



Installing Oracle9i™ Database with Real Application Clusters on IBM xSeries Model 365 running Red Hat Enterprise Linux AS version 3

April 2004

Author:

Gregory Verstraeten (gregory.verstraeten@fr.ibm.com)
with the Oracle/IBM JSC team and IBM/Oracle ICC team

Content

The aim of this document	2
Planning the installation	3
Hardware cabling and Firmware	4
Hardware configuration	4
Cabling	4
Firmware	4
Operating System setup	5
Configuring the storage	8
Software installation	8
Storage usage recommendations for Oracle datafiles	8
Storage configuration	8
Configuring LVM and creating raw devices	9
Installing the Cluster Manager	11
Oracle9i Database Installation	20
Creating a Cluster Database	30
Creating a Database Manually	31
Configuring High Availability and Load Balancing	32
listener configuration	32
tnsnames configuration	32
Administering Real Application Clusters Instances	34
References	35
Disclaimers	35
Trademarks	35



THE AIM OF THIS DOCUMENT

The aim of this document is to provide help to install Oracle9i Database™ with Real Application Clusters on IBM eServer xSeries Model 365 (x365) with FASTT700 running Red Hat Enterprise Linux AS version 3 (RHEL3) through a description of the different steps necessary to run Oracle9i RAC. In this document, we take in 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 documentations

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 taken the installation procedure "Step-By-Step Installation of RAC on Linux" from Metalink, and put it in this document. We have only 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



PLANNING THE INSTALLATION

Get the Red Hat EL 3.0 Advanced Server CDs with the latest release update (Update 1 at the time of writing)

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

[IBM eServer xSeries 365 - Flash BIOS Update version 1.03](#)

[IBM eServer xSeries 365 - Remote Supervisor Adapter II Firmware Update version 1.01 for Linux](#)

[IBM FASTT Fibre Channel Host Adapter BIOS/Firmware version 1.35 for QLA2342](#)

[Fibre Channel Solutions - IBM FASTT Host Adapter non-failover device driver version 6.06.60 for Linux](#)

[Fibre Channel Solutions - IBM FASTT Storage Manager 8.4 Fibre Channel Controller Firmware and NVSRAM for the FASTT 600, FASTT700 and FASTT700 version 05.40.06.01_701](#)

[Fibre Channel Solutions - IBM FASTT Storage Manager Version 8.4 for Linux](#)

[Fibre Channel Solutions - IBM FASTT MSJ Diagnostic and Configuration Utility version 43 for 32-bit Linux](#)

[IBM FASTT Storage Manager v8.4 download files matrix](#)

[IBM TotalStorage FASTT Linux RDAC Software Package version 8.4](#)

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 the latest release of Oracle9i R2 Enterprise Edition CDs (download from <http://otn.oracle.com>)

HARDWARE CABLING AND FIRMWARE

Hardware configuration

IBM eServer x365 each with:
 IBM FAStT Fibre Channel Host Adapter QLA2342
 IBM TotalStorage SAN Switch F16
 IBM FAStT700 Storage Server
 IBM FAStT EXP700 Storage Expansion Unit


Cabling

The chart below shows a 2 nodes configuration.

For more information, check the "Implementing IBM eServer xSeries SANs" Redbook for more information on SAN configurations.

In this chart, only one controller is connected to the SAN switch, the reason is that the FAStT RDAC was not available at the time the paper was written. Both controllers can be connected to the SAN switch, and therefore used by the nodes, if you install the RDAC.

The RDAC is now available at <http://www-307.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-54973>



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>):

x365 BIOS: 1.00
 FAStT FC-2 HBA BIOS: 1.35
 FAStT700 firmware: 8.24

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 FAStT firmware upgrade.

OPERATING SYSTEM SETUP

This chapter will guide you through the steps for the post-installation of RHEL3 to take in account the specifics of the IBM xSeries cluster.

Make sure you install RHEL3 on all nodes with "Development tools" and "kernel development" packages.

Make sure these packages are installed:

```
compat-db-4.0.14-5.i386.rpm
compat-libstdc++-7.3-2.96.123.i386.rpm
compat-libstdc++-devel-7.3-2.96.123.i386.rpm
compat-gcc-7.3-2.96.123.i386.rpm
compat-gcc-c++-7.3-2.96.123.i386.rpm
compat-glibc-7.x-2.2.4.32.5.i386.rpm
openmotif-2.2.2-16.i386.rpm
openmotif-devel-2.2.2-16.i386.rpm
setarch-1.3-1.i386.rpm
```

Follow the following steps on each node in the cluster.

Change the modules order of the devices

Edit `/etc/modules.conf`. The onboard SCSI adapter drivers, `mptbase` with `mptscsih`, have to be the first scsi module to be loaded, they are mapped to `scsi_hostadapter` and `scsi_hostadapter1`. The FASTT adapter comes after, it is mapped to `scsi_hostadapter2`, if you have a second FASTT adapter, it has to be mapped to `scsi_hostadapter3`, etc.

Append the following line to enable up to 128 SCSI devices per adapters

```
options scsi_mod max_scsi_luns=128
alias eth0 tg3
alias eth1 tg3
alias scsi_hostadapter mptbase
alias scsi_hostadapter1 mptscsih
alias scsi_hostadapter2 qla2300
alias usb-controller usb-ohci
options scsi_mod max_scsi_luns=128
```

`/etc/modules.conf`

Configure the network

Stop the network service with the command:

```
service network stop
```

The first onboard 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=10.10.10.1
NETMASK=255.255.255.0
```

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

The second onboard 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`

Recompile the FASTT adapter driver

```
cd /usr/src/linux-2.4
cp ./configs/kernel-2.4.21-i686-smp.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/2.4.21-9.ELsmp/kernel/drivers/addon/qla2200/
```

Recreate the initial RAM disk

```
rm /boot/initrd-2.4.21-9.ELsmp.img
mkinitrd /boot/initrd-2.4.21-9.ELsmp.img 2.4.21-4.EL-smp
```

Reboot the server.

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.

```
10.10.10.1 x365r1
10.10.10.2 x365r2
1.1.1.1 int-x365r1
1.1.1.2 int-x365r2
```

/etc/hosts

Configure the hosts equivalence.

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

```
x365r1
x365r2
int-x365r1
int-x365r2
```

/etc/hosts.equiv

Make the rlogin and rsh services started automatically.

```
chkconfig -level 345 rsh on
```

Create the Oracle user and directory

```
mkdir /oracle
addgroup dba
addgroup oinstall
adduser -g dba -G disk, oinstall -h /home/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=$ORACLE_BASE/9204
export ORACLE_TERM=xterm
export TNS_ADMIN=$ORACLE_HOME/network/admin
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 ORACLE_SID=test1
```

/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*
```

Change the kernel parameters for Oracle

```
# Kernel sysctl configuration file for Red Hat Linux
#
# For binary values, 0 is disabled, 1 is enabled.  See sysctl(8) and
# sysctl.conf(5) for more details.

# Controls IP packet forwarding
net.ipv4.ip_forward = 0

# Controls source route verification
net.ipv4.conf.default.rp_filter = 1

# Controls the System Request debugging functionality of the kernel
kernel.sysrq = 0

# Controls whether core dumps will append the PID to the core filename.
# Useful for debugging multi-threaded applications.
kernel.core_uses_pid = 1

# for a large SGA
kernel.shmmax = 3000000000

# default values
kernel.sem = 250 32000 32 128

# RAC tuning
net.core.rmem_default = 131070
net.core.rmem_max = 131070
net.core.wmem_default = 131070
net.core.wmem_max = 131070
```

/etc/sysctl.conf



CONFIGURING THE STORAGE

Software installation

Install IBM FASTT Storage Manager version 8.24 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 either raw devices or a cluster file system, such as Oracle Cluster File System, PolyServe Matrix Server.

Having datafiles on a cluster file system is obviously more convenient, as you can see the file you are using. Database backups are made easier.

Having datafiles on raw devices is the option to use for performance. Raw devices are used in benchmarks. Backup are made with Oracle RMAN. Logical Volume Manager (LVM), part of the RHEL3 distribution, brings manageability and flexibility to the raw devices management.

In this document, we will go through the configuration of raw devices with LVM.

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, what is the associated linux devices, keep the LVM information and it will you to remember the purpose of each raw devices.

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 name of the device for the LUN as seen by Linux, the fourth is the description of the datafile that will be put in the raw device, the fifth is the size of the datafile, the sixth is the LVM Volume group that contains the device, the seventh is the LVM logical volume, and the last is the raw device name.

