ORACLE TECHNOLOGY NETWORK

DBA: Linux

Install Oracle RAC 10*g* on Oracle Enterprise Linux Using VMware Server

by Vincent Chan

Evaluate Oracle Real Application Clusters 10*g* Release 2 on Oracle Enterprise Linux for free, using virtual machines

Published January 2007

DOWNLOAD

- VMware Server 1.0.1
- Oracle Enterprise Linux 4, Kernel 2.6.9-42.0.0.0.1.ELsmp
- Oracle Database 10g Release 2 for Linux x86 & Oracle Clusterware

TAGS

vmware, linux, opensource, All

For educational/evaluation purposes only; neither Oracle nor any other vendor will support this configuration

Introduced at Oracle OpenWorld 2006 in October, Oracle Unbreakable Linux aims to offer enterprise-class support services for Red Hat Linux, quicker bug fixes, and a significantly lower support prices. And Oracle's own Enterprise Linux, which is based on Red Hat Advanced Server Release 4 (Update 4) with additional bug fixes, is freely available for download.

As a side effect, it is now possible to evaluate Oracle Real Application Clusters (RAC) 10*g* running on Red Hat Advanced Server on home computers at no cost whatsoever, using VMware Server—a free virtualization environment from VMware.

VMware Server allows you to run multiple operating systems on a single physical machine. Each virtual machine is a self-contained operating environment with its own set of virtual components such as disk, processor, and memory. Virtualization technology is beneficial in a computing environment where software can be developed and tested in isolation on the same physical host machine to prevent data or software corruption. VMware software is widely used in server consolidation to reduce total cost of ownership and to accelerate application development and testing cycles.

In this guide, you will learn how to install and configure two nodes on running Oracle RAC 10*g* Release 2 on Enterprise Linux and VMware Server. Note that this guide is intended for educational/evaluation purposes only; neither Oracle nor any other vendor will support this configuration.

The guide is structured into the following sections:

1. Hardware Requirements and Overview

In this guide, you will install a 32-bit guest Linux operating system. A 64-bit guest operating system is supported only on the following 64-bit processors running on the host machines:

- AMD Athlon 64, revision D or later
- AMD Opteron, revision E or later
- AMD Turion 64, revision E or later
- AMD Sempron, 64-bit-capable revision D or later Intel EM64T VT-capable processors

If you decide to install a 64-bit guest operating system, verify that your processor is listed above. You would also have to verify that Virtualization Technology (VT) is enabled in your BIOS. A few mainstream manufacturers disable the field by default. Additional information on processor compatibility is available here.
To verify if your processor is supported, download the processor check compatibility tool from VMware.

Allocate a minimum of 700MB of memory to each virtual machine; reserve a minimum of 30GB of disk space for all the virtual machines.

An overview of the host operating system environment:

Host Name	OS	Processor	Memory	Disk	Network Card
pacu	Windows XP Professional Service Pack 2 (32-bit)	Intel Pentium 4 550, 3.4MHz, HT	2 GB DDR2 SDRAM, 533 MHz	250 GB, Ultra ATA/133, 7200 RPM	Intel Pro/1000 MT

An overview of guest operating system environment:

Host Name	OS	Processor	Memory
rac1	Oracle Enterprise Linux 4 (32-bit)	1	700 MB
rac2	Oracle Enterprise Linux 4 (32-bit)	1	700 MB

An overview of the virtual disk layout:

Virtual Disk on Host	Virtual Disk on	Virtual Device	Size (MB)	Description

	Guest	Node		
d:\vm\rac\localdisk.vmdk	/dev/sda1 /dev/sda2 /dev/sda3	SCSI 0:0	20	"/" mountpoint Swap space Oracle binaries
d:\vm\rac\sharedstorage\ocfs2disk.vmdk	/dev/sdb	SCSI 1:0	512	OCFS2 disk
d:\vm\rac\sharedstorage\asmdisk1.vmdk	/dev/sdc	SCSI 1:1	3072	ASM disk group 1
d:\vm\rac\sharedstorage\asmdisk2.vmdk	/dev/sdd	SCSI 1:2	3072	ASM disk group 1
d:\vm\rac\sharedstorage\asmdisk3.vmdk	/dev/sde	SCSI 1:3	2048	ASM flash recovery area

(To configure shared storage, the guest OS should not share the same SCSI bus with the shared storage. Use SCSI0 for the guest OS and SCSI1 for the shared disks.)

An overview of the RAC database environment:

Host Name	ASM Instance Name	RAC Instance Name	Database Name	Database File Storage	OCR & Voting Disk
rac1	+ASM1	devdb1	devdb	ASM	OCFS2
rac2	+ASM2	devdb2	devdb	ASM	OCFS2

You'll install the Oracle Home on each node for redundancy. The ASM and Oracle RAC instances share the same Oracle Home on each node.

2. Configure the First Virtual Machine

To create and configure the first virtual machine, you will add virtual hardware devices such as disks and processors. Before proceeding with the install, create the windows folders to house the virtual machines and the shared storage.

```
D:\>mkdir vm\rac\rac1
D:\>mkdir vm\rac\rac2
D:\>mkdir vm\rac\sharedstorage
```

Double-click on the VMware Server icon on your desktop to bring up the application:

1. Press CTRL-N to create a new virtual machine.
2. New Virtual Machine Wizard: Click on **Next**.
3. Select the Appropriate Configuration:

   a. Virtual machine configuration: Select **Custom**.
4. Select a Guest Operating System:
   a. Guest operating system: Select **Linux**.
   b. Version: Select **Red Hat Enterprise Linux 4**.
5. Name the Virtual Machine:
   a. Virtual machine name: Enter “rac1.”
   b. Location: Enter “d:\vm\rac\rac1.”
6. Set Access Rights:
   a. Access rights: Select **Make this virtual machine private**.
7. Startup / Shutdown Options:
   a. Virtual machine account: Select **User that powers on the virtual machine**.
8. Processor Configuration:
   a. Processors: Select One.
9. Memory for the Virtual Machine:
   a. Memory: Select **700MB**.
10. Network Type:
   a. Network connection: Select **Use bridged networking**.
11. Select I/O Adapter Types:
   a. I/O adapter types: Select **LSI Logic**.
12. Select a Disk:
   a. Disk: Select **Create a new virtual disk**.
13. Select a Disk Type:
   a. Virtual Disk Type: Select **SCSI (Recommended)**.
14. Specify Disk Capacity:
   a. Disk capacity: Enter “20GB.”
   b. Deselect **Allocate all disk space now**. To save space, you do not have to allocate all the disk space now.
15. Specify Disk File:
   a. Disk file: Enter “localdisk.vmdk.”
   b. Click on **Finish**.

Repeat steps 16 to 24 to create four virtual SCSI hard disks - ocfs2disk.vmdk (512MB), asmdisk1.vmdk (3GB), asmdisk2.vmdk (3GB), and asmdisk3.vmdk (2GB).

16. VMware Server Console: Click on **Edit virtual machine settings**.
17. Virtual Machine Settings: Click on **Add**.
18. Add Hardware Wizard: Click on **Next**.
19. Hardware Type:
   a. Hardware types: Select **Hard Disk**.
20. Select a Disk:
   a. Disk: Select **Create a new virtual disk**.
21. Select a Disk Type:
   a. Virtual Disk Type: Select **SCSI (Recommended)**.
22. Specify Disk Capacity:
   a. Disk capacity: Enter “0.5GB.”
   b. Select **Allocate all disk space now**. You do not have to allocate all the disk space if you want to save space. For performance reason, you will pre-allocate all the disk space for each of the virtual shared disk. If the size of the shared disks were to grow rapidly especially during Oracle database creation or when the database is under heavy DML activity, the virtual machines may hang intermittently for a brief period or crash in a few

rare occasions.

23. Specify Disk File:
    a. Disk file: Enter “d:\vm\rac\sharedstorage\ocfs2disk.vmdk.”
    b. Click on **Advanced**.
24. Add Hardware Wizard:
    a. Virtual device node: Select **SCSI 1:0**.
    b. Mode: Select **Independent**, **Persistent** for all shared disks.
    c. Click on **Finish**.

Finally, add an additional virtual network card for the private interconnects and remove the floppy drive, if any.

25. VMware Server Console: Click on E**dit virtual machine settings**.
26. Virtual Machine Settings: Click on **Add**.
27. Add Hardware Wizard: Click on **Next**.
28. Hardware Type:
    a. Hardware types: Ethernet Adapter.
29. Network Type:
    a. Host-only: A private network shared with the host
    b. Click on **Finish**.
30. Virtual Machine Settings:
    a. Select **Floppy** and click on **Remove**.
31. Virtual Machine Settings: Click on **OK**.

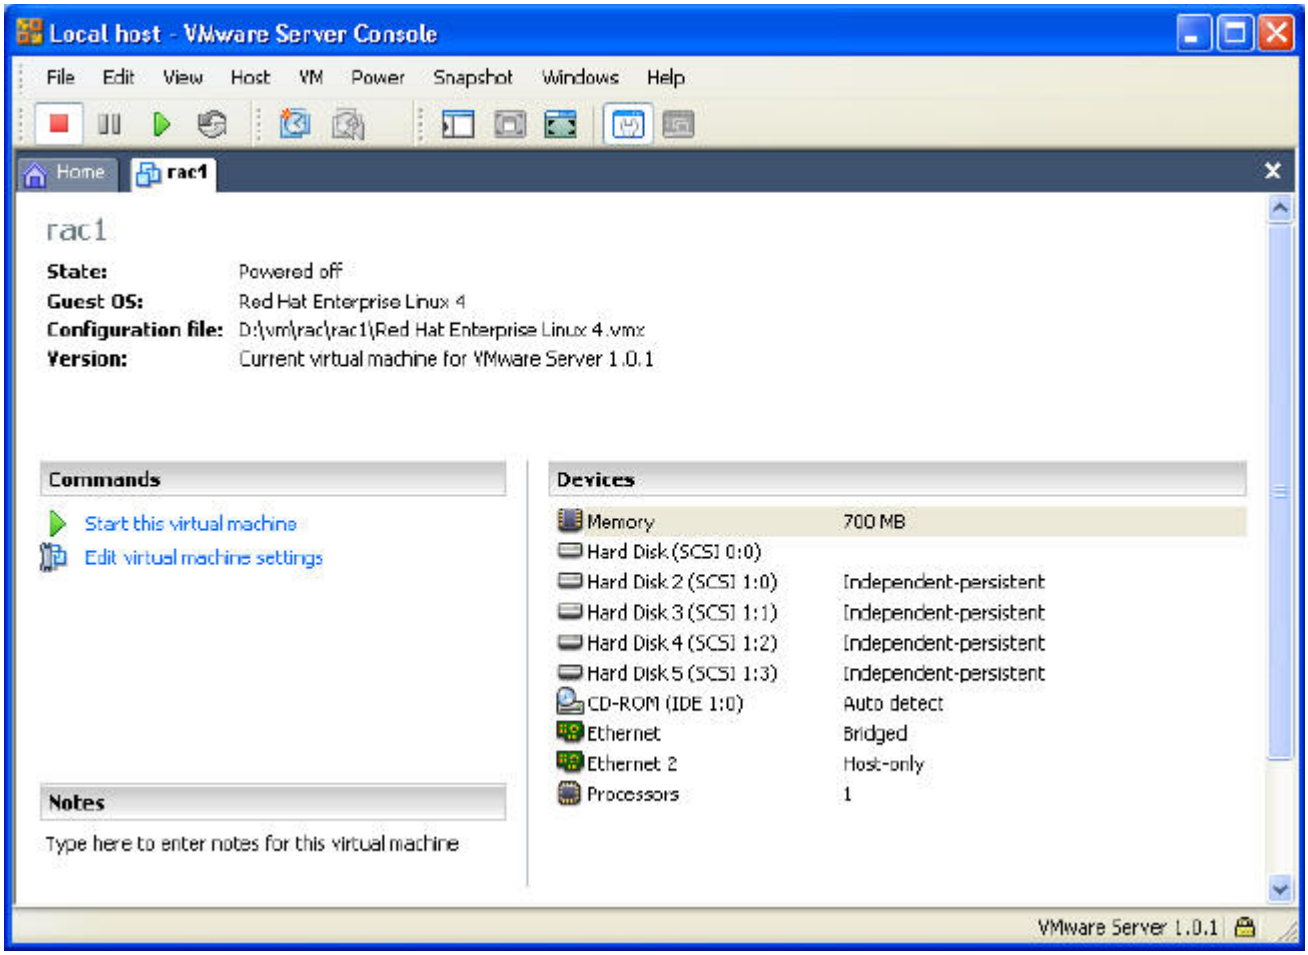

Modify virtual machine configuration file. Additional parameters are required to enable disk sharing between the two virtual RAC nodes. Open the configuration file, d:\vm\rac\rac1\Red Hat Enterprise Linux 4.vmx and add the bold parameters listed below.

```
config.version = "8"
virtualHW.version = "4"
scsi0.present = "TRUE"
scsi0.virtualDev = "lsilogic"
memsize = "700"
scsi0:0.present = "TRUE"
scsi0:0.fileName = "localdisk.vmdk"
ide1:0.present = "TRUE"
ide1:0.fileName = "auto detect"
ide1:0.deviceType = "cdrom-raw"
floppy0.fileName = "A:"
Ethernet0.present = "TRUE"
displayName = "rac1"
guestOS = "rhel4"
priority.grabbed = "normal"
priority.ungrabbed = "normal"

disk.locking = "FALSE"
diskLib.dataCacheMaxSize = "0"
scsi1.sharedBus = "virtual"

scsi1.present = "TRUE"
scsi1:0.present = "TRUE"
scsi1:0.fileName = "D:\vm\rac\sharedstorage\ocfs2disk.vmdk"
scsi1:0.mode = "independent-persistent"
scsi1:0.deviceType = "disk"
scsi1:1.present = "TRUE"
scsi1:1.fileName = "D:\vm\rac\sharedstorage\asmdisk1.vmdk"
scsi1:1.mode = "independent-persistent"
scsi1:1.deviceType = "disk"
scsi1:2.present = "TRUE"
scsi1:2.fileName = "D:\vm\rac\sharedstorage\asmdisk2.vmdk"
scsi1:2.mode = "independent-persistent"
scsi1:2.deviceType = "disk"
scsi1:3.present = "TRUE"
scsi1:3.fileName = "D:\vm\rac\sharedstorage\asmdisk3.vmdk"
scsi1:3.mode = "independent-persistent"
scsi1:3.deviceType = "disk"
scsi1.virtualDev = "lsilogic"
ide1:0.autodetect = "TRUE"
floppy0.present = "FALSE"
Ethernet1.present = "TRUE"
Ethernet1.connectionType = "hostonly"
```

3. Install and Configure Enterprise Linux on the First Virtual Machine

Download Enterprise Linux from Oracle and unzip the files:

- Enterprise-R4-U4-i386-disc1.iso
- Enterprise-R4-U4-i386-disc2.iso
- Enterprise-R4-U4-i386-disc3.iso
- Enterprise-R4-U4-i386-disc4.iso

1. On your VMware Server Console, double-click on the CD-ROM device on the right panel and

select the ISO image for disk 1, Enterprise-R4-U4-i386-disc1.iso.

