Veritas Storage Foundation™ for Oracle® RAC Application Note: Installing or upgrading to Oracle RAC 11g Release 2

AIX

5.0 Maintenance Pack 3 Rolling Patch 4

Installing or upgrading to Oracle RAC 11g Release 2

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 5.0 MP3 RP4

Document version: 5.0MP3RP4.0

Legal Notice

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

■ A range of support options that give you the flexibility to select the right amount of service for any size organization

■ Telephone and/or Web-based support that provides rapid response and up-to-the-minute information

■ Upgrade assurance that delivers software upgrades

■ Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis

■ Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our Web site at the following URL:

www.symantec.com/business/support/index.jsp

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

■ Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Documentation feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

sfha_docs@symantec.com

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Installing or upgrading to Oracle RAC 11g Release 2

This document includes the following topics:

■ Installing Oracle RAC 11g Release 2

■ Upgrading to Oracle RAC 11g Release 2

Installing Oracle RAC 11g Release 2

This section provides instructions for installing Oracle RAC 11g Release 2.

Note: If you want to install Oracle RAC 10g or Oracle RAC 11g Release 1, see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration Guide (5.0 Maintenance Pack 3)* for instructions.

Perform the following steps to install Oracle RAC 11g Release 2:

1. Complete the preparatory tasks before starting the Oracle RAC installation.

   See "Preparing to install Oracle RAC manually" on page 8.

2. Install Oracle RAC 11g Release 2.

   See "Installing Oracle Grid Infrastructure using the Oracle Universal Installer" on page 22.

   See "Configuring LLT links in the GPnP profile" on page 25.

   See "Installing the Oracle RAC database using the Oracle Universal Installer" on page 26.

3. Complete the following post-installation tasks:

- Relink the SF Oracle RAC libraries with Oracle RAC:
  See "Linking the ODM library" on page 28.

- Configure the CSSD resource.
  See "Configuring the CSSD resource manually" on page 28.

- Modify the Oracle RAC configuration to prevent automatic startup of
  Oracle Clusterware.
  See "Preventing automatic startup of Oracle Clusterware" on page 30.

- Create the Oracle RAC database.
  See "Creating the Oracle RAC database" on page 30.

- Modify the Oracle RAC configuration to prevent automatic startup of the
  Oracle database.
  See "Preventing automatic database startup" on page 30.

- Configure VCS service groups for Oracle RAC.
  See "Configuring VCS service groups for Oracle RAC" on page 31.

Preparing to install Oracle RAC manually

This section provides instructions for performing the pre-installation tasks for
Oracle RAC 11g Release 2.

The following preparatory tasks are similar to those documented in the version
5.0 MP3 guide and are therefore not described in this section. For instructions,
see the *Veritas Storage Foundation for Oracle RAC Installation and Configuration
Guide (5.0 Maintenance Pack 3)*, Chapter "Preparing to install Oracle RAC".

- Identify the public virtual IP addresses for use by Oracle.

- Set the kernel parameters.

- Verify the user "nobody" exists.

- Create Oracle user and groups.
  In Oracle RAC 11g Release 2, you need to create the grid user and the user
  home directory on each system in addition to the Oracle user and group.

- Create storage for OCR and voting disk.
  The Oracle Clusterware files for Oracle Cluster Registry (OCR) and voting disk
  in Oracle RAC 11g Release 2 must reside in directories in the cluster file system
  (CFS) or on ASM disk groups. You can create ASM disk groups using CVM raw
  volumes.

- Configure private IP addresses for Oracle RAC.

Perform the following tasks using the instructions in this section:

- Identify the SCAN IP address.
  See "Identifying the SCAN IP address" on page 9.

- Create Oracle Grid Infrastructure and Oracle database home directories manually.
  See "Creating Oracle Grid Infrastructure and Oracle database home directories manually" on page 9.

- Set up user equivalence on cluster nodes.
  See "Setting up user equivalence" on page 19.

- Synchronize time settings on cluster nodes.
  See "Synchronizing time settings on cluster nodes" on page 19.

- Edit the user profiles.
  See "Editing the user profiles" on page 19.

- Verify whether the Veritas Membership library is linked to the Oracle library.
  See "Verifying whether the Veritas Membership library is linked to Oracle libraries" on page 21.

Identifying the SCAN IP address

You need a public IP address available and configured as a Single Client Access Name (SCAN) address on the domain name server (DNS) for round robin resolution to three addresses (recommended) or at least one address. SCAN is used to connect to databases within the cluster irrespective of the nodes they are running on.

Creating Oracle Grid Infrastructure and Oracle database home directories manually

You can create the Oracle directories on the local file system, on a local Veritas file system, or on a Veritas cluster file system.

Note: Symantec and Oracle recommend that you install the Oracle Grid Infrastructure and Oracle database binaries local to each node in the cluster.

Table 1-1 lists the Oracle directories you need to create:

Table 1-1   List of directories

Directory	Description
Oracle Grid Infrastructure Home Directory (GRID_HOME)	The path to the home directory that stores the Oracle Grid Infrastructure binaries. The Oracle Universal Installer (OUI) installs Oracle Clusterware and Oracle ASM into this directory, also referred to as GRID_HOME.
	The directory must be owned by the installation owner of Oracle Grid Infrastructure (oracle or grid), with the permission set to 775.
	The path to the Grid home directory must be the same on all nodes. As the root user, create a path compliant with Oracle Optimal Flexible Architecture (OFA) guidelines, so that the OUI can select that directory during installation.
Oracle base directory (ORACLE_BASE)	The base directory that contains all the Oracle installations. Create separate Oracle base directories for the grid user and the Oracle user.
	It helps to ensure that installations of multiple databases maintain an Optimal Flexible Architecture (OFA) configuration.
	The path to the Oracle base directory must be the same on all nodes.
Oracle home directory (ORACLE_HOME)	The directory in which the Oracle RAC software is installed. The path to the Oracle home directory must be the same on all nodes.
	You can create the Oracle home directory on the local file system or on a cluster file system.

