



**Hewlett Packard**  
Enterprise

# **HPE Operations Manager i**

Software Version: 10.10

## **Database Guide**

Document Release Date: 18 December 2015  
Software Release Date: December 2015

## Legal Notices

### Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

### Restricted Rights Legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

### Copyright Notice

© Copyright 2015 Hewlett Packard Enterprise Development LP

### Trademark Notices

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

AMD and the AMD Arrow symbol are trademarks of Advanced Micro Devices, Inc.

Citrix® and XenDesktop® are registered trademarks of Citrix Systems, Inc. and/or one more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries.

Google™ and Google Maps™ are trademarks of Google Inc.

Intel®, Itanium®, Pentium®, and Intel® Xeon® are trademarks of Intel Corporation in the U.S. and other countries.

iPad® and iPhone® are trademarks of Apple Inc.

Java is a registered trademark of Oracle and/or its affiliates.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Microsoft®, Windows®, Lync®, Windows NT®, Windows® XP, and Windows Vista® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

NVIDIA® is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

Red Hat® is a registered trademark of Red Hat, Inc. in the United States and other countries.

SAP® is the trademark or registered trademark of SAP SE in Germany and in several other countries.

UNIX® is a registered trademark of The Open Group.

## Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

[https://softwaresupport.hp.com/group/softwaresupport/search-result?keyword=.](https://softwaresupport.hp.com/group/softwaresupport/search-result?keyword=)

This site requires an HP Passport account. If you do not have one, click the **Create an account** button on the HP Passport Sign in page.

## Support

Visit the HPE Software Support website at: <https://softwaresupport.hp.com>

This website provides contact information and details about the products, services, and support that HPE Software offers.

HPE Software Support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support website to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HPE support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to <https://softwaresupport.hp.com> and click **Register**.

To find more information about access levels, go to:

<https://softwaresupport.hp.com/web/softwaresupport/access-levels>

### HPE Software Solutions & Integrations and Best Practices

Visit HPE Software Solutions Now at <https://softwaresupport.hp.com/group/softwaresupport/search-result-/facetsearch/document/KM01702710> to explore how the products in the HPE Software catalog work together, exchange information, and solve business needs.

Visit the Cross Portfolio Best Practices Library at <https://hpln.hp.com/group/best-practices-hpsw> to access a wide variety of best practice documents and materials.

# Contents

Part I: Introducing the Database Environment .....	9
Chapter 1: Introduction to Preparing the Database Environment .....	10
Database Overview .....	10
Requirements .....	11
Database Character Set Configuration .....	11
Part II: Deploying and Maintaining the Microsoft SQL Server Databases .....	12
Chapter 2: Microsoft SQL Server Deployment Overview .....	13
About Microsoft SQL Server Deployment .....	13
System Requirements .....	13
Hardware Requirements .....	13
Software Requirements .....	13
Character Set Support in Microsoft SQL Server .....	14
Chapter 3: Manually Creating the Microsoft SQL Server Database .....	15
Microsoft SQL Server Databases Overview .....	15
Manually Creating the OMi Databases .....	15
Creating Databases .....	15
Creating the Objects .....	16
Creating the Management Objects .....	16
Creating the RTSM Objects .....	17
Creating the Event Objects .....	17
Chapter 4: Installing and Configuring Microsoft SQL Server .....	18
Workflow for Microsoft SQL Server Deployment .....	18
Installation Prerequisites .....	19
Notes and Limitations .....	19
Checklist for Support and Certification .....	20
Installing Microsoft SQL Server .....	20
Configuring Microsoft SQL Server .....	22
Service Configuration Options .....	22
Server Configuration Options .....	22
Verifying and Modifying Server and Database Settings .....	22
Port Allocation Options for Named Instances .....	24
Chapter 5: Creating and Configuring Microsoft SQL Server Databases .....	26
Creating Databases .....	26
Creating Microsoft SQL Server Named Instances .....	26
Database Permissions .....	28
Database File Layout .....	28
Data and Log Placement .....	28
File and Database Properties .....	29
tempdb Database Settings .....	29

File Groups .....	30
System Databases .....	30
Configuring Databases .....	31
Database File Configuration .....	31
Adding Files .....	31
Dropping Files .....	31
Changing File Properties .....	32
Database Configuration Options .....	32
Chapter 6: Maintaining Microsoft SQL Server Databases .....	35
Backing Up Databases .....	35
Full Backup .....	36
Differential Backup .....	36
Log Backup .....	36
File and File Group Backup .....	37
Maintenance Plan .....	37
Transaction Log Issues .....	37
Database Integrity and Fragmentation .....	37
Database Integrity .....	38
Understanding File System Fragmentation .....	38
Understanding Internal Fragmentation .....	38
Understanding External Fragmentation .....	39
Detecting and Handling Index Fragmentation .....	39
Using sys.dm_db_index_physical_stats to Detect Fragmentation .....	39
Handling Fragmentation .....	40
Supplied Utilities for Monitoring and Rebuilding Indexes .....	40
Utility to Rebuild All Indexes in Database .....	40
Utility to Rebuild Indexes Based on the Fragmentation Level of Each Index .....	41
Distribution Statistics .....	41
Collecting Statistics for RTSM .....	42
Utility to Refresh Statistics .....	42
Chapter 7: Using Windows Authentication to Access Microsoft SQL Server Databases .....	43
Enabling OMi to Work with Windows Authentication .....	43
Configuring Microsoft SQL Server to Use Windows Authentication .....	43
Launching the OMi service with a Windows User .....	43
Creating or Connecting to a Microsoft SQL Database Using Windows Authentication .....	44
 Part III: Deploying and Maintaining the Oracle Server Database .....	 45
Chapter 8: Oracle Server Deployment Overview .....	46
About Oracle Server Deployment .....	46
System Requirements .....	46
Hardware Requirements .....	47
Software Requirements .....	47
Oracle Instances .....	47
Character Set Support in Oracle .....	47
Chapter 9: Manually Creating the Oracle Schemas .....	49
Oracle Server Databases Overview .....	49

Creating OMi Schemas Manually .....	49
Creating a Management User Schema Manually .....	50
Creating an RTSM User Schema Manually .....	51
Creating an Event Schema Manually .....	51
OMi Oracle Schema Privileges .....	52
Chapter 10: Oracle Server Configuration and Sizing Guidelines .....	53
Oracle Parameter Settings .....	53
Oracle Tablespaces .....	56
Locally Managed Tablespaces .....	56
Creating an Oracle Tablespace .....	57
Oracle Tablespace Settings .....	57
Data Tablespace Settings .....	58
Data Tablespace Default Storage Settings .....	58
System Tablespace Settings .....	58
Temporary Tablespace Settings .....	59
Redo Log Settings .....	59
Undo Segment Settings .....	59
Using RAID Configuration .....	60
ASM Storage .....	61
Chapter 11: Maintaining an Oracle Server Database .....	62
Database Maintenance and Tuning .....	62
System Global Area (SGA) .....	62
Database Load Behavior .....	63
CPU and I/O .....	63
Oracle Alert File .....	63
Archive Log - File System .....	63
Tablespace Storage Space .....	64
Space Management .....	64
Collecting Statistics for Databases .....	64
Oracle Automated Statistics Collection .....	64
Collection Statistics for RTSM .....	65
Index Fragmentation .....	65
Index Maintenance Utility .....	66
Oracle Database Backup and Recovery .....	67
Available Backup Methods .....	67
Cold Backup .....	67
Hot Backup .....	68
Data Pump Export Import .....	68
Oracle Recovery Manager - RMAN .....	68
Chapter 12: Configuring the Oracle Client for OMi .....	70
Chapter 13: Oracle Summary Checklist .....	72
Part IV: Deploying and Maintaining the PostgreSQL Server Databases .....	76
Chapter 14: PostgreSQL Deployment Overview .....	77
About PostgreSQL Deployment .....	77

System Requirements .....	78
Hardware Requirements .....	78
Software Requirements .....	78
Character Set Support in PostgreSQL .....	79
Chapter 15: Creating the OMi Databases Manually .....	80
PostgreSQL Database Overview .....	80
Creating OMi Databases Manually .....	80
Creating a Management Database Manually .....	80
Creating an RTSM Database Manually .....	81
Creating an Event Database Manually .....	81
Chapter 16: Maintaining a PostgreSQL Database .....	83
Database Maintenance .....	83
PostgreSQL Database Backup and Restore .....	83
Backing up and Restoring an Embedded PostgreSQL Database .....	84
Embedded PostgreSQL Database Backup .....	84
Embedded PostgreSQL Database Restore .....	84
Backing up and Restoring an External PostgreSQL Database .....	85
SQL Dump .....	85
File System Level .....	85
Continuous Archiving .....	85
Resetting or Changing the User Password for the Embedded PostgreSQL Database .....	86
Chapter 17: Migrating from an Embedded PostgreSQL Database to an External PostgreSQL Database .....	88
Reasons for Migrating Your Embedded PostgreSQL Database .....	88
Procedure for Migrating from an Embedded PostgreSQL Database to an External PostgreSQL Database .....	88
<b>Part V: Appendixes .....</b>	<b>90</b>
Appendix A: Microsoft SQL Server Data Storage Recommendation .....	91
Creating Microsoft SQL Server File Groups .....	91
Benefits of Multiple File Groups .....	91
Appendix B: Microsoft SQL Server Configuration Options .....	94
Appendix C: Changing the Database Host .....	99
Appendix D: Support for Oracle Real Application Cluster .....	100
About Oracle Real Application Cluster (RAC) .....	100
Single Client Access Name .....	101
Client Side Configuration for Oracle RAC .....	101
Server Side Configuration .....	104
Recommendation for OMi Database .....	104
Create/Connect to One of the OMi Databases .....	105
<b>Send Documentation Feedback .....</b>	<b>106</b>





# Part I: Introducing the Database Environment

# Chapter 1: Introduction to Preparing the Database Environment

This chapter contains information on the types of databases used with HPE Operations Manager i (OMi).

## Database Overview

To work with OMi, you must set up the following databases:

- **Management.** For storage of system-wide and management-related metadata for the OMi environment. OMi requires one management database. You can create this database manually, or by using the configuration wizard.
- **Event.** For storage of events and related data, such as annotations, as well as for storage of configuration data, such as event correlation rules.
- **Run-time Service Model (RTSM).** For storage of configuration information that is gathered from the various OMi and third-party applications and tools. This information is used when building OMi views.

You can set up OMi databases on a Microsoft SQL Server, an Oracle Server, or a PostgreSQL Server, depending on the type of database server used in your organization:

- If you are working with a Microsoft SQL Server database, see ["Deploying and Maintaining the Microsoft SQL Server Databases" on page 12](#)
- If you are working with an Oracle Server database, see ["Deploying and Maintaining the Oracle Server Database" on page 45](#)
- If you are working with a PostgreSQL Server database, see ["Deploying and Maintaining the PostgreSQL Server Databases" on page 76](#).

The appendixes contain additional information that is pertinent to the Microsoft SQL Server, Oracle Server, and PostgreSQL Server databases.

**Note:** Database servers must be set to the same time zone, daylight savings settings and time as the OMi servers.

## Requirements

The following table describes the hardware (CPU and memory) requirements recommended for the OMi database server:

Deployment Size	Number of Processors	Physical Memory
Small (up to 2,000 nodes)	2 CPU cores	Minimum: 2 GB RAM Recommended: 4 GB RAM
Medium (up to 5,000 nodes)	2 CPU cores	Minimum: 2 GB RAM Recommended: 6 GB RAM
Large (more than 5,000 nodes)	Minimum 4 CPU cores	Minimum: 4 GB RAM Recommended: 8 GB RAM and up

For software requirements, see the OMi Support Matrix at:

<https://softwaresupport.hp.com/group/softwaresupport/search-result/-/facetsearch/document/KM323488>

## Database Character Set Configuration

When creating a database and specifying a character set, which determines what languages can be stored and represented in the database and OMi, make sure that you choose the one that will support the characters you need in your OMi environment. You must take special care when using languages with non-Latin characters or working in a multilingual environment that requires storing and processing data in several languages simultaneously.

For successful communication, both OMi and the database must use a compatible character set. Otherwise, data is not stored correctly and errors occur. In addition, using the same character set increases the performance because the need for character conversion is reduced or eliminated. For detailed information about considerations and current limitations when using non-Latin characters, see the OMi Administration Guide.

**Note:** You can use the Oracle and PostgreSQL server databases for storing multilingual data in OMi. However, with the current implementation, it is not possible for OMi to store multilingual data in the Microsoft SQL Server database—only one character set can be used at a time.

Depending on the type of database server, see one of the following sections:

- ["Character Set Support in Microsoft SQL Server" on page 14](#)
- ["Character Set Support in Oracle" on page 47](#)
- ["Character Set Support in PostgreSQL" on page 79](#)

# Part II: Deploying and Maintaining the Microsoft SQL Server Databases

# Chapter 2: Microsoft SQL Server Deployment Overview

You can use Microsoft SQL Server for deploying OMi databases. This chapter describes the following topics related to deploying Microsoft SQL Servers for use with OMi:

- ["About Microsoft SQL Server Deployment" below](#)
- ["System Requirements" below](#)
- ["Character Set Support in Microsoft SQL Server" on the next page](#)

## About Microsoft SQL Server Deployment

To deploy Microsoft SQL Server for use with OMi, perform the following procedures:

- **Install and configure Microsoft SQL Server.**  
For details on installing and configuring Microsoft SQL Server, see ["Installing and Configuring Microsoft SQL Server" on page 18](#).
- **Create databases on Microsoft SQL Server.**  
You can create OMi databases manually or you can use the configuration wizard to create the databases.  
For details on manually creating databases for OMi, see ["Manually Creating the OMi Databases" on page 15](#).  
For details on creating Microsoft SQL database, see ["Creating and Configuring Microsoft SQL Server Databases" on page 26](#).

## System Requirements

This section describes the system requirements for working with Microsoft SQL Server in conjunction with OMi.

### Hardware Requirements

For OMi hardware sizing guidelines, see ["Requirements" on page 11](#). For Microsoft SQL Server hardware requirements, see the installation guide for your Microsoft SQL Server release for your operating system.

### Software Requirements

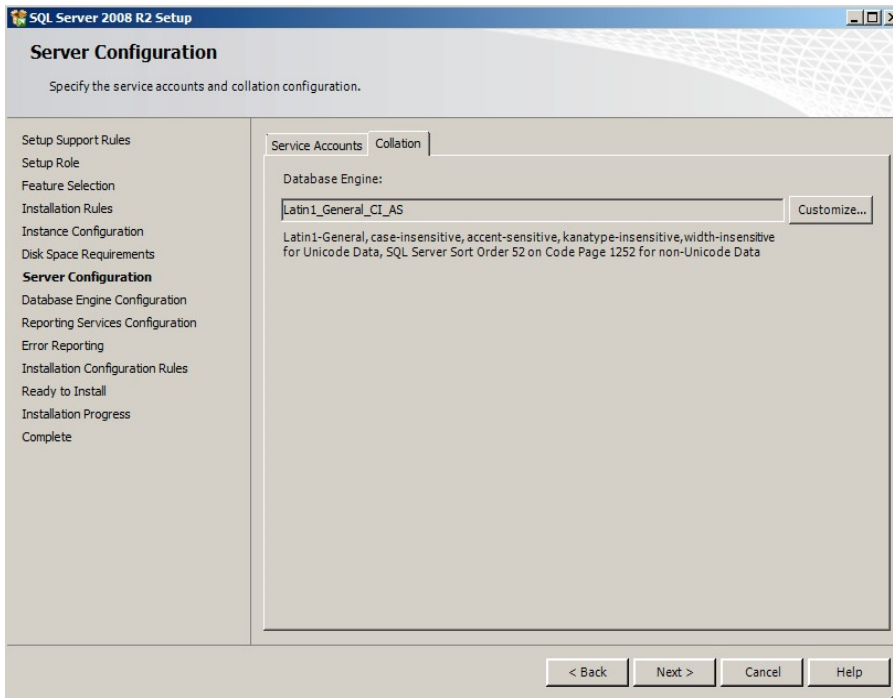
For a list of supported Microsoft SQL Server database versions, see the OMi Support Matrix at:

<https://softwaresupport.hp.com/group/softwaresupport/search-result/-/facetsearch/document/KM323488>

## Character Set Support in Microsoft SQL Server

Because Microsoft SQL Server does not support UTF-8 encoding (although data can be stored internally in the NVARCHAR fields as UCS or UTF-16), OMi cannot store multilingual data if a Microsoft SQL Server database is used. It is still, however, possible to use one or more languages with non-Latin characters, but they must share the same character set. For this, you must specify a collation that is in accordance with the OMi language and machine's regional settings.

During the database installation, the setup tool detects the current regional settings and suggests the most suitable collation.



The following table lists languages and their default collations:

Language	Default Collation
English	Latin1_General_CI_AS
French	French_CI_AS
Spanish	Modern_Spanish_CI_AS
German	Latin1_General_CI_AS
Russian	Cyrillic_General_CI_AS
Japanese	Japanese_CI_AS
Korean	Korean_Wansung_CI_AS
Simplified Chinese	Chinese_PRC_CI_AS

# Chapter 3: Manually Creating the Microsoft SQL Server Database

This chapter describes the procedure for manually creating OMi databases on a Microsoft SQL Server.

This chapter includes:

- ["Microsoft SQL Server Databases Overview" below](#)
- ["Manually Creating the OMi Databases" below](#)

## Microsoft SQL Server Databases Overview

OMi uses the Management, RTSM, and Event databases.

During the OMi setup, new databases can be set up automatically by the setup procedure, or already existing databases can be used. Existing databases can either be created manually in advance (for example, due to organization security restrictions), or can be created by a previous installation of the same release of OMi.

For details about installing OMi server, see the OMi Installation and Upgrade Guide.

## Manually Creating the OMi Databases

After you have installed and configured Microsoft SQL Server, you create OMi databases for the storage of data collected by OMi.

You can allow OMi to create OMi databases for you using the configuration wizard.

If you created the databases in advance (see below), you can connect OMi to the existing databases from the configuration wizard.

For information on the configuration wizard, see the OMi Installation and Upgrade Guide.

This section includes the following topics:

- ["Creating Databases" below](#)
- ["Creating the Objects" on the next page](#)

## Creating Databases

The database administrator should create the following Microsoft SQL Server databases, one for Management, one for RTSM, and one for Event.

To create a database, you must have **CREATE DATABASE** permissions. To connect to an existing database, the login account with which you are connecting must be mapped to **dbo** in the database.

**Note:** Members of the sysadmin server role automatically have **CREATE DATABASE**

permissions, and are also mapped to **dbo** in all databases. A database owner is automatically mapped to **dbo** in the database.

For details on creating databases manually, see "[Creating and Configuring Microsoft SQL Server Databases](#)" on page 26.

**Note:** When the OMi setup procedure automatically creates the databases, it uses two file groups for each database, one to hold the system tables and one for the application data.

## Creating the Objects

After creating each OMi database, you create objects by running SQL scripts. The object creation scripts are located in the following directory:

```
<OMi_HOME_GW>/AppServer/webapps/site.war/DataBases/SQL_Svr_DB_Utills
```

For the Management database, the SQL script is dynamically created by the `generate_create_script` utility:

```
<OMi_HOME_GW>/dbverify/bin/generate_create_script
```

**Note:** The SQL script files generated by the `generate_create_script` utility contain references to the [PRIMARY] files group. In order to use a different file group, manually open the .sql file in a text editor and replace all instances of [PRIMARY] with the desired file group.