2. VMware Server console:
    - Click on **Start this virtual machine**.
3. Hit **Enter** to install in graphical mode.
4. Skip the media test and start the installation.
5. Welcome to enterprise Linux: Click on **Next**.
6. Language Selection: <select your language preference>.
7. Keyboard Configuration: <select your keyboard preference>.
8. Installation Type: Custom.
9. Disk Partitioning Setup: Manually partition with Disk Druid.
    - Warning: Click on **Yes** to initialize each of the device – sda, sdb, sdc, sdd, and sde.
10. Disk Setup: Allocate disk space on sda drive by double-clicking on /dev/sda free space for the mount points (/ and /u01) and swap space. You will configure the rest of the drives for OCFS2 and ASM later.
    - Add Partition:
        - Mount Point: /
        - File System Type: ext3
        - Start Cylinder: 1
        - End Cylinder: 910

        - File System Type: Swap
        - Start Cylinder: 911
        - End Cylinder: 1170

        - Mount Point: /u01
        - File System Type: ext3
        - Start Cylinder: 1171
        - End Cylinder: 2610

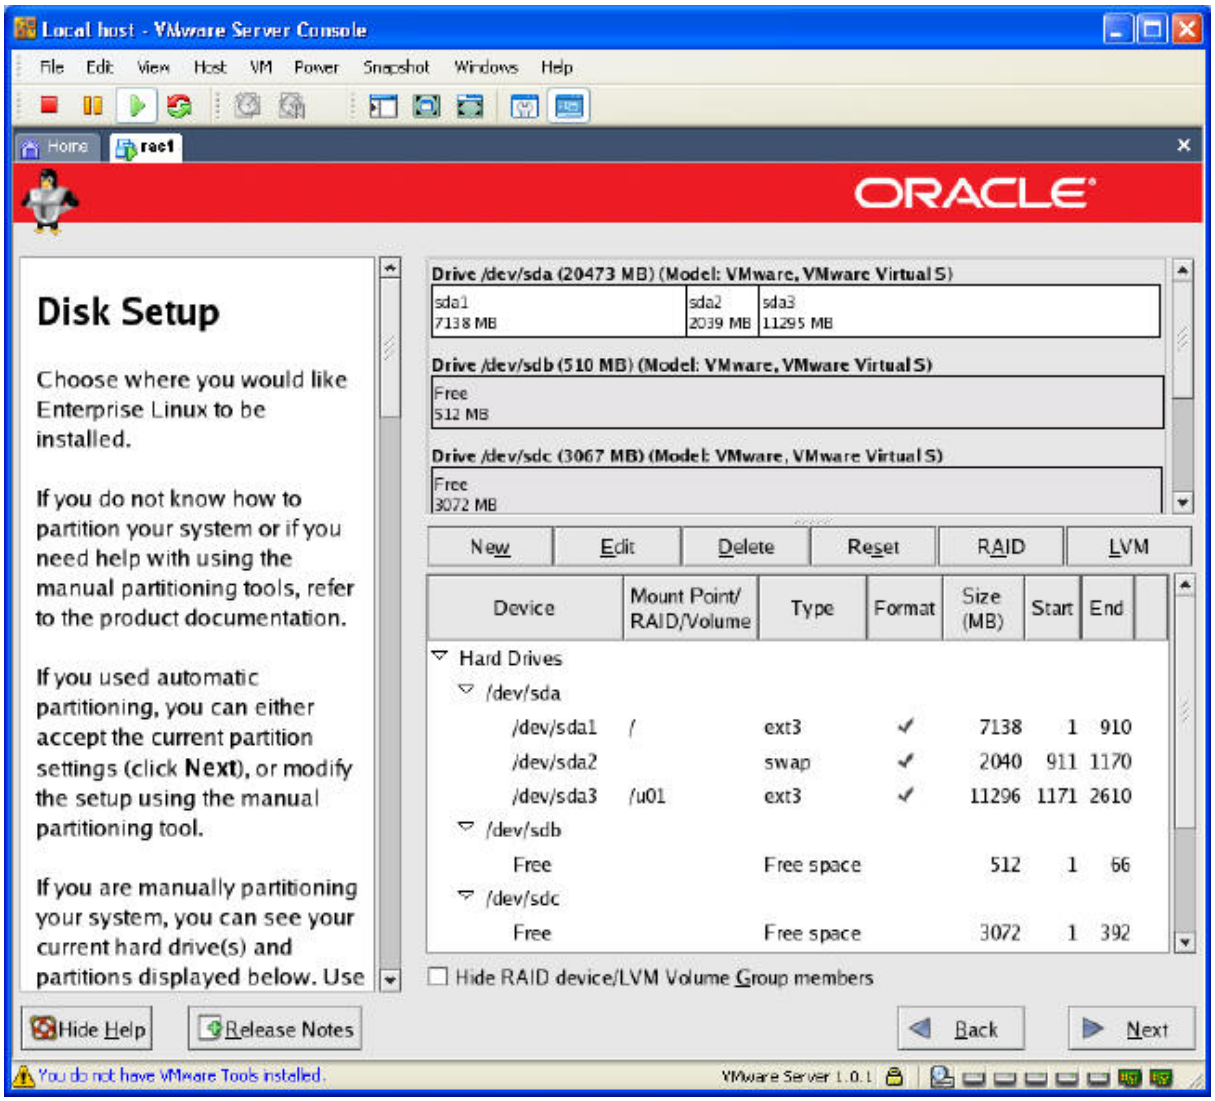

11. Boot Loader Configuration: Select only the default /dev/sda1 and leave the rest unchecked.
12. Network Configuration:
    a. Network Devices
        - Select and edit eth0
            1. De-select **Configure Using DHCP**.
            2. Select **Activate on boot**.
            3. IP Address: Enter “192.168.2.131.”
            4. Netmask: Enter “255.255.255.0.”
        - Select and edit eth1
            1. De-select **Configure Using DHCP**.
            2. Select **Activate on boot**.
            3. IP Address: Enter “10.10.10.31.”
            4. Netmask: Enter “255.255.255.0.”
    b. Hostname
        - Select **manually** and enter “rac1.mycorpdomain.com.”
    c. Miscellaneous Settings
        - Gateway: Enter “192.168.2.1.”
        - Primary DNS: <optional>
        - Secondary DNS: <optional>

13. Firewall Configuration:
    a. Select **No Firewall**. If firewall is enabled, you may encounter an error, "mount.ocfs2: Transport endpoint is not connected while mounting" when you attempt to mount ocfs2 file system later in the set up.
    b. Enable SELinux?: Active.
14. Warning – No Firewall: Click on **Proceed**.
15. Additional Language Support: <select the desired language>.
16. Time Zone Selection: <select your time zone>
17. Set Root Password: <enter your root password>
18. Package Group Selection:
    a. Select **X Window System**.
    b. Select **GNOME Desktop Environment**.
    c. Select **Editors**.
        - Click on **Details** and select your preferred text editor.
    d. Select **Graphical Internet**.
    e. Select **Text-based Internet**.
    f. Select **Office/Productivity**.
    g. Select **Sound and Video**.
    h. Select **Graphics**.
    i. Select **Server Configuration Tools**.
    j. Select **FTP Server**.
    k. Select **Legacy Network Server**.
        - Click on **Details**.
            1. Select **rsh-server**.
            2. Select **telnet-server**.
    l. Select **Development Tools**.
    m. Select **Legacy Software Development**.
    n. Select **Administration Tools**.
    o. Select **System Tools**.
        - Click on **Details**. Select the following packages in addition to the default selected packages.
            1. Select **ocfs-2-2.6.9-42.0.0.0.1EL** (driver for UP kernel), or select **ocfs-2-2.6.9-42.0.0.0.1ELsmp** (driver for SMP kernel).
            2. Select **ocfs2-tools**.
            3. Select **ocfs2console**.
            4. Select **oracle oracleasm-2.6.9-42.0.0.0.1EL** (driver for UP kernel) or select **oracleasm-2.6.9-42.0.0.0.1ELsmp** (driver for SMP kernel).
            5. Select **sysstat**.
    p. Select **Printing Support**.
19. About to Install: Click on **Next**.
20. Required Install Media: Click on **Continue**.
21. Change CD-ROM: On your VMware Server Console, press CTRL-D to bring up the Virtual Machine Settings. Click on the CD-ROM device and select the ISO image for disk 2, Enterprise-R4-U4-i386-disc2.iso, followed by the ISO image for disk 3, Enterprise-R4-U4-i386-disc3.iso.
22. At the end of the installation:
    a. On your VMware Server Console, press CTRL-D to bring up the Virtual Machine Settings. Click on the CD-ROM device and select **Use physical drive**.
    b. Click on **Reboot**.

23. Welcome: Click on **Next**.
24. License Agreement: Select **Yes, I agree to the License Agreement**.
25. Date and Time: Set the date and time.
26. Display: <select your desired resolution>.
27. System User: Leave the entries blank and click on **Next**.
28. Additional CDs: Click on **Next**.
29. Finish Setup: Click on **Next.**

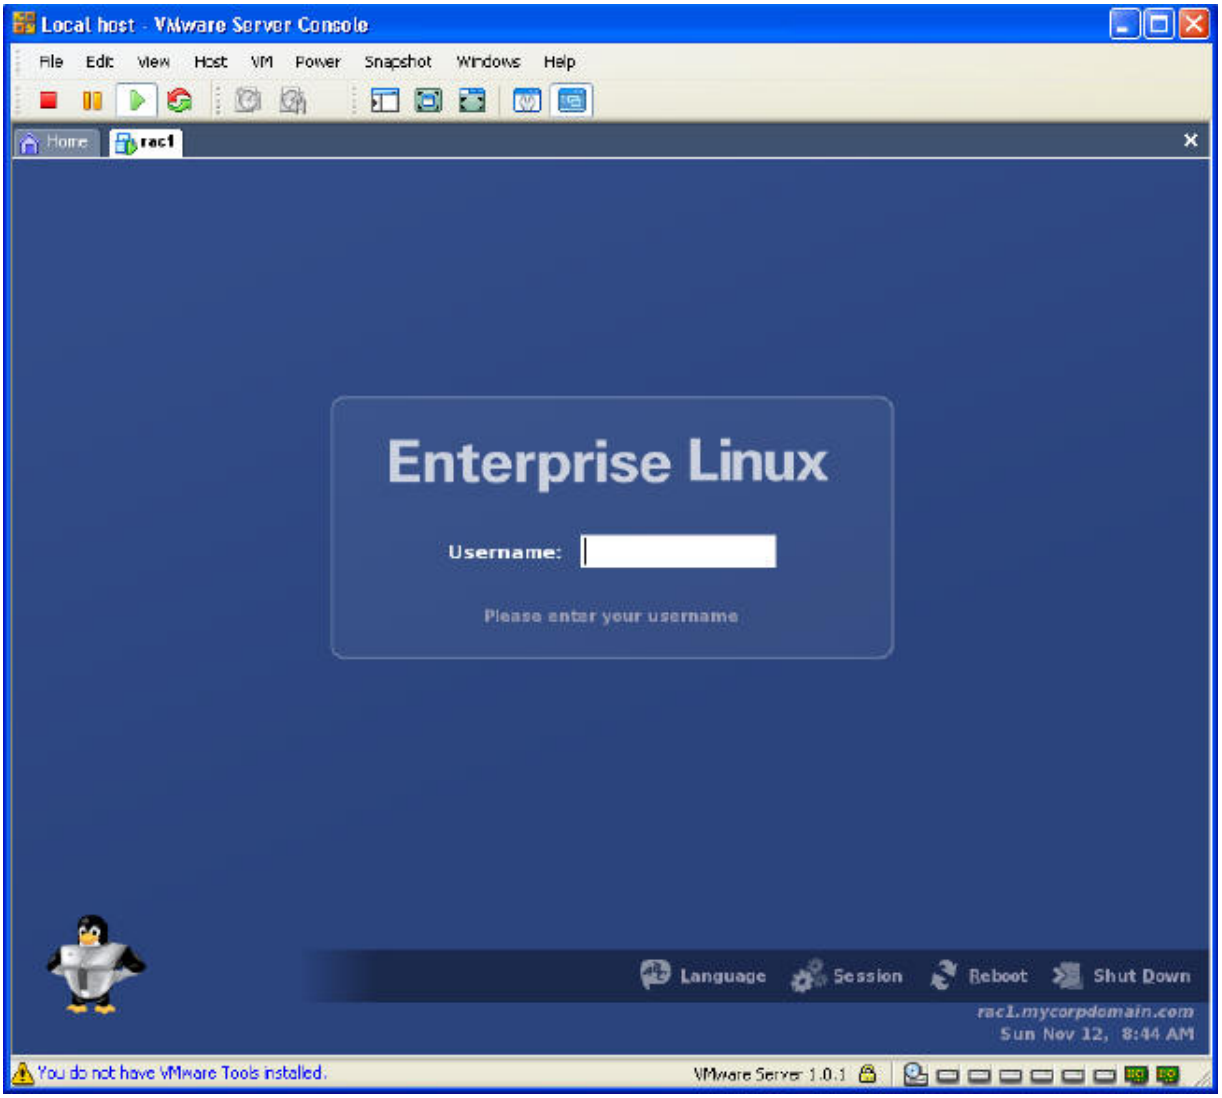

Congratulations, you have just installed Enterprise Linux on VMware Server!

Install VMware Tools. VMware Tools is required to synchronize the time between the host and guest machines.

On the VMware Console, log in as the root user,

1. Click on **VM** and then select **Install VMware Tools**.
2. rac1 – Virtual Machine: Click on **Install**.
3. Double-click on the VMware Tools icon on your desktop.
4. cdrom: Double-click on **VMwareTools-1.0.1-29996.i386.rpm**.

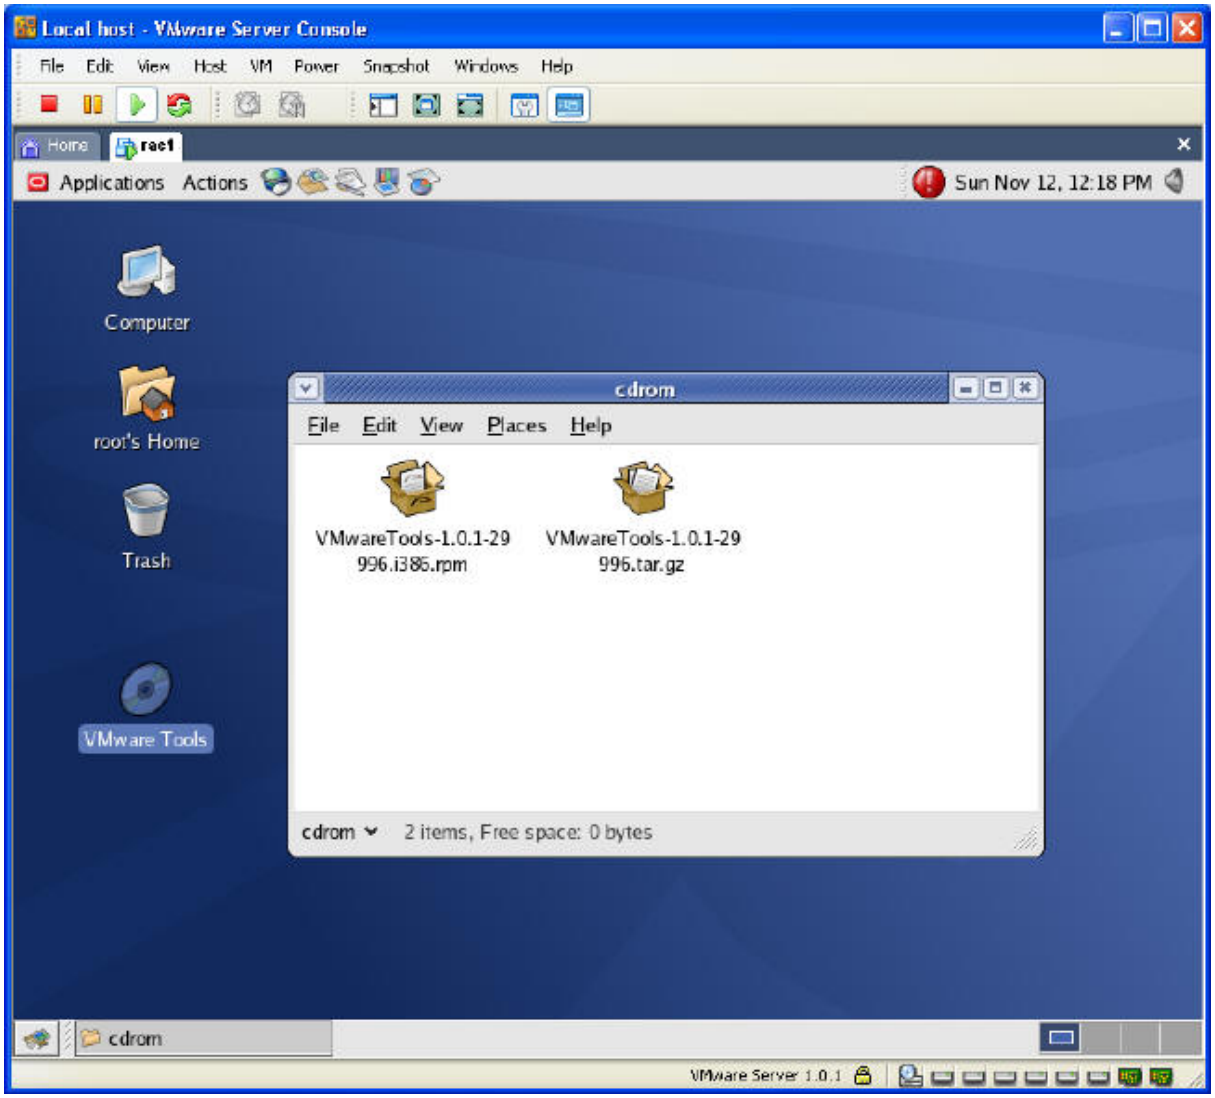

5. Completed System Preparation: Click on **Continue**.
6. Open up a terminal and execute **vmware-config-tools.pl**.
    - Enter the desired display size.

