Using Grid Control to Implement or Extend High Availability with Oracle Database 11*g* and Oracle Data Guard

Oracle Maximum Availability Architecture White Paper
March 2011

Maximum Availability Architecture

Oracle Best Practices For High Availability

ORACLE

Introduction

This white paper provides a step-by-step walkthrough of the tasks and processes involved in using Oracle Enterprise Manager Grid Control Release 10.2.0.5 to configure a complete Oracle Maximum Availability Architecture (MAA) implementation on Oracle Database 11g release 11.1.0.7.

The procedures in this paper chronicle how to use Grid Control release 10.2.0.5 automated processes to take a single-instance Oracle Database and configure an MAA environment, incurring minimal downtime. The ending MAA configuration included Oracle Database 11*g* release 1 (11.1.0.7), Oracle Data Guard, Oracle Clusterware, Oracle Real Application Clusters (Oracle RAC), Oracle Automatic Storage Management (Oracle ASM), and other Oracle high availability features. The MAA team used Grid Control provisioning and Deployment Procedures to automate each stage of configuration, and monitored and managed the complete Oracle IT infrastructure from a single Grid Control console.

Oracle Maximum Availability Architecture (MAA) is the Oracle best practices blueprint based on proven Oracle high availability technologies and recommendations. The goal of MAA is to achieve the optimal high availability architecture at the lowest cost and complexity. More white papers are published on the Oracle Technology Network (OTN) at http://www.otn.oracle.com/goto/maa. Also, for more detailed information about specific Oracle products and features, see the documents listed in the References section.

MAA Project Background

Prior to release 10.2.0.5, Grid Control incorporated many of the configuration tasks for implementing a highly available architecture, but it did not encompass all of the MAA best practices that deliver a near-zero downtime experience. With Oracle Enterprise Manager Grid Control Release 10.2.0.5 and Oracle Database 11*g* release 1 (11.1.0.7), you can migrate an existing

database to a complete MAA solution with minimal downtime, and perform virtually all steps in an automated fashion within the Grid Control environment. Moreover, this solution provides the flexibility to transition your configuration to MAA in phases, giving you the option of migrating components within the bounds of your organization's maintenance and budget requirements. Grid Control can be added to an existing environment at any phase of MAA migration to continue the process in this automated fashion.

This MAA white paper documents the tasks involved in migrating to an MAA configuration and the functionality available in Grid Control to perform the tasks. The point-and-click nature of Grid Control allows administrators to complete these tasks reliably and efficiently. The process is reentrant, allowing you to take an existing configuration that is at any stage of MAA implementation and use Grid Control to complete the work to extend its availability. The process enables you to:

- Perform complete, repeatable steps to deploy a single-instance database or an Oracle RAC database (with Oracle Clusterware, Oracle RAC, and Oracle ASM) in either a new or an existing environment. Grid Control processing handles all the steps necessary for software installation and configuration in a scheduled lights-out style of execution.
- Create and extend an MAA environment while incurring minimal downtime, with the capabilities and benefits described in the following table:

PROCESS	BENEFIT
Migrate to Oracle ASM by creating a standby with Oracle ASM and then switching over.	Downtime is limited to switchover to the standby database. Migrate to a more robust, flexible and available storage system.
Migrate from single instance to Oracle Real Application Cluster (Oracle RAC) by creating an Oracle RAC standby and then switching over.	Downtime is limited to switchover to the standby database.
Extend an existing Oracle RAC database to include more database instances and nodes by using the existing installation to copy software and database information to the new nodes.	Zero downtime is incurred. All settings from existing configuration are automatically applied to new nodes. No manual configuration is required.

Grid Control Overview

The starting configuration consisted of a single-instance Oracle database on local file system storage. This environment was extended to an ending configuration consisting of two dual-node Oracle RAC clusters running Oracle Data Guard and Oracle ASM. The workflow of all the software installation and configuration tasks that needed to be performed for each particular stage of activity was encapsulated in Grid Control Deployment Procedures, which is a process deployed to one or more targets through provisioning. The general workflow also took advantage of existing and new Grid Control wizards to complete many functions.

Besides describing how to deploy various configurations, this white paper helps you understand the architecture and flow of data among the Grid Control components. Based on this knowledge, you can make better decisions about how to configure Grid Control for your specific management requirements.

Grid Control Components

Although Grid Control is generally viewed as a single entity, technically, its architecture is composed of the following software components:

- Oracle Management Service (OMS)
- Oracle Management Agent (Management Agent)
- Oracle Management Repository (Management Repository)
- Oracle Enterprise Manager 10*g* Grid Control Console

Internally, the OMS orchestrates with Management Agents to discover targets (a logical entity in Grid Control that represents an object such as a database or listener), monitor and manage them, and store the collected information in the Management Repository for future reference and analysis. The OMS also renders the user interface for the Grid Control console, on which the health of the monitored targets is displayed.

Provisioning and Deployment Procedures

Provisioning is Grid Control's functionality that allows you to install and configure software and databases using Deployment Procedures. Deployment Procedures consist of a series of steps grouped together into a job format that perform pre-installation checking and setup, software installation, and post-installation configuration. These steps include performing privileged user (root) steps and creating databases or database instances as required or as requested.

The methodology outlined in this paper uses Grid Control Deployment Procedures[1] that provide an extremely flexible framework for automating tasks. The Deployment Procedures allow you the flexibility to:

- Modify the Deployment Procedures to include or remove steps, as necessary.
- Take into account shop security (for example, using the sudo command for privileged users or the software owner).
- Set an Error Handling Mode for each step of the procedure, controlling whether the procedure should continue to run if that particular step fails.

> **See Also:** *Oracle Enterprise Manager Advanced Configuration* [6] for complete information about using provisioning and Deployment Procedures

One of the building blocks of the Deployment Procedures is the Grid Control Components[2] element that is used by Deployment Procedures to mass deploy software and applications onto target servers. These elements allow you to use previously installed software that is configured to your environment as a source for future installations. Once the software is installed and configured to your liking, you upload the software into a software library (managed by the Grid Control OMS) and use it as a "gold" image with which you can deploy the software onto other servers. Configuring Components can entail applying patch sets or individual patches. The initial installation of each version and each software type must be performed manually.

> **Tip:** Saving the source software as a gold image is best suited for when you have a copy of stable, well-tested, and patched software installed in your environment to upload to the Software Library. The gold image allows for easier propagation of patched and tested software throughout your data center.

[1] Deployment procedures are licensed under the Provisioning and Patch Automation Pack.

[2] A component can represent operating system software, Oracle software, or any third-party software and applications. Software components are individually maintained in the Oracle Software Library. Versions, states, and maturity levels can be associated with each component. See the *Oracle Enterprise Manager Concepts* [7] documentation for more provisioning concepts.

Overview of the Process Flow

This white paper provides a comprehensive set of instructions and screenshots to help you use Oracle Grid Control to take a single-instance database and extend it to two two-node Oracle RAC clusters. You can perform the process in an end-to-end fashion, or you can use the steps as building blocks to take an existing configuration that is at any stage of MAA implementation and use Grid Control to complete the work to extend its availability.

This document provides a step-by-step guide for the configuration of an Oracle Database 11*g* MAA environment on UNIX/Linux. It is divided into modules that address specific tasks that guide you through the configuration process to implement an MAA environment either in stages to gradually extend the availability of any given configuration, or all at once for a complete transition to an MAA configuration. Table 1 describes the various starting and ending configurations.

Table 1: Starting and Target Configurations

		TARGET CONFIGURATION				
		SINGLE-INSTANCE DATABASE WITH ORACLE ASM	SINGLE-INSTANCE STANDBY DATABASE	ORACLE RAC	ORACLE RAC WITH A SINGLE-INSTANCE STANDBY DATABASE	COMPLETE MAA CONFIGURATION WITH RAC PRIMARY AND RAC STANDBY
STARTING CONFIGURATION	Single-Instance Database	Module 1	Module 1	Module 2	Module 2	Module 2
	Single-Instance Database with ASM	N/A	Module 1	Module 2	Module 2	Module 2
	Single-Instance Database with a Single-Instance Standby Database	Module 1	N/A	Module 2	Module 2	Module 2
	Oracle RAC Database	N/A	N/A	N/A	Module 1	Module 2
	Oracle RAC Database with a Single-Instance Standby	N/A	N/A	N/A	N/A	Module 2

As shown in Table 1, your specific situation determines which modules you must complete. The rest of this document is divided into the Grid Control preparation section and three modules,

with each module describing a different configuration scenario. A brief overview of each one follows:

- Prepare the Grid Control Environment

  This section sets up the environment, including installing Grid Control, provisioning an Oracle database home, and creating a single-instance database on the new server.

- Module 1: Creating a physical standby database from a single-instance database

  This module uses a Deployment Procedure to provision a second server with a single-instance database home and create an Oracle Data Guard physical standby database. The MAA example in this module uses the Add Standby Database wizard available in Grid Control.

- Module 2: Creating a physical standby database and converting it to an Oracle RAC configuration

  This module provisions a two-node cluster (with Oracle Clusterware, Oracle RAC, and Oracle ASM), uses the Add Standby Database wizard to create a single-instance physical standby database, and converts the physical standby database to Oracle RAC with no downtime or impact on the primary database during this phase of the process. When a maintenance window occurs, schedule a switchover to the newly created Oracle RAC database.

  Alternatively, you can use an existing single-instance physical standby database if multiple standby databases exist and application HA and DR requirements are still achieved during the conversion process. To use an existing physical standby database:

  1. Skip the step to create the standby database.
  2. Specify the newly installed Oracle RAC database home when converting the standby database to an Oracle RAC database.

  The Oracle RAC database home must be installed on the same server where the physical standby database is located.

See Also: Chapter 4, "High Availability Architecture and Solutions" in *Oracle Database High Availability Overview 11g Release 1*[11]

Prepare the Grid Control Environment

The following sections set up your Grid Control environment and prepare it for provisioning using the Grid Control Components feature and Deployment Procedures. The processes described in this paper require that you have already performed the following prerequisite tasks:

- Install and Set Up Grid Control
- Provision a Single-Instance Database Home

Install and Set Up Grid Control

To set up your Grid Control environment, you must perform the following one-time configuration activities:

- Install a database running Oracle Database 11*g* Release 1 to act as the Grid Control Repository[1]. Ensure your Oracle Database is running Oracle Database 11*g* Release 1, including the latest patch set release. See My Oracle Support (formerly Oracle*Metalink)* note 756388.1 for information about the latest patchset and recommended patches.

  The best practice is to use Oracle Universal Installer (OUI) to perform the initial installation of the Oracle Database 11*g* Release 1 software. Also, be sure to patch the installation with the latest patch set release and merge and patch bundles.

- Install the full release of Oracle Enterprise Manager Grid Control 10g Release 10.2.0.5 on the Oracle Database 11*g* Release 1 (11.1.0.7) server and patch it with the latest patch set release.

- See Note 872352.1 in My Oracle Support (formerly Oracle*Metalink*) for the latest support and required patch information related to this process.

  > **See Also:** "Installing and Configuring the Full Release, and then Patching" in *Oracle Enterprise Manager Grid Control Installation Guide* [5].

- Download and install a Management Agent on each server that is to be managed with the Grid Control console.

  You can download Oracle Management Agent either using the Download Agent Software application in the Grid Control console or from My Oracle Support (formerly Oracle*Metalink*) by means of the Grid Control console. For the latter option, you must set up My Oracle Support (formerly Oracle*Metalink*) credentials and proxies in Grid Control.

[1] To achieve full functionality, you must apply the relevant merge and bundle patches to the Oracle Database environment and the Grid Control environment.

See Also: *Oracle Enterprise Manager Grid Control Installation Guide* [5] and Chapters 3 and 4 in the *Oracle Enterprise Manager Advanced Configuration* [6] for these topics:

- "Downloading Management Agent Software Using Grid Control Console"
- "Installing Management Agent Using OUI"

- Set up a Software Library.

  For provisioning, you must perform a one-time activity of setting up a Software Library. Once configured, you can use the elements in the Software Library for any software-provisioning operation performed using the provisioning application.

  See Also: The "Using a Software Library" section in *Oracle Enterprise Manager Advanced Configuration* [6]

- Create Components

  Once the environment is ready, you can use the Grid Control Console to create components, directives, or images for deployment onto the target servers, as described in the "Install Software and Upload Components" section.

- Customize Copies of the Deployment Procedures

  Perform a one-time activity of creating copies of the Enterprise Manager supplied Deployment Procedures and customizing the deployment procedures for your environment, as described in the "Create Copies of Deployment Procedures" section.

Module 1: Create a Physical Standby Database

This module uses a Deployment Procedure to provision a second server with a single-instance database home and create an Oracle Data Guard physical standby database. The process can also be used to migrate a database to Oracle ASM managed storage. The MAA example in this module uses the Grid Control Add Standby Database wizard. The tasks in this module perform the following work:

Task 1: Prepare the Primary Single-Instance Database for Data Guard
Task 2: Provision an Oracle Home on the New Server
Task 3: Create a Physical Standby Database
Task 4: Perform a switchover (optional)

The tasks in this section assume that you have already created the gold images for the Components and copied and modified the Deployment Procedures as described in the Prepare the Grid Control Environment section. The tasks also assume there is a primary database (either an Oracle RAC database or a single-instance database) already running in the configuration.

The following table and Figure 1 explain the transitions in this module.

STATE	PRIMARY DATABASE	PHYSICAL STANDBY DATABASE
Initial	Single-instance database	None
Intermediate (before switchover)	Single-Instance database	Single-instance (Oracle ASM optional)
End (after optional switchover)	Single-Instance (Oracle ASM optional)	Single-instance database

Figure 1: Environment States for Module 1

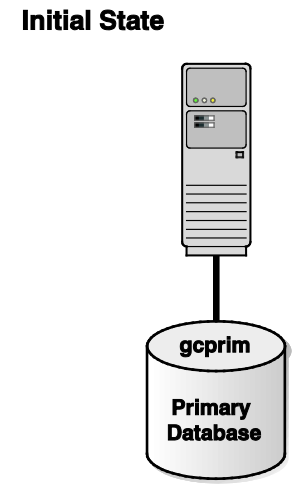

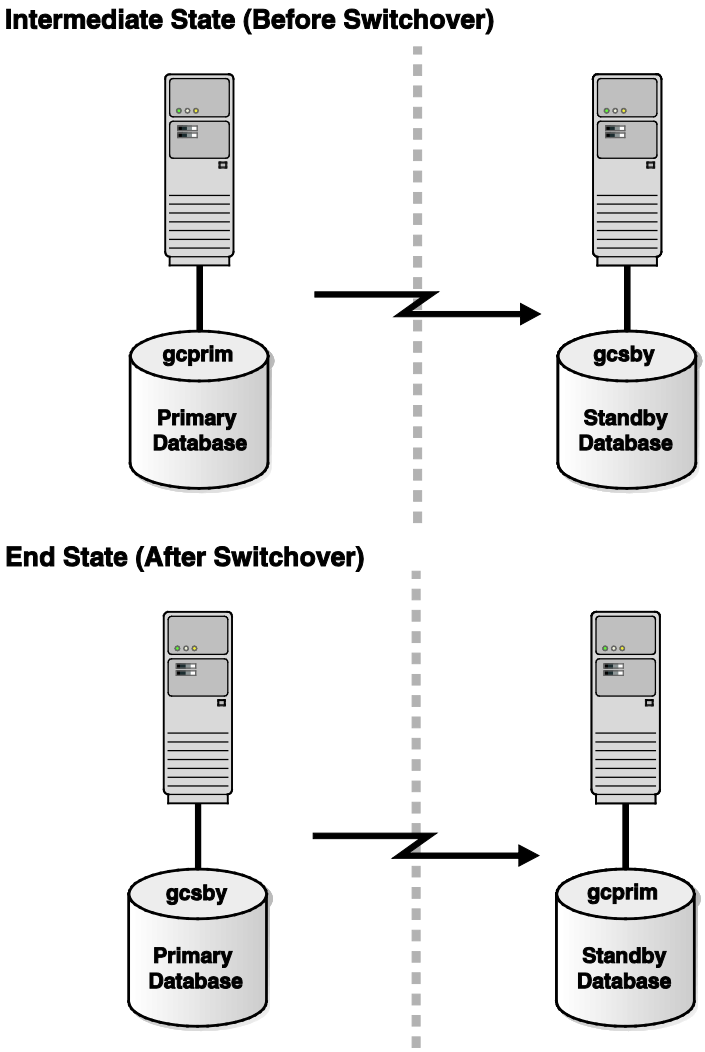

Task 1: Prepare the Primary Single-Instance Database for Data Guard

Configure the Oracle database for high availability following the MAA guidelines in *Oracle Database High Availability Best Practices* [3]. The best practices include enabling block checksums and block checking, fast-start fault recovery, a flash recovery area, and Flashback Database on the primary Management Repository database.

