Technical report:

Setting up OracleRAC10*g*R2 over NFS on IBM System Storage N series

Best practices for using automatic storage management

Document NS3572-0

June 22, 2008

Table of contents

Abstract

*Oracle10*g *RDBMS has a new facility called Automatic Storage Management (ASM) which provides integrated cluster file system and volume management features. ASM complements the Oracle10*g *RDBMS with both volume and disk management utilities, removing the need for third-party volume management tools while also reducing the complexity of the enterprise architecture. ASM provides simplicity of managing volumes that may be composed of block-based devices (e.g., in a storage area network, or SAN) or file-based devices (e.g., in a network file system, or). These devices are the underlying storage for the Oracle RDBMS. Although ASM over SAN storage has been the more popular option, a customer has the choice of deploying ASM over NFS as well. IBM System Storage N series systems offer substantial advantages in an NFS environemtn. This paper facilitates a customer NFS choice with best-practice recommendations for deploying Oracle10*g *RDBMS over NFS on an IBM System Storage N series.*

Introduction

This technical report covers the installation of Oracle Database 10*g* Release 2 with Real Application Clusters (OracleRAC10*g*R2) and Oracle Automatic Storage Management (ASM) with Oracle Clusterware over a network file system (NFS) on IBM® System Storage™ N series. This is now a certified configuration and therefore the components presented in this paper must be used in the same combination to gain support from all parties involved. The only exception is the application of certain patches as defined and required by all the vendors in this configuration. This report also covers the patches and recommendations for running OracleRAC10*g*R2 and ASM on N series Storage in an NFS environment.

Assumptions

We assume that readers are familiar with OracleRAC10*g*R2 and with the operation of IBM N series storage systems. We also assume that readers are familiar with the operation of the Enterprise Linux® operating-system environment and installation of Oracle Database 10*g* patches and any relevant Enterprise Linux Red Hat Package Managers (RPMs). It is also important to be familiar with all networking terminology and implementations.

System environment

The configuration presented in this document is based on the following OracleRAC10*g*R2 certification environment.

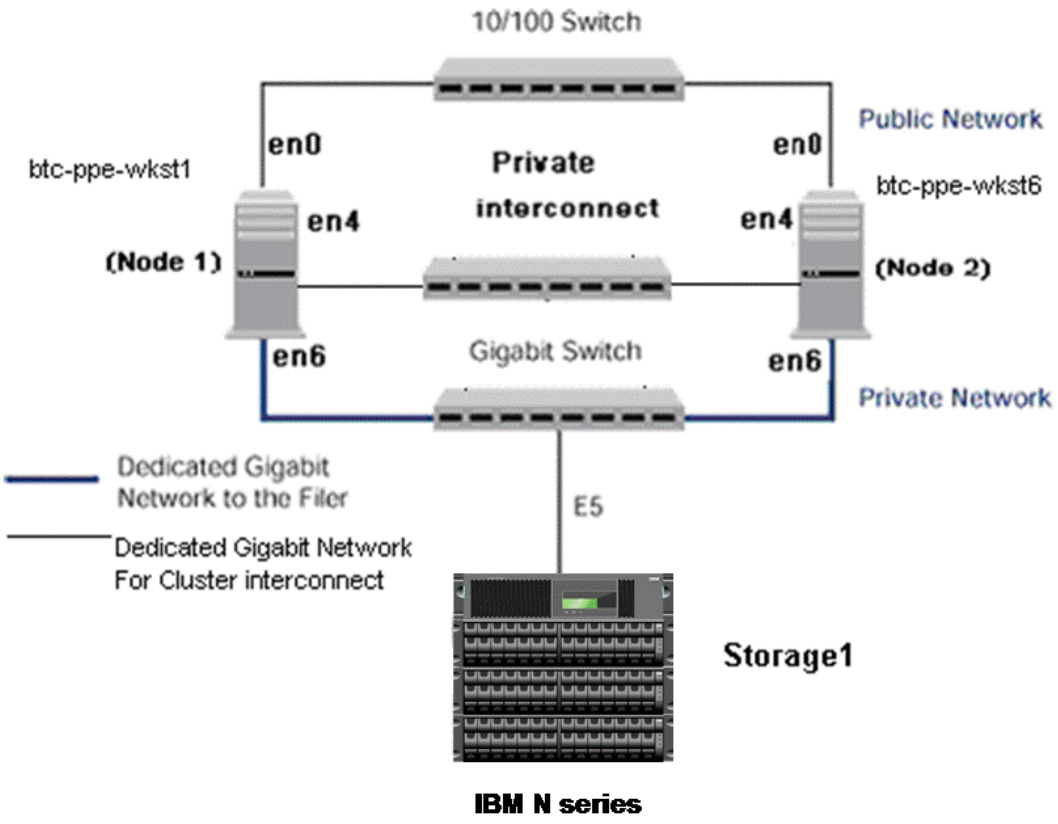

Figure 1) OracleRAC10gR2 on IBM xSeries (Intel x86) servers with IBM N series.

Figure 1 shows a typical configuration of OracleRAC10*g*R2 with N series storage and IBM xSeries® Intel® servers running the Oracle Enterprise Linux® (OEL4) operating system (OS). This configuration enables users to scale horizontally and internally in terms of processor, memory, and storage.

As shown in Figure 1, it is strongly recommended that you dedicate a private network connection between the OracleRAC10*g*R2 servers and the N series storage. This is accomplished by using a dedicated gigabit

network (with a gigabit switch) to the N series storage. A dedicated network connection is beneficial for the following reasons:

- In an OracleRAC10*g*R2 environment, it is important to eliminate contentions and latencies.
- Providing a separate network ensures security.

The cluster interconnect is an essential part of Oracle Database 10*g* clusters. Along with cache fusion, it is also used to monitor the heartbeat of the servers in the existing cluster group. This is a typical configuration that can be deployed in a customer's environment.

Requirements

Hardware

Cluster nodes:

- Two IBM xSeries 306 32-bit Intel® servers
- One 4-Port 10/100 Base-TX Ethernet PCI Adapter
- One 10/100/1000 Base-T Ethernet PCI Adapter (for private interconnect)
- One 10/100/1000 Base-T Ethernet PCI Adapter (connected to N series storage).

Storage Infrastructure:

- One IBM N series system with IBM System Storage N series with Data ONTAP® 7.2 or later
- One gigabit switch with at least four ports.
- One gigabit NIC in the system
- One or more disk shelves, based on the disk space requirements.

Software

For both nodes in the participating cluster unless specified otherwise:

- Oracle Enterprise Linux 4
- Oracle Database 10*g* Release 2 (10.2.0.1), with RAC license
- Oracle Database 10*g* Release 2 Patch 3.

Setup for IBM N series

1. Configure an IBM N series system running Data ONTAP 7.2 or above and also with NFS and IBM System Storage N series with SnapRestore® license keys.
2. Create and export volumes for storing database files on the storage:

   Create four volumes on the storage (Storage1) as listed below.

   orahome Shared Oracle and Cluster Registry Service (CRS) home (binaries)
   oradata Oracle data files and control files
   ora10g CRS files
   oralogs Database logs, a copy of control file and archive logs

   To create volumes, use the following command at the N series storage console:

   ```
   Storage1> vol create oradata 14
   ```

   Note: Volume oradata was created with 14 disks and volumes oralogs and orahome with eight disks each. You can create your volumes based on your workload needs.

   Add the following entries to the `/etc/exports` file on N series storage (Storage1):

   ```
   /vol/orahome -anon=0
   /vol/oradata -anon=0
   /vol/oralogs -anon=0
   /vol/ora10g  -anon=0
   ```

   Execute the following command at the storage system console:

   ```
   Storage1> exportfs -a
   ```

   Note: It is recommended that you use flexible volumes in your database environment. IBM System Storage N series with FlexVol™ technology pools storage resources automatically and enables you to create multiple flexible volumes on a large pool of disks. This flexibility means that you can simplify operations, gain maximum spindle utilization and efficiency, and make changes quickly and seamlessly.

   The database volume layout discussed in this document was defined for certification purposes, and your setup may vary depending upon requirements.

Operating-system configuration

Patches

Before your OracleRAC10*g*R2 installation, the following RPMs need to be applied on IBM servers. Some of these RPMs may already be applied to your system. Be sure to verify whether they already exist before applying them.

To determine whether the required RPMs are already installed and committed, enter a command similar to the following:

```
# rpm –qa | grep compat
```

If a patch is not installed and committed, then install it. Here is a list of required patches.

binutils-2.15.92.0.2-21
compat-libstdc++-296-2.96-132.7.2
compat-db-4.1.25-9
compat-libstdc++-33-3.2.3-47.3
make-3.80-6.EL4
glibc-common-2.3.4-2.25
glibc-2.3.4-2.25
glibc-devel-2.3.4-2.25
openmotif-2.2.3-10.RHEL4.5
setarch-1.6-1
control-center-2.8.0-12.rhel4.5
gcc-3.4.6-3.1
gcc-c++-3.4.6-3.1
gnome-libs-1.4.1.2.90-44.2
libstdc++-3.4.6-3.1
libstdc++-devel-3.4.6-3.1
libaio-0.3.105-2
pdksh-5.2.14-30.3
Note: Java 1.4.2 32-bit is installed with Oracle.