Synchronize Guest OS time with Host OS. When installing the Oracle Clusterware and Oracle Database software, the Oracle installer will initially install the software on the local node and then remotely copies the software to the remote node. If the date and time of both RAC nodes are not synchronized, you will likely receive errors similar to the one below.

```
"/bin/tar: ./inventory/Components21/oracle.ordim.server/10.2.0.1.0: time
stamp  2006-11-04 06:24:04 is 25 s in the future"
```

To ensure a successful Oracle RAC installation, the time on the virtual machines has to synchronize with the host machine. Perform the steps below to synchronize the time as the root user.

1. Execute “vmware-toolbox” to bring up the VMware Tools Properties window. Under the Options tab, select **Time synchronization between the virtual machine and the host operating system**. You should find the tools.syncTime = "TRUE" parameter appended to the virtual machine configuration file, d:\vm\rac\rac1\Red Hat Enterprise Linux 4.vmx.
2. Edit /boot/grub/grub.conf and add the options, "clock=pit nosmp noapic nolapic" to the line that reads kernel /boot/. You have added the options to both kernels. You are only required to make

the change to your specific kernel.

```
#boot=/dev/sda
default=0
timeout=5
splashimage=(hd0,0)/boot/grub/splash.xpm.gz
hiddenmenu
title Enterprise (2.6.9-42.0.0.0.1.ELsmp)
        root (hd0,0)
        kernel /boot/vmlinuz-2.6.9-42.0.0.0.1.ELsmp ro
root=LABEL=/ rhgb quiet clock=pit nosmp noapic nolapic
        initrd /boot/initrd-2.6.9-42.0.0.0.1.ELsmp.img
title Enterprise-up (2.6.9-42.0.0.0.1.EL)
        root (hd0,0)
        kernel /boot/vmlinuz-2.6.9-42.0.0.0.1.EL ro root=LABEL=/
rhgb quiet clock=pit nosmp noapic nolapic
        initrd /boot/initrd-2.6.9-42.0.0.0.1.EL.img
```

3. Reboot rac1.

```
# reboot
```

Create the oracle user. As the root user, execute

```
# groupadd oinstall
# groupadd dba
# mkdir -p /export/home/oracle /ocfs
# useradd -d /export/home/oracle -g oinstall -G dba -s /bin/ksh oracle
# chown oracle:dba /export/home/oracle /u01
# passwd oracle
New Password:
Re-enter new Password:
passwd: password successfully changed for oracle
```

Create the oracle user environment file.

/export/home/oracle/.profile

```
export PS1="`/bin/hostname -s`-> "
export EDITOR=vi
export ORACLE_SID=devdb1
export ORACLE_BASE=/u01/app/oracle
export ORACLE_HOME=$ORACLE_BASE/product/10.2.0/db_1
export ORA_CRS_HOME=$ORACLE_BASE/product/10.2.0/crs_1
export LD_LIBRARY_PATH=$ORACLE_HOME/lib
export PATH=$ORACLE_HOME/bin:$ORA_CRS_HOME/bin:/bin:
/usr/bin:/usr/sbin:/usr/local/bin:/usr/X11R6/bin
umask 022
```

Create the filesystem directory structure. As the oracle user, execute

```
rac1-> mkdir -p $ORACLE_BASE/admin
rac1-> mkdir -p $ORACLE_HOME
rac1-> mkdir -p $ORA_CRS_HOME
rac1-> mkdir -p /u01/oradata/devdb
```

Increase the shell limits for the Oracle user. Use a text editor and add the lines listed below to /etc/security/limits.conf, /etc/pam.d/login, and /etc/profile. Additional information can be obtained from the documentation.

/etc/security/limits.conf

```
oracle soft nproc 2047
oracle hard nproc 16384
oracle soft nofile 1024
oracle hard nofile 65536
```

/etc/pam.d/login

```
session required /lib/security/pam_limits.so
```

/etc/profile

```
if [ $USER = "oracle" ]; then
        if [ $SHELL = "/bin/ksh" ]; then
                ulimit -p 16384
                ulimit -n 65536
        else
                ulimit -u 16384 -n 65536
        fi
fi
```

Install Enterprise Linux software packages. The following additional packages are required for Oracle software installation. If you have installed the 64-bit version of Enterprise Linux, the installer should have already installed these packages.

- libaio-0.3.105-2.i386.rpm
- openmotif21-2.1.30-11.RHEL4.6.i386.rpm

Extract the packages from the ISO CDs and execute the command below as the root user.

```
# ls
libaio-0.3.105-2.i386.rpm  openmotif21-2.1.30-11.RHEL4.6.i386.rpm
#
# rpm -Uvh *.rpm
warning: libaio-0.3.105-2.i386.rpm: V3 DSA signature: NOKEY, key ID b38a8516
Preparing...
########################################### [100%]
   1:openmotif21
########################################### [ 50%]
   2:libaio
########################################### [100%]
```

Configure the kernel parameters. Use a text editor and add the lines listed below to /etc/sysctl.conf. To make the changes effective immediately, execute `/sbin/sysctl –p`.

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

Modify the /etc/hosts file.

```
# more /etc/hosts
127.0.0.1               localhost
192.168.2.131           rac1.mycorpdomain.com        rac1
192.168.2.31            rac1-vip.mycorpdomain.com    rac1-vip
```

```
10.10.10.31             rac1-priv.mycorpdomain.com   rac1-priv
192.168.2.132           rac2.mycorpdomain.com        rac2
192.168.2.32            rac2-vip.mycorpdomain.com    rac2-vip
10.10.10.32             rac2-priv.mycorpdomain.com   rac2-priv
```

Configure the hangcheck timer kernel module. The hangcheck timer kernel module monitors the system's health and restarts a failing RAC node. It uses two parameters, hangcheck_tick (defines the system checks frequency) and hangcheck_margin (defines the maximum hang delay before a RAC node is reset), to determine if a node is failing.

Add the following line in /etc/modprobe.conf to set the hangcheck kernel module parameters.

/etc/modprobe.conf

```
options hangcheck-timer hangcheck_tick=30 hangcheck_margin=180
```

To load the module immediately, execute "`modprobe -v hangcheck-timer`".

Create disk partitions for OCFS2 and Oracle ASM. Prepare a set of raw disks for OCFS2 (/dev/sdb), and for Oracle ASM (/dev/sdc, /dev/sdd, /dev/sde).

On rac1, as the root user, execute

 # fdisk /dev/sdb

```
Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-512, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-512, default 512):
Using default value 512

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

 # fdisk /dev/sdc

```
Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-391, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-391, default 391):
Using default value 391

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

 # fdisk /dev/sdd

```
Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-391, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-391, default 391):
Using default value 391

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

 # fdisk /dev/sde

```
Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-261, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-261, default 261):
Using default value 261

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

 # fdisk -l

```
Disk /dev/sda: 21.4 GB, 21474836480 bytes
255 heads, 63 sectors/track, 2610 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *           1         910     7309543+  83  Linux
/dev/sda2             911        1170     2088450   82  Linux swap
/dev/sda3            1171        2610    11566800   83  Linux

Disk /dev/sdb: 536 MB, 536870912 bytes
64 heads, 32 sectors/track, 512 cylinders
Units = cylinders of 2048 * 512 = 1048576 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1               1         512      524272   83  Linux

Disk /dev/sdc: 3221 MB, 3221225472 bytes
255 heads, 63 sectors/track, 391 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sdc1               1         391     3140676   83  Linux

Disk /dev/sdd: 3221 MB, 3221225472 bytes
255 heads, 63 sectors/track, 391 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

```
   Device Boot      Start         End      Blocks   Id  System
/dev/sdd1               1         391     3140676   83  Linux

Disk /dev/sde: 2147 MB, 2147483648 bytes
255 heads, 63 sectors/track, 261 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sde1               1         261     2096451   83  Linux
```

Install oracleasmlib package. Download the ASM library from OTN and install the ASM RPM as the root user.

```
 # rpm -Uvh oracleasmlib-2.0.2-1.i386.rpm
Preparing...
########################################### [100%]
   1:oracleasmlib
########################################### [100%]
```

At this stage, you should already have the following ASM packages installed.

```
[root@rac1 swdl]# rpm -qa | grep oracleasm
oracleasm-support-2.0.3-2
oracleasm-2.6.9-42.0.0.0.1.ELsmp-2.0.3-2
oracleasmlib-2.0.2-1
```

Map raw devices for ASM disks. A raw device mapping is required *only* if you are planning on creating ASM disks using standard Linux I/O. An alternative to creating ASM disks is to use the ASM library driver provided by Oracle. You will configure ASM disks using ASM library driver later.

Perform the following tasks to map the raw devices to the shared partitions created earlier. The raw devices have to bind with the block devices each time a cluster node boots.

Add the following lines in /etc/sysconfig/rawdevices.

/etc/sysconfig/rawdevices

```
/dev/raw/raw1 /dev/sdc1
/dev/raw/raw2 /dev/sdd1
/dev/raw/raw3 /dev/sde1
```

To make the mapping effective immediately, execute the following commands as the root user:

```
# /sbin/service rawdevices restart
Assigning devices:
           /dev/raw/raw1  -->   /dev/sdc1
/dev/raw/raw1:  bound to major 8, minor 33
           /dev/raw/raw2  -->   /dev/sdd1
/dev/raw/raw2:  bound to major 8, minor 49
           /dev/raw/raw3  -->   /dev/sde1
/dev/raw/raw3:  bound to major 8, minor 65
done

# chown oracle:dba /dev/raw/raw[1-3]
# chmod 660 /dev/raw/raw[1-3]
# ls -lat /dev/raw/raw*
crw-rw----  1 oracle dba 162, 3 Nov  4 07:04 /dev/raw/raw3
crw-rw----  1 oracle dba 162, 2 Nov  4 07:04 /dev/raw/raw2
crw-rw----  1 oracle dba 162, 1 Nov  4 07:04 /dev/raw/raw1
```

As the oracle user, execute

```
rac1-> ln -sf /dev/raw/raw1 /u01/oradata/devdb/asmdisk1
rac1-> ln -sf /dev/raw/raw2 /u01/oradata/devdb/asmdisk2
rac1-> ln -sf /dev/raw/raw3 /u01/oradata/devdb/asmdisk3
```

Modify /etc/udev/permissions.d/50-udev.permissions. Raw devices are remapped on boot. The ownership of the raw devices will change to the root user by default upon boot. ASM will have problem accessing the shared partitions if the ownership is not the oracle user. Comment the original line, "raw/*:root:disk:0660" in /etc/udev/permissions.d/50-udev.permissions and add a new line, "raw/*:oracle:dba:0660."

/etc/udev/permissions.d/50-udev.permissions

