

---

# Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

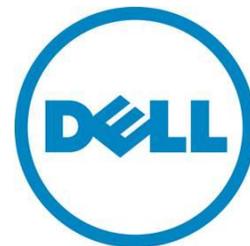
---

*Step-by-Step How-To Guide*

Authors

Roger Lopez, Dell Inc.

Adam Miller, Red Hat Inc.



**This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.**

© 2012 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell, the Dell logo, and PowerEdge are trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Microsoft, Windows, and Windows Server are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

June 2012 | Rev 1.0

## Contents

Overview .....	3
Software Requirements.....	4
Hardware Requirements.....	4
Network Requirements .....	5
Operating System Requirements .....	5
Oracle Prerequisites .....	5
Public Network Setup.....	6
Private Network Setup .....	7
IP Address and Name Resolution Requirements .....	9
For a Cluster Using DNS .....	9
Configuring a DNS Server .....	10
Configuring a DNS Client .....	10
Preparing Shared Storage for Oracle Installation .....	11
Identifying ASM disks to setup udev rules .....	11
Partitioning the Shared Disk .....	15
Installing Oracle 11gR2 Grid Infrastructure .....	16
Before You Begin .....	16
Configure the System Clock Settings for All Nodes .....	16
Configuring Node One .....	17
Configuring Shared Oracle Home using ACFS .....	20
Installing Oracle 11g R2 Database Software .....	21
Creating Diskgroups using ASM Configuration Assistant (ASMCA) .....	22
Creating Database using DBCA .....	23
Troubleshooting / Known Issues .....	24
Acknowledgements .....	25

## Overview

This document applies to IT and database administrators who are deploying Oracle Database 11g R2 database software on Red Hat Enterprise Linux 6.x. The paper is written as a step-by-step guide on how to install, configure, and build an Oracle Real Application Cluster (RAC) on Red Hat Enterprise Linux 6.x Server.

## Software Requirements

Software Installation Location	Size Required
Grid Infrastructure home	4.5GB of space
Oracle Database home	4GB of space
Shared Storage Disk Space	Sizes of Database and Flashback Recovery Area

## Hardware Requirements

- Oracle requires 1 gigabyte (GB) of physical memory
- Swap space is proportional to the amount of RAM allocated to the system

RAM	Swap Space
Up to 2GB	1.5 times the size of RAM
Up to 16GB	Equal the size of RAM
16GB or more	16GB

- Oracle's temporary space (/tmp) must be at least 1 GB in size
- A monitor that supports resolution of 1024 x 768 to correctly display the Oracle Universal Installer (OUI)
- For Dell supported hardware configurations, see the Software Deliverable List (SDL) for each Dell Validated Component at <http://www.dell.com/oracle>

## Network Requirements

- Ensure each node contains at least three network interface cards (NICs) are available. One NIC for public network and two NICs for the private node to node network traffic to ensure high availability of the Oracle RAC cluster.
- Public and private interface names must be the same on all nodes. For example, if eth0 is set as the public interface on node one, ensure all other nodes implement eth0 as the public interface.
- All public interfaces for each node should be able to communicate with all nodes within the cluster.
- All private interfaces for each node should be able to communicate with all nodes within the cluster.
- The hostname of each node must follow the RFC 952 standard (<http://www.ietf.org/rfc/rfc952.txt>). Hostnames that include an underscore ("\_") are not permitted.
- Each node in the cluster requires the following IP address:
  - One public IP address
  - Two private IP address
  - One virtual IP address
  - Three single client access name (SCAN) addresses for the cluster

## Operating System Requirements

- Red Hat Enterprise Linux 6.x Server x86\_64 (Kernel 2.6.32-71 or higher)

## Oracle Prerequisites

The required Oracle prerequisites are bundled together by Dell in a single Dell-Oracle-Deployment.tar.Gzip file. The Dell-Oracle-Deployment Tar Gzip install bundle was introduced to provide customers the freedom of installing their kernel of choice (RedHat or Oracle Linux) with additional value adds not found within the original oracle-validated rpm (such as the creation of the grid user) while assisting in the deployment of Oracle on Dell Tested and Validated Configurations. The Dell-Oracle-Deployment Tar Gzip can be found at: <http://tinyurl.com/7ye69af>. Please view the README for all necessary installation steps.

**NOTE:** A yum repository must be setup either through local installation media or online repository as the bundle requires additional packages that may have not been installed during OS installation.

