

Rapid Oracle RAC One Node  
Standby Deployment: Oracle  
Database 11g Release 2

*Oracle Maximum Availability Architecture White Paper  
September 2011*

# Maximum Availability Architecture

Oracle Best Practices For High Availability

**ORACLE**

Introduction .....	1
Initial Preparation and Setup .....	2
Primary Cluster .....	2
Standby Cluster .....	3
Process Overview .....	3
Deployment Example .....	4
Create and Configure the Primary Database .....	4
Collect Primary Database Information and Prepare the Standby Cluster .....	5
Use DBCA to Create an Oracle RAC One Node Database on the Standby Cluster ..	6
Update Network Configuration Files .....	8
Prepare the Standby Cluster .....	8
Use RMAN to Create the Standby Database .....	11
Start Redo Apply and Enable Oracle Flashback Database .....	13
Possible Next Steps .....	14
Conclusion .....	14
References .....	14

## Introduction

The previously released white paper [Rapid Oracle RAC Standby Deployment: Oracle Database 11g Release 2](#) introduces a rapid deployment process that efficiently creates an Oracle Real Applications Cluster (RAC) physical standby database from an existing Oracle RAC primary database. This paper describes how the same process can be easily extended for use with Oracle RAC One Node databases by including a small number of additional steps. By using the Oracle Database Configuration Assistant (DBCA) and Oracle Recovery Manager (RMAN) to perform most of the work, many otherwise manual configuration steps are eliminated and the time required to deploy the standby database is reduced. Because DBCA and RMAN automate standby database configuration steps across multiple cluster nodes, the rapid deployment process described in this paper provides the most benefit when used to deploy standby databases on larger clusters.

This document is part of a series of Oracle [Maximum Availability Architecture \(MAA\)](#) [1] white papers that provide best practices blueprints based on proven Oracle technologies and recommendations that remove complexity from the design of optimal high availability solutions.

The example described in this document is based on functionality provided in Oracle Database 11g Release 2 (11.2.0.3) or later.

*Oracle RAC One Node databases configured with Oracle Data Guard must be Oracle Database 11g Release 2 or later.*

## Initial Preparation and Setup

The rapid Oracle RAC One Node standby deployment process works best when the primary and standby databases are configured similarly on similar hardware. This makes configuration and administration easier and ensures that sufficient resources will be available if a role transition occurs due to planned maintenance or in response to a disaster.

The example configuration is deployed on two geographically separated two-node clusters with local storage on each cluster node for Oracle home directories and common iSCSI shared Automatic Storage Management (ASM) storage on each cluster for database data, log, and configuration files. When used in a production environment, the steps described in this document should be carefully reviewed and augmented with any necessary additional best practices.

### Primary Cluster

The instructions given in this document assume the following about the primary cluster:

- Primary cluster node names are node1 and node2 and together they form cluster clus1.
- ASM disk groups +DATA, and +RECO are configured and mounted on all primary cluster nodes.
- Oracle Clusterware 11g Release 2 (11.2) is installed and configured on a private disk on each primary node with home directory /private1/grid. The cluster OCR and Voting Disk files are located on shared ASM storage.
- Oracle Database 11g Release 2 (11.2) is installed on a private disk on each primary node with home directory /private1/oracle/product/11.2.0/db\_1.
- ORACLE\_BASE is set to /private2/oracle to provide a separate area for log and trace output.
- SCAN is clus1-s with SCAN listener port 1523. The local listener on each node is LISTENER with default port 1521.

*The SCAN and SCAN listener are new Oracle Database 11g Release 2 features.*

- The administrator managed RAC One Node primary database is called “east.us.oracle.com” with initial instance “east\_1” on host node1 and service “primary” for client access.
- The initialization parameters **db\_name** and **db\_unique\_name** are both set to “east”.
- ARCHIVELOG mode, FORCE LOGGING mode, and Oracle Flashback Database are enabled on the primary database.

*FORCE LOGGING mode protects against unlogged changes to the primary database that otherwise would not be propagated to the standby database.*

- A standby redo log with an appropriate number of standby redo log groups for each primary database thread has been added to the primary database. Note that a RAC One Node database has two threads (one thread for each of the two possible instances).
- The primary database data files, online log files, server parameter file, and one control file are located on data disk group +DATA. The primary database archive, standby redo, and flashback log files along with a second control file are located on fast recovery area disk group +RECO.

