



# Microsoft® System Center

## Operations Manager Field Experience

Danny Hermans, Uwe Stürtz, Mihai Sarbulescu  
Mitch Tulloch, Series Editor

PUBLISHED BY  
Microsoft Press  
A division of Microsoft Corporation  
One Microsoft Way  
Redmond, Washington 98052-6399

Copyright © 2015 by Microsoft Corporation All rights reserved.

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

Library of Congress Control Number  
ISBN: 978-0-7356-9582-5

Printed and bound in the United States of America.

First Printing

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Support at [mspinput@microsoft.com](mailto:mspinput@microsoft.com). Please tell us what you think of this book at <http://aka.ms/tellpress>.

This book is provided “as-is” and expresses the author’s views and opinions. The views, opinions and information expressed in this book, including URL and other Internet website references, may change without notice.

Some examples depicted herein are provided for illustration only and are fictitious. No real association or connection is intended or should be inferred.

Microsoft and the trademarks listed at <http://www.microsoft.com> on the “Trademarks” webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

**Acquisitions Editor:** Karen Szall  
**Developmental Editor:** Karen Szall  
**Editorial Production:** Megan Smith-Creed  
**Copyeditor:** Megan Smith-Creed  
**Cover:** Twist Creative • Seattle

# Contents

Introduction	vii
<b>Chapter 1 The role of Operations Manager</b>	<b>1</b>
Examining a typical Operations Manager installation	1
Understanding Operations Manager architecture and workflow	2
Extending Operations Manager with cloud services	3
Best practices for operating system configuration	4
Power management	4
Disk performance	5
Disk optimization for IaaS virtual machines	5
Management server placement	6
Antivirus software exclusions	7
Best practices for SQL Server configuration	7
Auto grow and auto shrink	7
Instant file initialization	8
Maximum degree of parallelism	8
Maximum memory	9
Splitting up TempDB files	9
Recovery model	10
Service Principal Names	10
Re-indexing	11
Block size	12
Collation	12
Operations Manager registry optimizations	13
Data Access Layer	14
Persistence Manager	14
Health Manager	15
Pool Manager	15
Group Calculation module	16
SNMP module	16
Management Configuration Service	17

---

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

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

<http://aka.ms/tellpress>

	Data Warehouse module	17
	Operations Manager environment sizing	18
	General recommendations	18
	Resource pool considerations	19
	ITIL and MOF	20
	ITIL	20
	MOF	20
	The role of Operations Manager	21
	Step 1: Changing	21
	Step 2: Operating	22
	Step 3: Supporting	23
	Step 4: Optimizing	24
<b>Chapter 2</b>	<b>Best practices for working with management packs</b>	<b>25</b>
	Understanding management packs	25
	Singleton vs. non-singleton classes	26
	Workflow targets	26
	Authoring classes	27
	State change events	28
	Module cookdown	29
	Type space	29
	Authoring groups	30
	Group calculation interval	30
	Sealed management packs	31
	Summary of best practices	31
	Tools for analyzing management packs	32
	Management Pack Viewer	32
	Override Explorer	39
	Workflow Analyzer explained	40
	Workflow Analyzer internals	41
	Using Workflow Analyzer	41
	Running workflow tracing manually	41
	Best practices for creating and using groups	43
<b>Chapter 3</b>	<b>Getting the most out of Operations Manager visualizations</b>	<b>45</b>
	Tuning using the out-of-the-box reports	45
	Alert reports	47
	Event reports	49
	Data Volume reports	51

Publishing visualizations	55
Publishing and scheduling reports	56
Publishing to SharePoint	57
Using Visio with Operations Manager	60
<b>Chapter 4 Troubleshooting your Operations Manager environment</b>	<b>63</b>
Understanding the HealthService process	63
MOMConnector	64
HealthServiceExecutive	64
PersistenceManager	65
ConnectorManager	65
ConfigurationManager	65
ExecutionManager	66
HealthManager	67
SecureStorageManager	67
DiscoveryManager	67
DataPublisherManager	67
JobManager	67
PoolManager	67
Understanding the Operational Database	70
ManagedType table	71
ManagedTypeProperty table	72
RelationshipType table	73
Workflow and basic data tables	73
Instance space and discovery tables	74
Overview of the Data Warehouse Database	76
Instance space	76
Datasets and their tables	76
General troubleshooting steps	80
Operations Manager event log	81
ETW tracing for Operations Manager	82
SQL Server Profiler Tracing	85
Performance troubleshooting	85
Data grooming settings	86
Event and performance data collection	86
State change event frequency (noisy monitors)	87

Top issues and best practices	89
Recommended operating system fixes for Operations Manager	89
Gray management servers or agents	89
Data missing from reports	89
Resource pool unavailable	89
<b>Chapter 5 Using Operations Manager in cloud environments</b>	<b>91</b>
Using Global Service Monitor	91
Understanding DevOps	91
Understanding global service monitoring	92
Installing Global Service Monitor	92
Getting started with Global Service Monitor	93
Troubleshooting Global Service Monitor	94
Using Azure Operational Insights (System Center Advisor)	95
Understanding Azure Operational Insights	95
Connecting to Azure Operational Insights	97
Planned additional capabilities of Azure Operational Insights	100
Comparing Azure Operational Insights with Operations Manager	100
Installing Operations Manager in Azure IaaS	102
Supported scenarios	102
Configuring the network	103
Configuring storage	104
The first virtual machine	107
Configuring the first virtual machine as a domain controller	111
Configuring the next virtual machine as the SQL Server server	112
Configuring the final virtual machine as the Operations Manager management server	113
Monitoring Azure IaaS	114
Connecting your on-premises Operations Manager environment to your Azure subscription	114
Using Operations Manager to monitor Azure virtual machines	116

---

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

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

<http://aka.ms/tellpress>

# Introduction

If you're responsible for designing, configuring, implementing, or managing a Microsoft System Center Operations Manager environment, then this book is for you. This book will help you understand what you can do to enhance your Operations Manager environment, and will give you the opportunity to better understand the inner workings of the product, even if you are a seasoned Operations Manager administrator.

This book assumes that you have a deep working knowledge of the Operations Manager product and its concepts, that you understand the concept of management packs, and that you are basically familiar with Microsoft Azure as an infrastructure-as-a-service platform. This is a book about best practices, design concepts, how-tos, and in-depth technical troubleshooting. It covers the role of the Operations Manager product, the best practices for working with management packs, how to use the reporting feature to simplify managing the product, how to thoroughly troubleshoot, and how to use and install Operations Manager in a Microsoft Azure Public Cloud environment.

## About the companion content

---

The companion content for this book can be downloaded from the following page:

*<http://aka.ms/OpsMgrFE/files>*

The companion content includes the following:

- The SQL query in Chapter 1 that you can run in SQL Server Management Studio to determine which collation settings you are using
- The series of commands used in the example in Chapter 2 to run workflow tracing manually
- The Windows PowerShell script used in Chapter 4 to view all TLMEs that exist order per resource pool and per current owning pool member (management server)
- The various SELECT queries included in Chapter 4
- A PDF file titled HealthService Event Reference that provides information about the events that Operations Manager can log to its event log from the HealthService features.

## Acknowledgments

---

We would like to thank Daniele Muscetta, Microsoft Program Manager for Azure Operational Insights, for his review and comments on the Azure Operational Insights section of Chapter 5;

Stefan Stranger, Microsoft Senior Premier Field Engineer, for the review of and his input on the remainder of Chapter 5; and Danny's loving wife, Vita Martinsone, for the pre-editing and formatting of our work.

## Free ebooks from Microsoft Press

---

From technical overviews to in-depth information on special topics, the free ebooks from Microsoft Press cover a wide range of topics. These ebooks are available in PDF, EPUB, and Mobi for Kindle formats, ready for you to download at:

*<http://aka.ms/mspressfree>*

Check back often to see what is new!

## Errata, updates, & book support

---

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

*<http://aka.ms/OpsMgrFE/errata>*

If you discover an error that is not already listed, please submit it to us at the same page.

If you need additional support, email Microsoft Press Book Support at *[mspinput@microsoft.com](mailto:mspinput@microsoft.com)*.

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to *<http://support.microsoft.com>*.

## We want to hear from you

---

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

*<http://aka.ms/tellpress>*

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

## Stay in touch

---

Let's keep the conversation going! We're on Twitter: *<http://twitter.com/MicrosoftPress>*.



# The role of Operations Manager

This chapter introduces the role of Microsoft System Center 2012 R2 Operations Manager in a typical enterprise environment. The chapter describes the architecture and workflow of an Operations Manager environment, best practices for operating system and Microsoft SQL Server configuration, and sizing guidelines.

## Examining a typical Operations Manager installation

---

Operations Manager is part of the Microsoft System Center suite of products. Within this suite, Operations Manager is the part that provides you with infrastructure monitoring and application performance monitoring.

**See also** *More information, trial versions, and support can be found at <http://www.microsoft.com/en-us/server-cloud/products/system-center-2012-r2/default.aspx>.*

Infrastructure monitoring is the comprehensive monitoring of physical, virtual, and cloud infrastructure. This includes the whole infrastructure: your servers and any network devices connecting your servers. It doesn't matter whether the servers are physical systems or virtual machines (VMs), whether they are hosted on-premises or in the cloud, or whether they are running Microsoft Windows or UNIX/Linux. When you apply network monitoring, you also get a view of how all the infrastructure devices are connected in the Network Vicinity dashboard.

Application performance monitoring provides deep insight into the health of your applications. This is done using a lightweight profiler that is installed to capture runtime information from the code of your application. With this information, for instance, you can see which line of code and which underlying T-SQL query is causing an unacceptable response time of a specific page within your application.

**NOTE** By default, every installation of Operations Manager is not registered; it's installed as an evaluation version. This is true even if you installed it from volume licensing media. To register your installed environment, you can use the Windows PowerShell cmdlet `Set-SCOMLicense`. For the full description of how to do this, see <http://support.microsoft.com/kb/2699998>.

# Understanding Operations Manager architecture and workflow

From an architectural point of view, the main features of Operations Manager are the databases, the management servers, and the agents. Additional features are the two SQL Reporting databases, ReportServer and ReportServerTempDB. Optionally, if you have installed Audit Collection Services (ACS), there is also the OperationsManagerAC database.

The two main Operations Manager databases are the Operational Database and the Data Warehouse Database.

- **Operational Database** This is the working database, used mainly as a repository for the configuration information and the raw monitoring data coming in from the agents.
- **Data Warehouse Database** This is the long-term storage location for all aggregated information needed for the reports.

The management servers connect directly to both databases. If you have only one management server, there is no failover and you will not have a highly available solution. Since Operations Manager was designed with built-in high availability when you have two management servers, having two or more management servers is recommended. That way, if one goes down, failover is possible. To determine which management server is down and which is still up and running, the server running the Operational Database serves as a watcher node, similar to a witness in a failover cluster, and has a majority in deciding which one is the functional management server. With three management servers, this additional role for the Operational Database server is not needed. However, this role can and will safely stay—even though it's possible, you should not remove this role.

The agents connect to the primary management server they are assigned to. From this management server, the agent receives management packs. They provide discoveries, rules, monitors, groups, and knowledge. Management packs are installed through the Operations Manager console and imported into the databases by the management server your console is connected to. By running the provided discovery management packs, the agent verifies which monitoring management packs are applicable to the system the agent is installed on. When the agent knows which management packs to run, the workflows from these management packs start running on the agent, locally on each machine that has the agent installed and where the monitoring management pack applies. Agents can also discover external managed objects if proxying is allowed, which would be needed, for instance, to discover Active Directory or clustering attributes. Although this poses a potential security risk since an agent can submit data on behalf of another agent, the risk is very low and no such cases have been reported. You can set Agent Proxying through the console or by using the Enable-SCOMAgentProxy Windows PowerShell cmdlet.

If the primary management server for the agent goes down, the agent tries to connect to one of the management servers defined as a failover. You can define a failover management

server through the console by using AD integration or by using the Set-SCOMParentManagementServer cmdlet with the –FailoverServer parameter. UNIX/Linux systems and network devices, unlike Windows-based servers, do not connect to a specific management server. They connect to a resource pool. Having more than one management server configured in this resource pool provides automatic failover.

The output of the workflows running on the agents is returned to the management server, which puts the data in both databases. The Operational Database contains the raw data, and the Data Warehouse Database contains aggregated data. Keep this in mind when creating dashboards or reports to fetch this data since aggregated data has fewer data points available. In other words, the raw data contains all data points and the aggregated data contains only changes to previous values. A line drawn from aggregated data will not be as smooth as a line drawn from raw data.

## Extending Operations Manager with cloud services

Apart from the built-in functions of Operations Manager and the management packs that you can download and install in the Operations Manager environment, you can also integrate cloud services running in Microsoft Azure. These Azure cloud-based extensions provide additional value to Operations Manager.

Global Service Monitor, or GSM, was the first extension released as a plug-in for Operations Manager. GSM is a cloud service that extends the application monitoring capabilities in Operations Manager beyond your organization's network boundary. With GSM, you can monitor applications from the perspective of the customers who use them. GSM uses Azure points of presence to provide a true reflection of an end-user's experience of a web application. Because GSM monitors from locations that are correlated to customer geographies, application owners can gain insight into customer experiences in addition to the separate problems related to external factors, such as Internet or network problems, from application or service problems. The GSM monitoring experience focuses on the application instead of the infrastructure or individual URL.

**See also** For more information on GSM, see <http://technet.microsoft.com/en-us/library/jj860368.aspx>.

The next extension released was System Center Advisor, now renamed as Azure Operational Insights. Using Azure Operational Insights, you can prevent server configuration problems before they impact performance or availability. This online service analyzes your workloads by collecting data from your installations. It generates alerts that identify potential issues (such as missing security updates) or deviations from identified best practices with regard to configuration and usage. Azure Operational Insights provides you with the combined knowledge of the Microsoft Support Engineers, who are responsible for adding rules to the product. These rules work like an additional management pack that is managed centrally by Microsoft. When the Microsoft Support Engineers notice recurring problems, they create a new rule to check whether the cause of the problem is present in your environment. If it is, you see

a new alert in the Operations Manager console, warning you of the potential problem and providing you with a Knowledge Base article explaining how you can mitigate it.

## Best practices for operating system configuration

---

This section includes a collection of best practices regarding the configuration of the Windows Server operating system on the servers that you use to run your Operations Manager environment.

### Power management

One of the lesser known influences on operating system performance is power management. Power management is useful for saving battery life on a portable system, but it can cause performance issues on servers.

By default, on Windows Server 2008 R2 and higher, power management is set to Balanced. In some cases, you may experience degraded overall performance on a Windows Server machine when running with the default power plan. This is most noticeable on the SQL server running the Operational Database, where the Balanced power setting results in slow console performance since most of the background actions in the console are SQL query-based. The issue may occur irrespective of platform and may be exhibited in both physical and virtual environments. The degraded performance can increase the average response time for some tasks and cause performance issues with CPU-intensive applications. The performance improvement you achieve when you reconfigure the power management setting, especially on physical hardware, can be extensive: up to 40 percent performance improvement when enabling High Performance.

Another power management setting to consider is described in the Knowledge Base article “Degraded overall performance on Windows Server 2008 R2” at <http://support.microsoft.com/kb/2207548>. Note that even though this article describes the problem in the context of Windows Server 2008 R2, the strategies described are also valid for later versions of Windows Server. The performance difference is simply larger for Windows Server 2008 R2.

You can find some important information about the power management setting on a network adapter at <http://support.microsoft.com/kb/2740020>. As stated in the Knowledge Base article, you might want to disable the Allow The Computer To Turn Off This Device To Save Power network adapter power management setting on servers. It is possible to configure some network interface cards (NICs) to turn off automatically when they are not in use. Choosing this option lowers the computer's power consumption, which is useful for laptop computers in particular. But enabling this option on a server that is expected to be connected to the network all the time can result in unexpected or inconsistent behavior.

## Disk performance

Operations Manager uses disk storage locally on every machine that has the agent installed. The agent collects data from all the workflows it is running, caching it locally until it can send it to the management server. The management server does the same for all the data that is coming from all the agents that are connected to the management server. All the data from all the management servers is stored as raw data in the Operational Database and synced with the Data Warehouse, where all the data is aggregated and stored for 400 days by default. Therefore, and because the Operations Manager console relies heavily on the Operational Database, it is very important to have good disk performance on at least the disks that are hosting the files for the Operational Database. Less critical is the disk performance on the Data Warehouse, management servers, and agents, although it will still have an impact on the overall performance of your Operations Manager environment.

**See also** *Find more information about which performance counters need to be investigated more closely in the “Performance counters to be investigated deeper” section of Chapter 4, “Troubleshooting your Operations Manager environment.”*

## Disk optimization for IaaS virtual machines

When Operations Manager is installed on VMs in Azure infrastructure as a service (IaaS)—for instance, to serve as a test environment for management pack development or disaster recovery scenarios or to have an Operations Manager environment close to your application servers in Azure—you can take certain measures to optimize the disk performance. For example, the disks of Basic tier machines have a 300 I/O operations per second (IOPS) limit, while Standard tier machines have a 500 IOPS limit. The limit of a single storage account is 20,000 IOPS. These are all maximum limits; they are not guarantees that you will have the maximum number of IOPS per disk. You can use the Operations Manager Sizing Helper Excel tool to calculate the actual number of IOPS you will need. This tool is described in the section titled “Operations Manager environment sizing” later in this chapter.

Because you can have multiple data disks connected to the larger IaaS VMs and because you can join these data disks together using Storage Spaces, you can surpass the 300- or 500-IOPS limit. This can make a big difference, especially for the SQL server hosting the Operational Database, which will also impact the speed of your Operations Manager console.

From a cost perspective, you can shut down the machines if they are no longer needed, as might be the case in a management pack or disaster recovery test environment. You pay only for what you use in Azure, so whether you put all data on one disk or spread it over multiple disks, you will see only the benefit of the increased performance.

The D drive on an Azure IaaS VM is a temporary disk, using local storage from the actual hardware that is hosting your VM. This means that everything on this drive will be lost in the case of a reboot, so don’t use it to store anything that you want to keep. Instead, put the operating system page file on this disk since it will be automatically re-created when the machine reboots, and the local storage might be faster than the blob storage that is hosting

your operating system and data disks, especially on the D and G series VMs that have local solid state disk (SSD) storage.

If you would like to put the TempDB of your SQL Server installation (or the Buffer Pool Extensions in the case of SQL Server 2014) on this D drive, make sure that you re-create the directory before SQL Server starts. To do this, set the startup of the SQL Server and SQL agent services to Manual. You can then create a script to re-create the necessary directory structure and to start the SQL services. This script can then run automatically at system startup.

**See also** Find more information about SQL Server at <http://blogs.technet.com/b/dataplatforminsider/archive/2014/09/25/using-ssds-in-azure-vm-to-store-sql-server-tempdb-and-buffer-pool-extensions.aspx>. The general information found in this blog post also applies to SQL Server: <http://blogs.msdn.com/b/mast/archive/2014/10/14/configuring-azure-virtual-machines-for-optimal-storage-performance.aspx>. The previous blog post includes a Windows PowerShell script to automate the creation of VMs in Azure with optimal disk performance.

**TIP** To test the speed of your disk subsystem, use the SQLIO Disk Subsystem Benchmark Tool from Microsoft, available at <http://www.microsoft.com/en-us/download/details.aspx?id=20163>.

## Management server placement

Because all management servers are connected in resource pools (groupings of management servers into functional roles), the latency should be less than 5 ms between the management servers and the SQL servers hosting the databases. This means that it's not recommended to place a management server in a remote subnet if the latency is more than 5 ms. In that case, you should have the agents connect over the slower link directly to the management servers or use a Gateway server.

**NOTE** Putting a Gateway server in a remote subnet to compress the outgoing data is no longer recommended. The agent itself does an equally good job of compressing the data in Operations Manager 2012 R2. However, the other reasons for installing a Gateway server in a remote subnet are still valid, for instance to reduce the administrative overhead and to minimize the number of certificates that are needed. More information can be found at <http://technet.microsoft.com/en-us/library/hh212823.aspx>.

**See also** A complete explanation of how to install a Gateway server can be found at <http://blogs.technet.com/b/pfesweplat/archive/2012/10/15/step-by-step-walkthrough-installing-an-operations-manager-2012-gateway.aspx>.

## Antivirus software exclusions

When you install Operations Manager on machines running antivirus software, you should configure the antivirus software so that the following directories are excluded:

- The Health Service State folder on every management server and every agent
- The data and log file directories where your databases are located

Excluding the actual binary files, such as MonitoringHost.exe, is not recommended.

**See also** A detailed overview of antivirus exclusions for Operations Manager can be found at <http://support.microsoft.com/kb/975931>.

## Best practices for SQL Server configuration

---

This section includes a collection of best practices for configuring the SQL Server installation that you use to run your Operations Manager environment.

**NOTE** The best way to configure SQL Server in your Operations Manager environment is to keep it simple. The default settings for Operations Manager should be left alone unless you have very specific reasons to change them. The sections that follow describe some specific optimizations to use and some things to avoid.

### Auto grow and auto shrink

Auto grow and auto shrink are database-level settings. Auto grow enables your database file to grow larger when needed, and auto shrink enables the file to be trimmed if there is too much free space in the database.

Neither auto grow nor auto shrink are recommended for the Operational Database because it needs 50 percent of free space at all times to perform maintenance and indexing tasks. If the database doesn't have enough free space, the scheduled maintenance tasks might fail. Operations Manager will alert you when there is less than 40 percent of free space.

Auto shrink also does not make sense for the Data Warehouse Database since it grows larger over time.

The SQL Server edition you are using also has an important role when you are considering auto grow. SQL Server Standard edition can cause the database tables to lock out when auto grow is configured. However, this does not occur with SQL Server Enterprise edition. This applies to both the Operational Database and the Data Warehouse Database.

Auto grow is supported (though not recommended), when enabled as an insurance policy against the database's file filling up. When using auto grow on the databases, it is better to set it to increase by a fixed amount rather than a percentage. The fixed increase amount should be no more than 500 MB or 1 GB in growth to limit the blocking that might occur during the

expansion process. It is also useful to configure a maximum possible size to prevent the databases from filling up the disk they reside on.

You can use the Operations Manager Sizing Helper (a Microsoft Excel file) to determine the disk space that is needed for your Operational Database and Data Warehouse Database. The Sizing Helper is described in detail at the end of this chapter.

**See also** *More information about auto grow and auto shrink can be found in the article at <http://support.microsoft.com/kb/315512/>.*

## Instant file initialization

To improve the performance of auto grow operations, you can configure Instant Data File Initialization by configuring the Perform Volume Maintenance Tasks local security policy. Data and log files are initialized to overwrite any existing data left on the disk from previously deleted files. Data and log files are first initialized by filling the files with zeros. In SQL Server, data files can be initialized instantaneously. This allows for fast running of file operations. Instant file initialization reclaims used disk space without filling that space with zeros. Instead, disk content is overwritten as new data is written to the files. Log files cannot be initialized instantaneously. Instant file initialization is available only if the SQL Server (MSSQLSERVER) service account has been granted the right to perform volume maintenance tasks (SE\_MANAGE\_VOLUME\_NAME). Members of the Windows Administrator group have this right and can grant it to other users by adding them to the Perform Volume Maintenance Tasks security policy.

**See also** *Find more information about instant file initialization at [http://blogs.msdn.com/b/sql\\_pfe\\_blog/archive/2009/12/23/how-and-why-to-enable-instant-file-initialization.aspx](http://blogs.msdn.com/b/sql_pfe_blog/archive/2009/12/23/how-and-why-to-enable-instant-file-initialization.aspx) and [http://sqlblog.com/blogs/tibor\\_karaszi/archive/2009/03/09/do-you-have-instant-file-initialization.aspx](http://sqlblog.com/blogs/tibor_karaszi/archive/2009/03/09/do-you-have-instant-file-initialization.aspx)*

## Maximum degree of parallelism

If you have a server with multiple cores, you should use them to spread the load. A maximum degree of parallelism (MaxDOP) setting of 1 on such servers will limit queries to one CPU core while the other cores might be idle. When changing this setting, verify the change before putting it into production and determine which setting works best for the workload in your specific environment.

The MaxDOP setting defines the number of processors that are used for running a query in a parallel plan. This option determines the computing and thread resources that are used for the query plan operators that perform the work in parallel. You can use a calculator to determine the optimal MaxDOP setting (available at <http://blogs.msdn.com/b/sqlsakthi/p/maxdop-calculator-sqlserver.aspx>). Since this is a complicated calculation that depends on your specific environment configuration, verify the new setting in a test environment before applying it on production machines.



**See also** You can find a Knowledge Base article describing the recommendations and guidelines for the MaxDOP configuration option in SQL Server at <http://support.microsoft.com/kb/2806535>.

## Maximum memory

By default, SQL Server is configured to use 2 TB of available memory, which is probably more than what is available on the SQL server hosting your Operations Manager databases. SQL Server will reserve all available memory that it determines it might need at some point. This is true for all separate instances hosted by your SQL server. Therefore, it is best to limit the memory SQL Server can access, calculated over all the instances the host machine is running.

As a general rule, set the combined value over all the instances to about 2 GB less than the actual memory available on the host. This will secure enough available memory for the operating system to function optimally. Also, it's a good idea to provide your SQL server with enough memory in the first place—running multiple instances on a total of 4 GB of memory is not an ideal configuration.

## Splitting up TempDB files

Another low-effort, high-reward action is splitting up the files that comprise the TempDB. There's only one TempDB per SQL Server instance, so it's often a performance bottleneck. Make sure that the disk subsystem that holds the TempDB files is up to the task. Increase the number of data files that make up your TempDB to maximize disk bandwidth and to reduce contention in allocation structures. Discover the best configuration for your specific environment by performing tests in a test environment.

Generally, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files; if contention continues, increase the number of data files by multiples of four (up to the number of logical processors) until the contention is reduced to acceptable levels or make changes to the workload/code. It is also best to spread these different files over multiple disk systems and to keep all files the same size. That way, SQL will use the different files in a round-robin fashion, which, of course, is exactly what you would like to accomplish—load balancing.

It is also recommended that you size the TempDB according to the Operations Manager environment. The default size for TempDB is 8 MB with a 1-MB log file. Every time you restart SQL, it will re-create this 8-MB file from the model database. With such a small size, your TempDB needs to grow considerably in the beginning, and during auto grow, your SQL server might be temporarily unavailable.

**NOTE** The log file for TempDB should remain a single file at all times.

**See also** TempDB optimization is explained in detail at [http://technet.microsoft.com/en-us/library/ms175527\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/ms175527(v=sql.105).aspx).

**See also** For SQL servers running in Azure IaaS, please refer to <http://msdn.microsoft.com/library/azure/dn248436.aspx> and <http://msdn.microsoft.com/en-us/library/azure/dn133149.aspx>.

## Recovery model

Some SQL teams automatically assume that all databases should be set to Full recovery model. This requires backing up the transaction logs on a regular basis, but gives the added advantage of restoring up to the time of the last transaction log backup. This approach does not make as much sense for Operations Manager because the data changing on an hourly basis is of little value compared to the complexity added by moving from Simple to Full recovery model. Also, changing to Full means transaction logs checkpoint only when a transaction log backup is performed. The only time to change to Full recovery model is when you are using an advanced replication strategy, like log shipping, AlwaysOn Failover Cluster, or AlwaysOn Availability Groups, which requires a Full recovery model.

**See also** Find more information about high availability options for SQL Server at [http://technet.microsoft.com/en-us/library/ms190202\(v=sql.110\).aspx](http://technet.microsoft.com/en-us/library/ms190202(v=sql.110).aspx).

## Service Principal Names

Service Principal Names (SPNs) are needed for Kerberos authentication. If SPNs are not available, there will be a fallback to Windows NT LAN Manager (NTLM) authentication instead of Kerberos authentication, which is less secure because there is no mutual authentication (among some other reasons).

An SPN provides the client connecting to the SQL Server service with certain information:

- Type of service (in this case, SQL Server, or MSSQLSvc)
- Name of the server
- Port
- Service account running the service

When the SQL Server service starts, it tries to register its own SPN. To understand how this works, it helps to understand how SQL Server handles Windows logins. Basically, SQL Server only handles the authentication of a SQL login, not that of a Windows login. Windows login authentication is passed off to the operating system via the Security Support Provider Interface (SSPI). If the SSPI says the login is good, SQL Server allows the login; if SSPI says it's bad, it doesn't. Since the default and recommended settings for SQL Server in an Operations Manager environment are Windows authentication, SQL Server uses Windows logins.

It is best practice to use a domain account to run your SQL Server service (MSSQLSvc). The problem with this is that if your SQL Server service is not running as either the server's system account or a domain administrator, SQL Server cannot register its Service SPN when the service is started. If the SQL Server service does not have sufficient rights, you can use the SETSPN tool manually as a domain administrator to register the necessary SPNs.

In Operations Manager, two services require a registered SPN: the Health Service (Microsoft Monitoring Agent service) and the SDK Service (System Center Data Access service). Both services try to register themselves in Active Directory at startup. Since you need one entry in the SPN list for the NetBIOS name and one entry for the fully qualified domain name (FQDN) of each service, four entries per Operations Manager management server are required.

**See also** Find more information about registering an SPN for Kerberos connections at <http://msdn.microsoft.com/en-us/library/ms191153.aspx>. More information about SPNs in Operations Manager can be found at <http://blogs.technet.com/b/kevinholman/archive/2011/08/08/opsmgr-2012-what-should-the-spn-s-look-like.aspx>.

## Re-indexing

By default, Operations Manager does self-maintenance. Since most Operations Manager administrators are not SQL Database Administrators (DBAs), Microsoft implemented several rules in Operations Manager to automatically keep the databases optimized. These maintenance tasks are defined as system rules in the Operations Manager management pack, one of the management packs installed by default when you install Operations Manager. Since these maintenance tasks run automatically, be careful that your own maintenance tasks do not conflict with the built-in system rules (if you or the DBA decide to implement additional maintenance).

**See also** For information about the schedules of these built-in system rules, see <http://technet.microsoft.com/en-us/library/hh212782.aspx>.

For the Operational Database, you can check the current level of fragmentation by running the DBCC SHOWCONTIG WITH FAST command in the SQL Server Management Studio query window. The output of this command will show that some tables are more fragmented than others. This is normal since Operations Manager does not optimize certain tables containing raw data. You can re-index these tables by running your own maintenance jobs. Be careful, however, that these additional maintenance jobs do not take too long, do not generate too much I/O, and do not conflict with the built-in maintenance rules; that could interfere with the normal operation of Operations Manager.

For the Operations Manager Data Warehouse, an automatic maintenance job runs every 60 seconds. This job, coming from the Standard Data Warehouse Data Set maintenance rule, does many things, of which re-indexing is only one. All the necessary tables are updated and re-indexed as needed. When a table is 10 percent fragmented, the job re-organizes it. When the table is 30 percent or more fragmented, the index is re-built. Therefore, especially since the built-in maintenance runs every 60 seconds, there is no need for a DBA to run any UPDATE STATISTICS or DBCC DBREINDEX maintenance commands against this database.

## Block size

By default, the block size of any disk less than 16 TB is 4 K. Since SQL Server reads in 64-K increments, it is best practice to format the disk containing the SQL data and log files with 64-K block size. You can only set this allocation unit size when you format the disk. To make sure that you have selected the correct block size, run the CHKDSK tool to verify the number of bytes in each allocation unit. The value should be 65,536, or 64 K. It should not be 4,096, since that would mean it is 4 K (the default). For cluster shared volumes (CSVs), you can run CHKDSK by following the instructions at <http://blogs.msdn.com/b/clustering/archive/2014/01/02/10486462.aspx>.

**See also** *Disk partition alignment is further explained at <http://blogs.msdn.com/b/jimmymay/archive/2014/03/14/disk-partition-alignment-for-windows-server-2012-sql-server-2012-and-sql-server-2014.aspx>.*

## Collation

The collation controls the sorting, storage, indexing, and comparisons of characters in a database. There are collation settings on both the SQL instance and database level. To ensure that the system functions correctly, it is essential to use the appropriate collation for both the instance and the database according to the type of data that you intend to store. If you use the wrong collation, searches may be less effective or not work at all, sorting might produce unexpected results, and other problems can happen when inserting or retrieving data.