## Public Network Setup

To configure the public network on each node:

1. Log in as root (or from a sudo user).
2. Edit the network device file `/etc/sysconfig/network-scripts/ifcfg-eth#` where # is the number of the network device:

**NOTE:** Ensure that the public IP address is a valid and routable IP address.

**NOTE:** Ensure that the Gateway address is configured for the public network interface. If the Gateway address is not configured, the Oracle Grid installation may fail.

```
DEVICE=eth0
ONBOOT=yes
IPADDR=<Public IP Address>
NETMASK=<Subnet mask>
BOOTPROTO=static
HWADDR=<MAC Address>
SLAVE=no
GATEWAY=<Gateway Address>
```

3. Edit the `/etc/sysconfig/network` file, and, if necessary, replace `localhost.localdomain` with the qualified public node name. For example, the command for node 1: `hostname=node1.domain.com`
4. Type `service network restart` to restart the network service.
5. Type `ifconfig` to verify that the IP addresses are set correctly.
6. To check your network configuration, ping each node's public IP address from a client on the LAN that is not a part of the cluster.
7. Connect to each node to verify that the public network is functioning. Type `ssh <public IP>` to verify that the secure shell (ssh) connection can be established.

## Private Network Setup

The grid infrastructure of Oracle 11gR2 (11.2.0.3) supports IP failover natively using a new feature introduced known as Redundant Interconnect. Oracle uses its ora.cluster\_interconnect.haip resource to communicate with Oracle RAC, Oracle ASM, and other related services. The Highly Available Internet Protocol (HAIP) has the ability to activate a maximum of four private interconnect connections. These private network adapters can be configured during the initial install process of Oracle Grid or after the installation process using the oifcfg utility.

Oracle Grid currently creates an alias IP (as known as virtual private IP) on your private network adapters using the 169.254.\*.\* subnet for the HAIP. If subnet range is already in use, Oracle Grid will not attempt to use it. The purpose of HAIP is to load balance across all active interconnect interfaces, and failover to other available interfaces if one of the existing private adapters becomes unresponsive.

### NOTE:

- Each of the two NIC ports for the private network must be on separate PCI buses.
- Configure different subnet for each of the private network you want to configure as a part of HAIP
- When adding additional HAIP addresses (maximum of four) after the installation of Oracle Grid, restart of your Oracle Grid environment to make sure these new HAIP addresses are active.

The example below provides step-by-step instructions on enabling redundant interconnect using HAIP on a fresh Oracle 11gR2 (11.2.0.3) Grid Infrastructure installation.

1. Edit the file, `/etc/sysconfig/network-scripts/ifcfg-ethX`, where X is the number of the eth device, `ifcfg-ethX` configuration files of the network adapters to be used for your private interconnect. The following example shows eth1 using 192.168.1.\* subnet and eth2 using 192.168.2.\* .

```
DEVICE=eth1
BOOTPROTO=static
HWADDR=00:1E:C9:4B:72:22
ONBOOT=yes
IPADDR=192.168.1.140
NETMASK=255.255.255.0
```

```
DEVICE=eth2
HWADDR=00:1E:C9:4B:71:24
BOOTPROTO=static
ONBOOT=yes
IPADDR=192.168.2.140
NETMASK=255.255.255.0
```

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

2. Once you have saved both the configuration files, restart your network service using `service network restart`. The completion of the steps above have now prepared your system to enable HAIP using the Oracle Grid Infrastructure installer. When you have completed all the Oracle prerequisites and are ready to install Oracle, you will need to select eth1 and eth2 as 'private' interfaces at the 'Network Interface Usage' screen.

**NOTE:** The following steps enable redundant interconnectivity once your Oracle Grid Infrastructure has successfully completed and is running.

3. To verify that your redundant interconnect using HAIP is running, you can test this feature using the `ifconfig` command. An example of the output is listed below.

```
ifconfig
eth1 Link encap:Ethernet HWaddr
00:1E:C9:4B:72:22
inet addr:192.168.1.140
Bcast:192.168.1.255 Mask:255.255.255.128
inet6 addr: fe80::216:3eff:fe11:1122/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:6369306 errors:0 dropped:0 overruns:0 frame:0
TX packets:4270790 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:3037449975 (2.8 GiB) TX
bytes:2705797005 (2.5 GiB)

eth1:1 Link encap:Ethernet HWaddr
00:1E:C9:4B:72:22
inet addr:169.254.167.163
Bcast:169.254.255.255 Mask:255.255.0.0
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1

eth2 Link encap:Ethernet HWaddr
00:1E:C9:4B:71:24
inet addr:192.168.2.140
Bcast:192.168.2.255 Mask:255.255.255.128
inet6 addr: fe80::216:3eff:fe11:1122/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:6369306 errors:0 dropped:0 overruns:0 frame:0
TX packets:4270790 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:3037449975 (2.8 GiB) TX bytes:2705797005 (2.5 GiB)

eth2:1 Link encap:Ethernet HWaddr
00:1E:C9:4B:71:24
inet addr:169.254.167.164
Bcast:169.254.255.255 Mask:255.255.0.0
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
```