Operating-system settings

On Enterprise Linux systems, the default `ulimits` for individual users are set in

`/etc/security/limits.conf`. As a root user, add the following entries using root users:

```
# Oracle specific settings
oracle soft nofile  4096
oracle hard nofile 65536
oracle soft nproc   2047
oracle hard nproc  16384
oracle soft memlock 3145728
oracle hard memlock 3145728
```

This needs to be done on all nodes of the cluster. It may be necessary to reboot the server to activate updated limits. After you modify the settings, the ulimit –a command should display the following:

```
# ulimit –a

core file size        (blocks, -c) 0
data seg size         (kbytes, -d) unlimited
file size                 (blocks, -f) unlimited
max locked memory     (kbytes, -l) unlimited
max memory size       (kbytes, -m) unlimited
open files                    (-n) 4096
pipe size          (512 bytes, -p) 8
stack size            (kbytes, -s) unlimited
cpu time             (seconds, -t) unlimited
max user processes            (-u) 15168
virtual memory        (kbytes, -v) unlimited
```

Verify the above setting for the `oracle` user.

Kernel settings

Add the following parameters for the shared memory and semaphores to the `/etc/sysctl.conf` file using root user.

```
kernel.shmall = 2097152
kernel.shmmax = 2147483648
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
fs.file-max = 65536
net.ipv4.ip_local_port_range = 1024 65000
net.core.rmem_default = 1048576
net.core.rmem_max = 1048576
net.core.wmem_default = 262144
net.core.wmem_max = 262144
```

Note: Set the parameter with the specified value in the /etc/sysctl.conf file:

`vm.lower_zone_protection (= 100)`

This parameter is to increases the default page threshold from 16 MB to 100 MB to ensure that kernel is available with sufficient memory during heavy online transaction processing (OLTP) operations on the system.

Preinstallation setup tasks (cluster nodes)

This information is intended for both nodes in the participating cluster unless specified otherwise.

1. Have two IBM xSeries Intel servers ready with the latest recommended patches and operating-system settings as described in another section of this report.
2. Install and configure NICs in the cluster nodes (three per node).
   a. Public IP: As indicated by name.
   b. Private interconnects: Connect one gigabit NIC to the gigabit switch of all the nodes for cluster interconnects.
   c. Server connection to N series storage: Connect one gigabit NIC to the gigabit switch, which will connect to the gigabit NIC on the N series storage.
3. Configure the network interfaces on each node.
   a. Configure the three network interfaces:

```
# btc-ppe-wkst1 (Host 1)
en0 – ip: 10.73.68.155, netmask 255.255.254.0
en6 - ip: 10.73.69.155, netmask: 255.255.255.0
en4 - ip: 192.168.73.1, netmask: 255.255.255.0
# btc-ppe-wkst6 (Host 2)
en0 – ip: 10.73.68.156, netmask: 255.255.254.0
en6 - ip: 10.73.69.156, netmask: 255.255.255.0
en4 - ip: 192.168.73.2, netmask: 255.255.255.0
```

   Where:

   Interface en0 is the public IP for each node.

   Interface en6 on both cluster nodes is connected to the gigabit switch for storage I/O.

   Interface en4 on both cluster nodes is connected to the gigabit switch for cluster private interconnects.

   b. Update the `/etc/hosts` file on the cluster nodes and add entries for public, private, and virtual IP (VIP) addresses.

   Note: In addition to the preconfigured public and private network, Oracle Database 10*g* requires additional IP addresses that will be mapped to the public address as VIPs. If a node fails when an application or user makes a connection using a VIP, the Oracle Clusterware will transfers the VIP address to another surviving instance. You should add the VIP to the `/etc/hosts` file on all nodes in the cluster as well as all nodes that access the database. VIP must have the same subnet as the public IP address of the database host.

   An example of `/etc/hosts` entries:

```
# Internet Address Hostname # Comments
10.73.68.155        btc-ppe-wkst1       btc-ppe-wkst1.btcppe.ibmn.com
10.73.69.155        btc-ppe-wkst1-en6
192.168.73.1        btc-ppe-wkst1-i     btc-ppe-wkst1-i.btcppe.ibmn.com
10.73.68.195        btc-ppe-wkst1-v     btc-ppe-wkst1-v.btcppe.ibmn.com
10.73.68.156        btc-ppe-wkst6       btc-ppe-wkst6.btcppe.ibmn.com
10.73.69.156        btc-ppe-wkst6-en6
192.168.73.2        btc-ppe-wkst6-I      btc-ppe-wkst6-i.btcppe.ibmn.com
10.73.68.196        btc-ppe-wkst6-v      btc-ppe-wkst6-v.btcppe.ibmn.com
10.73.69.105        Storage1
```

4. Use the `ping` command to ensure the connectivity of each interface (interconnects, public IPs, and storage).

5. Create NFS mount points and mount the volumes with the following mount options on all the cluster nodes. As a root user, update the `/etc/fstab` file on all server nodes and add the following entries:

```
Storage1:/vol/oradata  /oradata      nfs
hard,nointr,proto=tcp,suid,vers=3,rw,bg,rsize=32768,wsize=32768,actimeo=0,ti
meo=600

Storage1:/vol/orahome  /orahome      nfs
hard,nointr,proto=tcp,suid,vers=3,rw,bg,rsize=32768,wsize=32768,actimeo=0,ti
meo=600

Storage1:/vol/oralogs       /oralogs        nfs
hard,nointr,proto=tcp,suid,vers=3,rw,bg,rsize=32768,wsize=32768,actimeo=0,ti
meo=600

Storage1:/vol/ora10g        /ora10g         nfs
hard,proto=tcp,suid,vers=3,nointr,rw,bg,rsize=32768,wsize=32768,noac,timeo=6
00
```

Where:

Storage1 is the name of the N series storage system.

oradata, oralogs, orahome, and ora10g are the mount points on the cluster nodes.

ora10g is just a separate mount point for CRS files. CRS and cluster synchronization services (CSS) files (cluster registry file and voting disk file) can reside in the same `/ora10g` volume but they must be mounted with the `noac` mount option.

During the Oracle CRS installation, be sure to indicate a path starting with the `/ora10g` directory when prompted for the OCR and voting disk file location.

6. **Note:** If you dynamically mount the NFS volumes without adding entries in the /etc/fstab file, the Oracle installation will fail. Create the following mount points on all cluster nodes:

```
#mkdir /oradata
#mkdir /oralogs
#mkdir /orahome
#mkdir /ora10g
```

Mount exported volumes on the mount points created above on all the cluster nodes. It is always a good idea to verify mount options by using the `mount` command on each node. After the NFS volumes are mounted, change the ownership of these mounted volumes to `oracle` user and `dba group`.

Installation procedure

The following sections explain the steps to install OracleRAC10*g*R2 on Enterprise Linux.

Preparing to install the Oracle RAC on clusters nodes

1. This document assumes the `oracle` user account and the group to be `oracle` and `dba`, respectively, on both cluster nodes. The user ID and group name for the oracle account should be the same on both cluster nodes. A sample `oracle` user `.bash_profile` file is provided in the Appendix. Make sure that the user profile file exports at least the ORACLE_BASE, ORACLE_PRODUCT, ORACLE_HOME, ORACLE_SID, ORA_CRS_HOME, and PATH entries.
2. Grant appropriate permissions to the `oracle` user on all shared mounted volumes, /oradata, /orahome, and /oralogs:

   ```
   #chown -R oracle:dba /oradata
   #chmod -R 755 /oradata
   ```

   Repeat this step for the orahome and oralogs volumes.
3. Set up account equivalence between the cluster nodes for the `oracle` user account. Add the following entries to the `/etc/hosts.equiv` file on all cluster nodes:

   ```
   btc-ppe-wkst1 oracle
   btc-ppe-wkst6 oracle
   btc-ppe-wkst1-i oracle
   btc-ppe-wkst6-i oracle
   ```
4. Log in as `oracle` user from both cluster nodes and then test the `oracle` user account equivalence by using a remote shell utility such as `rsh`:

   ```
   btc-ppe-wkst1:

   #su – oracle
   $rsh btc-ppe-wkst1 pwd
   $rsh btc-ppe-wkst6 pwd
   $rsh btc-ppe-wkst1-i pwd
   $rsh btc-ppe-wkst6-i pwd

   btc-ppe-wkst6:

   #su – oracle
   $rsh btc-ppe-wkst1 pwd
   $rsh btc-ppe-wkst6 pwd
   $rsh btc-ppe-wkst1-i pwd
   $rsh btc-ppe-wkst6-i pwd
   ```

Installing OracleRAC10*g*R2 CRS

For detailed information on installing Oracle CRS on Linux, refer to the appropriate Oracle RAC installation and configuration guide (located, at the time of this report writing, at http://otn.oracle.com/docs/content.html). This section briefly describes the procedures for using Oracle Universal Installer (OUI) to install CRS.

Note: The CRS home that you identify in this phase of the installation is only for CRS software; this home cannot be the same home as the OracleRAC10*g* home. That is, ORACLE_HOME and CRS HOME must be different locations.

