

### *Compare performance data*

This is the performance threshold rule. The MOM agent compares samples of performance monitor data to set values (thresholds) and raises an alert when the sample shows that the counter has crossed the preset threshold.

There is only one type of alert rule and oddly enough, it doesn't create alerts. Alert rules are used to provide a common response for a number of event or performance rules. For example, in the Microsoft SQL Server\SQL Server 2000\SQL Server 2000 Event Collection\SQL Server Agent rule group there is an alert rule called Critical Error or higher - Database Administrators. This alert rule fires when a critical error alert or greater is generated by any of the rules on the SQL Server agent rule group. This alert rule runs a notification response to the database administrators notification group. Using alert rules saves you from configuring a response for every event and performance rule that generates an alert.

The examination and tuning process can be tedious and is made all the harder if there is no management pack guide for the management pack you are evaluating. Always check to see if there is a management pack guide. See the "[Additional Management Pack Configuration](#)" section later in this chapter. The management pack guides outline all the rules of a management pack, including the rule severity, the path to the rule, whether the rule is enabled or disabled, and, in some cases, the response.

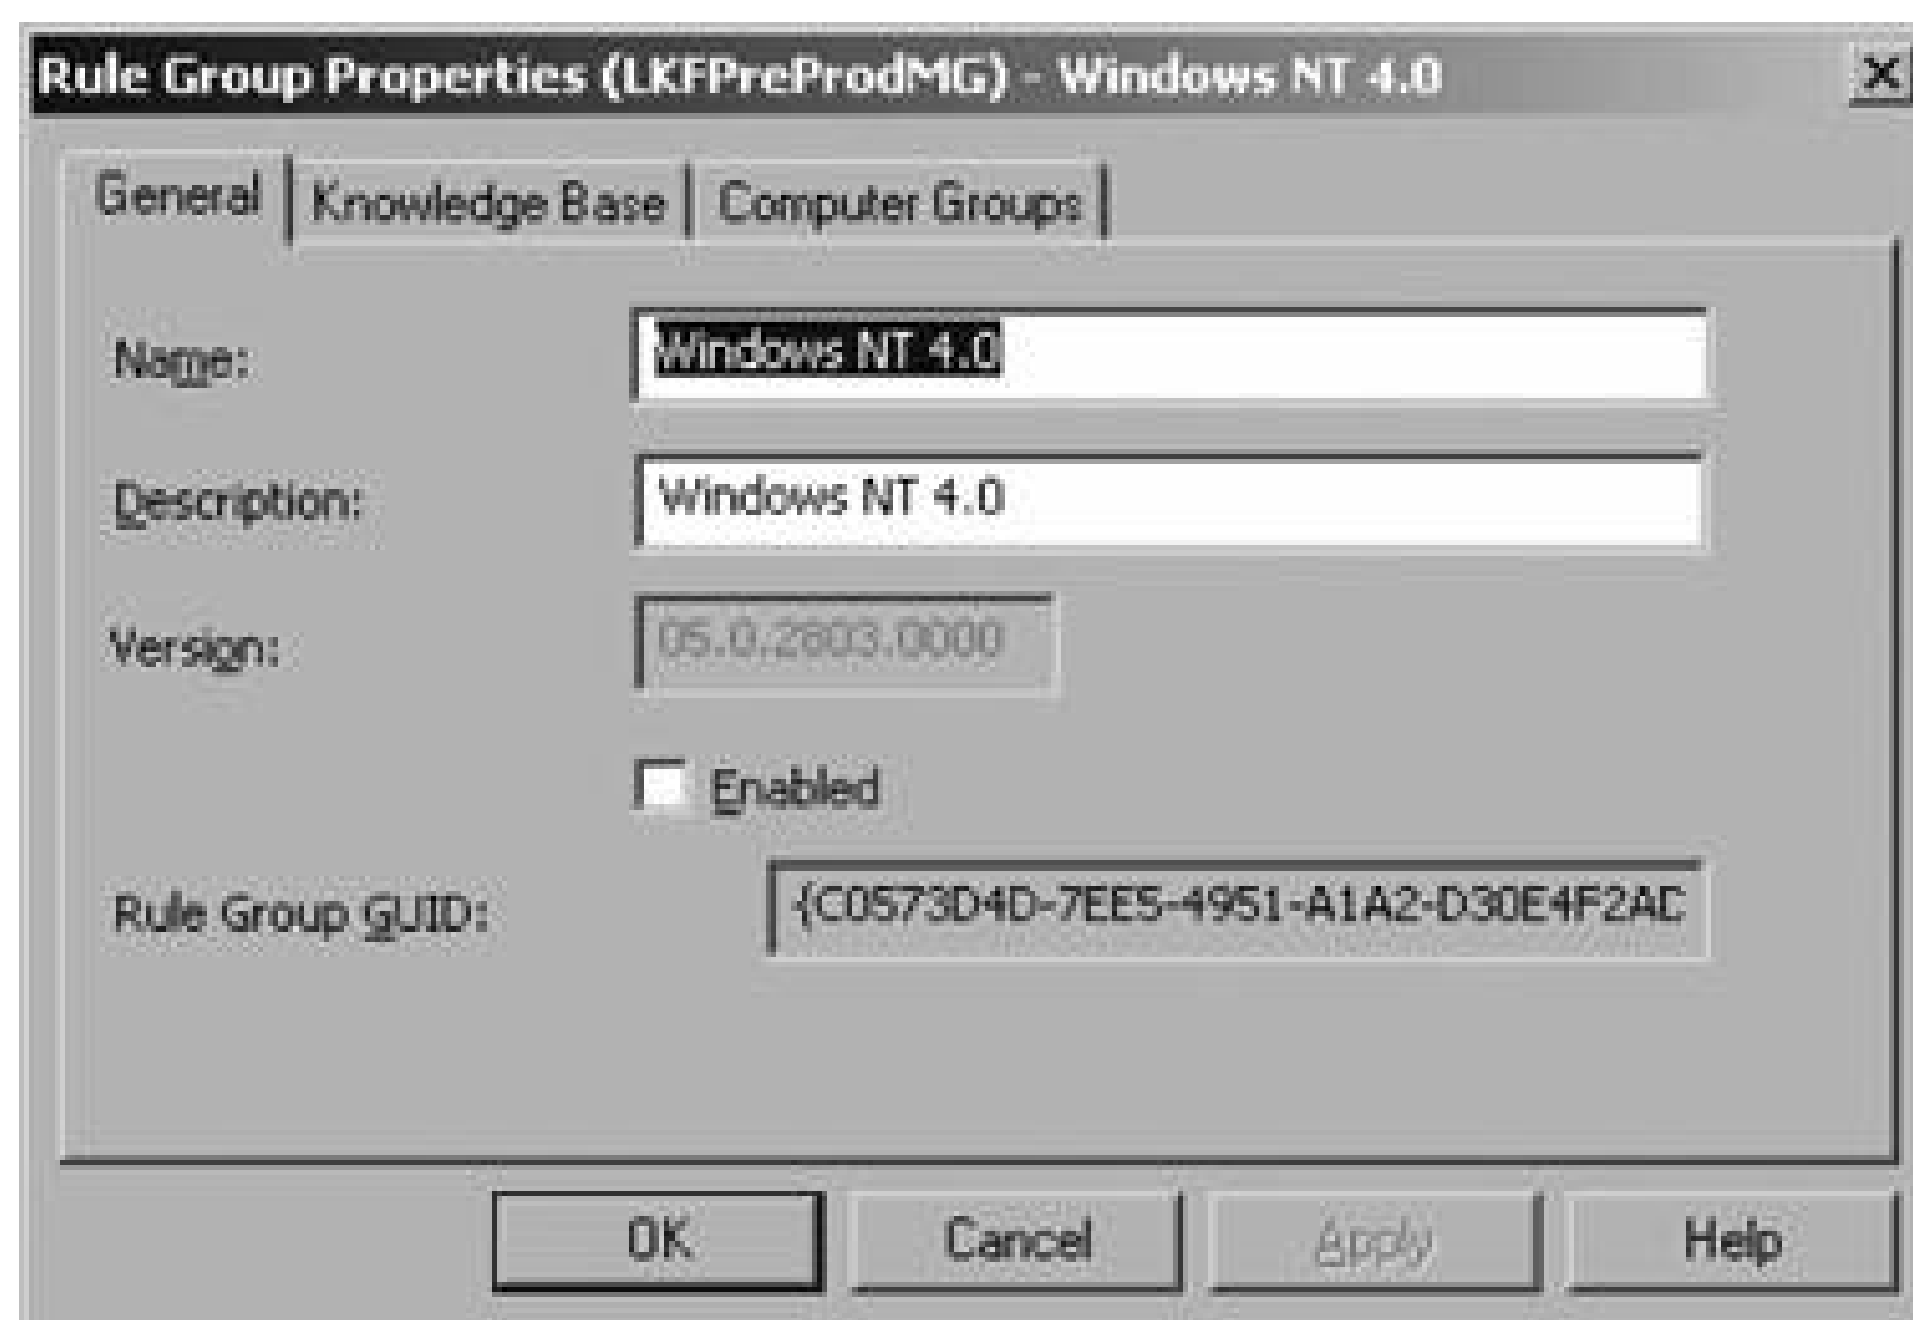
You will have to go through the rules, one at a time, in the Administrator console. Consider the following as you go through the rough-tuning process.

#### 4.2.3.1. Rule group application

Some rule groups contain rules for different versions of an application or OS that may not even be running in your environment. By disabling them, the processing load is reduced on the management server and the agent. The unneeded rule group could also be deleted, but that is a little drastic. If you ever need the functionality supplied by that rule group, you can simply re-enable it.

For example, the Base OS management pack contains rule groups for Windows NT 4.0. If there is no NT 4.0 server in your environment, disable the rule group and all child rule groups. To do this, bring up the context menu for that rule group, select Properties and then clear the Enabled checkbox ([Figure 4-7](#)).

Figure 4-7. Disabling a rule group by clearing the Enabled checkbox



Make sure you go to the Knowledge Base tab and edit the Company Knowledge base to include what modification was made, why it was made, when it was made, what environment it was made in, and who made it. This can be as simple as "This rule group was disabled in the preproduction environment because there are no NT4.0 servers to be monitored. 2/26/05, CJF." When a rule group has been disabled, its folder icon is disabled [\(Figure 4-8\)](#).

#### 4.2.3.2. Individual rule application

Once you have disabled the obviously unnecessary rule groups, go through the remaining rules and repeat the same process. There will likely be rules that you don't want to be active.

Figure 4-8. Disabled rule groups

For example, in the Leaky Faucet environment, Windows servers are patched via SMS, so there is no need to use the Automatic Update Service. Consequently, communication with that site is blocked at the firewall. When Max reviews the Windows Server Base Operating System management pack he discovers a rule called "Windows is unable to connect to the Automatic Update service" in the Core System Components and Services rule group. Because of the environmental configuration, this rule will raise an alert every time the automatic update service attempts to connect to the web site, which



it will do once a day unless it is disabled. To prevent generating unnecessary alerts from servers whose Automatic Update Service is still running (perhaps by mistake), Max disables this individual rule.

To disable the individual rules, clear the "This rule is enabled" checkbox on the General tab of the Rule Properties ([Figure 4-9](#)). Be sure to make note of the what, why, when, where, and who of the modification. When a rule is in a disabled state, a red X appears through the rule icon.

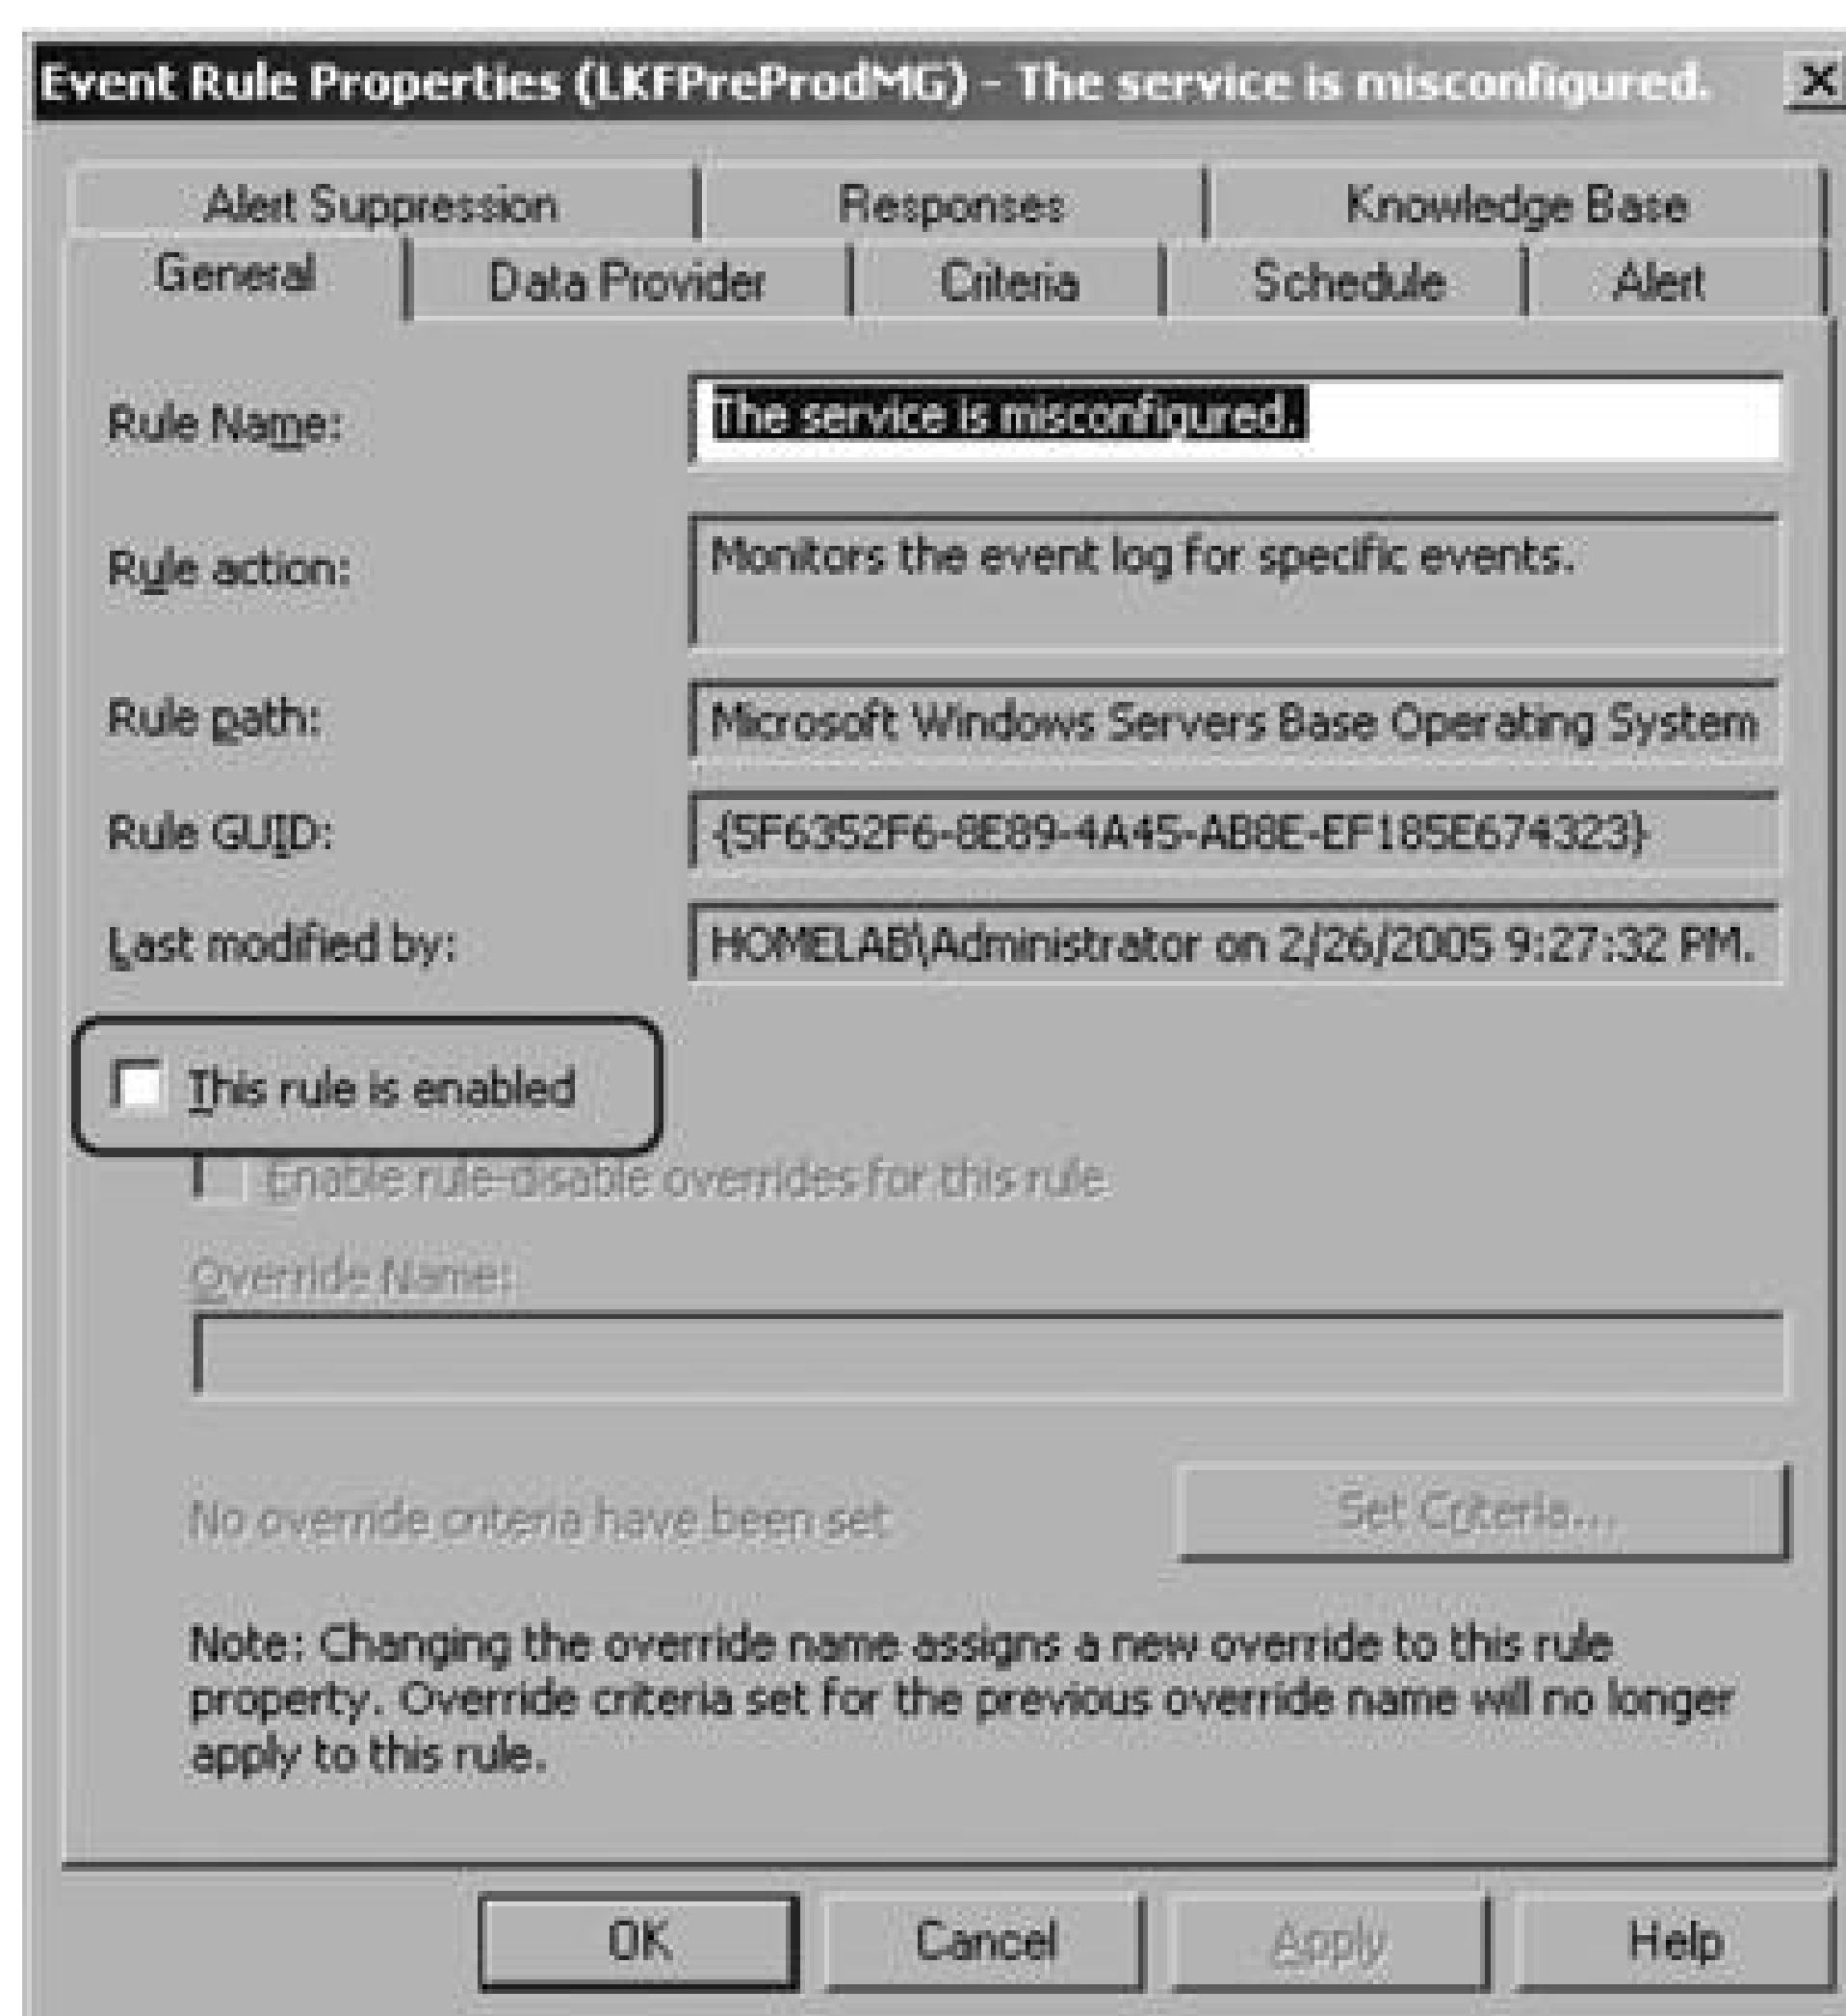
### 4.2.3.3. Overriding rule application

An enabled rule will be applied to all computers that are in the associated computer groups. Disabled rules are not applied to those computers. There will likely be cases when you want a rule applied to some computers but not others in the same computer group. To do this, you can use the "Enable rule-disable overrides for this rule" feature.

On the General tab of the Rule Properties tab, there is a checkbox for enabling overrides (see [Figure 4-10](#)) just below the "This rule is enabled" checkbox. An override is a way of allowing an exception to a rule without having to create an additional rule. For example, for the event rule shown in [Figure 4-10](#), by enabling the "Enable rule-disable overrides for this rule," the rule for a specific computer can be disabled, even if that computer is a member of the computer group that the rule is applied to. [Figure 4-10](#) shows the "Windows is unable to connect to the Automatic Updates service" rule.

Let's say that Leaky Faucet needs to keep this rule enabled because SMS has not been deployed to the remote site servers and the remote site administrators have implemented Software Update Services (SUS). However, the servers that Max and his team administer are using SMS. Rather than creating separate rule groups and computer groups, rule-disable overrides are enabled for this rule (point 1 in [Figure 4-10](#)). Once enabled, Max uses the Set Criteria button (point 2 in [Figure 4-10](#)), which defines a set of computers (his) that this rule is disabled for. This leaves the rule active for the remote site servers.

Figure 4-9. A disabled rule



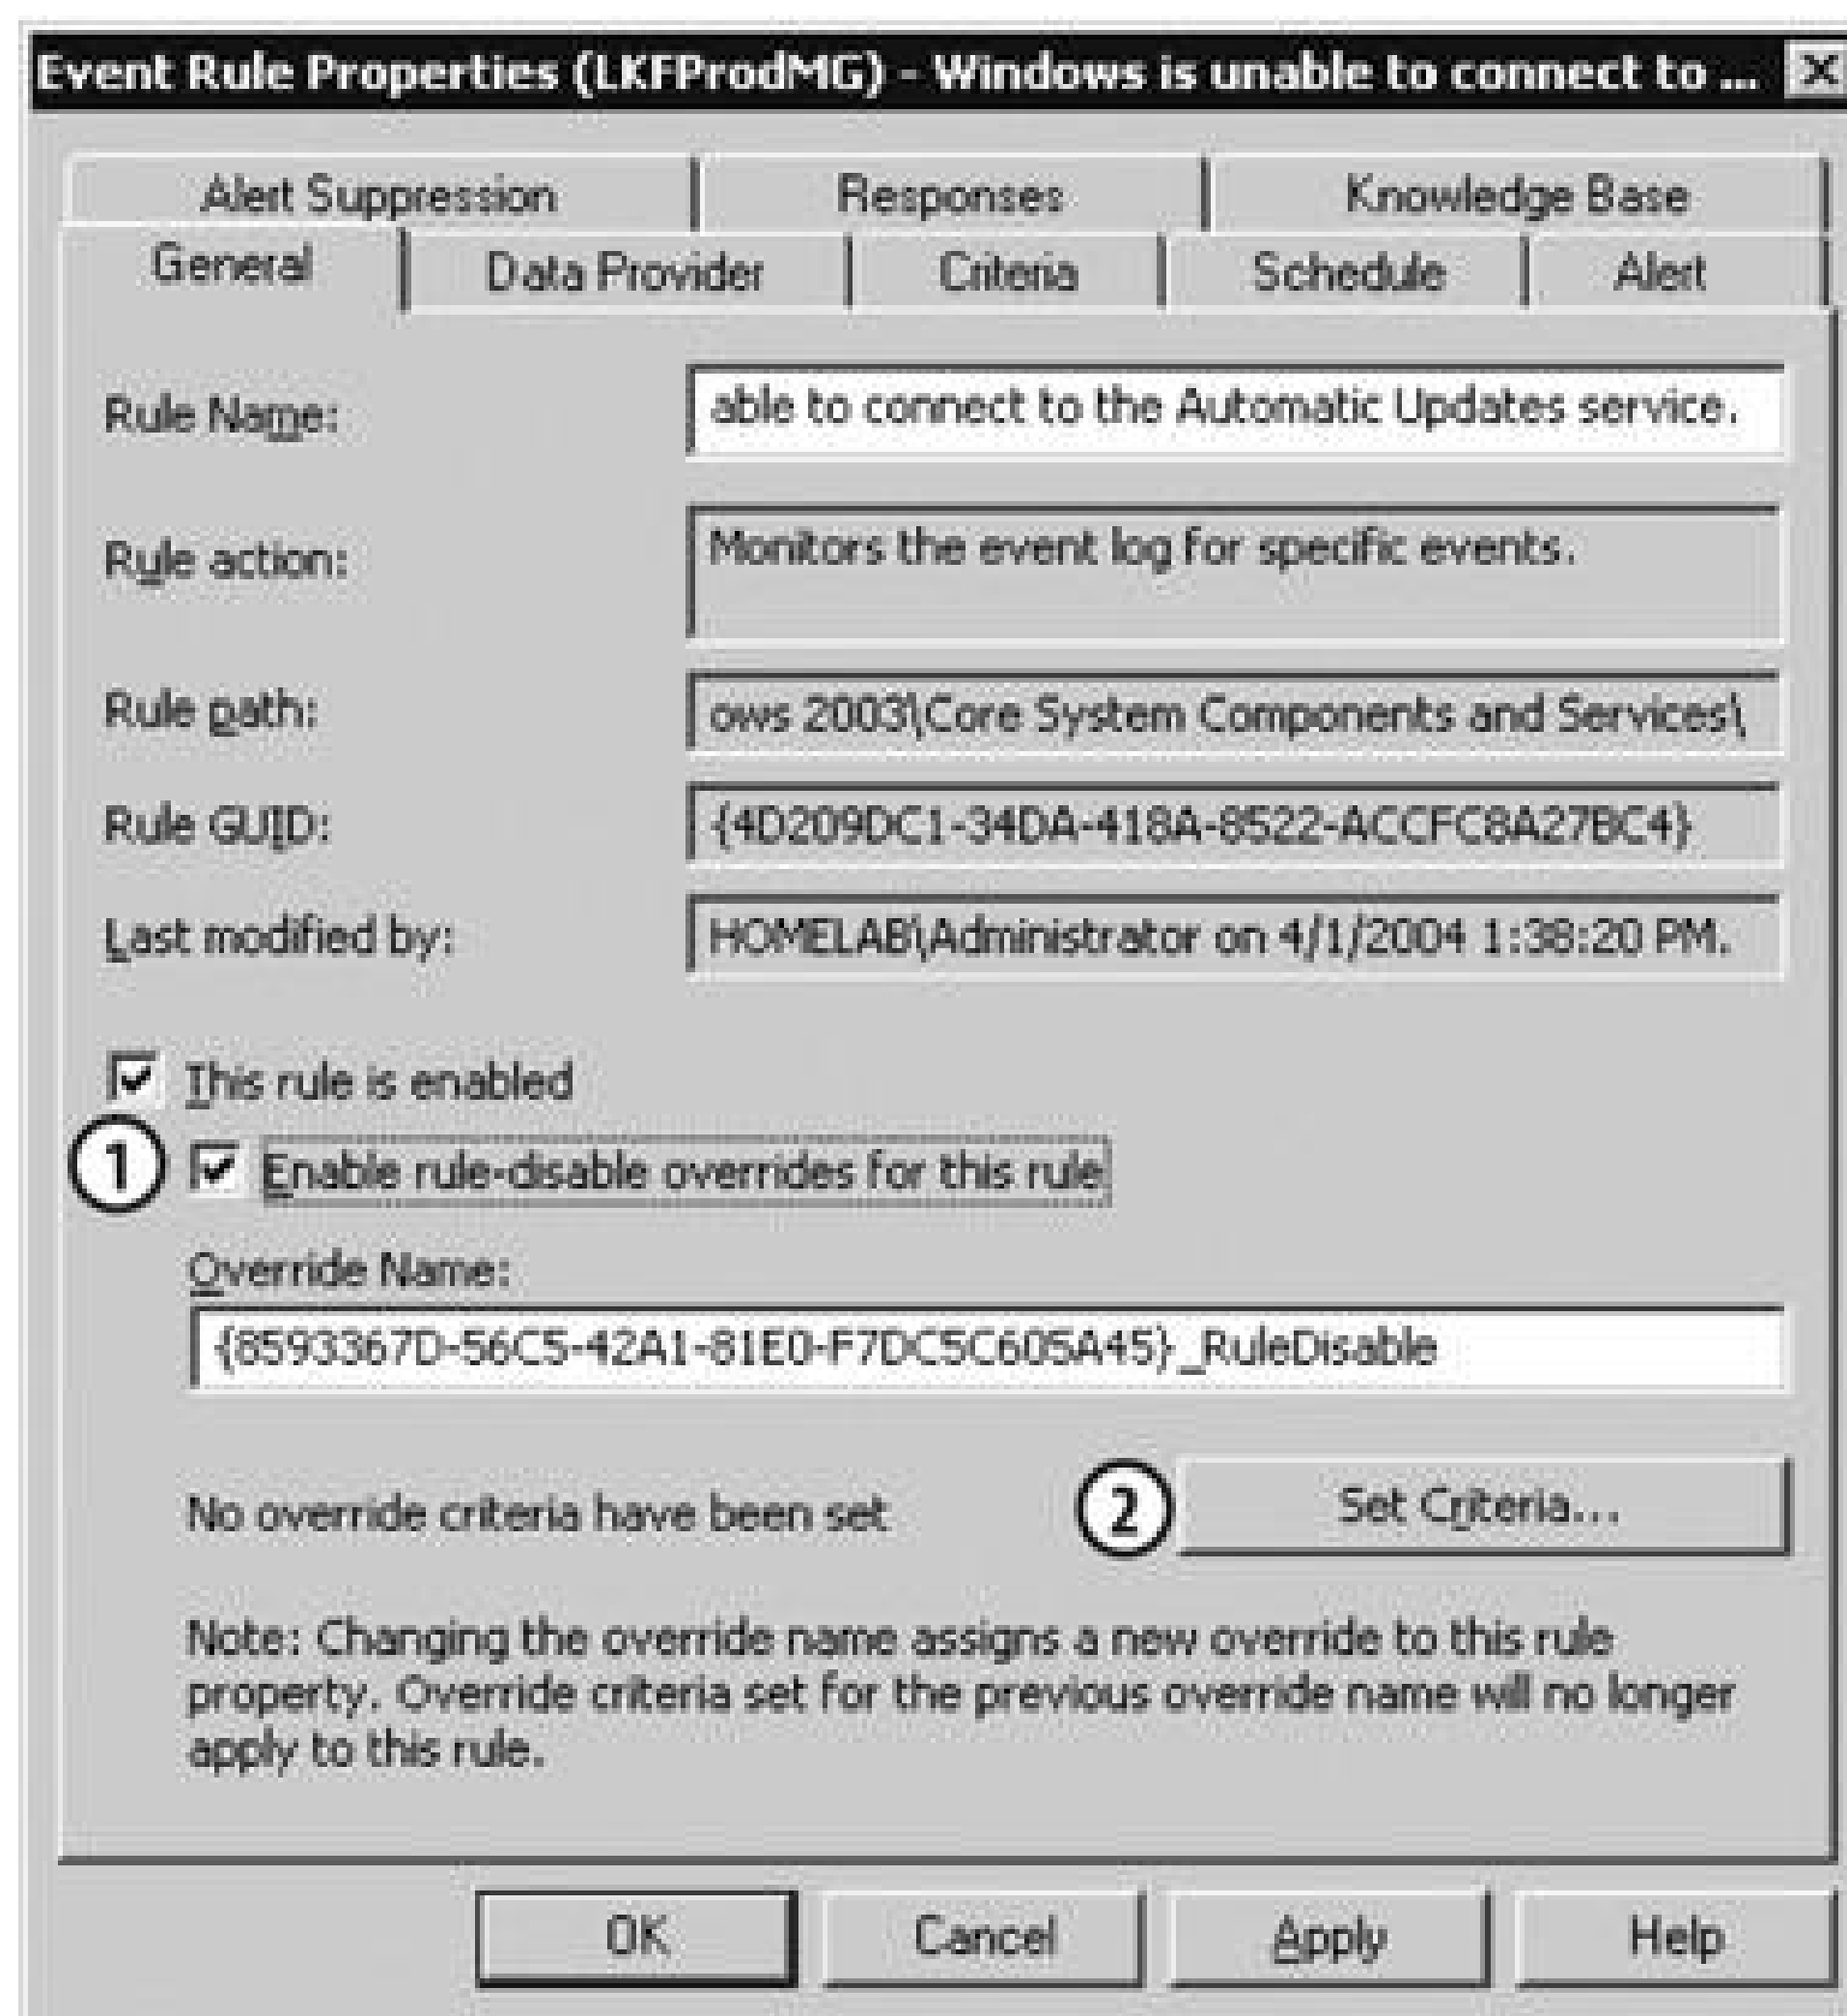
The defined criteria are simply a pairing of computers and/or computer groups with an enable/disable flag.

There are four types of overrides: rule disable, performance threshold, script parameter, and state alert severity. I will cover these more in the discussion on tuning management packs once they are in production (see "[Transfer the Management Pack to Production](#)" later in this chapter).

#### 4.2.3.4. Choosing alert severity

Not all rules generate alerts. For those that do, be sure to check the severity setting on the Alert tab of the Rule Properties. Make sure that the severity of the alert is what you want it to be by using the default as guidance. The management pack authors assigned the default severity when it was developed. So, if the generation of any particular event returns a critical error in any monitored application, you shouldn't change the alert severity to "information alert." If a change to the default severity level is necessary for your environment, do not make it here. Instead, bring up the context menu, copy the rule, and paste the new rule into the same rule group as the original. This will create a copy of the rule with the rule name prepended with "Copy of." This rule will have a different GUID, but otherwise it is identical to the original. Rename the rule by removing the "Copy of" prefix and replacing it with some specific tag that identifies it as a company-modified rule.

Figure 4-10. Enabling rule disable overrides

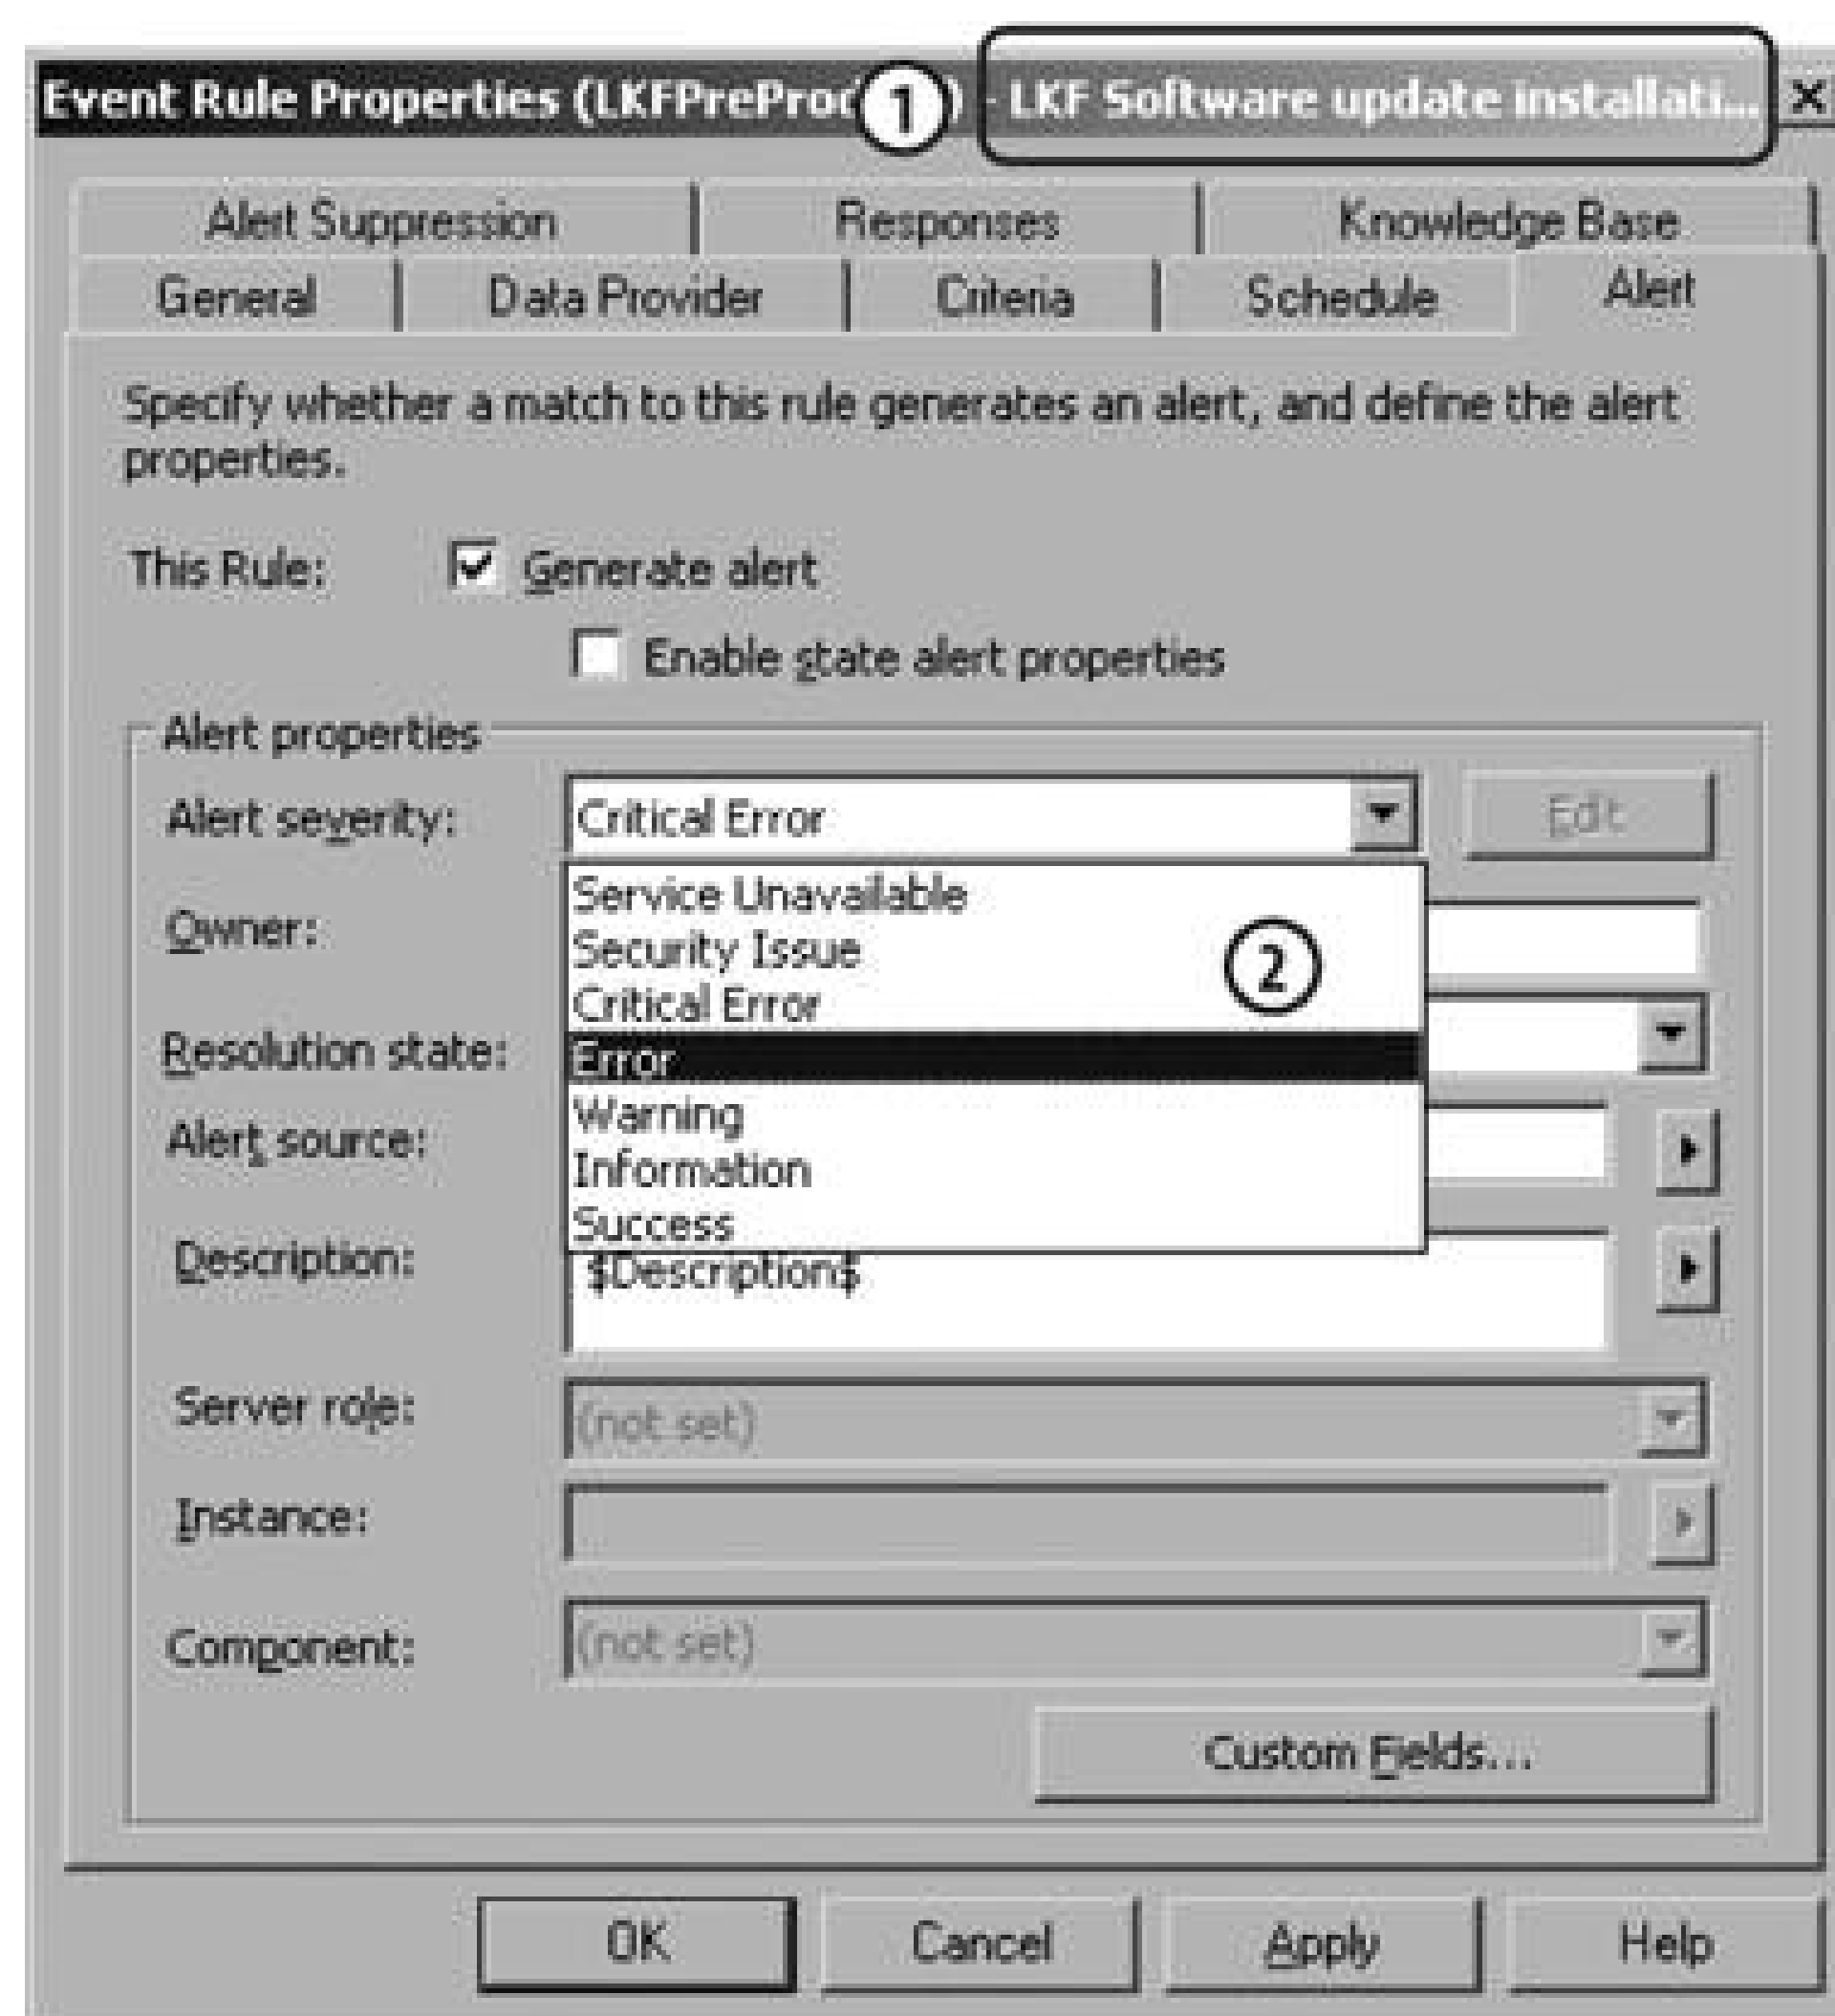


For example, in the Leaky Faucet environment, a copy of the "Software update installation failed" event rule has been made and prepended with the LKFPreProd tag (point 1 in [Figure 4-11](#)). The original rule is then disabled so that it will not be distributed to the agents. Adjustments must be made to the company-tagged version of the rule; in this case the alert severity is changed from Critical Error to Error (point 2 in [Figure 4-11](#)).

This process will protect your customized rules when you have to update the existing management pack, since user-modified rules are not overwritten by this type of import. Also, by keeping the rule in the same rule group, you take advantage of the pre-existing rule-group-to-computer-group association, which ensures the modified rule will be executed by the correct computers. This is where having two identical rules in the same rule group comes in handy.

Figure 4-11. User-modified event rule



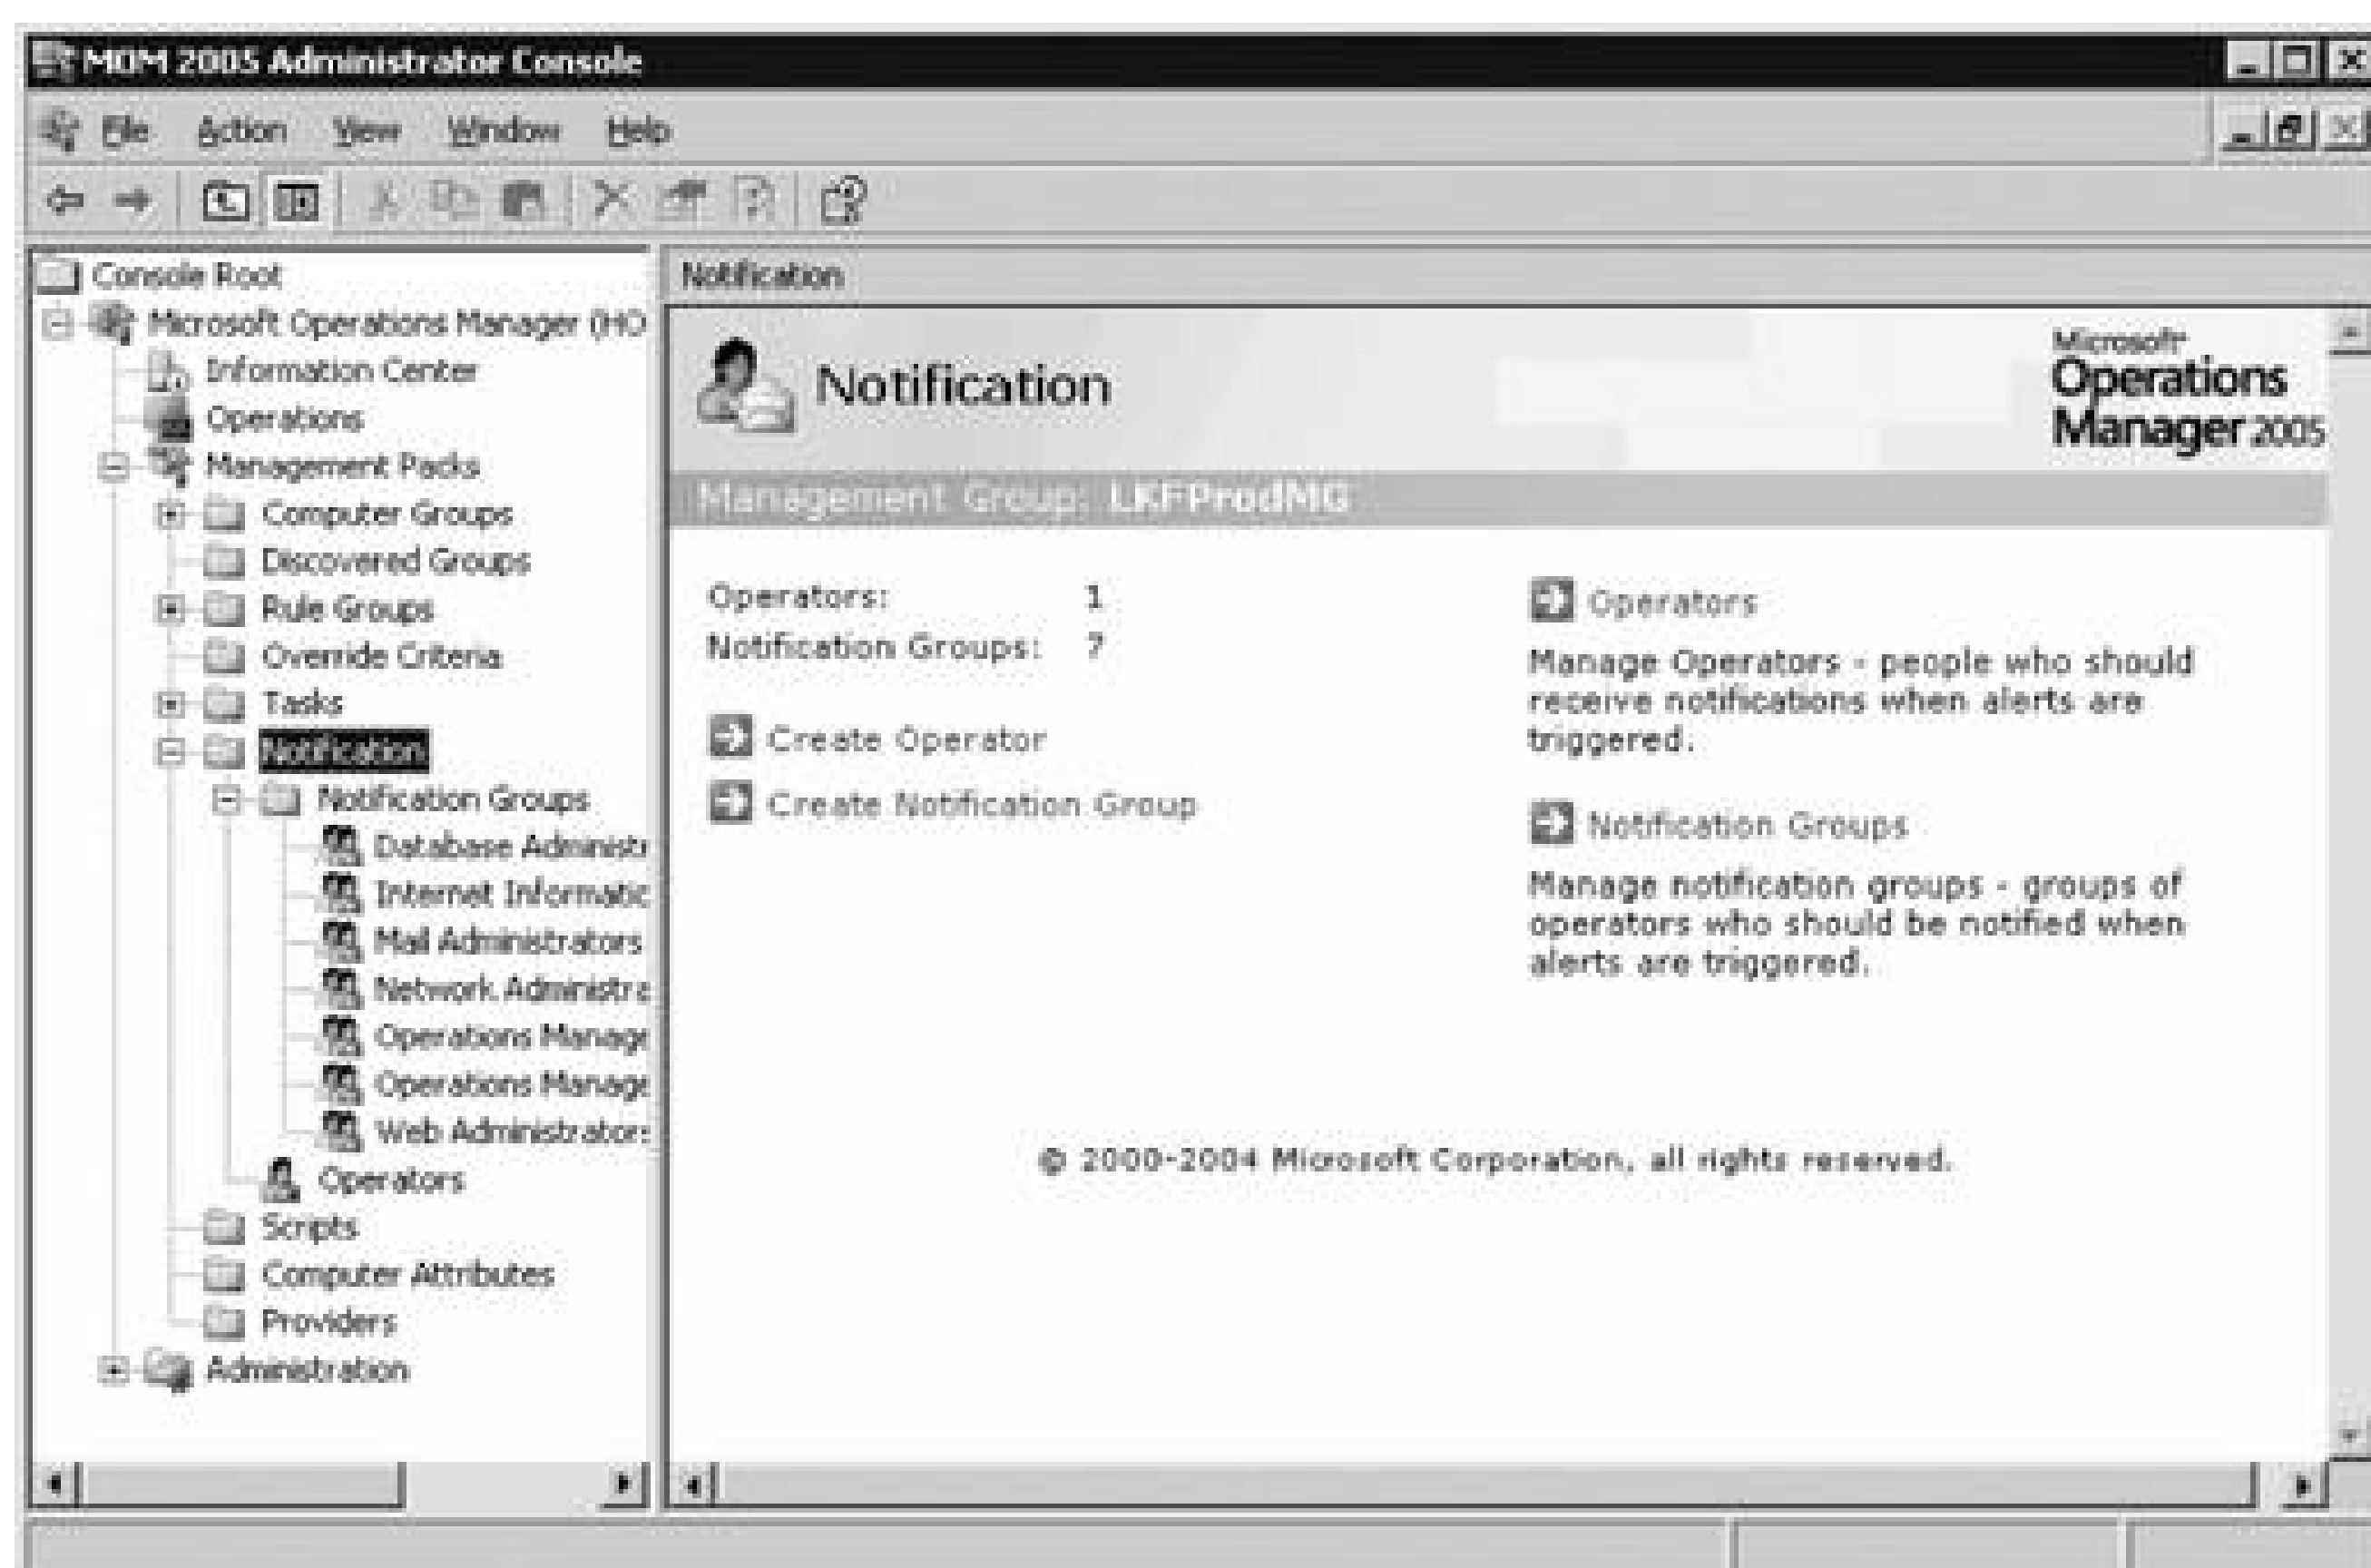


### 4.2.3.5. Alert notification

Alerts serve no purpose if they aren't visible to responders. You could ensure visibility by having someone watch the Operator console all day long, but that's not practical.

While the management pack is still in the preproduction phase, configure notifications in the Administrator console (see [Figure 4-12](#)). Notifications need to go to the people that can investigate or fix the issue. If an alert indicates an outage of a major system, such as an Exchange mailbox store going offline, you may want an executive sponsor of that system notified to communicate the outage to the rest of management. Ultimately, who you place in the notification groups is up to you. In the Notification node, you create and manage operator objects and notification groups. In the management pack, you will need to examine the alert rules response tab and adjust it to send a notification to a notification group. If the notifications are to be sent via SMTP, you must enter a valid SMTP server and return address in the Administrator console Administration node Global Settings Email Server tab properties. For the purposes of sending SMTP mail, any SMTP server will do; it doesn't have to be Exchange. The SMTP service in IIS works well and its licensing is included in the OS. Another benefit is that if the monitored Exchange SMTP server is down, the notification will still go out.

Figure 4-12. The notification node in the Administrator console



Leaky Faucet is concerned about unexpected service terminations on all Exchange servers. If an Exchange-related service stops, they want MOM to send a page to the Windows platform administrators who are Exchange administrators. If a non-Exchange related service stops, they want an email sent only to the help desk with an alert severity of Error, because the default severity (Service Unavailable) is too extreme. Leaky Faucet must modify the rules in the Exchange 2003 Unexpected Service Termination Rules rule group and configure notifications to accomplish this.

To modify the rules, the first step is to create an operator object. This is not a domain account, although it can appear that way. Instead of naming the objects for people, you can just as easily name a pager (e.g., Email On-call pager) or a distribution list with an SMTP address or a monitored shared public folder.

An operator object has only a few properties. It has a name, an SMTP address, a pager address (also SMTP), an entry for calling an external command, notification group membership, and it can be enabled or disabled (see [Figure 4-13](#)).

The "Email On-call pager" entry has been defined as an operator. It will receive notifications via the information entered into the Page Configuration tab of the Operator Properties. Help desk operators will receive an email when an alert is generated and will work to solve the problem.

The Email, Page, and Command tabs that are configurable via the operators object properties are almost identical.

[Figure 4-14](#) shows the Page tab for the "Email On-call pager" operator object in the Leaky Faucet environment.

Figure 4-13. The primary properties of an operator object

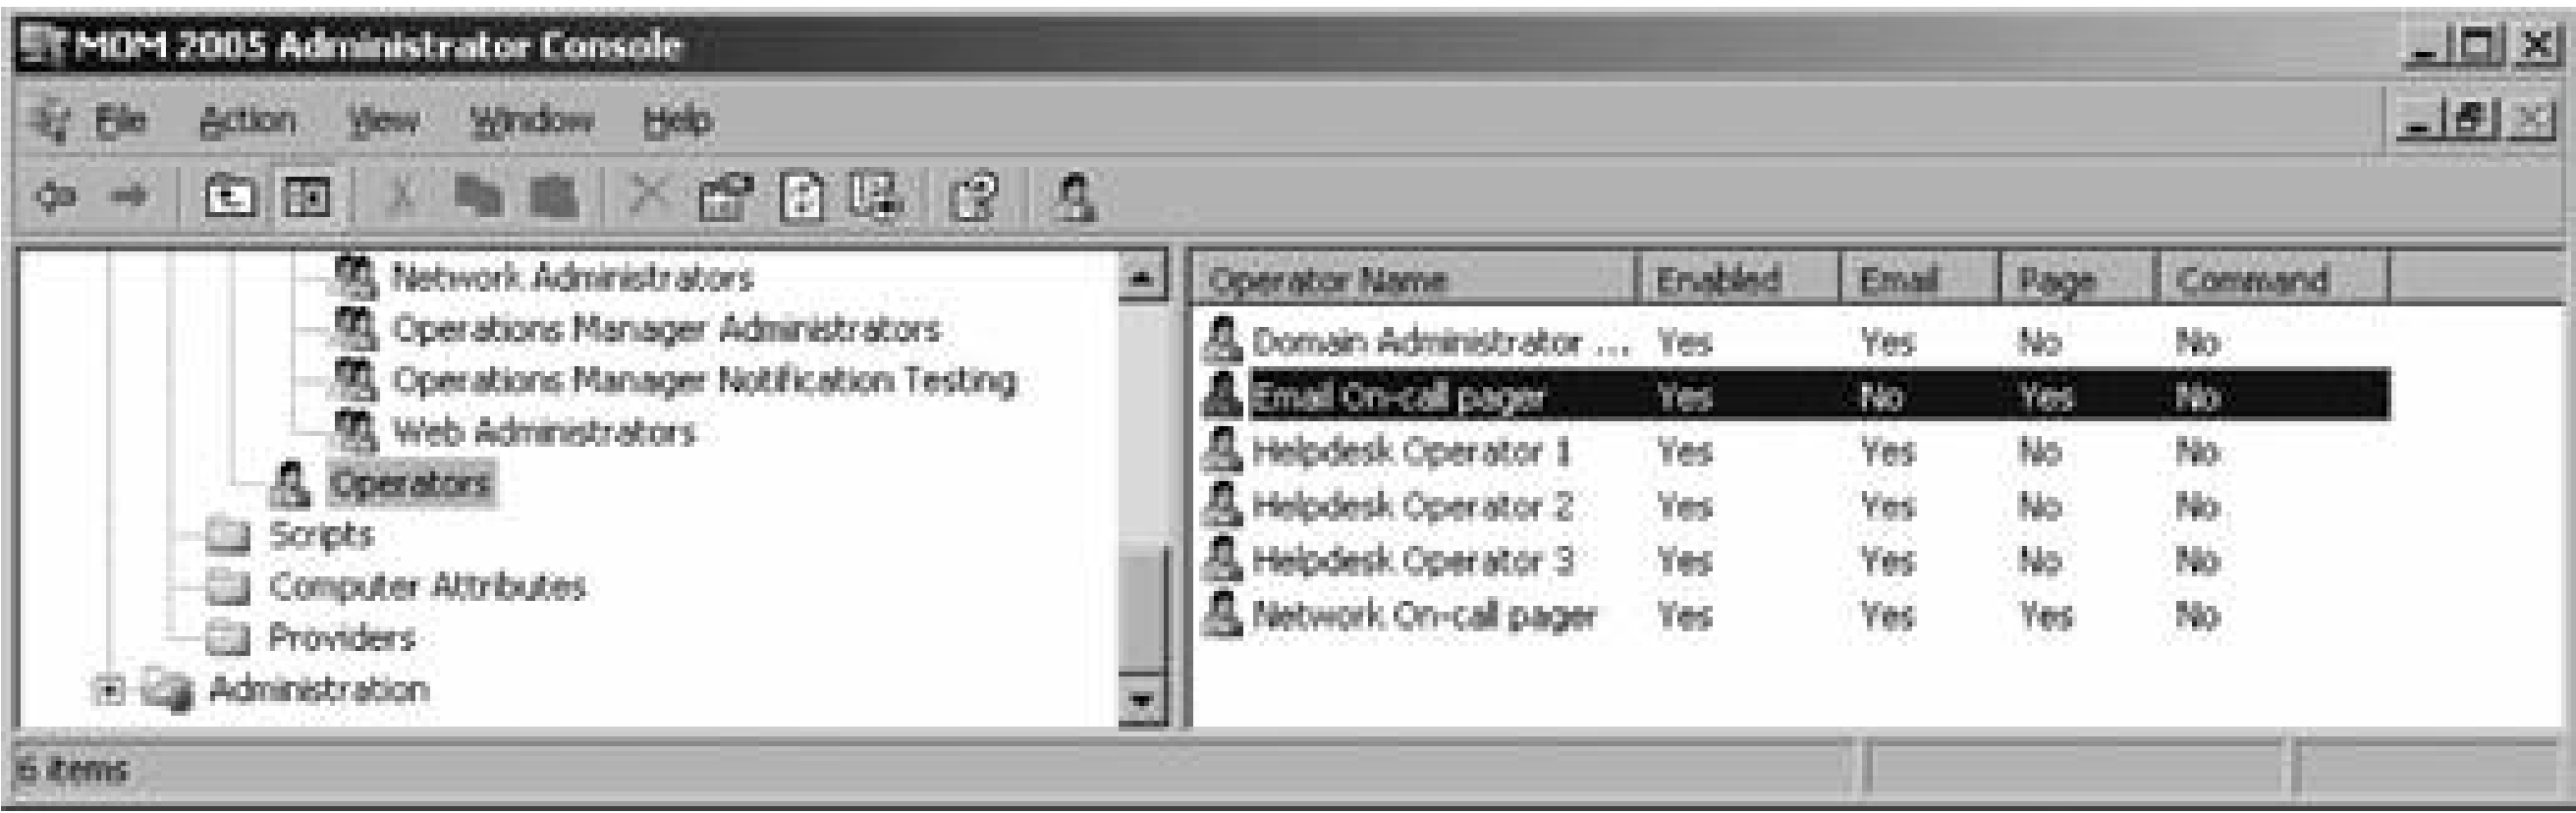
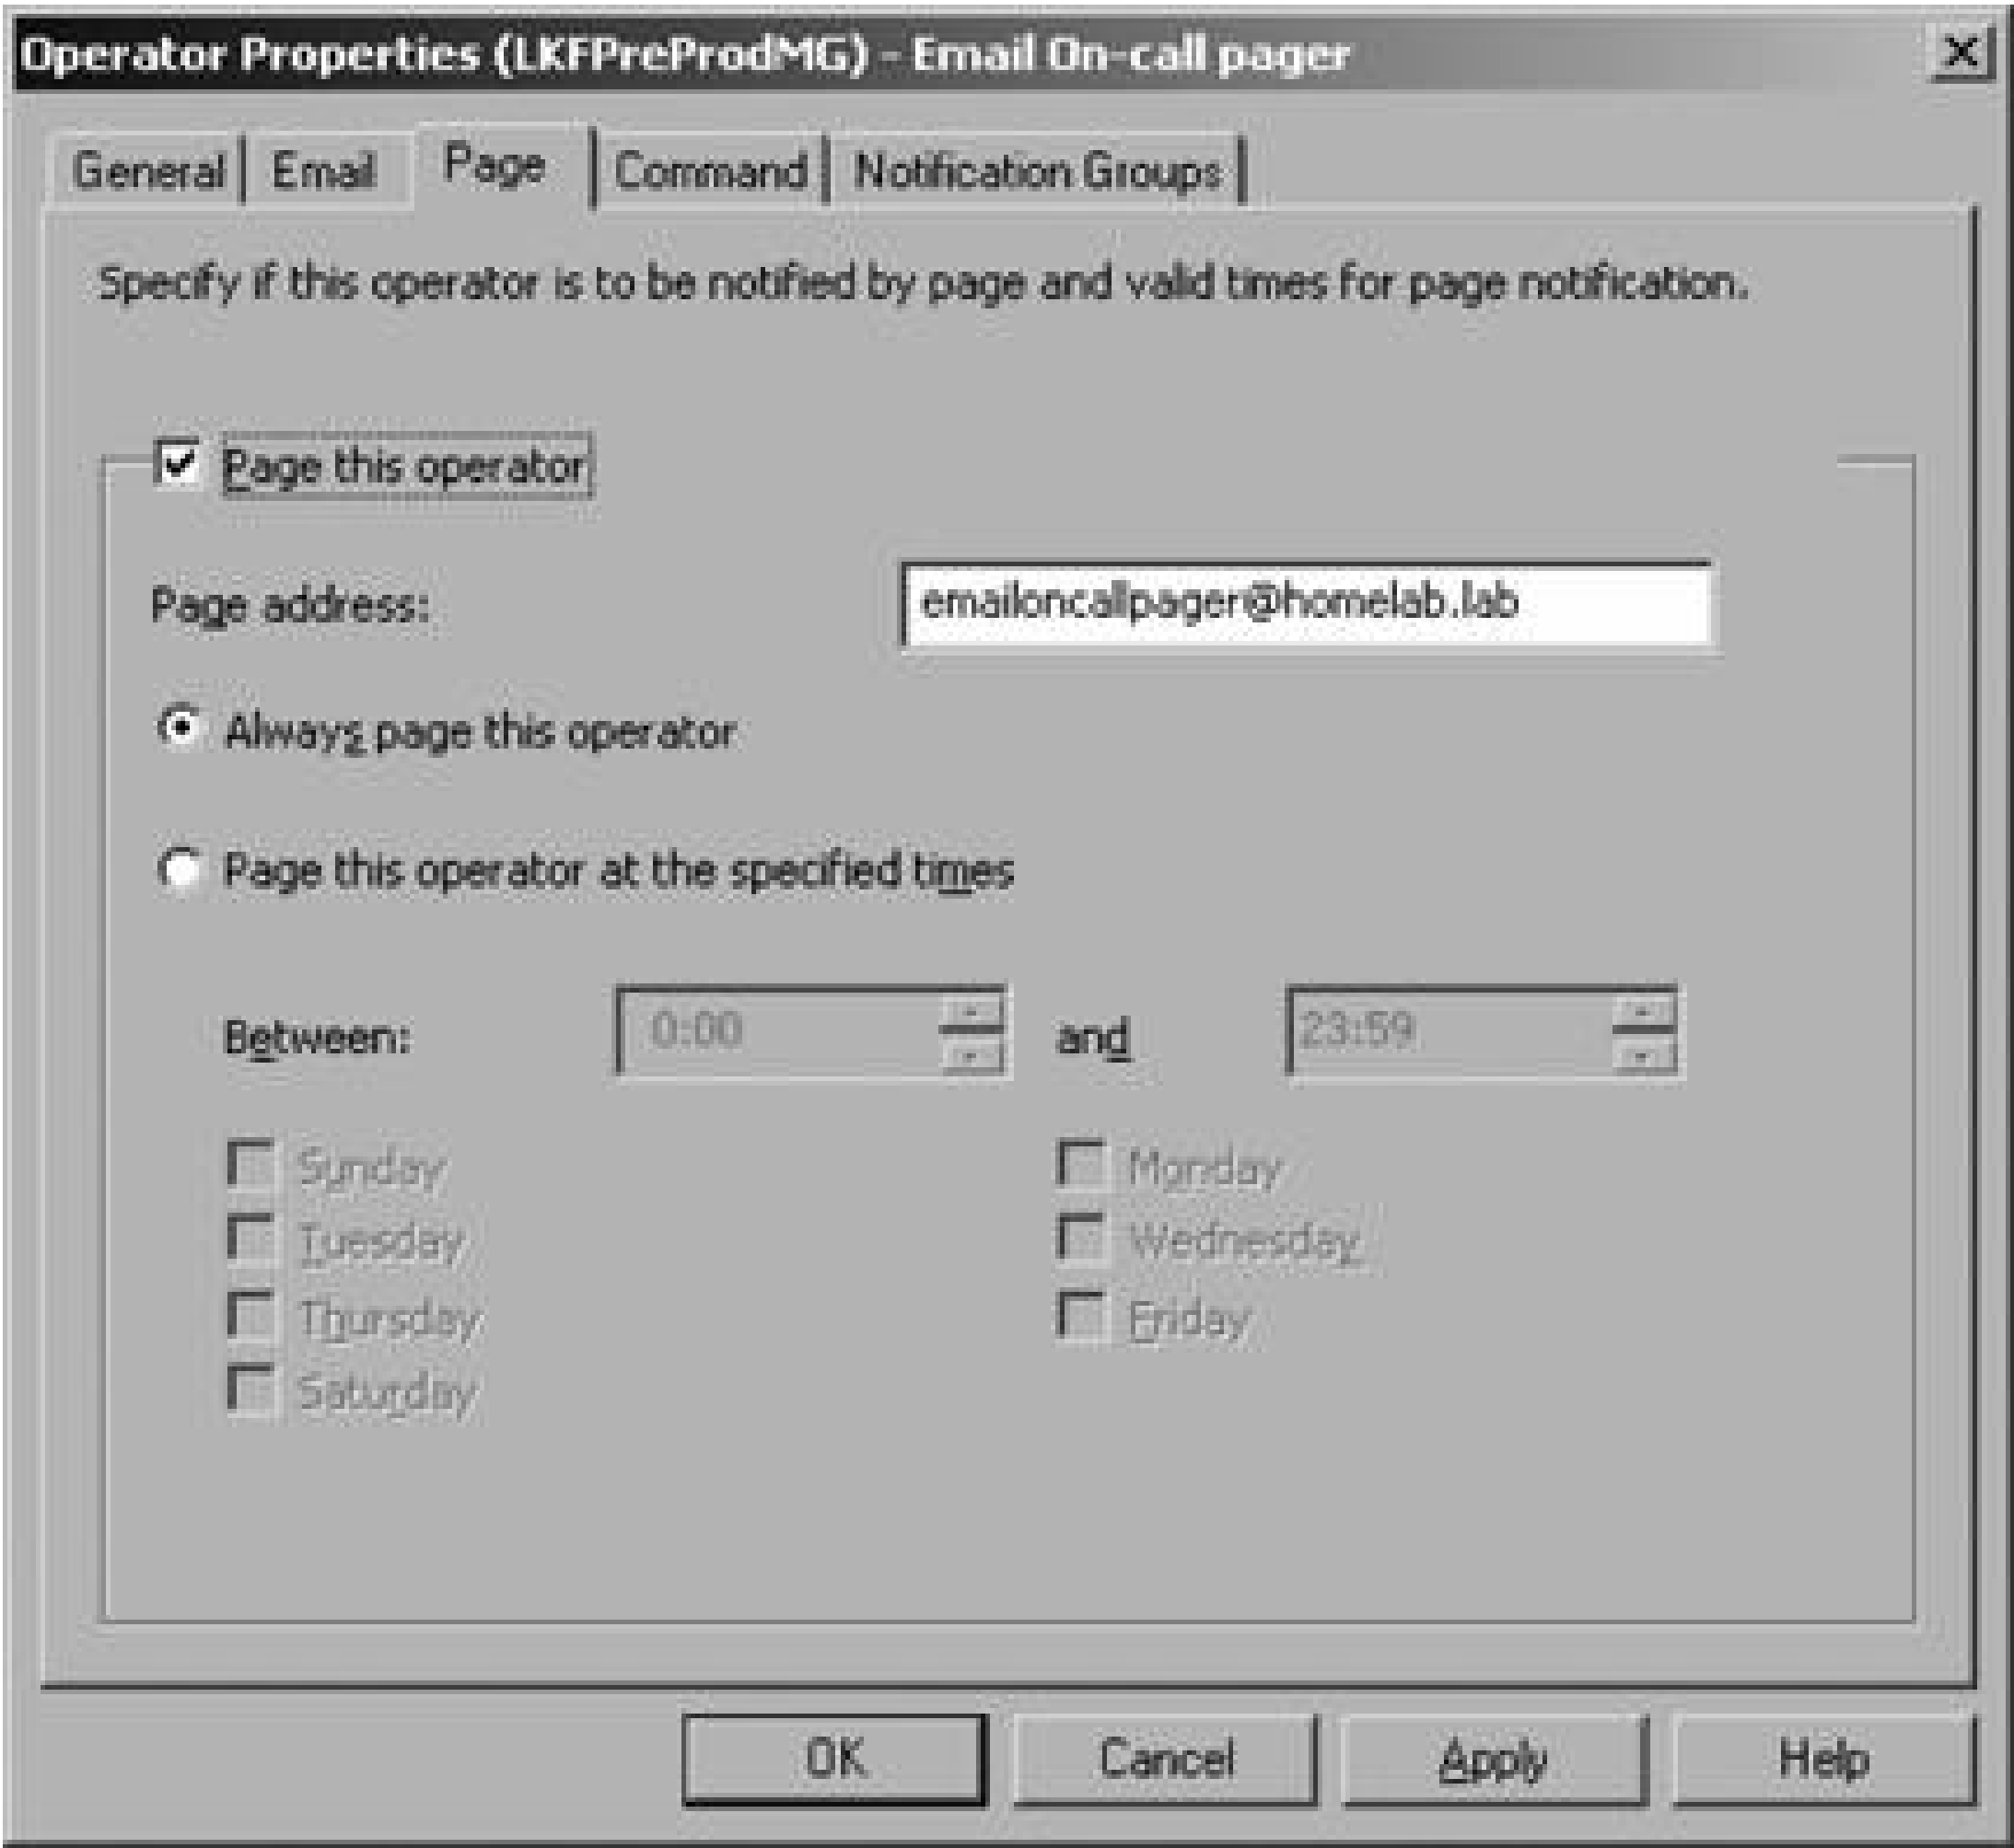


Figure 4-14. Pager configuration information for an operator object



On this tab and the email tabs, Max can enable this communication channel, provide the SMTP address to send the notification to, and specify if this channel is used at all times or according to a schedule. On the Command tab, Max provides the operator ID. Lastly, on the Notification Groups tab Max adds the operator to the desired notification group in this case the Mail Administrators notification group. The help desk operator objects have been added to the Level 1 Operators notification group.

To complete the configuration of this notification process, Max must modify the alert rule in the Unexpected Service Termination Rules rule group. To find this rule group, in the Administrator console he navigates to Rule Groups      Microsoft Exchange Server      Exchange 2003



Availability and State Monitoring ➔ Unexpected Service Termination      Unexpected Service Termination Rules. This rule group has two event rules and one alert rule, which are straightforward examples of event rules.

*Unexpected Exchange-related service termination event rule*

This rule monitors the Windows system event log (Data Provider tab) on an Exchange server for event IDs 7031 or 7034 from the service control manager that involve any of the Exchange critical services, such as IIS, SMTP, WWW, and Information store. The event-matching criteria can be found on the Criteria tab by clicking the Advanced button. When it finds one of these event IDs, the rule generates an alert of severity Service Unavailable as indicated on the Alert tab.

*Unexpected service termination (not Exchange-related) event rule*

This rule looks identical to the Exchange-related rule, except that in the criteria formula the value for parameter 1 (point 1 in [Figure 4-15](#)) is "doesn't match regular expression" whereas in the other the value is "matches regular expression."

Figure 4-15. Matching criteria for an event rule

When a rule can be described in a single sentence it follows this pattern: "The rule monitors *provider* for *event/value* and when it finds a match it *generates an alert/executes a response*." The parameters on all the other tabs simply serve to modify the data produced by three core components of an alert definition. The Exchange-related service termination rule needs no modification to function properly.

For the non-Exchange-related service termination rule, Max needs to change the alert severity to

Error. Following the rule modification procedure, there are three event rules, the original Exchange-related rule, the disabled original non-Exchange-related rule, and the modified copy of the non-Exchange rule, which is now tagged with the LKF prefix.

The job of the single alert rule in this rule group is to send a notification to the Mail Administrators group if an alert of severity Error or higher occurs. For an alert rule, the provider and criteria functionality is combined on the Alert Criteria tab. The notification is configured on the Responses tab. Since the event rules in this rule group will generate alerts of severity Error or Service Unavailable and an alert rule can only address one alert type, Leaky Faucet will need two enabled alert rules. Max will need to do the following:

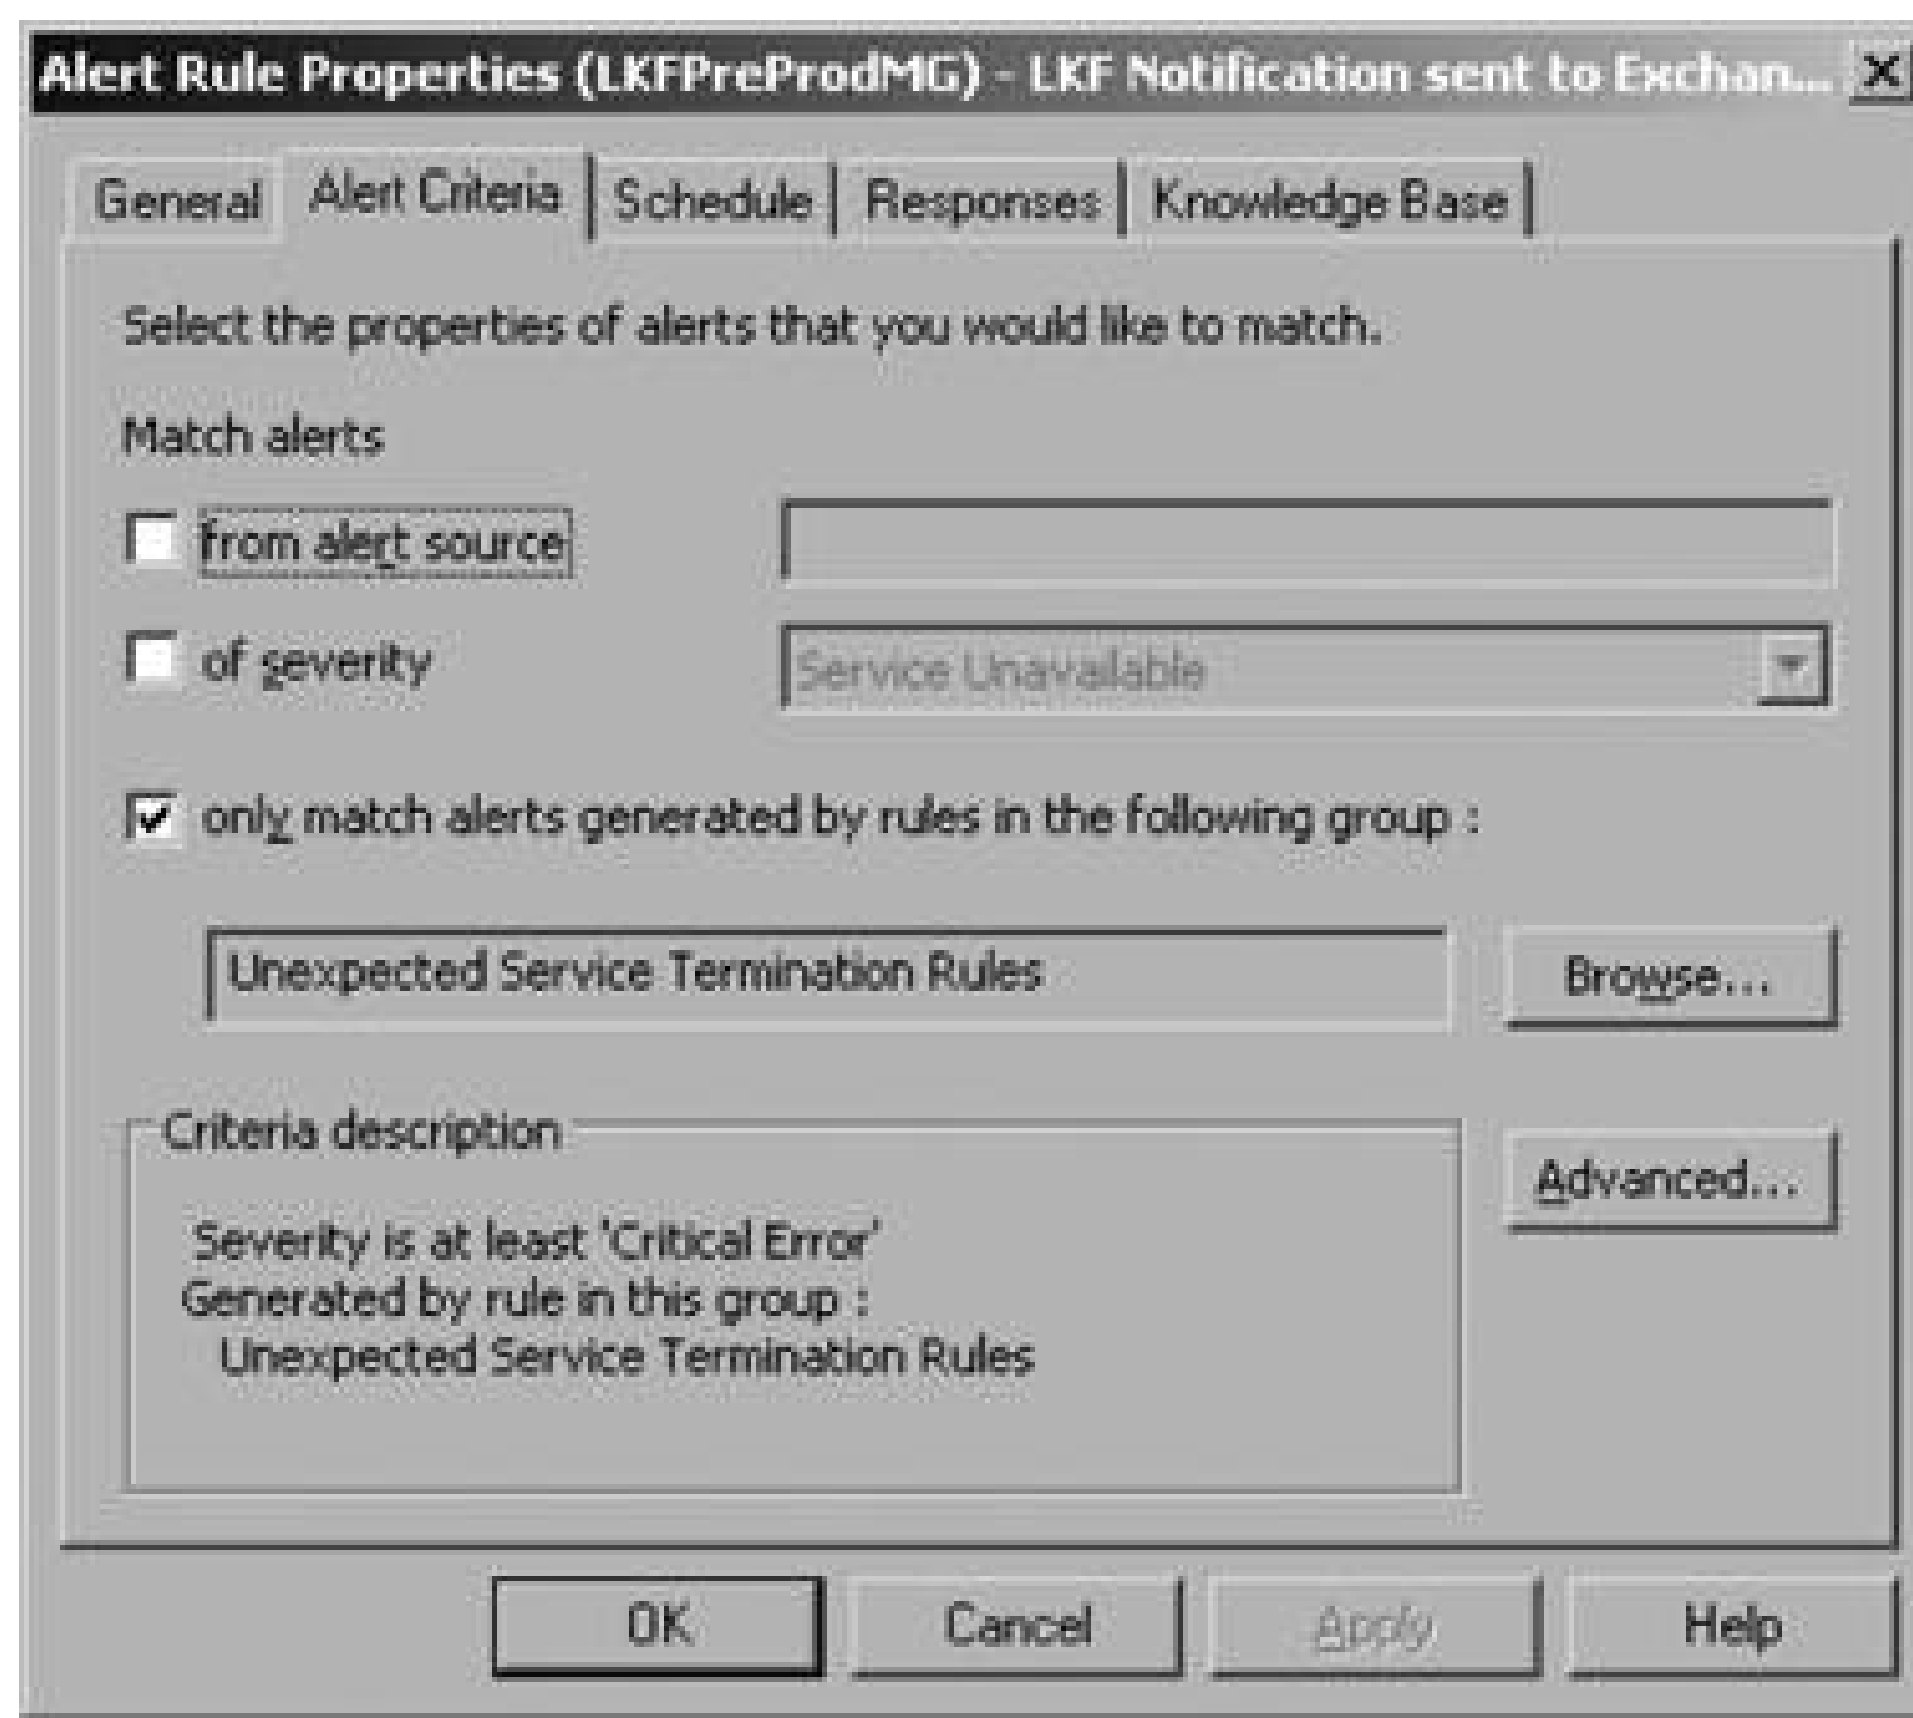
1. Copy the single alert rule and change the name to "LKF Notification sent to Exchange administrators if Severity Level is Critical Error or higher."
2. On the Alert Criteria tab, go to the Advanced Properties, select the criteria "Severity is at least Error" and remove it (see [Figure 4-16](#)). This will bring the Field, Condition, and Value parameters into the lower drop-down boxes for editing.

Figure 4-16. The alert processing rule advanced criteria

3. Change the Value field to Critical Error and add it back to the list. The criteria description on the Alert Criteria tab will reflect this change (as shown in [Figure 4-17](#)). No other change to this rule is required since the action on the Responses tab is to notify the mail administrator's notification group.
4. Copy the original alert rule again and rename it to "LKF Notification sent to Helpdesk Operators if Severity Level is at most Error."
5. Change the alert criteria from "Severity is at least Error" to "Severity is at most Error."
6. Modify the notification group on the Responses tab from the Mail Administrators group to the Level 1 Operators group and disable the original rule.

7. Commit the changes.

Figure 4-17. Updated alert criteria for notifying the Mail Administrators notification group



Once the configuration is complete, the sequence of events from a service stopping to the Windows administrator receiving an alert is:

1. The SMTP service on an Exchange server unexpectedly stops.
2. A 7031 or 7034 event is placed into the system event log on that server.
3. The MOM agent on the Exchange server detects the event, compares it to the Unexpected Exchange-related service termination event rule criteria, finds a match, and generates an alert of severity Critical Error.
4. The MOM agent detects the generation of the Critical Error alert from the Unexpected Service Termination Rules rule group, and compares it to the alert criteria in the alert rule "LKF Notification sent to Exchange administrators if Severity Level is Critical Error or higher." Finding a match, the agent executes the action configured on the Responses tab, which is to send a notification of the alert to the Mail Administrators notification group.
5. The alert is also forwarded to the management server, which passes it to the operations database and makes it available in the Operator console.



6. The agent on the management server reads the membership in the mail administrator's notification group, then reads the configuration for sending a notification to the "Email On-call pager" operator object. This is configured for sending the notification via page to the SMTP address [emailoncallpager@bigpagercompany.net](mailto:emailoncallpager@bigpagercompany.net).
7. The management server then formats the SMTP message according to the configuration defined on the "Notification sent to Exchange administrators if Severity Level is Error or higher" alert rule for the rule group. The page can either use a default format or a custom-defined one. This is on the Alert rule properties → Responses tab      Edit the responses      Page tab.
8. The SMTP email notification is sent to the [emailoncallpager@bigpagercompany.net](mailto:emailoncallpager@bigpagercompany.net) SMTP address via the SMTP server that the management group is configured to use.

The alert is available in the Operator console for the systems administrator to examine.

## 4.2.4. Additional Management Pack Configuration

Some of the more advanced management packs, such as Exchange, Active Directory, or SQL Server 2000, require additional configuration to be fully functional. If the additional configuration is not too extensive, the necessary steps will be presented in the Results pane when the desired rule group is selected ([Figure 4-18](#)). If the additional configuration is extensive or the management pack is particularly complex, there will be a management pack configuration guide. This is the case with the Exchange 2003 management pack for MOM 2005 . This management pack has a 130-page configuration guide and an additional tool that can be used to facilitate configuration of the management pack and all of its components.

Unfortunately, Microsoft does not bundle the management pack guide with the management pack download, so you have to search Microsoft's web site. The Exchange 2003 management pack guide is available at <http://www.microsoft.com/downloads/details.aspx?FamilyId=2215EEAB-41D7-423D-9F54-01F0DF4647E9>. For a listing of all management packs and their guides, go to the product documentation page, <http://www.microsoft.com/management/mma/catalog.aspx>.

Before deploying any Microsoft management pack into production, read its documentation to avoid any unnecessary errors that might arise from management pack misconfiguration.

## 4.2.5. Resource Kit Tools

Tuning management packs and processing rules is more of an art than a science, and you need to have experience with MOM 2005 and deep knowledge of your environment to be successful. When you are modifying rules and need to test a change, use the *Event Creator* (*EventCreator.exe*) tool in the MOM Resource Kit. This tool creates and places an event in the event log of your choice, as shown in [Figure 4-19](#). The target machine can be configured, as well as the log to create the event in; the source of the event; and the event ID, type, category, and username. You should know all of these parameters since the event processing rule for that event is already configured.

Figure 4-18. Basic management pack information and configuration steps

are included with the management pack itself

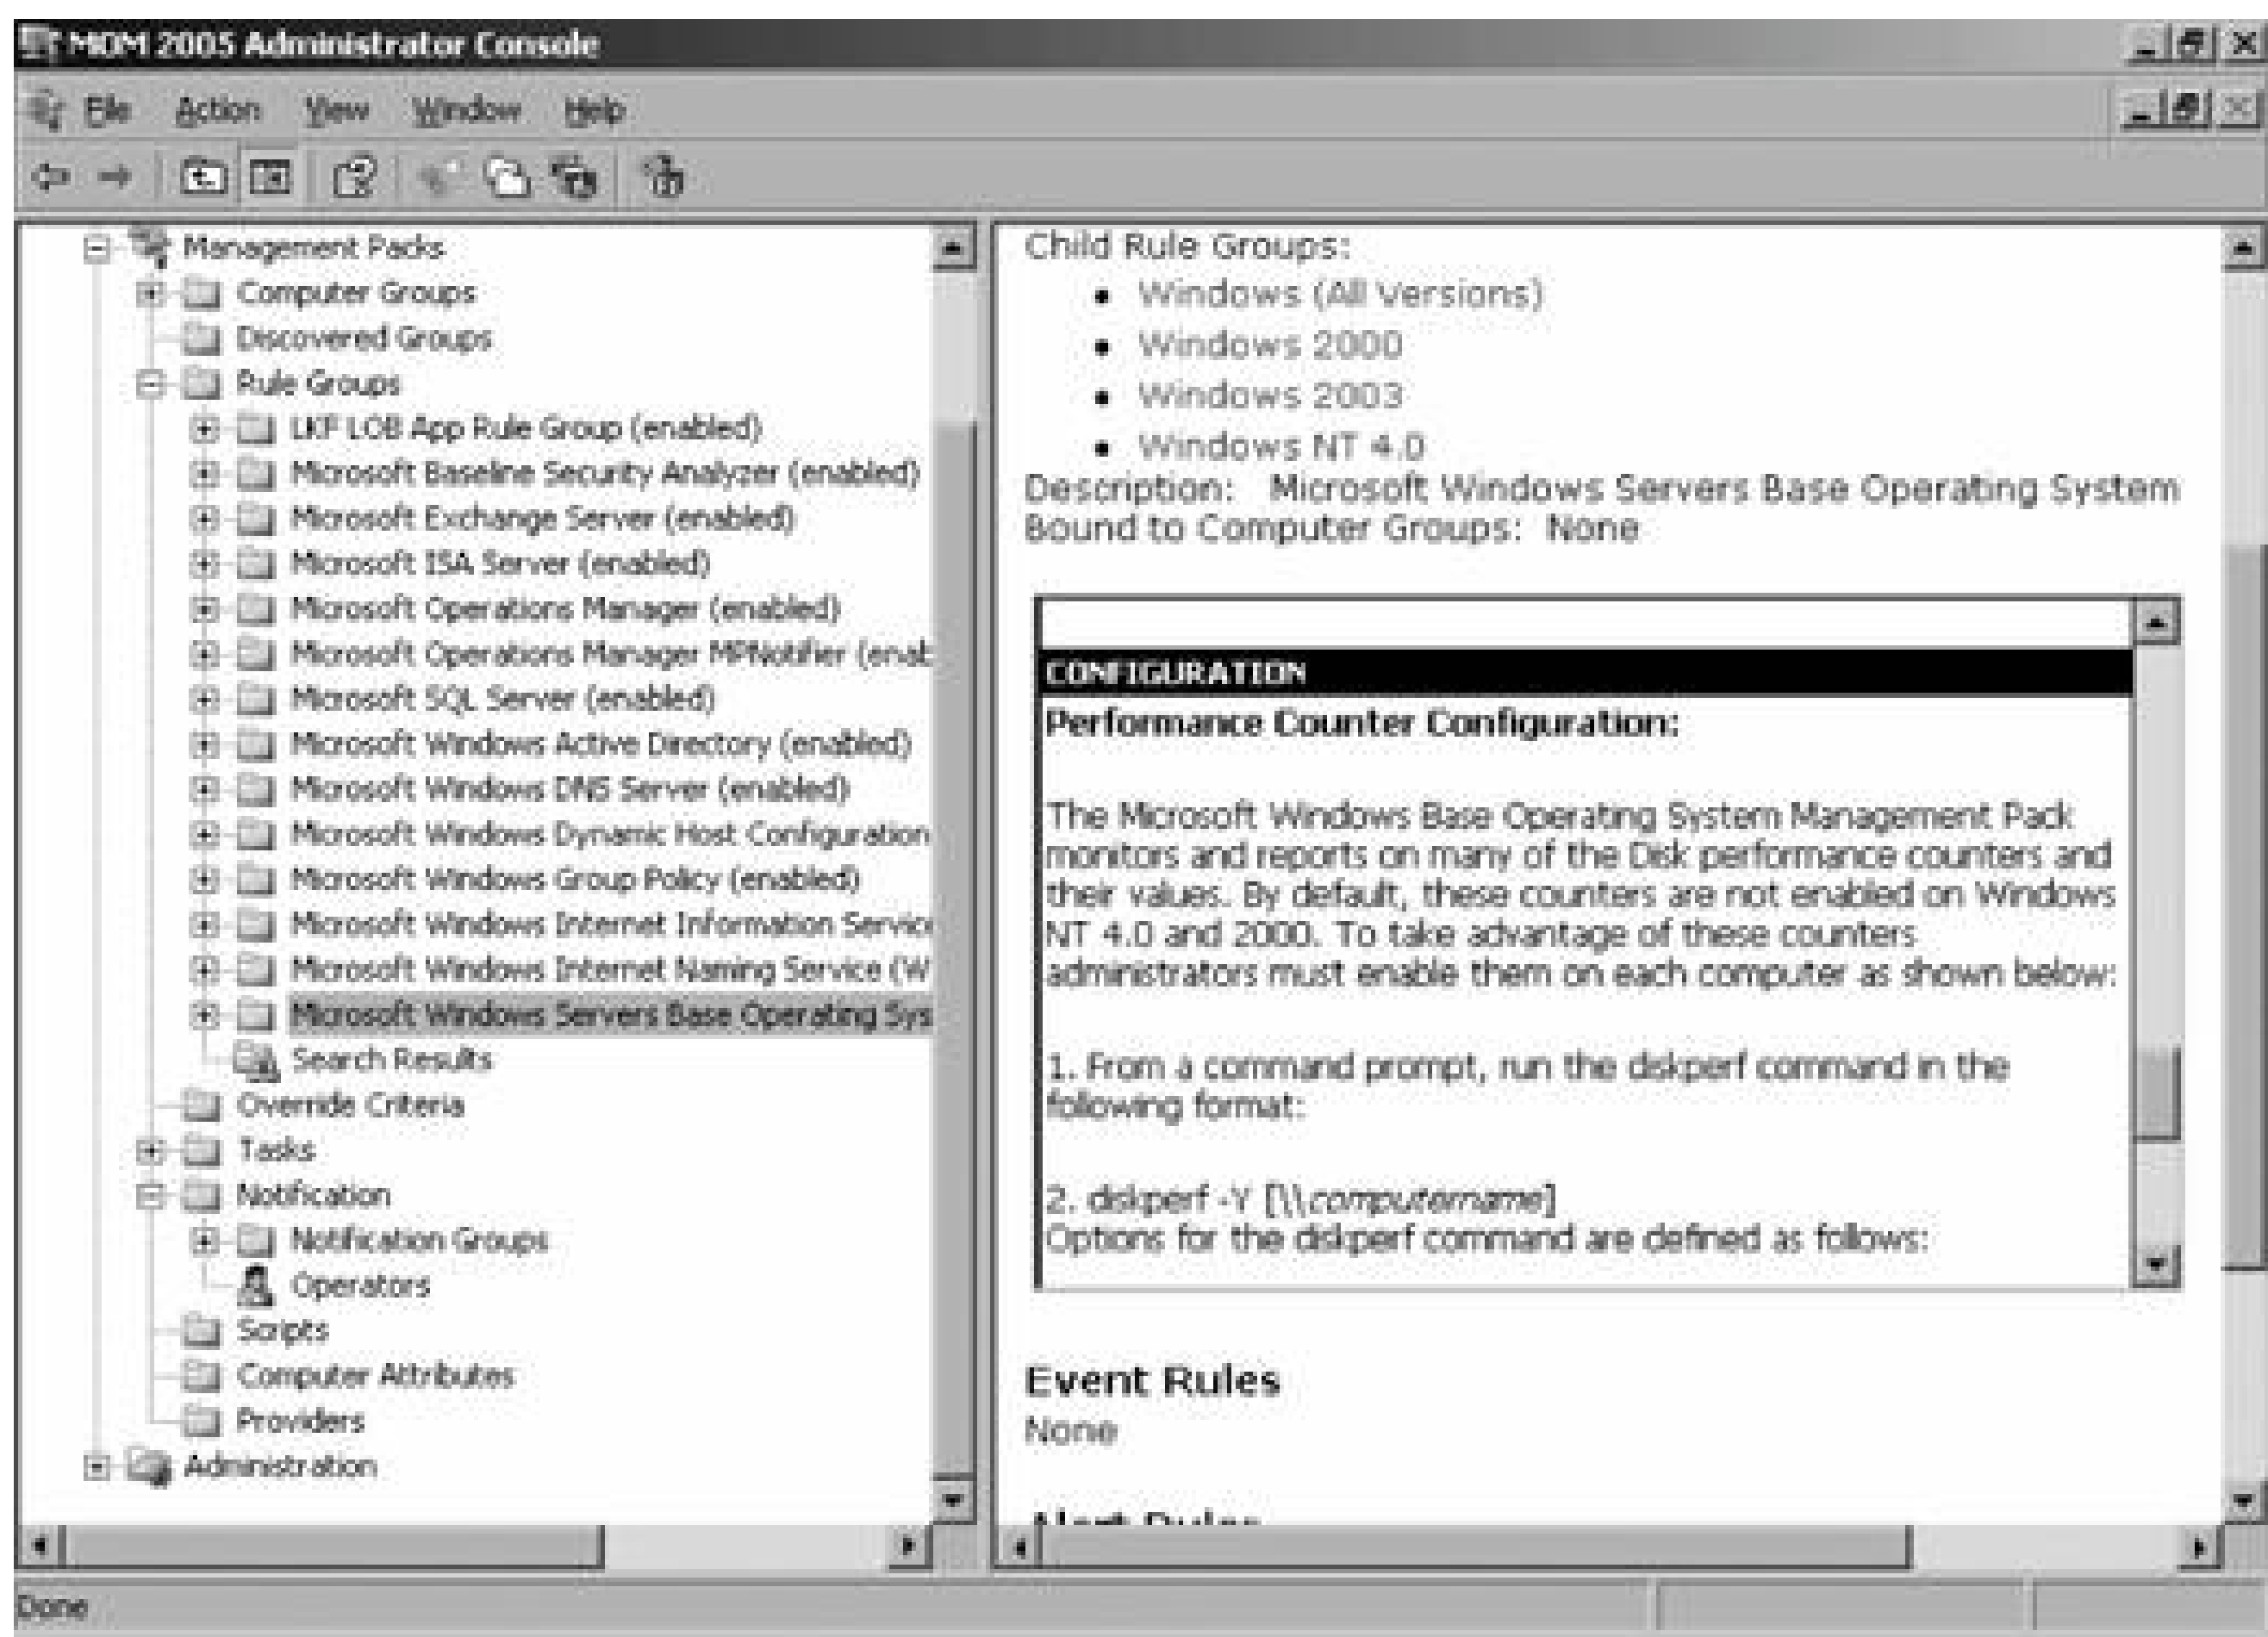


Figure 4-19. The Event Creator resource kit tool

Another useful tool to run in your preproduction environment is the Microsoft Operations Manager MPNotifier management pack . This simple management pack consists of two event rules, an alert rule, and a script.

Once every 24 hours (1,440 minutes) the Microsoft MPNotifier Version Check script is called, which collects an XML file from a Microsoft web site (<http://www.microsoft.com/management/mma/momnotifier.xml>). This file contains the current version of all the Microsoft-authored management packs. The versions of the existing management packs are compared to the current ones in the XML file and an error is generated if your existing management pack versions are behind the currently released version from Microsoft.

Many other rule parameters can be tweaked, but there is no reliable way to tell which rules and rule parameters will need further tuning or exactly how to tune them until the management pack is deployed into production .





## 4.3. Transfer the Management Pack to Production

The management pack is now in the MP.1.1 state in the management pack workflow diagram (Figure 4-1). Some rule groups have been disabled and some rules have been copied and modified. Operations data has been returned from the pilot agents that are multi-homed between LKFPreProdMG and the LKFProdMG management pack groups. Max can now export the management pack (point 3 in Figure 4-1) to an *.akm* file that can be imported into the production environment.

### 4.3.1. Export from Preproduction

The export process is relatively straightforward and accomplished through the Import/Export Management Pack tool. The main thing you need to know is that not all of the configuration information that is associated with a rule group gets exported. The tool will prompt you for the rule group that you want to export, as well as the views and tasks. The export process also includes computer group definitions, provider definitions, script attributes, and any managed code, in addition to the rule group.

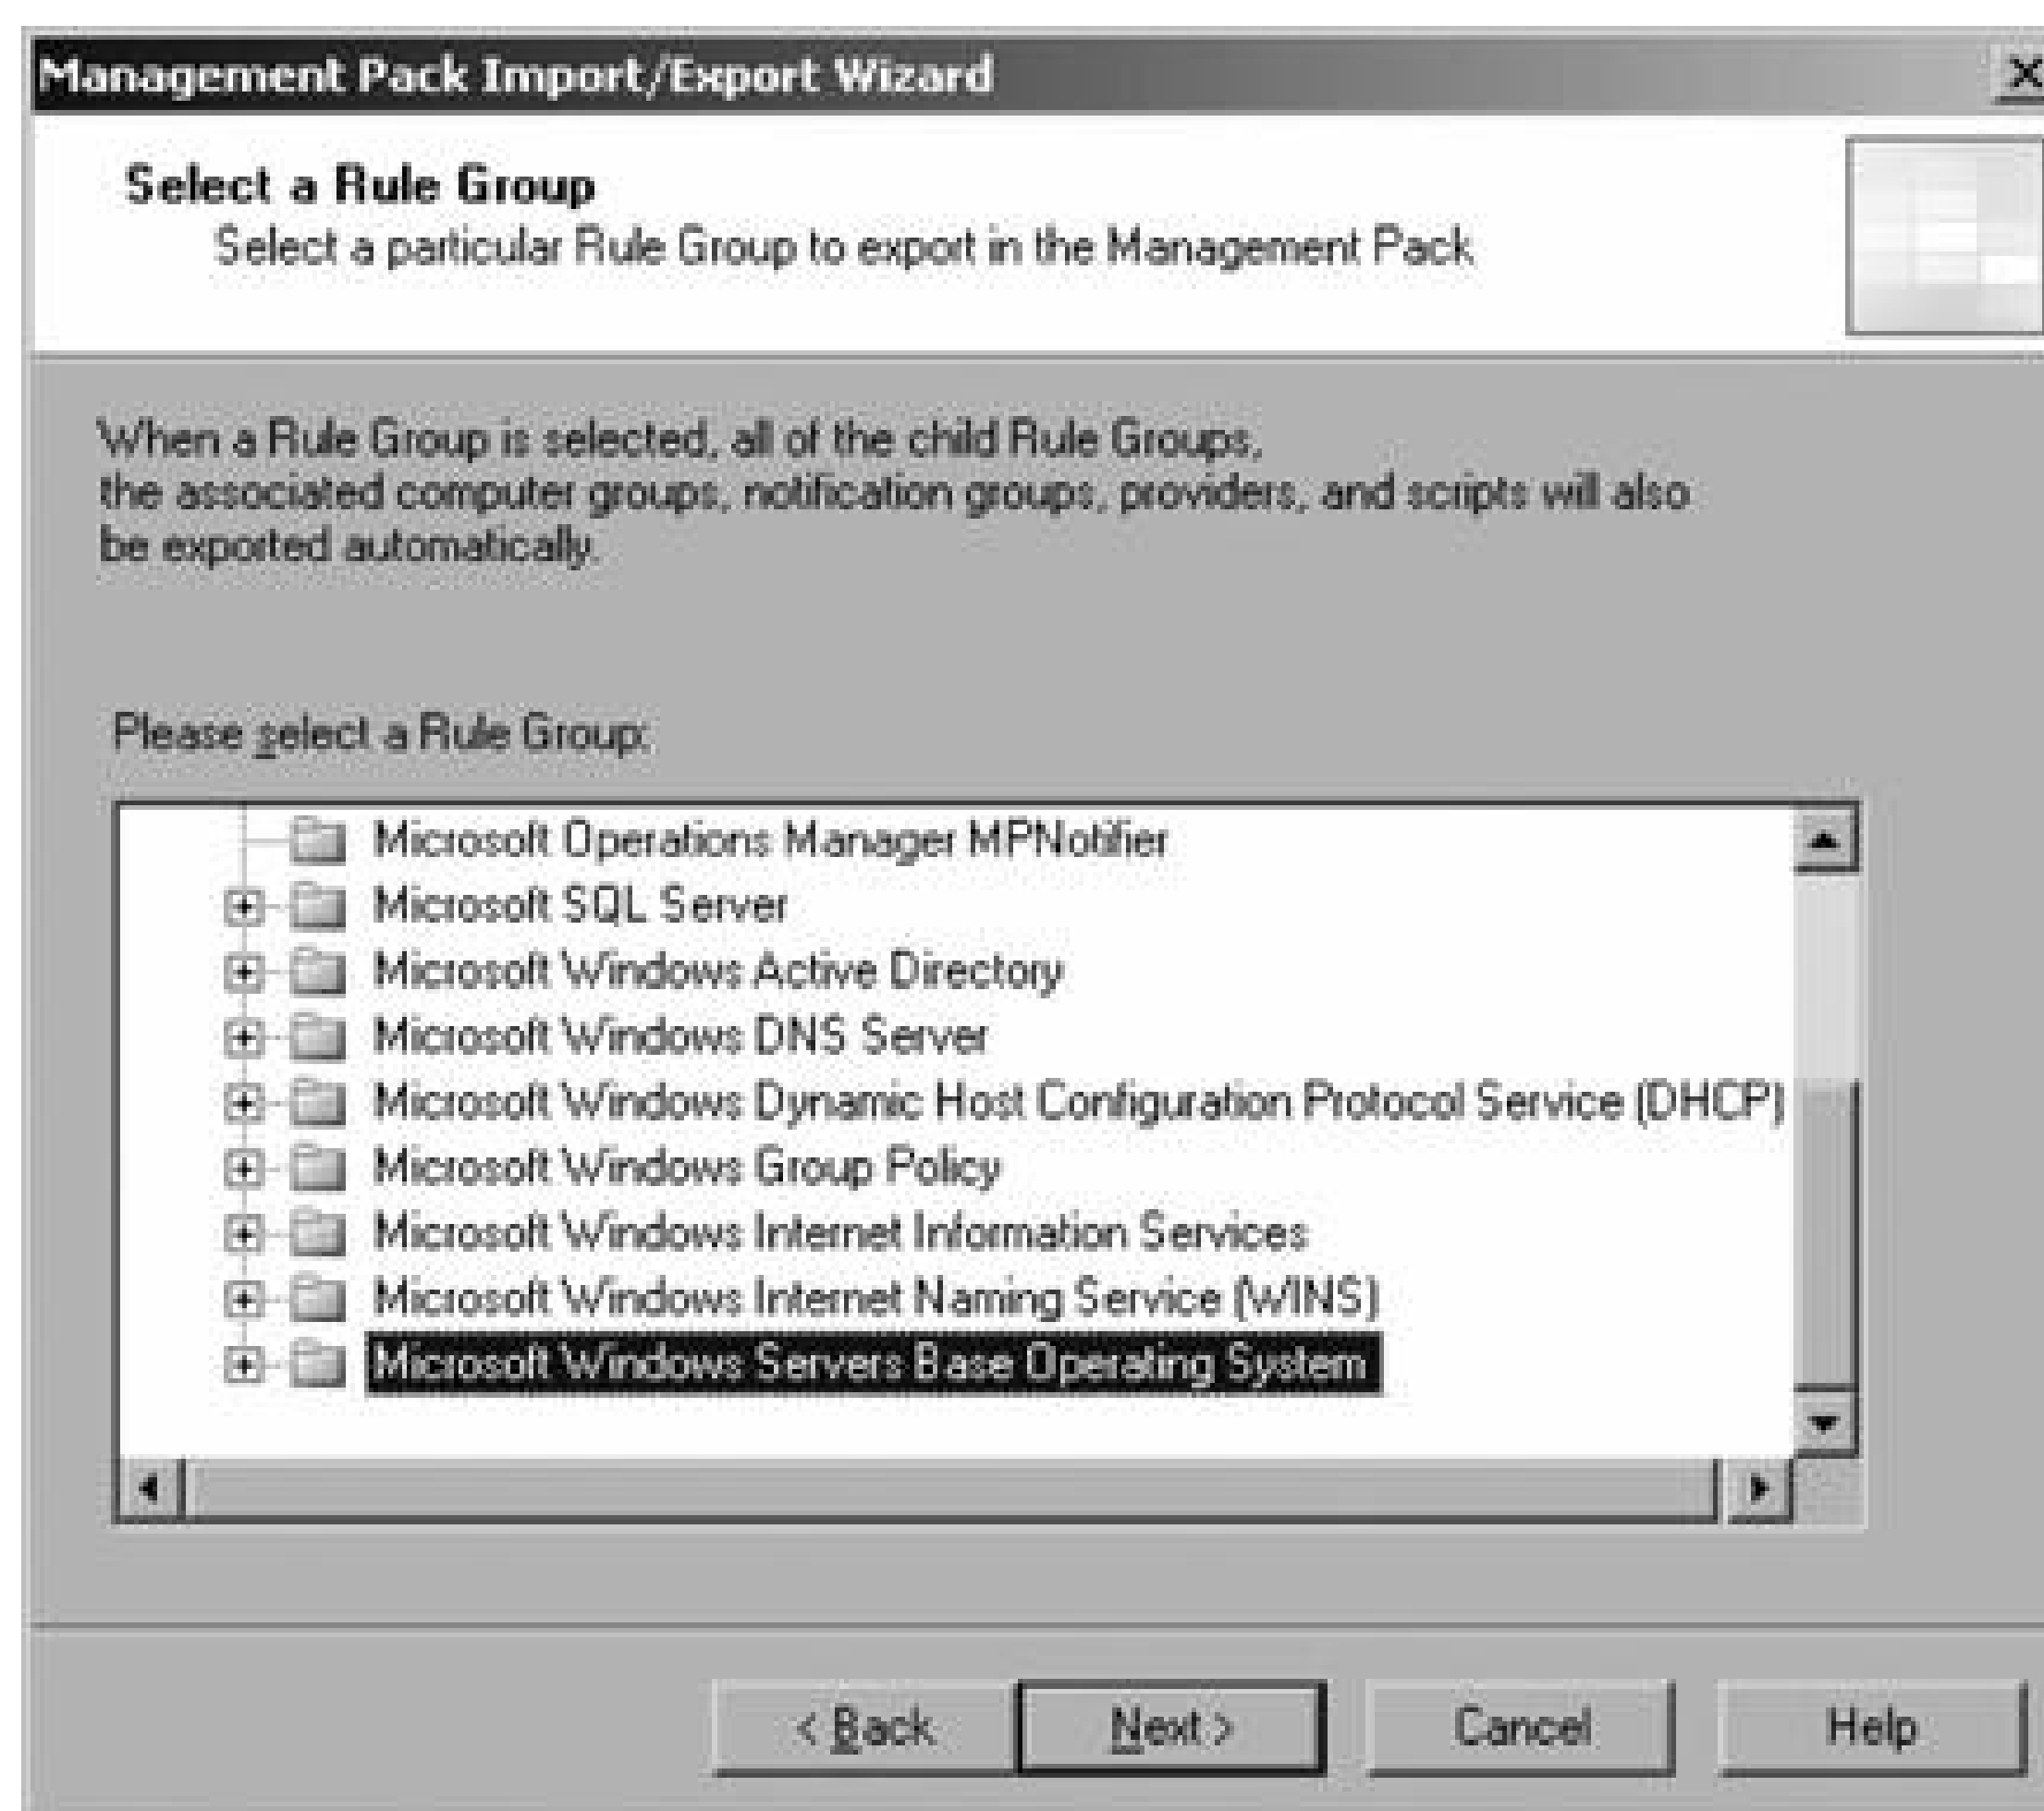
The export process doesn't export operator objects, notification groups, or overrides. If you have configured any of these, you will need to manually record these settings and reconfigure them in your production environment after you have imported the management pack.

To start the export process, launch the Import/Export Management Pack tool and select "to export a management pack." Then select the rule group that you want to export (Figure 4-20).

The best practice is to choose to export the top-level rule group folder, which automatically includes all sub-rule groups. In Figure 4-20, the top-level rule group folder Microsoft Windows Servers Base Operating System has been selected. This ensures that the management pack has everything needed and it can be redistributed.

On the next two pages of the wizard, you are prompted for the views and the tasks that you wish to include in the management pack (Figure 4-21). The folder structure for these is expandable, so you can select down to the individual object level.

Figure 4-20. Choosing to export the Base OS rule group



The best practice is to select the top-level object that belongs in the complete management pack. The two pages are pretty much the same thing except that one shows views and the other tasks.

On the next to last page, specify the directory and filename of the management pack (see Figure 4-22). This part is critical to your versioning process. Remember when you did the management pack import and specified the management pack backup directory (see Figure 4-4)? You pointed the tool to the *OldMP* directory in the file transfer folder and created the *.akm* file in this format: *ManagementPackName\_MM.DD.YY.HH.MM.SS*.

In the export portion of the tool, specify the *\\MPTransferFolder\\CurrentMP* directory. To be consistent with the tool's naming standard, call the newly created management pack *Windows.ServerBaseOS\_02.28.05.19.3*. The seconds portion of the timestamp can be dropped, because specificity to the minute is good enough. Use the 24-hour time format rather than 12-hour time format to ensure uniqueness. Since this is the first time you are exporting a management pack into this directory, there is no concern about earlier versions of the management pack being there. Otherwise, you would have moved the existing *.akm* file into the *OldMP* directory before executing the export process.

Figure 4-21. Select the views that you want to include in the management

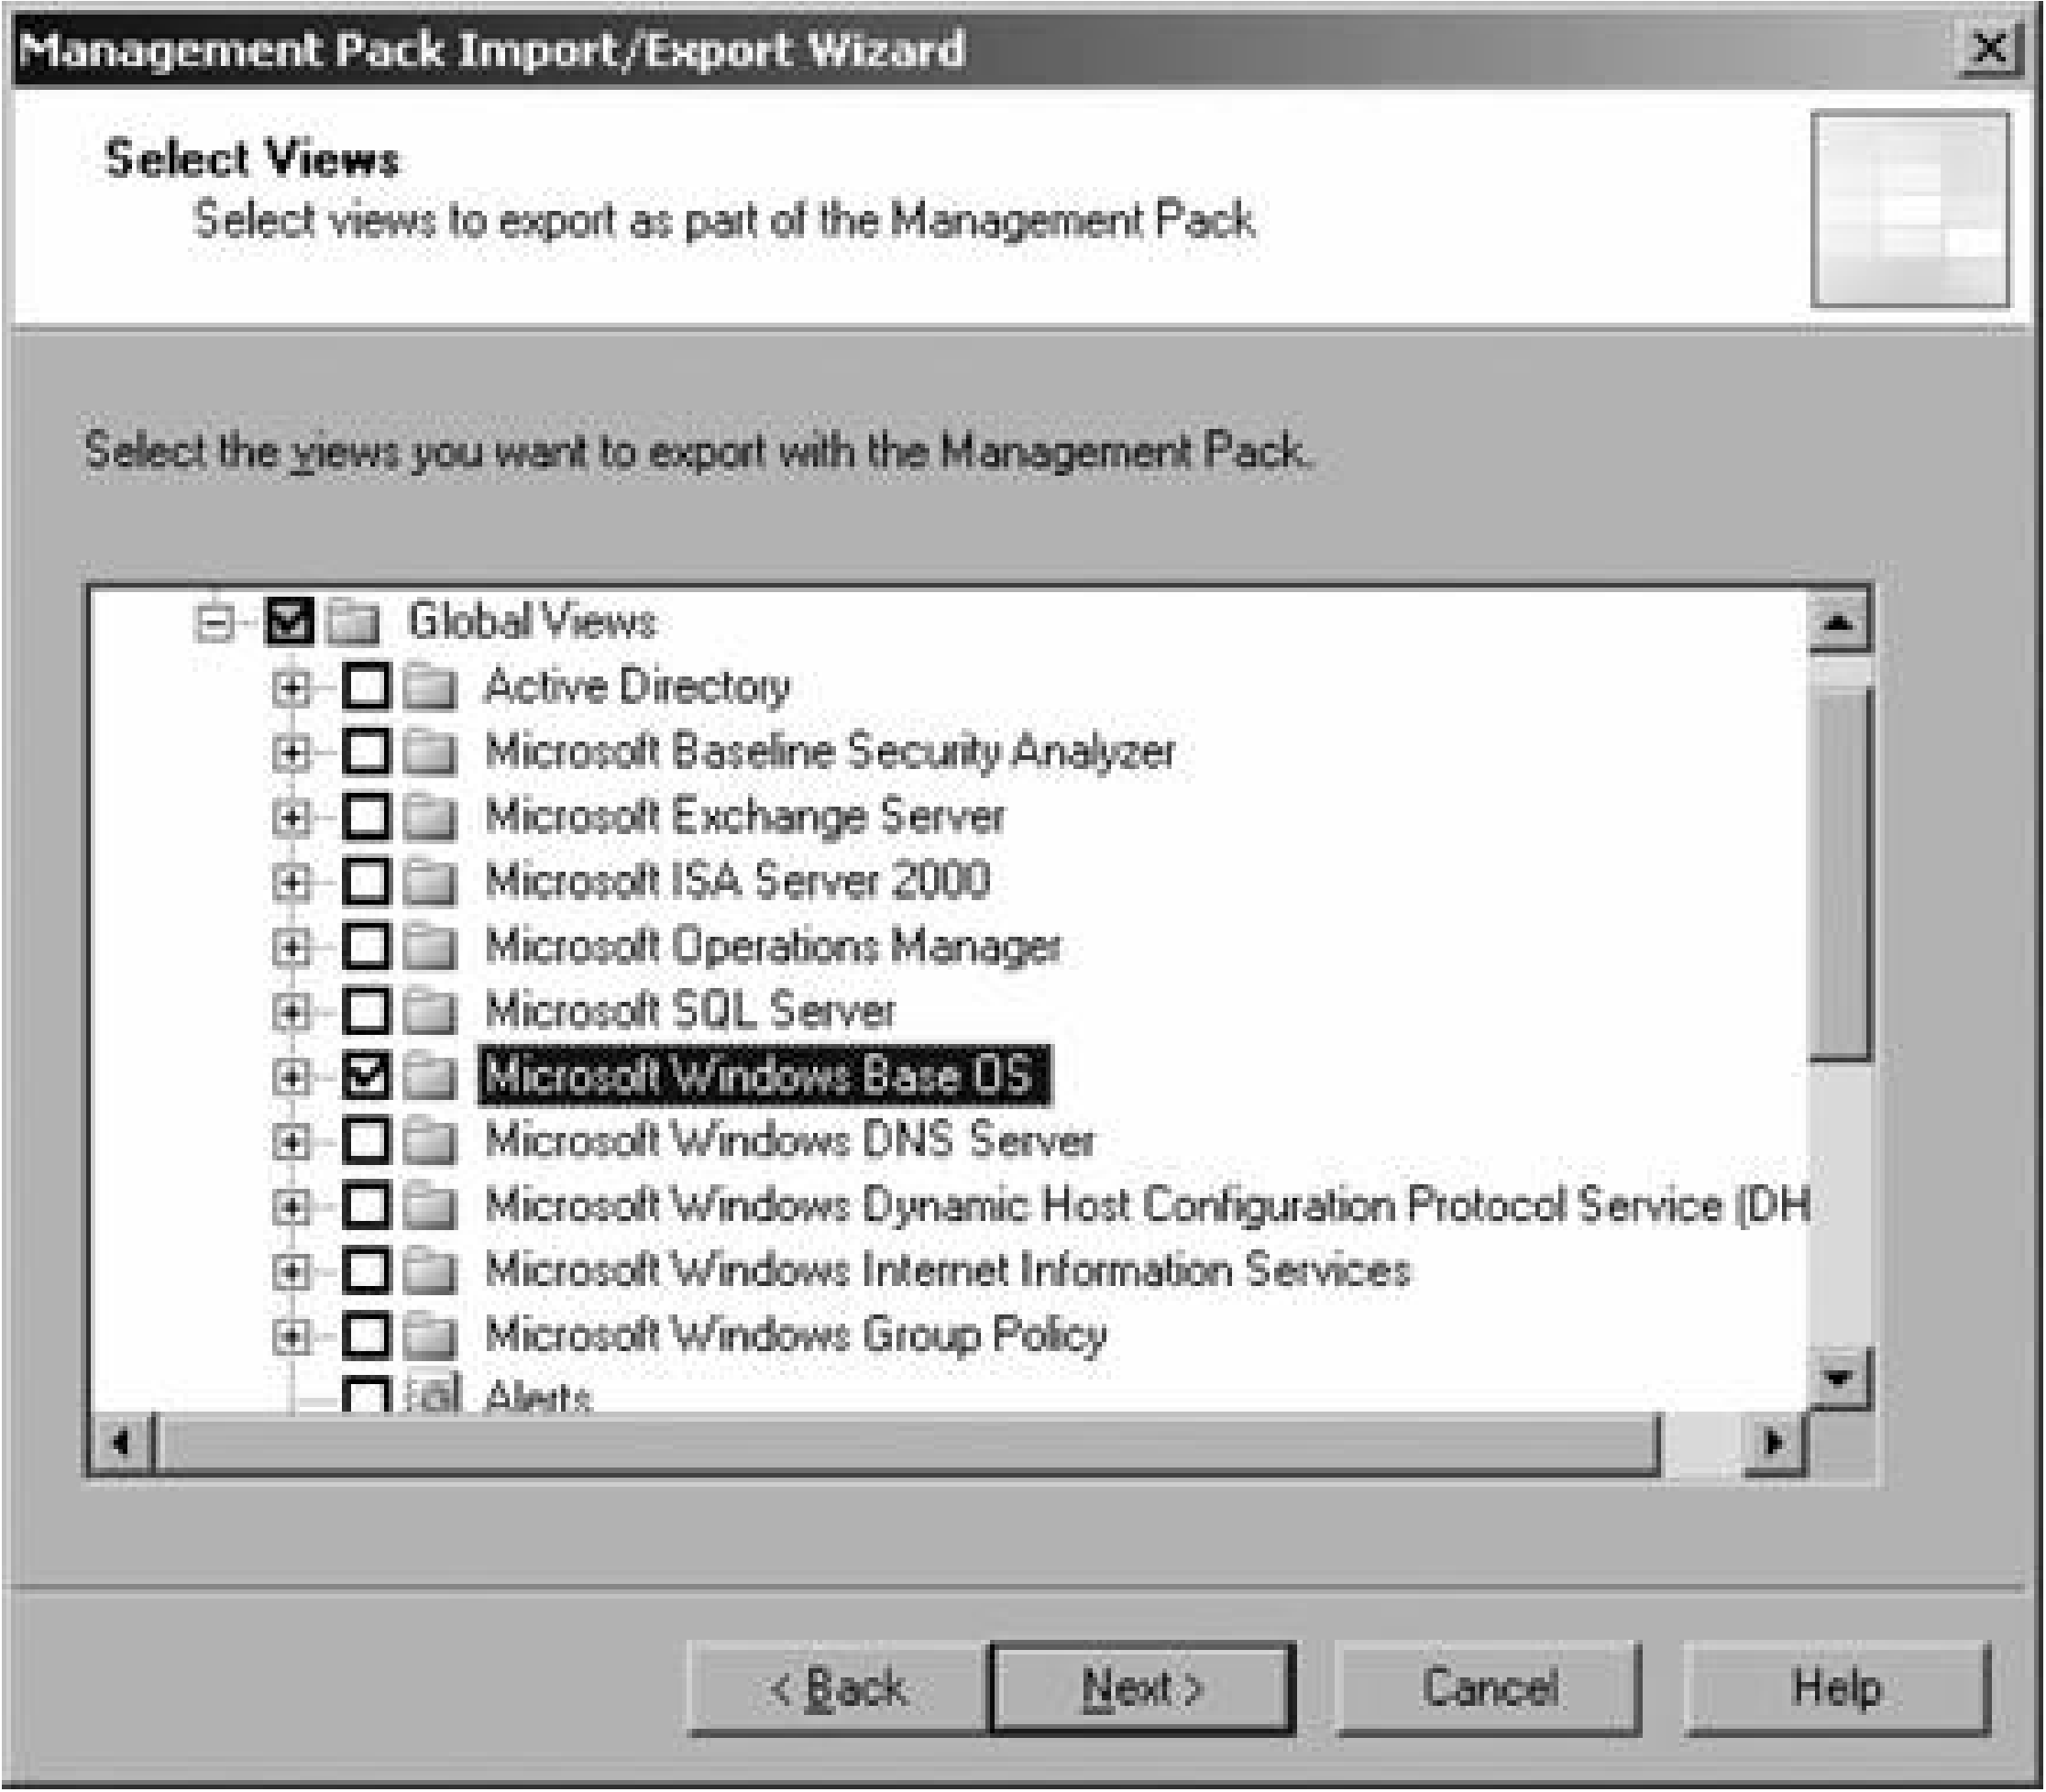
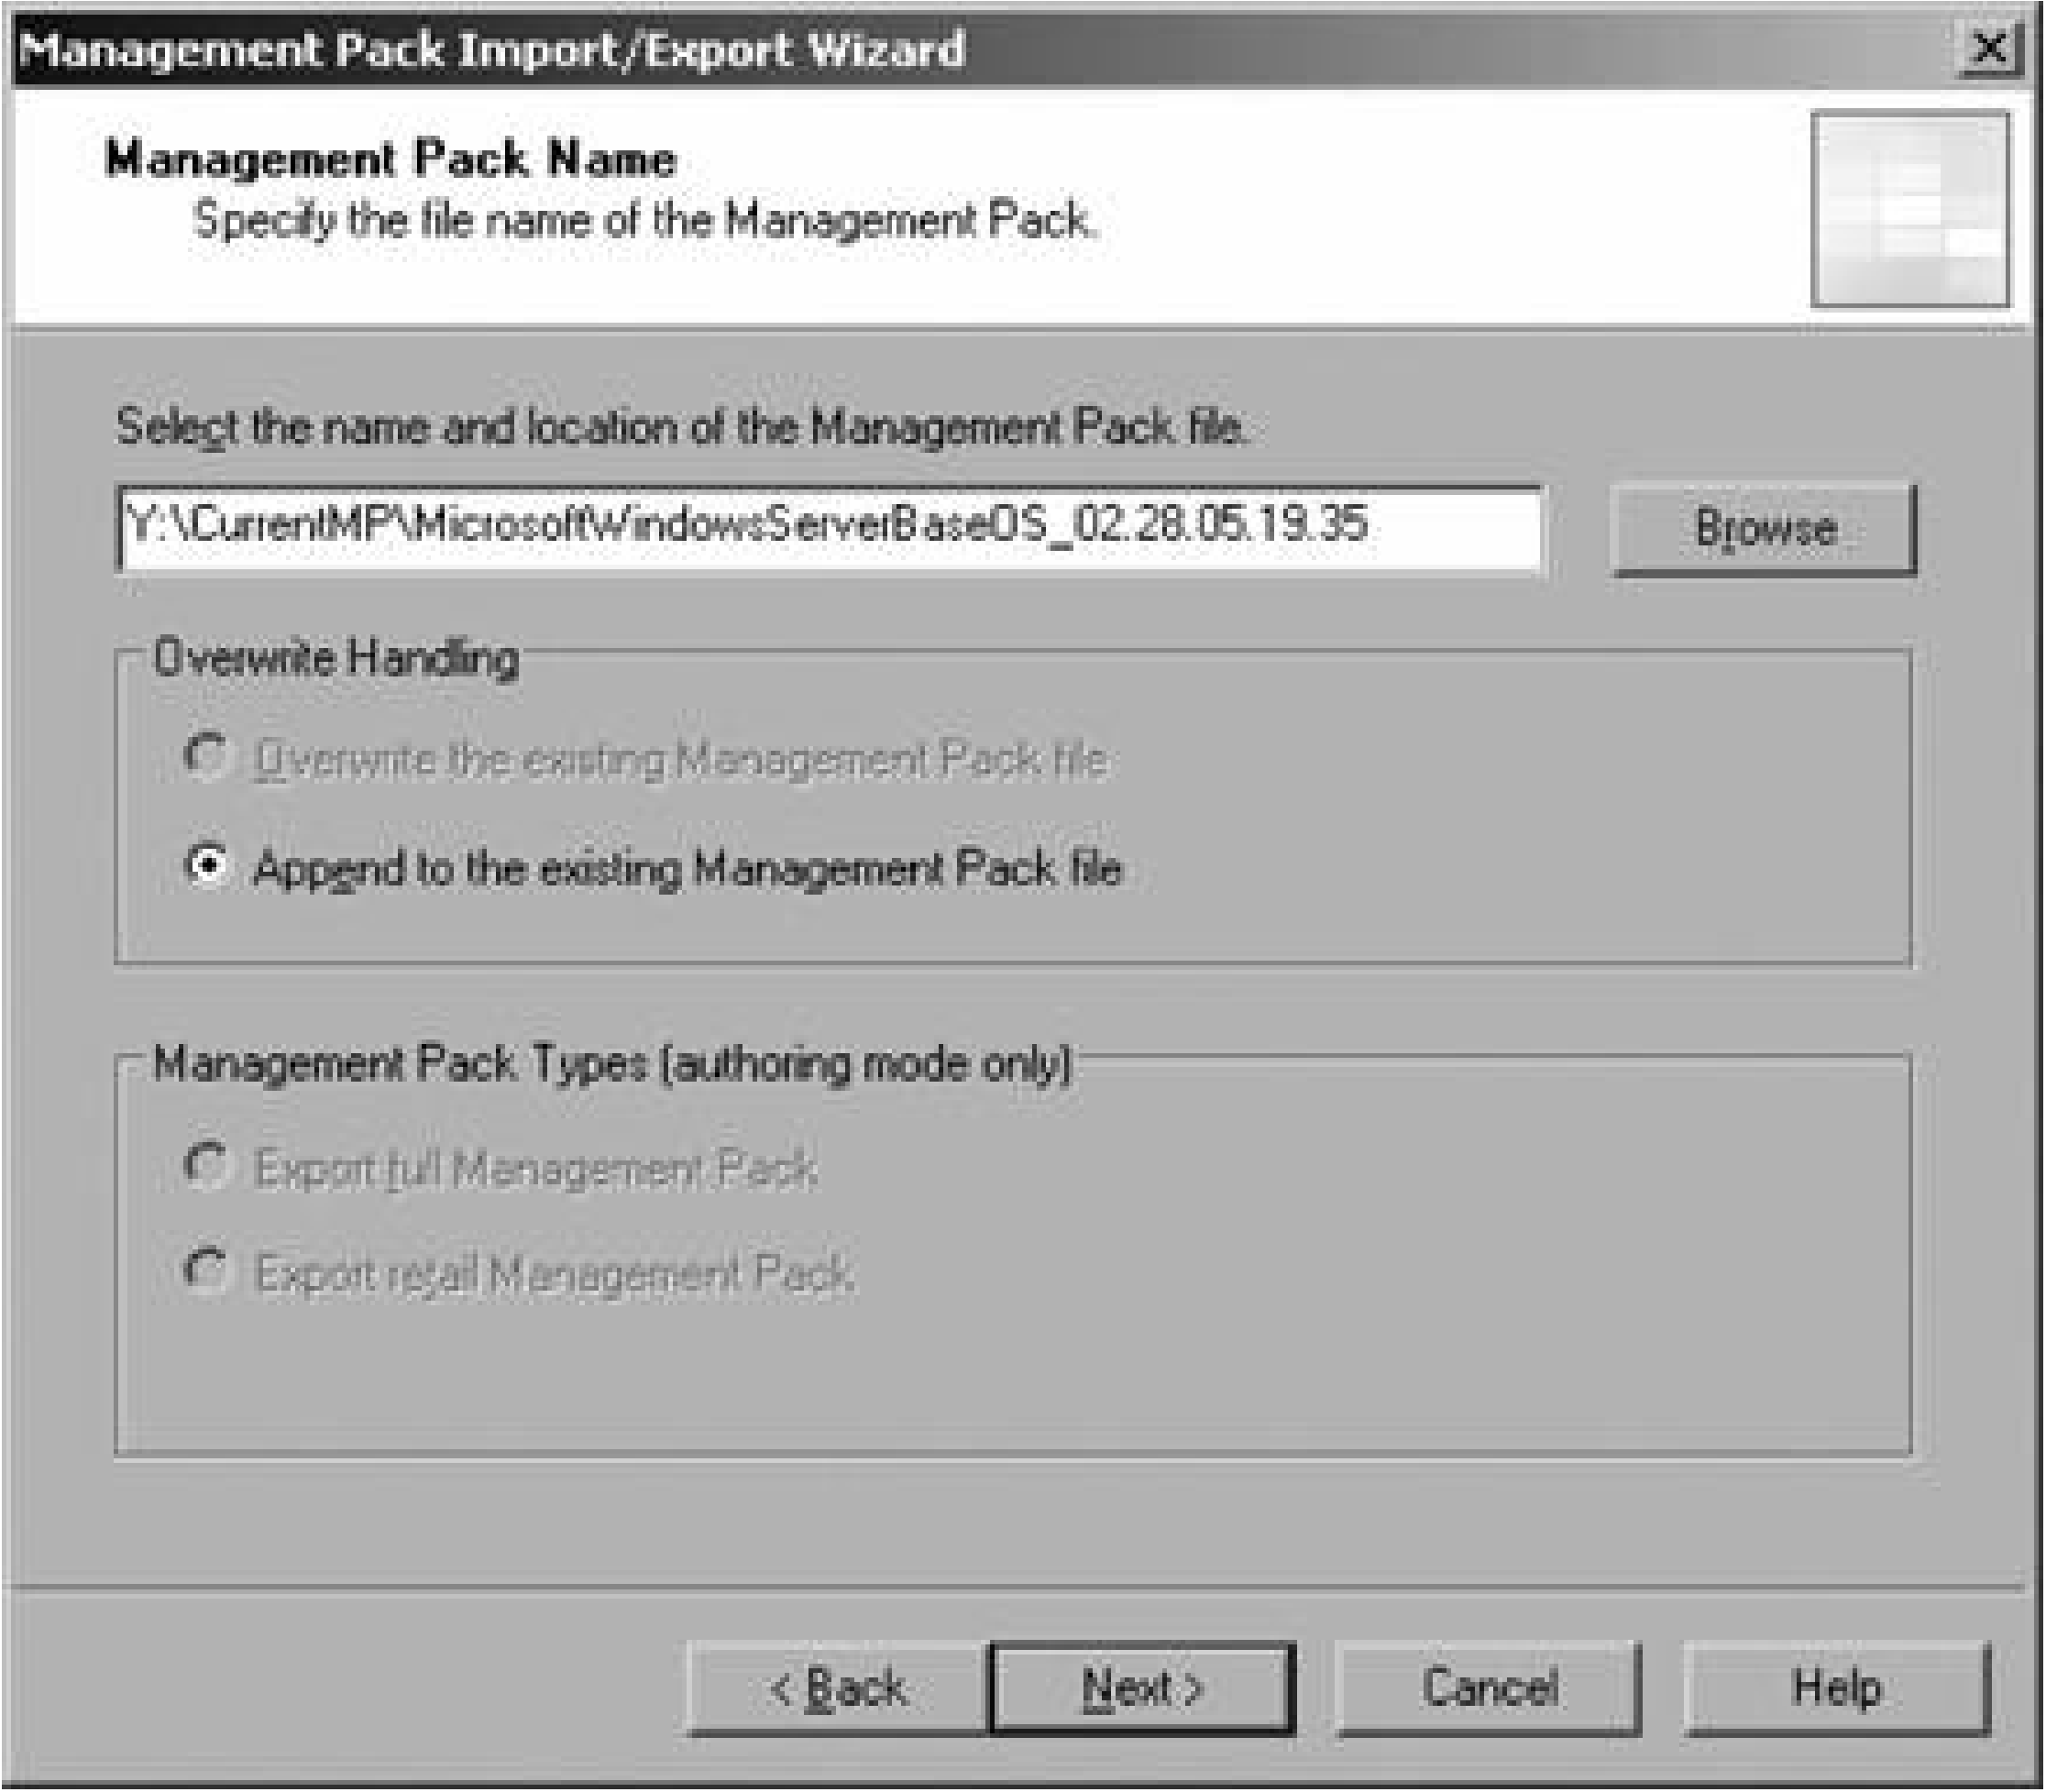


Figure 4-22. Specify the file location





If there had been an existing version of the management pack in the *CurrentMP* directory and the export the same name, the overwrite option in the Overwrite Handling box would have been enabled. Otherwise tool defaults to the Append option for creating new management packs.

Currently, the Management Pack Types controls are disabled because they aren't needed for the authori mode features to accomplish this export. Briefly, the authoring mode is enabled on the context menu of Rule Groups node in the Administrator console. When this feature is enabled, you gain edit access to the Knowledge tab of a created rule, and access to an Advanced Features tab that allows you to modify the relationships between parent and child rule groups. You also gain the ability to mark rule groups as dele they do not export and access the management pack's version ID field.

Clicking Next leads you to the summary page to confirm selections and finish the export process. Like th import portion of the tool, the export process displays a current status page that shows the progress of individual steps of the export. When this completes, a logfile of the export process can be created. Unless is an error during the export and you need the records for troubleshooting, it is not necessary to do this

In the management pack life cycle workflow, the management pack is at point 4 in Figure 4-1.

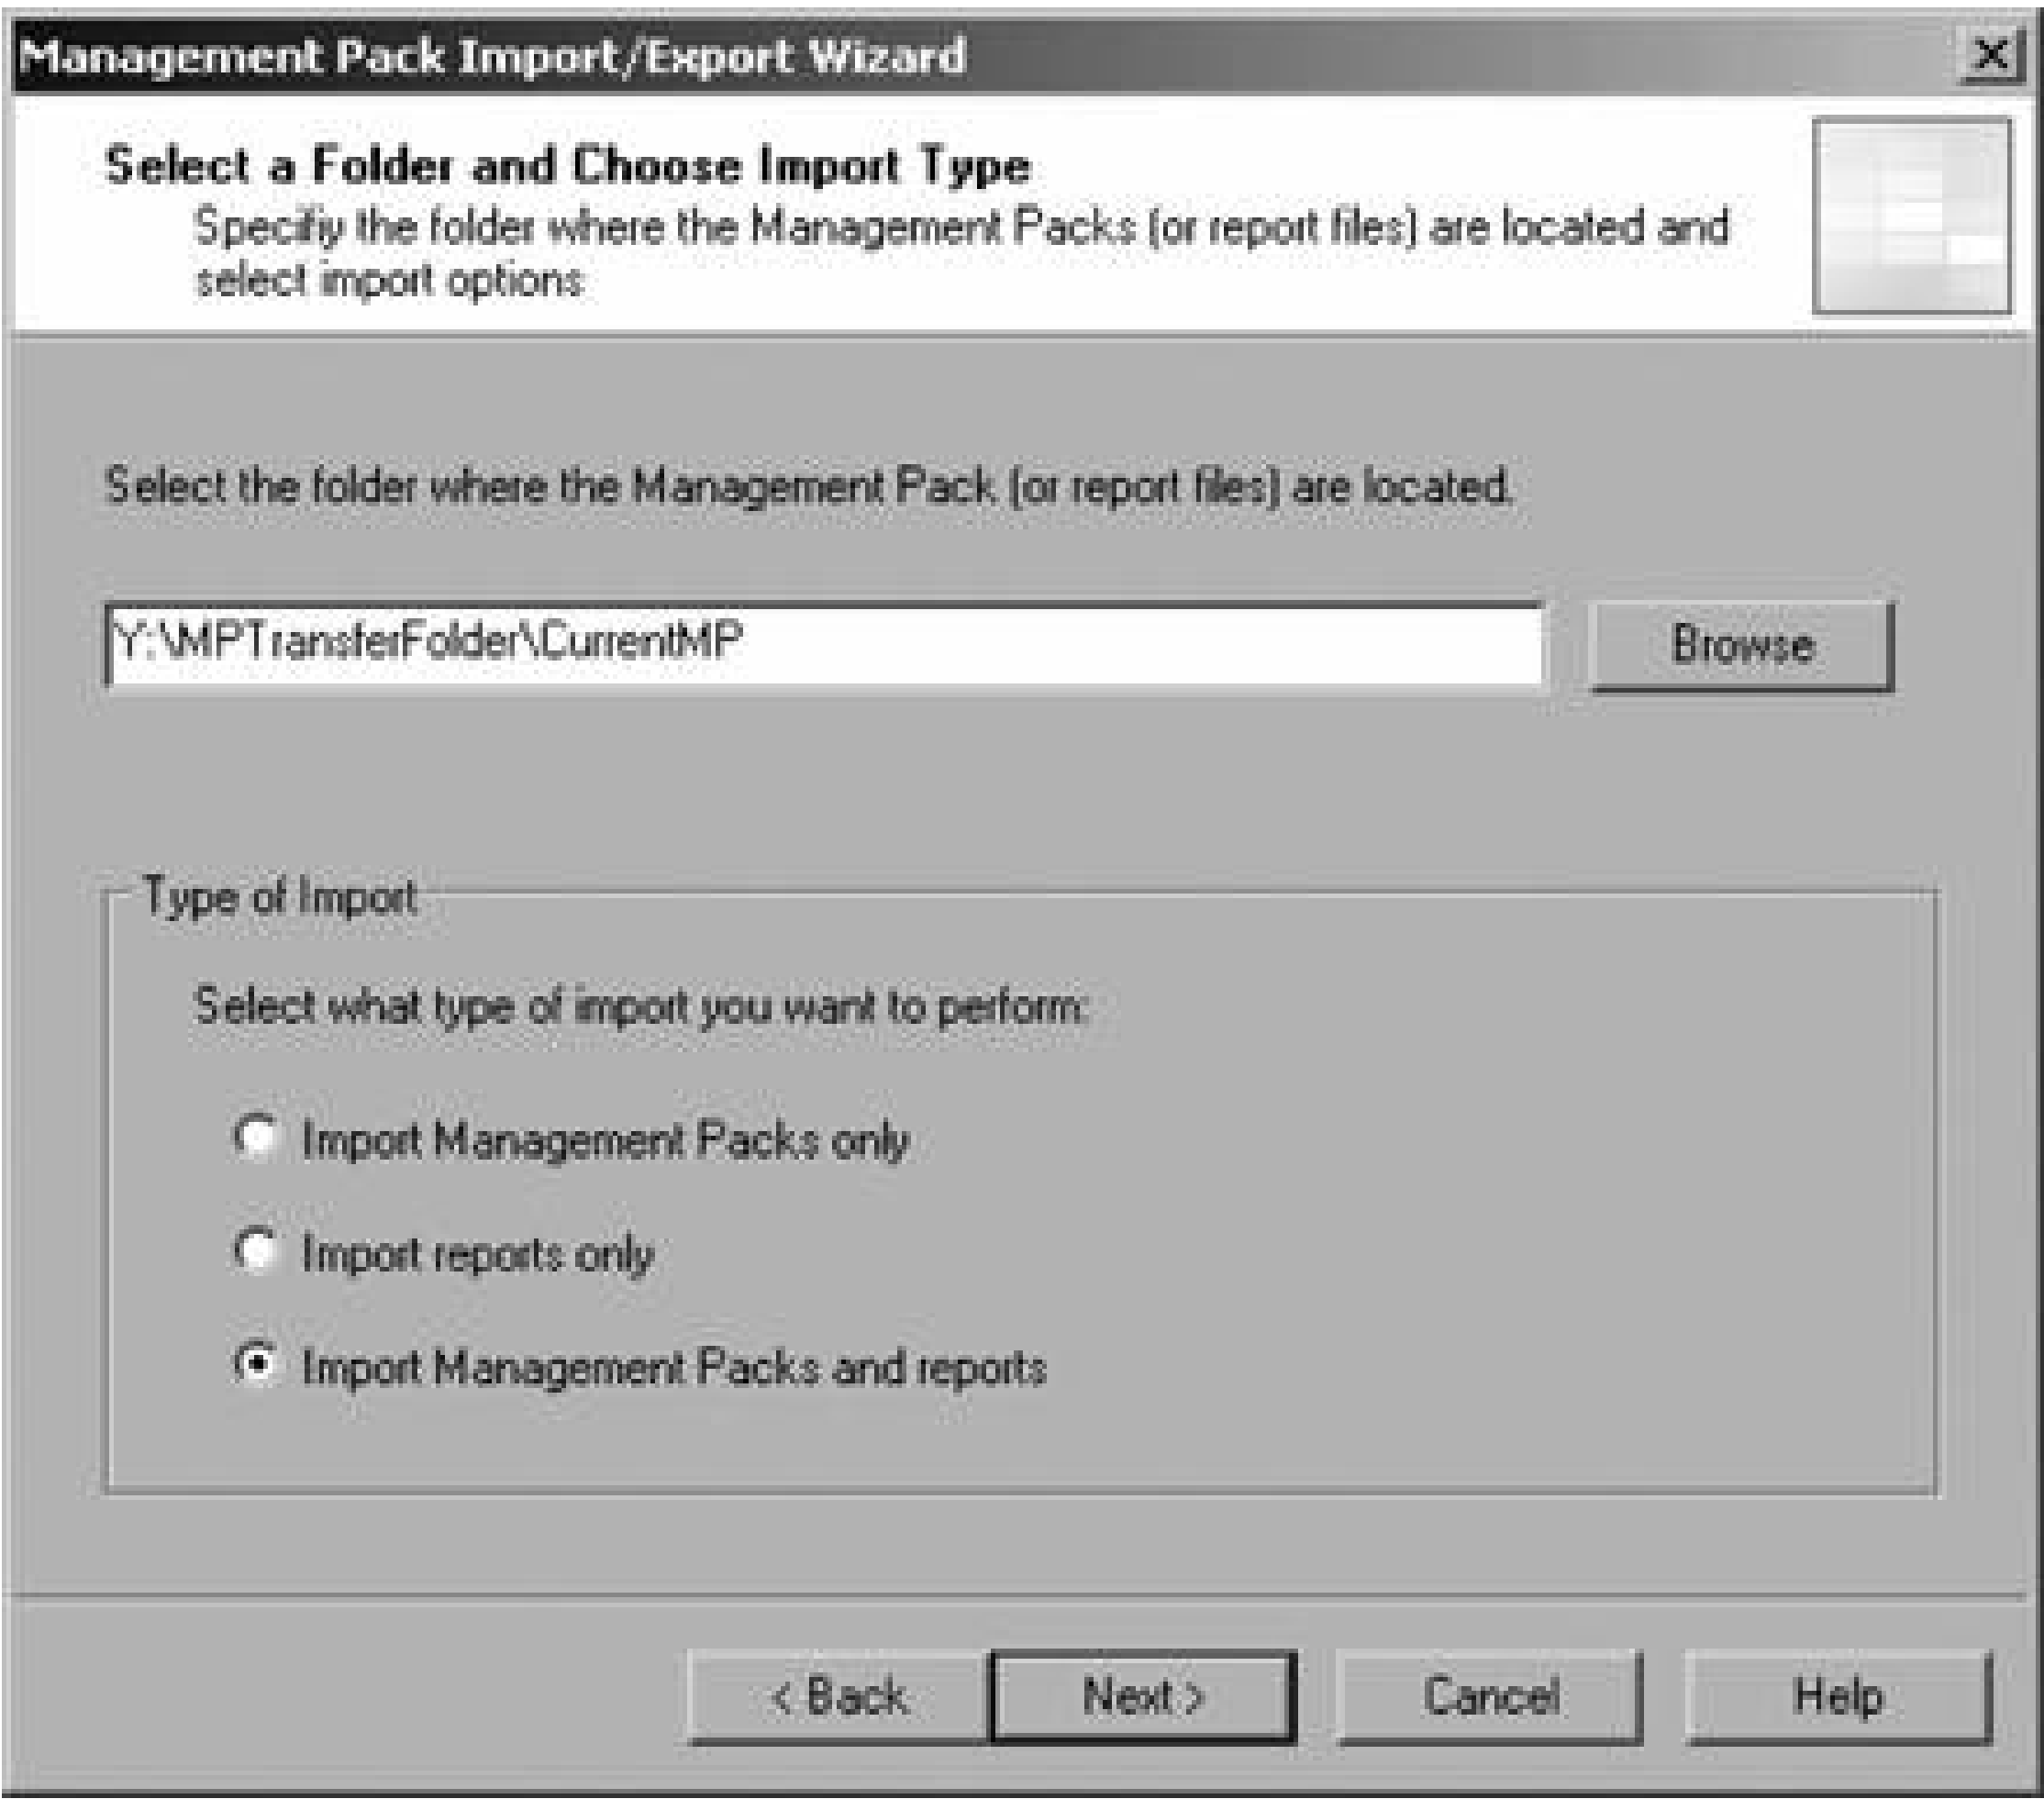
### 4.3.2. Import into Production

Importing the MP.1.1 version of the management pack into the production environment is only slightly c than importing it into the preproduction environment. The main difference between the two processes is the production management group (LKFPProdMG), has reporting services installed, so the management p report definitions can be imported at the same time as the management packs. Some management packl with a matching set of report definitions. They have the same name as the management pack, but end i

*Reports.xml*. These report definitions are still in the *IMPTransferFolder\VendorSupplied* folder, so if you to import them while you're running the Import/Export Wizard, you'll need to copy them to the *IMPTransferFolder\CurrentMP* folder, because the wizard can only browse files one directory at a time. Otherwise, the reports can be imported from their current location by running the import process again. 4-23 shows the choices page of the import process.

The next page has the update/replace/backup options (Figure 4-24) and is treated the same as importing Windows Server Base Operating System management pack into the preproduction environment. Select the Replace option because this is the first time that this management pack is being imported. The Replace completely overwrites the existing management pack, so all customizations and company knowledge will be lost. If this management pack already existed in the production management group, the Update option will preserve any modifications.

Figure 4-23. Importing a management pack into the production environment the first time



There might be situations where you need to completely overwrite the existing management pack. For example, if the existing management pack is no longer functioning correctly you may need to replace it with the known good backup to return to a functional state.

When a vendor-authored management pack is imported, MOM compares the vendor version numbers of the management packs. MOM will prompt you if you are about to overwrite a management pack with an older version, as shown in Figure 4-25. The version checking does not use your company's version numbers for comparison, so it will warn you if there is a conflict between two management packs that are based on the same vendor version number but have different company modifications embedded in them.

The next page has not been encountered in previous runs of the wizard; it simply allows you to select w report definition file to import (Figure 4-26). You can select multiple files if they are in the directory you browsing. You should always import the matched management pack and report definitions file at the sar The management pack contains event and performance rules that collect data that feed into the reports will see evidence of this about a day after you import a management pack. MOM 2005 will raise an infor alert telling you that data was collected for the reports and that the reports are ready to be viewed. Bas the management pack and the matching report definition are made to work together, so let them.

Figure 4-24. Choose the Replace option for the first-time import of a management pack into production

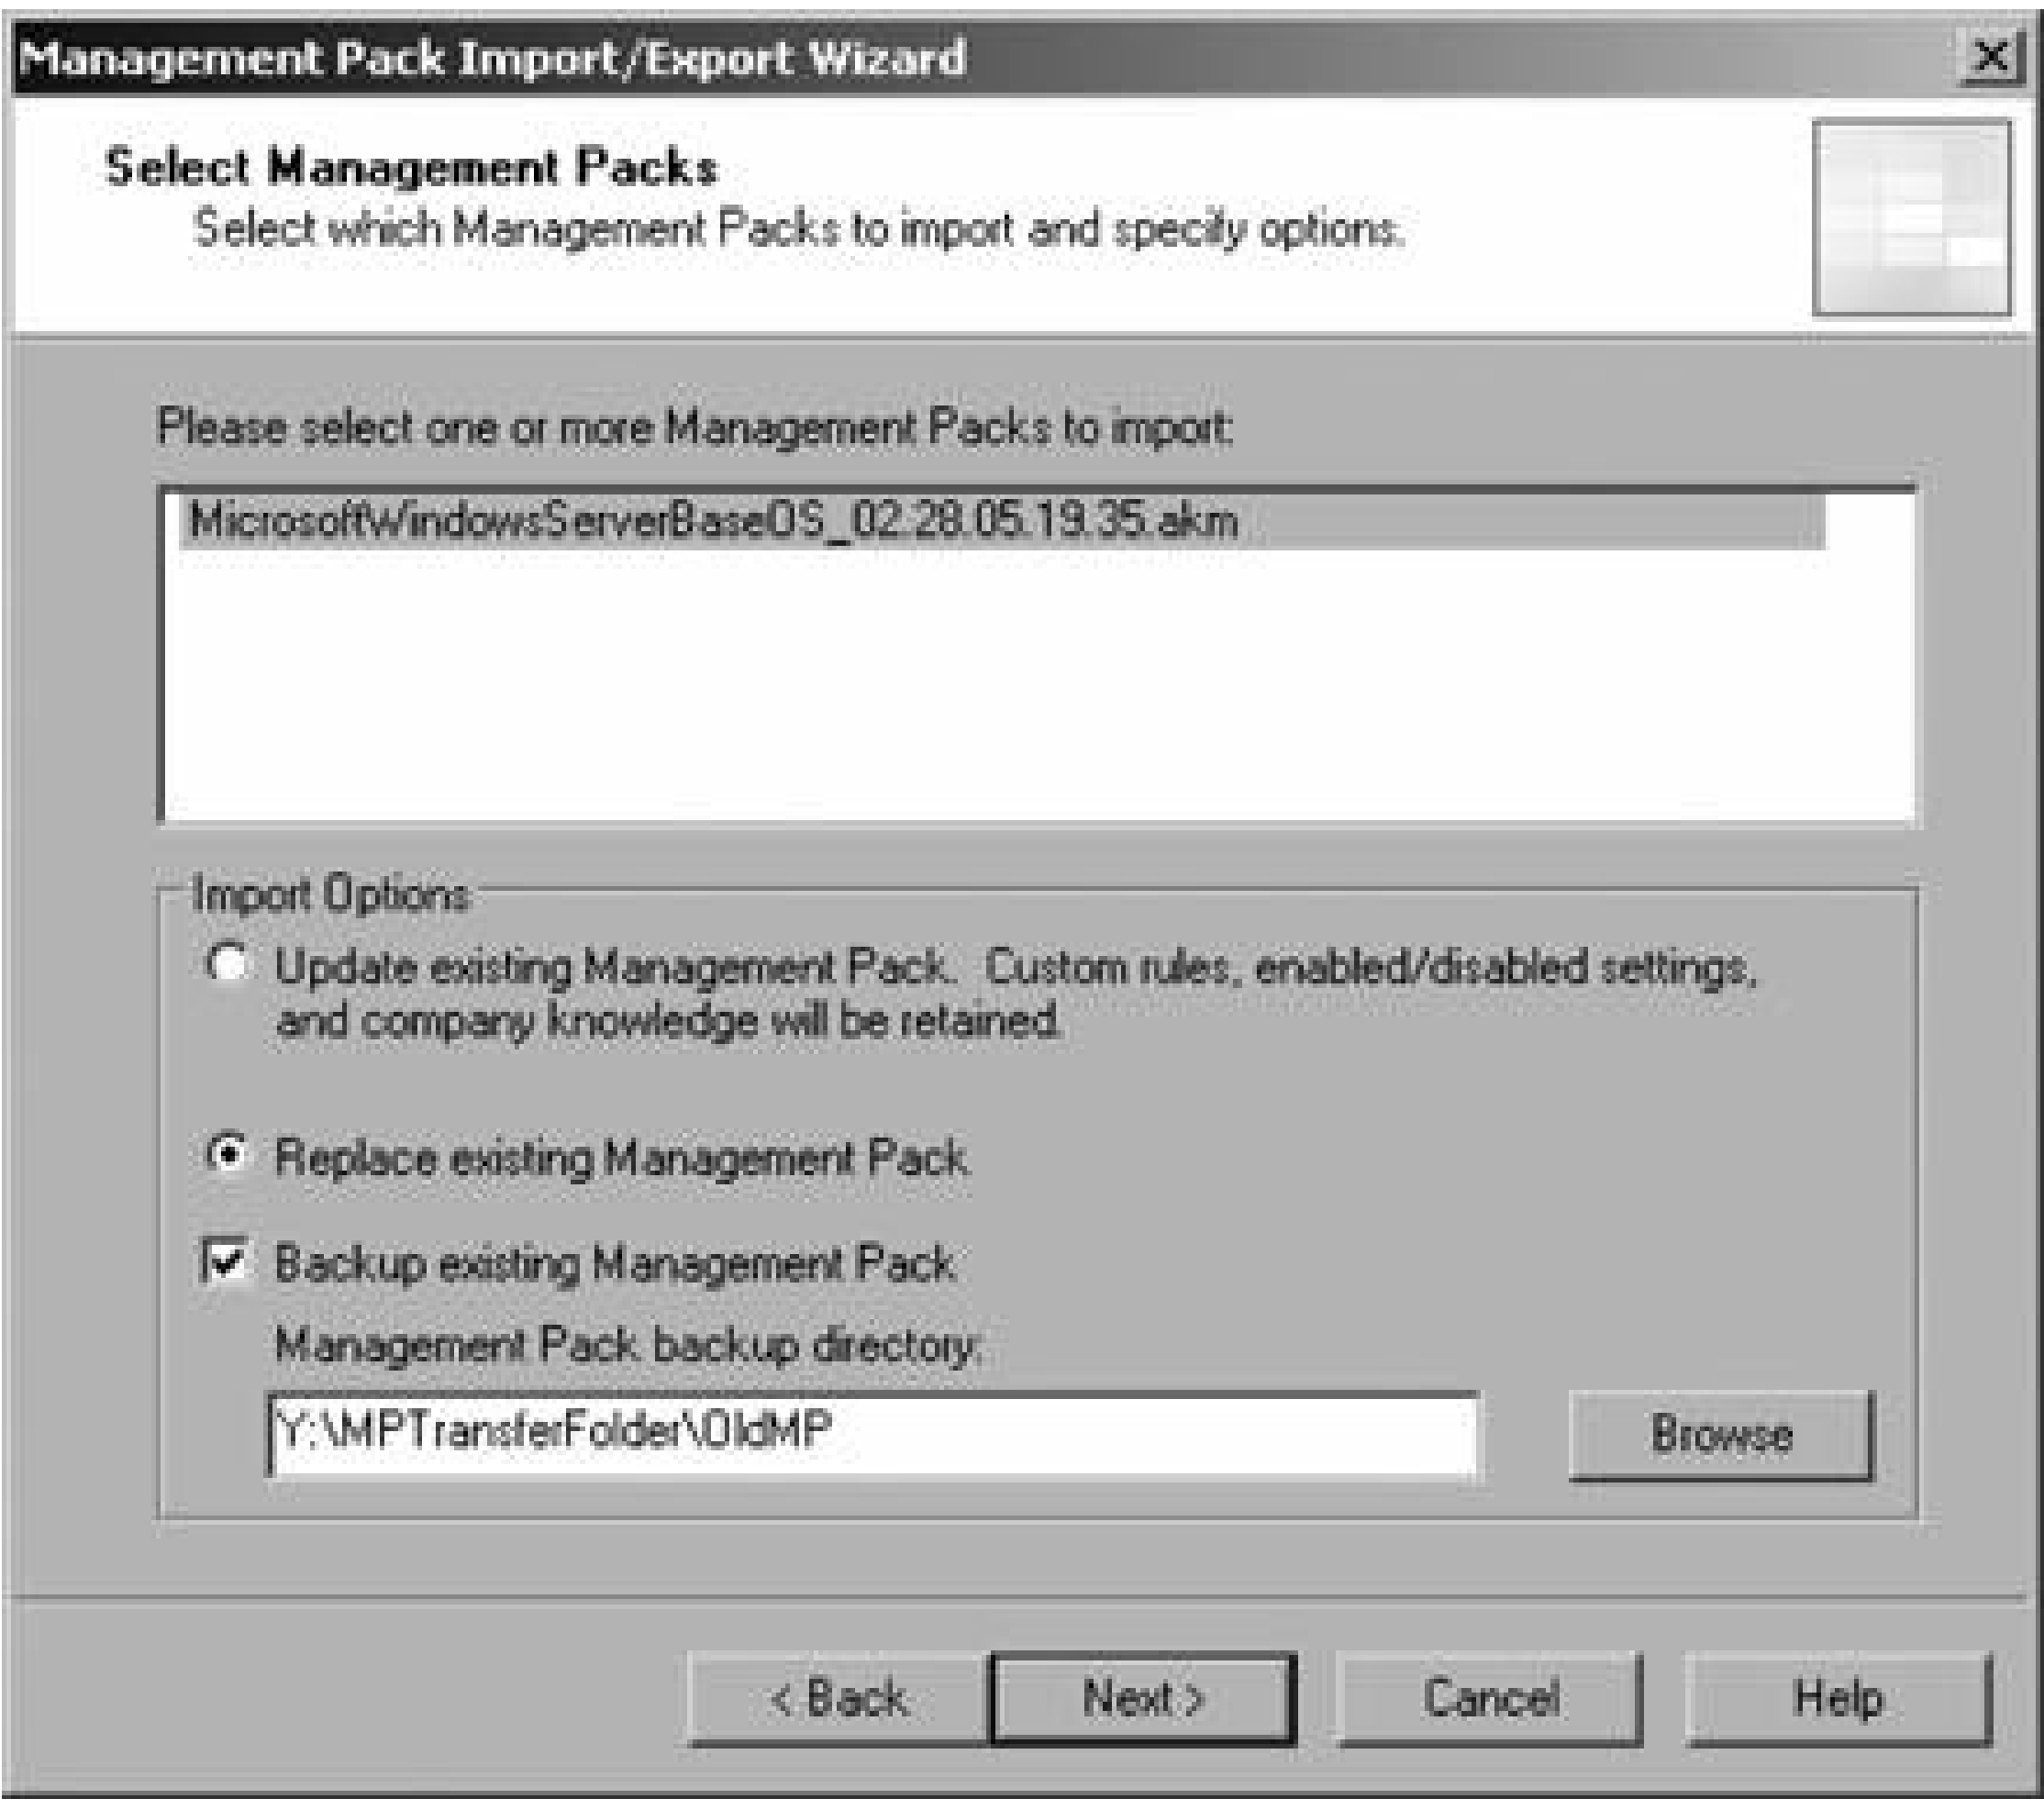


Figure 4-25. The import process automatically protects you from overwrit newer versions of rule groups based on the vendor version property





If the reporting services web site is not SSL-secured, and you have the "Always warn me about insecure connections" option enabled, you will receive a warning page. Since you are on the same network as the warehouse/reporting server, a secure connection is not really necessary. You can safely continue through warning.

After reviewing the summary page, allow the import process to finish. You will be presented with the now familiar Status Update page as the import process executes.

Figure 4-26. Select the reports that you want to import

The management pack is now active in the production environment. The registry key attribute discovery occurs and the managed computers will be sorted into computer groups. On the next agent configuration refresh cycle, the appropriate processing rules will be returned to the agents and monitoring will begin. The management pack is now at point 5 in the life cycle workflow (refer back to Figure 4-1).

### 4.3.3. Evolution of the Management Pack in Production

Once a management pack is brought into production, it will be changed by further tuning and capturing company troubleshooting knowledge. The longer it is in production, the more it will become specific to your environment and less like the vendor-authored original. Tuning production management packs helps eliminate noise alerts—those alerts that give incorrect, unusable, or undesired information. This is accomplished with the same techniques used for tuning the preproduction environment, with the addition of more extensive use of overrides. You can't perform all of the necessary management pack tuning in the preproduction environment simply because the two environments—including network, applications, and load—are not identical. So, you may find issues in production that you didn't see in preproduction and you have to respond to them there.

You can perform minor tweaks to rules to make changes, such as in the fields that are included in the alert description or by adding a notification group. These minor tweaks don't change the core parameters of a rule (provider, criteria, response, and vendor knowledge); they are, for the most part, cosmetic. If you need to change any of the core parameters of a rule, then you are no longer tweaking a rule but creating a new rule or modifying an existing one.

When a new management pack is first imported, it will start returning data and raising alerts. Expect the highest volume of alerts in the first 24 hours after a management pack has been installed. MOM 2005 will go through everything in your environment and tell you about issues that you weren't aware of.

As you go through and troubleshoot the issues that caused the alerts, capturing the troubleshooting steps for each alert's Company Knowledge tab, and resolving the alerts as outlined in Chapter 1, the volume of alerts will begin to decrease. A few weeks into the production deployment of the management pack, certain alerts will keep cropping up that don't provide any useful or actionable information. You may not want to be alerted about an event, or there may be nothing that can be done to resolve the issue in the immediate future. For example, you may not be interested in knowing about every time a user maps a printer, or having the baseline security analyzer tell you that your domain's NTLM security setting is too low when it has to be that way to accommodate Windows 95 machines. These types of alerts are "noise alerts" for your environment, and you may need to further tune the rules and rule groups to reduce them.

Be on the lookout for false positive alerts and false negative alerts. For example, a false positive alert (an alert that is generated) would be if MOM indicates that a server is down when it is not. A false negative (no state change) would be if in the State view a service health indicator remains in the success state when the service has actually failed.


The easiest way to eliminate a noise alert is either to disable or delete the rule that generated it. In case you really don't want the information, this is probably the best choice of action. Another way to reduce noise alerts is to configure overrides for a rule.

After the management pack is moved from preproduction to production, the first task is to complete the configuration of those items that are not included in the management pack, namely the operators object, notification groups, and overrides.

## 4.3.4. Overrides

Overrides provide additional flexibility by controlling how a rule is applied to target computers. Without overrides, a rule can only be enabled or disabled for all the computer groups it is applied to. Further, it can have a single criterion by which to judge if matching data is found. There are four types of overrides: rule enable/disable, performance threshold, script parameter, and state alert severity. Overrides direct MOM to apply a rule this way to these machines, but apply it a different way to those machines.





Be cautious with overrides and, as with all other things, test first.

#### 4.3.4.1. Rule Disable Override

The Rule Disable Override is available on each of the 10 types of rules. It allows you to selectively disable a rule that is enabled by default. You can also create an exception to the disablement override. The override criteria is simply the pairing of a target (a computer or a computer group) and a value. Multiple target and value pairs can be defined for a particular override rule. These criteria are placed in an ordered list and evaluated from top down. This way, a pair that is lower in the list will supersede all pairs that precede it in the list.

For example, in Figure 4-27 the "Enable rule-disable overrides for this rule" option has been enabled for the "LKF software update installation failed" rule. All this means is, "Do you want to allow the creation of override criteria for this rule?" This is an event rule in the Microsoft Windows Servers Base Operating System\Windows 2003\Core System Components and Services rule group. This rule group is bound to the Microsoft Windows 2003 Servers computer group. That computer group contains all Windows servers running the Server 2003 operating system regardless of the role (domain controller, cluster server, member server, and so on).

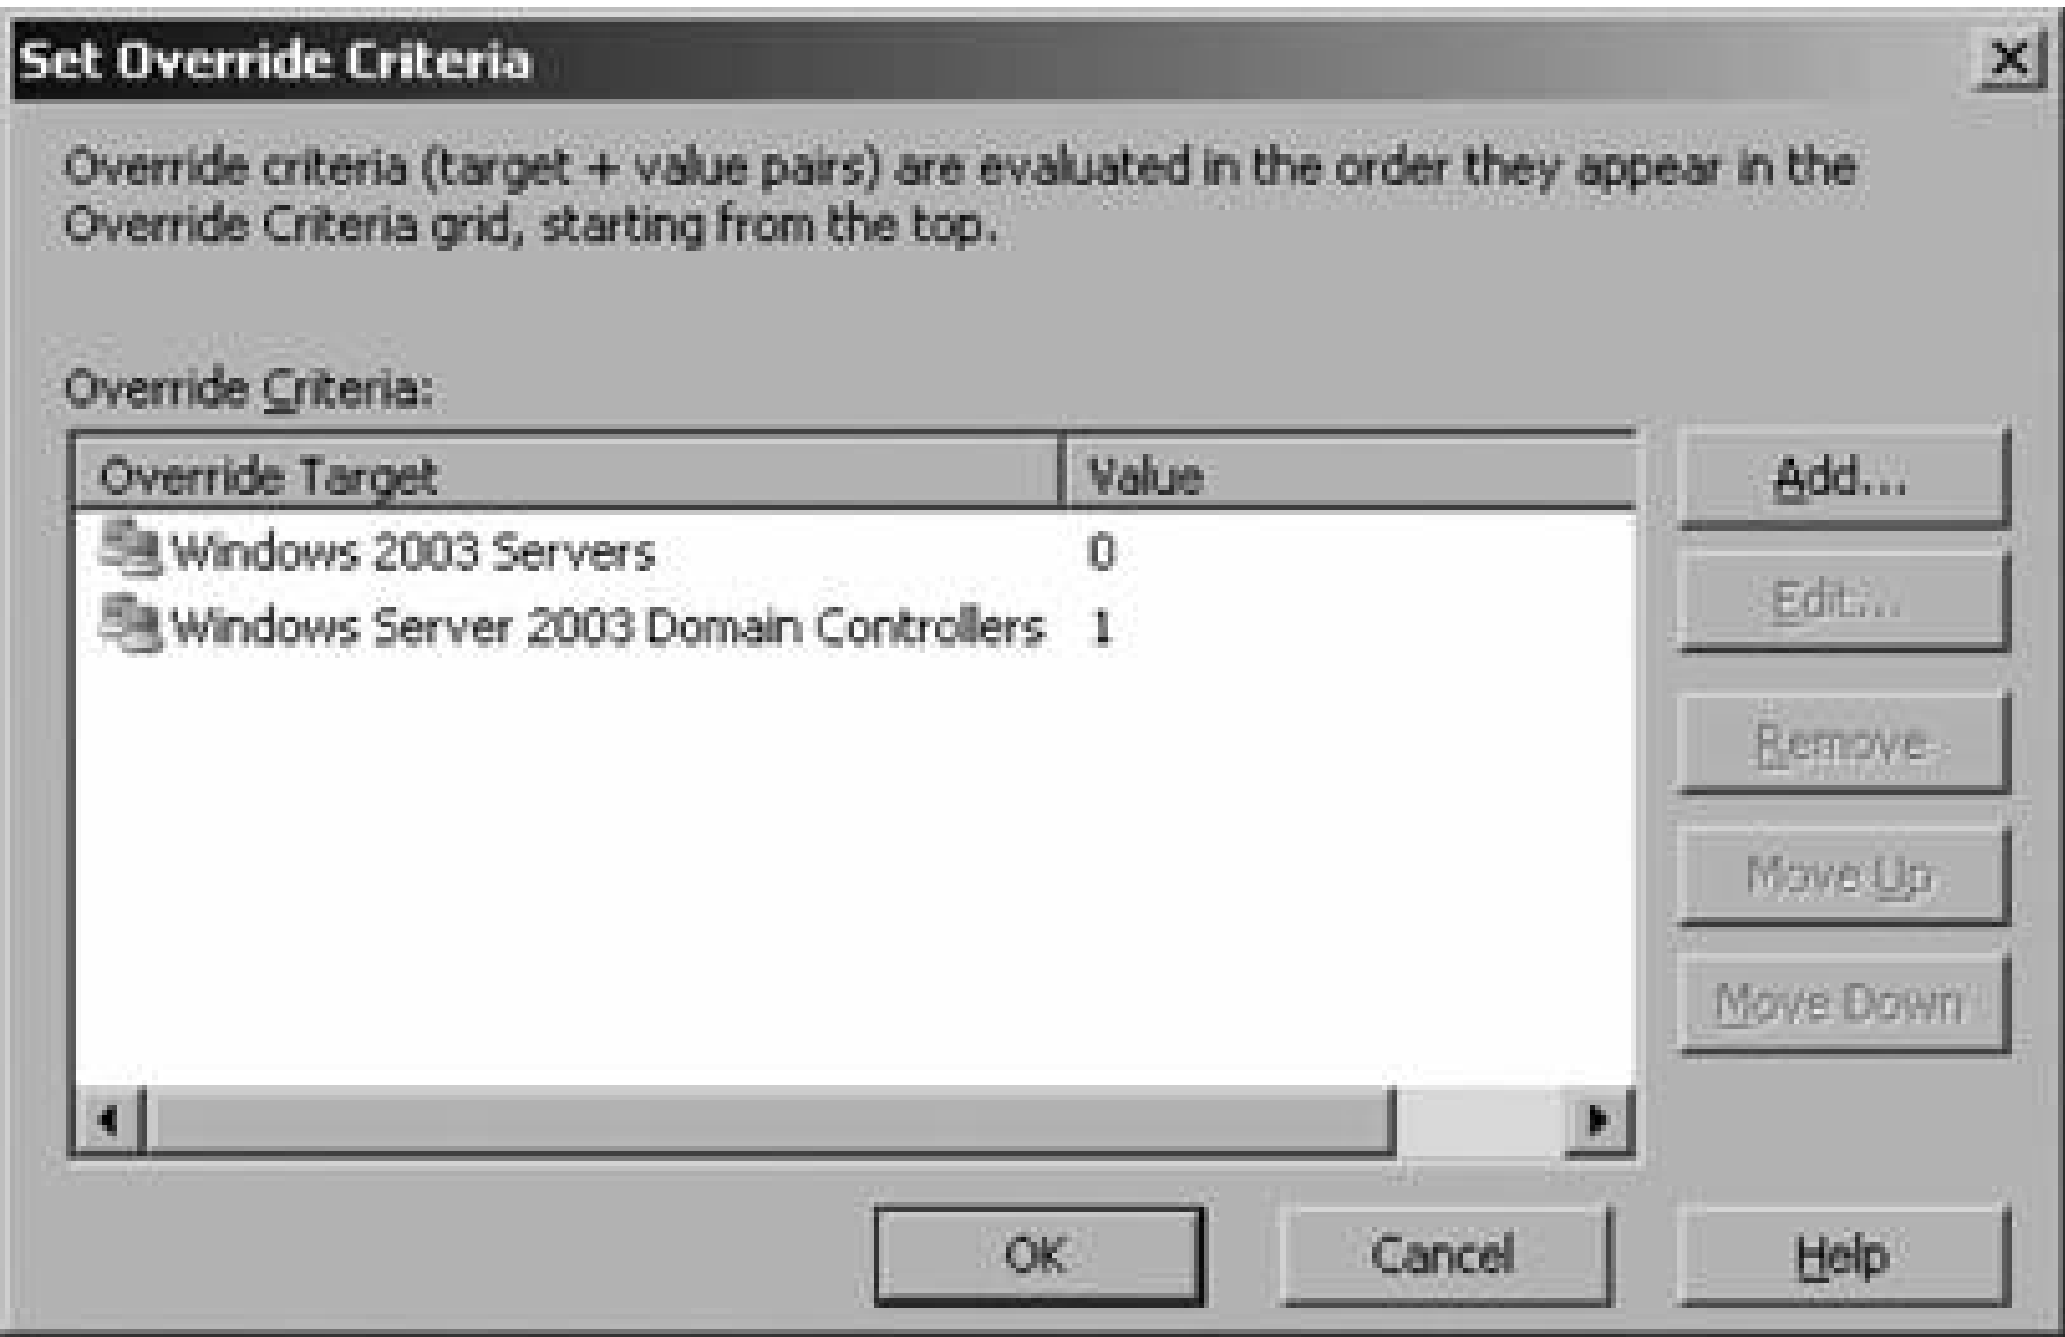
Figure 4-27. Rule-disable overrides are enabled for this rule

Once you've allowed disabling overrides, click on Set Criteria to bring up the page shown in Figure 4-28. The target + value pairs for the override are defined here.

Here is how this rule will be applied. First, because the rule is enabled, it will be distributed to the agent on all the computers in the associated computer groups. Then, if that server is a member of the Windows 2003 Servers computer group, the rule is disabled (the 0 value). If it is also a member of the Windows Server Domain Controllers computer group, then it is enabled (the 1 value). If the server is not in the Windows Server computer group, the rule will be applied.

Figure 4-28. Overlapping override criteria



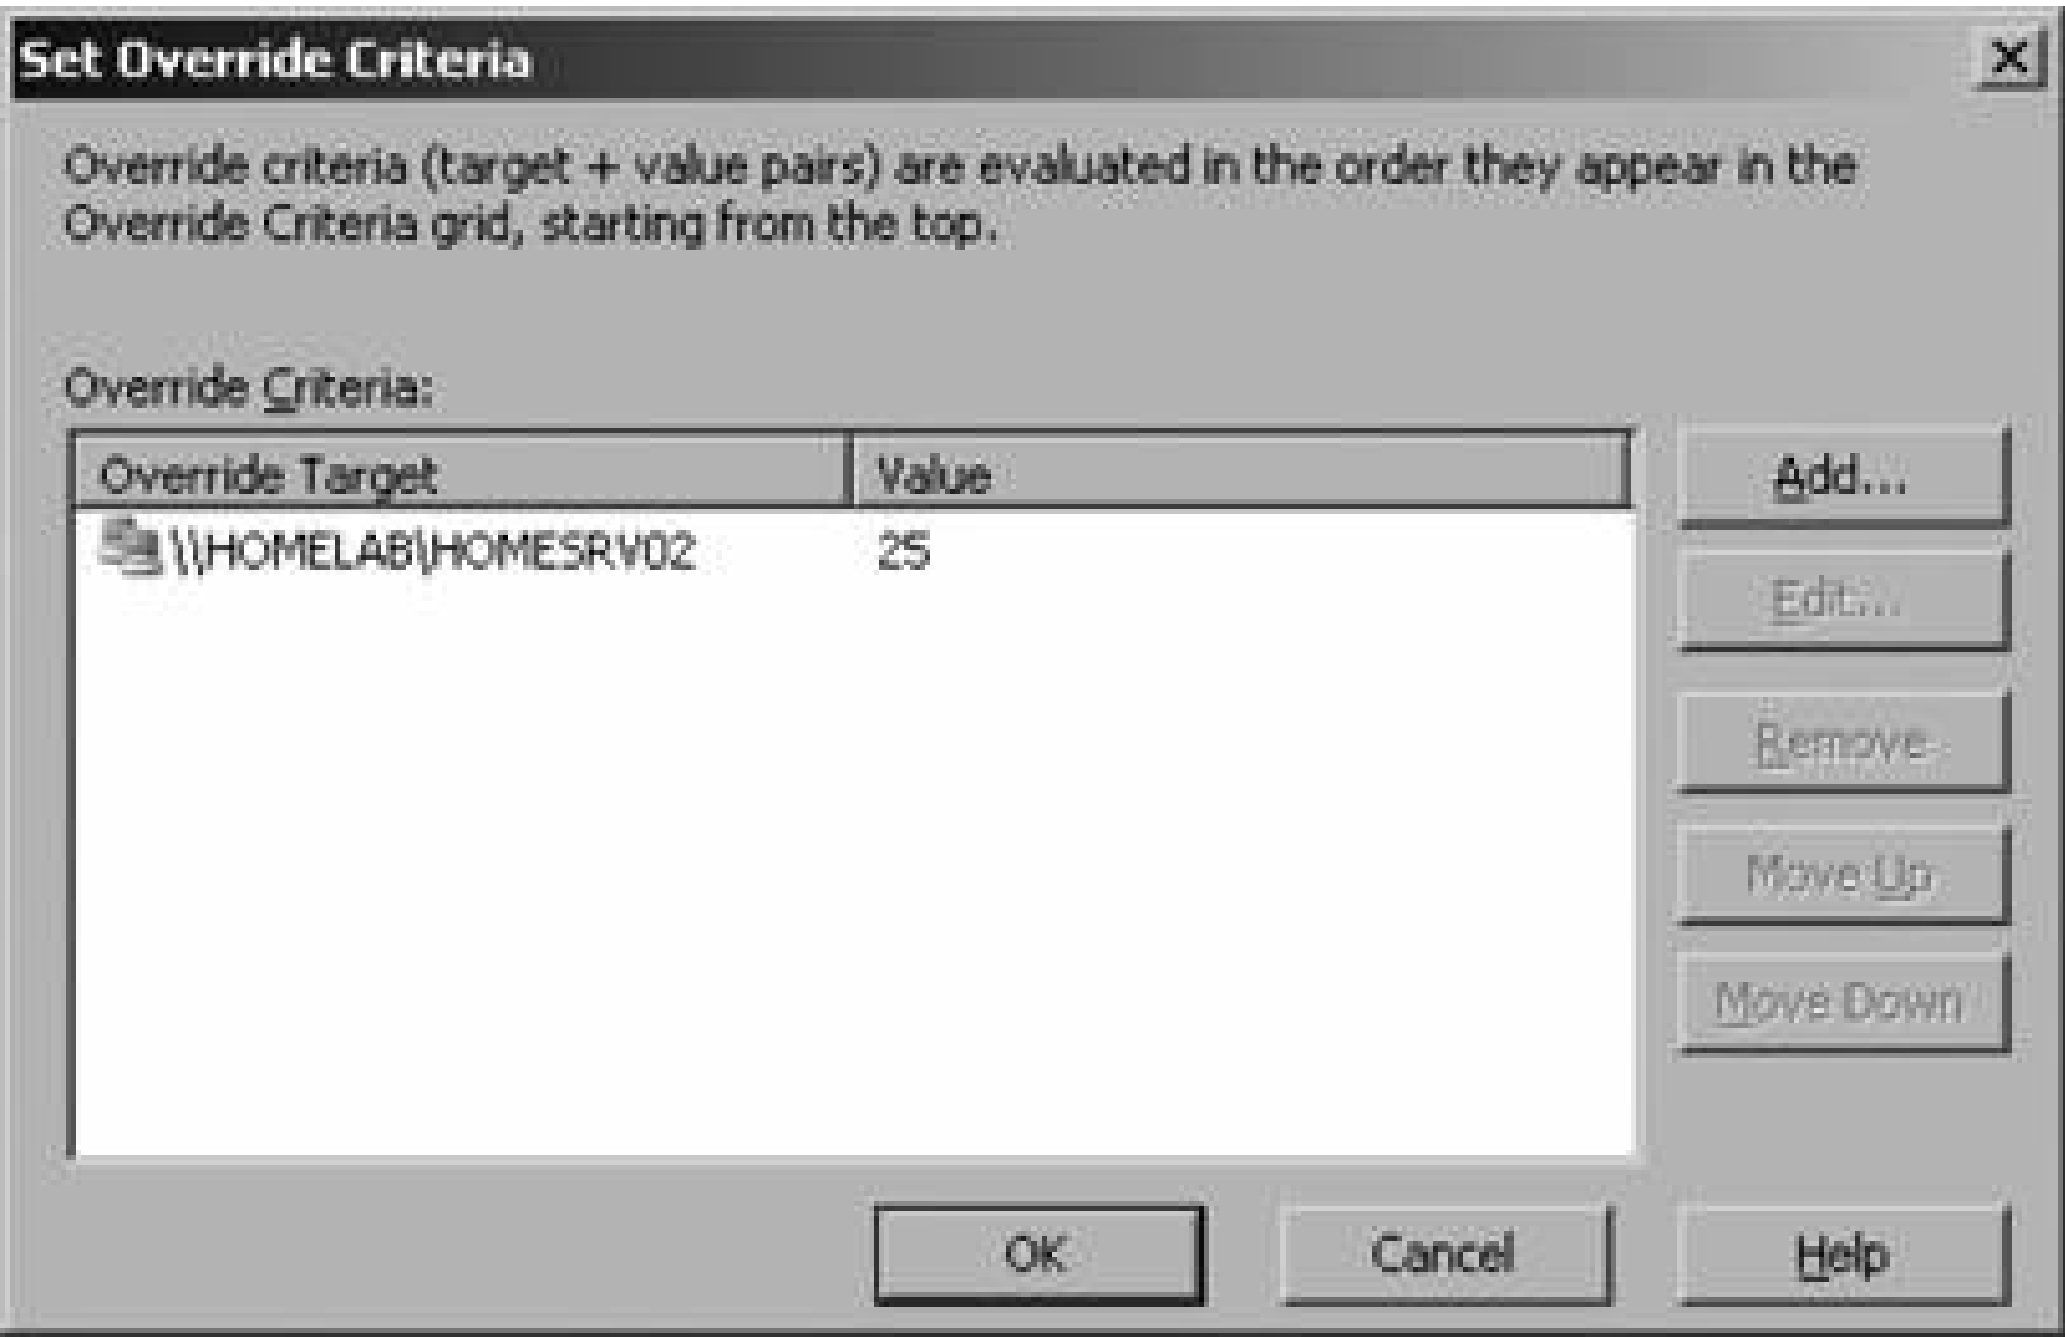


4.3.4.2. Performance threshold override

In addition to the enable/disable override, performance fsthreshold rules have another type of override that can be configured. Performance threshold rules generate an alert when a performance monitor counter value crosses a configured threshold, for example, when the actual processor utilization exceeds a certain percentage. Default threshold values can be overridden for computers or computer groups that are defined in a target value pair in the Threshold Override Set Criteria page. Figure 4-29 shows a performance threshold rule that generates an alert when the measured value exceeds 50 percent. For the sake of this example, it is not important what value is being tracked here. The "Enable threshold overrides for this rule" checkbox has been selected. The Override Name field is automatically populated with a machine-generated name. In the Set Override Criteria dialog box, a value of 25 has been specified for *homesrv02* (see Figure 4-30 ).

