Installing Oracle9*i* RAC on IBM @server xSeries running SuSE Linux Enterprise Server 7

Version 1.1

February 14, 2002

By:

Tim Spencer
Stephen Poon

IBM / Oracle International Competency Center
San Mateo, California

IBM

Contents

INTRODUCTION

The objective of this document is to provide help to install Oracle9*i* Real Application Clusters (RAC) on IBM eServer x440 with FAStT500 running SuSE Linux Enterprise Server 7 (SLES7) through a description of the different steps necessary to install Oracle9*i* RAC. In this document, we take into account the specifics of the IBM xSeries hardware, and provide you the necessary information to successfully install Oracle9*i* RAC on IBM xSeries.

The information contained in this paper resulted from:
- Oracle and IBM documentation
- Installation runs of Oracle9*i* RAC
- Certification and other work done by the EMEA Oracle/IBM Joint Solutions Center
- Contributions from Oracle and IBM specialists.

Please also refer to Oracle documentations for more information. (http://docs.us.oracle.com).

- Oracle9*i* Installation Guide Release 9.2.0.1 for UNIX Systems
- Oracle9*i* Administrator's Reference 9.2.0.1 for UNIX Systems
- Oracle9*i* Release Notes Release 2 (9.2.0.1.0) for Linux Intel
- Oracle9*i* Online Generic Documentation CD-ROM Installation and Usage Notes
- Oracle9*i* Real Application Clusters Installation and Configuration
- Oracle Enterprise Manager Configuration Guide

We used the xSeries x440 and FAStT500 Storage Controller to perform this installation. While illustrations may be specific to these products, the methods described generally apply to other xSeries servers and FAStT Storage Controllers. Where appropriate, we will point out areas that are specific only to the x440.

Once the hardware is configured and the environment ready, the installation of the Oracle product is virtually identical on all Linux platforms. We have based this installation procedure on the document “Step-By-Step Installation of RAC on Linux” from Oracle Metalink and “Installing Oracle9*i* on IBM eServer x360 with FAStT700 running Red Hat AS 2.1” by the EMEA Oracle/IBM Joint Solutions Center. We have illustrated it with the screen shots from our lab.

Your feedback is important for us. We want our technical papers to be as helpful as possible. Please send your comments about this document to the IBM/Oracle International Competency Center (IOICC). The IOICC’s email address is:

ibmoracl@us.ibm.com

Special thanks to the following for their assistance and contribution to this paper:

EMEA Oracle/IBM Joint Solutions Center, Montpellier, France

PLANNING THE INSTALLATION

Obtain the CDs for SuSE Linux Enterprise Server 7 (SLES7), SLES7 Patch CD2 and CD3.

Search and download the BIOS/drivers from http://www.pc.ibm.com:

- IBM eServer xSeries 440 - Flash BIOS Update version 1.06
- IBM eServer xSeries 440 - Diagnostic Flash Update version 1.03
- Servers - Fibre Channel Solutions: IBM FAStT Host Adapter BIOS version 1.76
- Servers and Fibre Channel Solutions - IBM FAStT Host Adapter failover device driver version 6.0 for Linux
- Fibre Channel Solutions - IBM FAStT Storage Manager version 8.21 Fibre Channel Controller firmware
- Fibre Channel Solutions - IBM FAStT Storage Manager version 8.21 for Red Hat Linux
- Fibre Channel Solutions - IBM FAStT MSJ Diagnostic and Configuration Utility version 2.0 release 33 for Linux

Search and download the Redbooks/Redpapers from http://www.redbooks.ibm.com:

- Implementing IBM eServer xSeries SANs, REDP0416
- Implementing Oracle9i RAC with Linux on IBM eServer xSeries servers, REDP0410

Obtain Oracle9*i* R2 Enterprise Edition CDs (Oracle9*i* can also be downloaded from http://technet.oracle.com).

HARDWARE CABLING AND FIRMWARE

HARDWARE CONFIGURATION

Figure 1 below shows our 2 node RAC cluster configuration, with xSeries 440's, IBM 2109-S16 SAN switch, FAStT500 Storage Controller and EXP500 Expansion Unit. Note that we had one SAN switch in our configuration. For higher availability, a second switch may be added. Refer to the "Implementing IBM eServer xSeries SANs" Redbook for more information on SAN configurations.

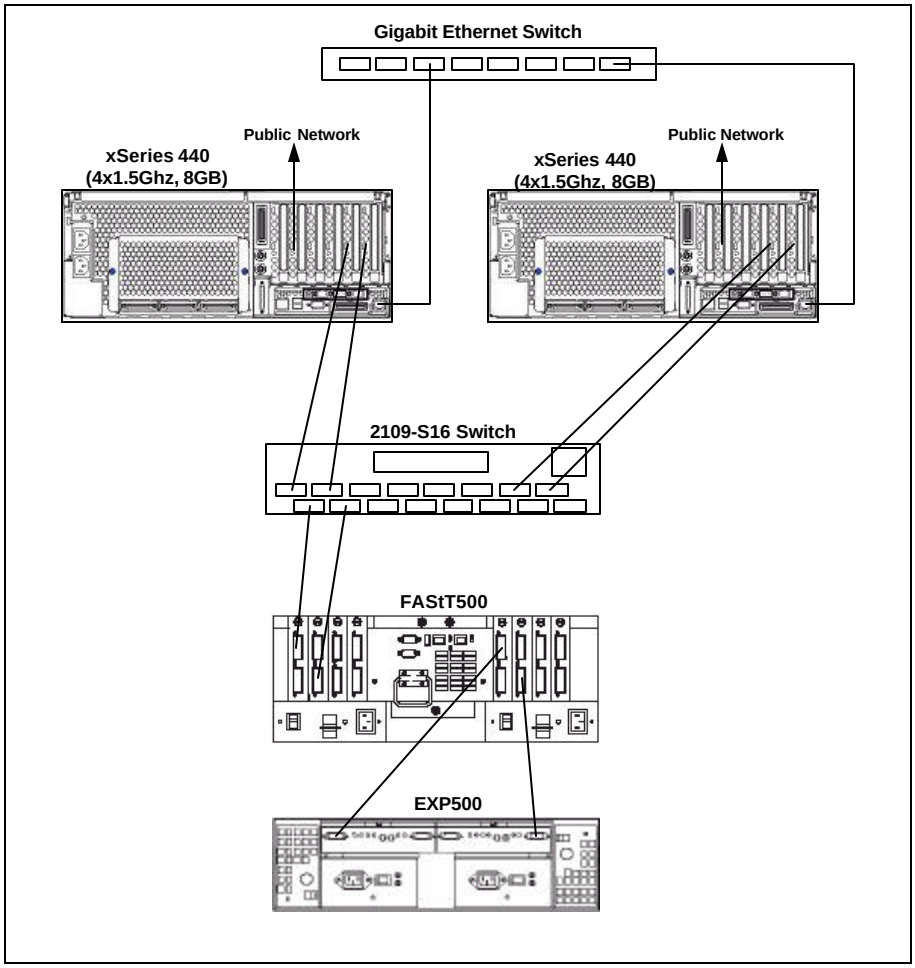

Figure 1. Oracle9i RAC configuration

FIRMWARE

The following are the minimum levels of firmware/BIOS for the different hardware parts (check the latest supported level on http://www.pc.ibm.com):

- XSeries 440 BIOS: 1.06
- ServeRaid BIOS: 4.84
- FAStT HBA BIOS: 1.76
- FAStT500 firmware: 8.21

Check the respective readme files from the downloaded files for the upgrade procedure. See the "PLANNING THE INSTALLATION" section for the downloads.

Also check the Implementing IBM eServer xSeries SANs Redbook for more information on the FAStT firmware upgrade.

OPERATING SYSTEM SETUP

This section will guide you through the steps for the post-installation of SLES7 to take in account the specifics of the IBM xSeries cluster.

After you have installed SLES7 on all the cluster nodes, follow these steps for each node in the cluster.

The oracle account (user ID) has been included in SLES7. In later sections, this account will be used to install, configure and run Oracle9i. You should change the password for the oracle user ID to secure your system. For the configuration steps described later, log in to X-Windows using this account. When necessary, we indicate the relevant steps for which you need to switch to the root user ID.

OPERATING SYSTEM PATCHES

Install SLES7 Patch CD2 and CD3.

1. Insert Patch CD2 into the CD drive. Open a terminal window. From oracle user's shell prompt, enter `xhost +` to allow connections to X-Windows.

2. Switch to root user, and check that the DISPLAY variable is set to "localhost:0.0". If not, set it, then enter `yast2`.

3. In the YaST2 Control Center window, click **Software** on the left hand panel, then **Patch CD Update** in the main panel as shown in Figure 2.

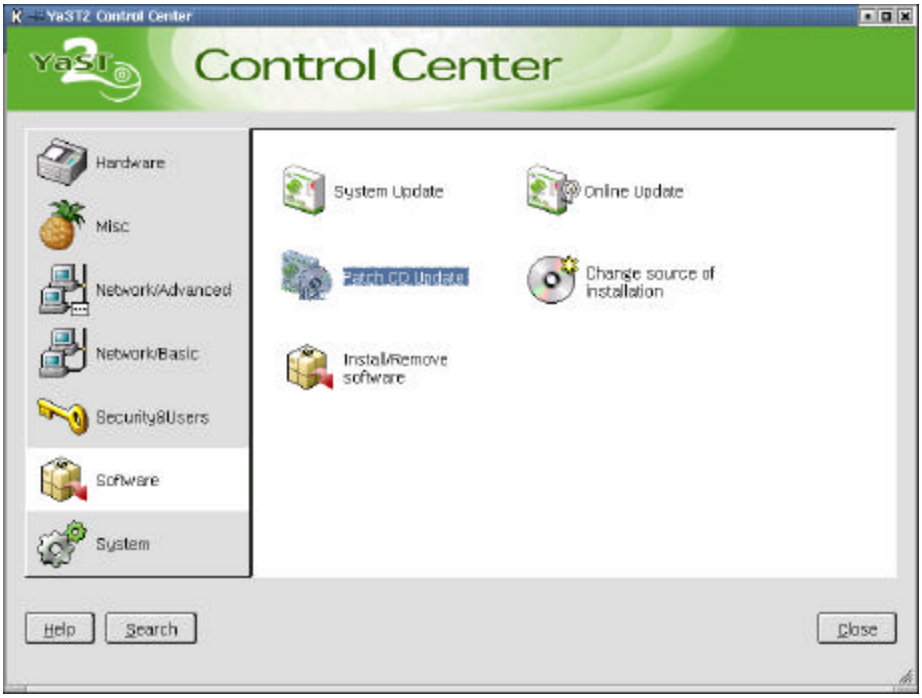

Figure 2. YaST2 Patch CD Update

Patches relevant to the installation will be automatically selected. With Patch CD2, you will install a YaST2 patch by itself first. Once this is done, exit and reload YaST2 in order to apply other patches.

4. Patch CD2 installs kernel 2.4.18, which contains Summit kernel patches for the x360 and x440. Be sure to select the Linux kernel source, patch-669, from CD2. Highlight the patch as shown in Figure 3 and

click **Apply** twice. You will see a "G", then the "X" appear in the status field, as shown in the figure. Click **Next** to install patches.

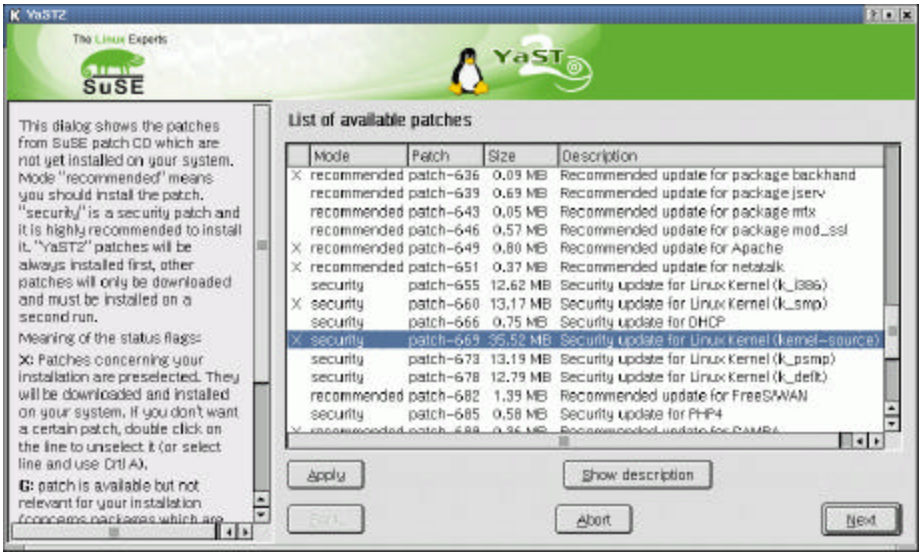

Figure 3. Install Linux kernel source

5. When completed, remove Patch CD2 from the CD drive and insert Patch CD3. Repeat the same procedure to install Patch-CD3. Exit yast2 when completed.

IMPORTANT: Be sure to execute mk_rdinit and lilo after installing patches from the Patch CDs. Do this from the root user's shell prompt.

NETWORK

In our example, the on-board Gigabit Ethernet device was used for the private LAN and an IBM 10/100 Ethernet Adapter 2 was used for the public LAN. The interfaces, IP addresses and hostnames are shown in Table 1.

Node	Public LAN			Private LAN		
	Hostname	Interface	IP address	Hostname	Interface	IP address
1	xrac1	eth1	192.168.10.10	xrac1i	eth0	1.1.1.10
2	xrac2	eth1	192.168.10.20	xrac2i	eth0	1.1.1.20

Table 1. Network settings

During our installation, we encountered problems configuring the second network adapter. To circumvent this problem, we modified the configuration text files to define the interfaces and used YaST2 to activate and define other network parameters. These are the steps we used:

1. Open a terminal window and switch to root user

2. Edit /etc/modules.conf and modify the alias statements for eth0 and eth1 as shown in Example 1 below for node 1 (xrac1):

```
#
# YaST2: Network card
#
alias eth0 bcm5700
alias eth1 e100
alias tr0 off
alias scsi_hostadapter off
alias fb0 off
```

Example 1. /etc/modules.conf

3. Edit /etc/rc.config. Modify the statements shown in bold highlighting in Example 2 for node 1.

```
# Number of network cards: "_0" for one, “_0 _1 _2 _3” for four cards
#
NETCONFIG="_0 _1"

#
# This variable contains all indices of active PCMCIA network devices
#
NETCONFIG_PCMCIA=""

#
# IP Adresses
#
IPADDR_0="1.1.1.10"
IPADDR_1="192.168.10.10"
IPADDR_2=""
IPADDR_3=""
#
# Network device names (e.g. "eth0")
#
NETDEV_0="eth0"
NETDEV_1="eth1"
NETDEV_2=""
NETDEV_3=""

#
# Parameters for ifconfig, simply enter "bootp" or "dhcpclient" to use the
# respective service for configuration.
# Sample entry for ethernet:
# IFCONFIG_0="192.168.81.38 broadcast 192.168.81.63 netmask 255.255.255.224"
#
IFCONFIG_0="1.1.1.10 broadcast 1.1.1.255 netmask 255.255.255.0"
IFCONFIG_1="192.168.10.10 broadcast 192.168.10.255 netmask 255.255.255.0"
IFCONFIG_2=""
IFCONFIG_3=""
```

Example 2. /etc/rc.config

4. From the root user’s shell prompt, enter yast2. Select **Network/Basic** in the left panel and then **Network card configuration** in the main panel. The Network basic configuration window, shown in Figure 4, is displayed:

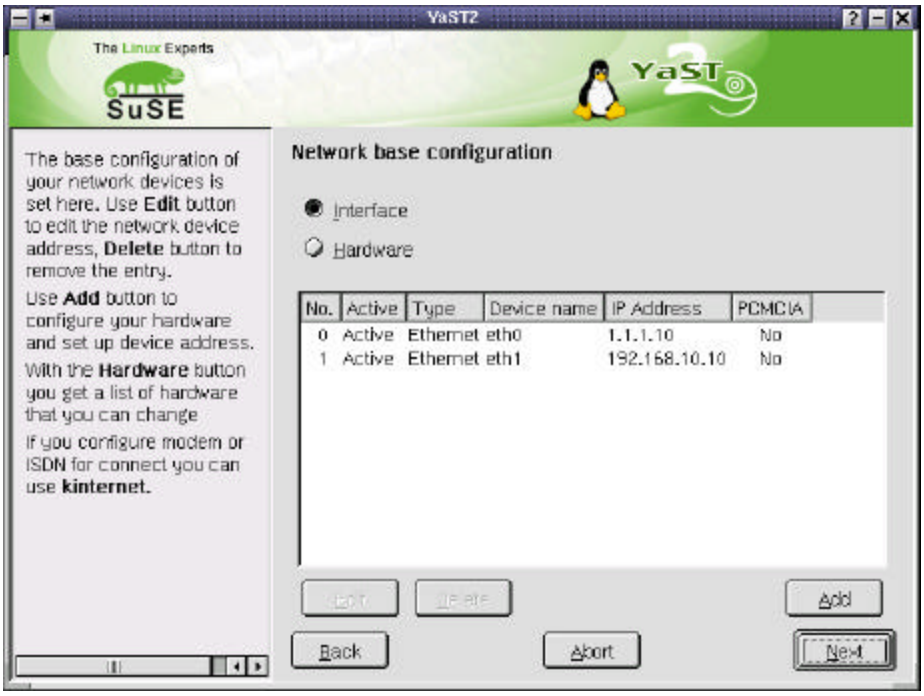

Figure 4. Network basic configuration window

5. Verify that the network devices in your system are shown correctly. If so, click **Next** to continue. Otherwise, go back and make sure that the previous steps have been done correctly.

6. Enable network services through the inetd server. From YaST2 Control Center window (Figure 5), select **Network/Basic**, then **Stop/Start services (inetd)**.

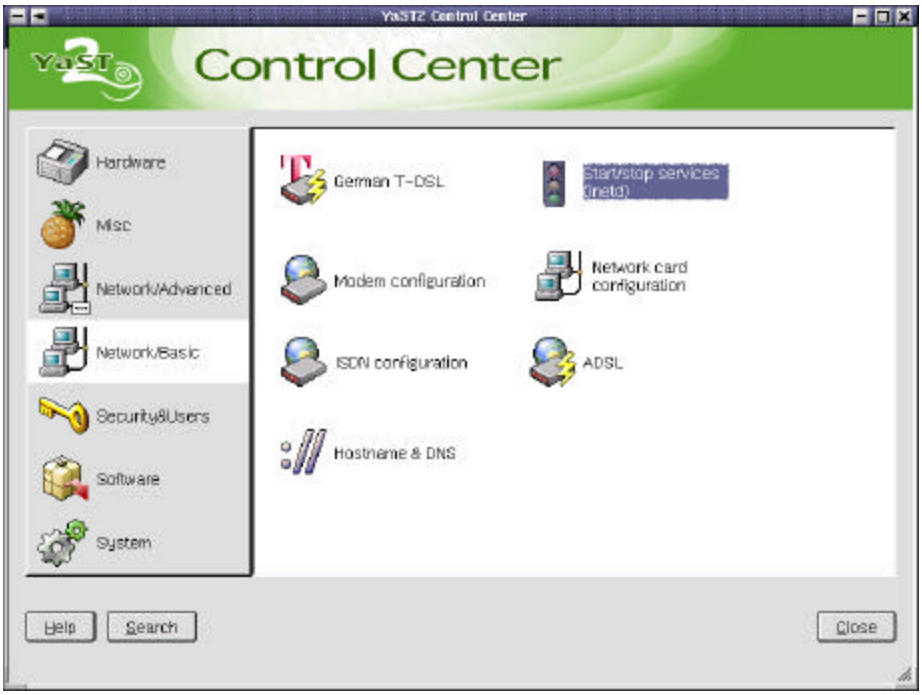

Figure 5. Select Network/Basic window

7. From the Network services windows, select **On with custom configuration** and click **Next**.

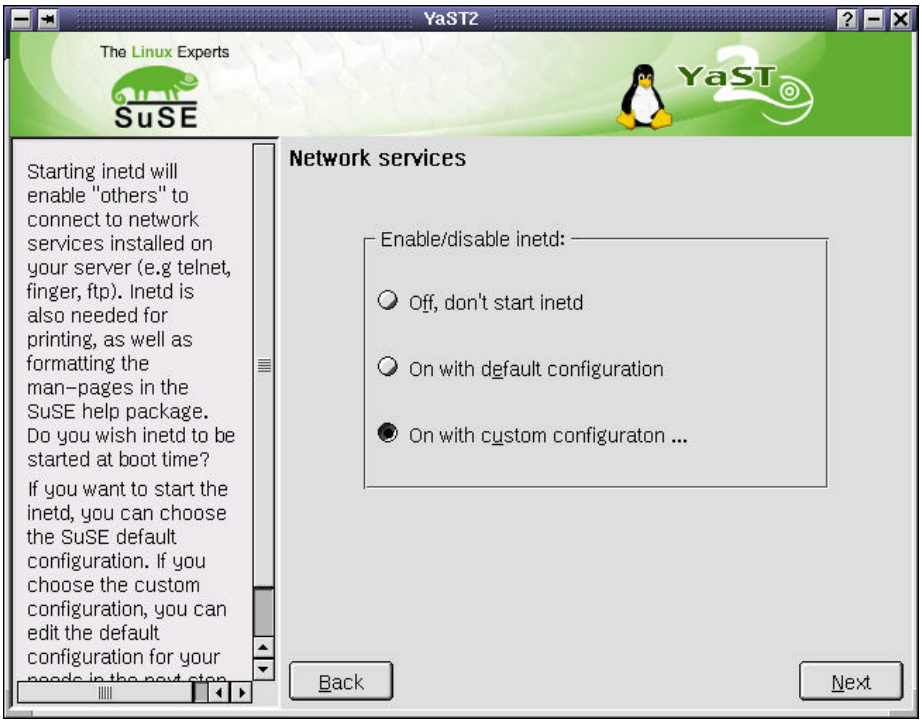

Figure 6. Network services window

8. The **Enable/disable network services** window is displayed as shown in Figure 7. Verify that the services shown in Table 2 are activated:

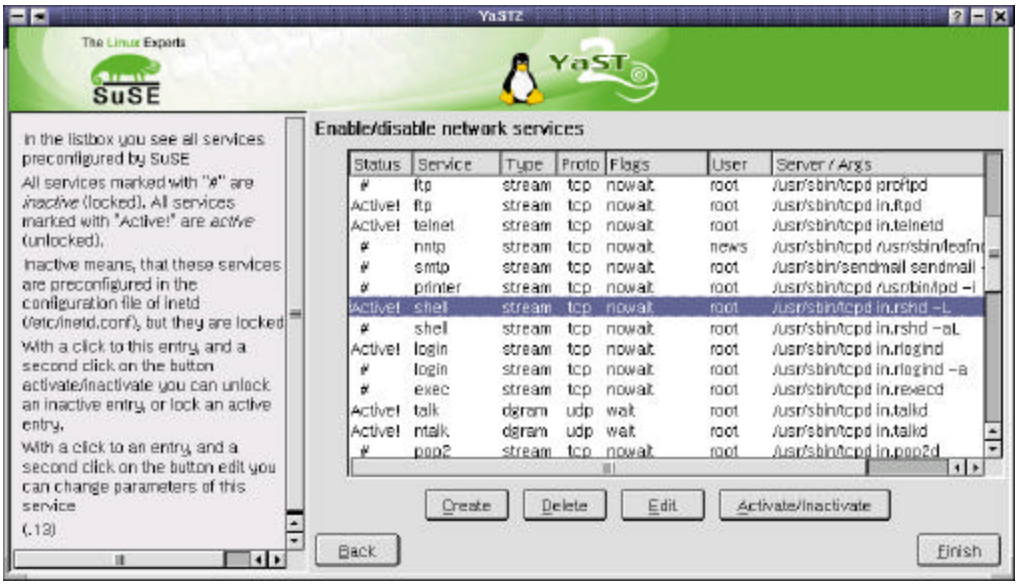

Figure 7. Enable/disable network services window

Service	Server / Args
ftp	/usr/sbin/tcpd in.ftpd
telnet	/usr/sbin/tcpd in.telnetd
shell	/usr/sbin/tcpd in.rshd -L
login	/usr/sbin/tcpd in.rlogind

Table 2. Network services

If **Active!** does not appear in the Status column for each of these services, select the service and click **Activate/inactivate**. Click **Finish** when done, then exit YaST2.