For more information on Redundant Interconnect and ora.cluster\_interconnect.haip, see [My Oracle Support note: 1210883.1](#).

## IP Address and Name Resolution Requirements

The steps below show how to setup your cluster nodes for using Domain Name System (DNS). For information on how to setup cluster nodes using GNS, see the wiki article <http://tinyurl.com/6tehajv>

### For a Cluster Using DNS

Set up an Oracle 11g R2 RAC using Oracle DNS (without GNS):

1. At least two interfaces must be configured on each node, one for the private IP address and one for the public IP address.
2. A SCAN NAME configured on the DNS for Round Robin resolution to three addresses (recommended) or at least one address. The SCAN addresses must be on the same subnet as virtual IP addresses and public IP addresses.

**NOTE:** For high availability and scalability, it is recommended that you configure the SCAN to use Round Robin resolution to three IP addresses. The name for the SCAN cannot begin with a numeral. For installation to succeed, the SCAN must resolve to at least one address.

The table below describes the different interfaces, IP address settings and the resolutions in a cluster.

Cluster Requirements for DNS		
Interface	Type	Resolution
Public	Static	DNS
Private	Static	Not Required
Node Virtual IP	Static	DNS
SCAN Virtual IP	Static	DNS

## Configuring a DNS Server

To configure changes on a DNS server for an Oracle 11g R2 cluster using a DNS (without GNS):

1. Configure SCAN NAME resolution on DNS server. A SCAN NAME configured on the DNS server using the Round Robin policy should resolve to three public IP addresses (recommended), however the minimum requirement is one public IP address.

For example:

```
scancluster IN A 192.0.2.1
            IN A 192.0.2.2
            IN A 192.0.2.3
```

Where scancluster is the SCAN NAME provided during Oracle Grid installation.

**NOTE:** The SCAN IP address must be routable and must be in public range.

## Configuring a DNS Client

1. Configure the `resolv.conf` on all the nodes in the cluster to contain name server entries that resolve to the appropriate DNS server. Provide an entry similar to the following:

```
/etc/resolv.conf:
search ns1.domainserver.com
nameserver 192.0.2.100
```

Where, 192.0.2.100 is a valid DNS server address in your network and ns1.domainserver.com is the domain server in your network.

2. Verify the order configuration. `/etc/nsswitch.conf` controls the name service order. In some configurations, the NIS can cause issues with Oracle SCAN address resolution. The recommendation is to place the NIS entry at the end of the search list and place the dns entry first. For example,

```
hosts: dns files nis
```

Once you have modified the `/etc/nsswitch.conf`, restart the `nscd` service by issuing the command:

```
# /sbin/service nscd restart
```

## Preparing Shared Storage for Oracle Installation

Oracle RAC requires shared LUNs for storing your Oracle Cluster Registry (OCR), voting disks, Oracle Home using ACFS, Oracle Database files, and Flash Recovery Area (FRA). To ensure high availability for Oracle RAC it is recommended that you have:

- Three shared volumes/ LUNs of 1 GB in size for normal redundancy or five volumes/LUNs for high redundancy are required for the Oracle clusterware.
- At least two shared disks to store your database. Each shared disk should be the same disk speed and size.
- At least two shared volumes/LUNs to store your Automatic Storage Management Cluster File System (ACFS). Each shared disk must be at least 10 GB, for a total size of 20 GB.
- At least two shared volumes/LUNs or volumes to store your FRA. Ideally, the FRA space should be large enough to copy all of your Oracle datafiles and incremental backups. For more information on optimally sizing your FRA, see My Oracle Support ID 305648.1 section "What should be the size of Flash Recovery Area?"

**NOTE:** In this section, the terms disk(s), volume(s), virtual disk(s), LUN(s) mean the same and are used interchangeably, unless specified otherwise.

**NOTE:** The use of device mapper multipath is recommended for optimal performance and persistent name binding across nodes within the cluster.