Guidance concerning which collations are supported and which are not can be found on the System Center Service Manager (not Operations Manager) blog at <http://blogs.technet.com/b/servicemanager/archive/2012/05/24/clarification-on-sql-server-collation-requirements-for-system-center-2012.aspx>. The most important information from this blog post concerns the collation for the Operations Manager Data Warehouse, which indicates, “The Operations Manager Data Warehouse installer will *always install the data warehouse with SQL\_Latin1\_General\_CP1\_CI\_AS* regardless of the SQL Server collation. Until this issue is fixed, please always install the Operations Manager Data Warehouse on a SQL server with the SQL\_Latin1\_General\_CP1\_CI\_AS collation. There are some compatibility issues when the Temp database on the DW SQL Server instance is anything other than SQL\_Latin1\_General\_CP1\_CI\_AS and the data warehouse DB is SQL\_Latin1\_General\_CP1\_CI\_AS.”

If a SQL Server collation other than SQL\_Latin1\_General\_CP1\_CI\_AS is specified when you create the database, you will have to reinstall Operations Manager and create another database to fix this problem because you cannot change the collation after installing Operations Manager. However, if the collation of the database itself is correct but the collation of the SQL Server instance is wrong, the database can be migrated to another SQL server with the correct collation configuration.

To determine which collation settings you are using, run the following queries in SQL Server Management Studio:

```
-- Get SQL Server collation
SELECT CONVERT (varchar, SERVERPROPERTY('collation')) as SQL_Server_Collation
-- Get OM DB Collation
SELECT DATABASEPROPERTYEX('OperationsManager', 'Collation') as OM_Database_SQLCollation;
-- Get DW DB Collation
SELECT DATABASEPROPERTYEX('OperationsManagerDW', 'Collation') as
DM_Database_SQLCollation;
```

**NOTE** There are some known issues with certain third-party management packs where reports are failing to generate when the collation is incorrect.

## Operations Manager registry optimizations

---

Many different features are used in the various Operations Manager processes. It is important to know these features and which ones can be configured using the registry. In general, for applications and some Windows features, certain registry keys have a default value hard-coded in the source code if applicable. When a particular setting (configurable via a registry entry) is needed, it is first read from the registry. If it exists, the value from the registry is used; if there's no value in the registry, the default value is used. Thus, you might discover that some registry keys or entries do not exist by default and that you need to create them before you can use them. The following sections describe some important registry settings and how they may need to be changed depending on your environment. This list is compiled from experience in the field working with customers. The default values for these settings may vary between System Center Operations Manager 2012 and 2012 R2 (update rollups included). The following information refers only to the default values in Operations Manager 2012 R2.

**NOTE** Some registry settings require a certain update rollup to be installed. This is noted where applicable.

**IMPORTANT** Making changes to the Windows registry can be dangerous if not done correctly. Please make sure to create a backup of the specific registry key before creating or deleting entries from it. Also make sure that you create the new entry in the correct key and that you give it the correct type and value. Pay attention when creating registry keys or entries using copy and paste from a rich text format source so that no unwanted space characters are entered upon pasting.

## Data Access Layer

The Data Access Layer is a feature used by the Health Service (through the Monitoring Host process where modules are loaded) and the Data Access Service to read and write data to the Operational Database and the Data Warehouse Database. The registry key path where settings for the Data Access Layer are included is:

*HKLM\SOFTWARE\Microsoft\System Center\2010\Common\DAL*

The DWORD setting, called `DALInitiateClearPool`, is used by the Data Access Service to control whether to reconnect to the database after a period of unavailability. The default value is 0 (disabled). The recommendation is to enable this feature by setting the value to 1 (decimal).

Directly related to the `DALInitiateClearPool` setting is the `DALInitiateClearPoolSeconds` DWORD setting. This setting controls the interval (in seconds) by which the Data Access Service should try to reconnect to the database after a period of unavailability. The default value of 60 (decimal) is recommended.

The Data Access Layer does have a retry mechanism in case it loses the connection to the database, but only for brief (a few seconds) unavailability scenarios. Any unavailability longer than a few seconds relies on the `DALInitiateClearPool` and `DALInitiateClearPoolSeconds` settings for the service to recover when the database is available again. For example, these settings are needed for installations where the databases are hosted on a clustered SQL Server instance. On fail over, if the settings do not exist and, thus, the feature is not enabled, the Data Access Service might not try to reconnect again (depending on how long the fail over takes).

## Persistence Manager

The Persistence Manager feature is used by the Health Service to read and write data to the local database. The local or cache database is called `HealthServiceStore.edb`, and it is a Microsoft Jet Database Engine database. The registry key path for settings belonging to this feature is:

*HKLM\SYSTEM\CurrentControlSet\Services\HealthService\Parameters*

The setting responsible for how often Persistence Manager writes data from memory to the disk is called `Persistence Checkpoint Depth Maximum` of type DWORD and is measured in bytes. The default value for this setting is 20971520 (decimal) bytes. On management servers that handle a large number of objects not managed directly by agents, such as SNMP Devices, Groups, URL Monitors, Cross-Platform Agents, and so on, you may need to increase this value to relieve disk pressure. The recommended value is 104857600 (decimal).

Another important setting is `Persistence Cache Maximum` of type DWORD. This setting controls the amount of memory in pages used by Persistence Manager for its data store on the local database. The default value for this is 262144 (decimal), which is also the recommended value. If you are running an older version of Operations Manager on management servers that manage a large number of objects, you should change the value to 262144 (decimal).

Persistence Manager also stores different versions of the data store. The DWORD Persistence Version Store Maximum setting controls how much memory (in pages) the version store can have. The default value is 131072 (decimal). On management servers where you have increased the value for Persistence Cache Maximum, you might also want to increase the value of Persistence Version Store Maximum.

**NOTE** Increasing the allowed memory usage of Persistence Manager enables larger amounts of data to be collected and relieves disk pressure, but also increases memory (RAM) usage for the Health Service.

## Health Manager

Health Manager is used by the Health Service to calculate and track the health state of each monitor of each object it monitors. The registry path for settings belonging to this feature is:

*HKLM\SYSTEM\CurrentControlSet\Services\HealthService\Parameters*

The important setting for the Health Manager is State Queue Items of type DWORD. This sets the maximum size (in bytes) of the state data queue. If the value is too small or if there are too many workflows running (based on the number of objects being managed), there could be possible state change data loss. The default value for this setting is calculated by the Health Service on startup based on how many objects it needs to manage. For agents in a small environment, this value is set to 1024 (decimal). The value is set to 10240 (decimal) on management servers in a mid-size environment. For large environments, on management servers that manage many objects, the default is 25600 (decimal). The recommendation is to double these default values, depending on where it is needed—for an agent that manages a lot of objects or a management server.

## Pool Manager

Operations Manager 2012 introduced a new feature called Pool Manager that offers management server high availability through resource pools. Pool Manager is used by the Health Service to calculate and track the Health Service's availability from the resource pools it is part of. The registry path for settings belonging to this feature is:

*HKLM\SYSTEM\CurrentControlSet\Services\HealthService\Parameters\PoolManager*

The DWORD PoolLeaseRequestPeriodSeconds setting controls the amount of time the pool member holds a lease, starting from the time the lease is requested until a new lease needs to be requested. The default value is 120 (decimal). In some large environments where there are many objects managed by the management servers, or maybe where there are sporadic network speed issues, this setting might need to be tweaked. A value no higher than 600 (decimal) is recommended.

The other setting involved in the resource pool member availability calculation is the DWORD PoolNetworkLatencySeconds setting. This setting controls the total worst case round trip time for a pool message (availability checks), and the default value is 30 (decimal). A value no higher than 120 (decimal) is recommended.

Setting values too high for these settings is not recommended because it will slow down the process in which a pool member detects an issue and stops processing workflows. If this happens, some objects managed by the pool may become managed by two Health Services (management servers) at the same time, which would cause duplicate workflows to be performed for that particular object.

**IMPORTANT** Do not change the settings for Pool Manager unless advised by Microsoft Support after a proper analysis of the environment, behavior of the resource pools, and load on the management servers. If these settings are changed, it is important to make sure that they are changed to the same value on *all* management servers in the environment.

## Group Calculation module

The Group Calculation module is used by the Health Service (through the Monitoring Host process) to periodically add or remove objects from groups based on the criteria of the group definitions. The registry path for settings belonging to this module is:

*HKLM\SOFTWARE\Microsoft\Microsoft Operations Manager\3.0*

Group Calculation is a workflow that runs per group. The more groups you have, the more workflows there will be. The more complicated the criteria, the more time and resources it will take for these workflows to finish. This can have a considerable impact on performance. The recommendation is to prefer explicit group membership rules over dynamic group membership rules whenever possible. For dynamic groups, it is also good practice to keep the group calculation criteria as simple as possible. Complex dynamic group calculation criteria rules can have a significant impact on management server performance. A way to tweak this is to set a bigger interval so that the group calculation workflows run less often. This is controlled by the DWORD GroupCalcPollingIntervalMilliseconds setting. The default value is 30000 (decimal); 90000 (decimal) is recommended for environments with many dynamic groups.

## SNMP module

The SNMP module is used by the Health Service (through the Monitoring Host process) to monitor SNMP devices (network devices) that have been discovered in the environment. The registry path for settings belonging to this module is:

*HKLM\SOFTWARE\Microsoft\Microsoft Operations Manager\3.0\Modules\Global\SNMP*



The DWORD `MaxOidsPerRequest` setting is used to identify the maximum number of object identifiers (OIDs) that will be queried per batch. If some OID values are too big, the data package with the results might exceed the buffer size. In this situation, you'll see the following error message on all or some of the management servers in the resource pool(s) responsible for monitoring SNMP devices: Incoming SNMP request (<NUMBER\_OF\_BYTES> bytes) from IP address <IP\_ADDRESS> and <PORT PORT\_NUMBER> Interface "inside" exceeds data buffer size, discarding the SNMP request. If this happens, you can decrease the value for this setting, basically make it smaller than the default, until the error message no longer appears. The default value for this setting is 50 (decimal).

## Management Configuration Service

The Management Configuration Service is responsible for configuration updates for the Health Service, ensuring that it has the latest changes from the database. The registry key path for this feature is:

*HKLM\SOFTWARE\Microsoft\Microsoft Operations Manager\3.0\Config Service*

The Management Configuration Service regularly checks the Operational Database for any changes. The registry setting that controls the frequency of these checks (in seconds) is `PollingIntervalSeconds` of DWORD type, and its default value is 30 (decimal). This process may have a significant performance impact on large environments with a large type space and frequent configuration changes. In such situations, the recommended value is 120 (decimal). This setting can bring a significant improvement in the overall performance, but setting the value too high can mean a delay in changes being loaded into the Health Service on the management servers, and this could cause delays when implementing new workflows, overrides, management packs, or discovered inventory.

Another important setting to address in the configuration update process is `CommandTimeoutSeconds` of type DWORD. The setting controls the timeout (in seconds) for the delta synchronization process (on configuration updates) and determines how long it runs. If the timeout is reached, the process is stopped and the configuration is not updated. In large environments with frequent configuration changes, a longer timeout is recommended. The default value is 30 (decimal), and the recommended value is 120 (decimal).

**NOTE** This registry setting is only available starting with Update Rollup 3 for Operations Manager 2012 R2.

## Data Warehouse module

The Data Warehouse module is used by the Data Access Service and Health Service (through the Monitoring Host process) to read and write data to the Data Warehouse Database. The registry key path for this feature is:

*HKLM\SOFTWARE\Microsoft\Microsoft Operations Manager\3.0\Data Warehouse*

One of the most important workflows performed using the Data Warehouse module is the StandardDatasetMaintenance workflow. This is performed for each data set (alert, event, performance, state, and custom data sets like the Exchange 2010 management pack creates). The StandardDatasetMaintenance workflow transforms the raw data from the Data Warehouse raw data set tables into daily and hourly data (called data aggregation) and stores it in the appropriate tables. This workflow is also responsible for grooming (deleting) the raw data after it is processed (aggregated). The setting that defines the timeout for this workflow is Command Timeout Seconds of type DWORD, and its default value (in seconds) is 300 (decimal). In large environments or environments where a large amount of data is collected for one or more data sets, the default timeout is usually not big enough, so in these cases, the recommended value is 1200 (decimal).

The Data Warehouse module is also used by the data synchronization workflow that moves new data from the various data sets from the Operational Database to the Data Warehouse Database. This workflow also has a certain timeout configured to set how long it runs. This can be controlled by the Bulk Insert Command Timeout Seconds setting of type DWORD. The default value (in seconds) for this setting is 30 (decimal). In environments where a large amount of data is collected for one or more data sets, this value should be increased. The recommended increased value is 90 (decimal).

## Operations Manager environment sizing

---

The Operations Manager 2012 Sizing Helper is an interactive Microsoft Excel tool designed to assist you with planning and sizing deployments of System Center 2012 Operations Manager. The sections that follow describe how you can use this tool.

**See also** *The Sizing Helper is further explained and can be downloaded from <http://blogs.technet.com/b/momteam/archive/2012/04/02/operations-manager-2012-sizing-helper-tool.aspx>. There is also a Windows Phone version of this tool available at [http://blogs.msdn.com/b/wei\\_out\\_there\\_with\\_system\\_center/archive/2014/09/09/opsmgr-introducing-the-om12-sizing-helper-app-for-windows-phone.aspx](http://blogs.msdn.com/b/wei_out_there_with_system_center/archive/2014/09/09/opsmgr-introducing-the-om12-sizing-helper-app-for-windows-phone.aspx).*

## General recommendations

You can use the Sizing Helper to plan the correct amount of infrastructure needed for a new Operations Manager 2012 deployment. The Sizing Helper can provide minimum hardware specifications for each server role, a topology diagram, and storage requirements.

**See also** *To prepare your environment for Operations Manager, follow the recommendations described at <http://technet.microsoft.com/en-us/library/dn249696.aspx>.*

The Sizing Helper helpfully provides a close approximation of the type and size of hardware you will need. Use this tool to calculate environment sizes. You can also use it to calculate the number of spindles you will need depending on the IOPS required to sustain the load the

databases will generate. Since disk performance is one of the main concerns in an Operations Manager environment, you might want to calculate whether your disk subsystem performs well enough. Remember that in an Azure IaaS environment, the disk IOPS limits (described previously) are maximum values and there is no guarantee that you will actually have those IOPS numbers.

As a general best practice, make sure the hardware (physical or virtual) is up to the task. For high availability, make sure that you have two or more management servers and that your SQL server is clustered using Failover Clustering or SQL AlwaysOn. The reporting part, based on SQL Server Reporting Services (SSRS), cannot be clustered. However, SSRS can be set up quickly on another server without losing any data, since the actual data is in the Data Warehouse Database.

## Resource pool considerations

A resource pool is a collection of management servers used to distribute work amongst the collection and take over work from a failed member. Resource pools cover only Health Service functionality; they are a collection of Health Services working together to manage instances assigned to the pool. Workflows targeted to the instances are loaded by the Health Service in the resource pool that ultimately manages that instance. If one of the Health Services in the resource pool fails, the other Health Services in the pool pick up the work that the failed member was running. Windows agents are targeted to specific management servers, but UNIX/Linux agents, web URL monitoring, and network devices are targeted to a resource pool.

If you are planning any UNIX/Linux, web URL, or network monitoring, you should create a separate resource pool for each and add dedicated management servers to this resource pool. Remove these management servers from all other resource pools, such as the All Management Servers Resource Pool, the Notifications Resource Pool, and the AD Assignment Resource Pool. Management servers in the All Management Servers Resource Pool are tasked with the following: group calculation, availability and dependency monitor calculations, health aggregation, and database grooming.

You can have only one network device discovery rule per resource pool. For multiple network discovery rules, you must create multiple network resource pools. The maximum number of network devices that can be managed using a resource pool containing at least three management servers is 1,000. The maximum number of devices per Operations Manager Management Group is 2,000 in two different resource pools containing at least three management servers.

To remove a server from the resource pools with automatic membership, first set the group membership to manual (automatic is the default). This can be done *only* from within Windows PowerShell as follows:

```
Get-ScomResourcePool -DisplayName "<Resource Pool Name>" | Set-SCOMResourcePool  
-EnableAutomaticMembership 0
```

After you run this command, you can then use either Windows PowerShell or the console to remove the management server.

## ITIL and MOF

---

Operations Manager supports two service-oriented frameworks: ITIL and MOF. Using a service-oriented framework helps you keep your Operations Manager environment documented and running smoothly. In the long run, applying a service-oriented framework saves time, even though in the beginning it might seem to add complexity.

### ITIL

The British Office of Government Commerce (OGC) provides their best practices on using information technology in service management and operations in the Information Technology Infrastructure Library, or ITIL. Many organizations are aware of the frameworks called ITIL and Microsoft Operations Framework (MOF). The networking and cooperation between the individual departments within the IT department is always pronounced. Therefore, it is important to employ an ITIL service manager, someone who knows the totality of all business processes. For Operations Manager engineering teams, the knowledge of ITIL and/or MOF is of enormous importance to the correct handling of processes and should be understood and applied as a base for all monitoring approaches.

ITIL provides best practices for IT service management in several publications maintained by the OGC. ITIL describes directions and guidance for delivering IT service in the same high quality manner as other business services in a company, therefore enhancing the business value generated by the IT department.

**See also** Find more information on ITIL at [http://en.wikipedia.org/wiki/Information\\_Technology\\_Infrastructure\\_Library](http://en.wikipedia.org/wiki/Information_Technology_Infrastructure_Library).

### MOF

Based on ITIL industry guidelines, Microsoft has adopted some features and created new ones specific to its operations framework. This effort is called the Microsoft Operations Framework, or MOF. Microsoft has created additional specific guidelines for using Microsoft technologies, especially guidance for the IT life cycle. This is done by taking advantage of the combined experiences from internal product and operations groups, as well as external knowledge from the hands-on experiences of partners, customers, support, and premier field engineers.

**See also** Find more information on MOF at [http://en.wikipedia.org/wiki/Microsoft\\_Operations\\_Framework](http://en.wikipedia.org/wiki/Microsoft_Operations_Framework) and <http://technet.microsoft.com/en-us/library/cc506049.aspx>.

# The role of Operations Manager

Both ITIL and MOF frameworks are guidelines that can help organizations create business-focused IT solutions that use a process model for service support and delivery. The activities and processes defined in both of these frameworks have specific standards, procedures, and tasks and can run simultaneously in an organization. Since they are guidelines, you don't have to follow them to the letter. You will need to adapt them to your specific environment in the same way you will have to adapt Operations Manager to your environment. Operations Manager is built on both of these service-oriented frameworks and supports you in the core areas of ITIL and MOF. The product knowledge and resolutions that come with each management pack will help you focus on your business needs and understand what you should do to keep your services and infrastructure running.

Keep in mind the important process structure surrounding Operations Manager. You can use this process structure to gain control of your monitored infrastructure. Consider the following short example to understand how this applies in a typical line of business (LOB) service. This example is enhanced with additional tips to fulfill the area tasks, which are the high-level steps defined in ITIL and MOF. In this example, the LOB service is integrated with Operations Manager. The explanation covers changing, operating, supporting, and optimizing. With this foundation in mind, you can be more process and service oriented and you can think proactively. If you are not thinking proactively, the process itself is there, but it is not standardized or written down. Adopt the best practice and operation guidance to get control of your monitoring business by remembering the following four areas of the frameworks.

## Step 1: Changing

When a new service solution is unfamiliar to the engineering team, you first need to determine if a management pack from Microsoft or a third-party vendor is available. If there is no available management pack, the team must develop its own. It is best practice for this custom created management pack to be based or developed on a service and health model since Operations Manager operates on a service-oriented, class-based structure in which all the different features of the product have relationships with and dependencies on each other. Creating the custom management pack oriented to a service and health model creates the hierarchy and descriptors of your business service. Authoring the custom management pack (and later modifications to the solution) is covered in the process-as-a-change area where you deploy something new in your lifecycle.

Keep in mind that a change management process should be in place to ensure that Operations Manager changes are implemented on a scheduled basis, and you can make sure that the impact on the user work environment is as low as possible. Adequate testing of all changes prevents negative influence on your business critical systems. You should have a separate environment to design and test management packs because Operations Manager is an infrastructure application. A separate test environment lets you import or export management packs multiple times as you write them without affecting production and without generating configuration changes throughout your environment.

## Step 2: Operating

If everything for your business service is imported and implemented correctly, the solution is in production and in day-to-day operations. Operations Manager can use built-in management pack tasks or automatic responses to work with your service. If there are no tasks or integrations implemented in the management pack, you can create your own tasks or integrate other applications specifically based on the context of your classes and objects. The tasks can run from the console as a command or a script or as agent tasks on the managed systems. To automate diagnostic or recovery actions to avoid having to run them manually, you can enhance your alerts with diagnostic actions that start to diagnose a problem when it occurs and recovery actions to fix a problem when it occurs. This is called the area of operating.

In addition, scheduling ongoing updates to all areas of the Operations Manager deployment (for example, hardware, operating system, Operations Manager, and so on) is recommended.

Think about creating recurring tasks (daily, weekly, and monthly) that should be performed for your Operations Manager environment. The following sections lists Microsoft's recommendations for recurring tasks.

### Daily tasks

- Use the imported management packs (general views, Management Group Health dashboard view and reports) to verify that the Operations Manager features are healthy.
- Check that alerts from the previous day are not still in state of New. Check the repeat counts and date created for your alerts.
- Check for any unusual alert or event noise; investigate further if required (for example, failing scripts, WMI issues, grey agents, and so on).
- Check the status of all agents for any state other than green. Verify that all managed computers are communicating.
- Review nightly backup jobs and database space allocation.
- Verify that predefined maintenance tasks scheduled to run daily are running successfully.
- Check the Operations Manager event logs on each management server for unusual behavior and error events.

## Weekly tasks

- Schedule weekly meetings with operational application owners to review previous most common alerts and events.
- Use the top-down approach to running the Most Common Alerts and Most Common Events reports. Investigate further where necessary.
- Run the Data Volume by Management pack and Data Volume by Workflow and Instance reports.
- Check available disk space on all Operations Manager database systems (data and log files).

## Monthly tasks

- Check for new management pack versions of any installed management packs. Also check the management pack guides of newly released management packs to determine whether they meet the requirements of your organization and are suitable for your environment.
- Review the baselines (performance counters) to assess the ongoing performance of the Operations Manager environment as new agents and management packs are added.
- Review the disaster recovery plan for any needed changes.

## Step 3: Supporting

Operations Manager, together with included and offered management packs, is focused on monitoring your infrastructure and the daily operations of your business services. The knowledge in the management packs is based on the product and company knowledge of Microsoft product groups, vendors, or your own company. By scoping user roles, you can involve different teams and departments of your organization to focus on your solution. It is possible to send notifications to or to forward alerts directly to ticketing systems (for example, using System Center Orchestrator as a connector from System Center Operations Manager to System Center Service Manager). In that fashion, the teams and departments of your organization can resolve issues in a timely manner with provided and verified resolutions for incidents and problems. In the Microsoft lifecycle, this is referred to as the area of support. At some point, products leave mainstream support. You will want to be familiar with these timelines. You can find additional information on Microsoft Support Lifecycle at <http://support.microsoft.com/gp/lifeselect>.

It is recommended that all unsupported customizations created for your Operations Manager deployment are well documented. Documenting your unsupported customizations allows the current Operations Manager ownership group to know what unsupported customizations have been implemented. It also allows for a smooth transition if another group takes over the Operations Manager implementation at some point in time. This documentation can be especially beneficial when contacting Microsoft Support with an issue, applying

Operations Manager update rollups, upgrading to a new version of Operations Manager, or, in the worst case, if dealing with disaster recovery. Unsupported customizations are often forgotten during these processes.

## Step 4: Optimizing

With context-based dashboards, performance views, distributed applications, and reporting capabilities in Operations Manager, you understand and control the normal or bad behavior of your solution. Again, this is built into the imported management packs. You can extend the provided solution by creating your own views, dashboards, distributed applications, or reports in the console or with authoring tools.

**See also** *Microsoft's Brian Wren has put together a very extensive authoring guide, which can be accessed at <http://social.technet.microsoft.com/wiki/contents/articles/15251.system-center-management-pack-authoring-guide.aspx>.*

If you understand the metrics of your application, you can control the business service and use it to increase efficiency. To help you understand the metrics of your application, you can define a distributed application that shows all the separate parts of one service with dependencies on each other. Dashboards and views show you the current state of your application and its features. Reports show you the historical state and allow you to analyze trends and capacity. This area is referred to as optimizing and is one of the main steps of the tuning exercise.



# Best practices for working with management packs

This chapter provides an overview of what is contained in a typical management pack, with a short explanation of each of the different parts. This chapter also explains several tools that you can use to view the contents of management packs and to analyze what the management pack does. The last part of this chapter explains using groups in Microsoft System Center 2012 R2 Operations Manager.

## Understanding management packs

---

A management pack is what makes Operations Manager work. It defines what to discover, what to monitor, and how to monitor it. It also identifies what data should be collected and can define visual elements, such as dashboards and views, as well.

The product group that creates the product also makes the management packs, so you will have the combined knowledge of the people who created the product to assist you with monitoring your applications in the most recommended way. The Operations Manager product group is responsible for making the management packs work optimally in Operations Manager. This gives you the best of both worlds. Brian Wren has done an outstanding job writing the System Center Management Pack Authoring Guide, which you can find at <http://social.technet.microsoft.com/wiki/contents/articles/15251.system-center-management-pack-authoring-guide.aspx>. There is also an MSDN Channel9 series about management packs available at <http://channel9.msdn.com/Series/System-Center-2012-R2-Operations-Manager-Management-Packs> and a Microsoft Virtual Academy series available at <http://www.microsoftvirtualacademy.com/training-courses/system-center-2012-r2-operations-manager-management-pack>.

In general, it is easy to create new groups, monitors, rules, and overrides with the Operations Manager console. More advanced elements can be created with the authoring console, which was designed for Operations Manager 2007 R2. However, keep in mind that the authoring console works only with management packs that have a v1.0 XML schema—Operations Manager 2012 has a v2.0 schema. You can also use the System Center 2012 Visual Studio Authoring Extensions or a third-party tool, such as MP Author from Silect Software.

It is important to understand how Operations Manager works with classes and groups. It is also important to understand what targets are before you create a monitor, rule, or override. Also, note that the terminology differs depending on whether you are working in the Operations Manager console as an operator/engineer or with authoring tools as a developer. In the console, the terms *targets*, *instances*, and *properties* are used; in authoring tools, the same items are referred to as *classes*, *objects*, and *attributes*.

In the console, each instance that you can see in discovered inventory must be discovered by targeting the parts that make up the application you want to monitor with Operations Manager. Every instance is considered a representation of a target that shares the same properties (the details) and a common means of being monitored. For authoring, a deeper knowledge is necessary, and you should have a profound knowledge of the complex interrelationships between classes and objects.

## Singleton vs. non-singleton classes

Many attributes of a class dictate how it is used. This is explained in detail in the TechNet article available at <http://technet.microsoft.com/en-us/library/hh457564.aspx>. Two class types—singleton and non-singleton—dictate how class instances are discovered and whether they are managed by agents or by management servers of a certain resource pool.

The *singleton* class is automatically created (discovered) with no discovery rule required. There can be only one instance of a singleton class. A group in Operations Manager is an example of a singleton class. It has only one instance (the group object itself) and is created during configuration through the Create Group Wizard or automatically when a management pack is installed. There is always a single instance of a given group, and groups are managed by the management servers from the All Management Servers Resource Pool.

A *non-singleton* class, on the other hand, can be managed either by agents or by management servers from any resource pool. There can be any number of instances of a non-singleton class. The Windows Computer class is an example of a non-singleton class. There needs to be as many instances of this class as there are Windows-based computers to be monitored. In this particular example, this class is managed by agents. Each agent installed on a Windows-based computer creates an instance of the Windows Computer class and manages it locally.

## Workflow targets

A *workflow*, such as a discovery, rule, monitor, override, or task, has a certain target defined. This target dictates what instances a particular workflow will run on. For example, if you create a monitor that you need to run only on computers with the Domain Controller role installed, you select the Domain Controller role as the target for this monitor. By doing so, you ensure that this monitor will run only on domain controllers. The target also defines which agents the management pack with this monitor is distributed to. This is important to note because some management packs can have embedded resources like dynamic-link libraries (DLLs) or other kinds of files that are automatically copied to the target as well.

It is best practice to always choose as specific a class as possible to ensure that the management pack and its workflows are downloaded only on computers where they are really needed. For example, to monitor something that exists only on a computer running SQL Server, select the SQL Database Engine class instead of a generic class like Windows Computer.

Also, when you create new monitors or rules, it is best to use an existing class instead of creating a new one. This keeps the type space smaller, which is better for performance. However, to extend monitoring for an entire application that has no management pack available for download, it is best to create new classes that specifically describe the application model and how you intend to monitor its various parts. Even though fewer classes is better for keeping the instance space smaller, the classes you create is trivial compared to the number of workflows that run on an agent. The classes you choose also influence how parts are displayed in views, dashboards, and reports. Depending on your needs, you can distinguish some parts by application version, or you can use a generic class that incorporates all versions, and then later you can build version-specific monitoring logic directly in the monitors and rules you create.

## Authoring classes

When building a class model for your application, you start with an initial, or base, class that needs to get discovered so that afterwards all the higher level classes are discovered based on that. This ensures that the management pack is downloaded only on the agents where that application exists. This also ensures that all the workflows that belong to parts of this application run only on those agents. This is why it is best practice to choose a very specific class that already exists as the target for the discovery rule that will discover this base class. Another option is to target this discovery rule to a more generic class that is a seed discovery class. This ensures that the discovery rule (workflow) that runs to discover the initial class is super lightweight, which is good for performance, and, ideally, runs on a wide interval (for instance, every 24 hours).

**See also** For more information, see the article "MP Best Practice: Using the Seed Pattern for Easy Discovery" on the TechNet Wiki at <http://social.technet.microsoft.com/wiki/contents/articles/1208.mp-best-practice-using-the-seed-pattern-for-easy-discovery.aspx>.

When defining classes, you should not use properties that can change frequently. For example, if you want to monitor important folders of an application and you discover these folders as classes, you should not define folder size as a property because, often, folder size changes frequently, and every time the discovery rule for this folder class runs, there will be a new value for the folder size property. This will cause a re-discovery of that class (to update the properties), and this will cause a configuration update on the agent(s) where this class is hosted. A configuration update is a costly operation if it happens too often and can have a significant impact on performance. This scenario is called a configuration churn and should be avoided.

**See also** For more information, see the article "How to detect and troubleshoot frequent configuration changes in Operations Manager" on Microsoft Support at <http://support.microsoft.com/kb/2603913>. See also Kevin Holman's article "What is config churn?" on his blog at <http://blogs.technet.com/b/kevinholman/archive/2009/10/05/what-is-config-churn.aspx>.

## State change events

The biggest difference between rules and monitors is that monitors also define a state. This is helpful, but note that every time a monitor changes its state (for example, from Healthy to Critical), it also inserts a state change event in the database. These entries are stored in the StateChangeEvent table in the Operational Database. Because the state of the objects is also displayed when you view data in the Operations Manager console, this table is used frequently in various queries used to get data from the database. The larger this table is, the slower the console becomes. It is important that monitors are created in such a way that they don't change state very often, for example, every couple of minutes. A monitor with this behavior is too sensitive and most likely is not reflecting the actual state of the part it monitors.

Ideally, such a monitor should be redesigned. If redesign is not possible, the monitor should be tuned. Tuning means changing the way the monitor works via the available overrides. For example, changing the thresholds for the different states for threshold-based monitors makes them less sensitive to changes. Too many state change events not only affects console performance, it affects management server performance because state change events write a lot of data in the database. Even with the state change event storm feature of the management servers, which prevents a new state change from being written to the database if it is part of a storm of changes to the same monitor, state change events still impact performance. Monitors that are very sensitive and generate a lot of state change events are known as flip-flopping or noisy monitors.

**See also** For more information, see Kevin Holman's article "Tuning tip: Do you have monitors constantly flip flopping?" on his blog at <http://blogs.technet.com/b/kevinholman/archive/2009/12/21/tuning-tip-do-you-have-monitors-constantly-flip-flopping.aspx>.

