

Table of Contents

<u>Rac11gR2OnSolaris</u>	1
<u>1. Introduction</u>	1
<u>1.1. Overview of new concepts in 11gR2 Grid Infrastructure</u>	1
<u>1.1.1. SCAN</u>	1
<u>1.1.2. GNS</u>	1
<u>1.1.3. OCR and Voting on ASM storage</u>	1
<u>1.1.4. Passwordless automatic SSH connectivity</u>	1
<u>1.1.5. Intelligent Platform Management interface (IPMI)</u>	1
<u>1.1.6. Time sync</u>	2
<u>1.1.7. Clusterware and ASM share the same Oracle Home</u>	2
<u>1.1.8. ACFS/ADVM</u>	2
<u>1.2. System Requirements</u>	2
<u>1.2.1. Hardware Requirements</u>	2
<u>1.2.2. Network Hardware Requirements</u>	4
<u>1.2.3. IP Address Requirements</u>	4
<u>1.2.4. Installation method</u>	5
<u>2. *Prepare the cluster nodes for Oracle RAC</u>	5
<u>2.1. User Accounts</u>	5
<u>2.2. Networking</u>	6
<u>2.3. Synchronizing the Time on ALL Nodes</u>	7
<u>2.4. Configuring Kernel Parameters</u>	7
<u>2.5. Create the Oracle Inventory Directory</u>	8
<u>2.6. Creating the Oracle Grid Infrastructure Home Directory</u>	8
<u>2.7. Creating the Oracle Base Directory</u>	8
<u>2.8. Creating the Oracle RDBMS Home Directory</u>	9
<u>2.9. Stage the Oracle Software</u>	9
<u>2.10. Check OS Software Requirements</u>	9
<u>3. Prepare the shared storage for Oracle RAC</u>	9
<u>4. Oracle Grid Infrastructure Install</u>	13
<u>4.1. Basic Grid Infrastructure Install (without GNS and IPMI)</u>	13
<u>5. Grid Infrastructure Home Patching</u>	32
<u>6. RDBMS Software Install</u>	32
<u>7. RAC Home Patching</u>	45
<u>8. Run ASMCA to create diskgroups</u>	45
<u>9. Run DBCA to create the database</u>	48

Rac11gR2OnSolaris

1. Introduction

1.1. Overview of new concepts in 11gR2 Grid Infrastructure

1.1.1. SCAN

The single client access name (SCAN) is the address used by all clients connecting to the cluster. The SCAN name is a domain name registered to three IP addresses, either in the domain name service (DNS) or the Grid Naming Service (GNS). The SCAN name eliminates the need to change clients when nodes are added to or removed from the cluster. Clients using SCAN names can also access the cluster using EZCONNECT.

- The Single Client Access Name (SCAN) is a domain name that resolves to all the addresses allocated for the SCAN name. Allocate three addresses to the SCAN name. During Oracle grid infrastructure installation, listeners are created for each of the SCAN addresses, and Oracle grid infrastructure controls which server responds to a SCAN address request. Provide three IP addresses in the DNS to use for SCAN name mapping. This ensures high availability.
- The SCAN addresses need to be on the same subnet as the VIP addresses for nodes in the cluster.
- The SCAN domain name must be unique within your corporate network.

1.1.2. GNS

In the past, the host and VIP names and addresses were defined in the DNS or locally in a hosts file. GNS can simplify this setup by using DHCP. To use GNS, DHCP must be configured in the subdomain in which the cluster resides.

1.1.3. OCR and Voting on ASM storage

The ability to use ASM diskgroups for Clusterware OCR and Voting disks is a new feature in the Oracle Database 11g Release 2 Grid Infrastructure. If you choose this option and ASM is not yet configured, OUI launches ASM configuration assistant to configure ASM and a diskgroup.

1.1.4. Passwordless automatic SSH connectivity

If SSH has not been configured prior the Installation, you can prompt the installer to this for you. The configuration can be tested as well.

1.1.5. Intelligent Platform Management interface (IPMI)

Intelligent Platform Management Interface (IPMI) provides a set of common interfaces to computer hardware and firmware that administrators can use to monitor system health and manage the system.

With Oracle Database 11g Release 2, Oracle Clusterware can integrate IPMI to provide failure isolation support and to ensure cluster integrity. You must have the following hardware and software configured to enable cluster nodes to be managed with IPMI:

- Each cluster member node requires a Baseboard Management Controller (BMC) running

firmware compatible with IPMI version 1.5, which supports IPMI over LANs, and configured for remote control.

- Each cluster member node requires an IPMI driver installed on each node.
- The cluster requires a management network for IPMI. This can be a shared network, but Oracle

recommends that you configure a dedicated network.

- Each cluster node's ethernet port used by BMC must be connected to the IPMI management network.

If you intend to use IPMI, then you must provide an administration account username and password to provide when prompted during installation.

1.1.6. Time sync

Oracle Clusterware 11g release 2 (11.2) requires time synchronization across all nodes within a cluster when Oracle RAC is deployed. To achieve this you should have your OS configured network time protocol (NTP). The new Oracle Cluster Time Synchronization Service is designed for organizations whose Oracle RAC databases are unable to access NTP services.

1.1.7. Clusterware and ASM share the same Oracle Home

The clusterware and ASM share the same home thus we call it Grid Infrastructure home (prior 11gR2 ASM and RDBMS could be installed either in same Oracle home or in separate Oracle homes)

1.1.8.ACFS/ADVM

ADVM (ASM dynamic volume manager) and ACFS (ASM cluster file system) are currently not available for Solaris. For details refer to note:

IS ACFS/ADVM SUPPORTED/CERTIFIED ON SOLARIS SPARC 64 PLATFORM? (Doc ID 973387.1)

1.2. System Requirements

1.2.1. Hardware Requirements

-Physical memory (at least 1.5 gigabyte (GB) of RAM)

/usr/sbin/prtconf | grep "Memory size"

-An amount of swap space equal the amount of RAM

/usr/sbin/swap -s

-Temporary space (at least 1 GB) available in /tmp

df -h /tmp

-A processor type (CPU) that is certified with the version of the Oracle software being installed

-At minimum of 1024 x 786 display resolution, so that Oracle Universal Installer (OUI) displays correctly

-All servers that will be used in the cluster have the same chip architecture, for example, all SPARC processors or all x86 64-bit processors

-Disk space for software installation locations

You will need at least 4.5 GB of available disk space for the Grid home directory, which includes both the binary files for Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) and their associated log files, and at least 4 GB of available disk space for the Oracle Database home directory.

-Shared disk space

An Oracle RAC database is a shared everything database. All data files, control files, redo log files, and the server parameter file (SPFILE) used by the Oracle RAC database must reside on shared storage that is accessible by all the Oracle RAC database instances. The Oracle RAC installation that is described in this guide uses Oracle ASM for the shared storage for Oracle Clusterware and Oracle Database files. The amount of shared disk space is determined by the size of your database.

-Check Operating System Packages:

Software Requirements List for Solaris Operating System (x86 64-Bit) Platforms:

Solaris 10 U6 (5.10-2008.10) or later

Packages and Patches for all installations

Solaris 10

SUNWarc

SUNWbtool

SUNWcsl

SUNWhea

SUNWlibC

SUNWlibm

SUNWlibms

SUNWsprot

SUNWtoo

SUNWi1of (ISO8859-1)

SUNWi1cs (ISO8859-15)

SUNWi15cs

SUNWxwfont

119961-05 or later

119964-14 or later

120754-06 or later

139556-08 or later

139575-03 or later

137104-02 or later

Software Requirements List for Solaris Operating System (SPARC 64-Bit):

Solaris 10 U6 (5.10-2008.10) or later

Packages and Patches for all installations

Solaris 10

SUNWarc

SUNWbtool
SUNWcsl
SUNWhea
SUNWlibC
SUNWlibm
SUNWlibms
SUNWspot
SUNWtoo
SUNWi1of (ISO8859-1)
SUNWi1cs (ISO8859-15)
SUNWi15cs
SUNWxfnt
119963-14 or later (SunOS₂ 5.10: Shared library patch for C++)
120753-06 or later (SunOS₂ 5.10: Microtasking libraries (libmtsk) patch)
139574-03 or later (SunOS₂ 5.10: file crle ldd stings elfdump patch, required for Oracle Clusterware))

1.2.2. Network Hardware Requirements

- Each node has at least two network interface cards (NIC), or network adapters.
- Public interface names must be the same for all nodes. If the public interface on one node uses the network adapter e1000g0, then you must configure e1000g0 as the public interface on all nodes.
- You should configure the same private interface names for all nodes as well. If e1000g1 is the private interface name for the first node, then e1000g1 should be the private interface name for your second node.
- The network adapter for the public interface must support TCP/IP.
- The network adapter for the private interface must support the user datagram protocol (UDP) using high-speed network adapters and a network switch that supports TCP/IP (Gigabit Ethernet or better).
- For the private network, the end points of all designated interconnect interfaces must be completely reachable on the network. Every node in the cluster should be able to connect to every private network interface in the cluster.
- The host name of each node must conform to the RFC 952 standard, which permits alphanumeric characters. Host names using underscores ("_") are not allowed.
- If you follow best practices and implement redundant Network adapters please review MOS Note:1069584.1 "11gR2 Grid Infrastructure Multiple Private Network Adapters"

1.2.3. IP Address Requirements

- A public IP address for each node
- A virtual IP address for each node
- Three single client access name (SCAN) addresses for the cluster.

(Define the SCAN in your corporate DNS (Domain Name Service) You must ask your network administrator to create a single name, that resolves to 3 IP addresses using a round robin algorithm. The IP addresses must be on the same subnet as your public network in the cluster.)



1.2.4. Installation method

This document details the steps installing a 2-node Oracle 11gR2 RAC cluster on Solaris:

-The Oracle Grid Homes binaries are installed on the local disk of each of the RAC nodes.

-The files required by Oracle Clusterware (OCR and Voting disks) are stored in ASM -The installation is explained without GNS and IPMI (additional Information for Installation with GNS and IPMI are explained)

2. *Prepare the cluster nodes for Oracle RAC

The guides include hidden sections, use the  and  image for each section to show/hide the section or you can Expand all or Collapse all by clicking these buttons. This is implemented using the [Twisty Plugin](#) which requires Java Script to be enabled on your browser.

2.1. User Accounts

1. Create OS groups using the command below Enter commands as the root user:

```
#!/usr/sbin/groupadd oinstall  
#!/usr/sbin/groupadd dba  
#!/usr/sbin/groupadd asmadmin  
#!/usr/sbin/groupadd asmdba  
#!/usr/sbin/groupadd asmoper
```

2. Create the users that will own the Oracle software using the commands:

```
#!/usr/sbin/useradd -g oinstall -G asmadmin,asmdba,asmoper -d /home/grid -m grid  
#!/usr/sbin/useradd -g oinstall -G dba,asmdba -d /home/oracle -m oracle
```

Note: you might have to disable autofs temporary to be able to create the home folders:

```
#svcadm disable autofs
```

3. Set the password for the oracle account using the following command. Replace password with your own password.

passwd oracle

Changing password for user oracle.

New UNIX password: **password**

retype new UNIX password: **password**

passwd: all authentication tokens updated successfully.

passwd grid

Changing password for user oracle.

New UNIX password: **password**

retype new UNIX password: **password**

passwd: all authentication tokens updated successfully.

4. Repeat Step 1 through Step 3 on each node in your cluster.

5. OUI can setup passwordless SSH for you, if you want to configure this yourself, refer to Note. 300548.1

2.2. Networking

NOTE: This section is intended to be used for installations NOT using GNS.

1. Determine your cluster name. The cluster name should satisfy the following conditions:

- The cluster name is globally unique throughout your host domain.
- The cluster name is at least 1 character long and less than 15 characters long.
- The cluster name must consist of the same character set used for host names: single-byte alphanumeric characters (a to z, A to Z, and 0 to 9) and hyphens (-).

2. Determine the public host name for each node in the cluster. For the public host name, use the primary host name of each node. In other words, use the name displayed by the hostname command for example: racnode1.

3. Determine the public virtual hostname for each node in the cluster. The virtual host name is a public node name that is used to reroute client requests sent to the node if the node is down. Oracle recommends that you provide a name in the format <public hostname>-vip, for example: racnode1-vip. The virtual hostname must meet the following requirements:

- The virtual IP address and the network name must not be currently in use.
- The virtual IP address must be on the same subnet as your public IP address.
- The virtual host name for each node should be registered with your DNS.

4. Determine the private hostname for each node in the cluster. This private hostname does not need to be resolvable through DNS and should be entered in the /etc/hosts file. A common naming convention for the private hostname is <public hostname>-pvt.

- The private IP should NOT be accessible to servers not participating in the local cluster.
- The private network should be on standalone dedicated switch(es).
- The private network should NOT be part of a larger overall network topology.
- The private network should be deployed on Gigabit Ethernet or better.
- It is recommended that redundant NICs are configured For Solaris either Sun Trunking (OS based) or Sun IPMP (OS based) More information: <<Note: 283107.1>>
- IPMP in general. When IPMP is used for the interconnect: <<Note: 368464.1>>

NOTE: If IPMP is used for public and/or cluster interconnect, critical merge patch 9729439 should be applied to both Grid Infrastructure and RDBMS Oracle homes.

5. Define a SCAN DNS name for the cluster that resolves to three IP addresses (round-robin). SCAN VIPs must NOT be in the /etc/hosts file, it must be resolved by DNS.

6. Even if you are using a DNS, Oracle recommends that you add lines to the /etc/hosts file on each node, specifying the public IP, VIP and private addresses. Configure the /etc/hosts file so that it is similar to the following example:

NOTE: The SCAN VIP MUST NOT be in the /etc/hosts file. This will result in only 1 SCAN VIP for the entire cluster.

```
# cat /etc/hosts
#
# Internet host table
#

# Public IPs

127.0.0.1 localhost
192.168.1.10 public loghost
192.168.1.11 public1

# Private IPs
10.10.10.10 public-priv
10.10.10.11 public1-priv

# VIPS

192.168.1.100 public-vip
192.168.1.101 public1-vip
```

2.3. Synchronizing the Time on ALL Nodes

Ensure that the date and time settings on all nodes are set as closely as possible to the same date and time. Time may be kept in sync with NTP or by using Oracle Cluster Time Synchronization Service (ctssd). For NTP with Solaris 10 the "slewalways yes" option in /etc/inet/ntp.conf should be used. See note 759143.1 for details.

2.4. Configuring Kernel Parameters

We need to set the following kernel parameters to values greater than or equal to the recommended values shown below.

```
set noexec_user_stack=1

set semsys:seminfo_semmni=100

set semsys:seminfo_semmns=1024

set semsys:seminfo_semmsl=256

set semsys:seminfo_semvmx=32767

set shmsys:shminfo_shmmax=4294967296

set shmsys:shminfo_shmmni =100
```

NOTE: You may skip to change the parameter if the default setting in your system is higher than Oracle requirement.

Since we are using Solaris 10, we are not required to make changes to the /etc/system file to implement the System V IPC. Solaris 10 uses the resource control facility for its implementation. However, Oracle recommends that you set both resource control and /etc/system/ parameters. Operating system parameters not replaced by resource controls continue to affect performance and security on Solaris 10 systems. For further information, contact your Sun vendor On all nodes:

```
# prctl -n project.max-sem-ids -v 100 -r -i project user.root
# prctl -n project.max-shm-ids -v 100 -r -i project user.root
# prctl -n project.max-shm-memory -v 4 gb -r -i project user.root
# vi /etc/system
```

```
set noexec_user_stack=1
set semsys:seminfo_semmni=100
set semsys:seminfo_semmns=1024
set semsys:seminfo_semmns=1024
set semsys:seminfo_semmns=1024
set semsys:seminfo_semmns=1024
set semsys:seminfo_semmns=1024
set shmsys:shminfo_shmmax=4294967296
set shmsys:shminfo_shmmni =100
```

```
# init 6
```

NOTE: OUI checks the current settings for various kernel parameters to ensure they meet the minimum requirements for deploying Oracle RAC.

2.5. Create the Oracle Inventory Directory

To create the Oracle Inventory directory, enter the following commands as the root user:

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

2.6. Creating the Oracle Grid Infrastructure Home Directory

To create the Grid Infrastructure home directory, enter the following commands as the root user:

```
# mkdir -p /u01/11.2.0/grid
# chown -R grid:oinstall /u01/11.2.0/grid
# chmod -R 775 /u01/11.2.0/grid
```

2.7. Creating the Oracle Base Directory

To create the Oracle Base directory, enter the following commands as the root user:

```
# mkdir -p /u01/app/oracle
# mkdir /u01/app/oracle/cfgtoollogs --needed to ensure that dbca is able to run after the rdbms installation.
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

2.8. Creating the Oracle RDBMS Home Directory

To create the Oracle RDBMS Home directory, enter the following commands as the root user:

```
# mkdir -p /u01/app/oracle/product/11.2.0/db_1  
# chown -R oracle:oinstall /u01/app/oracle/product/11.2.0/db_1  
# chmod -R 775 /u01/app/oracle/product/11.2.0/db_1
```

2.9. Stage the Oracle Software

It is recommended that you stage the required software onto a local drive on Node 1 of your cluster. Important. Ensure that you use the correct version, either SPARC or x86-64. For the RDBMS software download from OTN:

Oracle Database 11g Release 2 (11.2.0.1.0) for Solaris

For the Grid Infrastructure (clusterware and ASM) software download:

Oracle Database 11g Release 2 Grid Infrastructure (11.2.0.1.0) for Solaris

2.10. Check OS Software Requirements

The OUI will check during the install for missing packages and you will have the opportunity to install them at that point during the prechecks. Nevertheless you might want to validate that all required packages have been installed prior to launching the OUI.

NOTE: check on all nodes that the Firewall is disabled. Disable if needed:

```
#svcadm disable ipfilter
```

3. Prepare the shared storage for Oracle RAC

This section describes how to prepare the shared storage for Oracle RAC. Each node in a cluster requires external shared disks for storing the Oracle Clusterware (Oracle Cluster Registry and voting disk) files, and Oracle Database files. To ensure high availability of Oracle Clusterware files on Oracle ASM, you need to have at least 2 GB of disk space for Oracle Clusterware files in three separate failure groups, with at least three physical disks. Each disk must have at least 1 GB of capacity to ensure that there is sufficient space to create Oracle Clusterware files. Use the following guidelines when identifying appropriate disk devices:

-All of the devices in an Automatic Storage Management disk group should be the same size and have the same performance characteristics.

-A disk group should not contain more than one partition on a single physical disk device.

-Using logical volumes as a device in an Automatic Storage Management disk group is not supported with Oracle RAC.

-The user account with which you perform the installation (oracle) must have write permissions to create the files in the path that you specify. On Solaris 10, you can use format or smc utilities to carve disk or LUNs partitions/slices. It is very important to skip the first Cylinder on the disk to avoid ASM or Oracle Clusterware to overwrite the partition table. So you always start partitioning from cylinder number 3. failing to do so, you will find out after rebooting your machines that data on your disks is erased and Oracle Clusterware will not start and ASM will not be able to recognize any disks. Below I am running the format command from the first solaris node only. This formats the disk with solaris partitions, changes slice 4 to skip the first 3 cylinders and labels the disk.

format

Searching for disks...done

AVAILABLE DISK SELECTIONS:

0. c0d0 <DEFAULT cyl 2607 alt 2 hd 255 sec 63>

/pci@0,0/pci-ide@1,1/ide@0/cmdk@0,0

1. c2t12d0 <DEFAULT cyl 524 alt 2 hd 128 sec 32>

/iscsi/disk@0000iqn.2006-01.com.openfiler%3Atsn.ASM10001,0

2. c2t13d0 <DEFAULT cyl 524 alt 2 hd 128 sec 32>

/iscsi/disk@0000iqn.2006-01.com.openfiler%3Atsn.ASM20001,0

3. c2t14d0 <DEFAULT cyl 524 alt 2 hd 128 sec 32>

/iscsi/disk@0000iqn.2006-01.com.openfiler%3Atsn.ASM30001,0

4. c2t15d0 <DEFAULT cyl 524 alt 2 hd 128 sec 32>

/iscsi/disk@0000iqn.2006-01.com.openfiler%3Atsn.ASM40001,0

Specify disk (enter its number): 1

selecting c2t12d0

[disk formatted]

FORMAT MENU:

disk - select a disk

type - select (define) a disk type

partition - select (define) a partition table

current - describe the current disk

format - format and analyze the disk

fdisk - run the fdisk program

repair - repair a defective sector

label - write label to the disk

analyze - surface analysis

defect - defect list management

backup - search for backup labels

verify - read and display labels

save - save new disk/partition definitions

inquiry - show vendor, product and revision

volname - set 8-character volume name

!<cmd>* - execute *<cmd>*, then return*

quit

format> fdisk

No fdisk table exists. The default partition for the disk is:

a 100% "SOLARIS System" partition