Figure 4-29. Threshold override enabled

Figure 4-30. The threshold override criteria for a performance threshold processing rule



### 4.3.4.3. Script parameter override

If a script is run from a command prompt, parameters could be passed to the script when it is called if it was written to accept them. This is the same concept as passing parameters (variables) to an executable from a command prompt. A simple example can be seen in the ping command. When ping is run, you must provide the NetBIOS or FQDN name of the device that you want the tool to run against.

```
C:\>ping avalon.homelab.lab
```

The event and performance rules that use scripts often pass parameters to the script. A script parameter override allows you to define multiple parameters to pass to a script from a single rule. For example, the Microsoft Windows Servers Base OperatingSystem\Windows2003\State Monitoring and Service Discovery\Performance Threshold:Processor\%Processor Time threshold exceeded rule calls (Microsoft Windows Base OS CPU Overload Script) as you can see on the Responses tab in the Rule Properties (shown in Figure 4-31 ).

If you select the script name and click Edit, the Launch a Script dialog box appears and you can modify how the rule interacts with the script (see Figure 4-32). In the lower box, select the script parameter that you want to enable and click Edit Parameter. Here the CpuPercentage parameter is selected.

The Edit Script Parameter page has the now-familiar Enable Overrides checkbox and the Set Criteria button. The Set Criteria button takes you to the Set Override Criteria page, which is identical to that shown in Figure 4-31. In this case, you would enter the target and value criterion of `\\HOMELAB\\HOMESRV02` and `75` , and save the changes. The Launch a Script page updates the CpuPercentage variable with an icon (shown in Figure 4-32) to indicate that an override criterion has been specified for that parameter.

Figure 4-31. A performance threshold rule calls a script

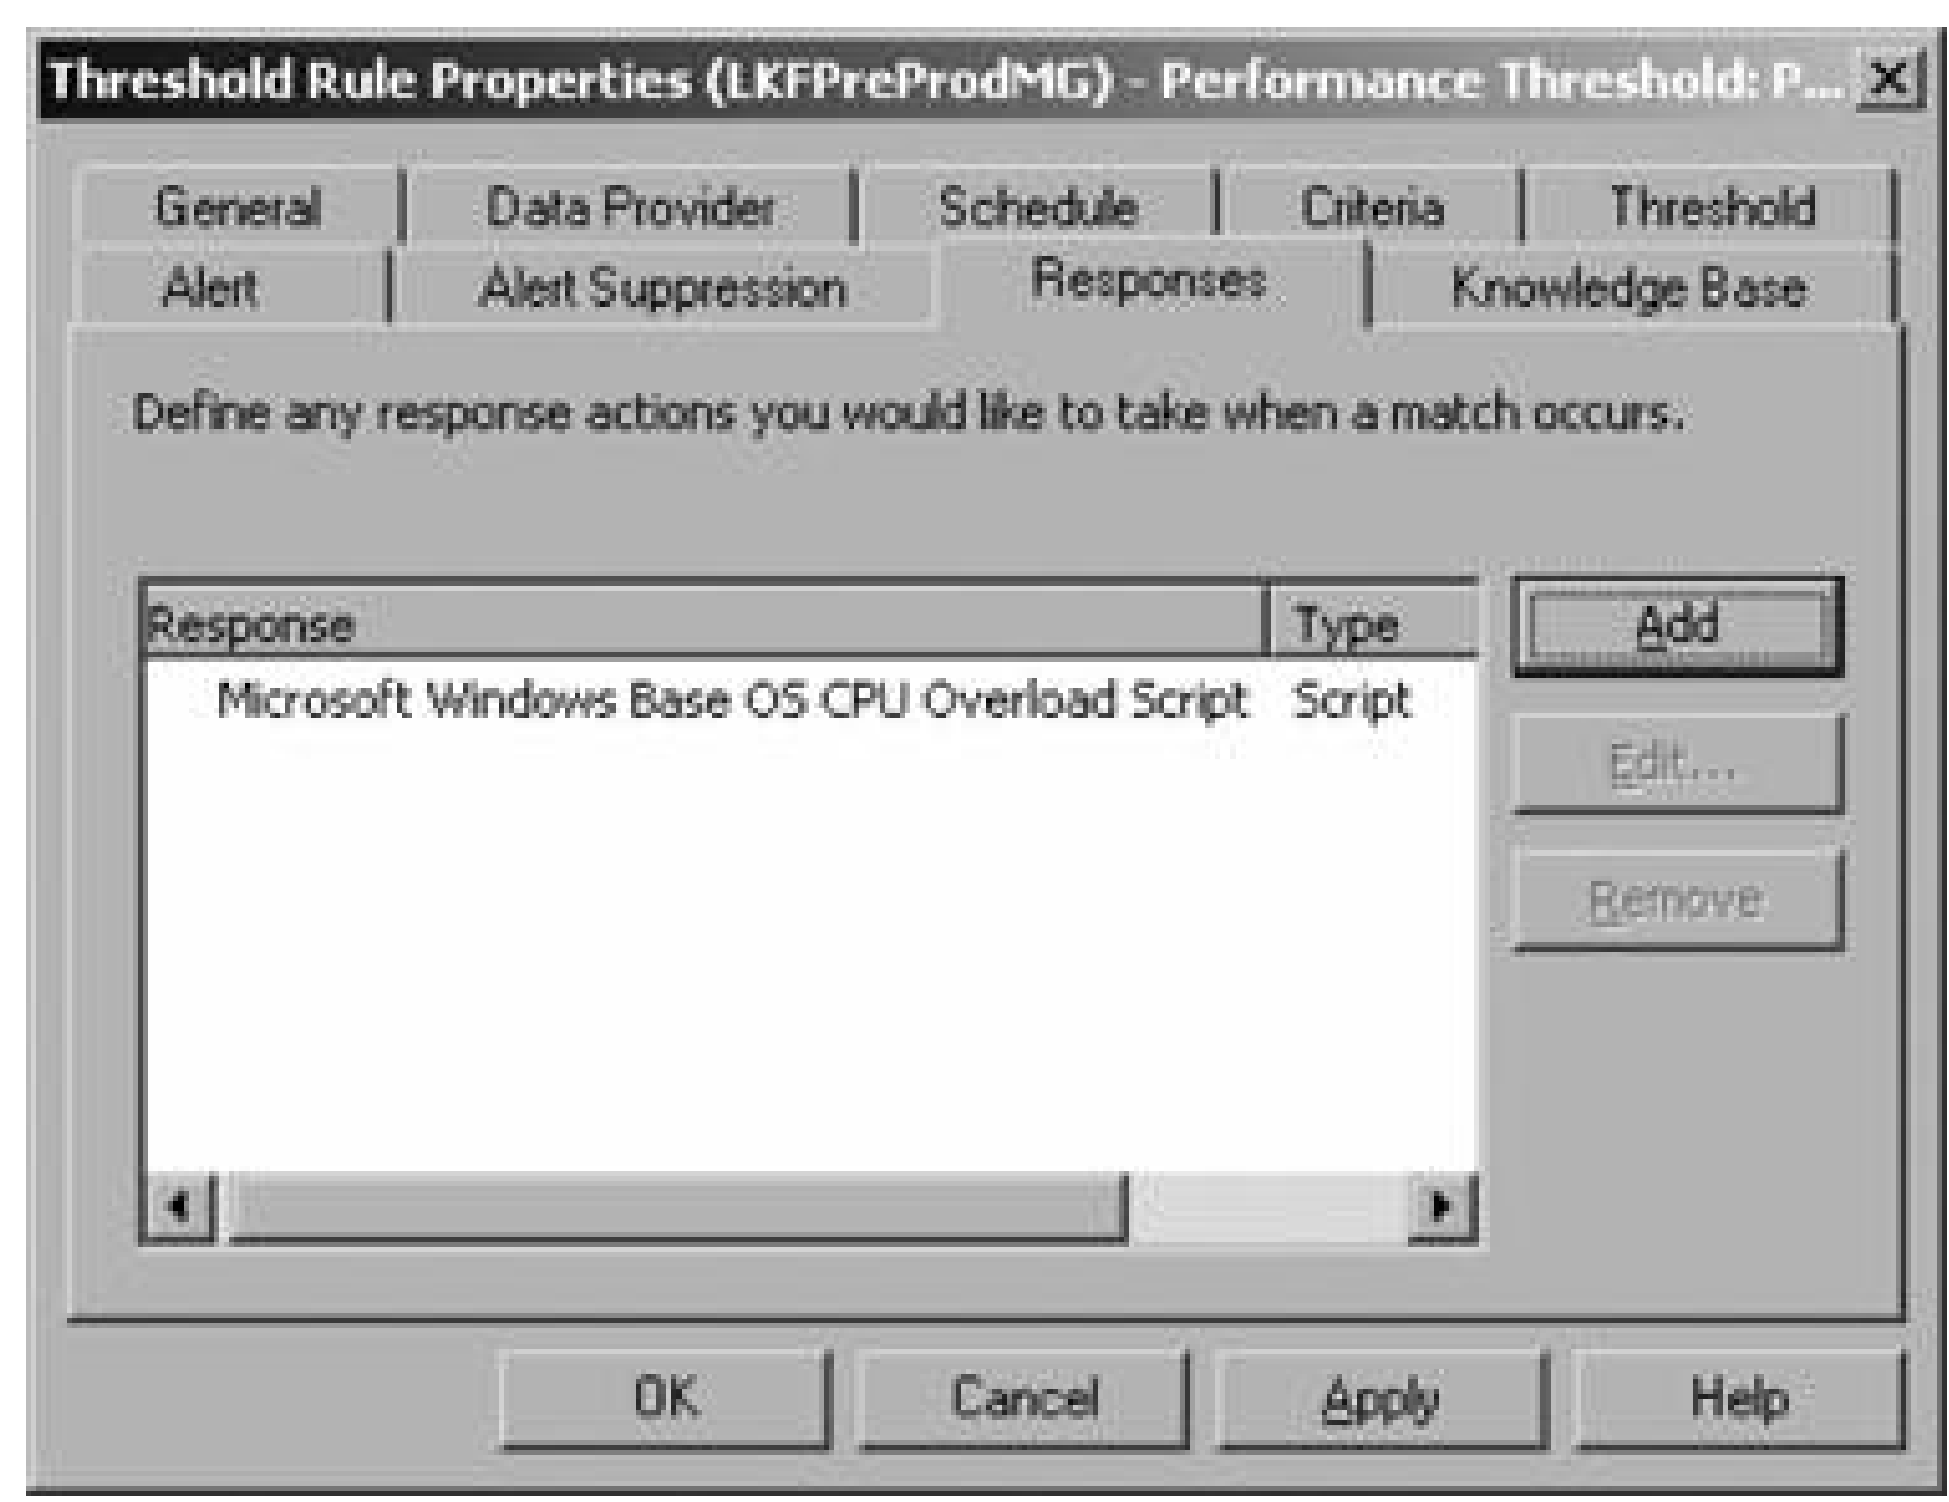
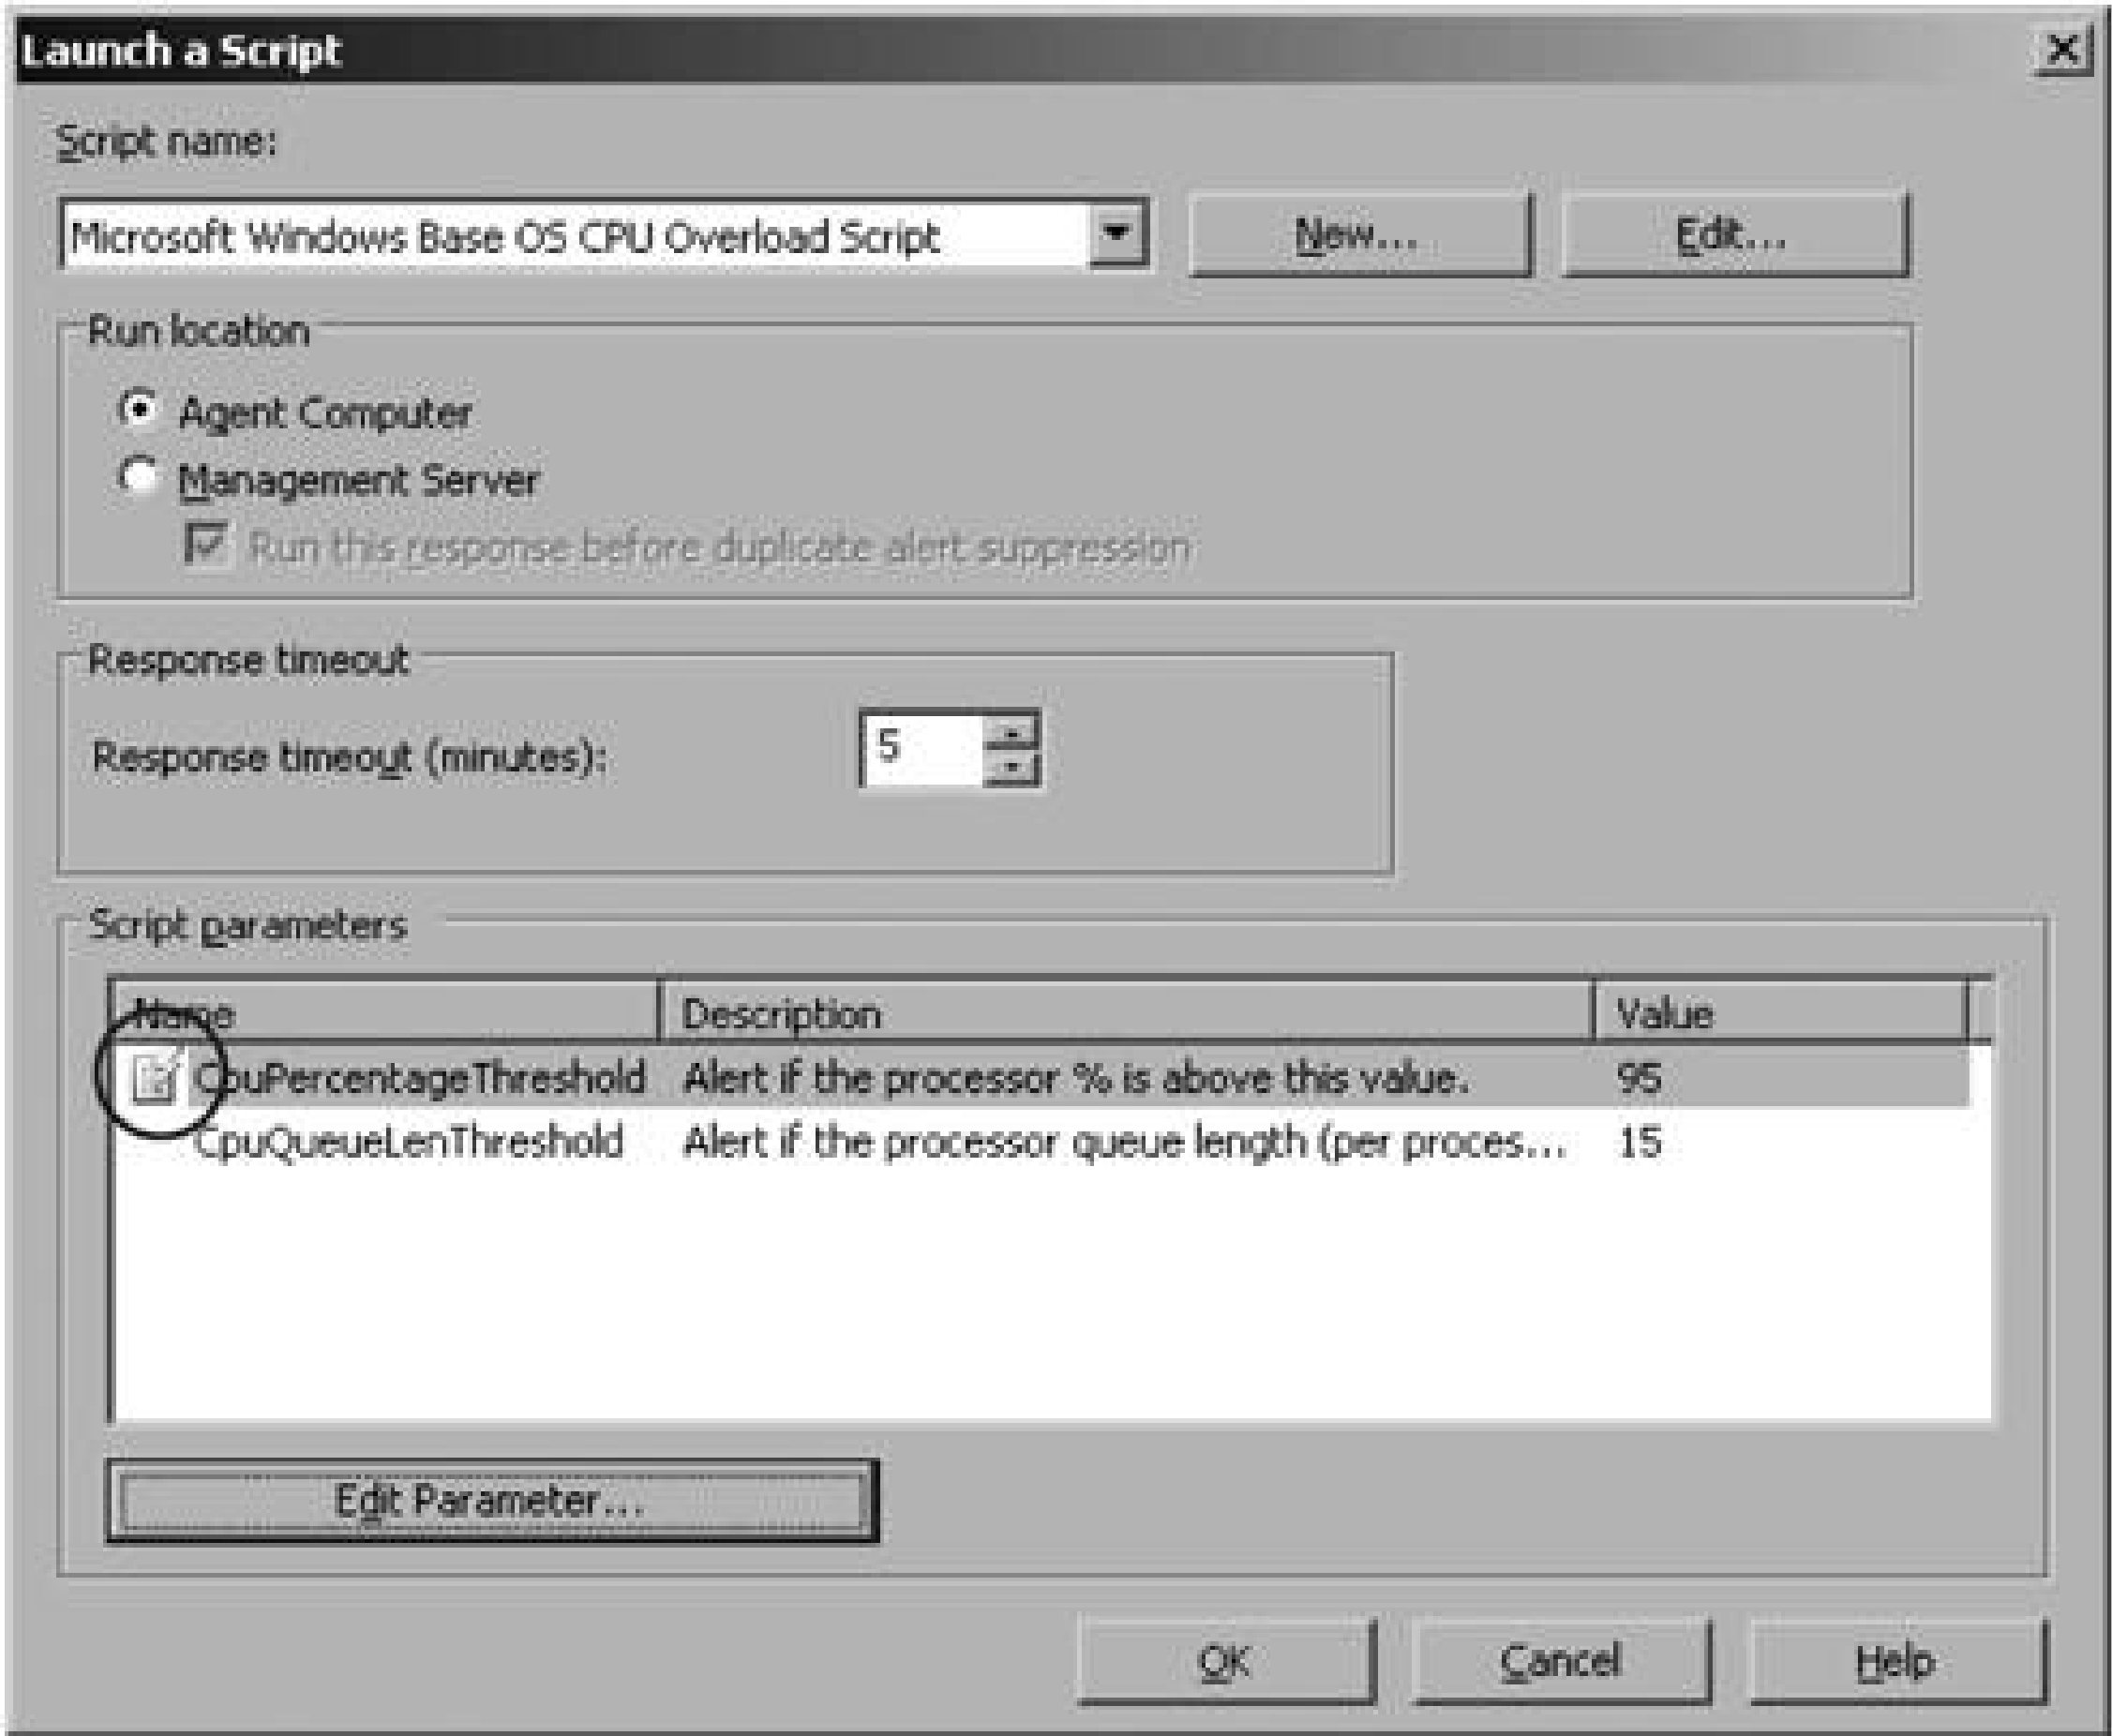


Figure 4-32. This rule generates an alert if the CPU utilization exceeds 95 pe

So now when the script runs, if the *homesrv02%*CpuUtilization counter exceeds 75 percent, an alert will  
For all other machines, the alert will fire off when the CpuUtilization counter exceeds 95 percent.



Figure 4-33. Note the icon that indicates an override has been configured; the same icon for all overrides



4.3.4.4. State alert severity override

Some rules contribute to the calculation of the health state of the application via the alerts that they generate. This is covered later in the "Health state roll-up" section of this chapter. The health state (as represented by green check marks, yellow warning triangles, and red X symbols) for an application or monitored computer is composed of the health of all its component parts. The best place to see these is in the State view of the Operator console.

The state alert severity override allows you to override the default values where a state alert is generated. For example, in the Microsoft Operations Manager 2005 Server rule group there is a performance threshold rule called "Performance Threshold: MOM Server Channel errors" that generates an alert. This rule is applied to MOM 2005 management servers. It has the "Enable state alert properties" checkbox selected (point 1 in Figure 4-34).

By clicking the Edit button (point 2 in Figure 4-34) you can edit the alert severity properties. For this alert, the severity properties (Figure 4-35) show that an alert of severity Error will be generated if the number of errors is greater than 10.

Figure 4-34. State alert properties enabled on an event rule

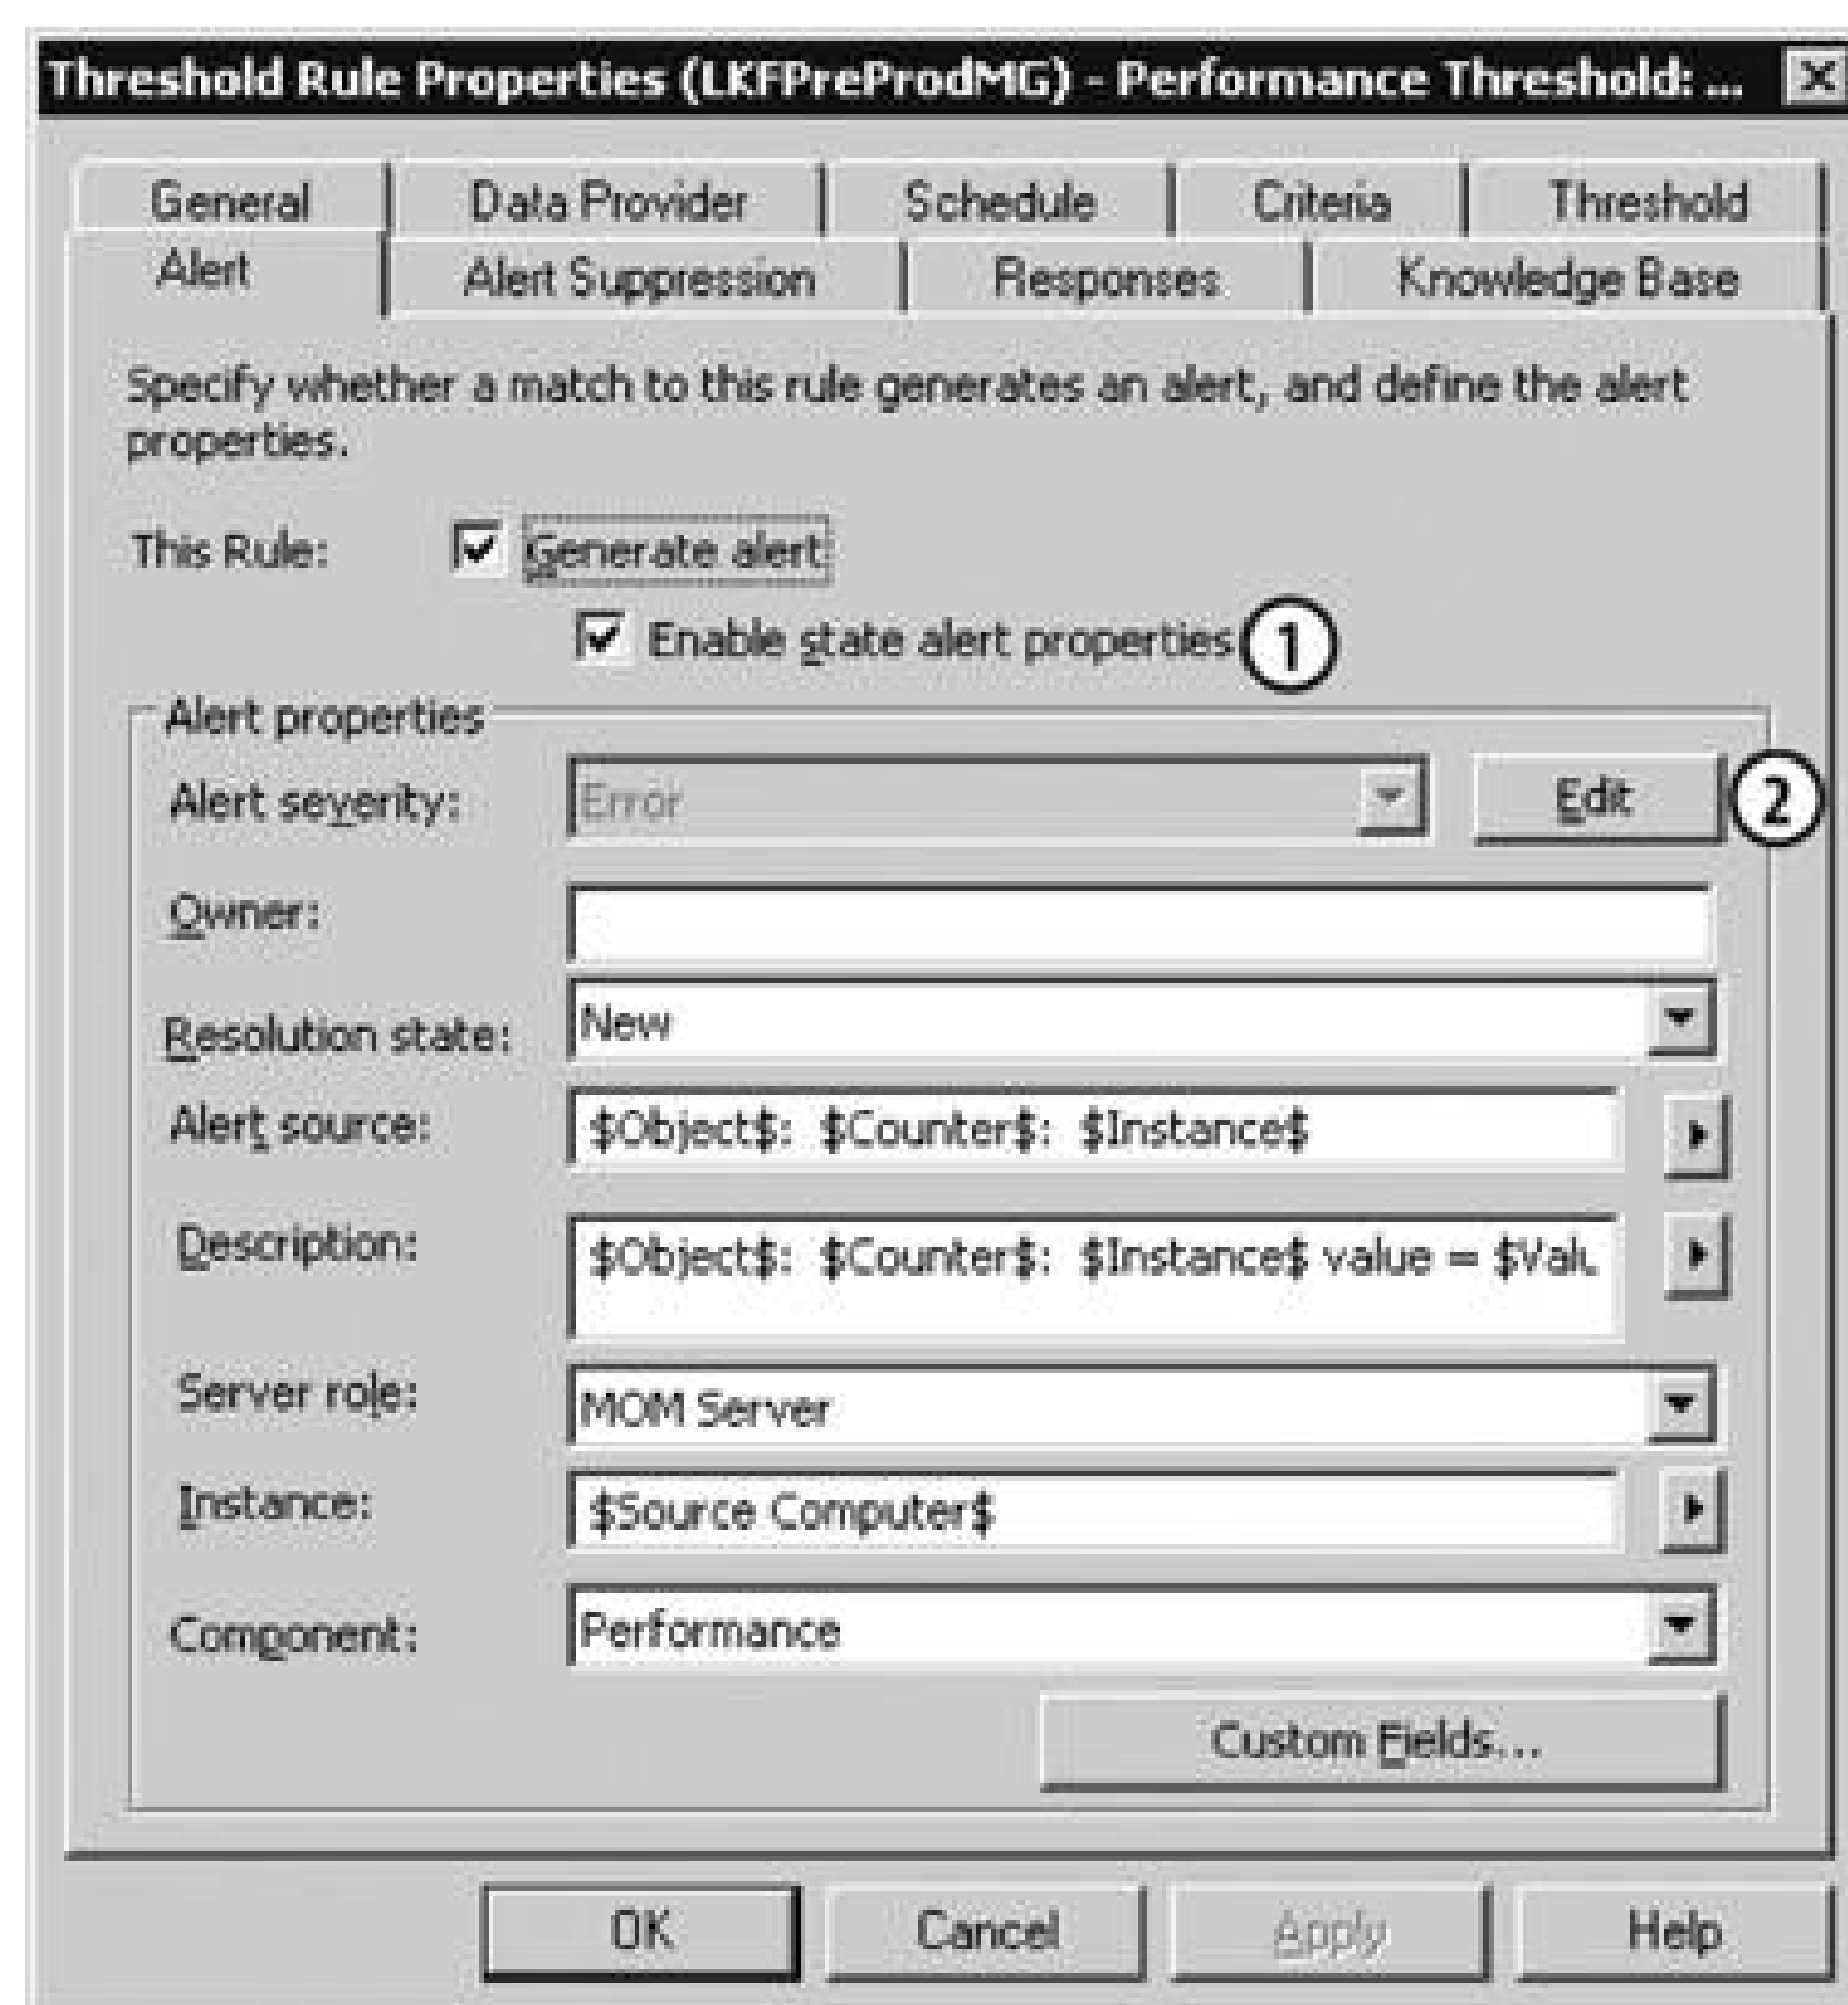


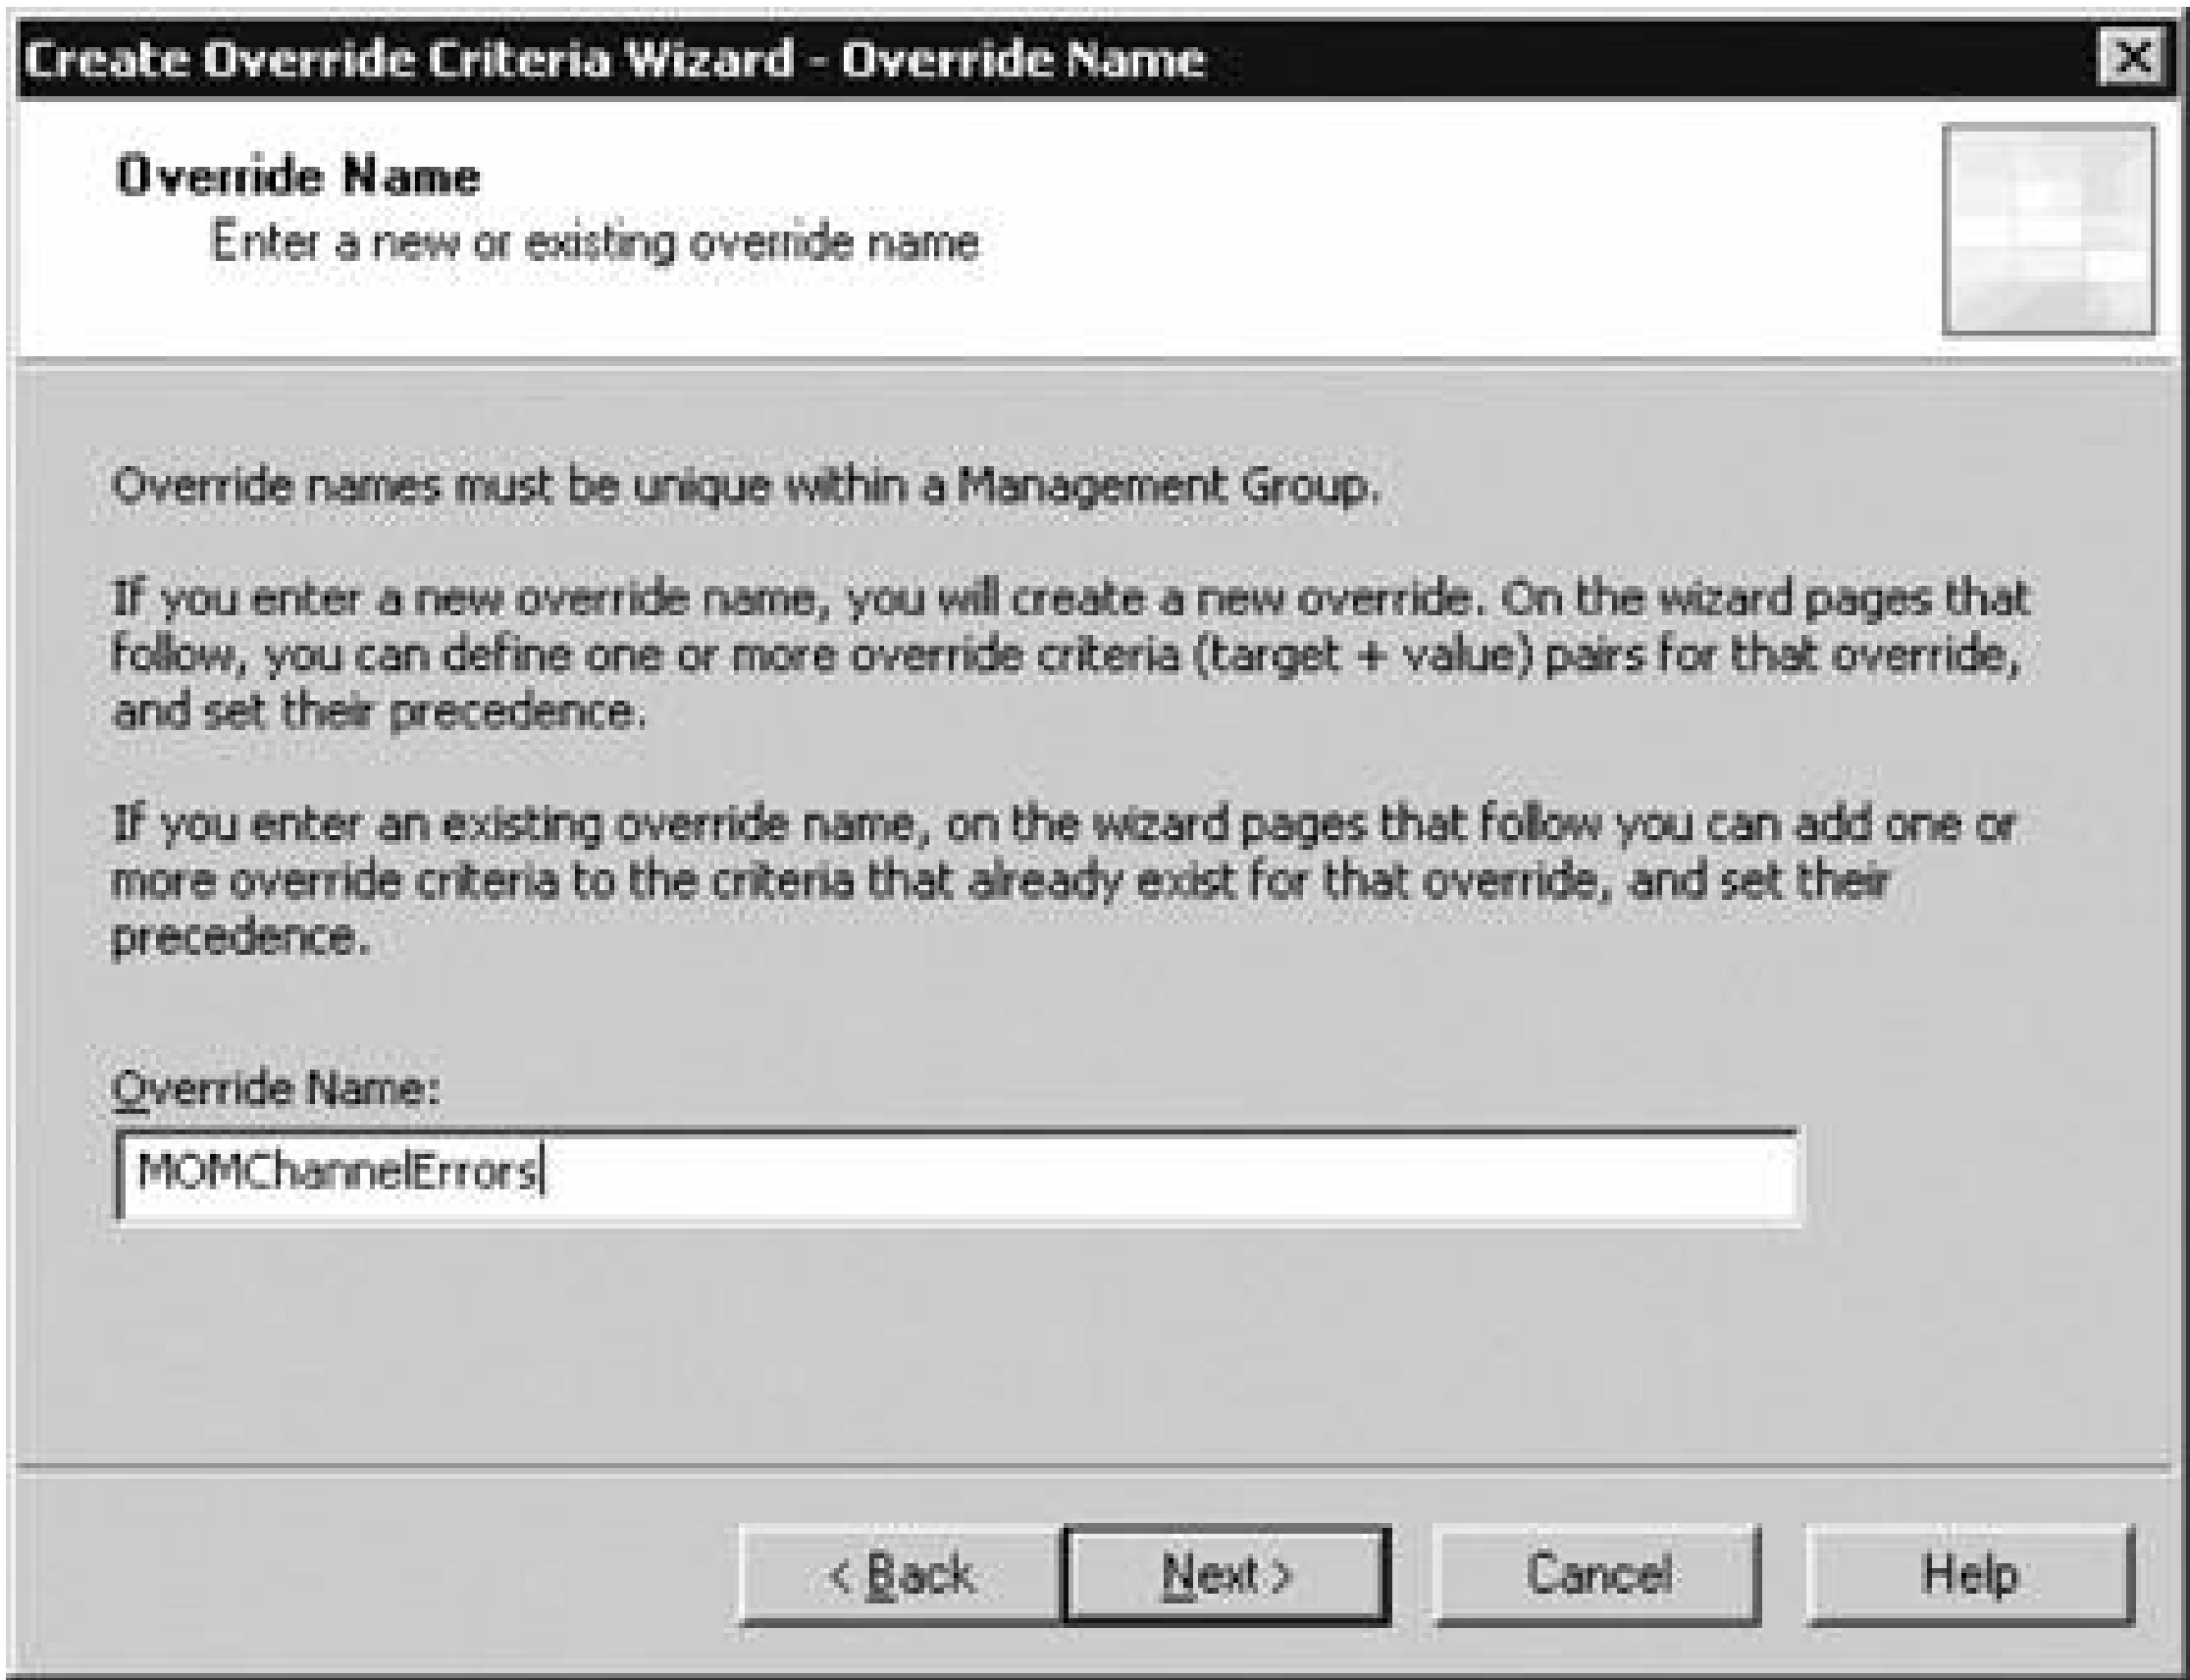
Figure 4-35. The default criteria and values for generating a state alert

To clarify, the value for (Value) is being collected from the performance counter and it points to the number of open communications channels on a management server. If this exceeds 10, an alert of severity Error will be generated, for example if there are multiple management servers in a management group and one needs to generate an error when the number of open channels exceeds 10, the rest of the servers should keep the default of 10. To do this, you need to do two things:

1. Create an override object in the Override Criteria folder in the Administrator console. It is named MOMChannelErrors and assigned a target + value pair of *homemomserver3* and *5*.
2. Rewrite the formula in the Alert Severity Calculation for State Rule box to include the MOMChannel override (*homemomserver3* and *5*) and the default value of 10.

In the Administrator console, right-click the Override Criteria folder and launch the Create Override Criteria wizard. On the Override Name page enter the name MOMChannelErrors (see Figure 4-36).

Figure 4-36. Creating override criteria



Click Next to bring up the Edit Override Criteria page, then click Add and enter the target + value pair (see Figure 4-37 ).

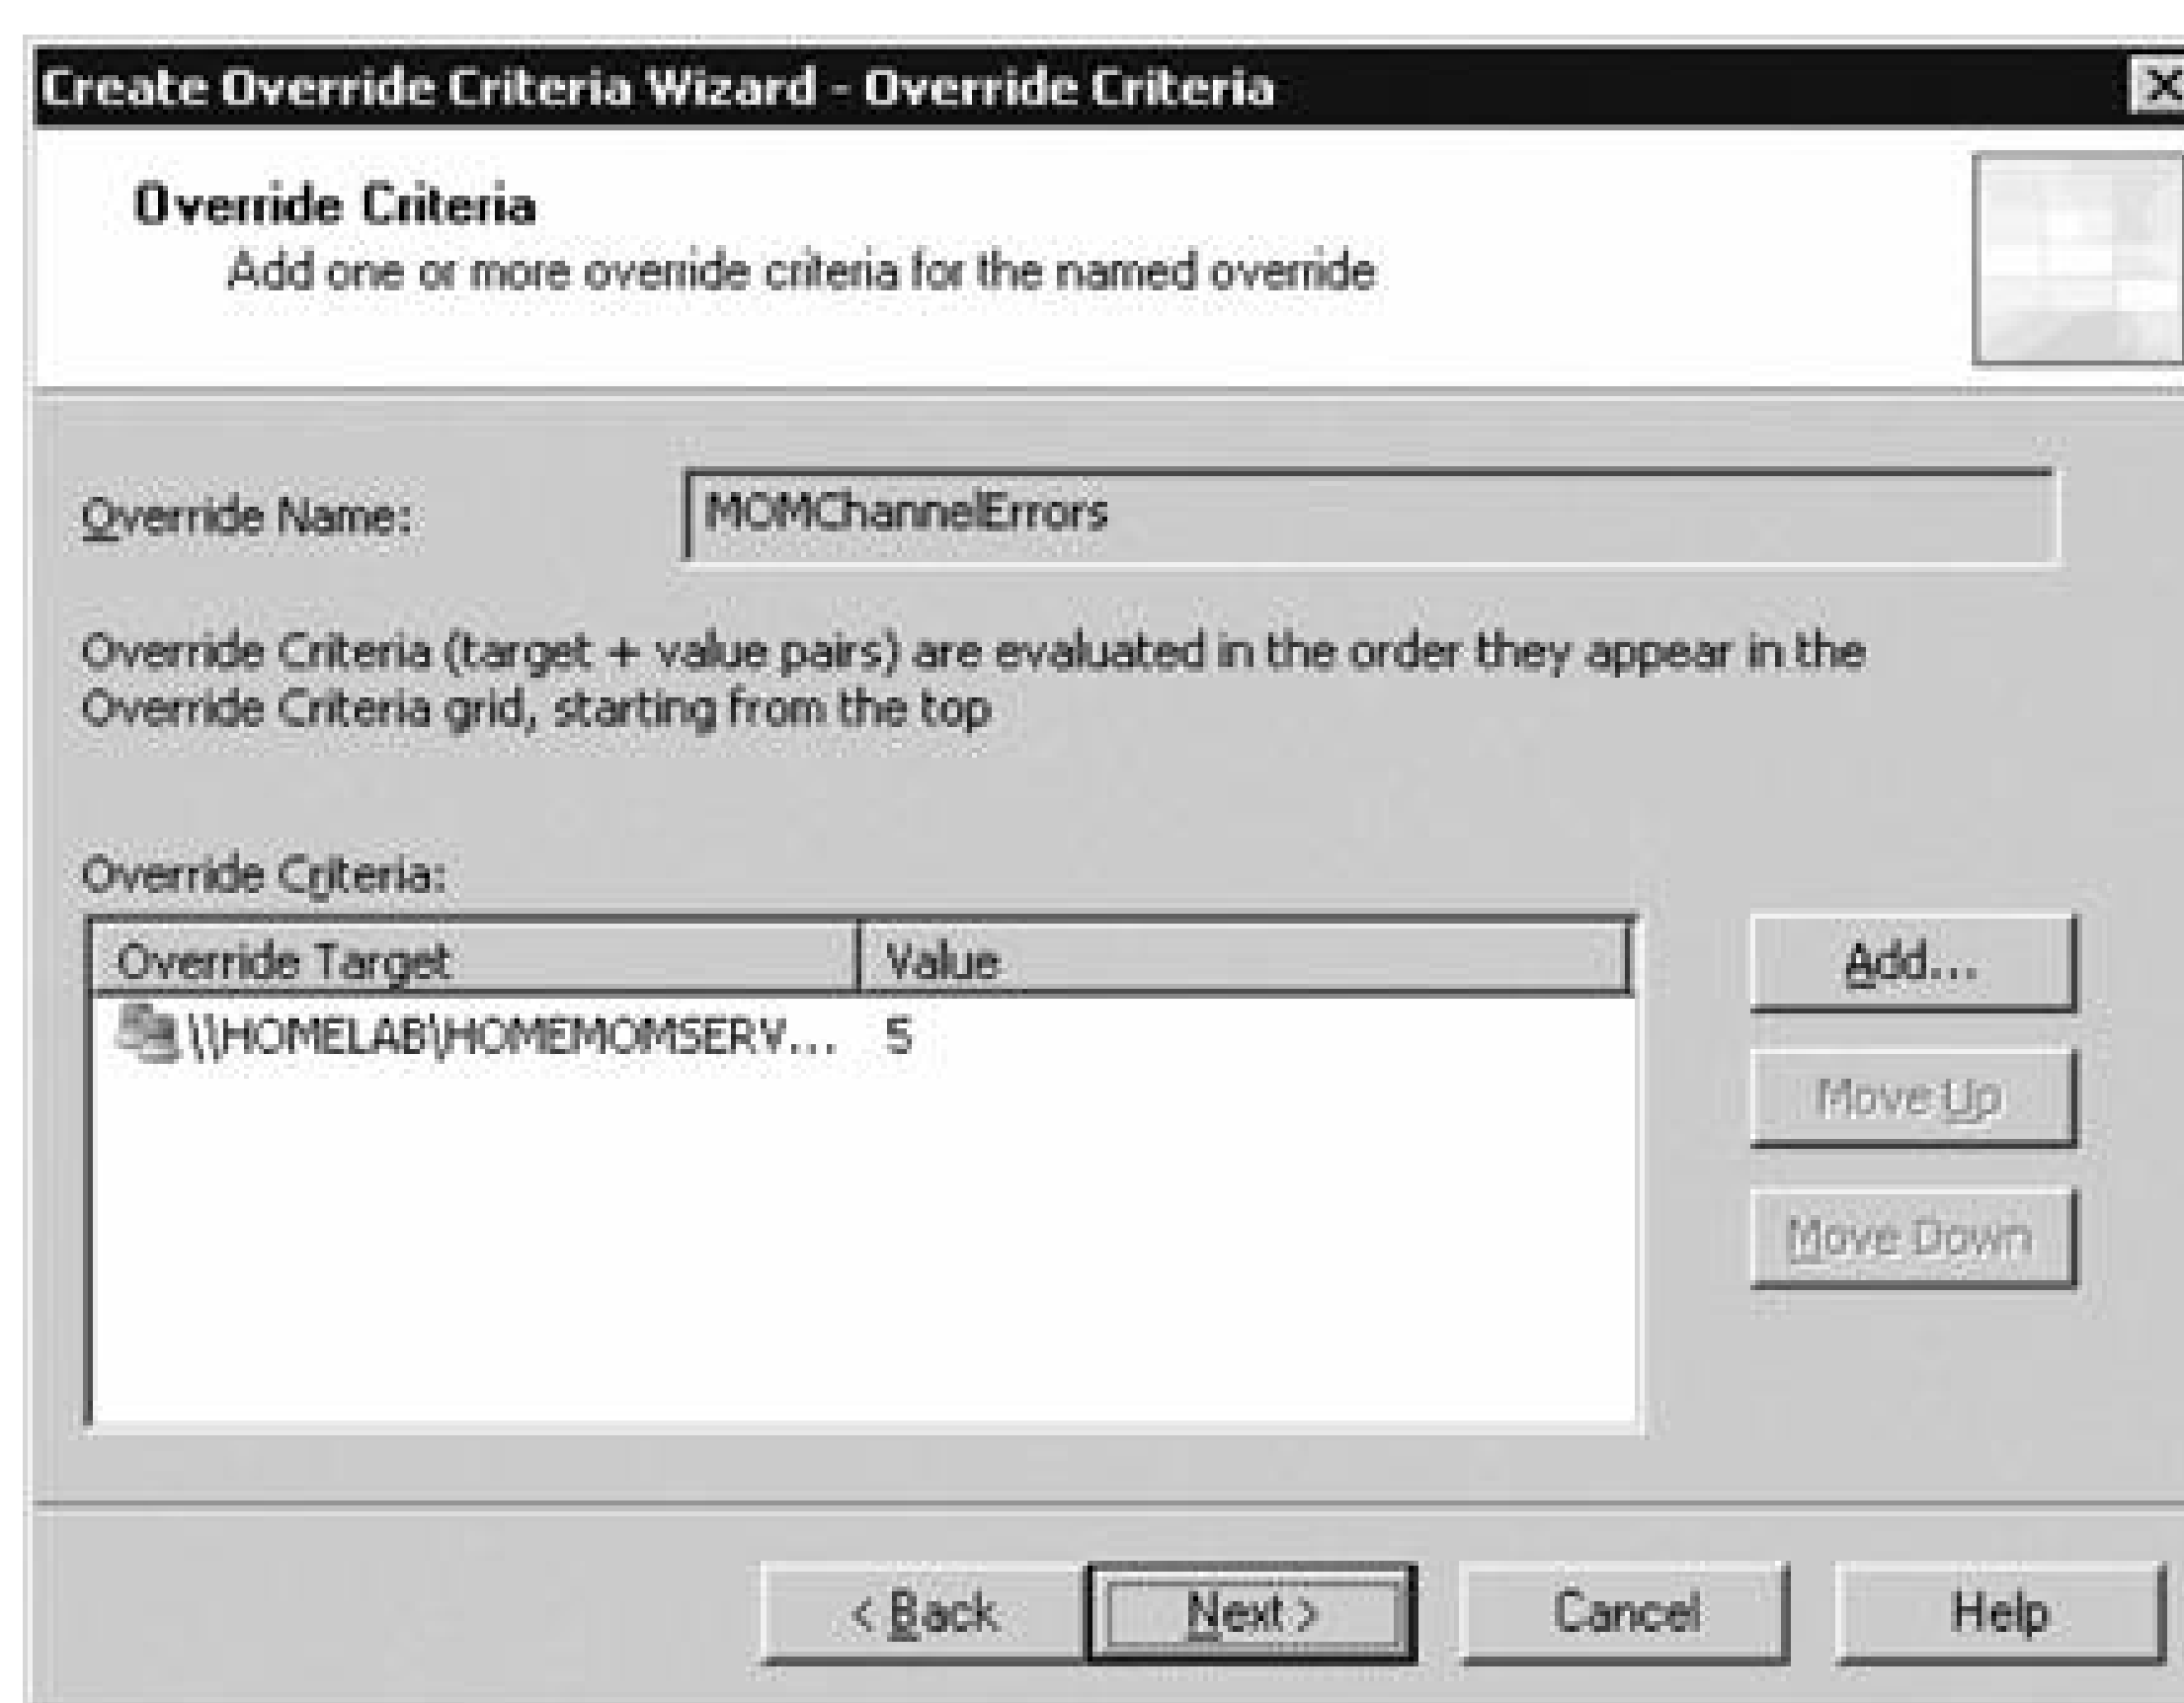
After entering the target + value pair, the Override Criteria page should look like Figure 4-38.

Figure 4-37. Creating the target + value pair for the override





Figure 4-38. The target and value settings for the override criteria page

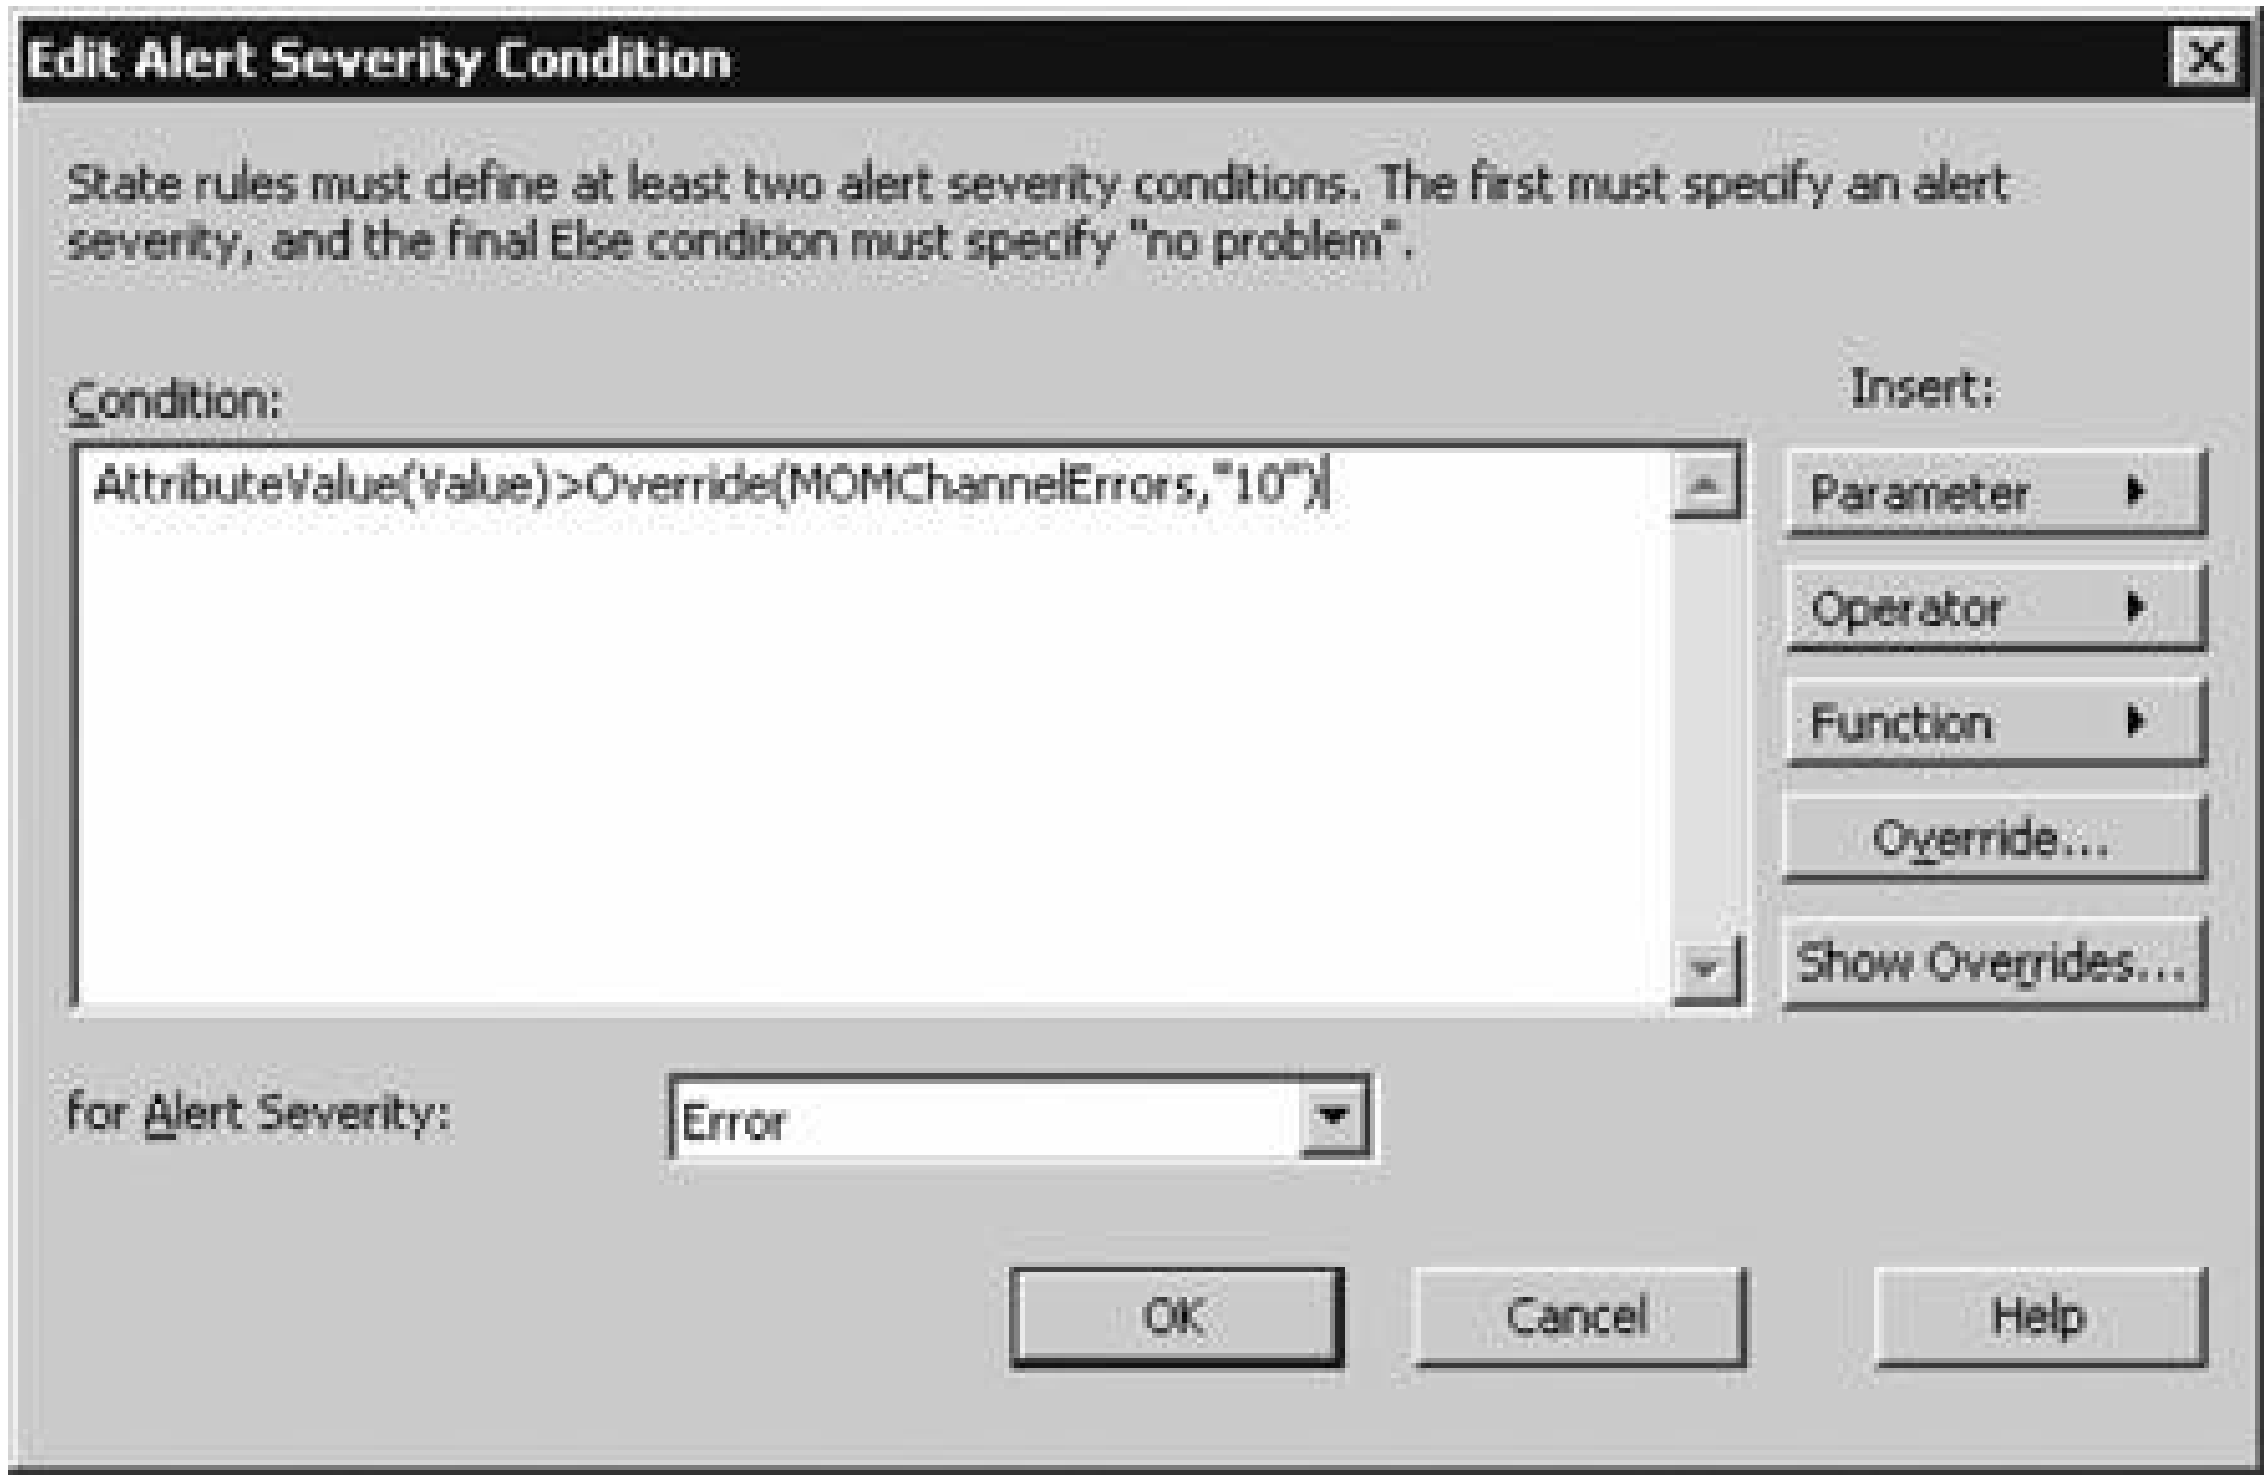


Clicking Next brings up the standard Finish page.

Now go back to the Performance Threshold: MOM Server Channel rule to the page shown in Figure 4-35. click the Edit button. Change the formula to read `AttributeValue(Value)>Override(MOMChannelErrors, ""` shown in Figure 4-39 .

Because you have already created the override object name MOMChannelErrors, the formula automatically the object. Click OK three times to close the rule. Now when this rule executes, if the number of open M channels exceeds five on *homemomserver3*, an error alert will be generated. For all other management in the management group, the error alert will not be generated until the number of open channels excee

Figure 4-39. Including the override object in the state alert severity evaluation formula



#### 4.3.4.5. Moving overrides from preproduction to production

The only bad part about using overrides is that while the identifying GUID will export from preproduction import into production, the criteria will not. This is the workaround.

When an override criterion for a rule is created, it is created as an object in the Override Criteria container under the Management Packs node in the Administrator console. Export the list of defined overrides to a delimited text file (right-click the Override Criteria container and select Export List). This captures the target value, override name (GUID), precedence, last modified date, and last modified by values. For example, override object data for the performance threshold rule is:

- Target: \\HOMELAB\HOMESERV02
- Value: 25
- Override name (GUID):  
 { 9DE93ECD-C739-41A7-BB4C-0BEDFD3AAFB8 }\_Threshold
- Last modified date: 02/27/2005 12:04:17 AM
- Last modified by: HOMELAB\Administrator

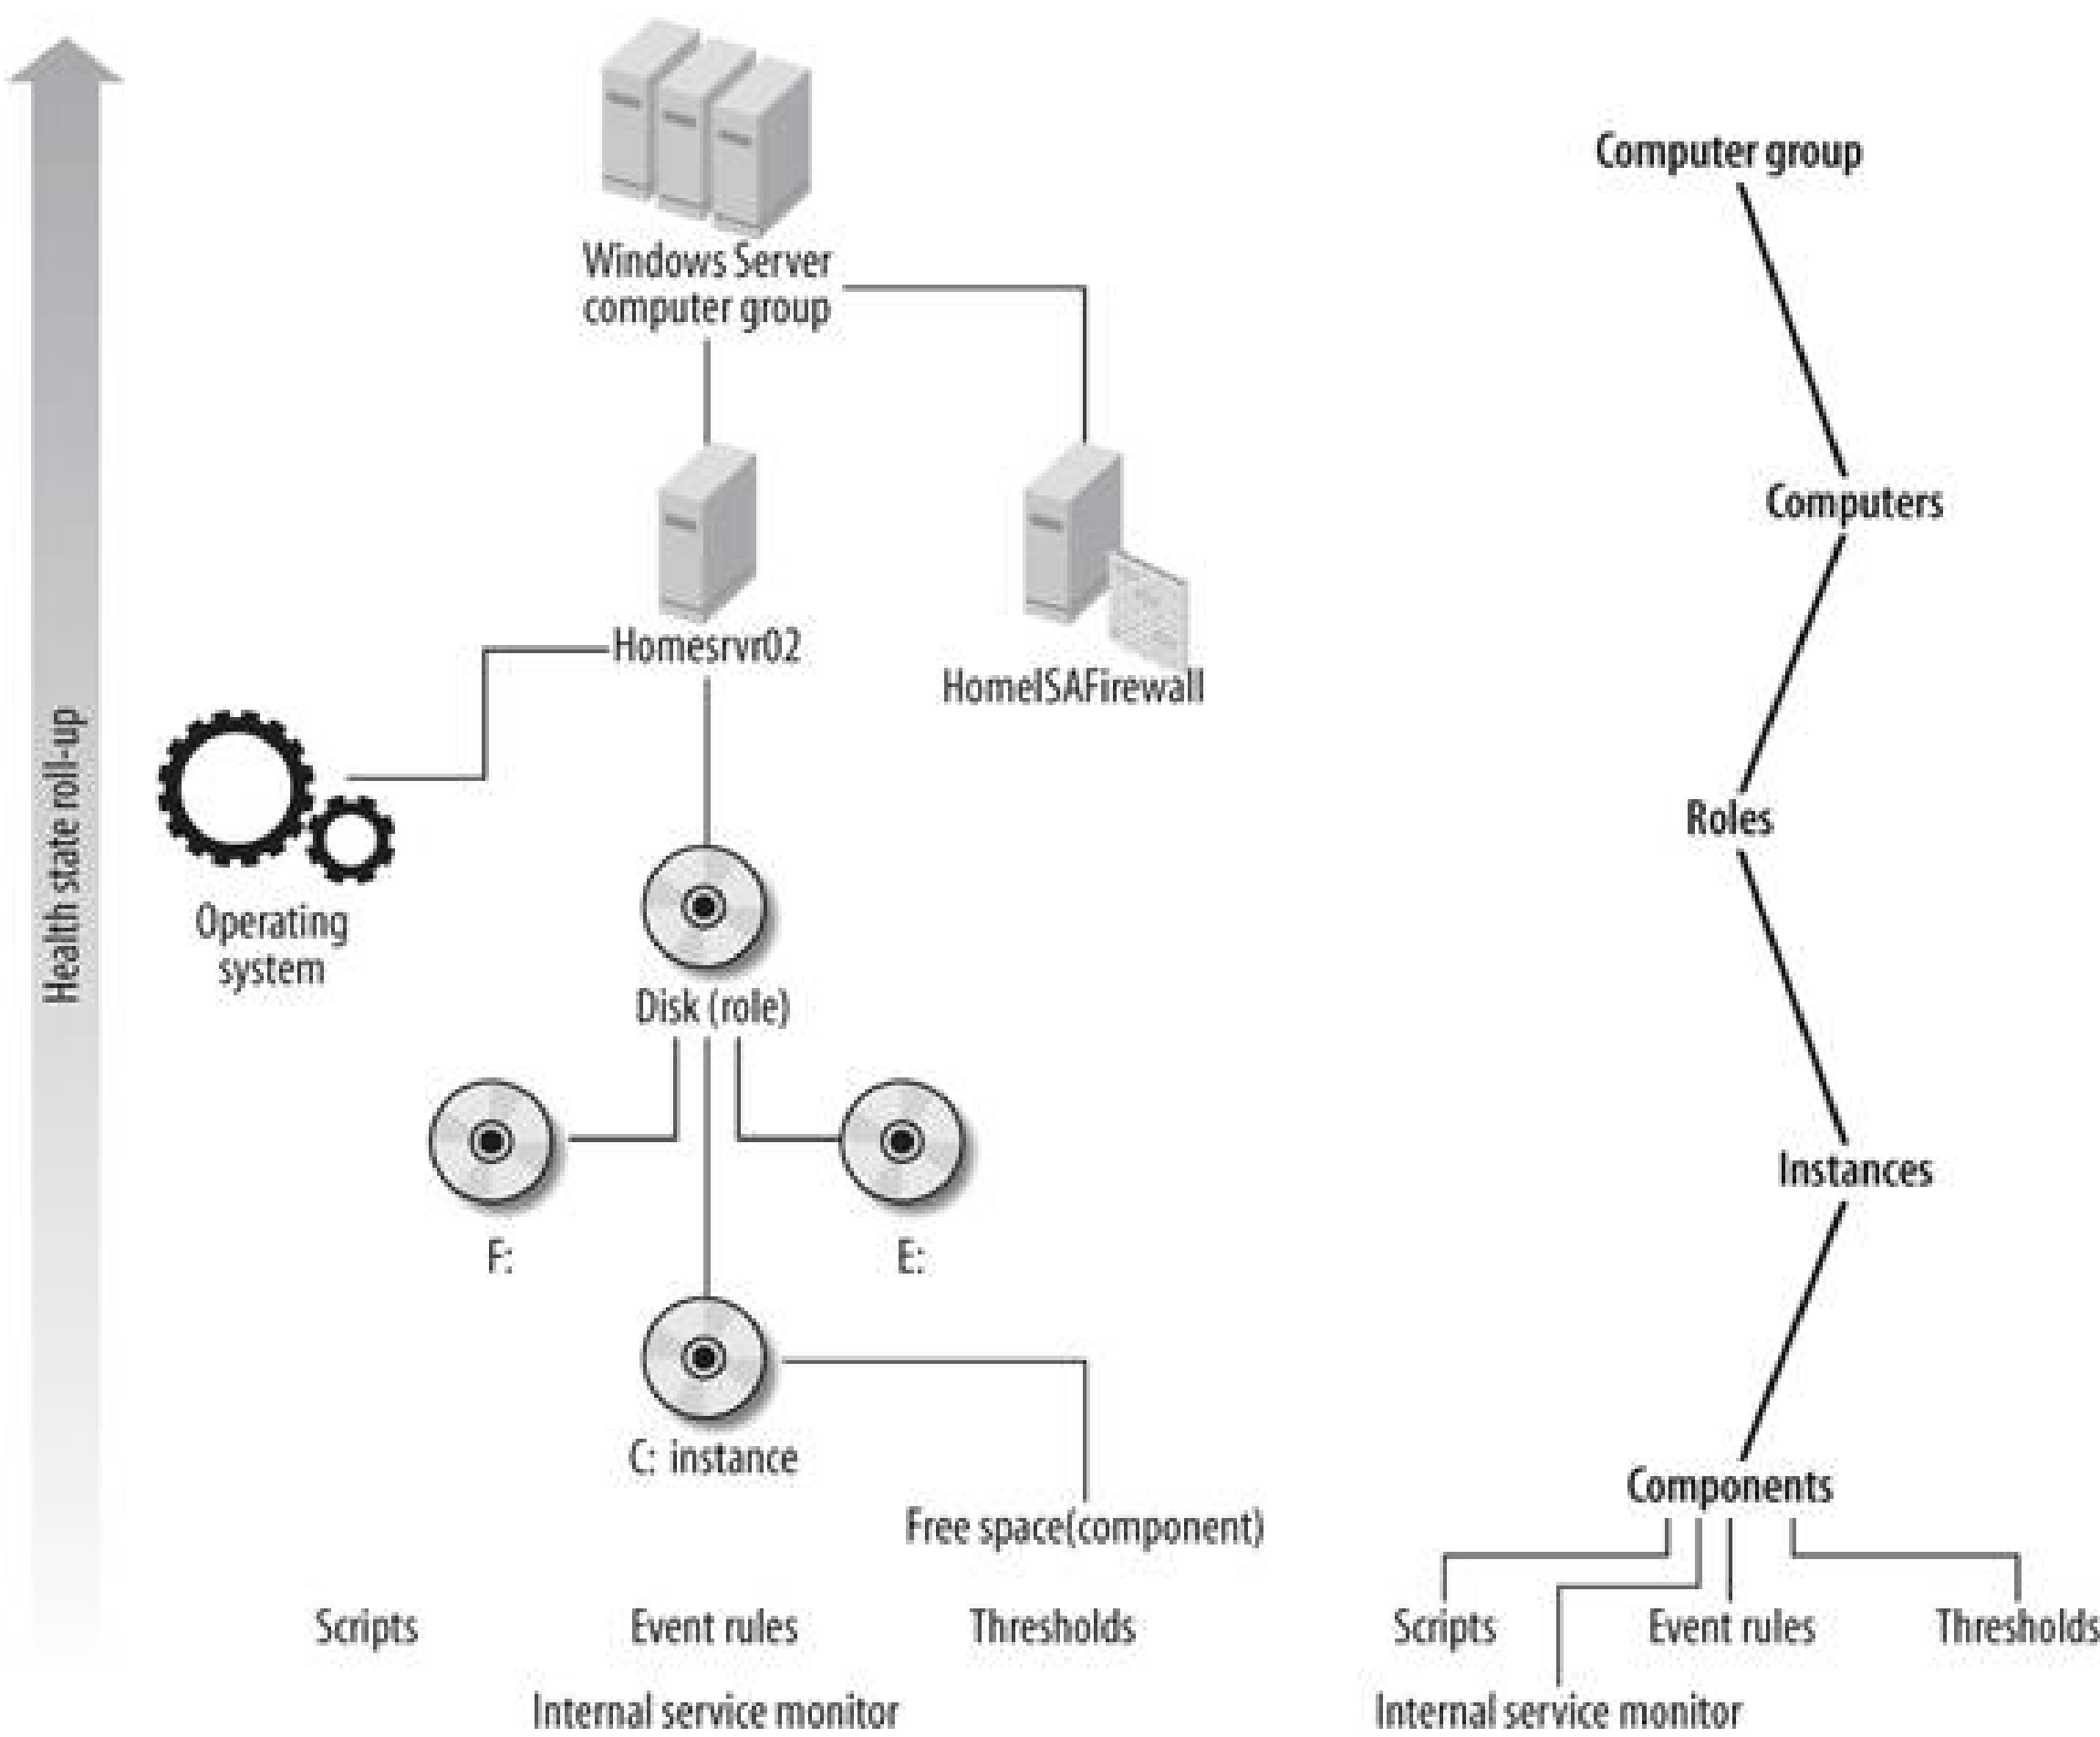
After you import the revised rule group into the production management group, use the Find Rules tool (click Rule Groups and select Find Rules) to perform a search for rules that contain the override name/GUID. Copy and paste the GUID from the text file into the search tool. Once you find the rule, open the appropriate Override Criteria dialog box and enter the criteria. Apply the changes and don't forget to commit configuration changes.

### 4.3.4.6. Health state roll-up

The last bit of fine-tuning for the management packs is to ensure the health state of a computer group is calculated and represented the way you want it to be.

State roll-up itself follows a hierarchy, as shown in Figure 4-40. Starting at the bottom, the health state variable of an individual component can be set by event rules, threshold rules, internal service monitoring scripts. In this case, the disk free space state is the component. The health of an instance (in this case the drive) then depends on the health of its components. The health of the role disk on this computer is dependent on the health of the instances of disk on the computer (C:, D:, and E:), which then feeds the health of the computer and rolls up to the computer group.

Figure 4-40. Health state of each object depends on the health of its individual components



This is a simplified view of health roll-up, as it shows only a single hierarchy and no inter-relationships. Remember that all computers are members of multiple computer groups and a top-level computer group has nested child computer groups. So, the health of an individual instance (the C: drive in this case) can influence the health state of more than one computer group.

For example, consider a computer group that consists of all Exchange servers. This computer group contains other computer groups like mailbox servers, bridgehead servers, public folder servers, and both 2000 and 2003 frontend servers. Now, say that you have a single physical machine that is both a mailbox and public folder server. The health of the C: drive on that one machine will impact the health of both the mailbox and public folder server.

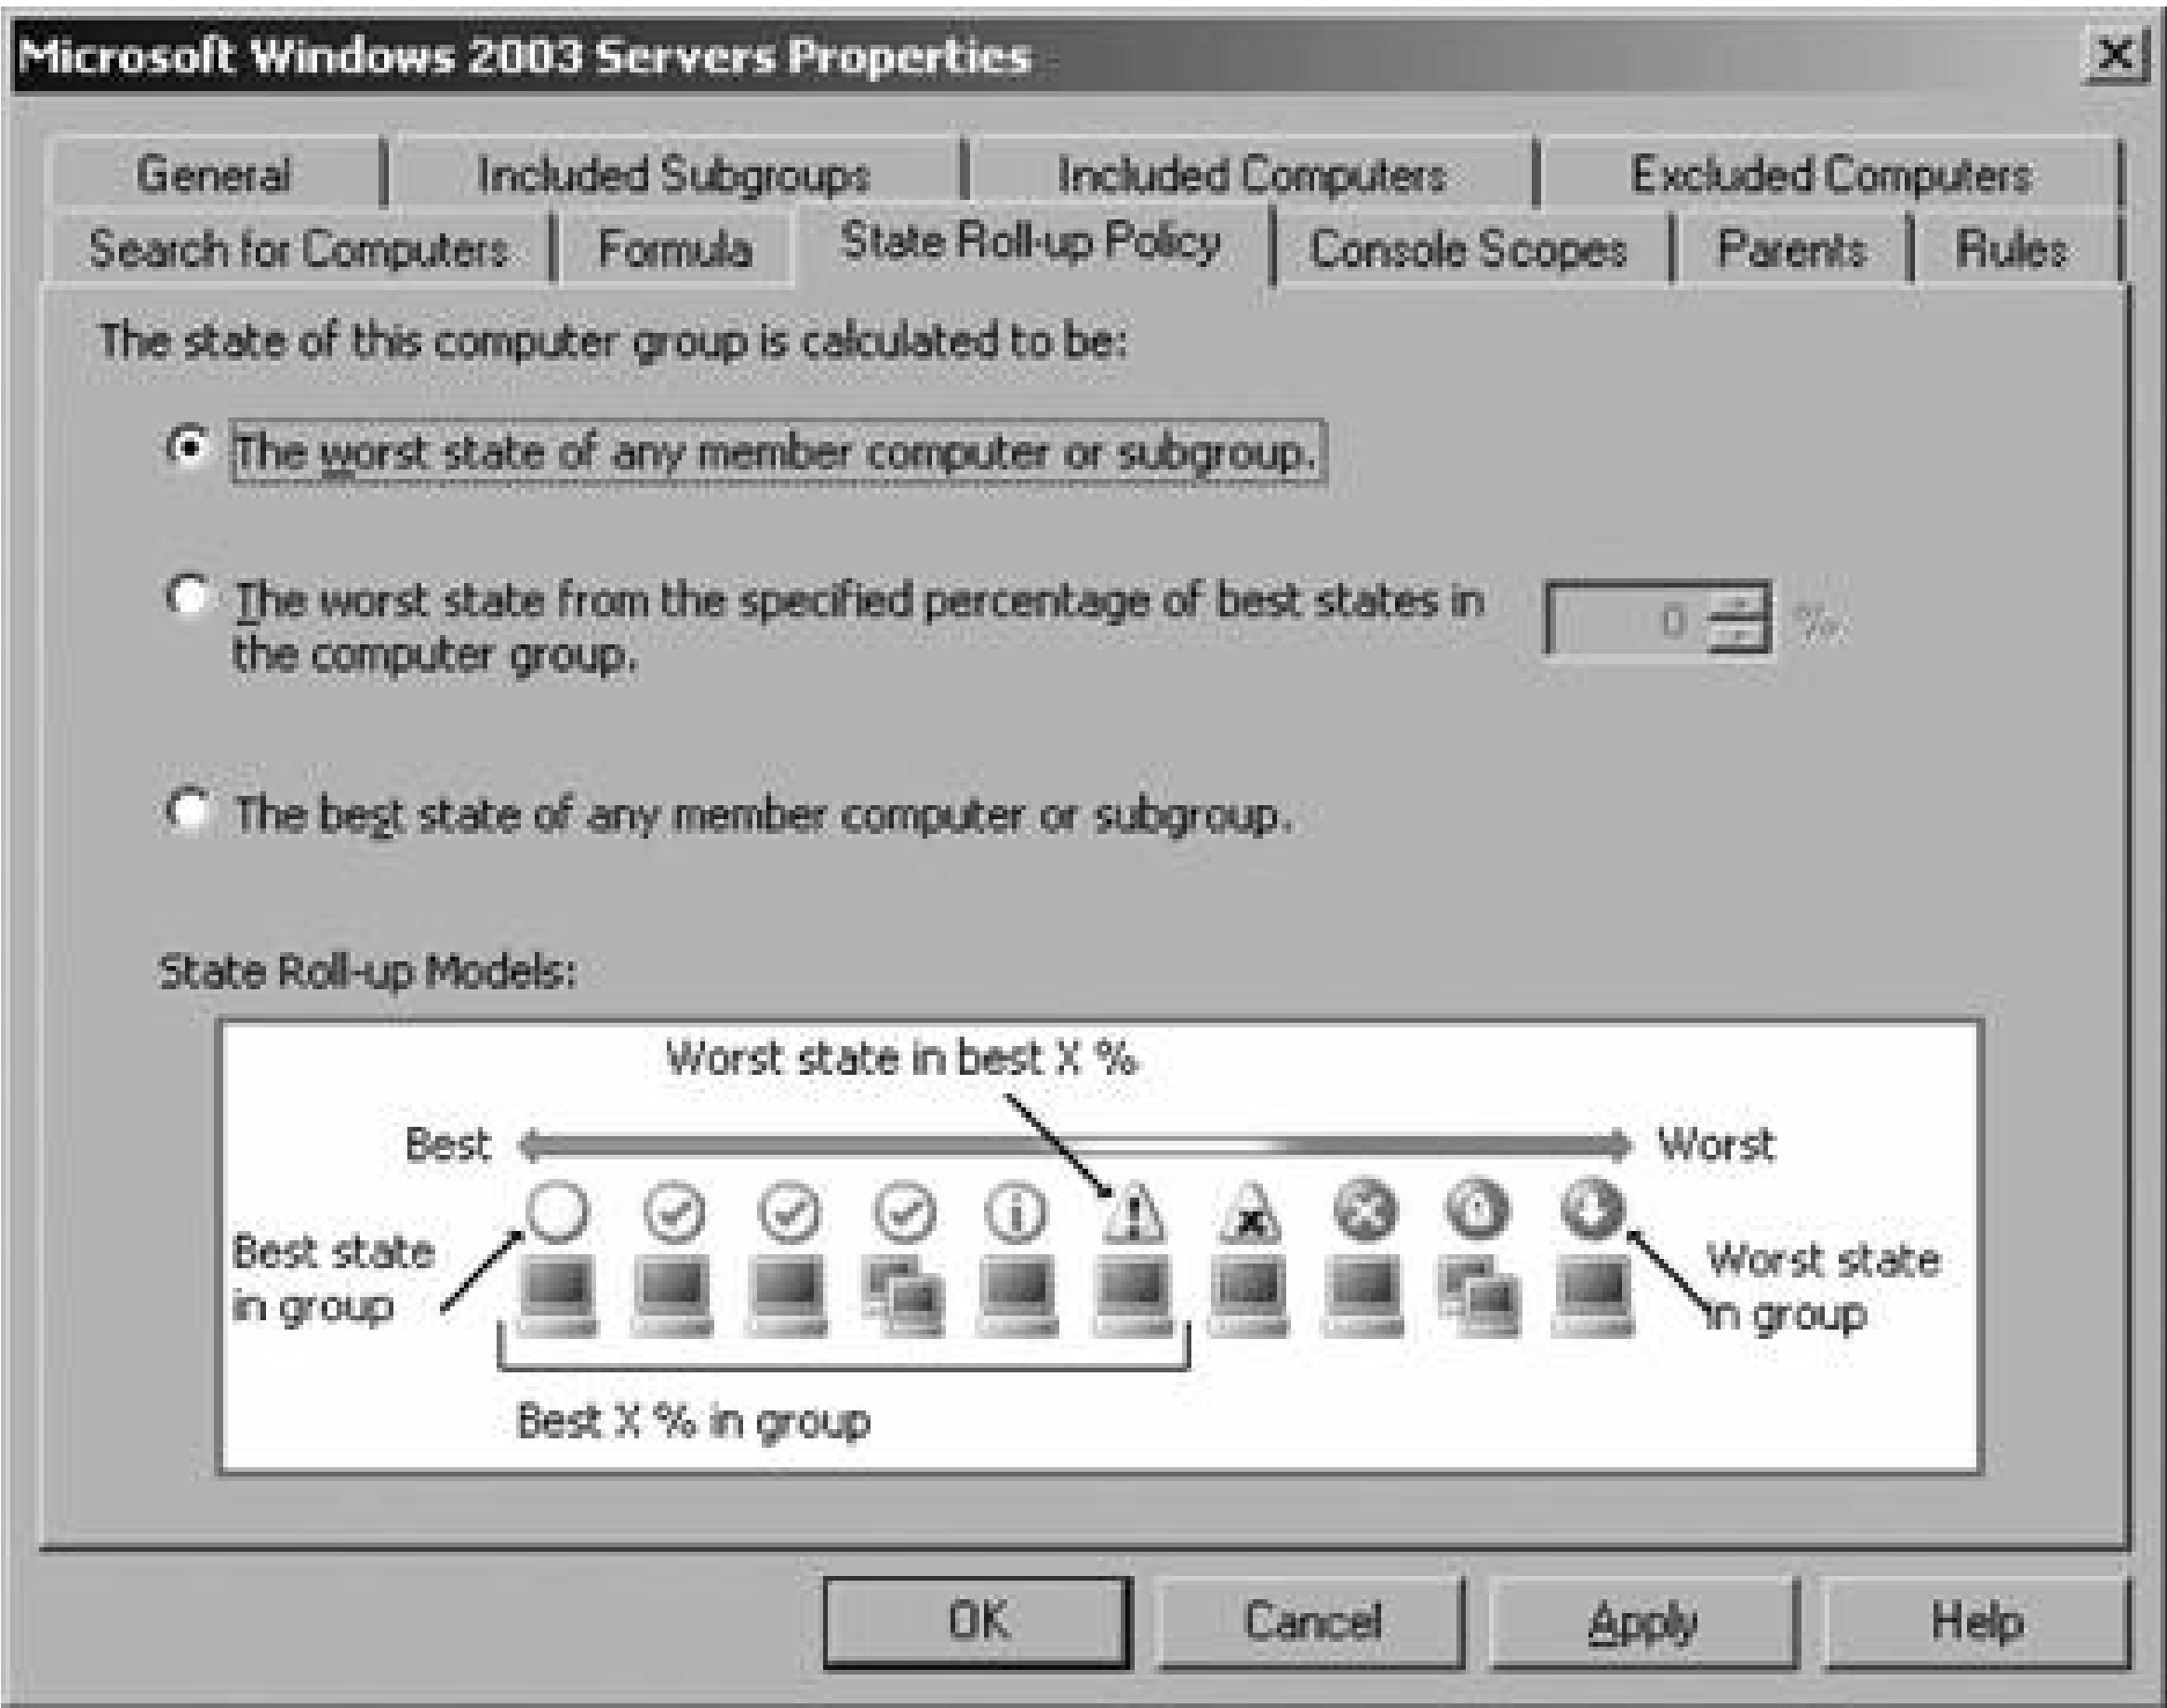


folder computer groups.

Trying to keep the inter-relationships straight between all the objects can be confusing. However, when overlaid in a single diagram, an unhealthy object will be apparent.

The health state roll-up policy is set at each computer group and regulates how the health state of an application or group of computers is represented in the Operator console (Figure 4-41). Although this dc involve tweaking an individual rule or management pack, it is included here because computer groups a created when management packs are imported, and how health state roll-up is configured controls how environment is represented.

Figure 4-41. Health state roll-up policy



There are three ways to configure this, but the default configuration is the most pessimistic way of view health of the computer or application:

*The worst state of any computer member or subgroup*

The top-level health state of a computer group reflects the worst of member states. Using this set make things look worse than they really are.

*The worst state from the specified percentage of best states in the computer group*

The top-level health state of a computer group reflects the best of the member states. This is the

optimistic, "glass half-full" scenario. The top-level health state can make things look better than they really are.

#### *The best state of any member computer or subgroup*

The top-level health state is the worst state of the top percentage of best states in the group. This basically tells MOM to rank in order the computers in a computer group from healthiest to unhealthiest. Then MOM looks at a certain percentage of the computers in a computer group, starting at the top of the scale. MOM then takes the health state of the unhealthiest one and uses that as the overall health state of the computer group. This configuration tries to most accurately represent the condition of the computer group as a whole.

The worst state in the group scenario tells you that if everything is successful at the top level, then all other components are successful across the board and there are no masked problems that could potentially be overlooked.

### 4.3.4.7. Tools

To further assist you in tuning your MOM environment, Microsoft has produced an Alert tuning solution accelerator . All the MOM 2005 solution accelerators are available at <http://www.microsoft.com/mom/evaluation/solutions/default.mspx> . This tool consists of report definitions that can be imported into the MOM 2005 Reporting Service and a description of how to use them. The reports include AlertCountByDates, AlertCountByDevice, and AlertCountByProcessingRules.

The idea behind the alert tuning accelerator is that by displaying alert counts by different parameters you can determine which rules are generating the most alerts and if those alerts are actionable. If an alert is not actionable, that is to say that it doesn't contain the information that is necessary to resolve the issue, then it is a noise alert and should be tuned to be useful, or disabled altogether.

The reports that this tool generates can be useful in any size environment to develop a profile of what is going on in your MOM implementation.

Referring back to Figure 4-1 , the Windows Server Base Operating Systems management pack has evolved to the MP.1.2 state and is tailored for your environment. But now there is a new issue to deal with: the version of the management pack in the production management group is different than the one in preproduction. In addition, a significant amount of work has gone into customizing the management pack and that needs to be protected.

### 4.3.5. Protecting Management Packs

Keeping production and preproduction management pack versions synchronized and backing up management packs can be accomplished with the same straightforward mechanism. You need to synchronize the preproduction environment with production so that when any rule group, view, or task testing is performed, the rule groups in the test environment are identical to those in the production environment (point 7 in Figure 4-1).

The Management Module Utility is the command-line version of the Import/Export wizard tool and includes all of the same functionality for importing and exporting management packs. It is in the MOM 2005 installation directory ( *C:\ProgramFiles\Microsoft Operations Manager 2005* by default) on any MOM management server computer on which the consoles have been installed. You can perform the export from production and reimport it into preproduction by using the Import/Export wizard, but that requires daily hands-on interaction. By



the Management Module Utility, you can automate this job. The Management Module Utility does not work with MOM 2005 report definition files for that, the *rptutil.exe* tool found in the same directory is used. Obviously since it runs from a command line it does not provide the ability to browse the directories where the *.akm* files are.

This tool, with minimal batch file or scripting skills, automates the management pack backups and keeps production and preproduction environments in sync.

The sample batch file, *ProdToPreProdSynch.bat*, does six things:

1. It prompts the user for the current date (MM.DD.YYYY) and time (HH.MM) in 24-hour format. This timestamp is embedded in the filename to ensure uniqueness.
2. It moves all the files in the *CurrentMP* folder to the *OldMP* folder. This preserves the file management practice that is in place.
3. It calls *ManagementModuleUtil.exe*, passing the *-o* (export) option, the source server, the name of rule group to export, the destination and name of the export file parameters, and the *-w* option (write new *.akm* file). This exports the rule group from production and writes it to the *CurrentMP* folder.
4. It exports tasks using the same command, except the *-o* option exports a rule group. This is changing the *-ot* option for a task and the *-w* option changes to *-A* so that the exported tasks are appended to the file.
5. It exports the views in the same fashion using the *-Ov* and *-A* options.
6. It imports (*-i*) the entire *.akm* file into the preproduction (*homemomserver3*) environment using the *-R* (replace existing management packs) option:

```
@echo off

rem get current date and time using . as a separator
rem varDATE and varTime are used to include time stamp in the file name for
rem uniqueness

set /p varDate=Please enter the date in MM.DD.YYYY format. Use periods only!
set /p varTime=Please enter the time in 24 hour time format e.g. 23.30 format. Use
periods only!

rem move all files from CurrentMP to OldMP

move \\homesrv02\mptransferfolder\currentmp\*. * \\homesrv02\mptransferfolder\OldMP

rem export the rule group (-O), tasks (-Ot), and views (-Ov). The -W writes a new
.akm file, the -A appends

ManagementModuleUtil.exe -O homemomserver "Microsoft Windows Servers Base Operating
System"
\\homesrv02\mptransferfolder\currentmp\MicrosoftWindowsServerBaseOperatingSystemMP%
Date%-%varTime%.akm -W
```



```
ManagementModuleUtil.exe -Ot homemomserver "Microsoft Windows Base OS\Share Configu
Query"
\\homesrv02\mptransferfolder\currentmp\MicrosoftWindowsServerBaseOperatingSystemMP%
Date%-%varTime%.akm -A
ManagementModuleUtil.exe -Ot homemomserver "Microsoft Windows Base OS\Telnet"
\\homesrv02\mptransferfolder\currentmp\MicrosoftWindowsServerBaseOperatingSystemMP%
Date%-%varTime%.akm -A
ManagementModuleUtil.exe -Ot homemomserver "Microsoft Windows Base OS\junk ping"
\\homesrv02\mptransferfolder\currentmp\MicrosoftWindowsServerBaseOperatingSystemMP%
Date%-%varTime%.akm -A
ManagementModuleUtil.exe -Ot homemomserver "Microsoft Windows Base OS\Networking\Pa
Ping"
\\homesrv02\mptransferfolder\currentmp\MicrosoftWindowsServerBaseOperatingSystemMP%
Date%-%varTime%.akm -A
ManagementModuleUtil.exe -Ov homemomserver "Microsoft Windows Base OS\Alerts"
\\homesrv02\mptransferfolder\currentmp\MicrosoftWindowsServerBaseOperatingSystemMP%
Date%-%varTime%.akm -A
ManagementModuleUtil.exe -Ov homemomserver "Microsoft Windows Base
OS\Performance\Performance Counter Request Errors"
\\homesrv02\mptransferfolder\currentmp\MicrosoftWindowsServerBaseOperatingSystemMP%
Date%-%varTime%.akm -A

rem now import the akm files. -R replaces existing management packs on import

ManagementModuleUtil.exe -I homemomserver3
\\homesrv02\mptransferfolder\currentmp\MicrosoftWindowsServerBaseOperatingSystemMP%
Date%-%varTime%.akm -R
```

Not all the tasks and views are included in this example, but you can see what is possible. With a little n work, this type of script could be scheduled to run with no user interaction.

The Management Module Utility includes a versioning option (-v ) that allows you to place a version num it. This is the same identifier that is found on the Properties tab of every rule group. It follows the forma ##.##.#####.#####. This is useful if you are exporting management packs for retail sale and you want t manually track and update this parameter, but it is overkill for everyday use.

Back up your production environment on a daily basis and pay particular attention to the operations dat. Back this up using standard SQL backup procedures as described in Chapter 7.

On the management pack life cycle workflow diagram, the Base OS management pack version MP.1.2, ( in Figure 4-1 ) is synchronized between production and preproduction, as well as backed up for immedia restoration to either environment if necessary.

### 4.3.6. Integrating Vendor Revisions


A management pack can also change if the vendor releases a new version. When this happens, you may to integrate your current version with the vendor's new version while still retaining your company knowl customized rules, and enabled/disabled rules. In Figure 4-1, the management pack version MP.2.0 repre new management pack from the vendor.

The easy answer to this problem is to import the MP.2.0 version into the preproduction environment, select the Update Existing Management Pack option during the import. Ultimately, that's the right way to do it. However, before you do the import and update, take the time to find out what the differences are between two packs (point 10 in Figure 4-1 ), rather than just importing and then going through the rules, tasks, and scripts in the consoles.

There are two tools that help with this: the Convert Management Packs to XML (*mp2xml.exe*) utility and Management Pack Differencing tool (MPDiff), which takes the management pack in XML format as its input. Both tools are found in the MOM 2005 resource kit and are not officially supported by Microsoft Support.

### 4.3.6.1. MP2XML

The MP2XML tool is run on a MOM 2005 management server from a command prompt, in this case the preproduction *homemomserver3*. When run, this tool reads the *.akm* version of the management pack and outputs all of the contents in an XML file. One way of staying organized is to keep the names of the output files the same as the folder that the management packs come from. For example, the original management pack output file will be *WindowsBaseOSVendorSupplied.xml* and the current one will be *WindowsBaseOSCurrentMP.xml*.



The original Windows Server Base Operating System management pack and the MP.1.2 versions are used for this example. This is because Microsoft has not updated this management pack.

The command syntax follows the format of:

```
C:\MOMRKtools>mp2xml
c:\momrktools\currentmp\MicrosoftWindowsServerBaseOperatingSystemMP03.06.2005-14.12.
akm c:\momrktools\WindowsBaseOSCurrentMP.xml
```

The same is done for the vendor-supplied version of the management pack. The last step is to copy both files to a common directory. Despite what the product documentation says, the MPDiff tool will not work if both XML files are in the same place.

### 4.3.6.2. Management Pack Differencing tool

The MPDiff tool compares the contents and settings of the XML versions in the management packs created by the MP2XML tool. It is found in the MOM 2005 Resource Kit. The comparison is performed between two versions of the same management pack, not between completely different management packs. The MPDiff can be launched from two different executable files. The *MPDiff.Console.exe* is the command-line version of the tool. When this version is used, you must pass a source XML file, a target XML file, and an output file. Using an old and a current version of the XML file generated, the command-line syntax looks like this:

```
C:\MOMRKtools\mpdiff.console.exe /src:WindowsBaseOSOldMP.xml /tgt:
WindowsBaseOSCurrentMP.xml /out:WindowsBaseOldToWindowsBaseOSCurrent.xml
```



To view the *WindowsBaseOSDiff.xml* file, simply open it with Internet Explorer.

To view the differences interactively, launch the *MPDiff.exe* from Windows Explorer; this opens the GUI view of the tool. The GUI tool can also save the differences to an XML file that is identical to what you would get if you had run the command-line version.

Be careful the order in which you provide the XML translations to the tool. In the command-line version, you must provide the source file (*/src* option) first and then the target file (*/tgt* option). Because this is a common way to think about it, compare the old or current file to the new one. Then save the output as *<oldmpname> To <newmpname>.xml*.

For example, Figure 4-42 shows the Summary/Changed view (point 1 in Figure 4-42) of the differences between the *WindowsBaseOSVendorSupplied.xml* and the *WindowsBaseOSCurrentMP.xml*.

MPDiff classifies changes into four categories (point 2 in Figure 4-42):

### *Functional*

A material change to the configuration of the selected object that changes how it performs its basic functions. The MPDiff tool looks at all the objects listed in the navigation pane in both management packs for differences. In Figure 4-42, the Summary/Changed view has been selected. A change in the functionality of an object means a substantial configuration change between the two, such as a change in the rules or criteria for a processing rule or the creation of an override.

Figure 4-42. View the differences between management packs interactively using the MPDiff tool



*Informational*

Changes listed in this column reflect cosmetic alterations to the selected object; for example, a change to the description (point 3 in Figure 4-42).

*Organizational*

If an object is moved or the relationship between objects changes, that would be included here. For example, the removal of a child rule group from a parent rule group is a relationship change.

*Deployment*

This includes any changes to how an object is deployed or applied to an agent. For example, if you change the association of a script from a rule in one rule group to a rule in another and those rule groups are associated with different computer groups, that would be reflected here.

If there is any descriptive text that is associated with the change, it will be displayed in the Details pane in Figure 4-42 ).

The other three columns contain the GUID (ID) of the object, the object type (area), and the name of the object as they appear in the Administrator console. You can use the values in these columns to look up the individual objects in the Administrator console and examine the details there for objects.

**4.3.6.3. Merge the management packs**

After examining the differences between the newly revised vendor management pack (Version MP.2.0) and the current one (Version MP.1.2), perform the import into the preproduction environment and select Update Existing Management Pack. You will be prompted for a location to save the backup to; this can be local or anywhere, because the real backup is in the *CurrentMP* directory.

This brings the management pack life cycle full circle. From a procedural perspective the MP.2.1 in Figure 4-42 is the same as MP.1.1. In the production to preproduction synchronization process you can only make a change when you are testing and tuning the MP.2.1 version. To prevent the current production version (MP.1.2) from overwriting the MP.2.1 version, you must avoid importing MP.1.2 into preproduction. However, MP.1.2 is backed up to the *CurrentMP* directory on a daily basis.

To do this, comment out the line in the *ProdToPreProdSynch.bat* file that performs the import by entering the beginning of the last line so that it now reads:

```
rem ManagementModuleUtil.exe -I homemomserver3
\\homesrv02\mptransferfolder\currentmp\MicrosoftWindowsServerBaseOperatingSystemMP%var
Date%varTime% .akm -R
```

## 4.4. Creating Simple Management Packs

Creating a management pack is an entirely different process than tuning an existing one. It can be quite involved and, depending on the functionality you want to include, it can require the knowledge of an application architect, as well as advanced programming skills for scripting and managed code responses. You need knowledge of all the indicators that can be used to define health states, services, events, performance counters, registry keys, and other identifying information.

However, odds are that you will never have to author a management pack for an application as complex as SQL, MOM, or Exchange. If you need to author complex management packs, refer to the Management Pack Development Guide from Microsoft. At some point you may have to provide basic monitoring for a homegrown application or a single service at your company. That's what this section is about, implementing a simple management pack to monitor the Windows Update Service.

### 4.4.1. Management Pack Wizard

To monitor an application, the application must either run as an installed service, have performance monitor objects, or have a *.dll* that writes events to any of the event logs. Without one of these three, it is very difficult to monitor the application.

The Management Pack Wizard (*mpwizard.exe*) is a MOM 2005 resource kit tool that walks you through the creation of new MOM roles and the processing rules that monitor that role based on defined components. It generates event rules that feed into health state and alerting, so the defined role will appear in the State view in the Operator console. If the application has performance monitor objects, the wizard can generate performance rules. The output of the wizard is the management pack *.akm* file and the XML version of the management pack.

The wizard does not create computer groups, views, or tasks for the Operator console; in fact, when imported, a management pack created with this tool is not associated with any computer groups. This means that its rules will not be deployed and service discovery will not be performed until the rule group is associated with a computer group. It also does not create any alert rules, so if you want an action taken in addition to an alert being displayed in the Operator console, you will have to configure those as well. The wizard only runs on a management server, although at the appropriate points it will allow you to browse the services, events, and performance monitor counters on remote machines if the application of interest is not installed on the local machines.

This example uses the tool to create a management pack for the Windows Update Service. This application installs a service named Automatic Updates and an associated *.dll* named *wuaueng.dll* that the event log messages can be pulled from. It doesn't have any performance monitor objects. This is an important point to remember.

To start the tool, navigate to the *MOM 2005 Resource Kit\Tools\Management Pack Wizard* directory and launch *mpwizard.exe*. This brings up the first page ([Figure 4-43](#)), where the three types of monitoring (service, performance, or event) for this management pack are listed.

Figure 4-43. Select which types of monitoring you want to build into this management pack



Knowing which types of monitoring will be performed beforehand is important because this tool presents pages based on the selections made here. Because the application does not have any associated performance monitor counters, don't bother selecting the Performance Threshold Monitoring options. However, since the application installs a service and generates events, select those options. Click Next to go to the Role Name creation page ([Figure 4-44](#)). The role name that is defined here will appear in the State view of the Operator console for every computer that is running this service. For this example, the role is called AutomaticUpdate.

Figure 4-44. Role creation page





Roles are the objects in MOM 2005 that reside immediately underneath the computer objects (as shown in [Figure 4-40](#)). The health state of each role on a computer is included in the health state roll-up calculation for the computer as a whole.

Click Next to bring up the Role Components creation page ([Figure 4-45](#)). Two components that match the types of monitoring indicated on the Welcome page have been created: AutoUpService and AutoUpEvent. On this page, enter a component name in the upper box and then click Add.

Click Next to go to the page where the details of the service are defined. It is initially blank, so click the Add button to select the service that you want to monitor ([Figure 4-46](#)). This application only installs one service, so there is only one to monitor. This is very different from building a management pack for an application like Exchange 2003 that has more than four services to monitor. However, you would never use this tool to create a management pack for something like Exchange.

On this page you can attach to a remote computer to get a list of the services there. You'd need to do this if the service you want to monitor is not on the local machine. Also, at this point the tool is only looking for a sample of the service to be monitored. Based on this, it will create the service discovery rules that will discover instances of this service on all the computers associated with this rule group.

Figure 4-45. Create the components that the role will be defined by



Figure 4-46. Select the service you want to monitor from the local machine

Click Add to go back to the Windows Service Monitoring page, which has now been populated with the Automatic Updates value in the Service column. Select the appropriate component that was created previously to associate with this service; in this case, the AutoUpService component (see [Figure 4-47](#)).

If you had selected the Performance option on the Welcome page, the wizard would go through the identification of performance objects and allow you to set thresholds and their associated red X and yellow warning triangle flags.

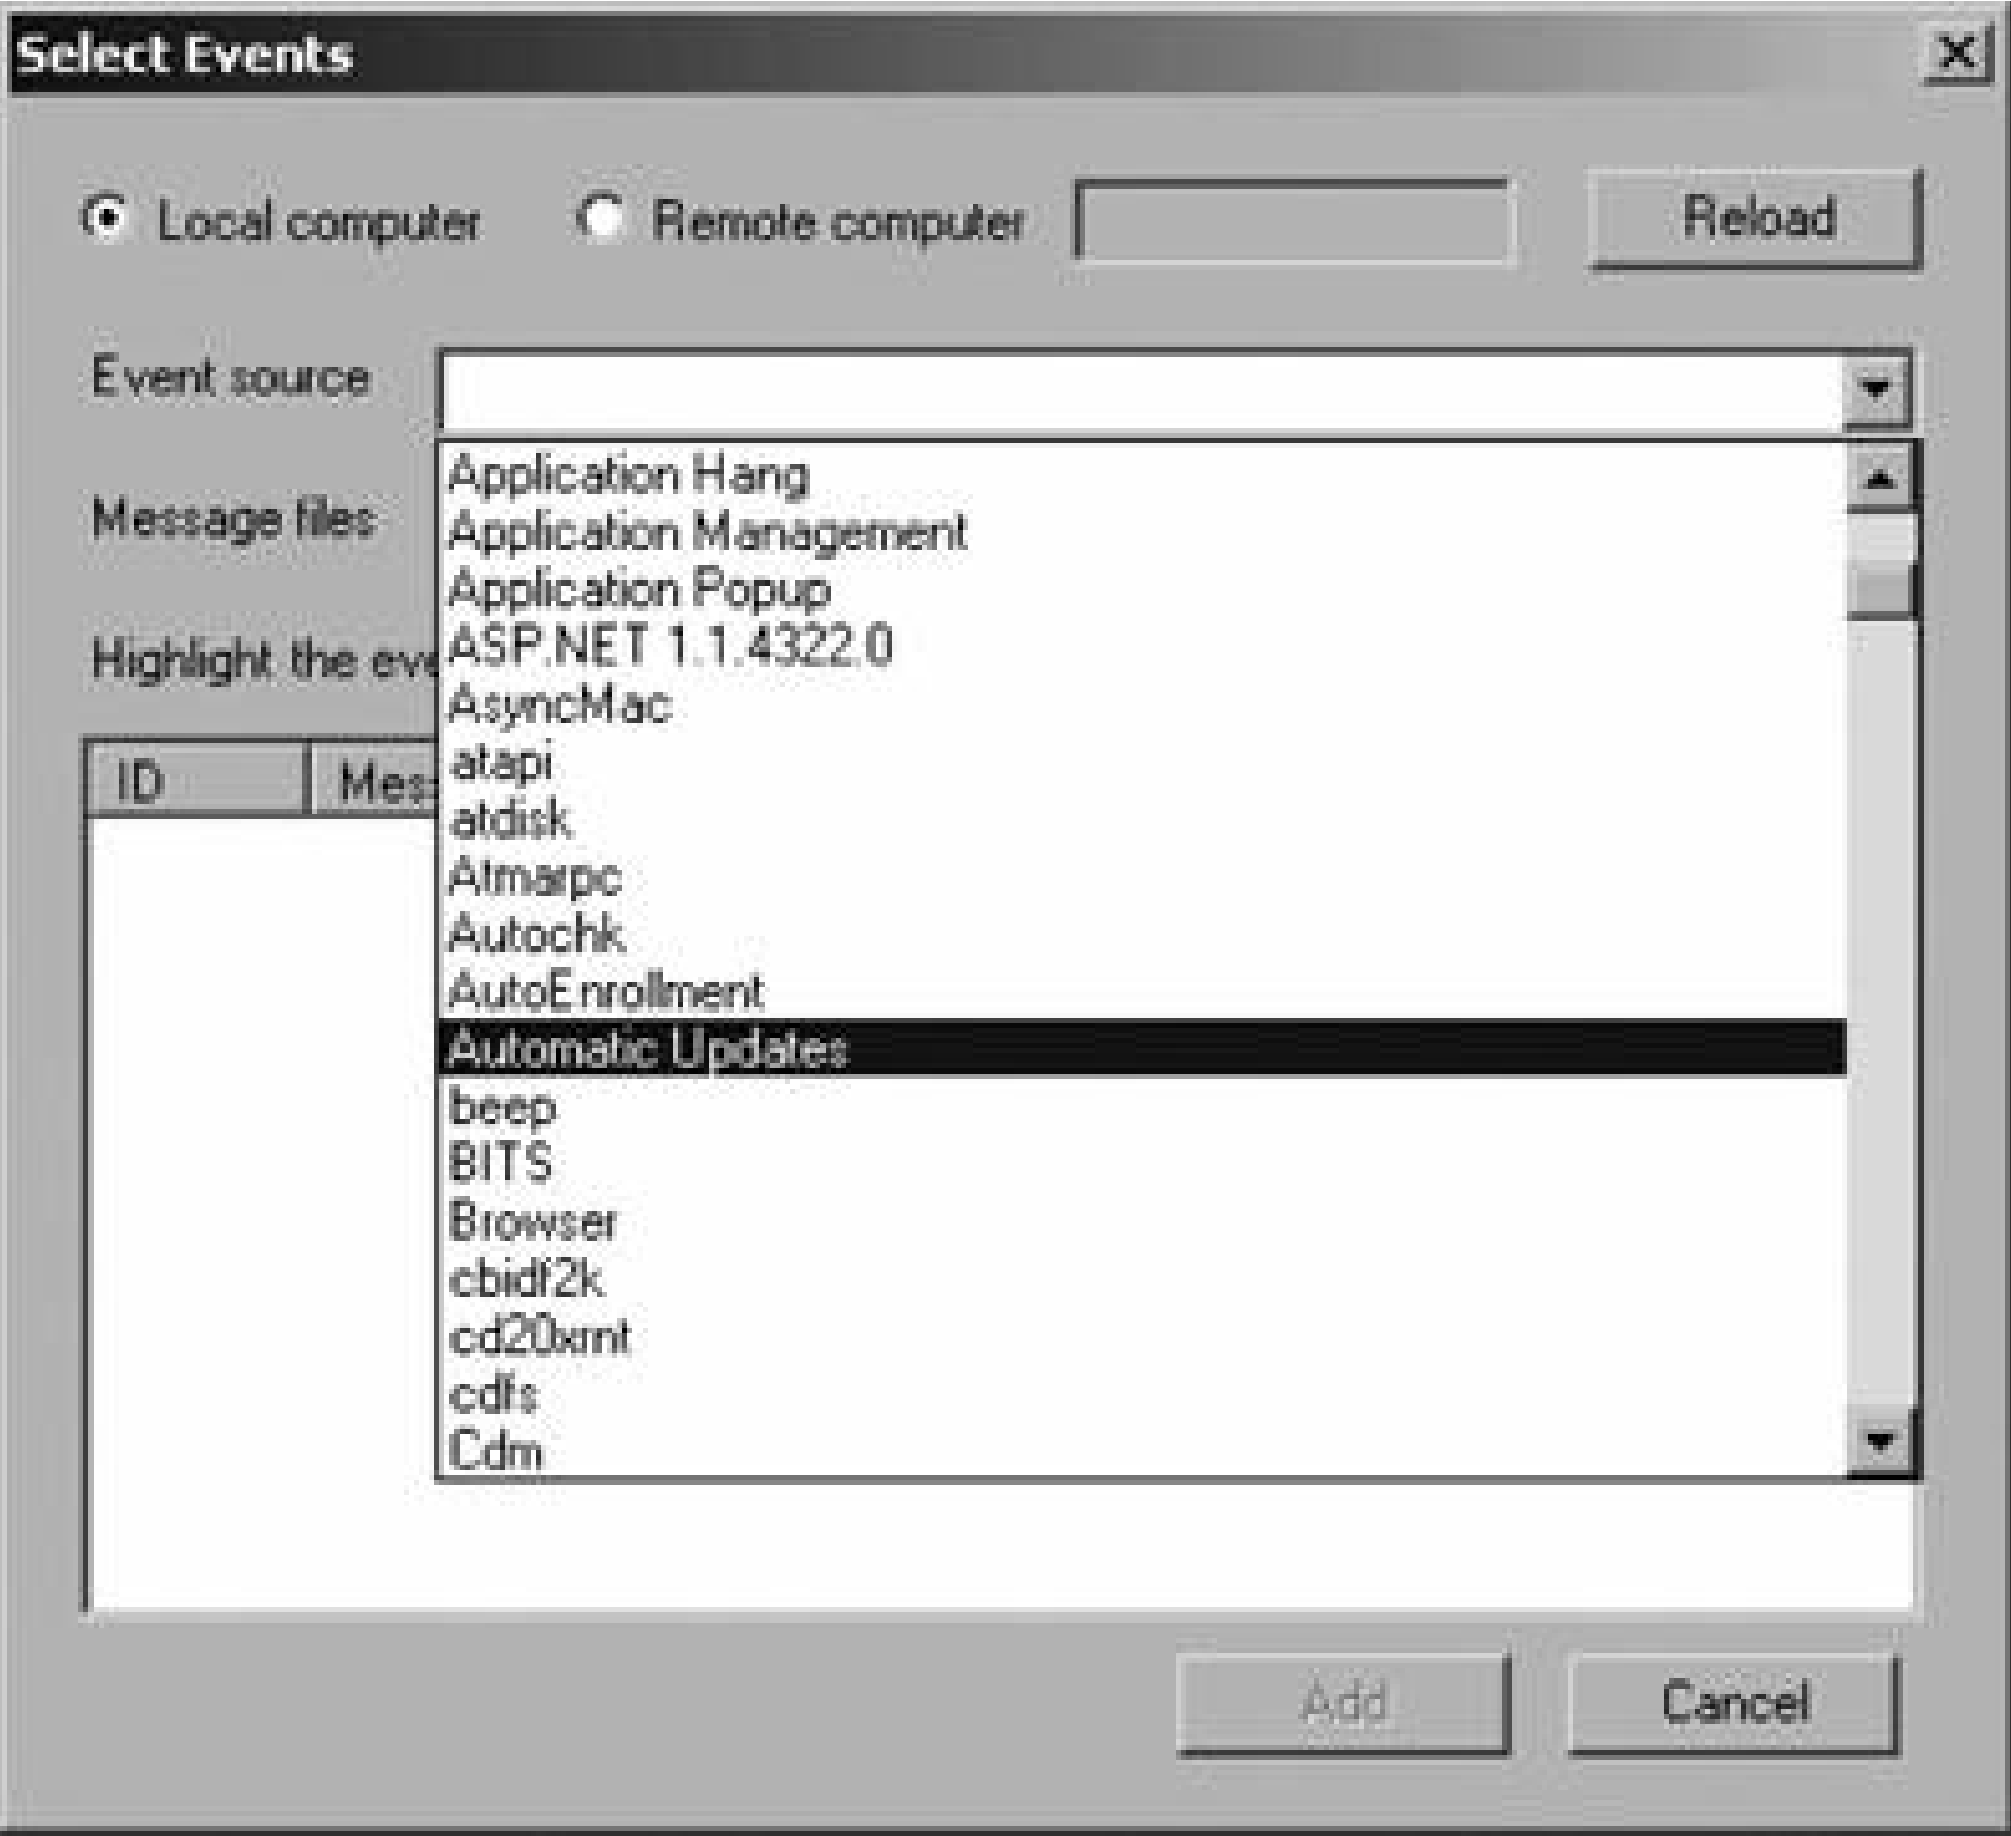
In this case though, since Automatic Updates has no performance objects, clicking Next brings you to the events identification (see [Figure 4-48](#)).

Figure 4-47. Create the association between the service to be monitored and the component that you created



Figure 4-48. Select the events to monitor





The Event source drop-down menu lists all the event sources that are available on the local machine. Note again that you can bind to a remote machine if necessary. Select Automatic Updates here and all of the event IDs along with their descriptions are populated in the lower box (see [Figure 4-49](#)).

Figure 4-49. A list of all the event IDs that the Automatic Updates application will generate

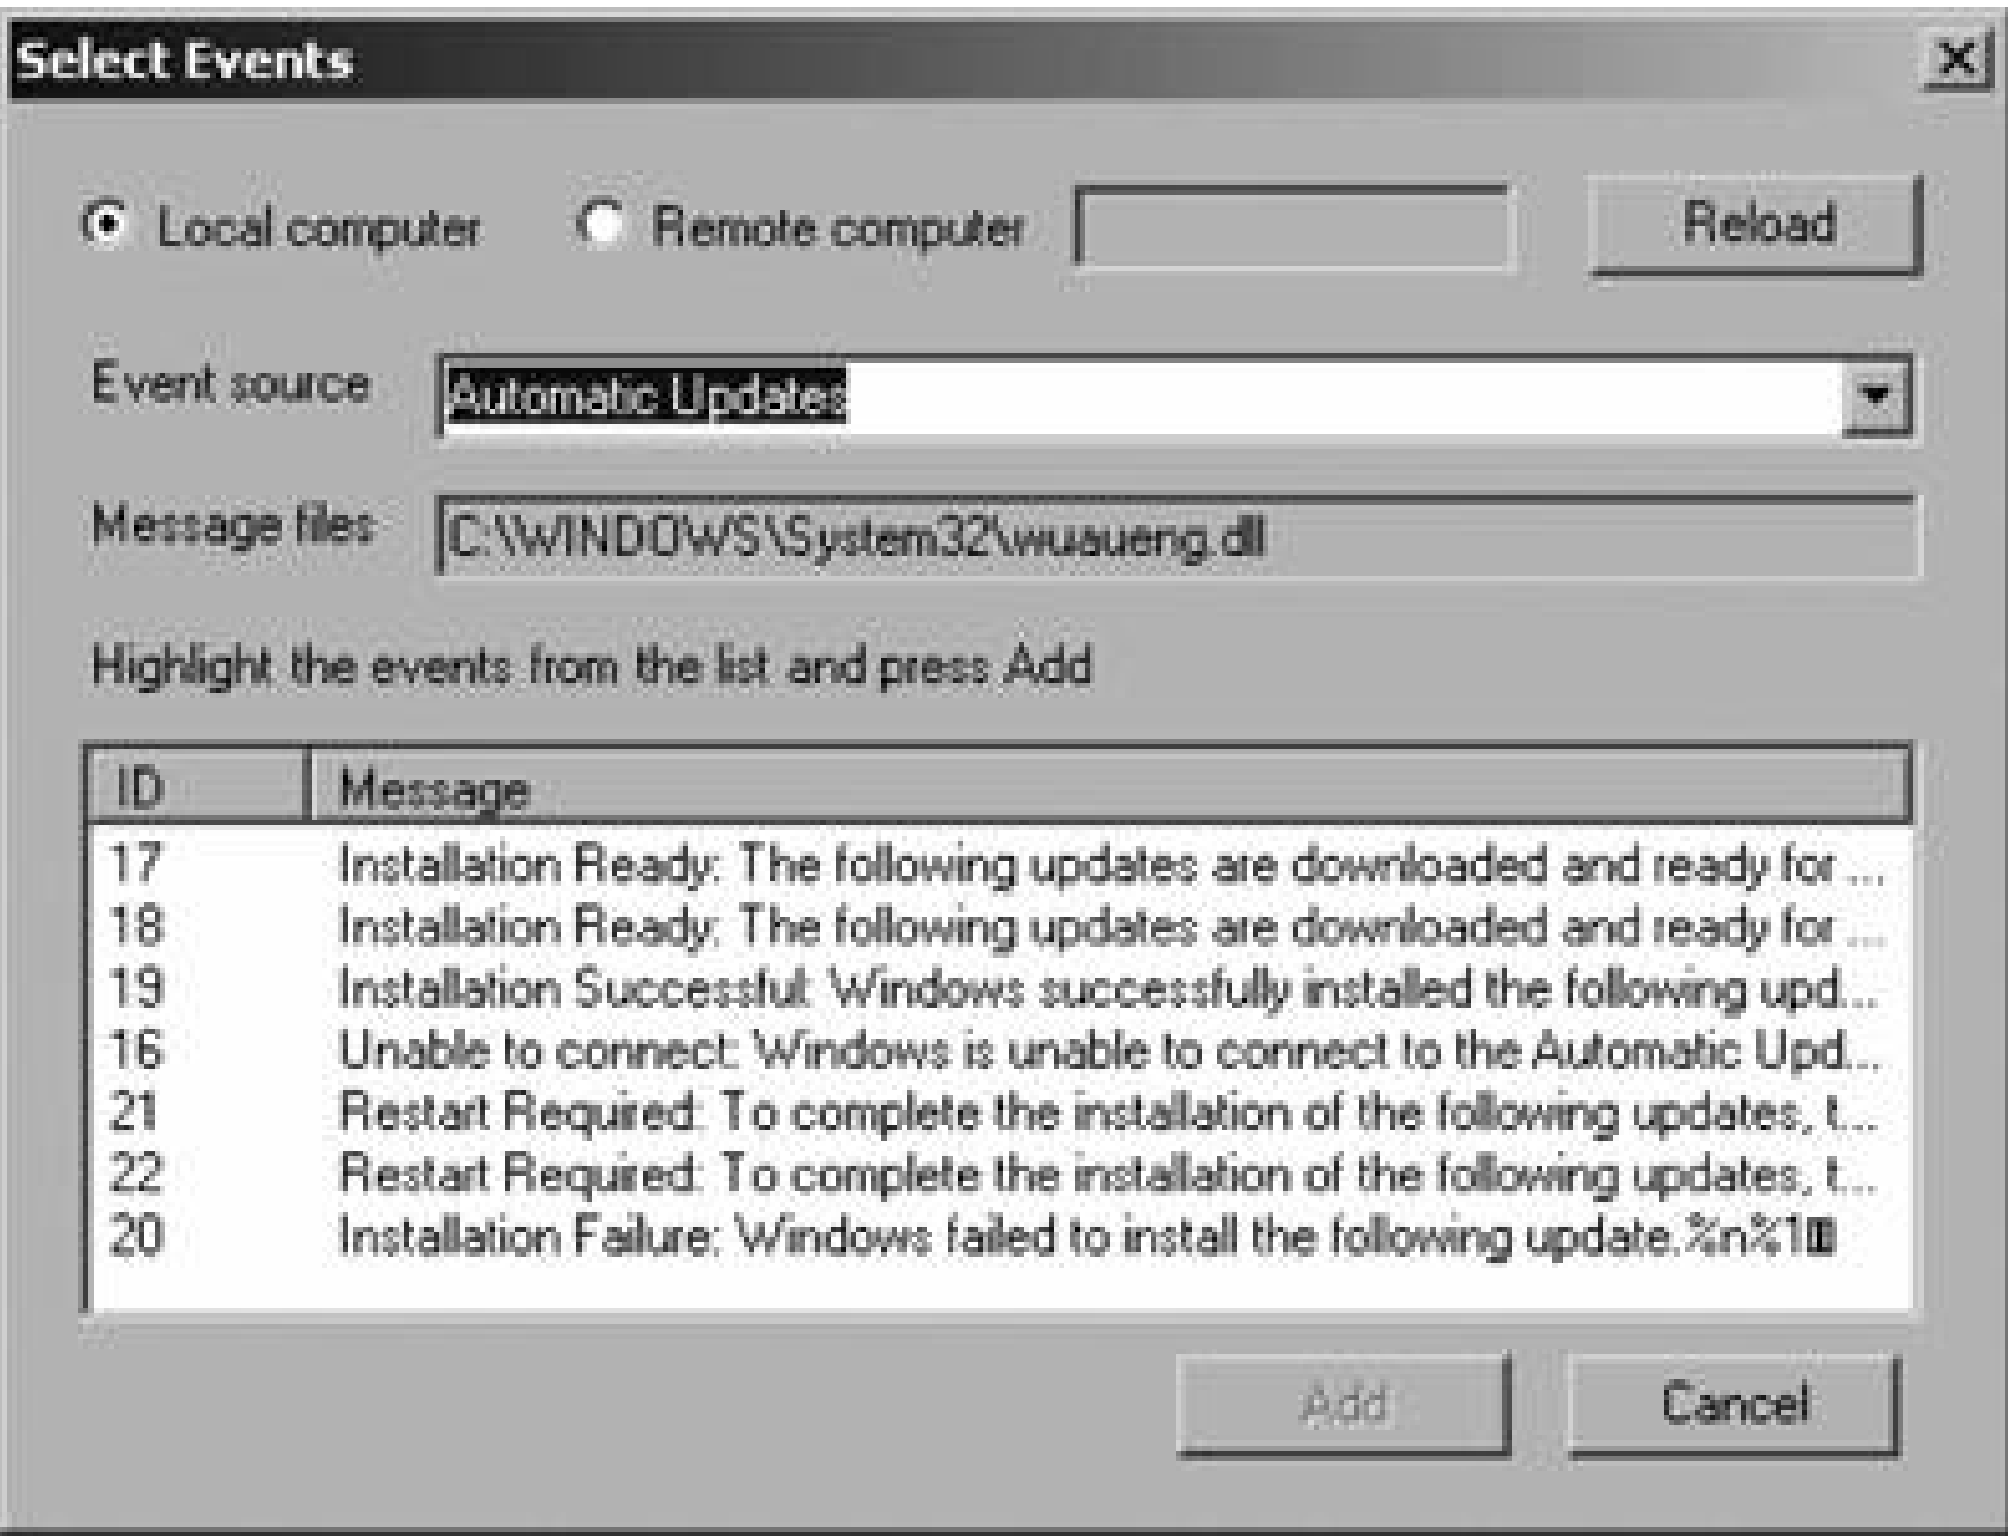


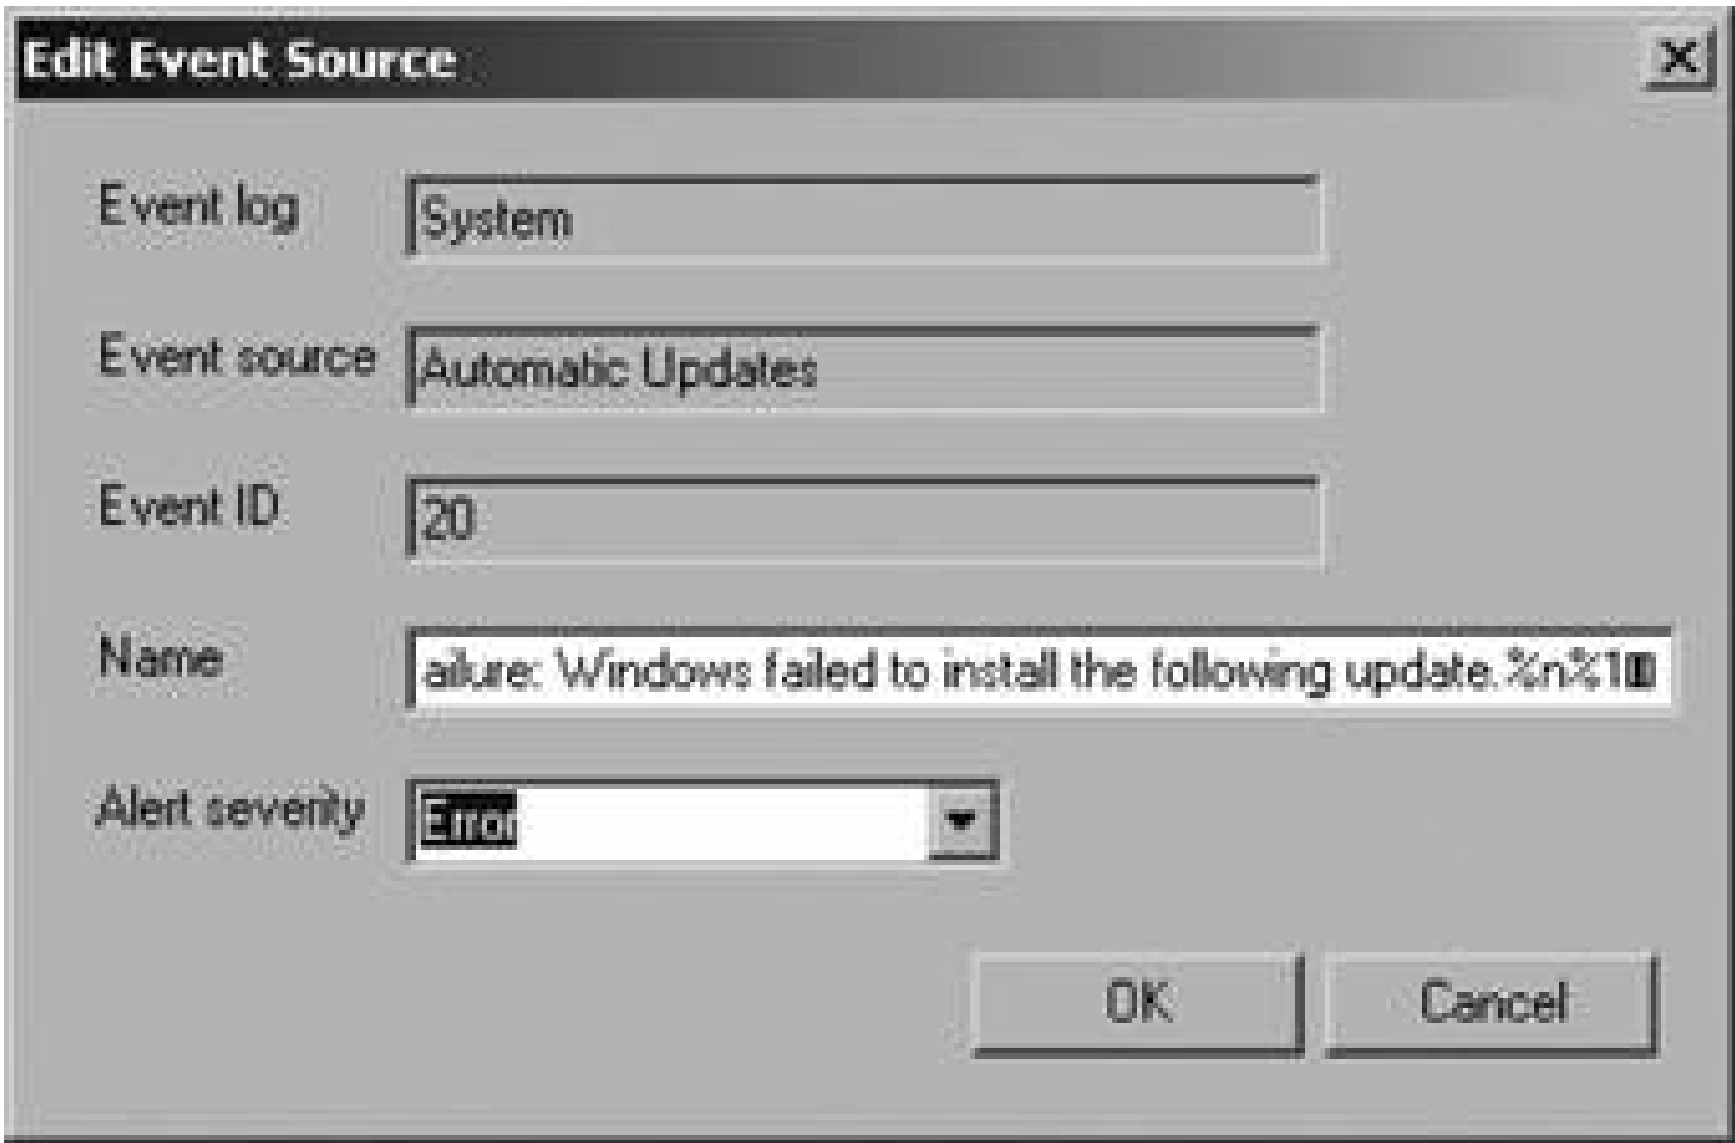
Figure 4-50. This page shows the event rules that will be created and the default alert severity that they will generate

This list is relatively short compared to some others (think of how many events are associated with the OS), so select all of them and click Add. This populates all the events into the Event Source Monitoring page with a default alert severity of Warning (see [Figure 4-50](#)).

Rules that indicate similar events, such as the two Installation Ready or the two Restart Require

rules, can be combined. When two event rules are combined into a single rule, the criteria in the resulting event rule will include both event numbers. This is a useful thing to do when two events are similar and you would want to generate the same alert for both of them. By selecting any individual rule and editing it, the default name and the severity level can be changed (see [Figure 4-51](#)).

Figure 4-51. Change the default alert severity from Warning to Error for a patch installation failure



Combine the Installation Ready and Restart Required rules and edit the default alert severities of Success and Failure rules to be Success and Error, respectively. This creates a framework for the event rules, shown in [Figure 4-52](#).

Figure 4-52. The framework for the Automatic Update rule group event rules



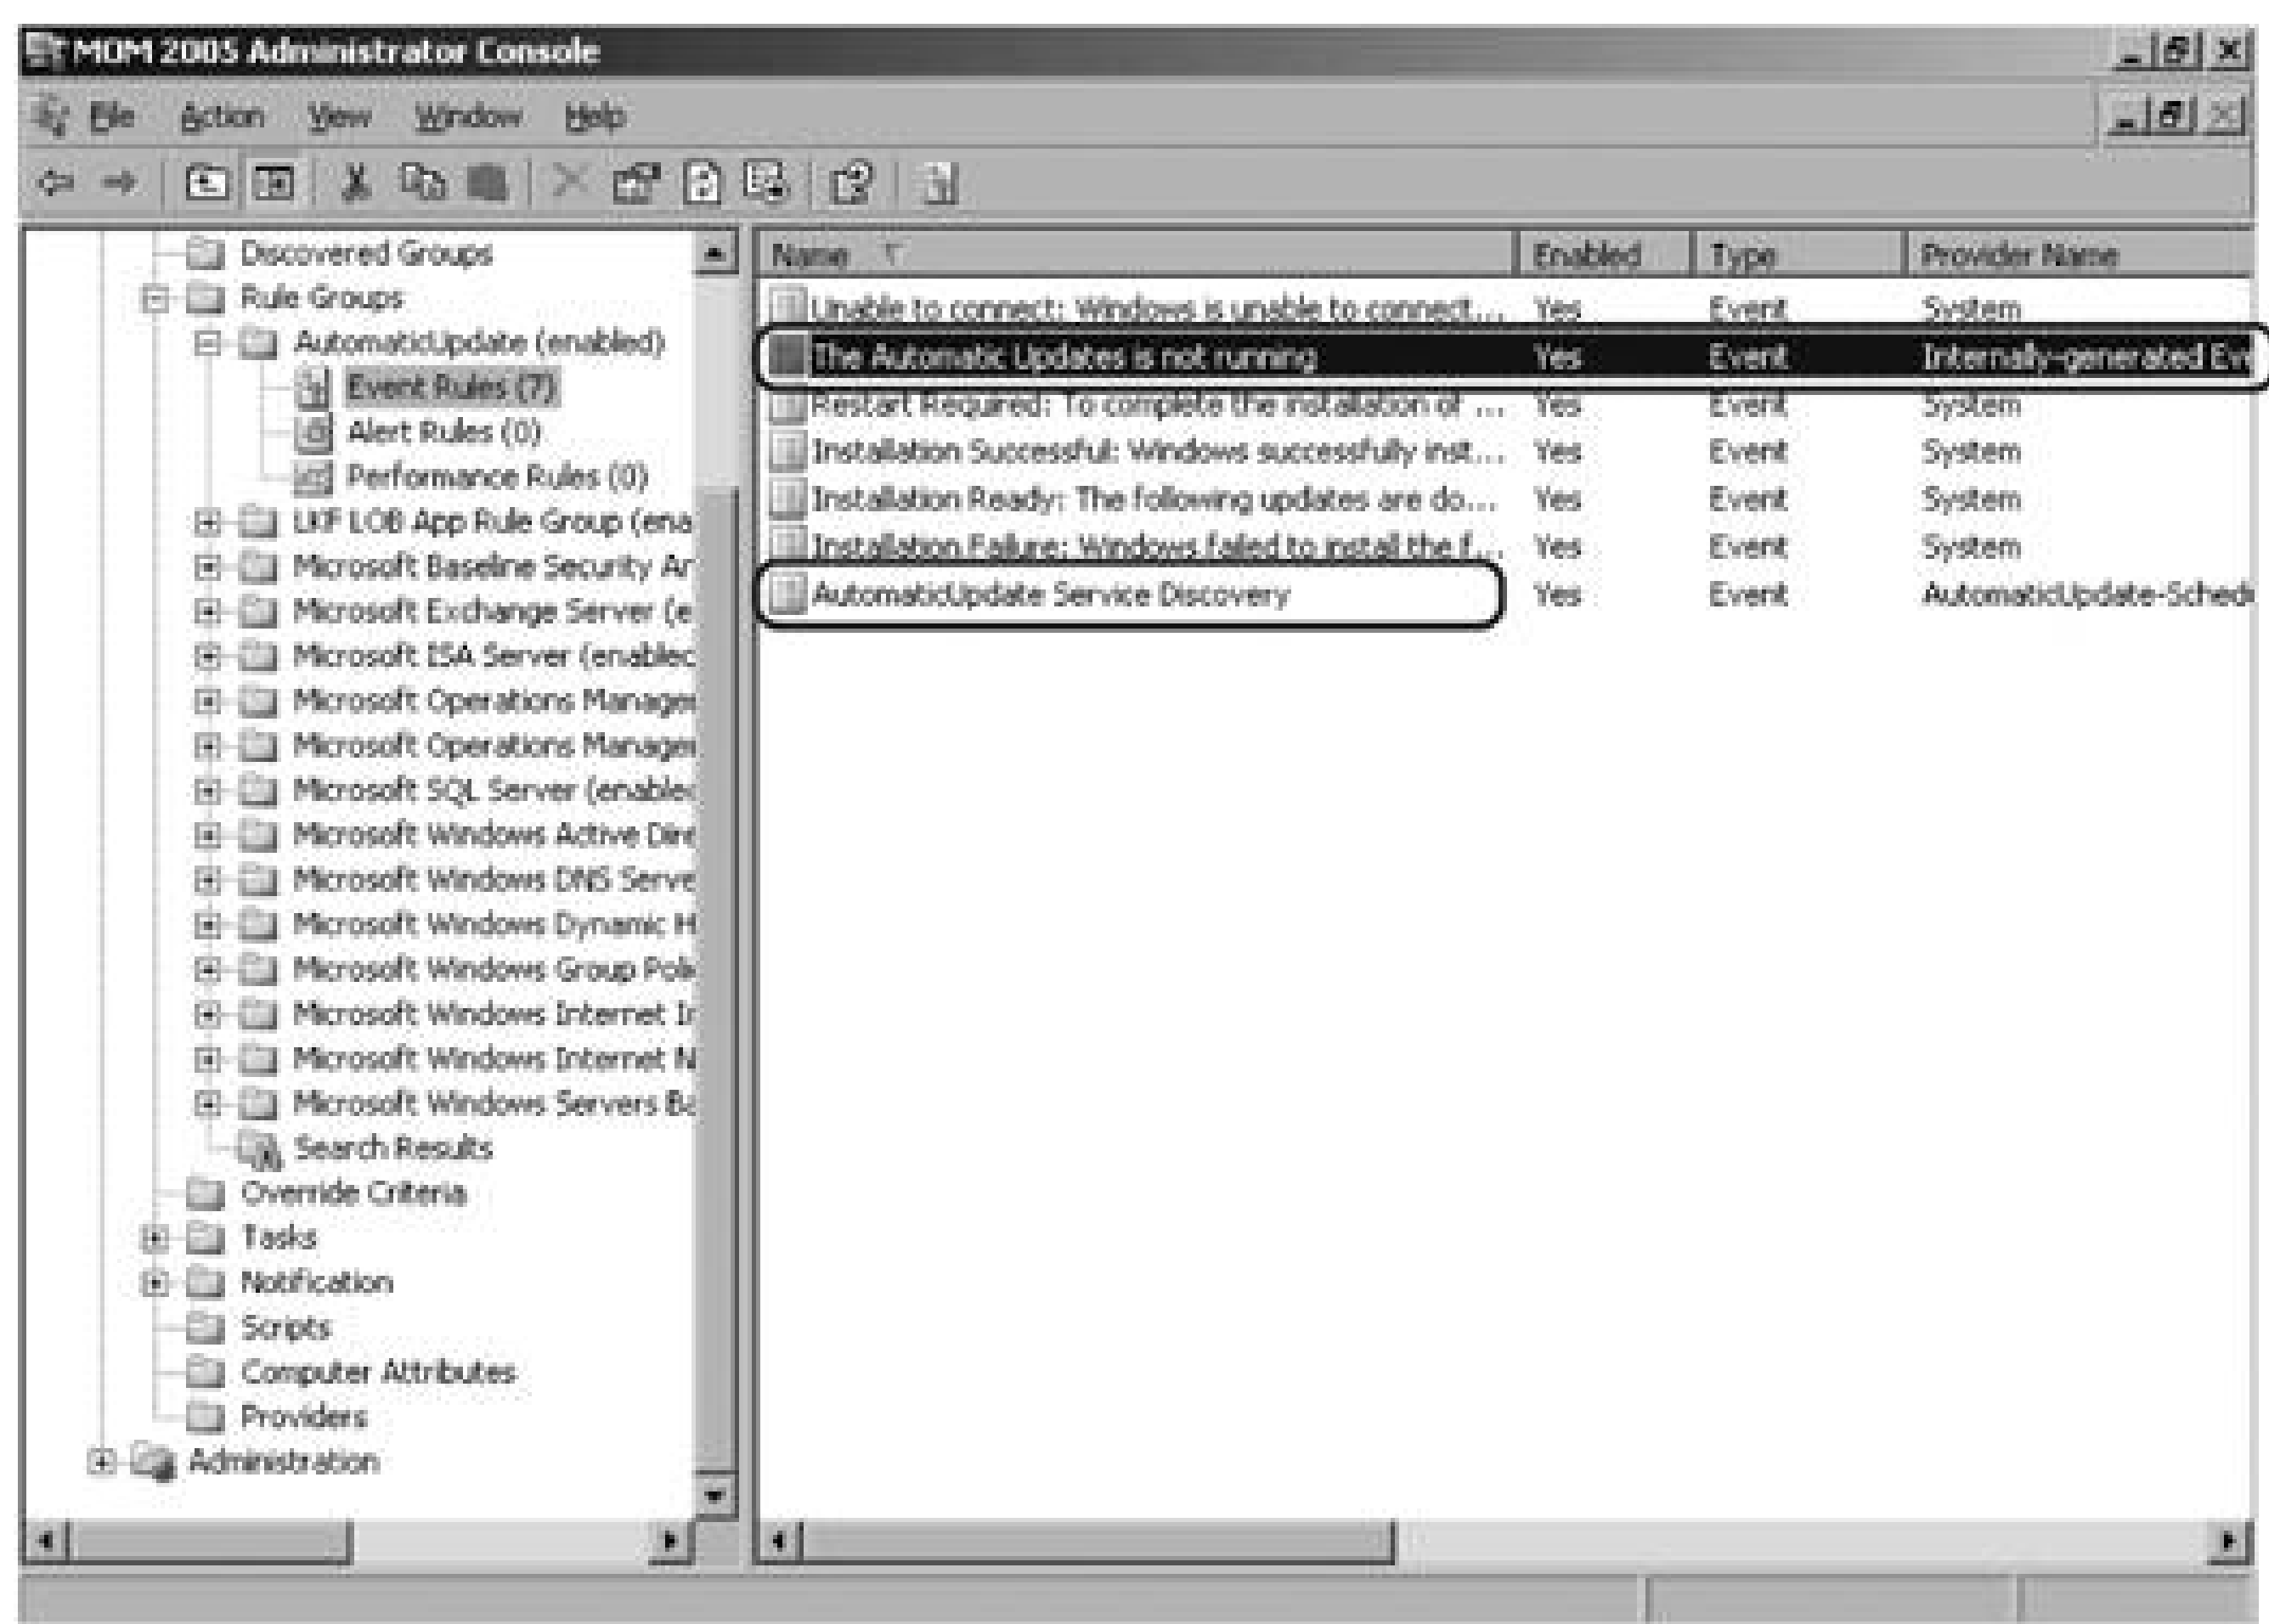


Click Next to go to a Summary page to review your choices. After that the management pack is built and the customary Finish page appears. It also tells you the location and the name of the newly minted management pack. The tool will create the *.akm* and XML files in the same directory that it was launched from and with the name you created for the role.

### 4.4.2. Importing the New Management Pack

Import the management pack into the preproduction environment using the Import/Export wizard. The rule group that has been created has seven event rules, no alert rules, and no performance rules ([Figure 4-53](#)). Five event rules map to the five shown in [Figure 4-52](#), the other two (circled) are the service discovery timed event rule and the health state alert rule.

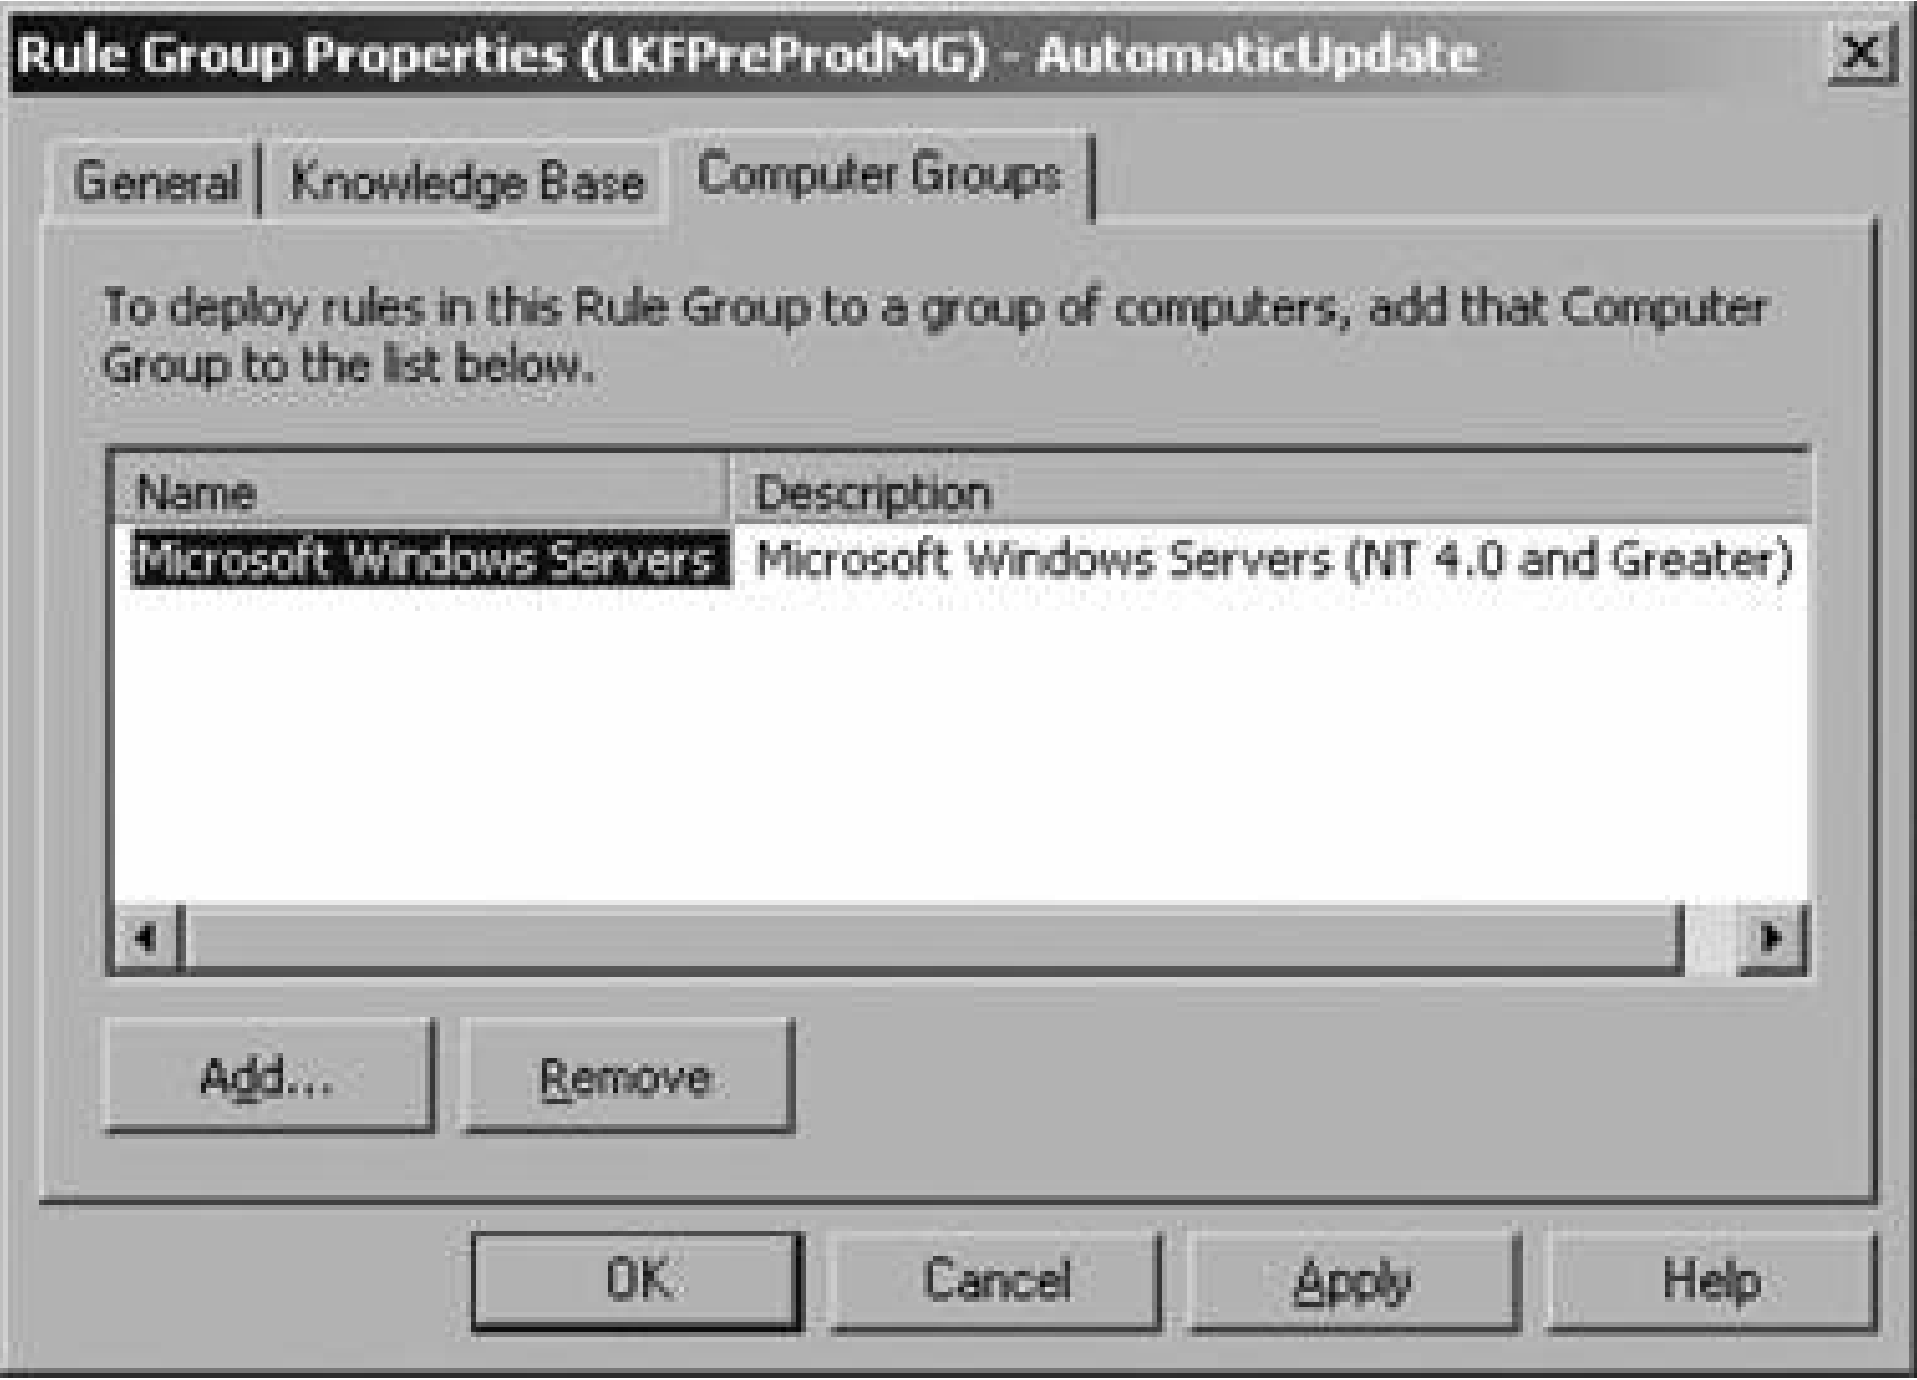
Figure 4-53. The Import/Export wizard gnerates service discovery and service state rules



All of the rules have been generated by the wizard. The script that the service discovery timed event calls was generated by the wizard and will run when the rule group is associated with a computer group. You already know what happens once the rule group is associated with a computer group.

On the context menu for the AutomaticUpdate rule group, associate it with a computer group. For the widest possible distribution for this rule group, associate it with the Microsoft Windows Servers computer group, which contains all Windows servers of version NT 4.0 and higher [Figure 4-54](#)).

Figure 4-54. Associating the Microsoft Windows Servers computer group with the AutomaticUpdate rule group



The AutomaticUpdate service discovery is a timed event that runs once an hour, so health state information does not show up immediately in the State view, but once it does, the AutomaticUpdate role is added to all computers that are running this service [\(Figure 4-55\)](#).

Figure 4-55. The State view in the Operator console offers proof that the management pack is collecting health state information for the AutomaticUpdate role



Microsoft Operations Manager 2005 - Operator Console - LKFPredProdMG

File Edit View Go Help

Tasks

Group: Microsoft Windows Servers

012

State Views

State

Active Directory

Microsoft Baseline Security

Microsoft Exchange Serve

Microsoft ISA Server 2000

Microsoft Operations Man

Microsoft SQL Server

Microsoft Windows Base C

Microsoft Windows DNS Se

Microsoft Windows Dynam

Alerts

State

Events

Performance

Computers and Groups

Diagram

My Views

Public Views

State

State	Domain	Computer	MOM Agent	AutoMaticUpdate	OS	Disk
Warning	HOMELAB	HOMEMOMSERVER3	✓	✓	✓	✓
Success	HOMELAB	HOMETSAFIREWALL	✓	✓	✓	✓
Success	HOMELAB	HOMESRV02	✓	✓	✓	✓

State Details - Computer: HOMEMOMSERVER3

Computer	Open Alerts	Events	Last Heartbeat
HOMEMOMSERV...	2	107	3/6/2005 11:25:...

Total: 3 Item(s) Selected: 1 Item(s)

Last refresh: 3/6/2005 11:25:25 PM

localhost

## 4.5. Summary

The entire MOM 2005 infrastructure exists to execute the instructions in management packs, and MOM 2005 serves no purpose without them. The management pack life cycle should be predictable. This chapter described the necessary steps to control the versioning, tuning, backup, and restore of management pack life cycles. Management packs consist of rule groups, providers, computer groups, scripts, attributes, and tasks, although only rule groups, computer attributes, and providers are required for the most basic management pack.

Tuning management packs occurs in two phases: the first is in preproduction where you enable and disable rule groups and individual rules based on their applicability to your environment. Ensure alerts are of the desired severity and that the right people are being notified. Once a management pack is deployed into production, the second phase of tuning occurs, focusing on alert reduction. Noise alerts are tuned out by creating overrides. Alerts that are generated by a single rule are more tailored to the different machines that they have been applied to and, therefore, produce more relevant information.

In addition to tuning out noise, once in production, a management pack is enriched with company knowledge gained from the troubleshooting efforts. This makes the management pack invaluable to anyone else that must troubleshoot an alert raised in MOM. The solution to this problem has already been found and captured.

To protect your company-specific management packs and to provide a test environment that is as close as possible to the production environment, this chapter covered how to back up the production management packs and synchronize them into preproduction at the same time.

Hopefully, vendors will update their management packs and release updates at regular intervals. When they do, you need to merge the new vendor management pack with your current knowledge-rich company version of the management pack. It pays to compare them so you know what the new vendor management pack will change and what to expect. The chapter covered some tools and provided guidance on how to do this.

Creating simple management packs based on event, performance, and service monitoring was covered. The management pack wizard utility will create a basic management pack based on services, events, and performance monitor counters. This includes the creation of a basic health state model that is reflected in the State view of the Operator console. For guidance on authoring advanced management packs that have complex health state models (beyond whether a single service is up or down), see the Management Pack Development Guide from Microsoft.

To round out the coverage of how MOM 2005 administrative tasks and functions work together to manage the management pack environment, [Chapter 5](#) details the global settings in a MOM 2005 management group that control functions of agents, management servers, security, communication, and the operations database.

# Chapter 5. Administering Global Settings

The Global Settings node in the Administrator console allows you to configure the default values for agents, management servers, and security for the whole management group. In some cases, you can override the settings for the agents and the management servers by using their properties. This was shown in [Chapters 2](#), [3](#), and [4](#).

There is another class of global settings that apply to the management group as a whole. They do not have a category name of their own, but are the rest of the settings (other than Agents and Management Servers in the Global Settings node), as shown in [Figure 5-1](#).

This chapter explores the functions of the other global settings that control management group behavior and when you would want to modify them.

Figure 5-1. All other global settings

These nine settings can be classified into three categories by function: alerts, connections and communication settings, and maintenance settings. The Security settings tab is covered in [Chapter 3](#).





# 5.1. Alerts

The first four settings relate to alerts and responses to alerts. The Notification Command Format, the Alert Resolution State, and the Custom Alert Field global settings work in concert with the event, performance, and alert rules in the rule groups. These settings set the default behavior of an alert, or modify the default information that is included in an alert at the time of its creation. All data that is included in an alert can be used in the Operator console to build filters for custom views.

## 5.1.1. Notification Command Format

When you configure responses in alert rules, you have the option to send a message to a notification group (see [Figure 5-2](#)). Here, the Network Administrators notification group has been selected.

Figure 5-2. Alert rule notification response options

One way to communicate with a notification group is to run an executable and pass application- and alert-specific fields to it. This is done on the Command Format tab (see [Figure 5-3](#)).

You can either define a custom command format for any particular alert rule or use the default command format. The default command format is set at the global level. When you select the Notification Command Format tab in the Global Settings container, you see the configuration page shown in [Figure 5-4](#).

The intended use for this feature is to call third-party paging applications. This feature would be used

for instance, if your MOM 2005 installation does not have access to an SMTP-based paging service or SMTP server for email, but you still need to send alerts to operators in addition to displaying them in the Operator console.

Figure 5-3. Alert rule notification response options Command Format tab

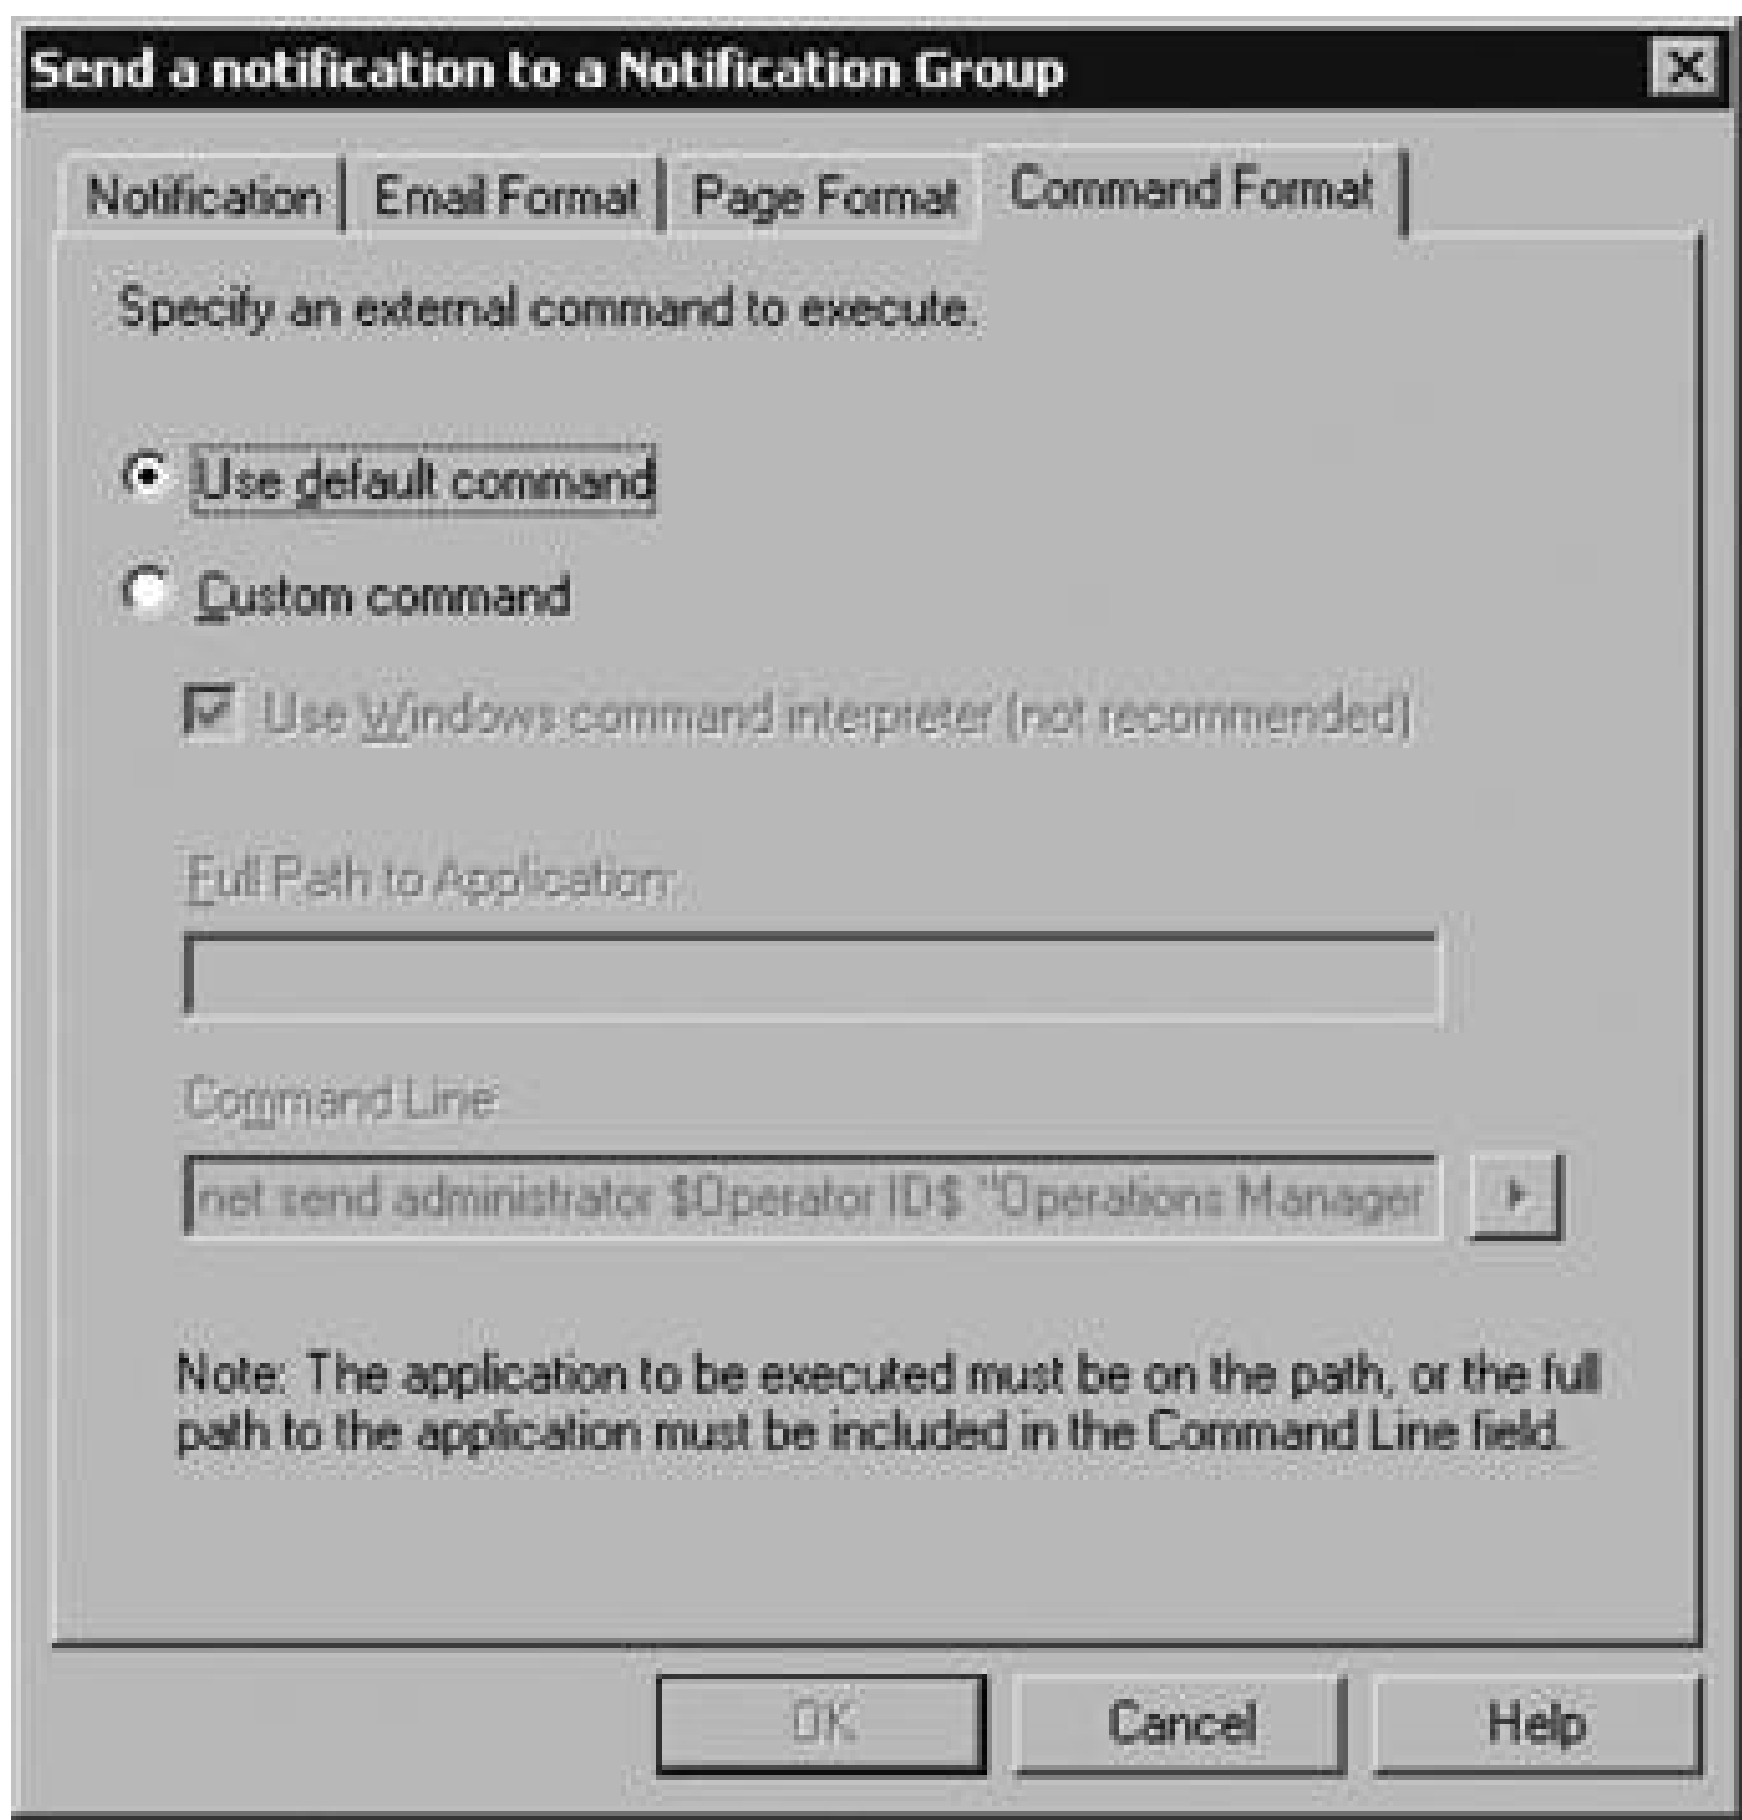


Figure 5-4. Global Settings Notification Command Format



One way to do this is to install a third-party paging program on the management server. When an alert is raised and the alert rule is configured to send out a message to a notification group, this option will launch the paging application to create the message and send it out via a direct attached modem or shared modem pool. When sending notifications out via modem, be aware that in the event of a notification flood, the modem may not be able to keep up with the number of notifications from MOM. In this case, make sure that only the most critical notifications are going out to reduce the load.

The best way to use this feature is to explicitly call an executable that is designated in the Full Path to Application text box. That means that you can only pass options to the executable that it understands. In addition to the executable parameters, event and alert fields can also be passed. Note that you will only receive the number of characters that the paging program supports.

Sending a page is not the only use, however, for the Notification Command Format. You can call any executable from here through the Windows command interpreter, as shown in [Figure 5-4](#).

For example, you might want to send a network notification to an operator who is logged onto the network, via the `net send` command. To do this, you need to check the "Use Windows command interpreter" box. This will pass the command line that you enter to `cmd.exe`. This is the exact same thing as typing the command-line syntax at a command prompt with the addition of MOM-specific variables.

Constructing the parameters string in the Command Line text box is simple. It consists of a combination of plain text for every notification and variables that are replaced with alert- or event-




specific data at the time the notification is generated. The variable fields are delineated with a dollar sign (\$) bracketing them. The plain text is simply written out.

You don't have to know all of the possible event and alert fields. The right-facing arrow to the right of the Command Line box (see [Figure 5-5](#)) branches into a list of alert fields and event fields. Select the variable you want to insert into the Command Line string and it is placed at the insertion point.

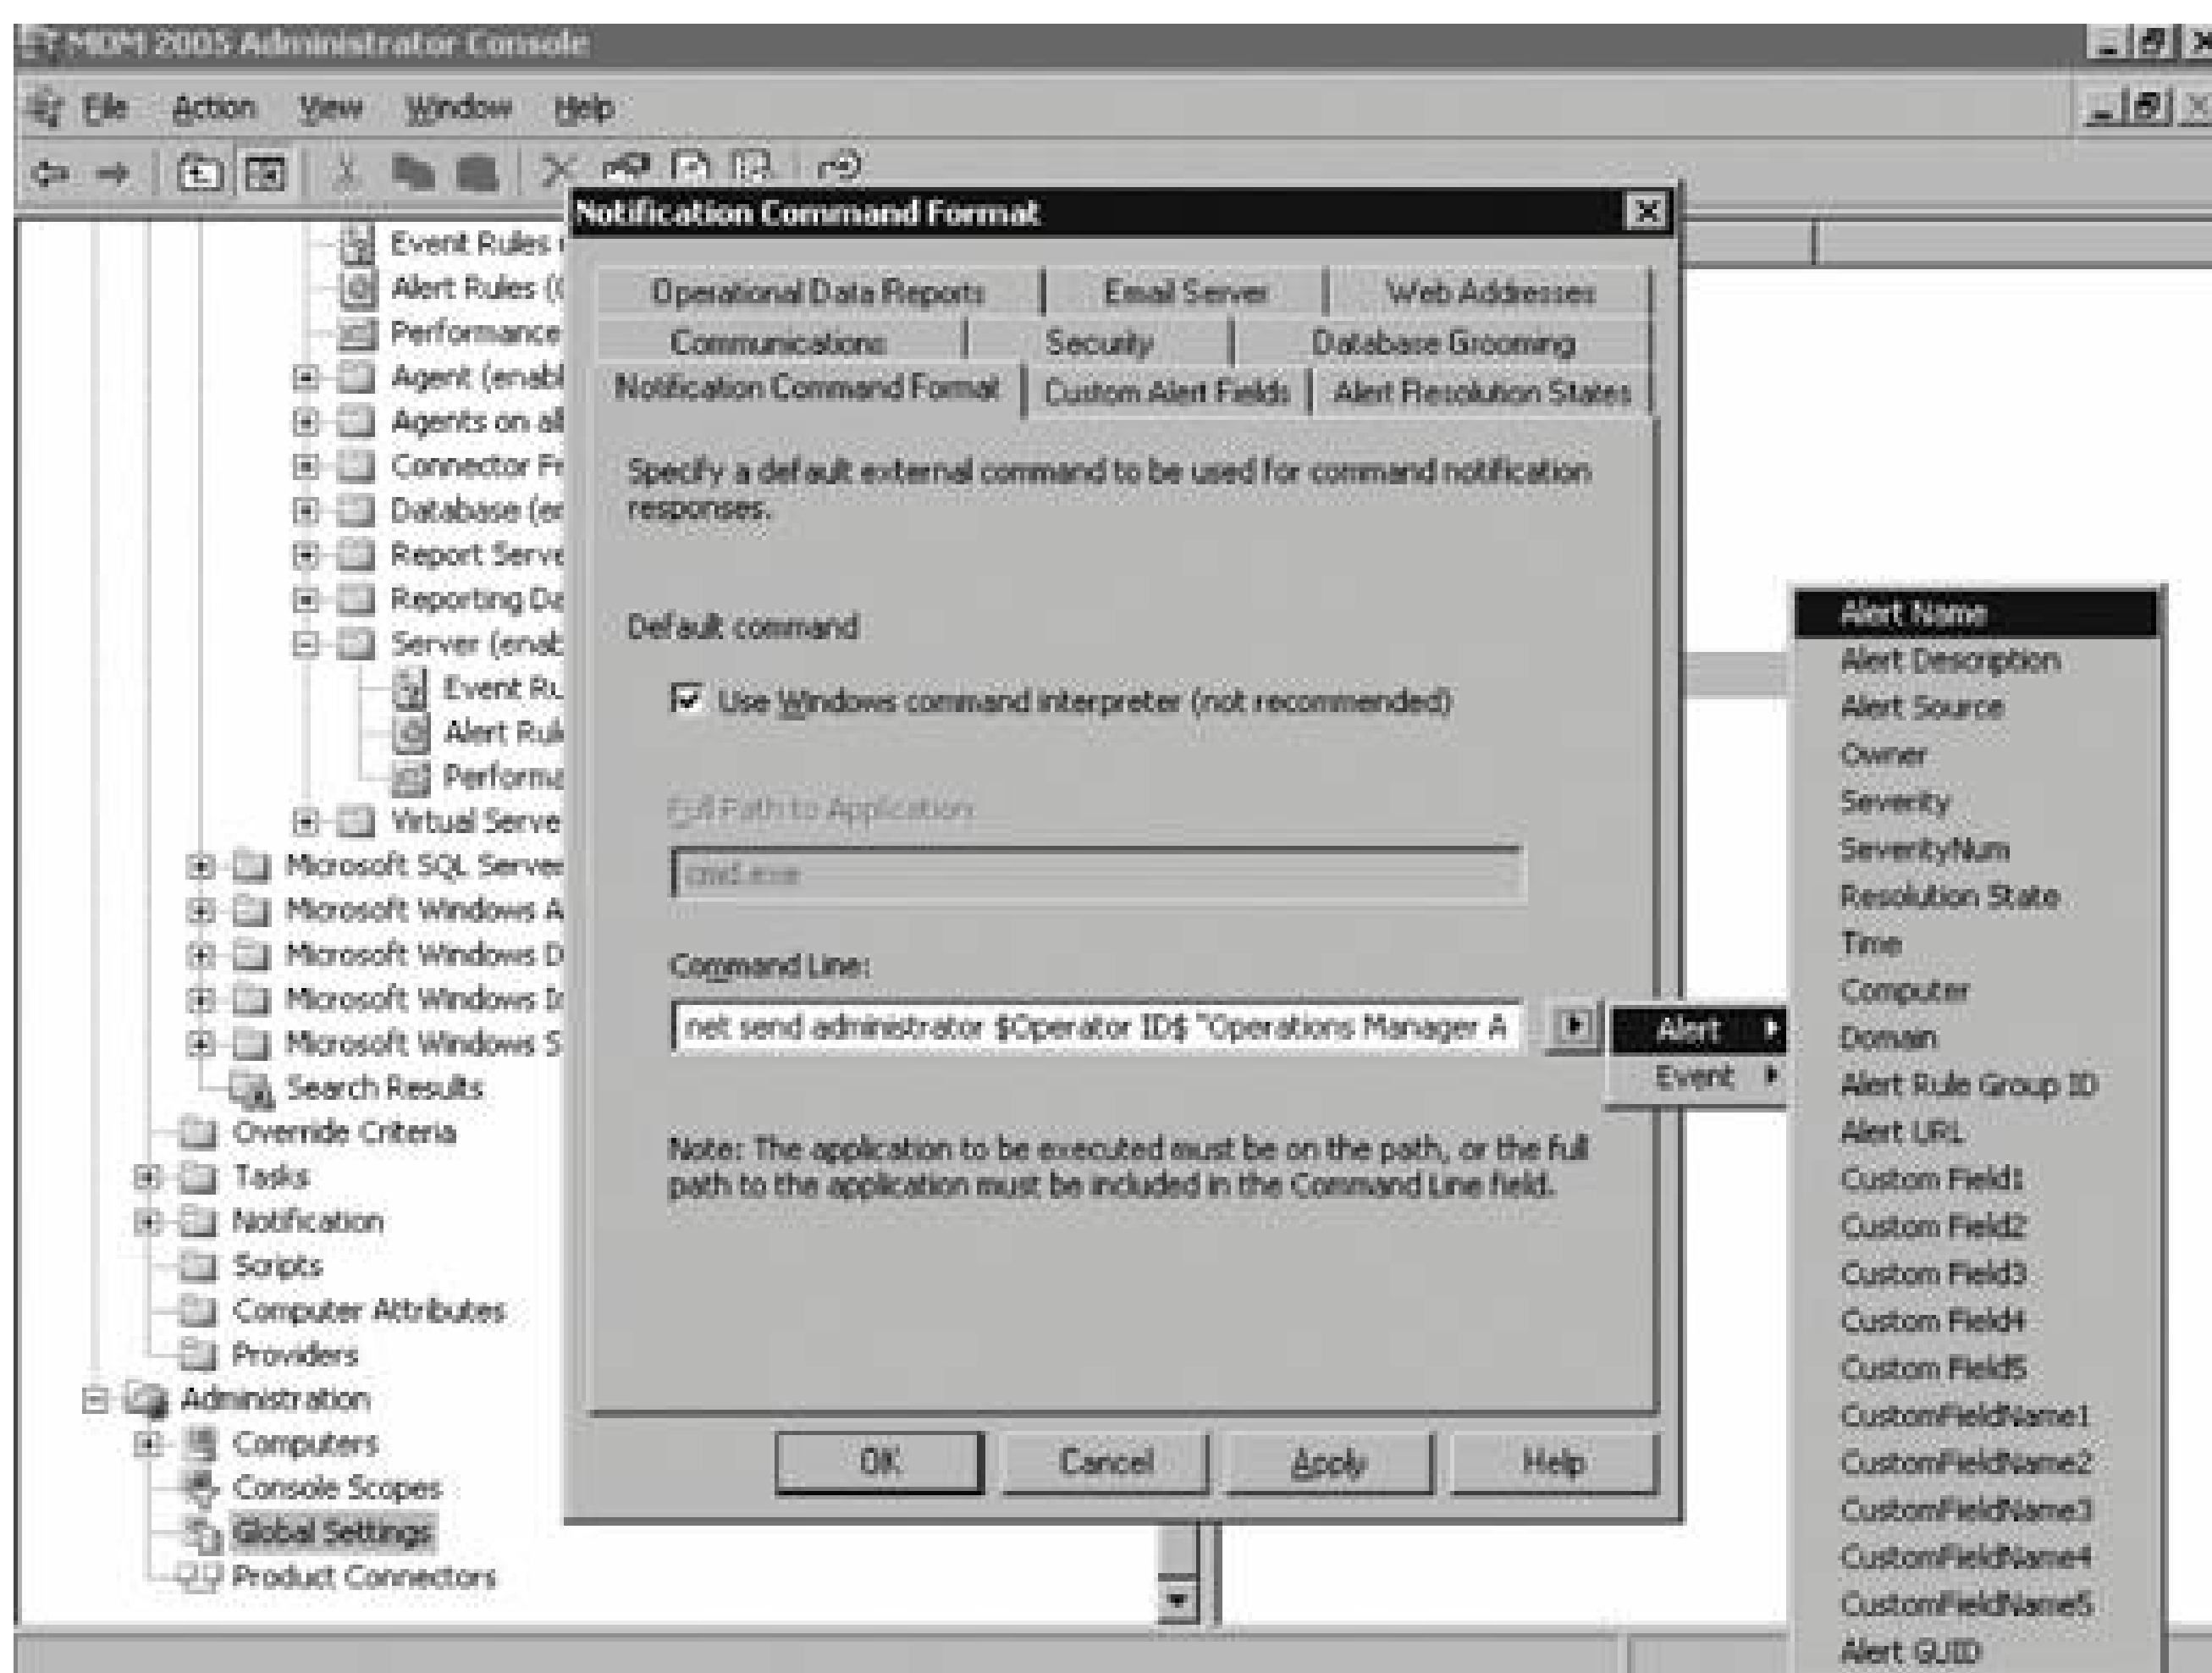
For example, if you want to use `net send`, which depends on the Windows messenger service and Windows Internet Naming Service (WINS) running in your environment, you would select the Windows command interpreter and construct a command line that looks like this:

```
Send $Operator ID$ "Operations Manager Alert on $Source Domain$\$Source Computer$:
$Description$ (view with $Alert URL$)"
```



`net send` is used as an illustration here. There are known exploits that make use of the messenger services so be careful about using this in production. See Microsoft article 330904 for more information (<http://support.microsoft.com/default.aspx?scid=kb;en-us;330904>).

Figure 5-5. Alert fields that are available for inclusion in the command-line parameters



When the alert fires off, it will generate a pop-up message like the one shown in [Figure 5-6](#).