RAC raw devices

Array name	LUN	Linux device	Database file	size	VG	Logical Volume	Raw device
Oracle	L1	/dev/sdb	quorum	110 MB	oravg	quorum_lv	/dev/raw/raw100
Oracle	L1	/dev/sdb	svrm	110 MB	oravg	svrm_lv	/dev/raw/raw101

Database raw devices

Array name	LUN	Linux device	Database file	size	VG	Logical Volume	Raw device
Oracle	L1	/dev/sdb	SPFILE	5 MB	oravg	spfile_lv	/dev/raw/raw1
Oracle	L1	/dev/sdb	SYSTEM	400 MB	oravg	system01_lv	/dev/raw/raw2
Oracle	L1	/dev/sdb	USERS	120 MB	oravg	users01_lv	/dev/raw/raw5
Oracle	L1	/dev/sdb	TOOLS	12 MB	oravg	tools01_lv	/dev/raw/raw6
Oracle	L1	/dev/sdb	TEMP	110 MB	oravg	temp01_lv	/dev/raw/raw7
Oracle	L1	/dev/sdb	Control file 1	110 MB	oravg	ctrl01_lv	/dev/raw/raw8
Oracle	L1	/dev/sdb	Control file 2	110 MB	oravg	ctrl02_lv	/dev/raw/raw9
Oracle	L1	/dev/sdb	INDEX	70 MB	oravg	index01_lv	/dev/raw/raw10
Oracle	L1	/dev/sdb	Redo T1 Grp1	120 MB	oravg	redot1g1_lv	/dev/raw/raw11
Oracle	L1	/dev/sdb	Redo T1 Grp2	120 MB	oravg	redot1g2_lv	/dev/raw/raw12
Oracle	L1	/dev/sdb	Redo T2 Grp1	120 MB	oravg	redot2g1_lv	/dev/raw/raw13
Oracle	L1	/dev/sdb	Redo T2 Grp2	120 MB	oravg	redot2g2_lv	/dev/raw/raw14
Oracle	L1	/dev/sdb	Undo T1	312 MB	oravg	undo01t1_lv	/dev/raw/raw15
Oracle	L1	/dev/sdb	Undo T2	312 MB	oravg	undo01t2_lv	/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.

Configuring LVM and creating raw devices

Use the LVM utilities to create partitions on the disk devices.

First of all, identify the SCSI device for the LUN you have created. List the devices available on the system with the following command:

```
[root@x365n1 root]# fdisk -l
Disk /dev/sda: 110.1 GB, 110109917184 bytes
254 heads, 63 sectors/track, 13439 cylinders
Units = cylinders of 16002 * 512 = 8193024 bytes
   Device Boot      Start         End      Blocks   Id  System
/dev/sda1    *           1           13     103981+   83   Linux
/dev/sda2                14          144     1048131   82   Linux swap
/dev/sda3                145        13439     106373295   83   Linux

Disk /dev/sdb: 437.2 GB, 437223161856 bytes
255 heads, 63 sectors/track, 53156 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk /dev/sdb doesn't contain a valid partition table
```

In our example, disk /dev/sda is the array for the internal drives, disk /dev/sdb is the LUN from the FASiT.

Then, initialize the LVM configuration on all nodes with the following command, you should see the associated output:

```
[root@x365n1 root]# vgscan
vgscan -- reading all physical volumes (this may take a while...)
vgscan -- "/etc/lvmtab" and "/etc/lvmtab.d" successfully created
vgscan -- WARNING: This program does not do a VGDA backup of your volume group
```

Define the SCSI device as an LVM physical volume with the following command, you should see the associated output:

```
[root@x365n1 root]# pvcreate /dev/sdb
pvcreate -- physical volume "/dev/sdb" successfully created
```

Create an LVM Volume Group name "oravg" for example containing the defined physical volume, with the following command, you should see the associated output:

```
[root@x365n1 root]# vgcreate oravg /dev/sdb
vgcreate -- INFO: using default physical extent size 4 MB
vgcreate -- INFO: maximum logical volume size is 255.99 Gigabyte
vgcreate -- doing automatic backup of volume group "oravg"
vgcreate -- volume group "oravg" successfully created and activated
```

Create each LVM Logical Volumes with the following command, you should see the associated output:

```
[root@x365n1 root]# lvcreate --size 100M oravg --name quorumlv
lvcreate -- doing automatic backup of "oravg"
lvcreate -- logical volume "/dev/oravg/quorumlv" successfully created
```

Run `vgscan` and `lvscan` on all nodes to make sure that all nodes have a consistent view of the volume groups and logical volumes definitions.

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/raw100 /dev/oravg/quorumlv
/dev/raw/raw101 /dev/oravg/srvconfiglv
```

`/etc/sysconfig/rawdevices`

Start the `rawdevices` service on all nodes 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
```





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. Install the Cluster Manager using the Oracle Universal Installer.

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 the first CD-ROM. You can also copy the content of the CDs in the stage directory on the hard disks. So you will not be prompted to switch CDs in the CD-ROM drive. Use a directory structure as follow:

```
/<mount point>/Disk1, /<mount point>/Disk2, /<mount point>/Disk3
```

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

Make sure your display is set accordingly to your environment.

If you are running VNC, the variable DISPLAY has to be set to ":1".

If you are using Xwindow from another workstation (named PC1 for instance), the variable DISPLAY has to be set to "PC1:0.0".

Apply the Oracle patch 3006854 (p3006854_9204_LINUX) on all nodes, this creates the file /etc/ld.so.preload.

Otherwise, at the next step you will get the following error message:

```
[oracle@x365n1 Disk1]$ Initializing Java Virtual Machine from
/tmp/OraInstall2004-04-08_03-50-02PM/jre/bin/java. Please wait...
Error occurred during initialization of VM
Unable to load native library: /tmp/OraInstall2004-04-08_03-50-
02PM/jre/lib/i386/libjava.so: symbol __libc_wait, version GLIBC_2.0 not defined
in file libc.so.6 with link time reference
```

Then you can launch the Oracle Universal Installer with "<cdrom_mount_point>/runInstaller" or "./runInstaller" from where you have mounted the CD-ROM.

Welcome screen.

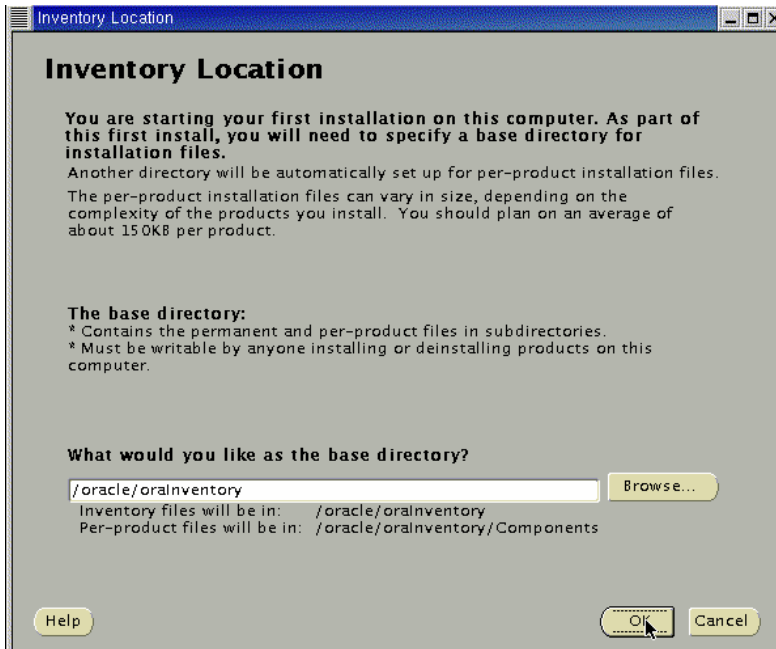
The Oracle Universal Installer loads the Java Runtime Environment (you may receive some messages on your terminal window regarding the fonts but this is normal, because the Java Runtime Environment may not find all the fonts it needs and then it substitutes them with other), and then shows you a welcome screen



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 locations, keep in mind you should be coherent about this location on the future.



UNIX Group screen.