Use one of the following options to create the directories:

Local file system	See "To create the directories on the local file system" on page 11.
Veritas File System	See "To create the file system and directories on the Veritas File system (VxFS) " on page 11.
Cluster File System	See "To create the file system and directories on cluster file system for Oracle Grid Infrastructure and Oracle database" on page 14.

To create the directories on the local file system

1 Log in as the root user on each node.

2 Create the Oracle base directory (ORACLE_BASE):

For grid user:

```
# mkdir -p /u02/app/grid
# chown -R grid:oinstall /u02/app/grid
# chmod -R 775 /u02/app/grid
```

For Oracle user:

```
# mkdir -p /u01/app/oracle
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

3 Create the Oracle Grid Infrastructure home directory (GRID_HOME):

Note: Ensure the grid home directory is not a subdirectory of the Oracle base directory. Installing Oracle Clusterware in an Oracle base directory causes installation errors.

```
# mkdir -p /u02/app/11.2.0/grid
# chown -R grid:oinstall /u02/app/11.2.0/grid
```

4 Create the Oracle database home directory (ORACLE_HOME):

```
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

To create the file system and directories on the Veritas File system (VxFS)

The sample commands in the procedure are for node galaxy. Repeat the steps on each node of the cluster.

1 As the root user, create a VxVM local disk group `bindg_`*`hostname`* on each node.

Make sure that the size of the disk group is approximately 20 GB; if not, add another disk to the disk group.

```
# vxdg init bindg_galaxy Disk_1
```

2 Create separate volumes for Oracle Grid Infrastructure (`crsbinvol`) and Oracle database (`orabinvol`):

```
# vxassist -g bindg_galaxy make crsbinvol 10G
# vxassist -g bindg_galaxy make orabinvol 8G
```

3 Create file systems with the volumes `crsbinvol` and `orabinvol`:

```
# mkfs -V vxfs /dev/vx/rdsk/bindg_galaxy/crsbinvol
# mkfs -V vxfs /dev/vx/rdsk/bindg_galaxy/orabinvol
```

4 Mount the file systems. Perform this step on each node.

```
# mount -V vxfs /dev/vx/dsk/bindg_galaxy/crsbinvol \
/u02
# mount -V vxfs /dev/vx/dsk/bindg_galaxy/orabinvol \
/u01
```

5 Create the following directories for Oracle, ORACLE_BASE, GRID_HOME, ORACLE_HOME.

The file system and directories created on shared storage in this procedure are based on the following layout:

$ORACLE_BASE	For grid user:
	/u02/app/grid
	For Oracle user:
	/u01/app/oracle
$GRID_HOME	*/u02/app/11.2.0/grid*
$ORACLE_HOME	*/u01/app/oracle/product/11.2.0/dbhome_1*

```
# mkdir -p /u02/app/11.2.0/grid
# mkdir -p /u01/app/oracle
# mkdir -p /u02/app/grid
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
```

6 Change the ownership and permissions on all nodes of the cluster.

Note: The ownership and permissions must be changed on all nodes of the cluster because /u02/app must be owned by grid:oinstall, otherwise /u02/app/oraInventory does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

```
# chown -R grid:oinstall /u02/app/grid
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
# chmod -R 775 /u02/app/grid
```

7 Add an entry for the filesystem in the `/etc/filesystems` file on each node:

```
/u01:
    dev = /dev/vx/dsk/bindg_galaxy/orabinvol
    vfs = vxfs
    mount = true
    check = true
    options = rw
    account = false
/u02:
    dev = /dev/vx/rdsk/bindg_galaxy/crsbinvol
    vfs = vxfs
    mount = true
    check = true
    options = rw
    account = false
```

8 Repeat all the steps on each node of the cluster.

Note: If you restart the nodes, the disk groups that were in deported state before restarting are not imported automatically.

Import the disk groups manually and mount the file system after the nodes restart:

\# `vxdg import dg_name`

\# `vxvol -g dg_name startall`

\# `mount -a`

To create the file system and directories on cluster file system for Oracle Grid Infrastructure and Oracle database

Perform the following steps on the CVM master node in the cluster.

1 As the root user, create a VxVM shared disk group `bindg`:

```
# vxdg -s init bindg Disk_1
```

2 Create separate volumes for Oracle Grid Infrastructure (`crsbinvol`) and Oracle database (`orabinvol`):

```
# vxassist -g bindg make crsbinvol 10G
# vxassist -g bindg make orabinvol 8G
```

3 Create the following directories for Oracle, ORACLE_BASE, GRID_HOME, ORACLE_HOME.

The file system and directories created on shared storage in this procedure are based on the following layout:

$ORACLE_BASE	For grid user:
	/u02/app/grid
	For Oracle user:
	/u01/app/oracle
	/u02/app/grid and /u01/app/oracle are on local storage.
$GRID_HOME	*/u02/app/11.2.0/grid*
	/u02/app/11.2.0 is on local storage.
	/u02/app/11.2.0/grid is on shared storage.
$ORACLE_HOME	*/u01/app/oracle/product/11.2.0/dbhome_1*
	/u01/app/oracle/product/11.2.0 is on local storage.
	/u01/app/oracle/product/11.2.0/dbhome_1 is on shared storage.

```
# mkdir -p /u02/app/11.2.0/grid
# mkdir -p /u01/app/oracle
# mkdir -p /u02/app/grid
# mkdir -p /u01/app/oracle/product/11.2.0/dbhome_1
```

4 Create file systems with the volumes crsbinvol and orabinvol:

```
# mkfs -V vxfs /dev/vx/rdsk/bindg/crsbinvol
# mkfs -V vxfs /dev/vx/rdsk/bindg/orabinvol
```

5 Mount the file systems. Perform this step on each node.

```
# mount -V vxfs -o cluster /dev/vx/dsk/bindg/crsbinvol \
/u02/app/11.2.0/grid
# mount -V vxfs -o cluster /dev/vx/dsk/bindg/orabinvol \
/u01/app/oracle/product/11.2.0/dbhome_1
```

6 Change the ownership and permissions on all nodes of the cluster.

Note: The ownership and permissions must be changed on all nodes of the cluster because /u02/app must be owned by grid:oinstall, otherwise/u02/app/oraInventory does not get created correctly on all the nodes. This can cause the Oracle Universal Installer to fail.

```
# chown -R grid:oinstall /u02/app/grid
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
# chmod -R 775 /u02/app/grid
```

7 Add the CVMVolDg and CFSMount resources to the VCS configuration.

See "To add the CFSMount and CVMVolDg resources to the VCS configuration using CLI" on page 16.

