

Installing Oracle9i RAC on IBM *@server* x360 with FAS*T*T700 running RedHat AS 2.1

Version 1.1b
February 28, 2002

Authors :

Jose Rodriguez Ruibal – Oracle/IBM Joint Solutions Center
Gregory Verstraeten – Oracle/IBM Joint Solutions Center

Many thanks to :

...
Oracle/IBM JSC team
IBM/Oracle ICC team

EMEA Oracle/IBM
Joint Solutions Center
MONTPELLIER - FRANCE



Installing Oracle 9.2 RAC on x360 and FAStT700 with RedHat Advanced Server 2.1

The aim of this document	3
Planning the installation	4
Hardware cabling and Firmware	5
Hardware configuration	5
Cabling	5
Firmware	5
UPDATE: February 28, 2003	6
x360 Bios Level and Red Hat Kernel.....	6
Oracle9i Patch Set 1.....	6
Operating System setup.....	7
Configuring the storage.....	11
Software installation	11
Storage usage recomandations for Oracle datafiles	11
Storage configuration	12
Installing the Cluster Manager	13
Oracle9i Database Installation	23
Creating a Cluster Database.....	38
Creating a Database Manually.....	47
Configuring High Availability and Load Balancing.....	48
listener configuration	48
tnsnames configuration	48
Administering Real Application Clusters Instances.....	50
References.....	51
Disclaimer	52
Trademarks	52
Special Notice.....	52

THE AIM OF THIS DOCUMENT

The aim of this document is to provide help to install Oracle9i Real Application Clusters on IBM eServer x360 with FAStT700 running Red Hat Advanced Server 2.1 through a description of the different steps necessary to run Oracle9i RAC. In this document, we take into account the specifics of the IBM xSeries hardware, and give you the necessary information to successfully install Oracle9i RAC on IBM xSeries.

The information contained in this paper resulted from:

- Oracle and IBM documentation
- Installation runs of Oracle9i RAC
- Certification work done on this cluster
- Workshop experiences done in the 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>).

- Oracle9i Installation Guide Release 9.2.0.1 for UNIX Systems
- Oracle9i Administrator's Reference 9.2.0.1 for UNIX Systems
- Oracle9i Release Notes Release 2 (9.2.0.1.0) for Linux Intel
- Oracle9i Online Generic Documentation CD-ROM Installation and Usage Notes
- Oracle9i Real Application Clusters Installation and Configuration
- Oracle Enterprise Manager Configuration Guide

IMPORTANT:

Once the hardware is configured and the environment ready, the installation of the Oracle product is identical on all Linux platforms. We have based this installation procedure on the document "Step-By-Step Installation of RAC on Linux" from Metalink. We have illustrated it with the screen shots from our labs.

Sections from "Hardware cabling and Firmware" to "Configuring the storage" are xSeries and RedHat specifics.

Sections from "Installing the Cluster Manager" to "Creating a Cluster Database" come from "Step-By-Step Installation of RAC on Linux" from Metalink.

Your comments are important for us. We want our technical papers to be as helpful as possible. Please send us your comments about this document to the Oracle/IBM Joint Solutions Center.

Use our email address :

oraclibm@fr.ibm.com

or our phone number :

+33 4 67 34 67 49

PLANNING THE INSTALLATION

Get the RedHat Advanced Server 2.1 CDs

Get the IBM ServeRAID 4.84 Support CD

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

- [IBM eServer xSeries 360 - Flash BIOS Update version 1.03](#)
- [IBM eServer xSeries 360 - Diagnostic Flash Update version 3.00](#)
- [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 documentation from <http://www.pc.ibm.com>:

- [IBM eServer xSeries 360 - Installing Red Hat Advanced Server Linux v2.1](#)

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

Get Oracle9i R2 Enterprise Edition CDs (they can also be downloaded from <http://technet.oracle.com>)

HARDWARE CABLING AND FIRMWARE


HARDWARE CONFIGURATION

Check the hardware configuration with the xSeries Oracle9i RAC Configuration Guide.

CABLING

The chart below shows a 2 nodes configuration.

Also, there are redundant SAN switches (2109 S16). The redundancy is not mandatory but it increases the availability of the cluster (check the "Implementing IBM eServer xSeries SANs" Redbook for more information on SAN configurations).



FIRMWARE

Here are the minimum levels of firmware/BIOS for the different hardware parts (Check the latest supported level on <http://www.pc.ibm.com>):

- x360 BIOS: 1.03
- ServeRaid BIOS: 4.84
- FAST FC-2 HBA BIOS: 1.76
- FAST700 firmware: 8.21

Check the readme of the downloaded files for the upgrade process. See the "PLANNING THE INSTALLATION" section for the downloads.

Also check the [Implementing IBM eServer xSeries SANs](#) Redbook for more information on the FAST700 firmware upgrade.

UPDATE: FEBRUARY 28, 2003

Since we performed our installation testing for this paper, a number of notable changes have occurred that affect the installation procedure detailed in the remainder of this paper. We have not gone back to perform installation testing to determine the detailed impact on our procedure, but we will list the changes here for your information.

X360 BIOS LEVEL AND RED HAT KERNEL

We performed our installation testing using BIOS level 1.03, which was the then-current BIOS level. The current BIOS level for the xSeries 360 is 1.06. If you have updated the BIOS code, you will need to run the Red Hat Enterprise kernel at 2.4.9-e.8 or later. Do not use the Summit kernel as described in "Operating System setup". Refer to the following url:

http://www.pc.ibm.com/qtechinfo/MIGR-43850.html?lang=en_NZ&page=brand&brand=IBM+PC+Server&doctype=&subtype=Cat

The current version of the Enterprise kernel is 2.4.9-e.12. Obtain the errata package named kernel-enterprise-2.4.9-e.12.i686.rpm and install using the `rpm -U` command.

To recompile the FASTT adapter driver, you will need the following packages:

- kernel-source-2.4.9-e.12.i386.rpm
- kernel-headers-2.4.9-e.12.i386.rpm
- glibc-devel-2.2.4-31.7.i386.rpm
- gcc-2.96-116.7.2.i386.rpm
- cpp-2.96-116.7.2.i386.rpm

These packages may be downloaded from <http://rhn.redhat.com>

ORACLE9/PATCH SET 1

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 Oracle9i 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. The hangcheck-timer module is available in Red Hat Advanced Server kernel errata e.12. If you installed the packages listed above, the hangcheck-timer module should be on your system.

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

OPERATING SYSTEM SETUP

This chapter will guide you through the steps for the post-installation of RedHat Advanced Server to take in account the specifics of the IBM xSeries cluster.

After you have installed RedHat Advanced Server 2.1 on all nodes (check the installation procedure "IBM eServer xSeries 360 - Installing Red Hat Advanced Server Linux v2.1" from <http://www.pc.ibm.com>), follow these steps on each nodes in the cluster:

1. Install the Summit kernel:
mount the 1st CDROM of RHAS, as root run:
`rpm -U /mnt/cdrom/RedHat/RPMS/kernel-summit-2.4.9-e.3.i686.rpm`
2. Change the Grub configuration to have the Summit kernel loaded as the default:
edit `/boot/grub/grub.conf`, the Summit kernel is the first kernel entry in the file, the entry default has to be set to **0**.

```
default=0
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz
title Red Hat Linux (2.4.9-e.3summit)
    root (hd0,0)
    kernel /vmlinuz-2.4.9-e.3summit ro root=/dev/sda3
    initrd /initrd-2.4.9-e.3summit.img
title Red Hat Linux Advanced Server (2.4.9-e.3enterprise)
    root (hd0,0)
    kernel /vmlinuz-2.4.9-e.3enterprise ro root=/dev/sda3
    initrd /initrd-2.4.9-e.3enterprise.img
```

grub.conf

If you use LILO and not GRUB as the boot loader, edit `/etc/lilo.conf`, and as root run `lilo` afterwards.

3. Change the modules order of the devices:
edit `/etc/modules.conf`. The ServeRaid adapter driver, `ipc`, has to be the first scsi module to be loaded, it is mapped to `SCSIXXXX`. The FASTT adapter comes after, it is mapped to `SCSIXXX1`, the others are mapped to `SCSIXXX2`, `SCSIXXX3`, etc.
append the following line to enable up to 32 SCSI devices per adapter:

```
options scsi_mod max_scsi_luns=32
alias parport_lowlevel parport_pc
alias eth0 eepro100
alias eth1 e1000
alias scsi_hostadapter ips
alias scsi_hostadapter1 qla2300
alias scsi_hostadapter2 qla2300
alias usb-controller usb-ohci
options scsi_mod max_scsi_luns=32
```

modules.conf

4. Configure the network:
Stop the network service with the command: `service network stop`
The motherboard Ethernet device is used for the public LAN. To configure this device, edit `/etc/sysconfig/network-scripts/ifcfg-eth0`.

```
DEVICE=eth0
BOOTPROTO=static
ONBOOT=yes
IPADDR=10.10.180.1
NETMASK=255.255.255.0
```

/etc/sysconfig/network-scripts/ifcfg-eth0

The Fiber Channel Gigabit Ethernet device is used for the private LAN. To configure this device, edit `/etc/sysconfig/network-scripts/ifcfg-eth1`.

```
DEVICE=eth1
ONBOOT=yes
BOOTPROTO=static
IPADDR=1.1.1.1
NETMASK=255.255.255.0
```

```
/etc/sysconfig/network-scripts/ifcfg-eth1
```

5. Install the necessary packages to recompile the FASTt adapter driver.

Mount the 1st CDROM of RHAS:

```
rpm -U /mnt/cdrom/RedHat/RPMS/kernel-header-2.4.9-e.3.i386.rpm
rpm -U /mnt/cdrom/RedHat/RPMS/binutils-2.11.90.0.8-12.i386.rpm
```

Mount the 2nd CDROM of RHAS:

```
rpm -U /mnt/cdrom/RedHat/RPMS/glibc-devel-2.2.4-26.i386.rpm
rpm -U /mnt/cdrom/RedHat/RPMS/cpp-2.96-108.1.i386.rpm
rpm -U /mnt/cdrom/RedHat/RPMS/gcc-2.96-108.1.i386.rpm
rpm -U /mnt/cdrom/RedHat/RPMS/kernel-source-2.4.9-e.3.i386.rpm
```

6. Recompile the FASTt adapter driver.

```
cd /usr/src/linux-2.4
cp ./configs/kernel-2.4.9-i686-summit.config .config
make dep
cd /root
```

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

```
tar xvfz 0103.tgz
```

The iXXXX will be created:

```
cd iXXXX
./libinstall
./drvsetup
make all SMP=1
```

Replace the existing drivers:

```
cp qla*.o /lib/modules/2.4.9-e.3summit/kernel/drivers/addon/qla2200/
```

Recreate the initial RAM disk:

```
rm /boot/initrd-2.4.9-e.3summit.img
mkinitrd /boot/initrd-2.4.9-e.3summit.img 2.4.9-e.3summit
```

7. Reboot the server.

8. Configure the Hosts file.

Edit /etc/hosts. The hosts name are mapped to the public LAN addresses. The int-<hosts name> are mapped to the private LAN addresses.

```
x360r1
x360r2
int-x360r1
int-x360r2
```

```
/etc/hosts
```

9. Configure the hosts equivalence.

Add the public and private hosts name to the /etc/hosts.equiv file.

```
x360r1
x360r2
int-x360r1
int-x360r2
```

```
/etc/hosts.equiv
```

10. Make the rlogin and rsh services started automatically.

```
chkconfig -level 3 rlogin on
chkconfig -level 4 rlogin on
chkconfig -level 5 rlogin on
chkconfig -level 3 rsh on
chkconfig -level 4 rsh on
chkconfig -level 5 rsh on
```

11. Create the Oracle user.

```
mkdir /oracle
addgroup dba
adduser -G dba -g disk -h /oracle oracle
```

Edit the profile for the Oracle user /oracle/.bash_profile.

```
# .bash_profile

# Get the aliases and functions
if [ -f ~/.bashrc ]; then
    . ~/.bashrc
```



```

fi

# User specific environment and startup programs

PATH=$PATH:$HOME/bin

export PATH
unset USERNAME

export ORACLE_BASE=/oracle
export ORACLE_HOME=$ORACLE_BASE/92
export ORACLE_TERM=xterm
export TNS_ADMIN=$ORACLE_HOME/network/admin
export ORA_NLS33=$ORACLE_HOME/ocommon/nls/admin/data
export PATH=$PATH:$ORACLE_HOME/bin
export
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/lib:$ORACLE_HOME/oracm/lib
export
CLASSPATH=$ORACLE_HOME/jre:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib:$ORACLE_HOME/network/jlib
export THREADS_FLAG=native

export ORACLE_SID=DB1

/oracle/.bash_profile

```

Change ownership for the oracle directory.

```
chown oracle:dba /oracle
```

Make sure the raw devices can be read and written by the group disk.