## Creating the Management Objects

To create the OMi Management objects, follow the procedure below to generate and run the Management database SQL script:

### To generate the Management database SQL script:

Run one of the following scripts file located in `<OMi_HOME_GW>/dbverify/bin` directory:

For Windows:

```
generate_create_script.bat [deployment_type] [Generated Script Name] [dbType]  
[sessionType]
```

For Linux:

```
./generate_create_script.sh [deployment_type] [Generated Script Name] [dbType]  
[sessionType]
```

The following parameters should be used for the Management schema:

- **deployment\_type.** EE (Enterprise Edition)
- **Generate Script Name.** The name of the SQL file to generate, for example, `Management.sql`.
- **dbType.** `SqlServer`
- **sessionType.** `Management`

**Example:**



```
/opt/HP/BSM/dbverify/bin/generate_create_script.sh EE Management.sql SqlServer  
Management
```

**To run the Management database SQL script:**

1. Create the Management database. For details, see ["Creating Databases" on page 15](#).
2. Connect to the Management database and run the script generated in the previous section to deploy the Management objects.
3. Run the following SQL script from the objects scripts directory:

```
management_sql_inserts.sql
```

## Creating the RTSM Objects

Create one database for storing the RTSM objects. No scripts are necessary to create the objects to populate these databases.

## Creating the Event Objects

Create the Event database for storing Event objects. Follow the procedure below to create Event objects:

**To generate the Event database SQL script:**

Run one of the following script files:

For Windows:

```
<OMi_HOME>\opr\bin\opr-schema-script-generator.bat
```

For Linux:

```
<OMi_HOME>/opr/bin/opr-schema-script-generator.sh
```

**Note:** The database type parameter, `-d`, should be set to **SQL**.

For example, on Windows, to generate a script called `event.sql` that can be used to create the necessary tables on the Microsoft SQL Server database, run the following command:

```
<OMi_HOME>\opr\bin\opr-schema-script-generator.bat -d SQL -o C:\tmp\event.sql
```

**To run the Event database SQL script:**

1. Create the Event database. For details, see ["Creating Databases" on page 15](#).
2. Connect to the Event database and run the script generated in the previous section to deploy the Operations objects.

**Note:** Perform the above procedures only if you are an experienced Microsoft SQL Server database administrator.

# Chapter 4: Installing and Configuring Microsoft SQL Server

This chapter contains information on the Microsoft SQL Server installation procedure and configuration settings.

This chapter includes:

- ["Workflow for Microsoft SQL Server Deployment"](#) below
- ["Installation Prerequisites"](#) on the next page
- ["Notes and Limitations"](#) on the next page
- ["Checklist for Support and Certification"](#) on page 20
- ["Installing Microsoft SQL Server"](#) on page 20
- ["Configuring Microsoft SQL Server"](#) on page 22
- ["Verifying and Modifying Server and Database Settings"](#) on page 22

## Workflow for Microsoft SQL Server Deployment

To deploy Microsoft SQL Server for use with OMi, perform the following steps:

1. **Review Microsoft SQL Server sizing guidelines**  
For details, see ["Database File Layout"](#) on page 28.
2. **Review installation prerequisites**  
For details, see ["Installation Prerequisites"](#) on the next page.
3. **Review server and database checklists**  
These checklists summarize the server and database configuration options supported and recommended for OMi.  
For details, see ["Checklist for Support and Certification"](#) on page 20.
4. **Install Microsoft SQL Server Enterprise edition**  
For details, see ["Installing Microsoft SQL Server"](#) on page 20.
5. **Configure Microsoft SQL Server**  
For details, see ["Configuring Microsoft SQL Server"](#) on page 22.
6. **Create OMi databases on Microsoft SQL Server**  
For details, see ["Creating Databases"](#) on page 15.
7. **Configure OMi databases**  
For details, see ["Configuring Databases"](#) on page 31.
8. **Verify Microsoft SQL Server and databases**  
For details, see ["Verifying and Modifying Server and Database Settings"](#) on page 22.
9. **Set up Windows authentication**

For details, see ["Using Windows Authentication to Access Microsoft SQL Server Databases"](#) on [page 43](#). This step is only relevant if you have selected Windows authentication instead of SQL Server authentication.

## Installation Prerequisites

- **System Requirements**

For a list of supported operating systems, see ["System Requirements"](#) on [page 13](#).

- **Software**

Software	Version Needed
Microsoft Windows Installer	4.5 or later

- **Database file placement**

For improved performance, it is recommended to spread out OMi databases among several disks or RAID systems.

- **Memory**

There must be at least 2 GB of RAM.

- **User Accounts**

- If you plan to perform activities outside the local machine (such as file copying to or from different servers, backing up to a shared remote repository, and so forth), you must have a domain account available for Microsoft SQL Server services.
- If you want to install OMi using the installation wizard, you must provide a user account that has database creator privileges. If you are going to create databases manually, supply OMi with a user account that is a part of **db\_datareader**, **db\_datawriter** and **db\_ddladmin** in each OMi database.

- **Miscellaneous**

- Verify that the disk where Microsoft SQL Server is to be installed is not compressed.
- Disable antivirus software and any other applications that may interfere while installing Microsoft SQL Server.

## Notes and Limitations

- Microsoft SQL Server must be installed on a dedicated machine to host the OMi databases.
- Do not change the default memory settings. Allow the Microsoft SQL Server to manage memory dynamically, except when you configure **awe enabled** support. The AWE feature was removed from SQL Server 2012, and is not recommended for use starting SQL Server 2008 R2.
- Microsoft SQL Server 64-bit versions include support for extended systems, also known as Windows on Windows (WOW64). WOW64 is a feature of 64-bit editions of Microsoft Windows that allows 32-bit applications to execute natively in 32-bit mode.
- Applications function in 32-bit mode even though the underlying operating system is running on the

64-bit platform.

- When working with Microsoft SQL Server, the RTSM collation and the SQL Server collation must be the same in the old server and new server for an upgrade to run properly. For more information, see ["Database Character Set Configuration" on page 11](#).

## Checklist for Support and Certification

The information provided in this section is for both supported and certified Microsoft SQL Server options. The certified options are recommended for working with OMi.

Subject	Microsoft SQL Server	
	Supported	Recommended
Instances	Default, Single	
Authentication Mode	Mixed, Windows	
Collation	Case-Insensitive. OMi does not support binary sort order and case sensitivity. Only case-insensitive order with a combination of accent, kana, or width settings is supported.	Use the Collation Settings dialog box to select the collation. Do not select the binary check box. Accent, kana, and width sensitivity should be selected according to the relevant data language requirements. The selected language must be the same as the OS Windows regional settings language.
Network Libraries	<b>Server:</b> TCP/IP and Named Pipes <b>Client:</b> TCP/IP and Named Pipes	<b>Server:</b> TCP/IP <b>Client:</b> TCP/IP
Server Configuration Options	Defaults, unless instructed otherwise	
Data File Properties	Manual file growth, or FILEGROWTH less than or equal to 100 MB	FILEGROWTH: ~30-100 MB
Collation Database Property	Server default	
Database Options	Defaults, unless instructed otherwise	
Recovery Model	Any	Full

## Installing Microsoft SQL Server

It is important that you familiarize yourself with all of the installation details so that you select the appropriate options. Selecting the default options may, in some cases, negatively affect the Microsoft

SQL Server's performance.

Select the following options in the installation dialog boxes:

- **Feature Selection dialog box configuration.**
  - Remove Full Text Search from the list since OMi does not use this indexing search feature.
  - Under Destination Folder, make sure that the Data Files directory is stored on a fault-tolerant disk system, for example, RAID 1. Even though these system databases are fairly small, they are essential for the operation of Microsoft SQL Server.
- **Instance Name dialog box configuration.** For details on named instances, see "[Port Allocation Options for Named Instances](#)" on page 24. You access a default instance by specifying the server name or IP address.
- **Service Account dialog box configuration.**
  - If all Microsoft SQL Server activities are outside the local machine (for example, file copying to or from a different server, backing up to shared remote repository, replication with other servers, ActiveX script job steps, CmdExec job steps, and so forth), choose **Use a Domain User account** and specify the user name, password, and domain of a user that is a member of the local machine's administrator group, and that has the appropriate permissions for network resources.
  - If all Microsoft SQL Server activities are limited to the local machine, choose **Use the built-in System account** and select **Local system**. This selection provides SQL Server administrative privileges on the local machine only.
- **Authentication Mode dialog box.** OMi works with both Windows authentication and with Microsoft SQL Server authentication (recommended). To enable Microsoft SQL Server authentication, do the following:
  - a. Choose **Mixed Mode (Windows Authentication and SQL Server Authentication)**.
  - b. Enter the password for user **sa**.

**Note:** To further secure your Microsoft SQL Server, it is important to enter a password.

- **Collation Settings dialog box.** Follow the recommendations found in the checklist in "[Checklist for Support and Certification](#)" on the previous page.

The above settings affect only the system databases and serve as the default settings for user databases. Databases can have different collation settings from the server's default settings, and a table column can have different collation settings from the database's default settings, but this is not recommended for dedicated OMi databases. Because of the flexibility in collation management in Microsoft SQL Server, you can restore or attach a database that has different collation settings.

**Important:** Changing any of the above settings requires scripting all system objects and routines (logins, user defined system messages, master stored procedures, and so forth), reinstalling Microsoft SQL Server (or running the RebuildM.exe utility) with the new settings, recreating all system objects from the saved scripts, and attaching the user databases. It is therefore recommended that you select the appropriate options during the installation process.

- **Data Directories dialog box.** Specify the destination folders for the data and log files of the user databases and temporary databases, as well as the backup directory. It is recommended to store

data and log files on high-performance storage systems.

Install the latest service pack for Microsoft SQL Server available at the Microsoft Download Center site at:

<http://www.microsoft.com/downloads/details.aspx?FamilyID=cb6c71ea-d649-47ff-9176-e7cac58fd4bc&DisplayLang=en>

For the latest service pack supported to work with OMi, see "[Software Requirements](#)" on page 13.

## Configuring Microsoft SQL Server

This section describes the service and server options you can configure once you have installed Microsoft SQL Server and includes the following topics:

- "[Service Configuration Options](#)" below
- "[Server Configuration Options](#)" below

### Service Configuration Options

If you installed **Full-Text Search**, ensure that it is disabled and set to manual mode (locate the service in the Services applet using Microsoft Search) so that no resources are wasted.

Unless you are using distributed transactions, ensure that the **Distributed Transactions Coordinator** service is also disabled or set to manual mode.

If the dynamic port option is used for Microsoft SQL Server instances, make sure the SQL Server browser service, as well as all the SQL Server instance services, are in automatic mode.

### Server Configuration Options

Most server configuration options are dynamically configured by Microsoft SQL Server. For OMi certification, you should not change the default options unless you are instructed to do so by HPE Software Support.

There are specific situations in which you may want to change the default settings. You can change these settings in the **sp\_configure** stored procedure, or in the various dialog boxes in Microsoft SQL Management Studio (mainly the Server Properties dialog box).

The default values should be used except for the following items:

- Agent XPs (A) = 1
- awe enabled (A) = 1. This should only be set when more than 4GB is required on a 32Bit machine. Note that this feature is deprecated in Microsoft SQL Server 2012.

For a full list of Microsoft SQL Server configuration options, see "[Microsoft SQL Server Configuration Options](#)" on page 94.

## Verifying and Modifying Server and Database Settings

The following table summarizes the procedures for verifying or modifying server and database settings:

Server/Database Setting	How to Verify/Modify the Setting
Default Instance	In the operating system's Services applet, a default Microsoft SQL Server instance and a named instance appear as SQL Server (Instance_Name).
Authentication Mode	In the Microsoft SQL Server Management Studio, right-click the server, choose <b>Properties</b> , and click the <b>Security</b> tab. Select <b>SQL Server and Windows Authentication mode (Mixed Mode)</b> .
Collation Settings	Run the following command:  <code>sp_helpsort</code>
Network Libraries	<p>On the Server, select <b>Start &gt; Programs &gt; Microsoft SQL Server &lt;release number&gt; &gt; Configuration Tools &gt; Configuration Manager</b>.</p> <p>At <b>SQL Native Client..</b>, choose <b>client protocols</b> and ensure the selected protocol is in the enabled state.</p> <ul style="list-style-type: none"> <li>• Supported Shared memory, TCP/IP ,and Named Pipes for both the server and client</li> <li>• Recommended Only TCP/IP for both the server and client</li> </ul>
View or change server configuration options	<ul style="list-style-type: none"> <li>• To allow the viewing of all options, run EXEC sp_configure 'show advanced options', 1 reconfigure with override.</li> <li>• To view the current values, run EXEC sp_configure.</li> <li>• To change a setting, run EXEC sp_configure '&lt;option&gt;', &lt;value&gt;.</li> </ul> <p>Some options take effect only after you run <b>reconfigure</b> with <b>override</b>, while others require restarting the Microsoft SQL Server service. For more details, see Microsoft SQL Server documentation.</p>
Check whether the user who is going to create OMi databases has CREATE DATABASE permissions	<p>Log in to the Microsoft SQL Server Management Studio with the login you want to check, and run the following:</p> <pre>USE master IF PERMISSIONS() &amp; 1 = 1 PRINT 'User has CREATE DATABASE permissions' ELSE PRINT 'User does not have CREATE DATABASE permissions'</pre>

Server/Database Setting	How to Verify/Modify the Setting
Check whether the OMi database user has enough permissions in the database	<ul style="list-style-type: none"> <li>Log in to the Microsoft SQL Server Management Studio with the user name you want to check.</li> <li>Change the context of the database to the required database.</li> <li>Open a new query and do the following in each database:                             <pre>select case when IS_MEMBER ('db_owner')=1                                 or IS_SRVROLEMEMBER ('sysadmin')=1                                 or (IS_MEMBER ('db_ddladmin') = 1 and                                     IS_MEMBER ('db_datareader')=1 and                                     IS_MEMBER ('db_datawriter')=1 and                                     IS_MEMBER ('db_denydatareader')=0 and                                     IS_MEMBER ('db_denydatawriter')=0 )                                 then 'User has enough permissions'                                 else 'User does not have enough permissions'                             end</pre> </li> </ul>
Data and log file destination directory is not compressed (only in NTFS)	Right click the directory, choose <b>Properties</b> , and then <b>Advanced</b> . Verify that the <b>Compression</b> check box is cleared.
Database and database file properties (including recovery model and collation properties)	<ul style="list-style-type: none"> <li>To view the database and database file properties, run:                             <pre>EXEC sp_helpdb &lt;database name&gt;</pre> </li> <li>To change the database properties, run:                             <pre>ALTER DATABASE &lt;database name&gt; SET &lt;option&gt; &lt;value&gt;</pre> </li> <li>To change the database file properties, run:                             <pre>ALTER DATABASE &lt;database&gt; MODIFY FILE (name =                                 &lt;filename&gt;, &lt;property&gt; = &lt;value&gt;)</pre> </li> </ul> <p>You can also view or change these properties from the Database Properties dialog box in the Enterprise Manager.</p>
Microsoft SQL Server service pack version and edition	Log into Management Studio and make the following query: <pre>select @@version.</pre>

## Port Allocation Options for Named Instances

When working with a named instance in SQL Server, there are two options for the instance's port allocation: dynamic port and static port.

- Dynamic Port.** When working with a dynamic port, a new port is assigned to the instance each time the instance is started. To enable clients to know the port when connecting to the instance, a service called SQL Server Browser needs to be started. The browser service listens to port 1434 and directs clients to the correct port according to the desired instance name.

You access the named instance by specifying the server name or IP address followed by `\<instance_name>`, for example, `server1\inst1`.

- Static Port.** When working with a static port, a specific port is assigned to the instance, and SQL



Server Browser service is not needed in order to connect to the instance (although you can still use it).

The default for SQL Server is a dynamic port.

**To work with a static port, follow the following steps:**

1. Select **Start > Programs > Microsoft SQL Server > Configuration Tools > SQL Server Configuration Manager**.
2. Expand Protocols for <instance name>, and double-click **TCP/IP**.
3. In the TCP/IP Properties dialog box, on the IP Addresses tab, several IP addresses appear, in the format IP1, IP2, up to IPAll. For each address:
  - a. If the TCP Dynamic Ports dialog is set to 0, indicating that the database engine is listening on dynamic ports, delete the 0.
  - b. In the TCP Port box, enter the port number on which you want this IP address to listen, and click **OK**.
4. In the console pane, click SQL Server Services.
5. In the details pane, right-click SQL Server (<instance name>) and click **Restart** to stop and restart SQL Server.

You access the named instance by specifying the server name or IP address followed by the port number, for example, **server1\1435**.

# Chapter 5: Creating and Configuring Microsoft SQL Server Databases

This chapter describes the creation and configuration of OMi databases on a Microsoft SQL Server.

This chapter includes:

- ["Creating Databases" below](#)
- ["Configuring Databases" on page 31](#)

## Creating Databases

This section includes the following topics:

- ["Creating Microsoft SQL Server Named Instances" below](#)
- ["Database Permissions" on page 28](#)
- ["Database File Layout" on page 28](#)
- ["System Databases" on page 30](#)

## Creating Microsoft SQL Server Named Instances

Before creating Microsoft SQL Server Named Instances, Microsoft SQL Server must be already installed on the system. For more information, see ["Installing Microsoft SQL Server"](#).

To create a Microsoft SQL Server Named Instance, follow these steps:

1. Logon as an administrator
2. Run the SQL Server installation files from the SQL server installation media (DVD).
3. Select **Installation**.
4. Select **New SQL Server stand-alone installation or add features to an existing installation**.
5. Click **OK**. The installation will search for the latest updates. You may choose to skip and check for updates later.
6. Click **Next**. An error message may appear stating that the windows firewall is active.  
You can enable the ports for remote database access later.
7. Select **Perform a new installation of SQL Server 2012** and click **Next**.
8. Type the product key and click **Next**.
9. Accept the license agreement and click **Next**.
10. Select all features with defaults and click **Next**.
11. In the Feature Selection page, remove Full Text Search from the list as OMi does not use the indexing search feature. Click **Next**.
12. In the Installation Rules page, click **Next**.

13. Select the following:

**Named Instance,**

**Instance name: <any name>**

**Instance ID**

**Note:** The Instance ID creates automatically when you click the field.

**Instance root directory: <path>**

14. Click **Next**

**Note:** Do not make any changes in the disk space requirements and the server configuration pages.

15. In the Database Engine Configuration page, select Mixed Mode (SQL Server authentication and Windows authentication) and specify the following:
  - a. Specify the password for the SA user
  - b. Specify SQL Server Administrators
16. Click **Next**
17. In the Analysis Service Configuration page, specify the Administrator with required privileges for Analysis Services.
18. In the Reporting Services Configuration page, select the following:
  - a. Reporting Services Native Mode
  - b. Install only (reporting services configuration is out of scope)
  - c. Reporting Services Sharepoint: Keep the default selection
19. Click **Next**.
20. Specify the administrators with privileges for the Distributed Replay Controller. Click **Next**
21. In the Distributed Replay Client page, type the name of the controller that the client computer will communicate with for the Distributed Replay client service.

**Note:** The name must be a fully qualified domain name (FQDN). For example, a host called server1 in the products hierarchy at Microsoft may have an FQDN of **server1.products.microsoft.com**. If you have already set up the controller, specify the name of the controller while configuring each client.

**Note:** If you have not set up the controller, you do not have to specify the controller name. However, you must manually specify the controller name in the client configuration file.

For the Working Directory and Result Directory fields, you can either keep the default path or choose a new path.