9. Add the host names to /etc/hosts.

```
192.168.10.10    xrac1.sanmateo.ibm.com      xrac1
192.168.10.20    xrac2.sanmateo.ibm.com      xrac2
1.1.1.10         xrac1i.local         xrac1i
1.1.1.20         xrac2i.local         xrac2i
```

Example 3. Entries for /etc/hosts

10. Configure host equivalence by adding the public and private hosts name to the /etc/hosts.equiv file, as shown in Example 4.

```
xrac1
xrac2
xrac1i
xrac2i
```

Example 4. /etc/hosts.equiv

Test host equivalence by executing rsh, rcp and rlogin commands on both xrac1 and xrac2. Example 5 shows the output from these commands on xrac1:

```
oracle@xrac1:~ > rsh xrac2 date
Mon Nov 27 17:29:05 PST 2002
oracle@xrac1:~ > rcp /etc/hosts xrac2:/tmp
oracle@xrac1:~ > rlogin xrac2
Last login: Wed Nov 27 16 17:29:03 from xrac1.sanmateo.ibm.com
Have a lot of fun...
oracle@xrac2:~ > ls /tmp/hosts
/tmp/hosts
oracle@xrac2:~ > exit
logout
rlogin: connection closed.
oracle@xrac1:~ >
```

Example 5. Testing host equivalence on xrac1

FAStT Host Adapter Driver

Install the IBM FAStT Host Adapter driver. We downloaded the package consisting ot two files, 01r1093.tgz and 01r1093.txt, from the IBM Support site. Use the following steps to install the driver.

1. Rebuild the dependencies for the kernel. Enter the following command as root user:

```
cd /usr/src/linux-2.4.18.SuSE
make dep
```

```
cd /root
```

2. Untar the IBM FAStT Host Adapter failover device driver version 6.0 for Linux.

```
tar xvfz 01r1093.tgz
```

The subdirectory named i2x00-v6.0 will be created.

3. Enter the following commands:

```
cd i2x00-v6.0
./libinstall
./drvsetup
make all OSVER=linux
```

4. Replace the existing drivers:

```
cp qla*.o /lib/modules/2.4.18-64GB-SMP/kernel/drivers/scsi/
```

5. Verify the order of the device modules:

Edit /etc/rc.config. Look for the line beginning with INITRD_MODULES. If your local disk drives are connected to the ServerRaid adapter, make sure that the ServerRaid adapter driver, ips, precedes the FAStT HBA driver (Example 6).

```
#
# This variable contains the list of modules to be added to the initial
# ramdisk by calling the script "mk_initrd"
# (like drivers for scsi-controllers, for lvm or reiserfs)
#
INITRD_MODULES="aic7xxx ips qla2200"
```

Example 6. Extract from /etc/rc.config – INITRD_MODULES

6. Recreate the initial RAM disk:

```
mk_initrd
lilo
```

7. Reboot the server.

8. After server reboot, log in as `oracle`. Switch to root user and verify that the FAStT HBA driver has loaded. Use the `lsmod` command to display the loaded kernel modules. You should see qla2200 in the displayed list of modules.

CONFIGURING SHARED STORAGE

SOFTWARE INSTALLATION

Install IBM FAStT Storage Manager version 8.21 for Red Hat Linux. This will allow you to configure the FAStT500 storage. To use this program, run as root: `SMclient`

Install IBM FAStT MSJ Diagnostic and Configuration Utility version 2.0 release 33. This will allow you to configure the FAStT Host Adapter failover driver and to monitor the FAStT adapters behavior. Install the agent and the GUI tool.

- To start the agent, run as root: `qlremote &`.
- To use the GUI program, run as root: `/opt/FAStT_MSJ/FAStT`

DATAFILES FOR ORACLE9*i* RAC

Oracle9*i* RAC stores the datafiles in raw devices. On platforms supporting an Oracle-certified clustered file system, these datafiles may be stored on the clustered file system. At the time we wrote this paper, Oracle Clustered File System (OCFS) became available for Red Hat Advanced Server 2.1, but it was not yet available on SLES7.

In order to configure the storage and to create the raw devices, draw a table like the one shown in Table 3. This will help you to determine how to configure the shared storage device, how many logical devices you have to create, what size to give the logical devices, and it will help you to remember the purpose of each raw device.

The first column is the description of the database file that will be on the raw device, the second is the size of the database file in our example, the third and fourth are the volume groups and logical volumes created using LVM, and the fifth is the raw device which is mapped to the logical volume.

Database file	Size (MB)	Volume Group	Logical Volume	Raw device
quorum	110	controlvg	quorumlv	/dev/raw1
srvm	110	controlvg	srvmlv	/dev/raw2
spfile	5	controlvg	spfilelv	/dev/raw3
control file 1	110	controlvg	control1lv	/dev/raw4
control file 2	110	controlvg	control2lv	/dev/raw5
control file 3	110	controlvg	control3lv	/dev/raw6
system	400	datavg	systemlv	/dev/raw7
temp	110	datavg	templv	/dev/raw8
tools	12	datavg	toolslv	/dev/raw9
users	120	datavg	userslv	/dev/raw10
index	70	indexvg	indexlv	/dev/raw11
redo log thread 1 group 1	120	redovg	redo1_1lv	/dev/raw12
redo log thread 1 group 2	120	redovg	redo1_2lv	/dev/raw13
redo log thread 2 group 1	120	redovg	redo2_1lv	/dev/raw14
redo log thread 2 group 2	120	redovg	redo2_2lv	/dev/raw15
undo tablespace thread 1	312	undovg	undo1lv	/dev/raw16
undo tablespace thread 2	312	undovg	undo2lv	/dev/raw17

Table 3. Database file correspondence and sizes

Configure the FAStT storage with the IBM Storage Manager client (`SMclient`), follow the instructions in the Implementing IBM eServer xSeries SANs Redbook. Once you have configured FAStT storage, reload the FAStT Host Adapter driver to force detection of the logical drives. Do this as root user on all server nodes, using the following commands:

```
rmmod qla2200
modprobe qla2200
```

To configure raw devices on SLES7, you may use the fdisk utility or Logical Volume Manager (LVM). LVM provides manageability and flexibility but is not required to implement Oracle9*i* RAC. For our installation, we used Logical Volume Manager.

1. Start YaST as root on one node of the cluster by typing the command `yast`. From the main menu, select **Installation settings**, then **Configure the Logical Volume Manager**. You will be asked to confirm that you want to configure LVM. Select **Yes** to start LVM setup.

2. We are going to create the Volume Group, representing one or more hard drives. First, we have to partition the logical drive we have created to make it usable by LVM. The window entitled ADMINISTRATION OF LVM VOLUME GROUPS is displayed as shown in Figure 8.

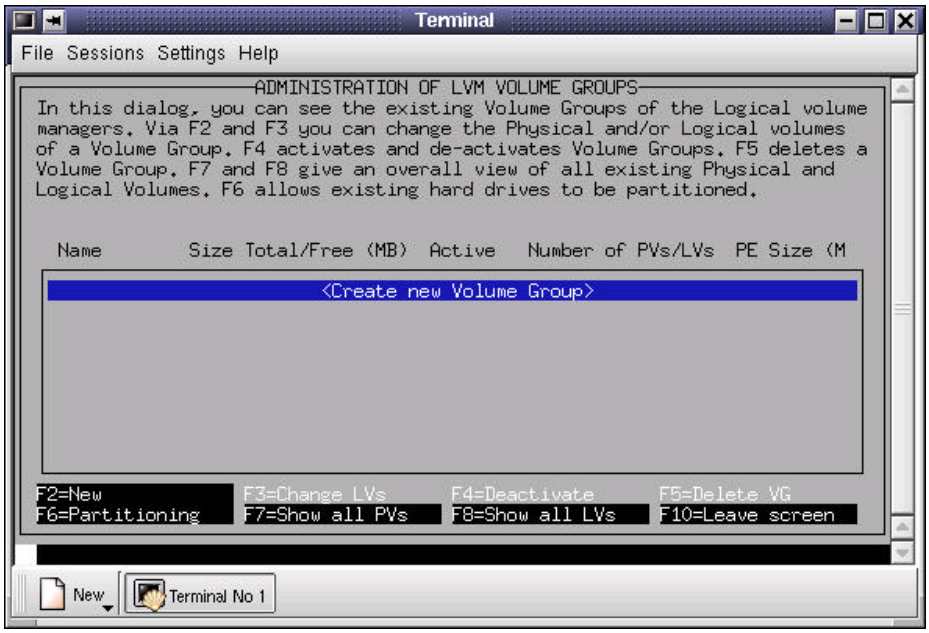

Figure 8. LVM Administration

a. Press F6 for Partitioning. You will be asked for confirmation. Press Enter to continue.

b. Select the device to be partitioned. In our case, it is /dev/sdb (see Figure 9):

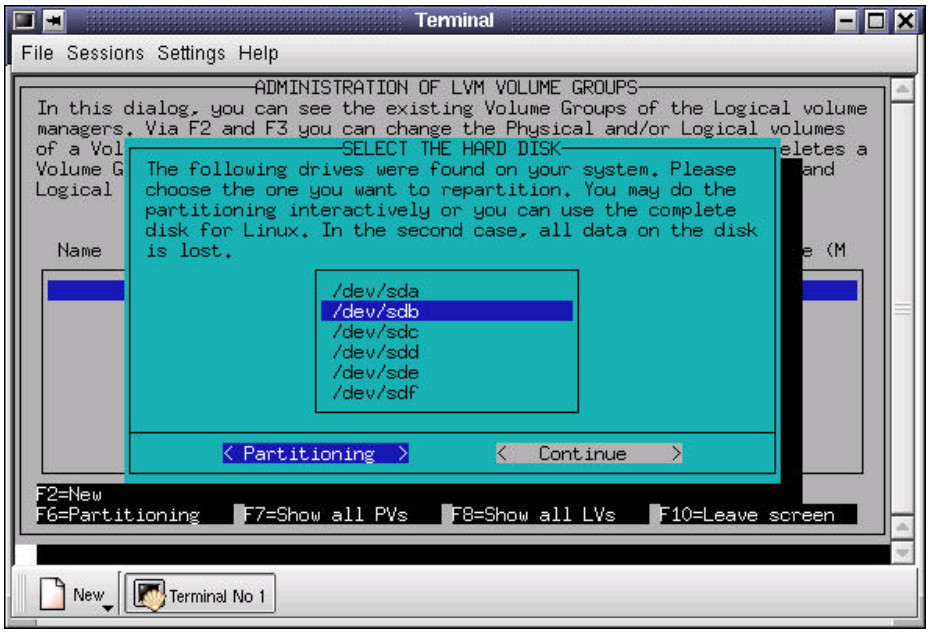

Figure 9. Select hard drive

c. Select **Partitioning** and press Enter.

d. Press F5 to create a new partition.

e. In the Partition Type window (see Figure 10), select **Primary partition**, then select **Continue** and press Enter.

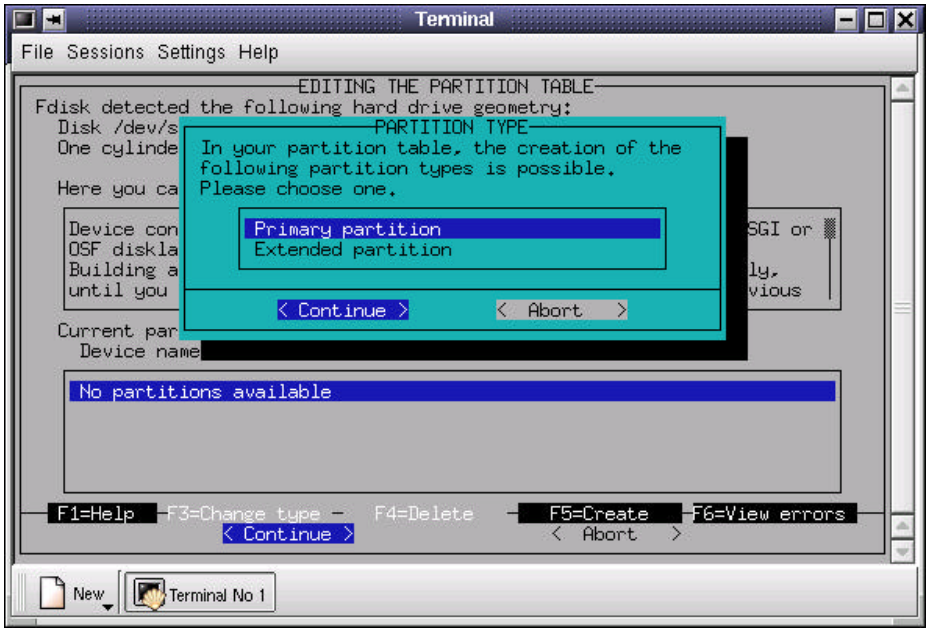

Figure 10. Partition type selection

f. In the Primary Partition window (Figure 11), select the first line (/dev/sdb1 in our system), then select **Continue** and press Enter.

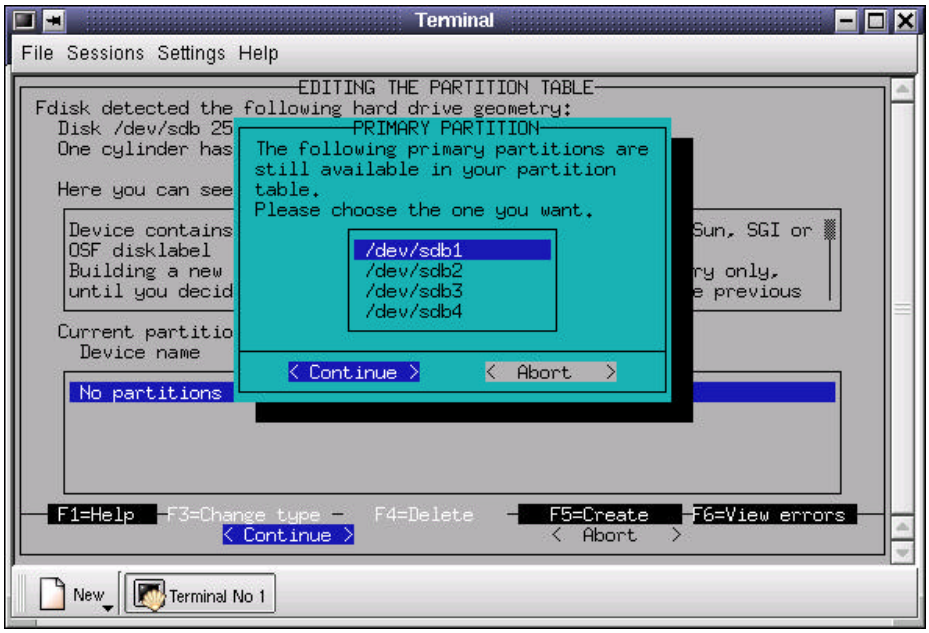

Figure 11. Partition number selection

g. In the Location of the Partition window (Figure 12), accept the default entries to create a single partition for the whole logical device. Select **Continue** and press Enter.

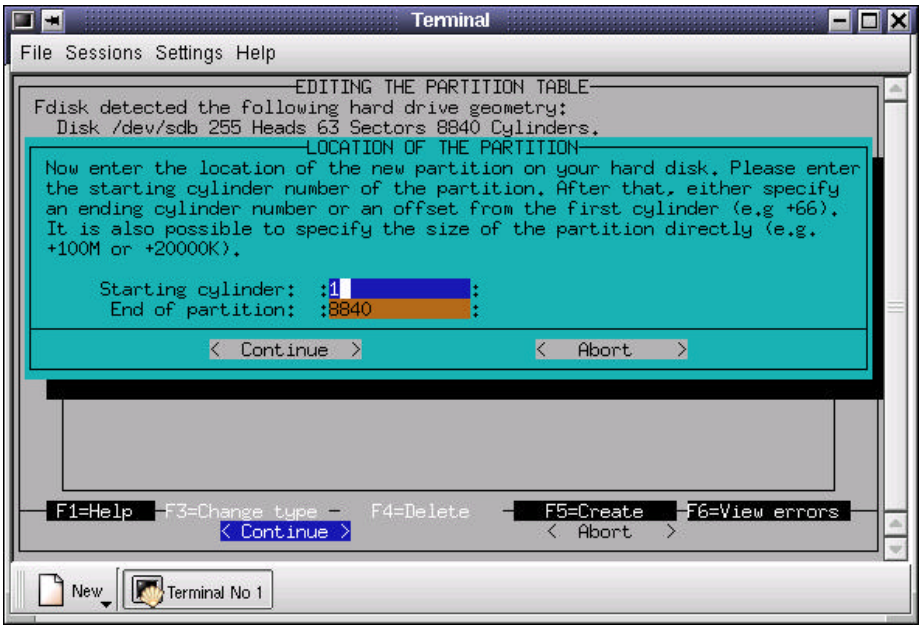

Figure 12. Location of the partition

h. The partition is now created. Press F3 to change its type.

i. In the Enter the Partition Type window (Figure 13), select **LVM Partition**, then select **Continue** and press Enter.

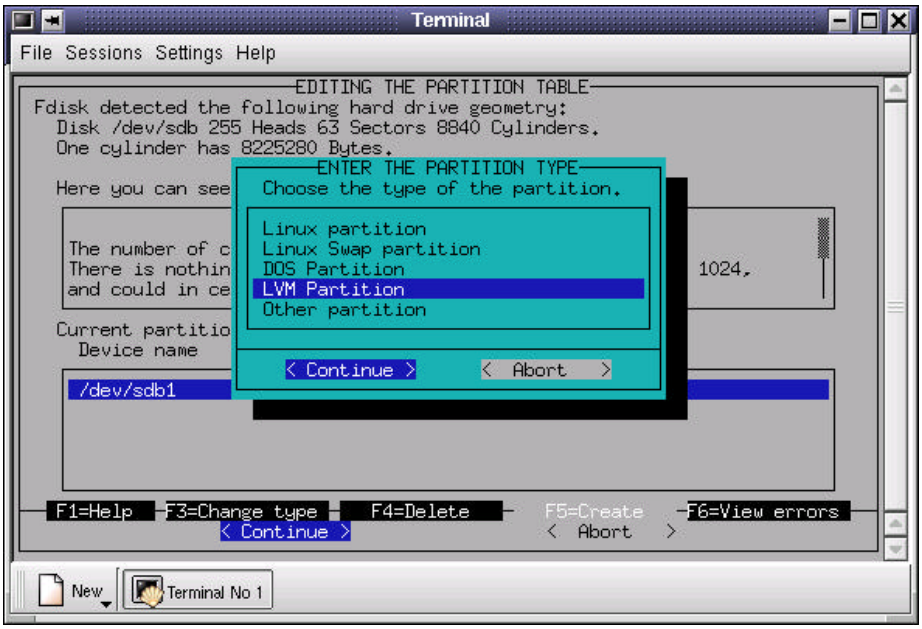

Figure 13. Partition type

j. The changes made are now reflected in the partition table shown in Figure 14.

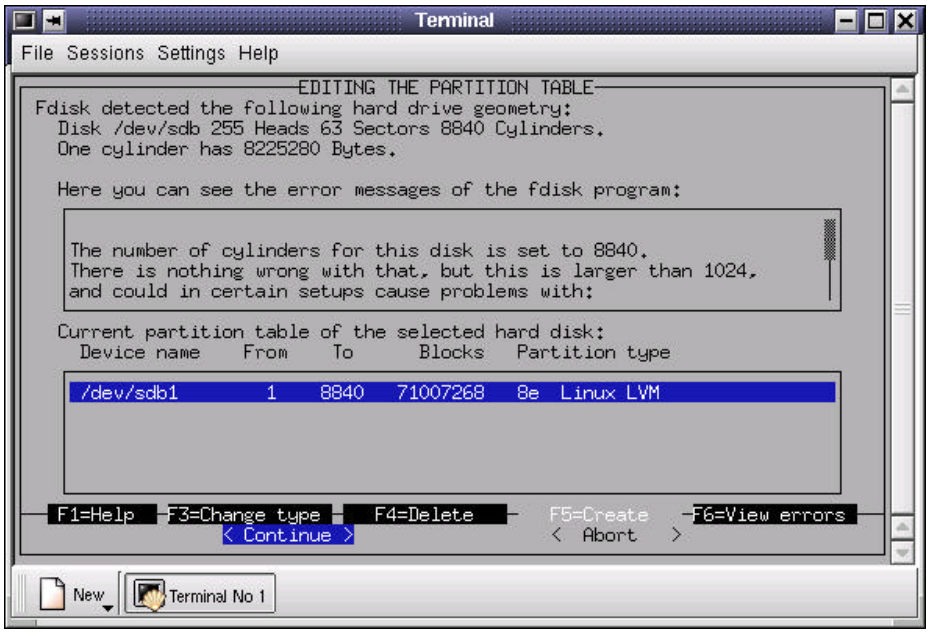

Figure 14. Partition table

k. Select Continue and press Enter. Press ESC to return to the LVM main window (see Figure 15).