### Standby Cluster

The instructions given in this document assume the following about the standby cluster:

- Standby cluster node names are node3 and node4 and together they form cluster clus2.
- ASM disk groups +DATA, and +RECO are configured and mounted on all standby cluster nodes.
- Oracle Clusterware 11g Release 2 (11.2) is installed and configured on a private disk on each standby node with home directory /private1/grid. The cluster OCR and Voting Disk files are located on shared ASM storage.
- Oracle Database 11g Release 2 (11.2) is installed on a private disk on each standby node with home directory /private1/oracle/product/11.2.0/db\_1.
- ORACLE\_BASE is set to /private2/oracle to provide a separate area for log and trace output.
- SCAN is clus2-s with SCAN listener port 1527. The local listener on each node is LISTENER with default port 1521.
- The administrator managed standby database will be called “west.us.oracle.com” with initial instance “west\_1” on node3 and service “primary” for client access in the primary role.
- The standby database data files, online log files, server parameter file, and one control file will be located on data disk group +DATA. The standby database archive, standby redo, and flashback log files along with a second control file will be located on fast recovery area disk group +RECO.

### Process Overview

The following steps present an overview of the process used to create the Oracle RAC One Node standby database:

- If the Data Guard configuration is being created from scratch and a primary database does not yet exist, create the initial Oracle RAC One Node primary database on the primary cluster using DBCA and ensure all relevant [Data Guard initialization parameters](#) are configured correctly.

- Collect primary database information needed to create the standby database and create any needed directories on the standby cluster.
- Use DBCA to configure an Oracle RAC One Node database (west.us.oracle.com) on the standby cluster with parameters and structure similar to the primary database. Initially, both the **db\_name** and **db\_unique\_name** initialization parameters will be set to the same value (west). Creating the initial Oracle RAC One Node database on the standby cluster can be done at the same time the primary database is created on the primary cluster if there is no initial existing primary database.

Using DBCA in this way automates many standby database, cluster, and network configuration steps that would otherwise need to be performed manually on each cluster node.

- Update the network configuration files and verify connectivity on both clusters and on any client systems to ensure that the network service names and static services required for use with Oracle Data Guard are correctly configured.
- Prepare the standby database cluster for standby database creation by:
  - shutting down the Oracle RAC One Node database on the standby cluster.
  - deleting the initial data files, control files, online redo log files, and archived log files.
  - changing the value of the **db\_name** initialization parameter in the standby database server parameter file from its initial value (“west”) to match that of the primary database (“east”).
- Use the Oracle Recovery Manager (RMAN) to create the standby database data files and standby control files from the primary database and to replace the existing standby database password file with a copy of the primary database password file.
- Start redo apply for the standby database
- Enable Oracle Flashback Database for the standby database.

## Deployment Example

The following example uses Oracle Database 11g Release 2 (11.2.0.3).

### Create and Configure the Primary Database

If this is a new deployment, use DBCA to create the primary database. Ensure that ARCHIVELOG mode, FORCE LOGGING mode and Oracle Flashback Database are enabled for the primary database and that a standby redo log is added. For this example configuration, the primary standby redo log was configured to use multiple standby redo log file groups in the fast recovery area by issuing the following commands while connected to the primary database as SYSDBA:

```
SQL> alter database add standby logfile thread 1 group 5 size 50M,
  2 group 6 size 50M, group 7 size 50M;
Database altered.
SQL> alter database add standby logfile thread 2 group 8 size 50M,
  2 group 9 size 50M, group 10 size 50M;
Database altered.
```

*Each thread of the standby redo log must have at least one more redo log group than the corresponding thread of the redo log at the redo source database. For the example configuration, each thread of the redo log is configured with two log groups, and each thread of the standby redo log is therefore configured with three log groups. For actively used databases, performance may be improved by adding additional redo log groups and standby redo log groups for each redo thread.*