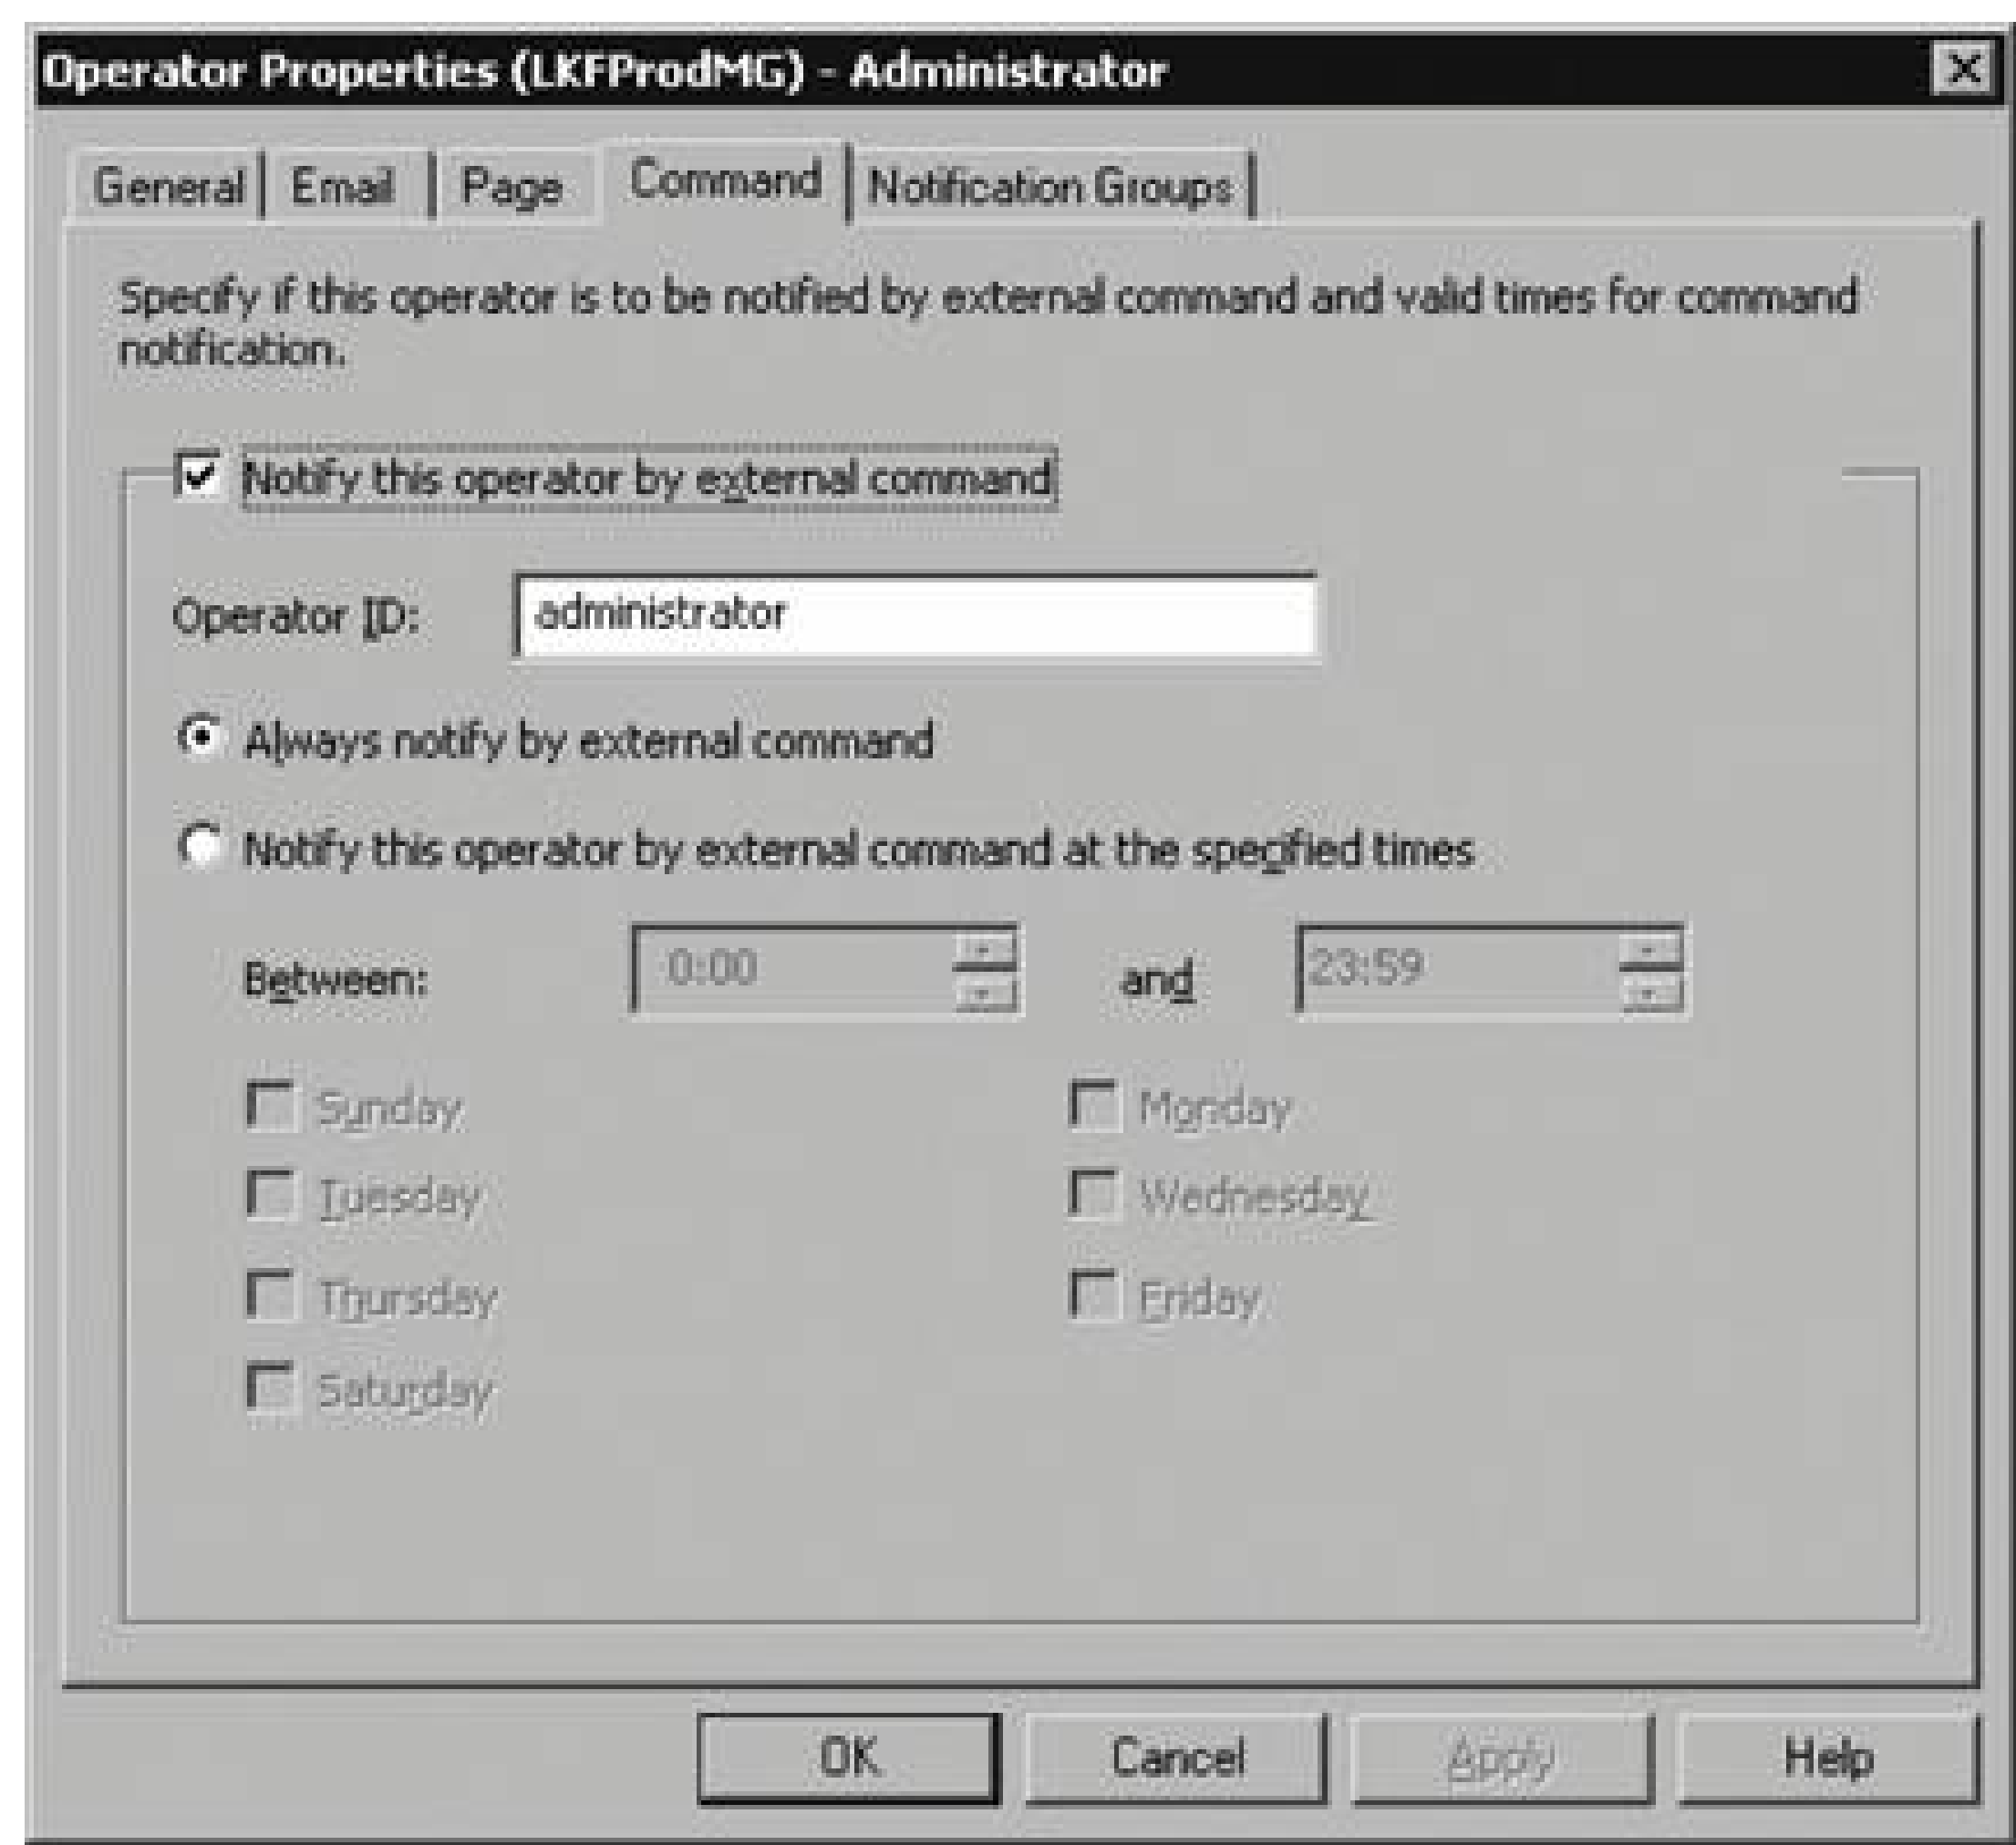
Figure 5-6. Messenger pop-up notification of an Agent heartbeat failure alert

It is a good idea to confirm what will be included in a notification message, no matter what transport you have chosen. The following example uses the MOM Agent heartbeat failure alert by following these steps:

1. Configure the Global Settings Notification Command Format (see [Figure 5-5](#)).

2. Create a notification operator and add it to the Network Administrators notification group. This operator must have the Command option enabled ([Figure 5-7](#)).

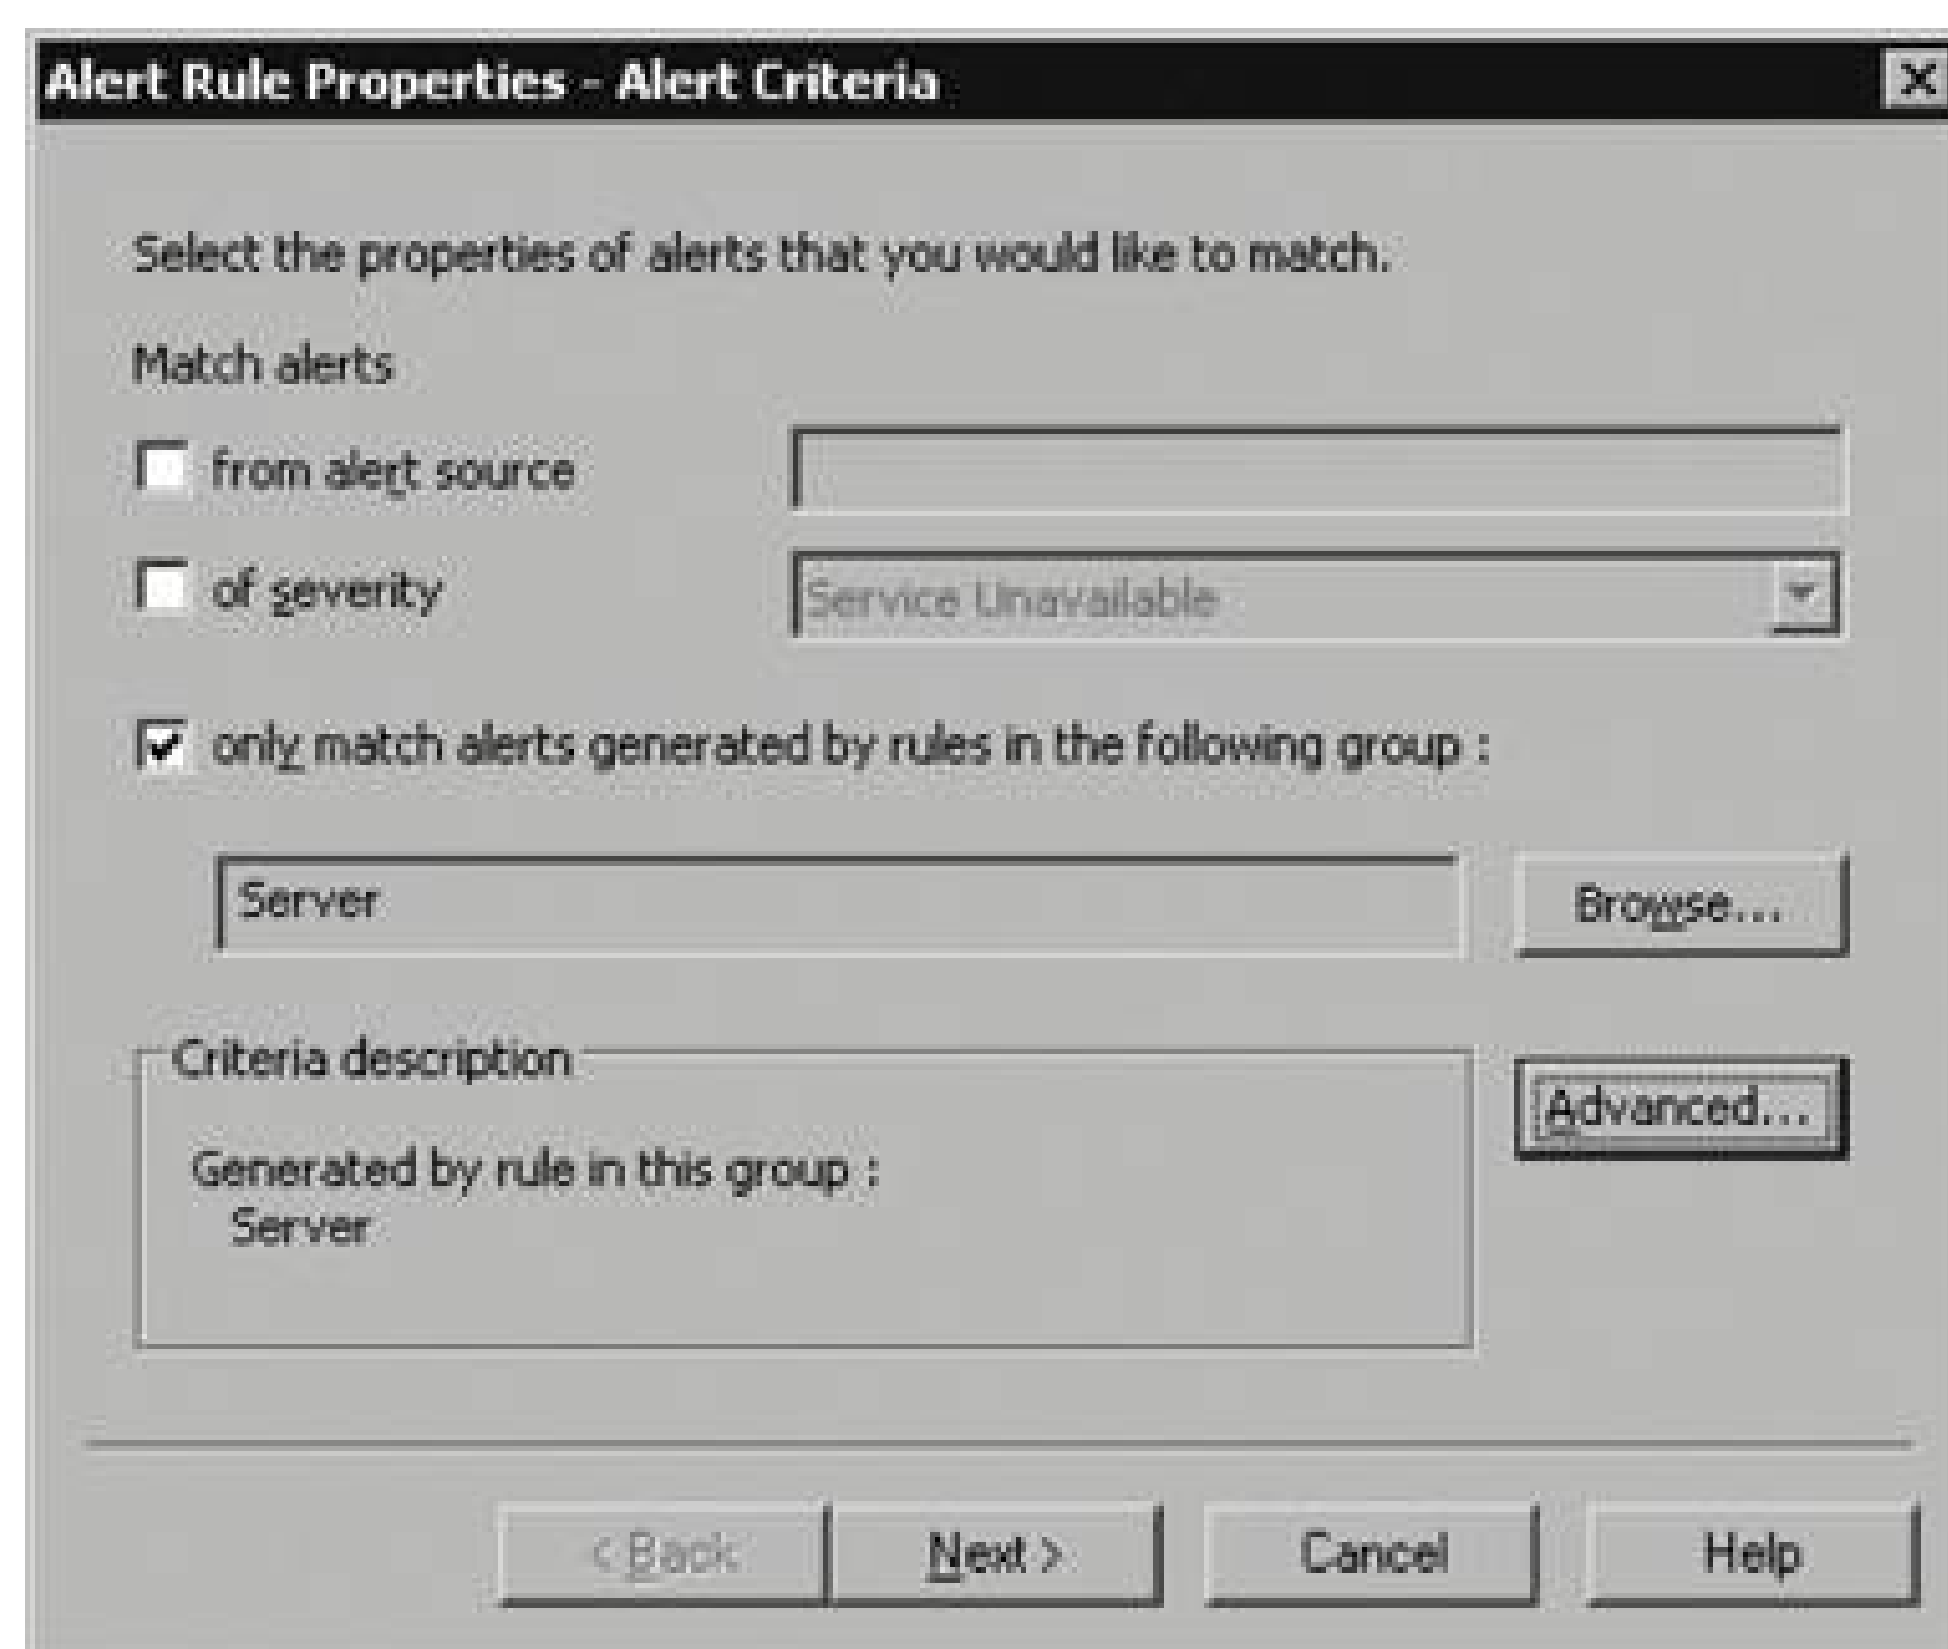
Figure 5-7. Enable the operator to receive notifications via command



3. Navigate to the Microsoft Operations Manager      Operations Manager 2005      Server Alert Rules group. Create an Alert rule (right-click, create rule) with Alert Criteria that only matches alerts generated by the Server rule group (see [Figure 5-8](#)).

Figure 5-8. Configuring the Alert Criteria to match all alerts generated in the Server rule group





4. Click Next and accept the default schedule, which is to always process data. Click Next again to proceed to the Responses configuration page.
5. On the Alert Rule Properties Responses page (see [Figure 5-9](#)), click Add and select the "Send a notification to a Notification Group" option. Then select the Network Administrators notification group from the drop-down list shown in [Figure 5-2](#).

Figure 5-9. Choosing to send an alert to a notification group

Click Next to bring up the Alert Rule Properties - General page, where you give the new alert rule a name ([Figure 5-10](#)). In this example, it is named NotificationTestRule. Click Finish.

To complete testing, log on as the operator account, which was added to the Network Administrators notification group, on any machine on the network (the messenger service must be enabled on both the management server and the computer you are logged onto). Then, on any of the agent-managec servers on the network, stop the MOM service and wait for the pop-up to appear.

## 5.1.2. Email Server

The Email Server tab in the Global Settings is where you indicate the SMTP server that all of the management servers in the management group communicate with when sending notifications. This tab consists of five fields:

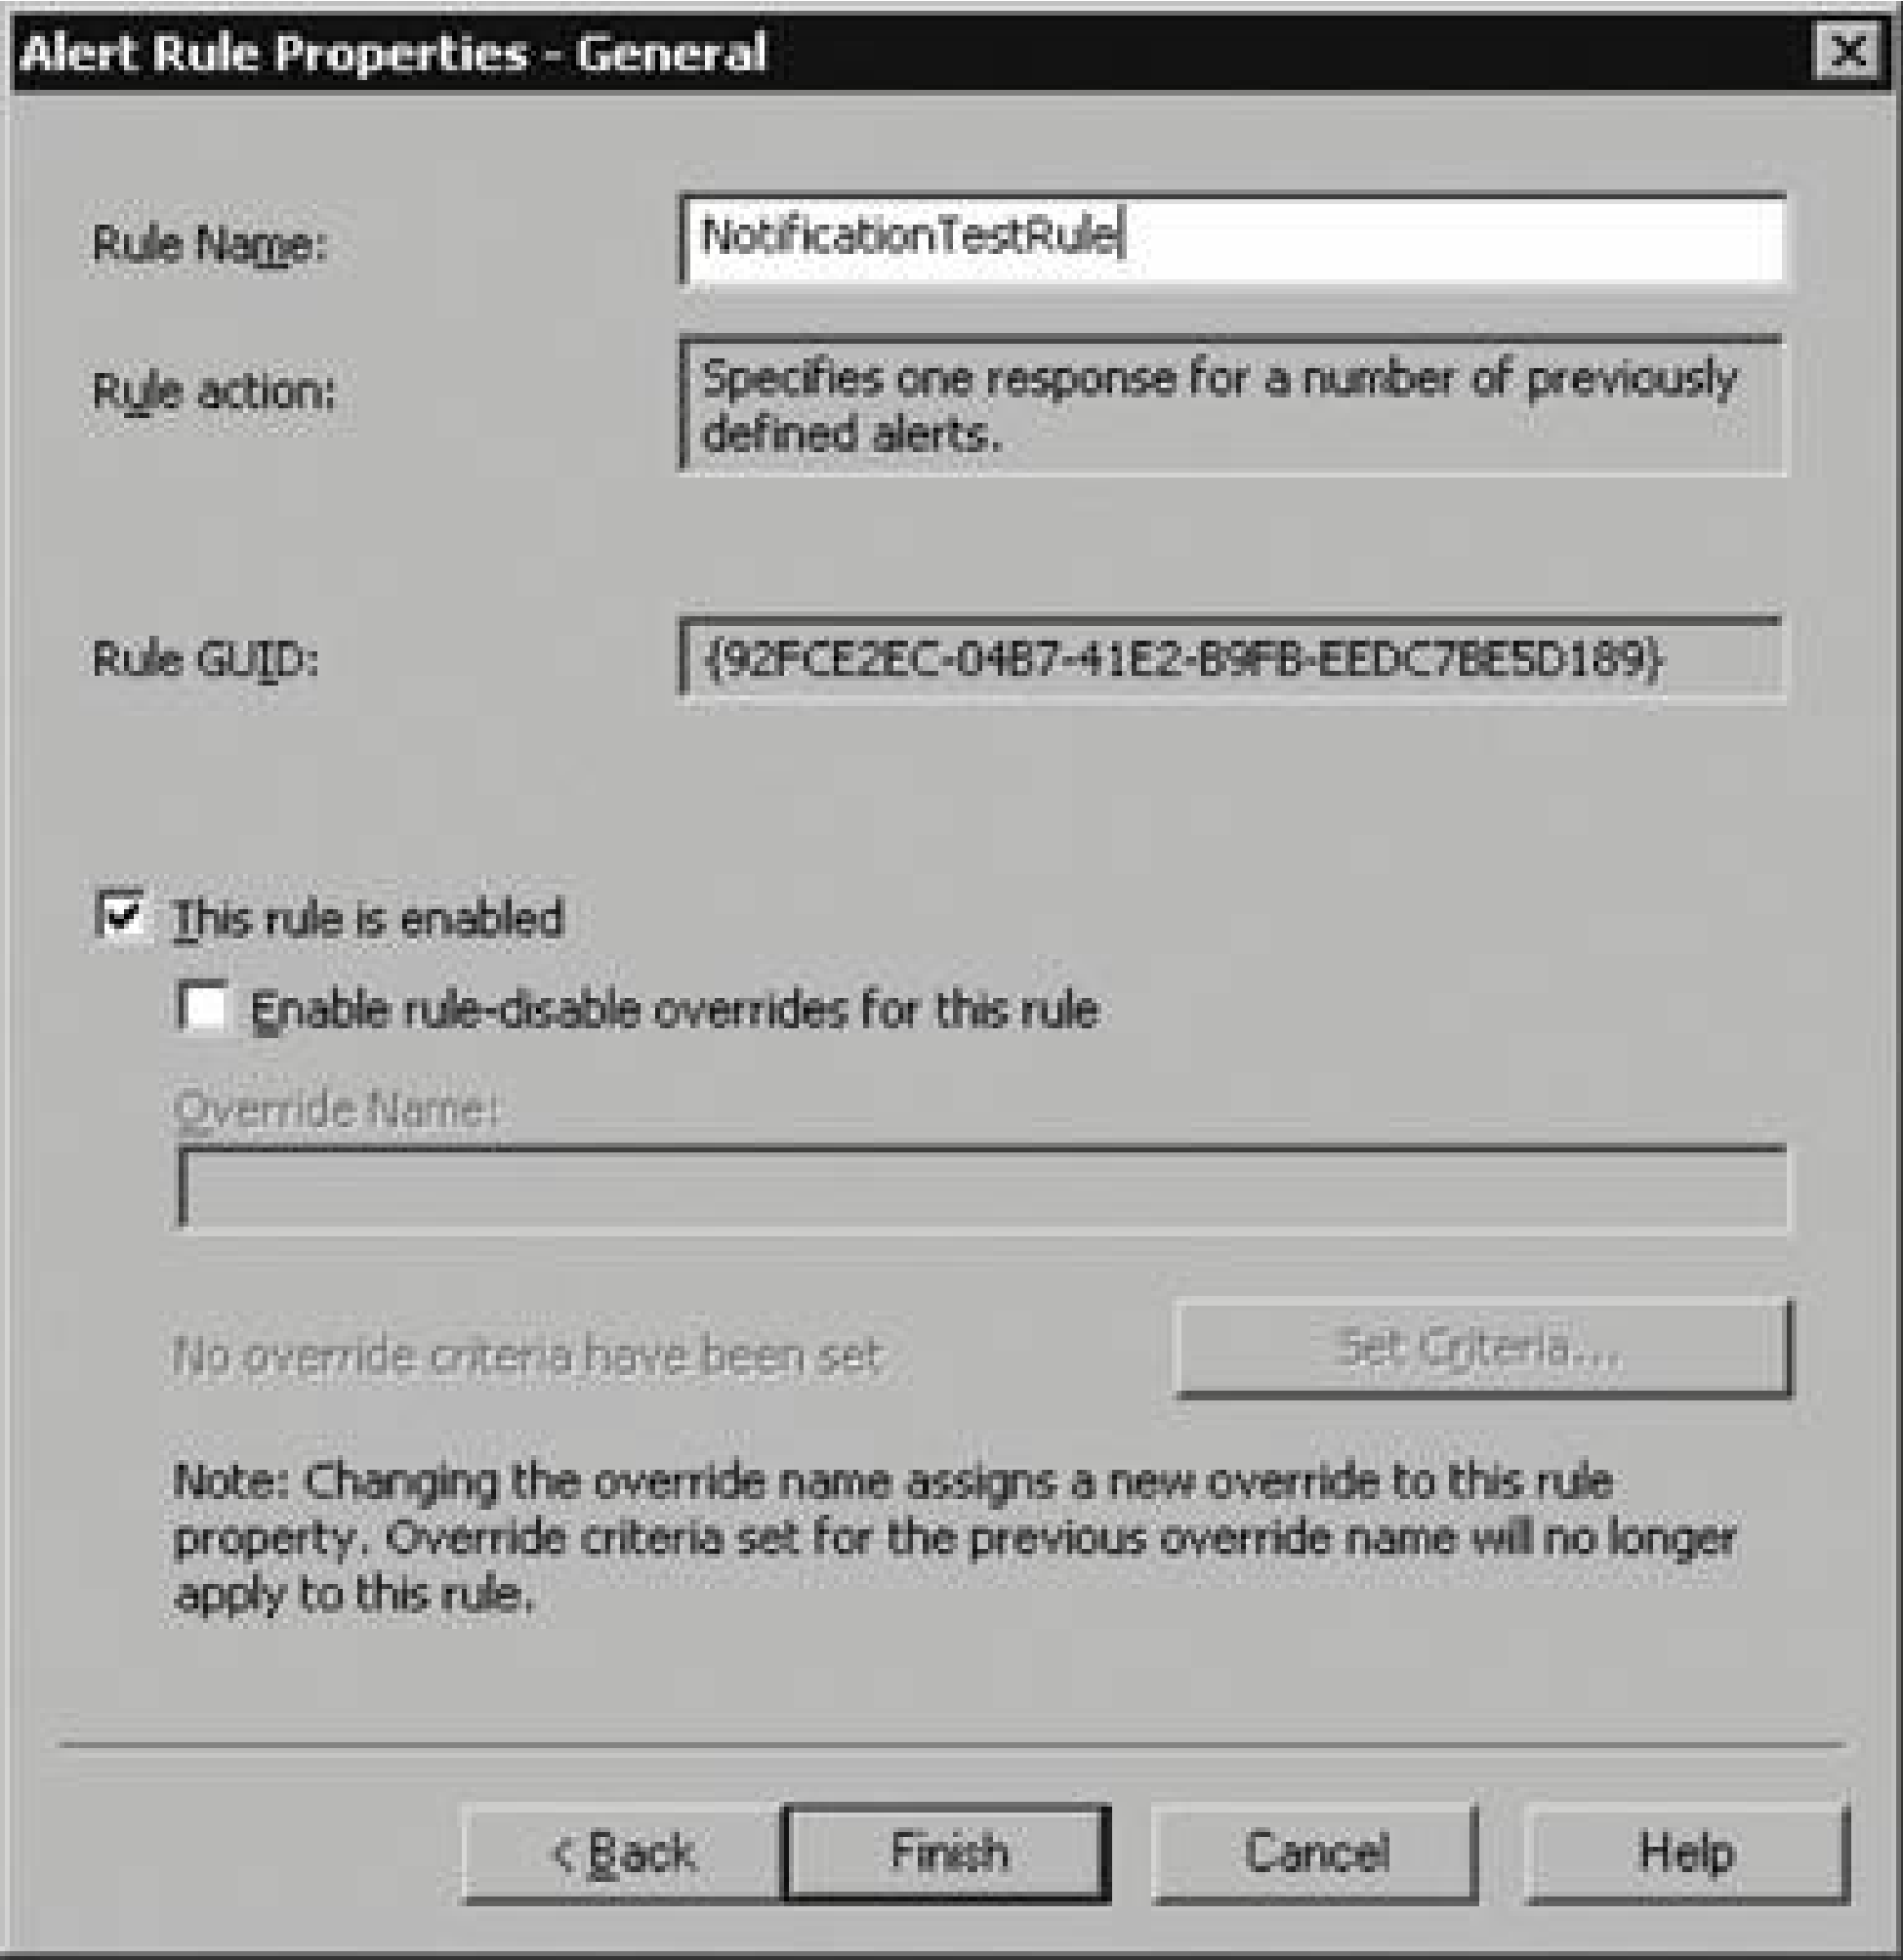
### *Transport*

This is read-only and is set to SMTP.

### *Server Name*

This can be either the FQDN or the NetBIOS name of the SMTP server.

Figure 5-10. Naming a newly created alert rule



*Return Address*

All SMTP communication sessions require that a properly formatted sender address be passed from the sender to the SMTP server. This is used as the From address on the email.

*Character Set*

This defines which character set to use for the email notifications. Keep the default of Unicode Transformation Format 8-bit (UTF-8).

*SMTP port*

The SMTP protocol defaults to TCP port 25. Unless you have a very specific reason to change this, I recommend accepting the default.

**5.1.3. Alert Resolution States**

Alert resolution states are covered in the "[The Life of a MOM 2005 Alert](#)" section in [Chapter 1](#). But to refresh, all alerts exist in one of several resolution states:

- New



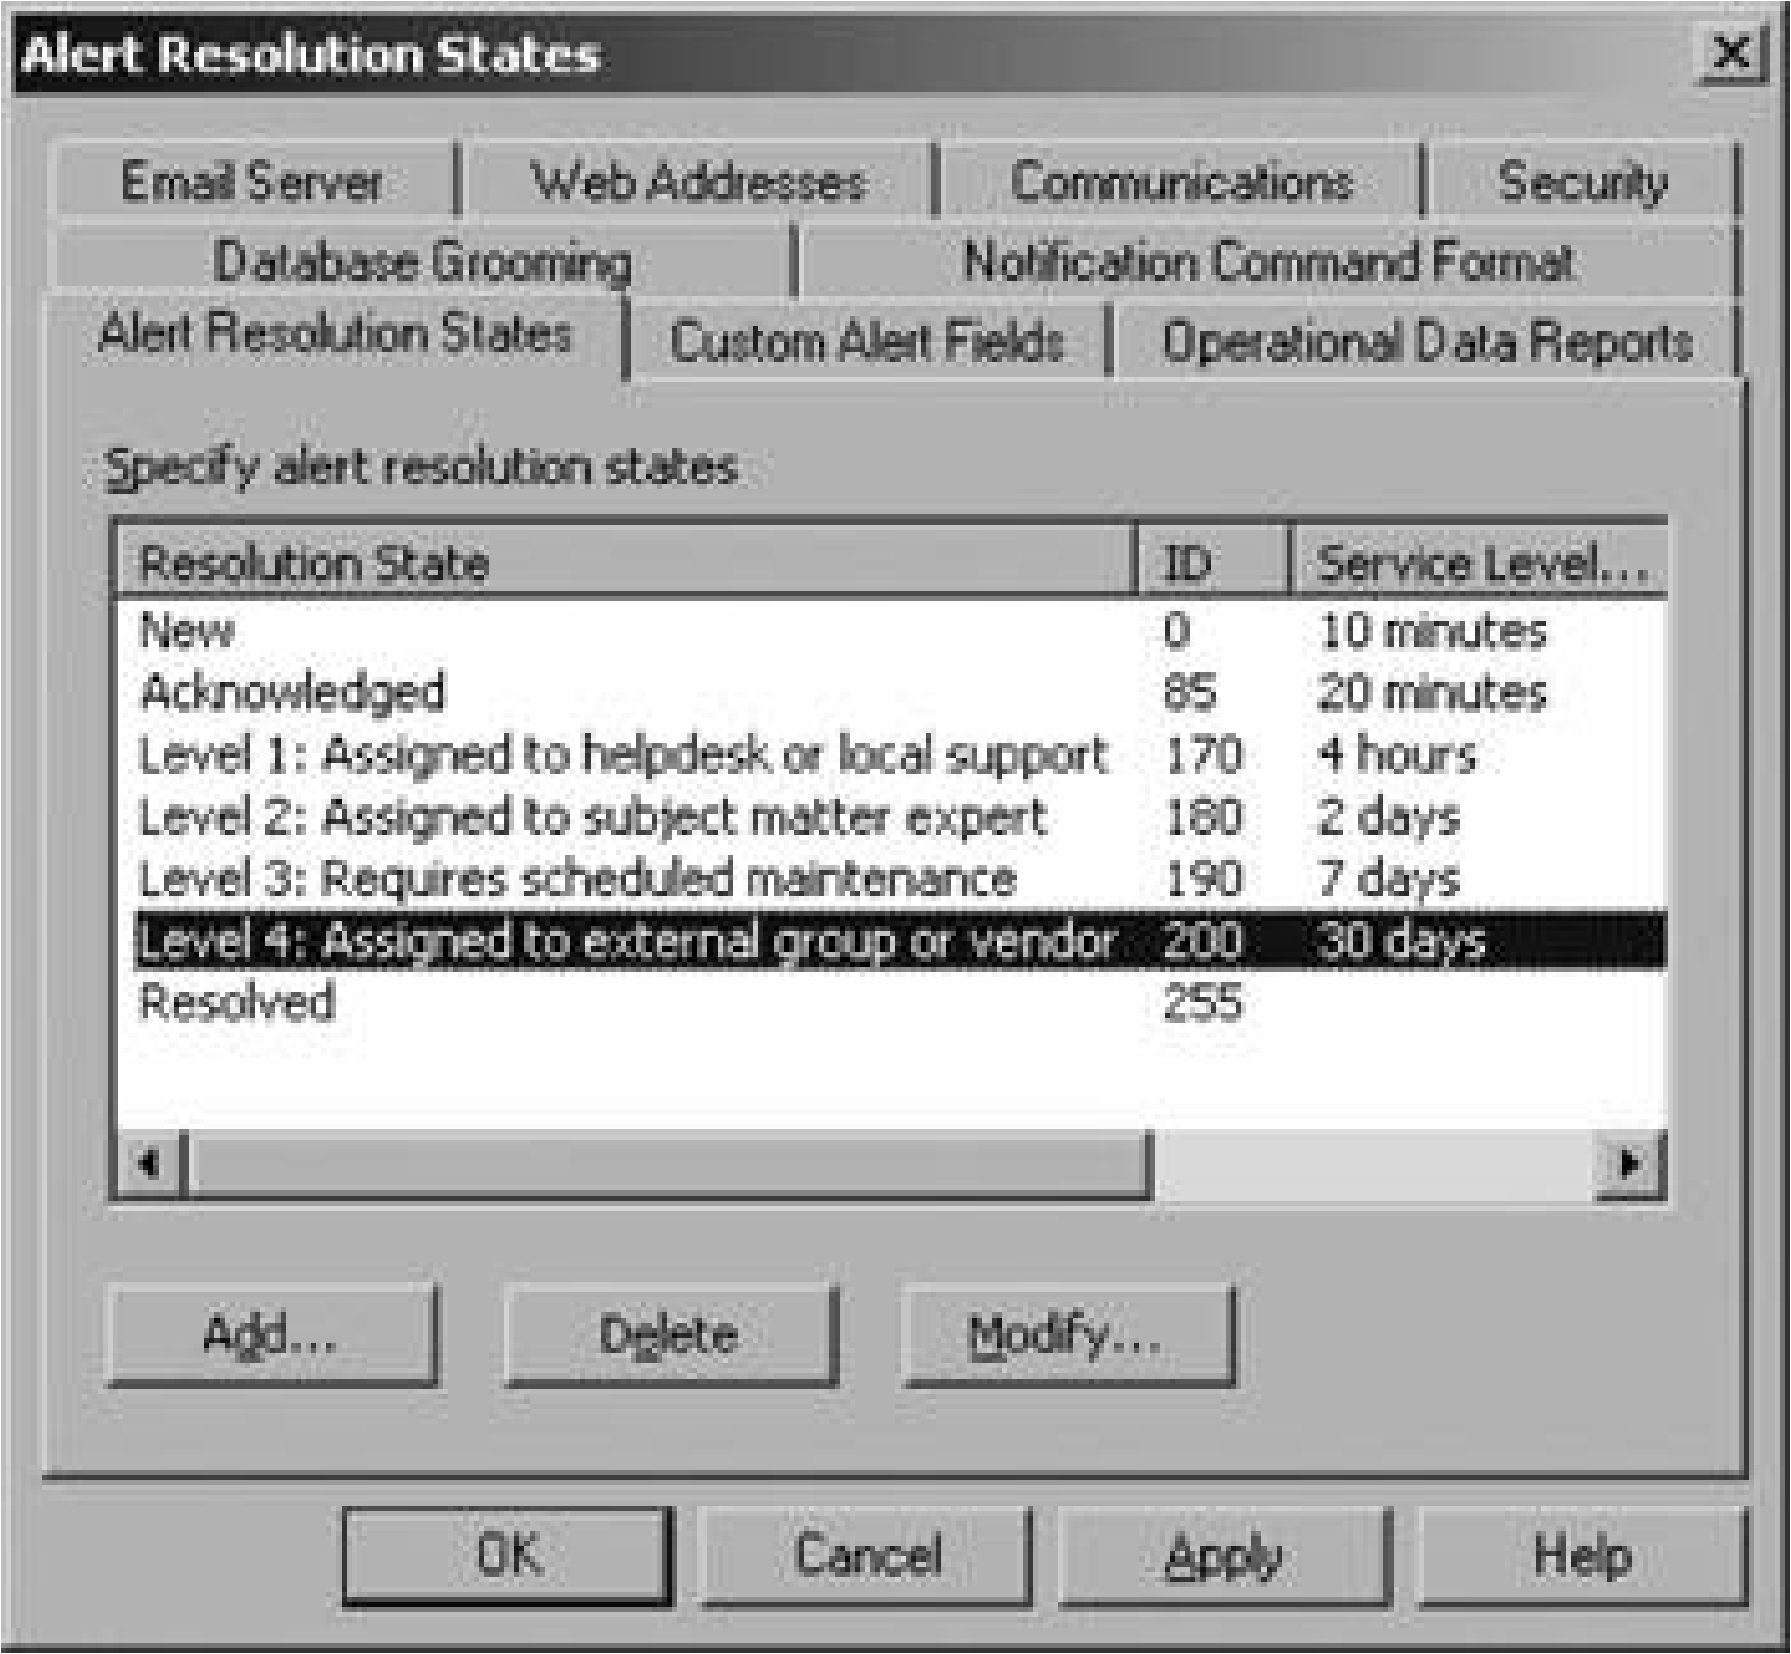
- Acknowledged
- Level 1: Assigned to help desk
- Level 2: Assigned to subject matter expert
- Level 3: Requires scheduled maintenance
- Level 4: Assigned to external group or vendor
- Resolved

By default, when an alert first appears in the Operator console, its resolution state is always set to New.

MOM times how long an alert is in any resolution state. When that time exceeds the value configured in the service-level agreement for that resolution state, the alert will then appear in the service-level Exceptions Alert view in the Operator console. This allows a company to track their ability to meet service-level agreements.

You can modify the default settings, shown in [Figure 5-11](#), and even delete them (except for the New and Resolved states). The most common modification made is to change the service-level agreement time period to suit your environment. The resolution states themselves are useful only if your support structure breaks down into these four levels. If your environment alerts are only sent to the Windows administrators, who then open a support case with an external vendor, then there is no need for the alert resolution states. You can delete the unnecessary alerts, or leave them and not use them, at your discretion. In addition, you can create new alert resolution states that suit your environment.

Figure 5-11. Default values for alert resolution states



When you select an alert state from the list and click Modify, you'll see the Edit Alert State dialog box shown in [Figure 5-12](#). An alert resolution state contains five fields:

*Name*

The purpose of the Name field is obvious, but keep in mind that the more descriptive it is, the more useful it is.

Figure 5-12. Attributes of an alert resolution state

✕

Edit Resolution State

Specify the properties and conditions of service level agreement for this resolution state.

Resolution state

Name:

ID:

Service level agreement

The amount of time Alerts can remain in this resolution state:

If alert remains in this resolution state longer than specified period it is marked as a service level exception.

Shortcut keystroke

Prefix:

Character: (A-Z, 0-9)

Display in console menus

☒ Users may set this alert resolution state in the Operator Console and Web Console.

Help

OK

Cancel

*ID*

Internally to MOM, all alert resolution state objects are actually identified by their ID number. This number must be between 0 and 255. The 0 ID is reserved for the New resolution state, and ID 255 is reserved for the Resolved resolution state. All other values (1 to 254) are available for use.

*Service level agreement*

This field can be set to any minute, hour, or day.

*Shortcut keystroke*

To help ease administration of an alert in the Operator console, you can assign a keystroke combination to an alert resolution state. For example, you can assign Ctrl-Alt-I to the Acknowledged state. After restarting the Operator console you can select an alert, press Ctrl-Alt-I, and its resolution state will be changed to Acknowledged.

*Display in console menus*

You can control whether operators use an alert resolution state in the Web or Operator

downloaded from: [lib.ommolkesab.ir](http://lib.ommolkesab.ir)



consoles. This is useful for restricting access to custom created states. If this checkbox is cleared, this particular resolution state will not be available in the Operator or Web consoles. You can still configure event and performance rules to create alerts with the unavailable state as the default state. The alerts will then have this state when they appear in the Operator console, but it would not appear as a resolution state option for that alert. However, you can still access and make use of this alert resolution state programmatically.

For example, to build a custom alert view in the My Views node in the Operator console, you could build a filter using this non-visible alert resolution state value to populate the view. This is particularly useful for alerts that you want to see in the Operator console, but do not require action. For example the alert generated by the Test End to End Monitoring task in [Chapter 3](#) confirmed that the round-trip communication from the management server to the agent was working. The following is how a Custom Alert view is built:

1. Create a new alert resolution state by clicking the Add button, as shown in [Figure 5-11](#).

The Edit Resolution State page will appear, as shown in [Figure 5-13](#).

Figure 5-13. Creating a new alert resolution state

2. Create a new resolution state, for example "Test alert - Ignore" with an ID value of 86.
3. Clear the "Users may set this alert resolution state in the Operator Console and Web Console" checkbox.
4. Set this custom resolution state as the default for the MOM End to End Monitoring event rule.
5. Navigate to the Rule Groups folder in the Administrator console and from there to Microsoft Operations Manager → Operations Manager 2005 → Agents on All MOM Roles rule group.
6. Open the properties of the MOM End to End Monitoring event rule.
7. Assign this new alert resolution state as the default state on the Alert tab, as shown in [Figure 5-14](#).

Whenever the Test End to End Monitoring task is fired off, if a success alert is generated it will have "Test alert - Ignore" as its starting resolution state.

Figure 5-14. Set the event rule to generate an alert with the newly created resolution state

8. Launch the Test End to End Monitoring task against an agent-managed computer.

The alert is generated and surfaces in the default Alerts view, as shown in [Figure 5-15](#).

- 9. Create a custom view in the My Views node that will filter all alerts that have the "Test alert - Ignore" state.
- 10. Navigate to the All My Views node in the Operator console, open the context menu, and select Create a New Alerts View (see [Figure 5-16](#)).

Figure 5-15. An alert with the custom "Test alert - ignore" state

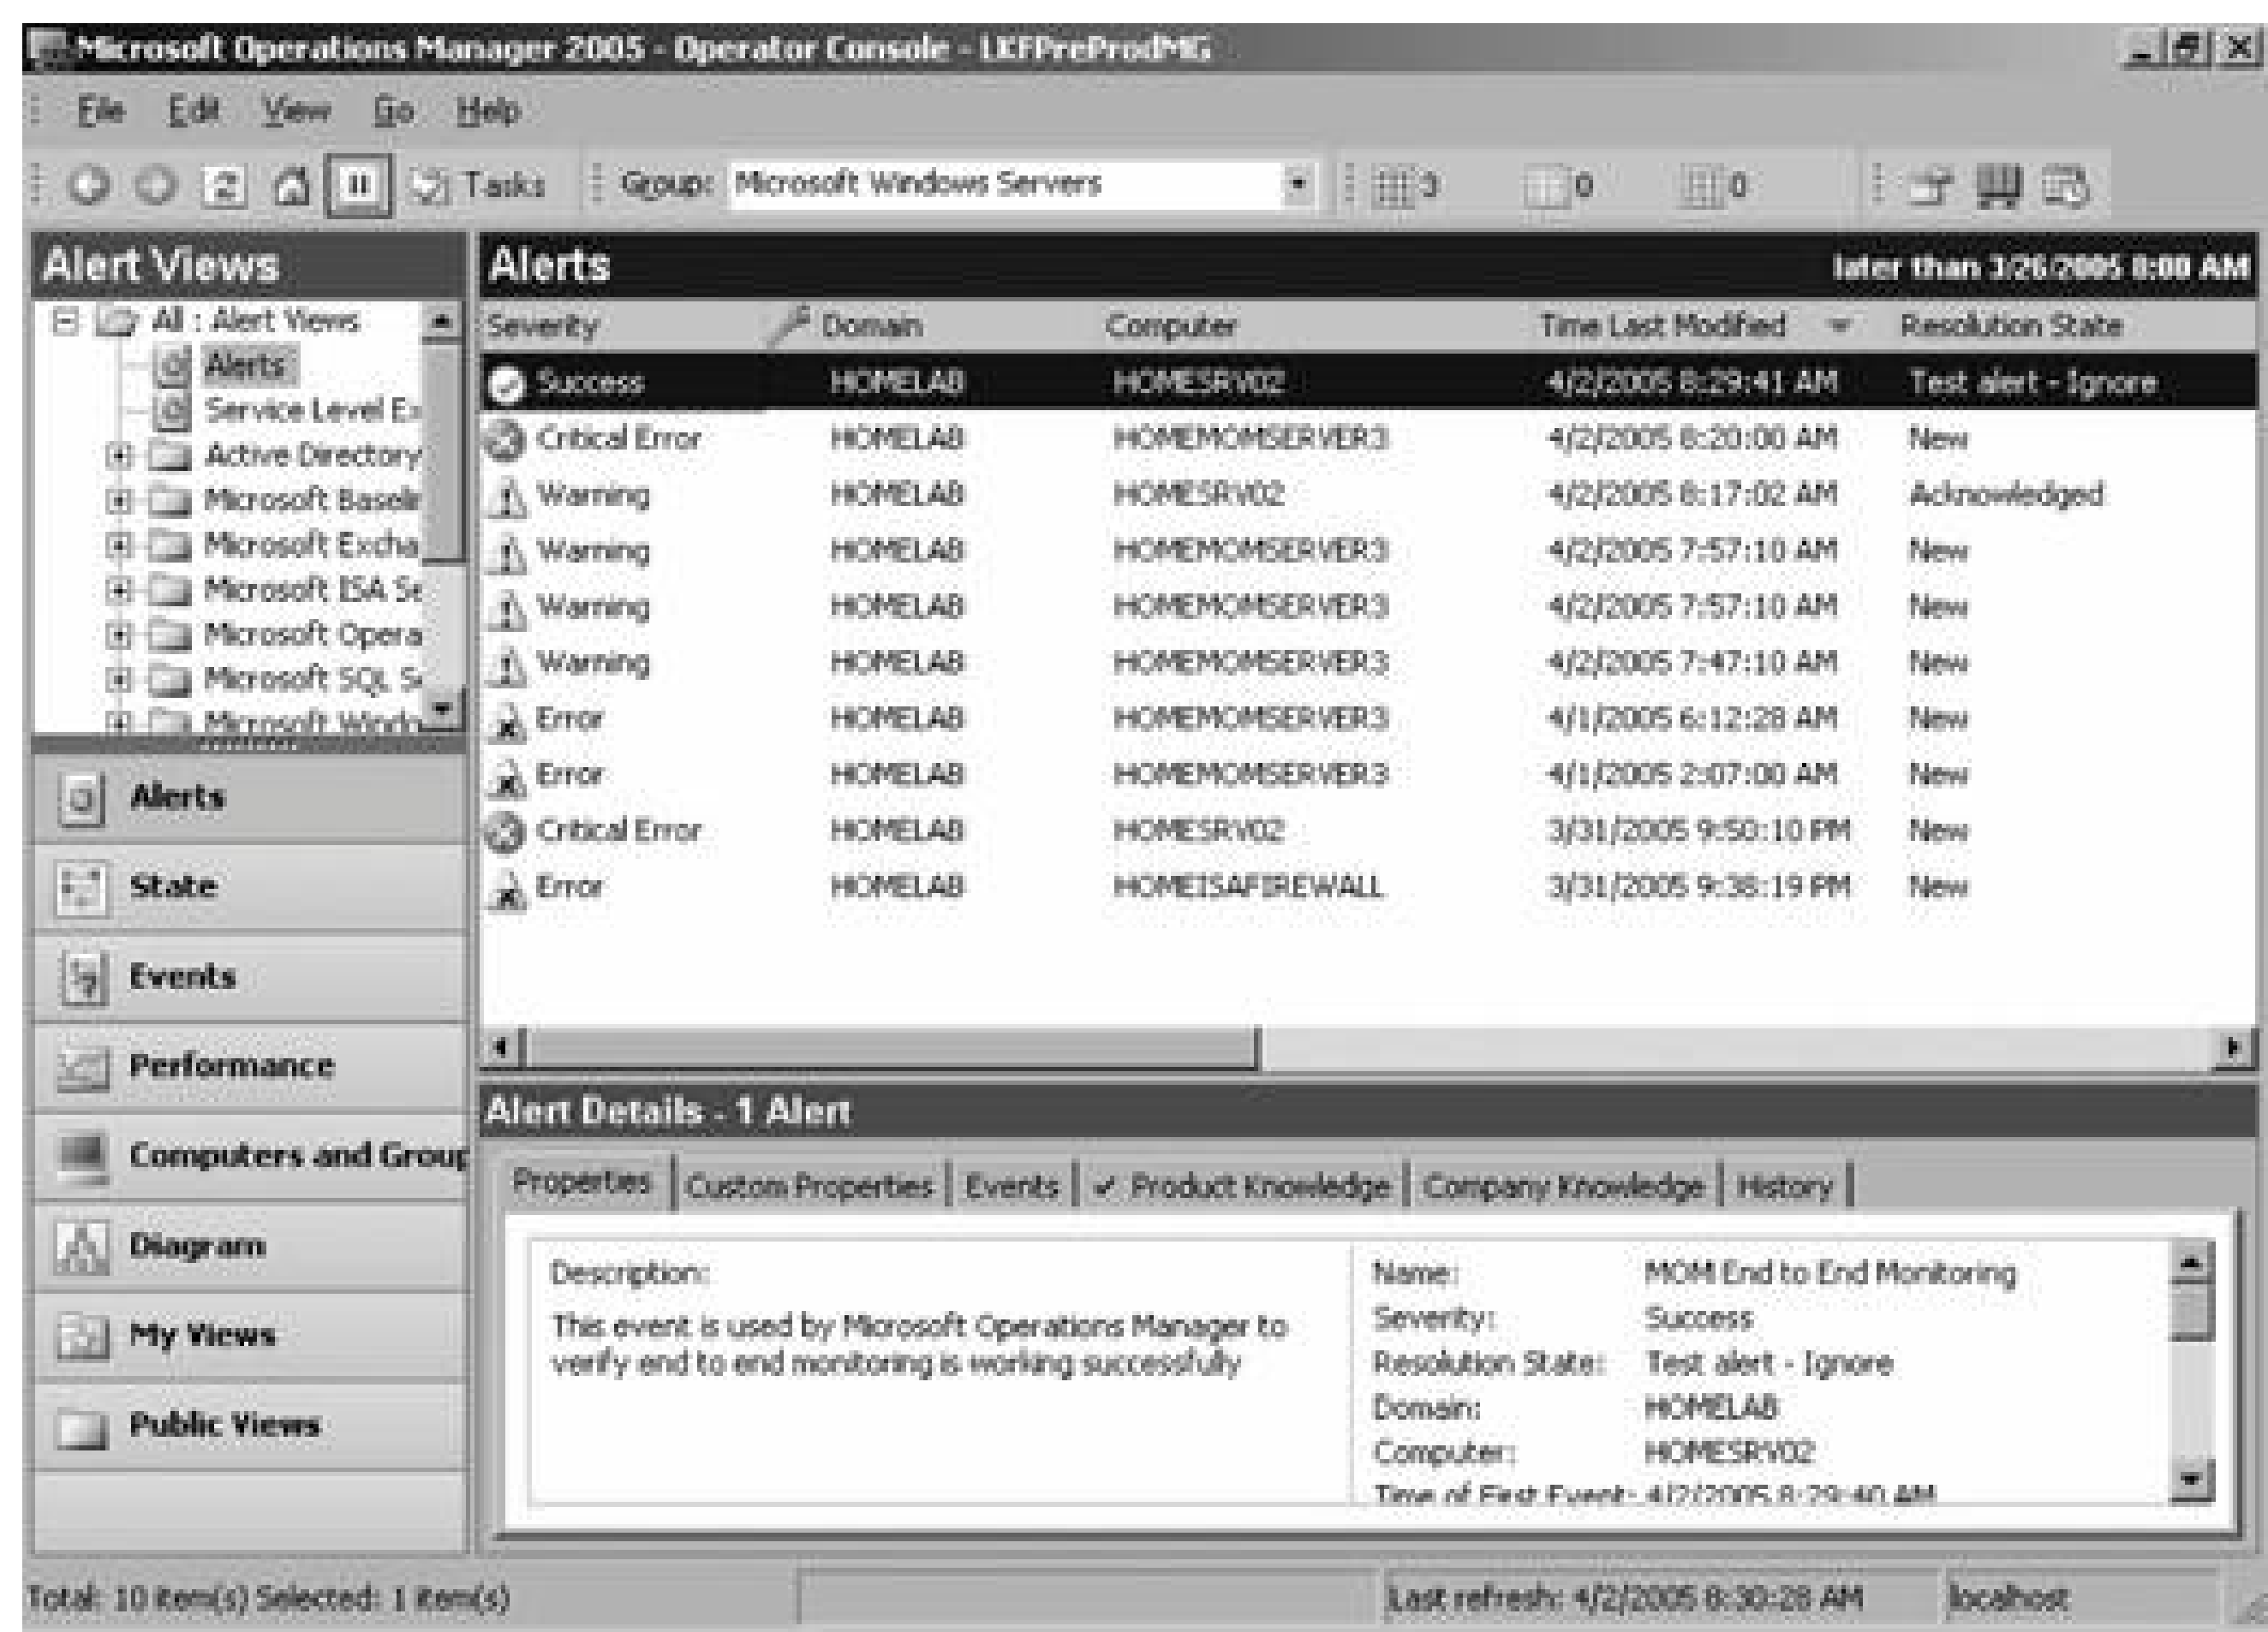



Figure 5-16. Creating a new Alerts view in the Operator console





This opens the "Which type of Alert view do you want to create?" page.

11. Select the "Alerts that satisfy specified criteria" option (see [Figure 5-17](#)).

Figure 5-17. Step 1 of creating a custom alert view



12. Proceed through the filter creation process. Click Next and select the "with specified resolution state" checkbox (see [Figure 5-18](#)).

Figure 5-18. Selecting to filter on a specified resolution state

13. Click on the underlined value to select the "Test alert - Ignore" value for the resolution state in the "View description" pane.

14. Select the "Test alert - Ignore" resolution state, as shown in [Figure 5-19](#).

Figure 5-19. Choosing the custom resolution state for the filter




15. Click OK, name the view, and finish the process.

The view is then built in the Operator console and is populated with the alert that that meets the criteria, as you can see in [Figure 5-20](#).

Figure 5-20. Custom view filtering for the "Test alert - ignore" resolution state



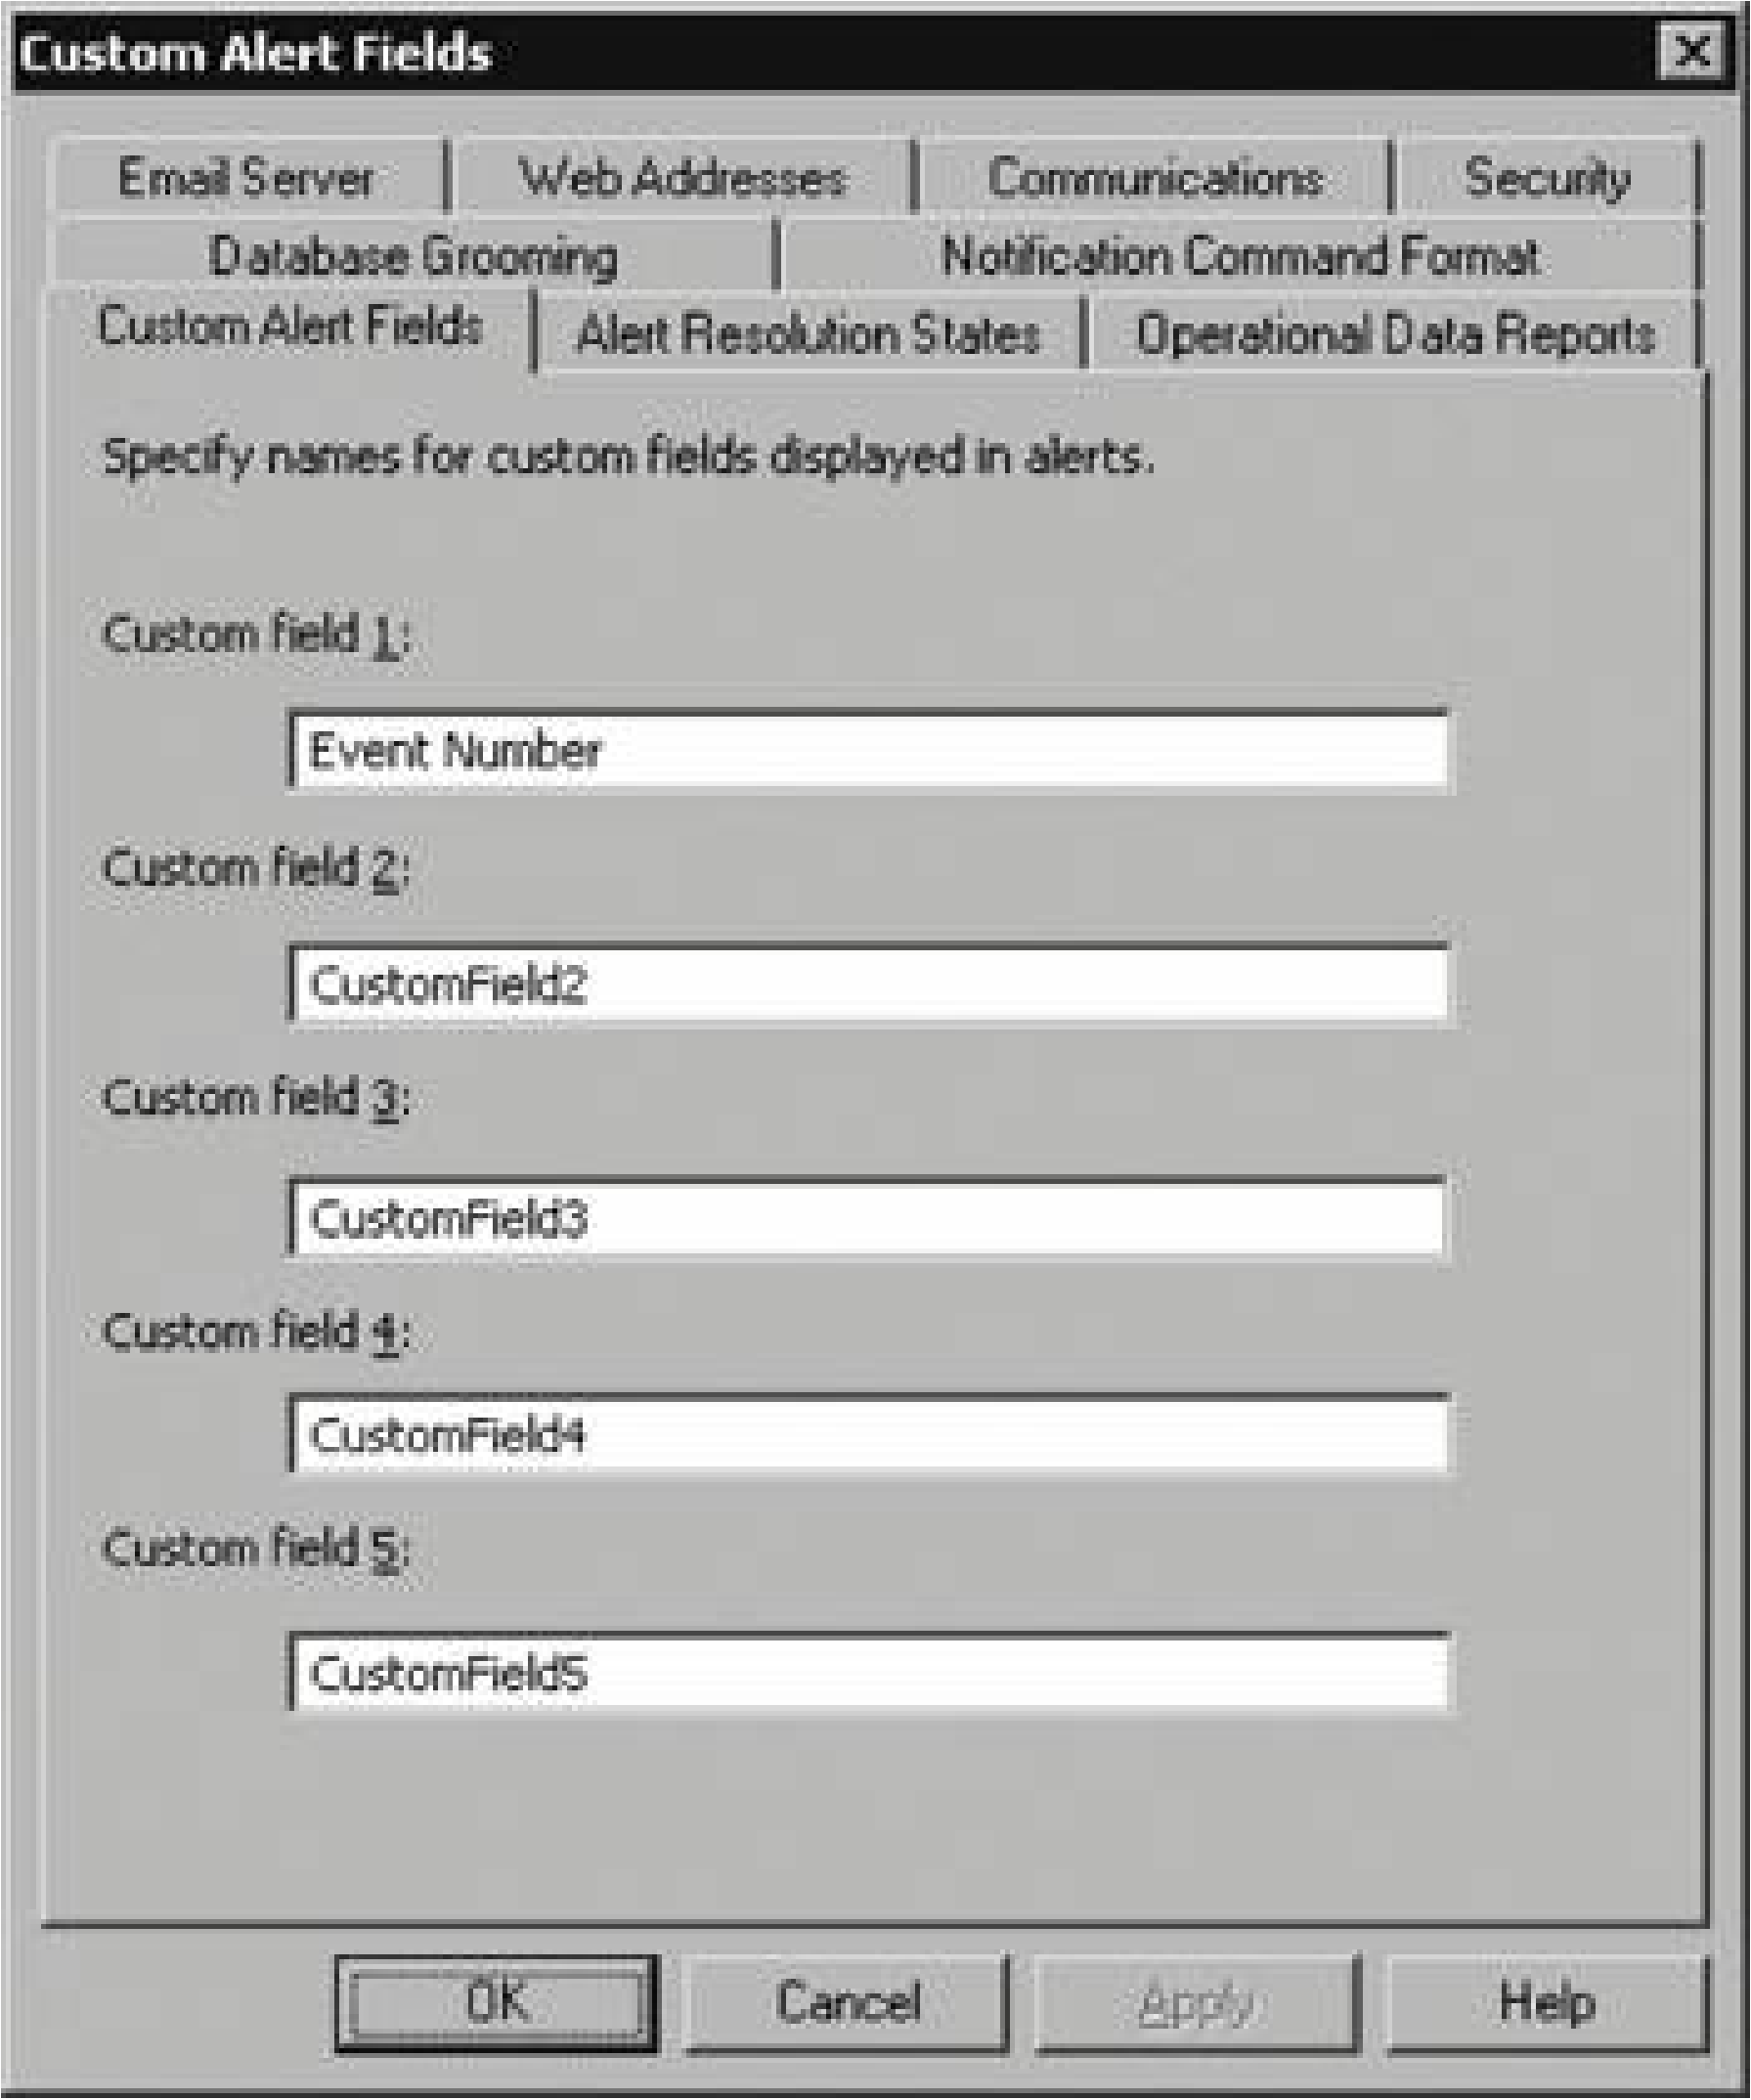


### 5.1.4. Custom Alert Fields

Each alert has five custom fields that enable you to modify the information embedded in an alert when it is created. At the global level, you can provide names for these fields. At the event and performance rule levels, you can populate the content of the named field much like the Notification Command Format command-line string is populated with text and variables.

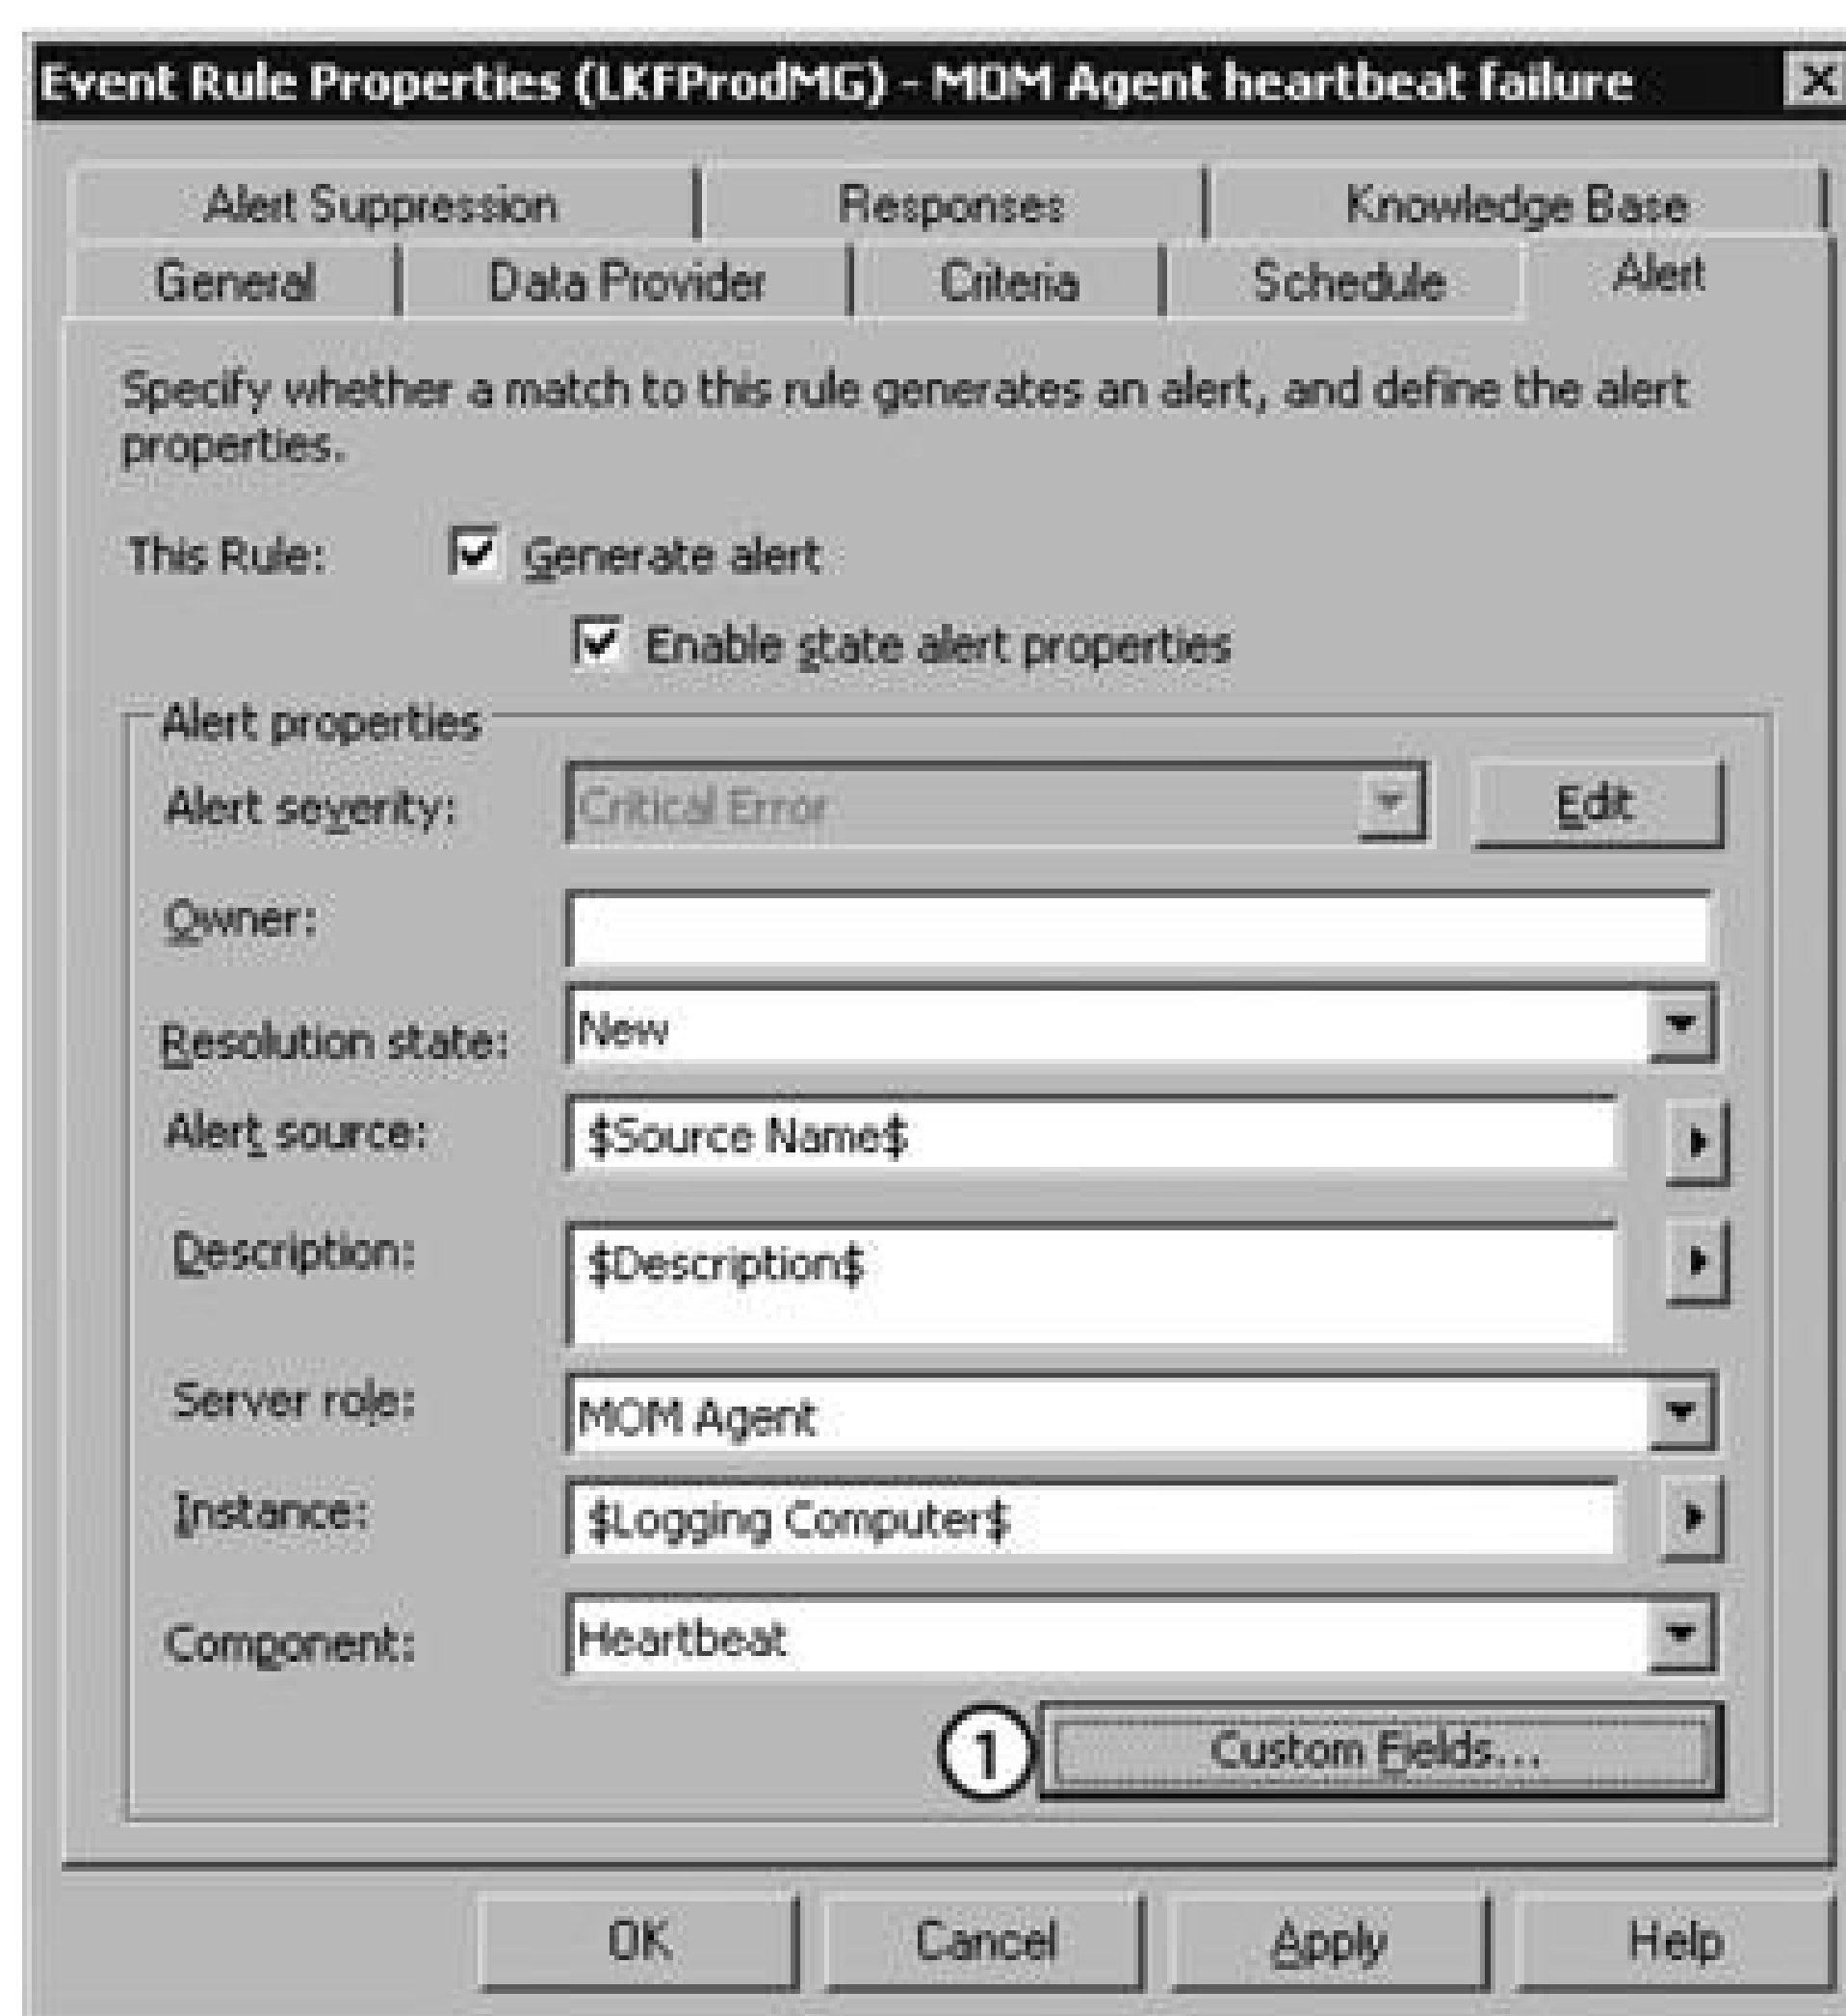
Out-of-the-box, a generated alert only shows about half of the available fields in an alert object. One of the fields that is not present by default is the Event Number. Ordinarily, to find out which event is associated with an alert, you need to navigate to the Events tab in the Details pane. To find the details of that event, such as the event number, open the associated event. An easier way to include this data in a generated alert is to name one of the custom alert fields with an appropriate name, like Event Number (see [Figure 5-21](#)), at the global level. Apply the change and then commit the configuration change on the management packs object in the Administrator console.

Figure 5-21. Renaming Custom field 1



To have this value included as a field in an alert, open the event or performance rule and click the Custom Fields button as shown in point 1 in [Figure 5-22](#).

Figure 5-22. Adding values to the custom fields of an event



This brings up the Custom Fields page, where the value of that custom field is set to the \$Event Number\$, as shown in [Figure 5-23](#). The \$Event Number\$ value is one of many values available, including Full Event Number, Event type, Message DLL, Source Name, Provider Name, Provider type, Description, and so on. For a complete listing of fields that are available for alerts, refer to the MOM SDK. The full list is available when you click on the right-facing arrow button pictured in [Figure 5-23](#) to the right of the \$Event Number\$ text box.

Once the configuration changes are set, all future alerts generated by that rule will include the \$Event Number\$ value. You can see this in the Operator console in the Alert Details pane for the alert.

One last note about custom fields: you are not restricted to using the MOM 2005 alert variables, you can add plain text as well. In [Figure 5-23](#), the value could have been constructed to read "The event number is \$Event Number\$." When the MOM Agent heartbeat alert is generated, the text reads "The event number is 21284," as shown in point 1 in [Figure 5-24](#).

Figure 5-23. Adding the event number parameter to CustomField1, renamed Event Number



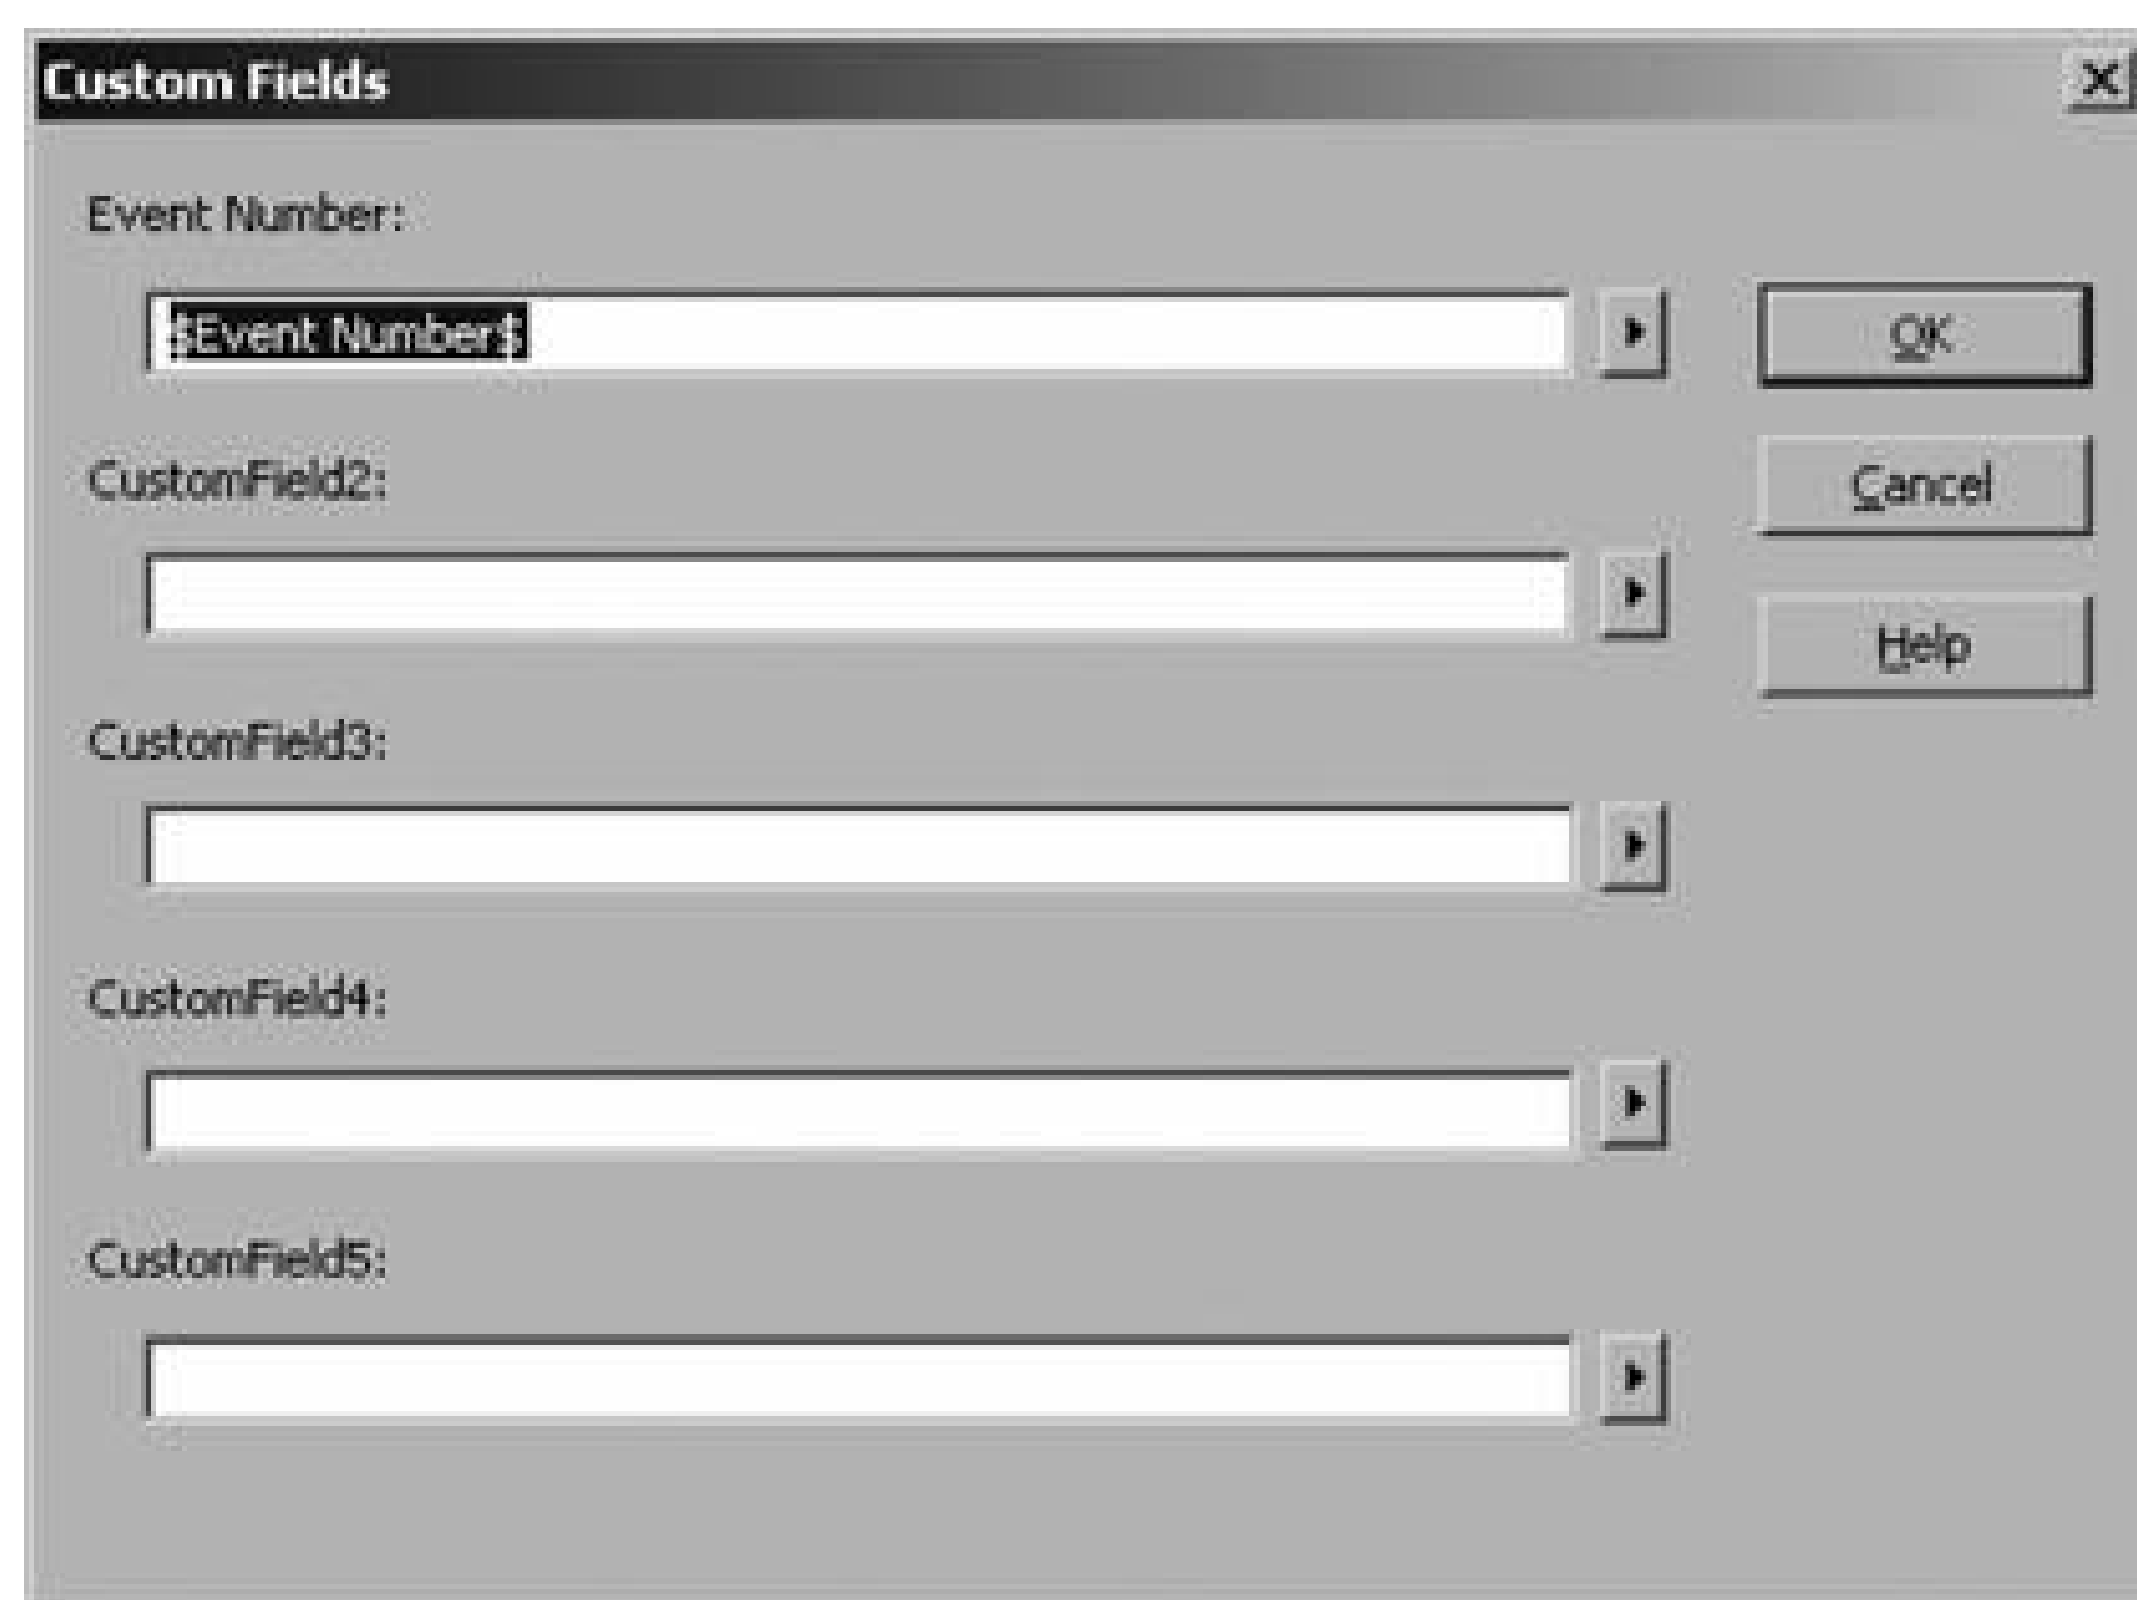


Figure 5-24. Concatenating plain text and alert variables into a custom field



## 5.2. Connections

Some inter-component communication channels are predefined and nonconfigurable, such as the protocols used for agent management or the connection between the reporting server and the operations database. Other types of communication and connections are configurable for the entire management group, such as the web addresses and the agent/management server TCP port. These global settings are covered in this section.

### 5.2.1. Web Addresses

Since Microsoft IIS is a prerequisite for the installation of MOM 2005, it is no surprise that MOM has several web-based interfaces. These interfaces are the Web console, the Online Product Knowledge Address, and the File Transfer Server, as shown in [Figure 5-25](#). There is also a field for the Online Company Knowledge Address, but this field, like all the fields on this tab, does not need to be configured for MOM 2005 to run satisfactorily. Management pack rules can make use of these default values that are defined at the global level shown in [Figure 5-25](#), or they can override them and use a custom value that is defined on the rule itself.

The Web Console Address is automatically populated if you choose to install the Web console. The entries on these tabs are merely text entries that can be used in alerts. Changing any configuration here does not update the actual web sites in IIS or DNS; you still have to do that manually.

If you have an online knowledge base or a help desk trouble ticketing system (or anything else with a URL that you want to reference) with web interfaces, you can make that link directly accessible in the Company Knowledge tab in the Details pane of all alerts by entering that URL in the Online Company Knowledge Address field.

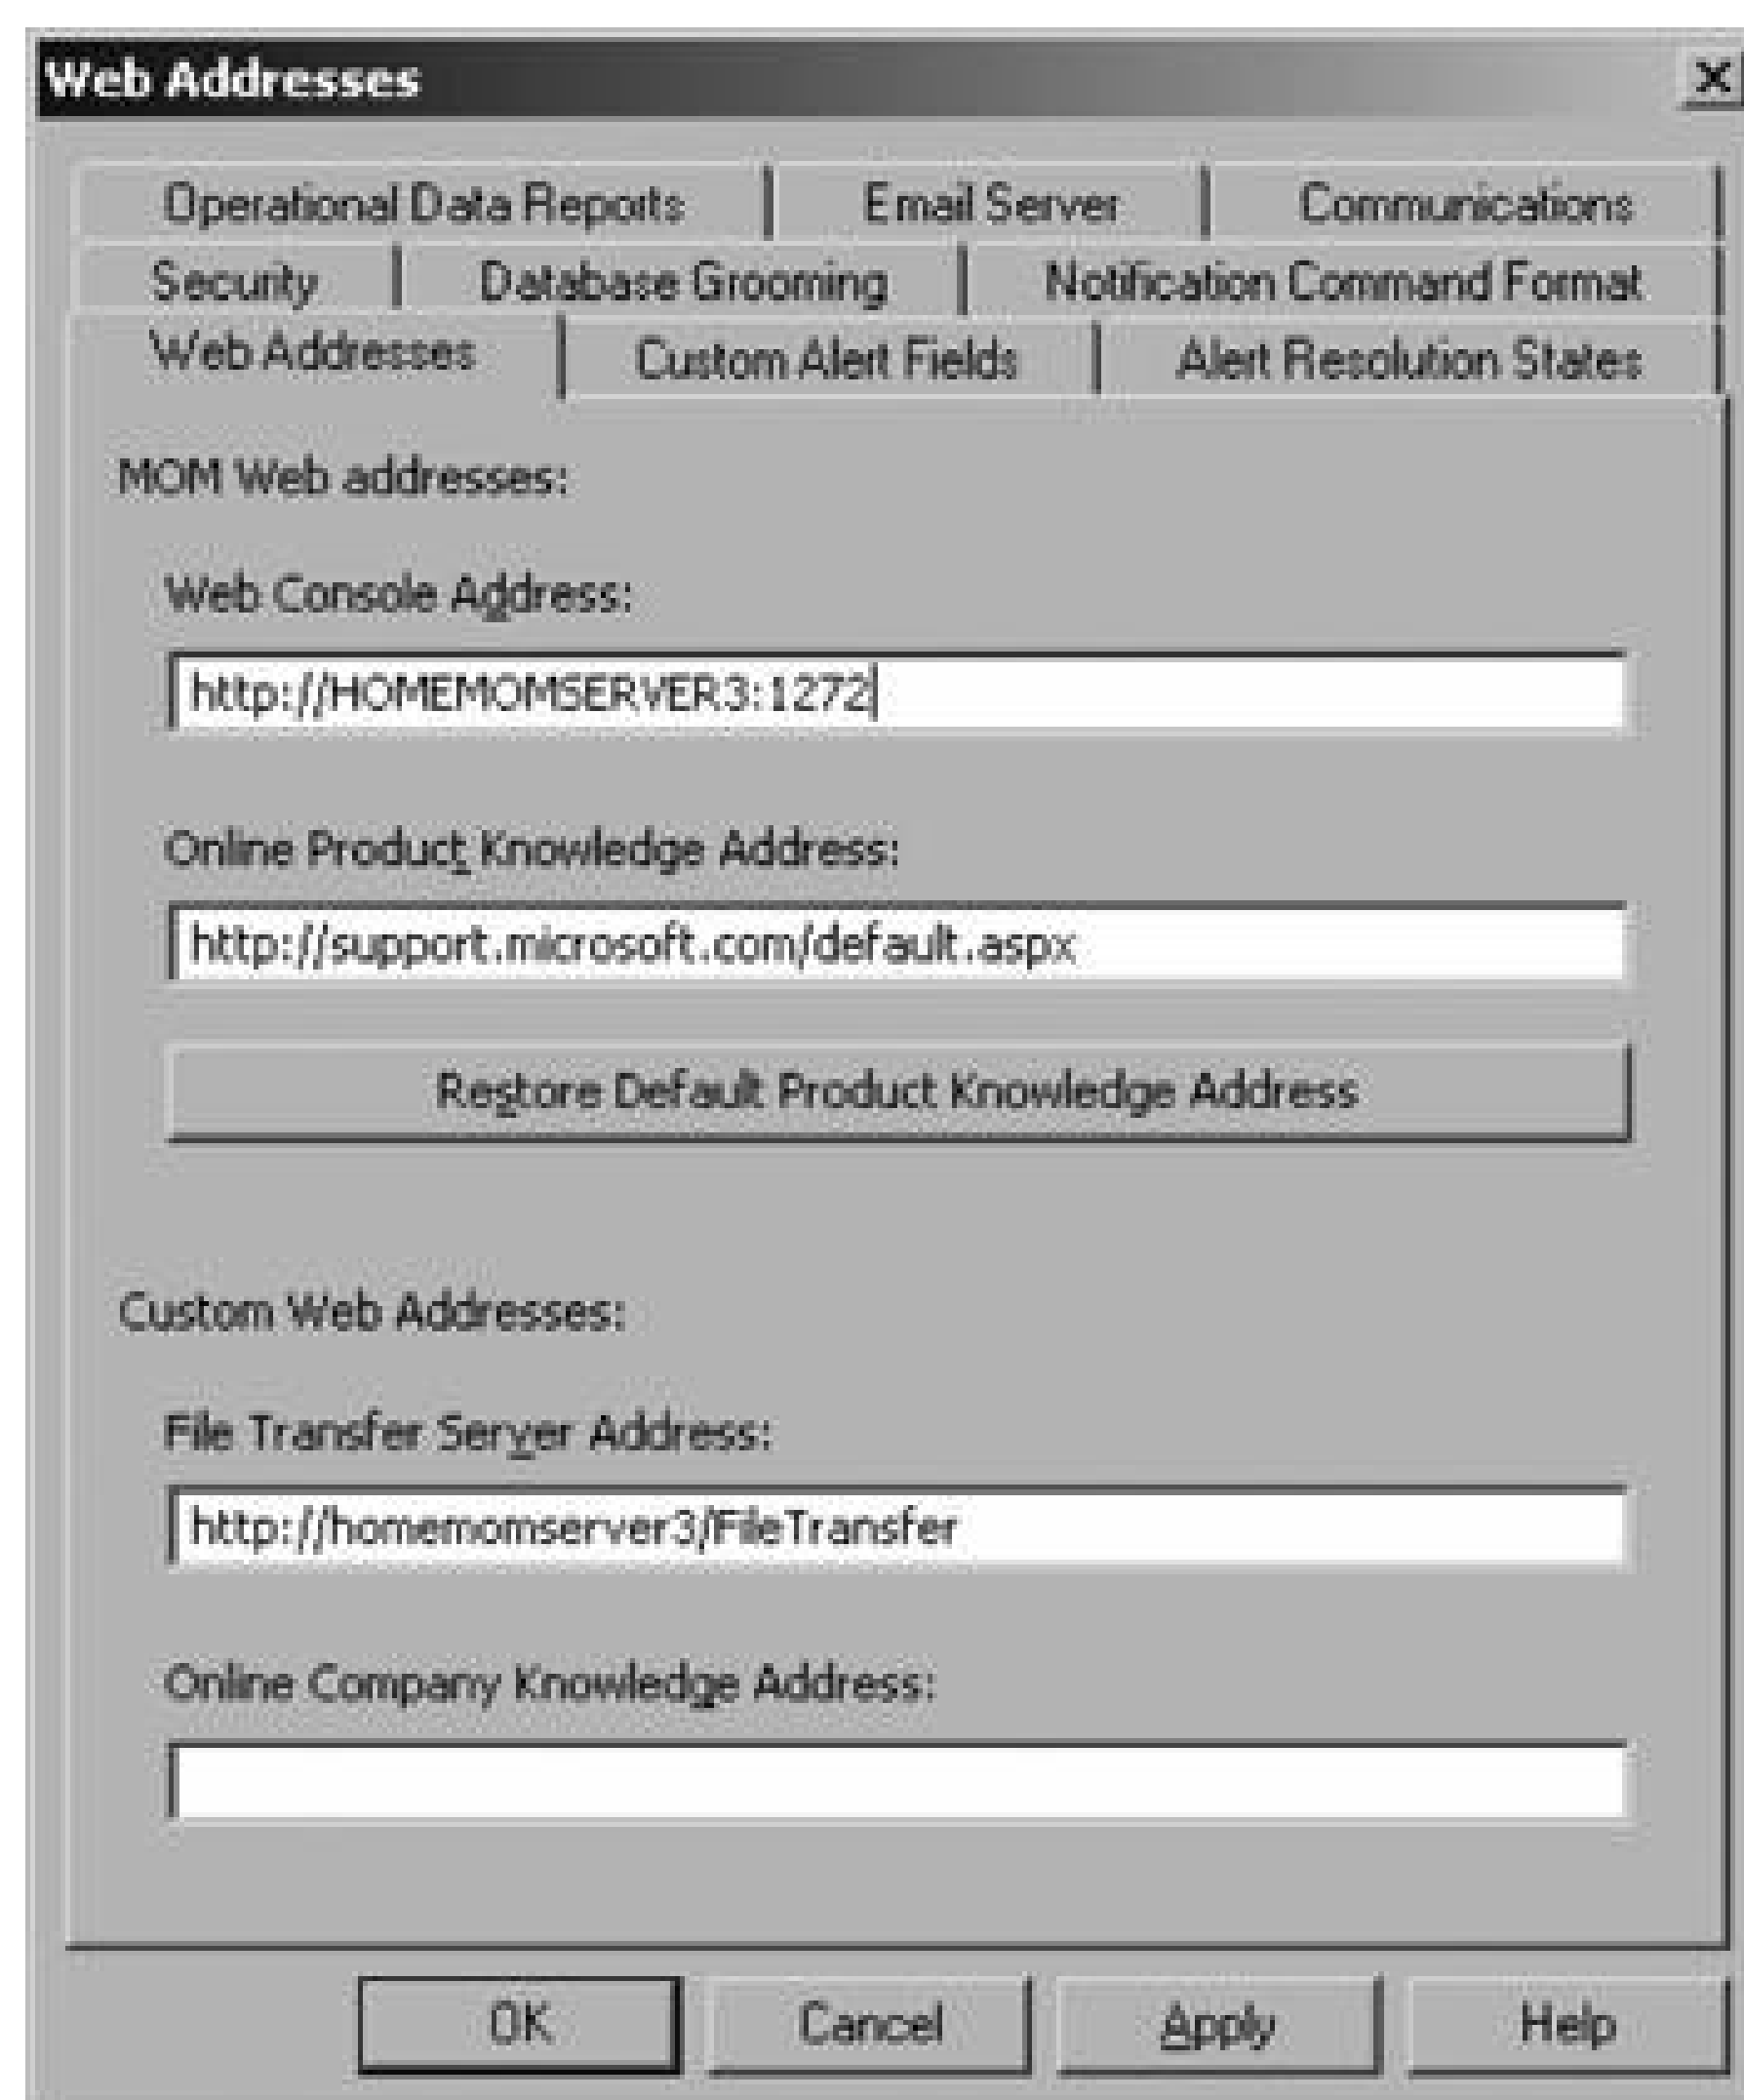
The file transfer response enables the bidirectional exchange of user-predefined files between an agent-managed computer and a file transfer server. On the file transfer server, you configure a virtual directory in IIS to be used as a location for transferring files. The transfer can occur on demand if it is launched from a task in the Operator console, or it can be triggered by event, alert, or performance rule criteria. The Microsoft Baseline Security Analyzer (MBSA) management pack uses this functionality.

It downloads the security profiles file (*mssecure.cab*) from the web and the agents request a download of the files for use in scanning.

The file transfer occurs over HTTP and the overall configuration of file transfer requires that the Background Intelligent Transfer Service (BITS) be installed on the file transfer server, along with IIS 5.0 or higher.

Figure 5-25. Specifying the management pack's web addresses





The configuration of the file transfer action is shown in [Figure 5-26](#). Once the MBSA scan of an agent-managed computer has been completed, every MBSA event is collected from the event logs, processed by the management server, and alerts are generated as appropriate.

The value that is entered here can be used by all objects that fire a file transfer response. Like most other global settings, this can be overridden at the rule and task level.

As shown in [Figure 5-26](#), there are four basic components that make up this configuration:

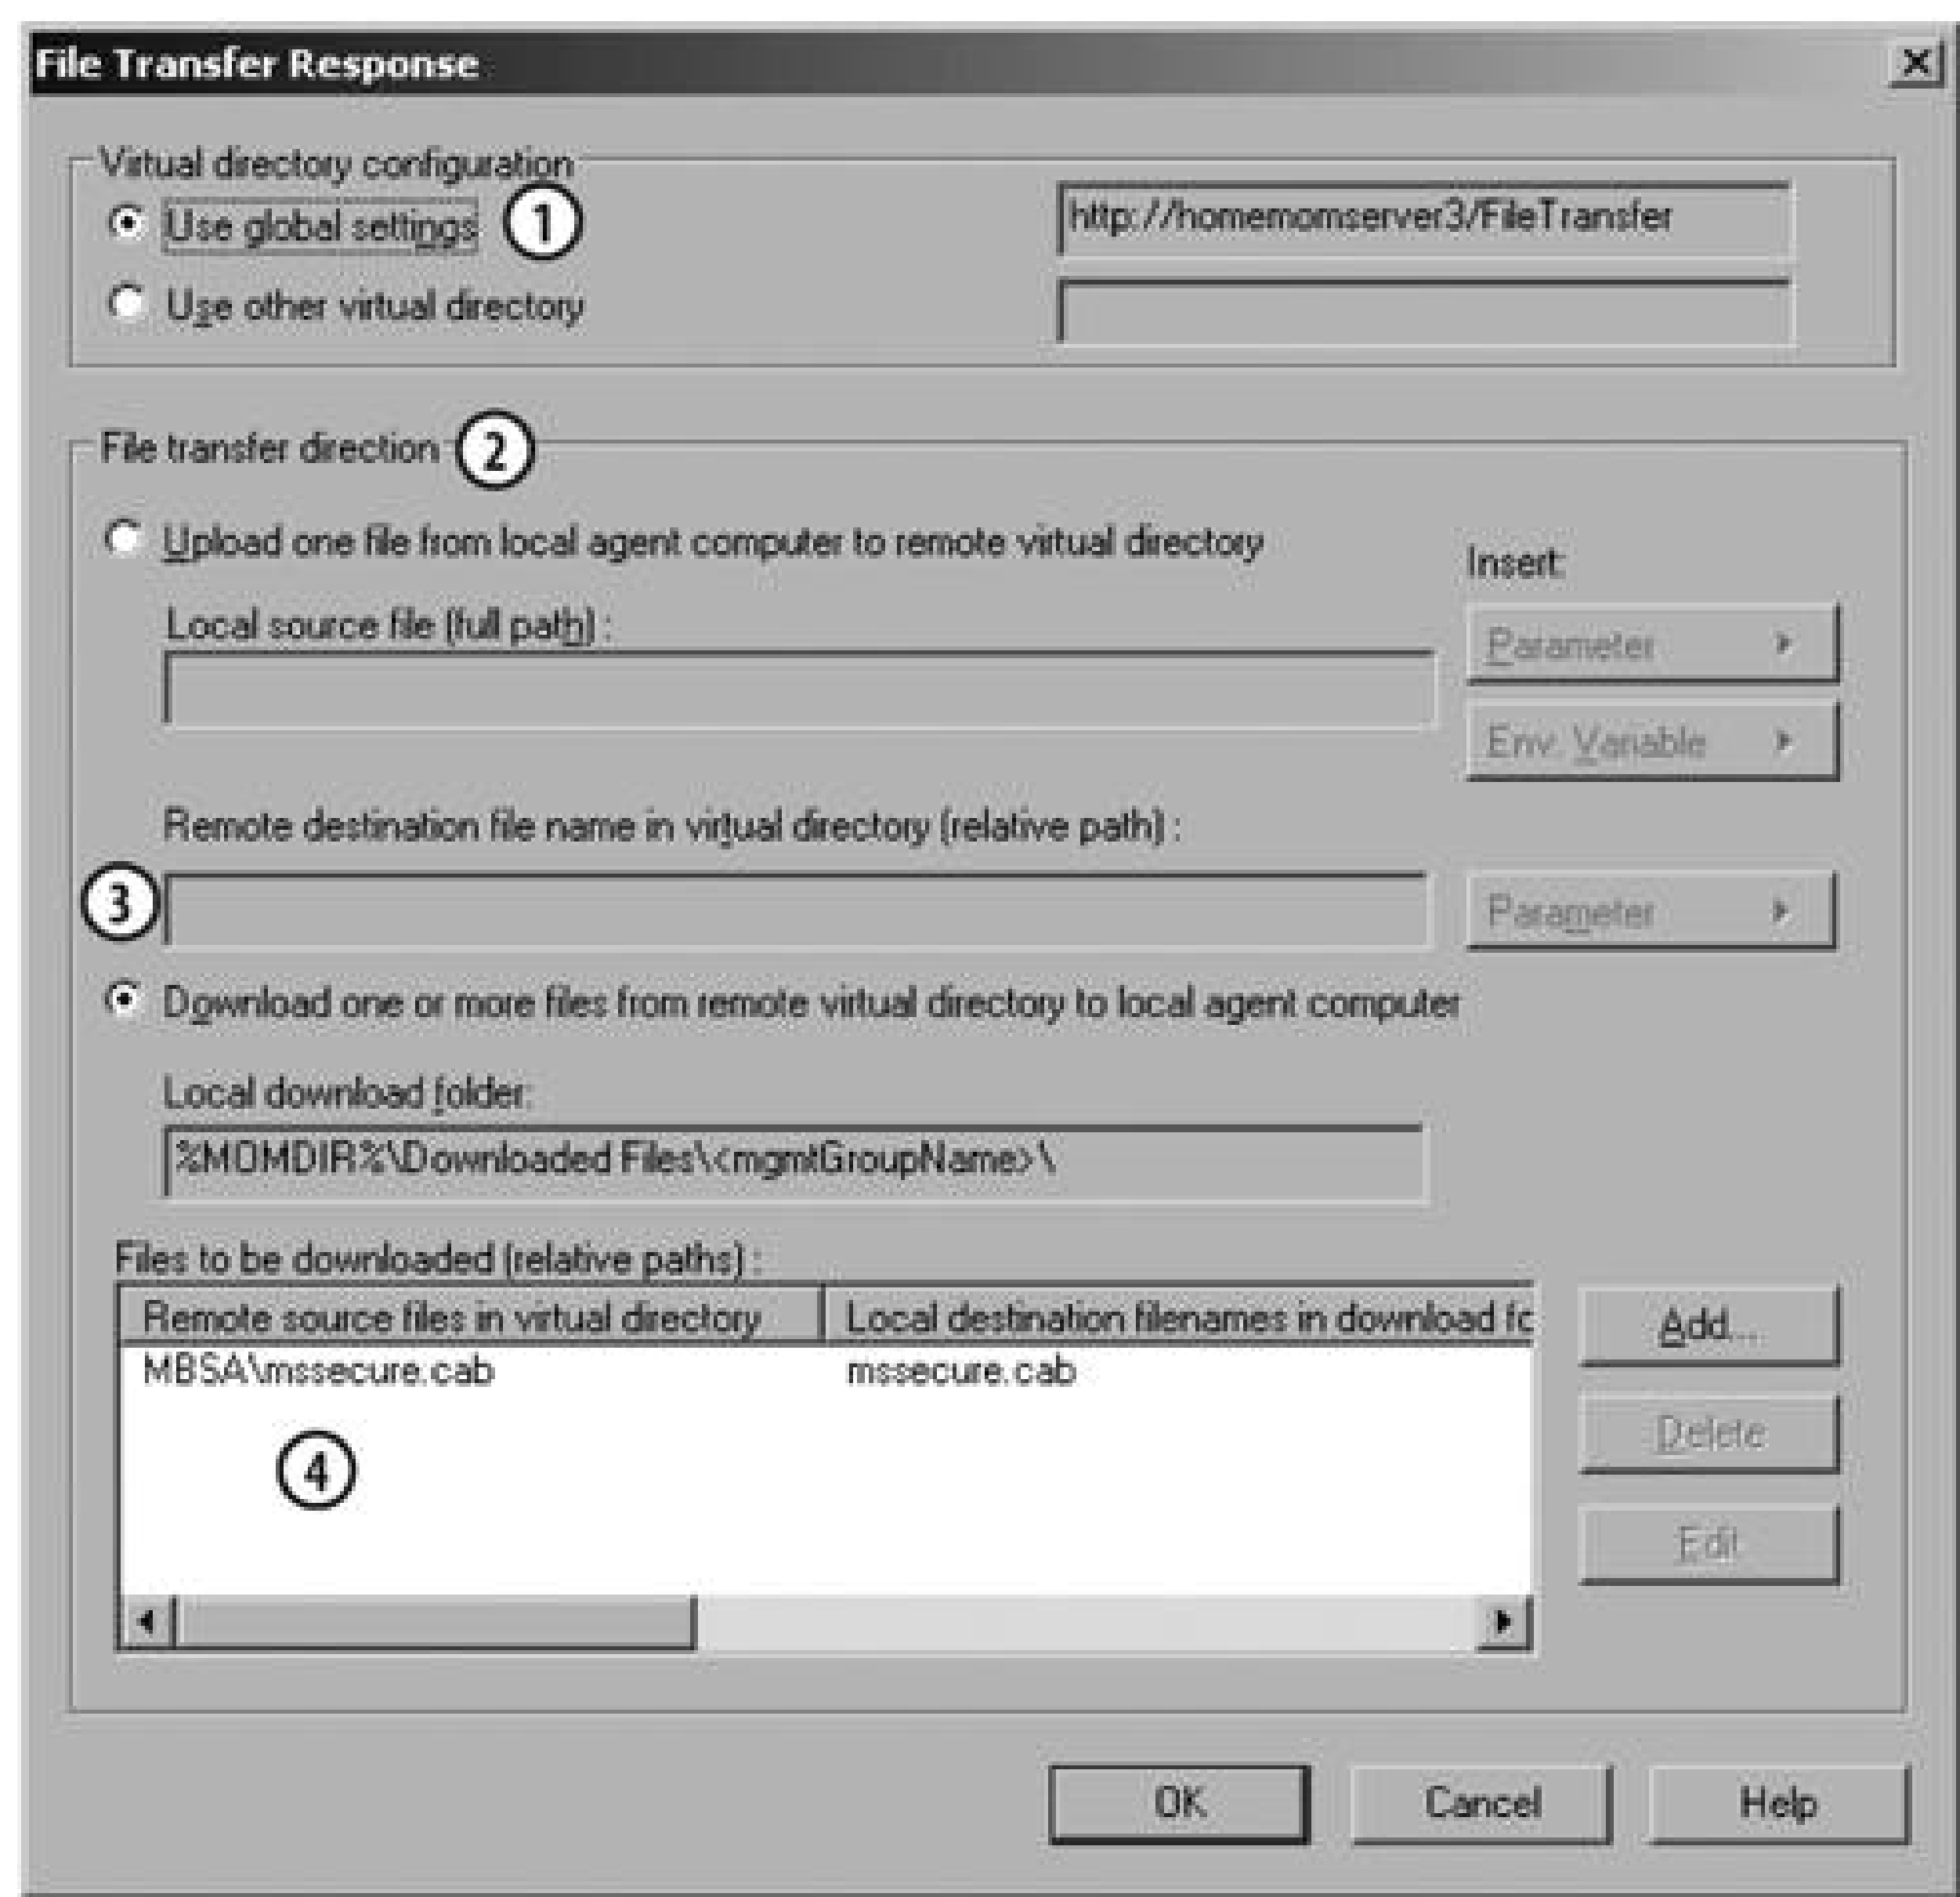
1. Choosing to use the default file transfer server URL or entering a custom URL.
2. Choosing the file transfer direction that this transfer action will execute.
3. Defining the directory to transfer to or from.
4. Identifying the file to be transferred .

## 5.2.2. Communications

The value on the Communications Settings tab of the Global settings is one of the few global settings

that cannot be overridden at a lower level, although it can be reconfigured (see the [Manual Agent Deployment](#)" section in [Chapter 3](#)).

Figure 5-26. An example of the file transfer configuration settings from the MBSA management pack



The default configuration is to use TCP port 1270 for all agent/management server communication. All communications over the port defined here are encrypted by default (see [Figure 5-27](#)). If a different port needs to be used for security or other reasons after the agents are deployed, make sure you update the existing agents.

## 5.3. Maintenance

The extent to which a MOM 2005 implementation is self-maintaining depends largely on the volume of data flowing into and being removed from the operations database. The ["Versions"](#) section in [Chapter 2](#) explains how data is groomed out of the operational database at configurable intervals. These intervals need to be coordinated with the DTS data transfer job that runs from the reporting server and copies (not deletes!) data from the operations database to the data warehouse database. In a nutshell, you want to be certain that data is copied to the data warehouse database more frequently than it is groomed (deleted) from the operations database. The default configuration is to copy data from the operations database to the data warehouse every day at 1:00 a.m. You can configure the grooming process in the Database Grooming section of the Global Settings node. The OnePoint database grooming will not occur if the daily DTS transfer has not completed successfully, so you are protected from losing data.

Figure 5-27. Configuring the default agent to management server communication port

### 5.3.1. Database Grooming

Deleting data from the operations database happens differently for alert data than it does for performance or event data. The difference is that alerts must be in a resolved state before they can be deleted out of the database. After all, it wouldn't be very helpful to delete an alert that has been



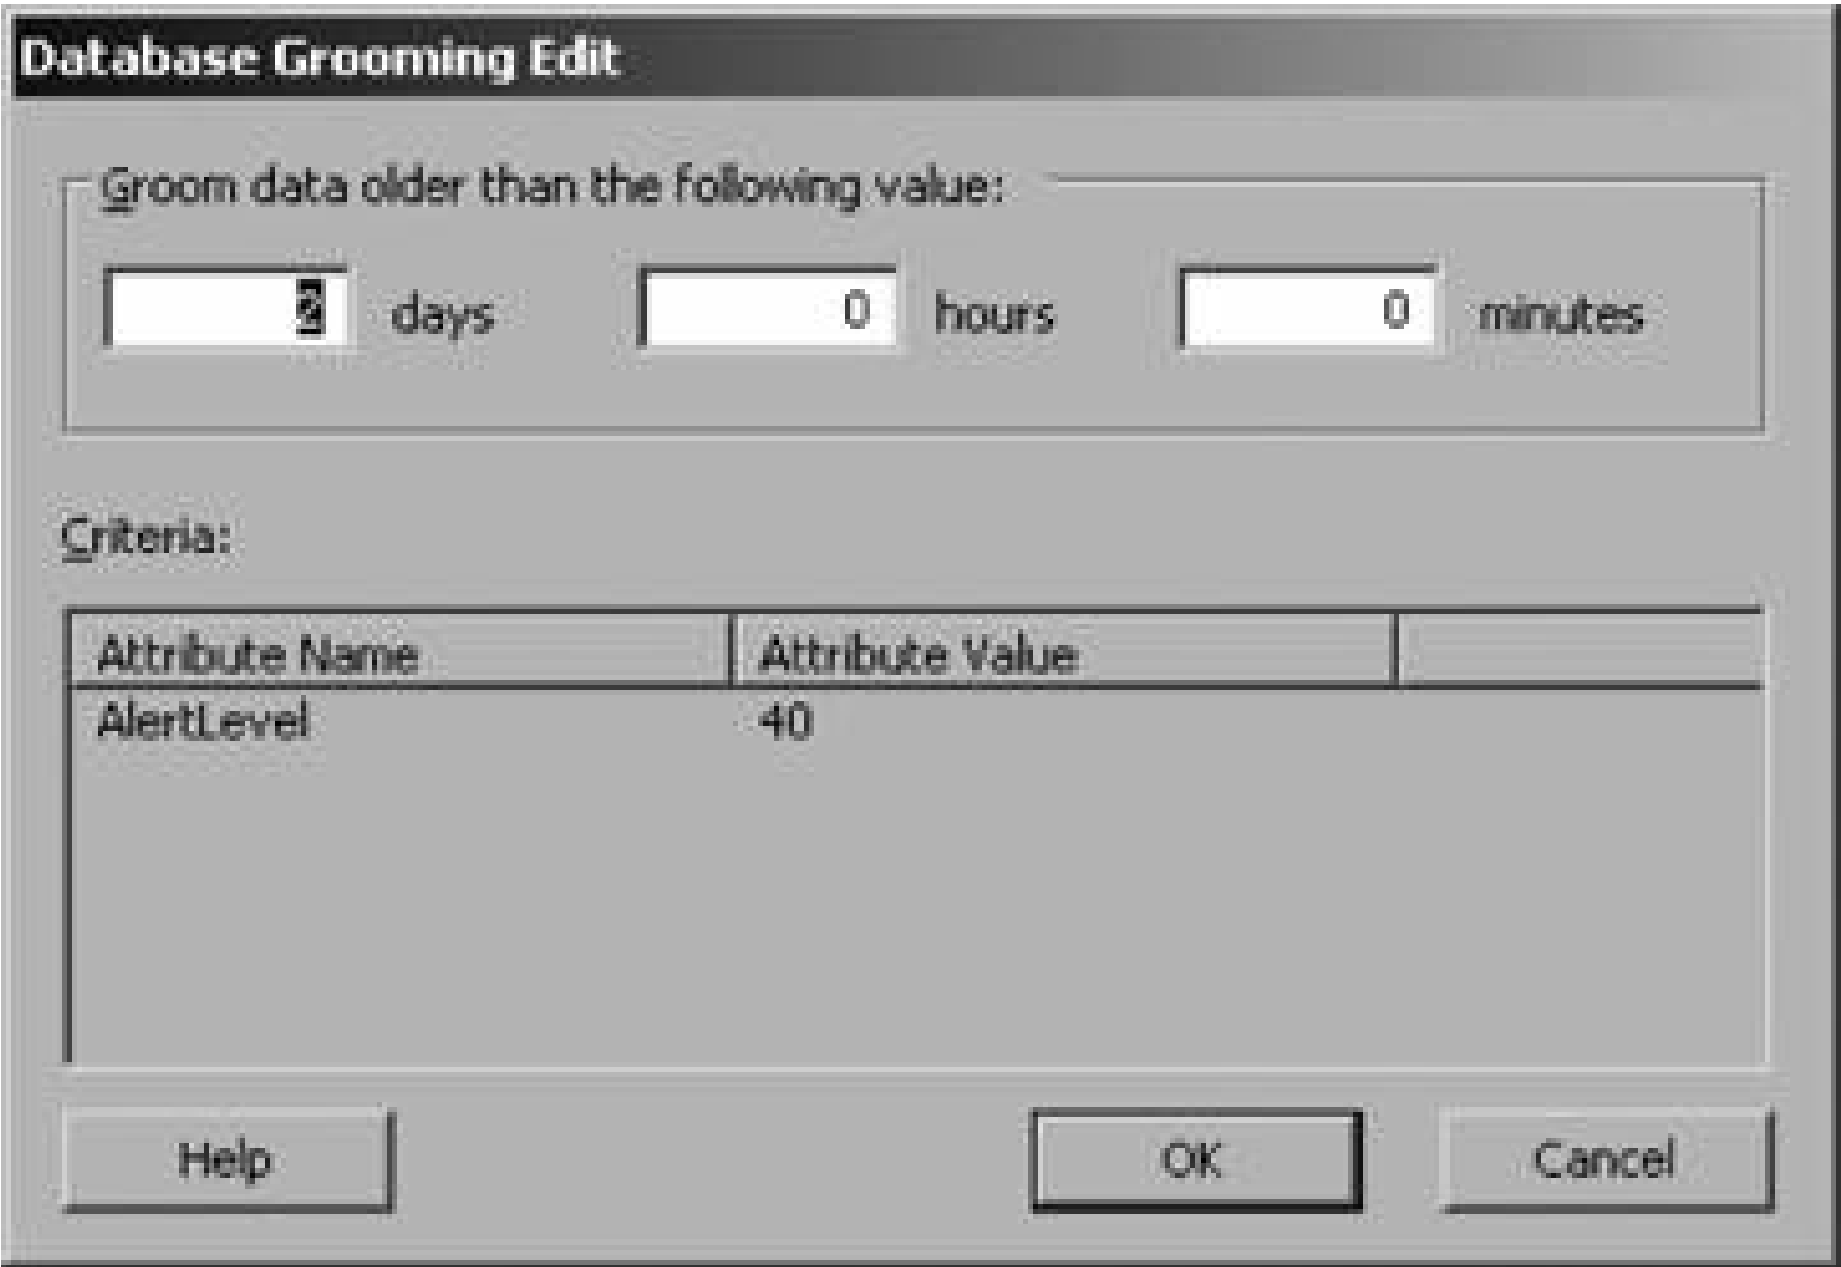
assigned to a vendor for resolution before the vendor fixes the problem. The default configuration for deleting data out of the operations database is to delete resolved alerts and event and performance data that is more than four days old. This is shown in the "Groom data older than the following number of days" setting in [Figure 5-28](#). Alert grooming is a little more complicated. MOM 2005 includes preconfigured SQL jobs that automatically resolve alerts after they have aged beyond the configured time limit defined in the "Specify when alerts are automatically resolved" box. For example, all alerts that are created with the default severity of Warning will have their alert resolution state changed from New to Resolved after one day unless an operator manually changes the alert resolution state from New to something else. Once the alert is in a Resolved state, it will be groomed out of the operations database after the four-day default limit.

You can edit individual auto-resolve time limits by selecting them in the Database Grooming tab and clicking Edit (see [Figure 5-29](#)). To determine if these values are adequate for your environment, you need to keep a close eye on your database size and the successful (or unsuccessful) completion of the database reindexing jobs and grooming jobs. If your operations database is rapidly growing to an unmanageable size, you can decrease the grooming time factor as well as the auto-resolve alert time factor.

Figure 5-28. Database grooming default settings

You may want to want to alter these values if you need more than four days to triage alerts. It really depends on how long you need to keep the data in the operations database "live." For details on database maintenance, see [Chapter 7](#).

Figure 5-29. Editing the time limits for auto-resolving alerts



Notice in [Figure 5-29](#) that the AlertLevel is referred to by its number in the Attribute Value field, not its name. In this case, the auto-resolve time limit for an Error alert is edited to have an Attribute Value of 40. [Table 5-1](#) shows the default alert severity levels and their corresponding numbers.

Table 5-1. Alert severities and their internal numerical values

Alert severity name	Alert severity number
Success	10
Information	20
Warning	30
Error	40
Critical Error	50
Security Issue	60
Service Unavailable	70

Notice that even though the Inactive state is listed as "Auto resolve inactive Alerts" in[Figure 5-28](#), it is not an alert severity level. This is a tip-off that MOM will not groom out any alert that is in an active state, regardless of its severity or age.

### 5.3.2. Operational Data Reports

Just as Microsoft included automated error reporting in their current OSes and Office applications, they have chosen to build in the generation and transmission of semi-scrubbed operational reports in MOM 2005. Basically, if you elect to send these reports to Microsoft (which they recommendgosh, there's a surprise), they will collect summary information on the number and type of alerts, your basic management group configuration, number of deployed agents, which management packs are deployed, and how long you are taking to resolve alerts.

View a sample of the data that Microsoft collects at <http://www.microsoft.com/technet/prodtechnol/mom/opreport/samplereport.mspx>.

If you don't have the MOM 2005 reporting feature installed, then it will take some effort to collect the same information on your environment for your own use. Microsoft should have provided this basic type of reporting in the Operator console out of the box.







## 5.4. Summary

This chapter reviewed the management group global settings that had not yet been covered yet in the book. In most cases these global settings merely provide values that can be used or disregarded at the agent and rule level. In this way, these settings are not effective by themselves but need to be configured in concert with alerts, rules, and tasks.

With the end of this chapter, almost all of the functions in the Administrator console have been introduced and discussed. The next chapter moves the focus of the discussion to the Operator console.



# Chapter 6. Operator Console

With the minor exception of [Chapter 1](#), this book has focused on aspects of MOM that are controlled and viewed through the Administrator console—specifically, how to configure and control the behavior of management groups and the information that they produce.

This chapter shifts focus to cover the Operator console and how to get the information you want out of it. As such, this chapter is for consumers of MOM information, not just MOM administrators. However, MOM administrators are not off the hook; you need to go through this chapter as well after all, who are the MOM users going to turn to for help with using the MOM 2005 Operator console?

The Operator console will tell you what is going on in your environment right now. Two types of user/console interactions are necessary to find that out. The first is MOM-initiated and it happens when a significant event occurs anywhere in the environment. It comes in the form of an alert or a state update. In this instance, MOM is saying, "Hey you, this thing that you wanted me to tell you about, well, it just happened over here." The top-level alert and Operator console State views provide a broad perspective of the whole environment. This can be limited to a portion of the environment if you've implemented console scopes.

The other use of the Operator console is to find out what happened or is happening in a smaller portion of the environment. For example you may need to find out, "What is the available disk space on the Exchange mailbox servers?" or "What was happening on the ISA when the DMZ web site was hacked this afternoon?"

This chapter could just as easily have been called "A million and one ways to look at your MOM data," because the Operator console is so flexible in how it can present the collected information. The simplest way to describe the Operator console is that it is a filtering tool used to find operational data in the operations database. Using the Operator console effectively is really a matter of learning how to manipulate its filtering functions.

This chapter teaches the use of the Operator console through several scenarios presented in the context of the Leaky Faucet environment. The information MOM Operator console users will see and the tasks they have access to is controlled by the console scope that is applied.



# 6.1. Console Scopes

To open the MOM 2005 Operator console, your account must have membership in one of the management server local MOM security groups, such as MOM Users, MOM Administrators, or MOM Authors. This can be by virtue of membership in a domain security group that is itself a member of one (or more) of these local groups or by being explicitly named in the local security group. When you launch the Operator console, these permissions are checked; if your account doesn't have permission, the operation fails.

In addition, your account is also associated with a *console scope*. Console scopes are composed of computer groups. If a computer group is in your assigned scope, you will be able to see that group and all of its information and associated tasks in the Operator console. An account can be associated with only one console scope at a time. [Figure 6-1](#) shows the default MOM Administrator Scope and some of the member computer groups.

Figure 6-1. The MOM Administrator Scope contains all computer groups

The Administrator, Operator, and User console scopes contain all the available computer groups by



default. You can change this configuration by adding or removing computer groups or users in the Administrator console in the Administration node under Console Scopes, but it is not recommended. Many times the path followed to troubleshoot an alert will cross computer groups, so the responsible person will need to have access across all computer groups to do her job. Besides composing a console scope, the computer groups are also used as the last step in filtering for the data you want when working with the Operator console views.

At Leaky Faucet, remote office administrators have been delegated limited authority over their organizational unit and they can only view a limited set of operational MOM data, so it is appropriate to create additional console scopes for those users.



## 6.2. Creating a Custom Console Scope


Before creating the console scope, you need two things: a computer group that contains the computers of interest, and user accounts. Administration of accounts is easier if they are in a security group, but they don't have to be.

The following are the steps that Leaky Faucet follows to grant the necessary accounts permissions to the management group:

1. Create a domain security group and populate it with the accounts that will be using the custom console scope. Leaky Faucet creates the user accounts LKFRemoteSiteAdmin1, 2, and 3, and places them in a newly created domain security group LKFRemoteSiteAdmins (see [Figure 6-2](#)).

Figure 6-2. Leaky Faucet's domain security group for remote site administrators

2. On all the management servers in the management group, add the security group that was created in step 1 to the MOM Users local group. This grants permission to launch the Operator console and assigns servers to the default console scope (MOM User). The assigned scope will be overridden when the individual accounts are associated with the custom console scope. Remember, a user account can only be associated with one scope at a time. In cases where an account might be associated with multiple scopes, such as the default and a custom scope, the most recently created account/scope association takes precedence.



Only user accounts can be named in a console scope, not security groups.

Creating the custom computer group and console scope is a little more involved, but is accomplished through the use of wizards in the Administrator console. The custom computer group is a very simple computer group because all included servers are explicitly named. Because this group will not be associated with any rule groups, state roll-up will also not be used.

[← PREV](#)



## 6.3. Creating a Custom Computer Group

As mentioned in [Chapters 3](#) and [4](#), computer group definitions are included in management packs and instantiated in a management group when the management pack is imported. Computer groups that are created through this method have already been associated with rule groups. When the computer discovery process runs, the computer groups are populated with computers according to the filter criteria for that group. Rules are then sent to the agents on those computers and monitoring begins.

For his own purposes, Max, the Leaky Faucet MOM administrator, will create a custom computer group that is not associated with any rule groups, but exists only to identify the computers that the remote office administrators will work with in the Operator console.

The following are the steps taken to create the custom group:

1. Name the group. In the Administrator console, navigate to the Computer Groups node, bring up the context menu, and create a new computer group. Proceed through the Welcome page to the name and description page. Max names the group LKF Remote Office Servers and puts in a description ([Figure 6-3](#)).
2. Add computers via explicit naming. Next, the computers explicitly wanted in this computer group are added (see [Figure 6-4](#)). Max is explicitly naming the servers because there is no criteria to calculate the group membership. For the sake of this example, *homesqlserver* and *homesrv02* are selected.

Figure 6-3. Name the custom computer group and describe its purpose

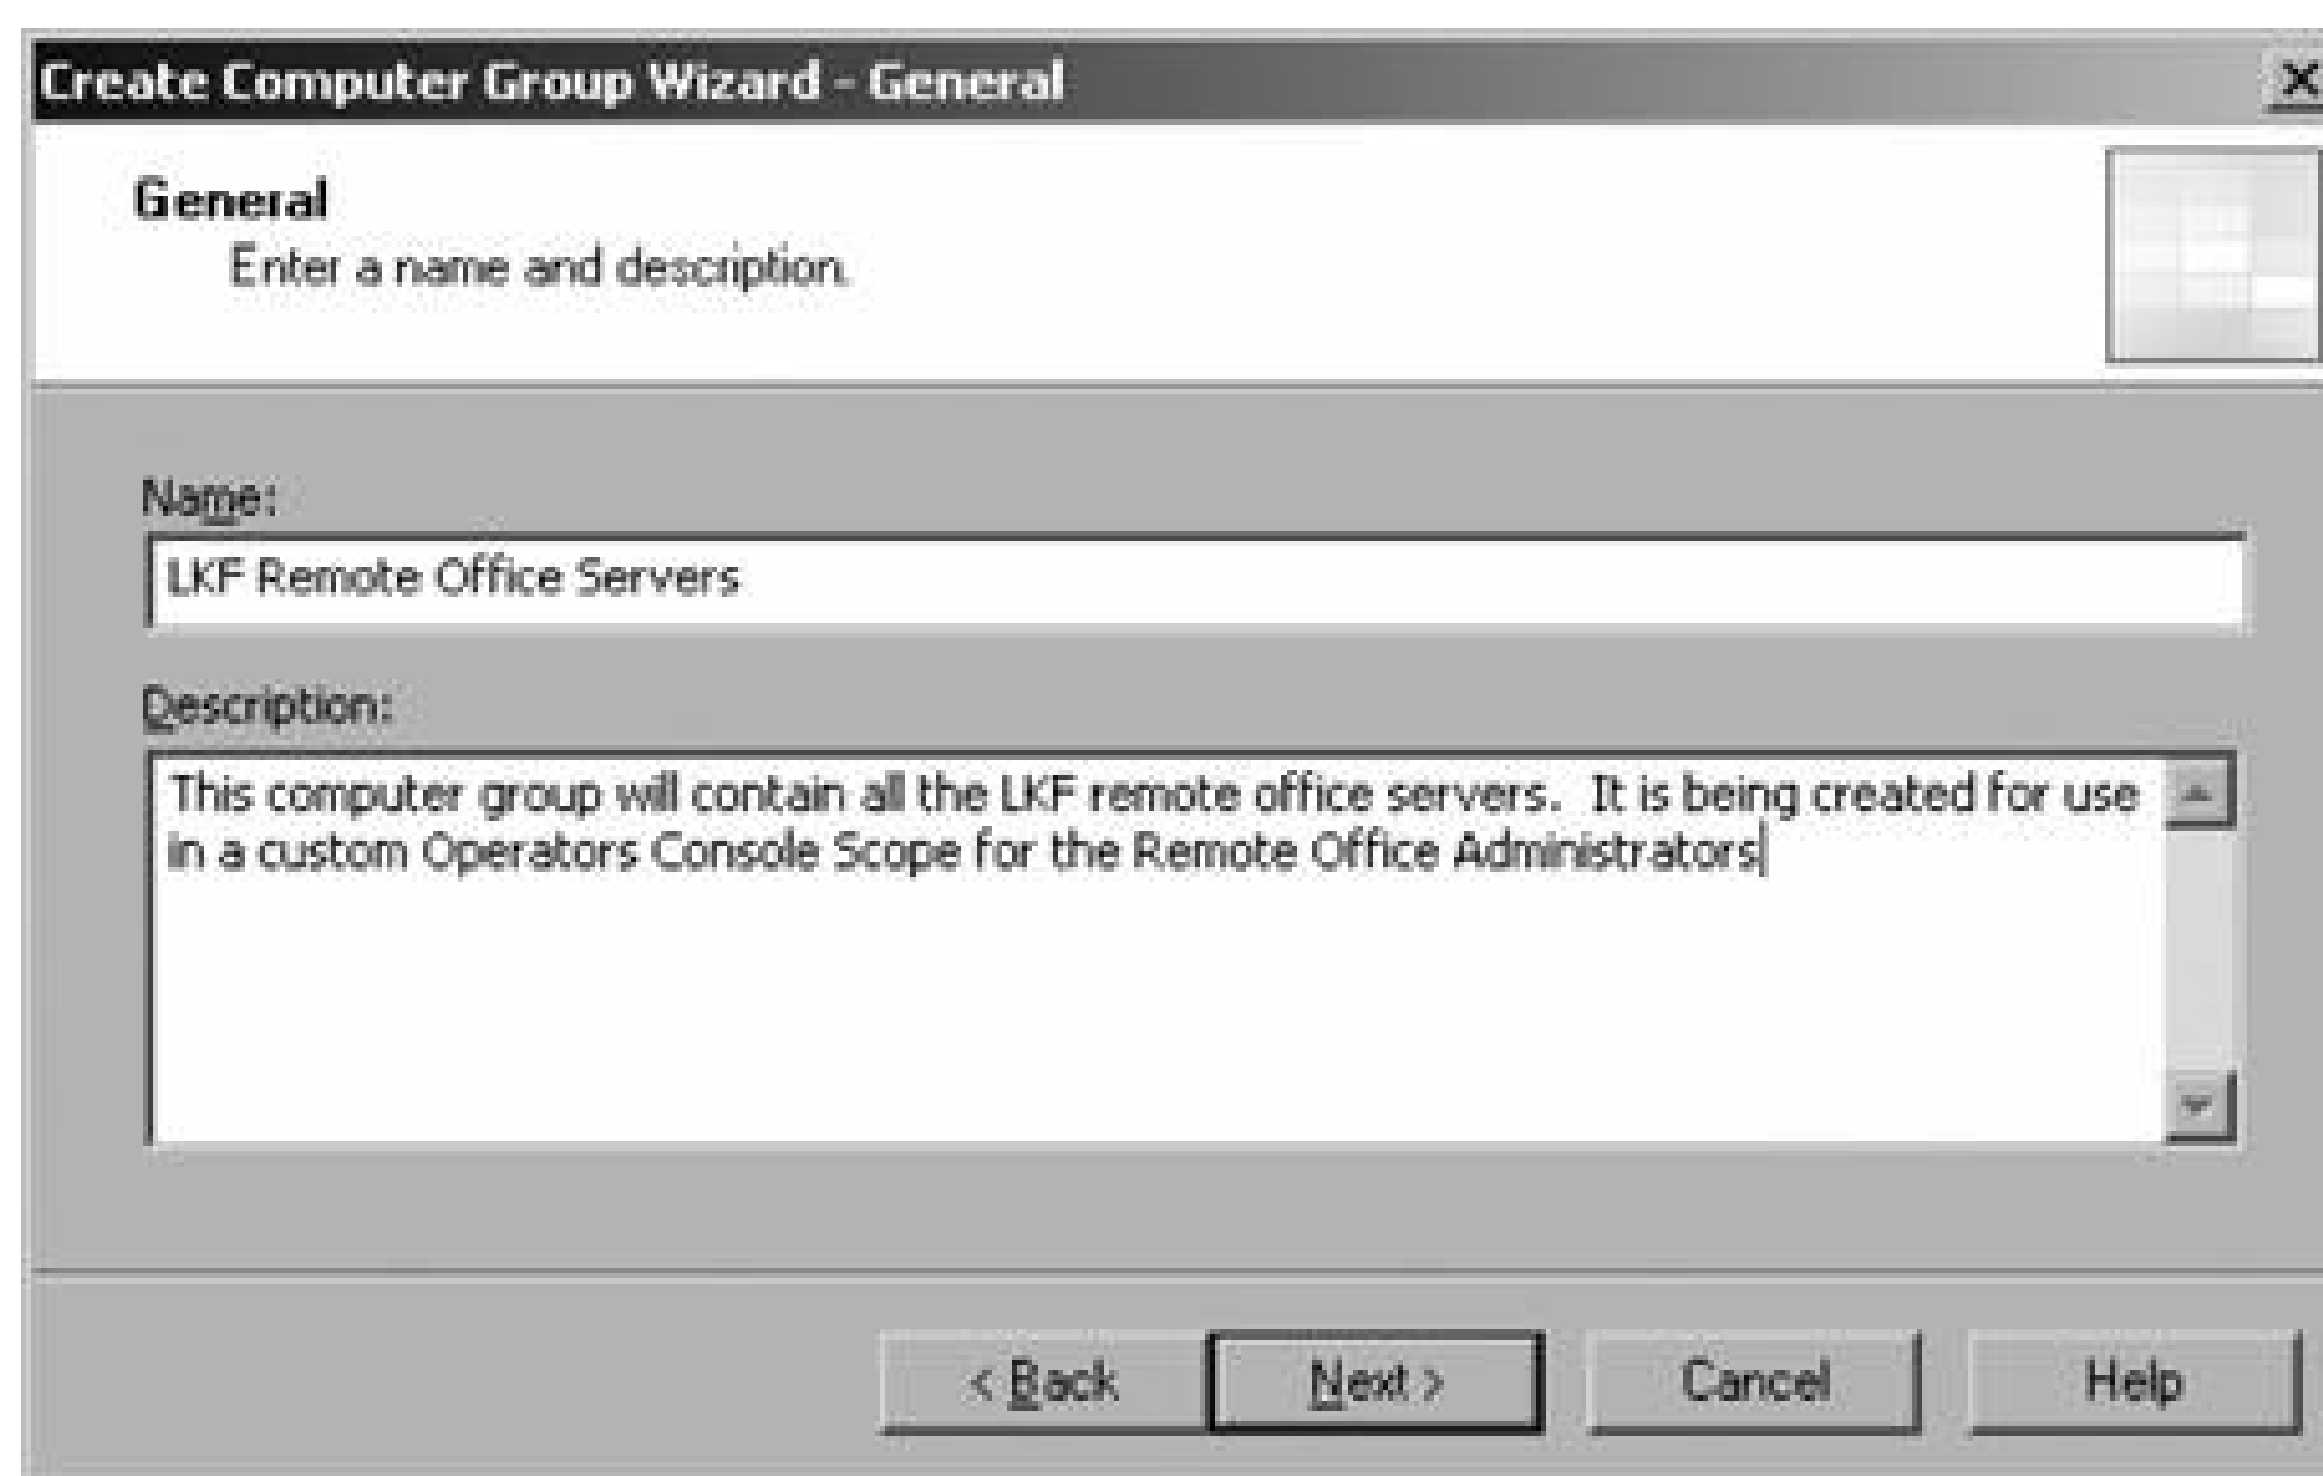



Figure 6-4. In the "Create Computer Group" wizard, you can explicitly name the computers to be included in addition to establishing search criteria

Otherwise, MOM could add computers to the group based on search criteria, such as a registry value or domain membership and NetBIOS name. Other computer groups in the management group could also be added during this step.

3. Define computer exclude, search, and formula criteria. As with overrides in processing rules where exception criteria are established for the rule, the Create Computer Group wizard will allow Max to exclude computers based on explicit naming or search criteria.

The next pages in the wizard allow Max to specify domain and computer name and type search criteria. Since Max is explicitly naming the group membership, neither the search or formula pages are used for determining membership in this computer group.



To practice setting search and formula criteria, do so in your preproduction management group. You can create unlimited test computer groups and populate them; just don't make the rule group association and the rules will not be sent to the agents. Delete your test groups when you are finished.

4. Set state roll-up policy. If Max was going to create associations between this computer group and rule groups, this is where he would define the overall state of the group. The wizard requires this field to be filled in, so Max accepts the default, which is "the worst state of any member computer or subgroup."
5. Confirm choices. [Figure 6-5](#) displays the summary of the choices made. The only information required for this simple computer group is the group name, the computers included, and the mandatory (but nonfunctional) roll-up policy. The description, while useful, is not technically required. Clicking Next brings up the Finish page with the Finish button.

Figure 6-5. Confirming your custom computer group settings



## 6.4. Creating the Console Scope


Max is now ready to create the custom console scope:

1. Start by navigating to the Console Scopes object under the Administration node in the Administrator console.
2. Right-click to bring up the context menu.
3. Select Create a New Console Scope.

This launches the Create Console Scope wizard. The information required is: a name for the scope, the computer groups to include in the scope, and the user accounts (not user groups) to associate with the scope.

- Give the scope a name, such as "Remote Office Admins Operators Console Scope," and a description such as "This Operator's console scope is being created for use by the Remote Office administrative staff. It will contain only the remote office servers."
- Add the desired computer group to the scope by clicking the Add button, which brings up the page shown in [Figure 6-6](#). Listed here are all of the computer groups in the management group.

Figure 6-6. Select the LKF Remote Office Servers computer group to include it in this console scope



For the purposes of this scope, only the LKF Remote Office Servers computer group is selected, yielding the computer group selection shown in [Figure 6-7](#).

- The Has All Computer Groups checkbox is selected by default for the MOM User, MOM Author, and MOM Administrator console scopes. For this scope, Max deselects this option.
- This brings up similar pages for adding user accounts to the console scope. The pages used to add the user accounts look different (see [Figure 6-8](#)). This is because they differ from the normal object picker used to add user accounts to a group. In MOM 2005, only the domain and account username text string is used to associate an account with a console scope, not the accounts Security Identifier (SID). Console scopes cannot be used to enforce security for individuals whose accounts have rights to MOM.

Figure 6-7. Including the LKF Remote Office Servers computer group in the Remote Office Admins Operators Console scope

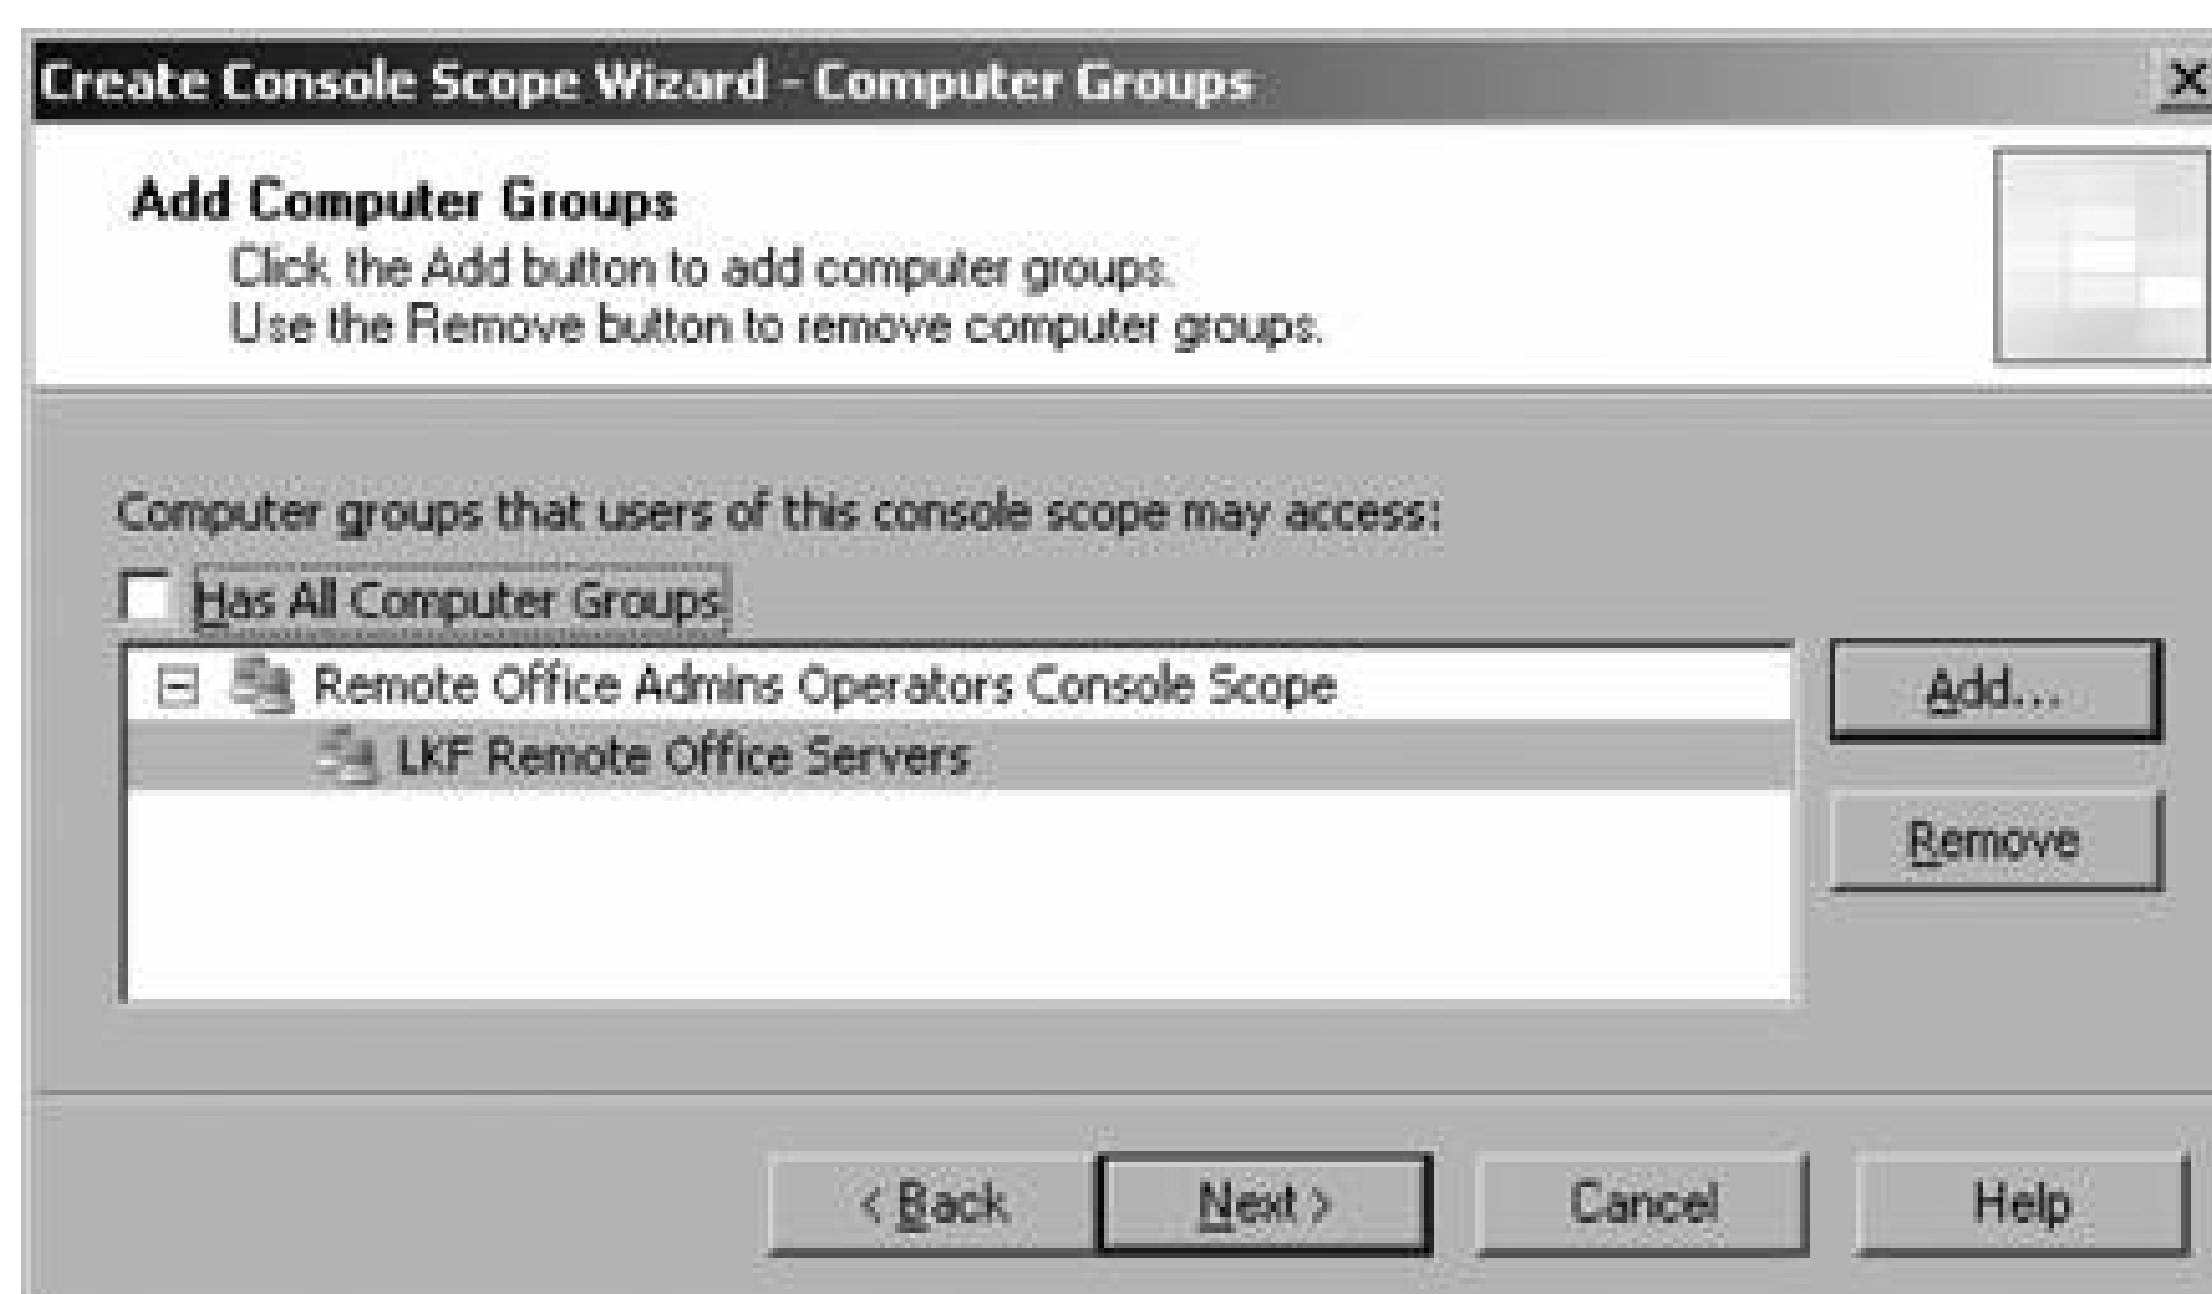


Figure 6-8. Naming user accounts to associate with the console scope

For example, the account chrisf (Christian Fowler) has MOM User rights via group membership and is associated with console scope A. If the chrisf (Christian Fowler) account is deleted, then it is removed from the MOM Users group but not from the scope definition. Another chrisf (Chris Fox) account could require MOM User access to more than scope A. Since console scope association is performed by account name evaluation only, the new chrisf (Chris Fox) account would, by default, be assigned the scope A console scope that was assigned to chrisf (Christian Fowler) even though they are two entirely different people.

Console scopes are useful only for filtering the computer groups that an Operator console user sees by default. As long as you stick to this use of console scopes, you won't get into trouble. If you need to provide a hard security boundary around the computer groups in the Operator console, you have to create an additional management group and multi-home selected computers into the second group. The next step is to grant MOM permissions to the second management group for the desired accounts and deny them access to the first management group. This is not very cost-effective, but it works.

Leaky Faucet adds the LKFRemoteSiteAdmin1, 2, and 3 accounts to this console scope (see [Figure 6-](#)

9).

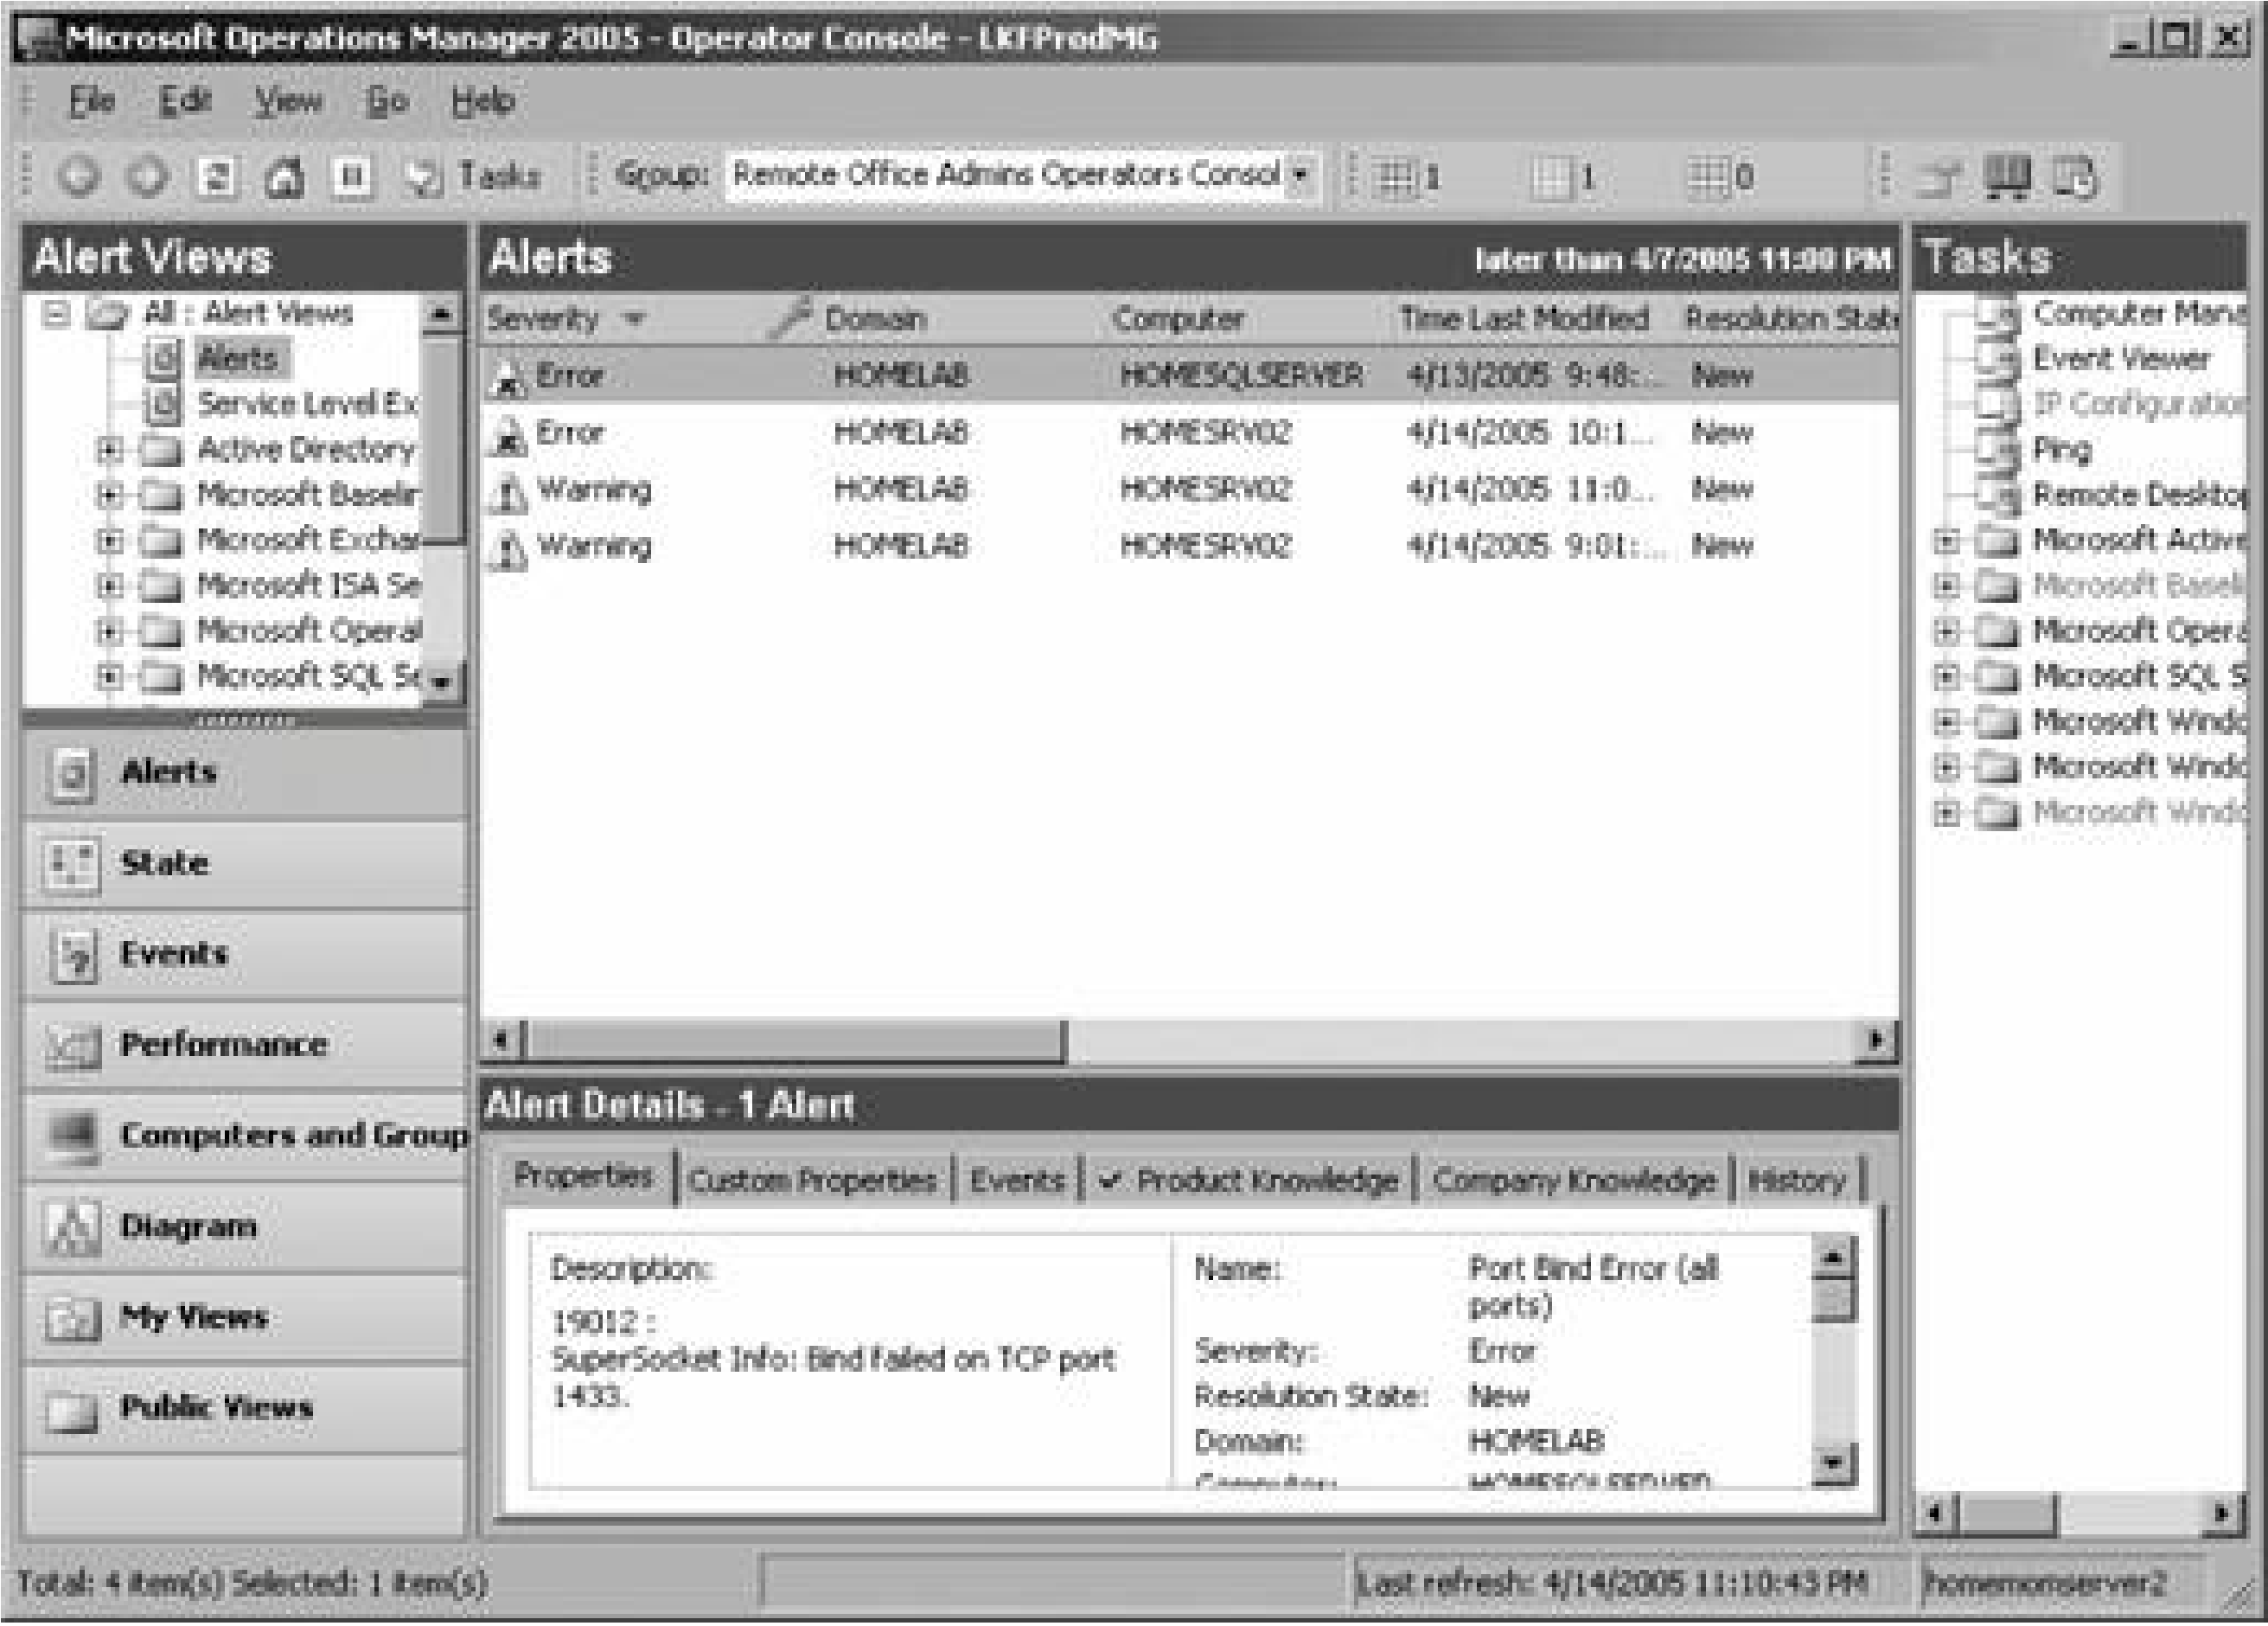
Figure 6-9. All the accounts associated with the custom console scope



Moving onto the next page finishes the wizard and the configuration is complete. Now, whenever any of the LKFRemoteSiteAdmin1, 2, or 3 accounts launch the Operator console, this console scope will appear by default. It cannot be changed and only data from the *homesrv02* and *homesqlserver* computers will be seen (see [Figure 6-10](#)).

Figure 6-10. Applied remote office console scope to logged-on LKFRemoteSiteAdmin3 user





Computer groups that belong to a scope can be used as filters. It is in this context that console scopes really shine. The next section demonstrates a specific methodology for building Operator console filters, and computer groups are a big part of that. Although the Operator console has a complex interface, following this three-step method will allow you to get the information with as little confusion as possible.

## 6.5. Using the Console

To understand the flow of filter-building in the Operator console, consider the navigation of the Outlook 2003 interface. There are very strong similarities between navigating Outlook 2003 and building a data filter in the MOM 2005 Operator console.

Basically, the Outlook 2003 interface has four panes (see [Figure 6-11](#)) to navigate and access email, calendar, tasks, and other types of information. These panes serve the exact same purpose in Outlook 2003 as they do in the Operator console.

To access information in Outlook 2003, you select the type of information you want, such as email, calendar, contact, tasks, and so on, from the folder buttons (point 1 in [Figure 6-11](#)) in the navigation pane. For this example, the mail folder button has been selected. The contents of that folder are displayed in the upper portion of the navigation pane, which also drills down a little to the subcontents of the top-level folder (point 2 in [Figure 6-11](#)). The Example Folder is selected, which brings a listing of the contents of that folder into the results pane (point 3 in [Figure 6-11](#)).

The contents of the results pane change if a different folder is opened in the navigation pane. The details pane (point 4 in [Figure 6-11](#)) displays the contents of the email object that has the focus in the results pane.

## 6.6. Building a Filter in the Operator Console

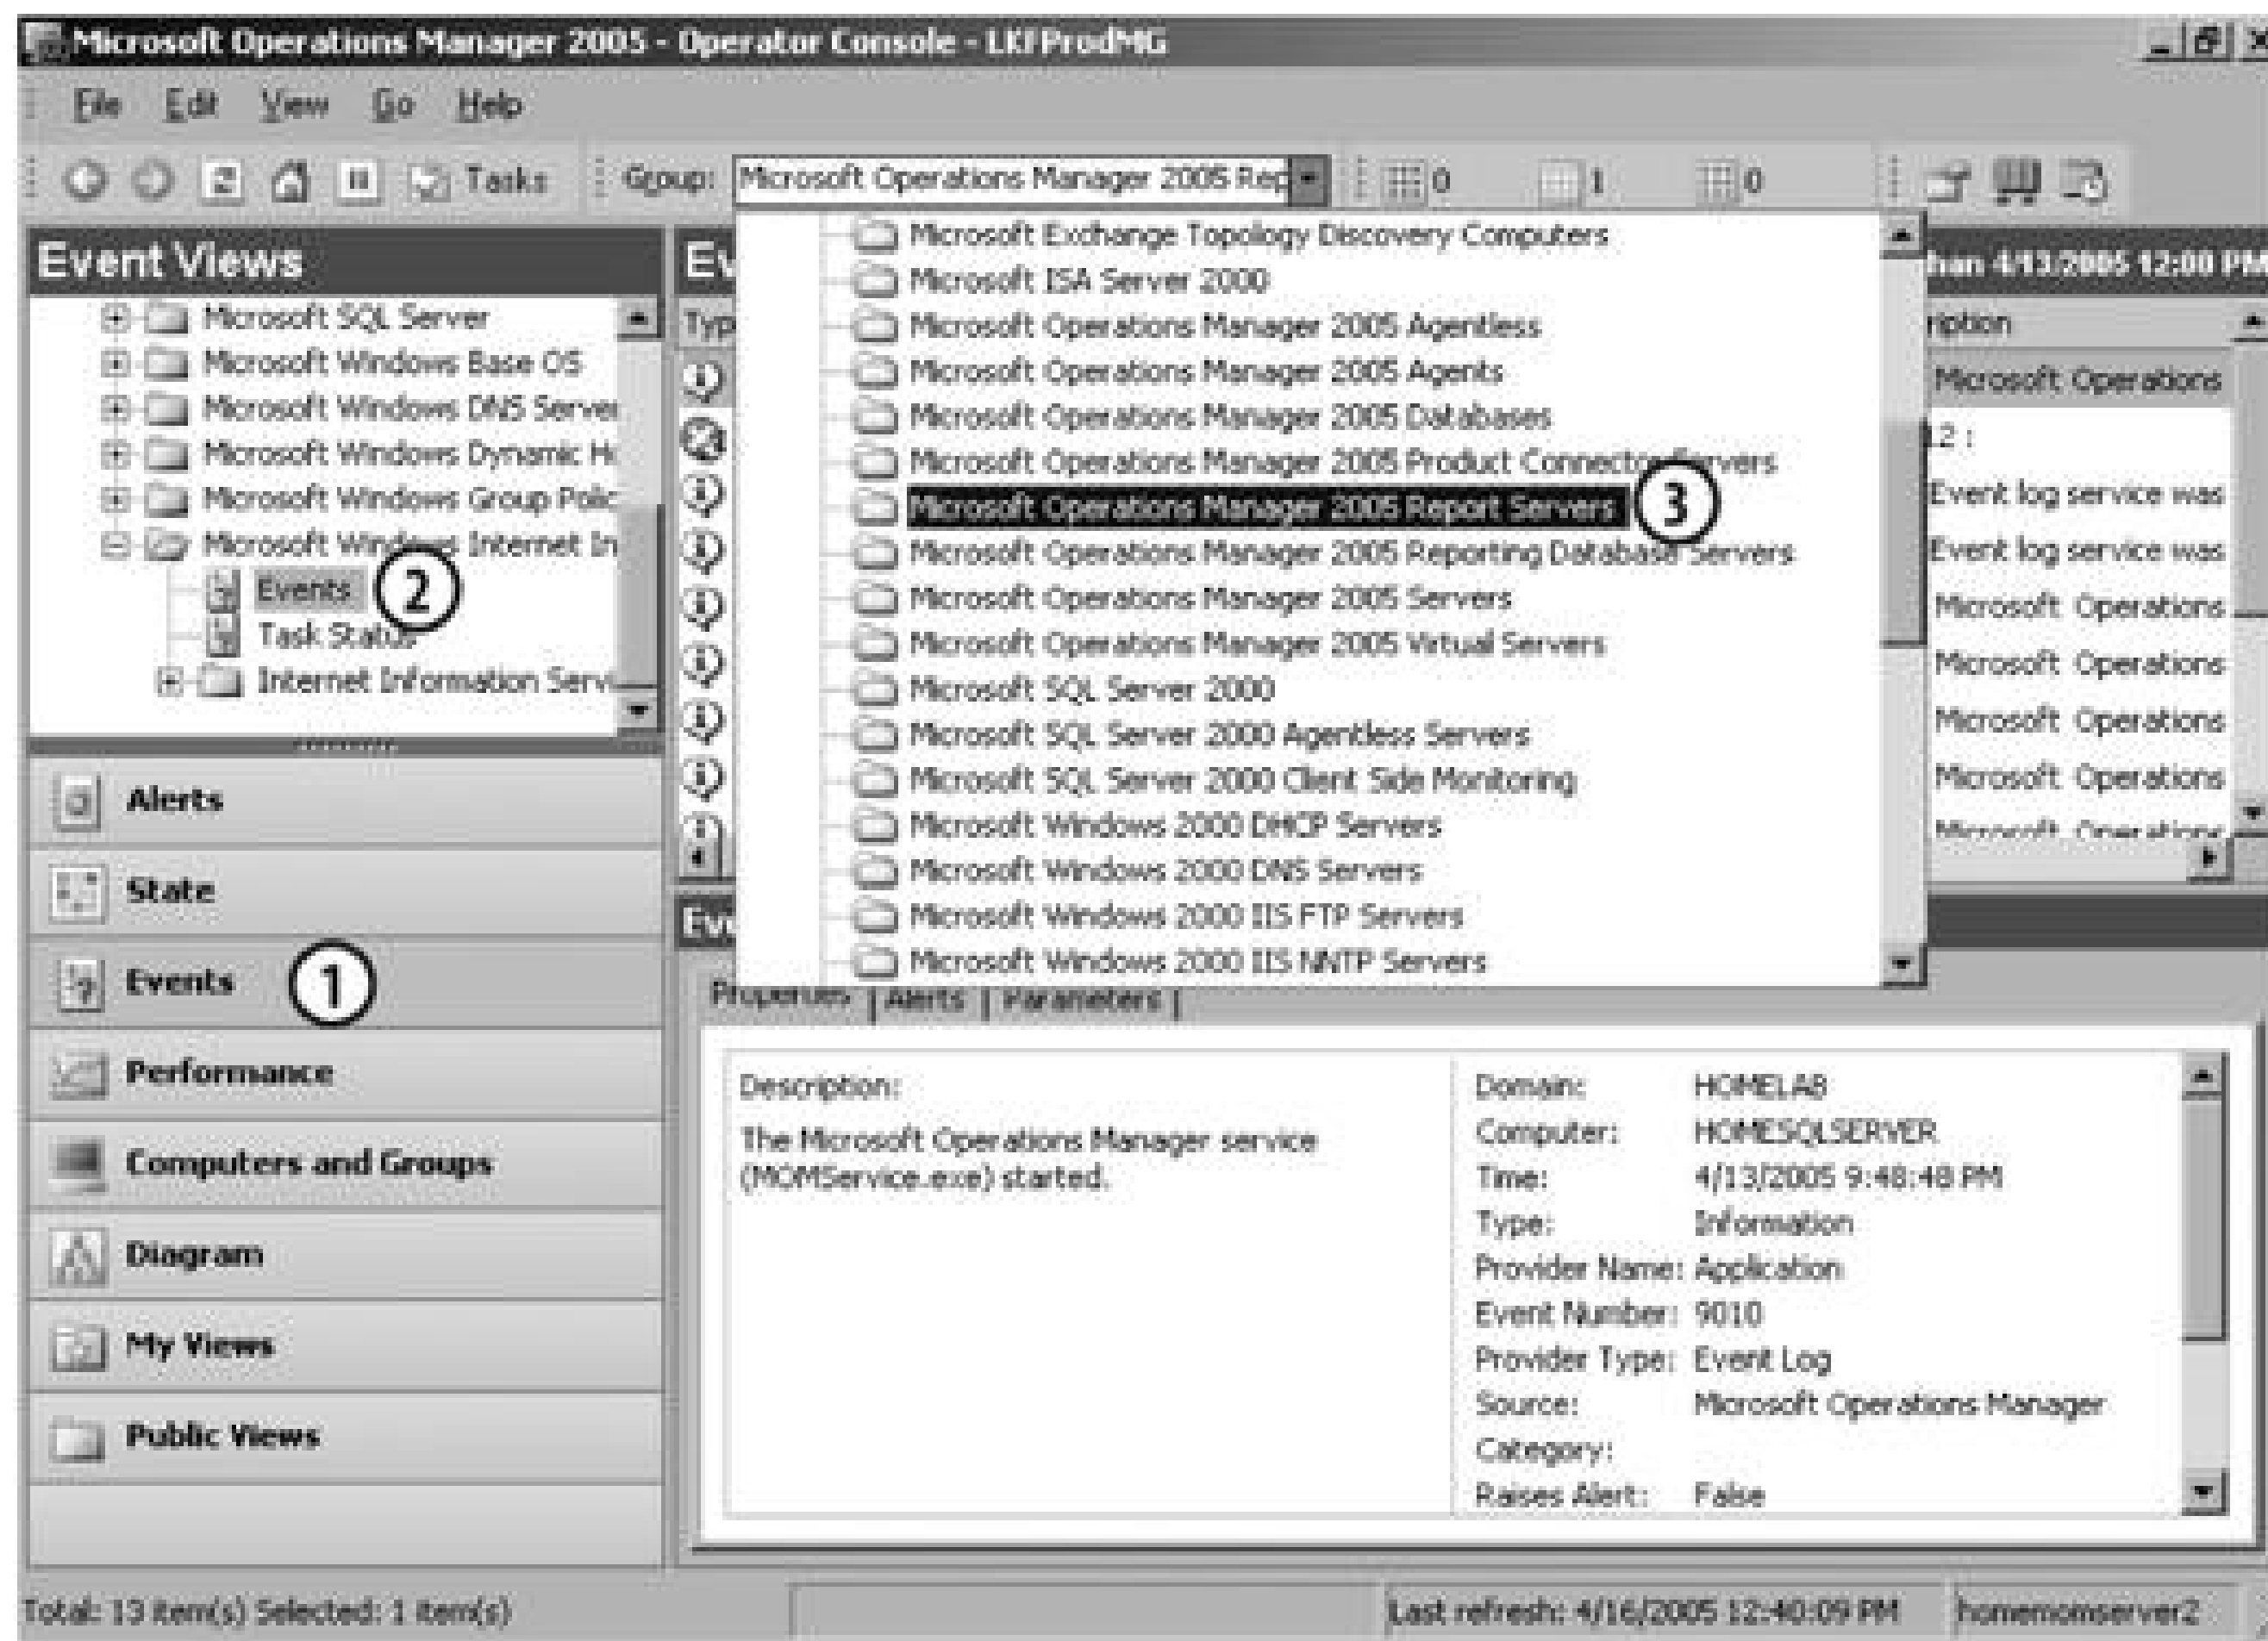
This 123 pattern of navigating has to change only slightly to get the results you want from the Operator console. In the Operator console in [Figure 6-12](#), points 1 and 2 are the same, and point 3 becomes the selection of the desired computer group from the console scope list. When this three-step process is followed, you'll see that the syntax of filter-building has a pattern and a grammatical quality that can be reflected in a sentence.

For example, Leaky Faucet had an intrusion incident on one of the MOM reporting servers that runs IIS. The administrators want to know what events occurred on that IIS server and others without having to go to each machine. To build this filter, the question asked is "What events (view button, point 1 in [Figure 6-12](#)) happened on the IIS servers (view object, point 2 in [Figure 6-12](#)) that are also MOM reporting servers (computer group, point 3 in [Figure 6-12](#))?"

Figure 6-11. These four panes in Outlook 2003 serve the exact same function in the Operator console

Figure 6-12. The controls that use are used to build a filter in the

## operator's console



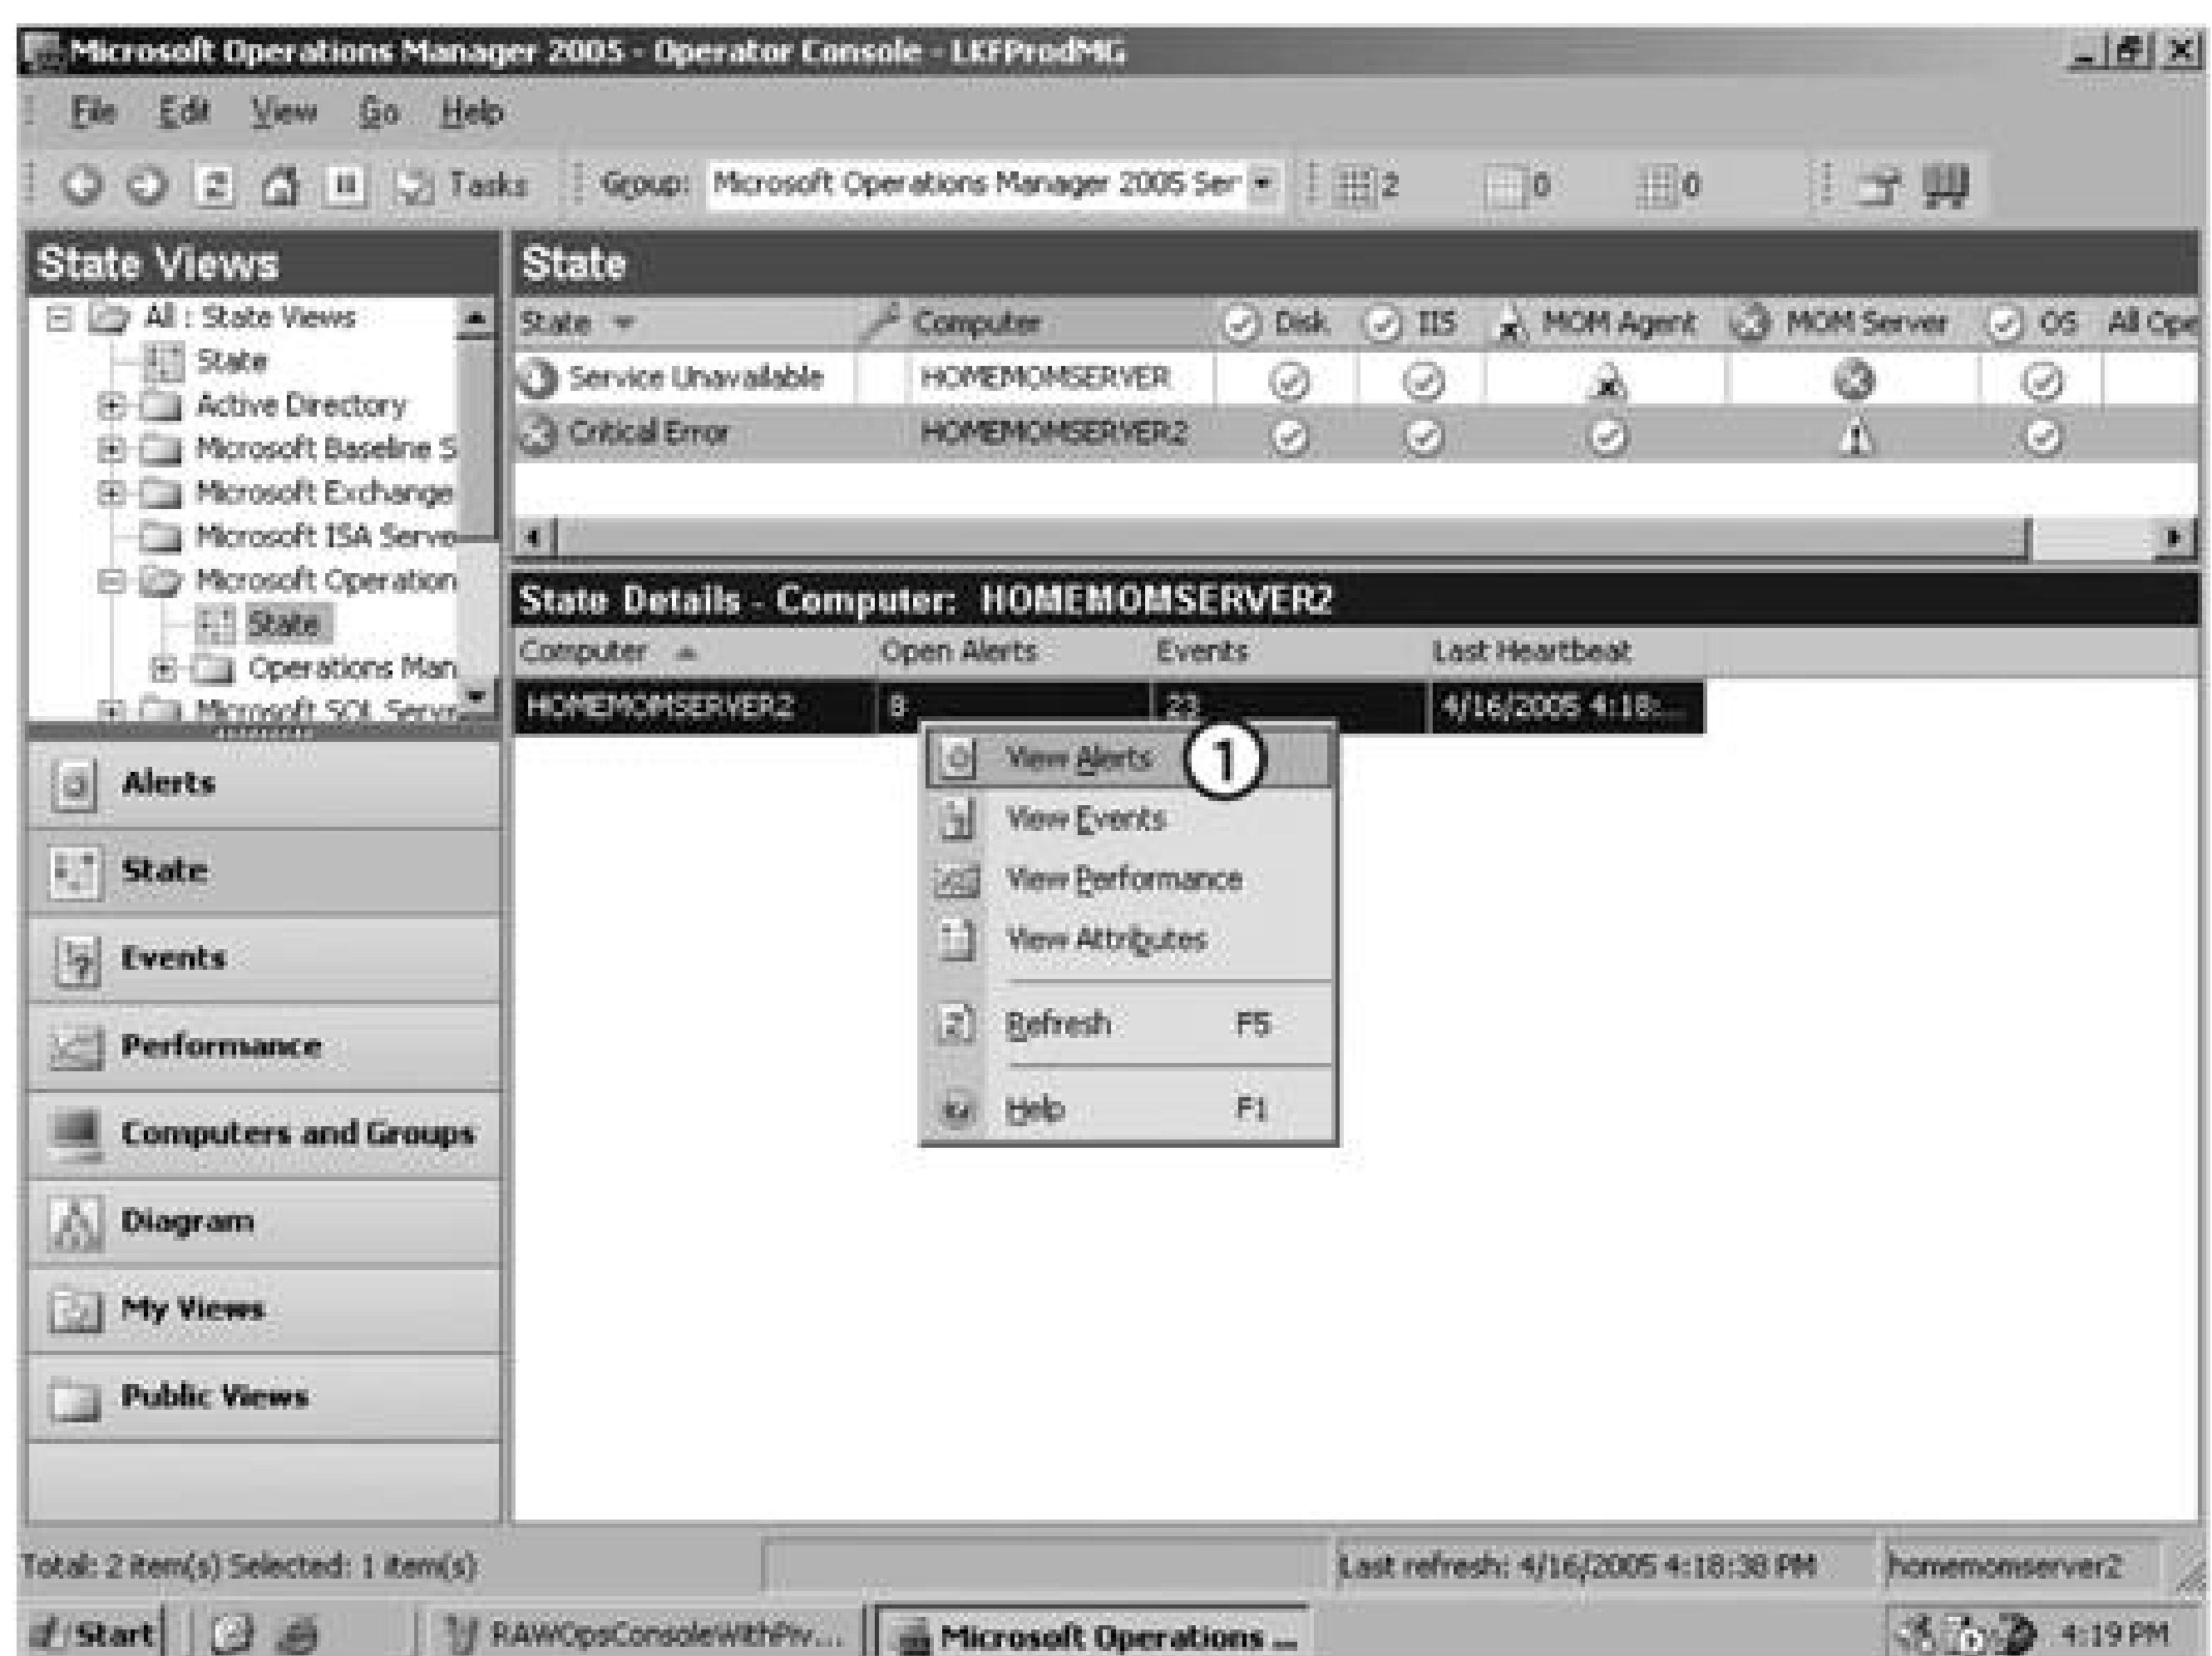
All of the events that were raised in the Windows IIS group that are also in the MOM 2005 reporting server computer group are brought forth. This method works for all view types.

