

ORACLE 11G REAL
APPLICATION CLUSTERS
FOR ADMINISTRATORS

ORACLE 11G REAL APPLICATION CLUSTERS FOR ADMINISTRATORS

Student Workbook

ORACLE 11G RAC FOR ADMINISTRATORS

Published by ITCourseware, LLC., 7245 South Havana Street, Suite 100, Centennial, CO 80112

Contributing Authors: John Mullins and Rob Roselius

Technical Editor: Danielle Hopkins

Editor: Ginny Jaranowski

Special thanks to: Many instructors whose ideas and careful review have contributed to the quality of this workbook, including Julie Johnson and John McAlister, and the many students who have offered comments, suggestions, criticisms, and insights.

Copyright © 2012 by ITCourseware, LLC. All rights reserved. No part of this book may be reproduced or utilized in any form or by any means, electronic or mechanical, including photo-copying, recording, or by an information storage retrieval system, without permission in writing from the publisher. Inquiries should be addressed to ITCourseware, LLC., 7245 South Havana Street, Suite 100, Centennial, Colorado, 80112. (303) 302-5280.

All brand names, product names, trademarks, and registered trademarks are the property of their respective owners.

CONTENTS

Chapter 1 - Course Introduction	11
Course Objectives	12
Course Overview	14
Using the Workbook	15
Suggested References	16
Chapter 2 - Cluster Architecture	19
Cluster Architecture	20
Shared Storage	22
Nodes and Interconnects	24
Oracle Cluster Registry (OCR)	26
The RAC Voting Disk	28
Virtual IP Addresses	30
SCAN Addresses and GNS	32
Oracle Software	34
Labs	36
Chapter 3 - Oracle Clusterware	39
What is Oracle Clusterware?	40
Oracle Clusterware Files	42
Oracle Clusterware Processes	44
The Grid Home	46
Oracle Clusterware Preinstallation Steps	48
Node Preparation	50
Configuring SSH User Equivalency	52
cluvfy — The Cluster Verification Utility	54
Installing the Grid Infrastructure	56
Oracle Clusterware Postinstallation Steps	58
Labs	60
Lab Datasheet	62

Chapter 4 - Installing DB Software and Creating a RAC Database	65
Installation Overview	66
Configuring the OS Environment	68
Verifying System Readiness with the CVU	70
Installing the Database Software	72
Troubleshooting Installation Setup	74
Recommended Postinstallation Tasks	76
Database Pre-creation Tasks	78
Creating a Cluster Database Using DBCA	80
Initialization Parameters	82
ORACLE_SID in the RAC Environment	84
Labs	86
Chapter 5 - RAC Database Architecture	89
Overview of RAC	90
Oracle Single-Instance Architecture	92
RAC Architecture	94
RAC Instances and Parameter Files	96
RAC Database Components	98
RAC Instance Background Processes	100
Global Resource Directory	102
Overview of Cache Fusion	104
Cache Fusion Components — GES and GCS	106
Cache Fusion Components — Resource Master and GRD	108
Application Connection to RAC	110
Labs	112
Chapter 6 - RAC Instance Management	115
Overview of RAC Instance Management	116
Starting and Stopping a RAC Database	118
Starting and Stopping a RAC Instance	120
RAC Database Identical Parameters	122
RAC Database Unique Parameters	124
Changing Parameter Values	126
Administering Undo Tablespaces in RAC	128
Administering Redo Logs in RAC	130
Labs	132

Chapter 7 - Managing Oracle Clusterware	135
About Oracle Clusterware	136
Adding and Removing Voting Disks	138
Backing Up and Recovering Voting Disks	140
The OCR	142
Changing the OCR Configuration	144
Adding and Removing an OCR Location	146
Backing Up and Recovering the OCR	148
Restoring from Automatic OCR Backups	150
Moving or Replacing the OCR	152
Repairing the OCR Configuration	154
Troubleshooting the OCR	156
Labs	158
Chapter 8 - RAC Utilities	161
The olsnodes Utility	162
The ocrcheck Utility	164
The ocrdump Utility	166
Querying Clusterware Resource Status	168
The crsctl Utility	170
The Server Control (SRVCTL) Utility	172
SRVCTLADD	174
SRVCTL CONFIG	176
SRVCTL ENABLE and DISABLE	178
SRVCTL GETENV	180
SRVCTL MODIFY	182
SRVCTL RELOCATE	184
SRVCTL STATUS	186
SRVCTL REMOVE	188
SRVCTL START	190
SRVCTL STOP	192
Labs	194
Chapter 9 - Services	197
Overview of Services	198
Types of Services	200
Creating Services with srvctl	202
Preferred and Available Instances for Services	204