To add the CFSMount and CVMVolDg resources to the VCS configuration using CLI

1 Change the permissions on the VCS configuration file:

```
# haconf -makerw
```

2 Configure the CVM volumes under VCS:

```
# hares -add crsorabin_voldg CVMVolDg cvm

# hares -modify crsorabin_voldg Critical 0

# hares -modify crsorabin_voldg CVMDiskGroup bindg

# hares -modify crsorabin_voldg CVMVolume -add crsbinvol

# hares -modify crsorabin_voldg CVMVolume -add orabinvol

# hares -modify crsorabin_voldg CVMActivation sw
```

3   Set up the file system under VCS:

```
# hares -add crsbin_mnt CFSMount cvm

# hares -modify crsbin_mnt Critical 0

# hares -modify crsbin_mnt MountPoint "/u02/app/11.2.0/grid"

# hares -modify crsbin_mnt BlockDevice  \
"/dev/vx/dsk/bindg/crsbinvol"

# hares -add orabin_mnt CFSMount cvm

# hares -modify orabin_mnt Critical 0

# hares -modify orabin_mnt MountPoint \
"/u01/app/oracle/product/11.2.0/dbhome_1"

# hares -modify orabin_mnt BlockDevice  \
"/dev/vx/dsk/bindg/orabinvol"
```

4   Link the parent and child resources:

```
# hares -link crsorabin_voldg cvm_clus

# hares -link crsbin_mnt crsorabin_voldg

# hares -link crsbin_mnt vxfsckd

# hares -link orabin_mnt crsorabin_voldg

# hares -link orabin_mnt vxfsckd
```

5   Enable the resources:

```
# hares -modify crsorabin_voldg Enabled 1

# hares -modify crsbin_mnt Enabled 1

# hares -modify orabin_mnt Enabled 1

# haconf -dump -makero
```

6 Verify the resource configuration in the main.cf file.

```
CFSMount crsbin_mnt (
              Critical = 0
              MountPoint = "/u02/app/11.2.0/grid"
              BlockDevice = "/dev/vx/dsk/bindg/crsbinvol"
              )


CFSMount orabin_mnt (
              Critical = 0
              MountPoint = "/u01/app/oracle/product/11.2.0/dbhome_1"
              BlockDevice = "/dev/vx/dsk/bindg/orabinvol"
              )


CVMVolDg crsorabin_voldg (
              Critical = 0
              CVMDiskGroup = bindg
              CVMVolume = { crsbinvol, orabinvol }
              CVMActivation = sw
              )
crsbin_mnt requires crsorabin_voldg
crsbin_mnt requires vxfsckd
orabin_mnt requires crsorabin_voldg
orabin_mnt requires vxfsckd
crsorabin_voldg requires cvm_clus
```

7 Verify that the resources are online on all systems in the cluster.

```
# hares -state crsorabin_voldg

# hares -state crsbin_mnt

# hares -state orabin_mnt
```

Note: At this point, the crsorabin_voldg resource is reported offline, and the underlying volumes are online. Therefore, you need to manually bring the resource online on each node.

To bring the resource online manually:

```
# hares -online crsorabin_voldg -sys galaxy

# hares -online crsorabin_voldg -sys nebula
```

Setting up user equivalence

You must establish Oracle user and grid user equivalence on all nodes to allow the Oracle Universal Installer to securely copy files and run programs on the nodes in the cluster without requiring password prompts.

Set up passwordless SSH communication between the cluster nodes for the Oracle user and the grid user.

For more information, see the Oracle documentation.

Synchronizing time settings on cluster nodes

Synchronize the time settings on all cluster nodes using NTP. For more information, see the Oracle documentation.

Editing the user profiles

Edit the Oracle user and grid user profile files to set the paths to ORACLE_BASE and ORACLE_HOME on each node.

In the following sample procedure, the shell environment is `ksh`, the Oracle user home directory is `/home/oracle`, and the Oracle grid user home directory is `/home/grid`.

To edit the Oracle user profile

1 As the Oracle user, set the proper environment variables on each node.

```
$ cat .profile
set -o vi

PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:\
/usr/bin/X11:/sbin:.

export PATH

if [ -s "$MAIL" ]          # This is at Shell startup.  In normal
then echo "$MAILMSG"       # operation, the Shell checks
fi                         # periodically.

ORACLE_HOME=/u01/app/oracle/product/11.2.0/dbhome_1; export ORACLE_HOME
ORA_NLS33=$ORACLE_HOME/ocommon/nls/admin/data; export ORA_NLS33
CLASSPATH=$CLASSPATH:$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:\
$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib; export CLASSPATH
DISPLAY=YOUR.DISPLAY.IP; export DISPLAY
GRID_HOME=/u02/app/11.2.0/grid; export GRID_HOME

LIBPATH=$ORACLE_HOME/lib:$ORACLE_HOME/lib32:$GRID_HOME/lib:\
$GRID_HOME/lib32:/usr/lib:/opt/VRTSdbms3/lib; export LIBPATH
PATH=$PATH:$ORACLE_HOME/bin:$GRID_HOME/bin:/opt/VRTS/bin:\
/opt/VRTSdbed/bin; export PATH
```