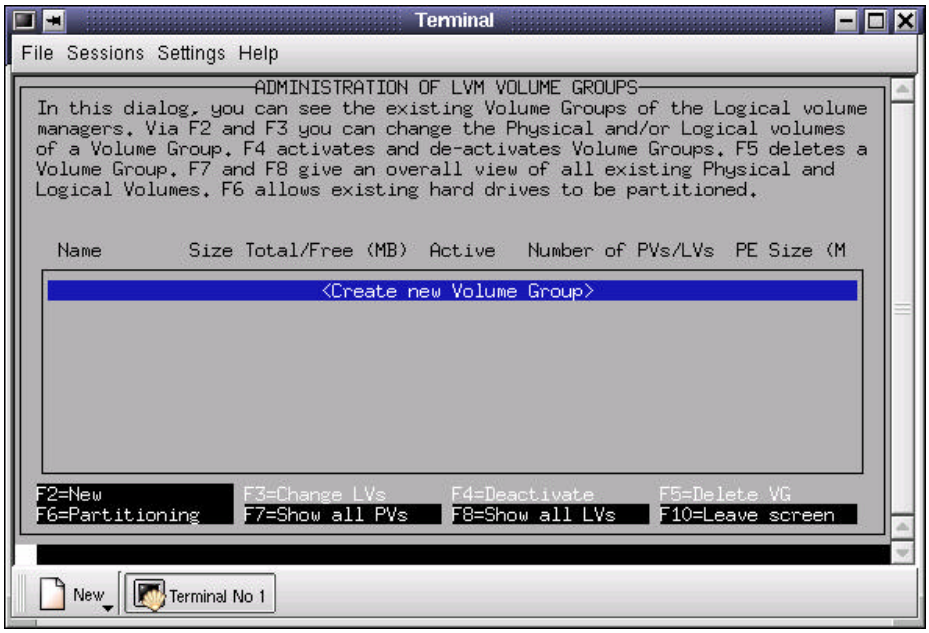

Figure 15. LVM main window

l. Press F2 to create a new Volume Group.

m. In the field Name of the Volume Group, enter the volume group name. Figure 17 shows the volume group `redovg` being created. Change or accept The size of a Physical extent, then highlight the appropriate device (/dev/sdc1) and select it by pressing the Spacebar. Select Create Volume group and press Enter.

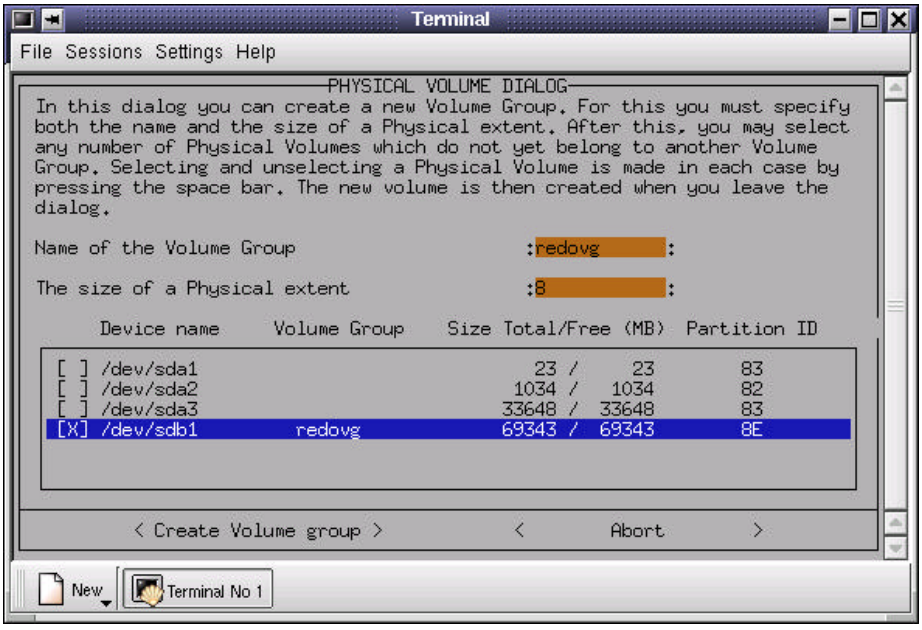

Figure 16. Physical Volume window

n. Repeat steps a - m to create the remaining volume groups. Once this is done, all the volume groups are shown in the Administration of LVM Volume Groups window (Figure 17).

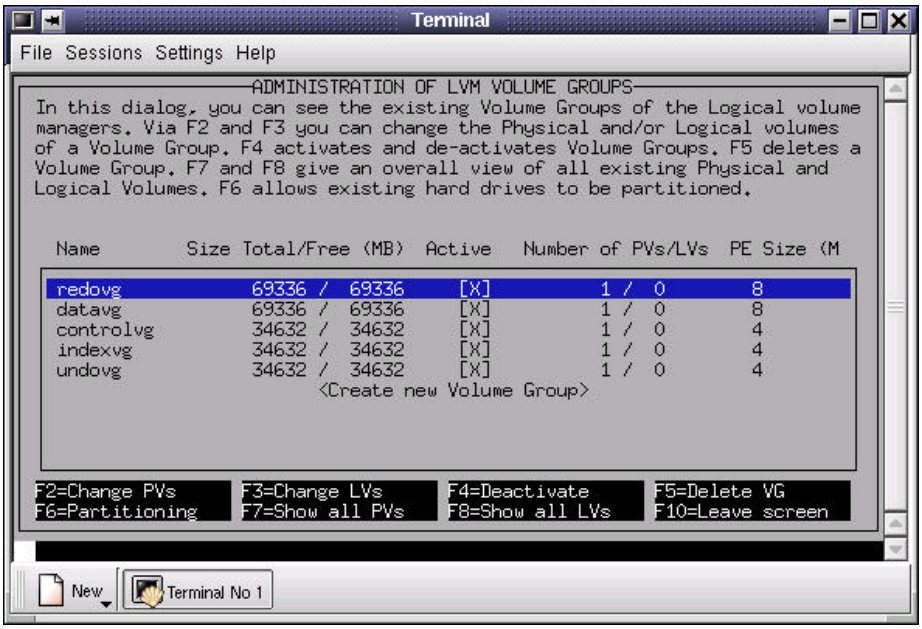

Figure 17. LVM main window

3. The next phase of the process is to create the logical volumes within the volume groups. The following steps illustrate the procedure for the redo1_1lv logical volume.
   a. Select a volume group (for example, redovg) in the LVM window. Press F3 to Change Logical Volumes. The Logical Volume window is displayed (Figure 18).

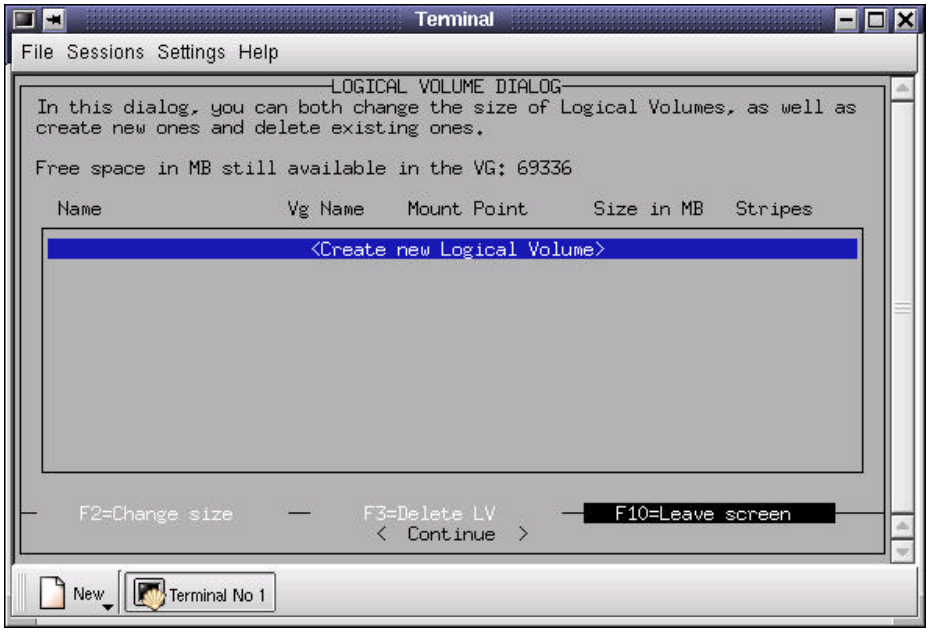

Figure 18. Logical Volume window

   b. Select Create new Logical Volume and press Enter.

   c. In the Create a New Logical Volume window (Figure 19), fill in the fields as follows:
      - Logical Volume name: redo1_1lv
      - Logical Volume size: 120M (per Table 3)
      - Number of stripes: 1

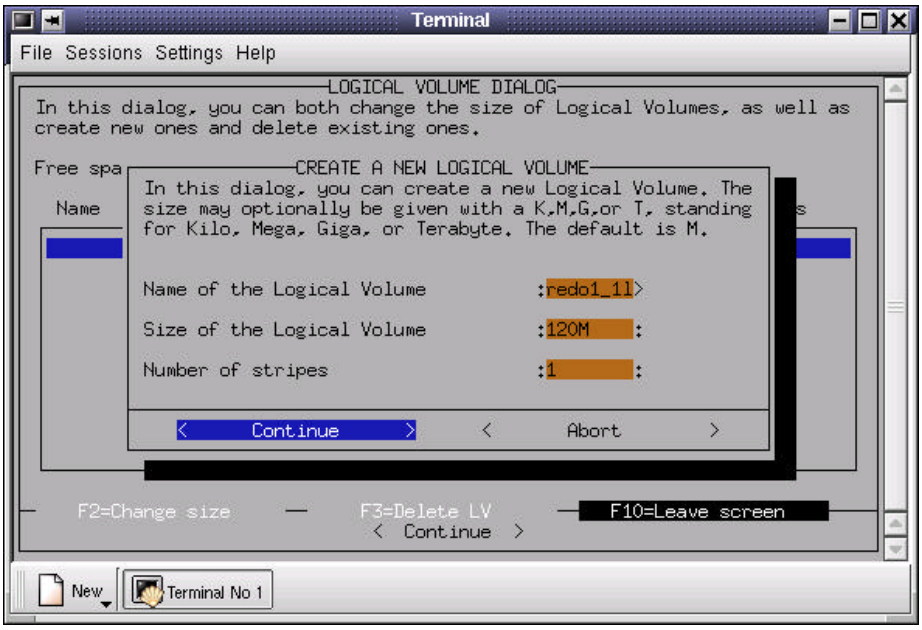

Figure 19. Create a new Logical Volume window

d. Select Continue and press Enter.

e. Repeat steps b - d to create the other logical volumes in the redovg volume group. When done, the Logical Volume window will list the logical volumes in the redovg volume group (Figure 20).

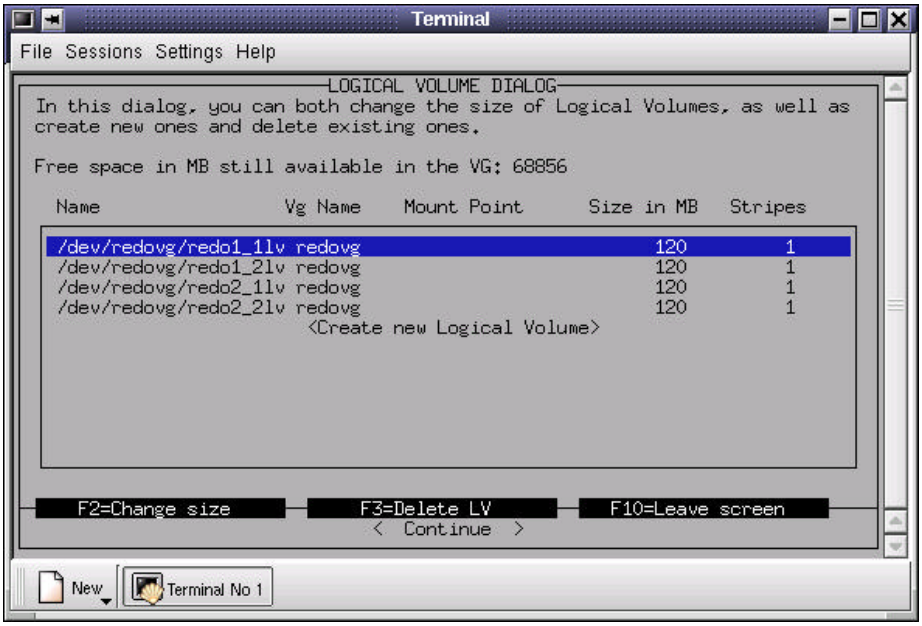

Figure 20. Logical Volume window – redovg volume group

f. Press F10 to return to the LVM main window. Repeat steps a - e to create logical volumes for the other volume groups - datavg, controlvg, indexvg, and undovg.

g. Press F10 to exit LVM.

h. Press Esc twice to exit YaST.

Note: Check the directories /dev/redovg, /dev/datavg, /dev/controlvg, /dev/indexvg, /dev/undovg. These directories represent the volume groups created. Within these directories, you will find files representing the logical volumes created.

4. At this point, only the sever node used to run LVM sees the LVM configuration. Activate the LVM configuration on the other nodes with the `vgscan` command as root user.

5. Map the logical volumes to raw devices. Perform the following steps on each server node:

   a. Create the script oracle_raw_devices in /etc/rc.d as shown in Example 7. This script must be owned by root and have execute permission.

```
# bind raw devices

/usr/sbin/raw /dev/raw1 /dev/controlvg/quorumlv
/usr/sbin/raw /dev/raw2 /dev/controlvg/srvmlv
/usr/sbin/raw /dev/raw3 /dev/controlvg/spfilelv

/usr/sbin/raw /dev/raw4 /dev/controlvg/control1lv
/usr/sbin/raw /dev/raw5 /dev/controlvg/control2lv
/usr/sbin/raw /dev/raw6 /dev/controlvg/control3lv

/usr/sbin/raw /dev/raw7 /dev/datavg/systemlv
/usr/sbin/raw /dev/raw8 /dev/datavg/templv
/usr/sbin/raw /dev/raw9 /dev/datavg/toolslv
/usr/sbin/raw /dev/raw10 /dev/datavg/userslv

/usr/sbin/raw /dev/raw11 /dev/indexvg/indexlv

/usr/sbin/raw /dev/raw12 /dev/redovg/redo1_1lv
/usr/sbin/raw /dev/raw13 /dev/redovg/redo1_2lv
/usr/sbin/raw /dev/raw14 /dev/redovg/redo2_1lv
/usr/sbin/raw /dev/raw15 /dev/redovg/redo2_2lv

/usr/sbin/raw /dev/raw16 /dev/undovg/undo1lv
/usr/sbin/raw /dev/raw17 /dev/undovg/undo2lv
```

Example 7. oracle_raw_devices

   b. To map the raw devices at server boot time, we configured the system to run the script at runlevel 5. To do this, create a symbolic link to the script file in /etc/rc.d/rc5.d. We used the 25th position in the startup sequence, using the following commands as root user:

```
cd /etc/rc.d/rc5.d
ln –s ../oracle_raw_devices S25raw_devices
```

   Note: The file named /etc/raw, together with a script named /etc/rc.d/raw and the symbolic link /etc/rc.d/rc5.d/S10raw provide a facility for mapping raw devices during startup. However, we did not use this facility because code within /etc/rc.d/raw prevents execution when invoked using S10raw.

   c. The raw devices have to be accessed by the oracle user id, so set permissions so that all the raw devices are owned by root and the group oinstall. Give the owner and group read/write access. This is done by executing the following commands as root user:

```
chown root:oinstall /dev/raw*
```

```
chmod 660 /dev/raw*
```

d. Execute the /etc/rc.d/oracle_raw_devices as root. Verify the raw device bindings with the `raw –qa` command.

6. Create symbolic links to the raw devices. Do this in each server node as oracle user. In our example, we placed these symbolic links in $ORACLE_HOME/oradata/XRAC, where XRAC is our database name. You may do this using the following script:

```
#  symbolic links to raw devices for database XRAC
mkdir $ORACLE_HOME/oradata
mkdir $ORACLE_HOME/oradata/XRAC
ln -s /dev/raw3 $ORACLE_HOME/oradata/XRAC/spfile
ln -s /dev/raw4 $ORACLE_HOME/oradata/XRAC/control01.ctl
ln -s /dev/raw5 $ORACLE_HOME/oradata/XRAC/control02.ctl
ln -s /dev/raw6 $ORACLE_HOME/oradata/XRAC/control03.ctl

ln -s /dev/raw7 $ORACLE_HOME/oradata/XRAC/system01.dbf
ln -s /dev/raw8 $ORACLE_HOME/oradata/XRAC/temp01.dbf
ln -s /dev/raw9 $ORACLE_HOME/oradata/XRAC/tools01.dbf
ln -s /dev/raw10 $ORACLE_HOME/oradata/XRAC/users01.dbf

ln -s /dev/raw11 $ORACLE_HOME/oradata/XRAC/indx01.dbf
ln -s /dev/raw12 $ORACLE_HOME/oradata/XRAC/redo1_1.log
ln -s /dev/raw13 $ORACLE_HOME/oradata/XRAC/redo1_2.log
ln -s /dev/raw14 $ORACLE_HOME/oradata/XRAC/redo2_1.log
ln -s /dev/raw15 $ORACLE_HOME/oradata/XRAC/redo2_2.log

ln -s /dev/raw16 $ORACLE_HOME/oradata/XRAC/undotbs1.log
ln -s /dev/raw17 $ORACLE_HOME/oradata/XRAC/undotbs2.log
```

Example 8. Create symbolic links for datafiles

Adding datafiles

If you create a new logical device on the SAN storage and add a logical volume (using LVM) or a partition (using fdisk), you have to reload the FAStT adapter driver. This means you lose the disk access for some time, and therefore the database instance running on the server has to be shut down.

To maximize database availability, it is recommended that spare partitions be created for use in the future as new datafiles.

When you run out of spare partitions, for each node in the cluster, you will have to shut down the instance, reload the FAStT adapter driver and restart the instance. This has to be done one node after another, in order to have the database always online.

ORACLE9*I* INSTALLATION

PREPARING FOR INSTALLATION

Oracle 9*i* Release 2 is supplied on three CDs, so it is necessary to switch product CDs during installation. Oracle Universal Installer (OUI) will manage the switching between CDs and ask you for the CD location when necessary.

In our installation, we copied the Oracle9*i* product CDs to the hard disk. The three product CDs were copied to three directories on the server named xrac1. These directories were named /opt/oracle/prod92/disk1, /opt/oracle/prod92/disk2 and /opt/oracle/prod92/disk3.

If you chose to install from the CDs, do not run OUI while the CD is the current directory, or you will not be able to switch CDs when prompted.

The following pre-installation tasks have to be performed on all nodes.

1. Modify the kernel parameters that affect resources used by Oracle9*i*: This must be performed as `root` user.
   a. Review kernel parameter recommendations in *Oracle9i Release Notes Release 2 (9.2.0.1.0) for Linux Intel (32-bit)* and *Installation Guide, Release 2 (9.2.0.1.0) for UNIX Systems: AIX-Based Systems, Compaq Tru64 UNIX, HP 9000 Series HP-UX, Linux Intel, and Sun Solaris*
   b. In our example, we increased SEMMNS to 1024 while retaining the minimum settings for SEMMSL, SEMOPM and SEMMNI. Example 9 shows a shell script containing the commands used to modify the kernel parameters.

```
export SEMMSL=100
export SEMMNS=1024
export SEMOPM=100
export SEMMNI=100
echo  $SEMMSL $SEMMNS $SEMOPM $ SEMMNI > /proc/sys/kernel/sem
export SHMMAX=2147483648
echo $SHMMAX > /proc/sys/kernel/shmmax
```

Example 9. Kernel parameters modification script

   These kernel parameters need to be initialized during system startup, so we named the script `oracle_rac` and stored it in /etc/rc.d.
   c. Add execute permission to this script and set it up to run at runlevel 5. Do this with the following commands:

```
chmod u+x /etc/rc.d/oracle_rac
cd /etc/rc.d/rc5.d
ln –s ../oracle_rac S26oracle_rac
```

   d. Execute /etc/rc.d/oracle_rac to modify the current kernel settings.

2. A number of environmental variables have to be set up for the `oracle` user ID. Do this as `oracle` user.
   a. Modify .bash_profile in oracle's home directory, adding the statements shown in Example 10.

```
# .bash_profile

xhost + > /dev/null 2>&1
export DISPLAY=localhost:0.0

export ORACLE_BASE=/oracle
export ORACLE_HOME=$ORACLE_BASE
export OH=$ORACLE_HOME
export ORACLE_TERM=xterm
export TNS_ADMIN=$ORACLE_HOME/network/admin
export NLS_LANG=AMERICAN_AMERICA.US7ASCII
export ORA_NL33=$ORACLE_HOME/ocommon/nls/admin/data
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/lib:$ORACLE_HOME/oracm/lib
export
CLASSPATH=$ORACLE_HOME/jre:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/netw
ork/jlib
export THREADS_FLAG=native

export PATH=$ORACLE_HOME/bin:$HOME/bin:$PATH
```

Example 10. Setting environmental variables in .bash_profile

b. Verify that the environmental variables are set correctly by loggin in using the oracle user ID and echoing a variable to the console. Here is an example:

```
oracle@xrac1:~ > echo $ORACLE_HOME
/oracle
oracle@xrac1:~ >
```

INSTALL ORACLE CLUSTER MANAGER

The first step in implementing Oracle9*i* RAC is to install and start Oracle Cluster Manager, which will let you install Oracle Database on every node on your cluster from a single node. Oracle Cluster Manager is installed using the Oracle Universal Installer (OUI).

Follow this procedure to install Oracle Cluster Manager and the Real Application Clusters software.. .

1. Log in as oracle on xrac1 and invoke the Oracle9*i* Universal Installer as follows:

```
cd /opt/oracle/prod92/disk1
./runInstaller
```

The Welcome screen is display as shown in The Oracle Universal Installer loads the Java Runtime Environment (you may receive some messages on your terminal window regarding the fonts but this is normal, because if the Java Runtime Environment does not find all the fonts it needs, then it substitutes them with others).

Figure 21. OUI Welcome screen

Click **Next** to go to the Inventory Location window shown in Figure 22. You will be prompted for the Inventory Location if this is the first time that you have started the Oracle Universal Installer on your system. The Inventory Location is the base directory into which Oracle Universal Installer will place the files needed to keep your Oracle installations centralized and up to date. The Oracle Inventory definition can be found in the file `/etc/oraInst.loc`. Accept the default values and then click **OK** to continue. If you choose another file location, keep in mind you should be consistent about this location in the future.