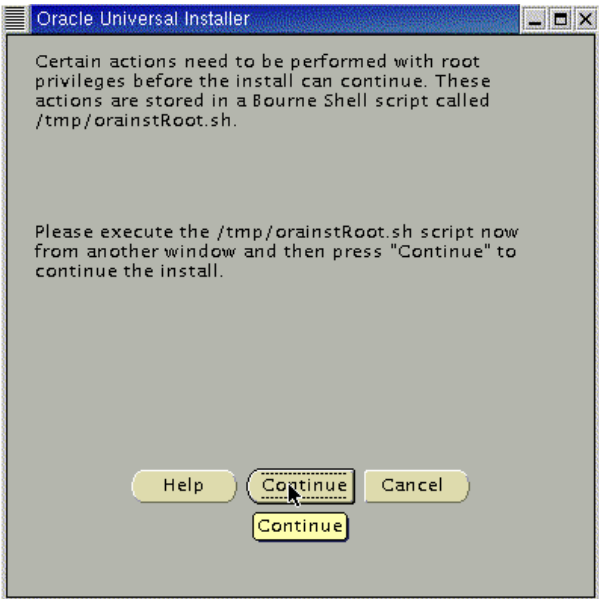
The next screen will prompt you for the UNIX group that the oracle user will use during the installation. This value should be "dba" or "oinstall" if you followed this document to install Oracle 9i RAC, or if you use the Step by Step procedure from Oracle that is available at Metalink's site. So, enter "oinstall" and once again, simply click on "Next" to continue. This screen will look like this:





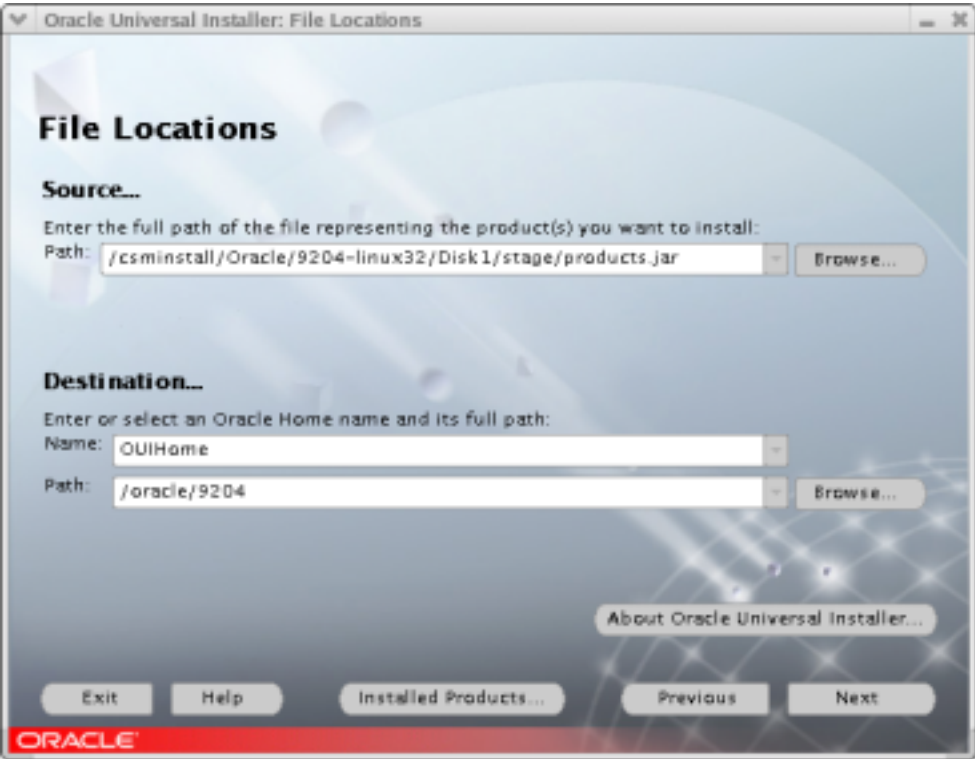
Shell script as root.

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 to let the installer work properly.



File locations screen.

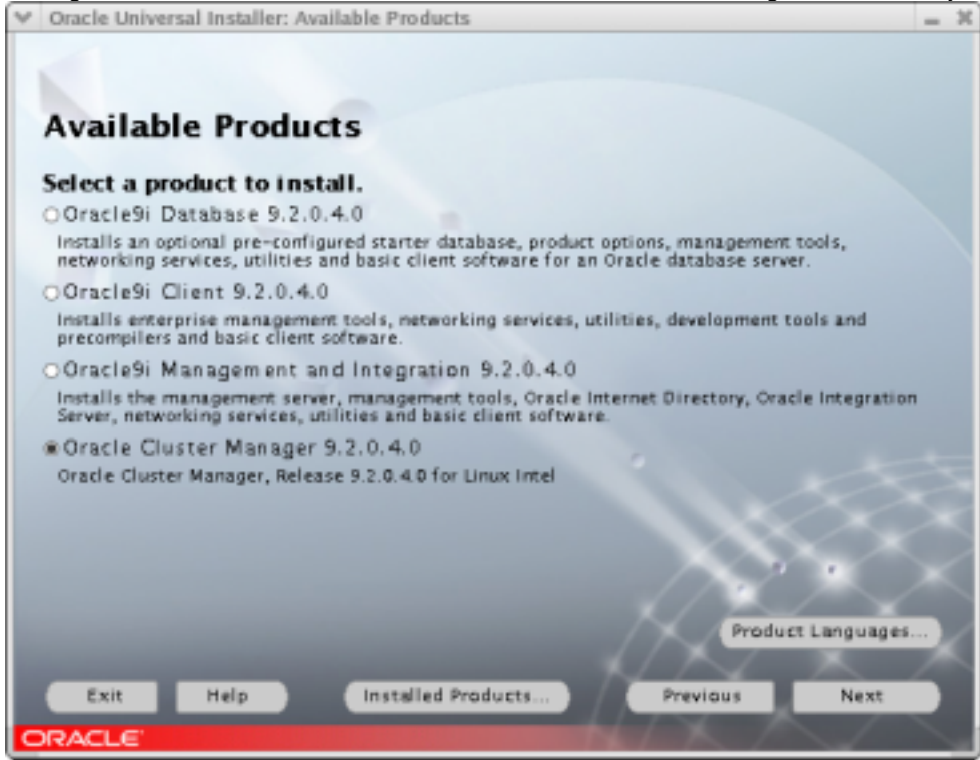
Once the shell script executed, you click “OK” and the File Location screen should appear. This screen asks you for some information about the source and destination paths for you installations. Do not change the source location. For the destination location, verify that the value that you have pre-selected 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 one to make Oracle Universal Installer configure the Real Application Clusters properly. In our case, we used two IBM eServer xSeries model 365 named “x365n1” and “x365n2”. Below is the configuration screen.





Oracle Universal Installer: Public Node Information

Public Node Information

Please enter the public node names in the network

Public Node 1:

Public Node 2:

Public Node 3:

Public Node 4:

Public Node 5:

Public Node 6:

Public Node 7:

Exit Help Installed Products... Previous Next

ORACLE

Private node information screen

On the next screen configure the Private Node addresses that will be used for the RAC 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



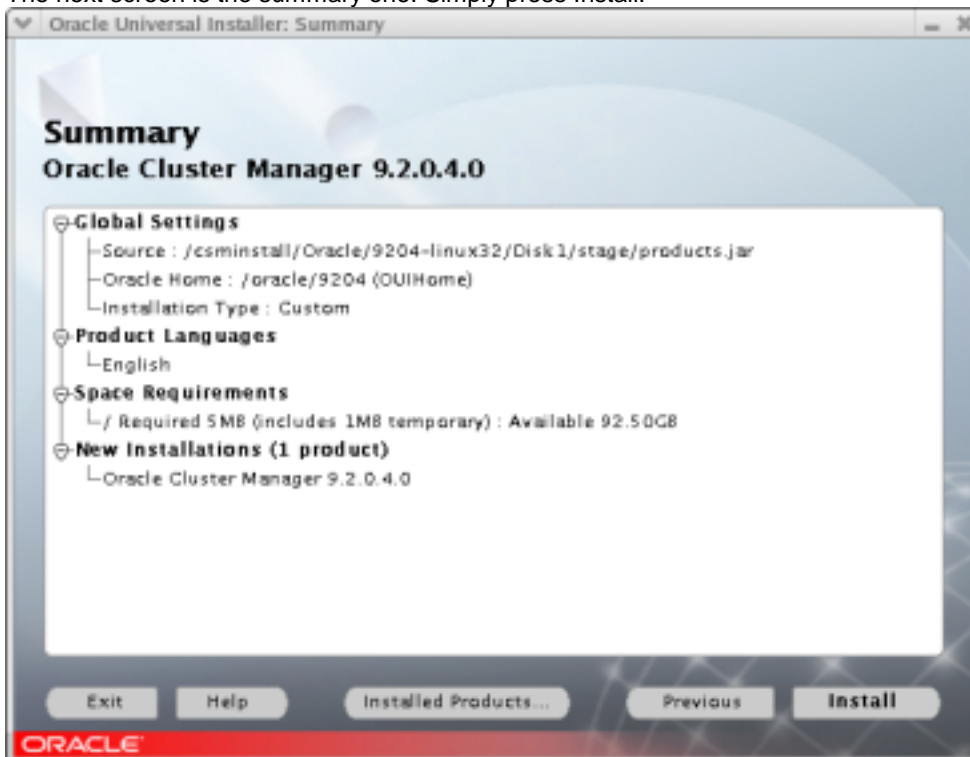
Quorum Disk Information

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/raw100`. Click Next to continue.