On the subject of state change events and their impact, it is also important to consider maintenance mode. When you put an agent into maintenance mode, each of its monitors generates a state change event, changing from its current state to the Not Monitored state. In return, when an agent exits maintenance mode, each monitor it uses sends a state change event from the Not Monitored state to the Healthy state. This is called *monitor initialization* and happens each time a monitor starts working. This functionality is crucial to the calculation of the availability of each part being monitored. However, it generates a significant number of state changes, and, therefore, it is best to avoid implementing scenarios where a large number of agents are put into and pulled out of maintenance mode frequently, such as each night. For example, in environments where provisioning servers reset to their image install every night, all of the agents are put into and pulled out of maintenance mode while the servers are being reverted. In scenarios like this, if maintenance mode cannot be avoided, then a good approach

is to reduce the Database Grooming settings for state change event data as much as possible (default 7 days). As a matter of fact, it is a good idea to reduce Database Grooming settings for all data types as much as the business allows and instead rely mostly on historical data that is available in the Data Warehouse Database through dashboards and reports.

## Module cookdown

When designing management packs, one of the most important tools to use is the module cookdown feature. A workflow (monitor, rule, and so on) contains more modules than it needs to function. Cookdown is a feature that saves memory and CPU time by re-using already loaded modules instead of loading and initializing new instances of those modules. For more information, see <http://technet.microsoft.com/en-us/library/ff381335.aspx>.

## Type space

The *type space* is the total number of management packs, classes, relationships, resources, enumerations, monitor types, views, and other internal definitions that exist in the environment (the Operational Database). A copy of the type space is held in memory by each Data Access Service on each management server. Each time a new class, workflow, view, and so on is created, modified, or deleted in the console, the Data Access Service of each management server reloads the type space. The bigger the type space is, the longer it takes to reload. In large environments, this might significantly impact performance on the management servers until the reload is finished.

It is better to have more management packs separated by application and other criteria, such as one management pack containing the definitions of classes, relationships, and discovery rules and a separate management pack containing the monitors, rules, views, and so on, than to have a very big management pack. Even so, it is best practice to import only management packs that you need. Likewise, when you author new management packs, it is best to make them as light as possible and to create and import only what you need. In large environments where a lot of management pack authoring is done and a lot of management packs are imported, it can have a significant effect on performance and memory usage on the management servers. The bigger the type space becomes, the bigger the impact on performance is. Each agent is able to handle many instances, but the impact to performance could be severe if the management group is not able to calculate configuration. In general, you can expect an average of 50 to 100 discovered instances hosted by an agent, which results in about 50,000 to 100,000 discovered objects, to be handled by a management group in a 1,000-agent environment.

It is also important to understand the impact type space size might have when you use Windows PowerShell scripts that connect to the Data Access Service to perform different actions, such as custom maintenance, custom monitoring, automatic overrides, and so on. Usually, such scripts consume a large portion of the type space loaded into memory from the Data Access Service, and in some situations, these scripts can load up to almost the entire type space, depending on what the script does. For example, a rule might connect to the Data

Access Service to get the list of all monitors and then, based on some criteria, take some action either on the monitors, on the objects to which these are tied to, or maybe on the alerts these have generated. In such a scenario, you might end up loading the monitor types, classes, or other parts of the type space into the memory of the associated MonitoringHost.exe instance that is running the Windows PowerShell script. This potentially causes high CPU usage and definitely causes high memory usage of that process. This might not be important in smaller environments or where there aren't a lot of these scripts, but in large environments with many of these scripts, the impact can be significant.

## Authoring groups

Groups are also an important element of authoring. Groups are singleton classes that are hosted by the All Management Servers Resource Pool. This means that management of groups is split between the management servers of this resource pool. The members of a group are dynamically calculated by workflows called Group Calculation workflows.

Static groups (groups with explicit membership) are much better for performance than dynamic groups (groups containing dynamic membership calculation rules). However, dynamic groups are much more resource intensive when processed. The more groups you have, and, specifically, the more dynamic groups you have, the bigger the performance impact is on the management servers of the All Management Servers Resource Pool. It is best to avoid creating new dynamic groups and to instead rely on classes for targeting or other scenarios where the desired functionality can be achieved using different methods. When dynamic groups are needed—and of course these will be needed and are a very important part of monitoring—try to use the simplest dynamic membership rules possible. For example, if regular expressions for the criteria is not required, use simpler criteria, even if it is harder to build initially. Doing this benefits performance in the long run and is worth the extra effort.

## Group calculation interval

Another good method for optimizing your Operations Manager environment is using and tuning the group calculation interval. Typically, there are many custom groups used for scoping user role views and dashboards or for filtering notifications or overrides. Discovery rules for your groups can impact the performance of your environment because the queries create multiple read operations to the Operations Manager database. Adding many dynamic groups with complex criteria to Operations Manager can negatively impact the overall performance.

Group calculations occur every 30 seconds by default. If you work in a corporate environment with many Operations Manager groups, then you should increase this value. On the other hand, increasing the calculation interval can affect the group membership discovery (the addition or removal of members of the group in a timely fashion). You have to find the optimal value that suits your environment. You can change the group calculation interval in the registry of the management server in the key GroupCalcPollingIntervalMilliseconds.

## Sealed management packs

Sealing a management pack changes it from an .xml file to an .mp file, which is a binary representation of the management pack. When you seal a management pack, the file is digitally signed by the provider and the user knows that it hasn't been modified since then.

To upgrade a sealed management pack, the same key must be used or the upgrade will fail. The sealed or the unsealed version of a management pack can be added to a management group, but never at the same time. Sealed management packs have version control when an updated version of the management pack is imported into a management group. If the management pack is sealed, only a newer version of the same management pack can be imported and only if the newer version successfully passes the backward compatibility check. For unsealed management packs, the new version is always imported regardless of its compatibility and regardless of its version.

Keep in mind that a management pack can reference another management pack only if the management pack that is referenced is sealed. Therefore, any modification to a management pack cannot break other management packs that reference it. If you configure typical parts that are used by other management packs, such as groups or modules, you must seal the management pack.

## Summary of best practices

In summary, here is a list of the most important things to consider when working with management packs:

- Class properties you choose should change values as seldom as possible, close to never.
- Don't use Operations Manager for software inventory (System Center Configuration Manager is built to do that), and don't collect too many properties.
- Monitors should change their state as seldom as possible. They should not be too sensitive, and the related issue that is described in the alert should be resolved in a more permanent manner.
- The type space should be kept as small as possible. Import or create only what you need and delete what you do not use.
- Windows PowerShell scripts that connect to the Data Access Service should be kept to a minimum. At least try to develop them in a way that loads as few objects as possible by using selection criteria for the Operations Manager cmdlets.
- Don't over-use maintenance mode. If there is no way around it, reduce database grooming settings for state change events data.
- Targets for workflows should be as specific as possible. Use seed classes with lightweight discovery rules for custom application monitoring.
- Tune existing workflows using overrides. Disable unneeded workflows, adjust thresholds, set higher run intervals, and so on.

- Prefer static groups instead of dynamic groups, or at least try to use lightweight criteria for your dynamic groups.
- Change the group calculation interval when there are many groups in the Operations Manager environment.
- Configure before you customize. Determine if you can use an existing workflow for what you need instead of creating a new one.
- Classes, groups, modules, and so on should be in a sealed management pack so that they are not unexpectedly modified and so that they can be referenced by content in other management packs.

## Tools for analyzing management packs

---

This section provides an overview of tools that are available to help you analyze the contents of a management pack and determine what actions the management pack performs.

### Management Pack Viewer

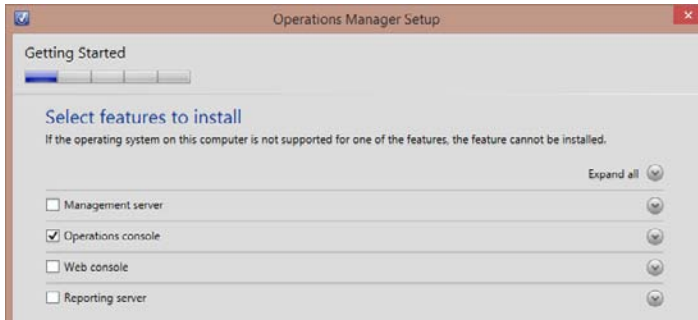
Management Pack Viewer (MP Viewer) was first developed by Boris Yanushpolsky and later updated for Operations Manager 2012 and management pack bundle (MPB) files by Daniele Muscetta. The download link for this tool is <http://blogs.msdn.com/b/dmuscett/archive/2012/02/19/boris-s-tools-updated.aspx>. No installation is required. This tool lets you view the contents of a management pack and export it to HTML or Microsoft Excel format for easier viewing and analyzing.

To use the MP Viewer tool, you need to have certain DLLs installed. For instance, if you download and start the application, you might get the following error message:

*System.IO.FileNotFoundException: Could not load file or assembly 'Microsoft.EnterpriseManagement.Core, Version=7.0.5000.0, Culture=neutral, PublicKeyToken=31bf3856ad364e35' or one of its dependencies. The system cannot find the file specified.*

To address this error, you can install Microsoft.EnterpriseManagement.Core.dll (and some more DLLs) into the assembly cache or you can install the Operations Manager console on the same system where you want to use MP Viewer. The latter method is the easiest. To install just the Operations Manager console, start the setup of Operations Manager from the installation media, and when asked to select the features to install, select only the operations console, as shown in Figure 2-1.





**FIGURE 2-1** Install the operations console to import the necessary DLL files into the assembly cache

To install the Operations Manager console, you also need to install the Microsoft Report Viewer 2012 runtime. The installation wizard informs you of this and provides the link to where you can download and install this prerequisite.

After installing the MP Viewer tool, double-click the executable file MPViewer.exe to run it. The tool prompts you to select a management pack file to open. This can be a management pack file (\*.MP), a management pack bundle file (\*.MPB), or an unsealed management pack file (\*.XML). For the remainder of this section, the System Center Management Pack for SQL Server version 6.5.1, which can be found at <http://www.microsoft.com/en-us/download/details.aspx?id=10631>, is used as an example.

To view the individual \*.MP files in this management pack, you need to download and install the management pack first. After installation, you will have all the files in C:\Program Files (x86)\System Center Management Packs\System Center Monitoring Pack for SQL Server\6.5.1.0. The SQL Server management pack is known to be one of the best and most extensive management packs, and so is the management pack guide (Word document) that describes the whole management pack, including monitoring scenarios, troubleshooting guidelines, security considerations, and more.

## The Discovery part of a management pack

Not all available management packs are divided into Discovery, Monitoring, and Presentation parts. If everything is in one management pack file, the following explanation is still valid. However, since dividing a management pack into these three different parts is best practice for building your own management packs, you should follow the example set by the SQL Server management pack.

When opening the Microsoft.SQLServer.2012.Discovery.mp, which is the management pack responsible for the discovery of SQL Server 2012 servers and instances, you first encounter the defined classes (if any), as shown in Figure 2-2. Classes are first in the list of types of items that can be defined in a management pack.

Name	Base Class	Is Abstract	Is Hosted	Is Singleton	Accessibility	Description
Database Critical Policy	SQL Server 2012 Custom User Policy	False	True	False	Public	Custom User F
Database Warning Policy	SQL Server 2012 Custom User Policy	False	True	False	Public	Custom User F
SQL Server 2012 Agent	SQL Agent	False	True	False	Public	Microsoft SQL
SQL Server 2012 Agent Job	SQL Agent Job	False	True	False	Public	All Microsoft Si
SQL Server 2012 Analysis Services	SQL Analysis Services	False	True	False	Public	An installation
SQL Server 2012 Computers	Microsoft.SystemCenter.ComputerGroup	False	False	True	Public	A group cortai
SQL Server 2012 Custom User Policy	SQL Component	True	False	False	Public	Custom User F
SQL Server 2012 DB	SQL Database	False	True	False	Public	Microsoft SQL
SQL Server 2012 DB Engine	SQL DB Engine	False	True	False	Public	An installation
SQL Server 2012 DB Engine Group	System Group	False	False	True	Public	A group cortai
SQL Server 2012 DB File	SQL DB File	False	True	False	Public	Microsoft SQL
SQL Server 2012 DB File Group	SQL DB File Group	False	True	False	Public	Microsoft SQL
SQL Server 2012 DB Log File	SQL DB File	False	True	False	Public	Microsoft SQL
SQL Server 2012 Distributor	SQL Distributor	False	True	False	Public	All Microsoft Si
SQL Server 2012 Installation Seed	Microsoft.Windows.ApplicationComponent	False	True	False	Public	It is a seed for
SQL Server 2012 Integration Services	SQL Integration Services	False	True	False	Public	An installation
SQL Server 2012 Publication	SQL Publication	False	True	False	Public	All Microsoft Si
SQL Server 2012 Publisher	SQL Publisher	False	True	False	Public	All Microsoft Si
SQL Server 2012 Replication Computers	Microsoft.SystemCenter.ComputerGroup	False	False	True	Public	A group cortai
SQL Server 2012 Reporting Services	SQL Reporting Services	False	True	False	Public	An installation
SQL Server 2012 Subscriber	SQL Subscriber	False	True	False	Public	All Microsoft Si
SQL Server 2012 Subscription	SQL Subscription	False	True	False	Public	All Microsoft Si

**FIGURE 2-2** The defined classes in a management pack are the first thing that is shown in MP Viewer

The Dependencies section lists all of the other management packs that this specific management pack relies on, for instance to extend a class that is defined in the other management pack. You cannot import this specific management pack if the dependent management packs, with the same version as defined in this management pack, have not been imported.

The Discoveries section shows all the discoveries defined in this management pack, together with their targets, whether they are enabled, their frequency, and more information. This is shown in Figure 2-3.

Name	Target	Enabled	Frequency	Remotable	Description
Discover Databases for a Database Engine	SQL Server 2012 DB Engine	True	14400	True	This object dis
Discover File Groups	SQL Server 2012 DB Engine	True		True	This object dis
Discover Files	SQL Server 2012 DB Engine	True		True	This object dis
Discover Replication Components	SQL Server 2012 DB Engine	False	14400	True	This object dis
Discover SQL Server 2012 Agent Jobs	SQL Server 2012 Agent	False	14400	True	This object dis
Discover SQL Server 2012 Analysis Services (Windows Server)	SQL Server 2012 Installation Seed	False	14400	True	This object dis
Discover SQL Server 2012 Database Engines (Windows Server)	SQL Server 2012 Installation Seed	True	14400	True	This object dis
Discover SQL Server 2012 Integration Services (Windows Server)	SQL Server 2012 Installation Seed	True	14400	True	This object dis
Discover SQL Server 2012 Reporting Services (Windows Server)	SQL Server 2012 Installation Seed	False	14400	True	This object dis
Discover SQL Server Agent for a DB Engine	SQL Server 2012 DB Engine	True	14400	True	This object dis
Populate Microsoft SQL Server 2012 Computer Group	SQL Server 2012 Computers	True		True	This object dis
Populate Microsoft SQL Server 2012 Instance Group	SQL Server 2012 DB Engine Group	True		True	This object dis
Populate Microsoft SQL Server 2012 Replication Computers Group	SQL Server 2012 Replication Computers	True		True	This object dis
SQL Server 2012 Database Custom User Policy Discovery	SQL Server 2012 DB Engine	False		True	Discovery of CI
SQL Server 2012 DB Installation Discovery Source	Microsoft.Windows.Server.Computer	True	14400	True	This object dis
SQL Server 2012 Replication Publications and Subscriptions Discovery Provider	SQL Server 2012 DB Engine	False	14400	True	This object dis

**FIGURE 2-3** The discoveries defined in Microsoft.SqlServer.2012.Discovery.mp

The Name column shows the name of the discovery. This name is also used in the Operations Manager console when you go to the Authoring pane, select Management Pack Objects, and then select Object Discoveries. The Target column shows which base class is used for this discovery. The least specific target, in this case Microsoft.Windows.Server.Computer, is the basic target. This means that this discovery runs on all the instances of this class; in other words, it runs on all of the computers running Windows Server. The subsequent discoveries will then use the seed discovery for more advanced discoveries.

**TIP** When creating your own management packs, you shouldn't use a broad class for all of your discoveries because it will negatively impact the performance of Operations Manager. Use a broad class only for the base discovery or the seed discovery.

The Enabled column indicates whether, by default, a certain discovery is enabled. In this case, the replication discovery is not enabled by default. This means that if you want to use this management pack for discovering your SQL Server replication configuration, you will have to enable this discovery. You can do that using overrides.

**See also** For best practices on how to configure overrides, please see the following Microsoft Knowledge Base article: <http://support.microsoft.com/kb/943239>.

The Frequency column shows how often a discovery will run. In this case, the discoveries run every 14,400 seconds, or every four hours. This is in line with best practices since you should not run discoveries too often. Running them too often can create config churn, which is explained further on Kevin Holman's blog <http://blogs.technet.com/b/kevinholman/archive/2009/10/05/what-is-config-churn.aspx> and in the Microsoft Knowledge Base article at <http://support.microsoft.com/kb/2603913>.

The Remotable column defines whether or not this workflow runs for agentless monitoring. The default value is True. The discovery management pack further defines the groups and relationships. Groups are used for targeting; relationships are defined between classes to indicate an association between a particular instance of one class and a particular instance of another. There are three types of relationships:

- Hosting relationship
- Containment relationship
- Reference relationship

**See also** More information on relationships can be found at <http://social.technet.microsoft.com/wiki/contents/articles/14256.operations-manager-management-pack-authoring-classes-and-relationships.aspx#Relationships>.

From the menu of the MP Viewer tool, you can unseal the currently loaded management pack as a plain XML file. You can also save the currently loaded management pack as an Excel or HTML file. This makes the tool valuable for further analyzing management packs. Using Excel, you can easily search and filter the contents of the sometimes very large management packs.

## The Monitoring part of a management pack

The Microsoft.SqlServer.2012.Monitoring.mp file contains the parts of a management pack responsible for the actual monitoring that Operations Manager performs. When opening this file using MP Viewer, notice that there is nothing in the Classes section. This is because the classes have been defined in the Discovery part of the management pack and don't need to be defined again in the Monitoring part.

Under Console Tasks, you can review the tasks that are defined in this management pack. These tasks appear on the right side of the Operations Manager console if you have selected the right targeted object in the console. To use these tasks, you need sufficient rights, and the application that is defined in the task must be installed on the same machine that is running the console.

Every management pack has dependencies on other management packs, so dependencies are indicated in the Monitoring part too, just as they are in the Discovery part of the management pack. However, the dependencies are different in each part of the management pack.

There are also linked reports, which are basically shortcuts to existing reports, but with different parameters. A linked report is derived from an existing report and retains the original's report definition. A linked report always inherits report layout and data source properties of the original report. All other properties and settings can be different from those of the original report, including security, parameters, location, subscriptions, and schedules. In this way, the SQL Server management pack can re-use existing reports from the Generic Report Library, and fill in the appropriate parameters to show SQL Server-related data.

Next are the Monitors: aggregated, dependency, and unit monitors. These are the building blocks of the monitoring that Operations Manager does, and monitors create the health state of monitored objects. Both aggregated and dependency monitors are rollup monitors that roll up the state of an object to the top level. Unit monitors are the base monitors. The Unit Monitors section includes many columns that define what a specific monitor does. The most important columns are the following:

- **Name** Shows the name of the monitor. This is the name that appears in the Operations Manager console in the Authoring pane, Management Pack Objects, Monitors.
- **Target** Identifies the actual monitoring objects that this monitor applies to.
- **Category** Reveals the type of monitoring, such as availability or performance.
- **Enabled** Indicates whether the monitor is enabled by default.
- **Generate Alert** Indicates whether the monitor generates an alert when the state of the monitor changes.
- **Alert Severity** Indicates whether the generated alert will be Informational, Warning, or Error. MatchMonitorHealth means that the monitor can generate different levels of alerts.
- **Alert Priority** Shows the priority as High, Normal, or Low. Combined with the three possible alert severities, this provides nine different alert levels.
- **Auto Resolve** Defines that the alert will be automatically closed when the health state returns to healthy (green).

The next part is Rules. Rules mainly collect performance counters and specified event IDs from event logs, but can also generate alerts. Monitors are state changing; they are what

makes your system go green, yellow, or red. Rules do not change the state of an object. This means that alerts coming from rules do not automatically go away; you should close those alerts manually. However, alerts coming from monitors should *not* be closed manually because that can cause the alert to not fire again when an alert condition is detected. Operations Manager will not generate a new alert for a monitor when there is an existing alert.

The most important columns are the following:

- **Name** Shows the name of the rule. This is name that appears in the Operations Manager console in the Authoring pane, Management Pack Objects, Rules.
- **Target** Identifies the actual monitoring objects that this rule applies to.
- **Category** Reveals the type of the rule, such as EventCollection or PerformanceCollection.
- **Enabled** Indicates whether the rule is enabled by default.
- **Counter Name** Shows the name of the collected performance counter.
- **Frequency** Indicates how often the performance collector is collected.
- **Event Log** Identifies the event log from which the event is collected.
- **Event ID** Shows the number of the event that is collected from the specified event log.
- **Event Source** Shows the event source of the event that is collected from the specified event log.
- **Generate Alert** Indicates whether it is an alert-generating rule.
- **Alert Severity** Indicates whether a generated alert will be Informational, Warning, or Error.
- **Alert Priority** Shows the priority as High, Normal, or Low. Combined with the three possible alert severities, this provides nine different alert levels.

The last part is Tasks. These are the tasks you can review in the Operations Manager console when you select one of the objects monitored by this management pack. A value of True in the Remotable column means that the task can run remotely against the monitored system.

Thresholds don't appear in the overview window of MP Viewer or in the Excel or HTML file when you export the management pack. To view them, select a rule or a monitor, and then click the Knowledge, Alert Description, or Raw XML tab (for monitors) or click the Knowledge or Raw XML tab (for rules) in the bottom right pane. When you select Raw XML, you will see the actual XML code that makes up the management pack. In this raw XML code, you can also see the thresholds, as is shown in Figure 2-4.

The screenshot shows the Management Pack Viewer 2012 interface. On the left is a tree view of classes. The main area displays a table with columns: Name, Target, Category, Enabled, Object Name, Counter Name, Frequency, and General. The 'Stolen Server Memory' item is highlighted. Below the table, the 'Raw XML' tab is selected, showing the XML configuration for the 'Stolen Server Memory' monitor. The XML includes settings for accessibility, priority, alert messages, parameters, operational states, and configuration values like interval, timeout, and threshold.

Name	Target	Category	Enabled	Object Name	Counter Name	Frequency	General
DB Log File Space	SQL Server 2012 DB Log File	PerformanceHealth	True				False
DB Space Percentage Change	SQL Server 2012 DB	PerformanceHealth	False				True
DB Total Space	SQL Server 2012 DB	PerformanceHealth	False				True
Destination Log Shipping	SQL Server 2012 DB	PerformanceHealth	True				True
Disk Read Latency	SQL Server 2012 DB	PerformanceHealth	False				True
Disk Write Latency	SQL Server 2012 DB	PerformanceHealth	False				True
Job Duration	SQL Server 2012 Agent Job	PerformanceCollection	True				False
Last Run Status	SQL Server 2012 Agent Job	AvailabilityHealth	True				False
Long Running Jobs	SQL Server 2012 Agent	PerformanceHealth	False				True
Page Life Expectancy	SQL Server 2012 DB Engine	PerformanceHealth	True				True
Page Verify Configuration	SQL Server 2012 DB	ConfigurationHealth	False				True
Recovery Model Configuration	SQL Server 2012 DB	ConfigurationHealth	False				True
Service Pack Compliance	SQL Server 2012 DB Engine	ConfigurationHealth	True				True
Service Principal Name Configuration Status	SQL Server 2012 DB Engine	AvailabilityHealth	True				True
Source Log Shipping	SQL Server 2012 DB	PerformanceHealth	True				True
SQL RA-Completion	SQL Server 2012 DB Engine	PerformanceHealth	False				True
SQL Server Agent Windows Service	SQL Server 2012 Agent	AvailabilityHealth	True				True
SQL Server Analysis Services Windows Service	SQL Server 2012 Analysis Services	AvailabilityHealth	True				True
SQL Server Full Text Search Service	SQL Server 2012 DB Engine	AvailabilityHealth	False				True
SQL Server Integration Services Windows Service	SQL Server 2012 Integration Services	AvailabilityHealth	True				True
SQL Server Reporting Services Windows Service	SQL Server 2012 Reporting Services	AvailabilityHealth	True				True
SQL Server Windows Service	SQL Server 2012 DB Engine	AvailabilityHealth	True				True
SQL Server Windows Service	SQL Server 2012 DB	AvailabilityHealth	True				False
SQL User Connections Performance	SQL Server 2012 DB Engine	PerformanceHealth	False				True
Stolen Server Memory	SQL Server 2012 DB Engine	PerformanceHealth	True				True
Thread Count	SQL Server 2012 DB Engine	PerformanceHealth	True				True
Transaction Log Free Space (%)	SQL Server 2012 DB	PerformanceHealth	False				True
Trustworthy Configuration	SQL Server 2012 DB	ConfigurationHealth	False				True

```

- <UnitMonitor ID="Microsoft.SQLServer.2012.DBEngine.StolenServerMemoryMonitor" Accessibility="Public" Enabled="true" System.Health.PerformanceState" Remotable="true" Priority="Normal" TypeID="Microsoft.SQLServer.2012.DBEngine" Category="PerformanceHealth" </Category>
- <AlertSettings AlertMessage="Microsoft.SQLServer.2012.DBEngine.StolenServerMemoryMonitor.AlertMessage">
  <AlertOnState>Error</AlertOnState>
  <AutoResolve>true</AutoResolve>
  <AlertPriority>Normal</AlertPriority>
  <AlertSeverity>Error</AlertSeverity>
- <AlertParameters>
  <AlertParameter1>$Target/Property[Type="SQL:Microsoft.SQLServer.ServerRole"]/InstanceName$</AlertParameter1>
  <AlertParameter2>$Target/Host/Property[Type="Windows:Microsoft.Windows.Computer"]/NetworkName$</AlertParameter2>
</AlertParameters>
- <OperationalStates>
  <OperationalState ID="Normal" MonitorTypeStateID="Normal" HealthState="Success" />
  <OperationalState ID="Error" MonitorTypeStateID="Error" HealthState="Error" />
</OperationalStates>
- <Configuration>
  <IntervalSeconds>300</IntervalSeconds>
  <TimeoutSeconds>200</TimeoutSeconds>
  <Threshold>70</Threshold>
  <NumSamples>6</NumSamples>
</Configuration>
</UnitMonitor>

```

FIGURE 2-4 The threshold can be viewed on the Raw XML tab

## The Presentation part of a management pack

Since Operations Manager 2012 SP1 Update Rollup (UR) 2 was released, it is considered best practice to separate the Presentation part in a separate management pack. Operations Manager 2012 SP1 UR2 (and Operations Manager 2012 R2) introduced new visualizations (widgets) that are not supported in earlier versions. As indicated in the dependencies of the Microsoft.SQLServer.2012.Presentation.mp file, Microsoft.SystemCenter.Visualization.Library is required.

The main part of the Presentation management pack is in Dashboards and Widgets, where you can find all the definitions of the types of items that make up a visualization in the management pack. The SQL Server management pack includes Microsoft.SQLServer.Generic.Dashboards.mp and a Microsoft.SQLServer.Generic.Presentation.mp part with more building blocks for visualizations.

## Other parts of a management pack

There might be other types of items that make up the complete management pack. The SQL Server management pack, for instance, includes a Library management pack, `Microsoft.SqlServer.Library.mp`. A Library management pack is a collection of common items across all the other management packs. In this case, it contains generic classes, discoveries, groups, and relationships along with common views and resources, such as bitmap images. Other management packs that need these resources reference the Library management pack (as a dependency) rather than defining all of these items themselves.

SQL Server 2008 Mirroring and SQL Server 2012 AlwaysOn have separate management packs in the same download package. To monitor these specific types of items, you need to import these additional management packs. Separate download packages are available for SQL Server Analysis Services and SQL Server Reporting (Native Mode).

**TIP** The MP Wiki found at <http://social.technet.microsoft.com/wiki/contents/articles/16174.microsoft-management-packs.aspx> contains all the Microsoft-provided management packs and their release dates. Check this page often to determine whether you have the latest version of the management packs installed.

## Override Explorer

Override Explorer is another tool in the same series as MP Viewer, created by Boris Yanushpolsky and later updated for Operations Manager 2012 by Daniele Muscetta. The download link for this tool is the same as for MP Viewer: <http://blogs.msdn.com/b/dmuscett/archive/2012/02/19/boris-s-tools-updated.aspx>.

Using Override Explorer, you can connect to your Operations Manager management group and get a list of all the overrides that have been applied in your environment. This includes the overrides that are imported with management packs by default, so even if you haven't deployed any overrides yourself, you will still find an extensive list of overrides. When the list of overrides is downloaded from your management group (this can take a while in a large environment), you can view them one by one or (just as for MP Viewer) you can export the whole list to XML or Excel format.

Apart from the export capability, three other functions in this tool add value. They can be accessed by selecting and right-clicking an override in the list. As shown in Figure 2-5, these three functions let you do the following:

- Easily change the target that an override applies to
- Move the override to a different management pack
- Delete the override

Parameter	New Value	Creation Time	Target	Is Enforced	Management Pack	Sealed MP	Description
GenerateAlert	False	1/23/2014 12:15:03 PM	All instances of type SQL DB Engine	False	DRH SQL overrides	False	

- Change Target
- Delete
- Move to different MP

**FIGURE 2-5** Override Explorer additional functions

You cannot change the target for an override using the Operations Manager console. Instead, you must note the changes you make in the specific override, delete it, and then re-create it with the new target. You would change the target when, for instance, you change the group that the override targets. Another method is to narrow the scope of the override from all instances to a specific group or instance.

In Operations Manager 2007 versions, overrides were stored in the default management pack by default. In later versions, you need to select in which management pack you want to store the override. Saving overrides in the default management pack was not a problem per se, but deleting the default management pack afterwards was a problem because of the many references to other management packs. These references get added every time you create an override to another management pack, and they do not get removed when you delete this override. This is also the reason why you should create a separate Override management pack for every Base management pack. You cannot use the Operations Manager console to move overrides to another management pack. But you can use Override Explorer to do so.

You can use the Operations Manager console to delete overrides; however, especially in large environments and when you need to delete a lot of overrides, the console might be too slow. Deleting them with Override Explorer is faster and easier.

You can use the Operations Manager console to create a large number of overrides, but it is a daunting task to do so. Instead, you can use another tool, Override Creator, for this task. Override Creator is available for download from [http://blogs.technet.com/b/scom\\_atlas/archive/2013/05/15/override-creator-for-2012.aspx](http://blogs.technet.com/b/scom_atlas/archive/2013/05/15/override-creator-for-2012.aspx). Similar to Override Explorer, to use this tool, you need to connect to your Operations Manager environment where the tool will download all the installed management pack contents and you can more conveniently create your overrides. Another method is to use Windows PowerShell.

## Workflow Analyzer explained

---

Tracing in Operations Manager is based on the Event Tracing for Windows technology (<http://msdn.microsoft.com/en-us/library/windows/desktop/bb968803%28v=vs.85%29.aspx>). Operations Manager comes with its own tracing providers for this functionality. The Workflow Analyzer tool can help with live tracing of a workflow (discovery, monitor, rule, or task) and can be accessed by installing the System Center Operations Manager 2007 R2 Authoring Resource Kit: <https://www.microsoft.com/en-us/download/details.aspx?id=18222>.



**NOTE** Even though the resource kit was created for the Operations Manager 2007 R2 version, it still works with Operations Manager 2012, Operations Manager 2012 SP1, and Operations Manager 2012 R2.

## Workflow Analyzer internals

Workflow Analyzer enables Event Tracing for Windows (ETW) tracing for Operations Manager on a specific workflow you choose from the tool's user interface. Workflow Analyzer uses an override to set the value for an internal property called TraceEnabled to TRUE for the chosen workflow using an override. The tool starts ETW tracing by giving the workflow a specific tracing provider with GUID c85ab4ed-7f0f-42c7-8421-995da9810fdd, which is called ModuleDebug. This starts tracing for all workflows that have the override and automatically traces all modules that the workflow will use.

## Using Workflow Analyzer

After installing the Operations Manager 2007 R2 Resource Kit, you can find the Workflow Analyzer executable file, called WFAalyzer.exe, in the install directory, which by default is %ProgramFiles%\System Center MP Authoring Console 2007. When you start Workflow Analyzer, you need to identify the Operations Manager management server it should connect to so that the tool can connect to the Data Access Service and get the list of existing workflows. In the same user interface, you can choose the agent (HealthService) on which Workflow Analyzer should run the workflow you will select later.