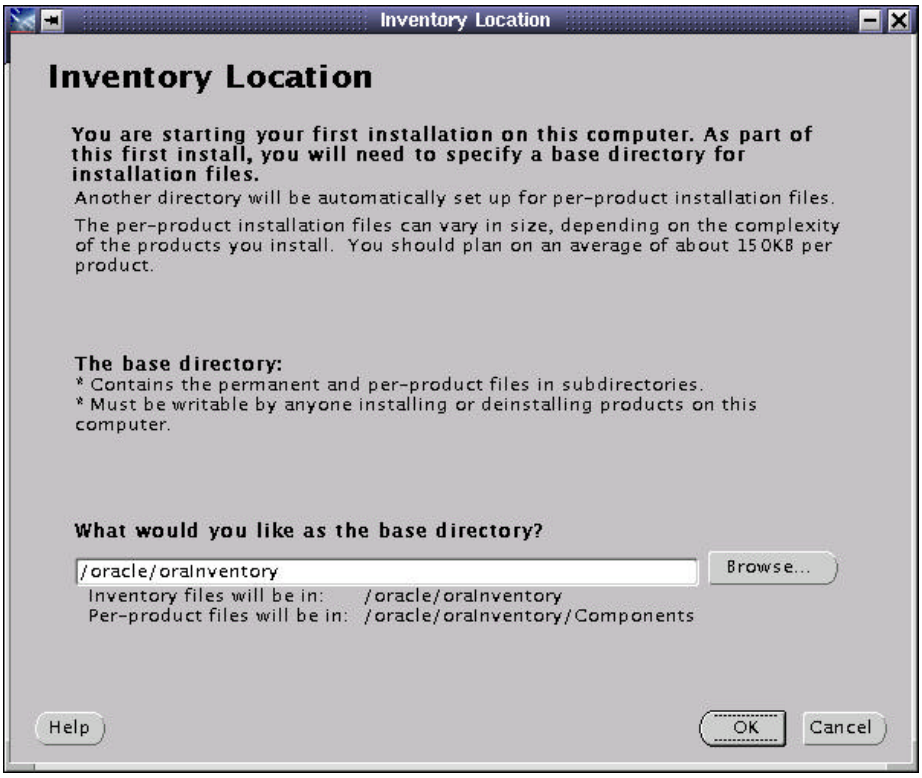

Figure 22. Inventory Location window

2. The next window (Figure 23) will prompt you for the UNIX group that the oracle user will use during the installation. This value should be "`dba`" or "`oinstall`" if you followed this document to install Oracle 9i RAC, or if you use the Step by Step procedure from Oracle that is available at Metalink's site. Enter "`oinstall`" and click **Next** to continue.

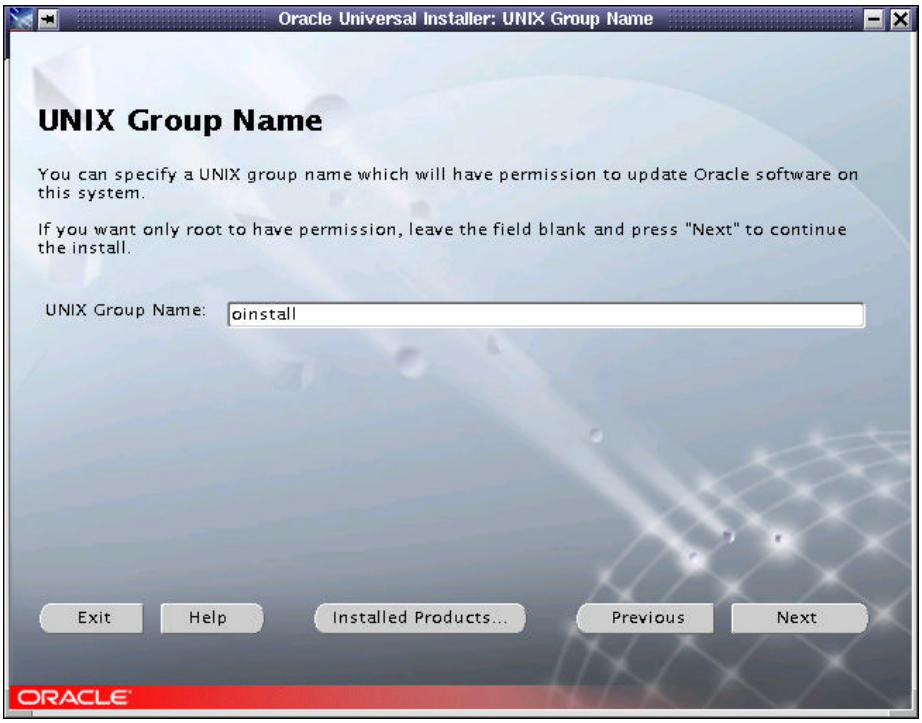

Figure 23. UNIX Group Name window

3. You will be then prompted to execute a shell script as root user (Figure 24). This script is located at "`/tmp/oraInst.sh`" and you must execute it in order for the installer to work properly. Open a terminal window, switch to root user and execute the script as directed.

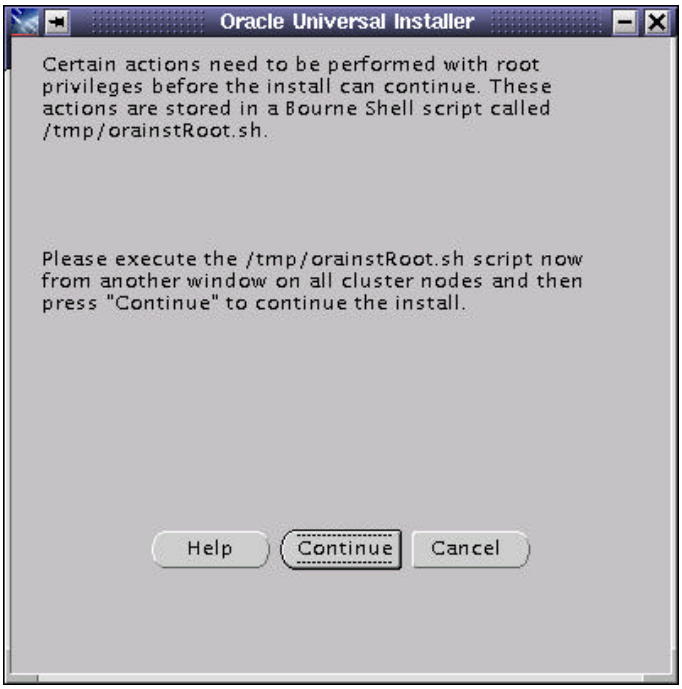

Figure 24. Pop-up window - Execute orainstRoot.sh

4. Once the shell script is executed, click **Continue** and the File Locations window appears (Figure 25). This window asks you for the source and destination paths for your installation. Do not change the source location. For the destination location, verify that the value shown is the ORACLE_HOME variable previously set in .bash_profile. If the values are correct, click **Next** to continue.

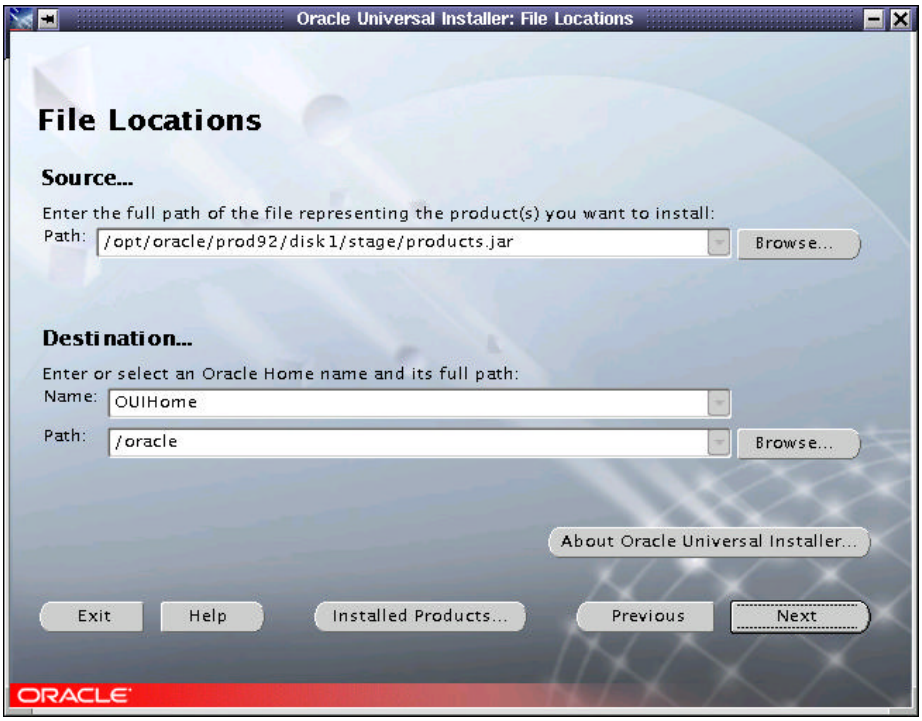

Figure 25. File Locations window

5. The Available Products window (Figure 26) shows the different products you can install from the CDs. Select “Oracle Cluster Manager” and click **Next** to start installation of the Cluster Manager.

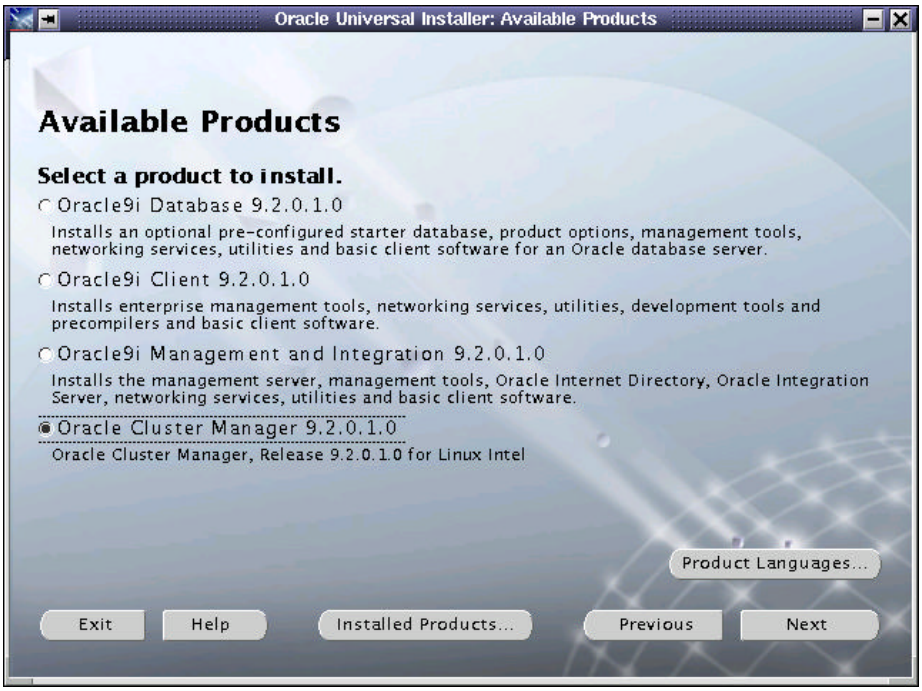

Figure 26. Available Products window

6. Oracle Universal Installer needs to know the public names or public IP addresses of the nodes of your cluster. The Public Node Information window will let you specify the node names. These nodes must be available from each other for the Oracle Universal Installer to configure the Real Application Clusters properly. In our case, our two xSeries 440s are named `xrac1` and `xrac2` (Figure 27).

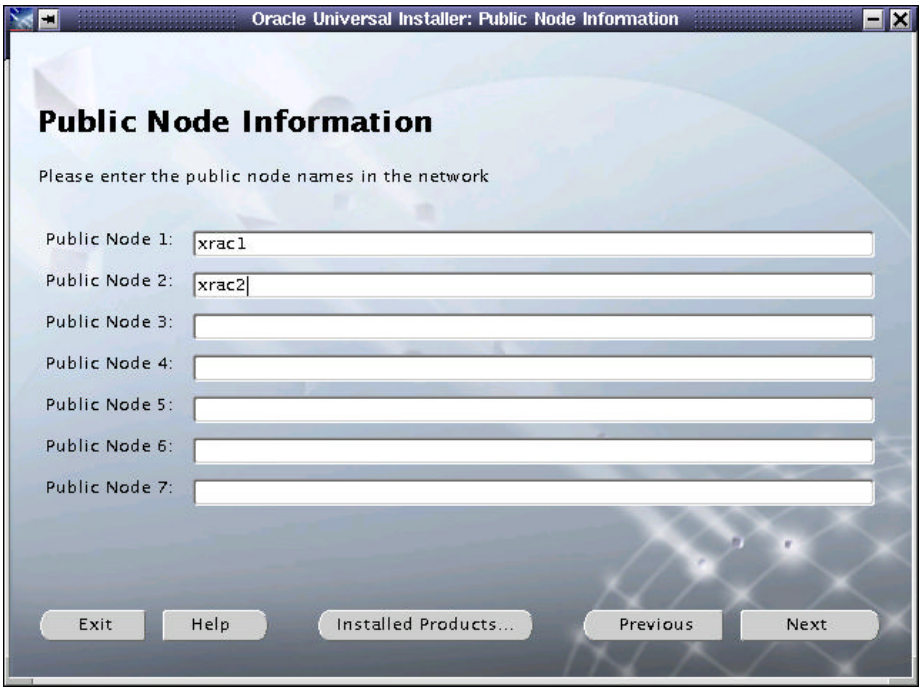

Figure 27. Public Node Information window

7. The Private Node Information window configures the Private Node addresses that will be used for the interconnect between the different nodes of your RAC cluster. In our example, the hostnames for the private interconnect are xrac1i and xrac2i. Enter the information as shown and click **Next** to continue.

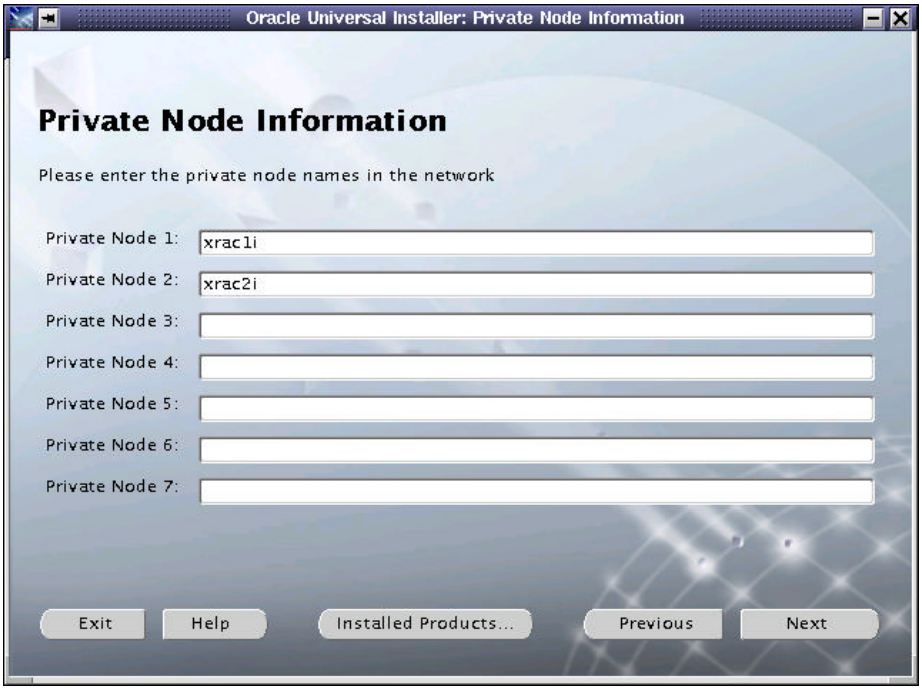

Figure 28. Private Node Information window

8. The Watchdog parameter window (Figure 29) asks you for the watchdog parameters. Simply accept the default value (`60000`) and click **Next** to continue.

Figure 29. WatchDog Parameter Information window

9. The Quorum Disk Information window displays next. Enter the full name of the raw device you have created for the node monitor for the Quorum disk information. In this case, our quorum device will be `/dev/raw1`. Enter /dev/raw1 and click **Next** to continue.

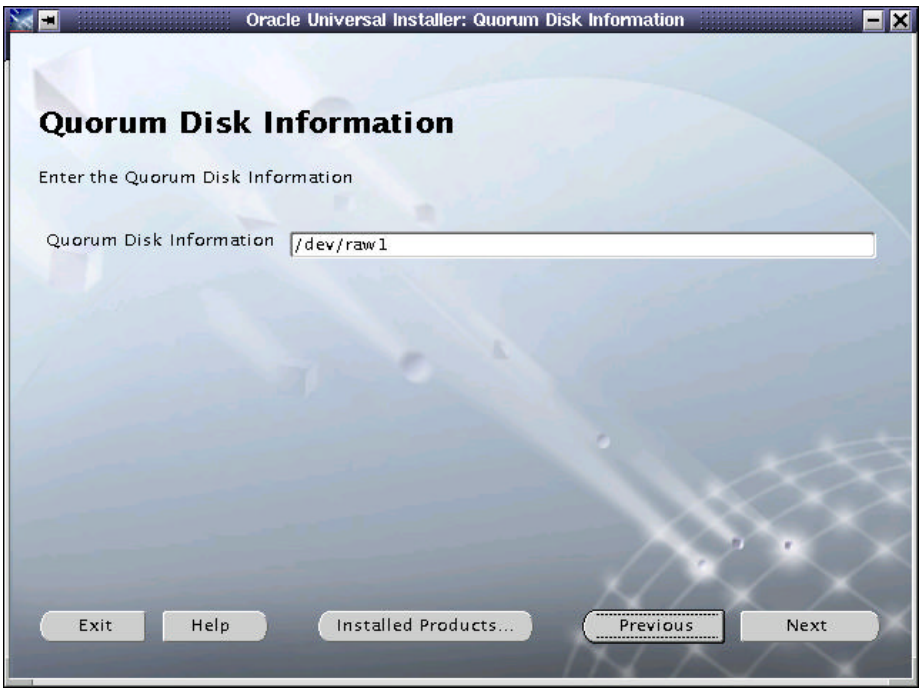

Figure 30. Quorum Disk Information window

10. The Summary window is displayed (Figure 31). Click **Install** to continue.

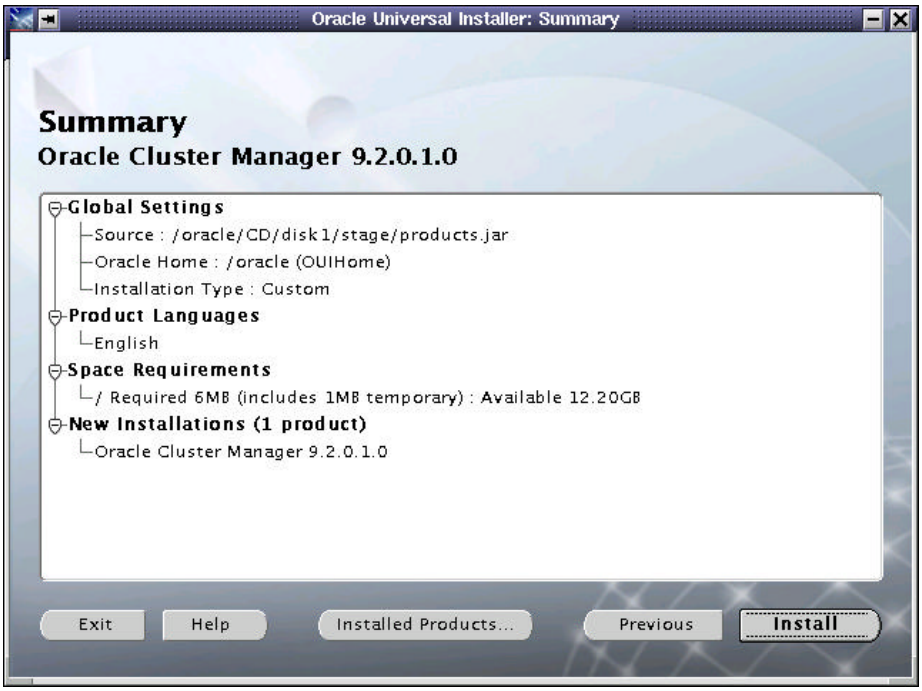

Figure 31. Summary window

11. An installation progress window (Figure 32) displays. Wait until the End of Installation window (Figure 33) displays.

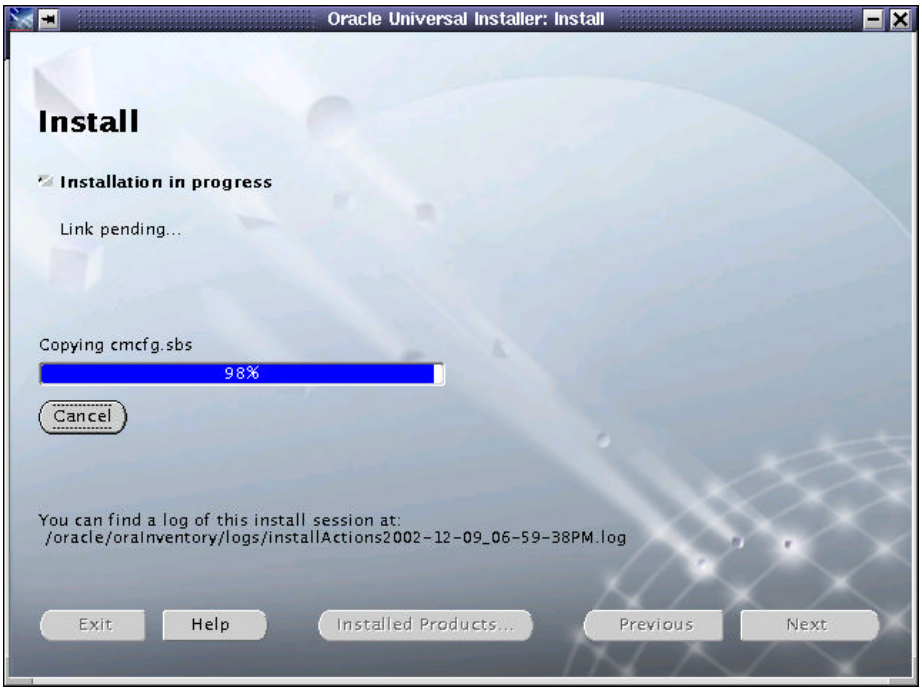

Figure 32. Install progress window

Figure 33. End of Installation window

12. Once the End of Installation window is displayed, click **Exit** and confirm by clicking **Yes**.

START ORACLE CLUSTER MANAGER

The next step is to start the Cluster Manager. This service has to run on each node of the cluster in order to install RAC.

1. Log in to one of the cluster nodes as oracle user and open a terminal window.

2. Before starting Oracle Cluster Manager, go to each cluster node and verify that the $ORACLE_HOME/oracm/log directory exists. If not, create the directory as oracle user. Do this with the following command:

```
mkdir $ORACLE_HOME/oracm/log
```

3. Switch to root on each node and execute the following commands:

```
export ORACLE_HOME=/oracle
$ORACLE_HOME/oracm/bin/ocmstart.sh
```

4. This will start the `watchdogd` daemon and the `oracm` process, which make up Oracle Cluster Manager. Verify you have started it on each node with the `ps -ef` command and check the oracm log: `$ORACLE_HOME/oracm/log/cm.log`

INSTALL ORACLE9*I* DATABASE

1. Start the Oracle Universal Installer again. Click Next on the Welcome screen and you will be prompted for the cluster nodes on which you would like to install the Database. This is because the Oracle Universal Installer detects the Cluster Manager. The Cluster Node Selection window is shown in Figure 34.