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

12. Create the Oracle script to be started automatically at boot time.

Edit /etc/init.d/oracle, and fill it as below

```

#!/bin/bash
#
# chkconfig: - 92 35
# description: Load and unload Oracle parameters. \

start() {
    echo -n "Load Oracle 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
}

stop() {
    echo -n "Unload Oracle parameters: "
}

# See how we were called.
case "$1" in
    start)
        start
        ;;
    stop)
        stop
        ;;
    restart|reload)
        stop
        start
        ;;
)

```

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

/etc/init.d/oracle
```

Make the script start automatically at runlevel 3.
chkconfig -level 3 oracle on

CONFIGURING THE STORAGE

SOFTWARE INSTALLATION

Install the IBM FAStT Storage Manager version 8.21 for Red Hat Linux. This will allow you to configure the FAStT700 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 monitor the FAStT adapters behavior. Install the agent and the GUI tool.

To start the agent, run as root: `service qlremote start`

To use the program, run as root: `/opt/FAStT_MSJ/FAStT`

STORAGE USAGE RECOMMENDATIONS FOR ORACLE DATAFILES

Oracle9i RAC stores the datafiles in raw devices.

There is one thing you have to keep in mind when configuring and using Oracle9i RAC with Red Hat. With AS 2.1, you should not use LVM. **Red Hat does not support LVM.**

LVM brings manageability and flexibility in the storage management, but it is not required to have Oracle9i RAC running.

Limitations for raw devices creation

Each LUN of the SAN storage is seen in Linux as the disk device.

A disk device can be partitioned.

A partition is mapped to a raw device.

Each disk device can be partitioned in a maximum of 3 primary partitions and 1 extended partition containing 10 logical partitions.

Therefore, for each LUN in the SAN you can only have 13 raw devices.

This has to be taken in account when you configure the SAN storage. When you create a LUN, you have to remember that you will be able to create only 13 raw devices out of it. You will not be able to put all the datafiles of the database on one huge LUN.

For example, if you need 15 datafiles for the data in your database, you will have to create 2 LUNs.

Creating raw devices

Use the `fdisk` utility to create partitions on the disk devices. Create 3 primary partitions and one extended partition, and afterwards create up to 10 logical partitions in the extended partitions.

These partitions will be seen in Linux as block devices (`/dev/sdb1`, `/dev/sdb2`, `/dev/sdb3`, `/dev/sdb5` ...).

To be able to use these fresh partitions, the FAStT adapter driver has to be reloaded with the commands below that you have to run as root:

```
rmmod qla2300
modprobe qla2300
```

To use the partitions as raw devices, you have to map the block devices to the raw devices. For this, fill in the file `/etc/sysconfig/rawdevices` as shown below:

```
# raw device bindings
# format: <rawdev> <major> <minor>
#         <rawdev> <blockdev>
# example: /dev/raw/raw1 /dev/sda1
#          /dev/raw/raw2 8 5

/dev/raw/raw1 /dev/sdo1
/dev/raw/raw2 /dev/sdo2
/dev/raw/raw3 /dev/sdo5
/dev/raw/raw4 /dev/sdo6
/dev/raw/raw5 /dev/sdo7
/dev/raw/raw6 /dev/sdo8
/dev/raw/raw7 /dev/sdo9
/dev/raw/raw8 /dev/sdo10
/dev/raw/raw9 /dev/sdo11
```

```

/dev/raw/raw10 /dev/sdp5
/dev/raw/raw11 /dev/sdp6
/dev/raw/raw12 /dev/sdp7
/dev/raw/raw13 /dev/sdp8
/dev/raw/raw14 /dev/sdp9
/dev/raw/raw15 /dev/sdp10
/dev/raw/raw16 /dev/sdp11

```

Start the rawdevices service to load the raw devices configuration, for this run the following command as root:

```
service rawdevices start
```

We recommend to start this service automatically in runlevel 3 (4 and 5). Use the following command:

```
chkconfig --level 3 rawdevices on
```

Adding datafiles when the database is online

If you add a partition with *fdisk* or create a new LUN on the SAN storage, you have to reload the FASTT adapter driver. This means you lose the disk access for some time, and therefore the database instance on the server has to be shutdown.

To maximize database availability, it is recommended to create some spare partitions to use them in the future as new datafiles.

When you run out of new partitions, for each node in the cluster, you will have to shutdown 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.

STORAGE CONFIGURATION

In order to configure the storage and to create the raw devices, draw a table like the following one. This will help you to determine how to configure the SAN storage, how many LUN you have to create, what size to give the LUNs, and it will help you to remember the purpose of each raw device.

The first column is the array you have to create on the FASTT, the second is the LUN name in the array, the third is the description of the datafile that will be put in the raw device, the fourth is the size of the datafile, the fifth is the block device created with *fdisk*, and the sixth is the raw device name.

RAC raw devices

Array name	LUN	Database file	size	Block device	Raw device
RAC	R1	quorum	110 MB	/dev/sdo1	/dev/raw/raw1
RAC	R1	srvm	110 MB	/dev/sdo2	/dev/raw/raw2

Database raw devices

Array name	LUN	Database file	size	Block device	Raw device
Data	D1	SPFILE	5 MB	/dev/sdo5	/dev/raw/raw3
Data	D1	SYSTEM	400 MB	/dev/sdo6	/dev/raw/raw4
Data	D1	USERS	120 MB	/dev/sdo7	/dev/raw/raw5
Data	D1	TOOLS	12 MB	/dev/sdo8	/dev/raw/raw6
Data	D2	TEMP	110 MB	/dev/sdo9	/dev/raw/raw7
Data	D2	Control file 1	110 MB	/dev/sdo10	/dev/raw/raw8
Data	D2	Control file 2	110 MB	/dev/sdo11	/dev/raw/raw9
Index	I1	INDEX	70 MB	/dev/sdp5	/dev/raw/raw10
Redo	R1	Redo log Thread 1 Group 1	120 MB	/dev/sdp6	/dev/raw/raw11
Redo	R1	Redo log Thread 1 Group 2	120 MB	/dev/sdp7	/dev/raw/raw12
Redo	R1	Redo log Thread 2 Group 1	120 MB	/dev/sdp8	/dev/raw/raw13
Redo	R1	Redo log Thread 2 Group 2	120 MB	/dev/sdp9	/dev/raw/raw14
Undo	U1	Undo tablespace Thread 1	312 MB	/dev/sdp10	/dev/raw/raw15
Undo	U1	Undo tablespace Thread 2	312 MB	/dev/sdp11	/dev/raw/raw16

Configure the FASTT storage with the IBM Storage Manager client (*SMclient*), follow the instructions in the [Implementing IBM eServer xSeries SANs Redbook](#).

Follow the instructions from the previous section to create the partitions and the raw devices.

INSTALLING THE CLUSTER MANAGER

First of all you must install the Cluster Manager, which will let you perform the installation of Oracle Database on every node on your cluster. You can easily install Cluster Manager using the Oracle Universal Installer that come with each Oracle product.

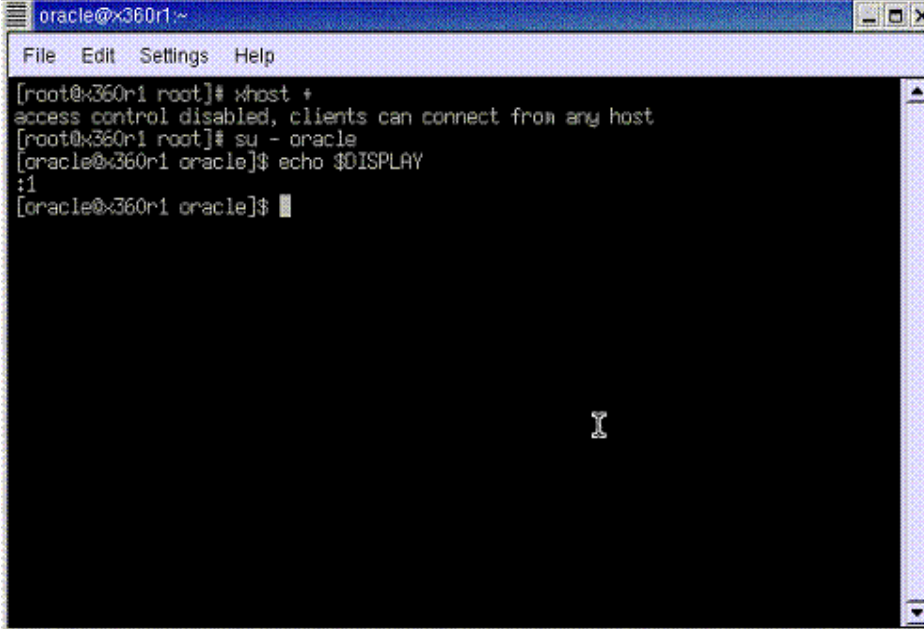
Follow these procedures to use the Oracle Universal Installer to install the Oracle Cluster Manager and the Real Application Clusters software. Oracle 9i RAC is supplied on multiple CD-ROM disks, so you must keep all the CDs mounted or available for the installer. During the installation process it may be necessary to switch between the CD-ROMs. Oracle Universal Installer will manage the switching between CDs and ask you for the CD-ROM location if necessary.

To install the Oracle Real Application Clusters Software, you should perform the following operations:

First of all, you should mount as `root` user at least the first CD-ROM. If you mount all the CDs on a directory it will be easier to switch from each CD. We will follow this procedure with a locally mounted CD-ROM, but you could do it with NFS (or equivalent) mounted media.

Allow the `oracle` user access to the X Window System server. To do so, do "`xhost +`" as root user and then log in as the "`oracle`" user with the "`su - oracle`" command.

Make sure your display is set with "`export DISPLAY=:0`" (or "`setenv DISPLAY :0`"). We used VNC while writing this document, so on the next captured screen you will see the `$DISPLAY` variable set to "`:1`" because it was our active display. Remember that the `oracle` user may not be able to access the screen if this variable is not correctly set.



```
oracle@x360r1 ~  
File Edit Settings Help  
[root@x360r1 root]# xhost +  
access control disabled, clients can connect from any host  
[root@x360r1 root]# su - oracle  
[oracle@x360r1 oracle]$ echo $DISPLAY  
:1  
[oracle@x360r1 oracle]$
```

Then you can launch the Oracle Universal Installer with "`<cdrom_mount_point>/runInstaller`" or "`./runInstaller`" from where you have mounted the CD-ROM. As we have mounted our CD-ROM on the `/mnt/cdrom` directory, we run the installation script from there as shown on the next figure.