**NOTE:** For more information on attaching shared LUNs/volumes, see the Wiki documentation found at: <http://tinyurl.com/6o7pf7a>

## Identifying ASM disks to setup udev rules

First, configure the `/etc/multipath.conf` (by default this file does not exist in RHEL6).

Run the following command as root to get a base point and adjust the device-mapper-multipath version for your particular release:

```
cp /usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf.defaults  
/etc/multipath.conf
```

Edit your `multipath.conf` as recommended below:

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

```
#----- start multipath.conf snippet -----#
defaults {
    udev_dir                /dev
    polling_interval        5
    path_selector            "round-robin 0"
    path_grouping_policy    failover
    getuid_callout          "/lib/udev/scsi_id --whitelisted --
device=/dev/%n"
    prio                    const
    path_checker            direction
    rr_min_io               1000
    rr_weight               uniform
    failback                manual
    no_path_retry           fail
    user_friendly_names     yes
}

blacklist {
    wwid <local_drives_that_won't_use_multipath>
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z]"
    devnode "^dcssblk[0-9]*"
}

#----- end multipath.conf snippet -----#
```

Next, the `multipaths` section of the `/etc/multipath.conf` requires creating an alias using the appropriate WWID of each disk. In order to capture the WWID, run the `scsi_id` command. In the example below, one disk is used for the database, OCR and voting disk.

**NOTE:** There will be two entries in `/dev/sd*` for each device because each block device has multiple paths.

**Example: Multiple paths for the same device.** Output below is a subset of running the command `cat /proc/partitions`

```
8      16  524288000  sdb
8      48  524288000  sdd
8      32  524288000  sdc
8      64  524288000  sde
```

**Example: Bash script to determine the WWID of each disk using the `scsi_id` command.**

```
[root@rhel6 ~]# for i in sdb sdc sdd sde; do \
    printf "%s %s\n" "$i" \
    "$(scsi_id --page=0x83 --whitelisted --device=/dev/$i)"; done
```

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

```
sdb 360026b900061855e000008a54ea5356a
sdc 360026b9000618571000008b54ea5360b
sdd 360026b900061855e000008a54ea5356a
sde 360026b9000618571000008b54ea5360b
```

Having identified the WWIDs of sdb,sdc,sdd, and sde, add them to the appropriate section within the `/etc/multipath.conf` file.

```
#----- start multipath.conf conf addition -----#

multipaths {

    multipath {

        wwid      360026b9000618571000008b54ea5360b

        alias     db

    }

    multipath {

        wwid      360026b900061855e000008a54ea5356a

        alias     ocrvote

    }

}

#----- end multipath.conf conf addition -----#
```

Once partitioning of the `/dev/mapper/` devices is complete, the permissions of those block devices must be set using udev rules.

1. Capture the WWIDs of each disk to be used within your ASM device using the `scsi_id` command described above.

2. To create udev rules that will be set to each device, create the file `/etc/udev/rules.d/20-dell_oracle.rules`. A separate KERNEL entry must exist for each storage device and will require adding the WWID to the "RESULT==" field.

**NOTE:** udev rules are executed in enumerated order, thus 10-\* rules are run prior to 20-\* rules. Device-mapper rules are executed as 10-\* rules, thus our udev rules require a prefix of 20-\*

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

```
#----- start udev rule contents -----#  
  
KERNEL=="dm-*", PROGRAM="scsi_id --page=0x83 --whitelisted --  
device=/dev/%k",RESULT=="360026b9000618571000008b54ea5360b", OWNER=="grid",  
GROUP=="asmadmin"  
  
KERNEL=="dm-*", PROGRAM="scsi_id --page=0x83 --whitelisted --  
device=/dev/%k",RESULT=="360026b900061855e000008a54ea5356a", OWNER=="grid",  
GROUP=="asmadmin"  
  
#----- end udev rule contents -----#
```

3. Reboot. A reboot is required because we want to ensure a clean block node creation setup.

4. Verify configuration, the `/dev/dm-*` should have uid:gid of `grid:asmadmin`. As root, run the command `ls -l /dev/dm-*`

Example:

```
[root@rhel6 ~]# ls -l /dev/dm-*  
brw-rw----. 1 root disk      253, 0 Dec  6 16:05 /dev/dm-0  
brw-rw----. 1 root disk      253, 1 Dec  6 16:05 /dev/dm-1  
brw-rw----. 1 grid asmadmin 253, 2 Dec  7 14:16 /dev/dm-2  
brw-rw----. 1 grid asmadmin 253, 3 Dec  7 14:16 /dev/dm-3  
brw-rw----. 1 root disk      253, 4 Dec  6 16:06 /dev/dm-4
```