```
# raw devices
ram*:root:disk:0660
#raw/*:root:disk:0660
raw/*:oracle:dba:0660
```

4. Create and Configure the Second Virtual Machine

To create the second virtual machine, simply shut down the first virtual machine, copy all the files in d:\vm\rac\rac1 to d:\vm\rac\rac2 and perform a few configuration changes.

Modify network configuration.

1. As the root user on rac1,
   `# shutdown –h now`
2. On your host system, copy all the files in rac1 folder to rac2.
   `D:\>copy d:\vm\rac\rac1 d:\vm\rac\rac2`
3. On your VMware Server Console, press CTRL-O to open the second virtual machine, d:\rac\rac2\Red Hat Enterprise Linux 4.vmx.
4. VMware Server console:
   - Rename the virtual machine name from rac1 to rac2. Right-click on the new **rac1** tab you have just opened and select **Settings**.
     - Select the **Options** tab.
       1. Virtual machine name: Enter "rac2."

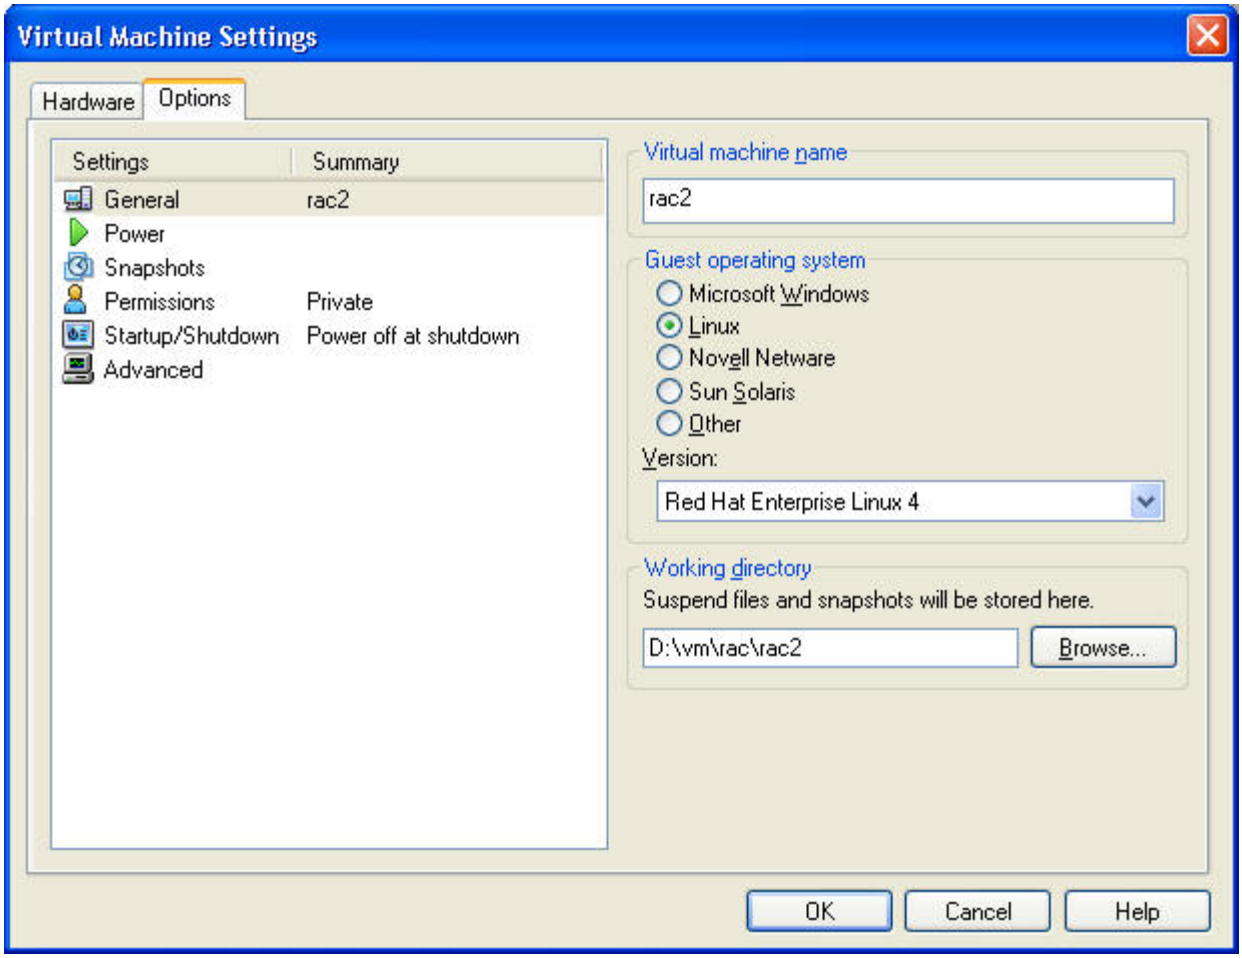

- Click on **Start this virtual machine** to start rac2, leaving rac1 powered off.
- rac2 – Virtaul Machine: Select **Create a new identifier**.

5. Log in as the root user and execute system-config-network to modify the network configuration.

   IP Address: Double-click on each of the Ethernet devices and use the table below to make the necessary changes.

Device	IP Address	Subnet mask	Default gateway address
eth0	192.168.2.132	255.255.255.0	192.168.2.1
eth1	10.10.10.32	255.255.255.0	<leave empty>

   MAC Address: Navigate to the **Hardware Device** tab and probe for a new MAC address for each of the Ethernet device.

   Hostname and DNS: Use the table below to make the necessary changes to the entries in the DNS tab and press CTRL-S to save.

Hostname	Primary DNS	Secondary DNS	DNS search path
rac2.mycorpdomain.com	Enter your DNS IP address or leave it empty.	Enter your DNS IP address or leave it empty.	Accepts the default or leave it empty.

Finally, activate each of the Ethernet device.

Modify /etc/hosts. Add the following entry in /etc/hosts.

127.0.0.1 localhost

VIPCA will attempt to use the loopback address later during the Oracle Clusterware software installation.

Modify /export/home/oracle/.profile. Replace the value of ORACLE_SID with devdb2.

Establish user equivalence with SSH. During the Cluster Ready Services (CRS) and RAC installation, the Oracle Universal Installer (OUI) has to be able to copy the software as oracle to all RAC nodes without being prompted for a password. In Oracle 10*g*, this can be accomplished using ssh instead of rsh.

To establish user equivalence, generate the user's public and private keys as the oracle user on both nodes. Power on rac1 and perform the following tasks on both nodes.
On rac1,

```
rac1-> mkdir ~/.ssh
rac1-> chmod 700 ~/.ssh
rac1-> ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/export/home/oracle/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /export/home/oracle/.ssh/id_rsa.
Your public key has been saved in /export/home/oracle/.ssh/id_rsa.pub.
The key fingerprint is:
87:54:4f:92:ba:ed:7b:51:5d:1d:59:5b:f9:44:da:b6 oracle@rac1.mycorpdomain.com
rac1-> ssh-keygen -t dsa
Generating public/private dsa key pair.
Enter file in which to save the key (/export/home/oracle/.ssh/id_dsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /export/home/oracle/.ssh/id_dsa.
Your public key has been saved in /export/home/oracle/.ssh/id_dsa.pub.
The key fingerprint is:
31:76:96:e6:fc:b7:25:04:fd:70:42:04:1f:fc:9a:26 oracle@rac1.mycorpdomain.com
```

On rac2,

```
rac2-> mkdir ~/.ssh
rac2-> chmod 700 ~/.ssh
rac2-> ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/export/home/oracle/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
```

```
Enter same passphrase again:
Your identification has been saved in /export/home/oracle/.ssh/id_rsa.
Your public key has been saved in /export/home/oracle/.ssh/id_rsa.pub.
The key fingerprint is:
29:5a:35:ac:0a:03:2c:38:22:3c:95:5d:68:aa:56:66 oracle@rac2.mycorpdomain.com
rac2-> ssh-keygen -t dsa
Generating public/private dsa key pair.
Enter file in which to save the key (/export/home/oracle/.ssh/id_dsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /export/home/oracle/.ssh/id_dsa.
Your public key has been saved in /export/home/oracle/.ssh/id_dsa.pub.
The key fingerprint is:
4c:b2:5a:8d:56:0f:dc:7b:bc:e0:cd:3b:8e:b9:5c:7c oracle@rac2.mycorpdomain.com
```

On rac1,

```
rac1-> cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys
rac1-> cat ~/.ssh/id_dsa.pub >> ~/.ssh/authorized_keys
rac1-> ssh rac2 cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys
The authenticity of host 'rac2 (192.168.2.132)' can't be established.
RSA key fingerprint is 63:d3:52:d4:4d:e2:cb:ac:8d:4a:66:9f:f1:ab:28:1f.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'rac2,192.168.2.132' (RSA) to the list of known hosts.
oracle@rac2's password:
rac1-> ssh rac2 cat ~/.ssh/id_dsa.pub >> ~/.ssh/authorized_keys
oracle@rac2's password:
rac1-> scp ~/.ssh/authorized_keys rac2:~/.ssh/authorized_keys
oracle@rac2's password:
authorized_keys                           100% 1716     1.7KB/s   00:00
```

Test the connection on each node. Verify that you are not prompted for password when you run the following the second time.

```
ssh rac1 date
ssh rac2 date
ssh rac1-priv date
ssh rac2-priv date
ssh rac1.mycorpdomain.com date
ssh rac2.mycorpdomain.com date
ssh rac1-priv.mycorpdomain.com date
ssh rac2-priv.mycorpdomain.com date
```

5. Configure Oracle Automatic Storage Management (ASM)

Oracle ASM is tightly integrated with Oracle Database and works with Oracle's suite of data management tools. It simplifies database storage management and provides the performance of raw disk I/O.

Configure ASMLib. Configure the ASMLib as the root user on both nodes.

```
# /etc/init.d/oracleasm configure
Configuring the Oracle ASM library driver.

This will configure the on-boot properties of the Oracle ASM library
driver.  The following questions will determine whether the driver is
loaded on boot and what permissions it will have.  The current values
will be shown in brackets ('[]').  Hitting  without typing an
answer will keep that current value.  Ctrl-C will abort.

Default user to own the driver interface []: oracle
Default group to own the driver interface []: dba
```

```
Start Oracle ASM library driver on boot (y/n) [n]: y
Fix permissions of Oracle ASM disks on boot (y/n) [y]: y
Writing Oracle ASM library driver configuration:           [  OK  ]
Loading module "oracleasm":                                [  OK  ]
Mounting ASMlib driver filesystem:                         [  OK  ]
Scanning system for ASM disks:                             [  OK  ]
```

Create ASM disks. Create the ASM disks on any one node as the root user.

```
# /etc/init.d/oracleasm createdisk VOL1 /dev/sdc1
Marking disk "/dev/sdc1" as an ASM disk:                   [  OK  ]
# /etc/init.d/oracleasm createdisk VOL2 /dev/sdd1
Marking disk "/dev/sdd1" as an ASM disk:                   [  OK  ]
# /etc/init.d/oracleasm createdisk VOL3 /dev/sde1
Marking disk "/dev/sde1" as an ASM disk:                   [  OK  ]
Verify that the ASM disks are visible from every node.
# /etc/init.d/oracleasm scandisks
Scanning system for ASM disks:                       [  OK  ]
# /etc/init.d/oracleasm listdisks
VOL1
VOL2
VOL3
VOL4
```

6. Configure Oracle Cluster File System (OCFS2)

OCFS2 is a general-purpose cluster file system developed by Oracle and integrated with the Enterprise Linux kernel. It enables all nodes to share files concurrently on the cluster file system and thus eliminates the need to manage raw devices. Here you will house the OCR and Voting Disk in the OCFS2 file system. Additional information on OCFS2 can be obtained from OCFS2 User's Guide.

You should already have the OCFS2 RPMs installed during the Enterprise Linux installation. Verify that the RPMs have been installed on both nodes.

```
rac1-> rpm -qa | grep ocfs
ocfs2-tools-1.2.2-2
ocfs2console-1.2.2-2
ocfs2-2.6.9-42.0.0.0.1.ELsmp-1.2.3-2
```

Create the OCFS2 configuration file. As the root user on rac1, execute

```
# ocfs2console
```

1. OCFS2 Console: Select **Cluster**, **Configure Nodes**.
2. "The cluster stack has been started": Click on **Close**.
3. Node Configuration: Click on **Add**.
4. Add Node: Add the following nodes and then click on **Apply**.
    - Name: rac1
    - IP Address: 192.168.2.131
    - IP Port: 7777
    - Name: rac2
    - IP Address: 192.168.2.132
    - IP Port: 7777
5. Verify the generated configuration file.

```
# more /etc/ocfs2/cluster.conf
node:
        ip_port = 7777
```

```
        ip_address = 192.168.2.131
        number = 0
        name = rac1
        cluster = ocfs2

node:
        ip_port = 7777
        ip_address = 192.168.2.132
        number = 1
        name = rac2
        cluster = ocfs2

cluster:
        node_count = 2
        name = ocfs2
```

6. Propagate the configuration file to rac2. You can rerun the steps above on rac2 to generate the configuration file or select **Cluster**, **Propagate Configuration** on the OCFS2 Console on rac1 to propagate the configuration file to rac2.

Configure the O2CB driver. O2CB is a set of clustering services that manages the communication between the nodes and the cluster file system. Below is a description of the individual services:

- NM: Node Manager that keep track of all the nodes in the cluster.conf
- HB: Heartbeat service that issues up/down notifications when nodes join or leave the cluster
- TCP: Handles communication between the nodes
- DLM: Distributed lock manager that keeps track of all locks, its owners, and status
- CONFIGFS: User space driven configuration file system mounted at /config
- DLMFS: User space interface to the kernel space DLM

Perform the procedure below on both nodes to configure O2CB to start on boot.

When prompted for a value for the heartbeat dead threshold, you have to specify a value higher than 7 to prevent the nodes from crashing due to the slow IDE disk drive. The heartbeat dead threshold is a variable used to calculate the fence time.

```
Fence time (seconds) = (heartbeat dead threshold -1) * 2
```

A fence time of 120 seconds works well in our environment. The value of heartbeat dead threshold should be the same on both nodes.

As the root user, execute

```
# /etc/init.d/o2cb unload
Stopping O2CB cluster ocfs2: OK
Unmounting ocfs2_dlmfs filesystem: OK
Unloading module "ocfs2_dlmfs": OK
Unmounting configfs filesystem: OK
Unloading module "configfs": OK

# /etc/init.d/o2cb configure
Configuring the O2CB driver.

This will configure the on-boot properties of the O2CB driver.
The following questions will determine whether the driver is loaded on
boot.  The current values will be shown in brackets ('[]').  Hitting
 without typing an answer will keep that current value.  Ctrl-C
will abort.
```

```
Load O2CB driver on boot (y/n) [y]: y
Cluster to start on boot (Enter "none" to clear) [ocfs2]:
Specify heartbeat dead threshold (>=7) [7]: 61
Writing O2CB configuration: OK
Loading module "configfs": OK
Mounting configfs filesystem at /config: OK
Loading module "ocfs2_nodemanager": OK
Loading module "ocfs2_dlm": OK
Loading module "ocfs2_dlmfs": OK
Mounting ocfs2_dlmfs filesystem at /dlm: OK
Starting O2CB cluster ocfs2: OK
```

Format the file system. Before proceeding with formatting and mounting the file system, verify that O2CB is online on both nodes; O2CB heartbeat is currently inactive because the file system is not mounted.

```
# /etc/init.d/o2cb status
Module "configfs": Loaded
Filesystem "configfs": Mounted
Module "ocfs2_nodemanager": Loaded
Module "ocfs2_dlm": Loaded
Module "ocfs2_dlmfs": Loaded
Filesystem "ocfs2_dlmfs": Mounted
Checking O2CB cluster ocfs2: Online
Checking O2CB heartbeat: Not active
```

You are only required to format the file system on one node. As the root user on rac1, execute

```
# ocfs2console
```

1. OCFS2 Console: Select **Tasks, Format**.
2. Format:
    - Available devices: /dev/sdb1
    - Volume label: oracle
    - Cluster size: Auto
    - Number of node slots: 4
    - Block size: Auto
3. OCFS2 Console: CTRL-Q to quit.

Mount the file system. To mount the file system, execute the command below on both nodes.

```
# mount -t ocfs2 -o datavolume,nointr /dev/sdb1 /ocfs
```

To mount the file system on boot, add the following line in /etc/fstab on both nodes.

/etc/fstab

```
/dev/sdb1 /ocfs ocfs2 _netdev,datavolume,nointr 0 0
```

Create Oracle Clusterware directory. Create the directory in OCFS2 file system where the OCR and Voting Disk will reside.

On rac1,

```
# mkdir /ocfs/clusterware
# chown -R oracle:dba /ocfs
```

You have completed the set up of OCFS2. Verify that you can read and write files on the shared cluster

file system from both nodes.

7. Install Oracle Clusterware

After downloading, as the oracle user on rac1, execute