**NOTE** If you have never used tracing before, then you need to copy the OpsMgrTraceTMF.cab file from the Tools\tmf directory directly into its parent directory, Tools, because the tool expects the path from Operations Manager 2007 R2.

When the management server and agent are selected, a new window appears showing the list of available workflows. Search for the workflow you are interested in tracing, and then right-click it and select Trace. The workflow starts running, showing you the details of what its modules are doing, what output they are getting, and so on. This information can help you understand how the workflow works and provides a great deal of detail for troubleshooting issues.

## Running workflow tracing manually

In some situations, Workflow Analyzer crashes with a System.InvalidOperationException immediately after you select the management server and agent you want to run the trace on. This error cannot be worked around because it is a Windows Forms error coming from .NET 2.0, which can only be fixed by changing the code to adapt the tool for new versions of software.

However, knowing how Workflow Analyzer works, you can manually perform the steps required to enable tracing. The only difference between using Workflow Analyzer and running workflow tracing manually is that instead of monitoring the trace output live in a user interface, you review it in a log file that is created after the workflow runs.

To manually trace a workflow, first you need to create the override on the workflow that you need scoped to the specific management server or agent you are troubleshooting with some manual XML editing in the Operations Manager console. From the console, on the Authoring tab, search for the workflow (discovery, rule, or monitor) you are interested in and create an override as you would normally do. When creating the override, choose a random property, for example Enabled, and specifically set its value to TRUE. Save this override in a new management pack, which you can name something like WorkflowTracingMP. Next, export this management pack from the console. Open the management pack (XML) in Notepad or an XML editor of your choice. Find the override section—it should look something like the following example of doing this for the Logical Disk Free Space monitor:

```
<Overrides>

  <MonitorPropertyOverride
    ID="OverrideForMonitorMicrosoftWindowsServer62LogicalDiskFreeSpaceForContextMicrosoftWin
    dowsServer62LogicalDiskIefd92a8658b43ae9144736d68035079"
    Context="Windows!Microsoft.Windows.Server.6.2.LogicalDisk" ContextInstance="141e7f-
    3f1b-93bf-b7f7-f7ca6ce15fd6" Enforced="false"
    Monitor="Windows!Microsoft.Windows.Server.6.2.LogicalDisk.FreeSpace"
    Property="Enabled">

    <Value>true</Value>

  </MonitorPropertyOverride>

</Overrides>
```

Next, change the name of the overridden property from its current value to TraceEnabled. In this example, you would change Property="Enabled" to Property="TraceEnabled". Next, either change the management pack version (<Version>1.0.0.0</Version>) to a higher value in the XML, save it, and then import it back into Operations Manager via the console, or simply save it, delete the original override management pack from the console, and import the updated one.

The configuration of the agent on which you want to trace the workflow in question will be updated in a couple of minutes and thus will get the new management pack containing the override and apply it to the workflow. On this agent, locally, from an elevated command prompt, change the working directory to the Tools folder as follows:

```
cd /D "%ProgramFiles%\Microsoft Monitoring Agent\Agent\Tools"
```

Next, start the ModuleDebug tracing by running the following command, which outputs the data to a file that you have passed in the command line (%windir%\Logs\OpsMgrTrace\WorkflowTrace.etl):

```
TraceLogSM.exe -start WorkflowTrace -flag 0x1F -level 6 -f  
"%windir%\Logs\OpsMgrTrace\WorkflowTrace.etl" -b 64 -ft 10 -cir 999 -guid #c85ab4ed-  
7f0f-42c7-8421-995da9810fdd
```

At this point, you either know when the workflow will run the next time or you can look at the output file (%windir%\Logs\OpsMgrTrace\WorkflowTrace.etl) and see when it starts getting bigger. When the workflow has finished running, stop the trace from the command prompt with the following command:

```
TraceLogSM.exe -stop WorkflowTrace
```

To read the trace file, you need to format it as readable text. To do so, run the following command from the command prompt, providing the path to the directory where the ETL trace file was created via the *-Path*: argument:

```
FormatTracing.cmd -Path:"%windir%\Logs\OpsMgrTrace"
```

After formatting is finished, open the trace log that is located in the same directory but has the .log extension: %windir%\Logs\OpsMgrTrace\WorkflowTrace.log.

When you are done troubleshooting, it is best to delete either the override directly or the entire override management pack, thus deactivating the trace entirely.

## Best practices for creating and using groups

---

Groups are mainly used for scoping views, alert notifications, and reports. Therefore, groups are often created based on Windows Computer objects because this class hosts most of the relevant monitors for a Windows application.

If groups are created with extended authoring tools (or directly in XML using your preferred XML editor), they can and should be based on Windows Computer objects hosting special applications, for instance, a Windows Computer group that contains only Windows computers based on a discovered custom special application class. For notifications, the corresponding Health Service Watcher objects could be added to the group. This is necessary because you need the Health Service Watcher objects for Operations Manager self-monitoring alerts like Heartbeat Failures or Computer Not Reachable to be included too. Also remember to add cluster objects (if you need cluster-based alerts), which are not hosted by Windows Computer. You could do that in the same way as explained for the watcher objects mentioned above.

**See also** *More information about building and understanding groups is covered in Kevin Holman's blog at <http://blogs.technet.com/b/kevinholman/archive/2010/07/27/authoring-groups-from-simple-to-complex.aspx>.*

In addition, groups are useful for creating overrides. Group-based overrides can be much easier to manage than putting overrides on specific instances. Nevertheless, carefully plan which groups you need and the purpose each group should serve.

It's recommended that you save groups in the same dedicated, custom, unsealed override pack you use for your application because you cannot reference objects or classes in a different unsealed management pack. Sealing the group management pack is also possible, but this has disadvantages based on comfort and editing, and sometimes it breaks compatibility. Having all parts of an application together lets you easily maintain the application parts in one management pack without having an influence on other management packs.

It is also important to set useful naming convention rules for your groups and your management packs. For instance, a naming convention like GRP\_xxx for the group name makes finding groups in the console easy. Custom management packs can have the same name as the base management pack with “– Override” added to the name, so that you can search for “override” to find your override management packs. For your own management packs, you can add a short version of your company name to the beginning of the name of the management pack, for instance, CONTOSO\_xxx.

# Getting the most out of Operations Manager visualizations

This chapter provides an overview of the reports and dashboards that you can use to check your Microsoft System Center Operations Manager environment for excessive data generation and collection. It also explains how to add more visualizations to your Operations Manager environment, including reports coming from an external source or reports and dashboards that you create. These visualizations can be published to Microsoft SharePoint, and you can connect Microsoft Visio with Operations Manager to create live-updating dashboards.

## Tuning using the out-of-the-box reports

---

A default installation of Operations Manager includes an extensive number of reports. The following reports are most useful for tuning:

- Alerts
- Event Analysis
- Most Common Alerts and Most Common Events
- Data Volume by Management Pack
- Data Volume by Workflow and Instance


All of the reports can be accessed via the Operations Manager console on the Reporting pane. The overview on the Reporting pane is the same as the overview you see when you access SQL Server Reporting Services (SSRS) directly using your browser. The URL is <http://<servername>/Reports>, where <servername> is the name of your SSRS server. On this page is the link to Report Builder, which is the SSRS application that you can use to author your own reports or to edit existing reports. The first time you click the link, the Microsoft SQL Server Report Builder application is installed.

The two folders in the Reporting pane that contain the reports described later in this chapter are Microsoft Generic Report Library and Microsoft ODR Report Library. Depending on the management packs that you have imported, you might see additional folders, such as SQL Server 2012 (Monitoring). These folders contain the reports that come with product management packs, such as Microsoft Exchange, SQL Server, Active Directory, and more.

You can also import reports created by third parties. As always, be careful when you import something into your environment. Test it first in a test environment. Make sure the reports are useful and working correctly. Some well-known third-party reports that are free to download (although you might need to register first) include the following:

- **System Center Central's SCC Health Check reports**  
*<http://www.systemcentercentral.com/opsmgr-database-hygiene-scc-health-check-reports-management-pack-by-oskar-landman-pete-zerger/>*
- **Veeam Report Library for System Center** *<http://www.veeam.com/report-library-system-center.html>*
- **The Approved Operations Manager Health Check Dashboard** (providing a single page overview of where tuning is needed) *[www.approvedconsulting.com/downloads](http://www.approvedconsulting.com/downloads)*

Most of the available reports contain a date and time picker field to select the start and end time for the report. You can select a fixed date and time, or you can click the down arrow to set more advanced options. In the advanced options, you can select a time offset to make the report independent of a fixed date. If you schedule the report with a fixed date, you will get the same report every time you run it because the start and end time will always be the same. If, on the other hand, you choose Today – 1 week as the start time for the report and Today as the end time, you will always get the report for the previous seven days (see Figure 3-1).



The image shows a screenshot of a date and time picker interface. At the top, there is a 'From' field with a dropdown menu set to 'Today'. Below it, there is a 'To' field with a dropdown menu set to 'Today', a green checkmark, and a red X. A 'minus' dropdown is set to '1', and a unit dropdown is set to 'week(s)'. The date 'Saturday, November 29, 2014' is displayed, along with the location 'Stockholm, Vienna'.

**FIGURE 3-1** Advanced date and time picker

## Alert reports

In the Microsoft Generic Report Library, you will find the Alerts and the Most Common Alerts reports. The Alerts report shows alerts raised during the selected report duration and for given filter parameters for selected objects. The Most Common Alerts report shows the most common alerts raised during the selected report duration and for given filter parameters for selected objects.

### Alerts report

To run or schedule the report, select a time frame and a group or an object to run the report against. If you want to select all the Windows-based computers in your Operations Manager environment, click Add Group and type **Windows computers** in the field next to Contains. Click Search and the results will show a group called All Windows computers. Click Add to add this group, and click OK to close the object picker dialog box. Click Run at the top left of the report window to run the report. You can sort the results by alert name or repeat count by clicking the column header. To sort the report results from low to high, click the header once. To sort them from high to low, click the header again. For the tuning purposes discussed in this chapter, the report should be sorted beginning with the highest repeat count first.

In the Alerts report, the alert name appears in the first column. This alert name is the same as the one that appears in the Operations Manager console in the Alerts view. To get to the Alerts view from the report, expand one of the alerts and click Alert View. This opens the corresponding Alert view in the Operations Manager console. The next column shows the repeat count, which is how many times the alert was generated in the time frame specified in the report parameters. The Object column identifies the Windows-based computer where the alert was generated. In the expanded view of the alert, you can click Alert Detail Report to show all occurrences of the alert during the specified time frame (see Figure 3-2).

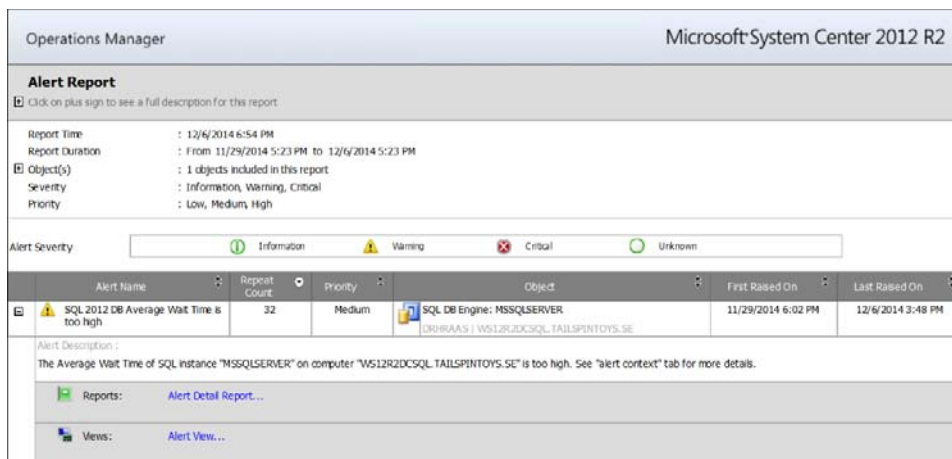
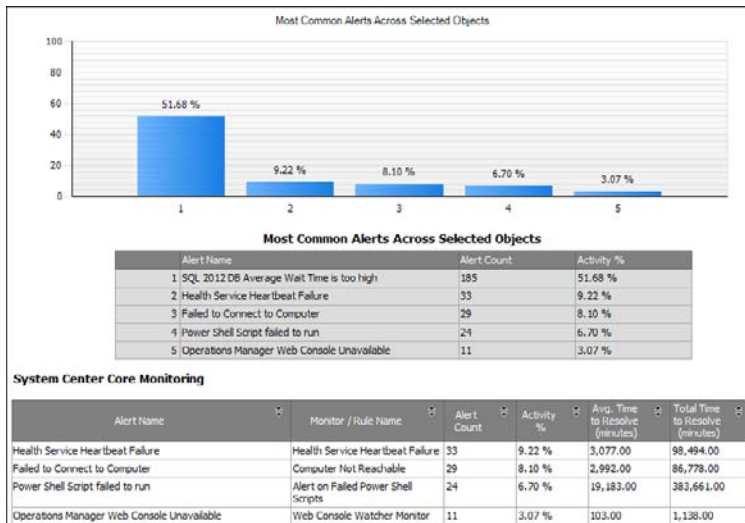


FIGURE 3-2 A sample Alerts report with an expanded view on one alert

## Most Common Alerts report

The Most Common Alerts report differs from the Alerts report because it shows the alerts per management pack whereas the Alerts report shows the alerts for a specified server. By default, the Most Common Alerts report shows the alerts coming from all the installed management packs. You can deselect management packs if you don't want to include them in the report. Since the report parameter is based on management packs, you don't need to specify the monitored systems it applies to.

When the report runs, it shows the most common alerts from the selected management packs. It also provides a drill down of the alerts per each management pack that is part of the top N most generated alerts. N is the number of most common alerts that you have chosen to show. By default, the report lists the top five most commonly generated alerts. You can change this in the report parameters. By scrolling down in the report, you can see how many alerts are coming from a specified management pack and which monitor or rule was responsible for generating the alert. The report details also show how long it took to resolve the alert on average and in total (see Figure 3-3).



**FIGURE 3-3** A sample Most Common Alerts report



## Event reports

In the Microsoft Generic Report Library, you will also find the Event Analysis and the Most Common Events reports. The Event Analysis report shows a table of events and a count by server filtered by all entered parameters. The Most Common Events report shows the most common events raised during the selected report duration and for given filter parameters for selected objects.

### Event Analysis report

The Event Analysis report is similar to the Alerts report, only it shows an analysis of the events that were collected rather than the alerts. To run this report, select the object to show collected events for, such as All Windows Computers, a more specific group of computers, or a specific object. You can filter the report by the source event log you want to see the collected events from, and by type, category, and event ID. You can sort report results by repeat count to view the most collected event IDs. In the details of each event, you have the option to open the Event view in the Operations Manager console or the Alert view for the object from which the event was collected. A sample Event Analysis report is shown in Figure 3-4.

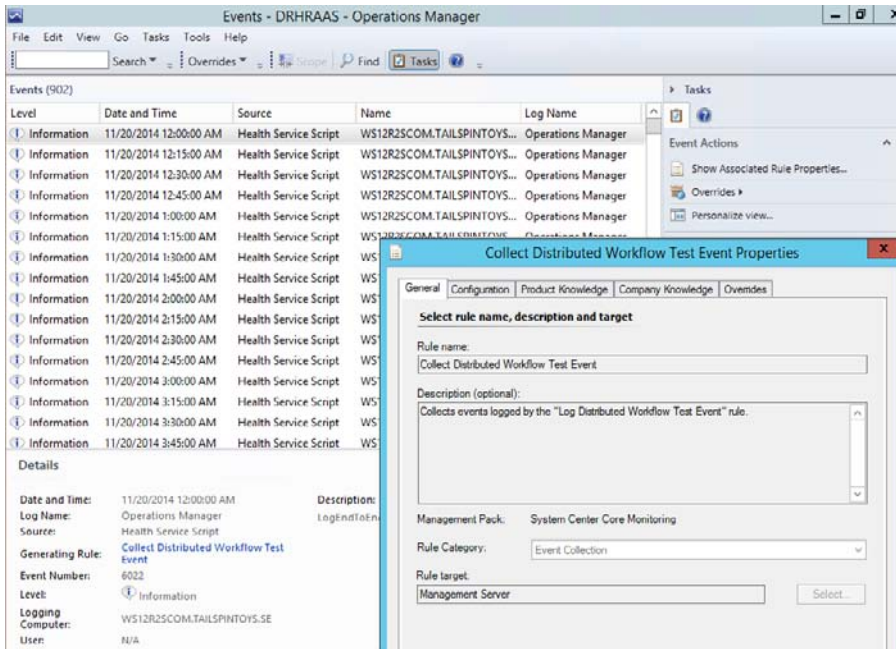
Object	Event Type	Time	Source	Repeat Count	Event ID	Rule
Health Service: WS12R25COM.TAILSPINTOYS.SE (DR RAAS) WS12R25COM.TAILSPINTOYS.SE	Information	11/29/2014 7:45:00 PM	Health Service Script	192	6022	Collect Distributed Workflow Test Event

Sample Description:  
LogEndToEndEvent.is : This event is logged to the Windows Event Log periodically to test an event collection.

Views:  
[Event view...](#)  
[Alert view...](#)

**FIGURE 3-4** A sample Event Analysis report with the detailed view for one event ID

If you click the Event view link, the report opens the specified Event view in the Operations Manager console. On the right side of the window, you can see additional tasks, navigation options, and event actions. Click Show Associated Rule Properties to open the properties of the rule that is collecting the specified event (see Figure 3-5). If necessary, you can create overrides directly in this dialog box, which makes it easier to tune the event collection rule since you don't have to search for it in the Authoring pane under Management Pack Objects and then Rules.



**FIGURE 3-5** The Event view in the Operations Manager console, opened from the report

## Most Common Events report

Similar to the Event Analysis report, the Most Common Events report shows the events per event log, source, and type of event. You need to specify which objects, such as All Windows Computers, you want to run the report for. You will get an overview of the most collected events according to the filters that you specified. For each event, you can also see a sample description of the event with the specified event ID in the details (see Figure 3-6). The option to go directly to the Event view in the Operations Manager console is not available in this report.

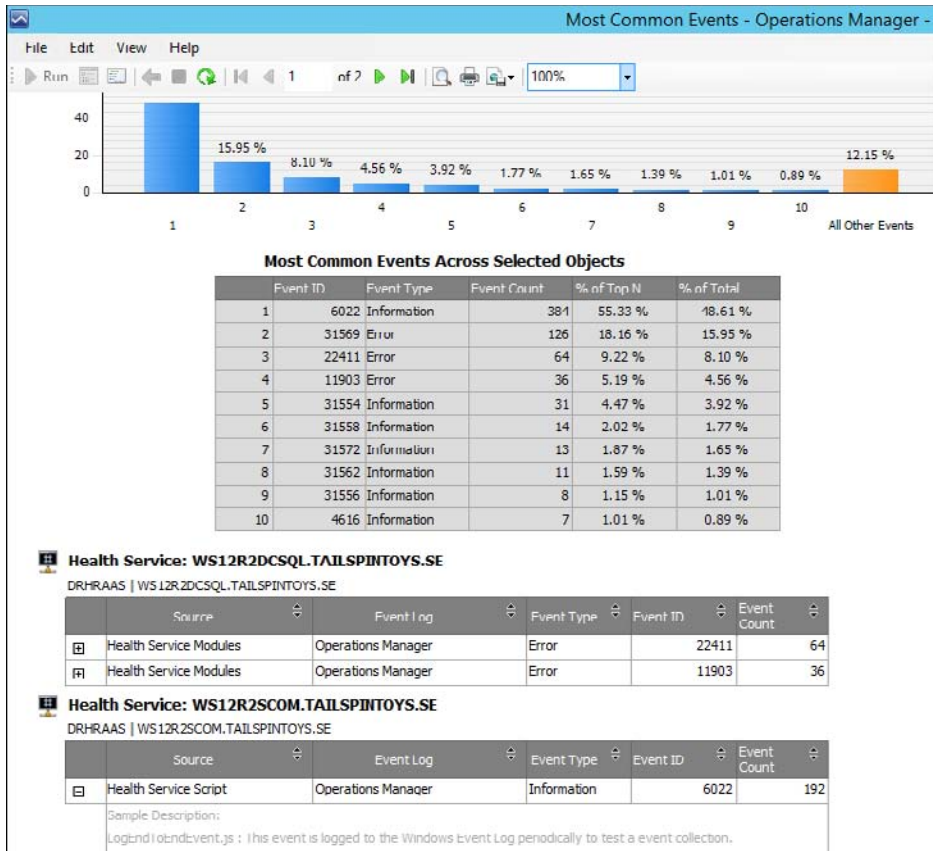


FIGURE 3-6 A sample Most Common Events report

## Data Volume reports

The Data Volume by Workflow and Instance and the Data Volume by Management Pack reports can be accessed in the Operations Manager console, in the Reporting pane, under System Center Core Monitoring Reports.

### Data Volume by Workflow and Instance report

The Data Volume by Management Pack report is a useful overview of the management packs generating the most data. In the report results, you can drill down by clicking each underlined link to generate a new report, tailored to the link you selected, using the Data Volume by Workflow and Instance report.

The Data Volume by Workflow and Instance report compiles information on the volume of data generated, broken down by workflows (discoveries, rules, monitors, and so on) as well as by instances. There are two ways to access this report: drilling down in the Data Volume by Management Pack report or running the report directly.

When running the report directly, you can select the management pack to show data from and the data type: discovery, alerts, events, performance, or state change data. By default, the report generates data for all data types and all installed management packs. This gives you another view of the overall data contained in your Operations Manager Operational database.

## Data Volume by Management Pack report


Figure 3-7 shows a sample Data Volume by Management Pack report.

Operations Manager		Microsoft System Center 2012 R2						
Data Volume by Management Pack								
Click on plus sign to see a full description for this report								
Time Span : 11/7/2014 11:06 AM - 12/7/2014 11:06 AM								
Aggregation : Daily								
Management Groups								
Counts by Management Pack								
Management Pack	% of total Data Vol.	Trend	Discovery Data	Alert Count	Perf Count	Event Count	State Changes	
<a href="#">SQL Server 2012 (Monitoring)</a>	88.12	↓	0	193	936209	0	642	
<a href="#">Data Warehouse Internal Library</a>	13.93	↓	0	12	186868	1422	338	
<a href="#">System Center Core Monitoring</a>	6.25	↓	0	126	79734	2731	9457	
<a href="#">Windows Server 2012 Operating System (Monitoring)</a>	3.54	↓	0	2	38481	178	1	
<a href="#">Microsoft Windows Server 2012 Internet Information Services (IIS)</a>	1.99	↓	46	2	27255	0	62	
<a href="#">Windows Server 2012 R2 Operating System (Monitoring)</a>	1.78	↓	0	0	24516	24	0	
<a href="#">Health Library</a>	1.58	↓	0	0	0	0	21251	
<a href="#">Windows Service Library</a>	0.87	↓	0	0	11904	0	0	
<a href="#">System Center Core Library</a>	0.58	↓	102	0	0	0	2760	
<a href="#">Microsoft System Center Operations Manager Summary Dashboard</a>	0.56	↓	0	0	7714	0	0	
<b>Grand Total</b>	<b>99.20</b>	<b>↓</b>	<b>239</b>	<b>335</b>	<b>1322781</b>	<b>4355</b>	<b>37080</b>	

**FIGURE 3-7** A sample Data Volume by Management Pack report

From the sample report shown in Figure 3-7, it is obvious that there is tuning to be done on the SQL Server 2012 (Monitoring) management pack. In this specific Operations Manager environment, 68 percent of all data in the Operations Manager database is coming from this management pack. You can click the name of a management pack to generate a more detailed report tailored to this specific management pack, using the Data Volume by Workflow and Instance report.

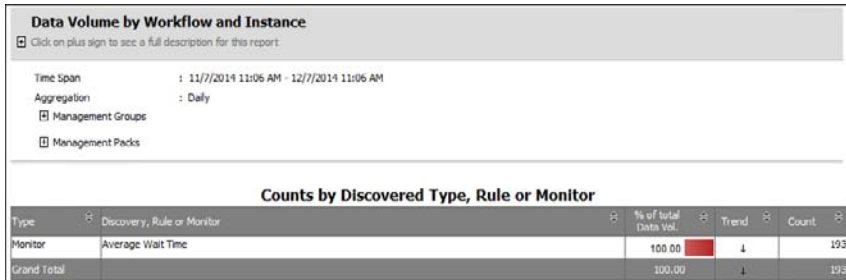
The best thing about the Data Volume by Management Pack report is that you can see on one page where tuning will be most useful. You can click any link to generate a drill-down report. When you're done with the drill-down report, you can return to the parent report by clicking the blue arrow on top of the child report, as is shown in Figure 3-8.



**FIGURE 3-8** Return to the parent report by clicking the blue arrow at the top menu in the child report

For example, tuning discovery data would not provide much benefit in this specific case. You can sort by value of the discovery data by clicking the up and down arrows next to the column title. If you encounter a problem with continuous rediscovery of objects, you can follow the guidance in the Microsoft Knowledge Base article at <http://support.microsoft.com/kb/2603913> and in Kevin Holman's blog about config churn at <http://blogs.technet.com/b/kevinholman/archive/2009/10/05/what-is-config-churn.aspx>.

In this case, there aren't many alerts. Clicking the number of alerts generated by the SQL Server 2012 (Monitoring) management pack (193 in this case) generates a new report that shows which rules or monitors in the management pack are generating these alerts. As shown in Figure 3-9, in this example, the Average Wait Time monitor is the culprit—all 193 alerts come from that monitor.



**FIGURE 3-9** A drill-down report showing the monitors and rules generating the alerts from the SQL Server 2012 (Monitoring) management pack

Based on the numbers in the parent report, it is obvious and expected that most data comes from performance counter collections. Since this is one of the main purposes of Operations Manager, this is nothing to worry about, per se. However, you need to make sure that you are collecting only the performance counters that you actually need. For example, Figure 3-10 shows the top performance counter collections from the SQL Server 2012 (Monitoring) management pack.

Counts by Discovered Type, Rule or Monitor				
Type	Discovery, Rule or Monitor	% of total Data Vol.	Trend	Count
Rule	MSSQL_2012: Collect DB Disk Read Latency (ms)	9.87	↓	92412
Rule	MSSQL_2012: Collect DB Disk Write Latency (ms)	9.87	↓	92412
Rule	MSSQL_2012: Collect DB Active Connections count	3.37	↓	31597
Rule	MSSQL_2012: Collect DB Active Requests count	3.37	↓	31595
Rule	MSSQL_2012: Collect DB Active Sessions count	3.37	↓	31595
Rule	MSSQL_2012: Collect Database Allocated Size (MB)	3.37	↓	31536
Rule	MSSQL_2012: Collect Database Total Free Space (MB)	3.37	↓	31536
Rule	MSSQL_2012: Collect Database Total Free Space (%)	3.37	↓	31536
Rule	MSSQL_2012: Collect DB Used Space (MB)	3.37	↓	31536
Rule	MSSQL_2012: Collect Transaction Log Free Space (%)	3.37	↓	31536
Grand Total		96.70	↓	437291

**FIGURE 3-10** Top performance counter collections from the SQL Server 2012 (Monitoring) management pack

It is always good to monitor free space. The Windows Server Operating System management pack also checks free space, but it is free disk space visible to the operating system. The performance counter collections shown in Figure 3-10 gather the free space within the actual database as seen by SQL Server.

For an example how this environment can be tuned, assume that it is not necessary to collect disk read and write latency, maybe because there is another management pack specifically designed to measure disk latency. In this case, it would be possible to override the collection rule and remove almost 20 percent of the collected performance counter data for this management pack (almost 10 percent multiplied by 2).

Apart from alerts and events and performance counter collections, continuous state changes can also have a big influence on the performance of your Operations Manager environment. For instance, when a monitored object changes from green to red, it changes state. Only monitors (not rules) can change the state of an object. Sometimes, the object changes state so fast that you don't see any alerts in the Operations Manager console. The Data Volume by Management Pack report shows whether there is a continuous state change problem by showing large numbers in the State Changes column. If this problem is present, try to determine why the object is changing state so frequently.

## Using Health Explorer

You can identify an affected object by drilling down in a report, and, using the Operations Manager console, you can view the affected object in Health Explorer. Within Health Explorer is a tab showing the state changes. On this tab, you can see the threshold being breached (see Figure 3-11). This can help you determine the root cause of a problem.

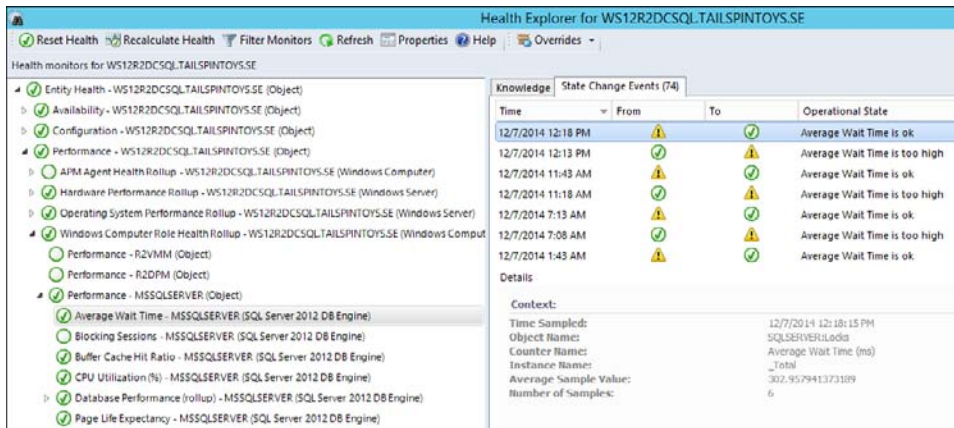


FIGURE 3-11 Health Explorer showing the state change events

**See also** On Kevin Holman's blog, you can find a further explanation about state changes and how to clean out old state changes from the Operations Manager Operational database. See <http://blogs.technet.com/b/kevinholman/archive/2009/12/21/tuning-tip-do-you-have-monitors-constantly-flip-flopping.aspx>.

## Publishing visualizations

You can publish reports and dashboards you create so that other people can easily access them. Publishing reports saves everyone time. You can also distribute reports automatically by email. Reports can be stored on a file share, where they can be collected and saved for later use, exported, scheduled, or saved to your favorites or a management pack, and you can print them, as shown in Figure 3-12.




FIGURE 3-12 Actions to take on a report

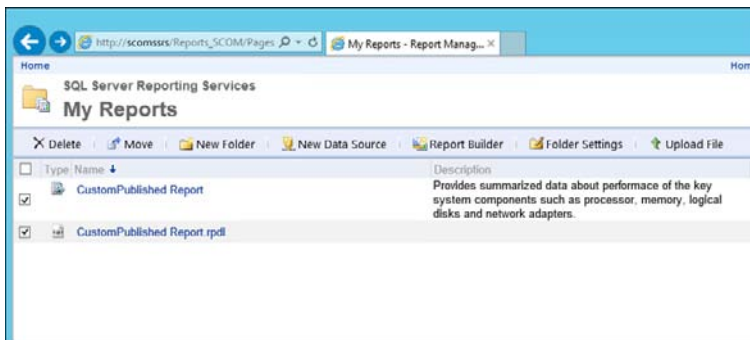


## Publishing and scheduling reports

After you generate a report with the necessary and correct parameters, you can save it for later use. Saving the report means you don't have to re-enter all the parameters since they are saved with the report.

You can save a report with your favorites, but that means only you can use and see it. Favorite reports are visible in the Operations Manager console, in the Reporting pane, under Favorite Reports. You can also publish reports. Publishing means you save the report in its current state, with all the parameters you selected, so that everyone else can access this report without having to select the parameters. This is also why it is a good idea to use a relative date and time for the report, as explained earlier. You can do this by selecting Advanced in the date and time picker and setting an offset date and time instead of a fixed date and time.