22. In the Error Reporting page, select the option if you want to send error reports to Microsoft.
23. In the Installation Configuration Rules page, click **Next**.
24. In the Ready to Install page, review the chosen configuration. After a few seconds, the summary page appears with the complete message. Click **Close**
25. Close the SQL server installation center.

## Database Permissions

To create a database, you must have CREATE DATABASE permissions. To grant CREATE DATABASE permissions to a user, the user's login must first be mapped to a database user in the master database.

**Note:** OMi login accounts should be mapped to **dbo** in the database. Members of the sysadmin server role automatically have CREATE DATABASE permissions, and are also mapped to **dbo** in all databases. A database owner is automatically mapped to **dbo** in the database.

To check whether a user has CREATE DATABASE permissions, log in to Management Studio with the login account of the user whose permissions you want to check, and run the following:

```
USE master
IF PERMISSIONS() & 1 = 1
PRINT 'User has CREATE DATABASE permissions.'
```

To check whether a user has enough permissions in the database, log in to Management Studio with the login account of the user whose mapping you want to check. Change the database context to the required database, and run the following command:

```
select case when IS_MEMBER ('db_owner')=1
or IS_SRVROLEMEMBER ('sysadmin')=1
or (IS_MEMBER ('db_ddladmin') = 1 and
IS_MEMBER ('db_datareader')=1 and
IS_MEMBER ('db_datawriter')=1 and
IS_MEMBER ('db_denydatareader')=0 and
IS_MEMBER ('db_denydatawriter')=0 )
then 'User has enough permissions'
else 'User does not have enough permissions'
end
```

## Database File Layout

When you create a database, it must consist of at least one data file (with an .mdf extension) and one transaction log file (with an .ldf extension). You can optionally create additional data files (.ndf), as well as additional log files (.ldf).

To enhance performance, you can create several data files. Microsoft SQL Server stripes the data among the data files, so that if you do not have RAID controllers that stripe your data, you can spread the data files over several regular physical disks and, in this way, have the data striped. The log, however, is read sequentially, so that there is no performance gain in adding more log files. An additional log file should be created on a different disk when your existing log is out of disk space.

## Data and Log Placement

**Important:**

- It is recommended not to place data or log files on the same disk that stores the page (swap) file.
- It is recommended that you place the data and log files on separate disk subsystems.
- **Log files.** Changes are not flushed to the database until they are written to the log, and the log architecture dictates serial writes. Therefore, it is advisable that there be as little interference as possible with the log activity. It is usually sufficient to place the log on a RAID 1 system because of the serial writes to the log. If you have processes reading from the log (for example, triggers accessing the inserted and deleted views which are formed from the log records or transactional replication), or several log files for different databases, consider placing the log file(s) on a RAID 1+0 (striped mirror) system.
- **Data files.** Data files should be placed on a RAID 1+0 system for optimal performance.

## File and Database Properties

When you create a database you can specify the following properties for each file (.mdf, .ndf, .ldf):

- **NAME.** The logical file name which you can use later when you want to alter one of the properties.
- **FILENAME.** The physical file path and name. Make sure the destination directory is not compressed (right-click the directory in Windows Explorer, select **Advanced**, and verify that the compression check box is not selected).
- **SIZE.** The initial file size.
- **MAXSIZE.** Determines the maximum size to which the file can grow. If this argument is omitted, or if you specify **Unlimited**, the file can grow until the disk is full.
- **FILEGROWTH.** The automatic growth increment of the file. This argument can be specified as either a percentage of the existing file size, or as a fixed size.

An autogrowth operation invoked by a modification sent by a client that timed out cannot be completed successfully. This means that the next time a client sends a modification, the autogrowth process starts at the beginning and may also time out.

To avoid this problem, it is recommended that you either expand the files manually every time the database nearly reaches full capacity (for example, 20 percent free), or set the growth increment to a fixed size that takes less time to be allocated than the client's timeout setting. Using a small growth increment is not recommended because it increases file system fragmentation. On the other hand, if you use a very large increment, modifications sent by clients might incur connection timeouts while waiting for the automatic expansion to finish. For large databases, a percentage growth increment can lead to exponential growth of the database and should be avoided.

## tempdb Database Settings

The frequent expansion of the tempdb system database can affect the database's performance, especially in large Microsoft SQL Server installations. The size of the tempdb, therefore, should be large enough to avoid the need for early expansion. The growth increment of the tempdb should be large enough to avoid fragmentation, yet not too large to expand in a reasonable amount of time. Create the tempdb with a minimum, initial size of 1 GB and with a growth increment of 50 MB. The tempdb database should be striped across several disks, ideally on a RAID 1+0 controller. It is recommended to move the tempdb database to its own set of disks.

To ensure that there is enough disk space for the tempdb to grow during times of heavy usage (for example, when aggregating or sorting data), it is recommended that you leave at least 20 GB free disk space on the drive where the tempdb is located.

## File Groups

File groups are logical groupings of data files. Each of the following objects can be placed in its own file group unit:

- A table's data
- A table's large objects (text, ntext, image columns)
- An index

Data is inserted proportionally into all files belonging to the file group in which the object is stored, according to the amount of free space in each file. The **.mdf** file is placed in a file group called **PRIMARY**, which is marked as **Default** when the database is created (the default file group for objects when no file group is specified). If you do not place other data files (**.ndf** files) in their own file groups, they are also placed in the **PRIMARY** file group. Note that you can change the **Default** file group later on.

File groups can be used for performance tuning or maintenance. For details, see Microsoft SQL Server Books Online at <http://www.microsoft.com/downloads>.

Following is an example that demonstrates how to use file groups for maintenance:

- **Partial Restoring.** Microsoft SQL Server does not support the restoration of a single table. Even if you place a single table in a file group, you cannot restore a file group to a point in time earlier than the rest of the data. Instead, you must apply all log file backups in order to synchronize the file group with the rest of the data. Microsoft SQL Server supports partial restoration to a database with a different name. A partial restoration allows you to restore a single file group, and supports point-in-time restoration. However, you must restore the **PRIMARY** file group because it contains the SYSTEM tables.

To restore a single table to a point in time if a logical error occurs, you need to design the file groups in your database as follows:

- Ensure that the **.mdf** file is the only file in the **PRIMARY** file group.
- Place each large table in its own file group.
- Place all small tables in a separate file group.

For additional data on creating file groups, see "[Microsoft SQL Server Data Storage Recommendation](#)" on page 91.

## System Databases

The following system databases are especially important for the smooth performance of Microsoft SQL Server:

- **tempdb.** Numerous Microsoft SQL Server activities—such as creating local and global temporary tables, creating work tables behind the scenes to spool intermediate query execution results, and sorting—implicitly or explicitly use the tempdb system database.

If your system is not configured properly, the tempdb database can become a performance bottleneck, so it is very important to determine the tempdb database's original size correctly.

For more information on setting database sizes, see ["tempdb Database Settings" on page 29](#).

To move tempdb's files, use the **ALTER DATABASE tempdb MODIFY FILE** command, and restart Microsoft SQL Server.

- **master, msdb, model.** These databases, although crucial for the operation of Microsoft SQL Server, are smaller than tempdb because they store only meta data.

It is strongly recommended to use a fault tolerant disk—ideally, RAID 1—for these databases.

**Note:** For OMi certification, place system databases on fault tolerant disks. It is recommended to use RAID 1 disks.

**To check the database's properties, run the following:**

```
EXEC sp_helpdb <database name>
```

## Configuring Databases

After you have created the necessary databases, you can add new files to the databases, change some of the existing database file properties, and set the database configuration options appropriately.

This section includes:

- ["Database File Configuration" below](#)
- ["Database Configuration Options" on the next page](#)

**Important:** If you enabled OMi to create OMi databases for you, it is strongly recommended that you configure them using the configuration instructions in this section.

## Database File Configuration

You can change certain database file properties, as well as add or drop files using either of the following methods:

- Use the Properties dialog box in Management Studio.
- Use the **ALTER DATABASE** command (for details, see Microsoft SQL Server Books Online).

## Adding Files

Data files can be added to an existing file group in a database, or to a new file group. There are no special restrictions or requirements.

## Dropping Files

To drop a file, you must first empty it using the **DBCC SHRINKFILE** command's **EMPTYFILE** option, which transmits the file data to all the other files in the file group. Once you empty the file, you can use the **ALTER DATABASE <database name> DROP FILE** command to drop it.

## Changing File Properties

You can change the size-related properties for all databases, as well as the filename property for the tempdb database (this takes effect after you restart Microsoft SQL Server). The `SIZE`, `MAXSIZE`, and `FILEGROWTH` properties can be changed using the `ALTER DATABASE tempdb MODIFY FILE` command. Note that the `SIZE` property can only be enlarged.

To shrink the file, use the `DBCC SHRINKFILE` command. For details and recommendations concerning file properties, see ["Creating Databases" on page 26](#).

## Database Configuration Options

Each database contains a set of configurable options that determine its behavior. You can view or change the database options using any one of the following utilities:

- The Options tab in the Management Studio's Properties dialog box
- The `EXEC sp_dboptions` stored procedure
- The `ALTER DATABASE <database name> SET` command

**Note:** Not all of the database configuration options are available in this dialog box.

The following table lists, in alphabetical order, the default configuration options, as well as the configuration settings required for OMi certification:

Configuration Option	Description	Default	OMi Certification in Microsoft SQL Server
ANSI NULL default (see note below)	Specifies whether the database columns are defined as NULL or NOT NULL, by default.	Not set	Not set
ANSI PADDING	Controls the way the column stores values shorter than the defined size of the column and the way the column stores values that have trailing blanks in char, varchar, binary, and varbinary data.	OFF	ON
Auto close	Specifies whether the database shuts down after its resources are freed and all users exit.	Not set	Not set If set, it may take a long time for the database to allocate resources every time a user connects, after the database is closed.



Configuration Option	Description	Default	OMi Certification in Microsoft SQL Server
Auto create statistics	Specifies whether missing statistics required by a query for optimization are built automatically during optimization.	Set	Set
Auto shrink	Specifies whether the database is automatically shrunk every hour, leaving 25% of free space.	Not set	Not set <b>Note:</b> If set, constant growth/ shrinkage may cause file system fragmentation.
Auto update statistics	Specifies whether out-of-date statistics required by a query for optimization are built automatically during optimization.	Set	Set
Compatibility level	The version of Microsoft SQL Server that the database appears to be (for the application).	The same version as the release installed	The same version as the release installed
Read only	The database is read only.	Not set (READ_WRITE)	READ_WRITE
Recovery	The database recovery model determines the recovery capabilities by controlling the amount of bulk operation logging (such as Select into, Bulk, Insert, Create index, LOB manipulation). The higher the recovery model, the higher the recovery capabilities. However, the amount of logging also increases, which may affect performance.	Full	Full (unless you are certain that the lower recovery capabilities are sufficient for your system)
Recursive triggers	Specifies whether recursive triggers are supported.	Not set	Not set
Restrict access	Only single users or members of the db_owner, dbcreator, or sysadmin groups can access the database.	Not set (MULTI_USER)	MULTI_USER
Page verify	Specifies whether incomplete pages can be detected.	Set	Set

Configuration Option	Description	Default	OMi Certification in Microsoft SQL Server
Quoted identifiers enabled	Specifies whether the Microsoft SQL Server enforces ANSI rules regarding quotation marks. Select this option to specify that double quotation marks be used only for identifiers, such as column and table names. Note that character strings must be enclosed in single quotation marks.	Not set	Not set

**Note:** Not all ANSI options can be set using Management Studio. The ANSI database configuration options include: ANSI\_NULLS, ANSI\_NULL\_DEFAULT, ANSI\_PADDING, ANSI\_WARNINGS, ARITHABORT, CONCAT\_NULL\_YIELDS\_NULL, NUMERIC\_ROUNDABORT, and QUOTED\_IDENTIFIER.

The options you set may not take effect, since these options can also be set at a higher level.

For example, if the session option **QUOTED\_IDENTIFIER** was turned on, the equivalent database configuration option is irrelevant. Some tools or database interfaces turn certain session options on or off, so that the database configuration options never take effect.

The following table summarizes the characteristics of each recovery model:

Model/Support	Allows Log Backup	Allows Point-in-Time/Log Mark Restoration	Allows Backup Log when Data Crashes (Saves changes until the crash point)	Amount of Bulk Operation Logging (can affect the performance of bulk operations)
Simple	No	No	No	Minimal
Bulk Logged	Yes	No	No	Minimal
Full	Yes	Yes	Yes	Full

To check your database's properties, run the command:

```
EXEC sp_helpdb <database name>
```

For information on SQL databases, see Microsoft SQL Server Books Online.

# Chapter 6: Maintaining Microsoft SQL Server Databases

This chapter describes the various maintenance tasks that are recommended for OMi databases created on Microsoft SQL Servers, such as backing up databases, checking database integrity and handling fragmentation, and monitoring databases.

This chapter includes:

- ["Backing Up Databases" below](#)
- ["Database Integrity and Fragmentation" on page 37](#)

## Backing Up Databases

Microsoft SQL Server supports three main types of database backup: full, differential, and log. It also supports file/file group backup, which is discussed in a separate section below. In order to develop a backup policy that provides the required recovery needs, it is important to thoroughly understand each backup type and the recovery model database configuration option explained in the previous section.

You can automate backup operations using Microsoft SQL Agent jobs. The Microsoft SQL Agent (represented by the SQLServerAgent service) is installed automatically when you install Microsoft SQL Server. Ensure that the Microsoft SQL Agent is configured to autostart in the operating system's Services applet when the server is started.

The following points are applicable to all backup types:

- The backup includes all changes made until the backup is complete.
- The backup can be performed online, but it is recommended to back up the database during periods of low activity, since the backup procedure can negatively impact on your system's performance.
- The following operations should not be performed during a backup procedure:
  - Adding or removing files
  - Shrinking the database
- The backup target can be a disk device (local or on a shared network that the Microsoft SQL Server service account needs permission to access) or tape (only local).

This section describes:

- ["Full Backup" on the next page](#)
- ["Differential Backup" on the next page](#)
- ["Log Backup" on the next page](#)
- ["File and File Group Backup" on page 37](#)
- ["Maintenance Plan" on page 37](#)
- ["Transaction Log Issues" on page 37](#)

## Full Backup

When you perform a full database backup, all information about the backed up database is contained within the backup, including data, metadata, and file information. A full backup is the basis for differential and log backups. With small databases, it is recommended to perform a full backup every day (for example, system databases that store mainly meta data). For large databases, it is generally recommended to have longer intervals between full backups (for example, once a week).

The storage requirements for a full backup are about the same as the storage requirements for the occupied data portion of the files. For example, if the total size of the data files is 20 GB, but only 15 GB are used (there are 5 GB of free space), the full backup size of the database should be approximately 15 GB.

## Differential Backup

You use a differential backup to back up the extents (blocks of eight contiguous 8K pages) that were changed since the last full backup. When restoring a database, you need only restore the last differential backup performed after the full backup.

After performing operations that affect large portions of data, such as index rebuilds or defragmentations, it is recommended that you perform a full backup. Otherwise, differential backups can become very large. For more information on index rebuilds and defragmentation, see ["Database Integrity and Fragmentation" on the next page](#).

Differential backup is usually scheduled at intervals between full backups. For example, if you perform a full backup once a week, you may want to perform a differential backup every day, or even several times a day.

The storage requirements for a differential backup are the total size of the extents (64 KB blocks) that were changed since the last full backup.

## Log Backup

A log backup—unlike full and differential backups which are mainly based on backing up an image of extents—backs up transactions from the transaction log and replays them upon restoration. In order to perform a log backup, the database must be set to the full or bulk-logged recovery model. If you want to perform a point-in-time or log mark restoration, or back up changes recorded in the log when the data crashes, you must set the database to the full recovery model. Otherwise, all changes made since the last performed backup are lost.

A log backup is incremental in nature, and backs up only the transactions performed since the previous log backup. When restoring a database, you must restore all log backups after the last differential (or full) backup you restored.

A log backup also marks the portion of the log that was backed up as available for reuse. In a database that is set to the full or bulk-logged recovery model, log portions that were not backed up cannot be reused. When the log is full, and Microsoft SQL Server cannot cycle to its beginning to reuse log space, it must expand. The frequency of your log backups, therefore, is a factor in determining the required size of the transaction log. Frequent log backups allow you to keep a smaller transaction log. It is recommended that you back up your log as frequently as possible, for example, every 30 minutes.

## File and File Group Backup

Instead of backing up the entire database, you can back up a file or file group. However, when you restore a single file or file group, you must apply all log backups up to and including the point of failure, in order to synchronize (same point-in-time) the file/file group with the rest of the database. This type of backup is generally useful with very large databases, for which you cannot frequently perform a full backup.

## Maintenance Plan

In the Microsoft SQL Server Management Studio, under the Management tree view, there is a graphic tool called Database Maintenance Plans. This tool allows you to define and automate common maintenance tasks (full and log backups, integrity checks, index rebuilds, and statistics collection).

## Transaction Log Issues

In terms of maintenance, the log is sensitive. When it is full, the log first tries to cycle and reuse inactive backed up log space, but if such space does not exist, the log tries to expand the file. If there is no room for the file to expand, Microsoft SQL Server rejects data modification requests.

To avoid log explosion, ensure that the log is large enough and that it is frequently backed up (ideally, by schedule). In addition, note that the active portion of the log starts with the oldest open transaction and continues until the current pointer in the log. The active portion cannot be reused or truncated. If a transaction remains open for a long time, it inevitably leads to log explosion at some point, even though the log is backed up.

To identify whether such a problem exists, run `DBCC OPENTRAN` to obtain the transaction that has been open for the longest period of time. To terminate the process running the transaction and roll back the transaction's activity, use the command:

```
KILL <process id>
```

**Note:** In Microsoft SQL Server, the `DBCCSHRINKFILE` command should always be successful.

## Database Integrity and Fragmentation

It is important to periodically check the physical integrity of your database objects, and to handle index fragmentation issues that are the main cause of performance degradation.

This section describes:

- ["Database Integrity" on the next page](#)
- ["Understanding File System Fragmentation" on the next page](#)
- ["Understanding Internal Fragmentation" on the next page](#)
- ["Understanding External Fragmentation" on page 39](#)
- ["Detecting and Handling Index Fragmentation" on page 39](#)

- ["Supplied Utilities for Monitoring and Rebuilding Indexes" on page 40](#)
- ["Distribution Statistics" on page 41](#)
- ["Collecting Statistics for RTSM" on page 42](#)
- ["Utility to Refresh Statistics" on page 42](#)

## Database Integrity

It is recommended that you run `DBCC CHECKDB` periodically to check the allocation and structural integrity of the objects in the database. You can automate and schedule the `DBCC CHECKDB` command using Microsoft SQL Agent jobs. Use the following command syntax:

```
DBCC CHECKDB ('database name')
```

**Note:** You can use the `WITH NO_INFOMSGS` option to reduce processing and tempdb usage. You can also run a quick physical-only test (page structure and record headers) using the `PHYSICAL_ONLY` option.

Because the Microsoft SQL Server database holds only schema locks (which prevent schema changes) and not data changes, the `DBCC CHECKDB` command can be run online. It is recommended, however, to run the `DBCC CHECKDB` command during periods of low activity, since it can negatively impact on your system's performance (`DBCC CHECKDB` is CPU- and disk-intensive, and uses tempdb for sorting).

## Understanding File System Fragmentation

File system fragmentation is relevant to all disk files, not just database files. It refers to the scattering of parts of the same disk file over different areas of the disk, as new parts of the file are added and existing parts are deleted. File system fragmentation slows disk access and degrades the overall performance of disk operations, although usually not severely.

