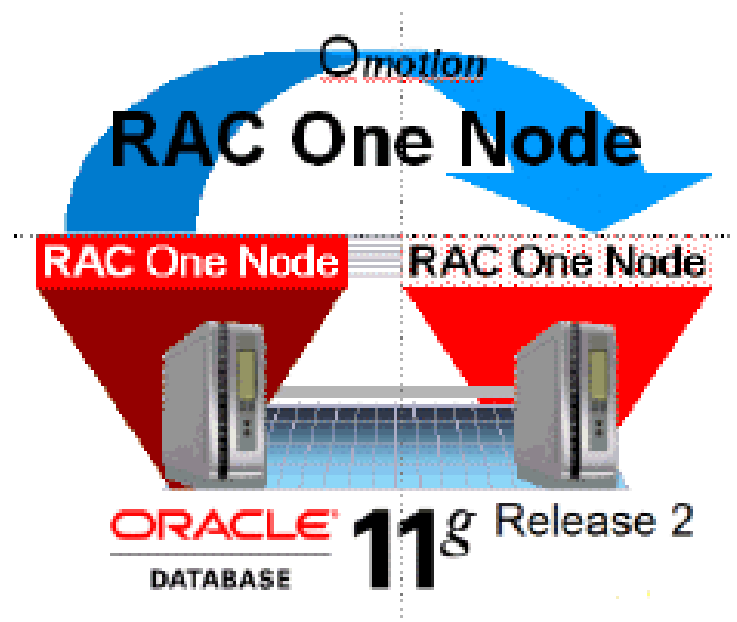


An Oracle White Paper
November 2009

Oracle RAC One Node 11g Release 2 User Guide



| | |
|---|----|
| Introduction | 1 |
| Software Installation | 3 |
| How to Configure an Oracle RAC One Node Database | 6 |
| Rolling Patch Application and Instance Relocation using Omotion | 6 |
| Checking the Status of Oracle RAC One Node..... | 9 |
| Fixing Configuration Issues | 10 |
| Scaling Oracle RAC One Node to Oracle RAC..... | 11 |
| Best Practices for Deploying Oracle RAC One Node..... | 12 |
| Conclusion | 13 |

Introduction

Oracle Database 11g Release 2 introduces a new option, 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. This option adds to the flexibility that Oracle offers for consolidation. Many databases can be consolidated into a single cluster with minimal overhead while providing the high availability benefits of failure protection, online rolling patch application, as well as rolling upgrades for the operating system and Oracle Clusterware.

You can limit CPU utilization of individual database instances within the cluster through Resource Manager Instance Caging¹. You can dynamically change this limit if needed. With Oracle RAC One Node, there is no limit to server scalability and if applications grow to require more resources than a single node can supply, then you can easily scale up your single instance database online to a full Oracle Real Application Clusters database. If the node that is running your Oracle RAC One Node becomes overloaded, you can migrate the instance to another node in the cluster using the Omotion utility with no downtime for application users.

For Oracle Database 11g (11.2.0.1) Oracle RAC One Node will be available on Linux only. Oracle RAC One Node will be supported on all platforms where Oracle RAC 11g Release 2 is certified with Patchset 1 (11.2.0.2). Oracle RAC One Node is not supported if you are using a 3rd Party Clusterware such as Veritas SFRAC, IBM PowerHA, Sun Solaris Cluster, HP Serviceguard. Oracle RAC One Node (11.2.0.1) is not supported with Oracle Data Guard. With Oracle RAC and Oracle RAC One Node, you can standardize your deployments across the data center, achieving the required level of scalability and high availability for your applications. Similar to Oracle RAC, Oracle RAC One Node will be certified on Oracle Virtual Machine (OVM). Oracle VM is a free, next-generation server virtualization and management solution that makes enterprise applications easier to deploy, manage, and support. Using Oracle RAC or Oracle RAC

¹ For more information on instance caging see the Oracle Database 11g Release 2 Administrators Guide
http://download.oracle.com/docs/cd/E11882_01/server.112/e10595/dbrm007.htm#CFHBDFAG

One Node with Oracle VM increases the benefit of Oracle VM with the high availability and scalability of Oracle RAC. If your VM is sized too small, then you can migrate online the Oracle RAC One instance to another Oracle VM node in your cluster (using Omotion), and then resize the Oracle VM. When you move the Oracle RAC One Node instance back to the newly resized Oracle VM node, you can dynamically increase any limits programmed with Resource Manager Instance Caging.