Verify the device mapper multipath aliases are pointing to the correct location. As root, run the command: `ls -l /dev/mapper/db /dev/mapper/vote`

Example:

```
[root@rhel6 ~]# ls -l /dev/mapper/db /dev/mapper/vote  
lrwxrwxrwx. 1 root root 7 Dec  6 16:06 /dev/mapper/db -> ../dm-3  
lrwxrwxrwx. 1 root root 7 Dec  6 16:06 /dev/mapper/vote -> ../dm-2
```

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

Verify that our multipaths are functional. As root, run the command: `multipath -ll`

```
[root@rhel6 ~]# multipath -ll
db (360026b9000618571000008b54ea5360b) dm-3 DELL,MD32xx
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
rdac' wp=rw
|+- policy='round-robin 0' prio=6 status=active
| `-- 3:0:1:1 sde 8:64 active ready running
`-+- policy='round-robin 0' prio=1 status=enabled
  `-- 3:0:0:1 sdc 8:32 active ghost running
    vote (360026b900061855e000008a54ea5356a) dm-2 DELL,MD32xx
      size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
rdac' wp=rw
  |+- policy='round-robin 0' prio=6 status=active
  | `-- 3:0:0:0 sdb 8:16 active ready running
  `-+- policy='round-robin 0' prio=1 status=enabled
    `-- 3:0:1:0 sdd 8:48 active ghost running
```

## Partitioning the Shared Disk

This section describes how to use Linux's native partition utility `fdisk` to create a single partition on a volume/virtual disk that spans the entire disk.

To use the `fdisk` utility to create a partition:

1. At the command prompt, type one of the following:

```
#> fdisk -cu /dev/mapper/<multipath_disk>
```

```
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF
disklabel
```

```
Building a new DOS disklabel with disk identifier 0x93302758.
```

```
Changes will remain in memory only, until you decide to write them.
```

```
After that, of course, the previous content won't be recoverable.
```

```
Warning: invalid flag 0x0000 of partition table 4 will be corrected by
w(rite)
```

```
Command (m for help): n # To create a new partition
```

```
Command action
```

```
e extended
```

```
p primary partition (1-4):
```

```
P # To create a primary partition
```

```
. Partition number (1-4): 1
```

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

- a. First sector (2048-xxxxxx, default 2048): 2) <Enter default value or return key>

Command (m for help): wq # write and quit

The system displays the following message:

```
The partition table has been altered!
```

```
Calling ioctl() to re-read partition table. Syncing disks.If you get a warning message instead saying the kernel still reads the old partition table then follow step 3 below for the kernel to be able to re-read the new partition table.
```

2. Repeat step 1 for all the disks that need to have a partition created.
3. Type the following to re-read the partition table and to be able to see the newly created partition(s)

```
#> kpartx -a /dev/mapper/<multipath_disk>
```

And

```
#> service multipathd restart
```

4. Reboot the system if your newly created partition(s) is not displayed properly.

## Installing Oracle 11gR2 Grid Infrastructure

This section provides installation guidelines for installing Oracle 11g R2 grid infrastructure in a cluster.

### Before You Begin

Before you install the Oracle 11g R2 RAC software on your system:

- Ensure that you have already configured your operating system, network, and storage based on the steps from the previous sections within this document.
- Locate your Oracle 11g R2 media kit.

## Configure the System Clock Settings for All Nodes

To prevent failures during the installation procedure, configure all the nodes with identical system clock settings. Synchronize your node system clock with the Cluster Time Synchronization Service

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

(CTSS) which is built in Oracle 11g R2. To enable CTSS, disable the operating system network time protocol daemon (ntpd) service using the following commands in this order:

1. `service ntpd stop`
2. `chkconfig ntpd off`
3. `mv /etc/ntp.conf /etc/ntp.conf.orig`
4. `rm /var/run/ntpd.pid`

## Configuring Node One

The following steps are for node one of your cluster environment, unless otherwise specified.