Summary Screen

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

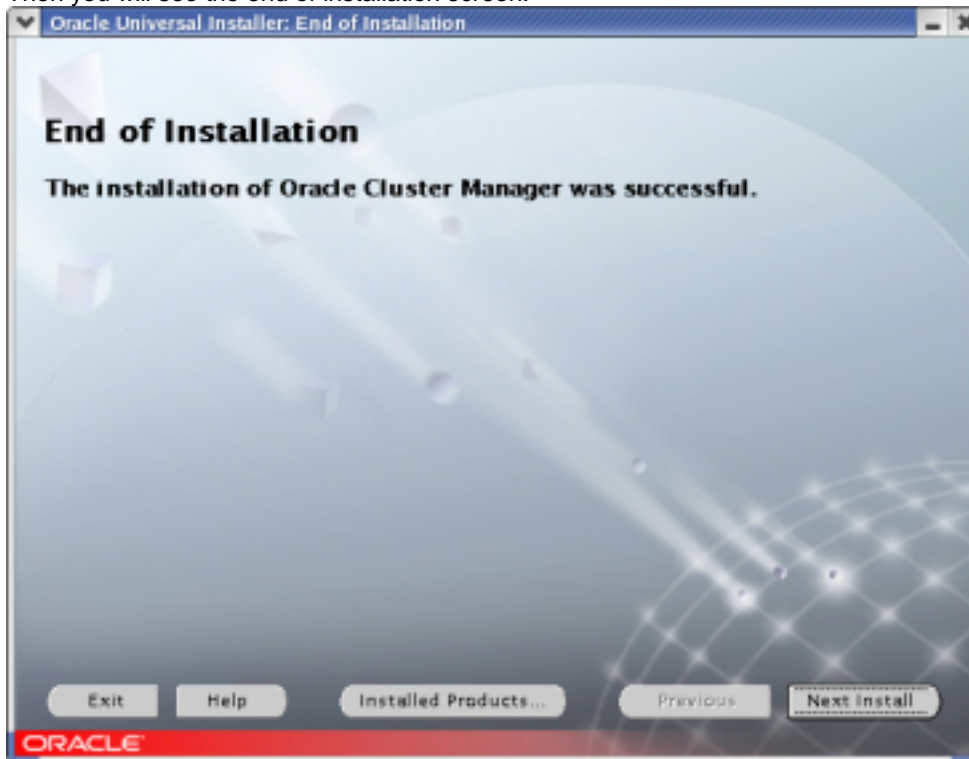


Cluster Manager Installation screen

You will then get a progress window, the installation is quick (less than 5 minutes).



Then you will see the end of installation screen.



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. You have to check the configuration file for the Cluster Manager, as root, edit `$ORACLE_HOME/oracm/admin/cmcfg.ora`, and do the appropriate changes (see the file below). The line with the `KernelModuleName` parameter has to be added.

```
ClusterName=Oracle Cluster Manager, version 9i
MissCount=20
PrivateNodeNames=int-x365n1 int-x365n2
PublicNodeNames=x365n1 x365n2
ServicePort=9998
CmDiskFile=/dev/raw/raw100
KernelModuleName=hangcheck-timer
```

`$ORACLE_HOME/oracm/admin/cmcfg.ora`

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

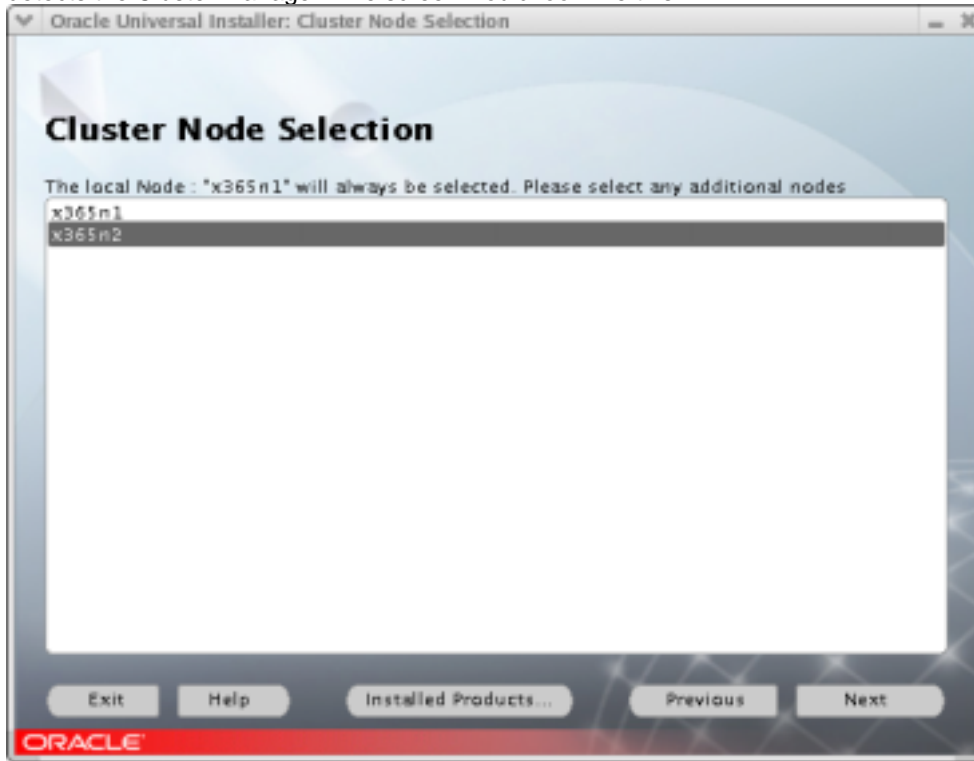
```
export ORACLE_HOME=/oracle/9204
$ORACLE_HOME/oracm/bin/ocmstart.sh
```

This will start the `hangcheck-timer` kernel loadable module and the `oracm` process which is the Cluster Manager. Verify you started it on each node with the `ps -ef` command and check the `oracm` log: `$ORACLE_HOME/oracm/log/cm.log`

ORACLE9i DATABASE INSTALLATION

Installing Oracle9i

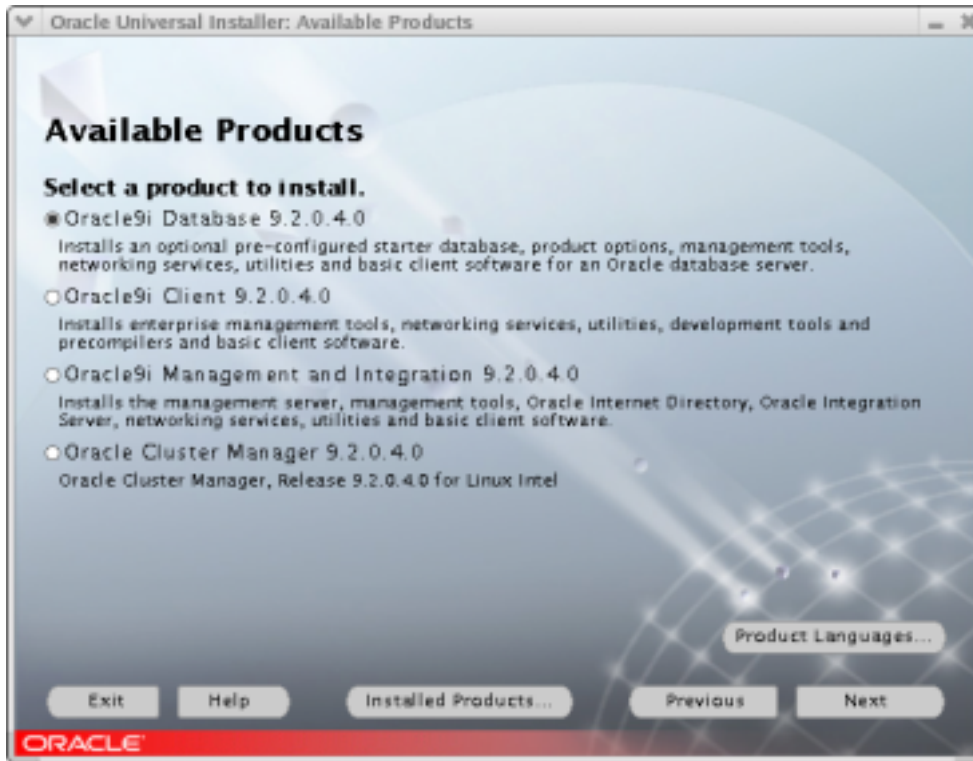
Start the Oracle Universal Installer again. If you click Next on the Welcome screen you will be prompted for the cluster nodes on which you would like to install the Database. This is because the Oracle Universal Installer detects the Cluster Manager. The screen would look like this:



Select all the nodes you will 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

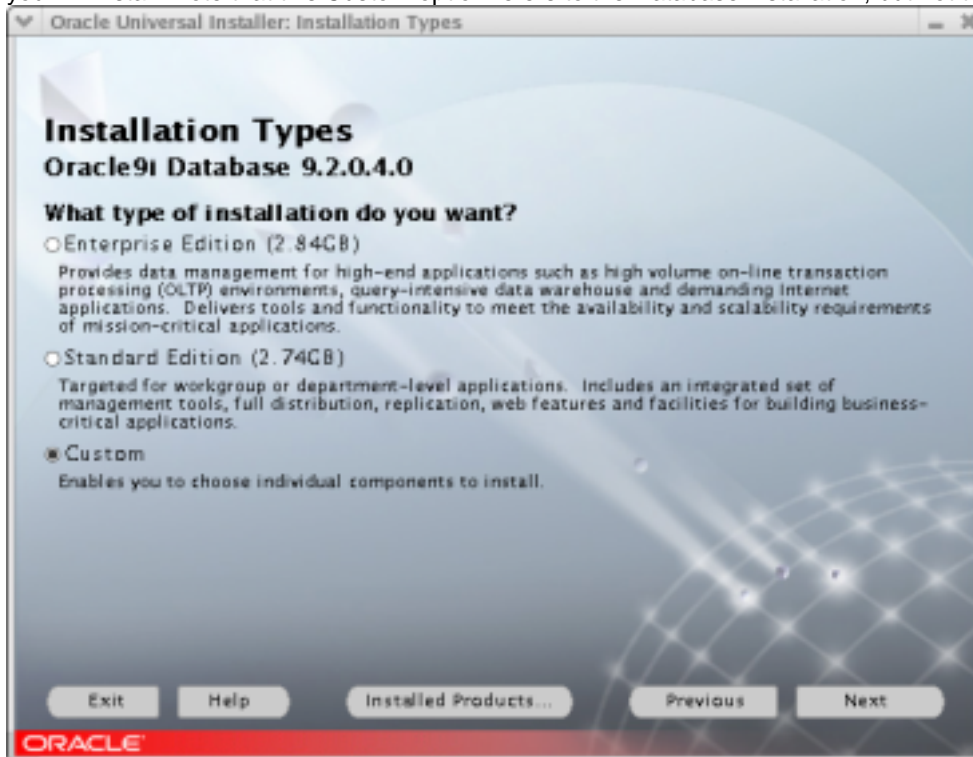
Select now on the Available Products screen the "Oracle Database" and click Next to continue as shown on the figure.



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

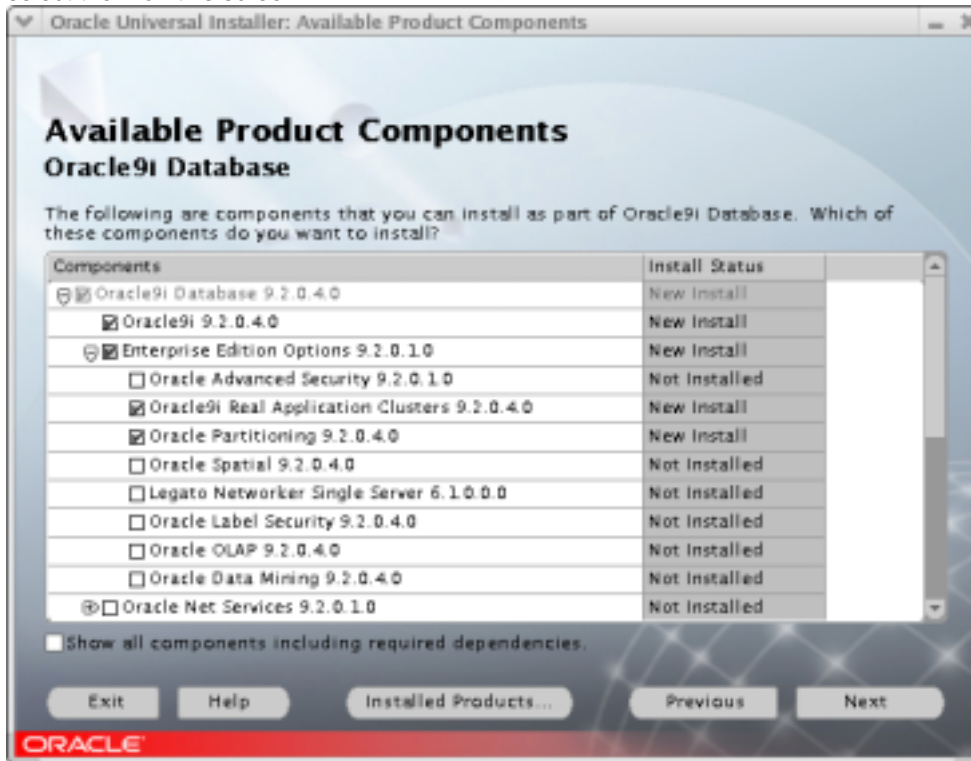
Installation Type screen.

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.



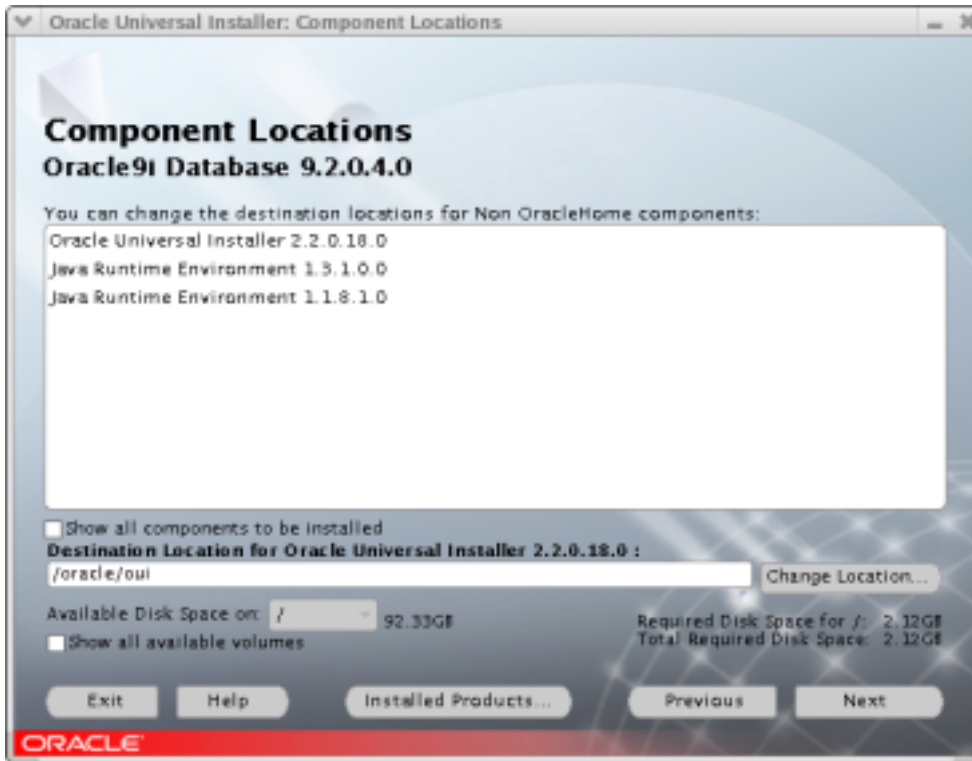
Components Selection Screen

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 show information about the components locations. Do not modify the values and click Next.