Task 2: Provision an Oracle Home on the New Server

1. Follow the steps in the Provisioning a Single-Instance Database Home section to create an Oracle database home on the server where the physical standby database will reside.

2. Optionally, you should provision an Oracle ASM home if you plan to migrate from file system storage to Oracle ASM shared storage and have not previously created an Oracle ASM instance.

You must provision a separate single instance database home for Oracle ASM using the same customized copy of the provisioning Deployment Procedure used to provision the single-instance database home (See the Create Copies of Deployment Procedures section to create a customized copy), and then manually create and start the Oracle ASM instance. After creation, you must configure Oracle ASM as a Grid Control target and provide Oracle ASM login credentials. Although the target is discovered during the execution of the deployment procedure, you must you provide the Oracle ASM credentials on the Management Agent page to make Oracle ASM a managed target in the Enterprise Manager console.

Task 3: Create a Physical Standby Database

This task uses the Enterprise Manager Console to create a physical standby database on the selected host server by using a backup of the primary database as the source. Perform the following steps to create a physical standby database using an existing backup of the primary database:

1. In Grid Control click **Targets**, and then click **Databases**.
2. On the Databases page, select the database that is to be used as the primary database and click **the database name** hyperlink. In the MAA example, the primary database is `gcprim`.
3. On the Database Instance page, click the **Availability** tab underneath the **Data Guard** heading and select **Add Standby Database** to invoke the Add Standby Database wizard.
4. In the "Add Standby Database" wizard dialog, respond as follows:
    a. On the first page of the wizard, select **Create a new physical standby database**.

       In the MAA example, the new standby database is `gcsby`.

    b. On the "Add Standby Database: Backup Type" page, select the method you want to use to create the standby, and choose to either perform a new backup (online or using a staging area) or use an existing backup. Then, click **Next**.

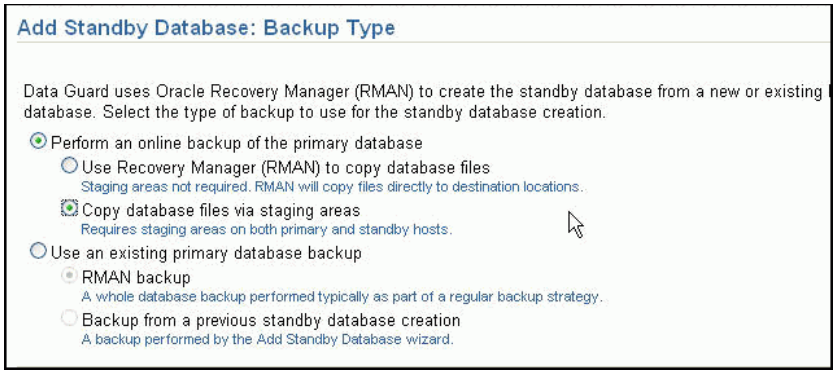

c. If you specified to use a staging area for the backup, on the "Add Standby Database: Backup Options" page, specify the staging area location on the primary site where RMAN will store the primary database backup files and the deletion or retention policy for this database backup.

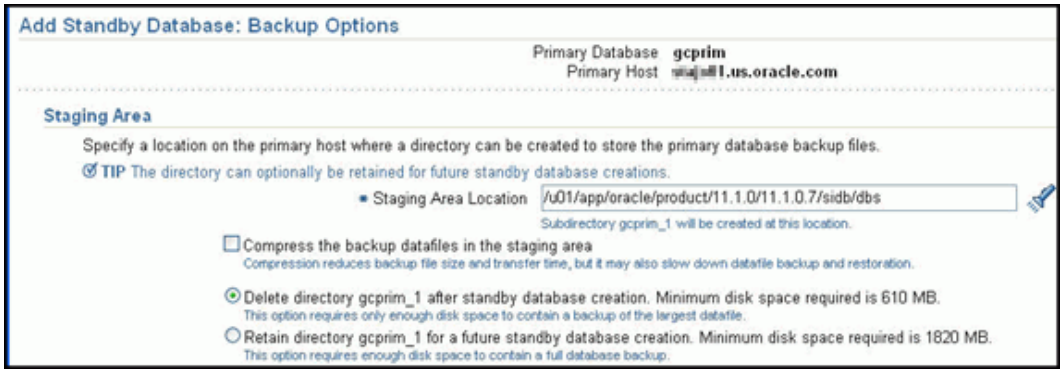

On this page, you can also specify the primary host credentials and whether you want to use Oracle Managed Files for the standby redo log files that are required by Data Guard. Click **Next**.

d. On the "Add Standby Database: Database Location" page, do the following:

- Specify the standby database SID name and storage type on the Standby Database Attributes section. You can specify a different Database Storage type for the standby database from what is being used on the primary database.

  See Also: *Oracle Enterprise Manager Concepts* for more information about how Enterprise Manager uses Oracle ASM.

- Click the torch icon next to the Host field to bring up a list of discovered Oracle homes on the standby database. Select the destination Host and Oracle home where you want the standby database to be created. Click **Select**.

  Click **Next**.

e. If you chose to configure Oracle ASM storage, then on the next page log into the Oracle ASM instance and click **Login**. This connects to the Oracle ASM instance. In the "Add Standby Database: File Locations" page, do the following:

   - In the Standby Host Backup **File** Access section, specify how the primary database backup files should be accessed on the standby database host.
   - For ASM storage, click the torch icon next to the "Configuration File Location" field to search for the listener configuration file in the Oracle ASM installation.

     Note: The best practice is to run the listener from the ASM Oracle Home when Oracle ASM is used.

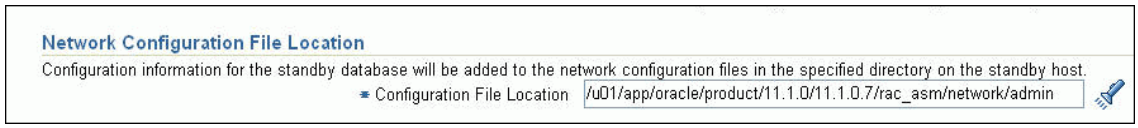

   - In the Standby Database File Locations section, select to use OFA or retain the primary database file names and locations for the standby database file locations.
   - Click the torch icon next to **Configuration File Location** field to specify the path where network configuration file information will be stored for the standby database.

     Click **Next**.

f. On the "Standby Database: Configuration" page, specify the configuration parameters (database unique name, Grid Control target name), the monitoring credentials (the username and password that will be used to access the database), and the Data Guard connect identifiers.

g. On the Review page, click **Finish** to start the job to create the physical standby database. The creation process runs as an Enterprise Manager job. The amount of execution time required for the job depends on the size of the primary database and the type of backup used to create the standby database.

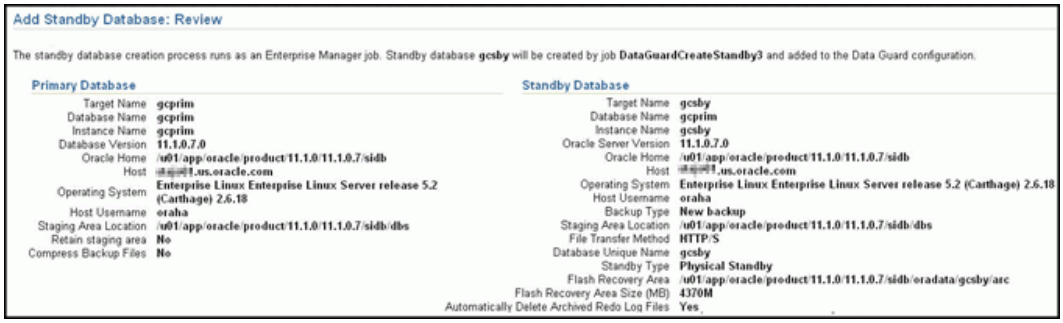

The standby database is created in the background, but you can monitor the progress on the Grid Control Job Activity Page by clicking **Job Activity Page** on the Grid Control **Jobs** tab and selecting the job name.

h. After the physical standby database creation job has completed, it is recommended to run **Verify Configuration** from the Data Guard page.

Task 4: Perform a Switchover (Optional)

Perform the following steps to switch over to the new standby database, if required (for example as a part of an Oracle ASM migration). It is a recommended best practice to switchover to verify the new configuration.

1. In Grid Control, click **Targets** and then click the **All Targets** subtab.
2. On the Data Guard page, select the Oracle physical standby database that you want to switch to the primary database role, and click **Switchover**.

   See Also: *Oracle Data Guard Broker 11g Release 1* Chapter 6 for more information about performing switchovers with Enterprise Manager.
3. After the switchover completes, the Data Guard page shows the database roles have been switched. If, as part of this process, you had created the physical standby database on Oracle ASM, your primary database will now be on Oracle ASM. Note that you could perform another switchover to return the databases to their original roles. In addition, you could now also migrate the original primary to Oracle ASM to provide resilient storage throughout your configuration.

Module 2: Convert a Single-Instance Physical Standby to Oracle RAC

The tasks in this module provision a two-node cluster, create an additional physical standby database that uses Oracle ASM or some type of shared storage on the new server, and convert the standby database to an Oracle RAC database. Finally, a switchover operation is performed to move the new Oracle RAC standby database to the primary database role.

> **Note:** The MAA recommended best practice is to add a standby database during this process even if there is already a standby database in place. This is to maintain the same level of protection from disasters during all phases of this process, with no compromise to the existing availability solution. Thus, if the starting configuration contained a standby database, the ending configuration would contain an Oracle RAC primary database and two single-instance standby databases. After the switchover is performed, either of the single-instance physical standby databases can be removed with no negative impact.

The tasks in this module perform the following work:

Task 1: Provision Oracle Clusterware, ASM, and Oracle RAC
Task 2: Create a Single Instance Physical Standby Database on the New Cluster
Task 3: Prepare the Environment Prior to Conversion to Oracle RAC
Task 4: Convert a Physical Standby Database to Oracle RAC
Task 5: Perform a switchover and Enable Additional Threads

The tasks in this module assume that you have already created the gold images for the Components and copied and modified the Deployment Procedures as described in the Prepare the Grid Control Environment section. The tasks also assume there is a primary database (either an Oracle RAC database or a single-instance database) already running in the configuration.

The following table shows the transitions in this module.

STATE	PRIMARY ROLE	STANDBY ROLE
Initial	Single-instance database	Single-instance database (already in place)
Intermediate 1 (create a new standby database for conversion to Oracle RAC)	Single-Instance database	Two single-instance databases
Intermediate 2 (after conversion to Oracle RAC, but before switchover)	Single-Instance database	An Oracle RAC standby database and a single-instance database
Intermediate 3 (after switchover)	Oracle RAC database	Two single-instance databases
End (optionally drop one standby database)	Oracle RAC database	Single-instance database

Figure 2 shows the initial configuration for this module.

Figure 2: Initial Environment State for Module 2

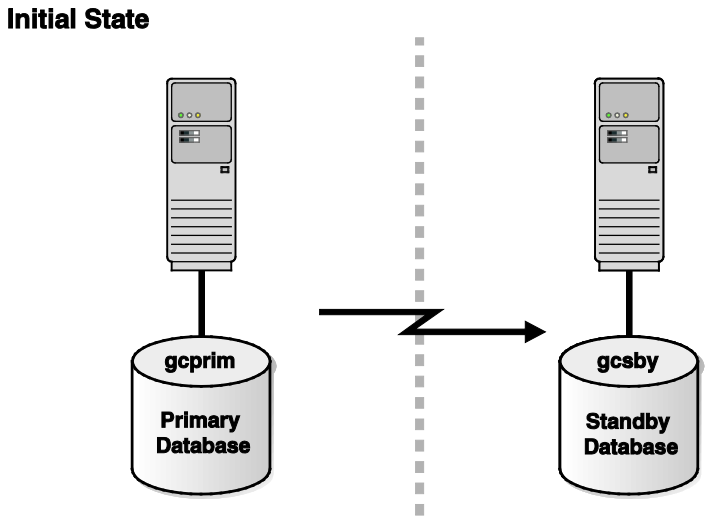

Task 1: Provision Oracle Clusterware, Oracle ASM, and Oracle RAC

This section describes using Deployment Procedures to provision a two-node cluster running Oracle Clusterware, Oracle ASM, and Oracle RAC. If you have already provisioned Oracle RAC software onto the new server, skip to Task 2 to create the physical standby database.

The process described in this section uses gold images of Oracle Clusterware, Oracle ASM, and Oracle RAC that were pre-created and stored in the Software Library.

Also, see "Provisioning Oracle RAC Using Gold Image" in the *Oracle Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching* for additional prerequisite information.

Perform the following steps:

1. In Grid Control, click the **Deployments** tab.
2. On the Deployments page, in the Deployment Procedure Manager section, click **RAC Provisioning Procedures**.
3. On the Deployment Procedure Manager page, in the Procedure subtab, select to run the Deployment Procedure for Oracle Clusterware and Oracle RAC that you created previously in the "Install Software and Upload Components" section.

   Note: Ensure the customized copy of the Deployment Procedure uses the appropriate commands (for example, `sudo`) for step execution.

   Click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Select Source page of the Deployment Procedure.