- a. Log in as root.
- b. If you are not in a graphical environment, start the X Window System by typing: `startx`
- c. Open a terminal window and type: `xhost +`
- d. Mount the Oracle Grid Infrastructure media.
- e. Log in as grid user, for example: `su - grid`.
- f. Type the following command to start the Oracle Universal Installer: `<CD_mountpoint>/runInstaller`
- g. In the Download Software Updates window, enter your My Oracle Support credentials to download the latest patch updates. If you choose not to download the latest patches, select Skip software updates.
- h. In the Select Installation Option window, select Install and Configure Grid Infrastructure for a Cluster and click Next.
- i. In the Select Installation Type window, select Advanced Installation option, and click Next.
- j. In the Select Product Languages window, select English, and click Next.
- k. In the Grid Plug and Play Information window, enter the following information:
  1. Cluster Name—Enter a name for your cluster.
  2. SCAN Name—Enter the named registered in the DNS server which is unique for the entire cluster. For more details on setting up your SCAN name see, "IP Address and Name Resolution Requirements". For details on enabling GNS, see the wiki article: <http://tinyurl.com/6tehajv>.

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

3. SCAN Port—Retain the default port of 1521.
  4. Configure GNS—Uncheck this option.
  5. Click Next.
- l. In the Cluster Node Information window, click Add to add additional nodes that must be managed by the Oracle Grid Infrastructure.
    1. Enter the public Hostname information
    2. Enter the Virtual IP name
    3. Repeat step for each node within your cluster
  - m. Click SSH Connectivity and configure your passwordless SSH connectivity by entering the OS Password for the grid user and click Setup.

NOTE: The default password set by the dell-validated and dell-oracle-utilities rpms is 'oracle' for both the grid user and oracle user.
  - n. Click Ok and then click Next to go to the next window.
  - o. In the Network Interface Usage window, make sure that the correct interface types are selected for the interface names. From the Interface Type drop-down list, select the required interface type. The available options are Private, Public, and Do Not Use. Click Next.
  - p. In the Storage Option Information window, select Automatic Storage Management (ASM) and click Next.
  - q. In the Create ASM Disk Group window, enter the following information:
    1. ASM diskgroup— Enter a name, for example: OCR\_VOTE
    2. Redundancy— For your OCR and voting disks, select High if five ASM disks are available, select Normal if three ASM disks are available, or select External if one ASM disk is available (not recommended).

**NOTE:** If no candidate disks are displayed, click Change Discovery Path and enter `/dev/mapper/*`

- r. In the Specify ASM Password window, choose the relevant option under Specify the passwords for these accounts and enter the relevant values for the password. Click Next.

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

- s. In the Failure Isolation Support window, select Do Not use Intelligent Platform Management Interface (IPMI). For information on enabling IPMI, see wiki article, <http://tinyurl.com/7j373ov>
- t. In the Privileged Operating Systems Groups window, select:
  - 1. asmdba for Oracle ASM DBA (OSDBA for ASM) Group
  - 2. asmoper for Oracle ASM Operator (OAOPER for ASM) Group
  - 3. asmadmin for Oracle ASM Administrator (OSASM) Group
- u. In the Installation Location window, specify the values of your Oracle Base and Software Location as configured within the Dell Oracle utilities RPM.  
**NOTE:** The default locations used within the Dell Oracle utilities RPM are:
  - 1. Oracle Base - /u01/app/grid
  - 2. Software Location - /u01/app/11.2.0/grid
  - 1.
- v. In the Create Inventory window, specify the location for your Inventory Directory. Click Next.  
**NOTE:** The default location based on the Dell Oracle utilities RPM for Inventory Directory is /u01/app/oraInventory
- w. In the Perform Prerequisite Checks window, check the overall status of all the prerequisites. If any of the prerequisites fail and have the status as Fixable, click Fix & Check Again button and execute the runfixup.sh script provided by the Oracle Universal Installer (OUI).  
  
**NOTE:** If there are other prerequisites that contain the status Error, repeat step 23 else select Ignore All, if proper requirements have been met and Error status still persists after all changes have been fixed.
- x. In the Summary window, select Install.
- y. After the installation is complete, the Execute Configuration Scripts wizard is displayed. Complete the instructions in the wizard and click Ok.
- z. In the Finish window, click Close.