Using Services	206
Managing Services	208
Service Views	210
Tracing with Services	212
Labs	214
Chapter 10 - Failover	217
Transparent Application Failover (TAF)	218
Client-Side vs. Server-Side TAF	220
Configuring TAF on the Client	222
Using OEM to Configure TAF	224
Using svrctl to Configure TAF	226
The DBMS_SERVICE Package	228
Connecting to the Database with TAF	230
Monitoring TAF Connections	232
Labs	234
Chapter 11 - RAC Backup and Recovery	237
Overview of RAC Backup and Recovery	238
Log Archiving in RAC	240
Undo Tablespaces in RAC	242
Using Flashback Features in RAC	244
Deploying a Fast Recovery Area in RAC	246
Performing RMAN Backups of a RAC Database	248
Performing Non-RMAN Backups	250
Preparing to Restore and Recover with RMAN	252
Recover the RAC Database with RMAN	254
Recovering Without RMAN	256
Labs	258
Chapter 12 - Cluster Management	261
Overview of Cluster Management Tasks	262
Extending the Clusterware Home Directory	264
Extending the Database Software Home Directory	266
Creating the New Instance	268
Verifying the New Instance	270
Removing a Node	272
Labs	274

Chapter 13 - Automatic Storage Management	277
ASM Overview	278
ASMLib	280
Installing ASM	282
Creating an ASM Instance	284
ASM Initialization Parameters	286
Accessing an ASM Instance	288
ASM Startup and Shutdown	290
Creating a Disk Group	292
Dynamic Performance Views for ASM	294
Changing a Disk Group	296
Oracle-Managed Files	298
Creating and Managing a Database	300
Creating and Managing a Tablespace	302
Creating and Managing Redo Logs	304
Creating and Copying Control Files	306
Backing up and Restoring Control Files	308
Creating and Managing Archive Logs	310
ASMCMD	312
Labs	314
 Chapter 14 - ASM Backup, Recovery, and Migration	 317
Overview of ASM Backup and Recovery	318
Overview of Recovery Manager (RMAN) and ASM	320
Migrating a Database to ASM	322
Migrating the Control File	324
Migrating the Database Files	326
Migrating the Online Redo Logs	328
Migrating the SPFILE	330
Labs	332
 Chapter 15 - Clusterware and RAC Troubleshooting	 335
The Oracle Clusterware Alert Log	336
Clusterware Component Log Files	338
Automatic Diagnostic Repository and ADRCI	340
Using crsctl To Diagnose Cluster Issues	342
Using diagcollection.pl	344
Checking Interconnect Settings	346
cluvfy — Verifying Clusterware Component Integrity	348

cluvfy — Verifying Cluster Registry Integrity	350
cluvfy — Verifying Cluster Integrity	352
ocrcheck — Verifying the Oracle Cluster Repository	354
RAC Database Alerts	356
Cluster Health Monitor (CHM)	358
The racdiag.sql Script	360
The oradebug Utility	362
Labs	364
Chapter 16 - RAC Tuning	367
RAC Tuning Methodology	368
Using Performance Views in RAC	370
Monitoring Cache Fusion	372
Global Cache Latencies	374
Monitoring Cache Transfers	376
OEM — Cluster Database Performance	378
OEM — RAC-Related Reports	380
Using AWR in the RAC Environment	382
Generating AWR Reports	384
Analyzing AWR Reports	386
Using ADDM in the RAC Environment	388
Analyzing ADDM Reports	390
RAC Tuning Tips and Best Practices	392
Labs	394
Appendix - Preparing a RAC Node	397
Checking the Hardware Requirements	398
Identifying Network Requirements	400
Configuring Operating System Users and Groups	404
Generating RSA and DSA Keys	406
Adding the Keys to an Authorized Key File	408
Configuring SSH User Equivalency	410
Configuring the Operating System Environment	412
Configuring the Network	414
Verifying the Network Configuration	416
Preparing the Operating System and Software	418
Configuring Installation Directories and Shared Storage	420
Choosing Directories	422
Labs	424

Solutions 427

Index 443

CHAPTER 1 - COURSE INTRODUCTION

COURSE OBJECTIVES

- * Explain the architecture of a RAC system.
- * Install and configure Oracle Grid Infrastructure and Oracle RAC Database Software in a cluster environment.
- * Add, remove, backup, and recover voting disks and the OCR.
- * Correctly set RAC initialization parameters.
- * Start and stop individual RAC instances and an entire RAC database.
- * Manage redo log files and undo tablespaces for a RAC database.
- * Configure and Manage Services in a RAC Environment.
- * Configure Transparent Application Failover for client load balancing, connect-time, and application failover.
- * Configure a RAC database for archiving redo logs and flashback database mode.
- * Backup and recover a RAC database and its components.
- * Prepare, configure, remove, and add nodes to an existing cluster.
- * Set up and manage an ASM instance.
- * Migrate a RAC database to ASM.
- * Locate and interpret the contents of critical RAC alert and RAC component log files.
- * Monitor and diagnose performance issues with built-in utilities, views, reports, and the OEM.

