

Dell PowerEdge Systems
Oracle 11g R2 Database on
Enterprise Linux x86_64
Getting Started Guide



Notes, Cautions, and Warnings



NOTE: A NOTE indicates important information that helps you make better use of your computer.



CAUTION: A CAUTION indicates potential damage to hardware or loss of data if instructions are not followed.



WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Information in this publication is subject to change without notice.

© 2011 Dell Inc. All rights reserved.

Reproduction of these materials in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: Dell™, the DELL logo, PowerEdge™, are trademarks of Dell Inc. Red Hat Enterprise Linux® and Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® is a registered trademark and SUSE™ is a trademark of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates.

Other trademarks and trade names may be used in this publication to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

Contents

1	Overview	5
	Before You Begin	5
	Hardware Requirements	5
	Network Requirements	5
	Operating System Requirements	6
2	Preparing Nodes for Oracle Installation	7
	Attaching to RHN/ULN Repository.	7
	Installing the Dell Validated RPM.	9
	Installing the Dell Oracle Utilities RPM.	10
	Oracle Software Binary Location.	11
	Setting up the Network	11
	Public Network	11
	Private Network	12
	IP Address and Name Resolution Requirements.	16
3	Preparing Shared Storage for Oracle RAC Installation	19
	Partitioning the Shared Disk	19
	Adjusting the Stripe Element Size on a	

	Primary Partition	22
	Installing and Configuring ASMLib	25
	Using ASMLib to Mark the Shared Disks as Candidate Disks	25
4	Installing Oracle 11g R2 Grid Infrastructure	29
	Before You Begin.	29
	Configure the System Clock Settings for All Nodes	29
	Configuring Node One	30
5	Configuring Shared Oracle Home for Database Binary Using ACFS	33
6	Installing Oracle 11g R2 Database (RDBMS) Software	35
7	Creating Diskgroup Using ASM Configuration Assistant (ASMCA)	37
8	Creating Database Using DBCA	39

Overview

This document applies to Oracle Database 11g R2 running on Red Hat Enterprise Linux 5.x AS x86_64 or Oracle Enterprise Linux 5.x AS x86_64.

Before You Begin

Hardware Requirements

- Oracle requires 1.5 gigabytes (GB) of physical memory.
- Swap space must be equal to the amount of RAM allocated to the system.
- 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 dell.com/oracle.

Table 1-1 describes the disk space required for an Oracle installation.

Table 1-1. Minimum Disk Space Requirements

Software Installation Location	Size Required
Grid Infrastructure home	4.5 GB of space
Oracle Database home	4 GB of space
Shared storage disk space	Sizes of database and Flashback Recovery

Network Requirements

- It is recommended that you ensure each node contains at least three network interface cards (NICs). One NIC for public network and two NICs for private network to ensure high availability of the Oracle Real Application Clusters (RAC).
- Public and private interface names must be the same on all nodes. For example, if *eth0* is used as the public interface on node one, all other nodes require *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 (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 5.x AS x86_64
- Oracle Linux 5.x AS x86_64

Preparing Nodes for Oracle Installation

Attaching to RHN/ULN Repository



NOTE: The documentation provided below discusses how to set up a local yum repository using your operating system installation media. If you would like to connect to the Red Hat Network (RHN)/Unbreakable Linux Network (ULN) channels, see the appropriate documentation. For Red Hat, see redhat.com/red_hat_network. For information relating to ULN network, see linux.oracle.com.

The recommended configuration is to serve the files over http using an Apache server (package name: `httpd`).

This section discusses hosting the repository files from a local filesystem storage. While other options to host repository files exist, they are outside of the scope of this document. It is highly recommended to use local filesystem storage for speed and simplicity of maintenance.

