

Oracle® Real Application Clusters

Installation Guide

11g Release 2 (11.2) for Microsoft Windows x64 (64-Bit)

E48195-02

February 2015

Oracle Real Application Clusters Installation Guide, 11g Release 2 (11.2) for Microsoft Windows x64 (64-Bit)
E48195-02

Copyright © 2007, 2015, Oracle and/or its affiliates. All rights reserved.

Primary Author: Janet Stern

Contributing Authors: Mark Bauer, Jonathan Creighton, Prakash Jashnani, Aparna Kamath, Reema Khosla, Markus Michalewicz, Douglas Williams

Contributors: Eric Belden, Karin Brandauer, Willard Cole, Bjørn Engsig, Mark Fuller, Alexander Keh, Roland Knapp, To-Choi Lau, Philip Newlan, Satish Panchumarthy, Srinivas Poovala, Rajesh Prasad, Sudhe Sampath, Mohammed Shahnawaz Quadri, Ravi Ramikissoo, Kevin Reardon, Trivikrama Samudrala, David Schreiner, Jacqueline Sideri, Janelle Simmons, James Stenoish, Richard Strohm, Preethi Vallam, Nitin Vengurlekar, Guanzhou Wang, Xiaofang Wang, Martin Widjaja, Sergiusz Wolicki

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface	xv
Audience	xv
Documentation Accessibility	xv
Related Documents	xv
Conventions	xvi
What's New in Oracle Real Application Clusters Installation?	xix
New Features for Oracle RAC 11g Release 2 (11.2.0.4)	xix
New Features for Oracle RAC 11g Release 2 (11.2.0.3)	xix
New Features for Oracle RAC 11g Release 2 (11.2.0.2)	xx
New Features for Oracle RAC 11g Release 2 (11.2)	xx
New Features for Oracle RAC 11g Release 1 (11.1)	xxiv
Deprecated Features in Oracle Database 11g Release 2 (11.2)	xxxii
Deprecated Features in Oracle Database 11g Release 1 (11.1)	xxxii
Desupported Features in Oracle Database 11g Release 2 (11.2)	xxxii
1 Preinstallation Checklist	
1.1 Before You Start: Planning Your Installation	1-1
1.1.1 Check My Oracle Support	1-2
1.1.2 Consider Using Software Updates Option	1-2
1.1.3 Check Oracle Technology Network	1-3
1.1.4 Review Installation Steps for Oracle RAC	1-3
1.1.5 Review How to Use Cluster Verification Utility (CVU)	1-4
1.1.6 Review and Back Up Existing Oracle Installations	1-4
1.1.7 Review Globalization Requirements	1-5
1.1.8 Review Documentation	1-6
1.2 Overview of Server Hardware, Network, and Operating System	1-7
1.2.1 Review Server Hardware and Software Requirements	1-7
1.2.2 Overview of Server Environment Configuration	1-8
1.2.3 Overview of Oracle Base	1-9
1.2.4 Overview of Network and Internet Configuration	1-10
1.2.4.1 Overview of Network Interface Configuration	1-10
1.2.4.2 Overview of Configuring Internode Communication	1-10
1.2.4.3 Overview of Browser Requirements	1-10
1.2.5 Overview of Platform-Specific Server Configuration	1-11

1.3	Overview of Oracle RAC Storage.....	1-11
1.3.1	General Storage Considerations for Oracle RAC.....	1-13
1.3.2	Overview of Oracle Database and Recovery File Options	1-13
1.4	Requirements for the Oracle Database Vault Option	1-15
1.5	Additional Information to Prepare for Installation.....	1-15
1.5.1	Oracle ASM Is Installed with Oracle Grid Infrastructure.....	1-15
1.5.2	Obtain the Oracle ASM ASMSNMP Password for Oracle Enterprise Manager	1-16
1.5.3	Oracle Clusterware and Vendor Clusterware	1-16
1.5.4	Shared Components of Oracle RAC Installations.....	1-16
1.5.5	Simplified Upgrade of TIMESTAMP WITH TIME ZONE Data.....	1-16
1.5.6	Oracle Database 11g Real Application Clusters Release Compatibility	1-17
1.6	Confirming Cluster Readiness Using Cluster Verification Utility	1-18
1.7	Installing Oracle Database 11g Release 2 on Systems with Existing Databases	1-19

2 Installing Oracle Database 11g with Oracle Real Application Clusters

2.1	Reviewing Oracle Application Express Requirements	2-2
2.2	Selecting an Installation Type	2-2
2.3	Selecting an Installation Option.....	2-2
2.4	Selecting an Oracle Database Configuration Type	2-3
2.4.1	Preconfigured Database Types Supplied with Oracle Database	2-3
2.4.2	Using the Advanced Database Configuration Type.....	2-4
2.4.3	Installing Oracle Database with Other Languages	2-4
2.5	Selecting the Oracle Database Vault Option.....	2-4
2.5.1	Overview of Oracle Database Vault Accounts.....	2-4
2.5.1.1	Oracle Database Vault Owner	2-4
2.5.1.2	Oracle Database Vault Account Manager.....	2-4
2.5.2	Overview of Oracle Database Vault Installation.....	2-5
2.6	Selecting the Oracle Configuration Manager Option	2-5
2.7	Selecting a Database Name.....	2-5
2.8	Creating Database Passwords.....	2-6
2.9	Understanding the Actions of OUI and DBCA During Installation	2-7
2.9.1	Managing Database Services After Installation	2-7
2.9.2	Installing Earlier NLS Compatibility Definition Files	2-7
2.10	Using Oracle Universal Installer to Install Oracle RAC.....	2-8
2.11	Using Oracle Universal Installer to Install Oracle RAC One Node.....	2-9
2.12	Installing Oracle RAC and Oracle Database Vault	2-10
2.13	Using DBCA to Configure Oracle RAC One Node.....	2-10
2.14	Updating Environment Variables on Remote Nodes.....	2-11

3 Creating Oracle Real Application Clusters Databases with Oracle Database Configuration Assistant

3.1	Using Oracle Database Configuration Assistant with Oracle RAC	3-1
3.2	Optimizing Your Database Environment by Using DBCA	3-2
3.3	Updating the Listener Used By an Earlier Release	3-3
3.4	Verifying Requirements for DBCA	3-3
3.5	Creating an Oracle RAC Database with DBCA.....	3-4
3.5.1	Configuring Disks for Use With Oracle ASM	3-4

3.5.2	Starting DBCA.....	3-5
3.5.3	Using DBCA to Create an Oracle RAC Database	3-5
3.6	Deleting an Oracle RAC Database with DBCA	3-8
3.7	Setting Environment Variables for Enterprise Manager DB Control.....	3-9

4 Oracle Real Application Clusters Postinstallation Procedures

4.1	Required Postinstallation Tasks.....	4-1
4.1.1	Determine If Any Patches Are Required For Your New Software.....	4-1
4.1.2	Types of Patches Available for Oracle Software	4-2
4.1.3	Downloading and Installing Patch Updates.....	4-2
4.1.4	Configuring Exceptions for the Windows Firewall.....	4-3
4.1.5	Configuring Oracle Products	4-4
4.1.5.1	Configuring Oracle Database Vault.....	4-4
4.1.5.2	Configuring Oracle Database Security Settings	4-4
4.1.6	Configuring Storage for External Tables, Shared Files, or Directory Objects.....	4-4
4.2	Recommended Postinstallation Tasks	4-5
4.2.1	Setting the Oracle User Environment Variables.....	4-5
4.2.2	Recompiling All PL/SQL modules.....	4-5
4.2.3	Setting Up User Accounts.....	4-5
4.2.4	Configuring Oracle Administration Assistant for Windows	4-6
4.2.5	Using Oracle9i Database Language and Definition Files with Oracle Database 11g	4-6
4.2.6	Logging in to Oracle Enterprise Manager Database Control.....	4-6
4.3	Oracle Configuration Manager Postinstallation Configuration for Oracle RAC	4-7
4.3.1	Additional Step for Oracle Enterprise Manager Grid Control.....	4-7
4.4	Enabling and Disabling Database Options	4-8
4.5	Using Earlier Oracle Database Releases with Oracle Grid Infrastructure.....	4-9
4.5.1	General Restrictions for Using Earlier Oracle Database Releases	4-9
4.5.2	Required Patches for Earlier Oracle Database Releases or Databases Being Upgraded... 4-10	
4.5.3	Pinning Cluster Nodes for Oracle Database Release 10.2 or 11.1.....	4-11

5 Configuring the Server Parameter File in Oracle Real Application Clusters Environments

5.1	About Parameter Files and Oracle RAC.....	5-1
5.2	Creating Server Parameter Files for Oracle RAC	5-2
5.3	Storage Requirements for the Server Parameter File in Oracle RAC	5-2
5.4	Parameter File Search Order in Oracle RAC.....	5-3
5.5	Migrating to an SPFILE in Oracle RAC Environments	5-3
5.6	Server Parameter File Errors in Oracle RAC.....	5-4

6 Configuring Server Pools and Oracle Enterprise Manager Database Control

6.1	Understanding Server Pools.....	6-1
6.1.1	Oracle RAC One Node and Server Pools	6-2
6.1.2	Policy-Based Cluster and Capacity Management.....	6-2
6.2	Understanding the Database Control Implementation for Oracle RAC.....	6-3
6.3	Understanding EMCA and Enterprise Manager Control with Oracle RAC.....	6-3

6.4	Configuring Database Control During Installation	6-3
6.5	Configuring Database Control with DBCA	6-4
6.6	Configuring Database Control After Installation Using EMCA	6-5
6.6.1	Using EMCA in Interactive Mode	6-5
6.6.2	Using EMCA in Silent Mode With an Input File	6-6
6.7	Managing Database Control Using EMCA	6-6
6.7.1	Changing the Database Control Implementation for Your Oracle RAC Database....	6-7
6.7.2	Updating Database Control Configuration When Adding and Removing Nodes....	6-8
6.7.2.1	Using EMCA After Adding Nodes to the Cluster	6-8
6.7.2.2	Using EMCA After Removing Nodes from the Cluster	6-8
6.7.3	Changing the Ports Used by Database Control.....	6-9
6.8	Enterprise Manager and EMCA Troubleshooting Tips	6-10
6.8.1	Updating Database Control After Changing the Database Listener Port	6-10
6.8.2	Upgrading Grid Control Agents for Oracle Database 11g Release 2	6-10
6.8.3	Updating Database Control After Changing the Host Name or IP Address.....	6-11
6.8.4	Using EMCA After Changing the TNS_ADMIN Configuration.....	6-11

7 Understanding the Oracle Real Application Clusters Installed Configuration

7.1	Understanding the Oracle RAC Configuration.....	7-1
7.2	Understanding Time Zone Settings on Cluster Nodes	7-1
7.3	Working with Multiple Oracle Home Directories on Windows.....	7-2
7.4	Understanding the Database Components Created by DBCA	7-2
7.4.1	Tablespaces and Data Files.....	7-3
7.4.2	Control Files.....	7-4
7.4.3	Online Redo Log Files	7-4
7.5	Managing Undo Tablespaces in Oracle RAC	7-5
7.6	Using Initialization Parameter Files	7-5
7.7	Oracle Database Vault Default Audit Policy and Initialization Parameters	7-5
7.8	Understanding Oracle Net Services Configuration Files	7-5
7.8.1	Database Services for an Oracle RAC Database.....	7-6
7.8.2	Configuring the Listener Configuration File (listener.ora).....	7-7
7.8.2.1	Grid Naming Service and Listeners.....	7-7
7.8.2.2	Local Listeners	7-7
7.8.2.3	How Oracle Database Uses the Listener File (listener.ora)	7-7
7.8.2.4	Configuring Multiple Listeners for Your Database Using DBCA.....	7-9
7.8.3	Net Service Names (tnsnames.ora File).....	7-9
7.8.3.1	Net Service Names for Database Connections	7-10
7.8.3.2	Net Service Name for Instance Connections	7-11
7.8.4	Net Services Profile (sqlnet.ora File)	7-12
7.9	Using SCAN with Oracle RAC Databases	7-12
7.9.1	Understanding SCAN	7-13
7.9.2	Understanding SCAN Listeners and Listener Registration	7-13
7.9.3	Understanding Failover, Load Balancing, and SCAN	7-14
7.9.4	Understanding Shared Server Use in an Oracle RAC Database.....	7-16

8 Removing Oracle Real Application Clusters Software

8.1	Overview of Deinstallation Tools Actions	8-1
-----	--	-----

8.2	Identifying All Instances On a Cluster	8-2
8.3	Deinstalling Oracle RAC Software	8-2
8.3.1	About the Deinstallation Tool.....	8-2
8.3.2	Using the Deinstallation Tool for Oracle RAC	8-6
8.3.3	Example of a Deinstallation Parameter File for Oracle RAC	8-6

A Troubleshooting Oracle Real Application Clusters Installations

A.1	Troubleshooting Oracle RAC Installation Errors.....	A-1
A.1.1	Verifying Requirements.....	A-1
A.1.2	Understanding How To Resolve Installation Errors	A-2
A.1.3	Locating the Log File For an Installation Session.....	A-2
A.2	Resolving Oracle RAC Installations Errors.....	A-2
A.2.1	General Installation Issues.....	A-2
A.2.2	Performing Cluster Diagnostics During Oracle RAC Installations.....	A-4
A.2.3	Reviewing the Log of an Installation Session	A-4
A.3	Troubleshooting Configuration Assistant Errors.....	A-4
A.3.1	Reviewing Configuration Assistant Errors.....	A-4
A.3.2	Identifying Configuration Assistant Failures	A-5
A.3.3	Resolving Irrecoverable Errors Reported by Configuration Assistants	A-5
A.4	Troubleshooting Silent or Noninteractive Installations	A-5
A.5	About Using CVU Cluster Healthchecks After Installation	A-6
A.6	Cleaning Up After a Failed Installation.....	A-7

B Using Scripts or Response Files to Create Oracle Real Application Clusters Databases

B.1	Using DBCA to Generate Installation Scripts for Oracle RAC.....	B-1
B.2	About DBCA Noninteractive (Silent) Configuration for Oracle RAC	B-2
B.3	Using DBCA Noninteractive (Silent) Configuration for Oracle RAC.....	B-2
B.4	Understanding How Response Files are Used During Installation	B-3
B.4.1	Deciding to Use Silent Mode or Response File Mode	B-3
B.4.2	General Procedure for Using Response Files	B-4
B.5	Preparing Response Files.....	B-4
B.5.1	Editing a Response File Template	B-4
B.5.2	Recording Response Files	B-6
B.6	Running Oracle Universal Installer Using a Response File.....	B-6
B.6.1	Viewing the Help Information for OUI.....	B-6
B.6.2	Running OUI Using a Response File for Input.....	B-7
B.7	Running Configuration Assistants Using Response Files	B-7
B.7.1	Running Oracle Database Configuration Assistant in Response File or Silent Mode.....	B-8
B.7.2	Running Oracle Net Configuration Assistant Using Response Files.....	B-9
B.7.3	About the Postinstallation Configuration File.....	B-9
B.7.4	Creating a Password Response File	B-10
B.7.5	Performing Postinstallation Configuration Using a Response File.....	B-11

C Converting to Oracle RAC and Oracle RAC One Node from Single-Instance Oracle Databases

C.1	Administrative Issues for Converting Databases to Oracle RAC	C-1
C.2	Converting to Oracle RAC and Oracle RAC One Node Using DBCA	C-2
C.2.1	Converting Oracle Database 11g Installations to Oracle RAC Using DBCA	C-2
C.2.1.1	Use DBCA to Create an Image of the Single-Instance Database	C-3
C.2.1.2	Complete the Oracle Clusterware Installation	C-3
C.2.1.3	Validate the Cluster	C-3
C.2.1.4	Copy the Preconfigured Database Image	C-3
C.2.1.5	Install Oracle Database 11g Software with Oracle RAC	C-3
C.2.2	Converting Single Instance on a Cluster to Oracle RAC One Node Using DBCA ...	C-4
C.2.3	Converting Single Instance on a Cluster to Oracle RAC Using DBCA	C-4
C.2.3.1	Single-Instance Database an Oracle RAC-Enabled Home.....	C-5
C.2.3.1.1	Automated Conversion Procedure Using DBCA	C-5
C.2.3.1.2	Manual Conversion Procedure	C-6
C.2.3.2	Single-Instance Database Using an Oracle RAC-Disabled Home	C-7
C.2.4	Converting an Oracle RAC One Node Database to an Oracle RAC Database	C-8
C.3	Preparing to Convert with rconfig and Oracle Enterprise Manager.....	C-9
C.3.1	Prerequisites for Converting to Oracle RAC Databases	C-9
C.3.2	Configuration Changes During Oracle RAC Conversion Using rconfig	C-9
C.3.3	Converting Databases to Oracle RAC Using rconfig or Enterprise Manager	C-10
C.3.4	Converting Databases to Oracle RAC Using Oracle Enterprise Manager	C-10
C.4	Converting Databases to Oracle RAC Using rconfig.....	C-12
C.5	Example of rconfig XML Input Files for ConvertToRAC	C-13
C.6	Postconversion Steps	C-15

D Directory Structure for Oracle Real Application Clusters Environments

D.1	Understanding the Oracle RAC Directory Structure.....	D-1
D.2	Directory Structures for Oracle RAC	D-1

E How to Stop Processes in an Existing Oracle Real Application Clusters Database

E.1	Back Up the Oracle RAC Database.....	E-1
E.2	Shut Down Oracle RAC Databases	E-1
E.3	Stop All Oracle Processes.....	E-2
E.3.1	Stopping Oracle Database Processes	E-2
E.3.2	Stopping Oracle Clusterware and Oracle ASM Instances	E-3

F Managing Oracle Database Port Numbers

F.1	About Managing Ports	F-1
F.2	Viewing Port Numbers and Access URLs	F-2
F.3	Port Numbers and Protocols of Oracle Components	F-2
F.4	Changing the Oracle Management Agent Port	F-4
F.5	Changing the Oracle Enterprise Manager Database Control Ports.....	F-4
F.6	Changing the Oracle XML DB Ports	F-5
F.7	Changing the Oracle Services for Microsoft Transaction Server Port.....	F-6

G Oracle Enterprise Manager Configuration Assistant Reference

G.1	About EMCA	G-1
G.1.1	Overview	G-1
G.1.2	Operational Notes.....	G-1
G.1.2.1	Usage Information.....	G-1
G.1.2.2	EMCA Parameters	G-2
G.1.2.3	Using Tool/Utility Help	G-3
G.2	EMCA Command Reference	G-3
G.2.1	addInst.....	G-4
G.2.1.1	Syntax	G-4
G.2.1.2	Parameters/Options	G-5
G.2.2	addNode.....	G-5
G.2.2.1	Syntax	G-5
G.2.2.2	Parameters	G-5
G.2.2.3	Examples	G-6
G.2.3	config all db	G-6
G.2.3.1	Syntax	G-6
G.2.3.2	Parameters	G-6
G.2.4	config centralAgent	G-7
G.2.4.1	Syntax	G-7
G.2.4.2	Parameters	G-7
G.2.5	config dbcontrol db	G-7
G.2.5.1	Syntax	G-7
G.2.5.2	Options.....	G-8
G.2.5.3	Examples.....	G-8
G.2.6	deconfig all db.....	G-8
G.2.6.1	Syntax	G-8
G.2.6.2	Parameters	G-9
G.2.7	deconfig centralAgent.....	G-9
G.2.7.1	Syntax	G-9
G.2.7.2	Parameters	G-9
G.2.7.3	Examples.....	G-10
G.2.8	deconfig dbcontrol db.....	G-10
G.2.8.1	Syntax	G-10
G.2.8.2	Parameters	G-10
G.2.8.3	Examples.....	G-11
G.2.9	deleteInst.....	G-11
G.2.9.1	Syntax	G-11
G.2.9.2	Parameters	G-11
G.2.10	deleteNode.....	G-11
G.2.10.1	Syntax	G-11
G.2.10.2	Parameters	G-12
G.2.10.3	Examples.....	G-12
G.2.11	displayConfig dbcontrol.....	G-12
G.2.11.1	Syntax	G-12
G.2.11.2	Parameters	G-12
G.2.12	reconfig dbcontrol.....	G-13

G.2.12.1	Syntax	G-13
G.2.12.2	Parameters	G-13
G.2.12.3	Examples	G-13
G.2.13	reconfig ports.....	G-13
G.2.13.1	Syntax	G-13
G.2.13.2	Parameters	G-13
G.2.13.3	Examples	G-13
G.2.14	restore	G-14
G.2.14.1	Syntax	G-14
G.2.14.2	Parameters	G-14
G.2.15	upgrade	G-14
G.2.15.1	Syntax	G-14
G.2.15.2	Parameters	G-15
G.2.16	version	G-15
G.2.16.1	Syntax	G-15

Index

List of Tables

1-1	Hard Disk Space Requirements for Oracle RAC.....	1-12
1-2	Supported Storage Options for Oracle Database and Recovery Files.....	1-14
4-1	Database Options for Chopt Tool Command.....	4-8
4-2	Patches Required for Oracle Database and Oracle Grid Infrastructure Release 11.2.0.3	4-10
6-1	EMCA Command Parameters for Configuring Port Assignments	6-9
7-1	Tablespace Names Used with Oracle Real Application Clusters Databases	7-3
7-2	Symbolic Link Names for Database Files on Windows Platforms	7-3
B-1	Response Files for Oracle Database.....	B-5
D-1	Directory Structure for a Sample OFA-Compliant Environment.....	D-2
F-1	Ports Used in Oracle Components	F-2
G-1	EMCA Optional Command-Line Parameters.....	G-2
G-2	EMCA Commands.....	G-3
G-3	addInst Options.....	G-5
G-4	addNode Options	G-5
G-5	config all db Options	G-6
G-6	config centralAgent Options	G-7
G-7	config dbcontrol Options	G-8
G-8	deconfig all db Options	G-9
G-9	deconfig centralAgent Options	G-9
G-10	deconfig dbcontrol db Options	G-10
G-11	deleteInst Options	G-11
G-12	deleteNode Options.....	G-12
G-13	displayConfig dbcontrol Options.....	G-12
G-14	reconfig dbcontrol Options.....	G-13
G-15	reconfig ports Options.....	G-13
G-16	restore Options	G-14
G-17	upgrade Options	G-15

List of Examples

4-1	How to Enable the Oracle Label Security Database Option Using Chopt	4-8
6-1	Configuring Database Control and Daily Database Backups using EMCA	6-5
6-2	Using a Response File with EMCA to Configure Database Control	6-6
6-3	Reconfiguring Database Control for an Oracle RAC Database.....	6-8
7-1	Example listener.ora File for an Oracle RAC Node	7-8
7-2	Example tnsnames.ora File.....	7-9
7-3	Example tnsnames.ora Entry for an Oracle RAC Database	7-10
7-4	Sample Net Service Name Entry for a Database Connection.....	7-11
7-5	Connecting to Oracle RAC Using the Easy Connect Naming Method.....	7-11
7-6	Sample Net Service Name Entry for an Instance Connection.....	7-11
8-1	Running the Deinstallation Tool Installed in an Oracle Home	8-6
8-2	Running the Standalone Version of the Deinstallation Tool	8-6
8-3	Sample Parameter File to use with the Deinstallation Tool to Uninstall Oracle RAC.....	8-6
B-1	Sample Password Response File for Oracle RAC Databases.....	B-10
B-2	Running Configuration Assistants in Response File Mode.....	B-11
C-1	Example rconfig ConvertToRAC XML file for Policy-Managed Databases	C-13
C-2	Example rconfig ConvertToRAC XML file for Administrator-Managed Databases.....	C-14

Preface

Oracle Real Application Clusters Installation Guide for Microsoft Windows x64 (64-Bit) explains how to install and configure Oracle Real Application Clusters (Oracle RAC). Before you use this guide, you must first complete successfully an installation of Oracle Clusterware, as described in *Oracle Grid Infrastructure Installation Guide* for your platform.

Audience

Oracle Real Application Clusters Installation Guide for Microsoft Windows x64 (64-Bit) provides database installation information for database administrators (DBAs) who install and configure Oracle RAC.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information, see the following documents in the Oracle Database 11g Release 2 (11.2) documentation set.

Oracle Real Application Clusters Documentation

Most Oracle error message documentation is available only in HTML format. If you have access only to the Oracle Documentation media, then browse the error messages by range. After you find a range, use your browser's "find in page" feature to locate a specific message. When connected to the Internet, you can search for a specific error message using the error message search feature of the Oracle online documentation.

Installation Guides

- *Oracle Diagnostics Pack Installation Guide*

- *Oracle Grid Infrastructure Installation Guide for IBM AIX on POWER Systems*
- *Oracle Grid Infrastructure Installation Guide for HP-UX*
- *Oracle Grid Infrastructure Installation Guide for Linux*
- *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)*
- *Oracle Grid Infrastructure Installation Guide for Solaris Operating System*

Operating System-Specific Administrative Guides

- *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems*
- *Oracle Database Platform Guide for Microsoft Windows*

Oracle Database 11g Release 2 (11.2) Real Application Clusters Management

- *Oracle Clusterware Administration and Deployment Guide*
- *Oracle Real Application Clusters Administration and Deployment Guide*
- *Oracle Database 2 Day DBA*
- *Oracle Database 2 Day + Real Application Clusters Guide*
- *Getting Started with the Oracle Diagnostics Pack*

Generic Documentation

- *Oracle Database Sample Schemas* describes the sample schemas provided in a basic installation of Oracle Database. Many of the examples in the Oracle Database Documentation Library use these schemas.
- *Oracle Database New Features Guide*
- *Oracle Database Net Services Administrator's Guide*
- *Oracle Database Concepts*
- *Oracle Database Reference*

To download free release notes, installation documentation, white papers, or other collateral, go to the Oracle Technology Network (OTN). You must register online before using OTN; registration is free and can be done at the following web site:

<http://www.oracle.com/technetwork/community/join/overview>

If you have a user name and password for OTN, then you can go directly to the documentation section of the OTN web site:

<http://www.oracle.com/technetwork/indexes/documentation/index.html>

Oracle error message documentation is available only in HTML. You can browse the error messages by range in the Documentation directory of the installation media. When you find a range, use your browser's "find in page" feature to locate a specific message. When connected to the Internet, you can search for a specific error message using the error message search feature of the Oracle online documentation.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

What's New in Oracle Real Application Clusters Installation?

This chapter describes new features as they pertain to the installation and configuration of Oracle Real Application Clusters (Oracle RAC). The topics are:

- [New Features for Oracle RAC 11g Release 2 \(11.2.0.4\)](#)
- [New Features for Oracle RAC 11g Release 2 \(11.2.0.3\)](#)
- [New Features for Oracle RAC 11g Release 2 \(11.2.0.2\)](#)
- [New Features for Oracle RAC 11g Release 2 \(11.2\)](#)
- [New Features for Oracle RAC 11g Release 1 \(11.1\)](#)
- [Deprecated Features in Oracle Database 11g Release 2 \(11.2\)](#)
- [Deprecated Features in Oracle Database 11g Release 1 \(11.1\)](#)
- [Desupported Features in Oracle Database 11g Release 2 \(11.2\)](#)

New Features for Oracle RAC 11g Release 2 (11.2.0.4)

The following is a list of new features that are available starting with Oracle RAC 11g Release 2 (11.2.0.4):

Windows Server 2012 and Windows Server 2012 R2 Support

Starting with Oracle Database 11g Release 2 (11.2.0.4), Oracle Database server, client, and Oracle RAC are supported on Windows Server 2012 (64-bit only) and Windows Server 2012 R2 (64-bit only).

Note: Windows Database Bundle Patch 11.2.0.4.11 is required for Windows Server 2012 R2 support.

New Features for Oracle RAC 11g Release 2 (11.2.0.3)

The following is a list of new features that are available starting with Oracle RAC 11g Release 2 (11.2.0.3):

Oracle Database Quality of Service Management

The Database Quality of Service (QoS) Management Server allows system administrators to manage application service levels hosted in Oracle Database clusters by correlating accurate run-time performance and resource metrics and analyzing with

an expert system to produce recommended resource adjustments to meet policy-based performance objectives.

See Also: *Oracle Database Quality of Service Management User's Guide*

New Features for Oracle RAC 11g Release 2 (11.2.0.2)

The following new features are available starting with Oracle Database 11g Release 2 (11.2.0.2):

New Software Updates Option

Use the Software Updates feature to dynamically download and apply software updates as part of the Oracle Database installation. You can also download the updates separately using the `downloadUpdates` option and later apply them during the installation by providing the location where the updates are present.

Oracle Real Application Clusters One Node (Oracle RAC One Node)

Oracle RAC One Node is a single instance of Oracle RAC running on one node in a cluster. You can use Oracle RAC One Node to consolidate many databases onto a single cluster with minimal overhead, and still provide the high availability benefits of failover protection, online rolling patch application, and rolling upgrades for the operating system and for Oracle Clusterware. With Oracle RAC One Node you can standardize all Oracle Database deployments across your enterprise.

You can use Oracle Database and Oracle Grid Infrastructure configuration assistants, such as Oracle Database Configuration Assistant (DBCA) and RCONFIG, to configure Oracle RAC One Node databases.

Oracle RAC One Node is a single Oracle RAC database instance. You can use a planned online relocation to start a second Oracle RAC One Node instance temporarily on a new target node, so that you can migrate the current Oracle RAC One Node instance to this new target node. After the migration, the source node instance is shut down. Oracle RAC One Node databases can also fail over to another cluster node within its hosting server pool if their current node fails.

Oracle RAC One Node is not supported if you use a third-party clusterware software, such as Veritas, SFRAC, IBMPowerHA, or HP Serviceguard. Oracle Solaris Cluster is currently not supported.

Starting with Oracle Database 11g release 2 (11.2.0.2), Oracle RAC One Node is supported on all platforms where Oracle Real Application Clusters (Oracle RAC) is certified. Oracle RAC One Node supports Oracle Data Guard starting with Oracle Database 11g release 2 (11.2.0.2).

New Features for Oracle RAC 11g Release 2 (11.2)

The following is a list of new features for Oracle RAC 11g release 2 (11.2):

- [Oracle Automatic Storage Management and Oracle Clusterware Installation](#)
- [Oracle Automatic Storage Management Cluster File System \(Oracle ACFS\)](#)
- [Cluster Time Synchronization Service](#)
- [Oracle Database Configuration Assistant and Listeners](#)
- [Daylight Savings Time Upgrade of `TIMESTAMP WITH TIMEZONE` Data Type](#)
- [Oracle Enterprise Manager Database Control Provisioning](#)

- Oracle Enterprise Manager and Oracle Clusterware Resource Management
- Oracle Enterprise Manager Patch Application
- Grid Plug and Play
- Improved Deployment, Deconfiguration, and Deinstallation
- SCAN for Simplified Client Access
- SRVCTL Command Enhancements for Patching
- Desktop Class and Server Class Installation Options
- Zero Downtime Patching for Oracle RAC
- Deinstallation Tool

Oracle Automatic Storage Management and Oracle Clusterware Installation

With Oracle Grid Infrastructure 11g release 2 (11.2), Oracle Automatic Storage Management (Oracle ASM) and Oracle Clusterware are installed into a single home directory, which is referred to as the Grid Infrastructure home. Configuration assistants that start after the installer interview process configure Oracle ASM and Oracle Clusterware.

The installation of the combined products is called Oracle Grid Infrastructure. However, Oracle Clusterware and Oracle Automatic Storage Management remain separate products.

Oracle Automatic Storage Management Cluster File System (Oracle ACFS)

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is a new multiplatform, scalable file system and storage management design that extends Oracle Automatic Storage Management (Oracle ASM) technology to support all application data. Oracle ACFS provides dynamic file system resizing, improved performance using the distribution, balancing, and striping technology across all available disks, and storage reliability through the Oracle ASM mirroring and parity protection.

Note: For Oracle ASM 11g Release 2 (11.2.0.1), Oracle ACFS is supported on only:

- Windows Server 2003 64-bit
- Windows Server 2003 R2 64-bit

Starting with Oracle Database 11g Release 2 (11.2.0.2), Oracle ACFS is also supported on:

- Windows Server 2008, 64-bit
- Windows Server 2008 R2 64-bit

Starting with Oracle Database 11g Release 2 (11.2.0.4), Oracle ACFS is also supported on:

- Windows Server 2012, 64-bit
 - Windows Server 2012 R2 64-bit
-
-

Oracle ASM Dynamic Volume Manager (Oracle ADVM) extends Oracle ASM by providing a disk driver interface to Oracle ASM storage allocated as Oracle ASM volume files. You can use Oracle ADVM to create virtual disks that contain file

systems. File systems and other disk-based applications issue I/O requests to Oracle ADVM volume devices as they would to other storage devices on a vendor operating system. The file systems contained on Oracle ASM volumes can support files beyond Oracle database files, such as executable files, report files, trace files, alert logs, and other application data files.

Cluster Time Synchronization Service

Cluster node times should be synchronized, particularly if the cluster is to be used for Oracle Real Application Clusters. With this release, Oracle Clusterware provides Cluster Time Synchronization Service (CTSS), which ensures that there is a synchronization service in the cluster. If neither Network Time Protocol (NTP) or Windows Time Service is not found during cluster configuration, then CTSS is configured to ensure time synchronization.

Oracle Database Configuration Assistant and Listeners

With this release, Oracle Database Configuration Assistant (DBCA) no longer sets the value for `LOCAL_LISTENER`. When Oracle Clusterware starts the database resource, it updates the instance parameters. The `LOCAL_LISTENER` is set to the virtual IP endpoint of the local node listener address. You should not modify the setting for `LOCAL_LISTENER`. New installation instances register only with single client access name (SCAN) listeners as remote listeners. SCANS are virtual IP addresses assigned to the cluster, rather than to individual nodes, so cluster members can be added or removed without requiring updates of clients served by the cluster. Upgraded databases continue to register with all node listeners, and additionally with the SCAN listeners.

Daylight Savings Time Upgrade of `TIMESTAMP WITH TIMEZONE` Data Type

When time zone files are updated to a new version, `TIMESTAMP WITH TIMEZONE` (TSTZ) data could become stale. In previous releases, database administrators ran the SQL script `utltzuv2.sql` to detect TSTZ data affected by the time zone version changes and then had to perform extensive manual procedures to update the TSTZ data.

With this release, TSTZ data is updated transparently with very minimal manual procedures using newly provided `DBMS_DST` PL/SQL packages. In addition, there is no longer a need for clients to patch their time zone files.

See Also:

- *Oracle Database Upgrade Guide* for information about upgrading time zone files
- *Oracle Database Globalization Support Guide* for information about how to upgrade the time zone file and TSTZ data
- *Oracle Call Interface Programmer's Guide* for information about performance effects of clients and servers operating with different versions of time zone files

Oracle Enterprise Manager Database Control Provisioning

Oracle Enterprise Manager Database Control 11g provides the capability to provision automatically Oracle Grid Infrastructure and Oracle RAC installations on new nodes, and then extend the existing Oracle Grid Infrastructure and Oracle RAC database to these provisioned nodes. This provisioning procedure requires a successful Oracle RAC installation before you can use this feature.

See Also: *Oracle Real Application Clusters Administration and Deployment Guide* for information about this feature

Oracle Enterprise Manager and Oracle Clusterware Resource Management

With this release, you can use Oracle Enterprise Manager Cluster home page to perform full administrative and monitoring support for both standalone database and Oracle RAC environments, using High Availability Application and Oracle Clusterware resource management. Such administrative tasks include creating and modifying server pools.

Oracle Enterprise Manager Patch Application

With this release, you can apply patches to the Oracle RAC database using Oracle Enterprise Manager. A new Oracle Enterprise Manager feature, the Provisioning Advisor Console, enables you to customize, monitor, and deploy patch applications to nodes on the cluster.

Grid Plug and Play

In the past, adding or removing servers in a cluster required extensive manual preparation. With this release, you can continue to configure server nodes manually, or you can use Grid Plug and Play to configure them dynamically as nodes are added or removed from the cluster.

Grid Plug and Play reduces the costs of installing, configuring, and managing server nodes by starting a Grid Naming Service (GNS) within the cluster to allow each node to perform the following tasks dynamically:

- Negotiate appropriate network identities for itself
- Acquire additional information it requires to operate from a configuration profile
- Configure or reconfigure itself using profile data, making host names and addresses resolvable on the network

Because servers perform these tasks dynamically, adding and removing nodes simply requires an administrator to connect the server to the cluster, and to allow the cluster to configure the node. Using Grid Plug and Play and best practice recommendations, you can add a node to the database cluster as part of the server restart, and remove a node from the cluster automatically when a server is turned off.

Improved Deployment, Deconfiguration, and Deinstallation

Oracle configuration assistants ensure the success of recommended deployments, and prevent configuration issues.

Oracle configuration assistants provide the capability of deconfiguring and deinstalling Oracle Real Application Clusters, without requiring additional manual steps.

SCAN for Simplified Client Access

With this release, the single client access name (SCAN) is the address to provide for all clients connecting to the cluster. SCAN is a domain name registered to at least one and up to three IP addresses, either in the domain name service (DNS) or the Grid Naming Service (GNS). SCAN eliminates the need to change clients when nodes are added to or removed from the cluster. Clients using SCAN can also access the cluster using Easy Connect Naming.

SRVCTL Command Enhancements for Patching

With this release, you can use the Server Control Utility (SRVCTL) to shut down all Oracle software running within an Oracle home, in preparation for patching. Oracle Grid Infrastructure patching is automated across all nodes, and patches can be applied in a multinode, multipatch method.

Desktop Class and Server Class Installation Options

In this release, there are two installation types: Desktop Class and Server Class.

The Desktop Class installation type is a simplified installation with a minimal number of manual configuration choices. The Desktop Class installation performs a full Oracle Database installation with a basic configuration.

The Server Class installation type allows for more advanced configuration options. Select this option when installing Oracle RAC, if you use Oracle Enterprise Manager Grid Control, or to configure database storage on Oracle ASM.

Zero Downtime Patching for Oracle RAC

The Oracle patch utility, Opatch, can apply patches in a multinode, multipatch method, and does not start instances that have a nonrolling patch applied to it if other instances of the database do not have that patch. Opatch also detects if the database schema is an earlier patch level than the new patch, and it runs SQL statements to upgrade the database schema to the current patch level.

See Also: *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX*

Deinstallation Tool

Oracle Universal Installer (OUI) no longer removes Oracle software. Use the new Deinstallation Tool (`deinstall.bat`) available on the installation media before installation, and in the Oracle home directory after installation. This tool can also be downloaded from Oracle Technology Network.

See Also: [Chapter 8, "Removing Oracle Real Application Clusters Software"](#) for more information

New Features for Oracle RAC 11g Release 1 (11.1)

This section describes Oracle Database 11g release 1 features as they pertain to the installation and configuration of Oracle Real Application Clusters (Oracle RAC).

The topics in this section are:

- [Changes in Installation Documentation](#)
- [Changes in the Installation Options](#)
- [New Components Available for Installation](#)
- [Enhancements and New Features for Installation](#)

Changes in Installation Documentation

With Oracle Database 11g release 1 (11.1), Oracle Clusterware can be installed or configured as an independent product. In addition, new documentation is provided for Oracle Database storage administration. For installation planning, note the documentation described in the following subsections.

Oracle Database 2 Day + Real Application Clusters Guide

This guide provides an overview and examples of the procedures to install and configure a two-node Oracle Clusterware and Oracle RAC environment.

Oracle Clusterware Installation Guide

This guide provides procedures either to install Oracle Clusterware as a standalone product, or to install Oracle Clusterware with either Oracle Database, or Oracle RAC. It contains system configuration instructions that require system administrator privileges.

Oracle Real Application Clusters Installation Guide

The guide that you are reading provides procedures to install Oracle RAC after you have successfully completed an Oracle Clusterware installation. It contains database configuration instructions for database administrators.

Oracle Database Storage Administrator's Guide

This guide provides information for database and storage administrators who administer and manage storage, or who configure and administer Oracle ASM.

Oracle Clusterware Administration and Deployment Guide

This is the administrator's reference for Oracle Clusterware. It contains information about administrative tasks, including those that involve changes to operating system configurations.

Oracle Real Application Clusters Administration and Deployment Guide

This is the administrator's reference for Oracle RAC. It contains information about administrative tasks. These tasks include database cloning, node addition and deletion, Oracle Cluster Registry (OCR) administration, use of Server Control (SRVCTL) and other database administration utilities.

Changes in the Installation Options

The following are installation option changes for Oracle Database 11g release 1 (11.1):

- **Oracle Application Express:** This feature is installed with Oracle Database 11g. It was previously named HTML DB, and was available as a separate Companion CD component.
- **Oracle Configuration Manager:** Oracle Configuration Manager (OCM) is integrated with OUI. However, it is an optional component with database and client installations, and you must select Custom Installation to enable it. Oracle Configuration Manager, used in previous releases as customer configuration repository (CCR), is a tool that gathers and stores details relating to the configuration of the software stored in the Oracle ASM and Oracle Database home directories.

See ["Oracle Configuration Manager for Improved Support"](#) on page 3-xxx for further information.
- **Oracle Data Mining:** The Enterprise Edition installation type selects Oracle Data Mining option for installation by default.
- **Oracle Database Vault:** This feature is offered during installation. It is an optional component for database installation, available through Custom installation.
- **Oracle SQL Developer:** This feature is installed by default with template-based database installations, such as General Purpose, Transaction Processing, and Data

Warehousing. It is also installed with database client Administrator, Run-Time, and Custom installations.

- **Oracle Warehouse Builder:** This information integration tool is now installed with both Standard and Enterprise Edition versions of Oracle Database. With Enterprise Edition, you can purchase additional extension processes. Installing Oracle Database also installs a previously seeded repository, OWBSYS, necessary for using Oracle Warehouse Builder.
- **Oracle XML DB:** Starting with Oracle Database 11g, Oracle XML DB is no longer an optional feature. It is installed and configured using Oracle Database Configuration Assistant (DBCA) for all database installations.

New Components Available for Installation

The following are the new components available while installing Oracle Database 11g:

- **Oracle Application Express:** Starting with Oracle Database 11g, HTML DB is no longer available as a Companion CD component. Renamed Oracle Application Express, this component is installed with Oracle Database 11g.
With Oracle Database 11g, Oracle Application Express replaces *iSQL*Plus*.
- **Oracle Configuration Manager:** This feature is offered during custom installation. It was previously named Customer Configuration repository (CCR). It is an optional component for database and client installations. Oracle Configuration Manager gathers and stores details relating to the configuration of the software stored in Oracle Database home directories.
- **Oracle SQL Developer:** This feature is installed by default with template-based database installations, such as General Purpose, Transaction Processing, and Data Warehousing. It is also installed with database client Administrator, Run-Time, and Custom installations.
- **Oracle Warehouse Builder:** This feature is now included as an option in the Oracle Database installation.
- **Oracle Real Application Testing:** This feature is installed by default with the Enterprise Edition installation type of Oracle Database 11g.

See Also: *Oracle Database Performance Tuning Guide* for more information about Oracle Real Application Testing

Enhancements and New Features for Installation

The following are the enhancements and new features for Oracle Database 11g release 1 (11.1).

Automatic Diagnostic Repository

Automatic Diagnostic Repository (ADR) is a feature added to Oracle Database 11g. The main objective of this feature is to reduce the time required to resolve bugs. ADR is the layer of the Diagnostic Framework implemented in Oracle Database 11g that stores diagnostic data and also provides service APIs to access data. The default directory that stores the diagnostic data is `ORACLE_BASE\diag`.

Automatic Diagnostic Repository implements the following:

- Diagnostic data for all Oracle products which is written into an on-disk repository
- Interfaces that provide easy navigation of the repository and the capability to read and write data

For Oracle RAC installations, if you use a shared Oracle Database home, then ADR must be located on a shared storage location available to all the nodes.

Oracle Clusterware continues to place diagnostic data in the directory `CRS_home\log`, where `CRS_home` is the Oracle Clusterware home.

Oracle Automatic Storage Management Fast Mirror Resync

Oracle ASM fast mirror resync quickly resynchronizes Oracle ASM disks within a disk group after transient disk path failures if the disk drive media is not corrupted. Any failures that render a failure group temporarily unavailable are considered transient failures. Disk path malfunctions, such as cable disconnections, host bus adapter or controller failures, or disk power supply interruptions, can cause transient failures. The duration of a fast mirror resync depends on the duration of the outage. The duration of a resynchronization is typically much shorter than the amount of time required to completely rebuild an entire Oracle ASM disk group.

See Also: *Oracle Automatic Storage Management Administrator's Guide*

Oracle ASM and Oracle RAC Deinstallation and Other Configuration Assistant Enhancements

Oracle Database Configuration Assistant (DBCA), Database Upgrade Assistant (DBUA), and Oracle Net Configuration Assistant (NETCA) have been improved. These improvements include the following:

DBCA

- Provides a command-line feature, `deleteASM`, that removes Oracle ASM instances.
- Provides the option to switch from a database configured for Oracle Enterprise Manager Database Control to Oracle Enterprise Manager Grid Control.

DBUA

- Includes an improved pre-upgrade script to provide space estimation, initialization parameters, statistics gathering, and new warnings. DBUA also provides upgrades from Oracle Database releases 9.0, 9.2, 10.1, and 10.2.
- Supports in-place patch set upgrades.
- Starts any services that were running before the upgrade.

New SYSASM Privilege for Oracle ASM Administration

This feature introduces a new SYSASM privilege that is specifically intended for performing Oracle ASM administration tasks. Using the SYSASM privilege instead of the SYSDBA privilege provides a clearer division of responsibility between Oracle ASM administration and database administration.