To defragment a file system, you rewrite parts of a file to contiguous sectors on a hard disk. This increases the speed of data access and retrieval. To avoid fragmentation of your database files, create the files with as large an initial size as possible (so that they can accommodate changes in the future), and manually expand them with large increments as they become full.

If you cannot anticipate the future size of a database file, to avoid small, fragmented parts, use a large value as the file growth increment. Do not use too large a value, however, because this leads to client request timeouts when the file autogrows (for more details, see ["Database Permissions" on page 28](#)). In addition, avoid using the autoshrink database option, because it increases the chances of fragmentation as the database files continually shrink and grow.

**Note:** It is recommended that you periodically run a defragmentation utility on the database.

## Understanding Internal Fragmentation

Internal fragmentation refers to the percentage of data contained in the pages. In environments such as the OMi system, which are characterized by transactions that frequently insert data, internal fragmentation is sometimes initiated in anticipation of new data in indexes and can be a positive

occurrence. By leaving a certain percentage of the index pages free, you can avoid page splits for a certain period of time. This is especially significant for clustered indexes, because they contain the actual data pages. You can achieve internal fragmentation by periodically rebuilding your indexes using the `CREATE INDEX` command, with the `DROP_EXISTING` and `FILLFACTOR` options, or the `ALTER INDEX REBUILD` command (online or offline) and `FILLFACTOR` option. The `FILLFACTOR` option specifies the fullness of the leaf level index pages.

## Understanding External Fragmentation

As page splits occur in your indexes, new allocated pages are acquired from the database file. Ideally, a page split should yield the allocation of a page contiguous to the one that split. However, in practice, the space contiguous to the split page is usually already occupied. The more page splits that occur, the less the index's linked list reflects the physical layout of the pages on disk, and the greater the amount of external fragmentation.

External fragmentation impacts negatively on the performance of ordered index scans because the disk arm needs to move back and forth in order to retrieve the pages from disk. Ideally, the linked list should reflect the physical layout of the pages on disk so that when an ordered index scan is performed, the disk arm moves in one direction as it retrieves the pages from disk.

You can handle external fragmentation proactively by initiating internal fragmentation and leaving a certain percentage of the leaf level index pages free, thus avoiding page splits for a certain period of time. As mentioned earlier, internal fragmentation can be achieved by periodically rebuilding your indexes using the `FILLFACTOR` option. You can also handle external fragmentation by checking the external fragmentation status of your indexes, and rebuilding the indexes.

## Detecting and Handling Index Fragmentation

### Using `sys.dm_db_index_physical_stats` to Detect Fragmentation

The dynamic management function `sys.dm_db_index_physical_stats` is used to determine the degree of fragmentation of an index. You can detect fragmentation in a specific index, in all indexes on a table or indexed view, in all indexes in a specific database, or in all indexes in all databases. For partitioned indexes, `sys.dm_db_index_physical_stats` also provides fragmentation information for each partition.

Fragmentation of a table occurs through the process of data modifications (`INSERT`, `UPDATE`, and `DELETE` statements) that are made against the table and to the indexes defined on the table. Because these modifications are not ordinarily distributed equally among the rows of the table and indexes, the fullness of each page can vary over time. For queries that scan part or all of the indexes of a table, this kind of fragmentation can cause queries to return slower.

The table syntax of `sys.dm_db_index_physical_stats` is:

```
sys.dm_db_index_physical_stats (  
  { database_id | NULL }  
  , { object_id | NULL }  
  , { index_id | NULL | 0 }  
  , { partition_number | NULL }  
  , { mode | NULL | DEFAULT }  
  )
```

Run this command as regular select statement from table, for example:

```
select * from sys.dm_db_index_physical_stats (DB_ID('<BAC_database>'),  
object_id('<Table_Name>'), NULL, NULL, 'SAMPLED')
```

The **Avg\_fragmentation\_in\_percent** column returned by `sys.dm_db_index_physical_stats` is the logical and extent fragmentation of the index.

For more information about **sys.dm\_db\_index\_physical\_stats**, see the Microsoft SQL Server documentation.

## Handling Fragmentation

The table column **Avg\_fragmentation\_in\_percent** returned by `sys.dm_db_index_physical_stats` reflects the degree of fragmentation per index. Based on this value, you can determine whether to handle the fragmentation and which method to use in handling it.

Use the following rough guidelines to determine the best method to correct the fragmentation:

- **Between 5% and 30%.** Use the `ALTER INDEX REORGANIZE` command to reorganize the index. Index reorganization is always executed online.
- **Greater than 30%.** Use the `ALTER INDEX REBUILD` command to rebuild the index. Index rebuilding can be executed online or offline. Rebuild the indexes online to achieve the availability similar to that of the Reorganize option.

Very low levels of fragmentation (less than 5%) should not be addressed by either of these commands due to the negligible benefit gained from removing such a small amount of fragmentation.

For more information on defragmenting indexes, see "Reorganizing and Rebuilding Indexes" under Designing and Implementing Structured Storage in SQL Server Books Online:

<http://technet.microsoft.com/en-us/library/ms189858.aspx>.

**Note:** It is strongly recommended that you create an automatic index rebuild task for each of the OMi databases. In the RTSM and event databases, run this task at least once a day.

## Supplied Utilities for Monitoring and Rebuilding Indexes

OMi provides two utilities that can be used to detect and rebuild fragmented indexes. The **rebuild\_fragmented\_indexes.bat** utility uses the Logical scan fragmentation and Scan Density criteria to detect, and if instructed to rebuild, fragmented indexes. The operation of listing the fragmented tables has a very small impact on system performance and can be executed on line. The operation of rebuilding the indexes usually hinders performance as tables are partially locked during the process, and CPU and I/O are heavily utilized. It is recommended to rebuild the indexes in a maintenance window. The utilities should be run by a database administrator.

The utilities for Microsoft SQL Server are located in the **<Gateway Server root directory>\AppServer\webapps\site.war\DataBases\SQL\_Svr\_DB\_Utills** directory on the OMi Gateway Server.

## Utility to Rebuild All Indexes in Database

The **rebuild\_indexes.bat** utility runs through all tables in the database and rebuilds related indexes.



### To run the `rebuild_indexes.bat` utility:

Execute `rebuild_indexes.bat` with the following parameters:

- SQL Server name
- Database name
- SA password

#### Example:

```
rebuild_indexes.bat SQL_SRVR_3 OMi_DB_3 ad%min52.
```

Output from the procedure is located in the `rebuild_indexes.log` file in the same directory.

## Utility to Rebuild Indexes Based on the Fragmentation Level of Each Index

The `rebuild_fragmented_indexes.bat` utility has two working modes:

- List fragmented tables. In this mode, a list of the fragmented tables (that is, tables with over 30% fragmentation) is returned, together with the commands needed to rebuild the tables at a later time.
- Rebuild fragmented tables. In this mode, all fragmented tables (that is, tables with over 30% fragmentation) are rebuilt.

### To run the `rebuild_fragmented_indexes.bat` utility:

Execute `rebuild_fragmented_indexes.bat` with the following parameters:

- SQL Server name
- Database name
- SA password
- Working mode – 0 to provide a rebuild script for later use; 1 to rebuild indexes automatically.

#### Example:

```
rebuild_fragmented_indexes.bat SQL_SRVR_3 OMi_DB_3 ad%min52 1
```

Output from the procedure (a list of fragmented tables and the rebuild commands) is located in the `rebuild_indexes.log` file in the same directory.

## Distribution Statistics

Microsoft SQL Server allows statistical information regarding the distribution of values in a column to be created. This statistical information can be used by the query processor to determine the optimal strategy for evaluating a query. When an index is being created, SQL Server automatically stores statistical information regarding the distribution of values in the indexed columns. The query optimizer in SQL Server uses these statistics to estimate the cost of using the index for a query. As the data in a column changes, index and column statistics can become out of date and cause the query optimizer to make less than optimal decisions on how to process a query.

It is recommended to update index statistics daily to provide the query optimizer with up to date information about the distribution of data values in the tables. This allows the query optimizer to make

better judgments about the best way to access data, as it has more information about the data stored in the database.

Whether the auto update statistics database option is enabled or disabled, it is strongly recommended that you create an automatic task to update statistics for each of the OMi databases on a daily basis, as the data is frequently changed. The job should execute the `sp_updatestats` API against the specific database.

## Collecting Statistics for RTSM

Unlike some databases where queries are predefined and can be modified according to the expected database size, the RTSM database constructs queries dynamically, according to the pattern views defined against its data model. This necessitates accurate statistics at all times. In addition to running a daily job to update statistics, it is recommended that you manually refresh statistics if major changes to the RTSM schema objects have occurred, usually as a result of bulk insert transactions.

An example of a scenario that warrants a manual refresh of the RTSM statistics is automated Data Flow Management (DFM) jobs. DFM is the process responsible for automatically detecting configuration items and inserting them into the RTSM.

## Utility to Refresh Statistics

The `update_statistics.bat` utility has two working modes:

- List tables with out of date statistics. In this mode, a list of the tables is returned, together with the commands needed to update statistics the tables at a later time.
- Update statistics on the tables. In this mode, all tables with outdated statistics are being updated.

**To run the `update_statistics.bat` utility:**

From `<Gateway Server root directory>\AppServer\webapps\site.war\DataBases\SQL_Svr_DB_Utills`, execute `update_statistics.bat` with the following parameters:

- SQL Server name.
- Database name.
- SA password.
- Working mode - 0 to provide a script for later use; 1 to update statistics automatically.

### **Example:**

```
update_statistics.bat SQL_SRVR_3 RTSM_DB_3 ad%min52 1
```

Output from the procedure (a list of tables and the update commands) is located in the `update_statistics.log` file in the same directory.

# Chapter 7: Using Windows Authentication to Access Microsoft SQL Server Databases

Unless configured otherwise, OMi uses Microsoft SQL Server authentication to access Microsoft SQL Server databases. However, Windows authentication can also be used.

This chapter describes how to enable OMi to use Windows authentication to access Microsoft SQL Server databases.

This chapter includes:

["Enabling OMi to Work with Windows Authentication" below](#)

## Enabling OMi to Work with Windows Authentication

You can enable OMi to use Windows authentication instead of Microsoft SQL Server authentication to access any of the OMi databases.

To enable OMi to use Windows authentication to access a Microsoft SQL database, you must:

- Configure the Microsoft SQL Server to use Windows authentication.
- Launch the OMi Server service on all the OMi servers with a Windows user that has the necessary permissions to access the Microsoft SQL database.
- Use the configuration wizard to create or connect to a Microsoft SQL database and specify the use of Windows authentication.

This section includes the following topics:

- ["Configuring Microsoft SQL Server to Use Windows Authentication" below](#)
- ["Launching the OMi service with a Windows User" below](#)
- ["Creating or Connecting to a Microsoft SQL Database Using Windows Authentication" on the next page](#)

## Configuring Microsoft SQL Server to Use Windows Authentication

In the SQL Server Enterprise Manager, select **Security > Logins**, right-click and choose **New Login**. Enter the desired domain account, including the domain name, in the following pattern: DOMAIN\USER (for example, MY\_DOMAIN\MY\_ACCOUNT).

In the **Server Roles** tab, select **System Administrators** and click **OK**.

## Launching the OMi service with a Windows User

By default, the OMi service is run as a system service. If you have configured your Microsoft SQL Server to use Windows authentication, you must change the user running the OMi service to the same Windows user you defined for the Microsoft SQL Server in order to enable the service user to access the database.

The user you assign to run the service must have the following permissions:

- Sufficient database permissions (as defined by the database administrator)
- Sufficient network permissions
- Administrator permissions on the local server

**To change the OMi service user:**

1. Disable OMi (**Start > Programs > HPE Operations Manager i > Administration > Disable Operations Manager i**).
2. In Microsoft's Services window, double-click **HPE Operations Manager i**. The HPE Operations Manager i Properties (Local Computer) dialog box opens.
3. Click the **Log On** tab.
4. Select **This account** and browse to choose the user you previously defined for your Microsoft SQL Server.
5. Enter the selected user's Windows password and confirm this password.
6. Click **Apply** to save your settings and **OK** to close the dialog box.
7. Enable OMi (**Start > Programs > HPE Operations Manager i > Administration > Enable Operations Manager i**).

## Creating or Connecting to a Microsoft SQL Database Using Windows Authentication

You create or connect to a database using the configuration wizard. To create or connect to a Microsoft SQL database using Windows authentication, you must select this option within the configuration wizard. For details on using the configuration wizard, see the OMi Installation and Upgrade Guide.

# Part III: Deploying and Maintaining the Oracle Server Database

# Chapter 8: Oracle Server Deployment Overview

You can set up OMi schemas on an Oracle Server. This chapter describes the following topics related to deploying Oracle Servers for use with OMi:

- ["About Oracle Server Deployment" below](#)
- ["System Requirements" below](#)
- ["Character Set Support in Oracle" on the next page](#)

## About Oracle Server Deployment

To deploy Oracle Server for use with OMi, perform the following procedures:

- **Install Oracle Server.**

For details on Oracle software installation, see the installation guide in the documentation for your specific Oracle platform. For software installation options, see ["Oracle Server Configuration and Sizing Guidelines" on page 53](#) and ["Oracle Summary Checklist" on page 72](#).

**Note:** If you are installing Oracle 12c, in the Typical Installation step, it is recommended that you clear the **Create as Container database** option. This ensures that the database is created without CDB and will be the same as in Oracle 11.

For a *silent* or *response file* setup, use the **createDatabase** operation (with the corresponding configuration section) instead of **createPluggableDatabase** in your DBCA response file.

```
OPERATION_TYPE = "createDatabase"  
[CREATEDATABASE]
```

- **Build a database on Oracle Server to store OMi data.**

For instance configuration and sizing guidelines, see ["Oracle Server Configuration and Sizing Guidelines" on page 53](#). For details on database instance installation, see the installation guide in the documentation for your specific Oracle platform.

- **Create one or more Oracle tablespaces to store OMi data.**

For details, see ["Oracle Tablespaces" on page 56](#).

- **Create an Oracle user schema for OMi schemas.**

You can create OMi user schemas manually, or you can use the configuration wizard to create the schemas for you. For details on creating an Oracle user schema for OMi, see ["Creating OMi Schemas Manually" on page 49](#).

## System Requirements

This section describes the system requirements for working with Oracle Server in conjunction with OMi.

This section includes the following topics:

- ["Hardware Requirements" below](#)
- ["Software Requirements" below](#)
- ["Oracle Instances" below](#)

## Hardware Requirements

For OMi hardware sizing guidelines, see ["Requirements" on page 11](#).

For Oracle hardware requirements, see the installation guide for your specific Oracle platform. Additional information is also available in the Oracle software distribution media as well as the online Oracle documentation. For Oracle documentation, see:

<http://otn.oracle.com/documentation/index.html>

## Software Requirements

For a list of supported Oracle database versions, see the OMi Support Matrix at:

<https://softwaresupport.hp.com/group/softwaresupport/search-result/-/facetsearch/document/KM323488>

## Oracle Instances

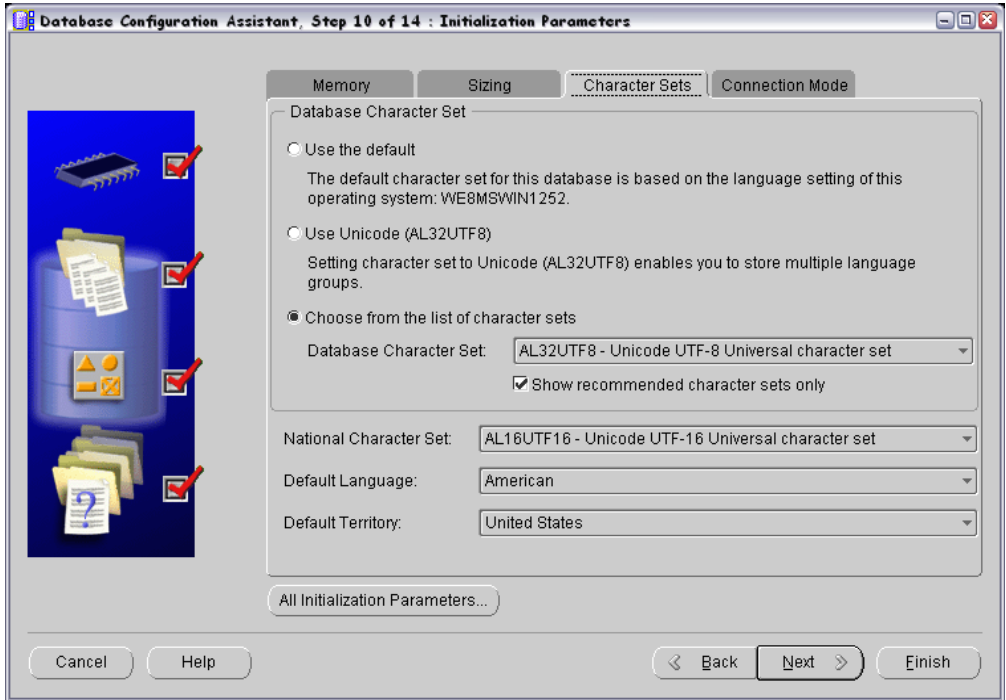
You can install more than one Oracle instance on a machine using the same Oracle database engine.

For OMi certification, do not use more than one Oracle instance. If you do use more than one instance for the OMi databases, ensure that all the instances are configured as described in this guide and that they all have the same characteristics (such as the same character set).

## Character Set Support in Oracle

When you create a new Oracle instance in an Oracle database, you must specify the character set for the instance. All character data, including the data in the data dictionary, is stored in the character set of the instance.

During the database creation, you can choose from the list of available character sets.



A UTF-8 character set is required for Oracle instances. There are two Oracle character sets that implement UTF-8 encoding, UTF8 and AL32UTF8.

It is recommended to set the client locale by setting the NLS\_LANG environment variables to appropriate values, so that database client applications can communicate with the database server.

Language	NLS_LANG Value
English	american_america.AL32UTF8
French	french_france.AL32UTF8
Spanish	spanish_spain.AL32UTF8
German	german_germany.AL32UTF8
Russian	russian_russia.AL32UTF8
Japanese	japanese_japan.AL32UTF8
Korean	korean_korea.AL32UTF8
Simplified Chinese	simplified chinese_china.AL32UTF8



# Chapter 9: Manually Creating the Oracle Schemas

This chapter describes how to manually create and populate OMi user schemas. It also contains information on Oracle permissions for OMi user schemas.

This chapter includes:

- ["Oracle Server Databases Overview" below](#)
- ["Creating OMi Schemas Manually" below](#)
- ["OMi Oracle Schema Privileges" on page 52](#)

## Oracle Server Databases Overview

OMi uses the Management, RTSM, and Event databases.

During the OMi setup, new databases can be set up automatically by the setup procedure, or already existing databases can be used. Existing databases can either be created manually in advance (for example, due to organization security restrictions), or can be created by a previous installation of the same release of OMi.

For details about installing OMi server, see the OMi Installation and Upgrade Guide.

## Creating OMi Schemas Manually

The OMi user schemas can be created manually or automatically using the configuration wizard. For more information on the configuration wizard, see the OMi Installation and Upgrade Guide.

**Note:** You must use the same password for all schemas, which you then enter in the Database Settings page of the configuration wizard when connecting OMi to the database.