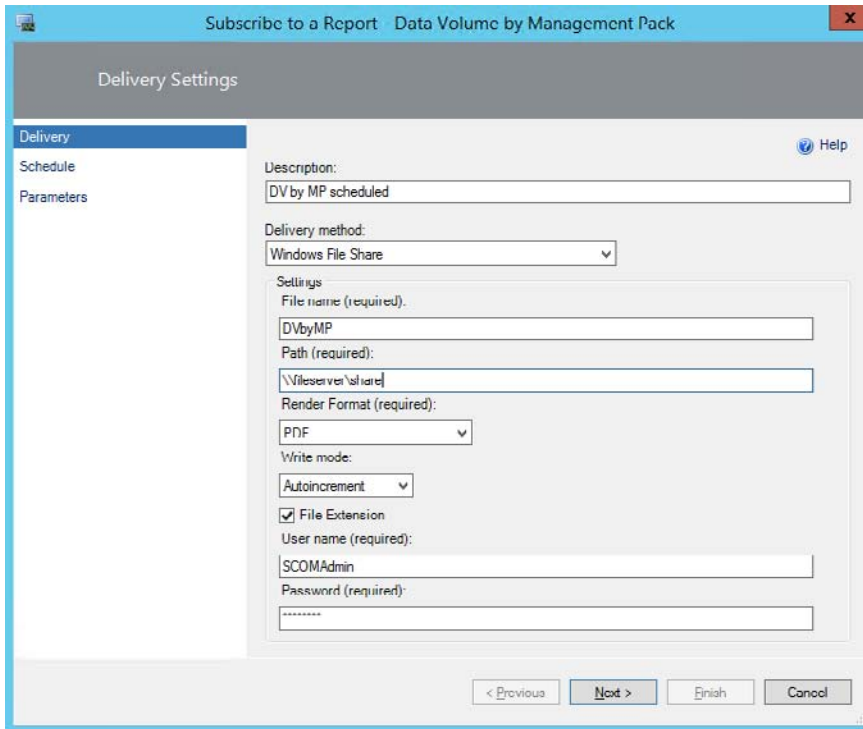
Published reports appear in the Operations Manager console, in the Reporting pane, under Authored Reports. However, they do not appear in the console for other users because this folder behaves like a personal profile, the same as the My Workspace folder in the Operations Manager console. This means that other user roles cannot access those reports by default. The advantage of having the reports in the Authored Reports folder is that you can move them easily in SQL Server Reporting Services (SSRS). To move the reports, open the SQL Server Reporting Services webpage, for instance <http://reportserver/reports/>, and create a new custom folder for the custom reports you want to share. Go to the My Reports folder, select the Show Details view, and move both files (the report .rdl file and the .rpd file) to the newly created custom folder by clicking Move (see Figure 3-13).



**FIGURE 3-13** Move custom published reports from the My Reports folder in SSRS to a custom folder



You can also schedule the report to run automatically on a regular basis. This auto-generates a new report according to the set schedule and it saves to a file. Figure 3-14 shows the options for scheduling a report.



**FIGURE 3-14** Scheduling a report

Another option is to email the scheduled report. Before you can schedule a report for email delivery, you must configure the email settings in the report server using the Reporting Server Configuration Manager. This process is explained at <http://technet.microsoft.com/en-us/library/hh212769.aspx>.

## Publishing to SharePoint

Reports and dashboards are accessible through the Operations Manager console. However, sometimes it is better to publish the reports outside of the Operations Manager environment, for instance so that managers can quickly review how things are going without having to log in to the Operations Manager environment.

## Publishing reports to SharePoint

Using the reports outside of the Operations Manager console, based on pre-populated parameters (for example, a published report) is possible through the `http://<servername>/Reportserver` URL. The URLs published on this SSRS webpage need to render without the smart header (the parameter control part). A URL report viewed on the SSRS webpage looks similar to the one shown in Figure 3-15.

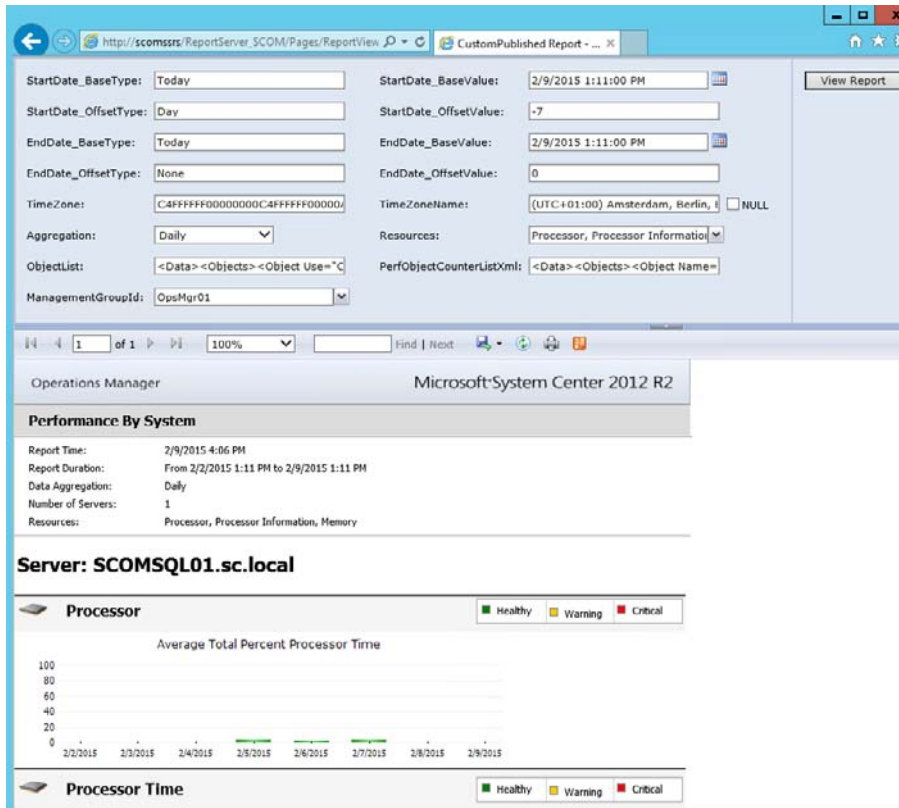


FIGURE 3-15 Rendering a report from the SSRS report server URL

You can copy the URL to a SharePoint Report Viewer web part where it can be easily administered using SSRS and SharePoint security. You can modify the URL to show or hide the toolbar or parameters by adding additional control options in the URL text. For example, you could hide the toolbar by editing the URL to include `&rc:toolbar=false` as shown in Figure 3-16, which shows the same report that's in Figure 3-15, but without the toolbar.

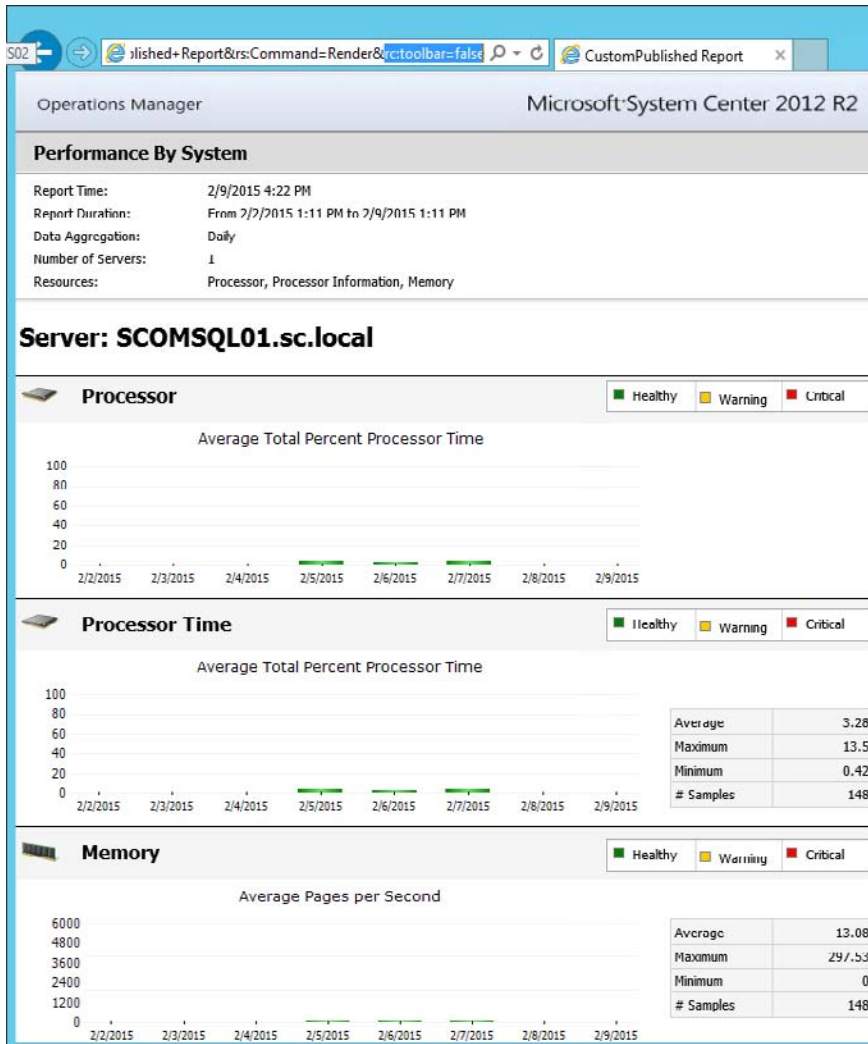


FIGURE 3-16 Rendering a report without the toolbar by updating the URL

Other URL examples to customize the appearance of reports include the following:

- **&rc:Parameters=collapsed** Parameter in the URL
- **&rc:Parameters=true** Shows the parameter bar (default)
- **&rc:Parameters=false** Hides the parameters
- **&rc:toolbar=true** Shows the toolbar

## Publishing dashboards to SharePoint

Update Rollup 2 for Operations Manager 2012 R2 and Update Rollup 6 for Operations Manager 2012 SP1 introduced newer and richer dashboard widgets. The list of new widgets now available are described at <http://social.technet.microsoft.com/wiki/contents/articles/24133.operations-manager-dashboard-widgets.aspx> and at <http://blogs.technet.com/b/momteam/archive/2014/04/24/new-widgets-and-dashboard.aspx>. A video with demos can be found at Channel 9 at <http://channel9.msdn.com/Events/TechEd/NorthAmerica/2014/DCIM-B329>.

A well-configured widget allows you to see the availability and performance metrics of the applications in your environment, and with SharePoint web parts, this view can be transferred out of Operations Manager. Typically, these SharePoint web parts are useful for providing holistic operational state views to everyone in your organization who does not have an Operations Manager role. Operations Manager includes a SharePoint web part that displays selected dashboards from the web console. The SharePoint farm must be running SharePoint 2013, SharePoint Server 2010 Standard, SharePoint Server 2010 Enterprise, or SharePoint Foundation 2010, and the procedure to configure it is described at <https://technet.microsoft.com/en-us/library/hh212924.aspx>.

## Using Visio with Operations Manager

To show health states of Operations Manager, you can also use the Visio 2010 add-in as an intermediate integration solution. You can export complex distributed applications (such as for your infrastructure applications or custom business services) or create diagrams directly in Visio and connect them with live health state views of Operations Manager. The add-in allows filtering of data, automatic generation of diagrams, and an automatic refresh of the associated data.

## Installing and using the Visio 2010 add-in with SharePoint

You can download the Visio 2010 extensions from <http://www.microsoft.com/en-us/download/details.aspx?displaylang=en&id=29268>. Although the add-in was designed for Visio 2010 and SharePoint 2010, it has been shown to work with Visio 2013 and SharePoint 2013. However, since it has not officially been tested by the Microsoft product group, its use with the 2013 versions is not officially supported.

You also need Visual Studio 2010 Tools for Office Runtime, which is required to run Microsoft Office-based solutions built using Microsoft Visual Studio 2010, 2012, and 2013 (<https://www.microsoft.com/en-us/download/details.aspx?id=44074>). In addition, you need Microsoft Visio Professional or Premium edition, the Operations Manager console, and Microsoft .NET Framework 4.0 to install the Visio 2010 add-in.

After you install the add-in and open Visio, a new tab for Operations Manager appears (see Figure 3-17). Click this tab to start configuring, adding your infrastructure and data links to the represented Operations Manager objects.




**FIGURE 3-17** The Operations Manager tab in Visio 2013

The Visio 2010 add-in and SharePoint 2010 Visio Services Data Provider offers the following features:


- Live health state information from distributed applications exported from Operations Manager
- Automatic updates of the current state information published in Visio documents or in hosted SharePoint document libraries
- Possibility to create new Visio documents with your available stencils and to link live data to the managed objects
- Option to create summary displays reflecting your company's infrastructure or process pictures in Visio documents and ability to transfer these documents to SharePoint with automatically refreshed health states
- Predefined data graphics to update shapes with colored health states representing the overall status of your services

A sample Visio drawing with objects connected directly to Operations Manager is shown in Figure 3-18. (More complex drawings can, of course, be made—this one simply shows the possibility.)



**FIGURE 3-18** A sample Visio drawing with some basic shapes showing the live connection with Operations Manager

Tim McFadden has developed a good description of how to install the SharePoint 2013 Integration extensions and use the Visio Services Data Provider and Visio Web Access to add the web part you created in a Visio 2013 document (saved as a Visio 2010 Web Drawing with extension \*.vdw). Tim's explanation can be found at <http://www.scom2k7.com/installing-scom-2012-visio-dashboards-in-sharepoint-2013/>. Figure 3-19 shows how to save a Visio 2013 created Web Drawing as a Visio 2010 Web Drawing (\*.vdw).



**FIGURE 3-19** Saving a Visio 2013 Web Drawing as a Visio 2010 Web Drawing

# Troubleshooting your Operations Manager environment

This chapter provides information about how Microsoft System Center Operations Manager works and how to troubleshoot general problems that might appear.

## Understanding the HealthService process

---

Understanding how Operations Manager works with respect to the HealthService process and the important part of the database schema is an essential step in troubleshooting scenarios. The HealthService process is a Windows service whose name depends on the version of Operations Manager you are using:

- **System Center Management** This is the HealthService process name used in Operations Manager 2012 and Operations Manager 2012 Service Pack 1.
- **Microsoft Monitoring Agent (MMA)** This is the HealthService process name used for Operations Manager 2012 R2 or the stand-alone Microsoft Monitoring Agent (MMA) used for Azure Operational Insights (formerly System Center Advisor) with or without an Operations Manager environment.

On management servers only, two other available services are involved:

- **System Center Management Configuration** This is used only for updating the configuration by loading it from the database at the request of HealthService.
- **System Center Data Access** This acts like a software development kit (SDK) and is used by client processes (the Operations Manager Console, Windows PowerShell, System Center Orchestrator, and so on) to interact with the infrastructure.

The HealthService also uses MOMPerfSnapshotHelper.exe to natively read performance counters from the Windows operating system.

Essentially, the main part of the HealthService process, the HealthServiceExecutive, manages all the worker parts, also known as managers. The other important part the HealthService

process uses is called the MOMConnector. The HealthService uses one or more external process called MonitoringHost to run the actual workflows.

The following is a general list of the most important managers the HealthService uses:

- PersistenceManager
- ConnectorManager
- ConfigurationManager
- ExecutionManager
- HealthManager
- SecureStorageManager
- PoolManager
- DiscoveryManager
- DataPublishManager
- JobManager

The HealthService utilizes a local cache folder called Health Service State, which is located in the installation folder under the Server or Agent sub-folder, depending on whether you are looking on an Agent or a Management/Gateway server.

Another important part HealthService uses is the local (cache) database, located under \Health Service State\Health Service Store\HealthServiceStore.edb. This database is a Windows part called the *Extensible Storage Engine (ESE)*. More information about the ESE can be found at [https://msdn.microsoft.com/en-us/library/gg269259\(v=exch.10\).aspx](https://msdn.microsoft.com/en-us/library/gg269259(v=exch.10).aspx).

Actual workflow execution is done in the MonitoringHost processes, which can use various available modules. These can be event log reader modules, WMI interaction modules, database write modules, OLE DB modules, log file reader modules, and so on. Custom modules can also be developed and used in custom management packs as needed. More information about implementing managed modules can be found at <https://msdn.microsoft.com/en-us/library/hh769912.aspx>.

## MOMConnector

MOMConnector is a standalone Dynamic Link Library (DLL) that the HealthService uses to receive and transfer data on TCP/IP from and to other Health Services. This is the part used by an Agent (child) HealthService to send and receive data from its Management Server (parent) HealthService.

## HealthServiceExecutive

HealthServiceExecutive is responsible for starting/stopping/managing its own sub-parts, also known as *managers*. HealthServiceExecutive can start, stop, suspend, and resume managers, while also managing any MonitoringHost process that needs to be started and managed.



## PersistenceManager

PersistenceManager is used by other managers to read and write data to the local ESE database.

## ConnectorManager

ConnectorManager is used by the HealthService to send and receive outside data through a connector. The connector used in Operations Manager is the MOMConnector. This design exists because the binaries were created so that the HealthService can be a standalone part that runs workflows and manages itself, while receiving and sending data through connectors. The connections must be mutually authenticated via Kerberos or certificates (TLS) in non-Active Directory or non-trusted environments. ConnectorManager gets data from DataPublisherManager, which is sent to any parent HealthService (there can be multiple Health Services in a multi-homed environment), and is also responsible for the send queue and data items send and drop priority.

A data item is an output or result of a workflow that must be passed to another module or parent HealthService. For example, the result of an Event Collection Rule needs to be sent to the parent management server, which in turn passes it to a database write module that actually writes the data to the Operational Database or the Data Warehouse Database. The following is an example of this type of data item:

```
<DataItem type="System.XmlData" time="SOME_DATETIME_STAMP"
sourceHealthServiceId="SOME_HEALTHSERVICE_GUID">
  <EventData xmlns="http://schemas.microsoft.com/win/2004/08/events/event">
    <Data>SPECIFIC_EVENT_INFORMATION_1</Data>
    <Data>SPECIFIC_EVENT_INFORMATION_2</Data>
    <Data>SPECIFIC_EVENT_INFORMATION_3</Data>
    <Data>SPECIFIC_EVENT_INFORMATION_X</Data>
    <Data>...</Data>
  </EventData>
</DataItem>
```

## ConfigurationManager

The System Center Management Configuration service is an external service to the HealthService. This service is responsible for pulling the latest configuration from the database at the request of the ConfigurationManager feature. ConfigurationManager uses a cookie that represents a delta hash between any types of configuration that might change, such as management packs, relationships, run-as profiles and accounts, and so on. This helps ConfigurationManager to request changes (updates) for the configuration from just the System Center Management Configuration instead of always requesting the full configuration. The configuration is in the form of an XML file that is located under \Health Service State\Connector Configuration Cache\<<MANAGEMENT\_GROUP\_NAME>\OpsMgrConnector.Config.xml, where <MANAGEMENT\_GROUP\_NAME> is the name of the

Operations Manager Management Group (there can be multiple in multi-homed environments). This XML file holds information about the local HealthService, such as its name, parent HealthService (if any), a list of fail-over HealthServices (if any), list of Resource Pools it is part of (if any), the list of management packs, the list of instances it manages locally (for example, logical disk instances it discovers locally), and a blob that contains an encrypted hash that holds information related to the user accounts it needs to load locally to run different workflows.

After ConfigurationManager processes the XML, it requests any new management packs. These are also available in XML format in the \Health Service State\Management Packs folder. Either on startup or on configuration update, ConfigurationManager loads all workflows from all management packs, like discovery rules, monitors, and every other workflow that should run locally depending on their target classes and what instances of these classes are locally managed by this HealthService. After loading the workflows, ConfigurationManager passes them to PersistenceManager, which writes the workflows to the local ESE database and passes them to ExecutionManager, which starts running the workflows and notifies HealthManager of any monitors that it needs to manage.

## ExecutionManager

ExecutionManager is responsible for actually running the workflows and the set of modules that are loaded and unloaded for each one. Depending on the workflow, ExecutionManager decides whether to run it in an existing or new MonitoringHost process. Depending on isolation level and the credentials the workflow needs to run under, ExecutionManager uses its internal manager, called HostManager, which is responsible for starting/managing an existing or new MonitoringHost process.

If there is any error in either of the modules used by a workflow, ExecutionManager unloads the entire workflow. The same goes for multiple workflow failures inside of a MonitoringHost process. If multiple workflows fail, ExecutionManager might decide to stop the entire MonitoringHost process, thus affecting/stopping any other workflows running under the same process, even if they were successful.

To reduce the footprint and overhead of workflow processing, ExecutionManager implements a technique known as module cookdown for identical modules from separate workflows. This means that if a module with the same input data was already loaded into some running workflow, ExecutionManager uses this existing module instead of loading a new instance of it.

More about how cookdown works and what to consider when authoring management packs can be found at <https://technet.microsoft.com/en-us/library/ff381335.aspx>.

**NOTE** Even though the TechNet article mentioned in the previous paragraph refers to Operations Manager 2007 R2, the same principles apply for Operations Manager 2012, Operations Manager 2012 SP1, and Operations Manager 2012 R2.

## HealthManager

HealthManager is responsible for managing the state of each existing monitor, for each instance that is local to its HealthService. For example, if you have a free space monitor targeting the LogicalDisk class, HealthManager manages the state of each logical disk monitor belonging to each locally discovered LogicalDisk instance. This information is stored in the local ESE database and therefore the HealthService knows the state of each instance. If a monitor successfully runs and a new state is resulted, then HealthManager locally stores this new state and creates a State Change data item. This data item is published and then sent by the ConnectorManager through the MOMConnector to its parent HealthService (management server), which sets the data item to its publish queue to be written to the database via database write modules in the workflow.

## SecureStorageManager

SecureStorageManager loads or unloads and resolves any RunAs profiles and locally associates their RunAs accounts by applying the SecureReferenceOverride. This is the equivalent of applying a normal override for a workflow, but it associates the user account configured in RunAs accounts that have been associated with the locally used RunAs profiles that exist in the various management packs.

## DiscoveryManager

DiscoveryManager handles the scheduling and processing of discovery rules. It also has a sub-feature that manages the schedules for each discovery rule.

## DataPublisherManager

DataPublisherManager publishes the data for each data source, for example, published data that needs to be written to the databases.

## JobManager

When running a task from the console, either on a management server or on an agent, HealthService uses JobManager to run the task.

## PoolManager

HealthService uses PoolManager to manage the resource pools it belongs to. PoolManager creates a new pool object for each resource pool. Each such pool also has three sub-parts:

- **Process** Sends lease requests
- **Observer** Handles lease requests
- **Client** Sends check requests and handles check ACKs

A lease is like a ticket that contains availability information and has a certain expiration period. If a lease expires before it is renewed, then the process stops acting in the pool, drops all the instances managed by it for that particular pool, and becomes unavailable. The length of time (in seconds) that the lease is valid is calculated using the PoolManager registry keys and the following formula:

$$\text{PoolLeaseRequestPeriodSeconds} + 2 * \text{PoolNetworkLatencySeconds}$$

The process for resource pool availability is calculated across the HealthServices in a pool using the PoolManager and its sub-parts. This process is as follows:

1. The process sends a lease request to all other observers in the pool (including its own observer).
2. Each observer answers with a lease grant, which renews the lease the process has for all other processes in the pool.
3. The client sends a check request to all observers in the pool (including its own observer) to verify that they still have a valid lease.
4. Each observer replies with a check ACK telling the requesting client if it still has a valid lease or if it has expired (and is considered unavailable).

In this system, if any of the messages fails to be sent or received, the party involved is considered unavailable. When a process or a client sends a request, the answering observers must build a quorum in order for the answer to be valid. For example, of a total of three observers (including its own observer), at least two of them must reply and have the same information about availability. If a quorum is not met, then the process is unavailable.

The lease grant must be received within a certain time of the lease request or it will also be invalid and the involved process will be set as unavailable. This period (in seconds) is controlled by the PoolNetworkLatencySeconds registry key.

Besides adding to the number of voting observers, there is another reason why the Operational Database also plays the part of observer for each pool. If you install a single management server in a pool, then in order for the pool to function and for the quorum to be met, at least two observers are needed: its own observer and the database. If either of these fail to answer in the pool, it does not matter because if the database is unavailable for some reason, then everything is down anyway and in a one management server pool, if the management server goes down, then the database cannot act as a process and run workflows anyway. The observer availability information for each existing pool is stored in the database in a table called AgentPoolLease. Note that the database observer is known as the default observer.

It is important to know that each process has a process value and a process counter, which are involved in the lease grant and request. When the HealthService becomes available in the pool again after being considered unavailable, (for example, after the service was stopped or after being in maintenance mode), the process value is incremented (+1). The process counter

is incremented after each successful lease request. These values are used to validate the lease grant and request.

A resource pool manages top level instances, also known as top level managed entities, or TLMEs for short. TLMEs are instances that are not managed by a specific HealthService (either an agent, gateway server, or management server), but instead are managed by a resource pool in order to obtain high availability. This does not include Windows agents--these are managed by a parent management server and if it becomes unavailable, they can be configured to fail over to another management server. Previously, in Operations Manager 2007 R2, the TLMEs were managed only by the root management server. So in the case of an outage of the root management server, a lot of TLMEs were not managed by anything, and thus a lot of workflows would no longer run. An example of such a TLME would be the Datawarehouse Synchronization Server. This instance (of its class) is the target class for a lot of important workflows that are related to writing data into the Data Warehouse Database, such as a workflow that synchronizes data (like events) between the Operational Database and the Data Warehouse Database.

In a resource pool, the available TLMEs are split equally, and each HealthService in the pool gets its share. A hashing algorithm based on the GUIDs of the HealthServices in combination with those of the TLME ensures that if all the involved pool members (HealthServices) are available, the exact same TLMEs are given to each member. If a member becomes unavailable, then the hashing is re-done and the entire list of TLMEs from the pool is split equally among the remaining available members. As soon as the unavailable member becomes available again, the same process follows, and it receives the exact same TLMEs it was managing when it became unavailable.

It is important to know that each time there is a pool member change and the TLMEs are reorganized among the available members, a configuration update takes place on all of the available members (HealthServices).

To view all TLMEs that exist per resource pool and per current owning pool member (management server), you can run the following Windows PowerShell script on one of the management servers. It uses Out-GridView to output the data so that you can sort and filter it easily:

```
Import-Module OperationsManager; New-SCOMManagementGroupConnection -ComputerName
"localhost"
$GetTLMEfromPoolTask = Get-SCOMTask -Displayname "Get Top Level Instances Monitored By A
Pool Member"
$HealthServiceClass = Get-SCOMClass -Name "Microsoft.SystemCenter.HealthService"
$ResourcePools = Get-SCOMResourcePool; $TLMEInstances = @()
foreach($pool in $ResourcePools) {
    foreach($ms in $pool.Members) {
        $hs = Get-SCOMClassInstance -Class $HealthServiceClass | ? { $_.DisplayName -eq
$ms.DisplayName }
        $param = @{ PoolId = $pool.Id.ToString() }
        $out = Start-SCOMTask -Task $GetTLMEfromPoolTask -Instance $hs -Override $param
```

```

-ErrorAction SilentlyContinue
    do { $batch = $out.BatchId; Start-Sleep -Seconds 3
    } while ($batch -eq $null)
    do {$result = Get-SCOMTaskResult -BatchId $batch -ErrorAction SilentlyContinue
        $status = $result.Status; Start-Sleep -Seconds 5
    }
    while($status -eq "Started")
    [xml]$output = $result.Output
    if($output -ne $null) {
        $TLMEs = $output.SelectNodes("//ManagedEntity")
        foreach($TLME in $TLMEs) {
            $TLMEInstance = Get-SCOMClassInstance -Id
$TLME.GetAttribute("managedEntityId")
            $TLMEClass = Get-SCOMClass -Id $TLMEInstance.MonitoringClassIds
            $hash = @{ ResourcePool = $pool.DisplayName; ManagementServer =
$ms.DisplayName
                TLMEFullName = $TLMEInstance.FullName; TLMEDisplayName =
$TLMEInstance.DisplayName
                TLMEClass = $TLMEClass; TLMEId = $TLMEInstance.Id
            }
            $obj = New-Object PSObject -Property $hash; $TLMEInstances += $obj
        }
    }
}
}
}
$TLMEInstances | Sort-Object ResourcePool, ManagementServer, TLMEClass | Out-GridView

```

## Understanding the Operational Database

---

The Operational Database is the main element of the Operations Manager infrastructure where all the configuration and available data is stored. As long as there is a backup of this database, the infrastructure can be rebuilt. In many troubleshooting scenarios, retrieving information from the Operational Database is straightforward. It is important to understand the core structure of this database, including the most important tables and the information they contain.

**NOTE** Microsoft does not support making any direct changes in the database. Reading data from the database is only allowed in troubleshooting scenarios.

## ManagedType table

The ManagedType table holds information about each existing class defined in every available management pack. Each time a management pack is imported or updated, if it holds new class definitions, these are stored in this table.

The fields of interest in this table include the following:

- **ManagedTypeId** A class unique GUID that is generated using the internal class name, management pack name, management pack version, and, if sealed, PublicKeyToken
- **TypeName** The internal name of the class
- **BaseManagedTypeId** The corresponding ManagedTypeId of the base class (the one from which this class has been directly derived)
- **ManagementPackId** The GUID of the management pack from the ManagementPack table in which this class is declared
- **ManagedTypeTableName** The name of the table that was generated to store class-specific property information (only for non-abstract classes)

A new table is also created for each non-abstract class when the management pack that contains the class definition is imported. These tables have the format MT\_CLASSNAME, where CLASSNAME is the internal name of the class. In most cases, periods in the name of the class are replaced with dollar signs in the table name due to SQL table naming convention/limitation. For example, for a class named Microsoft.SqlServer.DBEngine, the corresponding generated table name is MT\_Microsoft\$SqlServer\$DBEngine. A corresponding view is also created with the MTV\_ prefix, which in this example is MTV\_Microsoft\$SqlServer\$DBEngine, with some special exceptions related to internal Operations Manager management packs with the Operations Manager 2007 R2 naming convention.

Another useful piece of information to note is that each MT\_CLASSNAME table has a corresponding MT\_CLASSNAME\_Log table. Each MT\_CLASSNAME\_Log table gets an entry whenever a property of an instance of that class type gets changed. So these tables can be used for troubleshooting when you need to determine what changed for that instance in the past couple of days. These tables have a Pre\_PROPERTYNAME and Post\_PROPERTYNAME field for each property of the class the instance belongs to (for each property field in the corresponding MT\_CLASSNAME table).

Continuing the example of the Microsoft.SqlServer.DBEngine class, after looking at the MT\_Microsoft\$SqlServer\$DBEngine table, you can see that the property name field for the DB Engine is ServiceName\_B5502305\_665D\_0E02\_18D6\_4F0BCED0A3DE. With this information and knowing that you want to search for changes to a DB Engine called MSSQL\$TEST, you can use a SQL query like the following one to get the list of changes that have happened in the near past.

```

select *
from MT_Microsoft$SQLServer$DBEngine mt
join MT_Microsoft$SQLServer$DBEngine_Log lg
     on mt.BaseManagedEntityId = lg.BaseManagedEntityId
where mt.ServiceName_B5502305_665D_0E02_18D6_4F0BCED0A3DE = 'MSSQL$TEST'

```

## ManagedTypeProperty table

Each property of any class that has its own non-derived properties is stored in a table called ManagedTypeProperty.

The fields of interest in a ManagedTypeProperty table include the following:

- **ManagedTypePropertyId** A unique GUID generated using the property name, corresponding class internal name, management pack name, version, and PublicKeyToken (if sealed)
- **ManagedTypeId** The GUID of the class to which the property belongs (from the ManagedType table)
- **ColumnName** The name of the field in which this property is stored (from its corresponding generated table name [MT\_CLASSNAME])
- **ManagementPackId** The GUID of the management pack where this property is stored (from the ManagementPack table)

The following is an example query that can be used to view the properties of a class when you know the class name (the Microsoft.Windows.Computer class is used as an example):

```

select
    mp.MPName,
    mt.ManagedTypeId,
    mt.TypeName,
    mtp.ManagedTypePropertyId,
    mtp.ManagedTypePropertyName,
    mtp.ColumnName
from ManagedTypeProperty mtp
join ManagedType mt
     on mt.ManagedTypeId = mtp.ManagedTypeId
join ManagementPack mp
     on mp.ManagementPackId = mtp.ManagementPackId
where mt.TypeName = 'Microsoft.Windows.Computer'

```



## RelationshipType table

Each relationship type defined in each imported management pack is stored in the RelationshipType table. The fields of interest in this table include the following:

- **RelationshipTypeId** A unique GUID generated using the relationship type name, management pack name, version, and PublicKeyToken (if sealed)
- **RelationshipTypeName** The internal name of the relationship type
- **BaseRelationshipTypeId** The corresponding RelationshipTypeId of the base class (the one from which this relationship type has been directly derived)
- **SourceManagedTypeId** The GUID of the source class that is the ManagedTypeId of the class in the ManagedType table
- **TargetManagedTypeId** The GUID of the target class that is the ManagedTypeId of the class in the ManagedType table
- **ManagementPackId** The GUID of the management pack in which this relationship type is defined (from the ManagementPack table)
- **RelationshipTypeViewName** The name of the view that stores relationships that are created of this type (only for non-abstract relationship types)

The views that are created follow the same naming conventions as classes, except that views have a prefix of MTV\_ instead of MT\_. It is better to use the views instead of the MT\_ tables because tables aren't created for all relationship types, while dynamic views are.