Shared Configuration File screen

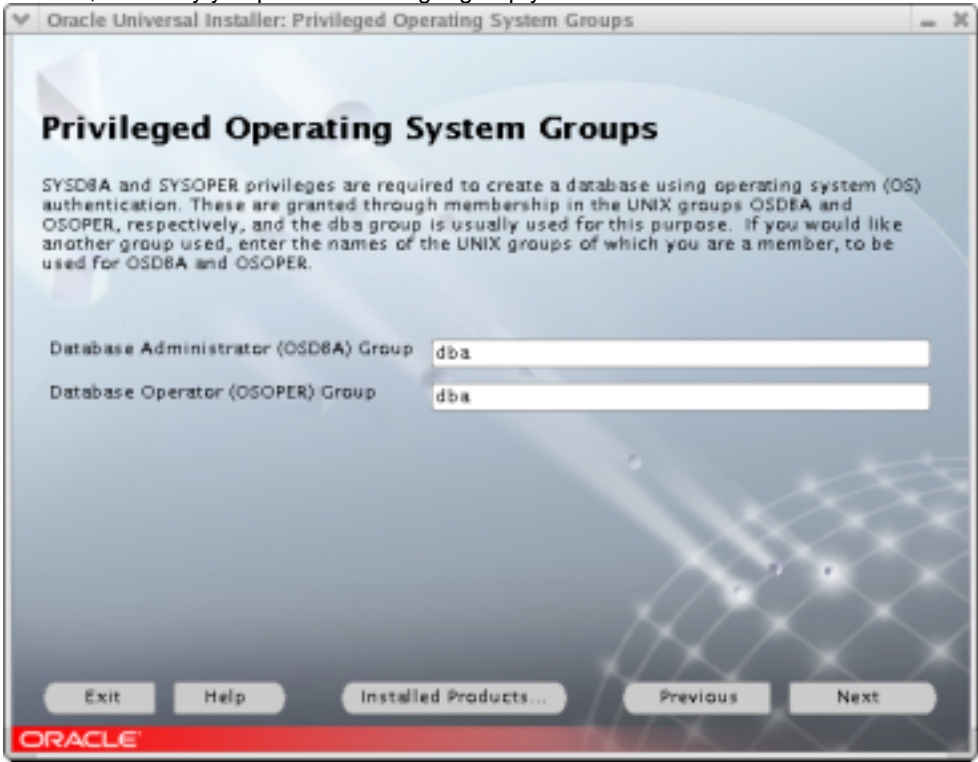
You will also be prompted for the raw device where you will have the shared configuration of your cluster. Enter your value there and click Next to continue.





Group Selection screen

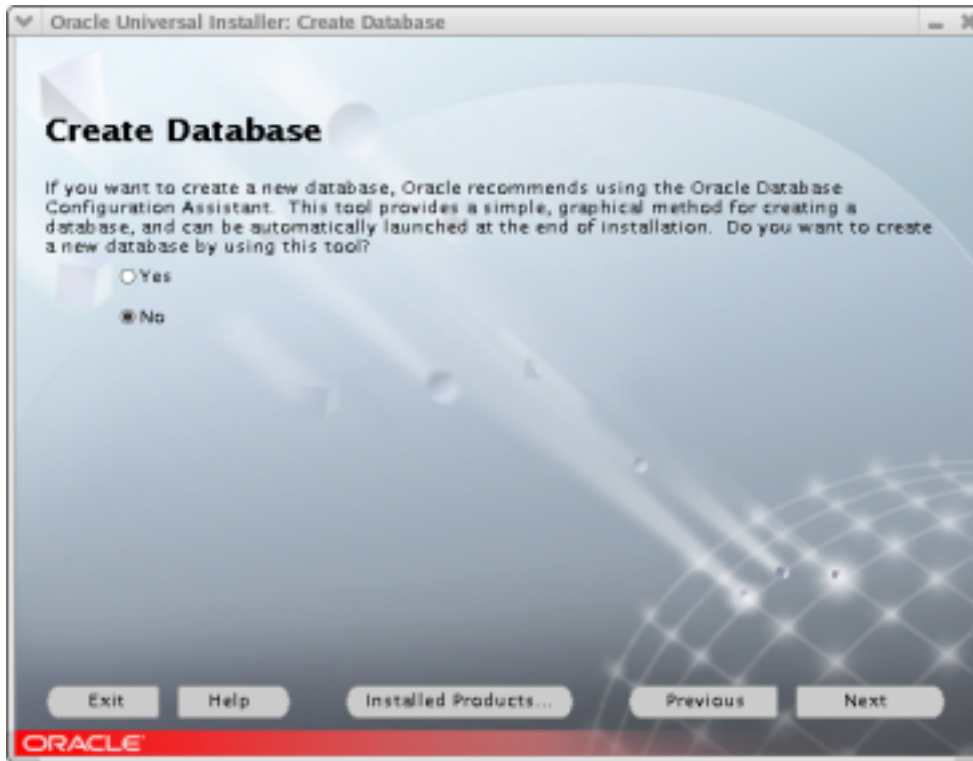
Then, select the group you will need for your database operator and administrator. Usually, these values are set to "dba", but verify you put there the right group you will use. Click Next to continue.



Database Creation screen

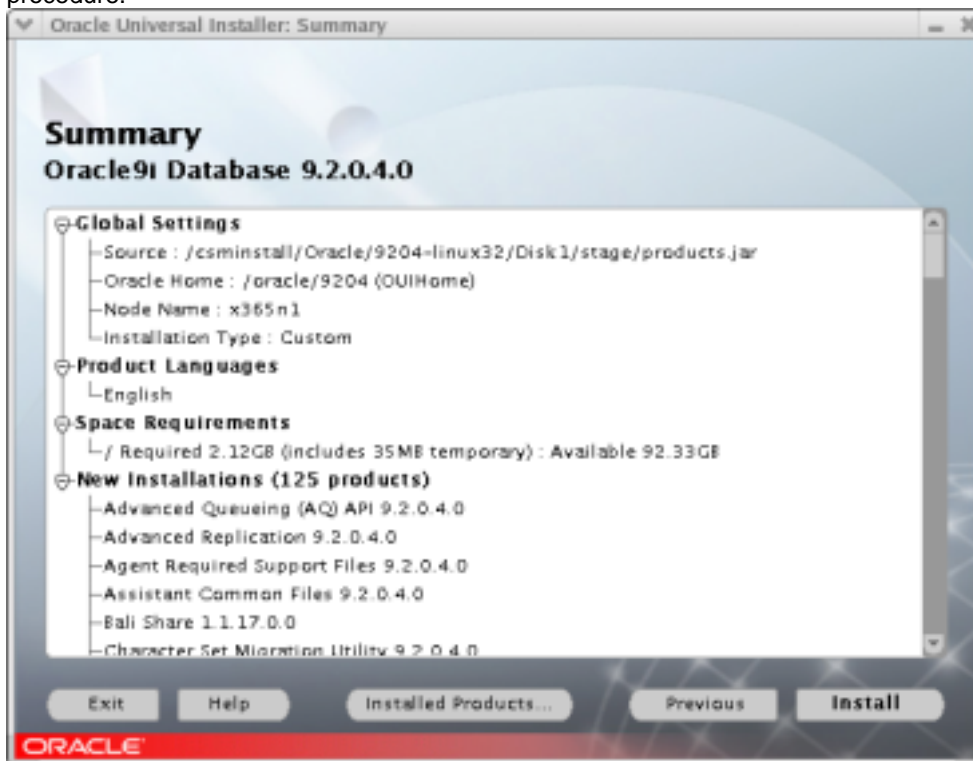
You will then be prompted for database creation. We will not create it right now, so select "No" and then Next to continue. If you wish to create a database during the installation, you should select "Yes" and follow the configuration steps. We will create the database later.






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:



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.