The configuration wizard prompts you to supply the following parameters:

- **Host name.** The name of the host machine on which Oracle Server is installed.
- **Port.** The Oracle listener port. OMi automatically displays the default port, **1521**.
- **SID.** The Oracle instance name that uniquely identifies the Oracle database instance being used by OMi.
- **Password.** The password of a user that is used by OMi to connect to and retrieve data from the OMi databases.
- **Admin user name and password.** (to connect as an administrator) The name and password of a user with administrative permissions on Oracle Server (for example, a System user).
- **Temporary tablespace.** The name of the temporary tablespace you assigned to the user schemas. The default Oracle temporary tablespace is **temp**.
- **Default tablespace.** The name of the dedicated default tablespace you created for the user

schemas. For details on creating a dedicated OMi tablespace, see ["Oracle Tablespaces" on page 56](#).

- **Database Names (Management, RTSM, and Event).** The names of the existing user schemas, or the names that you choose for the new user schemas (for example, OMi\_Management).

**Note:** If you are using the SQL\*Plus Oracle client, make sure you are aware of its limitations. For example, a database creation command may exceed the maximum command-line length supported by SQL\*Plus. For information about the limitations of SQL\*Plus, see the Oracle documentation.

This section includes the following topics:

- ["Creating a Management User Schema Manually" below](#)
- ["Creating an RTSM User Schema Manually" on the next page](#)
- ["Creating an Event Schema Manually" on the next page](#)

## Creating a Management User Schema Manually

It is recommended that a certified database administrator creates the management user schema manually (using your organization's database characteristic methodology). You then connect to the management user schema later in the configuration wizard. For more information on the configuration wizard, see the OMi Installation and Upgrade Guide.

**Note:** If you cannot submit database administrator connection parameters over a non-secure connection, you may need to configure a management user schema manually.

### To generate and run the Management schema SQL script:

1. Create the user schema:

Create a user schema with the permissions described in ["OMi Oracle Schema Privileges" on page 52](#).

2. Generate the Management schema SQL script:

Run the following script file located in the `<OMi_HOME_GW>/dbverify/bin` directory:

For Windows:

```
generate_create_script.bat [deployment_type] [Generated Script Name] [dbType] [sessionType]
```

For Linux:

```
./generate_create_script.sh [deployment_type] [Generated Script Name] [dbType] [sessionType]
```

The following parameters should be used for the Management schema:

- **deployment\_type.** EE (Enterprise Edition)
- **Generate Script Name.** The name of the SQL file to generate, for example, `management.sql`.

- **dbType.** Oracle
- **sessionType.** Management

**Example:**

```
/opt/HP/BSM/dbverify/bin/generate_create_script.sh EE Management.sql Oracle  
Management
```

3. Run the Management schema SQL script:
  - a. Connect to the Management schema and run the script generated in step 2 to deploy the Management objects.
  - b. Run the following SQL script located in `<OMi_HOME_GW>/AppServer/webapps/site.war/DataBases/ORA_DB_Utils:`  
  
`management_ora_inserts.sql`
  - c. When the script is finished running, commit the change.

## Creating an RTSM User Schema Manually

It is recommended that a certified database administrator creates the RTSM user schema manually (using your organization's database characteristic methodology). You then connect to the RTSM user schema later in the configuration wizard. For more information on the configuration wizard, see the OMi Installation and Upgrade Guide.

**To create a separate RTSM user schema manually:**

Create a user schema with the permissions described in "[OMi Oracle Schema Privileges](#)" on the next page.

## Creating an Event Schema Manually

Create an Event schema with the permissions described in "[OMi Oracle Schema Privileges](#)" on the next page.

**To generate the Event schema scripts:**

Run one of the following script files:

For Windows:

```
<OMi_HOME>\opr\bin\opr-schema-script-generator.bat
```

For Linux:

```
<OMi_HOME>/opr/bin/opr-schema-script-generator.sh
```

**Note:** The database type parameter, `-d`, should be set to **ORACLE**.

For example, on Windows, to generate a script called `event.sql` that can be used to create the necessary tables on the Oracle database, run the following command:

```
<OMi_HOME>\opr\bin\opr-schema-script-generator.bat -d ORACLE -o C:\temp\event.sql
```

**To run the Event database SQL script:**

1. Create the Event database.
2. Connect to the Event database and run the script generated in the previous section to deploy the Operations objects.

**Note:** Perform the above procedures only if you are an experienced Oracle Database administrator.

## OMi Oracle Schema Privileges

When using an Oracle user schema, you can perform any database action that is allowed by the Oracle permission granted to the user. The following list (taken from the `oracle_user_create.sql` script located in the `<Gateway Server root directory>/AppServer/webapps/site.war/DataBases/ORA_DB_Utils` directory) describes the required database permissions that must be granted to the OMi user. These permissions are also used by the OMi platform components to create a new Oracle user:

- `GRANT CONNECT TO <OMi Oracle user schema>`
- `GRANT CREATE SEQUENCE TO <OMi Oracle user schema>`
- `GRANT CREATE TABLE TO <OMi Oracle user schema>`
- `GRANT CREATE TRIGGER TO <OMi Oracle user schema>`
- `GRANT UNLIMITED TABLESPACE TO <OMi Oracle user schema>`
- `GRANT CREATE VIEW TO <OMi Oracle user schema>`
- `GRANT CREATE PROCEDURE TO <OMi Oracle user schema>`
- `ALTER USER <OMi Oracle user schema> DEFAULT ROLE ALL`

If the schema is being used for the RTSM database, then the following permission is also needed:

`GRANT CREATE TYPE TO <OMi Oracle user schema for RTSM>`

**Note:** OMi supports any user with higher permissions. For OMi certification, use an Oracle user that has the exact Oracle permissions described above.

# Chapter 10: Oracle Server Configuration and Sizing Guidelines

This chapter contains guidelines for the Oracle database configuration and storage settings that should be used when working with Oracle Server and OMi. Note that the recommended settings differ according to the size of your OMi deployment.

This chapter includes:

- ["Oracle Parameter Settings" below](#)
- ["Oracle Tablespaces" on page 56](#)
- ["Oracle Tablespace Settings" on page 57](#)
- ["Using RAID Configuration" on page 60](#)

## Oracle Parameter Settings

The table below describes the recommended values for a number of Oracle database initialization parameters, when working with the OMi database server.

**Note:** If any of the database initialization parameters listed below is deprecated in your Oracle version, you can safely ignore it.

Parameter Name	OMi Deployment		Remarks
	Small	Large	
DB_BLOCK_SIZE	8K	8K-16K	Should be a multiple of the operating system block size.
DB_CACHE_ADVICE	ON	ON	For gathering statistics when tuning is required.
SGA_TARGET	1 GB	Minimum 4 GB	<p>Setting this parameter configures Oracle to automatically determine the size of the buffer cache (<b>db_cache_size</b>), shared pool (<b>shared_pool_size</b>), large pool (<b>large_pool_size</b>), java pool (<b>java_pool_size</b>), and streams pool (<b>streams_pool_size</b>).</p> <p>The value configured for SGA_TARGET sets the total size of the SGA components.</p> <p>When SGA_TARGET is set (that is, its value is not 0), and one of the above pools is also set to a non-zero value, the pool value is used as the minimum value for that pool.</p>

Parameter Name	OMi Deployment		Remarks
	Small	Large	
MEMORY_TARGET	2 GB	Minimum 5 GB	Automatic Memory Management enables the entire instance memory to be automatically managed and tuned by the instance. The instance memory contains the System Global Area (SGA) and the Program Global Area (PGA). MEMORY_TARGET is the only required memory parameter to set. However, it is recommended to set SGA_TARGET or PGA_AGGREGATE_TARGET as well to avoid frequent resizing of the SGA and PGA components. The values entered for SGA_TARGET and PGA_AGGREGATE_TARGET serve as minimum values.
LOG_BUFFER	1 MB	5 MB	
DB_FILE_MULTIBLOCK_READ_COUNT	Oracle default value	Oracle default value	
PROCESSES	200	500	
SESSIONS	225	445	$(1.1 * PROCESSES) + 5$
OPTIMIZER_INDEX_COST_ADJ	100	100	When set to 100, the optimizer evaluates the index access path at the regular cost. When set to any other value (for example, 50), the access path is evaluated at that percentage of the regular cost.
TIMED_STATISTICS	True	True	
LOG_CHECKPOINT_INTERVAL	0	0	
LOG_CHECKPOINT_TIMEOUT	0 or $\geq 1800$ (greater than or equal to 1800)	0 or $\geq 1800$ (greater than or equal to 1800)	
OPTIMIZER_MODE	ALL_ROWS	ALL_ROWS	

Parameter Name	OMi Deployment		Remarks
	Small	Large	
CURSOR_SHARING	Exact	Exact	
OPEN_CURSORS	2000	2000	
COMPATIBLE	The same as the Oracle Server release installed	The same as the Oracle Server release installed	For example, 11.2.0 for Oracle 11g R2 or 12.0.0 for Oracle 12c.
UNDO_MANAGEMENT	Auto	Auto	
UNDO_RETENTION	Oracle default value	Oracle default value	Automatic tuning is performed.
RECYCLEBIN	Off	Off	
NLS_LENGTH_SEMANTICS	BYTE	BYTE	This parameter controls the length definition of varchar columns.
WORKAREA_SIZE_POLICY	AUTO	AUTO	
PGA_AGGREGATE_TARGET	400 MB	Minimum 1 GB	
STATISTICS_LEVEL	TYPICAL	TYPICAL	Enables tuning if required.
OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES	FALSE	FALSE	Controls Automatic Plan Capture as part of Oracle SQL Management Base (SMB).
AUDIT_TRAIL	NONE	NONE	Out-of-the-box database auditing is written to the SYS.AUD\$ audit trail table. It is advisable to change this value to none to avoid growth of the system tablespace.
BLANK_TRIMMING	False	False	

Parameter Name	OMi Deployment		Remarks
	Small	Large	
FIXED_DATE	<i>Do not set this value.</i>	<i>Do not set this value.</i>	OMi uses the SYSDATE function for generating system time as part of the application process.
_PARTITION_ LARGE_ EXTENTS	FALSE	FALSE	Relevant for Oracle 11.2.0.2 and Oracle 12c.  When this hidden parameter is set to TRUE, it affects the size of partitions in native partitioned tables. The initial extent allocated for each partition is very large, thus causing unwanted growth of database data files.
DEFERRED_ SEGMENT_ CREATION	FALSE	FALSE	Relevant for Oracle 11.2.0.* only.

## Oracle Tablespaces

An Oracle tablespace is an Oracle object that is a logical container of database objects, for example, tables, indexes, and so forth. When working with OMi, you must create one or more dedicated default tablespaces for your OMi user schemas. You may also want to create a dedicated temporary tablespace for OMi. To create a tablespace, you must provide specific operating system files that physically represent the tablespace, as well as extent parameters.

When mapping operating system files, there is an option to make the file auto-extendable. This feature is supported by OMi, but not certified for use with OMi, since it can cause the system to consume all available disk space.

**Note:** The following terms all refer to the identical database: Event schema = OPR Schema = Event database.

This section includes:

- ["Locally Managed Tablespaces" below](#)
- ["Creating an Oracle Tablespace" on the next page](#)

## Locally Managed Tablespaces

Locally managed tablespaces can have either uniform extent sizes, or variable extent sizes that are determined automatically by the system. When you create the tablespace, the **uniform** or **autoallocate** (system-managed) option specifies the type of allocation.

For system-managed extents, Oracle determines the optimal size of extents, with a minimum extent size of 64 KB. This is the default extent size for permanent tablespaces.

For uniform extents, you can specify an extent size, or use the default size, which is 1 MB. Temporary tablespaces that manage their extents locally can only use this type of allocation.



Note that the NEXT, PCTINCREASE, MINEXTENTS, MAXEXTENTS, and DEFAULT STORAGE storage parameters are not valid for extents that are managed locally.

All data and temporary tablespaces should be locally managed when working with OMi.

For information on locally managing temporary tablespace using TEMPFILE, see ["Temporary Tablespace Settings" on page 59](#).

## Creating an Oracle Tablespace

You create an OMi Oracle tablespace using the **oracle\_tablespace\_create.bat** (for a Windows installation) or the **oracle\_tablespace\_create.sh** script (for a Linux installation). Both scripts can be found in **<OMi server>\AppServer\webapps\site.war\DataBases\ORA\_DB\_Utills**.

**Note:** This script is a basic script for tablespace creation that contains one data file. You can edit this file according to the size of your OMi installation. For more information, see ["Data Tablespace Settings" on the next page](#).

### To create an OMi Oracle tablespace:

Run the following command from the directory in which the **oracle\_tablespace\_create.bat** or **oracle\_tablespace\_create.sh** script is located:

```
oracle_tablespace_create [admin_user] [admin_password] [tns_entry_name]  
[tablespace_name] [file_name] [file_size]
```

For example:

```
oracle_tablespace_create system manager topaz omi_tablespace  
h:\oracle\oradata\topaz01.ora 50M
```

- **[admin\_user]**. Name of user with administrative permissions on Oracle Server.
- **[admin\_password]**. Password of specified user.
- **[tns\_entry\_name]**. The TNS name specified in the tnsnames.ora file on the local Oracle Client.
- **[tablespace\_name]**. Name of tablespace to create.
- **[file\_name]**. File name to be created, including full path to file.
- **[file\_size]**. File size; use M for MB and K for KB. Note that the actual size needed varies depending on the amount of data OMi generates and the amount of time you keep historical data.

**Caution:** OMi does not support the table compression feature of Oracle Database. Therefore, when creating Oracle tablespaces, make sure table compression is disabled.

## Oracle Tablespace Settings

This section describes the storage settings and file sizing guidelines for data tablespaces, temporary tablespaces, redo logs, and undo tablespaces.

This section includes the following topics:

- ["Data Tablespace Settings" on the next page](#)
- ["System Tablespace Settings" on the next page](#)

- ["Temporary Tablespace Settings" on the next page](#)
- ["Redo Log Settings" on the next page](#)
- ["Undo Segment Settings" on the next page](#)

## Data Tablespace Settings

The following table specifies the recommended sizes for OMi tablespaces:

Tablespace	OMi Deployment		Remarks
	Small	Large	
Management	5 GB	10 GB	The specified size is a minimum requirement.
RTSM	1 GB	20 GB	The specified size is a minimum requirement.
Event	5 GB	20 GB	The specified size is a minimum requirement.

## Data Tablespace Default Storage Settings

The storage settings for data tablespaces should be:

- Locally managed tablespace
- Automatic segment space management
- Automatic local extent management

## System Tablespace Settings

The following table specifies the recommended settings for system tablespaces:

Tablespace	OMi Deployment	
	Small	Large
SYSTEM	2 GB	5 GB
SYSAUX	2 GB	5 GB

**Note:** Make sure you do not use the SYSTEM and SYSAUX tablespaces for the OMi schemas.

The System Tablespace storage default settings should be:

- Locally managed tablespace
- Segment space management:
  - **SYSAUX:** Automatic
  - **SYSTEM:** Manual
- Automatic local extent management

## Temporary Tablespace Settings

The following table specifies the recommended settings for temporary tablespaces:

Tablespace	OMi Deployment		Remarks
	Small	Large	
TEMP	1 GB	10 GB	Use multiple files with large tablespaces.
TEMP storage settings	Uniform allocation: 2 MB	Uniform allocation: 2 MB	<ul style="list-style-type: none"> <li>Should be locally managed (Uniform allocation).</li> <li>Tablespaces should be of a temporary type (use of TEMPFILE).</li> <li>Segment space management in temporary tablespaces is manual.</li> </ul>

## Redo Log Settings

The following table specifies the recommended settings for redo log files:

Setting	OMi Deployment	
	Small	Large
Redo log file size	200 MB	500 MB - 1 GB
Minimum number of groups	4	4
Minimum number of members per group	2	2

## Undo Segment Settings

The following table specifies the recommended Undo settings:

Setting	OMi System Profile		Remarks
	Small	Large	
Undo tablespace size	1 GB	10 GB	The number of segments, the minimum number of extents, and the rollback segment size (initial, next) are all set automatically by Oracle.
UNDO_ MANAGEMENT parameter	AUTO		Oracle default values.

Setting	OMi System Profile		Remarks
	Small	Large	
UNDO_ RETENTION parameter	Oracle default value		

The Undo Tablespace storage default settings should be:

- Locally managed tablespace
- Automatic local extent management
- Segment space management in undo tablespaces is manual

## Using RAID Configuration

The use of RAID is transparent to Oracle. All the features specific to RAID configurations are handled by the operating system and not by Oracle.

The use of RAID devices differs according to the Oracle file type. Data files and archive logs can be placed on RAID devices, since they are accessed randomly. Redo logs should not be put on RAID devices, since they are accessed sequentially and performance is enhanced by having the disk drive head near the last write location. However, mirroring of redo log files is strongly recommended by Oracle.

RAID is much easier to use than the Oracle techniques for data placement and striping.

Note the following RAID configuration recommendations:

- RAID usually impacts write operations more than read operations. This is especially true where parity needs to be calculated (RAID 3, RAID 5, and so forth).
- You can place online or archived redo log files on RAID 1 devices. Do not use RAID 5. In addition, place TEMP tablespace data files on RAID 1 devices, instead of RAID 5, because the streamed write performance of distributed parity (RAID 5) is not as good as that of simple mirroring (RAID 1).
- Swap space can be used on RAID devices without affecting Oracle.

The following table describes the RAID devices and types to be used with each Oracle file type:

RAID	Type of RAID	Control File	Database File	Redo Log File/ Temporary	Archive File
0	Striping	Avoid	OK	Avoid	Avoid
1	Shadowing	Recommended	OK	Recommended	Recommended
1+0	Striping and Shadowing	OK	Recommended	Avoid	Avoid
3	Striping with static parity	OK	OK	Avoid	Avoid

RAID	Type of RAID	Control File	Database File	Redo Log File/ Temporary	Archive File
5	Striping with rotating parity	OK	Recommended if RAID 1+0 is not available.	Avoid	Avoid

**Note:**

- RAID 0 does not provide protection against failures. It requires a strong backup strategy.
- RAID 1+0 is recommended for database files because it avoids hot spots and provides the best possible performance during a disk failure. The disadvantage of RAID 1+0 is its costly configuration.
- Use the highest RPM disks for temporary/redologs. Use as many controllers as you can in the array, and ensure that you place the redo log groups on different controllers.

## ASM Storage

In addition to the traditional file system option, Oracle Automatic Storage Management (ASM) can be used as a storage option for Oracle database files.

For details on configuring and administering ASM, see the Oracle Database Storage Administrator's Guide.

For an ASM Overview, best practices, and technical papers, see:

<http://www.oracle.com/technology/products/database/asm/index.html>

Pay attention to the following tasks:

- Preparing and mapping of the ASM disks, configuring the ASM disk groups, and setting the correct redundancy.
- Usage of the Oracle ASMLIB driver for improved discovery and management of ASM disks and for enhanced I/O processing (Linux only).
- Proper setting of ASM\_DISKSTRING according to your disks discovery path.
- Rebalance operation manual and automatic control options.

# Chapter 11: Maintaining an Oracle Server Database

This chapter describes the various maintenance and tuning procedures that are recommended for OMI databases created on Oracle Servers, as well as some of the available database backup and recovery methods.

This chapter includes:

- ["Database Maintenance and Tuning" below](#)
- ["Oracle Database Backup and Recovery" on page 67](#)

## Database Maintenance and Tuning

Poor database performance can be caused by the faulty configuration of the instance and database, or by abnormal resource consumption of an Oracle transaction, user, or process. It is essential for the database administrator to proactively monitor resource consumption, and correct any abnormalities before performance is affected.