The name of a view for an example class called Microsoft.SystemCenter.DataWarehouse.DataWarehouseContainsDataSet is MTV\_Microsoft\$SystemCenter\$DataWarehouse\$DataWarehouseContainsDataSet.

The RelationshipType table also has a corresponding RelationshipTypeProperty table, but it is not as relevant as the property table of the ManagedType table.

## Workflow and basic data tables

Like classes and relationship types, workflows are stored in their own tables:

- **Discovery rules** Stored in the Discovery table
- **Monitors** Stored in the Monitor table
- **Rules** Stored in the Rules table
- **Tasks** Stored in the Task table
- **Alerts** Stored in the Alert table
- **Monitor state changes** Stored in the StateChangeEvent table
- **Events** Stored in multiple partition tables and require the EventAllView view
- **Performance data** Stored in multiple partition tables and requires the PerformanceDataAllView view

## Instance space and discovery tables

The instance space naming refers to all instances that are discovered and stored in the database. This includes instances of classes and relationship types because, on an abstract level, they are the same thing: an instantiated object of an object schema definition. Like instances of classes, instances of relationship types are also discovered by a defined Discovery rule.

The discovery process and tables involved include the following:

1. An entry is made in the DiscoverySource table that ties the source to a corresponding Discovery rule from the Discovery table (which is the actual Discovery rule that discovered this instance).
2. For a class instance, an entry is made in the DiscoverySourceToTypedManagedEntity table; for a relationship type instance, an entry is made in the DiscoverySourceToRelationship table.
3. An entry is made in the TypedManagedEntity table for a class instance, or in the Relationship table for a relationship type instance.
4. For a class instance only (not a relationship type instance), each TypedManagedEntity entry has an associated entry in the BaseManagedEntity table (which either already exists or needs to be created, depending on the class definition).
5. For a class instance, an entry is also made in the corresponding MT\_CLASSNAME table; if it is a relationship instance and the corresponding relationship type has a MT\_RELATIONSHIPTYPENAME table, an entry is created there as well.

For example, to see this information about all discovered class instances for an agent (server) called Server.CONTOSO.com, you can run a SQL query like this one:

```
select
    tme.TypedManagedEntityId,
    mt.TypeName,
    d.DiscoveryName,
    bme.FullName,
    bme.DisplayName
from TypedManagedEntity tme
join BaseManagedEntity bme
    on bme.BaseManagedEntityId = tme.BaseManagedEntityId
join ManagedType mt
    on mt.ManagedTypeId = tme.ManagedTypeId
join DiscoverySourceToTypedManagedEntity dsttme
    on tme.TypedManagedEntityId = dsttme.TypedManagedEntityId
join DiscoverySource ds
    on ds.DiscoverySourceId = dsttme.DiscoverySourceId
join Discovery d
    on d.DiscoveryId = ds.DiscoveryRuleId
where bme.FullName like '%Server.CONTOSO.com%'
order by mt.TypeName, d.DiscoveryName
```

Relationship instances follow the same principle, except that the relationship type related tables are used. To look at all the discovered relationship instances where Agent Server.CONTOSO.com has a class instance as the source of the relationship instance, you can run a SQL query like this:

```
select
    rt.RelationshipTypeName,
    d.DiscoveryName,
    r.RelationshipId,
    bme1.FullName as [Source],
    bme2.FullName as [Target]
from Relationship r
join TypedManagedEntity tme1
    on r.SourceEntityId = tme1.TypedManagedEntityId
join BaseManagedEntity bme1
    on tme1.BaseManagedEntityId = bme1.BaseManagedEntityId
join TypedManagedEntity tme2
    on r.TargetEntityId = tme2.TypedManagedEntityId
join BaseManagedEntity bme2
    on tme2.BaseManagedEntityId = bme2.BaseManagedEntityId
join RelationshipType rt
    on rt.RelationshipTypeId = r.RelationshipTypeId
join DiscoverySourceToRelationship dstr
    on r.RelationshipId = dstr.RelationshipId
join DiscoverySource ds
    on ds.DiscoverySourceId = dstr.DiscoverySourceId
join Discovery d
    on d.DiscoveryId = ds.DiscoveryRuleId
where bme1.FullName like '%Server.CONTOSO.com%'
order by rt.RelationshipTypeName, d.DiscoveryName
```

**NOTE** You may notice that most of these tables also have an `IsDeleted` field. If the value of this field is inconsistent across the related entries in all of these tables, the database will be in a corrupt state. Furthermore, making direct changes in the database is not supported, so do not change the `IsDeleted` field manually. It is part of the internal discovery/undiscovery/grooming process and should not be manually changed.

Another important point about the instance space and discovery process is that an instance can have more than one discovery source. This is almost never the case for Operations Manager, but in a lot of situations true and relevant for System Center Service Manager (SCSM). In SCSM, you can have multiple discovery sources when you have multiple Connectors configured, and they can be of the same or different type. The SCSM architecture is based on the Operations Manager architecture and was adapted to fit the SCSM needs.

# Overview of the Data Warehouse Database

---

Even though the Data Warehouse Database contains only historical data and does not hold the entire configuration and critical information needed for the Operations Manager infrastructure, in Operations Manager 2012 and higher the Data Warehouse Database is a mandatory part because some features (such as dashboards showing performance data) cannot function without it. In a disaster recovery scenario, this database requires a fresh backup.

## Instance space

The tables containing the instance space from the Data Warehouse Database are simpler than the ones from the Operational Database. Almost every important table has a defined view, and it is much simpler to use the views for troubleshooting than the actual tables. When building custom reports, the same rule applies—in the custom report SQL queries, always use views instead of tables.

All the instances in the Data Warehouse Database can be found using the `vManagedEntity` view. However, because there can be more than one Operations Manager management group registered to a Data Warehouse Database, each instance is associated with a management group. The list of existing management groups can be found in the `vManagementGroup` view. For example, to see the instances that exist for the agent `Server.CONTOSO.com`, you would run the following SQL query:

```
select
    mg.ManagementGroupDefaultName,
    me.ManagedEntityGuid,
    me.FullName,
    me.Name,
    me.DisplayName,
    me.Path,
    me.DWCreatedDateTime
from vManagedEntity me
join vManagementGroup mg
    on me.ManagementGroupRowId = mg.ManagementGroupRowId
where me.FullName like '%Server.CONTOSO.com%'
```

## Datasets and their tables

Datasets are the backbone of the functionality of the Data Warehouse Database. They exist for each type of data that can be viewed in reports. Out of the box, the following datasets are available:

- **Performance dataset** This is the container for the performance data collected by the performance collection rules (usually via performance objects/counters).

- **State dataset** This is the container for the availability data, which is calculated based on state and state change information coming from monitors.
- **Alert dataset** This is the container for all alerts raised by monitors and rules.
- **Event dataset** This is the container for all events collected by event collection rules.
- **APM dataset** This is the container for performance and event data about monitored managed code applications collected by the Application Performance Monitoring feature.

New datasets might appear when you import various management packs. For example, the Microsoft Exchange 2010 management pack brings six new datasets. You can find a good description and schema reference for the Data Warehouse Database and its tables and datasets at <https://technet.microsoft.com/en-us/library/gg508713.aspx>.

**NOTE** Even though the TechNet article referenced in the previous paragraph is written for Operations Manager 2007 R2, it also applies to Operations Manager 2012, Operations Manager 2012 SP1, and Operations Manager 2012 R2 because the Data Warehouse Database schema did not really change in this area. The only big change in the 2012 versions is the addition of the APM dataset with its tables and views.

Each dataset comes with its own database schema. To view information about the datasets, use the vDataset view. To view additional information related to the datasets, you can create SQL queries using the Join parameter on the StandardDataset and StandardDatasetAggregation tables:

```
select
    ds.DatasetId,
    ds.DatasetDefaultName,
    sds.SchemaName,
    ds.ConfigurationXml,
    sds.DefaultAggregationIntervalCount,
    sds.StagingProcessorStoredProcedureName,
    sda.AggregationTypeId,
    sda.AggregationIntervalDurationMinutes,
    sda.AggregationStartDelayMinutes,
    sda.BuildAggregationStoredProcedureName,
    sda.DeleteAggregationStoredProcedureName,
    sda.GroomStoredProcedureName,
    sda.IndexOptimizationIntervalMinutes,
    sda.MaxDataAgeDays,
    sda.GroomingIntervalMinutes,
    sda.MaxRowsToGroom,
    sda.StatisticsMaxAgeHours,
    sda.StatisticsUpdateSamplePercentage,
    sds.LastOptimizationActionDateTime,
```

```

        sds.LastOptimizationActionSuccessfulCompletionDateTime,
        ds.InstallCompletedInd,
        ds.InstalledDateTime,
        sds.DebugLevel
from vDataset ds
join StandardDataset sds
    on ds.DatasetId = sds.DatasetId
join StandardDatasetAggregation sda
    on ds.DatasetId = sda.DatasetId
order by ds.DatasetDefaultName

```

Some datasets are aggregated, while others are stored in their raw format. Performance data and state (availability) data is aggregated, which means that it is transformed from its raw format (every collected value) and transformed into hourly and daily data. This means that you can see the raw (granular) performance and state data in the Operations Manager console views, but if you run reports, they will retrieve the data from the long-term storage (the Data Warehouse Database) and you will only be able to see hourly or daily aggregated data.

This type of transformation (aggregation) does not really make sense for information like alerts or events that were collected. This is why each event and alert that is collected is directly stored in its raw format.

Knowing this, you can understand that running the preceding SQL query provides some very useful information about the datasets. For example, in the vDataset view, a field called ConfigurationXML contains XML that indicates the dataset type (aggregated or stored). The vDataset view also reveals other useful information like the base table name of the dataset, the maximum rows a partition table is allowed to have before a new one is created, and the maximum rows to groom (delete) from the raw tables after the data has been processed (aggregated or stored in the final tables). The same information is also available in the StandardDatasetAggregation table. The configuration XML looks like this example from an alert dataset:

```

<Configuration>
  <Storage>
    <BaseTableName>Alert</BaseTableName>
    <MaxTableRowCount>1000000</MaxTableRowCount>
    <MaxDataAgeDays>400</MaxDataAgeDays>
    <GroomingIntervalMinutes>240</GroomingIntervalMinutes>
    <MaxRowsToGroom>50000</MaxRowsToGroom>
    <IndexOptimizationIntervalMinutes>240</IndexOptimizationIntervalMinutes>
  </Storage>
  <RawInsertTableCount>1</RawInsertTableCount>
  <BlockingMaintenanceDailyStartTime>01:00</BlockingMaintenanceDailyStartTime>
  <BlockingMaintenanceDurationMinutes>240</BlockingMaintenanceDurationMinutes>
</Configuration>

```

The partition tables are smaller and result in improved read/write performance. While Operations Manager is writing data to one partition tables (depending on the date/time of the data), it might also be reading data from other partition tables through the view that joins them, all while running a report. Although writing to a table is a blocking action, in this situation, there is no blocking because Operations Manager isn't reading from the table where it is currently writing data (and thus holding a no-read lock on it).

Whenever a dataset is processed (aggregated or stored), a different SQL stored procedure is used because each dataset is a different type of data and needs to be treated or transformed differently. A general SQL stored procedure runs on a fixed schedule for each dataset. This stored procedure is called `StandardDatasetMaintenance`, and it gets the `DatasetId` as a parameter.

The `StandardDatasetMaintenance` stored procedure processes, optimizes, and grooms the datasets by running other stored procedures. First, it runs `StandardDatasetProcessStaging`, which gets the specific stored procedure of the dataset that is currently being processed from the `StagingProcessorStoredProcedureName` field of the `StandardDataset` table, and runs it. The staging of the data means something different depending on the dataset and what its staging stored procedure does, but, in essence, staging means to take the data out of the incoming staging tables, prepare it, and insert it into the raw partition tables of the dataset.

Take for example the state dataset: its staging stored procedure is called `StateProcessStaging`, and it prepares the raw partition tables and then moves the data from the `State.StateStage` staging table to the actual raw partition tables. Depending on how much time has passed since the last time the dataset was worked on, the `StateProcessStaging` stored procedure starts the actual transfer or transformation of the data from the raw tables. These stored procedures are not all called every time, one after the other. It depends on how long in seconds the previous actions take. First Operations Manager calls the `StandardDatasetGroom` stored procedure, which deletes data that is already processed (moved to the long-term storage tables) from the raw partition tables, based on the grooming settings of each dataset.

Next, Operations Manager calls the `StandardDatasetAllocateStorage` stored procedure, which, based on the settings of each dataset for the storage partition tables, might create a new storage partition table if the current one is considered full. The storage partition tables are the long-term tables where the data is kept. These are also partitioned for the same performance improvement reasons mentioned earlier.

After finishing with storage allocation, Operations Manager continues to optimize the data by calling the `StandardDatasetOptimize` stored procedure. This runs SQL-specific optimization actions on the data, such as online optimization, index optimization, and so on.

Next, the actual data transformation begins with the `StandardDatasetAggregate` stored procedure. This stored procedure aggregates the unprocessed data from the raw partition tables by calling another stored procedure, which, like for the other, is different for each dataset and aggregation type.

There are three types of data aggregation:

- **0** No aggregation (for non-aggregated datasets)
- **20** Hourly aggregation
- **30** Daily aggregation

The aggregation stored procedure that is called for the current dataset comes from from the `BuildAggregationStoredProcedureName` field in the `StandardDatasetAggregation` table.

Finally, the data processing is finished and the aggregated data is available in the following views:

- **Alert.vAlert** The raw non-aggregated alert data (because Operations Manager doesn't aggregate alerts)
- **Event.vEvent** The raw non-aggregated event data (because Operations Manager doesn't aggregate events)
- **State.vStateHourly** The hourly aggregated state (availability) data
- **State.vStateDaily** The daily aggregated state (availability) data
- **Perf.vPerfHourly** The hourly aggregated performance data
- **Perf.vPerfDaily** The daily aggregated performance data
- **apm.\*** Different Application Performance Monitoring-related aggregated data about monitored managed code application events and performance

There are more views possible if you have more management packs installed that come with their own datasets. The Microsoft Exchange 2010 management pack example from earlier in this chapter is a good example of this as well because it brings many `Exchange2010.*` views that store the long-term Microsoft Exchange-related reporting data.

## General troubleshooting steps

---

There are a few basic troubleshooting steps that you need to take if you find or think that something is wrong with the Operations Manager environment itself or with any of the agents. Some of these steps are also recommended to be performed on a regular basis, such as during maintenance windows and patching downtime, to make sure that the Operations Manager environment is not suffering or starting to suffer from any critical failures or data overload, which would later manifest as massive performance issues.

Before going into the event logs, however, you always need to check that the service (on the agent) or services (on the management server) are started and running. It's also a good idea to monitor the alerts about the Operations Manager environment that come into the console and alternatively to occasionally look at the views and dashboards specific to the Operations Manager environment in the console (in the Operations Manager folder in the



Monitoring section). Operations Manager does a good job monitoring itself for failures as well. If there are still problems, a deeper analysis is required.

## Operations Manager event log

Every computer that has at least one Operations Manager part installed has a new event log present in the Event Viewer under Applications and Services Logs called Operations Manager. The Operations Manager event log is verbose. In most cases, it is all you need to figure out what is going wrong in your environment. In some situations, it at least sets you on the right path to an answer. There are very few situations where analyzing this event log thoroughly does not help you in any way.

The most basic thing to check is that the information event 6022 is being logged periodically, which indicates that the HealthService is running at least some workflows (through MonitoringHost processes) and is not in a hung state or something similar.

In some situations, it is a good idea to look at all the events without filtering anything out. It is more than enough to go through the events from the past 6 to 10 hours because if there is a failure at some point, that failure will repeat itself often. However, the situations where you need to look at all events are usually related to local performance issues or configuration updates.

Usually, you should first filter the event log just on Error and Warning events (Operations Manager never triggers a Critical level event). It is good to go through each Error or Warning event and make an analysis along these lines:

- What is the frequency of the event?
- What is the exact event description?
- For events with the same event ID, are these really the exact same event based on a careful comparison of the event description?
- If you see a problem event for some workflow that you know should run every 10 minutes, is the last such event fresh or is it too old, maybe indicating this was a one-time problem?
- Is there one or more events that seem to be spamming the event log? For example, do you see the same event 50 times in 1 second, or something similar?

The most important thing of all is to keep in mind that events that have the same event ID, or even a very similar event description, might actually be totally different failures. This is why you must look at the event description closely. There can be two (or more) events that have the same event ID and exact same event description, but with a very specific and important difference: a different error code in the description. Something like this is easy to miss. Checking the Application and System event logs is also a good idea for crashes or operating system failures.

A large variety of event IDs are usually grouped in ranges, from 1200 to 1210 for example. These event IDs can have different severities. However, a specific event ID can and always will have the same severity. The source of the events is also an important aspect because, after analyzing this event log for some time, you might start detecting patterns, recognizing event IDs or sources, and this can help you set more specific filters for the source or event ID range to speed up the analysis.

**NOTE** In very few situations, some of the same event IDs may be triggered by different scripts from certain workflows. The important difference is that the source will always be Health Service Script for these.

## ETW tracing for Operations Manager

Event Tracing for Windows (ETW) is a tracing technology that Windows uses. ETW tracing is also used in most Microsoft software and in Operations Manager as well. A tutorial about the ETW framework can be found at <https://technet.microsoft.com/en-us/library/jj714799.aspx>.

To start the trace with all available providers (trace everything), you can follow the Knowledge Base article at <http://support.microsoft.com/kb/942864>.

**NOTE** Even though the Knowledge Base article reference in the previous paragraph refers to Operations Manager 2007 R2, the same principles apply to Operations Manager 2012, Operations Manager 2012 SP1, and Operations Manager 2012 R2, except that you need to use the proper path to the Tools folder as appropriate for the product version.

Each module of each feature uses a certain trace provider. The list of providers with their GUIDs, names, and descriptions can be reviewed in `\Tools\Providers.xml`. The following is an example of a provider defined in the XML:

```
<TraceProvider name="Mom Modules" area="Modules" guid="B8530492-0105-4E89-844A-13AD3FE10E71" hidden="false" description="Key MOM Modules Eventlog Performance data Mapper modules Filter modules Correlation/Consolidation Numeric threshold modules Baselineing modules Application log modules WMI modules" />
```

The tool `\Tools\TraceLogSM.exe` can be used to start tracing for either one specific provider or multiple providers. Depending on the `-guid` parameter passed to this tool, you can either specify a file with multiple GUIDs or just a single GUID for a provider to be traced. The following are the parameters and how they need to be configured:

```
Usage: tracelog [actions] [options] | [-h | -help | -?]  
actions:  
-start <SessionName> Starts the trace session and enables its providers  
-stop <SessionName> Stops the trace session (disables providers, flushes  
buffers)  
-update <SessionName> Updates the trace session
```

```

-enable <SessionName> Enables providers for the trace session
-disable <SessionName> Disables providers for the <SessionName> session
-flush <SessionName> Flushes the trace session buffers
-q <SessionName> Displays the status of the trace session
-enumguid Lists ETW registered providers on the system
-l Lists all trace sessions
-x Stops all trace sessions
-h | -? |-help Displays usage

```

options:

```

-f [path/]<filename> Sends trace messages to specified trace log (.etl)
-rt Specifies a real-time trace session
-um Specifies a private (user mode process) trace session
-guid #<GUID>|<file> Specifies providers for the trace session.
-flag <Flag> Enables provider-defined trace flags
-level <Level> Enables provider-defined trace level
-b <SizeKB> Sets buffer size in KB.
-min <NumBuffers> Sets minimum number of buffers
-max <NumBuffers> Sets maximum number of buffers
-ft <NumberOfSeconds> Flushes trace buffers (in addition to flush when full)
-paged Use pageable memory for buffers
-seq <MaxFileSizeMB> Sequential trace log up to MaxFileSize in MB
-cir <MaxFileSizeMB> Circular trace log up to MaxFileSize in MB
-prealloc Pre-allocates trace log
-append Append trace messages to log specified by -f
-newfile <MaxFileSizeMB>
                                Creates new trace log at MaxFileSize (-f <filename> must

```

include %d)

```

-ls Generates local sequence numbers
-gs Generates global sequence numbers
-age <n> Specifies how long unused trace buffers are retained

```

(Windows 2000 only)

```

-kd Sends trace messages to KD or WinDBG

```

Additional options for NT Kernel Logger trace session:

\* use "NT Kernel Logger" for <SessionName> or omit <SessionName> \*

```

-noprocess Disable tracing of process start/end
-nothread Disable tracing of thread start/end
-nodisk Disable tracing of disk I/O
-nonet Disable tracing of TCP/IP and UDP
-fio Traces file I/O events
-pf Traces all page faults
-hf Traces hard page faults
-img Traces image load events
-cm Traces registry calls

```

NOTE: The following actions and options are not supported on Windows 2000.

```

-flush          -ls
-enumguid       -gs
-append         -paged
-newfile        -UsePerfCounter
-prealloc       -UseCPUCycle

```

To start the tracing just for the MOM Modules provider, start the trace using the following command line:

```

TraceLogSM.exe -start CustomTrace -flag 0x1F -level 6 -f
%windir%\Logs\OpsMgrTrace\CustomTrace.etl -b 64 -ft 10 -cir 999 -guid #B8530492-0105-
4E89-844A-13AD3FE10E71

```

After you have reproduced what you need, stop this trace, which you named CustomTrace per the command line, run the following command:

```

TraceLogSM.exe -stop CustomTrace

```

To create a trace file that includes more than a single provider, create a simple text file and give it whatever extension you want. Then copy this file to a folder such as %windir%\Logs\OpsMgrTrace\ and start a trace. For example, for the text file %windir%\Logs\OpsMgrTrace\CustomTraceFile.txt, start the trace as follows:

```

TraceLogSM.exe -start CustomTrace -flag 0x1F -level 6 -f
%windir%\Logs\OpsMgrTrace\CustomTrace.etl -b 64 -ft 10 -cir 999 -guid
%windir%\Logs\OpsMgrTrace\CustomTraceFile.txt

```

Because of the way ETW tracing works, a trace is an ETL file that needs a TMF file to be able to format the trace into human-readable format. As long as the custom trace file (ETL) is located under %windir%\Logs\OpsMgrTrace\, you can format any trace just by running the standard StartTracing.cmd batch file from the Tools folder. This batch file eventually calls the TraceFmtSM.exe tool to perform the actual formatting. The TMF files that are used to format the traces are located in the \Tools\TMF folder but are archived in CAB files. After the FormatTracing.cmd batch file runs at least once, it extracts every TMF from the CABs and concatenates every file into a single TMF file called all.tmf, which is saved under the Tools folder directly. To manually format a trace using the TraceFmtSm.exe tool, after the all.tmf file is created, you can use the following command line:

```

TraceFmtSM.exe %windir%\Logs\CustomTrace.etl -tmf all.tmf -o
%windir%\Logs\CustomTrace.log

```

Each trace line in the log has the following prefix:

```

[CPU]ProcessID.ThreadID::TimeStamp [TraceProvider] [Flags] [Level]
Component:Function{SourceAndLine} Trace Message

```

The following is an example of a trace line to further show this:

```

[2] [13196] [12068] [05/03/2015-13:04:37.919] [ModulesWMI] [] [Error] []
[CWMIAsyncProbe::OnTimerCallback] [WMIProbe_cpp232] Attempt 1705 to connect to WMI
failed - HRESULT=8004100E

```

## SQL Server Profiler Tracing

SQL Server Profiler Tracing is a feature of Microsoft SQL Server that allows you to trace different actions in SQL Server (see [https://msdn.microsoft.com/en-us/library/ms181091\(v=sql.110\).aspx](https://msdn.microsoft.com/en-us/library/ms181091(v=sql.110).aspx)). In some situations, when the actual error in Operations Manager is coming from an SQL query or when you want to generally understand what Operations Manager is doing when running a certain internal workflow, you can use SQL Server Profiler to create a trace that you can later analyze in detail.

A trace taken with SQL Server Profiler shows what SQL stored procedures and individual queries have been processed. In general, you use the trace to determine what stored procedures are running and then you look at the definition of those stored procedures to understand what they are doing (see <https://msdn.microsoft.com/en-us/library/ms345443.aspx#SSMSPProcedure>).

When starting a trace, you can include or exclude certain events or event fields. For standard analysis, these settings should be enough:

- Errors and Warnings – User Error Message – ApplicationName, SPID, DatabaseName, StartTime, TextData
- Stored Procedure – SP:Starting – ApplicationName, SPID, DatabaseName, StartTime, TextData
- Stored Procedure – SP:StmntStarting – ApplicationName, SPID, DatabaseName, StartTime, TextData
- TSQL – SQL:BatchStarting – ApplicationName, SPID, DatabaseName, StartTime, TextData
- TSQL – SQL:StmntStarting – ApplicationName, SPID, DatabaseName, StartTime, TextData

## Performance troubleshooting

---

In general, the overall performance of the Operations Manager environment, including management servers and the console, is tightly connected to the size of the tables in the databases and the amount of data that the management servers receive and write to the databases. One of the most important aspects of maintaining a healthy and performant Operations Manager environment is management pack tuning. Each time you import a new management pack, you need to monitor the data it collects and how it behaves in the following one to two weeks. How many alert, state change, event, and performance data a management pack collects is important with respect to performance. When dealing with a performance issue, you should always thoroughly analyze the Operations Manager event log of each management server to see any possible workflow failures involved or other events that might help determine what is going on and where the pressure is coming from.

## Data grooming settings

Data grooming settings can also play a role in improving the database performance. The default grooming setting for events and performance data for example, is seven days. This means that data as old as seven days is available in the Operational Database. This is a good setting for a standard environment, but if, for example, you don't need to see raw performance data in the console views for the last seven days and only need the data from the last three days, you should reduce the grooming setting to only three days. This ensures a smaller Operational Database and thus much better console performance. All of the performance data is still in the Data Warehouse Database, aggregated to hourly and daily data, which is usually more than enough for viewing the performance history of servers for some past timeframe.

## Event and performance data collection

Reducing data collection has a key role to play in the performance of not only the management servers, but also the database, and thus the console. Usually, a management pack comes with all event collection rules and all performance collection rules enabled by default. Most of the time, though, depending on business requirements, you don't need to collect all events and performance data that is collected. It is best to review all of these event and performance collection rules, determine what data they collect, and disable what you don't need.

The following SQL query can be used to view the list of performance collection rules so that you can disable what you don't need:

```
select
    mp.MPName,
    r.Name,
    dsv.DisplayName,
    dsv.Description
from RuleView r
join DisplayStringView dsv
    on dsv.LTStringId = r.Id
join ManagementPack mp
    on r.ManagementPackId = mp.ManagementPackId
where
    dsv.LanguageCode = 'ENU' and
    dsv.Description is not null and
    r.Category = 'PerformanceCollection'
order by mp.MPName
```

The following similar SQL query can be used to view the list of event collection rules so that you can disable what you don't need.

```

select
    mp.MPName,
    r.Name,
    dsv.DisplayName,
    dsv.Description
from RuleView r
join DisplayStringView dsv
    on dsv.LTStringId = r.Id
join ManagementPack mp
    on r.ManagementPackId = mp.ManagementPackId
where
    dsv.LanguageCode = 'ENU' and
    dsv.Description is not null and
    r.Category = 'EventCollection'
order by mp.MPName

```

## State change event frequency (noisy monitors)

The number of state change events that monitors generate can have one of the biggest performance impacts not only on the console performance, but also on the management servers that need to write a large amount of data into the databases. Another reason for a big StateChangeEvent table is state change events, which are older than the data grooming setting for state changes (default 7 days). This can happen if you manually (or via some automated/scripted method) close alerts without resetting the monitors that raised them. It is against best practice to do this because the grooming stored procedure to clean-up state changes does not also delete state changes that belong to a monitor that is not in the Healthy state. Additionally, a high number of state changes might cause the stored procedure to time out and not be able to delete everything.

To view the list of the top 100 monitors ordered by those that generated the most frequent state changes in the last 14 days, use the following SQL query:

```

select
    distinct top 100 count(sce.StateId) as StateChanges,
    dsv.DisplayName,
    m.Name,
    mt.TypeName
from StateChangeEvent sce
join state s
    on sce.StateId = s.StateId
join MonitorView m
    on s.MonitorId = m.Id
join DisplayStringView dsv
    on dsv.LTStringId = m.Id
join ManagedType mt
    on m.TargetMonitoringClassId = mt.ManagedTypeId

```

```

where
    m.IsUnitMonitor = 1 and
    dsv.LanguageCode = 'ENU' and
    sce.TimeGenerated > dateadd(dd, -14, getutcdate())
group by dsv.DisplayName, m.Name, mt.TypeName
order by StateChanges desc

```

To manually clean up the StateChangeEvent table and make sure you clean all data that is older than the data grooming setting, even if the corresponding monitors are still not in the Healthy state, you can run the following SQL query:

```

set ansi_nulls on
go
set quoted_identifier on
go
set nocount on
declare
    @Err int,
    @Ret int,
    @DaysToKeep tinyint,
    @GroomingThresholdLocal datetime,
    @GroomingThresholdUTC datetime,
    @TimeGroomingRan datetime,
    @MaxTimeGroomed datetime,
    @RowCount int = 1
set @TimeGroomingRan = getutcdate()
select @GroomingThresholdLocal = dbo.fn_GroomingThreshold(DaysToKeep, getdate())
from PartitionAndGroomingSettings
where ObjectName = 'StateChangeEvent'
exec p_ConvertLocalTimeToUTC @GroomingThresholdLocal, @GroomingThresholdUTC out
delete mjs
from MonitoringJobStatus mjs
join StateChangeEvent sce
    on sce.StateChangeEventId = mjs.StateChangeEventId
where sce.TimeGenerated < @GroomingThresholdUTC
while(@RowCount > 0) begin
    delete top(10000) sce
    from StateChangeEvent sce
    where TimeGenerated < @GroomingThresholdUTC
    set @RowCount = @@rowcount
end
update PartitionAndGroomingSettings
set
    GroomingRunTime = @TimeGroomingRan,
    DataGroomedMaxTime = @MaxTimeGroomed
where ObjectName = 'StateChangeEvent'

```



## Top issues and best practices

---

Some best practices and top issues are well documented in Microsoft Knowledge Base articles.

### Recommended operating system fixes for Operations Manager

The best practice is to always keep the Operations Manager management servers, gateway servers, and agents up to date with the update rollups for the Windows operating system, as well those for Operations Manager.

Many different issues on either management servers or agents are caused by known problems with certain versions of the Windows operating system. Because of this, a Knowledge Base article listing recommended Windows operating system (version dependent) hotfixes and update rollups is available at <http://support.microsoft.com/kb/2843219>.

### Gray management servers or agents

A management server or agent that is shown as gray in the console means that the corresponding HealthService is not healthy for some reason, as is explained at <http://technet.microsoft.com/en-us/library/hh212723.aspx>.

A good Knowledge Base article to help troubleshoot the different scenarios for agents that are displayed as gray is available at <http://support.microsoft.com/kb/2288515>. One of these scenarios also describes the presence of warning event 2115, which would most likely appear on management servers or gateway servers and may involve performance problems. Another great Knowledge Base article for troubleshooting this issue in detail is available at <http://support.microsoft.com/kb/2681388>.

### Data missing from reports

There are various reasons why data might be missing from reports, but most commonly missing information is caused by too much data being written to the Data Warehouse Database, which can cause a temporary stall in the data aggregation process. This is usually associated with the 31552 event that describes a TimeOut error that occurs during the StandardDataMaintenance stored procedure. A Knowledge Base article to help troubleshoot this type of issue is available at <http://support.microsoft.com/kb/2573329>.

### Resource pool unavailable

There are no real issues with the resource pools themselves. The resource pool high availability and load balancing technology works as long as the management servers from a pool can communicate with each other fast enough (before the lease expires) and as long as the configuration on them is successfully loaded and initialized.