```

oracle@x360r1:/mnt/cdrom
File Edit Settings Help
drwxr-xr-x  2 root  root    4096 Jun 21  2001  initrd
drwxr-xr-x  8 root  root    4096 Sep  9 11:30  lib
drwx----- 2 root  root   16384 Sep  4 16:10  lost+found
drwxr-xr-x  2 root  root    4096 Aug 29  2001  misc
drwxr-xr-x  5 root  root    4096 Sep  5 17:21  mnt
drwxr-xr-x  3 root  root    4096 Sep  4 16:18  opt
drwxr-xr-x  2 oracle dba    4096 Sep 10 14:06  oracle
dr-xr-xr-x 91 root  root     0 Sep  9 13:30  proc
drwxr-xr-x 12 root  root    4096 Sep 10 15:31  root
drwxr-xr-x  2 root  root    4096 Sep  4 16:19 /sbin
drwxr-xr-x  8 root  root    4096 Sep 10 10:45  tmp
drwxr-xr-x 15 root  root    4096 Sep  4 16:13  usr
drwxr-xr-x 20 root  root    4096 Sep  4 16:19  var
[oracle@x360r1 /]$ cd /mnt/cdrom
[oracle@x360r1 cdrom]$ ll
total 18
drwxr-xr-x  6 2840  562  2048 May 14 02:46  dbc
-rwxr-xr-x  1 2840  562  5688 May 14 02:46  index.htm
drwxr-xr-x  4 2840  562  2048 May 14 02:46  install
drwxr-xr-x  2 2840  562  2048 May 14 02:46  lgto
drwxr-xr-x  2 2840  562  2048 May 14 22:23  response
-rwxr-xr-x  1 2840  562  1163 May 14 02:46  runInstaller
drwxr-xr-x  7 2840  562  2048 May 14 03:30  stage
[oracle@x360r1 cdrom]$ ./runInstaller

```

You can also copy the contents of the CDs in the stage directory onto the hard disks, so that you will not be prompted to switch the CD in the CDRom drive. Use a directory structure as follows:
/<mount point>/Disk1, /<mount point>/Disk2, /<mount point>/Disk3

Welcome screen.

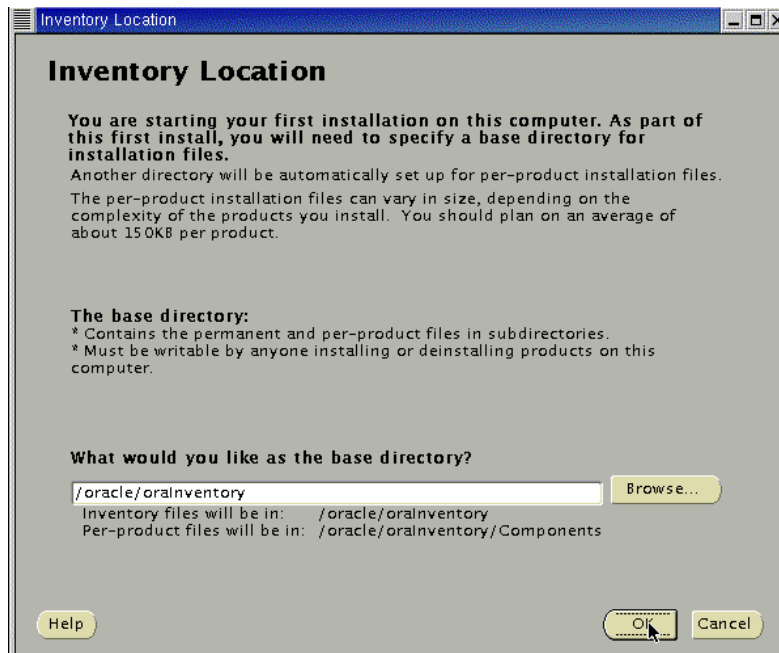
The Oracle Universal Installer loads the Java Runtime Environment (you may receive some messages on your terminal window regarding the fonts but this is normal, because if the Java Runtime Environment does not find all the fonts it needs, then it substitutes them with others), and then shows you a welcome screen like the one you can see on the next capture:



Click "Next" to go to the next screen.

Inventory Location screen.

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`. On this screen, which looks like the image you can see on the next capture, 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 on the future.



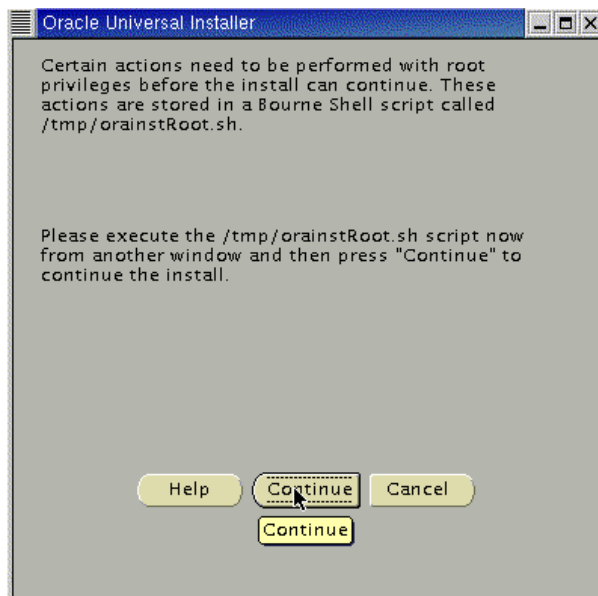
UNIX Group screen.

The next screen 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. So, enter "oinstall" and once again, simply click on "Next" to continue. This screen will look like this:



Shell script as root.

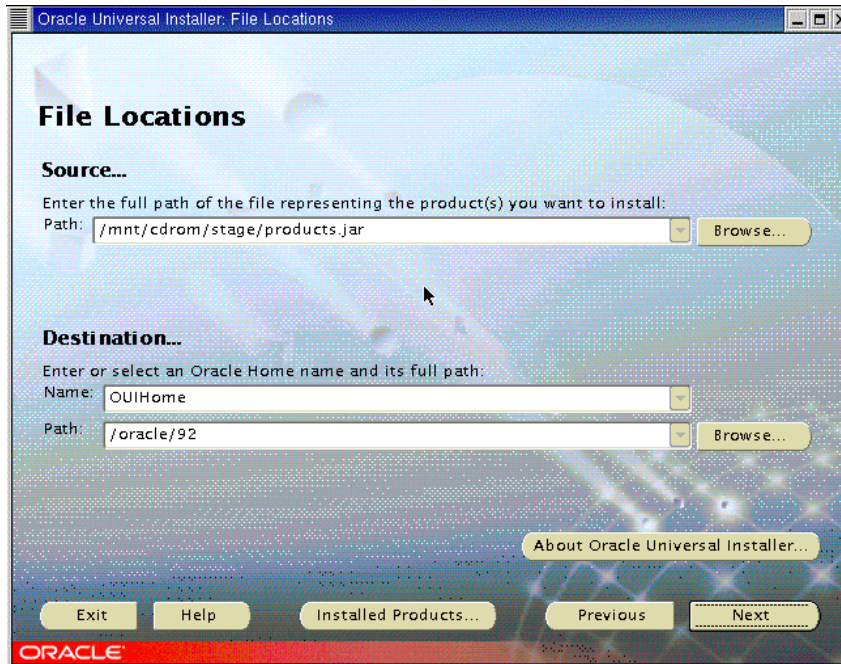
You will be then prompted to execute a shell script as root user. This script is located at `/tmp/orainst.sh` and you must execute it in order for the installer to work properly.



File locations screen.

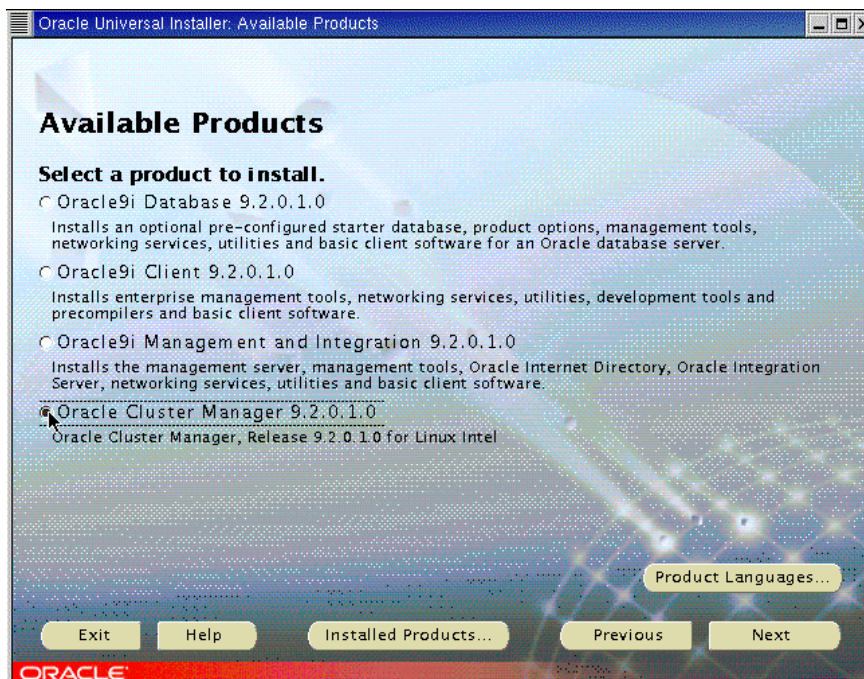
Once the shell script is executed, you click "OK" and the File Location screen should appear. This screen asks you for some information about the source and destination paths for you installations. Do not change the source location. For the destination location, verify that the value that you have pre-selected is the one you typed on the "ORACLE_HOME" variable that you configured on previous steps.

Once you verified that all the values are OK, then click on “Next” to continue. This screen will look like this:



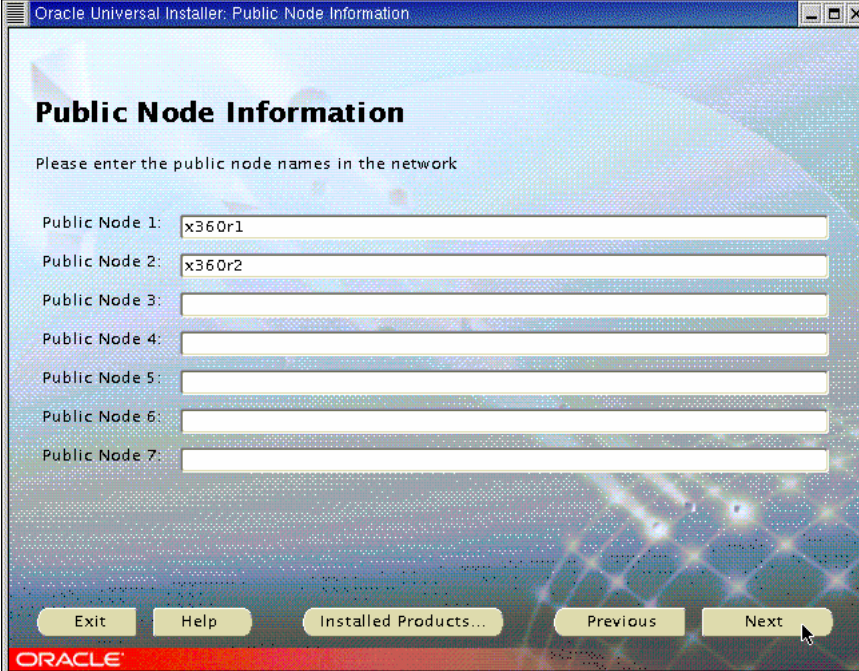
Product Selection screen.

The next screen will show you the different products you can install from the CDs. Select “Oracle Cluster Manager” and click “Next” to start the installation of the Cluster Manager. The screen you will see is this one:



Public nodes information screen.

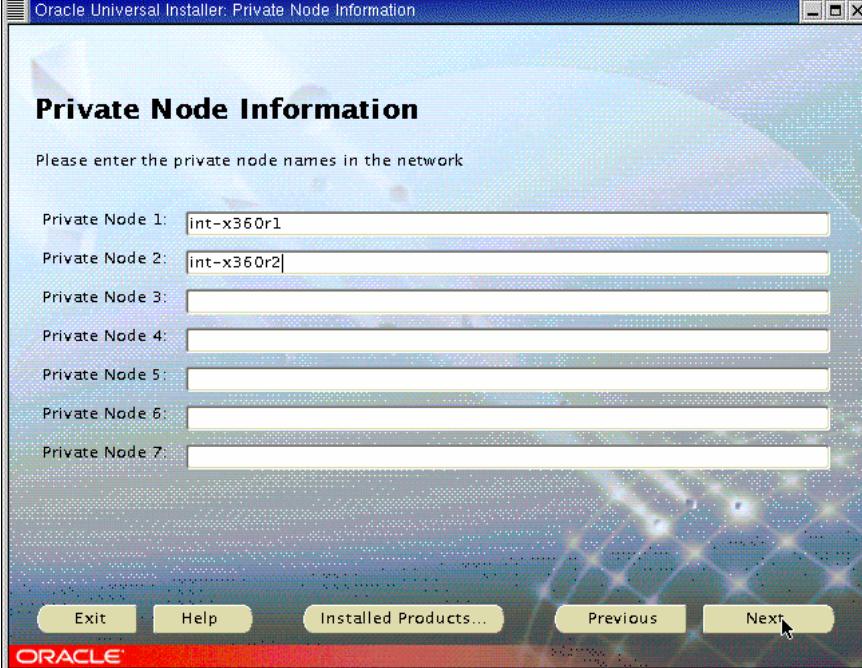
Oracle Universal Installer needs to know the public names or public IP addresses of the nodes of your cluster. This screen will let you introduce your node names. These nodes must be available from each other for the Oracle Universal Installer to configure the Real Application Clusters properly. In our case, we used two IBM eServer xSeries model 360 named "x360r1" and "x360r2". The next capture show this configuration screen.



The screenshot shows a window titled "Oracle Universal Installer: Public Node Information". The window has a blue header and a light blue background with a grid pattern. The main title is "Public Node Information". Below the title, it says "Please enter the public node names in the network". There are seven input fields labeled "Public Node 1:" through "Public Node 7:". The first two fields contain the text "x360r1" and "x360r2" respectively. The other five fields are empty. At the bottom of the window, there are five buttons: "Exit", "Help", "Installed Products...", "Previous", and "Next". The "Next" button is highlighted with a mouse cursor. The Oracle logo is visible in the bottom left corner of the window.

Private node information screen

On the next screen configure the Private Node addresses that will be used for the interconnect between the different nodes of your cluster. This screen is like this:



Oracle Universal Installer: Private Node Information

Private Node Information

Please enter the private node names in the network

Private Node 1:

Private Node 2:

Private Node 3:

Private Node 4:

Private Node 5:

Private Node 6:

Private Node 7:

Exit Help Installed Products... Previous Next

ORACLE

Watchdog parameter screen

The next screen asks you for the watchdog parameters. Simply accept the default value (60000) and click "Next" to continue.



Oracle Universal Installer: WatchDog Parameter Information

WatchDog Parameter Information

Enter the watchdog parameter value in milliseconds, default is 60000

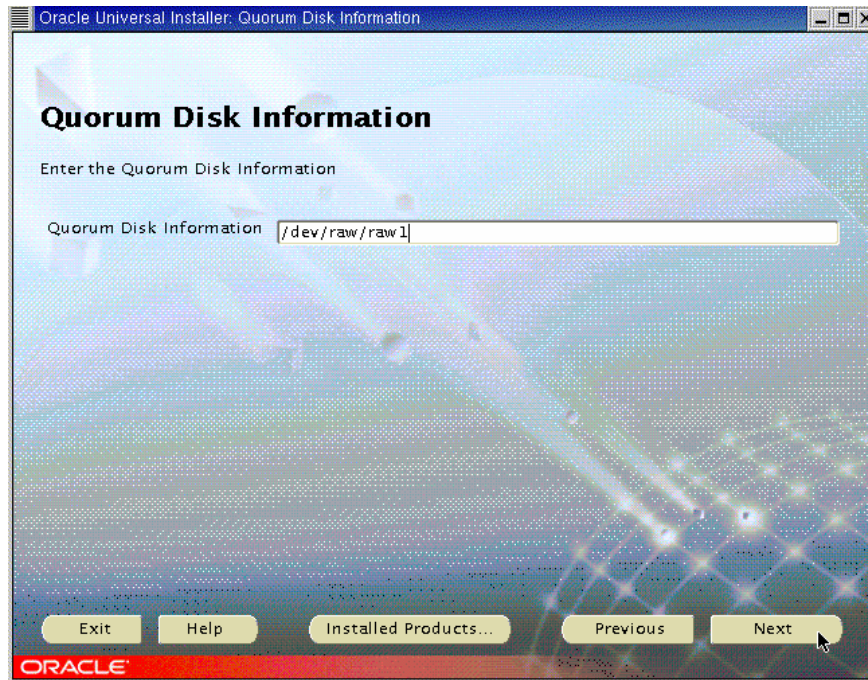
WatchDog Parameter Value

Exit Help Installed Products... Previous Next

ORACLE

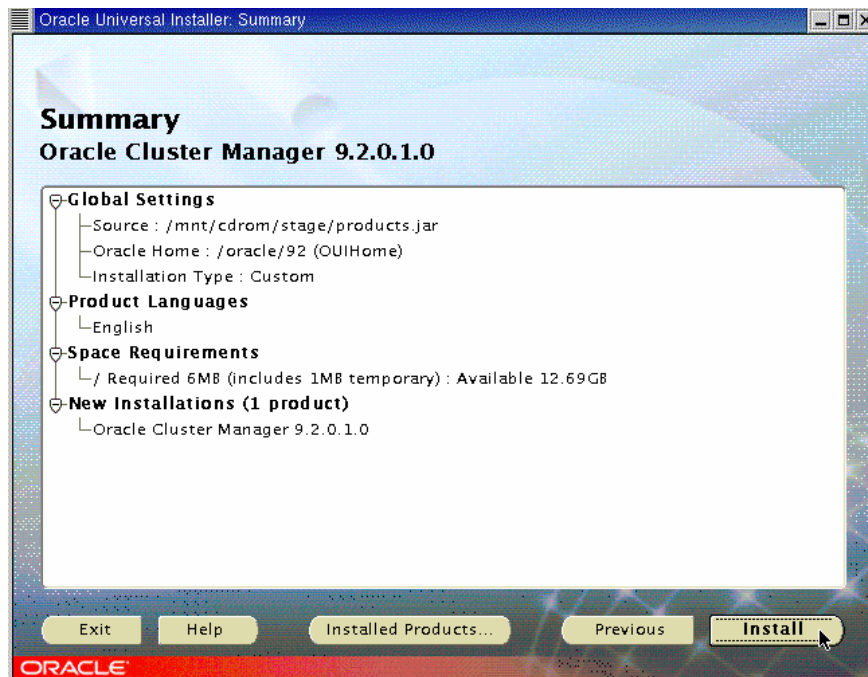
Quorum Disk Information

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/raw/raw1`, but do not forget to change it, because your configuration may be different. Click Next to continue.



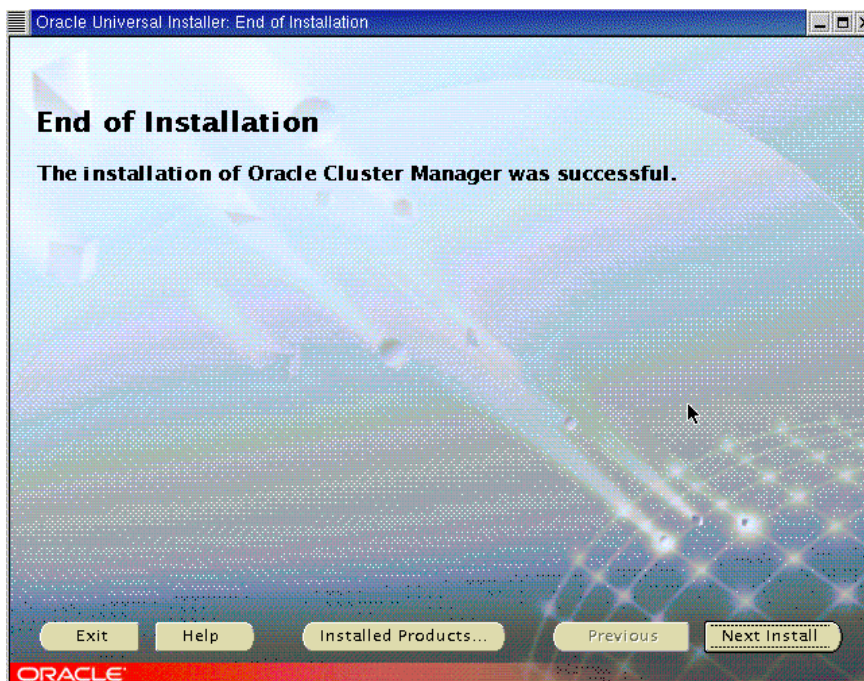
Summary Screen

The next screen is the summary one. Simply press Install.



Cluster Manager Installation screen

You will now briefly get a progress window followed by the end of installation screen. Click Exit. and confirm by clicking Yes.



Once this screen is shown, you can click Exit and confirm with Yes. Please, verify that you have successfully installed the Cluster Manager on each node.

Starting Cluster Manager

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

To do so, go as `oracle` user to `$ORACLE_HOME` on each node of your cluster and create the `$ORACLE_HOME/oracm/log` directory.

Make sure the the watchdog device `/dev/watchdog` exists on each of the nodes. If not, create it with the following commands, run them as root:

```
mkdev /dev/watchdog c 10 130  
chmod 600 /dev/watchdog
```

Then, log in as root on each node and run:

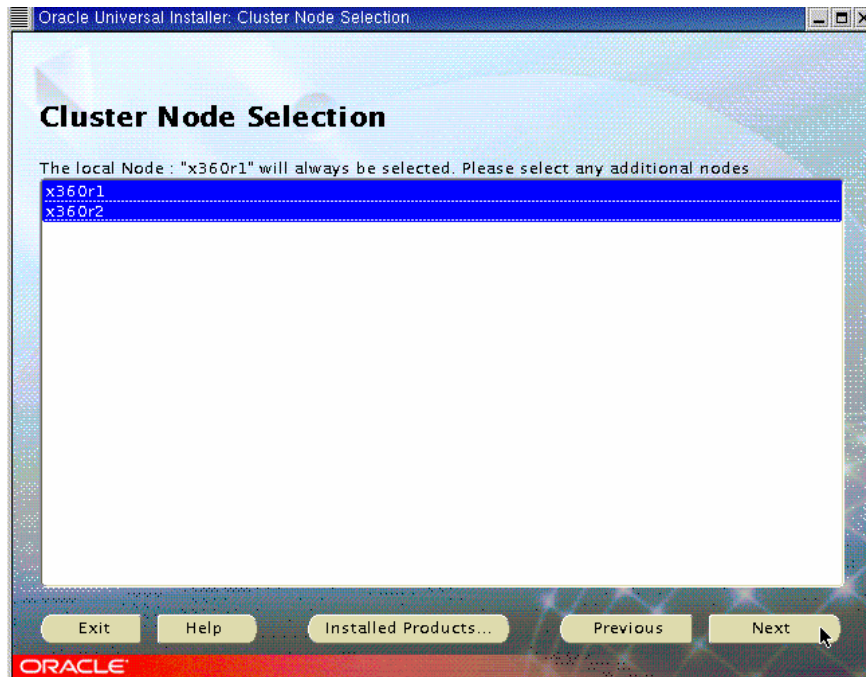
```
export ORACLE_HOME=/oracle/product/9.2.0  
$ORACLE_HOME/oracm/bin/ocmstart.sh
```

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

ORACLE9i DATABASE INSTALLATION

Installing Oracle9i

Start the Oracle Universal Installer again. If you click Next on the Welcome screen 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 screen would look like this:



Select all the nodes you would like to be on your cluster and click Next to continue. The next screen is the File Location Screen that we have already mentioned. Click on Next and continue.

Products Screen

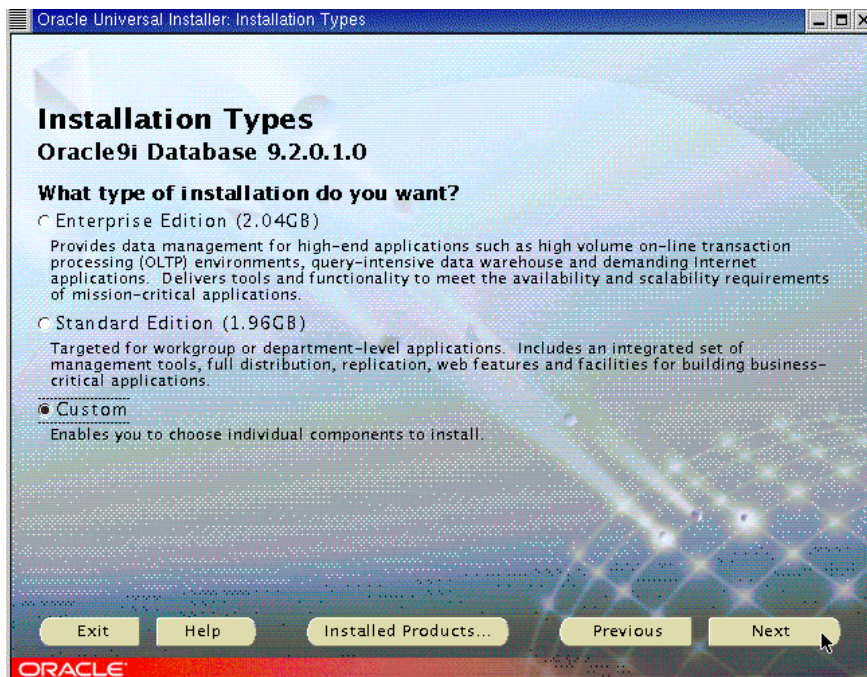
Select from on the Available Products screen the "Oracle Database" and click Next to continue as shown on the figure.



This will start the Oracle Database installation on each node of your cluster.

Installation Type screen.

The next screen is the Installation Type. Select the Custom option in order to have full control of each element you will install. Note that this Custom option refers to the Database installation, but not the configuration.



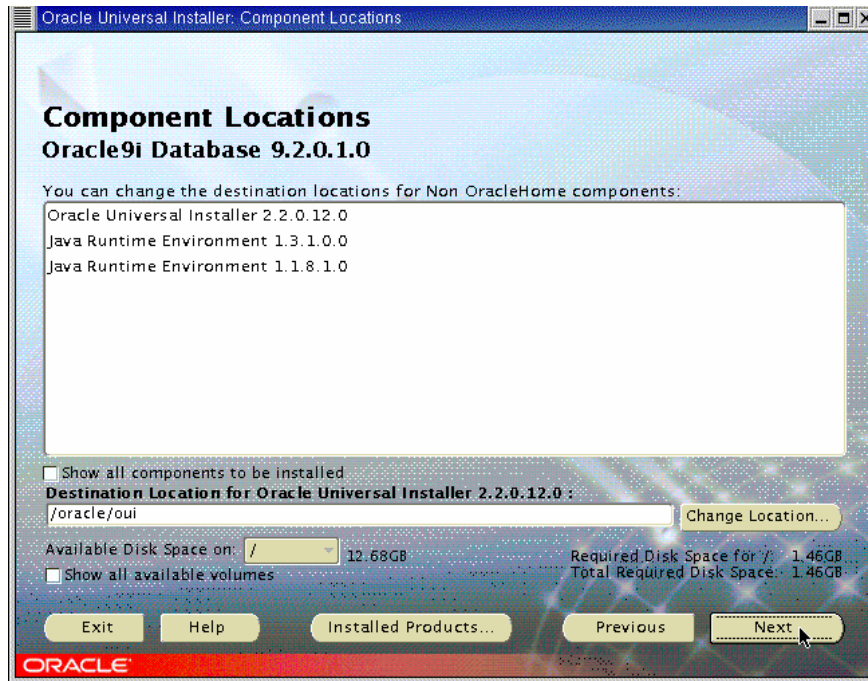
Components Selection Screen

On this screen, you can configure what elements will be installed. For a Real Application Cluster configuration, please be sure you install at least the Oracle Database, Real Application Clusters, Partitioning and Netlistener components and then click Next to start installation. If you need some additional components for your database, select them on this screen



Component Location Screen

This screen gives information about the components locations. Do not modify the values and click Next.



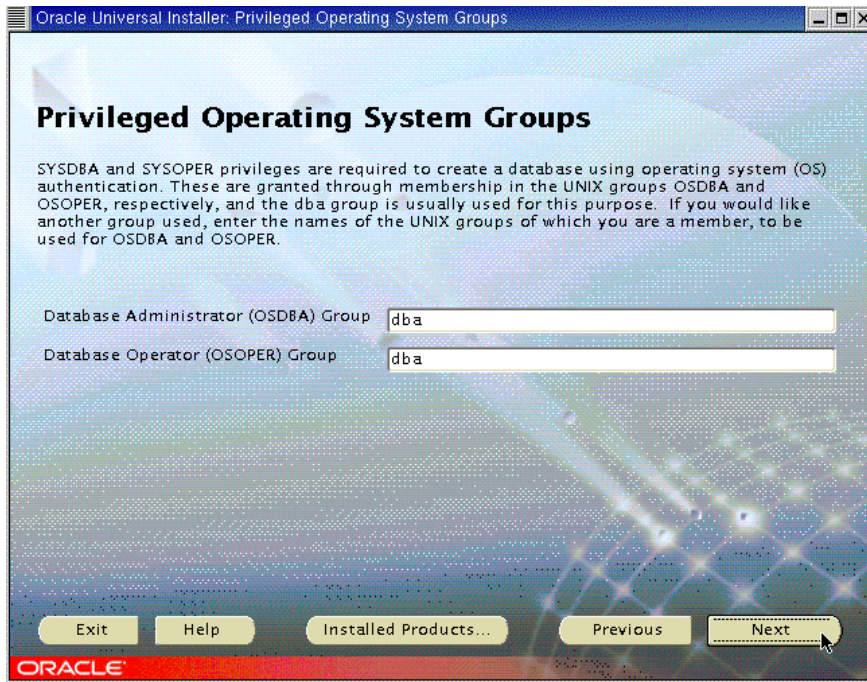
Shared Configuration File screen

You will also be prompted for the raw device where you will have the shared configuration of your cluster. Enter your value and click Next to continue.



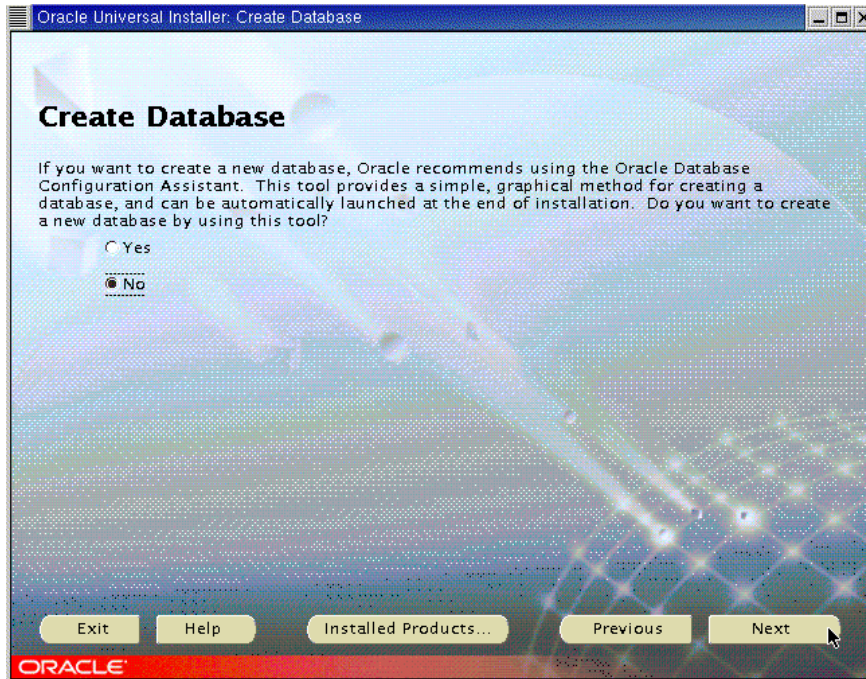
Group Selection screen

Then, select the group you will need for your database operator and administrator. Usually, these values are set to "dba", but verify that you have put there the group that you will use. Click Next to continue.



Database Creation screen

You will then be prompted for database creation. We will not create it right now, so select "No" and then 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.



Summary Screen

This latest screen shows you the configuration you have made. Verify it and if you need to add something else, you can always go back by clicking on "Back". If everything is OK, click "Install" to start the installation procedure.




Installation screen

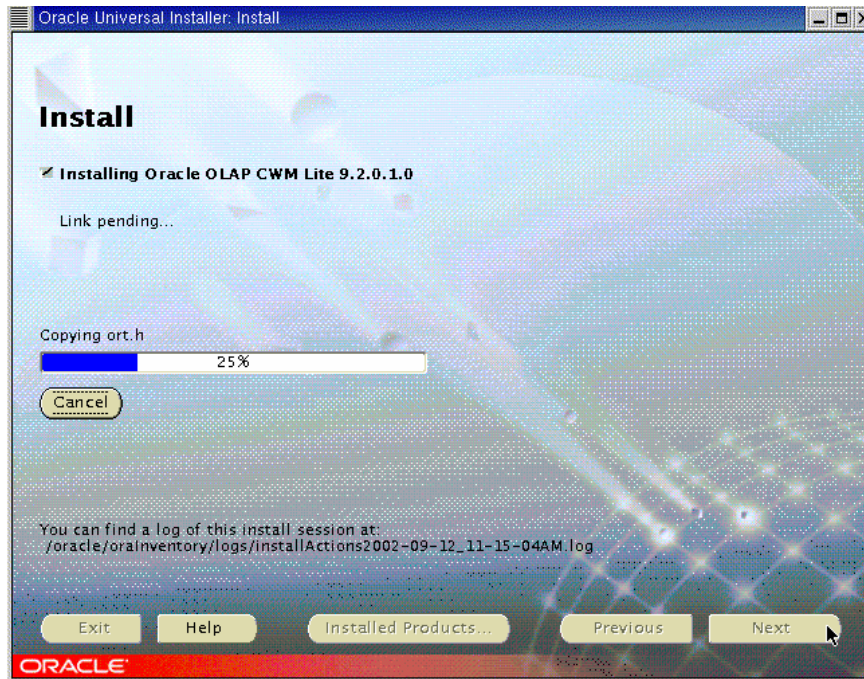
During the installation process, you will see a screen like this one:




You may be prompted for another CD-ROM. If it is your situation, unmount the CD-ROM, change it, and then mount the CD-ROM you have been asked for.



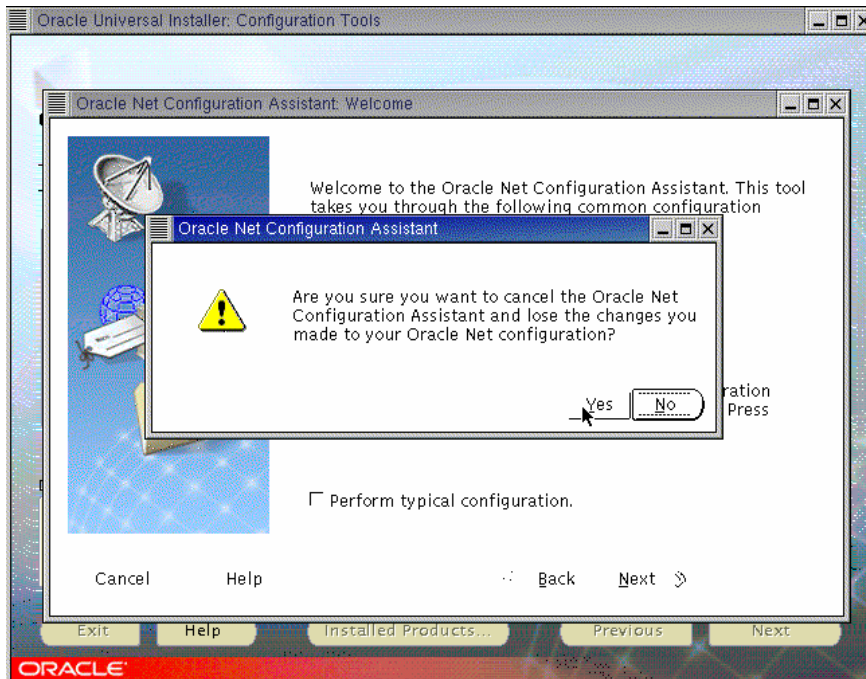
Then the installation procedure will continue.



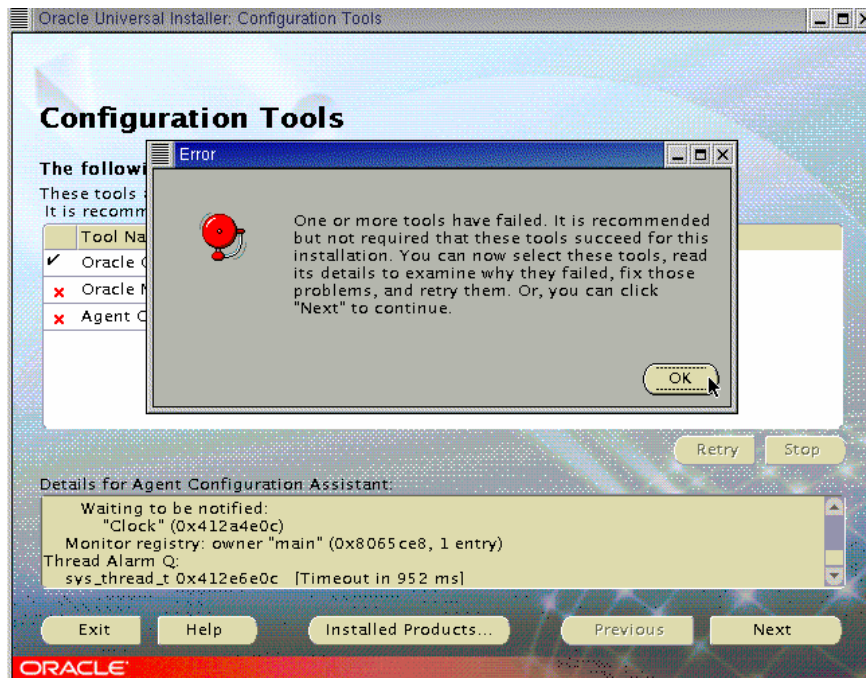
Once all the CDs are installed on the first node, the Oracle Universal Installer will install the software on each node on your cluster. Then, you will be prompted for a shell script execution. Do it on each node by accepting the default values and press OK to continue.



Then the configuration procedure will start. First, the Oracle Universal Installer will automatically configure the cluster and then ask you to configure the network. We will configure the listeners later, so you can cancel the network configuration procedure and confirm by clicking Yes.



You will get an error screen because we cancelled the network configuration. Accept the error by clicking OK and Next to finish the installation.



The Oracle Universal Installer shows you the End of Installation screen and warns you that you have got some components that were not configured. This is normal, so exit the installer with Yes and you are done.



Post-installation Configuration

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

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