4. On the Select Source page, do the following:

   a. In the **Select Source** section, select **Select from Software Library**.

   b. In the Source for Clusterware section, click the torch icon and select the Component that has the gold image of Oracle Clusterware. For example:

Select	Component Name	Location	Vendor	Revision	Status	Product Version
◉	ClusterwareGoldImage	Components/ClusterwareGoldImage	Oracle		Active	Unknown

   c. In the Source for RAC section, click the torch icon and select the Component that has the gold image of Oracle RAC. For example:

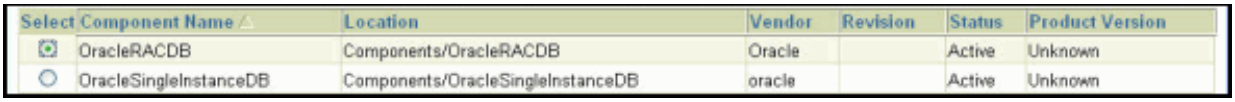

Select	Component Name	Location	Vendor	Revision	Status	Product Version
◉	OracleRACDB	Components/OracleRACDB	Oracle		Active	Unknown
○	OracleSingleInstanceDB	Components/OracleSingleInstanceDB	oracle		Active	Unknown

   d. In the Source for ASM section, choose whether or not you want to deploy Oracle ASM. The MAA team chose to deploy Oracle ASM using the same component for the ASM Oracle home as was used for the Oracle RAC Oracle home. For example:

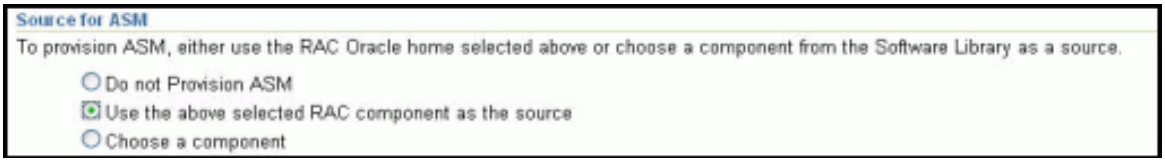

   Note: Ensure that you select only Components that are in "Active" status. Once you select the component name, the application automatically displays the component location.

5. On the Select Hosts page, perform the following:

   a. In the Hosts to Include in Cluster section, click **Add** and select the target hosts that should form the cluster.

   By default, the Show Suitable Hosts option is selected and the table lists only those hosts that are best suited for provisioning. If you do not find the host you want to add, then select **Show All Hosts** to view a complete list of hosts.

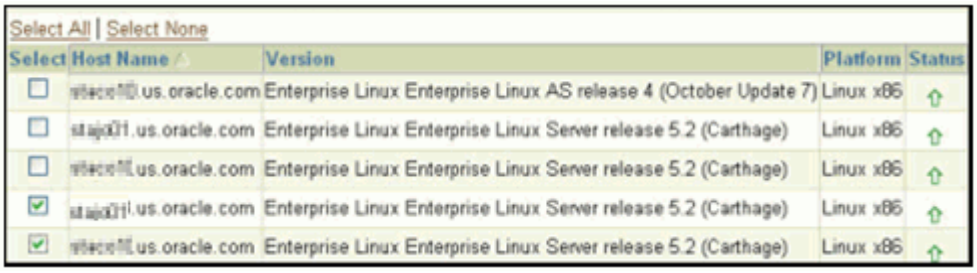

Select All | Select None

Select	Host Name	Version	Platform	Status
☐	[redacted].us.oracle.com	Enterprise Linux Enterprise Linux AS release 4 (October Update 7)	Linux x86	⇧
☐	[redacted].us.oracle.com	Enterprise Linux Enterprise Linux Server release 5.2 (Carthage)	Linux x86	⇧
☐	[redacted].us.oracle.com	Enterprise Linux Enterprise Linux Server release 5.2 (Carthage)	Linux x86	⇧
☑	[redacted].us.oracle.com	Enterprise Linux Enterprise Linux Server release 5.2 (Carthage)	Linux x86	⇧
☑	[redacted].us.oracle.com	Enterprise Linux Enterprise Linux Server release 5.2 (Carthage)	Linux x86	⇧

By default, the procedure automatically pre-fills the Private Host Name and Virtual Host Name fields with values. Ensure the correct Virtual Host Name for the VIP is used. If necessary, edit them to specify values that match your environment. Optionally, you can also specify IP addresses. For example:

b. In the Hosts to Include in Cluster section, configure the private and public network interfaces by clicking **Select Interfaces**. By default, the interfaces that have the same name and subnet for the selected target hosts are displayed. You can also choose Show Interface List to view all the interfaces for the selected target hosts. Select one of the existing interfaces or specify a completely new one. Click **OK**.

c. In the Network Interface Configuration section, review the details of the private and public interfaces.

Click **Next.**

6. On the Credentials/Schedule page:

   a. In the Hosts section, specify the Host Credentials (username and password that will be used to access the server during installation) to be used for the target.

      You can opt to retain the default selection (Preferred) so that the preferred credentials stored in the Management Repository are used, or override the preferred credentials to explicitly specify the host credentials.

   b. In the Schedule section, schedule the Deployment Procedure to run either immediately or later.

   c. Click **Next**.

7. On the Configure Cluster page:

   a. In the Cluster Name and Location section, review the default name and location details provided for Oracle Clusterware, Oracle RAC Database, and Oracle ASM and edit them if necessary. For example:

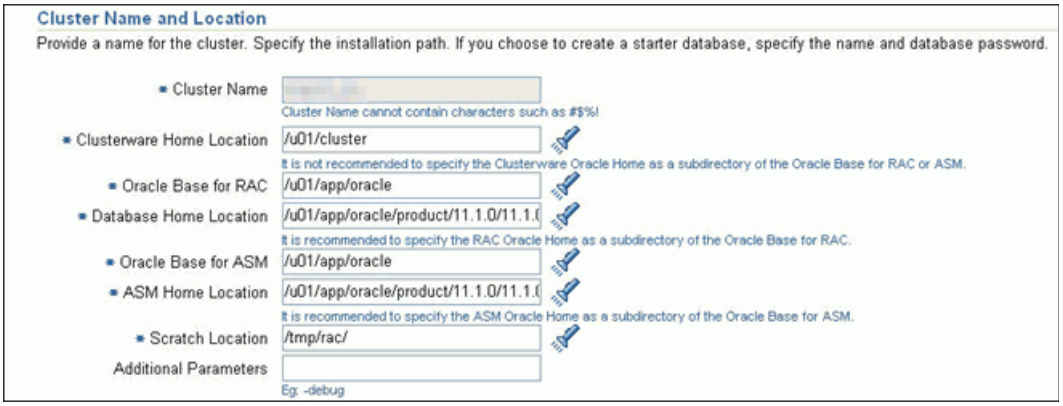

In this example, the default cluster name is based on the host cluster name you provided in the Agent Deploy application in Enterprise Manager Grid Control, while deploying Management Agents on a cluster. The scratch location is a temporary location on the target host where temporary files are placed before provisioning and configuring Oracle RAC.

For security purposes, the clusterware configuration sets the ownership of Oracle Clusterware home and all its parent directories to be owned by `root`. Hence, Oracle recommends that you install Oracle Clusterware outside the Oracle base of the Oracle RAC home.

b. In the Database Details section, by default, a starter database is created. However, because this environment will be used to create a standby database, no starter database is needed. Deselect the **Create Starter Database** checkbox.

c. In the ASM Instance Details section (that appears only if you had selected to deploy Oracle ASM), provide the password for the `SYS` user and the Oracle ASM disk string. For example:

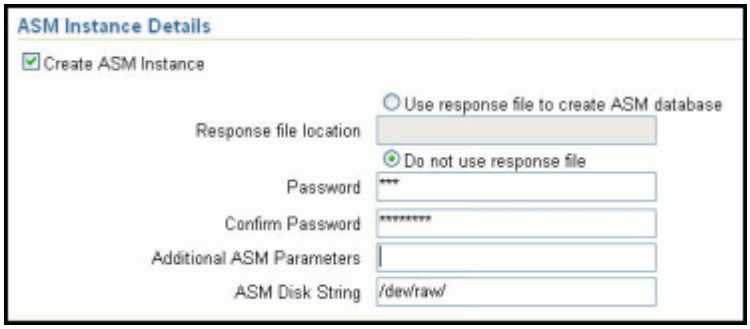

Click **Next**.

8. On the Storage page, do the following:

   a. On the Shared Storage Configuration section, specify locations in the partition name and the mount location fields for the Oracle Cluster Registry (OCR), voting disks, and data files. Also, select the mount format and a storage device for storing data. While partition name is the path to the location where the device is installed, mount location is the mount point that represents the partition location. For example:

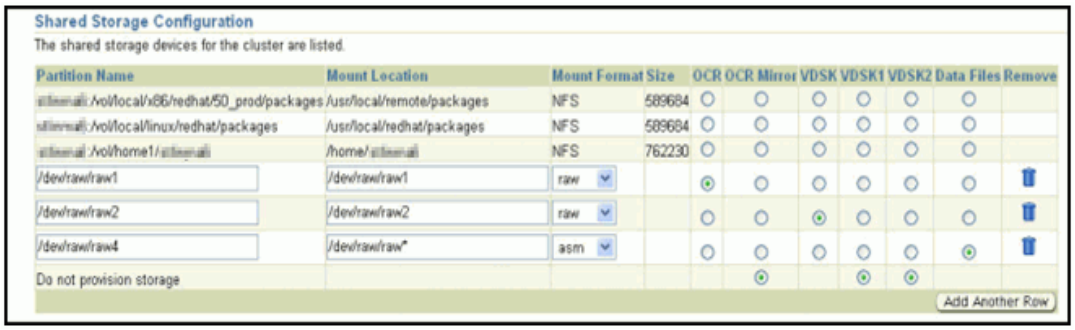

   If you are using raw devices, you can select the **Clear raw devices** checkbox under the table to clear the devices as a part of the installation. Doing so ensures the installation does not fail due to information remaining on the devices from previous installations.

   b. In the Advanced Options section, select a checkbox for ASM redundancy: None, Normal, or High.

   Click **Next**.

9. On the Advance Configuration page, do the following:

   a. In the Configure the Bonding Interface (Private Interconnect) section, if necessary. select **Configure Bonding Interface** to configure the bonding interface. For more information about the individual settings, click **Help** in the Enterprise Manager console.

   b. In the Sysctl File Configuration section, select **Configure Sysctl file** if you want to configure the `sysctl.conf` file. Specify the mode of editing the system configuration file and the location of the reference system configuration file used for modifying the kernel parameters. For more information about the individual settings, click **Help** in the Enterprise Manager console.

10. On the "Configure Oracle Home" page, you can optionally choose to install and initiate the configuration manager to receive security updates:

   - If the host where the database is being provisioned has a direct connection to the Internet, then specify an e-mail address and My Oracle Support password to install and initiate the configuration manager. An e-mail address is required so that security updates and install updates can be sent from My Oracle Support.

- If the host where the database is being provisioned has an indirect connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

  Click **Next**.

11. On the Review page, review the details you have provided for provisioning Oracle RAC and click **Finish** to submit the job to Enterprise Manager.

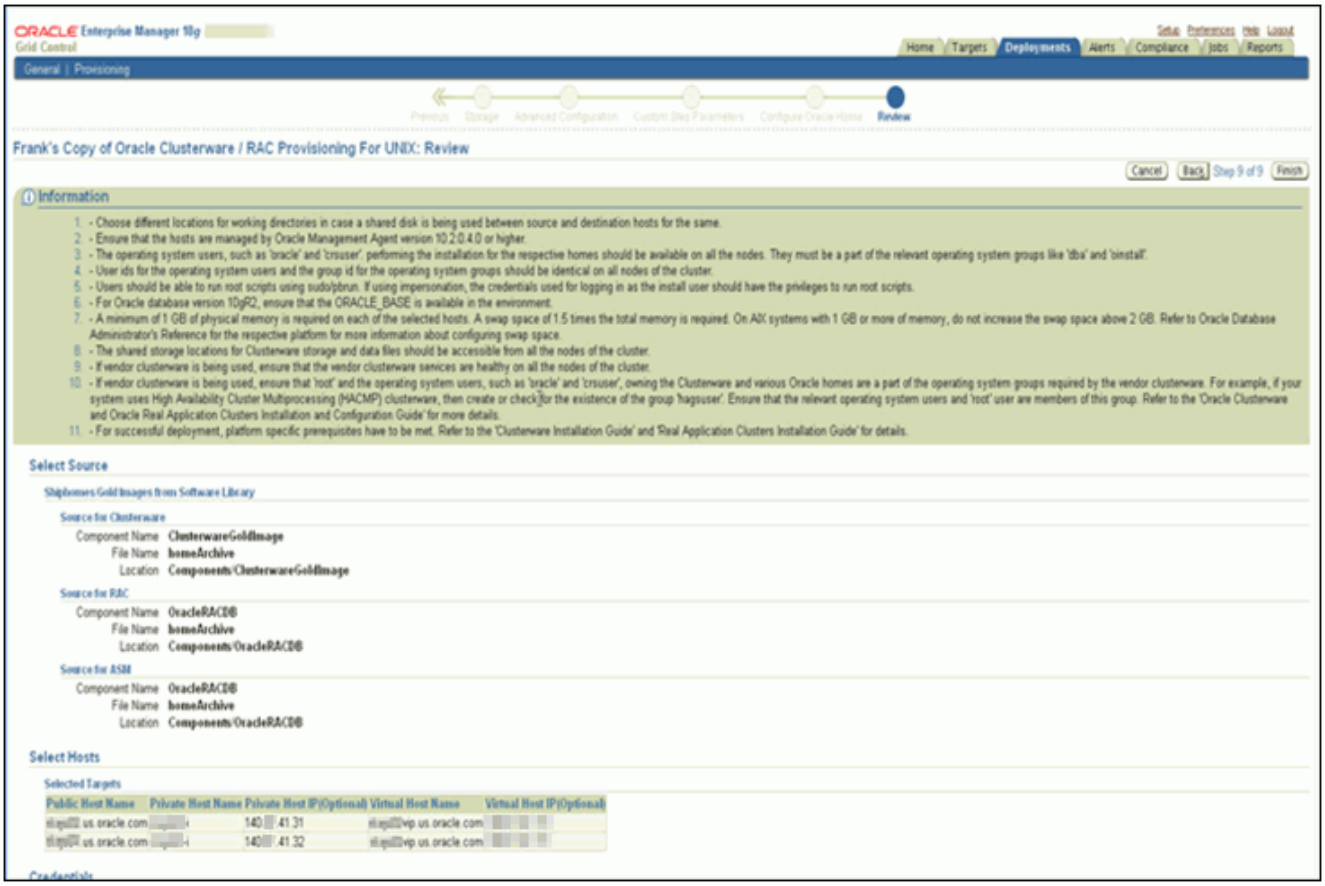

After the Deployment Procedure ends successfully you can verify the cluster configuration on the Target tab. In the following MAA example you can verify the clusterware home and the hosts included in the cluster:

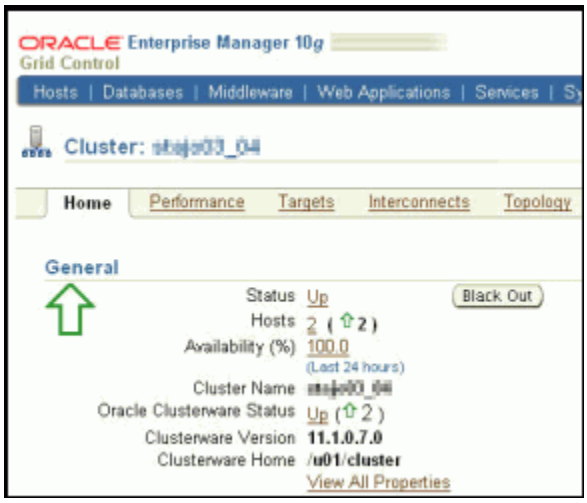

Task 2: Create a Single-Instance Physical Standby Database on the New Cluster

This module creates a new physical standby database using the Oracle RAC home on the new cluster. This new standby database will then be converted to Oracle RAC during Task 4. This module assumes any existing standby databases will be maintained during this process. During standby creation, create the database files on shared storage to facilitate conversion to Oracle RAC later. The MAA recommendation is to use Oracle ASM managed storage.