To facilitate future role transitions and other Data Guard operations, update the primary database server parameter file with appropriate values for the following Data Guard related initialization parameters (if they were not already configured when the primary database was created):

```
SQL> alter system set db_file_name_convert='/WEST/', '/EAST/', '/west/', '/east/'
  2 scope=spfile;
System altered.
SQL> alter system set dg_broker_config_file1='+DATA/east/drleast.dat' scope=both;
System altered.
SQL> alter system set dg_broker_config_file2='+DATA/east/dr2east.dat' scope=both;
System altered.
SQL> alter system set fal_server=west scope=both;
System altered.
SQL> alter system set log_archive_config='DG_CONFIG=(east,west)' scope=both;
System altered.
SQL> alter system set log_archive_dest_1='LOCATION=USE_DB_RECOVERY_FILE_DEST
  2 VALID_FOR=(ALL_LOGFILES,ALL_ROLES) DB_UNIQUE_NAME=east' scope=both;
System altered.
SQL> alter system set log_archive_dest_2='SERVICE=west ASYNC
  2 VALID_FOR=(ONLINE_LOGFILES,PRIMARY_ROLE) DB_UNIQUE_NAME=west' scope=both;
System altered.
SQL> alter system set log_archive_max_processes=30 scope=both;
System altered.
SQL> alter system set log_file_name_convert='/WEST/', '/EAST/', '/west/', '/east/'
  2 scope=spfile;
System altered.
SQL> alter system set standby_file_management=AUTO;
System altered.
```

The new values for the static initialization parameters **db\_file\_name\_convert** and **log\_file\_name\_convert** are set in the primary database server parameter file, but not yet in memory. Because these parameters are not used while the database remains in the primary role, updating the in memory values can be deferred until the next time the primary database is restarted (for example, during planned maintenance or during a role transition).

## Collect Primary Database Information and Prepare the Standby Cluster

Connect to the primary database as SYSDBA and issue a SQL SHOW PARAMETERS command to identify any primary-specific parameters with values that will require changes when configured for the standby database (for example, **db\_file\_name\_convert**, **log\_file\_name\_convert**, or other parameters containing a file path or derived from **db\_name** or **db\_unique\_name**).

Based on the collected primary database initialization parameter information, create any needed directories on the standby cluster. Then use the `srvctl` utility to verify that the standby SCAN and local listeners are running by issuing the following commands from one of the standby nodes:

```
$ srvctl status scan
$ srvctl status listener
```

Before proceeding to the next step, all Oracle software installations, ASM disk group creation, and other previously noted initial preparation and setup steps must be completed.

## Use DBCA to Create an Oracle RAC One Node Database on the Standby Cluster

If you are creating a deployment from scratch, you can perform this step in parallel with creating the primary database. DBCA automatically takes care of most network, Oracle Clusterware, Automatic Storage Management, and other configuration details for this database. Because the data files and control files will be replaced later with copies from the primary database, they do not need to match those used on the primary (so it is fine to use smaller files such as those in the General Purpose seed database to make the database creation process faster). However, because the goal is to create an Oracle RAC One Node database on the standby cluster with a configuration similar to that of the primary database, the parameter values in the standby server parameter file should parallel as much as possible those used for the primary database.

### Start DBCA on a Standby Node

From the command line on a standby node, issue appropriate operating system-specific commands to configure the database-related environment variables and start DBCA. For the example Linux deployment, issue commands similar to the following:

```
$ setenv ORACLE_BASE /private2/oracle
$ setenv ORACLE_HOME /private1/oracle/product/11.2.0/db_1
$ setenv PATH ${ORACLE_HOME}/bin:${ORACLE_HOME}/OPatch:${PATH}
$ dbca
```

### Enter Information for Initial DBCA Wizard Panels

For the initial DBCA panels, select **Oracle RAC One Node database, Create a Database**, and **General Purpose or Transaction Processing**. On the Database Identification Panel, specify **Admin-Managed** and enter the appropriate database and cluster node information. Enter the same SYS user password as used on the primary database. For the Database File Location panel, specify database file locations that parallel the primary database file locations. For the example configuration, the Storage Type is **Automatic Storage Management (ASM)** and Oracle-Managed Files are used with Database Area **+DATA**. Although it is possible to use file locations for the standby database that do not parallel the primary database file locations if appropriate file name conversion parameters are specified, it is easier to manage and configure the primary and standby databases if they are configured similarly. On the Recovery Configuration panel, specify the location (**+RECO** for the example) and size of the fast recovery area and enable archiving.