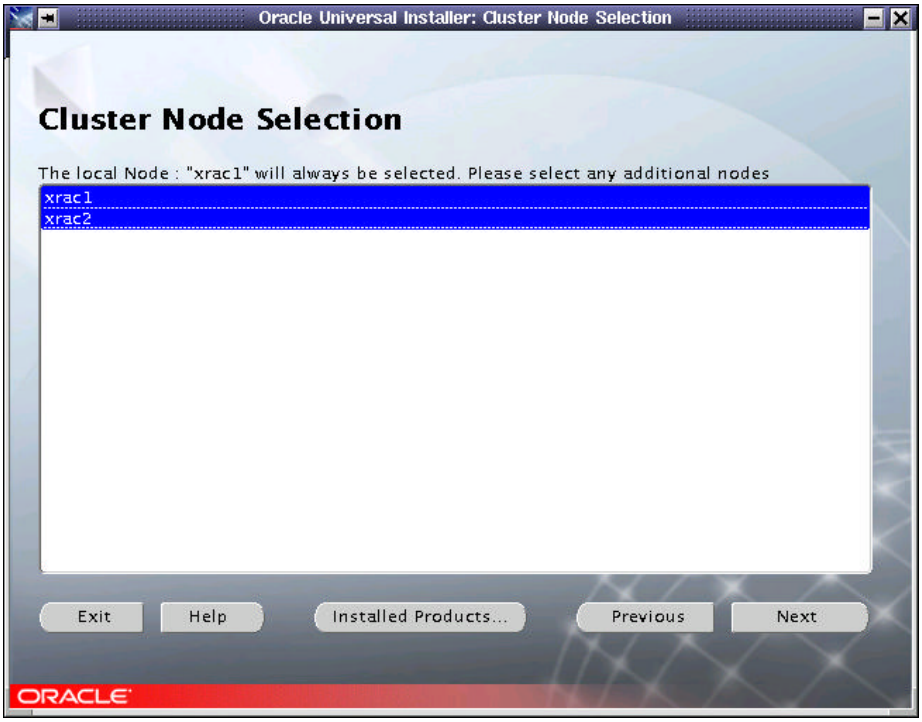

Figure 34. Cluster Node Selection window

Select all the nodes in your cluster and click **Next** to continue. The next window will be the File Location Screen that we have already seen in Figure 25. Click **Next** to continue.

2. Select **Oracle9*i* Database 9.2.0.1.0** from the Available Products window (Figure 35) and click **Next** to continue.

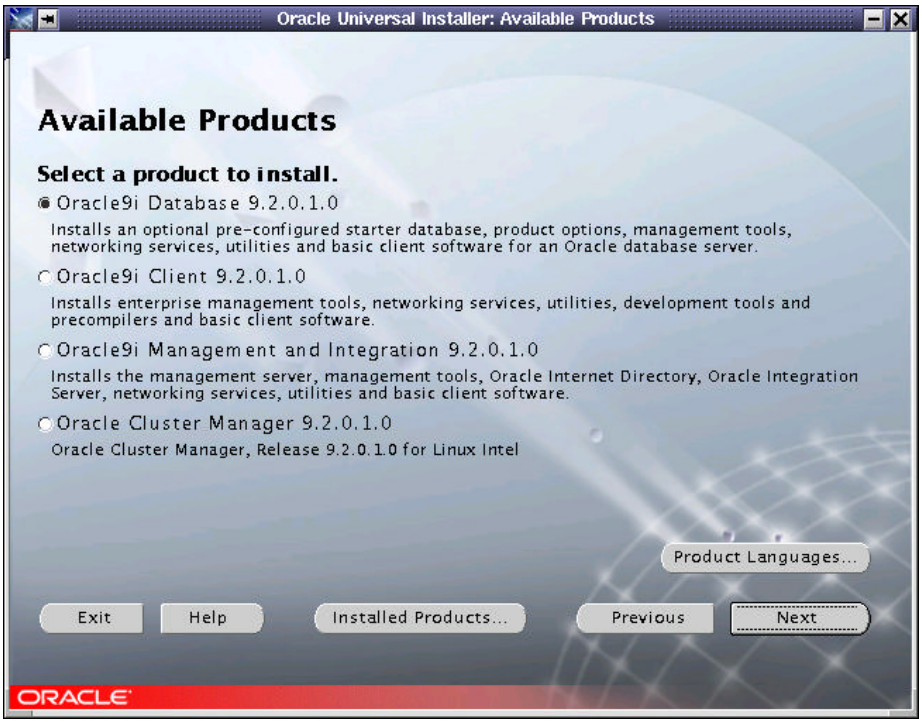

Figure 35. Available Products window

3. The next window displayed is the Installation Types window (Figure 36). Select **Custom** in order to have full control over the components you will install. Click **Next** to continue.

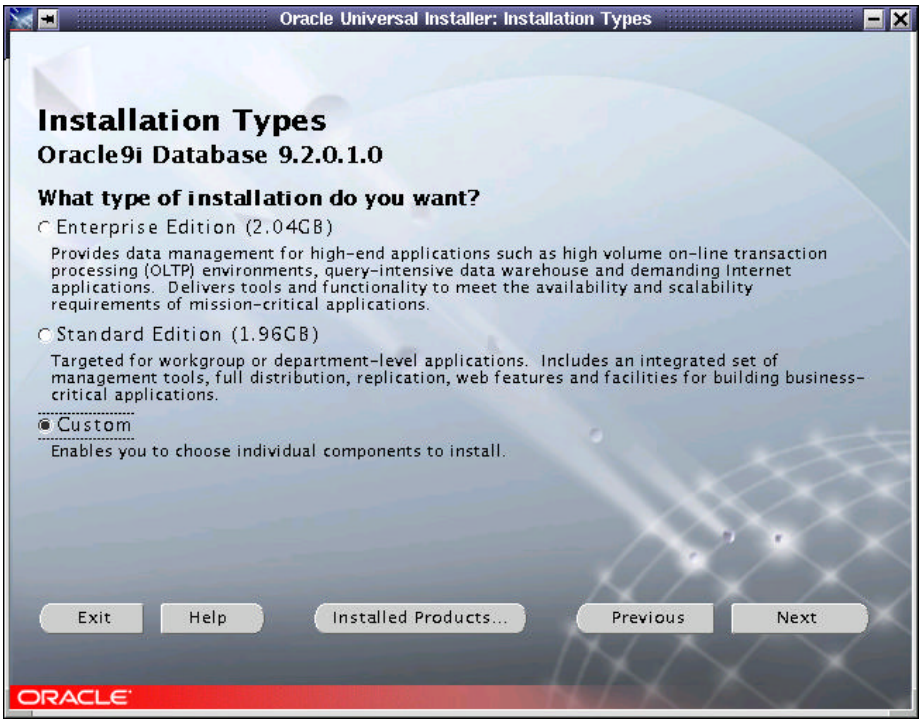

Figure 36. Installation Types window

4. The Available Product Components window is displayed (Figure 37). In this window, you can select components to be installed. For a Real Application Cluster configuration, be sure to install at least the Oracle Database, Real Application Clusters, Partitioning and Netlistener components. If you need additional components for your database, select them on this screen. Once you have selected all desired components, click **Next** to continue.

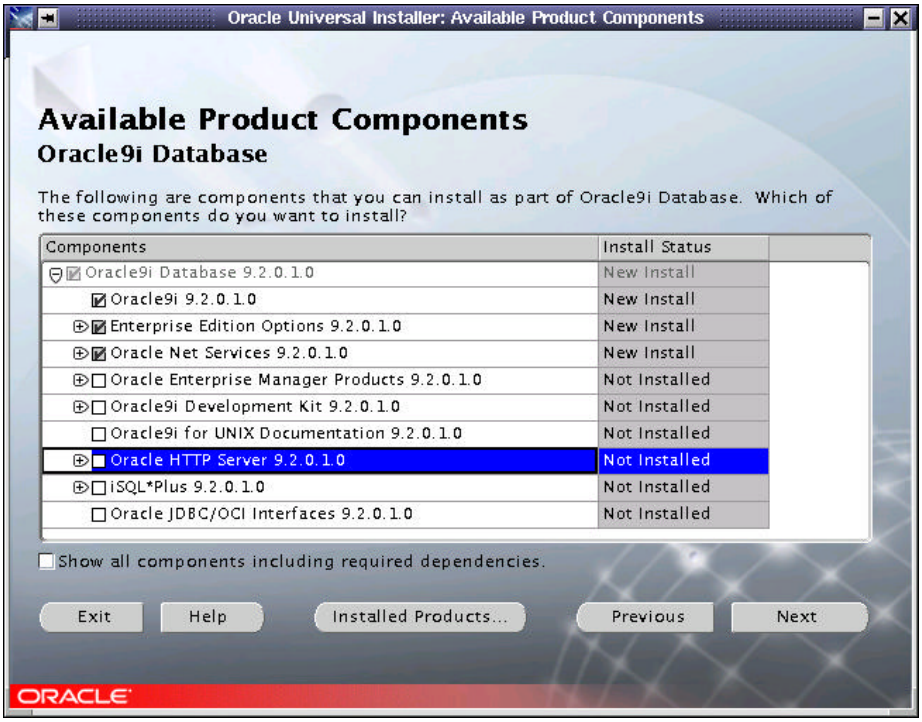

Figure 37. Available Product Components window

5. The Components Location window is displayed (Figure 38). Do not modify the values, click **Next** to continue.

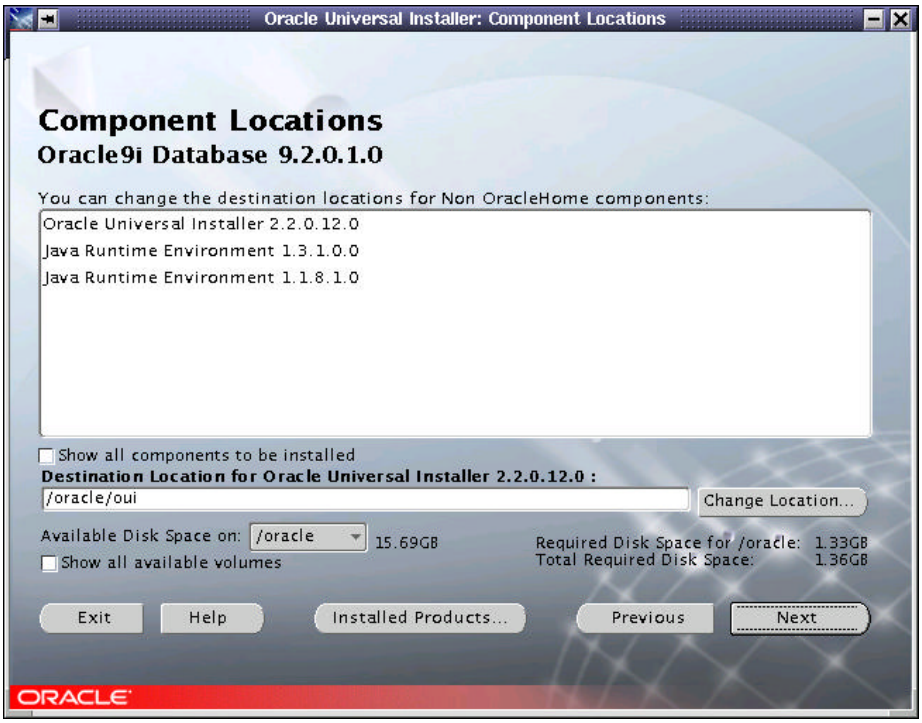

Figure 38. Component Locations window

6. You will also be prompted for the raw device where you will have the shared configuration of your cluster. Enter /dev/raw2 as shown in Figure 39. Click **Next** to continue.

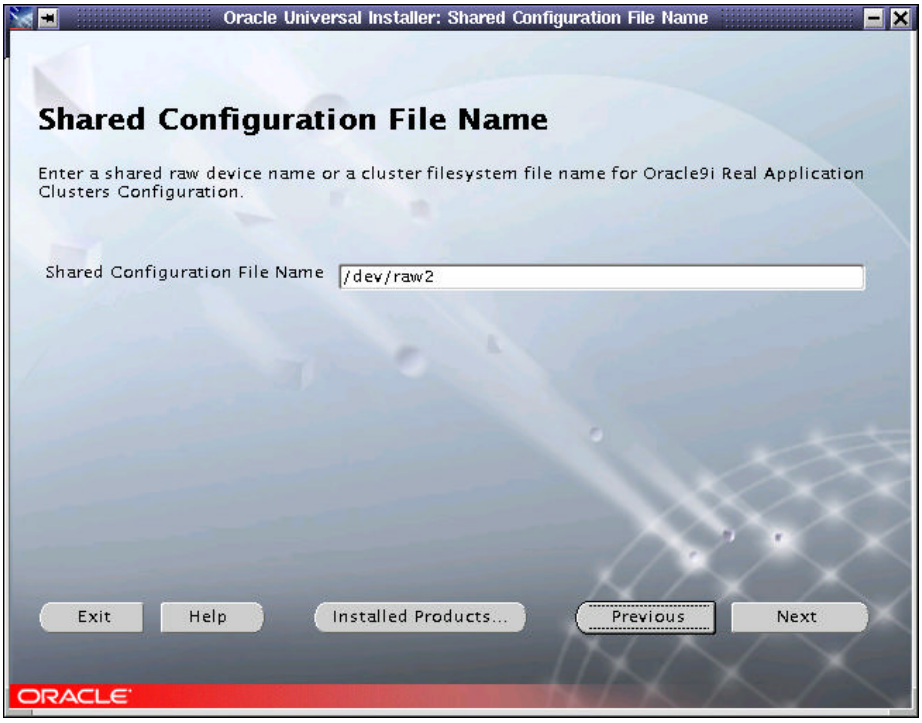

Figure 39. Shared Configuration File Name window

7. On the Privileged Operating System Groups window (Figure 40) select the group(s) for your database operator and administrator. Normally, these values are set to `dba`. Enter the appropriate information and click **Next** to continue.

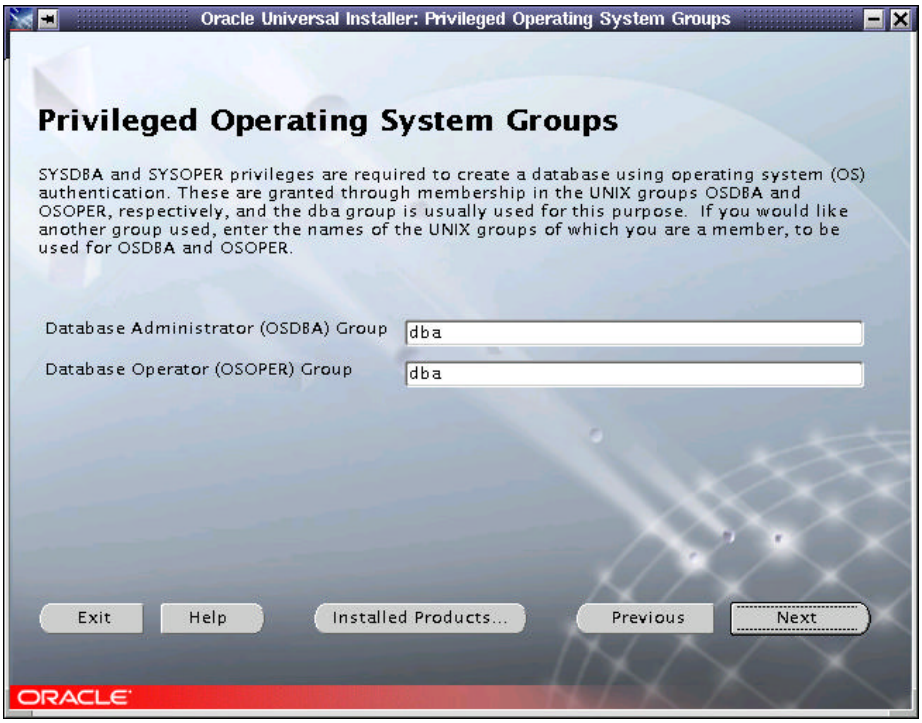

Figure 40. Privileged Operation System Groups window

8. You will then be prompted for database creation (Figure 41). We will not create it right now, so select **No** and then click **Next** to continue. If you wish to create a database during the installation, you should select "Yes" and follow the configuration steps. We will create the database later.

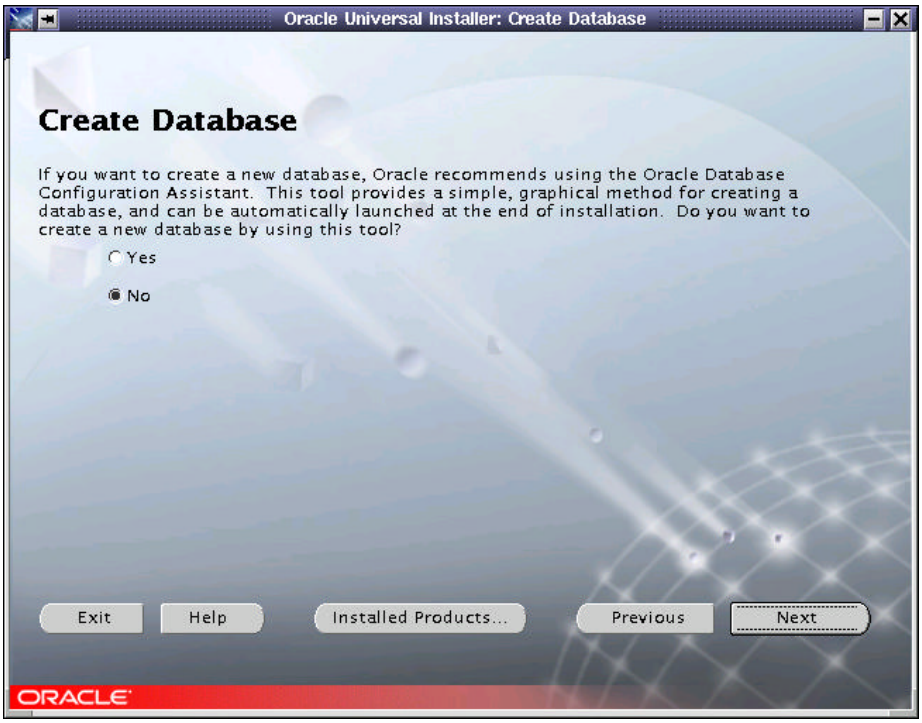

Figure 41. Create Database window

9. A Summary window shows you the product component selections you have made. Verify it and if you need to add something else, go back by clicking on **Back**. If everything is OK, click **Install** to start the installation procedure.

10. During the installation process, you will see an installation progress window such as shown in Figure 42.

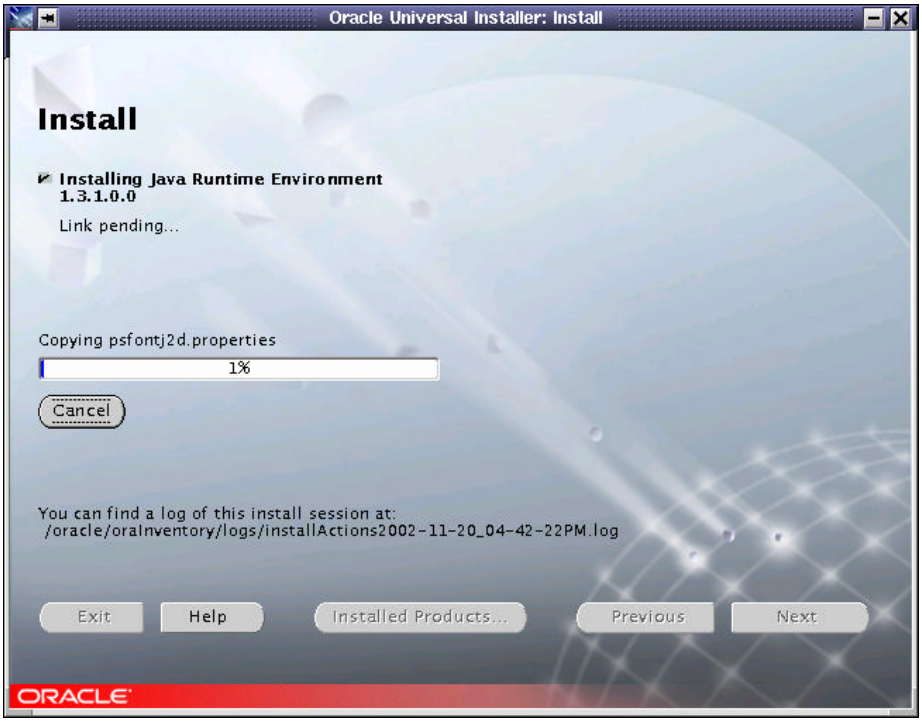

Figure 42. Install window

11. After a few minutes, a pop-up window entitled Disk Location displays (Figure 43). If you have copied the CDs to the disk, modify the path to indicate where the second CD's contents can be found. If you used a scheme like ours, change disk1 to disk2 and click **OK** to continue. If you are installing from the CDs, open another window, umount the CD, remove it, and insert the CD requested.

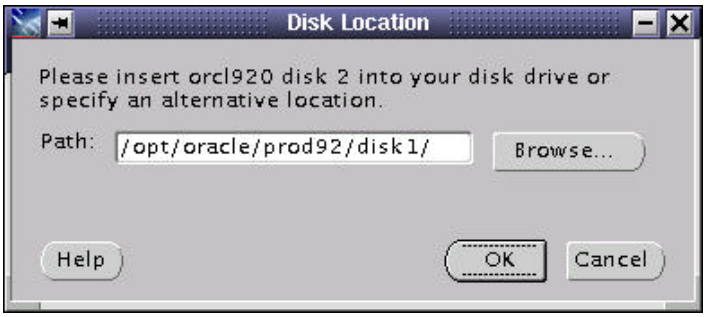

Figure 43. Disk Location window

12. After another few minutes, another Disk Location pop-up window is displayed to ask for the location of CD 3's contents. Modify the path as necessary and click OK to continue the installation.

13. Once all the CD's are installed on the first node, the Oracle Universal Installer will install the software on each node on your cluster. A pop-up window entitled Setup Privileges will ask you to execute `root.sh` (Figure 44).