COURSE OVERVIEW

- * **Audience:** Database administrators new to the RAC environment.
- * **Prerequisites:** *Oracle 11g Database Administration* and at least 6 months of administration experience recommended.
- * **Classroom Environment:**
 - A workstation per student.

USING THE WORKBOOK

This workbook design is based on a page-pair, consisting of a Topic page and a Support page. When you lay the workbook open flat, the Topic page is on the left and the Support page is on the right. The Topic page contains the points to be discussed in class. The Support page has code examples, diagrams, screen shots and additional information. **Hands On** sections provide opportunities for practical application of key concepts. **Try It** and **Investigate** sections help direct individual discovery.

In addition, there is an index for quick lookup. Printed lab solutions are in the back of the book as well as online if you need a little help.

The Topic page provides the main topics for classroom discussion.

The Support page has additional information, examples, and suggestions.

JAVA SERVLETS

THE SERVLET LIFE CYCLE

- * The servlet container controls the life cycle of the servlet.
 - > When the first request is received, the container loads the servlet class
 - The container uses a separate thread to call the `init()` method. After the `init()` method returns, the container calls the `destroy()` method.
 - As with Java's `finalize()` method, don't count on this being called.
- * Override one of the `init()` methods for one-time initializations, instead of using a constructor.
 - > The simplest form takes no parameters.


```
public void init() {...}
```
 - > If you need to know container-specific configuration information, use the other version.


```
public void init(ServletConfig config) {...}
```

 - Whenever you use the `ServletConfig` approach, always call the superclass method, which performs additional initializations.


```
super.init(config);
```

Page 16 Rev 2.0.0 © 2002 ITCourseware, LLC

Topics are organized into first (*), second (>), and third (▪) level points.

CHAPTER 2 SERVLET BASICS

Hands On:

Add an `init()` method to your `Today` servlet that initializes along with the current date:

Today.java

```
...
public class Today extends GenericServlet {
    private Date bornOn;
    public void service(ServletRequest request,
        ServletResponse response) throws ServletException, IOException
    {
        ...
        response.getWriter().println("This servlet was born on " + bornOn.toString());
        response.getWriter().println("It is now " + today.toString());
    }
}
```

The `init()` method is called when the servlet is loaded into the container.

© 2002 ITCourseware, LLC Page 17

Code examples are in a fixed font and shaded. The online file name is listed above the shaded area.

Callout boxes point out important parts of the example code.

Screen shots show examples of what you should see in class.

Pages are numbered sequentially throughout the book, making lookup easy.

SUGGESTED REFERENCES

Bach, Martin, Dyke, Julian and Shaw, Steve. 2010. *Pro Oracle Database 11g RAC on Linux: Installation, Administration, and Performance*. Apress, Berkeley, CA. ISBN 1430229586.

Gopalakrishnan, K. 2011. *Oracle Database 11g Real Application Clusters Handbook, Second Edition*. McGraw Hill Osborne Media, Emeryville, CA. ISBN 0071752625.

Jesse, Scott, Burton, Bill, and Vongray, Brian. 2011. *Oracle Database 11g Release 2 High Availability: Maximize Your Availability with Grid Infrastructure, RAC and Data Guard*. McGraw-Hill Osborne Media, Emeryville, CA. ISBN 0071752080

Vallath, Murali. 2006. *Oracle 10g RAC Grid, Services & Clustering*. Elsevier Digital Press, Burlington, MA. ISBN 1555583210.

The following books are available online at <http://www.oracle.com/pls/db112/homepage> :

- *Oracle Database Administrator's Guide 11g Release 2 (11.2)*
- *Oracle Database 2 Day + Real Application Clusters Guide 11g Release 2 (11.2)*
- *Oracle Clusterware Administration and Deployment Guide 11g Release 2 (11.2)*
- *Oracle Real Application Clusters Administration and Deployment Guide 11g Release 2 (11.2)*
- *Oracle Database Storage Administrator's Guide 11g Release 2 (11.2)*
- *Grid Infrastructure Installation Guide for Linux*
- *Real Application Clusters Installation Guide for Linux and Unix*

CHAPTER 7 - MANAGING ORACLE CLUSTERWARE

OBJECTIVES

- * Add, remove, back up, and recover voting disks.
- * Add, remove, back up, and recover the OCR.

ABOUT ORACLE CLUSTERWARE

- * Oracle Clusterware is software that is required for an Oracle RAC solution.
- * Oracle Clusterware is needed to bind multiple servers so that they operate as a single system.
- * Oracle Clusterware Components
 - The Oracle Clusterware Sack - background processes running on each node.
 - Voting Disk
 - Records node membership information
 - Oracle Cluster Registry (OCR)
 - Records cluster configuration information
 - Utilities for managing the clusterware components.
- * Once installed and configured, the clusterware integrates with the host operating system, starting automatically at system boot.
 - Many clusterware processes run with system (*root*) privileges.
 - Some clusterware management commands require *root* privileges.