1. Run the `runInstaller` command from the `/crs` subdirectory on the Oracle CRS R2 (10.2.0.1) CD-ROM or from the staging area where Oracle CRS software has been dumped. This is a separate CD that contains the CRS software. This document assumes that OUI is started from node 1 (btc-ppe-wkst1). When OUI displays the Welcome page, click Next.
2. On the Specify Inventory page, enter a nonshared location for Oracle Inventory. This is the only part of Oracle Database 10*g* that should not be shared. For this test, we used `/home/oracle/oraInventory` for the Oracle Inventory information. Click Next.
3. The Specify File Locations page contains predetermined information for the source of the installation files and the target destination information. Specify the destination path for the shared CRS home. The path should be on a shared file system and different from $ORACLE_HOME. In this exercise, the shared CRS home was `/orahome/ora10g/product/10.2.0/crs_1`
4. On the next screen, specify the cluster name, public interface names (hostnames), private interface names and virtual interface hostnames to be used for the cluster interconnect. In our case, the public names are `btc-ppe-wkst1` and `btc-ppe-wkst6`, the private names are `btc-ppe-wkst1-i` and `btc-ppe-wkst6-i`, and the virtual hostnames are `btc-ppe-wkst1-v` and `btc-ppe-wkst6-v`. Click next to continue.
5. On the Network Interface Usage page, specify the private network to be used for the cluster interconnect. This is an important step. Do not leave it set to the default, which is Do Not Use. In our case, `eth1 (btc-ppe-wkst1-i)` was used as the private interconnect; `eth0 (btc-ppe-wkst1)` was used as the public interface. Select the interface and click the Edit button to modify it. Click Next.
6. On the Oracle Cluster Registry page, specify the OCR file. Be sure to specify the full path to a shared location along with the name of the file. Do the same for a mirror file if you want normal redundancy. In our case, we used `/ora10g/ocrfile and /ora10g/ocrfile_mirror`. Click Next.
7. On the Voting Disk page, specify the CSS voting disk file location. We used `/ora10g/cssfile` for CSS services. In case of normal redundancy, specify the path along with name. Click Next to install the CRS.
8. When prompted, run the following script as root user starting from primary node:

```
/orahome/ora10g/orainventory/orainstRoot.sh
/oarhome/ora10g/product/10.2.0/crs_1/root.sh
```

9. In the Configuration Assistant window, you may see some warnings. Click OK to continue.
10. Run the `vipca` utility from the `$ORA_CRS_HOME/bin` directory as root user on the master Node (btc-ppe-wkst1). Click Next.
11. Select the Public Interface. Click Next.
12. Specify the VIP address and Subnet Mask of each node. Click Next.
13. Click Finish to continue VIPCA.
14. Click OK and then exit to finish VIPCA.
15. To verify your CRS installation by execute the `olsnodes` command from the `$CRS_HOME/bin` directory. The `olsnodes` command syntax is:

```
olsnodes [-n] [-l] [-v] [-g]
```

Where:

```
-n displays the member number with the member name
-l displays the local node name
-v activates verbose mode
-g activates logging
```

The output from this command should be a list of the nodes on which CRS was installed.

Installing Oracle10*g*R2 software

1. After making sure that Oracle CRS have started on the cluster nodes, start `runInstaller` from Disk1 of the Oracle10*g*R2 CDs or from the staging area where you have kept the Oracle Database 10*g* downloads.
2. On the Specify File Locations screen, enter the destination path for the shared ORACLE_HOME. This should be a different location than the shared CRS Home. For this exercise, the shared ORACLE_HOME was `/orahome/ora10g/product/10.2.0/db_1`.
3. On the next screen, select Cluster Installation and then select all the nodes in the cluster. For our exercise, the two cluster nodes were `btc-ppe-wkst1` and `btc-ppe-wkst6`. Click Next.

   Note: If the nodes are not displayed in the cluster node selection, then Oracle CRS is not configured or started on those cluster nodes.
4. For installation type, select Enterprise Edition and click Next.
5. On the Select Database Configuration screen, select Do not create a starter database. We used `dbca` to create a database later. Click Next.
6. Run the following scripts as root user starting from master node when prompted.

```
./$ORACLE_HOME/root.sh
```

7. Click exit to finish the database installation.

   Note: Install Oracle10*g*R2 Patch 3 on both CRS_HOME and ORACLE_HOME using OUI . For more details about the patch installation, refer to the appropriate Oracle patch installation guide.

Configuring a listener for an ASM instance

1. Run `netca` (Oracle Net Configuration Assistant) as `oracle` user to create a listener.
2. Select Cluster Configuration and click Next to select all the nodes displayed to create the listener.
3. Select the default listener name that is displayed.

Creating an ASM instance

Note: The ASM instance created manages the Oracle data on the NFS file system instead of on raw devices. Therefore, the use of ASMLib is not required.

To create the ASM instance, you will need device file or block partitions that Oracle will use to disk groups. Before creating an ASM instance, create device files to act as ASM disks, using the following commands as `oracle` user:

```
dd if=/dev/zero of=/oradata/_file_disk1 bs=1M count=30720
dd if=/dev/zero of=/oradata/_file_disk2 bs=1M count=30720
```

Note: You can create multiple volumes in N series storage for the database, according to your business requirements.

Follow these steps to create an ASM instance:

1. Start Database Configuration Assistant (DBCA) as `oracle` user and select Oracle Real Application Clusters Instance.
2. In the Operations screen, select Configure Automatic Storage Management and click Next.