Type "y" to accept the default partition, otherwise type "n" to edit the partition table.

y

format> partition

PARTITION MENU:

0 - change `0' partition

1 - change `1' partition

2 - change `2' partition

3 - change `3' partition

4 - change `4' partition

5 - change `5' partition

6 - change `6' partition

7 - change `7' partition

select - select a predefined table

modify - modify a predefined partition table

name - name the current table

print - display the current table

label - write partition map and label to the disk

!<cmd> - execute <cmd>, then return

quit

partition> 4

Part Tag Flag Cylinders Size Blocks

4 unassigned wm 0 0 (0/0/0) 0

Enter partition id tag[unassigned]:

Enter partition permission flags[wm]:

Enter new starting cyl[0]: 3

Enter partition size[0b, 0c, 3e, 0.00mb, 0.00gb]: 1gb

partition> l

Ready to label disk, continue? y

partition> q

FORMAT MENU:

disk - select a disk

type - select (define) a disk type

partition - select (define) a partition table

current - describe the current disk

format - format and analyze the disk

fdisk - run the fdisk program

repair - repair a defective sector

label - write label to the disk

analyze - surface analysis

defect - defect list management

backup - search for backup labels

verify - read and display labels

save - save new disk/partition definitions

inquiry - show vendor, product and revision

volname - set 8-character volume name

*!*cmd*> - execute *cmd*, then return*

quit format> q

#

Note: do the same for the other disks you want to use with ASM. Enter commands similar to the following on every node to change the owner, group, and permissions on the character raw device file for each disk slice that you want to add to a disk group, where *grid* is the grid infrastructure installation owner, and *asmadmin* is the OSASM group:

chown grid:asmadmin /dev/rdisk/cxydzs4

chmod 660 /dev/rdisk/cxydzs4

Verify the setting with:

ls -lL /dev/rdisk/cxydzs4

In this example, the device name specifies slice 4

4. Oracle Grid Infrastructure Install

4.1. Basic Grid Infrastructure Install (without GNS and IPMI)

As the *grid* user (grid infrastructure software owner) start the installer by running "runInstaller" from the staged installation media.

NOTE:

Be sure the installer is run as the intended software owner, the only supported method to change the software owner is to reinstall.

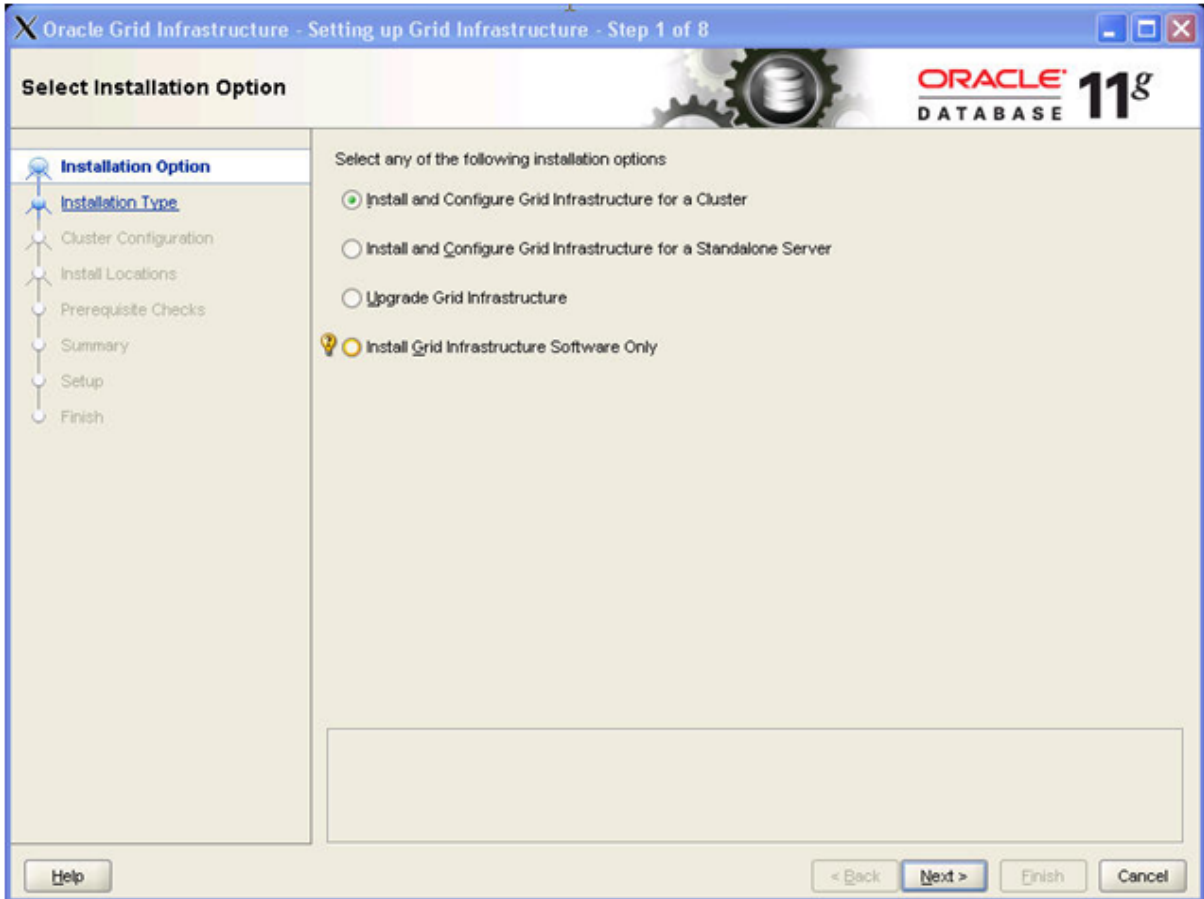
#xhost +

#su - grid

#DISPLAY=<ip address>:0.0; export DISPLAY

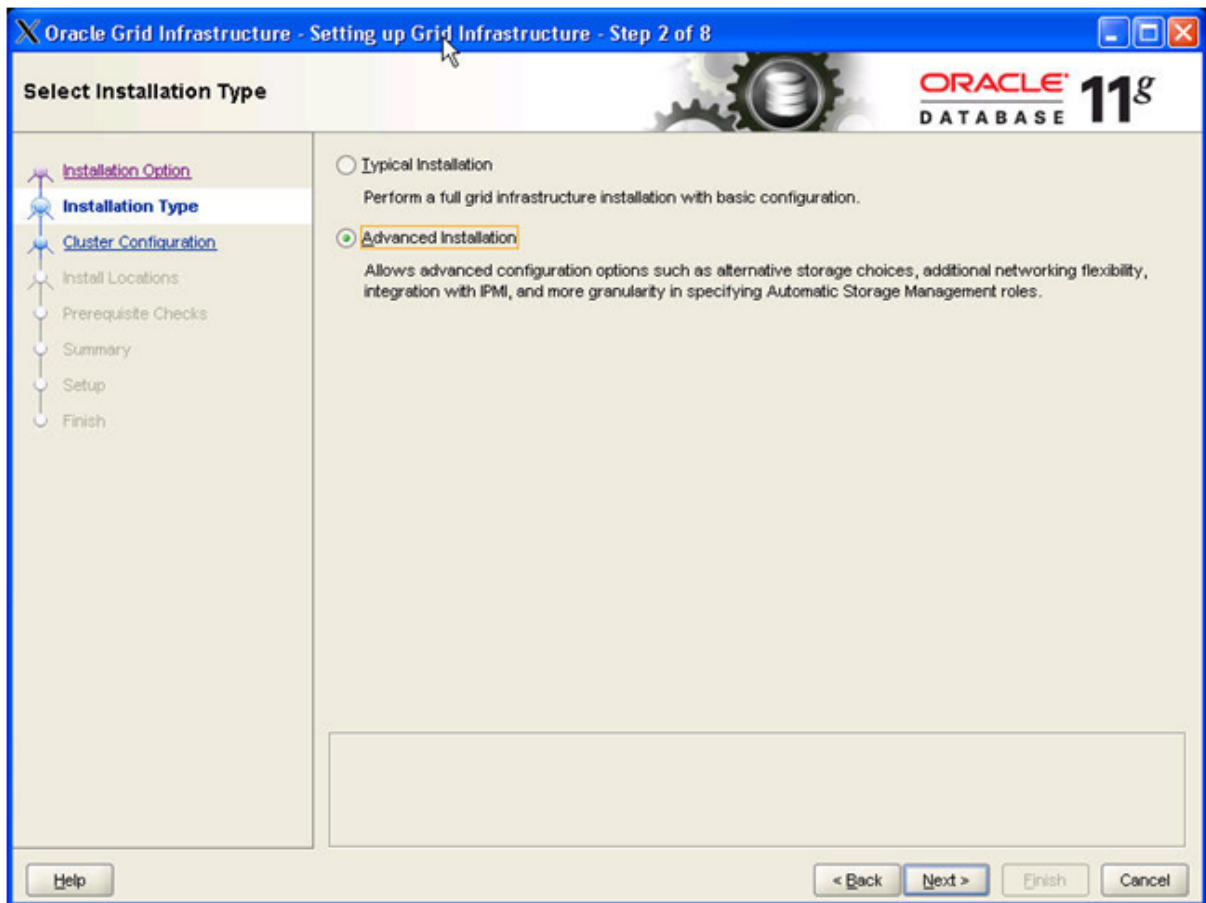
cd into the folder where you staged the grid infrastructure software

./runInstaller



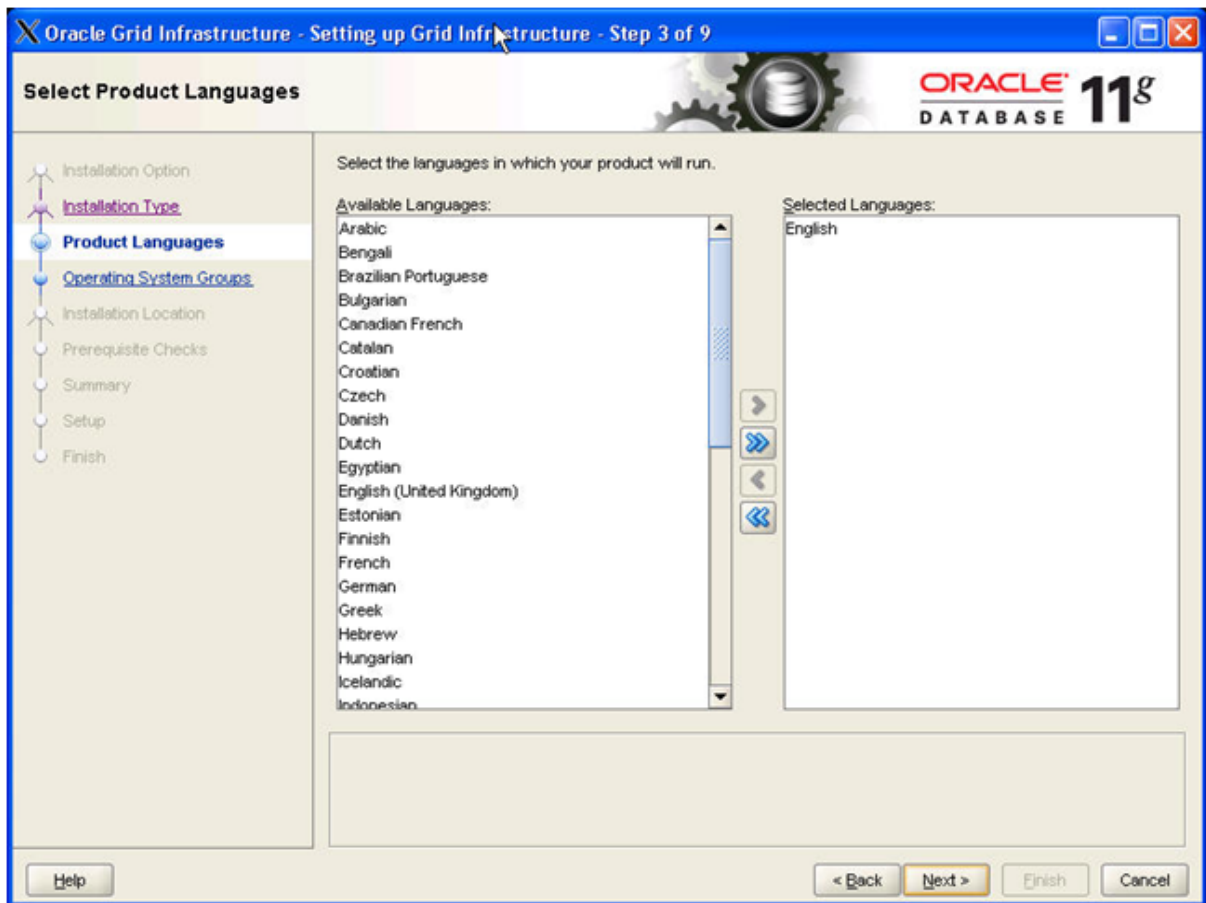
Action:

Select radio button 'Install and Configure Grid Infrastructure for a Cluster' and click 'Next>'

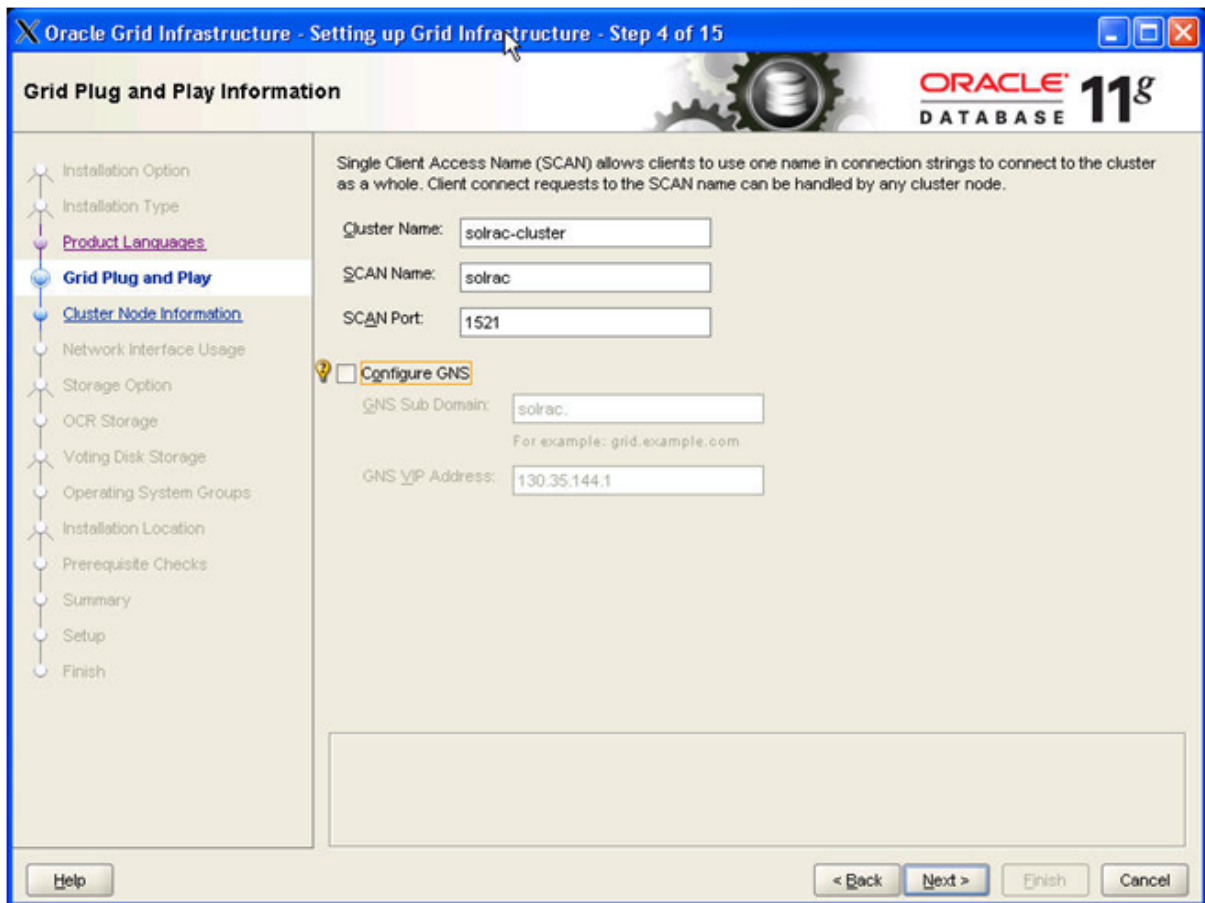


Action:

Select radio button 'Advanced Installation' and click ' Next> '



Action:
Accept 'English' as language' and click ' Next> '

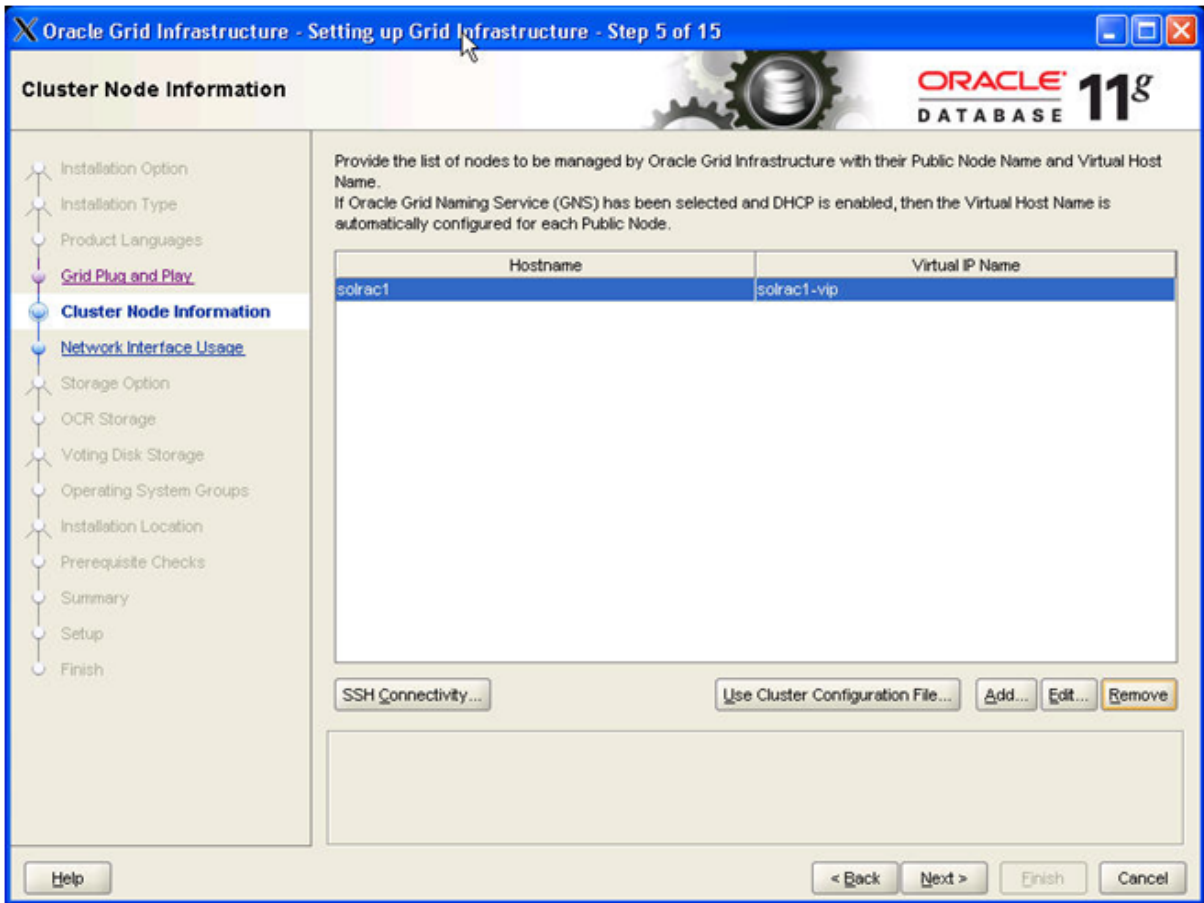


Action:

Specify your cluster name and the SCAN name you want to use and click ' Next> '

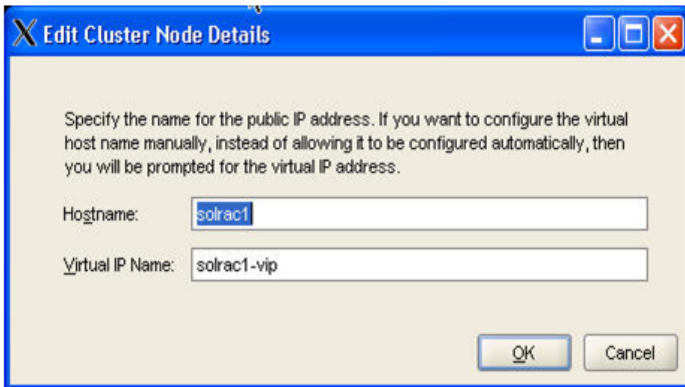
Note:

Make sure 'Configure GNS' is NOT selected.

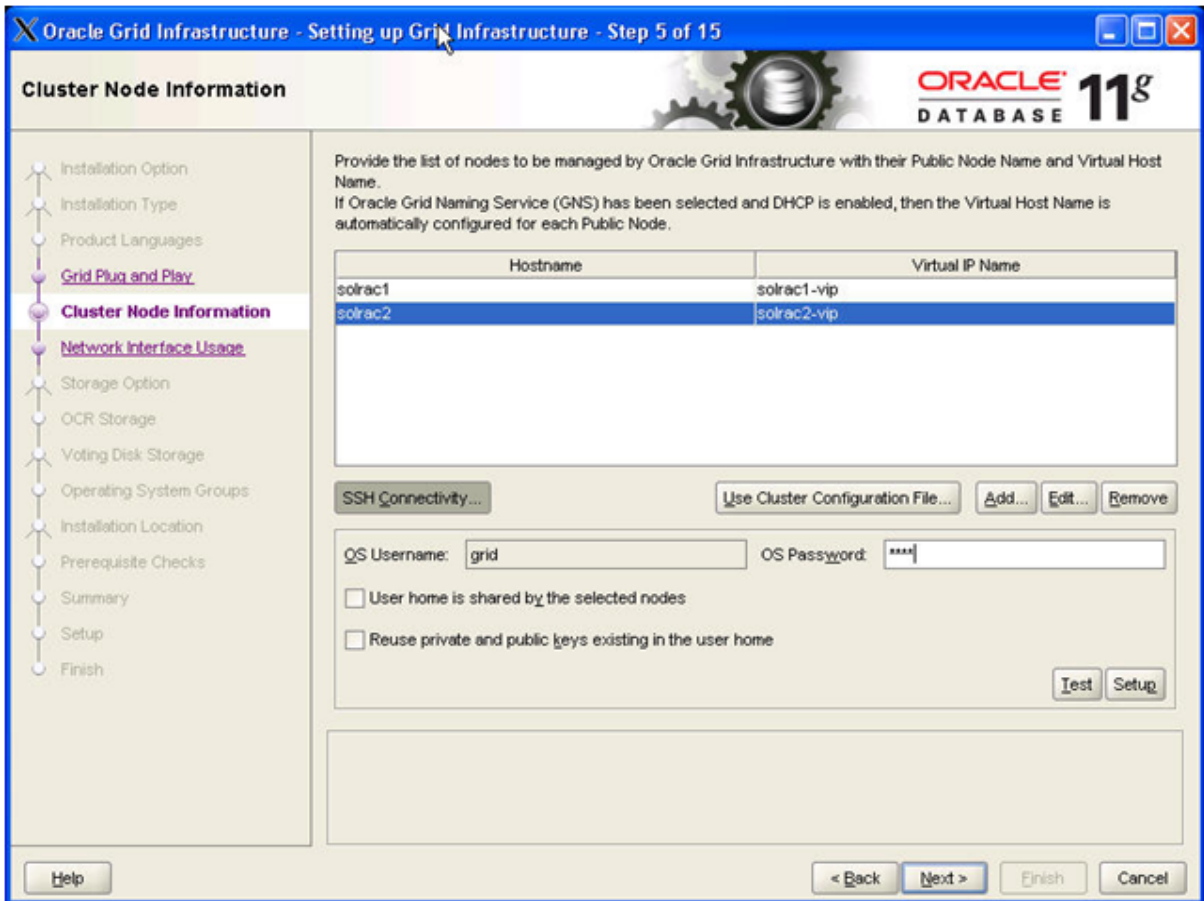


Action:

Use the Edit and Add buttons to specify the node names and virtual IP names you configured previously in your /etc/hosts file.

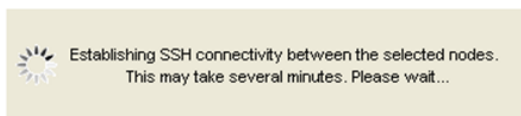


When finished click 'OK' and use the 'SSH Connectivity' button to configure/test the passwordless SSH connectivity between your nodes.



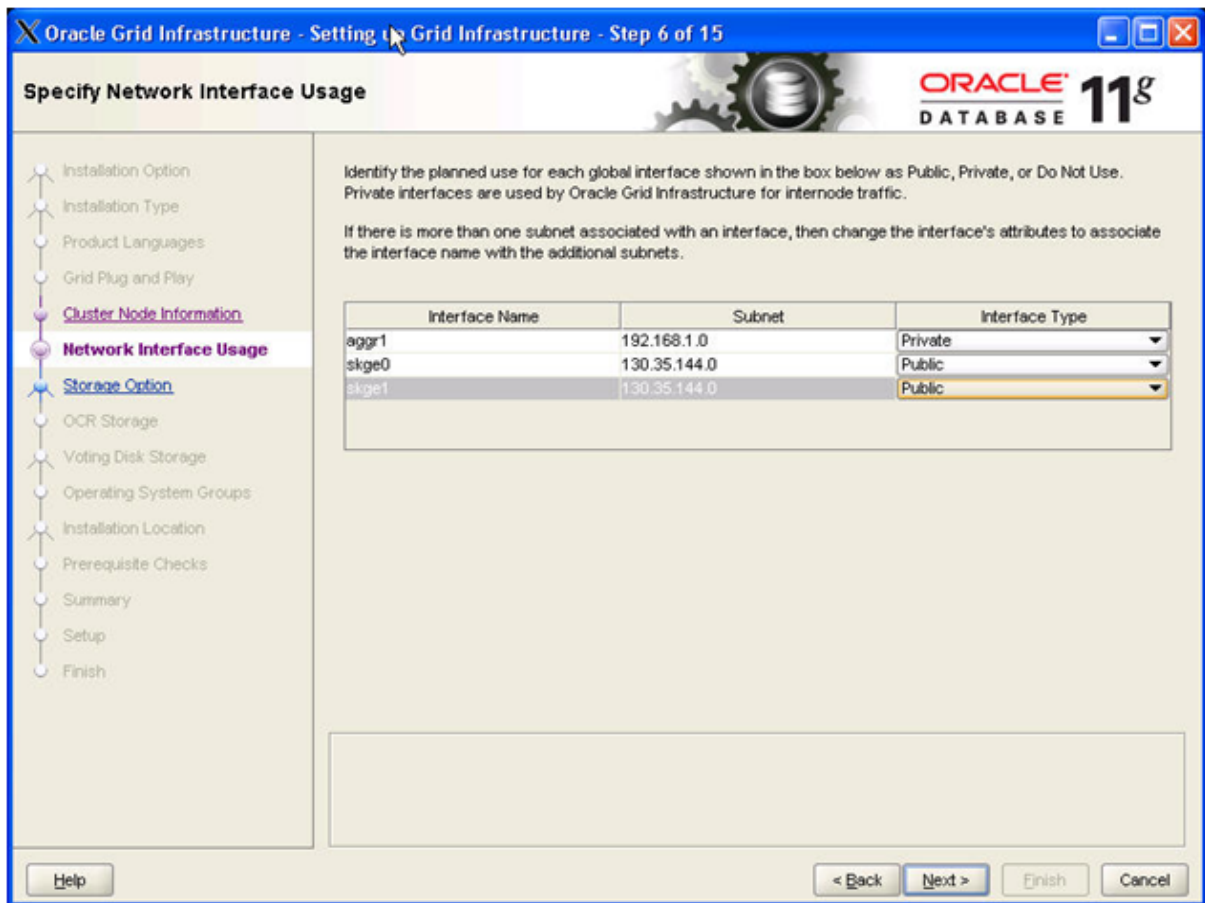
Action:

Type in the OS password for the user 'grid' and press 'Setup'



Action:

click 'OK'

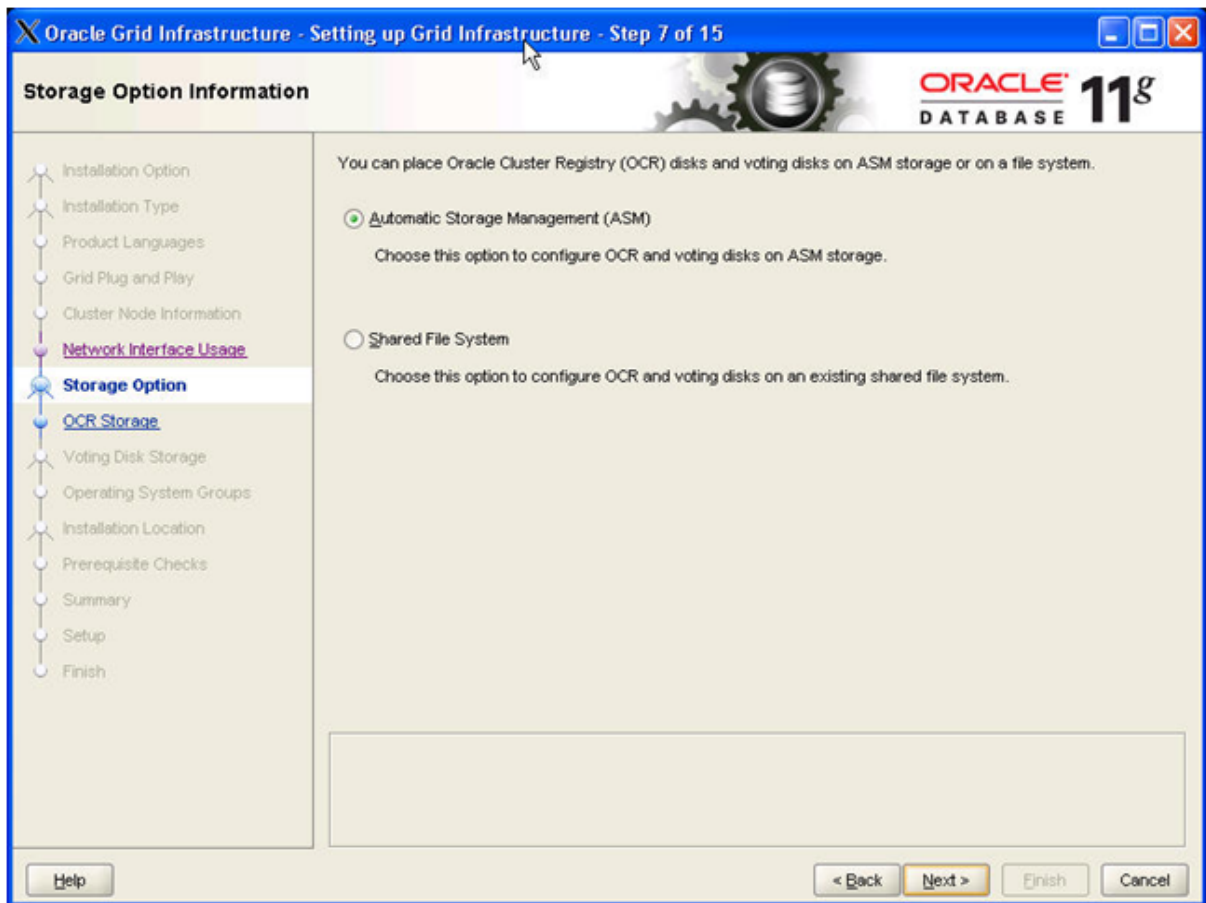


Action:

Click on 'Interface Type' next to the Interfaces you want to use for your cluster and select the correct values for 'Public', 'Private' and 'Do Not Use' . When finished click ' Next> '

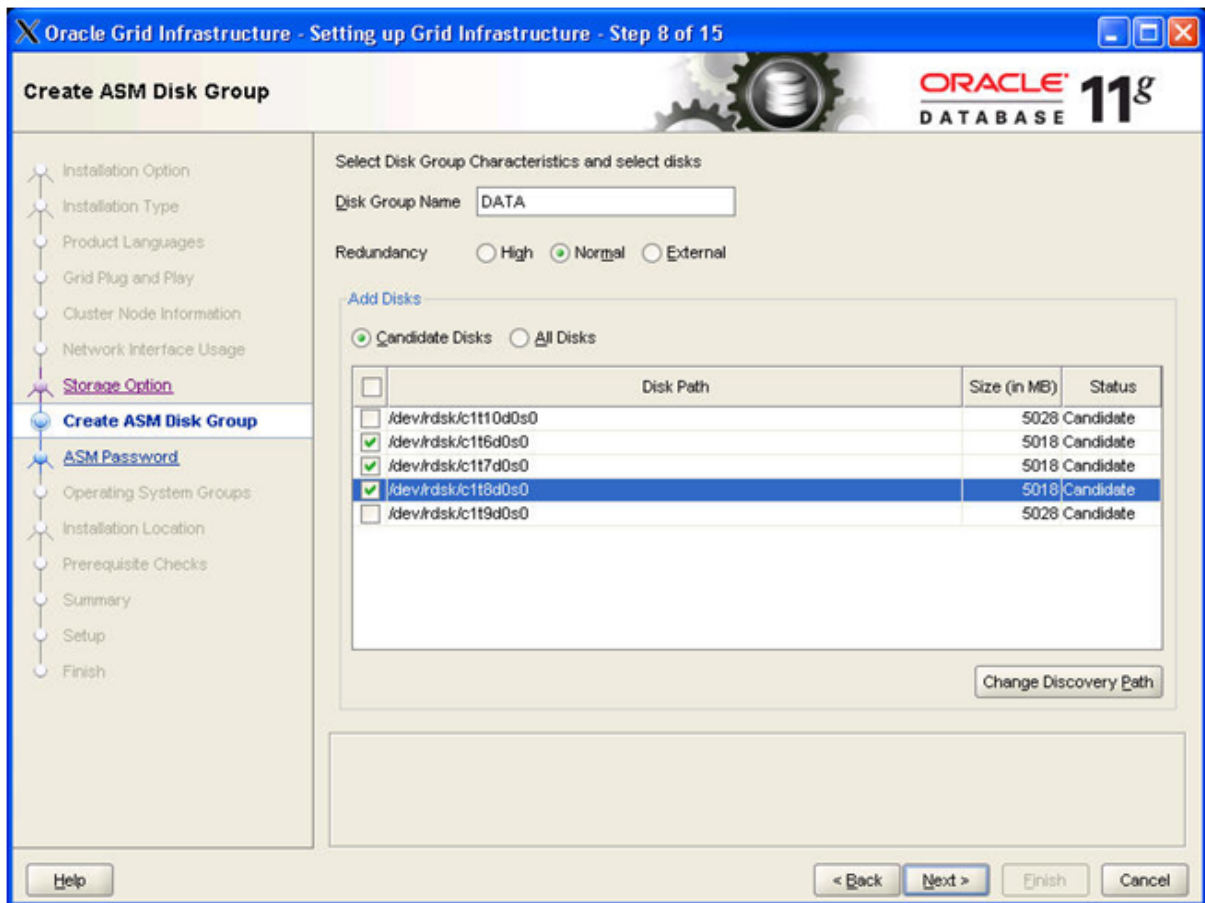
Note:

If you use multiple NIC's for redundancy the passive interfaces need to be selected here as well. In this example we are using IPMP for public network and Link Aggregation of private interconnect.



Action:

Select radio button 'Automatic Storage Management (ASM) and click ' Next> '



Action:

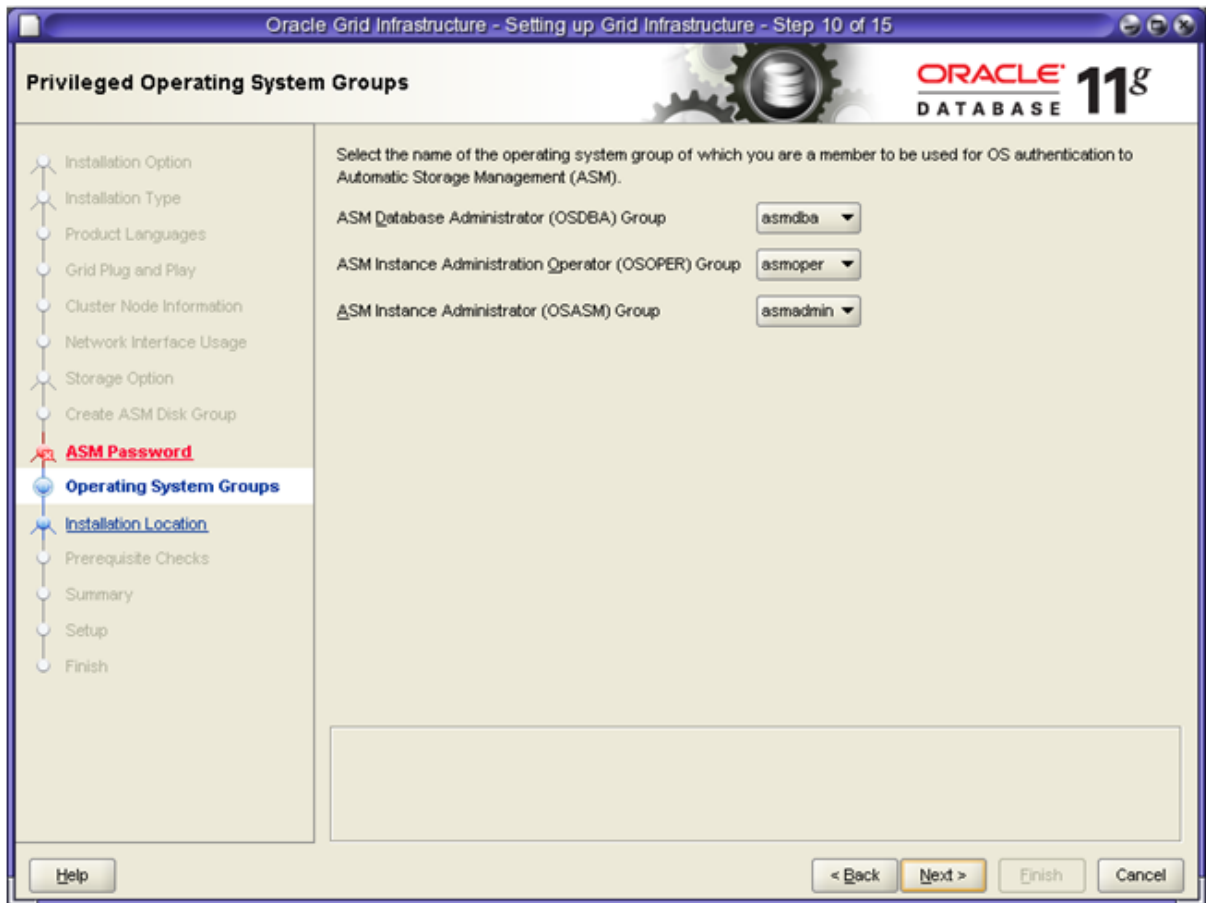
Type in a 'Disk Group Name' specify the 'Redundancy' and tick the disks you want to use, when done click ' Next> '

NOTE: The number of voting disks that will be created depend on the redundancy level you specify: external will create 1 voting disk, normal will create 3 voting disks, high will create 5 voting disks.



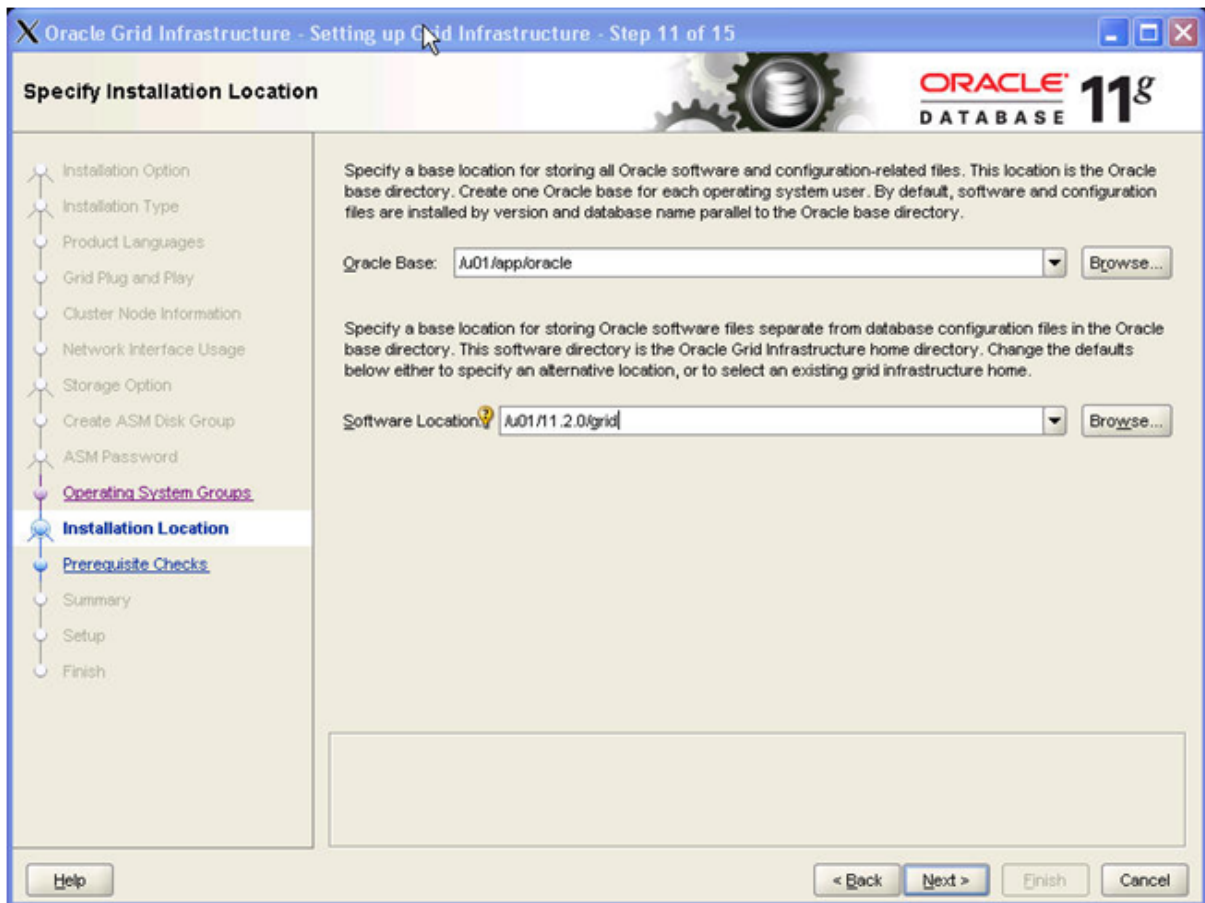
Action:

Specify and conform the password you want to use and click 'Next >'



Action:

Assign the correct OS groups for OS authentication and click 'Next>'

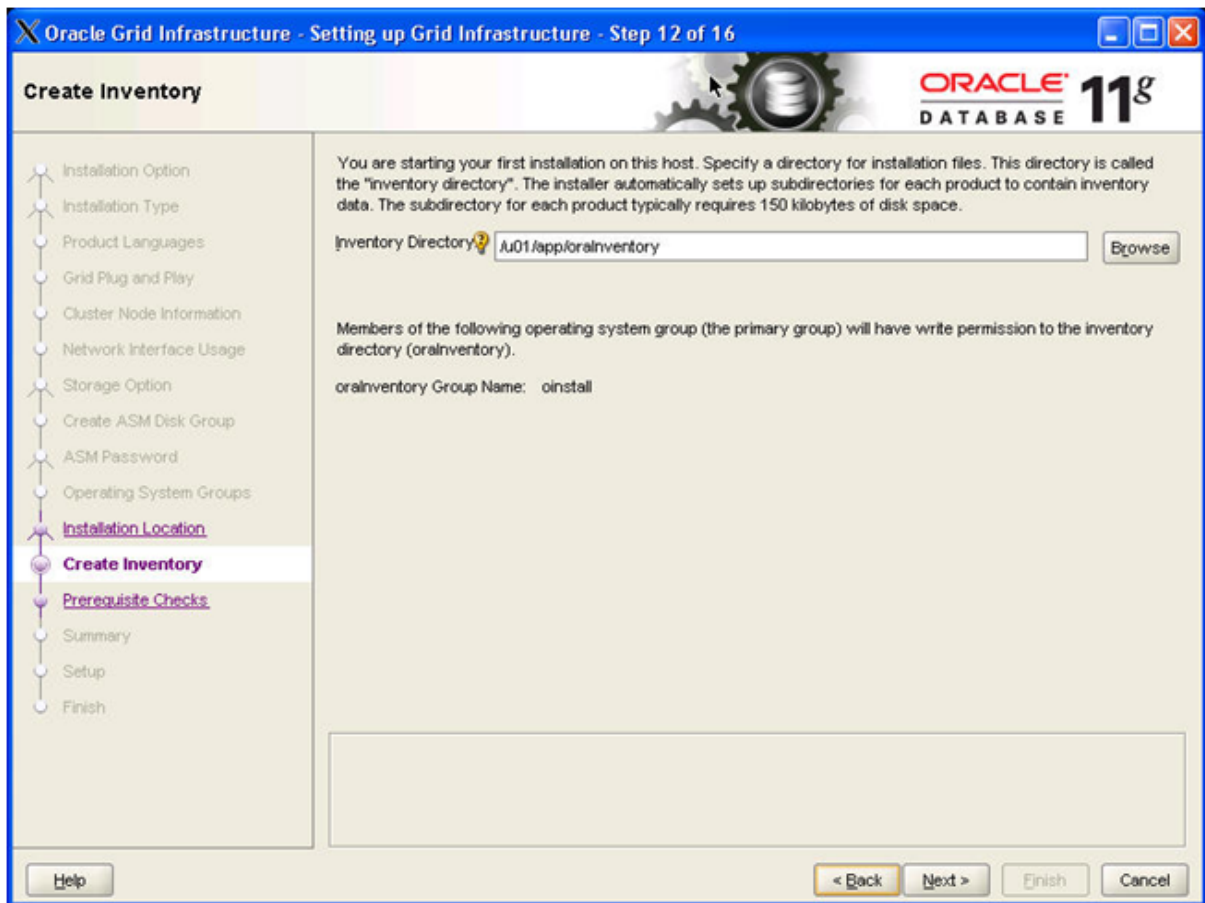


Action:

Specify the locations for your ORACLE_BASE(/u01/app/oracle) and for the Software location(/u01/11.2.0/grid) and click ' Next> '

Note:

We created the directories in step 2.6.

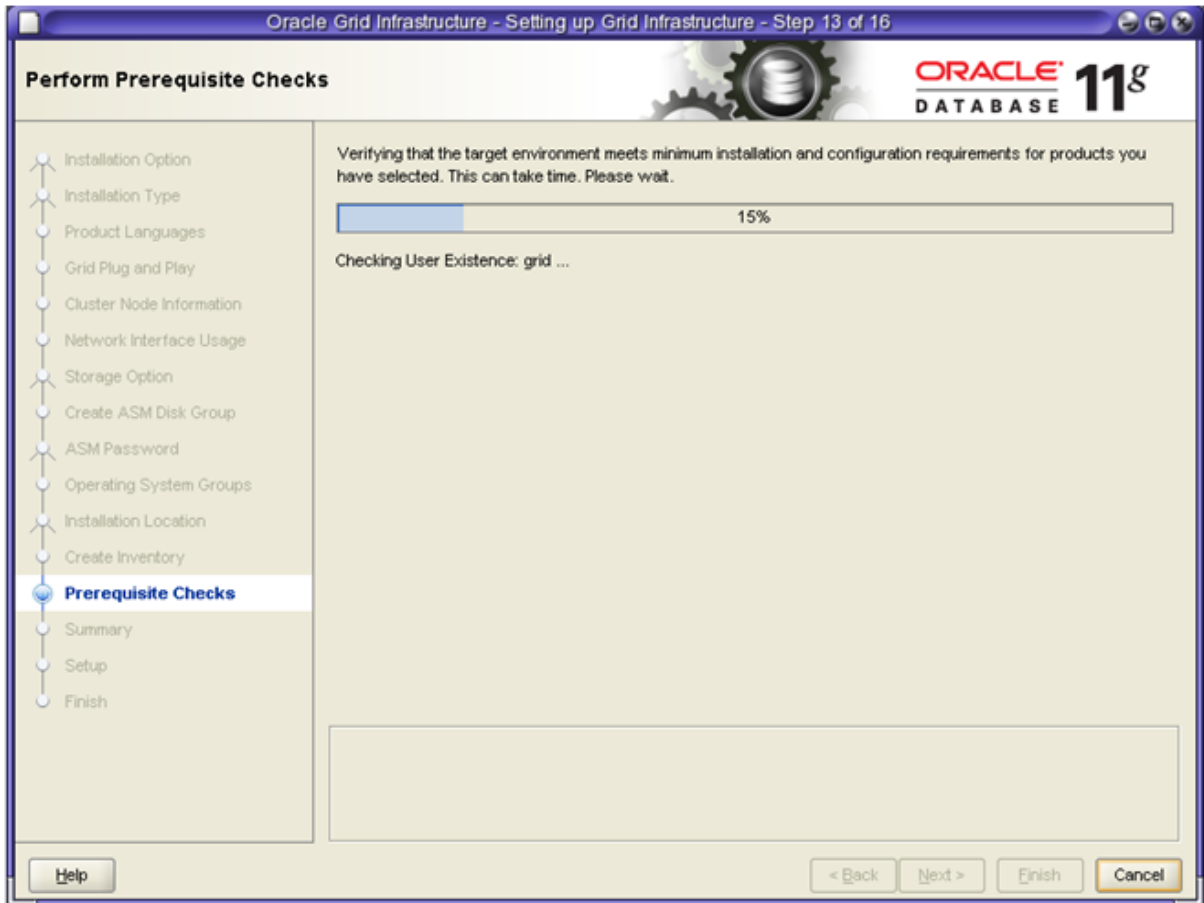


Action:

Specify the locations for your Inventory (/u01/app/orainventory) directory and click 'Next>'

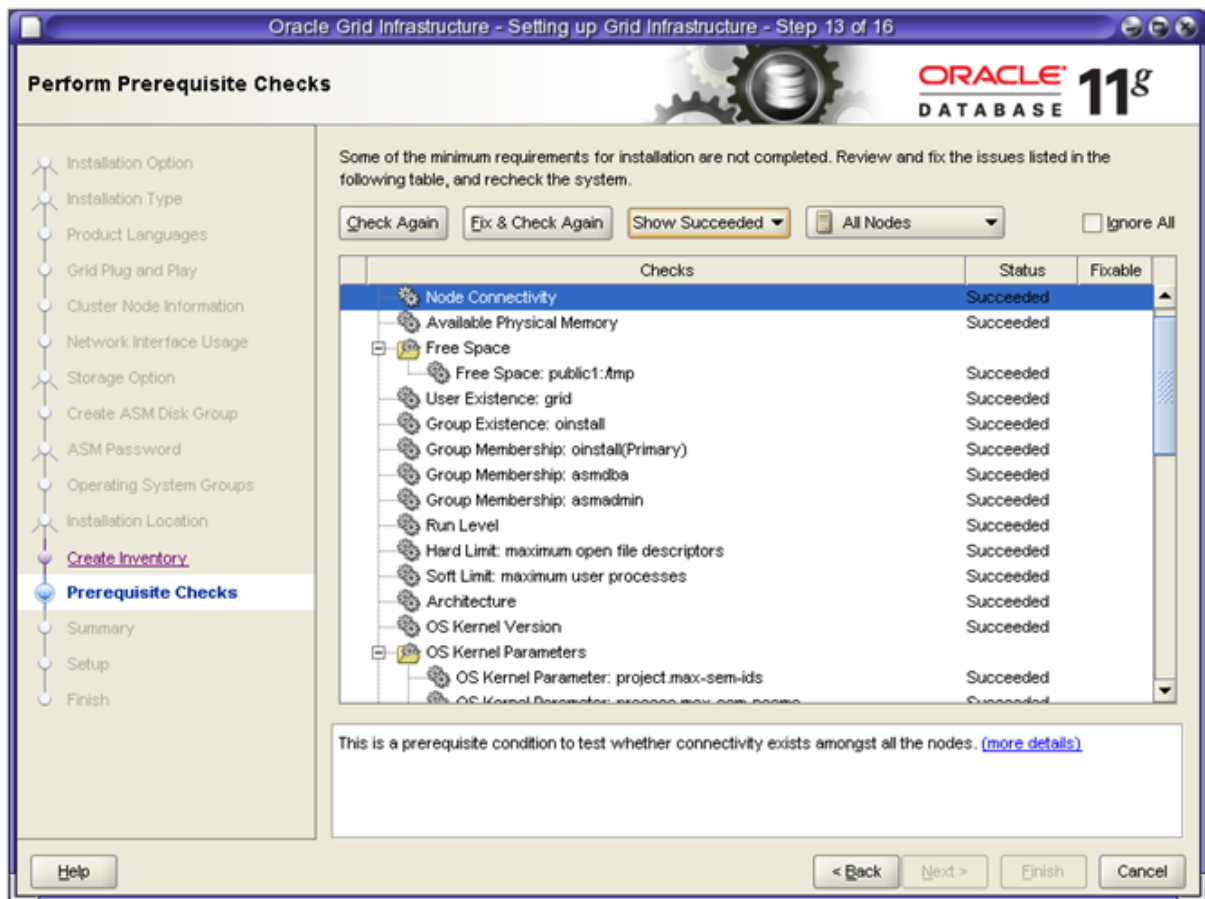
Note:

We created the directory in step 2.5.



Note:

OUI performs certain checks and comes back with the screen below

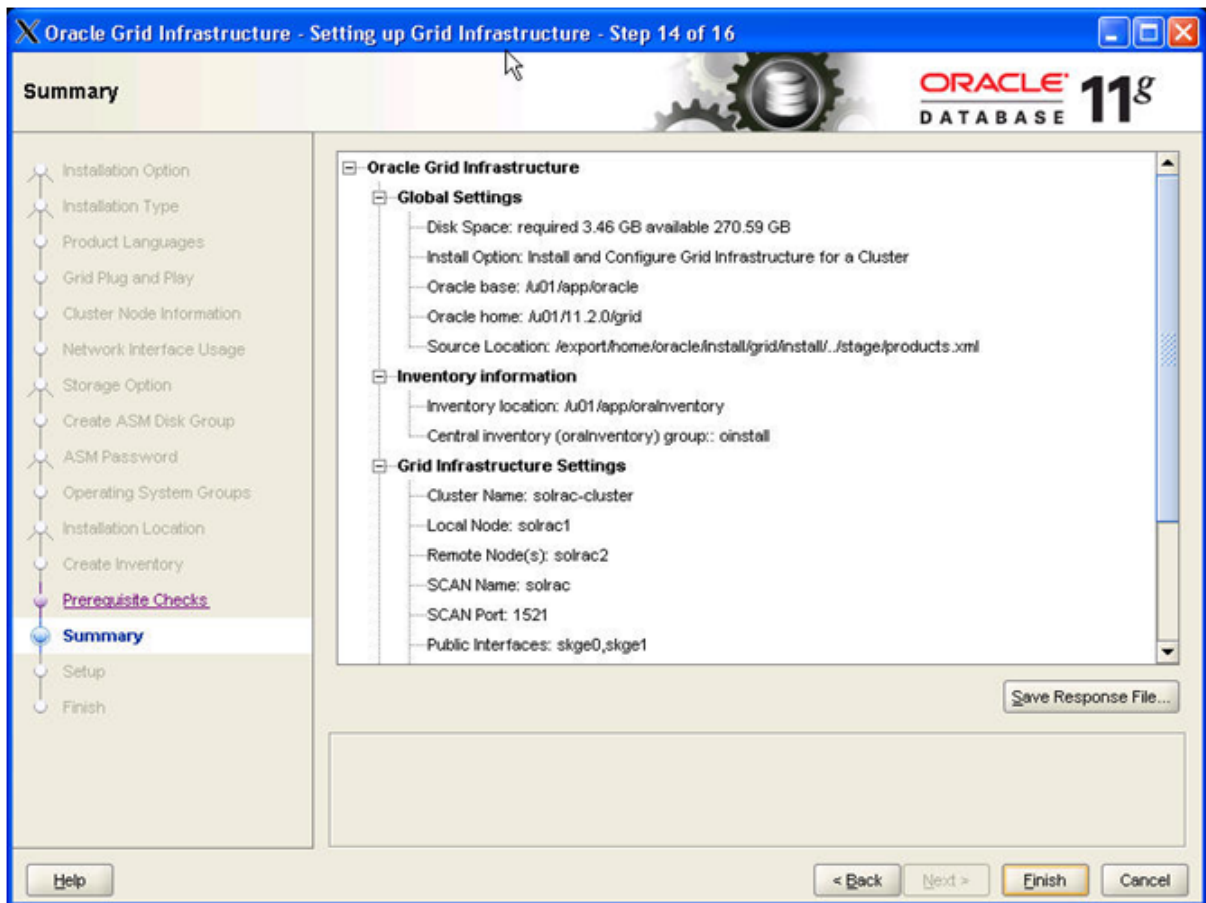


Action:

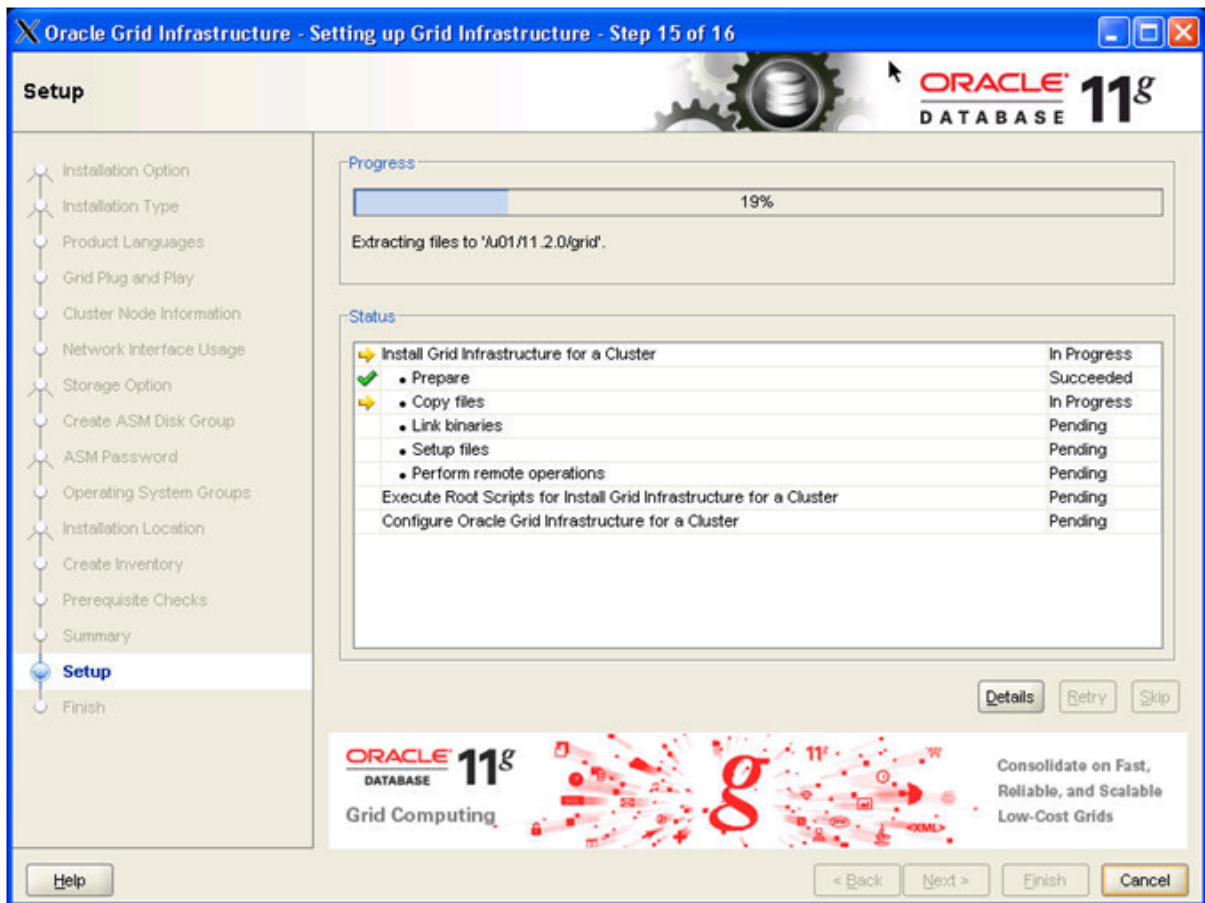
Check that status of all checks is Succeeded and click ' Next> '

Note:

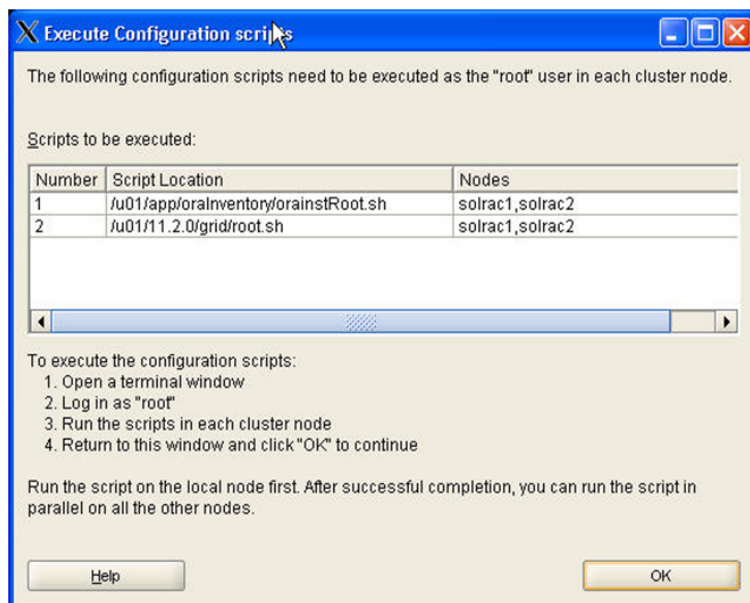
If you have failed checks marked as 'Fixable' click 'Fix & Check again'. This will bring up a window that instructs you to execute fixup scripts. Execute the runfixup.sh script as described on the screen as root user. Click 'Check Again' and if all checks succeeded click 'Next>'



Action:
Click ' Finish'

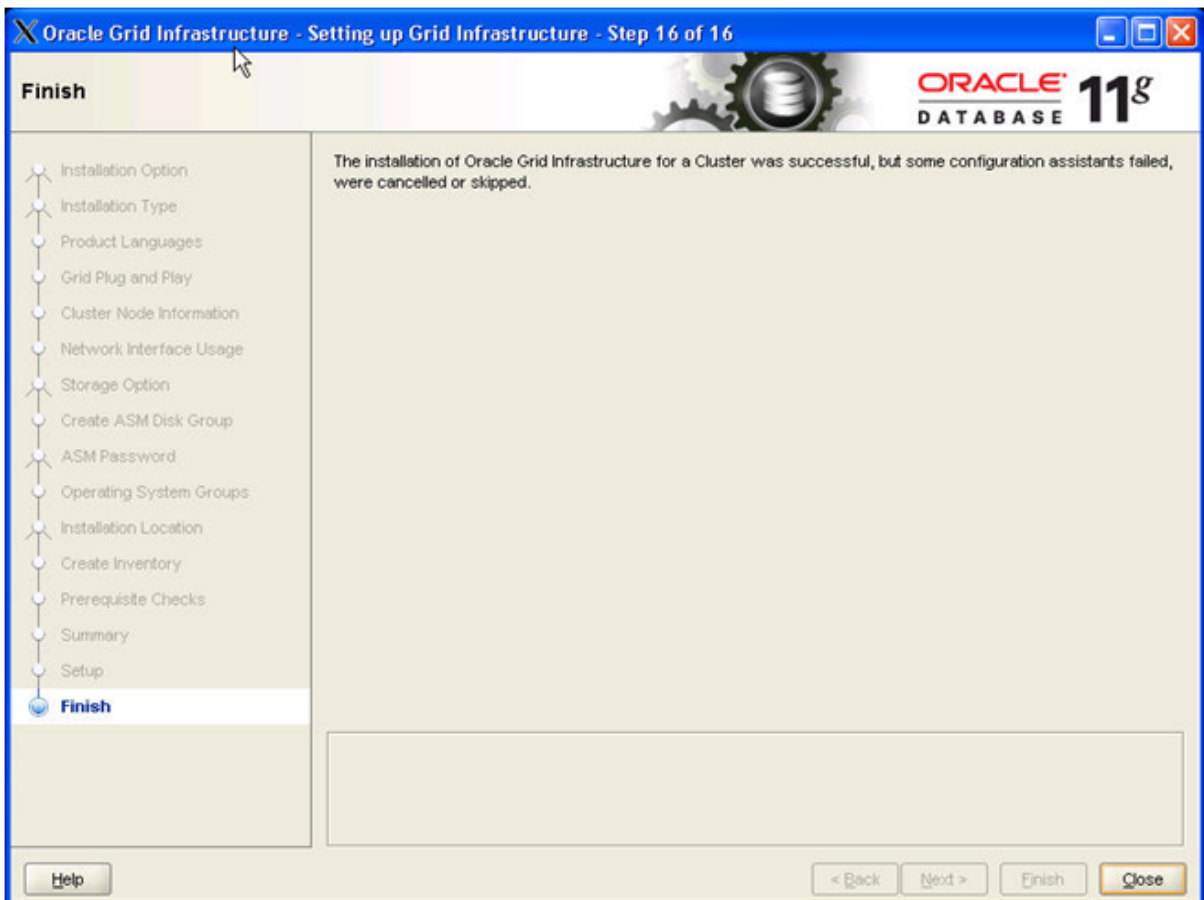
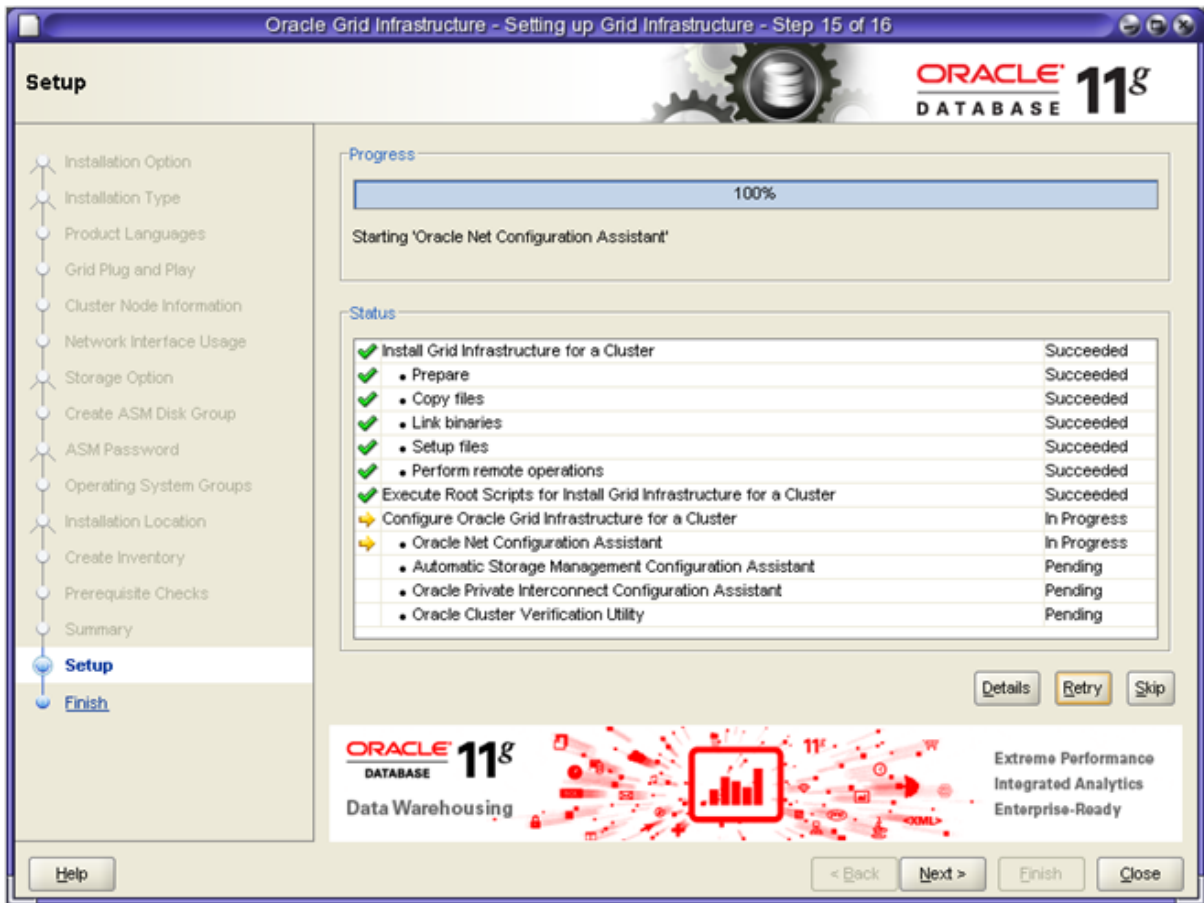


Action:
Wait for the OUI to complete its tasks



Action:
Follow the instructions on the screen running the orainstRoot.sh and root.sh scripts as root on all nodes before you click 'OK'

Note:
The required root scripts MUST BE RUN ON ONE NODE AT A TIME!



Action:

You should see the confirmation that the installation of the Grid Infrastructure was successful. Click 'Close' to finish the install.

5. Grid Infrastructure Home Patching

This Chapter is a placeholder

6. RDBMS Software Install

As the oracle user (rdbms software owner) start the installer by running "runInstaller" from the staged installation media.