Once you have the desired data in the results pane, you can drill into the details of it, manipulate it (if it is alert data), run tasks on the machine that has the focus, and use that machine as a pivot point to access other types of information that pertain to that machine. For example, say you want to know the current state of your MOM management servers. You would first select the State view button, then the State object for all Microsoft Operations Managers, and finally refine the filter by selecting the Microsoft Operations Manager Server 2005 Servers computer group.

For the Leaky Faucet management group, this process returns the state of all roles on *homemomserver* and *homemomserver2* computers. Max can then pivot (point 1 in [Figure 6-13](#)) into other data types by right-clicking on the computer of interest and bringing up the context menu. *homemomserver2* has been selected here.

Figure 6-13. From the State view of *homemomserver2*, you can pivot and choose to view all the alerts for that computer





There are two fundamental skills you need to use the Operator console effectively: you must be able to build a filter to get the required data in a predictable way, and you must learn how to use the information that is available. This is not possible until you know what information is available.

The next section provides an overview of the Operator console views and the data that they focus on. It also guides you through how the views are used in the two types of interactions in the Operator console, namely the MOM-initiated interaction (an alert or state change) and the user-initiated interaction (tell me what happened, when).

## 6.7. Views

Views, as represented by the view buttons, are collections of database queries that share a common theme. Each individual query in a view group extracts the same type of information from the operations database, but all views are all a variation on that theme. For example, alert view queries (also called view objects) all return alert data, but while one does so for alerts related to Active Directory, another does so for Exchange- or MOM-related alerts, and so on. You see these variations in the upper portion of the navigation panel.

To understand how a view query works, open its properties field. An individual view query is made up of a name, optional description, the data type being queried, and the parameters the query is using to filter. A good example of a simple view query is the All Alerts Query. [Figure 6-14](#) shows the properties of the All Alerts View - Alerts query.

Figure 6-14. The properties of the All Alerts View - Alerts

This alert query is looking for alerts with a specific value in the resolution state field. In this case, it is the Not Equals Resolved value. By clicking on the Not Equals Resolved text, you are taken to a dialog box where you specify the desired resolution state or states.

All view queries follow this same pattern and range from the very simple (see [Figure 6-14](#)) to the very complex, which include conditional statements and time ranges. Every view query in the

Operator console is editable, but it is not recommended that you modify the default queries. Instead, copy one of the existing view queries that is similar to the one that you want to create and paste it into the My Views group. Then open its properties and modify them to meet your needs. When you are getting started, this is a great way to learn the logic of the queries without the danger of affecting the default view queries that are being actively used.

## 6.7.1. Alerts view group

MOM initiates the interaction with the alerts view group, along with the state and diagram view groups. You, as a consumer of MOM information, can sit and watch the results pane for updates if you so choose. The types of alerts that you will see depend on the view query and the computer group from the scope console that is selected. Included in this view group are the Service Level Exception view queries. These are populated by regular alerts that have been in any of the given resolution states (except Resolved) for longer than the configured time (see the ['Alert Resolution States'](#) section in [Chapter 5](#)).

This is the only view type in which you can manipulate field values as necessary in the process of investigating and resolving an individual alert. The life cycle of investigating and resolving an alert was covered in ["The Life of a MOM 2005 Alert"](#) section in [Chapter 1](#).

## 6.7.2. State view group

Alerts are great for giving very specific information about one issue, usually on just one machine, and they are invaluable in the troubleshooting process. Because of their specificity, alerts can't give you a broad view of the overall health of your environment and they don't let you see the relationships between components, computers, and computer groups. This is what state views do. In a single glance, a state view will tell you the health of an individual computer or an entire computer group and its subgroups. It will tell you the health of each of the instances of a role on the computer (e.g., the C: drive is an instance of the disk role on a computer). Like all views, the state view data is frequently updated, once a minute by default. So, when your CIO walks into your cube and asks, "How is Exchange doing right now?" you can finally give an accurate answer.

State view information is generated by event and performance rules that have the "Enable state alert properties" value selected. Not every rule has this value selected, so the vendor is telling you how important a state alert properties enabled rule is as a health indicator merely by enabling it. If a rule has this value enabled, it also tells you that any alerts it generates are worthy of a higher level of attention.

When one of these rules fires an alert that will display in the alerts view, it also changes the state of the object that it is monitoring, which goes into the state calculation for the object instance, role, computer, and computer group. Rules that affect state calculations don't just change the state in a one-way fashion, such as from good to bad. Built into these types of rules is logic that changes the state from bad back to good when the appropriate criteria are satisfied. For example, there is a rule that monitors the DNS server service (*dns.exe*). It is located in the Microsoft Windows DNS Server Windows 2003 DNS State Monitoring and Service Discovery rule group and it is named DNS State Monitoring. If the DNS service changes from a started to a stopped state, this event rule fires an error alert and updates the state of the DNS role for that server. When that service changes from stopped to started, the same rule updates the state value of the DNS service instance for the DNS



role on that server to indicate that it is running again. The whole cycle is reflected in the state view by the icon for the DNS service changing from a green check mark to a yellow warning triangle and back again to the green check mark.

MOM also initiates interaction in the state view; this is seen in the constant updating of the state values. [Figure 6-15](#) shows the Microsoft Operations Manager view query of the MOM 2005 management server's computer group. The navigation pane and the tasks pane have been hidden for the sake of clarity.

Figure 6-15. Microsoft Operations Manager state view of the MOM 2005 management servers

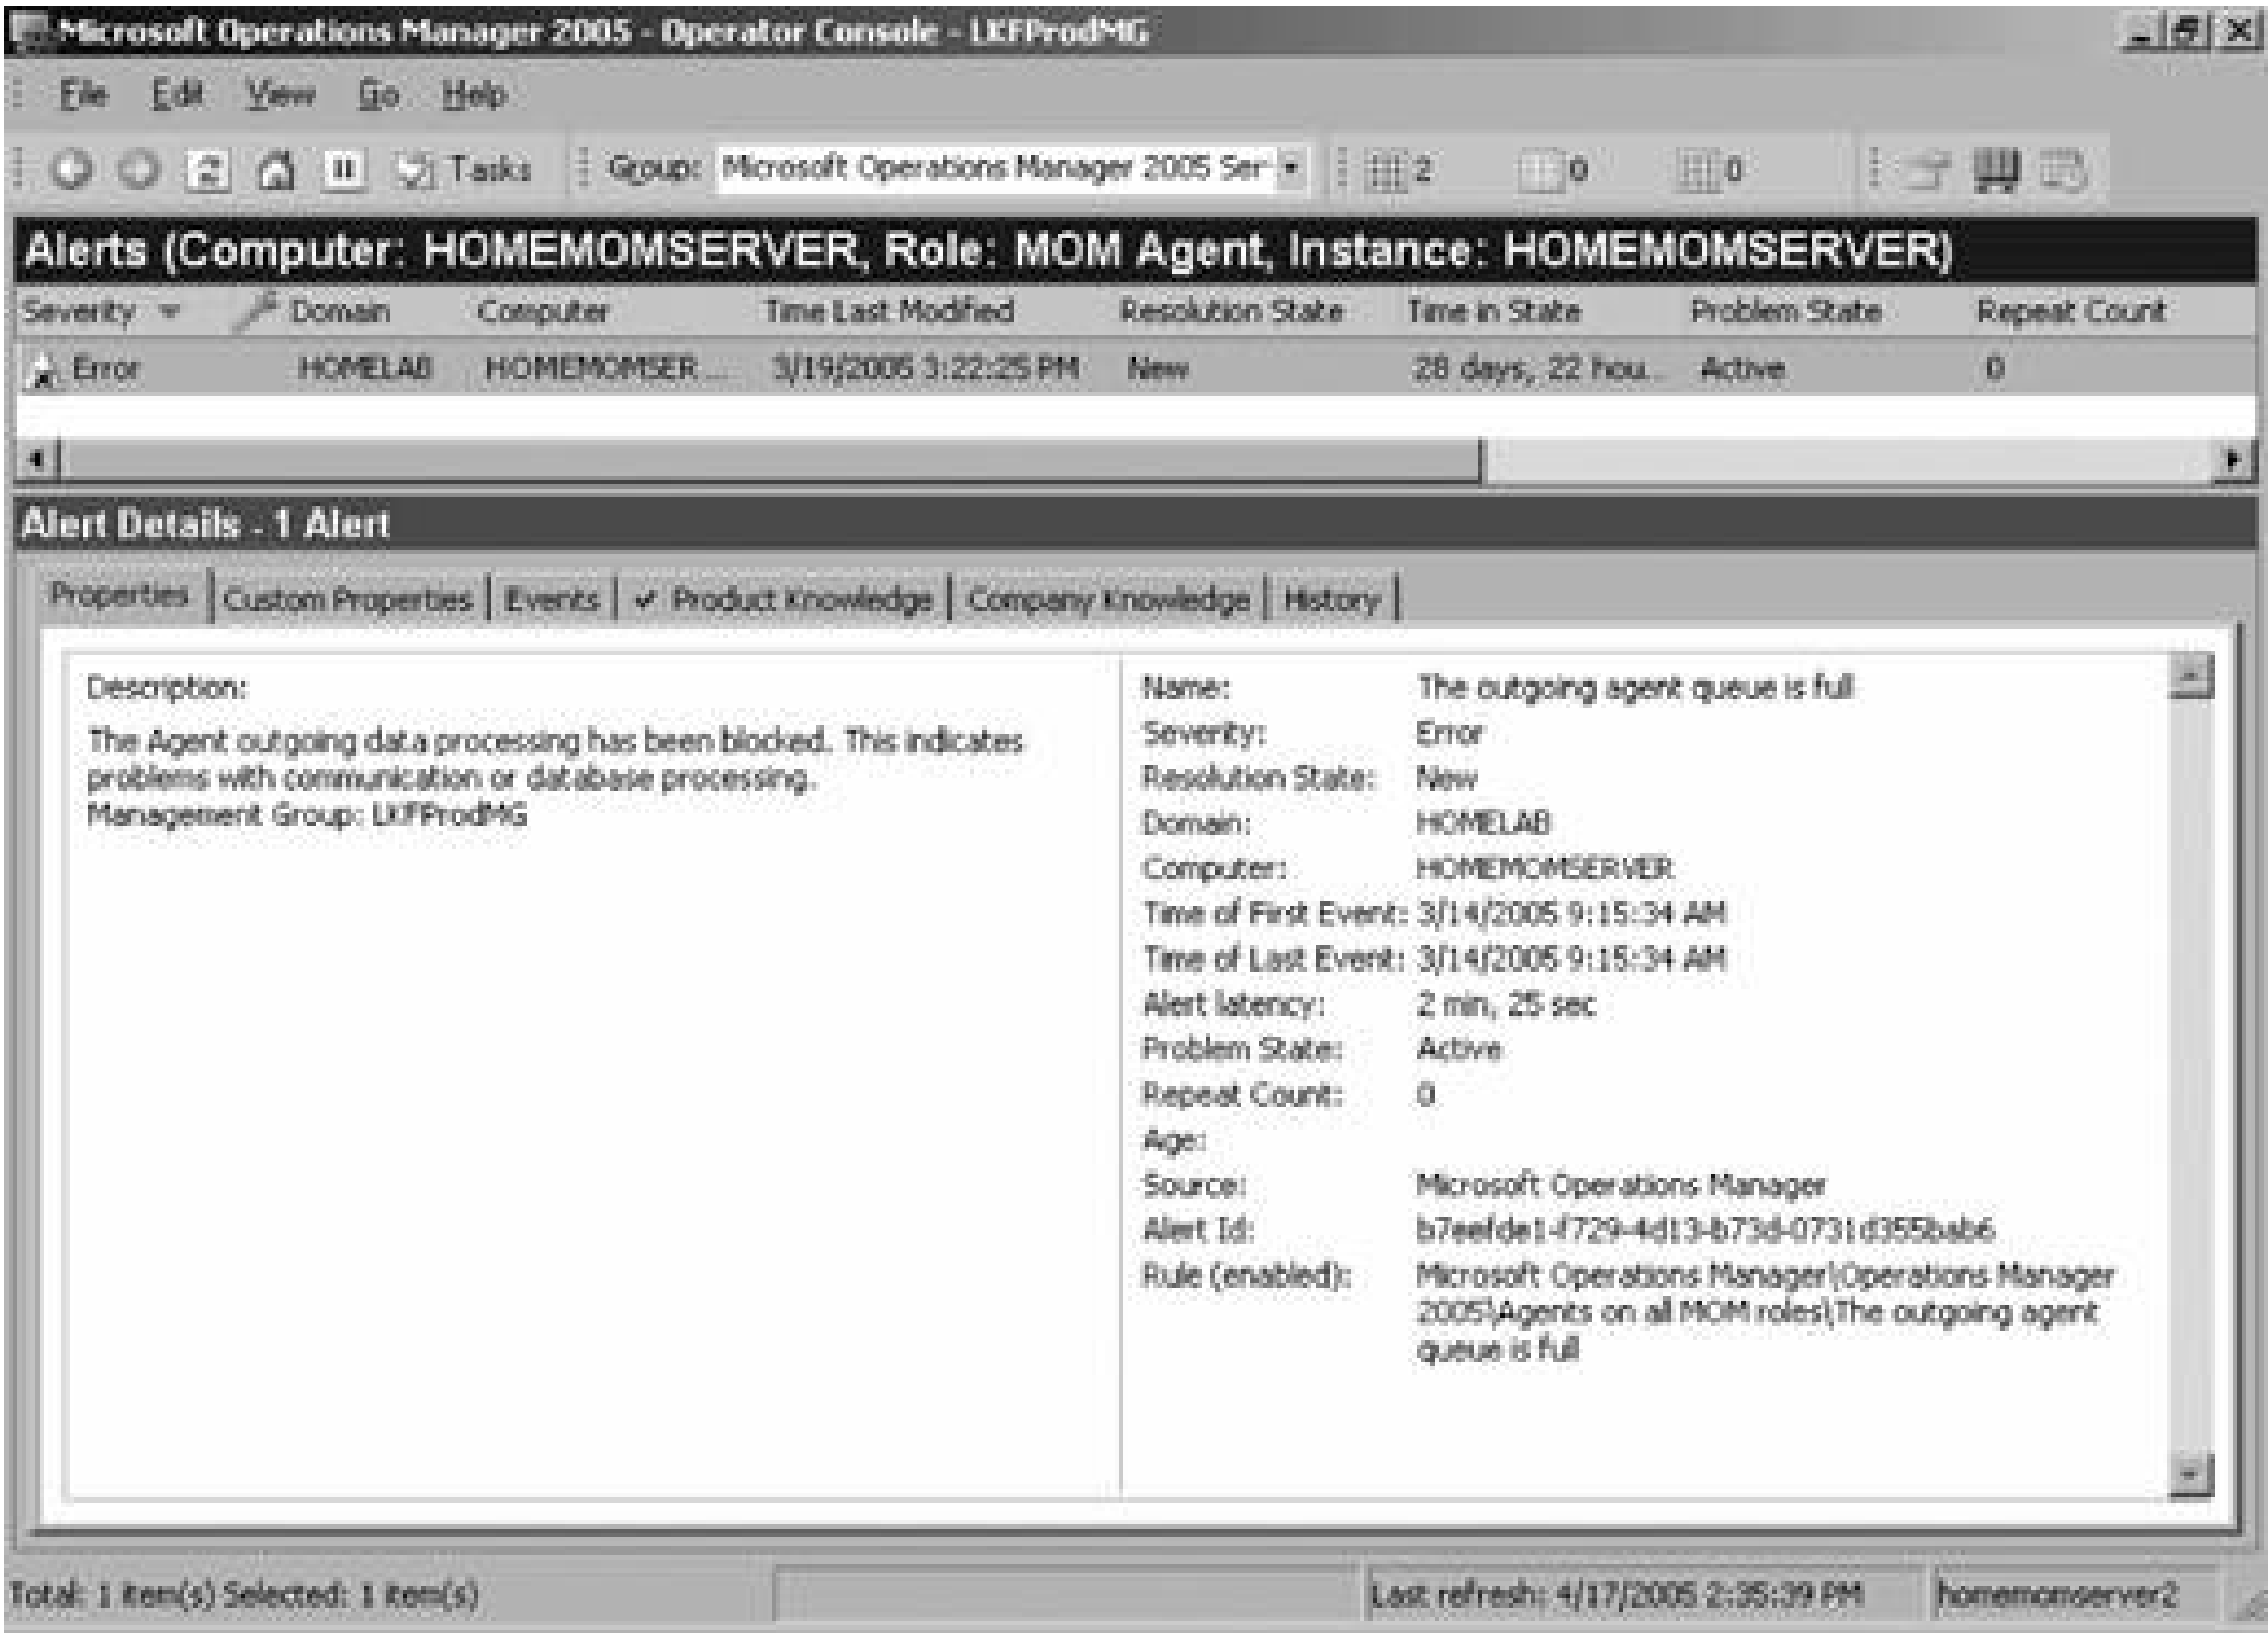


This state view query has returned the state of the roles for the two management servers, *homemomserver* and *homemomserver2*. In point 1 in [Figure 6-15](#), MOM Agent's cell for *homemomserver* has been selected. It is showing an error state. Unlike the alert view where a whole row is used to show a single alert and the details pane changes as you shift the focus between rows, in the state view an individual cell is used to show a state value and the contents of the details pane update as the focus shifts from cell to cell in the results pane. In [Figure 6-15](#), the details pane is showing the state details for *homemomserver*, MOM agent role (point 2 in [Figure 6-15](#)). The state of the MOM Agent role is calculated using the state values from the configuration, heartbeat, performance, queues, and WMI components. Of these, it is the MOM Agent queues that are in an error state (point 3 in [Figure 6-15](#)), so troubleshooting efforts should be concentrated there.

To find out more about what is causing this error state for the MOM agent queues on this box, select the queues error cell (point 3 in [Figure 6-15](#)) and right-click to bring up the context menu. You can view alerts, events, performance, and attribute information for this computer. Go to the alert view, which pinpoints the specific alert that the error indicated ([Figure 6-16](#)). From there, follow your regular process for troubleshooting the issue that caused the error alert. Once the issue is resolved, the state will automatically be set back to good. The alert is resolved, closing the troubleshooting loop.

Figure 6-16. MOM 2005 management server error for Computer: homemomserver, Role: MOM Agent, Instance: homemomserver





The events view would have shown all of the events for that computer in that computer group and would not have shown the alert that was affecting the health state. You probably would have found the specific event that caused the alert to be triggered, but you would have had to sift through all the events to find it. In other words, the event view creates too much noise.

The performance view contains all the performance monitor values for that computer. Since this error state was for a full queue, I probably could have found the queue that was full (if it had a performance counter object) and troubleshooted the issue, but that would not have been the shortest path to resolution. The view computer attributes choice would simply have displayed all of the attributes and their values that MOM collected for that computer. These are the same attributes (registry values and file versions) that are collected when the computer discovery process kicks off. These attributes sort computers into computer groups.

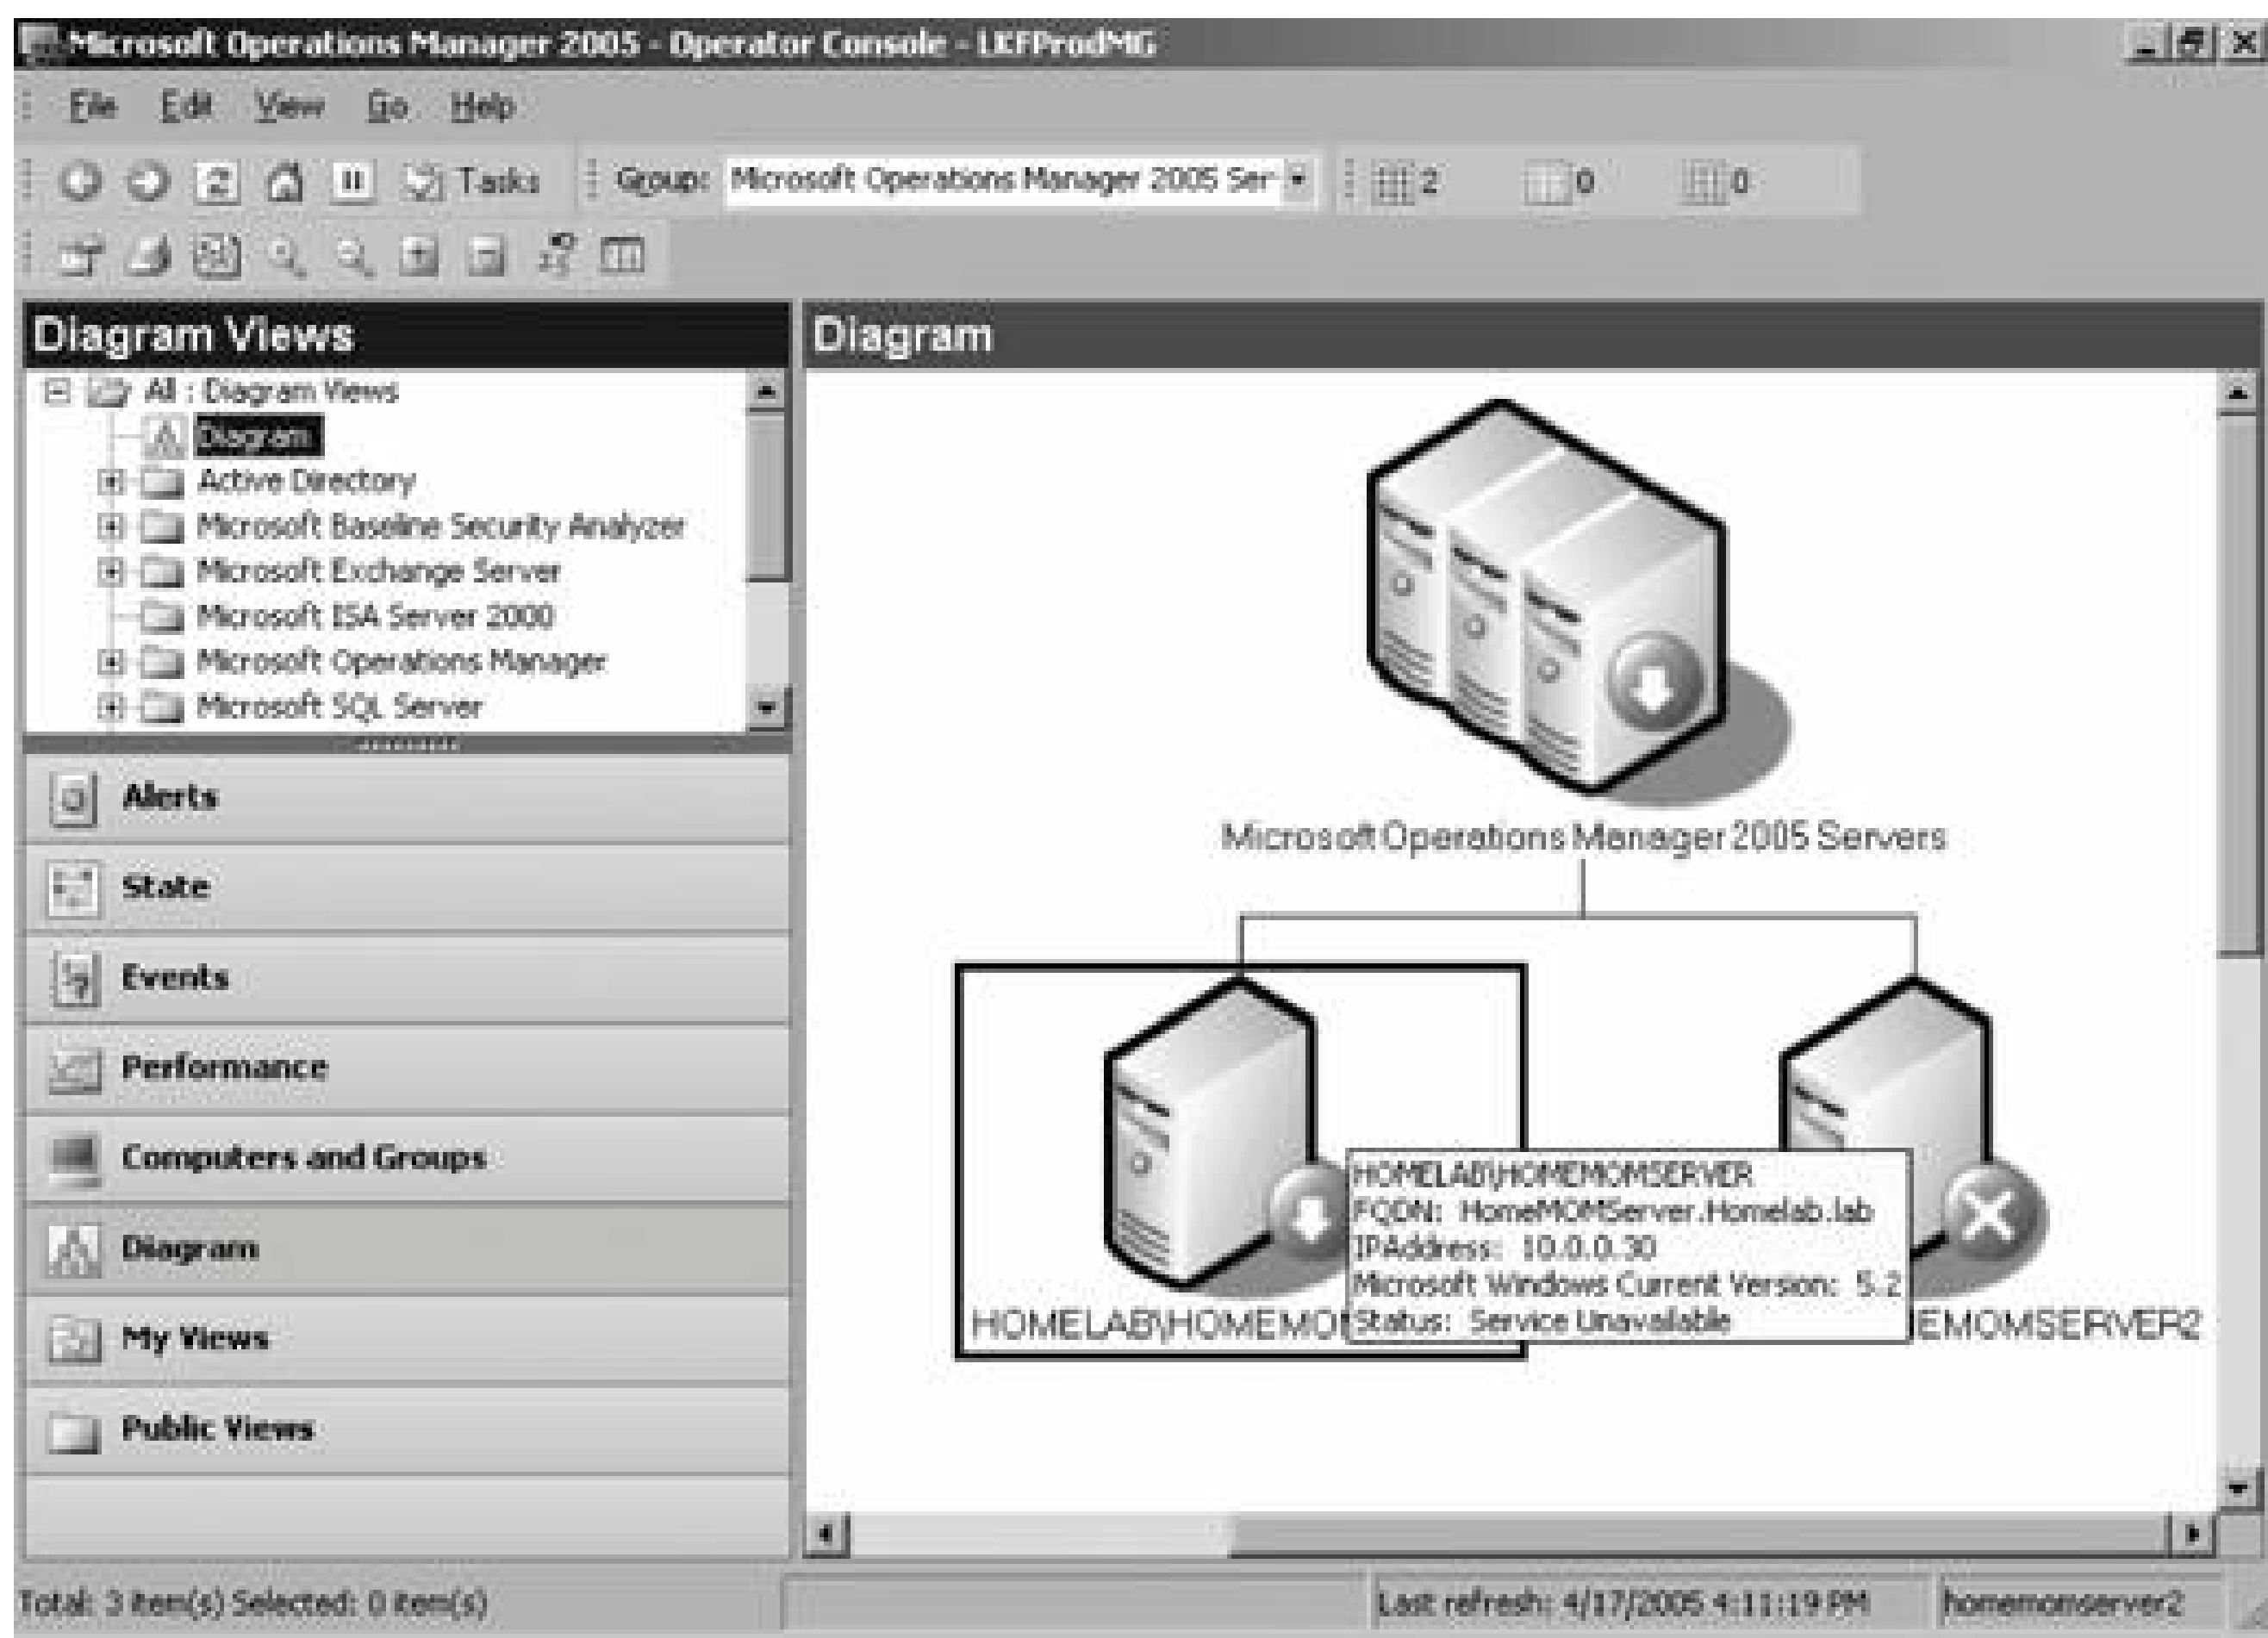
For example, one of the attributes of *homemomserver* is Microsoft Windows IIS Web Server, which is used to detect if the W3SVC is installed. If this registry value (HKLM\System\CurrentControlSet\Services\W3SVC\Start) is set to 2 then this server is sorted into the IIS Web Servers computer group. This is good to know, but it would not have helped in troubleshooting the error state.

### 6.7.3. Diagram view group

The diagram view is another way to present the state value of a computer group in the Operator console. [Figure 6-17](#) shows a very simple state representation of the MOM 2005 management servers computer group and its member servers. This diagram was generated by selecting the diagram view group button and the All: Diagram Views - Diagram view query. The computer group

selection out of the console scope was left at the Microsoft Operations Manager 2005 Servers value.

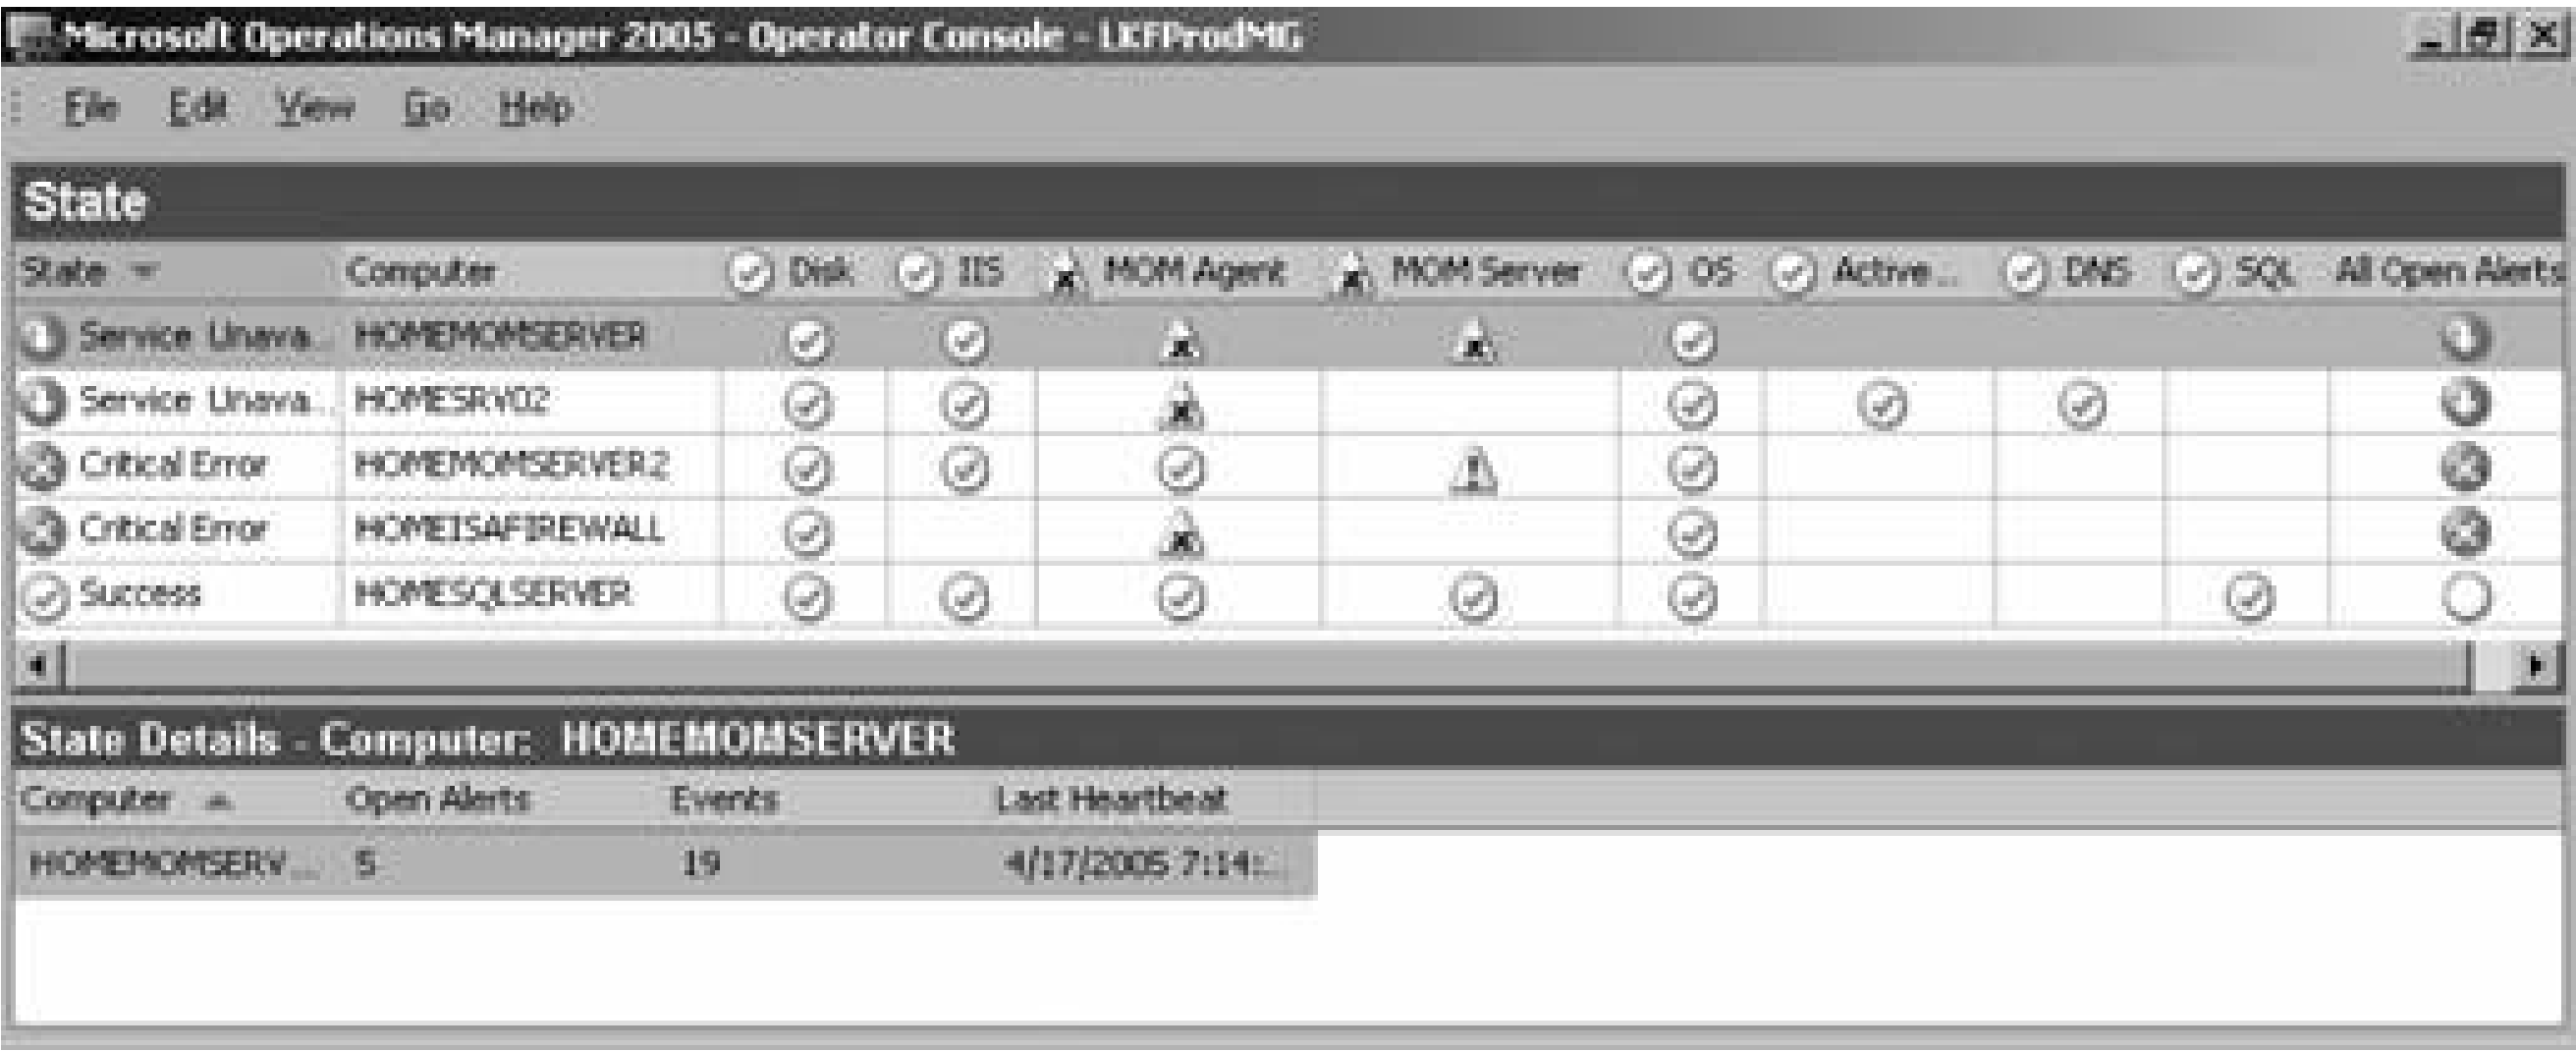
Figure 6-17. Diagram view of the Management Servers computer group



As in the state view, the status indicator is updated dynamically and some details about an object will appear in a pop-up box when you hover the mouse over the object. This helps to make up for the lack of a details pane in diagram views. As with the other views, the context menu for each icon allows you to pivot on the selected object into other data type views.

Getting a legible diagram based on your filter build can be a frustrating experience. This is because diagram views attempt to map out relationships between computers and computer groups, so even in a very simple environment, like the HOMELAB domain, this can yield an unusable drawing. For example, there are only five computers being monitored by this management group, which in the state view is easy to see and digest ([Figure 6-18](#)). Granted, this is only looking at the state for each machine and not at the computer group roll-up.

Figure 6-18. The state view for all computers in the management group is clean and legible



However, if you keep the same filter and switch to the diagram view, you get something that might be usable if printed out on a very large piece of paper from a plotter [Figure 6-19](#)).

Figure 6-19. The diagram view of all computers in the management group is not very usable for daily operations tasks

In this diagram, the icons that show three computers side by side represent computer groups; the individual machines are represented by single computer icons. As you can see, there are many more computer groups than individual machines. The complexity of this diagram is driven by individual



machines belonging to multiple computer groups. Fortunately, you can export these diagrams to Visio for further manipulation and final rendering. The point is you'll have to work harder to get what you want out of a diagram view than any of the other views, so be patient.

There are a few additional controls in the diagram view properties that you can manipulate. To access them, bring up the context menu in the diagram view and select Properties. The available properties are:

#### *Diagram type*

Some examples of diagram type are AD Connection Objects diagrams, AD Connection Objects in Error State, and AD Site links. These types of diagrams incorporate data from your AD Sites and Services structure into MOM. There is also a computer group containment and Exchange diagrams.

#### *Number of levels to display*

For diagrams that are based on computer group containment, you can display the hierarchy a maximum of 10 levels deep.

#### *Routing style*

This refers to how you want any connecting lines laid out. Values include straight, network, center-to-center, and so on.

#### *Node placement style*

Icons in the diagram are referred to as nodes. This property controls the layout of the icons in the diagram (e.g., north/south, east/west, and circular).

#### *Rendering quality*

This is a binary decision. You get either high speed or high quality. By the way, the rendering of the diagram is performed on the machine that is running the console, not the management server.

#### *Background image*

A background image file can be inserted behind the diagram view.

The diagram view sits at the top level in terms of summarizing the state of your environment, with alerts at the bottom, and state views in the middle. The renderings it gives will be the most useful to individuals that like to see things from the 50,000-foot view.

## 6.7.4. Computers and Groups view group



The views in this group focus on some basic demographic statistics of computer groups and the details of the computers in those groups. At the computer group level, the rule tracks data such as the number of open alerts in the group, the types of alerts, and the number of subgroups and computers in the group. For each computer in the group, the rule tracks the number and types of alerts and in the details pane, tracks each computer's attributes, the rule groups applied, the computer groups it belongs to, and the details of each of the roles on that computer.

The information in this rule group is used more for research than for tracking the current state or alerts for a computer or application. For example, [Table 6-1](#) shows the data that is tracked by this view for the Windows 2000 Domain Controllers computer group.

Table 6-1. Data tracked for the Windows 2000 Domain Controllers computer group

Field	Value
State	Service Unavailable
Computer Group	Windows 2000 Domain Controllers
Open Alerts	7
Subgroups	0
All Computers	1
Managed Computers	1
Agentless Managed Computers	0
Maintenance Mode Computers	0
Service Unavailable Computers	1
Security Issue Computers	0
Critical Error Computers	0
Unmanaged Computers	0
New Alerts	7
Error Computers	0
Warning Computers	0
Information Computers	0
Success Computers	0

You can export this information from the Operator console by selecting it and choosing "Copy formatted data" in the context menu. [Table 6-2](#) contains the default fields and values that are tracked for an individual computer, in this case *homesrv02*.

Table 6-2. Fields tracked for homesrv02

Field	Value
State	Service Unavailable
Maintenance Mode	False
Domain	HOMELAB
Name	homesrv02
Last Heartbeat	4/17/2005 9:54:14 PM
New Alerts	7
Service Unavailable	1
Security Issue	0
Critical Error	2
Operating System Type	Windows Primary Domain Controller
Computer ID	ffaf8f2a-a5b0-4959-b4fb-f5ab37a288df
FQDN	homesrv02.Homelab.lab
Error	3
Warning	1
Information	0
Success	0
Maintenance Mode End	1/1/0001 12:00:00 AM
Maintenance Mode User	
Maintenance Mode Reason	

In [Table 6-2](#), the last three entries need a little explaining. Maintenance Mode End has a date and time value assigned by default. As soon as this server is placed into maintenance mode for the first time, this value would be overwritten. Because this server has never been in maintenance mode, the last two entries have no value assigned to them.

One of the main uses you will have for this view is to get the details of the roles on any managed computer. For example, if you need to document the details for all instances of drives in the role of disks on the computer *homesrv02*, you would build a filter following the three-step method:

1. Select the Computers and Groups view group, then open the computers view query in the Active Directory folder.
2. Select MOM Administration Scope. This would place *homesrv02* as the only computer in the results pane (point 1 in [Figure 6-20](#)).
3. Select the Roles tab in the results pane (point 2 in [Figure 6-20](#)) to show the details of the Disk

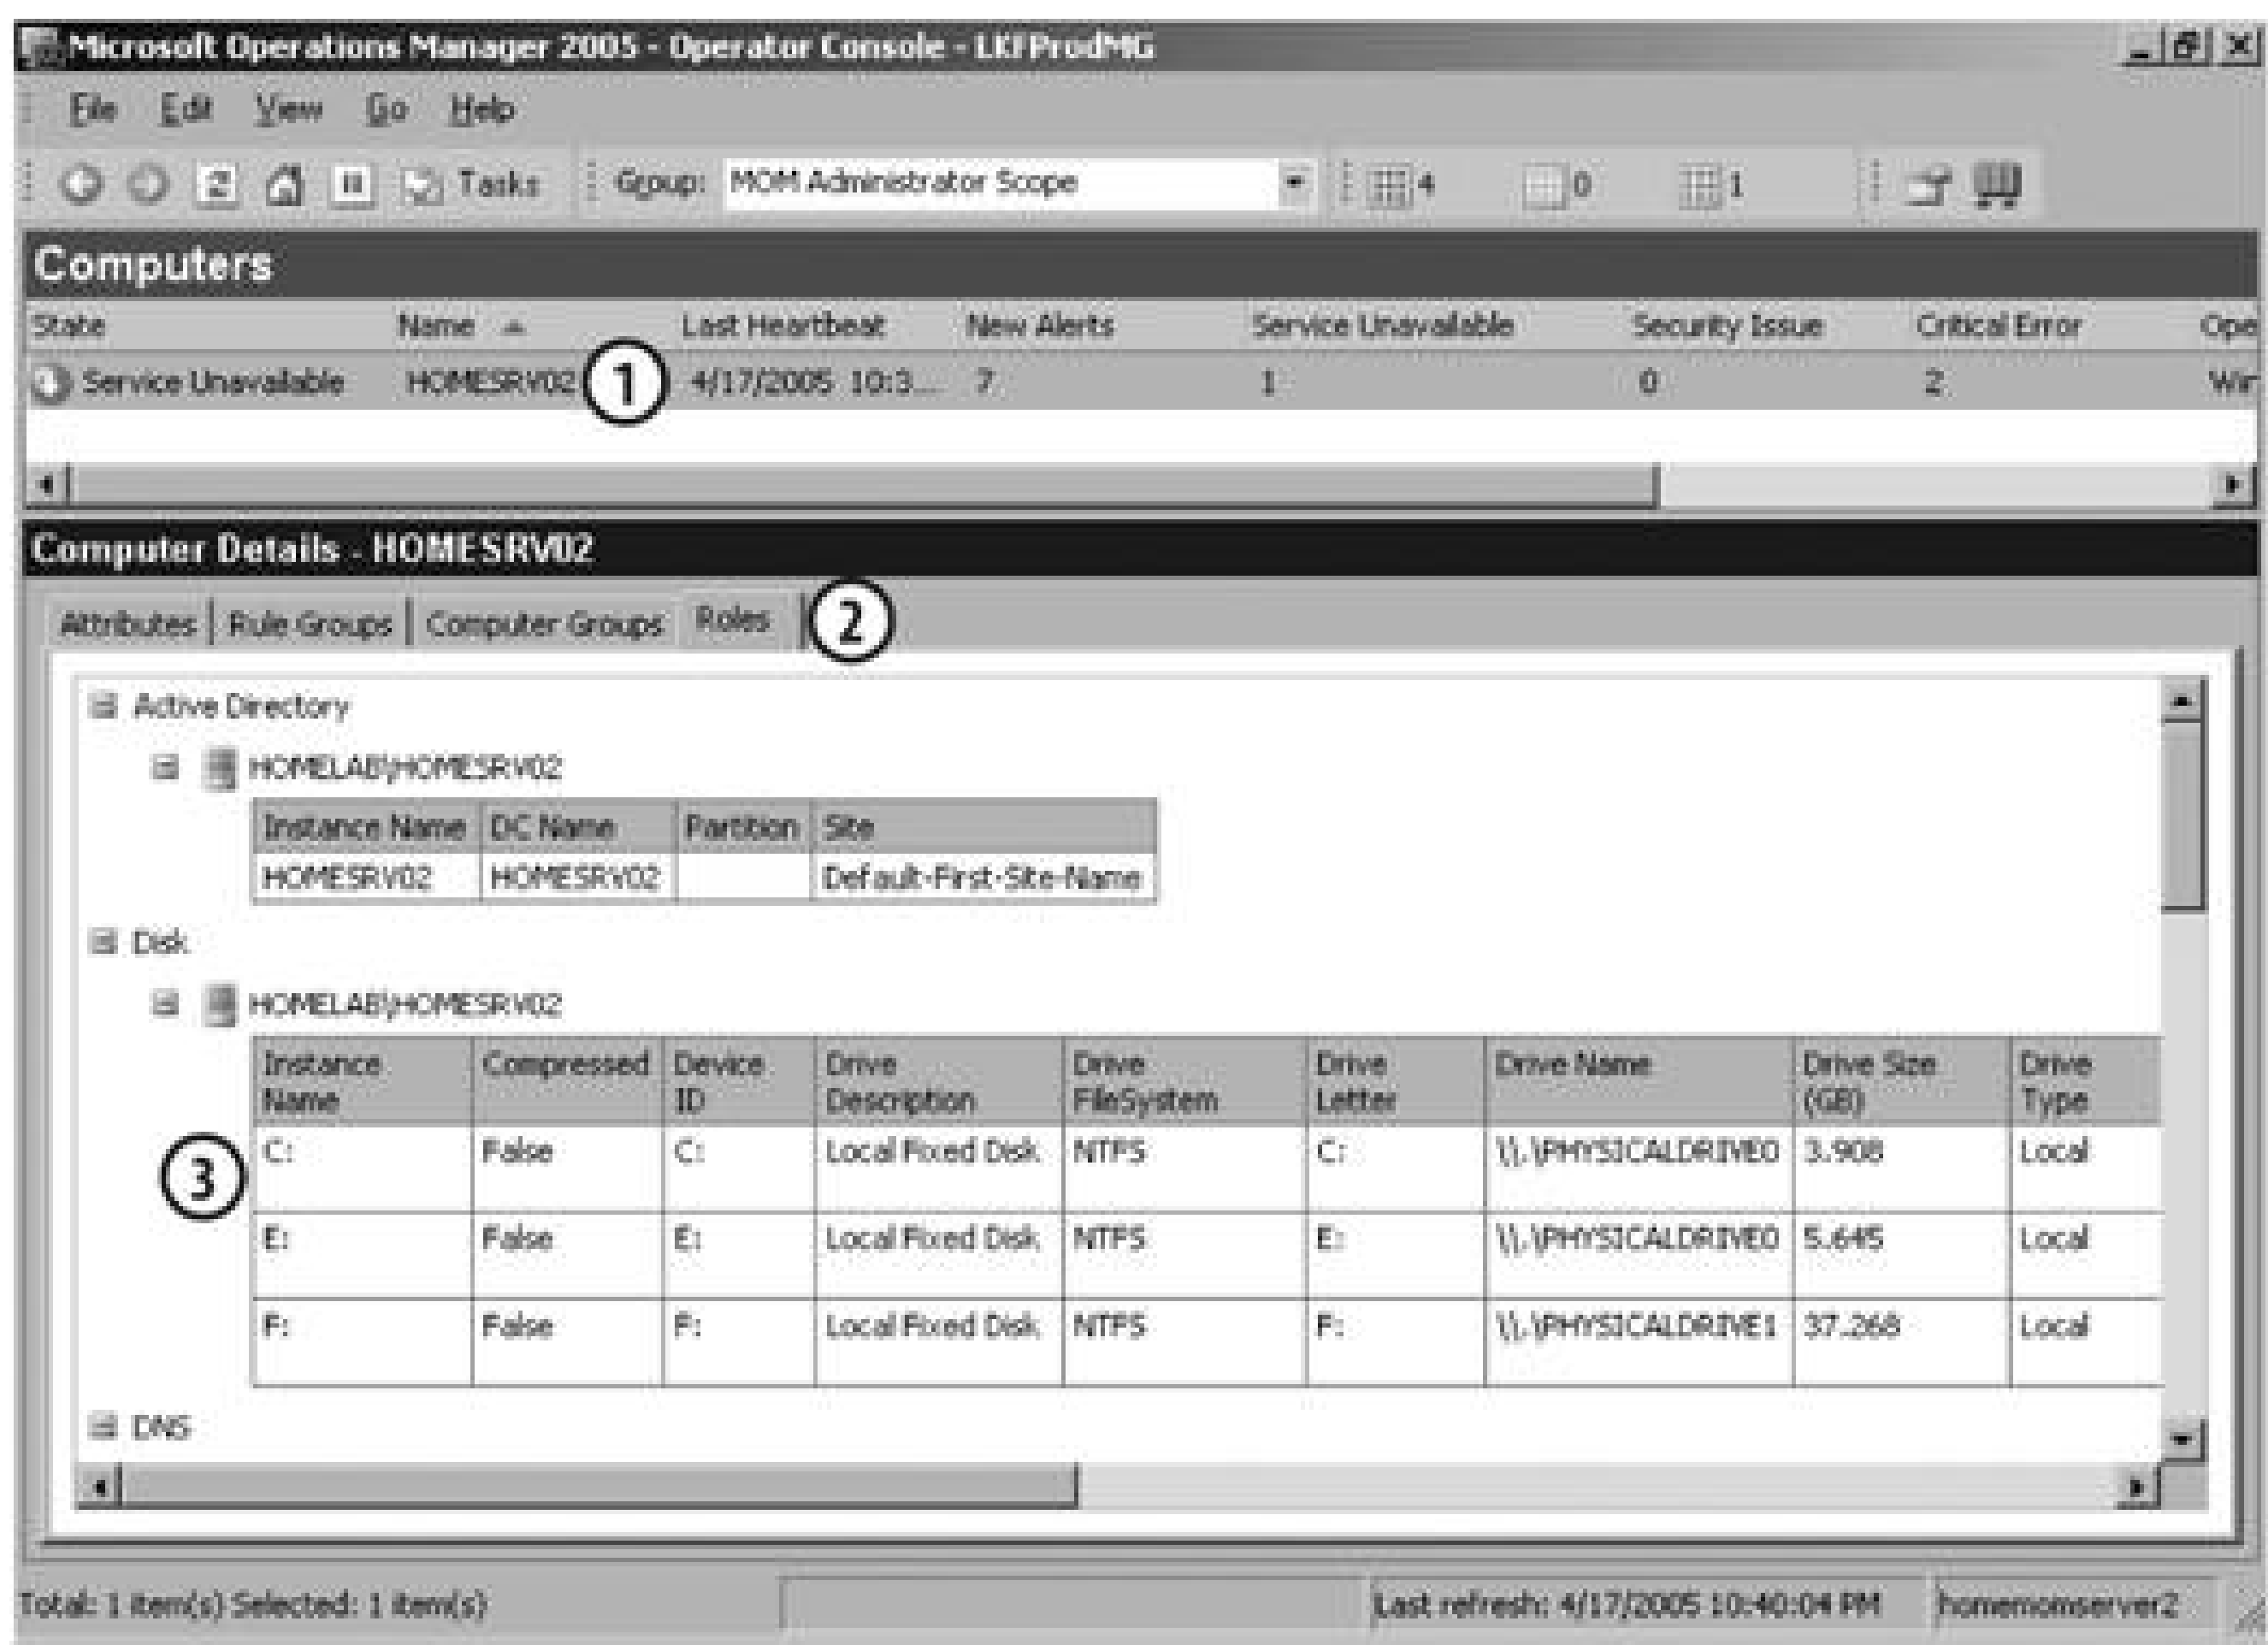
role on that computer (point 3 in [Figure 6-20](#)).

Much of the data in the details pane cannot be viewed, simply because there is so much. To work around this you can choose to select all the data in the Rules tab and copy it to an Excel spreadsheet for further use.

### 6.7.5. Events view group, performance view group , and public views group

These three view groups fall firmly into the "Tell me what was happening on XYZ computer" category for interactions. They are aggregations of the collected event logs and performance monitor counters from managed computers.

Figure 6-20. Disk detail information available in the Computers and Groups view group



The task status events track the milestones for both system- and user-initiated tasks. For example, every day at 2:05 a.m. the computer discovery task kicks off and its success or failure is recorded in the event logs on the management servers and is also reported in the task's status events. The computer discovery task also tracks task status for user-initiated tasks, such as on-demand computer discovery scan tasks in the Tasks pane of the Operator console. For some tasks, and you



have to determine this on a task-by-task basis, the output of the task only displays here. For example, in [Figure 1-9](#), the ping task was run against *homemomserver* and the output was presented in a pop-up box. For other tasks, such as "Test end-to-end monitoring," you must go to the "all tasks view" to see the output.

The performance monitor counter charting, while convenient, is no replacement for using the native performance monitor tool in the OS. With this view, you are limited to graphing the collected information. You do not have the ability to adjust the minimum/maximum values for the graph.

The public views group gathers all of the predefined view queries together into one view group.

## 6.7.6. My Views group

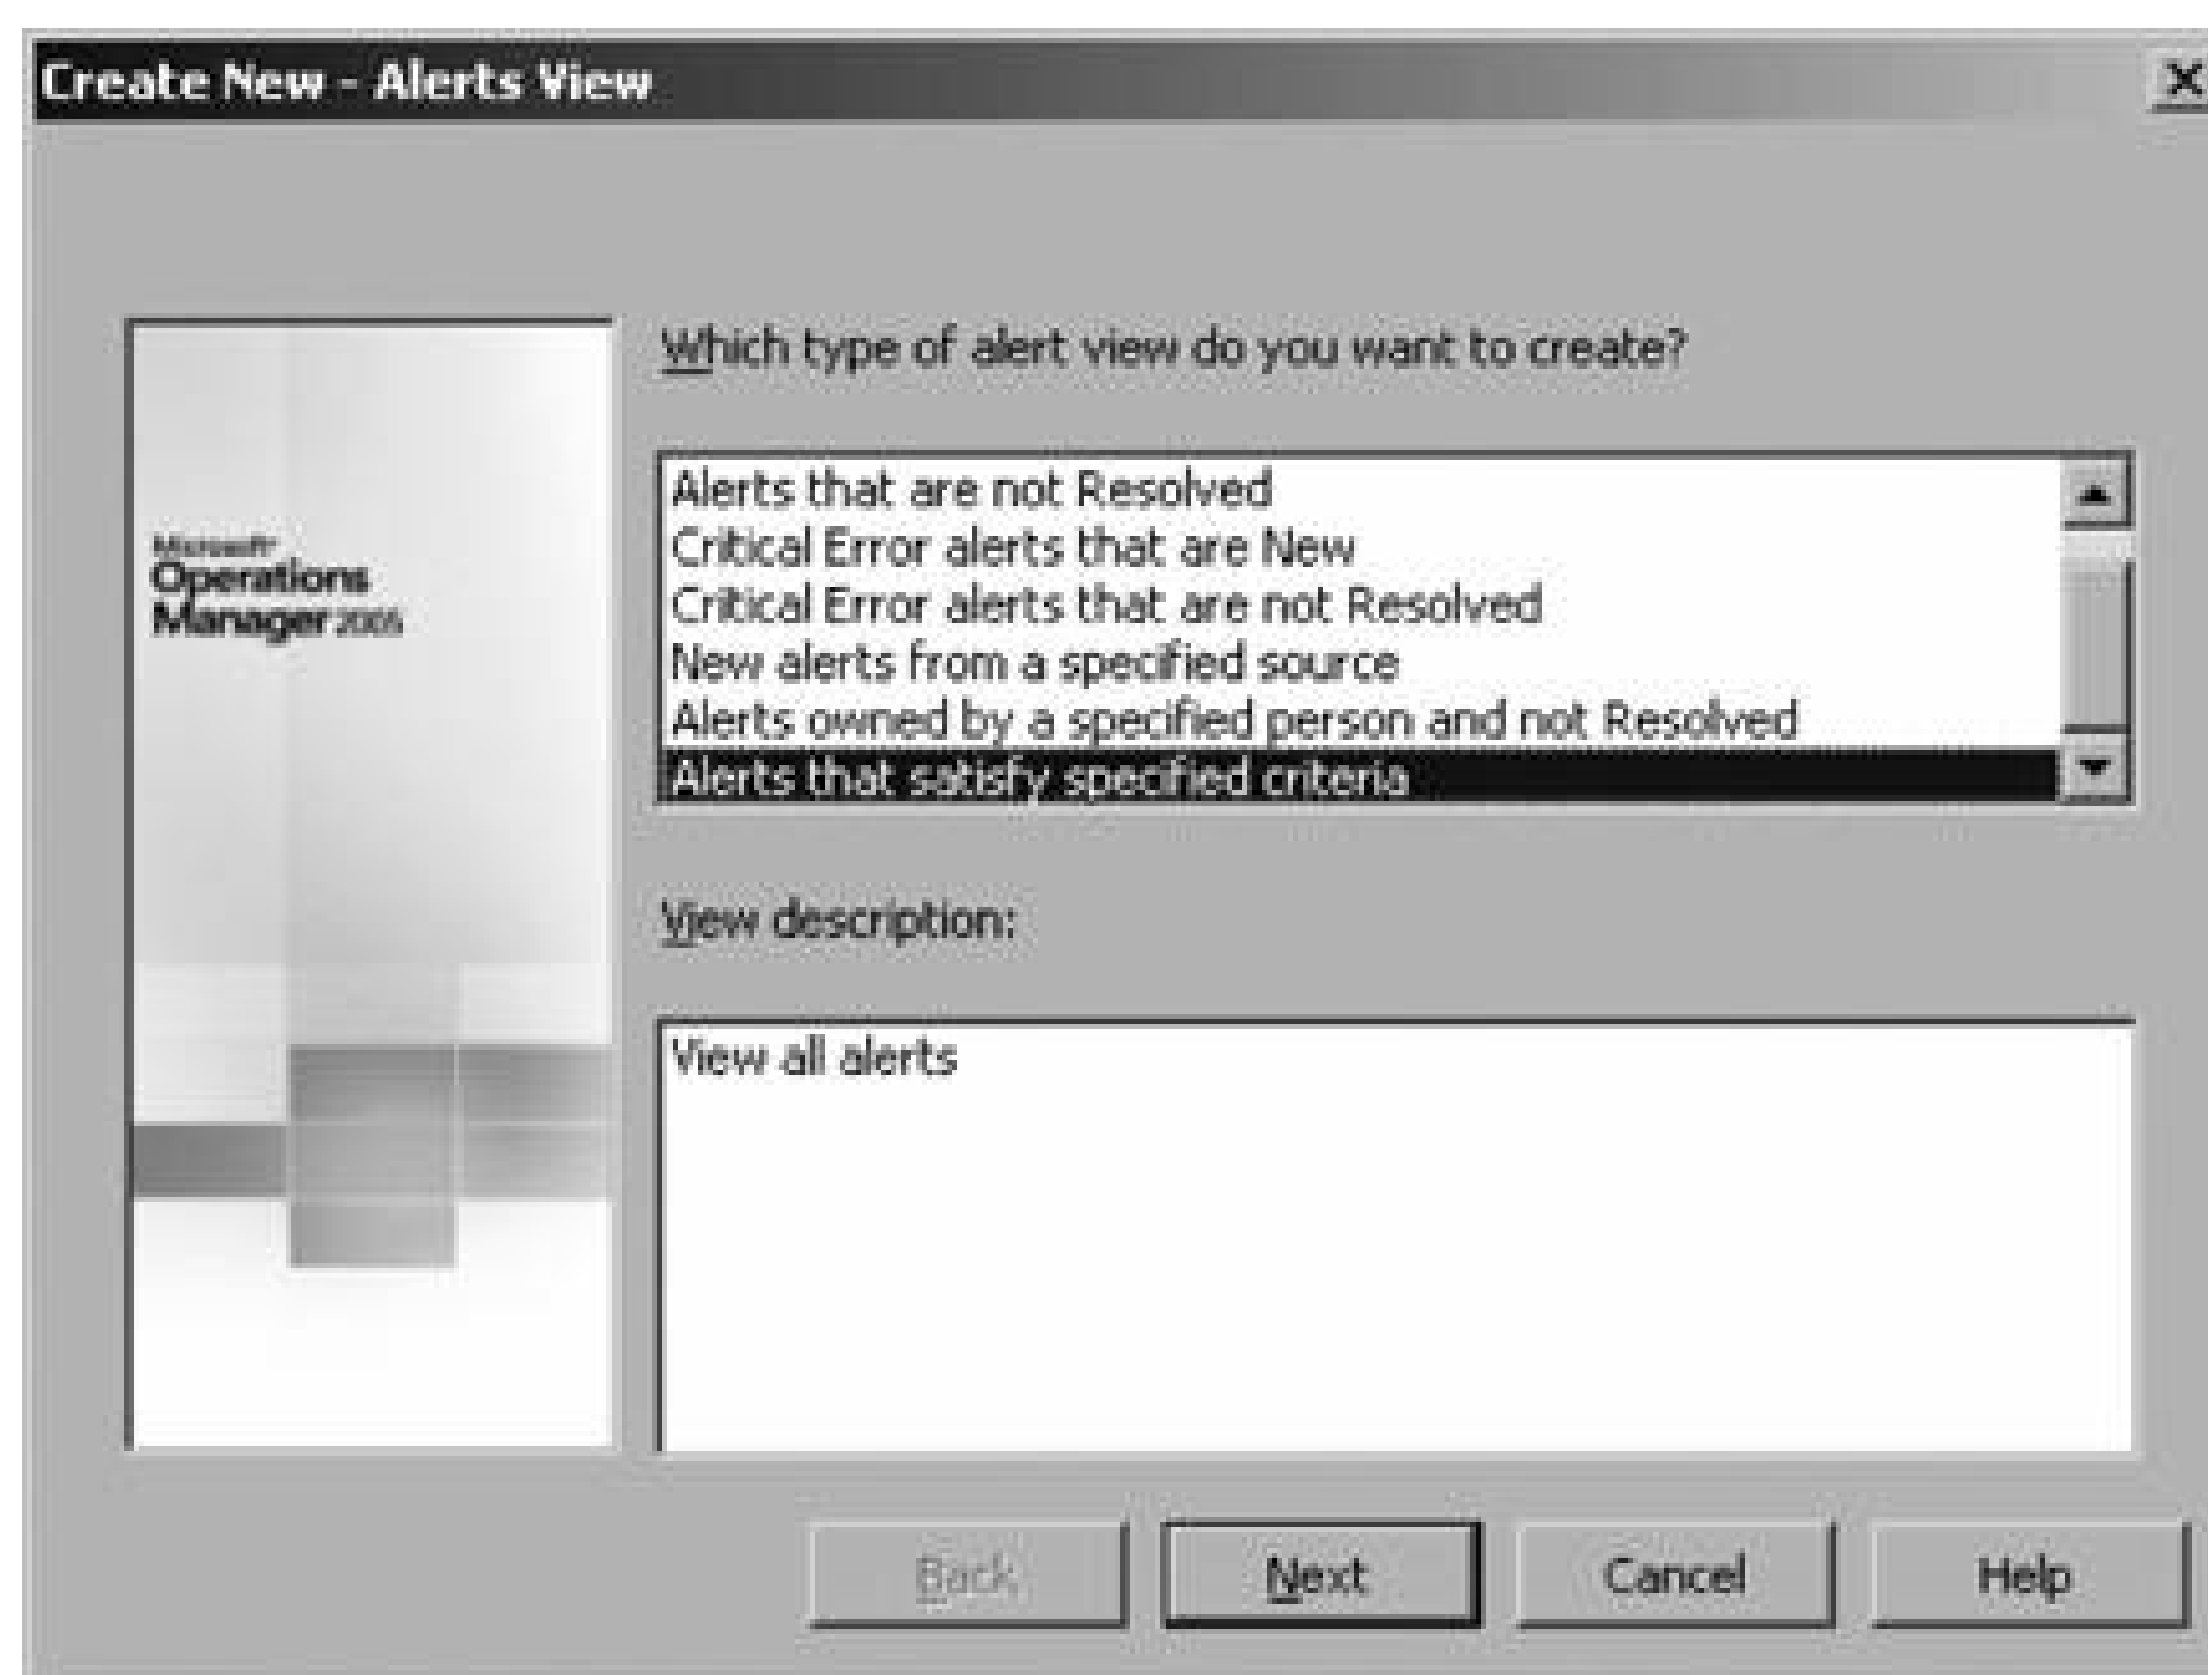
All the views and view queries other than the My Views group are available to all users of the console. And as many as there are, undoubtedly you will want to create ones that are tailored to a particular need at any given time. The My Views group is where you will customize existing views and create new ones from scratch. View queries that are created here are only available to you. To share a view query, it must be copied to any of the publicly viewable view groups.

In the beginning of this chapter, the Leaky Faucet administrator, Max, created a custom computer group called LKF Remote Office Servers, where he placed the remote site servers and added the *homesrv02* and *homesqlserver* servers. Now he needs to create custom views based on that and any other future Leaky Faucet-created computer groups. To make identifying Leaky Faucet-created computer groups easier, he will prepend the names with the LKF tag, just as he does with custom rules. Since Max is new to creating view queries, he will develop the folder and all queries in the My Views group. Once he is satisfied with the functionality, Max will move the queries to the Public Views group.

The first views that Max wants to create are state and alert views. He starts by creating an All LKF Views folder in the My Views group. This is done via the context menu for the All My Views folder. Then, in the All LKF Views folder, via the context menu he selects "Create a new alerts view query." Next, he selects the type of alerts he wants to view ([Figure 6-21](#)). In this case, it is alerts that satisfy specified criteria, which leaves the filter open to all alerts.

Figure 6-21. Create Alerts View wizard selection for alert types

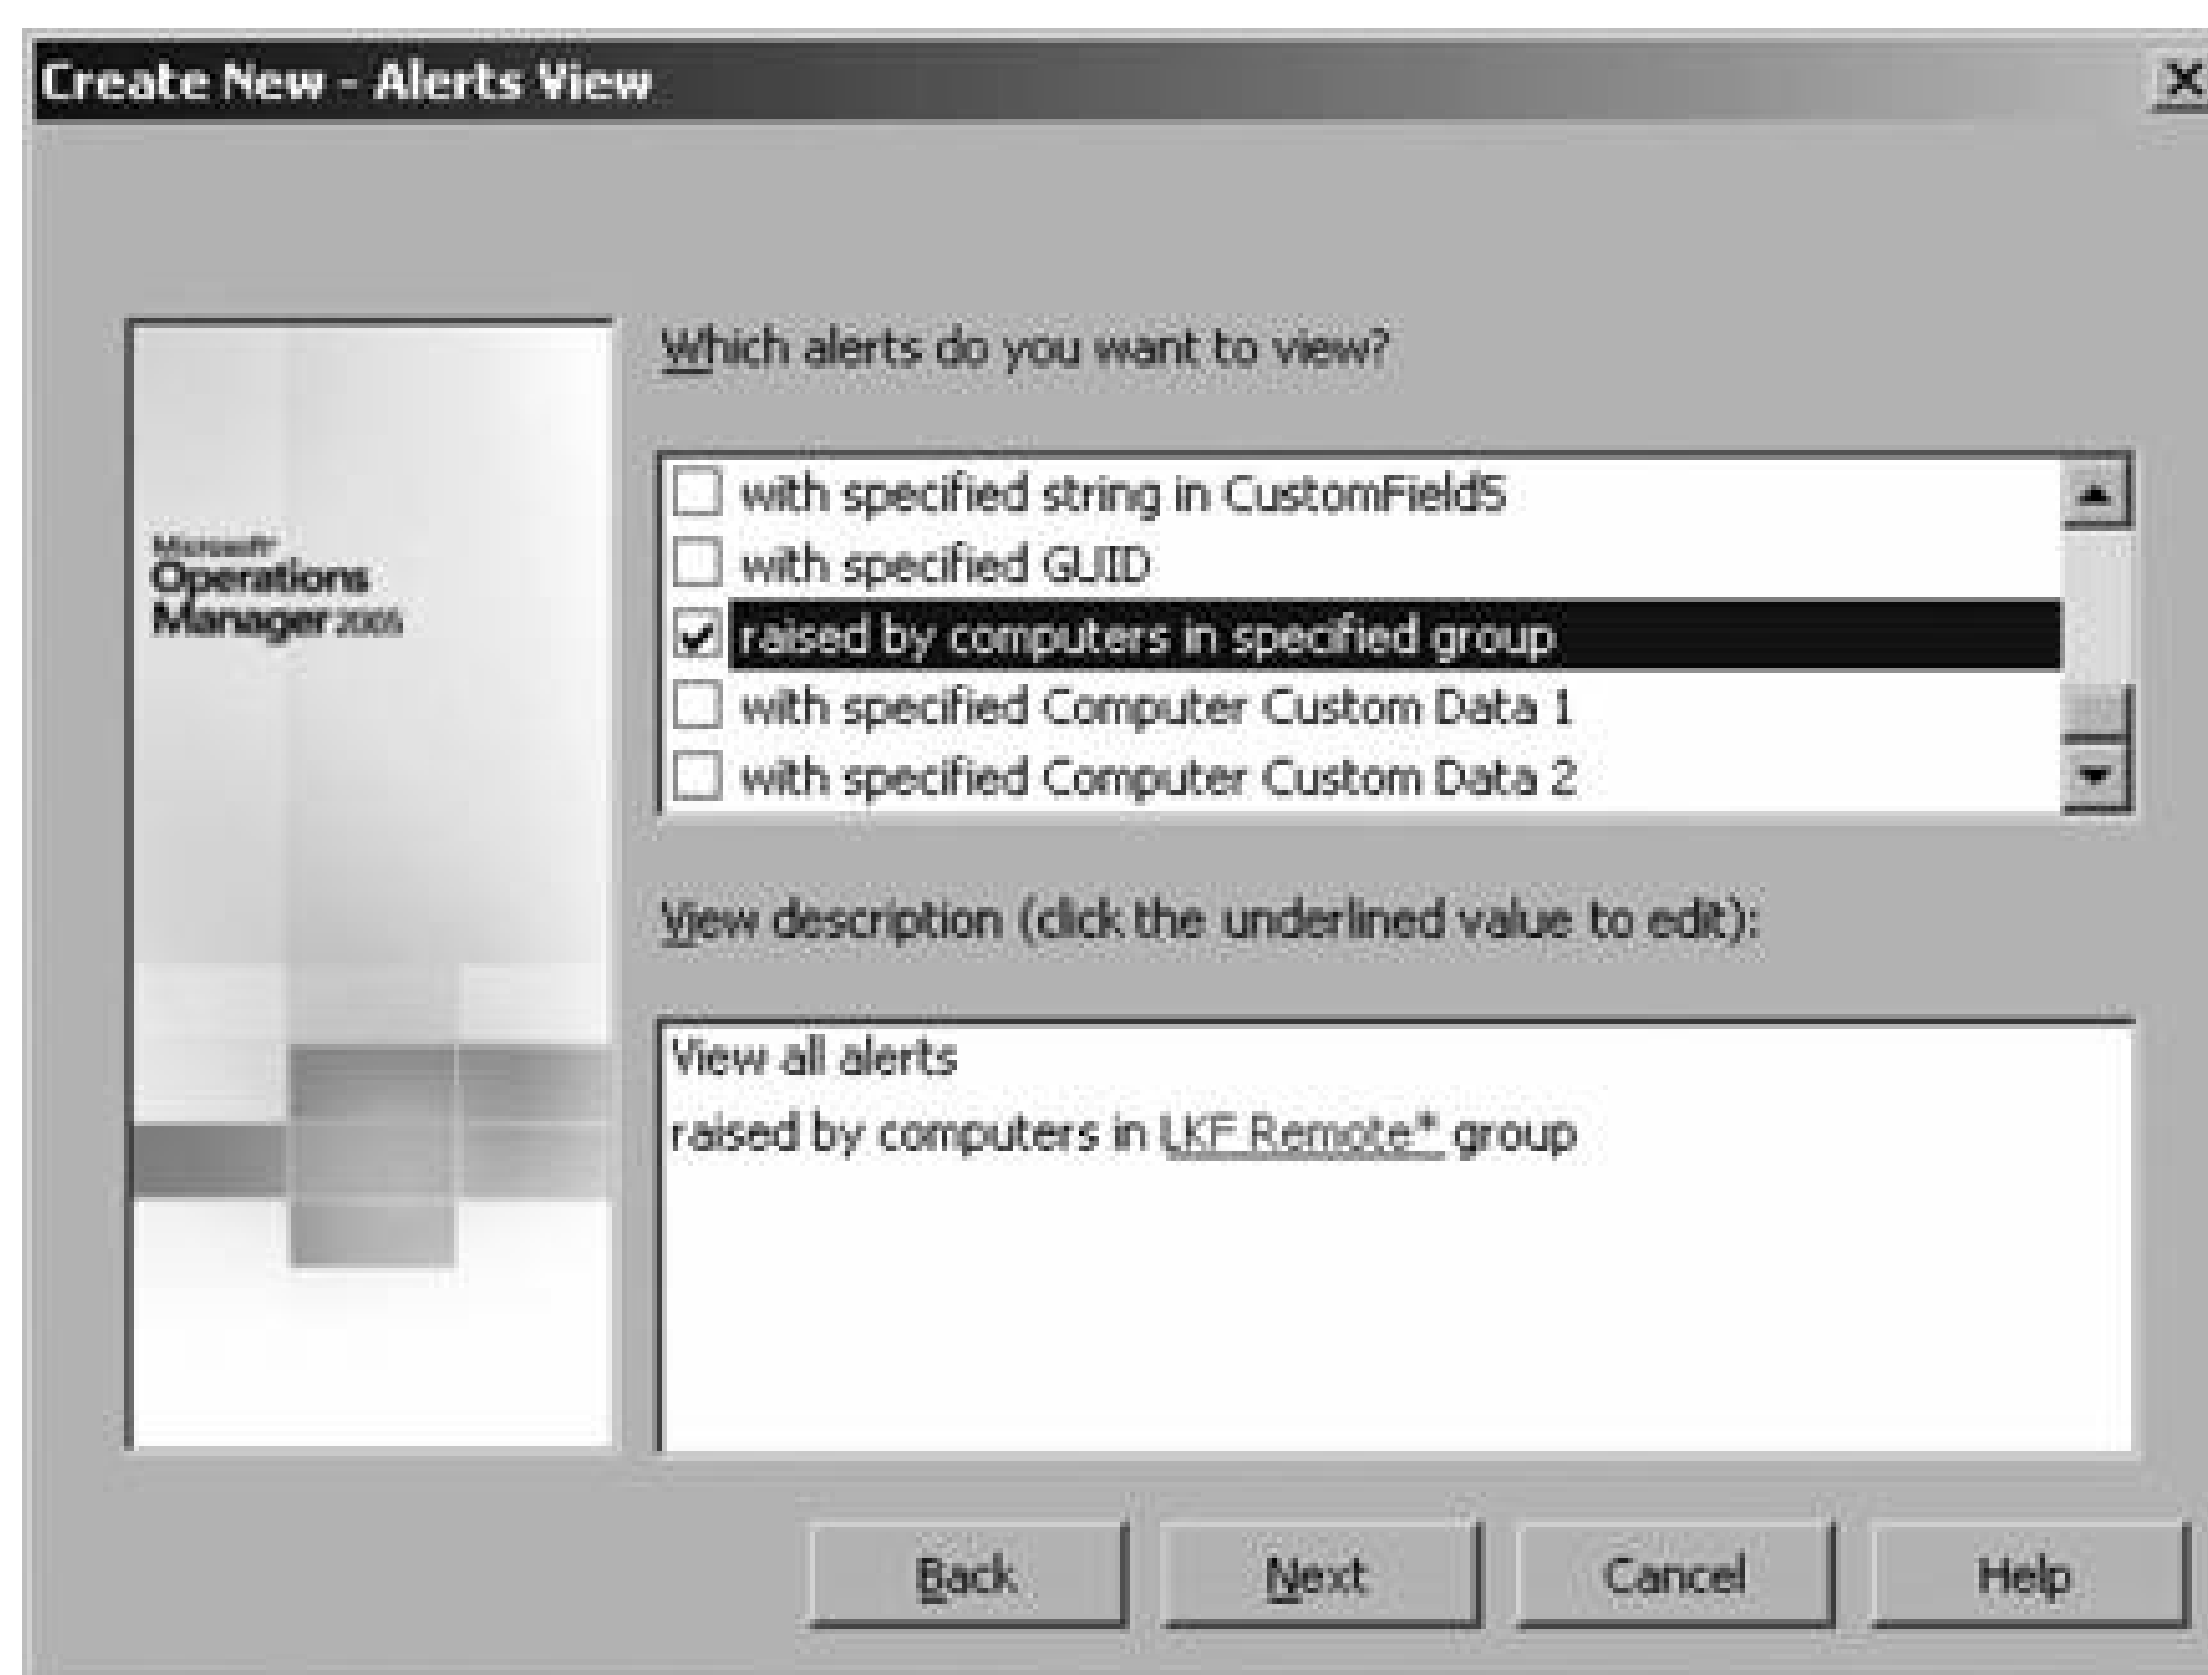




Moving to the next page, Max selects which alerts he wants included and enters the specified group text string. Note that you can use wildcard characters in the computer group name. This way, if Max creates any other computer groups for the remote offices, as long as he follows the naming convention, those groups will be included in the view ([Figure 6-22](#)).

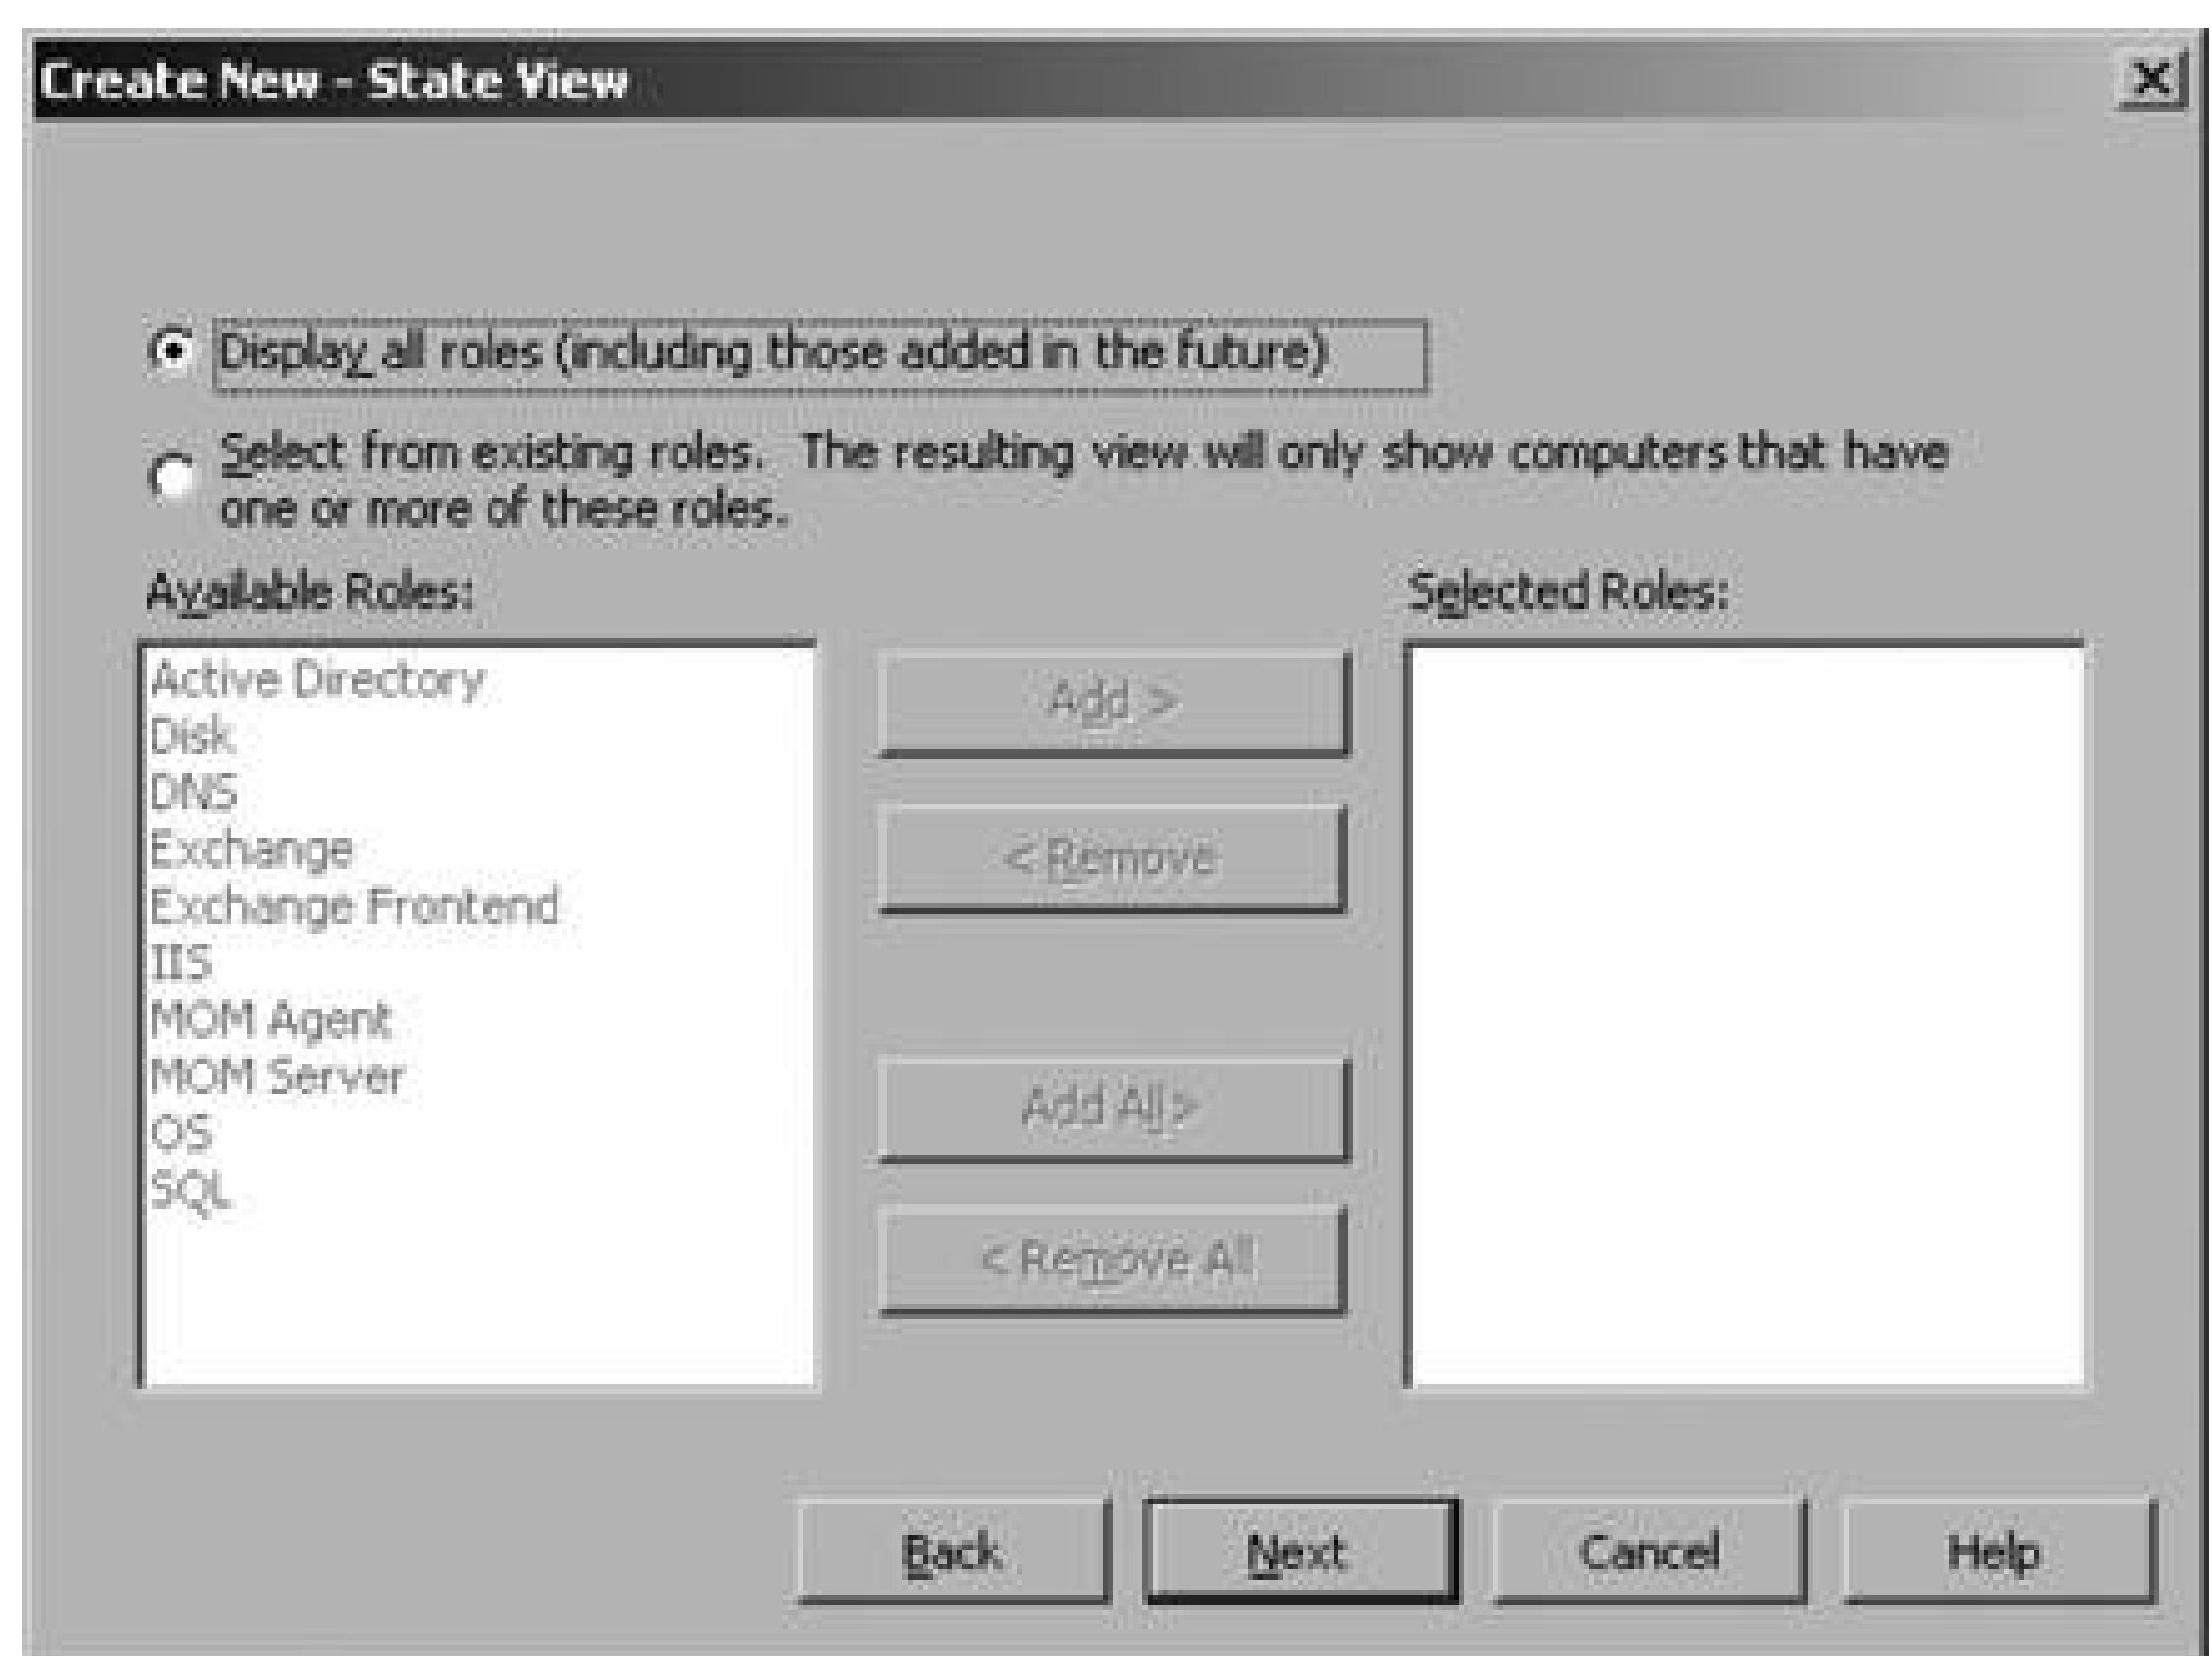
The wizard finishes by requiring a name and an optional description.

Figure 6-22. Specifying the computer group to view alerts from



Max then repeats this process to create a state view, with one notable exception. After he specifies the LKF Remote\* string for the computer group name, the wizard presents a page to select the roles he wants to include. The choices are to display all roles, including any roles that may be added in the future, or to pick specific roles. To minimize maintenance of the view, Max leaves the default, which is to include all roles ([Figure 6-23](#)).

Figure 6-23. Selecting to include all current and future roles that may be monitored on these computers



The wizard finishes and both view queries are immediately active. Unlike the Administrator console, where a configuration change like this would require a "commit configuration change" action, the Operator console does not. The whole console needs to be refreshed. This is done by either selecting View Console Refresh, or hitting Ctrl-F5. This has the same effect as closing and re-opening the console.

The last step to make these view queries available to the remote site administrators is to copy and paste them into the Public Views group. This step is necessary because these queries were developed in the My Views of an administrator, which is not accessible to anyone but that administrator. [Figure 6-24](#) shows the final resultthe All LKF Views folder published with the State View selected (looks like the MOM Agent Queues error alert finally got cleared up).

Figure 6-24. Publicly accessible custom view queries

Microsoft Operations Manager 2005 - Operator Console - LKFProdMG

File Edit View Go Help

Tasks

Group: MOM Administrator Scope

401

Public Views

All : Public Views

Alerts

Computer Groups

Computers

Diagram

Events

Performance

Service Level Exceptions

State

Task Status

Active Directory

All LKF Views

All LKF Remote Site Comput

All LKF Remote Site Comput

Microsoft Baseline Security Ana

Microsoft Exchange Server

Microsoft ISA Server 2000

Microsoft Operations Manager

Microsoft SQL Server

Microsoft Windows Base OS

Microsoft Windows DNS Server

Microsoft Windows Dynamic Ho

Microsoft Windows Group Polic

Microsoft Windows Internet Inf

All LKF Remote Site Computers State View

State

Computer

Active Directory

Disk

DNS

IIS

MOM Agent

Service Unava...	HOMESRV02					
Success	HOMESQLSERVER					

State Details - Computer: HOMESQLSERVER, Role: MOM Agent

Computer

Instance

Configuration

Heartbeat

Performance

Queues

WMI

HOMESQ...	HOMESQLSER...					
-----------	---------------	--	--	--	--	--

Total: 2 Item(s) Selected: 1 Item(s)

Last refresh: 4/18/2005 8:12:08 AM

homemomserver2

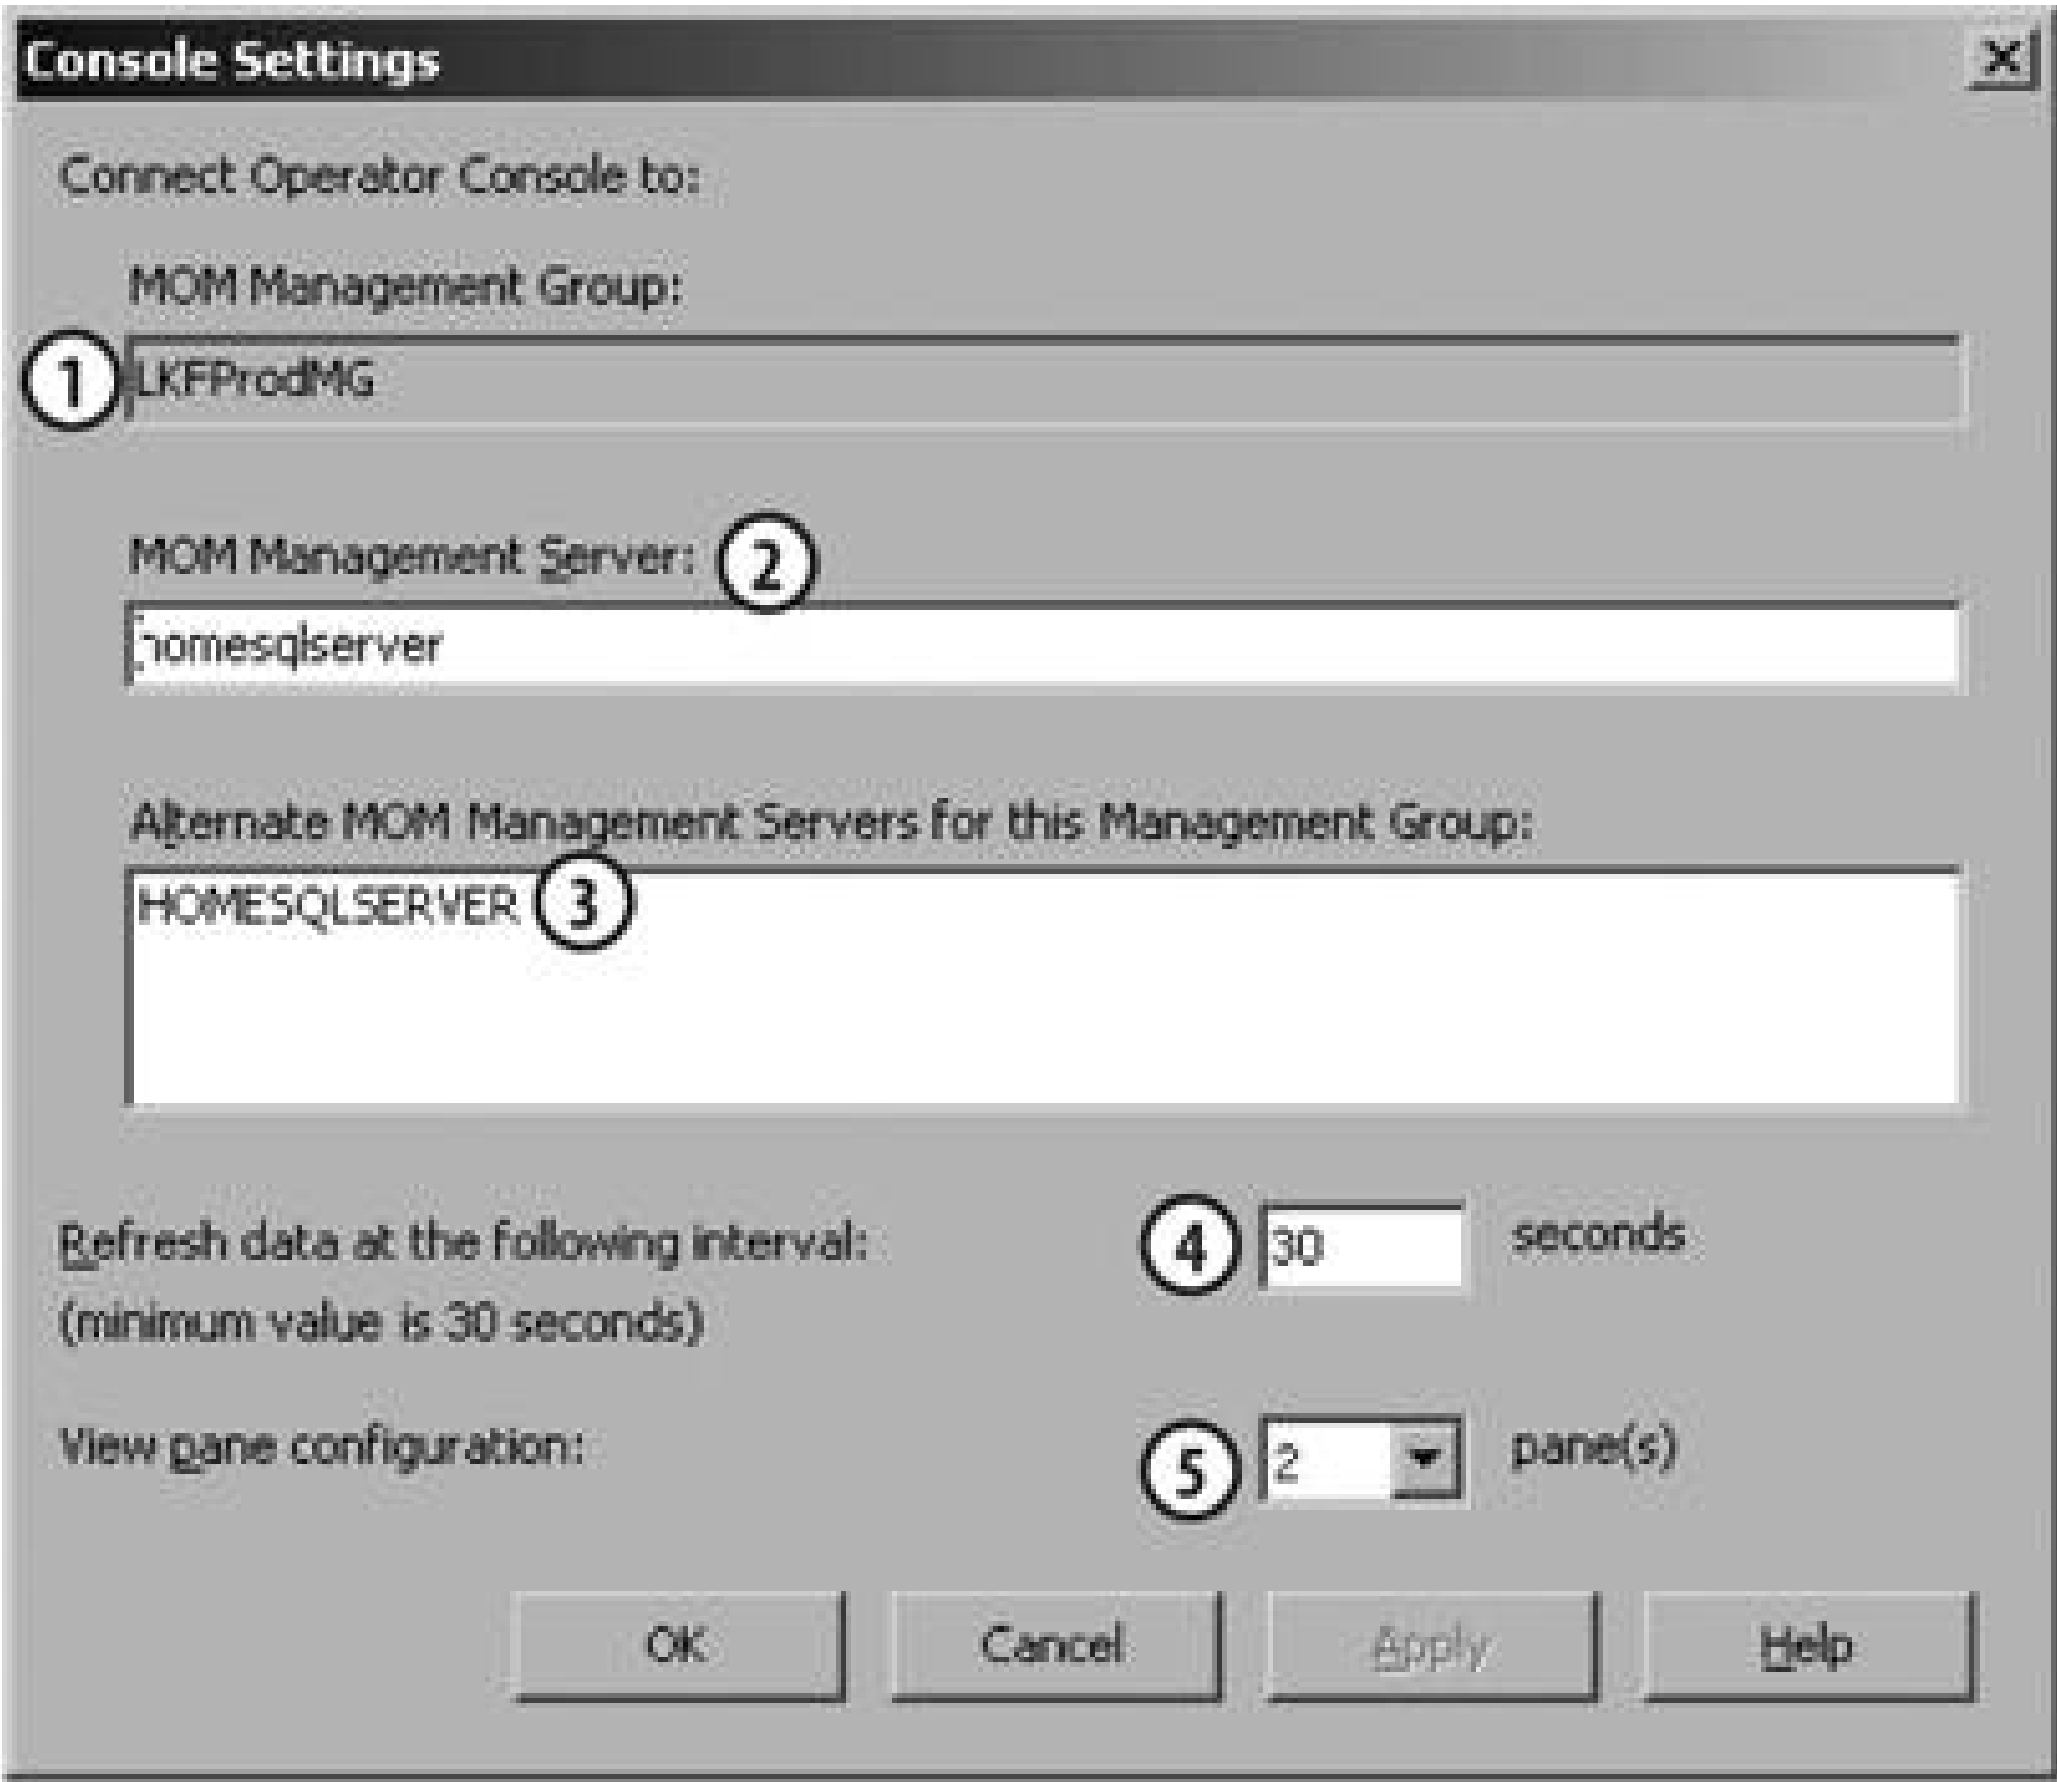


## 6.8. Customizing the Operator Console

The basic four-pane arrangement (five counting the Tasks pane) needn't be a static one. You can configure additional Results panes, each displaying different information based on the view group and selections in the Navigation panes. Additionally, for views that can display data with a time component, say events or alerts that have occurred between certain times, you can configure different time windows.

Additional Results panes and other settings are controlled from the File menu      Console Settings configuration page, as shown in [Figure 6-25](#):

Figure 6-25. Configuring Operator console settings



*MOM Management Group*

Point 1 in [Figure 6-25](#) indicates the management group that the Operator console is currently connected to.

### *MOM Management Server*

Point 2 in [Figure 6-25](#) is the text box that shows which management server you are connected to, and it is also where you can enter the name of a management server in a different management group when you want to change management groups.

### *Alternate MOM Management Servers*

All of the management servers in the management group are listed in the text box shown at point 3 in [Figure 6-25](#). You can switch between management servers by selecting the server name in the text box and clicking OK or Apply.

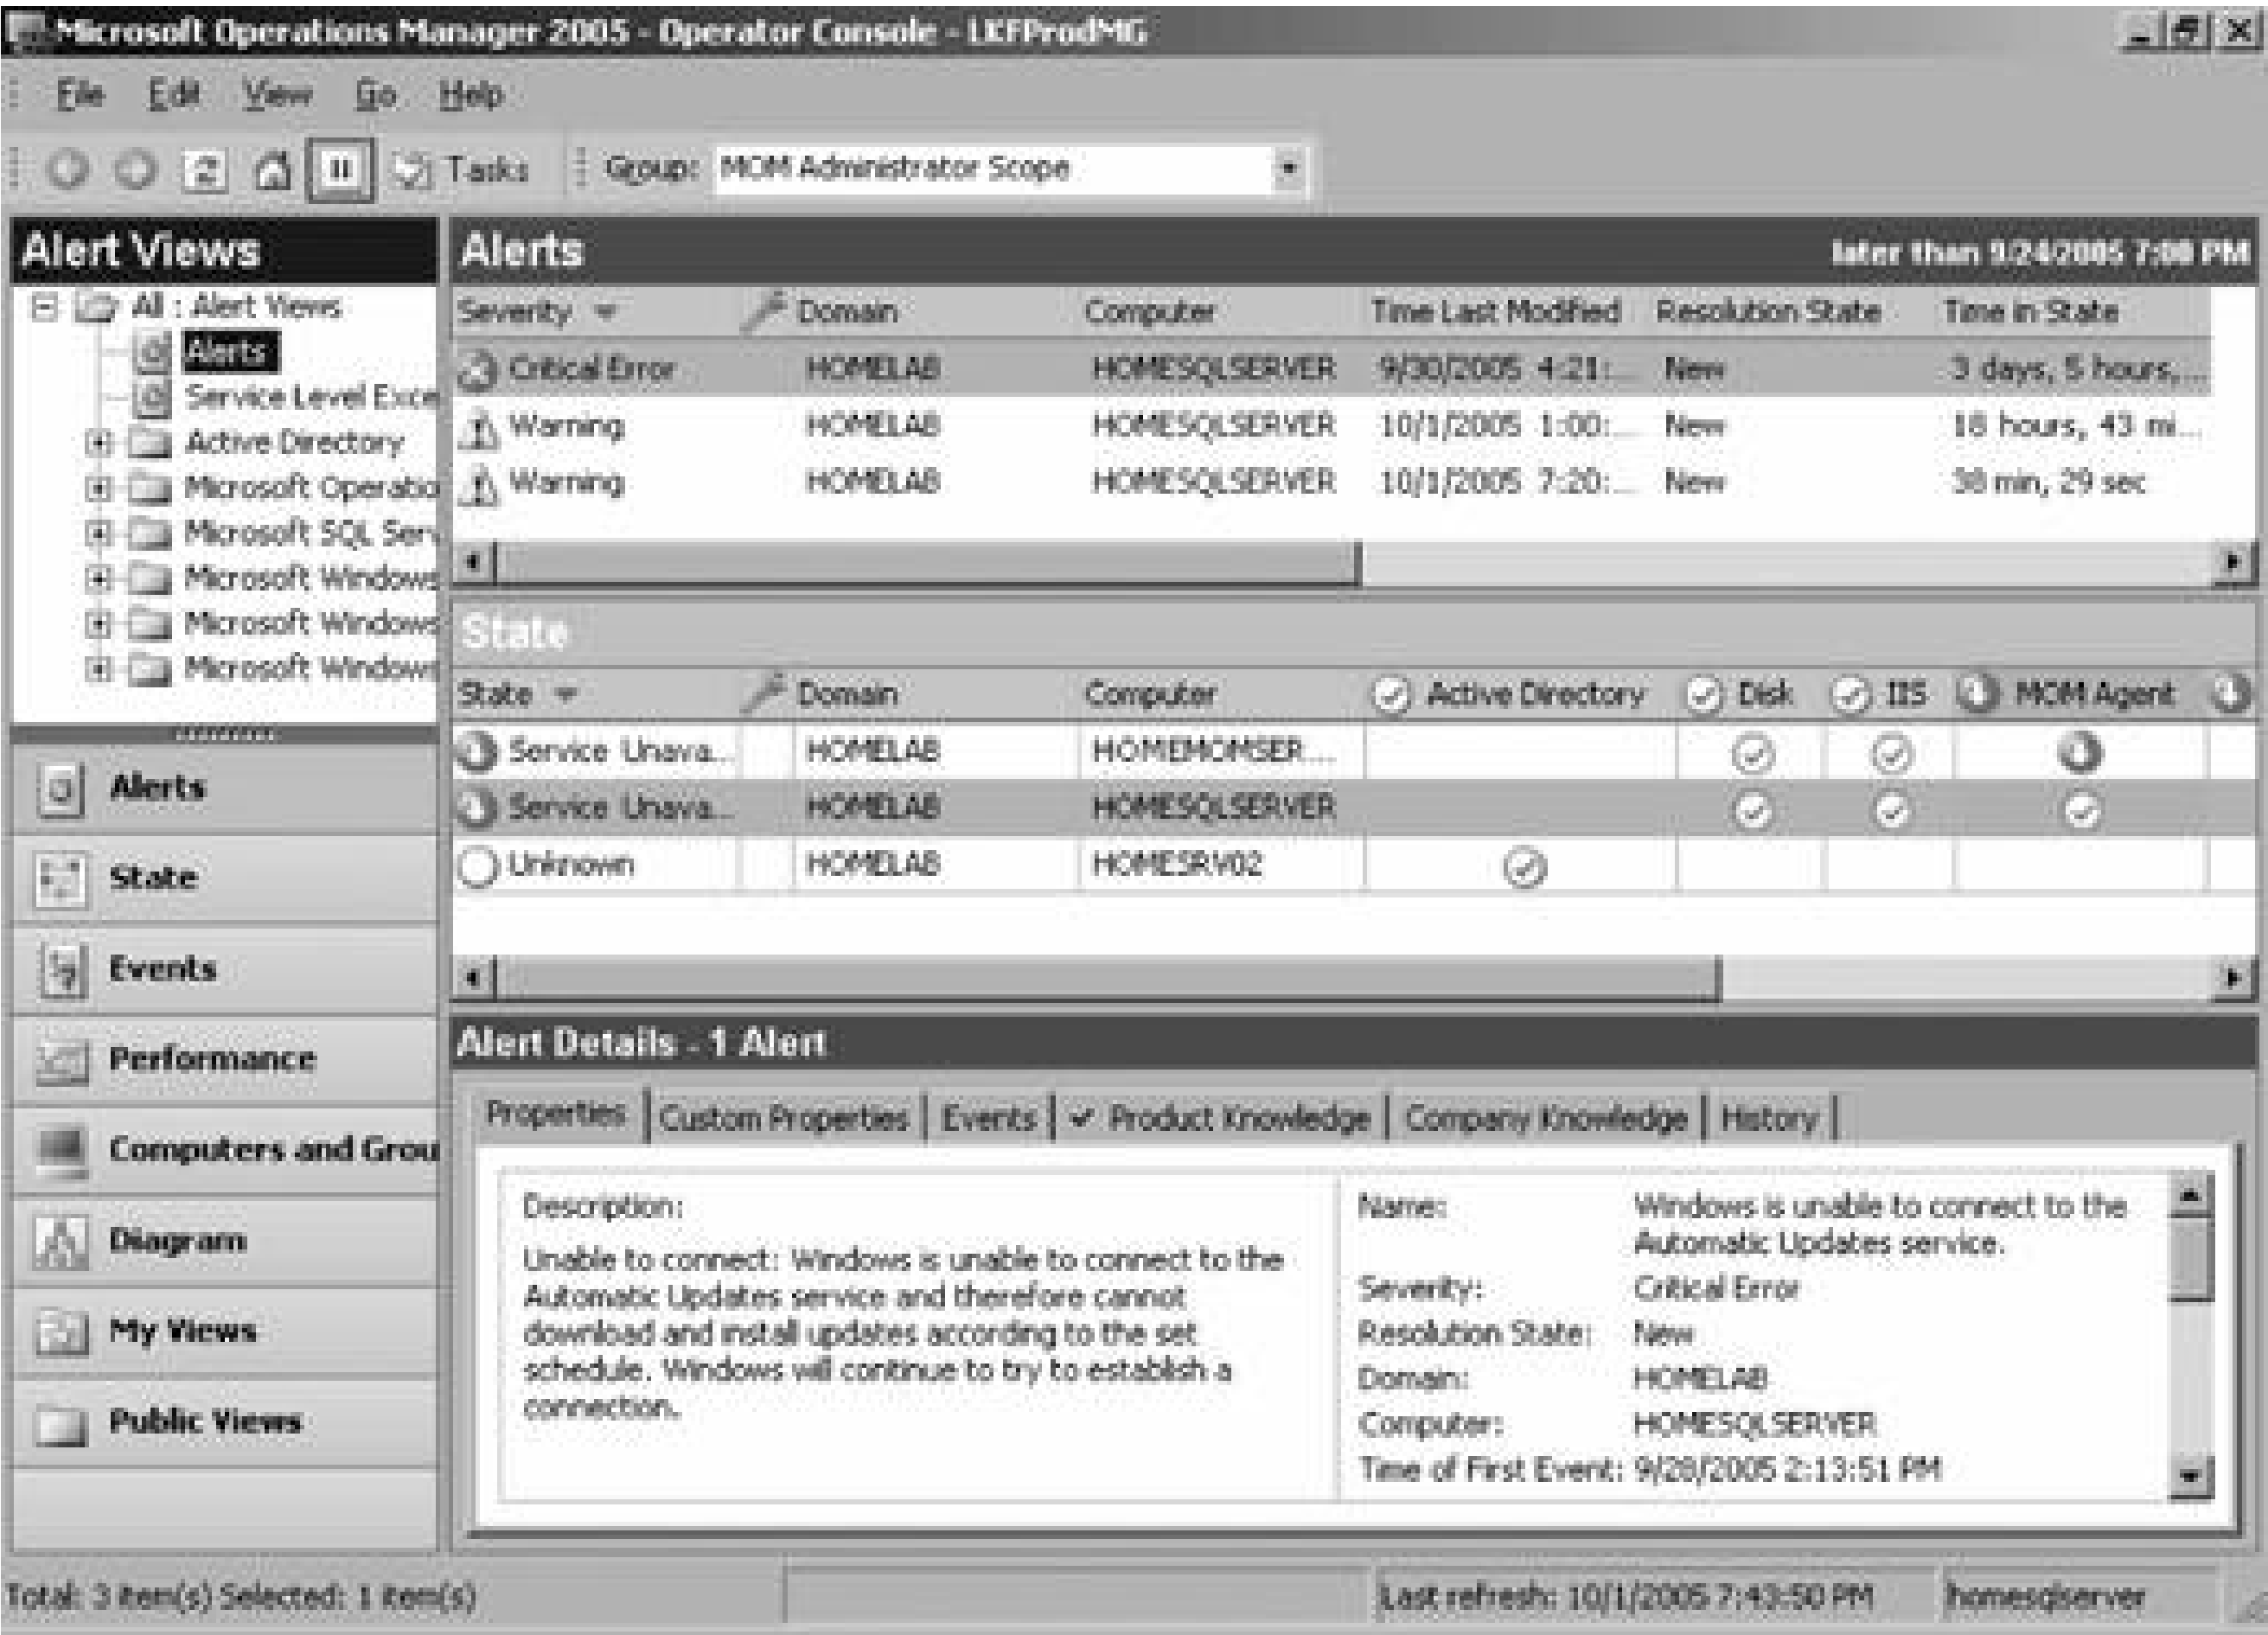
### *Refresh data*

The Operator console will refresh all data that it displays according to the interval that is configured at point 4 in [Figure 6-25](#). As noted in the figure, 30 seconds is the fastest that the refresh can occur. When there is a great deal of data to be displayed, it may take the Operator console a longer time to render the data than has been set here. The result is that the Operator console never gets to finish rendering its data before it has to start a refresh cycle and render from scratch. If you find yourself in this situation, increase this interval until the console can successfully render all its data before the refresh cycle starts.

### *View pane configuration*

Point 5 in [Figure 6-25](#) is a simple drop-down menu with the values of 1, 2, and 3. When you select 2 or 3, one or two additional Results panes are displayed. [Figure 6-26](#) shows the Operator console with two panes displayed.

Figure 6-26. The Operator console with multiple view panes configured



In [Figure 6-26](#), one Results pane is displaying MOM operational data in an Alerts view and the other one is showing a State view. The details pane always shows the data of a specific item in the Results panes; in this case it is the Critical Error alert on *homesqlserver*.

When you find a configuration that is particularly useful to you can save it as a file with an *.omc* extension (File menu → Save as) as shown in [Figure 6-27](#). Here you can see that a two- and three-pane configuration has been saved.

To control the time frame of the displayed data, use the View Date and Time Filter toolbar button. If it is not displayed, select it from the View menu → Toolbars → View menu. Once available, the icon for this filter appears on the toolbar as shown in [Figure 6-28](#).

When you select this, the View Date and Time Filter page is displayed, as shown in [Figure 6-29](#). This additional filter allows you to display data that occurred within a given period of time using the "Within the time range" options or simply within the last number of seconds, minutes, hours, or days using the "Within the last" options.

Figure 6-27. Saving an Operator console configuration




Figure 6-28. View Date and Time Filter toolbar button

Figure 6-29. Setting the time filter for the Operator console



View Date and Time Filter

Select the values to show based upon date and time ranges:

Within the time range:

After

9/28/2005 11:02 PM

Before

9/30/2005 11:02 PM

Within the last:

2

hour(s)

OK

Cancel

[← PREV](#)

## 6.9. Summary

The Operator console is used to consume and manipulate all of the data that MOM collects and has in the operational database; it does not have access to information that is in the data warehouse database. Out-of-the-box it provides a wide range of preconfigured views (queries) that have been sorted into groups according to the type of data they focus on.

The data that is accessible to you in the Operator console is first controlled by the console scope that you are a member of. A console scope is a non-secure association between a set of computer groups and user accounts. By default, all accounts in the MOM Users group are assigned to the MOM Users console scope. This scope has access to all computer groups.

To use the Operator console effectively, you must understand how to navigate it and then how to use the information in the view groups. Navigation, or building filters, is done by first selecting the view group, then the view query, and finally the computer group of interest. The view groups can be broken down into two major uses. One type is notification, where MOM is letting you know what is going on in your environment (alert, state, and diagram views). The other is more research-oriented, providing you with historical or configuration data (computers and groups, performance, and events views).

You can modify any existing view query or create your own in the My Views group. This is especially useful for developing custom views that you can later publish by copying to any of the other view groups, which are all publicly accessible.

Use the Operator console to find out what is going on in your environment right now or at least within the past four days if you have accepted the default grooming settings. If you want to make use of data that is older than that, you have to install and configure the MOM 2005 Reporting solution, which gives you a longitudinal look at your data.

Both the current and historical uses of the collected data depend entirely on the MOM 2005 databases. At a minimum, you must install the operations database; if the MOM 2005 Reporting solution is included, then the data warehouse and SQL Server Reporting database are added to the mix. The next chapter covers the basic components of the operational database and teaches you what you need to know to maintain and protect it.

# Chapter 7. MOM 2005 Database Fundamentals

The operational database is the choke point of every MOM 2005 management group. Before an alert is displayed in the Operator console, it is first written to the database. Before a notification on an alert can be sent, it goes through the database. The Operator console is a giant filter for all of the agent-collected information in the database. The Administrator console can't administer the management groups' configuration settings without interacting with this database.

The health of this database directly impacts the performance and health of the MOM 2005 management group. The MOM OnePoint database, which is the actual name of the operations database, doesn't operate alone. Other databases on the management groups' SQL Server perform supporting roles, so management group health is dependent on these as well.

This chapter teaches you how these databases interact, how to maintain them, and how to protect them (through backup and restore). It also addresses the functions and maintenance of the databases involved with MOM 2005 Reporting.

This chapter does not go into detail about SQL Server administration, but rather presents some the basic SQL Server tools and database administration tasks that will be relevant to the MOM 2005 administrator.

[Chapter 2](#) covered the sizing exercises that are necessary for the installation of the OnePoint database. Basically, this is to calculate the amount of data flowing into the database on a daily basis and the desired frequency of grooming, which deletes data from the database. If you estimate well and perform the calculations correctly, the database size will remain fairly constant.

[Chapter 5](#) covered the grooming settings that are configured in the Administrator console. These settings are adjustable, but you must be careful to strike the correct balance between removing the unnecessary data and keeping the data you need there.

What the setup and Administrator console interfaces don't give you are mechanisms to back up and restore the MOM 2005 database and to perform other administrative tasks interactively. They don't provide visibility into the SQL Server jobs that are used to groom and reindex the database. To perform these tasks, you have to use the SQL Server Enterprise Manager and the SQL Query Analyzer tools. Fortunately, for this type of maintenance you do not have to be a SQL database administrator (DBA) to use these tools effectively. Don't be intimidated by diving into the heart of MOM.

## 7.1. SQL Server Enterprise Manager

To start SQL Enterprise Manager on the MOM database server, select Start → Programs → Microsoft SQL Server → Enterprise Manager. Surprisingly enough, the first time you launch this tool, it will not have a connection to a database server you have to register a SQL Server in the console to manage it. To register a server:

1. Right-click the SQL Server Group object in the SQL Enterprise Manager and select New SQL Server Registration.

This launches the SQL Server Registration wizard. Proceed through the start page and the wizard presents you with a list of available SQL Servers.

2. Select either the "(local)" entry or enter the name of the server that you want to administer, highlight it, and click Add to move it to the "Added servers" column (see [Figure 7-1](#)).  
*homemomserver3* has been added.

Figure 7-1. Selecting the SQL servers to connect to

3. Click Next to bring up the Connect Using page to select the authentication to connect to the SQL Server that you specified in the previous step. For MOM database servers, you can always select

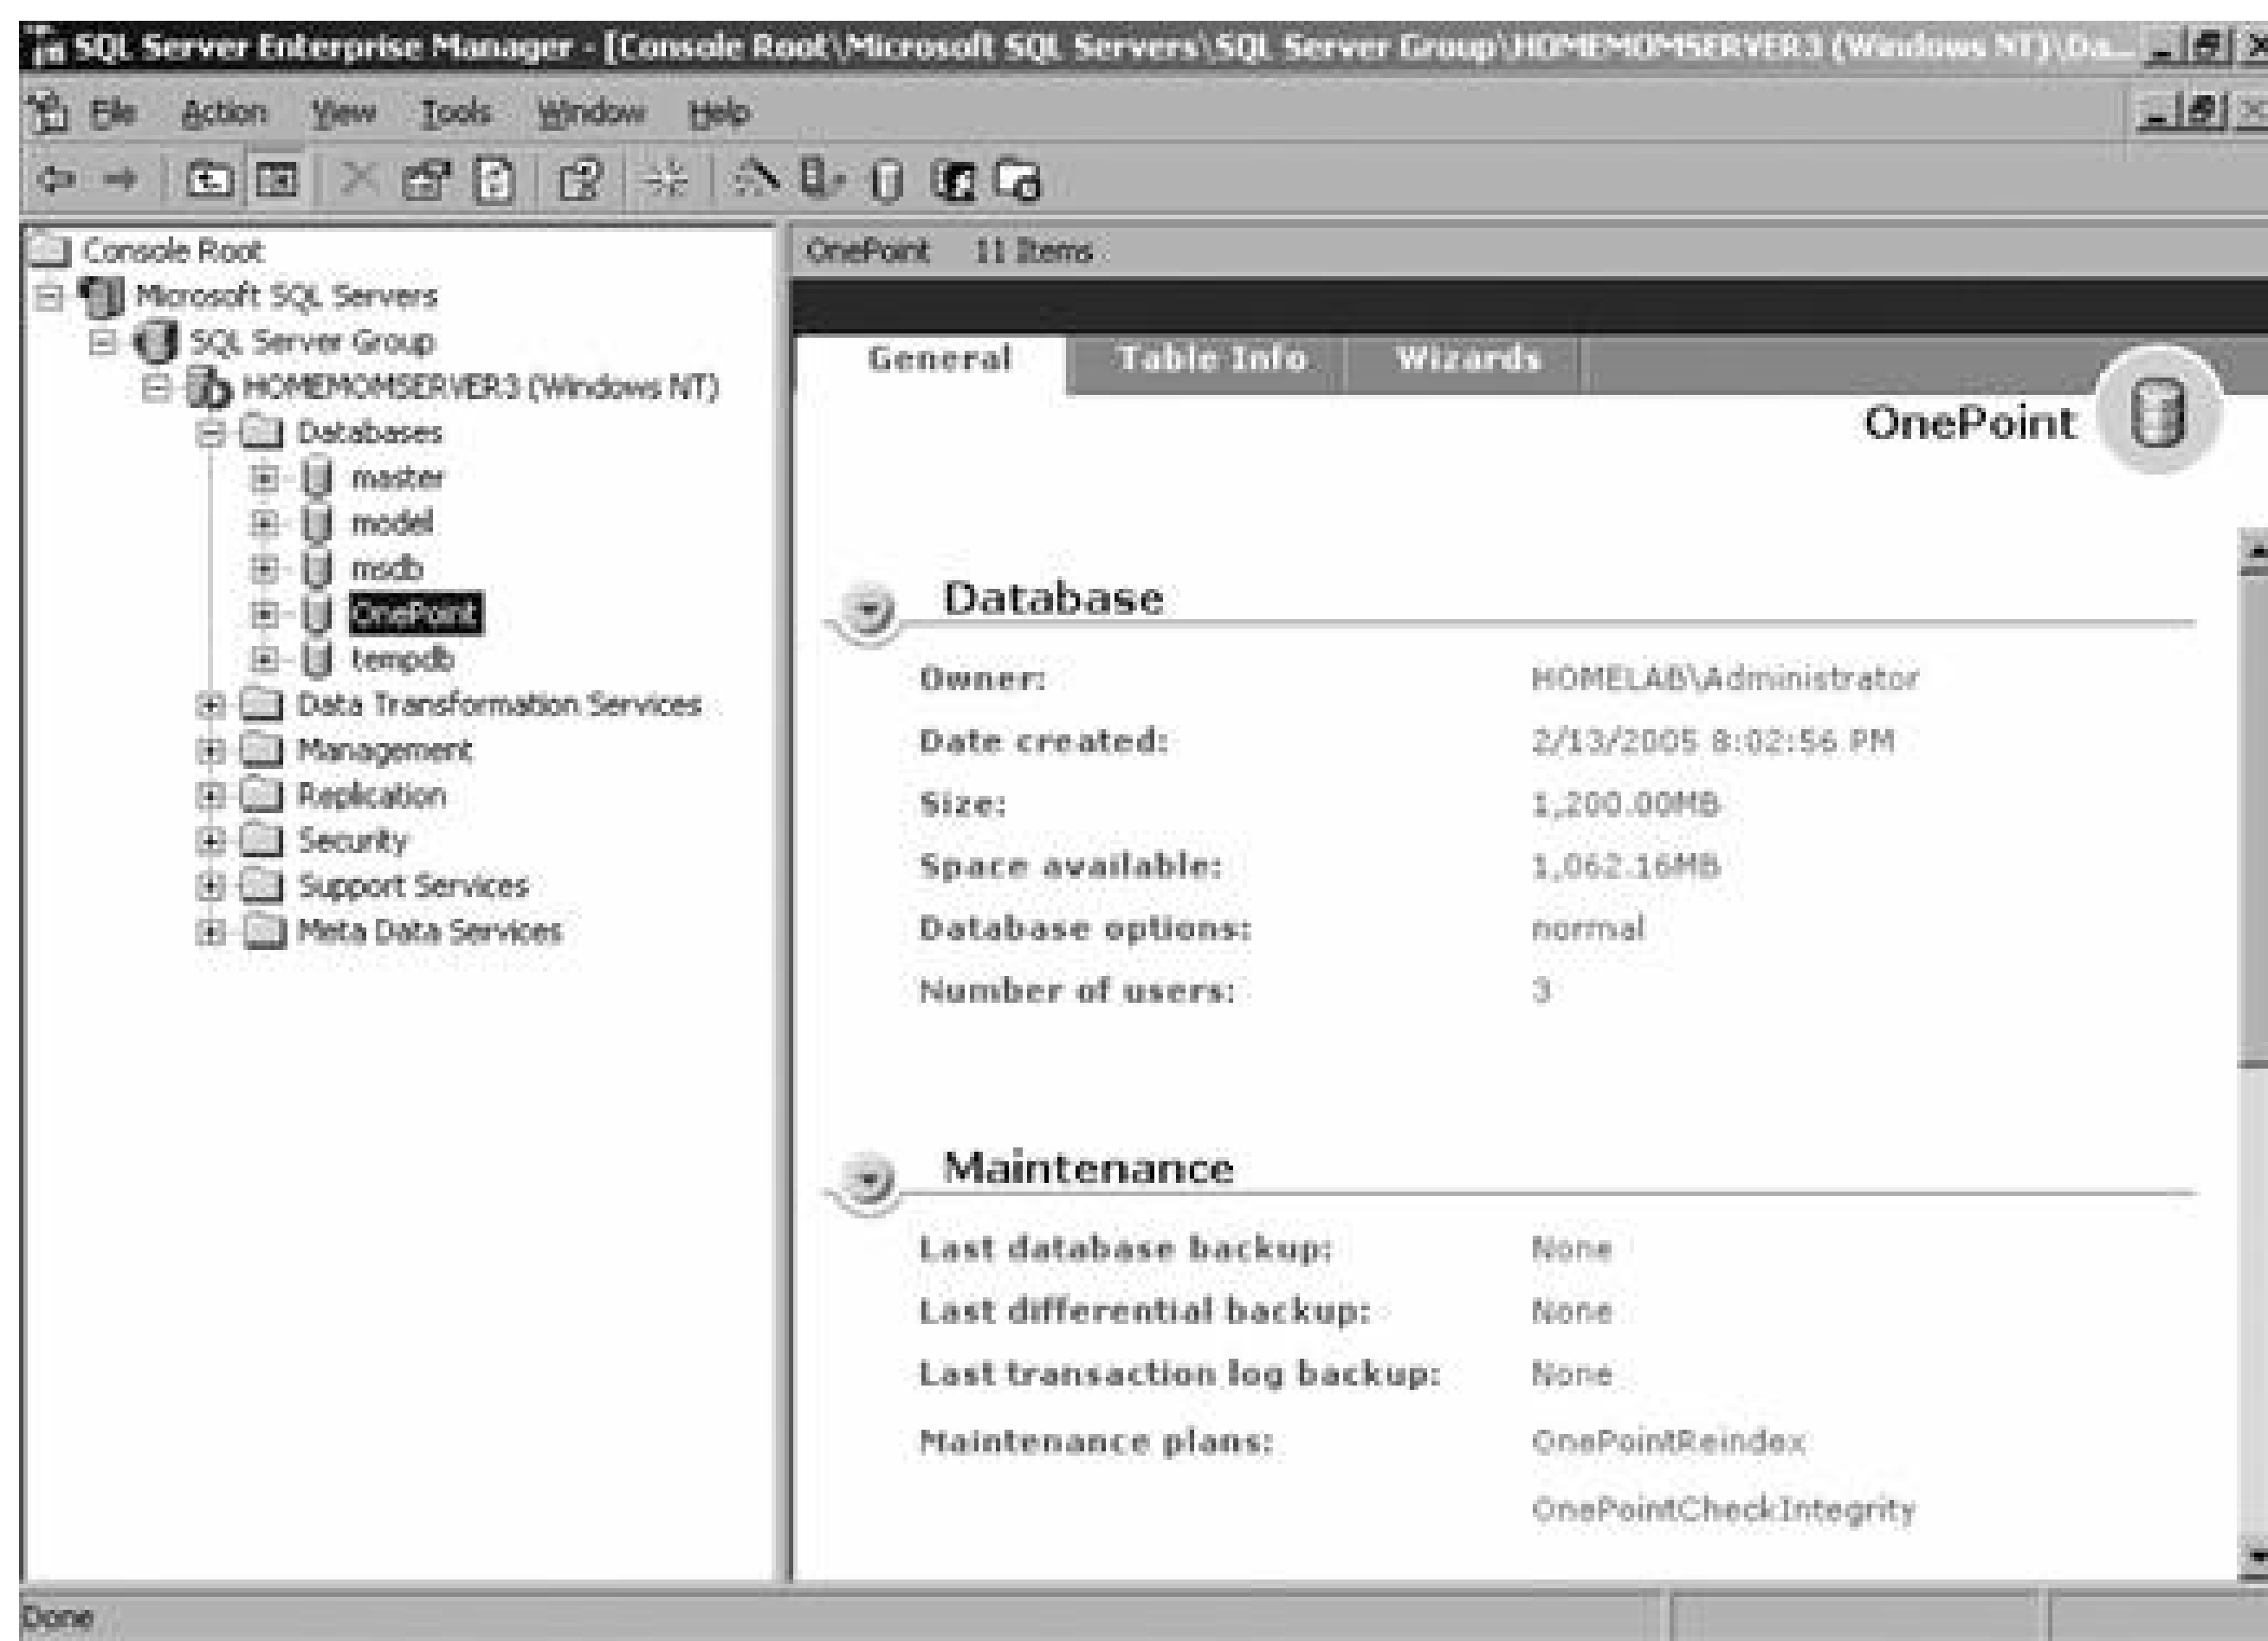


the Windows Authentication option, assuming that you are logged on with administrator rights.

4. Click Next to bring up the Select SQL Server Group to add the registered server. You can accept the default, which will add the server to the SQL Server Group top-level group. The next two pages are the standard Summary page and then the Successfully Completed page.

Once a server has been successfully registered in the console, you can expand it to the folder level and select the databases folder. In here you will see the OnePoint database as well as the master, msdb, model, and tempdb databases. These last four databases are the system databases and should not be tampered with. [Figure 7-2](#) shows the taskpad view of the OnePoint database in SQL Enterprise Manager. To open the taskpad view, right-click the OnePoint database and select View Taskpad view.

Figure 7-2. Taskpad view of the OnePoint database



Many of the administrative tasks are accessible in the taskpad view. The only other folder that you really need to be concerned about is the Management folder, which contains the SQL jobs that are used to maintain OnePoint, the backup utility, and the database maintenance plans. These will be covered in the "[Backing Up SQL Databases](#)" and "[Restoring SQL Databases](#)" sections later in this chapter.

In the taskpad view, you are immediately given valuable information regarding the size of the database, the space available, and the date and time of the last backups. In [Figure 7-2](#), the database size is listed as 1,200 MB, which is a combination of the database file and the transaction logfile.

## 7.1.1. Basic SQL Database Concepts

The OnePoint database actually consists of two files that are located in the SQL install directory *C:\Program Files\Microsoft SQL Server\MSSQL\Data*. The database file is *EeaData.mdf* and its transaction logfile is *EeaLog.ldf*. Both of these files should be the size specified at setup. When you start planning your backups, you must include the transaction logfile in the backups. To understand why, you need to understand a little about how transactional databases work.

The OnePoint database file, *EeaData.mdf*, consists of about 375 tables. For the sake of this discussion, you can think of a table as a spreadsheet. A row in the table is one record and a column in the table is a value that is a member of all the records. The relational database engine (SQL Server) is constantly reading and writing and working with this data as processes that are outside of SQL interact with the database.

If these processes were allowed to directly read, write, or delete records in the database and for some reason the connection was interrupted during the operation, the record that process was working with would become corrupted. This does not create a very stable database. So, instead of allowing direct access to the database tables, all interactions are first written to the transaction log. These transactions are then committed to the tables by the database engine itself. In this way, the transaction log proxies and buffers communications between the outside world and the database.

The way SQL Server uses the data in the transaction logs provides an additional layer of protection and reliability. Basically, each transaction is made up of multiple steps, which all need to be completed successfully for the whole transaction to be considered successful. This is called *atomic transaction*, which means that every step needed to complete a transaction is contained in the transaction itself.

Here's an example that, while not database-related, is useful for understanding the concept of self-containment or atomicity. Say you want to move a file from one drive to another in Windows Explorer. First you select the file, then you select Cut, then you navigate to the destination directory and select Paste. The OS then performs several separate actions to complete the move. The first step for the OS is to check the destination directory to see if there is an existing file with the same name and return an error if the file already exists. Assuming the filename is not found, the OS reads the file and writes it to the destination directory (which is actually a copy action), and then deletes it from the source directory. If any one of these steps fails, the operation fails. To protect the original document, it is not deleted from the source directory unless all the preceding steps have completed successfully. So, if the preceding steps fail it is as if the entire process never occurred at all.

SQL Server tracks the success or failure of the individual steps in committing a transaction to the database and if any of them fail, the transaction as a whole fails. SQL server will then basically undo each of the steps, thereby rolling back the database to its original state before the transaction started. This is how the database is kept in a non-corrupted or consistent state.

What all of this means is that at any given time, the data that represents the OnePoint database is actually split between the database file and its transaction logfile. This means you have to back them both up. Because these are physically separate files, for purposes of performance and robustness, it is wise to distribute these files across separate physical drives, preferably those that participate in one of the common fault-tolerant RAID array configurations.



## 7.1.2. System Databases

There are other databases created by the default installation of SQL Server as shown in [Figure 7-2](#). These are the four system databases. Of these four, the master and msdb databases and their transaction logs should be included in your backup plans. Briefly, here is what they do:

### *master*

The master database stores the configuration information for the system and all the other databases on the server, such as the location of database files. SQL server also maintains accounts that are separate from the local OS or AD. Accounts created in SQL are stored here and are referred to as SQL logins.

### *msdb*

SQL Server uses this database to record system-level operations that were performed on the server, when they were performed, and who performed them. Of particular interest is the record of all backups and restores. This history is used by SQL server when restoring a database to determine the sequence of backups to be used for the fastest restore. In addition, this database is used for scheduling tasks such as the MOM maintenance jobs.

The order in which the system and MOM 2005 databases need to be restored for a complete system restore is covered in the "[Restoring SQL Databases](#)" section later in this chapter.

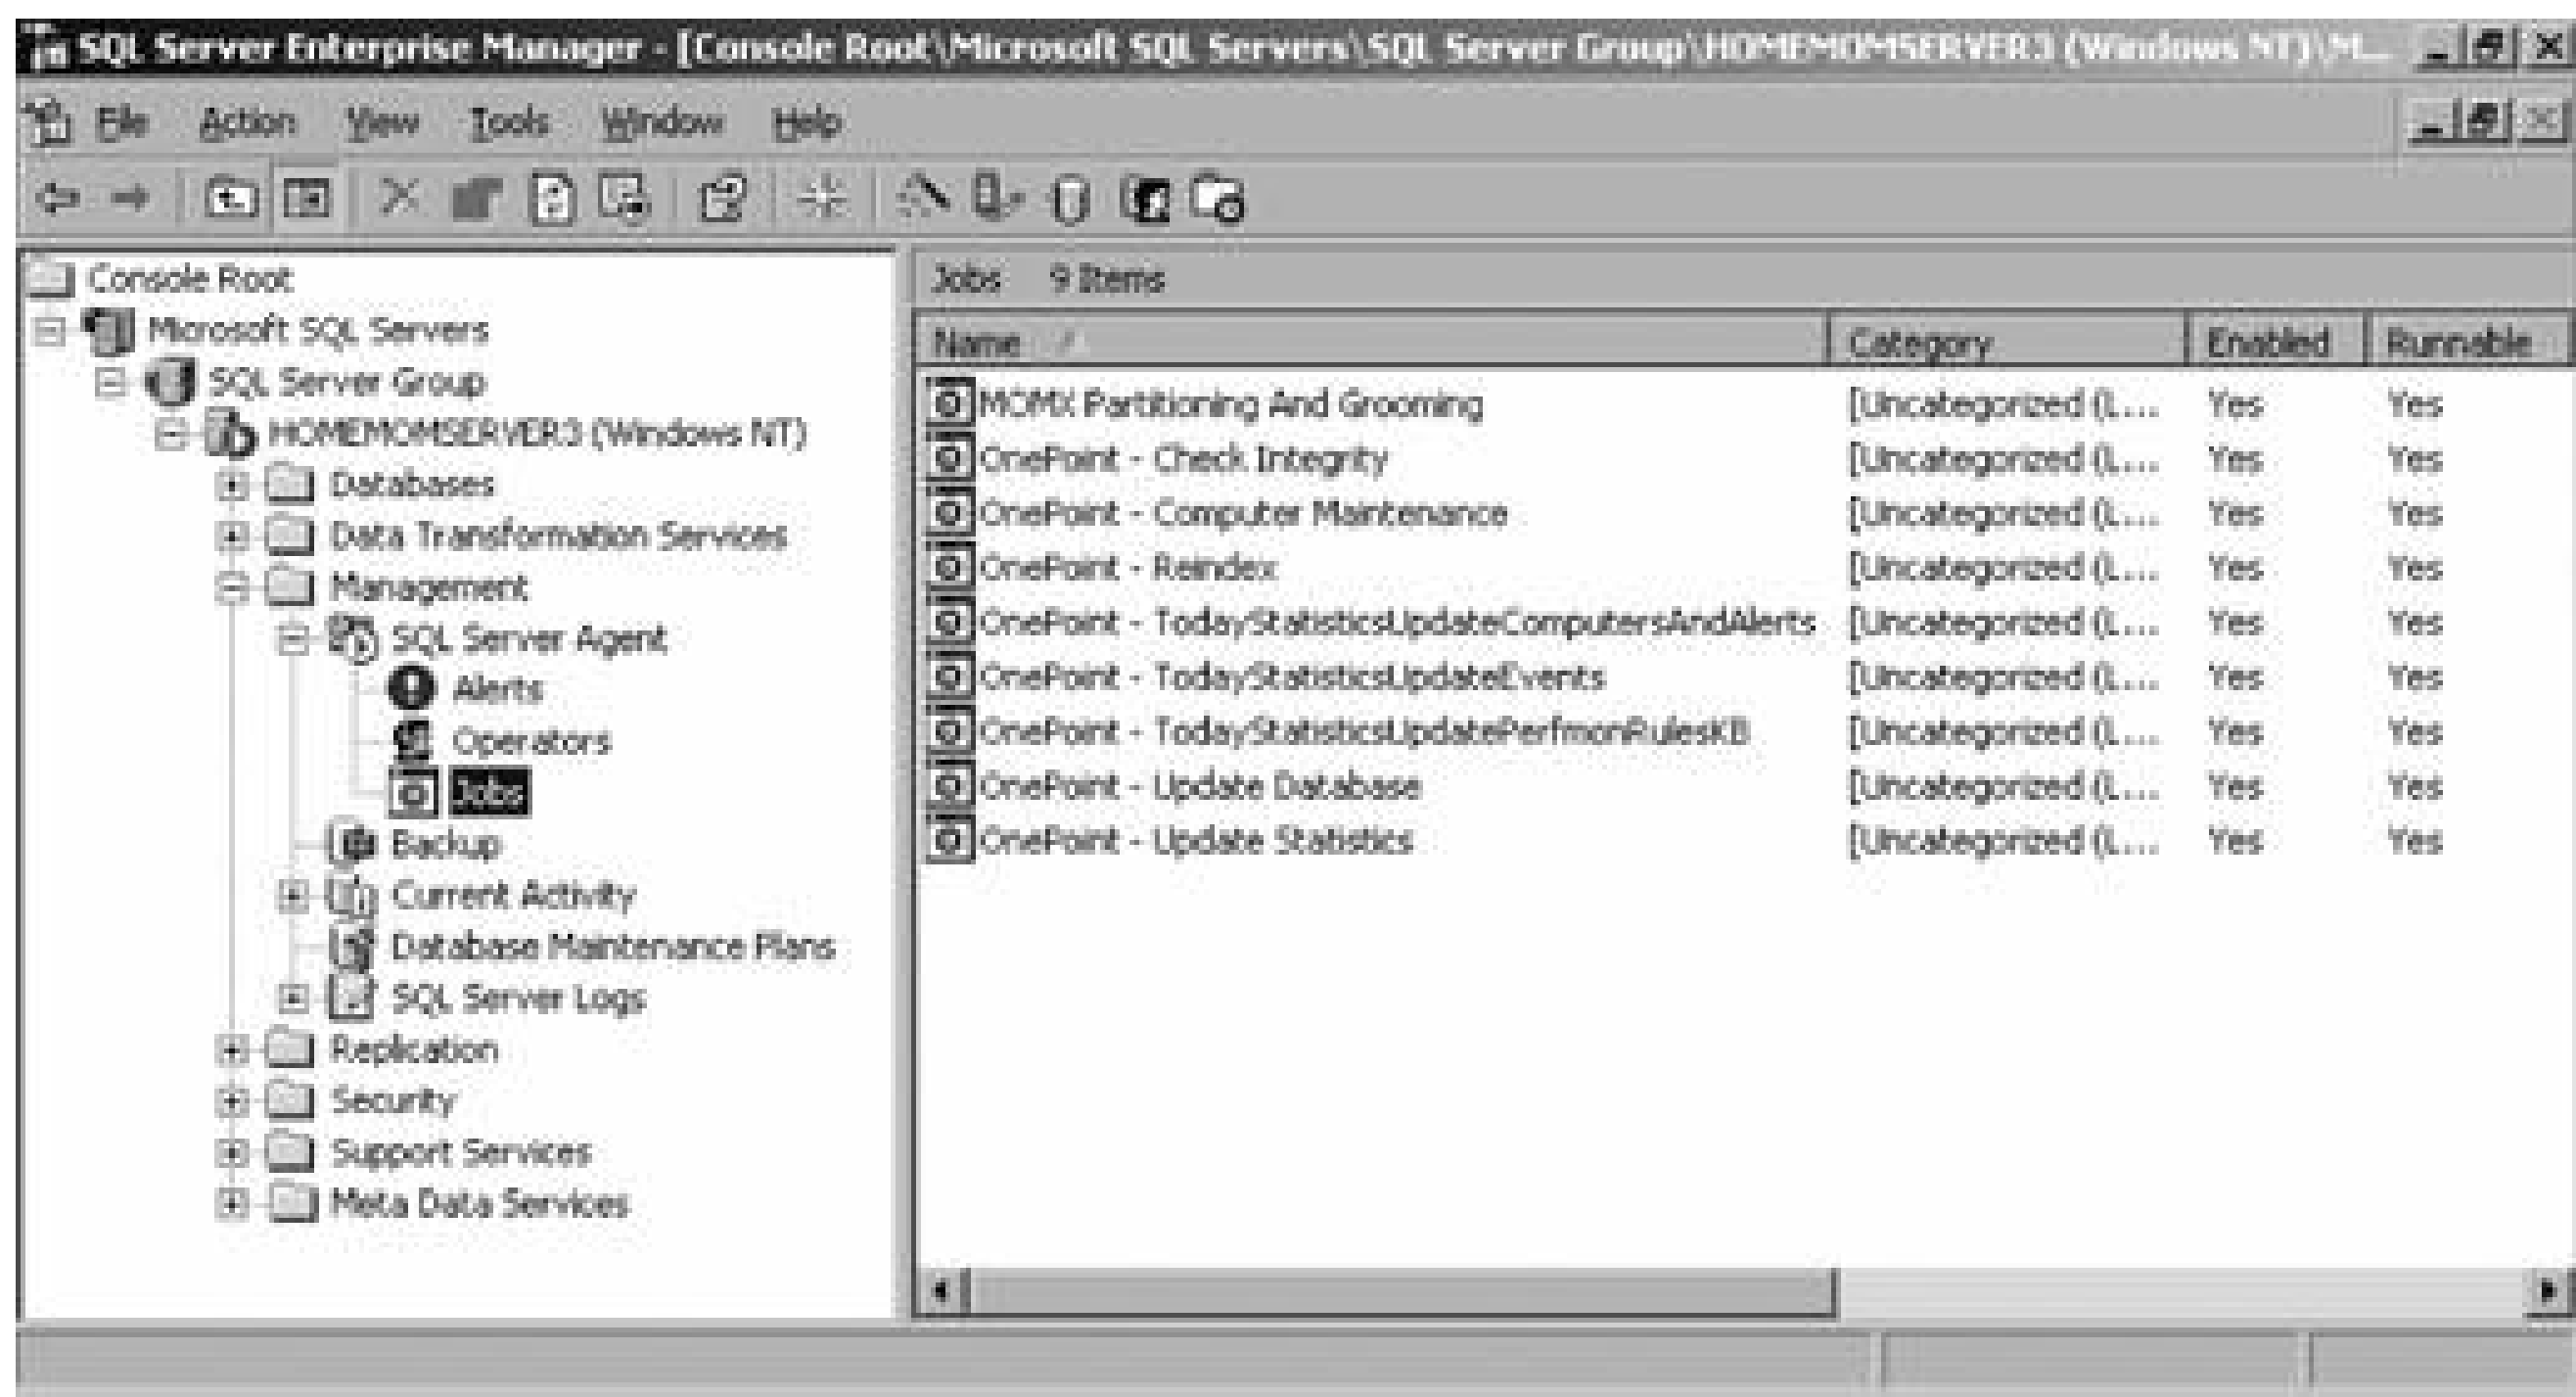
## 7.1.3. OnePoint SQL Maintenance Jobs

When the OnePoint database is installed, by default nine MOM-specific SQL jobs are created. They are located in the Management folder SQL Server Agent Jobs container ([Figure 7-3](#)). If you look in this container on a MOM database server that also has MOM Reporting installed, you will find that an additional grooming job is added for the data warehouse, bringing the total to 10.

Of these nine jobs, the MOMX Partitioning And Grooming, the OnePoint - Reindex, and the Update Database are the most important to you.

SQL jobs are basically wrappers around code that you want executed under certain conditions. Those conditions could be in response to a SQL alert (which is entirely different than a MOM alert and not relevant to this discussion), on a schedule, or when the SQL server processor reaches a certain user-defined level of utilization. For example, OnePoint jobs are all scheduled, although they can all be run manually, such as a task in the Windows Task Scheduler. A single job can consist of multiple steps (with each step consisting of a different set of instructions) that occur in a specific order. The execution of one step can be made conditional on the success (or failure, depending on the logic built into the job) of a prior step. At the completion of each step, the job reports the success or failure to a SQL operator. The OnePoint jobs are all configured to write an event to the Windows event log. As you can probably guess, these specific events are being watched for by the MOM agent on the database server and generate an alert if a failure is reported.

Figure 7-3. SQL jobs for maintaining OnePoint



The steps in these jobs all call SQL-stored procedures, which are located back under the OnePoint database object. For example, the MOMX Partitioning and Grooming job consists of two steps: partitioning and grooming. Each of these steps calls a SQL-stored procedure with a simple command like `EXEC dbo.MOMXPartitioning`. This Transact SQL statement is very straightforward and tells the SQL Server Agent to execute the MOMX Partitioning-stored procedure. The stored procedure then is a much longer set of instructions also written in Transact SQL syntax.

### 7.1.3.1. OnePoint - Update Database

By default, the OnePoint database runs every hour on the half hour and it automatically resolves alerts that have aged beyond the time limit defined in the grooming global settings, which by default is four days. Once an alert's resolution state is set to resolved, it can be deleted from the OnePoint database by the MOMX Partitioning And Grooming job.

### 7.1.3.2. MOMX Partitioning And Grooming

This is the job that actually deletes data from the OnePoint database, the yin to the yang of the inflow of event, alert, and performance dataflow. If the database grooming is configured correctly, there will be a harmonious balance and your database will achieve a Zen-like state of serenity. The job first partitions the data in the database by date, then calculates the cutoff time, and drops the data older than the cutoff point by partitions. This makes for very fast grooming because SQL does not need to search for individual records to delete, it just drops whole tables.

By default, this job runs at midnight daily. If it fails, an event ID 208 is written to the application event log and is picked up by MOM. If for some reason you find that your database has exceeded 60% utilized space and that the weekly reindex job is failing due to lack of free space, you can run



this job manually by right-clicking on it and selecting Start job. In an emergency grooming situation, before you run this job you should reset the "Groom data older than the following number of days" value in the Administrator console from the default of four days to some lesser amount depending on how much data you wish to delete and how much you wish to retain. Don't forget to do a commit configuration change and then reset the value back to the pre-emergency state.

### 7.1.3.3. OnePoint - Reindex

To speed searching, SQL builds indexes for the tables in the OnePoint database. These indexes are rebuilt weekly by this job. This job requires 40% free space in the database, and if it fails there is no great harm done to the functionality of the OnePoint database in the short run. However, if it continually fails, database response time will increase. This job is scheduled for 3:00 a.m. on Sundays, so look for specific alerts in the Operator console referring to application event ID 208 that cite the failure of this job when you get into work on Monday mornings. This job refers to the database maintenance plan OnePointReindex.

It is unlikely that you will have to touch or run the other jobs manually. Briefly, here is what they do:

#### *OnePoint - Computer Maintenance*

This job, which runs every five minutes, checks for computers that have completed their maintenance mode time period, need to come out of maintenance mode, and then do so.

#### *OnePoint - Check Integrity*

This job calls a database maintenance plan called OnePointCheckIntegrity every Saturday night at 10:00 p.m. This job checks for broken links between tables, ensures that there are no broken references, and confirms that all links, references, and indexes are valid.

## 7.1.4. Tuning OnePoint Database Grooming

[Chapter 5](#) introduced the database grooming controls that are available in the global settings of the Administrator console. The default setting is to groom out data that is more than four days old. To figure out if this setting is correct for your organization, you will need to perform a few tasks on a daily basis.

The first thing is to define what "correct" is, and you should already have a good idea based on database sizing calculations. In those calculations, you estimated the amount of data that would be flowing into the database on a daily basis and stated the amount of time you wanted to retain that data before deleting it. This number is an absolute limit because the database is configured not to allow automatic growth, so monitoring the total database size will not tell you how much of the database is actually used. The value that you need to watch is the percentage used. If your grooming values are correctly set for your environment, the percentage used will remain fairly constant over time. Of course, there will be spikes, especially if you deploy a large number of new agents or deploy a new management pack. You may need to make adjustments to the database size for these

deployments.

The MOM 2005 management pack includes performance rules and event rules that monitor the database free space, the status of the database, and the success/failure status of the SQL maintenance jobs; see [Table 7-1](#).

Table 7-1. Out-of-the-box MOM OnePoint database monitoring rules

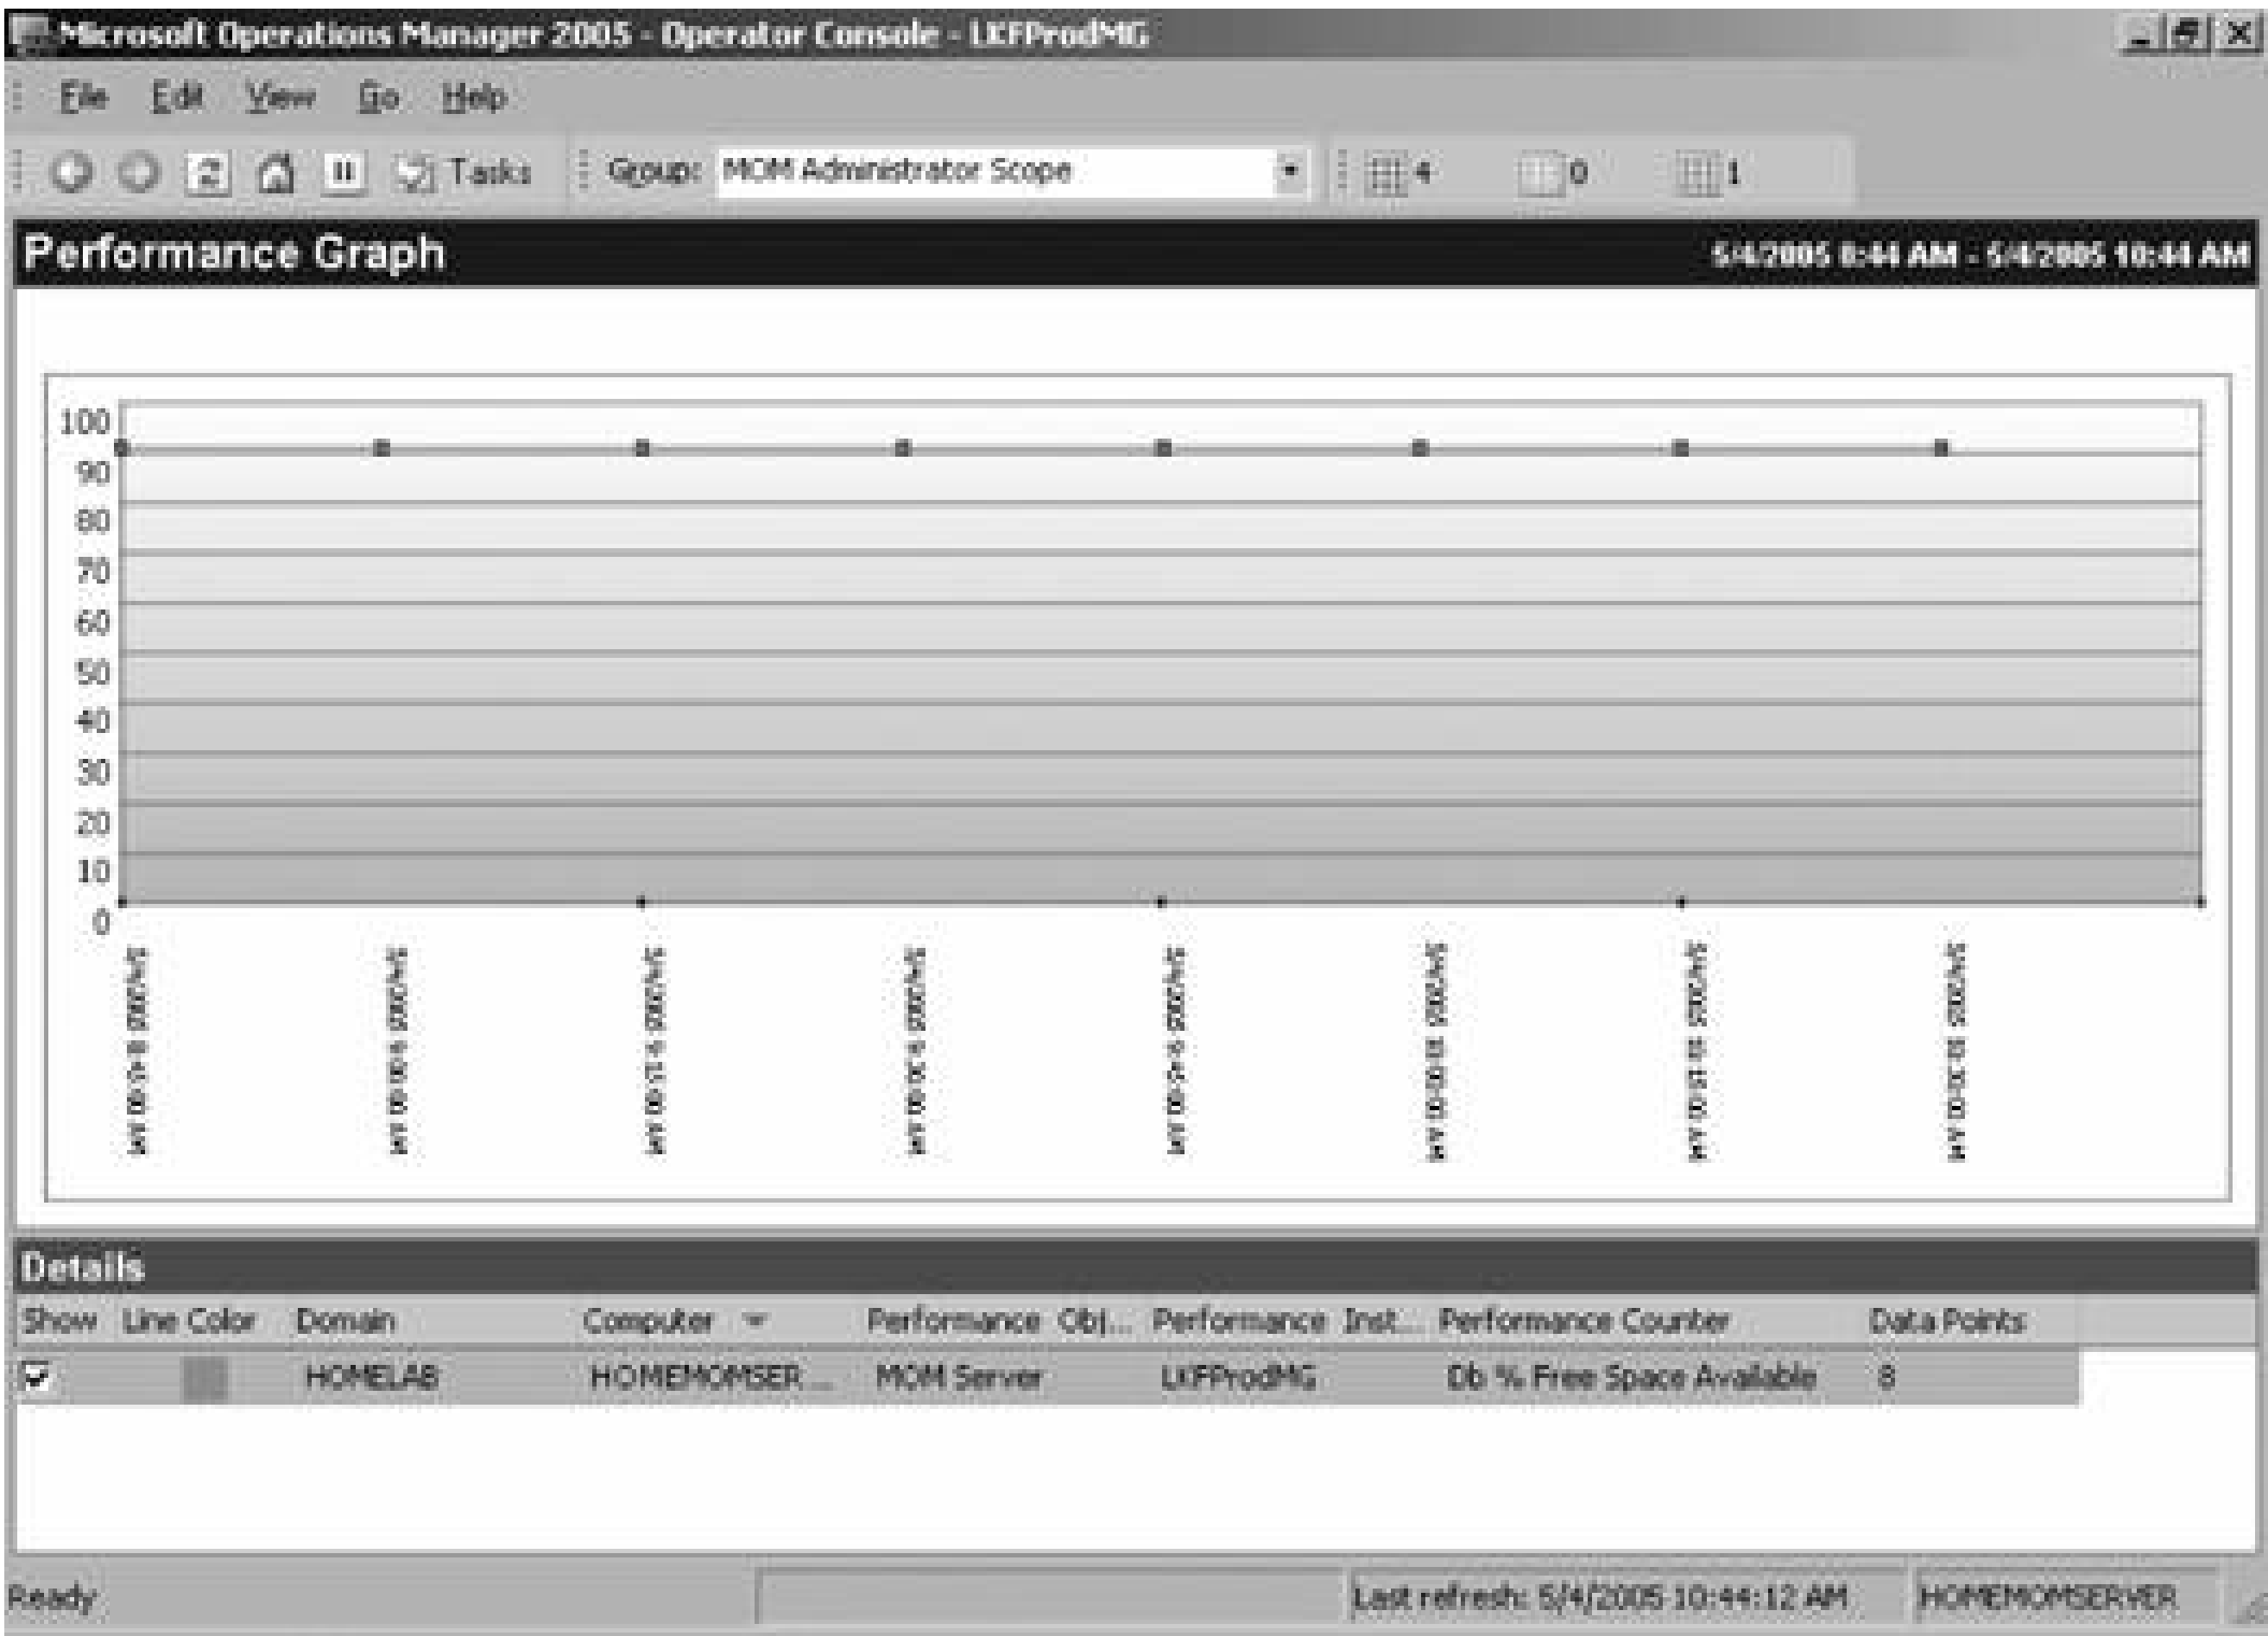
Rule name	Description/Criteria
Performance Threshold: MOM Database Free Space - Warning Threshold	Generates an alert with a severity of Warning when the database free space falls below 40%. Measured every 15 minutes.
Performance Threshold: MOM Database Free Space - Error Threshold	Generates an alert with a severity of Warning when the database free space falls below 20%. Measured every 15 minutes.
A MOM Grooming SQL Server Agent job failed	This event rule looks for event ID 208, with a source name of SQLServerAgent or SQLAgent* where event parameter 1 is MOMX Partitioning And Grooming and parameter 3 is Failed. This rule generates an alert with a severity of Critical Error.
A MOM SQL Server Agent job failed	This event rule looks for event ID 208, with a source name of SQLServerAgent or SQLAgent* where event parameter 1 matches the wildcard OnePoint* and parameter 3 is Failed. This rule generates an alert with a severity of Error.
Check the status of the MOM database	This rule looks for event ID 2600 in the database server application log. Occurrence of this event ID means that the database is for one reason or another not available or that the management server could not interact with the database.
MOM Database State Monitoring	This rule calls a script that confirms four itemsthat the database server can be connected to, that SQL Server is installed and the OnePoint database exists, that the autogrowth database setting is off, and that the authentication mode is set to Windows-only.

These monitoring rules are valuable to have, and you should never ignore an alert they generate, but rules can only respond to a condition that has already occurred. These rules do not give you the information needed to proactively tune your grooming settings, which is what the database free space measurement is.

Obviously, the quickest way to find out your database free space is to look it up in the taskpad view of the SQL Enterprise manager, but that only gives you a single point in time view. Preferable is a preconfigured performance view (see [Figure 7-4](#)) in the Operator console, which is found in the Performance view group      Microsoft Operations Manager      Operations Manager 2005  
Database Performance      Operational Database Free Space.



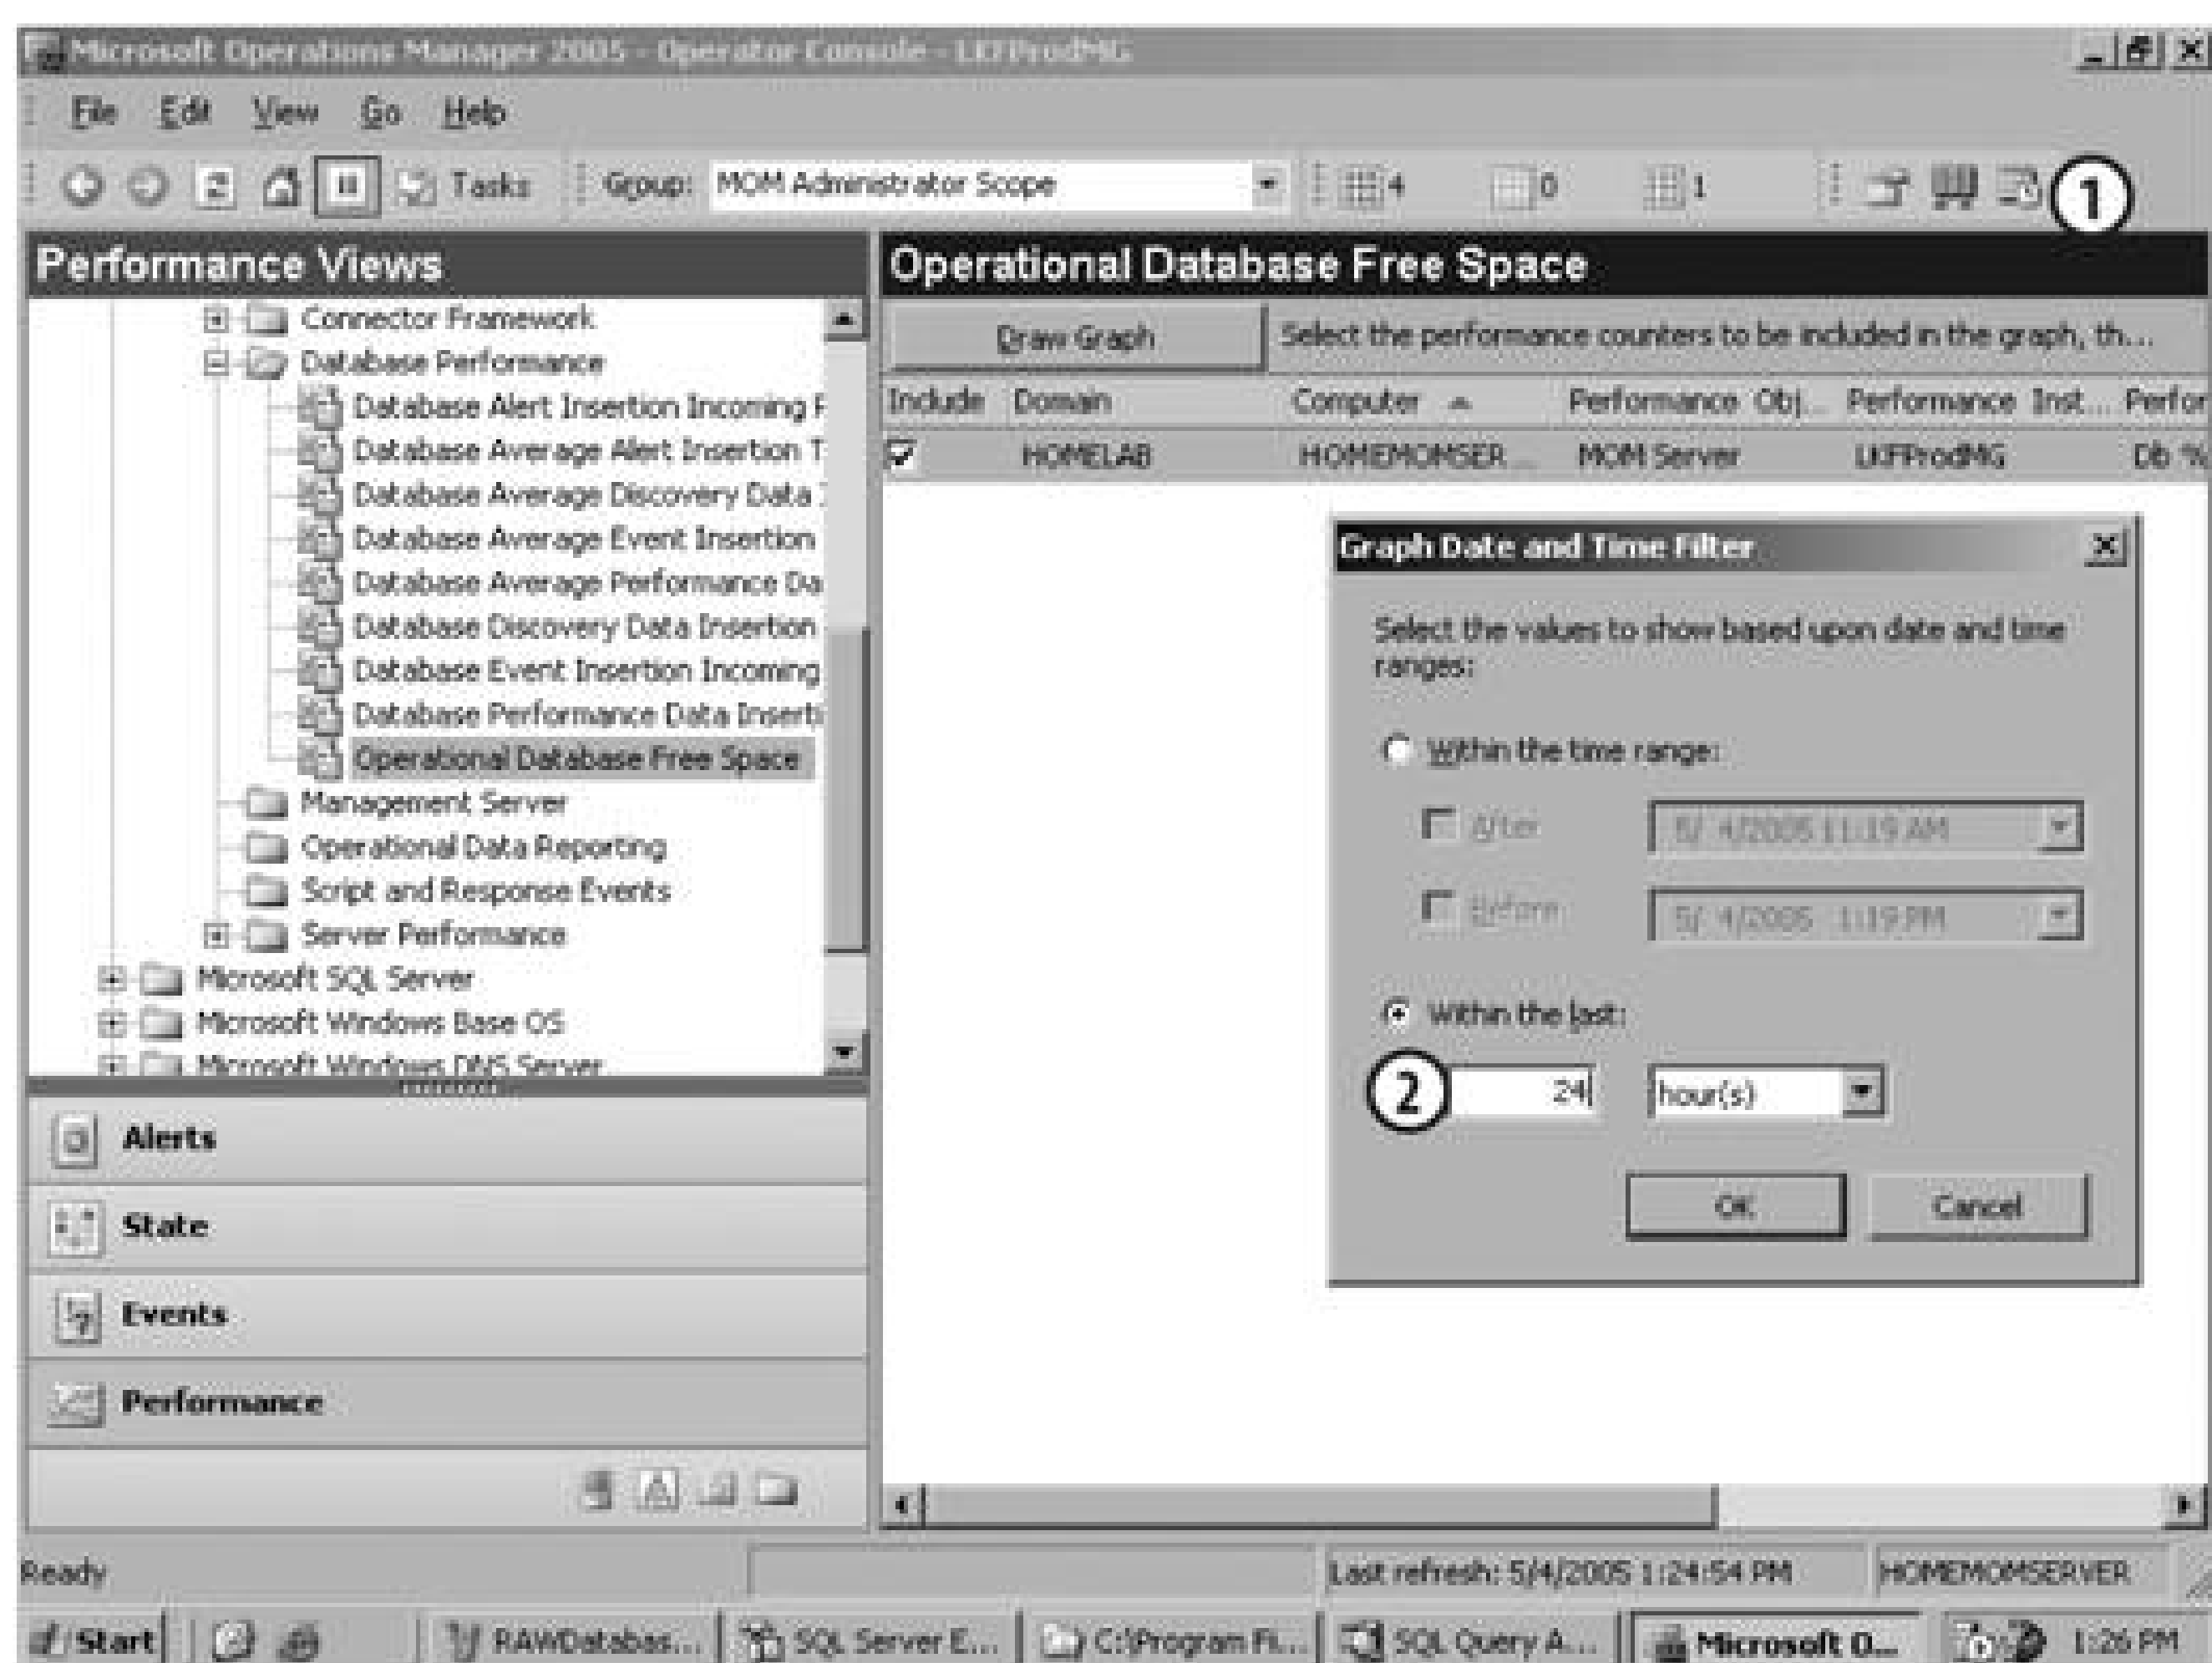
Figure 7-4. A view over the past two hours of the OnePoint database free space



This performance chart displays the last two hours of data, but you can display more by modifying the time filter (see [Figure 7-5](#)). You can go back as far as the data is retained in the OnePoint database. This control (point 1 in [Figure 7-5](#)) is available on the view toolbar. With this tool you can set the time range for the interval between two points or simply for the past number of hours up to the limit of the retained data. The past 24 hours of data is viewed here (point 2 in [Figure 7-5](#)).

You can save these graphs to the clipboard (right-click on the graph and select "copy to clipboard") if you want to keep a historical record of database free space to establish a baseline to measure against.

Figure 7-5. Adjust the time range of the data you want to see with this filter



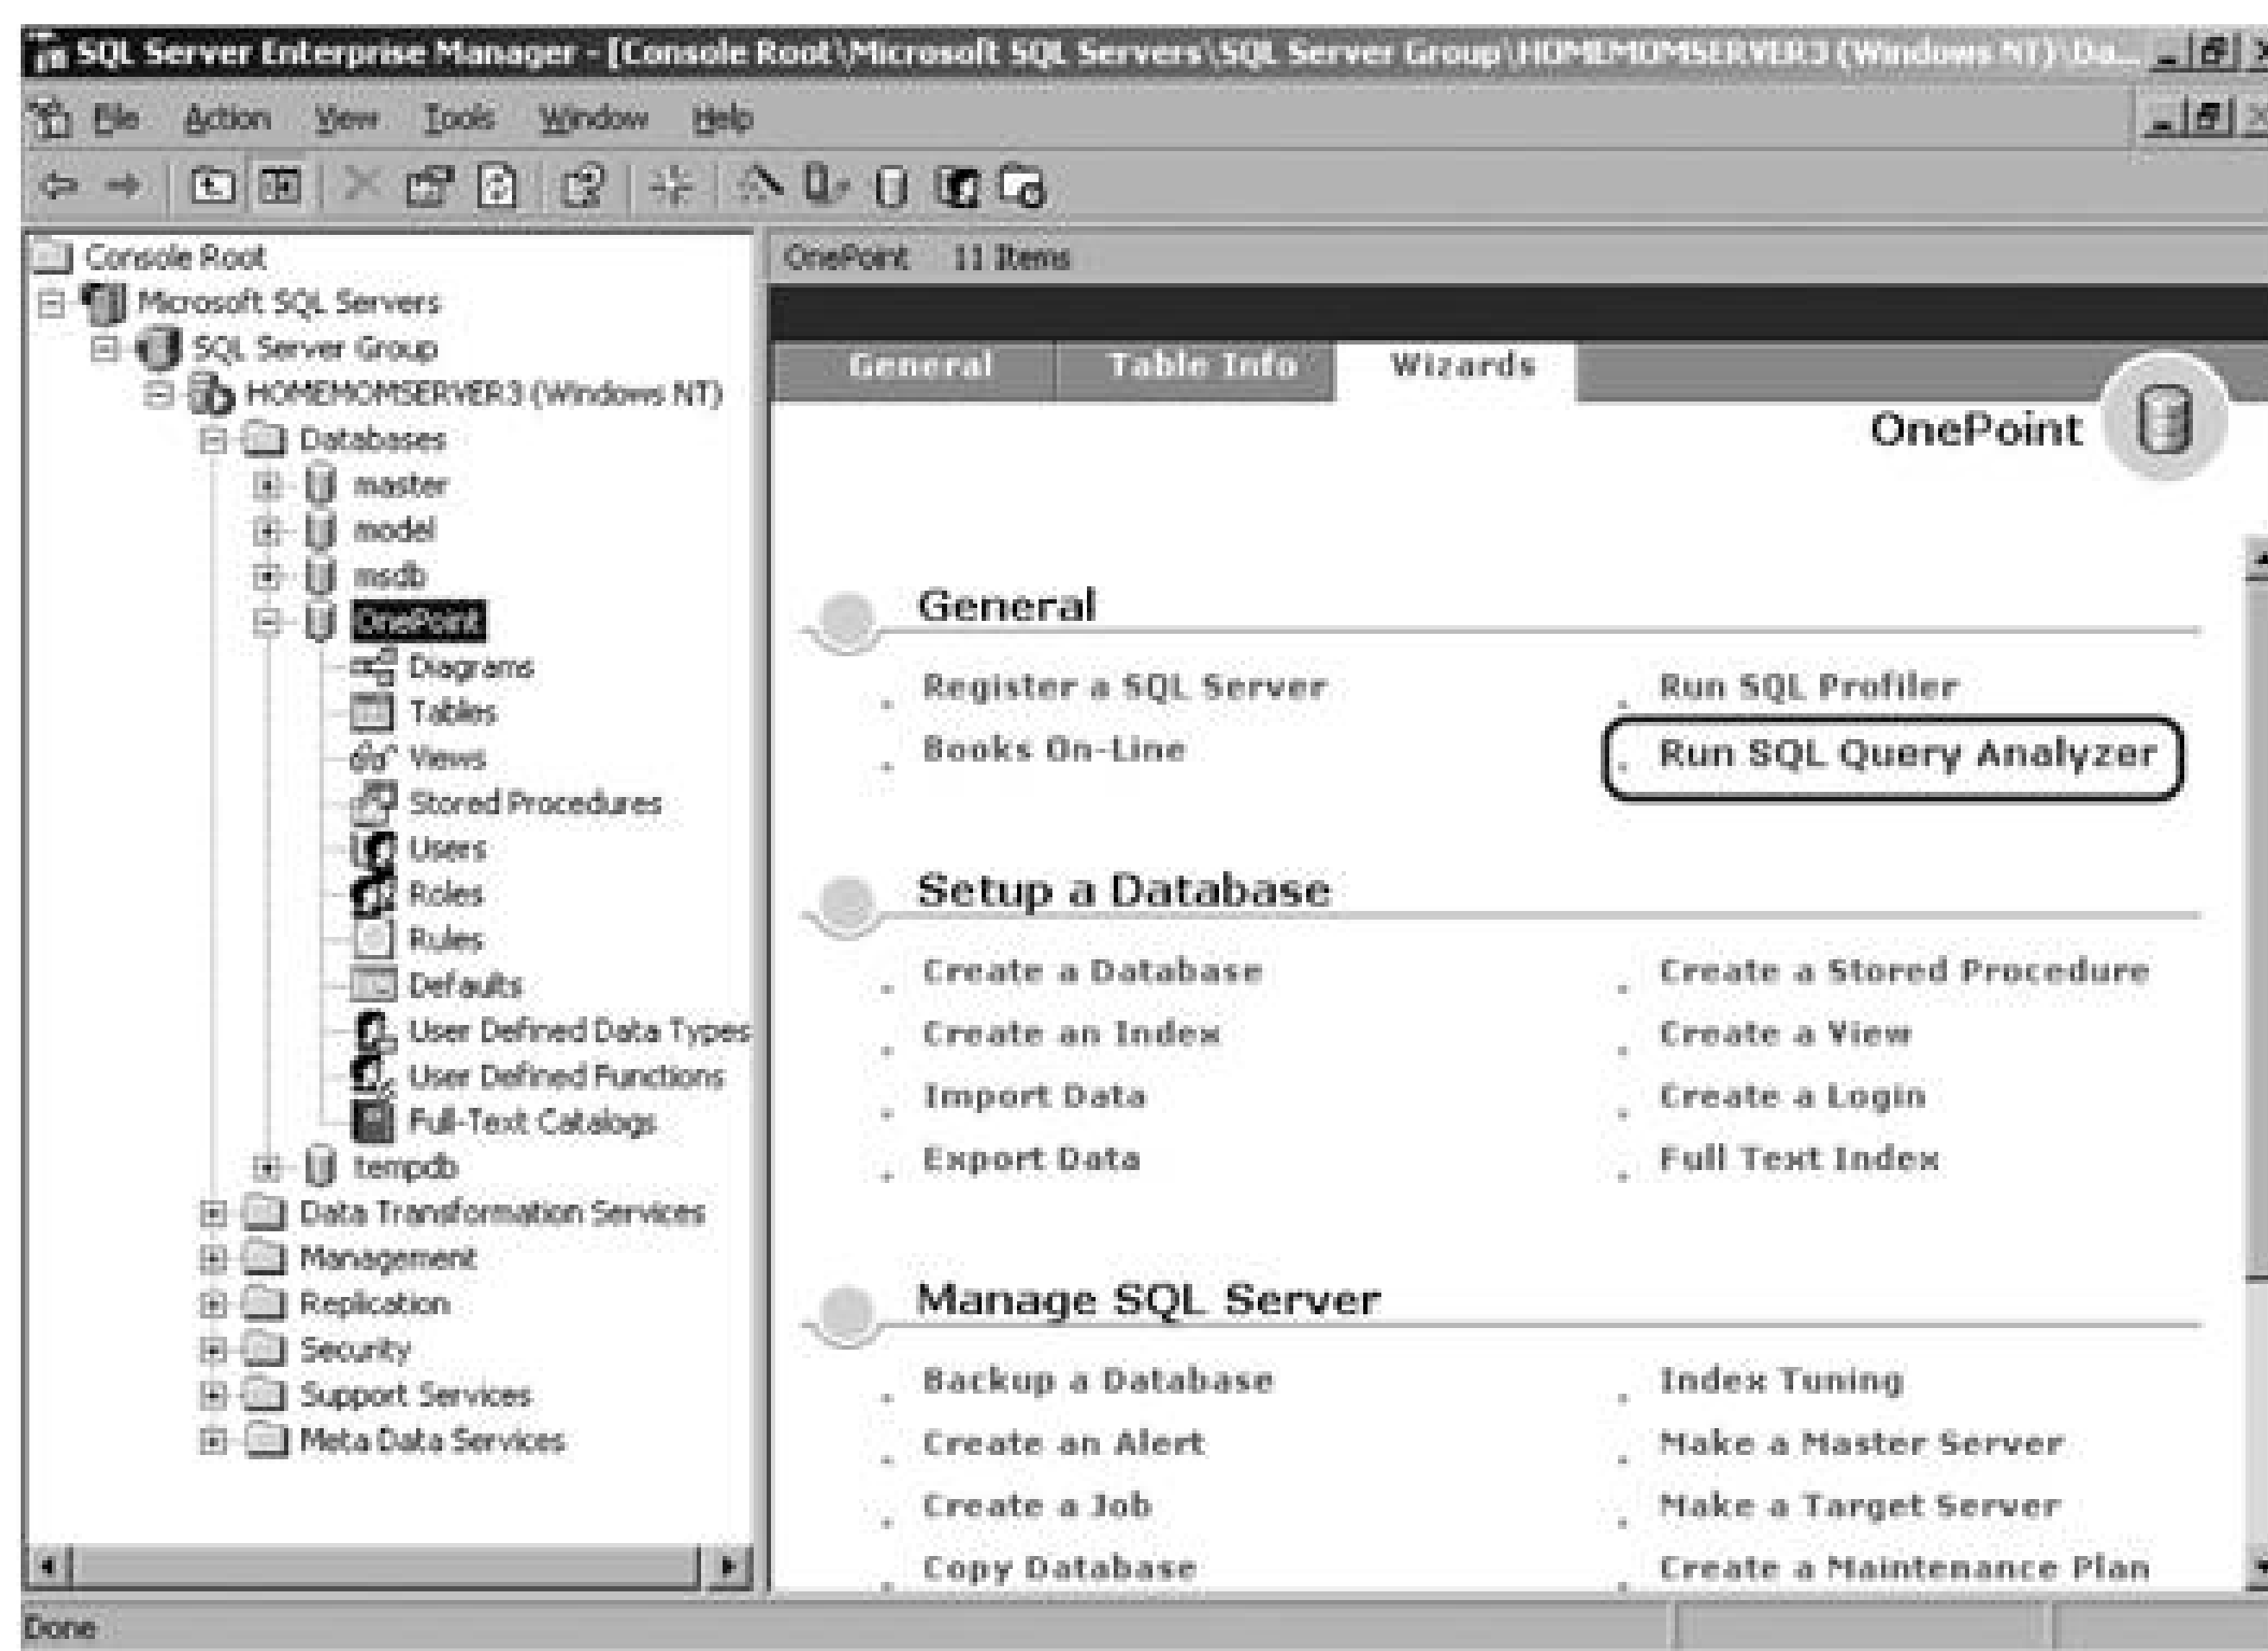
If you find that over time the amount of free space is steadily shrinking and it is starting to approach the 40% free space threshold, then this is an indication that more data is flowing in than is being groomed out on a nightly basis. Start by reducing the grooming setting by one day (to three) and carefully watch the results.