This document is a guide to creating an Oracle RAC One Node database and a user's guide to the Omotion utility as well as other management utilities that comes with Oracle RAC One Node. Since Oracle RAC One Node is a specialized version of Oracle RAC, the Oracle Real Application Clusters Installation Guide and the Deployment and Administration Guide should be used for normal database management activities.

To use Oracle RAC One Node, you must install Oracle Database 11g Release 2 Enterprise Edition with the Oracle Real Application Clusters Option.

Software Installation

The pre-requisite for installing an Oracle RAC One Node database is that you have a cluster. This requires that you have installed the grid infrastructure and have a cluster up and running. The installation of Oracle RAC One Node is a hybrid installation of Oracle Real Application Clusters and you should follow the Oracle Grid Infrastructure Installation Guide² and Oracle Real Application Clusters Installation Guide³.

Recommendation: Complete a software only install of Oracle Real Application Clusters on all nodes in the cluster where you want to deploy Oracle RAC One Node, or nodes that are candidates for Omotion or cold failover.

To use Oracle RAC One Node Utilities discussed in this document, you must apply patch 9004119 to your 11.2.0.1 database home.

NOTE: You must have Oracle RAC software home on any node in the cluster where you could relocate (using Omotion) the Oracle RAC One Node instance or where it could failover to if the original node fails. You will be asked for this candidate list of nodes during the initialization of Oracle RAC One Node.

Creating and Configuring Oracle RAC One Node Database

Oracle RAC One Node is currently only supported as an Administrator Managed Database with Oracle Real Application Clusters 11g Release 2. Creating an Oracle RAC One Node database is the same as any other Oracle RAC database however this database has only one instance. From the Oracle RAC database home on the node where you want to run Oracle RAC One Node, start DBCA (the database

² http://www.oracle.com/pls/db112/to_toc?pathname=install.112/e10812/toc.htm

³ http://www.oracle.com/pls/db112/to_toc?pathname=install.112/e10813/toc.htm

configuration assistant). During Step 3 Database Identification, ensure the radial button for “Admin-Managed” is selected. Choose a Global Database Name, SID prefix and the node where you want to run the database. For Oracle RAC One Node, the Global Name must be the same as the SID prefix and they must be less than or equal to 8 characters in length.