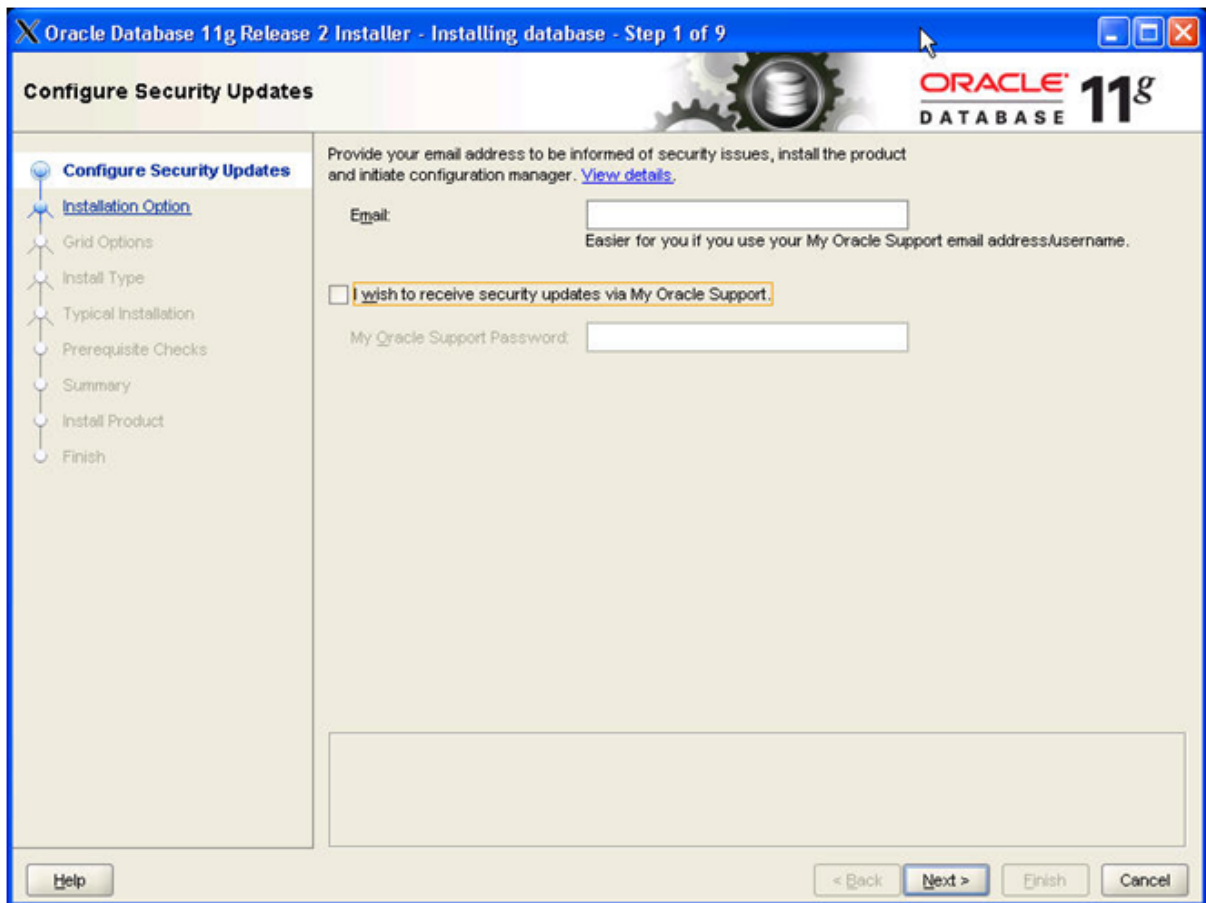
NOTE: Be sure the installer is run as the intended software owner, the only supported method to change the software owner is to reinstall.

#su - oracle

change into the directory where you staged the rdbms software

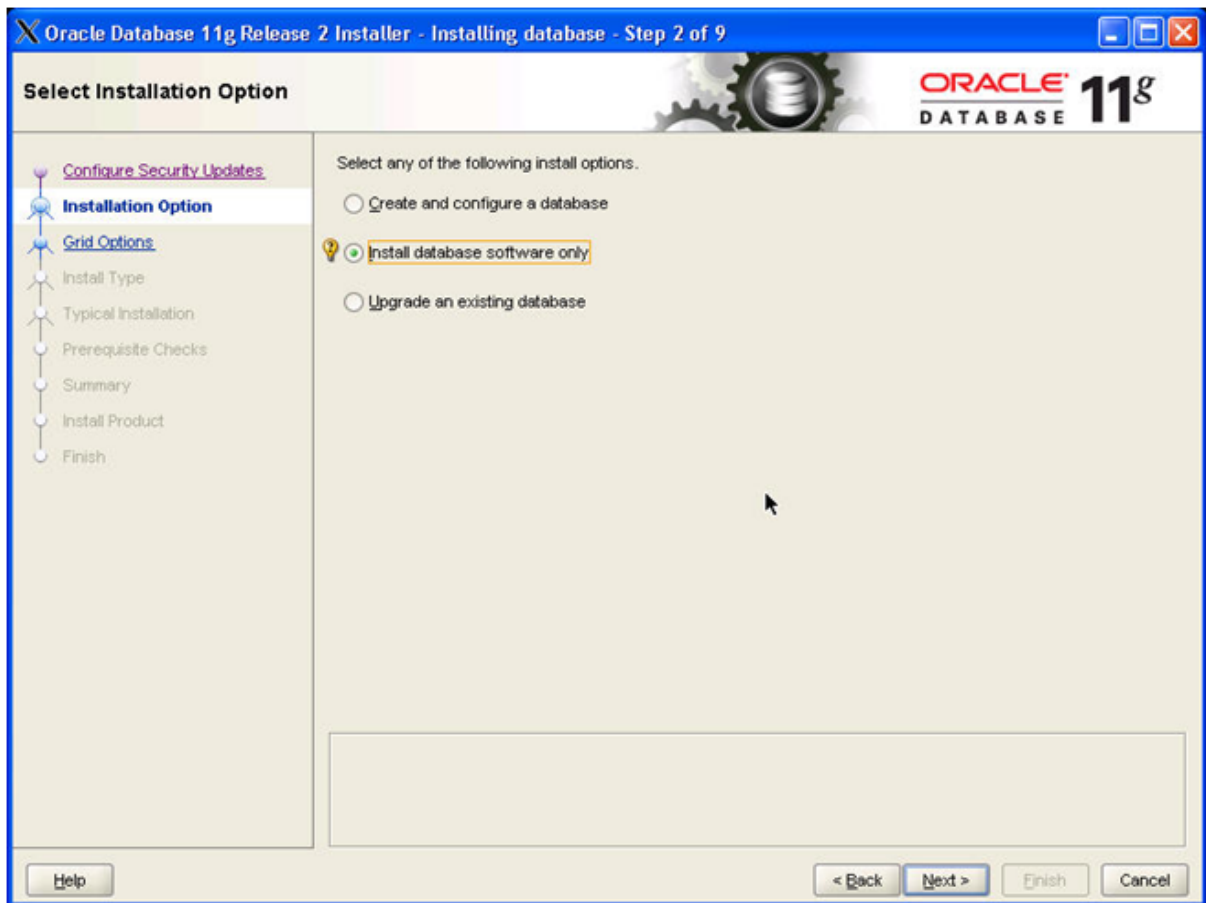
./runInstaller





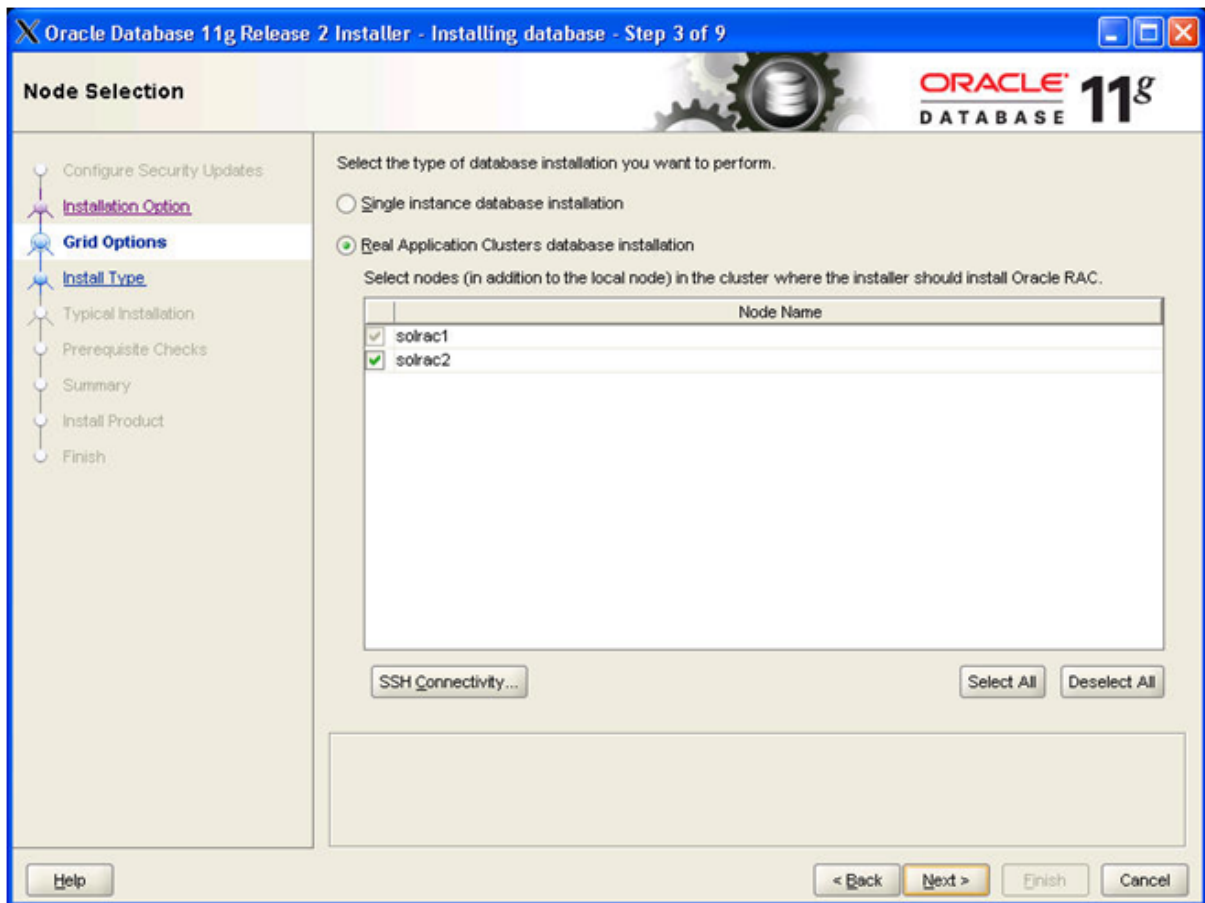
Action:

Provide your e-mail address, tick the check box and provide your Oracle Support Password if you want to receive Security Updates from Oracle Support, after click ' Next> '



Action:

Select the option 'Install Database software only' and click 'Next>'

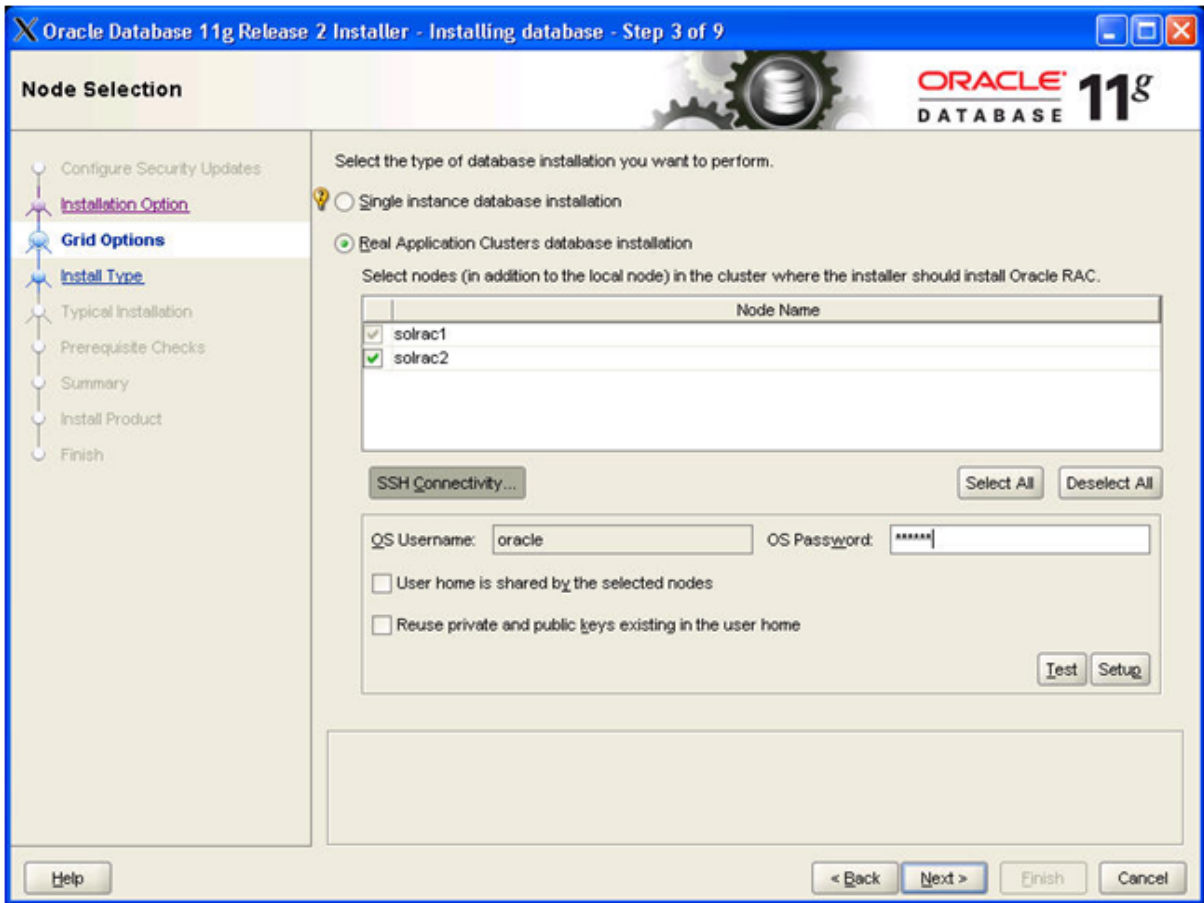


Action:

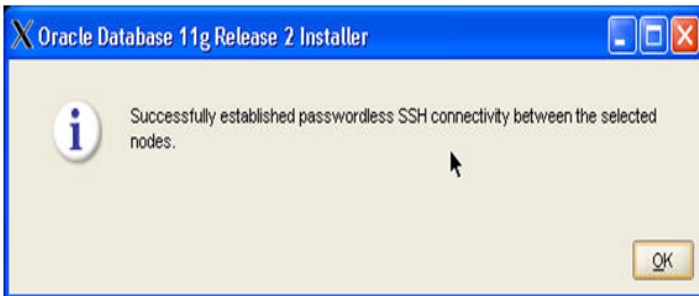
Select 'Real Application Clusters database installation', and select all nodes. If User Equivalent is not configured, click the 'SSH Connectivity' button to configure/test the passwordless SSH connectivity between your nodes.

Note:

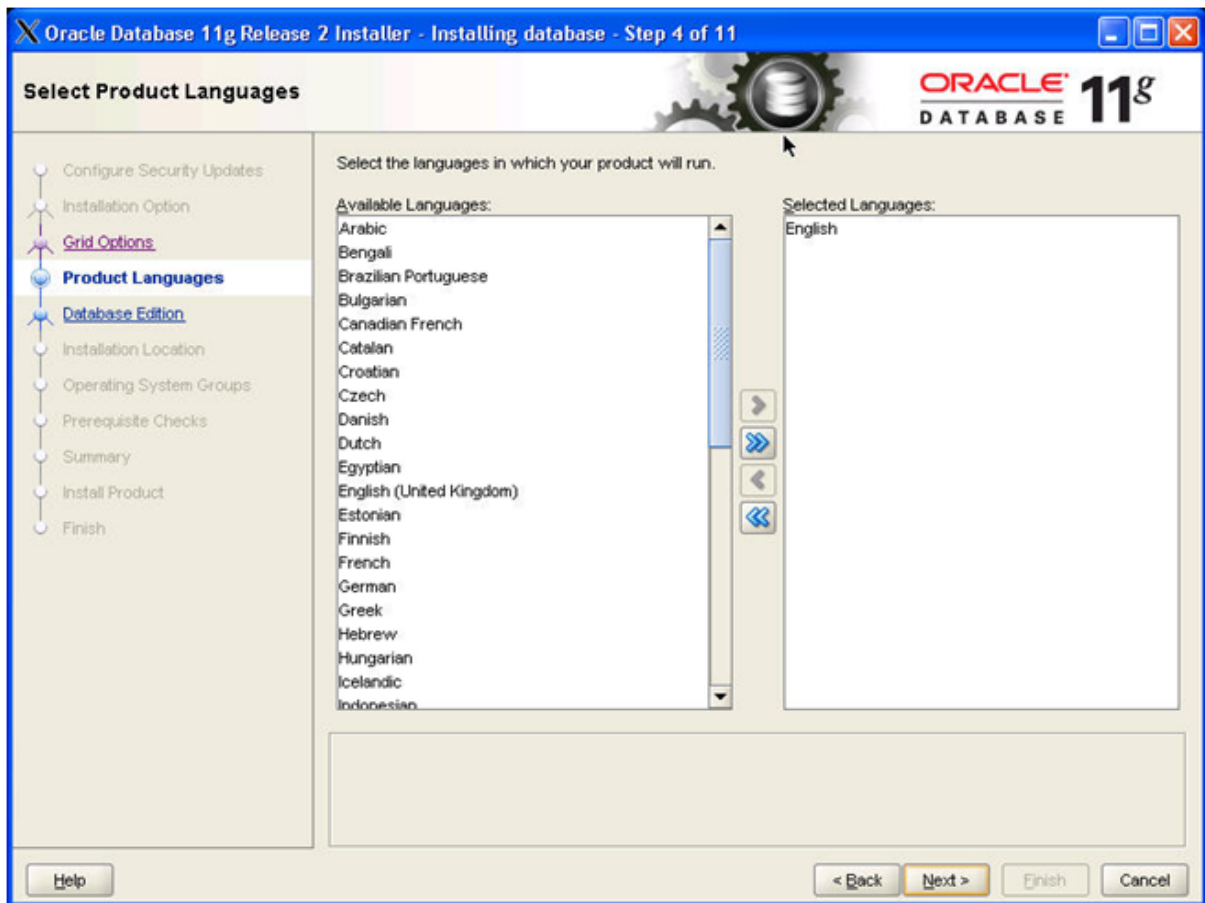
During the Grid Infrastructure installation you configured SSH for the grid user. If you install RDBMS with a different user (recommended) you have to configure it for this user now.



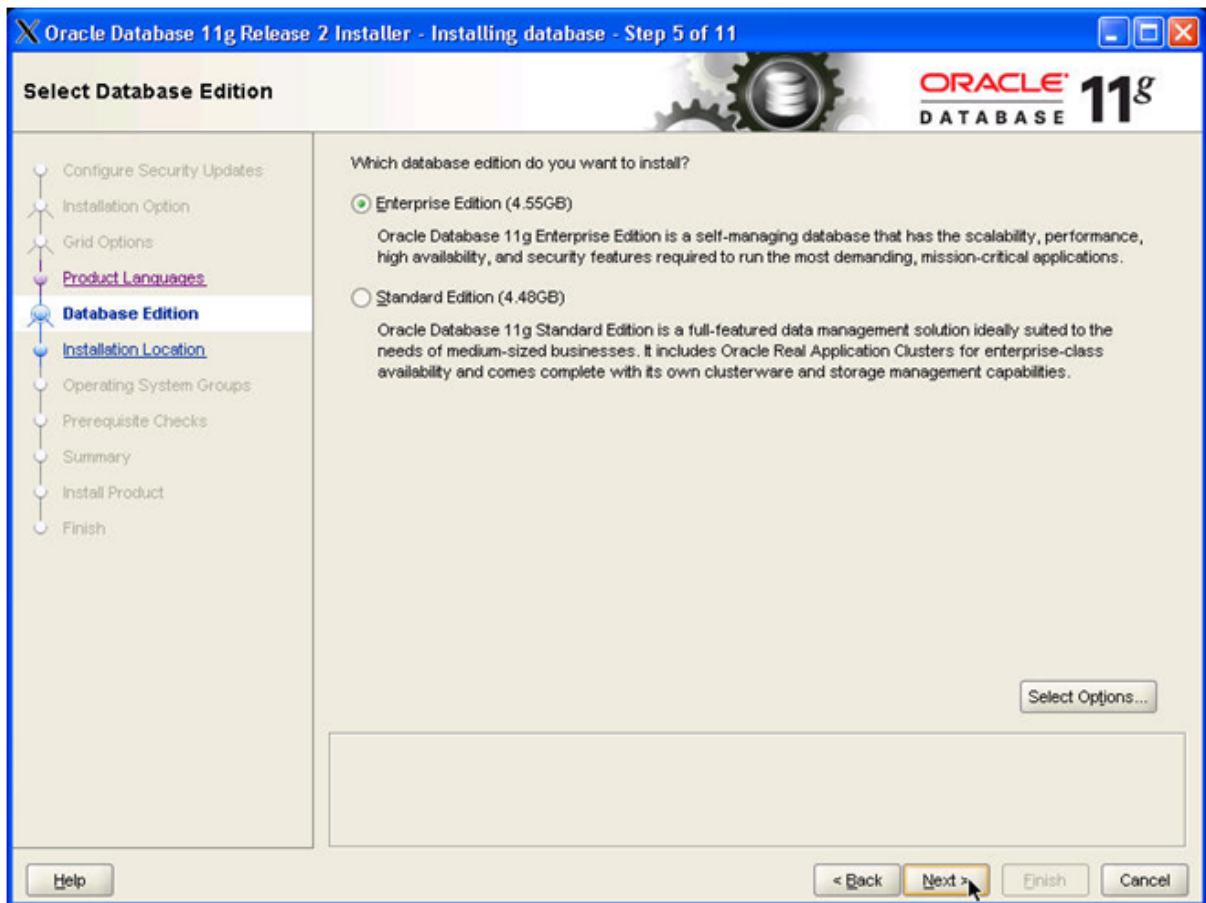
Action:
Type in the OS password for the oracle user and click 'Setup'



Action:
click 'OK' and ' Next> '

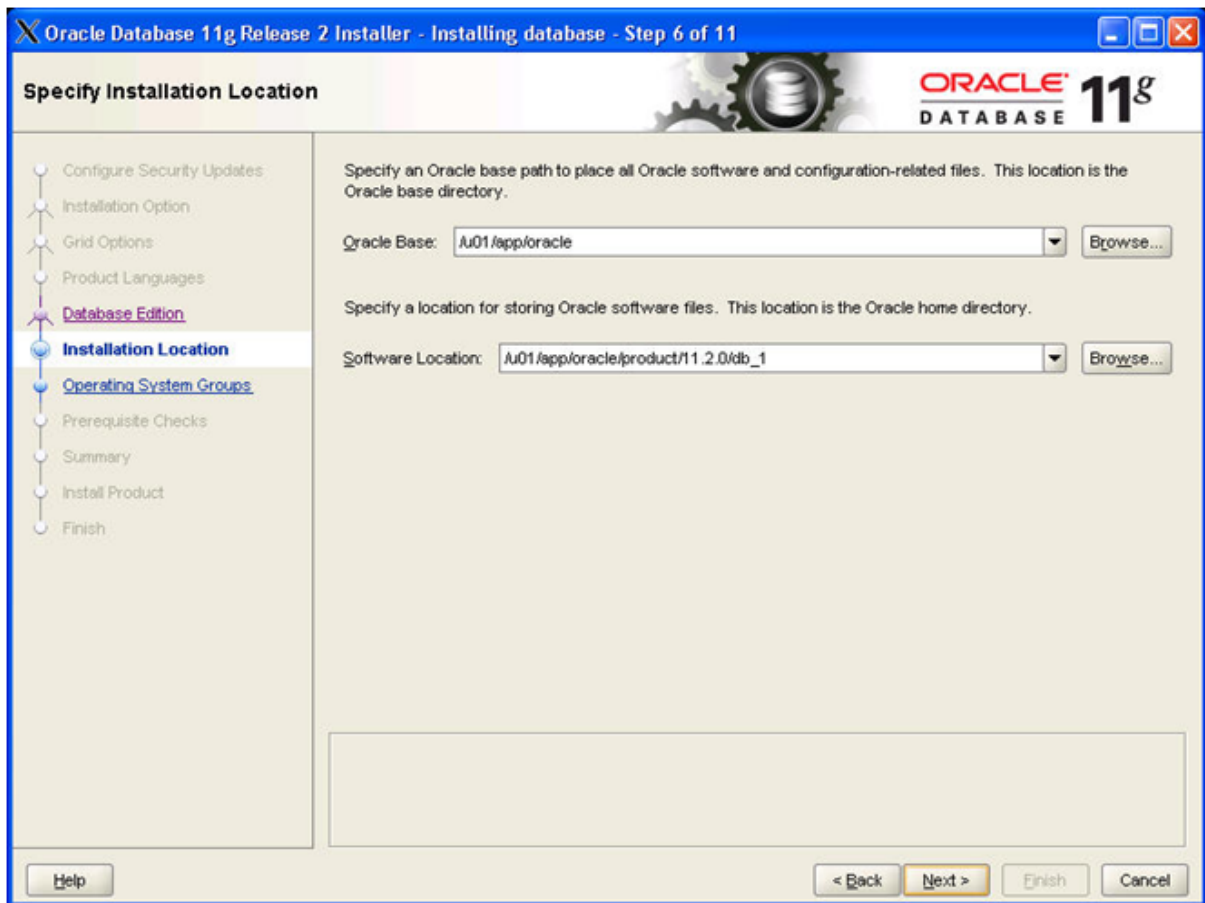


Action:
To confirm English as selected language click ' Next> '



Action:

Make sure radio button 'Enterprise Edition' is ticked, click 'Next>'



Action:

Specify path to your Oracle Base and below to the location where you want to store the software (Oracle home). Click ' Next> '

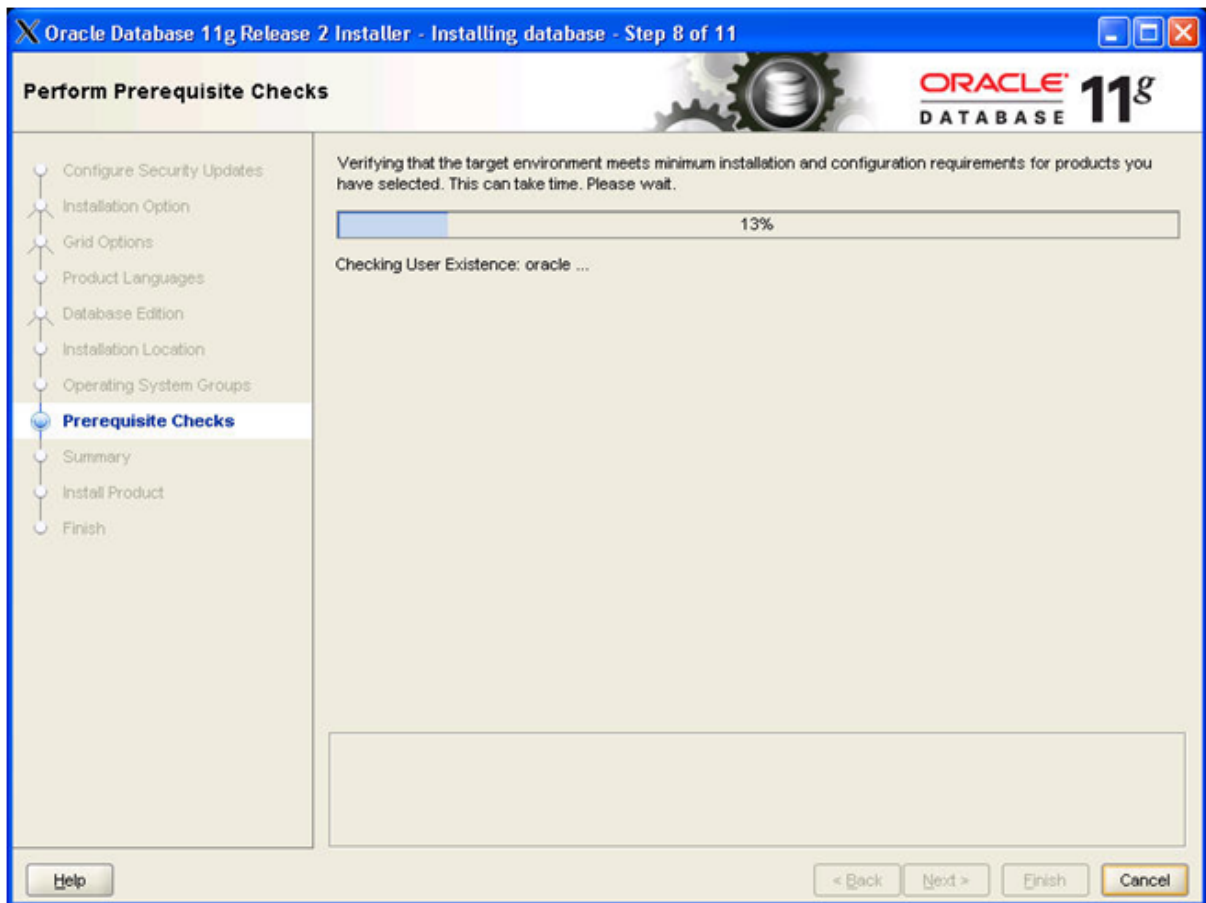
Note:

We created the directories in steps 2.7 and 2.8

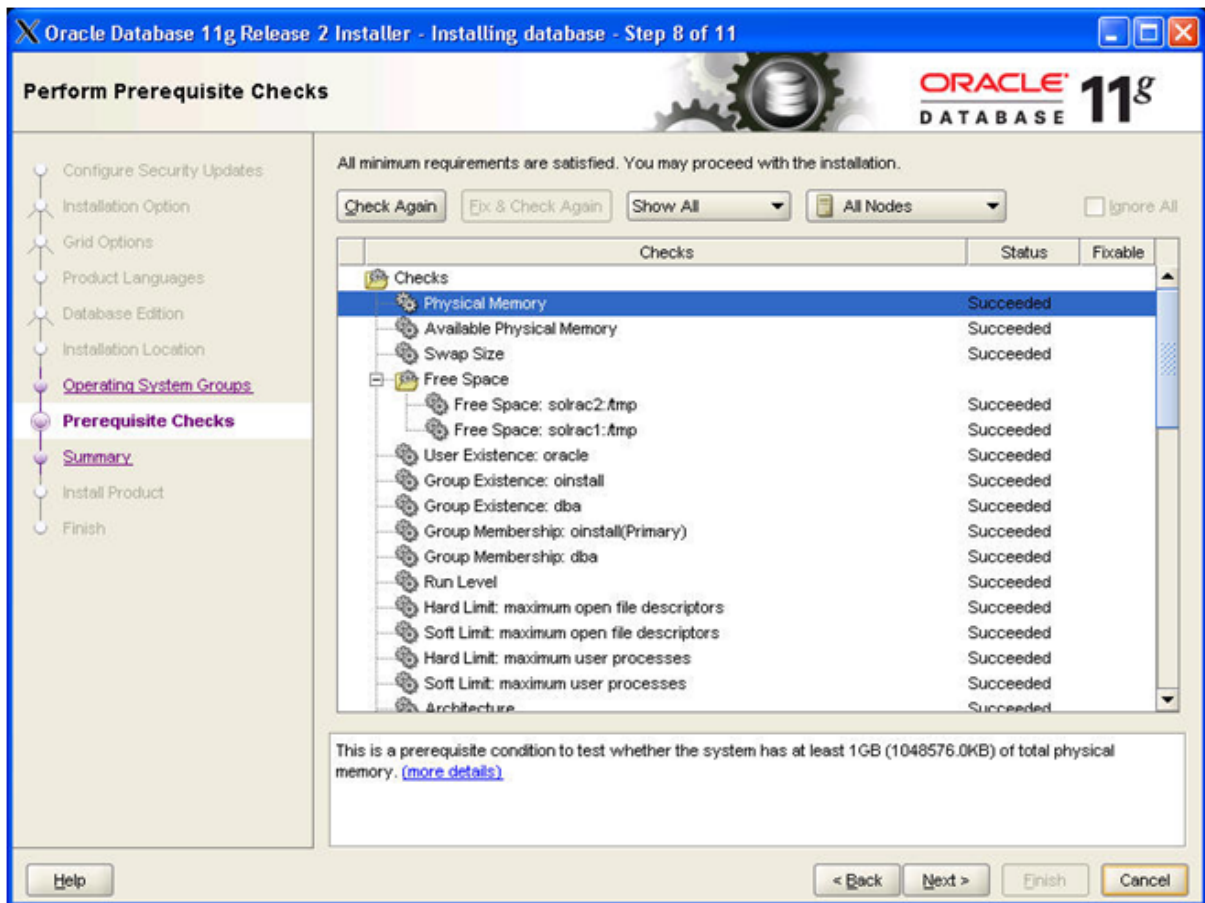


Action:

Use the drop down menu to select the names of the Database Administrators and Database Operators group and click Next> '



Note:
Oracle Universal Installer performs prerequisite checks.

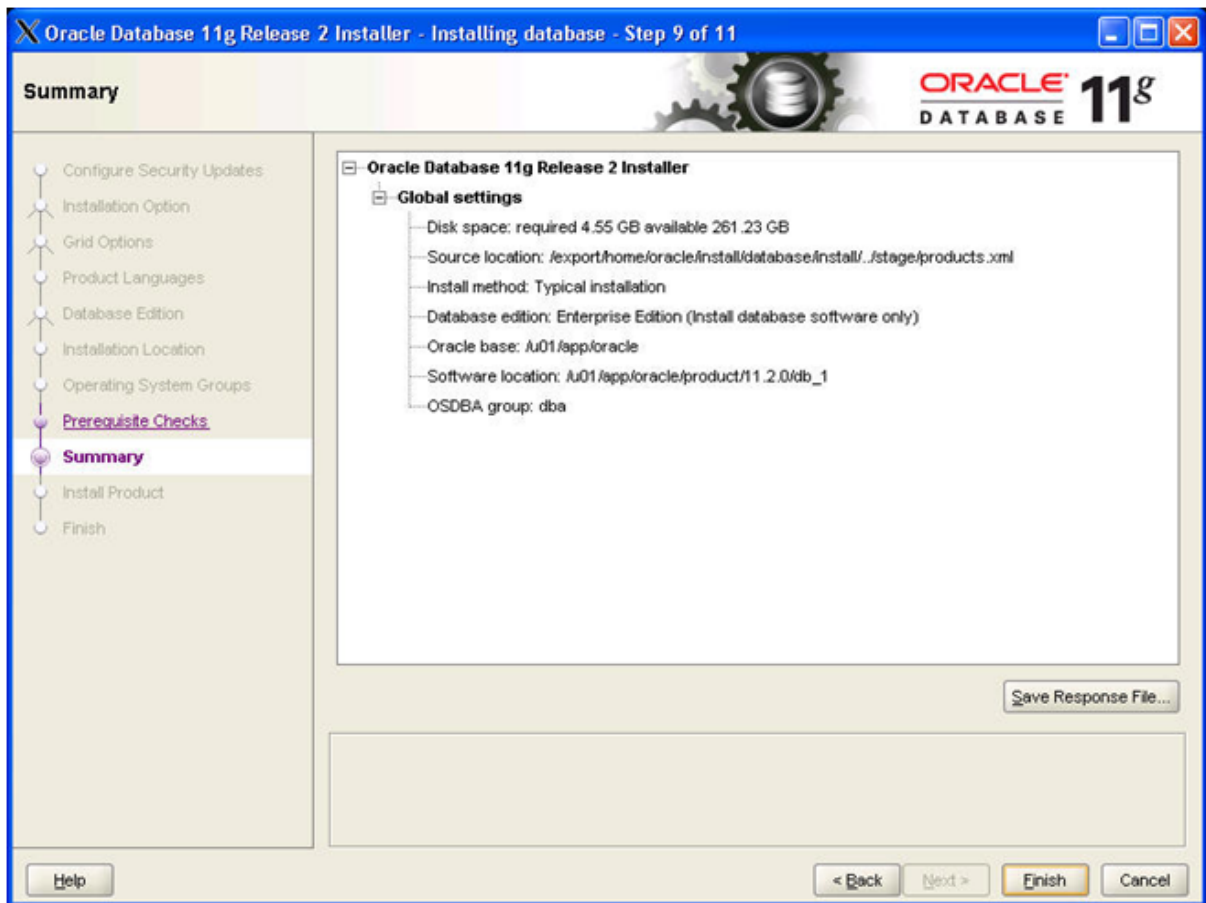


Action:

Check that status of all checks is Succeeded and click ' Next> '

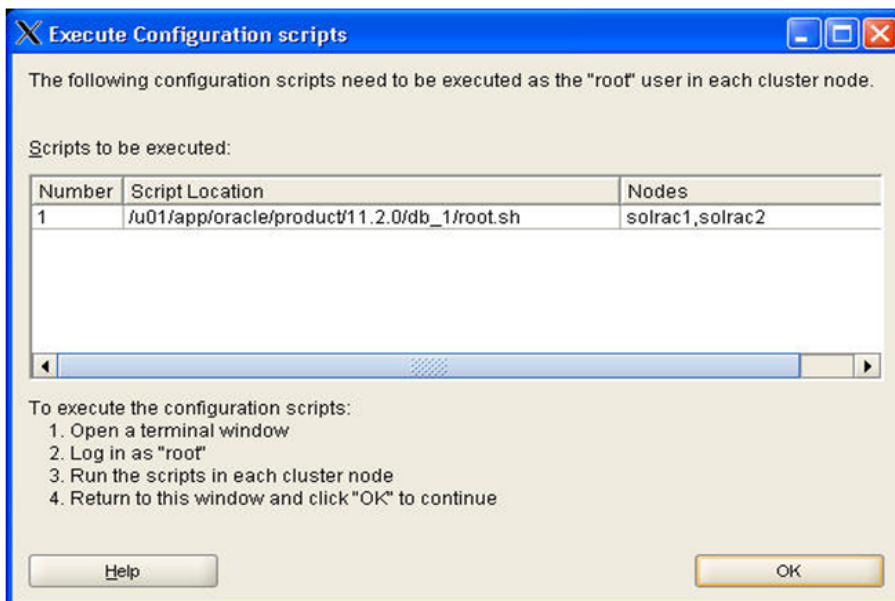
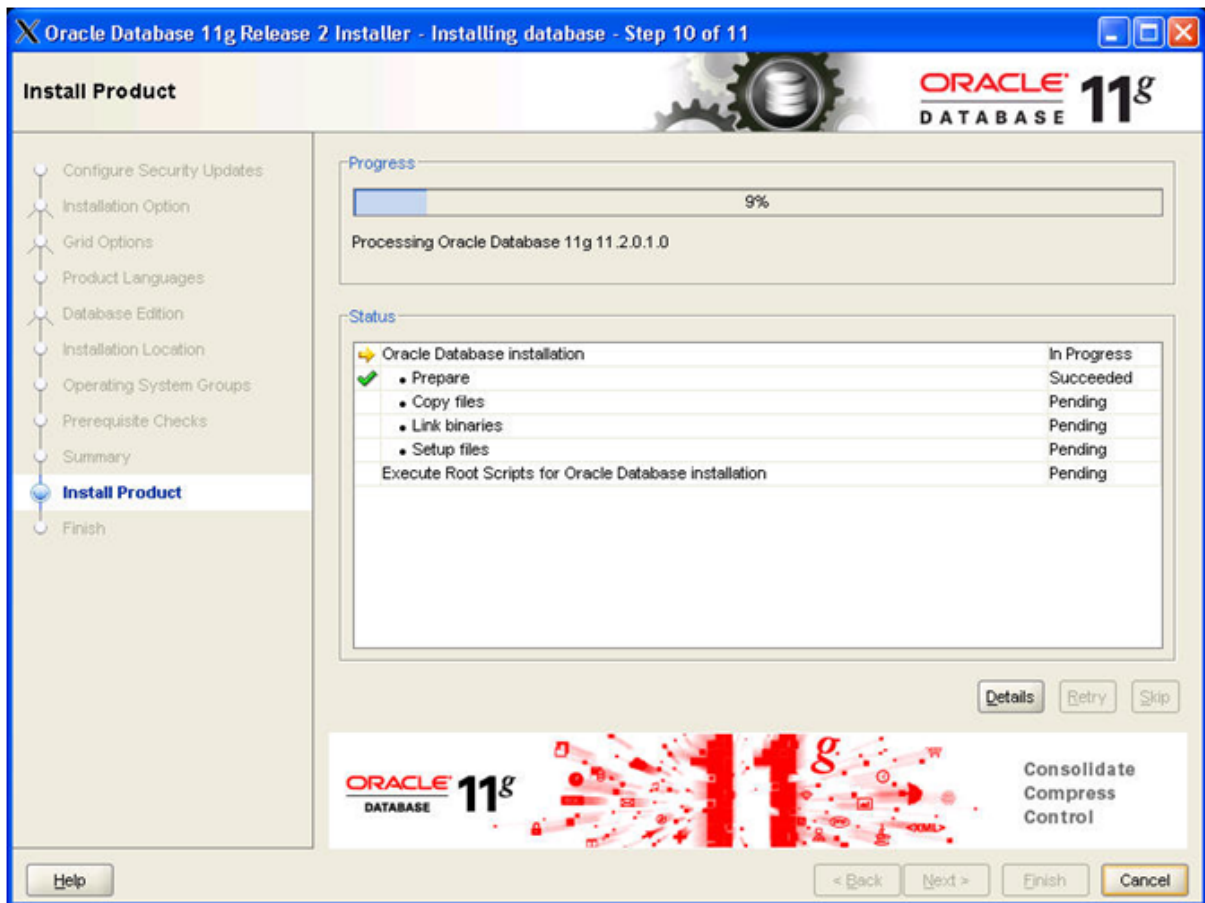
Note:

If you have failed checks marked as 'Fixable' click 'Fix & Check again'. This will bring up a window that instructs you to execute fixup scripts. Execute the runfixup.sh script as described on the screen as root user. Click 'Check Again' and if all checks succeeded click 'Next>' If you are sure the unsuccessful checks can be ignored tick the box 'Ignore All' before you click ' Next> '



Action:

Perform a last check that information on the screen is correct before you click 'Finish'

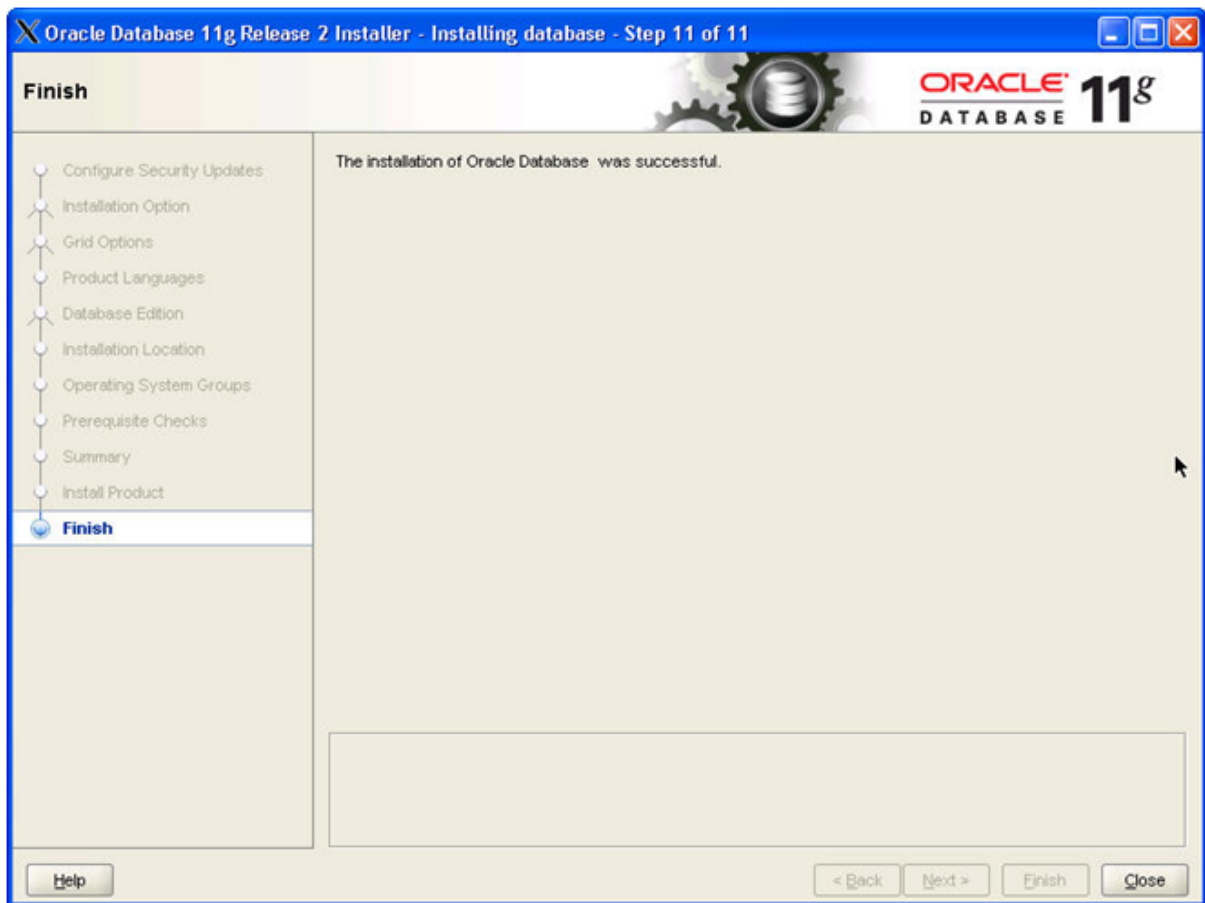


Action:

Log in to a terminal window as root user and run the root.sh script on the first node. When finished do the same for all other nodes in your cluster as well. When finished click 'OK'

Note:

root.sh should be run one node at a time.



Action:

Click 'Close' to finish the installation of the RDBMS Software.

7. RAC Home Patching

This Chapter is a placeholder

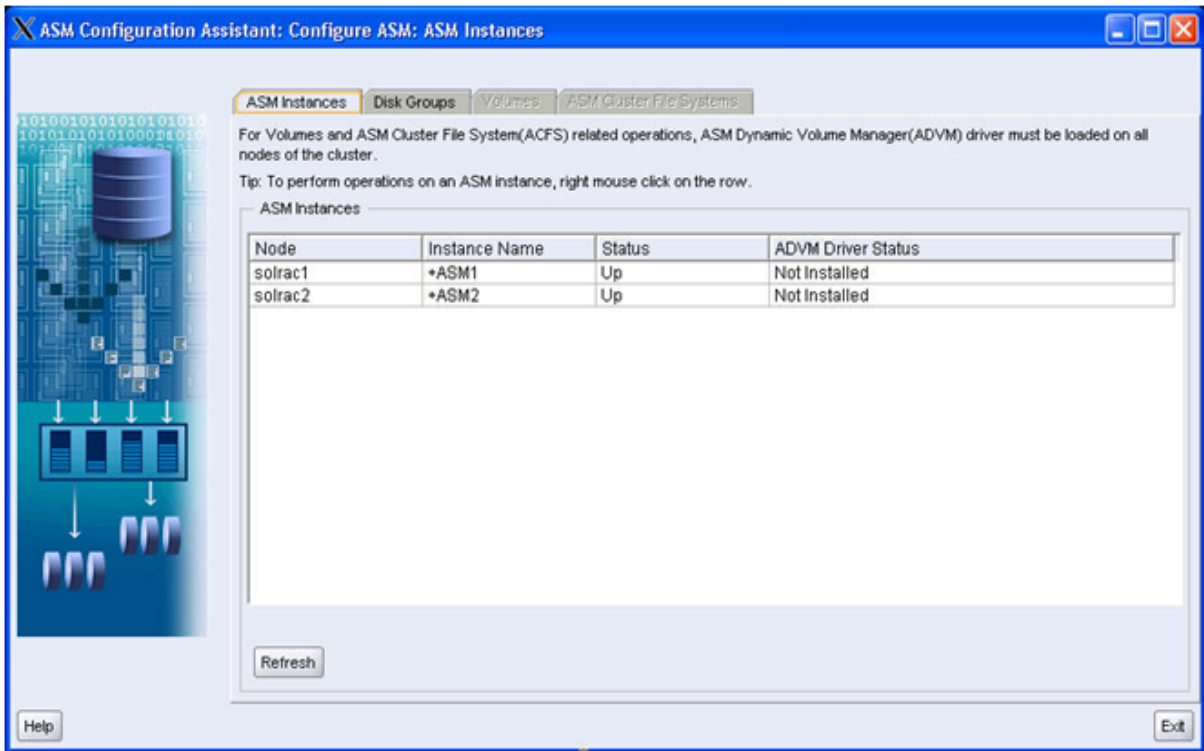
8. Run ASMCA to create diskgroups

As the grid user start the ASM Configuration Assistant (ASMCA)