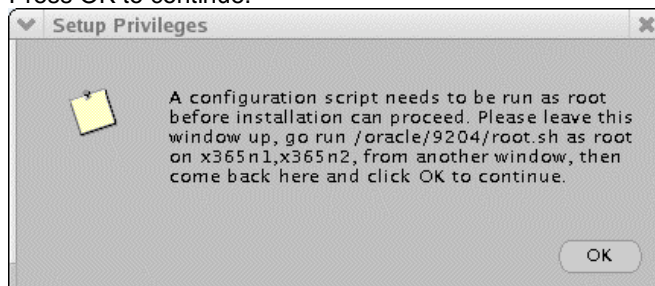
Execute the shell script as root in a separate xterm accepting the default values:

```
[root@x365n1 root]# /oracle/9204/root.sh
Running Oracle9 root.sh script...
The following environment variables are set as:
ORACLE_OWNER= oracle
ORACLE_HOME= /oracle/9204
```

```
Enter the full pathname of the local bin directory: [/usr/local/bin]:
Copying dbhome to /usr/local/bin ...
Copying oraenv to /usr/local/bin ...
Copying coraenv to /usr/local/bin ...
```

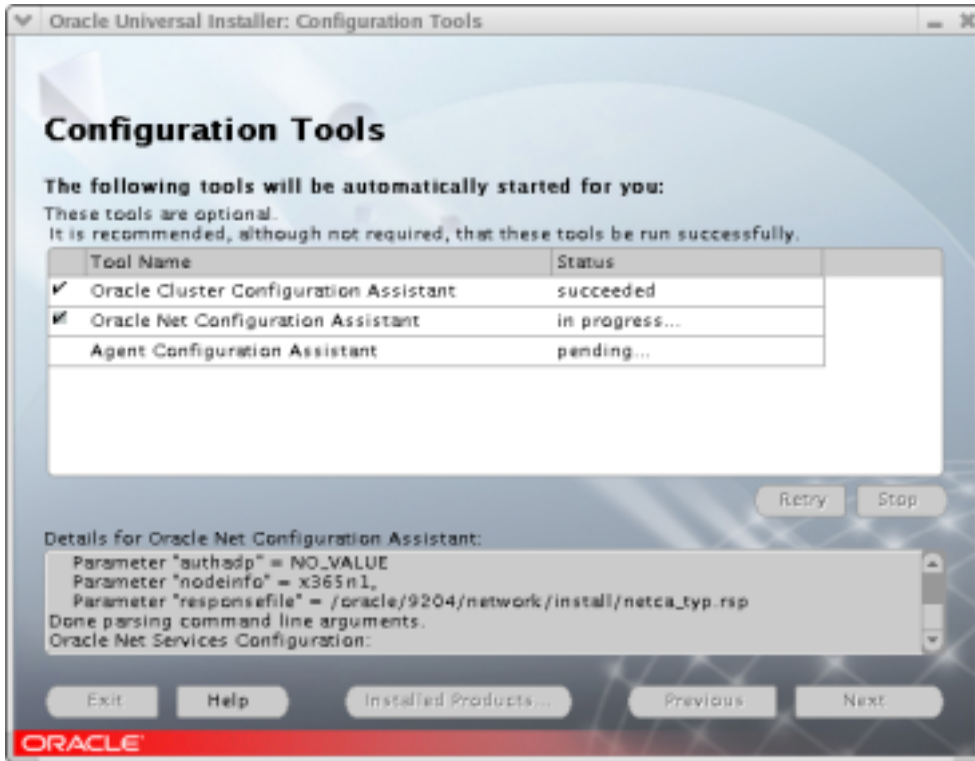
```
Creating /etc/oratab file...
Adding entry to /etc/oratab file...
Entries will be added to the /etc/oratab file as needed by
Database Configuration Assistant when a database is created
Finished running generic part of root.sh script.
Now product-specific root actions will be performed.
```

Press OK to continue.





Then the configuration procedure will start. First, the Oracle Universal Installer will automatically configure the cluster, configure the network listener, and finally configure the agent.

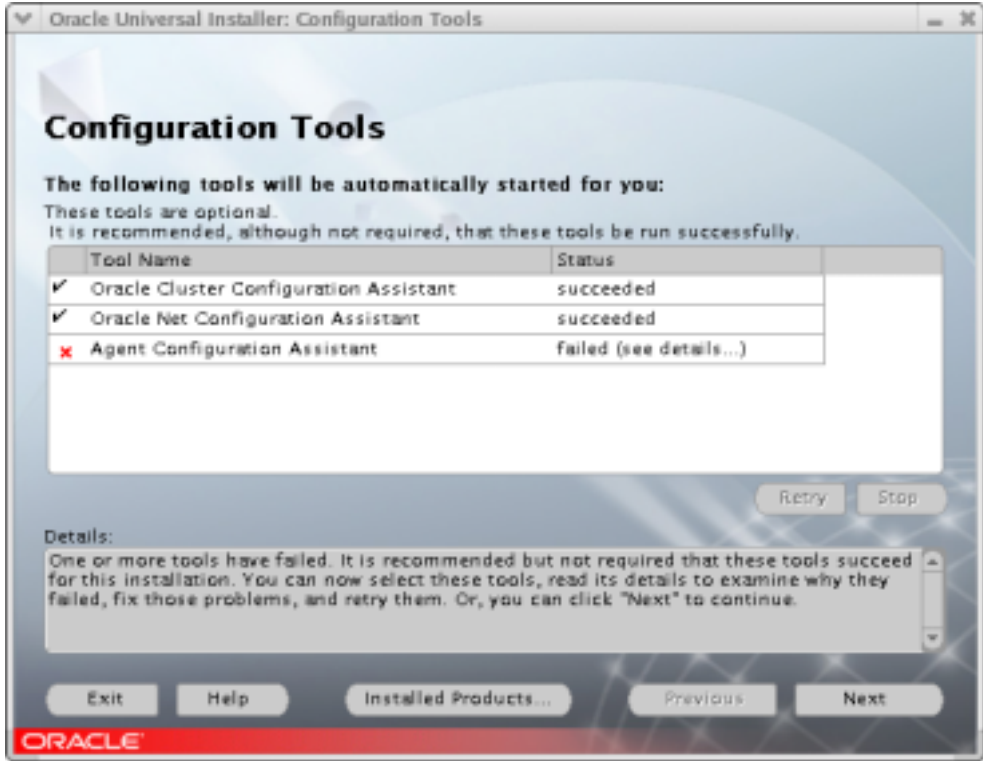
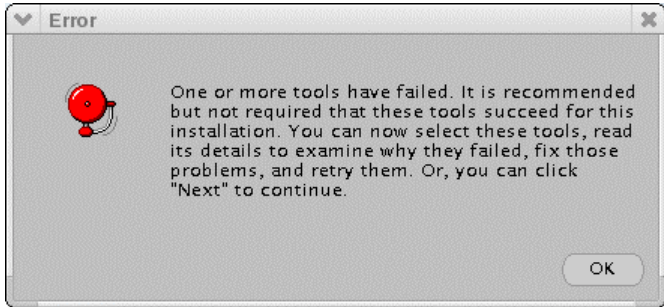


For the Oracle Net configuration, you will be prompted for details. Check "Perform typical configuration" and click Next.



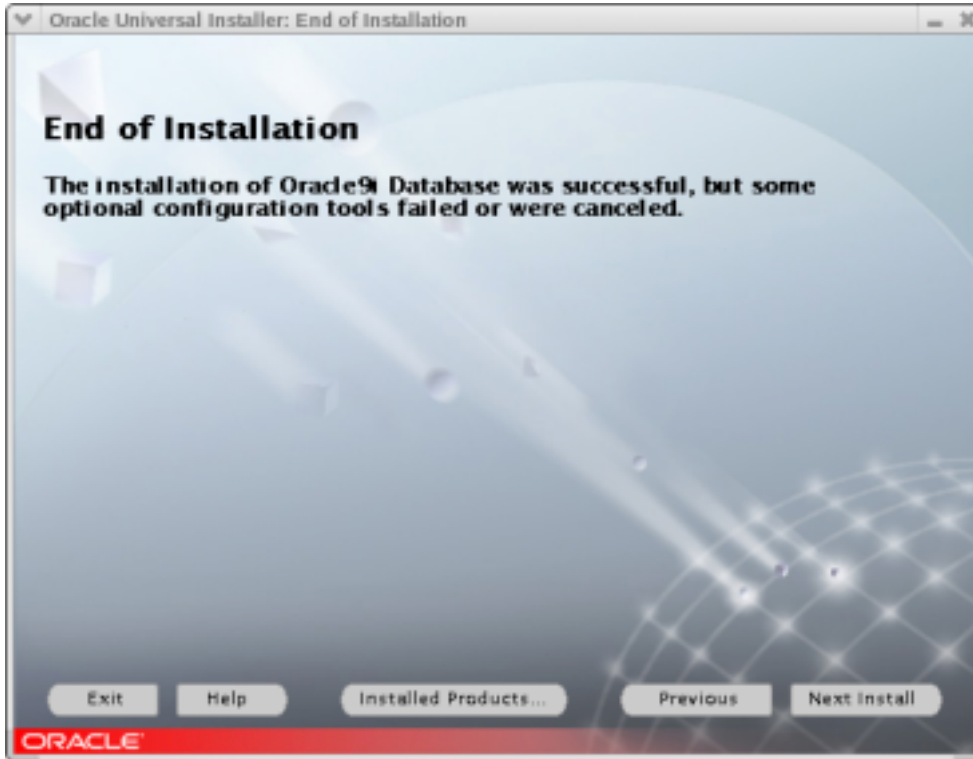
The Agent configuration will fail. This is a known issue. You can ignore it, and later on, fix it with the Oracle patch p3119415_9204_LINUX. Click Ok.