- 1 Mount the DVD image by either using the physical media or by using the ISO image.
 - Using the physical media: Insert the DVD into the server and it should auto-mount into the `/media` directory.
 - Using the ISO image: Run the following command as root, substituting the path name of your ISO image for the field `myISO.iso`:

```
mkdir /media/myISO  
mount -o loop myISO.iso /media/myISO
```

- 2 To install and configure the http daemon, configure the machine that will host the repository for all other machines to use the DVD image locally. Create the file `/etc/yum.repos.d/local.repo` and enter the following:

```
[local]

name=Local Repository

baseurl=file:///media/myISO/Server

gpgcheck=0

enabled=0
```

- 3 Install the Apache service daemon using the following command which temporarily enables the local repository for dependency resolution:


```
yum -y install httpd --enablerepo=local
```

After the Apache service daemon is installed, start the service and set it to start up for us next time we reboot. Run the following commands as root:

- `service httpd start`
- `chkconfig httpd on`

To use Apache to serve out the repository, copy the contents of the DVD into a published web directory. Run the following commands as root (make sure to switch myISO with the name of your ISO) command:

- `mkdir /var/www/html/myISO`
- `cp -R /media/myISO/* /var/www/html/myISO`

 **NOTE:** The command `createrepo` is used for creating custom repositories, but it is not required as the DVD already holds the repository information.

- This step is only necessary if you are running SELinux on the server that hosts the repository. Run the following command as root to restore the appropriate SELinux context to the copied files:
`restorecon -Rvv /var/www/html/.`
- The final step is to gather the DNS name or IP of the server that is hosting the repository. The DNS name or IP of the hosting server will be used to configure your yum repository repo file on the client server.

The following is an example of the configuration using the RHEL 5.x Server media in the configuration file

```
/etc/yum.repos.d/myRepo.repo
```

```
[myRepo]
```

```
name=RHEL5.5 DVD
```

```
baseurl=
```

```
http://reposerter.mydomain.com/RHEL5_5/Server
```

```
enabled=1
```

```
gpgcheck=0
```



NOTE: Replace reposerter.mydomain.com with your server's DNS name or IP address.



NOTE: You can also place the configuration file on the server hosting the repository for all other servers such that it can also use the repository as a more permanent solution to what was done in step 2.

Installing the Dell Validated RPM

Once your nodes have attached to the appropriate yum repository, install the Dell Validated RPM package. The Dell Validated RPM package automates certain pieces of the installation process required for the installation of Oracle RAC.

The process to install the Dell Validated RPM package is as follows:

- 1 Download the latest Dell Oracle Deployment tar file from http://en.community.dell.com/dell-groups/enterprise_solutions/m/oracle_db_gallery/default.aspx




NOTE: The filename will follow the convention: Dell-Oracle-Deployment-OS version-year-month.tar, for example: Dell-Oracle-Deployment-Lin-2011-07.tar

- 2 Copy the Dell Oracle Deployment tar file to a working directory of all your cluster nodes.
- 3 To go to your working directory, enter the following command:

```
# cd </working/directory/path>
```
- 4 Untar the Dell-Oracle-Deployment release using the command:

```
# tar -xvf Dell-Oracle-Deployment-o-y-m.tar
```

 **NOTE:** Where, *o* is the operating system version, *y* is the year, and *m* is the month of the tar release.

- 5 Change directory to *Dell-Oracle-Deployment-o-y-m*
- 6 Install the Dell Validated RPM package on all your cluster nodes using the following command:


```
# yum localinstall dell-validated* --nogpgcheck
```

Installing the Dell Oracle Utilities RPM

The Dell Oracle utilities RPM is designed to do the following Dell and Oracle recommended settings:


- Create Grid Infrastructure directories, set ownership, and permissions.
- Create grid user.
- Create Oracle Database (RDBMS) directories, set ownership, and permissions.
- Create the Oracle base directories, set ownership, and permissions.
- Set pam limits within (*/etc/pam.d/login*).
- Setup */etc/profile*.
- Set SELinux to **Disabled**.
- Install the Dell PowerEdge system component drivers if applicable.
- Set kernel parameters.