### Configure Standby Database Initialization Parameters

On the Initialization Parameters panel, enter appropriate information for memory use, click **All Initialization Parameters**, and then click **Show Advanced Parameters**.

Specify appropriate values, such as those shown in *Table 1*, for the **db\_unique\_name** and **standby\_file\_management** parameters, file name conversion parameters, broker configuration file location parameters, and any other parameters that may need adjustment based on your environment. To ensure that the Oracle Data Guard broker can be enabled successfully if later desired, specify a location on shared storage accessible to all database instances (in this example, on the ASM disk +DATA) for the broker configuration files. The example configuration also configures the fast recovery area as a local destination for archived log files because it is an Oracle recommended best practice. Alternatively, you can also configure these parameters using SQL after the database has been created.

INITIALIZATION PARAMETER	VALUE USED FOR EXAMPLE DEPLOYMENT
<b>db_file_name_convert</b>	'/EAST/','/WEST/','/east/','/west/'
<b>db_unique_name</b>	west
<b>dg_broker_config_file1</b>	+DATA/west/dr1west.dat
<b>dg_broker_config_file2</b>	+DATA/west/dr2west.dat
<b>fal_server</b>	east
<b>log_archive_config</b>	DG_CONFIG=(east,west)
<b>log_archive_dest_1</b>	LOCATION=USE_DB_RECOVERY_FILE_DEST VALID_FOR=(ALL_LOGFILES,ALL_ROLES) DB_UNIQUE_NAME=west
<b>log_archive_dest_2</b>	SERVICE=east ASYNC VALID_FOR=(ONLINE_LOGFILES,PRIMARY_ROLE) DB_UNIQUE_NAME=east
<b>log_archive_max_processes</b>	30
<b>log_file_name_convert</b>	'/EAST/','/WEST/','/east/','/west/'
<b>standby_file_management</b>	AUTO

*Table 1: Standby Database Initialization Parameters*

*DBCA automatically updates the CONTROL\_FILES initialization parameter with the correct Oracle-Managed ASM disk locations when the database is created and no manual changes to this parameter are required*

### Complete Remaining DBCA Steps

Click through the remaining DBCA panels and make any needed changes to the default values. On the final panel (or sooner if no additional changes are needed), click **Finish** and then, on the Summary screen, click **OK** to begin creating Oracle RAC One Node database west.

## Update Network Configuration Files

At this point there are two similarly configured but different Oracle RAC One Node databases, each in the primary role. To allow access to the primary and standby databases from all systems in the configuration, the `tnsnames.ora` file in the database home on each primary, standby, and client system must be updated with appropriate service name entries. Once a `tnsnames.ora` file has been updated with all relevant entries on one cluster node, it can be copied to the appropriate location on all other systems.

### Tnsnames.ora Files

Ensure that the following service name entries are included in the `tnsnames.ora` file for each database Oracle home. DBCA automatically adds a correctly formatted service name entry to the `tnsnames.ora` file on each cluster node that hosts an instance of a database at the time it creates the database. The `HOST` and `PORT` clauses in each service name entry contain, respectively, the `SCAN` and `SCAN` listener port associated with the cluster that hosts the database. Typically you will need to add entries for the primary database to the `tnsnames.ora` files on the standby cluster nodes and entries for the standby database to the `tnsnames.ora` files on the primary cluster nodes.

```
EAST =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = clus1-s)(PORT = 1523))
    (CONNECT_DATA =
      (SERVER = DEDICATED)(SERVICE_NAME = east.us.oracle.com)
    )
  )
WEST =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = clus2-s)(PORT = 1527))
    (CONNECT_DATA =
      (SERVER = DEDICATED)(SERVICE_NAME = west.us.oracle.com)
    )
  )
```

## Prepare the Standby Cluster

With the basic cluster and network configuration now in place, the next step is to prepare the standby cluster for standby database creation. Shut down the Oracle RAC One Node database on the standby cluster, delete any files that will be replaced later, and update the server parameter file to change the `db_name` initialization parameter value to match that of the primary database.