Some clusterware commands don't need special privileges and can be run by the *grid* user. Commands for managing clusterware components (startup and shutdown of the clusterware, manipulation of clusterware files and backups, etc.) require *root* privilege on Unix and Linux systems. If you are not allowed *root* access to the systems you support, you will have to get System Administration support each time you need to run these commands.

Alternatively, the system administrator can grant the *grid* user privileged access to clusterware commands via **sudo**:

```
/etc/sudoers:
```

```
grid    ALL = NOPASSWD: /u01/app/11.2.0/grid/bin/
```

```
[grid@node2]$ sudo ocrconfig -manualbackup
```

ADDING AND REMOVING VOTING DISKS

- * It is important to have mirrored copies of the voting disk files to help eliminate down time due to a clusterware issue.
- * A new voting disk file can be added to your RAC configuration using the **crsctl** utility.

- Add a voting disk to your RAC configuration:

```
# crsctl add css votedisk /u03/storage/vdsk2
Now formatting voting disk: /u03/storage/vdsk2
```

- Identify your new voting disk file:

```
# crsctl query css votedisk
## STATE      File Universal Id      File Name Disk group
--  -
  1. ONLINE    813322a180244f1bbf13e990f601bcc5 (/u02/storage/vdsk1) []
  2. ONLINE    95a2c6f51c474fb3bfed4e9936de29e6 (/u03/storage/vdsk2) []
Located 2 voting disk(s).
```

- * The **crsctl** utility can also be used to remove a voting disk file.
- Remove a voting disk from your RAC configuration:

```
# crsctl delete css votedisk /u03/storage/vdsk2
successful deletion of votedisk /u03/storage/vdsk2.
```

- * If you chose to store your voting disk under ASM, then redundancy is managed automatically.

With 11g Release 2, the voting disk can be stored in an ASM diskgroup. To migrate the voting disk into ASM diskgroup named **MYDATA**:

```
# crsctl replace votedisk +MYDATA
```

To migrate a voting disk from ASM to a shared filesystem path, say */u04/storage*:

```
# crsctl replace votedisk /u04/storage/vdisk1
```

BACKING UP AND RECOVERING VOTING DISKS

- * The clusterware automatically backs up the voting disk contents to the OCR.
- * Use the **crsctl** utility to identify your active voting disk files.

```
[root@c17node1 c17db]# crsctl query css votedisk
## STATE      File Universal Id      File Name Disk group
--  -
  1. ONLINE    813322a180244f1bbf13e990f601bcc5 (/u02/storage/vdisk1) []
  2. ONLINE    95a2c6f51c474fb3bfed4e9936de29e6 (/u03/storage/vdisk2) []
  3. ONLINE    4033f4d78b004fe8bf8258341bec1d3b (/u04/storage/vdisk3) []
Located 3 voting disk(s).
```

- * The clusterware backs up the voting disk information to the OCR any time a configuration change is made.
 - This information is used whenever you add a new voting disk.
- * In the event of *all* voting disks being lost or corrupted, you can restore the voting disks from the OCR (after recovering the OCR, if necessary).

1. List the votedisks, with their File Universal IDs (FUIDs).

```
# crsctl query css votedisk
```

2. Stop the clusterware, if it's running.

```
# crsctl stop crs
```

3. On one node, as **root**, start the clusterware in exclusive mode.

```
# crsctl start crs -excl
```

4. Delete the previous votedisk records using the FUIDs.

```
# crsctl delete css votedisk 813322a180244f1bbf13e990f601bcc5
```

5. Create new vote disks; the clusterware will use the OCR backup.

```
# crsctl add css votedisk /u02/storage/vdisk1
```

6. Stop and restart the clusterware.

```
# crsctl stop crs
```

```
# crsctl start crs
```

Prior to Oracle 11g Release 2, voting disk backups were made manually with the **dd** command (if the voting disk was on a raw device) or **cp** (if the voting disk was an ordinary file). Restoring it involved simply copying it back from the backup copy after stopping the clusterware, then restarting the clusterware.