For a Management Server to be available in a pool, it must be able to successfully load and initialize its configuration. If there seems to be an issue with a resource pool, the first thing to check is whether there are any issues with the configuration update on any of the management servers and to ensure sure that the event 1210 is present in the event log when the configuration is updated successfully.

Issues with a resource pool can be caused either by network connectivity issues or by slow connectivity issues, or, if one or all management servers from the pool suffer from performance issues, it might be due to too much data coming in, too many workflows, too many failures, or a critical failure in a workflow, which causes them to unload every workflow.

A management server suffering from performance issues like this constantly fluctuate between being available and unavailable in the pool. To determine if any management server is suffering from an issue like this, you can use the following SQL query on the Operational Database and review the process version:

```
select
    bme2.DisplayName as [Resource Pool],
    bme1.DisplayName as [HealthService],
    apl.LeaseExpirationDateTimeUtc,
    apl.ProcessCounter,
    apl.ProcessVersion
from AgentPoolLease apl
join BaseManagedEntity bme1
    on apl.ProcessId = bme1.BaseManagedEntityId
join BaseManagedEntity bme2
    on apl.AgentPoolId = bme2.BaseManagedEntityId
order by bme2.DisplayName, bme1.DisplayName
```

# Using Operations Manager in cloud environments

This chapter provides an overview of the cloud-based add-ins that are available for Operations Manager. The first part of the chapter explains Global Service Monitor, which allows you to perform web testing from locations all over the world. The second part explains Azure Operational Insights, the Operations Manager add-in that started life as System Center Advisor. The third part is a complete description of how to install Operations Manager in Microsoft Azure Infrastructure as a Service (IaaS). The last part of this chapter describes how you can use Operations Manager to monitor your cloud-based deployments.

## Using Global Service Monitor

---

Global Service Monitor is a Microsoft Software Assurance benefit of the System Center suite. You can sign up for a free trial account and use Global Service Monitor for free for up to 90 days. Beyond the 90-day free trial period, Global Service Monitor is available only to customers with active Software Assurance coverage for their System Center 2012 R2 server management licenses.

## Understanding DevOps

Operations teams use isolated operations tools and workflows that they understand, but developers are usually unfamiliar with these tools. The developers use their own tools and processes, which the operations teams do not always understand. Users or monitoring software detect errors in production applications, but developers don't necessarily have access to the production environment. So when a problem happens in the production environment, there is no direct way of debugging the application, which is what a developer would do if the application was installed in a development environment. Instead, the detected problem is assigned to the developers, but since they don't have actual debugging information, the mean time to repair is high. Furthermore, since the development environment is not identical to the production environment, the detected error is sometimes not reproducible, and the developer just does not know where to start troubleshooting.

The main challenge in DevOps is reconciling the different perspectives of developer and operations teams. Operations does not know the code; development cannot test on the production environment. Developers and operations people also use different tools: developers use Team Foundation Server (TFS) and Visual Studio, while operators use System Center. The focus points are also different: developers are concerned about bugs and TFS work items, while operators are concerned about their service level agreements (SLAs). Developers worry about code, while operators worry about availability and latency. This is where application performance monitoring in Operations Manager comes in. Global Service Monitor is one part of the 360-degree application performance monitoring that enables DevOps. Using TFS integration in Operations Manager, you can redirect alerts from Operations Manager to the developer's TFS environment so that a fix can be developed. With the included IntelliTrace logging, the developer can see the actual line of code where the application failed.

## Understanding global service monitoring

Global Service Monitor enables you to externally monitor customer services, such as websites. It runs availability and performance tests for web applications using remote probes and, with these probes, offers 360-degree application monitoring. These probes can be located in both a private cloud (internal locations) and the public cloud (external locations). From an application monitoring perspective, global service monitoring is the mid-mile monitoring part, whereas internal synthetic transactions (made using distributed applications in Operations Manager) is the first mile, and client-side monitoring is the final mile. You use application performance monitoring in Operations Manager to monitor the code from your applications in your environment.

Global Service Monitor is managed by Microsoft in Microsoft Azure. You can use it to monitor your external-facing applications and websites from a worldwide perspective regardless of where these sites are hosted, so you can also monitor externally provided websites with Global Service Monitor. For example, this can be beneficial if your e-commerce website relies on a third-party website for currency conversions—your site might be functional, but the third-party website might have issues that reflect on your website.

## Installing Global Service Monitor

You can start testing with Global Service Monitor by signing up for the free 90-day trial subscription. You can also sign up with an existing organizational ID (OrgId). If your organization has Software Assurance, Global Service Monitor is provided for free. The link to sign up for the trial is <https://orgaccount.microsoft.com/signup?offerIds=BE2A46EF-0639-43a4-8323-BB5E1D4340D3>.

The next step in the sign-up process is downloading and installing the required management packs, the Alert Attachment and Global Service Monitor management packs. Without these management packs, you won't see any reference to Global Service Monitor in Operations Manager. These management packs can be found at <http://www.microsoft.com/en-us/download/details.aspx?id=36422>.

The Alert Attachment management pack is also required, but is not included in this download; it can be found on the Operations Manager installation media. It is located in the ManagementPacks directory and is named Microsoft.SystemCenter.AlertAttachment.MPB. Since this management pack might be updated in an update rollup, check for a newer version in the update rollup that you have installed in your Operations Manager environment. This management pack is needed for the TFS integration with Visual Studio Web Test monitoring to work. In this management pack, you configure how to forward the IntelliTrace logs to TFS. For this process to work, further configuration of this management pack is needed and is explained at <https://technet.microsoft.com/en-us/library/jj899889.aspx>.

To start your subscription, you need to provide your onmicrosoft.com account (OrgId) in the Administration pane of the Operations Manager console. This is where the configuration page for Global Service Monitor appears after you successfully install the required management packs. If you haven't already done so, you will also need to install Windows Identity Foundation. This prerequisite is automatically selected on the Global Service Monitor configuration page in the Operations Manager console and is automatically installed when you click the link. This installation is necessary on all management servers that are in the resource pool to be used for Global Service Monitor. Although you are prompted to reboot the machine, you only need to restart the System Center management service (the Healthservice) on the management servers to complete the installation.

After entering the OrgId and restarting, you can continue with the Global Service Monitor configuration. Select a resource pool to be used for the monitoring, and specify a proxy server if that is needed for Internet access. With this information supplied, you can start the subscription and begin configuring web tests for Global Service Monitor to perform.

## Getting started with Global Service Monitor

After configuring Global Service Monitor in the Administration pane of the Operations Manager console, you can start creating web tests within the Authoring pane. This pane is located under Management Pack Templates. Select either Web Application Availability Monitoring or Visual Studio Web Test Monitoring.

With Web Application Availability monitoring, you can input the URLs to monitor, or you can import them from a comma-separated value file. Next, you can choose an Operations Manager agent or an Operations Manager resource pool to do the monitoring from an internal location, or you can select an external location if you have Global Service Monitor configured correctly. The Web Application Availability Monitoring Wizard in the Authoring pane of the console is where you start Global Service Monitor by choosing to monitor from external locations, which gives you a list of all the available points of presence from Microsoft.

Visual Studio Web Test monitoring is the other option for Global Service Monitor. It gives you the ability to import more extensive Global Service Monitor web tests that have been created using Visual Studio. Using Visual Studio Web Test monitoring, you can record actions to take against your external-facing applications and validate against multiple criteria and multiple websites at the same time. With Visual Studio Web Test monitoring, you can import a

web test that was built by a developer using Visual Studio. Transactions are supported, as are authentication actions.

The results of all the web tests come into the Operations Manager console as alerts and views that allow you to monitor the test results. You can also use Health Explorer to see more details about the health status of a web application availability test running against a URL from a particular location.

Furthermore, there is a summary dashboard, showing a map with the locations that you have selected to monitor from. For a deeper view, there is also a detailed dashboard showing six key metrics of location and the tests in that location that you can zoom in to. This detailed dashboard is available only for Web Application Availability monitoring. It is not available for Visual Studio Web Test monitoring.

## Troubleshooting Global Service Monitor

Some common difficulties in configuring Global Service Monitor include the following:

- Global Service Monitor needs a proxy or a server that has ports opened to the Internet. If your proxy server needs authentication, you will need to follow the steps described in the Microsoft Knowledge Base article at <http://support.microsoft.com/kb/2900136/en-us>. From a security perspective, everything that is sent over the Internet is encrypted and is also stored encrypted on the Microsoft Azure watcher nodes that are managed by Microsoft.
- To upgrade from the 90-day trial subscription to the full Software Assurance benefit, go to the same portal where you requested the OrgId and the trial subscription. Click the Global Service Monitor tile and select Use As Software Assurance Benefit.
- Even in the full subscription, there is a limit of 25 web tests. This can be changed by contacting Microsoft support. A link for contacting support can also be found on the same portal. When you contact support, you will need your Software Assurance ID for the support request form. More test maximums and locations are explained at <https://technet.microsoft.com/library/jj860368.aspx>.
- To monitor an internal service that is accessible only from the local network, you shouldn't use Global Service Monitor. Instead, specify an internal watcher node in the Web Application Availability Monitoring Wizard in the Authoring pane of the Operations Manager console, where you create your Global Service Monitor web test. An internal watcher node can be any Operations Manager agent or a dedicated resource pool for URL monitoring. These internal watcher nodes can reach and monitor the internal-only exposed endpoints.
- You can create a map view of all the agents, whether they are internal agents or Global Service Monitor external agents. Internal agents can be added to the map using Windows PowerShell, which is explained in the blog post at [http://blogs.technet.com/b/marcin\\_jastrzebski/archive/2012/06/01/3499135.aspx](http://blogs.technet.com/b/marcin_jastrzebski/archive/2012/06/01/3499135.aspx).

- When your management servers need to connect to the Internet using a proxy server, and that proxy server needs authentication, you need to create or add the required Run As account to the Global Service Monitor Run As profile. By default, proxy access is unauthenticated.
- If you get an error stating “Failed to discover Global Service Monitor locations,” make sure you have KB931125 installed, which updates the Microsoft root certificates. More information can be found on Mihai Sarbulescu’s blog at <http://blogs.technet.com/b/mihai/archive/2013/10/10/global-service-monitor-not-working-because-the-remote-certificate-is-invalid-according-to-the-validation-procedure.aspx>.

## Using Azure Operational Insights (System Center Advisor)

---

This section explains what Azure Operational Insights is and how you can use it to enhance the monitoring capabilities of Operations Manager. System Center Advisor is now part of the new Azure Operational Insights (existing data from System Center Advisor is automatically migrated to Azure Operational Insights) and can be found at <https://preview.opinsights.azure.com/>. General information about Azure Operational Insights is available at <http://msdn.microsoft.com/en-us/library/azure/dn884662.aspx>.

## Understanding Azure Operational Insights

Azure Operational Insights is a cloud-based service that can also be an extension to Operations Manager, similar to Global Service Monitor. Data is sent to the cloud continuously from the Operations Manager management server or directly from the Microsoft Monitoring Agent (MMA).

Certain tasks are difficult to accomplish or not available in the on-premises version of Operations Manager, and that is where Azure Operational Insights provides additional functionality using intelligence packs (IPs). IPs function like management packs in System Center Operations Manager or integration packs in System Center Orchestrator.

You can use direct agents, connecting to Azure Operational Insights directly, or you can connect your Operations Manager management group to Azure Operational Insights. When you connect Operations Manager, the data is sent from the Operations Manager management server, and the generated alerts from Azure Operational Insights are visible in the Operations Manager console. With the Alert Management IP, you can also see all of your Operations Manager alerts in Azure Operational Insights. By default, this IP shows you the top critical and warning alerts, which is useful for tuning.

Using Azure Operational Insights, you can collect, combine, correlate, and visualize all of your machine data. The combination and correlation logic is built into Azure Operational Insights. This logic is not present in Operations Manager. With some simple queries, you can

view the collected alert data from Operations Manager in a more logical way. Azure Operational Insights collects additional data for the numerous new scenarios that are enabled. Azure Operational Insights also sorts through millions of lines of data faster than Operations Manager. This helps find the root cause of operational issues more quickly and easily.

Azure Operational Insights does not gather the same or existing data that you are already collecting in Operations Manager. IPs and the Log Management collection policy define new and different data to be gathered and sent to the service. This data is sent directly to the cloud without persisting it in the on-premises Operations Manager database. This is particularly important for scenarios such as log management or security event investigations where the volume of data is very high, and the on-premises databases traditionally became the performance bottleneck. Azure Operational Insights was built with a scalable backend in mind to allow those new high-volume monitoring scenarios that were not that easy to accomplish with Operations Manager.

The only exception to this is some configuration data that is pre-existing in Operations Manager through System Center Virtual Machine Manager, which is used by the Capacity IP, and the Alert Management IP, which synchronizes existing Operations Manager alerts. Operations Manager has limited capacity planning capabilities. The necessary data is in the Data Warehouse Database, but you need specialized reports (that are not available out of the box) to start planning for capacity, such as storage space. With Azure Operational Insights, you can see where your virtual machine infrastructure needs more resources and where it's under-utilized. You can also use "what-if" scenarios to enhance your planning options.

**NOTE** For the Capacity IP, the connection between Operations Manager and Virtual Machine Manager must be configured.

Azure Operational Insights also analyzes the systems that you have enabled for the latest updates. This includes security updates as well as missing hotfixes or service packs. This functionality is not present in Operations Manager. In Azure Operational Insights, this functionality is provided by the System Update Assessment IP. Additionally, using the Malware Assessment IP, you can verify that the necessary malware protection is present on your enabled systems. This functionality is also not present in Operations Manager.

An important part of troubleshooting is knowing when a configuration change happens to a system, such as newly added software or hardware. These changes are tracked in Azure Operational Insights with the Change Tracking IP and can help you pinpoint the root cause of a problem. For instance, installing a certain hotfix might cause another application to function erratically. Operations Manager does not have this functionality.

Certain combinations of software or versions can cause issues. When these kinds of problems are detected by the Microsoft support teams, new rules are created in Azure Operational Insights so that your configuration is checked for these combinations. In this way, known problems can be avoided, and you will be warned, for instance, that you need to install a certain service pack to avoid the problem. These rules are maintained by Microsoft and



function in the same way as a separate management pack in Operations Manager. However, you do not need to update this management pack yourself; updates are handled by Azure Operational Insights and the Microsoft support teams.

Another added functionality is a smartphone application where you can see the current status of your environment. Using third-party solutions, you could have a similar functionality for your Operations Manager environment. With the Alert Management IP, you can see your Operations Manager alerts in Azure Operational Insights and thus on your phone also. This is called Operational Insights on the Go.

Using the Log Management IP, you can easily analyze large log files. This is especially useful for W3C format Internet Information Services (IIS) logs. All event logs can be collected, except the Security event log. For the Security event log, there is a separate IP. Operations Manager also collects events from event logs, but only when a management pack has a rule to actually grab that event. With Azure Operational Insights, all events from all selected event logs are captured. In this way, you might see other events that need investigation that were not captured using Operations Manager.

The SQL Assessment IP gives you a continuously updated view on the health of your SQL Server servers. This includes security and compliance, availability and business continuity, performance and scalability, upgrade migration and deployment, operations and monitoring, and change and configuration management information. The SQL management pack in Operations Manager can give you an already extensive view of the health of your SQL server environment, but this SQL Assessment IP goes much deeper than that.

## Connecting to Azure Operational Insights

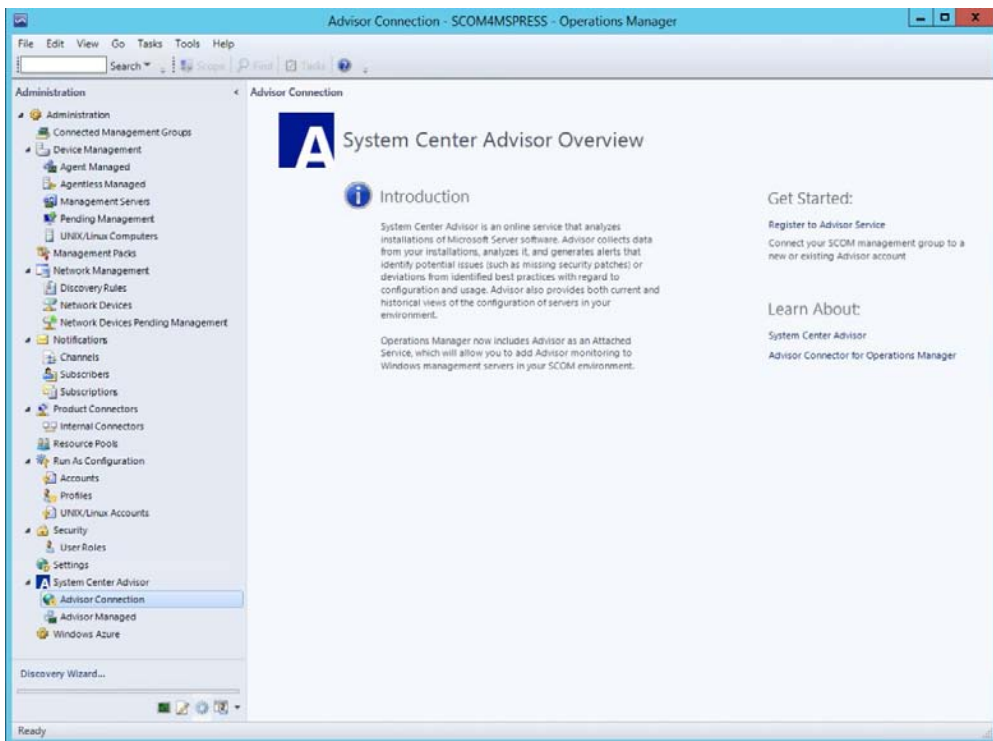
Following the instructions at <https://preview.opinsights.azure.com/instructions>, there are three possible ways to connect to Azure Operational Insights:

- By using directly connected agents that have the Microsoft Monitoring Agent (MMA) installed.
- By using Operations Manager (where the MMA is used as the Operations Manager agent). This gives you integration between Operations Manager and Azure Operational Insights so that you can see the alerts that Azure Operational Insights generates directly in the Operations Manager console.
- By connecting Azure Operational Insights to an Azure Storage Account.


To connect Azure Operational Insights with Operations Manager, you need to open the Operations Manager console and click the Administration tab. If you have not installed Update Rollup 5 (UR5) or later for Operations Manager 2012 R2, on the left side of the screen, select System Center Advisor, and on the right side, select Register To Advisor Service. If you have installed the latest Update Rollup (UR5 or later), select Operational Insights on the left side of the screen and click Operational Insights Connection. Figure 5-1A shows how to connect Operations Manager to Azure Operational Insights prior to Operations Manager 2012 R2 UR5,

while Figure 1B shows how to do this when using Operations Manager 2012 R2 UR5 or Operations Manager 2012 SP1 UR9 or later.

If this procedure fails, check if your server clock (time) is at the correct time and that there's no firewall or proxy blocking the communication. The communication paths that should be opened can be found at <http://blogs.technet.com/b/momteam/archive/2014/05/29/advisor-error-3000-unable-to-register-to-the-advisor-service-amp-onboarding-troubleshooting-steps.aspx>. If a proxy server is needed for communication, you can configure one in the console under System Center Advisor, Advisor Connection (before UR5) or Operational Insights, Operational Insights Connection (UR5 or later). If credentials are needed for this proxy, they can be specified in the console in the Administration pane, click Run As Configuration, click Profiles, and there find an account called System Center Advisor Run As Profile Proxy.



**FIGURE 5-1A** Connecting Operations Manager to Azure Operational Insights (before Operations Manager 2012 R2 UR5)



**FIGURE 5-1B** Connecting Operations Manager to Azure Operational Insights (Operations Manager 2012 R2 UR5 or Operations Manager 2012 SP1 UR9 or later)

When you click Register To Advisor Service or Operational Insights, a new window opens and asks for your credentials. These can be either your Microsoft account (such as *yourname@outlook.com*) or your organizational account. After that, you need to name the new Operational Insights account or select an existing one to continue. This can be changed later if necessary by selecting Re-configure Advisor (Operational Insights) from the Advisor (Operational Insights) Connection menu in the Operations Manager console.

When your Operations Manager environment is connected to the Advisor (Azure Operational Insights) service, select the systems to be monitored using this cloud-based service. You can also manage the alert rules. By default, all rules from the service are enabled. This means that, out of the box, Azure Operational Insights checks your environment for Windows operating system, domain controller, Exchange, Hyper-V, Lync, SharePoint, and SQL Server configuration settings. As with new management packs, do not enable all rules on all of your systems at once—start small, with a few systems, and work your way up to the rest of the environment. Also, test in a test environment before you enable the new rules in your production environment.

To add systems to Azure Operational Insights, click Add A Computer/Group in the Administration pane, under System Center Advisor, Advisor Managed. At the time of writing, the Operations Manager console still shows System Center Advisor, but depending on the installed update rollup, this could be renamed. It might take a while before the newly added systems start to be monitored. To keep track of that in the Monitoring pane, click System Center Advisor and then click Advisor Health State, which shows the active state of the management server and the Azure Operational Insights connected systems. Under Active Alerts, you can see the alerts from your enabled systems.

## Planned additional capabilities of Azure Operational Insights

Since Azure Operational Insights is still in preview at the time of writing, additional capabilities are continuously being added. You can keep track of the changes that are planned on the UserVoice page at <http://feedback.azure.com/forums/267889-azure-operational-insights>, on Twitter at @Oplnsights, on the Operations Manager blog at <http://blogs.technet.com/b/momteam/>, or on Daniele Muscetta's blog at <http://blogs.msdn.com/b/dmuscett/>. This last one is the most useful, and you can find a complete walkthrough series there on how to do searches in Azure Operational Insights. A query cheat sheet was published by Stefan Roth at <https://gallery.technet.microsoft.com/Azure-Operational-Insights-de86f8fe>.

One of the upcoming IPs is the Active Directory Assessment IP. This one is similar to the SQL Assessment IP, and it will give you a complete overview of the health and risks within your Active Directory environment. This is not available as Operations Manager functionality.

Another IP is the Security IP. This IP will complement the Log Management IP, capturing the security event log and firewall logs. Within Operations Manager, this functionality is partly available with the optional Audit Collection Services (ACS). Since ACS in most cases is a separate environment that is large and difficult to manage, the Security IP removes much of the complexity from the on-premises product and adds capabilities, such as firewall log monitoring.

## Comparing Azure Operational Insights with Operations Manager

Table 5-1 provides an overview of the capabilities offered out of the box by Azure Operational Insights and Operations Manager. Operations Manager is able to offer many of the same solutions, but only with custom development.

**TABLE 5-1** Comparison of out-of-the-box capabilities of Azure Operational Insights and Operations Manager

CAPABILITY	SYSTEM CENTER 2012 R2 OPERATIONS MANAGER	MICROSOFT AZURE OPERATIONAL INSIGHTS
<b>Operational Visibility and Management</b>		
Proactive smart alerts	X	X
Comprehensive operations dashboards and reporting	X	X
Real-time and customizable monitoring	X	X
Customizable dashboards	X	X

<b>Performance Monitoring and Analytics</b>		
Monitoring of OS Resources (CPU, disk, memory, network) for Windows and Linux systems	X	X
On-premises server performance monitoring	X	X
Azure server performance monitoring		X
SAN Storage analytics		X
<b>Capacity Planning</b>		
Forecast resource utilization trends	X	X
Optimize virtual machine placement, investigate “what-if” scenarios, pinpoint capacity shortages, identify stale and over-allocated VMs		X
Identify storage bottlenecks		X
<b>Configuration and Change Tracking</b>		
Detect potential configuration issues or deviations from identified best practices		X
Monitor software, Windows Services, Registry keys, Group Policy, and file changes		X
<b>Security</b>		
Security log collection	X	X
Breach and threat detection		X
Deep forensic analysis		X
Malware detection and software update status		X
<b>Log Management</b>		
Universal log collection and analysis		X
Unlimited data retention		X
Adding structure to all types of unstructured data		X

Real-time monitoring, search, and log analytics		X
Dashboards powered by search queries		
Third-party and community-based intelligence packs		

## Installing Operations Manager in Azure IaaS

---

This section shows you how to deploy Operations Manager in Microsoft Azure Infrastructure as a Service (IaaS). It also describes best practice configuration steps for getting the best performance out of your virtual machines in the Azure environment.

### Supported scenarios

Running Operations Manager on Azure virtual machines is basically the same as running Operations Manager in an on-premises environment. When you first log in to Azure IaaS, it is like coming to a new datacenter: you have the racks with host machines for your virtual machines, power supplies are connected, storage is available, and there are network switches in the racks. Installing software, configuring everything, and creating virtual machines is up to you.

System Center 2012 R2 Operations Manager runs on virtual machines in Microsoft Azure just as it does on physical computer systems. It does require specific configurations, however. This includes network and storage setup, just as you would configure in a physical datacenter.

Microsoft recommends using Operations Manager with Azure virtual machines in the following scenarios:

- Scenario 1: You can run Operations Manager on a Azure virtual machine and use it to monitor other Azure virtual machines.
- Scenario 2: You can run Operations Manager on a Azure virtual machine and use it to monitor instances that are not running on Azure.
- Scenario 3: You can run Operations Manager on-premises and use it to monitor Azure virtual machines.

Microsoft tested Operations Manager by installing and using it in an Azure virtual machine. The standard sizing and supported configuration for Operations Manager applies to Azure virtual machines.

**TIP** You can get an Azure subscription for free, as a trial account, or you can use the pay-as-you-go model, your enterprise agreement (EA), the Open Volume License Program, or you can use the subscription that comes with MSDN, MPN, or BizSpark. In the case of a trial subscription, if it expires, your virtual machines will be gone, but the virtual hard disks (VHDs) will still be available. If you need to add more administrators to a subscription, you can do that in the Settings screen in the Azure portal, where you can add them in the Administrators field.

Many actions that you need to take in Azure can be performed much easier and faster using Windows PowerShell with the Microsoft Azure PowerShell module. This is an additional module for the already installed Windows PowerShell on your system and can be downloaded and installed from <http://azure.microsoft.com/en-us/documentation/articles/install-configure-powershell/>. The easiest way to connect Windows PowerShell to your Azure subscription is to use the Azure AD method by typing `Add-AzureAccount`. The Azure AD method is the recommended authentication method since it makes it easier to manage access to a subscription. It works with the Azure Resource Manager API as well. After typing `Add-AzureAccount`, in the authentication window that opens, enter your account and password to connect to the Azure subscription.

**See also** *The complete Azure cmdlet reference can be found at <http://msdn.microsoft.com/en-us/library/azure/jj554330.aspx>. Using `(Get-module -Name Azure).Version`, you can check if you have the latest version installed, since this module gets updated often*

## Configuring the network

The first step in configuring your Azure datacenter is setting up the network. This can be done using the management portal at <https://manage.windowsazure.com/> or the preview portal at <https://portal.azure.com/> or by using Windows PowerShell. You need to create a virtual network subnet so that the virtual machines can communicate with each other.

**IMPORTANT** Azure uses DHCP to assign IP addresses to the virtual machines you create. This should *not* be changed to a fixed address since that will make your virtual machine inaccessible.

The first machine to be started in the address range you specify will have an internal IP address of x.y.z.4. The addresses x.y.z.1 to x.y.z.3 are reserved for Azure internal usage. So if you use a subnet with a classless inter-domain routing (CIDR) notation of /29, which includes eight IP addresses, there's already three fewer IP addresses. Since you need Active Directory to set up an Operations Manager environment, you will need to install a domain controller. This domain controller will also have a DHCP-assigned IP address. To make sure the domain controller keeps the IP address, there are several options. One option is to always start this

virtual machine first in your network space so that it gets the address x.y.z.4 again. You can also use Windows PowerShell to assign a static IP address with Set-AzureStaticVNetIP.

Your virtual machines in a cloud service will also have an external IP address, a virtual IP address (VIP), so that they can be accessed from outside the Azure environment and for Internet connection. Since this VIP is accessible to anyone on the Internet, you need to put a strong administrator password on your virtual machine. The management portal will block you from using "Administrator: as the username or "P@ssw0rd" as the password. The VIP of a cloud service will not change as long as it has a provisioned virtual machine within the cloud service. One way to ensure the VIP is never lost is to keep a virtual machine provisioned at all times. To keep the cost down, you can use the smallest virtual machine for your domain controller and keep that one running at all times. However, Microsoft has created another option: the ability to reserve VIPs for an Azure subscription. These VIPs can then be used with cloud services. In addition, the VIPs are kept for the lifetime of your Azure subscription.

You can use Windows PowerShell to create a reserved IP address. For example, to reserve an IP address, you can use the following command:

```
New-AzureReservedIP -ReservedIPName "YourVIP" -Label "YourVIPLabel" -Location  
"YourLocation"
```

**See also** *Further explanations about the different types of IP addresses can be found in the blog post at [http://blogs.msdn.com/b/lalitesh\\_kumar/archive/2014/10/06/static-ip-reserved-ip-and-instance-level-ip-in-azure.aspx](http://blogs.msdn.com/b/lalitesh_kumar/archive/2014/10/06/static-ip-reserved-ip-and-instance-level-ip-in-azure.aspx).*

To further protect your virtual machine, you can put access control lists (ACLs) on the endpoints. These are the endpoints exposed to the outside world for your virtual machine. You can also use a jump box configuration: one virtual machine that has the Windows PowerShell and Remote Desktop Protocol (RDP) endpoints enabled and has a very strong password and ACLs set. All the other virtual machines are then accessible only from this jump box virtual machine since they don't have enabled endpoints.

**See also** *Guidelines for Deploying Windows Server Active Directory on Azure Virtual Machines can be found at <http://msdn.microsoft.com/en-us/library/azure/jj156090.aspx>.*

It's easy to create a network range for your virtual machines. In the management portal, click Networks in the left pane, and at the bottom left of the screen, click New to launch the Create a Virtual Network Wizard. You can use the Quick Create or Custom Create option. For example, using the Custom Create option, on the first page, specify the name of the network and its location (which datacenter). On the second page, specify the DNS server address if you want to use your own DNS server. You can also select the option to create a Site-to-Site (S2S) VPN or a Point-to-Site (P2S) VPN. (This can also be configured later.) On the third and final page, shown in Figure 5-2, enter the address range of your network and the subnet you want to create in this address range.

Alternatively, you can do the same thing using Windows PowerShell. You can check for existing virtual networks in the selected subscription with this command:

```
Get-AzureVNetConfig -ExportToFile "$ScriptPath\NetCfg.xml"
```



This saves the current virtual networks to an XML file. Use Set-AzureVNetConfig to create the virtual network using an XML configuration file.

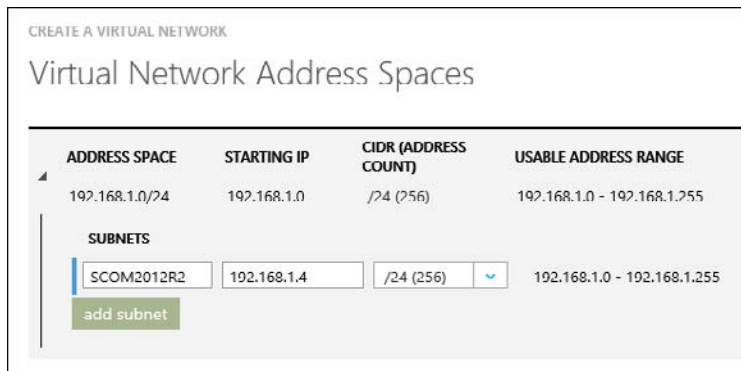


FIGURE 5-2 Custom create a network in Azure

## Configuring storage

As with a physical datacenter, after configuring the network, you need to configure your storage. The number of Input/Output Operations per Second (IOPS) and the number of disks are limited in Azure storage accounts. The IOPS maximum at the time of writing is 20,000 per storage account. There is also an IOPS limit per disk: 300 IOPS for a basic disk and 500 IOPS for a standard disk. Premium storage features much higher performance figures, but of course it costs more.

**See also** *The current limits can be found at <http://azure.microsoft.com/en-us/documentation/articles/azure-subscription-service-limits/#subscription>.*

Since there is an IOPS limit per disk and per storage account, you shouldn't put too many disks in one storage account. If you do, you might configure your virtual machines with enough disks to attain the performance level you want, but reach the storage account performance limit, meaning all the disks in that storage account will throttle down. You can set up monitoring for this throttling condition, according to <http://blogs.msdn.com/b/mast/archive/2014/08/02/how-to-monitor-for-storage-account-throttling.aspx>. To work around this limitation, you can have more than one storage account and, for instance, use a separate storage account per virtual machine. The storage account limit at the time of this writing is 100.