Follow the instructions in Module 1: Task 3 to create a single-instance physical standby database on some form of shared storage using the newly installed Oracle RAC home that you created in Module 2: Task 1. In this example, the new physical standby database is called `racsby`. The standby database is created on Oracle ASM managed storage.

Figure 3 Intermediate State 1 for Module 2

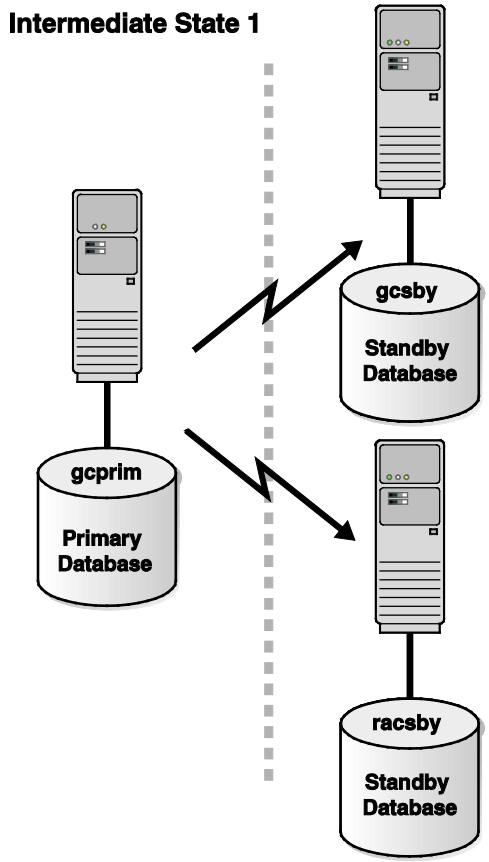

Task 3: Prepare the Environment Prior to Conversion

Prepare the environment prior to conversion to Oracle RAC by performing the following tasks:

1. Verify that the `STANDBY_FILE_MANAGEMENT` initialization parameter is set to `AUTO` on both the primary database and on the standby database that is to be converted.
2. Run the `$ORACLE_HOME/rdbms/admin/catclust.sql` script on the primary database to install the cluster database views into the environment.

   Note: This step is required only if the configuration does not already include a RAC database.
3. Manually create a second undo tablespace on the primary database to support the new database instance being created. The following SQL statement is an example:

```
create undo tablespace undotbs2 datafile
'/u01/app/oracle/oradata/gcprim/undotbs02.dbf' size 500M autoextend on
retention guarantee;
```

Note: This step is required only if there are not enough undo tablespaces already created to support the new Oracle RAC database. You must create an undo tablespace for each new database instance to be added. The Convert to Cluster wizard (executed in Task 4) will display an error if there are not enough undo tablespaces created.

4. Remove the standby database to be converted from the Data Guard broker configuration.

   Note: It is necessary to temporarily remove the new standby database from broker management in order to update the broker configuration file initialization parameters. The next step describes how to change the parameters to re-create the broker configuration files on Oracle ASM shared storage (instead of the current location on file system storage).

Perform the following steps to prepare the standby database for conversion to an Oracle RAC standby database. In the following examples, the standby database to be converted is called `racsby`.

1. In Grid Control, click the **Targets** tab, and then click the **Databases** subtab.
2. On the Databases page, select the primary database (`gcprim`) from the table.
3. On the `gcprim` database home page, click **Primary** in the High Availability section. For example:

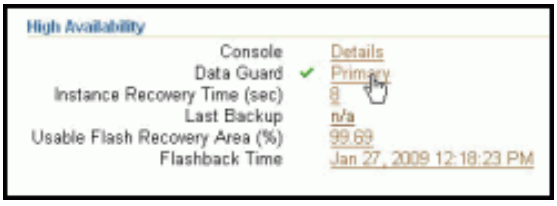

4. On the Data Guard page for `gcprim`, select the standby database that is to be converted (this is the `racsby` database in the MAA examples), and click **Remove**. For example:

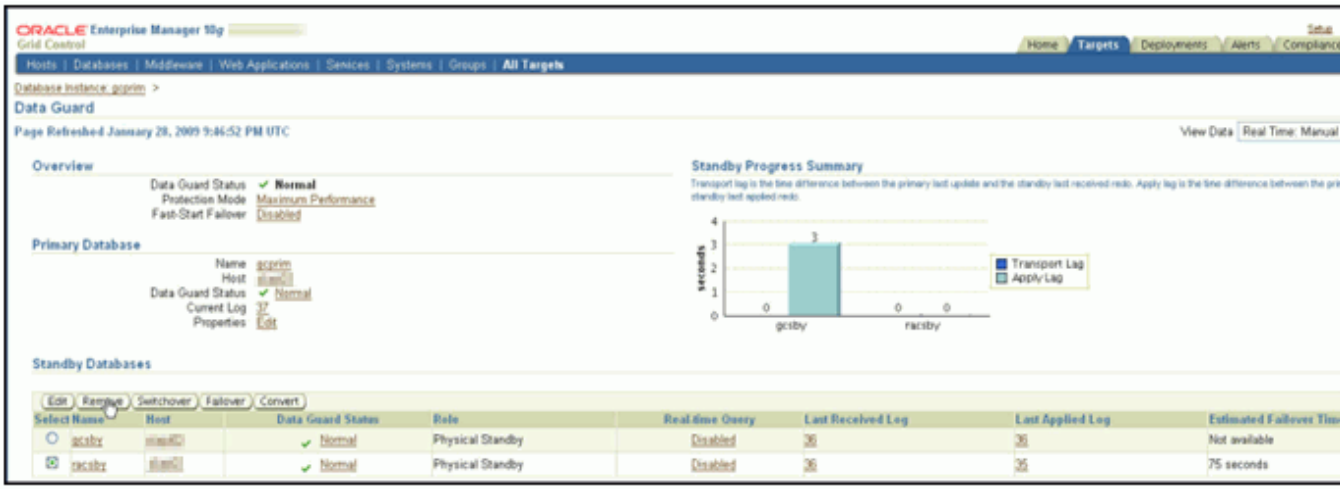

You should click the check box to "Preserve the destination corresponding to this standby" to continue shipping redo data to the standby.

The Remove Standby Database wizard asks you to confirm the removal and then proceeds to remove the racsby standby database from the broker configuration.

Note: Removing the standby database from the broker configuration only removes the entries from the broker configuration files. The removal does not delete the database from the Data Guard configuration. In fact, the current state of the physical standby database does not change.

5. Modify the following broker initialization parameters for the racsby standby database by using **Initialization Parameters** option on the Server tab:

   - The broker configuration files must be on shared storage when used in conjunction with a RAC database. Edit the `DG_BROKER_CONFIG_FILE1` and `DG_BROKER_CONFIG_FILE2` parameter values to set them to a location on Oracle ASM shared storage. The recommended Oracle ASM location is in the same diskgroup as the datafiles for the database. Note this directory must exist prior to adding the standby database back into the broker configuration.
   - Set the `DG_BROKER_START` initialization parameter to `FALSE` to disable the broker. This is necessary for the parameter changes for the broker configuration file locations to take effect.

   Click **Save to File**.

6. Add the `racsby` standby database back into the broker configuration. This step completes the action of relocating the broker configuration files into the new Oracle ASM shared storage locations.

   a. In Grid Control click **Targets**, and then click **Databases**.

b. On the Databases page, select the primary database (`gcprim` in this example).

c. On the Database Instance page for the primary database, click **Primary** in the list under the High Availability section of the page.

d. On the Data Guard page for the primary database, click **Add Standby Database** to invoke the Add Standby Database wizard.

> **Note:** The following steps automatically re-enable the broker management of the `racsby` standby database.

The following sequence of steps is similar to the process described in Module 1: Create a Physical Standby Database except you will be enabling broker management only for an existing standby database rather than creating a new standby. Respond to the wizard dialog as follows:

- On the Add Standby Database page, select "Manage an existing standby database" and click **Continue**.
- On the Select Existing Standby Database page, select the standby database from the table. In this example, this is the `racsby` database. Click **Next**.
- On the Configuration page, click **Next**.
- On the Review page, review the configuration settings and click **Finish** to add the standby database back into the broker configuration. The following screenshot shows the broker configuration running with the `gcprim` primary database and the two physical standby databases: `gcsby` and `racsby`.

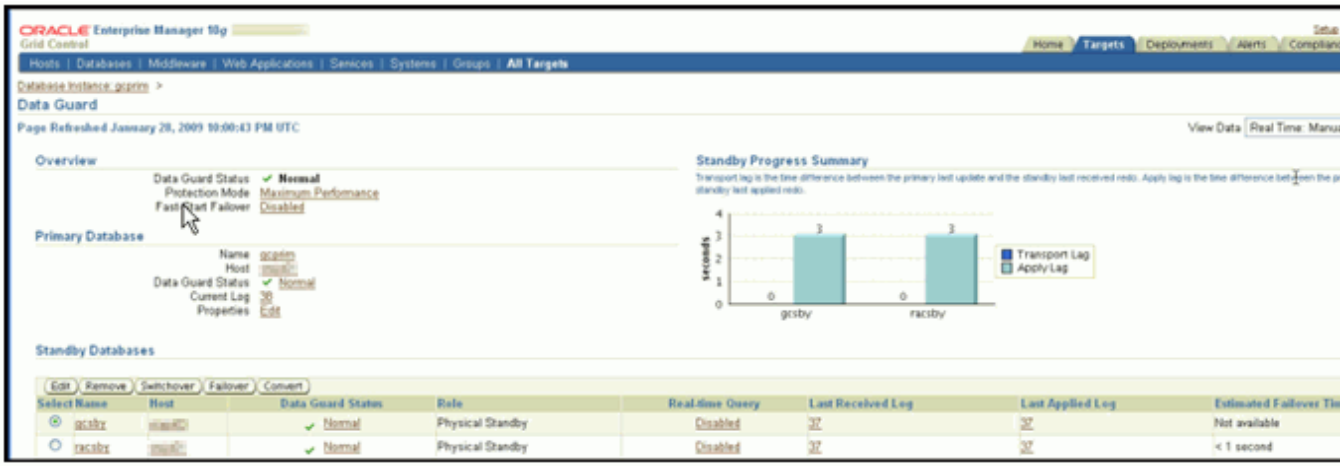

Task 4: Convert the Physical Standby Database to an Oracle RAC Database

This task converts the newly created physical standby database into an Oracle RAC database, with no downtime occurring to the primary database.

Figure 4 shows the state of the configuration after you perform the steps to convert the new `racsby` physical standby database to an Oracle RAC standby database. The configuration contains a single-instance primary database, a single-instance physical standby database, and an Oracle RAC physical standby database.

Figure 4: Intermediate State 2 for Module 2 After Conversion to Oracle RAC

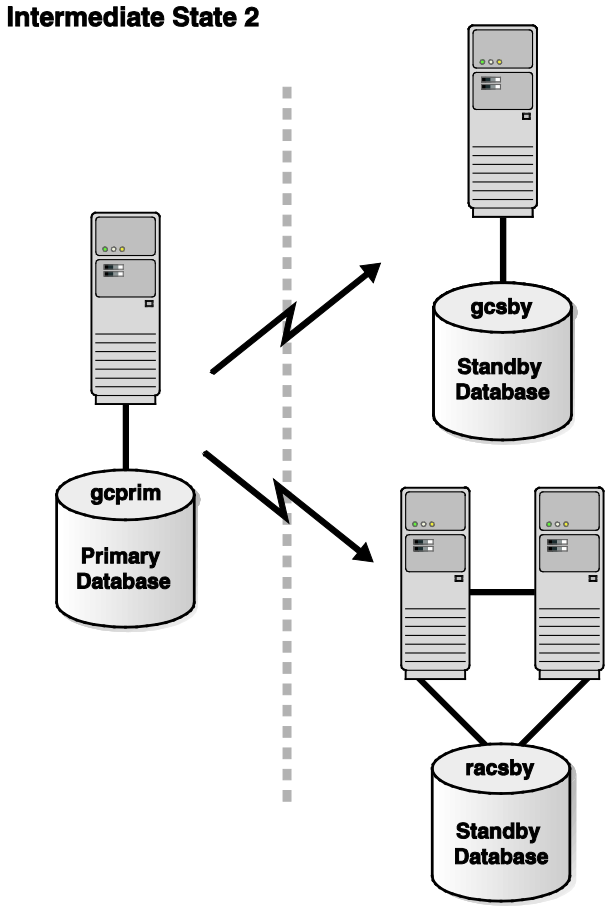

Perform the following steps to convert the `racsby` standby database to Oracle RAC:

1. In Grid Control click **Targets**, and then click **Databases**.
2. On the Databases page, select the standby database (`racsby` in our example) from the table and then on the standby database home page, click the **Server** subtab.
3. On the Server page in the Change Database section, click **Convert to Cluster**. The Convert to Cluster Database wizard starts.
4. Respond to the Convert to Cluster Database wizard, as follows:
   a. On the Cluster Credentials page, specify the credentials for the Oracle RAC and Oracle ASM homes. For example:

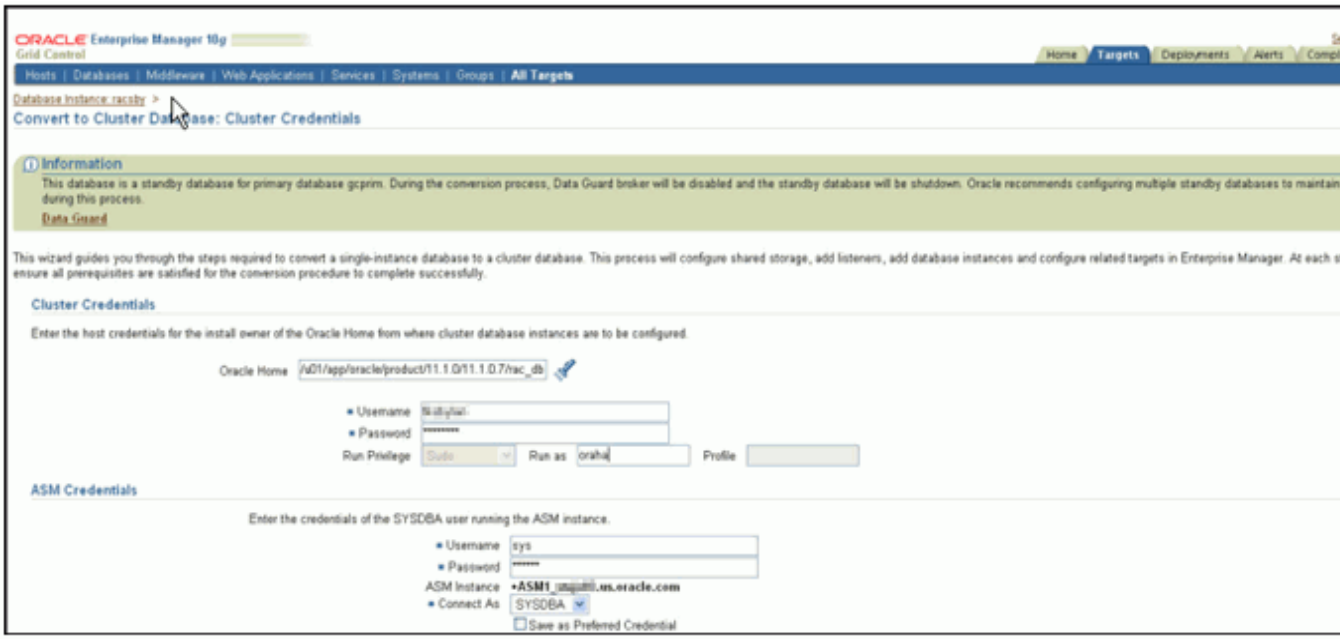

   Note: In the Information section of the page, you can see the wizard recognizes that the `racsby` database is a standby database for the `gcprim` primary database.