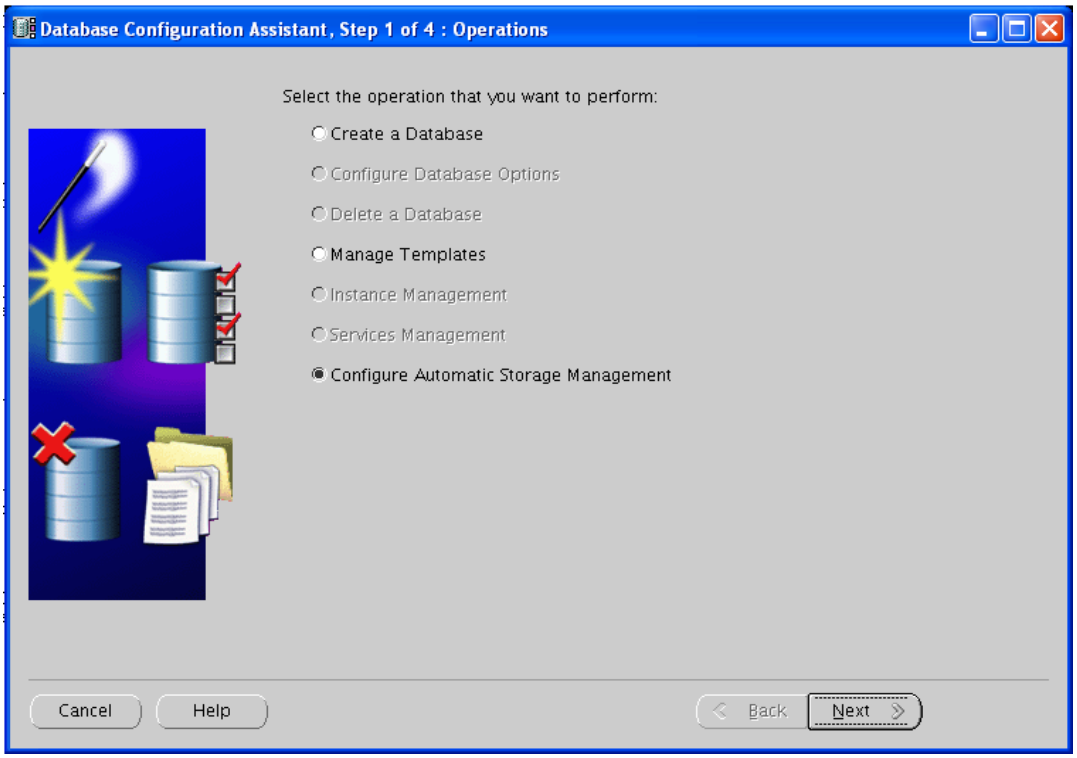

3. In the Node Selection screen, select all the nodes and click Next.

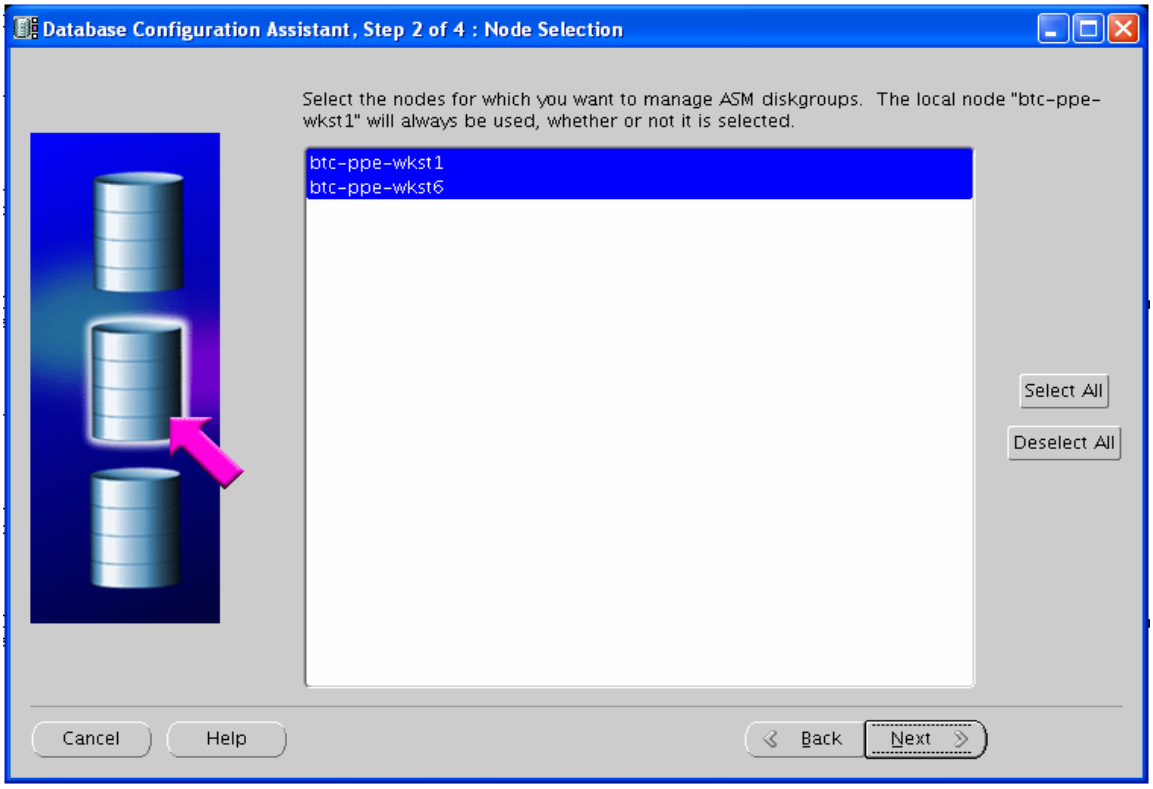

4. In the Create ASM Instance screen, enter the password for the SYS `user` and click Next.

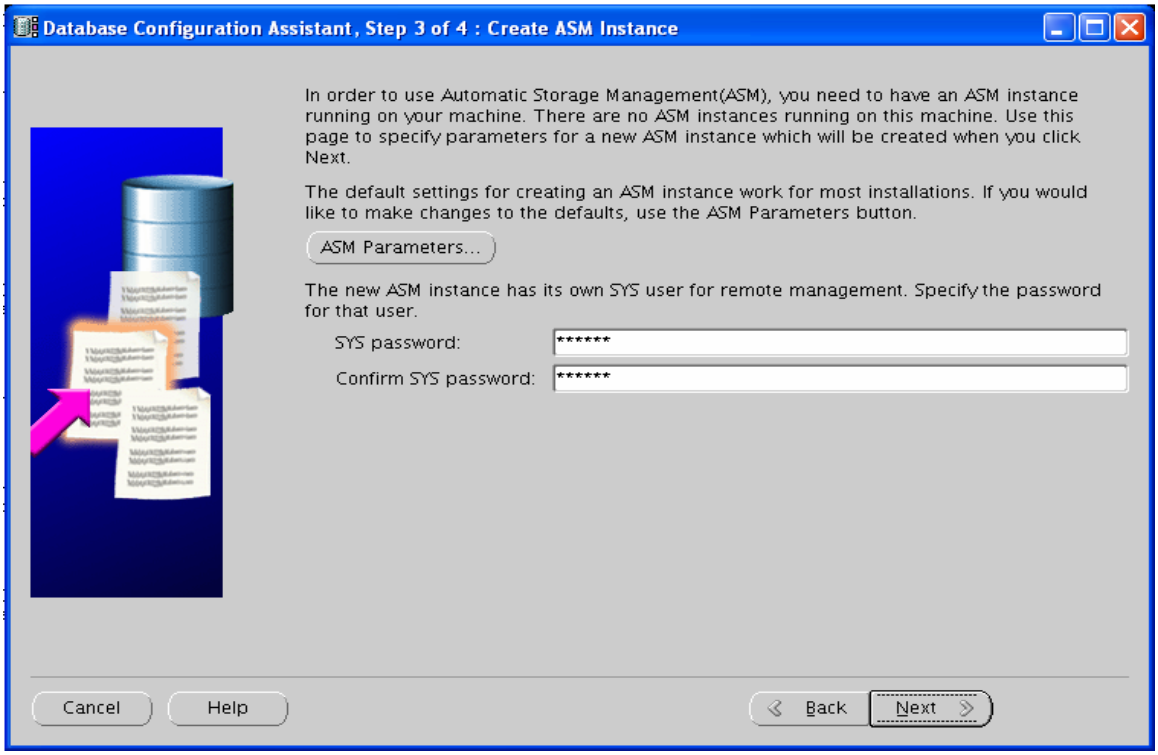

5. DBCA prompts you to start the creation of ASM instance. Click OK to continue.

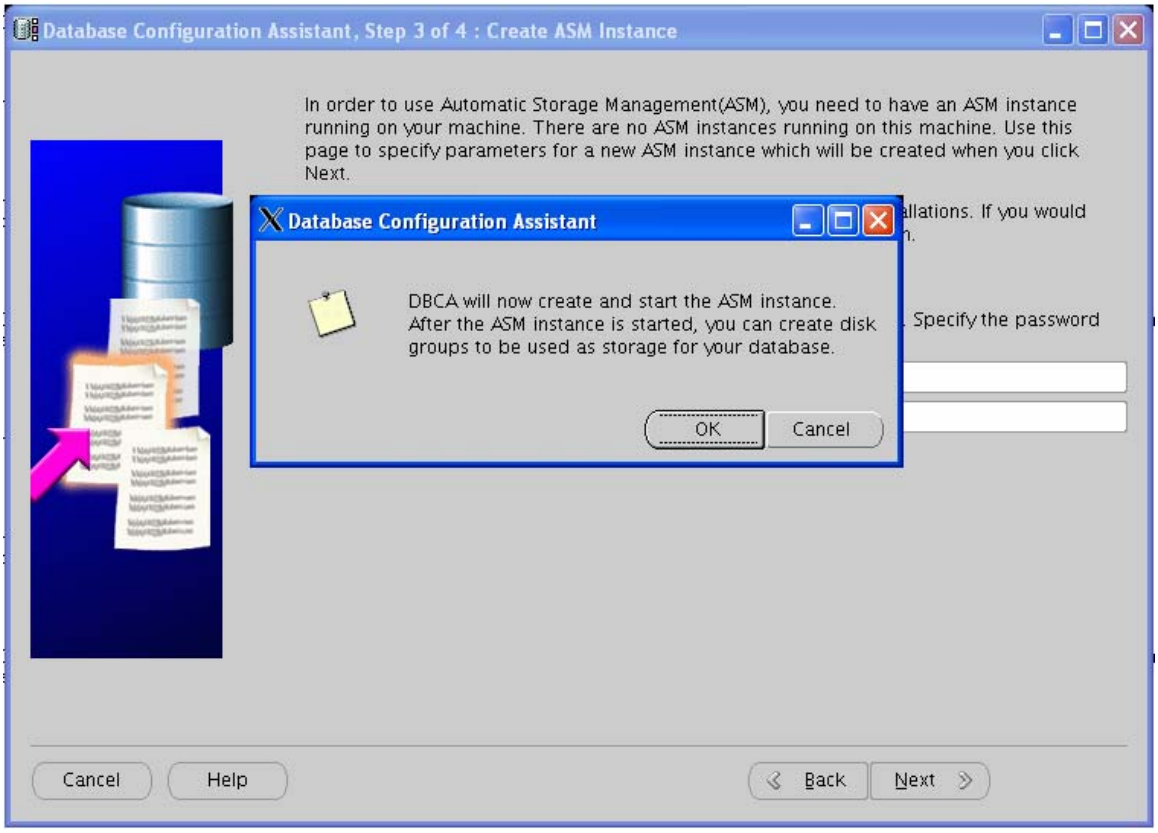

6. After the instance is created, DBCA prompts you to create a disk group. Click Create New to create a disk group. Click Change Disk Discovery Path to change the disk device location. Provide the following disk string path:

```
/oradata/*
```

Note: If you have multiple volumes on N series storage that you have mounted on the servers to store the database, you can specify those mounted directories in the disk string path putting a "-,-" as separator.

Example: `/oradata1/*, /oradata2/*, /oradata3/*, and so on.`

All the disks become visible.