The process to install the Dell Oracle utilities RPM is as follows:

- 1 Download the latest Dell Oracle Deployment tar file from http://en.community.dell.com/dell-groups/enterprise_solutions/m/oracle_db_gallery/default.aspx
 **NOTE:** The filename will follow the convention: Dell-Oracle-Deployment-OS version-year-month.tar, for example: Dell-Oracle-Deployment-Lin-2011-07.tar
- 2 Copy the Dell Oracle Deployment tar file to a working directory of all your cluster nodes.
- 3 To go to your working directory, type the following command:

```
# cd </working/directory/path>
```

- 4 To Untar the Dell-Oracle-Deployment release type the following command: `# tar -xvf Dell-Oracle-Deployment-o-y-m.tar`


 **NOTE:** Where, *o* is the operating system version, *y* is the year, and *m* is the month of the tar release.

- 5 Change directory to *Dell-Oracle-Deployment-o-y-m*
- 6 To install the Dell Oracle utilities RPM package on all your cluster nodes type the following command:

```
# yum localinstall dell-oracle-utilities* --nogpgcheck
```

- 7 Once the rpm is installed, run the *dodeploy* script to set up the environment as follows: `# dodeploy -g -r 11gR2`

For more information about the Dell Oracle utilities RPM and its options, check the man pages using the command: `# man 8 dodeploy`


 **NOTE:** The Dell-Oracle-Deployment tar contains the latest supported drivers provided from our Software Deliverable List (SDL). Consult the README file found within the Dell-Oracle-Deployment tar for installation instructions of the latest drivers.

Oracle Software Binary Location

The Oracle software binaries should be located on node one of your cluster. It is important to note that starting with Oracle 11g R2 (11.2.0.2), Oracle Database patch sets are full installation of the Oracle software. For more information on how this impacts future Oracle deployments, see My Oracle Support note: *1189783.1 Important Changes to Oracle Database Patch Sets Starting with 11.2.0.2*.


Setting up the Network

Public Network

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

To configure the public network on each node:

- 1 Log in as *root*.
- 2 Edit the network device file `/etc/sysconfig/network-scripts/ifcfg-eth#` where *#* is the number of the network device:

 **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 will be: `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 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) command is working.

Private Network

 **NOTE:** Each of the two NIC ports for the private network must be on separate PCI buses.

The grid infrastructure of Oracle 11gR2 (11.2.0.2) supports IP failover natively using a newly introduced feature 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 the subnet range is already in use, Oracle Grid does 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: When adding additional HAIP addresses (maximum of four) after the installation of Oracle Grid, restart your Oracle Grid environment to make these new HAIP addresses active.

The example below provides step-by-step instructions on enabling redundant interconnect using HAIP on a fresh Oracle 11gR2 (11.2.0.2) 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 and eth2 using a 192.168.0.* subnet.

```
DEVICE=eth1  
  
BOOTPROTO=static  
  
HWADDR=00:1E:C9:4B:72:22  
  
ONBOOT=yes  
  
IPADDR=192.168.0.140  
  
NETMASK=255.255.255.0
```

```
DEVICE=eth2  
  
HWADDR=00:1E:C9:4B:71:24  
  
BOOTPROTO=static  
  
ONBOOT=yes  
  
IPADDR=192.168.0.141
```

```
NETMASK=255.255.255.0
```

- 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.

This step enables redundant inter connectivity 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:192.168.0.140
Bcast:192.168.0.255  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:192.168.0.141
Bcast:192.168.0.255  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 metalink note: 1210883.1 at support.oracle.com.

IP Address and Name Resolution Requirements

The steps below show how to set up your cluster nodes for using Domain Name System (DNS). For information on how to set up cluster nodes using GNS, see the Wiki article http://en.community.dell.com/dell-groups/enterprise_solutions/w/oracle_solutions/1416.aspx.

For a Cluster Using DNS

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

- 1 Configure at least two interfaces on each node, one for the private IP address and one for the public IP address.
- 2 Configure the SCAN NAME 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.

Table 2-1 describes the different interfaces, IP address settings and the resolutions in a cluster.

Table 2-1. 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

To configure the changes required on the cluster nodes for name resolution:

- 1 Configure the `resolv.conf` on all the nodes in the cluster to contain name server entries that are resolvable 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. It is recommended that you 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 RAC Installation

 **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. Similarly, the terms Stripe Element Size and Segment Size both can be used interchangeably.

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 LUNs, each of 1 GB in size for normal redundancy or five LUNs for high redundancy 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 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 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:** It is recommended to use the device mapper multipath 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://en.community.dell.com/dell-groups/enterprise_solutions/w/oracle_solutions/3-storage.aspx

Partitioning the Shared Disk

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

 **CAUTION:** In a system running the Linux operating system, align the disk prior to being written to the Volume/Virtual Disk (VD). Failure to do so will cause all data on the disk to be destroyed.