   On this page, only login credentials were needed to be supplied because the database was already configured to use Oracle RAC and Oracle ASM homes.

   Click **Next**.

   b. On the Hosts page, select the hosts from the table on which you want to run the converted Oracle RAC database. **Note:** The current host for the standby database is always selected.

   Click **Next**.

c. On the Options page, select to use either an existing listener or create a new listener, and specify a prefix to be used to name the cluster database instance (`ORACLE_SID`).

d. On the Shared Storage page, specify the data file and `FLASH_RECOVERY_AREA` storage locations. If the database is to be converted to Oracle RAC *in-place* (that is, the files are already located on shared storage), then you can use the existing locations. Otherwise specify the target disk groups for the data files and the `FLASH_RECOVERY_AREA`.

e. Review the settings and click **Submit** to run the Convert Cluster Database job in Enterprise Manager.

5. After the Convert Cluster Database job has completed successfully, remove the original (non Oracle RAC) standby database definition (racsby) from Grid Control.

   Warning: When you remove the old database from Grid Control, all the monitoring history is deleted. Remove a database only when this data is no longer needed.

   a. In Grid Control, click **Targets**.

   b. On the Databases page, notice that there are two entries for same standby database: one is for the original single-instance standby database and the other is for the new clustered standby database. Select the single-instance standby database definition from the table and click **Remove**. For example:

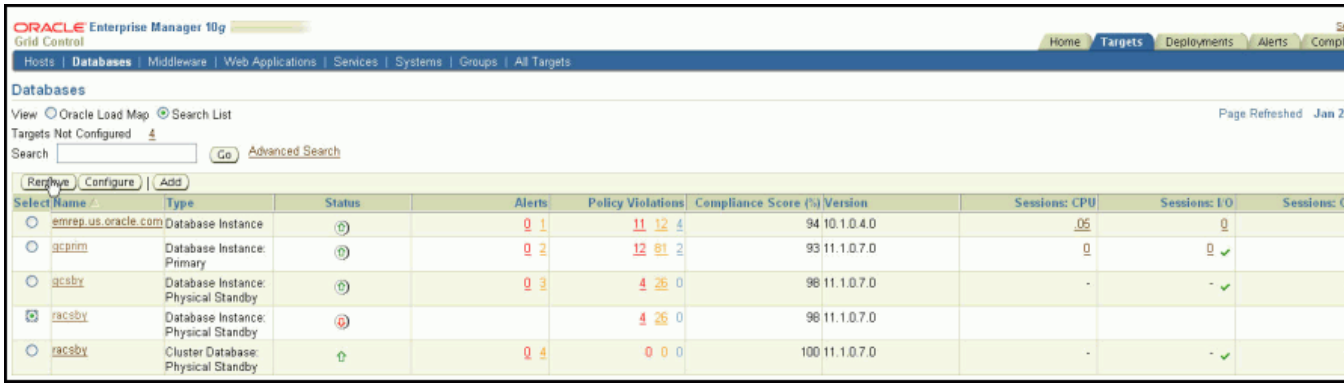

Task 5: Perform a Switchover and Enable Additional Threads

This section describes how to perform a switchover to complete the database conversion to Oracle RAC and to enable the new thread on the Oracle RAC database.

In the MAA example used in this section, the primary database is `gcprim` and has a single-instance standby database with the database unique name `gcsby`. The primary database is running on a file system. The new Oracle RAC database unique name is `racsby` and resides on

Oracle ASM. After the switchover, `racsby` is the Oracle RAC primary database, and `gcprim` is a single-instance physical standby database.

Perform a Switchover

Perform the following steps to switchover the newly converted Oracle RAC standby database (`racsby`) to run in the primary database role.

1. In Grid Control, click **Targets** and then click the **All Targets** subtab.
2. Select the `gcprim` database hyperlink.
3. Select **Details** in the High Availability section.
4. On the Data Guard page, select the Oracle RAC physical standby database that you want to switch to the primary database role, and click **Switchover**.

After the switchover completes, the Data Guard page shows the roles have been switched between `racsby` (now an Oracle RAC primary database), and `gcprim` (now a single-instance physical standby database).

Enable Additional Threads

Manually enable additional threads on the Oracle RAC primary database to make this a two-node Oracle RAC database. For example:

```
SQL> ALTER DATABASE ENABLE PUBLIC THREAD 2;
Database altered.
```

Note: Repeat this step for each additional thread added.

Figure 5 shows the configuration that results after you have completed the steps to perform a switchover and enable the additional database threads. The configuration contains an Oracle RAC primary database and two single-instance physical standby databases.

Figure 5: Intermediate State 3 for Module 2 after Switchover

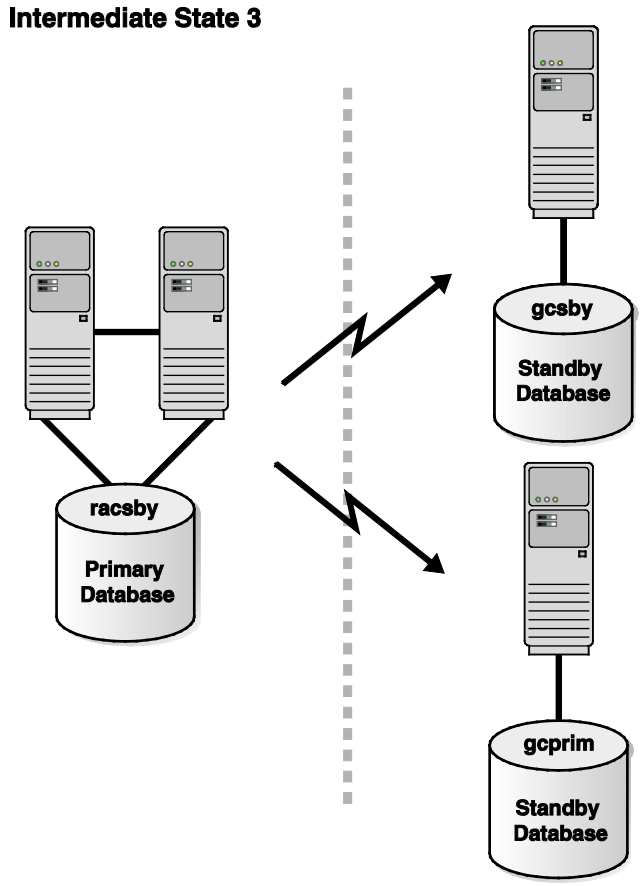

At this point, you can optionally remove the additional physical standby database. Figure 6 shows the ending configuration containing an Oracle RAC primary database and one single-instance physical standby database.

Figure 6: Ending State for Module 2 after Removing One Standby Database

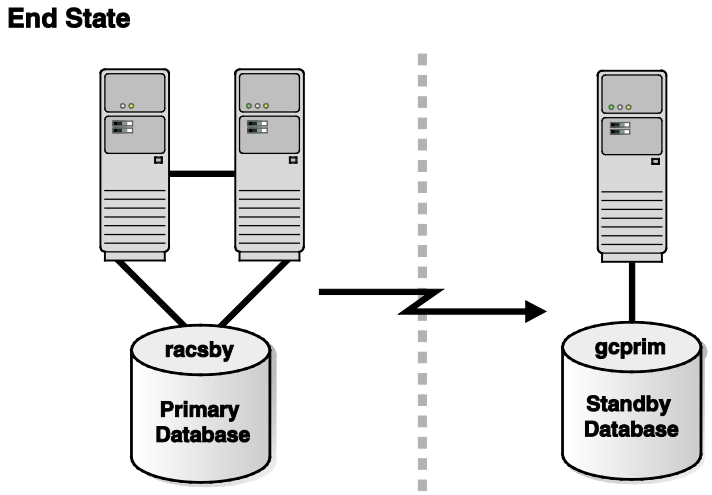

Appendix A: Create Components and Customize Deployment Procedures

This appendix describes how to create and upload Components for each type of software you plan to deploy, and customize deployment procedures for your needs. The tasks in this appendix are one-time activities.

Task 1: Install Software and Upload Components

Before deploying the configurations described in this white paper, you must have already configured Components for a single-instance Oracle Database 11*g* Release 1 (11.1.0.7), Oracle Clusterware, and Oracle RAC.

To install software and configure it as a Component for future installations, perform the following configuration activities for each type of software.

Note: Software that has been uploaded to the Grid Control software library cannot be modified directly in the library. To update software in the library, you must manually apply new patches or patchsets to a software installation, then re-upload the installation to the software library.

Manually Install Software

Use the OUI and OPatch to install, patch, and configure the software that you intend to deploy in your environment (such as a single-instance Oracle database home, Oracle Clusterware home, and Oracle RAC home). You will use these installations later as the source for future installations performed by Deployment Procedures. Note that configuring the Oracle software components will include applying any relevant patch sets or individual patches.

Create a Component for Each Type of Software

Upload each software installation as a "Component" into the Grid Control software library. The example used in the following steps shows how to create a component for an Oracle RAC database home.

Note: The components and configurations described in this white paper are provided as examples only. The actual Grid Control configurations that you deploy in your own environment will vary depending upon the needs of your organization.

1. In Grid Control, click **Deployments** and select Provisioning on the menu bar.
2. In the Components tab, select the folder under which you want to create the component and click **Create Component**. For example, select Components and Oracle Components.

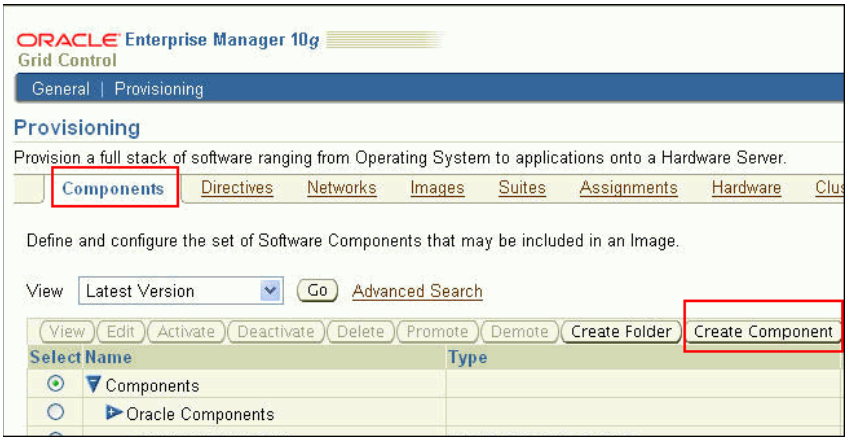

3. In the "Create Component: Describe" page, describe the component you are creating:

   a. Select the type of component you are creating from the list (Required)

   b. Name the component (Required)

   c. Provide a description of the component (Optional)

   d. Provide a product name/patch number (Optional)

   e. Supply the product version (Optional)

   f. Supply the vendor name (Optional)

   For example:

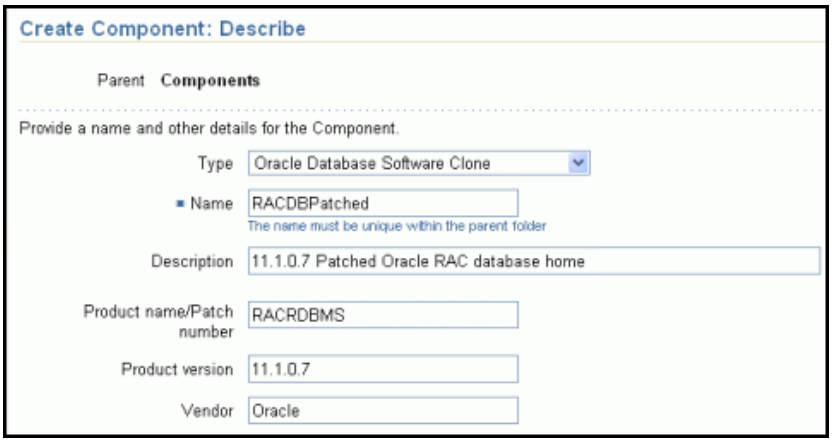

   Click **Next**.

   The dialog in the next steps of the Create Component wizard may vary depending on the type of component you have selected to create. Ensure that you associate the correct version of the component type when creating the component.

4. In the "Create Component: Configure" page, configure the Source for the Component.

   The source software must be installed on a host on which the Management Agent is already running. In our example to create a database software clone, the wizard asks you to:

   a. Select the Host where the software is currently installed and click **OK**.

      Note: Once you have selected the source Oracle home, the remaining fields are completed automatically by the wizard.

   b. Select the Home Location directory where the software is currently installed.

   c. Specify the Host Credentials to be used.

      You can opt to retain the default selection (Preferred) so that the preferred credentials stored in the Management Repository are used, or override the preferred credentials to explicitly specify the host credentials.

   d. Specify or select the Working Directory for file staging.

   e. List the log and trace files to be excluded.

5. In the "Create Component: Review" page, review all the information you have provided and click **Finish** to submit the component creation job.

6. Optionally, in the "Confirmation" page, click the Job Execution Id if you want to monitor the job. Note that the component has already been added to the Components list in the Software Library even though the job has not completed.

7. Go to the directory in the Software Library where you saved the component and verify that the Component has been created.

8. When the component upload has completed, activate the component

   a. Return to the Components page.

   b. Select the just uploaded Component.

   c. Click **Activate**.