```
#su - grid
```

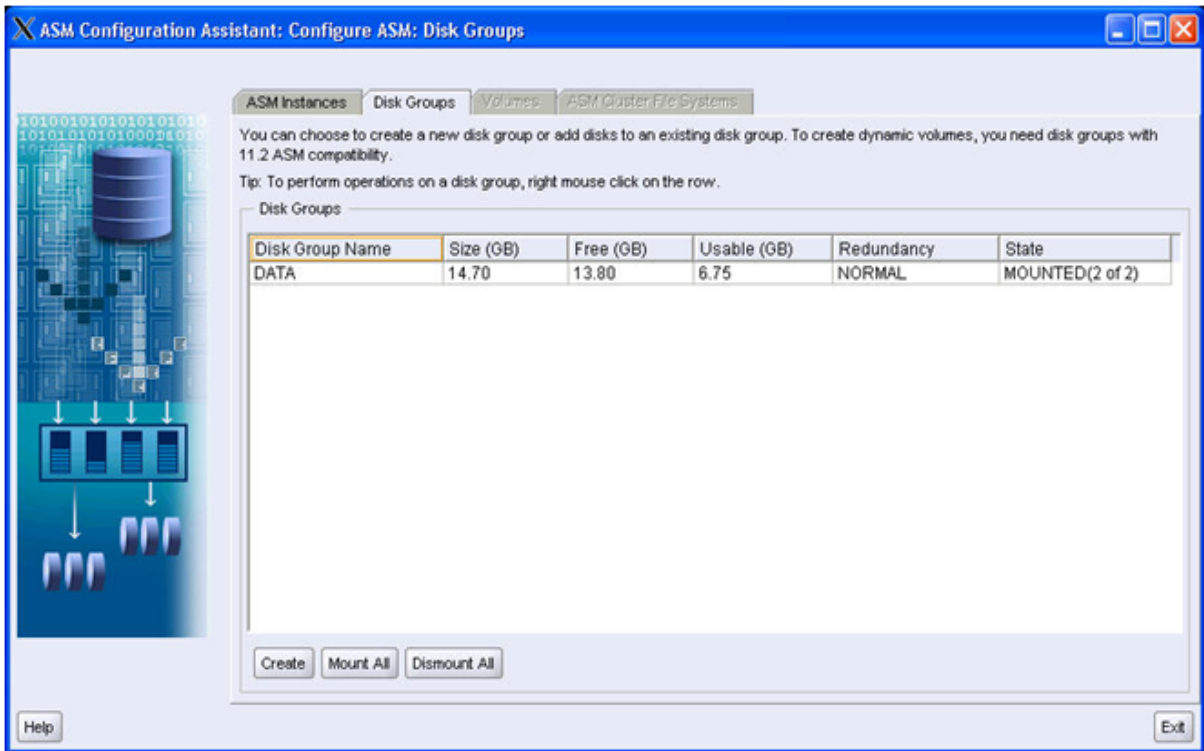
```
cd /u01/11.2.0/grid/bin
```