Oracle ASM Preferred Read Disk Groups

In previous releases, Oracle ASM used the disk with the primary copy of a mirrored extent as the preferred disk for data read operations. With this release, using the new initialization file parameter `asm_preferred_read_failure_groups`, you can specify disks located near a specific cluster node as the preferred disks from which that node obtains mirrored data. This option is presented in DBCA, or you can configure it after installation. This change facilitates faster processing of data with widely distributed shared storage systems or with extended clusters (clusters whose nodes are geographically dispersed), and improves disaster recovery preparedness.

Oracle Automatic Storage Management Rolling Migration

Rolling migration for Oracle ASM enables you to upgrade or patch Oracle ASM instances on clustered Oracle ASM nodes without affecting database availability. Rolling migration provides greater availability and a smoother migration of Oracle ASM software from one release to the next. This feature applies to Oracle ASM configurations that run on Oracle Database 11g release 1 (11.1) and later. In other words, you must have Oracle Database 11g release 1 (11.1) installed before you can perform rolling migrations.

Note: You cannot change the owner of the Oracle ASM or Oracle Database home during an upgrade. You must use the same Oracle software owner that owns the existing Oracle ASM or Oracle Database home.

See Also: *Oracle Automatic Storage Management Administrator's Guide*

Conversion of Single-Instance Oracle ASM to Clustered Oracle ASM

Using either Oracle Enterprise Manager Grid Control or the `rconfig` script, you can convert an existing Oracle ASM instance from a single-instance storage manager to a clustered storage manager. You can convert Oracle ASM release 11.1 instances directly, and convert releases prior to 11.1 by upgrading the instance to 11.1, and then performing the conversion.

Data Mining Schema Creation Option

In Oracle Database 11g, the data mining schema is created when you run the SQL script `catproc.sql` as the SYS user. Therefore, the data mining option is removed from the Database Features screen of DBCA.

Oracle Disk Manager Network File System Management

Oracle Disk Manager (ODM) can manage a network file system (NFS) on its own, without using the operating system kernel NFS driver. This is referred to as Direct NFS. Direct NFS implements the NFS version 3 protocol within the Oracle Database kernel. This change enables monitoring of NFS status using the ODM interface. The Oracle Database kernel driver tunes itself to obtain optimal use of available resources.

This feature provides the following:

- Ease of tuning, and diagnosability, by giving the Direct NFS client control over the I/O paths to the network file system, and avoiding the need to tune network performance at the operating system level.
- A highly stable, highly optimized NFS client for database operations.
- Use of the Oracle buffer cache, rather than the file system cache, for simpler tuning.
- A common, consistent NFS interface, capable for use across Linux, UNIX and Windows platforms.
- No requirement for additional configuration of NFS mounts.

With operating system NFS drivers, NFS drives have to be mounted with the option `noac` (No Attribute Caching) to prevent the operating system NFS driver from optimizing the file system cache (by keeping file attributes locally). ODM automatically recognizes Oracle RAC instances, and performs appropriate

operations for data files without requiring additional reconfiguration from users, system administrators, or DBAs.

Optimal Flexible Architecture (OFA) Simplified

With the development of stripe and mirror everything (SAME) architecture, and improved storage and throughput capacity for storage devices, the original OFA role of enhancing performance has shifted to a role of providing well-organized Oracle installations with separated software, configuration files, and data. This separation enhances security, and simplifies upgrading, cloning, and other administrative tasks.

Oracle Database 11g release 1 (11.1) incorporates several improvements to OFA to address this changed purpose.

As part of this shift in roles, the following features have been added:

- During Oracle RAC installation, you are prompted to accept the default, or select a location for the Oracle base directory, instead of the Oracle home directory. This change facilitates installation of multiple Oracle home directories in a common location, and separates software units for simplified administration. For this release, you are not required to use the Oracle base directory, but this may become a requirement in a future release.
- With this release, as part of the implementation of Automatic Diagnostic Repository (ADR), the following admin directories are changed:
 - `bdump` (location set by the `background_dump_dest` initialization parameter; storage of Oracle background process trace files)
 - `cdump` (location set by the `core_dump_dest` initialization parameter; storage of Oracle core dump files)
 - `udump` (location set by the `user_dump_dest` initialization parameter; storage of Oracle user SQL trace files)

By default, the location of these trace and core files is in the `\diag` directory, which is in the path `ORACLE_BASE\diag`.

- The initialization parameters `background_dump_dest` and `user_dump_dest` are deprecated. They continue to be set, but you should not set these parameters manually.
- A new initialization parameter is introduced. `diagnostic_dest` contains the location of the ADR base directory, which is the directory under which one or more ADR homes are kept. Oracle documentation commonly refers to these homes as ADR homes. Each database instance has an ADR home, which is the root directory for several other directories that contain trace files, the alert log, health monitor reports, and dumps for critical errors. You can also view alert and trace files with the SQL statement `select name, value from v$diag_info`.
- The default fast recovery area (previously known as the Flash Recovery area) is moved from `ORACLE_HOME\.. \flash_recovery_area` to `ORACLE_BASE\flash_recovery_area`.
- The default data file location is moved from `ORACLE_HOME\.. \oradata` to `ORACLE_BASE\oradata`.
- A new utility, ADR Command Interpreter (ADRCI), is introduced. ADRCI facilitates reviewing alert log and trace files.

For Oracle RAC installations, Oracle requires that the fast recovery area and the data file location are on a location shared among all the nodes. Oracle Universal Installer

(OUI) confirms that this is the case during installation. This change does not affect the location of trace files for Oracle Clusterware.

See Also: *Oracle Database Administrator's Guide* for detailed information about these changes, and *Oracle Database Utilities* for information about viewing alert log and list trace files with ADRCI

Oracle Configuration Manager for Improved Support

During a custom installation, you are asked if you want to install Oracle Configuration Manager (OCM). OCM is an optional tool that enables you to associate your configuration information with your My Oracle Support account. This can facilitate handling of service requests by ensuring that server system information is readily available.

Configuring the OCM tool requires that you have the following information from your service agreement:

- Customer Support Identification (CSI) Number
- My Oracle Support user account name
- Country code

In addition, you are prompted for server proxy information if the host system does not have a direct connection to the Internet.

Support for Large Data Files

Large data file support is an automated feature that enables Oracle to support larger files on Oracle ASM more efficiently and to increase the maximum file size.

See Also: *Oracle Automatic Storage Management Administrator's Guide*

Switching a Database from Database Control to Grid Control Configuration

In previous releases, Oracle Database Configuration Assistant (DBCA) contained the functionality to configure databases while creating them either with Database Control or with Grid Control, or to reconfigure databases after creation. However, to change the configuration from Database Control to Grid Control required significant work. With Oracle Database 11g, DBCA enables you to switch configuration of a database from Database Control to Grid Control by running the Oracle Enterprise Manager Configuration Plug-in.

New Features in Oracle Real Application Clusters Release 11.1

Oracle Real Application Clusters release 11.1 includes the following:

- ODP .NET configuration improvements:
 - Developers can now configure ODP .NET using configuration files, including application or web config, or `machine.config`.
 - The settings for specific versions of ODP.NET can be configured several ways for specific effects on precedence. For example, `machine.config` settings are .NET framework-wide settings that override the Windows registry values. The application or web config file settings are application-specific settings that override the `machine.config` settings and the Windows registry settings.
- Performance enhancements, such as the following:
 - Improved parameter context caching

This release enhances the existing caching infrastructure to cache ODP .NET parameter contexts. This enhancement is independent of database version and it is available for all the supported database versions. This feature provides significant performance improvement for the applications that execute the same statement repeatedly.

- Efficient large object (LOB) retrieval

This release improves the performance of small-sized LOB retrieval by reducing the number of round-trips to the database. This enhancement is available only with Oracle 11g release 1 or later database releases.

This enhancement is transparent to the developer. No code changes are needed to use this feature.

Deprecated Features in Oracle Database 11g Release 2 (11.2)

The following are deprecated with Oracle Database 11g Release 2:

- Oracle Cluster File System (OCFS) for Windows
- Installing files on raw devices is no longer an option available during installation. You must use a shared file system, or use Oracle ASM. If you are upgrading from a previous release and currently use raw devices, then your existing raw devices can continue to be used. After upgrade is complete, you can migrate to Oracle ASM or to a shared file system if you choose.

Deprecated Features in Oracle Database 11g Release 1 (11.1)

The following components that were part of Oracle Database 10g release 2 (10.2) are not available for installation with Oracle Database 11g:

- *iSQL*Plus*
- Oracle Workflow
- Data Mining Scoring Engine
- Oracle Enterprise Manager Java console

Desupported Features in Oracle Database 11g Release 2 (11.2)

The following features are no longer supported with Oracle Database 11g Release 2 (11.2):

- The SYSDBA privilege of acting as administrator on the Oracle ASM instance is removed with this release.
- The `-cleanupOBASE` flag of the deinstallation tool is desupported. There is no replacement for this flag.

Preinstallation Checklist

Most Oracle Real Application Clusters (Oracle RAC) installation errors are due to a failure to complete all the required steps before starting Oracle Universal Installer (OUI).

This chapter assists database administrators and the system and storage administrators with coordinating the installation and planning the installation tasks in preparation for a complete installation of Oracle RAC.

This chapter contains the following topics:

- [Before You Start: Planning Your Installation](#)
- [Overview of Server Hardware, Network, and Operating System](#)
- [Overview of Oracle RAC Storage](#)
- [Requirements for the Oracle Database Vault Option](#)
- [Additional Information to Prepare for Installation](#)
- [Confirming Cluster Readiness Using Cluster Verification Utility](#)
- [Installing Oracle Database 11g Release 2 on Systems with Existing Databases](#)

1.1 Before You Start: Planning Your Installation

Oracle recommends that you complete the following list of tasks before you start an Oracle RAC installation. Whether your location is a Tier IV data center with a large project team of system, storage, network, and database administrators and third-party hardware and software vendors, or you are a project team of one, planning is important to help ensure that your installation proceeds smoothly.

It is beyond the scope of this documentation set to advise how to determine hardware sizing or capacity planning for your installation. Note that with Oracle Clusterware and Oracle RAC, you can add nodes and instances as needed in response to testing, or in response to increased workloads.

Review and complete the following steps as part of your installation plan:

- [Check My Oracle Support](#)
- [Consider Using Software Updates Option](#)
- [Check Oracle Technology Network](#)
- [Review Installation Steps for Oracle RAC](#)
- [Review How to Use Cluster Verification Utility \(CVU\)](#)
- [Review and Back Up Existing Oracle Installations](#)

- [Review Globalization Requirements](#)
- [Review Documentation](#)

1.1.1 Check My Oracle Support

Before you decide whether you want to install Oracle Database 11g Release 2 (11.2) on existing hardware or purchase additional server and storage hardware for an installation, log on to My Oracle Support:

<https://support.oracle.com>

My Oracle Support contains guidance about supported hardware options that can assist you with your purchasing decisions and installation planning. The My Oracle Support certifications page contains more detailed information about certified hardware and has information specific to each release and platform. My Oracle Support is available at the following URL:

<http://support.oracle.com/>

You must register online before using My Oracle Support. Use the following steps to locate the certification information for your platform:

1. After logging in, click the **Certifications** tab.
2. In Certification Search box, enter information about your software and hardware. For example, you might use the following:
 - Product: **Oracle Database**
 - Release: **11.2.0.3**
 - Platform: **Microsoft Windows x64 (64-bit) 2008 R2**

After entering the required information, press **Search**. A new window appears with the product certification details.

3. On the Certification Details page, click the links to view certification information and ensure that your hardware configuration is supported for use with Oracle Clusterware and Oracle RAC.

Note: Contact your Oracle sales representative if you do not have a My Oracle Support account.

Also, you may want to visit Oracle.com (<http://www.oracle.com>) for additional resources about planning for specific implementation scenarios, best practices, and other information that can help you with your installation plan. In particular, see the Oracle Real Application Clusters pages on the Oracle Technology Network:

<http://www.oracle.com/technetwork/database/clustering/overview/>

The Oracle RAC Technologies Certification Matrix provides details about supported technology that you can use to deploy Oracle RAC.

1.1.2 Consider Using Software Updates Option

OUI can install critical patch updates, system requirement updates for supported operating systems, and other significant updates that can ensure that your installation succeeds. Oracle recommends that you enable the Software Updates option during installation.

If you choose to enable the Software Updates option, then during installation you must provide a valid My Oracle Support user name and password, so that OUI can download the latest updates, or you must provide a path to the location of a Software Updates option package that you have downloaded previously.

In addition, if you have a proxy realm, then be prepared to provide the destination proxy realm, and the user authentication (user name and password) required to authenticate access through that realm to the Internet, so that OUI can obtain software updates. Check with your network administrator before installation if you do not have this information.

If you plan to run the installation in a secured data center, then you can download updates before starting the installation by starting OUI on a system that has Internet access. To start OUI to download updates only, enter the following command:

```
setup.exe -downloadUpdates
```

Provide the My Oracle Support user name and password, and provide proxy settings if needed. After you download updates, transfer the update file to a directory on the server where you plan to run the installation.

1.1.3 Check Oracle Technology Network

Oracle Technology Network (OTN) contains white papers about deployment options, capacity planning, best practices, and extended clusters deployments, which are not addressed in this guide. You can review available papers at the following web site:

<http://www.oracle.com/technetwork/database/clustering/overview/>

1.1.4 Review Installation Steps for Oracle RAC

Installing Oracle RAC consists of the following steps:

1. Prepare servers (system, storage, and network administration):
 - Install the operating system, operating system packages, and patches to the required version.
 - Create the required groups, users, and software homes.
 - Set up the domain name forwarding for Grid Naming Service (GNS) if you plan to deploy GNS, and set up the network addresses in the domain name service (DNS) and on the server as needed.
 - Set up the required storage.
 - (optional) Stage all the software on one node for installation (also called the *local node*).

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about using OIFCFG to modify interconnects

2. Install Oracle Grid Infrastructure for a cluster, which includes Oracle Clusterware and Oracle Automatic Storage Management (system and storage administration):
 - Install Oracle Grid Infrastructure for a cluster.
 - Patch Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) to the latest patchset.
3. Install Oracle RAC (database administration):

- Install Oracle RAC.
- Patch Oracle RAC to the latest patchset.
- Complete the postinstallation configuration of the Oracle RAC database.

1.1.5 Review How to Use Cluster Verification Utility (CVU)

Oracle provides Cluster Verification Utility (CVU) to perform system checks in preparation for an installation, patch updates, or other system changes.

Before installing Oracle Clusterware or Oracle RAC, use CVU to ensure that your cluster is prepared for an installation. CVU is incorporated into OUI, so CVU performs its checks when you start an Oracle RAC installation. However, you can use CVU to ensure that any packages or configuration information that is required for Oracle RAC are in place before you begin your Oracle RAC installation.

Using CVU can help system, storage, and database administrators ensure that all system configuration and preinstallation steps are completed, so that installations, updates, or patches also complete successfully. You can obtain the latest version of CVU at the following URL:

<http://www.oracle.com/technetwork/database/options/clustering/downloads/cvu-download-homepage-099973.html>

If you have vendors performing hardware or operating system configuration steps, then ask the vendors to complete the relevant CVU checks of the cluster after they complete their work to ensure that your system is configured correctly. Database administrators should review "[Confirming Cluster Readiness Using Cluster Verification Utility](#)" on page 1-18 to confirm that their system is prepared for installation before they start an Oracle RAC installation.

1.1.6 Review and Back Up Existing Oracle Installations

If you have an existing Oracle installation, then document version numbers, patches, and other configuration information, and review upgrade procedures for your existing installation. Review the Oracle upgrade documentation before proceeding with your installation, to decide how you want to proceed.

As with any system change, back up your existing database before attempting to install new software.

Note that to install Oracle RAC 11g Release 2 (11.2), you must have Oracle Clusterware and Oracle ASM 11g Release 2 (11.2) installed on your cluster. The Oracle Clusterware version must be equal to or more recent than the Oracle RAC version that you plan to install.

See Also: *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)*

For the most current information and best practices about pre-upgrade, post-upgrade, compatibility, and interoperability discussions, see "Oracle Upgrade Companion", which is available through Note 785351.1 on My Oracle Support:

<https://support.oracle.com/>

For upgrades, note the following:

- You can have only one version of Oracle Clusterware running on a cluster at a time. The Oracle Clusterware version must be the most recent release of any software (Oracle Clusterware, Oracle Database, Oracle RAC, and Oracle ASM)

running on the cluster. Before upgrading your database to Oracle Database 11g release 2, all nodes in the cluster must be upgraded to Oracle Clusterware and Oracle ASM 11g Release 2 (11.2).

- You can have multiple Oracle homes for Oracle databases on your cluster. However, the Oracle RAC database software in these homes must be from a version that is equal to or prior to the version of Oracle Clusterware that is installed; you cannot have a version of Oracle Database running on Oracle Clusterware that was released after the version of Oracle Clusterware that you are running.

For example:

- If you have Oracle Clusterware 11g release 2 installed as your clusterware, then you can have an Oracle Database 10g release 2 single-instance database running on one node, and separate Oracle RAC 10g release 1, release 2, and Oracle RAC 11g release 1 or release 2 databases also running on the cluster.
- You cannot have Oracle Clusterware 10g release 2 installed on your cluster, and install Oracle RAC 11g.
- Starting with releases 10.1.0.6 and 10.2.0.3, you can use Database Upgrade Assistant (DBUA) for patch set upgrades with Oracle RAC. You can also use DBUA to upgrade between major point releases of Oracle RAC (for example, from 10.2 to 11g, or 11.1 to 11.2).
- You cannot change the owner of the Oracle Database home during an upgrade. You must use the same Oracle software owner that owns the existing Oracle Database home.
- For the Oracle RAC database being upgraded, each cluster member node that hosts an instance of the database must be a member of the upgraded installation. For example, if you have an existing Oracle RAC database running on a three-node cluster, then you must upgrade Oracle RAC on all three nodes. You cannot upgrade only two nodes of the cluster, and then remove the third node and instance in the upgrade.

Note: If you are upgrading an existing Oracle RAC installation, then you must use the same type of Oracle home that you have in your existing installation. For example, if you have a shared Oracle home in your existing installation, then you must upgrade to a shared Oracle home with Oracle RAC 11g Release 2 (11.2).

Similarly, if you have local Oracle homes on cluster nodes, then you must upgrade to local Oracle homes on cluster nodes.

See Also: *Oracle Database Upgrade Guide* for information about how to prepare for upgrading an existing database

1.1.7 Review Globalization Requirements

Before you start an installation on which you want to support languages other than English, review *Oracle Database Globalization Support Guide*.

Note the following:

- Oracle recommends that you use Unicode AL32UTF8 as the database character set.

Unicode is the universal character set that supports most of the currently spoken languages. Unicode also supports many historical scripts (alphabets). Unicode is the native encoding of many technologies, including Java, XML, XHTML, ECMAScript, and LDAP. Unicode is ideally suited for databases connected to the Internet and the global economy.

- The locale setting of your operating system session determines the language of the user interface and the globalization behavior for components such as OUI, Oracle Net Configuration Assistant (NETCA), and Oracle Database Configuration Assistant (DBCA). The locale setting also determines the globalization behavior of Oracle Database sessions created by a user application through the JDBC driver, unless overridden by the application.
- The character set specified by the `NLS_LANG` environment variable determines the language of the user interface and the globalization behavior for components such as SQL*Plus, `exp`, and `imp`. This variable also sets the language and territory used by the client application and the database. The `NLS_LANG` environment variable also declares the character set for entering and displaying data by the client application.

The `NLS_LANG` environment variable should reflect the setting of the operating system character set of the client. For example, if the database character set is `AL32UTF8` and the client is running on a Windows operating system, then you should not set the `NLS_LANG` environment variable to `AL32UTF8` because there are no UTF-8 WIN32 clients. Instead, the `NLS_LANG` setting should reflect the code page of the client. For example, on an English Windows client, the code page is 1252, so an appropriate setting for `NLS_LANG` is `AMERICAN_AMERICA.WE8MSWIN1252`.

Note: *Oracle Database Installation Guide* for your platform contains a more detailed discussion of database character sets used with different languages, and provides more information about installing and configuring Oracle Database globalization support.

1.1.8 Review Documentation

Before installing Oracle RAC, review the *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)*, particularly the preinstallation, and the storage chapters, to complete all the required steps for an installation.

Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit) also contains most tasks requiring Administrator privileges or storage administrator privileges that must be completed before starting an Oracle RAC installation.

In addition, review the Release Notes and My Oracle Support (<https://support.oracle.com>) to ensure that you have the most current information about system requirements and other information that can affect your installation. The small amount of time that this review takes can save a much greater amount of time required to track down causes of installation errors later. Also, verify that you have the most current version of this document; Oracle documentation is updated after the initial release and posted to the following web site:

<http://docs.oracle.com/>

Oracle recommends that you install a web browser on at least one of your cluster nodes to enable Oracle Enterprise Manager and Oracle Application Express for Oracle RAC, and to access online documentation as needed. Online documentation is available in PDF and HTML formats and can be downloaded.

See Also:

- Oracle Database Concepts for an overview of Oracle Database
- *Oracle Real Application Clusters Administration and Deployment Guide* for additional information about Oracle Clusterware or Oracle RAC configuration and deployment
- *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)* for server and storage configuration information for Oracle RAC

1.2 Overview of Server Hardware, Network, and Operating System

Oracle Clusterware must be installed successfully as part of an Oracle Grid Infrastructure for a cluster installation before attempting to install Oracle RAC. To complete installations successfully, ensure that required hardware, network, and operating system preinstallation steps for Oracle software are performed as required. Failure to complete the required preinstallation steps is the most common reason for unsuccessful installations.

Before Oracle Clusterware is installed as part of an Oracle Grid Infrastructure for a cluster installation, you must have completed installing and configuring CPUs, memory, shared storage, local disks, network cards, host bus adapters, interconnects, and any other networking or server hardware. You should also have installed the operating system, and any required packages or third-party software. Review your vendor documentation to complete these tasks, and if relevant, work with your vendor to complete the Oracle preinstallation steps that are listed in this document to confirm that the vendor hardware and software is correctly configured.

Server and network preparation for installation includes the following:

- [Review Server Hardware and Software Requirements](#)
- [Overview of Server Environment Configuration](#)
- [Overview of Oracle Base](#)
- [Overview of Network and Internet Configuration](#)
- [Overview of Platform-Specific Server Configuration](#)

1.2.1 Review Server Hardware and Software Requirements

The following summary of server hardware and software configuration requirements and recommendations will enable you to prepare for a successful installation of Oracle RAC.

Required for Oracle RAC

Each node in a cluster requires the following:

- Supported server hardware, including processors and system configuration.

Review My Oracle Support before starting an installation on your current hardware and before purchasing new hardware to ensure that the hardware is supported by Oracle Clusterware with Oracle RAC 11g Release 2 (11.2).

Also review the preinstallation chapters in *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)* for more details about the supported configurations.

Note: You must use the same operating system on each node in the cluster. Oracle strongly recommends that you use the same software configurations on each node of the cluster. Oracle Clusterware and Oracle RAC do not support heterogeneous platforms (each server must run the same Oracle software binaries) in the same cluster.

- Operating system updates listed in the system requirements.
Oracle may require updates to the operating system, such as service packs and individual patches. If such requirements exist, then they are stated in the Release Notes for a particular release. You can also apply other operating system patches as recommended by Microsoft, if there are no "certification exceptions" listed in the Release Notes.
- A supported interconnect software protocol on each node, to support Oracle Clusterware voting disk polling, and to support Cache Fusion with Oracle RAC. Your interconnect *must* be certified by Oracle for your platform.

Recommended for Oracle RAC: System Administrator Tasks

Oracle recommends the following tasks to simplify server installation and maintenance, and to prevent service issues:

- Utilizing a time protocol, such as Network Time Protocol (NTP), to ensure that all nodes in the cluster use the same reference time. With Oracle Clusterware 11g Release 2 (11.2), if neither the Network Time Protocol or Windows Time Service is enabled at installation, then the Oracle Clusterware installation enables the Cluster Time Synchronization Service (CTSS).
- Configuring redundant switches, for all cluster sizes.
- Using identical server hardware on each node, to simplify server maintenance.
- Avoiding resource contention issues by not installing Oracle RAC on a primary domain controller or backup domain controller.

Additional Options: System Administrators and Vendors

After you have set up server hardware, review "Checking the Hardware Requirements" in *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)* to ensure that your system has enough RAM, that the `TEMP` environment variable points to a location that has enough available space for the installation, and that your system meets other hardware requirements.

1.2.2 Overview of Server Environment Configuration

Configure the users and user environments as described in the preinstallation chapters of *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)*. These include the following tasks:

- Creating operating system users to install Oracle software
- Configuring the Oracle software owner user environments

See Also: *Oracle Database Platform Guide for Microsoft Windows* for information about creating user groups

To install the Oracle software, you must be a user that is a member of the Administrators group. If you use a local user account for the installation, then the user

account must exist on all nodes in the cluster and the user name and password must be the same on all nodes.

If you use a domain account for the installation, then the domain user must be explicitly declared as a member of the local Administrators group on each node in the cluster. It is not sufficient if the domain user has inherited membership from another group. The user performing the installation must be in the same domain on each node. For example, you cannot have a dba1 user on the first node in the DBADMIN domain and a dba1 user on the second node in the RACDBA domain.

Assume that you have one Oracle installation owner, and the user name for this Oracle installation owner is oracle. The oracle user must be either a local Administrator user or a domain user, and the same user must exist (same user name, password, and domain) on each node in the cluster.

When installing the Oracle Database software, the user performing the installation is automatically added to the ORA_DBA group. If you use a domain user, then after installation you must ensure that the domain user on each node is a member of the local Administrator's group and the ORA_DBA group.

Members of the local ORA_DBA group automatically receive the SYSDBA and SYSASM privileges.

Note: If you use a domain user to perform the installation, then in some cases you must manually add the user to ORA_DBA group of the remote nodes after installation completes.

Oracle strongly recommends the use of domain users for Oracle RAC installations.

See Also: *Oracle Database Installation Guide for Microsoft Windows* for more information about the ORA_DBA group

1.2.3 Overview of Oracle Base

The Oracle base directory is the location where Oracle software and configuration files are stored. By default, the Oracle software binary files are installed by version and Oracle home name in a subdirectory of the Oracle base directory. An Oracle base directory can be used for multiple installations.

Every Oracle home can have its own Oracle base. Typically, the Oracle base path is based on the installation owner name. If you have separate installation owners for the Oracle Grid Infrastructure installation and the Oracle RAC installation, then you have two Oracle base paths that are in accordance with Optimal Flexible Architecture (OFA) guidelines. For example, if the user grid is the installation owner of the Oracle Grid infrastructure installation and the user oracle is the installation owner of the Oracle Database installation, then you have two Oracle base directories:

- X:\app\grid—the Oracle base for the user grid, which owns the Oracle Grid Infrastructure binaries
- X:\app\oracle—the Oracle base for the user oracle, which owns the Oracle Database binaries

In preparation for installation, Oracle recommends that you only set the ORACLE_BASE environment variable to define paths for Oracle binaries and configuration files. Oracle Universal Installer (OUI) creates other paths and environment variables as

necessary, in accordance with the Optimal Flexible Architecture (OFA) rules for well-structured Oracle software environments.

For example, with Oracle Database 11g, Oracle recommends that you do not set an Oracle home environment variable, and instead enable OUI to create it. If the Oracle base path is `E:\app\oracle`, then by default, OUI creates the following Oracle home path:

```
X:\app\oracle\product\11.2.0\dbhome_n
```

The variable *n* is the Oracle home number. The first time you create an Oracle home as the oracle user, the default Oracle home location is

```
C:\app\oracle\product\11.2.0\dbhome_1.
```

Ensure that the paths that you select for Oracle software, such as Oracle home paths and the Oracle base path, use only ASCII characters. Because installation user names are used by default for some path, this ASCII character restriction applies to user names, file names, and directory names.

1.2.4 Overview of Network and Internet Configuration

You must configure each server so it can communicate with the other nodes in the cluster. Also, you should configure one or more nodes to access the internet and support a web browser.

1.2.4.1 Overview of Network Interface Configuration

During the installation of Oracle Grid Infrastructure for a cluster, the person completing the installation identifies the planned use for each global interface, identifying each interface as a *Public* interface type (used with public IP addresses and virtual IP addresses), a *Private* interface type (used with interconnects between cluster member nodes), or a *Do not use* interface type, which Oracle Clusterware and Oracle RAC should ignore. For example, an interface used as a dedicated interface for a network file system should be marked as a *Do not use* interface type.

Additional network configuration is not required during an Oracle RAC configuration.

1.2.4.2 Overview of Configuring Internode Communication

The Oracle software is installed on one node in the cluster and then OUI copies the software to the other nodes in the cluster. Verify that each node in your cluster can communicate with the other nodes using the `net use` command, for example, on `node1` you would use the following command:

```
C:\> net use \\node2\c$  
The command completed successfully.
```

On `node2`, you would use the following command:

```
C:\> net use \\node1\c$  
The command completed successfully.
```

1.2.4.3 Overview of Browser Requirements

You need a web browser to access documentation, and to use Oracle Enterprise Manager and Oracle Application Express. Web browsers must support JavaScript and the HTML 4.0 and Cascading Style Sheets (CSS) 1.0 standards. The following browsers meet these requirements:

- Microsoft Internet Explorer 10 (supports Oracle Enterprise Manager Database Control 11.2.0.3 or higher)
- Microsoft Internet Explorer 9.0
- Microsoft Internet Explorer 8.0
- Microsoft Internet Explorer 7.0 SP1
- Microsoft Internet Explorer 6.0 SP2
- Firefox 21.0 (supports Oracle Enterprise Manager Database Control 11.2.0.4)
- Firefox 17.0.6 ESR (supports Oracle Enterprise Manager Database Control 11.2.0.4)
- Firefox 3.6
- Firefox 3.5
- Safari 4.0.x
- Safari 3.2
- Safari 3.1
- Google Chrome 27.0 (supports Oracle Enterprise Manager Database Control 11.2.0.4)
- Google Chrome 4.0
- Google Chrome 3.0

For a list of browsers supported with Oracle Application Express, see *Oracle Application Express Installation Guide*.

Note: You do not require a browser to install Oracle RAC.

1.2.5 Overview of Platform-Specific Server Configuration

Oracle Clusterware and Oracle RAC are tested with specific operating system versions, and specific operating system components. Oracle requires that you use the operating system versions and components that are certified for this release.

Oracle recommends that you or your system administrator review the system requirements carefully in *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)* before beginning installation, to ensure that your system meets these requirements. If your system does not meet minimum hardware, operating system version, and component requirements, then your installation may fail to complete, or other errors may develop during Oracle Clusterware or Oracle Database run time.

In addition to the standard system requirements configuration, deployment on specific server hardware can include additional operating system configuration steps. Review the preinstallation chapter in *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)*, and check the My Oracle Support Certify page to ensure that you are aware of any additional requirements or recommendations for your specific hardware and platform configuration.

1.3 Overview of Oracle RAC Storage

For both the Enterprise and Standard Editions of Oracle RAC, the hard disk requirements for Oracle Database components include 32 MB required to install Java

Runtime Environment (JRE) and OUI on the disk partition where the operating system is installed. If sufficient space is not detected, then the installation fails and an error message appears. [Table 1–1](#) lists the hard disk space requirements, including the requirements for a starter database.

Table 1–1 Hard Disk Space Requirements for Oracle RAC

Installation Type	TEMP space	%SYSTEMDRIVE%\Program Files\Oracle	Oracle Database Home	Data Files	Total
Standard Edition	500 MB	100 MB	5 GB	3 GB	8.6 GB
Enterprise Edition	500 MB	100 MB	5.5 GB	3 GB	9.1 GB

In [Table 1–1](#), the Data Files column refers to the contents of the `admin`, `flash_recovery_area`, and `oradata` directories in the Oracle base directory. The actual size of these files can be greater than what is specified in the table, depending on the installation options selected, such as languages or additional components. If you later plan to install Oracle Database and Oracle Real Application Clusters with automated backups enabled, then include at least an additional 2 GB for data file disk space.

To install Oracle RAC, you must configure shared storage for the database files and optionally the recovery files. If you instruct DBCA to configure automatic backups, the backups require a database recovery area that must be shared.

See Chapter 3 in *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)* to review storage options for installation planning. Storage and system administrators can see this chapter to configure storage for database files for an Oracle RAC database.

See Also: *Oracle Database Administrator's Guide* for information about recovery areas

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is a new multiplatform, scalable file system, and storage management technology that extends Oracle ASM functionality to support customer files maintained outside of the Oracle Database. Files supported by Oracle ACFS include database and application executable files, trace files, alert logs, application reports, BFILES, and configuration files. Oracle ACFS leverages Oracle ASM functionality to enable the following:

- Oracle ACFS dynamic file system resizing
- Maximized performance through direct access to Oracle ASM disk group storage
- Balanced distribution of Oracle ACFS across Oracle ASM disk group storage for increased I/O parallelism
- Data reliability through Oracle ASM mirroring protection mechanisms

The Oracle ASM Dynamic Volume Manager (Oracle ADVN) provides volume management services for disks and a standard disk device driver interface to clients (users and applications). File systems and other disk-based applications issue I/O requests to Oracle ADVN volume devices as they would to other storage devices on a vendor operating system. Oracle ADVN extends Oracle ASM by providing a disk driver interface to Oracle ASM storage allocated as Oracle ASM volume files. You can use Oracle ADVN to create virtual disks that contain file systems. These file systems contained on Oracle ASM volumes can support files beyond Oracle database files, such as executable files, report files, trace files, alert logs, and other application data files.

See Also :

- The Certify page on My Oracle Support for the most current information about storage options:
<https://support.oracle.com>
- *Oracle Automatic Storage Management Administrator's Guide* for an overview of storage configuration administration

1.3.1 General Storage Considerations for Oracle RAC

For all installations, you must choose the storage option to use for Oracle RAC database files. To enable automated backups during the installation, you must also choose the shared storage option to use for recovery files (the fast recovery area). You do not have to use the same storage option for each file type.

Use the following guidelines when choosing the storage options to use for each file type. You can choose any combination of the supported shared storage options for each file type if you satisfy all requirements listed for the chosen storage options

- Oracle recommends that you choose Oracle ASM as the shared storage option for database and recovery files.
- For Standard Edition Oracle RAC installations, Oracle ASM is the only supported shared storage option for database or recovery files.
- If you do not have a storage option that provides external file redundancy, then you must configure at least three voting disk areas to provide voting disk redundancy.
- If you intend to use Oracle ASM with Oracle RAC and you did not use Oracle ASM for the Oracle Clusterware files (Oracle Cluster Registry (OCR) and voting disk), then you must configure the Oracle ASM instance with Oracle ASM Configuration Assistant (ASMCA) before starting OUI. All nodes in the cluster must have the Oracle Database 11g release 2 version of Oracle Grid Infrastructure installed (Oracle Clusterware and Oracle ASM).
- If you intend to upgrade an existing Oracle RAC database, then you must ensure that your system meets the following conditions:
 - OUI and DBCA are run on the node where the Oracle RAC database instance is located.
 - The Oracle RAC database instance is running on the *same* nodes that you intend to make members of the new cluster installation. For example, if you have an existing Oracle RAC database running on a three-node cluster, then you must install the upgrade on all three nodes. You cannot upgrade only two nodes of the cluster, removing the third instance in the upgrade.

See Also: *Oracle Database Upgrade Guide* for information about how to prepare for upgrading an existing database

1.3.2 Overview of Oracle Database and Recovery File Options

There are two ways of storing Oracle Database and recovery files:

Note: When creating a new Oracle RAC database using OUI, storing data files and recovery files on raw devices (unformatted partitions) is not supported. Place data files and recovery files on a shared file system or on Oracle ASM. If you are upgrading a database that uses raw devices, they are still supported with Oracle Database 11g release 2.

- **Oracle Automatic Storage Management:** Oracle ASM is an integrated, high-performance database file system and disk manager for Oracle Database files. It performs striping and mirroring of database files automatically.

Note: For Standard Edition Oracle Database installations using Oracle RAC, Oracle ASM is the only supported storage option.

- **A supported shared file system:** Supported file systems include the following:
 - **Oracle Cluster File System (OCFS) for Windows:** Note that if you intend to use OCFS for Windows for your database files, then you should create partitions large enough for the all the database and recovery files when you create partitions for use by Oracle Database.

Note: To have a shared Oracle Database home directory for all nodes, you must use a supported cluster file system.

- **Direct Network File Systems (NFS):** You can configure Oracle Database to access NFS version 3 (NFSv3) servers directly using Direct NFS. Direct NFS is an optimized NFS client that provides faster and more scalable access to NFS storage located on NAS storage devices (accessible over TCP/IP). Direct NFS is built directly into the database kernel and provides faster performance than what can be provided by the operating system's NFS driver. When accessing NFS servers, Oracle Database bypasses the operating system and generates exactly the requests it needs (no user configuration or tuning required).

Table 1–2 shows the storage options supported for storing Oracle Database files and Oracle Database recovery files. Oracle Database files include data files, control files, redo log files, the server parameter file (SPFILE), and the password file.

Table 1–2 Supported Storage Options for Oracle Database and Recovery Files

Storage Option	File Types Supported	
	Database	Recovery
Oracle ASM	Yes	Yes
Oracle ASM Cluster File System (Oracle ACFS)	No	No
OCFS for Windows	Yes	Yes
Direct NFS	Yes	Yes
Local storage	No	No

Table 1–2 (Cont.) Supported Storage Options for Oracle Database and Recovery Files

Storage Option	File Types Supported	
	Database	Recovery
Shared unformatted partitions	Not supported by OUI or ASMCA, but supported by Oracle RAC. Raw partitions can be added or removed after installation.	Not supported by OUI or ASMCA, but supported by Oracle RAC. Raw partitions can be added or removed after installation.

1.4 Requirements for the Oracle Database Vault Option

Oracle Database Vault requires the following:

- An Oracle Database Enterprise Edition installation. Oracle Database Vault cannot be used with Oracle Database Standard Edition.
- The initialization parameter `db_block_size` for the database set to 4096 or higher. You can specify the initial value for this parameter during installation.
- Oracle Enterprise Manager Database Control.
- Oracle Clusterware installed and running.
- The password file authentication parameter, `REMOTE_LOGIN_PASSWORDFILE`, set to `EXCLUSIVE` or `SHARED` for the database.

You can specify the initial value for this parameter during installation, or in the Oracle Database initialization file. To create and manage password files, use the `orapwd` utility.

See Also: *Oracle Database Administrator's Guide* for more information about creating and maintaining a password file

You can install Oracle Database Vault more than once on the same system, in different Oracle homes for different databases.

1.5 Additional Information to Prepare for Installation

The following information about Oracle Clusterware, Oracle ASM, and Oracle RAC may be helpful to your installation planning team when deciding how you want to configure your installation. It contains the following topics:

- [Oracle ASM Is Installed with Oracle Grid Infrastructure](#)
- [Obtain the Oracle ASM ASMSNMP Password for Oracle Enterprise Manager](#)
- [Oracle Clusterware and Vendor Clusterware](#)
- [Shared Components of Oracle RAC Installations](#)
- [Simplified Upgrade of TIMESTAMP WITH TIME ZONE Data](#)
- [Oracle Database 11g Real Application Clusters Release Compatibility](#)

1.5.1 Oracle ASM Is Installed with Oracle Grid Infrastructure

In past releases, Oracle ASM was installed as part of the Oracle Database installation. With Oracle Database 11g release 2 (11.2), Oracle ASM is part of an Oracle Grid

Infrastructure installation. To upgrade an existing Oracle ASM installation, you must upgrade Oracle ASM by running an Oracle Grid Infrastructure upgrade.

If you did not select Oracle ASM as the storage option for the Oracle Clusterware files, then you must first use Oracle ASM Configuration Assistant (ASMCA) to configure Oracle ASM, create an Oracle ASM instance, and create a disk group to use for your Oracle Database storage.

1.5.2 Obtain the Oracle ASM ASMSNMP Password for Oracle Enterprise Manager

Obtain the password for the Oracle ASM user ASMSNMP from the Oracle Grid Infrastructure administrator.

During Oracle Grid Infrastructure for a cluster installation, the ASMSNMP account is created and its status is set to OPEN. It is granted the SYSDBA for ASM privilege. Oracle Enterprise Manager requires access to the ASMSNMP account to monitor Oracle ASM instances and retrieve data from ASM-related data dictionary views.

See Also: *Oracle Automatic Storage Management Administrator's Guide* for additional information about the ASMSNMP account

1.5.3 Oracle Clusterware and Vendor Clusterware

Oracle Clusterware provides clustering services. Currently, there are no supported clusterware products other than Oracle Clusterware for the Microsoft Windows platforms. If you intend to install Oracle RAC, then you must first install Oracle Clusterware.

1.5.4 Shared Components of Oracle RAC Installations

For Oracle RAC, you and your system administrator should note that all instances in Oracle RAC environments share the control file, SPFILE, redo log files, and all data files. These files must be placed on a shared file systems, and all the cluster database instances must have access to them. Each instance also has its own set of redo log files. During failures, shared access to redo log files enables surviving instances to perform recovery.

1.5.5 Simplified Upgrade of TIMESTAMP WITH TIME ZONE Data

As part of the installation of Oracle Database 11g Release 2 (11.2), time zone files from versions 1 to 11 are installed in the path *Oracle_home\oracore\zoneinfo*. You can continue to use the current time zone file or upgrade to the latest version. Oracle recommends that you upgrade the server to the latest version of the time zone file. Upgrading to a new version of the time zone file may cause existing TIMESTAMP WITH TIME ZONE (TSTZ) data to become stale. Using the newly provided DBMS_DST PL/SQL package, you can update the TSTZ data transparently, with minimal manual procedures and system downtime.

All instances of an Oracle RAC database must use the same time zone. The Oracle RAC database time zone defaults to the time zone setting of the Oracle Grid Infrastructure *grid* user, unless an instance is started with SQL*Plus. When you use SQL*Plus, you must be sure to use the same time zone setting for the database instance that is used for Oracle Clusterware. You can change the time zone Oracle Clusterware uses for a database by using the following command, where *time zone* is the time zone to which you want to change:

```
srvctl setenv database -T "TZ=time zone"
```

Time zone files are also installed on the client computers. Starting with Oracle Database 11g release 2, you no longer have to upgrade Oracle Client time zone files immediately. Upgrades can be done at a time when it is most convenient to the system administrator. However, there could be a small performance penalty when client and server use different time zone versions.

See Also:

- *Oracle Database Upgrade Guide* for information about preparing to upgrade TSTZ data
- *Oracle Database Globalization Support Guide* for information about how to upgrade the time zone file and TSTZ data
- *Oracle Call Interface Programmer's Guide* for information about performance effects of clients and servers operating with different versions of time zone files

1.5.6 Oracle Database 11g Real Application Clusters Release Compatibility

You can install and operate different releases of Oracle Database software on the same computer when all the following conditions are met:

- If you have Oracle Clusterware installed and different releases of other Oracle software installed, then the Oracle Clusterware release must be later than or equal to the Oracle Database software release. Oracle Clusterware and Oracle ASM are both upgraded to release 11.2.0.1 when you perform an Oracle Grid Infrastructure 11g Release 2 (11.2) installation.
- If you have an existing Oracle Database home, then you can create a new Oracle Database home and install Oracle Database 11g Release 2 (11.2) into the new Oracle home. You should ensure that Oracle Clusterware is in a separate Oracle Grid Infrastructure home (referred to as the Grid home). Oracle Grid Infrastructure installations cannot be installed in the Oracle base directory for Oracle Database.
- If you are running the Oracle⁹ⁱ release of Oracle RAC, and you want to continue to use that release, then you must run cluster software that is compatible with that release, such as Oracle Cluster Manager. Oracle Clusterware release 11g can be installed on the same system as Oracle⁹ⁱ Database software, but Oracle⁹ⁱ Database software cannot be supported by Oracle Clusterware 11g.
- If OUI detects a previous database release, then OUI asks you about your upgrade preferences. You have the option to upgrade a previous release database with DBUA or to create a new database using DBCA. The information collected during this OUI dialog is passed to DBUA or DBCA after the software is installed.

If OUI detects a previous Oracle Clusterware release, then you are asked to upgrade the existing Oracle Clusterware installation. Only one Oracle Clusterware version can be active on a server, and a server must be a member of only one cluster.

Note: Do not move Oracle binaries from the Oracle home to another location. Doing so can cause dynamic link failures.

- You cannot install Oracle Grid Infrastructure for a standalone server (Oracle Restart) and then install Oracle Grid Infrastructure for a cluster. If you have Oracle

Restart installed, then you must remove that installation before you can install Oracle Grid Infrastructure for a cluster.

Note: Do not move Oracle executable files from the directory in which they were installed to another location. To move the executable files to a new location, you must reinstall the software.

- You can run different releases of Oracle Database and Oracle ASM. For example, you can use Oracle ASM release 11g Release 2 (11.2) with an Oracle Database 10g release 2 (10.2) database.

Note: When using different Oracle ASM and Oracle Database releases, the functionality of each is dependent on the functionality of the earlier software release. For example, an Oracle Database 10g release 10.2 database using an Oracle ASM 11g Release 2 (11.2) instance is not able to use new features available for Oracle ASM 11g Release 2 (11.2) release, but instead it can use only Oracle ASM 10.2 features.

See Also:

- *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)* for information about using earlier Oracle Database releases with Oracle Grid Infrastructure

1.6 Confirming Cluster Readiness Using Cluster Verification Utility

Before you start your Oracle RAC installation, use CVU to ensure that your system is prepared for installing Oracle RAC. If any checks fail, then fix the errors reported, or contact your system or storage administrator to have the cause of the errors addressed.