The Oracle Universal Installer will show you the End of Installation screen and warns you that you have got some components that were not configured (Oracle Agent). Exit the installer by clicking "Exit".





Oracle Global Service Daemon startup

Start the Global Services Daemon on all nodes of the cluster.

As oracle user, on each node, type:

```
gsdctl start
```

You should have the following output

```
[oracle@x365n1 Disk1]$ gsdctl start
Successfully started GSD on local node
```

This has to be done each time the nodes boot up.

Oracle Listener startup

Start the Oracle listener on all nodes of the cluster.

As oracle user, type "lsnrctl start" on each node.

Whenever you change the configuration of the listener, you will have to restart it:

```
lsnrctl stop ; lsnrctl start
```

This has to be done each time the nodes boot up.

CREATING A CLUSTER DATABASE

You may have your own database creation script, otherwise the database creation can be done with the DBCA tool.

DBCA means Database Creation Assistant and it is a graphical tool to make easier the database creation which 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, named `mappingDBCA.cfg`, which 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
```

DBCA is started with the `dbca` executable.

We will not go through the detail of DBCA screens. The process is pretty straight forward.

Depending on the level of Oracle9i, you may experience compatibility problems with Red Hat EL 3.0 and the new thread architecture (NTPL). In this case, you will have to set the variable `LD_ASSUME_KERNEL` to "2.4.1" before starting any Oracle assistants and installers.

CREATING A DATABASE MANUALLY

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

- o on one node :


```
srvctl add database -d <db_name> -o <oracle_home>
```
- o for each instance of the database :


```
srvctl add instance -d <db_name> -i <SID> -n <node_name>
```
- o on each node, check the configuration :


```
srvctl config
```
- o on each node, create (or update) oratab file (in /etc directory) with the following line :


```
<db_name>:<${ORACLE_HOME}>:N
```
- o create the udump, cdump and bdump directories.
- o set the SID in .bash_profile of oracle user
- o create the init<SID>.ora file for each node. For a RAC environment, the following parameters have to be set in the init.ora for each instances (these are the setting for the first instance, test1):

```
*.cluster_database = TRUE
*.cluster_database_instances = 4
test1.undo_tablespace=undots1
test1.instance_name=bench
test1.instance_number=1
test1.thread=1
test1.local_listener=LISTENER_TEST1
```

- o create the database creation script from one node, here is an example:

```
CREATE DATABASE test
  CONTROLFILE REUSE
  LOGFILE
  GROUP 1 ('/dev/raw/raw40') SIZE 100 M,
  GROUP 2 ('/dev/raw/raw41') SIZE 100 M,
  GROUP 3 ('/dev/raw/raw42') SIZE 100 M
  MAXLOGFILES 30
  MAXLOGMEMBERS 3
  MAXDATAFILES 500
  MAXINSTANCES 8
  CHARACTER SET WE8ISO8859P1
  NATIONAL CHARACTER SET UTF8
  DATAFILE '/dev/raw/raw30' SIZE 1000 M REUSE
  DEFAULT TEMPORARY TABLESPACE temp TEMPFILE '/dev/raw/raw31' SIZE 1000
  M REUSE
  UNDO TABLESPACE undotbs1 DATAFILE '/dev/raw/raw32' SIZE 1000 M REUSE
  AUTOEXTEND OFF;
```

- o On all nodes, create a password file. Under \$ORACLE_HOME/dbs execute:


```
orapwd file=orapw<SID> password=###.
```
- o On the first instance, for each other instances, you will have to create instances specific Undo tablespaces and redo logs. Finally, you will have to enable the instances. Here is an example for a second instance:

```
CREATE UNDO TABLESPACE undotbs2 DATAFILE '/dev/raw/raw33' SIZE 1000 M
  AUTOEXTEND OFF;
ALTER DATABASE ADD LOGFILE THREAD 2
  GROUP 4 (/dev/raw/raw50) SIZE 100 M,
  GROUP 5 (/dev/raw/raw51) SIZE 100 M,
  GROUP 6 (/dev/raw/raw52) SIZE 100 M;
ALTER DATABASE ENABLE PUBLIC THREAD 2;
```

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 =x365n1) (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

The file \$ORACLE_HOME/network/admin/tnsnames.ora has to be configured to enable load balancing and Transparent Application Failover. See an example below:

```

LISTENERS_TEST =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL =TCP) (HOST =x365r1) (PORT =1521))
      (ADDRESS = (PROTOCOL =TCP) (HOST =x365r2) (PORT =1521))
    )
  )
TEST =
  (DESCRIPTION =
    (LOAD_BALANCE=yes)
    (ADDRESS= (PROTOCOL=TCP) (HOST=x335n1) (PORT=1521))
    (ADDRESS= (PROTOCOL=TCP) (HOST=x335n2) (PORT=1521))
    (FAILOVER=on)
    (CONNECT_DATA=
      (SERVICE_NAME=TEST)
    )
    (FAILOVER_MODE=
      (TYPE=SELECT)
      (METHOD=BASIC)
      (RETRIES=30)
      (DELAY=1)
    )
  )
TEST1 =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL =TCP) (HOST =x365r1) (PORT =1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME =TEST)
      (INSTANCE_NAME =TEST1)
    )
  )

```



```
)  
)  
TEST2 =  
  (DESCRIPTION =  
    (ADDRESS_LIST =  
      (ADDRESS = (PROTOCOL =TCP) (HOST =x365r2) (PORT =1521))  
    )  
    (CONNECT_DATA =  
      (SERVICE_NAME =TEST)  
      (INSTANCE_NAME =TEST2)  
    )  
  )  
)  
LISTENER_TEST1 =  
  (DESCRIPTION=  
    (ADDRESS_LIST=  
      (ADDRESS=(PROTOCOL=TCP) (HOST=x365n1) (PORT=1521))  
    )  
  )  
LISTENER_TEST2 =  
  (DESCRIPTION=  
    (ADDRESS_LIST=  
      (ADDRESS=(PROTOCOL=TCP) (HOST=x365n2) (PORT=1521))  
    )  
  )  
)
```

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 article, `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 test  
$ srvctl status database -d test  
Instance test1 is running on x365n1  
Instance test2 is running on x365n2
```

```
$ srvctl stop database -d test  
$ srvctl status database -d test  
Instance test1 is not running on x365n1  
Instance test2 is not running on x365n2
```

```
$ srvctl start instance -d test -i test1  
$ srvctl status instance -d test -i test1  
Instance test1 is running on x365n1
```

```
$ srvctl status database -d test  
Instance test1 is running on x365n1  
Instance test2 is not running on x365n2
```

```
$ srvctl stop instance -d test -i test1
```

For further information on `srvctl` see the Oracle9i Real Application Clusters Administration Release 2 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.

Disclaimers

(c) 2004 IBM Corp. All rights reserved.

Visit www.ibm.com/pc/safecomputing periodically for the latest information on safe and effective computing. Warranty Information: For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven. This publication was developed for products and services offered in the United States. IBM may not offer the products, services or features discussed in this document in other countries. Information is subject to change without notice. Consult your local IBM representative for information on offerings available in your area.

This document reflects the Oracle/IBM Joint Solutions Center 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 Oracle/IBM Joint Solutions Center, IBM Product and Solutions Support Center and others.

The examples given in this paper are hypothetical examples of how a customer can use the products described herein and examples of potential cost or efficiency savings are not based on any actual case study. There is no guarantee of comparable results. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

MB, GB, and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated. Some numbers given for storage capacities give capacity in native mode followed by capacity using data compression technology.

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 Oracle/IBM Joint Solutions Center (oraclibm@fr.ibm.com).

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM®
IBM.COM™
xSeries™

For a list of additional IBM trademarks, please see <http://www.ibm.com/legal/copytrade.shtml>



Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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