THE OCR

- * All clusterware-managed resources are recorded in the *Oracle Cluster Registry* (OCR), a file on shared storage accessed and maintained by clusterware processes and utilities on all nodes.
 - Information in the OCR includes the list of cluster nodes and various clusterware components and node applications, RAC databases and services, instance-to-node mapping, VIPs, listeners, and ASM instances and diskgroups.
- * You can find the location of the OCR file in the *ocr.loc* file, or by running the **ocrcheck** utility.
 - On UNIX and Linux systems the *ocr.loc* file is in */etc/oracle*.

```
# more /etc/oracle/ocr.loc
ocrconfig_loc=/u02/storage/ocr1
local_only=FALSE
```

- * The clusterware on each node also maintains its own local registry in the Grid_Home on its local hard drive: the *Oracle Local Registry* (OLR).
 - Information in the OLR includes internal configuration and status information for the clusterware processes (including ASM), lists of OCR backups, and version information.

```
# more /etc/oracle/olr.loc
olrconfig_loc=/u01/app/11.2.0/grid/cdata/c20node4.olr
crs_home=/u01/app/11.2.0/grid
```

The **ocrcheck** utility can also be used to identify the locations of your OCR files.

```
# /u01/app/11.2.0/grid/bin/ocrcheck
Status of Oracle Cluster Registry is as follows :
Version                :                3
Total space (kbytes)   :           262120
Used space (kbytes)    :             3236
Available space (kbytes) :         258884
ID                     : 1587080569
Device/File Name       : /u02/storage/ocr1
                        Device/File integrity check succeeded
                        Device/File not configured
                        Device/File not configured
                        Device/File not configured
                        Device/File not configured
                        Device/File not configured

Cluster registry integrity check succeeded

Logical corruption check succeeded

# /u01/app/11.2.0/grid/bin/ocrcheck -local
Status of Oracle Local Registry is as follows :
Version                :                3
Total space (kbytes)   :           262120
Used space (kbytes)    :             2212
Available space (kbytes) :         259908
ID                     : 1054057231
Device/File Name       : /u01/app/11.2.0/grid/cdata/c20node4.olr
                        Device/File integrity check succeeded

Local registry integrity check succeeded

Logical corruption check succeeded
```

CHANGING THE OCR CONFIGURATION

- * The **ocrconfig** utility is used to change the configuration of the OCR.
- * **ocrconfig** tasks to change the OCR's configuration:
 - Adding, removing, moving, repairing an OCR location.

ocrconfig Configuration Tasks for the OCR:

OCR Task	ocrconfig Option
Add an OCR location	-add
Remove an OCR location	-delete
Change the location of an OCR	-replace
Repair an OCR configuration	-repair
Take a manual OCR backup	-manualbackup
Restore the OCR from backup	-restore
List available OCR backups	-showbackup
Export OCR contents to a file	-export
Import OCR contents from file	-import

Note:

You can also use the **-help** option with the **ocrconfig** utility to get further assistance on using its options.

You must be logged in as the **root** user to run the **ocrconfig** utility.

ADDING AND REMOVING AN OCR LOCATION

* Protecting the OCR contents is very important to the functionality of the Oracle Clusterware.

➤ Mirroring is used to protect those contents.

* An OCR location can be added or removed using the **ocrconfig** utility.

➤ When adding a new OCR file location, you must first create an empty file and assign it appropriate ownership and permissions:

```
# touch /u03/storage/ocr2
# chown root:oinstall /u03/storage/ocr2
# chmod 640 /u03/storage/ocr2
```

➤ Run the **ocrconfig** utility to add an OCR mirror location:

```
# ocrconfig -add /u03/storage/ocr2
```

➤ Run the following command on any node in the cluster to remove an OCR location:

```
# ocrconfig -delete /u02/storage/ocr1
```

▪ You cannot remove the last remaining OCR location.

See the *Oracle Clusterware Administration and Deployment Guide* for a complete reference of the **ocrconfig** utility.

BACKING UP AND RECOVERING THE OCR

- * Oracle Clusterware automatically creates OCR backups every four hours.
 - The latest three backup copies of the OCR are retained.
 - Four hours old
 - One day old
 - One week old
 - You cannot customize the backup frequencies or the number of files that Oracle Clusterware retains.
 - Use **ocrconfig -showbackup** to determine the backup location.
 - The automatically generated backups should be copied to tape.
- * Use **ocrconfig** to find the most recent backup of the OCR.
 - The **ocrdump** utility can also be used to identify the latest backups of the OCR.
- * You can create manual backups using **ocrconfig -manualbackup**.
 - You might want to perform manual OCR backups before and after making significant changes to cluster and resource configuration.
- * In addition to the automatic and manual backups to disk, you can use **ocrconfig** to export the contents of the OCR to give you more flexibility in the event that a restore/recovery is needed.
 - An OCR export should only be performed when the clusterware is not running.

Use the **ocrconfig** utility to show the most recent backups:

```
# ocrconfig -showbackup
c20node1    2010/11/09 10:45:17    /u01/app/11.2.0/grid/cdata/c20/backup00.ocr
c20node1    2010/11/09 06:45:16    /u01/app/11.2.0/grid/cdata/c20/backup01.ocr
c20node1    2010/11/09 02:45:14    /u01/app/11.2.0/grid/cdata/c20/backup02.ocr
c20node1    2010/11/08 06:45:05    /u01/app/11.2.0/grid/cdata/c20/day.ocr
c20node1    2010/11/01 03:43:37    /u01/app/11.2.0/grid/cdata/c20/week.ocr
c20node1    2010/10/26 15:01:43    /u01/app/11.2.0/grid/cdata/c20/backup_20101026_150143.ocr
```

Use **ocrdump** to dump an outline of contents, including backup information, to a text file:

```
# ocrdump ocrtextdump
```

In the above example, *ocrtextdump* is the resulting file name.