CVU is available in the Grid home, in the bin directory. For example, if the Grid home is C:\app\11.2.0\grid, then CVU is located in the C:\app\11.2.0\grid\bin directory. To start CVU, go to the bin directory in the Grid home, and use a command similar to the following:

```
cluvfy.bat stage -pre dbinst -fixup -n nodelist -r release -verbose
```

In the preceding command, *nodelist* is a comma-delimited list of node names and *release* is the version of the Oracle Database software being installed. The *-fixup* and *-verbose* flags are optional.

For example, for a two node cluster with *node1* and *node2*, where you are testing the cluster to prepare to install Oracle Database 11g Release 2 (11.2) with Oracle RAC, the following command checks for system readiness:

```
cluvfy.bat stage -pre dbinst -fixup -n node1,node2 -r 11gR2 -verbose
```

For more information about CVU commands, run the command `cluvfy.bat -help`.

See Also: *Oracle Clusterware Administration and Deployment Guide* for detailed information about CVU

1.7 Installing Oracle Database 11g Release 2 on Systems with Existing Databases

If you are planning an installation on a system where you have an existing Oracle RAC or Oracle Database installation, then you must perform additional tasks to prepare your system for installation.

If your system has an Oracle Database 10g installation, and you install Oracle Database 11g Release 2 (11.2) either to coexist with or to upgrade the Oracle Database 10g release 2 (10.2) installation, then most installation types configure and start a default Oracle Net listener using TCP/IP port 1521 and the Inter-process communication (IPC) key value `EXTPROC`. One of the following actions occurs:

- During a coexisting installation, DBCA automatically migrates the listener and related files from the Oracle Database 10g Oracle home to the Oracle Database 11g Oracle home.
- During an upgrade, DBUA automatically locates the Oracle Database 10g listener, and migrates the listener to Oracle Database 11g.

Note: For Oracle Database 11g Release 2 (11.2) to coexist with Oracle Database 10g release 1 (10.2) databases using Oracle ASM, the 10.2 database must be release 10.2.0.2 or later.

See Also: *Oracle Database Upgrade Guide* for additional information about preparing for and performing upgrades

Installing Oracle Database 11g with Oracle Real Application Clusters

This chapter describes phase two of the installation procedures for installing Oracle Database 11g release 2 (11.2) with Oracle Real Application Clusters (Oracle RAC). Phase one is the completion of the Oracle Clusterware installation, as described in *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)*.

See Also: *Oracle Database Installation Guide for Microsoft Windows* for additional information, such as how to open compressed installation files, and other information about setting up files for installation

Also, during installation, click **Help** to find explanations of options and prompts for each installation screen, and click **Details** to see the log file.

This chapter contains the following topics:

- [Reviewing Oracle Application Express Requirements](#)
- [Selecting an Installation Type](#)
- [Selecting an Installation Option](#)
- [Selecting an Oracle Database Configuration Type](#)
- [Selecting the Oracle Database Vault Option](#)
- [Selecting the Oracle Configuration Manager Option](#)
- [Selecting a Database Name](#)
- [Creating Database Passwords](#)
- [Understanding the Actions of OUI and DBCA During Installation](#)
- [Using Oracle Universal Installer to Install Oracle RAC](#)
- [Using Oracle Universal Installer to Install Oracle RAC One Node](#)
- [Installing Oracle RAC and Oracle Database Vault](#)
- [Using DBCA to Configure Oracle RAC One Node](#)
- [Updating Environment Variables on Remote Nodes](#)

2.1 Reviewing Oracle Application Express Requirements

With Oracle Database 11g, Oracle Application Express is provided as an Oracle Database option, rather than as an Oracle Application option. If you intend to install Oracle Database after installing Oracle Clusterware, and you have an existing Oracle Database installation that you are upgrading, then review the Oracle Application Express requirements in *Oracle Application Express Installation Guide*.

2.2 Selecting an Installation Type

When you run Oracle Universal Installer (OUI), you can select the **Typical** or the **Advanced** installation type.

The Typical installation type installs a default configuration of Oracle Database, with basic configuration choices. Oracle recommends that most users select Typical as their installation type.

The Advanced installation type is for customized installations, and should be used only when you have a specific requirement for it, such as adding specific components to your installation, requiring different passwords for the SYS, SYSMAN, SYSTEM and DBSNMP accounts, using a different database character set than is in use on your servers, changing product languages, or other nonstandard configurations.

See Also:

- *Oracle Database Globalization Support Guide* for detailed information about character sets and language configuration
- *Oracle Database Installation Guide for Microsoft Windows* for information about configuring globalization

You can install Oracle RAC One Node on two or more nodes in the cluster. An Oracle RAC One Node installation starts up an instance on one of the nodes you select as an Oracle RAC One Node pool member. If that instance goes down, then the Oracle RAC One Node instance is started up on another pool member using Online Database Relocation. This feature migrates database instances and connections to other cluster nodes for high availability.

See Also:

- *Oracle Database High Availability Overview* for information about Oracle RAC One Node
- [Appendix C, "Converting to Oracle RAC and Oracle RAC One Node from Single-Instance Oracle Databases"](#) for information about how to convert single-instance databases to Oracle RAC

2.3 Selecting an Installation Option

On the Select Installation Option page, you can select one of the following options:

- **Create and Configure a Database:** Provides you with the option to create a database using a preconfigured database template designed for particular system load demands, such as an online transaction processing (OLTP) database, or a decision support or data warehouse database.

If you select the Oracle Automatic Storage Management (Oracle ASM) option for storage, then you must have previously configured Oracle ASM. If you did not select Oracle ASM as the storage option for the Oracle Clusterware files when

installing Oracle Grid Infrastructure, then you must first use Oracle ASM Configuration Assistant (ASMCA) from the Grid home to configure Oracle ASM. You must create an Oracle ASM instance and create a disk group to use for your Oracle Database storage before you can select Oracle ASM as a storage option for your Oracle RAC database.

- **Install database software only:** Installs Oracle Database software; you must complete the database configuration after the installation completes using the installed utilities.
- **Upgrade an existing database:** Upgrades an existing database.

If you are installing Oracle Database software, then Oracle recommends that you use a preconfigured database option, or select the **Advanced** option on the Select Configuration page, and configure a custom starter database. See "[Selecting an Oracle Database Configuration Type](#)" for more information about the different preconfigured database options.

See Also: If you have an existing Oracle installation, then write down the version numbers, patches, and other configuration information, and review upgrade procedures for your existing installation. Review *Oracle Database Upgrade Guide* before proceeding with the installation, to decide how you want to proceed.

For late-breaking updates and best practices about pre-upgrade, post-upgrade steps, compatibility, and interoperability discussions, see "Oracle Upgrade Companion," which is available through Note 785351.1 on My Oracle Support:

<https://support.oracle.com/>

2.4 Selecting an Oracle Database Configuration Type

When you run OUI, you can select the **General Purpose/Transaction Processing**, **Data Warehouse**, or **Advanced** database configuration type.

For the first two configuration types, you can complete additional procedures that are described later in this chapter.

If you select Advanced configuration, then you can use Oracle Database Configuration Assistant (DBCA) to create the database as described in [Chapter 3, "Creating Oracle Real Application Clusters Databases with Oracle Database Configuration Assistant"](#). Oracle recommends that you use DBCA to create your database.

2.4.1 Preconfigured Database Types Supplied with Oracle Database

The General Purpose and Transaction Processing type and the Data Warehouse type use preconfigured database templates optimized for each type of database.

During installation, OUI starts Oracle Net Configuration Assistant (NETCA) and DBCA, and installs the preconfigured database without further input. During database installation, OUI displays a progress indicator.

DBCA processing for these two configuration types creates a starter database, and configures the Oracle network services.

2.4.2 Using the Advanced Database Configuration Type

Advanced configuration options available using this installation type include Oracle RAC, Automatic Storage Management, backup and recovery configuration, integration with Enterprise Manager Grid Control, and more fine-grained memory tuning, as well as other options.

2.4.3 Installing Oracle Database with Other Languages

To use languages other than the default (English), either for the database or for applications running on the database, you must use the Advanced Installation method.

See Also: *Oracle Database Globalization Support Guide*

2.5 Selecting the Oracle Database Vault Option

Installing and configuring Oracle Database Vault requires actions during and after installation, as described in the following topics:

- [Overview of Oracle Database Vault Accounts](#)
- [Overview of Oracle Database Vault Installation](#)

2.5.1 Overview of Oracle Database Vault Accounts

Oracle Database Vault prompts for two accounts that you can create during installation. These are the Oracle Database Vault Owner and the Oracle Database Vault Account Manager accounts. Creating an Oracle Database Vault Account Manager is optional, but Oracle recommends that you create this account for better separation of duty.

You must supply a database user name and password for the Database Vault Owner during installation. Do not use the names of Oracle Database Vault roles (such as DV_OWNER and DV_ACCTMGR) for these two user accounts. See *Oracle Database Vault Administrator's Guide* for a full list of Database Vault roles.

2.5.1.1 Oracle Database Vault Owner

The Oracle Database Vault Owner account is granted the DV_OWNER role. This account can manage Oracle Database Vault roles and configuration.

The Oracle Database Vault Owner user name can be a minimum of 2, and a maximum of 30 characters. The account password can be a minimum of 8, and a maximum of 30 characters.

The password that you choose for the Oracle Database Vault Owner account must be secure. The following password restrictions are enforced:

- The password must include at least one alphabet, one digit, and one nonalphanumeric character (symbol).
- The password cannot be the same as the account name.
- The password cannot contain any consecutive repeating characters.

2.5.1.2 Oracle Database Vault Account Manager

The Oracle Database Vault Account Manager is granted the DV_ACCTMGR role. This account is used to manage database user accounts. The Oracle Database Vault Account Manager is created to facilitate separation of duties. If you do not choose to create the

Oracle Database Vault Account Manager account, then the DV_ACCTMGR role is granted to the Oracle Database Vault Owner account by default.

The same password restrictions that apply to Oracle Database Vault Owner are applicable to Oracle Database Vault Account Manager as well.

2.5.2 Overview of Oracle Database Vault Installation

Oracle Database Vault is installed with the Advanced Install option of Oracle Database Enterprise Edition. After the Oracle Database Vault software has been installed, you configure Oracle Database Vault using DBCA. Oracle Database Vault is not installed or configured by default.

2.6 Selecting the Oracle Configuration Manager Option

During installation, you are prompted to provide information needed to enable Oracle Configuration Manager (OCM). This option enables you to associate information about your Oracle RAC configuration with your My Oracle Support (formerly Oracle*Metalink*) account. In the event that you must place a service request with Oracle Support Services, that configuration information can help provide a more rapid resolution to the service issue.

You can enable Oracle Configuration Manager during or after installation, or choose not to enable it. Configuring the OCM tool requires that you have the following information available from your Oracle Support Services service agreement:

- My Oracle Support e-mail address or user name
- Password

In addition, you are prompted for server proxy information if the host system does not have a direct connection to the Internet.

Go to My Oracle Support (<https://support.oracle.com>) if you encounter registration failures.

See Also: *Oracle Configuration Manager Installation and Administration Guide* for more information

2.7 Selecting a Database Name

The database name input field is used to set the following Oracle initialization parameter values:

- DB_NAME
- DB_UNIQUE_NAME
- DB_DOMAIN

In Oracle RAC environments, the database name (DB_UNIQUE_NAME) portion is a string of no more than 30 characters that can contain alphanumeric, underscore (_), dollar (\$), and pound (#) characters, but must begin with an alphabetic character. No other special characters are permitted in a database name. The DB_NAME parameter for a database is set to the first 8 characters of the database name.

The domain portion of the global database name (DB_DOMAIN) can be no more than 128 characters. Domain names using underscores (_) are not allowed. The values for DB_UNIQUE_NAME.DB_DOMAIN in its entirety must be unique within the enterprise.

For example, if your database has a global database name of `orl$racprod2551.example.com` which you supplied during installation, then the following values are used for initialization parameters:

Initialization Parameter	Example Value
DB_UNIQUE_NAME	orl\$racprod2551
DB_DOMAIN	example.com
DB_NAME	orl\$racp

The SID prefix is the first 8 characters of the database name. The SID prefix can contain only the characters a-z, A-Z, and 0-9. The SID prefix cannot contain operating system special characters, so if you use special characters in the first 8 characters of the database name, then these special characters are omitted in the SID prefix. There is a single SID prefix for every database. The SID prefix for a database must be unique within the cluster.

For an Oracle RAC database, each instance has a unique identifier, `ORACLE_SID`, which consists of the SID prefix and an instance number. The `ORACLE_SID` for Oracle RAC database instances is generated differently, depending on how you choose to manage the database. If you select a policy-managed database, then Oracle generates the SID in the format `name_#`, where `name` is the first eight alphanumeric characters of `DB_UNIQUE_NAME`, and `#` is the instance number. If you select an admin-managed database, then DBCA generates the SID for the instance names in advance, and the SID is in the format `name#`.

For example, if the `DB_UNIQUE_NAME` for a database is `orl$racprod2551`, then the following SID values are used:

Database or Instance Type	Value Used for ORACLE_SID
Single-instance Oracle database	orlracpr
Policy-managed Oracle RAC instance	orlracpr_1
Admin-managed Oracle RAC instance	orlracpr1

2.8 Creating Database Passwords

Apply the following guidelines when specifying passwords:

- Passwords should be between 8 and 30 characters long.
- Passwords must be from the ASCII character set.
- Passwords must not start with a numeral.
- Passwords should not be the same as the user name.
- Passwords must not be Oracle reserved words.
- The SYS account password should not be `change_on_install`.
- The SYSTEM account password should not be `manager`.
- To use the same password for all the accounts, the password should not be `change_on_install`, `manager`, `sysman`, or `dbsnmp`.
- Passwords should have at least 1 alphabetic, 1 numeric, and 1 punctuation character.

- Passwords should not be simple or obvious words, such as `welcome`, `account`, `database`, `oracle`, or `user`.

2.9 Understanding the Actions of OUI and DBCA During Installation

After installing Oracle Database, OUI runs DBCA to create your database, in accordance with Optimal Flexible Architecture (OFA) guidelines. An installation created following OFA guidelines means that DBCA creates your database files, including the default server parameter file (SPFILE), using standard file naming and file placement practices.

Note: NETCA is run as part of an Oracle Grid Infrastructure installation.

The primary tasks of DBCA are the following:

- Create the database.
- Configure the Oracle network services, if necessary.
- Start the listeners and database instances.

You can also use DBCA in standalone mode to create or delete a database, or to switch database management from Oracle Enterprise Manager Database Control to Oracle Enterprise Manager Grid Control. See "[Managing Database Control Using EMCA](#)" on page 6-6.

See Also: *Oracle Database Net Services Administrator's Guide* if you experience problems (for example, with the listener configuration), and for more information about Lightweight Directory Access Protocol (LDAP) support

2.9.1 Managing Database Services After Installation

Starting with Oracle Database 11g release 1 (11.1), you cannot use DBCA to manage database services for Oracle RAC databases. Use the Server Control Utility (SRVCTL), Oracle Enterprise Manager Database Control or Oracle Enterprise Manager Grid Control for all administration and monitoring of database services for an Oracle RAC database.

See Also: *Oracle Database 2 Day + Real Application Clusters Guide* for service management using Oracle Enterprise Manager

2.9.2 Installing Earlier NLS Compatibility Definition Files

In Oracle Database 10g and later releases, some database language and territory definition files have been updated to better reflect locale conventions used in associated locales.

If the resulting changes in default National Language Support (NLS) parameter values adversely affect your existing applications, then you can revert the changes by installing Oracle9i Database compatibility definition files. To install Oracle9i Database files in place of Oracle11g Database files, you must run OUI from a command line, as described in "[Using Oracle Universal Installer to Install Oracle RAC](#)" on page 2-8, and use the following statement to set the `b_cr9idata` variable to true:

```
setup.exe oracle.rsfnlsrtl_rsf:b_cr9idata=true
```

See Also: *Oracle Database Globalization Support Guide* for additional information about updates to the Oracle Database language and territory definition files

2.10 Using Oracle Universal Installer to Install Oracle RAC

To install Oracle Database 11g release 2 (11.2) with Oracle RAC:

1. Log in to Windows as a user with Administrator privileges.
2. Verify that you have Administrator privileges on the other nodes. To do this, enter the following command for each node that is a part of the cluster where *node_name* is the name of the remote node:

```
net use \\node_name\C$
```

3. To use the Software Updates option to obtain the most current system requirement updates, critical patch updates, and other recommended updates for installation, use a valid My Oracle Support username and password, or provide the path to a downloaded Software Updates option package when prompted.
4. Start the `setup.exe` command from the base directory of the Oracle Database 11g release 2 (11.2) installation media, and select the database type to install.
5. Provide information when prompted by OUI. If you need assistance during installation, then click **Help**. If you encounter problems during installation, then examine the OUI actions recorded in the installation log file. The log file is located in the Oracle Inventory directory with a name that includes the time stamp (*date_time*) of the installation process, as shown in this example:

```
C:\Program Files\Oracle\Inventory\logs\installActionsdate_time.log
```

In the preceding example, the variables *date* and *time* represent the date and the time of the log file.

Caution: The Oracle home name and path that you provide during database installation *must be different* from the home that you used during Oracle Grid Infrastructure installation. **Do not** install Oracle Database 11g with Oracle RAC software into the same home in which you installed Oracle Grid Infrastructure software.

If you are upgrading a database, then you must use the same type of Oracle home. For example, you should not change from local Oracle homes to a shared Oracle home.

Note that because Oracle ASM is part of the Grid Infrastructure home, it cannot be upgraded during an Oracle RAC installation.

When you have completed this second and final phase of the installation, proceed to [Chapter 4, "Oracle Real Application Clusters Postinstallation Procedures"](#) to perform the postinstallation tasks.

Caution: After installation is completed and you have created the database, if you decide to install additional Oracle Database 11g release 2 (11.2) products in the database, then you must stop *all* processes running in the Oracle home before you attempt to install the additional products. See [Appendix E, "How to Stop Processes in an Existing Oracle Real Application Clusters Database"](#) for additional information.

The following is a list of additional information to note about installation:

- On the Select Database Management Option page, if you completed the Grid Control Management Agent installation, then you can select either Oracle Enterprise Manager Grid Control or Database Control. If you did not complete the Grid Control Management Agent installation, then only Database Control for database management is supported for Oracle RAC. When you use the Oracle Enterprise Manager Database Control console, you can choose the e-mail option for alerts and enter the outgoing mail server name and e-mail address.
- If you perform an installation that does not include Oracle Enterprise Manager (for example, an Advanced installation without configuring Oracle Enterprise Manager, an installation with no Oracle Enterprise Manager configuration, or a database creation with your own scripts), then you can configure Oracle Enterprise Manager later using OUI, DBCA, or the Oracle Enterprise Manager Configuration Assistant (EMCA) utility. See "[Managing Database Control Using EMCA](#)" on page 6-6.

2.11 Using Oracle Universal Installer to Install Oracle RAC One Node

To install Oracle RAC One Node:

1. Login to Windows as a user with Administrative privileges.
2. Start the `setup.exe` command from the base directory of the Oracle Database 11g release 2 (11.2) installation media, and select the database type to install.
3. Select to configure an Oracle RAC One Node database, and select to install Oracle RAC on all nodes where Oracle RAC One Node will run.
4. Provide information when prompted by OUI. If you need assistance during installation, click **Help**. Click **Details** to see the log file.

If you encounter problems during installation, then examine the OUI actions recorded in the installation log file. The log file is located in the Oracle Inventory directory with a name that includes the timestamp (*date_time*) of the install process, as shown in this example:

```
C:\Program Files\Oracle\Inventory\logs\installActionsdate_time.log
```

Caution: The Oracle home name and path that you provide during database installation *must be different* from the home that you used during the Oracle Grid Infrastructure installation in phase one. You *must not* install Oracle Database 11g with Oracle RAC software into the same home in which you installed the Oracle Grid Infrastructure for a cluster software.

If you are upgrading a database, then you must use the same type of Oracle home. For example, you should not change from local Oracle homes to a shared Oracle home.

Note that because Oracle ASM is part of the Oracle Grid Infrastructure home, it cannot be upgraded during an Oracle RAC installation.

2.12 Installing Oracle RAC and Oracle Database Vault

To install Oracle Database Vault while installing Oracle RAC:

1. Start OUI and select **Advanced installation**.
2. Select **Enterprise Edition**, and then select **Options**.
3. In addition to the default installation components, select **Oracle Label Security (OLS)** and **Oracle Database Vault**.
4. Continue the rest of the installation process as you normally would.

Note: You must set the database initialization parameter `db_block_size` to 4096 or higher. You can configure this parameter during the installation of Oracle RAC.

After you have installed Oracle RAC with the Oracle Database Vault option you must complete the configuration of Oracle Database Vault, as described in *Oracle Database Vault Administrator's Guide*.

2.13 Using DBCA to Configure Oracle RAC One Node

If you have selected to install Oracle RAC software only on cluster nodes, you can use Oracle Database Configuration Assistant (DBCA) to configure Oracle RAC One Node.

After installation, start up DBCA, and on the Welcome page, select **Oracle RAC One Node database**. DBCA starts on only one node of the set of nodes you select. You can add additional nodes to the cluster later.

Selecting one node deploys Oracle RAC One Node on a single node. Oracle recommends that you select all nodes in the cluster to which you want Oracle RAC One Node to be able to fail over.

Caution: By default, any named user may create a server pool. To restrict the operating system users that have this privilege, Oracle strongly recommends that you add specific users to the CRS Administrators list. See *Oracle Clusterware Administration and Deployment Guide* for more information about adding users to the CRS Administrators list.

If you select less than two nodes, or create server pools with a cardinality of 2 or more, then DBCA posts a warning message that the configuration you select will not support failover of the Oracle RAC One Node instance.

When you create an administrator-managed Oracle RAC One Node database, note that while the database is started on only one of the pool of nodes you installed the binaries, all the candidate servers are placed into the Generic server pool. If the servers are not already in Generic or Free, then this may result in stopping resources that are running on candidate servers.

See Also: Oracle Technology Network for more information about Oracle RAC One Node:

<http://www.oracle.com/technetwork/database/clustering/overview/>

2.14 Updating Environment Variables on Remote Nodes

During the installation, the Windows registry and environment variables are modified on each node on which you installed Oracle RAC. The new registry entries and environment variable settings are visible on the node where the installation was performed. However, the new settings on the remote nodes are not immediately available to your user session. Attempting to run scripts or applications from the newly installed Oracle home can produce errors similar to the following:

- ORACONFIG.exe - Unable to Locate Component
- OCI.dll not found

To make the modified environment variables available on the remote nodes, you can close your current session on the remote nodes (log off) and then log on to the remote nodes to create a new session, or perform the following actions:

1. From the Start menu, right-click **My Computer**, then select **Properties**.

Alternatively, enter **SYSDM.CPL** in the Run window.

2. Select the **Advanced** tab.
3. Click **Environment Variables**.

The modified environment variables are now visible.

4. Click **OK** to close the System Properties window.

Creating Oracle Real Application Clusters Databases with Oracle Database Configuration Assistant

This chapter describes how to use Oracle Database Configuration Assistant (DBCA) in standalone mode to create and delete Oracle Real Application Clusters (Oracle RAC) databases.

This chapter contains the following topics:

- [Using Oracle Database Configuration Assistant with Oracle RAC](#)
- [Optimizing Your Database Environment by Using DBCA](#)
- [Updating the Listener Used By an Earlier Release](#)
- [Verifying Requirements for DBCA](#)
- [Creating an Oracle RAC Database with DBCA](#)
- [Deleting an Oracle RAC Database with DBCA](#)
- [Setting Environment Variables for Enterprise Manager DB Control](#)

See Also: *Oracle Real Application Clusters Administration and Deployment Guide* for information about using DBCA to add and delete instances

3.1 Using Oracle Database Configuration Assistant with Oracle RAC

DBCA has the following primary database functions:

- Create and delete databases.
- Add and delete database instances.
- Set up network configurations for databases and their instances.
- Register databases in Oracle Enterprise Manager Grid Control or configure Oracle Enterprise Manager Grid Database Control.
- Configure database options, such as Oracle Database Vault.
- Start up databases and their instances.

Note: Cluster Managed Services are no longer managed through DBCA. Instead, use the Cluster Managed Services page in Oracle Enterprise Manager Database Control (accessible from the Cluster Database Availability page). For more information, see *Oracle Real Application Clusters Administration and Deployment Guide*.

See Also:

- ["Creating an Oracle RAC Database with DBCA"](#) on page 3-4 for more information about using DBCA in standalone mode
- *Oracle Database Net Services Administrator's Guide* to resolve problems—for example, with the listener configuration—and for more information about Lightweight Directory Access Protocol (LDAP)-compliant directory support
- *Oracle Database 2 Day + Real Application Clusters Guide* for more information about using Oracle Enterprise Manager

3.2 Optimizing Your Database Environment by Using DBCA

Oracle recommends that you use DBCA to create your Oracle RAC database, because preconfigured databases optimize your environment for features such as the server parameter file (SPFILE) and automatic undo management. If you use Oracle Automatic Storage Management (Oracle ASM) or cluster file system storage, then DBCA also configures automated backup, which uses the Fast Recovery Area.

DBCA enables you to create both policy-managed and administrator-managed databases. With DBCA, you can create site-specific tablespaces as part of database creation. If you have data file requirements that differ from those offered by DBCA templates, then create your database with DBCA and modify the data files later. You can also run user-specified scripts as part of your database creation process.

DBCA also configures your Oracle RAC environment for various Oracle high availability features, such as cluster administration tools. DBCA also starts any database instances required to support your defined configuration.

You can use DBCA to create a database from templates supplied by Oracle, or from templates that you create. The templates contain settings optimized for a particular type of workload.

Oracle ships templates for the following two workload types:

- General purpose or transaction processing
- Data warehouse

For more complex environments, you can select the **Custom Database** option. This option does not use templates and results in a more extensive interview, which means that it takes longer to create your database.

Click **Show Details** to see the configuration for each type of database. Select the template suited to the type of workload your database will support. If you are not sure which to choose, then select the default **General Purpose or Transaction Processing** template.

3.3 Updating the Listener Used By an Earlier Release

If your system has an Oracle Database 10g or 11g installation, and you install Oracle Database 11g release 2 (11.2) either to coexist with or to upgrade the Oracle Database 10.1, 10.2, or 11.1 installation, then most installation types automatically migrate the existing Oracle Database listener to the 11g release 2 (11.2) Oracle home. During migration, the upgrade process configures and starts a default Oracle Net listener using the same TCP/IP port as the existing listener, with the IPC key value.

During the Oracle Clusterware upgrade, the default listener (`LISTENER_NODENAME`) is migrated to the Oracle Grid Infrastructure home (Grid home). DBCA always uses the default listener.

The listener migration process stops the listener in the existing Oracle home, and restarts it in the new Oracle home. If the database uses the default listener (`LISTENER_NODENAME`), then the listener is migrated automatically to the Oracle Clusterware home by Oracle Net Configuration Assistant (NETCA) as part of the Oracle Clusterware upgrade. If the database uses a nondefault listener, then DBUA migrates the nondefault listener to the Oracle home.

Note: During migration, client applications may not be able to connect to any databases that are registered to the listener that is being migrated.

3.4 Verifying Requirements for DBCA

To use DBCA to change the database configuration, use Cluster Verification Utility (CVU) to verify that your system is prepared for configuration changes using the following command syntax:

```
Grid_home\bin\cluvfy stage -pre dbcfg -n node_list -d Oracle_home [-verbose]
```

In the preceding example, the variable `Grid_home` is the Grid home, the variable `node_list` is the list of nodes in your cluster, separated by commas, and the variable `Oracle_home` is the path for the Oracle home directory where OUI creates or modifies the database.

For example, to verify that your system is prepared for an Oracle Database with Oracle RAC installation on a two-node cluster with nodes `node1` and `node2`, with the Grid home path `C:\app\11.2.0\grid`, and with the Oracle home path `C:\app\oracle\product\11.2.0\dbhome_1`, enter the following command:

```
C:\app\11.2.0\grid\bin> cluvfy stage -pre dbcfg -n node1,node2 \
-d C:\app\oracle\product\11.2.0\dbhome_1
```

You can select the option `-verbose` to receive progress updates as CVU performs its system checks, and detailed reporting of the test results.

If the CVU summary indicates that the cluster verification fails, then review and correct the relevant system configuration steps, and run the test again.

The command `cluvfy stage -pre dbcfg` verifies the following:

- Node Reachability: All the specified nodes are reachable from the local node.
- User Equivalence: User equivalence exists on all the specified nodes.
- Total Memory: Each node specified has at least the minimum required memory configured.

- Available Memory: Each node specified has at least the minimum amount of available memory.
- Swap Space: Each node has enough swap space available.
- Free Disk Space: Each node has sufficient free disk space available in the path for the Oracle home, the Grid home, and the temporary directory.
- System Architecture: Each node must be a x64 (64-bit) system.
- Length of the Value of the Environment Variable PATH: The PATH environment variable is set on each node and does not contain more than 1023 characters.
- Oracle Clusterware Integrity: All the components of Oracle Clusterware are fully operational.
- Node Application Existence: All the specified nodes have the following node resources registered with Oracle Clusterware and whether the node resources are online:
 - virtual IP (VIP)
 - Oracle Notification Services (ONS) (not required)
 - Global Services Daemon (GSD) (not required)
 - EONS (not required)
 - NETWORK (not required)
- Time Zone Consistency: All the specified nodes have the same time zone setting.

3.5 Creating an Oracle RAC Database with DBCA

To create an Oracle RAC database with DBCA without Oracle ASM or a cluster file system, you must have configured shared storage devices. In addition, you must have run NETCA to configure your Oracle Net listener.ora file.

This section contains the following topics:

- [Configuring Disks for Use With Oracle ASM](#)
- [Starting DBCA](#)
- [Using DBCA to Create an Oracle RAC Database](#)

3.5.1 Configuring Disks for Use With Oracle ASM

On Windows-based systems, if you plan to use Oracle ASM storage, then before you use DBCA to create a database, you must first create logical partitions without primary partitions on the disk drives. After creating the logical disk partitions you must delete the drive letters for these partitions on all nodes. Alternatively, you can stamp these partitions with `asmtool`. See *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)* for more information on configuring disk devices for use with Oracle ASM. After you have configured the disks to be used by Oracle ASM, you must create the disk groups that will be used by the database prior to starting DBCA. You can create disk groups by using SQL*Plus, Oracle Enterprise Manager, or Oracle ASM Configuration Assistant (ASMCA).

When using ASMCA to create a disk group, if you do not see the storage devices to add, then click **Change Disk Discovery Path** to alter the search path used by ASMCA to find available storage devices, or click **Stamp Disks** to start the `asmtoolg` GUI tool. You can select storage devices with a status of Candidate (the storage device was

never used in an Oracle ASM disk group) or Former (the storage device is no longer in a disk group). To add storage devices that still have Oracle ASM disk headers, but the disk group is no longer in use (a case that can occur if you are selecting storage devices after a terminated installation attempt, you attempted deinstallation without dropping the disk group, or other configuration problems), use the `Force` command.

See Also:

- *Oracle Database Installation Guide for Microsoft Windows* for more information about `asmtoolg`
- *Oracle Automatic Storage Management Administrator's Guide* for more information about creating disk groups

3.5.2 Starting DBCA

To start DBCA, connect as the user with Administrator privileges that was used to install the Oracle RAC software (for example, `oracle`) on one of your nodes in your cluster, and enter the command `dbca` from the `Oracle_home\bin` directory.

On Windows-based platforms, you can also click **Start**, and select **Programs, Oracle - Oracle_home name, Configuration and Migration Tools**, and then **Database Configuration Assistant**.

Note: To run DBCA, you no longer need to set operating system environment variables `%ORACLE_HOME%` to the Oracle RAC database home, or `%ORACLE_UNQNAME%` to the database unique name.

3.5.3 Using DBCA to Create an Oracle RAC Database

When you start DBCA, the first page that is displayed is the Welcome page for Oracle RAC, which includes the option to select an Oracle RAC database. DBCA displays this Oracle RAC Welcome page only if the Oracle home from which it is started is installed on a cluster. Enter information as prompted by DBCA. Click **Help** if you need assistance.

If DBCA does not display the Welcome page for Oracle RAC, then DBCA was unable to detect if the Oracle home is installed on a cluster. In this case, verify that the OUI inventory is correctly located in the directory `C:\Program Files\Oracle\Inventory\oraInst.loc`, and that the inventory files are not corrupted. Also, perform clusterware diagnostics by using the following CVU command syntax, where `C:\app\11.2.0\grid` is the location of your Grid home:

```
C:\app\11.2.0\grid\bin\cluvfy stage -post crsinst -n nodelist
```

For example, if your cluster consists of nodes `node1` and `node2`, run the following command:

```
C:\app\11.2.0\grid\bin\cluvfy stage -post crsinst -n node1,node2
```

Note the following important information when using DBCA:

- If nodes that are part of your cluster installation do not appear on the Node Selection page, then run the `lsinventory` command of OPatch to perform inventory diagnostics and run CVU to perform clusterware diagnostics.
- In Oracle RAC environments, the database name (`DB_UNIQUE_NAME`) portion is a string of no more than 30 characters that can contain alphanumeric, underscore (`_`), dollar (`$`), and pound (`#`) characters, but must begin with an

alphabetic character. No other special characters are permitted in a database name. The `DB_NAME` parameter for a database is set to the first 8 characters of the database name.

- The domain portion of the global database name (`DB_DOMAIN`) can be no more than 128 characters. Domain names using underscores (`_`) are not allowed. The values for `DB_UNIQUE_NAME.DB_DOMAIN` in its entirety must be unique within the enterprise.
- A system identifier (SID) is used as a unique identifier for each Oracle instance. The SID prefix, or first 8 characters of the database name, can contain only the characters `a-z`, `A-Z`, and `0-9`. The SID prefix cannot contain operating system special characters, so if you use special characters in the first 8 characters of the database name, then these special characters are omitted in the SID prefix. There is a single SID prefix for every database. The SID prefix for a database must be unique within the cluster. On some platforms, the SID is case-sensitive.

For an Oracle RAC database, each instance has a unique identifier, `ORACLE_SID`, which consists of the SID prefix and an instance number. The `ORACLE_SID` for Oracle RAC database instances is generated differently, depending on how you choose to manage the database. If you select a policy-managed database, then Oracle generates the SID in the format `name_#`, where `name` is the first eight alphanumeric characters of `DB_UNIQUE_NAME`, and `#` is the instance number. If you select an admin-managed database, then DBCA generates the SID for the instance names in advance, and the SID is in the format `name#`.

Note: It is common practice to set the SID to be equal to the database name. For more information, see the discussion of the `DB_NAME` initialization parameter in *Oracle Database Reference*.

- On the Management Options page, you are provided with selections for Oracle Enterprise Manager monitoring and management interfaces.
Enterprise Manager Database Control is always an option. If DBCA discovers Grid Control agents on the cluster, then it also provides you with the option Enterprise Manager with the Grid Control. Choose between Database Control and Grid Control interfaces.
You can enable daily backup operations. For daily backups, you enter the backup time and operating system credentials for the user that performs backup operations.
You can no longer set up e-mail notification either from DBCA or OUI starting with Oracle Database 11g Release 2 (11.2.0.2).
- To use a Fast Recovery Area, Oracle recommends that you create at least two separate Oracle ASM disk groups on shared storage: one for the database area and one for the recovery area. Oracle recommends that you place the database area and the recovery area in separate failure groups.
A failure group is defined by shared hardware, such as a controller shared between two storage devices, or two disks that are on the same spindle. If two storage devices share hardware that could fail, making both storage devices unavailable, then these storage devices are said to be in the same failure group.
If you do not use Oracle ASM, then Oracle recommends that the data files and the Fast Recovery Area are located outside of the Oracle home and in separate

locations, as with separate Oracle ASM failure groups, so that a hardware failure does not affect availability.

See Also:

- *Oracle Database Concepts* for more information about using a Fast Recovery Area
- *Oracle Automatic Storage Management Administrator's Guide* for information about failure groups and best practices for high availability and recovery
- On the Recovery Configuration page, if you are using Oracle ASM or cluster file system storage, then you can also select the Fast Recovery Area and the size of the recovery area on the Recovery Configuration page. If you are using Oracle ASM, then the Fast Recovery Area defaults to an Oracle ASM disk group. If you are using OCFS for Windows, then the Fast Recovery Area defaults to *Oracle_base\fast_recovery_area*.
- If you intend to add more nodes to your cluster than you have during the current DBCA session, then click **All Initialization Parameters**, and change the parameter `CLUSTER_DATABASE_INSTANCES` to the total number of nodes that you plan to add to the cluster.

In addition, if you click **All Initialization Parameters**, note that if your global database name is longer than 8 characters, then the database name value (in the `DB_NAME` parameter) is truncated to the first 8 characters, and the `DB_UNIQUE_NAME` parameter value is set to the global name.

See Also: *Oracle Database Administrator's Guide* for information about initialization parameters

- When you are prompted to create the Database Vault user accounts, do not use the names of Database Vault roles, such as `DV_OWNER` and `DVACCTMGR`. See *Oracle Database Vault Administrator's Guide* for a full list of roles.

After you respond to DBCA prompts, review the Summary dialog information and then click **OK**. DBCA does the following:

- Starts the Oracle services if you are on a Windows-based platform
- Configures the network for the cluster database
- Migrates Oracle Database listeners from a previous release and related files to the Oracle Database 11g Oracle home
- Creates an Oracle RAC database and its instances
- Creates the Oracle RAC data dictionary views
- Configures Oracle Enterprise Manager Database Control or Grid Control
- Starts the listeners and database instances, and then starts the high availability services

Caution: After you have created the database, if you decide to install additional Oracle Database products in the database you have created, then you must stop *all* processes running in the Oracle home before you attempt to install the additional products. See [Appendix E, "How to Stop Processes in an Existing Oracle Real Application Clusters Database"](#) for additional information.

3.6 Deleting an Oracle RAC Database with DBCA

Deleting an Oracle RAC database using DBCA involves first deleting the database, and then removing the database's initialization parameter files, instances, Optimal Flexible Architecture (OFA) structure, and the Oracle network configuration for the database. However, this process does not remove data files if you placed the files on raw devices or on raw partitions.

To delete a database using DBCA:

1. Start DBCA on one of your cluster nodes:

On Windows-based platforms, click **Start**, and select **Programs, Oracle - Oracle_ home name, Configuration and Migration Tools**, and then **Database Configuration Assistant**.

The DBCA Welcome page appears.

2. Select **Oracle Real Application Clusters database**, then click **Next**.

After you click Next, DBCA displays the Operations page.

3. Select **Delete a database**, and click **Next**.

DBCA displays the List of Cluster Databases page.

4. If your user ID and password are not operating-system authenticated, then the List of Cluster Databases page displays the user name and password fields. If these fields appear, then enter a user ID and password for a user account that has SYSDBA privileges.

5. Select the database to delete, and click **Finish**.

After you click Finish, DBCA displays a dialog box to confirm the database and instances that DBCA is going to delete.

6. Click **OK** to begin the deletion of the database and its associated files, services, and environment settings, or click **Cancel** to stop the operation.

When you click OK, DBCA continues the operation and deletes all the associated instances for this database. DBCA also removes the parameter files and password files.

At this point, you have accomplished the following:

- Deleted the selected database from the cluster
- Deleted the Oracle services for the selected database
- Deleted high availability services that were assigned to the database
- Deleted the Oracle Net configuration for the database
- Deconfigured Database Control for that database
- Deleted the OFA directory structure for that database from the cluster
- Deleted the data files if the data files were not on raw devices or raw partitions

3.7 Setting Environment Variables for Enterprise Manager DB Control

In previous releases of Oracle Database, you were required to set the environment variables `ORACLE_HOME` and `ORACLE_SID` to start, stop, and check the status of Enterprise Manager. With Oracle Database 11g release 2 (11.2) and later, you need to set the environment variables `ORACLE_HOME` and `ORACLE_UNQNAME` to use Enterprise Manager. For example, on each node, enter commands similar to the following to set these values for the `oracle` user, where the Oracle home is

`C:\app\oracle\product\11.2.0.1\dbhome_1`, and where the database unique name is `sales`:

```
C:\> set ORACLE_HOME=C:\app\oracle\product\11.2.0.1\dbhome_1
C:\> set ORACLE_UNQNAME=sales
```

Place these environment variables in the `oracle` user profile file on each cluster member node to ensure that the environment variable settings are available after system restarts.

Oracle Real Application Clusters Postinstallation Procedures

This chapter describes how to complete the postinstallation tasks after you have installed the Oracle Database 11g release 2 (11.2) with Oracle Real Application Clusters (Oracle RAC) software. This chapter contains the following topics:

- [Required Postinstallation Tasks](#)
- [Recommended Postinstallation Tasks](#)
- [Oracle Configuration Manager Postinstallation Configuration for Oracle RAC](#)
- [Enabling and Disabling Database Options](#)
- [Using Earlier Oracle Database Releases with Oracle Grid Infrastructure](#)

Note: This chapter describes only basic configurations. See also:

- *Oracle Database Platform Guide for Microsoft Windows* and the product administration and tuning guides for more detailed configuration and tuning information
 - *Oracle Database Installation Guide for Microsoft Windows* for additional postinstallation configuration information
-
-

4.1 Required Postinstallation Tasks

You must perform the following tasks after completing your installation:

- [Determine If Any Patches Are Required For Your New Software](#)
- [Types of Patches Available for Oracle Software](#)
- [Downloading and Installing Patch Updates](#)
- [Configuring Exceptions for the Windows Firewall](#)
- [Configuring Oracle Products](#)
- [Configuring Storage for External Tables, Shared Files, or Directory Objects](#)

4.1.1 Determine If Any Patches Are Required For Your New Software

After installing the Oracle software, sometimes you will need to apply one or more patches to enable certain software features to work correctly. To determine if any patches are required for your system, review the *Oracle Database Readme* and the Oracle Database Release Notes 11g Release 2 (11.2.0.3) for Microsoft Windows.

See Also: ["Required Patches for Earlier Oracle Database Releases or Databases Being Upgraded"](#) on page 4-10 for information about patches required for Oracle Database software release 11.1 or earlier

4.1.2 Types of Patches Available for Oracle Software

On a regular basis Oracle provides patch sets that include generic and port specific fixes encountered by customers since the base product was released. Patch sets increment the 4th digit of the version number e.g. 10.2.0.1.0 to 10.2.0.4.0, these patch sets are fully regression tested in the same way that the base release is (i.e. 10.2.0.1.0). Customers are encouraged to apply these fixes.

If a customer encounters a critical problem that requires a fix prior to the next patch set becoming available, they can request that a one off fix is made available on top of the latest patch set. This delivery mechanism is similar to the Microsoft Hot Fixes and is known as an Oracle patch set exception (or interim patch). Unlike Unix platforms these patch set exceptions are delivered in a **patch set exception bundle** (cumulative patch bundle), which includes all fixes since the current patch set. For example, bug 12393432 is a patch set exception bundle, Patch 12, for Oracle Database release 11.2.0.1 for Microsoft Windows (x64). You should always apply the latest patch bundle available for your release.

The patch set exception bundles also include the fixes for the CPU (Critical Patch Update), DST (Daylight Saving Time), PSU (Patch Set Update) and Recommended Patch Bundles. It is not required to have previous security patches applied before applying the patch set exception bundle. However, you must be on the stated patch set level for a given product home before applying the patch set exception bundle for that release.

See Also: *Oracle Database High Availability Overview* for more information about the various types of patches.

4.1.3 Downloading and Installing Patch Updates

Go to the My Oracle Support web site for required patch updates for your installation.

Note: Browsers require an Adobe Flash plug-in, version 9.0.115 or later to use My Oracle Support. Check your browser for the correct version of the Flash player by going to the Adobe Flash checker page, and installing the latest version of Adobe Flash.

If you do not have Adobe Flash installed, then download the latest version of the Flash Player from the Adobe web site:

<http://www.adobe.com/go/getflashplayer>

To download required patch set updates:

1. Use a web browser to view the My Oracle Support web site:

<https://support.oracle.com>

2. Log in to My Oracle Support.

Note: If you are not a My Oracle Support registered user, then click **Register for My Oracle Support** and register.

3. On the main My Oracle Support page, click the **Patches & Updates** tab.
4. On the left side of the page, in the Patching Quick Links section, under the heading Oracle Server / Tools, click **Latest Patchsets**.

A new browser window or tab is opened, displaying the Quick Links page.

5. In the Patch Sets for Product Bundles section, cursor over the link for Oracle Database. In the floating list that appears, move your cursor to highlight your operating system, for example Microsoft Windows x64 (64-bit). In the drill-down list that appears, click the desired patch set version, for example, 11.2.0.2.

A Search Results page appears, with the selected patch set displayed at the bottom. Patch sets for Oracle databases are identified in the Description column as Oracle Database Family:Patchset *x.x.x.x* PATCH SET FOR ORACLE DATABASE SERVER.

- a. Click the number of the desired patch set in the first column of the search results. On the patch set information page, click **View Readme** to view the Patch Set Notes, also called the README file, for that patch. The README page contains information about the patch set and how to apply the patches to your installation.
 - b. To download the patch to your local file system, click the **Download** button on the patch set information page.
6. To locate all the available patches instead of patch sets, use your browser navigation buttons to return to the Search Results page. In the middle of the Search Results page, in the Patch Type action list, change the current setting to **Patch**. Below the search fields, click the **Go** button.

The previous search is repeated, but this time, a list of all available patches is displayed instead of a list of patch sets. To access the patch information, perform the tasks listed in Step 5.a. and Step 5.b.