2 Apply the profile changes:

```
$ . /home/oracle/.profile
```

To edit the Oracle grid user profile

1 As the Oracle grid user, set the proper environment variables on each node.

```
$ cat .profile
set -o vi

PATH=/usr/bin:/etc:/usr/sbin:/usr/ucb:$HOME/bin:\
/usr/bin/X11:/sbin:.

export PATH

if [ -s "$MAIL" ]          # This is at Shell startup.  In normal
then echo "$MAILMSG"       # operation, the Shell checks
fi                         # periodically.

ORACLE_HOME=/u01/app/oracle/product/11.2.0/dbhome_1; export ORACLE_HOME
ORA_NLS33=$ORACLE_HOME/ocommon/nls/admin/data; export ORA_NLS33
CLASSPATH=$CLASSPATH:$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:\
$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib; export CLASSPATH
DISPLAY=YOUR.DISPLAY.IP; export DISPLAY
GRID_HOME=/u02/app/11.2.0/grid; export GRID_HOME

LIBPATH=$ORACLE_HOME/lib:$ORACLE_HOME/lib32:$GRID_HOME/lib:\
$GRID_HOME/lib32:/usr/lib:/opt/VRTSdbms3/lib; export LIBPATH
PATH=$PATH:$ORACLE_HOME/bin:$GRID_HOME/bin:/opt/VRTS/bin:\
/opt/VRTSdbed/bin; export PATH
```

2 Apply the profile changes:

```
$ . /home/grid/.profile
```

Verifying whether the Veritas Membership library is linked to Oracle libraries

The Veritas Membership library (VCSMM) must be linked with Oracle libraries to enable coordinated exchange of cluster membership information and protection of data integrity. Oracle uses the linked skgxn library (libskgxn) to make ioctl calls to VCSMM, which in turn obtains membership information for clusters and instances.

To verify whether the Veritas Membership library is linked to Oracle libraries

1 Verify that the library `/opt/ORCLcluster/lib/libskgxn2.so` is linked to the
`/opt/VRTSvcs/rac/lib64/libvcsmm.so` library:

```
# ls -l /opt/ORCLcluster/lib
libskgxn2.so -> /opt/VRTSvcs/rac/lib64/libvcsmm.so
```

If the link does not exist, create a symbolic link to the Oracle library as follows:

```
$ ln -s /opt/VRTSvcs/rac/lib64/libvcsmm.so \
/opt/ORCLcluster/lib/libskgxn2.so
```

2 Create the required sub-directories and set appropriate permissions:

Create the `/usr/sbin/cluster/utilities/` directory:

```
# mkdir -p /usr/sbin/cluster/utilities/
```

Set the cluster ID in the `cldomain` file:

```
# lltstat -C > /usr/sbin/cluster/utilities/cldomain
```

Change the permissions on the file:

```
# chmod 744 /usr/sbin/cluster/utilities/cldomain
```

3 Verify that the library `/opt/ORCLcluster/lib/libskgxnr.so` is linked to
the `/opt/VRTSvcs/rac/lib64/libvcsmm.so` library:

```
# ls -l /opt/ORCLcluster/lib
libskgxnr.so -> /opt/VRTSvcs/rac/lib64/libvcsmm.so
```

If the link does not exist, create a symbolic link to the Oracle library as follows:

```
$ ln -s /opt/VRTSvcs/rac/lib64/libvcsmm.so \
/opt/ORCLcluster/lib/libskgxnr.so
```

Installing Oracle Grid Infrastructure using the Oracle Universal Installer

This section provides instructions for installing the Oracle Grid Infrastructure
software using the Oracle Universal Installer. The software is installed on each
node in the Oracle Grid Infrastructure home directory.

To install Oracle Grid Infrastructure using the Oracle Universal Installer

1 The *ncargs* attribute controls the maximum command line length, including the environment variables on the AIX nodes. Verify the current value of the *ncargs* attribute.

    # **lsattr -l sys0 -a ncargs -E**

If the value is less than 128, change the attribute to a larger value before invoking the Oracle Installer.

    # **chdev -l sys0 -a ncargs=128**

2 Log in as the Oracle grid user. On the first node, set the DISPLAY variable.

- For Bourne Shell (bash), type:

    $ **DISPLAY=*10.20.12.150:0.0*;export DISPLAY**

  where 10.20.12.150 is the IP address of X client where you want to export the display for the installer.

- For C Shell (csh or tcsh), type:

    $ **setenv DISPLAY *10.20.12.150:0.0***

  where 10.20.12.150 is the IP address of X client where you want to export the display for the installer.

3 Start the Oracle Universal Installer on the first node.

    $ **cd */dvd_mount***

    $ **./runInstaller**

4 Enter the following information when prompted by the Oracle Universal Installer:

- Select installation option
  Select the option **Install and Configure Grid Infrastructure for a Cluster**.

- Select installation type
  Select the option **Advanced Installation**.

- Specify cluster configuration
  Enter the SCAN name for the cluster that will be used by the database clients to connect to databases within the cluster.

- Grid Plug and Play information

  Provide the following information:

- Name of the cluster
- SCAN name
  The SCAN address on the domain name server (DNS) must resolve to three addresses (recommended) or at least one address.
- SCAN port

■ Specify network interface usage
Identify the planned use for each interface: Public, Private, or Do Not use.

Note: Make sure that the same private interfaces that you specified at the time of configuring PrivNIC and MultiPrivNIC are listed on the screen.

Note: Mark the interfaces for the subnet containing the private IP addresses managed by the PrivNIC/MultiPrivNIC agents as 'Private'.

The interfaces that are **Private** are stored in GPnP profile as a 'cluster_interconnect' for Oracle Clusterware communication and database cache fusion traffic.

■ Storage option information
Select the option **Shared File System**.

■ OCR storage option
Enter the full path of the location where you want to store the OCR information.
For example, if you are storing the OCR information on CFS, enter:
`/ocrvote/ocr`.

Note: Select the option **External Redundancy**. Mirror the OCR volumes using CVM.

■ Voting Disk storage option
Enter the full path of the location where you want to store the voting disk information.
For example, if you are storing the voting disk information on CFS, enter:
`/ocrvote/vote`

Note: Select the option **External Redundancy**. Mirror the voting disk volumes using CVM.

- ■ Specify installation location
  Enter the full path to the Oracle base directory and the Oracle Grid
  Infrastructure home directory.

- ■ Create inventory
  Enter the full path to the Oracle inventory directory where you want to
  store the installation files.

5 Review the configuration summary presented by the Oracle Universal
  Installer. The Oracle Universal Installer begins the Oracle Grid Infrastructure
  installation.

6 Run the orainstRoot.sh script as prompted by the Oracle Universal Installer.

7 Run the root.sh script on each node as prompted by the Oracle Universal
  Installer:

   # **cd $GRID_HOME**

   # **./root.sh**

   The Oracle Clusterware daemons are started on the node.