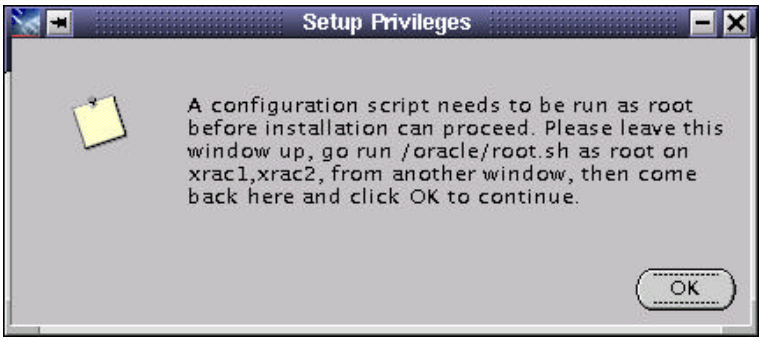

Figure 44. Setup Privileges window

14. Before executing `root.sh`, create the directory `/var/opt/oracle` and make `oracle` its owner on ***both*** nodes. Execute the following commands from the root shell prompt:

```
mkdir /var/opt/oracle
chown oracle:oinstall /var/opt/oracle
```

15. Oracle Universal Installer will automatically configure the cluster and then ask you to configure the network. We will configure the listeners later, so click **Cancel** on the Oracle Net Configuration Assistant: Welcome window and confirm by clicking **Yes** (Figure 45).

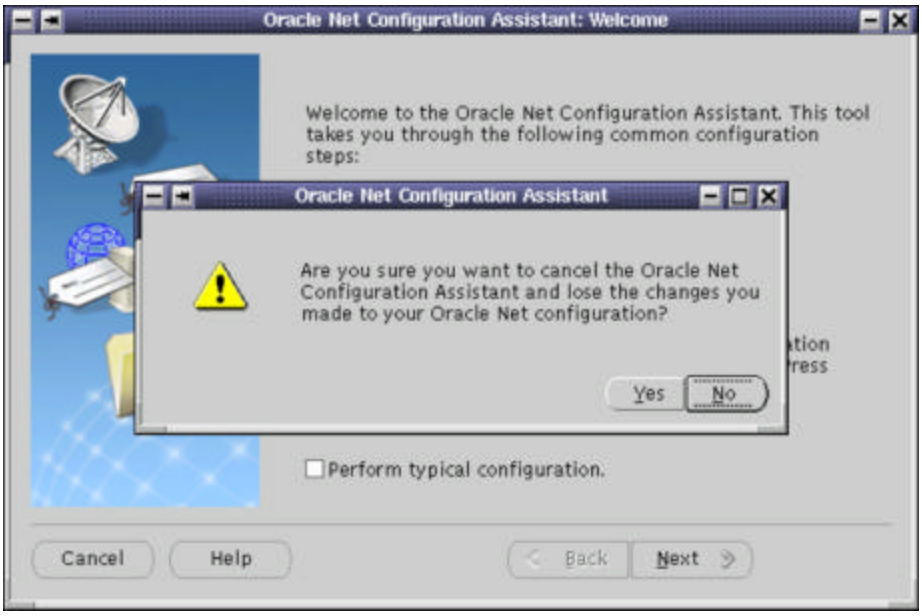

Figure 45. Oracle Net Configuration Assistant window

16. An Error window (Figure 46) is displayed because we cancelled the network configuration. Accept the error by clicking **OK** and **Next** to finish the installation.

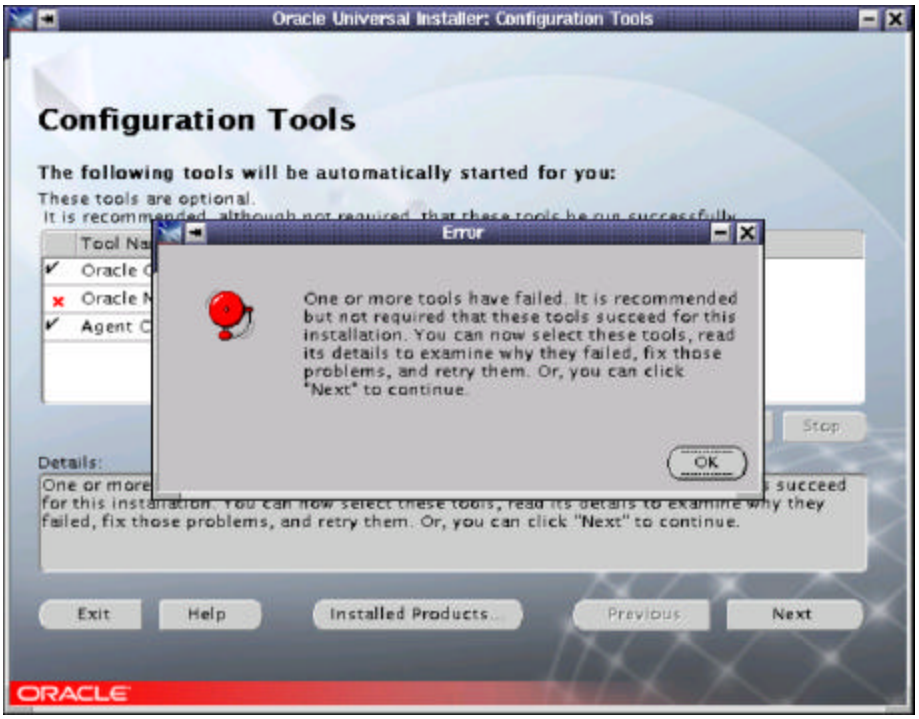

Figure 46. Error window

17. The Oracle Universal Installer displays the End of Installation (Figure 47) window and warns you that you have got some components that were not configured. Since we cancelled network configuration, this is expected, so exit the installer by clicking **Exit**.

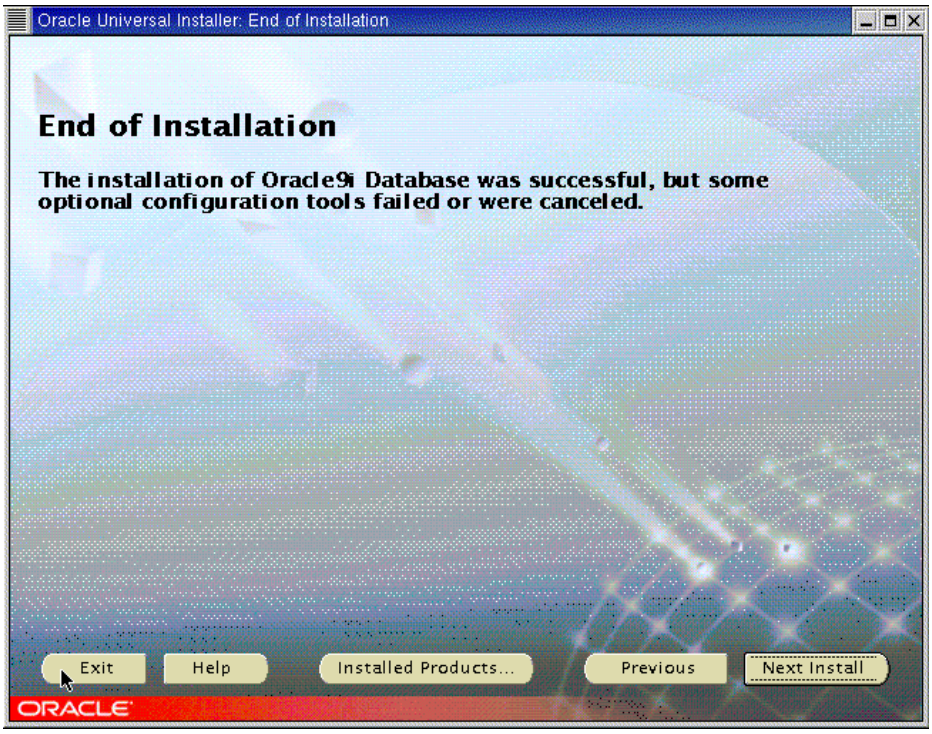

Figure 47. End of Installation window

18. Some of the log directories may not have been created on all nodes during the installation. Make sure then that the following directories exist on each node:
    - `$ORACLE_HOME/rdbms/audit`
    - `$ORACLE_HOME/rdbms/log`
    - `$ORACLE_HOME/network/log`
    - `$ORACLE_HOME/network/trace`

    If they do not exist, please create them in order to have full logs on your system.

19. Start the Global Services Daemon. You should also add it to your Oracle start-up script in order to have it running each time you start your machine. As `oracle` user, type "`gsdctl start`" on each node of your cluster to start the Global Services Daemon.

POST-INSTALLATION

Listeners configuration

1. Before the database creation we will create the listeners that we will use later and which will let us configure our database without problems. Make sure your "LANG" variable is not set by typing "`unset LANG`". Then you can start the Network Configuration Assistant. This tool is the one we stopped during the installation procedure. To start NCA, type "`netca`".

2. The TOPSWelcome window shown in Figure 48 is displayed. Select **Cluster Configuration** and click **Next**.

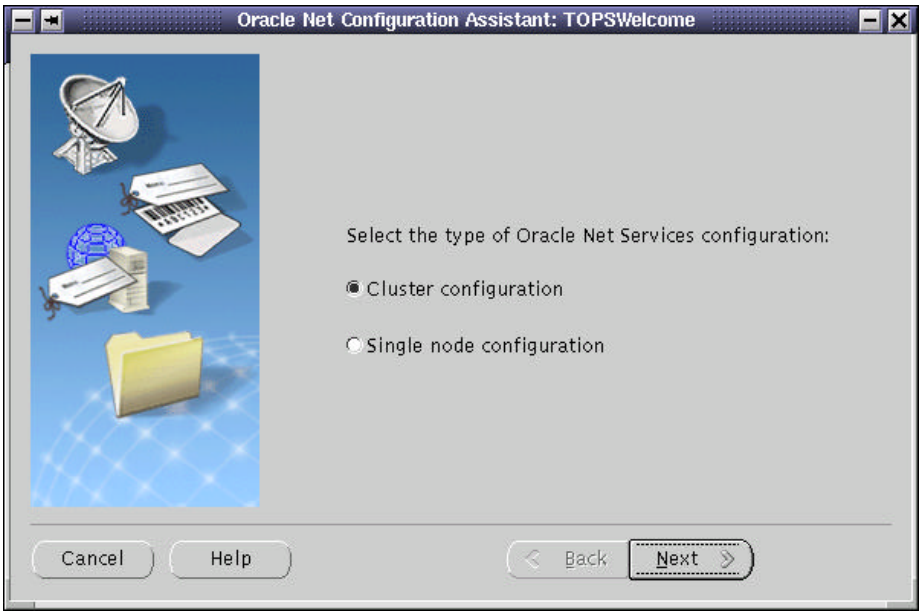

Figure 48. Net Configuration Assistant TOPSWelcome window

3. The TOPSNodes window is displayed (Figure 49). Select all the nodes on the cluster as shown and click **Next**.

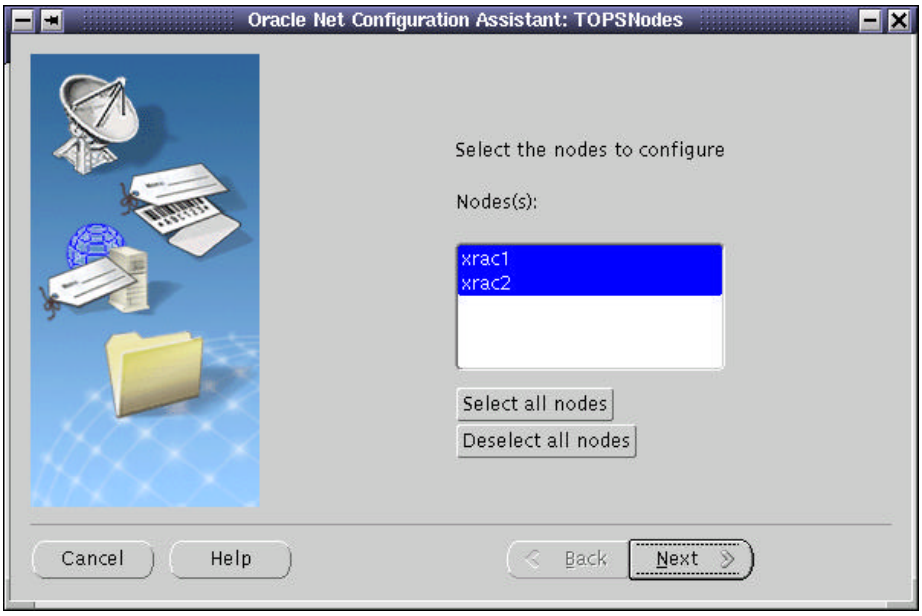

Figure 49. Net Configuration Assistant TOPSNodes window

4. The Welcome window is displayed (Figure 50). Select **Listener Configuration** and click **Next**.

Figure 50. Net Configuration Assistant Welcome window

5. The Listener Configuration window displays (Figure 51. Select **Add** and click **Next** to continue.

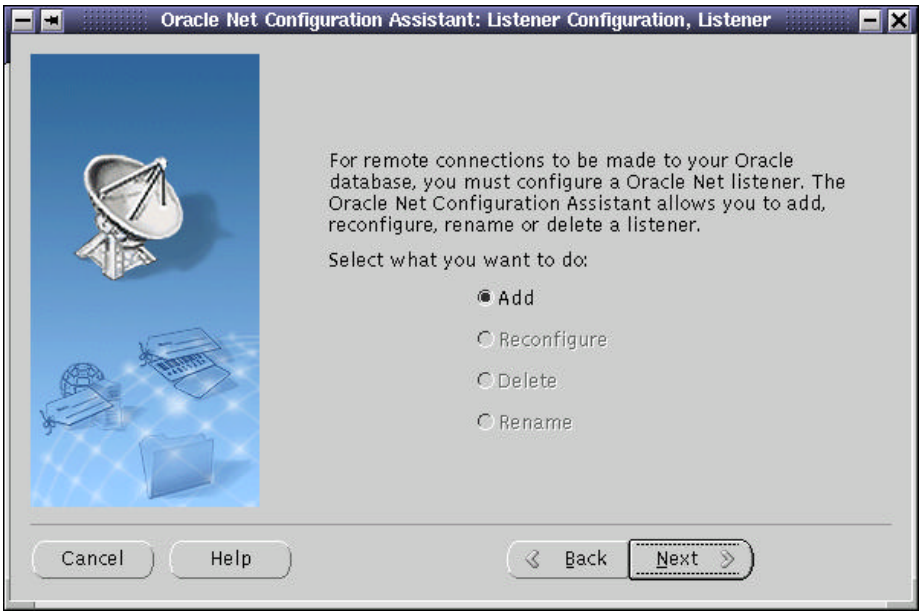

Figure 51. Listener Configuration window

6. Accept the default Listener name shown on the **Listener Name** window (Figure 52). Click **Next** to continue.

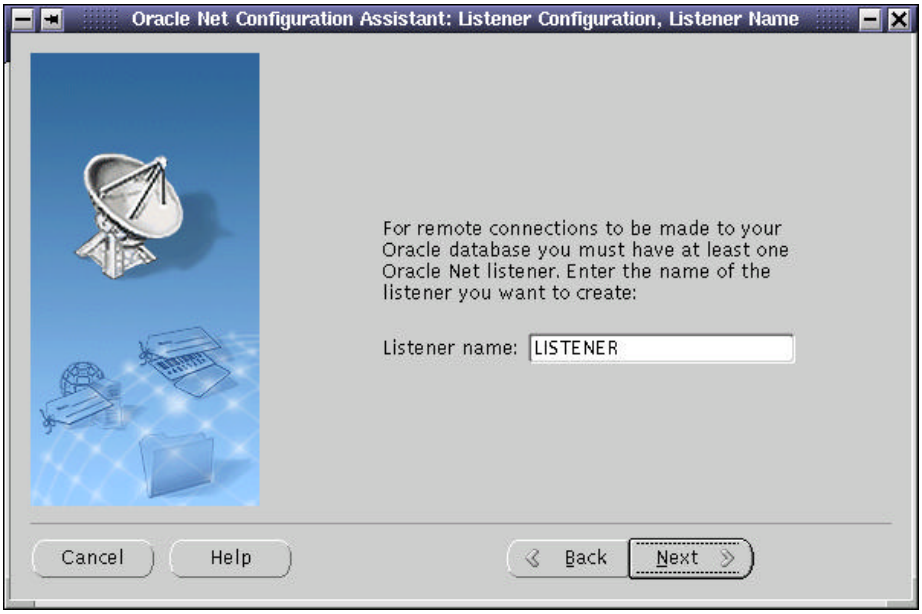

Figure 52. Listener Name window

7. The next window (Figure 53) will let you select network protocols to configure for the listener. TCP/IP should already be selected, but we will also add IPC support. Select **IPC** and click "**>**". Click **Next** to continue.

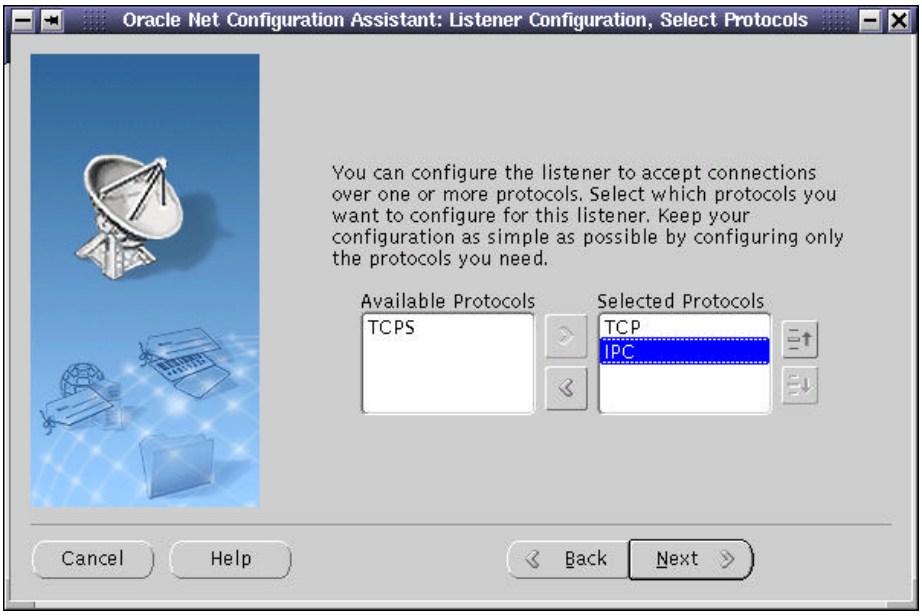

Figure 53. Select Protocols window

8. The next window (Figure 54) will prompt for the TCP/IP port number to be used by the listener. Accept the default port value (set to 1521) and go to the next window by clicking **Next**.

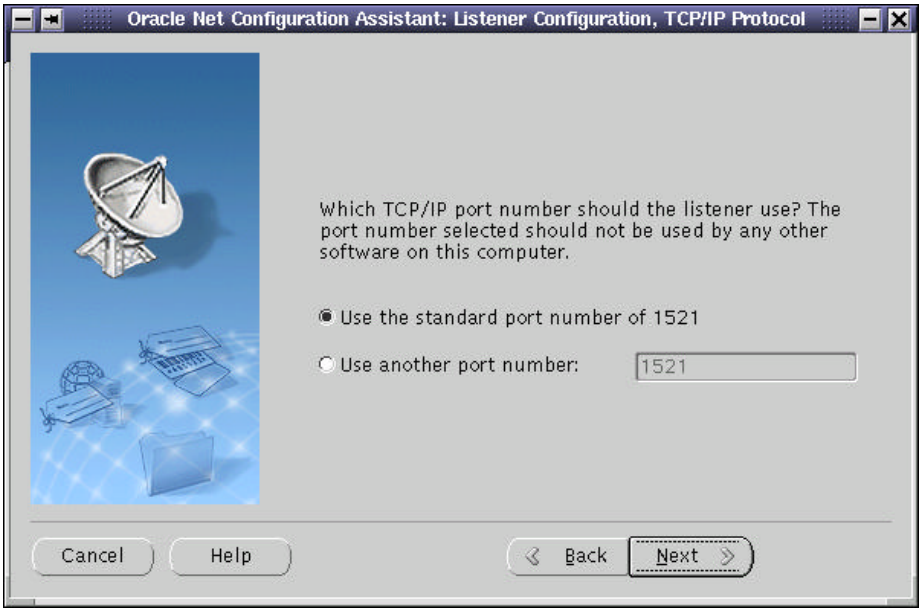

Figure 54. Listener TCP/IP port number window

9. The next window (Figure 55) will ask for the IPC key to be used. Enter your database name as key and then click **Next**. In our example, we used “DB”.

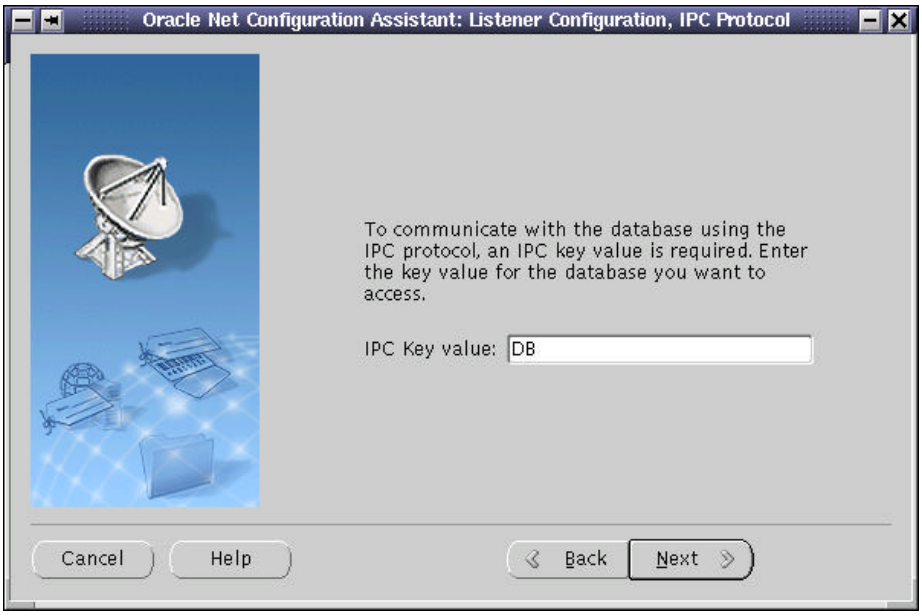

Figure 55. IPC protocol window

10. This should complete listener configuration. Oracle Universal Installer will then ask if you want to create another listener (Figure 56). Select **No** and click **Next**.