Export the contents of the OCR to a file, with a resulting file name of *ocrexpert*:

```
# ocrconfig -export ocrexpert
```

Note:

If needed, the OCR export file can then be used with the **ocrconfig -import** command to correct an OCR problem. However, the OCR export does not guarantee an internally consistent copy of the OCR if it's taken while the clusterware is running; Oracle recommends shutting down the clusterware to obtain an internally consistent export.

Export files and backup files are in different, incompatible formats and can't be used interchangeably.

RESTORING FROM AUTOMATIC OCR BACKUPS

- * Before recovering the OCR, make sure it is the root of the problem.
 - Use the **cluvfy** or **ocrcheck** utility to confirm the integrity of the OCR.
- * Steps to recover the OCR using an automatic OCR backup:
 1. Determine which backup from which you'll restore the OCR; make sure you're on the node where the backup file is located.

```
ocrconfig -showbackup
```
 2. Stop the Oracle Clusterware using the **crsctl** utility on all nodes.

```
crsctl stop crs
```
 3. Use the **ocrconfig** utility to restore the OCR from a backup.

```
ocrconfig -restore ocrbackupfile
```
 4. Start the Oracle Clusterware using the **crsctl** utility on all nodes.

```
crsctl start crs
```
 5. Verify the integrity of the new OCR with the **cluvfy** utility or **ocrcheck** utility.

MOVING OR REPLACING THE OCR

* To change the location of your OCR, you can replace it with the **-replace** option of the **ocrconfig** utility.

* The following are steps to change (replace) the location of an OCR file.

1. Verify that you have a good, online OCR available other than the one being moved (replaced), using the **ocrcheck** utility.
2. Make sure the Oracle Clusterware is running with **crsctl**.

```
# crsctl check crs
CRS-4638: Oracle High Availability Services is online
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
```

3. Create an empty file, with correct ownership and permissions, in the new location.

```
# touch /u04/storage/ocr1
# chown root:oinstall /u04/storage/ocr1
# chmod 640 /u04/storage/ocr1
```

4. Use **ocrconfig -replace** to perform the move (replace).

```
# ocrconfig -replace /u02/storage/ocr1 \
            -replacement /u04/storage/ocr1
```

Removing or replacing an OCR location does not remove the old OCR file from the disk.

REPAIRING THE OCR CONFIGURATION

- * If a node is down in your cluster and changes are made on other nodes that affect the OCR (that is, addition, deletion, or move of an OCR location), you will need to repair the OCR on the node that was down during the changes.

- * To repair an OCR, use the **ocrconfig -repair** option:
 1. Shut down the Oracle Clusterware on the affected node if it is not already down.
 2. Repair the OCR on the affected node.

```
# ocrconfig -repair -add /u04/storage/ocr3
```

3. Start the Oracle Clusterware on the affected node.

```
# crsctl start crs
Attempting to start CRS stack
The CRS stack will be started shortly
```


TROUBLESHOOTING THE OCR

- * If you experience error messages online or in various log files regarding the OCR, use the **ocrcheck** utility to verify the OCR's integrity.

```
# ocrcheck -help
Name:
  ocrcheck - Displays health of Oracle Cluster Registry.

Synopsis:
  ocrcheck

Description:
  prompt> ocrcheck
  Displays current usage, location and health of the
  cluster registry
```

- * You can run **cluvfy** in "component" mode to check the OCR file locations and their permissions.

```
$ cluvfy comp ocr
```

- You cannot run **cluvfy** as a privileged user (**root**).
- * To check the OLR (Oracle Local Registry), run **ocrcheck -local**.
- You can also run **cluvfy comp olr**.

Verify the integrity of the OCR:

```
# ocrcheck
Status of Oracle Cluster Registry is as follows :
  Version                :                3
  Total space (kbytes)   :            262120
  Used space (kbytes)    :                3236
  Available space (kbytes) :            258884
  ID                     : 1587080569
  Device/File Name       : /u02/storage/ocr1
                        Device/File integrity check succeeded
                        Device/File not configured
                        Device/File not configured
                        Device/File not configured
                        Device/File not configured

  Cluster registry integrity check succeeded

  Logical corruption check succeeded

# su - grid
$ cluvfy comp ocr -n all -verbose
```

ocrcheck reports the status of the primary OCR file and up to four mirrors, if any are configured. At least one OCR file must show a successful integrity check.