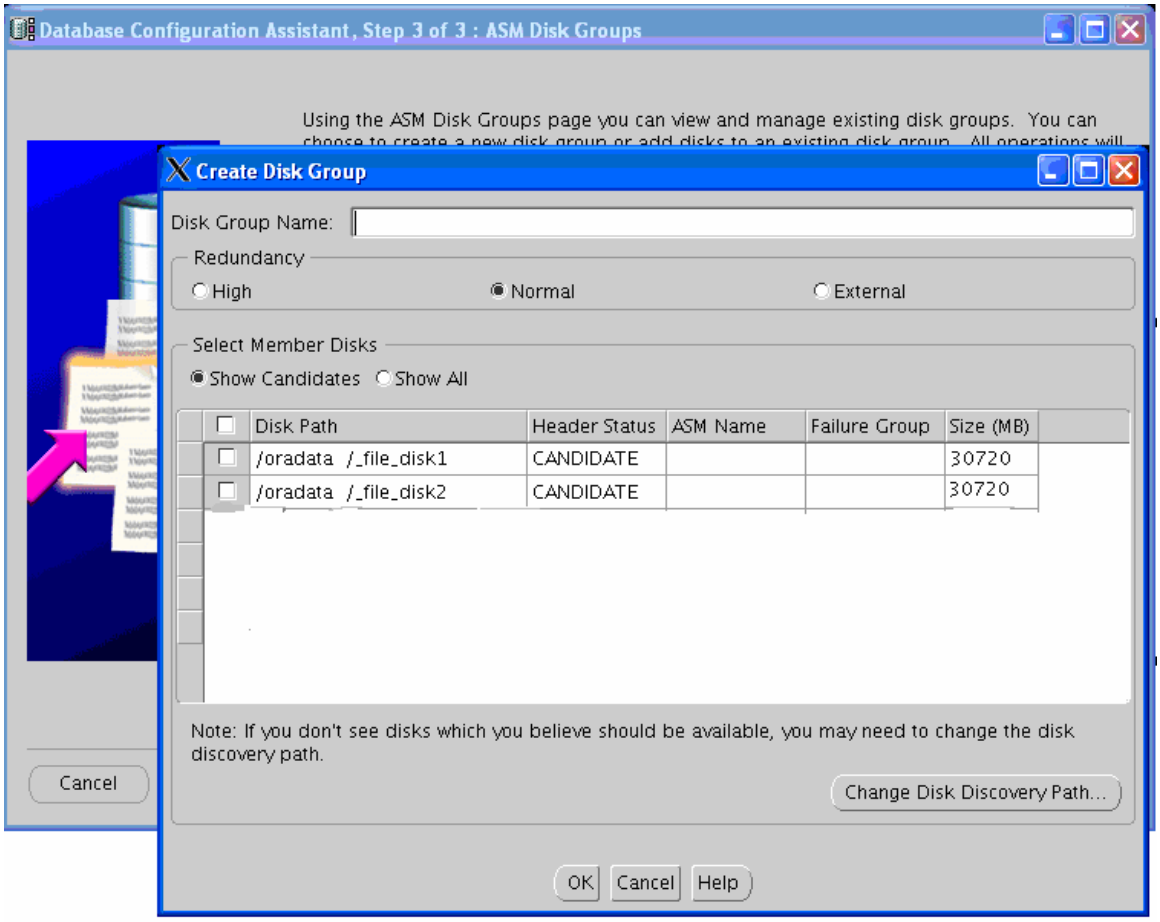

7. Enter a disk group name and select a redundancy option.

   Select the disks to include in the disk group and click OK. DBCA creates the disk group and mounts it on the node from where DBCA is running. Then it mounts the disk group on other nodes of the cluster.

   Note: If the ASM instance is not registered with the listener then DBCA returns a listener connectivity error when trying to mount the disk group on the other node. In that case, you can login in to the respective ASM instance and manually mount the disk group from the SQLPlus prompt using the "`alter diskgroup dsgrp1 mount`"; command.

8. Repeat step 7 to create the required disk groups.

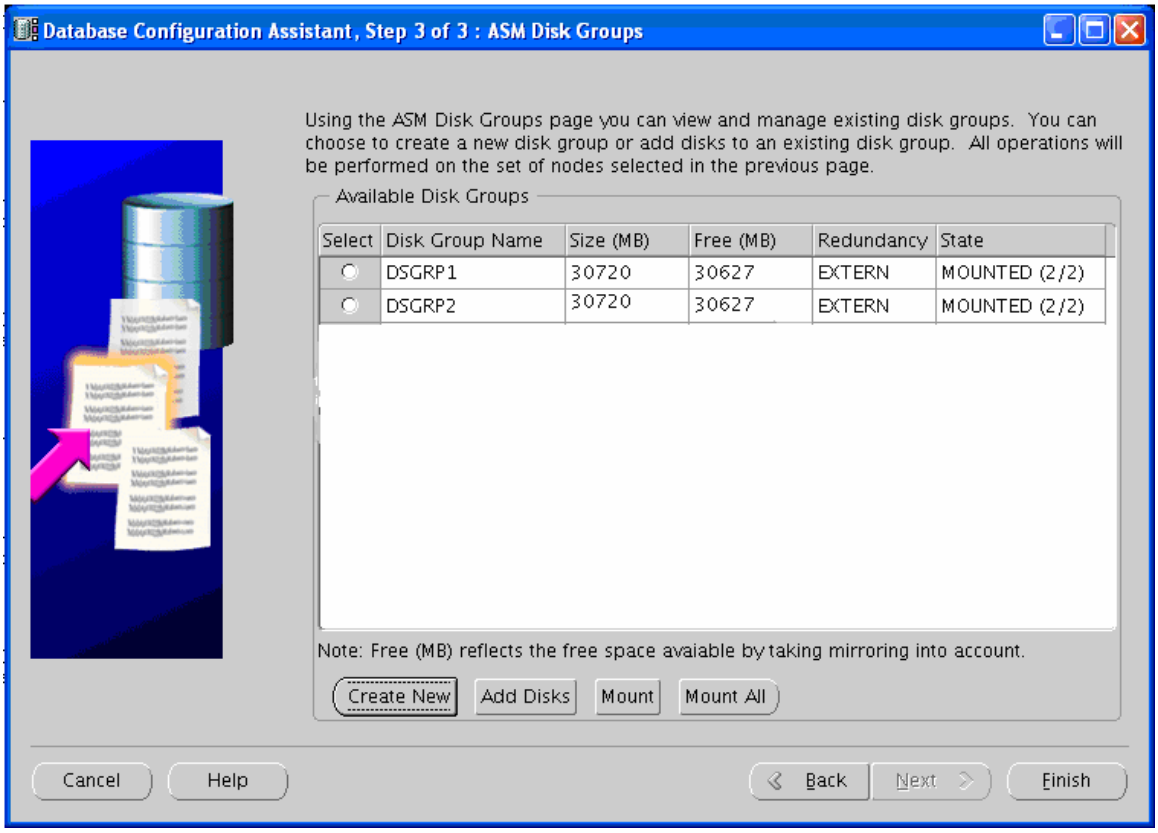

9. When all disk groups have been created, click Finish to exit DBCA.

Note: If DBCA does not mount the disk group on the other ASM instance, logon to the ASM instance as user SYS on the other node and check the status of the disk groups:

```
Select name, state from v$asm_diskgroup;
```

The output of this command will shows the status (mounted or unmounted) of the disk groups. If the status of the disk groups is unmounted, then execute the following command to mount the disk groups:

```
Alter disk group <diskgroup_name> mount;
```

Creating a database using ASM disk groups

ASM disk groups can be used to store all types of database files, including binary and test files. We chose ASM disk group to store data files, control files, online redo logs, flash recovery area, and parameter files. We kept archive log and trace file destinations on file system.

Follow these steps to create an Oracle Database using ASM disk groups:

1. Start DBCA and select Create a Database.

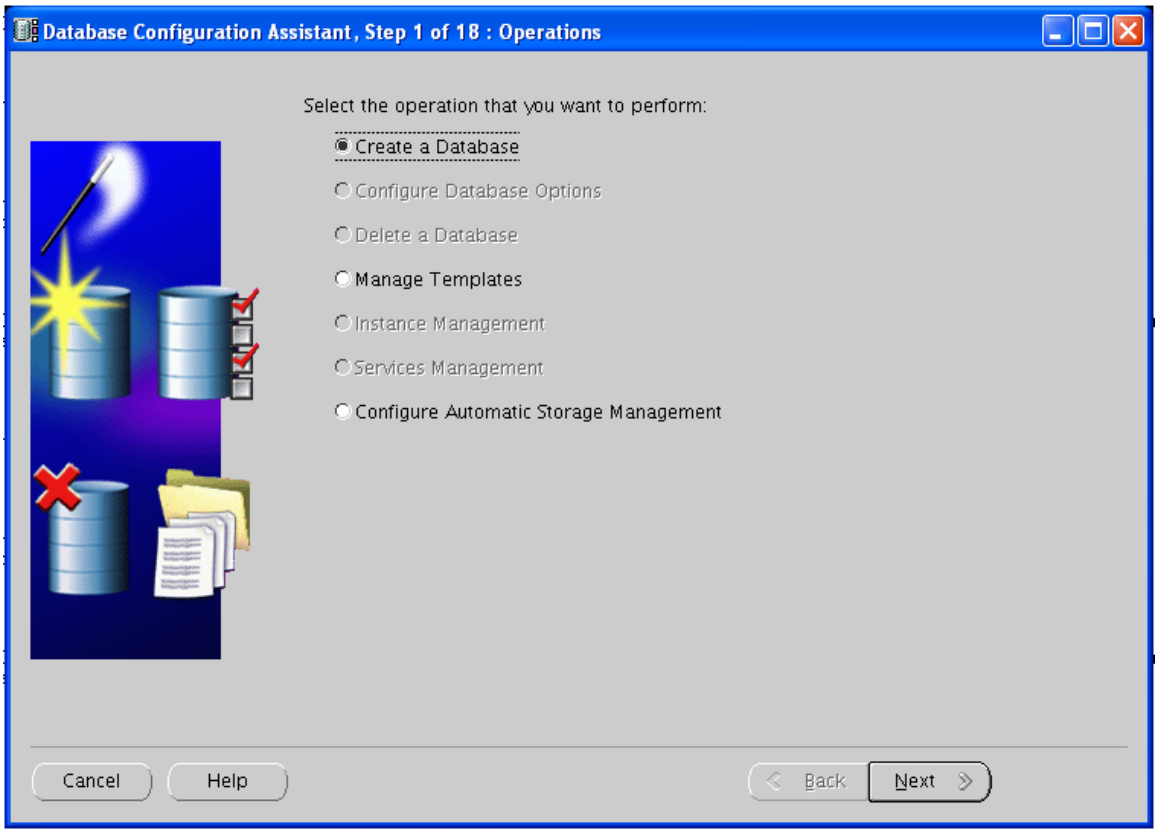

2. In the Node Selection screen, select all the nodes and click Next.

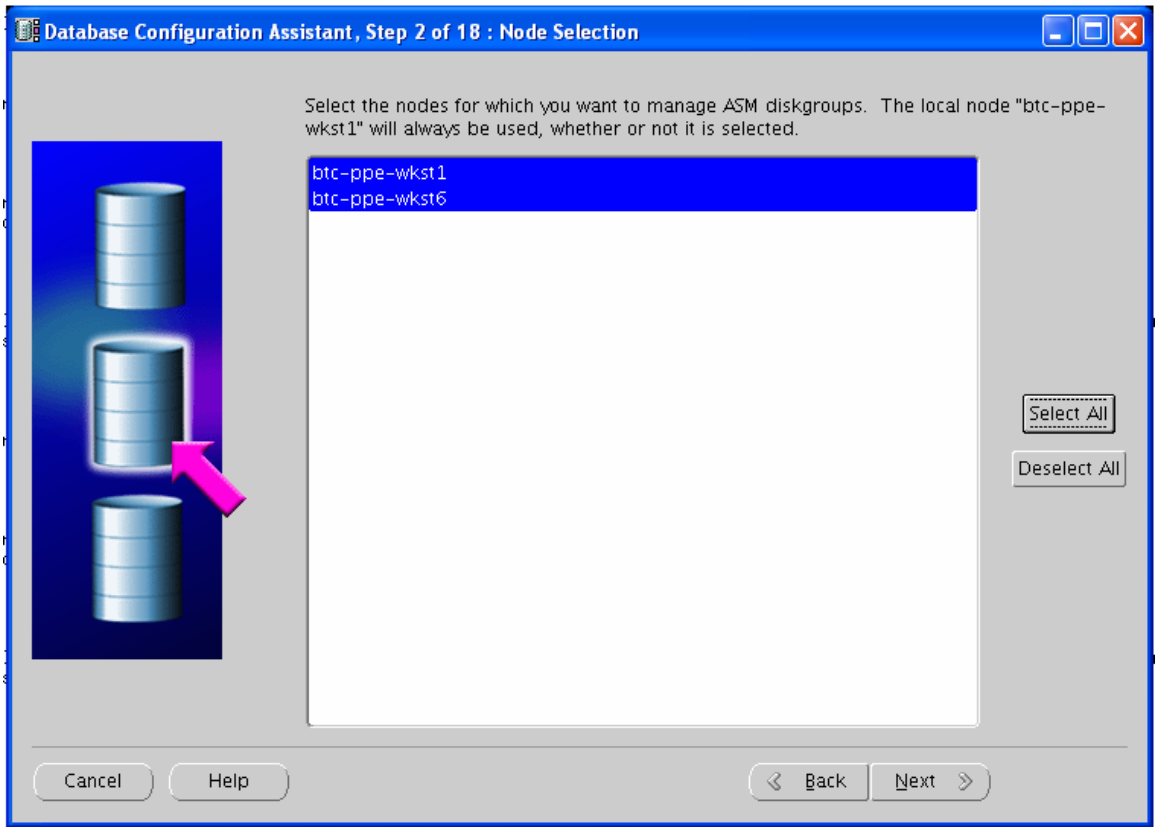

3. In the Database Templates screen, select General Purpose and click Next.
4. In the Database Identification screen, enter a Global Database Name and SID prefix. Click Next.
5. In the Management options, select Configure the Database with Enterprise Manager to configure the DB console and its repository. Click Next.
6. In the Database Credentials screen, enter a password for SYS and SYSTEM and click Next.

Note: If you have already created a RAC listener, DBCA prompts you to register the database with the listener in the Network Configuration screen. If you want to create separate listener for database, then skip this option.