Configuring LLT links in the GPnP profile

Update the GPnP profile to include the remaining LLT links that were not added
to the profile during the Oracle Grid Infrastructure installation.

To configure the LLT links in the GPnP profile

1   View the currently configured interfaces:

```
# $GRID_HOME/bin/oifcfg getif
en0 10.2.156.0        global       public
en1 192.168.12.0      global       cluster_interconnect
```

The interfaces that are currently stored in the GPnP profile, their subnets, and their role (public or cluster_interconnect) are displayed.

2   Add the remaining LLT links to the GPnP profile:

```
# $GRID_HOME/bin/oifcfg setif -global \
en2/192.168.12.0:cluster_interconnect
```

If you are using multiple IP addresses on different subnet for cluster interconnect (for load balancing), add the remaining interface subnets to the GPnP profile.

```
# $GRID_HOME/bin/oifcfg setif -global \
en2/192.168.2.0:cluster_interconnect
# $GRID_HOME/bin/oifcfg setif -global \
en1/192.168.2.0:cluster_interconnect
```

3   Verify that the correct interface subnet is in use:

```
# $GRID_HOME/bin/oifcfg getif
en0 10.2.156.0        global       public
en1 192.168.12.0      global       cluster_interconnect
en2 192.168.12.0      global       cluster_interconnect
en1 192.168.2.0       global       cluster_interconnect
en2 192.168.2.0       global       cluster_interconnect
```

Make sure all the LLT links are configured and listed in the GPnP profile.

Installing the Oracle RAC database using the Oracle Universal Installer

The following procedure describes how to install the Oracle RAC database using the Oracle Universal Installer. Symantec recommends that you install the Oracle RAC database locally on each node.

To install Oracle RAC database using the Oracle Universal Installer

1   Log in as the Oracle user. On the first node, set the DISPLAY variable.

   ■ For Bourne Shell (bash), type:

```
$ DISPLAY=10.20.12.150:0.0;export DISPLAY
```

■ For C Shell (csh or tcsh), type:

```
$ setenv DISPLAY 10.20.12.150:0.0
```

2 Start the Oracle Universal Installer.

```
$ cd /dvd_mount
```

```
$ ./runInstaller
```

3 Open another console and run the rootpre.sh script as the root user. The location of the script is indicated by the Oracle Universal Installer.

4 Return to the console where you started the installation and confirm the execution of the script.

```
Has 'rootpre.sh' been run by root? [y/n] (n) y
```

5 Enter the following information when prompted by the Oracle Universal Installer:

Select installation option	Select the option **Install database software only**.
Node selection	Select **Real Application Clusters database installation**.
	Select the nodes on which the Oracle RAC database software must be installed.
Select database edition	Select **Enterprise Edition**.
Specify installation location	Review or enter the ORACLE_BASE and ORACLE_HOME directory paths.

The Oracle Universal Installer runs product-specific prerequisite checks. Any items that are flagged must be manually checked and configured.

6 Review the configuration summary presented by the Oracle Universal Installer. The Oracle Universal Installer begins the Oracle database installation.

7 Run the root.sh script as prompted by the Oracle Universal Installer.

```
# cd $ORACLE_HOME
```

```
# ./root.sh
```

Linking the ODM library

Perform the steps in the procedure on each node if the Oracle libraries are on local storage. If the Oracle libraries are installed on shared storage, copy the libraries on one node only. Use the mount command to check that the file system containing the Oracle libraries are mounted.

To link the Veritas ODM library

1   Log in as the Oracle user.

2   Change to the `$ORACLE_HOME/lib` directory:

```
$ cd $ORACLE_HOME/lib
```

3   Back up Oracle's ODM library:

```
$ mv libodm11.so libodm11.so.oracle-`date +%m_%d_%Y-%H_%M_%S`
```

4   Link the Veritas ODM library with Oracle's `libodm` library:

```
$ ln -s /opt/VRTSodm/lib/libodm64.so libodm11.so
```

Configuring the CSSD resource manually

Add the `cssd` resource to the VCS configuration and set CSSD dependencies on the resources that manage OCR and voting disk and the private IP addresses for Oracle Clusterware.

Note: It is recommended that the OCR, voting disk, and PrivNIC/MultiPrivNIC resources be configured in the same VCS group as that of the cssd resource. If the resources are not in the same group, set the appropriate dependencies between the service groups.

To configure the CSSD resource

1   Change the permission on the VCS configuration file to read-write mode:

```
# haconf -makerw
```

2   Add the CSSD resource to the cvm group:

```
# hares -add cssd Application cvm
```

3 Modify the CSSD resource attributes:

```
# hares -modify cssd StartProgram /opt/VRTSvcs/rac/bin/cssd-online
# hares -modify cssd StopProgram  /opt/VRTSvcs/rac/bin/cssd-offline
# hares -modify cssd MonitorProgram /opt/VRTSvcs/rac/bin/cssd-monitor
# hares -modify cssd CleanProgram /opt/VRTSvcs/rac/bin/cssd-clean
# hares -modify cssd Critical 0
# hares -override cssd OnlineWaitLimit
# hares -modify cssd OnlineWaitLimit 5
```

4 Enable the CSSD resource:

```
# hares -modify cssd Enabled 1
```

5 Set the dependency of the CSSD resource on the CFSMount or CVMVolDg resources that manage OCR and voting disk.

If you configured OCR and voting disk on CVM raw volumes:

```
# hares -link cssd ocr_voldg_ocrvotedg
```

```
# hares -link cssd vote_voldg_ocrvotedg
```

If you configured OCR and voting disk on CFS:

```
# hares -link cssd ocrvote_mnt_ocrvotedg
```

6 Set the dependency of the CSSD resource on the PrivNIC or MultiPrivNIC resources that manage the private IP address for Oracle Clusterware.