```
./asmca
```



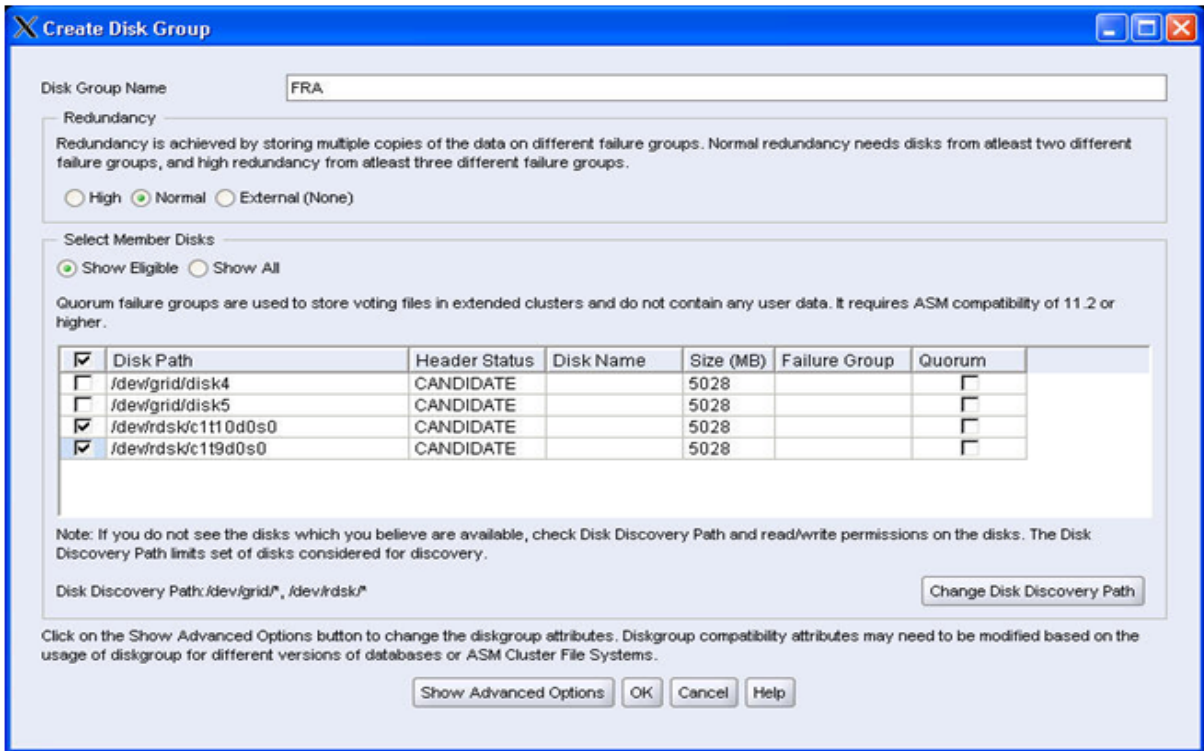
Action:

Click 'Disk Groups' tab



Action:

Click 'Create' to create a new diskgroup



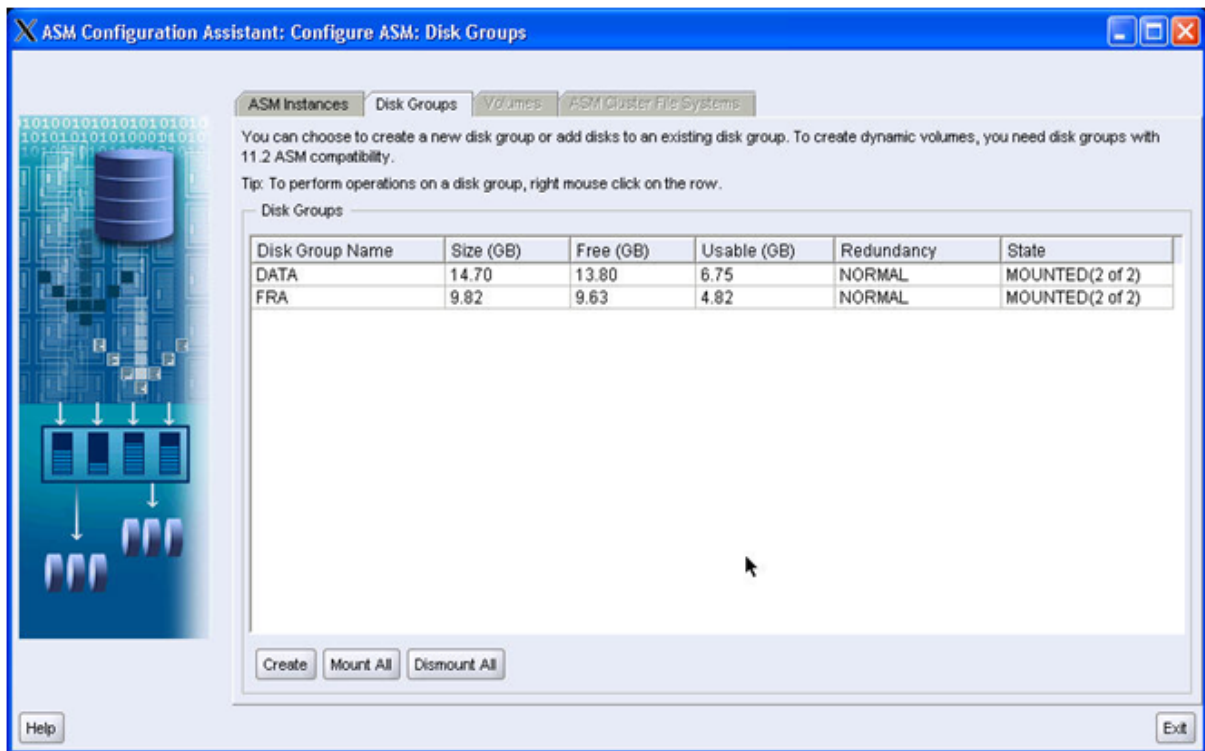
Action:

Type in a name for the disk group, select the redundancy you want to provide and mark the tick box for the disks you want to assign to the new disk group.




Action:

Click 'OK'



Action:

Click 'Exit'



Action:

Click 'Yes'

Note:

It is Oracle's Best Practise to have an OCR mirror stored in a second disk group. To follow this recommendation add an OCR mirror. Mind that you can only have one OCR in a diskgroup.

Action:

1. To add OCR mirror to an Oracle ASM disk group, ensure that the Oracle Clusterware stack is running and run the following command as root:
2. # ocrconfig -add +FRA
3. # ocrcheck

9. Run DBCA to create the database

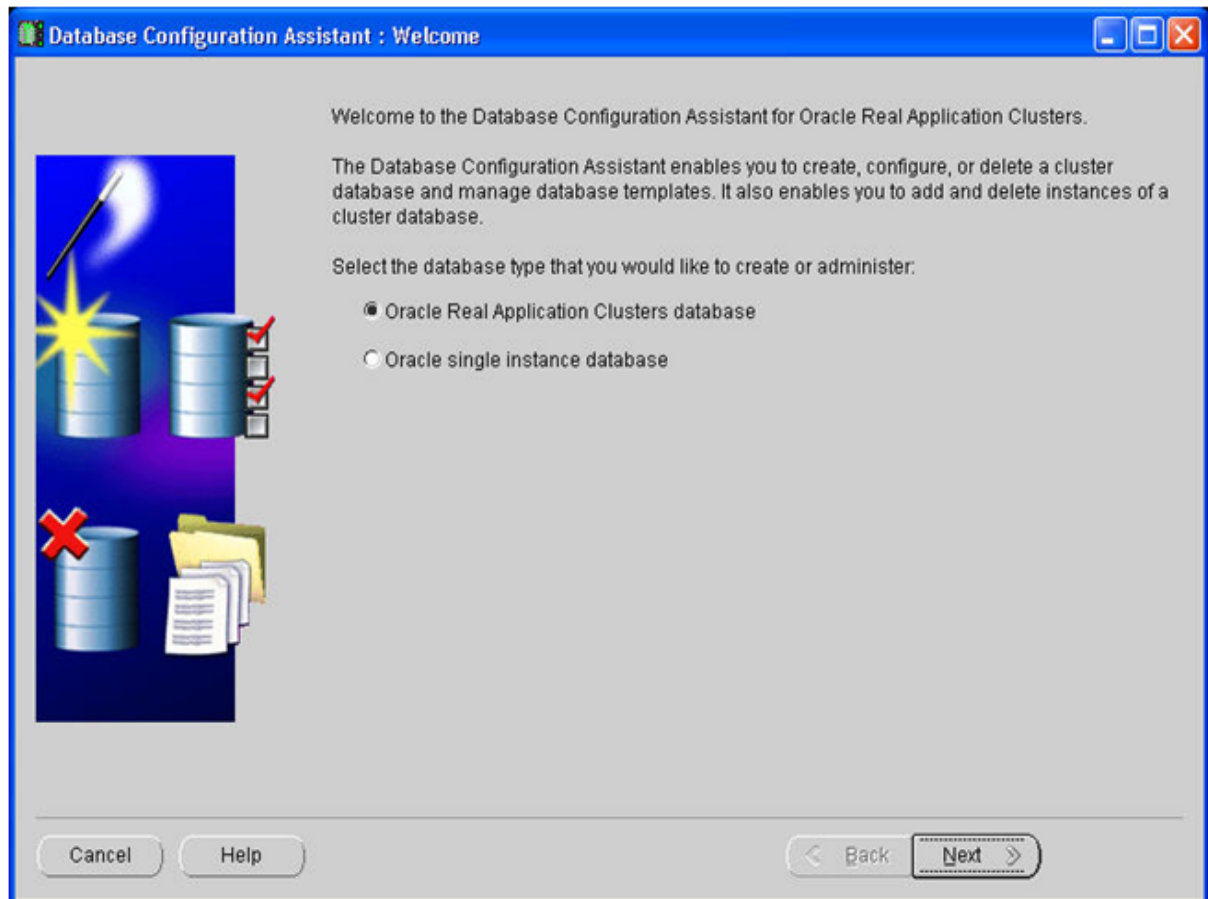
As the oracle user start the Database Configuration Assistant (DBCA)

#su - oracle

9. Run DBCA to create the database

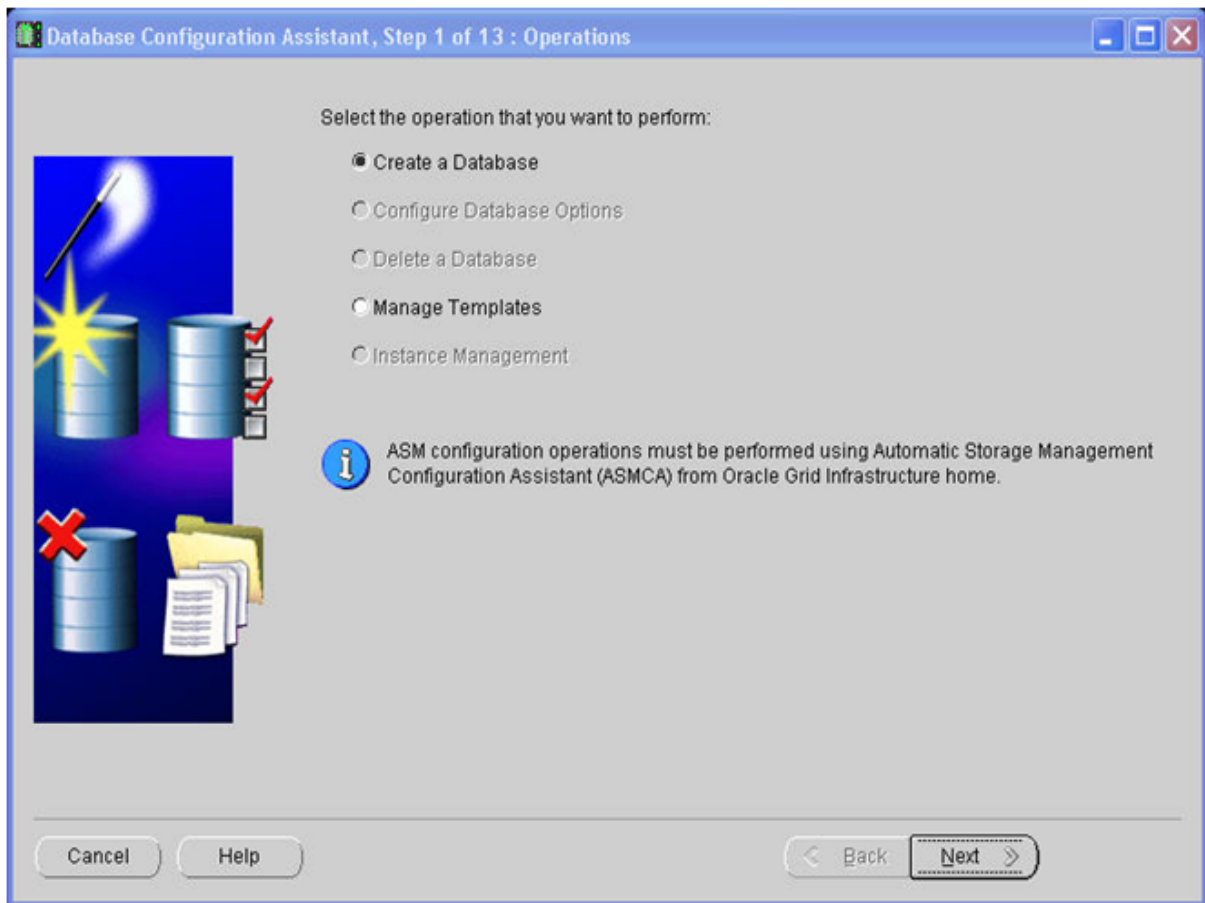
\$cd /u01/app/oracle/product/11.2.0/db_1/bin

./dbca

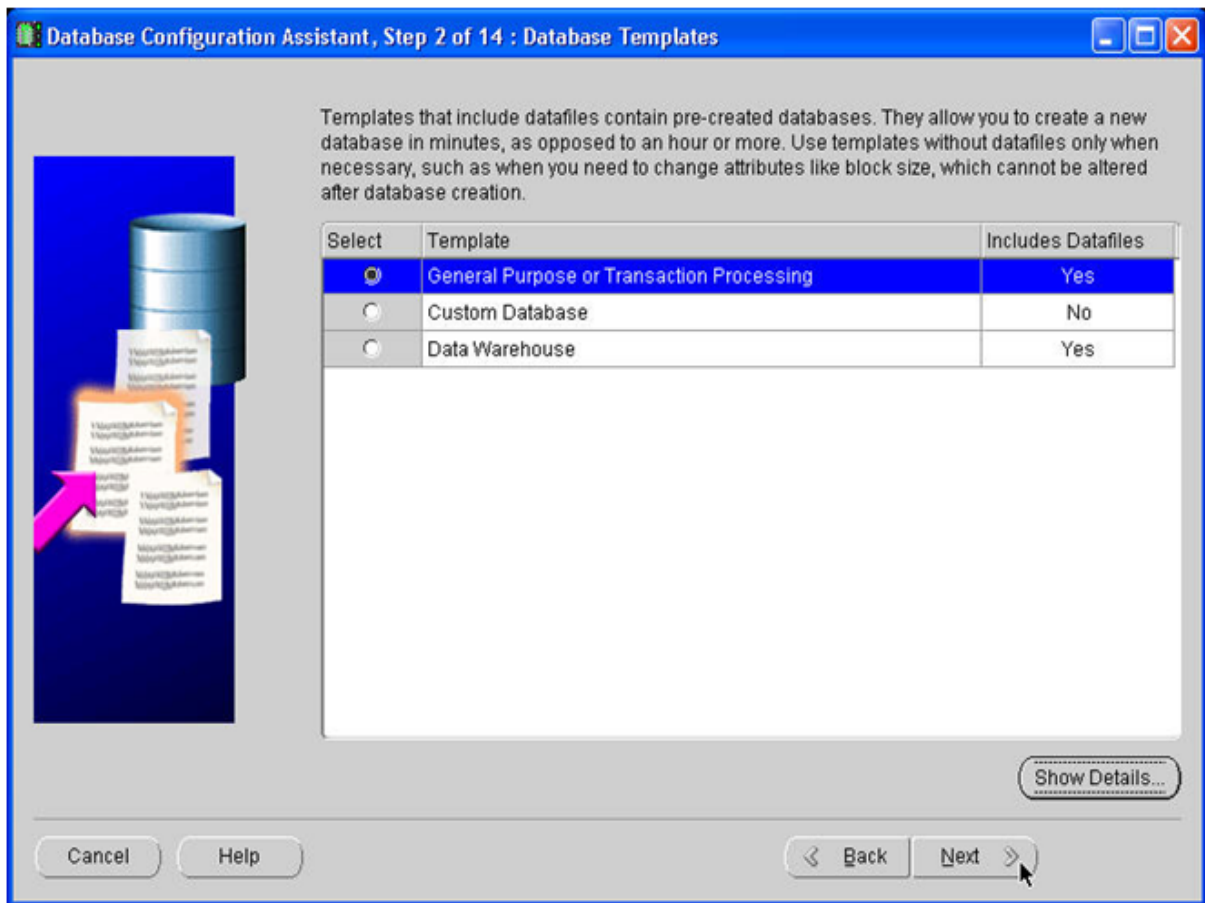


Action:

Select 'Oracle Real Application Clusters database' and click 'Next'

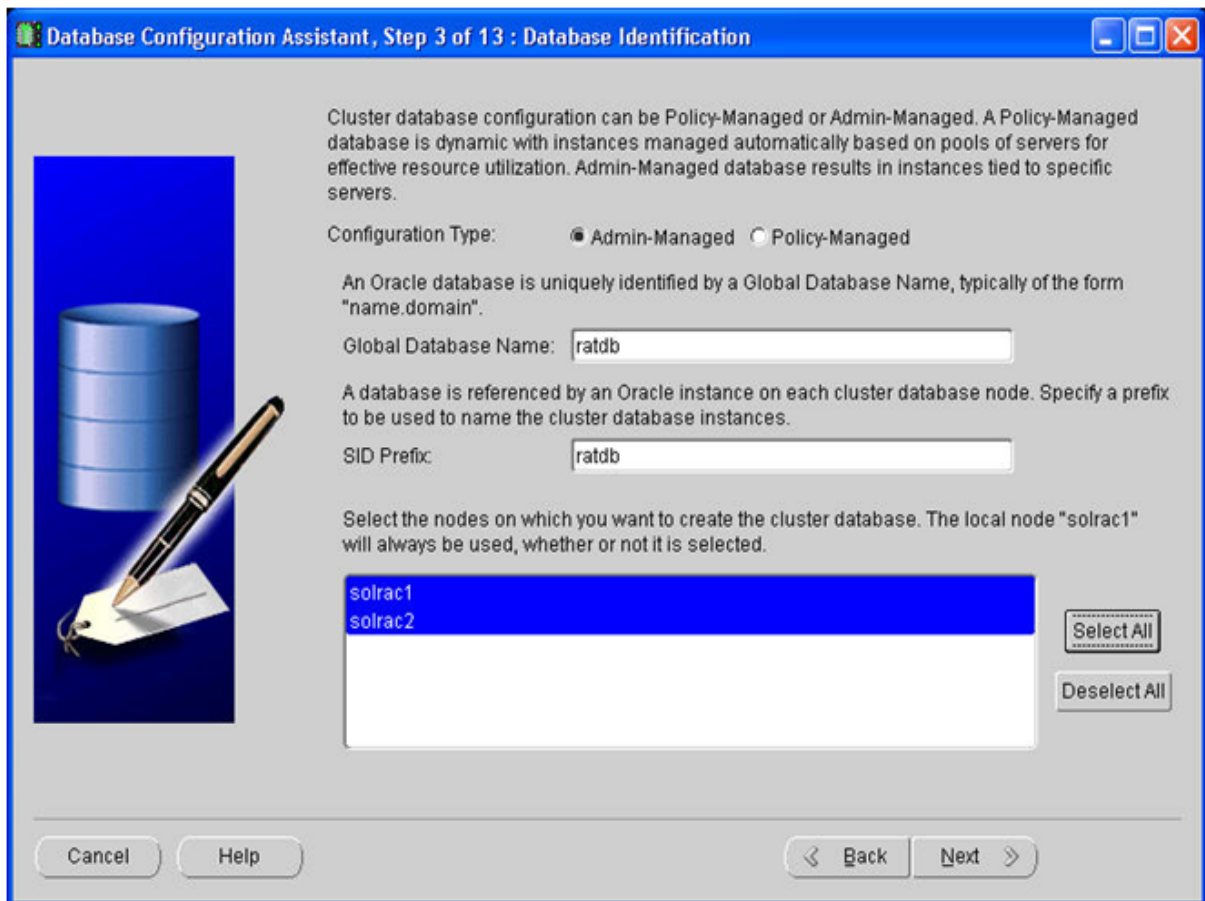


Action:
choose option 'Create a Database' and click 'Next'



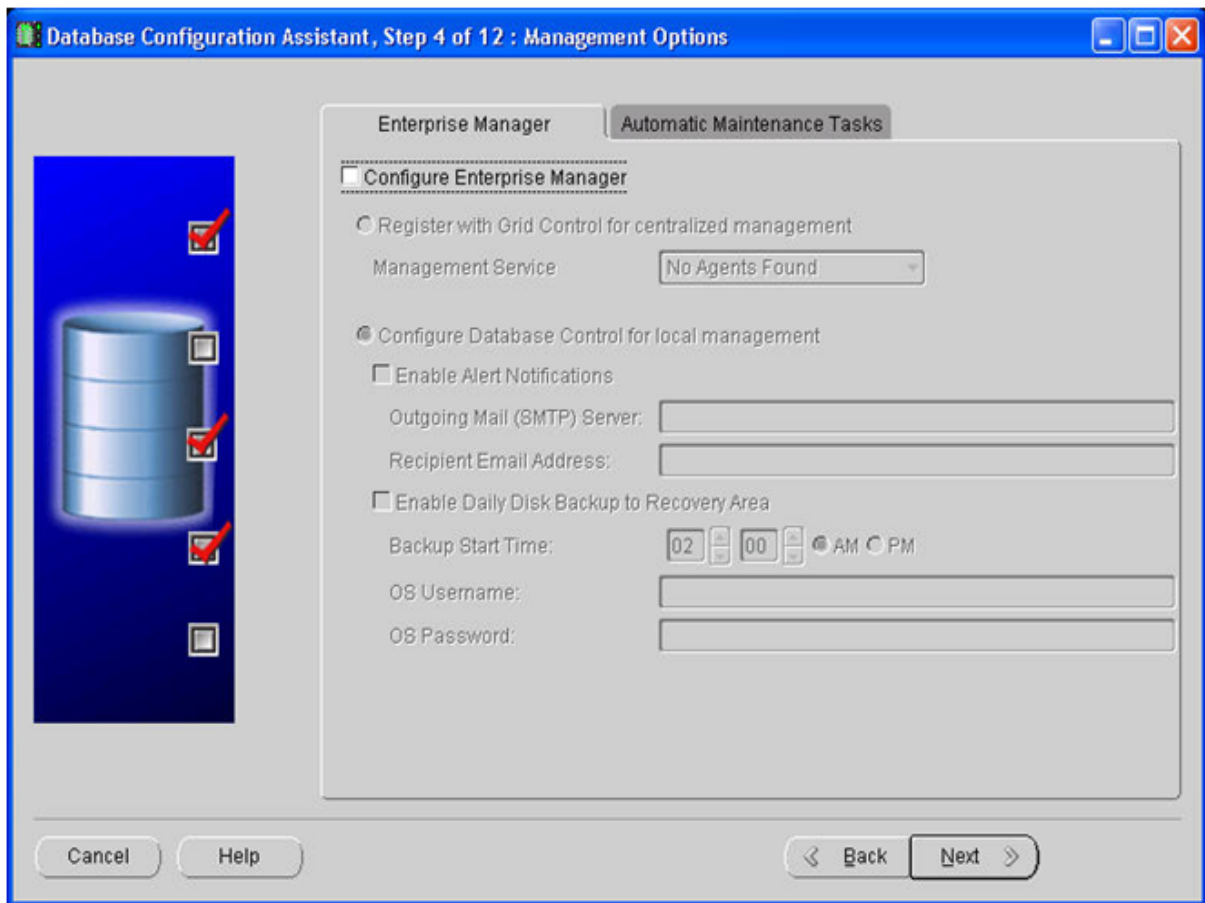
Action:

Select the database template that you want to use for your database and click 'Next'

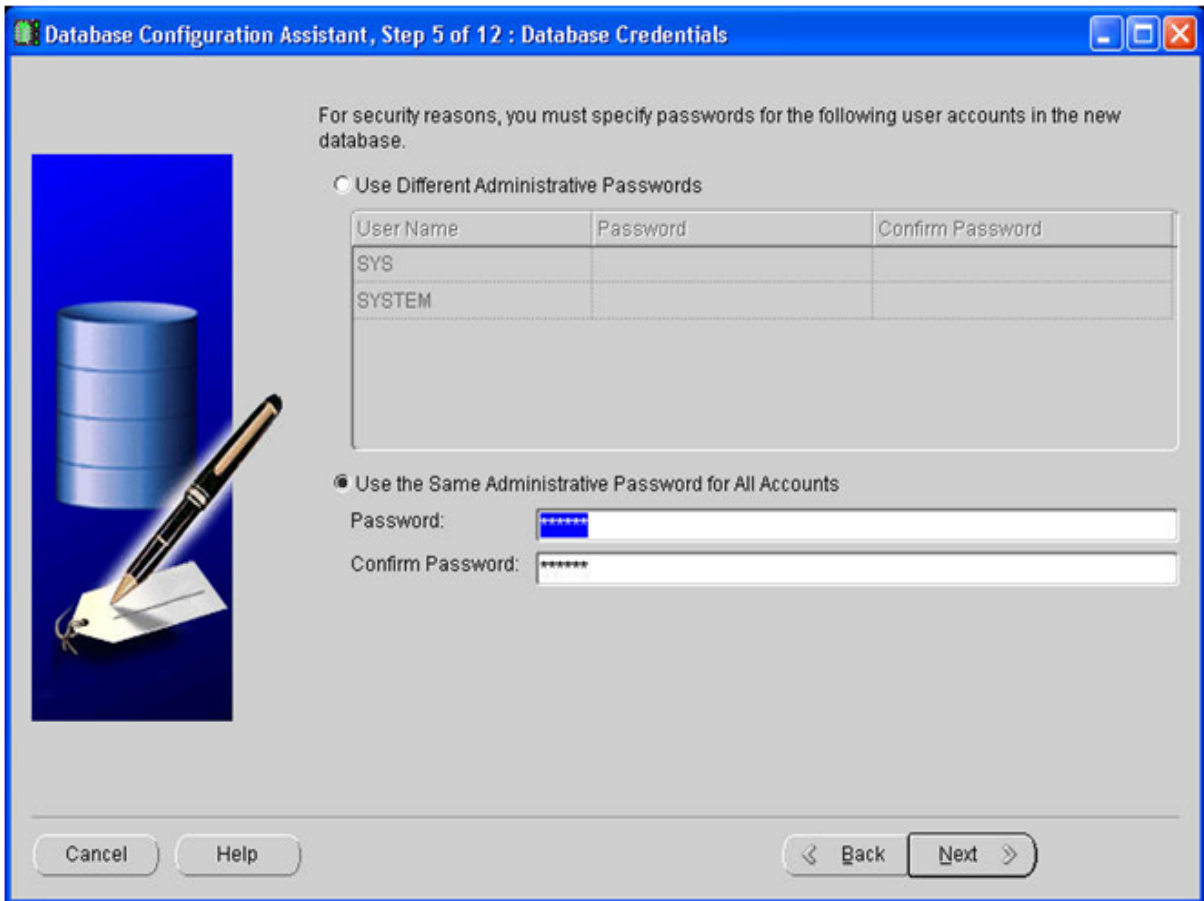


Action:

Type in the name you want to use for your database and select all nodes before you click 'Next'