## Configuring Shared Oracle Home using ACFS

The following steps are only applicable for the first node of your cluster environment, unless otherwise specified:

1. Login as root and type: `xhost +`
2. Login as grid user and run the asmca utility by typing:  
`$(GRID_HOME)/bin/asmca`
3. In the ASM Configuration Assistant window, select the Disk Groups tab, click Create, and perform the following steps:
  - a. Enter a name of the disk group. For example, ORAHOME.
  - b. Select the External Redundancy, and then select the ASM stamped disk that you want to use for the shared database home.
2. NOTE: If no candidate disks are displayed, click Change Discovery Path and enter `ORCL:*or/dev/oracleasm/disks/*`
3. NOTE: Ensure that you have marked your Oracle ASM disks. For more information, see "Using ASMLib to Mark the Shared Disks as Candidate Disks".
4. Click Ok.
5. Right-click the disk group you have created for the shared Oracle home, and select Create ACFS for Database Home.
6. The Create ACFS Hosted Database Home Volume option is displayed.
  - a. Enter name for the volume (for example, ORAHOME.)
  - b. Enter name for the mount point for Database Home (for example `/u01/app/oracle/acfsorahome`).
  - c. Enter the Database Home Size (must equal at least 20 GB).
  - d. Enter the name of the Database Home Owner. (for example: oracle).
  - e. Enter the name of the Database Home Owner Group (for example: oinstall).
  - f. Click Ok.
7. As root, run the `acfs_script.sh` mentioned in the RUN ACFS Script window. This automounts the new ACFS Home on all nodes.
8. Click Close to exit ACFS script window.

## Installing Oracle 11g R2 Database Software

The following steps are for node 1 of your cluster environment, unless otherwise specified.

1. Log in as root and type: `xhost +`.
2. Mount the Oracle Database 11gR2 media.
3. Log out as root user and log in as Oracle user by typing:  
`su - oracle`
4. Run the installer script from your Oracle database media:  
`<CD_mount>/runInstaller`
5. In the Configure Security Updates window, enter your My Oracle Support credentials to receive security updates, else click Next.
6. In the Download Software Updates window, enter your My Oracle Support credentials to download patch updates available after the initial release. If you choose not to update at this time, select Skip software updates and click Next.
7. In the Select Installation Option window, select Install database software only.
8. In the Grid Installation Options window:
  - a. Select Oracle Real Application Clusters database installation and select all the nodes by clicking the Select All button.
  - b. Click SSH Connectivity and configure your passwordless SSH connectivity by entering the OS Password for the oracle user and selecting Setup. Click Ok and click Next to go the next window.

**NOTE:** The default password set by the dell-validated and dell-oracle-utilities rpms is oracle for both the grid user and oracle user.

9. In the Select Product Languages window, select English as the Language Option and click Next.
10. In the Select Database Edition window, select Enterprise Edition and click Next.
11. In the Installation Location window,
  - a. Specify the location of your Oracle Base configured within the Dell oracle utilities RPM.
  - b. Enter the ACFS shared Oracle home address for Software Location.

**NOTE:** The default locations used within the Dell Oracle utilities RPM are as follows:

- c. Oracle Base—`/u01/app/oracle`
  - d. Software Location—`/u01/app/oracle/product/11.2.0/db_1`
12. In the Privileged Operating System Groups window, select dba for Database Administrator (OSDBA) Group and asmoper for Database Operator (OSOPER) Group and click Next.
  13. In the Perform Prerequisite Checks window, check the overall status of all the prerequisites.

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

- a. If any prerequisites fail and have the status as Fixable, click the Fix & Check Again button.
- b. Execute the runfixup.sh script provided by the Oracle OUI.

**NOTE:** If there are other prerequisites that display status Error, repeat step 13, If the Error status still persists after all changes have been fixed, select Ignore All.

14. In the Summary window, select Install.
15. On completion of the installation process, the Execute Configuration scripts wizard is displayed. Follow the instructions in the wizard and click Ok.

**NOTE:** Root.sh should be run on one node at a time.

16. In the Finish window, click Close.

## Creating Diskgroups using ASM Configuration Assistant (ASMCA)

This section contains procedures to create the ASM disk group for the database files and Flashback Recovery Area (FRA).

1. Log in as grid user.
2. Start the ASMCA utility by typing:  
`$<GRID_HOME>/bin/asmca`
3. In the ASM Configuration Assistant window, select the Disk Groups tab.
4. Click Create.
5. Enter the appropriate Disk Group Name, for example: DBDG.
6. Select External for Redundancy.
7. Select the appropriate member disks to be used to store your database files, for example:  
`/dev/mapper/dbp1, /dev/mapper/db2p1`
8. **NOTE:** If no candidate disks are displayed, click Change Discovery Path and type:  
`/dev/mapper/*`
9. Click Ok to create and mount the disks.
10. Repeat step 4 to step 8 to create another disk group for your Flashback Recovery Area (FRA). Click Exit to exit the ASM Configuration Assistant.