**TIP** To check the number of disks in a storage account, run the following Azure PowerShell command:

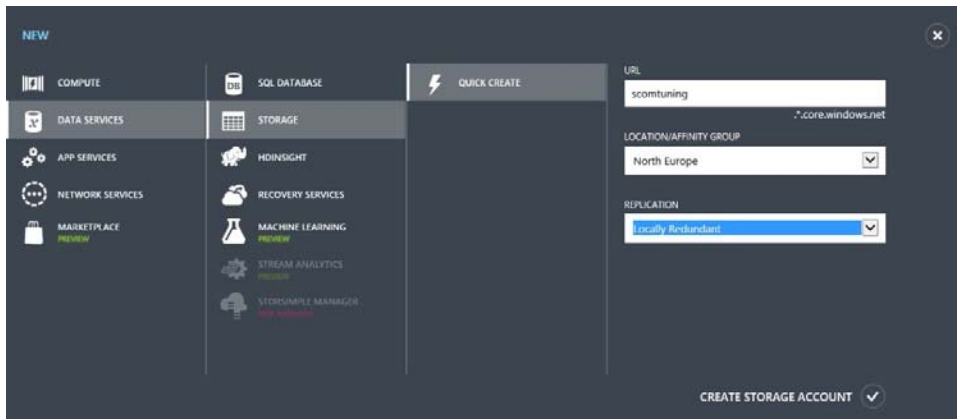
```
Get-AzureDisk | Where-Object { $_.AttachedTo } | Group-Object  
{$_ .Medialink.Host.Split('.')[0]} -NoElement
```

When you create a storage account, you select the datacenter for the storage and indicate if the storage will be geo-replicated or local to the selected datacenter. Geo-replication is the default, which means your storage will be in the local datacenter and replicated to another datacenter at least 500 kilometers away. To keep the cost and the latency as low as possible, you can change this to local replication only. Not replicating the data to a distant location also keeps the latency of your storage lower. Even with local replication, you still have three copies of your data, so it is protected against hardware failures.

**See also** More information about storage accounts can be found at <http://azure.microsoft.com/en-us/documentation/articles/storage-what-is-account/> and further storage account performance optimization is discussed at <http://blogs.msdn.com/b/mast/archive/2014/10/14/configuring-azure-virtual-machines-for-optimal-storage-performance.aspx> and at <http://azure.microsoft.com/en-us/documentation/articles/storage-performance-checklist/>.

To create a storage account, click Storage in the left pane of the management portal, and click New at the bottom of the window. Enter an account name, which can contain only lower case characters, and select the location and the type of replication. In the example shown in Figure 5-3, North Europe is selected as the location and locally redundant replication is indicated instead of geo-replicated replication.

**See also** The complete overview of all virtual machine and cloud service sizes can be found at <http://msdn.microsoft.com/en-us/library/azure/dn197896.aspx> or by typing `Get-AzureRoleSize` in an Azure PowerShell window.



**FIGURE 5-3** Storage account creation options

One of the (many) tools for managing your storage account is Azure Storage Explorer, which can be found at CodePlex at <http://azurestorageexplorer.codeplex.com/>. To connect this (or any other) tool to your Azure subscription's storage accounts, you need to get the storage account key from your subscription. In the management portal, click Storage in the left pane, select the storage account from which to retrieve the access key, and then click Manage Access

Keys at the bottom of the window to open the Manage Access Keys dialog box, shown in Figure 5-4.

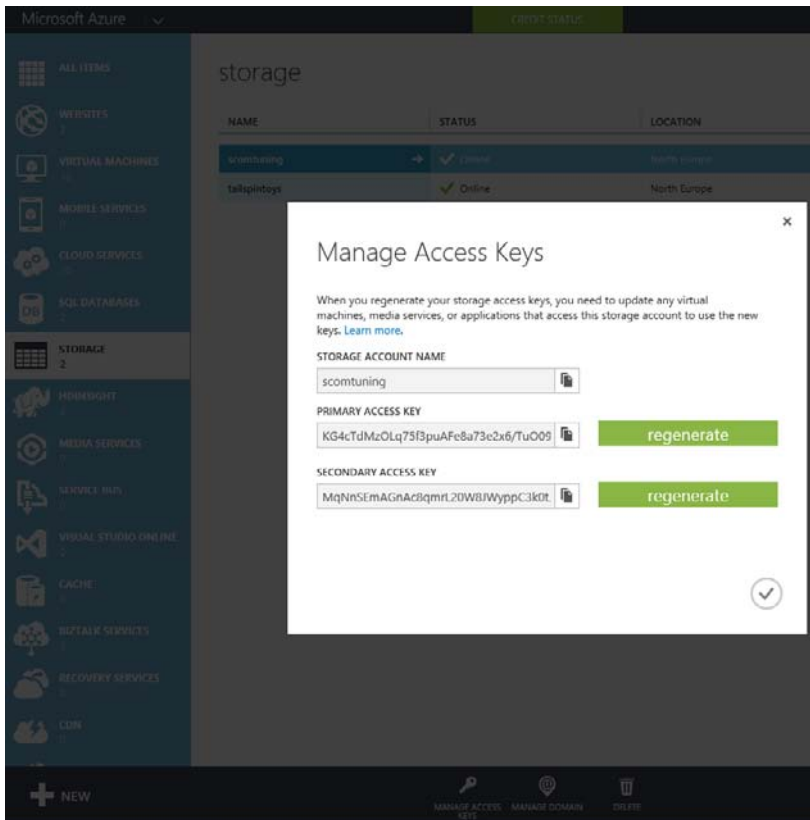
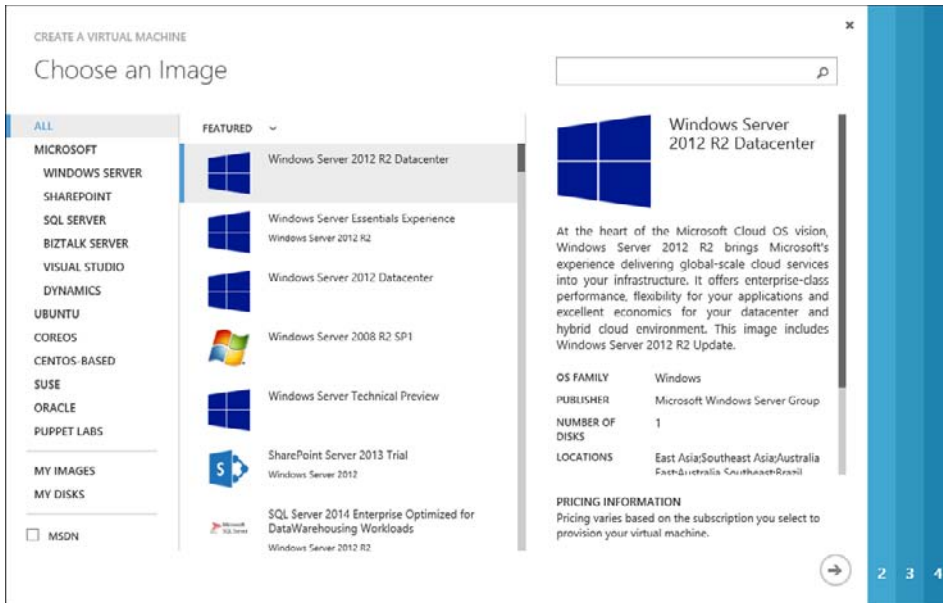


FIGURE 5-4 Manage Access Keys dialog box

## The first virtual machine

After creating the virtual network and the storage account, you can start creating virtual machines. To create a virtual machine, select Virtual Machines in the management portal, and click New at the bottom of the window. This opens the Create a Virtual Machine Wizard. You can select the Quick Create option or the From Gallery option. In the example in Figure 5-5, the From Gallery option was selected and then the Windows Server 2012 R2 Datacenter image is shown as selected.



**FIGURE 5-5** Choose an image from the gallery

A virtual machine can have five different states, as described in Table 5-2. When you shut down a virtual machine from Windows, it goes to Stopped state. If you shut down a virtual machine from the Azure management portal, it goes to Stopped (Deallocated) state. It is important to shut down the virtual machine from the Azure management portal to make sure the virtual machine is deallocated and not billed.

**TABLE 5-2** Virtual machine state and billing

STATE	BILLED	DETAILS
Starting	Yes	The initial starting state of VMs as they go through the boot cycle.
Running (Started)	Yes	The running state of the VM.
Stopped (Allocated)	Yes	For allocated stopped (from within the machine) VMs, the cores are still billed.
Stopped (Deallocated)	No	The VM is stopped (from the portal or Windows PowerShell) and un-mounted from the host.
Deleted	No	The VM has been deleted and is no longer occupying cores.

After choosing the template for your virtual machine, give it a name, select the size of the machine, and enter the administrator account and password. Since the first virtual machine is a domain controller, it can be small, such as the A0 size (with a shared CPU core and 768 MB of memory) from the basic tier. The smaller the machine, the less it costs. The disks in the basic tier are limited to 300 IOPS each; in the standard tier, the limit is 500 IOPS.

Next, select the network and storage options for the virtual machine. For a highly available setup, also select the option to create an availability set. You need at least two virtual machines in an availability set. Without an availability set, no Microsoft Azure SLA is applicable.

**See also** *More information about availability sets can be found at <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-manage-availability/>.*

The screenshot displays the 'Virtual machine configuration' page in the Azure portal. The page is titled 'CREATE A VIRTUAL MACHINE' and 'Virtual machine configuration'. It features several configuration options:

- CLOUD SERVICE:** A dropdown menu with 'Create a new cloud service' selected.
- CLOUD SERVICE DNS NAME:** A text input field containing 'MSPRESSDC01' and a domain suffix dropdown set to '.cloudapp.net'.
- REGION/AFFINITY GROUP/VIRTUAL NETWORK:** A dropdown menu with 'SCOM\_Network' selected.
- VIRTUAL NETWORK SUBNETS:** A dropdown menu with 'SCOM2012R2(192.168.1.4/24)' selected.
- STORAGE ACCOUNT:** A dropdown menu with 'scomtuning' selected.
- AVAILABILITY SET:** A dropdown menu with '(None)' selected.
- ENDPOINTS:** A table with columns for NAME, PROTOCOL, PUBLIC PORT, and PRIVATE PORT.
 

NAME	PROTOCOL	PUBLIC PORT	PRIVATE PORT
Remote Desktop	TCP	AUTO	3389
PowerShell	TCP	5986	5986

On the right side, there is a preview of the 'Windows Server 2012 R2 Datacenter' image. The preview includes a description: 'At the heart of the Microsoft Cloud OS vision, Windows Server 2012 R2 brings Microsoft's experience delivering global-scale cloud services into your infrastructure. It offers enterprise-class performance, flexibility for your applications and excellent economics for your datacenter and hybrid cloud environment. This image includes Windows Server 2012 R2 Update.' Below the description, it lists: OS FAMILY: Windows; PUBLISHER: Microsoft Windows Server Group; NUMBER OF DISKS: 1. At the bottom right of the preview, there is a 'PRICING INFORMATION' section stating: 'Pricing varies based on the subscription you select to provision your virtual machine.'

**FIGURE 5-6** Virtual machine configuration options

The first option you select is the cloud service. The cloud service is the name of the virtual machine that you can connect to from the Internet. In a load balanced scenario, such as with an IIS Server web farm, there is a single cloud service with multiple machines. In this example, a new cloud service is created for this virtual machine. This gives you the option to connect directly to it from the Internet.

**See also** *More information on cloud services can be found at <http://azure.microsoft.com/en-us/documentation/articles/cloud-services-what-is/>.*

The name you entered when you created the virtual machine automatically appears as the name of the cloud service, but you can change it on this page. This will be the name you use to connect to the cloud service later. The network and subnet created earlier in the chapter are

used for the virtual machine in this example. You can select a previously created network or subnet from the top of the drop-down lists.

A storage account was also created earlier in this chapter. This storage account is selected for this virtual machine. In this example, an availability set is not selected or created. By default, there are two exposed endpoints, one for Remote Desktop so you can use it to connect to the virtual machine and one for PowerShell so you can use Windows PowerShell to manage the virtual machine remotely. You can also add or remove endpoints on this page.

The next step is to configure the VM extensions. The VM agent is installed by default, and you can add extensions such as Puppet or Chef, or security extensions such as Microsoft Antimalware. Click the check mark to create the virtual machine you just configured.

**See also** *The complete documentation for creating a virtual machine in Azure running Windows, for both the current and the preview Azure management portal, can be found at <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-windows-tutorial/>.*

Of course, you can also create a virtual machine using only Windows PowerShell. The first step is to get the list of images that are available (the templates to create your virtual machine from) using the following command:

```
Get-AzureVMImage | Select ImageName
```

After you select the image to use, create the new Azure virtual machine configuration object using the following command:

```
New-AzureVMConfig -Name $VMName -InstanceSize $SizeOfTheVM -Image $ImageName
```

The user name and password can be set using the following command:

```
Add-AzureProvisioningConfig -Windows -AdminUserName $AdminUser -Password $AdminPwd
```

You can also set other options such as the endpoints for Windows PowerShell and Remote Desktop or the time zone settings using the Add-AzureProvisioningConfig cmdlet.

When all the settings are configured, create the actual virtual machine using the following command:

```
New-AzureVM -ServiceName $CloudService -Location $Location
```

You can also set `-AffinityGroup` instead of `-Location`.

A quicker way to create the virtual machine with fewer configuration options, similar to the quick create option in the management portal, is to use the New-AzureQuickVM Windows PowerShell cmdlet. In the script center at <http://azure.microsoft.com/en-us/documentation/scripts/>, you can find more scripts that can automatically create virtual machines with multiple data disks.

Alternatively, you can use your existing Hyper-V images by uploading them to your Azure storage account. This means you run the Sysprep tool on your existing virtual machine image and upload the processed VHD file using the following command:

```
Add-AzureVhd -Destination "<BlobStorageURL>/<YourImagesFolder>/<VHDName>.vhd" -  
LocalFilePath <PathToVHDFile>
```

Then you add the VHD as an image that you can use in your Azure subscription, just as you would use the Microsoft-provided images, with the following command:

```
Add-AzureVMImage -ImageName $MyImageName -MediaLocation  
"https://yourstorage.blob.core.windows.net/vhds/$MyImageName.vhd"
```

**See also** *A more extensive explanation of how to create and upload Windows Server VHD files to Microsoft Azure is here <http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-create-upload-vhd-windows-server/>.*

## Configuring the first virtual machine as a domain controller

After creating the first virtual machine, you can configure it as a domain controller. By default, a virtual machine in Azure has an operating system disk C and a temporary disk D

**NOTE** The temporary disk, as the name says, is really temporary. Don't place anything on this disk, apart from the page file for the operating system. In case of a reboot, this disk will be recreated, and everything on it will be lost. Microsoft Support cannot help you restore any data lost from this temporary drive.

The operating system disk has read and write caching enabled. This is the best setting for the operating system, but for the Active Directory database (NTDS.DIT) and log files, it is recommended that you attach another disk with only read caching enabled to the virtual machine and put the Active Directory files on that drive. You can add a disk to a virtual machine using the management portal or by using the Windows PowerShell command `Add-AzureDataDisk`.

Using the management portal, you can select a recently created the virtual machine, and then select `Attach and Attach Empty Disk` from the bottom of the page. Indicate the size of the new disk and the cache preference. For this example, since the disk will contain the Active Directory database and log files, the `Read Only` option is preferred.

**See also** *More information about adding data disks can be found at <http://azure.microsoft.com/en-us/documentation/articles/storage-windows-attach-disk/>.*

To connect to the virtual machine that was just created, `Connect` in the management portal. This generates and downloads a Remote Desktop file that you can use to connect to the virtual machine in the same way as you would connect to a (virtual) machine in an on-premises datacenter.

When a disk is added, it needs to be initialized and formatted before it can be used. Then you can promote this server to a domain controller and choose the newly attached and formatted drive to store your Active Directory files.

The internal (to Azure) IP address of this virtual machine corresponds to the address space you chose; for example, if 192.168.1.0 is the address space, the IP address for the virtual

machine is 192.168.1.4. You can enter that internal IP address as the DNS server address in your virtual network. To do this, click Networks in the management portal, select your virtual network, and click Configure. Provide the name of the DNS server and the IP address.

To make sure the virtual machine keeps this IP address, which would be best since this virtual machine will be your domain controller and DNS server, there are several options:

- You can keep this virtual machine running at all times, which costs money. If you set the size of this domain controller virtual machine to Basic-Small, then the cost is minimal.
- You can set this IP address as a static address using Set-AzureStaticVNetIP. This is possible for a virtual machine that you are creating as well as for an existing virtual machine. This static address reservation also costs money, although it is less expensive than having the virtual machine running constantly.
- You can, if all the machines are in the same subnet, make sure that you start them in exactly the same order every time. This way, the first virtual machine that is started will always get x.y.z.4 as its address, the next one will get x.y.z.5, and so on. This costs nothing, but you need to make sure not to make any mistake.

**See also** *Static internal IP addresses are explained at <http://msdn.microsoft.com/en-us/library/azure/dn630228.aspx>.*

## Configuring the next virtual machine as the SQL Server server

The next server you configure is the SQL Server server that will be used for the Operational Database, the Data Warehouse Database, and SQL Server Reporting Services (SSRS). The procedure is the same as for the first virtual machine.

You can create a new virtual machine with only the operating system image and install SQL Server yourself, or you can choose an image with SQL Server already installed. You cannot use the Azure SQL databases because they do not conform to the requirements for Operations Manager. The SQL Server software can be uploaded in a VHD to the Azure blob storage, where you can connect this VHD to every virtual machine that needs the software. Stefan Stranger describes how to do this in his blog article at [http://blogs.technet.com/b/stefan\\_stranger/archive/2015/02/05/opsmgr-installation-in-azure-vm.aspx](http://blogs.technet.com/b/stefan_stranger/archive/2015/02/05/opsmgr-installation-in-azure-vm.aspx).

If you have an MSDN subscription, you can also download the software directly from MSDN. Since the Azure datacenters and the MSDN download servers are very well connected, this download goes very fast.

To get the best performance from a storage perspective, add multiple data disks and spread the load over all of them. If you would like to use the temporary D drive to store the TempDB database, you need to use a startup script to re-create the folders to hold the TempDB at machine startup time. If you don't do that, the folders will not exist when the machine starts up. Performance guidance for SQL Server in Azure virtual machines can be found at <http://msdn.microsoft.com/en-us/library/azure/dn248436.aspx>.



The image that includes SQL Server comes with many pre-configured options, such as SQL Server Analysis Services. You can integrate System Center Virtual Machine Manager with SQL Server Analysis Service (SSAS) to provide forecasting reports. In that case, you need to make sure that SSAS is installed on the Operations Manager Reporting server. The full list of installed features on the image that includes SQL Server can be found at <http://msdn.microsoft.com/library/azure/dn133151.aspx>. When you want to use this image with just Operations Manager, many of these optional features are not needed, and you can safely uninstall them. Further configuration, such as configuring Reporting Services and starting the SQL Agent, is also required. You also need to configure the firewall so that the Operations Manager management server installation can contact the SQL Server.

When you add the SQL Server virtual machine to the same virtual network as the domain controller, both machines can see each other. This new virtual machine can be joined to the domain that was created when you installed the domain controller.

## Configuring the final virtual machine as the Operations Manager management server

Microsoft tested Operations Manager by installing and using it in an Azure virtual machine. The standard sizing and supported configuration for Operations Manager applies to Azure virtual machines as well.

The same principles apply for this virtual machine as for the SQL Server and domain controller virtual machines: do not put anything on the temporary drive, use multiple disks to increase performance, and turn off the machine when not in use to save money on your subscription.

When installing multiple management servers, you can put them in an availability group. Availability groups are used to group similar computer roles together.

You can install the prerequisites automatically using Windows PowerShell. There are many examples of how to do this on the Internet. One that installs Operations Manager unattended can be found at <http://scug.be/christopher/2014/03/10/scom-2012-r2-unattended-installation-command-line/>.

Another option under development is PowerShell Desired State Configuration (DSC). PowerShell DSC is a declarative management system inside Windows PowerShell 4.0 that enables servers to self-provision themselves during initial deployment and also self-remediate their configuration if it should fall out of compliance with their assigned "desired state." More information about PowerShell DSC can be found on Keith Mayer's blog at <http://blogs.technet.com/b/keithmayer/archive/2014/10/24/end-to-end-iaas-workload-provisioning-in-the-cloud-with-azure-automation-and-powershell-dsc-part-1.aspx> and on the Windows PowerShell blog at <http://blogs.msdn.com/b/powershell/archive/2014/08/07/introducing-the-azure-powershell-dsc-desired-state-configuration-extension.aspx>. Using DSC and the SCOM module that is located at <https://gallery.technet.microsoft.com/xSCOM-PowerShell-Desired-052fc73c>, you

can install a complete Operations Manager environment. All the other modules, for instance the module that can install a SQL Server server, can be found at <https://gallery.technet.microsoft.com/DSC-Resource-Kit-All-c449312d>.


## Monitoring Azure IaaS

This section provides an example of how you can use Operations Manager to monitor your Azure subscription using the Azure management pack for Operations Manager. Operations Manager can also monitor the virtual machines you deploy in Microsoft Azure Infrastructure as a Service (IaaS). They can be monitored from within Azure or from your on-premises environment.

## Connecting your on-premises Operations Manager environment to your Azure subscription


The Azure management pack provides a view and monitoring of your Azure subscription. This includes the virtual machines that you create, but not the details of what is running on these virtual machines. This management pack can be downloaded from <http://www.microsoft.com/en-us/download/details.aspx?id=38414>.

To connect your on-premises Operations Manager environment to your Azure subscription, you need to generate a certificate for authentication. The easiest way to create a self-signed certificate is to use an Operations Manager management server with IIS (the web console) installed. Within IIS, you can generate a self-signed certificate as shown in Figure 5-7.



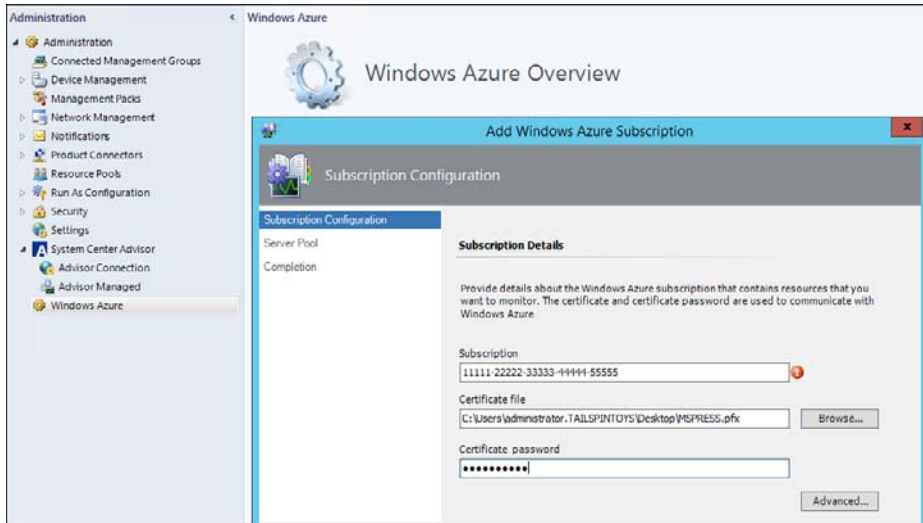
**FIGURE 5-7** Using IIS Manager to generate a self-signed certificate

After the self-signed certificate is generated, you need to open the Microsoft Management Console and add the Certificates snap-in for the computer account. From the personal certificate store, you can then export the certificate, as shown in Figure 5-8, first as a .PFX file (with the public and private keys) and then as a .CER file (with the public key).



**FIGURE 5-8** Exporting the self-signed certificate

The PFX file is needed by the Azure management pack to be able to authenticate itself to your Azure subscription. In the Administration pane, when you select Windows Azure (this will be renamed to Microsoft Azure in the future), you can add your subscription by specifying the subscription ID, the PFX file, and the password for this certificate file, as shown in Figure 5-9. The subscription ID can be found in the Azure management portal. Click Settings at the bottom of the list of available services for your subscription. In the Settings menu, under Subscriptions, you can find the subscription ID.



**FIGURE 5-9** Import the PFX certificate into Operations Manager

The CER file is needed in the Azure subscription. To add your CER certificate to your subscription, click Settings at the bottom of the list of available services. From the Settings menu, select Management Certificates. Select the CER file you just exported and click Upload.

Further configuration steps are explained in the management pack guide that you can download from the URL mentioned earlier. After you install the Azure management pack, in the Authoring pane of the management console, you will find a new management pack template called Windows Azure Monitoring (this also will be renamed in the future). Using this

template, you can configure monitoring of cloud services in production and staging slots, virtual machines, and storage.

## Using Operations Manager to monitor Azure virtual machines

Apart from using the Azure management pack for Operations Manager, you can also use Operations Manager to directly monitor the virtual machines you have deployed in Azure by installing a Microsoft Monitoring Agent on the virtual machines and connecting this agent to your Operations Manager environment. It does not matter to Operations Manager if that environment is on-premises or in the Cloud.

There are two important considerations for monitoring virtual machines in Azure. The first one is that cross-premises connectivity must be configured between your corporate network and the Azure network in order for the Operations Manager management server to communicate with the agents that are deployed on the Azure virtual machines. The appropriate ports in the firewall must be opened to make this communication possible. This can be established by enabling a Site-to-Site (S2S) or Point-to-Site (P2S) VPN connection between your corporate network and the Azure network. You can define your S2S or P2S VPN in the virtual network settings, as explained earlier in this chapter.

While the S2S connectivity requires you to have a VPN device, the P2S connectivity allows you to set up VPN connections between individual computers and an Azure virtual network without the need for a VPN device.

The second consideration is that there must be mutual authentication between the Operations Manager management group and the Azure IaaS virtual machines. If the virtual machines are part of a trusted environment, then Kerberos can be used for authentication. If this is not the case, certificates should be used for mutual authentication. The certificates deployed to the Azure IaaS virtual machines enable mutual authentication between the Operations Manager management group and the Azure IaaS virtual machines. If you have more than a handful of virtual machines to monitor in Azure, then deploy an Operations Manager Gateway Server in the Azure network and configure certificates on this server for mutual authentication with your on-premises Operations Manager management group. The Operations Manager Gateway server then uses Kerberos to manage all the downstream IaaS virtual machines in Azure.

**See also** *An extensive explanation of the Azure Monitoring Gateway was written by Cameron Fuller, MVP. This series can be found on his blog at <http://blogs.catapultsystems.com/cfuller/archive/2013/12/04/operations-manager-and-azure-better-together-introducing-the-azure-monitoring-gateway-%5b%5d.aspx>.*

# About the authors



**DANNY HERMANS** is a Senior Partner Technical Advisor for System Center for Broad Commercial Services EMEA based in Sweden, with almost 20 years of experience in various consulting projects, ranging from green field design over optimization to migration assignments. Since joining Microsoft three years ago, he has been a Senior Premier Field Engineer (PFE), Tech Lead, and Risk Assessment Program (RAP) as a Service Developer for Operations Manager. Throughout his career, Danny has developed a love of fine-tuning Microsoft products and seeks to enhance performance wherever possible to add value to his customer outcomes. Danny's blog can be found at <http://blogs.technet.com/b/dhermans>.



**UWE STÜRTZ** has worked for Microsoft since 2007 as a Premier Field Engineer (PFE) for System Center Technology in Global Business Support Germany. Mainly focused on System Center Operations Manager and other monitoring solutions, he has a background of more than 18 years in system management. He leads the Microsoft German PFE-SCOM community and is a member of the EMEA SCOM Tech Lead team. Uwe is responsible for leading the SCOM PFE RAP as a Service process and teaches the Operations Manager engineers in EMEA to deliver this proactive service to Microsoft premier customers with the right technology and quality skills.



**MIHAI SARBULESCU** is a Support Escalation Engineer on the Cloud & Datacenter Management team for the EMEA region. He is based in Romania and has been working for Microsoft for the past five years. Before working at Microsoft, he was a web developer. Mihai enjoys debugging, programming, computer games, and martial arts. His blog can be found at <http://blogs.technet.com/b/mihai/>.

# About the series editor



**MITCH TULLOCH** is a well-known expert on Windows Server administration and cloud computing technologies. He has published hundreds of articles on a wide variety of technology sites and has written, contributed to or been series editor for over 50 books. Mitch is one of the most popular authors at Microsoft Press—the almost two dozen ebooks on Windows Server and System Center he either wrote or was Series Editor on have been downloaded more than 2.5 million times! For a complete list of

free ebooks from Microsoft Press, visit the Microsoft Virtual Academy at <http://www.microsoftvirtualacademy.com/ebooks>.

Mitch has repeatedly received Microsoft's Most Valuable Professional (MVP) award for his outstanding contributions to supporting the global IT community. He is a ten-time MVP in the technology area of Windows Server Software Packaging, Deployment & Servicing. You can find his MVP Profile page at <http://mvp.microsoft.com/en-us/mvp/Mitch%20Tulloch-21182>.

Mitch is also Senior Editor of WServerNews, a weekly newsletter focused on system admin and security issues for the Windows Server platform. With almost 100,000 IT pro subscribers worldwide, WServerNews is the most popular Windows Server–focused newsletter in the world. Visit <http://www.wservernews.com> and subscribe to WServerNews today!

Mitch also runs an IT content development business based in Winnipeg, Canada, that produces white papers and other collateral for the business decision maker (BDM) and technical decision maker (TDM) audiences. His published content ranges from white papers about Microsoft cloud technologies to reviews of third-party products designed for the Windows Server platform. Before starting his own business in 1998, Mitch worked as a Microsoft Certified Trainer (MCT) for Productivity Point.

For more information about Mitch, visit his website at <http://www.mtit.com>. You can also follow Mitch on Twitter @mitchtulloch.



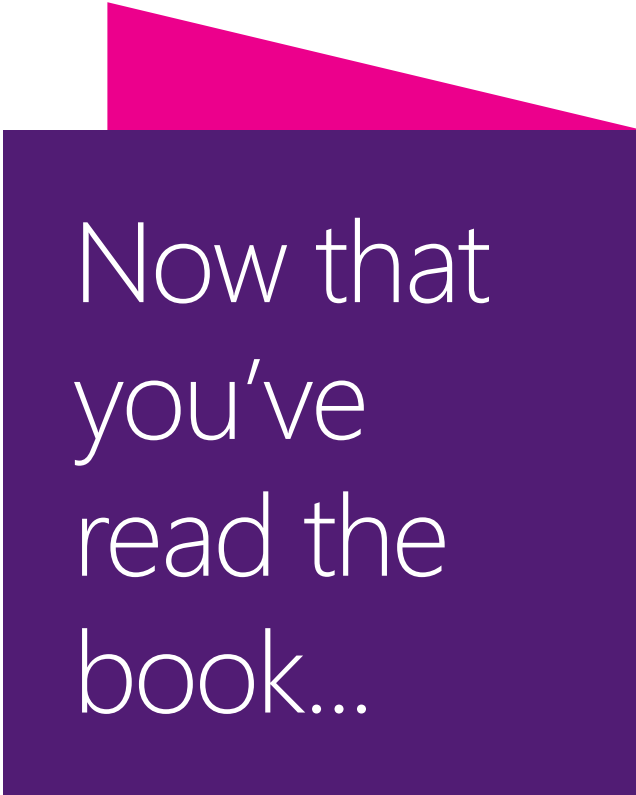
From technical overviews to drilldowns on special topics, get *free* ebooks from Microsoft Press at:

**[www.microsoftvirtualacademy.com/ebooks](http://www.microsoftvirtualacademy.com/ebooks)**

Download your free ebooks in PDF, EPUB, and/or Mobi for Kindle formats.

Look for other great resources at Microsoft Virtual Academy, where you can learn new skills and help advance your career with free Microsoft training delivered by experts.

**Microsoft Press**



Now that  
you've  
read the  
book...

Tell us what you think!

Was it useful?

Did it teach you what you wanted to learn?

Was there room for improvement?

**Let us know at <http://aka.ms/tellpress>**

Your feedback goes directly to the staff at Microsoft Press,  
and we read every one of your responses. Thanks in advance!