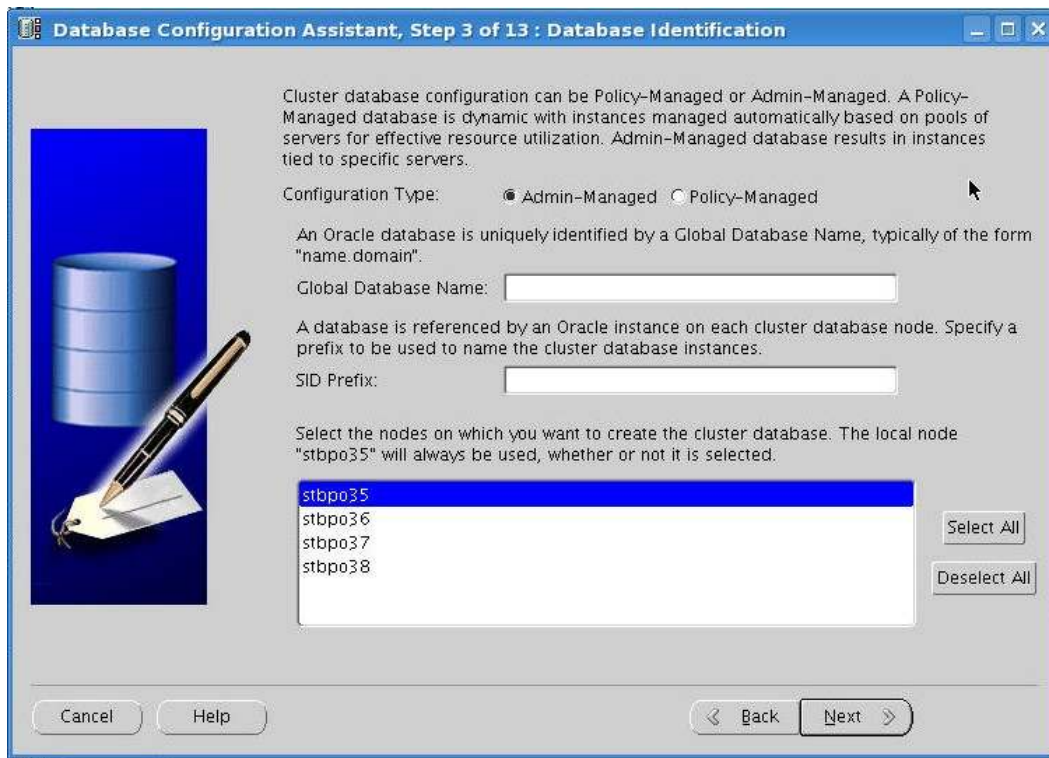


Figure 1: Choose Administrator Managed Database for Oracle RAC One Node

Once the database creation is complete, create any services you need for the application that will be using the Oracle RAC One Node database.

Recommendation: Create services *before* running the Oracle RAC One Node initialization utility, using the Cluster Managed Services page (found under the Availability section) in Oracle Enterprise Manager Database Control.

Oracle RAC One Node implements a specialized profile for the Oracle Clusterware resource for both the database and services. The profile can only be modified by the initialization utility. If you add services after you have run the initialization utility, you must run the initialization utility after creating the service.

The screenshot displays the Oracle Enterprise Manager Database Control interface for Cluster Managed Services. The page title is "Oracle Enterprise Manager (SYS) - Cluster Managed Database Services - Mozilla Firefox". The browser address bar shows "https://stbpo47.oracle.com:1158/em/console/rac/racServices?ev". The page content includes a navigation bar with "Cluster" and "Database" tabs, and a "Database Control" header. Below the header, there is a section for "Cluster Managed Database Services" with a "Refresh" button and a "Page Refreshed 1/30/09 2:26 PM" timestamp. A table of services is displayed with the following data:

| Select | Service Name | Status | Running Servers | Server Pool | Response Time (ms) | % CPU Load | Service related alerts among all Instances | Status Details |
|--------------------------|--------------|--------|------------------------------------|----------------------------|--------------------|------------|--|---|
| <input type="checkbox"/> | erp | ↑ | stbpo49, stbpo50 | ora.backoffice | 0.00 0.00 0.00 | 0.00 | 0 1 | UNIFORM Service is not running on all servers of the server pool. |
| <input type="checkbox"/> | hr | ↑ | stbpo49, stbpo50 | ora.backoffice | 0.00 0.00 0.00 | 0.00 | 0 1 | UNIFORM Service is not running on all servers of the server pool. |
| <input type="checkbox"/> | qos1 | ↑ | stbpo47, stbpo49, stbpo50, stbpo48 | ora.backoffice, ora.online | 0.00 2.45 7.80 | 0.21 | 0 0 | Default Database Service is running. |
| <input type="checkbox"/> | sales | ↑ | stbpo47, stbpo48 | ora.online | 0.00 0.00 0.00 | 0.00 | 0 1 | UNIFORM Service is not running on all servers of the server pool. |

Below the table, there is a "TIP" indicating that Response Time and %CPU Load data is averaged over the last 5 minutes. The page also includes a "Return" button and a footer with copyright information and the URL "stbpo47.oracle.com:1158".

Figure 2: Oracle Enterprise Manager Database Control - Cluster Managed Services Page

Pre-requisites for using the Oracle RAC One Node Utilities

The following pre-requisites must be in place before executing any of the Oracle RAC One Node utilities.