7. In the Network Configuration screen, select Register this database with all listeners and click Next.

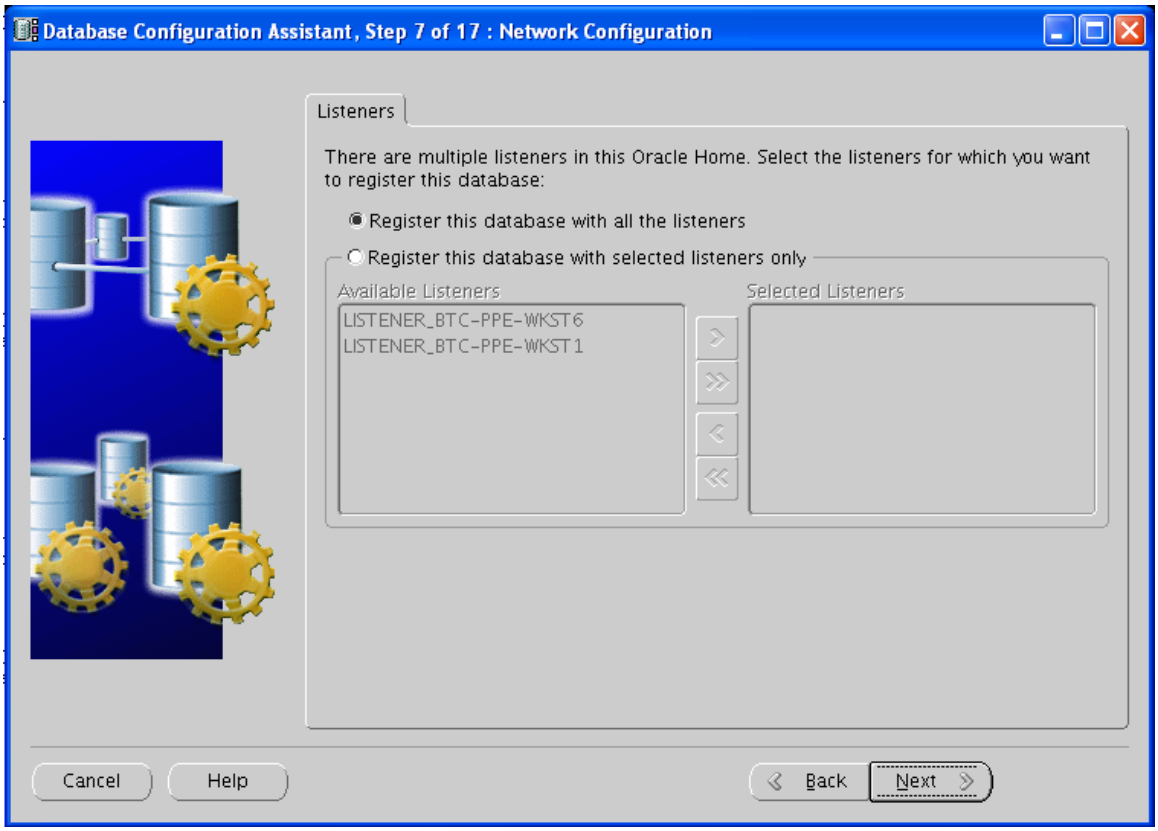

8. In the Storage Options screen, select Automatic Storage Management (ASM) and click Next.

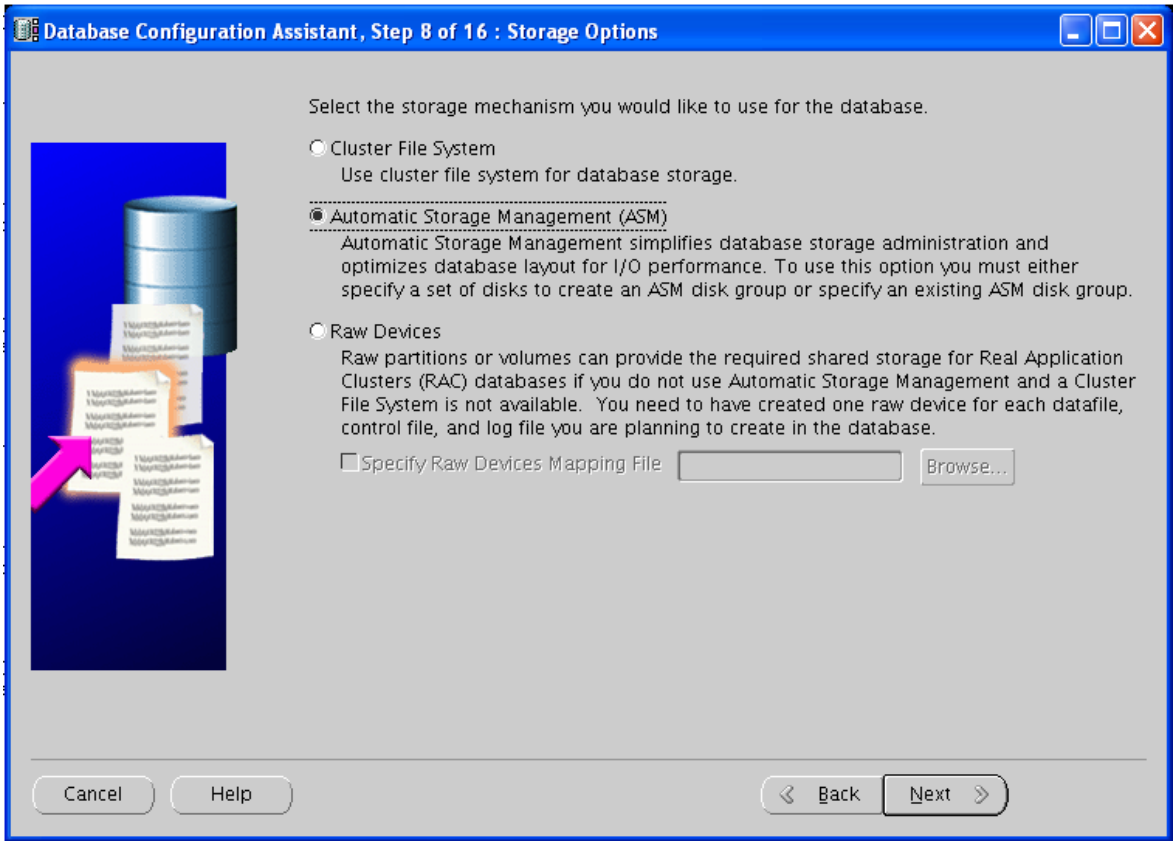

9. In the ASM Disk Groups screen, select the disk groups where you want to create or store the database files. In this example, disk group DSGRP1 is selected to store system files, control files and online redo log files. Click Next.

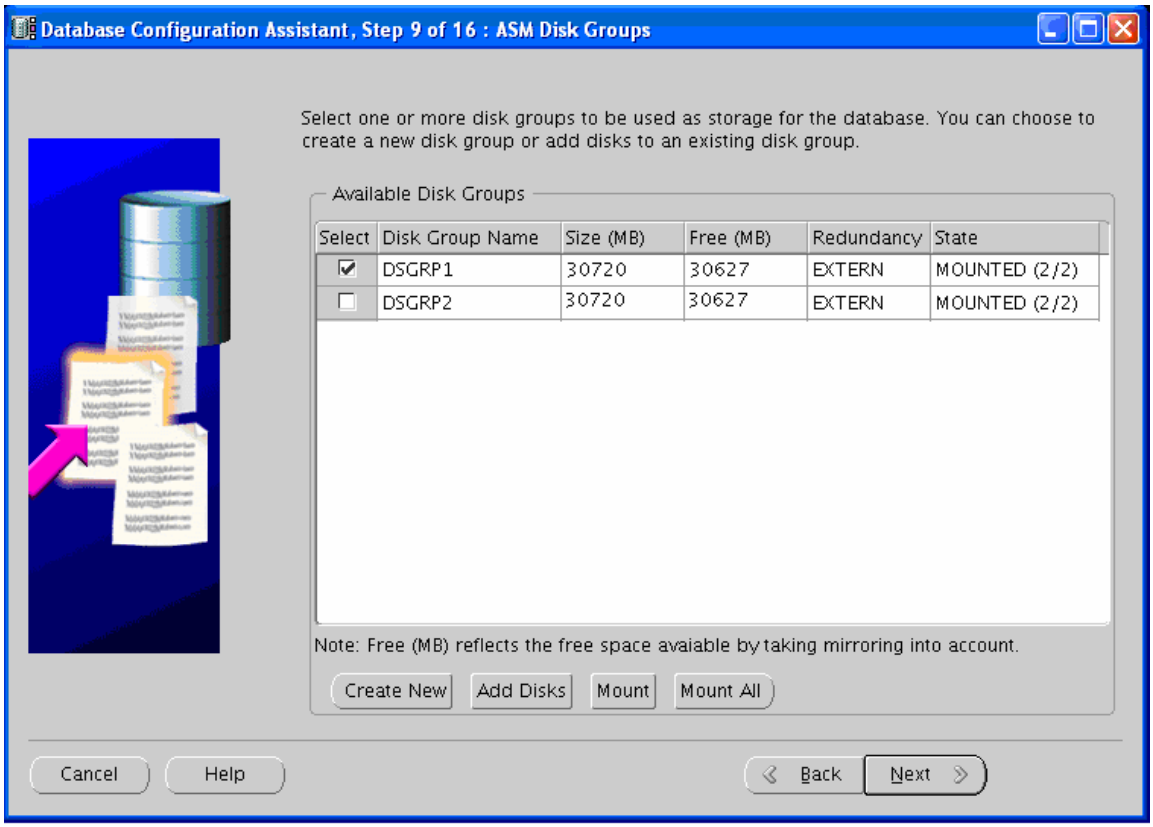

10. In the Database File Locations screen, select an option. In this test process, we selected Use Oracle-Managed files. Enter the disk group in which DBCA will create all the database files. Click Next.

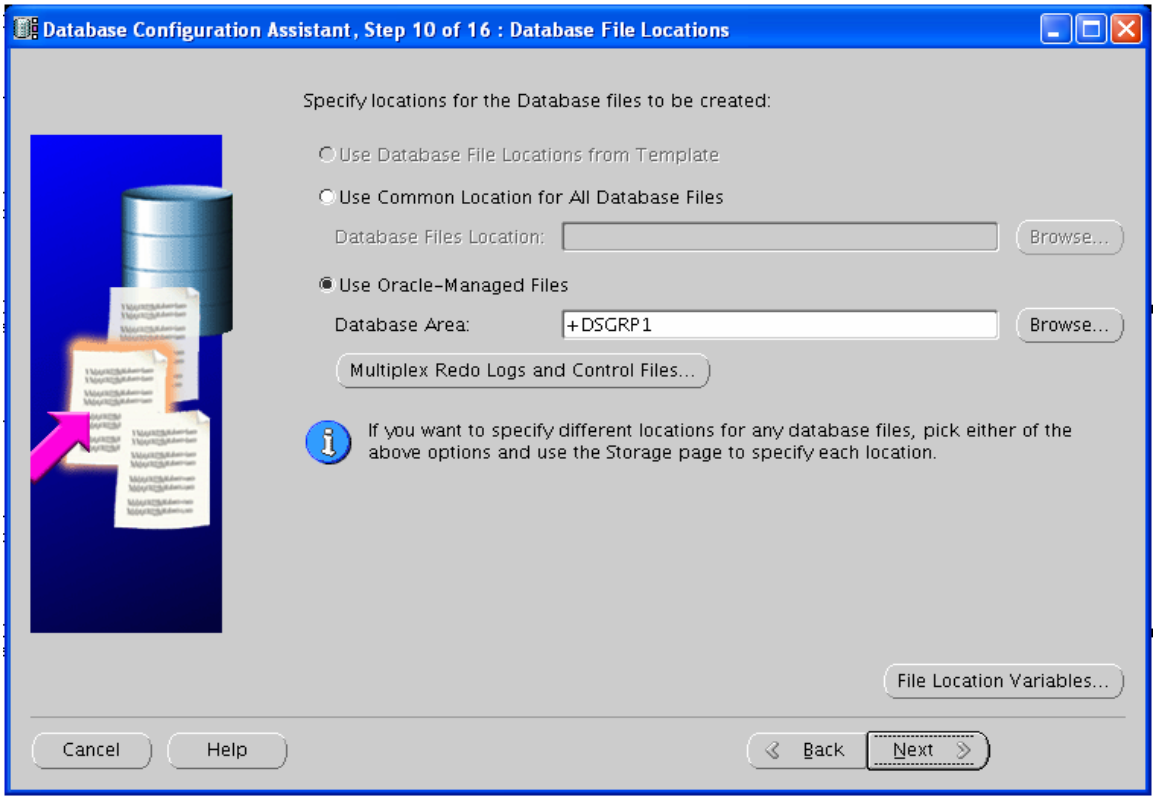

11. In the Recovery Configuration screen, select Specify Flash Recovery Area. Enter disk group DSGRP2 as the flash recovery area destination and enter 2048 as the size. Click Next.