In addition to monitoring the status of the database, the success or failure of jobs, and the amount of free space in the database, you should know what types of data are waiting to be groomed out and how much data there is. With this knowledge, you can decide to adjust the auto-resolve values for alerts in the database grooming settings. You do this by interacting directly with the database via the SQL Query Analyzer . This tool allows the targeting of Transact SQL (TSQL) statements against specific databases or tables.

You can access this tool either through the start menu or through the wizard page of the taskpad view of the database in the SQL Enterprise Manager. If you launch the SQL Query Analyzer from the OnePoint taskpad view (see [Figure 7-6](#)), the focus of the analyzer is automatically set to the OnePoint database. Using this method ensures that the analyzer is connected to the desired database on the desired server, which is an advantage in an environment with multiple SQL Servers, each with many databases.

Figure 7-6. Launching the SQL Query Analyzer from the SQL Enterprise Manager guarantees correct focus



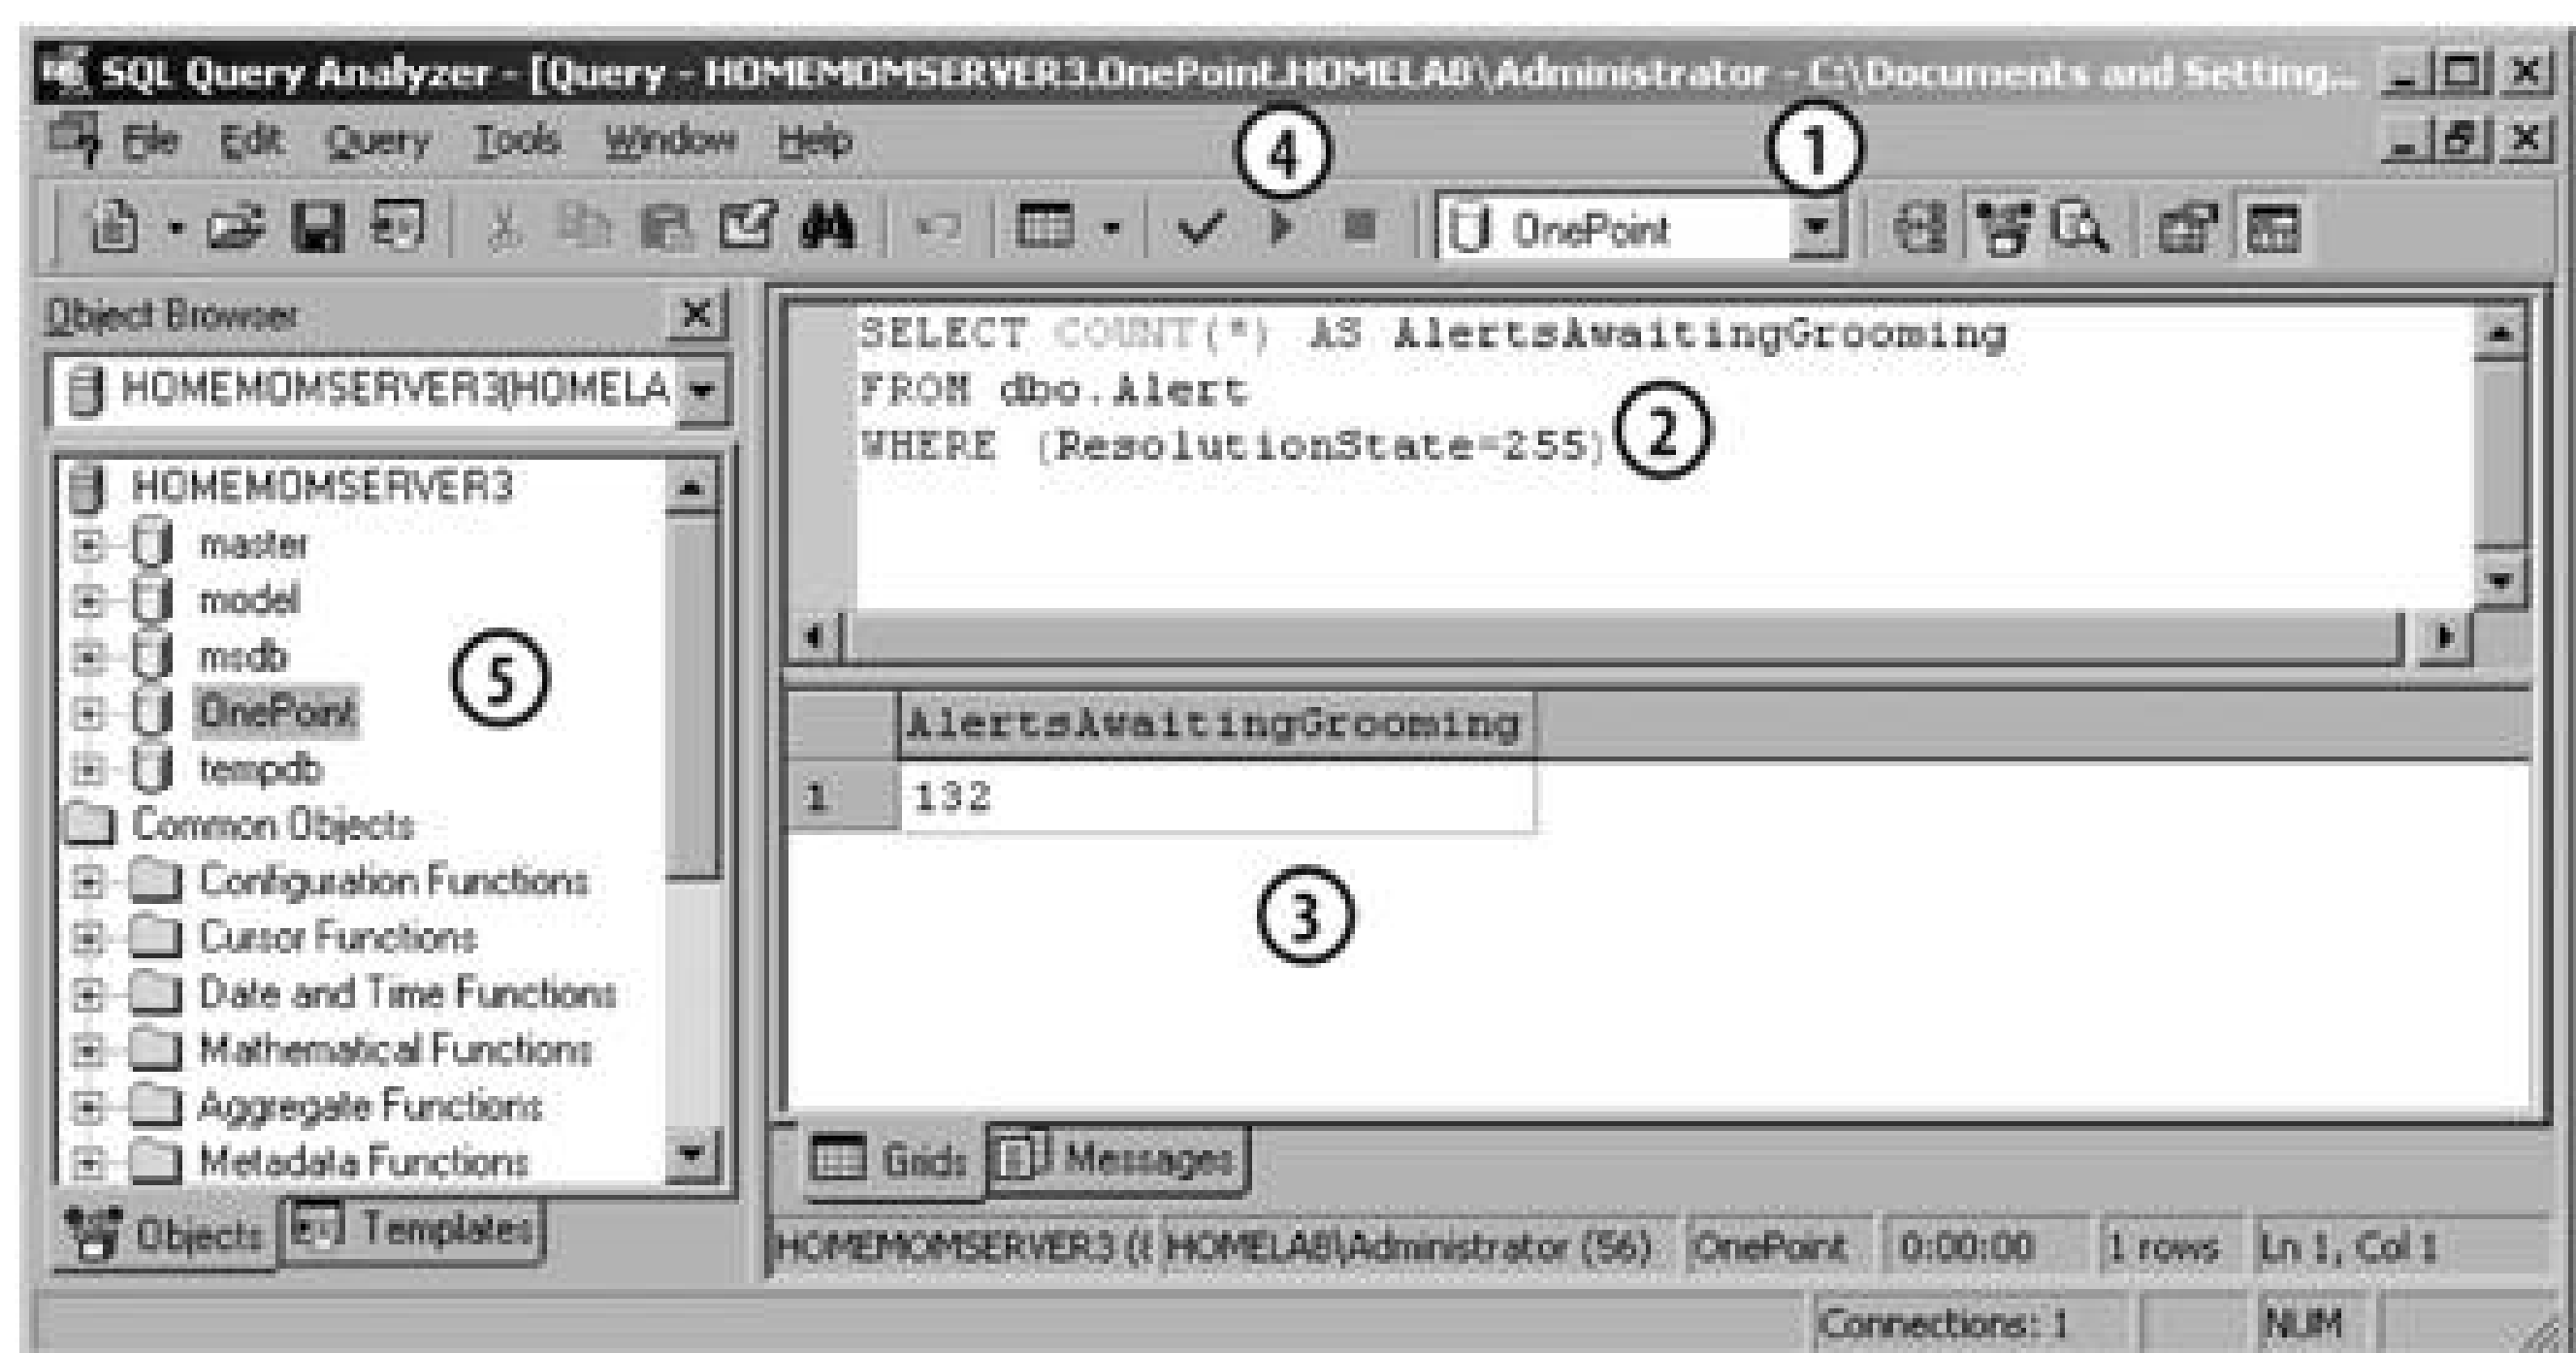


The SQL Query Analyzer is a sophisticated tool with powerful features, but as a MOM administrator, there are five basic parts you need to be familiar with (see [Figure 7-7](#)):

1. Current database. This drop-down list shows you the database that has the focus of the analyzer and allows you to switch focus to a different database.
2. Query pane. TSQL statements to be evaluated are entered here.
3. Results pane. Any output from running the query in the query pane is displayed here.
4. Execute Query button. Clicking on this green arrow runs the code in the query pane against the selected database
5. Object browser. This opens by default and you really don't need to use it. Just don't get confused into thinking that you can change the selected database in this pane.

In the MOM 2005 Operations Guide and in the MOM 2005 Resource Kit, Microsoft has included a few SQL queries that you can use to find out the number of alerts, events, and individual alerts that are waiting to be groomed. In [Figure 7-7](#), the TSQL statement shows how many alerts are ready to be groomed out from the OnePoint database on *homemomserver3*.

Figure 7-7. The SQL Query Analyzer interface



This chapter does not teach you TSQL, but it is not too hard to decipher a simple query like the one in point 2:

- Line 1 is telling the analyzer to count the total number of rows in the table that match the query parameters and prepare to display them as the text string AlertsAwaitingGrooming. This text string could be anything, as long as it is text.

```
SELECT COUNT(*) AS AlertsAwaitingGrooming
```

- Line 2 is telling the analyzer which database table to perform the `SELECT` and `COUNT(*)` action on; in this case it's the `dbo.Alert` (a DBO is a database owner). You will find the Alert table object in the OnePoint database-Tables in SQL Enterprise Manager.

```
FROM dbo.ALERT
```

- Line 3 is the qualifying statement of the query and is telling the analyzer to include only those rows in the count that have a value of 255 (which is the "resolved" state) in the ResolutionState column of the table:

```
WHERE (ResolutionState=255)
```

This query happens to find that there are 132 alerts waiting to be groomed out (point 3 in [Figure 7-7](#)). You can run this type of query at any time, or you could even set up a SQL job to execute these on a regular basis. Collecting this type of information over time is useful because you can use it as a baseline for spotting trends in the type of data being groomed out. A sudden spike in the number of AlertsAwaitingGrooming can be a tip-off that one of two things has occurred: there was a spike in the

number of alerts generated, or something has gone wrong in the grooming job and further investigation is required.

All queries can be saved as a text file with a *.sql* extension so you don't have to type the TSQL statements every time you want to run them. By default, they are saved to the My Documents folder of the machine you are running the SQL Query Analyzer on. The "Enumeration of all alerts awaiting grooming" is another query that can give you further insight on the types and volume of data waiting to be groomed out. This query outputs the name, description, alert level, repeat count, who resolved the alert, and the resolution time for each individual alert that is waiting to be groomed.

```
SELECT Name, Description, AlertLevel, RepeatCount, ResolvedBy, TimeResolved
FROM dbo.Alert
Where (ResolutionState=255)
```

Once you have your grooming setting tuned, you only need to keep tabs regularly on these values for trending purposes. You shouldn't adjust the grooming settings unless the data you are collecting indicates that you should.

The OnePoint database and other critical information on your MOM database server should be backed up on daily basis. The next section in this chapter outlines what you need to back up and why, how to perform the backups, and how to perform a restore.

## 7.1.5. Backing Up SQL Databases

Everyone knows it is important to have a backup of your data/OS/systems configuration in case of hardware failure, data corruption, etc. The backup restores the system to its pre-failure state and hopefully picks up exactly where you left off. So, consider this the standard "Do your backups" admonishment. Telling an IT professional that he needs to do backups is like having to tell a builder that she needs to put a roof on a house.

The curious thing about recovery with MOM is that it is more important to preserve the configuration of the management group and all of its components than it is to preserve the operational data. This is because all of the operational data was generated by the monitored environment, so if the OnePoint database is lost, all of the conditions that caused the alerts to be generated would still exist and the alerts would generate again once MOM is running again. On the other hand, the configuration data tuned the management packs, and all of the historical data only exists in the OnePoint database. If you lose it and you have no backups saved anywhere else (such as in a synchronized preproduction environment), then you will have to start your MOM implementation from scratch.

That being said, since the two types of data are in the same database, you get both for the cost of a single database backup. In addition to backing up the OnePoint database, you must also capture backups of your management packs and report definitions (if you have reporting installed), the master and msdb systems databases, custom files like *Manual/MC.txt*, and any file transfer server files. Critically important, but seldom remembered, is to capture the OS configuration for the management servers and database servers. Specifically, the physical and logical drive configuration, the accounts and permissions with local rights, and the accounts used for the DAS account and the management server action account. Also, you need to track the permissions granted to SQL logins.



The next section talks about the planning steps for a backup and how to actually perform a backup using a combination of SQL and system backup tools. The restore process is discussed and the necessary steps outlined. Now, for the standard backup/restore disclaimer: you have to adapt the following information for your environment. Test your backup plan by taking your production backups and using them to re-create your management groups in an isolated environment. If you don't do this for practice, then the only time you will test these procedures is in production during an outage, when you are under quite a bit of stress.

### 7.1.5.1. Backup tools

To get a complete backup of all components that will be needed for recovery, you need to use both the SQL backup utility and an OS-level backup utility. For the purposes of this discussion, the utilities that are built into the products will be used. There are many third-party SQL/OS backup solutions from companies such as BMC, VERITAS (recently purchased by Symantec), CA, IBM (Tivoli), and so on. You should consider these tools when you need to back up databases in addition to the OnePoint and reporting databases.

### 7.1.5.2. What to back up

The first step in planning your MOM backups is to identify what you need to back up:

#### *Server OS configuration*

The starting place for this information should be your OS installation documentation. Be sure to record major variances from this build standard over the life of the server. This is likely to be a manual process. Things you should capture here include: the logical drive configuration, application installation paths, server names and IP configuration, OS service pack level, installed hotfixes, and any optional OS components. This is also a good place to record the management group name and the MOM account information (DAS and management server action account).

#### *OnePoint database*

Use the SQL backup utility to make scheduled backups that dump to a tape device or a file on disk. If you dump to disk, you can then use the OS-level backup utility to back up that file. SQL backups automatically get both the targeted database and its transaction logs. This will also capture all installed management packs, so those will be available at restore time.

#### *SQL master and msdb databases*

These are both necessary to restore SQL and the SQL configuration. Use the SQL backup utility.

#### *Management packs*



Management packs are the most important objects to protect, since you will put more time into configuring and tuning them than any other part of the management group. If you follow the procedures in [Chapter 4](#), you will have a solid backup plan for the management packs. It is more likely that you will have to roll back individual management packs than restore the entire OnePoint database, so having a separate backup routine for these is valuable.

### *Report definitions*

If you change the report definitions from the defaults, a process that requires coding in Visual Studio, then they need to be backed up. You can export them individually using the report utility (*rptutil.exe*), or simply back up the reporting database to capture all report definitions.

### *ManualMC.txt files*

Remember that the *ManualMC.txt* file is used to list the computers that you want a management server to be responsible for and install agents on. This text file is placed on the management server and is scanned during the nightly computer discovery scan (2:05 a.m.). If you use this method for installing agents, then you should capture these files through a file-level backup or simply by copying them off. You could even keep them with current management packs in the MPTransfer folder.

## 7.1.5.3. When not to back up

In general, you want to avoid scheduling your database backups during times when the database is especially busy, or when a process puts a lock on a database table, such as when the grooming SQL jobs run. The following are the jobs and their schedules that you want to avoid. These are the default times, so if any of these have changed in your environment be sure to check the job schedules in SQL Enterprise Manager Management SQL Server Agent Jobs.

- MOMX Partitioning And Grooming: runs daily at midnight.
- OnePoint - Check Integrity: Runs every Saturday night at 10:00 p.m.
- OnePoint - Reindex: Runs every Sunday at 3:00 a.m.
- OnePoint - Update Database: Runs every hour on the half-hour. You can't really schedule around this one, but try to avoid performing backups during the time periods when all the other SQL jobs are running as well as this one.
- SCDWGroomJob: This is the grooming job for the data warehouse. You only need be concerned with this on the data warehouse server. It runs daily at 3:00 a.m.
- Daily DTS transfer: If MOM reporting has been installed, this task is scheduled in the Windows Task Scheduler on the MOM Reporting server. It transfers data from the OnePoint database to the data warehouse daily at 1:00 a.m.

### 7.1.5.4. Creating a SQL backup job

Looking at the times listed in the previous section, it seems that the best time for the SQL database backups to start is around 8:00 p.m. For the sake of speed, the SQL backups will be configured to write to disk-based backup devices on the SQL Server itself. This means that there must be enough local disk space to accommodate this. It is a best practice to write the SQL backups to a different logical disk than the ones the database and transaction logs are on.

After the backups complete successfully, create and schedule an OS-level backup that captures the SQL backups along with any other custom files that may exist on the SQL Server or management servers. This OS-level backup will write to a network share, thus copying the SQL backups off of the SQL Server itself.

Before creating and scheduling the SQL backup jobs, some preparation work needs to be done. This involves defining the backup routine and creating SQL *backup devices*. In SQL, a backup device refers to a physical object like a tape drive or a file on the local hard drive. Regardless if you are backing up to disk (a file) or a tape (an actual physical device), it is useful to create backup devices because in SQL they are recorded in the master database and are, therefore, persistent and reusable. You will be able to refer to them when you create the backup jobs rather than having to supply a fully qualified path to a file along with the filename as the destination for the backup every time. [Table 7-2](#) summarizes what will be backed up, the backup name, the frequency of the backup, and the backup device name. The plan is to take a full backup of the OnePoint database nightly and on Sunday nights take a full backup of OnePoint, master, and msdb databases.

This is a relatively simple backup plan; you may want to alter it based on your company's needs. For example, to avoid data loss throughout the day, each database's transaction logs can be backed up at regular intervals.

Each nightly backup will have its own SQL job. This pattern will be repeated weekly, with each job overwriting the previous week's files. That means there should never be more than seven SQL backup files in the destination directory.

Table 7-2. MOM database server SQL backup routine

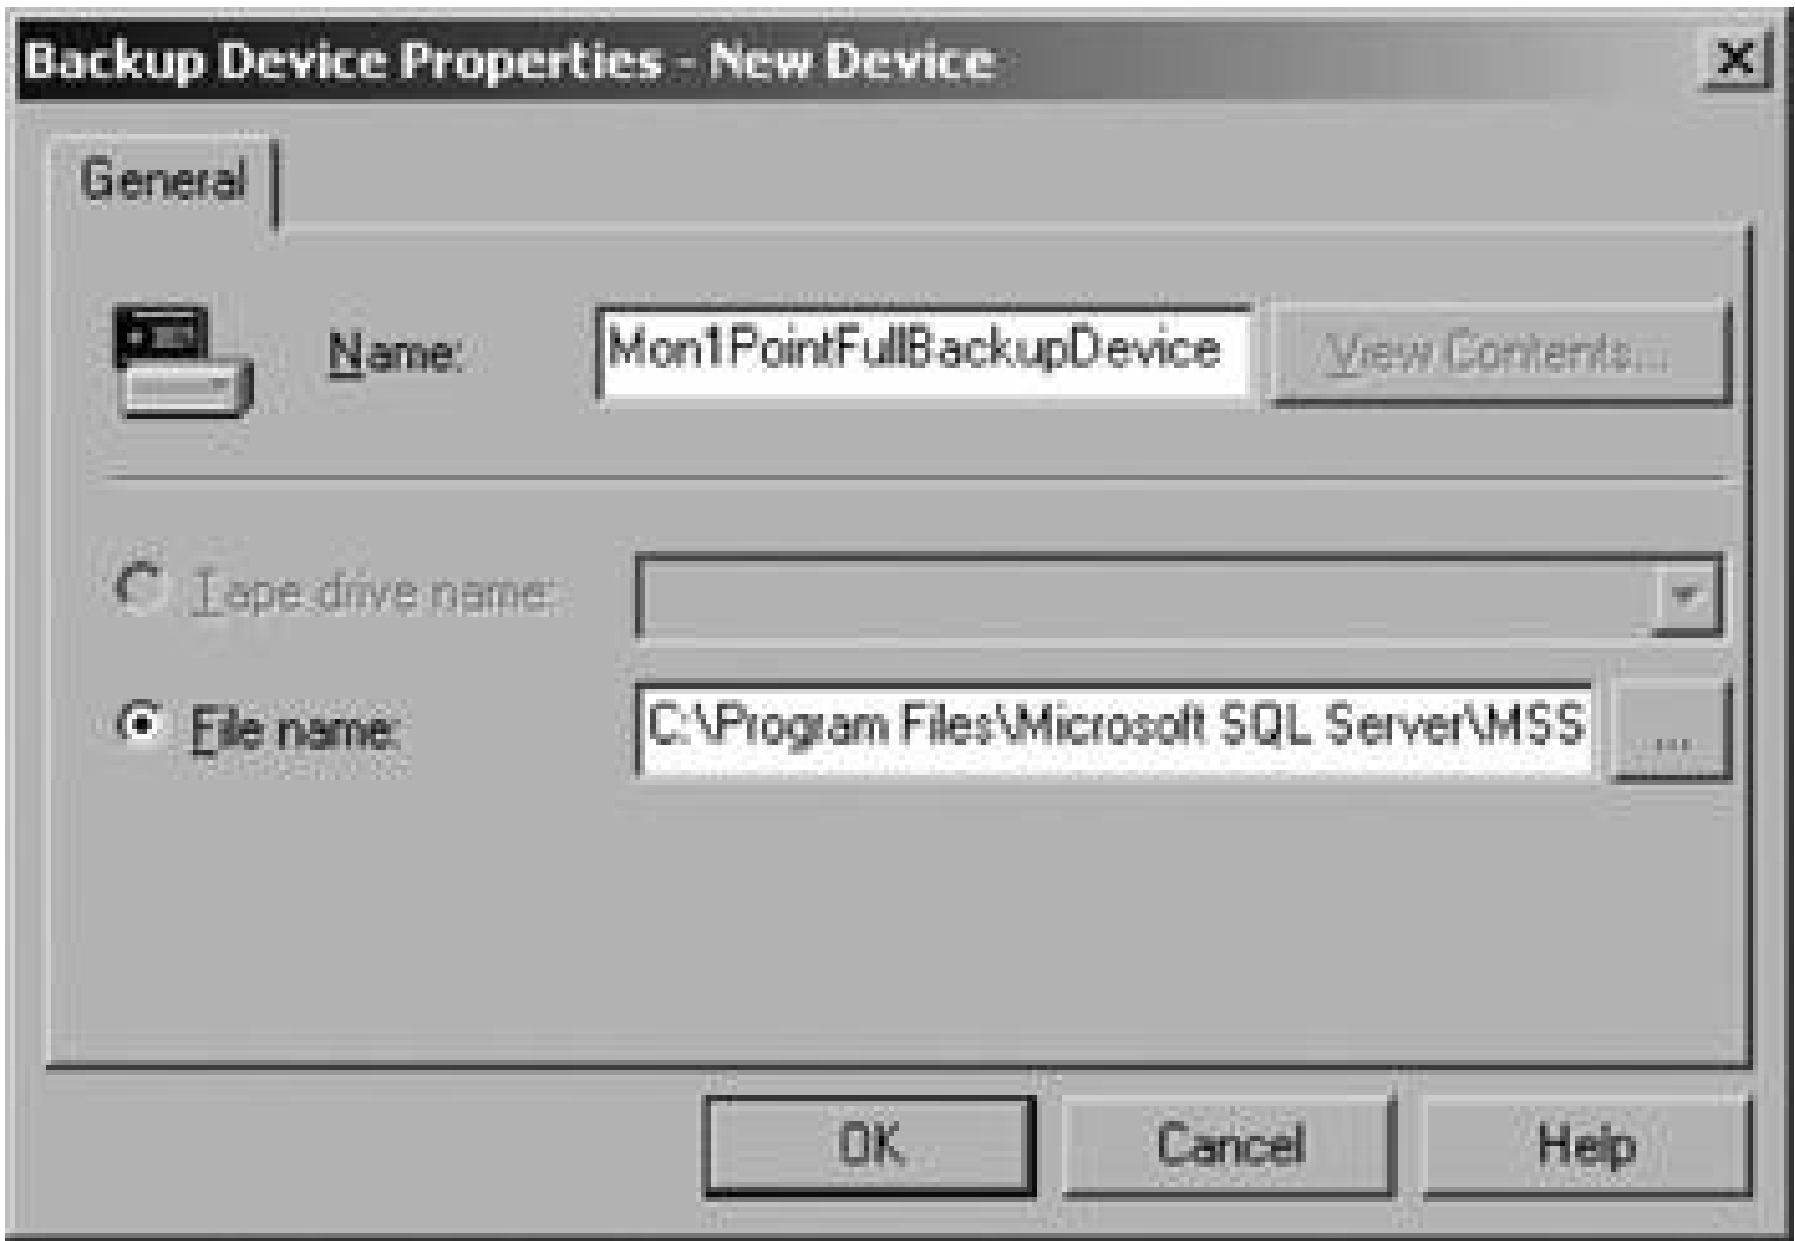
Backup name/SQL job name	Backup device name
Mon1PointFullBackup	Mon1PointFullBackupDevice
Tues1PointFullBackup	Tues1PointFullBackupDevice
Weds1PointFullBackup	Weds1PointFullBackupDevice
Thurs1PointFullBackup	Thurs1PointFullBackupDevice
Fri1PointFullBackup	Fri1PointFullBackupDevice
Sat1PointFullBackup	Sat1PointFullBackupDevice
Sun1PointFullBackup	Sun1PointFullBackupDevice



Backup name/SQL job name	Backup device name
WeeklyMasterFullBackup	WeeklyMasterFullBackupDevice
WeeklyMSDBFullBackup	WeeklyMSDBFullBackupDevice

To create a backup device, open the SQL Enterprise Manager and navigate to the SQL Server Management  $\rightarrow$  Backup container. Open the context menu for the Backup container and select to create a new backup device. This brings up the backup device creation page (see [Figure 7-8](#)). Here, enter the name of the backup device as specified in [Table 7-2](#). For this example, the Mon1PointFullBackupDevice is created. The default path for all backups, whether to a file or a device on disk is *<InstallationDrive> \Program Files\Microsoft SQL Server\MSSQL\Backup\<file or device name>*. For a device on disk, a file is created consisting of the device name appended with the *.bak* extension.

Figure 7-8. Creating a reusable SQL backup device on disk



Repeat this process for each of the backups defined in [Table 7-2](#), resulting in nine separate backup devices.

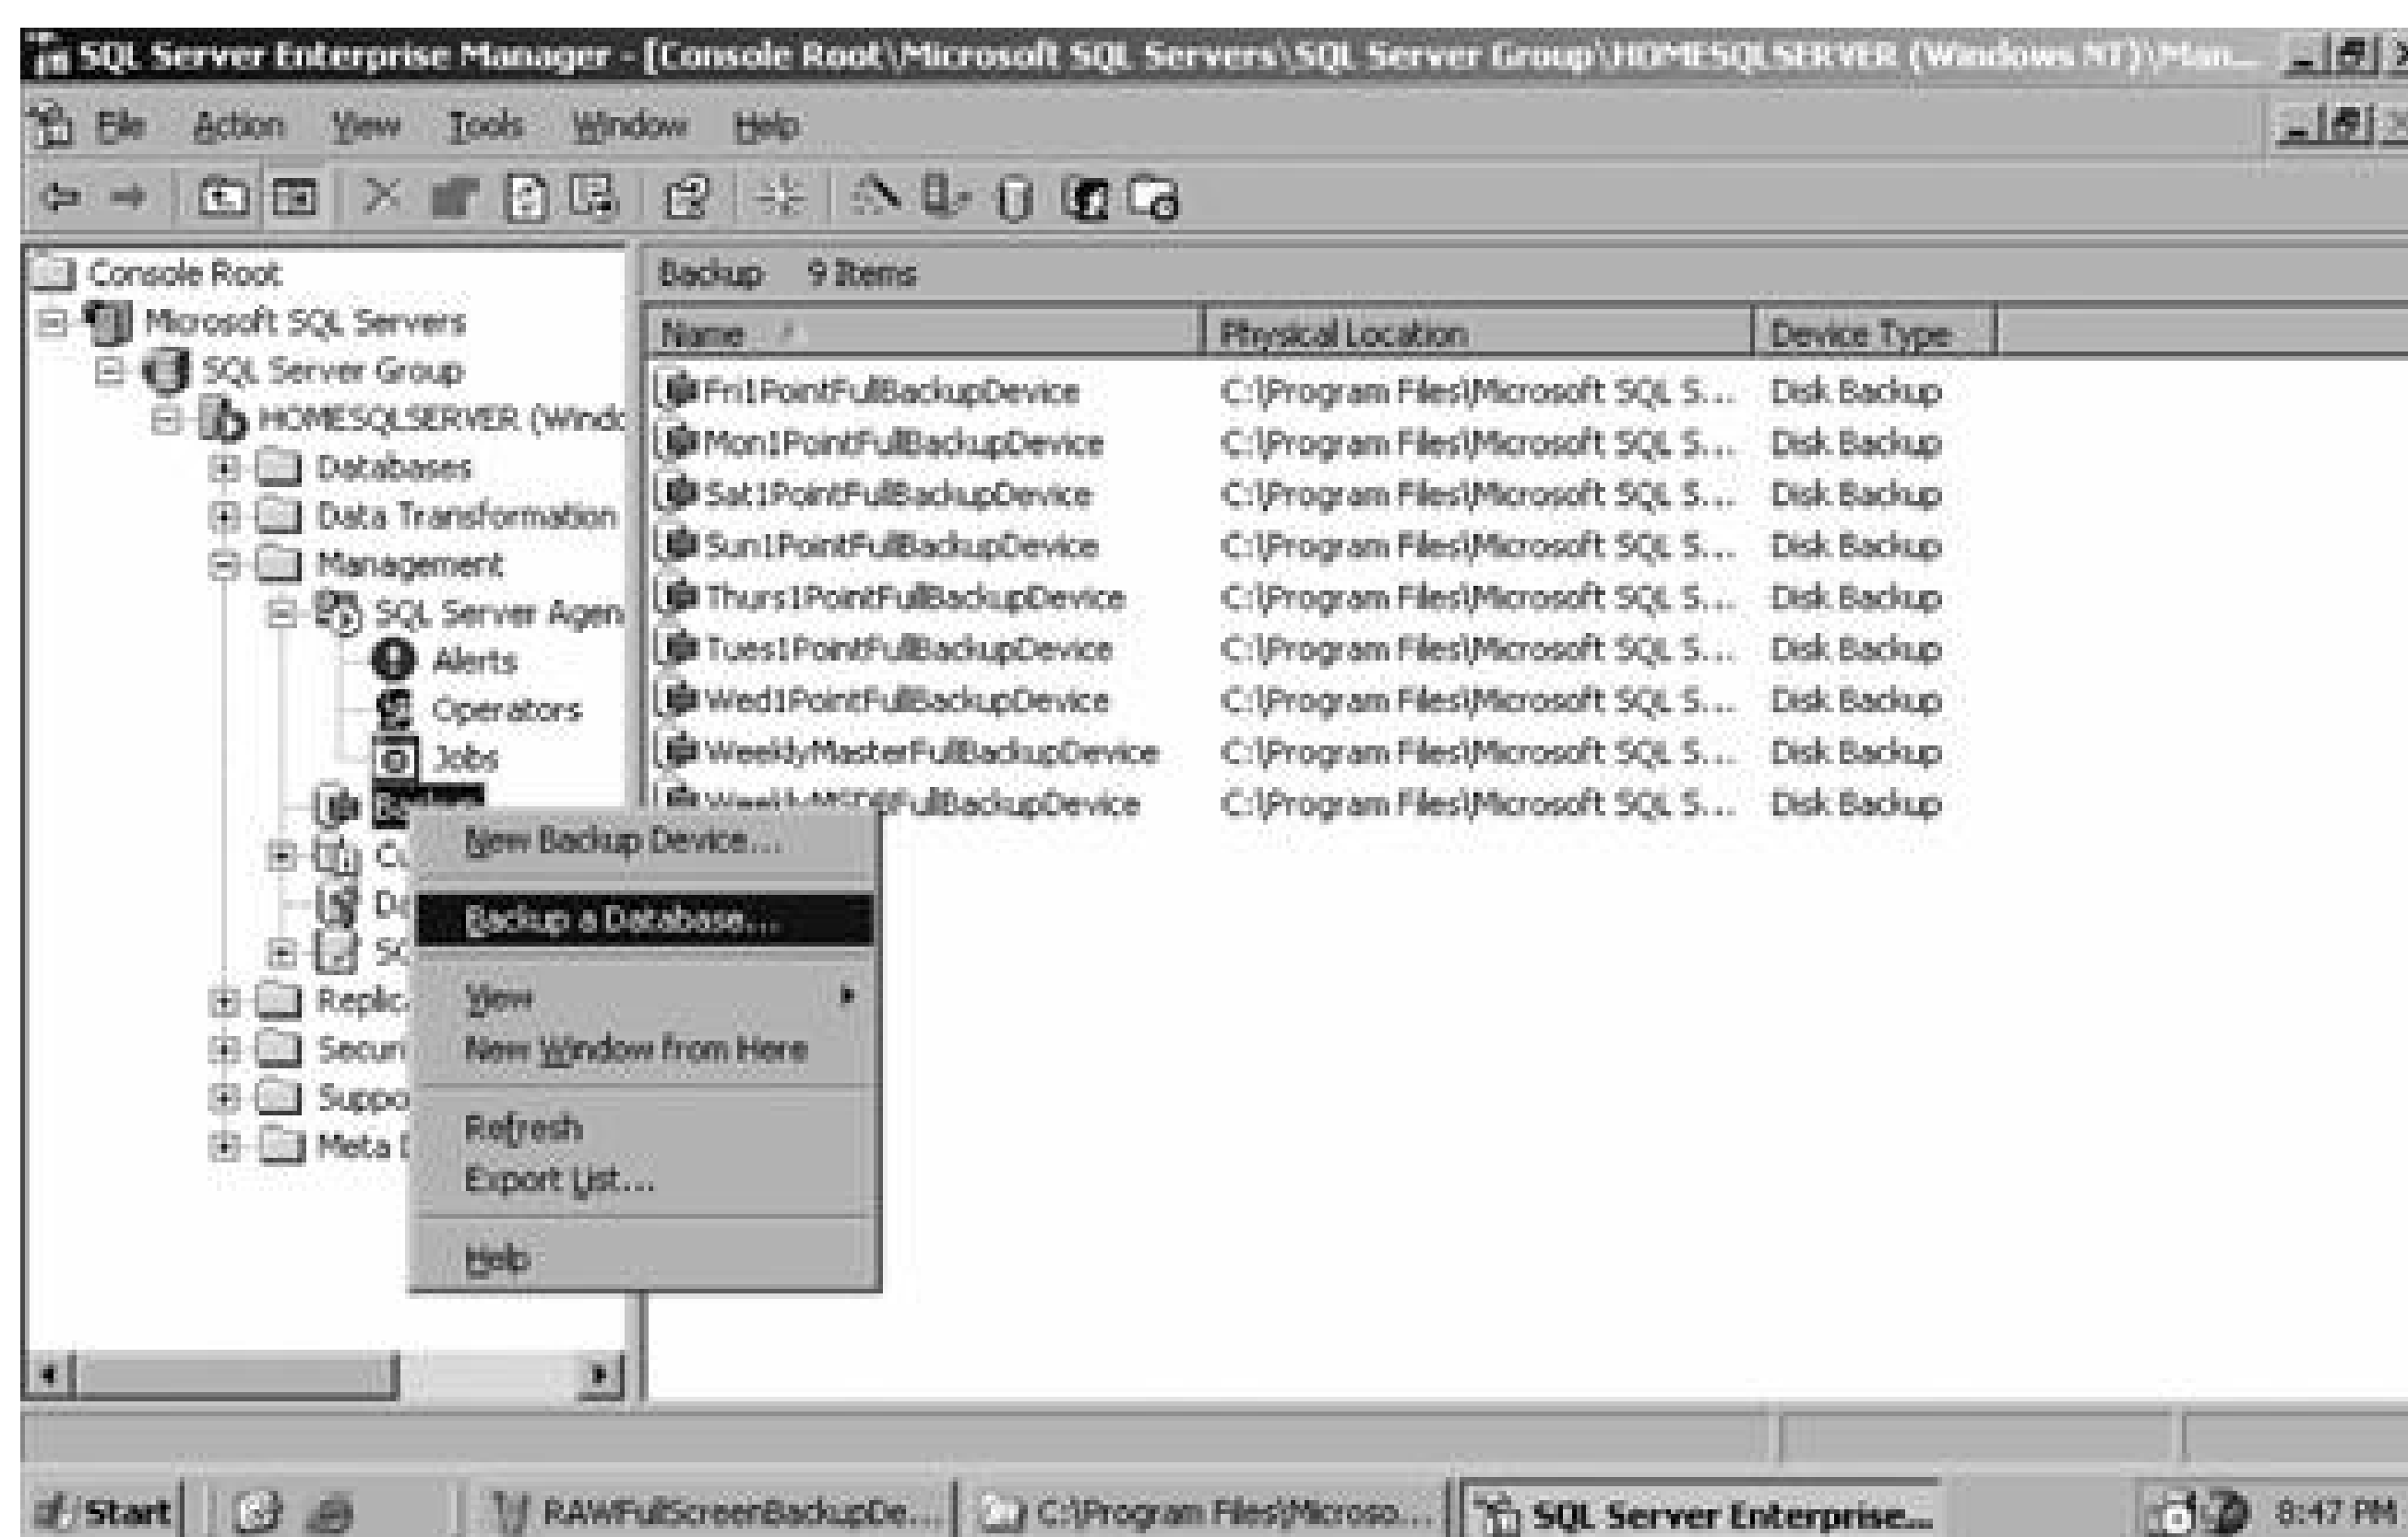
Now that the devices are defined, you can create the backup jobs. In SQL, the definition of what to back up, where to back it up to, and how often to perform the backup is all part of the same process. There are several different ways to create and schedule a backup. You can create the backup from the context menu of the database object, from the Database Taskpad View Wizards Backup a Database or from the Tools Wizards Management Backup a Database. For the sake of this example, the backup is created by bringing up the context menu for the Backup container and selecting the Backup a Database option (see [Figure 7-9](#)).

This brings up a two-tab page in which you define all of the backup parameters. On the general tab (see [Figure 7-10](#)) make selections at the following points:



1. From the drop-down list, select the database that you want to target with this backup job.
2. Enter the name for this backup as defined in [Table 7-2](#). The text entered here will be the name of the job when it appears in the Jobs container.

Figure 7-9. Creating a backup definition in the same place the backup devices were created



3. Enter an appropriate description.
4. Select the desired type of backup. Performing a full backup every night may seem like overkill, but it really is the best option. Selecting differential (which only backs up what has changed since the previous backup) won't save much time if you are backing up a database that is restricted in size to local disk (which is very fast).
5. In the "Backup to" destination box, select the Add button and then designate the Mon1PointFullBackupDevice (see [Figure 7-11](#)).
6. Since this backup routine calls for keeping one week's worth of backups, select to overwrite the previous week's version of the backup.
7. Select the schedule and then click on the ellipsis to set it. This takes you to the Edit Schedule page ([Figure 7-12](#)) where you enter a name for the schedule and then click Change to set the parameters in the Edit Recurring Job Schedule page (see [Figure 7-13](#)).
8. On the Edit Recurring Job Schedule page, configure the job to run weekly, on Mondays at 8:00 p.m., with no defined end date.

Once finished with the General tab, the values on the Options tab need to be edited (see [Figure 7-](#)

14).

To be sure that this backup can be read and written to successfully, select to verify the backup. In this example, the additional amount of time required for this is not a concern because the file size will be consistent and this is all happening on a fast disk.

Figure 7-10. The general tab for creating a backup

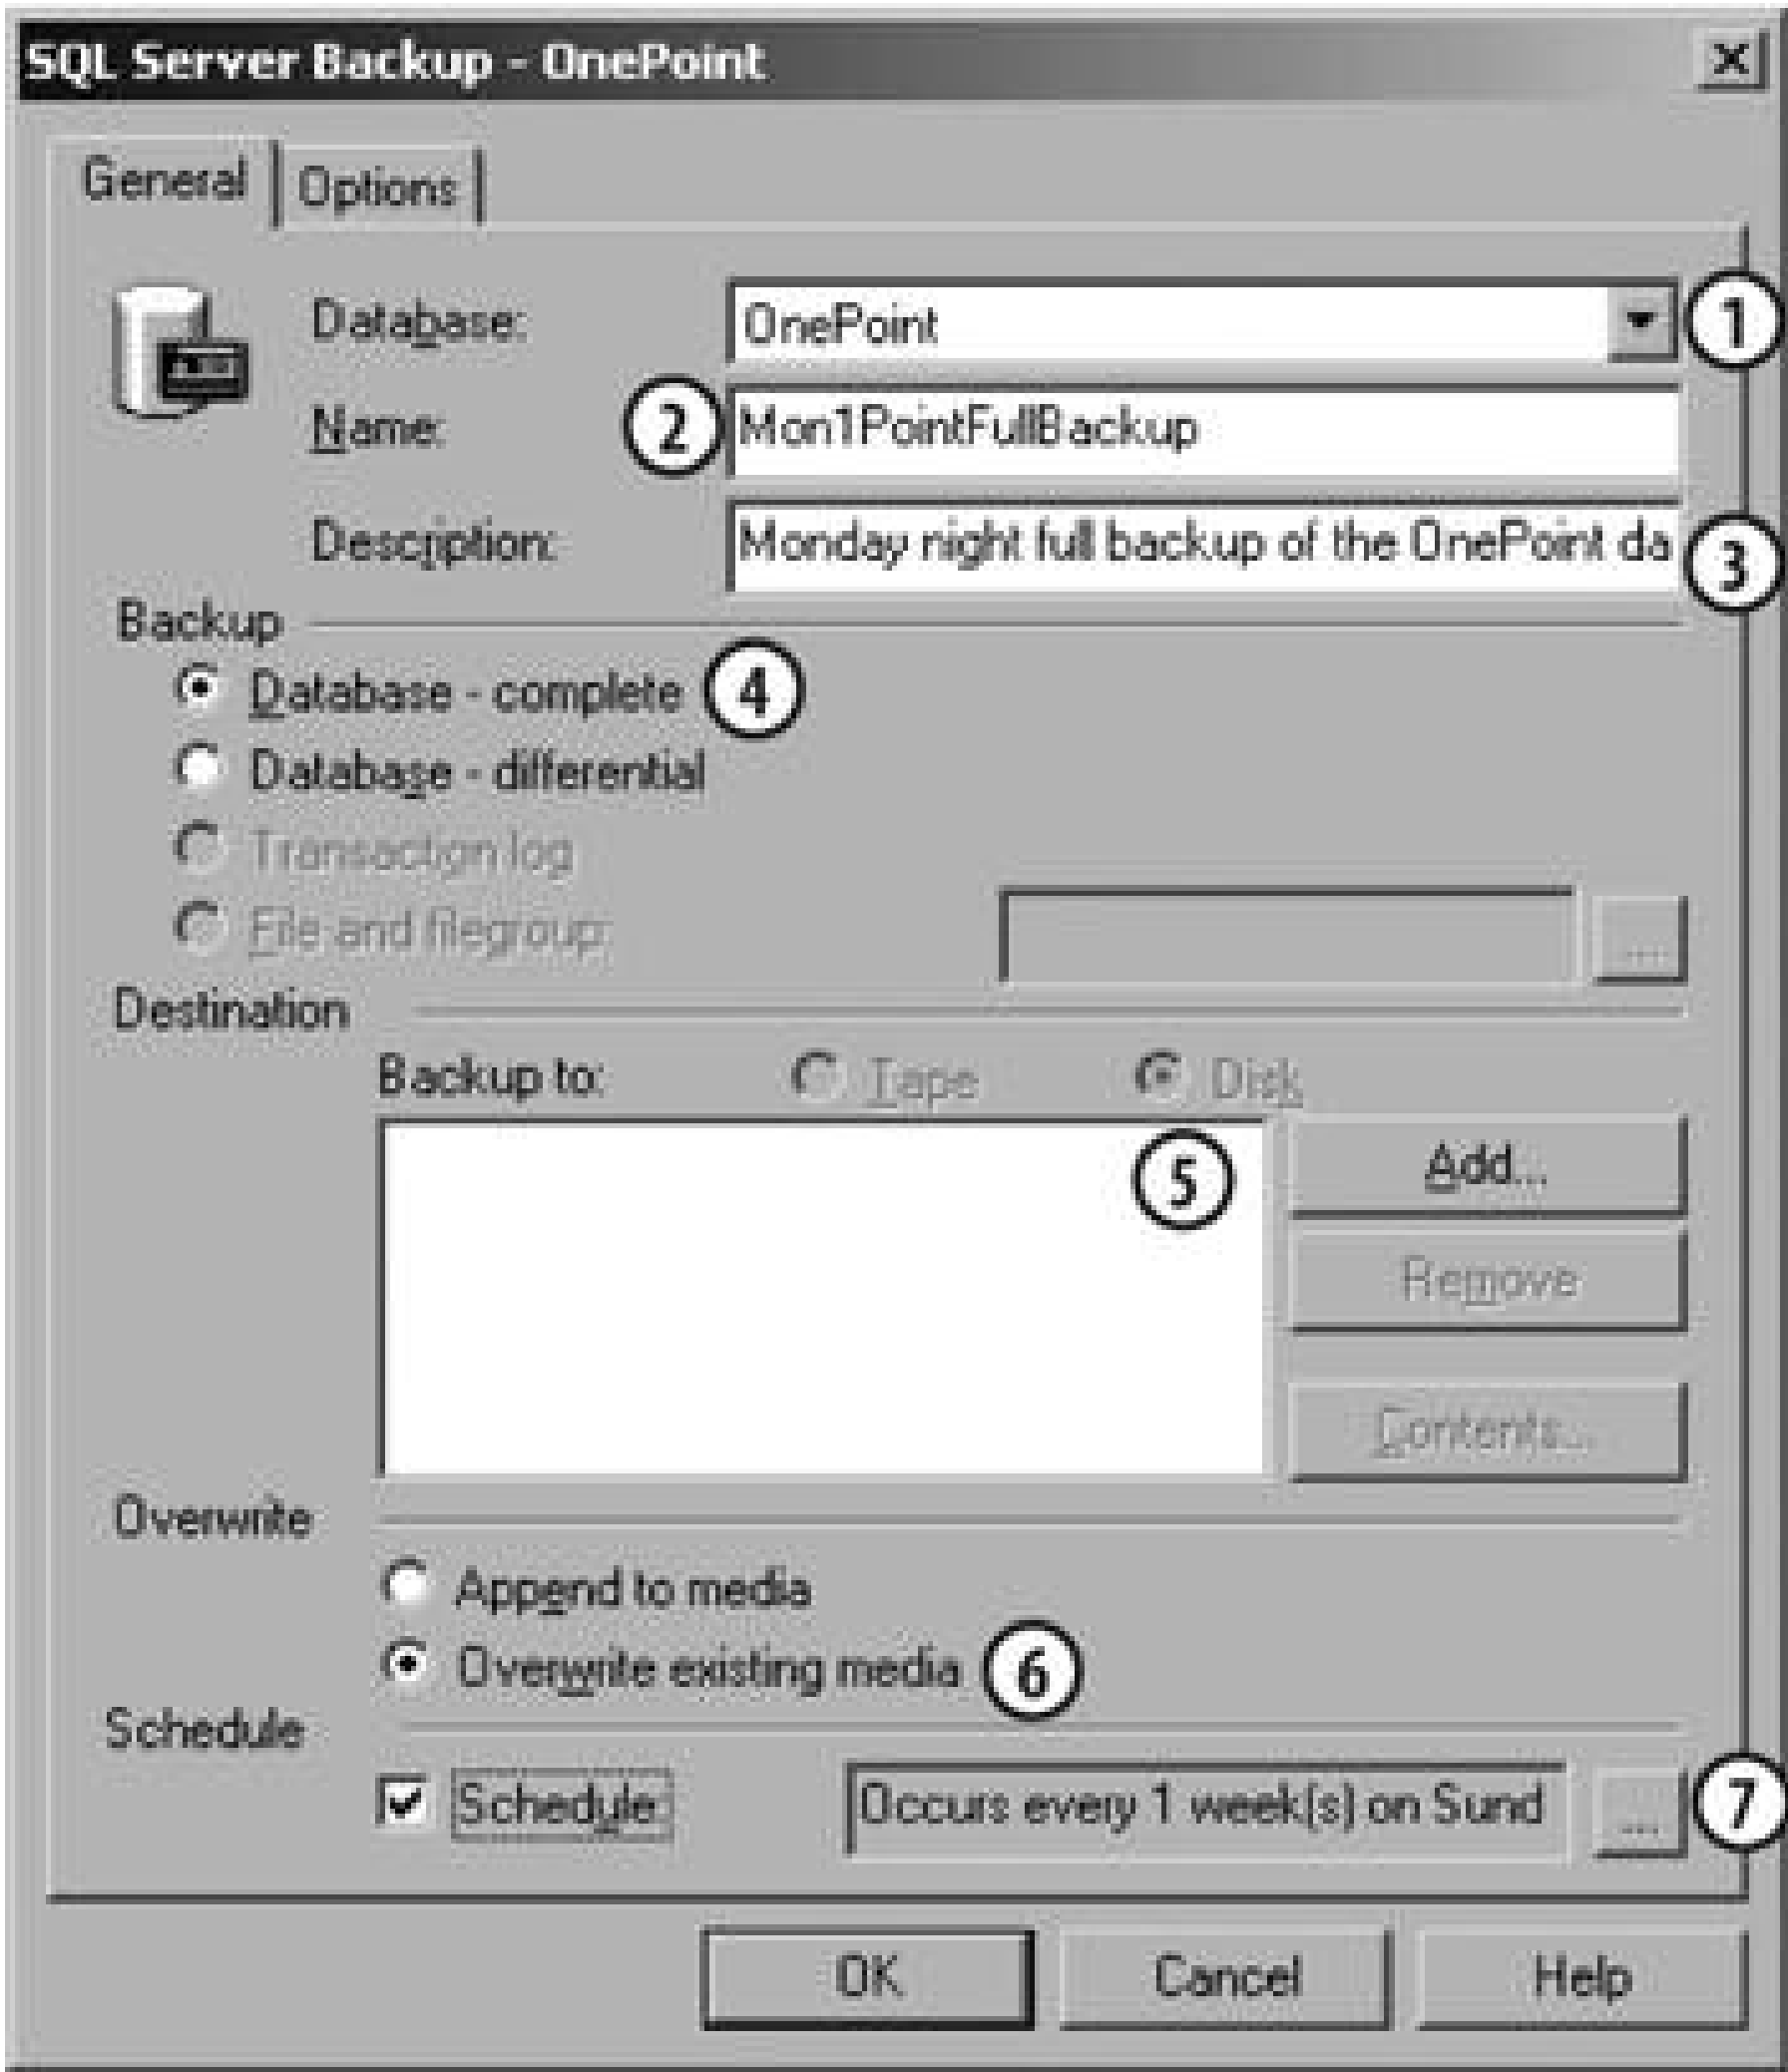
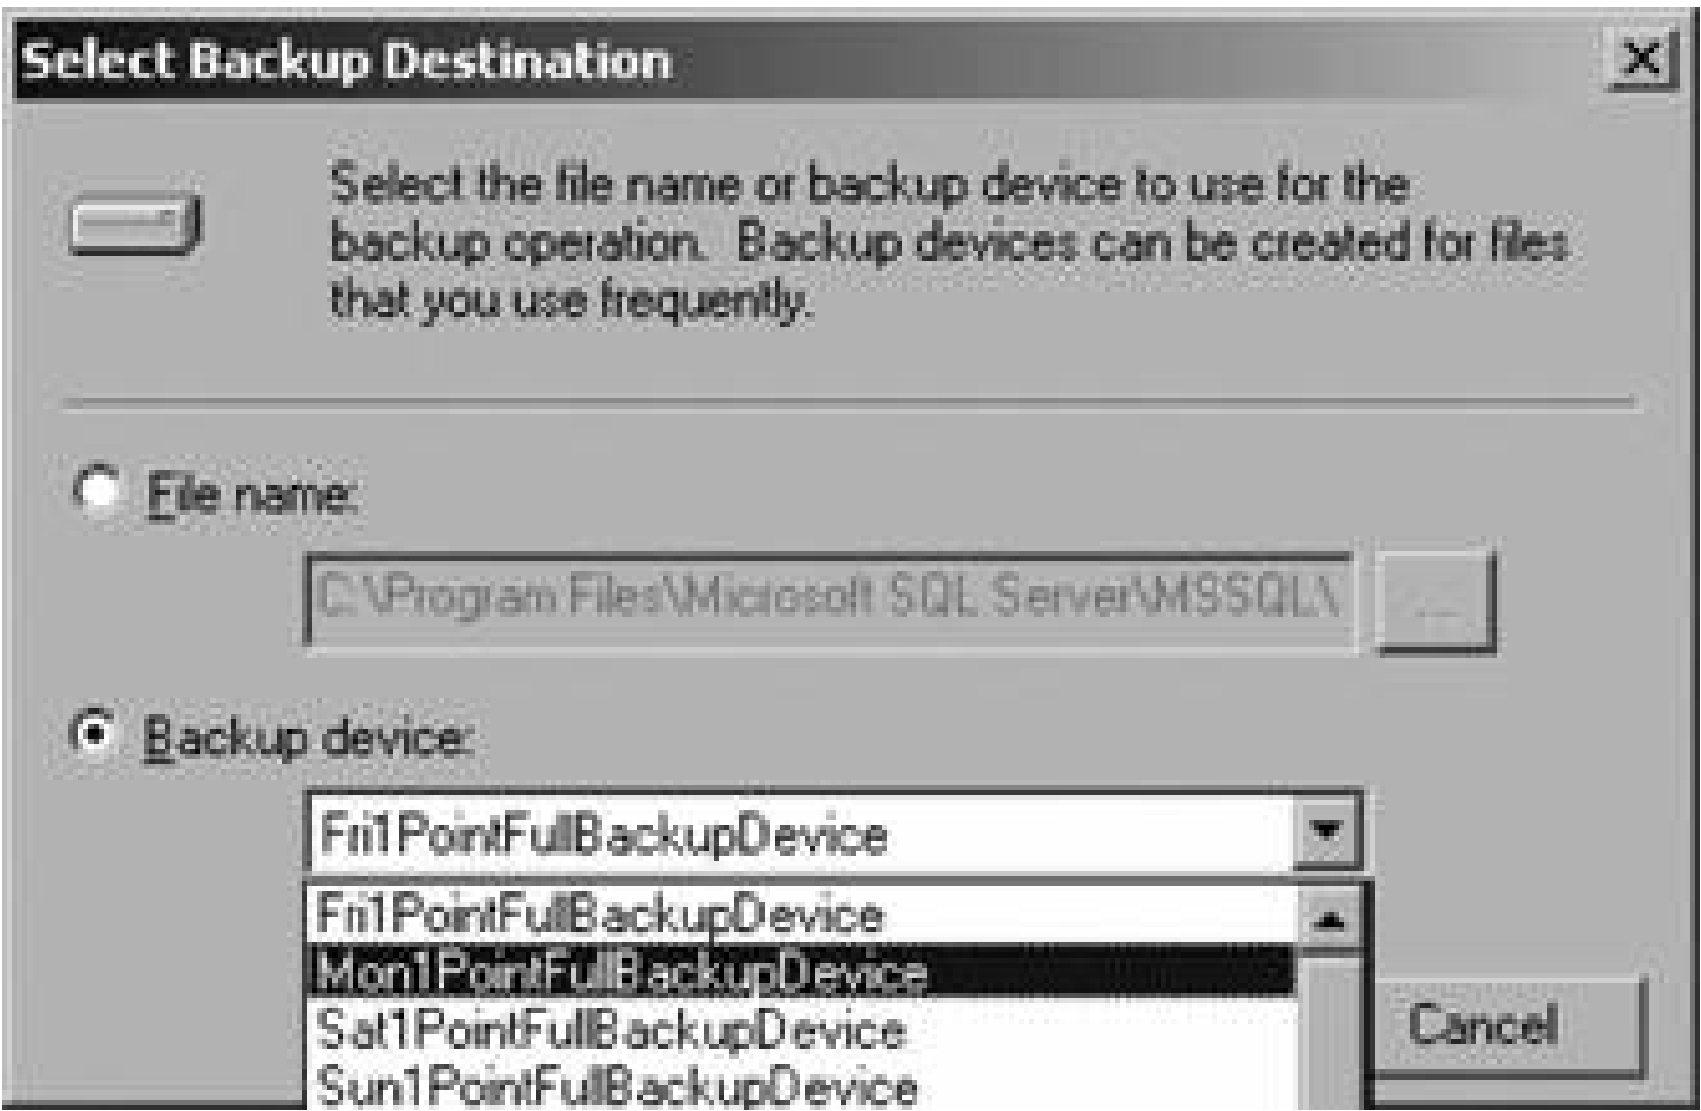


Figure 7-11. Selecting the backup device for the Monday night OnePoint full backup

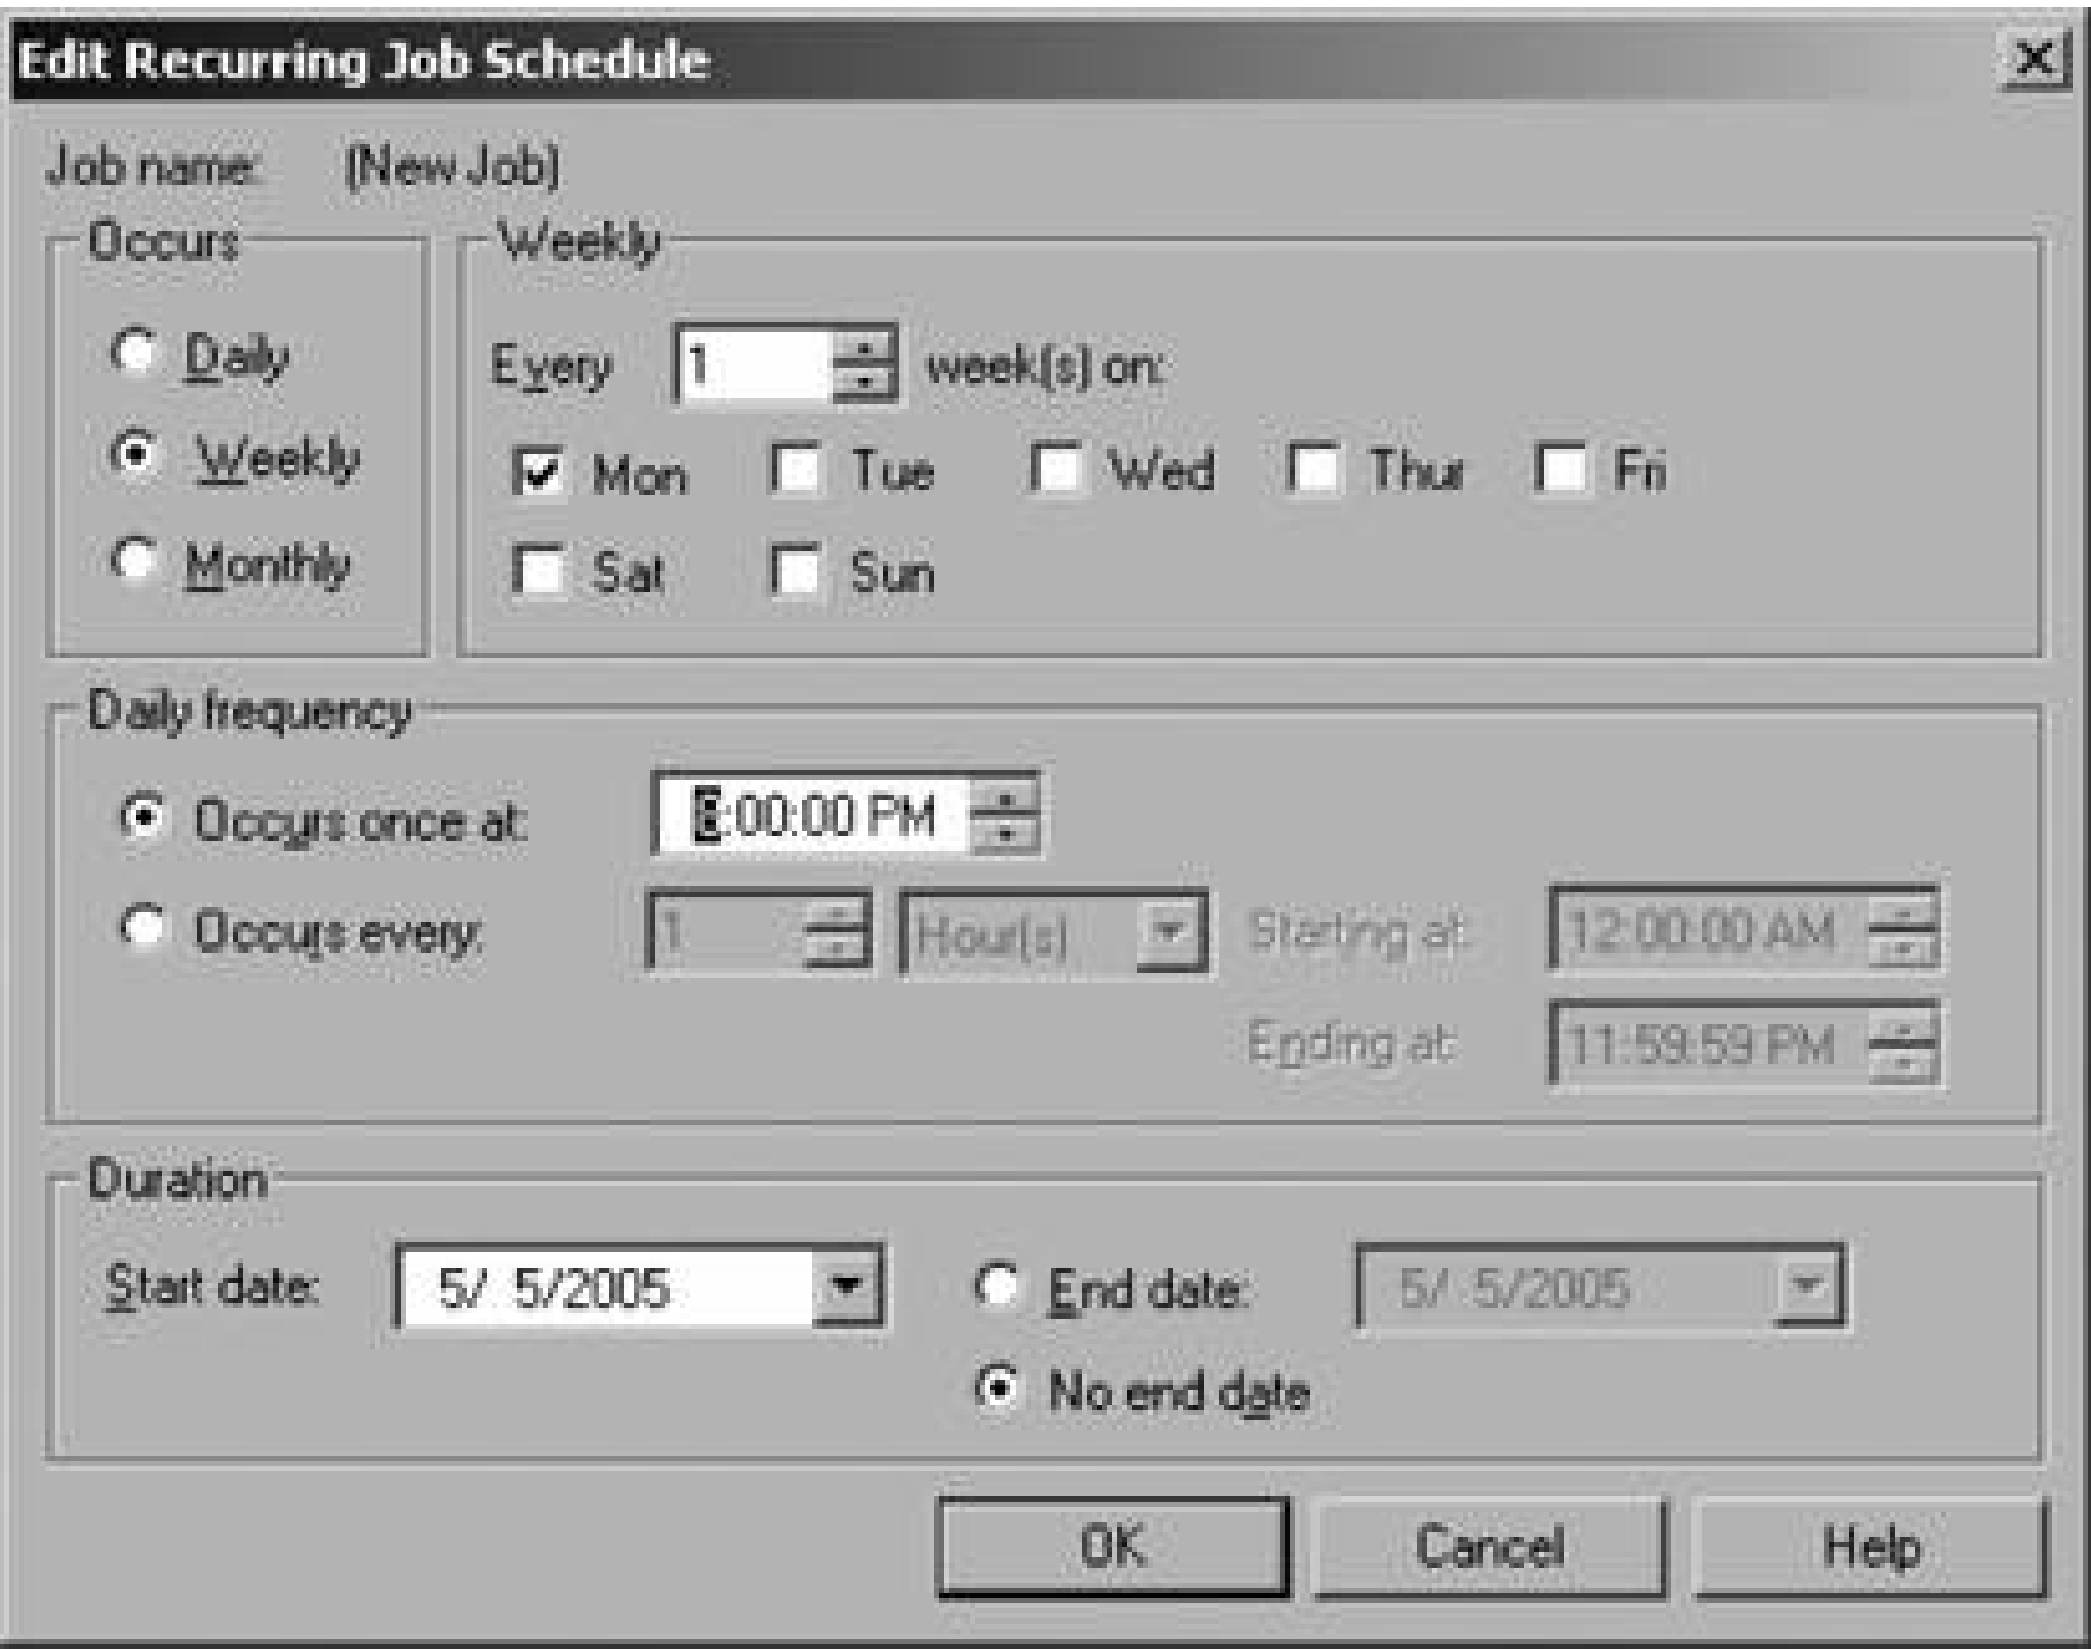


To help keep the OnePoint transaction logs clean of transactions that have been successfully committed to the database, remove them from the transaction log.

Figure 7-12. Name the schedule on this page, select recurring, and click next

Figure 7-13. Define the recurring schedule here



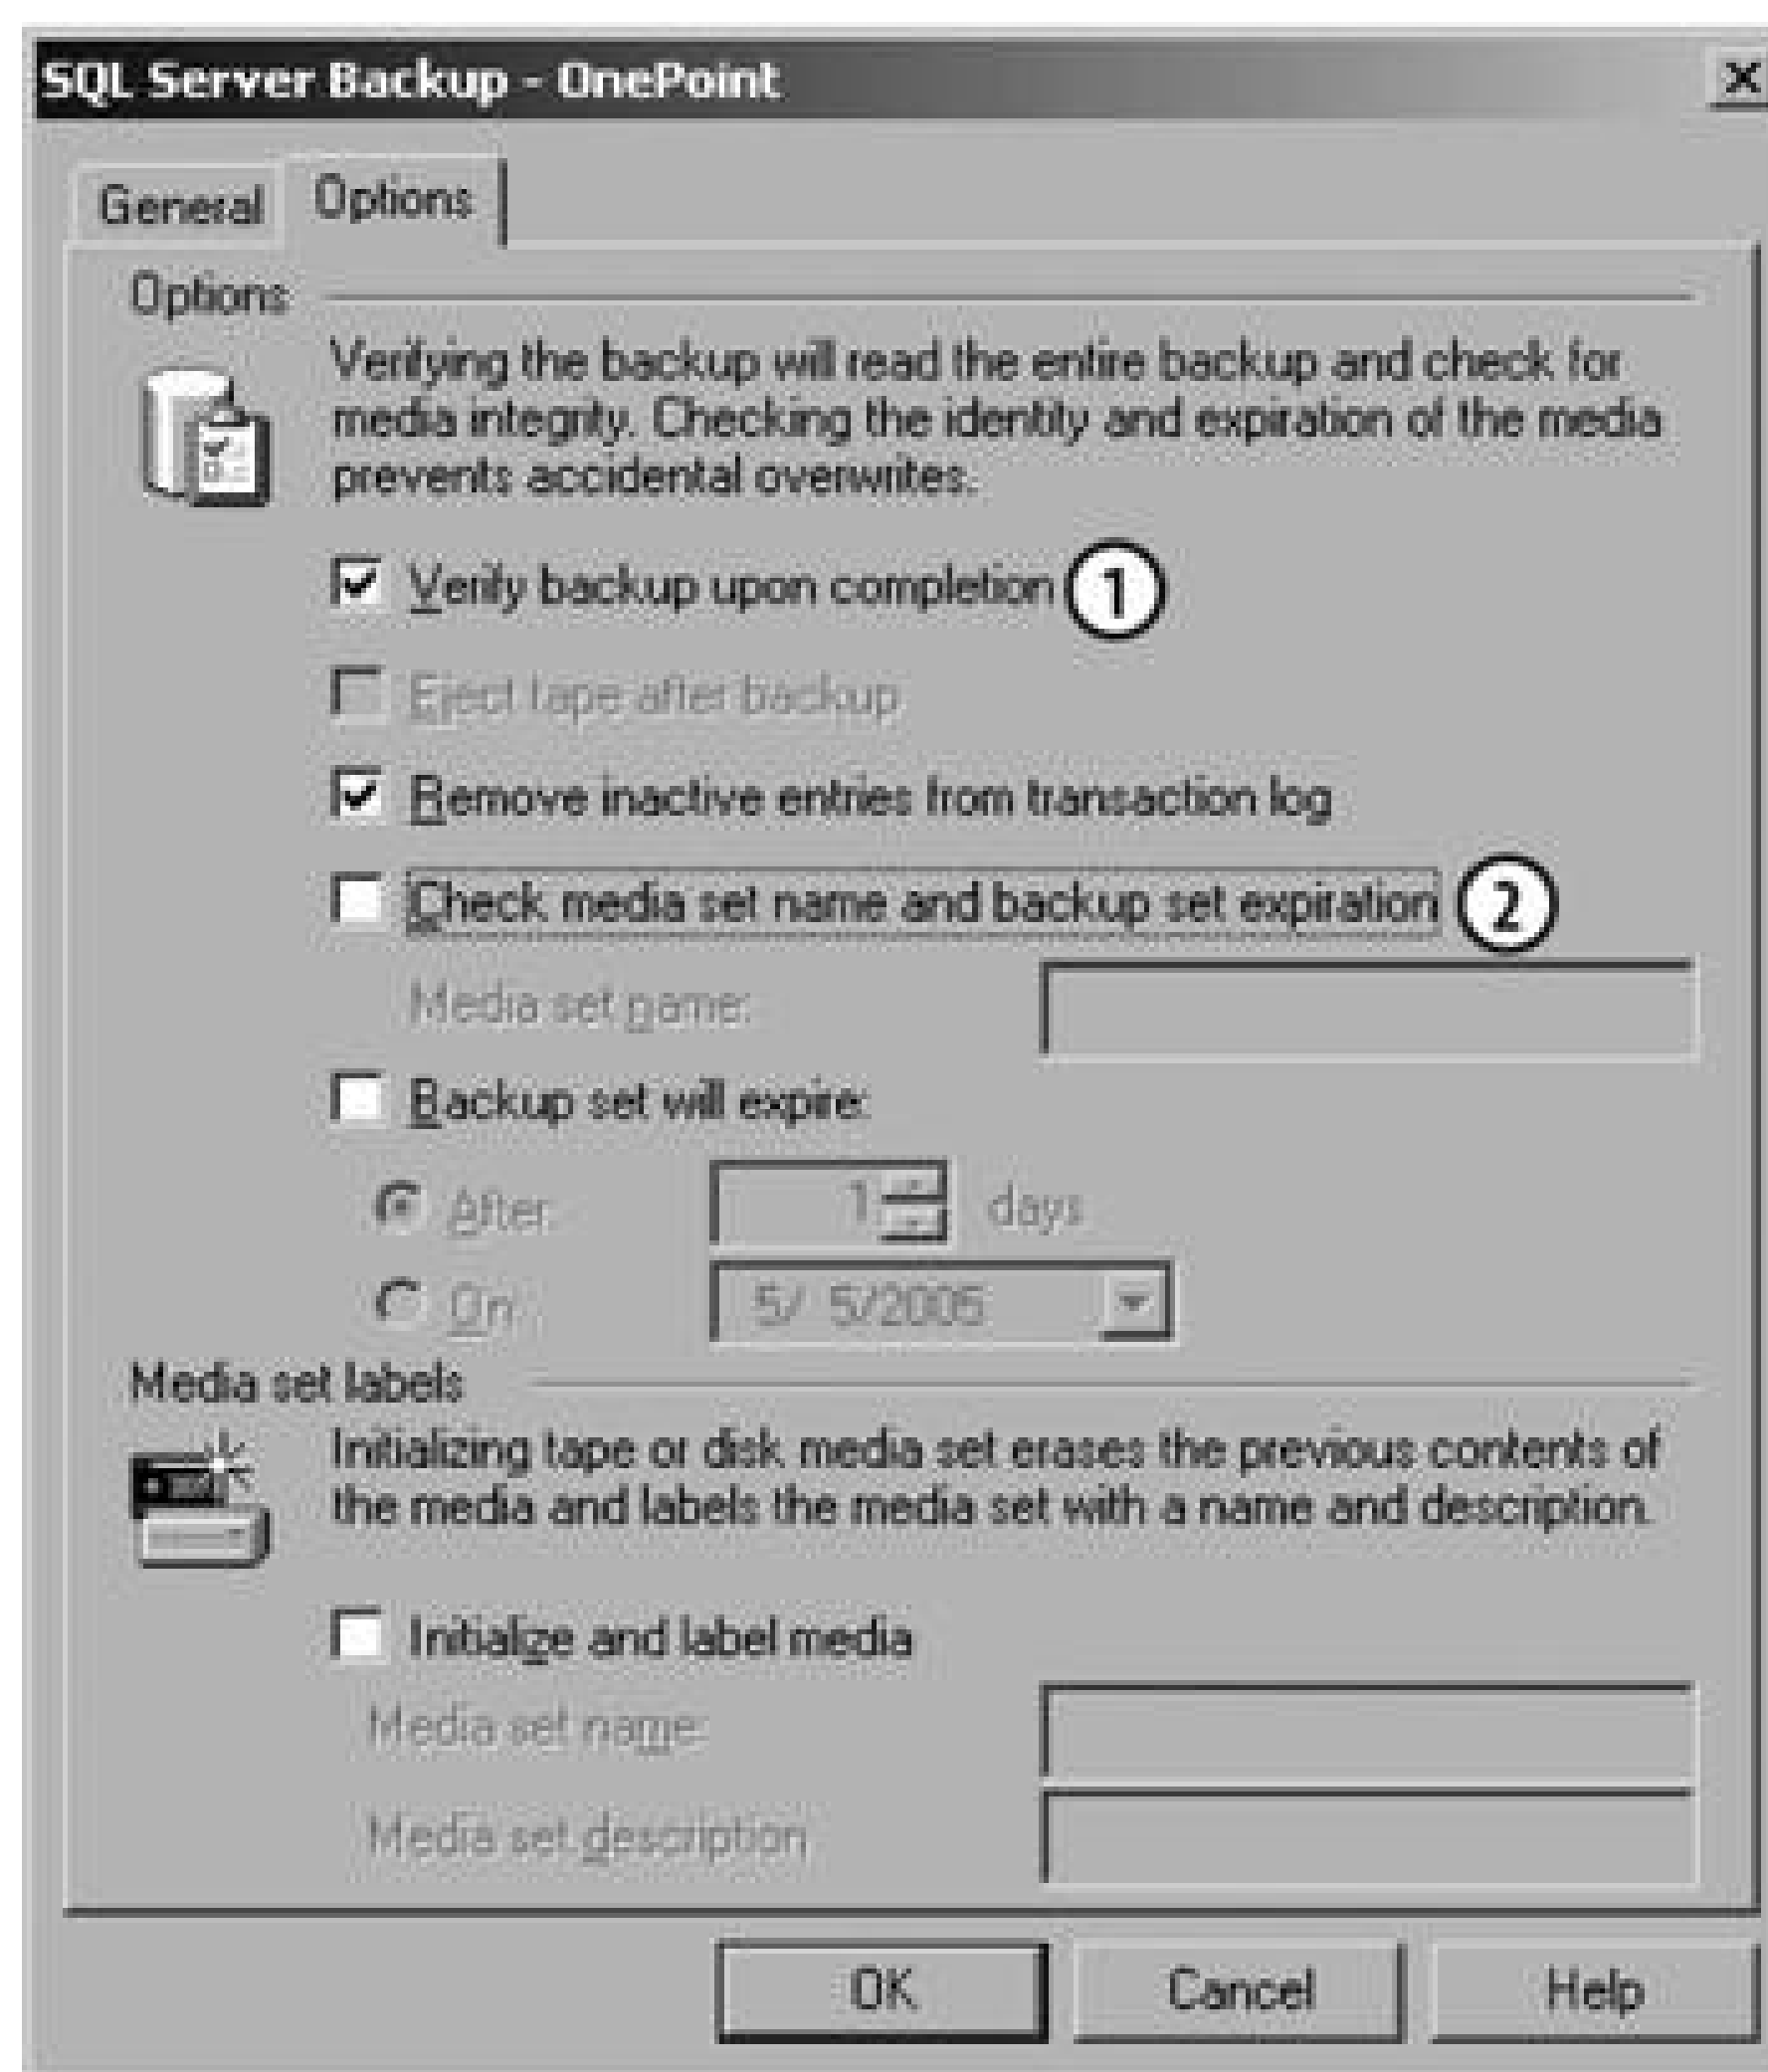


The other settings that were not selected are not that relevant to this backup scenario; they are more useful if backing up to a tape device. But briefly here is what they do:

*Check media set, name, and backup set for expiration*

If you are backing up to tape, it is likely that multiple tapes will be required to complete a single backup. This set of tapes is called a *media set*. To protect against overwriting a mistakenly inserted tape that still has a valid backup on it, SQL allows you to set an expiration date on the set. This way, if the expiration date for the media set has not passed, the backup job will not overwrite it or initialize it.

Figure 7-14. Define the backup options here



### *Backup set will expire*

This is where you set the expiration date for the media set that is being created by this backup

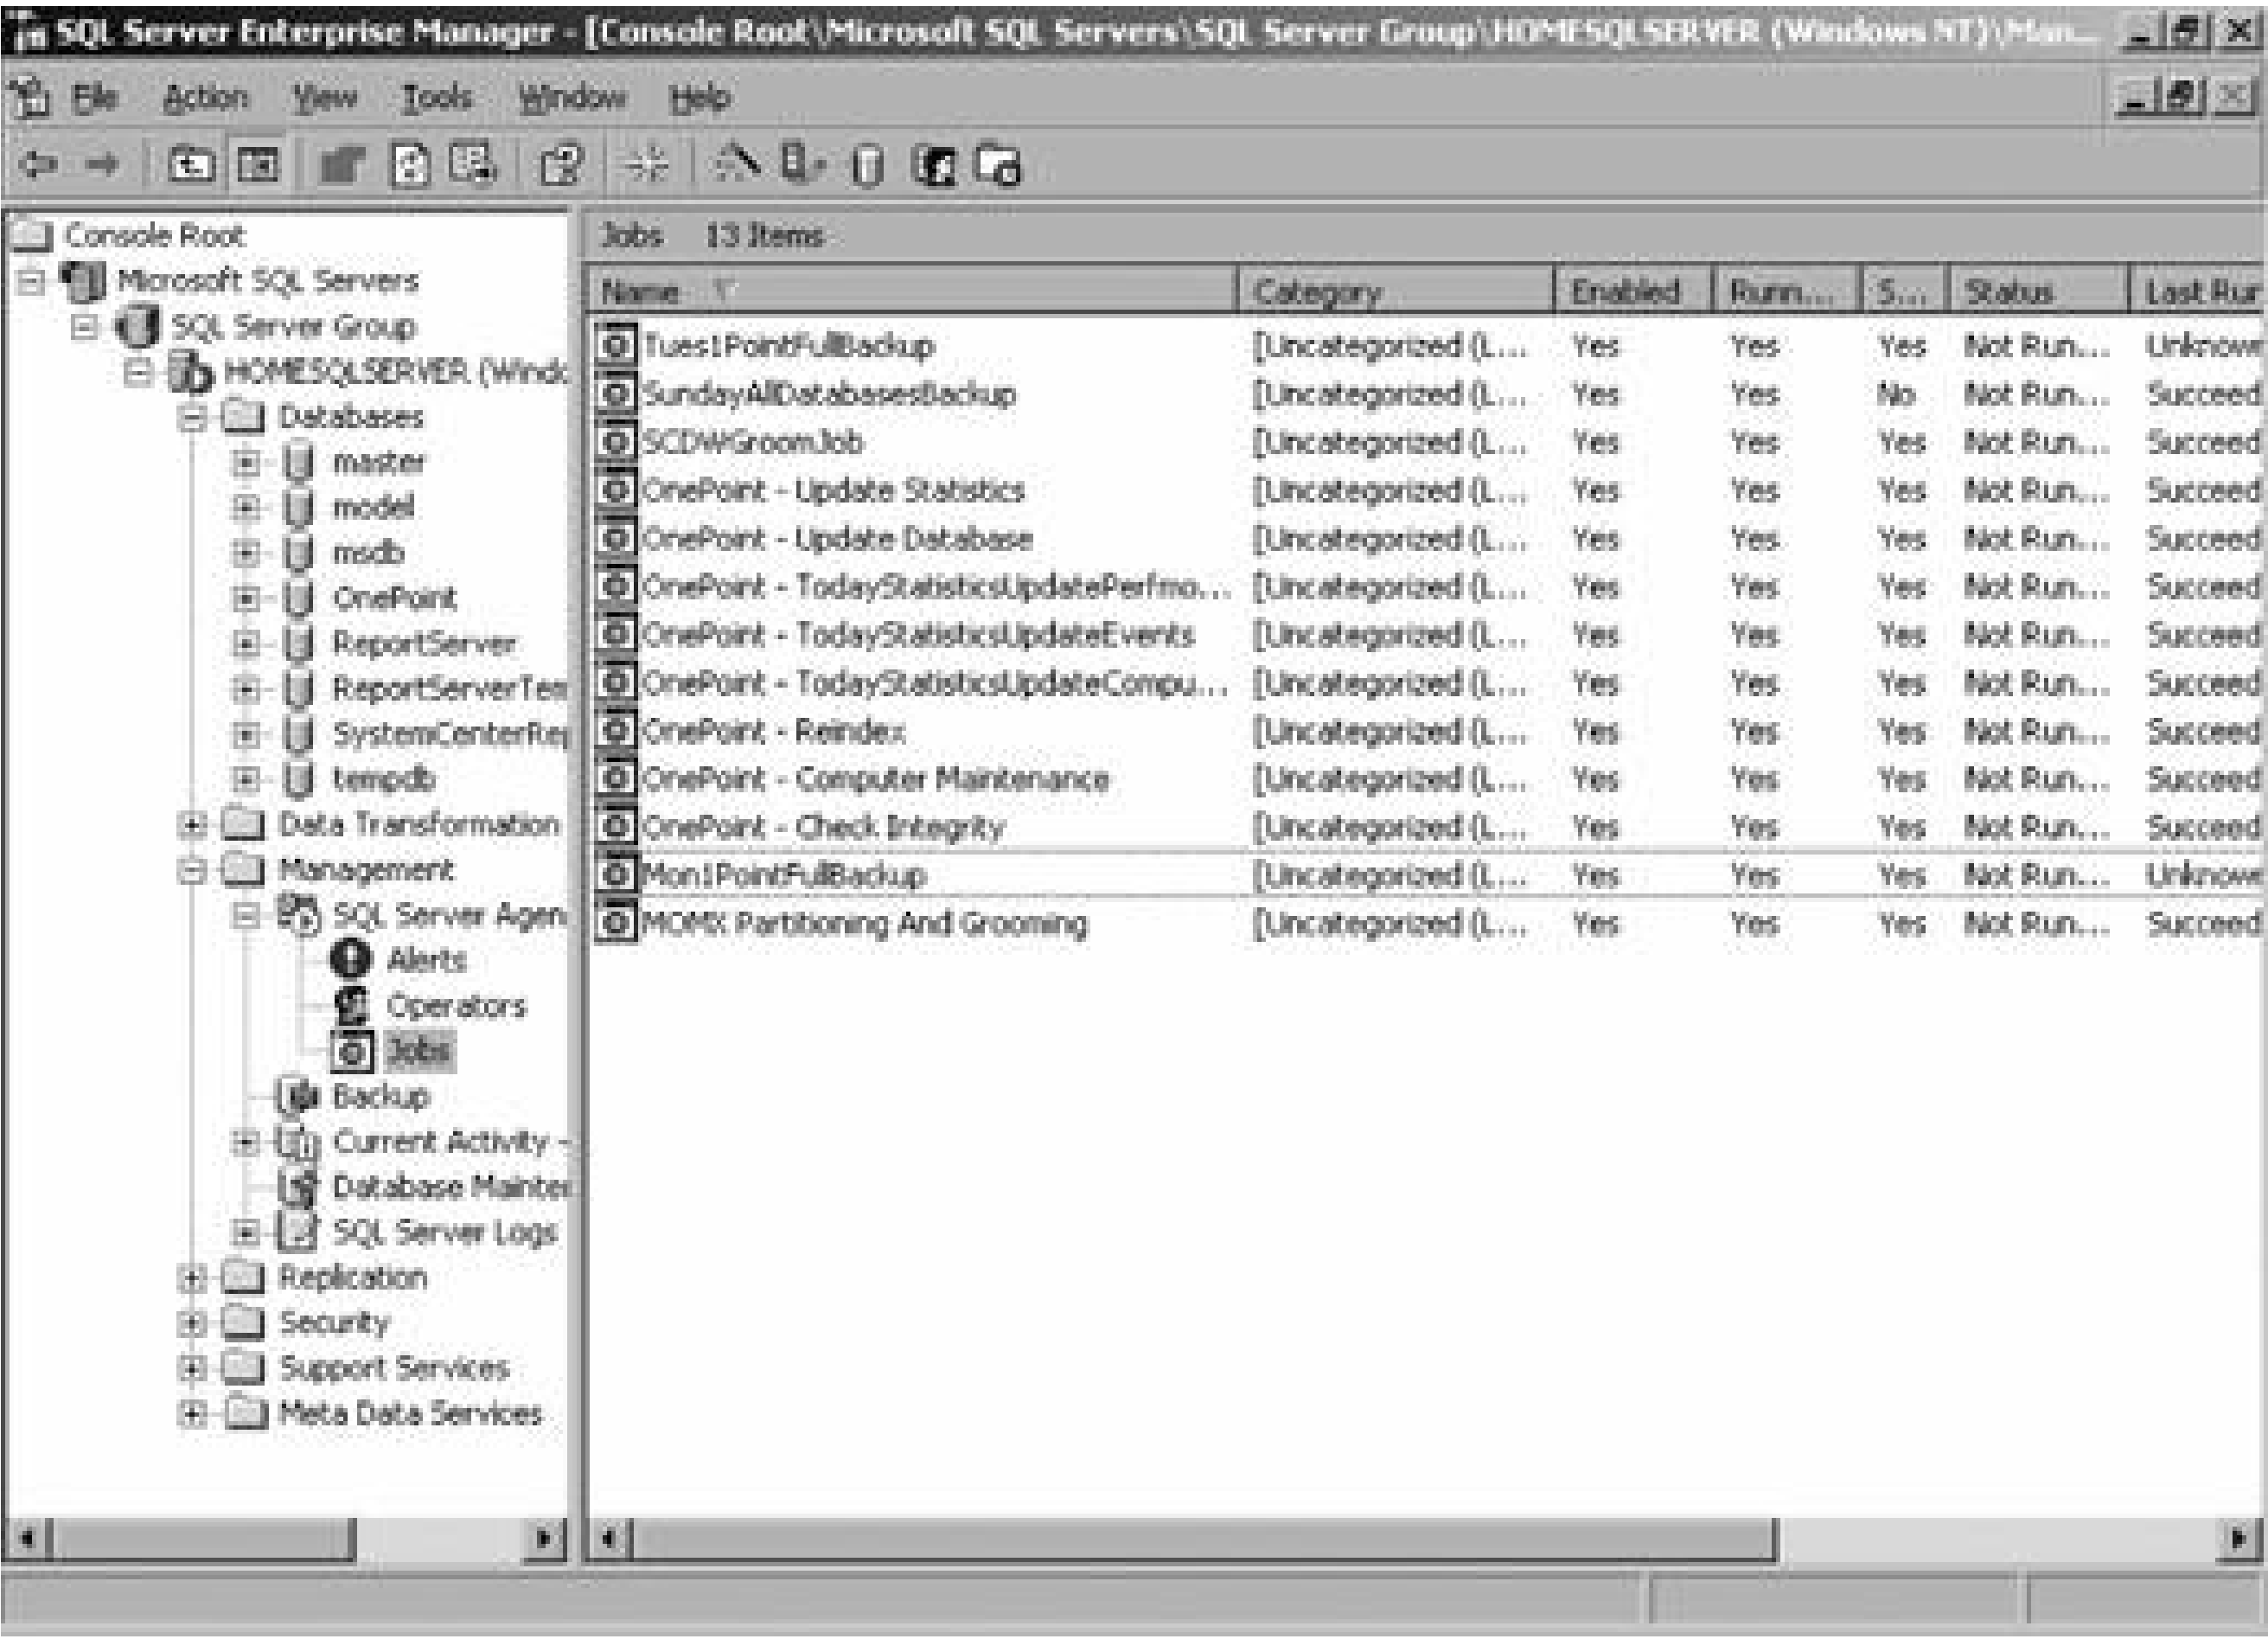
### *Initialize and label media*

Just as you need to format a floppy or a new partition prior to using it, tape media must be formatted. Here you can select to perform this format and to write a header on the first tape of the media set that includes its name and description.

When the configuration process is completed, SQL will create a job and place it in the Jobs container along with all the other SQL jobs (see [Figure 7-15](#)).

The next step is to create the same jobs for the rest of the week for the OnePoint backups. To speed this process up and to avoid manually re-creating these backup definitions should they be deleted or lost, SQL will generate a script that will create the same backup job when run from the SQL Query Analyzer against this server. Once the script generates, edit it with Notepad (you can also use the SQL Query Analyzer tool if that is your preference). Perform a find on Mon1PointFullBackup and replace it with Tues1PointFullBackup. Replace the Monday in the description with Tuesday and so on for the rest of the days in the week.

Figure 7-15. The completed Mon1PointFullBackup job



To do this, open the context menu for the Mon1PointFullBackup object in the Jobs container and select "All tasksGenerate SQL script." This brings up the Generate SQL Script page (see[Figure 7-16](#)), where all you need to enter is a path and filename to save the script as.

Figure 7-16. Saving the backup script



Click the ellipsis and accept the default path. The file is named Mon1PointFullBackup and the *.sql* extension is automatically appended.

Next, perform the editing of the text file and save one for each day of the week. Open and run the Tues, Weds, Thurs, Fri, Sat, and Sun versions of the files in the SQL Query Analyzer and the other jobs are created automatically. The last step is to open the job object and set the scheduled day of the week to the desired setting. This feature is a real time saver.

To round out the backup job creation, create the backup for the master and msdb databases by using the process shown in [Figures 7-10](#) through [7-14](#).

These jobs will run on schedule and output the *<jobname>.bak* file on the local disk. If you look at these, you will see that they are significantly smaller than the total database size, which is a good thing because the next step is to create an OS-level backup job that grabs these files and backs them up across the network to a share.

### 7.1.5.5. Creating OS-level backups

This portion of the backup scheme uses the Windows Backup utility to back up the *.bak* files (SQL backup devices) from the SQL server to a network share. The use of this interface is already well documented so this next example will just go through the creation of one of these tasks. To start, on the MOM database server open the Windows Backup utility in the Backup wizard mode.

1. On the What to Backup page, select "Backup selected files, drives, or network data."
2. On the Items to Backup page, browse to the directory that contains the SQL backup files. By default, this is *C:\Program Files\Microsoft SQL Server\MSSQL\Backup*. Select the file or files that match the day that you are making the backup for. For example, if you are creating the Windows backup scheduled task for a Monday, then select the *Mon1PointFullBackup.bak* file.
3. On the Backup Type, Destination, and Name page, enter the UNC path to the network share for the destination directory and then enter a name for the backup. For example, [\\homesrv02\SQL backups\Mon1PointFullBackup](#).
4. On the Completing the Backup Wizard page, select the Advanced button.
5. On the Type of Backup page, select Normal backup from the drop-down list (this is the same as a full backup).
6. On the How to Backup page, select the "Verify data after backup" checkbox.
7. On the Backup Options page, select the "Replace the existing backups" option. This will overwrite backup files of the same name that already exist in the destination directory.
8. On the When to Backup page, select the Later radio button, give the backup job a name, and then click Set Schedule.

9. On the Schedule Job page, select "weekly" from the schedule task drop-down list, the day of the week, and the time. Be sure to allow sufficient time for the SQL backups to complete. Also, since this is a file-level backup and not a database backup, you don't need to be concerned with impacting the OnePoint grooming operations.
10. At the Set Account Information prompt, enter the credentials under which to run this scheduled task.

Completing this process will create a job in the scheduled tasks. When this runs (and you can kick it off manually by right-clicking the task and selecting Run), the backup cycle for the databases (and any other files that have been included) will be completed.

## 7.1.6. Restoring SQL Databases

Restoring a database in SQL, when SQL is healthy, is about as painless a process as you can imagine. By using the method outlined in the last section, SQL has recorded every backup action taken. This includes the name and location of the backup file/device, the type of backup, when it was done, and under what credentials it executed. SQL tracks this information for each database in the msdb system database. Thus, when you start the restore process for any given database, SQL will tell you which backup set or sets to use, where those backups are, and in what order to apply them.

Being able to restore the OnePoint database this easily is only possible if both the msdb and master system databases are up and healthy. As a MOM administrator, you must be prepared to recover from a OnePoint-only failure and a failure that impacts the SQL system databases.

### 7.1.6.1. Restoring OnePoint

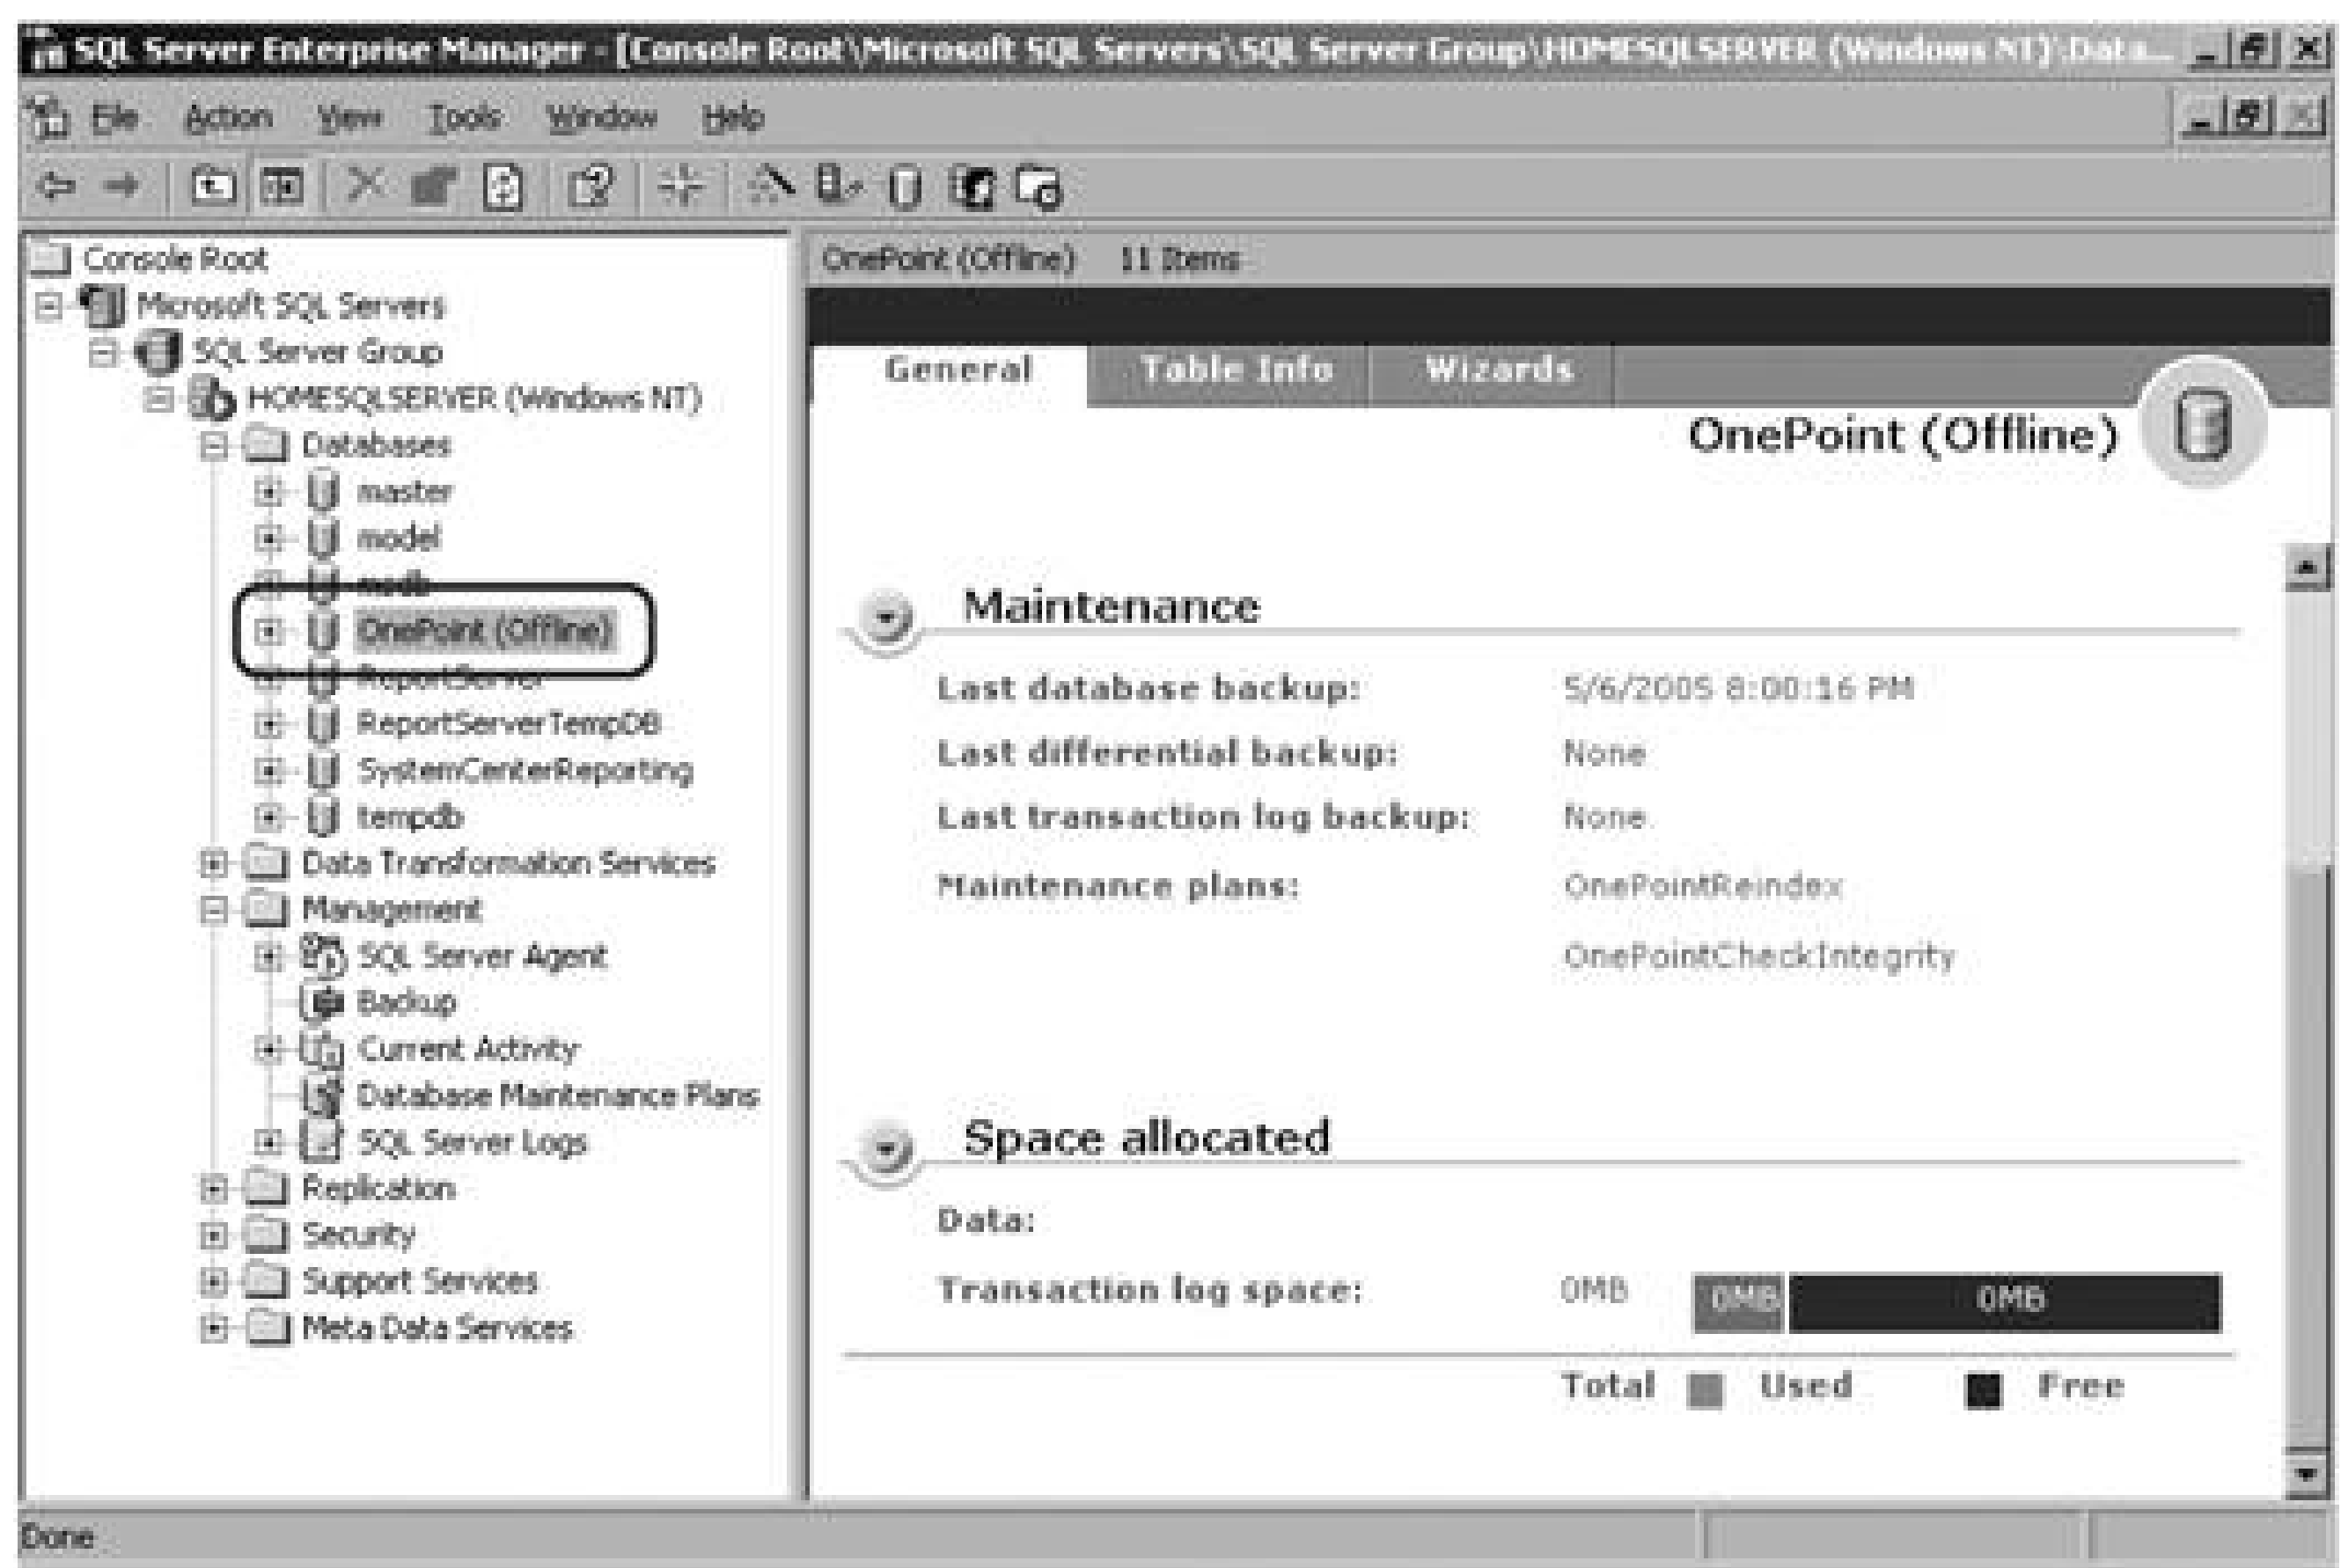
The first step in the process of restoring the OnePoint database and its transaction log is to stop all communications with the OnePoint database while the restore is in progress. You can do this by stopping the MOM service on all management servers and on the database server. You should also make sure that the reporting DTS scheduled task is disabled if the restore is performed around the time that the operations-to-warehouse data copy is supposed to take place.

Keep in mind that none of the agents are stopped on the managed computers. They will continue to process rules and responses as well as cache operations data locally up to a configurable limit (3 MB by default) until the agents are able to communicate with their management servers again.

Another way to isolate the OnePoint database from the rest of the management groups is to take it offline. You do this by bringing up the context menu for the OnePoint database in SQL Enterprise Manager, selecting All Tasks and selecting "Take offline." The database icon will then be grayed out and be marked as offline (see [Figure 7-17](#)).

Figure 7-17. The OnePoint database taken offline





This method of database isolation allows the management servers to continue receiving operational data from their agents. In addition, if there are multiple management servers in the management group, agent failover is configured automatically. So, if both management servers become unavailable at the same time, the agents will continually try to failover to the alternate management server. This will generate a great deal of event log data on the agent-managed machines as the agents ping-pong between the management servers trying to find one that is available. By keeping both of them up, communications between the agents and management servers are never interrupted.

Once the OnePoint database is offline, bring up the context menu, and select All Tasks → Restore Database. SQL will then retrieve the backup information from the msdb database and the list of backups that represent the quickest path to a fully restored and functioning database will automatically be listed and selected for the restore operations (see [Figure 7-18](#)).

There are no options to be configured on the General tab, but on the Options tab select the "Force restore over existing database" option to ensure that the old, damaged database is overwritten (see [Figure 7-19](#)).

Figure 7-18. Restoring the OnePoint database



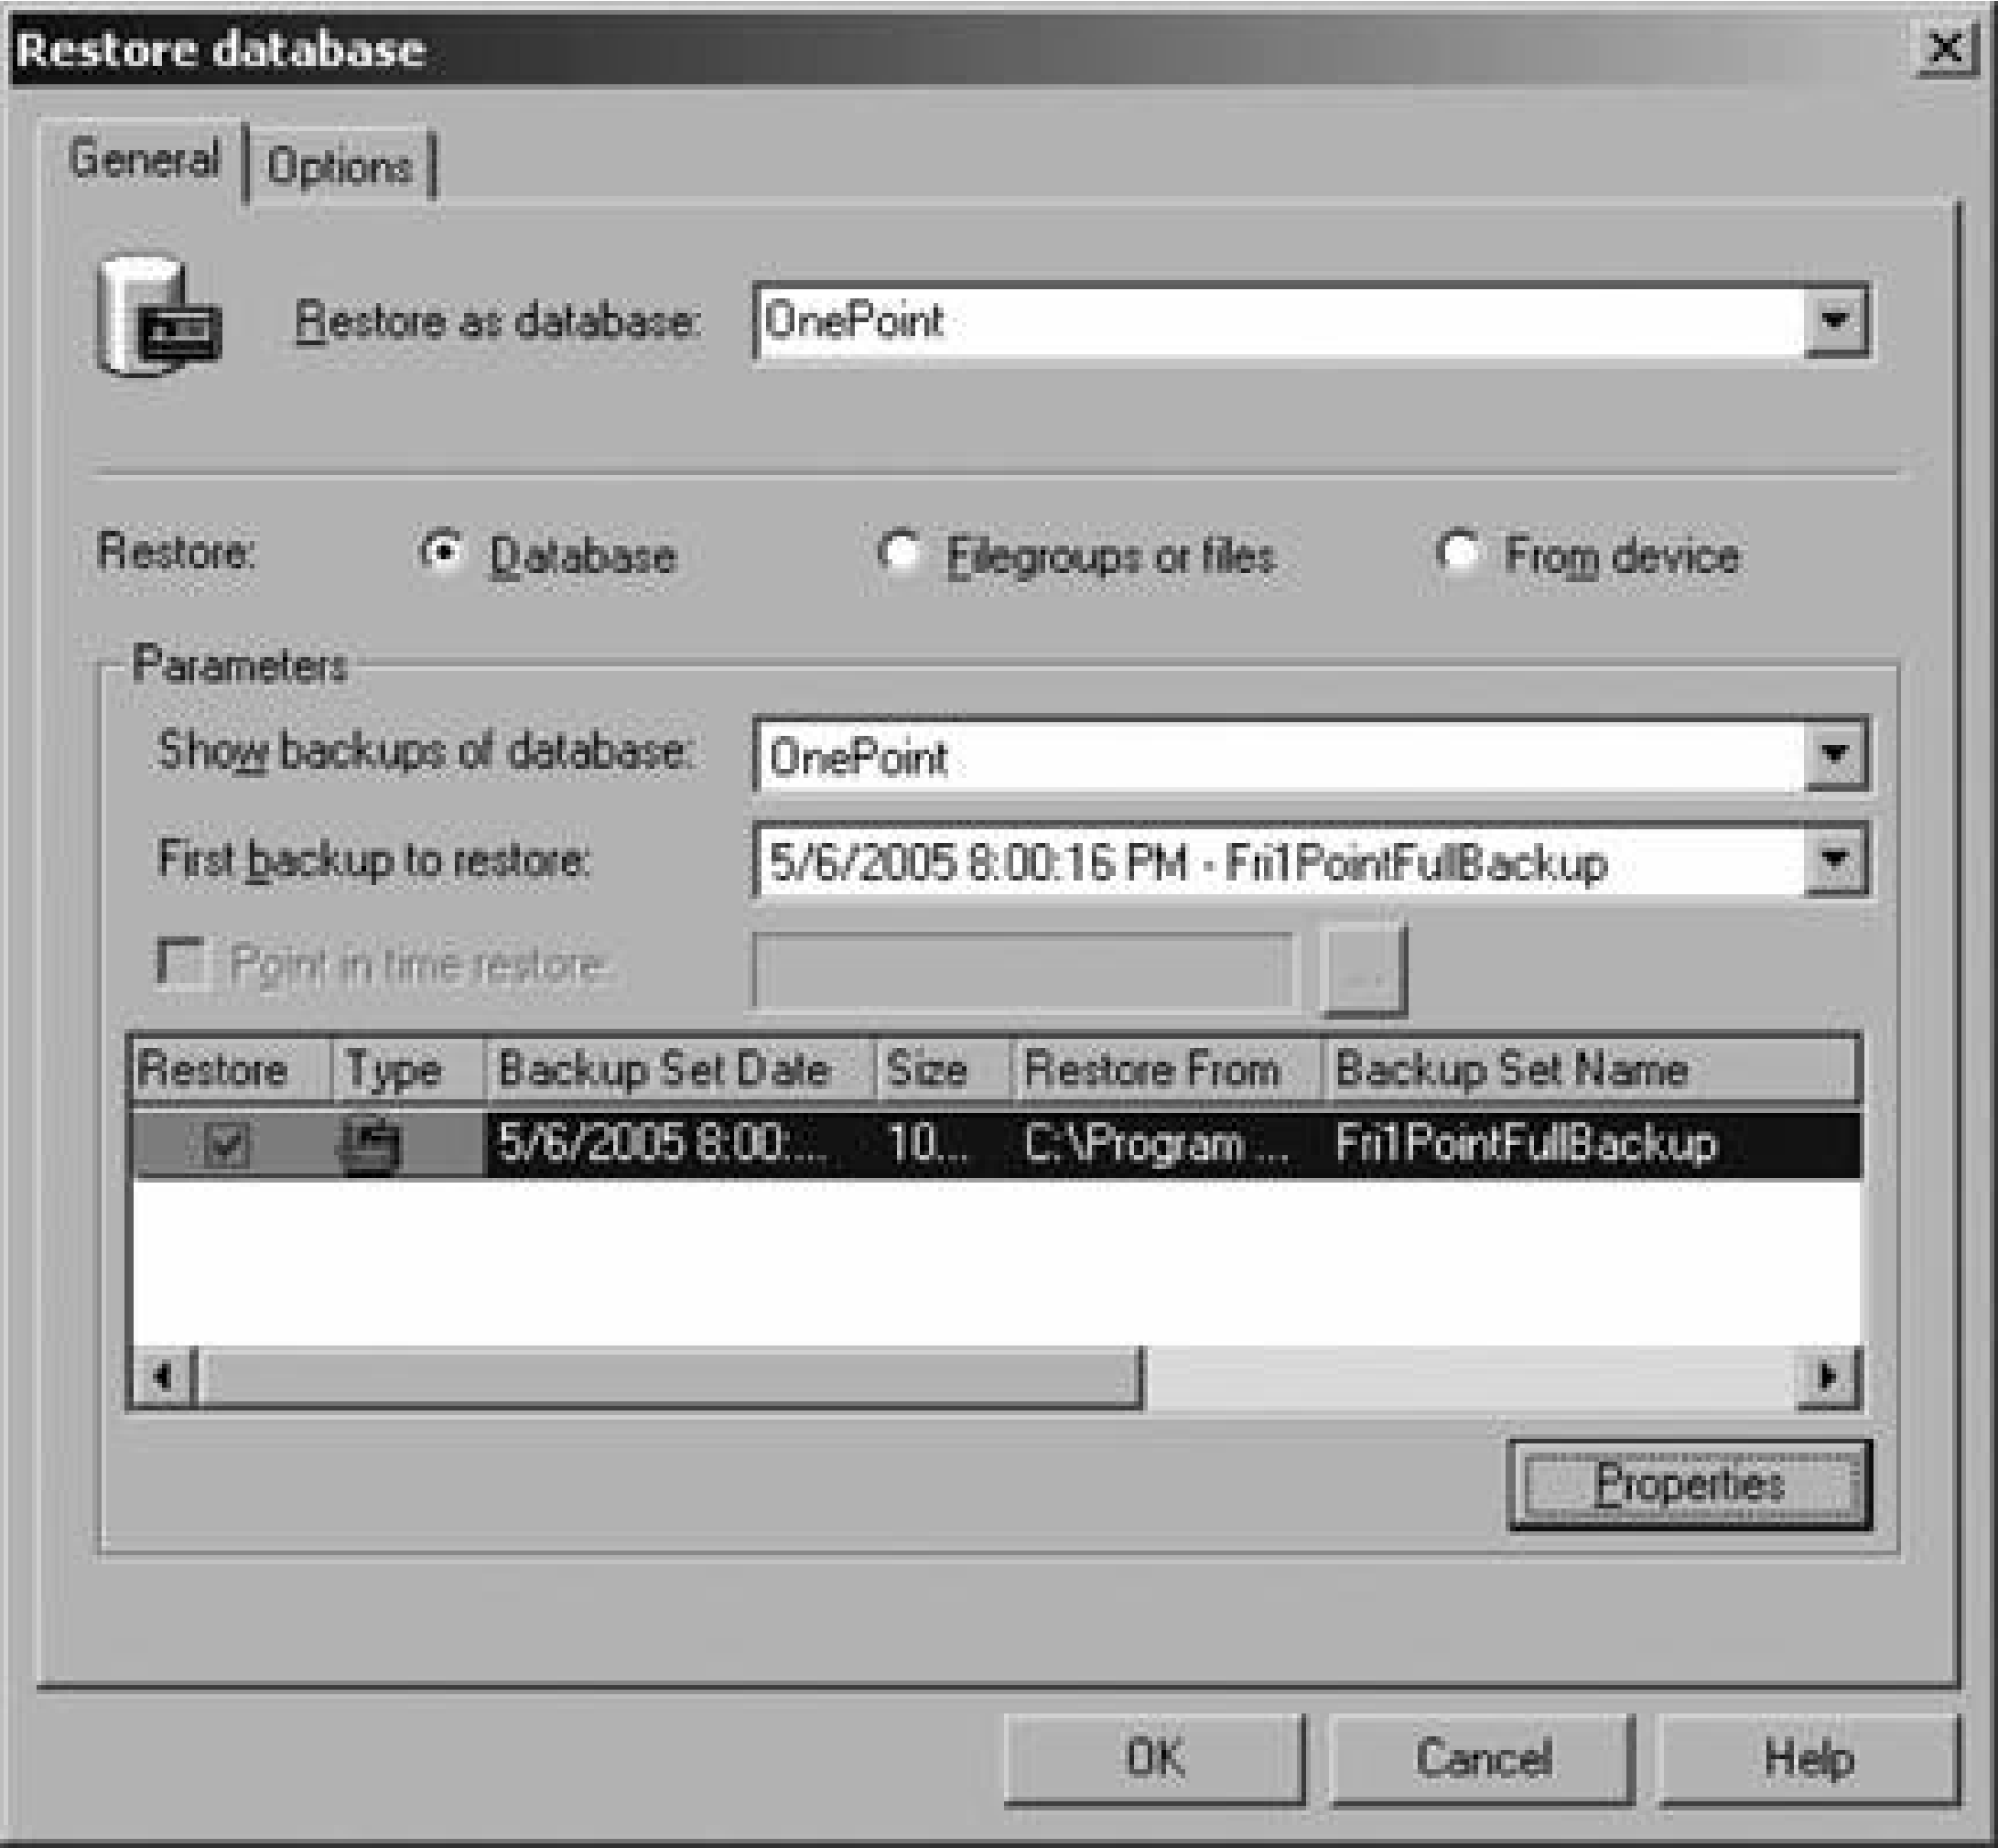
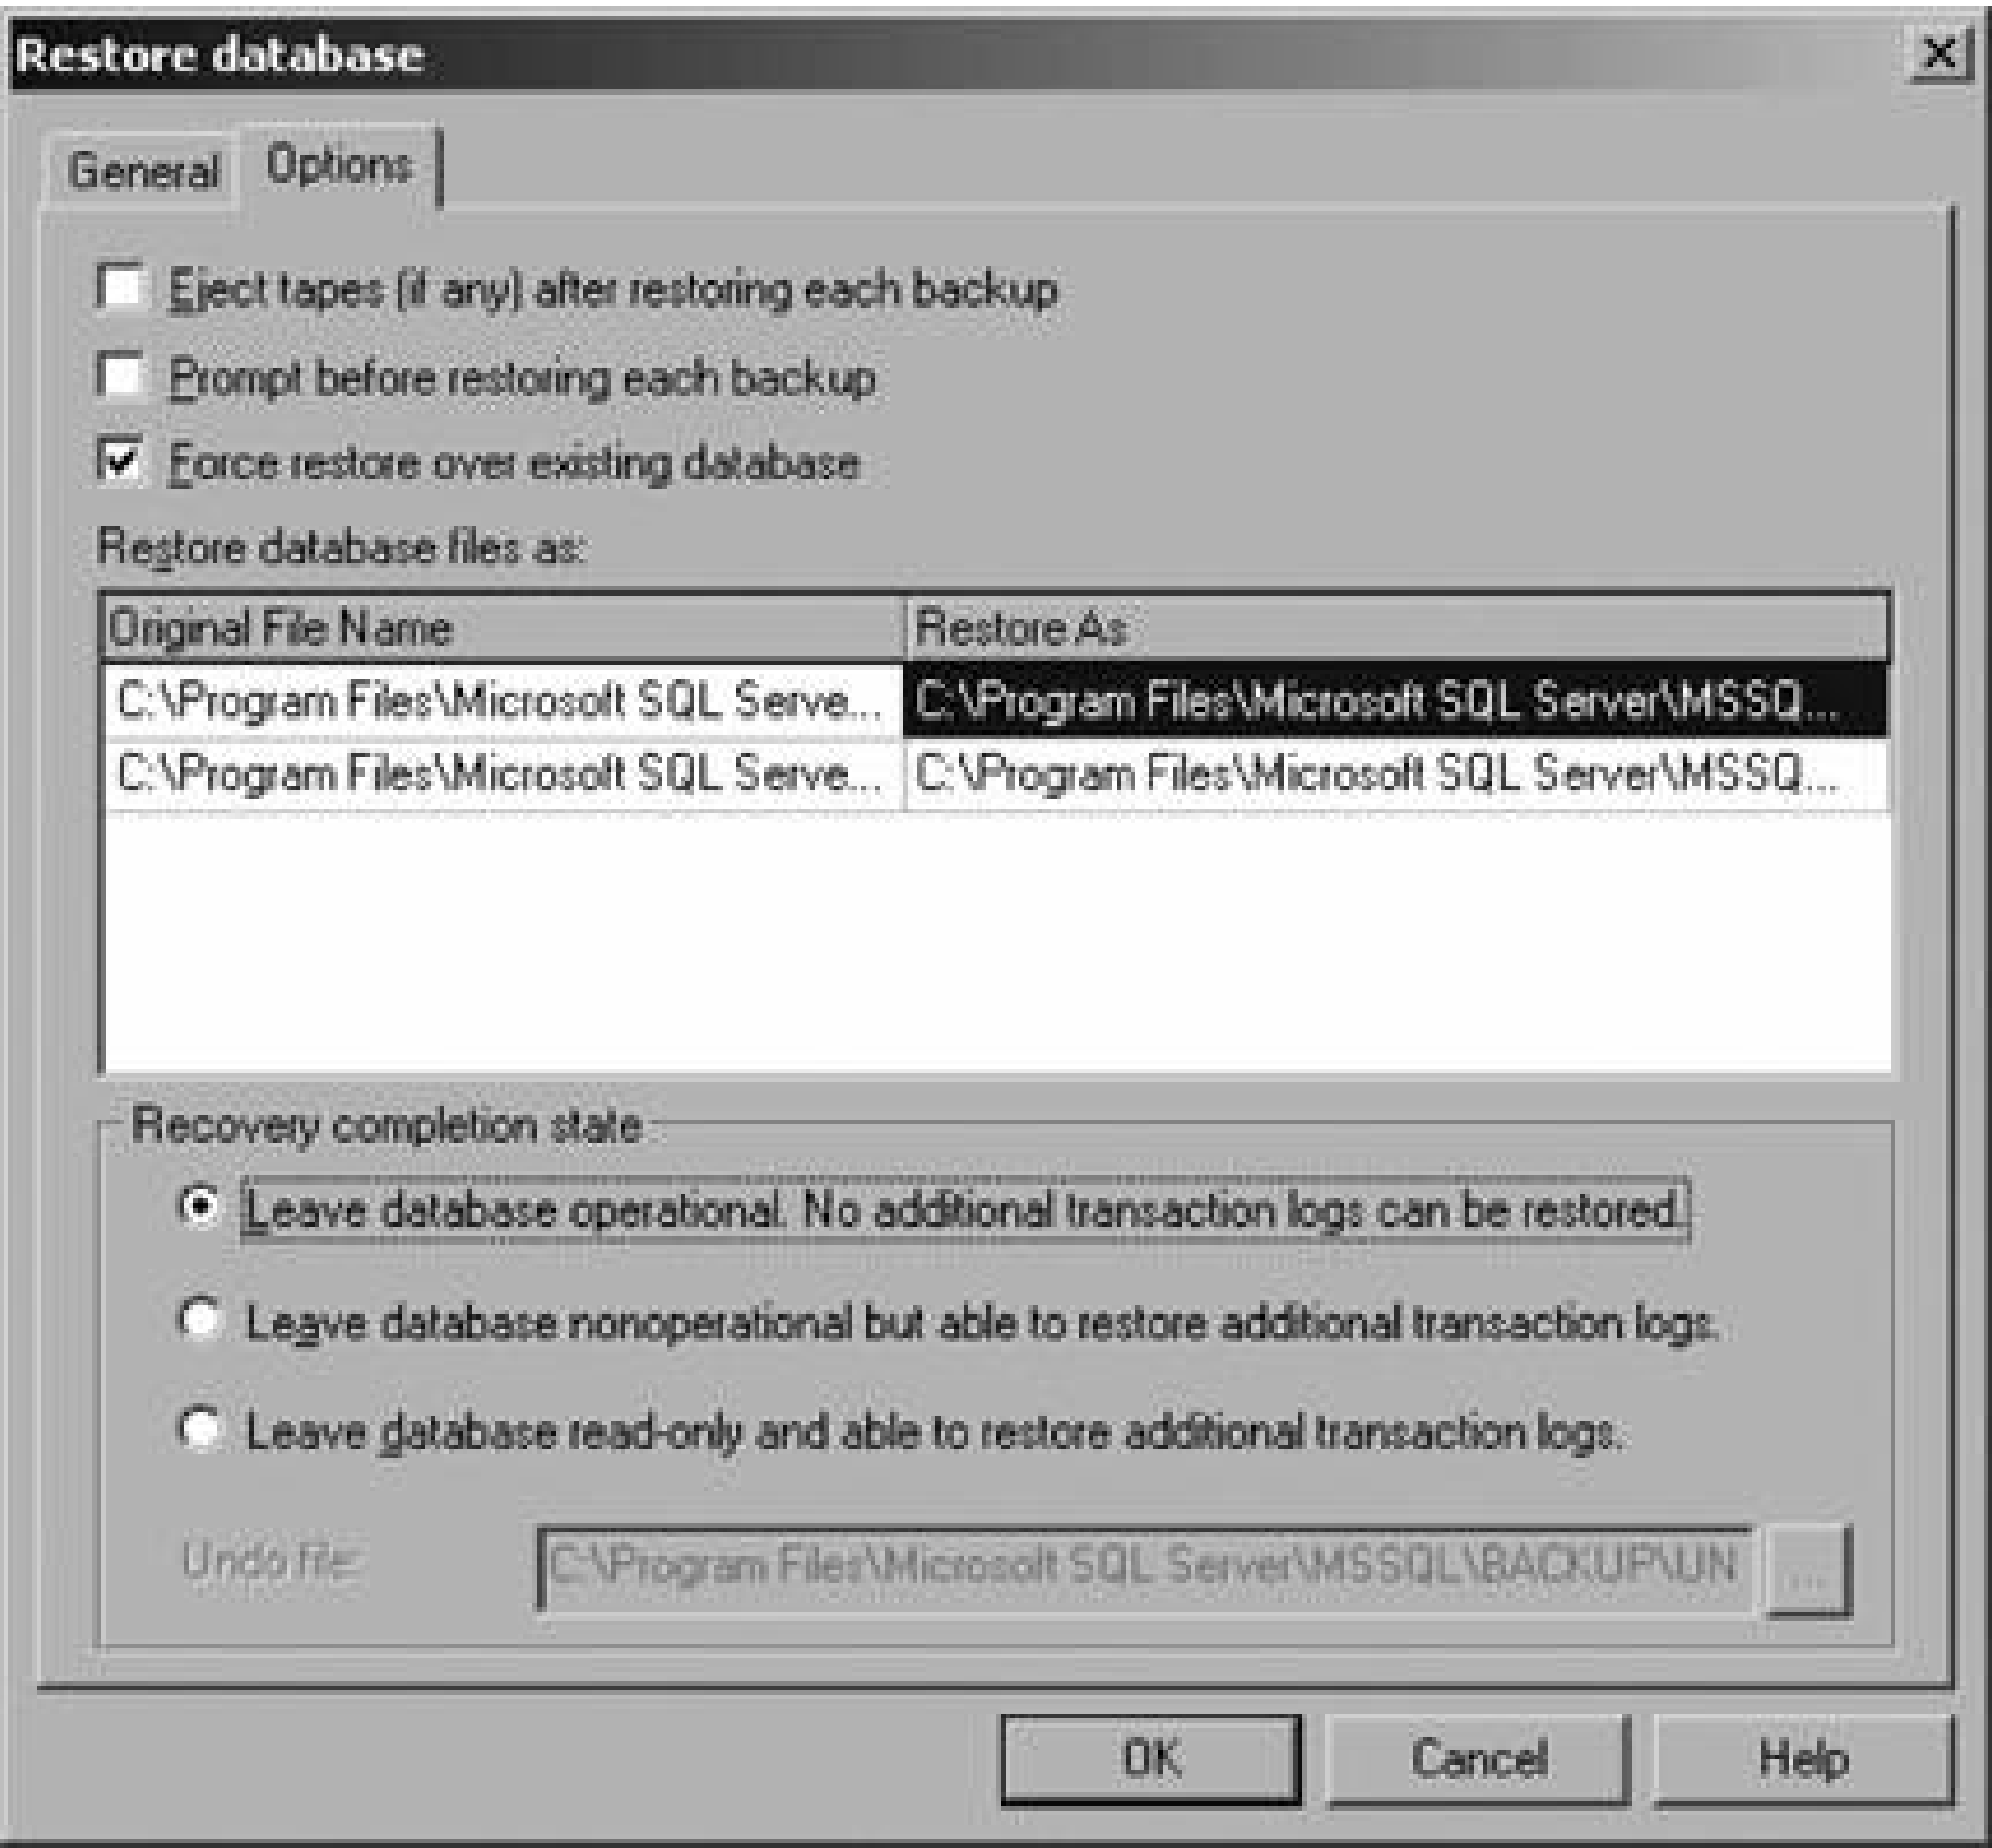


Figure 7-19. Options for a restore job



By selecting the "Leave database operational" option in the "Recovery completion state" box, the OnePoint database will automatically be brought back online when the restore completes successfully. To start the restore, click OK. SQL Enterprise Manager will display a progress box for the restore process and notify you when it has completed successfully.

### 7.1.6.2. Restoring system databases

Restoring the system databases is only slightly more complex than restoring user databases like OnePoint. This type of restore is required if the msdb and/or the master databases become corrupted or you are recovering from the loss of a whole server. Detailed discussion of the steps for whole server recovery is beyond the scope of this book, but basically the steps include:

1. Hardware configuration with sufficient capacity to serve in the role of the lost server.
2. OS installation and configuration to logically reproduce the lost server. This includes drive configuration, installed service packs and patches, security and account configuration, and server name.
3. Domain membership in the same domain as the lost server.
4. Identical SQL installation including features, installation paths, and service packs.

5. Any other configured features such as scheduled jobs or shares that were on the lost machine.

You can restore the system databases once you have reached the point in the rebuild process that you can start SQL, even if the master and msdb are corrupted and SQL never went down.

When you install SQL on the recovery server, new copies of the master and msdb in their default configurations are created. Therefore, the master will be lacking any SQL logins that had been defined and it will have no record of the other user databases. The msdb will only contain records of any default jobs. So, even though they are there, they are not really in a usable state. The restore process is below:

1. Start by stopping the SQL Server service (Start → All Programs → Administrative Tools → Services → MSSQLSERVER). This will also stop the SQLSERVERAGENT service.
2. Next, from a command prompt on the recovery server, navigate to the *<InstallDrive>\Program Files\Microsoft SQL Server\MSSQL\bin* directory. From here run **SQLSERVER.EXE -c -m**. This starts SQL Server in single-user mode (see [Figure 7-20](#)).
3. Leaving the command prompt window open, open SQL Enterprise Manager and navigate to the database you are restoring, which is the master database for this example.

Figure 7-20. Single-user mode startup messages

4. Invoke a restore operation on the master database. You will have to provide the path to the las



master backup file and select the "Force restore over existing database" option. Running the restore yields a message box stating, "The master database has been successfully restored. Shutting down SQL server. SQL Server is terminating this process." See [Figure 7-21](#).

5. Close the command prompt once SQL has been shut down and start the MSSQLSERVER and SQLSERVERAGENT processes.

Repeat this process for the other system databases if required and then perform the restore of the OnePoint database following the previously outlined procedures.

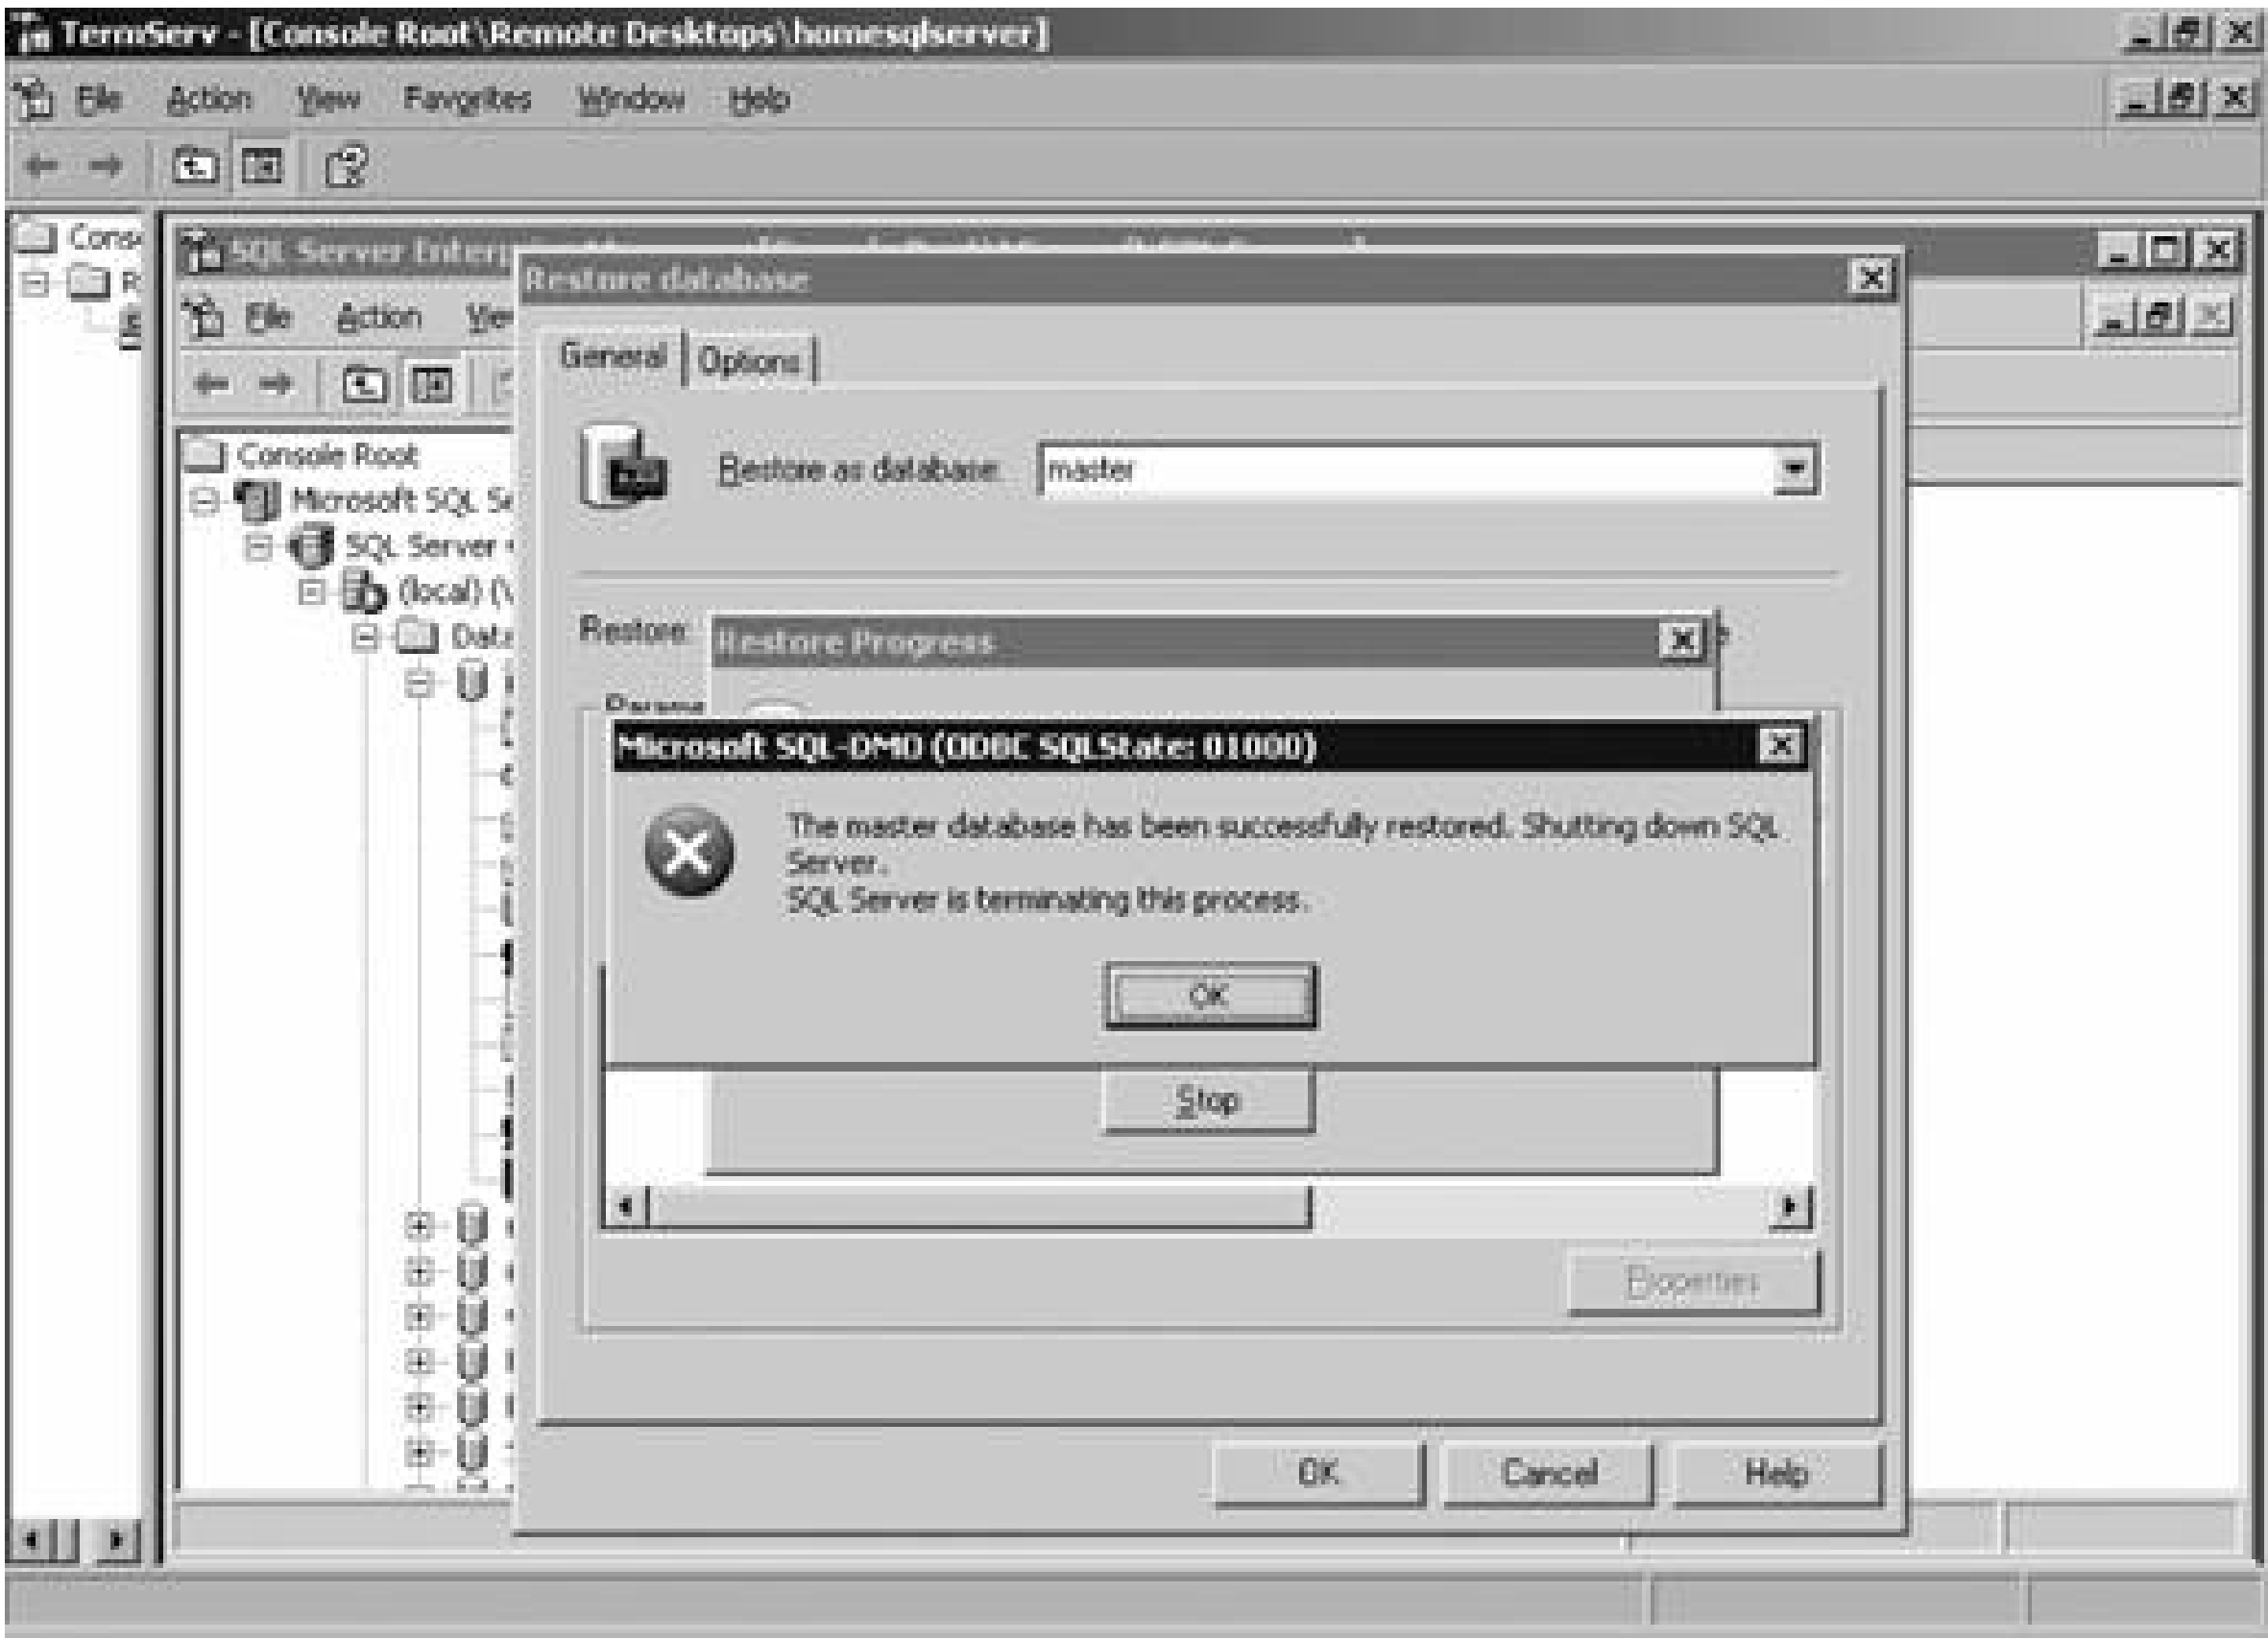


# 7.2. Data Transformation Service

DTS is a SQL Server 2000 tool that takes data from a source database, manipulates it, and inserts it into a destination database. When MOM Reporting is installed, a scheduled task is created on the reporting server that calls a DTS package-creating executable. The executable, located in *<installdrive> \Program Files\Microsoft System Center Reporting\Reporting*, is named *MOM.Datawarehousing.DTSPackageGenerator.exe*. This DTS package then takes operational data, transforms it into the correct format for the System Center Reporting database, and inserts the data into that database.

All the configuration options are set in this command-line call of the scheduled task:

Figure 7-21. Successful restore of the master database



```
MOM.Datawarehousing.DTSPackageGenerator.exe /silent /srcserver:homesqlserver /srcdb:
OnePoint /dwserver:HOMESQLSERVER /dwdb:SystemCenterReporting /product:"Microsoft
Operations Manager"
```

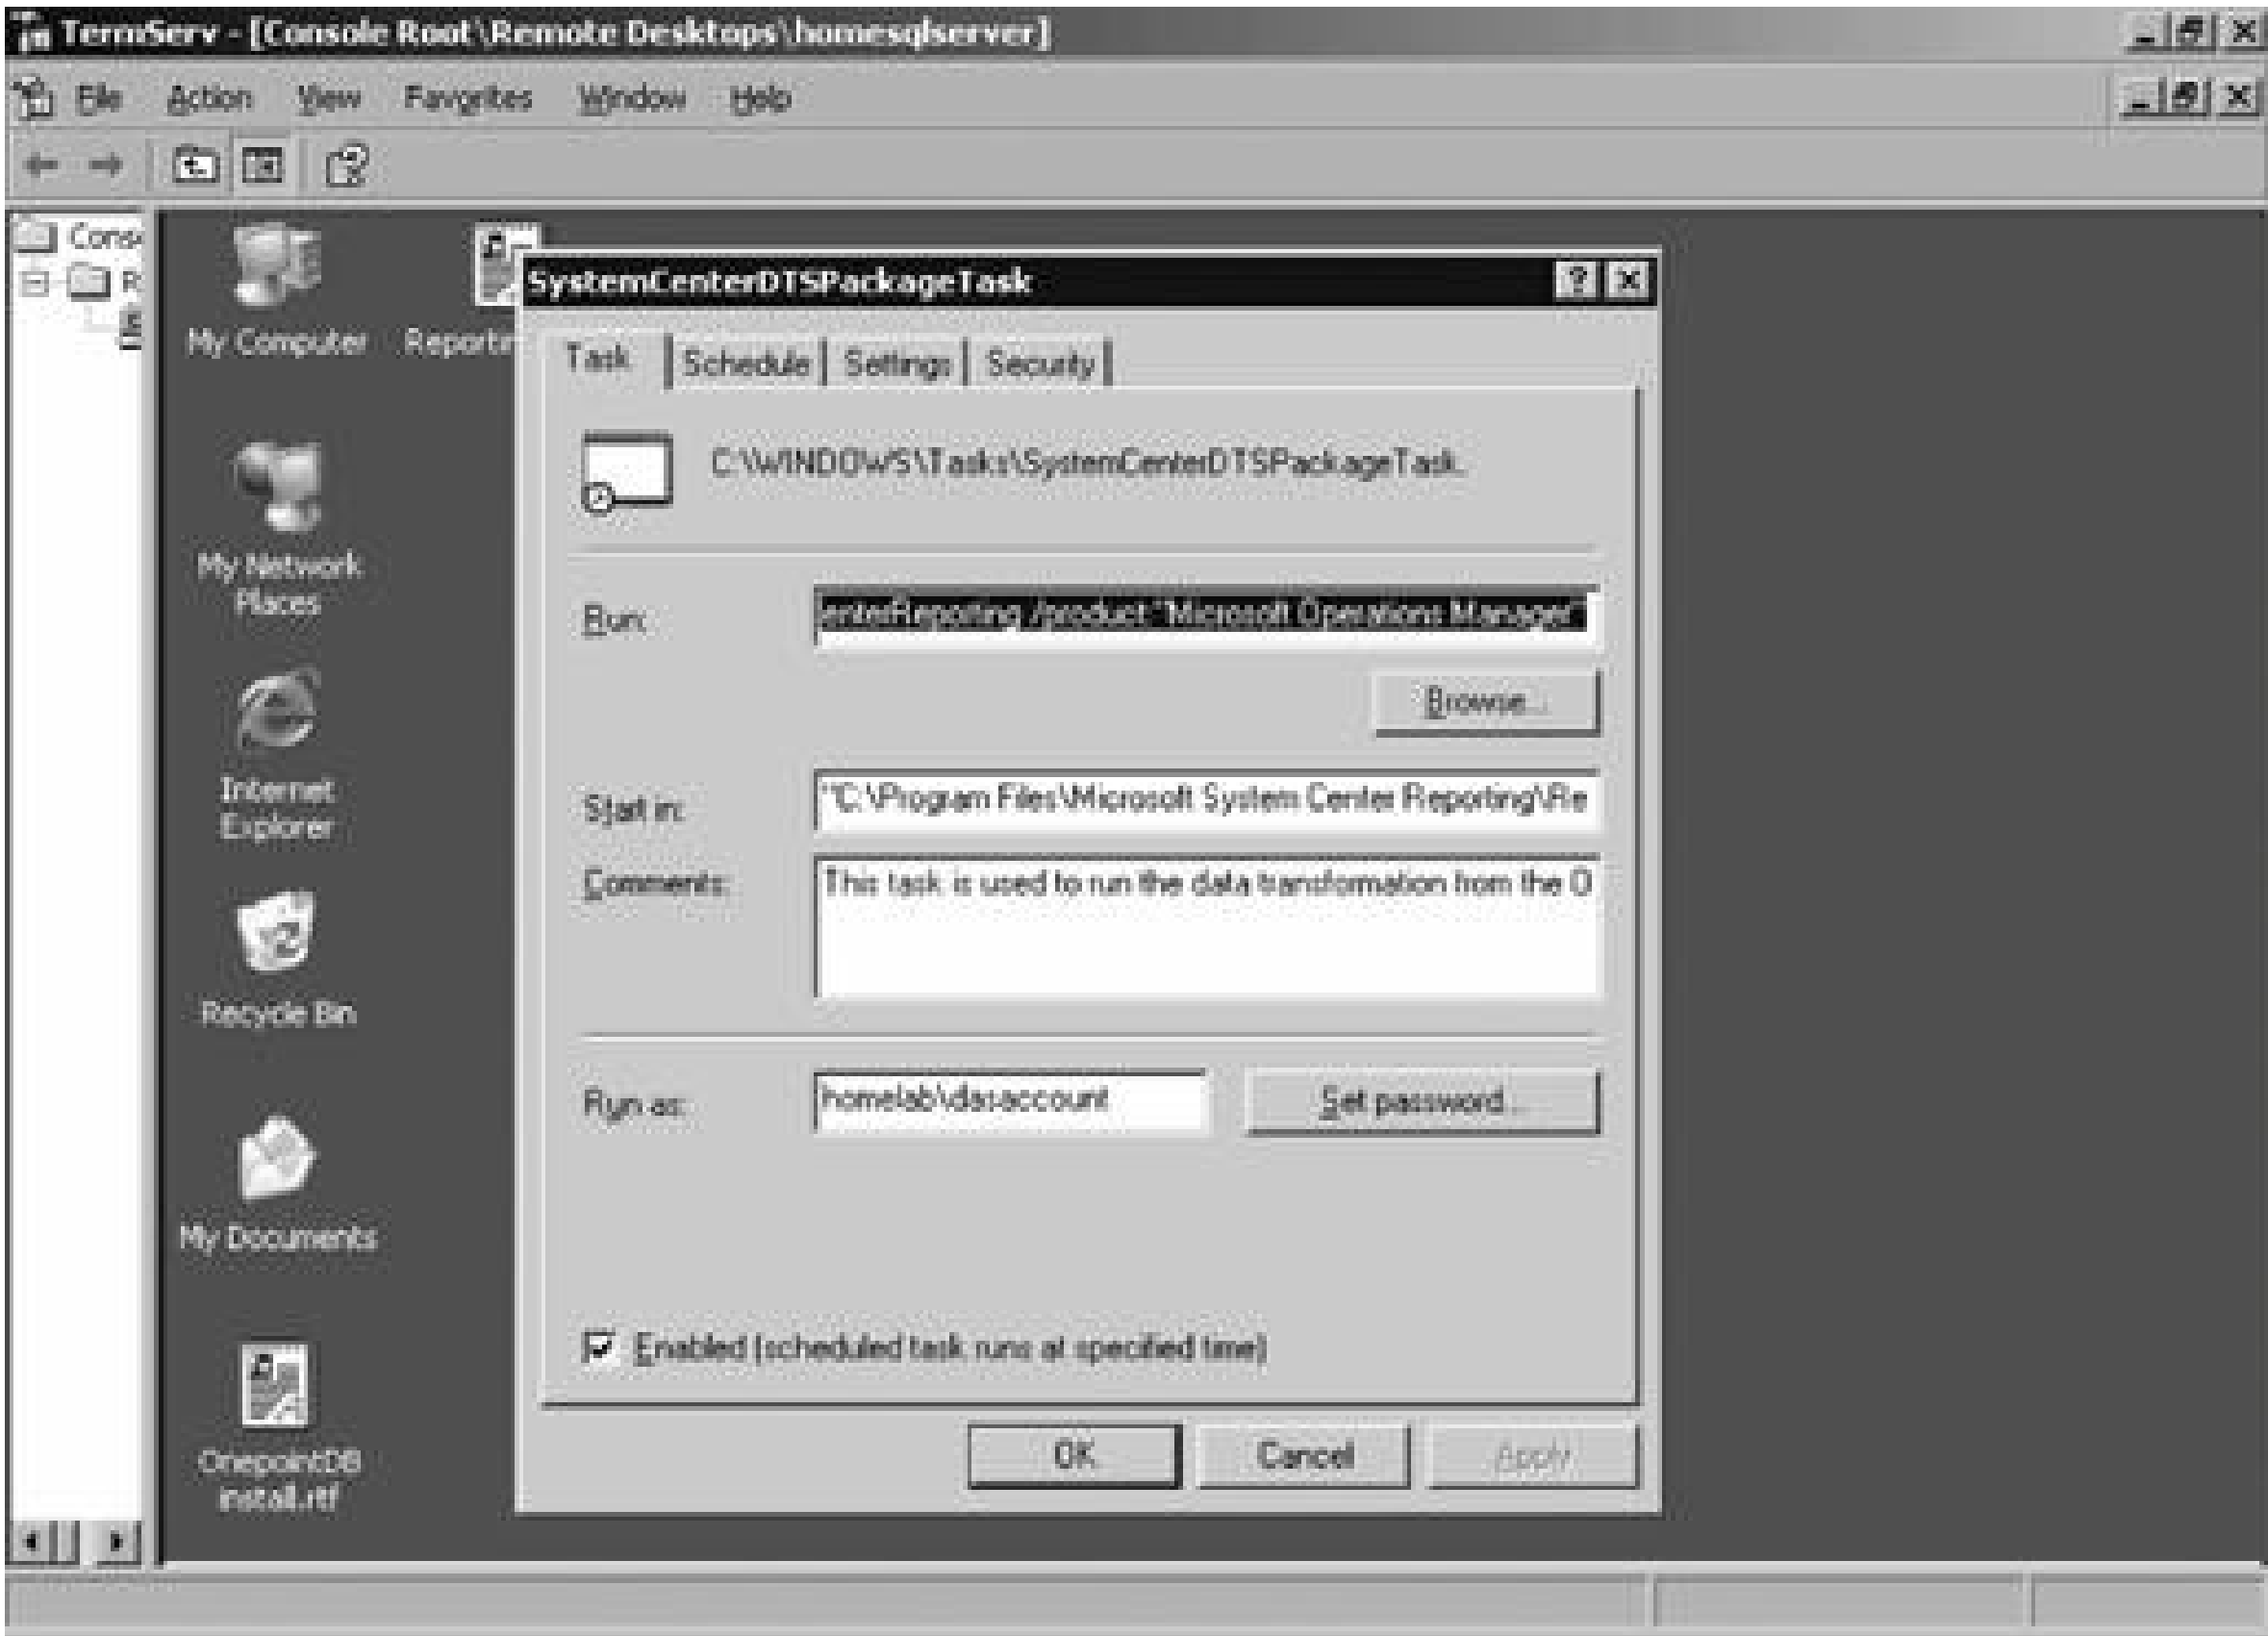
There is nothing else that you should need to do. However, since DTS touches the OnePoint database on a daily and on-demand basis, you should be familiar with it. There may come a time when you want to

trigger a data transfer, such as before performing maintenance on the OnePoint database server that requires an outage. You want to make sure that the most current data has been copied to the data warehouse before proceeding.

By default, the scheduled task (see Figure 7-22) runs once a day at 1:00 a.m. and transfers data that is more than five minutes old and has been newly written to the OnePoint database since the last transfer. Chapter 8 covers the installation of MOM Reporting Services, but as you can see in Figure 7-22, this tasks needs a security context to run in, which is supplied during install. Since this account must have access to OnePoint and System Center Reporting databases, it makes sense to use the DAS Account, bu you don't have to; any account with the appropriate permissions will do.

Deselecting the Enabled checkbox lets you disable this task when restoring the OnePoint database.

Figure 7-22. DTS OnePoint to System Center Reporting data transform task







## 7.3. MOM 2005 Reporting Databases

The last two databases that need to be included in your MOM backup and restore plans are the SystemCenterReporting database and the ReportServer database. As mentioned in the previous section, the SystemCenterReporting database is the actual data warehouse where alert, performance, and event data reside for long-term storage. The ReportServer database stores the report definitions, metadata, and any cached reports. In SQL Server Enterprise Manager you will also see the ReportServerTempDB, which is used during the reporting process but does not store any relevant data long term.

### 7.3.1. SystemCenterReporting Database

This database houses three primary types of data: dimension data, fact data, and snapshot fact data. Each data type is groomed differently

Dimension records are never groomed out of the database, so this data will always exist. Think of this data type as tracking every type of record that has ever existed in its management group. For example, there is a dimension table in the database called `SC_AlertResolutionStateDimension_Table`. This table contains information about the available alert resolution states. If you create a new alert resolution state a record (a row in the table) is created here. If you later delete that custom alert resolution state, that change would be reflected in the OnePoint database, but not in the SystemCenterReporting database. This is necessary because even after the custom resolution state is removed from OnePoint, there will still be historical records in the SystemCenterReporting database that were set to that custom resolution state.

All of the events, alerts, and performance monitor data fall into the fact data type. By default it is retained in the SystemCenterReporting database for 395 days (13 months). This can be altered on a per table basis through the use of a stored procedure called `p_updategroomdays`. To change the retention interval, run the following query in the SQL Query Analyzer:

```
exec p_updategroomdays 'TableName', DaysToRetainData
```

`TableName` is the name of the targeted table and `DaysToRetainData` is the number of days for which to retain the data. The tables are:

- `AC_AlertFact_Table`
- `SC_AlertHistoryFact_Table`
- `SC_AlertToEventFact_Table`
- `SC_EventFact_Table`
- `SC_EventParameterFact_Table`

- SC\_SampledNumericDataFact\_Table

Only the most current information is placed in the snapshot fact data tables, since the existing data is overwritten with every DTS cycle. If it is not overwritten, the data is groomed out after three days.

Since the SystemCenterReporting database is expected to grow quite large, be sure to allow sufficient time for backups to complete when creating your backup schedule. While the database is small, the backup to device-on-disk method will be usable. But as the database grows larger it may not be practical to maintain local disk space that is sufficient for the backups. You will be forced to go to multiple tapes or to a storage area network (SAN) that does have sufficient space.

Because both the SystemCenterReporting and ReportServer databases hold primarily unchanging historical data or report definition data, it is appropriate to perform a set of differential backups after you have obtained a good full backup of each, and then a full backup on the weekends. In this way, only changes in the databases since the last backup are captured, resulting in a smaller volume of data and quicker backups. Also, since these databases are not very dynamic, backing them up on a less than daily frequency is a good idea.

## 7.3.2. ReportServer Database

Report definitions can be imported at the same time that the management pack definitions are imported into the management group. Therefore, you can restore the default report definitions at any time. Like the management packs, you can customize the available reports (see Chapter 8). Therefore, you will need a method to protect those changes from SQL or server failures.

There are two ways to protect the data that is contained in this database. You can easily include it in your standard SQL database backup routine, but just as all of the management packs are contained in the OnePoint definition, so too are all the reports contained in the ReportServer backup. To recover a single report, you would have to restore the entire ReportServer database, which may not be the most practical thing to do.

Instead, there is a reporting equivalent of the ManagementModuleUtility called the Report Utility (*rptutil.exe*). It is installed on the MOM reporting server in the *<installdrive> \Program Files\Microsoft System Center Reporting\Reporting* directory. With this tool you can import and export reports in groups or individually from the ReportServer database.

This is a command-line tool that, when run without any parameters, will export all the report definitions in the database and save them to a single file (*MOMReports.xml*) on the root of the computer that it is run from. This tool gives more granular control over imports and exports by specifying the path to an individual report or report folder.

For example, to back up the report definitions for Microsoft Operations Manager Management Alert Tuning reports you would specify the following command line:

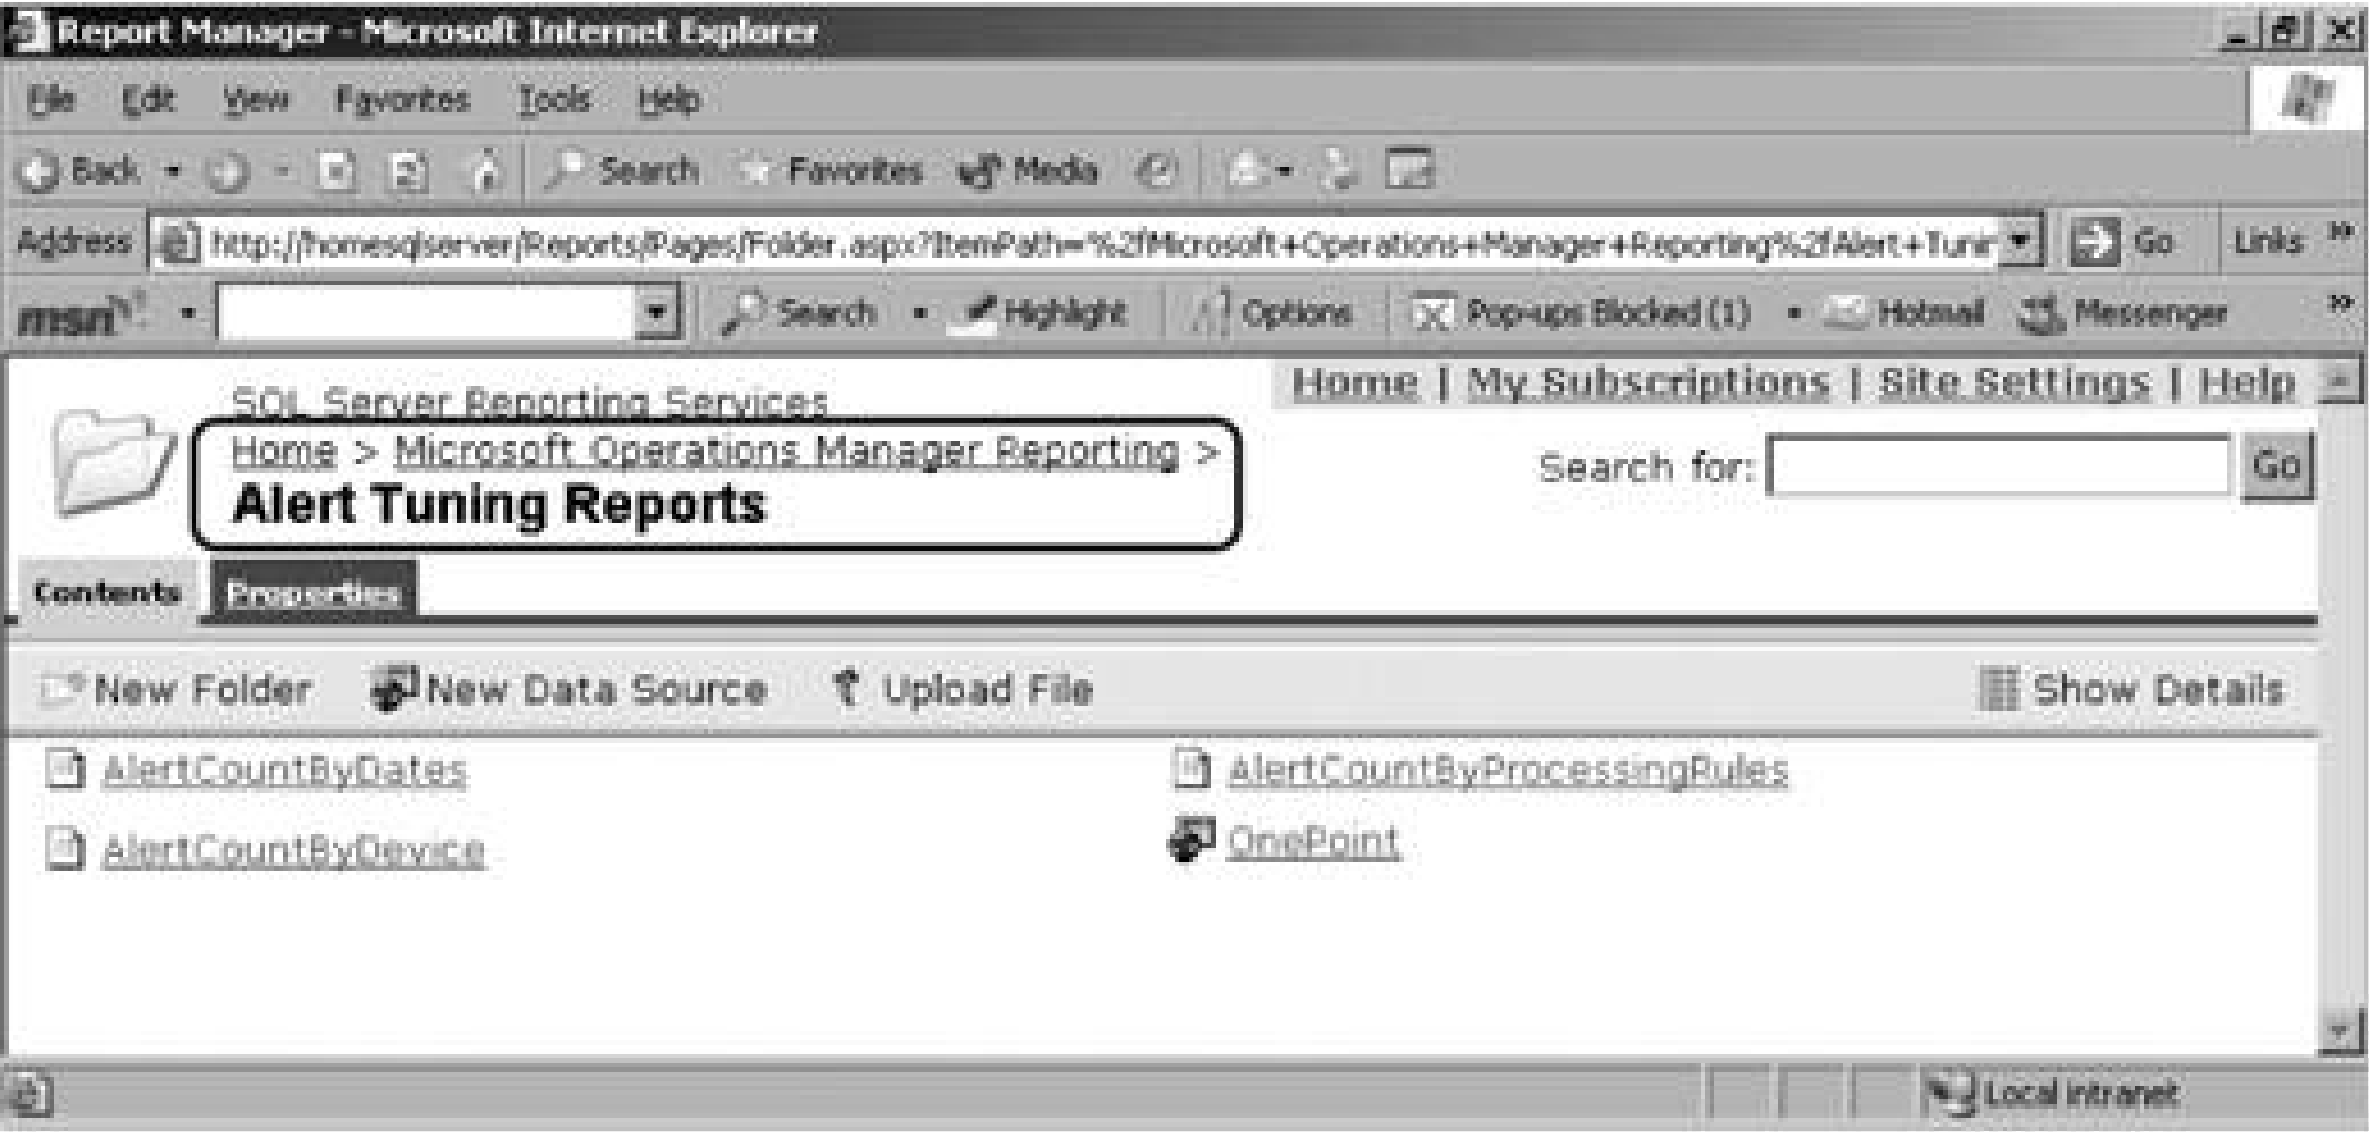
```
Rptutil /file:c:\MOMAlertTuningReports.xml /reportpath:"Microsoft Operations Manager Reporting/Alert Tuning Reports"
```

This would capture all reports in the Alert Tuning reports folder to the *MOMAlert TuningReports.xml* file.

The easiest way to find the path to an individual report or report folder is to open the reporting console (*http:// <reportservername>/Reports*). Navigate to the report folder of interestthe path to the folder is tracked in the header of the page and updated as you navigate through the console. For example, in Figure 7-23 the path to the Alert Tuning Reports folder is circled.

In the command-line example, it is not necessary to include the top-level Home folder in the path. If you want to target an individual report at the command line, simply continue the path with the report name.

Figure 7-23. Path to the Alert Tuning Reports folder







# 7.4. Summary

Effective administration of MOM 2005 includes working with the OnePoint and reporting databases through the SQL Enterprise Manager and the SQL Query analyzer. These tools will be used to obtain information about the databases, to configure and run backups, and to restore databases as needed.

The next chapter addresses the installation, administration, and navigation of the reporting console. It touches on how to modify existing reports and how to help reporting customers get the most out of your MOM data.



# Chapter 8. MOM 2005 Reporting

The previous chapter introduced the database that is involved with storing MOM 2005 data for long periods (SystemCenterReporting ) and the database used for generating reports on that data (ReportServer). Although they work together to deliver the MOM 2005 Reporting solution, they are two different products. The ReportServer database is created when SQL 2000 Reporting Service (SRS) is installed and the SystemCenterReporting is created when MOM 2005 Reporting is installed. SQL Reporting Services on SQL 2000 is a standalone product that MOM 2005 Reporting builds on.

This chapter addresses the planning, implementation, and administration of SQL Reporting Services and MOM 2005 Reporting in support of the MOM 2005 Reporting Solution. It introduces the Report Manager, which is the primary tool used by administrators to manage and configure reports and reporting services, as well as by all users to access reports and to configure subscriptions to reports.

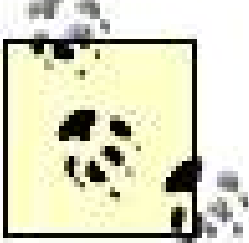
Authoring reports is not addressed since it requires Microsoft Visual Studio 2003 or newer. See *Microsoft SQL Server 2000 Reporting Services* (Osborne/McGraw-Hill). As a MOM 2005 administrator, your daily tasks would not involve creating new reports and although that is not a difficult task, it's beyond the scope of what you need to know to successfully administer MOM 2005. For a decent tutorial on Reporting Services see *Microsoft SQL Server 2000 Reporting Services Step by Step* (Microsoft Press).

There is a sample set of reports (on the SQL reporting server CD under the Extras folder) that can be used for monitoring the longest and shortest time it takes to run a report, but it requires the MOM administrator to use Visual Studio. See *Programming C#* (O'Reilly) for an overview of Visual Studio.

Because MOM 2005 Reporting builds on SQL Server 2000 Reporting Services, there is a whole string of items that must be installed in the correct order before importing MOM 2005 reports and making them available to users. The planning for this installation was completed in [Chapter 2](#) when you calculated the starting size of the data warehouse database and its transaction logs and created the DAS account. With this out of the way, the next step is to begin installing the prerequisites.

# 8.1. Installation

In addition to calculating the SystemCenterReporting database starting size, you must also choose the configuration of SQL 2000 Reporting Services and install it. For the scope of this chapter, one server will host all of the MOM database roles. Specifically, *homemomserver3* will be the database and reporting server.



SQL Reporting Services is a separate product from MOM 2005 and it enjoys great popularity. SQL Reporting Services scales from a single server configuration with all components existing on a single machine, to a web farm in which components are distributed across multiple load- balanced web frontend and database backend machines. A single instance of SQL Reporting can support the reporting needs of multiple applications, resulting in a shared SQL Reporting environment. However, planning a distributed, shared SQL Reporting installation is beyond the scope of this book.

SQL Reporting services come in four versions: *standard* for an all-in-one machine deployment, *enterprise* for large-scale deployment or when all features are required, *developer* for authoring reports, and an *evaluation* edition. The evaluation and developer editions cannot be licensed for production use, so don't even consider those in your planning. In addition to providing support for a scalable web farm solution, the enterprise edition has the ability to generate reports and send them to various destinations (e.g., email and file shares) in varying formats based on address entries that you create and maintain in SQL. This is called a *data-driven subscription*. For this chapter, the standard edition of SQL Reporting is used (this can install on SQL 2000 Developer, Standard, or Enterprise) and all examples will be based on it. Fortunately, there is only one version of MOM 2005 reporting to install on top of SQL Reporting, so there is no decision to be made there.

## 8.1.1. SQL 2000 Reporting Services

All of the services necessary to support MOM Reporting will be installed on one dedicated server. This server will host both the ReportServer and SystemCenterReporting databases. Since this server will also host the Reports web site, IIS 6.0 is installed with ASP.NET support.

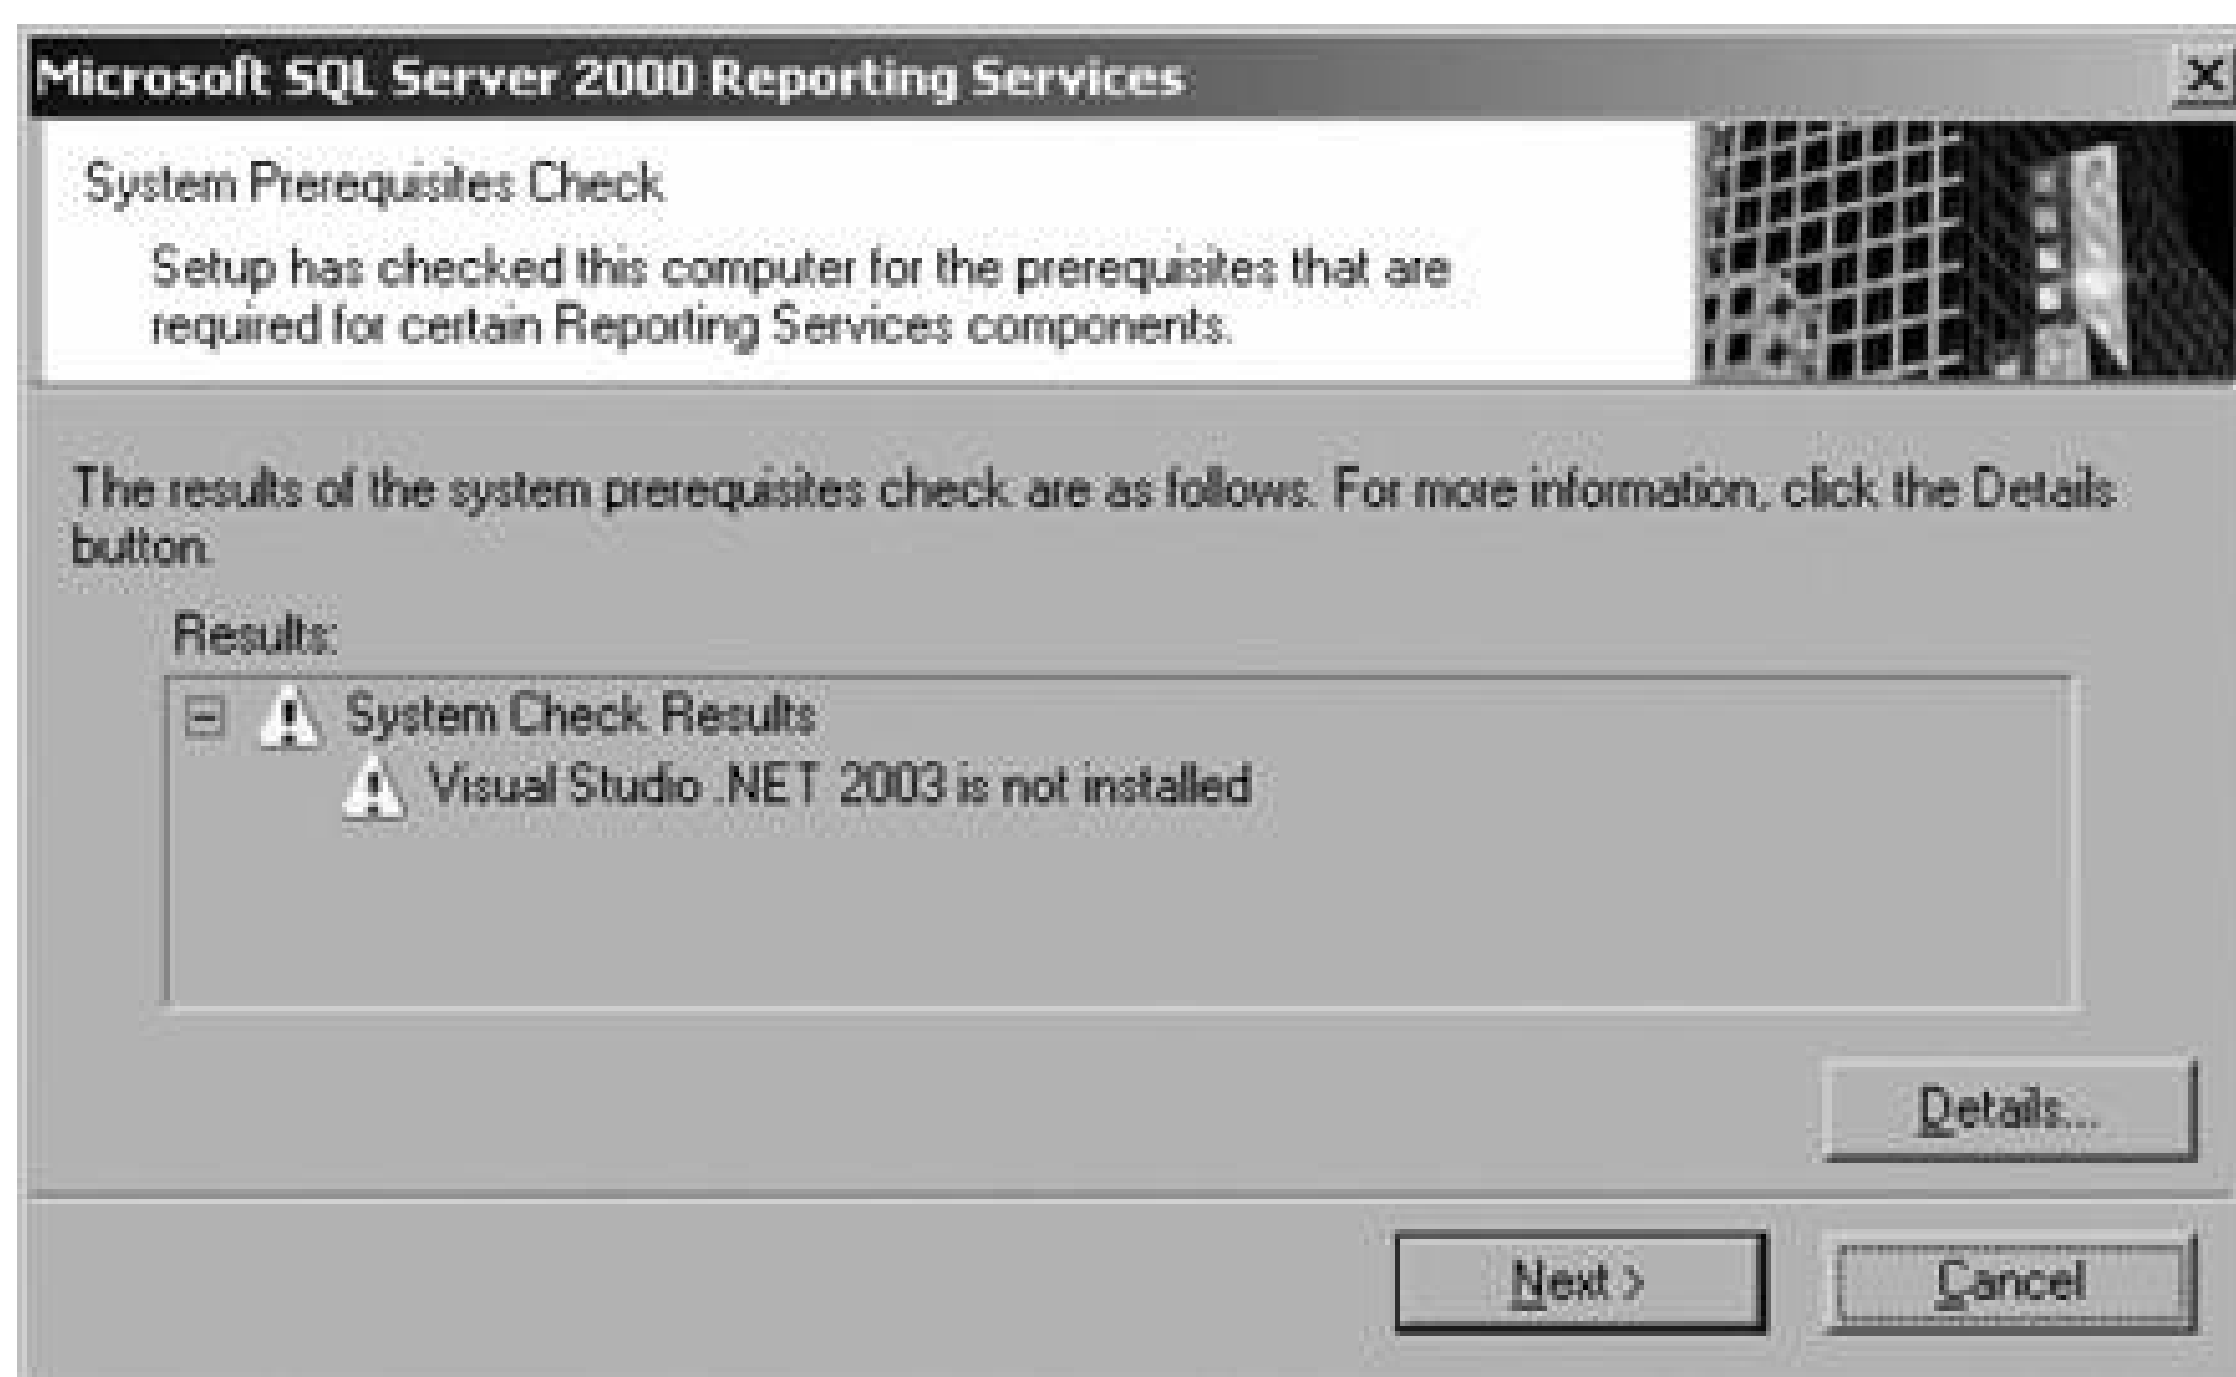
At the time of this writing, SQL 2000 SP4 and SQL Reporting Services SP2 are released. If you build your all-in-one reporting server with these service packs installed, the MOM 2005 Reporting component will not install and will fail at the prerequisite checker. It is not capable of recognizing patch levels above SQL SP3A or SQL Reporting Services SP1 until MOM 2005 SP1 is installed. The easiest upgrade path is to install MOM and MOM Reporting on the SQL SP3A and SQL Reporting SP1, apply MOM 2005 SP1, and then upgrade to SQL 2000 SP4 and SQL Reporting SP2.



Once SQL 2000 is installed and patched to SP3A, you can begin the SRS installation. Surprisingly enough, the SRS installation doesn't begin with the typical welcome page, so don't be thrown off by that. It begins with the end-user license agreement (EULA) and then proceeds through a component update and prerequisite checker for pages two and three. Page four then is the expected welcome page.

During the prerequisite check, if Visual Studio .NET 2003 has not been installed on the reporting server (it does not have to be), Setup will display a warning telling you that it is not present, as shown in [Figure 8-1](#).

Figure 8-1. SQL Reporting Services install prerequisite checker warning



This warning can be safely ignored. Visual Studio .NET 2003 is used to author and alter reports and is, therefore, not necessary on the production reporting server. Instead, it is better to install the SRS report designer and other client components on a workstation that already has Visual Studio .NET 2003 installed. When this is done, a report project wizard is made available in Visual Studio.