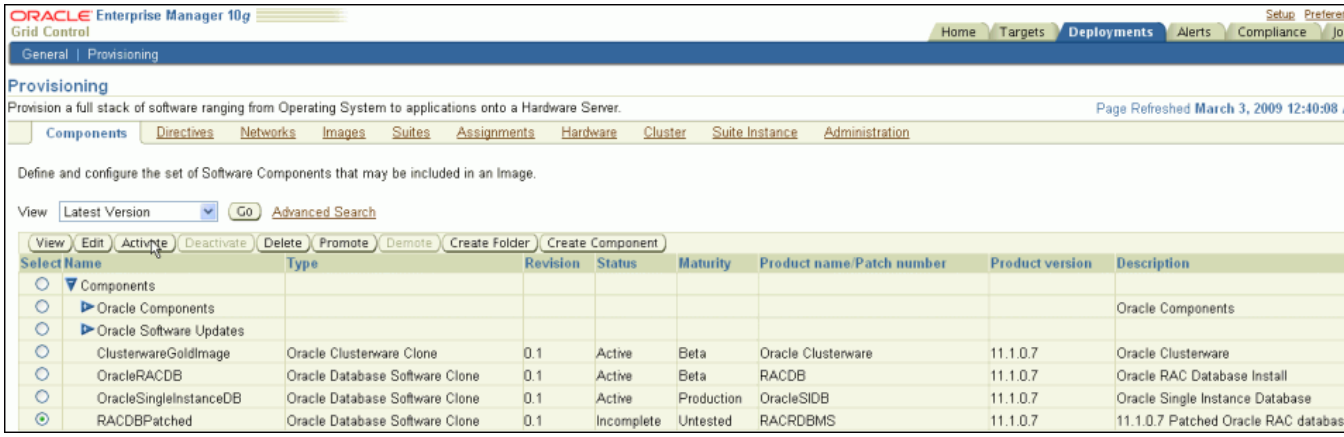

The component can now be used as a source in Deployment Procedures to provision other servers.

Note: Repeat the steps in this section for each type of software installation that you want to deploy with Grid Control.

Task 2: Create Copies of Deployment Procedures

The first step towards customizing a Deployment Procedure is to create a copy of the default Deployment Procedure that is supplied by Enterprise Manager Grid Control and customize the steps for your environment. You should edit the Deployment Procedures for customizations and to modify them for sudo credentials, if necessary, to perform operations as the software owner. While editing a Deployment Procedure, you can choose to run any step using SUDO. You can specify the sudo commands to run and also set environment variables and the preferred command interpreter for them.

The MAA team created copies of the original Deployment Procedures (for example, to create a single-instance database, to create a cluster, and to create an Oracle RAC database).

> **Note**: Only a copy can be edited and customized with your changes; the default Deployment Procedures cannot be directly edited. You should need to create only one copy of the Deployment Procedure and re-use it for subsequent executions. Deployment Procedure customization is therefore usually a one-time only process.

To create a copy of a default Deployment Procedure, follow these steps:

1. In Grid Control, click **Deployments**.
2. On the Deployments page, from the Deployment Procedure Manager section, click **Deployment Procedures**.
3. On the Deployment Procedure Manager page, select the Deployment Procedure you want to customize, and click **Create Like**.

4. On the Create Like Procedure page, edit the procedure to customize it according to your needs:

    a. Provide a new **Name** to personalize the copy

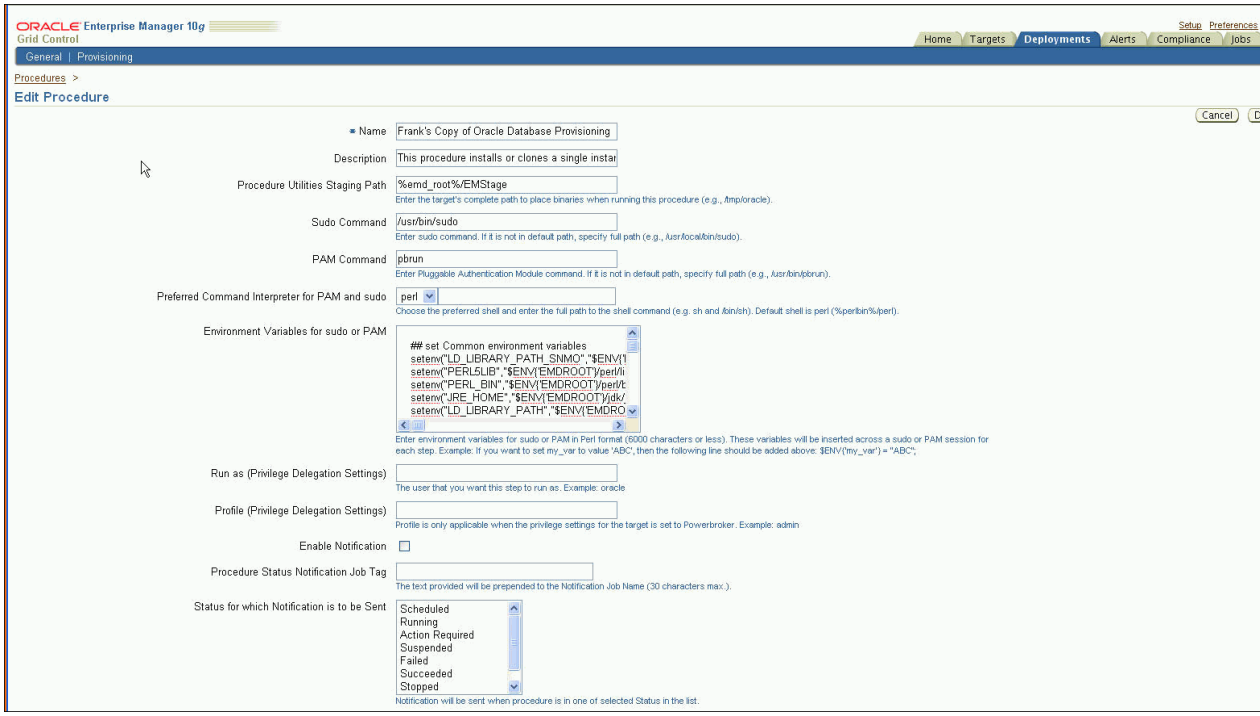

    b. All steps are enabled by default. You are allowed to **Disable, Delete** or **Edit** each step. You can **Enable** a step if it has previously been disabled. You can also **Insert** your own steps into the Deployment Procedure.

    c. Each step has a Run Privilege associated with it to determine how the step will execute.

        i. A Run Privilege of **Normal** will run under the USERID provided in the **Credentials** when scheduling/submitting the procedure. You would use **Normal** when the credentials you specify are for those of the software owner.

        ii. **sudo**, **PAM,** and **Privilege Delegation** are used for steps to be executed by a different or privileged user, for example when a step needs to be executed by the root user.

            Each step requiring Run Privileges must be modified individually. For example, if your software owner is a locked account and you have implemented sudo in your environment , every step would need to be modified to use the Run Privilege of sudo and the **Run**

Privilege Command/Privilege Delegation column will be filled with the command format. For steps to be executed by the software owner the Run Privilege command would contain something similar to:

/usr/bin/sudo –h oracle

Steps to be executed by root would contain something similar to:

/usr/bin/sudo –h root

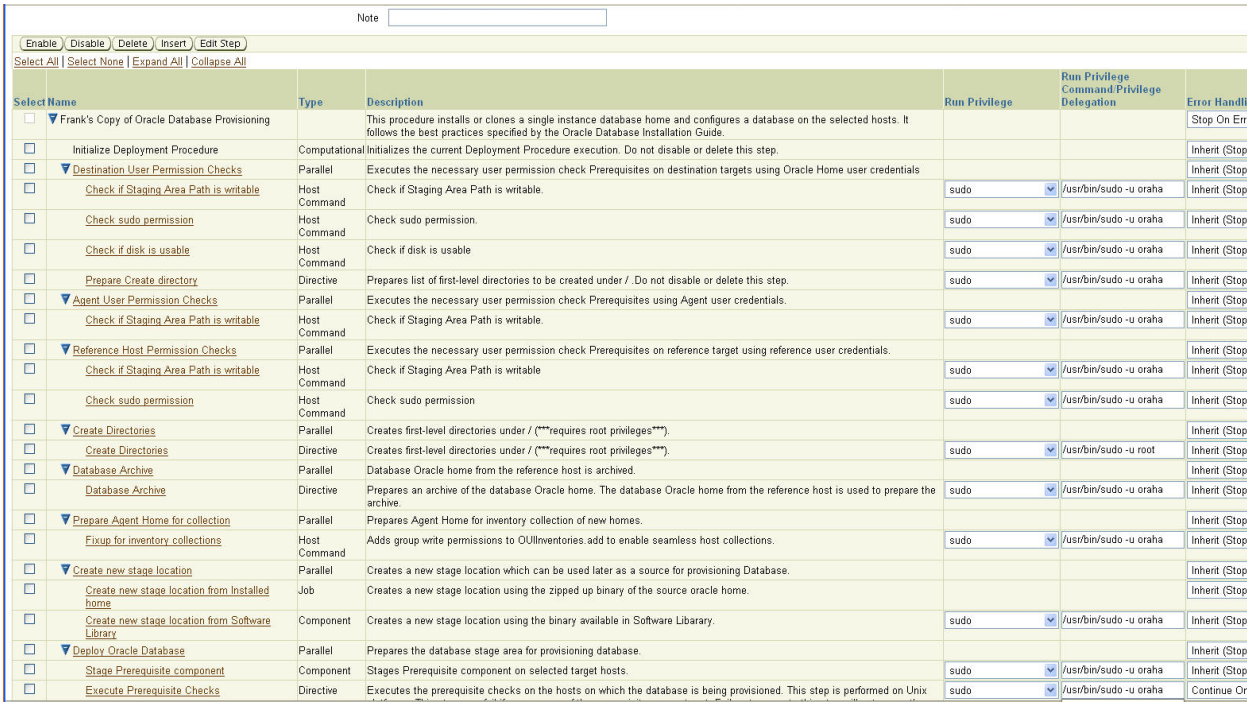
Note

Enable | Disable | Delete | Insert | Edit Step

Select All | Select None | Expand All | Collapse All

Select	Name	Type	Description	Run Privilege	Run Privilege Command/Privilege Delegation	Error Handli
	Frank's Copy of Oracle Database Provisioning		This procedure installs or clones a single instance database home and configures a database on the selected hosts. It follows the best practices specified by the Oracle Database Installation Guide.			Stop On Err
	Initialize Deployment Procedure	Computational	Initializes the current Deployment Procedure execution. Do not disable or delete this step.			Inherit (Stop
	Destination User Permission Checks	Parallel	Executes the necessary user permission check Prerequisites on destination targets using Oracle Home user credentials			Inherit (Stop
	Check if Staging Area Path is writable	Host Command	Check if Staging Area Path is writable.	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Check sudo permission	Host Command	Check sudo permission.	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Check if disk is usable	Host Command	Check if disk is usable	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Prepare Create directory	Directive	Prepares list of first-level directories to be created under / .Do not disable or delete this step.	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Agent User Permission Checks	Parallel	Executes the necessary user permission check Prerequisites using Agent user credentials.			Inherit (Stop
	Check if Staging Area Path is writable	Host Command	Check if Staging Area Path is writable.	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Reference Host Permission Checks	Parallel	Executes the necessary user permission check Prerequisites on reference target using reference user credentials.			Inherit (Stop
	Check if Staging Area Path is writable	Host Command	Check if Staging Area Path is writable	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Check sudo permission	Host Command	Check sudo permission	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Create Directories	Parallel	Creates first-level directories under / (***requires root privileges***).			Inherit (Stop
	Create Directories	Directive	Creates first-level directories under / (***requires root privileges***).	sudo	/usr/bin/sudo -u root	Inherit (Stop
	Database Archive	Parallel	Database Oracle home from the reference host is archived.			Inherit (Stop
	Database Archive	Directive	Prepares an archive of the database Oracle home. The database Oracle home from the reference host is used to prepare the archive.	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Prepare Agent Home for collection	Parallel	Prepares Agent Home for inventory collection of new homes.			Inherit (Stop
	Fixup for inventory collections	Host Command	Adds group write permissions to OUIInventories.add to enable seamless host collections.	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Create new stage location	Parallel	Creates a new stage location which can be used later as a source for provisioning Database.			Inherit (Stop
	Create new stage location from Installed home	Job	Creates a new stage location using the zipped up binary of the source oracle home.			Inherit (Stop
	Create new stage location from Software Library	Component	Creates a new stage location using the binary available in Software Libarary.	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Deploy Oracle Database	Parallel	Prepares the database stage area for provisioning database.			Inherit (Stop
	Stage Prerequisite component	Component	Stages Prerequisite component on selected target hosts.	sudo	/usr/bin/sudo -u oraha	Inherit (Stop
	Execute Prerequisite Checks	Directive	Executes the prerequisite checks on the hosts on which the database is being provisioned. This step is performed on Unix	sudo	/usr/bin/sudo -u oraha	Continue On

d. When all editing has been completed, click **Save.**

e. When using the 1-Click Extend deployment procedure to add a node to an existing RAC cluster, the 'Copy Packages' and 'Copy Archives' step will only work with the Privilege Delegation feature if working with a locked account

f. More information on Run Privileges can be found in Chapter 23 of Oracle® Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching.

5. View details of a Deployment Procedure.

To view the default configuration settings of a Deployment Procedure and the steps involved in it, follow these steps:

a. In Grid Control, click **Deployments**, and in the Deployment Procedure Manager section, click **Deployment Procedures**.

b. On the Deployment Procedure Manager page, in the Procedures tab, from the table, select the Deployment Procedure for which you want to view details, and click **View**.

Grid Control displays the View Procedure page that shows the default configuration settings and steps involved in the selected Deployment Procedure.

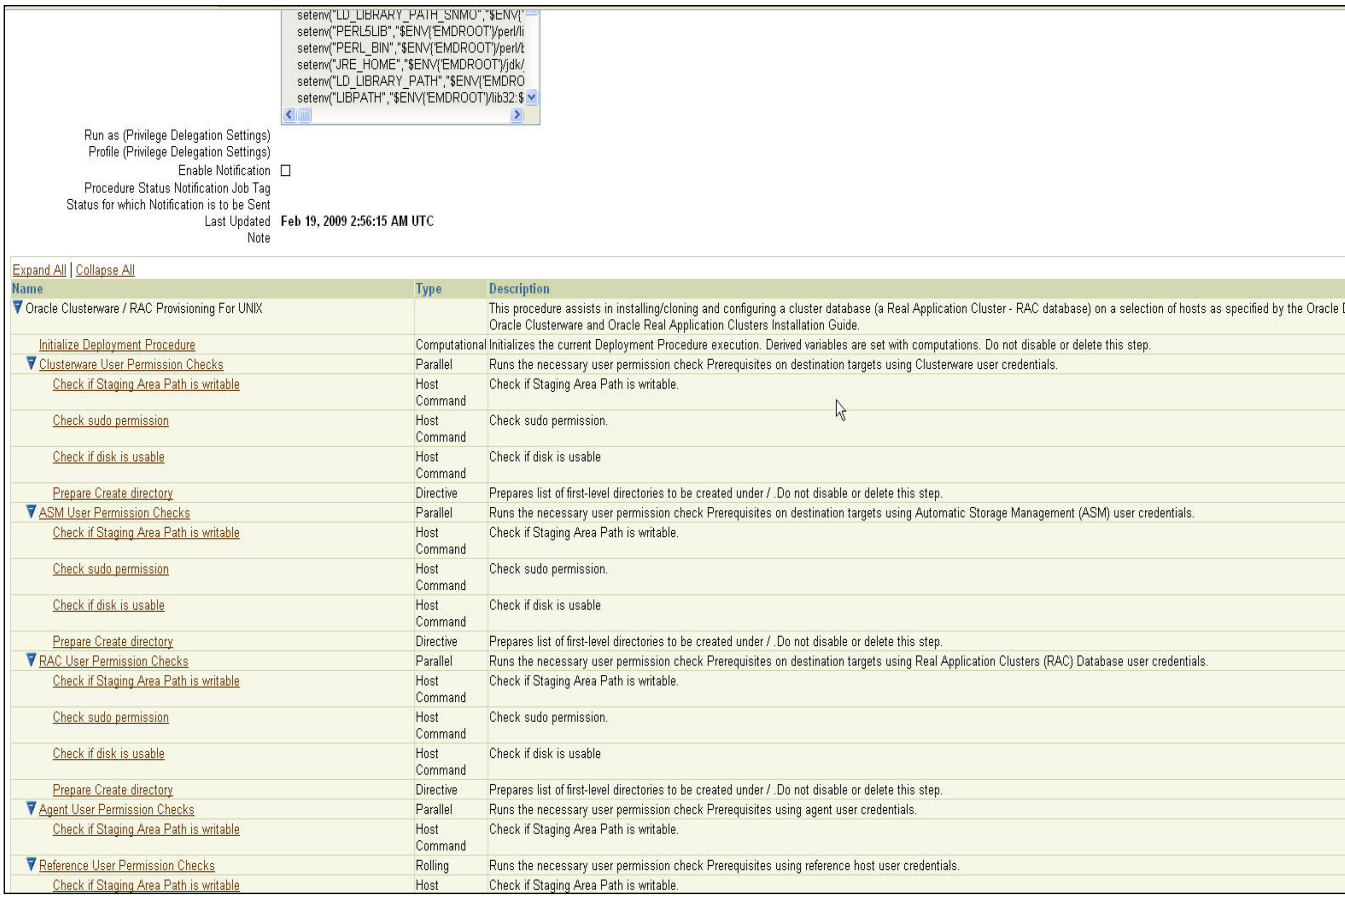

Appendix B: Converting a Primary Database to Oracle RAC

The MAA team chose to convert a physical standby database to Oracle RAC because this option incurs minimal downtime. Alternately, you can convert a primary database to an Oracle RAC database. However, this option can require additional downtime than the Convert Standby to RAC method.