### Shut Down the Oracle RAC One Node Database on Standby Cluster

From a node on the standby cluster, use the `srvctl` utility to stop the Oracle RAC One Node database (“west”) on the standby cluster:

```
$ srvctl stop database -d west -o abort
```

### Delete Existing Data, Control, and Log Files

The existing Oracle RAC One Node database on the standby cluster was created as a placeholder so that DBCA could preconfigure the environment for the still-to-be-created Oracle RAC One Node physical standby database. This approach also allows network access to be verified before the standby database is created. The existing data, control, and log files from the placeholder Oracle RAC One Node database will not be used by the standby database and must be deleted. Because the standby database is configured with ASM storage, the ASM command line utility (ASMCMD) from the Oracle Clusterware home must be used to perform these tasks. For the example configuration, the following commands ensure that empty versions of the standby directories exist on the ASM disk groups (issued here from the second standby cluster node, node4, which hosts ASM instance +ASM2):

```
$ setenv ORACLE_BASE /private2/oracle
$ setenv ORACLE_HOME /private1/grid
$ setenv PATH ${ORACLE_HOME}/bin:${ORACLE_HOME}/OPatch:${PATH}
$ setenv ORACLE_SID +ASM2
$ asmcmd
ASMCMD> cd +DATA/WEST
ASMCMD> rm CONTROLFILE/*
You may delete multiple files and/or directories.
Are you sure? (y/n) y
ASMCMD> mkdir CONTROLFILE
ASMCMD> rm DATAFILE/*
You may delete multiple files and/or directories.
Are you sure? (y/n) y
ASMCMD> mkdir DATAFILE
ASMCMD> rm ONLINELOG/*
You may delete multiple files and/or directories.
Are you sure? (y/n) y
ASMCMD> mkdir ONLINELOG
ASMCMD> rm TEMPFILE/*
You may delete multiple files and/or directories.
Are you sure? (y/n) y
ASMCMD> mkdir TEMPFILE
ASMCMD> cd +RECO/WEST
ASMCMD> rm ARCHIVELOG/**
You may delete multiple files and/or directories.
Are you sure? (y/n) y
ASMCMD> mkdir ARCHIVELOG
ASMCMD> rm CONTROLFILE/*
You may delete multiple files and/or directories.
Are you sure? (y/n) y
ASMCMD> mkdir CONTROLFILE
ASMCMD> rm ONLINELOG/*
You may delete multiple files and/or directories.
Are you sure? (y/n) y
ASMCMD> mkdir ONLINELOG
ASMCMD> exit
```

*The mkdir commands are required only if ASM or OMF also automatically deleted the directory when the contents were deleted.*

### Update the Standby Server Parameter File and Start the Standby Database

The initial values of the **db\_name** and **db\_unique\_name** initialization parameters were used by DBCA to determine default values for other initialization parameters and to identify where to locate many database related files on disk. Both parameters were initially set to a value of “west” to ensure that the directory structures and file names created on the standby cluster would parallel the directory structures and file names that were created using “east” on the primary cluster. However, because the physical standby database will be an exact copy of the primary database, the value for **db\_name** in the standby server parameter file (currently “west”) must now be changed to match that of **db\_name** on the primary (“east”). Note that the value for **db\_unique\_name** must remain “west” to uniquely identify the standby database.

For the example configuration, issue the following commands from a standby cluster node:

```
$ setenv ORACLE_HOME /private1/oracle/product/11.2.0/db_1
$ setenv PATH ${ORACLE_HOME}/bin:${ORACLE_HOME}/OPatch:${PATH}
$ srvctl start database -d west -o nomount
$ srvctl status database -d west
Instance west_1 is running on node node3
Online relocation: INACTIVE
```

Then, from the node that is currently hosting instance west\_1:

```
$ setenv ORACLE_HOME /private1/oracle/product/11.2.0/db_1
$ setenv PATH ${ORACLE_HOME}/bin:${ORACLE_HOME}/OPatch:${PATH}
$ setenv ORACLE_SID west_1
$ sqlplus /nolog
SQL> connect / as sysdba
Connected.
SQL> alter system set db_name='east' scope=spfile;
System altered.
SQL> exit
$ srvctl stop database -d west
$ srvctl start database -d west -o nomount
```

### Create a Static Service and Temporary Service Name Entry for the Standby Instance

RMAN must connect to an unmounted standby database instance to execute the DUPLICATE database command. A static listener service is required to make this connection. From the command line on the cluster node currently hosting the active RAC One Node standby instance, issue appropriate operating system-specific commands to configure the Oracle Clusterware related environment variables and then use the netmgr utility to temporarily add a static listener service stby.us.oracle.com to the local listener in the listener.ora file in the Oracle Clusterware home:

```
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (GLOBAL_DBNAME = stby.us.oracle.com)
      (ORACLE_HOME = /private1/oracle/product/11.2.0/db_1)
      (SID_NAME = west_1)
    )
  )
)
```