```
rac1-> /u01/staging/clusterware/runInstaller
```

1. Welcome: Click on **Next**.
2. Specify Inventory directory and credentials:
   - Enter the full path of the inventory directory: /u01/app/oracle/oraInventory.
   - Specify Operating System group name: oinstall.
3. Specify Home Details:
   - Name: OraCrs10g_home
   - /u01/app/oracle/product/10.2.0/crs_1
4. Product-Specific Prerequisite Checks:
   - Ignore the warning on physical memory requirement.
5. Specify Cluster Configuration: Click on **Add**.
   - Public Node Name: rac2.mycorpdomain.com
   - Private Node Name: rac2-priv.mycorpdomain.com
   - Virtual Host Name: rac2-vip.mycorpdomain.com
6. Specify Network Interface Usage:
   - Interface Name: eth0
   - Subnet: 192.168.2.0
   - Interface Type: Public
   - Interface Name: eth1
   - Subnet: 10.10.10.0
   - Interface Type: Private
7. Specify Oracle Cluster Registry (OCR) Location: Select **External Redundancy**.
   For simplicity, here you will not mirror the OCR. In a production environment, you may want to consider multiplexing the OCR for higher redundancy.
   - Specify OCR Location: /ocfs/clusterware/ocr
8. Specify Voting Disk Location: Select **External Redundancy**.
   Similarly, for simplicity, we have chosen not to mirror the Voting Disk.
   - Voting Disk Location: /ocfs/clusterware/votingdisk
9. Summary: Click on **Install**.
10. Execute Configuration scripts: Execute the scripts below as the root user sequentially, one at a time. Do not proceed to the next script until the current script completes.
    - Execute /u01/app/oracle/oraInventory/orainstRoot.sh on rac1.
    - Execute /u01/app/oracle/oraInventory/orainstRoot.sh on rac2.
    - Execute /u01/app/oracle/product/10.2.0/crs_1/root.sh on rac1.
    - Execute /u01/app/oracle/product/10.2.0/crs_1/root.sh on rac2.

    The root.sh script on rac2 invoked the VIPCA automatically but it failed with the error "The given interface(s), "eth0" is not public. Public interfaces should be used to configure virtual IPs." As you are using a non-routable IP address (192.168.x.x) for the public interface, the Oracle Cluster Verification Utility (CVU) could not find a suitable public interface. A workaround is to run VIPCA manually.
11. As the root user, manually invokes VIPCA on the second node.
    ```
    # /u01/app/oracle/product/10.2.0/crs_1/bin/vipca
    ```

12. Welcome: Click on **Next**.
13. Network Interfaces: Select **eth0**.
14. Virtual IPs for cluster nodes:
    - Node name: rac1
    - IP Alias Name: rac1-vip
    - IP address: 192.168.2.31
    - Subnet Mask: 255.255.255.0
    - Node name: rac2
    - IP Alias Name: rac2-vip
    - IP address: 192.168.2.32
    - Subnet Mask: 255.255.255.0
15. Summary: Click on **Finish**.

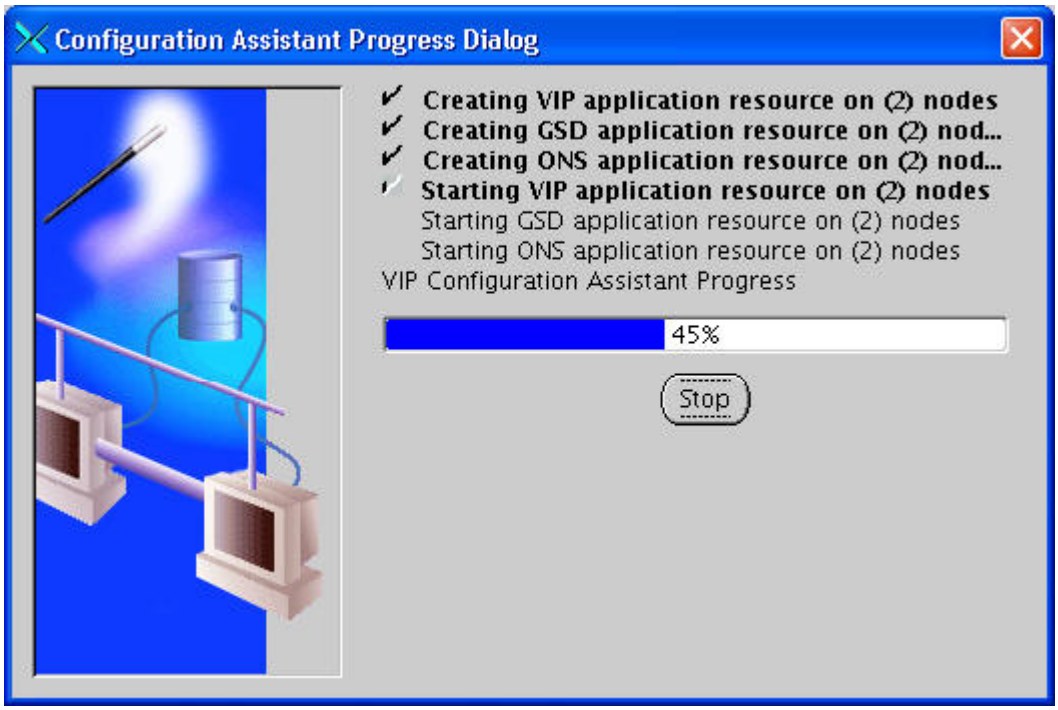

16. Configuration Assistant Progress Dialog: After the configuration has completed, click on **OK**.
17. Configuration Results: Click on **Exit**.
18. Return to the Execute Configuration scripts screen on rac1 and click on **OK**.

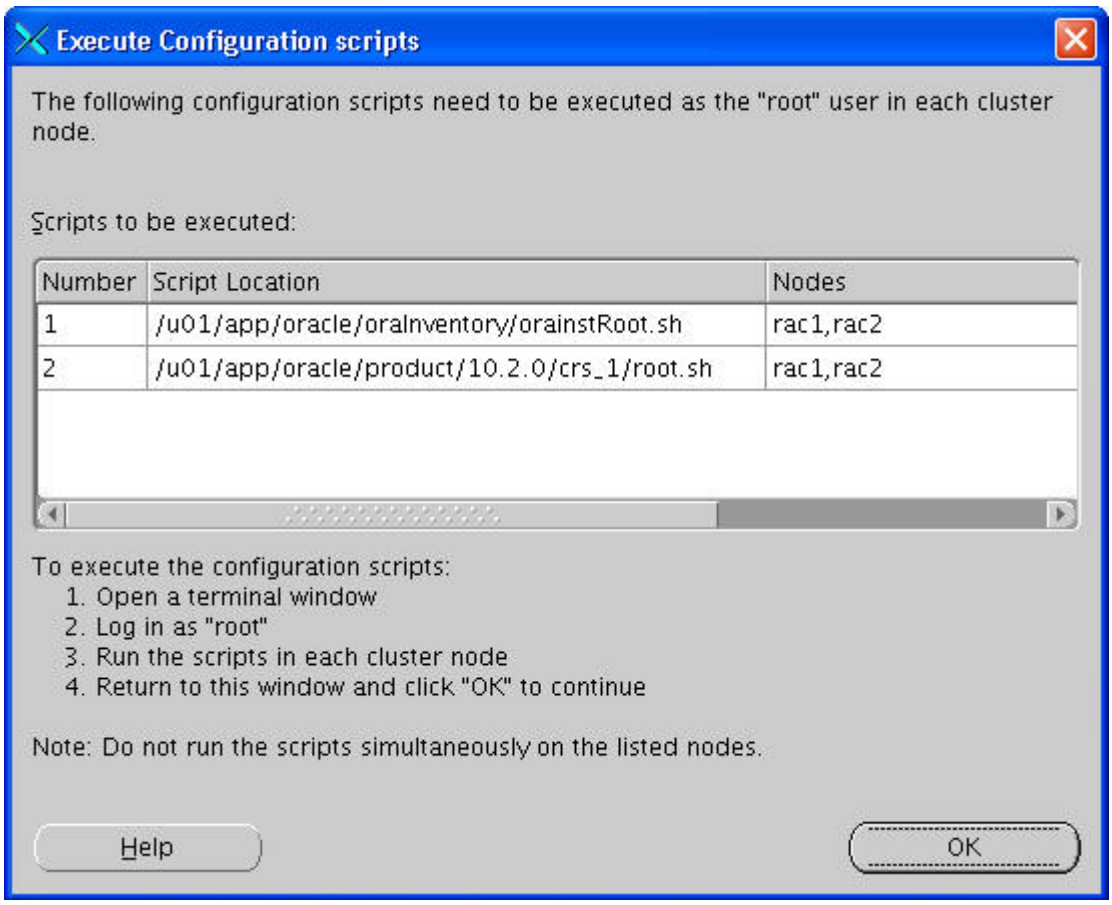

19. Configuration Assistants: Verify that all checks are successful. The OUI does a Clusterware post-installation check at the end. If the CVU fails, correct the problem and re-run the following command as the oracle user:

```
rac1-> /u01/app/oracle/product/10.2.0/crs_1/bin/cluvfy stage
-post crsinst -n rac1,rac2

Performing post-checks for cluster services setup

Checking node reachability...
Node reachability check passed from node "rac1".

Checking user equivalence...
User equivalence check passed for user "oracle".

Checking Cluster manager integrity...

Checking CSS daemon...
Daemon status check passed for "CSS daemon".

Cluster manager integrity check passed.

Checking cluster integrity...

Cluster integrity check passed

Checking OCR integrity...

Checking the absence of a non-clustered configuration...
All nodes free of non-clustered, local-only configurations.

Uniqueness check for OCR device passed.
```

```
Checking the version of OCR...
OCR of correct Version "2" exists.

Checking data integrity of OCR...
Data integrity check for OCR passed.

OCR integrity check passed.

Checking CRS integrity...

Checking daemon liveness...
Liveness check passed for "CRS daemon".

Checking daemon liveness...
Liveness check passed for "CSS daemon".

Checking daemon liveness...
Liveness check passed for "EVM daemon".

Checking CRS health...
CRS health check passed.

CRS integrity check passed.

Checking node application existence...

Checking existence of VIP node application (required)
Check passed.

Checking existence of ONS node application (optional)
Check passed.

Checking existence of GSD node application (optional)
Check passed.

Post-check for cluster services setup was successful.
```

20. End of Installation: Click on **Exit**.

8. Install Oracle Database 10*g* Release 2

After downloading, as the oracle user on rac1, execute