```

oracle@x360r2:~/92
File Edit Settings Help
[oracle@x360r2 92]$ pwd
/oracle/92
[oracle@x360r2 92]$ mkdir rdbms/audit
[oracle@x360r2 92]$ mkdir rdbms/log
[oracle@x360r2 92]$ mkdir network/log
[oracle@x360r2 92]$ █

```

Cluster Manager configuration

We will now discuss how to configure the Cluster Manager, the Node Monitor and the SRVCTL configuration file.

First of all, launch 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.

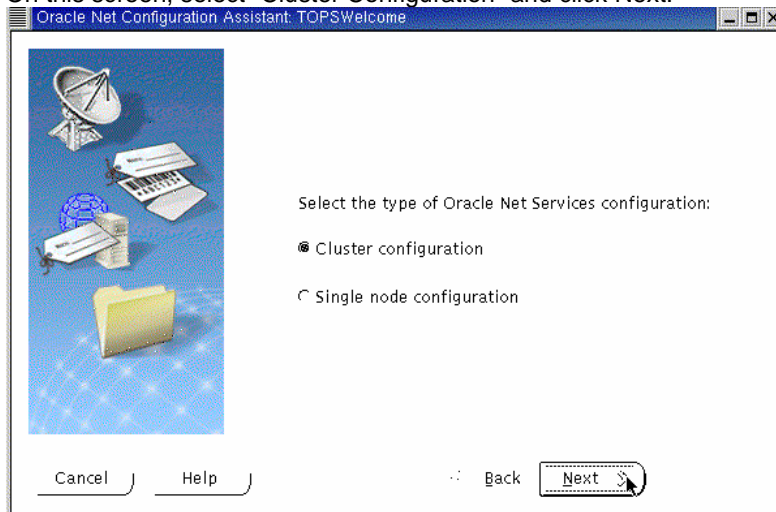
As oracle user, type “gsdctl start” on each node of your cluster to start the Global Services Daemon.

Listeners configuration

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”.

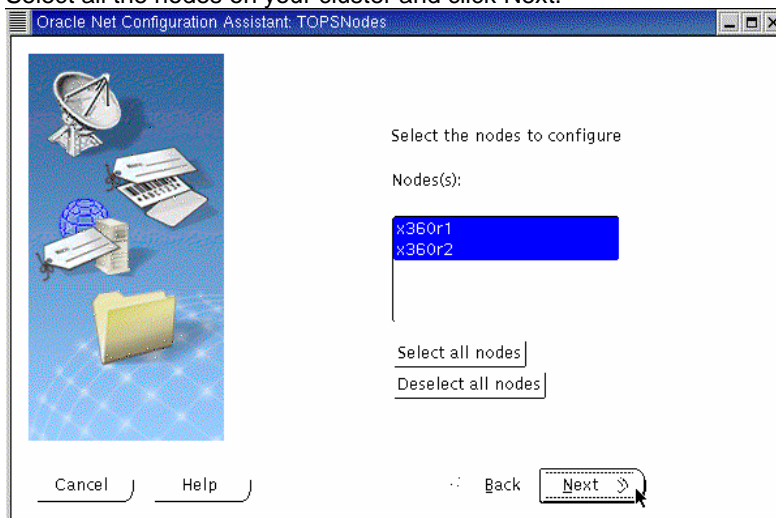
Welcome Screen

On this screen, select “Cluster Configuration” and click Next.



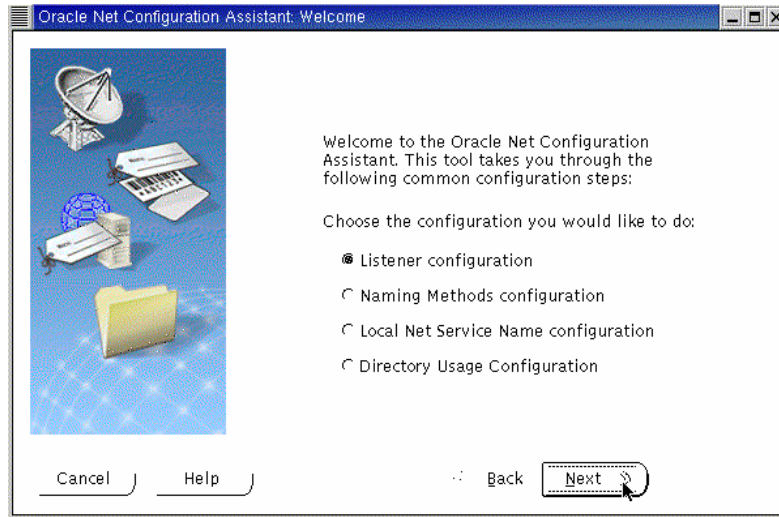
Nodes Screen

Select all the nodes on your cluster and click Next.

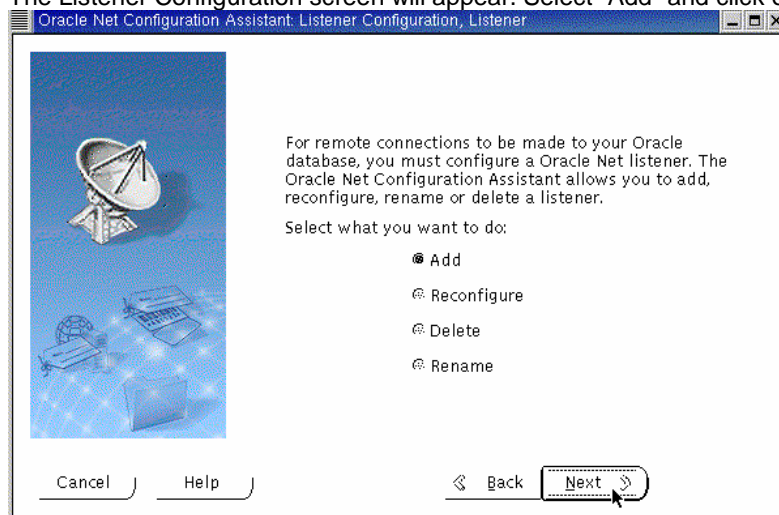


Listener Configuration screens

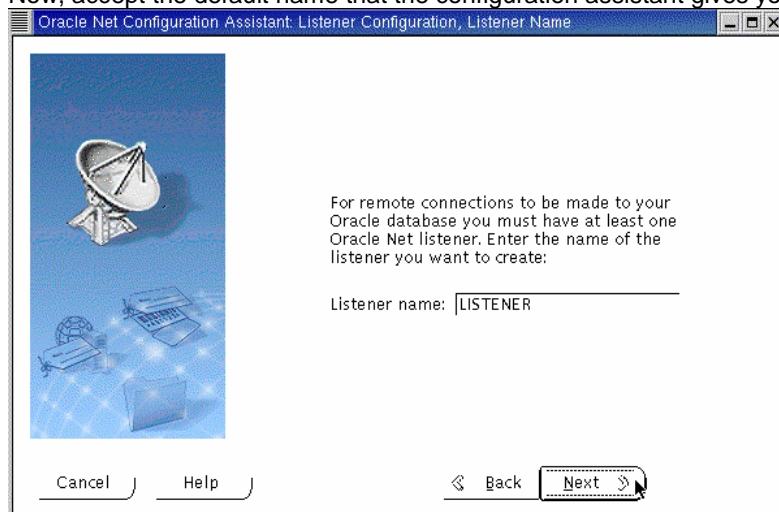
The next screen will offer you some configuration possibilities: Listener Configuration, Naming Method Configuration, Local Net service Name Configuration and Directory Usage Configuration. Select “Listener Configuration” and click Next.



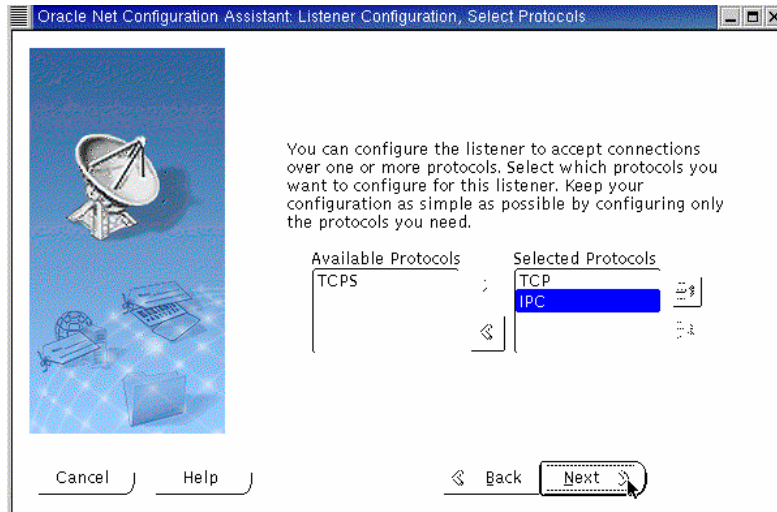
The Listener Configuration screen will appear. Select "Add" and click on Next to continue.



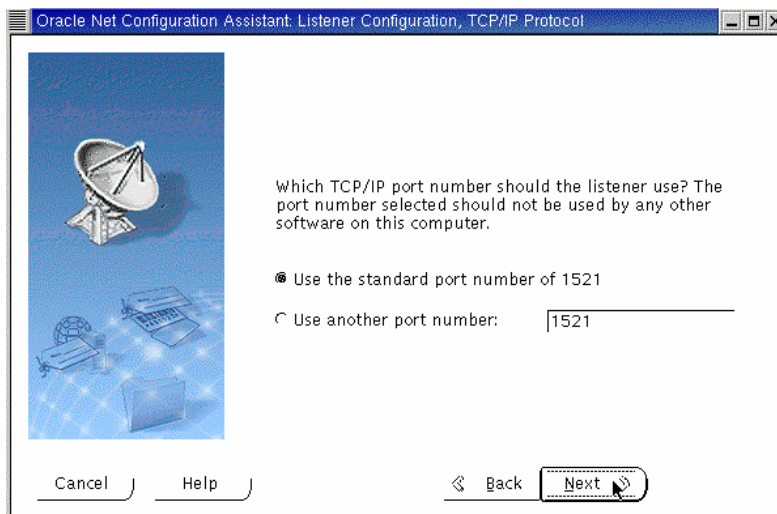
Now, accept the default name that the configuration assistant gives you and click Next again.



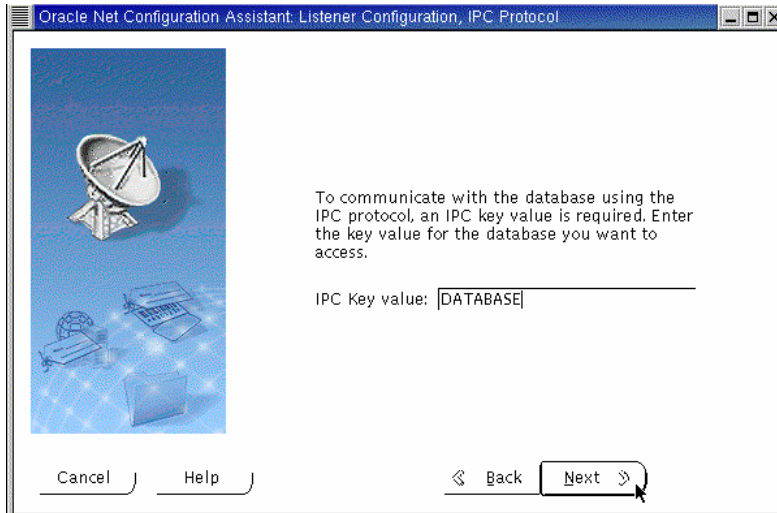
The next screen will let you choose between some network protocols. TCP/IP should be selected, but we will also add IPC support. Click on "IPC" and then on ">". To continue, click on Next.



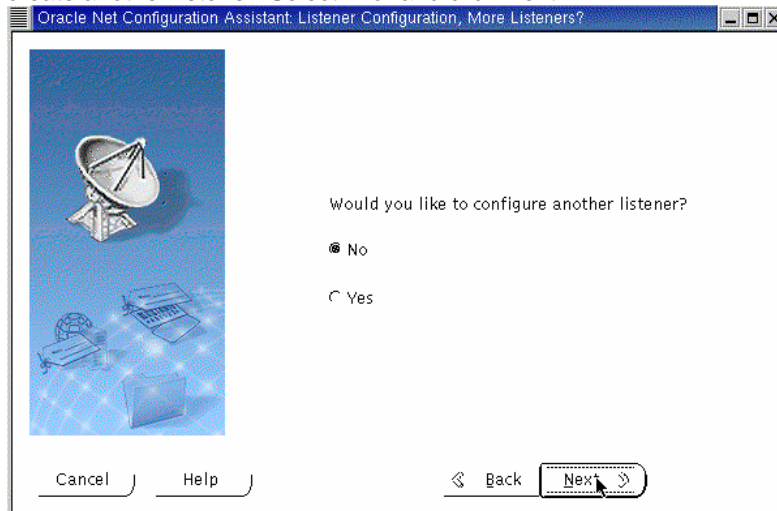
You will be then prompted for the port number you would like to use. Accept the default port value (set to 1521) and go to the next screen clicking on Next.



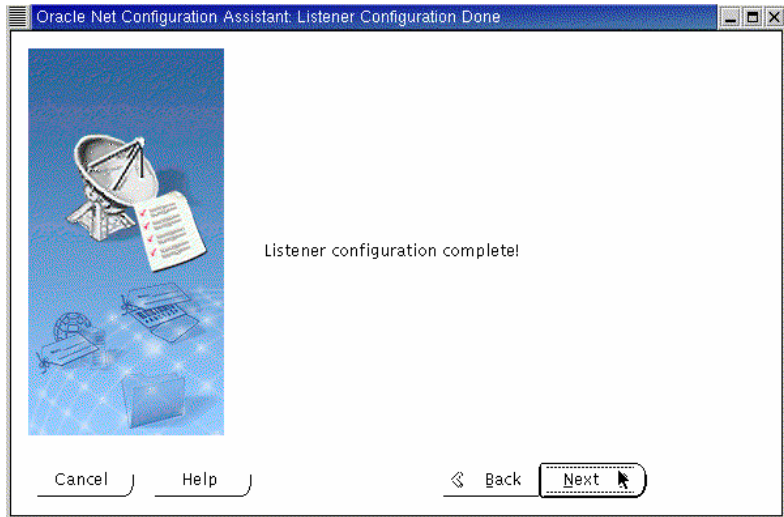
The next screen is the IPC Configuration one. On this screen you will be asked to insert your database name as key and then click Next.



This should create the listener. At this point, Oracle Universal Installer will ask you if you want to create another listener. Select "No" and click Next.



The Listener Configuration is now complete. You can now exit the Network Configuration Assistant unless you want to create another listener.



CREATING A CLUSTER DATABASE

For the database creation, we will use the DBCA tool. DBCA means Database Creation Assistant and it 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 Oracle9i features such as the server parameter file and automatic undo management. The DBCA also enables you to define arbitrary 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.

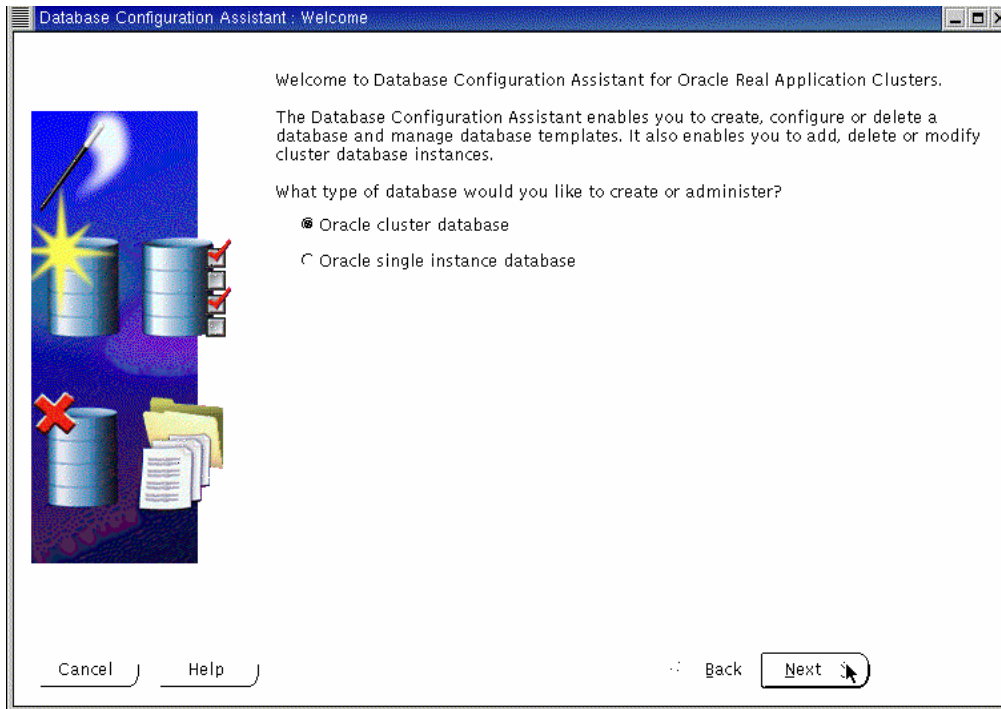
The Database Configuration Assistant (DBCA) makes use of a file, which is named mapping DBCA.cfg, that contains the raw device names for the system datafiles. You should create this file to suit your own tablespaces and associated raw devices and make sure it is owned by the oracle user. See below an example of the file.