The following tables shows the transitions performed in this appendix.

STATE	PRIMARY DATABASE	PHYSICAL STANDBY DATABASE
Initial	Single-instance database	Optional
End (after switchover)	Oracle RAC database	Optional

The "Convert to Cluster" wizard allows the conversion to be performed only onto the same hardware. In other words, the server on which the database currently resides must have Oracle Clusterware, Oracle RAC, and Oracle ASM shared storage available, as described in Module 2 Task 1: Provision Oracle Clusterware, Oracle ASM, and Oracle RAC.

To convert a primary database to Oracle RAC, perform the following steps:

1. In Grid Control click **Targets**, and then click **Databases**.
2. On the Databases page, select the primary database to be converted (`gcprim` in our example) from the table and then on the primary database home page, click the **Server** subtab.
3. On the Server page in the Change Database section, click **Convert to Cluster**. The Convert to Cluster Database wizard starts.
4. Respond to the Convert to Cluster Database wizard, as follows:
   a. On the Cluster Credentials page, modify the Oracle home field to point to the Oracle RAC home, and supply the host credentials and Oracle ASM instance credentials. For example:

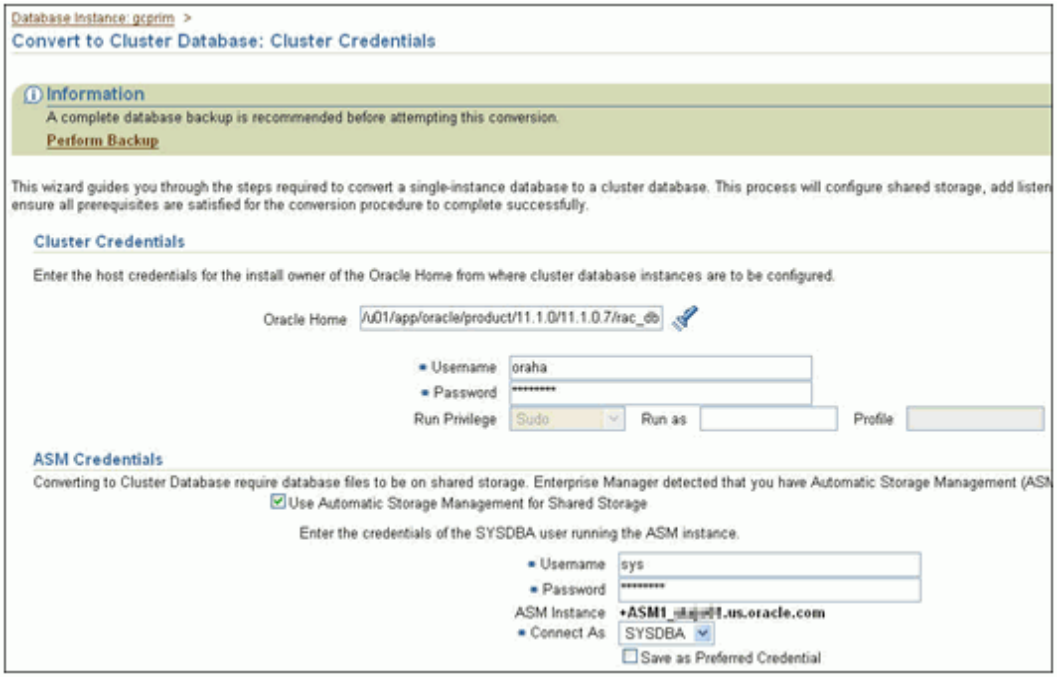

Note: In the Information section of the Cluster Credentials page, the message indicates the wizard recognizes that it is converting a single-instance database to Oracle RAC, not a standby database.

b. On the Hosts page, select the host from the table on which you want to run the converted Oracle RAC database. **Note:** The list will contain all servers in the cluster and the current host is selected by default.

   Click **Next**.

c. On the Options page, select to use either an existing listener or create a new listener, and specify a prefix to be used to name the cluster database instance (`ORACLE_SID`).

   Click **Next**.

d. On the Shared Storage page, specify the data file and `FLASH_RECOVERY_AREA` storage locations. If the database is to be converted to Oracle RAC in-place (that is, the files are already located on shared storage), then you can use the existing locations. Otherwise specify the target disk groups for the data files and the `FLASH_RECOVERY_AREA` (to move to shared storage).

   Click **Next**.

e. On the Review page, review the settings and click **Submit** to run the Convert Cluster Database job in Enterprise Manager.

The following screenshot shows the results of the MAA team's successful conversion.

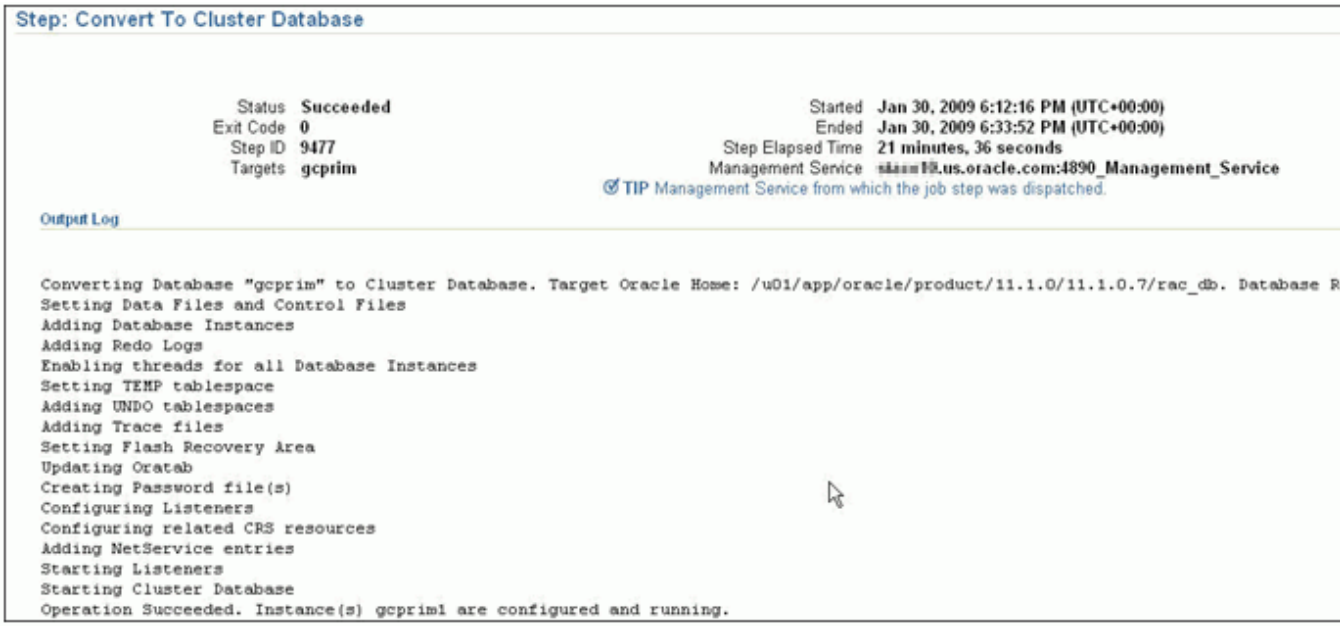

Upon successful completion of these steps, remove the original database target definition from Grid Control.

> **Warning:** When you remove the old database target definition, all log and metric data stored in OMS is deleted. Remove a database only when the data is no longer needed.

Appendix C: Extend an Oracle RAC Primary Database to Add Additional RAC Nodes

This module describes using a Deployment Procedure to extend the Oracle RAC primary database to more nodes in the cluster. Enterprise Manager Grid Control supplies the "One Click Extend Cluster Database" Deployment Procedure for adding Oracle RAC to other nodes.

The following table and Figure 7 show the state transitions that occur during this module.

STATE	PRIMARY DATABASE	PHYSICAL STANDBY DATABASE
Initial	Oracle RAC on *X* nodes	Optional
End (after switchover)	Oracle RAC on *X* plus *Y* additional nodes	Optional

Figure 7: Environment States for Module 3

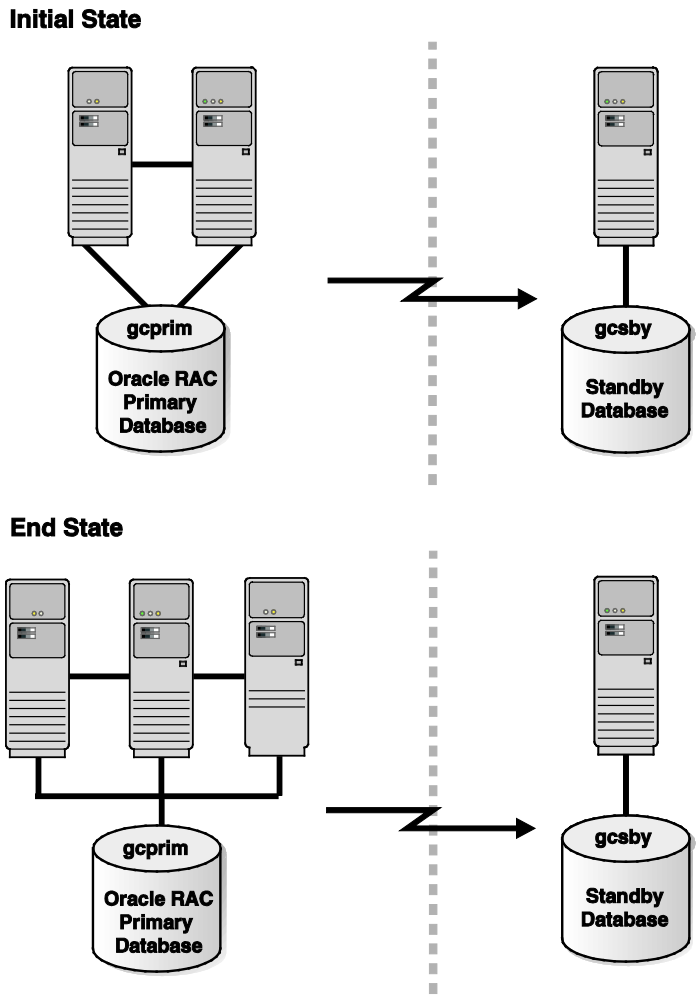

Task 1: Prerequisites

Ensure your environment meets the following prerequisites:

- Ensure that the credentials being used to run this operation along with the group ID are the same on all nodes of the selected cluster.
- Do not use an NIS-based operating system user.
- Ensure that you use an operating system user that has the privileges to run the Deployment Procedure and its commands on the target hosts.
- Ensure that the shared storage (used by the existing cluster nodes) is accessible to the nodes you want to add.

Task 2: Use the One Click Extend Cluster Database Deployment Procedure

This section shows how to extend an already existing cluster database. The provisioning process uses deployment procedures to complete the steps in this task with no downtime to the source database. It will copy, deploy, and configure everything necessary to deploy software and extend the cluster database on the new node.

To extend an existing Oracle RAC, perform the following steps:

1. In Grid Control, click the **Deployments** tab.
2. On the Deployments page, in the Deployment Procedure Manager section, click **RAC Provisioning Procedures**.
3. On the Deployment Procedure Manager page, in the Procedures subtab, from the table, select your customized copy of the **One Click Extend Cluster Database** Deployment Procedure (as described in the Create Copies of Deployment Procedures section).

   Click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Extend Real Application Clusters page.
4. On the Extend Real Application Clusters page, do the following:
   a. In the Select Real Application Clusters (RAC) section, select the Oracle RAC database you want to extend and specify the host credentials for the Oracle homes to install. Once you select a source cluster database, Grid Control determines what has been installed and deploys it across the cluster. All you need to do is provide host credentials for the source and target systems. The associated clusterware and Oracle ASM are also extended if they do not already exist on the cluster to which you are extending the Oracle RAC database. In the MAA example, `gcprim` is a single-node RAC database that we want to extend to a second node.

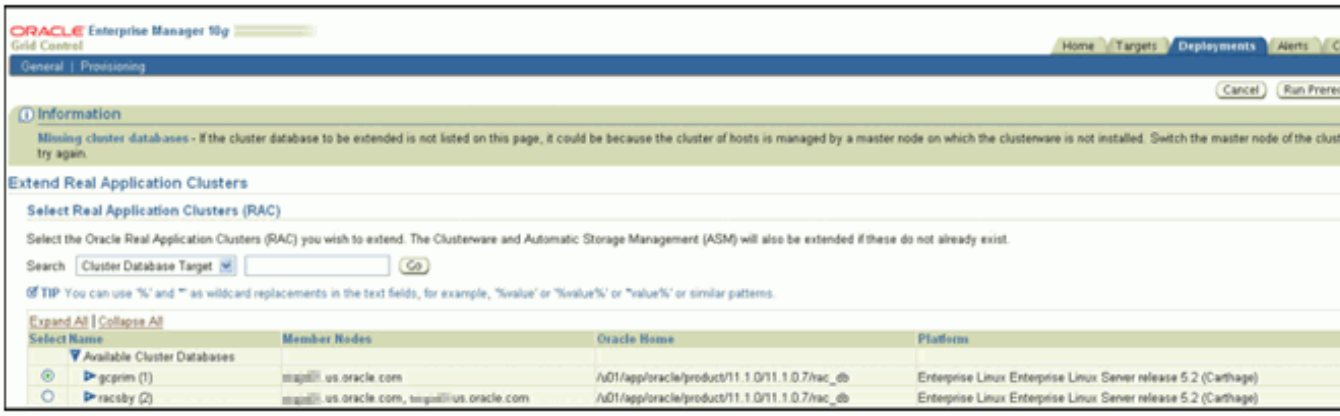

Tip: You can use the Search section to search for a particular Oracle RAC. From the Search list, select the target type based on which you want to search, and click **Go**. You can use wildcards such as % and *.

b. In the Select New Nodes section, click **Add** to add the nodes to which you want to extend the Oracle RAC configuration. After adding the node or nodes, specify the virtual node name for each new node and verify the values displayed by default.

c. In the User Credentials section, specify the credentials for the Oracle Clusterware, Oracle RAC, and Oracle ASM. You can opt to retain the default selection (Preferred) so that the preferred credentials stored in the Management Repository are used, or override the preferred credentials to explicitly specify the host credentials.

d. Review the settings and click **Submit**.