When run as a privileged user (**root**, on UNIX and Linux systems), **ocrcheck** performs an additional "logical corruption check". When run as a non-privileged user (such as **grid** or **oracle**), this check is skipped.

LABS

- ❶ If necessary, create storage locations on shared storage for new voting disk and OCR mirror locations. Make sure ownership and permissions are set appropriately.
- ❷ Locate your existing voting disk files. Using storage provided, add a mirror copy of the voting disk to your RAC configuration.
- ❸ Locate your existing OCR file(s). Using storage provided, add an OCR location on shared storage for your cluster.
- ❹ Using storage provided, move the original OCR to a new location.
- ❺ List your current OCR backups. Create a manual OCR backup. Determine the location of the backup file (including which node it was created on.) Go to the backup file location and examine it with "ls-l".

CHAPTER 12 - CLUSTER MANAGEMENT

OBJECTIVES

- * Add a node to an existing cluster.
- * Remove a node from an existing cluster.

OVERVIEW OF CLUSTER MANAGEMENT TASKS

- * Once your cluster is in place, there are some typical cluster management tasks that may need to be performed, such as:
 - Adding a new node to the cluster.
 - The new node should have the same version of the operating system as the existing nodes in your cluster.
 - Removing a node from the cluster.
- * You may need to add a new node to your cluster due to the the need for additional resources.
 - Additional resources (memory, cpu) are not always available on an existing server.
- * The need to remove a node from your cluster may arise due to:
 - Node failure.
 - The node is underutilized and needed elsewhere.
- * The operating system and hardware architecture of the new node must match the existing nodes.
 - The new nodes hardware and operating system must have been fully prepped for Oracle Clusterware installation.

General steps for adding a new node to your cluster:

- Check for Dependencies and Prerequisites.
 - Operating system versions, packages, user accounts, groups.
- Configure Network Components.
 - Add network information to the */etc/hosts* file:
 - Public node names and ip addresses.
 - Private node names and ip addresses.
 - VIP node names and ip addresses.
 - SSH user equivalence.
- Install Oracle Clusterware.
 - Run the Cluster Verification Utility (CVU) before installing.
- Configure Oracle Clusterware.
 - Run the *addNode.sh* script provided by Oracle.
 - The script is located in the *Grid_home/oui/bin* directory.
- Install Oracle Software.
- Add New Instance(s).
 - Execute the DBCA utility.
- Do Housekeeping Chores.
 - Set up environment variables.
 - Configure Oracle Net.

General steps to remove a node from your cluster:

- Remove the instance.
 - Run the DBCA or **SRVCTL** utility.
- Remove the node from the cluster.
- Reconfigure the OS and remaining hardware.
 - If necessary, remove all Oracle installation files from the removed node.

EXTENDING THE CLUSTERWARE HOME DIRECTORY

- * Once your new node has been configured for Oracle Clusterware, use the OUI to add an Oracle Grid Infrastructure home to the new node.
- * To extend your Grid Home directory you must have a successfully installed Oracle Clusterware version on another node.

To extend the Oracle Clusterware installation to include your new node:

As **grid**:

1. Verify the new node has been fully prepped to join the cluster.

```
cluvfy stage -pre nodeadd -n node3 -verbose
```

2. Go to *Grid_home/oui/bin* and run the *addNode.sh* script.

```
cd /u01/app/11.2.0/grid/oui/bin
./addNode.sh -silent "CLUSTER_NEW_NODES={node3}" \
              "CLUSTER_NEW_VIRTUAL_HOSTNAMES={node3-vip}"
```

3. Run the *orainstRoot.sh* script on the new node.
4. Run the *Grid_home/root.sh* script on the new node. This starts the Clusterware.

Verify the status of the Clusterware extension on your new node by running *cluvfy* on an existing node.

```
cluvfy stage -post nodeadd -n node3 -verbose
```

EXTENDING THE DATABASE SOFTWARE HOME DIRECTORY

- * After extending the Oracle Clusterware and ASM homes to your new node, you must extend the Oracle Database home (RDBMS) to your new node.
- * The steps to add an Oracle RDBMS home to the new node are very similar to the steps you just completed for extending ASM to the new node.

To extend the Oracle RDBMS installation to include the new node:

Note:

First make sure that you have successfully installed the Oracle RDBMS software on at least one node in your cluster environment. As **oracle**:

1. Go to the `$ORACLE_HOME/oui/bin` directory on **node1** and run the `addNode.sh` script.

```
./addNode.sh -silent "CLUSTER_NEW_NODES={node3}"
```

2. Run the `root.sh` script on the new node from the `$ORACLE_HOME` directory on that node.

CREATING THE NEW INSTANCE

- * Once your new node is in place and you have installed the Oracle software and configured the listener, you are ready to create a new instance for your cluster.
- * To create a new instance on your new node use the DBCA utility.
- * The DBCA utility will update the OCR with the new instance's information.