If you configured the PrivNIC resource:

```
# hares -link cssd ora_priv
```

If you configured the MultiPrivNIC resource:

```
# hares -link cssd multi_priv
```

7   If the Oracle Clusterware and the Oracle database binaries are on CFS, set
    the dependencies between the CSSD resource and the CFSMount resources
    for the binaries manually:

    ```
    # hares -link cssd crsbin_mnt
    # hares -link cssd orabin_mnt
    ```

8   Change the permission on the VCS configuration file to read-only mode:

    ```
    # haconf -dump -makero
    ```

Preventing automatic startup of Oracle Clusterware

The use of the CSSD agent is mandatory to ensure adequate handling of service
group inter-dependencies and thereby prevent the premature startup of Oracle
Clusterware. Therefore, disable automatic startup of Oracle Clusterware when
the system starts.

To prevent automatic startup of Oracle Clusterware

1   Log in as the root user on each node in the cluster.

2   Disable automatic startup of Oracle Clusterware:

    ```
    # $GRID_HOME/bin/crsctl disable crs
    ```

    Where $GRID_HOME is the Oracle Grid Infrastructure home directory.

Creating the Oracle RAC database

Create the Oracle RAC database on shared raw volumes or cluster file systems.

To create the requisite database storage, see the instructions in the *Veritas Storage
Foundation for Oracle RAC Installation and Configuration Guide (5.0 Maintenance
Pack 3)*.

For instructions on creating the database, see the Oracle RAC documentation.

Preventing automatic database startup

Configure the Oracle RAC database for manual startup if you want the Oracle RAC
database to be managed by VCS using the Oracle agent. Before you configure the
VCS service groups for Oracle, you need to prevent the Oracle database from
starting automatically. The Oracle Clusterware and Oracle agent may attempt to
start the database instance at the same time if the database mount is available.
To prevent the Oracle database from starting automatically, you must change the

management policy for the database from automatic to manual using the Oracle SRVCTL command. The command changes the AUTO_START attribute of the Oracle database and instance resources.

To prevent automatic database startup

1   Register the database, if not already registered:

```
$ srvctl add database -d db-name -o oracle_home \
-p location-of-parameterfile -y manual
```

2   Once the database is registered, change the management policy for the database to manual:

```
$ srvctl stop database -d db-name
$ srvctl modify database -d db-name -y manual
```

3   Start the database:

```
$ srvctl start database -d db-name
```

Configuring VCS service groups for Oracle RAC

You can set up the Oracle database to be managed by VCS or by Oracle Clusterware. Symantec recommends that the Oracle database be configured under VCS.

Note: The VCS Oracle agent is supported only for administrator-managed databases. For policy-managed databases, use the service group configuration without the VCS Oracle agent.

When the database is configured under VCS:

- You can choose to configure the service group in a way that insulates all the databases from failure in any of the databases in the group.

- VCS manages the start and stop sequence of the applications and the database.

Figure 1-1 illustrates a service group configuration with the VCS Oracle agent.

Figure 1-1   Service group configuration with the VCS Oracle agent
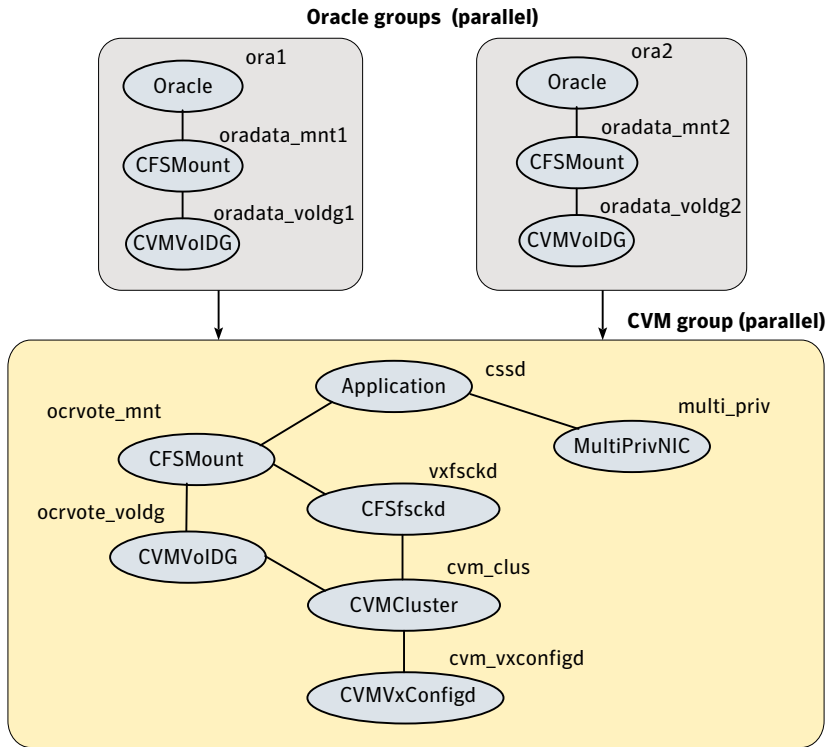
Figure 1-2 illustrates a service group configuration without the VCS Oracle agent.

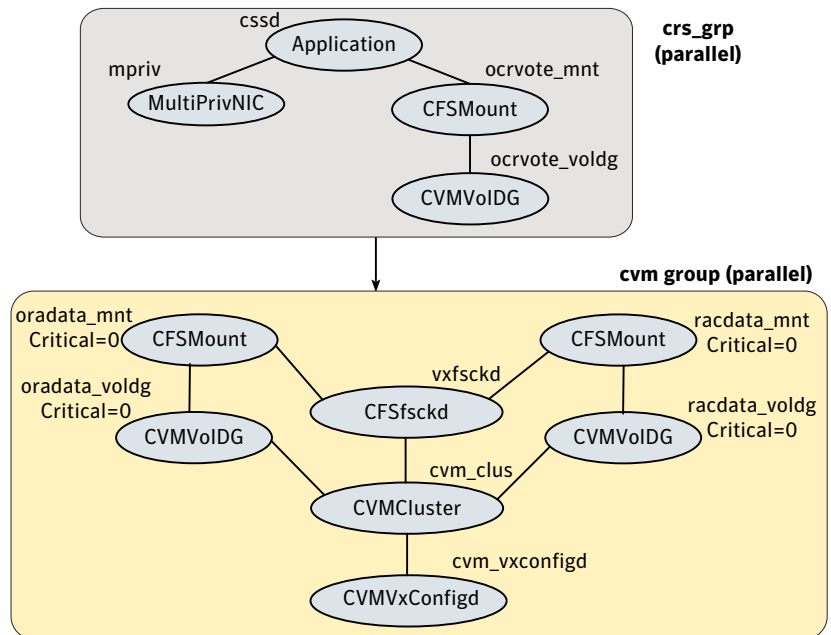**Figure 1-2**   Service group configuration without the VCS Oracle agent

In a service group configuration without the VCS Oracle agent, Oracle Clusterware controls the database. An online local firm dependency exists between the Oracle Clusterware group and the CVM group. When the system starts, the CVM group brings up the volume and mount points for the databases. The Oracle Clusterware group brings up the OCR and voting disk, configures the private IP address for Oracle Clusterware, and starts Oracle Clusterware. Oracle Clusterware starts the database and the application is brought online.