As a MOM administrator, you shouldn't have to create reports; the reports that come with the management packs have all been predefined and you will only need to import them. It is helpful, though, to understand the three components of a report:

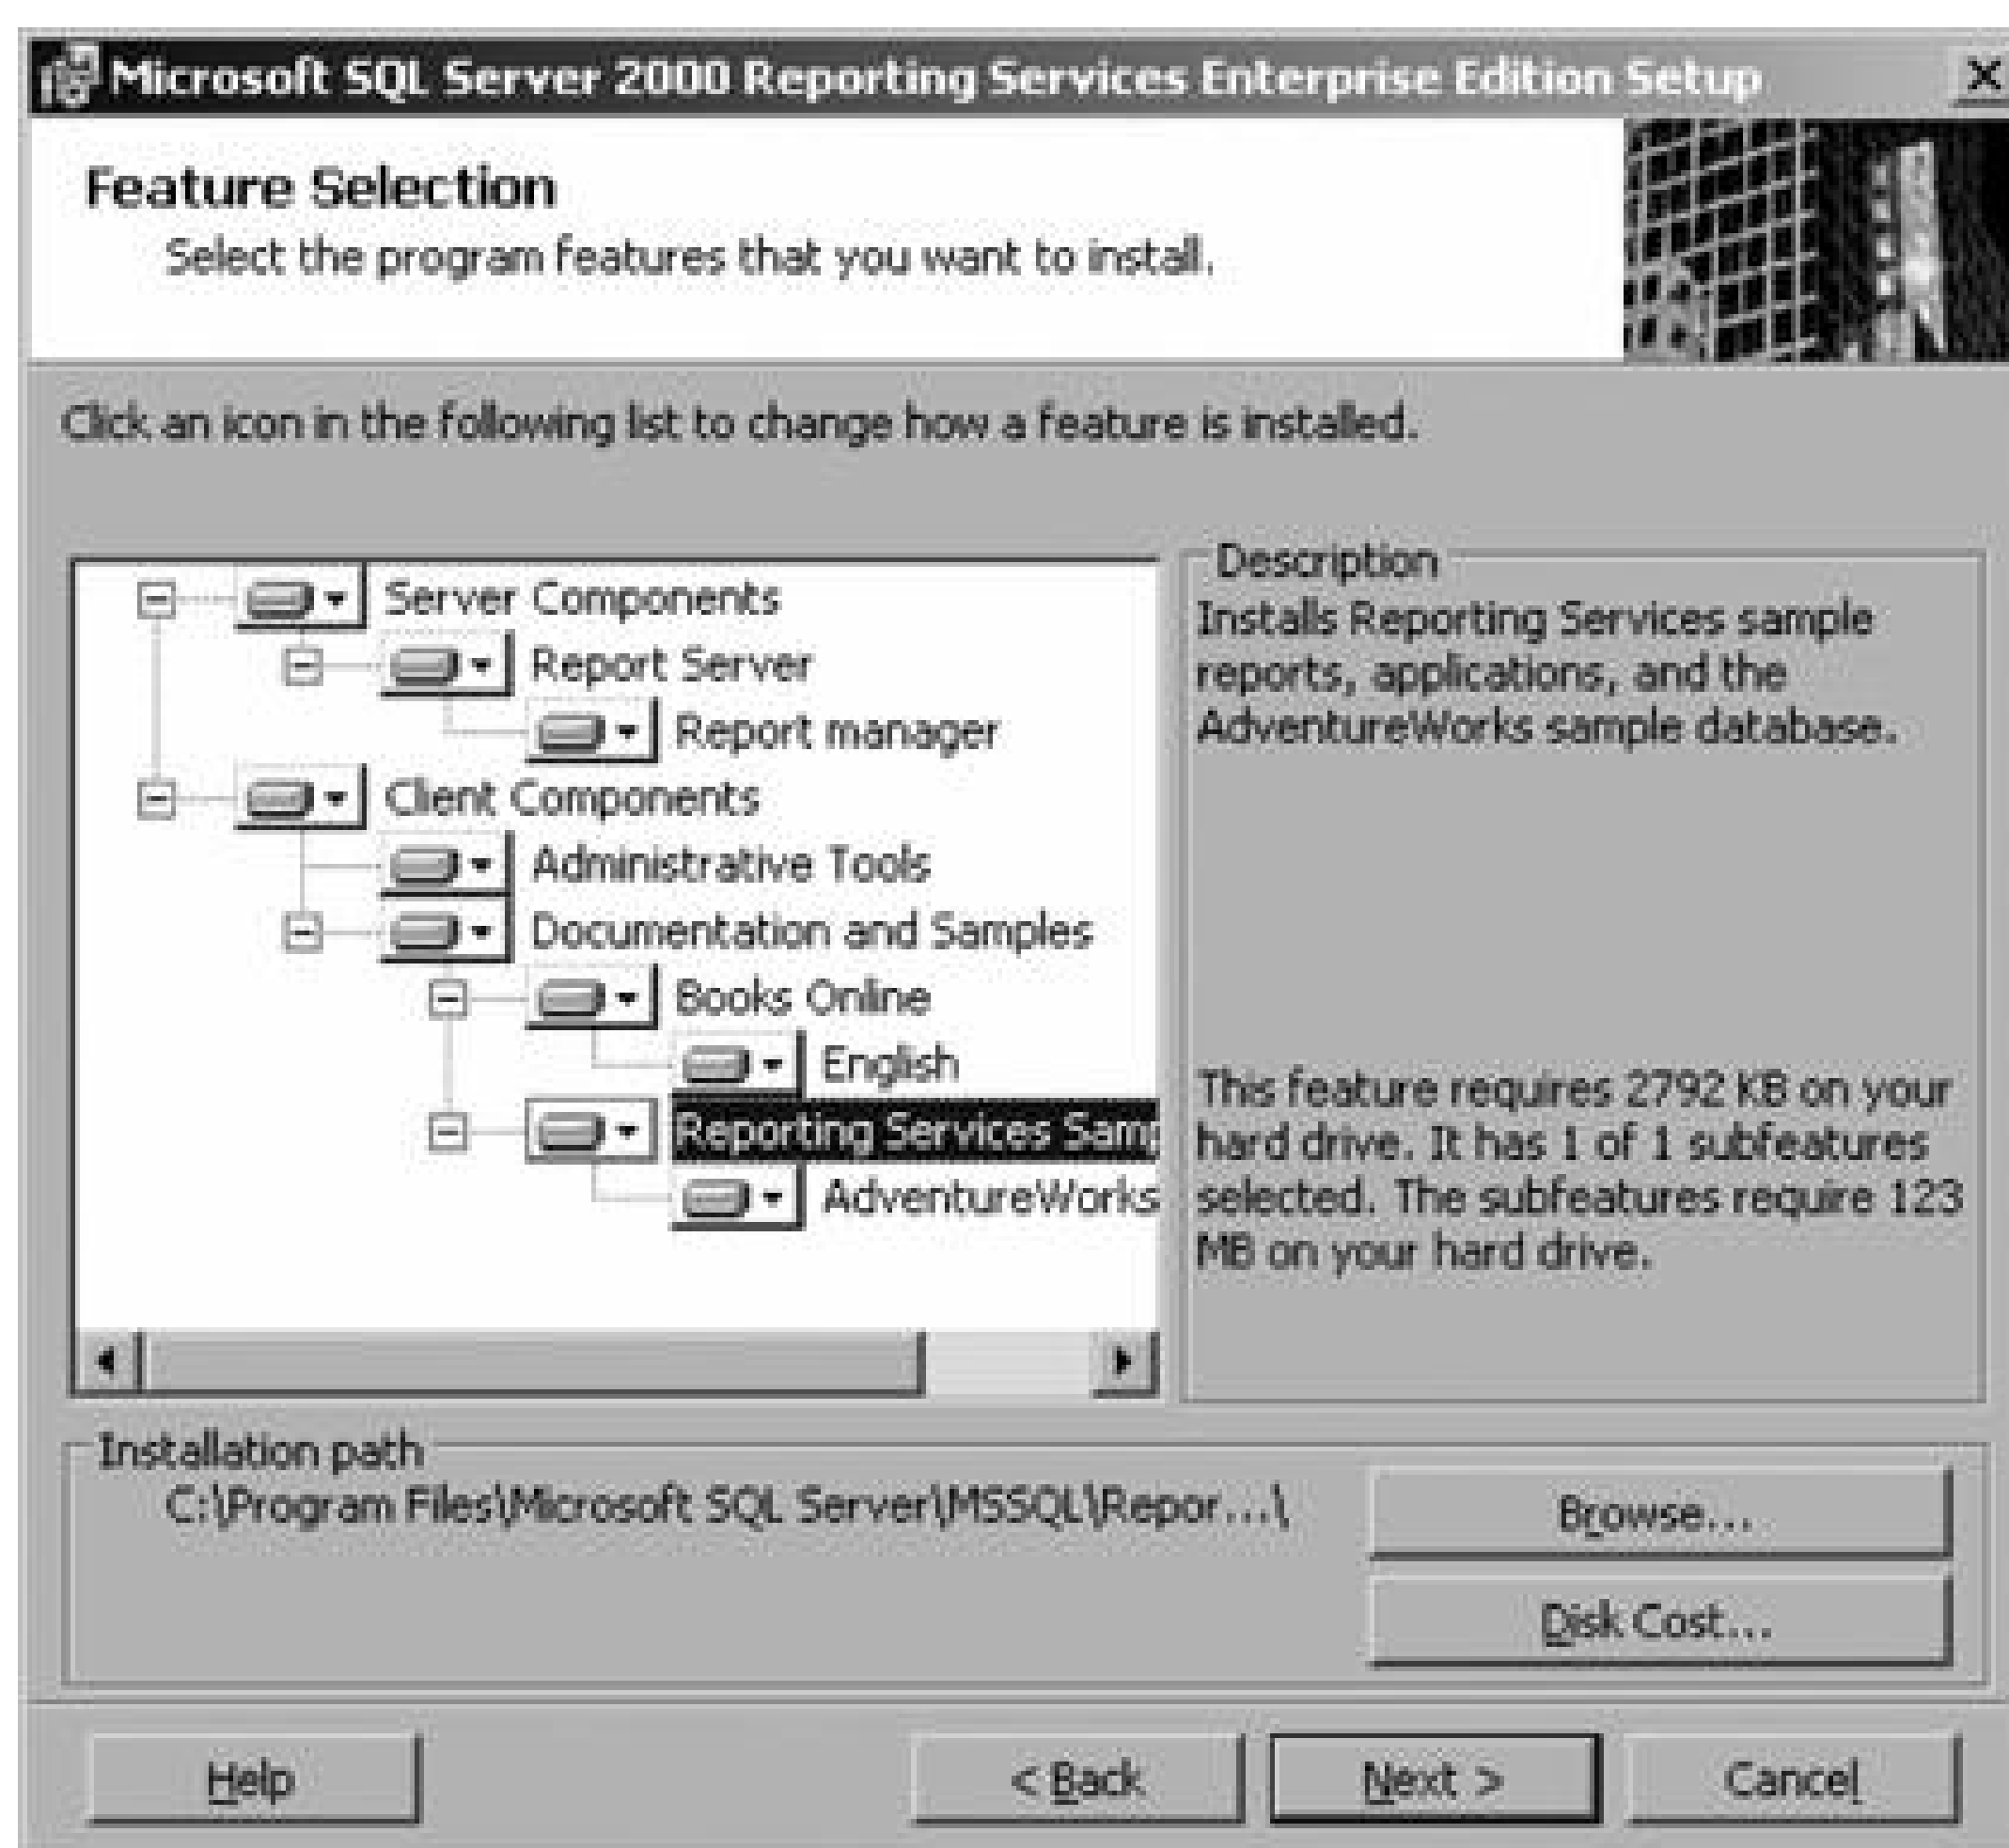
- The database that SRS will connect to in order to extract the data. This is called the *data source*.
- The SQL query returns a set of data that reports are built from. These queries can be very complex and require a deep knowledge of both the query language and the structure of the data source.

- Define the layout of the report. There are a number of templates available to assist you with this.

All three components can then be wrapped up into a report definition language file (*rdl*) and imported directly into the ReportServer database using the Report Manager, or placed into XML format, as has been done for the MOM management pack reports. Report definitions that are in XML format need to be imported using the management pack import/export wizard or the *rptutil.exe* tool.

Following the welcome and registration pages is the component selection page (see [Figure 8-2](#)).

Figure 8-2. SRS component installation for an all-in-one configuration



The SRS installation consists of server components and client components. The server components are the report server services and the Report Manager, which is the reporting web site. The client components include the report designer (not pictured here because Visual Studio is not installed), administrative tools, and the online books and samples. For this installation everything but the report designer will be installed.

On the next page you are prompted for the first of two sets of credentials. This first set is used to run the SRS service on the reporting machine (in this case *homemomserver3*). In [Figure 8-3](#), you have the option to use the network service account (because this is on Windows Server 2003), the local system account, or a domain service account. For this example, a domain service account was created (SRSSA) and given local administrator rights on *homemomserver3*. Also, you can (and should) configure the service to start automatically.

Figure 8-3. Selecting the credentials that the reporting service will run as



The next page (see [Figure 8-4](#)) allows you to change the default configuration for the virtual directories for report server, to redirect browsing to the default web sites, and to require SSL connections. Because this will be a dedicated server, the default values for the virtual directories are unchanged. In addition, because this server is not exposed to the Internet, SSL encryption will not be required.

On the next page, you identify the server that will house the report server database, the name of the database, and the credentials that the report service will use to connect to the report server database. In [Figure 8-5](#), the local SQL Server instance has been accepted as well as the default database name. For the credentials used to connect to the database, the previously created homelab\dasaccount is specified. There is only one account for all database access across all of the MOM servers.

Figure 8-4. Configuring virtual directories



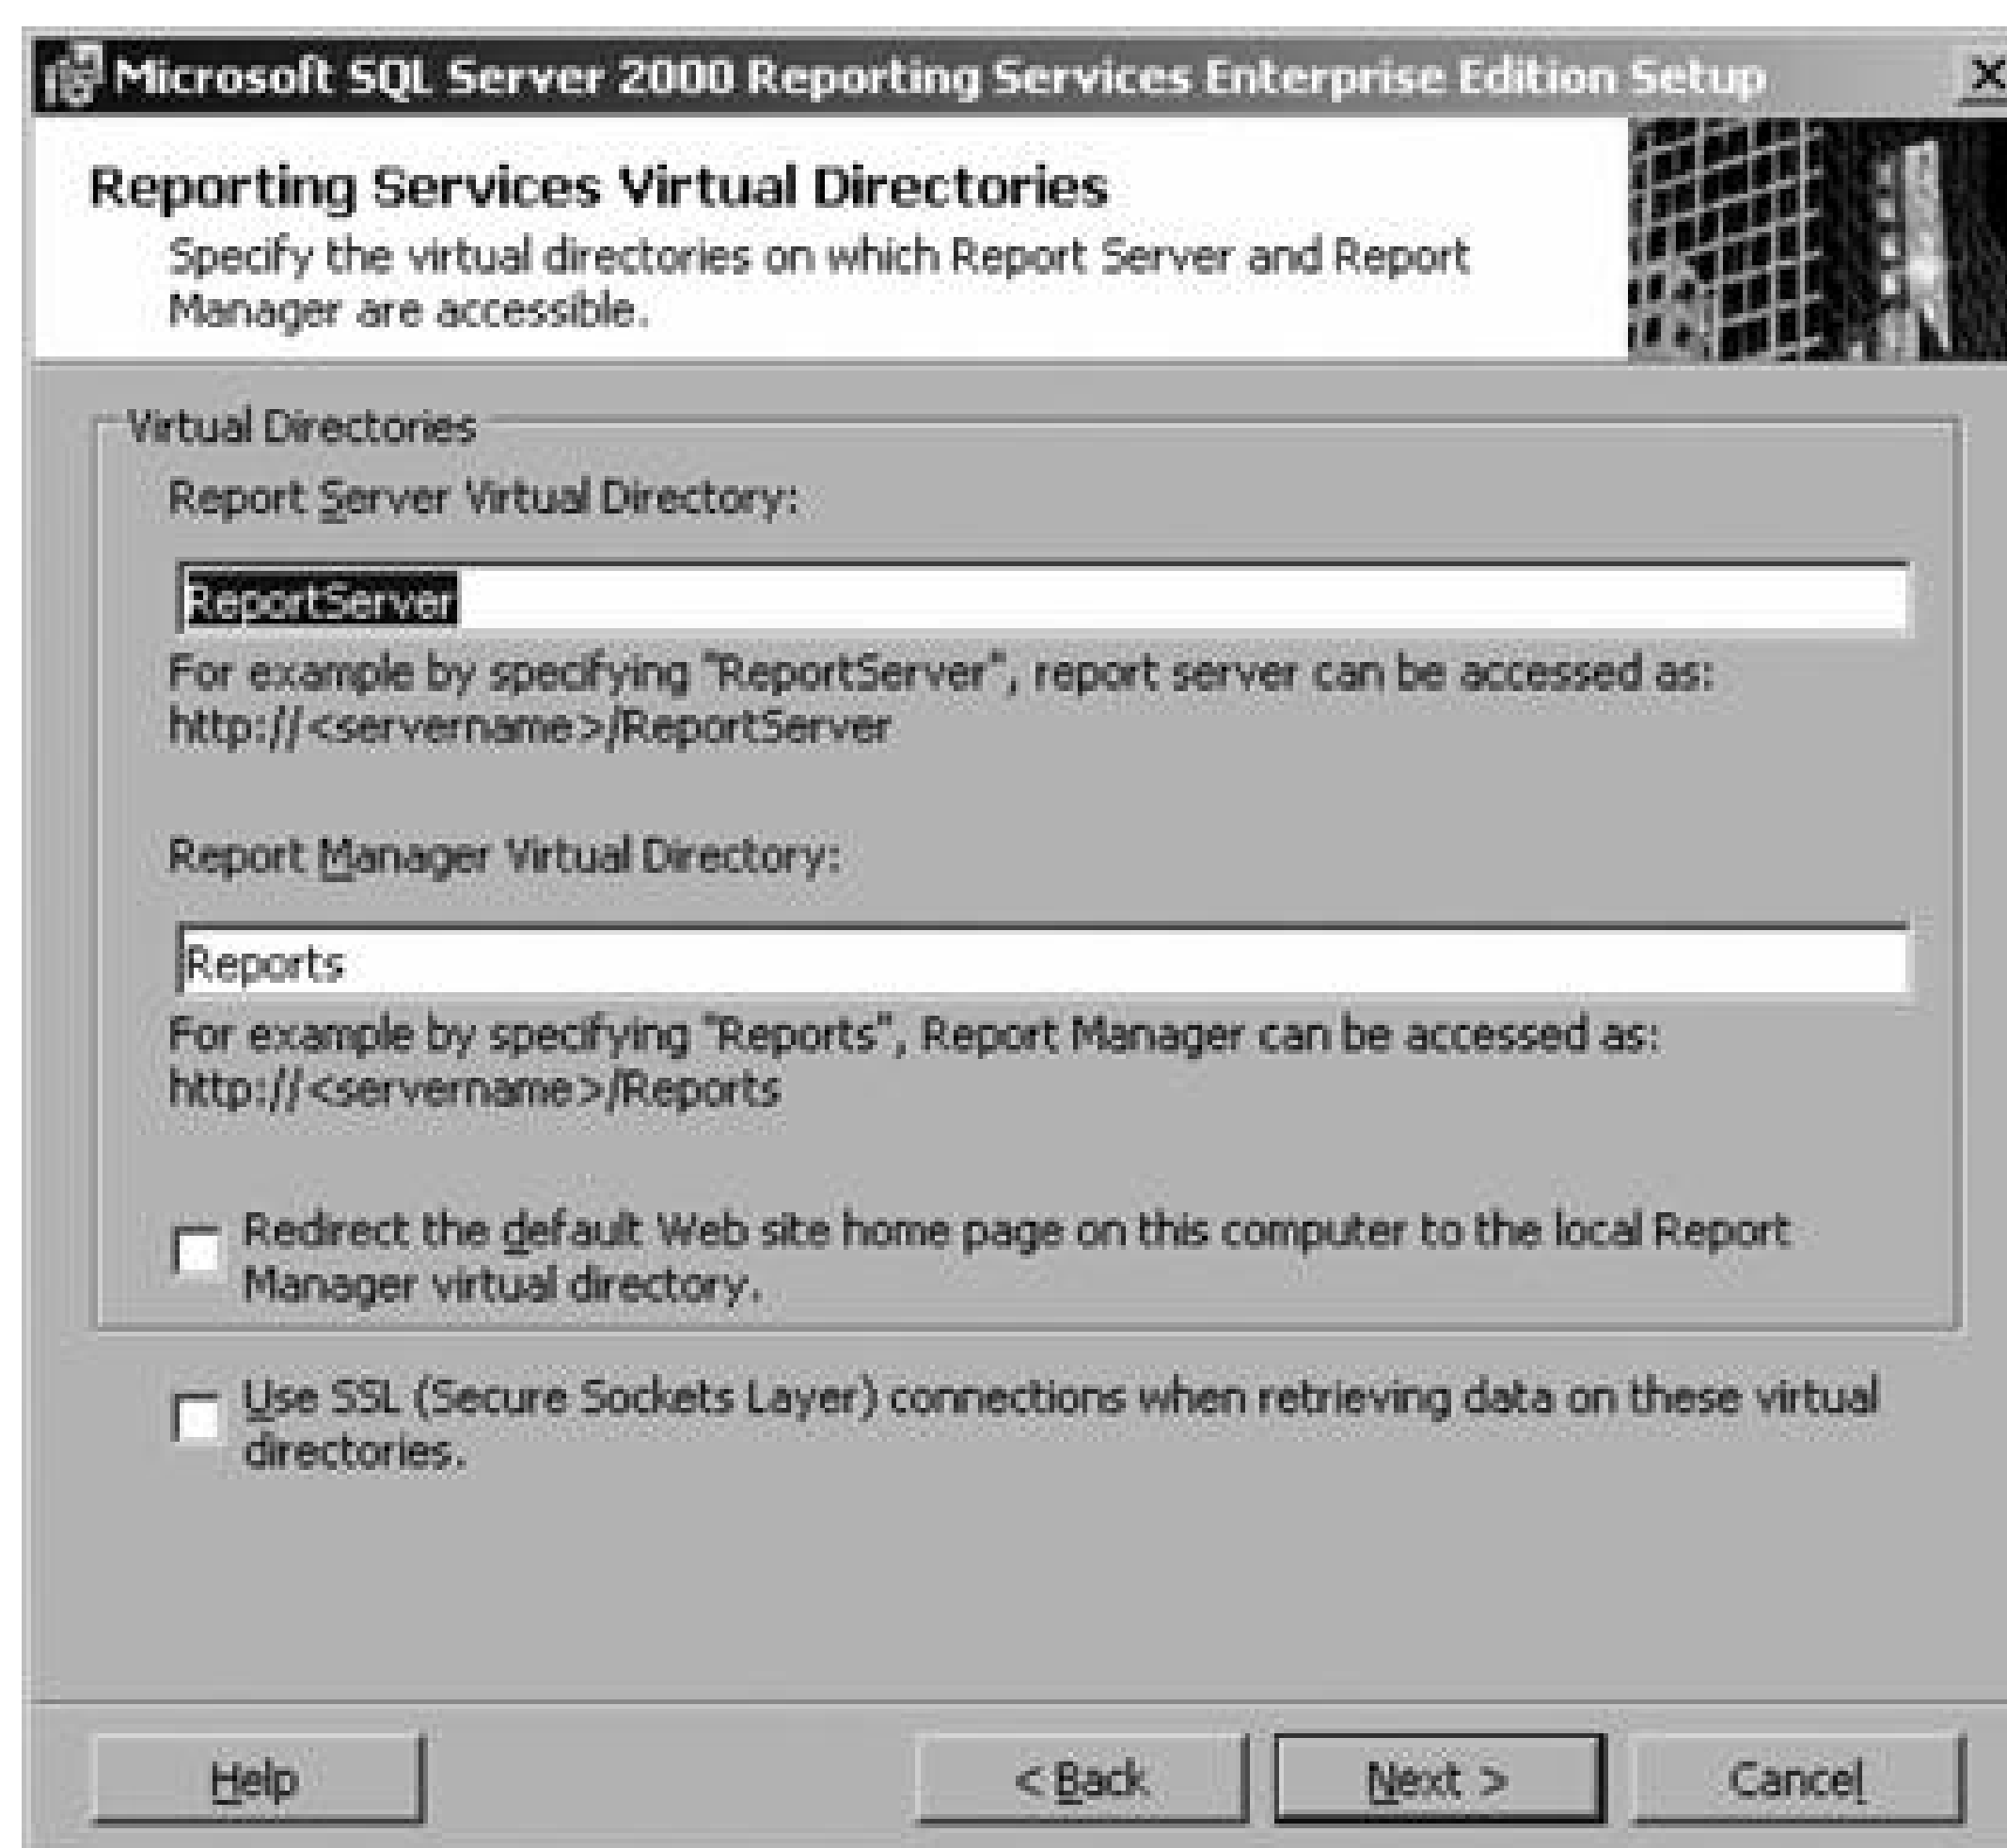
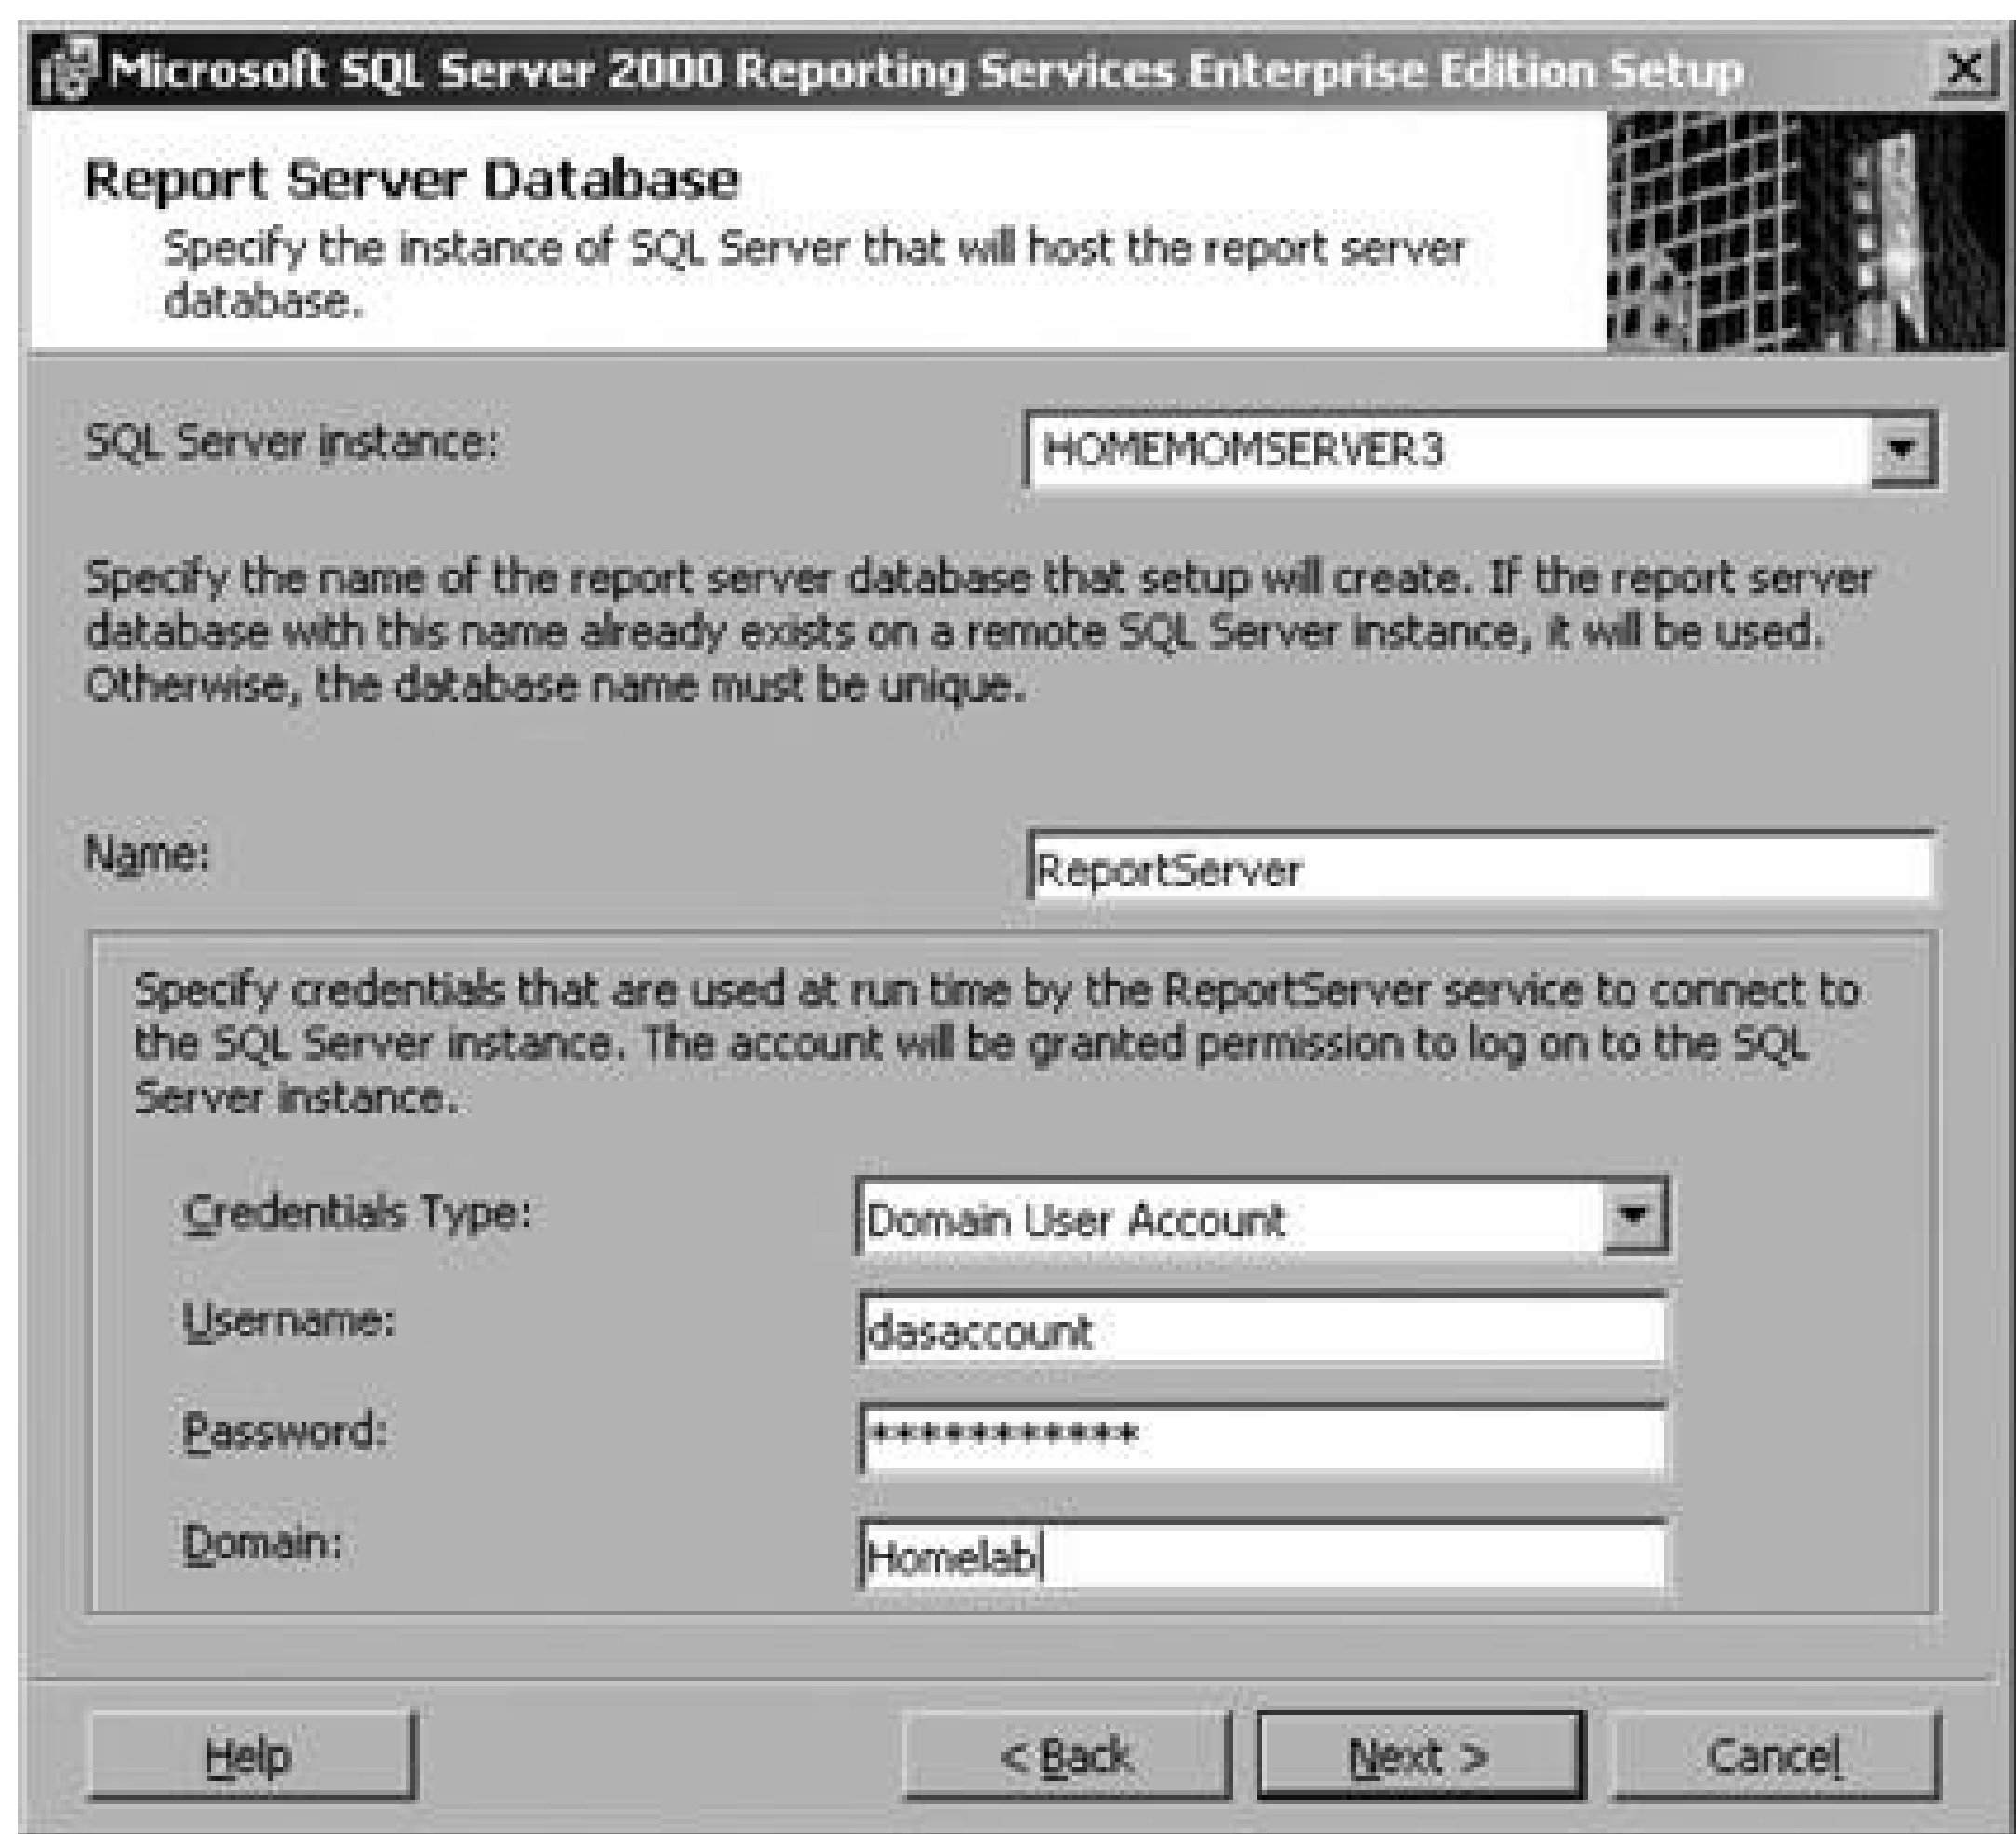


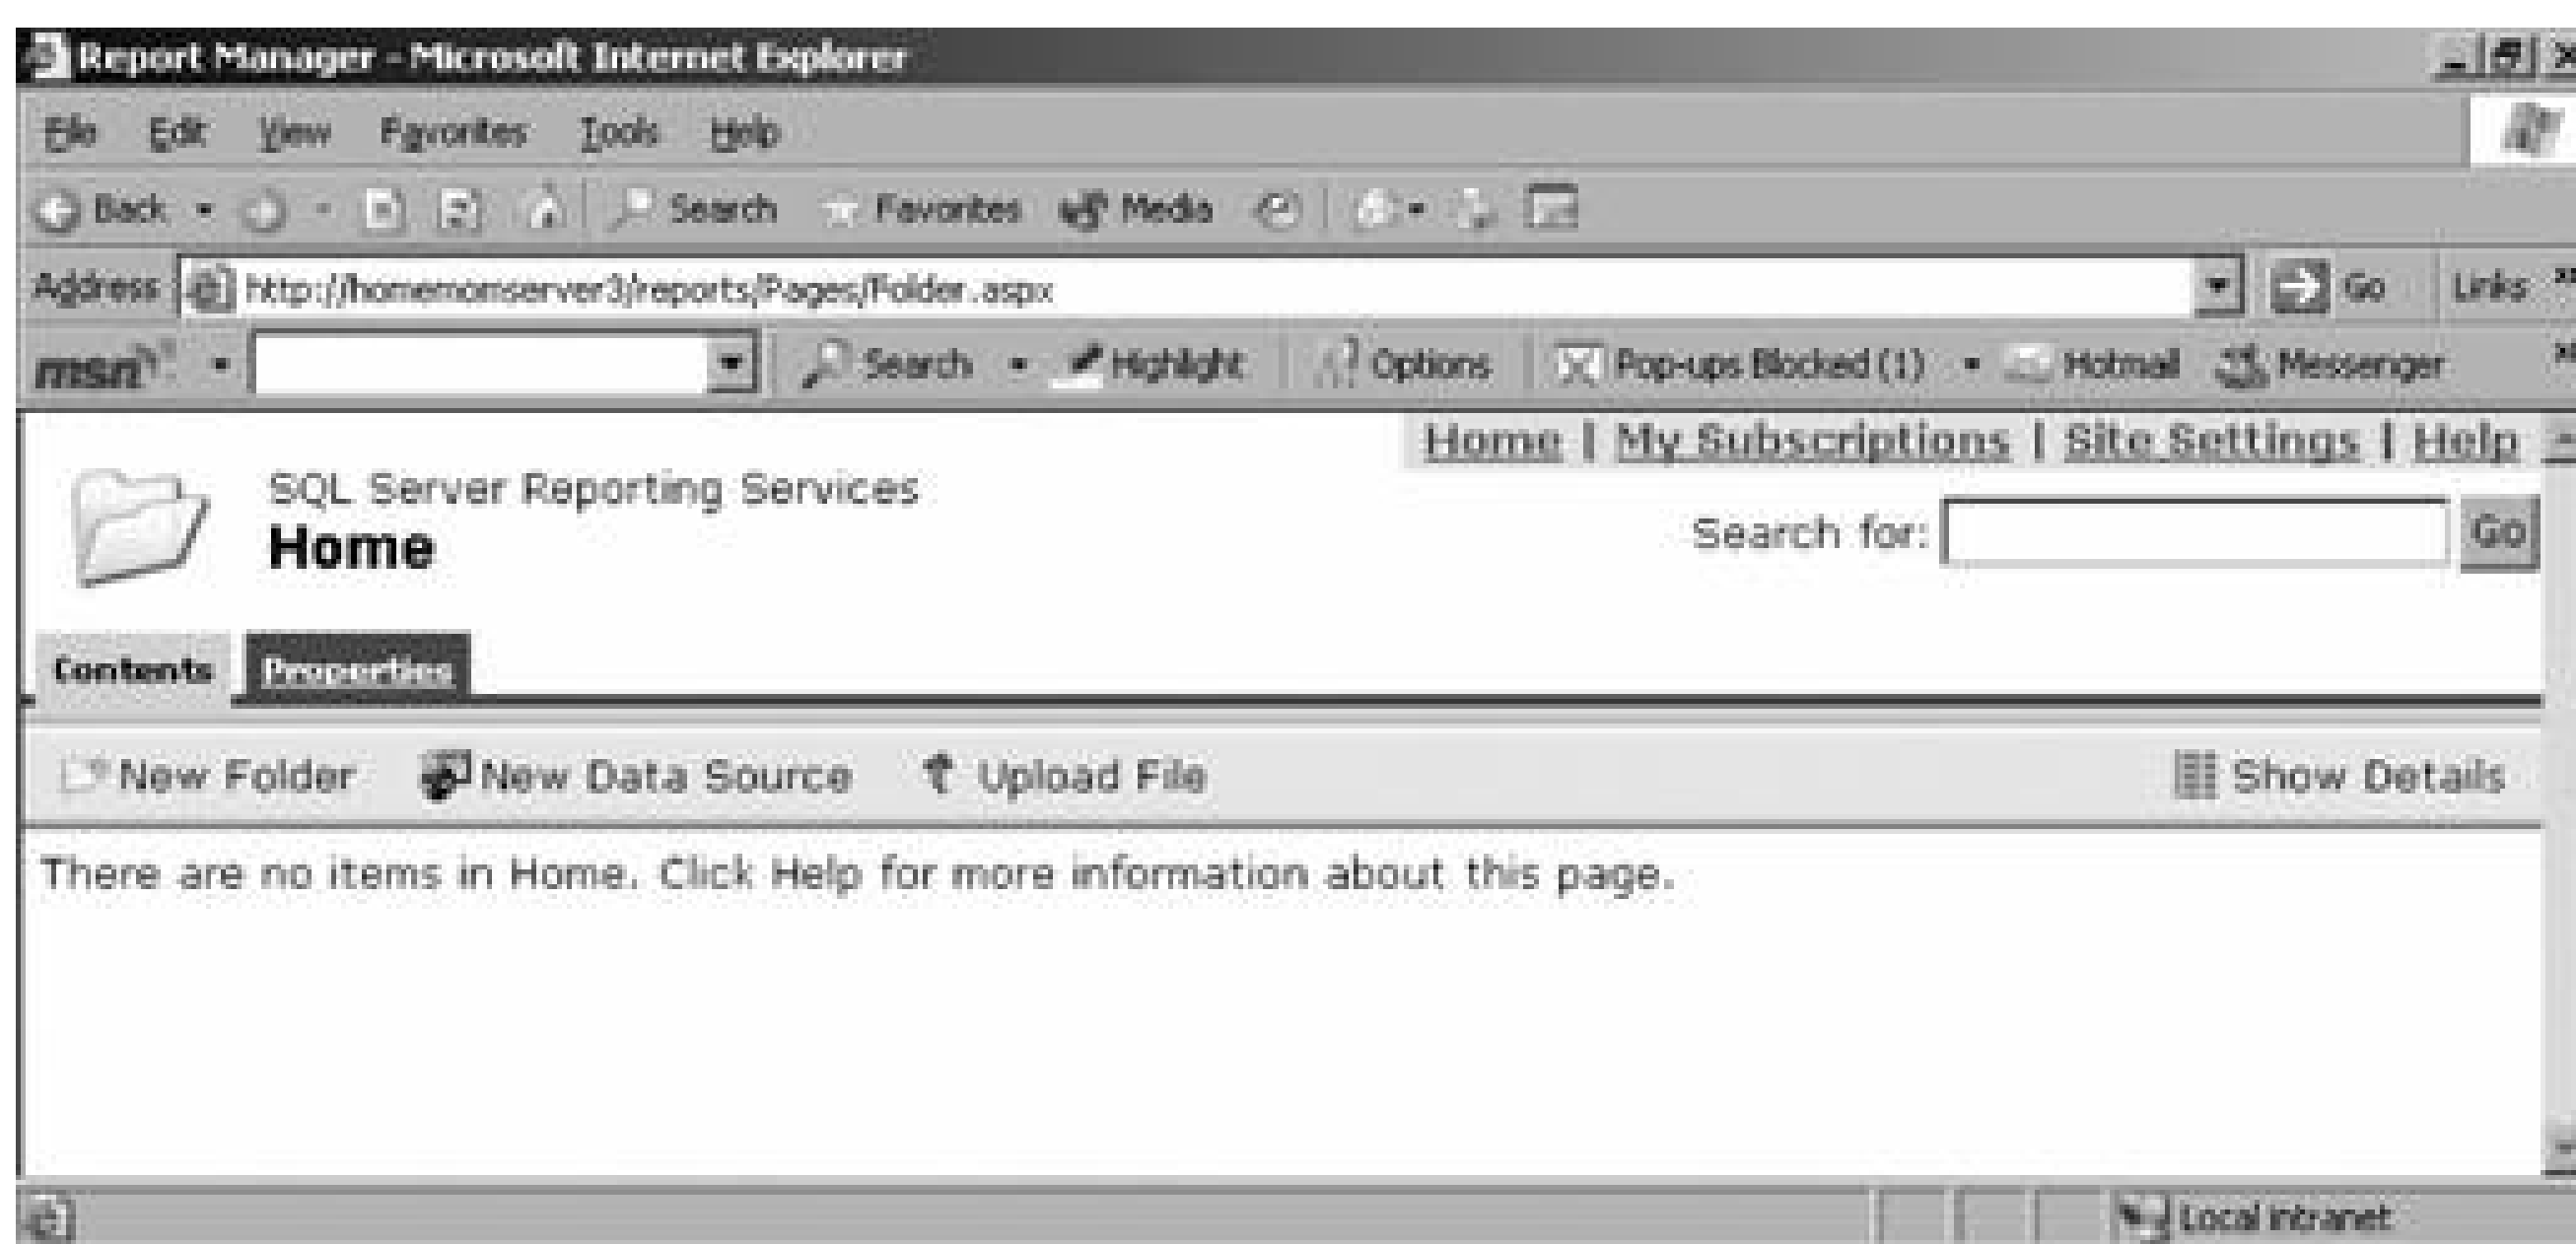
Figure 8-5. Specifying the SQL instance, the database name, and the credentials for the report sever database



On the next page, you designate the SMTP server to be used for forwarding subscribed reports as well as the From address. This can be any server that is running SMTP and is configured to perform SMTP relay. Exchange servers can do this just fine. For this all-in-one installation, the SMTP component of IIS on *homemomserver3* is installed (the reporting server) and the From address was set to [reportmaster@homelab.lab](mailto:reportmaster@homelab.lab). The remaining pages cover choices on licensing (per seat or per processor) and the usual Ready to Install and Completed Successfully pages.

Once setup completes successfully, launch your browser and open <http://<servername>/reports> to confirm that the report manager is accessible. Since connectivity is just being tested, there won't be any reports or folders here yet. [Figure 8-6](#) shows the Report Manager page on *homemomserver3* at the end of SRS installation.

Figure 8-6. The empty, but functioning, reporting services site



At this point, don't forget to install SRS SP1 on your reporting server. After installing SP1, you can confirm that the patch applied correctly by looking at the SRS version number. To see this, browse to <http://<servername>/reportserver>. If the upgrade was applied successfully, version number 8.00.878.00 will be displayed. If not, the RTM version number 8.00.0743.00 is displayed.

## 8.1.2. MOM 2005 Reporting

The installation process instantiates the SystemCenterReporting database and creates the necessary links between the OnePoint, SystemCenterReporting, and Report Server databases. Because it touches two machines (at a minimum) and three different databases, you should perform the install using a domain administrator account. You will have to provide several pieces of information for the setup process to complete successfully. The steps in the installation process are summarized in the following list.

1. To start the installation, run the setup from the MOM 2005 CD and select the Install Microsoft Operations Manager 2005 Reporting link on the setup Tasks tab.
2. Proceed through the Welcome page.
3. Accept the terms of the EULA and proceed through.
4. On the Registration information page, enter the username and organization information and proceed through.
5. On the Destination folder page, specify where you set the installation path for the MOM 2005 Reporting to be installed. By default this is *C:\Program Files\Microsoft Systems Center Reporting*. In preparation for the merging of data warehouse data from MOM and Microsoft Systems Management Sever (SMS) into a single product called Systems Center, the MOM data warehouse and reporting solution have already been branded as System Center (<http://www.microsoft.com/windowsserversystem/systemcenter/default.aspx>).



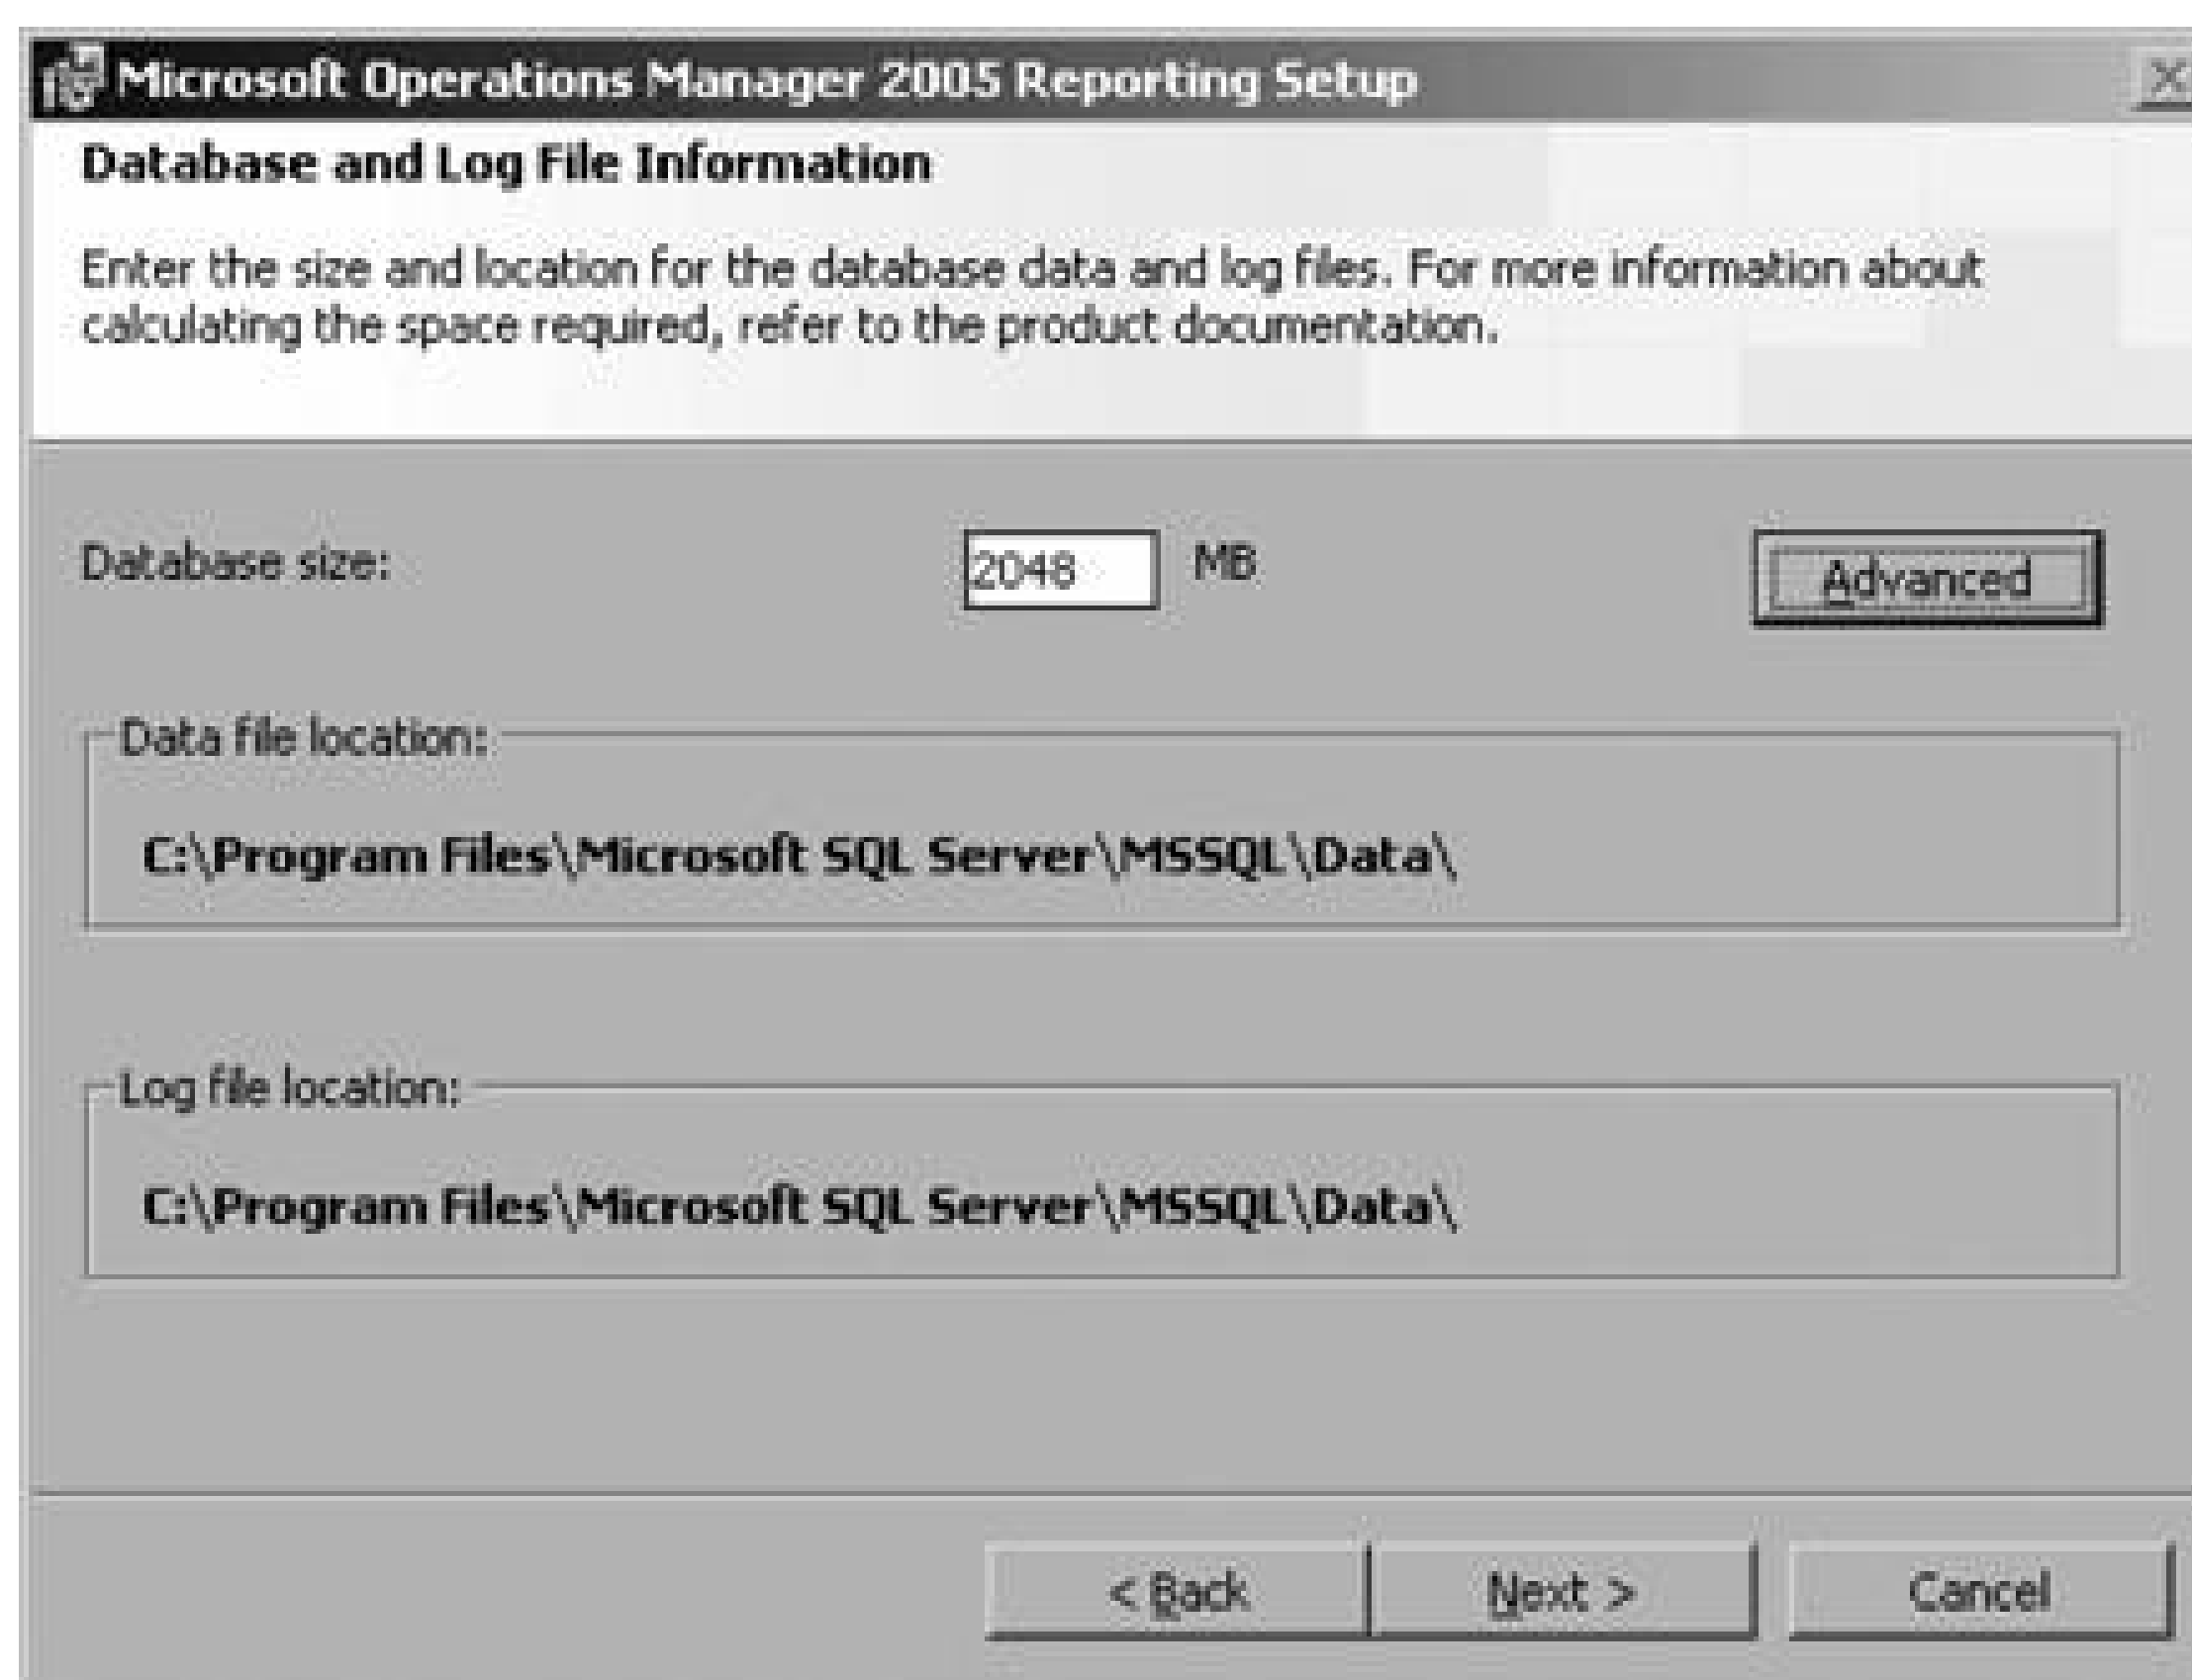
6. On the SQL Server Reporting Services Server page in [Figure 8-7](#), you need to indicate the name of the Reporting Services server. In this case it is *homemomserver3*. Here you also instruct the setup process to either auto-detect the report server virtual directories (Reports and ReportServer) or if you have customized these values during the SRS installation, enter the correct paths here.

Figure 8-7. Identifying the SQL Reporting server to the MOM Reporting Service installation process




7. The Prerequisite Check page checks for the correct versions of OS, SQL, SRS; if there is sufficient processor, memory, and disk space; and for other items like MDAC version and presence of the correct .NET framework version. It always generates a log that you can view regardless of the pass/fail status. If the prerequisite checker fails, the log will tell you specifically at what point failure occurred.
8. On the MOM Database Server instance page, enter the name of the server and the SQL instance (if there is one) where the OnePoint database is. This database will be the source for all data that flows via DTS to the SystemCenterReporting database. Again, since this is an all-in-one configuration, *homemomserver3* is named.
9. On the SQL Server Database instance page, indicate the SQL server and instance where the SystemCenterReporting database is to be created (*homemomserver3*).
10. On the Database and Log File Information page in [Figure 8-8](#), refer back to the data warehouse sizing calculations performed in [Chapter 2](#). The size of the Systems Center Reporting database is set to 2,048 MB.

Figure 8-8. Configuring the SystemCenterReporting database



11. Next is the Data Transfer Task Account page. [Chapter 7](#) introduced the DTS package that is responsible for copying and transforming data from the OnePoint database to the SystemsCenterReporting database. This task must have a security context to run in. Since the dasaccount already has access to all the databases concerned (OnePoint and SystemCenterReporting), this is the logical choice. You are prompted for the account name, password, and the domain name if it is a domain account, see [Figure 8-9](#). A local account can also be used.

Figure 8-9. Identifying the account to be used for executing the DTS package

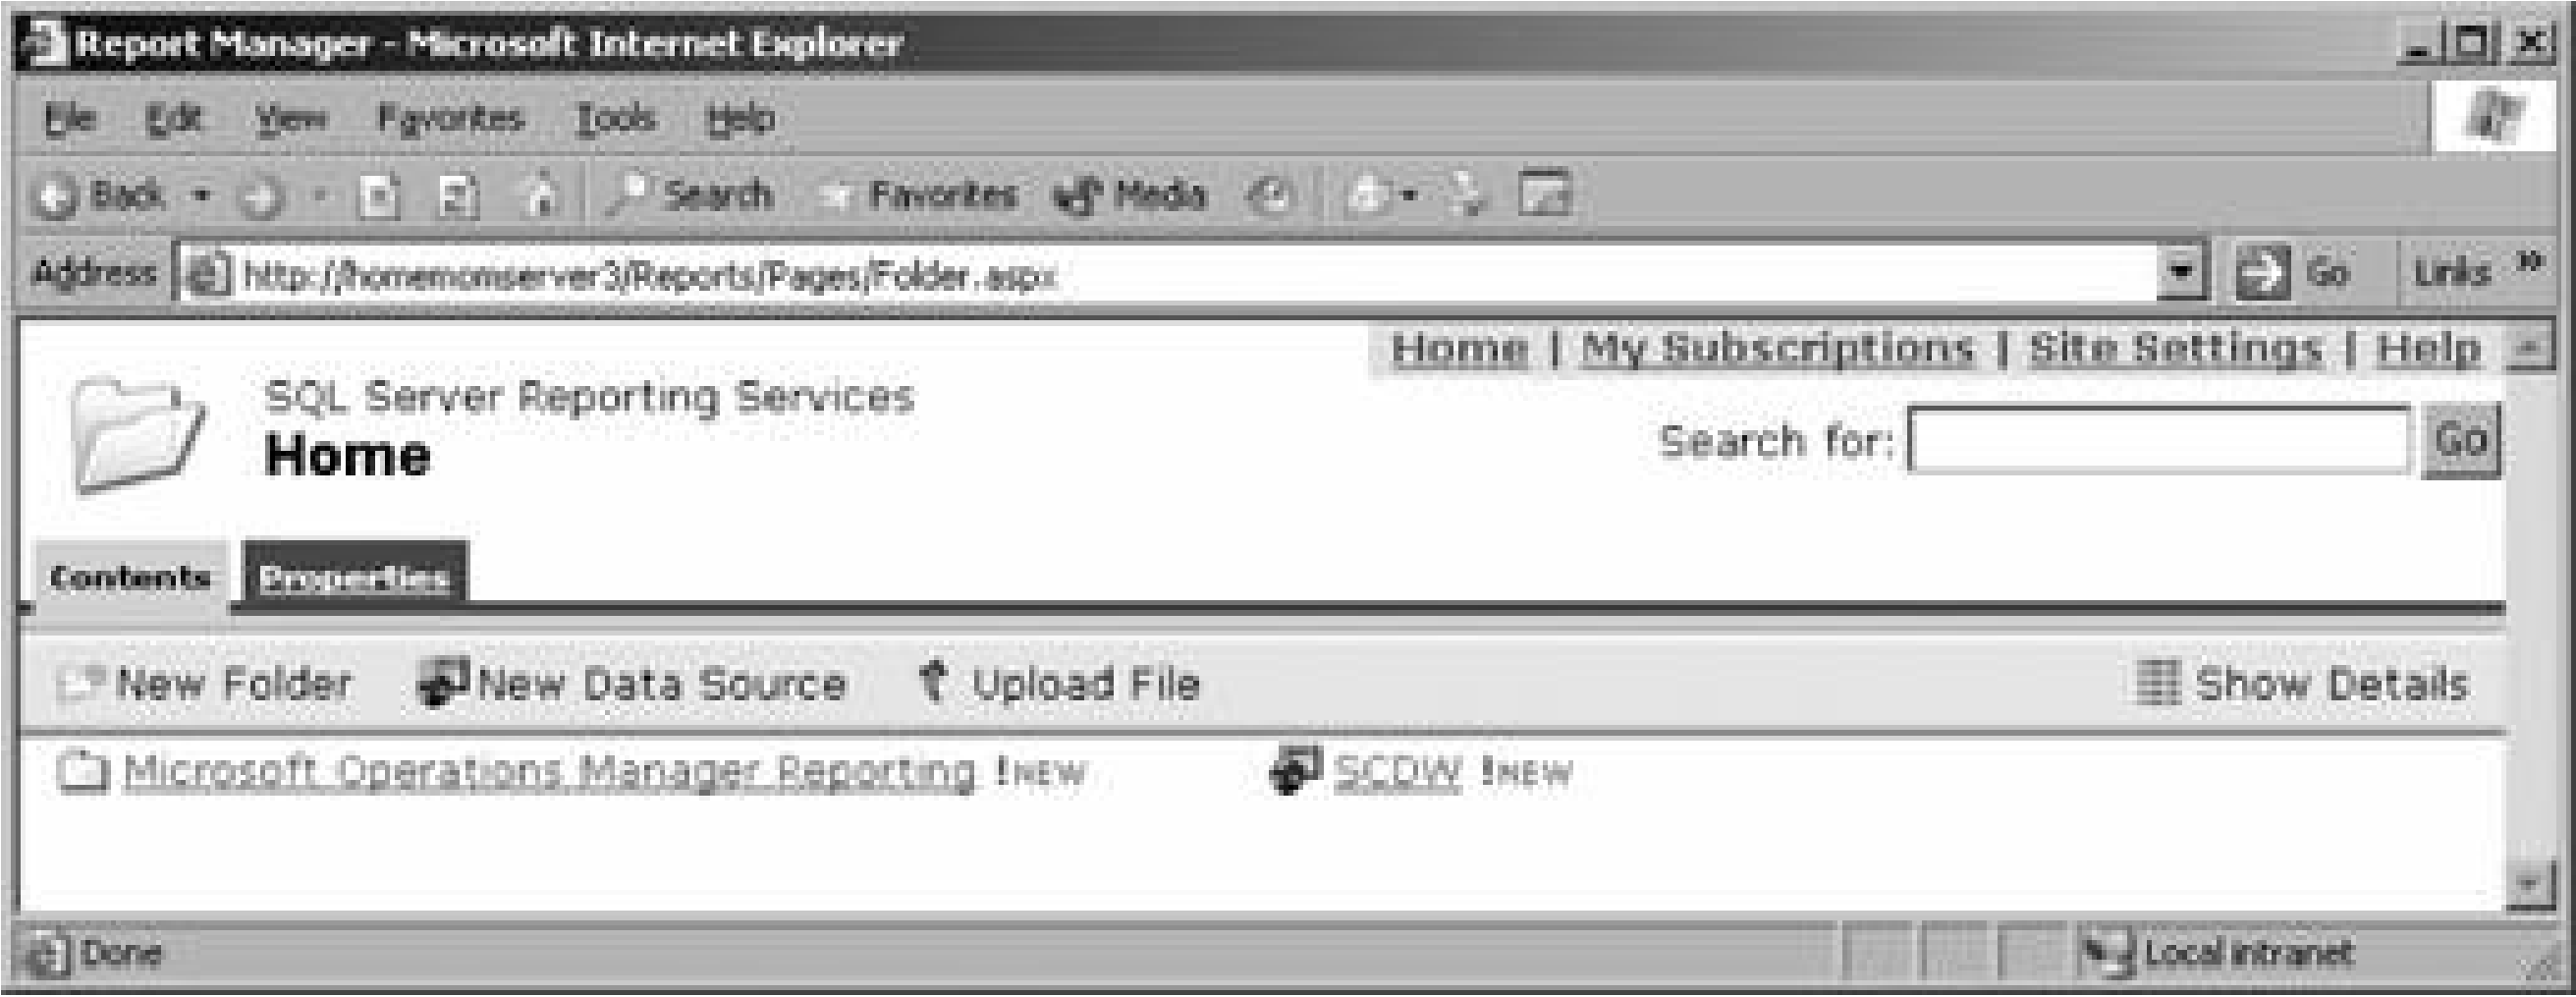


12. The Reporting User Account page is next. Just as the reporting service needs a set of credential to connect to the ReportServer database, a set of credentials is required to connect to the SystemCenterReporting database. The homelab\dasaccount is again used here. This page is almost identical to the Data Transfer Task Account page.
13. On the Operational Data Reports Settings page you can choose to have operational reports forwarded to Microsoft. There is no direct benefit to doing so other than the satisfaction of knowing that information about your MOM environment will be used to improve future releases of MOM.
14. Proceed through the Ready to Install page.
15. On the Installation Completed successfully page, selecting the "Transfer data from the MOM database server now" option invokes the scheduled DTS task on the reporting server. Select this to have the available operational data copy over, and the reports will immediately contain something instead of waiting for the 1:00 a.m. transfer. This can take a long time if it is the first transfer and there is a large amount of data. The transfer could fail, but there is a fix for this in MOM 2005 SP1.

Upon successful completion of the setup, browse to the Report Manager Home page (<http://<servername>/reports>) to confirm that the Microsoft Operations Manager Reporting folder and the SCDW data source objects have been created. The Report Manager page should now look like [Figure 8-10](#).

Figure 8-10. The SQL Reporting Services Report Manager Home page after successful installation of MOM 2005 Reporting





← PREV

## 8.2. Administering MOM 2005 Reporting

If you skipped to this section from the beginning of [Chapter 4](#) to install MOM 2005 Reporting before importing your management packs and reports at the same time, congratulations! You have saved yourself some re-work. Now that MOM 2005 Reporting has been installed and is functioning correctly import the management packs with the processing rules and reports.

The report import process brings the report definition into the ReportServer database and creates the correct fields in the SystemCenterReporting database so the operational data is received when it comes over from OnePoint. It also creates a folder hierarchy that is seen in Report Manager in the Home folder as a peer-level folder to the Microsoft Operations Manager Reporting folder in [Figure 8-10](#).

SRS is a product in its own right. Fortunately, just as in the case of SQL 2000, you do not need to be fully versed in SRS as a MOM administrator. In fact, the number and complexity of administrative tasks that must be performed to provide basic functionality are few and are really quite simple. There are a number of additional tasks that you will probably need to do, so this next section covers those as well. Once you get beyond enabling rudimentary functionality, how much more you provide depends on your business needs.

### 8.2.1. Required Administrative Tasks

By default, at the end of a successful installation and DTS transfer of data, the only group that can access the available reports is the Local Administrators group, which is secure but not very functional. To permit your users to access reports, each user's AD account must be associated with an SRS *item-level role*. Item-level roles are used to control what actions can be performed by a designee on SRS folders and their contents. The association can be created directly between the account and the role or between an AD security group in which the account has membership and a role. In SRS, a role is defined by the tasks it can perform in the context of the web site. Roles do not exist outside of the Report Manager web site.

If you are familiar with Windows Sharepoint Services, an SRS role is identical to a site group. In fact, much of the functionality in the Report Manager site and the security model is based on functionality in Windows Sharepoint Services.

In SRS, two groups of tasks are used to define roles to which user accounts are assigned. Some of the roles include System Administrator, System User, Content Manager, My Reports, and Publisher. There are item-level tasks and system-level tasks. [Tables 8-1](#) and [8-2](#) provide the details of the different tasks and the roles that they have been used to define. This is viewable in Report Manager under Site Settings → Configure System - Role definitions.

Table 8-1. System-level role definitions

Role	Task	Description
This role is not assigned by default	Generate events	Provide an application with the ability to generate events within the report server namespace
System Administrator	Manage jobs	View and cancel running jobs
System Administrator	Manage report server properties	View and modify properties that apply to the report server and to items managed by the report server
System Administrator	Manage report server security	View and modify system-wide role assignments
System Administrator	Manage roles	Create, view, modify, and delete role definitions
System Administrator	Manage shared schedules	Create, view, modify, and delete shared schedules used to run or refresh a report
System User	View report server properties	View properties that apply to the report server
System User	View shared schedules	View a predefined schedule that has been made available for general use

Table 8-2. Item-level role definitions

Role	Task	Description
Content Manager, My Reports, Publisher	Create linked reports	Create linked reports and publish them to a report server folder
Content Manager	Manage all subscriptions	View, modify, and delete any subscriptions regardless of who owns the subscription
Content Manager, My Reports, Publisher	Manager data sources	Create and delete shared data source items; modify data source properties
Content Manager, My Reports, Publisher	Manage folders	Create, view, and delete folders; view and modify folder properties
Browser, Content Manager, My Reports	Manage individual subscriptions	Each user can create, view, modify, and delete subscriptions that he owns
Content Manager, My Reports	Manage report history	Create, view, and delete report history snapshots; modify report history properties
Content Manager, My Reports, Publisher	Manage reports	Create, view, and delete reports; modify report properties
Content Manager, My Reports, Publisher	Manage resources	Create, modify, and delete resources; view and modify resource properties



Role	Task	Description
Content Manager	Set security for individual items	View and modify security settings for reports, folders, resources, and shared data sources
Content Manager, My Reports	View data sources	View shared data source items in the folder hierarchy; view data source properties
Browser, Content Manager, My Reports	View folders	View folder items in the folder hierarchy; view folder properties
Browser, Content Manager, My Reports	View reports	View reports and linked reports in the folder hierarchy; view report history snapshots and report properties
Brower, Content Manager, My Reports	View resources	View resources in the folder hierarchy; view resource properties

The best way to grant users permission to reports is to add their AD accounts to the domain-level SCDW readers group that you created in [Chapter 2](#). Then, add that domain group to the local SCDW readers group on the reporting server. This grants those users the necessary SQL permissions to use SRS. The last step is to assign the local SCDW readers group to an item-level role. To do this, log onto Report Manager as a local administrator, select the Properties tab of the Home folder (point 1 in [Figure 8-11](#)) and then select the New Role Assignment link (point 2 in [Figure 8-11](#)).

This chapter emphasizes using the SCDW readers local group for assigning users to roles, but you can create associations between any domain, local user, or group and any role. In addition, you can create new roles simply by creating a grouping of items or system-level tasks into a named role in the site settings page under the configure item-level or system-level role definitions links. What you can't do is create new item- and system-level tasks.

Figure 8-11. Creating a new role assignment

Out of the item-level roles defined in [Table 8-2](#), you must pick one to assign the bulk of your users to. I strongly recommend that you don't assign the average MOM report consumer to either the Content Manager role or the Publisher role. There is simply no need for this group of users to have that much power in your MOM 2005 reporting solution. That leaves the Browser and My Reports roles, and there are significant differences between the two.

### *Browser*

Users that are assigned this role will be able to execute all reports and manipulate the parameters of those reports. They can also create a subscription to that report for themselves.

### *My Reports*

Think of this role as the power user role for SRS. It allows the user to do everything the Browser role can, plus create and maintain persistent customized reports in their own My Reports folder.

As shown in [Figure 8-12](#):

1. Enter the `<machine>/sc dw`, in this case `homemomserver3/sc dw reader` (point 1).
2. Select the My Reports role (point 2).
3. Click OK (point 3).

All user accounts can now perform the tasks of the My Reports role across all folders in the hierarchy. Just as with a filesystem folder hierarchy, all permissions flow down from the top level unless you specifically configure a child folder not to inherit the permissions from its parent folder. This also holds true for items within a folder.

Remember that a report consists of a data source definition, a query, and a layout or report definition. When a user clicks on a report to view it, depending on the report, certain selections will have to be made to fill out the query's value fields. For example, [Figure 8-13](#) shows the Agent Configuration report. To navigate to this from the Home folder, select Microsoft Operations Manager Reporting      Microsoft Operations Manager      Agent Configuration.

Figure 8-12. Completing the initial role assignment for your users

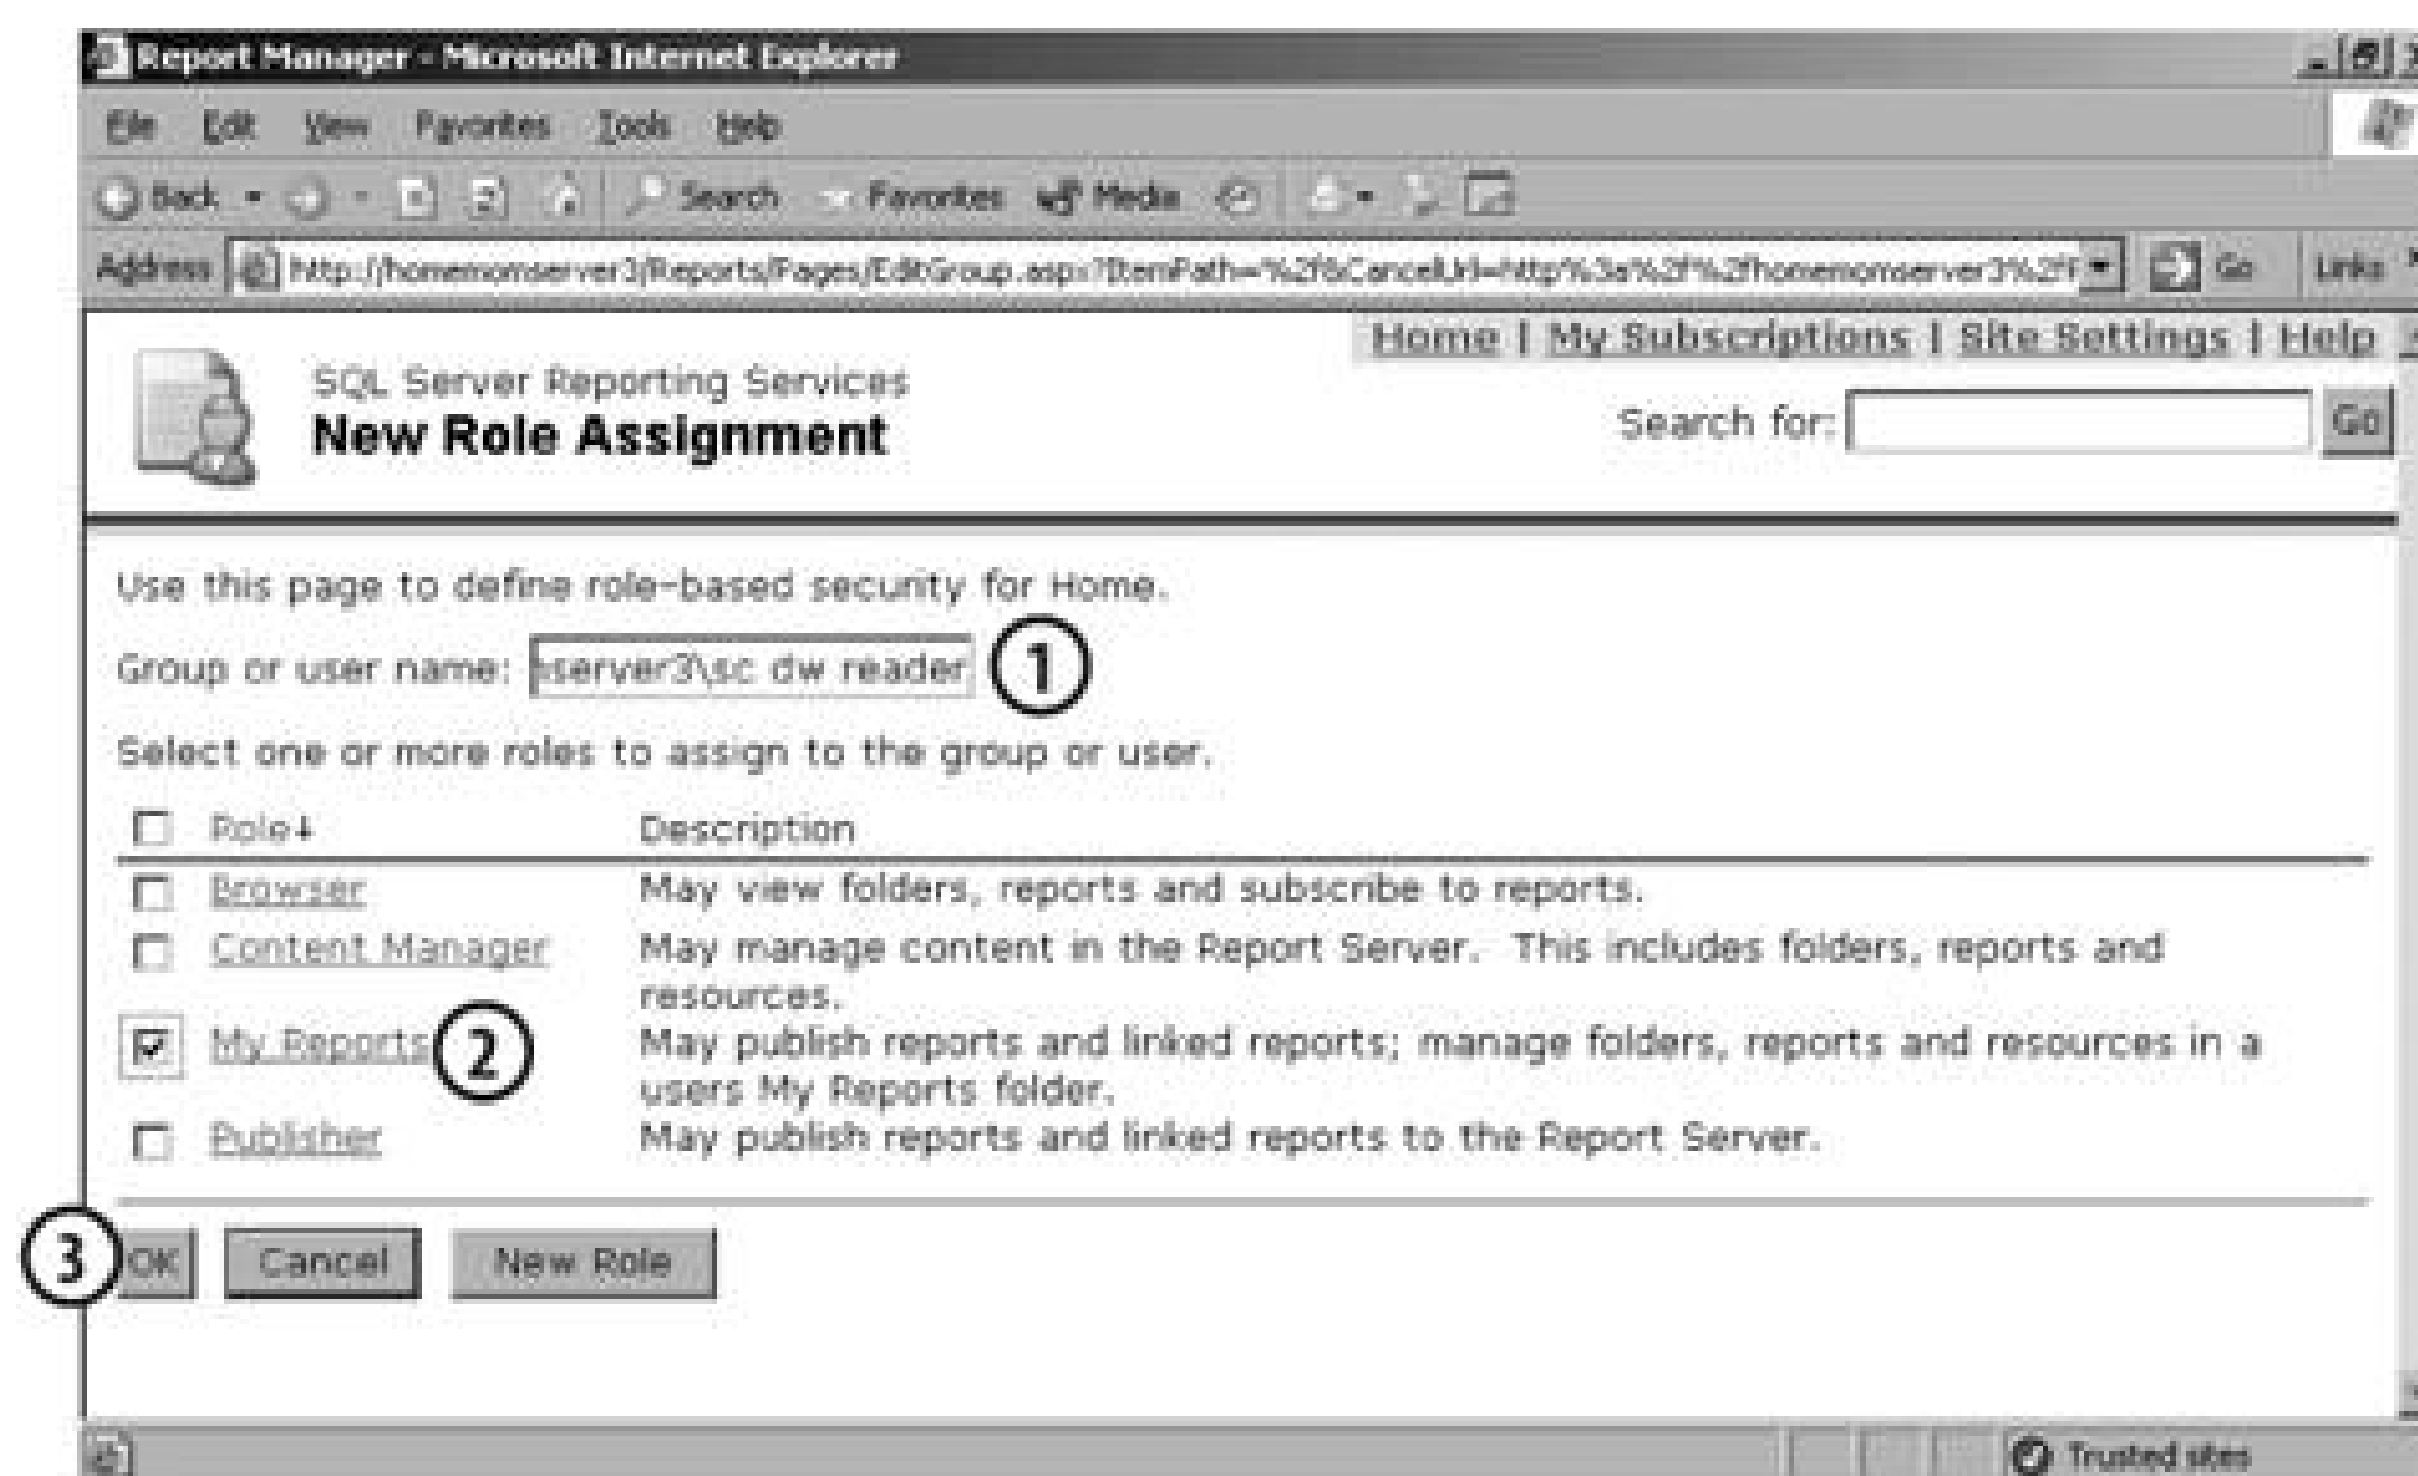


Figure 8-13. Selecting the report parameters

Before the query behind this report can run (*execute*) against the SystemsCenterReporting database, values (*parameters*) must be supplied for the management group, agent, and primary server fields. Defaults have been supplied for the Sort Order and Sort By fields, but as you can see by the drop-down boxes for those parameters, they can be changed. For the report in [Figure 8-13](#):

1. Select LKFProdMG Management Group for the Management Group field (point 1).
2. Leave <ALL> for the agent field.
3. Select *homemomserver* for the primary server field and then click View Report (point 2).



This will then execute the report and display the results in the pane beneath the parameters pane.

Users will have their own preferences for report parameters based on their needs. Once a user in the Browser role finds those preferences, she may *subscribe* to that report (point 3 in [Figure 8-13](#)). A report subscription requires the definition of all the same report parameters, as well as the definition of the desired format of the report, the delivery mechanism, schedule, and credentials if necessary. SRS will then render the report, using the defined parameters, in the defined format, and deliver it via the selected mechanism (to a mailbox or a file share). You cannot manipulate the parameters of the delivered report.

The Browser role is good for users who don't want to navigate the Report Manager interface on a daily basis and set report parameters. Typically, this type of user just wants the report to appear in his email inbox or in the same filesystem folder on a regular basis with no questions asked. If this description fits the bulk of your user population, then assign the Browser role to the SCDW readers group.

If your reporting business requirements are met by assigning your users the Browser role, then you have completed all of the MOM 2005 Reporting administrative tasks that must be done. However, this is unlikely to be the case. Additional business requirements, as demonstrated in the Leaky Faucet environment, include ensuring that certain reports are not publicly accessible and that the executives can get access to reports with minimal interaction with IT. To design a solution to meet these needs, your users must be assigned to other roles.

## 8.2.2. Additional Administrative Tasks

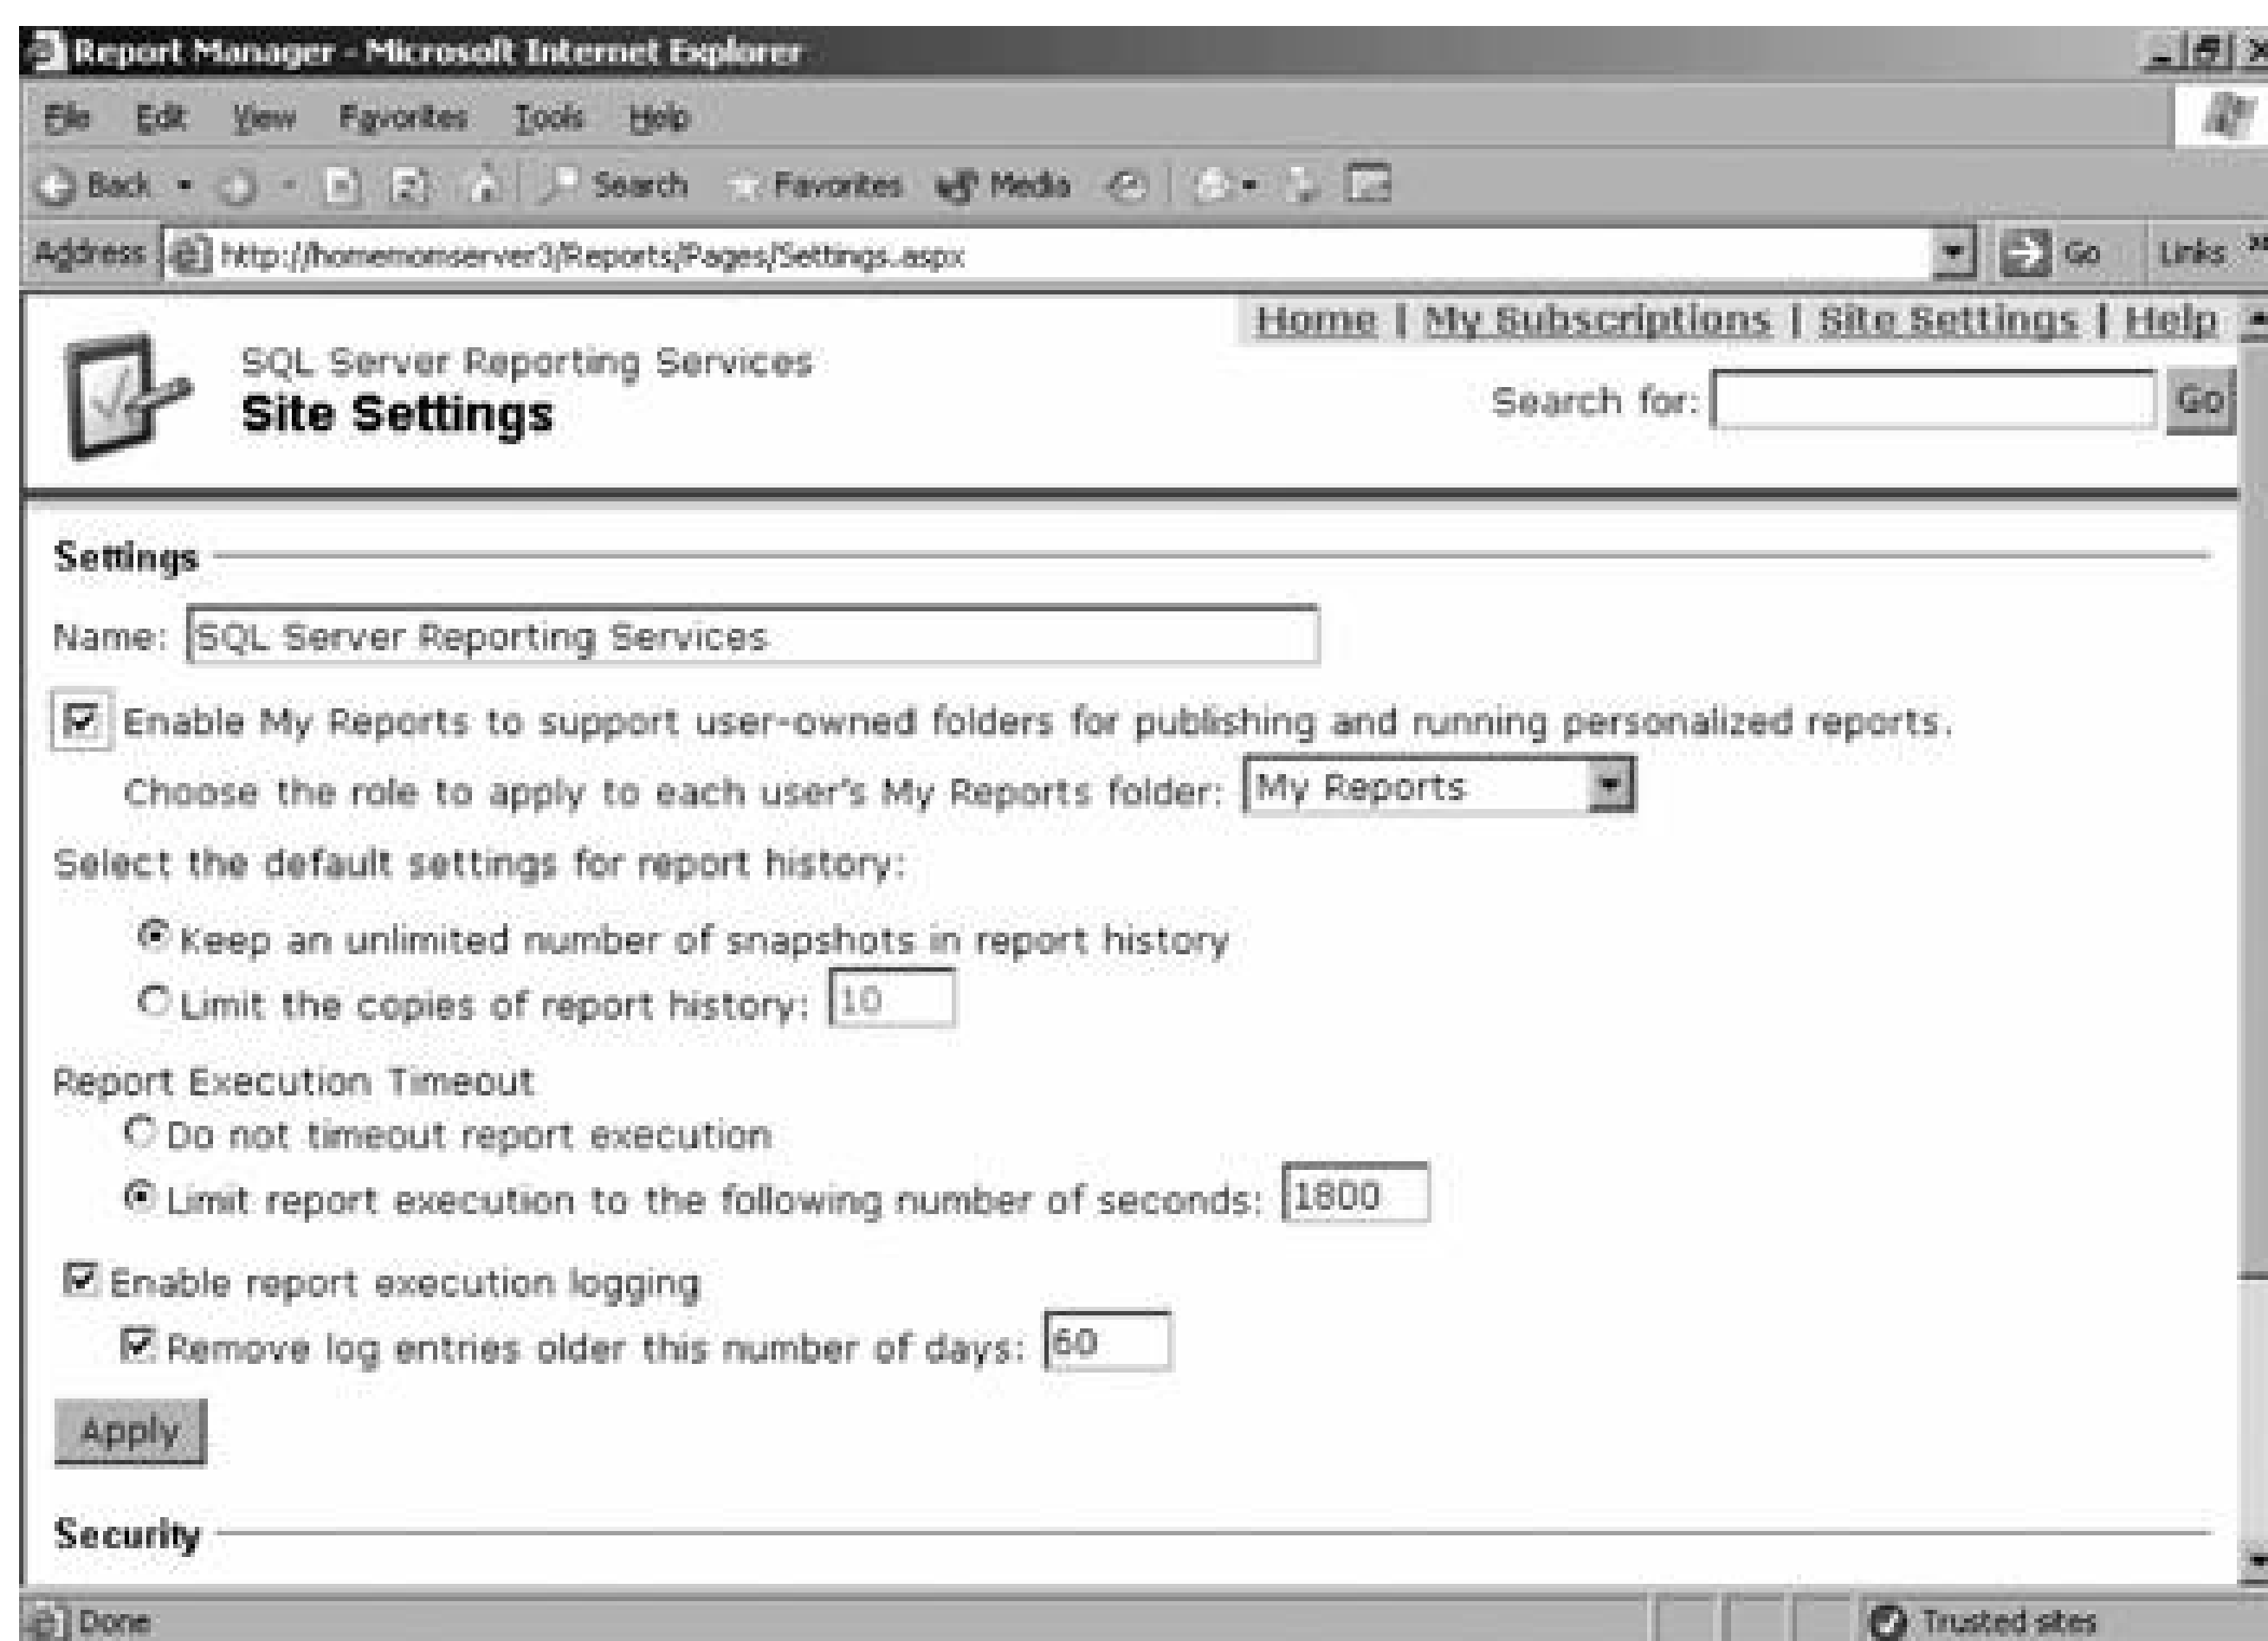
You may want to assign the My Reports role to the bulk of your users. What distinguishes this role from the Browser role is its ability to create folders and linked reports. These two abilities, along with enabling the My Reports feature, can be used together to deliver a good deal of customization in your reporting solution.

### 8.2.2.1. Enabling the My Reports feature

When the My Reports feature is enabled, a new folder (titled My Reports) will appear in the Home folder for every user that has access to MOM 2005 reporting. The idea behind the My Reports folder is to give users a private place to store reports that they are working with. This folder is not viewable by other users, but administrators will have a Users folder created on their home page that will allow them full access to all users' My Reports folders and their contents. To enable the My Reports feature:

1. Open Report Manager as an administrator of the Reporting Services machine.
2. Select the Site Settings link in the top right-hand corner of the page. This takes you to the Site Settings page (see [Figure 8-14](#)).
3. Select the checkbox to enable My Reports, leave the default role setting, and click Apply at the bottom of the page.

Figure 8-14. Enabling the My Reports feature



To take advantage of the My Reports folder, your users must be assigned to the My Reports role as shown in [Figure 8-12](#). Again, you don't have to use the SCDW readers group for this role assignment if you have already created an association between it and another role. Simply create another group, add your users, and create the association.

Each user's My Reports folder can be used similarly to the My Views container in the Operator console. In the My Views container, you create and store customized versions of view objects; in the My Reports folder, you store customized versions of existing reports, called *linked reports*. However, with the My Reports folder you can also upload supporting files, also called *resources*. This is another feature taken directly from Windows SharePoint Services document libraries.


### 8.2.2.2. Creating a linked report in a My Reports folder

A linked report is a personal version of a public report. In the linked version, which is best stored in the My Reports folder, you can define the parameters for the report. This specific report configuration is then retained in the linked report so it becomes the default configuration for that linked report. The advantage of this over the subscription report version is that you can still interact with the report because it has not been rendered into a static format.

Linked reports are created from the Properties tab of an existing report. In [Figure 8-15](#), browse to the Home Microsoft Operations Manager Reporting Microsoft Windows Base OS folder

Operating System configuration report and select the Properties tab. Additional entries (point 1 in [Figure 8-15](#)) are available here that would not be available via the Browser role. The security link is missing because the Content Manager role is required to access that. Security for this individual report is configured at that point. Click the Create Linked Report link to start the creation process (point 2 in [Figure 8-15](#)).

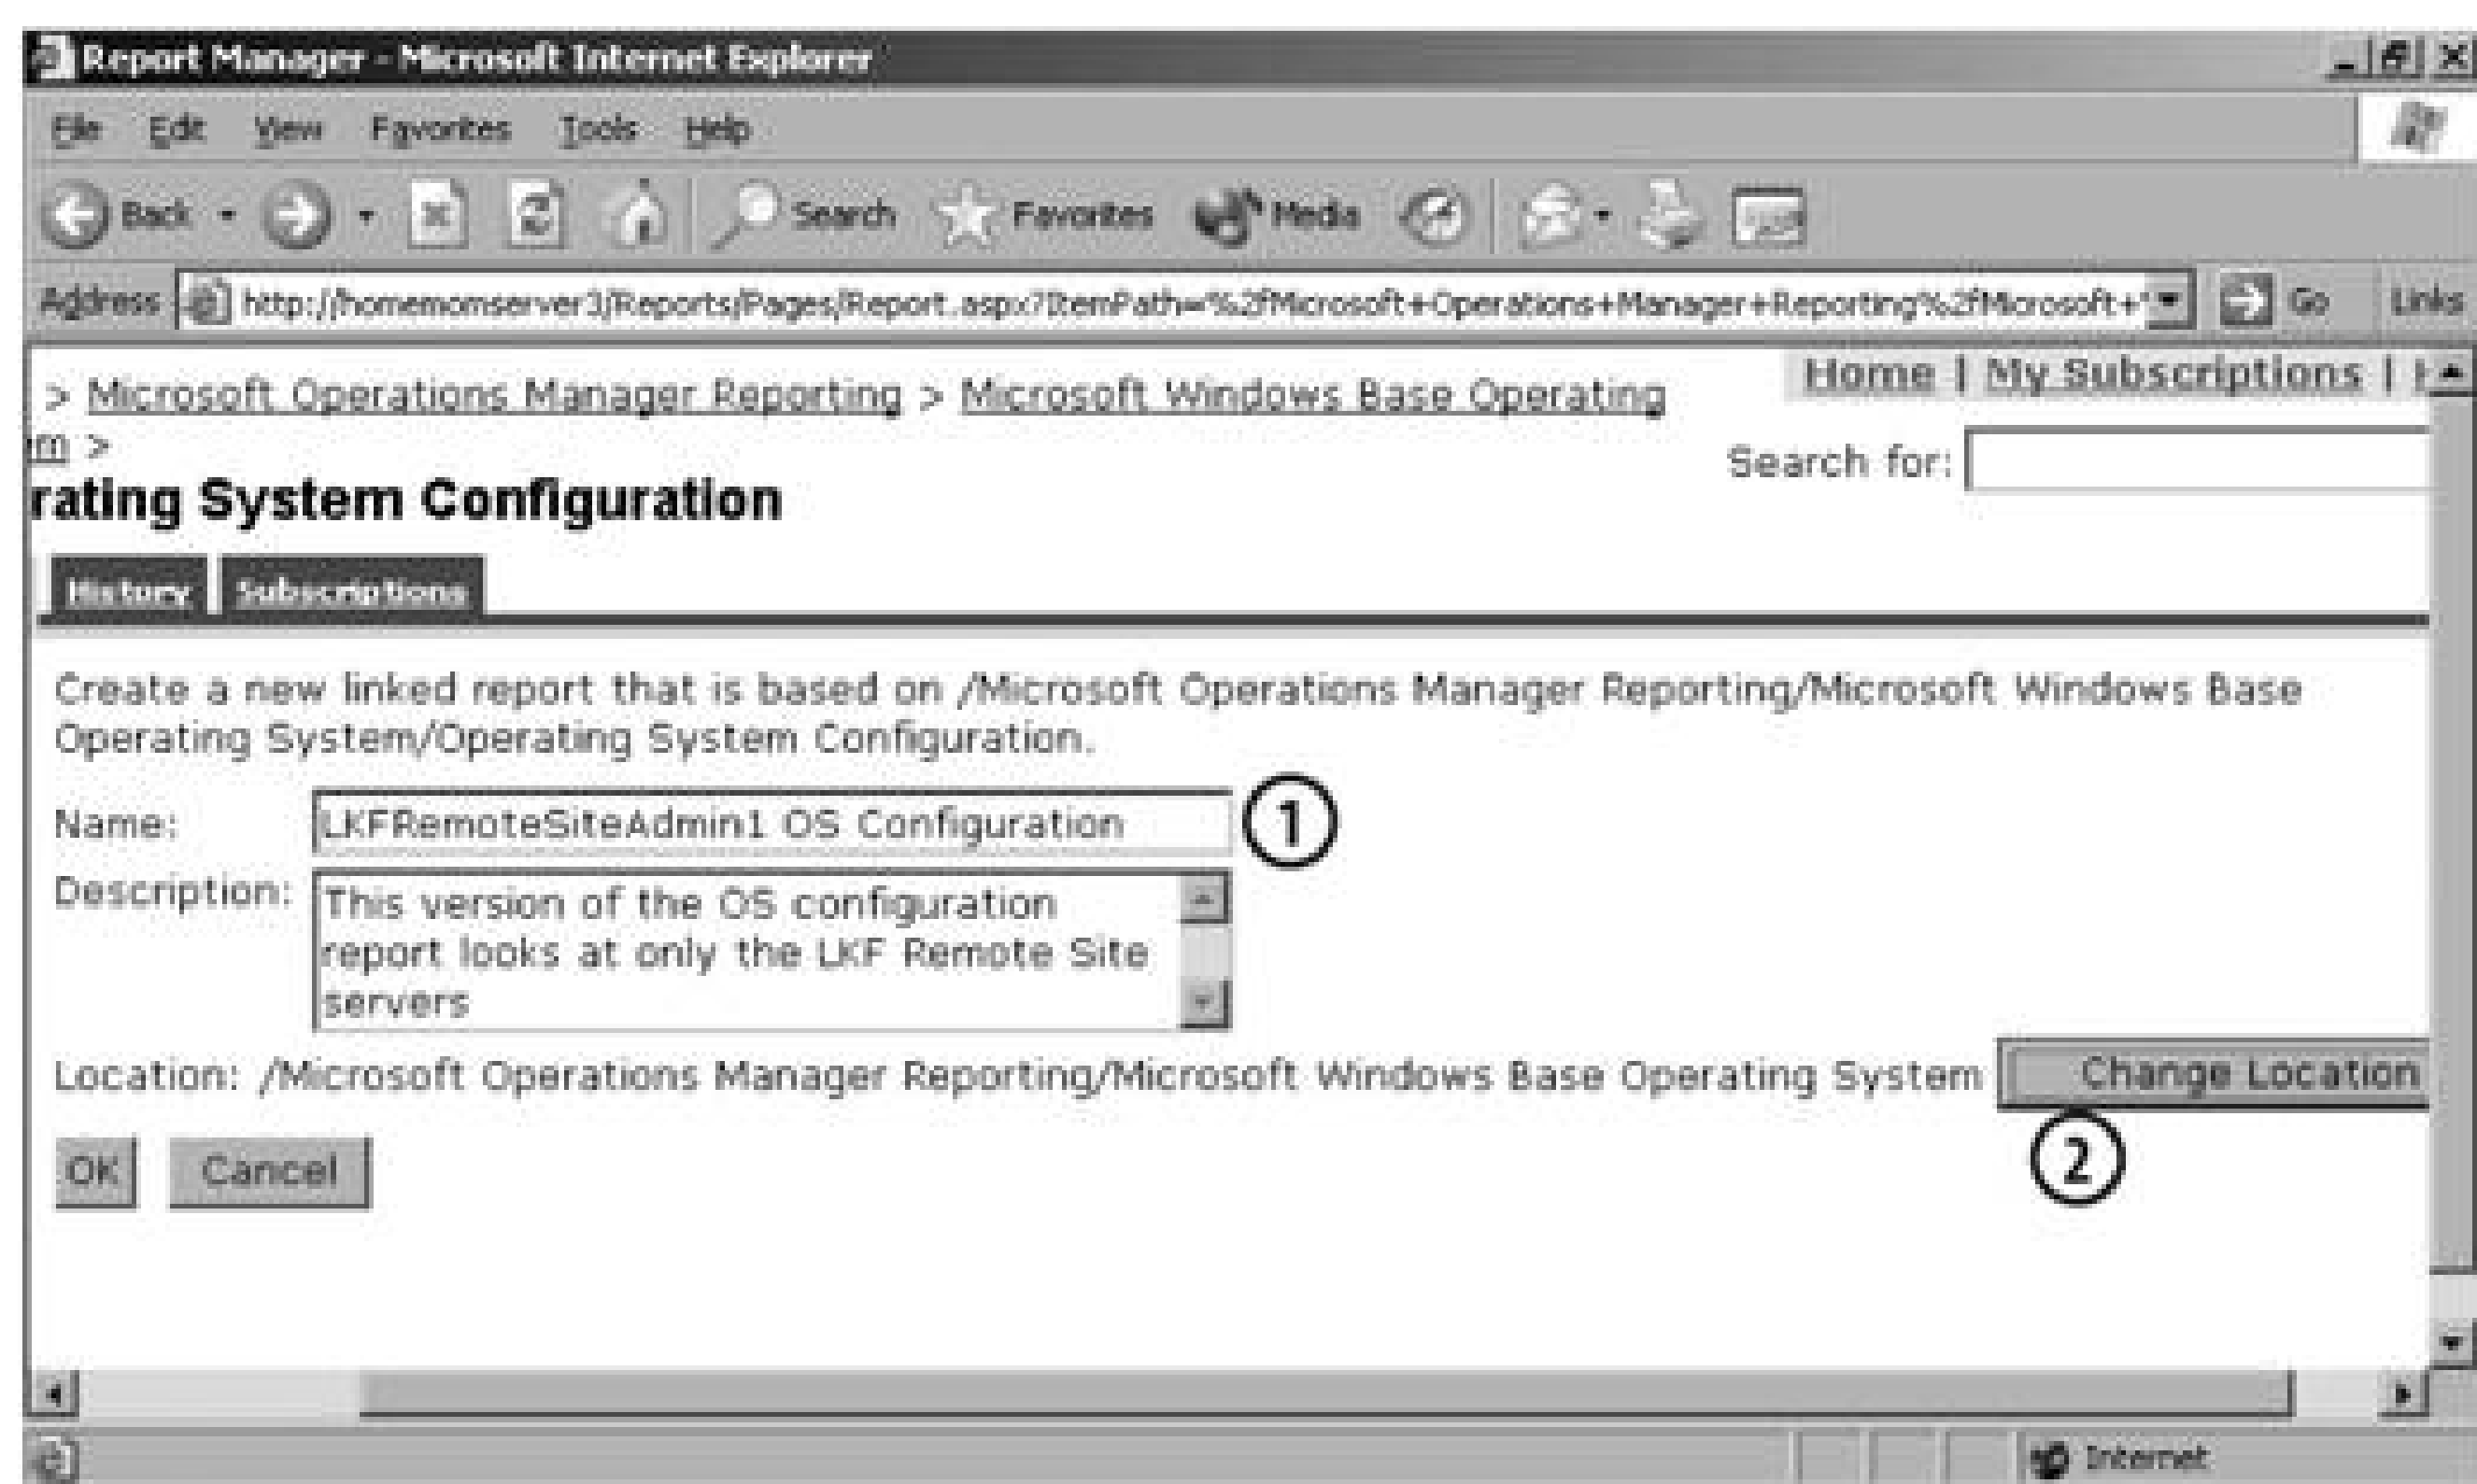
Figure 8-15. Creating a linked report from a report's Properties tab, General page



On the next page, you enter a new name for the linked report, prepending the user name to the report name (point 1 in [Figure 8-16](#)). In the description field, it is a good idea to record what distinguishes this report from the default report. Select the Change Location button (point 2 in [Figure 8-16](#)) so that the newly created linked report can be placed into the My Reports folder.

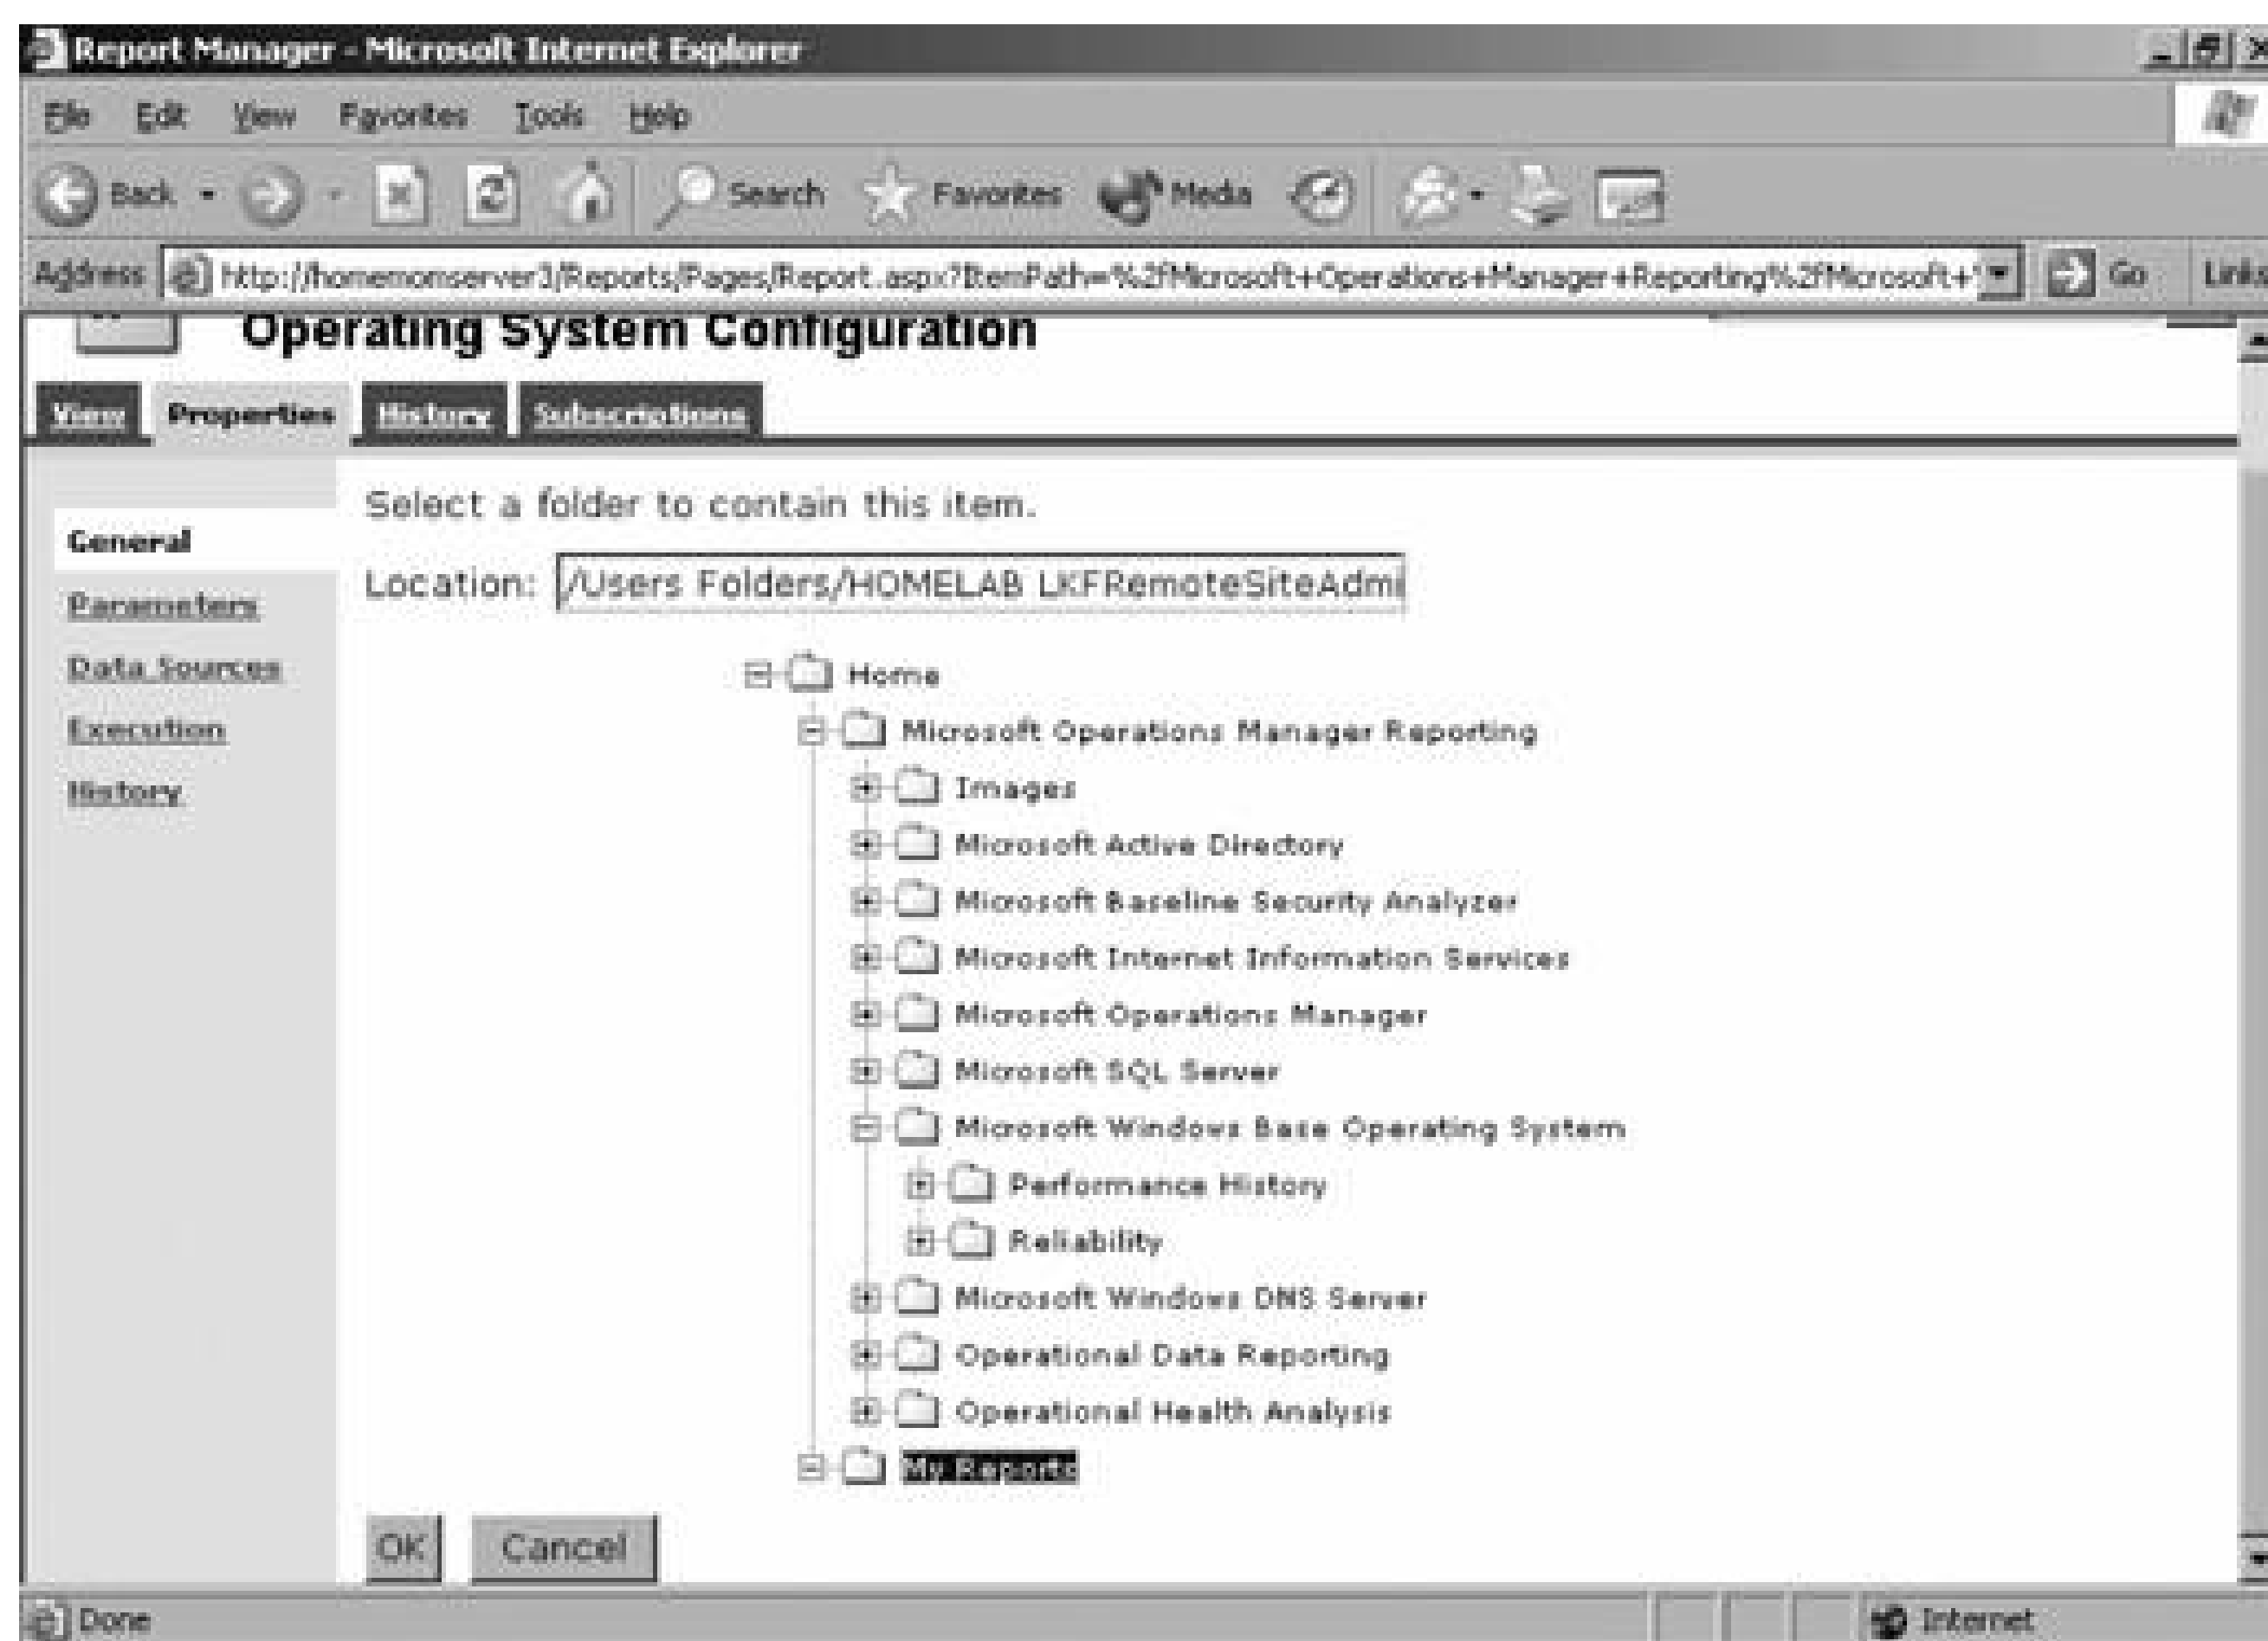
Figure 8-16. Naming a linked report and changing its location





On the change location page, the entire folder hierarchy is displayed (see [Figure 8-17](#)). A user that is in the My Reports role can place this report in any of the existing folders, just as a custom view can be moved from the My Views container to a publicly available view folder. Navigate to the My Reports folder for LKFRemoteSiteAdmin1 and click OK.

Figure 8-17. Placing the linked report in the My Reports folder

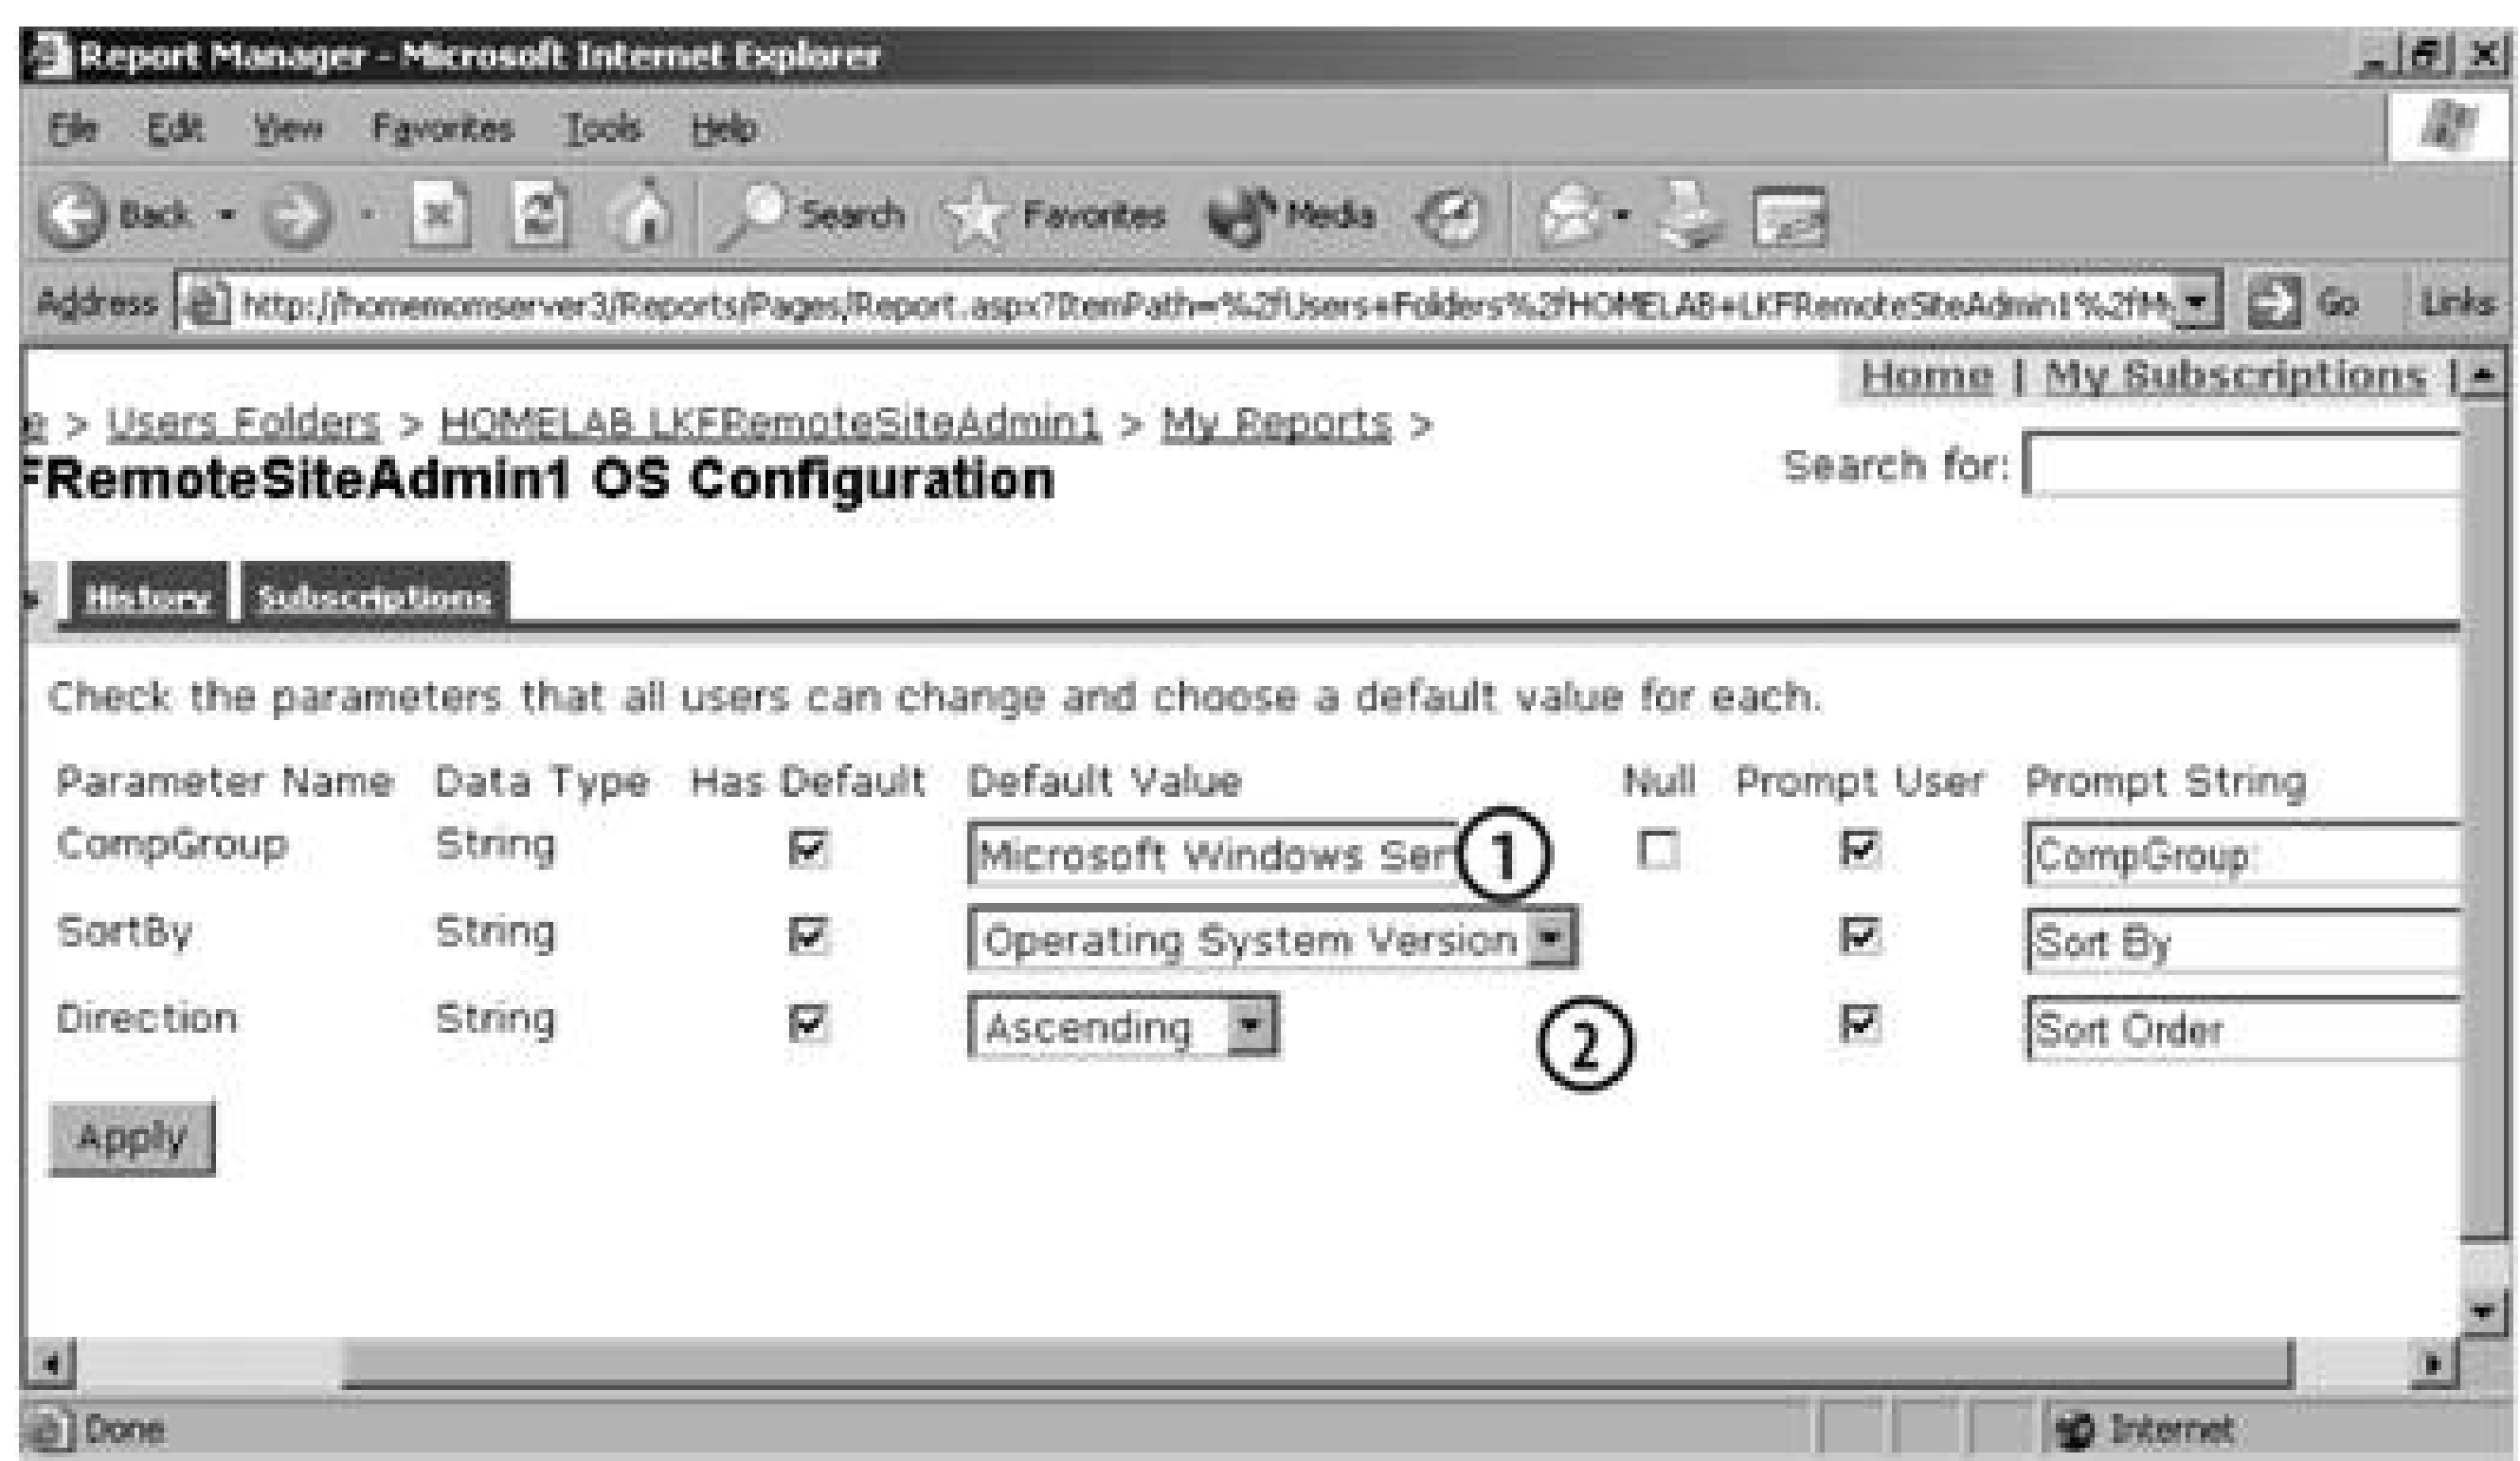


In [Figure 8-17](#), the page has changed slightly. The Location string now points to `/Users Folders/Homelab LKFRemoteSiteAdmin/My Reports`. Click OK and navigate to the My Reports folder. Here, select the report, the Properties tab, and then the Parameters page.

On this page ([Figure 8-18](#)), change the value for the CompGroup and the SortBy values. The default settings for these fields are:

- CompGroup: Blank, and the Null checkbox is selected
- SortBy: Server name
- Direction: Ascending

Figure 8-18. Configuring custom default parameters for a report

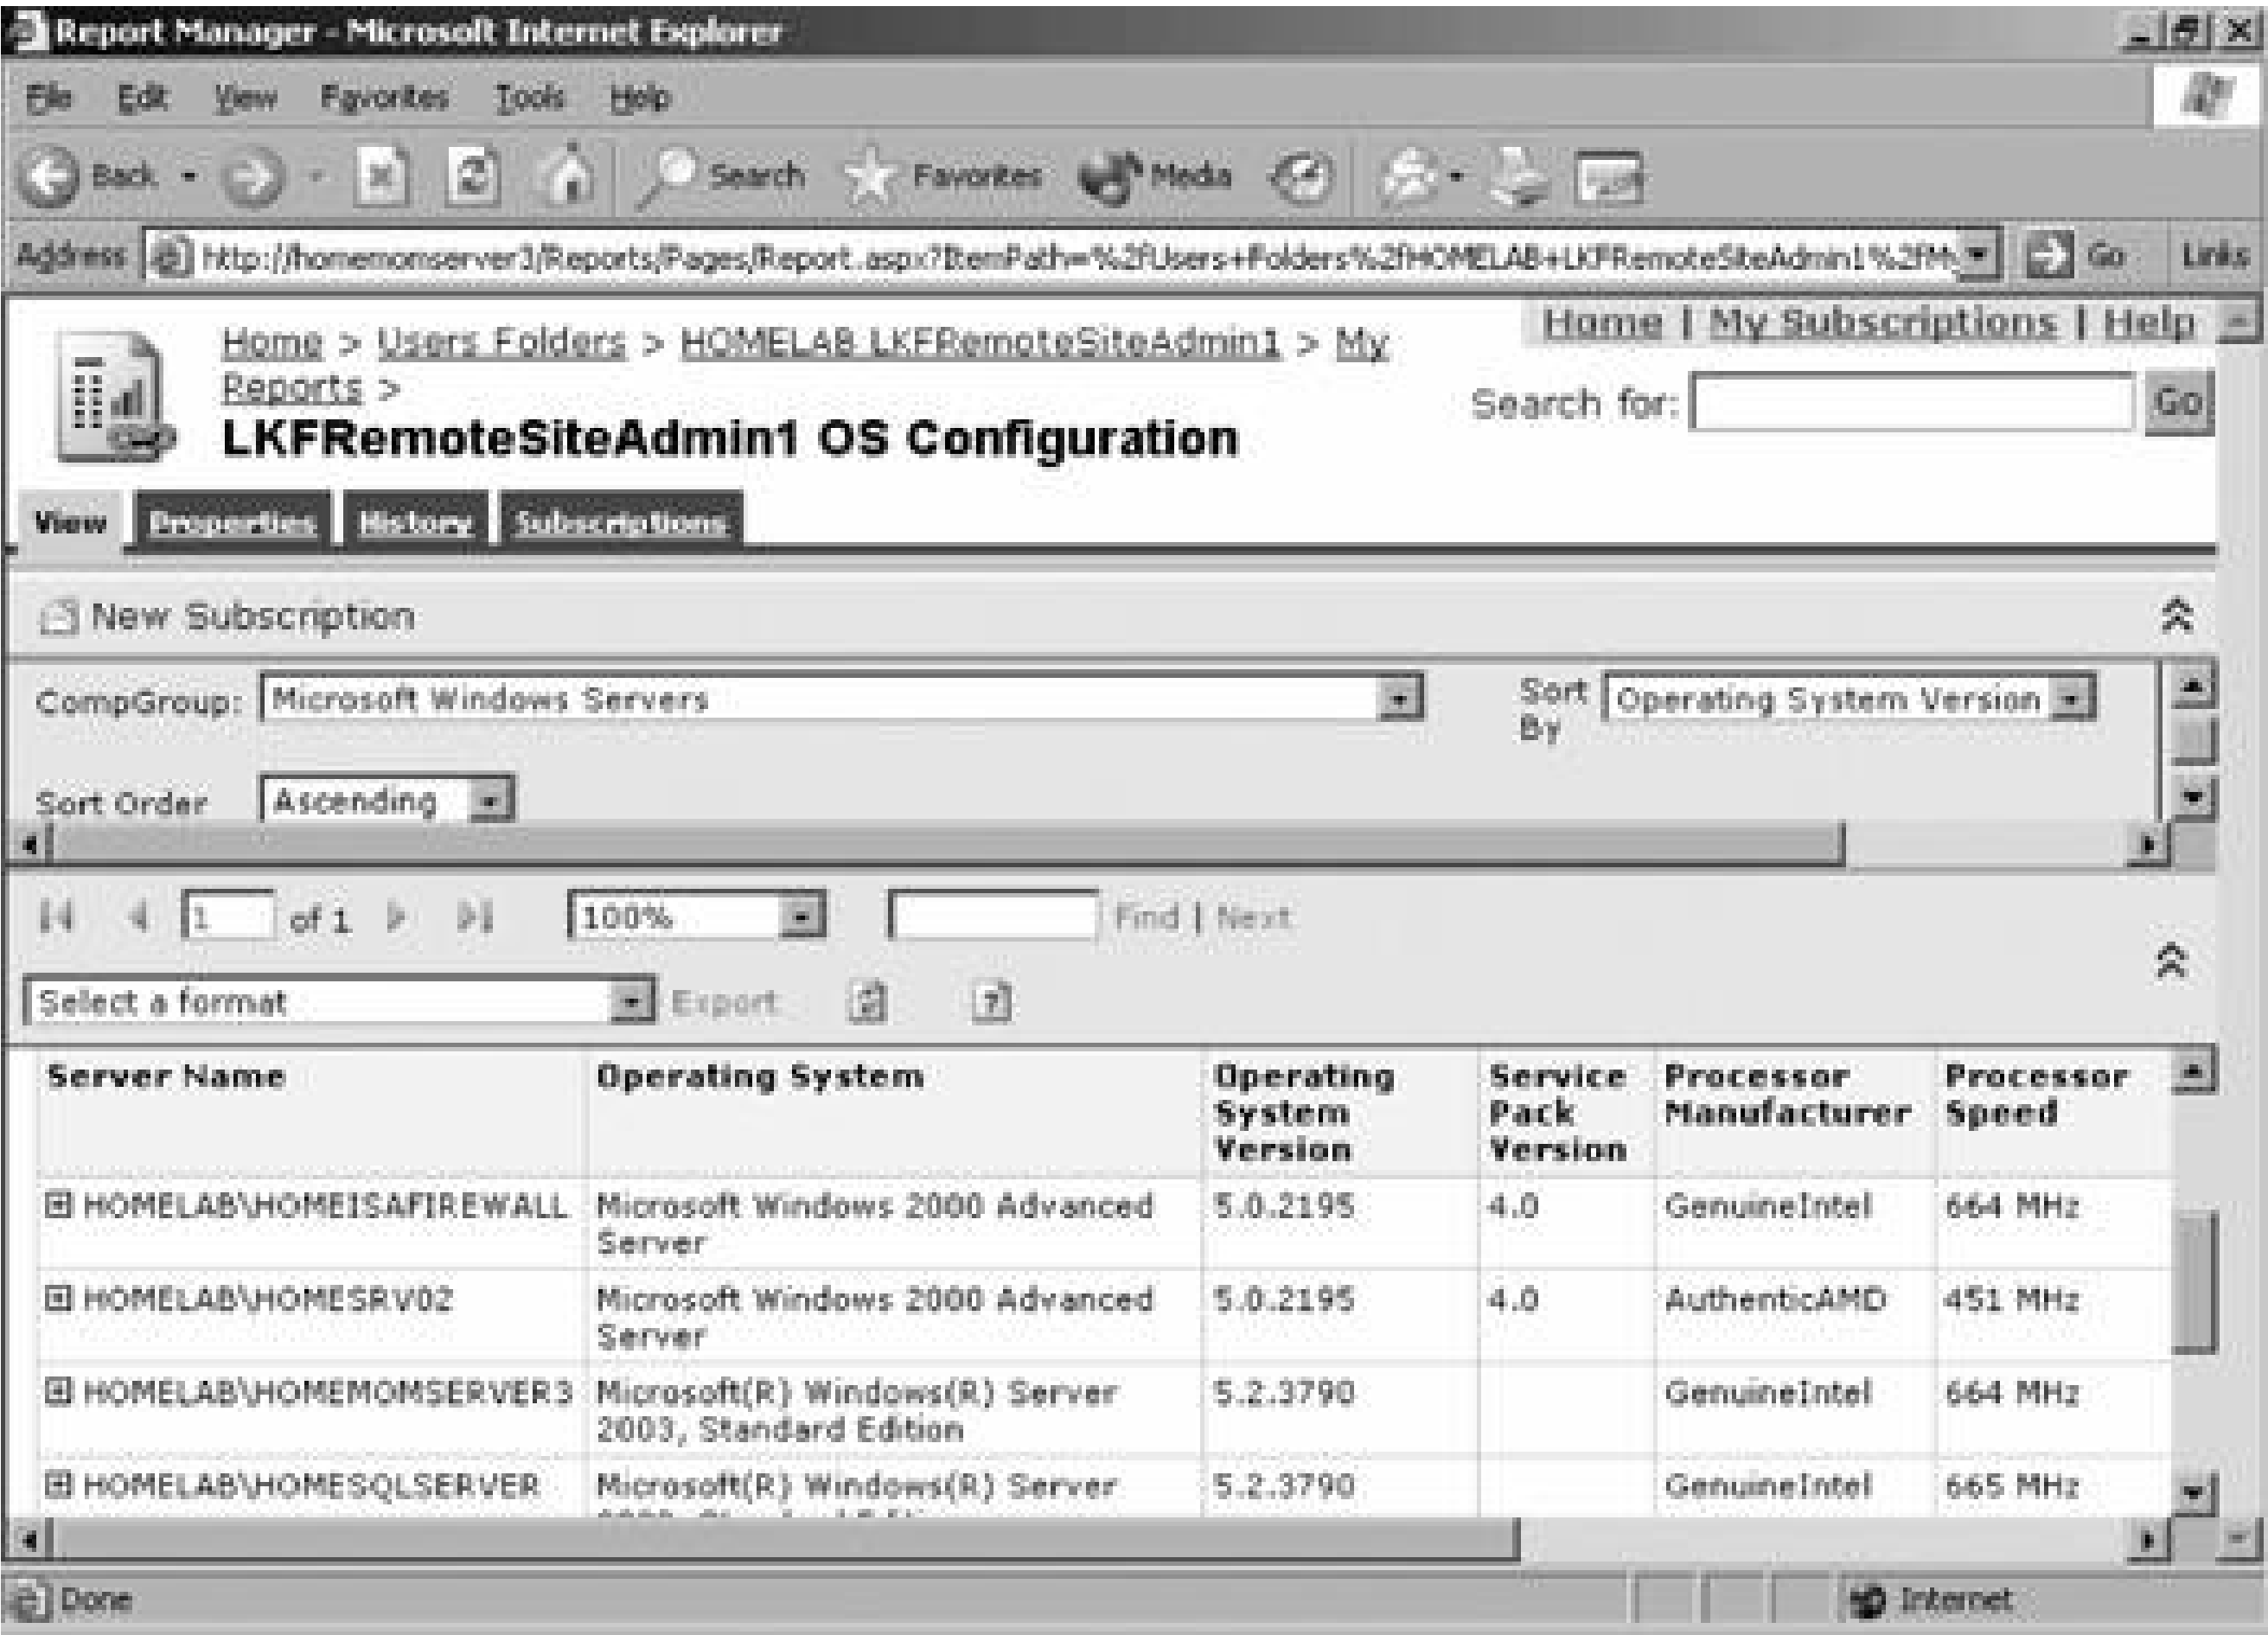


To set the CompGroup value to be Microsoft Windows Servers by default, clear the Null checkbox (point 1 in [Figure 8-18](#)) and enter the desired value. The report is also to be sorted by Operating System Version (point 2 in [Figure 8-18](#)), so that value is changed. Notice the Prompt User checkboxes to the right of points 1 and 2. By leaving these checked, the report will still allow the values in these fields to be changed from the defaults on the View page. If these checkboxes are cleared, the report will automatically execute using the defined default values when the View page is opened and there is no prompting for values. [Figure 8-19](#) shows the rendered report using the new default values.

The report icon in the upper left-hand corner has changed and now includes a chain, symbolizing the linked nature of the report. So, why go through all that trouble when you can select the desired values for CompGroup and SortBy on the public version of the report? The answer is that the linked report will always start with the desired default values, whereas the public version of the report will not. Basically, you are always presented with the version of the report that you want right from the start.

Figure 8-19. The LKFRemoteSiteAdmin1 OS Configuration report with the new default values





Linked reports are entirely dependent on the base report that they are linked to. If the base report is deleted or otherwise becomes unavailable, the linked report will cease to function.

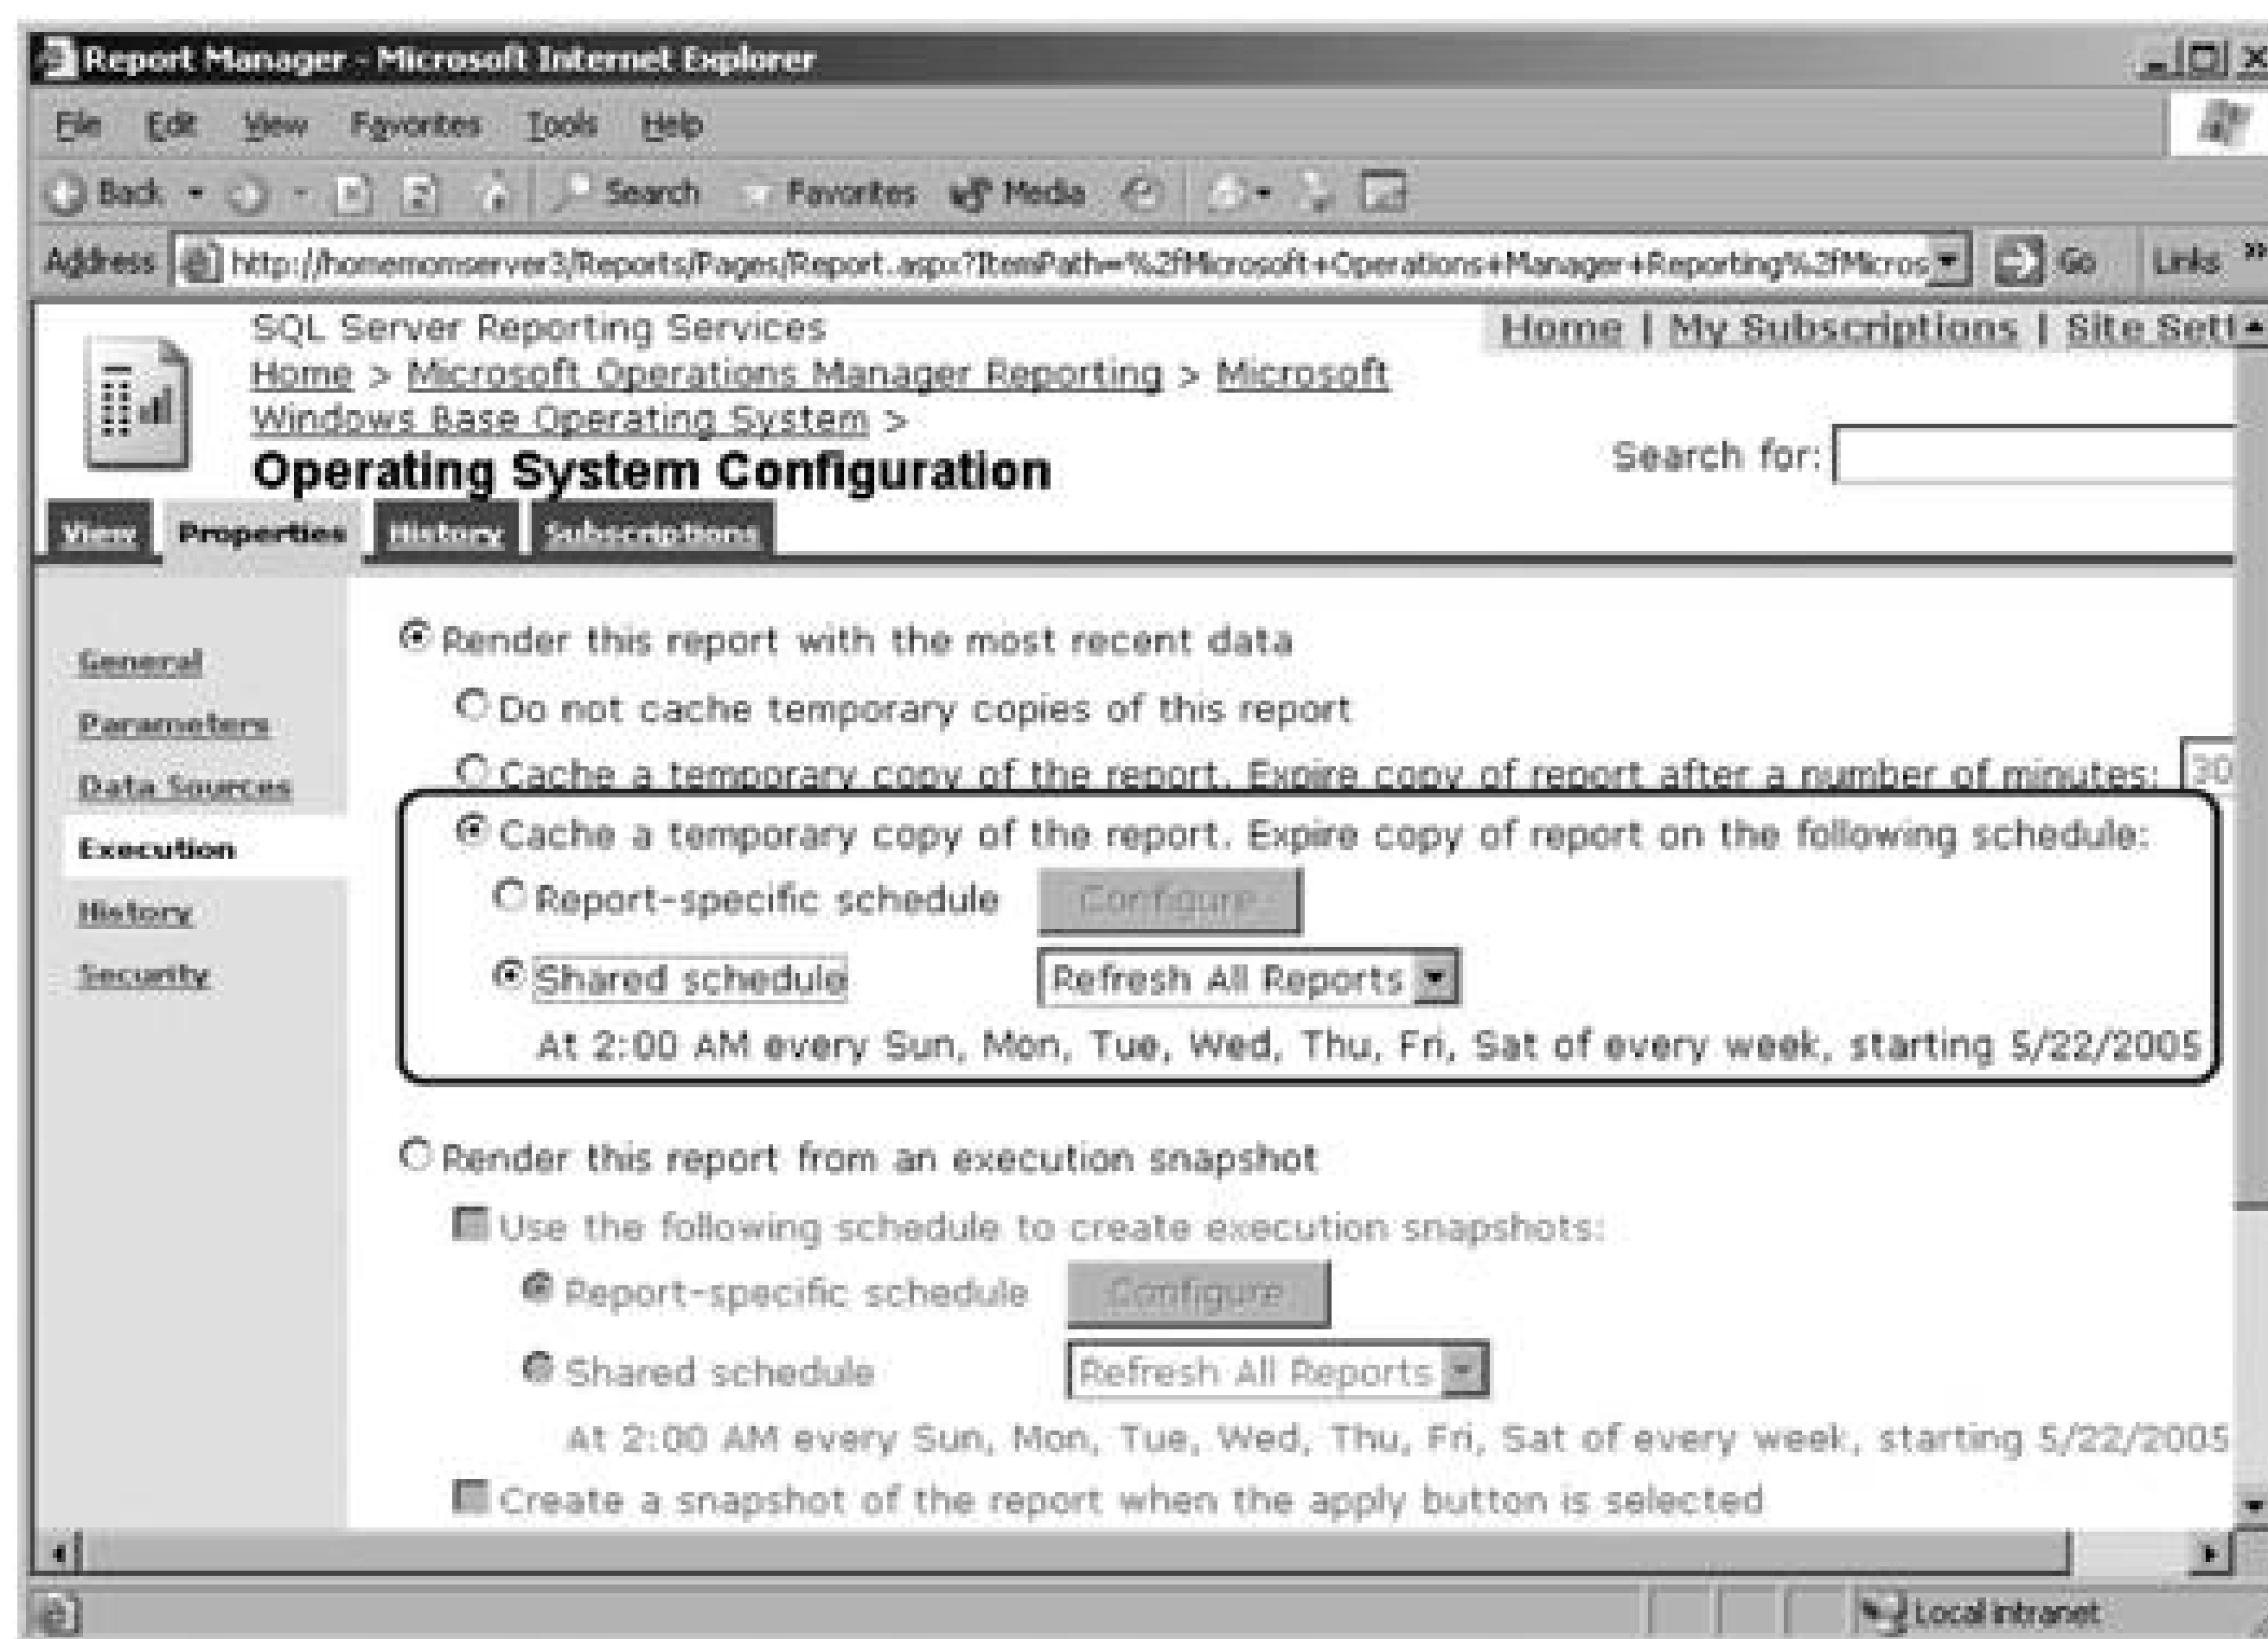
### 8.2.2.3. Controlling report execution

*On-demand execution* is when reports execute each time that they are opened and viewed. Since the SQL DTS job transfers data from the OnePoint database to the SystemCenterReporting database once a day at 1:00 a.m. (which is the default), there will be no change in the reports throughout the workday. Therefore, it is unnecessary to have the reports hit the SystemCenterReporting database each time they are viewed.

To reduce the load on the reporting server and the SystemCenterReporting database, reports can be configured to be cached for a certain period of time in the ReportServerTempDB after they are opened. Subsequent viewings of the report (within the configured time limit) access the cached version and do not generate another database hit. Once the cache time limit has expired, the report will be executed in an on-demand fashion, which then refreshes the cached version and restarts the cache report Time-To-Live countdown.

[Figure 8-20](#) shows the execution page for the public Operating System Configuration report with the cache expiration set for 2:00 a.m. daily, according to a shared schedule. The report will be executed in on-demand mode the first time it is executed every day after 2:00 a.m.

Figure 8-20. Configuring a report to refresh at 2:00 a.m. daily




#### 8.2.2.4. Shared schedules

Shared schedules are exactly what their name states. They are created by an administrator on the Site Settings page Managed Shared Schedules link (see [Figure 8-21](#)). These schedules can be used for report caching, generating report subscriptions, or report snapshots.

A shared schedule consists of a schedule name, start and end dates (which can be left open-ended), and the schedule details. For each shared schedule created, all reports that make use of it are listed on the Reports page. Execution of a shared schedule can be paused, thereby pausing all jobs that are dependent on them. This is done on the Shared Schedule Summary page, Site Settings Manage Shared Schedules.

Figure 8-21. Configuring a shared schedule



SQL Server Reporting Services  
**Scheduling**

[Home](#) | [My Subscriptions](#) | [Site Settings](#) | [Help](#)

Search for:

Schedule

Reports

Use this page to create or modify a schedule.

Schedule Name:

**Schedule details**

Define a schedule that runs on an hourly, daily, weekly, monthly, or one time basis. All times are expressed in (GMT -05:00) Central Daylight Time.

☐ Hour  
☒ Day  
☐ Week  
☐ Month  
☐ Once

**Daily Schedule**

☒ On the following days:  
☒ Sun ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☒ Sat

☐ Every weekday

☐ Repeat after this number of days:

Start time:  :  ☒ AM ☐ PM

**Start and end dates**

Specify the date to start and optionally end this schedule.

Begin running this schedule on:

## 8.2.2.5. Snapshot execution

On-demand reports query the SystemCenterReporting database, returning the current values of whatever data was queried. When the daily DTS transfer runs, all data that is more than five minutes old and that has not copied in the previous transfer cycle (basically the past 24 hours of data) is added to the SystemCenterReporting database and made available to the report queries. Because the queried data is changing, the reports by their very nature must change as well. If on-demand reports were the only way to extract data from the SystemCenterReporting database, it would be very difficult to keep a historical record of reports as they existed at any given point in time.

To enable historical tracking of reports, you can take snapshots of reports and keep them in the report's history on the reporting server, or deposit them to a file share or email mailbox via a subscription. Snapshot reports can be generated and stored on a scheduled basis or (ironically) on demand. Reports that are executed on a scheduled basis do so without human interaction, which means that all report parameters must be predefined and a set of credentials that can access the data source must be provided. Snapshot reports, like rendered subscription reports, are non-interactive. Scheduled snapshot execution can be done according to a report-specific schedule or a shared schedule. If you want your users to make use of a shared schedule for generating snapshots (or subscriptions for that matter), make sure they have been assigned a My Reports item role and System User system-level role. System-level roles are assigned on the Site Settings → Configure Site Wide Security page. Create the association between the domain/local user or group and the System User role. Otherwise, you (as the administrator) will need to configure a snapshot report to use a shared schedule.

To take and save a snapshot of a report to that report's history folder, first navigate to the report of interest, then to the Properties tab (see [Figure 8-22](#)):

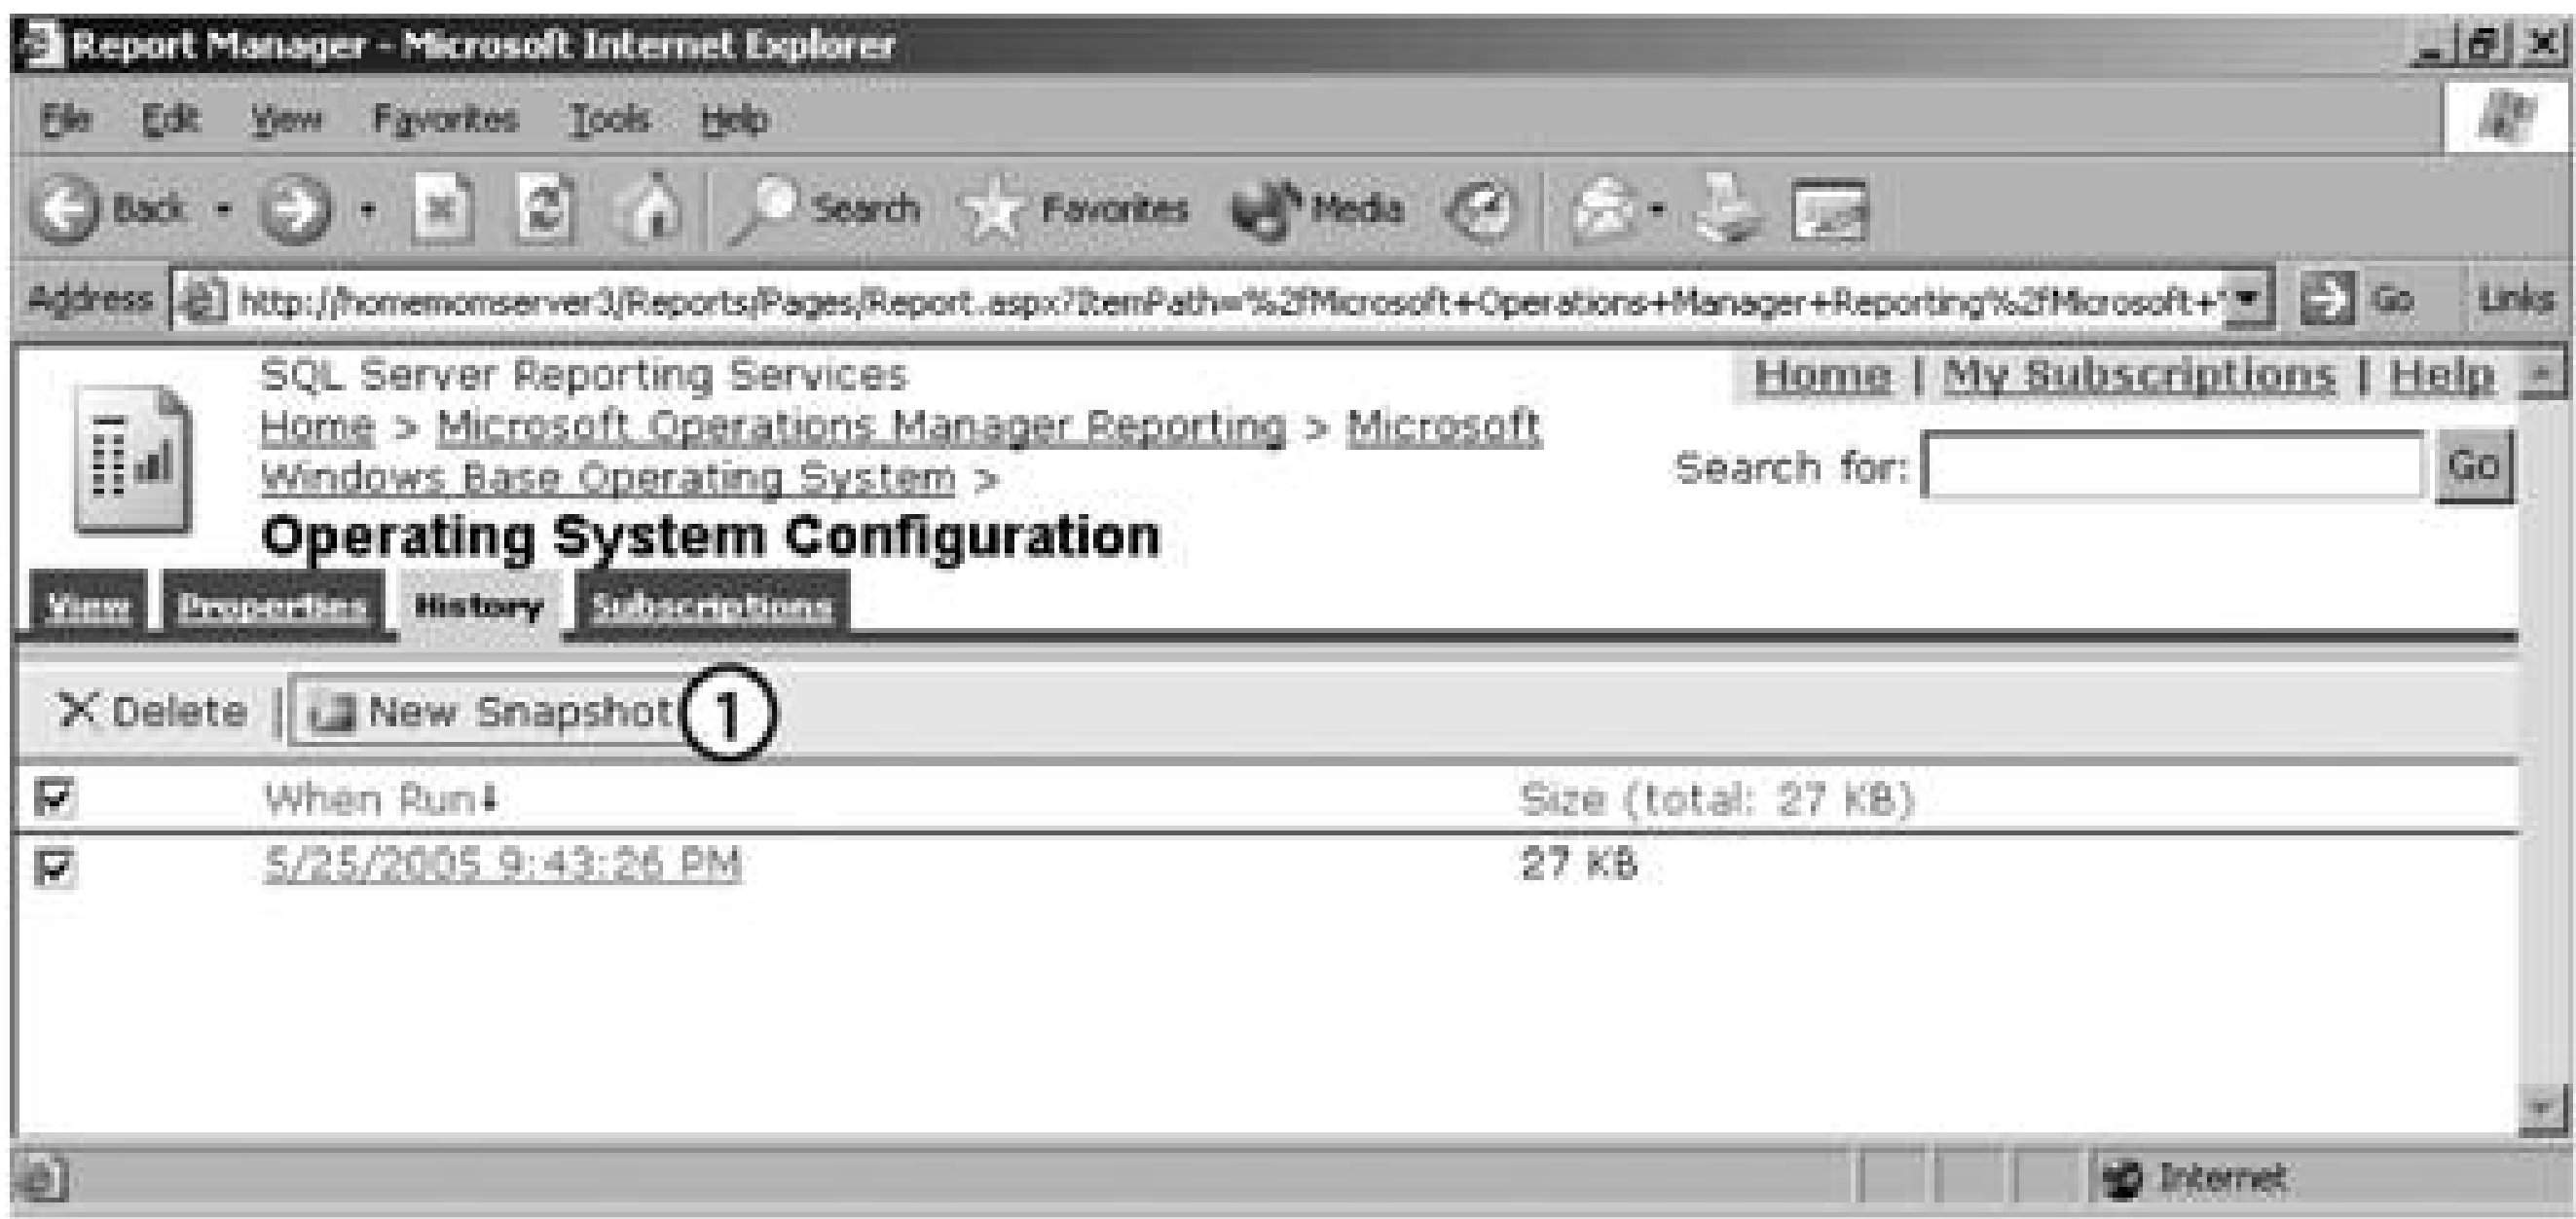


1. Select the History link along the left side of the page (point 1). This opens the report history configuration page.
2. Select the "Allow history to be created manually" checkbox (point 2). This enables the New Snapshot button on the History tab.
3. Select the "Store all report execution snapshots in history" checkbox (point 3) to keep all of the snapshots in the reports history.
4. Select "Shared Schedule" since this snapshot will be collected at 2:10 a.m. every day. It has already been configured in the previous section (point 4).
5. Select to keep an unlimited amount of reports in the history (point 5).
6. Click Apply (point 6).

Figure 8-22. Configuring report snapshots on the history configuration page

At this point, the scheduled time for snapshot creation has not passed, so there will be no entries on the History tab. To generate an on-demand snapshot, click the History tab, then click the New Snapshot button (point 1 in [Figure 8-23](#)).

Figure 8-23. Creating a report snapshot on demand

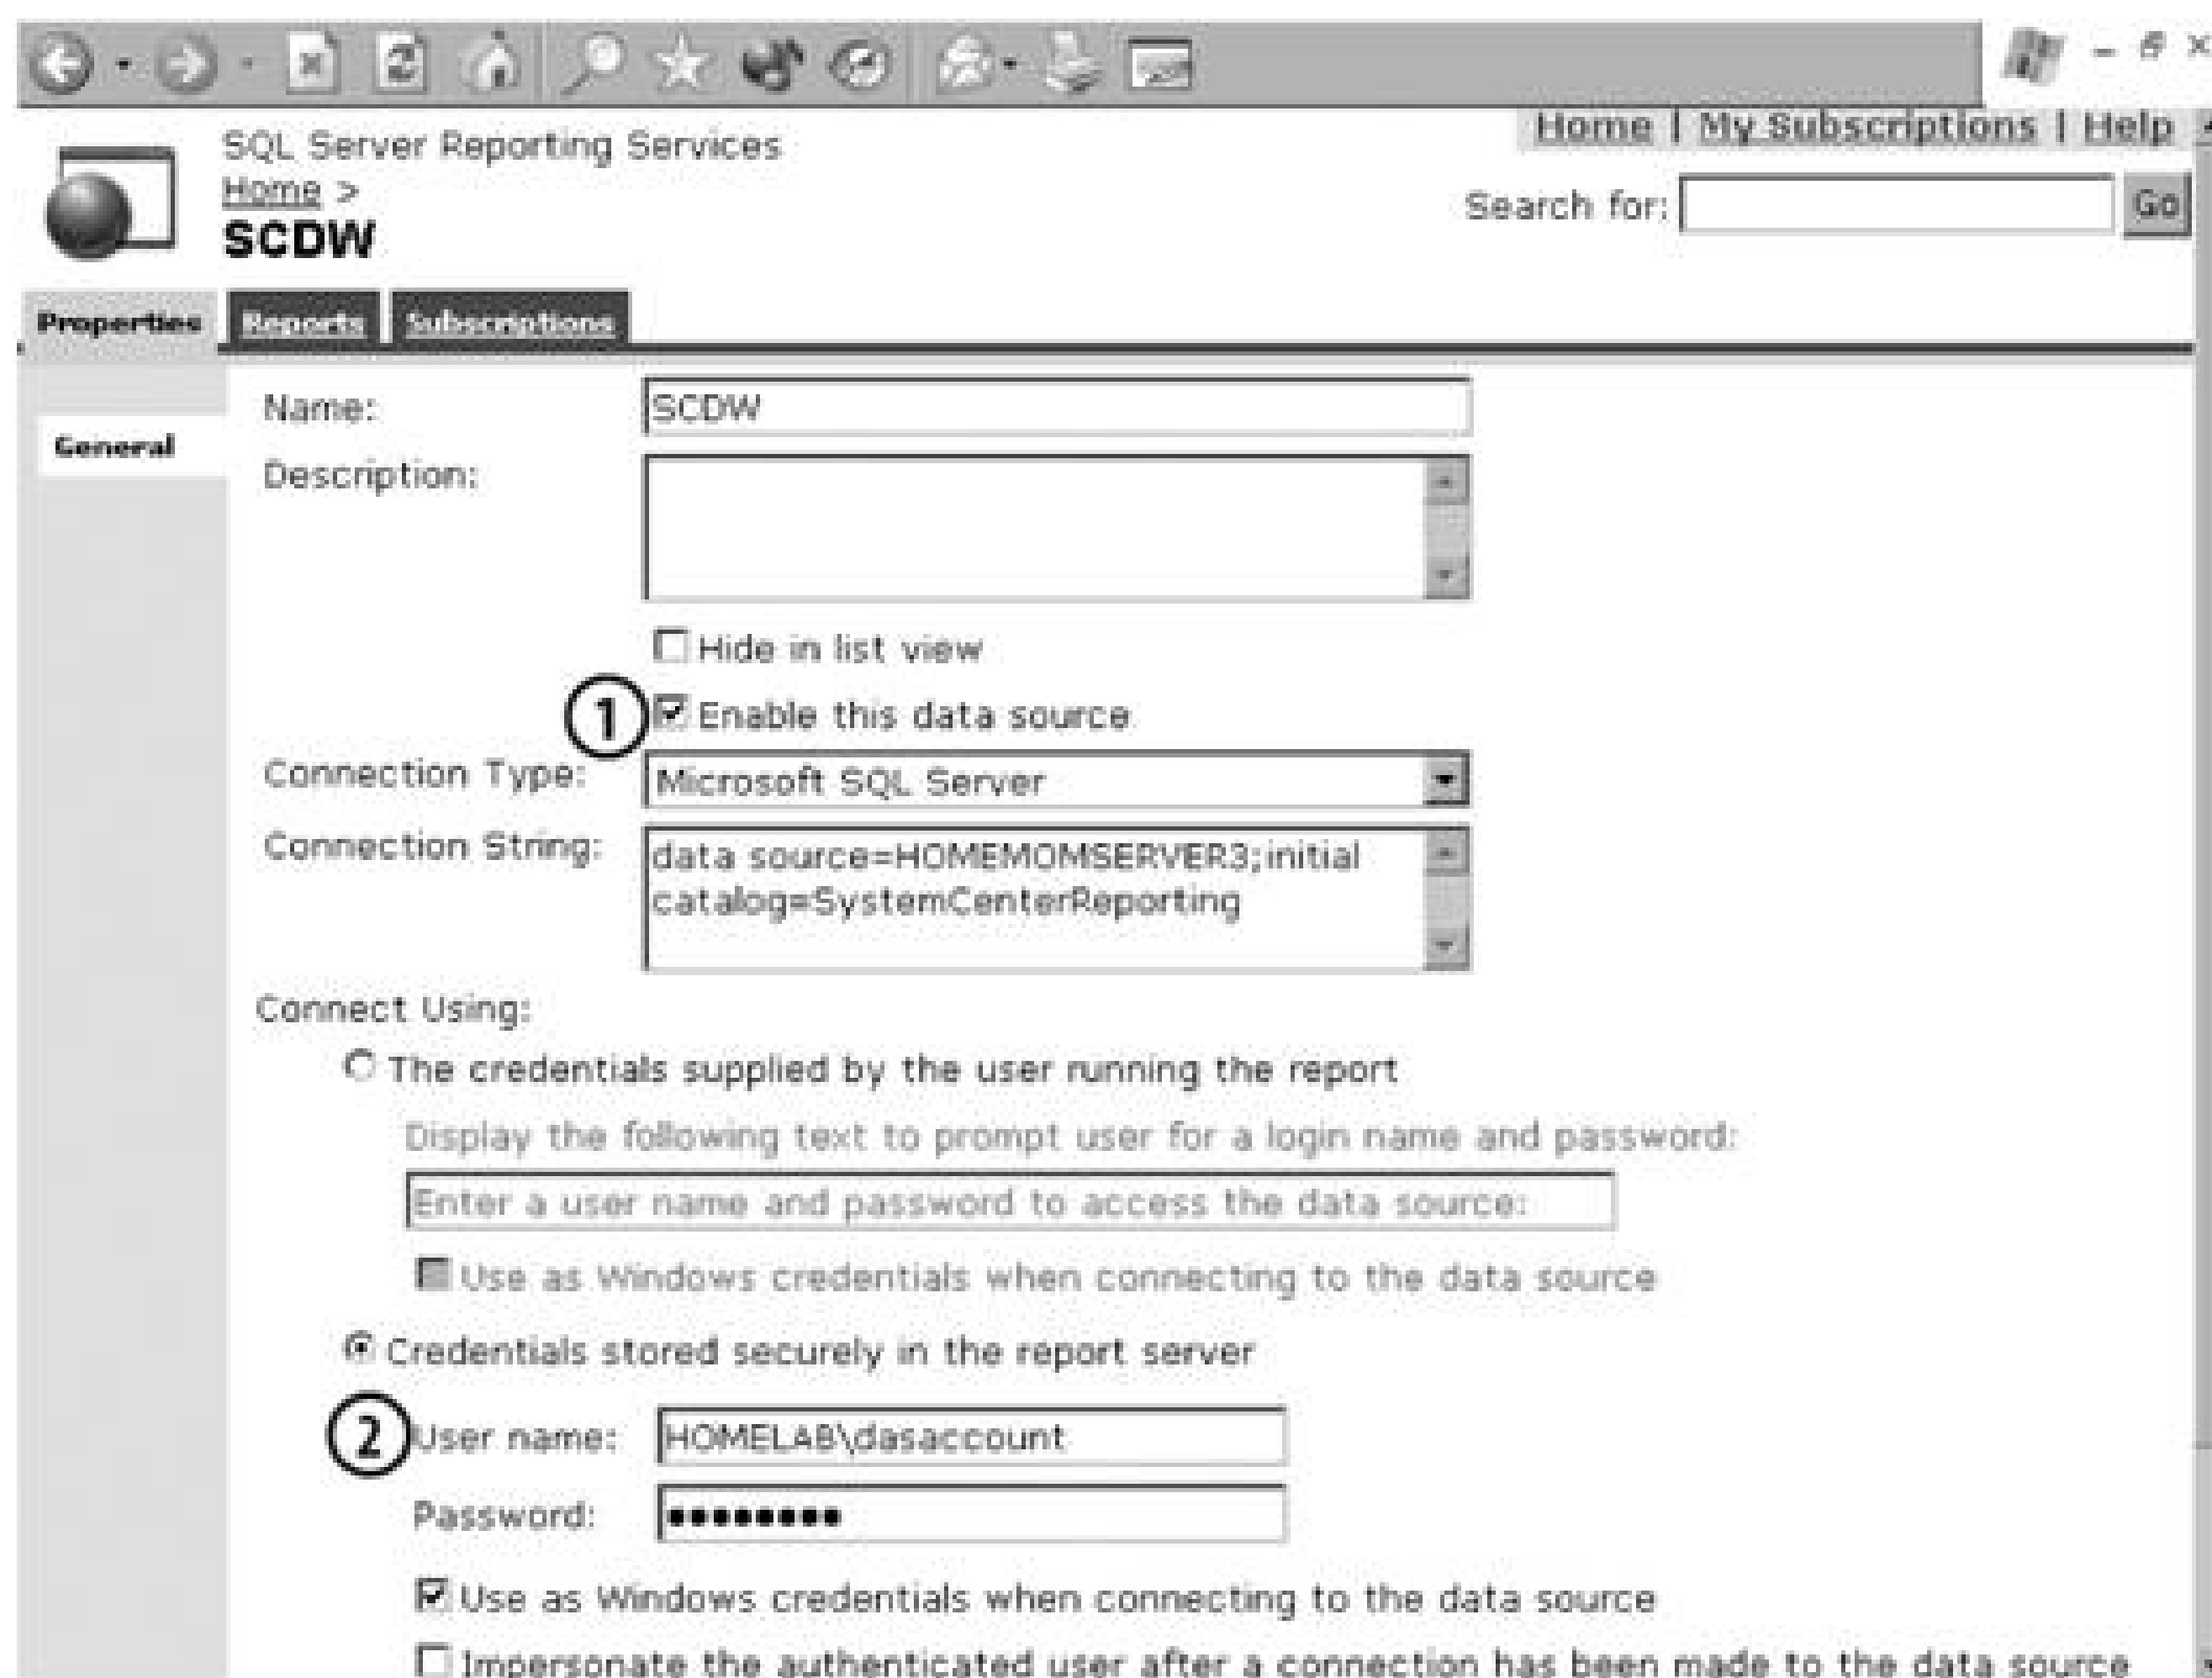


Snapshot reports that are run according to a schedule must have a set of credentials stored in the data source. For MOM 2005 reports, this is not something you really need to worry about because they will all use a shared data source, which is the SCDW object that you see in most report folders, including the Home folder (see [Figure 8-24](#)).

Figure 8-24. SRS data source definition

There are a few things that you need to be aware of in this shared data source. In point 1 in [Figure 8-25](#) the "Enable this data source" checkbox is the single control point that you can use to stop the execution of all reports. If you disable this data source, all of the reports listed on the Reports tab will cease to function until the data source is re-enabled. This shared data source already has a stored set of credentials that are used to access the SystemCenterReporting database, the homelab/dasaccount (point 2 in [Figure 8-25](#)).

Figure 8-25. Properties of a shared data source



### 8.2.2.6. Managing subscriptions

You will probably be asked to create a subscription to certain reports that are configured with certain parameters for a user or group of users. As a MOM administrator, you can configure two types of subscriptions: a standard subscription, in which every recipient of the report receives it in the same format and at the same location (their inbox or a filesystem folder), or a data-driven subscription.

Data-driven subscriptions for reports are probably the most useful because they allow you to mix report configurations (e.g., parameters, recipients, and delivery type) in a single SQL table, which is read from when the subscription runs. Each recipient then receives the report within her parameters (the report she wants), rendered in the format she wants (e.g., .PDF, .XLS, .HTM, and .CSV) and delivered where she wants it (e.g., inbox or file share).

To create a standard subscription for a user as an administrator, navigate to the desired report, then click the Subscriptions tab, and select the New Subscriptions button.

This brings up the Report Delivery Options page (see [Figure 8-26](#)), which is configured as follows:

- Delivered by: Report email server
- To: SMTP addresses, [LKFRemoteSiteAdmin1@homelab.lab](mailto:LKFRemoteSiteAdmin1@homelab.lab), [LKFCFO@homelab.lab](mailto:LKFCFO@homelab.lab), and



[LKFMOMAdministrators@homelab.lab](mailto:LKFMOMAdministrators@homelab.lab)

- Reply to: [reportmaster@homelab.lab](mailto:reportmaster@homelab.lab)
- Rendered: As a PDF file and included in the email as an attachment
- Run: According to a shared schedule called "Create and send subscriptions"

Figure 8-26. Configuring standard subscription options

Operational Data Reporting > **Subscription: Management Group Health** Search for:  Go

**Report Delivery Options**

Specify options for report delivery.

Delivered by:

To:

Cc:

Bcc:

(Use ";" to separate multiple e-mail addresses.)

Reply-To:

Subject:

☒ Include Report    Render Format:

☐ Include Link

Priority:

Comment:

**Subscription Processing Options**

Specify options for subscription processing.

Run the subscription:

☐ When the scheduled report run is complete.

At 8:00 AM every Mon of every week, starting 5/25/2005

☒ On a shared schedule:

At 6:00 AM every Mon, Tue, Wed, Thu, Fri of every week, starting 5/25/2005

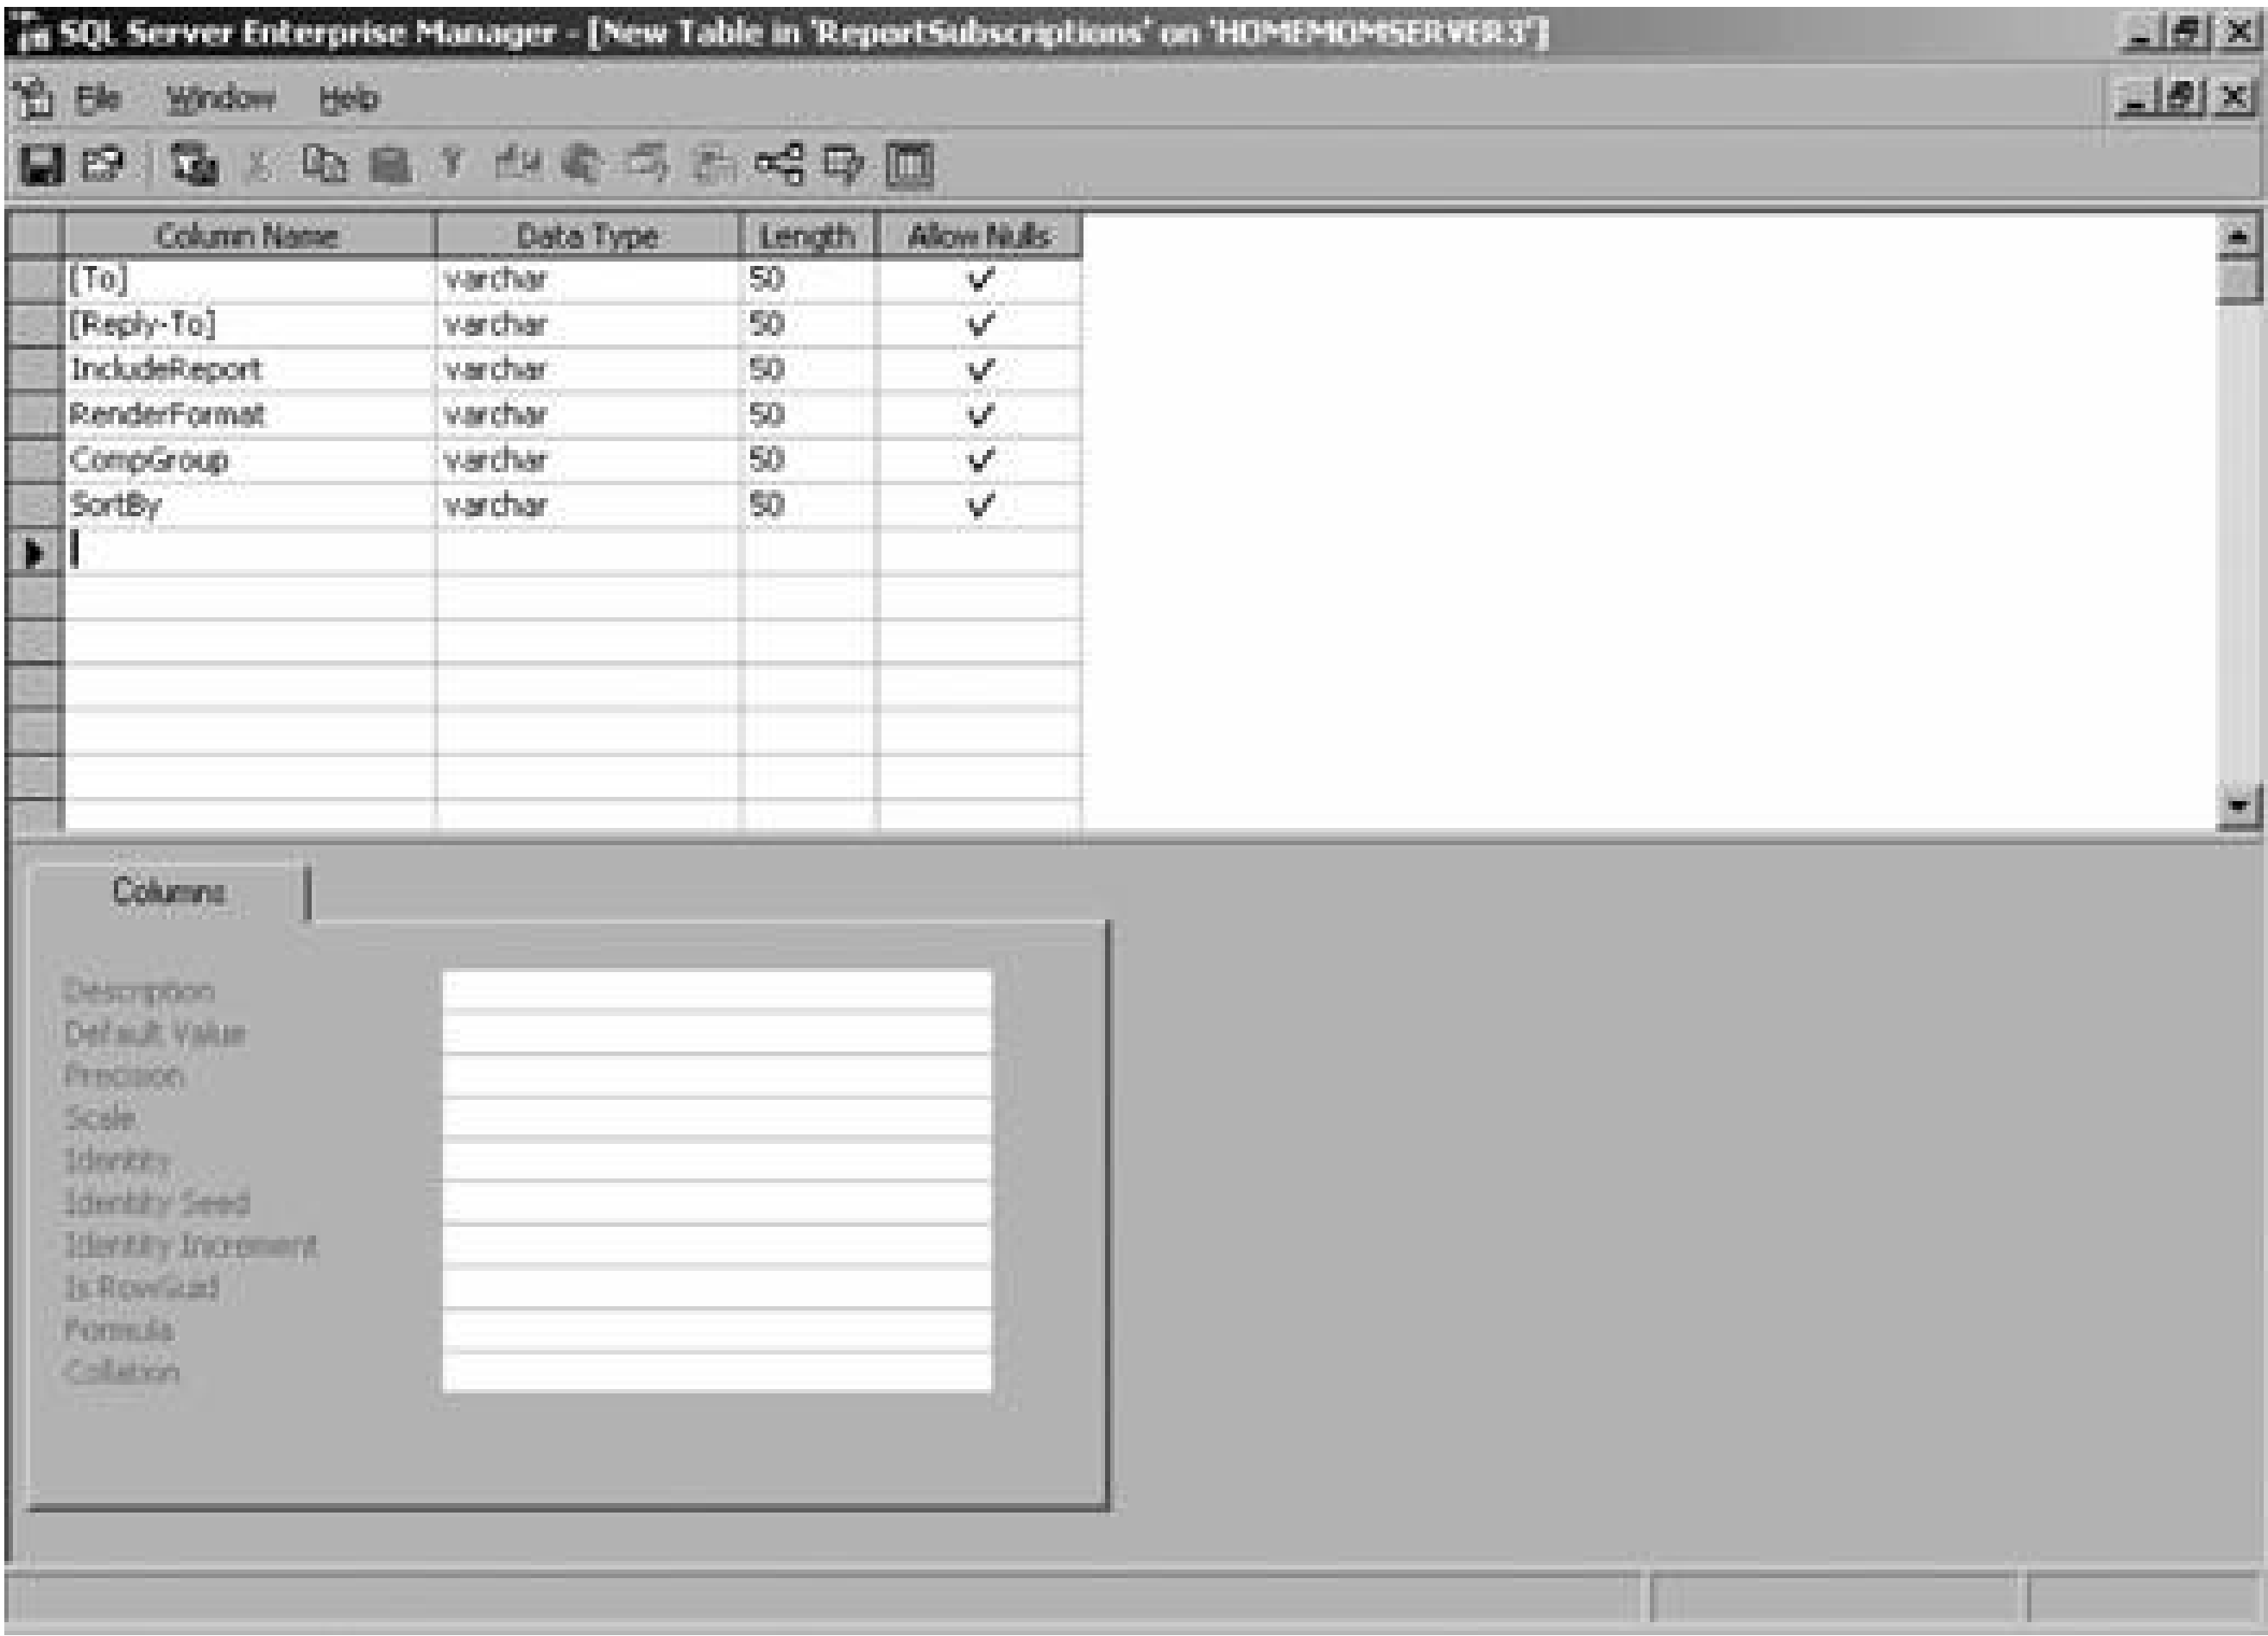
This report will be rendered as a PDF file and emailed to the recipients every morning at 6:00 a.m.