```
rac1-> /u01/staging/database/runInstaller
```

1. Welcome: Click on **Next**.
2. Select Installation Type:
   - Select **Enterprise Edition**.
3. Specify Home Details:
   - Name: OraDb10g_home1
   - Path: /u01/app/oracle/product/10.2.0/db_1
4. Specify Hardware Cluster Installation Mode:
   - Select **Cluster Installation**.
   - Click on **Select All**.
5. Product-Specific Prerequisite Checks:
   - Ignore the warning on physical memory requirement.
6. Select Configuration Option:
   - Create a database.

7. Select Database Configuration:
    - Select **Advanced**.
8. Summary: Click on **Install**.
9. Database Templates:
    - Select **General Purpose**.
10. Database identification:
    - Global Database Name: devdb
    - SID Prefix: devdb
11. Management Options:
    - Select **Configure the Database with Enterprise Manager**.
12. Database Credentials:
    - Use the Same Password for All Accounts.
13. Storage Options:
    - Select **Automatic Storage Management (ASM)**.
14. Create ASM Instance:
    - SYS password: <enter SYS password>.
    - Select **Create initialization parameter file (IFILE)**.
15. ASM Disk Groups:
    - Click on **Create New**.
16. Create Disk Group:
    Create two disk groups – DG1 and RECOVERYDEST.
    - Disk Group Name: DG1
    - Select **Normal** redundancy.
    - Select Disk Path, ORCL:VOL1 and ORCL:VOL2. If you have configured the ASM disks using standard Linux I/O, you will select /u01/oradata/devdb/asmdisk1 and /u01/oradata/devdb/asmdisk2 instead.
    - Click on **OK**.

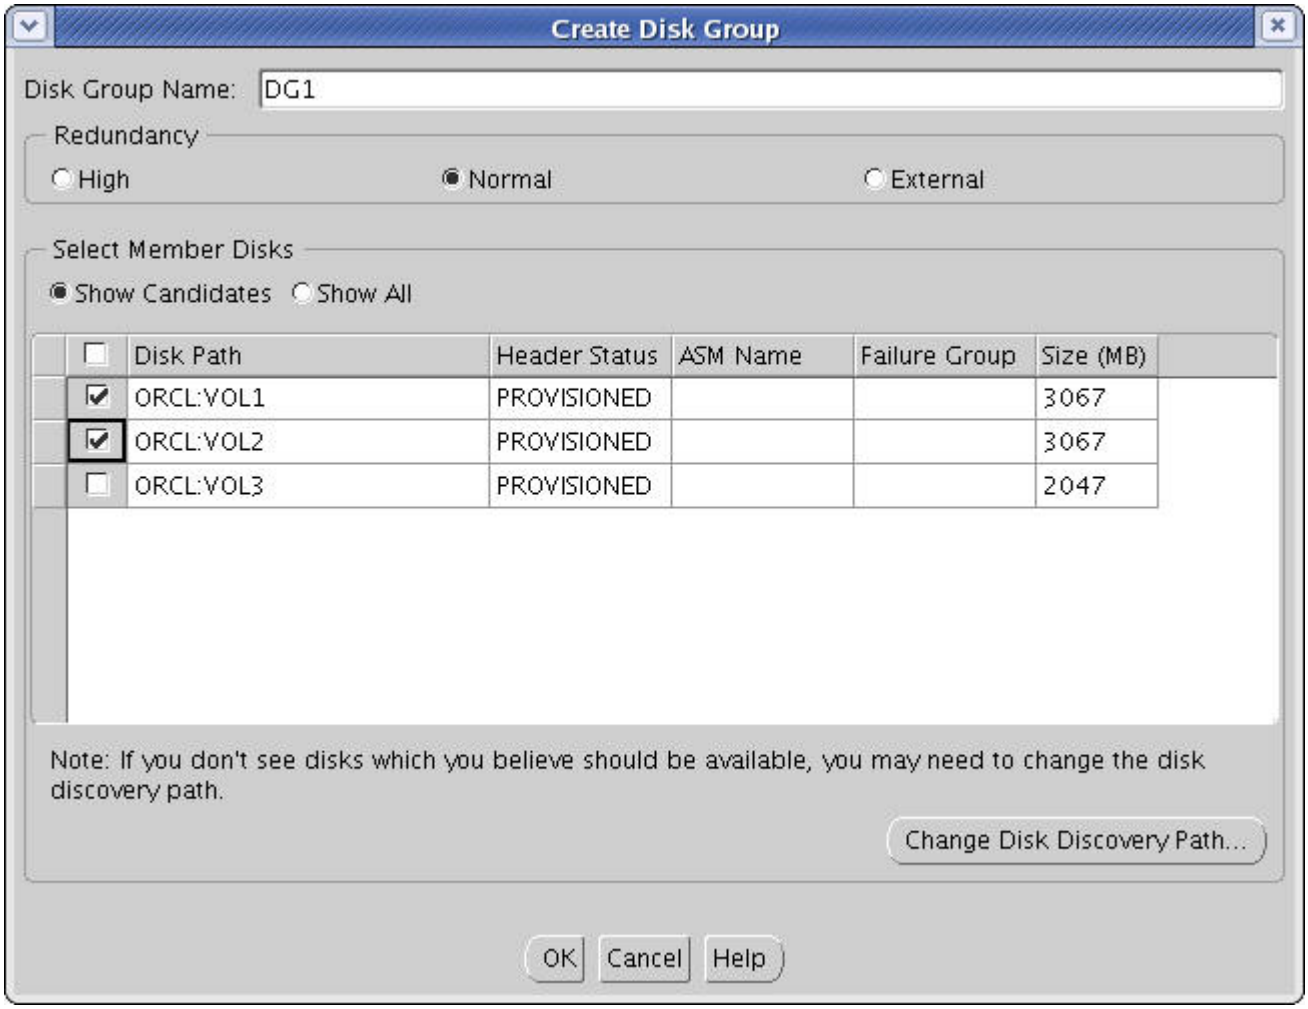

- Disk Group Name: RECOVERYDEST.
- Select **External** redundancy.
- Select Disk Path, ORCL:VOL3. If you have configured the ASM disks using standard Linux I/O, you will select /u01/oradata/devdb/asmdisk3 instead.
- Click on **OK**.

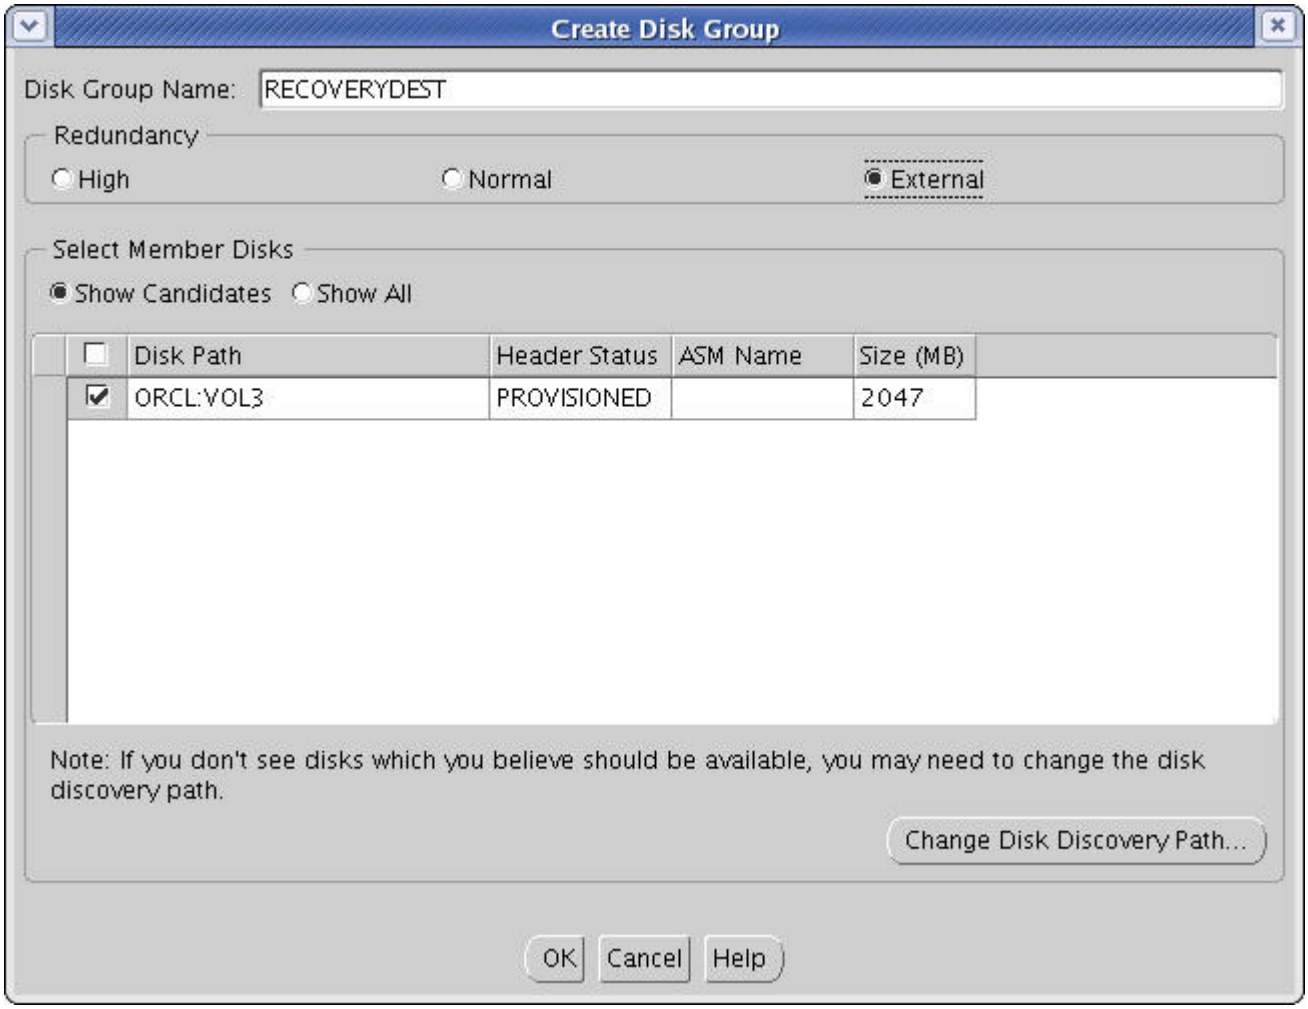

17. ASM Disk Groups: Click on **Next**.

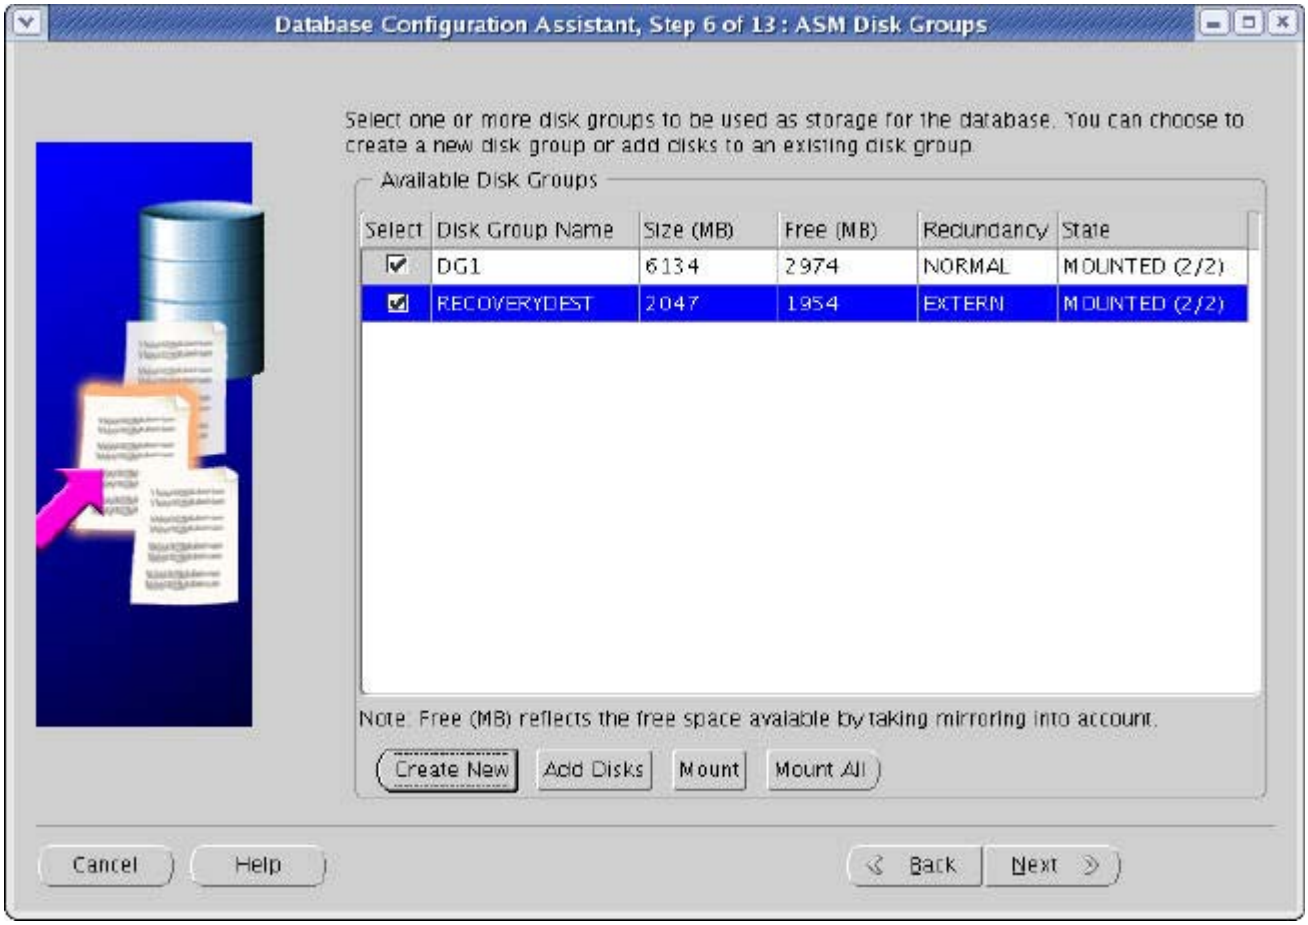

18. Database File Locations:
    - Select **Use Oracle-Managed Files.**
        - Database Area: +DG1
19. Recovery Configuration:
    - Select **Specify Flash Recovery Area**.
        - Flash Recovery Area: +RECOVERYDEST
        - Flash Recovery Area Size: 1500M
    - Select **Enable Archiving**.
20. Database Content:
    - Select or deselect the sample schemas.
21. Database Services:
    - Click on **Next**. You can always create or modify additional services later using DBCA or srvctl.
22. Initialization Parameters:
    - Select **Custom**.
        - Shared Memory Management: Automatic
        - SGA Size: 200MB
        - PGA Size: 25MB
    - b. Modify the rest of the parameters as necessary.
23. Database Storage: Click on **Next**.
24. Creation Options:
    - Select **Create Database**.
    - Click on **Finish**.

25. Summary: Click on **OK**.
26. Database Configuration Assistant: Click on **Exit**.

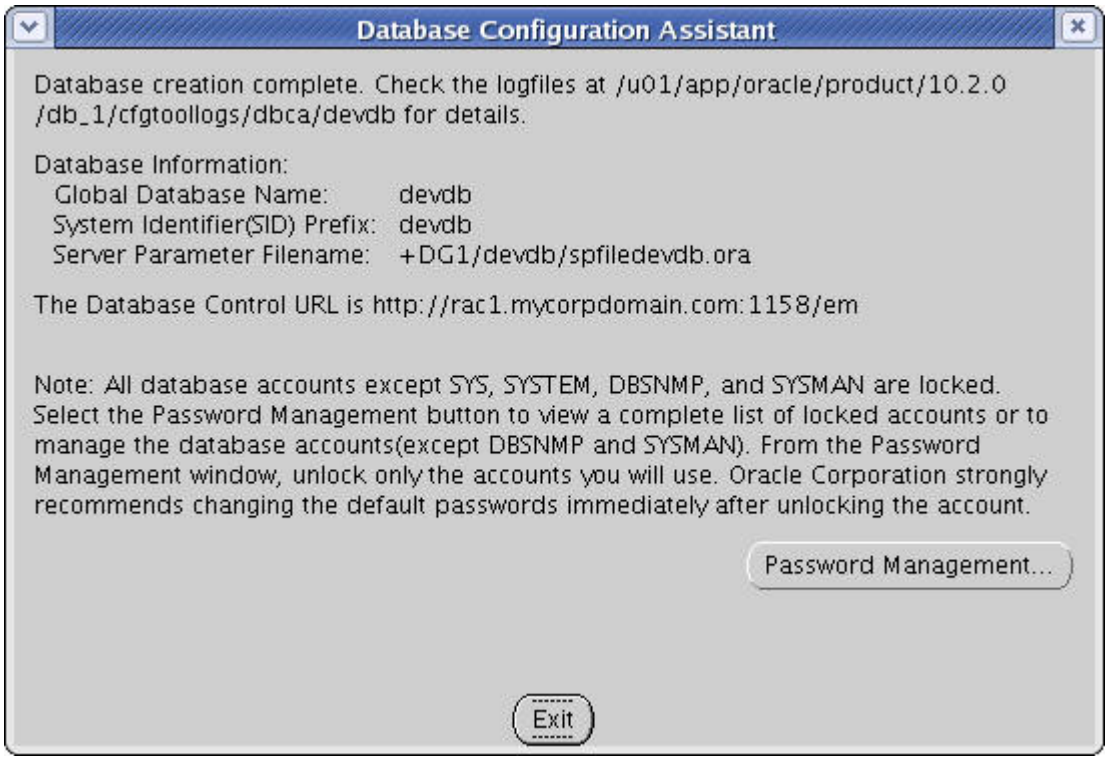

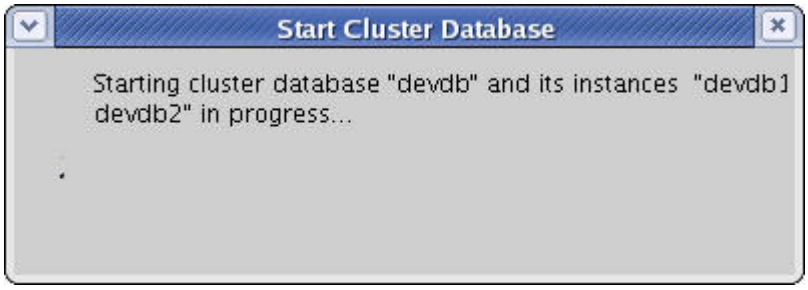

27. Execute Configuration scripts: Execute the scripts below as the root user.
    - Execute /u01/app/oracle/product/10.2.0/db_1/root.sh on rac1.
    - Execute /u01/app/oracle/product/10.2.0/db_1/root.sh on rac2.
28. Return to the Execute Configuration scripts screen on rac1 and click on **OK**.
29. End of Installation: Click on **Exit**.

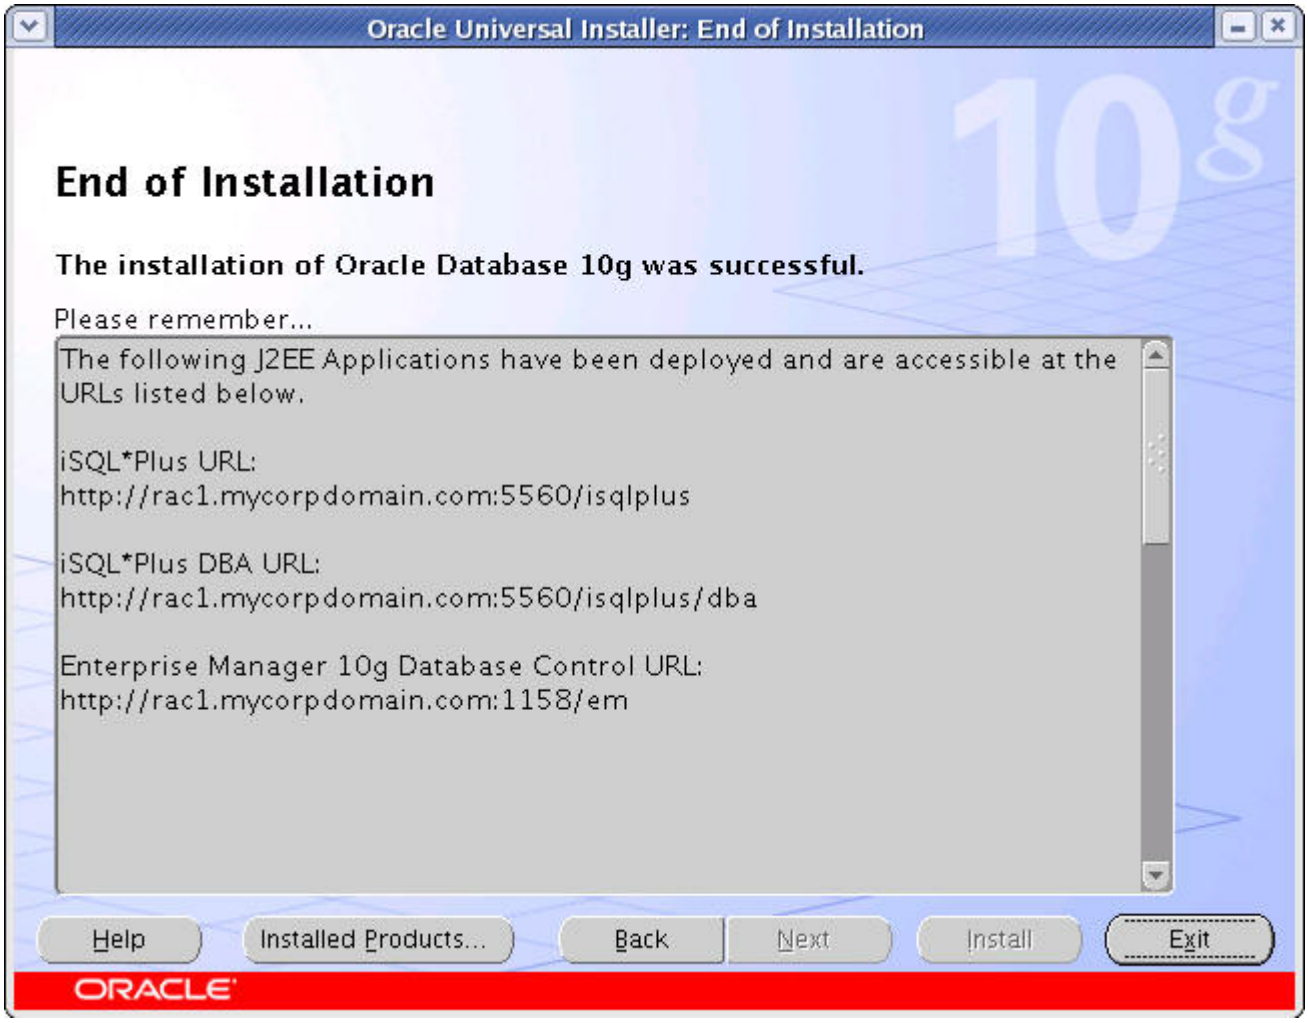

Congratulations, you have completed the installation of Oracle RAC Database 10*g* on Enterprise Linux!

9. Explore the RAC Database Environment

Now that you have successfully installed a virtual two-node RAC database, it's time to do a little exploration of the environment you have just set up.

Check the status of application resources.