After exiting netmgr, reload the local listener to enable the static service by issuing the command:

```
$ lsnrctl reload LISTENER
```

From the command line, issue appropriate operating system-specific commands to configure the database-related environment variables and use netmgr to add a service name entry STBY to the tnsnames.ora file in the network/admin directory of the active standby instance Oracle home. Use the physical host name instead of the SCAN in the HOST clause. For the example configuration, the local listener is listening on the default port 1521 and the relevant tnsnames.ora entry is:

```
STBY =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST =node3.us.oracle.com)(PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)(SERVICE_NAME = stby.us.oracle.com)
      (INSTANCE_NAME = west_1)
    )
  )
)
```

Repeat this process (or simply copy the modified tnsnames.ora file) to update the tnsnames.ora file in the Oracle home for the active primary database instance and verify that connections from the nodes hosting active database instances can be made using the STBY service name entry.

## Use RMAN to Create the Standby Database

The Oracle Recovery Manager is used to complete the rapid Oracle RAC One Node standby deployment process. An RMAN script file is created and executed to configure the standby database data, control, and log files and to replace the existing password file for the standby database with a copy of the primary database password file. Some additional steps are then performed to clean up the environment and complete the standby database configuration.

### Create the RMAN Script File

Using a text editor of your choice, create and save a text file named *create\_standby.txt* containing the RMAN commands required to create the standby database from the active primary database:

```
run {
  allocate channel prmy1 type disk;
  allocate channel prmy2 type disk;
  allocate channel prmy3 type disk;
  allocate channel prmy4 type disk;
  allocate auxiliary channel stby type disk;
  duplicate target database for standby from active database;
}
```

### Execute RMAN Commands

Start RMAN on the node where instance west\_1 is running and issue the commands required to create the standby database. All RMAN connections to the primary database during the standby creation process must be through a service name that always connects to the same instance and

that is not load balanced over other instances. Because there is only one active instance for the RAC One Node primary database east, the tnsnames.ora service **east** can be used for the primary connection. The commands should look similar to the following

```
$ rman
RMAN> connect target sys/*****@east
connected to target database: EAST (DBID=3664244207)
RMAN> connect auxiliary sys/*****@stby
connected to auxiliary database: EAST (not mounted)
RMAN> @create_standby.txt
```

### Copy the Updated Password File

RMAN has replaced the password file located in the dbs directory of the active RAC One Node standby database with a renamed copy of the password file from the primary database east. Copy the updated password file (orapwwest) from the dbs directory of the active RAC One Node instance to the dbs directory of each additional instance on the standby cluster. This will ensure that the Data Guard environment continues to function if the standby database is ever started on a different cluster node. This step is only required when a separate copy of the database password file exists on each cluster node (for example, when ASM storage is used).

*If the password file is located on a shared cluster file system (such as OCFS2) RMAN will update the single shared copy of the password file and it is not necessary to manually copy any files. In this case, the password file entries in the dbs directories are actually symbolic links (for example, orapwwest -> /dbshare/oradata/west/orapwwest) that each point to the single shared password file located on the cluster file system.*

### Update the Role of the Standby Database Recorded in the Oracle Clusterware Repository

The standby database was originally created by DBCA in the primary role; to update Oracle Clusterware with the new role to ensure that the correct role-based services are started whenever the standby database is opened, issue the following command on a node on the standby cluster:

```
$ srvctl modify database -d west -r physical_standby
```

*The Data Guard broker automatically updates Oracle Clusterware with any role changes during broker initiated role transitions.*

### Remove the Static Service and Temporary Service Name Entry

The stby.us.oracle.com static service and STBY service name are no longer needed after the standby is created. Use netmgr to remove the static service entry from the listener.ora file in the Oracle Clusterware home on the active standby database system and to remove the temporary STBY service name entry from the tnsnames.ora files in the database homes of the active instances of the primary and standby databases.