1. You must have applied patch 9004119 to the Oracle Database home
2. Set the ORACLE_HOME environment variable to the database home used for the Oracle RAC One Node database
3. You must be logged in as the owner of the Oracle Database home
4. SSH must be set for the user running the utility

The Oracle RAC One Node utilities can be run from any node in the cluster. The default location for the logs created by the utilities is the \$ORACLE_HOME/log/racone/, if the user has permissions to write to the directory, where ORACLE_HOME is the database home for the Oracle RAC One Node database. Otherwise the logs will be in the users \$HOME directory I.E /home/oracle/racone/*.log

How to Configure an Oracle RAC One Node Database

As the user who is the owner of the Oracle RAC Database Home (Oracle in the example below), run the configuration utility. First choose the database from the list that you want to configure. You will then be prompted to enter the server(s) where this database is allowed to run. This list will be used for cold failover of the database as well as candidate servers when the Omotion utility is run. The utility will provide you with a list of servers that are part of the cluster, which you can use for this list. The servers must currently either be in the FREE pool or in the GENERIC pool. Servers that are part of a user defined server pool are not eligible as candidate servers for Oracle RAC One Node. The server where you created the database will not be seen in the list and **DO NOT** enter it in your response. The server where the instance is running will always be in the candidate list. If you enter the server where the instance is running, the utility will exit with an error.

Note: the configuration utility must be run in order for the other Oracle RAC One Node utilities to work.

```

$ ./raconeinit

Candidate Databases on this cluster:

#      Database      RAC One Node      Fix Required
===      =====      =====
[1]      adm1              YES                N
[2]      adm2              YES                N
[3]      RacOne1           YES                N

Enter the database to initialize [1]: 1

Database adm1 is now running on server lnx3

Candidate servers that may be used for this DB:  lnx1 lnx2

Enter the names of additional candidate servers where this DB may
run (space delimited): lnx1

Database configuration modified

```

Figure 3: Sample of Oracle RAC One Node configuration utility

The raconeinit utility will modify the instance name to dbname_1. i.e. for the example above the instance name is changed from adm1 to adm_1.

Rolling Patch Application and Instance Relocation using Omotion

The Omotion utility allows you to move the Oracle RAC One Node instance from one node to another in the cluster. There are several reasons you may want to move the instance such as the node is overloaded so you need to balance the workload by moving the instance, or you need to do some operating system maintenance on the node but you want to eliminate the outage for application users by

moving the instance to another node in the cluster. To migrate an instance without client interruption, the application connections must use client connections that are integrated with Fast Application Notification (FAN). For information on using FAN, please read the Oracle Real Application Clusters Administration and Deployment Guide as well as the Automatic Workload Management Technical White Paper available on OTN⁴. You can also use connections with Transparent Application Failover (TAF) enabled (When using TAF, you should always enable FAN). For more information on TAF, please read the Oracle Net Services documentation⁵ and the Oracle Call Interface Programmers Guide⁶. If you do not use FAN or TAF, any in-flight transactions will be allowed to complete (as long as they complete within the timeout period entered which has a maximum of 30 minutes) then clients will receive an error when their session is terminated due to the shutdown of the Oracle RAC One Node instance (E.G. ORA-3113 End of Line on communication channel). Since the new instance will be running, the client can immediately login again.

Note: If the shutdown of the original instance takes longer than the set timeout, we will abort this database instance. The new instance will then perform recovery to cleanup any transactions that were aborted due to the shutdown.