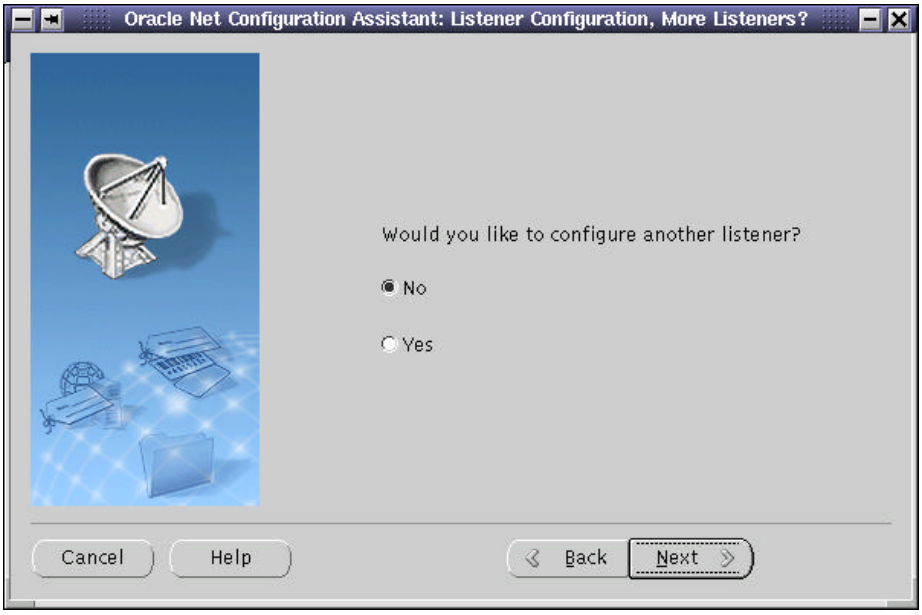

Figure 56. New Listeners window

11. The Listener Configuration is now complete. **Click** Next to exit the Network Configuration Assistant.

Figure 57. Listener Configuration Done window

INSTALL ORACLE9*I* PATCH SET 1

Prior to the end of our project, Oracle released *Oracle9i Release 2 Database Server Patch Set 1 with Cluster Manager Patch for Linux-32.* This patch set is not a complete software distribution and requires an existing Oracle9*i* Release 2 Oracle Server installation.

Patch Set 1 includes patches to the Oracle Cluster Manager on Linux. The watchdogd implementation has been removed and replaced by a Linux kernel module named hangcheck-timer. At the time we wrote this paper, the hangcheck-timer module was only available for Red Hat Advanced Server 2.1. We obtained a development version of the hangcheck-timer module from Oracle for our tests. Oracle recommends loading the hangcheck-timer, however, the oracm process is not prevented from running if the hangcheck-timer is not loaded.

IMPORTANT as of December 16, 2002: If you are installing Oracle9*i* RAC on xSeries 440s with more than 4 CPUs, contact the IBM/Oracle International Competency Center for latest information regarding open issues and latest patches. The IBM/Oracle International Competency Center can be reached by email at: ibmoracl@us.ibm.com.

1. Download Patch Set 1 from Oracle MetaLink. The patch set number is 2632931 and the download file is named p2632931_9202_LINUX.zip.

2. Check MetaLink to see if the hangcheck-time for SLES7 is available for download. The patch number for the Red Hat Advanced Server module is 2594820. If not, contact Oracle Support for assistance.

3. Review the README.html file shipped with Patch Set 1 for notes that apply to your implementation. Follow the detailed procedure for installing Patch Set 1 in the README file.

UPDATE STARTUP SCRIPT

After Patch Set 1 is installed, we can now complete the startup script that was started in Example 9 on page 24. The script shown in Example 11 now initializes kernel parameters, loads the hangcheck timer, starts Oracle Cluster Manager, and Global Services Daemon.

```
#! /bin/sh
#
. /opt/oracle/.bash_profile

case "$1" in
  start)
#
      echo "Modifying kernel parameters"
      export SEMMSL=100
      export SEMMNS=1024
      export SEMOPM=100
      export SEMMNI=100
      echo  $SEMMSL $SEMMNS $SEMOPM $ SEMMNI > /proc/sys/kernel/sem
      export SHMMAX=2147483648
      echo $SHMMAX > /proc/sys/kernel/shmmax

      echo "Loading hangcheck-timer module"
      /sbin/insmod hangcheck-timer hangcheck_tick=30 hangcheck_margin=180

      echo "Starting Oracle Cluster Manager"
      if [ -f /oracle/oracm/log/ocmstart.ts ]; then
         rm /oracle/oracm/log/ocmstart.ts
      fi
      /oracle/oracm/bin/ocmstart.sh

      echo "Starting Global Services Daemon"
      su - oracle -c "gsdctl start"
      ;;

  stop)
      echo "Stopping Global Services Daemon"
      su - oracle -c "gsdctl stop"
      echo "Stopping Oracle Cluster Manager"
      killall oracm
      /sbin/rmmod hangcheck-timer
      ;;

  restart)
      $0 stop
      $0 start
      ;;

  *)
      echo "Usage: $0 {start|stop|restart}"
      exit 1
      ;;
esac
```

Example 11. Oracle startup script - /etc/rc.d/oracle_rac

CREATING A CLUSTER DATABASE

USING DATABASE CONFIGURATION ASSISTANT

For the database creation, we will use the Database Configuration Assistant (DBCA) tool. DBCA is a graphical tool that makes the database creation easier. It uses the Optimal Flexible Architecture (OFA). This means that DBCA creates the database files, including the default server parameter file, using standard file naming and file placement practices. In order to run DBCA, you must have the GSD daemon started. If not, please be sure you run it before starting database creation.

Verify that you have correctly configured all the shared disks for each tablespace (for non-cluster file system platforms) and then you can start database creation.

Oracle Corporation recommends that you use the DBCA to create your database. This is because the DBCA preconfigured databases optimize your environment to take advantage of Oracle9*i* features such as the server parameter file and automatic undo management. The DBCA also enables you to define your own tablespaces as part of the database creation process. So even if you have datafile requirements that differ from those offered in one of the DBCA templates, use the DBCA. You can also execute user-specified scripts as part of the database creation process.

The Database Configuration Assistant (DBCA) makes use of a file that maps database tablespaces to datafiles. Create this file to suit your own tablespaces and make sure it is owned by the oracle user. See Example 12 below for a listing of the file we used in our example, which we named mappingDBCA.cfg.

```
#  DBCA mappings for XRAC
spfile=/oracle/oradata/XRAC/spfile

control1=/oracle/oradata/XRAC/control01.ctl
control2=/oracle/oradata/XRAC/control02.ctl
control3=/oracle/oradata/XRAC/control03.ctl

system=/oracle/oradata/XRAC/system01.dbf
temp=/oracle/oradata/XRAC/temp01.dbf
tools=/oracle/oradata/XRAC/tools01.dbf
users=/oracle/oradata/XRAC/users01.dbf

indx=/oracle/oradata/XRAC/indx01.dbf
redo1_1=/oracle/oradata/XRAC/redo1_1.log
redo1_2=/oracle/oradata/XRAC/redo1_2.log
redo2_1=/oracle/oradata/XRAC/redo2_1.log
redo2_2=/oracle/oradata/XRAC/redo2_2.log

undotbs1=/oracle/oradata/XRAC/undotbs1.log
undotbs2=/oracle/oradata/XRAC/undotbs2.log
```

Example 12. DBCA mapping file – mappingDBCA.cfg

1. The name of this file is passed to DBCA by the variable DBCA_RAW_CONFIG. As oracle user, set this varable by entering:

   ```
   export DBCA_RAW_CONFIG=~/mappingDBCA.cfg
   ```

2. Start DBCA with the `dbca` command. A Welcome screen is displayed as shown in Figure 58. Select **Oracle Cluster Database** and click **Next**.

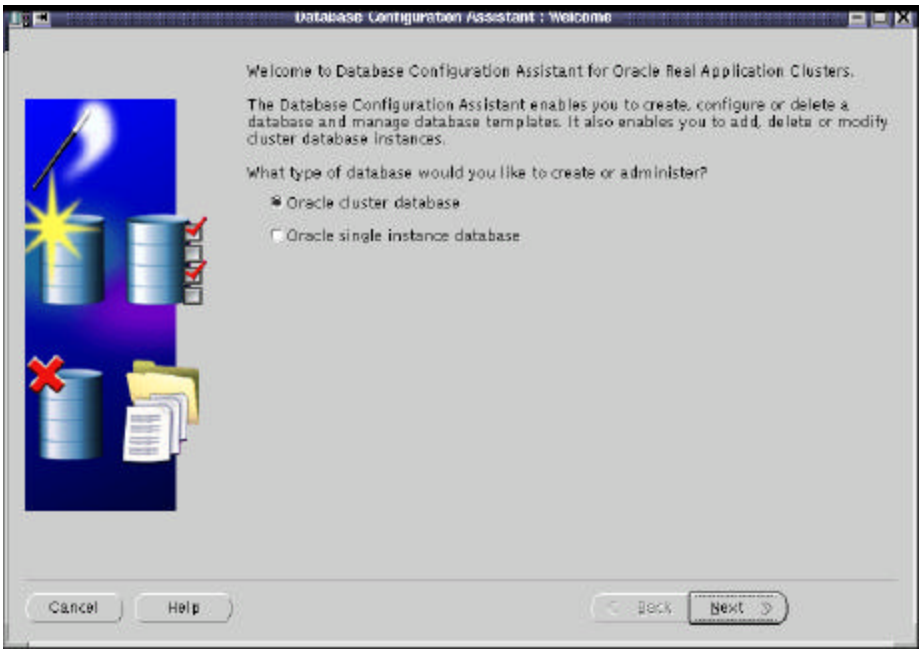

Figure 58.DBCA Welcome window

3. The Operations window (Figure 59) is displayed. Select **Create Database** and click **Next**.

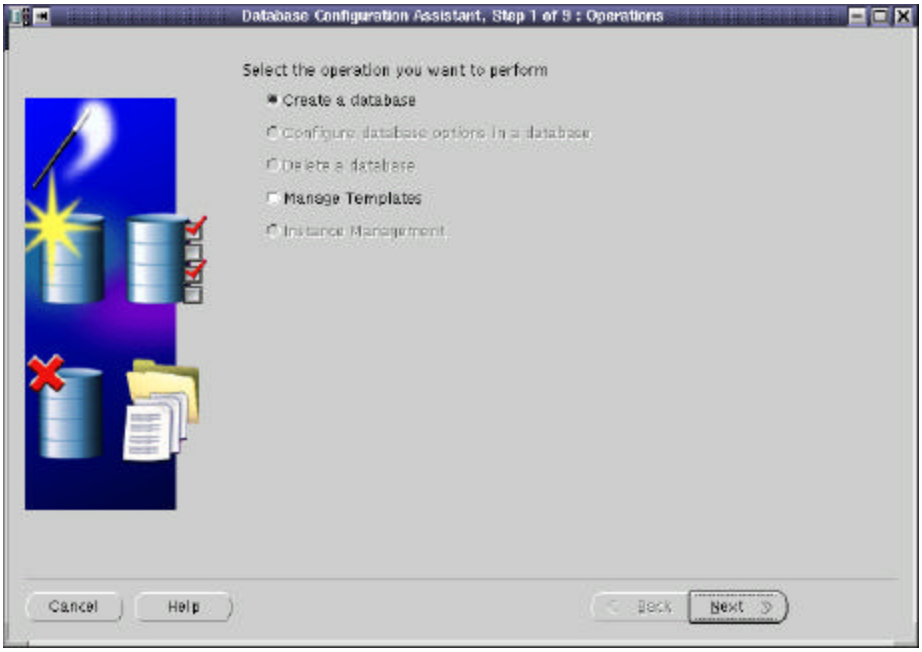

Figure 59. Operations window

4. The Node Selection window is displayed (Figure 60). Select all the nodes that apply and click **Next** to continue. In our example, we selected both xrac1 and xrac2 as shown.

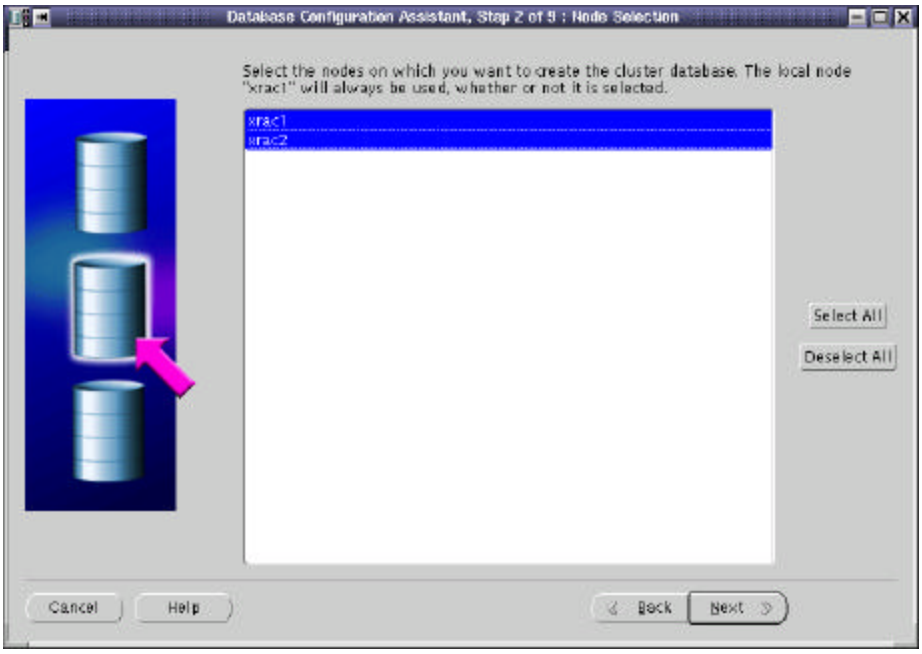

Figure 60. Node Selection window

If some of your nodes are missing, there is a cluster diagnostic tool called `lsnodes` which will let you know what nodes are currently active. Resolve the problem and try again the process from the beginning. In our tests, we used two nodes: "xrac1" and "xrac2", and "lsnodes" shows both of them (Example 13).

```
oracle@xrac1:~ > lsnodes
xrac1
xrac2
oracle@xrac1:~ >
```

Example 13. lsnodes command

5. The next window is the Database Templates window (Figure 61). All of the templates, other than New Database, include datafiles. Select **New Database** and click **Next**.

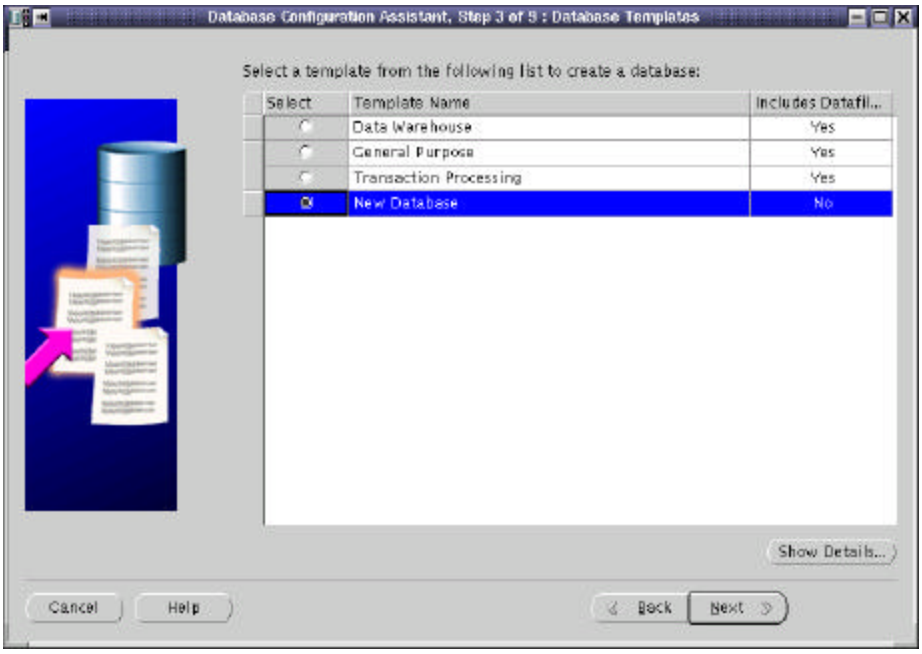

Figure 61. Database Templates window

6. This next window is the Database Identification window (Figure 62). You should enter the Global Database Name (typically on the name.domain form) and the Oracle System Identifier (SID). The SID will identify an instance and DBCA will suggest you use the same name that you entered on the Database Name field. The SID identifier will be used on RAC databases as a prefix for the instance number. In our example, we used XRAC for the Global Database Name and took the suggested SID Prefix. Click **Next** to continue.

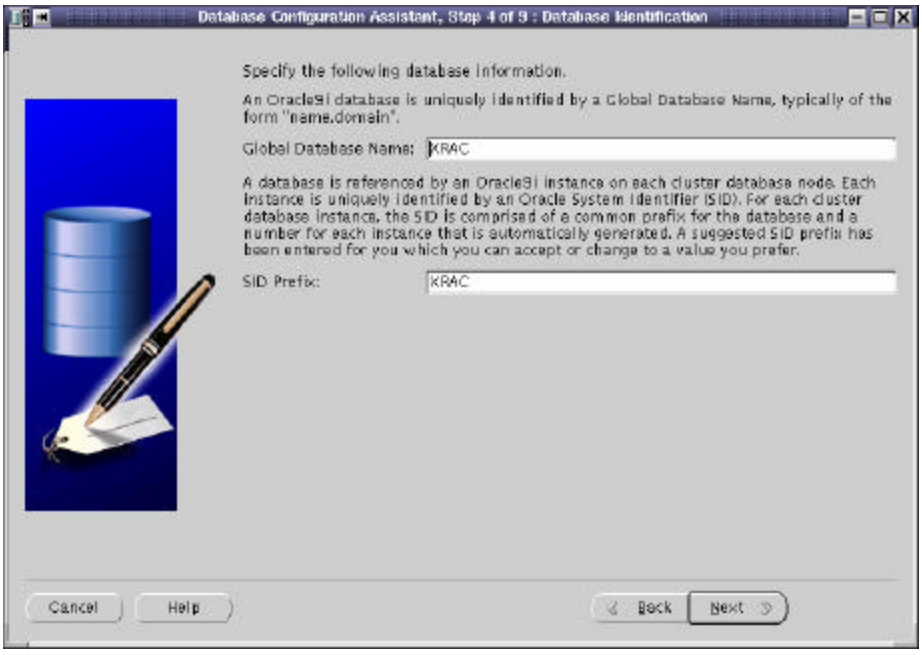

Figure 62. Database Identification window

7. The Database Features window displays (Figure 63). Select the appropriate database features. In our example, we deselected every feature.

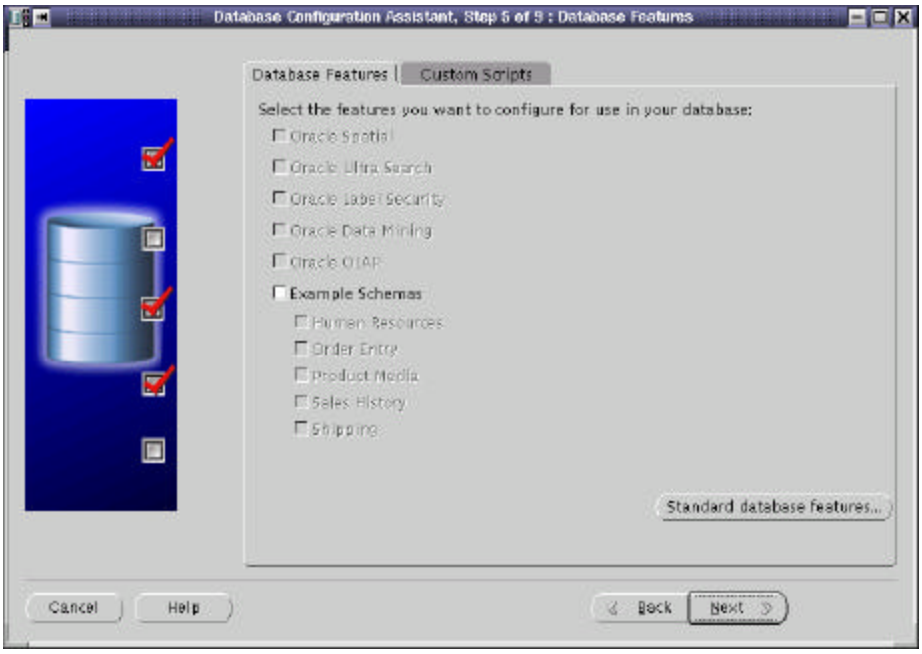

Figure 63. Database Features window

8. Click on **Standard database features** to add or remove some standard options. The Standard database features window shown in Figure 64 displays. In our example, we deselected every standard option as shown. Click **OK** to continue.

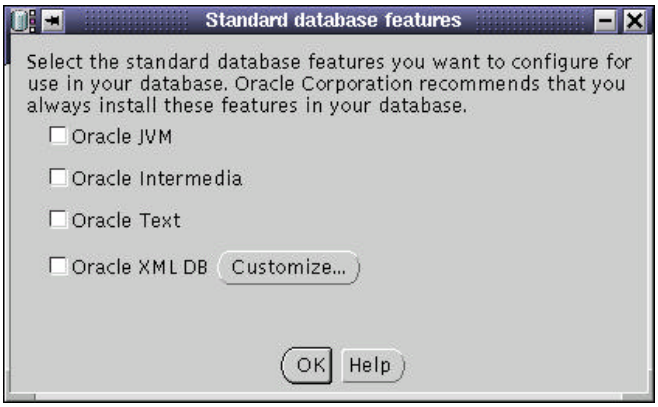

Figure 64. Standard Database Features window

9. The Database Connection Options window displays (Figure 65). Select **Dedicated Server Mode** and click **Next** to continue.

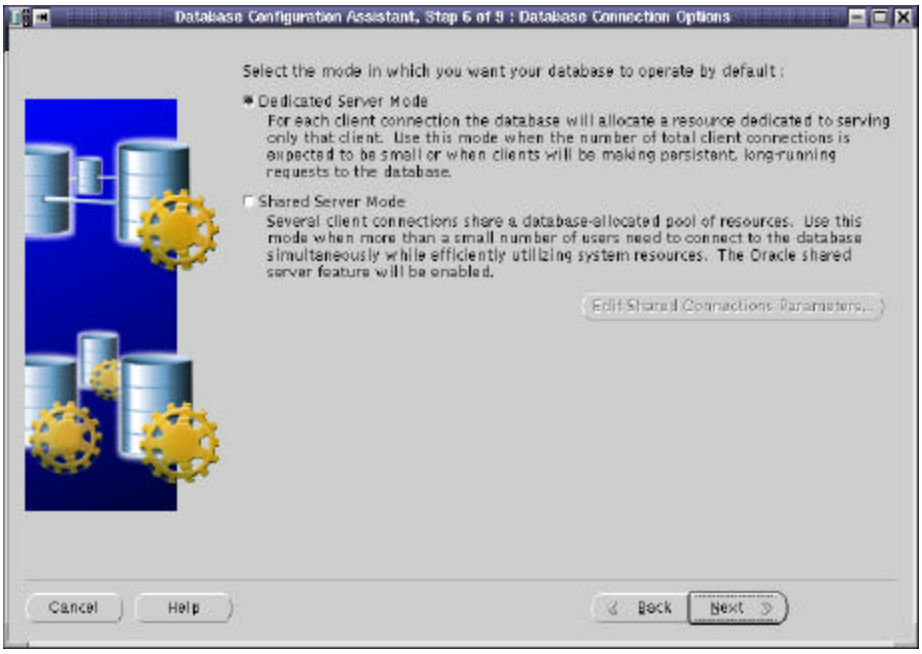

Figure 65. Database Connection Options window

10. The next window will let you configure the Initialization Parameters. On this window you will see multiple page tabs. The first tab is the Memory page, shown in Figure 66. Modify the memory setting if desired.