## Start Redo Apply and Enable Oracle Flashback Database

The primary and standby Oracle RAC One Node databases are now configured, but redo apply is not started. To start redo apply, issue the following SQL command when connected to the standby database as SYSDBA:

```
SQL> alter database recover managed standby database using current logfile disconnect;
Database altered.
```

After redo apply has been started, switch logs on the primary database and verify by viewing the standby database alert log that logs are received from the primary database and applied on the standby database. Transport lag and apply lag also can be monitored by connecting to the standby database as SYSDBA and issuing the following query:

```
SQL> select name, value from v$dataguard_stats where name like '%lag';
```

### Enable Oracle Flashback Database

Enabling Oracle Flashback Database is an Oracle recommended best practice. If you plan to use Oracle Data Guard broker's fast-start failover feature, Oracle Flashback Database must be enabled on both the primary and standby databases.

Oracle Flashback Database should already have been enabled on the primary database as part of the initial setup. To enable Oracle Flashback Database on the newly created standby database, connect to the standby database as SYSDBA and issue the following commands after the standby database has become current with the primary:

```
SQL> alter database recover managed standby database cancel;
Database altered.
SQL> alter database flashback on;
Database altered.
SQL> alter database recover managed standby database using current logfile disconnect;
Database altered.
```

### Database Relocation Considerations

The `srvctl relocate database` command can be used to move the active RAC One node instance to another cluster node. For administrator managed databases, the `-n` option is used to specify the new host; for policy managed databases, the `-n` option is not used and Oracle Clusterware determines the new host based on user-specified policies. When used with standby databases, `srvctl relocate database` is supported only for the maximum performance or maximum availability protection modes. Attempts to relocate a standby database in a maximum protection configuration will not succeed and the following messages will be displayed:

```
$ srvctl relocate database -d west
PRCD-1222 : Online relocation of database west failed but database was restored to
its original state
PRCD-1246 : Online relocation of RAC One Node Data Guard standby database west was
rejected because the database is in Data Guard maximum protection mode
```

To relocate the standby database in this case, lower the protection mode to maximum availability, issue the relocation command, and then restore the protection mode to maximum protection. There are no restrictions when using `srvctl relocate database` on a primary database.

Note that redo apply is not automatically restarted on a standby database after a relocation or restart unless the database is enabled as part of an Oracle Data Guard broker configuration.

### Possible Next Steps

- Refer to the [Oracle Data Guard Concepts and Administration](#) manual for more information about:
  - Opening a physical standby database for read-only access (and also optionally enabling the Oracle Active Data Guard option).
  - Converting a physical standby database into a snapshot standby database.
- Refer to the [Oracle Data Guard Broker](#) manual for more information about:
  - Enabling the Oracle Data Guard broker and creating a broker configuration.
  - Configuring role-based services and database connection failover for client applications.
  - Configuring fast-start failover and specifying the conditions under which automatic failover to the standby database will occur.

### Conclusion

The rapid Oracle RAC One Node standby deployment process allows you to directly create an Oracle RAC One Node standby database without any primary database downtime. In addition, most of the creation and configuration steps are automated by using DBCA and RMAN. This automation eliminates the risk of many potential errors during deployment and significantly reduces the time required to create an Oracle RAC One Node physical standby database. When deploying a new configuration from scratch, the creation of the primary database and initial RAC One Node database on the standby cluster can be performed in parallel.

### References

1. Oracle Maximum Availability Architecture Web Site on OTN  
<http://www.oracle.com/technetwork/database/features/availability/maa-090890.html>
2. Oracle Data Guard Web Site on OTN  
<http://www.oracle.com/technetwork/database/features/availability/dataguardoverview-083155.html>
3. Oracle Real Application Clusters Web Site on OTN  
<http://www.oracle.com/technetwork/database/clustering/overview/index.html>

4. Oracle Clusterware Web Site on OTN  
<http://www.oracle.com/technetwork/database/clusterware/overview/index.htm>
5. Oracle Automatic Storage Management Administrator's Guide  
[http://download.oracle.com/docs/cd/E11882\\_01/server.112/e16102/toc.htm](http://download.oracle.com/docs/cd/E11882_01/server.112/e16102/toc.htm)



Rapid Oracle RAC One Node Standby  
Deployment:  
Oracle Database 11g Release 2  
September, 2011  
Author: Laurence Clarke

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 © 2011, 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.