Appendix D: Provision a Single-Instance Database Home

Grid Control can be used to provision single-instance database homes and create databases to run in them. This appendix shows an example of using a Deployment Procedure to provision an Oracle home for a single-instance database and creates a database on the selected host server. The following steps assume that you have already created the source Components and stored the images in the gold software library (as described in the Install Software and Upload Components section), and prepared the Deployment Procedures (as described in Create Copies of the Deployment Procedures).

Perform the following steps to provision a single-instance database home:

1. In Grid Control, click the **Deployments** tab.
2. On the Deployments page, in the Deployment Procedures Manager section, click **Database Provisioning Procedures**.
3. On the Deployment Procedures Manager page, in the Procedures subtab, select your copy of the **Oracle Database Provisioning** deployment procedure from the list in the table. Then click **Schedule Deployment**.

   Enterprise Manager Grid Control displays the Select Source and Destination page.
4. On the "Oracle Database Provisioning: **Select Source and Destination**" page, select **Software Library** in the Select Source section. Then click the torch icon next to the **Component** field and select the generic component that has the gold image. Ensure that you select only components that are in "Active" status.

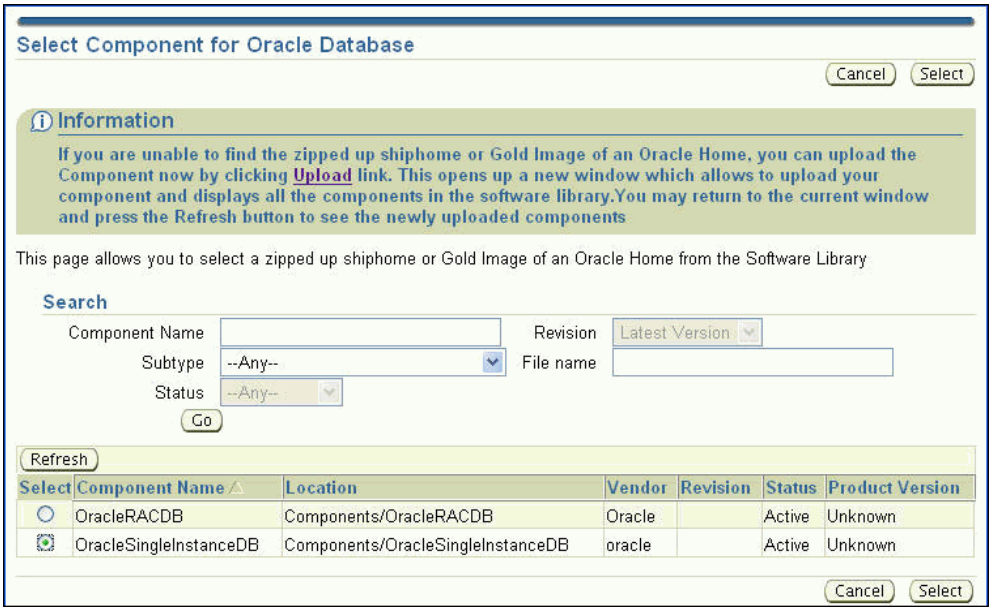

In our testing we selected the `OracleSingleInstanceDB` Component gold image created previously and stored in the Software Library.

Click **Select**.

5. On the "Oracle Database Provisioning: Select Source and Destination" page, do the following:

   - In the "Specify Destination Host Settings" section, click **Add** and select the target host on which you want to install the gold image of Oracle Database.

     By default, **Oracle Base**, **Oracle Home**, and **Working Directory** are pre-filled with sample values. Edit them to specify values that match your environment. If you specify directories that do not exist, then the Deployment Procedure creates them.

   - From the Credentials list, you can opt to retain the default selection (Preferred) so that the preferred credentials stored in the Management Repository are used, or Override preferred credentials to explicitly specify the host credentials.

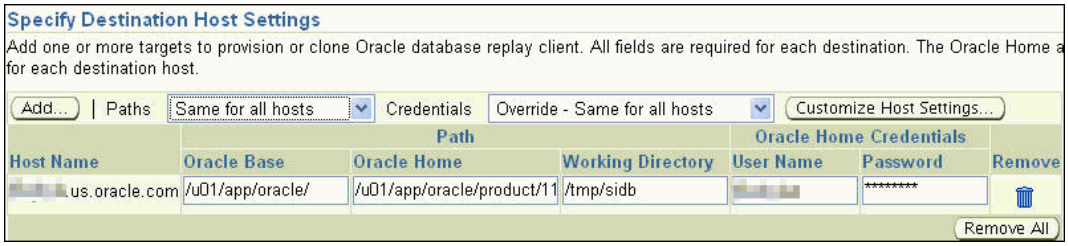

     Click **Next**.

   - In the "Create Database" section at the bottom of the page, select the check box to create the database as a step during provisioning. Alternatively, you can create the database later using Oracle Database Configuration Assistant (DBCA. In our test, we chose to manually create the database using DBCA later. Thus, the steps associated with the database creation will be skipped during the Deployment Procedure execution.

     Tip: Installing the database using DBCA is recommended to make sure that you inherit all Oracle installation best practices.

   Click **Next**.

6. On the "Configure Oracle Home" page, you can optionally choose to install and initiate the configuration manager to receive security updates:

   - If the host where the database is being provisioned has a direct connection to the Internet, then specify an e-mail address and My Oracle Support password to install and initiate the configuration manager. An e-mail address is required so that security updates and install updates can be sent from My Oracle Support.
   - If the host where the database is being provisioned has an indirect connection to the Internet through a proxy server, then specify an e-mail address and My Oracle Support password, and then in the Connection Details section, specify the proxy server details.

   Click **Next**.

7. On the Review page, review the details you have provided for provisioning a standalone Oracle Database, and click Finish. The Deployment Procedure starts running.

   Optionally, you can monitor the Deployment Procedure while it is running, as shown in the following screenshot.

Deploy Oracle Database	Succeeded	Parallel	Prepares the database stage area for provisioning database.		
Stage Prerequisite component	Succeeded	Component	Stages Prerequisite component on selected target hosts.	sudo	/usr/bin/sudo -u oraha
Execute Prerequisite Checks	Failed	Directive	Executes the prerequisite checks on the hosts on which the database is being provisioned. This step is performed on Unix platforms. This step may fail if one or more of the prerequisites are not met. Failure to execute this step will not cause the deployment procedure to fail.	sudo	/usr/bin/sudo -u oraha
Check if Fixes are Required	Succeeded	Computational	Check if the nodes need to be fixed.	Normal	
Execute System Fixes	Succeeded	Directive	Performs system level fixes based on the prerequisite check failures in the previous step. This step is performed on Unix operating systems. This step is skipped if the 'Execute Prerequisite Checks' step is disabled or if execution is successful. (*** requires sudo privileges as root ***)	sudo	/usr/bin/sudo -u root
Verify Prerequisite Checks	Succeeded	Directive	Executes the prerequisite checks on the nodes to provision the cluster. Any failure in this step indicate that one or more prerequisite checks have not been fixed automatically and need manual intervention. It is recommended that you fix the problem and then retry this step in the failed nodes. This step is performed on Unix operating systems. This step is skipped if the 'Execute Prerequisite Checks' step is disabled or if execution is successful.	sudo	/usr/bin/sudo -u oraha
Sysctl Configuration		Component	Modifies sysctl.conf by updating IPV4/IPV6 and kernel parameters (*** requires sudo privileges as root ***)	sudo	/usr/bin/sudo -u root
Copy Clone Database Archive	Skipped	Job	Copies the database archive from the reference host to the target hosts that require database Oracle home for clone option.		
Stage database Archive	Skipped	Directive	Stages the database archive to the database Oracle home location on target hosts.	sudo	/usr/bin/sudo -u oraha
Setup Database Stage Area	Succeeded	Component	Sets up the Stage Area for a database install from shiphome.	sudo	/usr/bin/sudo -u oraha
Verify Prerequisite Checks	Succeeded	Directive	Executes the prerequisite checks on the nodes to provision the cluster. Any failure in this step indicate that one or more prerequisite checks have not been fixed automatically and need manual intervention. It is recommended that you fix the problem and then retry this step in the failed nodes. This step is performed on Unix operating systems. This step is skipped if the 'Execute Prerequisite Checks' step is disabled or if execution is successful.	sudo	/usr/bin/sudo -u oraha
Execute pre-installation root scripts	Succeeded	Directive	Executes pre-installation root scripts as part of system preparation. This step is executed for specific platforms like AIX (***requires root privileges***).	sudo	/usr/bin/sudo -u root
Install/Clone Oracle Database	Succeeded	Directive	Lays down Oracle database software bits.	sudo	/usr/bin/sudo -u oraha
Attach home to central inventory		Directive	Registers the Oracle home with the central inventory. This is a special step for cases where environments are hardened and compilers and linkers are not available on production machines. Note: This step can only be used for UNIX as long as cloning is performed from the source Oracle home whose path is the same as that of the target. This step does not work for different paths.	sudo	/usr/bin/sudo -u oraha
Update Oracle Home Installation Time	Succeeded	Component	Updates Oracle home installation time with current time stamp.	sudo	/usr/bin/sudo -u oraha
Root scripts	Succeeded	Directive	Executes root scripts	sudo	/usr/bin/sudo -u root
Configure Oracle home	Succeeded	Job	Installs and initiates configuration manager for security updates	sudo	/usr/bin/sudo -u oraha
Stage Response File component	Skipped	Component	Stages response file component on selected target hosts.	sudo	/usr/bin/sudo -u oraha
Stage Template File component	Skipped	Component	Stages template file component on selected target hosts.	sudo	/usr/bin/sudo -u oraha
Copy template file	Skipped	Job	Copies the database template file from the reference host to the target hosts.		
Run Database Configuration Tools	Skipped	Component	Executes database config tools on the target hosts.	sudo	/usr/bin/sudo -u oraha
Create SPFile	Skipped	Directive	Create the SPfile for the database instance.	sudo	/usr/bin/sudo -u oraha
Database Verification Run	Skipped	Directive	Executes the database verification utility on each host.	sudo	/usr/bin/sudo -u oraha
Clean up Database Stage Area	Succeeded	Directive	Executes a script that cleans up the temporary database stage area that was set up for database install/clone.	sudo	/usr/bin/sudo -u oraha
Targets Discovery	Succeeded	Parallel	Performs a target discovery and registers new targets with Enterprise Manager.		
Register target	Succeeded	Directive	Registers DB target with agent.	sudo	/usr/bin/sudo -u oraha
Refresh Host Configuration	Succeeded	Job		Normal	
Targets Discovery	Skipped	Directive	Performs a target discovery and registers new targets with Enterprise Manager.	sudo	/usr/bin/sudo -u oraha
Register Database targets to Enterprise Manager	Skipped	Parallel			
Register Database targets to Enterprise Manager		Computational	Registers database targets to Enterprise Manager. It helps in displaying the database instances as up in Enterprise Manager, otherwise the provisioned database instances are displayed with Metric collection error status.	Normal	
Clean up source Database Stage Area	Skipped	Parallel	Cleans up the working directory on the reference host target using the source home credentials.		
Clean up source Database Stage Area		Directive		sudo	/usr/bin/sudo -u oraha

Delete | Schedule Deployment... | Stop | Suspend | Resume

In the example screenshot, note that a step in the procedure has failed although the Deployment Procedure continues to run:

- The step that failed checks prerequisites for the software installation, such as kernel settings, user limits, and permissions. If all prerequisites had been satisfied, the step would have completed with a Succeeded step status.

  Note: Every step in a Deployment Procedure is preconfigured with an error handling mode that indicates how the Deployment Procedure will behave when the phase or step encounters an error. In the far right-hand column for the failed step, the *Continue On Error* error handler allows processing to continue even if an error is encountered.

- The three steps after the failed step determine if changes should be made. If so, then the procedure will make the changes and ensure the changes complete successfully and confirm that the system is now ready for the installation.

Also, in this example, note that because the MAA team did not elect to create a database, the steps associated with the database creation were skipped.

8. After the Deployment Procedure ends, click **Done**.
9. If you did not create the database as a step in the Deployment Procedure, then create it now using DBCA.

References

1. Oracle Maximum Availability Architecture Web site
   http://www.oracle.com/goto/maa
2. *Oracle Database High Availability Overview (Part# B14210)*
   http://otn.oracle.com/pls/db111/db111.to_toc?partno=b28281
3. *Oracle Database High Availability Best Practices (Part# B25159)*
   http://otn.oracle.com/pls/db111/db111.to_toc?partno=b28282
4. *Oracle Database Storage Administrator's Guide 11g Release 1 (11.1)*
   http://otn.oracle.com/pls/db111/db111.to_toc?partno=b31107
5. *Oracle Enterprise Manager Grid Control Installation Guide* 10*g* Release 5 (10.2.0.5.0)
   http://download.oracle.com/docs/cd/B16240_01/doc/install.102/e10953/toc.htm
6. *Oracle Enterprise Manager Advanced Configuration 10g Release 5 (10.2.0.5)* at
   http://download.oracle.com/docs/cd/B16240_01/doc/em.102/e10954/toc.htm
7. *Oracle Enterprise Manager Concepts 10g Release 5 (10.2.0.5)*
   http://download.oracle.com/docs/cd/B16240_01/doc/em.102/b31949/toc.htm
8. "Using Enterprise Manager to Achieve Grid Automation With Deployment Procedures" white paper at
   http://www.oracle.com/technology/products/oem/pdf/grid-automation-deployment-procedures.pdf
9. *Oracle Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching 10g Release 5 (10.2.0.5.0)*
   http://download.oracle.com/docs/cd/B16240_01/doc/doc.102/e14500/toc.htm
10. Oracle Data Guard Broker
    http://otn.oracle.com/pls/db111/db111.to_toc?partno=b28295
11. *Oracle Database High Availability Overview 11g Release 1 (Part# B28281-03)* at
    http://download.oracle.com/docs/cd/B28359_01/server.111/b28281/toc.htm
12. *Oracle Data Guard Concepts and Administration 11g Release 1 (Part# B28294-03)* at
    http://download.oracle.com/docs/cd/B28359_01/server.111/b28294/toc.htm

Using Grid Control to Implement an MAA Environment on Oracle Database 11*g*
March 2011
Authors: Frank Kobylanski and James Viscusi
Contributors: Venkat Maddali, Bharat Paliwal, Vivian Schupmann

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

0109