```
rac1-> crs_stat -t
Name           Type           Target    State     Host
------------------------------------------------------------
ora.devdb.db   application    ONLINE    ONLINE    rac1
ora....b1.inst application    ONLINE    ONLINE    rac1
ora....b2.inst application    ONLINE    ONLINE    rac2
ora....SM1.asm application    ONLINE    ONLINE    rac1
ora....C1.lsnr application    ONLINE    ONLINE    rac1
ora.rac1.gsd   application    ONLINE    ONLINE    rac1
ora.rac1.ons   application    ONLINE    ONLINE    rac1
ora.rac1.vip   application    ONLINE    ONLINE    rac1
ora....SM2.asm application    ONLINE    ONLINE    rac2
ora....C2.lsnr application    ONLINE    ONLINE    rac2
ora.rac2.gsd   application    ONLINE    ONLINE    rac2
ora.rac2.ons   application    ONLINE    ONLINE    rac2
ora.rac2.vip   application    ONLINE    ONLINE    rac2
```

```
rac1-> srvctl status nodeapps -n rac1
VIP is running on node: rac1
GSD is running on node: rac1
Listener is running on node: rac1
ONS daemon is running on node: rac1

rac1-> srvctl status nodeapps -n rac2
VIP is running on node: rac2
GSD is running on node: rac2
Listener is running on node: rac2
ONS daemon is running on node: rac2

rac1-> srvctl status asm -n rac1
ASM instance +ASM1 is running on node rac1.

rac1-> srvctl status asm -n rac2
ASM instance +ASM2 is running on node rac2.

rac1-> srvctl status database -d devdb
Instance devdb1 is running on node rac1
Instance devdb2 is running on node rac2

rac1-> srvctl status service -d devdb
rac1->
```

Check the status of Oracle Clusterware.

```
rac1-> crsctl check crs
CSS appears healthy
CRS appears healthy
EVM appears healthy

rac2-> crsctl check crs
CSS appears healthy
CRS appears healthy
EVM appears healthy
```

Execute `crsctl` on the command line to check out all the available options.

List the RAC instances.

```
SQL> select
  2  instance_name,
  3  host_name,
  4  archiver,
  5  thread#,
  6  status
  7  from gv$instance;

INSTANCE_NAME  HOST_NAME             ARCHIVE  THREAD# STATUS
-------------- --------------------- ------- -------- ------
devdb1         rac1.mycorpdomain.com STARTED        1 OPEN
devdb2         rac2.mycorpdomain.com STARTED        2 OPEN
```

Check connectivity.

Verify that you are able to connect to the instances and service on each node.

```
sqlplus system@devdb1
sqlplus system@devdb2
sqlplus system@devdb
```

Check database configuration.

```
rac1-> export ORACLE_SID=devdb1
rac1-> sqlplus / as sysdba

SQL> show sga

Total System Global Area  209715200 bytes
Fixed Size                  1218556 bytes
Variable Size             104859652 bytes
Database Buffers          100663296 bytes
Redo Buffers                2973696 bytes


SQL> select file_name,bytes/1024/1024 from dba_data_files;

FILE_NAME                                 BYTES/1024/1024
----------------------------------------- ---------------
+DG1/devdb/datafile/users.259.606468449                 5
+DG1/devdb/datafile/sysaux.257.606468447              240
+DG1/devdb/datafile/undotbs1.258.606468449              30
+DG1/devdb/datafile/system.256.606468445              480
+DG1/devdb/datafile/undotbs2.264.606468677              25


        SQL> select
          2  group#,
          3  type,
          4  member,
          5  is_recovery_dest_file
          6  from v$logfile
          7  order by group#;

GROUP# TYPE    MEMBER                                                  IS_
------ ------- ------------------------------------------------------- ---
     1 ONLINE  +RECOVERYDEST/devdb/onlinelog/group_1.257.606468581 YES
     1 ONLINE  +DG1/devdb/onlinelog/group_1.261.606468575          NO
     2 ONLINE  +RECOVERYDEST/devdb/onlinelog/group_2.258.606468589 YES
     2 ONLINE  +DG1/devdb/onlinelog/group_2.262.606468583          NO
     3 ONLINE  +DG1/devdb/onlinelog/group_3.265.606468865          NO
     3 ONLINE  +RECOVERYDEST/devdb/onlinelog/group_3.259.606468875 YES
     4 ONLINE  +DG1/devdb/onlinelog/group_4.266.606468879          NO
     4 ONLINE  +RECOVERYDEST/devdb/onlinelog/group_4.260.606468887 YES


rac1-> export ORACLE_SID=+ASM1
rac1-> sqlplus / as sysdba

SQL> show sga

Total System Global Area   92274688 bytes
Fixed Size                  1217884 bytes
Variable Size              65890980 bytes
ASM Cache                  25165824 bytes


SQL> show parameter asm_disk

NAME                           TYPE        VALUE
------------------------------ ----------- ------------------------
asm_diskgroups                 string      DG1, RECOVERYDEST
asm_diskstring                 string


SQL> select
  2  group_number,
  3  name,
  4  allocation_unit_size alloc_unit_size,
```

```
  5  state,
  6  type,
  7  total_mb,
  8  usable_file_mb
  9  from v$asm_diskgroup;

                      ALLOC                          USABLE
 GROUP                 UNIT                  TOTAL    FILE
NUMBER NAME             SIZE STATE   TYPE       MB      MB
------ ------------ -------- ------- ------ ------ -------
     1 DG1           1048576 MOUNTED NORMAL   6134    1868
     2 RECOVERYDEST  1048576 MOUNTED EXTERN   2047    1713


SQL> select
  2  name,
  3  path,
  4  header_status,
  5  total_mb free_mb,
  6  trunc(bytes_read/1024/1024) read_mb,
  7  trunc(bytes_written/1024/1024) write_mb
  8  from v$asm_disk;

NAME  PATH       HEADER_STATU    FREE_MB    READ_MB   WRITE_MB
----- ---------- ------------ ---------- ---------- ----------
VOL1  ORCL:VOL1  MEMBER             3067        229       1242
VOL2  ORCL:VOL2  MEMBER             3067        164       1242
VOL3  ORCL:VOL3  MEMBER             2047         11        354
```

Create a tablespace.

```
SQL> connect system/oracle@devdb
Connected.
SQL> create tablespace test_d datafile '+DG1' size 10M;

Tablespace created.

SQL> select
  2  file_name,
  3  tablespace_name,
  4  bytes
  5  from dba_data_files
  6  where tablespace_name='TEST_D';

FILE_NAME                                TABLESPACE_NAME      BYTES
---------------------------------------- --------------- ----------
+DG1/devdb/datafile/test_d.269.606473423 TEST_D            10485760
```

Create an online redo logfile group.

```
SQL> connect system/oracle@devdb
Connected.
SQL> alter database add logfile thread 1 group 5 size 50M;

Database altered.

SQL> alter database add logfile thread 2 group 6 size 50M;

Database altered.

SQL> select
  2  group#,
  3  thread#,
  4  bytes,
  5  members,
  6  status
```

```
  7  from v$log;

    GROUP#    THREAD#      BYTES    MEMBERS STATUS
---------- ---------- ---------- ---------- ----------------
         1          1   52428800          2 CURRENT
         2          1   52428800          2 INACTIVE
         3          2   52428800          2 ACTIVE
         4          2   52428800          2 CURRENT
         5          1   52428800          2 UNUSED
         6          2   52428800          2 UNUSED

SQL> select
          2  group#,
          3  type,
          4  member,
          5  is_recovery_dest_file
          6  from v$logfile
          7  where group# in (5,6)
          8  order by group#;

GROUP# TYPE    MEMBER                                                  IS_
------ ------- ------------------------------------------------------- ---
     5 ONLINE  +DG1/devdb/onlinelog/group_5.271.606473683              NO
     5 ONLINE  +RECOVERYDEST/devdb/onlinelog/group_5.261.606473691     YES
     6 ONLINE  +DG1/devdb/onlinelog/group_6.272.606473697              NO
     6 ONLINE  +RECOVERYDEST/devdb/onlinelog/group_6.262.606473703     YES
```

Check flash recovery area space usage.

```
SQL> select * from v$recovery_file_dest;

NAME          SPACE_LIMIT SPACE_USED SPACE_RECLAIMABLE NUMBER_OF_FILES
------------- ----------- ---------- ----------------- ---------------
+RECOVERYDEST  1572864000  331366400                 0               7


SQL> select * from v$flash_recovery_area_usage;

FILE_TYPE    PERCENT_SPACE_USED PERCENT_SPACE_RECLAIMABLE NUMBER_OF_FILES
------------ ------------------ ------------------------- ---------------
CONTROLFILE                .97                         0               1
ONLINELOG                   20                         0               6
ARCHIVELOG                   0                         0               0
BACKUPPIECE                  0                         0               0
IMAGECOPY                    0                         0               0
FLASHBACKLOG                 0                         0               0
```

Start and stop application resources.

Follow the steps below to start and stop individual application resource.

```
srvctl start nodeapps -n <node1 hostname>
srvctl start nodeapps -n <node2 hostname>
srvctl start asm -n <node1 hostname>
srvctl start asm -n <node2 hostname>
srvctl start database -d <database name>
srvctl start service -d <database name> -s <service name>
crs_stat -t

srvctl stop service -d <database name> -s <service name>
srvctl stop database -d <database name>
srvctl stop asm -n <node1 hostname>
srvctl stop asm -n <node2 hostname>
srvctl stop nodeapps -n <node1 hostname>
srvctl stop nodeapps -n <node2 hostname>
```

```
crs_stat -t
```

10. Test Transparent Failover (TAF)

The failover mechanism in Oracle TAF enables any failed database connections to reconnect to another node within the cluster. The failover is transparent to the user. Oracle re-executes the query on the failed over instance and continues to display the remaining results to the user.

Create a new database service. Let's begin by creating a new service called CRM. Database services can be created using either DBCA or the srvctl utility. Here you will use DBCA to create the CRM service on devdb1.

Service Name	Database Name	Preferred Instance	Available Instance	TAF Policy
CRM	devdb	devdb1	devdb2	BASIC

As the oracle user on rac1, execute

```
rac1-> dbca
```

1. Welcome: Select **Oracle Real Application Clusters database**.
2. Operations: Select **Services Management**.
3. List of cluster databases: Click on **Next**.
4. Database Services: Click on **Add**.
    - Add a Service: Enter "CRM."
        - Select devdb1 as the Preferred instance.
        - Select devdb2 as the Available instance.
        - TAF Policy: Select **Basic**.
    - Click on **Finish**.

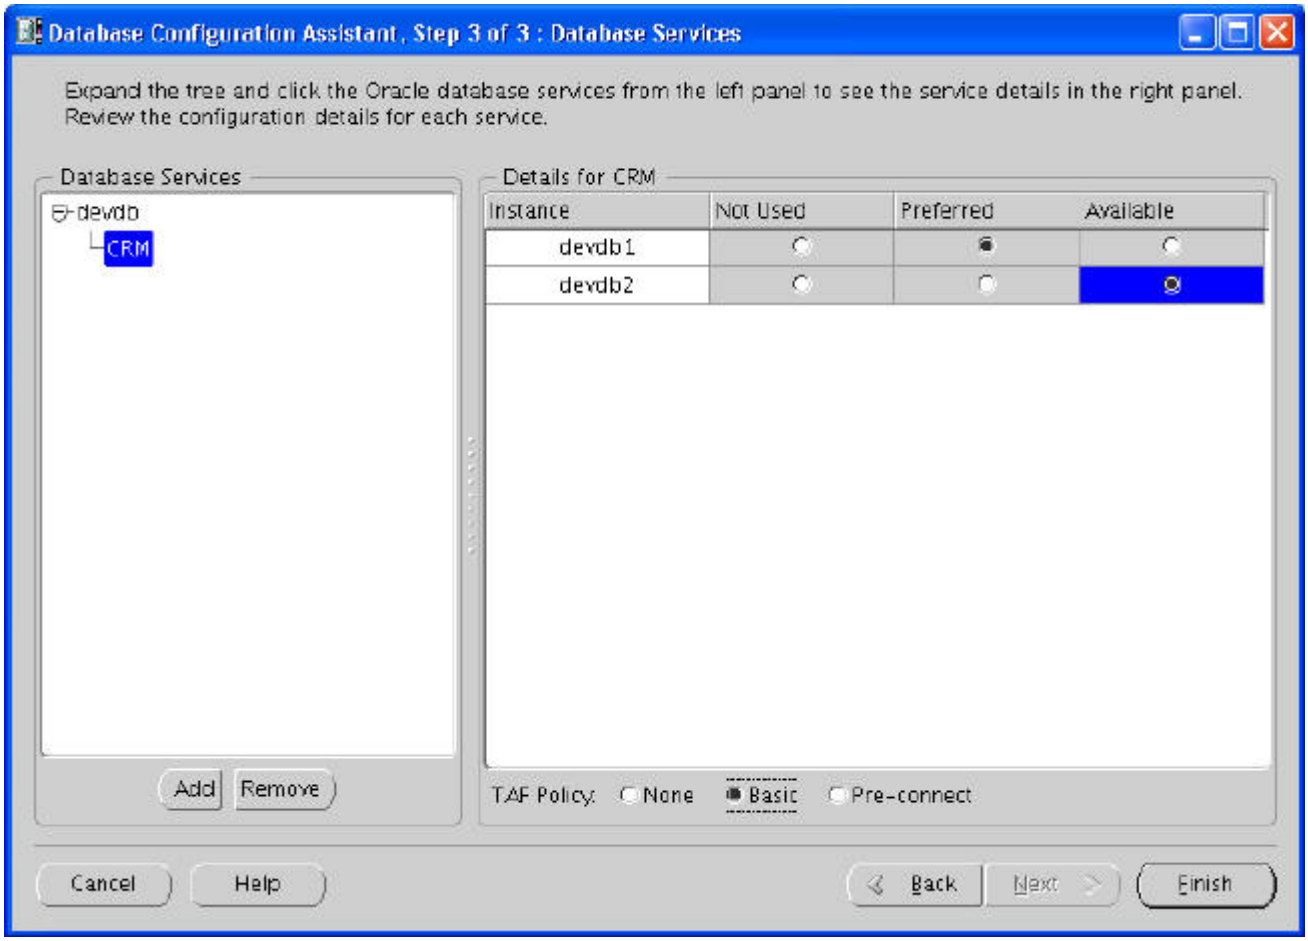

5. Database Configuration Assistant: Click on **No** to exit.

The Database Configuration Assistant creates the following CRM service name entry in tnsnames.ora:

```
CRM =
          (DESCRIPTION =
            (ADDRESS = (PROTOCOL = TCP)(HOST = rac1-vip)(PORT = 1521))
            (ADDRESS = (PROTOCOL = TCP)(HOST = rac2-vip)(PORT = 1521))
            (LOAD_BALANCE = yes)
            (CONNECT_DATA =
        (SERVER = DEDICATED)
        (SERVICE_NAME = CRM)
        (FAILOVER_MODE =
                  (TYPE = SELECT)
                  (METHOD = BASIC)
                  (RETRIES = 180)
                  (DELAY = 5)
        )
            )
          )


SQL> connect system/oracle@devdb1
        Connected.
        SQL> show parameter service

NAME                           TYPE        VALUE
------------------------------ ----------- ------------------------
service_names                  string      devdb, CRM
```

```
SQL> connect system/oracle@devdb2
        Connected.
        SQL> show parameter service

NAME                           TYPE        VALUE
------------------------------ ----------- ------------------------
service_names                  string      devdb
```

Connect the first session using the CRM service. If the returned output of failover_type and failover_mode is 'NONE', verify that the CRM service is configured correctly in tnsnames.ora.