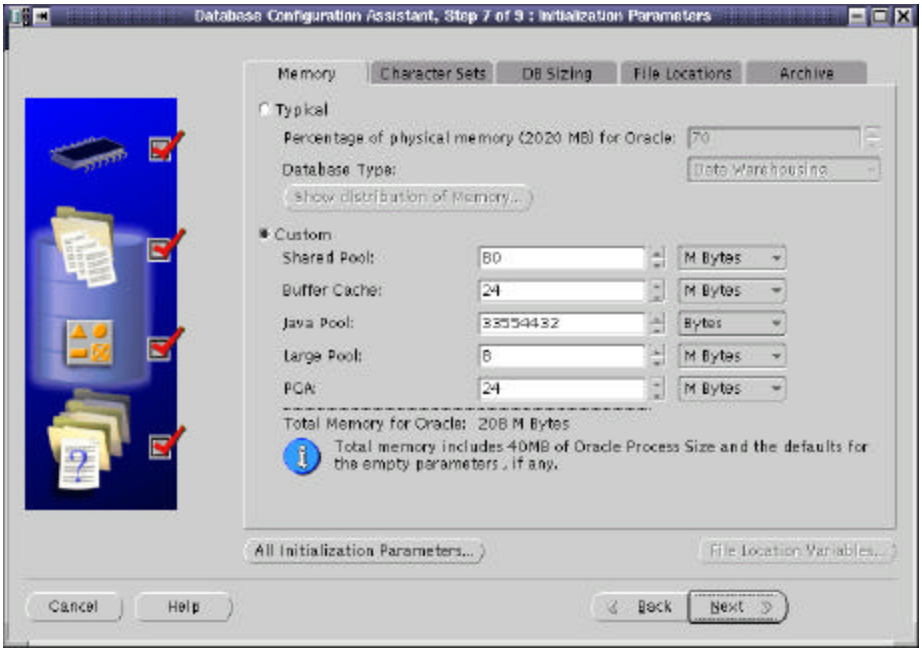

Figure 66. Initialization Parameters window - Memory tab

11. Click the **File Locations** tab. The File Locations page is displayed in Figure 67. Select the "File Locations" tab. The option "Create Persistent Initialization" is selected by default. Server Parameters Filename should point to the correct datafile. If not, either DBCA_RAW_CONFIG is not set correctly or the file contains errors.

    Click on the othe tabs and verify all the configuration options according to your requirements. When this is done, click **Next** to continue.

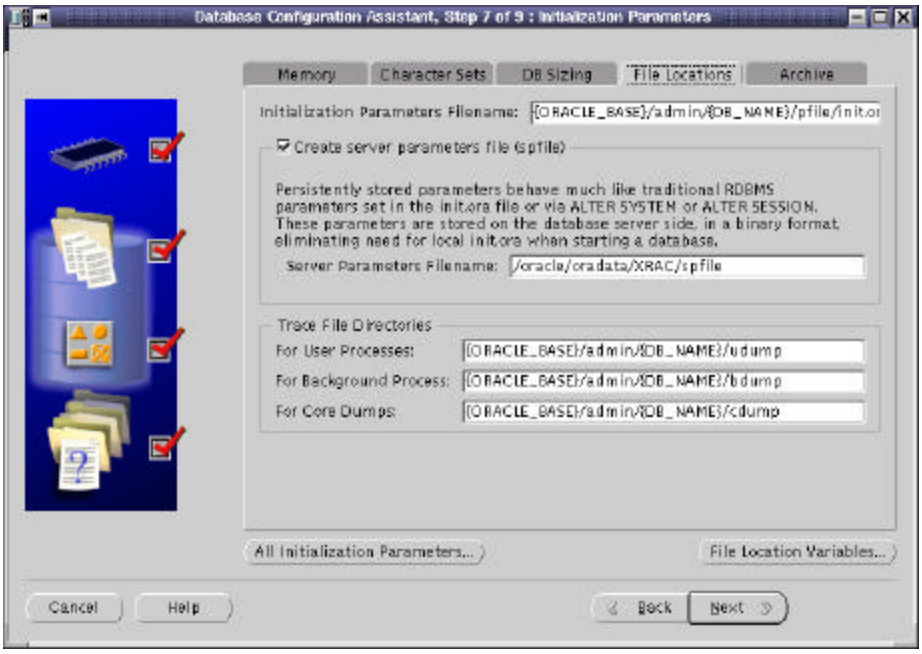

Figure 67. Initialization Parameters window – File Locations tab

12. DBCA now shows you the Database Storage Window (Figure 68). The tablespaces and datafiles should be as specified in the mapping file (mappingDBCA in Example 12). Verify the settings and click **Next** when you're finished.

Figure 68. Database Storage window

13. The Database Creation Options window is displayed. Ensure that the option **Create Database** is checked and click **Finish**.

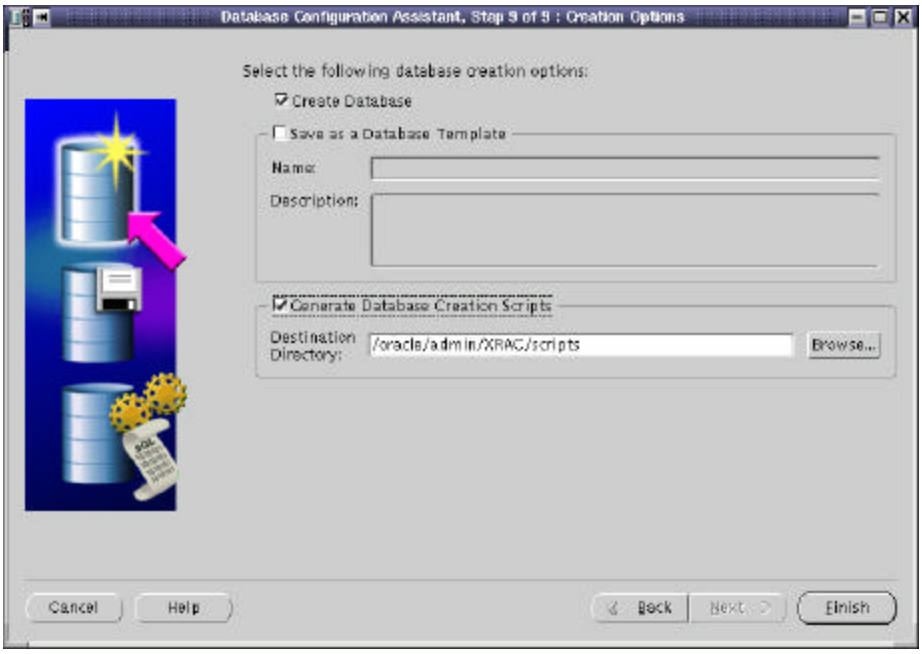

Figure 69. Creation Options window

14. The DBCA Summary window (Figure 70) is displayed. Using the scroll bar, review all information and then click **OK** to continue.

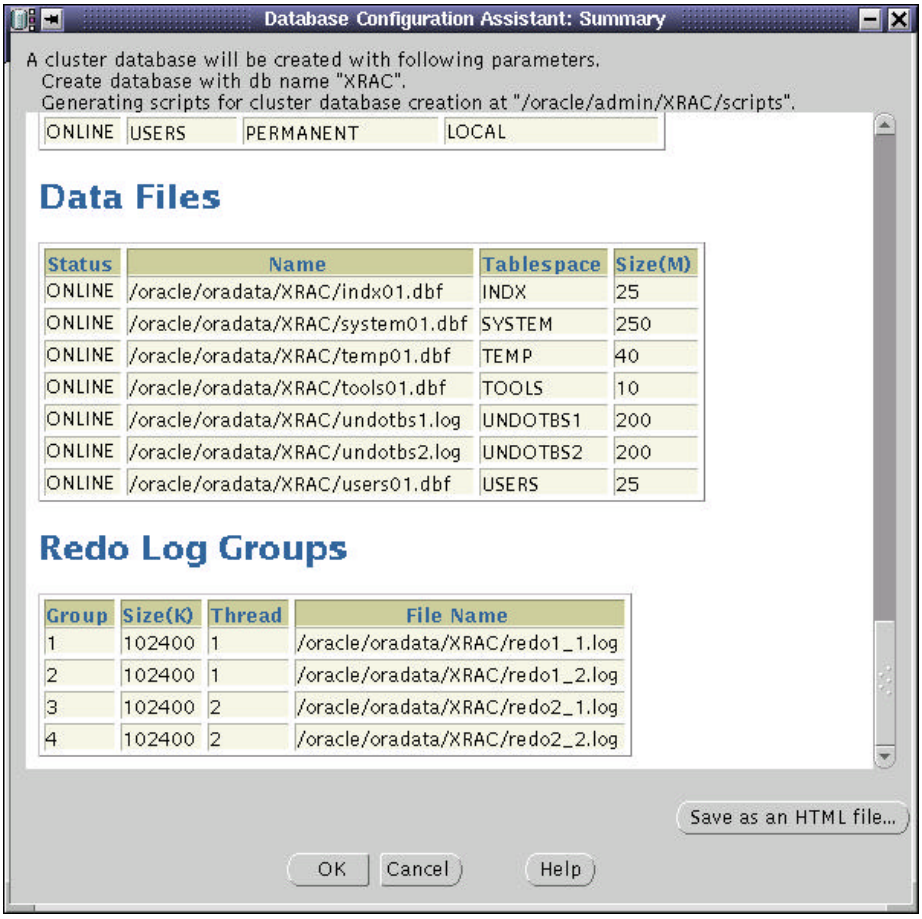

Figure 70. DBCA Summary window

15. DBCA begins to create the database. Since we selected Generate Database Creation Scripts (Figure 69), a pop-up window will display to indicate successful script generation (Figure 71).

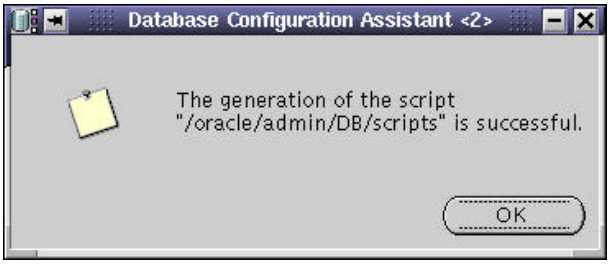

Figure 71. Pop-up window - successful script generation

16. DBCA starts the database creation. A progress window is displayed (Figure 72).

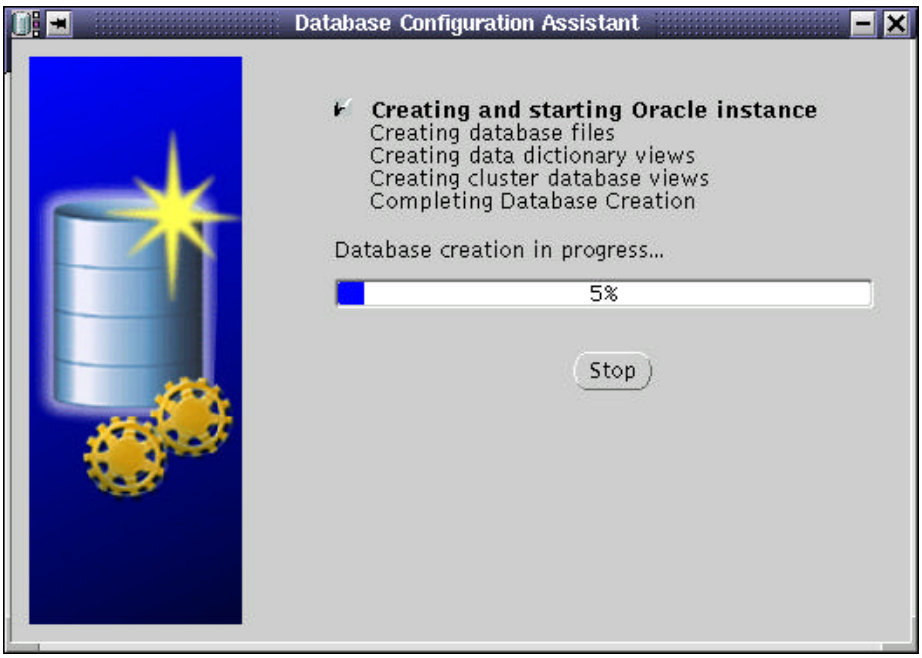

Figure 72. DB creation progress window

17. When DBCA has finished database creation, the completion window shown in Figure 73 is displayed. DBCA prompts you to change the passwords for SYS and SYSTEM database users. Enter the desired passwords and click **Exit**.

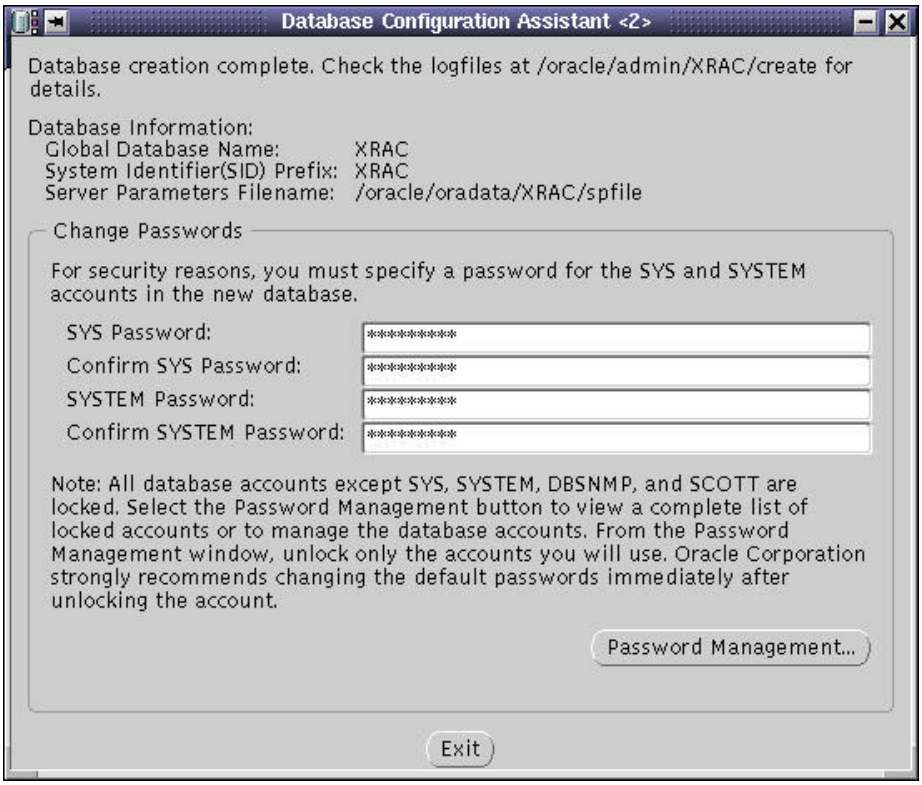

Figure 73. Database creation complete window

18. Once this step finished, you have successfully created the clustered database. It can now be accessed by Oracle SQL*PLUS or other applications designed to work with an Oracle RAC database. You should make sure to execute start-up scripts during system startup as root and other scripts or commands as oracle user as described earlier in this document.

CREATING A DATABASE MANUALLY

You can also create the databases manually. To do so, follow the next steps:

- on one node : **srvctl add db -p <db_name> -o <oracle_home>**
- for each instance of the database :

  srvctl add instance –p <db_name> -i <SID> -n <node_name>

  (it is advised to set the SID to db_name plus instance_number)
- on each node, check the configuration : **srvctl config –p <db_name>**
- on each node, create (or update) oratab file (in /etc directory) with the following line :

  <db_name>:<$ORACLE_HOME>:N
- create the udump, cdump and bdump directories.
- set the SID in .profile of oracle user

- create the init.ora file for each node. It is possible to prefix parameter which are local to the instance with the instance name. For example, thread, instance_name, rollback_segments…
  It is possible to create a spfile with "**create spfile='/dev/rawspfile' from pfile=/…/init.ora**".

- create the database creation script. You can use $ORACLE_HOME/srvm/clustdb.sql script as sample. The script must be adapted to your environment. Be careful, there are some errors in this file.

- you can create a password file. Under $ORACLE_HOME execute **orapwd file=orapw password=###.**

- as sysdba, execute the script.

For further information, refer to Step-by-step Installation of RAC on Linux, Oracle Metalink Note184821.1.

UPDATE ORACLE USER PROFILE

The database has been created and each node has its own instance. Update each node's .bash_profile (created in Example 10 on page 25) for the `oracle` user with the correct ORACLE_SID value. The instance name for a node is the SID prefix concatenated with the instance number. For example, the instance on xrac1 will be XRAC1 and the instance on xrac2 will be XRAC2.

For example, in xrac1, we added the line

```
export ORACLE_SID=XRAC1
```

to oracle user's .bash_profile on xrac1. Source the new .bash_profile in your current shell by issuing the command:

```
. ~/.bash_profile
```

CONFIGURING HIGH AVAILABILITY AND LOAD BALANCING

LISTENER CONFIGURATION

The listener should be properly configured by **netca** and **dbca** but it is worth checking that the contents of $ORACLE_HOME/network/admin/listener.ora are similar to those in the following listener.ora for the first instance:

```
LISTENER =
(DESCRIPTION_LIST =
   (DESCRIPTION =
      (ADDRESS_LIST =
         (ADDRESS =(PROTOCOL =TCP)(HOST =xrac1)(PORT =1521))
      )
   )
)
SID_LIST_LISTENER =
(SID_LIST =
   (SID_DESC =
      (SID_NAME =PLSExtProc)
      (ORACLE_HOME =/oracle)
      (PROGRAM =extproc)
   )
   (SID_DESC =
      (ORACLE_HOME =/oracle)
      (SID_NAME = XRAC1)
   )
)
```

Example 14. listener.ora for the first instance

TNSNAMES CONFIGURATION

```
LISTENERS_XRAC =
   (DESCRIPTION =
      (ADDRESS_LIST =
         (ADDRESS =(PROTOCOL =TCP)(HOST =xrac1)(PORT =1521))
         (ADDRESS =(PROTOCOL =TCP)(HOST =xrac2)(PORT =1521))
      )
   )
XRAC =
   (DESCRIPTION =
      (ADDRESS_LIST =
         (ADDRESS =(PROTOCOL =TCP)(HOST =xrac1)(PORT =1521))
         (ADDRESS =(PROTOCOL =TCP)(HOST =xrac2)(PORT =1521))
      )
      (FAILOVER =yes)
      (LOAD_BALANCE =yes)
      (CONNECT_DATA =
         (SERVICE_NAME =XRAC)
         (FAILOVER_MODE=
         (TYPE=SELECT)
         (METHOD=BASIC)
         (RETRIES=30)
         (DELAY=1)
      )
   )
)
XRAC1 =
   (DESCRIPTION =
      (ADDRESS_LIST =
         (ADDRESS =(PROTOCOL =TCP)(HOST =xrac1)(PORT =1521))
      )
      (CONNECT_DATA =
         (SERVICE_NAME =XRAC)
         (INSTANCE_NAME =XRAC1)
      )
   )
XRAC2 =
   (DESCRIPTION =
      (ADDRESS_LIST =
         (ADDRESS =(PROTOCOL =TCP)(HOST =xrac2)(PORT =1521))
      )
      (CONNECT_DATA =
         (SERVICE_NAME =XRAC)
         (INSTANCE_NAME =XRAC2)
      )
   )
```

Example 15. Extract from tnsnames.ora

ADMINISTERING REAL APPLICATION CLUSTERS INSTANCES

Oracle Corporation recommends that you use SRVCTL to administer your Real Application Clusters Database environment. SRVCTL manages configuration information that is used by several Oracle tools.

Before using SRVCTL, ensure that your Global Services Daemon (GSD) is running. To use SRVCTL, you must have already created the configuration information for the database that you want to administer.

You must have done this either by using the Oracle Database Configuration Assistant (DBCA), or by using the `srvctl add` command. If you have followed the instructions in this document, `dbca` will have added your database and instances.

```
$ srvctl config database -d XRAC
xrac1 XRAC1 /oracle
xrac2 XRAC2 /oracle
```

Examples of starting and stopping RAC follow:-

```
$ srvctl start database -d XRAC
$ srvctl status database -d XRAC
Instance XRAC1 is running on xrac1
Instance XRAC2 is running on xrac2

$ srvctl stop database -d XRAC
$ srvctl status database -d XRAC
Instance XRAC1 is not running on xrac1
Instance XRAC2 is not running on xrac2

$ srvctl start instance -d XRAC -i XRAC1
$ srvctl status instance -d XRAC -i XRAC1
Instance XRAC1 is running on xrac1

$ srvctl status database -d XRAC
Instance XRAC1 is running on xrac1
Instance XRAC2 is not running on xrac2

$ srvctl stop instance -d XRAC -i XRAC1
```

For further information on `srvctl` see the Oracle9*i* Real Application Clusters Administration Release 1 (9.0.1) manual.

REFERENCES

- Oracle9i Installation Guide Release 2 for UNIX Systems: AIX-Based Systems, Compaq Tru64 UNIX, HP 9000 Series HP-UX, Linux Intel, and Sun Solaris, Part No. A96167-01
- Oracle9i Release Notes Release 2 (9.2.0.1.0) for Linux Intel (32-bit), Part No. A97349-02
- Oracle9*i* Real Application Clusters Setup and Configuration Release 2 (9.2), March 2002, Part No. A96600-01
- Step-by-step Installation of RAC on Linux, Oracle Metalink Note184821.1.
- Installing Oracle9i RAC on IBM x360 with FAStT700 running RedHat AS 2.1, Version 1.0, September 2002, EMEA Oracle/IBM Joint Solutions Center

DISCLAIMER

This document reflects the IBM / Oracle International Competency Center's understanding about Oracle products running on IBM hardware servers and storage. It was produced and reviewed by the members of the IBM organization called the IBM / Oracle International Competency Center and others.

This document is presented "As-Is" and IBM does not assume responsibility for the statements expressed herein. It reflects the experiences of the IBM / Oracle International Competency Center. If you have questions about the contents of this document, please direct them to the IBM / Oracle International Competency Center (ibmoracl@us.ibm.com).

TRADEMARKS

All brand names are the trademarks of their respective owners.

SPECIAL NOTICE