7. After you have downloaded the patch and transferred it to your server, use the UNZIP utility provided with Oracle Database 11g release 2 (11.2) to uncompress the Oracle patch updates. The UNZIP utility is located in the `Oracle_home\bin` directory.

Note: You can click the link on the bottom of the patch set information page to obtain an UNZIP utility for most operating systems.

4.1.4 Configuring Exceptions for the Windows Firewall

If the Windows Firewall feature is enabled on one or more of the nodes in your cluster, then virtually all TCP network ports are blocked to incoming connections. As a result, any Oracle product that listens for incoming connections on a TCP port does not receive any of those connection requests and the clients making those connection requests report errors.

If the Windows Firewall feature is enabled on one or more nodes of your Oracle RAC cluster, you must create exceptions for Oracle RAC applications and ports as described in *Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)*.

4.1.5 Configuring Oracle Products

Many Oracle products and options must be configured before you use them for the first time. Before using individual Oracle Database 11g release 2 (11.2) products or options, refer to the appropriate manual in the product documentation library, which is available in the DOC directory on the Oracle Database installation media, or on the Oracle Technology Network (OTN) web site at:

<http://www.oracle.com/technetwork/indexes/documentation/>

4.1.5.1 Configuring Oracle Database Vault

If you installed Oracle Database Vault during the Oracle RAC installation, then you must register Oracle Database Vault with the database and create database user accounts.

See Also: *Oracle Database Vault Administrator's Guide* for more information about configuring Oracle Database Vault after installation

4.1.5.2 Configuring Oracle Database Security Settings

To enable or disable the database security configuration after installation, you must use command-line Oracle Database Configuration Assistant (DBCA) options. By design, the DBCA graphical user interface (GUI) does not have the option to enable or disable secure configuration. For example, to enable the security settings after installation, you would use a command of the following form, where myRACdb1.example.com is either the name of the local database instance or the DB_UNIQUE_NAME:

```
dbca -configureDatabase -sourceDB myRACdb1.example.com -SID  
-enableSecurityConfiguration true
```

See Also: *Oracle Database Installation Guide for Microsoft Windows* for more information about enabling and disabling the database security option

4.1.6 Configuring Storage for External Tables, Shared Files, or Directory Objects

If your Oracle RAC database uses files that are external to the database, then the external files should be located on shared storage that is accessible to all nodes. Each node should use the same mount point to access the file. Acceptable shared file systems include Database File System (DBFS), Oracle ASM Cluster File System (Oracle ACFS), Oracle Cluster File System (OCFS) for Windows, or a supported network file system (NFS).

The database directory object used to write and read files external to the database must point to a shared storage location, and each node must use the same mount point for the same shared storage location. For example, each node might have a directory object called DPUMP for the mount point C:\app\acfsmounts\dpump, which accesses Oracle ACFS shared storage.

Note: There is no checking of the contents of the external files or directory object specified as part of the external table to ensure that the directory contents are consistent on each node. To avoid unpredictable results, you must ensure that the same file is accessed from all nodes, or that the same file is used on all nodes.

4.2 Recommended Postinstallation Tasks

Oracle recommends that you complete the following tasks after completing an Oracle RAC installation:

- [Setting the Oracle User Environment Variables](#)
- [Recompiling All PL/SQL modules](#)
- [Setting Up User Accounts](#)
- [Configuring Oracle Administration Assistant for Windows](#)
- [Using Oracle9i Database Language and Definition Files with Oracle Database 11g](#)
- [Logging in to Oracle Enterprise Manager Database Control](#)

4.2.1 Setting the Oracle User Environment Variables

Unlike on other platforms, you should not set `ORACLE_HOME` as a fixed environment variable when running Oracle on Windows operating systems. This is because the Oracle software determines where executable files reside at run time.

When you invoke an Oracle executable program on Windows, for example `sqlplus.exe`, the `ORACLE_HOME`, `ORACLE_BASE`, and `ORACLE_SID` variables are determined by the `PATH` environment variable and the location of the executable program (which Oracle home it resides in). If you want to use SQL*Plus to manage a different database or Oracle ASM instance, then you can click the Windows Start button, select the correct Oracle Home for the instance you want to manage, and then select the SQL*Plus utility.

You can use Oracle Universal Installer (OUI) to specify an Oracle home as the default Oracle home and update the `PATH` environment variable to point to that Oracle home. Refer to [Section 7.3, "Working with Multiple Oracle Home Directories on Windows"](#) for detailed instructions on how to change the default Oracle home.

4.2.2 Recompiling All PL/SQL modules

Oracle recommends that you run the `utlrp.sql` script after creating or upgrading a database. This script recompiles all PL/SQL modules that might be in an invalid state, including packages, procedures, and types. This is an optional step, but Oracle recommends that you do it immediately following installation, not at a later date.

1. Complete setting up the Oracle software owner user account environment (for example, `oracle`), as described in ["Setting the Oracle User Environment Variables"](#).
2. Start SQL*Plus, as follows:

Click **Start**, select **Programs** (or **All Programs**), then **Oracle - HOME_NAME**, then **Application Development**, and then **SQL*Plus**.

3. Run the `utlrp.sql` script, where `Oracle_home` is the Oracle home path:

```
SQL> @Oracle_home\rdbms\admin\utlrp.sql
```

See Also: *Oracle Database Administrator's Guide* for information about connecting to the database using SQL*Plus

4.2.3 Setting Up User Accounts

For information about setting up additional optional user accounts, see *Oracle Database Platform Guide for Microsoft Windows*.

When you complete these procedures, you are ready to perform the initial configuration tasks described in [Chapter 5, "Configuring the Server Parameter File in Oracle Real Application Clusters Environments"](#).

4.2.4 Configuring Oracle Administration Assistant for Windows

Oracle Administration Assistant for Windows requires the Microsoft Management Console and HTML Help 1.2 or later to run. Microsoft Management Console (MMC) version 2.0 ships with Windows Server 2003; version 3.0 of MMC is available with Windows Server 2003 R2 and later releases. Oracle recommends that you use the latest version of MMC that is available.

See Also:

- *Oracle Database Platform Guide for Microsoft Windows*
- Microsoft documentation at:
<http://www.microsoft.com/>

4.2.5 Using Oracle9i Database Language and Definition Files with Oracle Database 11g

You can use Oracle9i Database language and territory definition files with Oracle Database 11g release 2 (11.2).

To enable this functionality:

1. Run the `cr9idata.pl` script, by default located in `Oracle_home\nls\data\old`.
Alternatively, before you install Oracle Database, you can run the Oracle Universal Installer (OUI) setup command with the `b_cr9idata` variable set to true, as follows:

```
setup.exe oracle.rsfnlsrtl_rsf:b_cr9idata=true
```
2. Set the `ORA_NLS10` environment variable to point to the directory where you installed the new language and territory definition files, which by default are in `Oracle_home\nls\data`.
3. Restart the Oracle database.

4.2.6 Logging in to Oracle Enterprise Manager Database Control

If you configured Oracle Enterprise Manager Database Control during installation, then you can use it to manage your database. Oracle Enterprise Manager Database Control provides a web-based user interface that you can use to monitor, administer, and maintain an Oracle database, including Automatic Storage Management.

See Also:

- *Oracle Database Installation Guide for Microsoft Windows* for information about logging in to Enterprise Manager Database Control
- ["Managing Database Control Using EMCA"](#) on page 6-6

4.3 Oracle Configuration Manager Postinstallation Configuration for Oracle RAC

If you have installed Oracle Configuration Manager, then you must run a script to create a database account to collect database configuration collections. You must create this account in both *Connected* and *Disconnected* modes. The database account stores the PL/SQL procedures that collect the configuration information, and the account owns the database management system (DBMS) job that performs the collection. After the account has been set up, the account is locked because login privileges are no longer required.

To configure the database for configuration collection, run the following script:

```
Oracle_home\ccr\admin\scripts\installCCRSQL.exe collectconfig -s SID -r \
SYSDBA-USER
```

The script `installCCRSQL.exe` creates an Oracle Configuration Manager user and loads the PL/SQL procedure into the database defined by the `ORACLE_SID` environment variable. You can also specify the database SID by using the `-s` option in the command line, as in the following example, where the SID is `orcl`:

```
Oracle_home\ccr\admin\scripts\installCCRSQL.exe collectconfig -s orcl
```

For Oracle RAC, you must run the database script against only one instance, such as the local instance on which you performed the installation. However, Oracle Configuration Manager must be installed in all instance homes.

By default, the connection to the database is through operating system authentication, `"/as sysdba."` To specify a different user name and password, you can use these options:

- r `SYSDBA-USER`: The login name for the user with the SYSDBA privilege
- p `SYSDBA-PASSWORD`: The password for the user with the SYSDBA privilege

Note:

- If you specify the user name without specifying the password, then you are prompted to enter the password.
 - If you specify only the password without specifying the user name, then the user `SYS` is used by default.
-
-

See Also:

- *Oracle Configuration Manager Installation and Administration Guide* for more information
- *Oracle Database Vault Administrator's Guide* for additional configuration tasks

4.3.1 Additional Step for Oracle Enterprise Manager Grid Control

If the database is used as a repository for Oracle Enterprise Manager Grid Control, then you must also run the following script when configuring Oracle Configuration Manager:

```
Oracle_home\ccr\admin\scripts\installCCRSQL.exe collectemrep
```

When you run this script, the application prompts you for the `SYSMAN` password. To automate the installation, run the `installCCRSQL` script to specify the `SYSMAN` password. For example:

```
Oracle_home\ccr\admin\scripts\installCCRSQL.exe collectemrep
-e SYSMAN PASSWORD
```

You can add the `-s SID` command to specify the `SID` of the Oracle Enterprise Manager Grid Control database instance. You must run this script from the `Oracle_home` in which the Oracle Enterprise Manager Grid Control database has been hosted.

If you are not using operating system authentication to connect to the database, then you must use the `-r` and `-p` parameters to specify the following:

`-r SYSDBA-USER`: The login name for the user with the `SYSDBA` privilege

`-p SYSDBA-PASSWORD`: The password for the user with the `SYSDBA` privilege

If you do not specify the `-p` parameter, then you are prompted to enter the password for the specified user.

4.4 Enabling and Disabling Database Options

When you install Oracle Database, some options are enabled and others are disabled. You can view the enabled Oracle Database options by querying the `V$OPTION` view using `SQL*Plus`.

See Also: *Oracle Database Administrator's Guide* for information about connecting to the database using `SQL*Plus`

To enable or disable a particular database feature for an Oracle home, use the `chopt` tool. The `chopt` tool is a command-line utility that is located in the `Oracle_home\bin` directory. The syntax for `chopt` is as follows:

```
chopt [ enable | disable] db_option
```

The possible values for `db_option` are described in [Table 4-1](#):

Table 4-1 Database Options for Chopt Tool Command

Value	Description
dm	Oracle Data Mining Database Files
dv	Oracle Database Vault
lbac	Oracle Label Security
olap	Oracle OLAP
partitioning	Oracle Partitioning
rat	Oracle Real Application Testing
ode_net_2	Oracle Database Extensions for .NET 2.0

Before you run the `chopt` tool, shut down any Oracle databases running from the Oracle home being modified.

Example 4-1 How to Enable the Oracle Label Security Database Option Using Chopt

To enable the Oracle Label Security option in your Oracle executable files, perform the following tasks:

1. Shut down the Oracle RAC database, and any other databases running out of the Oracle home directory, using the Server Control Utility (SRVCTL).

```
srvctl stop database -d myRACdb
```

2. Use the Windows Services console to stop the service for each Oracle database that runs from the Oracle home being modified.

For example, if your Oracle RAC database is named `myRACdb`, then on the first node of the cluster you would stop the `OracleServiceMYRACDB1` service.

3. Change directories to the `Oracle_home\bin` directory:

```
cd %ORACLE_HOME%\bin
```

4. Run the `chopt` tool, as follows:

```
chopt enable lbac
```

5. Use the Windows Services console to restart the services you stopped in Step 2.

6. Use SRVCTL to restart the databases you stopped in Step 1, for example:

```
srvctl start database -d myRACdb
```

4.5 Using Earlier Oracle Database Releases with Oracle Grid Infrastructure

The following topics describe using earlier Oracle Database releases with Oracle Clusterware 11g release 2 (11.2) installations:

- [General Restrictions for Using Earlier Oracle Database Releases](#)
- [Required Patches for Earlier Oracle Database Releases or Databases Being Upgraded](#)
- [Pinning Cluster Nodes for Oracle Database Release 10.2 or 11.1](#)

4.5.1 General Restrictions for Using Earlier Oracle Database Releases

You can use Oracle Database release 10.2 and release 11.1 with Oracle Clusterware and Oracle ASM release 11.2.

Note: Refer to *Oracle Automatic Storage Management Administrator's Guide* for details on configuring disk group compatibility for databases using Oracle Database versions 11.1 or earlier with Oracle Grid Infrastructure 11.2.

If you upgrade an existing release of Oracle Clusterware and Oracle ASM to Oracle Grid Infrastructure release 11.2 (which includes Oracle Clusterware and Oracle ASM), and also upgrade your Oracle RAC database to release 11.2, then the required configuration of existing databases is completed automatically.

If you plan to use Oracle Clusterware and Oracle ASM release 11.2 with Oracle RAC 10.2 or 11.1, then you must complete additional configuration tasks as described in the following sections. You also must apply patches to the Oracle RAC 10.2 or 11.1 software installations before Oracle RAC 10.2 and 11.1 will work correctly with Oracle

Clusterware 11.2 and Oracle ASM 11.2. See the *Oracle Database Readme* for information about specific patches.

Note: If you are upgrading Oracle RAC or Oracle Database from releases 11.1.0.7, 11.1.0.6, or 10.2.0.4, and you have upgraded to Oracle Clusterware and Oracle ASM release 11.2, then Oracle recommends that you check for the latest recommended patches for the release you are upgrading from, and install those patches as needed before upgrading to Oracle RAC or Oracle Database release 11.2.

For more information about recommended patches, see "Oracle Upgrade Companion," which is available through Note 785351.1 on My Oracle Support:

<https://support.oracle.com>

You can also see Oracle Support Notes 756388.1 and 756671.1 for the current list of recommended patches for each release.

See Also: "Downloading and Installing Patch Updates" on page 4-2

4.5.2 Required Patches for Earlier Oracle Database Releases or Databases Being Upgraded

Before you use DBCA to create an Oracle RAC or Oracle Database 10.x or 11.1 database on an Oracle Clusterware and Oracle ASM release 11.2 installation, you must install patches to the Oracle RAC or Oracle Database home. See *Oracle Database Readme* for information about specific patches.

See Also: "Downloading and Installing Patch Updates" on page 4-2

Before applying any patches to your Oracle Database 10g release 2 or Oracle Database 11g release 1 software, you must first stop the `OracleRemExecService` Windows service on all the nodes in your cluster. If you do not stop this process before applying patches to your Oracle Database 10g release 2 or Oracle Database 11g release 1 database, then you will receive errors during the patching operation and the Oracle software will not be patched correctly.

For Oracle Database release 11.2.0.3, if you plan to create or upgrade Oracle databases or release 11.2.0.2 (or earlier) after upgrading to or installing Oracle Grid Infrastructure release 11.2.0.3, then you must first apply the patches listed in [Table 4-2](#).

Table 4-2 Patches Required for Oracle Database and Oracle Grid Infrastructure Release 11.2.0.3

Environment	Patches	Apply to Oracle Grid Home?	Apply to Oracle RAC Home?
Create an Oracle RAC release 11.1 database to run on Oracle Grid Infrastructure 11.2.0.3	8288940 10636465	No	Yes
Upgrade Oracle Grid Infrastructure and Oracle Database release 11.2.0.1 to release 11.2.0.3	10100101	Yes, after creating the Oracle Database release 11.2.0.1 database	Yes, after creating the Oracle Database release 11.2.0.1 database

Table 4–2 (Cont.) Patches Required for Oracle Database and Oracle Grid Infrastructure Release 11.2.0.3

Environment	Patches	Apply to Oracle Grid Home?	Apply to Oracle RAC Home?
Create an Oracle Database release 11.2.0.1 database to run on a fresh install of Oracle Grid Infrastructure 11.2.0.3	10100101	No	Yes, after creating the Oracle Database release 11.2.0.1 database
Starting an Oracle Database 10g Release 2 database (10.2.0.5) running on Oracle Grid Infrastructure release 11.2.0.3	12799034	No	Yes

4.5.3 Pinning Cluster Nodes for Oracle Database Release 10.2 or 11.1

When Oracle Database release 10.2 or 11.1 is installed on a new Oracle Grid Infrastructure for a cluster configuration, it is configured for dynamic cluster configuration, in which some or all Internet Protocol (IP) addresses are provisionally assigned, and other cluster identification information is dynamic. This configuration is incompatible with earlier database releases, which require fixed addresses and configurations.

You can change the nodes where you want to run the older database to create a persistent configuration. Creating a persistent configuration for a node is called **pinning a node**.

To pin a node in preparation for installing an earlier Oracle Database version, use `CRS_home\bin\crsctl` with the following command syntax, where *nodes* is a space-delimited list of the names of one or more nodes in the cluster whose configuration you want to pin:

```
crsctl pin css -n nodes
```

For example, to pin nodes `node3` and `node4`, log in as an Administrator user and enter the following command:

```
C:\> crsctl pin css -n node3 node4
```

To determine if a node is in a pinned or unpinned state, use `CRS_home\bin\olsnodes` with the following syntax:

```
olsnodes -t -n
```

For example, to list all pinned nodes, use the following command:

```
C:\> app\11.2.0\grid\bin\olsnodes -t -n
node1 1      Pinned
node2 2      Pinned
node3 3      Pinned
node4 4      Pinned
```

To list the state of a particular node, use the `-n` option as shown in the following example:

```
C:\> app\11.2.0\grid\bin\olsnodes -t -n node3
node3 3      Pinned
```

See Also: *Oracle Clusterware Administration and Deployment Guide* for more information about pinning and unpinning nodes

Configuring the Server Parameter File in Oracle Real Application Clusters Environments

This chapter describes server parameter file (SPFILE) placement and configuration in Oracle Real Application Clusters (Oracle RAC) environments. This chapter contains the following topics:

- [About Parameter Files and Oracle RAC](#)
- [Creating Server Parameter Files for Oracle RAC](#)
- [Storage Requirements for the Server Parameter File in Oracle RAC](#)
- [Parameter File Search Order in Oracle RAC](#)
- [Migrating to an SPFILE in Oracle RAC Environments](#)
- [Server Parameter File Errors in Oracle RAC](#)

See Also: *Oracle Real Application Clusters Administration and Deployment Guide* for more information about parameters, and for a discussion of parallel execution-related parameters in Oracle RAC data warehouse environments

5.1 About Parameter Files and Oracle RAC

Oracle Database uses parameter settings in parameter files to determine how to control various database resources. You can use two types of files for parameter administration: the SPFILE or one or more traditional client-side parameter files (PFILES). Oracle recommends that you manage parameters using an SPFILE. If you use client-side parameter files, then Oracle does not preserve database or instance parameter changes made for performance tuning after you shut down the instance.

When you create the database, Oracle Database creates an SPFILE. The SPFILE can be stored on an Oracle Automatic Storage Management (Oracle ASM) disk group, a cluster file system, or a shared raw device. The Oracle ASM instance also uses an SPFILE.

All instances in the cluster database use the same SPFILE at startup. Because the SPFILE is a binary file, do not directly edit the SPFILE with an editor. Instead, change SPFILE parameter settings using Oracle Enterprise Manager or `ALTER SYSTEM SQL` statements.

Note: Oracle recommends that you avoid modifying the values for self-tuning parameters; overriding these settings can affect performance adversely.

With Oracle Database 11g release 2 and later databases, Database Agent automatically maintains the SPFILE based on the value you specify in the command `srvctl add/modify database -d db_unique_name -p spfile`, where *spfile* is the full path name of the SPFILE and *db_unique_name* is the unique name for your Oracle RAC database. You do not have to create or maintain `initORACLE_SID.ora` files.

The default location of the database SPFILE is:

```
%ORACLE_HOME%\database\SPFILE%ORACLE_SID%.ORA
```

Using local storage for the default location of the SPFILE is inappropriate for Oracle RAC databases and Oracle ASM instances, because all database instances must use the same SPFILE. See ["Storage Requirements for the Server Parameter File in Oracle RAC"](#) on page 5-2.

See Also: *Oracle Database Administrator's Guide* for more information about managing and modifying parameter files

5.2 Creating Server Parameter Files for Oracle RAC

When DBCA is used to create an Oracle RAC database, DBCA creates the SPFILE in either a cluster file system or in the disk group that is assigned for data files depending on whether data files are placed in a cluster file system or in an Oracle ASM disk group.

If you manually create your database, then Oracle recommends that you create an SPFILE from a PFILE.

If you are upgrading from a previous Oracle Database release, then review the section ["About Parameter Files and Oracle RAC"](#) on page 5-1.

5.3 Storage Requirements for the Server Parameter File in Oracle RAC

All instances in an Oracle RAC environment must use the same SPFILE. Place the SPFILE on shared storage, such as in an Oracle ASM disk group, a cluster file system, or on a network file system (NFS). Otherwise, place the SPFILE on a shared raw device that is at least 5 MB, although this is not recommended.

If you must use a parameter file stored on local storage on each node, then Oracle recommends that you use a PFILE. The PFILE should be located and named as follows:

```
%ORACLE_HOME%\database\init%ORACLE_SID%.ora
```

This path is valid for each instance and it refers to a single, nonshared PFILE. The PFILE should contain an entry similar to one of the following, where *dbname* is the name of the Oracle RAC database:

Shared storage accessed using the Universal Naming Convention (UNC) format:

```
SPFILE='\\server\share_name\dbname\spfile.ora'
```

Shared storage mapped to a network drive:

```
SPFILE='X:\dbname\spfile.ora'
```

The SPFILE entry directs each instance to read from a single, shared SPFILE that is located on shared storage.

5.4 Parameter File Search Order in Oracle RAC

Oracle RAC searches for your parameter file in the following order:

1. %ORACLE_HOME%\database\spfile%ORACLE_SID%.ora
2. %ORACLE_HOME%\database\spfile.ora
3. %ORACLE_HOME%\database\init%ORACLE_SID%.ora

If you do not have a shared Oracle home for your Oracle RAC database, then the default locations for the parameter file are not shared. Therefore, you must create a nondefault SPFILE for use with Oracle RAC.

5.5 Migrating to an SPFILE in Oracle RAC Environments

If your Oracle RAC database does not use an SPFILE, then you can switch to using an SPFILE by creating and editing a single PFILE using the procedures described in this section.

1. Create a single PFILES for all the instances of your Oracle RAC database using the name `initdbname.ora` file, where `dbname` is the name of your Oracle RAC database. The file should be located in the `Oracle_home\database` directory.
2. Copy all instance-specific parameter definitions from the PFILES files using the following syntax, where the variable `sid` is the ORACLE_SID for the instance:

```
sid.parameter=value
```

For example, if your Oracle RAC database is named `sales`, and the values of ORACLE_SID for the instances are `sales1`, `sales2`, and so on, then the instance-specific parameters copied from each PFILE would have the following format:

```
sales1.LICENSE_MAX_USERS=value
sales2.LICENSE_MAX_USERS=value
...
```

3. Copy all shared IFILE content *as is*. All parameters defined in your IFILES are global. Any parameters that have the same value for every instance are also global parameters. Therefore, create entries for these parameters in the format `parameter=value`, without ORACLE_SID prefixes.

Note: The IFILE entry in a PFILE is used to embed another parameter file within the current parameter file. These additional parameter files are called IFILES.

4. Shutdown your Oracle RAC database.
5. Create an SPFILE from the combined PFILE using the `CREATE SPFILE` statement in SQL*Plus, as shown in the following example, where `dbname` is the name of the Oracle RAC database and `X`: represent a shared storage location:

```
CREATE SPFILE='X:\database\spfiledbname.ora'
FROM PFILE='%ORACLE_HOME%\database\initdbname.ora'
```

If you use Oracle ASM for shared storage, then use the following SQL statement to create the SPFILE, where *disk_group_name* is the name of the Oracle ASM disk group, *db_unique_name* is the database unique name, *sid* is the ORACLE_SID for the Oracle RAC database, *oracle_home* is the Oracle home path, and *initdbname.ora* is the file you created in Step 1:

```
CREATE SPFILE='+disk_group_name\db_unique_name\spfiledbname.ora'  
FROM PFILE='%ORACLE_HOME%\database\initdbname.ora'
```

The CREATE SPFILE SQL statement reads the *initdbname.ora* file that you created by merging your PFILES and IFILES, and then transfers the settings for the parameters from the PFILE into an SPFILE.

6. Update the local *initsid.ora* file on each node to contain an SPFILE entry that points to the location of the new SPFILE.

Because the SPFILE is located in a nondefault location, you must start up the database by first using a traditional parameter file that contains a single entry, SPFILE.

7. Start the Oracle RAC database using the STARTUP statement for the database and specify that the local PFILE should be used by the instance. For example, if you have set the environment variables ORACLE_HOME and ORACLE_SID for your Oracle RAC database, then you would use a SQL*Plus statement like the following, where ORACLE_SID is the name of the local Oracle RAC database instance:

```
STARTUP PFILE=%ORACLE_HOME%\database\init%ORACLE_SID%.ora
```

When you use this STARTUP command, then the Oracle Database uses the SPFILE parameter setting in the PFILE to locate the SPFILE in the nondefault location on shared storage.

8. Finally, update the Oracle RAC database configuration so that you do not have to use a local PFILE when starting the database. To update the Oracle RAC database configuration, use the following command to define the location of the new SPFILE, where *db_unique_name* is the database unique identifier and *spfile* is the full path name of the SPFILE (the same value you used in Step 5):

```
%ORACLE_HOME%\bin> srvctl modify database -d db_unique_name -p spfile
```

Note: With Oracle ASM 11g release 2 and later releases, the SPFILE for the Oracle ASM instance can be stored in Oracle ASM. You can also place an Oracle ASM instance SPFILE on a shared file system, or you can use a traditional client-side PFILE for each Oracle ASM instance.

See Also: *Oracle Automatic Storage Management Administrator's Guide* for more information about administering an Oracle ASM instance

5.6 Server Parameter File Errors in Oracle RAC

Oracle Database reports errors that occur during SPFILE creation, or while reading the file during startup. If an error occurs during a parameter update, then Oracle records the error in the database alert log file, and ignores subsequent parameter updates to the file. If this happens, then do either of the following:

- Shut down the instance, recover the SPFILE, and restart the instance.
- Enable the instance to continue running without regard for subsequent parameter updates.

Oracle Database displays errors for parameter changes that you attempt when you incorrectly use the `ALTER SYSTEM SET` statement. Oracle Database does this when an error occurs while reading from or writing to the SPFILE.

See Also: *Oracle Real Application Clusters Administration and Deployment Guide* for more information about backing up the SPFILE

Configuring Server Pools and Oracle Enterprise Manager Database Control

This chapter describes server pool and Oracle Enterprise Manager Database Control configuration in Oracle Real Application Clusters (Oracle RAC) environments.

This chapter contains the following topics:

- [Understanding Server Pools](#)
- [Understanding the Database Control Implementation for Oracle RAC](#)
- [Understanding EMCA and Enterprise Manager Control with Oracle RAC](#)
- [Configuring Database Control During Installation](#)
- [Configuring Database Control with DBCA](#)
- [Configuring Database Control After Installation Using EMCA](#)
- [Managing Database Control Using EMCA](#)
- [Enterprise Manager and EMCA Troubleshooting Tips](#)

6.1 Understanding Server Pools

Server pools are logical divisions of a cluster into groups of servers, which are allocated to host databases or other applications. Server pools are managed using `crsctl` and `srvctl` commands.

Each server pool name must be unique within the cluster. Two server pools cannot use the same name.

Caution: By default, any named user may create a server pool. To restrict the operating system users that have this privilege, Oracle strongly recommends that you add specific users to the CRS Administrators list. See *Oracle Clusterware Administration and Deployment Guide* for more information about adding users to the CRS Administrators list.

There are two types of server pool management:

- **Administrator-managed:** Database administrators define the servers on which database resources run, and place resources manually as needed. This is the management strategy used in previous releases.

- **Policy-managed:** Database administrators specify in which server pool (excluding the generic or free server pools) the database resource will run. Oracle Clusterware is responsible for placing the database resource on a server.

The server pool name is a required attribute. You can also provide values for the following attributes; if you do not specify them, then they are set to the default value:

- **MIN_SIZE:** Minimum number of servers on which you want a resource to run. The default is 0.
- **MAX_SIZE:** Maximum number of servers on which you want a resource to run. The default is -1, which indicates that resources can run on all available nodes in the cluster.
- **IMPORTANCE:** The relative importance of the server pool, used to determine how to reconfigure servers when a node joins or leaves the cluster. The default is 0.

Note: Policy-managed databases that use Oracle Database 11g release 2 or later and are deployed in on or more server pools should access the database using the Single Client Access Name (SCAN). Using SCAN is highly recommended because the instances of the policy-managed database can run on different servers at different times; using a virtual IP (VIP) address to connect to a particular node or instances for a policy-managed database is not possible.

6.1.1 Oracle RAC One Node and Server Pools

Note the following about Oracle RAC One Node and server pools:

- Oracle RAC One Node runs only in one server pool. This server pool is treated the same as any other server pool.
- Online relocation of an Oracle RAC One Node database instances permits planned migrations of an Oracle RAC One Node database from one node to another node. Relocations must always be within a server pool.
- An Oracle RAC One Node database can be a Data Guard physical standby.

6.1.2 Policy-Based Cluster and Capacity Management

Using policy-based server pool management, database administrators define resource requirements for the workload, including the following:

- Defining the number of servers or instances preferred for the application. Oracle Clusterware automatically allocates resources up to the number specified as long as they are physically available, and relocates a resource to another node as necessary.
- Defining the number of instances required to support the application.

Resources are no longer defined as belonging to a specific instance or node. Instead, the priority of resource requirements is defined.

See Also: *Oracle Real Application Clusters Administration and Deployment Guide* for more information about policy-based management and server pools

6.2 Understanding the Database Control Implementation for Oracle RAC

When you configure Database Control for Oracle RAC, you configure the Database Control for each instance in the cluster. However, by default, the Database Control console will only start on the local node, the node on which the Oracle RAC installation was started. On every other node of the cluster, only the Oracle Management Agent starts. This is because the Database Control console opens a number of connections to the database. If an instance of the console is running on every host in the cluster, then you can easily exceed the maximum number of permitted open connections on a 32-node or 64-node environment.

Each of the remote agents upload their respective data to the console running on the local node, where you can monitor and manage all the nodes and instances in the cluster. On each node in the Oracle home of the Oracle RAC database, the following subdirectories are created, where *nodename* is the name of a node in the cluster, and *DBName* is the unique name of the Oracle RAC database.

```
%ORACLE_HOME%\nodename_DBName
%ORACLE_HOME%\nodename_DBName
```

If you upgrade an existing Oracle RAC 10g release 1 database configured with Database Control to the current release, then the existing Database Control configuration is retained. A Database Control console runs on each cluster node associated with the upgraded Oracle RAC database.

See Also: [Section 6.7.1, "Changing the Database Control Implementation for Your Oracle RAC Database"](#)

6.3 Understanding EMCA and Enterprise Manager Control with Oracle RAC

When you use Enterprise Manager Configuration Assistant (EMCA) to configure Oracle Enterprise Manager Database Control for Oracle RAC with the `-cluster` option, you configure Database Control for each instance in the cluster. However, by default, the Database Control console starts only on the local node.

If the Database Control console is started on the local node, then you can use the Enterprise Manager Control (EMCTL) commands `start dbconsole` and `stop dbconsole` to start and stop the Database Control console. The commands `emctl start dbconsole` and `emctl stop dbconsole` start and stop only the Management Agent on the other nodes where the Database Control console is not running.

See Also: [Appendix G, "Oracle Enterprise Manager Configuration Assistant Reference"](#)

6.4 Configuring Database Control During Installation

If you create a database while installing Oracle RAC 11g release 2 (11.2), you have the option of configuring your database so it can be managed by Oracle Enterprise Manager Grid Control or by Oracle Enterprise Manager Database Control.

To select Grid Control as your management option, the Oracle Management Service must be installed on a network host. In addition, the Management Agent must be installed on the host where you are installing the database. Otherwise, the Grid Control option is unavailable and you must instead manage your database with Database Control.

For most of the Oracle RAC 11g release 2 (11.2) installation types, you must choose either Database Control or Grid Control as your management option when you create a database during the installation.

If you do not want to configure Enterprise Manager for your Oracle RAC database, then you must use either the Advanced database configuration option during an Enterprise or Standard Edition installation or run Oracle Database Configuration Assistant (DBCA) after the installation to create your Oracle RAC database. If you do not configure Database Control during the Oracle RAC 11g release 2 (11.2) installation, then the `nodename_DBName` directory is not created in the Oracle home for your database.

See Also: [Section 2.2, "Selecting an Installation Type"](#)

6.5 Configuring Database Control with DBCA

The primary method for configuring an existing Oracle Database 11g release 2 (11.2) database so it can be managed with Database Control is to use DBCA. You can use DBCA to create a new database or to reconfigure an existing database.

See Also:

- "Installing Oracle Software and Building the Database" in *Oracle Database 2 Day DBA* for more information about using DBCA to create a new database instance
- *Oracle Database 2 Day + Real Application Clusters Guide* for Oracle RAC examples using DBCA

To use DBCA to configure your database so it can be managed with Database Control:

1. Log in to the database host as an Administrator user.
2. Click **Start**, and select **Programs**, then **Oracle - Oracle_home name**. In the Configuration and Migration Tools folder, select **Database Configuration Assistant**.
The DBCA Welcome page appears.
3. Click **Next** until the Operations page is displayed. Select **Configure Database Options**.
4. Click **Next** until the Database page is displayed. Select the database you want to configure.
5. Click **Next** until the Management Options page is displayed. Select the following options:
 - **Configure the Database with Enterprise Manager**
 - **Use Database Control for Database Management**
6. Optionally, select the options for enabling daily backups.
For more information about Oracle Enterprise Manager daily backups, click **Help** on the Management Options page.
7. Click **Next** until the Finish button is available, then click **Finish** to configure the database to use Database Control.

After DBCA configures the database, a new subdirectory appears in the Oracle home. This directory contains Database Control configuration and state files specific to the database you just configured. For cluster databases, the directories are named

`nodename_DBName`. For example, if the name of the local node is `node1.example.com` and the database unique name is `myNewDB`, then the directory name is:

```
node1.example.com_myNewDB
```

6.6 Configuring Database Control After Installation Using EMCA

EMCA provides a command-line interface for configuring Database Control. EMCA is available in the `bin` directory of the Oracle home of your database.

Caution: During configuration of Database Control using EMCA, the database may become unavailable and users may not be able to connect to the database or perform operations on the database.

6.6.1 Using EMCA in Interactive Mode

To use EMCA to configure your database for management using Database Control:

1. Set the default Oracle home to the location of the home directory for the database you are modifying. See [Section 7.3, "Working with Multiple Oracle Home Directories on Windows"](#) for details.
2. Change the directory to `%ORACLE_HOME%\bin`.
3. Start EMCA by entering the following command with any of the optional command-line parameters shown in [Table G-1, "EMCA Optional Command-Line Parameters"](#):

```
%ORACLE_HOME%\bin> emca.bat
```

Depending upon the parameters you include on the EMCA command line, you might be prompted for additional information required to configure Database Control.

Note: To configure Database Control for a single-instance database using Oracle Automatic Storage Management (Oracle ASM), you do not have to include extra parameters in the EMCA command. Run the following command to configure Database Control for a new installation. EMCA automatically detects the Oracle ASM instance:

```
emca -config dbcontrol db -repos create
```

Example 6-1 Configuring Database Control and Daily Database Backups using EMCA

Enter the following commands to configure Database Control for your Oracle RAC database and also to enable automatic daily backups of your database:

```
X:\> cd %ORACLE_HOME%\bin
```

```
X:\..\bin> emca.bat -config dbcontrol db -backup
```

See Also: [Appendix G, "Oracle Enterprise Manager Configuration Assistant Reference"](#)

6.6.2 Using EMCA in Silent Mode With an Input File

Instead of answering a series of prompts when you run EMCA, you can use the `-respFile` parameter to specify an input file. The input file you create must be in a format similar to the following example:

```
PORT=1521
SID=DB
DBSNMP_PWD=xpE234D
SYSMAN_PWD=KD0dk432
```

After you create an EMCA input file, you can use it on the command line as follows, where `input_file_path` is the full path name of the input file:

```
emca -config dbcontrol db -respFile input_file_path
```

Example 6–2 Using a Response File with EMCA to Configure Database Control

To configure Database Control to perform daily backups and create the Oracle Management Repository, create an input file similar to the one shown in the following example:

```
PORT=1521
SID=DB
DBSNMP_PWD=dow31224
SYSMAN_PWD=squN3243
HOST_USER=johnson
HOST_USER_PWD=diTf32of
SYS_PWD=q1Kj4352
BACKUP_SCHEDULE=06:30
```

After you have created the input file, enter the following command at the operating system prompt from the `bin` directory of the Oracle home, where `input_file_path` is the full path name of the input file:

```
X:\..\bin> emca -config dbcontrol db -repos create -backup -respFile
input_file_path
```

See Also: [Appendix G, "Oracle Enterprise Manager Configuration Assistant Reference"](#)

6.7 Managing Database Control Using EMCA

You can use DBCA to configure Grid Control or Database Control for your Oracle RAC database. However, if you prefer to use a command-line interface to configure Database Control, then you can use the Oracle Enterprise Manager Configuration Assistant (EMCA) utility.

The following sections describe how Database Control can be configured after installation:

- [Changing the Database Control Implementation for Your Oracle RAC Database](#)
- [Updating Database Control Configuration When Adding and Removing Nodes](#)
- [Using EMCA in Silent Mode With an Input File](#)
- [Changing the Ports Used by Database Control](#)
- [Enterprise Manager and EMCA Troubleshooting Tips](#)

See Also: [Section 6.5, "Configuring Database Control with DBCA"](#)

6.7.1 Changing the Database Control Implementation for Your Oracle RAC Database

By default, when you configure Database Control for an Oracle RAC 11g database, the Database Control console starts only on the local node. If you upgrade an existing Oracle RAC database configured with Database Control to the current release, then the existing Database Control configuration is retained. For example, if you have an Oracle RAC 10g release 1 (10.1) database, then the existing Database Control configuration has a Database Control console running on each cluster node. If you upgrade this database to Oracle Database 11g release 2, then the Database Control console is still started on each individual node.

You can modify the number of Database Control consoles started for your Oracle RAC database. If you have one Database Control console running in the cluster, or more than one, you can still monitor and manage every node in the cluster from a single Database Control console.

To modify the implementation of Database Control for your database, use the following command, where *nodename* is the public name of the node, and *node_list* is a comma-delimited list of node names:

```
emca -reconfig dbcontrol -cluster -EM_NODE nodename -EM_NODE_LIST node_list
```

Caution: If you use the command `emca -cluster` to configure Database Control for Oracle RAC, then check the `TNS_ADMIN` environment variable on all cluster nodes. If you have a different `TNS_ADMIN` set for each node, then the listener for the target node cannot be configured correctly. In this case, set the same `TNS_ADMIN` on all cluster nodes before running the `emca -cluster` command.

You can use the `emca -reconfig dbcontrol` command to start a Database Control console on multiple nodes. This command performs the following actions in sequence:

1. Starts a Database Control console on the node specified by *nodename*, if there a Database Control console is not started.
2. Redirects the Management Agents monitoring the database instances in *node_list* so that they upload their data to the Database Control console running on *nodename*. Also, Management Agents monitoring the Oracle Database instances on *nodename* upload their data to the local console. Note that if you do not enter the command options `-EM_NODE` or `-EM_NODE_LIST` at the command line, you will be prompted for them.

Note: `-EM_NODE` defaults to the local node if not specified when prompted. If not specified, `-EM_NODE_LIST` defaults to all nodes where the database is configured.

For information about the current cluster configuration, you can run the following command:

```
emca -displayConfig dbcontrol -cluster
```

This command prompts for the database unique name for the cluster database. It prints the current configuration to the screen, indicating the nodes that have Database Control consoles running on them and the Database Control consoles to which each Management Agent is uploading its data.

Example 6–3 Reconfiguring Database Control for an Oracle RAC Database

This example demonstrates how to start the Database Control console on more than one node in a cluster. If you have an eight-node cluster with nodes `node1`, `node2`, `node3`, `node4`, `node5`, `node6`, `node7`, and `node8`, then you can run the following commands in succession:

```
emca -reconfig dbcontrol -cluster -EM_NODE node1 -EM_NODE_LIST node2,node3,node4
emca -reconfig dbcontrol -cluster -EM_NODE node5 -EM_NODE_LIST node6,node7,node8
```

After you execute these commands, there are two Database Control consoles running, one on `node1` and the other on `node5`. From either of these consoles, you can manage and monitor all of the targets in the cluster. Additionally, the Management Agents on nodes `node1`, `node2`, `node3`, and `node4` upload their information to the Database Control console running on `node1`. Similarly, the Management Agents on nodes `node5`, `node6`, `node7`, and `node8` upload their information to the Database Control console running on `node5`.

See Also: [Appendix G, "Oracle Enterprise Manager Configuration Assistant Reference"](#)

6.7.2 Updating Database Control Configuration When Adding and Removing Nodes

For an administrator-managed database, a common operation is the addition and deletion of cluster nodes.

6.7.2.1 Using EMCA After Adding Nodes to the Cluster

After you add a new node, run EMCA to configure Oracle Enterprise Manager so that you can manage the newly added instance in a way consistent with the rest of the cluster database instances. Use EMCA to configure a Database Control console or Management Agent for that node by executing the following command on a node that currently has Enterprise Manager configured:

```
emca -addNode db
```

Because the `-addNode` option can be used only in an Oracle Real Application Clusters environment, you do not have to use the `-cluster` option on the command line. After running the command `emca -addNode db`, enter the node name and database unique name for the node and Oracle RAC database, for example:

```
Node name: node6
Database Unique Name: myNewDB
```

When configuring Oracle Enterprise Manager for a new node, run the EMCA command only after you have added the node to your cluster and Oracle RAC database. Also, run the command from a node in the cluster that has Oracle Enterprise Manager configured, because the Oracle Enterprise Manager configuration settings on the existing node are propagated to the new node when you run the command.

See Also: [Appendix G, "Oracle Enterprise Manager Configuration Assistant Reference"](#)

6.7.2.2 Using EMCA After Removing Nodes from the Cluster

To deconfigure Oracle Enterprise Manager for a specific node you are deleting from the cluster, log in to a different node from the one you plan to delete, and use the following command:

```
emca -deleteNode db
```

Ensure that you run the EMCA command before you delete the actual cluster node. Because this option can be used only in an Oracle Real Application Clusters environment, you do not have to use the `-cluster` option on the command line.

Running this command removes the Oracle Enterprise Manager configuration so that you can not use Oracle Enterprise Manager to manage any instances on that node.

See Also: [Appendix G, "Oracle Enterprise Manager Configuration Assistant Reference"](#)

6.7.3 Changing the Ports Used by Database Control

When you initially install Oracle RAC 11g release 2 (11.2) or configure Database Control with EMCA, Database Control uses a set of default system ports. For example, by default, you access Database Control in Oracle Database 11g release 2 (11.2) using a web address that consists of the local host name (*host.domain*) and port 1158:

```
https://host.domain:1158/em
```

Port 1158 is the default port assigned to Database Control by the Internet Assigned Numbers Authority (IANA). Similarly, the default Management Agent port, as assigned by the IANA, is 3938.

To use ports other than the default ports, use EMCA command-line arguments when configuring Database Control using EMCA. [Table 6–1](#) summarizes the EMCA command-line parameters that control the standard Database Control port assignments.

Table 6–1 EMCA Command Parameters for Configuring Port Assignments

Option	Description
<code>-DBCONTROL_HTTP_PORT port_number</code>	<p>This port number is used in the Database Control console URL. For example, if you set this port to 5570, then you can then access the Database Control console using the following URL:</p> <pre>http://host.domain:5570/em</pre>
<code>-RMI_PORT port_number</code>	<p>This port number is used by the Remote Method Invocation (RMI) system, which is part of the Java EE software required by Database Control. The default port can be changed if the user wants to configure a specific port for Database Control console.</p> <p>When a port other than the default port (1521) is used, use the <code>-RMI_PORT</code> parameter along with the <code>emca reconfig</code> command.</p>
<code>-JMS_PORT port_number</code>	<p>This port is used by the Oracle Containers for Java EE (OC4J) and Java Message Service (JMS), which is part of the Java EE software required by Database Control. The default port can be changed if the user wants to configure a specific port for Database Control console.</p> <p>When a port other than the default port (1521) is used, use the <code>-JMS_PORT</code> parameter along with the <code>emca reconfig</code> command.</p>

Table 6–1 (Cont.) EMCA Command Parameters for Configuring Port Assignments

Option	Description
<code>-AGENT_PORT port_number</code>	This port is used by Oracle Management Agent, which monitors and manages the database for the Database Control console.