**Note:** Memory, CPU, and I/O are the three most common system resources consumed by Oracle.

There are a number of third-party tools that you can use to monitor database behavior and assist you in identifying bottlenecks in your system. Use the following guidelines to help you.

This section includes the following topics:

- ["System Global Area \(SGA\)" below](#)
- ["Database Load Behavior" on the next page](#)
- ["CPU and I/O" on the next page](#)
- ["Oracle Alert File" on the next page](#)
- ["Archive Log - File System" on the next page](#)
- ["Tablespace Storage Space" on page 64](#)
- ["Space Management" on page 64](#)
- ["Collecting Statistics for Databases" on page 64](#)
- ["Oracle Automated Statistics Collection" on page 64](#)
- ["Collection Statistics for RTSM" on page 65](#)
- ["Index Fragmentation" on page 65](#)

## System Global Area (SGA)

Always configure your SGA to fit physical memory and avoid using swapping. It is recommended that you not set the SGA for more than 70 percent of system physical memory, leaving enough memory for

additional system and client processes. The parameters `MEMORY_TARGET` or `MEMORY_MAX_TARGET` cannot be greater than the shared memory file system (`/dev/shm`) on your operating system.

## Database Load Behavior

Oracle AWR (Automatic Workload Repository) reports can be used to monitor performance bottlenecks and database behavior. For details, see Oracle Metalink Note 276103.1: PERFORMANCE TUNING USING ADVISORS AND MANAGEABILITY FEATURES.

It is also recommended that you monitor I/O load on the system to identify I/O contention. Once you determine which disk is most loaded, you can use the AWR output to determine which particular Oracle data file is the cause of the contention and consider changing the I/O storage configuration.

## CPU and I/O

It is recommended that you monitor the CPU and file system, which are the main resources consumed by the database server. CPU usage should not exceed 70 percent and the I/O wait should not be higher than 10 percent.

To monitor the above resources, you can use **perfmon** on Windows or **top** on Linux, and the storage system management tools.

## Oracle Alert File

Oracle registers abnormal events in the **alert.log** file, whose location is defined by the `BACKGROUND_DUMP_DEST` parameter. As of Oracle 11g, the `BACKGROUND_DUMP_DEST` parameter is replaced by `DIAGNOSTIC_DEST` as part of Oracle Automatic Diagnostic Repository (ADR).

It is recommended that you check this file regularly to identify abnormalities that should be corrected, for example, ORA-XXXXX errors.

## Archive Log - File System

When using the archivelog mode, monitor your `ARCHIVE_DUMP_DEST` location for disk usage. These files should be backed up and deleted regularly to leave sufficient disk space for new archive files.

The archive file is usually the same size as the redo log file. To determine the size of a redo log file, use the operating system command or the following query:

```
SQL> select GROUP#, BYTES
        from V$LOG;
```

To determine the number of archive files generated over a period of time, for example, a day, you can use the following query after the system is stable:

```
SQL> alter session set NLS_DATE_FORMAT = 'DD-MON-YYYY';
SQL> select FIRST_TIME as "Day",
        COUNT(*) as "Number of files"
        from V$LOG_HISTORY
        group by FIRST_TIME
        order by 1 asc;
```

## Tablespace Storage Space

To avoid space errors caused by data growth, monitor your tablespace usage regularly.

If you run out of space in one of your tablespaces, you can add one or more data files to it by using the `ALTER TABLESPACE <tablespace name> ADD DATAFILE...` command.

## Space Management

Free space in Oracle tablespaces is composed of newly created extents, or space that was used and freed by operations such as update and delete. If some of the free space in a tablespace is composed of extents that were used and freed, your tablespace may become fragmented. You can use the Oracle Segment Advisor to determine whether objects have unused space that can be released. Objects that are found to be fragmented can be reorganized to reclaim wasted space and compact the segment. For an overview on reclaiming wasted space, see the Oracle Database Administrator's Guide for your Oracle release.

## Collecting Statistics for Databases

The OMi platform is planned and built to work with the Oracle Cost Base Optimizer. For the Optimizer to work properly, you must periodically collect statistics for all schema tables.

During the initial phase of OMi deployment, it is recommended that you collect statistics for all OMi objects (tables and indexes).

When working with large OMi environments, it is recommended that you collect statistics only for objects for which the amount of data changes significantly during the day, or for new objects that are created (such as the tables and indexes created by the Partitioning and Purging Manager). Once your OMi system is stable, you should collect statistics once a day.

## Oracle Automated Statistics Collection

The Automatic Optimizer Statistics Collection job, `GATHER_STATS_JOB`, is managed by the Automatic Maintenance Tasks framework instead of the Oracle Scheduler framework. Automatic Maintenance Tasks are predefined tasks that perform maintenance operations on the database. These tasks run in a maintenance window, which is a predefined time interval that is intended to occur during a period of low system load. You can manually customize maintenance windows based on the resource usage patterns of your database, or disable certain default windows from running.

According to the Oracle documentation, the preferred way to manage or change configuration of the Automatic Optimizer Statistics Collection job is to use the Automatic Maintenance Tasks screens in Enterprise Manager Database Control and Grid Control. For an overview of the Automatic Maintenance Tasks, see the Oracle Database Administrator's Guide.

### **To manually collect statistics for all OMi objects on a database schema:**

1. Log in to the OMi relevant schema using SQL\*Plus (a third party Oracle client).
2. Run the following command:



```
Exec DBMS_STATS.GATHER_SCHEMA_STATS (ownname => '<name of Oracle schema>',  
options => 'GATHER AUTO');
```

#### To collect statistics for specific schema tables and their indexes:

1. Log in to the schema using SQL\*Plus.
2. For each table, collect statistics by running the following command:

```
Exec DBMS_STATS.GATHER_TABLE_STATS (ownname => '<name of Oracle schema>',  
tabname => '<Name of table for which you want to collect statistics>',  
estimate_percent => 5, cascade => TRUE);
```

#### Note:

- Cascade => True instructs the Oracle database to analyze all the indexes in the table.
- Collecting statistics is a resource-consuming operation that can take a long time. It is therefore recommended that you collect statistics during special maintenance hours.

## Collection Statistics for RTSM

Unlike some databases, where queries are predefined and can be tuned according to the expected database size, RTSM database constructs queries dynamically, according to Pattern Views defined against its data model. This necessitates accurate statistics at all times. In addition to running a daily job to update statistics for RTSM, it is recommended that you manually refresh statistics if major changes to the RTSM schema objects have occurred, such as bulk insert transactions triggered by automated DFM jobs. DFM is the process responsible for automatically detecting configuration items (CIs) and inserting them into RTSM.

## Index Fragmentation

OMi schemas consist of Oracle B-tree indexes for enhancing searches on table columns.

It is recommended to validate the structure of the schema indexes periodically (at least every week for active systems), and if necessary to rebuild the indexes found to be fragmented.

The main reasons for indexes becoming fragmented are:

- **Row deletes.** When rows in a table are deleted, Oracle index nodes are not physically deleted, nor are the entries removed from the index. Rather, Oracle logically deletes the index entry and leaves dead nodes in the index tree, where they may be reused if another adjacent entry is required. However, when large numbers of adjacent rows are deleted, it is highly unlikely that Oracle is able to reuse the deleted leaf rows. In addition to wasting space, large volumes of deleted leaf nodes cause index scans to take more time.

Over time, following row deletes from schema tables, there may be a need to rebuild some of the schema indexes.

- **Index height.** The height of an index refers to the maximum number of levels encountered within the index. As the number of levels in an index increases, more block reads are needed when searching the index. When a large amount of rows are added to a table, Oracle may create additional levels of an index to accommodate the new rows, thereby causing an index to reach four levels, although only in those areas of the index tree where massive inserts have occurred. While Oracle indexes can

support many millions of entries in three levels, any Oracle index that has four or more levels can benefit from rebuilding.

For OMi tables, it is recommended to rebuild any index that has more than three levels.

## Index Maintenance Utility

OMi's index maintenance utility (**maintain\_indexes.bat**) can be used to identify and rebuild indexes that have more than three levels, or that have 100,000 values or more with 10% of deleted values.

You can set a flag when running the utility to instruct it to rebuild indexes identified as being fragmented automatically, although it is recommended that you rebuild indexes manually.

When run, the utility produces a log file (**index\_stats.log**) that contains the following entries:

- An alphabetical list of indexes that were identified as candidates for rebuilding. For each index listed, statistics are shown such as the height of the index and the percentage of deleted rows.
- Rebuild commands for each index listed that can be used to rebuild the indexes manually.

The utility also creates a table called **TEMP\_STATS** in the target schema that contains all the indexes and their related statistics (not only the indexes listed as candidates for rebuilding). The table remains in the schema until it is manually dropped, to enable inspection of the results at a later stage.

**Warning:** The index maintenance utility is resource intensive, as it analyzes all indexes in the schema. It can also cause locks on database objects, or skip indexes that are locked by other sessions. It is recommended to run the index maintenance utility during maintenance hours only.

### To run the index maintenance utility:

1. Copy the following files from the **\<Gateway Server root directory>\AppServer\webapps\site.war\DataBases\ORA\_DB\_Utills** directory to a Windows machine that has Oracle database client installed:
  - **maintain\_indexes.bat**
  - **maintain\_indexes.sql**
2. On the machine to which you copied the files, open a DOS command window and move to the location in which you copied the files.
3. Run the index maintenance utility with the following command:

```
maintain_indexes.bat <schema> <password><db alias> (rebuild flag)
```

where:

- **schema.** The name of the database schema user of the schema for which you are running the utility.
- **password.** The password for the database schema user.
- **db alias.** The db alias for connecting to the target database as specified in the **tnsnames.ora** file. Ensure that there is an entry in the **tnsnames.ora** file for the target server.
- **rebuild flag.** The flag to instruct the utility to rebuild indexes automatically. Set the flag to 0 if

you don't want the utility to rebuild indexes automatically and to 1 if you do. The default setting is 0.

When the index maintenance utility has finished running, check the `index_stats.log` file in the directory into which you copied the files in step 1, for the list of indexes that are candidates for rebuilding and the rebuild commands to be used.

**Note:** The execution time of the index maintenance utility depends on the size of the indexes and the load on the system when being run.

## Oracle Database Backup and Recovery

Your backup strategy is tested when a failure occurs and data is lost. You can lose or corrupt data in several ways, such as a logical application error, an instance failure that prevents Oracle from starting, or a media failure caused by a disk crash. In addition to your scheduled backups, it is important to perform a backup when the database structure changes (for example, when a data file is added to the database), or before you upgrade your software or hardware.

When choosing a backup strategy, consider several factors, such as the system workload, the usage schedule, the importance of the data, and the hardware environment of the database.

Oracle backups can be performed using scripts executing SQL commands combined with operating system commands to copy files, or using Oracle RMAN (Recovery Manager) commands.

It is recommended that you maintain updated records of backups performed on your database so that you can use them for recovery on demand. If you are using RMAN, catalog information is available from the catalog.

This section describes:

- ["Available Backup Methods" below](#)
- ["Oracle Recovery Manager - RMAN" on the next page](#)

### Available Backup Methods

This section describes the various backup methods that are available.

#### Cold Backup

Cold backup, also known as offline backup, is a database level backup. It normally requires that the database be shut down before the backup is started. The amount of downtime is dependent on the database size, the backup media (disk or tape), the backup software, and the hardware in use.

Once the instance is down, all its data files, log files, control files, and configuration files should be copied either to disk or other media. After the copy is complete, the instance can be restarted.

This backup method enables recovery to a the point in time in the past at which the database snapshot was taken.

For more information, see the Oracle Backup and Recovery Guide for your Oracle release.

## Hot Backup

Hot backup, also known as online backup, enables you to run a backup while the instance is running and users are connected to the database. This backup method is a tablespace backup level and requires the database to operate in archivelog mode, which enables Oracle to track changes over time by generating redo log file copies called archive files. The generated archive files are written to the archive destination specified by the LOG\_ARCHIVE\_DEST (or LOG\_ARCHIVE\_DEST\_NN) parameter in the instance parameter files.

**Note:** If the Oracle Flash Recovery Area is used, then the default destination for archiving is the value set in the `db_recovery_file_dest` parameter.

Once you start the backup, all the data files, control files, archive files, and configuration files should be copied either to disk or other media. This method enables recovery to any point in time. Note that working in archivelog mode requires additional disk space to contain incremental archive files, which can influence database performance. During the backup process, OMi may also experience some performance degradation due to disk load.

For more information, see the Oracle Backup and Recovery Guide for your Oracle release.

## Data Pump Export Import

In addition to the cold and hot physical backup methods, you can use the backup method known as Data Pump. Data Pump APIs for moving data and metadata between databases and schemas replaced Oracle's original export/import utilities.

The Oracle Data Pump export utility dumps schema structure and contents into one or more Oracle binary files called a dump file. This method can be used to transfer data between two schemas in the same database, or between two separate Oracle databases. To load exported data back into the database, use the Data Pump import utility. For more information on Oracle Data Pump utilities, see Oracle Database Utilities in the Oracle documentation for your release.

**Note:** OMi does not require you to use a specific backup method; however, it is recommended that your backup method accommodate OMi's use of more than one database user schema.

## Oracle Recovery Manager - RMAN

Recovery Manager (RMAN) is a generic Oracle tool that enables you to back up and restore your target database. When working with RMAN, you can choose to work with the RMAN catalog schema. The catalog is managed within the Oracle schema and stores information on the registered database structure and backups performed using RMAN. It can be queried to produce backup reports and copy availability. A single catalog can manage backup information from one or more target databases.

The RMAN catalog is usually placed on a different database instance than the operational database and has a backup strategy of its own. It need only be available during the backup or recovery process.

The RMAN tool can be used in conjunction with third-party backup software for a complete backup and recovery solution.

The following are some advantages of RMAN:

- Minimizes backed up data by compressing backed up files to exclude empty data blocks, thereby saving time and space.
- Supports incremental backups.
- Supplies the user with backup status reporting ability.
- Supports parallel backup and recovery processes when possible.
- Can be used with a third-party backup media tool.

For more information on RMAN, see the Oracle Recovery Manager User's Guide for your release.

# Chapter 12: Configuring the Oracle Client for OMi

This chapter describes how to configure the Oracle Client for OMi database offline utilities. Installing an Oracle client on the application server is not a requirement. However, it is recommended to have one to operate the various offline database scripts such as manually creating users or collecting statistics against OMi databases.

## Oracle Client Versions and Operating System Platforms

The Oracle Client versions and operating system platforms that are supported and recommended for working with OMi are identical to the Oracle Server versions and platforms. For details, see the Oracle Server requirements table in ["Software Requirements" on page 47](#).

## Oracle Client Installation

To install the Oracle Client, see the Oracle documentation.

If you choose the custom installation option during the installation process, ensure that you install the following components (under **Oracle Client**):

- Oracle Net (Including TCP/IP Adaptor)
- Oracle Database Utilities
- SQL\*Plus
- Oracle Call Interface (OCI)

## Oracle Client Configuration

In order to work with OMi, you must configure the **tnsnames.ora** file that is located in the **<ORACLE\_HOME>\network\admin** directory. Make sure to specify the name or IP of the Oracle Server host machine, the Oracle Server listener port (by default, usually 1521), and the SID (by default, ORCL) or service\_name. The following is an example of a **tnsnames.ora** file.

```
# TNSNAMES.ORA Network Configuration File: D:\oracle\ora81\network\admin\tnsnames.ora
# Generated by oracle configuration tools.

MY_COMPANY.CO =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = london)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SID = london)
    )
  )
```

It is recommended that you use the Oracle Net Configuration Assistant Oracle tool built to configure the **tnsnames.ora** file. For more information, see the Oracle documentation.

Ensure that the Oracle Client configuration, such as SID and port settings, matches the Oracle Server configuration. To test the connection from the Oracle Client machine to the Oracle Server machine, use the **tnsping** utility.

If you are using a RAC deployment, you can configure your tnsnames entry to work with the RAC. For an example of a RAC entry, see "[Support for Oracle Real Application Cluster](#)" on page 100.

**Note:**

- The OMi servers access the Oracle Server using the JDBC thin driver. The JDBC thin driver does not support a firewall connection that is net\*8/9 compliant, and therefore allows for SQL data transmission only.
- The Easy Connect Naming method eliminates the need to look up names in the tnsnames.ora files for TCP/IP environments. For details, see the Oracle Database Net Services Administrator's Guide in the Oracle documentation set of your release.

## Chapter 13: Oracle Summary Checklist

This chapter contains a checklist summarizing the requirements for OMi support and certification.

For more information on the Oracle database configuration settings that should be used when working with Oracle Server and OMi, see ["Oracle Server Configuration and Sizing Guidelines" on page 53](#).

### Checklist for OMi Support and Certification

The information provided in this section is for both supported and certified Oracle options. The certified options are recommended for working with OMi.

Option	Supported	Recommended	Remarks	For more information, see
Oracle edition	Enterprise	Enterprise		
Dedicated OMi server	Not necessary	Not necessary It is recommended to dedicate an instance for OMi.		
Use of multiple Oracle instances	Yes	No	The configuration of all instances must match in a certified environment.	<a href="#">"Oracle Instances" on page 47</a>
Use of non-default port	Yes	Yes		
Undo management	Automatic; Manual	Automatic	Set UNDO_MANAGEMENT parameter to AUTO in a certified environment.	
Oracle Shared Servers connection method	Yes	No	OMi uses a connection pool architecture. Use the dedicated server connection method in a certified environment.	
Oracle replication	No full support	No		



Option	Supported	Recommended	Remarks	For more information, see
Operating System file compression	No	No	Not supported by Oracle; Causes abnormal behavior and affects performance.	
Database control files required	Greater than or equal to 2	3	Preferable on different disks.	
Redo log groups	Greater than or equal to 3	4	Oracle enables the software mirroring of redo log files. This is achieved by creating at least two members of redo log in each group. Members of the same group should reside on different disks.	
Character set	UTF8; AL32UTF8	AL32UTF8		
OPEN_CURSORS	2000	2000		
Working in archive log mode	True; False	True		
Autoextend option in tablespace files	Yes	No		
Locally managed data tablespace	Yes	Yes		<a href="#">"Locally Managed Tablespaces" on page 56</a>
Tablespace extent management	Local automatic for application schemas; Local uniform for TEMP tablespace	Local automatic for application data tablespaces; Local uniform for TEMP tablespace		<a href="#">"Oracle Tablespace Settings" on page 57</a>

Option	Supported	Recommended	Remarks	For more information, see
Automatic Segment Space Management Tablespace (ASSM)	Yes	Yes		

### Oracle Client Configuration

In order to work with OMi, you must configure the **tnsnames.ora** file that is located in the **<ORACLE\_HOME>\network\admin** directory. Make sure to specify the name or IP of the Oracle Server host machine, the Oracle Server listener port (by default, usually 1521), and the SID (by default, ORCL) or service\_name. The following is an example of a **tnsnames.ora** file.

```
# TNSNAMES.ORA Network Configuration File: D:\oracle\ora81\network\admin\tnsnames.ora
# Generated by oracle configuration tools.

MY.COMPANY.CO =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = london)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SID = london)
    )
  )
```

It is recommended that you use the Oracle Net Configuration Assistant Oracle tool built to configure the **tnsnames.ora** file. For more information, see the Oracle documentation.

Ensure that the Oracle Client configuration, such as SID and port settings, matches the Oracle Server configuration. To test the connection from the Oracle Client machine to the Oracle Server machine, use the **tnsping** utility.