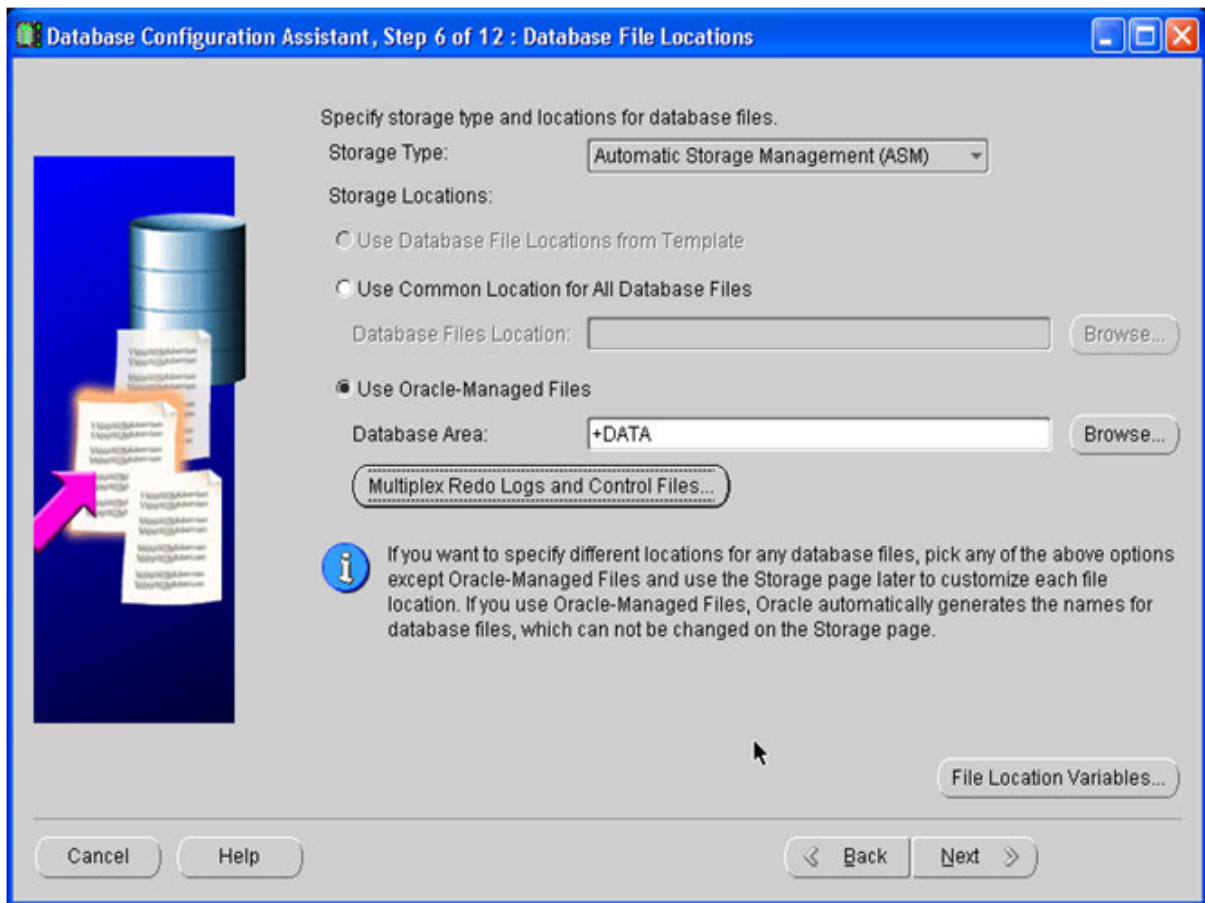


Action:
select the options you want to use to manage your database and click 'Next'



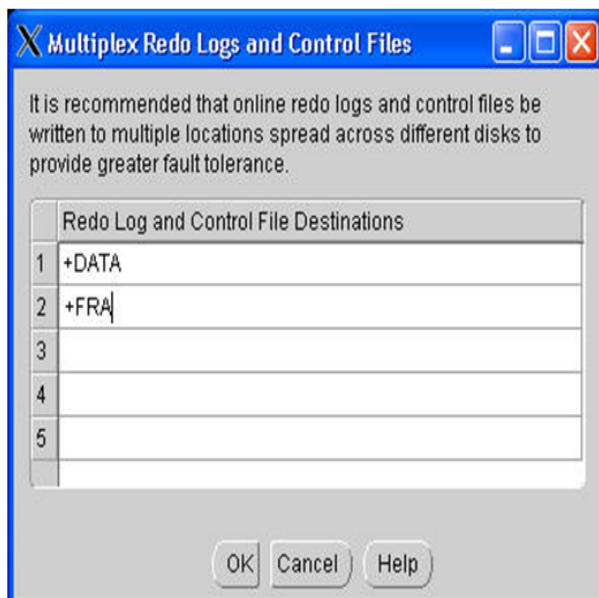
Action:

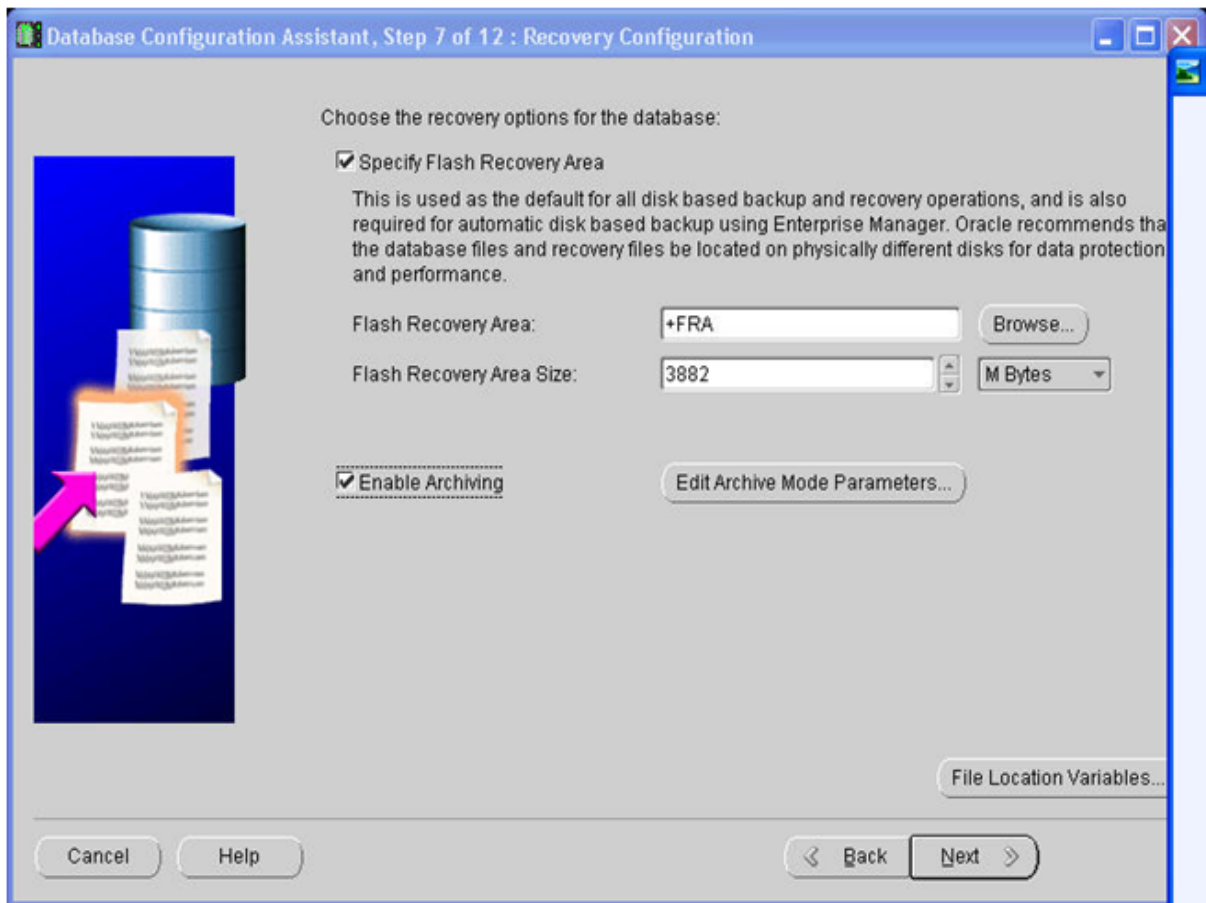
Type in the passwords you want to use and click 'Next'



Action:

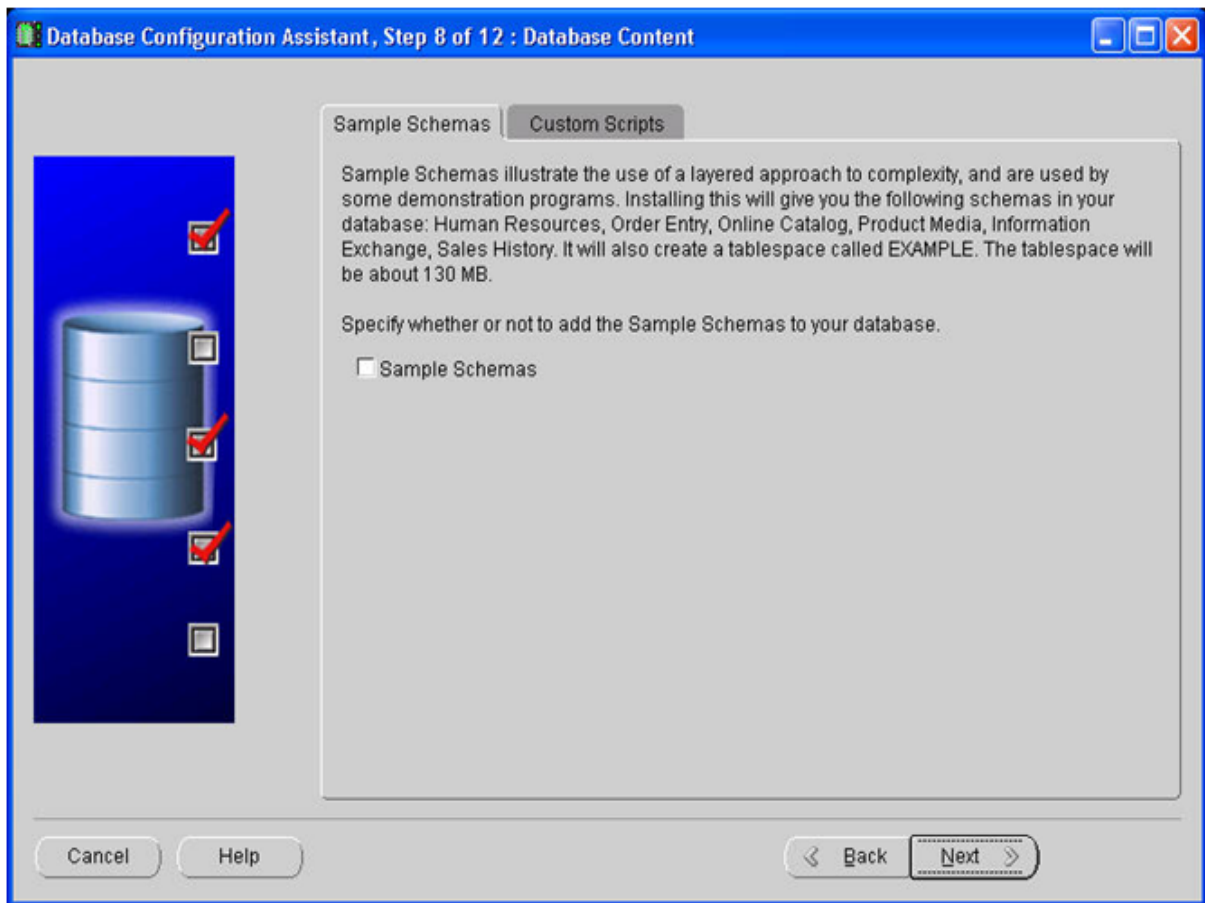
Select the diskgroup you created for the database files and click 'Multiplex Redo Logs and Control Files'. In the popup window define the diskgroup that should contain controlfiles and redo logfile and the diskgroup that should contain the mirrored files.





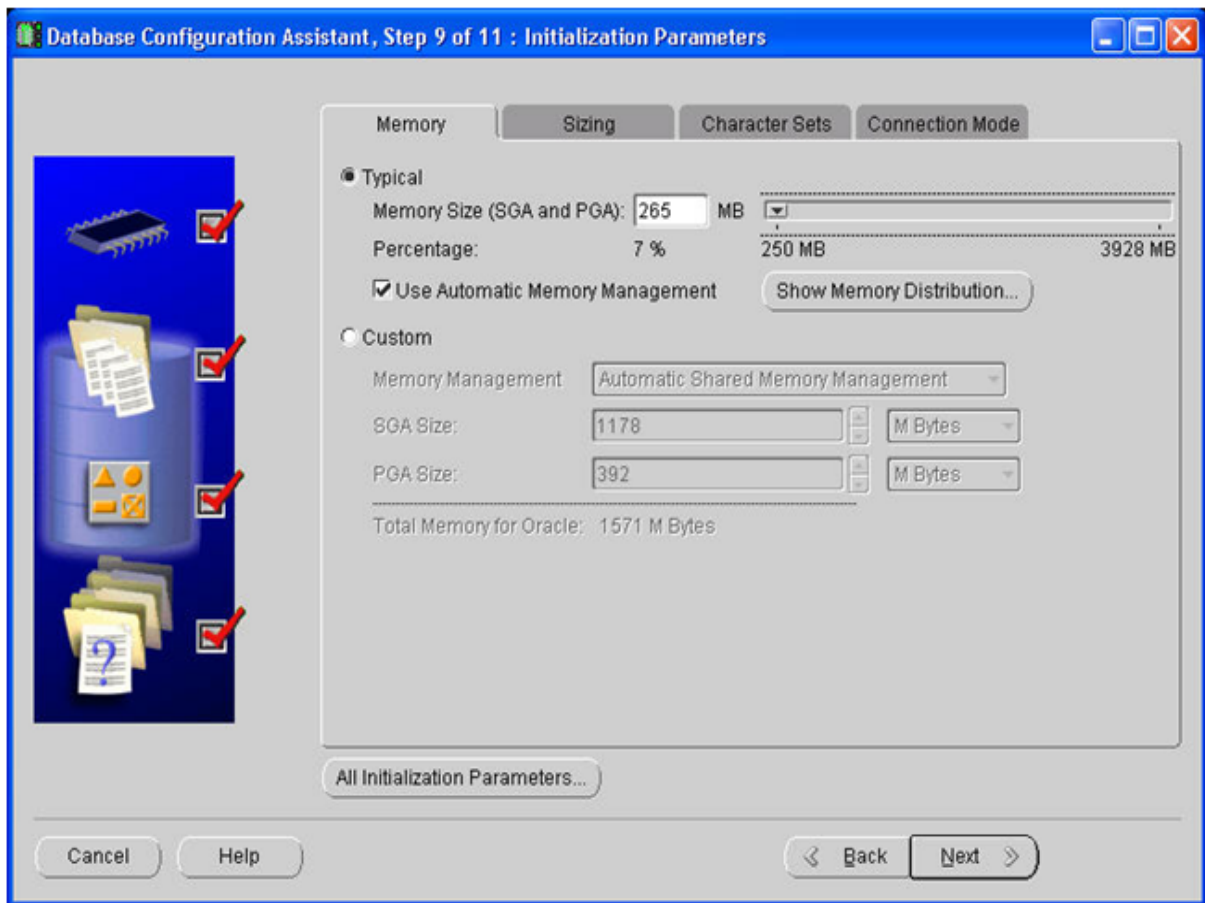
Action:

Specify the diskgroup that was created for the flash recovery area and define the size. If the size is smaller than recommended a warning will popup.



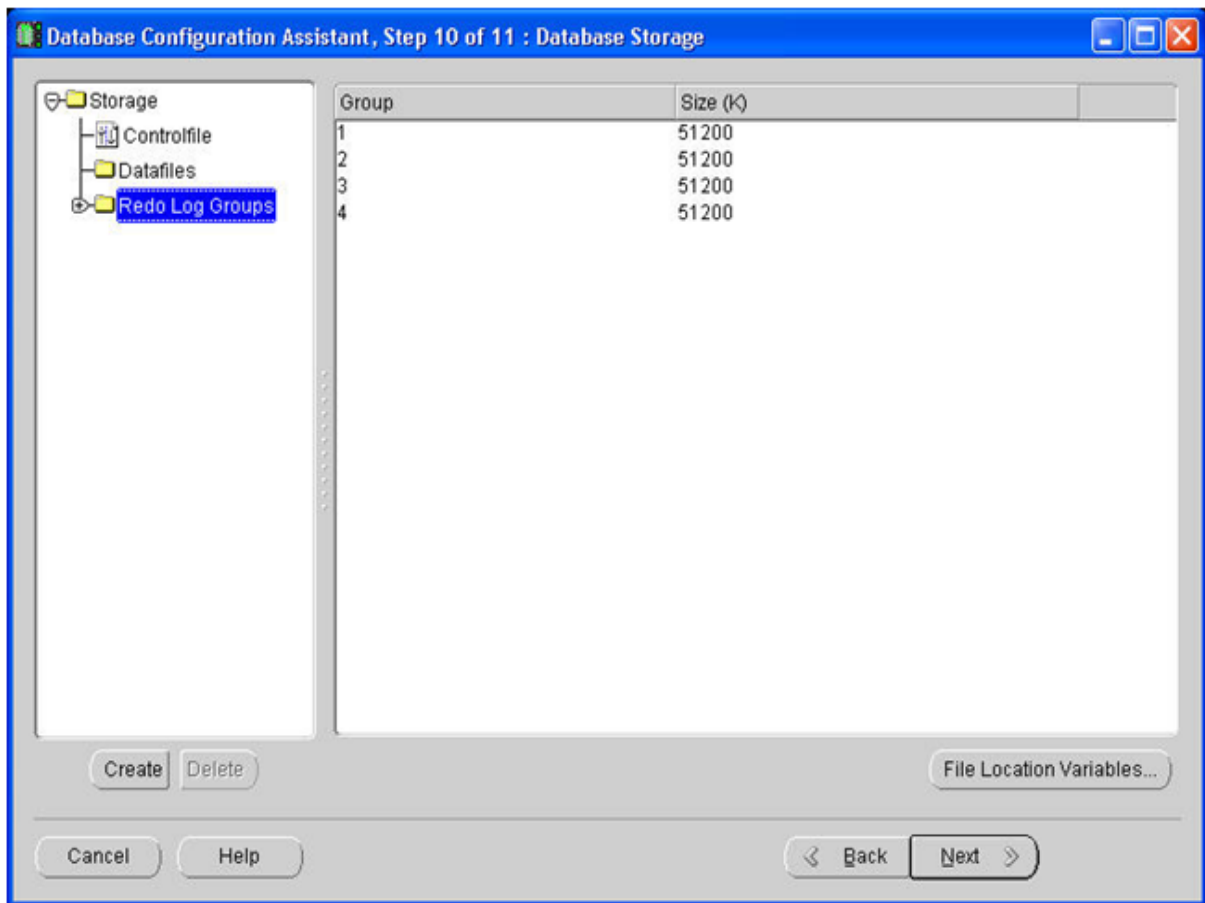
Action:

Select if you want to have sample schemas created in your database and click 'Next'

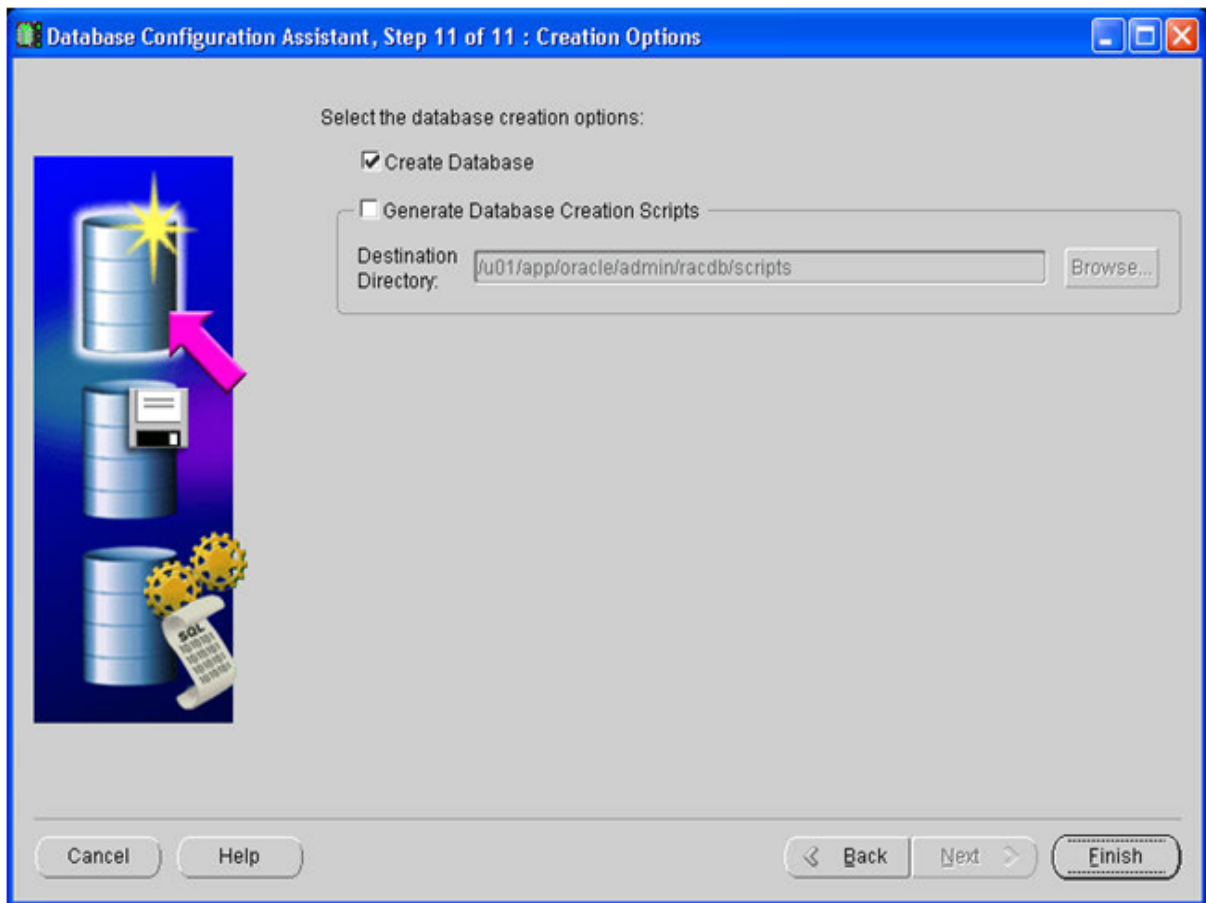


Action:

Review and change the settings for memory allocation, characterset etc. according to your needs and click 'Next'

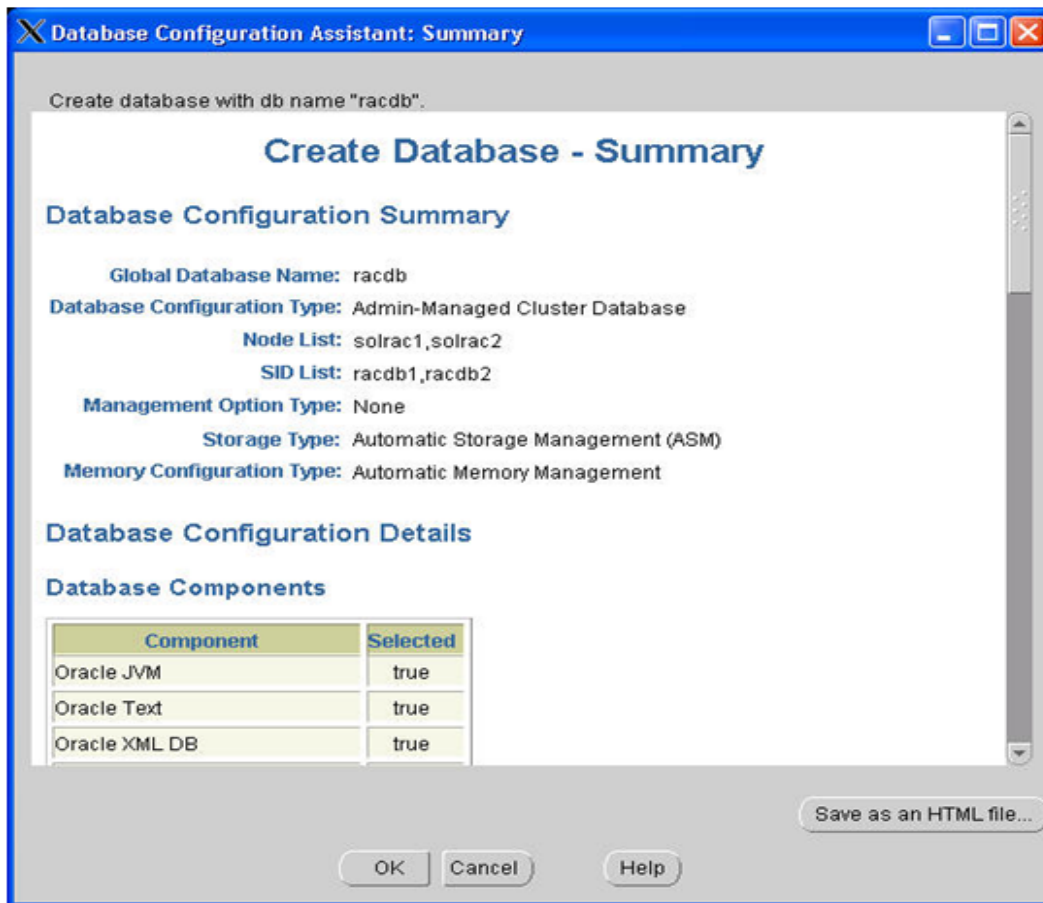


Action:
Review the database storage settings and click 'Next'



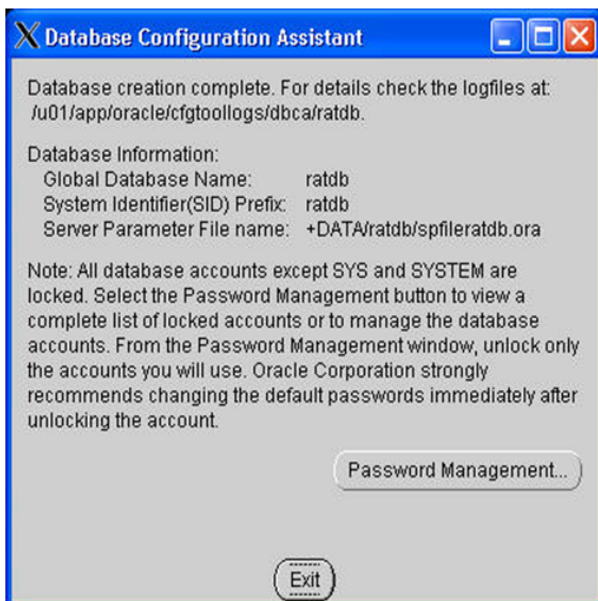
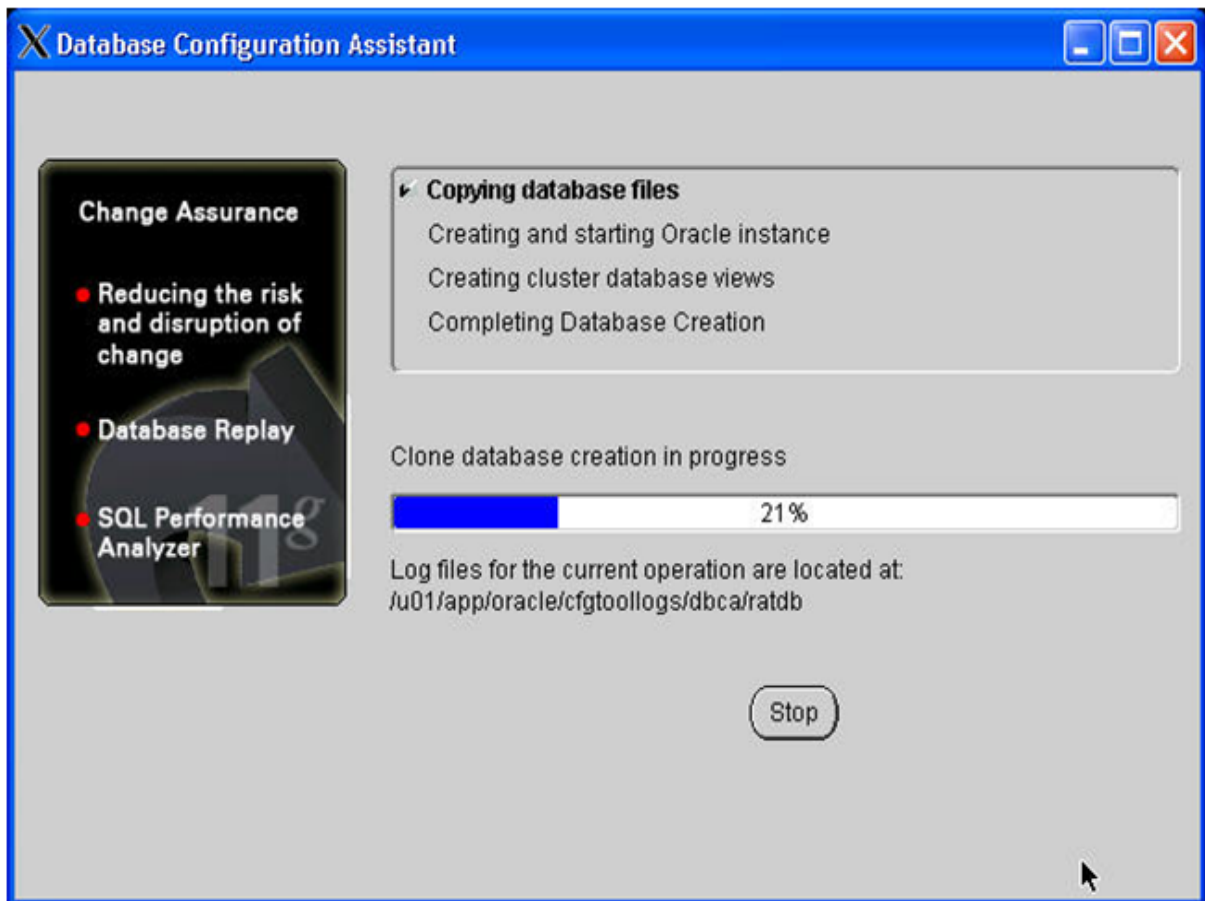
Action:

Ensure the tickbox 'Create Database' is ticked and click 'Finish'



Action:

Review again the database configuration details and click 'OK'



Action:

The database is now created, you can either change or unlock your passwords or just click Exit to finish the installation.