Recommendation: Tune instance recovery to meet the needs of your organization.⁷

From any node in the cluster, run the Omotion utility. The utility will prompt you to choose the database to be migrated, the maximum amount of time you want to wait for transactions to complete before shutting down the current Oracle RAC One Node instance, and the node in the cluster where you want to move the instance to. This is time in minutes to a maximum of 30 minutes:

```

$ ./Omotion

RAC One Node databases on this cluster:

#          Database          Server          Fix Required
===          =====          ==============          ==============
[1]          adm1          lnx3          N
[2]          adm2          lnx2          N
[3]          RacOne1          lnx2          N

Enter number of the database to migrate [1]: 2

```

⁴ <http://www.oracle.com/technology/products/database/clustering/pdf/awmrac11g.pdf>

⁵ http://www.oracle.com/pls/db112/portal.portal_db?selected=4&frame=#network_management

⁶ http://download.oracle.com/docs/cd/E11882_01/appdev.112/e10646/toc.htm

⁷ http://download.oracle.com/docs/cd/E11882_01/server.112/e10821/instance_tune.htm#PFGRF13010

```

Specify maximum time in minutes for migration to complete (max
30) [30]: 5

Available Target Server(s) :
#           Server           Available
===          =====          =====
[1]          lnx1             Y
[2]          lnx3             Y

Enter number of the target node [1]: 1

Omotion Started...
Starting target instance on  lnx1
Migrating sessions ...
Stopping source instance on  lnx2
Omotion Completed...

=== Current Status ===
Database adm2 is running on node lnx1

```

Figure 4: Sample Output from Omotion

If the migration of the instance was for a short time in order to complete maintenance, you can run Omotion after the maintenance is complete to move the instance back to its original node.

The Omotion utility provides a verbose option (-v) that provides additional details during the migration.

```

$ ./Omotion -v

RAC One Node databases on this cluster:

#           Database           Server           Fix Required
===          =====          =====          =====
[1]          adm1             lnx3             N
[2]          adm2             lnx2             N
[3]          RacOne1         lnx2             N

Enter number of the database to migrate [1]: 2

Specify maximum time in minutes for migration to complete (max
30) [30]: 5

adm2 Database is administrator managed .
adm2 Database is running in adm2 Server Pool.
Current Running instance: adm2_1

Current Active Server      : lnx2

```

```

Available Target Server(s) :
#           Server           Available
===          =====          =====
[1]         lnx1             Y
[2]         lnx3             Y

Enter number of the target node [1]: 1

Omotion Started...
Starting target instance on lnx1
Migrating sessions ...
Stopping source instance on lnx2
Omotion Completed...

=== Current Status ===
Database adm2 is running on node

```

Figure 5: Sample Output from Omotion with Verbose option

You can also provide the database name and migration time as inputs on the command line.

```

$ ./Omotion -d adm2 -m 5

Available Target Server(s) :
#           Server           Available
===          =====          =====
[1]         lnx1             Y
[2]         lnx3             Y

Enter number of the target node [1]: 1

Omotion Started...
Starting target instance on lnx1
Migrating sessions ...
Stopping source instance on lnx2
Omotion Completed...

=== Current Status ===
Database adm2 is running on node lnx1

```

Figure 6: Sample Output from Omotion with Command Line Inputs

Checking the Status of Oracle RAC One Node

To make it easy to see the status of your Oracle RAC One Node databases, a status utility is included. This can be run on any node in the cluster to provide a summary of the Oracle RAC One Node databases in the cluster. The output displays the node on which the databases are currently running, whether a fix is

required as a result of a failure, and the candidate servers for Omotion for instance migration for each database.

```

$ ./raconestatus

RAC One Node databases on this cluster:

Database UP Fix Required Current Server Candidate Server Names
===== == =====
      adm1 Y      N      lnx3      lnx3 lnx2 lnx1
      adm2 Y      N      lnx2      lnx2 lnx3
      RacOne1 Y      N      lnx2      lnx2 lnx1

Available Free Servers:

```

Figure 7: Sample Output from Oracle RAC One Node Status utility

Fixing Configuration Issues

When you run any of the Oracle RAC One Node utilities, they will check the configuration and exit if they find any problem with the configuration. A problem in the configuration is generally a result of something happening in the cluster (e.g. a node rebooted) while the Omotion utility was running. To reset the configuration, run the raconfix utility. The raconfix utility will display all RAC One Node databases running on the cluster and whether there is a configuration issue. If the “Need Fix” column has a Y, then raconfix utility needs to be run. Enter the number of the database to be fixed when prompted.

```

$ ./raconefix

RAC One Node databases on this cluster:

#          Database          Server          Fix Required
===          =====          ==============          ==============
[1]          adm1             lnx3             N
[2]          adm2             lnx2, lnx1       Y
[3]          RacOne1         lnx2             N

Enter number of the database to fix [1]: 2

rac2a Service Server Names is different than the DB Server Names

This Database need to be fixed.

RACOneNode Database adm2 fixed.