If you are using a RAC deployment, you can configure your tnsnames entry to work with the RAC. For an example of a RAC entry, see ["Support for Oracle Real Application Cluster" on page 100](#).

**Note:**

- The OMi servers access the Oracle Server using the JDBC thin driver. The JDBC thin driver does not support a firewall connection that is net\*8/9 compliant, and therefore allows for SQL data transmission only.
- The Easy Connect Naming method eliminates the need to look up names in the tnsnames.ora files for TCP/IP environments. For details, see the Oracle Database Net Services Administrator’s Guide in the Oracle documentation set of your release.

### Oracle Server and Client Requirements

For the Oracle Server and Client versions and operating system platforms that are supported for working with OMi, see ["Software Requirements" on page 47](#).

## Setting Oracle Initialization Parameters

For a list of recommended settings for Oracle server initialization parameters, see "[Oracle Parameter Settings](#)" on page 53.

# Part IV: Deploying and Maintaining the PostgreSQL Server Databases

# Chapter 14: PostgreSQL Deployment Overview

You can use PostgreSQL for deploying OMi databases. This chapter describes the following topics related to deploying PostgreSQL for use with OMi:

- ["About PostgreSQL Deployment" below](#)
- ["System Requirements" on the next page](#)
- ["Character Set Support in PostgreSQL" on page 79](#)

When configuring OMi, you can choose between an external PostgreSQL database (**Postgres**) and an embedded PostgreSQL database (**PostgresEmbedded**). You can always migrate from your embedded database to an external one. For details, see ["Migrating from an Embedded PostgreSQL Database to an External PostgreSQL Database" on page 88](#).

If you use the embedded PostgreSQL database, you can set up only one Data Processing Server. No additional Data Processing Servers can be added to an OMi deployment.

**Note:** The location of the embedded PostgreSQL database files is as follows:

- **Windows:** %OvDataDir%\databases\omidb\
- **Linux:** /var/opt/OV/databases/omidb/

## About PostgreSQL Deployment

To deploy PostgreSQL for use with OMi, perform the following procedures:

- **Install and Configure PostgreSQL.**

For details on PostgreSQL software installation, see the installation guide in the documentation for your specific PostgreSQL version.

**Caution:** Make sure you configure the `pg_hba.conf` file on the PostgreSQL server to accept remote connections. Otherwise, the connection to the PostgreSQL database cannot be established when running the OMi configuration wizard.

- **Create Databases on the PostgreSQL Server.**

You can create OMi databases manually, or you can use the configuration wizard to create the databases for you.

The configuration wizard prompts you to supply the following parameters:

- **Host name.** The name of the host machine on which PostgreSQL is installed.
- **Port.** The PostgreSQL listening port. OMi automatically displays the default port for the external PostgreSQL database, 5432.

**Note:** If you use the embedded PostgreSQL database, OMi uses port 5433 to connect to the database.

- **User name and password.** The name and password of a user that OMi uses to connect to and retrieve data from the OMi databases.
- **Use SSL.** If selected, OMi will use an encrypted connection to communicate with the database server.

**Note:** If you use SSL, make sure that the PostgreSQL server is started with SSL enabled and a client certificate for use by OMi is provided. For details on configuring PostgreSQL for SSL, see the PostgreSQL documentation.

- **Admin user name and password.** (*Only when creating a new database.*) The name and password of a user with administrative permissions on the PostgreSQL server.
- **Database names (Management, RTSM, and Event).** The names of the databases (for example, OMi\_Management).

**Note:** If you use the embedded PostgreSQL database, you are asked to provide only the password of a user that OMi uses to connect and retrieve data from the OMi databases.

For details on creating the databases for OMi manually, see "[Creating the OMi Databases Manually](#)" on [page 80](#). For more information on the configuration wizard, see the OMi Installation and Upgrade Guide.

## System Requirements

This section describes the system requirements for working with PostgreSQL in conjunction with OMi.

This section includes the following topics:

- "[Hardware Requirements](#)" below
- "[Software Requirements](#)" below

## Hardware Requirements

For OMi hardware sizing guidelines, see "[Requirements](#)" on [page 11](#).

For PostgreSQL hardware requirements, see the PostgreSQL documentation available at:

<http://www.postgresql.org/docs/manuals/>

## Software Requirements

For a list of supported PostgreSQL database versions, see the OMi Support Matrix at:

<https://softwaresupport.hp.com/group/softwaresupport/search-result/-/facetsearch/document/KM323488>

## Character Set Support in PostgreSQL

When you create a new PostgreSQL cluster in a PostgreSQL database, you can specify the character set during the cluster creation by using the `-E` parameter with the `initdb` command:

```
initdb -E UTF8
```

To create a database inside a cluster with a different locale, use one of the following commands:

- From a terminal: `createdb -E UTF8 <dbname>`
- In SQL: `CREATE DATABASE <dbname> WITH ENCODING 'UTF8';`

**Note:** To check the encoding of the database server, run the `psql -l` command or type `\l` inside the `psql` terminal.

When you use the PostgreSQL client application or any other client application to manually connect to the database, it is recommended that you specify the correct character set with the following SQL command:

```
SET NAMES "UTF8";
```

# Chapter 15: Creating the OMi Databases Manually

This chapter describes the procedures for creating OMi databases on a PostgreSQL server manually.

This chapter includes:

- ["PostgreSQL Database Overview" below](#)
- ["Creating OMi Databases Manually" below](#)

## PostgreSQL Database Overview

OMi uses the Management, RTSM, and Event databases.

During the OMi configuration, new databases can be created automatically by the configuration wizard, or already existing databases can be used. Existing databases can either be created manually in advance (for example, due to organization security restrictions), or can be created by a previous installation of the same release of OMi.

For details about installing the OMi server, see the OMi Installation and Upgrade Guide.

## Creating OMi Databases Manually

This section describes how to create each of the OMi databases manually:

- ["Creating a Management Database Manually" below](#)
- ["Creating an RTSM Database Manually" on the next page](#)
- ["Creating an Event Database Manually" on the next page](#)

**Note:** It is recommended that a certified database administrator creates the required database in compliance with the organization's database strategy.

## Creating a Management Database Manually

To create a Management database manually, follow these steps:

1. Create the Management database for storing Management objects.
2. Register database extensions by running the following command:

```
<OMi_HOME>/AppServer/webapps/site.war/DataBases/Postgres_DB_Utils/postgres_database_extensions.sql
```

**Note:** You must perform this step because the Management database requires the `uuid-ossp` extension to be registered.



3. Create the Management objects as follows:
  - a. Generate the Management database SQL script by using the following script file:

- **For Windows:**

```
<OMi_HOME>\dbverify\bin\generate_create_script.bat
```

- **For Linux:**

```
/opt/HP/BSM/dbverify/bin/generate_create_script.sh
```

Use the following parameters for the Management database:

- **deployment\_type.** EE (Enterprise Edition)
- **Generate Script Name.** The name of the SQL file to generate, for example, `management.sql`.
- **dbType.** Postgres
- **sessionType.** Management

**Example:**

```
/opt/HP/BSM/dbverify/bin/generate_create_script.sh EE management.sql  
Postgres Management
```

- b. Connect to the Management database, and then run the generated script to deploy the Management objects.
- c. Run the following SQL script located in `<OMi_HOME>/AppServer/webapps/site.war/DataBases/Postgres_DB_Utils/management_postgres_inserts.sql`

## Creating an RTSM Database Manually

Create a database for storing RTSM objects. Unlike the Management and Event databases, the RTSM database does not require running any scripts for creating the objects to populate it.

## Creating an Event Database Manually

To create an Event database manually, follow these steps:

1. Create the Event database for storing Event objects.
2. Register database extensions by running the following command:

```
<OMi_HOME>/AppServer/webapps/site.war/DataBases/Postgres_DB_Utils/postgres_  
database_extensions.sql
```

**Note:** You must perform this step because the Event database requires the `uuid-osspl` extension to be registered.

3. Create the Event objects as follows:

- a. Generate the Event database SQL script by using the following script file:

- **For Windows:**

```
<OMi_HOME>\opr\bin\opr-schema-script-generator.bat
```

- **For Linux:**

```
<OMi_HOME>/opr/bin/opr-schema-script-generator.sh
```

**Note:** The database type parameter, `-d`, should be set to **Postgres**.

For example, on Windows, to generate a script called `event.sql` that can be used to create the necessary tables on the PostgreSQL database, run the following command:

```
<OMi_HOME>\opr\bin\opr-schema-script-generator.bat -d Postgres -o  
C:\temp\event.sql
```

- b. Connect to the Event database, and then run the generated script to deploy the Event objects.

# Chapter 16: Maintaining a PostgreSQL Database

This chapter describes the maintenance procedures that are recommended for OMi databases created on PostgreSQL servers.

This chapter includes:

- ["Database Maintenance" below](#)
- ["PostgreSQL Database Backup and Restore" below](#)
- ["Resetting or Changing the User Password for the Embedded PostgreSQL Database" on page 86](#)

## Database Maintenance

Achieving optimum database performance requires that certain maintenance tasks are performed regularly. One of the required PostgreSQL database maintenance tasks is called **vacuuming** and it is performed by the autovacuum daemon, a built-in PostgreSQL tool that is enabled in OMi and executes the following recurring maintenance tasks:

- Recovering disk space by running VACUUM commands
- Updating planner statistics

You can also perform database vacuuming manually. This is especially useful after running a large amount of INSERT, UPDATE, or DELETE statements. For this purpose, use the built-in PostgreSQL tool, `vacuumdb`, which you can find at the following location:

```
<OMi_HOME>/pgsql/bin/
```

**Note:** The built-in tool may not be fully compatible with the version of the external PostgreSQL database server. Therefore, if you use an external PostgreSQL database, make sure to use the `vacuumdb` tool from the database server instead.

Although indexes are maintained regularly by PostgreSQL, sometimes a manual rebuild of an index can be done to improve query performance (for example, after running a large amount of INSERT, UPDATE, or DELETE statements). To do so, use the built-in tool for rebuilding indexes in a PostgreSQL database, `reindexdb`.

**Note:** Rebuilding indexes by using `reindexdb` may block table access. Therefore, run this tool only when the database server is under low load.

## PostgreSQL Database Backup and Restore

A well-defined backup and restore strategy is crucial for every database because you can lose or corrupt data due to various reasons. Whatever might be the cause for data loss or corruption, this is when your backup strategy, the key to successfully recovering your data, is put to the test. In addition to your scheduled backups, it is important to perform a backup whenever the database structure changes or before you upgrade your software or hardware.

Depending on whether you use an embedded or external PostgreSQL database, see one of the following sections about how to back up and restore your data:

- ["Backing up and Restoring an Embedded PostgreSQL Database" below](#)
- ["Backing up and Restoring an External PostgreSQL Database" on the next page](#)

## Backing up and Restoring an Embedded PostgreSQL Database

OMi provides you with the `opr-pgctl` tool that you can use to manage, back up, and restore your embedded PostgreSQL database. You can find this tool at the following location:

- **Windows:** `<OMi_HOME>\bin\opr-pgctl.bat`
- **Linux:** `<OMi_HOME>/bin/opr-pgctl.sh`

The `opr-pgctl` tool is available only on the Data Processing Server.

### Embedded PostgreSQL Database Backup

To perform a full backup of an embedded PostgreSQL database, run the following command:

```
opr-pgctl -backup <destination> [-z]
```

In this instance, `<destination>` is the name of the backup output file and you must make sure that it does not exist already. If you use the `-z` parameter, the output file will be compressed using `gzip`.

The file that is created by the backup operation includes the information for recreating all databases, user roles and groups, access permissions, and tablespaces of the embedded PostgreSQL database cluster. It does not, however, include the information about the configuration file (that is, `postgresql.conf`). If you want to back up the complete cluster data directory, including all the configuration information of the database, you can use the continuous archiving backup method. For details, see the PostgreSQL documentation available at:

<http://www.postgresql.org/docs/manuals/>

**Note:** Backing up the embedded PostgreSQL database requires no OMi downtime.

### Embedded PostgreSQL Database Restore

To restore an embedded PostgreSQL database using the backup file that you created during the backup operation, run the following command:

```
opr-pgctl -restore <source> [-z]
```

In this instance, `<source>` is the name of the backup file. Make sure that you provide the `-z` parameter if you compressed the backup file using `gzip`.

Restoring the embedded PostgreSQL database is performed in the following order:

- Existing cluster data directory of the embedded PostgreSQL database is removed.
- Cluster data directory is reinitialized using the credentials entered in the configuration wizard.
- Databases are recreated using the specified backup file.

**Caution:** You must stop OMi before restoring the embedded PostgreSQL database.

## Backing up and Restoring an External PostgreSQL Database

When backing up an external PostgreSQL database, you can choose among the following backup methods:

- SQL dump
- File system level
- Continuous archiving

This section includes a brief description of all three methods. For more details as well as for advantages and disadvantages of each of these backup methods, see the PostgreSQL documentation.

### SQL Dump

With the SQL dump method, a file containing the necessary SQL commands to recreate the database is generated. You can choose between the following utility programs to generate the backup file:

- `pg_dump`  
The backup output file contains the information about tables of a single database. It does not, however, include the information about all databases, user roles and groups, access permissions, and tablespaces, which means that you must recreate this information manually before you restore your database.
- `pg_dumpall`  
The backup output file contains all the information inside a database cluster, including the information about all databases, user roles and groups, access permissions, and tablespaces.

**Note:** The configuration files are not backed up with the SQL dump backup method. Therefore, you must recreate them or back them up manually.

### File System Level

With the file system level backup method, you directly copy the files that your PostgreSQL database uses for storing the data. For this purpose, you can use tools such as `copy` or `tar`.

**Caution:** Make sure that you shut down the database server before copying the files used for storing the data.

### Continuous Archiving

This method involves the continuous archiving of the write ahead logs (WALs) that are used to record all changes made to the data files inside a PostgreSQL database. By default, these log files are periodically overwritten with new information. However, you can configure PostgreSQL to keep copies of the old WALs, so that it is possible to restore a database to a specific point in time by replaying the log entries.

## Resetting or Changing the User Password for the Embedded PostgreSQL Database

To reset a lost user password or to change a user password for the embedded PostgreSQL database, on the Data Processing Server, perform the following steps:

1. Stop OMi.
2. Temporarily disable the password requirement for the administrative PostgreSQL user `postgres`. To do so, edit the `pg_hba.conf` file by adding the following line or lines at the beginning of the file:

- **Windows:**

```
host all postgres ::1/128 trust
host all postgres 127.0.0.1/32 trust
```

- **Linux:**

```
local all postgres trust
```

**Note:** You can find the `pg_hba.conf` file at the following location:

- **Windows:** `%OvDataDir%\databases\omidb\`
- **Linux:** `/var/opt/OV/databases/omidb/`

3. Start the embedded PostgreSQL database by running the following command:

```
<OMi_HOME>/bin/opr-pgctl -start
```

4. Connect to the database by running the following command:

- **Windows:**

```
<OMi_HOME>\pgsql\bin\psql.exe -U postgres -p 5433
```

- **Linux:**

```
<OMi_HOME>/pgsql/bin/psql -U postgres -p 5433
```

**Note:** You are not prompted to provide a password because the password requirement was temporarily disabled for the `postgres` user.

5. Change the password of the `postgres` and `hpbsm` database users (both users must share the same password) by running the following SQL commands in an open PostgreSQL client session:

```
\password postgres
```

```
\password hpbsm
```

You must enter a new password twice for each command.

6. Terminate the PostgreSQL client session by entering `\q`.
7. Stop the embedded PostgreSQL database by running the following command:

```
<OMi_HOME>/bin/opr-pgctl -stop
```

8. Enable the password requirement for the postgres user by removing the following from the pg\_hba.conf file:
  - **Windows:**  
host all postgres ::1/128 trust  
host all postgres 127.0.0.1/32 trust
  - **Linux:**  
local all postgres trust

After resetting the lost user password or changing the user password, you must reconfigure OMi to use the new password. To do so, on all Data Processing Servers and Gateway Servers, perform the following steps:

1. Run the OMi configuration wizard. For information on the configuration wizard, see the OMi Installation and Upgrade Guide.
2. Select **Connect to an existing database or user schema**.
3. Choose the **Embedded Postgres** database type.
4. Enter your new password.

# Chapter 17: Migrating from an Embedded PostgreSQL Database to an External PostgreSQL Database

This chapter describes how to migrate your embedded PostgreSQL database to an external PostgreSQL database.

This chapter includes:

- ["Reasons for Migrating Your Embedded PostgreSQL Database" below](#)
- ["Procedure for Migrating from an Embedded PostgreSQL Database to an External PostgreSQL Database" below](#)

## Reasons for Migrating Your Embedded PostgreSQL Database

The embedded PostgreSQL database server supplied with OMi allows for a quick and simple configuration procedure and it is sufficient for small environments. However, as the managed environment grows in complexity, it may be necessary for you to move to a dedicated external PostgreSQL database server.

The external PostgreSQL database is more suitable for advanced environments because the administrator has increased control over the configuration of the database server so that the database adjusts better to the needs of the OMi environment. With the external database, it is also possible to set up replication and load balancing, which highly increases the performance and availability of the database server, and helps reduce downtime in case of failure.

**Note:** With the external PostgreSQL database, it is the database administrator's responsibility to keep the database up and running all the time, as well as to ensure that it is responsive and performs well.

## Procedure for Migrating from an Embedded PostgreSQL Database to an External PostgreSQL Database

To migrate from an embedded PostgreSQL database to an external PostgreSQL database, follow these steps:

1. Create a backup of the existing embedded PostgreSQL database as described in ["Backing up and Restoring an Embedded PostgreSQL Database" on page 84](#).



2. Recreate the OMi databases on the external PostgreSQL database cluster by running the following command:

```
psql -f <backup file> -h <db host> -p <db port> -U <admin user>
```

In this instance, *<backup file>* is the name of the backup file that you created during the backup procedure. If you used the *-z* parameter with the *opr-pgctl* command, make sure to unzip the backup file manually by using tools such as 7-Zip or *gzip*.

**Note:** The *psql* tool is available with OMi and you can find it in the *<OMi\_HOME>/pgsql/bin* directory. However, if you use an external PostgreSQL database, it is highly recommended to use the tool from the database server instead.

3. Reconfigure OMi to connect to the already existing PostgreSQL database.

The old OMi databases are removed and the following databases are created on the specified external PostgreSQL database cluster:

- *mgmt*
- *rtsm*
- *event*

# Part V: Appendixes

# Appendix A: Microsoft SQL Server Data Storage Recommendation

This appendix contains recommendations for creating Microsoft SQL Server file groups in which to store OMi data. This information applies to advanced users only.

This appendix includes:

["Creating Microsoft SQL Server File Groups" below](#)

## Creating Microsoft SQL Server File Groups

When OMi databases are created automatically by the configuration wizard, OMi stores database objects in two file groups: the **PRIMARY** file group for system tables and the **userdata001** file group for user tables. If you are setting up a large deployment of OMi, it is recommended that you create additional file groups in which to store OMi database objects.