```
spfile1=/dev/raw/raw3
system1=/dev/raw/raw4
users1=/dev/raw/raw5
tools1=/dev/raw/raw6
templ=/dev/raw/raw7
control1=/dev/raw/raw8
control2=/dev/raw/raw9
indx1=/dev/raw/raw10
redo1_1=/dev/raw/raw11
redo1_2=/dev/raw/raw12
redo2_1=/dev/raw/raw13
redo2_2=/dev/raw/raw14
undotbs1=/dev/raw/raw15
undotbs2=/dev/raw/raw16
```

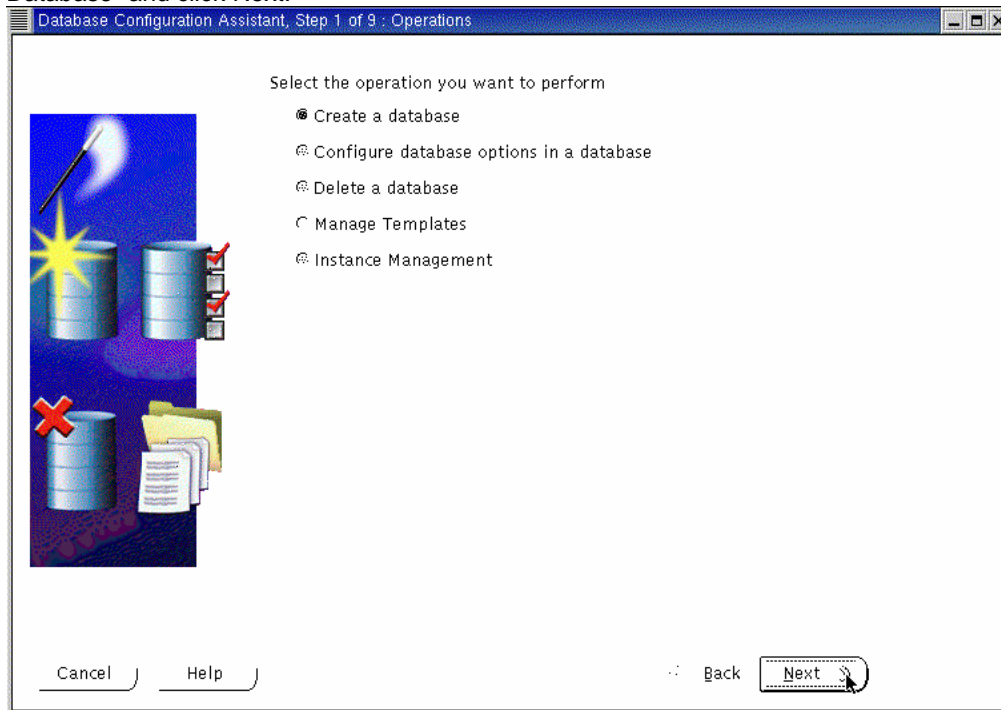
DBCAmapping.cfg

The location of this file is passed to DBCA by the variable DBCA_RAW_CONFIG. Set this variable with:
`export DBCA_RAW_CONFIG=~/.mappingDBCA.cfg`

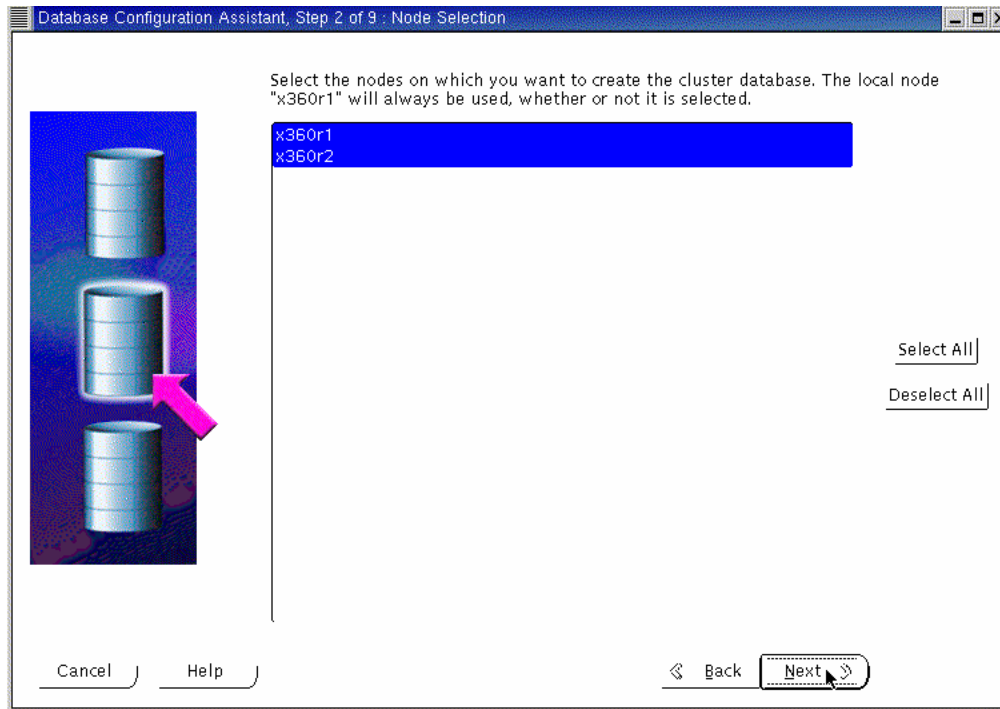
Start DBCA with the “dbca” command. You will see a welcome screen that looks like this one:



Select "Oracle Cluster Database" and click Next. The Operations screen will appear. Select "Create Database" and click Next.



You will be prompted for node selection. Select all the nodes that apply and click Next to continue.

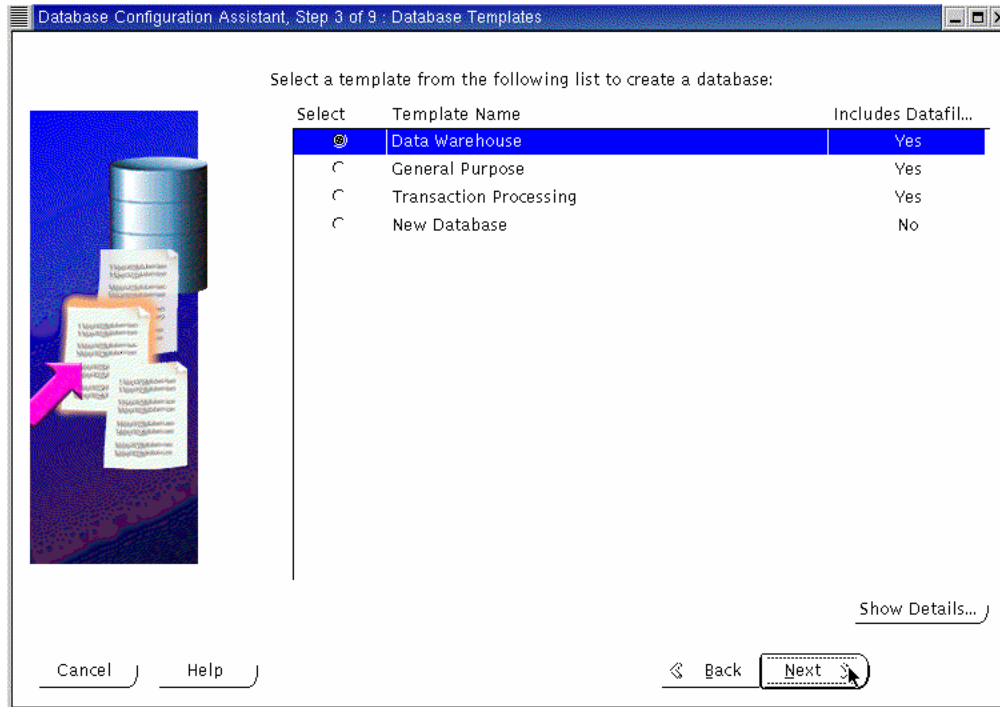


If some of your nodes are missing, there is a cluster diagnostic tool called "lsnodes" which will let you know what nodes are active at this moment. Resolve the problem and try again the process from the beginning. In our tests, we used two nodes: "x360r1" and "x360r2", and "lsnodes" shows both of them.

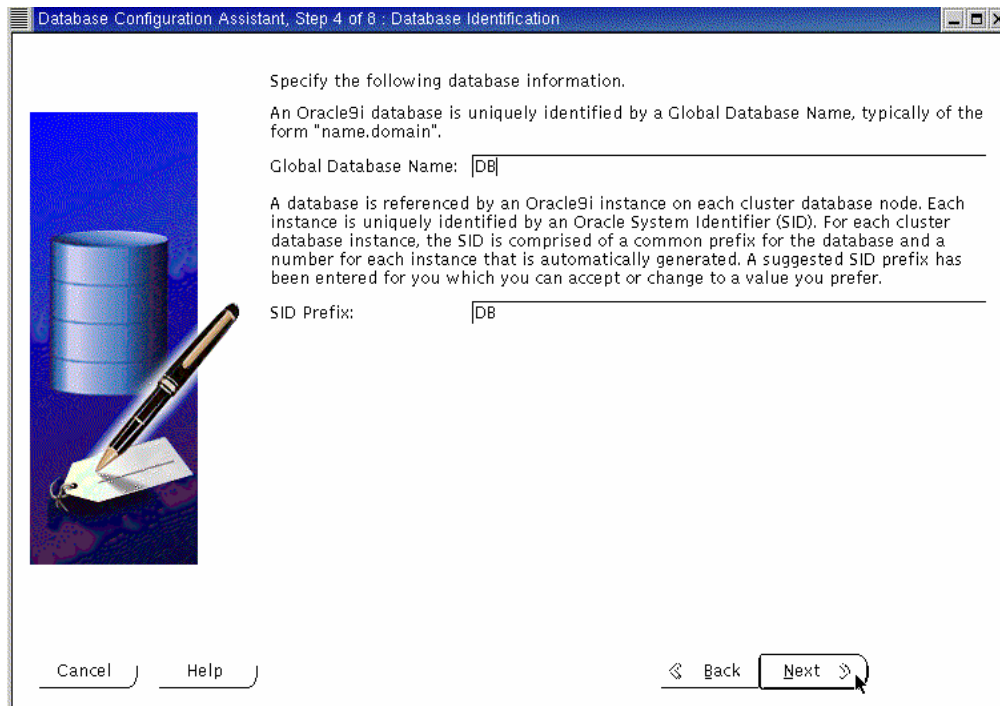


The next screen that you will see is the Database Templates one. All of the templates other than "New Database", include datafiles.

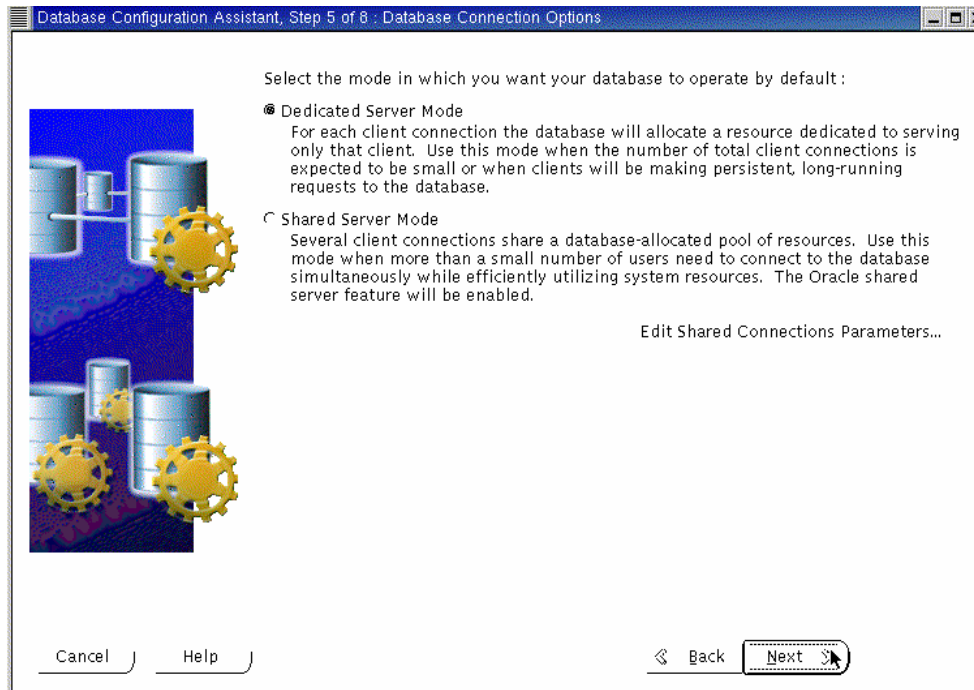
Select "New Database" and click Next.



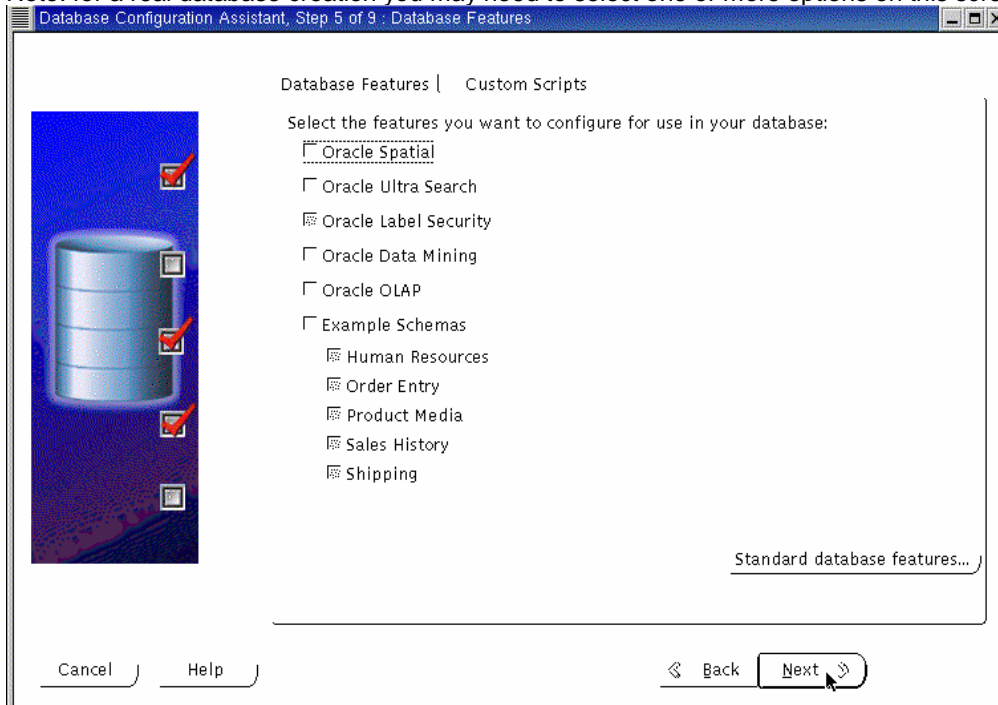
This next screen is the Database Identification page. You should enter the Global Database Name (typically on the name.domain form) and the Oracle System Identifier (SID). The SID will identify an instance and DBCA will suggest you use the same name that you entered on the Database Name field. The SID identifier will be used on RAC databases as a prefix for the instance number.



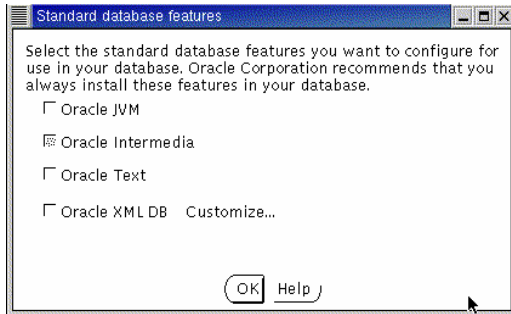
Select "Dedicated Server Mode" on the Database Connection screen.



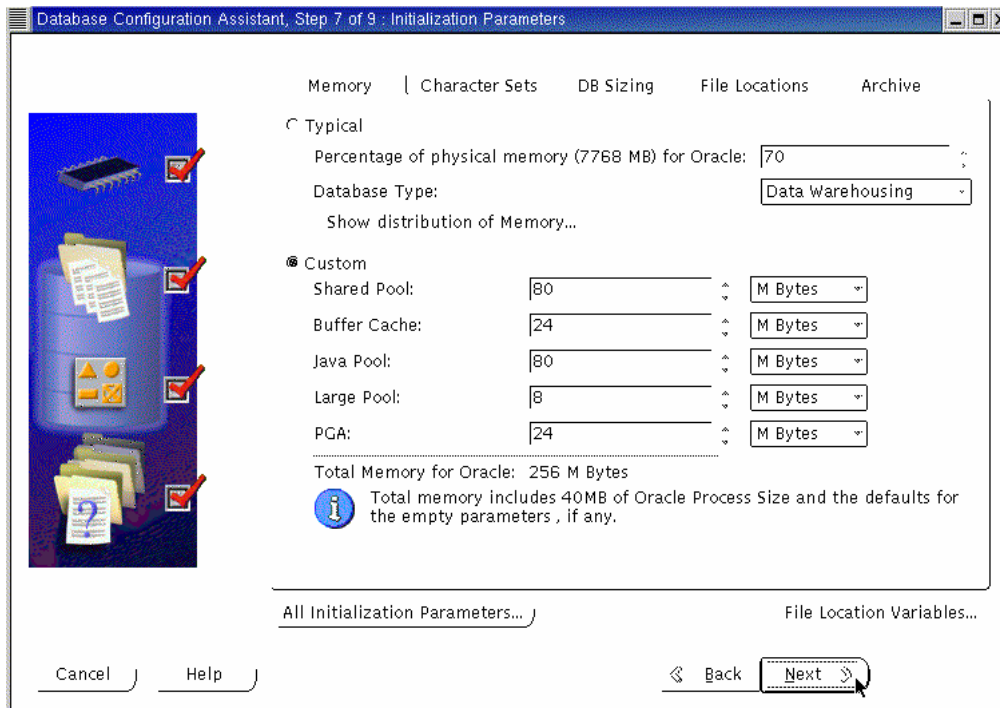
You will be prompted for Database Features. Unselect all the options and continue by clicking on Next.
Note: for a real database creation you may need to select one or more options on this screen.



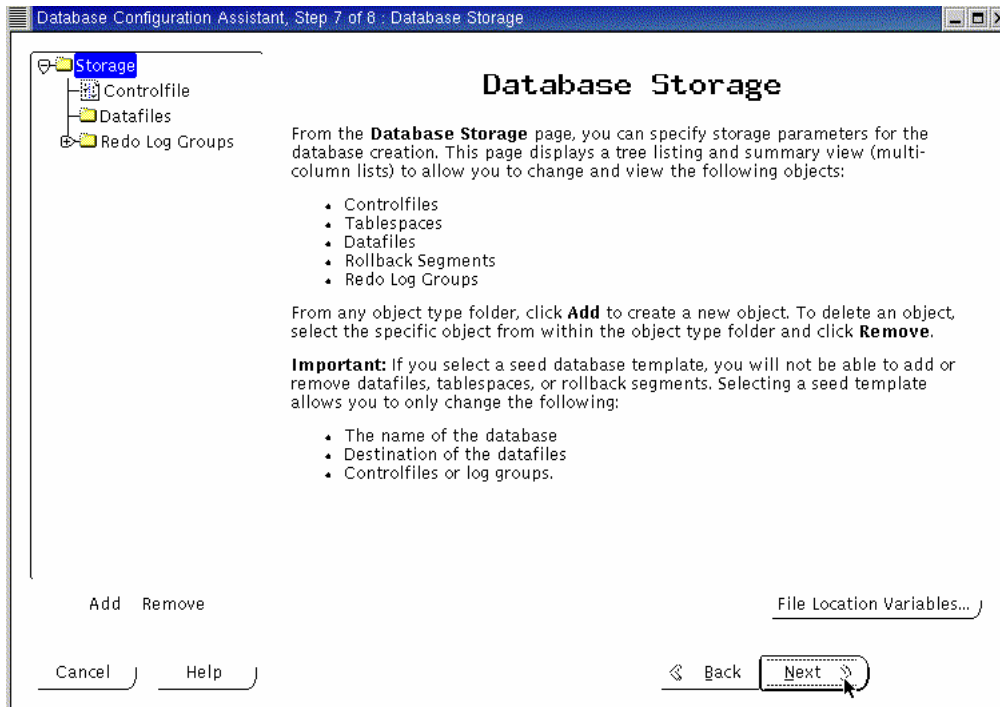
You can click on the “Standard database features” to add or remove some standard options. For this procedure, we deselected every standard option:



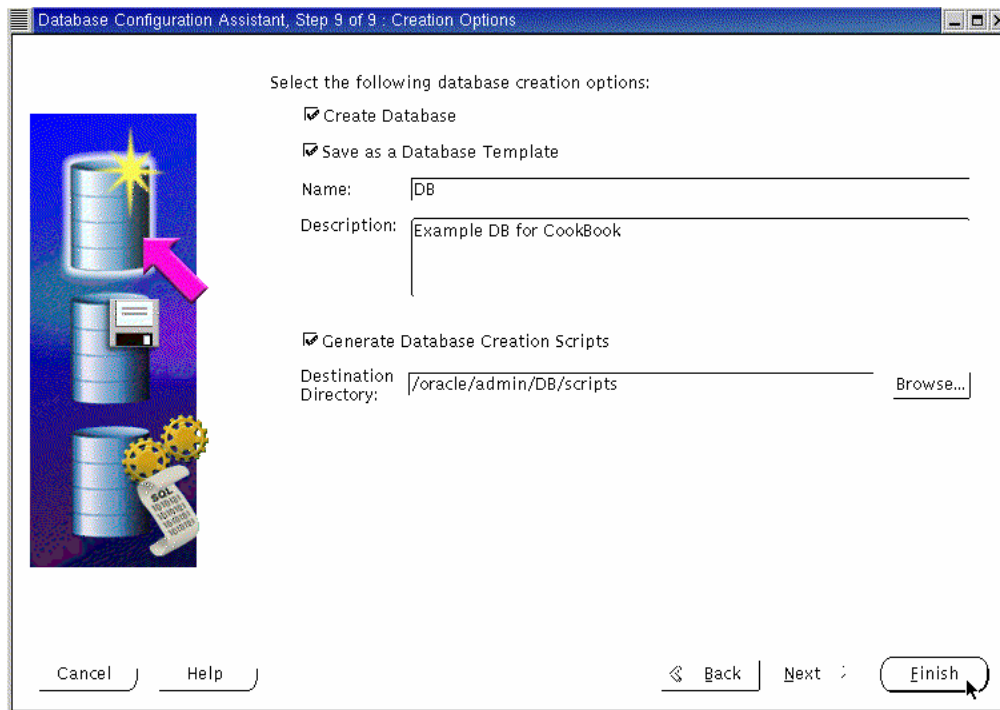
The next screen will let you configure the Initialization Parameters. On this screen you will see some tabular pages. Modify the memory setting if desired and then select the “File Locations” tab. The option “Create Persistent Initialization” is selected by default. The file name should point to the correct database link. Otherwise it will mean that you didn’t set the DBCA_RAW_CONFIG correctly or the file contains errors. Verify all the configuration options according to your needs and make sure the listener values are set to the ones that you previously created. Then, click Next.



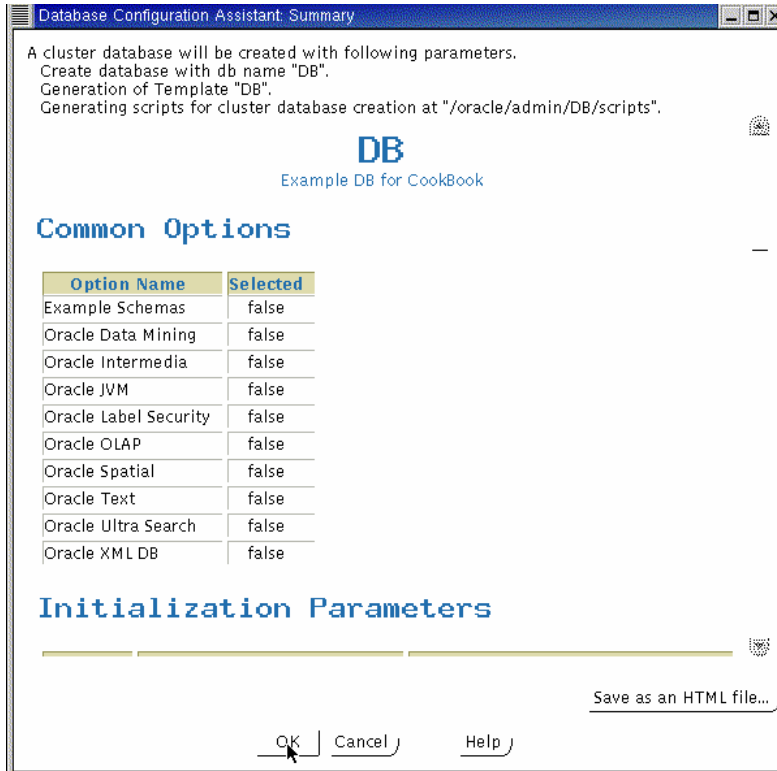
DBCA now shows you the Database Storage Window. This screen will let you enter file names for each tablespace in your database. They should be pointing to the right raw device or database link. Make sure you correct the datafile sizes to the sizes that you have planned to use and remember that the datafile should be smaller than the raw device. Click Next when you’re finished.




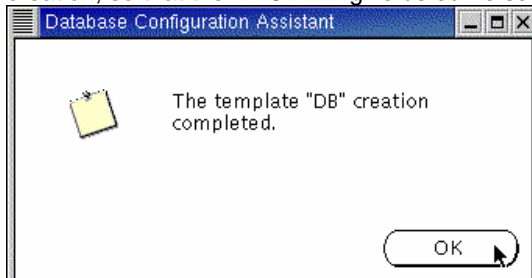
The Database Creation Options page is displayed. Ensure that the option Create Database is checked and click Finish.



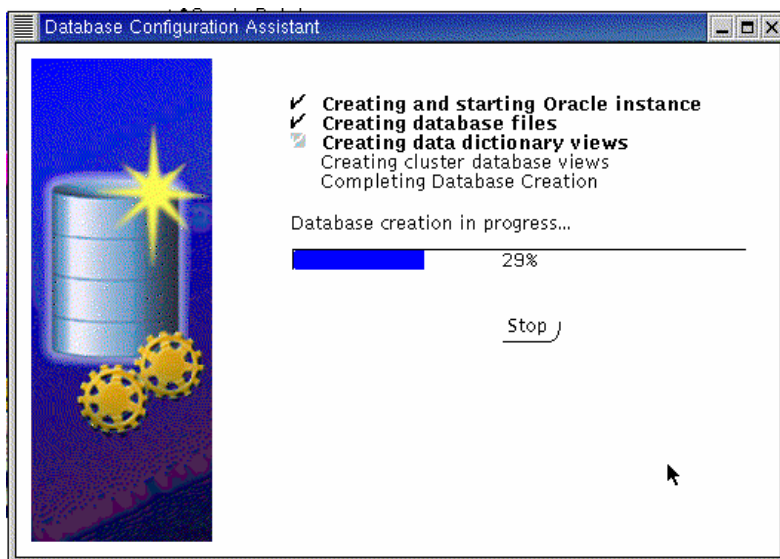
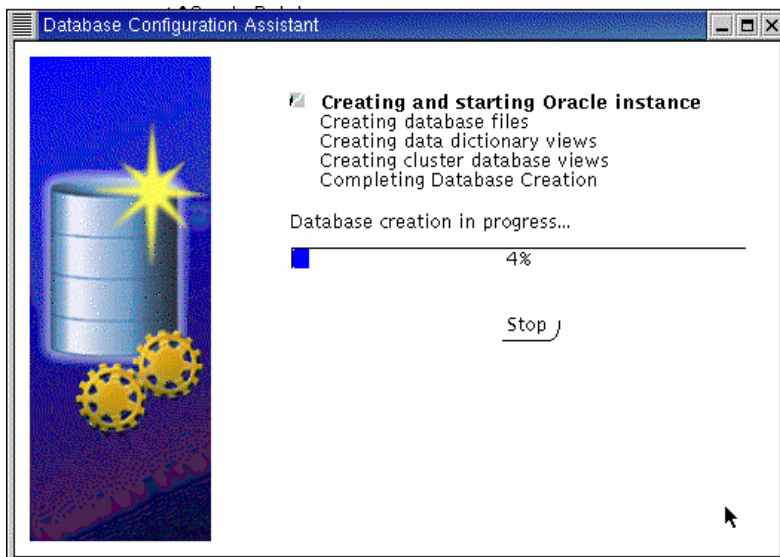
The DBCA Summary window is displayed. Review this information and then click OK.



Once the Summary screen is closed using the OK option, DBCA begins to create the database according to the values specified. We selected to create a template and the scripts for manual creation, so that the DBCA will give us some screens when the work is done.



Then, DBCA starts the database creation.



Once this step finished, you have successfully created the clustered database. It can now be accessed via Oracle SQL*PLUS or other applications designed to work with an Oracle RAC database. You should make sure to execute your start-up scripts at system startup as root and the other as user oracle as shown 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.

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 =test1)(PORT =1521))
)
)
)
SID_LIST_LISTENER =
(SID_LIST =
(SID_DESC =
(SID_NAME =PLSExtProc)
(ORACLE_HOME =/oracle/product/9.0.1)
(PROGRAM =extproc)
)
(SID_DESC =
(ORACLE_HOME =/oracle/product/9.0.1)
(SID_NAME = test1)
)
)
)