Note: In a service group configuration without the VCS Oracle agent, when the system starts, all volumes and mount points MUST be online for the dependent service groups to be online.

To configure the Oracle database under VCS, create Oracle service groups after installing Oracle RAC and creating a database.

You can create Oracle service groups by modifying the VCS configuration file in one of the following ways:

- Edit the main.cf file
  See "Editing the main.cf file to configure VCS service groups for Oracle RAC" on page 34.

■ Use the CLI
  See "Using the CLI to configure VCS service groups for Oracle RAC" on page 36.

Editing the main.cf file to configure VCS service groups for Oracle RAC

This section describes how to manually edit the
`/etc/VRTSvcs/conf/config/main.cf` file to configure VCS service groups for
Oracle RAC.

Note: Before you edit the main.cf file, change the cluster configuration to read-write
mode: # `haconf -makerw`

You need to perform the steps in the following procedures:

■ Configure VCS service groups for Oracle RAC
  See "To configure VCS service groups for Oracle RAC" on page 34.

■ Start VCS after modifying the configuration file
  See "To start VCS after modifying the configuration file" on page 36.

To configure VCS service groups for Oracle RAC

1 Log in as the root user to one of the systems.

2 As root user, save your existing configuration before you modify main.cf:

   # `haconf -dump -makero`

   If the configuration is not writable, a warning appears: "Cluster not writable."
   You may safely ignore the warning.

3 Stop the VCS engine on all systems and leave the resources available:

   # `hastop -all -force`

4 Back the `main.cf` file:

   # `cd /etc/VRTSvcs/conf/config`

   # `cp main.cf main.orig`

5 Using vi or another text editor, edit the main.cf file to add the Oracle service groups.

For example:

```
group oradb_grp (
        SystemList = { galaxy = 0, nebula = 1 }
        AutoFailOver = 0
        Parallel = 1
        AutoStartList = { galaxy, nebula }
        )

        Oracle oradb (
                Critical = 0
                Sid @galaxy = oradb1
                Sid @nebula = oradb2
                Owner = oracle
                Home = "/u01/app/oracle/product/11.2.0/dbhome_1"
                StartUpOpt = "SRVCTLSTART"
                ShutDownOpt = "SRVCTLSTOP"
                )

        CFSMount oradata_mnt (
                Critical = 0
                MountPoint = "/oradbdata"
                BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
                )

        CVMVolDg oradata_voldg (
                Critical = 0
                CVMDiskGroup = oradatadg
                CVMVolume = { oradatavol }
                CVMActivation = sw
                )

        requires group cvm online local firm
        oradb requires oradata_mnt
        oradata_mnt requires oradata_voldg
```

6 Save and close the main.cf file.

7 Verify the syntax of the file :

```
# cd /etc/VRTSvcs/conf/config
```

```
# hacf -verify .
```

For additional information and instructions on modifying the VCS configuration by editing the main.cf file:

See the *Veritas Cluster Server User's Guide*.

Starting VCS after modifying the configuration file

You need to start VCS for the new configuration to take effect.

To start VCS after modifying the configuration file

1 Start VCS on the node on which you modified the configuration file and check its status:

```
# hastart
```

```
# hastatus
```

2 When "LOCAL_BUILD" is listed in the message column, start VCS on the other systems:

```
# hastart
```

3 On one of the systems, enter the following command to verify that the service groups for Oracle RAC are brought online:

```
# hagrp -display
```

4 Check the status of the groups.

```
# hagrp -state
```

5 Check the status of the resources.

```
# hares -state
```

Using the CLI to configure VCS service groups for Oracle RAC

This section describes how to configure the Oracle service group using the CLI.

The following procedure assumes that you have created the database.

To configure the Oracle service group using the CLI

1 Change the cluster configuration to read-write mode:

> # **haconf -makerw**

2 Add the service group to the VCS configuration:

> # **hagrp -add** *oradb1_grp*

3 Modify the attributes of the service group:

> # **hagrp -modify** *oradb1_grp* **Parallel 1**
>
> # **hagrp -modify** *oradb1_grp* **SystemList** *galaxy* **0** *nebula* **1**
>
> # **hagrp -modify** *oradb1_grp* **AutoStartList** *galaxy nebula*

4 Add the CVMVolDg resource for the service group:

> # **hares -add** *oradata_voldg* **CVMVolDg** *oradb1_grp*

5 Modify the attributes of the CVMVolDg resource for the service group:

> # **hares -modify** *oradata_voldg* **CVMDiskGroup oradatadg**
> # **hares -modify** *oradata_voldg* **CVMActivation sw**
> # **hares -modify** *oradata_voldg* **CVMVolume oradatavol**

6 Add the CFSMount resource for the service group:

> # **hares -add** *oradata_mnt* **CFSMount** *oradb1_grp*

7 Modify the attributes of the CFSMount resource for the service group:

> # **hares -modify** *oradata_mnt* **MountPoint "/oradata"**
> # **hares -modify** *oradata_mnt* **BlockDevice **
> **"/dev/vx/dsk/oradatadg/oradatavol"**