To use the *fdisk* utility to create a partition and set the alignment:

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


- #> `fdisk -u /dev/<block_device>`
- `fdisk -u /dev/mapper/<multipath_disk>`

Where, *<block_device>* is the name of the block device that you are creating and aligning a partition on.


For example, if the block device is */dev/sdb*, type: `fdisk -u /dev/sdb`

If multiple paths to a shared disk are being used and device mapper is the multipath software. The system displays the following message:

The number of cylinders for this disk is set to 8782.

 **NOTE:** The number of cylinder is larger than 1024, and could in certain setups cause problems with:

- Software that runs at boot time (old versions of LILO)
- Booting the Partitioning software from other operating systems (for example, DOS FDISK, OS/2 FDISK)

 **NOTE:** The value of the number of cylinders in your display message may be different depending on the size of your disk.

- a Command (m for help): n # To create a new partition
- b Command action extended primary partition (1-4): P # To create a primary partition
- c Partition number (1-4): 1
- d First sector (63-xxxxxxx, default 63): <Stripe Element Size or Segment Size in terms of Sectors>


Where the *Stripe Element Seizes*) or the *Segment Size* (SS) is the amount of disk space that is consumed on a single physical disk by a Stripe Element as part of Stripe.

For example, a stripe that contains 256 KB of disk space and has 64 KB

of data residing on each disk in the stripe. Here the stripe element size is 64 KB and the stripe size is 256 KB.

Use the following formula to set the value above:

Stripe Element Size in Sectors = Stripe Element Size in KB * 2
First Sector = Stripe Element Size in Sectors

 **NOTE:** The above formula assumes that 1 Sector = 512 Bytes or 0.5 KB.

Set the above value to the following, if the SES/SS was left at the Storage Controller's default value:

- For Dell PowerVault MD30xx/MD30xxi, set First sector to: 128 (default 64 KB * 2)
- For Dell PowerVault MD32xx/MD32xxi, set First sector to: 256 (default 128 KB * 2)
- For Dell EqualLogic PS-Series, set First sector to: 128 (default 64 KB * 2)

If the SES/SS for the disk/Volume/VD is set to a non-default value of the storage array, for example 512KB in case of MD32xx, set the First sector value to 1024.

Last sector or +size or +sizeM or +sizeK (1024-xxxxx, default xxxxxx):
<Enter default value or return key> # Default Value so the single partition spans the entire disk

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 for the kernel to be able to re-read the new partition table.

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

```
#> service multipathd restart Or
#> kpartx -a /dev/mapper/<multipath_disk>
```

- 4 Verify that the partition has been aligned by running the one of the following command:

- #> fdisk -ul /dev/<block_device>
- #> fdisk -ul /dev/mapper/<multipath_device>

Where, <block_device> or <multipath_device> is the name of the disk whose partition was aligned.

The following example is of a sample output of the above command on a block device that has been aligned. If your partition is properly aligned then you will see the desired starting sector that you set in step 1 under the Start column against your partition.

```
Disk /dev/mapper/mpath70: 53.6 GB, 53697576960
bytes
```

```
255 heads, 63 sectors/track, 6528 cylinders, total
104878080 sectors
```

```
Units = sectors of 1 * 512 = 512 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/mapper/mpat		1024	104872319	52436096	83	Linux
h70p1						


- 5 Reboot the system if your newly created and aligned partition is not displayed properly.

Adjusting the Stripe Element Size on a Primary Partition

To use the *fdisk* utility to adjust a disk partition, perform the following steps:



NOTE: This article assumes that the disk to be aligned already contains a single primary partition. If you want to create a primary partition, follow the steps in "Partitioning the Shared Disk" on page 19.

 **CAUTION: In a system running the Linux operating system, align the partition table before the data is written to the Volume. Failure to follow the precaution could lead to all the data on the Volume to be destroyed.**

At the command prompt, type:

1 #> fdisk -u /dev/<block_device>

where <block_device> is the name of the block device that you are adjusting. For example, if the block device is /dev/mapper/db, type:
fdisk /dev/mapper/db.

The system displays the following message:

The number of cylinders for this disk is set to 8782. There is nothing wrong with that, but this is larger than 1024, and could in certain setups cause problems with:

1)software that runs at boot time (e.g., old versions of LILO)


2)booting and partitioning software from other OSs (e.g., DOS FDISK, OS/2 FDISK)

a. Command (m for help): x # To enter the expert Mode


b. Expert command (m for help): b # To move beginning of data in a partition

c. Partition number (1-4): 1 # The partition number to be aligned

d. New beginning of data (128-xxxxx, default 128): 128

 **NOTE:** 1 block = 512 bytes; 128 blocks * 512 bytes = 64KB

e. Expert command (m for help): w # write

 **NOTE:** 128 blocks/64 KB is the default Stripe Element Size of EqualLogic PS Series and 256 blocks/256 KB is the default Stripe Element Size of the PowerVault MD 32xx/32xxi Line of storage arrays.

2 Repeat step 1 for all the disks that need to be aligned.

- 3 Run the following command to re-scan all the partitions on node one when using device mapper:

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

On all other nodes run:

```
#> kpartx -l /dev/mapper/devicename>
```



NOTE: If your device name does not end in "p1", reboot your system. Proper name convention would append a "p1" and display as /dev/mapper/ACFSp1.

- 4 Verify that the partition has been aligned by running the following command:

```
#> fdisk -ul /dev/<block_device>
```

Where, *<block_device>* is the name of the block device.

The following example of a sample output of the command executed on a block device that has been aligned. If your partition is properly aligned then 128 is displayed under the *Start* column against your partition.

```
Disk /dev/mapper/mpath70: 53.6 GB, 53697576960 bytes
```

```
255 heads, 63 sectors/track, 6528 cylinders, total 104878080 sectors
```

```
Units = sectors of 1 * 512 = 512 bytes
```


Device	Boot	Start	End	Blocks	Id	System
/dev/mapper/mpat		128	104872319	52436096	83	Linux
h70p1						


- 5 Reboot the system if your newly created and aligned partition is not displayed.

Installing and Configuring ASMLib

- 1 Use oracle.com/technetwork/server-storage/linux/downloads/rhel5-084877.html to download the following files:


- oracleasm-support
- oracleasmlib
- oracleasm

 **NOTE:** If your current OS distribution is Oracle Linux, you can obtain the software from the Unbreakable Linux Network using ULN.

 **NOTE:** Download the latest versions of *oracleasm-support* and *oracleasmlib* but the version of *oracleasm* must match the current kernel used in your system. Check this information issuing the command `uname -r`.

- 2 Enter the following command as *root*:


```
rpm -Uvh oracleasm-support-* \  
oracleasm-lib-* \  
oracleasm-$(uname -r)-*
```

 **NOTE:** Replace *** by the correct version numbers of the packages or you can leave them in place of the command ensuring that there are no multiple versions of the packages in the shell's current working directory.

Using ASMLib to Mark the Shared Disks as Candidate Disks

- 1 To configure ASM use the init script that comes with the *oracleasm-support* package. The recommended method is to run the following command as *root*:

```
# /usr/sbin/oracleasm configure -i
```


 **NOTE:** Oracle recommends using the *oracleasm* command found under `/usr/sbin`. The `/etc/init.d` path has not been deprecated, but the *oracleasm* binary provided by Oracle in this path is used for internal purposes.