Upon starting up the DBCA to create a new instance, it will prompt you for whether you are using a RAC database or a single instance database.

- Choose RAC database.
- From here, you will have eight steps to complete.
 1. Choose Instance Management.
 2. Choose Add an Instance.
 3. Select the cluster database to which the instance will be added.
 4. Review list of existing RAC instance names.
 5. Enter the name of your new instance and which node to create it on.
 6. For each database service defined, specify its preferred and available instance. You can also specify a service to not be used for a particular instance.
 7. Enter the storage information for your new instance.
 - Redo log storage specifics.
 - Undo Tablespace.
 8. Review the summary of your input and then the instance will be created.

Note:

If you had been using ASM in the existing cluster, the DBCA will prompt you for whether you wish to create an ASM instance on your new node also.

VERIFYING THE NEW INSTANCE

- * The DBCA should automatically add your new instance to the RAC cluster, but it is always a good idea to check for yourself.
- * To verify that the RAC cluster recognizes your new instance:
 - Use the **SRVCTL** utility.
 - Use the **ocrdump** utility.

The **SRVCTL** utility will report the status of all recognized instances in your cluster.

```
$ srvctl status database -d xyz
```

```
Instance xyz1 is running on node node1  
Instance xyz2 is running on node node2  
Instance xyz2 is running on node node3
```

REMOVING A NODE

- * A node may need to be removed from your cluster for various reasons.
 - There is a problem with the node.
 - The node is not needed due to under-utilization.
- * To remove a node, several components need to be taken care of:
 - Relocate or remove database services from instances on the node.
 - Remove any instances on the node from their database configurations.
 - Deconfigure the node from the clusterware.
- * To return the node platform to its original state:
 - Deconfigure the clusterware from the node's operating system.
 - De-install the Oracle Database software (ORACLE_HOME).
 - De-install the Grid Infrastructure software (Grid_HOME).

Steps to Remove a Node:

Steps for removing a node differ depending on the exact version of Oracle software, whether your database is administrator- or policy-managed, and other factors. See the *Real Application Clusters Administration and Deployment Guide* and the *Clusterware Administration and Deployment Guide*. Carefully study the *Adding and Deleting Nodes* chapters of both guides before beginning node removal. These example steps are for an administrator-managed, 11g Release 2 database. The node being removed is *node3*.

1. As **oracle**, on a node you are NOT deleting. Stop and remove the database instance. Before doing this, if necessary relocate any services from the instance to remaining instances.

```
srvctl stop instance -d xyz -i xyz3
dbca          #Use "Instance Management" to remove the instance
```

2. As **oracle**, on the node you ARE deleting. Update the local Oracle inventory to include only the name of the node to be deleted. Make sure you include ONLY the node you are deleting, and be careful the syntax of the CLUSTER_NODES parameter is correct. Incorrect syntax could result in software being removed on other nodes!

```
$ORACLE_HOME/oui/bin/runInstaller -updateNodeList \
ORACLE_HOME=$ORACLE_HOME "CLUSTER_NODES={node3}" -local
```

3. As **oracle**, on the node you ARE deleting. Delete the local Oracle Software. This removes the ORACLE_HOME from the node(s) listed in the local inventory; if the previous step was performed incorrectly, that could be *all* nodes!

```
$ORACLE_HOME/deinstall/deinstall -local
```

4. As **oracle**, on a node you are NOT deleting. Update the remaining nodes' Oracle inventories to include only the names of the remaining nodes.

```
$ORACLE_HOME/oui/bin/runInstaller -updateNodeList \
ORACLE_HOME=$ORACLE_HOME "CLUSTER_NODES={node1,node2}" "
```

5. As **root**, on the node you ARE deleting. Disable the Clusterware for the node being removed.

```
$GRID_HOME/crs/install/rootcrs.pl -deconfig -force
```

6. As **root**, on a node you are NOT deleting. Delete the node from the cluster configuration.

```
$GRID_HOME/bin/crsctl delete node -n node3
```

7. As **grid**, on the node you ARE deleting. Delete the Clusterware Software.

```
$GRID_HOME/oui/bin/runInstaller -updateNodeList \
ORACLE_HOME=$GRID_HOME "CLUSTER_NODES={node3}" CRS=TRUE \
-silent -local
$GRID_HOME/deinstall/deinstall -local
```

8. As **grid**, on a node you are NOT deleting. Update the remaining nodes' Oracle inventories to include only the names of the remaining nodes.

```
$GRID_HOME/oui/bin/runInstaller -updateNodeList CRS=TRUE \
ORACLE_HOME=$GRID_HOME "CLUSTER_NODES={node1,node2}" "
```

LABS

- ❶ Extend Oracle Clusterware to your newly prepared node.
- ❷ Extend Oracle RAC Database to your new node.
- ❸ Establish and verify an instance on the new node.