Alternatively, you can explicitly assign ports after configuring Database Control using the following command, where *parameters* refers to the command-line parameters listed in [Table 6–1](#):

```
emca -reconfig ports [-cluster] [parameters]
```

See Also: [Appendix G, "Oracle Enterprise Manager Configuration Assistant Reference"](#)

6.8 Enterprise Manager and EMCA Troubleshooting Tips

The following topics describe some troubleshooting tips to consider when using EMCA and Database Control or Grid Control:

- [Updating Database Control After Changing the Database Listener Port](#)
- [Upgrading Grid Control Agents for Oracle Database 11g Release 2](#)
- [Updating Database Control After Changing the Host Name or IP Address](#)
- [Using EMCA After Changing the TNS_ADMIN Configuration](#)

See Also: [Appendix G, "Oracle Enterprise Manager Configuration Assistant Reference"](#)

6.8.1 Updating Database Control After Changing the Database Listener Port

If you change the listener port of the database after you have configured Database Control, then the database status appears as down in the Database Control console. To reconfigure Database Control to use the new listener port, run EMCA using the following command:

```
emca.bat -config dbcontrol db [-cluster]
```

6.8.2 Upgrading Grid Control Agents for Oracle Database 11g Release 2

When upgrading an Oracle Database 10g release 1 database instance that was configured for Oracle Enterprise Manager (either Database Control or a Grid Control Management Agent) to an Oracle Database 11g release 2 (11.2) database, all resources monitored by Oracle Enterprise Manager targets associated with the upgraded instances are updated automatically. This is because the upgrade involves altering the Oracle home for the instance, the port numbers, or other instance- and node-associated properties.

If an upgraded instance is managed by a Grid Control Agent, then the associated resource information in Grid Control might not be updated successfully during the upgrade. To update the Grid Control target information, perform the following steps:

1. Use Grid Control to access the home page of the upgraded instance.
2. On the home page for the upgraded instance, select **Monitoring Configuration**.

3. On the Monitoring Configuration page, update the properties to the correct values as needed, for example, Oracle home, listener port, and so on.

6.8.3 Updating Database Control After Changing the Host Name or IP Address

After changing the host name (including the domain name) or the IP address for a node, you must deconfigure Database Control on that node and then reconfigure it using EMCA with the `-repos` flag, as shown in the following examples:

Reconfigure Enterprise Manager, deleting the repository when removing the configuration

```
emca -deconfig dbcontrol db -repos drop
emca -config dbcontrol db -repos create
```

Reconfigure Enterprise Manager, re-creating the repository during reconfiguration

```
emca -deconfig dbcontrol db
emca -config dbcontrol db -repos recreate
```

6.8.4 Using EMCA After Changing the TNS_ADMIN Configuration

After you change the `TNS_ADMIN` environment variable to specify a new configuration, you must update Database Control. Set the `TNS_ADMIN` environment variable to the new value, then run the following command:

```
emca -config dbcontrol db
```

Understanding the Oracle Real Application Clusters Installed Configuration

This chapter describes the configuration of an Oracle Real Application Clusters (Oracle RAC) installation.

This chapter contains the following topics:

- [Understanding the Oracle RAC Configuration](#)
- [Understanding Time Zone Settings on Cluster Nodes](#)
- [Working with Multiple Oracle Home Directories on Windows](#)
- [Understanding the Database Components Created by DBCA](#)
- [Managing Undo Tablespaces in Oracle RAC](#)
- [Using Initialization Parameter Files](#)
- [Oracle Database Vault Default Audit Policy and Initialization Parameters](#)
- [Understanding Oracle Net Services Configuration Files](#)
- [Using SCAN with Oracle RAC Databases](#)

7.1 Understanding the Oracle RAC Configuration

Oracle Net Configuration Assistant (NETCA) and Oracle Database Configuration Assistant (DBCA) configure your environment to meet the requirements for database creation and to enable discovery of Oracle RAC databases by Oracle Enterprise Manager.

Note: Configuration files are created on each node of your cluster database.

Avoid changing host names after you complete the Oracle RAC installation, including adding or deleting domain qualifications. Node names are created from the host names during an Oracle Clusterware installation and are used extensively with database processes. Nodes with changed host names must be deleted from the cluster and added back with the new host names.

7.2 Understanding Time Zone Settings on Cluster Nodes

Oracle RAC requires that all cluster nodes have the same time zone setting. During an Oracle Clusterware installation, the installation process determines the time zone

setting of the owner of the Oracle Grid Infrastructure installation on the node where Oracle Universal Installer (OUI) runs, and uses that time zone value on all of the nodes as the default time zone setting for all processes that Oracle Clusterware manages. This default setting is used for databases, Oracle Automatic Storage Management (Oracle ASM), and any other managed processes.

If you use SQL*Plus to start an instance, then you must ensure that the time zone value used by Oracle RAC is the same as the time zone value configured in Oracle Clusterware for the database. You can change the time zone that Oracle Clusterware uses for a database by running the following command, where *time zone* is the time zone to which you want to change:

```
srvctl setenv database -T "TZ=time zone"
```

7.3 Working with Multiple Oracle Home Directories on Windows

Each Oracle product should be installed in its own Oracle home. In other words, each product should be installed into a different directory structure from other Oracle products. The value for %ORACLE_BASE% is stored in the registry (for example, in HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE\HOME0). The values for ORACLE_HOME and ORACLE_SID are also stored in the registry. Symbolic links for these directories, like those used on UNIX platforms, are not supported on Windows platforms.

See Also: *Oracle Database Platform Guide for Microsoft Windows* for more information about the registry keys

Oracle Universal Installer (OUI) records the values for environment variables such as ORACLE_BASE, ORACLE_HOME, and ORACLE_SID in the registry and also updates the value for the PATH environment variable for the user performing the installation. In Linux and UNIX systems, you must manually set these environment variables in the user session or user profile.

To change the current setting for %ORACLE_HOME% (the default Oracle home), perform the following steps:

1. Run OUI, release 10.*n* or later.
2. Click the **Installed Products** button.
3. Click the **Environment** tab at the top of the window.
4. Move the Oracle home directory that you want as your default to the top of the list.
5. Apply the changes, and exit the installer.

This procedure changes the value of the default ORACLE_HOME variable in the registry to the value you selected. It also ensures that the %ORACLE_HOME%\bin directories for each product are listed in the correct order in your PATH environment variable.

7.4 Understanding the Database Components Created by DBCA

The database components that DBCA creates include:

- [Tablespaces and Data Files](#)
- [Control Files](#)
- [Online Redo Log Files](#)

7.4.1 Tablespaces and Data Files

In both single-instance and cluster database environments, an Oracle database is divided into smaller logical areas of space known as **tablespaces**. Each tablespace corresponds to one or more data files on the shared storage. [Table 7-1](#) shows the tablespace names used by an Oracle RAC database and the types of data they contain.

Table 7-1 Tablespace Names Used with Oracle Real Application Clusters Databases

Tablespace Name	Contents
SYSTEM	A mandatory tablespace that consists of the data dictionary, including definitions of tables, views, and stored procedures needed by the database. Oracle Database automatically maintains information in this tablespace.
SYSAUX	A mandatory, auxiliary system tablespace that is used by many Oracle Database features and products. This tablespace contains content that was previously stored in the DRSYS, CWMLITE, XDB, ODM, OEM_REPOSITORY, and SYSTEM tablespaces.
USERS	An user-created tablespace that consists of application data. As you create and enter data into tables, Oracle Database fills this space with your data.
TEMP	A mandatory tablespace that contains temporary tables and indexes created during SQL statement processing. You may have to expand this tablespace if you run SQL statements that involve significant sorting, such as <code>ANALYZE COMPUTE STATISTICS</code> on a very large table, or the constructs <code>GROUP BY</code> , <code>ORDER BY</code> , or <code>DISTINCT</code> .
UNDOTBS n	System-managed tablespaces that contain undo data for each instance. Each Oracle RAC instance uses a different value for n in the tablespace name. These tablespaces are used for automatic undo management.
RBS	A system tablespace that contains rollback segments. If you do not use automatic undo management, then you must configure the RBS tablespace. The RBS tablespace should only be used when needed for compatibility with earlier versions of Oracle Database.

See Also: *Oracle Database Administrator's Guide* for more information about the SYSTEM, SYSAUX, and other tablespaces

You cannot alter these tablespace names if you select a preconfigured database templates, or database configuration type when installing Oracle RAC. However, you can change the names of the tablespaces if you select the Advanced database configuration type during the Oracle RAC installation. See "[Selecting an Oracle Database Configuration Type](#)" on page 2-3 for more information about database configuration types.

As mentioned, each tablespace has one or more data files. The data file names created by the preconfigured database configuration types vary by storage type such as Oracle Automatic Storage Management (Oracle ASM), or a file system.

Windows platforms use symbolic link names for the data files and other database files as shown in [Table 7-2](#). You can specify different symbolic link names when you select the Advanced database configuration type during installation.

Table 7-2 Symbolic Link Names for Database Files on Windows Platforms

Windows Symbolic Link Name	Tablespace or Other Database File Name
<code>db_name_system</code>	SYSTEM
<code>db_name_SYSAUX</code>	SYSAUX
<code>db_name_users</code>	USERS
<code>db_name_temp</code>	TEMP

Table 7–2 (Cont.) Symbolic Link Names for Database Files on Windows Platforms

Windows Symbolic Link Name	Tablespace or Other Database File Name
<code>db_name_undotbs1</code>	UNDOTBS1
<code>db_name_undotbs2</code>	UNDOTBS2
<code>db_name_rbs</code>	RBS
<code>db_name_example</code>	EXAMPLE
<code>db_name_indx</code>	INDX
<code>db_name_spfile</code>	SPFILE
<code>db_name_control1</code>	First control file
<code>db_name_control2</code>	Second control file
<code>db_name_redoinstance_number_log_number</code>	Online redo log files
Where <i>instance_number</i> is the number of the instance and <i>log_number</i> is the redo log number (1 or 2).	Each database instance must have at least two online redo log files. If the database name is <code>db</code> , then the symbolic link names for online redo log files for the first database instance are:
	<code>db_redo1_1</code>
	<code>db_redo1_2</code>
	The link names for the online redo log files for the second instance are:
	<code>db_redo2_1</code>
	<code>db_redo2_2</code>

7.4.2 Control Files

The database is configured with two control files that must be stored on shared storage. Every database must have one unique control file; any additional control files configured for the database are identical copies of the original control file.

If a control file becomes unusable, then the database instance fails when it attempts to access the damaged control file. By multiplexing (creating multiple copies of) a control file on different disks, the database can achieve redundancy and thereby avoid a single point of failure.

See Also: *Oracle Database Concepts* for more information about the control file

7.4.3 Online Redo Log Files

Each database instance must have at least two online redo log files. The online redo log files for a database instance are called the **redo thread**. Each Oracle RAC database instance has its own redo thread to avoid contention for a single set of online redo log files. In case of instance failure, the online redo log files must be accessible by the surviving instances. Therefore, the online redo log files for an Oracle RAC database must be placed on shared storage or Oracle ASM. If you use a file system for storage, then the file system must be a shared or cluster file system.

The generated file names created by the preconfigured database configuration types for the online redo log files can vary depending on the type of storage used. Unless you are using a cluster file system or Oracle ASM, during installation or database creation, you must specify a raw device name for the storage location of each online redo log file.

See Also:

- *Oracle Database Concepts* for more information about the online redo log files
- *Oracle Real Application Clusters Administration and Deployment Guide* for more information about storage for online redo log files

7.5 Managing Undo Tablespaces in Oracle RAC

Oracle Database stores rollback or undo information in undo tablespaces. To manage undo tablespaces, Oracle recommends that you use **automatic undo management**, which is an automated management mode for the undo tablespace that makes the undo tablespace easier to administer.

When Oracle ASM and Oracle Managed Files (OMF) are used along with automatic undo management, an instance that is started for the first time, and thus does not have an undo tablespace, has its undo tablespace created for it by another instance automatically. The same is also true for online redo logs.

See Also:

- *Oracle Database Administrator's Guide* for more information about automatic undo management
- *Oracle Real Application Clusters Administration and Deployment Guide* for more information about managing undo tablespaces

7.6 Using Initialization Parameter Files

Oracle recommends that you use a server parameter file (SPFILE) for storing Oracle Database initialization parameters. Oracle recommends that you store all SPFILES on Oracle ASM, including the Oracle ASM SPFILE. SPFILES must be located on shared storage so that all instances in a cluster database can access this parameter file.

See Also: [Chapter 5, "Configuring the Server Parameter File in Oracle Real Application Clusters Environments"](#) for more information about the creation and use of parameter files

7.7 Oracle Database Vault Default Audit Policy and Initialization Parameters

Oracle Database Vault installs a baseline database auditing policy. This policy covers the access control configuration information stored in Oracle Database Vault tables, information stored in the Oracle data dictionary (rollback segments, tablespaces, and so on), the use of system privileges, and Oracle Label Security configuration. When you install Oracle Database Vault, the security specific database initialization parameters are initialized with default values.

See Also: *Oracle Database Vault Administrator's Guide* for more information about how installing Oracle Database Vault affects the Oracle RAC database configuration

7.8 Understanding Oracle Net Services Configuration Files

Networking elements for the Oracle Database server and clients are preconfigured for most environments. The Easy Connect naming method is enabled by default and does

not require a repository. If you use a naming method other than Easy Connect, then additional configuration of Oracle Net Services may be required.

The following sections the Oracle Net Services configuration for an Oracle RAC database:

- [Configuring the Listener Configuration File \(listener.ora\)](#)
- [Net Service Names \(tnsnames.ora File\)](#)
- [Net Services Profile \(sqlnet.ora File\)](#)

7.8.1 Database Services for an Oracle RAC Database

Each database is represented by one or more services. A service is identified by a service name, for example, `sales.example.com`. A client uses a service name to identify the database it must access. The information about the database service and its location in the network is transparent to the client.

The information needed to use a service name to create a database connection is stored in a repository, which is represented by one or more naming methods. A **naming method** is a resolution method used by a client application to resolve a connect identifier (such as the service name) to a connect descriptor when attempting to connect to a database service. Oracle Net Services offers several types of naming methods that support localized configuration on each client, or centralized configuration that can be accessed by all clients in the network.

An Oracle Database 11g release 2 (11.2) database service automatically registers with the listeners specified in the `LOCAL_LISTENER` and `REMOTE_LISTENER` parameters. During registration, PMON sends information such as the service name, instance names, and workload information to the listeners. This feature is called **service registration**.

Services coordinate their sessions by registering their **workload**, or the amount of work they are currently handling, with the local listener and the SCAN listeners. Clients are redirected by the SCAN listener to a local listener on the least-loaded node that is running the instance for a particular service. This feature is called **load balancing**. The local listener either directs the client to a dispatcher process (if the database was configured for shared server), or directs the client to a dedicated server process.

When a listener starts after the Oracle instance starts, and the listener is available for service registration, registration does not occur until the next time the Oracle Database process monitor (PMON) starts its discovery routine. By default, the PMON discovery routine is started every 60 seconds. To override the 60-second delay, use the SQL statement `ALTER SYSTEM REGISTER`. This statement forces PMON to register the service immediately.

Note: Oracle recommends that you create a script to run this statement immediately after starting the listener. If you run this statement when the instance is registered are services are currently registered, or while the listener is down, then the statement has no effect.

See Also: *Oracle Database Net Services Administrator's Guide* for more information about service registration

7.8.2 Configuring the Listener Configuration File (listener.ora)

You can configure Oracle Net Services and listeners as described in the following sections:

- [Grid Naming Service and Listeners](#)
- [Local Listeners](#)
- [Configuring Multiple Listeners for Your Database Using DBCA](#)
- [How Oracle Database Uses the Listener File \(listener.ora\)](#)

7.8.2.1 Grid Naming Service and Listeners

If you enable GNS, then you do not have to manually configure the listener.

In GNS configurations, there is one GNS daemon for the cluster. Service requests to the cluster domain that GNS manages are routed to the GNS VIP address, which routes these requests to the GNS daemon.

The GNS daemon listens for registrations. When a SCAN VIP starts on a node, it registers its addresses with GNS. When GNS receives a request from a DNS for the SCAN, it returns the registered addresses to the DNS.

7.8.2.2 Local Listeners

Starting with Oracle Database 11g release 2 (11.2), the local listener, or default listener, is located in the Grid home when you have Oracle Grid Infrastructure installed. The `listener.ora` file is located in the `Grid_home\network\admin` directory. You can use a non-default location for the `listener.ora` file by setting the `TNS_ADMIN` environment variable or registry value to point to the directory that contains the Oracle Net Services configuration files.

During Oracle Database creation, the `LOCAL_LISTENER` parameter is automatically configured to point to the local listener for the database. The Database Agent process (`oraagent.exe`, previously known as `raccimon`) sets the `LOCAL_LISTENER` parameter to a connect descriptor that does not require an Oracle Net service name.

You can set a value manually for `LOCAL_LISTENER`. If you set `LOCAL_LISTENER`, then the Database Agent process does not automatically update this value. Oracle recommends that you leave the parameter unset so that the Database Agent process can maintain it automatically. If you do not set `LOCAL_LISTENER`, then the Database Agent process automatically updates the database associated with the local listener in the Grid home, even when the ports or IP address of that listener are changed.

See Also:

- ["Net Service Names \(tnsnames.ora File\)"](#) on page 7-9 for more information about listener associations defined in the `tnsnames.ora` file
- *Oracle Database Net Services Reference* for more information about the `listener.ora` file
- *Oracle Database Net Services Administrator's Guide* for information about understanding and configuring listeners

7.8.2.3 How Oracle Database Uses the Listener File (listener.ora)

Local listeners are configured to respond to database connection requests, and to nondatabase connection requests, such as external procedures or Oracle XML Database (XDB) requests. The `listener.ora` file is the configuration file for a listener.

It can include the protocol addresses it is accepting connection requests on, a list of the database and nondatabase services it is listening for, and control parameters used by the listener. You can modify the configuration of the listeners used by Oracle Clusterware and Oracle RAC with Server Control Utility (SRVCTL) commands, or by using NETCA. Manual editing of the `listener.ora` file is not required.

Before you install Oracle RAC, during the Oracle Grid Infrastructure installation, NETCA creates and starts a default listener in the Grid home called `LISTENER`. The listener is configured with default protocol listening addresses. The listener is configured to respond to connection requests that are directed to one protocol address specified during installation.

During the Oracle RAC installation, the Oracle RAC database uses the listener in the Grid home, and configures service information about the Oracle RAC database. The database services are registered automatically and do not require configuration in the `listener.ora` file. Dynamic service registration eliminates the need for static configuration of database services. However, static service configuration is required if you plan to use Oracle Enterprise Manager.

Each listener is configured with one or more protocol addresses that specify its listening endpoints. The protocol address defines the protocol the listener listens on, and any other protocol specific information, such as the address, host, and port. Starting with Oracle Database 11g release 2, the `listener.ora` file now contains only an IPC key and the following information:

```
(ADDRESS = (PROTOCOL=TCP) (HOST=) (PORT=1521))
```

In the previous example, the protocol `ADDRESS` refers implicitly to the `HOST` endpoint of the local node. The `listener.ora` file is the same on every node for an Oracle RAC database. Listening endpoints, such as the port numbers, are dynamically registered with the listener.

Example 7-1 Example `listener.ora` File for an Oracle RAC Node

The following is an example `listener.ora` file as it would appear after installation, with an entry for a node named `node1` and a `SCAN` listener.

```
LISTENER_SCAN1=(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=IPC) (KEY=LISTENER_SCAN1))))
# line added by Agent
LISTENER_NODE1=(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=IPC)
(KEY=LISTENER))))
# line added by Agent
# listener.ora.mycluster Network Configuration File:
C:\app\oracle\product\11.2.0\dbhome_1\network\admin\listener.ora.mycluster
# Generated by Oracle configuration tools.

LISTENER_NODE1 =
  (DESCRIPTION_LIST =
    (DESCRIPTION =
      (ADDRESS = (PROTOCOL = IPC) (KEY = EXTPROC1521))
    )
  )

ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER_NODE1=ON# line added by Agent
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER_SCAN2=ON# line added by Agent
ENABLE_GLOBAL_DYNAMIC_ENDPOINT_LISTENER_SCAN1=ON# line added by Agent
```

7.8.2.4 Configuring Multiple Listeners for Your Database Using DBCA

It is possible to configure multiple listeners, each with unique name, in one listener.ora file. Multiple listener configurations are possible because each of the top-level configuration parameters has a suffix of the listener name or is the listener name itself.

Note: Oracle recommends running only one listener for each node in most customer environments.

You can use DBCA to configure which listeners your Oracle RAC database registers with. If DBCA detects more than one listener on a node, then it displays a list of the listeners. You can select one or all of these listeners with which to register your database.

Note: To administer Oracle Database 11g release 2 (11.2) local and SCAN listeners using the `lsnrctl` command, set your `ORACLE_HOME` environment variable to the path for the Grid home. Do not attempt to use the `lsnrctl` commands from Oracle home locations for earlier releases, because they cannot be used with Oracle Database 11g release 2 (11.2).

7.8.3 Net Service Names (tnsnames.ora File)

A `tnsnames.ora` file is created on each node and acts as a repository of net service names. Each net service name is associated with a connect identifier. A **connect identifier** is an identifier that maps a user-defined name to a connect descriptor. A connect descriptor contains the following information:

- The network route to the service, including the location of the listener through a protocol address
- The `SERVICE_NAME` parameter, with the value set to the name of a database service

Note: The `SERVICE_NAME` parameter you use in the `tnsnames.ora` file is singular, because you can specify only one service name. There is also the `service_names` parameter in the database initialization file. The `service_names` database parameter defaults to the global database name, a name comprising the `db_name` and `db_domain` parameters in the initialization parameter file.

Example 7-2 Example tnsnames.ora File

The following is an example of a `tnsnames.ora` file. It contains a single net service name entry. The connect identifier for the net service name is `sales`. The database service associated with the `sales` connect identifier is `myApp`. The connect descriptor is the information in the line that starts with `ADDRESS`.

```
# tnsnames.ora Network Configuration File:
C:\app\grid\11.2.0\network\admin\tnsnames.ora
# Generated by Oracle configuration tools.

sales =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = myscan.mycluster.example.com)(PORT = 1522))
    (CONNECT_DATA =
```

```

        (SERVER = DEDICATED)
        (SERVICE_NAME = myApp)
    )
)

```

The `tnsnames.ora` file is located in both the `Grid_home\network\admin` and `Oracle_home\network\admin` directories. By default, the `tnsnames.ora` file is read from the Grid home when Oracle Grid Infrastructure is installed.

Example 7-3 Example `tnsnames.ora` Entry for an Oracle RAC Database

The following is an example of a net service name stored in the `tnsnames.ora` file for an Oracle RAC 11g release 2 (11.2) database after installation, where the database name is `RACDB`, the cluster name is `mycluster`, and GNS is configured for the cluster:

```

RACDB =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST=myscan.mycluster.example.com) (PORT=1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = RACDB)
    )
  )
)

```

With Oracle Clusterware 11g release 2 and later, the listener association no longer requires `tnsnames.ora` file entries. The listener associations are configured as follows:

- DBCA no longer sets the `LOCAL_LISTENER` parameter. The Oracle Clusterware agent that starts the database sets the `LOCAL_LISTENER` parameter dynamically, and it sets it to the actual value, not an alias. So `listener_alias` entries are no longer needed in the `tnsnames.ora` file.
- The `REMOTE_LISTENER` parameters is configured by DBCA to reference the `SCAN` and `SCAN` port, without any need for a `tnsnames.ora` entry. Oracle Clusterware uses the Easy Connect naming method with `scanname:scanport`, so no listener associations for the `REMOTE_LISTENER` parameter are needed in the `tnsnames.ora` file.

See Also: *Oracle Database Net Services Administrator's Guide* for more information about the `tnsnames.ora` file

DBCA creates net service names for connections as described in the following sections:

- [Net Service Names for Database Connections](#)
- [Net Service Name for Instance Connections](#)

7.8.3.1 Net Service Names for Database Connections

Clients that connect to any instance of the database should use `SCAN` in the connect descriptor. You could also use a net service name to connect to the database. The net service name created by DBCA enables Oracle Enterprise Manager to discover an Oracle RAC database.

Three `SCAN` addresses are configured for the cluster, and allocated to servers. When a client issues a connection request using `SCAN`, the three `SCAN` addresses are returned to the client. If the first address fails, then the connection request to the `SCAN` name fails over to the next address. Using multiple addresses allows a client to connect to an instance of the database even if the initial instance has failed.

Example 7-4 Sample Net Service Name Entry for a Database Connection

This example shows a connect descriptor that is used in a `tnsnames.ora` file. The connect identifier in this case is the same as the database name, `mycluster.example.com`. Instead of specifying an individual server, virtual Internet Protocol (VIP) address, or cluster node name, the connect descriptor uses SCAN, which is `myscan.mycluster.example.com`.

Connections that use the net service name `mycluster.example.com` are connected to any of the database instances of the `mycluster` database that runs the `myApp` database service.

```
mycluster.example.com =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = host=myscan.mycluster.example.com)
      (PORT = 1522))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = myApp)
    )
  )
```

The net service name does not need to know the physical address of the server on which the database, database instance, or listener runs. SCAN is resolved by DNS, which returns three IP addresses to the client. The client then tries each address in succession until a connection is made.

Example 7-5 Connecting to Oracle RAC Using the Easy Connect Naming Method

Clients can use SCAN and the Easy Connect naming method to connect to an Oracle RAC database without configuring the `tnsnames.ora` file. If the Oracle RAC database runs on a cluster for which SCAN is `myscan.mycluster.example.com`, then you could use a connection request for the database service `RACDB.example.com` using a connect descriptor similar to the following:

```
username/@myscan.mycluster.example.com:1521/RACDB.example.com
```

7.8.3.2 Net Service Name for Instance Connections

Clients that connect to a particular instance of the database use the net service name for the instance. The `tnsnames.ora` entry in [Example 7-6](#) enables Oracle Enterprise Manager to discover the database instances in the cluster. These entries are also used to start and stop instances.

Example 7-6 Sample Net Service Name Entry for an Instance Connection

This example shows a net service name that is used to connect to a specific database instance. The connect identifier is the same as the instance name, `mycluster1.example.com`. The connect descriptor uses SCAN to locate the instance. Connections that use the net service name `mycluster1.example.com` are connected to the `mycluster1` database instance of the `mycluster` database, without needing to know on which node the instance is running.

```
mycluster1.example.com=
  (DESCRIPTION=
    (ADDRESS=(PROTOCOL=TCP) (HOST=myscan.mycluster.example.com) (PORT=1521))
    (CONNECT_DATA=
      (SERVICE_NAME=mycluster.example.com)
      (INSTANCE_NAME=mycluster1)
    )
  )
```

)

7.8.4 Net Services Profile (sqlnet.ora File)

OUI starts NETCA after the database is created. NETCA creates the Oracle Net Services profile, or the `sqlnet.ora` file. In an Oracle Grid Infrastructure installation, the `sqlnet.ora` file is located in the following directory by default:

```
Grid_home\network\admin
```

In the `%ORACLE_HOME%\network\admin` directory, in addition to a default `sqlnet.ora` file, you also can find a sample `sqlnet.ora` file in the subdirectory `sample`.

During installation, NETCA creates the following entries in the `sqlnet.ora` file, where `%ORACLE_BASE%` is the path to the Oracle base directory for the Oracle RAC installation:

```
SQLNET.AUTHENTICATION_SERVICES= (NTS)
NAMES.DIRECTORY_PATH=(TNSNAMES, EZCONNECT)
ADR_BASE =%ORACLE_BASE%
```

The `AUTHENTICATION_SERVICES` parameter specifies the method by which users are authenticated for database access. The value `NTS` indicates that Microsoft Windows native operating system authentication should be used to authorize access to the database. The parameter `NAMES.DIRECTORY_PATH` specifies the priority order of the naming methods to use to resolve connect identifiers to connect descriptors. The `ADR_BASE` parameter specifies the base directory into which tracing and logging incidents are stored when Automatic Diagnostic Repository (ADR) is enabled for the database.

The Easy Connect naming method eliminates the need to lookup service names in the `tnsnames.ora` files for TCP/IP environments. With Easy Connect, clients use a connect string for a simple TCP/IP address, consisting of a host name and optional port and service name. If you use this method, then no naming or directory system is required. See [Example 7-5, "Connecting to Oracle RAC Using the Easy Connect Naming Method"](#) for an example.

See Also:

- *Oracle Database Net Services Administrator's Guide* for more information about the `sqlnet.ora` file
- *Oracle Database Concepts* for more information about authentication
- *Oracle Database Client Installation Guide for Microsoft Windows* for more information about configuring clients for database connectivity
- *Oracle Database Administrator's Guide* for more information about ADR

7.9 Using SCAN with Oracle RAC Databases

Oracle RAC databases provide the important benefits of connection load balancing and failover. Standalone Oracle databases perform load balancing by distributing connections among the shared server dispatcher processes.

Oracle RAC databases increase load balancing by using the single client access name (SCAN) for the cluster to balance connections among all instances in a cluster database. SCAN also adds location independence for the databases, so that client

configuration does not have to depend on which nodes are running a particular database or database instance.

Review the following sections to understand Oracle RAC load-balancing features:

- [Understanding SCAN](#)
- [Understanding SCAN Listeners and Listener Registration](#)
- [Understanding Failover, Load Balancing, and SCAN](#)
- [Understanding Shared Server Use in an Oracle RAC Database](#)

7.9.1 Understanding SCAN

SCAN is a fully qualified name (*host name.domain name*) that is configured to resolve to all the addresses allocated for the SCAN listeners. The default value for SCAN is *cluster_name.GNS_sub_domain*, or *cluster_name-scan.domain_name* if GNS is not used. For example, in a cluster that does not use GNS, if your cluster name is `salesRAC`, and your domain is `example.com`, then the default SCAN address is `salesRAC-scan.example.com:1521`.

SCAN is configured in DNS to resolve to three IP addresses, and DNS should return the addresses using a round-robin algorithm. This means that when SCAN is resolved by DNS, the IP addresses are returned to the client in a different order each time.

Note: Configuring SCAN in DNS using a round-robin algorithm does not ensure failover of connections. Instead, failover of connections is handled by the Oracle Client. Oracle recommends that you use Oracle Client 11g release 2 for connections that use SCAN.

SCAN provides location independence for the databases, so that the client configuration does not have to include the address of each node that hosts a particular database or database instance. For example, if you configure policy-managed server pools for a cluster, then SCAN enables connections to databases in these server pools regardless of which nodes are allocated to the server pool.

During the installation of Oracle Grid Infrastructure, several Oracle Clusterware resources are created for SCAN:

- A SCAN VIP is created for each IP address that SCAN resolves to
- A SCAN listener is created for each SCAN VIP
- A dependency on the SCAN VIP is configured for the SCAN listener

SCAN VIPs function like node VIPs, except that SCAN VIPs can run on any node in the cluster. Also, if you have three SCAN VIPs but only two nodes in your cluster, you can have two SCAN VIPs running on the same server. Clients (users or applications) that connect using SCAN instead of the node VIPs do not have to update the list of VIP addresses in their local `tnsnames.ora` file when nodes are added to or removed from the cluster.

7.9.2 Understanding SCAN Listeners and Listener Registration

During Oracle Grid Infrastructure installation, SCAN listeners are created for as many IP addresses as there are addresses assigned to resolve to the SCAN. Oracle recommends that the SCAN resolves to three addresses, to provide high availability and scalability. If the SCAN resolves to three addresses, then there are three SCAN listeners created.

The addresses for the SCAN listeners resolve with either the domain name service (DNS), or within the cluster for the Grid Naming Service (GNS), using a round-robin method. SCAN listeners can run on any node in the cluster.

The database parameter `LOCAL_LISTENER` specifies the listening endpoint of the local database listener, and the database parameter `REMOTE_LISTENER` parameter identifies the SCAN listeners. The database registers with the local and SCAN listeners by using the connect description information contained in these parameters. Oracle Database 11g release 2 and later instances only register with SCAN listeners as remote listeners. Upgraded databases register with SCAN listeners as remote listeners, and also continue to register with all node listeners.

Starting with Oracle Database 11g release 2, the `REMOTE_LISTENER` parameter is always set to the SCAN address. Do not set it to an Oracle Net alias that has a single address that uses SCAN for the host name (`HOST=scan`). For example, if SCAN for the cluster is `myscan`, and the GNS subdomain for the cluster is `mycluster.example.com`, then the `REMOTE_LISTENER` parameter has the following value:

```
myscan.mycluster.example.com:1521
```

The Database Agent process (previously known as `racgimon`) configures the `LOCAL_LISTENER` initialization parameter automatically for Oracle RAC databases by specifying a connect descriptor for the node VIP address of the local node. You do not have to alter this setting. The `REMOTE_LISTENER` parameter should be set to the SCAN and SCAN port. If you change the SCAN port, then you must change the value for the `REMOTE_LISTENER` parameter using the SQL statement `ALTER SYSTEM SET REMOTE_LISTENER`.

7.9.3 Understanding Failover, Load Balancing, and SCAN

Oracle RAC provides failover with the node VIP addresses by configuring multiple listeners on multiple nodes to manage client connection requests for the same database service. If a node fails, then the service connecting to the VIP is relocated transparently to a surviving node, enabling fast notification of the failure to the clients connecting through the VIP. If the application and client are configured with transparent application failover options, then the client is reconnected to the surviving node.

Clients configured to use VIP addresses for Oracle RAC releases prior to Oracle Database 11g release 2 can continue to use their existing connection addresses; using SCAN is not required. When an earlier release of Oracle Database is upgraded, it registers with the SCAN listeners, and clients can start using SCAN to connect to that database.

There are two types of load balancing that you can implement for an Oracle RAC database: client-side and server-side load balancing. Client-side load balancing balances the connection requests across the listeners. With server-side load balancing, the SCAN listener directs a connection request to the best instance currently providing the service by using the load balancing advisory.

Based on the environment, the following actions occur when you use SCAN to connect to an Oracle RAC database using a service name. The numbered actions correspond to the arrows shown in [Figure 7-1](#).

1. The PMON process of each instance registers the database services with the default listener on the local node and with each SCAN listener, which is specified by the `REMOTE_LISTENER` database parameter. The listeners are dynamically updated on the amount of work being handled by the instances and dispatchers.

- The client issues a database connection request using a connect descriptor of the form:

```
orausr/@scan_name:1521/sales.example.com
```

Note: If you use the Easy Connect naming method, then ensure the `sqlnet.ora` file on the client contains `EZCONNECT` in the list of naming methods specified by the `NAMES.DIRECTORY_PATH` parameter.

- The client uses DNS to resolve `scan_name`. After DNS returns the three addresses assigned to SCAN, the client sends a connect request to the first IP address. If the connect request fails, then the client attempts to connect using the next IP address.
- When the connect request is successful, the client connects to a SCAN listener for the cluster which hosts the `sales` database. The SCAN listener compares the workload of the instances `sales1` and `sales2` and the workload of the nodes on which they are running. Because `node2` is less loaded than `node1`, the SCAN listener selects `node2` and sends the address for the listener on that node back to the client.
- The client connects to the local listener on `node2`. The local listener starts a dedicated server process for the connection to the database.
- The client connects directly to the dedicated server process on `node2` and accesses the `sales2` database instance.

Figure 7-1 Load Balancing Actions for Oracle RAC Connections That Use SCAN

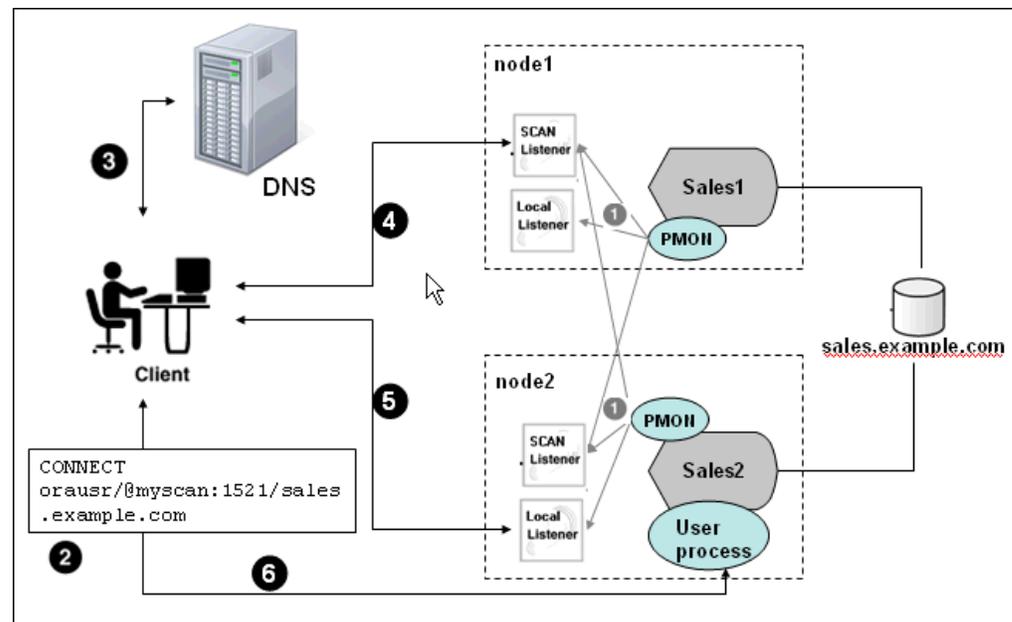


Figure 7-1 is an illustration of how load balancing works on an Oracle RAC system when SCAN is used. The contents of the illustration were described in the preceding numbered list.

See Also:

- *The Oracle Grid Infrastructure Installation Guide for Microsoft Windows x64 (64-Bit)* for more information about SCAN and its configuration
- *Oracle Real Application Clusters Administration and Deployment Guide* for more information about failover, load balancing, and the load balancing advisory

7.9.4 Understanding Shared Server Use in an Oracle RAC Database

By default, DBCA configures your Oracle RAC database with dedicated servers, not shared servers. However, if you select the shared server option when using DBCA, then DBCA configures shared servers. Oracle RAC uses both dedicated and shared server processing when shared servers are configured.

See Also:

- *Oracle Database Net Services Administrator's Guide* for more information about cross-instance registration, shared and dedicated server configurations, and connection load balancing
- *Oracle Real Application Clusters Administration and Deployment Guide* for more information about connection load balancing in an Oracle RAC environment

Removing Oracle Real Application Clusters Software

This chapter describes how to completely remove all Oracle databases, instances, and software from an Oracle Real Application Clusters (Oracle RAC) Oracle home directory.

This chapter contains the following topics:

- [Overview of Deinstallation Tools Actions](#)
- [Deinstalling Oracle RAC Software](#)

See Also: The product-specific documentation, for example *Oracle Database Client Installation Guide for Microsoft Windows*, if you want to remove a single product to ensure that you are aware of requirements and restrictions that apply for that product

8.1 Overview of Deinstallation Tools Actions

The Deinstallation tool removes Oracle RAC or standalone Oracle Database installations from your server. It can also remove the Oracle Grid Infrastructure installation, which consists of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM). The Deinstallation tool performs the follows actions when removing software from an Oracle home directory:

- Identifies all instances associated with the Oracle home
- Shuts down processes
- Removes listeners installed in the Oracle Database home
- Removes database instances
- Removes Oracle ASM release 11.1 or earlier installations
- Removes the Oracle Grid Infrastructure installation

Note:

- You cannot remove, uninstall, or deinstall the Oracle Database Vault option. However, you can disable it. See *Oracle Database Vault Administrator's Guide* for more information.
 - You can remove Oracle Configuration Manager. See *Oracle Configuration Manager Installation and Administration Guide* for more information.
 - Starting with Oracle Database 11g release 2 (11.2), you cannot remove, uninstall, or deinstall Oracle ASM separately from Oracle Clusterware or install it in a location outside of the Grid home.
-
-

Caution: If any cluster member node has more than one database with the same global database name (GDN) on a server, then you cannot use the Deinstallation tool to remove one database only. For example, if you have a standalone database on one of your cluster nodes with the GDN `mydb.example.com`, and your Oracle RAC database GDN is also `mydb.example.com`, then both databases on that node are removed.

8.2 Identifying All Instances On a Cluster

To identify all instances and services associated with the Oracle home that you want to remove, perform the following steps:

1. Log in to a cluster node as a user with Administrator privileges.
2. Locate the Oracle services related to the Oracle home.

Use the Windows Services Control Manager to manage the services. Look for any Oracle services (their names begin with `Ora`) that access the Oracle home you are removing and have the status Started.

To determine which Oracle home a service is associated with, on Windows 2003, check "Path to Executable" for a service to see the directory where the executable for the service is located.

8.3 Deinstalling Oracle RAC Software

The Deinstallation tool removes Oracle Clusterware and Oracle ASM from your server, as well as Oracle Database installations, for both Oracle RAC and single-instance databases. The topics describe the Deinstallation tool, and provide information about additional options you can use:

- [About the Deinstallation Tool](#)
- [Using the Deinstallation Tool for Oracle RAC](#)
- [Example of a Deinstallation Parameter File for Oracle RAC](#)

8.3.1 About the Deinstallation Tool

The Deinstallation tool (`deinstall.bat`) is available in the installation media before installation, and is available in Oracle home directories after installation in the directory `%ORACLE_HOME%\deinstall`. The Deinstallation tool is also available for

download from Oracle Technology Network (OTN) (<http://www.oracle.com/technetwork/>). You can download it with the complete Oracle Database 11g release 2 software, or as a separate archive file.

The `deinstall.bat` command uses the information you provide, plus information gathered from the software home to create a parameter file. You can alternatively supply a parameter file generated previously by the `deinstall.bat` command using the `-checkonly` and `-o` flags. You can also edit a response file template to create a parameter file.

The Deinstallation tool stops Oracle services, instances, and processes, and removes Oracle software and configuration files on the operating system for a specific Oracle home. At the end of the deinstallation process, you are prompted to run the `rootcrs.pl` script as a user that is a member of the Administrators group.

Caution: When you run the `deinstall` command, if the central inventory contains no other registered homes besides the home that you are deconfiguring and removing, then the Deinstallation tool removes the following files and directory contents in the Oracle base directory of the Oracle RAC installation owner:

- admin
- cfgtoollogs
- checkpoints
- diag
- oradata
- flash_recovery

Oracle strongly recommends that you configure your installations using an Optimal Flexible Architecture (OFA) configuration, and that you reserve Oracle base and Oracle home paths for exclusive use of Oracle software. If you have any user data in these locations in the Oracle base that is owned by the user account that owns the Oracle software, then the `deinstall` command deletes this data.

The `deinstall.bat` command uses the following syntax, where variable content is indicated by italics:

```
deinstall.bat [-home complete path of Oracle home] [-silent] [-checkonly] [-local]
[-paramfile complete path of input parameter property file]
[-params name1=value name2=value ...]
[-o complete path of directory for saving files]
[-tmpdir complete path of temporary directory to use]
[-logdir complete path of log directory to use] [-help]
```

Note: You cannot use the Deinstallation tool with the `-local` option for a shared Oracle home.

The default method for running the Deinstallation tool is from the `deinstall` directory in the Oracle home. For example:

```
C:\> cd %ORACLE_HOME%\deinstall
```

```
C:\..\dbhome_1\> deinstall.bat
```

In addition, you can run the Deinstallation tool from other locations, or with a parameter file, or select any of the following options when using Deinstallation tool:

- `-home`

Use this flag to indicate the home path of the Oracle home to check or deinstall. To deinstall Oracle software using the `deinstall.bat` command located within the Oracle home being removed, provide a parameter file in a location outside the Oracle home, and do not use the `-home` flag.

If you run `deinstall.bat` from the `%ORACLE_HOME%\deinstall` path, then the `-home` flag is not required because the tool knows from which home it is being run. If you use the standalone version of the tool, then the `-home` flag is mandatory

- `-silent`

Use this flag to run the command in noninteractive mode. This option requires one of the following:

- A working system that the Deinstallation tool can access to determine the installation properties
- A properties file that contains the configuration values for the Oracle home that is being deinstalled or deconfigured

To create a properties file, modify the template file `deinstall.rsp.tmpl`, located in the `response` folder of the Deinstallation tool home or the Oracle home.

If you have a working system, you can generate a properties file to use or modify by running the `deinstall.bat` command with the `-checkonly` flag. The `deinstall.bat` command then discovers information from the Oracle home to deinstall and deconfigure. It generates the properties file, which you can then use with the `-silent` option.

- `-checkonly`

Use this flag to check the status of the Oracle software home configuration. Running the command with the `checkonly` flag does not remove the Oracle configuration. The `checkonly` flag generates a properties file that you can use with the `deinstall.bat` command.

When you run `checkonly` to generate a properties file, you are prompted to provide information about your system. You can accept the default value the Deinstallation tool has obtained from your Oracle installation, indicated inside brackets ([]), or you can provide different values. To accept the defaults, click **Enter** at each prompt.

- `-local`

Use this flag in a multinode environment to deconfigure Oracle software on a local node of a cluster.

When you run the `deinstall.bat` command with this flag, it deconfigures and deinstalls the Oracle software on the local node (the node on which you run the `deinstall.bat` command) for non-shared Oracle home directories.

Note: You cannot use the Deinstallation tool with the `-local` option for a shared Oracle home.

- `-paramfile` *complete path of input parameter property file*

This is an optional flag. Use this flag to run the `deinstall.bat` command with a parameter file in a location other than the default. When you use this flag, provide the complete path where the parameter file is located. If you are running the `deinstall.bat` command from the Oracle home that you plan to deinstall, then you do not need to use the `-paramfile` option.

The default location of the parameter file depends on the location of the `deinstall.bat` command you use:

Location of <code>deinstall.bat</code>	Default Location of the Parameter File
In the installation media or stage location:	<code>X:\staging_location\deinstall\response</code>
From an unzipped archive file downloaded from OTN	<code>X:\ziplocation\deinstall\response</code> where <code>X:\ziplocation</code> refers to the directory into which the downloaded archive file was extracted
After installation, in the installed Oracle home	<code>%ORACLE_HOME%\deinstall\response</code>

- `-params [name1=value name2=value name3=value ...]`

Use this flag with a parameter file to override one or more values in a parameter file that you created.