```

listener.ora for the first instance

TNSNAMES CONFIGURATION

```

LISTENERS_TEST =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL =TCP)(HOST =x360r1)(PORT =1521))
(ADDRESS = (PROTOCOL =TCP)(HOST =x360r2)(PORT =1521))
)
)
TEST =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL =TCP)(HOST =x360r1)(PORT =1521))
(ADDRESS = (PROTOCOL =TCP)(HOST =x360r2)(PORT =1521))
)
(FAILOVER =yes)
(LOAD_BALANCE =yes)
(CONNECT_DATA =
(SERVICE_NAME =TEST)
)
(FAILOVER_MODE=
(TYPE=SELECT)
(METHOD=BASIC)
(RETRIES=30)
(DELAY=1)
)
)
)
TEST1 =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL =TCP)(HOST =x360r1)(PORT =1521))
)
(CONNECT_DATA =
(SERVICE_NAME =TEST)
)
)

```



```
(INSTANCE_NAME =TEST1)
)
)
TEST2 =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS =(PROTOCOL =TCP)(HOST =x360r2)(PORT =1521))
)
(CONNECT_DATA =
(SERVICE_NAME =TEST)
(INSTANCE_NAME =TEST2)
)
)
```

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 racdb1
racnode1 racinst1 /oracle/product/9.2.0
racnode2 racinst2 /oracle/product/9.2.0
```

Examples of starting and stopping RAC follow:-

```
$ srvctl start database -d racdb1
$ srvctl status database -d arldb
Instance racinst1 is running on racnode1
Instance racinst2 is running on racnode2
```

```
$ srvctl stop database -d racdb1
$ srvctl status database -d arldb
Instance racinst1 is not running on racnode1
Instance racinst2 is not running on racnode2
```

```
$ srvctl start instance -d racdb1 -i racinst1
$ srvctl status instance -d racdb1 -i racinst1
Instance racinst1 is running on racnode1
```

```
$ srvctl status database -d arldb
Instance racinst1 is running on racnode1
Instance racinst2 is not running on racnode2
```

```
$ srvctl stop instance -d racdb1 -i racinst1
```

For further information on `srvctl` see the Oracle9i Real Application Clusters Administration Release 1 (9.0.1) manual.

REFERENCES

- Oracle9i Installation Guide for UNIX Systems: AIX-Based Systems, Compaq Tru64 UNIX, HP 9000 Series HP-UX, Linux Intel and Sun SPARC Solaris
- Oracle9i Installation Checklist for Linux Intel
- Oracle9i Quick Installation Procedure for Linux Intel Part Number
- Oracle9i Real Application Clusters Installation and Configuration
- Oracle/IBM CookBook on installing Oracle 9i RAC on AIX
- Step-by-step Installation Guide from Oracle that is available on Metalink.

DISCLAIMER

This document reflects the EMEA Oracle / IBM Joint Solutions 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 EMEA Oracle / IBM Joint Solutions 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 Oracle / IBM Joint Solutions Center. If you have questions about the contents of this document, please direct them to the EMEA Oracle / IBM Joint Solutions Center (oraclibm@fr.ibm.com).

TRADEMARKS

All brand names are the trademarks of their respective owners.

SPECIAL NOTICE

The material in this document was produced by a joint effort of IBM and Oracle specialists. The material herein is copyrighted by both IBM and Oracle.