```
SQL> connect system/oracle@crm
Connected.
SQL> select
  2  instance_number instance#,
  3  instance_name,
  4  host_name,
  5  status
  6  from v$instance;

 INSTANCE# INSTANCE_NAME    HOST_NAME              STATUS
---------- ---------------- ---------------------- ------------
         1 devdb1           rac1.mycorpdomain.com OPEN


SQL> select
  2  failover_type,
  3  failover_method,
  4  failed_over
  5  from v$session
  6  where username='SYSTEM';

FAILOVER_TYPE FAILOVER_METHOD FAILED_OVER
------------- --------------- ----------------
SELECT        BASIC           NO
```

Shut down the instance from another session. Connect as the sys user on CRM instance and shut down the instance.

```
rac1-> export ORACLE_SID=devdb1
rac1-> sqlplus / as sysdba

SQL> select
  2  instance_number instance#,
  3  instance_name,
  4  host_name,
  5  status
  6  from v$instance;


 INSTANCE# INSTANCE_NAME    HOST_NAME              STATUS
---------- ---------------- ---------------------- ------------
         1 devdb1           rac1.mycorpdomain.com OPEN

SQL> shutdown abort;
ORACLE instance shut down.
```

Verify that the session has failed over. From the same CRM session you opened previously, execute the queries below to verify that the session has failed over to another instance.

```
SQL> select
  2  instance_number instance#,
  3  instance_name,
  4  host_name,
```

```
  5  status
  6  from v$instance;


INSTANCE# INSTANCE_NAME    HOST_NAME              STATUS
---------- ---------------- --------------------- ------------
         2 devdb2           rac2.mycorpdomain.com OPEN


SQL> select
  2  failover_type,
  3  failover_method,
  4  failed_over
  5  from v$session
  6  where username='SYSTEM';


FAILOVER_TYPE FAILOVER_METHOD FAILED_OVER
------------- --------------- ----------------
SELECT        BASIC           YES
```

Relocate the CRM service back to the preferred instance. After devdb1 is brought back up, the CRM service does not automatically relocate back to the preferred instance. You have to manually relocate the service to devdb1.

```
rac1-> export ORACLE_SID=devdb1
rac1-> sqlplus / as sysdba

SQL> startup
ORACLE instance started.

Total System Global Area  209715200 bytes
Fixed Size                  1218556 bytes
Variable Size             104859652 bytes
Database Buffers          100663296 bytes
Redo Buffers                2973696 bytes
Database mounted.
Database opened.
SQL> show parameter service

NAME                           TYPE        VALUE
------------------------------ ----------- -------------------------
service_names                  string      devdb


rac2-> export ORACLE_SID=devdb2
rac2-> sqlplus / as sysdba

SQL> show parameter service

NAME                           TYPE        VALUE
------------------------------ ----------- -------------------------
service_names                  string      devdb, CRM


rac1-> srvctl relocate service -d devdb -s crm -i devdb2 -t devdb1


SQL> connect system/oracle@devdb1
Connected.
SQL> show parameter service

NAME                           TYPE        VALUE
------------------------------ ----------- -------------------------
service_names                  string      devdb, CRM
```

```
SQL> connect system/oracle@devdb2
Connected.
SQL> show parameter service

NAME                           TYPE        VALUE
------------------------------ ----------- ------------------------
service_names                  string      devdb
```

11. Database Backup and Recovery

The backup and recovery procedure of an Oracle RAC database using Oracle Recovery Manager (RMAN) is no different than that of a single instance database.

In this section you will follow a very simple backup and recovery scenario:

1. Perform a full database backup.
2. Create a table, mytable in the test_d tablespace.
3. At time t1, insert the first record into mytable.
4. At time t2, insert the second record into mytable.
5. At time t3, drop the table, mytable.
6. Recover the test_d tablespace to a point in time.
7. Verify the recovery.

Perform a full database backup.

```
rac1-> rman nocatalog target /

Recovery Manager: Release 10.2.0.1.0 - Production on Mon Nov 13 18:15:09 2006

Copyright (c) 1982, 2005, Oracle.  All rights reserved.

connected to target database: DEVDB (DBID=511198553)
using target database control file instead of recovery catalog

RMAN> configure controlfile autobackup on;

RMAN> backup database plus archivelog delete input;
```

Create a table, mytable in the test_d tablespace.

```
19:01:56 SQL> connect system/oracle@devdb2
Connected.
19:02:01 SQL> create table mytable (col1 number) tablespace test_d;

Table created.
```

At time, t1, insert the first record into mytable.

```
19:02:50 SQL> insert into mytable values (1);

1 row created.

19:02:59 SQL> commit;

Commit complete.
```

At time, t2, insert the second record into mytable.

```
19:04:41 SQL> insert into mytable values (2);

1 row created.

19:04:46 SQL> commit;

Commit complete.
```

At time, t3, drop the table, mytable.

```
19:05:09 SQL> drop table mytable;

Table dropped.
```

Recover the test_d tablespace to a point in time.

Create an auxiliary directory for the auxiliary database.

```
rac1-> mkdir /u01/app/oracle/aux


RMAN> recover tablespace test_d
2> until time "to_date('13-NOV-2006 19:03:10','DD-MON-YYYY HH24:MI:SS')"
3> auxiliary destination '/u01/app/oracle/aux';

RMAN> backup tablespace test_d;

RMAN> sql 'alter tablespace test_d online';
```

Verify the recovery.

```
19:15:09 SQL> connect system/oracle@devdb2
Connected.
19:15:16 SQL> select * from mytable;

      COL1
----------
         1
```

12. Explore Oracle Enterprise Manager (OEM) Database Console

Oracle Enterprise Manager Database Console provides a really nice integrated and comprehensive GUI interface to administering and managing your cluster database environment. You can perform virtually any tasks from within the console.

To access the Database Console, open a Web browser and enter the URL below.
Log in as sysman and enter the password you have chosen earlier during the database installation.

http://rac1:1158/em

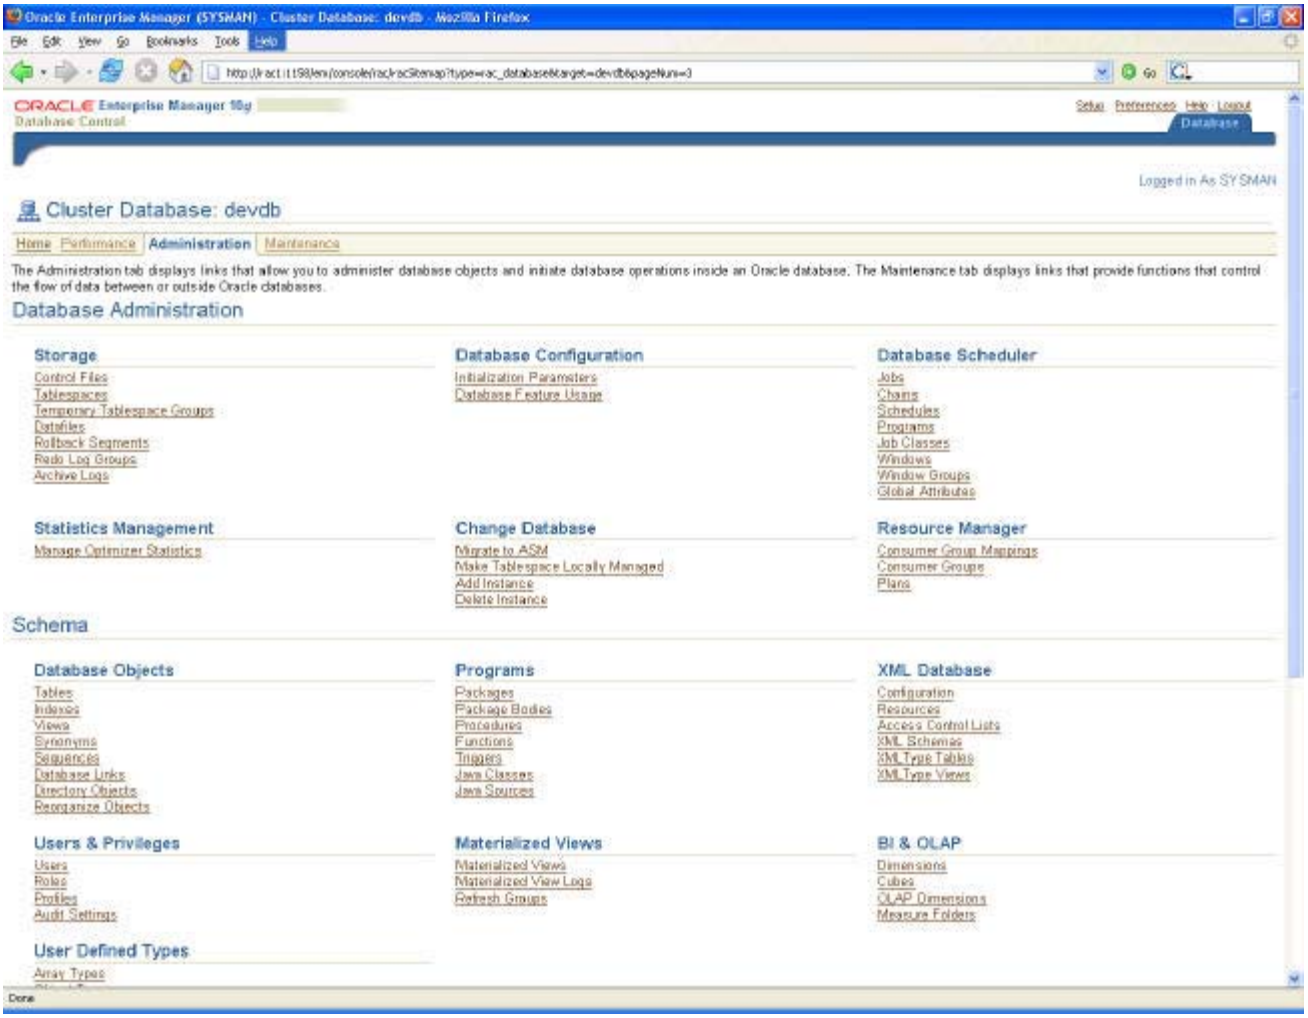

Start and stop the Database Console.

```
rac1-> emctl stop dbconsole
TZ set to US/Eastern
Oracle Enterprise Manager 10g Database Control Release 10.2.0.1.0
Copyright (c) 1996, 2005 Oracle Corporation.  All rights reserved.
http://rac1.mycorpdomain.com:1158/em/console/aboutApplication
Stopping Oracle Enterprise Manager 10g Database Control ...
 ...  Stopped.

        rac1-> emctl start dbconsole
TZ set to US/Eastern
Oracle Enterprise Manager 10g Database Control Release 10.2.0.1.0
Copyright (c) 1996, 2005 Oracle Corporation.  All rights reserved.
http://rac1.mycorpdomain.com:1158/em/console/aboutApplication
Starting Oracle Enterprise Manager 10g Database Control
.................. started.
------------------------------------------------------------------
Logs are generated in directory
/u01/app/oracle/product/10.2.0/db_1/rac1_devdb1/sysman/log
```

Verify the status of Database Console.

```
rac1-> emctl status dbconsole
TZ set to US/Eastern
Oracle Enterprise Manager 10g Database Control Release 10.2.0.1.0
```

```
Copyright (c) 1996, 2005 Oracle Corporation.  All rights reserved.
http://rac1.mycorpdomain.com:1158/em/console/aboutApplication
Oracle Enterprise Manager 10g is running.
------------------------------------------------------------------
Logs are generated in directory
/u01/app/oracle/product/10.2.0/db_1/rac1_devdb1/sysman/log

rac1-> emctl status agent
TZ set to US/Eastern
Oracle Enterprise Manager 10g Database Control Release 10.2.0.1.0
Copyright (c) 1996, 2005 Oracle Corporation.  All rights reserved.
---------------------------------------------------------------
Agent Version     : 10.1.0.4.1
OMS Version       : 10.1.0.4.0
Protocol Version  : 10.1.0.2.0
Agent Home        : /u01/app/oracle/product/10.2.0/db_1/rac1_devdb1
Agent binaries    : /u01/app/oracle/product/10.2.0/db_1
Agent Process ID  : 10263
Parent Process ID : 8171
Agent URL         : http://rac1.mycorpdomain.com:3938/emd/main
Started at        : 2006-11-12 08:10:01
Started by user   : oracle
Last Reload       : 2006-11-12 08:20:33
Last successful upload                       : 2006-11-12 08:41:53
Total Megabytes of XML files uploaded so far :     4.88
Number of XML files pending upload           :        0
Size of XML files pending upload(MB)         :     0.00
Available disk space on upload filesystem    :    71.53%
---------------------------------------------------------------
Agent is Running and Ready
```

13. Common Issues

Below is a summary list of issues and resolutions you may find useful.

Issue 1: Cannot activate Ethernet devices.
Error message, “Cannot activate network device eth0! Device eth0 has different MAC address than expected, ignoring.”

Resolution:
The MAC address reported by “ifconfig” does not match /etc/sysconfig/network-scripts/ifcfg-eth0. You can either update the file with the new MAC address or simply probe for the new MAC address via the `system-config-network` tool.

Issue 2: Cannot generate OCFS2 configuration file.
Error message, “Could not start cluster stack. This must be resolved before any OCFS2 filesystem can be mounted” when attempting to generate OCFS2 configuration file.

Resolution:
Execute ocfs2console as the root user instead of the oracle user.

Issue 3: Cannot install Oracle Clusterware or Oracle Database software on remote node.
Error message, “ /bin/tar: ./inventory/Components21/oracle.ordim.server/10.2.0.1.0: time stamp 2006-11-04 06:24:04 is 25 s in the future” during Oracle Clusterware software installation.

Resolution:
Synchronize the time between guest OS and host OS by installing VMware Tools and include the options, “clock=pit nosmp noapic nolapic” in /boot/grub/grub.conf. Refer to Section 3 for more information.

Issue 4: Cannot mount OCFS2 file system.
Error message, “mount.ocfs2: Transport endpoint is not connected while mounting” when attempting to mount the ocfs2 file system.

Resolution:
Execute `/usr/bin/system-config-securitylevel` to disable firewall.

Issue 5: Cannot start ONS resource.
Error message, “CRS-0215: Could not start resource ‘ora.rac2.ons’” when VIPCA attempts to start ONS application resource.

Resolution:
ONS attempts to access localhost but cannot resolve the IP address. Add the following entry in /etc/hosts.

127.0.0.1 localhost

Conclusion

Hopefully this guide has provided you a quick and free method of building a clustered Oracle database environment using VMware Server. Take advantage of the freely available software, and start learning and experimenting with Oracle RAC on Enterprise Linux!

Vincent Chan (vkchan99@yahoo.com) is a Senior Consultant at MSD Inc. He is an Oracle Certified Master DBA with more than ten years of experience architecting and implementing Oracle solutions for various clients.