- `-o complete path of directory for saving response files`

Use this flag to provide a path other than the default location where the properties file (`deinstall_Oracle_home_name.rsp`) is saved.

The default location of the properties file depends on the location of the Deinstallation tool:

Location of <code>deinstall.bat</code>	Default Location of the Properties File
In the installation media or stage location:	<code>X:\staging_location\deinstall\response</code>
From an unzipped archive file downloaded from OTN	<code>X:\ziplocation\deinstall\response</code> where <code>X:\ziplocation</code> refers to the directory into which the downloaded archive file was extracted
After installation, in the installed Oracle home	<code>%ORACLE_HOME%\deinstall\response</code>

- `-tmpdir complete path of temporary directory`

Specifies a non-default location where the deinstallation tool writes the temporary files for the deinstallation.

- `-logdir complete path of log directory`

Specifies a non-default location where the deinstallation tool writes the log files for the deinstallation.

- `-help`

Use this flag (`-help` or `-h`) to obtain additional information about the command options and flags.

If you use the `deinstall.bat` command located in an Oracle home, or the `deinstall.bat` command downloaded from OTN (not installed in an Oracle home),

then the command writes log files in the C:\Program Files\Oracle\Inventory\logs directory. If, however, you are using the `deinstall.bat` command to remove the last Oracle home installed on the server, then the log files are written to:

- `%TEMP%\OraDeinstall<timestamp>\logs` if you use the `deinstall.bat` command located in the Oracle home
- `X:\ziplocation\deinstall\logs` if you use the `deinstall.bat` command downloaded from OTN

8.3.2 Using the Deinstallation Tool for Oracle RAC

If you use the separately downloaded version of the Deinstallation tool, then when the `deinstall.bat` command runs, you are prompted to provide the home directory of the Oracle software to remove from your system. Provide additional information as prompted.

If you run the `deinstall.bat` command outside of the `%ORACLE_HOME%\deinstall` folder, then help output is displayed, unless you use the `-home` flag and provide a path to the Oracle home.

If you run the `deinstall.bat` command from the `%ORACLE_HOME%\deinstall` folder, then deinstallation starts without prompting you for a home address.

Example 8–1 Running the Deinstallation Tool Installed in an Oracle Home

To run the `deinstall.bat` command located in an Oracle Database home in the path C:\app\oracle\product\11.2.0\dbhome_1, enter the following command while logged in as a member of the Administrator group:

```
C:\> cd app\oracle\product
C:\app\oracle\product> 11.2.0\dbhome_1\deinstall\deinstall.bat
```

To completely remove the Oracle Database installation, you must run the `deinstall.bat` command from a prompt that is located outside of the Oracle home. If you run the `deinstall.bat` command after changing into the `Oracle_home\deinstall` directory, then the Deinstallation tool will not be able to remove all the Oracle software files and directories.

Example 8–2 Running the Standalone Version of the Deinstallation Tool

Assume you have installed the Deinstallation tool in the path C:\app\oracle\product\deinstall. Use the `deinstall.bat` command in this directory to deinstall an Oracle Database home using a parameter file. In this example, the parameter file is located in the user directory C:\user\oracle:

```
C:\> cd app\oracle\product\deinstall
C:\..deinstall> deinstall.bat -home C:\app\oracle\product\11.2.0\dbhome_1
-paramfile C:\user\oracle\my_db_paramfile.tmpl
```

8.3.3 Example of a Deinstallation Parameter File for Oracle RAC

You can run the `deinstall` command with the `-paramfile` option to use the values you specified in a parameter file instead of responding to individual prompts.

Example 8–3 Sample Parameter File to use with the Deinstallation Tool to Uninstall Oracle RAC

The following is an example of a parameter file created by the Deinstallation tool. The parameter file contains the following information:

- The Oracle Database software owner is oracle
- The Oracle Database home (Oracle home) is in the path
C:\app\oracle\product\11.2.0\dbhome_1
- The Oracle base is C:\app\oracle
- The Grid home is C:\app\11.2.0\grid
- The central Oracle Inventory home is C:\Program Files\Oracle\Inventory
- The virtual IP address (VIP) is 192.0.2.1
- The name of the local node (the node on which you run the Deinstallation tool) is node1
- The name of the remote node is node2
- The database name is orcl and it is an Oracle RAC database
- The database uses Oracle ASM for storage

```
#Copyright (c) 2005, 2010 Oracle Corporation. All rights reserved.
LOCAL_NODE=node1
HOME_TYPE=RACDB
ORACLE_BASE=C:\app\oracle
ObaseCleanupPtrLoc=C:\DOCUME~1\oracle\LOCALS~1\Temp\OraDeinstall2010-03-31_
  15-24-53\utl\orabase_cleanup.lst
DB_TYPE.orcl=RAC_DB
silent=false
ORA_CRS_HOME=C:\app\11.2.0\grid\
EMCP.CONFIG_MODE.orcl=CONFIG_DBCONTROL
EMCA_LOG.orcl=C:\app\oracle\cfgtoollogs\emca\orcl
LOGDIR=C:\Program Files\Oracle\Inventory\logs\
DATAFILE_LOC.orcl=
ORACLE_BINARY_OK=true
DUMP_DESTINATION_LOC.orcl=C:\app\oracle\admin\orcl
EMCP.DB_UNIQUE_NAME_LIST=orcl
CLUSTER_NODES=node1,node2
REMOTE_NODES=node2
SPFILE_LOC.orcl=+DATA/orcl/spfileorcl.ora
ARCHIVE_LOG_DESTINATION_LOC.orcl=
ORACLE_BASE.orcl=C:\app\oracle
LOCAL_SID.orcl=orcl1
local=false
INVENTORY_LOCATION=C:\Program Files\Oracle\Inventory
NODE_LIST.orcl=node1,node2
RAW_MAPPING_FILE.orcl=
DIAG_DEST.orcl=C:\APP\ORACLE\diag\rdbms\orcl
ORACLE_HOME=C:\app\oracle\product\11.2.0\dbhome_1\
STORAGE_TYPE.orcl=ASM
CCR_CONFIG_STATUS=CCR_DEL_HOME
CRS_HOME=false
EMCP.CHECK_POINT.orcl=WINDOWS_SERVICE_REGISTRY_CREATE,DBCONTROL_START
CONFIGFILE_LOC.orcl=+DATA/orcl/controlfile/current.260.712754807,C:\APP\ORACLE
  \PRODUCT\11.2.0\DBHOME_1\DATABASE\DR1ORCL.DAT,C:\APP\ORACLE\PRODUCT\11.2.0\
  DBHOME_1\DATABASE\DR2ORCL.DAT
SID_LIST.orcl=orcl1,orcl2
OLD_ACTIVE_ORACLE_HOME=
CREATION_MODE.orcl=y
COMPS_TO_REMOVE=ode.net.ntoledb,oramts
DB_UNIQUE_NAME_LIST=orcl
FLASH_RECOVERY_LOC.orcl=
```

Troubleshooting Oracle Real Application Clusters Installations

This appendix provides troubleshooting information for installing Oracle Real Application Clusters (Oracle RAC).

See Also: The Oracle Database 11g release 2 (11.2) Oracle Real Application Clusters documentation set included with the installation media in the Documentation directory:

- *Oracle Real Application Clusters Administration and Deployment Guide*
- *Oracle Clusterware Administration and Deployment Guide*

This appendix contains the following topics:

- [Troubleshooting Oracle RAC Installation Errors](#)
- [Resolving Oracle RAC Installations Errors](#)
- [Troubleshooting Configuration Assistant Errors](#)
- [Troubleshooting Silent or Noninteractive Installations](#)
- [About Using CVU Cluster Healthchecks After Installation](#)
- [Cleaning Up After a Failed Installation](#)

A.1 Troubleshooting Oracle RAC Installation Errors

The following information explains how to proceed if you encounter an error during installation of Oracle RAC.

A.1.1 Verifying Requirements

Before you try any of the troubleshooting steps in this appendix, do the following:

- Check [Chapter 1, "Preinstallation Checklist"](#) to ensure that your system meets the requirements and that you have completed all of the preinstallation tasks.
- Read the release notes for the product on your platform before installing it. The release notes are available on the Oracle Database installation media. You can find the latest version of the release notes on the Oracle Technology Network web site:

<http://www.oracle.com/technetwork/indexes/documentation/>

A.1.2 Understanding How To Resolve Installation Errors

If you encounter an error during installation:

- Do not exit Oracle Universal Installer (OUI).
- If you clicked **Next** after you entered incorrect information in an installation window, then click **Back** to return to the window and correct the information.
- If you encounter an error while OUI is copying files, then see [Section A.1.3, "Locating the Log File For an Installation Session"](#) for interactive installations or [Section A.4, "Troubleshooting Silent or Noninteractive Installations"](#) for more information.
- If you encounter an error while a configuration assistant is running, then see the [Section A.3, "Troubleshooting Configuration Assistant Errors"](#) section.
- If you cannot resolve the problem, then remove the failed installation by following the steps listed in [Section A.6, "Cleaning Up After a Failed Installation."](#)

A.1.3 Locating the Log File For an Installation Session

During an installation, OUI records all the actions that it performs in a log file. If you encounter problems during the installation, then review the log file for information about possible causes of the problem.

By default, the log files are located in the following directory:

```
%SYSTEMDRIVE%\Program Files\Oracle\Inventory\logs
```

Log file names from interactive installations have the form:

```
installActionsdate_time.log
```

For example, if an interactive installation occurred at 9:00:56 a.m. on October 14, 2009, then the log file would be named:

```
installActions2009-10-14_09-00-56AM.log
```

Note: Do not delete or manually alter the `Inventory` directory or its contents. Doing so can prevent OUI from locating products that you install on your system.

See Also: [Section A.4, "Troubleshooting Silent or Noninteractive Installations"](#)

A.2 Resolving Oracle RAC Installations Errors

This section contains the following topics:

- [General Installation Issues](#)
- [Performing Cluster Diagnostics During Oracle RAC Installations](#)
- [Reviewing the Log of an Installation Session](#)

A.2.1 General Installation Issues

The following list contains examples of the types of errors that can occur during installation:

- Nodes unavailable for selection from the OUI Node Selection screen
- Node nodename is unreachable
- PROT-8: Failed to import data from specified file to the cluster registry
- Time stamp is in the future
- Administrative user unable to log in to SQL*Plus using the SYSDBA privilege

Nodes unavailable for selection from the OUI Node Selection screen

Cause: Oracle Clusterware was either not installed, or the Oracle Clusterware services were not active or started.

Action: Install Oracle Clusterware, or review the status of your Oracle Clusterware. Consider restarting the nodes, because doing so may resolve the problem.

Node *nodename* is unreachable

Cause: The IP host was unavailable.

Action: Attempt the following:

1. Run the command `ipconfig /all`. Compare the output of this command with the contents of the `C:\WINDOWS\system32\drivers\etc\hosts` file to ensure that the node IP is listed.
2. Run the command `nslookup` to verify that the host is reachable.

PROT-8: Failed to import data from specified file to the cluster registry

Cause: Insufficient space in an existing Oracle Cluster Registry raw device partition, which caused a migration failure at the end of the installation. To confirm this issue, look for the error "Not enough space in the backing store" in the log file `Grid_home\log\hostname\client\ocrconfig_pid.log`.

Action: Identify a raw disk partition that has 280 MB or more available space. Copy the contents of the existing raw partition from to the new device using the command `XCOPY`.

Time stamp is in the future

Cause: One or more nodes had a different clock time than the local node. If this was the case, then you may see output similar to the following:

```
time stamp 2019-04-04 14:49:49 is 106 s in the future
```

Action: Ensure that all member nodes of the cluster have the same clock time.

Administrative user unable to log in to SQL*Plus using the SYSDBA privilege

Cause: When Oracle Database on Microsoft Windows was installed, Oracle Universal Installer created a Windows local group called `ORA_DBA`, and then added the user name of the user performing the installation to this group. Members of `ORA_DBA` automatically receive the SYSDBA privilege. However, for cluster installations, OUI does not add the user to `ORA_DBA` if they have performed the installation remotely. As a result, this user could not log in to SQL*Plus using the SYSDBA privilege.

Action: Manually add remote users to `ORA_DBA`.

See Also: *Oracle Database Administrator's Guide* for information about connecting to the database using SQL*Plus

A.2.2 Performing Cluster Diagnostics During Oracle RAC Installations

When installing Oracle RAC, if OUI does not display the Node Selection page, then perform clusterware diagnostics by running the `olsnodes -v` command from the binary directory in your Oracle Clusterware home (*Grid_home*\bin) and analyzing its output. See *Oracle Clusterware Administration and Deployment Guide* if the detailed output indicates that your clusterware is not running.

In addition, use the following command to check the integrity of the Cluster Manager, where the variable *node_list* is the list of nodes in your cluster, separated by commas:

```
cluvfy comp clumgr -n node_list -verbose
```

A.2.3 Reviewing the Log of an Installation Session

During an installation, OUI records all of the actions that it performs in a log file. If you encounter problems during the installation, then review the log file for information about possible causes of the problem.

To view the log file, follow these steps:

1. Go to the `logs` subdirectory of the Oracle Inventory directory:

```
C:\> cd Program Files\Oracle\Inventory\logs
```

2. View the file details for the folder to determine the most recent file. Installer log files have names similar to the following, where *date_time* indicates the date and time that the installation started:

```
installActionsdate_time.log
```

3. Use a text editor to view the most recent entries in the log file, where information about a problem is most likely to appear.

A.3 Troubleshooting Configuration Assistant Errors

To troubleshoot an installation error that occurs when a configuration assistant is running, do the following:

- Review the installation log files listed in the section [Section A.3.1, "Reviewing Configuration Assistant Errors."](#)
- Review the specific configuration assistant log file located in the Oracle base directory, in the path `%ORACLE_BASE%\cfgtoollogs`. Try to fix the issue that caused the error.
- If you see the "Fatal Error. Reinstall" message, then look for the cause of the problem by reviewing the log files. See the section [Section A.3.3, "Resolving Irrecoverable Errors Reported by Configuration Assistants"](#) for further instructions.

A.3.1 Reviewing Configuration Assistant Errors

To troubleshoot an installation error that occurs when a configuration assistant, such as Oracle Net Configuration Assistant (NETCA), is running:

1. Review the installation log files as described in the section [Section A.2.3, "Reviewing the Log of an Installation Session."](#)

2. Review the specific configuration assistant log file as described in [Section A.3.2, "Identifying Configuration Assistant Failures."](#) Try to fix the issue that caused the error.
3. If you see the "Fatal Error. Reinstall" message, then look for the cause of the problem by reviewing the log files. See [Section A.3.3, "Resolving Irrecoverable Errors Reported by Configuration Assistants"](#) for more information.

A.3.2 Identifying Configuration Assistant Failures

If an Oracle configuration assistant is not able to complete its configuration tasks, then it reports a status of 'Failed'. Oracle configuration assistant failures are noted at the bottom of the installation screen. The configuration assistant interface displays additional information, if available.

The configuration assistant execution status is stored in the following file:

```
C:\Program Files\Oracle\Inventory\logs\installActionsdate_time.log
```

Completion status codes are listed in the following table:

Status	Result Code
Configuration assistant succeeded	0
Configuration assistant failed	1
Configuration assistant canceled	-1

More information about errors related to the configuration assistant can be found in the following directory, where the ORACLE_BASE directory is the Oracle base for the user performing the Oracle RAC installation:

```
ORACLE_BASE\cfgtoollogs
```

A.3.3 Resolving Irrecoverable Errors Reported by Configuration Assistants

If you receive an irrecoverable (or *fatal*) error while a configuration assistant is running, then you must complete the following tasks:

1. Remove the failed installation as described in [Section A.6, "Cleaning Up After a Failed Installation."](#)
2. Correct the cause of the irrecoverable error.
3. Reinstall the Oracle software.

See Also: [Chapter 8, "Removing Oracle Real Application Clusters Software"](#)

A.4 Troubleshooting Silent or Noninteractive Installations

To determine whether a silent or noninteractive installation succeeds or fails, check the silentInstallActionsdate_time.log file, located in `DRIVE_LETTER:\Program Files\Oracle\Inventory\logs`, as described in [Section A.2.3, "Reviewing the Log of an Installation Session."](#)

A silent or noninteractive installation fails if:

- You do not specify a response file.

- You specify an incorrect or incomplete response file.

For example, a common problem is that while all the product-specific data is correct, the staging area location may be incorrect. If this is the case, then check the `FROM_LOCATION` variable and ensure that it points to the `products.xml` file in the installation media. In the installation media, this `products.xml` is in `database\stage`.

- OUI encounters an error, such as insufficient disk space.

OUI or a configuration assistant validates the response file at run time. If the validation fails, then the silent or noninteractive installation or configuration process ends. OUI treats values for parameters that are of the wrong context, format, or type as if no value had been specified in the file.

See Also: [Section A.1.3, "Locating the Log File For an Installation Session"](#)

A.5 About Using CVU Cluster Healthchecks After Installation

Starting with Oracle Grid Infrastructure 11g release 2 (11.2.0.3) and later, you can use the CVU healthcheck command option to check your Oracle Clusterware and Oracle Database installations for their compliance with mandatory requirements and best practices guidelines, and to ensure that they are functioning properly.

Use the following syntax to run the healthcheck command option:

```
cluvfy comp healthcheck [-collect {cluster|database}] [-db db_unique_name]
[-bestpractice|-mandatory] [-deviations] [-html] [-save [-savedir directory_path]]
```

For example:

```
$ cd /home/grid/cvu_home/bin
$ ./cluvfy comp healthcheck -collect cluster -bestpractice -deviations -html
```

The options are:

- `-collect [cluster|database]`

Use this flag to specify that you want to perform checks for Oracle Clusterware (cluster) or Oracle Database (database). If you do not use the collect flag with the healthcheck option, then `cluvfy comp healthcheck` performs checks for both Oracle Clusterware and Oracle Database.

- `-db db_unique_name`

Use this flag to specify checks on the database unique name that you enter after the `db` flag.

CVU uses JDBC to connect to the database as the user `cvusys` to verify various database parameters. For this reason, if you want checks to be performed for the database you specify with the `-db` flag, then you must first create the `cvusys` user on that database, and grant that user the CVU-specific role, `CVUSAPP`. You must also grant members of the `CVUSAPP` role `SELECT` permissions on system tables.

The SQL script `cvusys.sql` is included in the `CVU_home/cv/admin` directory to facilitate the creation of this user. Use this SQL script to create the `cvusys` user on all the databases that you want to verify using CVU.

If you use the `db` flag but do not provide a database unique name, then CVU discovers all the Oracle Databases on the cluster. If you want to perform best practices checks on these databases, then you must create the `cvusys` user on each

database, and grant that user the CVUSAPP role with the SELECT privileges needed to perform the best practice checks.

- [-bestpractice | -mandatory] [-deviations]

Use the `bestpractice` flag to specify best practice checks, and the `mandatory` flag to specify mandatory checks. Add the `deviations` flag to specify that you want to see only the deviations from either the best practice recommendations or the mandatory requirements. You can specify either the `-bestpractice` or `-mandatory` flag, but not both flags. If you specify neither `-bestpractice` or `-mandatory`, then both best practices and mandatory requirements are displayed.

- -html

Use the `html` flag to generate a detailed report in HTML format.

If you specify the `html` flag, and a browser CVU recognizes is available on the system, then the browser is started and the report is displayed on the browser when the checks are complete.

If you do not specify the `html` flag, then the detailed report is generated in a text file.

- -save [-savedir *dir_path*]

Use the `save` or `-save -savedir` flags to save validation reports (`cvuchecdkreport_timestamp.txt` and `cvucheckreport_timestamp.htm`), where *timestamp* is the time and date of the validation report.

If you use the `save` flag by itself, then the reports are saved in the path `CVU_home/cv/report`, where *CVU_home* is the location of the CVU binaries.

If you use the flags `-save -savedir`, and enter a path where you want the CVU reports saved, then the CVU reports are saved in the path you specify.

A.6 Cleaning Up After a Failed Installation

If an installation fails, then you must remove the Oracle home directory and remove all files that OUI created during the attempted installation. Perform the following steps to clean up the failed installation:

1. Follow the instructions in [Chapter 8, "Removing Oracle Real Application Clusters Software"](#) to run OUI to deinstall Oracle RAC.
2. Manually remove the directory that was used as the Oracle home directory during the installation.
3. Remove the following Windows Registry keys created by OUI during the previous installation attempt:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ODBC\ODBCINST.INI
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Oracle11\Performance
  \KEY_ORACLE_HOME_NAME
```

After you have completed these steps, you can start the installation again.

Using Scripts or Response Files to Create Oracle Real Application Clusters Databases

This appendix describes the steps for noninteractive installations during which you can create Oracle Real Application Clusters (Oracle RAC) databases using scripts.

Note: The scripts generated by Oracle Database Configuration Assistant (DBCA) are for reference purposes only. Oracle strongly recommends that you use DBCA to create a database.

See Also: *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for detailed information about generating installation scripts

This appendix contains the following topics:

- [Using DBCA to Generate Installation Scripts for Oracle RAC](#)
- [About DBCA Noninteractive \(Silent\) Configuration for Oracle RAC](#)
- [Using DBCA Noninteractive \(Silent\) Configuration for Oracle RAC](#)
- [Understanding How Response Files are Used During Installation](#)
- [Preparing Response Files](#)
- [Running Oracle Universal Installer Using a Response File](#)
- [Running Configuration Assistants Using Response Files](#)

B.1 Using DBCA to Generate Installation Scripts for Oracle RAC

To generate scripts to create an Oracle RAC database, create a database using the generated scripts, and prepare the database for use, complete the following steps:

1. Start DBCA and select your preferred options to build the Oracle RAC database.

On the Creation Options page of your DBCA session, deselect Create Database and select **Generate Database Creation Scripts** before you click **Finish**. You can accept the default destination directory for the scripts, or browse for a different location. In either case, you should note the path name for use in the next step.

See Also: [Chapter 3, "Creating Oracle Real Application Clusters Databases with Oracle Database Configuration Assistant"](#) for additional information about running a DBCA session

2. Go to the directory where DBCA created the scripts, and review the SQL scripts to ensure that they contain the statements to build a database with the characteristics you require. If they do not, then Oracle recommends that you rerun DBCA to create scripts with the desired configuration rather than editing the scripts yourself.
3. On each cluster node you identified during your DBCA session, run the script *sid.bat*, where *sid* is the SID prefix that you entered on the DBCA Database Name page.
4. Set the initialization parameter, *cluster_database*, to the value `TRUE` in your *SPFILE* by entering an `ALTER SYSTEM` statement in SQL*Plus, or by uncommenting the parameter in the *PFILE* for each instance.
5. Configure Oracle Net Services to support your new database and instances as described in [Chapter 7, "Understanding the Oracle Real Application Clusters Installed Configuration"](#).
6. Set the `REMOTE_LISTENER` parameter to the `SCAN` (using the Easy Connect Naming syntax *scanname:scanport*) in your *SPFILE* by entering an `ALTER SYSTEM` statement in SQL*Plus, or by uncommenting the parameter in the *PFILE* for each instance.
7. Run the Server Control Utility (SRVCTL) to configure and start database and instance applications as described in *Oracle Real Application Clusters Administration and Deployment Guide*.

See Also: *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for additional information about creating and using scripts to install Oracle software with OUI

B.2 About DBCA Noninteractive (Silent) Configuration for Oracle RAC

You can perform a noninteractive, or *silent* configuration of Oracle RAC using DBCA. To perform a silent configuration, you must have completed an Oracle Grid Infrastructure (Oracle Clusterware and Oracle Automatic Storage Management) installation and defined the Oracle home directory environment variable

You can use DBCA to create a database from templates supplied by Oracle, or from templates that you create. The templates contain settings optimized for a particular type of workload.

Oracle provides templates for the following two workload types:

- General purpose or transaction processing
- Data warehouse

For more complex environments, you can select the Custom Database option. This option does not use templates and results in a more extensive installation interview, which means that it takes longer to create your database.

B.3 Using DBCA Noninteractive (Silent) Configuration for Oracle RAC

The following command creates an Oracle RAC database using the general purpose template, placing the data files in an existing Oracle ASM disk group. The clusters on which an Oracle RAC database instance are created are nodes *node1* and *node2*, the disk group name is *+ASMgrp1*, and *password* is a placeholder for a password. The passwords can be all the same password or different passwords each time:

```
%ORACLE_HOME%\bin\dbca -silent -createDatabase -templateName General_Purpose.dbc  
-gdbName %DBNAME% -sid %ORACLE_SID% -sysPassword password -systemPassword password
```

```
-sysmanPassword password -dbsnmpPassword password -emConfiguration LOCAL
-storageType ASM -diskGroupName +ASMgrp1 -datafileJarLocation %ORACLE_HOME%\
assistants\dbca\templates -nodeinfo node1,node2 -characterset WE8MSWIN1252
-obfuscatedPasswords false -sampleSchema false -asmSysPassword password
```

B.4 Understanding How Response Files are Used During Installation

When you start OUI, you can use a response file to automate the installation and configuration of Oracle software, either fully or partially. OUI uses the values contained in the response file to provide answers to some or all installation prompts.

Typically, the installer runs in interactive mode, which means that it prompts you to provide information in graphical user interface (GUI) screens. When you use response files to provide this information, you run the installer from a command prompt using either of the following modes:

- Silent

If you include responses for all of the prompts in the response file and specify the `-silent` option when starting the installer, then it runs in silent mode. During a silent mode installation, the installer does not display any screens. Instead, it displays progress information in the terminal that you used to start it.

- Response file

If you include responses for some or all of the prompts in the response file and omit the `-silent` option, then the installer runs in response file mode. During a response file mode installation, the installer displays all the screens, screens for which you specified information in the response file, and also screens for which you did not specify the required information in the response file.

You define the settings for a silent or response file installation by entering values for the variables listed in the response file. For example, to specify the Oracle home name, supply the appropriate value for the `ORACLE_HOME` variable:

```
ORACLE_HOME="OraDBHome1"
```

Another way of specifying the response file variable settings is to pass them as command-line arguments when you run the installer. For example:

```
-silent "ORACLE_HOME=OraDBHome1" ...
```

This method is particularly useful if you do not want to embed sensitive information, such as passwords, in the response file. For example:

```
-silent "s_dlgRBOPassword=binks342" ...
```

Enclose the variable name and its value in double quotation marks.

See Also: *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for more information about response files

B.4.1 Deciding to Use Silent Mode or Response File Mode

The following table provides reasons for running the installer in silent mode or response file mode:

Mode	Reasons to Use
Silent	<p>Use silent mode for the following installations:</p> <ul style="list-style-type: none"> ■ To complete an unattended installation, which you schedule using operating system utilities ■ To complete several similar installations on multiple systems without user interaction <p>OUI displays progress information on the terminal that you used to start it, but it does not display any of the installer screens.</p>
Response file	<p>Use response file mode to complete similar Oracle software installations on more than one system, providing default answers to some, but not for all of the OUI prompts.</p> <p>If you do not specify information required for a particular OUI screen in the response file, then OUI displays that screen. OUI suppresses screens for which you have provided all of the required information.</p>

B.4.2 General Procedure for Using Response Files

The following are the general steps to install and configure Oracle products using the installer in silent or response file mode:

Note: You must complete all required preinstallation tasks on a system before running the installer in silent or response file mode.

1. Prepare a response file.
2. Run OUI in silent or response file mode.
3. If you completed a software-only installation, then run Oracle Net Configuration Assistant (NETCA) and DBCA in silent or response file mode.

These steps are described in the following sections.

B.5 Preparing Response Files

The following are methods you can use to prepare response files for silent mode or response file mode installations:

- [Editing a Response File Template](#)
- [Recording Response Files](#)

B.5.1 Editing a Response File Template

Oracle provides response file templates for each product and installation type and for each configuration tool. For Oracle Database, these files are located in the `database\response` directory on the installation media or in the `Oracle_home\inventory\response` directory after the software has been installed.

Note: If you copied the installation media to a directory on a local disk (referred to as the `staging_dir` directory), then the response files are located in the directory `staging_dir\database\response`.

Table B-1 lists the response files provided with the Oracle RAC software.

Table B-1 Response Files for Oracle Database

Response File	Used For
db_install.rsp	Silent installation of Oracle Database 11g
dbca.rsp	Silent creation and configuration of an Oracle database using DBCA
netca.rsp	Silent configuration of Oracle Net using NETCA

Caution: When you modify a response file template and save a file for use, the response file may contain plain text passwords. Oracle strongly recommends that database administrators or other administrators delete or secure the response files when they are not in use.

To copy and modify a response file:

1. Copy the response file from the response file directory to a directory on your system.

If you have copied the installation media to a hard drive, then you can edit the file in the `staging_dir\database\response` directory.

2. Open the response file in a text editor.

During installation, you can specify sensitive information, such as passwords, at the command line rather than within the response file. The section [Section B.4, "Understanding How Response Files are Used During Installation"](#) explains this method.

See Also: *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for detailed information about creating response files

3. Follow the instructions in the file to edit it.

Note: OUI or the configuration assistants fail if you do not correctly configure the response file.

4. Secure the response file.

Ensure that only the user that installed the Oracle software can view or modify response files. Consider deleting the modified response file after the installation succeeds.

Note: A fully specified response file for an Oracle Database installation contains the passwords for database administrative accounts and for a user who is a member of the operating system groups `ORA_DBA` and `Administrators` (required for automated backups).

B.5.2 Recording Response Files

You can use OUI in interactive mode to record response files, which you can then edit and use to complete silent mode or response file mode installations. This method is useful for Advanced or software-only installations.

Starting with Oracle Database 11g release 2 (11.2), you can save all the installation steps into a response file during installation by clicking **Save Response File** on the Summary page. You can use the generated response file for a silent installation later.

When you record the response file, you can either complete the installation, or you can exit from OUI before OUI starts to copy the software to the local disk.

If you use record mode during a response file mode installation, then OUI records the variable values that were specified in the original source response file into the new response file.

Note: You cannot save passwords while recording the response file.

To record a response file:

1. Complete preinstallation tasks as for a standard installation.
When you use OUI to record a response file, it checks the system to verify that it meets the requirements to install the software. For this reason, Oracle recommends that you complete all of the required preinstallation tasks and record the response file while completing an installation.
2. Log in as a user that is a member of the local Administrators group and start OUI.
3. On each installation screen, specify the required information.
4. When OUI displays the Summary screen, perform the following:
 - a. Click **Save Response File**. In the pop up window, specify a file name and location for the new response file. Click **Save** to write the settings you have entered to the file.
 - b. Click **Finish** to continue with the installation.
Click **Cancel** if you do not want to continue with the installation. The installation stops, but the recorded response file is retained.
5. If you do not complete the installation, then delete the Oracle home directory that the installer created using the path you specified in the Specify File Locations screen.
6. Before you use the saved response file on another system, edit the file and make any required changes. Use the instructions in the file as a guide when editing it.

B.6 Running Oracle Universal Installer Using a Response File

Now, you are ready to run OUI at the command line, specifying the response file you created, to install the Oracle Database software. The OUI executable file, `setup.exe`, provides several options.

B.6.1 Viewing the Help Information for OUI

For information about all the available options for OUI, run the `setup.exe` command with the `-help` option. For example, if the installation software is accessed through a DVD disk drive mounted as `E:`, then you would use the following commands:

```
C:\> E:
E:\> database\setup.exe -help
```

If the installation software has been extracted from a downloaded ZIP file to the directory *X:\staging_dir*, where *X*: is a drive on your system, then you would use the following command:

```
C:\> X:
X:\> staging_dir\database\setup.exe -help
```

The help information appears in your session window after a brief delay.

B.6.2 Running OUI Using a Response File for Input

To use OUI with a response file to install the Oracle Database software, complete the following steps:

1. Complete the preinstallation tasks as for a typical installation.
2. Log in as a user with Administrator privileges.
3. Start OUI in silent or response file mode, by entering a command similar to the following:

```
X:\staging_dir> database\setup.exe [-silent] [-noconfig] \
-responseFile responsefilename
```

Note: Do not specify a relative path to the response file. If you specify a relative path, then OUI fails. Use the full path location for the response file.

In this example:

- *X:staging_dir* is either the location of the installation media (for example, a DVD drive) or the path of staging directory on the local node.
- `-silent` runs OUI in silent mode.
- `-noconfig` suppresses running the configuration assistants during installation, and a software-only installation is performed instead.
- *responsefilename* is the full path and file name of the installation response file that you configured previously. See [Section B.5.1, "Editing a Response File Template."](#)

B.7 Running Configuration Assistants Using Response Files

You can run configuration assistants in response file or silent mode to configure and start Oracle software after it is installed on your system. To run configuration assistants in response file or silent mode, you must first copy and edit a response file template.

Note: If you copied the software to a hard disk, then the response file template is located in the directory *X:\staging_dir\database\response*

This section contains the following topics:

- [Running Oracle Database Configuration Assistant in Response File or Silent Mode](#)
- [Running Oracle Net Configuration Assistant Using Response Files](#)
- [About the Postinstallation Configuration File](#)
- [Creating a Password Response File](#)
- [Performing Postinstallation Configuration Using a Response File](#)

B.7.1 Running Oracle Database Configuration Assistant in Response File or Silent Mode

To run Oracle Database Configuration Assistant (DBCA) in response file or silent mode:

1. Oracle provides a response file template named `dbca.rsp` in the `response` directory on the installation media. Copy the `dbca.rsp` response file template from the `response` file directory to a directory on your system.

If you have copied the software to a hard drive, you can edit the file in the `X:\staging_dir\database\response` directory if you prefer.

Note: As an alternative to editing the response file template, you can also create a database by specifying all required information as command-line options when you run DBCA. For information about the list of options supported, enter the following command:

```
C:\> %ORACLE_HOME%\bin\dbca -help
```

2. Log in as a user that is a member of the local Administrators group. Set the `%ORACLE_HOME%` environment variable to the correct Oracle home directory. See [Section 7.3, "Working with Multiple Oracle Home Directories on Windows."](#)
3. Open the response file in a text editor.
4. Edit the file, following the instructions in the file.

In response file mode, DBCA uses values that you specify in the response file or as command-line options, to create a database.

Note: DBCA fails if you do not correctly configure the response file.

5. Open a command-line window. Change directories to the location of the Oracle home directory.
6. Use a command similar to the following example to run DBCA in silent or response file mode using a response file:

```
C:\> %ORACLE_HOME%\bin\dbca {-progressOnly | -silent} -responseFile \  
local_dir\dbca.rsp
```

In this example:

- The `-silent` option runs DBCA in silent mode, suppressing user prompts.
- The `-progressOnly` option runs DBCA in response file mode.
- `local_dir` is the full path of the directory where the `dbca.rsp` response file is located.

As DBCA configures and starts the database, it displays a window that contains status messages and a progress bar. The window that DBCA displays is the same window that is displayed when you choose to create a preconfigured database during an Oracle Database Enterprise Edition or Standard Edition installation.

B.7.2 Running Oracle Net Configuration Assistant Using Response Files

You can run Oracle Net Configuration Assistant (NETCA) in silent mode to configure and start an Oracle Net listener on the system, configure naming methods, and configure Oracle Net service names. To run NETCA in silent mode, you must copy and edit a response file template. Oracle provides a response file template named `netca.rsp` in the `response` subdirectory of:

- The `Oracle_home\inventory\response` directory after a software-only installation
- The `database\response` directory on the installation media or staging area

To run NETCA using a response file:

1. Copy the `netca.rsp` response file template from the response file directory on the installation media to a directory on your system.

If the software is staged on a hard drive, or has already been installed, then you can edit the file in the `response` directory located on the local disk instead.

2. Open the response file in a text editor.
3. Follow the instructions in the file to edit it.

Note: NETCA fails if you do not correctly configure the response file.

4. Log in as a user that is a member of the local Administrators group. Set the `%ORACLE_HOME%` environment variable to the correct Oracle home directory. See [Section 7.3, "Working with Multiple Oracle Home Directories on Windows."](#)
5. Enter a command similar to the following to run NETCA in silent mode:

```
C:\> Oracle_home\bin\netca -silent -responsefile X:\local_dir\netca.rsp
```

In this command:

- The `-silent` option runs NETCA in silent mode.
- `X:\local_dir` is the full path of the directory where you copied the `netca.rsp` response file template, where `X` is the drive on which the file is located and `local_dir` is the folder on that drive.

B.7.3 About the Postinstallation Configuration File

When you perform an installation using silent more or response file mode, you provide information about your servers in a response file that you otherwise provide manually using a graphical user interface. However, the response file does not contain passwords for user accounts that configuration assistants require after software installation is complete. To run the configuration assistants after the installation completes in silent mode, you must provide the passwords used by the assistants in a password file.

The configuration assistants are started with a script called `configToolAllCommands`. You can run this script in silent mode by using a password response file. The script uses the passwords in the file to run the configuration tools in succession to complete

the software configuration. If you keep the password file to use when cloning installations, then Oracle strongly recommends that you store the password file in a secure location.

You can also use the password file to restart a failed installation. If you have to stop an installation to fix an error, then you can rerun the configuration assistants using `configToolAllCommands` and a password response file.

The `configToolAllCommands` password response file has the following options:

- `internal_component_name` is the name of the component that the configuration assistant configures.
- `variable_name` is the name of the configuration file variable.
- `value` is the desired value to use for configuration.

The command syntax is as follows:

```
internal_component_name|variable_name=value
```

For example, to set the password for the SYS user of Oracle ASM:

```
oracle.assistants.asm|S_ASMPASSWORD=myPassWord
```

Oracle Database configuration assistants require the SYS, SYSTEM, SYSMAN, and DBSNMP passwords for use with DBCA. You might need to specify the following additional passwords, depending on your system configuration:

- If the database is using Oracle ASM for storage, then you need to specify a password for the `S_ASMSNMPPASSWORD` variable. If you are not using Oracle ASM, then leave the value for this password variable blank.
- If you selected to configure Oracle Enterprise Manager, then you must set the value of the `S_HOSTUSERPASSWORD` variable to the password of the Oracle software installation owner. If you do not want to enable Oracle Enterprise Manager, the leave the value for this password variable blank.

Example B-1 Sample Password Response File for Oracle RAC Databases

This example provides a template for a password response file to use with the database configuration assistants.

```
oracle.assistants.server|S_SYSPASSWORD=password
oracle.assistants.server|S_SYSTEMPASSWORD=password
oracle.assistants.server|S_SYSMANPASSWORD=password
oracle.assistants.server|S_DBSNMPPASSWORD=password
oracle.assistants.server|S_HOSTUSERPASSWORD=password
oracle.assistants.server|S_ASMSNMPPASSWORD=password
```

B.7.4 Creating a Password Response File

To create a password response file to use with the configuration assistants, perform the following steps:

1. Create a response file that has a name of the format `filename.properties`.
2. Open the file with a text editor, and cut and paste the sample password file contents, as shown in [Example B-1, "Sample Password Response File for Oracle RAC Databases"](#), modifying as needed.
3. If the file is stored on a volume formatted for Windows New Technology File System (NTFS), then modify the security permissions to secure the file.

B.7.5 Performing Postinstallation Configuration Using a Response File

To run configuration assistants with the `configToolAllCommands` script in silent mode or response file mode, perform the following steps:

1. Change directory to `%ORACLE_HOME%\cfgtoollogs`.
2. Run the configuration script using the following syntax:

```
configToolAllCommands RESPONSE_FILE=\path\name.properties
```

Example B-2 Running Configuration Assistants in Response File Mode

Assume you created a password response file in the `C:\users\oracle\db` directory with a name of `cfg_db.properties`. To run the configuration assistants in response file mode to configure the Oracle software after installation, enter commands similar to the following:

```
C:\> cd %ORACLE_HOME%\cfgtoollogs
C:\..\cfgtoollogs> configToolAllCommands RESPONSE_FILE=C:\users\oracle\db\cfg_db.properties
```

Converting to Oracle RAC and Oracle RAC One Node from Single-Instance Oracle Databases

This appendix describes the procedures for converting from Oracle Database 11g release 2 (11.2) single-instance databases to Oracle Real Application Clusters (Oracle RAC) and Oracle RAC One Node databases.

The procedures in this appendix assume that your original single-instance database and the target Oracle RAC database are using the same release, and running on the same platform. If you are upgrading from an earlier release of Oracle RAC to Oracle RAC 11g release 2, then use the Oracle Database Upgrade Assistant (DBUA).

This appendix contains the following topics:

- [Administrative Issues for Converting Databases to Oracle RAC](#)
- [Converting to Oracle RAC and Oracle RAC One Node Using DBCA](#)
- [Converting Databases to Oracle RAC Using rconfig](#)
- [Preparing to Convert with rconfig and Oracle Enterprise Manager](#)
- [Converting Databases to Oracle RAC Using rconfig](#)
- [Example of rconfig XML Input Files for ConvertToRAC](#)
- [Postconversion Steps](#)

Note: You must use clustered Oracle Automatic Storage Management (Oracle ASM) instances for Oracle RAC databases.

See Also: *Oracle Database Licensing Information* for compliance requirements that apply to the license you have purchased

C.1 Administrative Issues for Converting Databases to Oracle RAC

Note the following administrative considerations before converting single-instance databases to Oracle RAC:

- Backup procedures should be available before converting from a single-instance Oracle Database to Oracle RAC. This includes taking a backup of your existing database before converting to Oracle RAC and being prepared to backup your Oracle RAC database immediately following the conversion.
- For archiving with Oracle RAC environments, the archive file format requires a thread number.

- The archived logs from all instances of an Oracle RAC database are required for media recovery. Because of this, if you archive to a file and you do not use a cluster file system, or some other means to provide shared file systems, then you require a method of accessing the archive logs from all nodes on which the cluster database has instances.
- By default, all database files are migrated to Oracle Managed Files (OMF). This feature simplifies tablespace creation, ensures data file location consistency and compliance with Oracle Flexible Architecture (OFA) rules, and reduces human error with data file management.

See Also: *Oracle Database 2 Day + Real Application Clusters Guide* for information about redo logs and redo threads

C.2 Converting to Oracle RAC and Oracle RAC One Node Using DBCA

You can use Oracle Database Configuration Assistant (DBCA) to convert from single-instance Oracle databases to Oracle RAC or Oracle RAC One Node databases. DBCA automates the configuration of the control file attributes, creates the undo tablespaces and the redo logs, and creates the initialization parameter file entries for cluster-enabled environments. DBCA also configures Oracle Net Services, Oracle Clusterware resources, and the configuration for Oracle RAC database management for use by Oracle Enterprise Manager or the Server Control utility (SRVCTL).

Before you use DBCA to convert a single-instance database to an Oracle RAC or an Oracle RAC One Node database, ensure that your system meets the following conditions:

- Your system uses supported hardware and operating system software. Your system is configured properly to support an Oracle RAC database.
- The nodes have access to shared storage; either Oracle Cluster File System for Windows (OCFS for Windows) or Oracle ASM is available and accessible from all nodes.
- Your applications have no design characteristics that preclude their use with cluster database processing.

If your platform supports a cluster file system, then you can use it for Oracle RAC. You can also convert to Oracle RAC and use a nonshared file system. In either case, Oracle strongly recommends that you use Oracle Universal Installer (OUI) to perform an Oracle Database 11g release 2 (11.2) installation that sets up the Oracle home and inventory in an identical location on each of the selected nodes in your cluster.

This section describes the following scenarios:

- [Converting Oracle Database 11g Installations to Oracle RAC Using DBCA](#)
- [Converting Single Instance on a Cluster to Oracle RAC One Node Using DBCA](#)
- [Converting Single Instance on a Cluster to Oracle RAC Using DBCA](#)
- [Converting an Oracle RAC One Node Database to an Oracle RAC Database](#)

C.2.1 Converting Oracle Database 11g Installations to Oracle RAC Using DBCA

To convert from a single-instance Oracle Database that is on a nonclustered computer to Oracle RAC, perform the procedures described in the following sections, and in the order shown:

- [Use DBCA to Create an Image of the Single-Instance Database](#)

- [Complete the Oracle Clusterware Installation](#)
- [Validate the Cluster](#)
- [Copy the Preconfigured Database Image](#)
- [Install Oracle Database 11g Software with Oracle RAC](#)

C.2.1.1 Use DBCA to Create an Image of the Single-Instance Database

Use DBCA to create a preconfigured image of your single-instance database by using the following procedure:

1. Go to the `bin` directory in `%ORACLE_HOME%`, and start DBCA.
2. At the Welcome page, click **Next**.
3. On the Operations page, select **Manage Templates**, and click **Next**.
4. On the Template Management page, select **Create a database template** and **From an existing database (structure as well as data)**, and click **Next**.
5. On the Source Database page, select the database name in the Database instance list, and click **Next**.
6. On the Template Properties page, enter a name for your template in the Name field. Oracle recommends that you use the database name.

By default, the template files are generated in the directory `%ORACLE_HOME%\assistants\dbca\templates`. You can enter a description of the file in the Description field and change the template file location in the Template data file field.

When you have completed the entries, click **Next**.

7. On the Location of Database Related Files page, select **Maintain the file locations**, so that you can restore the database to the current directory structure, and click **Finish**.

DBCA generates two files: a database structure file (`template_name.dbc`), and a database preconfigured image file (`template_name.dfb`).

C.2.1.2 Complete the Oracle Clusterware Installation

Complete the installation of Oracle Clusterware, as described in *Oracle Grid Infrastructure Installation Guide* for your platform.

C.2.1.3 Validate the Cluster

Validate the cluster configuration using Cluster Verification Utility (CVU). See [Section 1.6, "Confirming Cluster Readiness Using Cluster Verification Utility"](#) for instructions on how to use CVU.

C.2.1.4 Copy the Preconfigured Database Image

Copy the preconfigured database image. This includes copying the database structure `*.dbc` file and the database preconfigured image `*.dfb` file that DBCA created in [Section C.2.1.1, "Use DBCA to Create an Image of the Single-Instance Database"](#) to a temporary location on the node in the cluster from which you plan to run DBCA.

C.2.1.5 Install Oracle Database 11g Software with Oracle RAC

1. Run OUI to perform an Oracle Database installation with Oracle RAC.

2. Select **Cluster Installation Mode** on the Specify Hardware Cluster Installation page of OUI, and select the nodes to include in your Oracle RAC database.
3. On the OUI Database Configuration Types page, select the **Advanced** installation type.

After installing the Oracle Database software, OUI runs postinstallation configuration tools, such as Oracle Net Configuration Assistant (NETCA), DBCA, and so on.

4. On the DBCA Template Selection page, use the template that you copied to a temporary location in [Section C.2.1.4, "Copy the Preconfigured Database Image."](#) Use the browse option to select the template location.

Select the option that you want to deploy. Your choices are the following: Oracle RAC database; Oracle RAC One Node database; or Oracle single instance database.

5. After creating the Oracle RAC database, DBCA displays the Password Management page on which you must change the passwords for database users who have SYSDBA and SYSOPER privileges. When DBCA exits, the conversion process is complete.

Note: Raw devices are not supported with OUI or DBCA in Oracle Database 11g release 2 (11.2). To use raw storage for your database, you must follow the instructions in [Section C.2.3.1.2, "Manual Conversion Procedure."](#)

See Also:

- *Oracle Database Platform Guide for Microsoft Windows* for more information about configuring user roles
- [Chapter 3, "Creating Oracle Real Application Clusters Databases with Oracle Database Configuration Assistant"](#) for more information about using DBCA

C.2.2 Converting Single Instance on a Cluster to Oracle RAC One Node Using DBCA

Use DBCA to convert a single-instance Oracle Database to Oracle RAC One Node by using the following procedure:

1. Go to the directory %ORACLE_HOME%\bin.
2. Start DBCA:

```
c:\..\bin> dbca
```
3. From the Welcome window, select **Oracle Real Application Clusters** database.
4. Use the template that you selected to deploy in [Section C.2.1.5, "Install Oracle Database 11g Software with Oracle RAC."](#)

C.2.3 Converting Single Instance on a Cluster to Oracle RAC Using DBCA

There are three scenarios in which a single-instance database can exist on a server that is a node in a cluster:

- Scenario 1: The Oracle home for the single-instance database was installed on a cluster node and has Oracle RAC enabled. The steps to perform a conversion to

Oracle RAC in this scenario are described in [Section C.2.3.1, "Single-Instance Database an Oracle RAC-Enabled Home"](#).

- Scenario 2: The Oracle home for the single-instance database was installed on a cluster node, but the Oracle RAC feature is disabled for this Oracle home. Refer to [Section C.2.3.2, "Single-Instance Database Using an Oracle RAC-Disabled Home"](#) for instructions on how to convert the database to Oracle RAC in this environment.
- Scenario 3: The Oracle home for the single-instance database was installed on only the local node in a cluster. This happens when you select the Local Installation option on the OUI Specify Hardware Cluster Installation page during the Oracle Database 11g release 2 (11.2) installation. The procedure to convert the single-instance database to Oracle RAC in this scenario is the same as that described in [Section C.2.1, "Converting Oracle Database 11g Installations to Oracle RAC Using DBCA."](#) In the last step, ["Install Oracle Database 11g Software with Oracle RAC"](#), make sure that you select a new Oracle home, not the one from which the single-instance database is running.

C.2.3.1 Single-Instance Database an Oracle RAC-Enabled Home

Perform the following procedures to convert a single-instance database on a cluster node running from an Oracle home that has the Oracle RAC option enabled.

1. Use DBCA to create a preconfigured image of your single-instance database as described in [Section C.2.1.1, "Use DBCA to Create an Image of the Single-Instance Database."](#) To perform the conversion manually, shut down the single-instance database.
2. To add nodes to your cluster, follow the instructions to add and connect these nodes to the cluster as described in "Adding a Node to a Cluster on Windows Systems" in the *Oracle Clusterware Administration and Deployment Guide*. Ensure that all nodes can access the shared storage used by Oracle Clusterware and Oracle RAC.
3. From the existing Oracle home, extend this home to the new nodes using the procedure "Adding Oracle RAC to Nodes with Oracle Clusterware Installed" as described in the *Oracle Real Application Clusters Administration and Deployment Guide*.
4. From a newly added node, configure the listeners on the additional nodes using Oracle Net Configuration Assistant (NETCA). Choose the same port number and protocol that you used on the existing node. If NETCA displays the existing node in the node list page, then do not select this node, because the listener is already configured on it.
5. Convert the database using one of the following procedures:
 - [Automated Conversion Procedure Using DBCA](#)
 - [Manual Conversion Procedure](#)

C.2.3.1.1 Automated Conversion Procedure Using DBCA

If you used DBCA to create a preconfigured image of your single-instance database as described in [Section C.2.1.1, "Use DBCA to Create an Image of the Single-Instance Database,"](#) then perform the following steps to complete the conversion:

1. Start DBCA from the initial node. Select the names of the nodes to include as part of your cluster database. On the Template Selection page, select the preconfigured

template that you created. Enter the database name and respond to the remaining DBCA prompts.

2. Specify the shared storage location for the Oracle Database data files. To use raw devices for the data files, you cannot use DBCA. Instead, follow the steps documented in [Section C.2.3.1.2, "Manual Conversion Procedure."](#)

After creating the Oracle RAC database, DBCA displays the Password Management page on which you must change the passwords for the database users who have SYSDBA and SYSOPER privileges. When DBCA exits, the conversion process is complete.

C.2.3.1.2 Manual Conversion Procedure

If you did not use DBCA to create a preconfigured image of your single-instance database as described in [Section C.2.1.1, "Use DBCA to Create an Image of the Single-Instance Database,"](#) then perform the following steps to complete the conversion:

1. Create the OFA directory structure on each of the nodes that you have added.

See Also: [Section D.2, "Directory Structures for Oracle RAC"](#) for more information about OFA

2. If you are converting single-instance database files on a file system to raw devices, then copy the database data files, control files, redo logs, and SPFILE to their corresponding raw devices using the `ocopy` command. Otherwise, continue to the next step.
3. Re-create the control files by running the `CREATE CONTROLFILE SQL` statement with the `REUSE` keyword and specify `MAXINSTANCES` and `MAXLOGFILES`, and so on, as needed for your Oracle RAC configuration. The `MAXINSTANCES` recommended default is 32.
4. Shut down the database instance.
5. If your single-instance database was using an SPFILE, then create a temporary parameter file (PFILE) from the SPFILE using the following SQL statement:

```
CREATE PFILE='pfile_name' from spfile='spfile_name'
```

6. Set the `CLUSTER_DATABASE` parameter to `TRUE`, set the `INSTANCE_NUMBER` parameter to a unique value for each instance, using the `sid.parameter=value` syntax.

If you optimized memory usage on your single-instance database, then adjust the size of the System Global Area (SGA) to avoid swapping and paging when you convert to Oracle RAC. You should make this adjustment because Oracle RAC requires about 350 bytes for each buffer to accommodate the Global Cache Service (GCS). For example, if you have 10,000 buffers, then Oracle RAC requires about 350 multiplied by 10,000 bytes more memory. Therefore, adjust the size of the SGA by changing the `DB_CACHE_SIZE` and `DB_nK_CACHE_SIZE` parameters accordingly.

7. Start the database instance using the PFILE created in Step 5.
8. If your single-instance database was using automatic undo management, then create an undo tablespace for each additional instance using the `CREATE UNDO TABLESPACE SQL` statement. If you are using raw devices, then ensure that the data file for the undo tablespace is on the raw device.
9. Create redo threads that have at least two redo logs for each additional instance. If you are using raw devices, then ensure that the redo log files are on raw devices.

Enable the new redo threads by using an `ALTER DATABASE SQL` statement. Then, shut down the database instance.

10. Copy the Oracle password file from the initial node, or from the node on which you are working, to the corresponding location on the additional nodes on which the cluster database will have an instance. Replace the `ORACLE_SID` name in each password file appropriately for each additional instance.
11. Set the `REMOTE_LISTENER` parameter to the single client access name (SCAN) and port.
12. Configure the net service entries for the database and instances, and address entries for the `LOCAL_LISTENER` for each instance and for the `REMOTE_LISTENER` in the `tnsnames.ora` file. When you have done this, copy the `tnsnames.ora` file to all nodes.
13. Create the SPFILE from the PFILE using the procedures in [Section 5.5, "Migrating to an SPFILE in Oracle RAC Environments."](#) If you are not using Oracle ASM or a cluster file system, then ensure that the SPFILE is on a raw device that is accessible from every node in the cluster.

14. Create the `Oracle_home\database\initsid.ora` file that contains the following entry, where `spfile_path_name` is the complete path name of the SPFILE:

```
spfile='spfile_path_name'
```

15. On the local node, use SQL*Plus to run `catclust.sql`. This script creates the dictionary views needed for Oracle RAC databases. For example:

```
SQL> start %ORACLE_HOME%\rdbms\admin\catclust.sql
```

16. Add the configuration for the Oracle RAC or Oracle RAC One Node database and its instance-to-node mapping using `SRVCTL`.

- a. To add the configuration of an Oracle RAC database, use the following commands:

```
srvctl add database -d db_name -o Oracle_home -p spfile_path_name
srvctl add instance -d db_name -i inst1_name -n node1_name
srvctl add instance -d db_name -i inst2_name -n node2_name
...
```

- b. To add the configuration of an Oracle RAC One Node database, use the following command:

```
srvctl add database -d db_name -c RACONENODE -o Oracle_home
-p spfile_path_name
```

17. Start the Oracle RAC or Oracle RAC One Node database using `SRVCTL`:

```
C:\> srvctl start database -d db_name
```

After starting the database with `SRVCTL`, your conversion process is complete. You can run the following SQL statement to see the status of all the instances in your Oracle RAC database:

```
SQL> select * from v$active_instances;
```

C.2.3.2 Single-Instance Database Using an Oracle RAC-Disabled Home

You can create a single-instance database on a cluster using an Oracle home with the Oracle RAC option disabled. To create an Oracle home on a cluster with Oracle RAC

disabled, you can select **local** and **non-cluster** on the Node Selection Page of Oracle Universal Installer (OUI) when installing the Oracle Database software. You can also performed a one-node cluster (with Oracle RAC) installation, but later disable the Oracle RAC option.

Perform the following procedures to convert this type of single-instance database to an Oracle RAC or Oracle RAC One Node database:

1. Use DBCA to create a preconfigured image of your single-instance database as described in [Section C.2.1.1, "Use DBCA to Create an Image of the Single-Instance Database."](#) To perform the conversion manually, shut down the single-instance database.
2. Go to the `bin` subdirectory in the Oracle home.
3. Enable the Oracle RAC option by renaming the `orarac11.dll.dbl` file using the following command:

```
C:\..\bin> copy /Y orarac11.dll.dbl orarac11.dll
```
4. Continue with step 2 in [Section C.2.3.1, "Single-Instance Database an Oracle RAC-Enabled Home."](#)

C.2.4 Converting an Oracle RAC One Node Database to an Oracle RAC Database

You can convert an Oracle RAC One Node database to an Oracle RAC database by logging in as the Oracle RAC One Node database owner and entering the following SRVCTL command:

```
C:\> srvctl convert database -c RAC
```

Caution: By default, any named user may create a server pool. To restrict the operating system users that have this privilege, Oracle strongly recommends that you add specific users to the CRS Administrators list. See *Oracle Clusterware Administration and Deployment Guide* for more information about adding users to the CRS Administrators list.

After you run this command, you must create server pools for each database service, in addition to the database server pool. The values for `SERVER_NAMES` for the server pools used by the database services must be set to the node that you converted from an Oracle RAC One Node to an Oracle RAC node. You can use the `CRSCTL` utility or Oracle Enterprise Manager to create and configure the server pools.

Converting an administrator-managed Oracle RAC One Node database to an Oracle RAC database configures all database services so that the single instance is the preferred instance for that service. After you convert the database, you can add instances to your database by using the `SRVCTL add instance` command.

Converting a policy-managed Oracle RAC One Node database to an Oracle RAC database sets all database services to UNIFORM cardinality. It also results in reusing the server pool in which the database currently runs. The conversion reconfigures the database to run on all of the nodes in the server pool. The command does not start any additional instances but executing the `SRVCTL start database` command starts the database on all of the nodes in the server pool.

See Also: The *Oracle Real Application Clusters Administration and Deployment Guide* for more information about policy-managed and administrator-managed databases, server pools, and database services.

C.3 Preparing to Convert with rconfig and Oracle Enterprise Manager

You can use `rconfig`, or Oracle Enterprise Manager to assist with converting a single-instance database installation to an Oracle RAC database. The first of these, `rconfig`, is a command-line utility. Oracle Enterprise Manager Grid Control database administration option, Convert to Cluster Database, provides a GUI-based conversion tool. The following sections describe how to use these conversion tools:

- [Prerequisites for Converting to Oracle RAC Databases](#)
- [Configuration Changes During Oracle RAC Conversion Using rconfig](#)
- [Converting Databases to Oracle RAC Using rconfig or Enterprise Manager](#)
- [Converting Databases to Oracle RAC Using Oracle Enterprise Manager](#)

Note: Before you start the conversion, back up your existing database; you should take a backup of your database before starting any major change.

C.3.1 Prerequisites for Converting to Oracle RAC Databases

Before you convert a single-instance database to an Oracle RAC database, ensure that the following conditions are met for each cluster node that you intend to make an Oracle RAC database node:

- Oracle Clusterware 11g release 2 (11.2) is installed, configured, and running.
- Oracle RAC 11g release 2 (11.2) software is installed.
- The Oracle software has the Oracle RAC option enabled.
- Shared storage, either Oracle Cluster File System or Oracle ASM, is available and accessible from all nodes.
- User equivalence exists for the `oracle` account, or the user account used to install the Oracle software.
- If you intend to use Oracle Enterprise Manager, then the Oracle Management Agent on each node is configured and running, and is configured with cluster and host information.
- You have backed up your existing database.

Note: You must use clustered Oracle ASM instances for Oracle RAC databases. Oracle ASM is installed with the Oracle Grid Infrastructure software.

C.3.2 Configuration Changes During Oracle RAC Conversion Using rconfig

The following changes occur when you convert a single-instance database to Oracle RAC using the `rconfig` utility"

- During the conversion, `rconfig` places the target Oracle RAC database into archive log mode, and enables archiving for the database. If you do not plan to use archive log space, then you can disable archive logging after the conversion has completed.
- For the Shared Storage Type value, if you enter **CFS**, and you use a cluster file system for your single-instance database storage, then `rconfig` converts the environment to use Oracle Managed Files (OMF) for database storage, and places the data files in a subdirectory located under the shared storage location.
- During the conversion, `rconfig` moves database files to a specified shared location, and configures them using OMF.

To avoid using OMF with your converted database, the single-instance database files must be located in a shared file system, and you must indicate that `rconfig` should not move the files.

C.3.3 Converting Databases to Oracle RAC Using rconfig or Enterprise Manager

The following list describes scenarios for converting a single-instance Oracle database to an Oracle RAC database:

- Converting a single-instance Oracle Database 11g release 2 (11.2) database to an Oracle RAC 11g release 2 (11.2) database, running out of the *same* Oracle home and using the same data files as the single-instance database.

In this scenario, run the `rconfig` utility from the Oracle RAC database home, or use the Convert to RAC option on the single-instance database target of Oracle Enterprise Manager Grid Control.

- Converting a single-instance database that uses a release of Oracle Database earlier than Oracle Database 11g release 2 (11.2) to an Oracle RAC 11g release 2 (11.2) database, running out of the same Oracle home and using the same data files as the single-instance database

In this scenario, use OUI and DBUA to update the single-instance database to Oracle Database 11g release 2 (11.2). Then, use `rconfig` or the Oracle Enterprise Manager Convert to RAC option, as described in the preceding scenario.

- Converting a single-instance Oracle Database 11g release 2 (11.2) to an Oracle RAC 11g release 2 (11.2) database, running out of a *different* Oracle home and using the same data files as the single-instance database.

In this scenario, run the `rconfig` utility in the target database home, or use the Convert to RAC option in the single-instance database target of Oracle Enterprise Manager Grid Control. Provide the file storage location when prompted.

- Converting a single-instance Oracle Database 11g release 2 (11.2) database to an Oracle RAC 11g release 2 (11.2) database, running out of a *different* Oracle home, and where the host on which the single-instance database runs is *not* a node used by the Oracle RAC database.

In this scenario, create a clone image of the single-instance database, and move the clone image to a host that is one of the nodes used by the Oracle RAC database. Then, use `rconfig` or the Oracle Enterprise Manager Convert to RAC option, as described in the preceding scenario.

C.3.4 Converting Databases to Oracle RAC Using Oracle Enterprise Manager

You can use Oracle Enterprise Manager Grid Control to convert a single-instance database to an Oracle RAC database. To use this feature, complete the following steps:

See Also: *Oracle Database Upgrade Guide* for information about upgrading existing Oracle Database installations

1. Log in to Grid Control. From the Grid Control Home page, click the **Targets** tab.
2. On the Targets page, click the **Databases** secondary tab, and click the link in the Names column of the database to convert to Oracle RAC.
3. On the Database Instance Home page, click the **Change Database** secondary tab.
4. On the Administration page, in the Database Administration Change Database section, click **Convert to Cluster Database**.
5. Log in as the database user SYS with SYSDBA privileges to the database you want to convert, and click **Next**.
6. On the Convert to Cluster Database: Cluster Credentials page, provide a user name and password for the `oracle` user and password of the target database to be converted. If the target database is using Oracle ASM, then also provide the SYSASM user and password, and click **Next**.
7. On the Hosts page, select the host nodes in the cluster to be cluster members in the installed Oracle RAC database. When you have completed your selection, click **Next**.
8. On the Convert to Database: Options page, select whether you want to use the existing listener and port number, or specify a new listener and port number for the cluster. Also provide a prefix for cluster database instances on the cluster.

When you have finished entering information, click **Next**, or click **Help** if you need assistance in deciding how to enter information.

9. On the Convert to Cluster Database: Shared Storage page, either select the option to use your existing shared storage area, or select the option to have your database files copied to a new shared storage location. Also, decide if you want to use your existing fast recovery area, or copy your recovery files to a new fast recovery area using files managed by Oracle Database.

If you use Oracle ASM, then Oracle recommends that you place the data files and the recovery files in separate failure groups. A failure group is defined by shared hardware, such as a controller shared between two disks, or two disks that are on the same spindle. If two disks share hardware that could fail, making both disks unavailable, then these disks are said to be in the same failure group. If you do not use Oracle ASM, then Oracle recommends that the data files and the recovery files are stored in separate locations, as with separate Oracle ASM failure groups, so that a hardware failure does not affect availability.

See Also: *Oracle Automatic Storage Management Administrator's Guide* for more information about failure groups

When you have finished entering information, click **Next**, or click **Help** if you need assistance in deciding how to enter information.

10. On the Convert to Cluster Database: Review page, review the options you have selected. Click **Submit Job** to proceed with the conversion. To change any of the options you have selected, click **Back**. To cancel the conversion, click **Cancel**.
11. On the Confirmation page, click **View Job** to check the status of the conversion.

C.4 Converting Databases to Oracle RAC Using rconfig

You can use the command-line utility `rconfig` to convert a single-instance database to an Oracle RAC database, or to convert it to an Oracle RAC One Node database, depending on the values you provide in the `ConvertToRAC.xml` file. To use this feature, complete the following steps:

See Also: *Oracle Database Upgrade Guide* for information about upgrading existing Oracle Database installations

1. As the `oracle` user, go to the directory `%ORACLE_HOME%\assistants\rconfig\sampleXMLs`, and open the file `ConvertToRAC.xml` using a text editor.
2. Review the `ConvertToRAC.xml` file, and modify the parameters as required for your system. The XML sample file contains comments that provide instructions for how to configure the file.

Caution: Set the `convert verify="ONLY"` option to perform a test conversion to ensure that a conversion can be completed successfully.

When you have finished modifying parameters, save the file with the a name of the format `filename.xml`. Make a note of the name you select.

3. Go to the directory `%ORACLE_HOME%\bin`, and use the following command to run `rconfig`, where `input.xml` is the name of the XML input file you configured in Step 2:

```
C:\..\bin> rconfig input.xml
```

For example, if you create an input XML file called `convert.xml`, then you would use the following command

```
C:\...\bin> rconfig.bat convert.xml
```

Note: The `Convert verify` option in the `ConvertToRAC.xml` file has three options:

- `Convert verify="YES"`: `rconfig` performs checks to ensure that the prerequisites for single-instance to Oracle RAC conversion have been met before it starts conversion
 - `Convert verify="NO"`: `rconfig` does not perform prerequisite checks, and starts conversion
 - `Convert verify="ONLY"` `rconfig` only performs prerequisite checks; it does not start conversion after completing prerequisite checks
-
-

If performing the conversion fails, then use the following procedure to recover and reattempt the conversion:

1. Attempt to delete the database using the DBCA delete database option.
2. Restore the source database.

3. Review the conversion log, and fix any problems reported by rconfig that may have caused the conversion failure. The rconfig log files are under the rconfig directory in %ORACLE_BASE%\cfgtoollogs.
4. Reattempt the conversion.

C.5 Example of rconfig XML Input Files for ConvertToRAC

The following are two examples of an XML ConvertToRAC input file for the rconfig utility. [Example C-1](#) is an XML input file to convert a single-instance database with Oracle ASM to a policy-managed Oracle RAC database (using server pools) on Oracle ASM storage. [Example C-2](#) is an XML input file to convert a single-instance database with Oracle ASM to an administrator-managed Oracle RAC database.

Example C-1 Example rconfig ConvertToRAC XML file for Policy-Managed Databases

```
<?xml version="1.0" encoding="UTF-8"?>
<n:RConfig xmlns:n="http://www.oracle.com/rconfig"
          xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
          xsi:schemaLocation="http://www.oracle.com/rconfig">
  <n:ConvertToRAC>
    <!-- Verify does a precheck to ensure all pre-requisites are met, before the
    conversion is attempted. Allowable values are: YES|NO|ONLY -->
    <n:Convert verify="YES">
      <!--Specify current OracleHome of non-rac database for SourceDBHome -->
      <n:SourceDBHome>C:\app\oracle\product\11.2.0\db_1</n:SourceDBHome>
      <!--Specify OracleHome where the rac database should be configured. It can be same
      as SourceDBHome -->
      <n:TargetDBHome>C:\app\oracle\product\11.2.0\db_1</n:TargetDBHome>
      <!--Specify SID of non-rac database and credential. User with sysdba role is
      required to perform conversion -->
      <n:SourceDBInfo SID="sales">
        <n:Credentials>
          <n:User>sys</n:User>
          <n>Password>oracle</n>Password>
          <n:Role>sysdba</n:Role>
        </n:Credentials>
      </n:SourceDBInfo>
      <!--Specify the list of existing or new server pools which are used by the Policy
      Managed Cluster Database. -->
      <n:ServerPoolList>
        <n:ExistingServerPool name="custom"/>
        <n:NewServerPool name="newpool" cardinality="2"/>
      </n:ServerPoolList>
      <!--Specify RacOneNode along with servicename to convert database to RACOne Node
      -->
      <!--n:RacOneNode servicename="salesrac1service"/-->
      <!--InstancePrefix is not required for Policy Managed database. If specified, it
      will be ignored. Instance names are generated automatically based on db_unique_
      name for Policy Managed dababase.-->
      <!-- Listener details are no longer needed starting 11.2. Database is registered
      with default listener and SCAN listener running from Oracle Grid Infrastructure
      home. -->
      <!--Specify the type of storage to be used by rac database. Allowable values are
      CFS|ASM. The non-rac database should have same storage type. ASM credentials are
      not needed for conversion. -->
      <n:SharedStorage type="ASM">
        <!--Specify Database Area Location to be configured for rac database.If this field
        is left empty, current storage will be used for rac database. For CFS, this field
```

```

will have directory path. -->
    <n:TargetDatabaseArea>+ASMDG</n:TargetDatabaseArea>
<!--Specify Fast Recovery Area to be configured for rac database. If this field
is left empty, current recovery area of non-rac database will be configured for
rac database. If current database is not using recovery Area, the resulting rac
database will not have a recovery area. -->
    <n:TargetFlashRecoveryArea>+ASMDG</n:TargetFlashRecoveryArea>
    </n:SharedStorage>
    </n:Convert>
</n:ConvertToRAC>
</n:RConfig>

```

Example C-2 Example rconfig ConvertToRAC XML file for Administrator-Managed Databases

```

<?xml version="1.0" encoding="UTF-8"?>
<n:RConfig xmlns:n="http://www.oracle.com/rconfig"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://www.oracle.com/rconfig rconfig.xsd">
    <n:ConvertToRAC>
<!-- Verify does a precheck to ensure all pre-requisites are met, before the
conversion is attempted. Allowable values are: YES|NO|ONLY -->
<n:Convert verify="YES">
<!--Specify current OracleHome of non-rac database for SourceDBHome -->
    <n:SourceDBHome>C:\app\oracle\product\11.2.0\db_1</n:SourceDBHome>
<!--Specify OracleHome where the rac database should be configured. It can be same
as SourceDBHome -->
    <n:TargetDBHome>C:\app\oracle\product\11.2.0\db_1</n:TargetDBHome>
<!--Specify SID of non-rac database and credential. User with sysdba role is
required to perform conversion -->
    <n:SourceDBInfo SID="sales">
    <n:Credentials>
    <n:User>sys</n:User>
    <n>Password>oracle</n>Password>
    <n:Role>sysdba</n:Role>
    </n:Credentials>
    </n:SourceDBInfo>
<!--Specify the list of nodes that should have rac instances running for the Admin
Managed Cluster Database. LocalNode should be the first node in this nodelist. -->
    <n:NodeList>
    <n:Node name="node1"/>
    <n:Node name="node2"/>
    </n:NodeList>
<!--Specify RacOneNode along with servicename to convert database to RACOne Node
-->
<!--n:RacOneNode servicename="salesrac1service"/-->
<!--Instance Prefix tag is optional starting with 11.2. If left empty, it is
derived from db_unique_name.-->
    <n:InstancePrefix>sales</n:InstancePrefix>
<!-- Listener details are no longer needed starting 11.2. Database is registered
with default listener and SCAN listener running from Oracle Grid Infrastructure
home. -->
<!--Specify the type of storage to be used by rac database. Allowable values are
CFS|ASM. The non-rac database should have same storage type. ASM credentials are
no needed for conversion. -->
    <n:SharedStorage type="ASM">
<!--Specify Database Area Location to be configured for rac database.If this field
is left empty, current storage will be used for rac database. For CFS, this field
will have directory path. -->
    <n:TargetDatabaseArea>+ASMDG</n:TargetDatabaseArea>

```

```
<!--Specify Fast Recovery Area to be configured for rac database. If this field
is left empty, current recovery area of non-rac database will be configured for
rac database. If current database is not using recovery Area, the resulting rac
database will not have a recovery area. -->
  <n:TargetFlashRecoveryArea>+ASMDG</n:TargetFlashRecoveryArea>
</n:SharedStorage>
  </n:Convert>
</n:ConvertToRAC>
</n:RConfig>
```

C.6 Postconversion Steps

After completing the conversion, note the following recommendations for Oracle RAC environments, as described in the Oracle RAC documentation:

- Follow the recommendations for using load balancing and transparent application failover, as described in *Oracle Real Application Clusters Administration and Deployment Guide*.
- Use locally managed tablespaces instead of dictionary managed tablespaces to reduce contention and manage sequences in Oracle RAC as described in *Oracle Real Application Clusters Administration and Deployment Guide*.
- Follow the guidelines for using automatic segment space management as described in *Oracle Real Application Clusters Administration and Deployment Guide*.

The buffer cache and shared pool capacity requirements in Oracle RAC are slightly greater than the requirements for single-instance Oracle databases. Therefore, you should increase the size of the buffer cache by about 10%, and the size of the shared pool by about 15%.

Directory Structure for Oracle Real Application Clusters Environments

This appendix describes the directory structures for Oracle Real Application Clusters (Oracle RAC) software environments. This appendix contains the following topics:

- [Understanding the Oracle RAC Directory Structure](#)
- [Directory Structures for Oracle RAC](#)

D.1 Understanding the Oracle RAC Directory Structure

When you install Oracle Database 11g release 2 (11.2) with Oracle RAC, all subdirectories except for the Oracle Inventory directory, the Oracle Automatic Storage Management (Oracle ASM) home (if applicable), and the Oracle Clusterware home, are under a top-level Oracle base directory. The Oracle home and `admin` directories are also located under the Oracle base directory.

See Also: *Oracle Database Installation Guide for Microsoft Windows* for more information about the Oracle home and `admin` directories

D.2 Directory Structures for Oracle RAC

[Table D-1](#) shows the hierarchical directory tree of a sample Optimal Flexible Architecture (OFA)-compliant database for Oracle RAC:

Table D-1 Directory Structure for a Sample OFA-Compliant Environment

Root	Second-Level	Third-Level	Fourth-Level	Description
%ORACLE_BASE%				C:\app\oracle The default ORACLE_BASE directory, where the software was installed by the user oracle.
	installation type			C:\app\oracle\product\11.2.0\db The type of installation under the Oracle base directory. For example, when installing Oracle Database 11g release 2, the value for installation type is product\11.2.0\db.
		%ORACLE_HOME%		The location in which the Oracle Database software is installed. You can also choose to add a counter, if you have multiple installations of the software. For example, if you are creating a second Oracle home for the Oracle Database 11g release 2 software, then the path is as follows: C:\app\oracle\product\11.2.0\dbhome_2 Note that the Oracle database home is under the Oracle base directory path. An Oracle Clusterware directory must <i>not</i> be under the Oracle base directory path where the database executable files are located.
			database	The directory in which the local initialization parameter file is stored for the database.
	admin			The administrative directory. Note that with Oracle Database 11g and later releases, bdump, cdump, and udump files are relocated to the directory associated with %ADR_BASE%.
		db_unique_name		The database unique name; this is the same as dbname when the database name is 8 or fewer characters in length. For example, if your database name is sales, the directory path is: C:\app\oracle\admin\sales
			/hdump /pfile	The dump destinations for the database server.
	%ADR_BASE%		/bdump /cdump /udump	Automatic Diagnostic Repository dump destination trace files. Note that this directory path is set by the initialization parameter DIAGNOSTIC_DEST, and that the path for the Automatic Diagnostic Repository must be located on a shared storage location available to all the nodes. By default, this path is a subset of the Oracle base directory, in the following path: %ORACLE_BASE%\diag\
Oracle Grid Infrastructure for a cluster home (Grid home)				An OFA-compliant path for the Oracle Clusterware home. The default value is: C:\app\11.2.0\grid During the Oracle Grid Infrastructure for a cluster installation, Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) software is installed.
	bin			The subtree for Oracle Clusterware and Oracle ASM executable files.
	network			The subtree for Oracle Net Services configuration files and utilities.

See Also: *Oracle Database Installation Guide for Microsoft Windows* for more information about the Optimal Flexible Architecture standard used on Microsoft Windows platforms

How to Stop Processes in an Existing Oracle Real Application Clusters Database

This appendix describes how to stop all processes in an Oracle Real Application Clusters (Oracle RAC) database, in preparation for adding additional products to an existing database, or in preparation for patch updates. This appendix contains the following topics:

- [Back Up the Oracle RAC Database](#)
- [Shut Down Oracle RAC Databases](#)
- [Stop All Oracle Processes](#)

E.1 Back Up the Oracle RAC Database

Before you make any changes to the Oracle software, Oracle recommends that you create a backup of the Oracle Database installation.

See Also:

- *Oracle Database Upgrade Guide*
- *Oracle Database Backup and Recovery User's Guide*
- *Oracle Database Platform Guide for Microsoft Windows*

E.2 Shut Down Oracle RAC Databases

Shut down any existing Oracle Database instances on each node, with normal or immediate priority. You can use Oracle Enterprise Manager or the Server Control (SRVCTL) utility to shut down the Oracle RAC databases and their instances.

Note: If you are using Oracle Clusterware or Oracle Restart, then you must shut down all Oracle Database instances on all cluster nodes before modifying the Oracle software. If you are performing a patch update, then review the instructions in the Patch Set Notes for detailed instructions.

See Also:

- *Oracle Database Administrator's Guide* for more information about Oracle Restart
- [Section 4.1.3, "Downloading and Installing Patch Updates"](#) for more information about patches and patch sets

E.3 Stop All Oracle Processes

Stop all listener and other processes running in the Oracle home directories where you want to modify the database software.

Note: Before you can shut down Oracle ASM instances, Oracle Clusterware, or Oracle Restart, you must first shut down all database instances that use them.

E.3.1 Stopping Oracle Database Processes

This section provides an overview of what must be done before adding additional products to Oracle Database 11g release 2 (11.2). If you are updating only the Oracle Database software, then you only need to stop the databases instances that run from the Oracle home being updated. If you are performing a patch upgrade, then refer to the Database Patch Set Notes for the patch for additional instructions.

Note: You must perform these steps in the order listed.

1. Shut down any processes in the Oracle home on each node that can access a database.

Note: Before you shut down any processes that are monitored by Oracle Enterprise Manager Grid Control, set a blackout in Grid Control for the processes that you intend to shut down. This is necessary so that the availability records for these processes indicate that the shutdown was planned downtime, rather than an unplanned system outage.

For example:

```
C:\> %ORACLE_HOME%\bin\emctl stop dbconsole
```

2. Shut down all Oracle RAC instances on all nodes. To shut down all Oracle RAC instances for a database, enter the following command, where *db_name* is the name of the database:

```
C:\> %ORACLE_HOME%\bin\srvctl stop database -d db_name
```

See Also:

- *Oracle Database Administrator's Guide* for information about configuring blackouts with Oracle Enterprise Manager
- [Section 4.1.3, "Downloading and Installing Patch Updates"](#) for more information about patches and patch sets

E.3.2 Stopping Oracle Clusterware and Oracle ASM Instances

If you are updating the Oracle Clusterware or Oracle ASM software, then you must first stop all Oracle RAC database instances and all Oracle database instances that use Oracle ASM for storage. After stopping the database instances, shut down the Oracle ASM instances and Oracle Clusterware, as described in the following steps:

Note: You must perform these steps in the order listed.

1. Shut down any processes in the Oracle home on each node that might be accessing a database, for example, Oracle Enterprise Manager Database Control.

Note: Before you shut down any processes that are monitored by Oracle Enterprise Manager Grid Control, set a blackout in Grid Control for the processes that you intend to shut down. This is necessary so that the availability records for these processes indicate that the shutdown was planned downtime, rather than an unplanned system outage.

2. Shut down all Oracle RAC instances on all nodes. To shut down all Oracle RAC instances for a database, enter the following command, where *db_name* is the name of the database:

```
Oracle_home\bin\svrctl stop database -d db_name
```

3. Stop all Oracle Clusterware node applications and resources on all nodes. To stop node applications running on a node, enter the following command, where *C:\app\11.2.0\grid* is the Grid home, and *node* is the name of the node where the applications run:

```
C:\> cd app\11.2.0\grid\bin
C:\..\bin> svrctl stop nodeapps -n node
```

4. Shut down the Oracle Clusterware processes and the Oracle ASM instance on each node by entering the following command on all nodes as a user with Administrator privileges:

```
C:\app\...\bin> crsctl stop cluster
```

To stop Oracle Clusterware and Oracle ASM instances on all nodes in the cluster, you can enter a single command: `crsctl stop cluster -all`.

Note: The `crsctl stop crs` command is used to stop Oracle Restart services for single-instance databases.

5. Shut down the Oracle Clusterware processes from the Services window.

Access the Services window by clicking **Start**, selecting **Control Panel**, selecting **Administrative Tools**, and then selecting **Services**. In the Services window, shut down the following services:

- Oracle Object Service
- OracleClusterVolumeService
- OracleOHService

Note: Depending on your configuration, your nodes might not be running all of the services listed.

Managing Oracle Database Port Numbers

This appendix lists the default port numbers and describes how to change the assigned port after installation. This appendix contains the following topics:

- [About Managing Ports](#)
- [Viewing Port Numbers and Access URLs](#)
- [Port Numbers and Protocols of Oracle Components](#)
- [Changing the Oracle Management Agent Port](#)
- [Changing the Oracle Enterprise Manager Database Control Ports](#)
- [Changing the Oracle XML DB Ports](#)
- [Changing the Oracle Services for Microsoft Transaction Server Port](#)

F.1 About Managing Ports

During installation, Oracle Universal Installer (OUI) assigns port numbers to components from a set of default port numbers. Many Oracle Real Application Clusters (Oracle RAC) components and services use ports. As an administrator, it is important to know the port numbers used by these services, and to ensure that the same port number is not used by two services on your system.

Most port numbers are assigned during installation. Every component and service has an allotted port range, which is the set of port numbers Oracle RAC attempts to use when assigning a port. Oracle RAC starts with the lowest number in the range and performs the following:

- Is the port used by another Oracle Database installation on the system?
The installation can be either active or inactive at the time; Oracle Database can still detect if the port is used.
- Is the port used by a process that is currently running?
This could be any process on the host, including processes other than Oracle Database processes.

If the answer to any of the preceding questions is yes, then Oracle RAC moves to the next highest port in the allotted port range and continues checking until it finds a free port.

F.2 Viewing Port Numbers and Access URLs

In most cases, the Oracle Database component's port number is listed in the tool used to configure the port. In addition, ports for some Oracle Database applications are listed in the `portlist.ini` file. This file is located in the directory `%ORACLE_HOME%\install`.

If you change a port number after installation, then it is not updated in the `portlist.ini` file, so you can rely on this file only immediately after installation. To find or change a port number, use the methods described in this appendix.

F.3 Port Numbers and Protocols of Oracle Components

[Table F-1](#) lists the port numbers and protocols used by components that are configured during the installation. By default, the first port in the range is assigned to the component, if it is available.

Table F-1 Ports Used in Oracle Components

Component and Description	Default Port Number	Port Range	Protocol	Used Only On Interconnect
Cluster Manager The port number is assigned automatically during installation. You cannot view or modify it afterward.	Dynamic	Dynamic	TCP	Yes
Cluster Synchronization Service daemon (CSSD) The Cluster Synchronization Service (CSS) daemon uses a fixed port for node restart advisory messages. This port is used on all interfaces that have broadcast capability. Broadcast occurs only when a node eviction restart is imminent.	42424/ Dynamic	Dynamic	TCP	Yes
Oracle Data Guard Shares the Oracle Net listener port and is configured during installation. To reconfigure this port, use Oracle Net Configuration Assistant (NETCA) to reconfigure the listener.	1521 (same value as the listener)	1521	TCP	No
Oracle Connection Manager Listening port for Oracle client connections to Oracle Connection Manager. You can configure Oracle Connection Manager after installation using NETCA.	1630	1630	TCP	No
Cluster Ready Services daemon (CRSD) Oracle Clusterware Cluster Ready Services (CRS) daemon internode connection. The port number is assigned automatically during installation. You cannot view or modify it afterward.	Dynamic	Dynamic	TCP	Yes
Oracle Cluster Registry The port number is assigned automatically during installation. You cannot view or modify it afterward.	Dynamic	Dynamic	TCP	Yes

Table F-1 (Cont.) Ports Used in Oracle Components

Component and Description	Default Port Number	Port Range	Protocol	Used Only On Interconnect
Oracle EM Database Control - HTTP HTTP port for Oracle Enterprise Manager Database Control. The port is configured during installation. Section F.5, "Changing the Oracle Enterprise Manager Database Control Ports" explains how to modify this port number.	1158	5500–5519	TCP/HTTP	No
Oracle EM Database Control - Java RMI Java Remote Method Invocation (RMI) port for Oracle Enterprise Manager Database Control. The port is configured during installation. Section F.5, "Changing the Oracle Enterprise Manager Database Control Ports" explains how to modify this port number.	5520	5520–5539	TCP	No
Oracle EM Database Control - JMS Oracle Java Message Service (JMS) port for Oracle Enterprise Manager Database Control. The port is configured during installation. Section F.5, "Changing the Oracle Enterprise Manager Database Control Ports" explains how to modify this port number.	5540	5540–5559	TCP	No
Oracle Event Manager (EVM) Generates events for Oracle Clusterware. The port number is assigned automatically during installation. You cannot view or modify it afterward.	Dynamic	Dynamic	TCP	Yes
Oracle HA Services daemon (OHASD) The Oracle High Availability Services (OHAS) daemon starts the Oracle Clusterware stack.	42424	Dynamic	TCP	Yes
Oracle Management Agent HTTP port for Oracle Management Agent, which is part of Oracle Enterprise Manager. The port is configured during installation. Section F.4, "Changing the Oracle Management Agent Port" explains how to modify this port number.	3938	1830–1849	HTTP	No
Oracle Net Listener Allows Oracle clients to connect to the database by using Oracle Net Services. You can configure this port during installation. To reconfigure this port, use NETCA.	1521	1024-65535	TCP	No
Oracle Notification Services (ONS) Port for ONS, used for the publish and subscribe service for communicating information about Fast Application Notification (FAN) events. The FAN notification process uses system events that Oracle Database publishes when cluster servers become unreachable or if network interfaces fail. Use <code>srvctl</code> to modify ONS ports.	6100 (local) 6200 (remote)	Configured manually	TCP	No

Table F-1 (Cont.) Ports Used in Oracle Components

Component and Description	Default Port Number	Port Range	Protocol	Used Only On Interconnect
<p>Oracle Real Application Clusters</p> <p>The port number is assigned automatically during installation. You cannot view or modify it afterward.</p>	Dynamic	Dynamic	TCP	Yes
<p>Oracle Services for Microsoft Transaction Server</p> <p>The port number for Microsoft Transaction Server is configured automatically by Oracle Universal Installer (OUI) the first time you install the software on a particular server. If you install the software in multiple Oracle homes on the same server, then OUI uses the same port number for all installations.</p> <p>In most cases, you do not have to reconfigure the port number. Section F.7, "Changing the Oracle Services for Microsoft Transaction Server Port" explains how to change this port number.</p>	Dynamic	49152 to 65535	TCP	No
<p>Oracle XML DB - HTTP</p> <p>The Oracle XML DB HTTP port is used if web-based applications must access an Oracle database from an HTTP listener. The port is configured during installation, and you <i>cannot</i> view it afterward.</p> <p>Section F.6, "Changing the Oracle XML DB Ports" explains how to change this port number.</p>	0	Configured manually	HTTP	No
<p>Oracle XML DB - FTP</p> <p>The Oracle XML DB FTP port is used when applications must access an Oracle database from an FTP listener. The port is configured during installation, and you <i>cannot</i> view it afterward.</p> <p>Section F.6, "Changing the Oracle XML DB Ports" explains how to change this port number.</p>	0	Configured manually	FTP	No

F.4 Changing the Oracle Management Agent Port

To find the current setting for the Management Agent port, search for `EMD_URL` in the file `%ORACLE_HOME%\host_sid\sysman\config\emd.properties`, where `host_sid` is a string that contains the local host name and the SID for the Oracle RAC database.

To change the Management Agent HTTP port, use the `emca -reconfig ports` command, as shown in this example:

```
emca -reconfig ports -AGENT_PORT 1831
```

F.5 Changing the Oracle Enterprise Manager Database Control Ports

To find the current HTTP, RMI, and JMS port settings, search in the following files, where `host_sid` is a string that contains the local host name and the SID for the Oracle RAC database:

- **HTTP port:** Search for `REPOSITORY_URL` in the `%ORACLE_HOME%\host_sid\sysman\config\emd.properties` file
- **RMI port:** Search for the port attribute in the `rmi-server` tag in the file `%ORACLE_HOME%\oc4j\j2ee\OC4J_DBConsole_host_sid\config\rmi.xml` file

- **JMS port:** Search for the port attribute in the `jms-server` tag in the `%ORACLE_HOME%\oc4j\j2ee\OC4J_DBConsole_host_sid\config\jms.xml` file

To change the Oracle Enterprise Manager Database Control ports, use the `emca -reconfig ports` command:

```
C:\> %ORACLE_HOME%\bin\emca -reconfig ports option setting
```

In the previous example, `option` specifies one or more of the following ports and `setting` is the new port value:

Option	Description	Example
DBCONTROL_HTTP_PORT	Sets the HTTP port	<code>emca -reconfig ports -DBCONTROL_HTTP_PORT 1820</code>
RMI_PORT	Sets the RMI port	<code>emca -reconfig ports -RMI_PORT 5520</code>
JMS_PORT	Sets the JMS port	<code>emca -reconfig ports -JMS_PORT 5521</code>

You can specify multiple port settings in one line, for example:

```
emca -reconfig ports -DBCONTROL_HTTP_PORT 1820 -AGENT_PORT 1831 -RMI_PORT 5520
```

F.6 Changing the Oracle XML DB Ports

By default, the FTP and HTTP (including HTTPS) ports for Oracle XML DB are set to 0, which disables FTP or HTTP access to Oracle XML DB. To change the FTP and HTTP ports for Oracle XML DB, you must run the `catxdbdbca.sql` script, which in a default installation is located in `%ORACLE_HOME%\rdbms\admin`.

To change the Oracle XML DB ports:

1. Check that the Oracle listener is running. In the Windows Services Control Manager, ensure that the Oracle listener service (for example, `OracleOraDb11g_home1TNSListener`) is set to Started.

If you cannot start the listener, then see *Oracle Database Net Services Administrator's Guide*.

2. Log in to SQL*Plus as the SYS or XDB user using the SYSDBA privilege. For example, to log in to SQL*Plus as SYS:

```
SQL> sqlplus sys/ as sysdba
```

3. Run the `catxdbdbca.sql` script.

For example, assuming your Oracle home is located in the `C:\app\oracle\product\11.2.0\db_1` directory, to use 2200 for the FTP port and 8200 for the HTTP port, you would enter the following SQL statement:

```
SQL> @C:\app\oracle\product\11.2.0\db_1\rdbms\admin\catxdbdbca.sql 2200 8200
```

4. Exit SQL*Plus.

See Also: *Oracle Database Administrator's Guide* for information about connecting to the database using SQL*Plus

F.7 Changing the Oracle Services for Microsoft Transaction Server Port

In most cases, you are not required to reconfigure the port number for the Oracle Services for Microsoft Transaction Server. If you must change the port number, then you can use the Registry Editor to edit its value in the `HKEY_LOCAL_MACHINE\SOFTWARE\ORACLE\OracleMTSRecoveryService\Protid_0` Windows Registry key to any available port within the range 1024 to 65535.

During installation, Oracle Universal Installer takes the value for the port from the key, if it exists. Otherwise, a free port ranging from 49152 to 65535 is chosen automatically.

Oracle Enterprise Manager Configuration Assistant Reference

Oracle Enterprise Manager Configuration Assistant (EMCA) provides a command-line interface for configuring Database Control. This appendix contains the following topics:

- [About EMCA](#)
 - [Overview](#)
 - [Operational Notes](#)
- [EMCA Command Reference](#)

See Also: [Section 6.7, "Managing Database Control Using EMCA"](#) for more information using EMCA

G.1 About EMCA

This section contains topics that relate to using the EMCA command-line interface.

- [Overview](#)
- [Operational Notes](#)

G.1.1 Overview

The primary method for configuring an existing Oracle Database 11g release 2 (11.2) database to be managed with Oracle Enterprise Manager Database Control is to use DBCA, but EMCA can be used if a nongraphical user interface is preferred.

Caution: During the database configuration using EMCA, the database is unavailable and users cannot connect to the database or perform operations on the database.

G.1.2 Operational Notes

This section contains information about how to use the EMCA command-line interface.

G.1.2.1 Usage Information

Every EMCA command can be used in Oracle RAC environments, and certain commands are applicable only in cluster setups. To perform cluster commands using

EMCA, use the `-cluster` flag that is available in almost every EMCA operational mode.

EMCA commands are of the form:

```
emca [operation] [mode] [flags] [parameters]
```

For information about the current cluster configuration, you can run the following command:

```
emca -displayConfig dbcontrol -cluster
```

G.1.2.2 EMCA Parameters

Depending upon the parameters you include on the EMCA command line, EMCA prompts you for additional information required to configure Database Control.

Each EMCA command can include optional parameters to specify the additional information required by EMCA. Parameters used with EMCA are of the form `[-parameterName parameterValue]`. Multiple parameters can be used in combination at the command line. [Table G-1](#) lists the optional parameters that can be used with EMCA.

Note: You can view a list of the available parameters by entering `emca -help` at the command line.

Table G-1 EMCA Optional Command-Line Parameters

Parameter	Description
<code>-respFile</code>	The path of an input file that lists parameters for EMCA to use while performing its configuration operation. For more information, see Section 6.6.2, "Using EMCA in Silent Mode With an Input File."
<code>-SID</code>	Database system identifier.
<code>-PORT</code>	Port number for the listener servicing the database.
<code>-ORACLE_HOME</code>	Database Oracle home, as an absolute path.
<code>-LISTENER_OH</code>	Oracle home from where the listener is running. If the listener is running from an Oracle home other than the one on which the database is running, then the parameter <code>LISTENER_OH</code> must be specified.
<code>-HOST_USER</code>	Host computer user name (for automatic backup).
<code>-HOST_USER_PWD</code>	Host computer user password (for automatic backup).
<code>-BACKUP_SCHEDULE</code>	Schedule in the form of "HH:MM" (for daily automatic backups).
<code>-EMAIL_ADDRESS</code>	E-mail address for notifications.
<code>-MAIL_SERVER_NAME</code>	Outgoing mail server, or Simple Mail Transfer Protocol (SMTP) server, for notifications.
<code>-ASM_OH</code>	Oracle Automatic Storage Management (Oracle ASM) home.
<code>-ASM_SID</code>	System identifier for Oracle ASM instance.
<code>-ASM_PORT</code>	Port number for the listener servicing the Oracle ASM instance.
<code>-ASM_USER_ROLE</code>	User role for connecting to the Oracle ASM instance.
<code>-ASM_USER_NAME</code>	User name for connecting to the Oracle ASM instance.

Table G–1 (Cont.) EMCA Optional Command-Line Parameters

Parameter	Description
-ASM_USER_PWD	Password for connecting to the Oracle ASM instance.
-DBSNMP_PWD	Password for DBSNMP user.
-SYSMAN_PWD	Password for SYSMAN user.
-SYS_PWD	Password for SYS user.
-SRC_OH	Oracle home of the database with Oracle Enterprise Manager configuration to be upgraded or restored.
-DBCONTROL_HTTP_PORT	The port you use to display the Database Control console in your web browser. For more information, see Section 6.7.3, "Changing the Ports Used by Database Control."
-AGENT_PORT	The Oracle Management Agent port for Database Control. For more information, see Section 6.7.3, "Changing the Ports Used by Database Control."
-RMI_PORT	The RMI port for Database Control. For more information, see Section 6.7.3, "Changing the Ports Used by Database Control."
-JMS_PORT	The JMS port for Database Control. For more information, see Section 6.7.3, "Changing the Ports Used by Database Control."
-CLUSTER_NAME	Cluster name (for cluster databases).
-DB_UNIQUE_NAME	Database unique name (for cluster databases).
-SERVICE_NAME	Database service name (for cluster databases).
-EM_NODE	Node from which Database Control console is to be run (for cluster databases). For more information, see Section 6.7.1, "Changing the Database Control Implementation for Your Oracle RAC Database."
-EM_NODE_LIST	Comma-separated list of nodes for agent-only configurations, uploading data to -EM_NODE. For more information, see Section 6.7.1, "Changing the Database Control Implementation for Your Oracle RAC Database."

G.1.2.3 Using Tool/Utility Help

To see a list of EMCA options, start EMCA using any of the following commands:

```
emca -h
emca --h
emca -help
emca --help
```

These commands display the Help message for the EMCA utility. These help commands also list the command arguments described in [Table G–2](#), and the optional parameters described in [Table G–1](#).

G.2 EMCA Command Reference

This section describes the commands available with the EMCA utility.

Table G–2 EMCA Commands

Command	Description
addInst on page G-5	Configures Oracle Enterprise Manager for a new cluster instance of a database or Oracle ASM storage.

Table G-2 (Cont.) EMCA Commands

Command	Description
addNode on page G-5	Configures Oracle Enterprise Manager for a new cluster instance of a database or Oracle ASM storage.
config all db on page G-6	Configures both Database Control and central agent management for a database.
config centralAgent on page G-7	Configures central agent management for a database or an Oracle ASM instance.
config dbcontrol db on page G-7	Configures Database Control for a database.
deconfig all db on page G-8	Deconfigures both Database Control and central agent management for a database.
deconfig centralAgent on page G-9	Deconfigures central agent management for a database or an Oracle ASM instance.
deconfig dbcontrol db on page G-10	Deconfigures Database Control for a database.
deleteInst on page G-11	Deconfigures Oracle Enterprise Manager for a specific instance of a cluster database or Oracle ASM storage.
deleteNode on page G-11	Deconfigures Oracle Enterprise Manager for a node getting deleted.
displayConfig dbcontrol on page G-12	Displays information about the current deployment configuration of Database Control in a cluster environment.
reconfig dbcontrol on page G-13	Reconfigures the Database Control deployment for a cluster database.
reconfig ports on page G-13	Explicitly reassigns Database Control ports.
restore on page G-14	Restores the current release of Oracle Enterprise Manager configuration to an earlier release.
upgrade on page G-14	Upgrades the configuration of an release version of Oracle Enterprise Manager to the current release.
version on page G-15	Prints the release information associated with EMCA.

G.2.1 addInst

Use this command to configure Oracle Enterprise Manager for a new cluster instance of a database or Oracle ASM.

G.2.1.1 Syntax

There are two forms for this command, depending on whether you are adding an Oracle database instance or an Oracle ASM instance:

```
emca -addInst db [-silent] [parameters]
```

```
emca -addInst asm [-silent] [parameters]
```

G.2.1.2 Parameters/Options

Table G–3 *addInst Options*

Parameter	Description
db	<p>Performs the operation for a database (including cluster databases).</p> <p>Use this option for databases that use Oracle ASM to store the data files. If a database is using Oracle ASM, then all the EMCA configuration commands (except for <code>-upgrade</code> and <code>-restore</code>) detect this automatically and apply the changes to both the database and Oracle ASM instances.</p>
asm	Performs the operation for an Oracle ASM-only instance (including cluster Oracle ASM instances).
-silent	<p>Performs the operation without prompting for additional information.</p> <p>If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.</p>
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.2 addNode

Use this command to configure Oracle Enterprise Manager for a new cluster instance of a database or Oracle ASM.

For more information, see [Section 6.3, "Understanding EMCA and Enterprise Manager Control with Oracle RAC."](#)

G.2.2.1 Syntax

```
emca -addNode db [-silent] [parameters]
```

```
emca -addNode asm [-silent] [parameters]
```

G.2.2.2 Parameters

Table G–4 *addNode Options*

Parameter	Description
db	<p>Performs the operation for a database (including cluster databases).</p> <p>Use this option for databases that use Oracle ASM to store the data files. If a database is using Oracle ASM, then all the EMCA configuration commands (except for <code>-upgrade</code> and <code>-restore</code>) detect this automatically and apply the changes to both the database and Oracle ASM instances.</p>
asm	Performs the operation for an Oracle ASM-only instance (including cluster Oracle ASM instances).

Table G–4 (Cont.) addNode Options

Parameter	Description
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.2.3 Examples

This command is discussed further in [Section 6.7.2, "Updating Database Control Configuration When Adding and Removing Nodes."](#)

G.2.3 config all db

Use this command to configure both Database Control and Oracle Management Agent for a database. Options include creating (or re-creating) the Database Control repository, configuring automatic backups, and performing these operations on a cluster database.

This operation configures the database so that it can be centrally managed by the Oracle Enterprise Manager 10g Grid Control console. To use this option, you must have previously installed the Oracle Management Service component of Oracle Enterprise Manager on a network host. In addition, the Oracle Management Agent must be installed on the host where you run the database.

G.2.3.1 Syntax

```
emca -config all db [-repos (create | recreate)] [-cluster] [-silent]
      [-backup] [parameters]
```

G.2.3.2 Parameters

Table G–5 config all db Options

Parameter	Description
-repos create	Creates a new Oracle Management Repository.
-repos recreate	Deletes the current Oracle Management Repository and then re-creates a new one.
-cluster	Performs the operation for a cluster database or Oracle ASM instance.
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.4 config centralAgent

Use this command to configure the Management Agent for a database or an Oracle ASM instance. Options include performing this operation on a cluster environment.

This operation configures the database so that it can be centrally managed by the Oracle Enterprise Manager 10g Grid Control console. To use this option, you must have previously installed the Oracle Management Service component of Oracle Enterprise Manager on a network host. In addition, the Oracle Management Agent must be installed on the host where you run the database.

G.2.4.1 Syntax

```
emca -config centralAgent (db | asm) [-cluster] [-silent] [parameters]
```

G.2.4.2 Parameters

Table G-6 *config centralAgent Options*

Parameter	Description
db	Performs the operation for a database (including cluster databases). Use this option for databases that use Oracle ASM to store the data files. If a database is using Oracle ASM, then all the EMCA configuration commands (except for <code>-upgrade</code> and <code>-restore</code>) detect this automatically and apply the changes to both the database and Oracle ASM instances.
asm	Performs the operation for an Oracle ASM-only instance (including cluster Oracle ASM instances).
-cluster	Performs the operation for a cluster database or Oracle ASM instance.
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File." for more information.
<i>parameters</i>	See Table G-1, "EMCA Optional Command-Line Parameters" .

G.2.5 config dbcontrol db

Use this command to configure Oracle Enterprise Manager Database Control for a database. Options include creating (or re-creating) the Oracle Management Repository, configuring automatic backups, and performing these operations on a cluster database.

G.2.5.1 Syntax

```
emca -config dbcontrol db [-repos (create | recreate)] [-cluster] [-silent]
[-backup] [parameters]
```

G.2.5.2 Options

Table G–7 *config dbcontrol Options*

Parameter	Description
-repos create	Creates a new Management Repository.
-repos recreate	Deletes the current Management Repository and then creates a new one.
-cluster	Performs the operation for a cluster database or Oracle ASM instance.
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
-backup	Configures automatic backup for a database. EMCA prompts for daily automatic backup options. The default Oracle Enterprise Manager settings are used to back up the database files. If you use this option, then EMCA uses the value of the <code>db_recovery_file_dest</code> initialization parameter to identify the fast recovery area for the automated backups. If that parameter is not set, then EMCA generates an error. You can modify these settings later using the Maintenance page in Database Control.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.5.3 Examples

Example 1 Configuring Database Control for a Single Oracle ASM Instance

To configure the Database Control console for single-instance database using Oracle ASM, run the following command to configure the Database Control console and automatically detect the Oracle ASM instance:

```
emca -config dbcontrol db -repos create
```

Example 2 Configuring Automatic Backups of Your Database

Enter the following command to configure Database Control so it performs automatic daily backups of your database:

```
emca -config dbcontrol db -backup
```

G.2.6 deconfig all db

Deconfigures both Database Control and the Oracle Management Agent for a database. Options include deleting the Oracle Management Repository and performing these operations on a cluster database.

G.2.6.1 Syntax

```
emca -deconfig all db [-repos drop] [-cluster] [-silent] [parameters]
```

G.2.6.2 Parameters

Table G–8 *deconfig all db Options*

Parameter	Description
-repos drop	Deletes the current Management Repository.
-cluster	Performs the operation for a cluster database or Oracle ASM instance
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.7 deconfig centralAgent

Use this command to deconfigure the Oracle Management Agent for a database or an Oracle ASM instance. Options include performing this operation on a cluster environment.

For example, you might use this command to remove the Management Agent configuration from a database you are planning to delete. In such a scenario, remove the Management Agent configuration before physically deleting the database. This operation does not remove the actual database or its data files.

G.2.7.1 Syntax

```
emca -deconfig centralAgent (db | asm) [-cluster] [-silent] [parameters]
```

G.2.7.2 Parameters

Table G–9 *deconfig centralAgent Options*

Parameter	Description
db	Performs the operation for a database (including cluster databases). Use this option for databases that use Oracle ASM to store the data files. If a database is using Oracle ASM, then all the EMCA configuration commands (except for <code>-upgrade</code> and <code>-restore</code>) detect this automatically and apply the changes to both the database and Oracle ASM instances.
asm	Performs the operation for an Oracle ASM-only instance (including cluster Oracle ASM instances).
-cluster	Performs the operation for a cluster database or Oracle ASM instance.

Table G–9 (Cont.) deconfig centralAgent Options

Parameter	Description
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.7.3 Examples

Example 1 Removing the Management Agent Configuration from a Database Being Deleted

If you are planning to delete the `testdb` database, then you would use a command similar to the following to remove the Management Agent configuration before physically deleting the database:

```
emca -deconfig centralAgent db -ORACLE_HOME C:\app\oracle\10.2\db1
-SID testdb
```

G.2.8 deconfig dbcontrol db

Use this command to deconfigure Database Control for a database. Options include deleting the Oracle Management Repository and performing these operations on a cluster database.

This operation does not remove the actual database or its data files

G.2.8.1 Syntax

```
emca -deconfig dbcontrol db [-repos drop] [-cluster] [-silent] [parameters]
```

G.2.8.2 Parameters

Table G–10 deconfig dbcontrol db Options

Parameter	Description
-repos drop	Deletes the current Management Repository.
-cluster	Performs the operation for a cluster database or Oracle ASM instance.
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.8.3 Examples

Example 1 Removing the Database Control Configuration from a Database Being Deleted

If you are planning to delete the `testdb` database, then you would use a command similar to the following to remove the Database Control configuration before physically deleting the database:

```
emca -deconfig dbcontrol db -repos drop -ORACLE_HOME C:\app\oracle\10.2\db1
-SID testdb
```

G.2.9 deleteInst

Use this command to deconfigure Oracle Enterprise Manager for a specific instance of a cluster database or Oracle ASM.

G.2.9.1 Syntax

```
emca -deleteInst (db | asm) [-silent] [parameters]
```

G.2.9.2 Parameters

Table G–11 *deleteInst Options*

Parameter	Description
db	Performs the operation for a database (including cluster databases). Use this option for databases that use Oracle ASM to store the data files. If a database is using Oracle ASM, then all the EMCA configuration commands (except for <code>-upgrade</code> and <code>-restore</code>) detect this automatically and apply the changes to both the database and Oracle ASM instances.
asm	Performs the operation for an Oracle ASM-only instance (including cluster Oracle ASM instances).
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
parameters	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.10 deleteNode

Use this command to deconfigure Oracle Enterprise Manager for a node being deleted. For more information, see [Section 6.3, "Understanding EMCA and Enterprise Manager Control with Oracle RAC."](#)

G.2.10.1 Syntax

```
emca -deleteNode (db | asm) [-silent] [parameters]
```

G.2.10.2 Parameters

Table G–12 *deleteNode Options*

Parameter	Description
db	Performs the operation for a database (including cluster databases). Use this option for databases that use Oracle ASM to store the data files. If a database is using Oracle ASM, then all the EMCA configuration commands (except for <code>-upgrade</code> and <code>-restore</code>) detect this automatically and apply the changes to both the database and Oracle ASM instances.
asm	Performs the operation for an Oracle ASM-only instance (including cluster Oracle ASM instances).
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.10.3 Examples

This command is discussed further in [Section 6.7.2, "Updating Database Control Configuration When Adding and Removing Nodes."](#)

G.2.11 displayConfig dbcontrol

Use this command to display information about the current deployment configuration of Database Control in a cluster environment. This command must be used with the `-cluster` option.

G.2.11.1 Syntax

```
emca -displayConfig dbcontrol -cluster [-silent] [parameters]
```

G.2.11.2 Parameters

Table G–13 *displayConfig dbcontrol Options*

Parameter	Description
-cluster	Performs the operation for a cluster database or Oracle ASM instance.
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.12 reconfig dbcontrol

Use this command to reconfigure Database Control deployment for a cluster database. This command must be used with the `-cluster` option.

G.2.12.1 Syntax

```
emca -reconfig dbcontrol -cluster [-silent] [parameters]
```

G.2.12.2 Parameters

Table G–14 *reconfig dbcontrol Options*

Parameter	Description
<code>-cluster</code>	Performs the operation for a cluster database or Oracle ASM instance.
<code>-silent</code>	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.12.3 Examples

This command is discussed further in [Section 6.7.1, "Changing the Database Control Implementation for Your Oracle RAC Database."](#)

G.2.13 reconfig ports

Use this command to explicitly reassign Database Control ports. Options include performing this operation on a cluster environment.

G.2.13.1 Syntax

```
emca -reconfig ports [-cluster] [parameters]
```

G.2.13.2 Parameters

Table G–15 *reconfig ports Options*

Parameter	Description
<code>-cluster</code>	Performs the operation for a cluster database or Oracle ASM instance.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.13.3 Examples

This command is discussed further in [Section 6.7.3, "Changing the Ports Used by Database Control."](#)

G.2.14 restore

Use this command to restore the current release of Oracle Enterprise Manager configuration to an earlier release. This is the inverse of the `-upgrade` command. Executing this command reverses any changes that resulted from an `-upgrade` operation. This operation can be performed for database, Oracle ASM, or database and Oracle ASM instances simultaneously.

G.2.14.1 Syntax

```
emca -restore (db | asm | db_asm) [-cluster] [-silent] [parameters]
```

G.2.14.2 Parameters

Table G–16 *restore Options*

Parameter	Description
db	Performs the operation for a database (including cluster databases).
asm	Performs the operation for an Oracle ASM-only instance (including cluster Oracle ASM instances).
db_asm	Performs the restore operation for a database and an Oracle ASM instance. Restoring an Oracle ASM instance does not require restoring the database instances it services. Therefore, the Oracle Enterprise Manager configuration can be restored separately for a database and its respective Oracle ASM instance.
-cluster	Performs the operation for a cluster database or Oracle ASM instance.
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.15 upgrade

Use this command to upgrade the configuration of an earlier release of Oracle Enterprise Manager to the current release. This operation can be performed for database, Oracle ASM, or database and Oracle ASM instances simultaneously.

This command does not upgrade the actual database or Oracle ASM instances, nor does it upgrade the Oracle Enterprise Manager software. Instead, it upgrades the configuration files for the specified instance so that they are compatible with the current release of the Oracle Enterprise Manager software. EMCA attempts to upgrade all instances of the specified database or Oracle ASM target on the host, across all Oracle homes (because it is likely that certain target properties, such as listener port or Oracle home, have changed).

G.2.15.1 Syntax

```
emca -upgrade (db | asm | db_asm) [-cluster] [-silent] [parameters]
```

G.2.15.2 Parameters

Table G–17 *upgrade Options*

Parameter	Description
db	Performs the operation for a database (including cluster databases).
asm	Performs the operation for an Oracle ASM-only instance (including cluster Oracle ASM instances).
db_asm	Performs the upgrade operation for a database and an Oracle ASM instance. Upgrading an Oracle ASM instance does not require upgrading the database instances it services. Therefore, the Oracle Enterprise Manager configuration can be upgraded separately for a database and its respective Oracle ASM instance.
-cluster	Performs the operation for a cluster database or Oracle ASM instance.
-silent	Performs the operation without prompting for additional information. If this mode is specified, then all the required parameters must be entered at the command line or specified in an input file using the <code>-respFile</code> argument. See Section 6.6.2, "Using EMCA in Silent Mode With an Input File" for more information.
<i>parameters</i>	See Table G–1, "EMCA Optional Command-Line Parameters" .

G.2.16 version

Use this command to print the release information associated with EMCA.

G.2.16.1 Syntax

```
emca -version
```


A

ACFS. *See* Oracle ACFS.

Advanced

database configuration type, 2-3, 2-4, C-4

installation type, 2-2, 2-4, 2-5, C-4

architectures

different chip architectures, 1-8

Optimal Flexible Architecture (OFA), xxix, 1-9,
2-7, 3-8, C-2, C-6, D-1

stripe and mirror everything (SAME), xxix

archive logs

and converting to cluster databases, C-1

ASM Dynamic Volume Manager (AVMD), 1-12

ASM. *See* Oracle ASM.

ASMSNMP, 1-16

AUTHENTICATION_SERVICES, 7-12

Automatic Diagnostic Repository (ADR), 7-12

Automatic Storage Management. *See* Oracle ASM.

automatic undo management, 7-5

and converting a single-instance database to

Oracle RAC, C-6

and DBCA, 3-2, C-2

tablespaces, 7-3

AVMD. *See* ASM Dynamic Volume Manager
(AVMD).

B

b_cr91data, 2-7

backups

and converting to cluster database, C-1

automated, 1-12, 1-13, 6-4, 6-5, 6-6, B-5, G-2, G-6,
G-7

configuring with EMCA, G-8

space requirements when using DBCA, 1-12

browsers. *See* web browsers.

C

candidate disks, 3-4

changing

database listener port, 6-10

Oracle Enterprise Manager Database Control
ports, F-4

Oracle Management Agent port, F-4

Oracle XML DB ports, F-5

product languages, 2-2

changing host names, 6-11, 7-1

cluster databases

archive logs, C-2

converting to using DBCA, C-2

converting to using Enterprise Manager Grid
Control, C-10

converting to using rconfig, C-9

creating using DBCA, C-5

installed configuration, 7-3

load balancing, 7-12

password file, C-7

SCAN, 7-12

SPFILE, 5-1, 7-5

cluster file systems

See also OCFS for Windows

See also Oracle ACFS

storage options for data files, 1-13

Cluster Manager

ports, ranges, and protocols, F-2, F-4

Cluster Verification Utility

checking the integrity of the Cluster
Manager, A-4

DBCA database creation stage readiness
check, 3-3

download location, 1-4

incorporated into OUI, 1-4

installed location, 1-18

overview, 1-4

verifying readiness for database installation, 1-18

CLUSTER_DATABASE_INSTANCES, 3-7

clusterware diagnostics, A-4

commands

crsctl, 4-11

emca, 4-6

olsnodes, 4-11, A-4

components

created when using DBCA, 7-2

new for Oracle Database 11g release 1, xxvi

Oracle Configuration Manager, xxv

Oracle Database Vault, xxv

configuration assistants

failed status, A-5

troubleshooting, A-4

configuration types

- Advanced, 2-3
- connect descriptors, 7-6
- connection load balancing, 7-12
- Connection Manager. *See* Oracle Connection Manager
- control files, 1-16
 - described, 7-4
 - installed configuration, 7-4
- converting to Oracle RAC
 - administrative considerations, C-1
 - from nonclustered systems, C-2
 - from single-instance databases, B-1, C-1, C-4, C-9, C-10
 - postconversion steps, C-15
 - reasons not to convert, C-2
 - with rconfig, C-9
- CPU, 4-2
- creating Oracle RAC databases
 - using DBCA, 2-7, 3-1, 3-5
 - using scripts, B-1
- Critical Patch Updates, 4-2
- Customer Configuration Repository. *See* Oracle Configuration Manager.

D

- data files, 1-16
 - and DBCA, 7-3
 - described, 7-3
 - storage options, 1-13
- Data Warehouse
 - database configuration type, 2-3
- database
 - components created by DBCA, 7-2
 - configuration types, 2-2, 2-3
 - configuring automatic daily backups, 6-5
 - creation using scripts, B-1
 - services, 2-7, 7-6
 - supported storage options for files, 1-14
- Database Agent, 7-14
 - and listeners, 7-7
 - and server parameter files, 5-2
- Database Agent process, 7-7
- Database Configuration Assistant. *See* Oracle Database Configuration Assistant (DBCA).
- Database Control
 - automatic daily backups with EMCA, 6-5
 - changing the host name or IP address, 6-11
 - changing the listener port, 6-10
 - configuration for Oracle Real Application Clusters, 6-6
 - configuring after installation, 6-5, 6-6
 - configuring on new nodes using EMCA, 6-8
 - configuring on remote nodes using EMCA, 6-7
 - configuring using DBCA, 6-4
 - configuring with EMCA, 6-5
 - displaying the current configuration, G-2
 - email notifications, 3-6
 - modifying the existing configuration, 6-7
 - port used, 6-9
 - port used by the Oracle Management Agent, 6-10

- reconfiguring after changing TNS_ADMIN, 6-11
- removing from a node using EMCA, 6-8
- specifying ports, 6-9
- Database Upgrade Assistant (DBUA), 3-3, C-1
- daylight savings time
 - and TIMESTAMP WITH TIME ZONE data, 1-16
- DB_NAME, 3-7
- DB_UNIQUE_NAME, 3-7
 - initialization parameters, 3-7
- dbca.rsp file, B-5
- DBCA. *See* Oracle Database Configuration Assistant (DBCA).
- DBSNMP user
 - password requirements, 2-6
- dedicated servers, 7-16
- Deinstallation tool, 8-1
 - example of using the, 8-6
 - log files, 8-5
- deleting databases
 - using DBCA, 3-8
- deprecated features
 - for Oracle Database 11g release 1, xxxi
 - initialization parameters, xxix
 - service management, 2-7
- desupported
 - raw devices, xxxi
- diagnostics, A-4
- directory objects, 4-4
- directory structures, D-1
- domain users
 - used in installation, 1-9
- downloading installer updates before installing from a secure data center, 1-3
- DST patches, 4-2
- DV_ACCTMGR, 2-4
- DV_OWNER, 2-4

E

- Easy Connect Naming, xxiii, 7-12
- email notifications
 - with Database Control, 3-6
- EMCA. *See* Enterprise Manager Configuration Assistant (EMCA).
- Enterprise Manager Configuration Assistant (EMCA) and Oracle RAC, 6-3
 - changing the host name or IP address, 6-11
 - concepts, 6-6
 - configuring automatic daily backups, 6-5
 - configuring Database Control on newly added nodes, 6-8
 - configuring Database Control on remote nodes, 6-7
 - configuring Database Control with, 6-5
 - displaying current Database Control configuration, G-2
 - reconfiguring after changing TNS_ADMIN, 6-11
 - reconfiguring Database Control after changing the listener port, 6-10
 - removing Database Control from a node, 6-8

- specifying port assignments, 6-9
- troubleshooting problems with Database Control, 6-10
- using an input file for EMCA parameters, 6-6
- Enterprise Manager Control (EMCTL), 6-3
- enterprise.rsp file, B-5
- error messages
 - during installation, A-2
 - silent mode, A-5
- external files, 4-4
- external tables, 4-4

F

- failover
 - and service registration, 7-14
- failure group, 3-6
- fast recovery area, 3-6
 - and files managed by Oracle, C-11
- features, new, xix, xxiv
- file systems
 - storage options for data files, 1-13
 - supported options for Oracle Database, 1-13
- files
 - dbca.rsp, B-5
 - enterprise.rsp, B-5
 - response files, B-4

G

- General Purpose/Transaction Processing database configuration type, 2-3
- global database names
 - selecting, 2-5
- globalization
 - support for, 1-5, 2-2
- Grid Control
 - updating targets after upgrade to Oracle Database 11g release 2, 6-10
- Grid Plug and Play, xxiii
- groups
 - ORA_DBA, 1-9, A-3, B-5

H

- high availability
 - SCAN listeners, 7-13
 - stopping services, E-3
- host names
 - changing, 7-1
 - reconfiguring Database Control after changing, 6-11

I

- initialization parameter files, 7-5
 - for instances, 5-1
 - listeners
 - parameters, 7-7
 - See also* parameter files (PFILE).
 - See also* server parameter files (SPFILE).

- initialization parameters
 - CLUSTER_DATABASE_INSTANCES, 3-7
 - DB_NAME, 3-7
 - DISPATCHERS, 7-16
 - LOCAL_LISTENER, 7-14
 - REMOTE_LISTENER, 7-14
- installActions.log files, A-2
- installation
 - and globalization, 1-5, 2-2
 - and Oracle Configuration Manager, 2-5
 - and Oracle Database Vault, 1-15
 - directory structure, D-1
 - listener.ora file, 7-7
 - noninteractive, B-1
 - noninteractive mode error handling, A-6
 - of additional products after installation is completed, 3-8
 - response files, B-4
 - errors, A-5
 - preparing, B-4, B-6
 - templates, B-4
 - silent mode, B-7
 - silent mode error handling, A-6
 - tnsnames.ora files, 7-9
- instances
 - initialization parameter files, 5-1

J

- Java Message Service (JMS)
 - port used with Database Control, 6-9
- Java Runtime Environment
 - Windows requirements, 1-11

L

- language. *See* globalization.
- language support and installation, 2-2
- List of Cluster Databases page, 3-8
- listeners
 - automatic migration from earlier releases, 3-3
 - configuring
 - listener.ora file, 7-7
 - local listeners, 7-14
 - remote listeners, 7-14
 - default configuration of listener.ora, 7-7
 - files
 - listener.ora, 7-7
 - sqlnet.ora, 7-12
 - parameters
 - LOCAL_LISTENER, 7-10, 7-14
 - REMOTE_LISTENER, 7-10, 7-14
 - registration, 7-6
 - service registration, 7-8
- load balancing
 - and service registration, 7-12
- local administrator users, 1-9
- local groups
 - ORA_DBA, 1-9
- local listeners, 7-7

local PFILE, 5-2
log files
 how to see the log files during installation, 2-1
 troubleshooting with, A-2, A-4

M

Management Options page
 in DBCA, 6-4, G-1
managing tablespaces, 7-5
Metalink. *See* My Oracle Support web site.
migration
 from single-instance databases to Oracle RAC, C-2
multiple Oracle homes
 and Oracle Clusterware, 1-5
multiple Oracle RAC databases
 clusterware requirements for, 1-5
My Oracle Support web site, 4-2
 about, 1-2
 accessing, 1-2
 and storage requirements, 1-11
 check before installation, 1-2

N

naming methods, 7-6
 Easy Connect Naming, 7-12
National Language Support (NLS)
 compatibility files, 2-7
Net Configuration Assistant. *See* Oracle Net Configuration Assistant (NETCA)
net service names, 7-10
NETCA. *See* Oracle Net Configuration Assistant (NETCA).
netca.rsp file, B-5
network attached storage (NAS), 1-13
networks
 configuration files, 7-7
 sqlnet.ora, 7-12
 tnsnames.ora, 7-9
 directories, D-2
 planned interface uses, 1-10
new features, xix, xxiv
noninteractive mode. *See* response file mode.
NTS operating system authentication, 7-12

O

OCFS for Windows, 1-14, 3-7
olsnodes command, A-4
operating system
 preinstallation requirements for Oracle RAC, 1-11
Optimal Flexible Architecture (OFA), xxix, 1-9, 2-7, 3-8, C-2, C-6, D-1
ORA_DBA group, 1-9, A-3, B-5
Oracle ACFS, xxi, 1-12
Oracle Application Express
 requirements for installation, 2-2
 web browsers supported for, 1-10

Oracle ASM
 and candidate disks, 3-4
 and failure groups, 3-6
 and fast recovery area, 3-6
 change disk discovery path, 3-4
 Oracle Database releases compatible with this release, 1-19
 storage options for data files, 1-13
Oracle ASM Dynamic Volume Manager (Oracle ADVM), xxi
Oracle Cluster Registry port, F-2
Oracle Clusterware
 and multiple database releases, 1-5
Oracle Configuration Manager
 information needed to enable, 2-5
 installing, xxx
 overview, xxv, xxvi
 postinstallation, 4-7
Oracle Connection Manager
 ports, ranges, and protocols, F-2
Oracle Data Guard
 ports, ranges, and protocols, F-2
Oracle Database
 data file storage options, 1-13
 releases that can coexist using Oracle ASM, 1-19
 supported storage options for, 1-13
 upgrades of, 1-4, 2-3
Oracle Database Configuration Assistant (DBCA)
 components created by, 7-2
 configuring Database Control, 6-4
 control files, 7-4
 creating Oracle Real Application Clusters database after installation, 3-5
 during installation, 2-7, 3-1
 data files, 7-3
 deleting databases, 3-8
 deleting Oracle RAC databases, 3-8
 initialization parameter files, 7-5
 List of Cluster Databases page, 3-8
 Management Options page, 6-4, G-1
 no longer sets LOCAL_LISTENER and REMOTE_LISTENER, 7-10
 Operations page, 3-8
 Recovery Configuration page, 3-7
 redo log files, 7-4
 response file, B-5
 rollback segments, 7-5
 running in silent mode, B-7
 starting, 6-4
 Summary dialog, 3-8
 tablespaces, 7-3
 templates, 3-2
 using, 3-1
 Welcome page, 3-5
Oracle Database Vault
 and multiple Oracle homes, 1-15
 configuration, 1-15
 Oracle Database Vault Account Manager (DV_ACCTMGR), 2-4
 Oracle Database Vault Owner (DV_OWNER), 2-4

- postinstallation steps, 4-4
- Oracle Enterprise Manager
 - ASMSNMP, 1-16
 - Database Control, 4-6
 - changing ports, F-4
 - ports, ranges and protocols, F-3
 - logging in to Database Control, 4-6
 - using to convert single-instance databases to
 - Oracle Real Application Clusters, C-10
- Oracle Event Manager
 - ports, ranges, and protocols, F-3
- Oracle home
 - ASCII path for, 1-10
 - multiple, 7-2
- Oracle Managed Files (OMF), C-2
- Oracle Management Agent
 - ports
 - changing, F-4
 - ranges and protocols, F-3
- Oracle Net Configuration Assistant (NETCA), 2-7
 - response file, B-5
- Oracle Net Services listener
 - ports, ranges and protocols, F-3
- Oracle Net Services listeners
 - ports, ranges and protocols, F-3
- Oracle RAC One Node, 2-2
- Oracle RAC Technologies Certification Matrix, 1-2
- Oracle Real Application Clusters
 - databases, deleting, 3-8
 - operating system requirements for, 1-11
 - overview, 7-1
 - ports
 - ranges and protocols, F-4
 - raw device setup, 7-4
 - shared files, 1-16
 - storage requirements for, 1-11
- Oracle Real Application Testing, xxvi
- Oracle Universal Installer (OUI)
 - response files
 - list of, B-5
- Oracle Upgrade Companion, 1-4, 2-3
- Oracle XML DB
 - ports
 - changing, F-5
 - ranges and protocols, F-4
- ORACLE_BASE
 - default value, 1-9
 - setting the environment variable, 4-5
- ORACLE_HOME
 - default value, 1-9
 - setting the environment variable, 4-5
- ORACLE_SID
 - setting the environment variable, 4-5
- OracleMetalink. *See* My Oracle Support web site.
- OracleRemExecService Windows service, 4-10
- OUI
 - See* Oracle Universal Installer (OUI).

P

- parameter files (PFILE)
 - local, 5-2
 - search order, 5-3
- parameters
 - See also* initialization parameters
- passwords, 2-6
 - Oracle Database Vault Account Manager, 2-5
 - Oracle Database Vault Owner, 2-4
 - specifying for response files, B-3
 - See also* security
- patch bundles
 - contents, 4-2
 - description, 4-2
- Patch Set Updates, 4-2
- patch sets
 - description, 4-2
- patches
 - and My Oracle Support, 4-2
 - and upgrades, E-2
 - applying to older software releases, 4-10
 - CPU, 4-2
 - downloading patch updates, 4-2
 - DST, 4-2
 - installing, 4-2
 - operating system patches for Oracle RAC, 1-11
 - PSU, 4-2
- PFILE. *See* parameter files (PFILE).
- PMON process
 - and listener registration, 7-6
 - discovery routine, 7-6
- policy-managed databases
 - and SCAN, 6-2
- portlist.ini file, F-2
- ports
 - access URLs, F-2
 - Cluster Manager
 - ranges and protocols, F-2, F-4
 - configured for applications, F-2
 - Database Control, 6-9
 - default ranges, F-1
 - Java Message Service (JMS), 6-9
 - Oracle Cluster Registry, F-2
 - Oracle Connection Manager
 - ranges and protocols, F-2
 - Oracle Data Guard
 - ranges and protocols, F-2
 - Oracle Enterprise Manager Database Control
 - changing, F-4
 - ranges and protocols, F-3
 - Oracle Event Manager
 - ranges and protocols, F-3
 - Oracle Management Agent, 6-10
 - changing, F-4
 - ranges and protocols, F-3
 - Oracle Net Services listener
 - ranges and protocols, F-3
 - Oracle Real Application Clusters
 - ranges and protocols, F-4
 - Oracle XML DB

- changing, F-5
- ranges and protocols, F-4
- reconfiguring Database Control after changing the listener port, 6-10
- Remote Method Invocation (RMI), 6-9
- specifying Database Control ports, 6-9
- postinstallation
 - patch download and install, 4-2
 - product configuration, 4-4
 - user accounts setup, 4-5
- preconfigured database installation types, 2-3
- preinstallation
 - operating system requirements, 1-11
 - raw device creation, 7-4
- processes
 - stopping existing, 1-19
- PSU, 4-2

R

RAC. *See* Oracle Real Application Clusters.

raccimon, 7-14

raw devices

- identifying on Windows
 - db_name_idx1 symbolic link, 7-4
 - db_name_rbs1 symbolic link, 7-4
 - db_name_redo_thread_number symbolic link, 7-4
 - db_name_spfile1 symbolic link, 7-4
 - db_name_system1 symbolic link, 7-3
 - db_name_temp1 symbolic link, 7-3
 - db_name_users1 symbolic link, 7-3
 - undotbs1 and undotbs2 symbolic links, 7-4
- setting up, 7-4
- supported for upgrades only, 1-13

raw devices desupported, xxxi

RBS tablespace

- description, 7-3

rconfig, C-9

readme.txt file, F-2

recovery files

- See also* fast recovery area
- supported storage options, 1-14

redo log files, 1-16

- described, 7-4
- installed configuration, 7-4

Remote Method Invocation (RMI)

- port used with Database Control, 6-9

requirements

- for Java Runtime Environment on Windows, 1-11

response file mode

- about, B-3
- error handling, A-5
- installation
 - preparing, B-4
 - templates, B-4
- reasons for using, B-4
- See also* response files, silent mode

response files

- about, B-3

- creating with template, B-4
- dbca.rsp, B-5
- enterprise.rsp, B-5
- error handling, A-5
- general procedure, B-4
- netca.rsp, B-5
- Oracle Net Configuration Assistant (NETCA), B-9
- passing values at the command-line, B-3
- passwords, B-3
- security, B-3
- specifying with Oracle Universal Installer, B-6
 - See also* silent mode
- rollback segments
 - described, 7-5

S

SCAN addresses

- returned to the client, 7-10

SCAN. *See* single client access name (SCAN).

SCAN VIP, 7-13

scripts to create an Oracle Real Application Clusters database, B-1

secure data center

- downloading installer updates for, 1-3

security

- configuring after installation, 4-4
- for Oracle Database Vault Account Manager password, 2-5
- for Oracle Database Vault owner password, 2-4
- selecting passwords, 2-6
- using DBCA to enable or disable configuration, 4-4
- See also* passwords

server parameter files (SPFILE)

- about, 7-5
- and DBCA, 3-2
- and shared storage, 1-16
- configuring, 5-1
- default location, 5-2
- errors, 5-4
- managing on raw device, 5-2
- migrating to, 5-3

server pools

- configuration, 6-1
- SCAN required for policy-managed, 6-2
- XML conversion files for, C-13

service registration, 7-6

- about, 7-8
- configuring, 7-12

shared servers, 7-16

SID. *See* system identifiers (SID).

silent mode

- about, B-3
- error handling, A-5
- errors, A-5
- installation, B-7
- reasons for using, B-4
- See also* response files, response mode

- single client access name (SCAN)
 - description, xxiii
 - use required for clients of policy-managed databases, 6-2
- single-instance databases
 - converting, C-10
- SPFILE. *See* server parameter files (SPFILE).
- sqlnet.ora files, 7-12
- Standard Edition of Oracle Database
 - supported storage options, 1-14
- storage options
 - for Enterprise Edition installations, 1-13
 - for Oracle RAC, 1-11
 - for Standard Edition installations, 1-13
 - supported options, 1-13, 1-14
- Summary dialog, 3-8
- supported technology, 1-2
- symbolic links
 - db_name_indx1, 7-4
 - db_name_rbs1, 7-4
 - db_name_redo_thread_number, 7-4
 - db_name_spfile1, 7-4
 - db_name_system1, 7-3
 - db_name_temp1, 7-3
 - db_name_users1, 7-3
 - undotbs1 and undotbs2, 7-4
- SYS user
 - password requirements, 2-6
 - used by default, 4-7
- SYSAUX tablespace
 - description, 7-3
- SYSDBA privilege, 1-9, A-3
- SYSMAN user
 - password requirements, 2-6
- system identifier (SID)
 - selecting, 2-5
 - used in DBCA, 3-6
- SYSTEM tablespace
 - description, 7-3
- SYSTEM user
 - password requirements, 2-6

T

- tablespaces
 - and DBCA, 7-3
 - expanding for large sorts, 7-3
 - RBS, 7-3
 - SYSAUX, 7-3
 - SYSTEM, 7-3
 - TEMP, 7-3
 - undo tablespaces for automatic undo management, 7-3
 - USERS, 7-3
- TEMP tablespaces
 - described, 7-3
- templates
 - and DBCA, 3-2
- time zone version files, 1-16
- TIMESTAMP WITH TIME ZONE (TSTZ) data, 1-16

- TNS_ADMIN
 - and the listener.ora file, 7-7
 - changing and reconfiguring Database Control, 6-11
- tnsnames.ora files
 - and VIP addresses, 7-13
 - default configuration, 7-9
- Transaction Processing
 - database configuration type, 2-3
- troubleshooting
 - configuration assistant errors, A-4
 - disk space errors, 1-10
 - log files, 2-1, A-2, A-4
 - node is unreachable, A-3
 - nodes unavailable for selection, A-3
 - problems starting or configuring Database Control, 6-10
 - PROT-8
 - Failed to import data from specified file to the cluster registry, A-3
 - reviewing a log file, A-2, A-4
 - time stamp in the future, A-3
 - with EMCA, 6-10
- TSTZ. *See* time zone version files.
- Typical installation type, 2-2

U

- undo management, 7-5
- upgrades
 - and SCANS, 7-14
 - and using same Oracle home, 1-5
 - and using same Oracle home type, 2-10
 - of existing Oracle databases, 1-4, 2-3, 3-3
 - of listeners, 3-3
 - of Oracle Database using a patch set, E-2
 - of TIMESTAMP WITH TIME ZONE data, 1-16
 - Oracle Clusterware release requirement for, 1-4
 - preupgrade steps, 1-4
 - stopping processes, 1-19
 - using the same Oracle home, 2-8
- user authentication, 7-12
- users
 - domain, 1-9
 - local Administrator, 1-9
 - postinstallation user account setup, 4-5
- USERS tablespace
 - described, 7-3

V

- VIP addresses, 7-13

W

- web browsers
 - supported with Oracle Application Express, 1-10
 - supported with Oracle Enterprise Manager, 1-10