8 Add the Oracle RAC database instance to the service group:

> # **hares -add ora1 Oracle oradb1_grp**

9 Modify the attributes of the Oracle resource for the service group:

```
# hares -modify ora1 Owner oracle
# hares -modify ora1 Home "/app/oracle/orahome"
# hares -modify ora1 StartUpOpt SRVCTLSTART
# hares -modify ora1 ShutDownOpt SRVCTLSTOP
```

10 Localize the Sid attribute for the Oracle resource:

```
# hares -local ora1 Sid
```

11 Set the Sid attributes for the Oracle resource on each system:

```
# hares -modify ora1 Sid vrts1 -sys galaxy
# hares -modify ora1 Sid vrts2 -sys nebula
```

12 Set the dependencies between the CFSMount resource and the CVMVolDg resource for the Oracle service group:

```
# hares -link oradata_mnt oradata_voldg
```

13 Set the dependencies between the Oracle resource and the CFSMount resource for the Oracle service group:

```
# hares -link ora1 oradata_mnt
```

14 Create an online local firm dependency between the oradb1_grp service group and the cvm service group:

```
# hagrp -link oradb1_grp cvm online local firm
```

15 Enable the Oracle service group:

```
# hagrp -enableresources oradb1_grp
```

16 Change the cluster configuration to the read-only mode:

```
# haconf -dump -makero
```

17 Bring the Oracle service group online on all the nodes:

```
# hagrp -online oradb1_grp -any
```

For more information and instructions on configuring the service groups using the CLI:

See the *Veritas Cluster Server User's Guide*.

Location of VCS log files

You may want to review the log files at `/var/VRTSvcs/log/engine_A.log` for errors or status messages. When large amounts of data are written, multiple log files may be written, such as engine_B.log, engine_C.log, and so on. The engine_A.log contains the most recent data.

Upgrading to Oracle RAC 11g Release 2

This section provides instructions on upgrading to Oracle RAC 11g Release 2.

Note: Starting with Oracle RAC 11g Release 2, the Oracle Clusterware software must be upgraded to run in a new directory called the Oracle Grid Infrastructure home. If you plan to upgrade ASM to Release 11g R2, make sure that ASM also resides in the same Oracle Grid Infrastructure home directory.

Even if the existing Oracle Clusterware software is installed in a shared home directory, you can use a local home directory for installing the Oracle Grid Infrastructure software.

For information on supported upgrade paths:

See "Supported upgrade paths" on page 39.

Complete the following steps to upgrade to Oracle RAC 11g Release 2:

1. Prepare to upgrade Oracle RAC.

   See " Preparing to upgrade Oracle RAC" on page 40.

2. Upgrade Oracle RAC binaries.

   See "Upgrading Oracle RAC binaries" on page 41.

3. Migrate the Oracle RAC database.

   See "Migrating the Oracle RAC database" on page 41.

4. Complete the post-upgrade tasks.

   See "Performing post-upgrade tasks" on page 41.

Supported upgrade paths

Table 1-2 lists the upgrade paths for Oracle RAC.

Table 1-2   Supported upgrade paths for Oracle RAC

From current version	Upgrade to
Oracle RAC 10g Release 2	Oracle RAC 11g Release 2
Oracle RAC 11g Release 1	Oracle RAC 11g Release 2

Note: When you upgrade to a different version of Oracle RAC, make sure that the full path of the Oracle Grid Infrastructure home directory and the Oracle database home directory is different from the path where the existing version of Oracle RAC resides.

The upgrade procedure assumes that the beginning configuration includes the following components, and that these components are running on the cluster nodes:

■   SF Oracle RAC 5.0 MP3 RP4

■   A supported version of the operating system

Preparing to upgrade Oracle RAC

To prepare for upgrading Oracle RAC

1   Take a hot or cold backup of the existing database.

2   Back up the existing Oracle home and central inventory.

3   If the Oracle RAC database is under VCS control, freeze the Oracle service groups to prevent VCS from reporting the resource as faulted when Oracle RAC stops and starts the database during the upgrade:

```
# haconf -makerw
```

```
# hagrp -freeze oracle_group -persistent
```

4   Freeze the cvm service group to prevent VCS from reporting the resource as faulted when Oracle Clusterware is stopped and started during the upgrade:

```
# hagrp -freeze cvm_group -persistent
```

```
# haconf -dump -makero
```

Upgrading Oracle RAC binaries

Review your Oracle installation manuals and the appropriate Oracle support Web sites before upgrading Oracle RAC.

Note: Make sure that you upgrade to Oracle RAC 11g Release 2 with the same user credentials as that of the existing installation.

To upgrade Oracle RAC binaries

1 Upgrade Oracle Clusterware to a new directory called the Oracle grid infrastructure home (GRID_HOME).

Note: Starting with Oracle RAC 11g Release 2, ASM must reside in the Oracle Grid Infrastructure home directory. If you plan to upgrade ASM to Release 11g R2, make sure that you upgrade it to run in the Oracle Grid Infrastructure home directory.

For instructions, see the Oracle RAC documentation.

2 Make sure that Oracle Clusterware is running.

3 Install the Oracle RAC database binaries.

For instructions, see the Oracle RAC documentation.

4 Relink the SF Oracle RAC libraries with Oracle RAC.

See "Linking the ODM library" on page 28.

Migrating the Oracle RAC database

For instructions on migrating the existing Oracle RAC database, see the Oracle metalink documentation.

After migrating the database, complete the post-upgrade tasks:

See "Performing post-upgrade tasks" on page 41.

Performing post-upgrade tasks

Perform the steps in the following procedure to complete the upgrade.

To perform post-upgrade tasks

1   Change the cluster configuration to read-write mode:

    # **haconf -makerw**

2   Modify the Oracle RAC configuration to prevent automatic startup of Oracle
    Clusterware.

    See "Preventing automatic startup of Oracle Clusterware" on page 30.

3   Modify the Oracle RAC database configuration to prevent automatic database
    startup if you want the Oracle RAC database to be managed by VCS using the
    Oracle agent.

    See "Preventing automatic database startup" on page 30.

4   Unfreeze the VCS service groups that were frozen earlier.

    As root user, enter:

    # **hagrp -unfreeze *oracle_group* -persistent**

    # **hagrp -unfreeze *cvm_group* -persistent**

    # **haconf -dump -makero**