```
Default user to own the driver interface []: grid
```

```
Default group to own the driver interface []:  
asmadmin
```


```
Start Oracle ASM library driver on boot (y/n) [n]:  
y
```

Fix permissions of Oracle ASM disks on boot (y/n)
[y]: y

 **NOTE:** In this setup the default user is set to *grid* and the default group is set to *asmadmin*. Ensure that the oracle user is part of the asmadmin group. You can do so by using the dell-validated and dell-oracle-utilities rpms.


The boot time parameters of the Oracle ASM library are configured and a sequential text interface configuration method is displayed.

- 2 Set the ORACLEASM_SCANORDER parameter in `/etc/sysconfig/oracleasm`

 **NOTE:** When setting the ORACLEASM_SCANORDER to a value, specify a common string associated with your device mapper pseudo device name. For example, if all the device mapper device had a prefix string of the word "asm", (`/dev/mapper/asm-ocr1`, `/dev/mapper/asm-ocr2`), populate the ORACLEASM_SCANORDER parameter as: `ORACLEASM_SCANORDER="asm"`. This would ensure that oracleasm will scan these disks first.


- 3 Set the ORACLEASM_SCANEXCLUDE parameter in `/etc/sysconfig/oracleasm` to exclude non-multipath devices.


For example: `ORACLEASM_SCANEXCLUDE=<disks to exclude>`

 **NOTE:** If we wanted to ensure to exclude our single path disks within `/dev/` such as `sda` and `sdb`, our ORACLEASM_SCANEXCLUDE string would look like: `ORACLEASM_SCANEXCLUDE="sda sdb"`

- 4 To create ASM disks that can be managed and used for Oracle database installation, run the following command as root:

```
/usr/sbin/oracleasm createdisk DISKNAME  
/dev/mapper/diskpartition
```

 **NOTE:** The fields DISKNAME and `/dev/mapper/diskpartition` should be substituted with the appropriate names for your environment respectively.

 **NOTE:** It is highly recommended to have all of your Oracle related disks to be within Oracle ASM. This includes your OCR disks, voting disks, database disks, and flashback recovery disks.

- 5 Verify the presence of the disks in the ASM library by running the following command as root:

```
/usr/sbin/oracleasm listdisks
```

All the instances of DISKNAME from the previous command(s) are displayed.

To delete an ASM disk, run the following command:

```
/usr/sbin/oracleasm deletedisk DISKNAME
```

- 6** To discover the Oracle ASM disks on other nodes in the cluster, run the following command on the remaining cluster nodes:

```
/usr/sbin/oracleasm scandisks.
```


Installing Oracle 11g R2 Grid Infrastructure

This section gives you the installation information for Oracle 11g R2 grid infrastructure for 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 (CTSS) which is built in Oracle 11g R2. To enable CTSS, disable the operating system network time protocol daemon (ntpd) service using the commands in the following 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.

- 1 Log in as *root*.
- 2 If you are not in a graphical environment, start the **X Window System** by typing: `startx`
- 3 Open a terminal window and type: `xhost +`
- 4 Mount the Oracle Grid Infrastructure media.
- 5 Log in as *grid user*, for example: `su - grid`.
- 6 Type the following command to start the **Oracle Universal Installer**:
`<CD_mountpoint>/runInstaller`
- 7 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**.
- 8 In the **Select Installation Option** window, select **Install and Configure Grid Infrastructure for a Cluster** and click **Next**.
- 9 In the **Select Installation Type** window, select **Advanced Installation** option, and click **Next**.
- 10 In the **Select Product Languages** window, select **English**, and click **Next**.
- 11 In the **Grid Plug and Play Information** window, enter the following information:
 - **Cluster Name**—Enter a name for your cluster.
 - **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" on page 16. For details on enabling GNS, see the Wiki article: en.community.dell.com/dell-groups/enterprise_solutions/w/oracle_solutions/1416.aspx.
 - **SCAN Port**—Retain the default port of 1521.
 - **Configure GNS**—Uncheck this option.
 - Click **Next**.


- 12 In the **Cluster Node Information** window, click **Add** to add additional nodes that must be managed by the **Oracle Grid Infrastructure**.
 - Enter the public **Hostname** information
 - Enter the **Virtual IP name**
 - Repeat step 12 for each node within your cluster
- 13 Click **SSH Connectivity** and configure your password less 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.
- 14 Click **Ok** and then click **Next** to go to the next window.
- 15 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**.
- 16 In the **Storage Option Information** window, select **Automatic Storage Management (ASM)** and click **Next**.
- 17 In the **Create ASM Disk Group** window, enter the following information:
 - **ASM diskgroup**—Enter a name, for example: `OCR_VOTE`
 - **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 `ORCL:*` or `/dev/oracleasm/disks/*`. Ensure that you have marked your Oracle ASM disks, for more informations see, "Using ASMLib to Mark the Shared Disks as Candidate Disks" on page 25.
- 18 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**.
- 19 In the **Failure Isolation Support** window, select **Do Not use Intelligent Platform Management Interface (IPMI)**.
For information on enabling IPMI, see Wiki article, http://en.community.dell.com/dell-groups/enterprise_solutions/w/oracle_solutions/1414.aspx.


- 20 In the **Privileged Operating Systems Groups** window, select:
 - *asmdba* for Oracle ASM DBA (OSDBA for ASM) Group
 - *asmoper* for Oracle ASM Operator (OAOPER for ASM) Group
 - *asmadmin* for Oracle ASM Administrator (OSASM) Group
- 21 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:
 - Oracle Base - **/u01/app/grid**
 - Software Location - **/u01/app/11.2.0/grid**
- 22 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**.
- 23 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** and execute the **runfixup.sh** script provided by the Oracle Universal Installer (OUI).
 **NOTE:** For 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.
- 24 In the **Summary** window, select **Install**.
After the installation is complete, the **Execute Configuration Scripts** wizard is displayed.
- 25 Complete the instructions in the wizard and click **Ok**.
- 26 In the **Finish** window, click **Close**.

Configuring Shared Oracle Home for Database Binary Using ACFS

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

- 1 Log in as *root* and type: `xhost +`
- 2 Log in 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:
 - Enter a name of the disk group. For example, ORAHOME.
 - Select the **External Redundancy**, and then select the ASM stamped disk that you want to use for the shared database home.


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

 **NOTE:** Ensure that you have marked your Oracle ASM disks. For more information, see "Using ASMLib to Mark the Shared Disks as Candidate Disks" on page 25.
- 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 In the **Create ACFS Hosted Database Home Volume** screen:
 - Enter name for the volume (for example, ORAHOME.)
 - Enter name for the mount point for Database Home (for example, `/u01/app/oracle/acfsorahome`).
 - Enter the Database Home Size (must be at least 20 GB).
 - Enter the name of the Database Home Owner. (for example: `oracle`).
 - Enter the name of the Database Home Owner Group (for example: `oinstall`).


- 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 (RDBMS) Software


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


- 1 Log in as *root* and type: `xhost +`.
 - 2 Mount the Oracle Database 11g R2 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:
 - Select **Oracle Real Application Clusters database installation** and select all the nodes by clicking the **Select All** button.
 - 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,
 - Specify the location of your **Oracle Base** configured within the Dell oracle utilities RPM.
 - Enter the ACFS shared Oracle home address for **Software Location**.

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

 - Oracle Base—**/u01/app/oracle**.
 - 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.
 - If any prerequisites fail and have the status as *Fixable*, click the **Fix & Check Again** button.
 - 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:** Run **Root.sh** on one node at a time.
- 16 In the **Finish** window, click **Close**.

Creating Diskgroup 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: ORCL:DB1, ORCL:DB2. enter ORCL:* or /dev/oracleasm/disks/*

 **NOTE:** If no candidate disks are displayed, click **Change Discovery Path** and type: ORCL:* or /dev/oracleasm/disks/*

 **NOTE:** Please ensure you have marked your Oracle ASM disks. For more information, see "Using ASMLib to Mark the Shared Disks as Candidate Disks" on page 25.


- 8 Click **Ok** to create and mount the disks.
- 9 Repeat step 4 to step 8 to create another disk group for your Flashback Recovery Area (FRA).


 **NOTE:** Make sure that you label your FRA disk group differently than your database disk group name. For labeling your Oracle ASM disks, see "Using ASMLib to Mark the Shared Disks as Candidate Disks" on page 25

- 10 Click **Exit** to exit the ASM Configuration Assistant.

Creating Database Using DBCA

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

- 1 Log in 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**.
-  **NOTE:** For more information on Policy-Managed configuration, see the Wiki article http://en.community.dell.com/dell-groups/enterprise_solutions/w/oracle_solutions/1418.aspx.
- 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:
 - **Automatic Storage Management (ASM)** for **Storage Type**.
 - **Use Oracle-Managed Files** for **Storage Location**.
 - 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:
 - Select **Custom**.
 - For the **Memory Management** section, select **Automatic Shared Memory Management**.
 - Specify appropriate values for the **SGA Size** and **PGA Size**.
 - 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.
 **NOTE:** Database creation can take some time to complete.
- 16** Click **Exit** on the **Database Configuration Assistant** window after the database creation is complete.