### Benefits of Multiple File Groups

Placing the system catalog in its own file group and creating several user-defined file groups for data storage has the following benefits:

- **Maintenance.** Backing up and restoring databases can be performed on a file group level.
- **Partial restoration.** You can partially restore a database on the file group level; however when you do so, you must restore the **PRIMARY** file group as well, as it contains the system catalog. By keeping only the **.mdf** file in the **PRIMARY** file group and setting another user file group as the default file group, the **PRIMARY** file group remains small and does not pose difficulties for a partial restoration. This suggested file group structure creates a single user file group, which is a step toward multiple file groups that can be created by the database administrator or by OMi in the future.
- **Read-only.** A file group can be marked as read-only. Because data cannot be changed, Microsoft SQL Server does not have to acquire shared locks and performance is enhanced.
- **Recovery.** If the data portion of a database crashes, you can save the changes made before the crash even if the **.mdf** file is damaged, provided the log file and master database are undamaged.
- **Performance.** Numerous tuning options are available, such as hot tables on fast disks, the separation of table data and indexes into different file groups, and so forth.
- **Flexibility and Control.** The use of multiple file groups provides the database administrator with greater flexibility and control. In addition, the use of a fixed size for database growth prevents situations in which the 10 percent default autogrowth increment in large databases cannot finish before the timeout setting of OMi components and rolls back database activity, causing an infinite loop of inserts.

You can create an additional file group (together with an additional Microsoft SQL Server database) by running the following scripts.

**To create a database with default settings:**

```
CREATE DATABASE [testdb]
```

**To modify the .mdf (catalog) file SIZE parameter:**

```
ALTER DATABASE [testdb]
MODIFY FILE
(
    NAME = N'testdb',
    SIZE = 5MB
)
```

**To modify the .mdf (catalog) file FILEGROWTH parameter:**

```
ALTER DATABASE [testdb]
MODIFY FILE
(
    NAME = N'testdb',
    FILEGROWTH = 5MB
)
```

**To modify the .ldf (log) file SIZE parameter:**

```
ALTER DATABASE [testdb]
MODIFY FILE
(
    NAME = N'testdb_log',
    SIZE = 10MB
)
```

**To modify the .ldf (log) file FILEGROWTH parameter:**

```
ALTER DATABASE [testdb]
MODIFY FILE
(
    NAME = N'testdb_log',
    FILEGROWTH = 50MB
)
```

**To return the path of the database's default data directory:**

```
SELECT LEFT(filename,
            LEN(filename) - CHARINDEX(N'\', REVERSE(filename))) AS path
FROM master.dbo.sysdatabases
WHERE name = N'testdb'
```

Use the **'testdb'** name returned above in the following script.

**To create a user file group, add a data file to it, and set it as the default file group:**

```
ALTER DATABASE [testdb]
ADD FILEGROUP USERDATA001

ALTER DATABASE [testdb]
ADD FILE
(
```

```
        NAME = N'testdb_Data001',  
        FILENAME = N'<path to data directory>',  
        SIZE = 10MB,  
        FILEGROWTH = 50MB  
    )  
    TO FILEGROUP USERDATA001
```

**To set the userdata001 file group as the default file group:**

```
ALTER DATABASE [testdb] MODIFY FILEGROUP USERDATA001 DEFAULT
```

**Note:**

When creating file groups, the following configuration rules are important for proper file recovery, fault tolerance, performance, and archiving:

- You should create two separate data files, in two separate file groups: the .mdf file in the **PRIMARY** file group and the .ndf file in the **userdata001** group.
- The **userdata001** group should be set as the default file group.
- The .mdf file should be the only file in the **PRIMARY** file group that contains the system catalog. The .mdf file should be created with an initial **SIZE** parameter value of 5 MB and a **FILEGROWTH** parameter value of 5 MB. The other data file (.ndf) and log file (.ldf) should be created with an initial **SIZE** parameter value of 10 MB and a **FILEGROWTH** parameter value of 50 MB. In all cases, the **FILEGROWTH** parameter value should be a fixed size (not a percentage).

# Appendix B: Microsoft SQL Server Configuration Options

Most server configuration options are dynamically configured by Microsoft SQL Server. For OMi certification, you should not change the default options unless you are instructed to do so by HPE Software Support.

There are specific situations in which you may want to change the default settings. You can change these settings in the **sp\_configure** stored procedure, or in the various dialog boxes in Microsoft SQL Management Studio (mainly the Server Properties dialog box).

The default values should be used except for the following items:

- Agent XPs (A) = 1
- awe enabled (A) = 1. This should only be set when more than 4GB is required on a 32Bit machine. Note that this feature is deprecated in Microsoft SQL Server 2012.

Configuration Option	Default Value for SQL Server 2012 and 2014	Default Value for SQL Server 2008 and 2008 R2	OMi Certification
access check cache bucket count (A)	0	0	Default
access check cache quota (A)	0	0	Default
ad hoc distributed queries (A)	0	0	Default
affinity I/O mask (A, RR)	0	0	Default
affinity64 I/O mask (A, only available on 64-bit version of SQL Server)	0	0	Default
affinity mask (A)	0	0	Default
affinity64 mask (A, RR), only available on 64-bit version of SQL Server	0	0	Default

Configuration Option	Default Value for SQL Server 2012 and 2014	Default Value for SQL Server 2008 and 2008 R2	OMi Certification
Agent XPs (A)	0 (Changes to 1 when SQL Server Agent is started. Default value is 0 if SQL Server Agent is set to automatic start during Setup)	0 (Changes to 1 when SQL Server Agent is started. Default value is 0 if SQL Server Agent is set to automatic start during Setup)	1
allow updates (Obsolete. Do not use. Will cause an error during reconfigure.)	0	0	Default
awe enabled (A, RR)		0	Default, Unless more than 4GB memory is required on a 32Bit machine. This feature was deprecated in Microsoft SQL Server 2012.
backup compression default	0	0	Default
blocked process threshold (A)	0	0	Default
c2 audit mode (A, RR)	0	0	Default
clr enabled	0	0	Default
common criteria compliance enabled (A, RR)	0	0	Default
contained database authentication	0	N/A	Default
cost threshold for parallelism (A)	5	5	Default
cross db ownership chaining	0	0	Default
cursor threshold (A)	-1	-1	Default
Database Mail XPs (A)	0	0	Default
default full-text language(A)	1033	1033	Default
default language	0	0	Default

<b>Configuration Option</b>	<b>Default Value for SQL Server 2012 and 2014</b>	<b>Default Value for SQL Server 2008 and 2008 R2</b>	<b>OMi Certification</b>
default trace enabled (A)	1	1	Default
disallow results from triggers (A)	0	0	Default
EKM provider enabled	0	0	Default
filestream_access_level	0	0	Default
fill factor (A, RR)	0	0	Default
ft crawl bandwidth (max), see ft crawl bandwidth(A)	100	100	Default
ft crawl bandwidth (min), see ft crawl bandwidth(A)	0	0	Default
ft notify bandwidth (max), see ft notify bandwidth(A)	100	100	Default
ft notify bandwidth (min), see ft notify bandwidth(A)	0	0	Default
index create memory (A, SC)	0	0	Default
in-doubt xact resolution (A)	0	0	Default
lightweight pooling (A, RR)	0	0	Default
locks (A, RR, SC)	0	0	Default
max degree of parallelism (A)	0	0	Default
max full-text crawl range(A)	4	4	Default
max server memory (A, SC)	2147483647	2147483647	Default
max text repl size	65536	65536	Default



Configuration Option	Default Value for SQL Server 2012 and 2014	Default Value for SQL Server 2008 and 2008 R2	OMi Certification
max worker threads (A)	0 Zero auto-configures the number of max worker threads depending on the number of processors, using the formula $(256 + (<processors> - 4) * 8)$ for 32-bit SQL Server and twice that for 64-bit SQL Server.	0 Zero auto-configures the number of max worker threads depending on the number of processors, using the formula $(256 + (<processors> - 4) * 8)$ for 32-bit SQL Server and twice that for 64-bit SQL Server.	Default
media retention (A, RR)	0	0	Default
min memory per query (A)	1024	1024	Default
min server memory (A, SC)	0	0	Default
nested triggers	1	1	Default
network packet size (A)	4096	4096	Default
Ole Automation Procedures (A)	0	0	Default
open objects (A, RR, obsolete)	0	0	Default
optimize for ad hoc workloads (A)	0	0	Default
PH_timeout (A)	60	60	Default
precompute rank (A)	0	0	Default
priority boost (A, RR)	0	0	Default
query governor cost limit(A)	0	0	Default
query wait (A)	-1	-1	Default
recovery interval (A, SC)	0	0	Default
remote access (RR)	1	1	Default
remote admin connections	0	0	Default

<b>Configuration Option</b>	<b>Default Value for SQL Server 2012 and 2014</b>	<b>Default Value for SQL Server 2008 and 2008 R2</b>	<b>OMi Certification</b>
remote login timeout	10	20	Default
remote proc trans	0	0	Default
remote query timeout	600	600	Default
Replication XPs Option (A)	0	0	Default
scan for startup procs (A, RR)	0	0	Default
server trigger recursion	1	1	Default
set working set size (A, RR, obsolete)	0	0	Default
show advanced options	0	0	Default
SMO and DMO XPs (A)	1	1	Default
SQL Mail XPs (A)	N/A	0	Default
transform noise words (A)	0	0	Default
two digit year cutoff (A)	2049	2049	Default
user connections (A, RR, SC)	0	0	Default
User Instance Timeout (A, only appears in SQL Server 2008 Express.)	N/A	60	Default
user instances enabled (A, only appears in SQL Server 2008 Express.)	N/A	0	Default
user options	0	0	Default
xp_cmdshell (A)	0	0	Default

# Appendix C: Changing the Database Host

This section describes the procedure to migrate your OMi database schema to a different server.

1. Stop all OMi servers. This is required to prevent data loss and schema structure mismatch in the target environment.
2. Duplicate the OMi schemas you wish to migrate.  
This can be performed by an offline database backup or a database level export. For details, see one of the following sections, depending on your database type:
  - Microsoft SQL: ["Backing Up Databases" on page 35](#)
  - Oracle: ["Oracle Database Backup and Recovery" on page 67](#)
  - PostgreSQL: ["PostgreSQL Database Backup and Restore" on page 83](#)
3. Modify sessions table in your management database according to your new configuration.
4. Run the configuration wizard, specifying the location of the new database host in the Database Settings page.
5. Start all OMi servers.

# Appendix D: Support for Oracle Real Application Cluster

This appendix contains the configuration that needs to be done for OMi to work with Oracle Real Application Cluster. This information applies to advanced users only.

This appendix includes:

- ["About Oracle Real Application Cluster \(RAC\)" below](#)
- ["Single Client Access Name" on the next page](#)
- ["Client Side Configuration for Oracle RAC" on the next page](#)
- ["Server Side Configuration" on page 104](#)
- ["Create/Connect to One of the OMi Databases" on page 105](#)

## About Oracle Real Application Cluster (RAC)

A cluster is a collection of interconnected servers that appear as one server to the end user and to applications. Oracle Real Application Cluster (RAC) is Oracle's solution for high availability, scalability, and fault tolerance. It is based on clustered servers that share the same storage.

Oracle RAC is a single Oracle database installed on a cluster of hardware servers. Each server runs an instance of the database and all the instances share the same database files.


For more details about Oracle RAC, see the Oracle Clusterware Guide and the Oracle Real Application Clusters Administration and Deployment Guide in the Oracle documentation set of your release.

In this appendix, the following Oracle RAC example is used:

- Oracle RAC database name: OMiRAC
- Machine names: Server1, Server2
- On each machine there is an Oracle instance of OMiRAC:
  - SID on Server1: OMiRAC1
  - SID on Server2: OMiRAC2
- On each machine there is a virtual IP (Server1-Vip and Server2-Vip):
  - Server1-Vip is assigned to Server1
  - Server2-Vip is assigned to Server2

The virtual IP is in addition to the static IP assigned to the machine.

- The listeners on both servers are listening on the default port 1521 and support the database service OMiRAC.



## Single Client Access Name

In release 11g, Oracle introduced the Single Client Access Name (SCAN), as a preferred access method for clients connecting to the RAC. In this method clients are not required to configure individual nodes in the RAC, rather they use a single virtual IP known as the SCAN or the SCAN VIP.

The SCAN is a single network name defined for the cluster either in your organization's Domain Name Server (DNS) or in the Grid Naming Service (GNS) that rotates between several IP addresses reflecting multiple listeners in the cluster. The SCAN eliminates the need to change clients when nodes are added to or removed from the cluster.

The SCAN and its associated IP addresses provide a stable name for clients to use for connections, independent of the nodes that make up the cluster. Database server SCAN addresses, virtual IP addresses, and public IP addresses must all be on the same subnet.

**Note:** When using Oracle 11g RAC, it is recommended to use the SCAN method.

## Client Side Configuration for Oracle RAC

OMi uses DataDirect's JDBC driver to connect to regular Oracle databases and to Oracle RAC databases.

**Note:** OMi also supports Oracle Data Guard.

Make the following changes in OMi's configuration files before you create the management database or connect to an existing one on Oracle RAC:

1. **On all servers, create the file <OMi\_HOME>\conf\omi-tnsnames.ora.**

The format of **omi-tnsnames.ora** is the same as **tnsnames.ora**:

```
<DB service name> =  
(DESCRIPTION =  
  (ADDRESS_LIST =  
    (ADDRESS = (PROTOCOL = TCP)(HOST = <first instance virtual ip> )  
      (PORT = <first instance's listener port>))  
    (ADDRESS = (PROTOCOL = TCP)(HOST = <second instance virtual ip> )(PORT =  
<second instance's listener port>))  
    (... entry for each instance...)  
    (LOAD_BALANCE = on)  
    (FAILOVER = on)  
  )  
  (CONNECT_DATA =  
    (SERVER = DEDICATED)  
    (SERVICE_NAME = <DB service name>)  
  )  
)
```

where:

- **<DB service name>** is the name of a service the listeners support. It is the same one used in the CONNECT\_DATA part.
- **ADDRESS\_LIST** contains an address entry for each node in the RAC environment. In the case of Oracle using SCAN, it contains only the SCAN virtual IP. The address contains all the details needed for connecting to the node:
  - **HOST** contains the virtual-IP for that instance. It is important to use the virtual IP and not the static IP of the node for faster failure detection.
  - **PORT** is the port on which the listener is configured to listen on that given node.
- **FAILOVER** set to **on** allows the driver to try to connect to another node after failing to connect to one of the nodes. The connection attempts continue until a connection is successfully established to one of the nodes or until none of the nodes can be reached.

An attempt to connect to another node takes place only if the connection to the current node fails. If the connection is refused by the node (for example, communication to the node was established, but the connection was rejected), no attempt is made to connect to another node.

**Important:** Failover is for connection attempts only. If a connection fails during a database transaction, there is no failover of the transaction to another machine to continue the transaction.
- **LOAD\_BALANCE** set to **on** instructs the driver to distribute connection requests between the nodes to prevent overloading any single node. The order in which the driver accesses the nodes is random.
- **SERVER** is the connection mode you want to use. There are two possible connection modes: **dedicated** and **shared**. Configure this according to the Server configuration you support.
- **SERVICE\_NAME** is the name of a service that the listeners support. It is the same one you gave in **<DB service name>**.

**Note:** Keep in mind that when connecting to a single Oracle database system, you can specify either a service name or a SID.

In the above example, **omi-tnsnames.ora** would be configured as:

```
OMiRAC =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = Server1-Vip)(PORT = 1521))
    (ADDRESS = (PROTOCOL = TCP)(HOST = Server2-Vip)(PORT = 1521))
    (LOAD_BALANCE = on)
    (FAILOVER = on)
  )
  (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = OMiRAC)
  )
)
```

**Note:** If OMi schemas are deployed in more than one Oracle database or service, **omi-tnsnames.ora** should contain one entry for each database or database service. When all schemas are under the same RAC service, only one entry is required.

2. **On all servers, edit the file <OMi\_HOME>\conf\jdbc.drivers.properties.**

- a. Find the section named **#ddoracle**. In this section, there is a line for **ddoracle.url** similar to:

```
ddoracle.url=jdbc:mercury:oracle://${host}:${port};sid=${sid}
```

- b. Replace this line with either of the following lines (depending on the operating system):

■ **Windows.**

```
ddoracle.url=jdbc:mercury:oracle:TNSNamesFile=<OMi_HOME>\\conf\\omi-tnsnames.ora;TNSServerName=${sid}
```

where **\${sid}** correlates to the service name entered in the Service and Database Configuration step. This parameter is saved per OMi schema; therefore, if there is more than one service name in the system, make sure that **omi-tnsnames.ora** includes an entry for each service name.

Note that each back slash (\) in the path of **<OMi\_HOME>** is doubled.

■ **Linux.**

```
ddoracle.url=jdbc:mercury:oracle:TNSNamesFile=<OMi_HOME>/conf/omi-tnsnames.ora;TNSServerName=${sid}
```

3. **On all servers, open the directory <OMi\_HOME>\odb\conf. Locate the jdbc.properties file.**

- a. Find the line starting with **cmdb.url**.

- b. Replace this with the following line:

```
cmdb.url = jdbc:mercury:oracle:TNSNamesFile=<OMi_HOME>\\conf\\omi-tnsnames.ora;TNSServerName=<SERVICE NAME>
```

where <SERVICE\_NAME> is the entry in **omi-tnsnames.ora**, equivalent to the RAC service name.

- c. If the file does not exist, create an empty **jdbc.properties** file under the above folder and add the following entry:

```
Oracle = ddoracle  
cldb.url =  
jdbc:mercury:oracle:TNSNamesFile=<OMi_HOME>\\conf\\omi-  
tnsnames.ora;TNSListenerName=<SERVICE_NAME>
```

where <SERVICE\_NAME> is the entry in **omi-tnsnames.ora**, equivalent to the RAC service name.

- d. If your server is running Linux OS, replace all the double back slashes with single slashes.

## Server Side Configuration

In Oracle RAC, the Oracle listeners are always balancing the connection between the nodes according to either of the following algorithms:

- **Load Based (Default).** The listener redirects the connection according to the run queue length on the nodes. The least loaded node, CPU-wise, is connected.
- **Session based.** The listener balances the number of sessions between the nodes.

The first algorithm is optimized for short-lived connections, and less optimized for long-lived connections like the ones used with connection pools.

The entire load balancing is done during the connection time and not after it. This means that an established connection does not move to another node once connected.

## Recommendation for OMi Database

It is recommended to use the session based algorithm since OMi uses connection pools. One connection can be used for different purposes and has a long lifetime.

To make the Oracle listener use the session based algorithm, a listener parameter must be added to the listener.ora parameter file of every listener (each node has its own listener, so this change must be done on all listeners):

```
PREFER_LEAST_LOADED_NODE_<LISTENER_NAME> =OFF
```

where:

<LISTENER\_NAME> is the name of the listener. The default is LISTENER\_<node name>.

For example:

The default listener name on **Server1** is **LISTENER\_SERVER1**. In this case, you need to add the following to Server1's listener.ora file:

```
PREFER_LEAST_LOADED_NODE_LISTENER_SERVER1=OFF
```



## Create/Connect to One of the OMi Databases

When creating a new database schema or connecting to an existing schema in the configuration wizard, fill all the parameters as you usually would except for the following:

Parameter	Needed Value
Oracle Machine	One of the Virtual IPs. The SCAN virtual IP can be used.
Port	The local listener port on the Oracle machine or the port of the SCAN listener.
SID	The service name of the database. <b>Note:</b> This must be the same as <SERVICE_NAME> entered in <b>omi-tnsnames.ora</b> .

In the example, the parameters would be:

Parameter	Value
Oracle Machine	Server1-Vip
Port	1521
SID	OMiRAC

# Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

**Feedback on Database Guide (Operations Manager i 10.10)**

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to [ovdoc-asm@hpe.com](mailto:ovdoc-asm@hpe.com).

We appreciate your feedback!



Go OMi!