Creating a data-driven subscription requires more work up front, but saves you from creating and maintaining a standard subscription for every user who wants a report delivered to him in a different format or with different fields. In a nutshell, you create a database that has a table that houses all the fields required for generating the report, such as To, Reply-To, IncludeReport, IncludeLink, RenderFormat, and any other required parameter fields. These values can all be recorded in a single table with each row in the table representing a different recipient. Then, in the subscription creation process, you access this table and map the table fields to the report fields and set a delivery trigger. Whenever the subscription triggers, each recipient receives the report he wants, in the format that he wants it. For every report that you want to create a data-driven subscription for, you must gather the parameters in the subscriber's database along with the recipient's information, and any other applicable options. These will be used to create columns in the subscriber's information table.

Here's how to create a data-driven subscription for the operating system configuration report:

1. Log on as an administrator on the reporting server, open SQL Enterprise Manager, and navigate to the Databases folder.
2. Right-click and select New Database.
3. Enter a name, for example ReportSubscriptions. Click OK.
4. Right-click on the ReportSubscriptions database. Select New Table.
5. Enter the subscription data so that it looks like the table in [Figure 8-27](#). You can type directly in the cells. When you enter the varchar data type, the length will automatically be set to 50 .

Figure 8-27. Entering the subscriber's information for a data-driven subscription

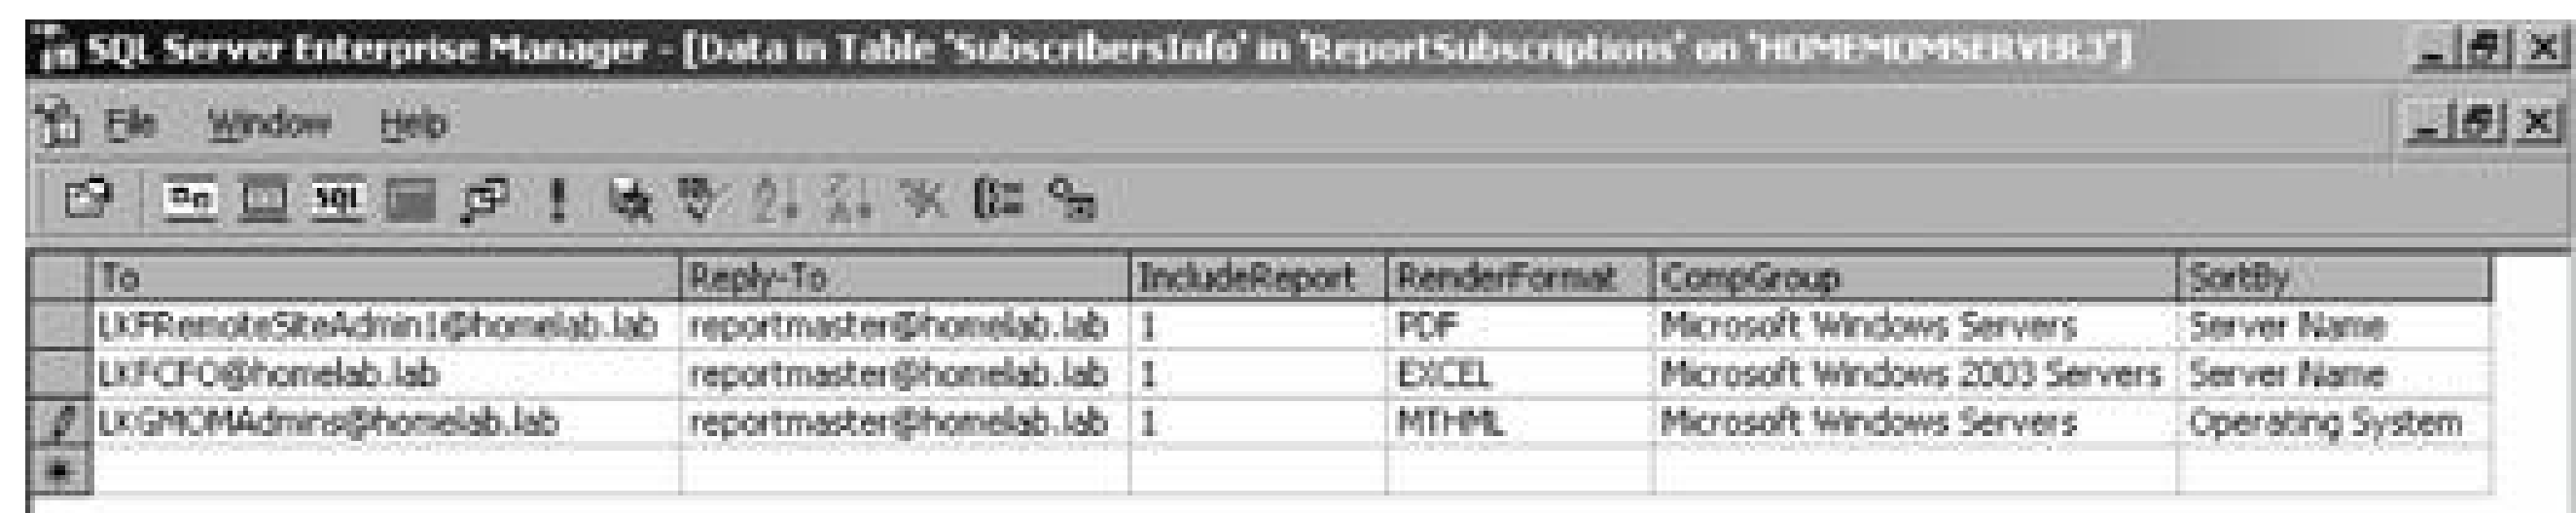


6. Click Save (the icon of the floppy disk in the upper lefthand corner). Name the table SubscribersInfo.
7. Select the SubscribersInfo table, right-click, and select Open Table Return All Rows. This opens the table so you can enter the values in each of the columns. For this example, the data entered is:

- To: [LKFRemoteSiteAdmin1@homelab.lab](mailto:LKFRemoteSiteAdmin1@homelab.lab), [LKFMOMAdmins@homelab.lab](mailto:LKFMOMAdmins@homelab.lab), and [LKFCFO@homelab.lab](mailto:LKFCFO@homelab.lab)
- Reply-To: [reportmaster@homelab.lab](mailto:reportmaster@homelab.lab)
- IncludeReport: 1
- RenderFormat: PDF, Excel, MHTML
- CompGroup: Microsoft Windows Servers, Microsoft Windows 2003 Servers, Microsoft Windows Servers
- SortBy: Server Name, Server Name, Operating System

When you're done entering data, the SubscribersTable looks like [Figure 8-28](#). Close the table and the values are automatically saved.

Figure 8-28. Enter the desired values in the SubscribersInfo table

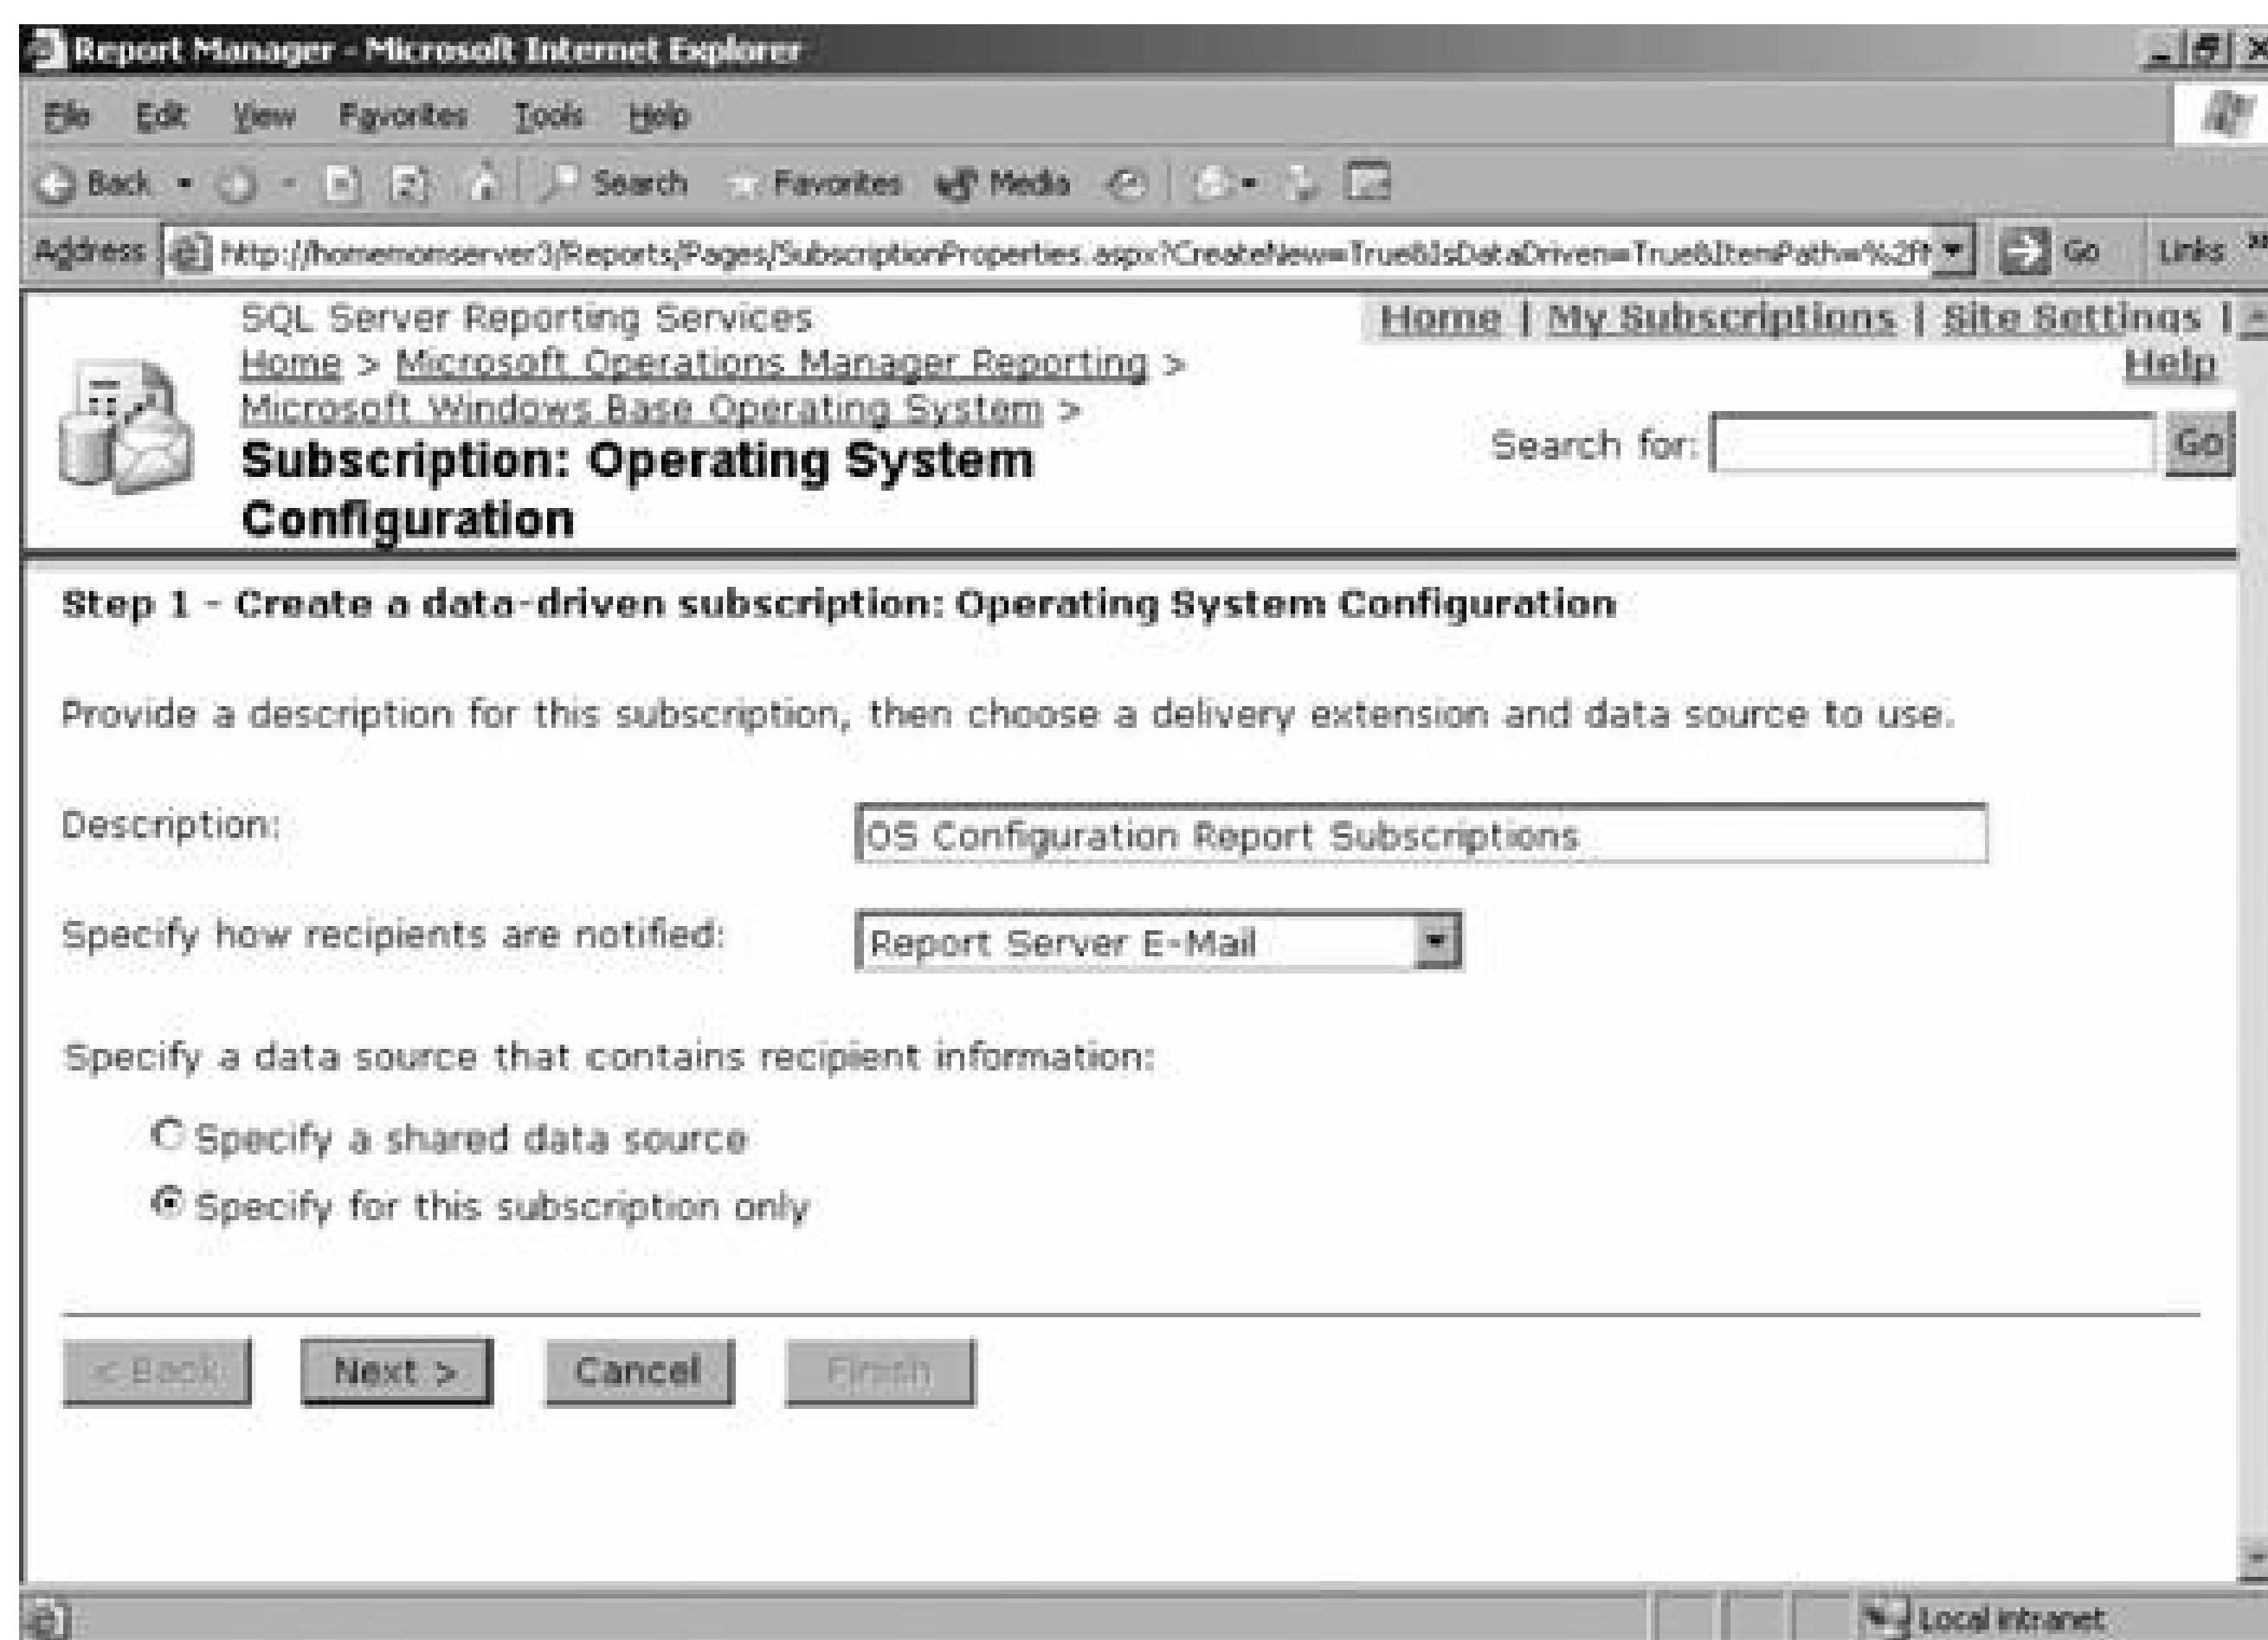


To	Reply-To	IncludeReport	RenderFormat	CompGroup	SortBy
LKFRemoteSiteAdmin1@homelab.lab	reportmaster@homelab.lab	1	PDF	Microsoft Windows Servers	Server Name
LKFCFO@homelab.lab	reportmaster@homelab.lab	1	EXCEL	Microsoft Windows 2003 Servers	Server Name
LKGMOMAdmins@homelab.lab	reportmaster@homelab.lab	1	MHTML	Microsoft Windows Servers	Operating System

8. Open SQL Query Analyzer, make sure the ReportSubscriptions database has the focus, and type **SELECT \* FROM SubscribersInfo**. The query should return the table and all the data that has been entered.
9. Open Report Manager and navigate to Microsoft Operations Manager Reporting      Microsoft Base Operating System      Operating System Configuration report      Subscriptions tab and select "New data-driven subscription."
10. Enter a description on the Step 1 page ([Figure 8-29](#)), select a notification method (in this case, email), and select "Specify for this subscription only" (which is the default). Click Next.
11. Access the ReportSubscriptions database by using the connecting string on the Step 2 page (see [Figure 8-30](#)). It is in the format of data source=<databaseservername>; initial catalog= <database>.

Figure 8-29. Step 1 of a data-driven subscription





In this case, the data source is *homemomserver3* and the initial catalog is ReportSubscriptions, as shown in [Figure 8-30](#). Provide credentials to connect to the database; use the same login credentials you used when you created the database. Select the "Use as Windows credentials" checkbox and click Next.

12. Enter a query to return the list of recipients, in this case **SELECT \* FROM SubscribersInfo**. Click the Validate button on the bottom of the Step 3 page. When the query successfully validates, click the Next button.
13. Map the report values to values in the table on the Step 4 page (see [Figure 8-31](#)). Select the value in the drop-down box that matches the report field on the left-hand side of the page as shown in [Figure 8-31](#). Here, the report field is Reply-To and the mapped value from the table is set to Reply-To. Do this for all appropriate values on the page and click Next.
14. Apply the same report parameter to the table value mapping in Step 5 as you did on the Step 4 page. This will map the CompGroup, SortBy, and SortOrder fields (see [Figure 8-32](#)). Click Next.
15. Select the "On a shared schedule" option on the Step 6 page and choose the "Create and send subscriptions" shared schedule that was previously created.

This returns you to the Subscriptions tab for the report, which now contains an entry for the data-driven subscription.

Figure 8-30. Setting the connection string and credentials

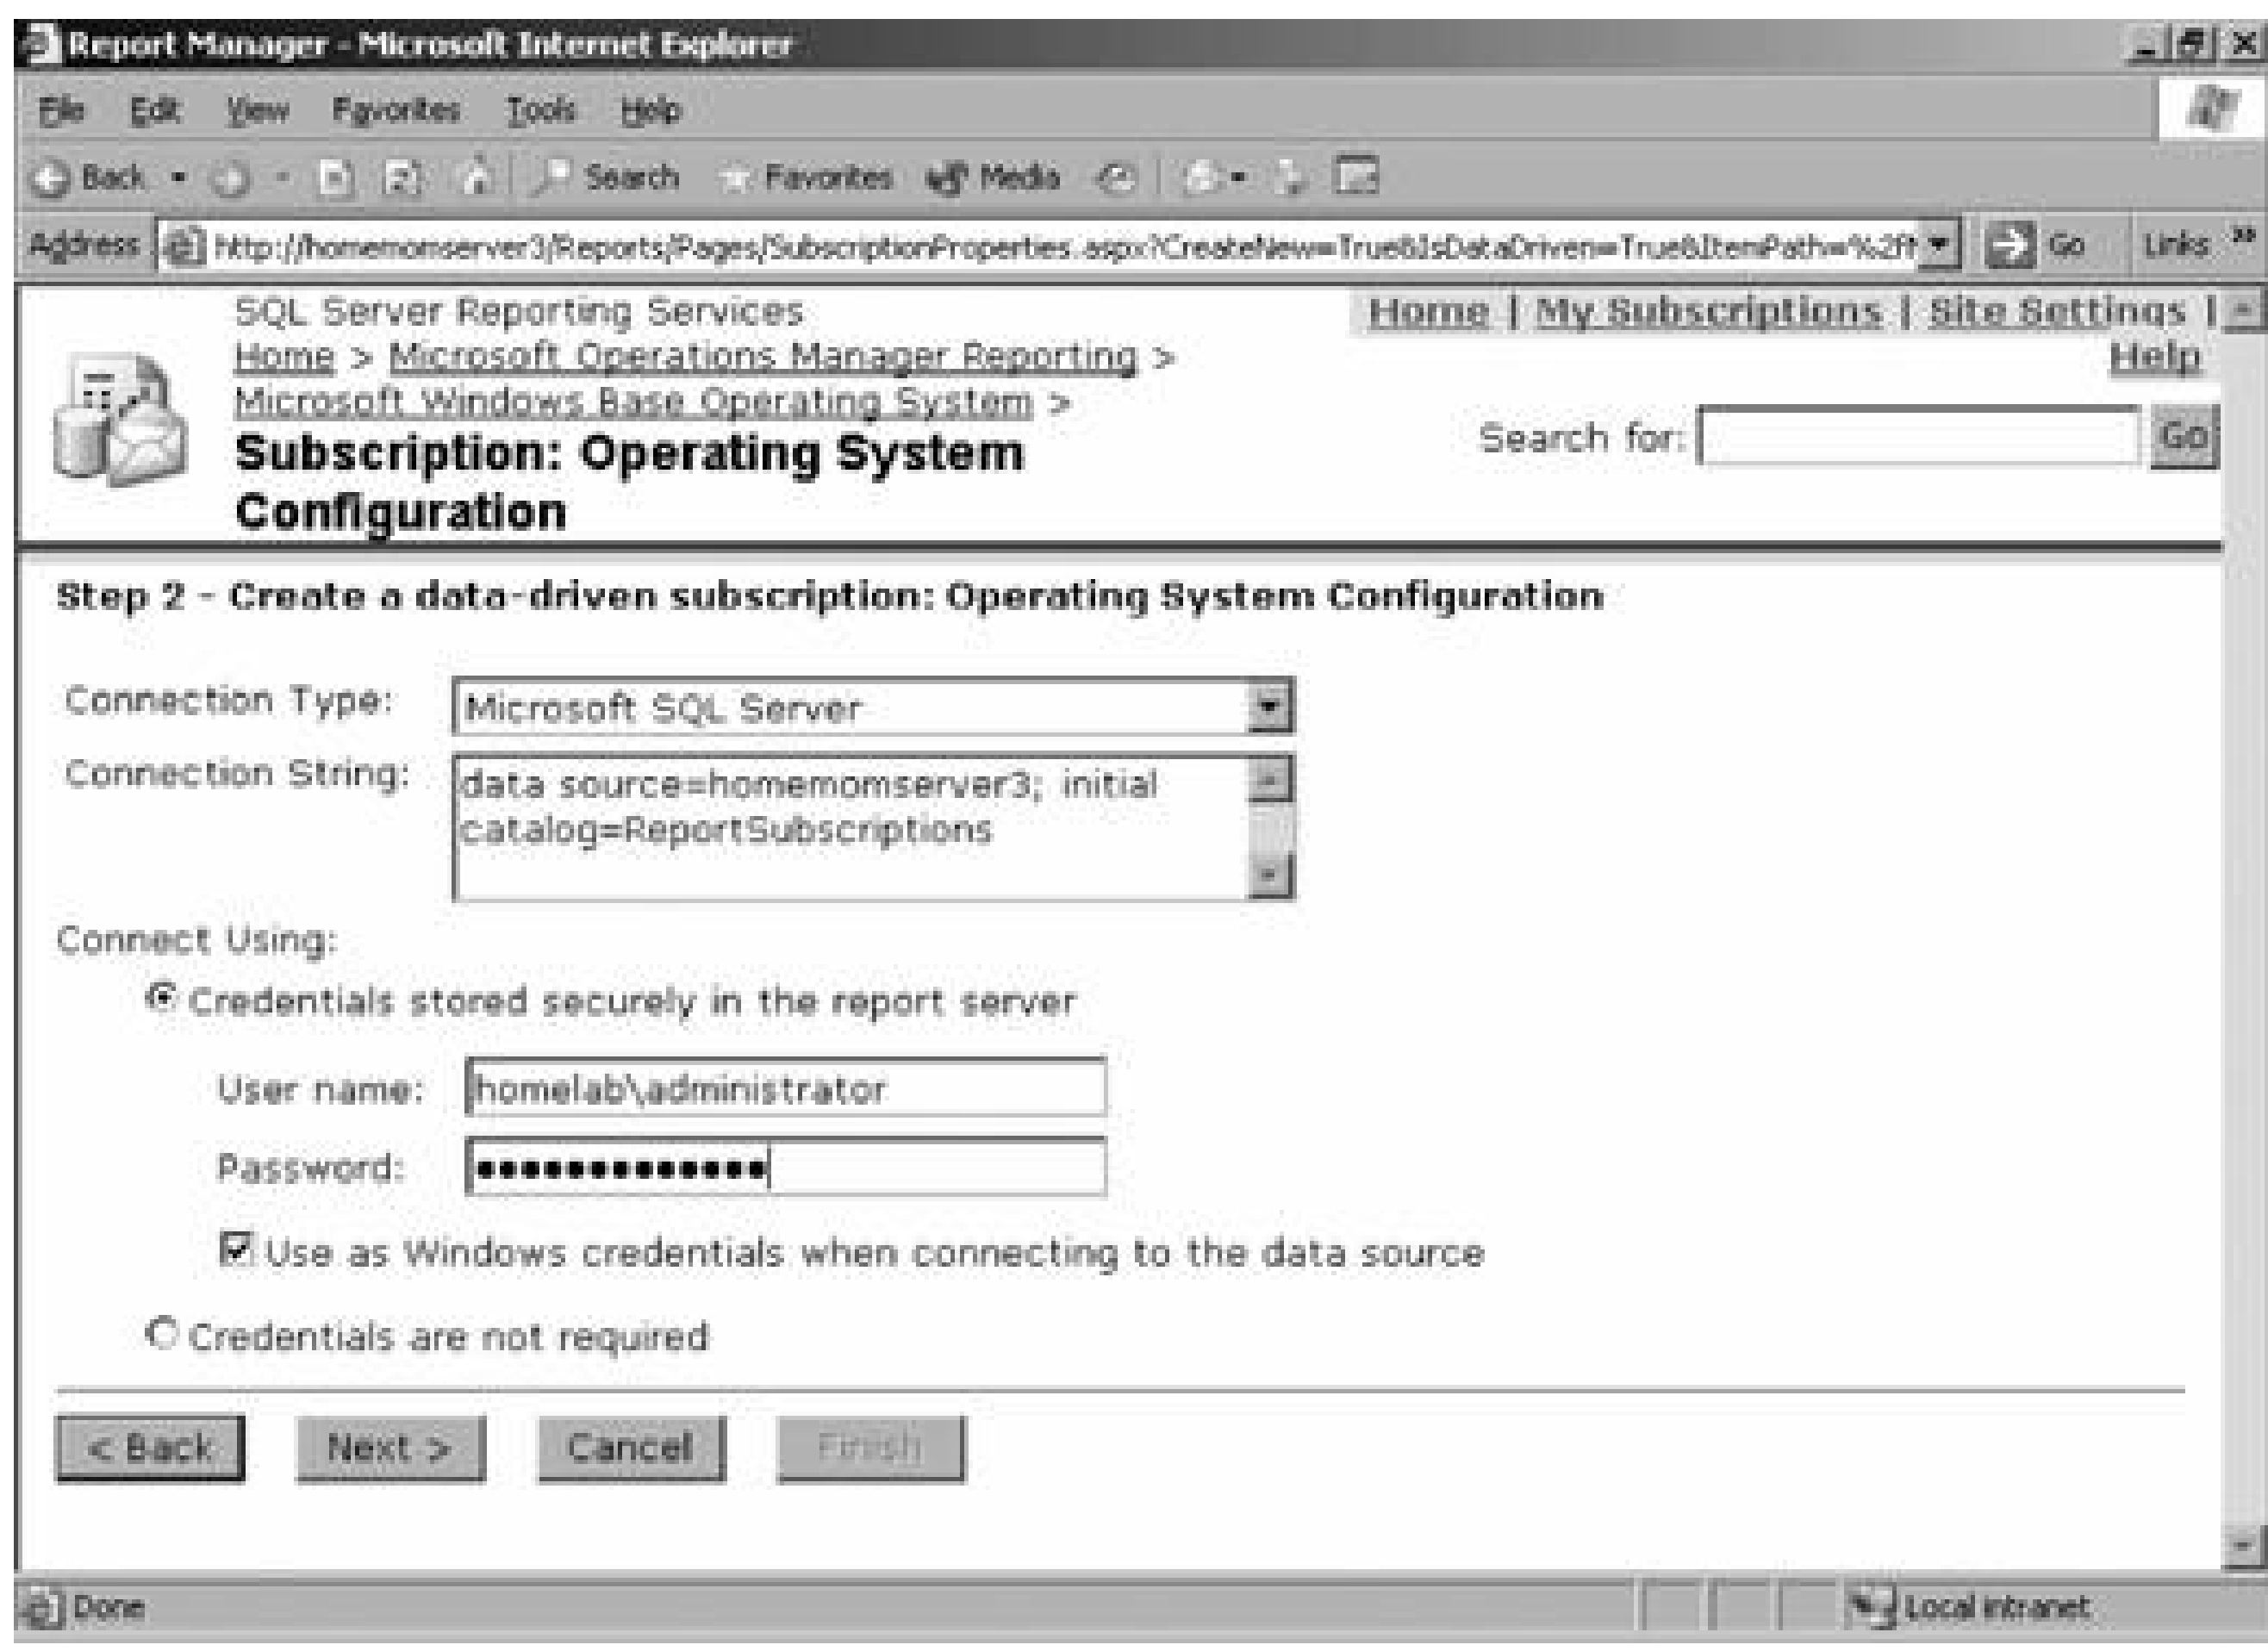



Figure 8-31. Mapping report values to table values

Figure 8-32. Mapping field parameters

SQL Server Reporting Services

Home | My Subscriptions | Site Settings | Help

Home > Microsoft Operations Manager Reporting > Microsoft Windows Base Operating System >

 **Subscription: Operating System Configuration**

Search for:

**Step 5 - Create a data-driven subscription: Operating System Configuration**

Specify report parameter values for Operating System Configuration

**Computer Group**

☐ Specify static value:  ☐ NULL ☐ Use Default

☒ Get the value from the database:

**Sort By**

☐ Specify static value:  ☐ Use Default

☒ Get the value from the database:

**Sort Order**

☒ Specify static value:  ☒ Use Default

☐ Get the value from the database:



## 8.3. Summary

This chapter covered the necessary planning and prerequisites for installing SQL 2000 Reporting Services and MOM 2005 Reporting. Once the reporting solution is installed and data is transferring from the OnePoint database to the SystemsCenterReporting database, you will have to grant your users access to the MOM 2005 Report Manager interface and reports. This is most easily done by granting rights at the root folder level.

Although this is the minimum amount of administrative tasks that you must perform (other than backing up the reporting solution databases as covered in [Chapter 7](#)), you will more than likely be asked to perform a long list of other tasks by your user community.

The next chapter delves into some of the more complex configurations of MOM 2005 management groups, including creating multitiered management group configurations and the tools used to connect MOM with other management frameworks. This is accomplished primarily through the MOM-to-MOM Product Connector (MMPC).

◀ PREV

# Part III: MOM 2005 Enterprise Integration

[Chapter 9, Connecting MOM 2005](#)

[Chapter 10, Extending Monitoring](#)

◀ PREV

# Chapter 9. Connecting MOM 2005

To this point, we've limited the MOM 2005 discussion to the smallest self-contained unit, the management group. In the preproduction and production management group architecture, the two management groups are independent peers of each other. If one became nonfunctional, the other would continue to chug along none the wiser.

Likewise, all discovery and operational data is restricted to the management group. Reporting data, in the form of published reports, can be consumed outside of the reporting console in a variety of formats, but the data itself remains contained in the MOM 2005 Reporting solution. Basically, all MOM data originates within the management group from the agents and ultimately finds its way into the bit bucket via grooming operations, never having left the boundaries of the management group.

This architecture is a great solution as long as your environment isn't too big for a single management group, you aren't required to monitor platforms other than Windows-based servers, or you don't have administrative and security boundaries between you and untrusted groups. Under these conditions and others, the peer management group architecture won't meet your needs. Your MOM architecture may need to scale to multiple management groups to monitor many thousands of agents and still provide a single console for resolving alerts. You may need to collect data from or send data to other management tools like those for Unix or mid-range and mainframe machines.

This chapter and the next introduce the tools and methods that are available for connecting a MOM 2005 management group to other entities, including other management groups. This can be for the purpose of forwarding operational data out of the management group to another entity, or collecting monitoring data for the management group through a route other than agent- or agentless-managed computers. There are many reasons for needing more than one management group, and by understanding those reasons, you can make an informed decision to invest in more management groups.



## 9.1. Partitioning

Dividing operations management tasks across multiple management groups is called *partitioning*. The peer production and preproduction management group configuration is an example of partitioning operations management tasks across management groups for reasons of functionality.

For larger, more complex environments that have tens of thousands of machines to be monitored, or hundreds of machines to be monitored that are separated by slow WAN links, monitoring duties would probably be divided across multiple management groups. Each management group is responsible for monitoring some portion of the environment. That portion could be a section of the network, a geographic location, or a particular set of applications.

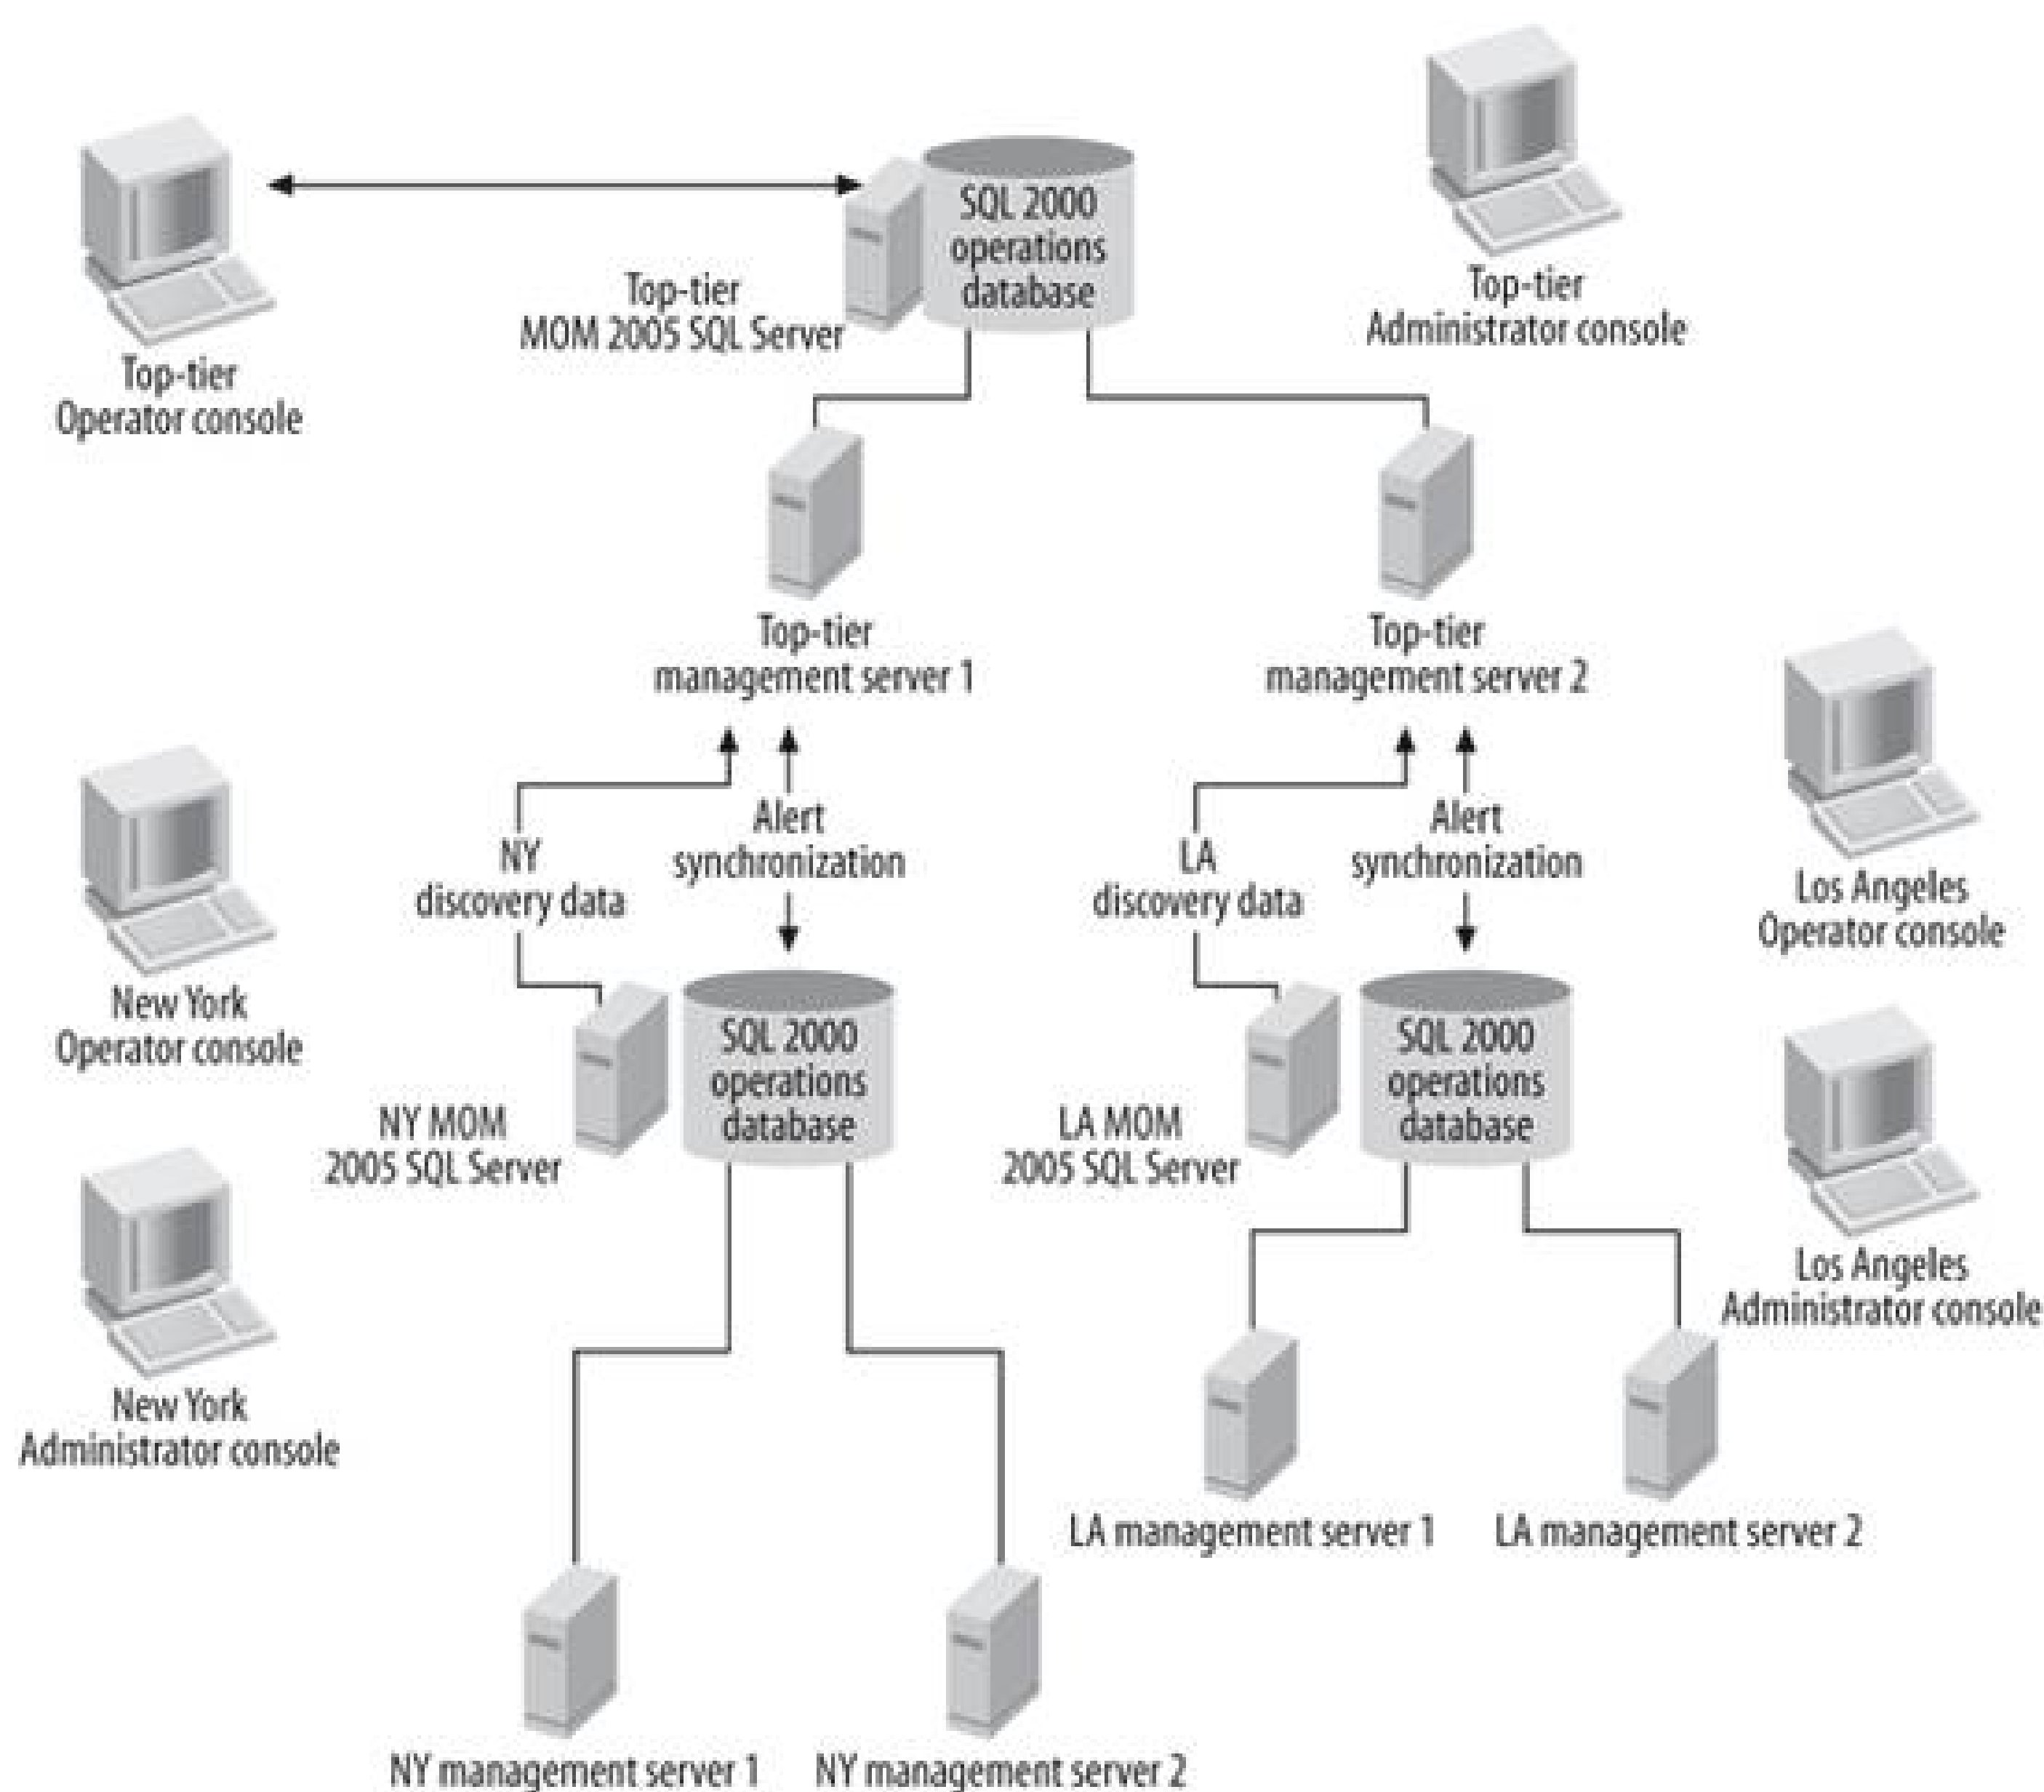
For example, say there is a company called Drippy Sink with two major offices in New York and Los Angeles. There are 1,000 servers to monitor in New York and 1,500 in Los Angeles. From a business perspective, the two locations operate fairly autonomously, but from an IT perspective they are administered as a single environment. All IT administration is performed in the Los Angeles office and the two offices are linked via a single heavily used T1 circuit, so available bandwidth is limited. Even though there are only 2,500 servers to monitor, a quantity that fits comfortably into the capacity of a single management group, the network connection cannot support the traffic of 1,000 reporting agents. So, Drippy Sink deploys production management groups in Los Angeles and New York and a preproduction management group in Los Angeles.

Drippy Sink has three independent management groups that must be managed. There are two separate Operator consoles and no single interface that shows what is going on in each one's environment right now. To make matters worse, there are AD domain controllers, Exchange servers, and shared ISA servers used by staff from both sites. So, for an administrator to monitor AD, she has to watch both Operator consoles, which only show a fraction of the operations management picture.

This example illustrates a situation where centralized monitoring is needed. MOM 2005 supports a tiered architecture, where one destination management group can receive the alerting and discovery data from up to 10 source management groups. Between management groups, this is done via the MOM-to-MOM Product Connector (MMPC). To permit centralized monitoring of distributed management groups, Drippy Sink implements a fourth management group that both New York and Los Angeles production management groups will forward their alerts to. [Figure 9-1](#) is a schematic of this architecture.

The New York and Los Angeles management groups are the *source management groups* that are forwarding their alert and discovery data to the top-tier *destination management group*. When an administrator changes the resolution state of an alert in the top-tier management group, this change can be synchronized back to the source management group that it came from. The diagnostic tools in the top-tier Operator console can be used against the managed computers in the New York and Los Angeles source management groups as well. Configuration and functionality details for the MMPC are covered in the "[Connecting MOM to MOM](#)" section later in this chapter.

Figure 9-1. An example of a tiered architecture at Drippy Sink



This is just one example of a situation in which partitioning is useful. Some of the other reasons for partitioning include capacity, administration, functionality, and configuration.


### 9.1.1. Capacity

The need for additional operational data processing capacity is the number-one reason that companies choose to create multiple management groups and arrange them in a tiered architecture. If there is insufficient processing capacity in your overall system (and this includes available network bandwidth between agents and their management servers), you can deploy additional management groups that are appropriately placed to overcome this. The insufficient capacity issue is another way of saying that you have found a bottleneck between your agents and Operator console that is slowing down operational data processing to an unacceptable level.

The best way to tell if your management group is at, or near, capacity is to track the Alert Logging Latency and Event Logging Latency reports in the reporting console. These measurements take into



account the capacity of the agent, management server, and OnePoint database system as a whole. Many factors influence the performance of this system and determine the volume of data coming in and how fast your management server and OnePoint database can process that data.

These reports are found in the Home  Microsoft Operations Manager Reporting Microsoft Operations Manager Report folder. The Alert Logging Latency and Event Logging Latency reports measure the time difference between when events/alerts are raised and when they were stored in the database. As a rule of thumb, average alert latency should be less than 60 seconds in a well-performing system, but an event latency of between 60 to 90 seconds is also acceptable. You can, of course, decide that longer latency thresholds are acceptable in your environment, which is appropriate if you are accommodating agents across unreliable links and you have adjusted the heartbeat or maximum amount of data to send per second values for the agents (see [Table 3-1](#) in [Chapter 3](#) for these settings).

Before deciding to deploy a tiered architecture, be diligent about reducing capacity bottlenecks in the management servers and, most importantly, the database server. Deal with hardware bottlenecks first—for example, a slow disk subsystem on the database machine, a high percent processor utilization, or a server having a high number of page faults/second, which can indicate a memory bottleneck. Then, look at the percent utilization and overall size of the OnePoint database. If it is near or has exceeded the 60%/18 GB limit, you need to configure more aggressive grooming. Remember, the smaller the database is, the better it will perform.

Other good indicators that can help determine if you have reached capacity are the repeated occurrence of alerts that say: "The incoming MOM server queue is full," "The outgoing MOM server queue is full," and "The outgoing agent queue is full." The outgoing "MOM server queue is full" alert is an indicator that the OnePoint database isn't moving data quickly enough. If you have already tuned the database server and the OnePoint database, and you are still getting these alerts, then it is time to look at deploying additional management groups.

In initial planning stages, there are obvious indicators to keep track of. For example:


- If you know that you are going to deploy more than 4,000 agents, you will need multiple management groups.
- If you are planning on deploying more than 10 management servers (as might be the case if you have more than 9 remote sites, and each will have a management server), then you will need additional management groups.
- If you are deploying a management pack to a computer group that is generating a high volume of data, such as Exchange or Active Directory, or collecting Security events, and you want to guarantee high performance, then consider deploying these management packs to a dedicated management group. This type of deployment requires that the agents be multi-homed between the management group that is performing the high data volume monitoring and the management group that is performing all the other monitoring.

When partitioning for reasons of capacity, remember to build in some extra bandwidth so that spikes in data volume, which can occur with the deployment of a new or greatly revised management pack, don't tank your system. You don't need to wait until the queues are consistently full and alert latency is at 300 seconds to justify the deployment of multiple management groups in a tiered architecture.

## 9.1.2. Administration

MOM 2005 management groups can be partitioned along administrative divisions. This is done when peer IT groups in a company require autonomy in how they run their MOM 2005 management group. Looking at the tiered hierarchy of Drippy Sink, let's say the administrative model changes and the Los Angeles IT shop has full administrative control of the Los Angeles office and the New York IT shop has full control of the New York office. In addition to this, the New York and Los Angeles offices are in separate AD domains in the same tree. [Figure 9-2](#) shows the domain configuration.

Figure 9-2. The Drippy Sink AD domain hierarchy



Because domain administrative responsibilities have been divided between the two locations, they want to divide their MOM management groups along the same lines. This enables Drippy Sink to use different DAS and management server action accounts and different global settings at the management group level and still synchronize alerts and forward discovery data to the top-tier management group for a centralized, all-encompassing view of the environment. In addition, each group of administrators can work with their own alerts at the source management group layer if they choose to. When alert synchronization is configured to be bidirectional, the alert can be resolved in one layer and that change will be reflected in the other layer.

For consolidated reporting purposes, it is possible to configure multiple DTS jobs from the Reporting server that transfer the operational data to the reporting database from the source management group's operational databases. Officially, this is not supported by Microsoft. Microsoft published the Multiple Management Group Roll-up Solution Accelerator, which is available on the MOM web site. It contains information on consolidating reporting data for up to 10 source management groups.

Partitioning for administrative reasons can just as easily be in the same domain rather than between domains. The main requirement is that there are administrative needs, not necessary technical needs, for the partitioning.



For example, another common monitoring practice is the creation of a separate management group for the purposes of collecting security events for auditing and control purposes. This security and audit management group is run by a group of IT security administrators and the managed servers are multihomed into it. Its sole reason for existence is to collect events from the Windows Security event log. Because it has been implemented as an IT audit control, the Windows administration team has no access to it, thus enabling external, independent audit control. This partitioning also makes sense for capacity planning because it is very likely that with auditing enabled the Windows security logs will generate a high volume of data.

To fill out this configuration, the security management group may have its own separate MOM 2005 Reporting solution and forward any alerts to a destination management group.

### 9.1.3. Functionality

The top-tier management group is another example of partitioning management groups along functionality lines. The top-tier group is used as a concentration point for discovery and alert data, as well as a point for centralized management. Because it has this special purpose, it isn't normally used to manage computers, except for its own management servers and database server.

### 9.1.4. Configuration

A MOM 2005 management group has certain configuration settings that are global they apply across all the management servers and managed computers that are in the management group. Most of these settings at the global level can be overridden at the individual, agent, or management server level. These types of global settings are like defaults that you can choose to accept or override on an agent-by-agent (or management server) basis.

However, some of these settings cannot be overridden. These management group global settings are the equivalent of the password group policy that is set at the domain level in Active Directory. If you have established a password expiration policy and applied it to all of the organizational units, then all passwords on all accounts in that domain will expire according to that policy. The only way to accommodate any other password expiration policy is to create a different AD domain and set the desired policy there. Similarly, if requirements exist for more than one configuration for these settings, a different management group must be created to meet those requirements. The following are the management group global settings that that can't be overridden:

*Mutual Authentication (Administrator console      Global settings      Security tab)*

In a MOM 2005 management group, agents and management servers can be required to authenticate each other before data and configuration information is exchanged. This is a Kerberos v5 authentication and serves the same purpose in MOM as it does in ADit blocks against man-in-the-middle attacks.

*Block Legacy Agents (Administrator console      Global settings      Security tab)*

This setting must be considered if you have MOM 2000 or 2000 SP1 Agents reporting to a MOM 2005 management server. If enabled, the MOM 2005 management servers will not communicate with any agents that were created before MOM 2005 was installed.

*Communications Ports (Administrator console      Global settings      Communications tab)*

By default, the management server and agents communicate over port 1270. This communication is encrypted for added security. This port can be changed; however, to allow for managed computer failover between management servers, all management servers and agents must use the same port. If there are requirements for a few managed computers to use a different port, as might be the case if your firewall disallows port 1270 communication, then an additional management group must be created. You would then manage all of the internal, port 1270 communicating machines in one group and the non-port 1270 communicating machines in another management group.

*Enable/Disable Server-Side Responses (Administrator console      Global settings      Security tab)*

The management server can execute an action against a managed computer based on a generated alert or some other criteria. These responses are called server-side responses because they are run from the management server. The MOM 2005 management packs have predefined server-side responses that are not affected by this setting. Custom server-side responses can be developed and placed in an event, alert, or threshold rule. The Enable/Disable Server-Side Responses global setting only enables or disables the custom developed server-side responses. The default setting is to disable the custom server-side responses. If there is a management pack that must execute custom server-side responses but is not allowed to by security policy, then an additional management group must be created to accommodate it.

Any one of these reasons is enough to justify additional management groups. For example, if you need to create an additional management group for added capacity, that additional management group will also serve a specific function as well to monitor the extra servers that can't be monitored by the first management group alone. Use that observation when you need to build your business case to justify the additional cost.



## 9.2. Connecting MOM to MOM

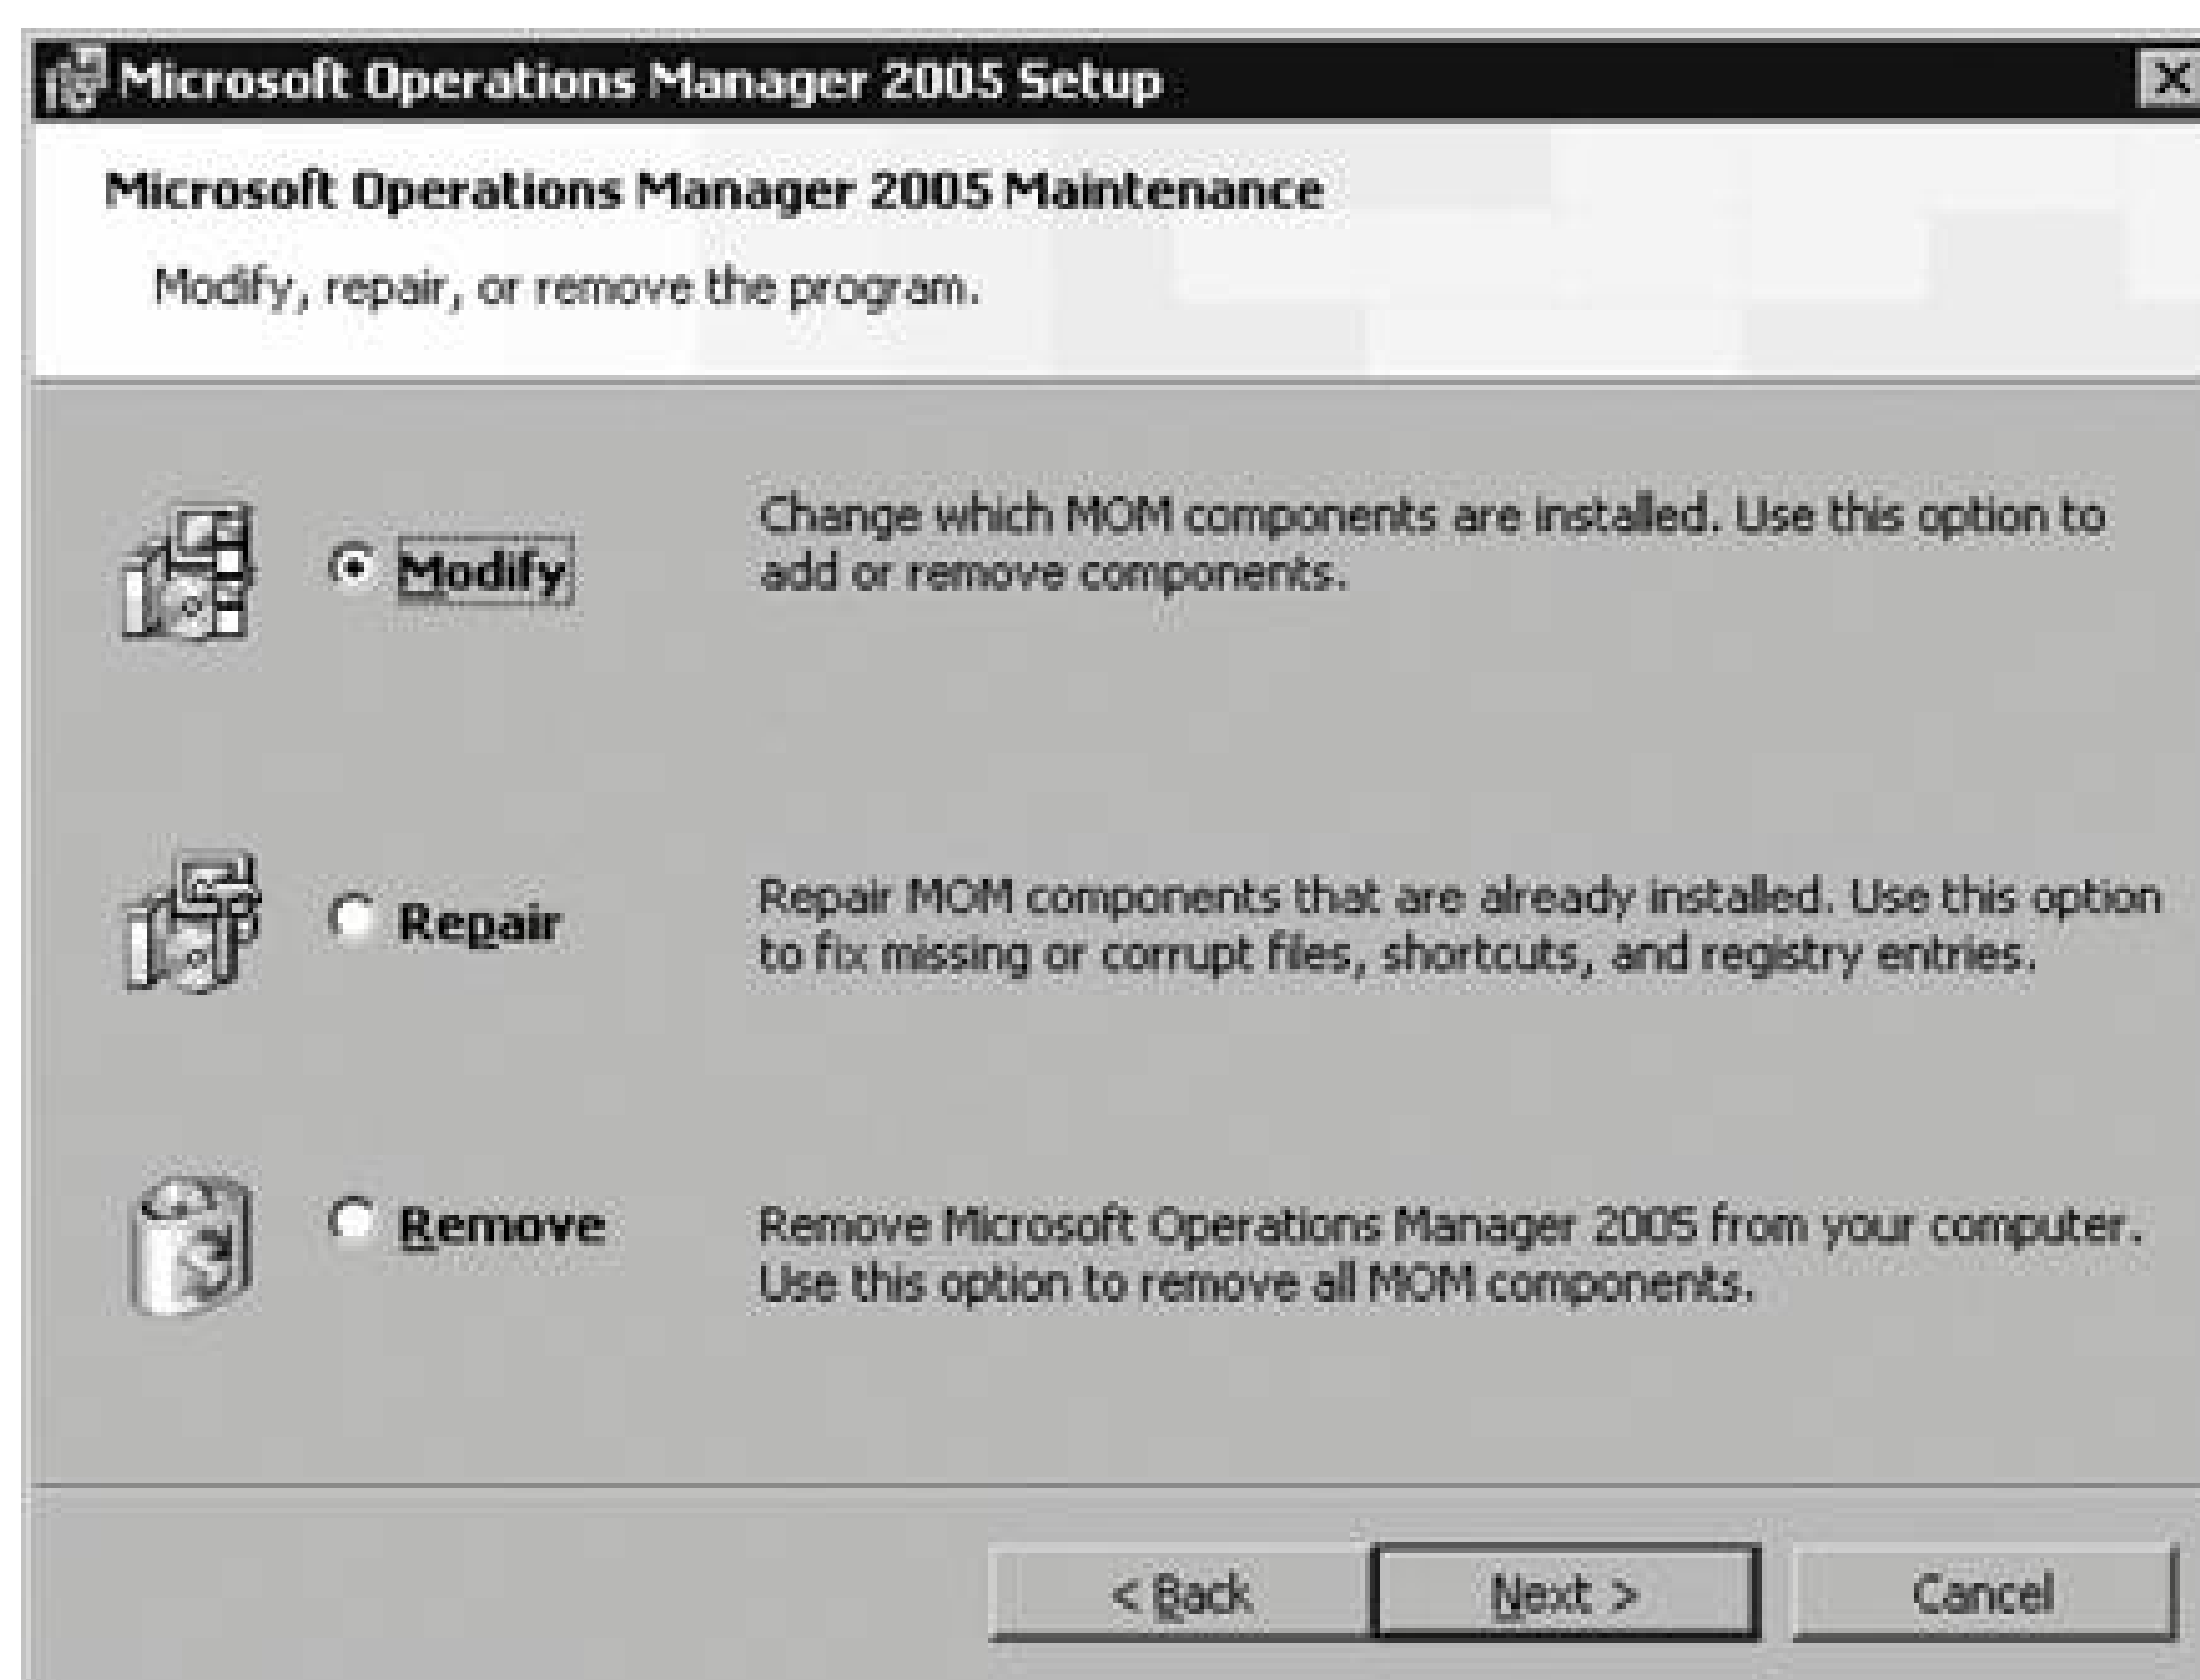
The MOM-to-MOM Product Connector is one type of product connector, and it is only used to send discovery and alert data between MOM 2005 management groups. There are other types of product connectors for MOM that allow bidirectional communication between a MOM management group and other management products or trouble ticketing systems. For example, Microsoft ships a MOM to a Tivoli TEC connector, a MOM to an HP OpenView connector, and a MOM to an HP Network Node Manager connector. These connectors are built on the MOM Connector Framework (MCF), which is one of the APIs that is included in the MOM SDK. There are other APIs in the MOM SDK, namely the MOM Managed Code Library (MCL) and the MOM Runtime interface.

The MMPC is a web service and runs as a service on management servers. The executable is *momconn.exe*. A web service is like a web page that contains a good deal of logic, but it has no interface to browse and can be accessed over a normal protocol (HTTP). It is used for passing data back and forth between applications. To make use of an MMPC, you must create a connector on the management servers in the source management groups. Because all of the configuration and operational information for a connector is stored in the OnePoint database, all of the management servers in the source management group can support the MMPC. This means that if the source management server that is hosting the MMPC fails, the MMPC function will failover to another source management server.

### 9.2.1. Creating a Tiered Configuration

Creating the source-to-destination relationship between two management groups requires that the MCF and MMPC are installed in the source management group, and that the MCF is installed on the management servers in the destination management group. These components can be installed when the management group is created or any time afterwards, as is the case here. The management groups used here to illustrate the installation and configuration process are in the same domain and use the same DAS account and Microsoft action account. The installation starts with the destination management group running the MOM 2005 Setup wizard. Proceed through the Welcome page to get to the Modify, Repair, Remove page as shown in [Figure 9-3](#).

Figure 9-3. Choosing to modify an existing installation of MOM 2005



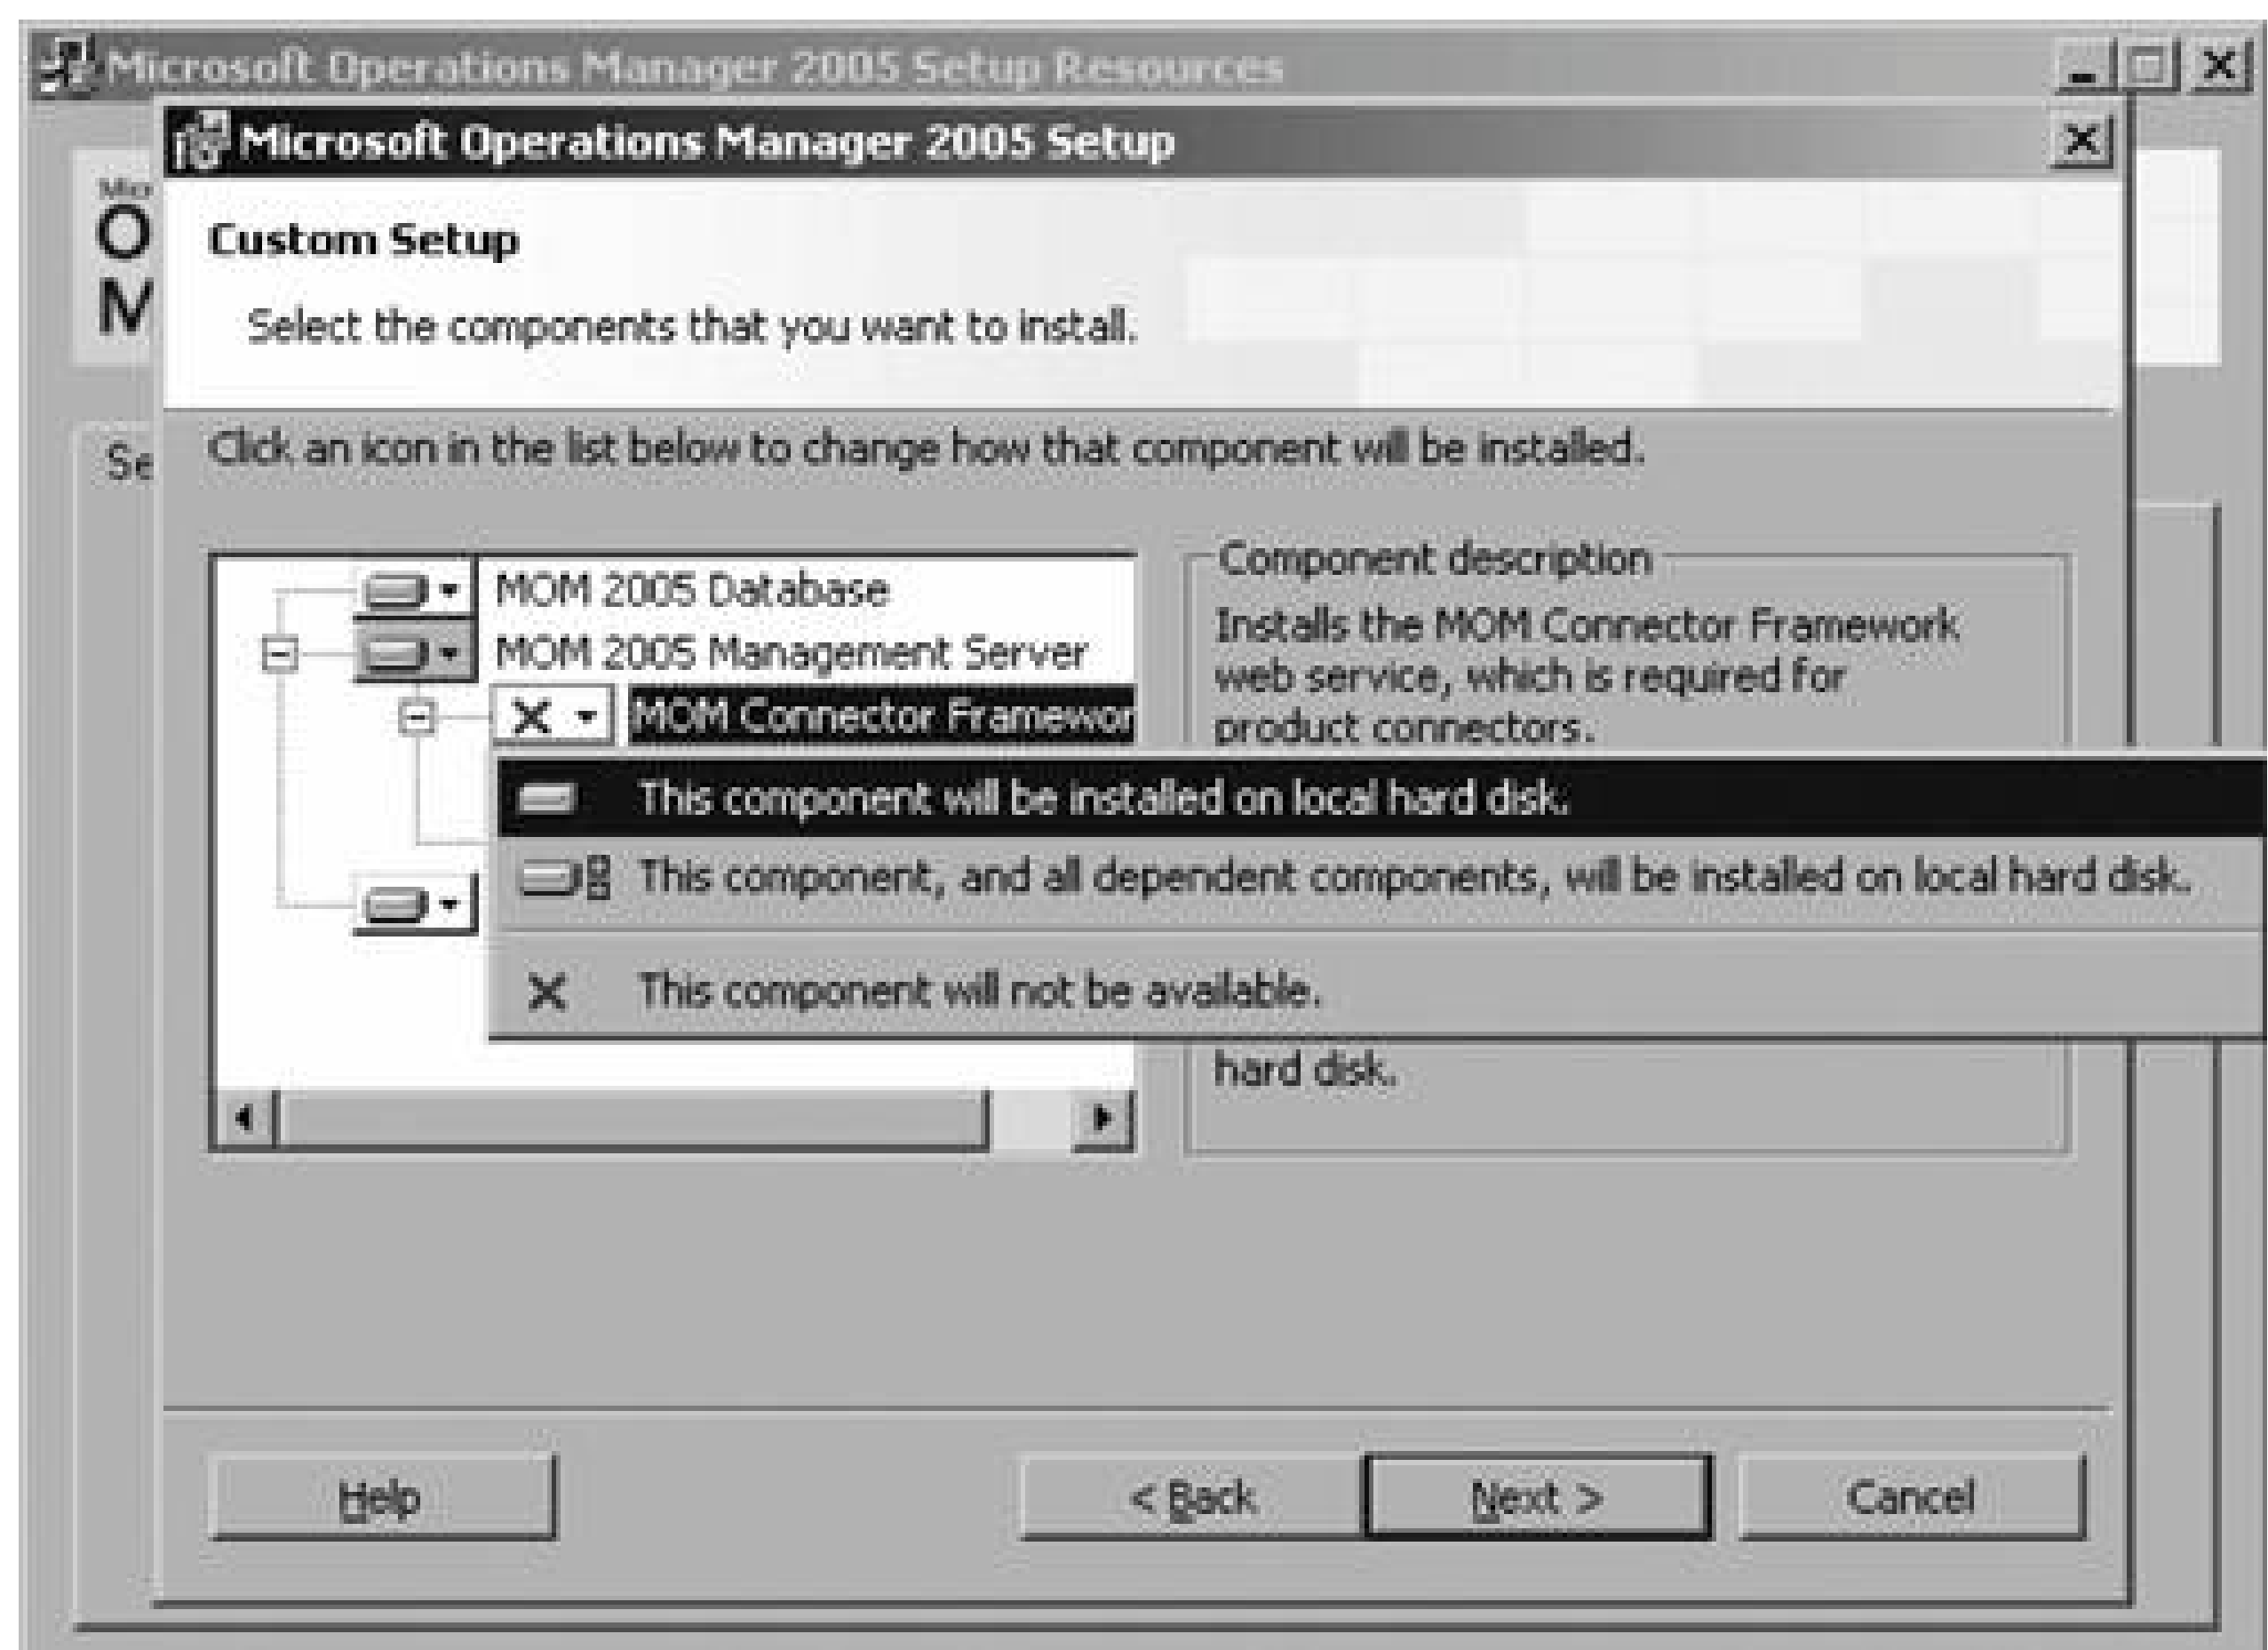
When you launch the setup wizard, the Modify, Repair, Remove page is presented because the wizard detected the previous installation. You can launch the setup wizard from the MOM 2005 CD or from the Control Panel's Add and Remove Programs feature; however, if you do the latter you may still be required to provide the source media. Select the Modify option and click Next to bring up the MOM 2005 Custom Setup page in [Figure 9-4](#).

Under the MOM 2005 Management Server object, select to install the MOM Connector Framework with the "This component will be installed on local hard disk" option, and click Next. If you select the "This component and all dependent components" option, the MMPC will be installed as well. This doesn't hurt anything, but the MMPC is not required on the destination management group so it serves no purpose since only the MCF is required.

Click Next to proceed through the prerequisite checker, which will pass because MOM 2005 has already been successfully installed. Click Next again to bring up the Ready to Modify page where you click Install. The MMPC uses the MCF methods `Connector.GetData` and `Connector.InsertAlert` to retrieve new alerts or updated alerts from the source management group and insert them into the OnePoint database in the destination management group. To access both the source and destination OnePoint databases, the MMPC uses the DAS account from the source management group. To get access to the destination OnePoint database, the source DAS account must be placed in the MOM Service local group on all of the destination management servers. As I mentioned in [Chapter 2](#), the MOM Service group is empty by default, so the source DAS account should be the only one in there. However, if you have created a domain-level MOM Service group, you can add the source DAS account to that and then add the domain group to the local group; either way will work. The source and destination management groups do not have to be in the same domain to set up this relationship. The only requirement is that you add the DAS account of the source management group to the MOM Service local group on the destination management servers.



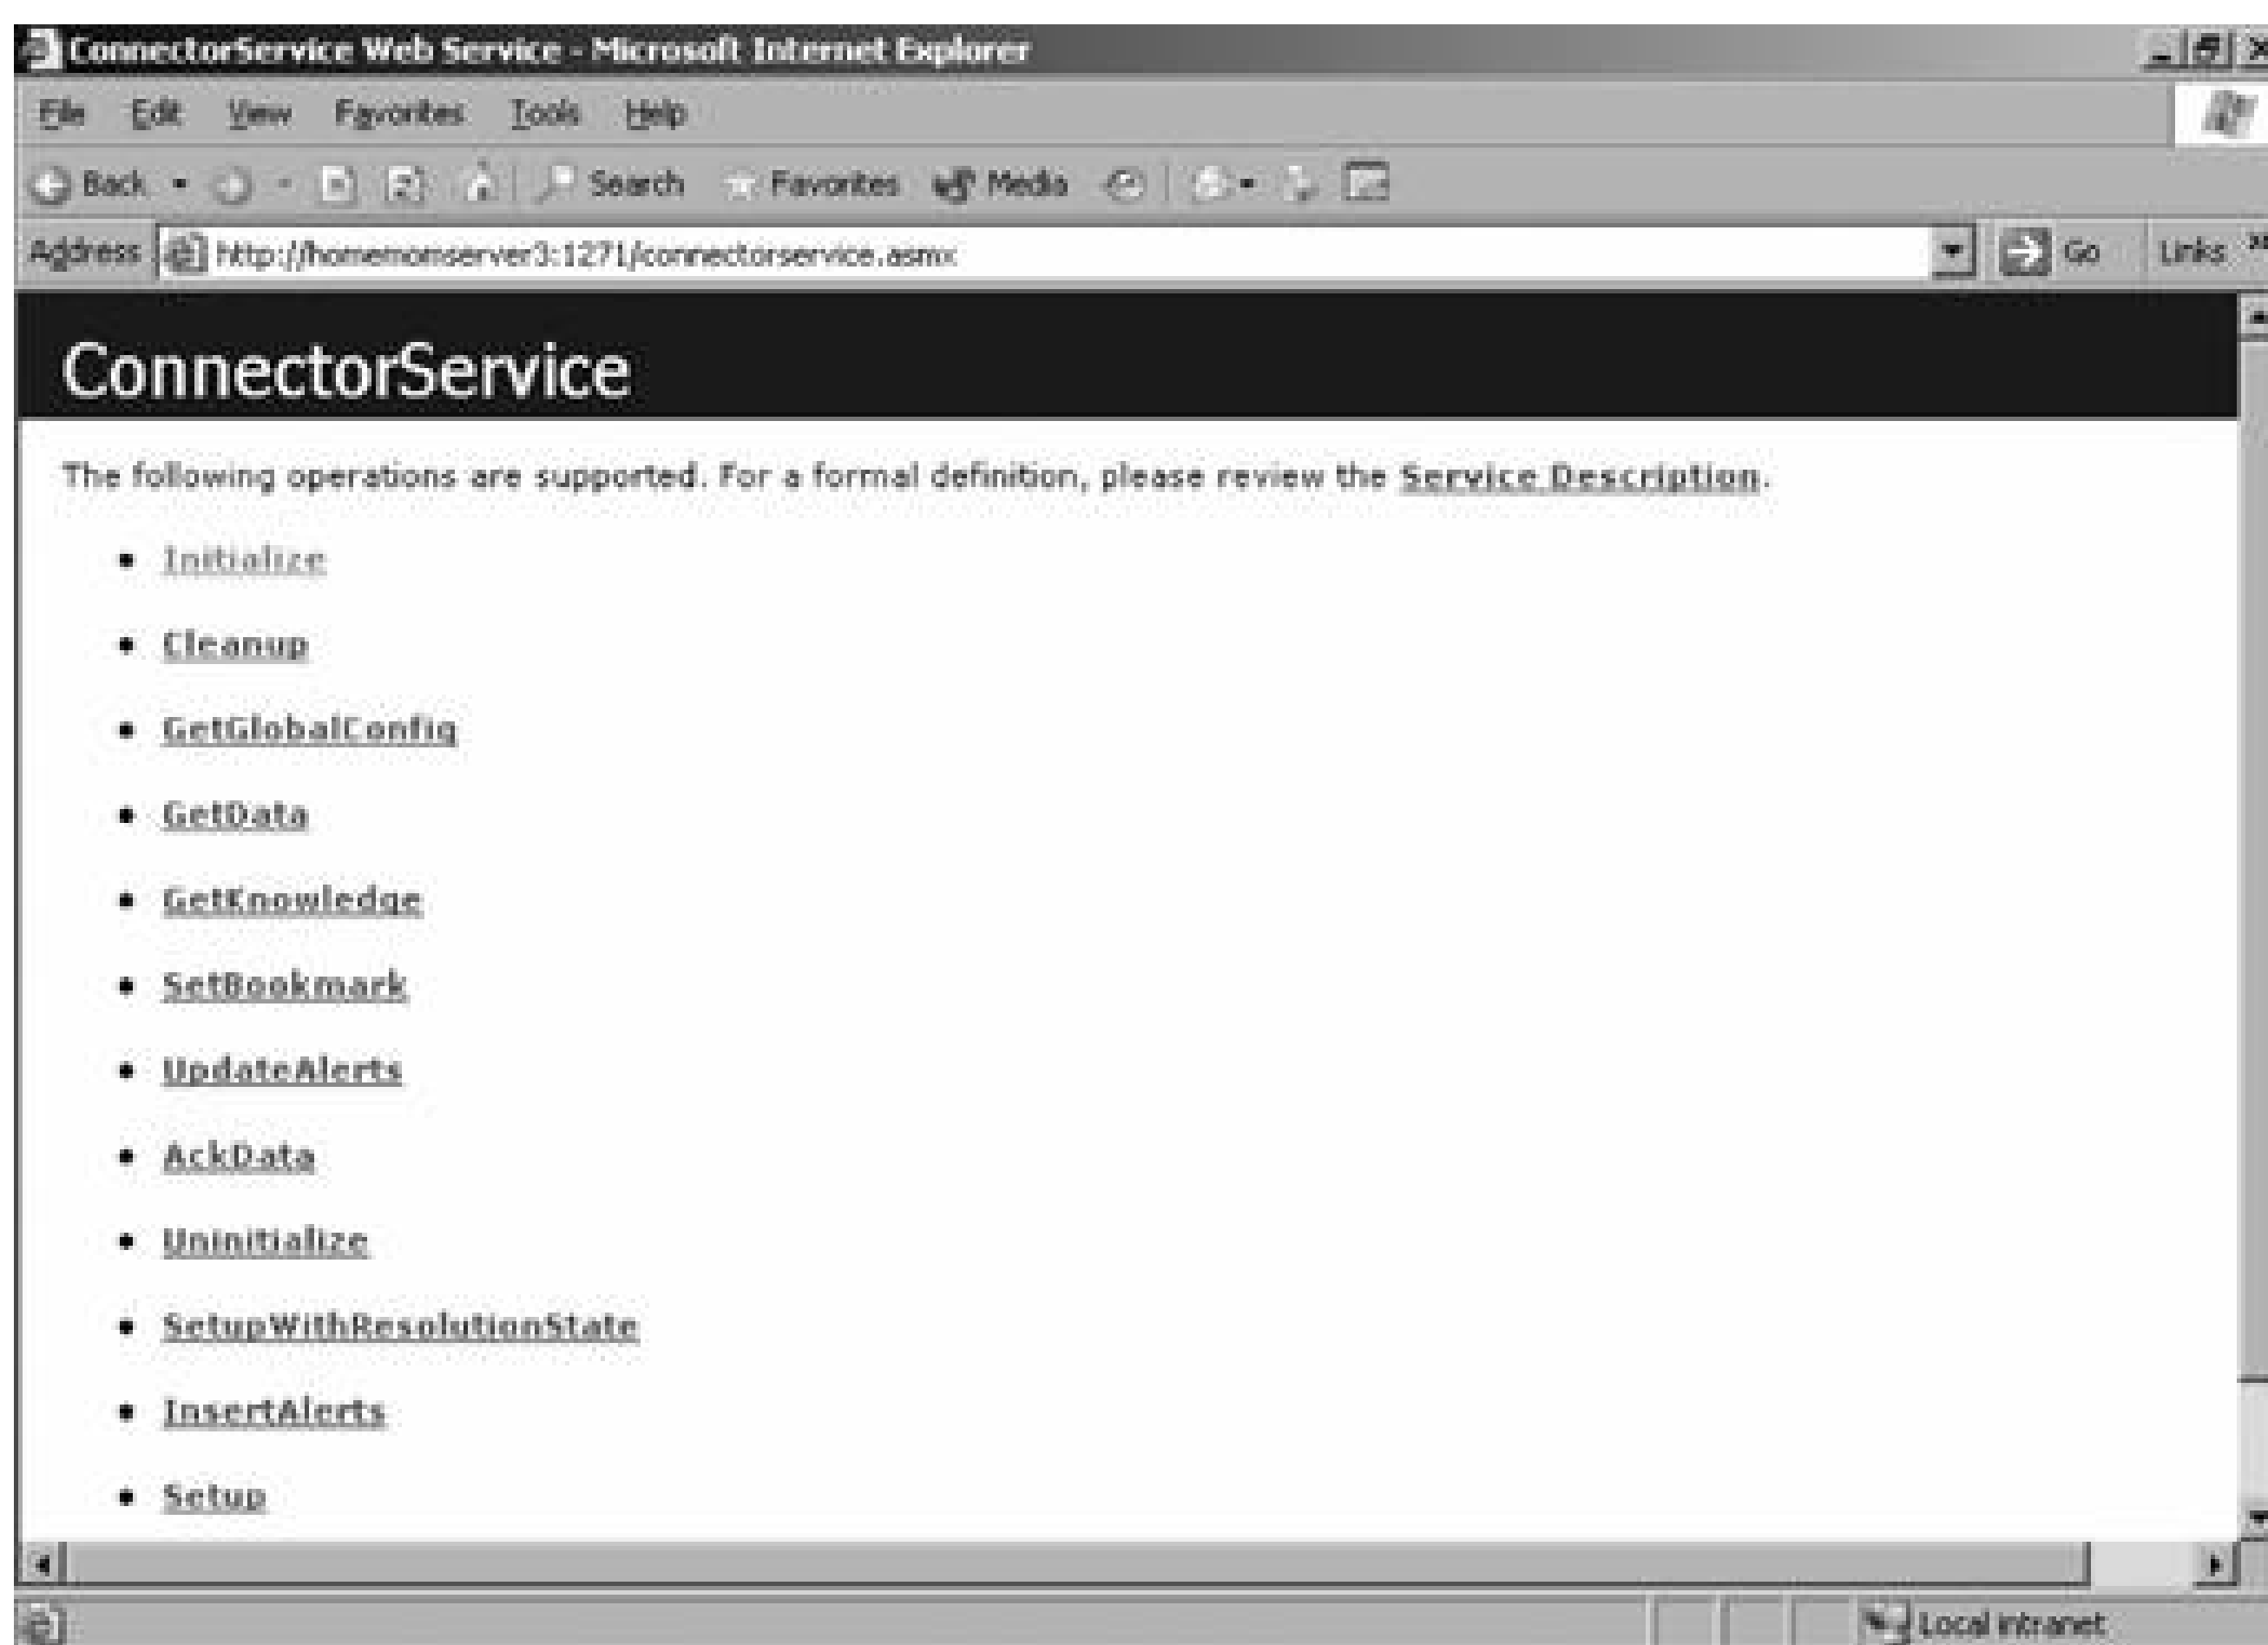
Figure 9-4. Installing the MCF on the destination management group



Once the MCF installation completes, you should confirm that it is functioning correctly. On the management server where you performed the installation, open IE and browse to <http://<computername>:1271/connectorsevice.asmx> where *<computername>* is the name of the server that you installed the MCF on. The number 1271 in the address signifies the default port that the MCF service is listening on. If everything is working correctly, you will get the page shown in [Figure 9-5](#).

If the source-to-destination connection traverses untrusted network connections, you can SSL-secure the communication by placing a web site certificate on the destination management group's management servers.

Figure 9-5. MCF Connector Service page



[Figure 9-5](#) shows some of the methods that are available in the MCF; by clicking on them you can bring up pages that show the code that makes up the method. For example, [Figure 9-6](#) shows some of the code page for the `Getdata` method.

Most of the work done by the MMPC is through the `Setup`, `Initialize`, `GetTData`, `AckData`, and `UpdateAlerts` methods:

### `Setup`

This method is called when you are creating an instance of the connector on the source management server. It sets up the connector.

### `Initialize`

This method starts the connector that was created when `Setup` was called.

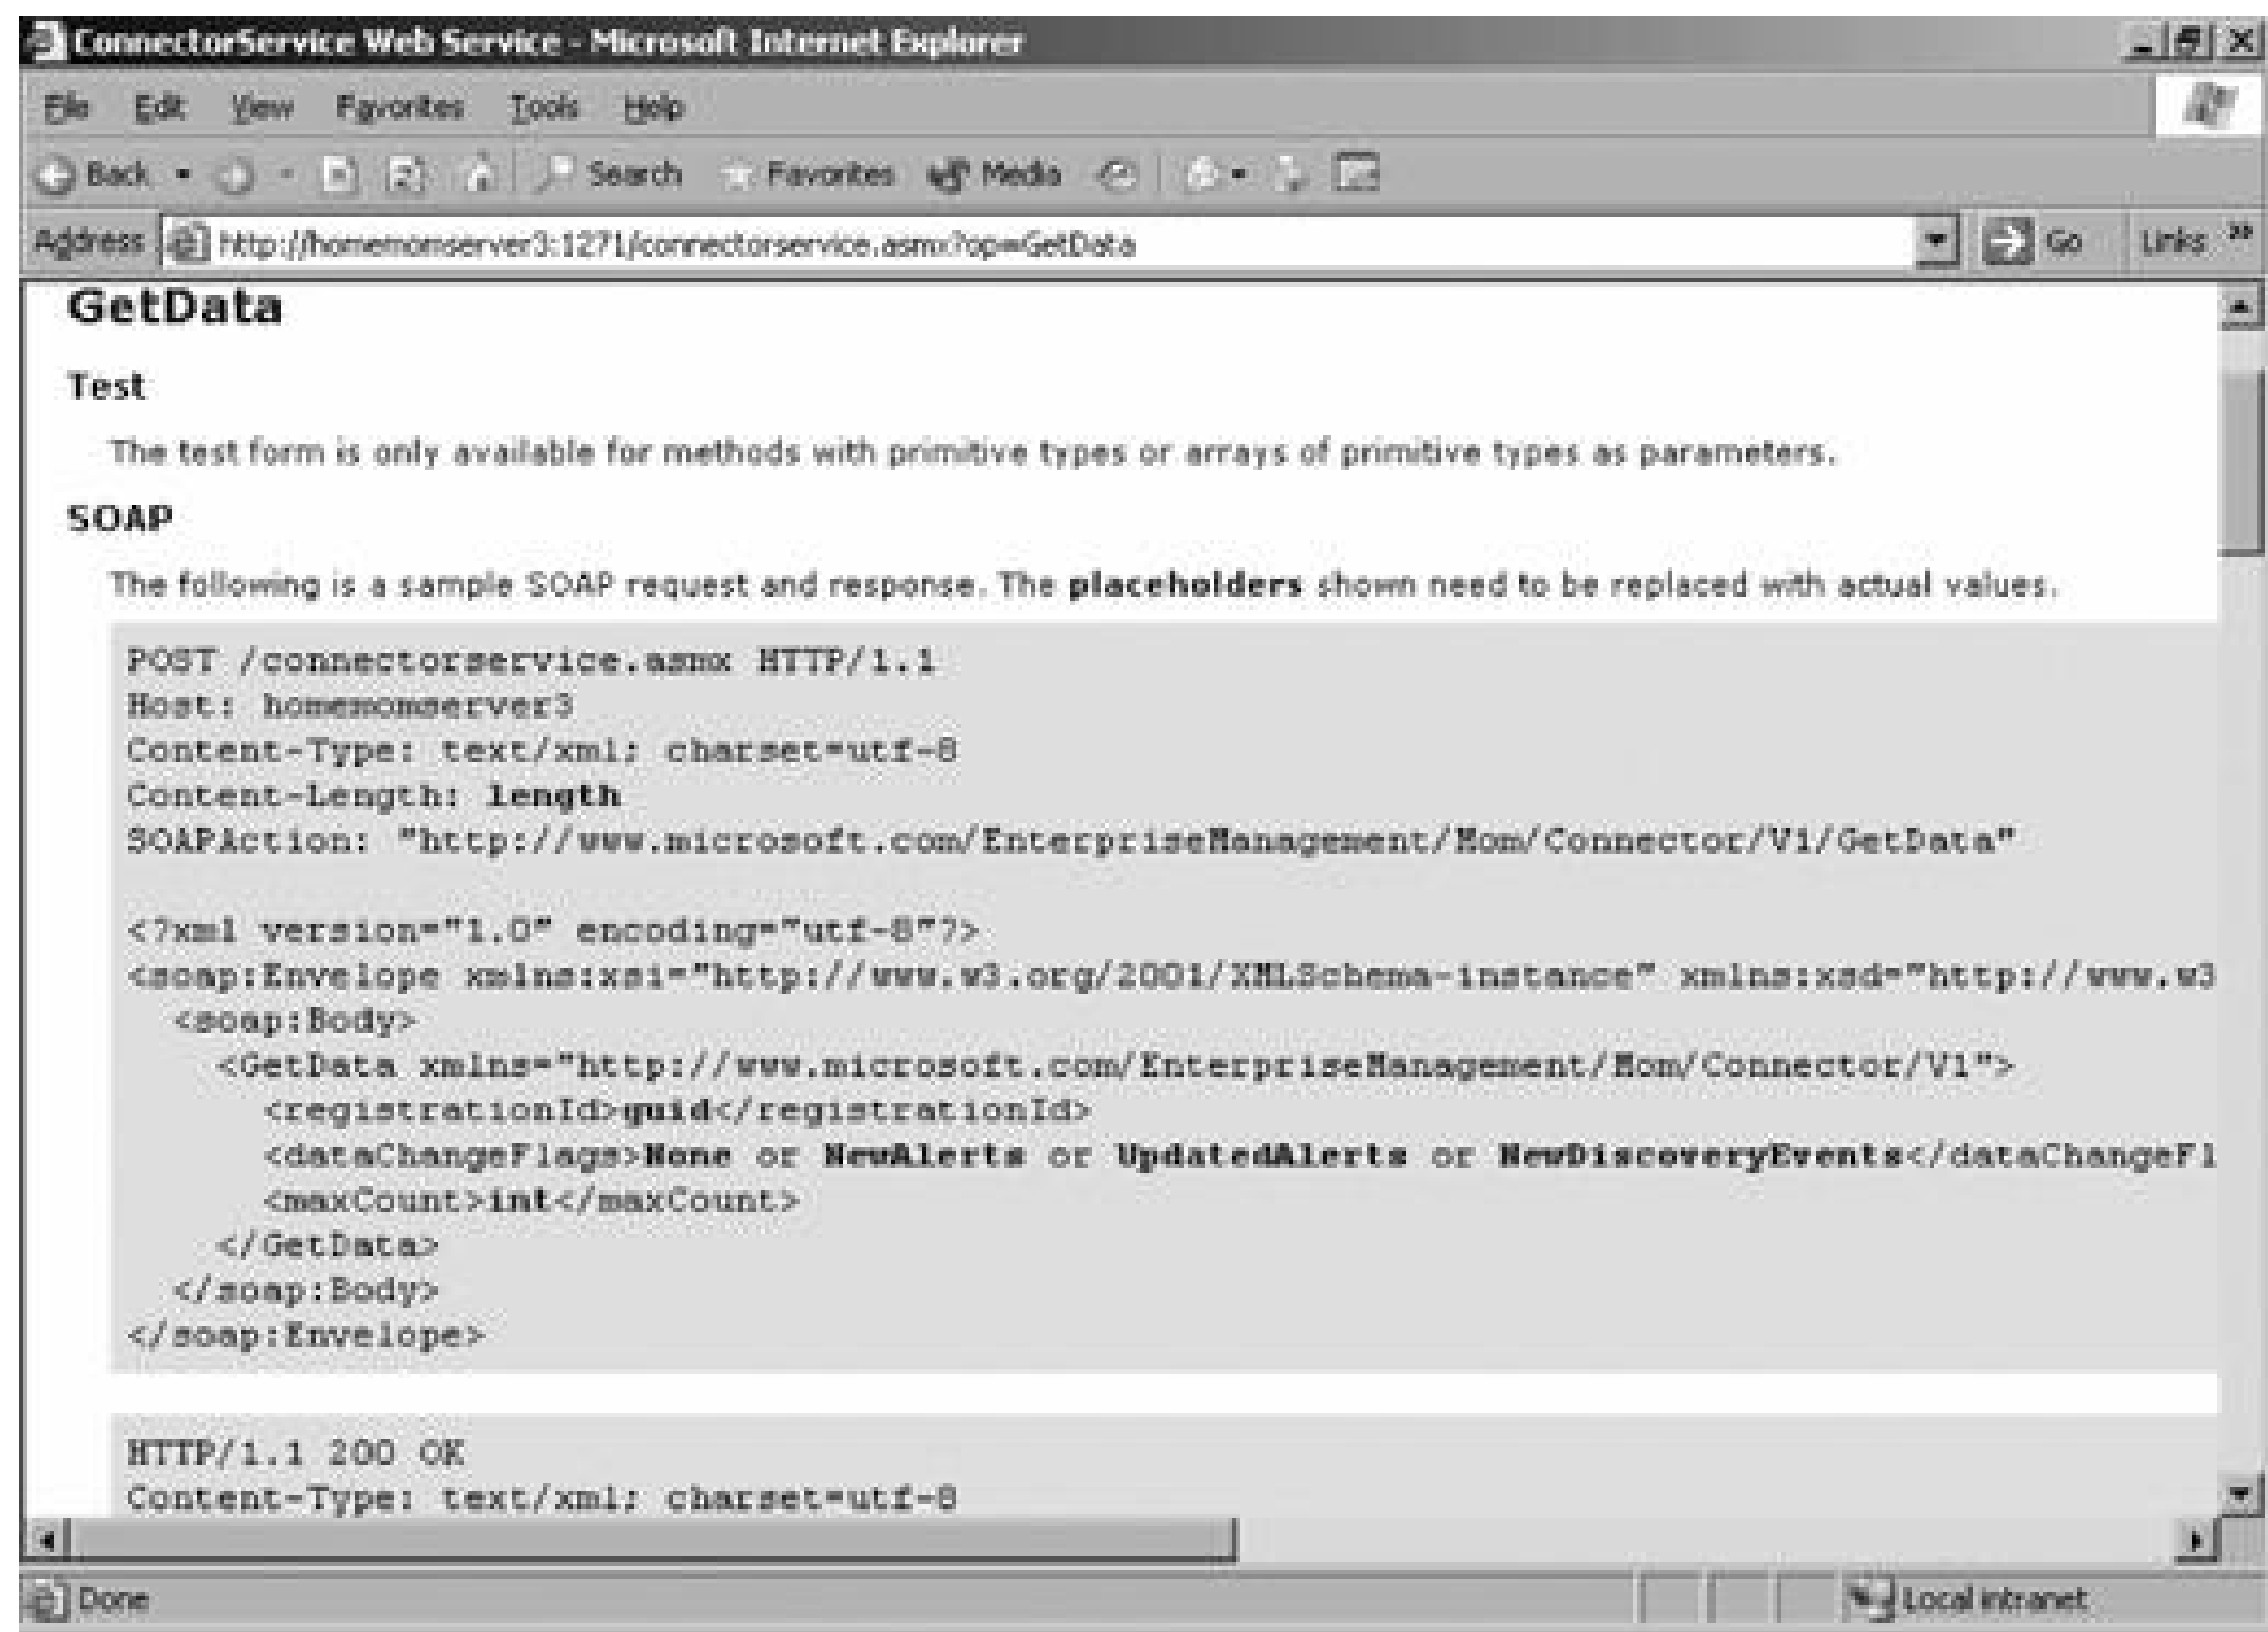
### `GetData`

This call retrieves alerts and discovery data from the source management group OnePoint database. This method is called every 30 seconds, making it the most commonly called method in the MCF.

### `AckData`

After an alert has been inserted into the destination group OnePoint database, the destination management group acknowledges receipt of the alert back to the source management group.

Figure 9-6. Some of the code page for the GetData method



### UpdateAlerts

This method updates the information in an alert in either the source or destination management group when it is changed in the connected management group.

The next step in creating the tiered configuration is to install both the MFC and the MMPC on the source management group's management servers. For a previously existing management server, launch MOM 2005 Setup and select the Modify option. On the custom setup page (see [Figure 9-4](#)), instead of selecting the "Install this component" option, select "This component, and all dependent components." Then finish the Setup wizard just as you did for the MCF setup on the destination management group management servers.

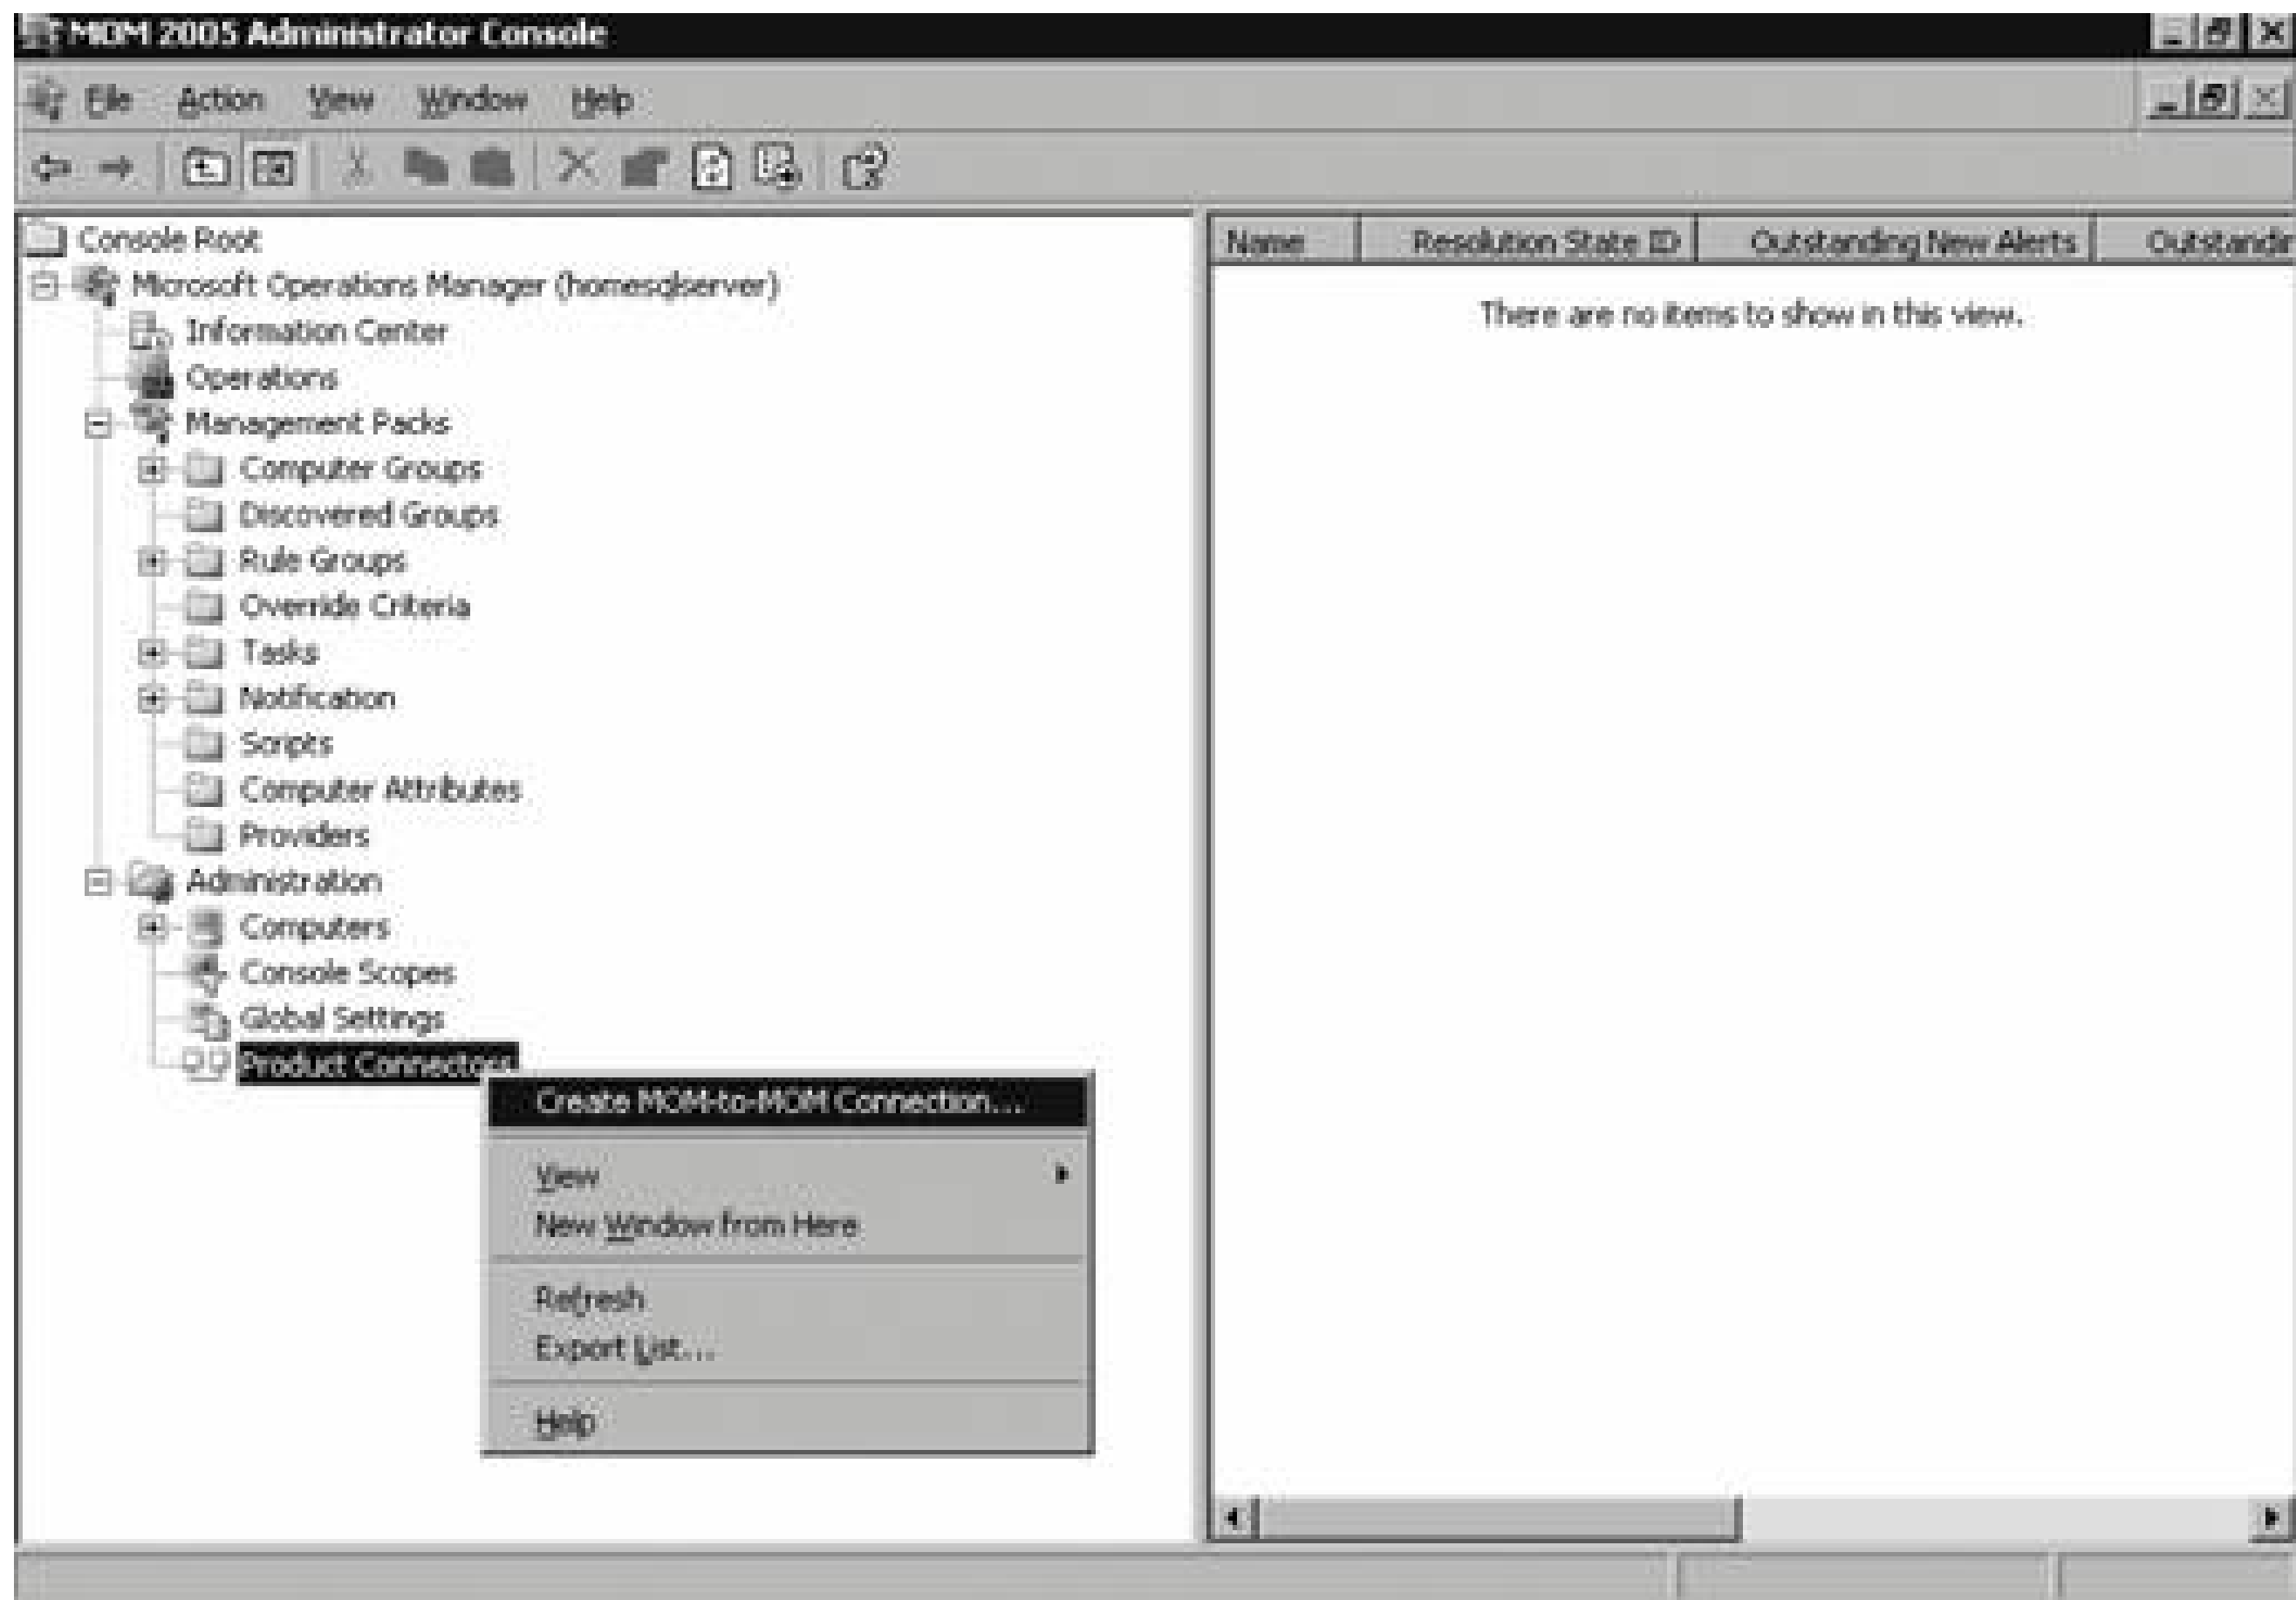
Now you are ready to create the connector between the source and destination management groups. This example uses the LKFProdMG *homesql/server* as the source and a new management group called TopTier that has *homemomserver3* as its management server. The TopTier management group does not monitor any servers other than its own management server and database server, in keeping with the special purpose of the destination management group.

To create an MMPC, open the Administrator console and navigate to Administration Product



Connectors and right-click to bring up the context menu. From the context menu, select Create MOM to-MOM Connection as shown in [Figure 9-7](#).

Figure 9-7. Creating a new MOM-to-MOM product connector



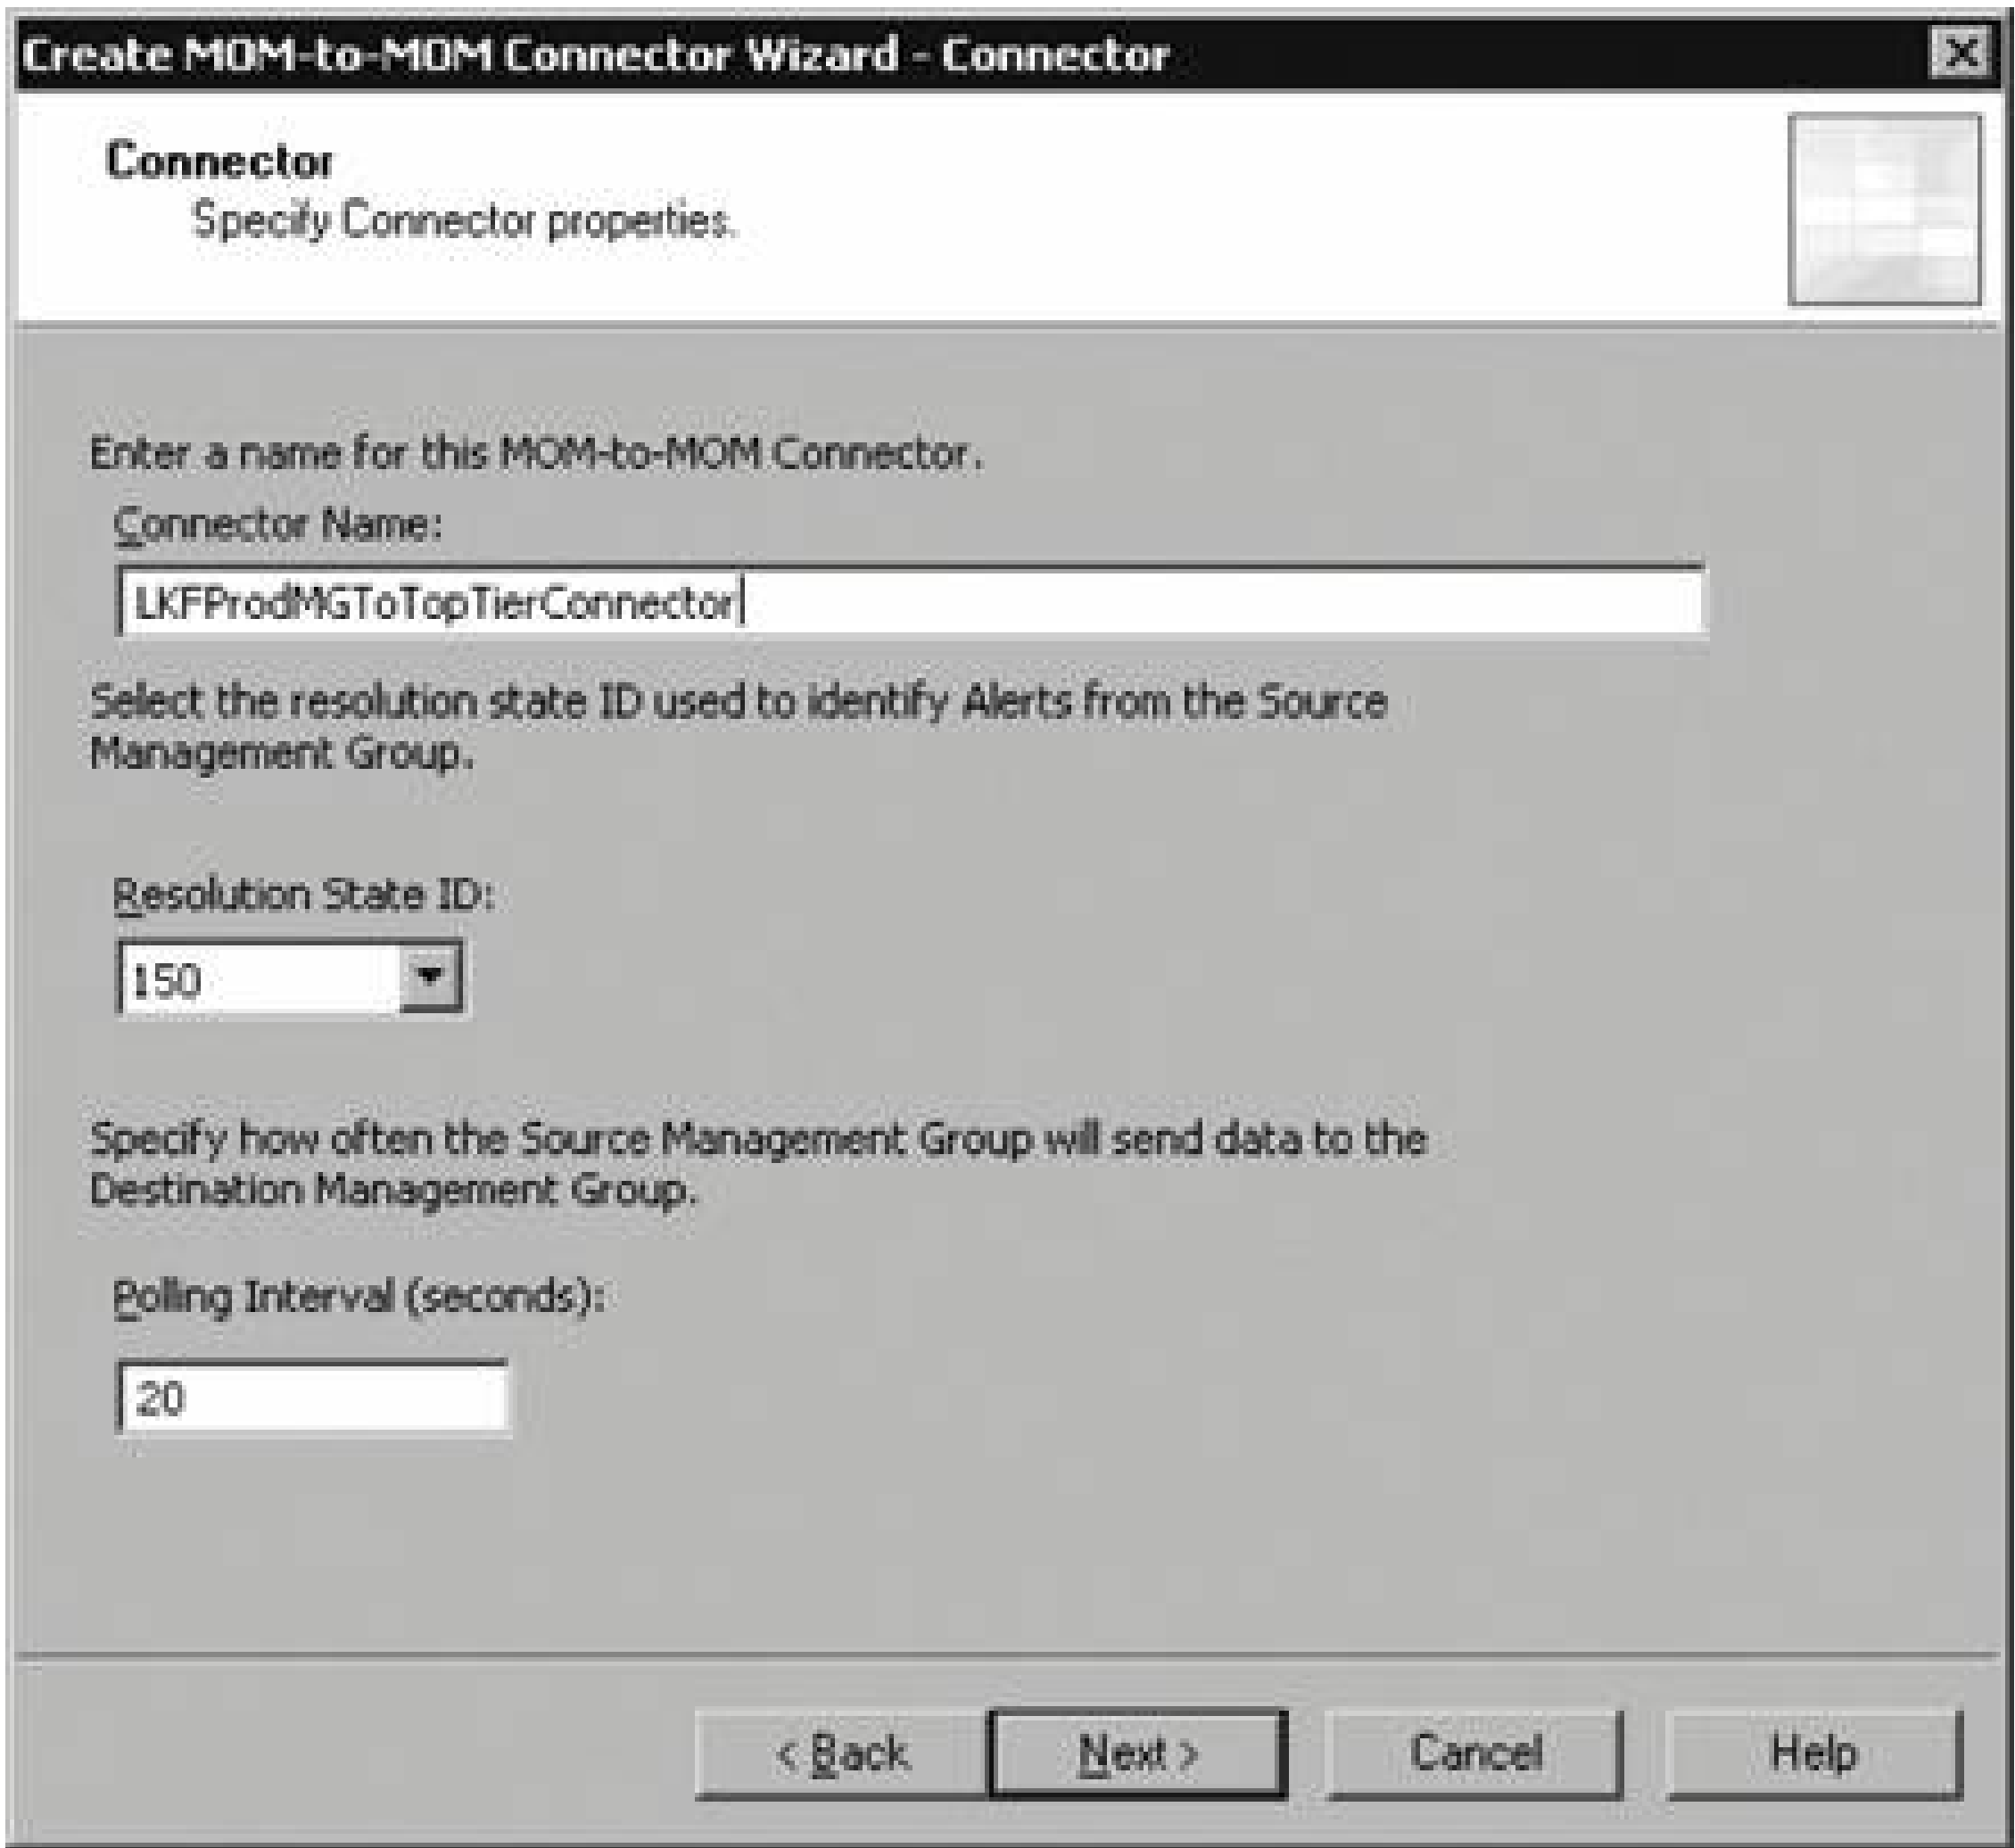
This starts the Create MOM-to-MOM Connector wizard. Click through the Welcome page to get to the Connector page, shown in [Figure 9-8](#).

Provide the connector a meaningful, distinguishing name, in this case, "LKFPProdMGToTopTierConnector." In the Resolution State ID field you can accept the default of 150 or supply your own value. Each MMPC registers a specific resolution state in the configuration portions of the OnePoint database. When any alert is set to this state, which appears as the connector name in the Resolution State drop-down list, that alert is forwarded to the connector for synchronization with the destination management group. Remember from [Chapter 5](#) that every alert exists in one of several resolution states at any given point in time and that internal to MOM, the states are represented as numerical values.

For example, the resolution states of New and Resolved are reserved and have the values of 0 and 255 respectively. When you create a new MMPC, a new resolution state is created with the default value of 150. So when you forward an alert to the destination management group, you can either manually select the MMPC name (e.g., LKFPProdMGToTopTierConnector) in the Resolution State list, or you can programmatically set the resolution state to 150 via an alert rule. This way, the forwarding will occur automatically, or you can have the alerts created with the default resolution state set to the name of the connector.

The Polling Interval refers to how often synchronization bits flow across the wire between the source and destination management group. It has nothing to do with the frequency that the `Connector.GetData` method is called, which is every 30 seconds. This means that if you leave the default settings, it could take up to 50 seconds for an alert to be synchronized from source to destination *after* it appears in the source Operator console.

Figure 9-8. Setting the properties for the MOM-to-MOM connector



Click Next to bring up the Add MOM Master Management Group page; see [Figure 9-9](#).

At this point in creating an MMPC, you only have the option to designate one management server or web service. Later in the setup, you can enter additional web services to allow for failover at the destination management group level. Click Next to open the Forwarding Properties page; see [Figure 9-10](#).

Synchronization across an MMPC can be one way (from source to destination), or two way (starting with the source, going to the destination and then back again) when the alert has been changed in the destination management group. By selecting only the first checkbox in the Alert Forwarding Properties section, you configure one-way synchronization. Selecting "Receive alert updates from Destination back to Source" configures two-way synchronization.

In the "Forward Discovery Information from Source to Destination" section, by selecting the "Forward

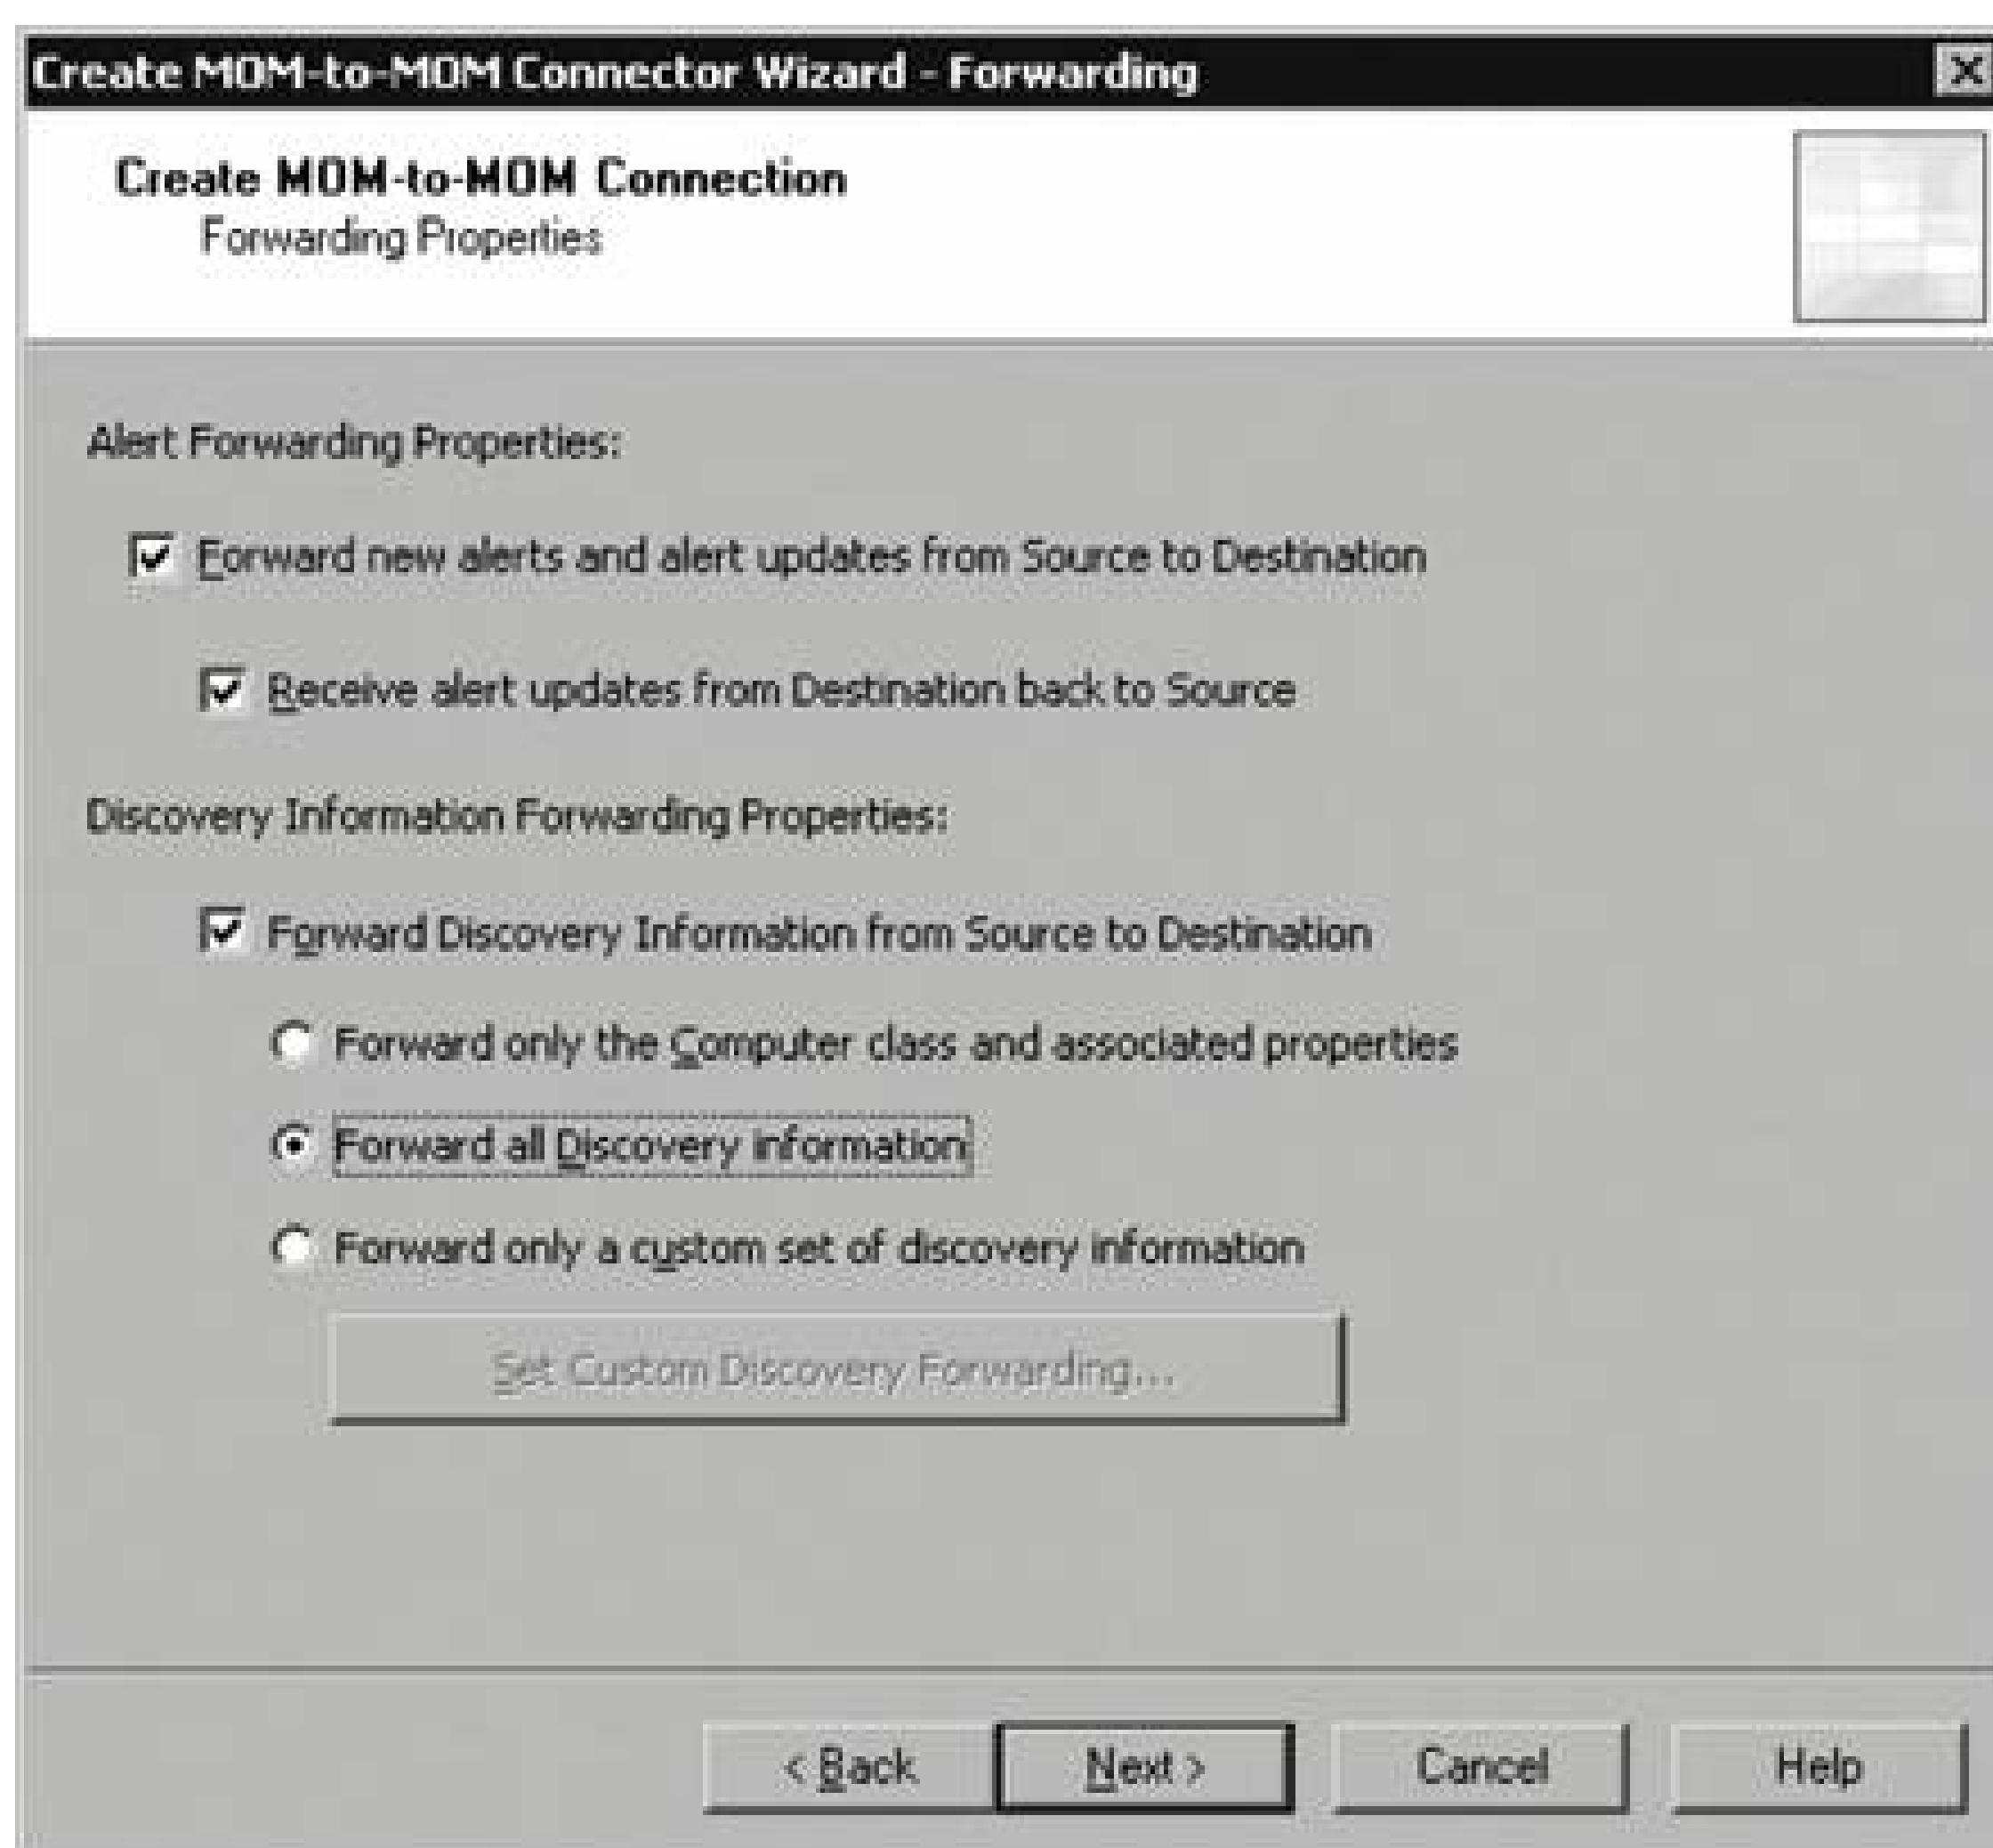
all Discovery information" option, you will have the richest, most congruent experience between the two Operator consoles.

Figure 9-9. Identifying the destination management group's management server that the connector will communicate with



Figure 9-10. Configuring alert forwarding properties and discovery information forwarding





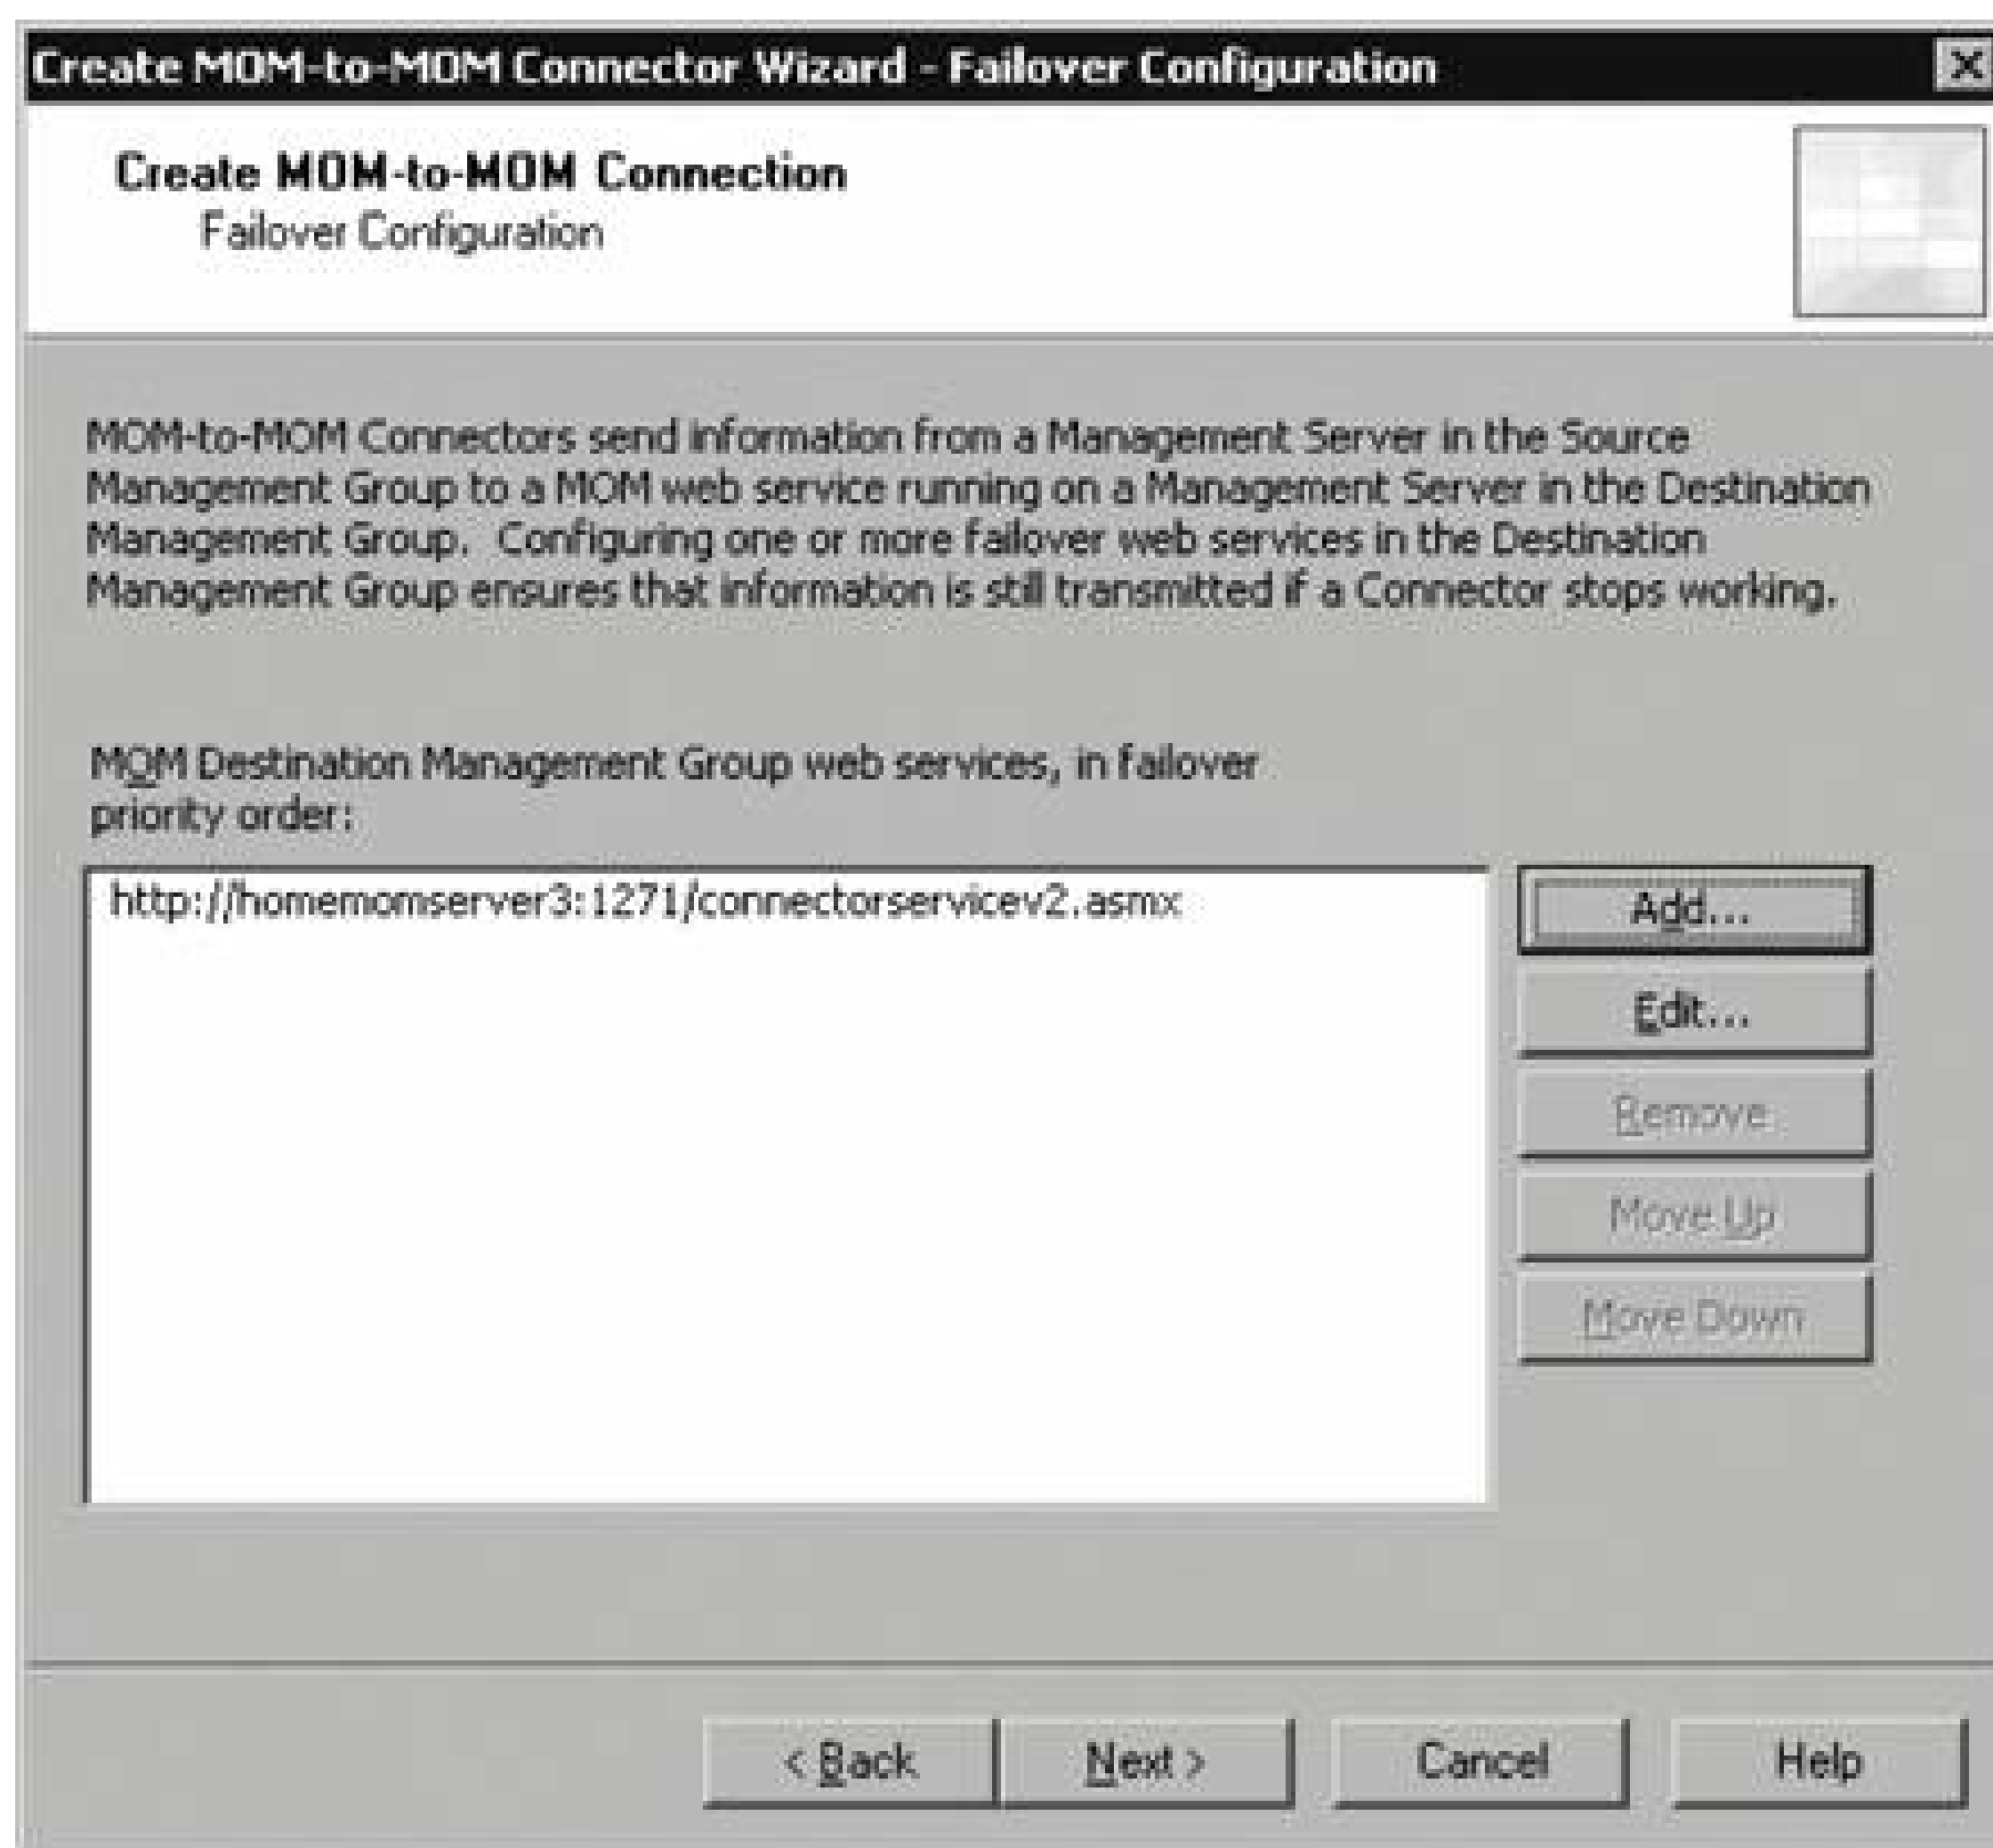
The MCF exposes MOM discovery data, which can be forwarded to the destination management group. This is very valuable, because otherwise only alerts and therefore only the Alert view would show data from the source management group. Since the discovery data is available, the State view, Computer Groups and Computers view, and the Diagram view all work just as they do in the respective source management groups. Because the computer grouping data is forwarded to the destination management group, you can execute destination Operator console tasks against servers that are managed by the source management group.

Now picture an architecture where there are several different source management groups synchronizing their data with a single destination management group. The destination management group provides a one-stop shop where you can view all of the alerts, the state of all your monitored machines and applications, and perform troubleshooting via tasks and resolve alerts.

Not all the data types in the source management group can be forwarded to the destination management group. For example, performance and events are not forwarded over the MMPC.

Click Next to open the Failover Configuration page, as shown in [Figure 9-11](#).

Figure 9-11. Configuring destination management group MCF web service failover

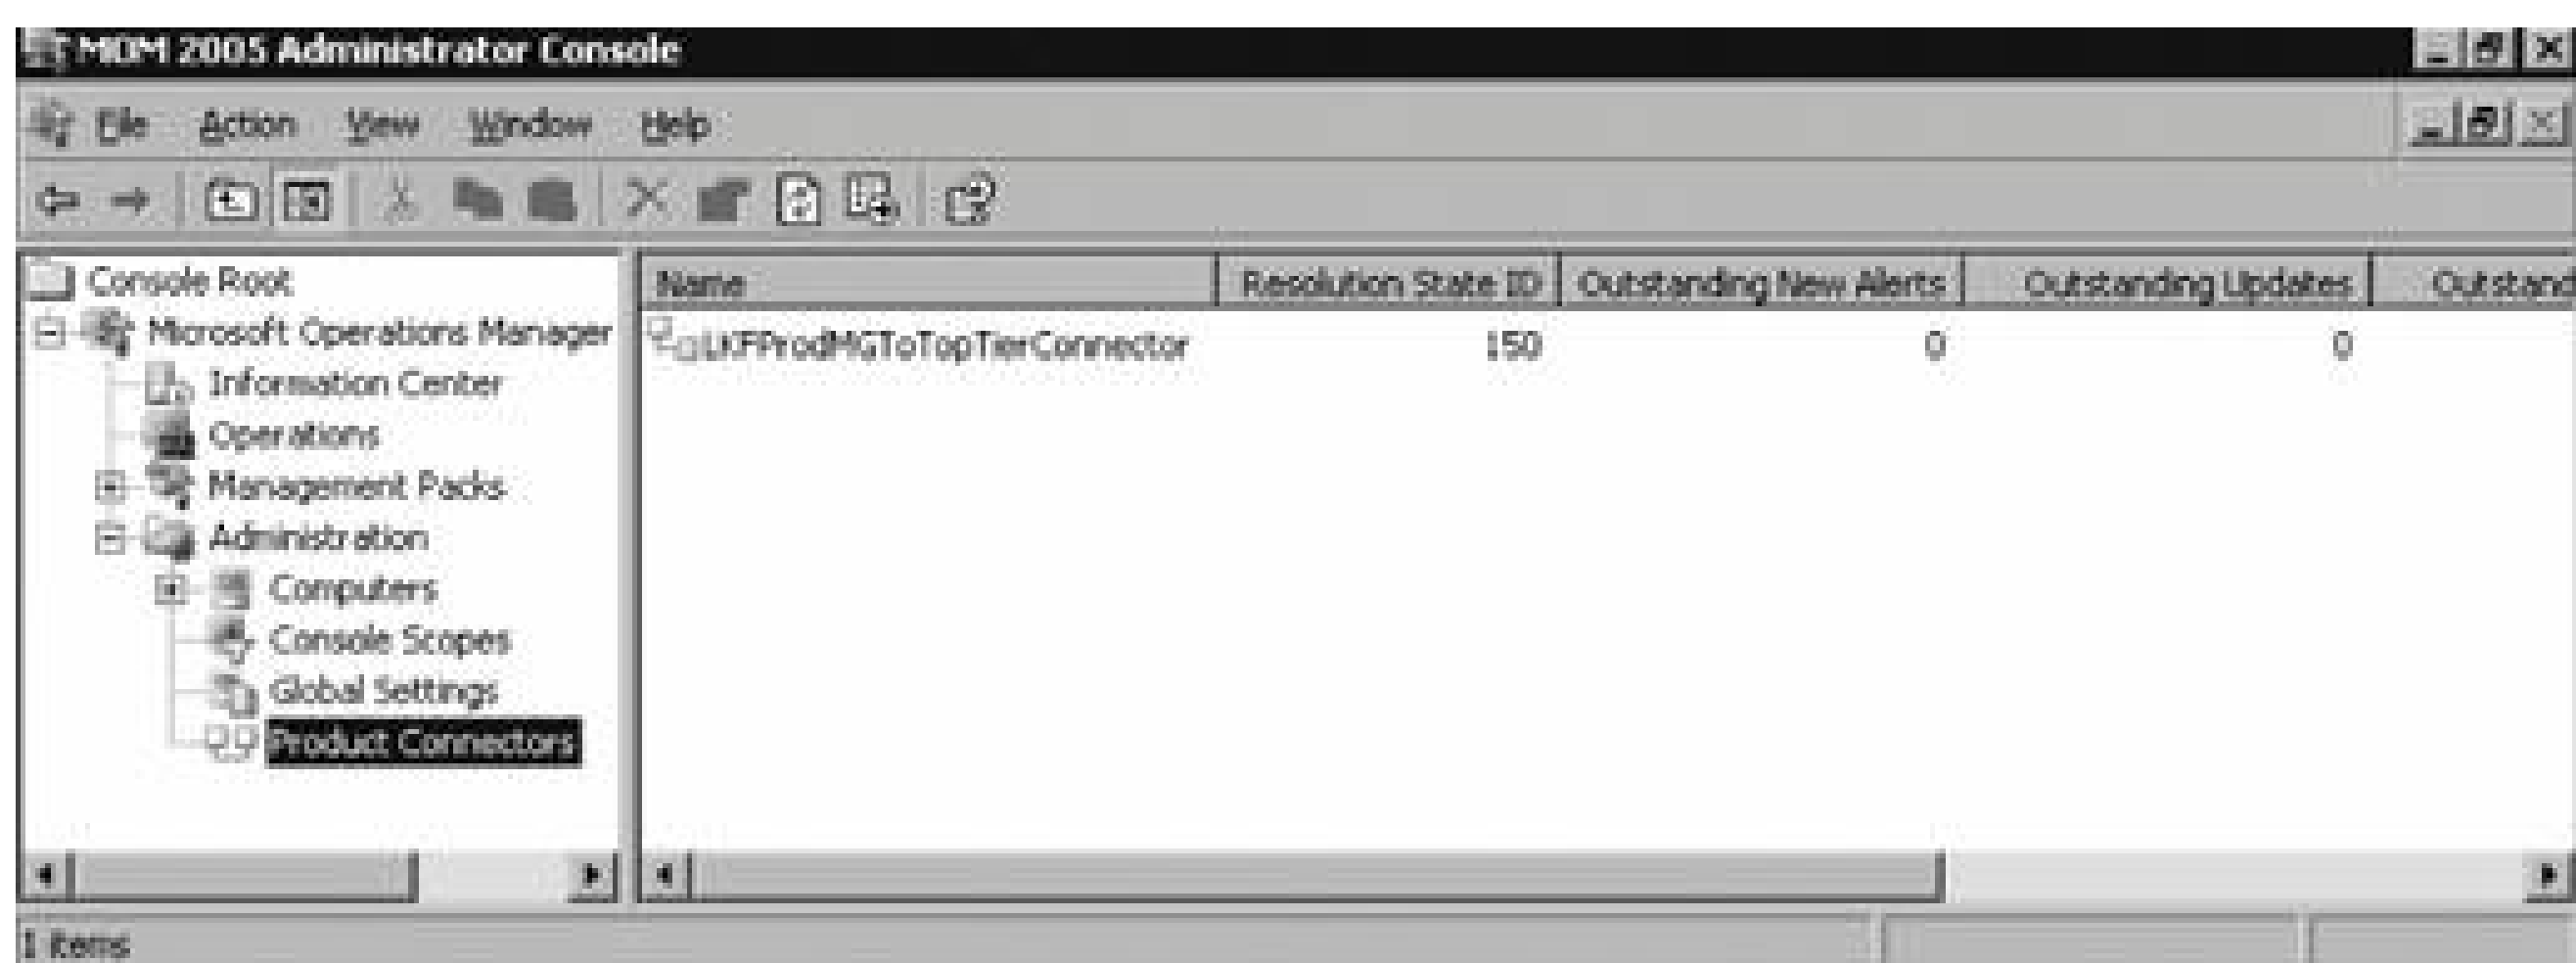


If there are multiple management servers in the destination management group, you should install the MCF on every one of them. Once the MCF is installed, each destination management group management server will have its own instance of the connector service web service. You can add the management servers here and place them in the order that you want failover to occur from the top down. When you create an MMPC on a source management server, you can set one of the destination management servers as the primary target for communication and the others in the destination management group in whatever order you choose. This is somewhat like the process for configuring which management servers are available to an agent for failover.

Click Next to bring up a summary/review page that shows all of the configuration choices that have been made, and then complete the installation.

In the Administrator console of the source management group, you should now see a functional MMPC with the name you specified (e.g., LKFProdMGToTopTierConnector) as shown in [Figure 9-12](#). The MMPC object is only editable from the Administrator console in the source management group; however, you should see it in the destination management groups Administrator console as well.

Figure 9-12. The MMPC as viewed in the source management group's Administrator console




In [Figure 9-12](#), the LKFPProdMGToTopTierConnector MMPC has a registered resolution state ID of 150. If any alerts are assigned that resolution state they are immediately sent to the MMPC for synchronization. All the parts are now in place and the next step is to configure the management groups for synchronization.

## 9.2.2. Configuring and Using an MMPC

For an alert to be successfully forwarded in a tiered configuration, the management packs must be synchronized between the source and destination management groups. This enables the destination management group to recognize a forwarded alert. The destination management group must contain a match to the rule that generated it from the source management group. The rules are matched based on their rule GUID, not on the management packs' version number. This means that the matching rule in the destination management group can be a different version as long as the GUID has not been changed. The easiest way to do this is to export all of the management packs from the source management group and then import them into the destination management group. One way to do this easily would be to modify the batch file that is used to synchronize and back up the management packs from the production to the preproduction environments. A sample batch file can be found in the "[Protecting Management Packs](#)" section in [Chapter 4](#). By adding a line that imports the management packs into the destination management group the end of the batch file, you can ensure that your source and destination management packs are always in sync. In this example line, the server *homemomserver3* is in the destination management group and the current version of the Microsoft Windows Base Operating System Management Pack is in the *mptransferfolder\currentmp* directory on *homesrv02*.

```
ManagementModuleUtil.exe -I homemomserver3 \\homesrv02\mptransferfolder\
currentmp\
MicrosoftWindowsBaseOperatingSystemMP%varDate%-%varTime%.akm -R
```





Admittedly, since only the rule GUIDs must match, synchronizing the management packs on a daily basis is overkill. However, having this process in place only requires a small amount of work and the benefits of automating this task far outweigh any concerns of unnecessary processing overhead.

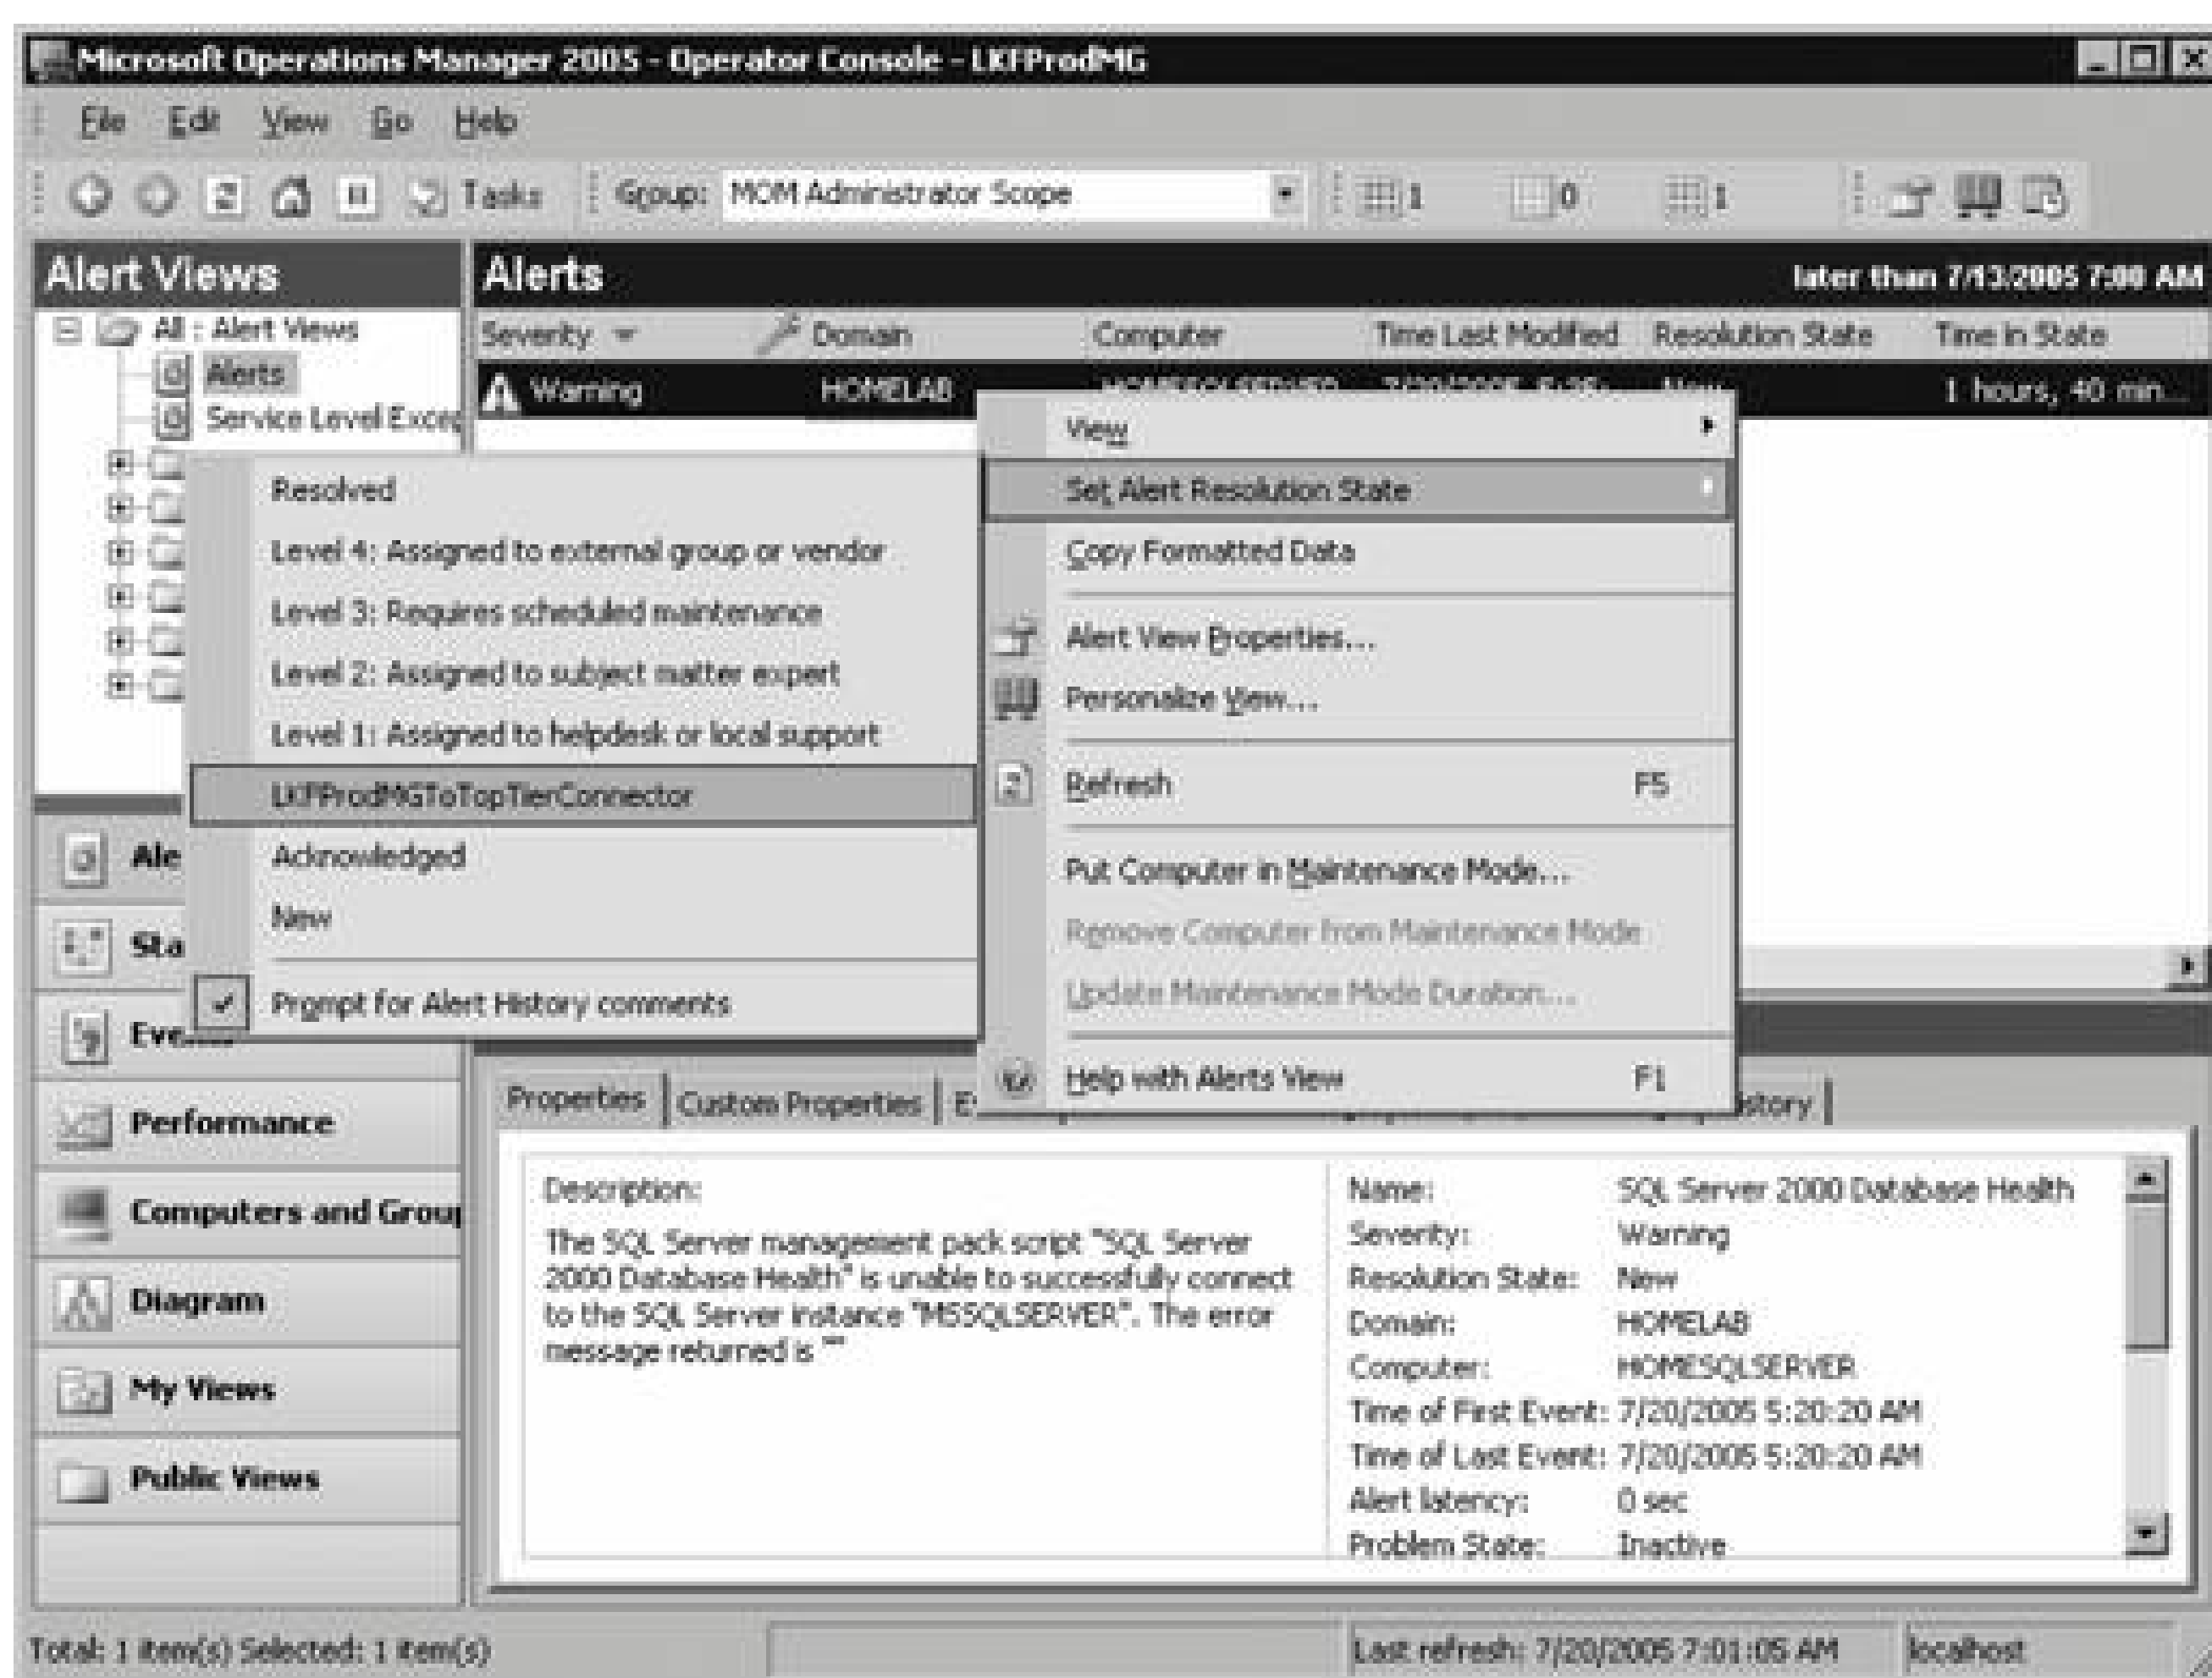
When the same management pack exists in more than one source management group, you only need to be concerned if the rule GUIDs differ between the management packs. It is fine to have a management pack Version 1.0 in one source group, Version 2.0 in another source, and Version 1.5 in the destination group.

For this example, all of the management packs have been exported from the LKFProdMG source management group and imported into the top-tier destination management group. Now alerts can be forwarded between the two groups.

One way to forward alerts is to manually set their resolution state to the name of the MMPC. The name of the MMPC will appear in the list of available resolution states, as shown in [Figure 9-13](#).

This step actually sets that alert resolution state to be 150, which is the resolution state value associated with the connector. The alert is then sent to the MMPC and synchronized to the destination management group. The alert appears in the destination group Operator console exactly as it does in the source Operator console. This alert resolution state only exists in the source management group; you cannot forward an alert from the destination back to the source. Since two-way synchronization is in place, any changes made to the alert in one console will also be made in the other. Therefore, resolving the alert in the destination Operator console resolves it in the source Operator console.

Figure 9-13. Manually setting the resolution state to the name of the MMPC



Manually forwarding alerts from source to destination is not a scalable solution. There is a preconfigured alert rule in the MOM management pack that addresses this issue. To enable automated forwarding of alerts, you must enable the alert forwarding rule and associate its rule group with the computer groups that then forward the alert to the destination group. In the Administrator console, navigate to Management Packs > Rule Groups node > Microsoft Operations Manager > Operations Manager 2005 > Connector Framework > "Mark Alerts for forwarding to MOM Master management group" rule group. In the alert rules, open the "Mark Alerts for forwarding to the MOM Master management group alert" rule. Select the "This rule is enabled" checkbox on the General tab. This alert rule fires for all alerts generated in the associated computer groups with a severity more than Error, as shown in [Figure 9-14](#).

This alert rule has a single response, which is to call a VBScript named "MOM Mark Alerts for forwarding to MOM Master management group." This is a very simple VBScript that changes the alert resolution state to 150:

```
' -----
' <company>Microsoft Corporation</company>
' <copyright>Copyright (c) Microsoft Corporation 2003</copyright>
' <name>MOM Mark alerts for forwarding to MOM Master management group</name>
' -----

Option Explicit

Sub Main( )
    Dim myAlert
```

```
'change resolution state
Set myAlert = ScriptContext.Alert
myAlert.ResolutionState = 150

End Sub
```

Figure 9-14. Criteria for the "Mark Alerts for forwarding to MOM Master management group" alert rule



By default, this rule group is not associated with any computer groups. In the example, the alerts with a severity greater than Error are forwarded for all computers that the source management group manages to the destination management group. This is done by associating the "Mark Alerts for forwarding to MOM Master management group" rule group to all the managed computers via computer groups. The computer groups chosen here are the Microsoft Operations Manager 2005 Agents and Microsoft Operations Manager 2005 Agentless computer groups. The combined membership of these two groups represents all of the managed computers that the source group is monitoring.

To create a rule group/computer group association, open the context menu of the rule group and select the "Associate with Computer Group" option as shown in [Figure 9-15](#).

This opens the rule group's Properties page with the Computer Groups tab on top. From here you can click Add and select the computer groups that you want to create the association with; see [Figure 9-16](#).



Figure 9-15. Creating a rule group/computer group association

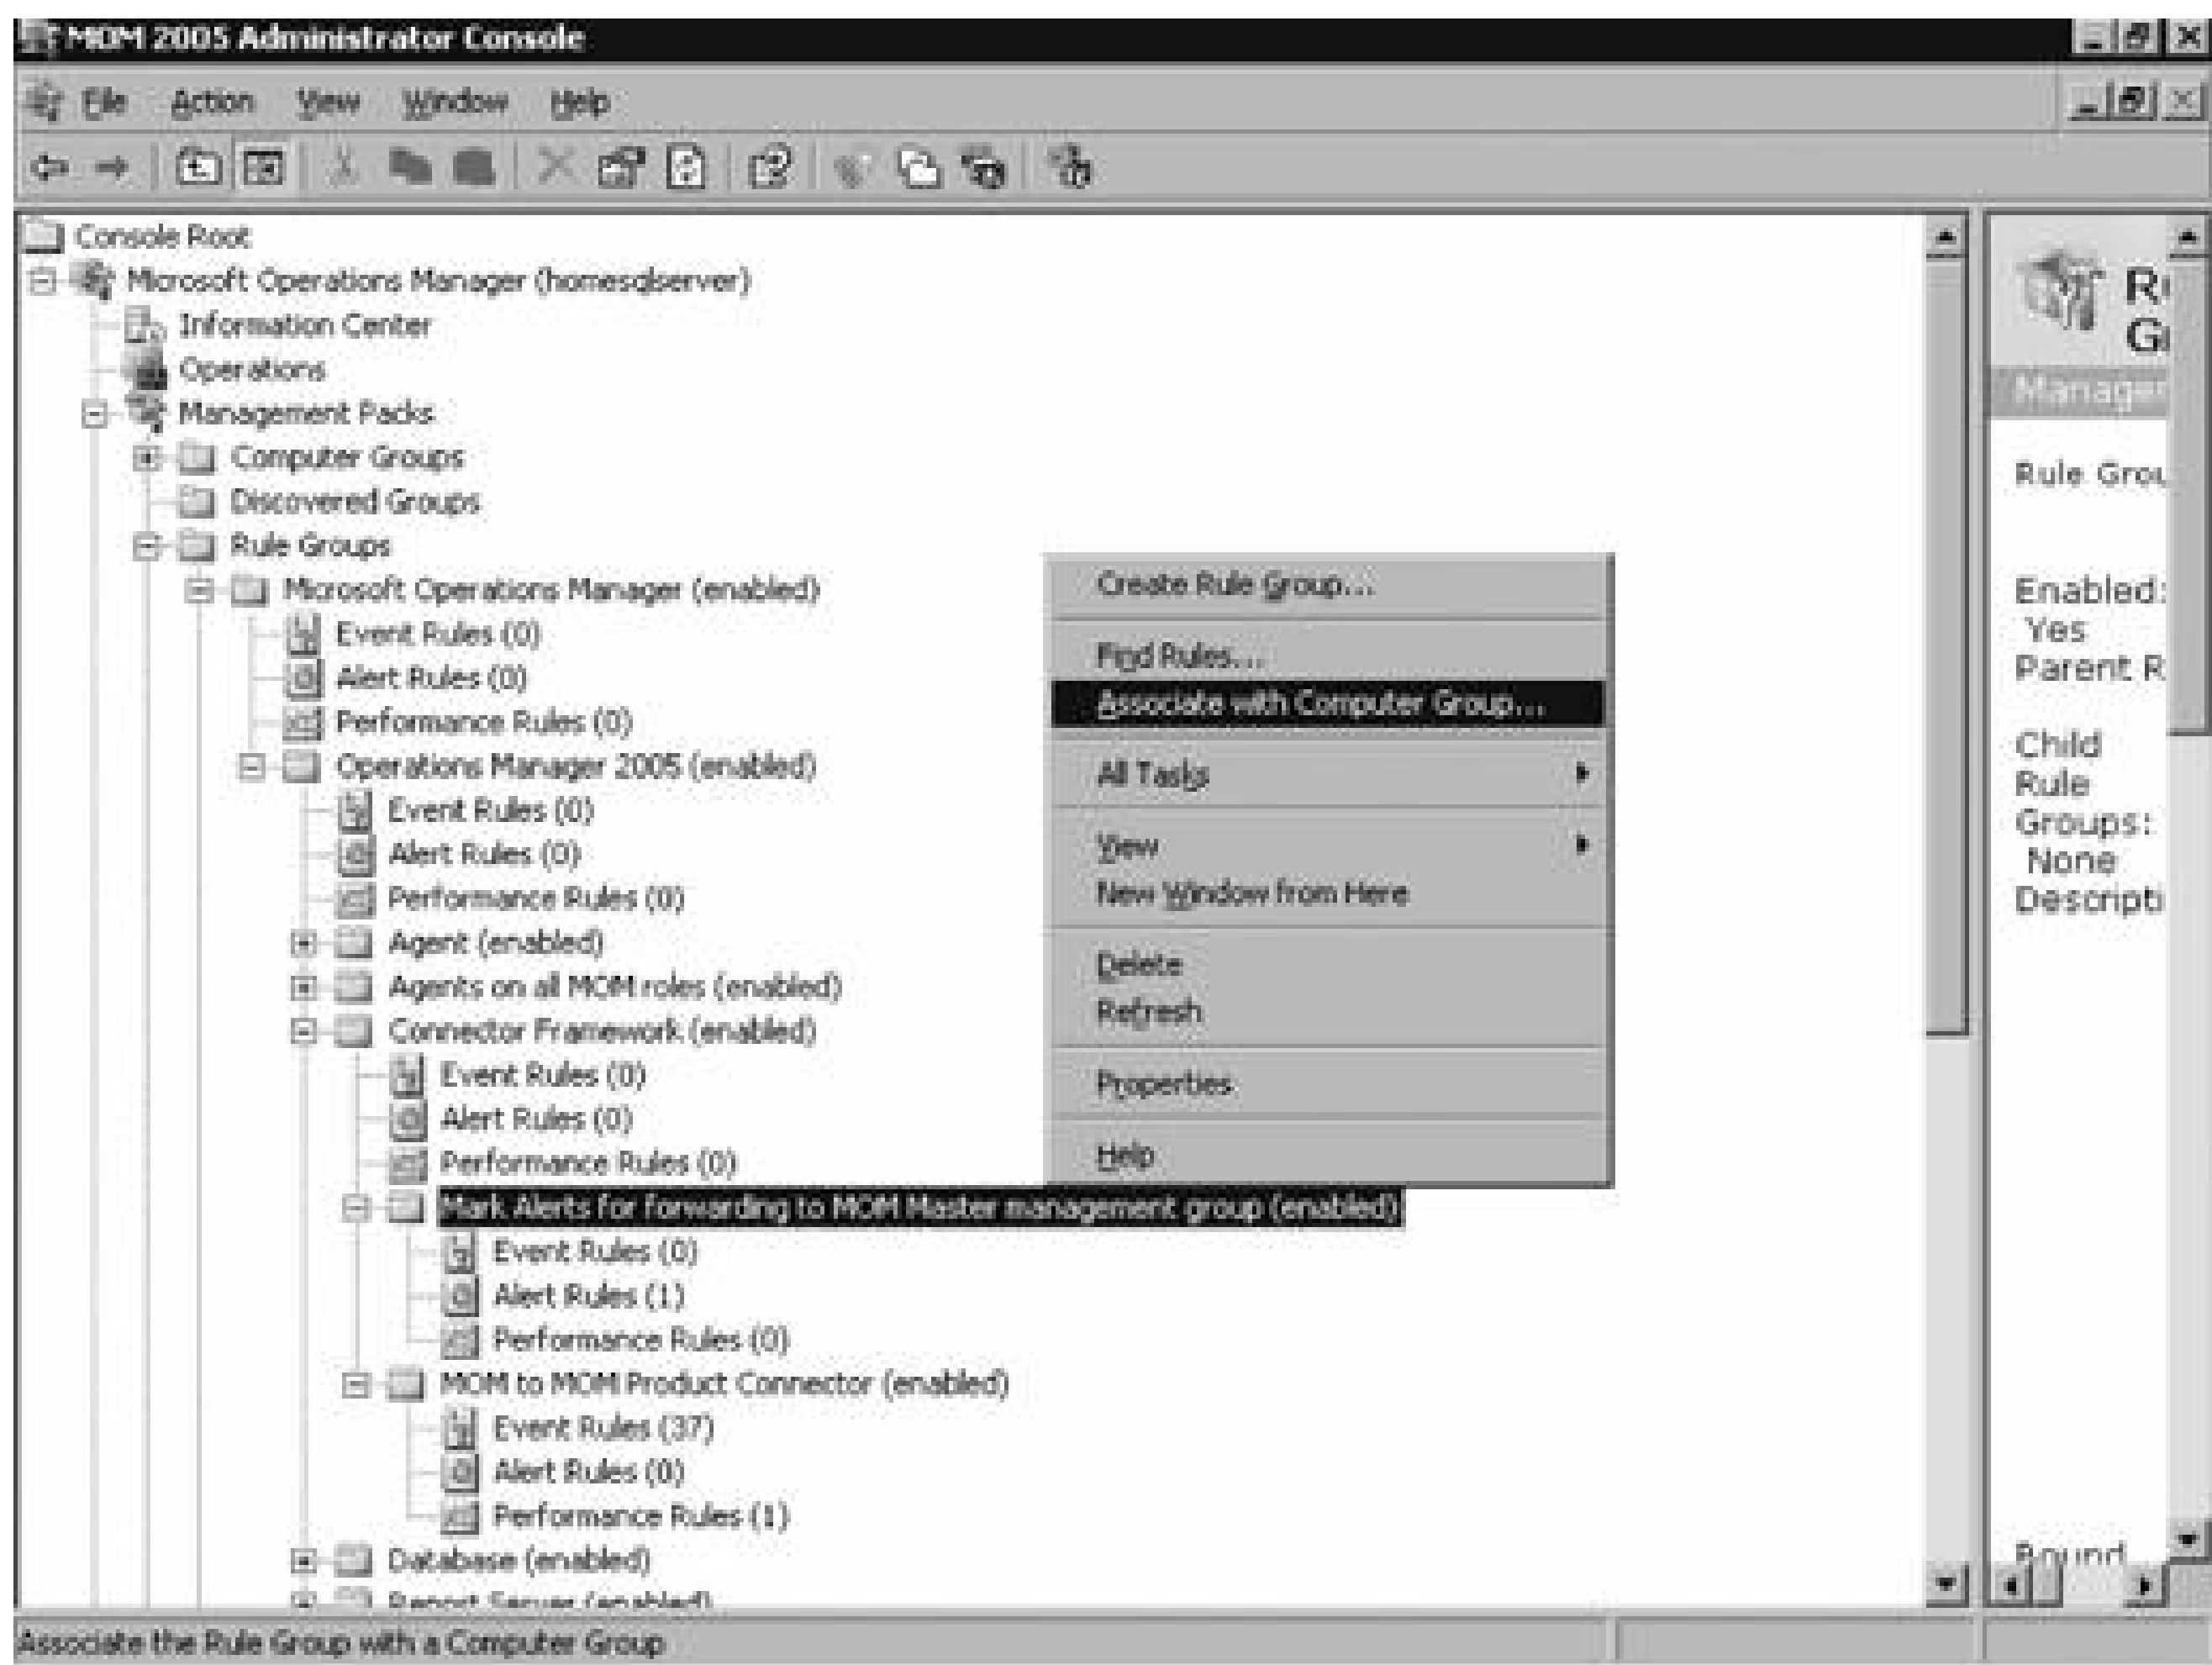


Figure 9-16. Selecting the computer groups to create the association with

Click OK, then click Apply to close the rule group properties. Don't forget to commit the configuration change so the modifications go into effect and are distributed by all of the agents.