## Creating Database using DBCA

The following steps are applicable for node 1 of your cluster environment, unless otherwise specified:

1. Login as oracle user.
2. From `$(ORACLE_HOME)`, run the DBCA utility by typing:  
`$(ORACLE_HOME)/bin/dbca &`
3. In the Welcome window, select Oracle Real Application Cluster Database and click Next.
4. In the Operations window, select Create Database, and click Next.
5. In the Database Templates window, select Custom Database, and click Next.
6. In the Database Identification window:
  - a. Select Admin-Managed for Configuration Type.
  - b. Enter appropriate values for Global Database Name and SID Prefix.
  - c. In the Node Selection list box, select All Nodes.
  - d. Click Next.
7. In the Management Options window, select the default values and click Next.
8. In the Database Credentials window, enter the appropriate credentials for your database.
9. In the Database File Location window, select:
  - a. Automatic Storage Management (ASM) for Storage Type.
  - b. Use Oracle-Managed Files for Storage Location.
  - c. Browse to select the ASM disk group that you created to store the database files (DBDG)for Database Area.
10. In the Recovery Configuration window:
  - a. Select Specify Flash Recovery Area.
  - b. Browse and select the ASM disk group that you created for Flash Recovery Area.
  - c. Enter a value for Flash Recovery Area Size.
  - d. Select Enable Archiving.
  - e. Click Next.
11. In the Database Content window, click Next.
12. In the Initialization Parameters window:
  - a. Select Custom.
  - b. For the Memory Management section, select Automatic Shared Memory Management.
  - c. Specify appropriate values for the SGA Size and PGA Size.

- d. Click Next.
13. In the Database Storage window, click Next.
14. In the Creation Options window, click Finish.
15. In the Summary window, click Ok to create database.
16. Click Exit on the Database Configuration Assistant window after the database creation is complete.

## Troubleshooting / Known Issues

1. Currently there is one known issue with the installation of Oracle Database version 11.2.0.3 64 bit with regards to reporting that packages “elfutils-libelf-devel-0.9.7” and “pdksh-5.2.14” are missing. The My Oracle Support article ID is **1454982.1**

The cause of the issue above is related to how the Oracle Universal Installer (OUI) invokes the rpm command to see what version of the Linux distribution is currently being used. The result of the OUI query returns no output thus using the default Linux distribution “OEL4” and uses those preexisting prerequisite checks. The issue is easily fixed by modifying the “cvu\_config” file under directory <path>/database/stage/cvu/cv/admin. Within the cvu\_config file, look for the parameter labeled “CV\_ASSUME\_DISTID=OEL4” and change it to “CV\_ASSUME\_DISTID=OEL6”. Once changed, follow the steps above to invoke the runInstaller. This will guarantee now that the OUI will perform the correct prerequisite checks for your RHEL 6.x distribution.

**NOTE:** Oracle Linux 6.x and RHEL 6.x prerequisite checks are the same.

2. SQLplus error on libnnz11.so: cannot restore segment prot after reloc  
My Oracle Support article ID: 454196.1

The cause of the issue is incorrect context when SELinux is set to enforcing. In order to keep RHEL6 with SELinux set to enforcing, one must change the context of the affected Oracle libraries with the following command:

Example of one Oracle affected library file:

```
/usr/sbin/semanage fcontext -a -t textrel_shlib_t $ORACLE_HOME/lib/libnnz11.so  
/sbin/restorecon libnnz11.so
```

Solution 2:

Set SELinux to “Permissive” mode. Setting SELinux to permissive mode will allow SELinux to continue to run and log, but it will not deny any violations. Oracle recommends setting SELinux into permissive mode until Development resolves the issue above.

To change the SELinux mode immediately:

```
[root@r820-node1 ~]$ getenforce  
Enforcing  
[root@r820-node1 ~]# setenforce 0
```

## Deploying Oracle RAC 11g R2 Database on Red Hat Enterprise Linux 6.x

```
[root@r820-node1 ~]# getenforce
```

```
Permissive
```

To change SELinux mode persistently:

```
[root@r820-node1 ~]# vim /etc/sysconfig/selinux
```

Change the “SELINUX” variable from “enforcing” to “permissive” and save the file.

## Acknowledgements

Adam and I would like to thank those who contributed sections of the deployment guide. Thanks to Naveen Iyengar, Phani MV, and Balamurugan B.