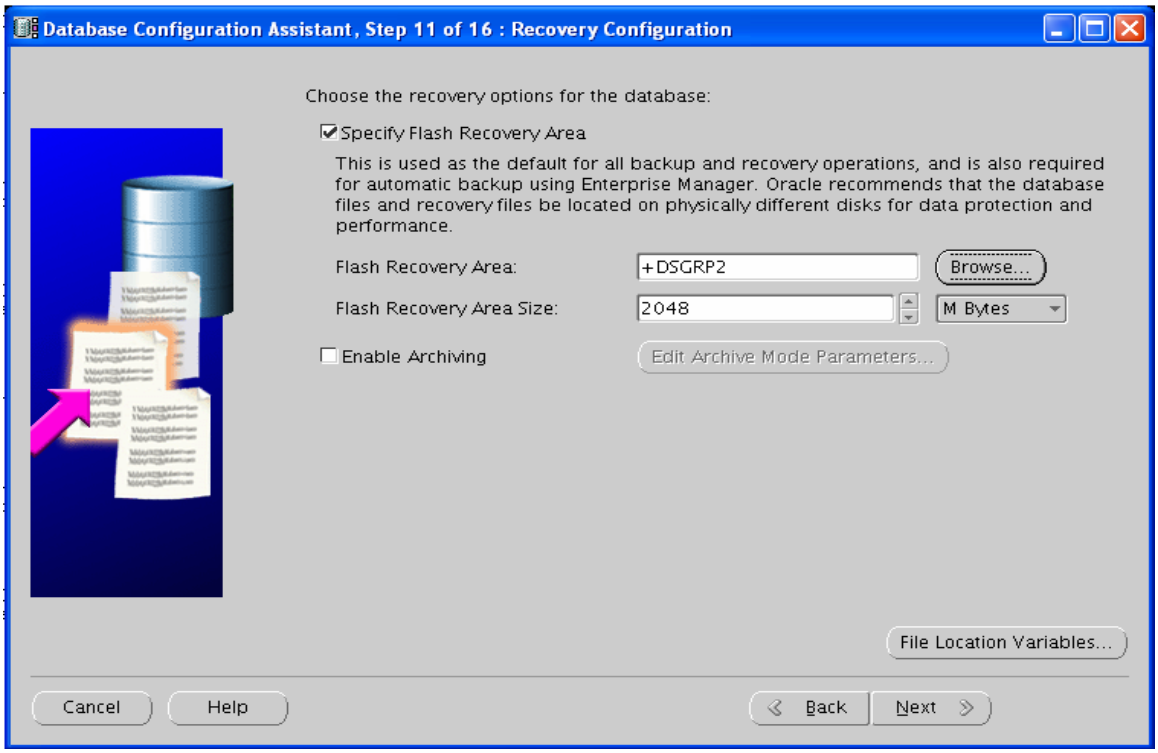

12. In the next screen, accept the default options. In the Creation Options screen, click Finish to start the database creation. After successful creation of the database, DBCA shows the completion message.

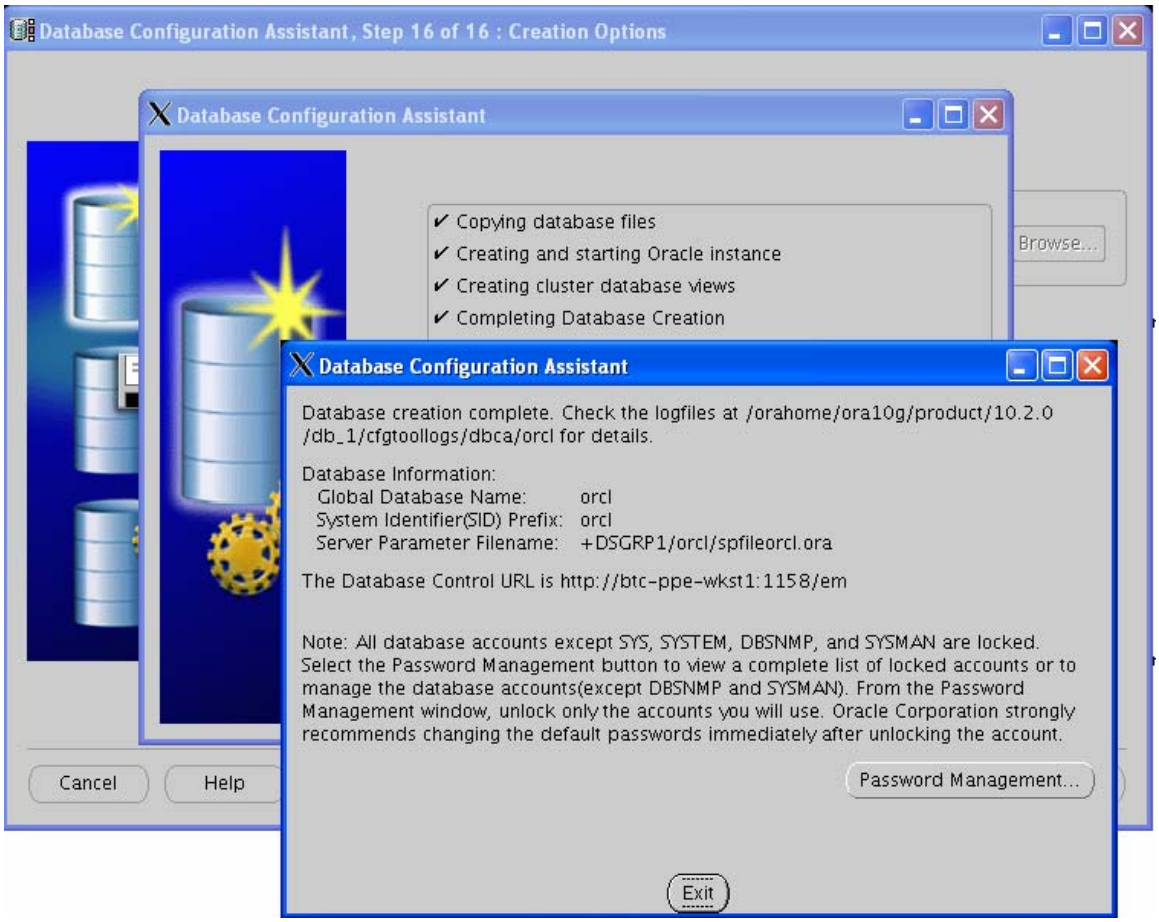

Database settings for I/O performance

Logon as `oracle` user, export the ORACLE_SID and connect to the database instance as the SYS user. Check the parameter settings in `spfile` as follows:

```
[root@btc-srv1 ~]# su – oracle
[oracle@btc-srv1 ~]$ export ORACLE_SID=orcl1
[oracle@btc-srv1 ~]$ sqlplus "/ as sysdba"
```

Check for the `disk_asynch_io` and `filesystemio_options` parameter settings in `spfile`:

```
SQL > show parameter disk_asynch_io

NAME                                 TYPE        VALUE
------------------------------------ ----------- ------------------------------
disk_asynch_io                       boolean     FALSE
```

If `disk_asynch_io` is FALSE , change this parameter value to TRUE with the following command:

```
SQL > alter system set disk_asynch_io=TRUE scope=spfile;

SQL > show parameter filesystemio_options

NAME                                 TYPE        VALUE
------------------------------------ ----------- ------------------------------
filesystemio_options                 string      NONE
```

If the `filesystemio_options` setting is NONE, change this parameter value to `setall` with the following command:

```
SQL > alter system set filesystemio_options = setall scope=spfile;
```

Bring down all the database instances and ASM instance on all the cluster nodes.

Relink the oracle binaries in all cluster nodes as follows:

```
[oracle@btc-srv1 ~]$ cd $ORACLE_HOME/rdbms/lib

[oracle@btc-srv1 ~]$ make PL_ORALIBS=-laio -f ins_rdbms.mk async_on
```

Repeat this command in all cluster nodes. It should relink the Oracle binaries without any error.

Start up the ASM instance and database instance on all cluster nodes.

To confirm whether both direct I/O and asynch I/O are enabled, issue the following commands as `oracle` user:

```
[oracle@btc-srv1 ~]$ cat  /proc/slabinfo | grep kio

kioctx      53   75  256  15  1 : tunables  120  60   8 : slabdata    5     5     0
kiocb       47   64  128  31  1 : tunables  120  60   8 : slabdata    3     5     0
```

`Non-zero` values indicate that both direct I/O and asynch I/O are enabled.

Appendix: Sample .bash_Profile file for the oracle user

```
export ORACLE_BASE=/orahome/ora10g;
export ORACLE_PRODUCT=$ORACLE_BASE/product;
export ORACLE_HOME=$ORACLE_PRODUCT/10.2.0/db_1;
export ORACLE_CRS=$ORACLE_PRODUCT/10.2.0/crs_1;
export ORACLE_SID=orcl;
export
LD_LIBRARY_PATH=$ORACLE_HOME/lib:$ORACLE_CRS/lib:$ORACLE_HOME/lib32:$LD_LIBRARY_PATH;
export LIBPATH=$ORACLE_HOME/lib:$ORACLE_CRS/lib:$ORACLE_HOME/lib32:$LIBPATH
export PATH=$PATH:$ORACLE_HOME/bin:$ORACLE_HOME:$ORACLE_CRS/bin:/usr/java14/bin;
```

Trademarks and special notices

Data ONTAP, FlexVol, Network Appliance, the Network Appliance logo, SnapRestore are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.