```

Figure 8: Sample output from raconefix

Scaling Oracle RAC One Node to Oracle RAC

When your application requirements grow either in the availability requirements or the scalability requirements, it is a simple process to migrate your Oracle RAC One Node database to an Oracle RAC database. Since you are already running in a cluster, you need to run the `raconetorac` utility and then use DBCA to add new instances to the database. Note: During the migration, the Oracle RAC One Node instance will be shutdown and restarted.

```

$ ./racone2rac

Candidate Databases on this cluster:

#          Database          RACOne Node    Fix Required
===          =====          ==============          ==============
[1]          Admin             YES             N

Enter the database to convert to one node RAC database [1]:

You are converting Admin to single instance Oracle RAC database
running on lnx2.

```

```

Do you want to Continue? [Y]es or [N]o? Y
Please wait, this may take a few minutes to finish.....
Database configuration modified.
Run dbca to add as much instances as you want to this database.

```

Figure 9: Sample execution of the raconetorac utility

Best Practices for Deploying Oracle RAC One Node

Most database administration activities do not change with an Oracle RAC One Node database. There are a few items that you should be aware of if you are moving from a single instance database on a standalone server.

In an Oracle RAC database, each instance has its own thread of redo, its own redo log files, and its own undo tablespace. When you initiate Omotion, we bring up a second instance and as long as you are using Oracle Managed Files (OMF), we will automatically create an additional thread of redo, the appropriate redo logs, and the undo tablespace required for the new instance. The new redo and undo will be built using the same characteristics as the initial instance.

Recommendation: Size the redo and undo on the original instance to meet your applications requirements before initiating an Omotion. Note that if you modify them after the first Omotion has been executed; you will need to update all the redo logs and all the undo tablespaces.

All database connections should use the SCAN⁸. This will allow the client connections to easily connect to the database instance independently of which node in the cluster it is currently running on. Do not use node VIPs or Hostnames in the connection string since this will require modifying the client connect string whenever a failover or Omotion takes place.

In an Oracle RAC database, instances can have different parameter settings. When Omotion is used to move the database instance, we bring up a second instance and shutdown the initial instance. In order to ensure that all instances have the same parameter settings, you must use SCOPE= both and SID=*? when modifying parameters.

⁸ Refer to Oracle 11g Release 2 Clusterware Administration and Deployment Guide for details on SCAN

Recommendation: To start and stop the Oracle RAC One Node database, use Oracle Enterprise Manager or `srvctl` (`srvctl start database -d dbname`). For full syntax of `srvctl`, refer to the Oracle Real Application Clusters Administration and Deployment Guide Appendix A.⁹

Conclusion

Oracle RAC One Node provides an excellent way to reduce your server footprint, improve availability, better manage workloads, reduce maintenance outages, streamline your database administration processes, and create an "on-ramp" to upgrade to a full multi-node Oracle RAC implementation sometime in the future.

This document provides a guide to the utilities that are delivered as part of Oracle RAC One Node 11g Release 2 (11.2.0.1). Since Oracle RAC One Node is a specialized implementation of Oracle RAC, database administrators should follow the standard Oracle Documentation.

⁹ http://download.oracle.com/docs/cd/E11882_01/rac.112/e10718/srvctladmin.htm#CDCGICIF



Oracle RAC One Node 11g Release 2

User Guide

November 2009

Author: Barb Lundhild

Contributing Authors: Aiman Al-Khammash, Kirk McGowan

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



| Oracle is committed to developing practices and products that help protect the environment

Copyright © 2009, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.