The process an alert goes through, from the time it is generated in the source management group to its ultimate resolution, is recorded in the Alerts History tab. The text from the Alert History Tab of an Agent heartbeat failure alert is as follows. On the History tab, the events are inserted at the top rather than at the bottom so if you start reading at the top, you are seeing the most recent events first. For sake of clarity, that order has been revised here, presenting the oldest events first.

```
=====
7/20/2005 8:09:56 PM: NT AUTHORITY\NETWORK SERVICE
Alert is created in management group LKFProdMG.
=====
7/20/2005 8:09:57 PM: HOMELAB\dasaccount
HOMESQLSERVER:
[Forwarding] Attempting to forward new alert to LKFProdMGToTopTierConnector.
=====
7/20/2005 8:09:57 PM: HOMELAB\dasaccount
HOMESQLSERVER:
[Forwarding] Successfully forwarded alert change occurring at Jul 21 2005 1:09AM
(gmt) to LKFProdMGToTopTierConnector.
=====
7/20/2005 8:09:58 PM: HOMELAB\dasaccount
HOMESQLSERVER:
[Forwarding] Attempting to forward updated alert to LKFProdMGToTopTierConnector
=====
7/20/2005 8:09:58 PM: HOMELAB\dasaccount
HOMESQLSERVER:
[Forwarding] Successfully forwarded alert change occurring at Jul 21 2005 1:09AM
(gmt) to LKFProdMGToTopTierConnector.
=====
7/20/2005 8:12:21 PM: NT AUTHORITY\NETWORK SERVICE
Alert is created in management group LKFProdMG.
=====
7/20/2005 8:12:21 PM: HOMEMOMSERVER3\Administrator
HOMEMOMSERVER3: Changed 'Resolution State' from 'New' to 'Resolved'.
=====
```

Although this history does not capture every change made, it does track the important changes. In the first entry, at 8:09:56 p.m., the alert is created in the source management group. At 8:09:57, MOM attempts to forward the alert to the connector and at the next timestamp (also 8:09:57), the successful forward to the connector is noted. After the event is successfully forwarded, there is an attempt to forward an update to the alert to the connector (8:09:58) and its corresponding success at the same timestamp. Next, evidence of an update coming from the destination management group is logged at 8:12:21, and the final change of the resolution state from new to resolved at the same timestamp.

### 9.2.3. Monitoring the MCF and MMPC

MOM will monitor the MCF and MMPC for throughput (performance rules) and general errors (event rules). There is only one performance rule provided out-of-the- box and that samples the New Alert Forwarding Rate counter from the MOM-to-MOM Connector Service object. There is another counter, the Total New Alerts Forwarded, that you can use to create a performance rule to sample and use in a custom report.

Through the MOM Connector Framework performance object, you can monitor Alert Insert Rate and Alert Update Rate. These rules measure the rate, in alerts per second, that alerts are inserted into a management group's OnePoint database. The rules are used on the destination management group or any management group that receives operational data from an outside entity.

### 9.2.4. Reporting from Multiple Source Management Groups

Tiered configurations don't allow for consolidated reporting because none of the event or performance data is available in the destination management groups' OnePoint databases. You'd think that Microsoft would have built some sort of reporting data consolidation feature into MOM 2005, but they haven't. To configure reporting, the MOM 2005 Reporting server must still pull data from source layer management groups.

Recall from [Chapter 7](#) that when the MOM 2005 reporting solution is installed, a Windows scheduled task is created that runs a command line like this:

```
MOM.Datawarehousing.DTSPackageGenerator.exe /silent /srcserver
:homesqlserver /srcdb
:
OnePoint /dwserver
:HOMESQLSERVER /dwdb
:SystemCenterReporting /product:"Microsoft
Operations Manager"
```

To consolidate operational data from multiple source management groups, create additional scheduled tasks that run the `MOM.Datawarehousing.DTSPackageGenerator.exe`. Do this by creating a copy of the original task and editing the command line so that the desired *srcserver* (source server to connect to), *srcdb* (source database on the source server to pull data from), *dwserver* (the MOM 2005 Reporting server), and *dwdb* (the database into which data is inserted) flags to it. Change the time that the task will be run, because you don't want more than one DTS package running concurrently.

Combining reporting data from multiple source groups into one reporting server is not officially supported by Microsoft.

### 9.2.5. Connecting MOM to Other Management Frameworks



Connecting to other management frameworks can be a huge topic. The MCF API was created so that connectors could be written that would allow synchronization between MOM and any other operations management platforms. Creating a connector to a third-party product requires a detailed knowledge of the APIs available for that product. Instruction on how to write product connectors that sit between MOM and other operations management products is beyond the scope of this book. There are other alternatives though.

Microsoft has written connectors that provide two-way synchronization between MOM and HP OpenView, HP Network Node Manager, and Tivoli TEC. These connectors are written entirely on the MCF and were actually released for use with the previous version of MOM (MOM 2000, SP1). These connectors can be downloaded from the Microsoft web site.

There is also an active development community that has written a number of MCF- based connectors. Here is a partial listing of products that you can connect MOM to using these connectors:

- Aprisma SPECTRUM
- BMC Impact
- BMC Patrol
- CA Solve
- CA Unicenter
- HP OpenView
- MetiLinx
- Micromuse Netcool
- NetIQ AppManager
- Quest InTrust Connector for Windows
- Remedy ARS
- Tivoli TEC

For a more complete and current listing, see the MOM Management Pack and Product Connector Catalog at <http://www.microsoft.com/management/mma/catalog.aspx>.

## 9.3. Summary

A MOM architecture made up of a single or multiple management groups that don't exchange operations data is a great starting place for your monitoring and alerting solution. However, there are many situations that require the deployment of multiple management groups that exchange operations management data. You may need partitioning of management duties for scalability reasons, administrative needs, or because of security or configuration requirements. Management groups can be arranged in a tiered configuration where a source management group performs the bulk of the monitoring and alerting tasks and then forwards that data to a top-tier destination management group for centralized consumption and resolution. The communication between the management groups can be one way (source to destination) or two way (source to destination and back again).

The MCF is one of three MOM APIs included in the MOM SDK. It is the preferred API for exposing operational and discovery data and for building connectors that communicate data from a source management group to a destination management group or to another operations management product. MOM ships with a built-in connector called MMPC that an administrator with no programming skills can configure to forward alert and computer discovery data between management groups.

In addition to the MMPC, Microsoft has created connectors that integrate MOM data into HP OpenView, Tivoli TEC, and HP Network Node Manager. There are many ISV-built connectors available for purchase.

The next chapter looks at extending MOM's management capabilities beyond Microsoft platforms to SNMP-enabled devices and those that make use of syslog. This is done using MOM's out-of-the-box abilities and the Windows OS.

# Chapter 10. Extending Monitoring

You can purchase a third-party solution to integrate operational data from other platforms (e.g., Unix) or from network devices (e.g., routers and switches) into MOM. Some of the companies that build MOM connectors also create solutions for monitoring non-Microsoft devices. For example, eXc (<http://www.excsoftware.com>) sells both types of solutions, as does Vintela(now owned by Quest Software, <http://www.vintela.com> or <http://www.quest.com>), Skywire (<http://www.skywiresoftware.com>), and Jalasoft (<http://www.jalasoft.com>).

Third-party solutions provide the richest possible integration experience, but are usually expensive. If you don't have the option to buy a pre-made solution or to have a solution custom developed, and you only need basic integration between MOM and Unix or MOM and other networked devices, then you can use some native MOM 2005 capabilities and the Windows OS. Unix-based systems track operations data in *syslog files*, which a MOM management server can receive and parse, generating alerts based on the content of those files. Network devices and computers can make use of the Simple Network Management Protocol (SNMP), which is both a communication protocol and a data storage format to track configuration information and send traps (which are kind of like an alert without any diagnostic, historic, or resolution data) to central monitoring consoles.

This chapter teaches you how to configure MOM and the Windows OS to make use of SNMP and syslogs so that you can integrate other platforms into your MOM monitoring solution.



## 10.1. Understanding SNMP

SNMP is part of the TCP/IP suite of protocols and is used for communicating monitoring data (called an SNMP trap) from SNMP agents to an SNMP console over port 162. It is also used to gather configuration information from a device and to write configuration data to a device over port 161. When you discuss monitoring using SNMP, you're usually referring to network devices (switches and routers). Most operations management solutions are capable of receiving and sending SNMP traps to be backward compatible. MOM is no exception and uses the Windows SNMP Windows Management Instrumentation (WMI) provider.

An SNMP trap is similar to an alert it is triggered by a predefined event, such as a reboot, on the SNMP-monitored device. It contains information about the event and is sent from the SNMP agent to a central console. Unlike a MOM alert, the only thing you can do with an SNMP trap is to acknowledge it in the SNMP console. You could keep it for historical purposes, but the only thing a trap really gives you is a message from the managed device saying "this event happened at this time."

SNMP data is arranged in a hierarchy, much like the DNS hierarchy. At the top level of the hierarchy, public identifiers are defined by Internet authorities. Where DNS uses domain names like .com, .gov, or .org at the top and allows registration of sub-domain names, the SNMP namespace uses a dotted decimal notation to assign numbers that map to Internet entities and sub-entities. A complete SNMP data identification string looks very much like an IP address, except much longer. In DNS, the complete path to an object in the public DNS namespace, such as [homemomserver.homelab.lab.com](http://homemomserver.homelab.lab.com), is the FQDN. In SNMP, every attribute of a device or an event can be described in the dotted decimal notation and the whole string is called an *object identifier* (OID). For example, this OID string is for a successful network logon to a Windows server:

```
.1.3.6.1.4.1.311.1.13.1.9999.1.0
```

In this OID string, the numbers map to these fields respectively:

```
.iso.org.dod.internet.private.enterprises.microsoft.software.13.1.9999.1.0
```

The complete mapping of fields to actual values for a device or application is done in a management information block (MIB) file for any SNMP device or application. SNMP management applications (consoles) need the mappings in MIB files to decode the OIDs in SNMP traps and to read and write information to an SNMP device. The management applications compile raw MIB files, which are just text files of a specific format, into a format that is used by the management application. For example for MOM to catch SNMP traps from a Cisco router you would need to get the MIB file for that device from Cisco and compile it into the Windows WMI namespace. The traps could be translated from OID format into something that is readable by the Windows OS, MOM, and humans. SNMP traps come in

three versions: v1, v2, and v3. The SNMP-monitored device and the SNMP console must speak the same version of SNMP to communicate. The versions are differentiated by increasing functionality starting from v1. In the context of the Windows OS and MOM, you will only be working with v1 and v2 traps.

MOM can also generate SNMP traps as a response to an alert. The MOM-generated SNMP traps can be sent to another application that speaks SNMP. This is another way that MOM alerts can be integrated into other operations management systems. The MOM MIB file is called *MicrosoftOperationsManager.mib*, and it is in the MOM 2005 SDK.

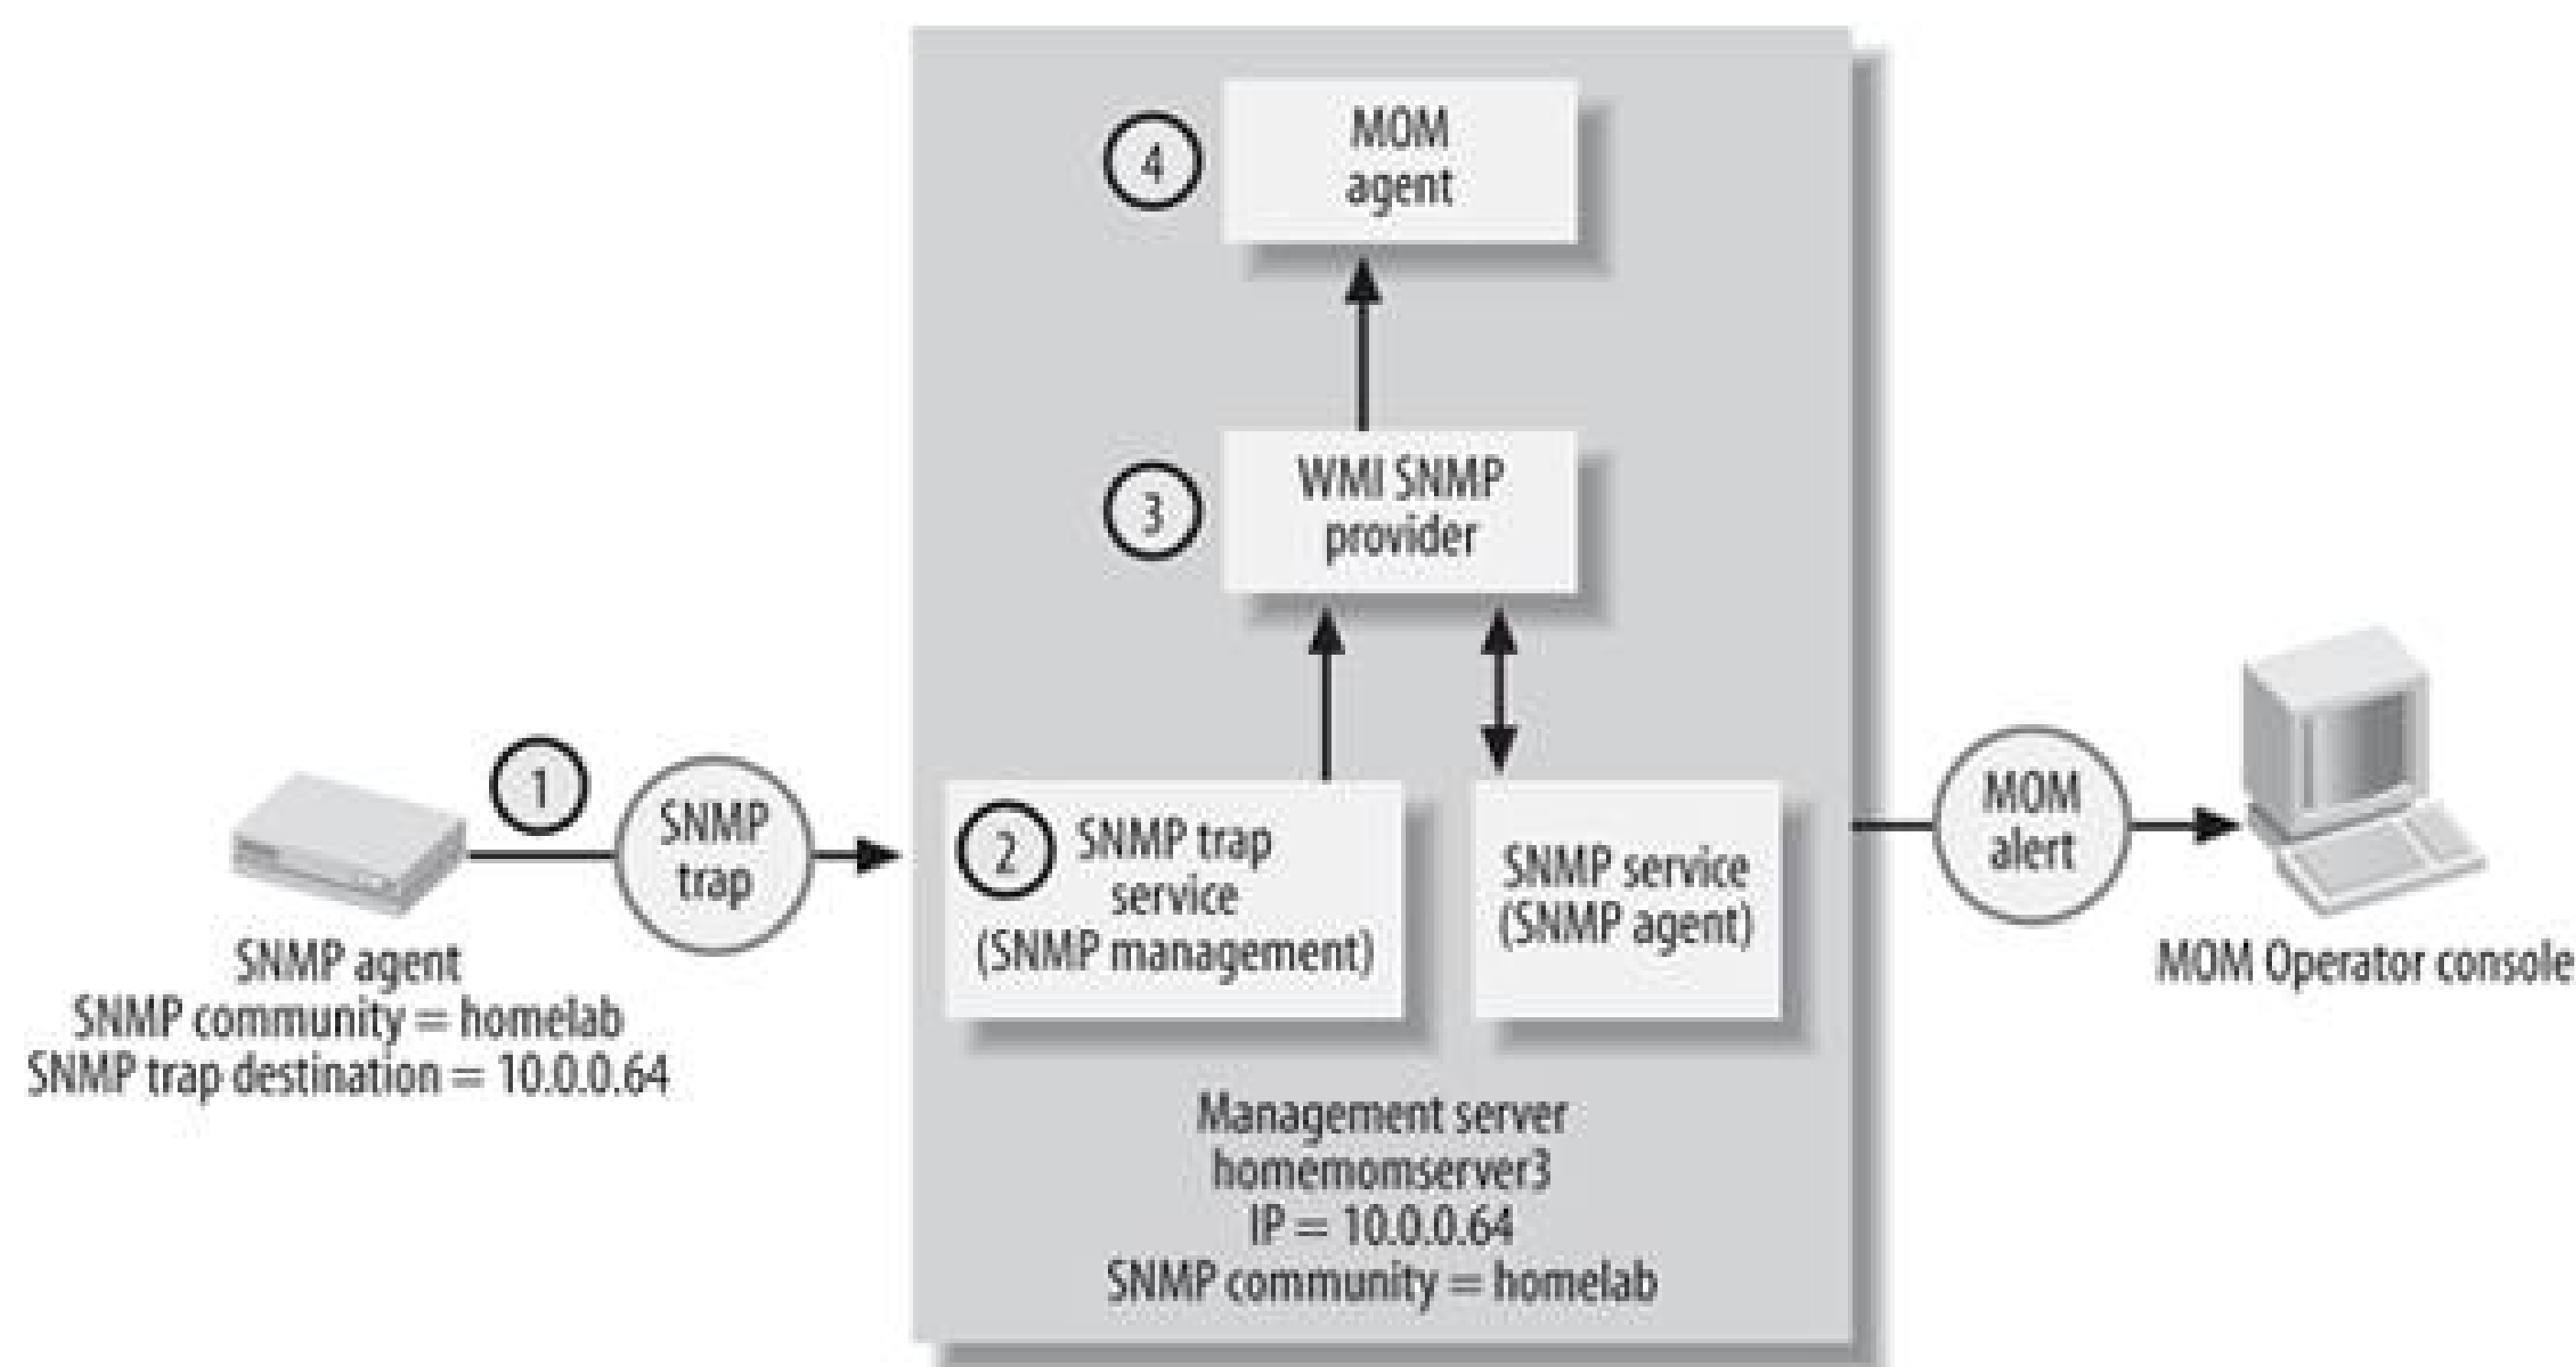
If you read the *MicrosoftOperationsManager.mib* file, you will notice that it is called the *Mission\_Critical\_MIB* file. Before Microsoft acquired the product that eventually became MOM, it was owned by a company called Mission Critical Software that created the initial MIB. Mission Critical was issued a namespace in the OID hierarchy and the MIB was created with that information embedded. This fact is irrelevant to the product's functions, but it is a little bit of interesting history and it answers what can be a puzzling question.



## 10.2. Windows and MOM Implementation of SNMP

By itself, MOM cannot receive or send SNMP. In fact, it can't speak SNMP at all; it relies on the SNMP Windows OS, the SNMP service, and SNMP Trap service for translation and sending or receiving SNMI data. The SNMP service communicates to MOM through WMI, which MOM monitors via the WMI SNMI provider. [Figure 10-1](#) shows the flow and conversion of an SNMP trap from an SNMP device to a management server.

Figure 10-1. SNMP trap and data flow



In this flow, a preconfigured event occurs and the SNMP agent on the SNMP-enabled device generates a trap (point 1 in [Figure 10-1](#)). In this example, the device is a member of the SNMP community named Homelab. SNMP communities are insecure boundaries that are created for administrative purposes. The device is configured to send its traps to 10.0.0.64. SNMP trap communication uses UDP port 162, and all other SNMP management communication occurs over UDP port 161. Therefore, SNMP information is not secure by default and delivery is not guaranteed.

The trap is received (point 2 in [Figure 10-1](#)) by the management server, which is running the SNMP Trap service and the SNMP service. The SNMP Trap service allows the management server to receive and interpret SNMP traps. When the SNMP Windows component is installed on a server, the SNMP Trap service and SNMP service are both installed. The SNMP service is essentially an SNMP agent. It is used when the management server needs to generate and send a trap of its own (as in the SNMP alert response option). The SNMP Trap service uses whatever MIB the sending device formatted the



trap in and sends the trap data into the WMI namespace via the WMI SNMP provider (point 3 in [Figure 10-1](#)).

Once the trap data is in the WMI namespace, it is in WMI format and is fully accessible to a MOM agent via WMI (point 4 in [Figure 10-1](#)). The management server agent applies a rule that uses the WMI Extended SNMP Trap Catcher provider. The agent can match the incoming data to the appropriate rule criteria and then generate an alert.

MOM alerts that are created from SNMP traps are not as rich as MOM alerts that are created natively from Windows, but there is a lot of useful information in them.

When generating SNMP traps from MOM to be sent to another SNMP management tool, the flow is almost the exact opposite of the one in [Figure 10-1](#):

1. A MOM rule is configured with a response to send an SNMP trap that uses the content of the alert as its source.
2. MOM calls the SNMP service (through WMI) and passes the data to it.
3. The SNMP service takes the alert data and translates it into SNMP format, using the *MicrosoftOperationsManager.mib* as a template.
4. The SNMP service then sends the trap to whatever community and destination IP address have been configured.

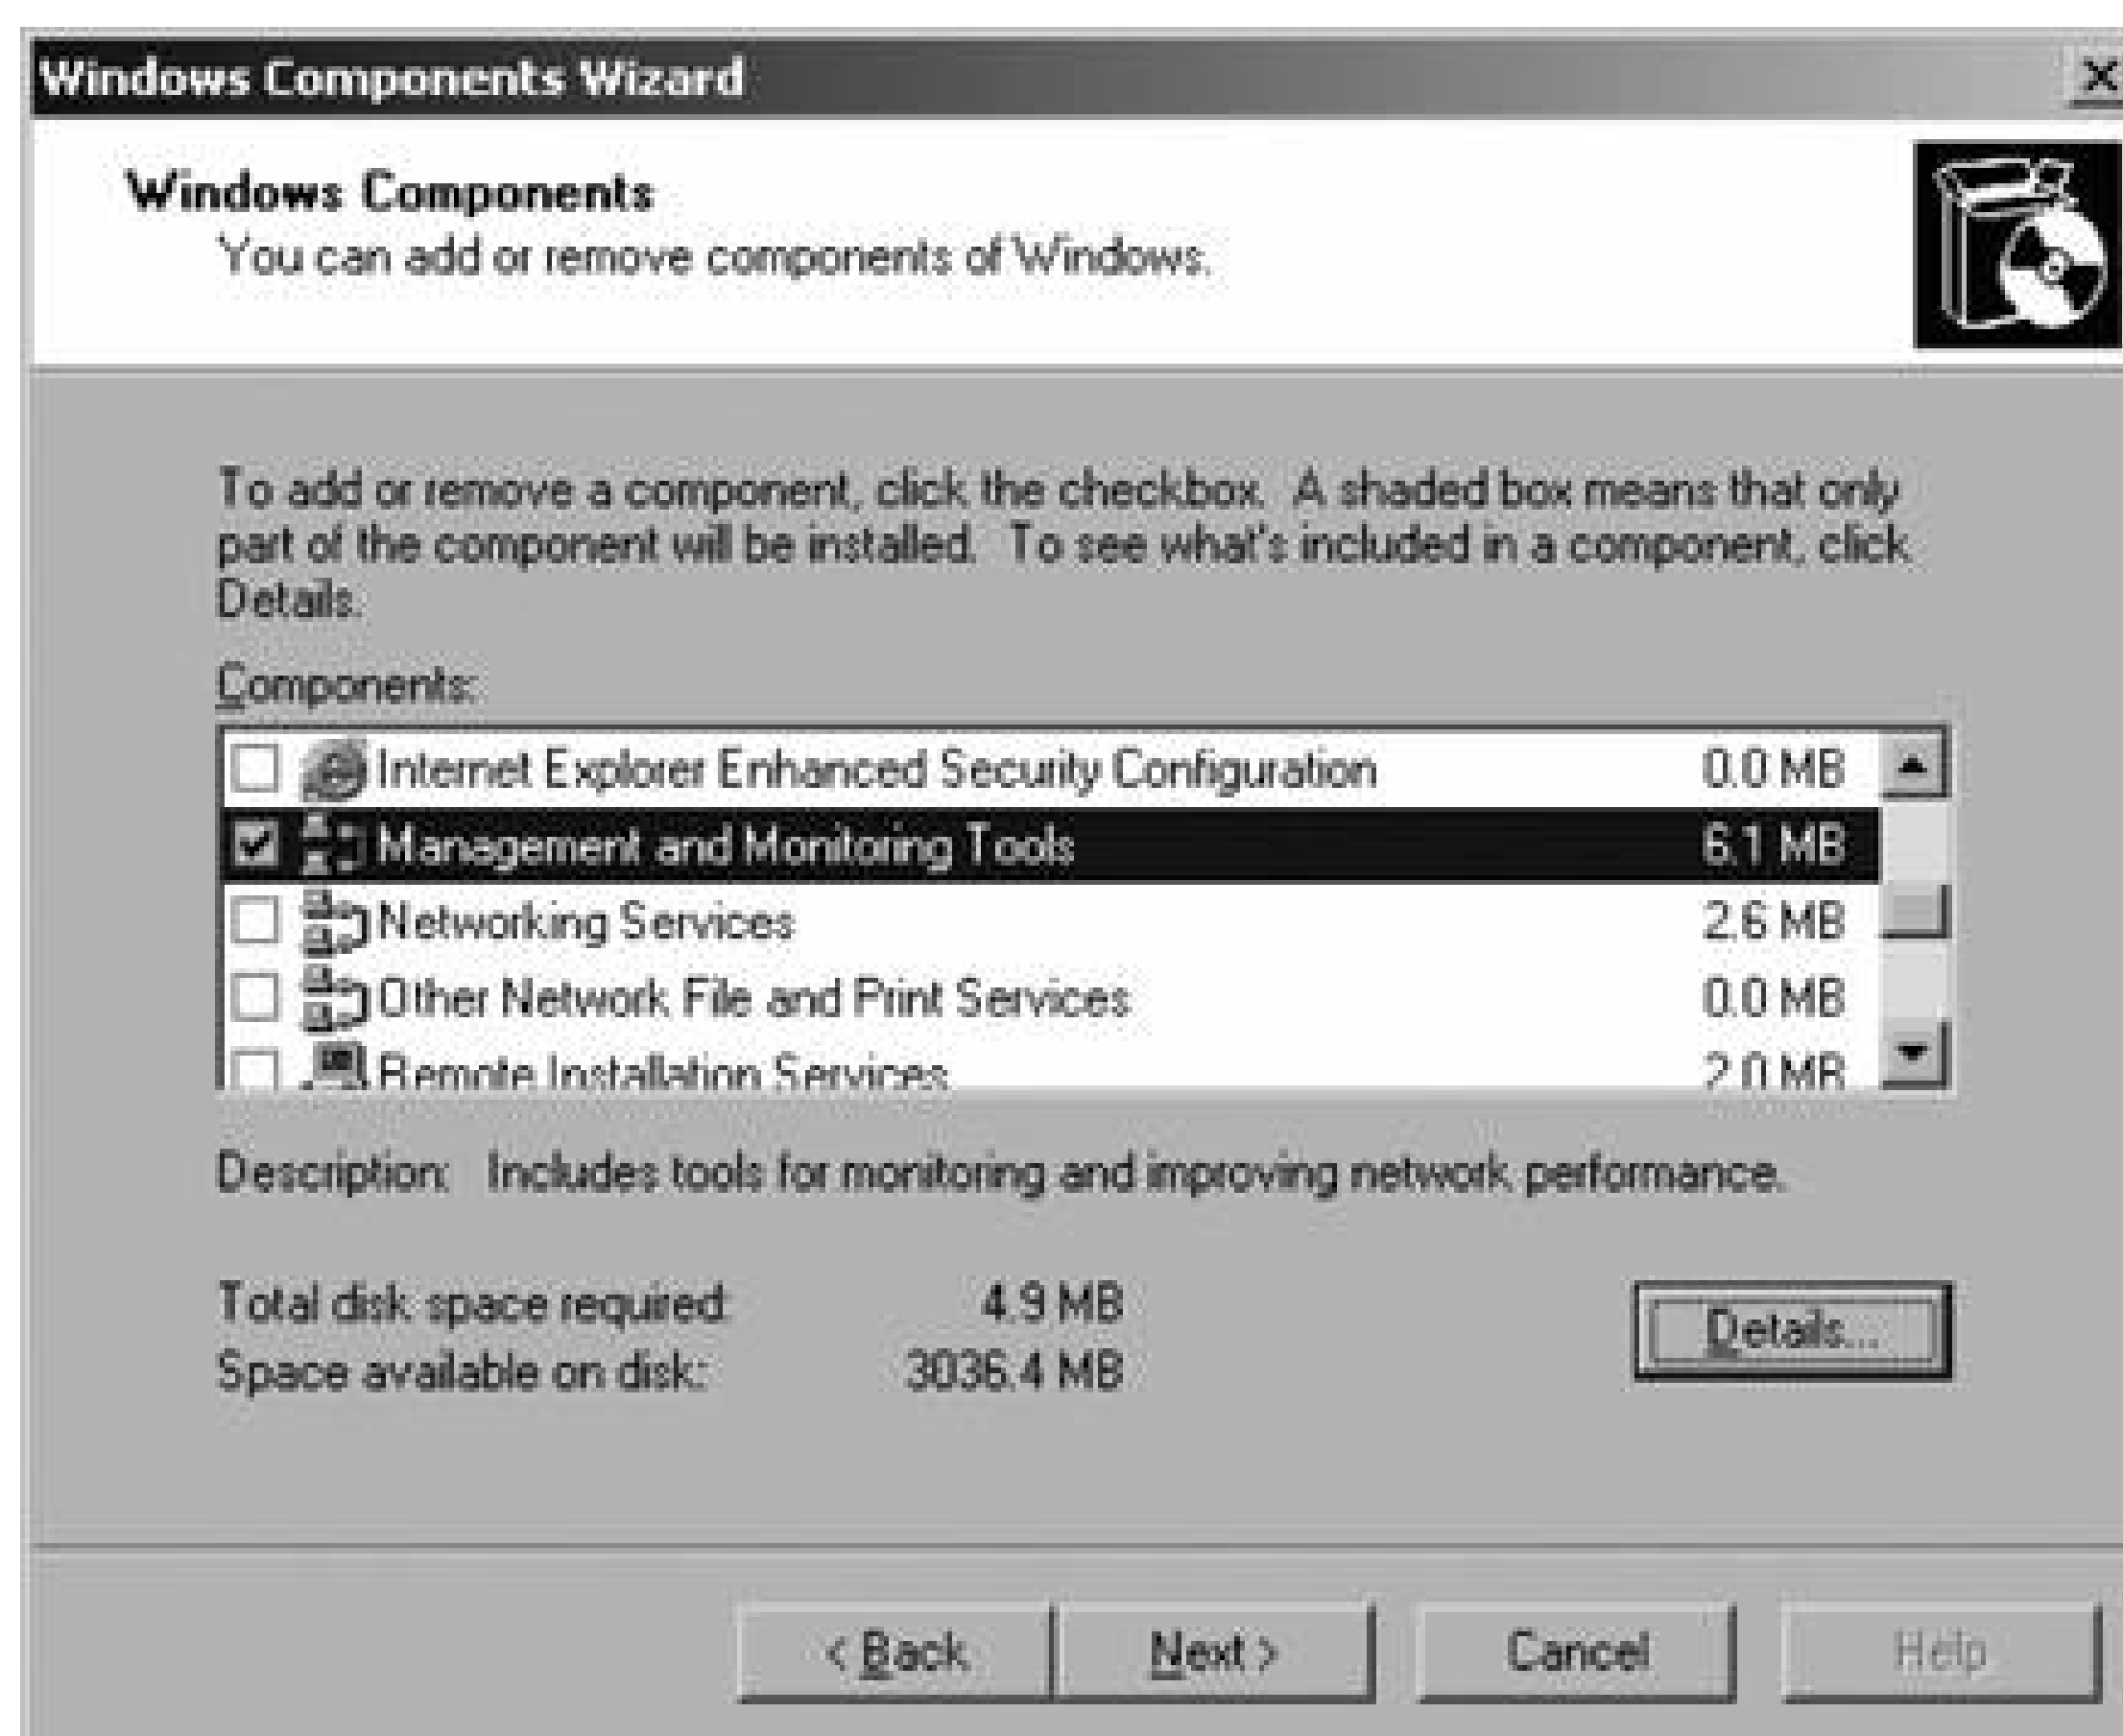
## 10.3. Configuring Windows and MOM for SNMP

To get SNMP data into and out of MOM, you have to touch the OS, MOM, and the SNMP devices that you want to receive traps from. If you are going to translate alerts into traps, then you need to configure the SNMP service on the management server to send traps to your other SNMP management system.

SNMP management systems and agents are related to each other by an SNMP *community name*, which is a simple text string and can be anything you want. Be careful though, because text strings are case-sensitive. If you have worked with SNMP, then you will have heard of the private and public SNMP communities. Many SNMP products come preconfigured to work with these two community names. In the example for this chapter, the community name Homelab is used. All SNMP agents need to be told which community they are members of and SNMP management consoles will only manage SNMP agents that are in their community (although they can receive traps from any SNMP agent if so configured). So, your first two steps in configuring SNMP are to pick a community name and a computer to which all traps will be sent. [Figure 10-1](#) shows the top-tier management server, *homemomserver3*, at 10.0.0.64. With those two decisions out of the way, most of the work in preparing for SNMP is in the OS.

The first step is to install the SNMP service and the WMI SNMP provider's Windows OS components. To do this, open the Add or Remove Programs utility and select the Add/Remove Windows Components. This opens the Windows Components page (see [Figure 10-2](#)).

Figure 10-2. The top-level Windows Components page



Here you select the Management and Monitoring Tools option and click the Details button. This brings up all the tools that are under the Management and Monitoring Tools heading (see [Figure 10-3](#)).

Here you select the Simple Network Management Protocol and the WMI SNMP Provider. Click OK to close the Details dialog, and click Next on the Components page to finish the component install wizard. If the install directory (i386) is not available to the local machine via network share or on the hard drive, you will be prompted to supply the OS source media.

Now you need to configure the SNMP service and the SNMP Trap service. On the management server, open the Computer Management tool → Services and Applications → Services container. Then open the SNMP Service Properties and select the Agent tab ([Figure 10-4](#)).

Figure 10-3. Select the SNMP components



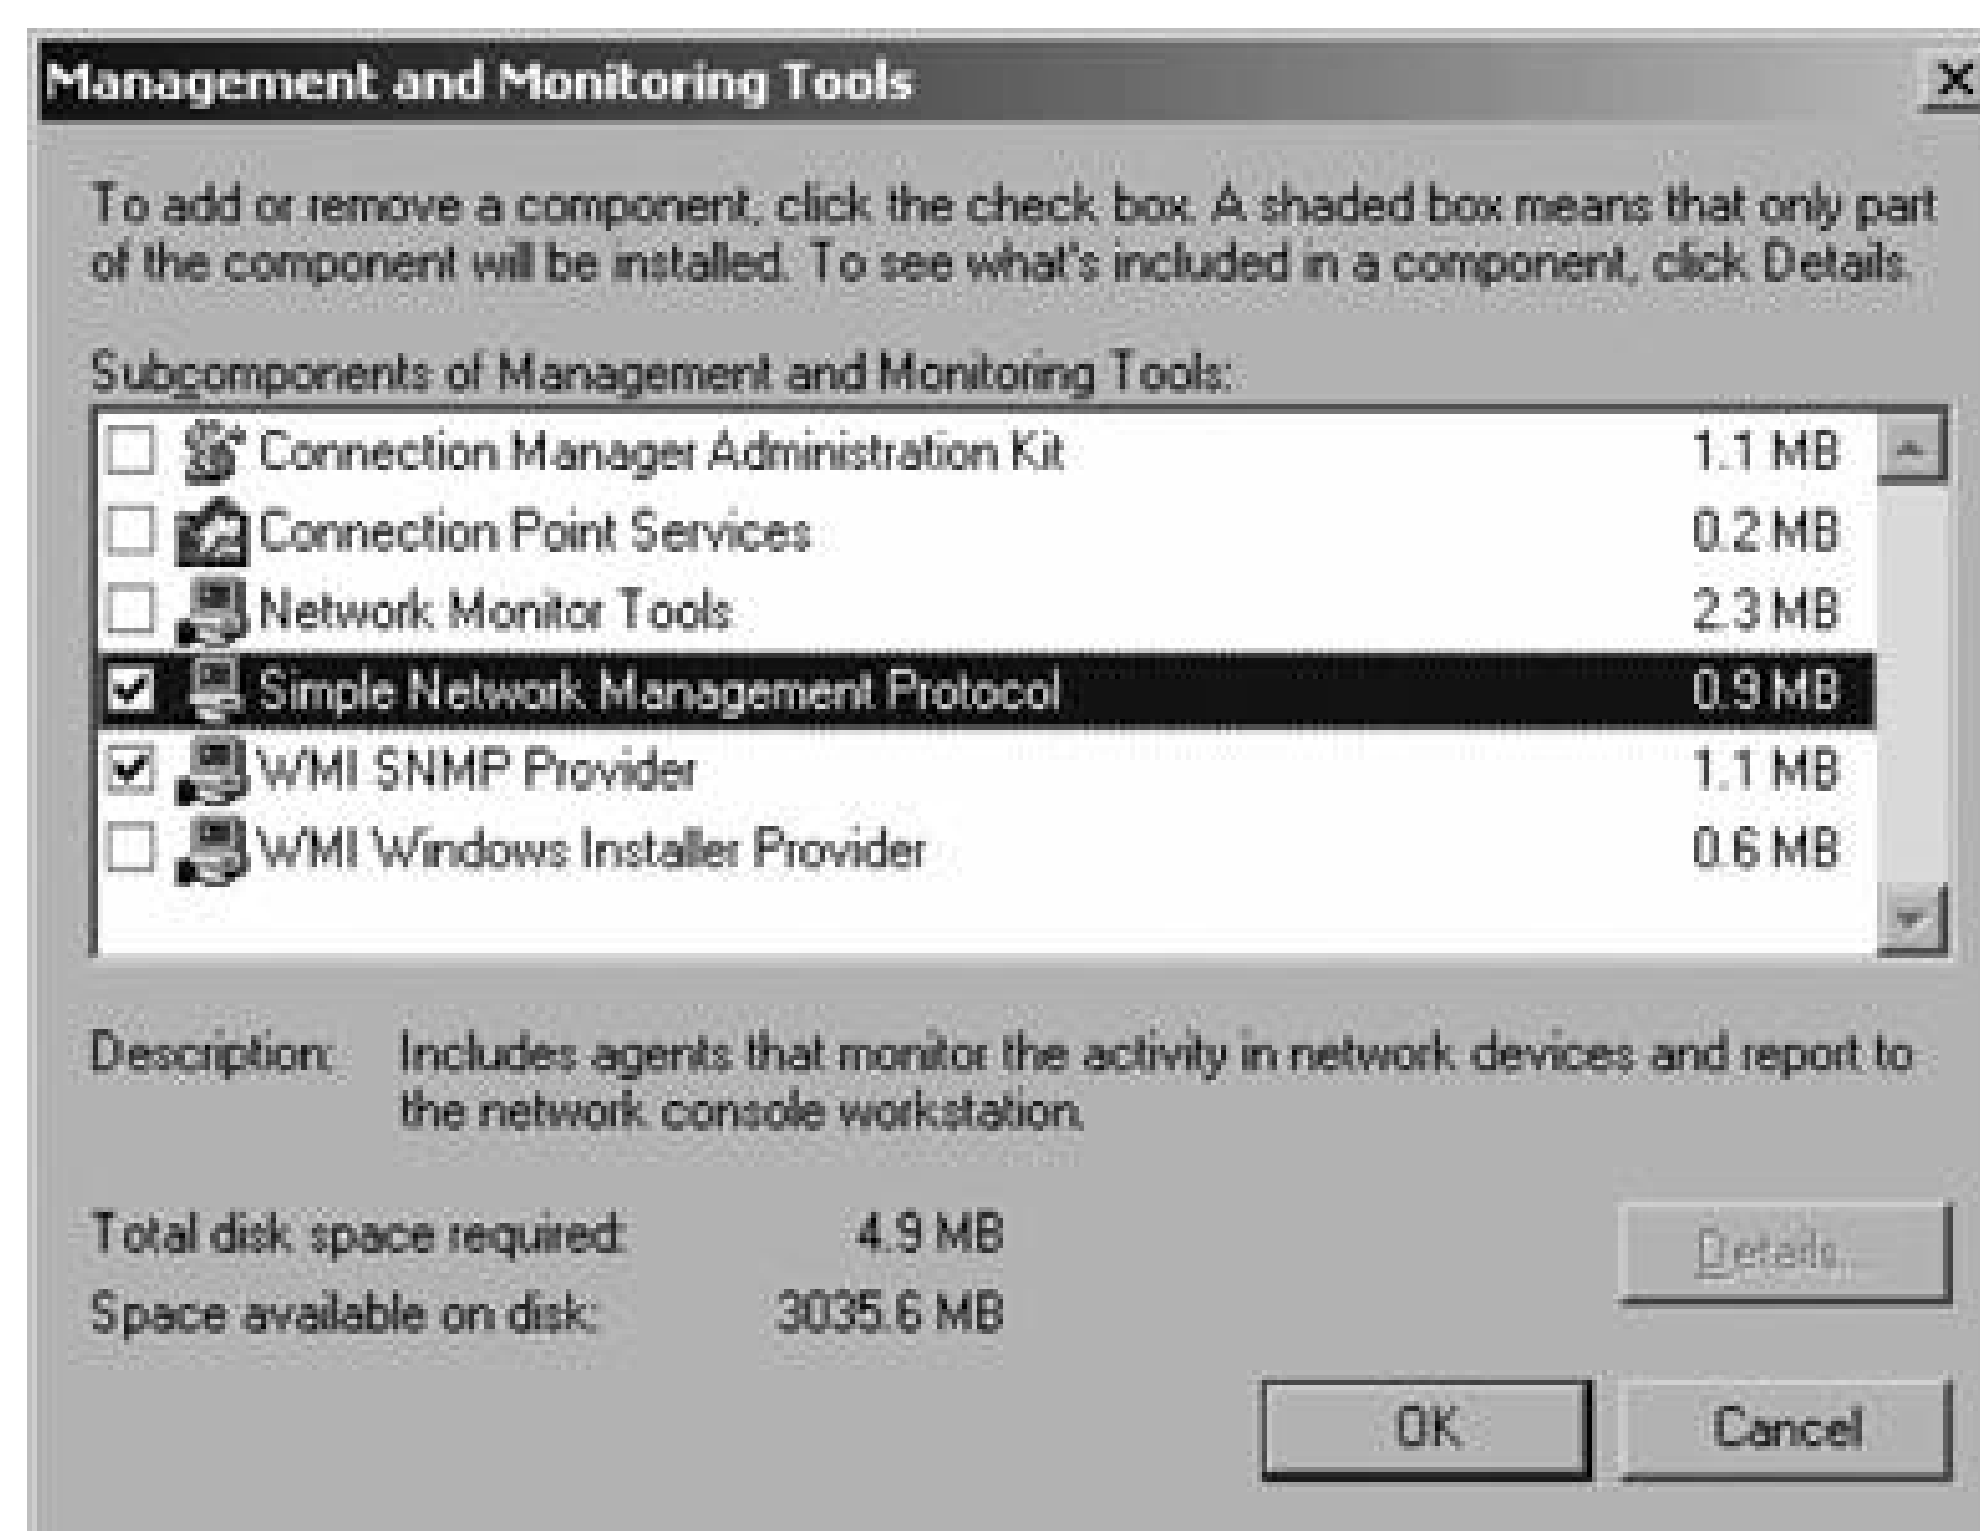


Figure 10-4. The Agent configuration tab of the SNMP service

The Contact and Location data are optional, but in a widely distributed network with a lot of devices, it is easy to forget the physical location of a given computer.

The rest of the properties in the Service box on this tab apply to the role that this SNMP agent plays in your overall SNMP configuration:

### *Physical*

Select this if this computer will be used for SNMP monitoring of physical devices such as fans or power supplies.

### *Applications*

Always select this option. It enables monitoring of any application that uses TCP/IP.

### *Datalink and subnetwork*

This is used for monitoring bridging devices.

### *Internet*

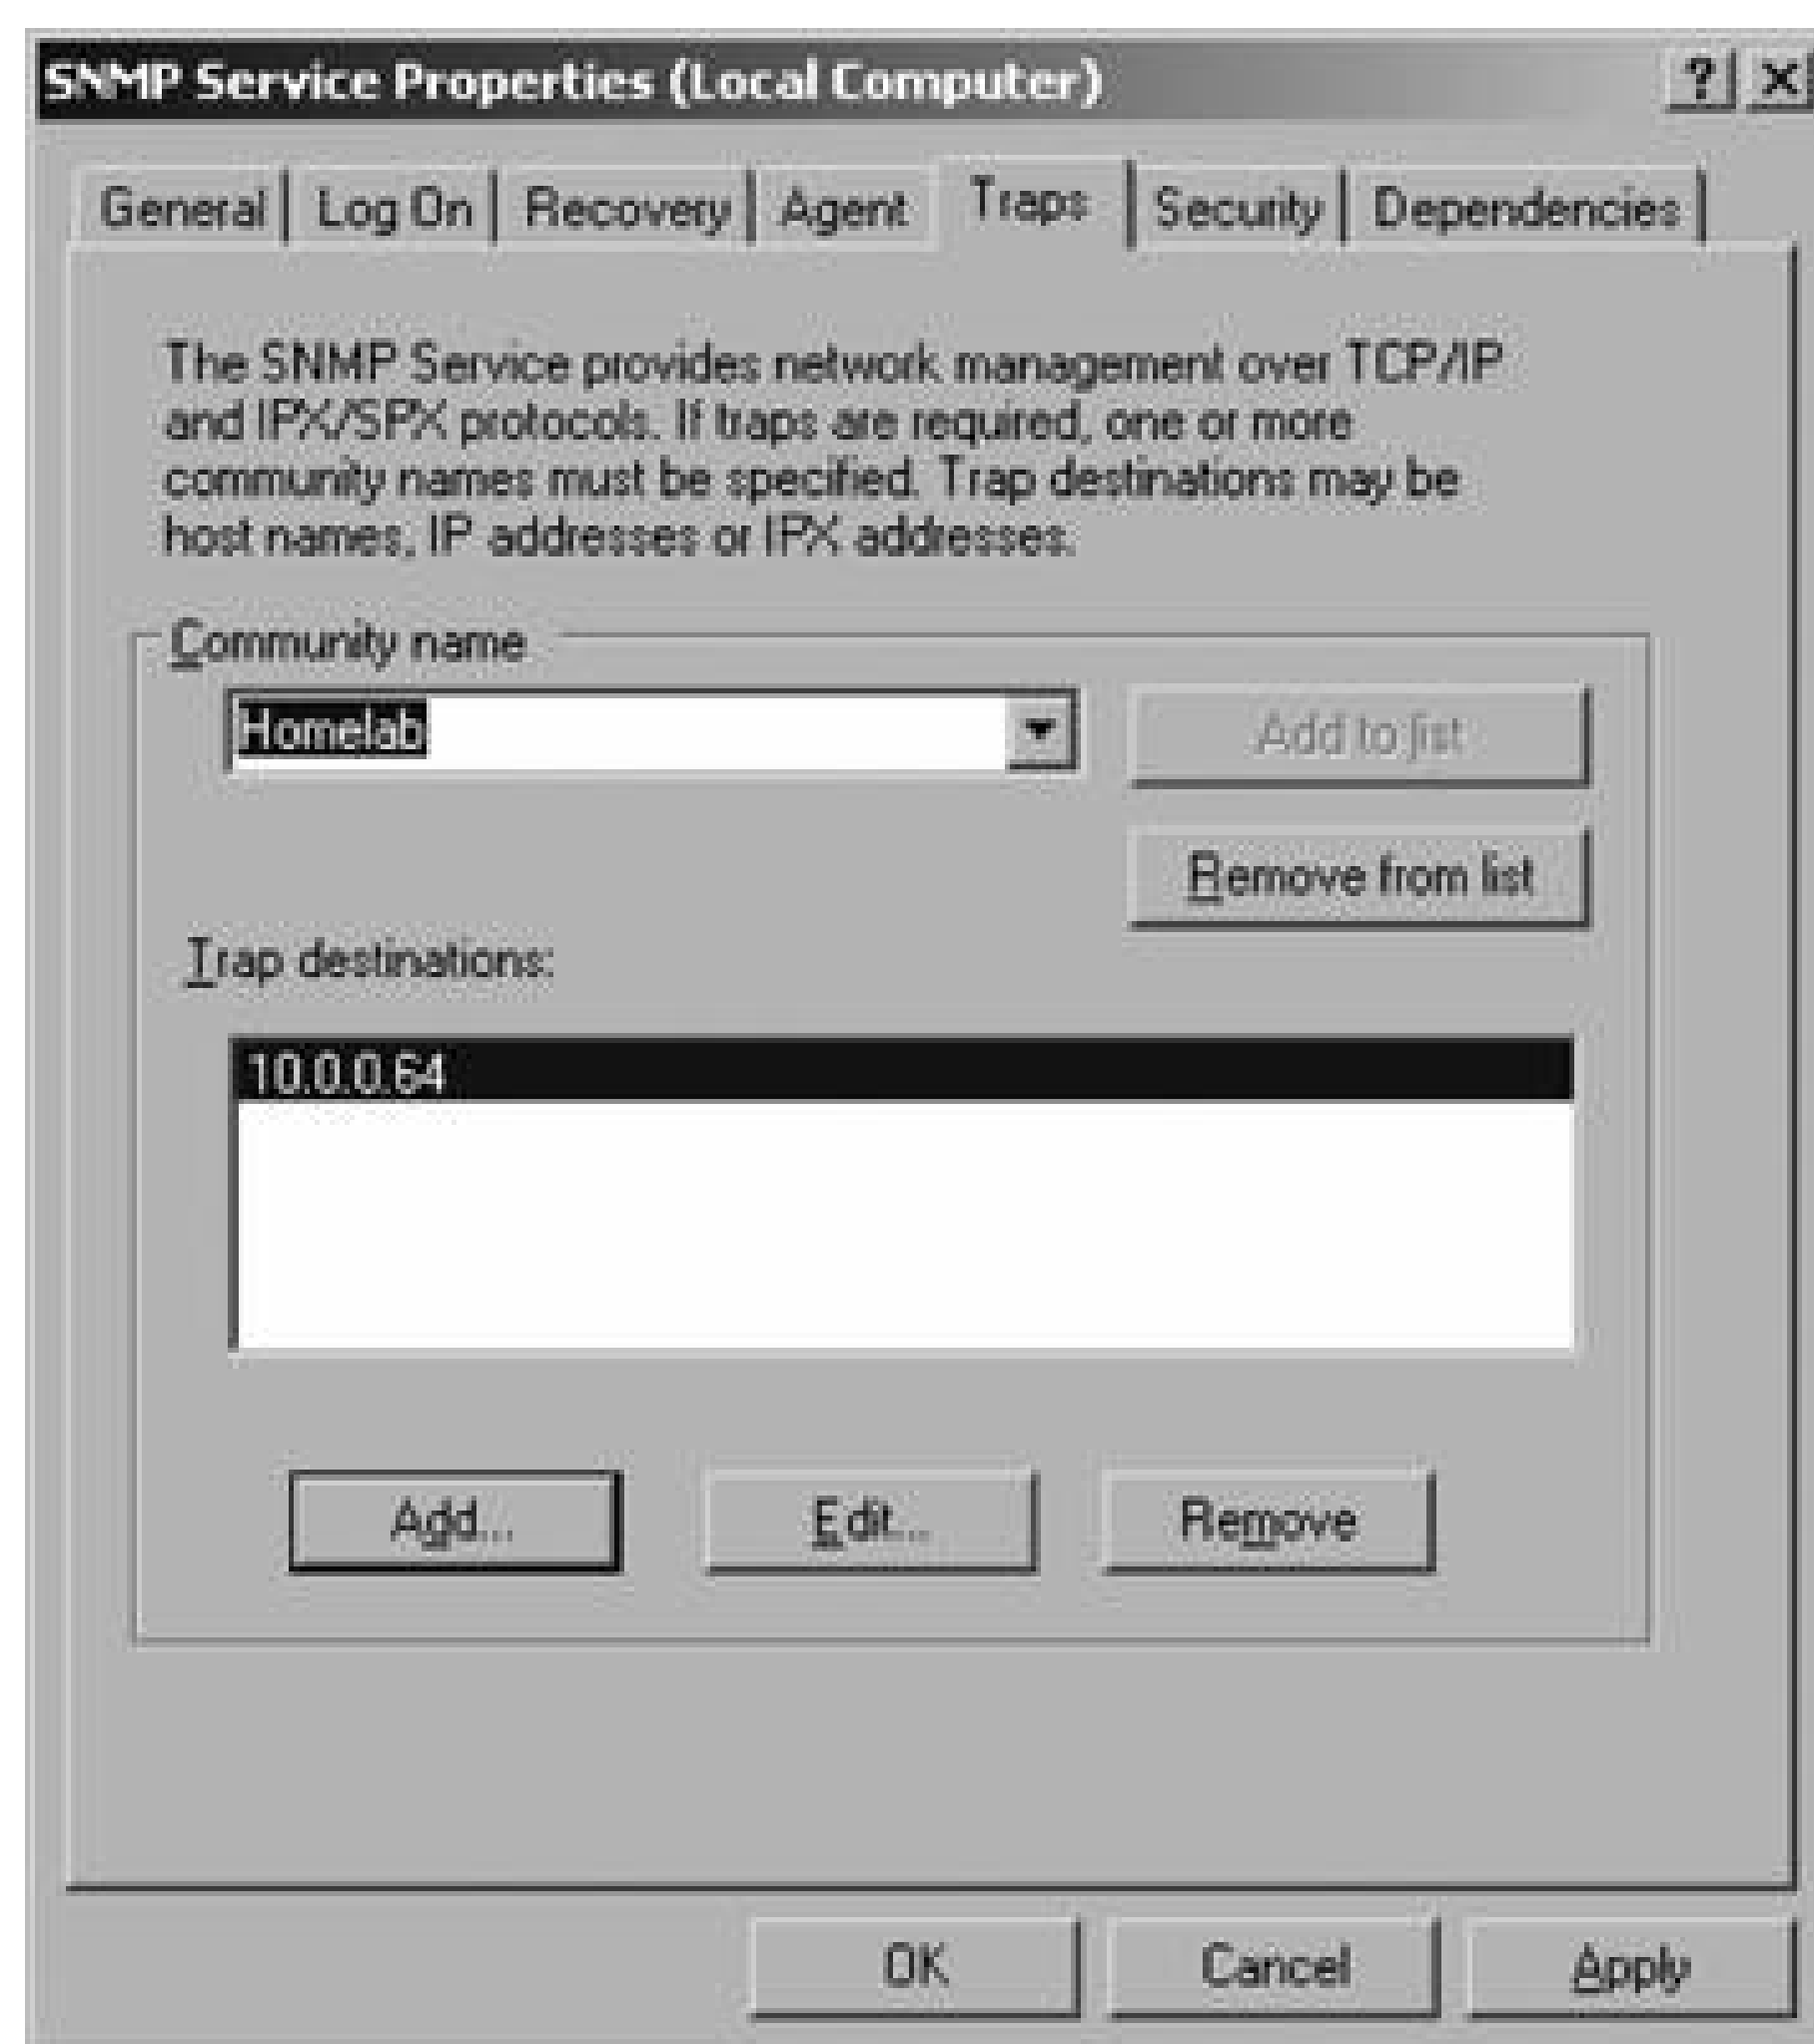
Windows servers can be configured to act as routers between IP networks. Select this option if this is how the computer is configured.

### *End-to-end*

Like the Applications option, this should always be selected. It runs if the server is an IP host. To configure:

- a. Select the Traps tab ([Figure 10-5](#)). The Windows SNMP agent can belong to multiple communities for the purposes of sending traps. Start by entering the desired "Community name," then click "Add to list."

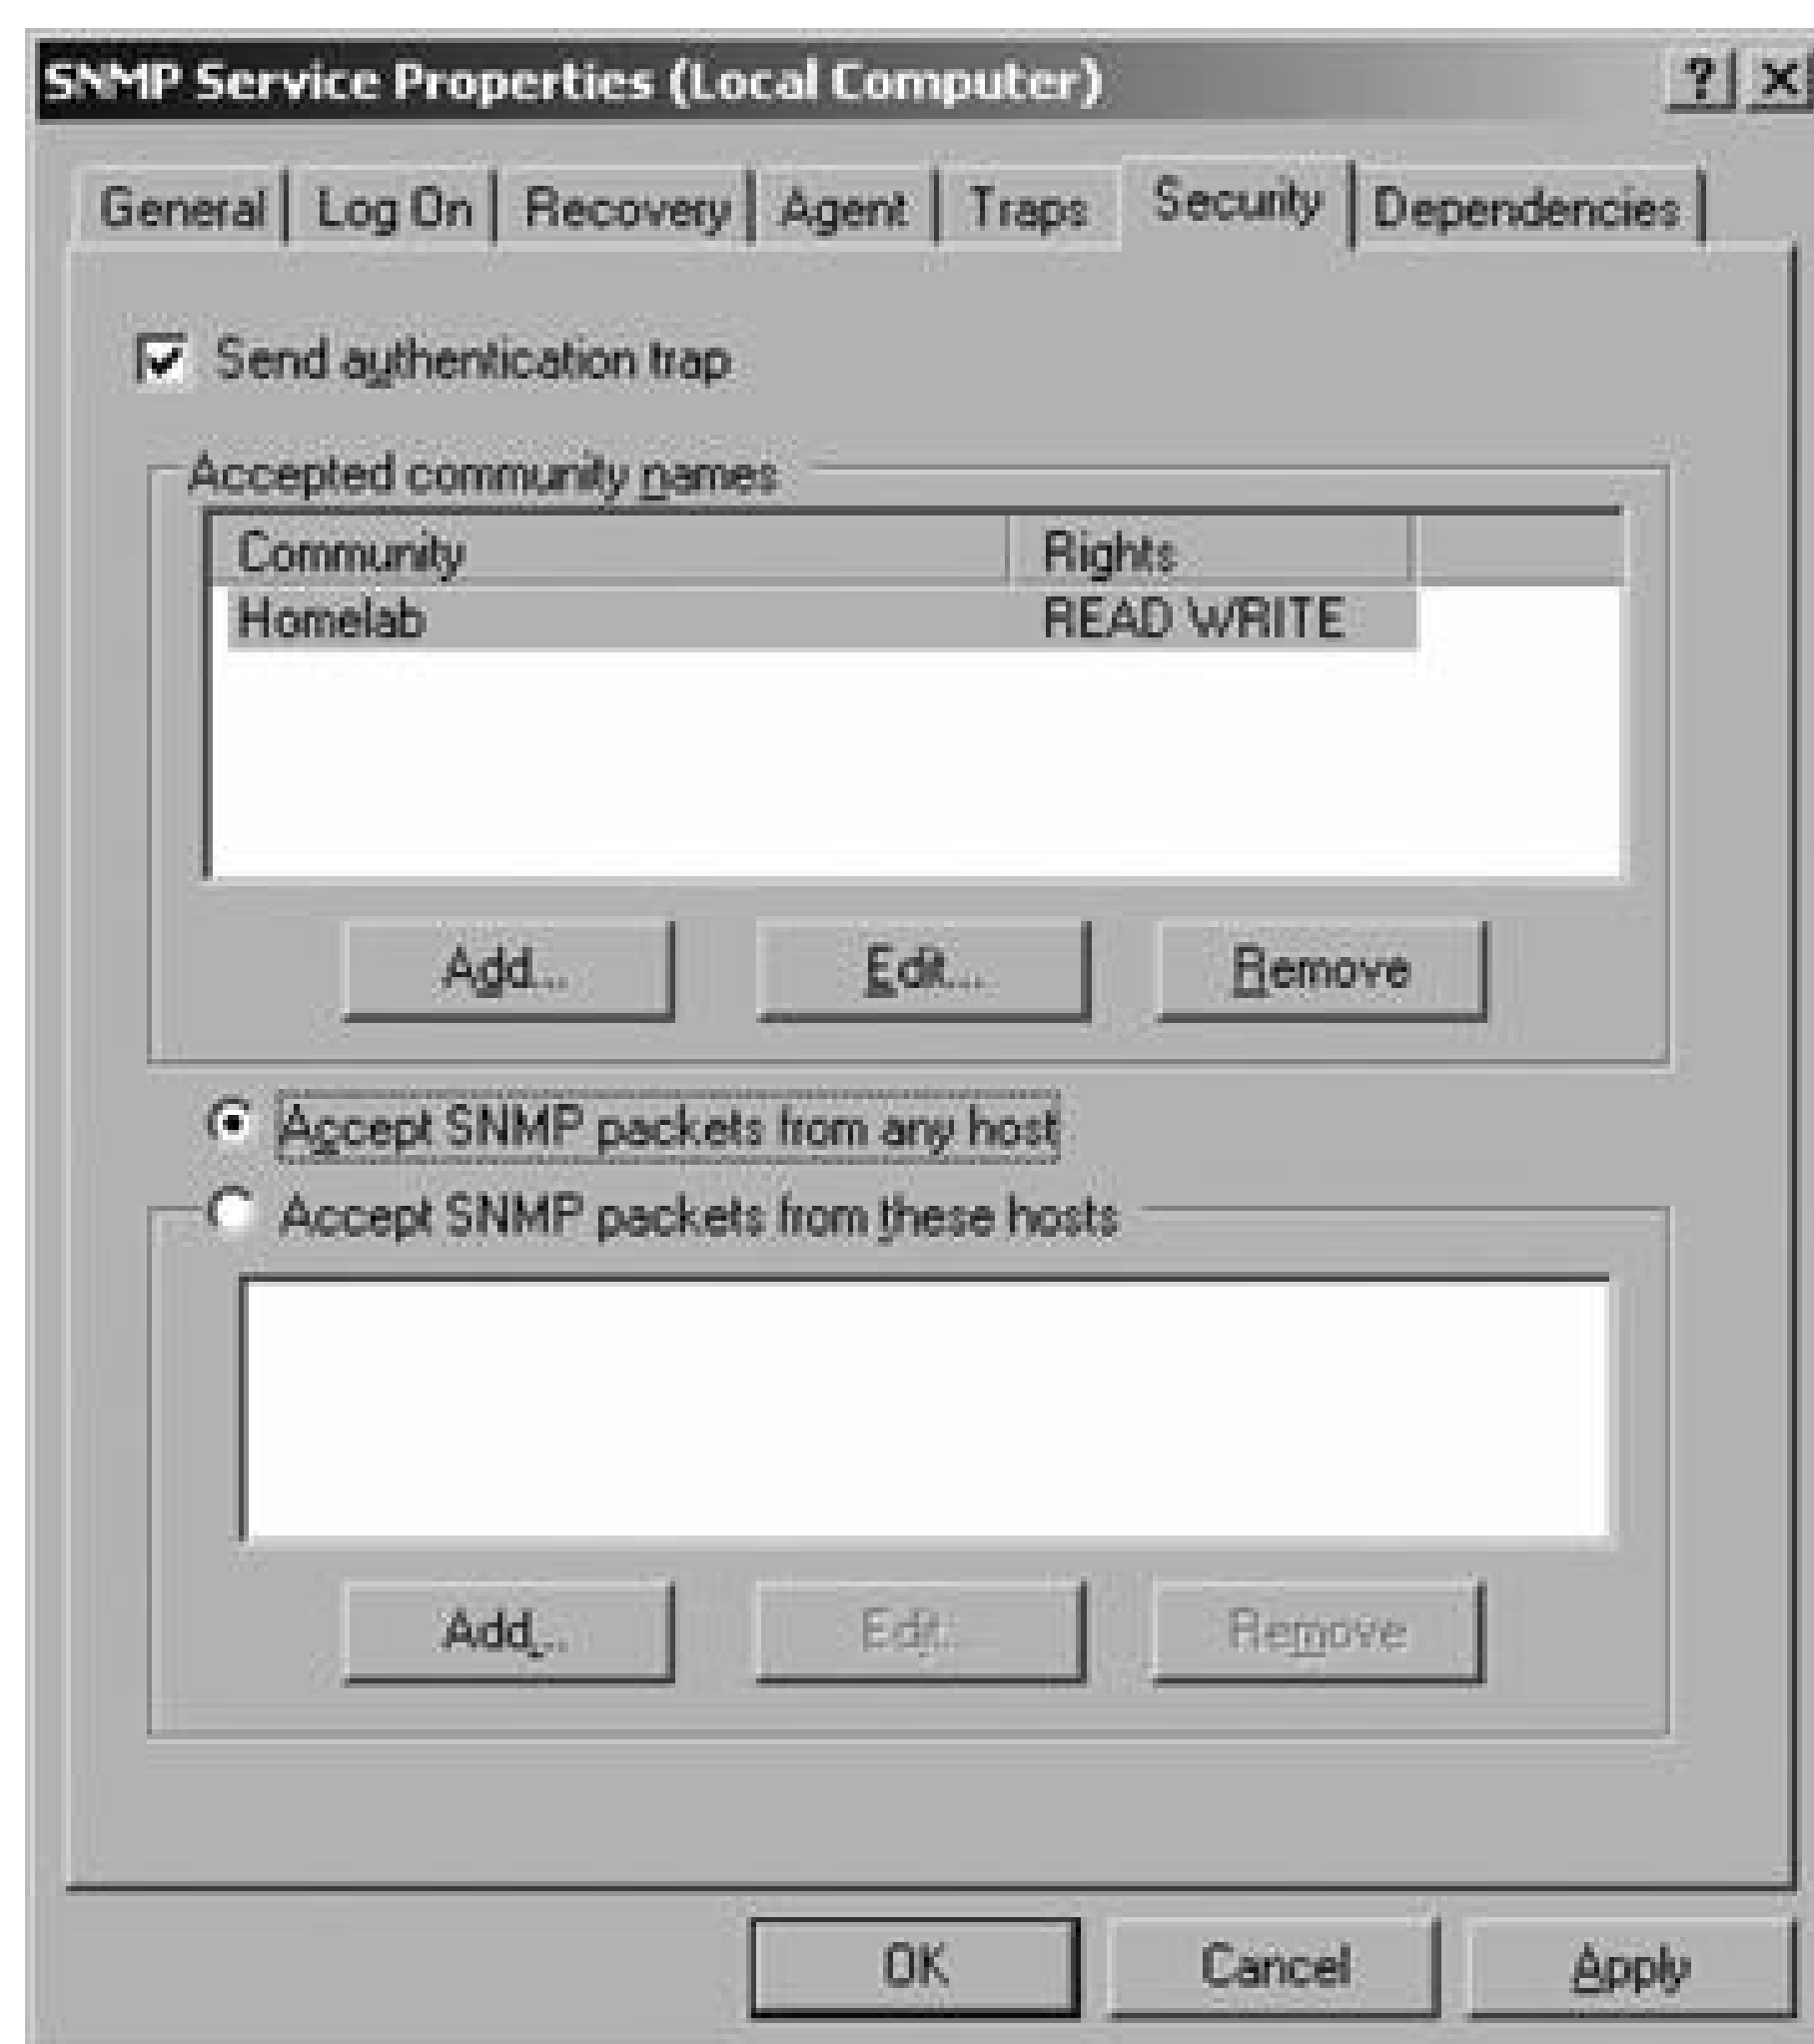
Figure 10-5. Defining the SNMP community name and destination host IP address



- b. Click the Add button under the "Trap destinations" text box.
- c. Enter the IP address of the host that you want the traps to go to.
- d. In this example, the agent machine points to itself since it is also the SNMP management host.
- e. Select the Security tab ([Figure 10-6](#)).

Figure 10-6. Configuring basic security for the SNMP host





The settings on this tab cover the reception of traps. Remember that SNMP is not secure by default. There is no Kerberos authentication between agents and management servers, there is no encryption of packets on the wire, and all communications are sent in plain text. The private and public community names are well known, so using those provides no security.

However, there are some minimal steps that you can take to improve security:

1. Send an authentication trap. Every trap that comes in includes the community name of the agent that sent the trap. If that community name does not match the community name of the management host, its authentication is considered to have failed. The agent on the management host can generate a trap (which it sends to itself) if this happens. Select this checkbox if you want to receive the "Authentication failed" traps.
2. Create accepted community names. There are two types of SNMP communication: management communication and trap communication. Management communication uses port 161 and it involves read and write requests that originate from the management host and are sent to an agent. Through the management communications, the management host can query an SNMP device's configuration information that is stored in the device's MIB. It can also write new values to the device if write access has been granted. For this example, the community name that this agent will accept management communication from and the type of permission to grant to it has been entered.
3. Accept SNMP packets from any host/these hosts. The default setting for this is to only accept packets from the local host. For production networks, I strongly recommend that you explicitly define the host IP addresses to accept SNMP packets from. This will help lock down SNMP.

communications, as all other hosts are excluded. In the case of the management host itself (*homemomserver3*), it accepts packets from any of the agents, and since it operates on a private network with a limited number of hosts, the network is not unduly exposed by selecting the "Accept from any host" option.

### 10.3.1. Compiling MIBs into Windows with SMI2SMIR

For MOM to be able to interpret the data sent in the SNMP traps, the data mappings of the OIDs in the MIB must be translated into the WMI namespace, where a device-specific sub-namespace is created for SNMP and SNMP devices. The WMI namespace, like the MIB namespace, is arranged in a hierarchy with the *\root* namespace at the top. If you have ever done any scripting against WMI, you are probably familiar with the *\root\cimv2* namespace. The common information model version 2 (*cimv2*) contains local configuration information that can be can queried. The SNMP namespace is *\root\snmp\localhost*.

Every vendor that sells an SNMP-enabled device must provide an MIB file to go with it. These MIB files are compiled and written into WMI by the *SMI2SMIR.exe* (pronounced smee-to-smear) command-line tool. This tool is installed when you install the SNMP and WMI SNMP Provider Windows components. It is in the *%SystemRoot%\system32\wbem\snmp* directory. SMI2SMIR has quite a few flags and not all are that well documented, but you are only interested in the */l*, */a*, and */t* options.

For example, on the *homemomserver3* server the MIB file of a Linksys device needs to be compiled to receive traps:

```
C:\Windows\system32\wbem\snmp\smi2smir.exe /a /t linksys.mib
```

The */a* flag checks the syntax of the MIB file and the */t* flag generates the notification classes or extended notification classes respectively. Either way, you will probably see some errors from SMI2SMIR; however, don't worry as long as the errors are labeled warning and at the end of the compilation the tool returns "smi2smir : Loaded "<mibfilename>.mib" successfully into the SMIR."

To check that the MIB file has been successfully loaded, run the SMI2SMIR command with the */l* option, which lists all of the MIB files that have been compiled and loaded. In addition to the ones you load yourself, you should see RFC1213\_MIB. This is a generic MIB file.

### 10.3.2. Confirming SNMP Traps with SNMPUTIL

When configured for SNMP, MOM is being used strictly as a trap reader. It cannot query (unless you write some code to do so) an SNMP-enabled device for its OID values. There are tools that provide extensive parameter read, write, create, and trap analysis capabilities. These full-featured tools usually come with hundreds of pre-compiled MIBs and are meant for managing very large SNMP implementations; Castle Rock SNMPc Network Manager and OidView are good examples of SNMP management tools. These two products offer full-featured trial versions that are fairly straightforward to use.



Microsoft provides a very simple SNMP monitoring tool called SNMPUTIL. It is a command-line tool that you can use to query an SNMP device and to listen for all incoming traps on a server that is configured for SNMP. This tool provides a quick, no-nonsense way to confirm that a server is receiving and interpreting SNMP traps and that any device that needs to be monitored via SNMP is listening and will respond to queries from the server.

SNMPUTIL functions in two modes. The first is query mode, where you specify the `get`, `getnext`, or `walk` options; the SNMP device to query; the community name; and the OID that you want to query.

```
usage: snmputil [get|getnext|walk] agent community oid [oid ...]
        snmputil trap
```

## get

Along with the agent IP address and OID, it returns the value of that single OID.

## getnext

Returns the value of the OID that comes after the query with the `get` flag.

## walk

Starts at the OID provided and returns all of the values for all OIDs of a device.

Using `get` from *homemomserver3*, SNMP returns the name of a network device listed in OID 3955.1.1.0:

```
C:\SNMPMIB>snmputil get 10.0.0.23 Homelab .iso.org.dod.internet.private.enterprises
.3955.1.1.0
Variable = .iso.org.dod.internet.private.enterprises.3955.1.1.0
Value = String Linksys BEFSR81v3
```

By using `walk` and stopping the OID definition at `enterprises`, SNMPUTIL returns all of the OID values for the Linksys BEFSR81 v3.

Invoking SNMPUTIL with the `TRap` flag starts the utility in listening mode. Here is an example of the traps generated by a Cisco 675 router when the device is rebooted:

```
C:\Documents and Settings\Administrator.HOMELAB>snmputil trap
snmputil: listening for traps...
Incoming Trap:
    generic = 0
    specific = 0
```



```
enterprise = .iso.org.dod.internet.private.enterprises.9.10.1.1
agent = 10.0.0.1
source IP = 10.0.0.1
community = Homelab
```

By using SNMPUTIL with the `trap` flag, you can confirm that the server can successfully receive and translate these traps.

### 10.3.3. Testing SNMP into WMI with WBEMTEST

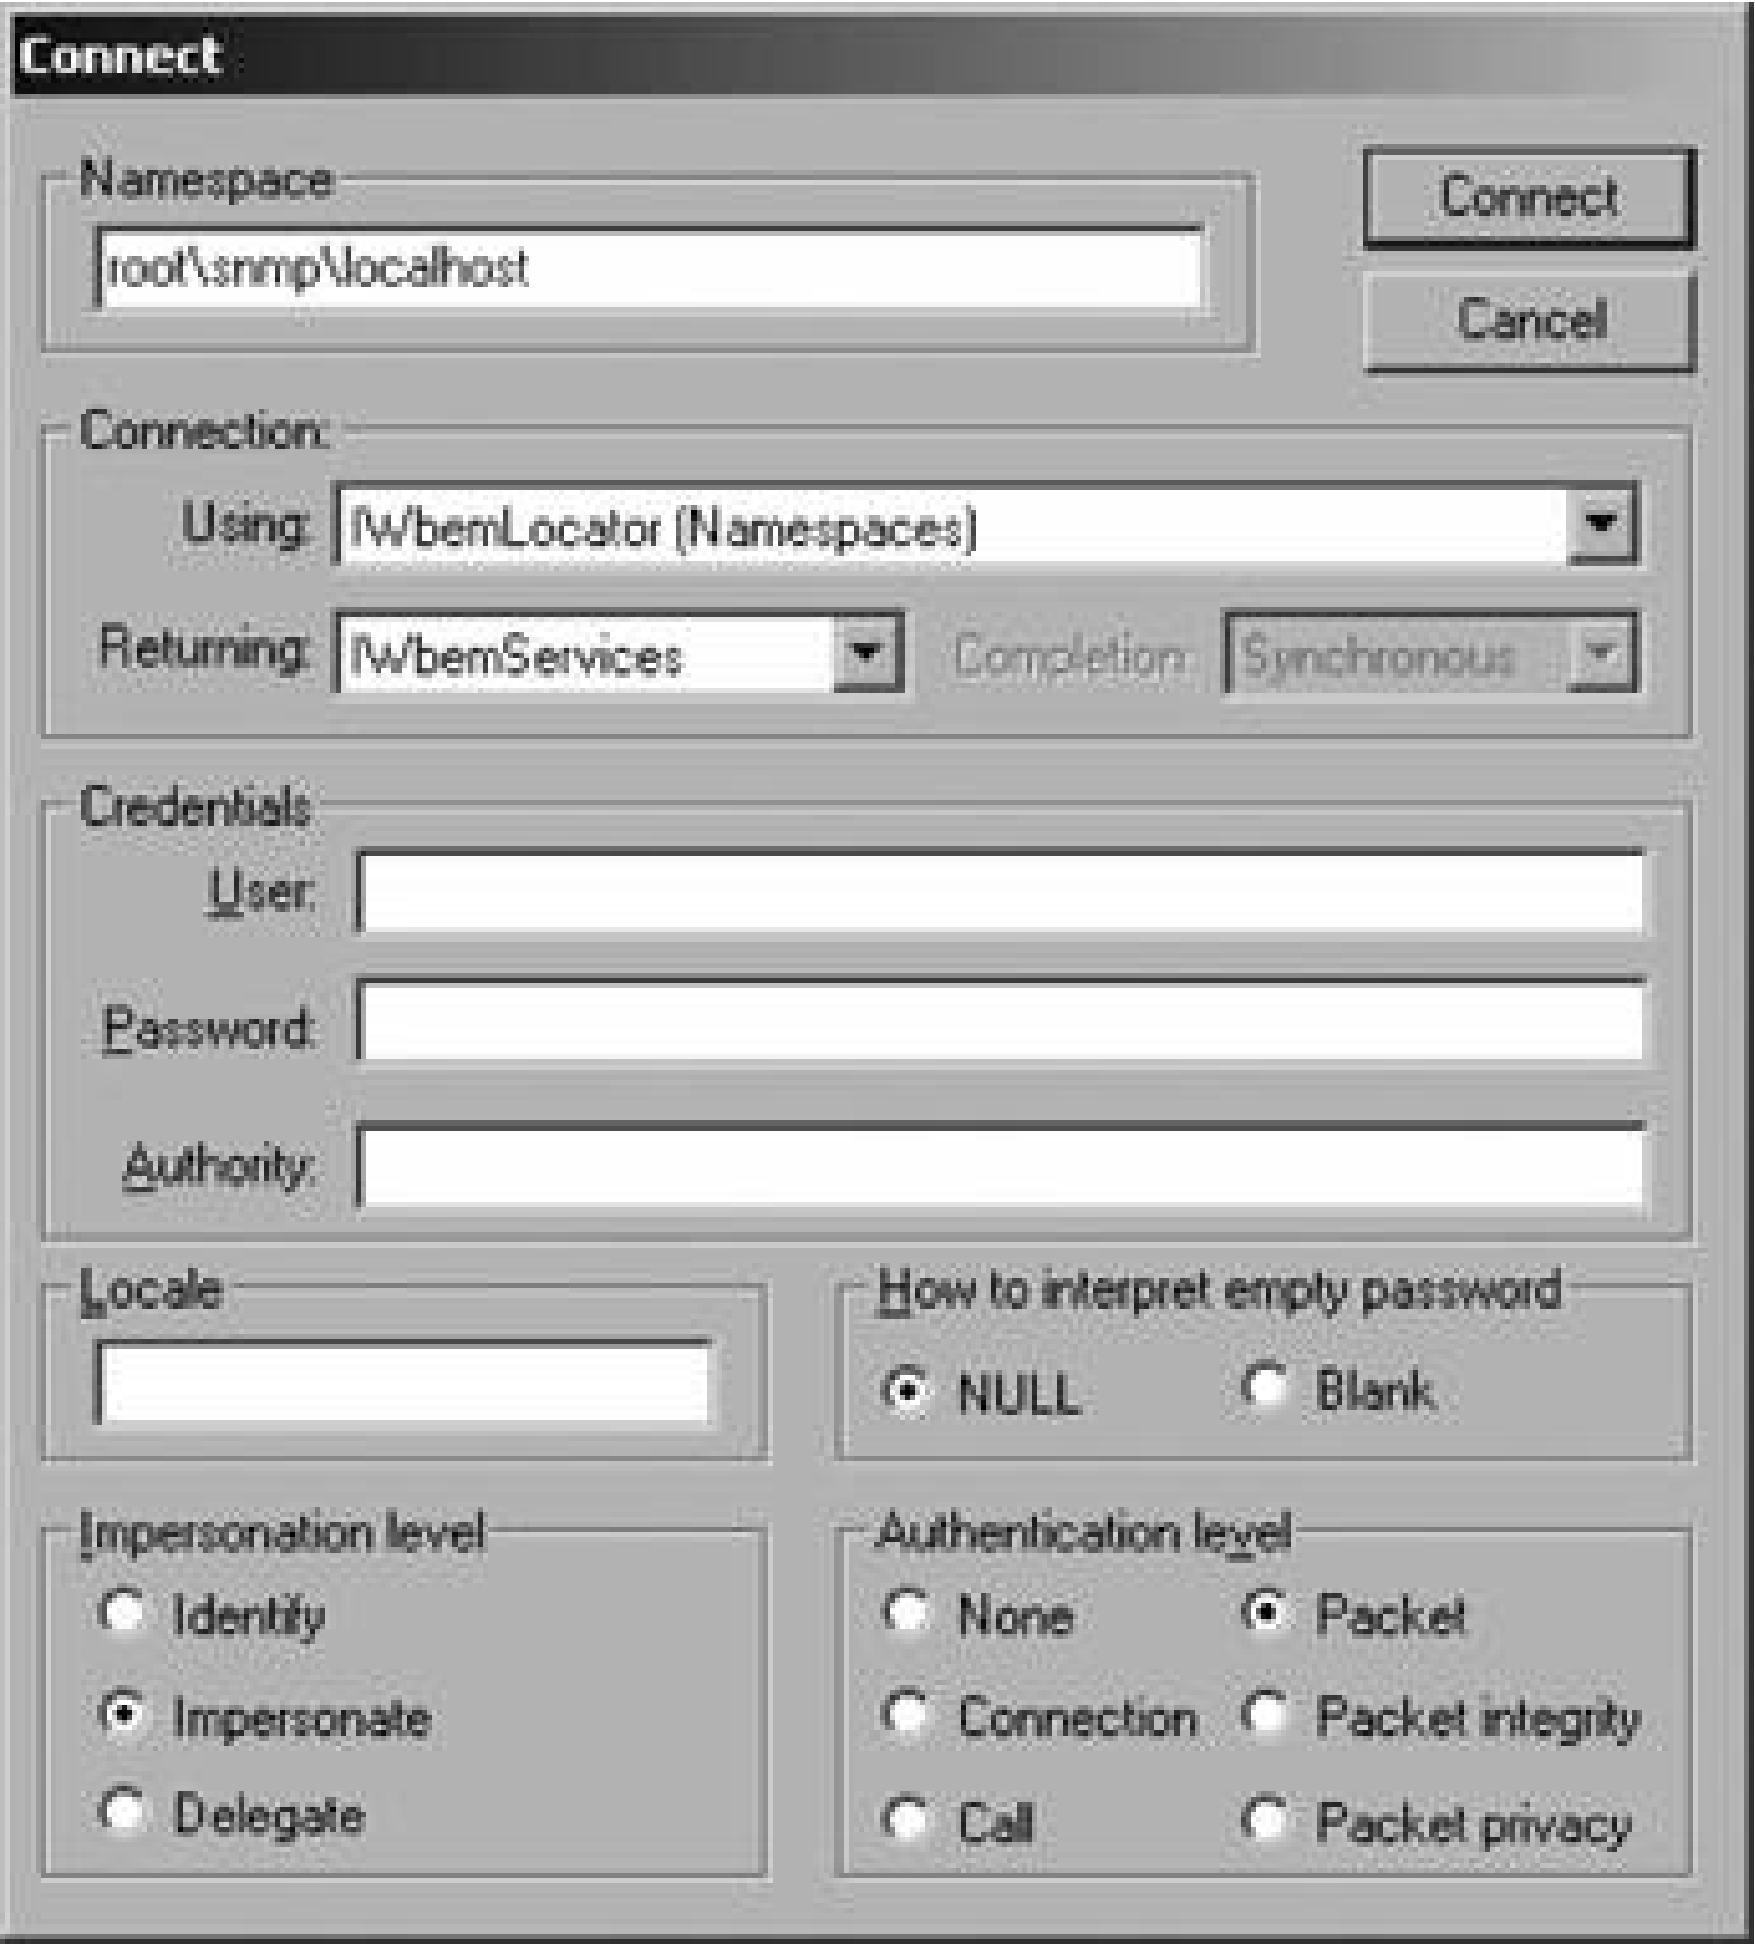
The WMI Tester tool is a graphical interface that is used to interact with objects in the WMI namespace. Its use here is to confirm that SNMP traps are making it through the SNMP service layer and into the WMI namespace. Once there, the MOM SNMP Extended Trap Catcher provider will make it available to the MOM agent on *homemomserver3*.

To start the WMI Tester tool from the command line, enter **WBEMTEST**. This brings up the tool shown in [Figure 10-7](#).

Figure 10-7. The WBEMTEST tool

To confirm that traps are making it to the WMI SNMP namespace, you must first connect to SNMP. Click on the Connect button and enter **root\snmp\localhost** as shown in [Figure 10-8](#).

Figure 10-8. Configuring the WMI namespace to connect to



Click Connect again to return to the main page, which will now show the WMI namespace that the tool is connected to, as shown in [Figure 10-9](#).

The WMI Tester tool can monitor the SNMP namespace in real time, notifying you when traps come through. Click on the Notification Query button to open the Query page [Figure 10-10](#). Enter the query **SELECT \* FROM SnmpExtendedNotification**.

Click Apply and the tool will monitor the namespace for all traps that come through. To test this, the Cisco 675 is rebooted to generate a cold boot trap, as shown in [Figure 10-11](#).

The traps are sitting in the WMI SNMP namespace and are waiting to be picked up by MOM.

### 10.3.4. MOM Configuration

Converting the traps into MOM alerts is now a matter of creating a rule that will monitor the MOM SNMP provider and generate alerts from the traps. You can create this rule in an existing rule group or in its own rule group. Since converting traps to alerts is a new ability, create a new rule group with new rules, then create a new computer group with the MOM SNMP Trap Catcher server (in this example, *homemomserver3*) as the only member. Once that computer group and rule group are

associated, the traps will start flowing into MOM. Using one of the existing computer groups that has broad membership is not a best practice because then the SNMP rule would be applied to machines that cannot make use of it.

Figure 10-9. Successfully connected to the SNMP namespace on the localhost

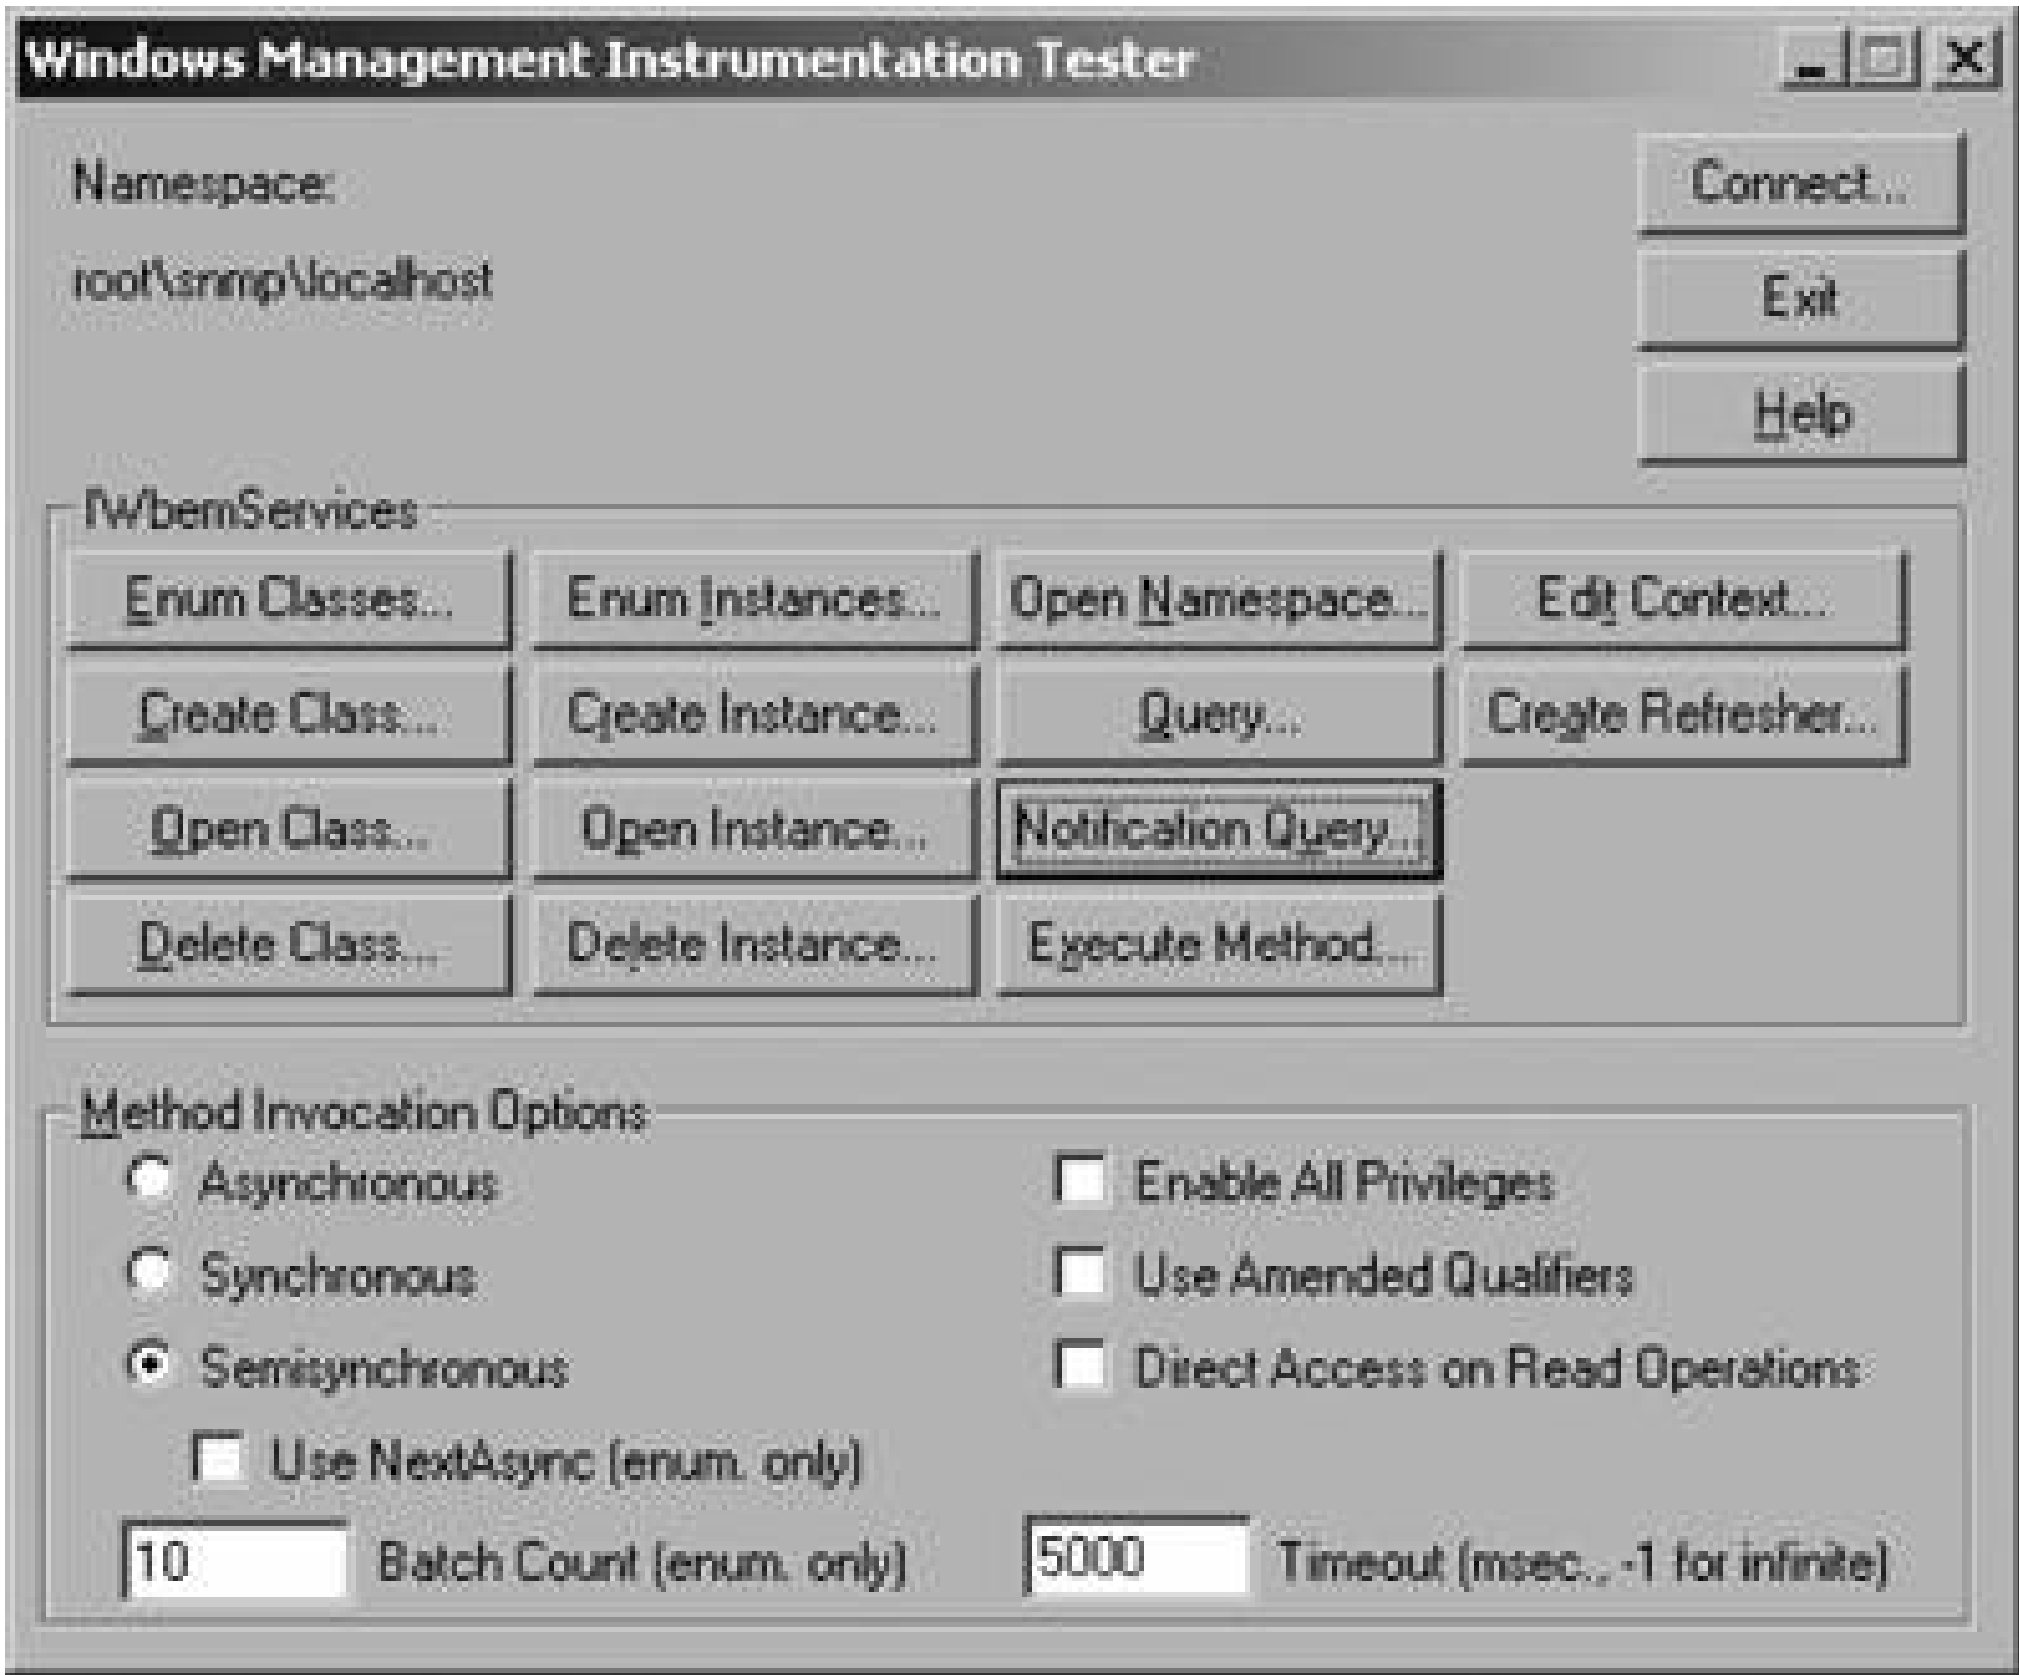
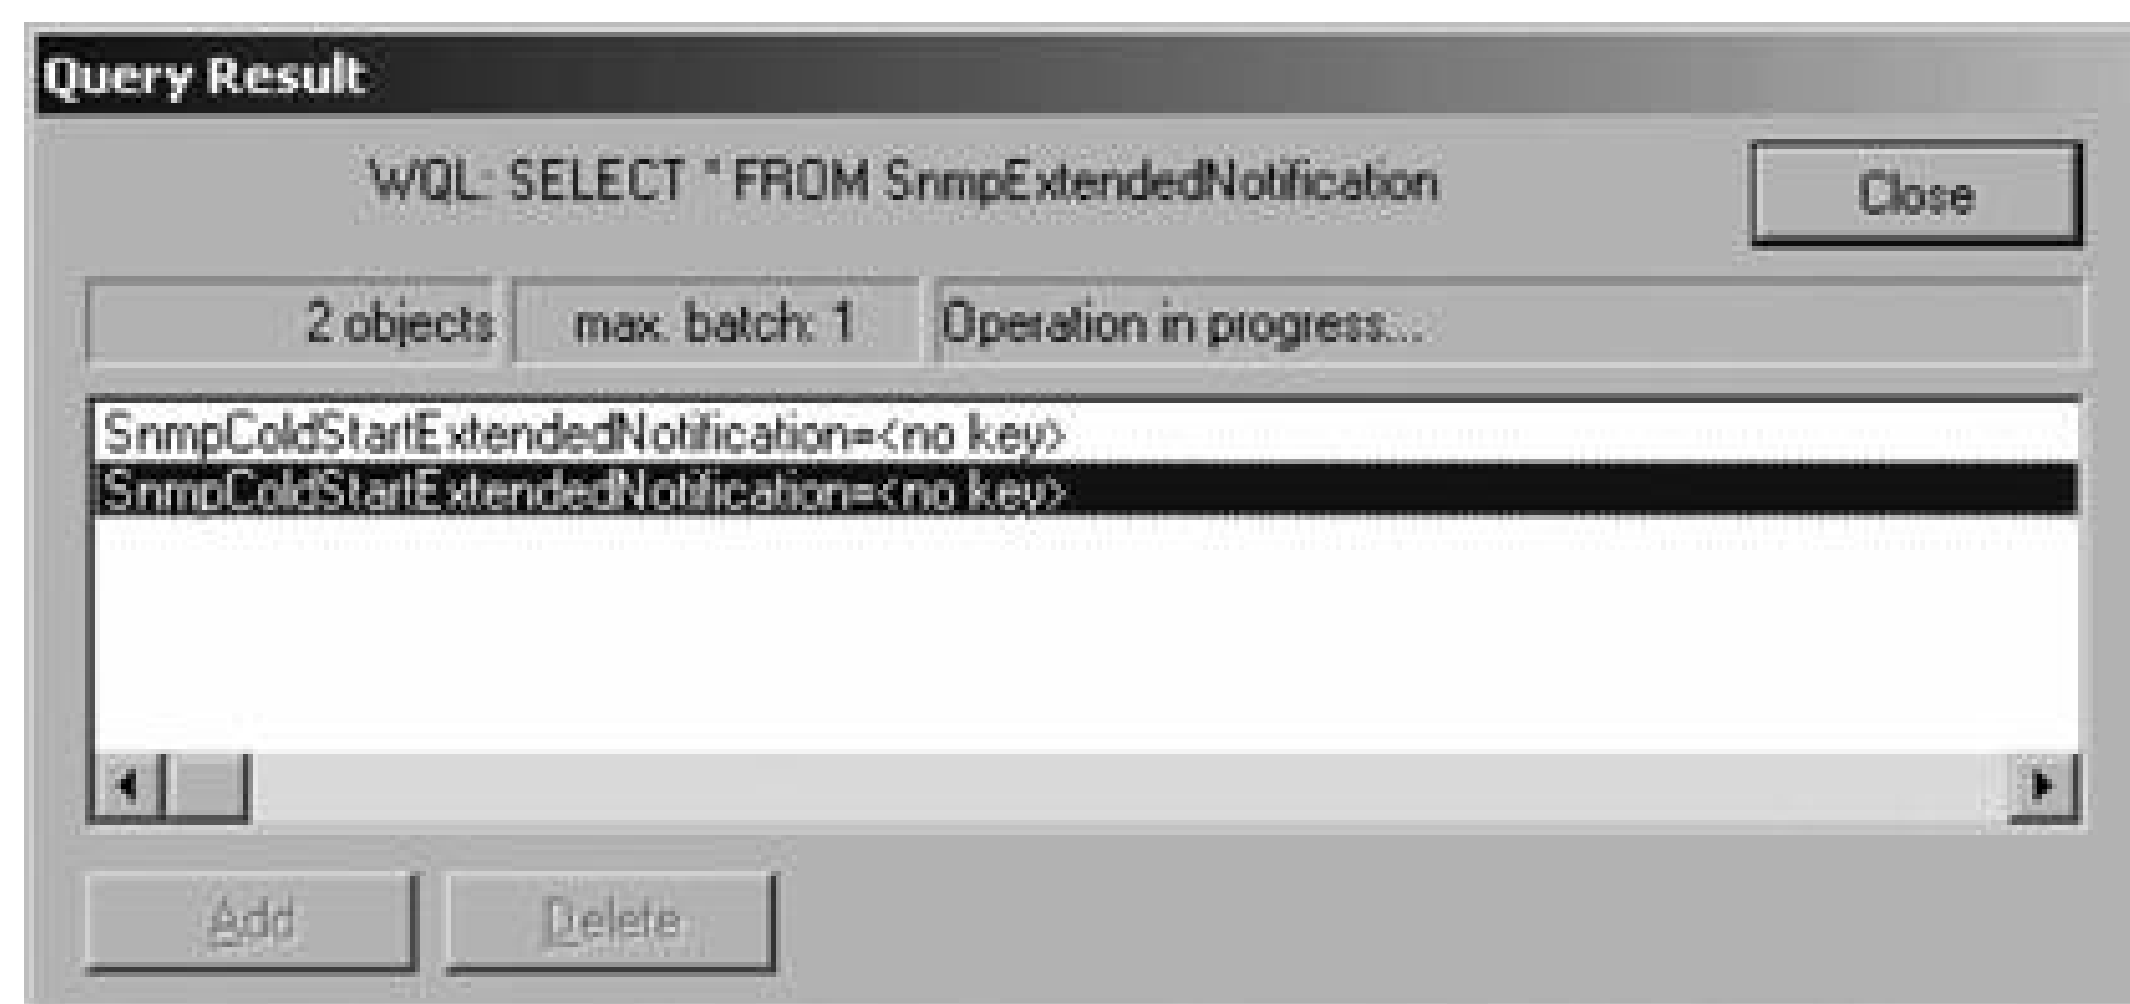


Figure 10-10. Configuring a query in the WMI Tester tool

Figure 10-11. SNMP trap passing through the WMI SNMP namespace





The first step is to create a new rule group in the Administrator console. Do this by bringing up the context menu for the Rule Groups folder and select the "Create Rule group" option. Give the rule group a name, such as SNMP, and a description. Click Next and then click Finish. Open the event rules object in the newly created SNMP rule group and create a new event rule as follows:

1. Select the rule type of "Alert on or Respond" to Event (see [Figure 10-12](#)), and click Next (see the section "[Types of Rules](#)" in [Chapter 4](#)).

Figure 10-12. Creating an event rule for SNMP traps

This brings up the Data Provider page, shown in [Figure 10-13](#). MOM comes with two preconfigured SNMP data providers: the SNMP Trap Catcher and the SNMP Extended Trap Catcher.

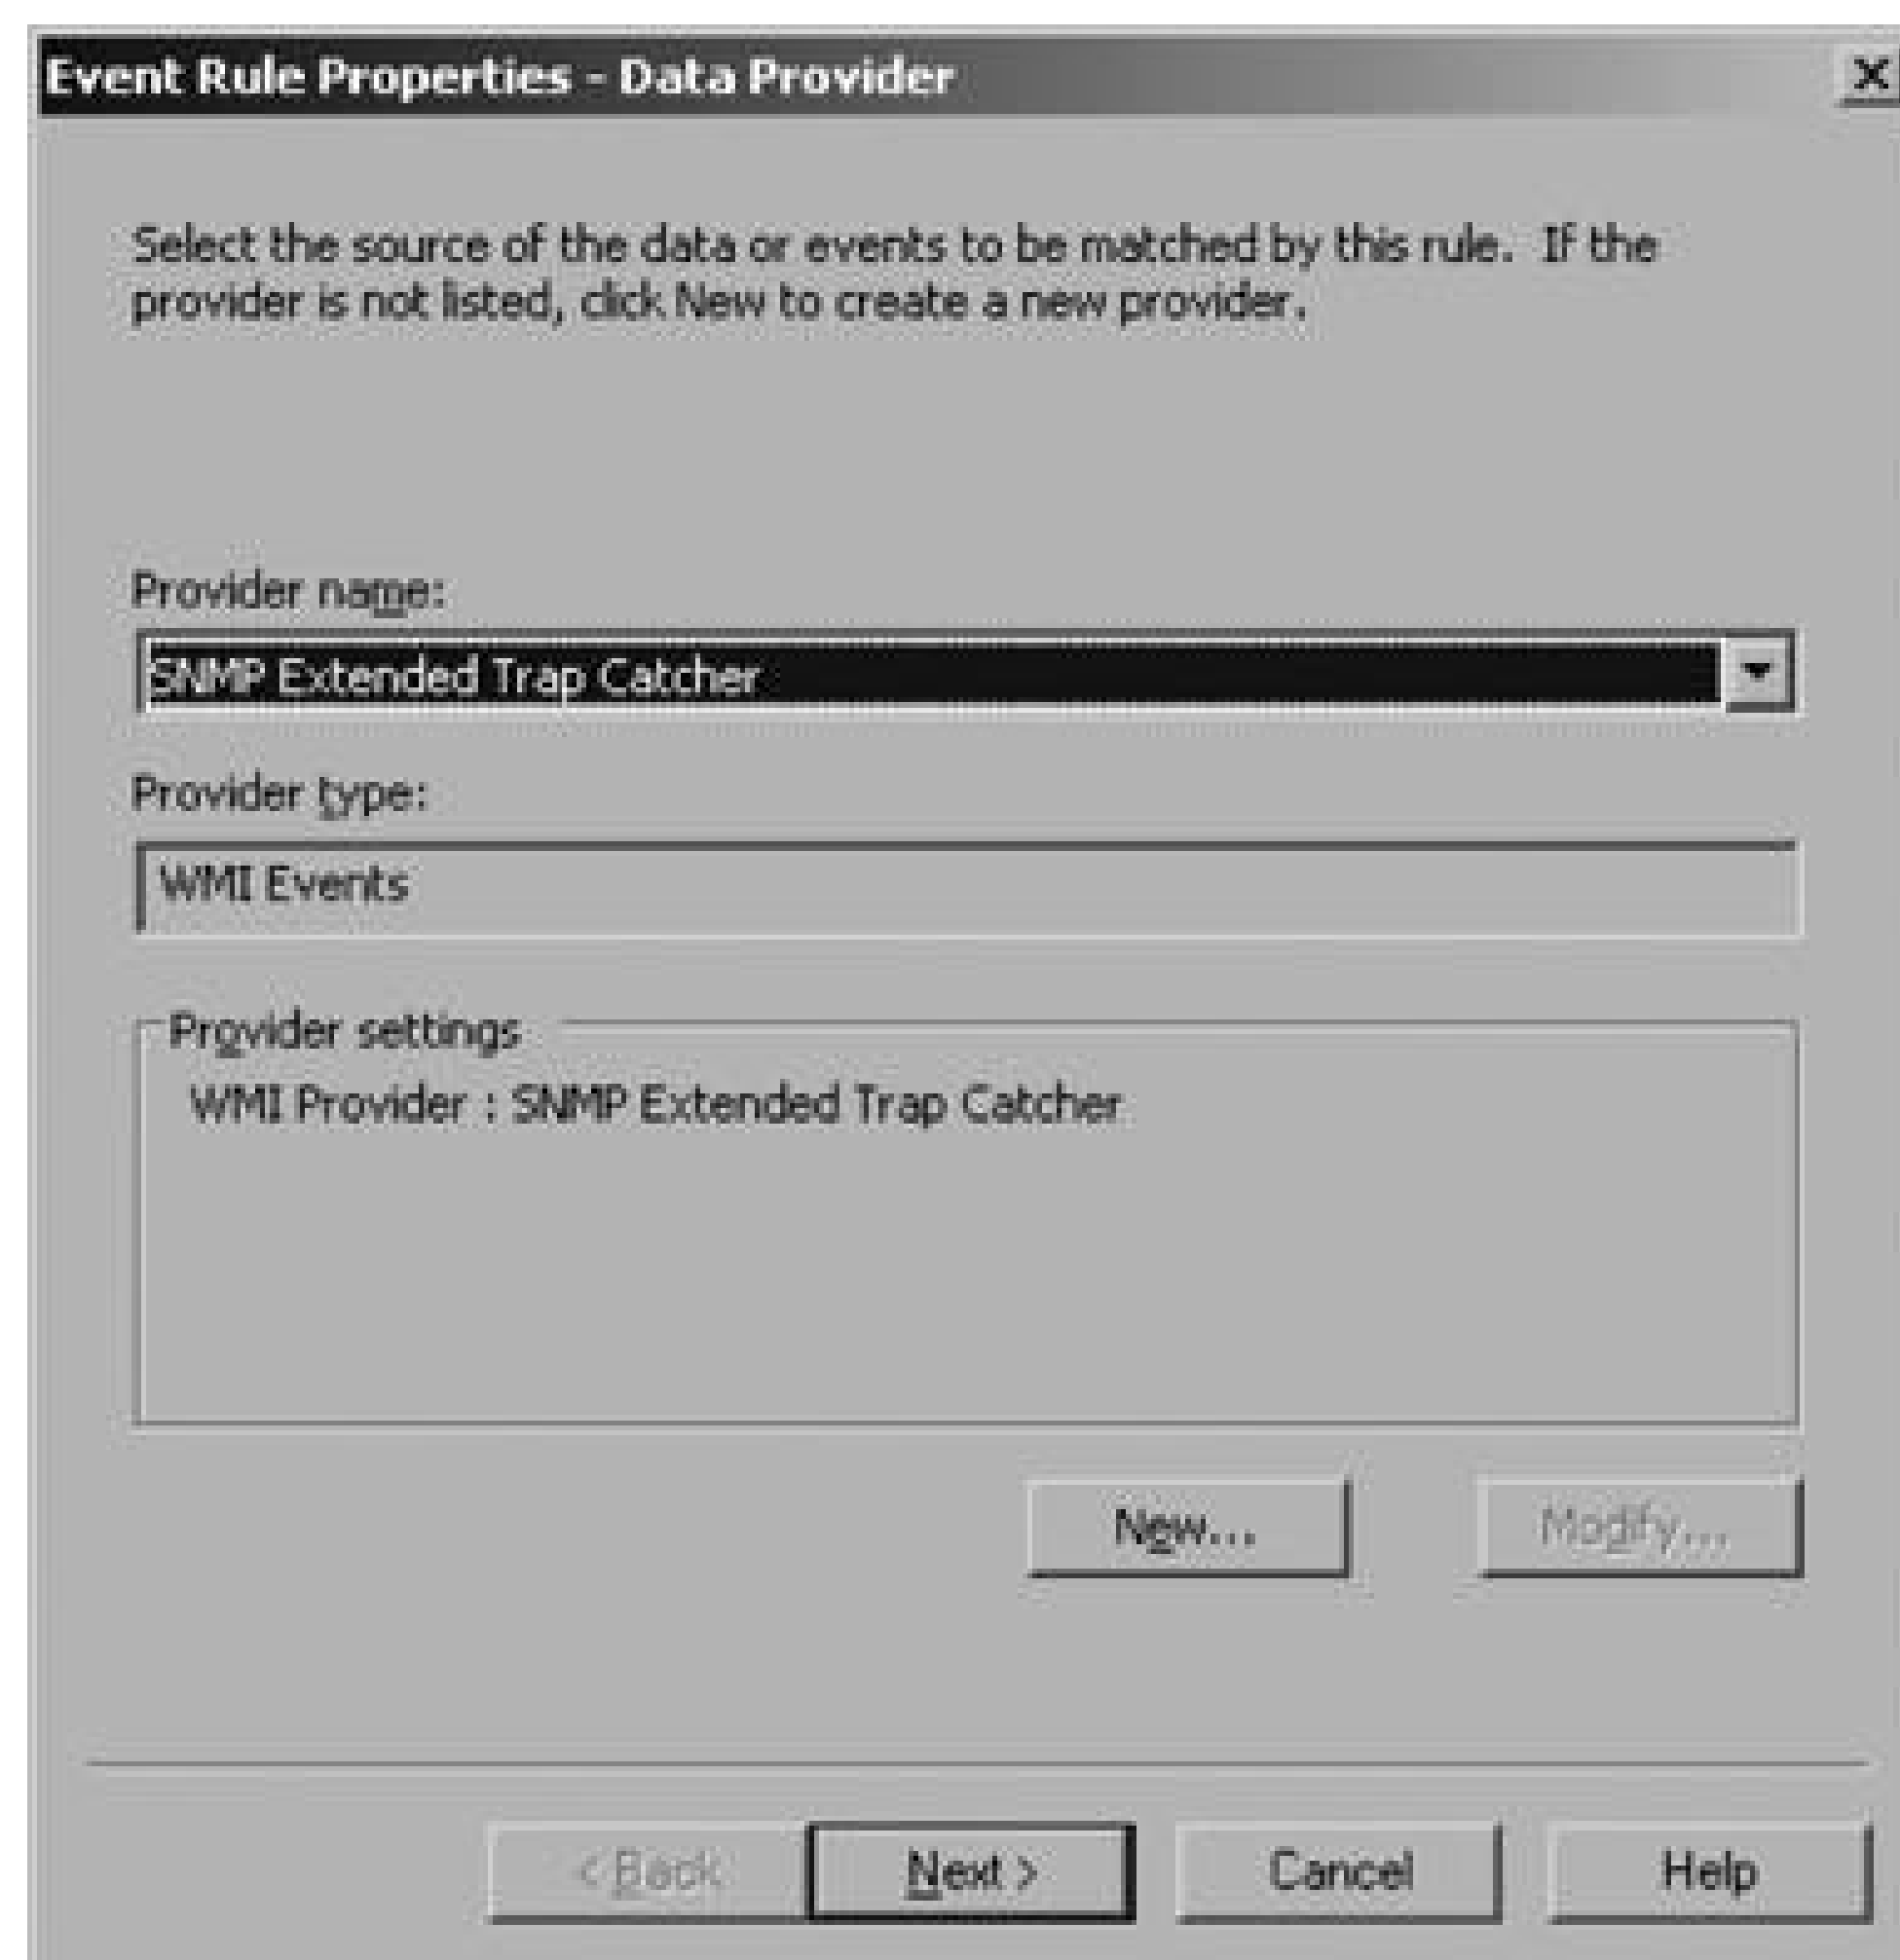
2. Select the Extended option to capture SNMP v1 and v2 traps. The non-extended version only catches SNMP v1 traps. Notice that the "Provider type" is WMI events. If you open this provider object (Administrator Console Management Packs Providers SNMP Extended Trap Catcher) you will see that the query string is the same one that you used in the WMI Tester

tool:

```
SELECT * FROM SnmpExtendedNotification
```

3. Click Next to bring up the Criteria page. Leave this at the default (blank) so that this rule will respond to all SNMP traps that come through. Click Next to bring up the Schedule page.
4. Accept the default on the schedule page, which is to "always process data." Click Next to open the Alert page (see [Figure 10-14](#)).

Figure 10-13. Selecting the SNMP Extended Trap Catcher



On the Alert configuration page, make sure that the rule is generating an alert. Don't bother with the "Enable state alert properties" option since there is no health model for devices generating the SNMP traps.

5. Select the "Alert severity" that you want. This example generates error alerts from all SNMP traps. Click Next to open the Alert Suppression page (see [Figure 10-15](#)).
6. Make some adjustments to the "Suppress duplicate alerts" setting of the SNMP rule. By default, an alert is considered to be a duplicate if:

- a. There is an existing alert with a resolution state other than Resolved in the Operator console.
- b. The alert comes from the same machine.
- c. The alert has been generated by the same rule.

When these three conditions are met, the new alert is suppressed and the repeat count of the existing alert is incremental. All SNMP trap alerts that occurred after the initial one would never be seen except as an increment in the alert count, even if the trap occurred on an entirely different device.

Figure 10-14. Alert configuration page

On the Alert Suppression tab, you need to make sure that "Suppress duplicate alerts" is enabled and only the Alert Description box is checked.

7. Click Next twice to proceed through the Responses page and the "Knowledge base" page. This brings up the General page where you give the rule a name; [Figure 10-16](#) shows it here as All SNMP Traps.
8. Click Finish to complete the rule creation process and commit the configuration change.



9. Create the association between the rule group that you just created and the computer group that was created earlier.
10. Commit the configuration change and get ready to watch the SNMP alerts roll in.

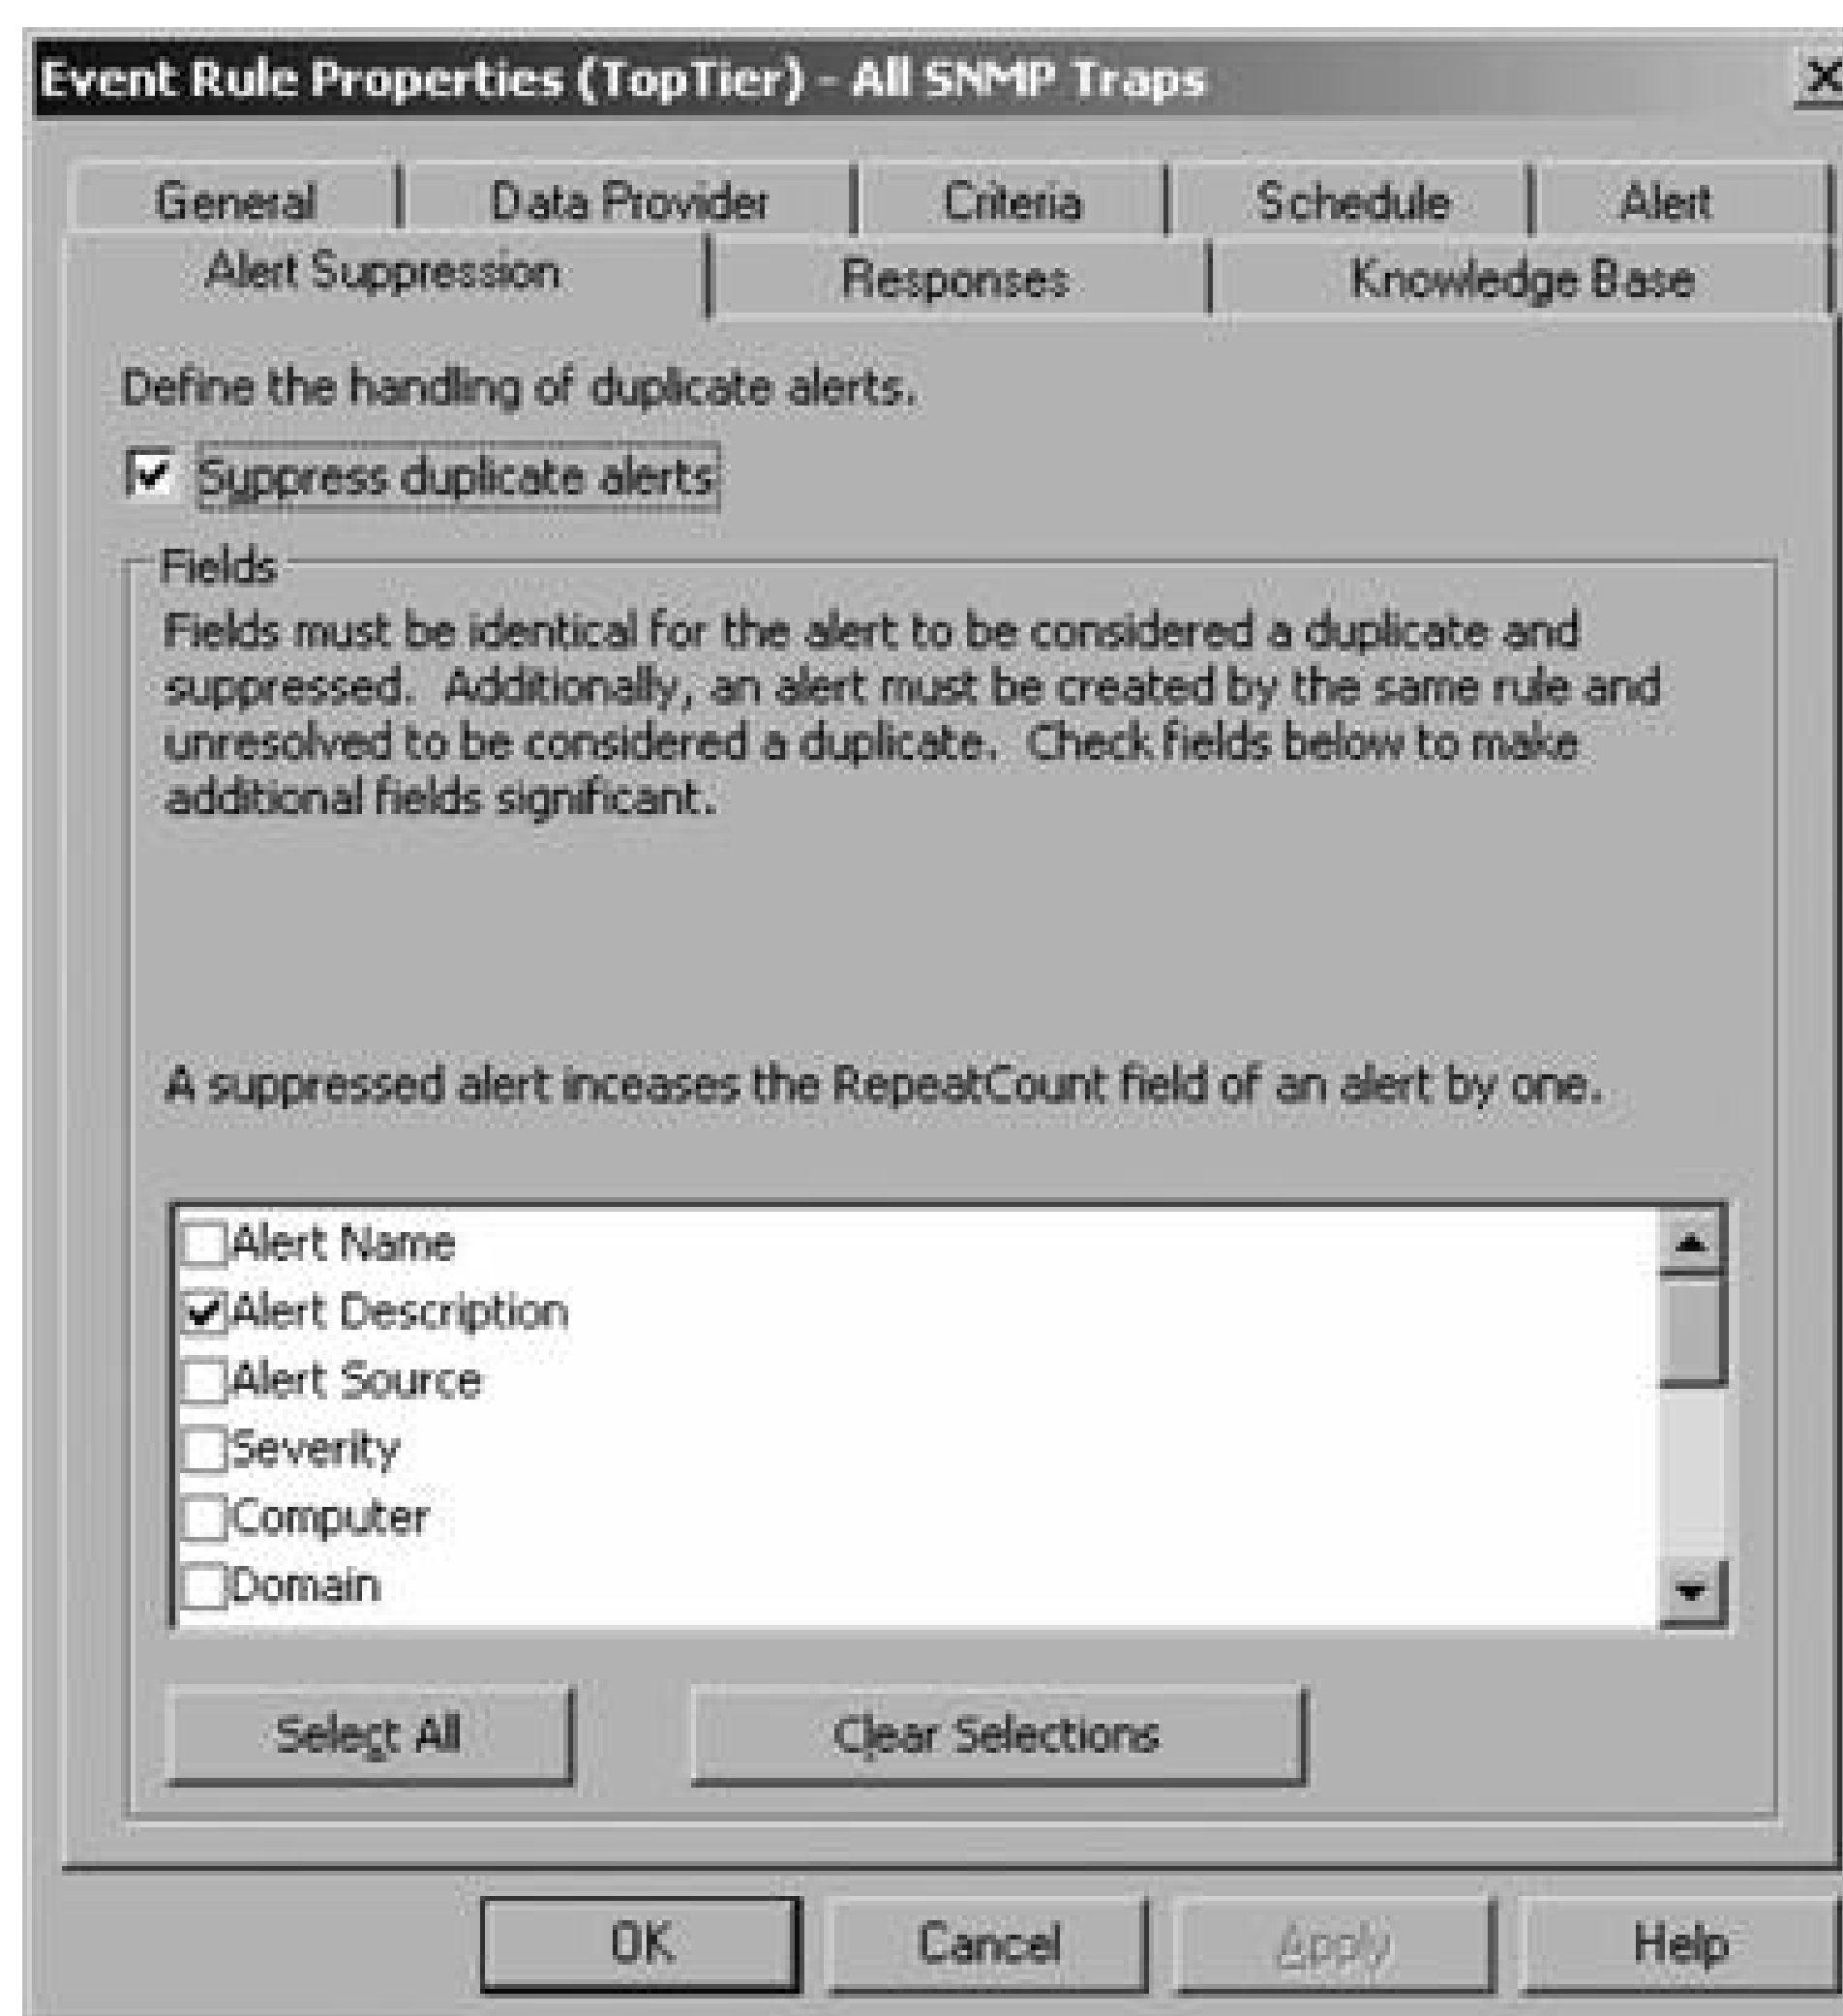
### 10.3.5. Making Use of Alerts from Traps

Now you are collecting all this wonderful information from your SNMP-enabled devices, but what can you get from it and what can you do with it? Traps merely indicate that an event happened. In this sense, they are much like an event in the Windows event log and only remotely like an alert.

Inferring the current state of an SNMP-enabled device from a single trap is possible. However, if you are looking for a larger health-state-oriented understanding of the state of a device or a group of devices, you are going to have to rely on your own knowledge and experience with your devices.

Traps are just bits of evidence that when considered together will let the experienced administrator draw some conclusions. In this case, the experienced administrator is much like the doctor from [Chapter 1](#) who considered her patient's symptoms (high fever, aching muscles, fatigue, and headache) as bits of evidence and concluded that the patient had the flu. [Figure 10-17](#) shows an alert that was generated by a trap, with some interesting landmarks pointed out. The alert contains a good deal of information that has been made more readable by the presence of an MIB on the management server (see [Figure 10-17](#)):

Figure 10-15. Selecting to suppress on Alert Description



- Point 1. `_CLASS=SNMPColdStartExtendedNotification` indicates that this device was powering up from a complete shutdown.
- Point 2. This alert was generated on a Cisco router with the IP address of 10.0.0.1. This alert was generated by unplugging the router and then plugging it back in.
- Point 3. The alert itself was generated by the server *homemomserver3*. All alerts generated from SNMP traps will come from *homemomserver3*, which is why using the machine that generated the alert as a distinguishing field for duplicate alert suppression will fail miserably.
- Point 4. This alert has occurred two more times. Since the duplicate alert suppression is based on the alert description (all the text on the lefthand side of the details pane), other cold start traps from other devices (which will have different IP addresses) will also be uniquely displayed in the Operator console.

Figure 10-16. Giving the event rule a name

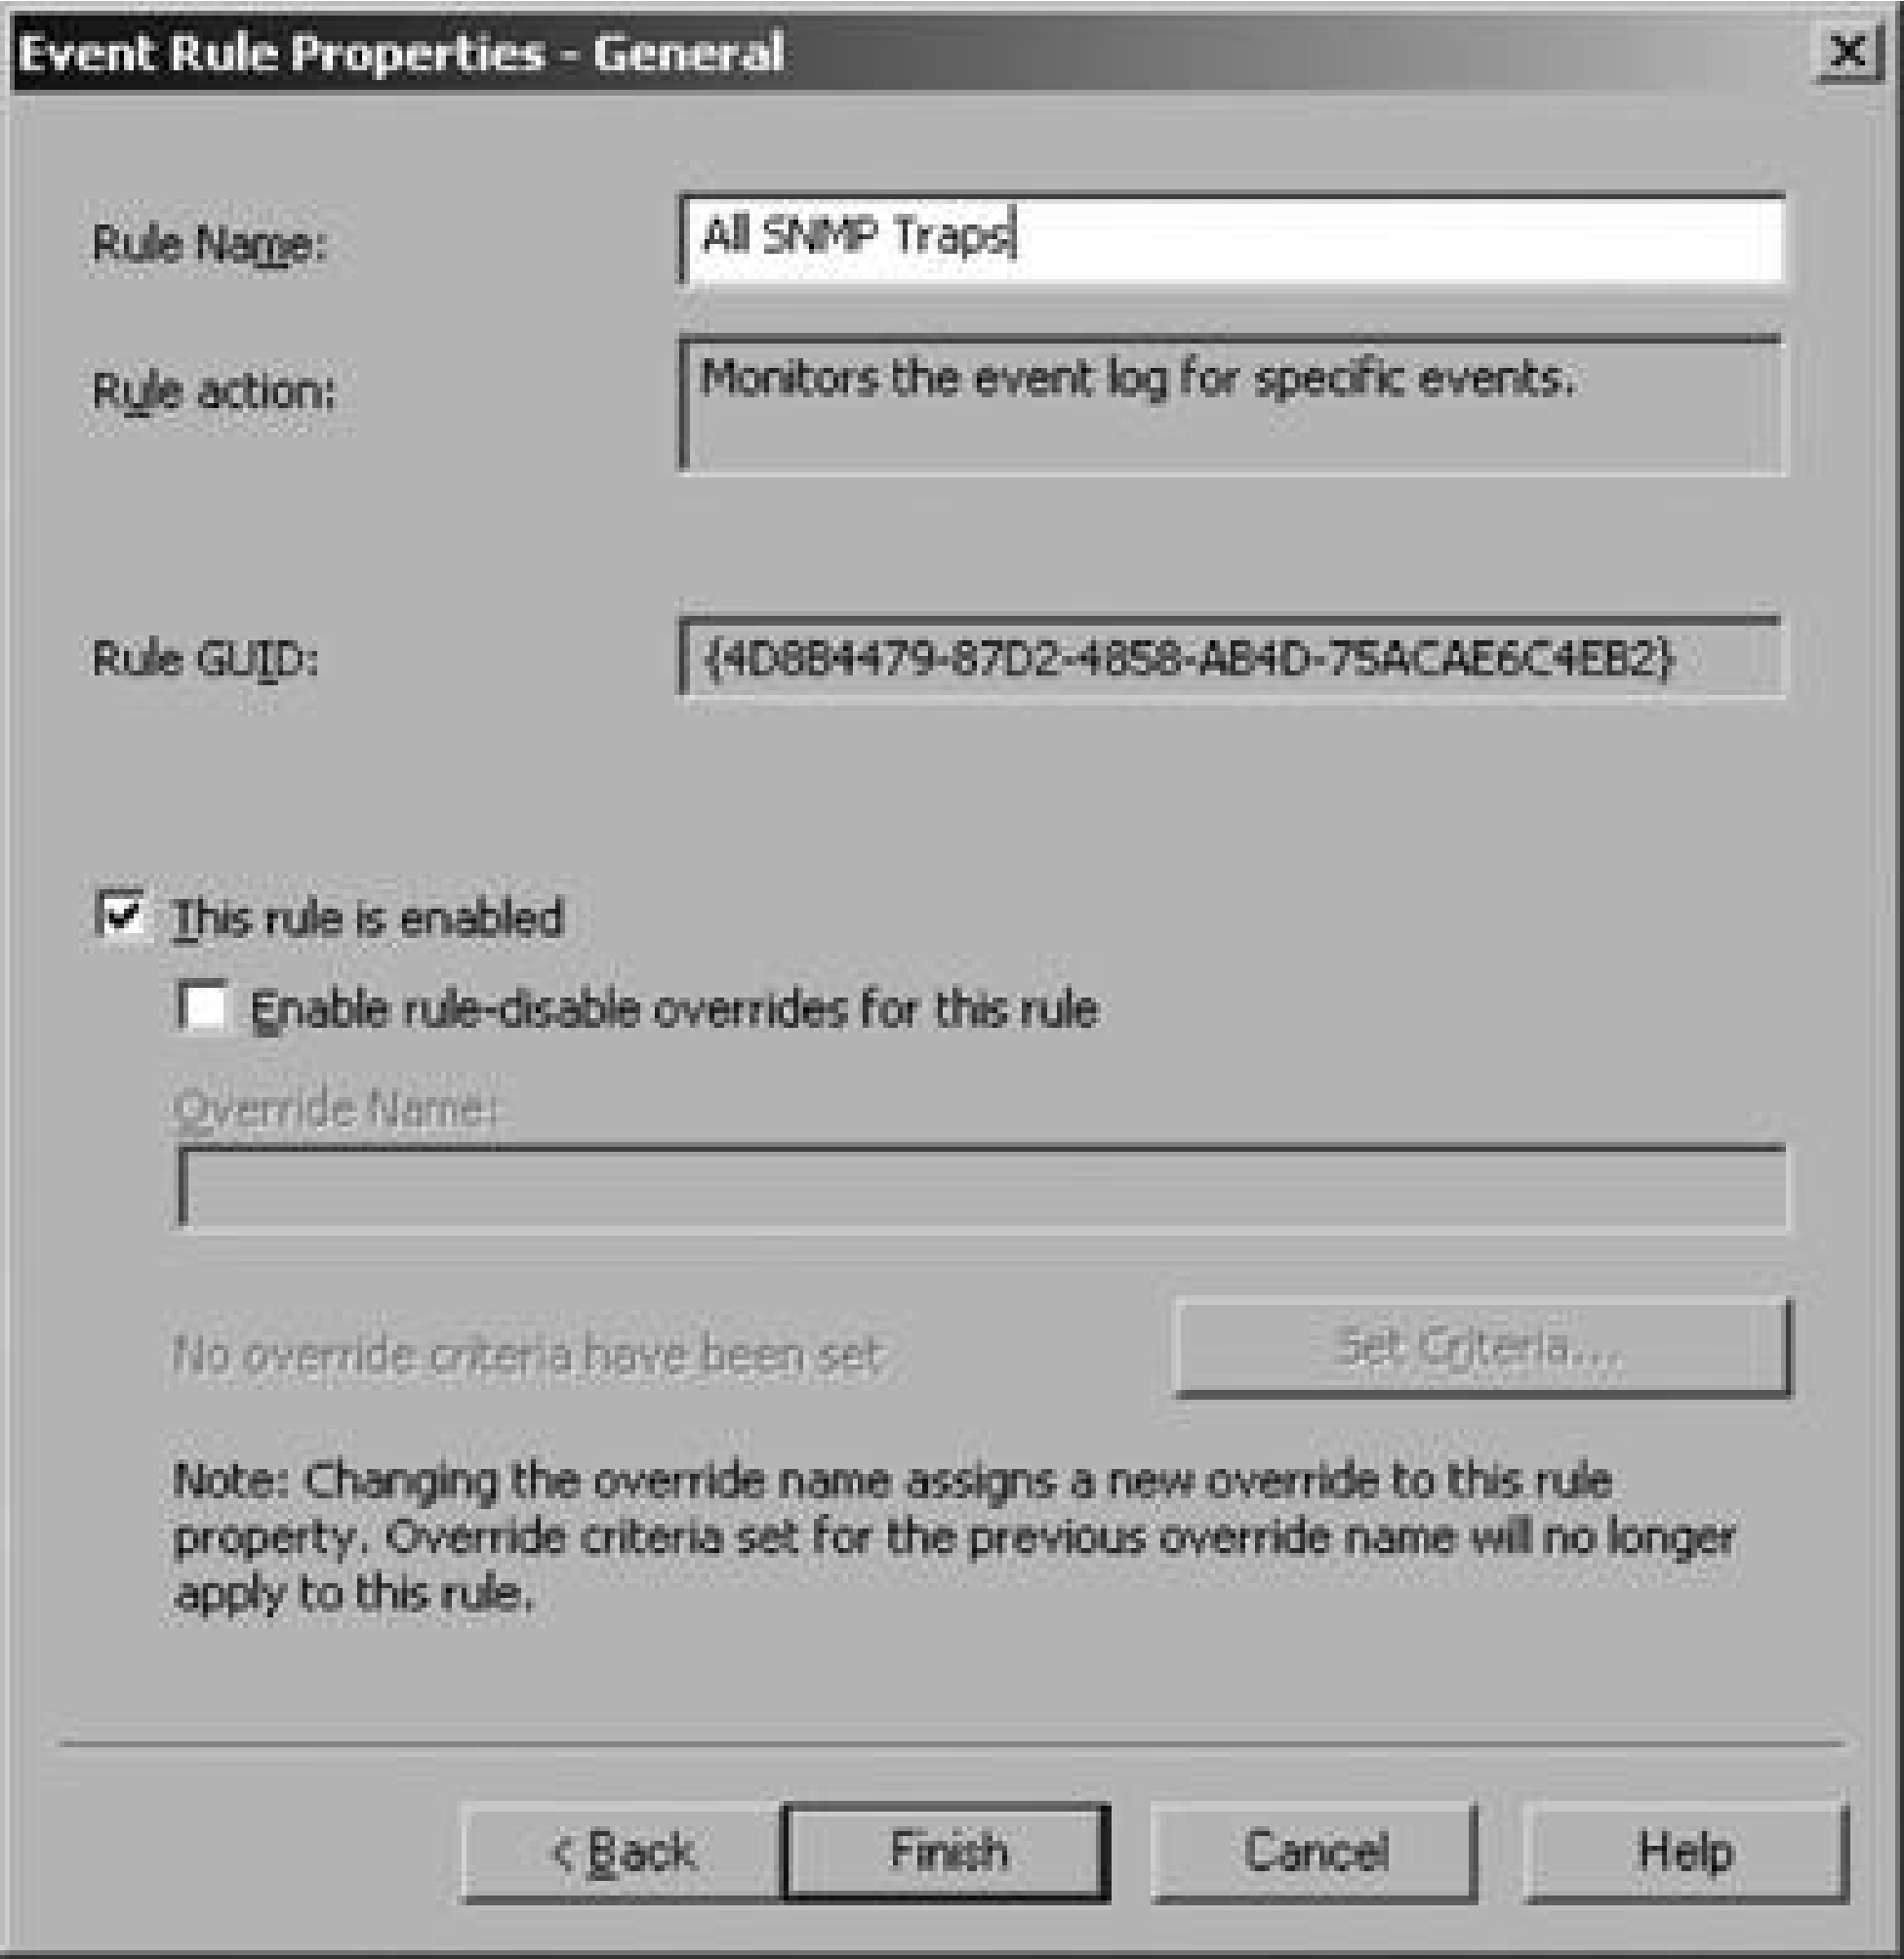


Figure 10-17. An alert generated from a trap





# 10.4. Syslog

All Unix variants record significant system and application events to text files called syslog files. Syslog is also the name of the Unix daemon (the equivalent of a Windows service) that performs the logging function. Syslogging isn't restricted to Unix servers: some network devices (from companies like Cisco) also implement the syslog feature.

## 10.4.1. Sending Syslog Events to MOM

Syslog supports redirection of the events to syslog files on other computers. This is how you get the syslog events into MOM. Controlling the behavior of syslog is done through the entries in the *syslog.conf* file, typically located in the */root/etc* directory on a Unix system. Below is the default *syslog.conf* file from a Linux variant.

```
1 # /etc/syslog.conf - Configuration file for syslogd(8)
2 #
3 # For info about the format of this file, see "man syslog.conf".
4 #
5
6 #
7 #
8 # print most on tty10 and on the xconsole pipe
9 #
10 kern.warning;*.err;authpriv.none    /dev/tty10
11 kern.warning;*.err;authpriv.none    | /dev/xconsole
12 *.emerg                             *
13
14 # enable this, if you want that root is informed
15 # immediately, e.g. of logins
16 #*.alert                             root
17
18
19 #
20 # all email-messages in one file
21 #
22 mail.*                               -/var/log/mail
23 mail.info                            -/var/log/mail.info
24 mail.warning                         -/var/log/mail.warn
25 mail.err                             /var/log/mail.err
26
27 #
28 # all news-messages
29 #
```

```

30 # these files are rotated and examined by "news.daily"
31 news.crit          -/var/log/news/news.crit
32 news.err           -/var/log/news/news.err
33 news.notice        -/var/log/news/news.notice
34 # enable this, if you want to keep all news messages
35 # in one file
36 #news.*             -/var/log/news.all
37
38 #
39 # Warnings in one file
40 #
41 *.=warning;*.=err   -/var/log/warn
42 *.crit              /var/log/warn
43
44 #
45 # save the rest in one file
46 #
47 *.*;mail.none;news.none -/var/log/messages
48
49 #
50 # enable this, if you want to keep all messages
51 # in one file
52 #*.*                -/var/log/allmessages
53
54 #
55 # Some foreign boot scripts require local7
56 #
57 local0,local1.*     -/var/log/localmessages
58 local2,local3.*     -/var/log/localmessages
59 local4,local5.*     -/var/log/localmessages
60 local6,local7.*     -/var/log/localmessages

```

If Unix systems are not where you spend the majority of your time, notice the text in line 3, [man syslog.conf](#). Entering this at a Unix command prompt opens the help file for *syslog.conf*. The help files are very detailed, so plan on spending some time going through it if you want to know more. The pertinent information has been extracted and it will give you what you need to perform the *syslog.conf* editing here.

Next, look at lines 21 and 22. Line 21 consists of only a `#` sign; this is the comment symbol for the syslog file that tells the syslog daemon to ignore the text on this line. Line 22, lacking the `#` sign will be interpreted and starts with `mail.*`. In syslog syntax, this means all events from the mail system. If you wanted error events, you would change this to `mail.err` as shown on line 25. Line 22 continues with a tab, which is ignored by the interpreter, and then `-/var/log/mail`. This is the default location that all mail events will be written to. Hopefully, the rest of the file now makes more sense.

To get all of the syslog events into MOM, replace all of the default locations (`/var/log/mail`) with the IP address of the MOM management server you will be collecting these at, preceded by a `@` sign. So, you would change `-/var/log/mail` to `@10.0.0.64`. After editing, the *syslog.conf* file looks like this:

```
# /etc/syslog.conf - Configuration file for syslogd(8)
#
# For info about the format of this file, see "man syslog.conf".
#

#
#
# print most on tty10 and on the xconsole pipe
#
kern.warning;*.err;authpriv.none    @10.0.0.64
kern.warning;*.err;authpriv.none    @10.0.0.64
*.emerg                             *

# enable this, if you want that root is informed
# immediately, e.g. of logins
#*.alert                            @10.0.0.64

#
# all email-messages in one file
#
mail.*                               @10.0.0.64
mail.info                           @10.0.0.64
mail.warning                        @10.0.0.64
mail.err                            @10.0.0.64

#
# all news-messages
#
# these files are rotated and examined by "news.daily"
news.crit                           -@10.0.0.64
news.err                            -@10.0.0.64
news.notice                         -@10.0.0.64
# enable this, if you want to keep all news messages
# in one file
news.*                              -@10.0.0.64

#
# Warnings in one file
#
*.=warning;*.=err                   -@10.0.0.64
*.crit                              @10.0.0.64

#
# save the rest in one file
#
*.*;mail.none;news.none             -@10.0.0.64

#
# enable this, if you want to keep all messages
# in one file
```



```
#*.*                                -@10.0.0.64

#
# Some foreign boot scripts require local7
#
local0,local1.*                    -@10.0.0.64
local2,local3.*                    -@10.0.0.64
local4,local5.*                    -@10.0.0.64
local6,local7.*                    -@10.0.0.64
```

Save the file and reboot the Unix server (which restarts the syslog daemon). All of the syslog events will flow into MOM.

## 10.4.2. Configuring MOM for Syslogs

Preparing MOM to receive syslog events is much simpler than preparing MOM to receive SNMP because there is nothing in the OS to configure; all configuration is done in MOM.

The rule that will generate alerts from syslog messages must have a data provider to monitor, so the first step is to create that provider. In the Administrator console open the Management Packs Providers container, bring up the context menu for the Providers container and select to create a new provider. This opens the Select Data Provider Type page, as shown in [Figure 10-18](#).

Figure 10-18. Creating a syslog port provider in MOM

Select the Application Log data provider type and click Next. Since syslogs are just plain text files, it make sense that MOM would treat them the same as any other application log, such as the IIS logs. This opens the Application Log Provider Properties page (see [Figure 10-19](#)).

Syslog runs over UDP port 541, and MOM has a Provider log type already created for this. Select the "Syslog port" Provider type, give the provider a name, and click Finish.

Now you can create a syslog rule group, associate it with a computer group and make a new event rule just as you did for the SNMP rule. In this case though, for duplicate alert suppression, select computer, domain, and Alert Description as the qualifying fields. Don't forget to commit your configuration change.

When MOM starts receiving syslog messages, you will see entries for these computers in the Administrator console, in the All Computers container (see [Figure 10-20](#)). [Figure 10-21](#) shows what a syslog message generated alert looks like in the Operator console. Because MOM is able to classify the domain as Unix Systems and give each alert a unique name its IP address the alerts generated by these messages are shown in the Operator console as coming from those machines. This happens despite the fact that all of the syslog-to-alert alerts are generated on the management server (point 1 in [Figure 10-21](#)).

Figure 10-19. Selecting the application log type

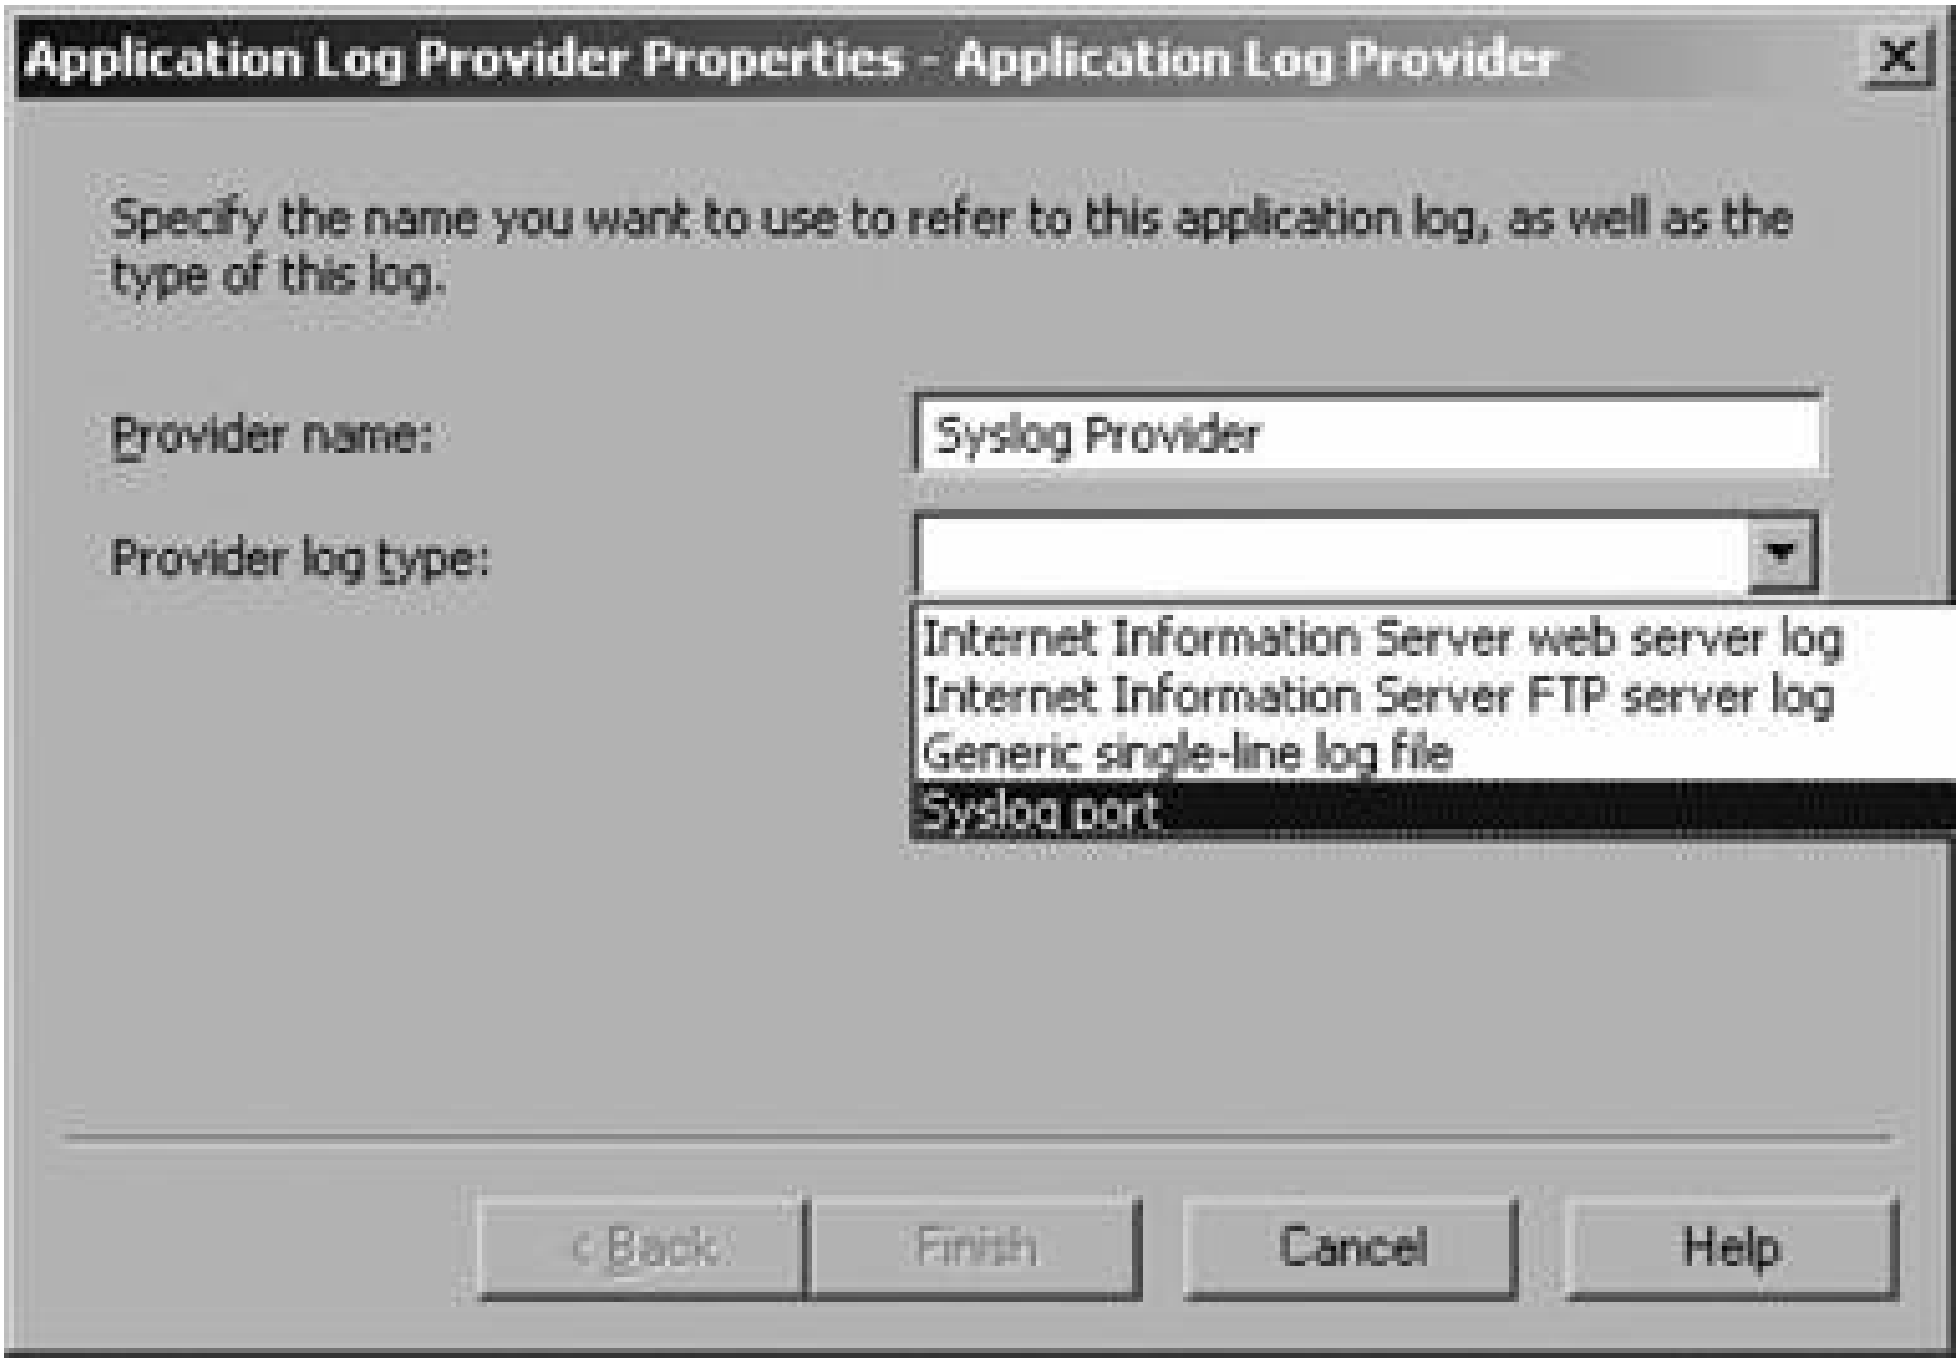


Figure 10-20. Unix systems in the Administrator console

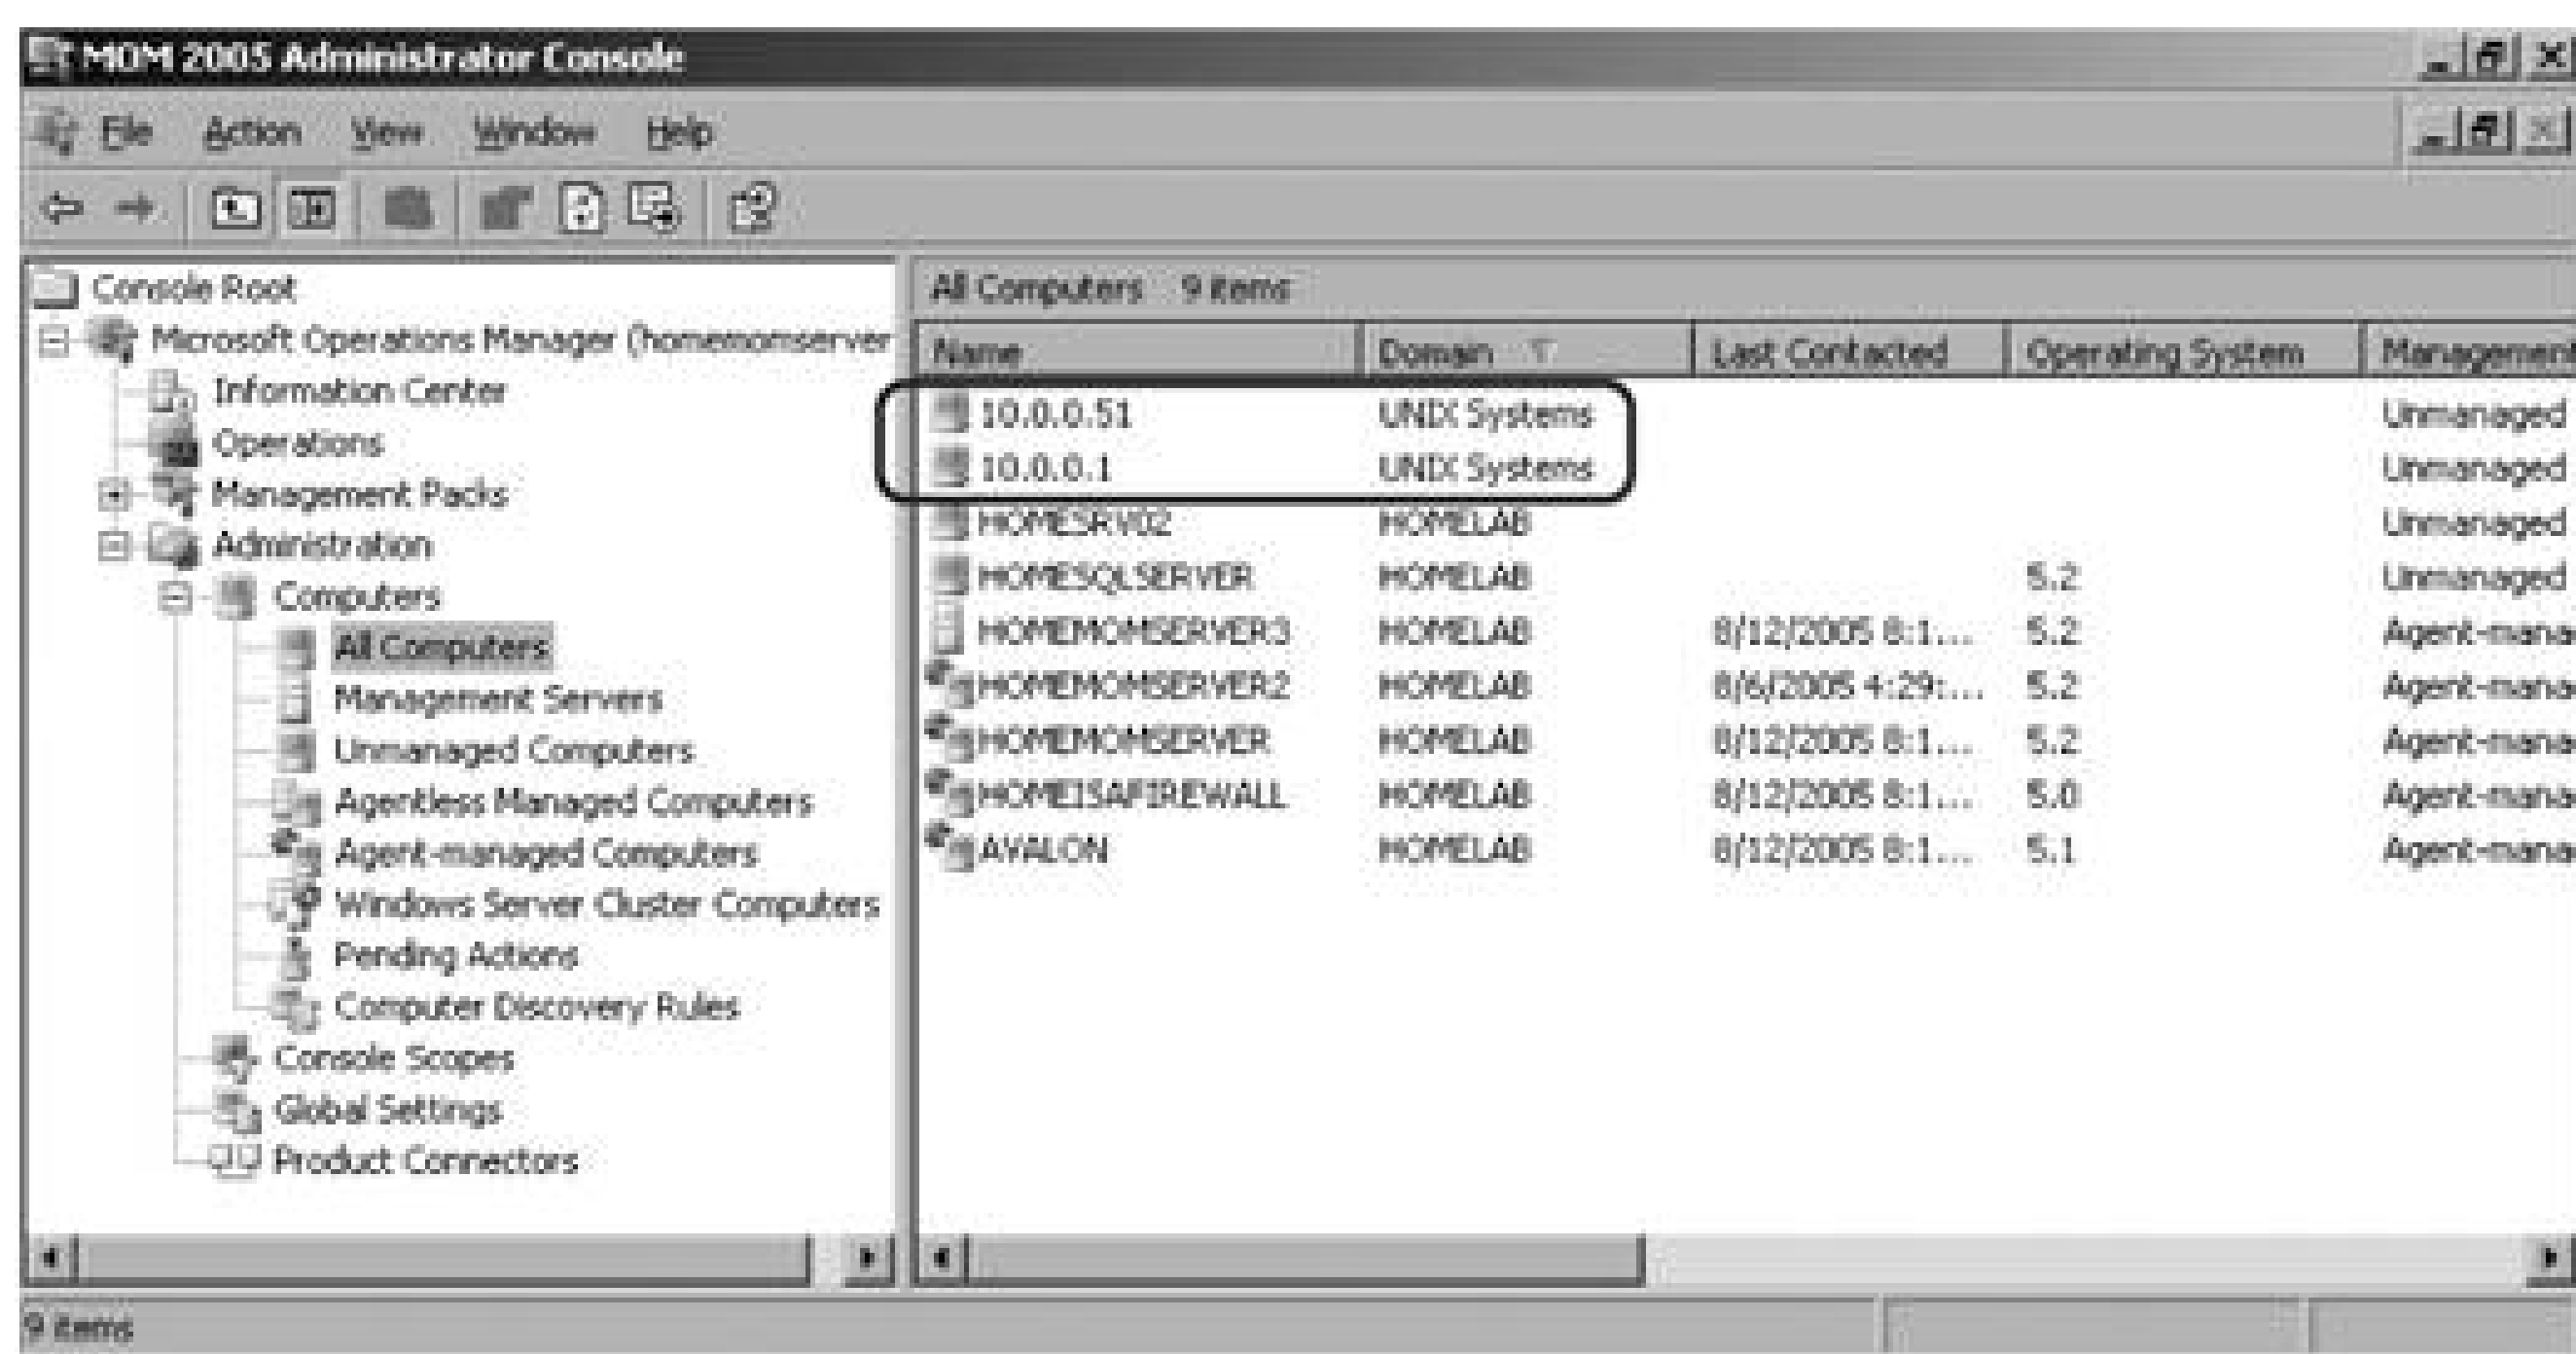


Figure 10-21. A syslog message generated alert

Point 2 in [Figure 10-21](#) shows that the syslog message has the severity of Error and is from a system daemon named windbindd. Point 3 shows a repeat count of 2 for this alert. By itself, this is not remarkable, but if you look in the results pane (under point 1 in [Figure 10-21](#)) you can see that there is a previous alert from the 10.0.0.51 system. This proves that duplicate alert suppression is working via the repeat count being incremented to 2 and that alerts that differ based on Alert Description are not being suppressed.





# 10.5. Summary

It is a rare IT environment that consists solely of Windows servers and workstations, which are in MOM's sweet spot for management. Routers and switches are critical components of all but the smallest of networks that support Windows servers, Unix hosts, or other platforms. Using out-of-the-box functionality, SNMP traps and syslog messages can be brought into MOM and presented in the Operator console. This chapter teaches the procedures you need to do this. Third-party products will provide a richer integration experience but have a commensurate cost. By using native capabilities you effectively extend your MOM operations management environment beyond the Windows space with no out-of-pocket cost and little invested time.





## About the Author

Chris Fox has been working with MOM since the release of MOM SP1 in 2002. He has architected and implemented MOM 2000 and 2005 installations for over 500 managed nodes. He currently works as a Portals Technology Specialist for Microsoft Corporation.





# Colophon

The animal on the cover of *Essential Microsoft Operations Manager* is a beaver. The beaver (*Castor canadensis*) is found all over North America, except in northernmost Canada and the warmer southern parts of the United States. The largest rodent in North America, it is characterized by dark brown fur, long incisors, and small ears and eyes. It can measure four feet in length and weigh between 40 to 100 pounds in adulthood. Its most noticeable feature is its long, flat tail, which it uses both as a rudder for swimming and as a balancing aid when standing on its hind legs.

Beavers are industrious builders and spend a great deal of time constructing dams and lodges. Their dams are dome-shaped and measure as high as 10 feet. Their purpose is to raise the surrounding water level two to three feet so the beavers can build a lodge. Lodges typically have two underwater entrances, and the water must be deep enough so that the entrances will not be blocked by ice.

Beaver pelts were once highly prized by North American settlers, and it was the beaver trade that drove the exploration of the continent. Pelts were traded as currency and were considered fashionable for top hats and as the trim for royal robes. The beaver pelt trade nearly drove the animals to extinction in the 1800s, and it is estimated that the North American beaver population today is only 5 percent of the size of the original population.

Humans see beavers as both a blessing and a curse. The dams that beavers create can help purify water by breaking down pesticides. However, dams can also flood roads and farmlands. Beavers are tenacious and difficult to outsmart, but environmentalists are learning ways to prevent beaver damage by creating structures that utilize the beaver's natural building tendencies rather than deterring them.

The cover image is from *Wood's Illustrated Natural History*. The cover font is Adobe ITC Garamond. The text font is Linotype Birka; the heading font is Adobe Myriad Condensed; and the code font is LucasFont's TheSans Mono Condensed.



[← PREV](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

[← PREV](#)

 [PREV](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

[.akm file extension](#)

[.bak files](#)

[.sql extension](#)

[12-step build process 2nd](#)

 [PREV](#)



# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

- [AckData method](#)
- [Administrator console](#)
  - [Global Settings node](#)
- [agent action accounts](#)
- [agents 2nd 3rd](#)
  - [agent failover configuration](#)
  - [Agent Helper](#)
  - [agent management](#)
  - [caching](#)
  - [data volumes generated per day](#)
  - [default values](#)
  - [handling low bandwidth](#)
  - [rejecting manually installed agents](#)
  - [tools](#)
  - [uninstalling](#)
- [Alert Logging Latency reports](#)
- [Alert tuning solution accelerator](#)
- [alert-oriented management packs](#)
- [alerts](#)
  - [agents, generation by](#)
  - [Alert Details pane](#)
    - [Events tab](#)
    - [Product Knowledge tab](#)
  - [alert grooming \(database maintenance\)](#)
  - [alert notification](#)
  - [alert rules](#)
  - [alert severity settings](#)
  - [alert/response rules](#)
  - [alerts view group](#)
  - [custom alert fields](#)
  - [global settings](#)
  - [Properties tab](#)
  - [resolution states](#)
    - [fields](#)
  - [Service Level Exceptions Alert view](#)
  - [size in bytes](#)
  - [SNMP traps, conversion to](#)
  - [SNMP traps, generation](#)
  - [Windows event log events, compared to](#)
- [atomic transactions](#)
- [authentication configuration](#)







# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

backups

- [backup jobs, creating](#)
- [backup tools](#)
- [creating and scheduling](#)
- [critical data](#)
- [databases](#)
- [management packs](#)
- [ManualMC.txt](#)
- [OnePoint database](#)
- [OS-level backups, creating](#)
- [report definitions](#)
- [restoring](#)
  - [OnePoint database](#)
  - [system databases](#)
- [scheduling](#)
- [Server OS configuration](#)
- [SQL backup routine](#)

# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

- [collection event rules](#)
- [community names](#)
- [compare performance data rules](#)
- [Component Object Model \(COM+\) applications](#)
- computer discovery
  - [trusted LANs, operation over](#)
- [computer discovery rules 2nd](#)
  - [agent installation on trusted LANs](#)
  - [Computer Discovery Rules node](#)
  - [trusted LANs](#)
    - [installing agents](#)
- computer groups
  - [computer security groups \(Active Directory\) and MOM groups](#)
  - [console scopes and](#)
  - [custom groups, creating](#)
  - [membership and naming](#)
- configuration
  - [pre-installation choices](#)
    - [agent action account](#)
    - [agent installation and management](#)
    - [DAS account](#)
    - [groups and roles](#)
    - [management server action account](#)
    - [operations database](#)
    - [reporting databases](#)
    - [Reporting Server database](#)
    - [security](#)
    - [service accounts](#)
  - [Connector.GetData method](#)
  - [Connector.InsertAlert method](#)
  - [console scopes](#)
    - [creating](#)
    - [custom scopes, creating](#)
  - [consolidate similar events rules](#)



# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

[DAS \(Data Access Server\)](#)

[DAS accounts](#)

[data providers](#)

[data sources for reports](#)

[data-driven subscriptions 2nd](#)

[databases](#)

[data warehouse database](#)

[database grooming](#)

[OnePoint database grooming](#)

[databases folder](#)

[MOM 2005 reporting databases](#)

[ReportServer database](#)

[restoring](#)

[rows and columns](#)

[saving queries](#)

[SQL database concepts](#)

[SQL Query Analyzer](#)

[SQL Server system databases](#)

[SystemCenterReporting database](#)

[tables](#)

[transaction logs](#)

[transactional databases](#)

[dbo.Alert](#)

[deploying agents](#)

[firewalls and WANs](#)

[multiple management groups](#)

[preparation](#)

[computer discovery rules](#)

[trusted LANs](#)

[agent management](#)

[computer discovery rules](#)

[confirming functionality](#)

[Install/Uninstall Agents Wizard](#)

[managed computers, preparation](#)

[management group preparation](#)

[ManualMC.txt, discovery and installation using](#)

[untrusted environments](#)

[destination management groups](#)

[detect missing event rules](#)

[diagram view group](#)

[domain accounts](#)



[DTS \(Data Transformation Service\)](#)

[DTS transfers](#)

[reporting database size and  
storage requirements](#)

[dwdb](#)

[dwserver](#)





# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

- [EeaData.mdf](#)
- [EeaLog.ldf](#)
- [Email Server global settings](#)
- [Event Logging Latency reports](#)
- [EventCreator.exe](#)
- events
  - [event rules 2nd](#)
  - [events view group](#)
  - [size in bytes](#)
- [eXc](#)
- [Exchange 2003 management pack for MOM 2005](#)

[← PREV](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

- [Failover tab](#)
- [File transfer folder](#)
- [filter rules](#)

[← PREV](#)



# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

- [GetData method](#)
- [global settings](#)
  - [alert resolution states](#)
    - [fields](#)
  - [alerts](#)
  - [connections](#)
    - [communications](#)
    - [web addresses](#)
  - [custom alert fields](#)
- [Email Server](#)
- [maintenance](#)
  - [database grooming](#)
  - [operational data reports](#)
- [grooming \(database maintenance\)](#)
- [groups view group](#)



[← PREV](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

[heartbeat checking](#)

[← PREV](#)



# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

- [implementation](#)
  - [pre-implementation testing](#)
- [Initialize method](#)
- [Install/Uninstall Agents Wizard](#)
- [installation](#)
  - [all-in-one versus custom option](#)
  - [installation order, MOM and its dependencies](#)
  - [operations database](#)
  - [setup program pages, management server and operations database](#)
  - [Setup Resources page](#)
- [item-level roles](#)
  - [Browser role](#)
  - [My Reports](#)

[← PREY](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

[Kerberos](#)

[← PREY](#)

◀ PREV

# Index

[SYMBOL] [A] [B] [C] [D] [E] [F] [G] [H] [I] [K] [L] [M] [N] [O] [P] [Q] [R] [S] [T] [U] [V] [W]

[life cycles of management packs](#)

[linked reports](#)

[local groups](#)

◀ PREV





# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

- managed computers
  - [maintenance mode](#)
  - [pinging](#)
  - [trusted LANs, preparation for use on](#)
- [management groups](#)
  - [account permissions, granting of](#)
  - [agents, deploying from multiple groups](#)
  - [data processing capacity](#)
- group interconnections
  - [administrative partitioning](#)
  - [configuration](#)
  - [distributed groups, monitoring](#)
  - [functionality-based partitioning](#)
  - [MMPC and MCF dependency](#)
  - [partitioning of tasks](#)
  - [source and destination groups](#)
- [managed computers](#)
- [management server](#)
- [naming](#)
- [operations database](#)
- [topology planning](#)
  - [requirements](#)
- [trusted LANs, preparation for](#)
- [Management Module Utility](#)
- [management packs, xvi 2nd](#)
  - [advanced pack configuration](#)
  - [alert rules](#)
  - [alert-oriented versus state-oriented](#)
  - [backups](#)
  - [creating](#)
    - [importing into preproduction](#)
  - [File transfer folder](#)
  - [GUIDs versus names](#)
  - [importing into preproduction](#)
  - [individual rule application](#)
  - [life cycles](#)
  - [management pack guides](#)
  - [Management Pack Wizard](#)
  - [management pack workflow](#)
  - [merging](#)
  - [Microsoft server products and](#)

[overriding rule application](#)

[processing](#)

[Management Packs node](#)

[production](#)

[evolution in production](#)

[exporting from preproduction](#)

[importing into production](#)

[resource kit tools](#)

[rule groupings](#)

[tuning](#)

[health state roll-up](#)

[tools](#)

[vendor revisions, integrating](#)

[version control](#)

**management servers**

[action accounts](#)

[installation](#)

[management server agent](#)

[MOMHost.exe](#)

[topology planning](#)

[ManualMC.txt](#)

[backups](#)

[master database](#)

[MCF \(MOM Connector Framework\)](#)

[installation](#)

[management group servers, installation on](#)

[methods](#)

[MOM discovery data and](#)

[monitoring](#)

[product connectors and](#)

[media sets](#)

[Microsoft Management Console \(MMC\) snap-ins](#)

[Microsoft Operations Manager MPNotifier management pack](#)

[MicrosoftOperationsManager.mib](#)

[MMPC \(MOM-to-MOM Product Connector\) 2nd](#)

[alert retrieval via MFC methods](#)

[configuration and use](#)

[distributed management groups, monitoring](#)

[management groups, connecting](#)

[monitoring](#)

[primary methods](#)

[reporting](#)

[tiered configuration, creating](#)

[MOM \(Microsoft Operations Manager\) 2005, xv](#)

[component setup](#)

[MOM 2005 server](#)

[MOM 2005 service](#)

[MOM administrators group](#)

[MOM agent](#)

[MOMHost.exe](#)

[MOM Agent heartbeat failure alert](#)

[MOM Authors](#)

[MOM authors group](#)

[MOM users group](#)

[SNMP trap generation](#)

[structure](#)

[testing](#)

[versions](#)

[MOM 2005 Reporting](#)

[administration](#)

[item-level roles](#)

[permissions, granting](#)

[queries](#)

[reports components](#)

[SystemCenterReporting](#)

[MOM.Datawarehousing.DTSPackage Generator.exe](#)

[momagent.msi](#)

[momconn.exe](#)

[MOMHost.exe 2nd 3rd 4th](#)

[MOMService group](#)

[MOMService.exe 2nd](#)

[MOMX partitioning and grooming](#)

[monitoring](#)

[distributed management groups](#)

[MCF \(MOM Connector Framework\)](#)

[MMPC \(MOM-to-MOM Product Connector\)](#)

[rules, OnePoint database](#)

[MP2XML tool](#)

[MPDiff tool](#)

[MPDiff.Console.exe](#)

[mpwizard.exe](#)

[msdb database](#)

[My Reports 2nd](#)

[enabling](#)

[linked reports, creating](#)

[My Views group](#)

 [PREV](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

[network service account](#)  
[Notification Command Format](#)

 [PREV](#)





# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

- [object identifiers \(OIDs\)](#)
- [on-demand execution of reports](#)
- [OnePoint](#)
- [OnePoint database](#)
  - [backups](#)
  - [computer maintenance check](#)
  - [database grooming, tuning](#)
  - [DTS and](#)
  - [MOMX partitioning and grooming](#)
  - [monitoring rules](#)
  - [OnePointCheckIntegrity](#)
  - [OnePointReindex](#)
  - [restoring](#)
  - [SQL Enterprise Manager and](#)
  - [SQL jobs](#)
  - [SQL maintenance jobs](#)
  - [updates](#)
- [operating capacity](#)
- [operational data reports](#)
- [operations databases 2nd](#)
  - [configuration choices](#)
  - [data volumes added daily](#)
  - [installation](#)
  - [OnePoint](#)
  - [optimum size](#)
  - [topology planning](#)
- [operations management, xv](#)
  - [versus system administration](#)
- [Operator console 2nd 3rd](#)
  - [customizing](#)
  - [filters](#)
    - [building](#)
  - [OnePoint database and](#)
  - [Outlook 2003 interface, similarity to](#)
  - [panes, adding](#)
  - [performance chart](#)
  - [permissions](#)
  - [Repeat Count field](#)
  - [summary](#)
  - [Tasks pane](#)
  - [using](#)

[overrides](#)

[health state roll-up](#)

[moving from preproduction to production](#)

[performance threshold override](#)

[Rule Disable Override](#)

[script parameter overrides](#)

[state alert severity override](#)

[tools](#)





# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

- [partitioning](#)
- performance monitor data
  - [sample size in bytes](#)
- [performance rules 2nd](#)
- [performance view group](#)
- [pinging](#)
- [prerequisite checker](#)
- [product connectors](#)
  - [supported products](#)
  - [third-party solutions](#)
- [Product Knowledge tab](#)
  - [Causes section](#)
  - [Resolutions section 2nd](#)
- [product registration](#)
- [public view group](#)

[← PREY](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

[queries, saving](#)

[← PREY](#)





# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

- [RAW MP objects](#)
- Remote Web console
  - [secure configuration](#)
- [report definitions, backups](#)
- reporting
  - [multiple source management groups](#)
- [Reporting console](#)
- Reporting Server
  - [Reporting Server database](#)
    - [configuration choices](#)
    - [optimum size](#)
  - [topology planning](#)
- reports
  - [on-demand execution](#)
  - [shared schedules](#)
  - [subscription management](#)
  - [SystemCenterReporting](#)
- [reports components](#)
- [Reports.xml file extension](#)
- [ReportServer database 2nd](#)
  - [Report Utility](#)
- [resources](#)
- [rptutil.exe](#)
- [rules](#)
  - [alert notification](#)
  - [alert rules](#)
  - [alert severity settings](#)
  - [alert/response rules](#)
  - [collection event rules](#)
  - [compare performance data rules](#)
  - [consolidate similar event rules](#)
  - [customizing](#)
  - [individual rule application](#)
  - [overriding rule application](#)
  - [rule group application](#)
  - [rule overrides](#)
    - [health state roll-up](#)
    - [performance threshold override](#)
    - [script parameter override](#)
    - [state alert severity override](#)
  - [tools](#)

[rule overrides Rule Disable Override](#)  
[SNMP rules](#)





# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

- [sample performance data rules](#)
- [SCDW \(System Center Data Warehouse\)](#)
- [script parameter overrides](#)
- security
  - [configuration choices](#)
- [Server OS configuration, backups](#)
- service accounts
  - [configuration](#)
- [service discovery rules](#)
- [Setup method](#)
- [shared schedules for reporting](#)
- [Skywire](#)
- [SMI2SMIR.exe](#)
- [SNMP \(Simple Network Management Protocol\)](#)
  - [hierarchical data structure](#)
  - [management information block \(MIB\) files](#)
  - [MOM and Windows](#)
  - [MOM and Windows configuration](#)
    - [alerts from traps, using](#)
    - [community names](#)
    - [MIB files, compiling](#)
    - [SNMP service configuration](#)
    - [SNMP service installation](#)
    - [SNMP, testing](#)
    - [SNMPUTIL, confirming traps with](#)
    - [traps, converting to alerts](#)
    - [traps, generation by MOM](#)
- [SNMPUTIL](#)
- [solution accelerators](#)
- [source management groups](#)
- [SQL 2000 Reporting Services](#)
  - [installation](#)
  - [patches and MOM compatibility](#)
  - [reports components](#)
  - [versions](#)
- [SQL Enterprise Manager](#)
  - [databases folder](#)
  - [taskpad view](#)
- [SQL jobs](#)
  - [SQL backup jobs](#)
  - [SQL backup routine](#)

[SQL Query Analyzer](#)

[MOM database administration](#)

SQL Server

[SQL Server DTS package](#)

[SQL Server Reporting Services 2nd](#)

[system databases](#)

[transaction management](#)

[srcdb](#)

[srcserver](#)

[standard subscriptions](#)

[state alert severity override](#)

[state view group](#)

[state-oriented management packs](#)

[Syslog](#)

[syslog files](#)

system databases

[restoring](#)

[SystemCenterReporting](#)

[SystemCenterReporting database](#)

[Systems Center Data Warehouse \(SCDW\) readers group](#)





# Index

[[SYMBOL](#)] [[A](#)] [[B](#)] [[C](#)] [[D](#)] [[E](#)] [[F](#)] [[G](#)] [[H](#)] [[I](#)] [[K](#)] [[L](#)] [[M](#)] [[N](#)] [[O](#)] [[P](#)] [[Q](#)] [[R](#)] [[S](#)] [[T](#)] [[U](#)] [[V](#)] [[W](#)]

- [tables](#)
- [Tasks pane](#)
- [testing MOM 2005 prior to implementation](#)
- [tiered MOM infrastructure](#)
- [topology planning](#)
  - [initial design](#)
  - [management groups](#)
  - [management servers](#)
  - [operations databases](#)
    - [RAID versus clustered configuration](#)
  - [Reporting Server](#)
  - [requirements, determining](#)
  - [soliciting user input](#)
- [Transact SQL \(TSQL\)](#)
- [transaction logs](#)
- [transactional databases](#)
- [TSQL \(Transact SQL\)](#)

[← PREV](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

[Unix and MOM connectivity](#)

[UpdateAlerts method](#)

[← PREV](#)



# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

[views](#)

- [alerts view group](#)
- [diagram view group](#)
- [events, performance, and public views groups](#)
- [groups view group](#)
- [My Views group](#)
- [preconfigured performance view](#)
- [state view group](#)
- [view queries](#)

[Vintela](#)

[Visual Studio .NET 2003](#)

 [PREV](#)

# Index

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#)

- [warehouse database](#)
- [WBEMTEST](#)
- [Web console](#)
- [web-based interfaces](#)
- [Windows event log events, compared to MOM alerts](#)
- Windows Server 2003
  - [network service account](#)
- [WMI Tester tool](#)

 [PREV](#)