



Oracle Databases on VMware RAC Deployment Guide

December 2011

© 2011 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at <http://www.vmware.com/download/patents.html>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com

Contents

1.	Introduction	5
2.	Infrastructure Environment.....	6
2.1	Hardware.....	6
2.2	Software Environment.....	6
2.3	VMware Virtual Machines	6
2.4	Virtual Disk Layout	7
2.5	Oracle RAC Database Environment	8
2.6	Network Configuration	8
2.7	Physical Architecture	10
2.8	Logical Architecture	11
3.	Overview of Deployment Steps	12
4.	Configure NTP Client on All ESX Hosts	13
5.	Prepare a Virtual Machine for the First RAC Node.....	15
5.1	Create a Virtual Machine	15
5.2	Add Two NICs	15
5.3	Add CRS and Voting Disk.....	16
5.4	Install the Guest Operating System	20
5.5	Install VMware Tools.....	21
5.6	Timekeeping with the Linux Guest Operating System.....	22
6.	Prepare Subsequent Virtual Machines/RAC Nodes	23
6.1	Clone the First Virtual Machine/RAC Node	23
6.2	Change the Hostname and Assign an IP Address	23
6.3	Configure CSR and Voting Disks.....	24
6.4	Format the Added Disks Using fdisk.....	26
7.	Install Oracle Grid Infrastructure	27
8.	Verify the Grid Infrastructure Installation	29
9.	Add DATA and REDO Disks to Virtual Machine RAC Nodes	31
10.	Create DATA and REDO ASM Disk Groups as Grid User	35
11.	Install and Create the Oracle Database 11g R2 RAC	36
12.	References	37
	Appendix A: Deployment of Oracle RAC with RDM	38

1. Introduction

This document describes an installation example of Oracle Real Application Clusters (RAC) on the VMware vSphere® platform, for Oracle DBAs planning to virtualize Oracle RAC on the VMware platform. The reader is assumed to have some hands-on experience with VMware vSphere® Client™. A database administrator should work with a VMware and storage administrator to successfully install Oracle RAC on VMware technology.

Some aspects of the virtualized installation are the same as with a physical installation:

- After virtual machines are created and correctly configured, the installation of the guest operating system and Oracle RAC software is the same as with a physical installation.
- Storage layout guidelines for Oracle Database on VMware technology, with respect to the number of LUNs and maximizing I/O performance, are the same as with a physical installation.

Because installation details are specific to the software versions used in this deployment, there might be differences from this document if using different versions. In some cases the installation steps in this document are at high level – in these situations the detailed procedures can be found in referenced Oracle and VMware installation guides and knowledge base articles.

The sizes of VMware ESX® virtual machines (memory, CPU, and virtual CPU) used in this document are not based on any specific business scenario. These configuration parameters are different in customer deployments and depend on customer-specific sizing requirements.

The major software components used in this installation are:

- VMware ESXi™ 5.0 – Abstracts processor, memory, storage, and networking resources into multiple virtual machines, increasing hardware utilization.
- VMware vCenter Server™ 5.0 Update 1 – Delivers comprehensive virtualization, management, resource optimization, application availability and operational automation into an integrated product suite.
- vSphere Client 5.0 – Delivers comprehensive virtualization, management, resource optimization, application availability and operational automation into an integrated product suite.
- Oracle Database 11g Release 2 Grid Infrastructure (GRID) – Oracle RAC uses multiple Oracle instances on multiple nodes (servers) to connect to single database.
- Oracle Database 11g Release 2 – The latest RDBMS Server from Oracle.
- Oracle Automatic Storage Management (ASM) – ASM integrates the file system with the volume manager designed for Oracle files.
- Oracle Enterprise Linux x86_64 5.5.

For the VMware Oracle support policy, go to <http://www.vmware.com/support/policies/oracle-support.html>.

2. Infrastructure Environment

This section summarizes the infrastructure environment and covers hardware, software versions, virtual machines deployed, storage, and network configuration.

2.1 Hardware

Component	Quantity	Configuration
ESXi Hosts	4	Cisco UCS blades 2x 6-core CPUs, 196GB RAM, and 4 NICs
Storage	1	EMC VNX5500 Unified Storage
Ethernet Switches	4	Cisco 3750
FC Switches	2	Cisco MDS 9134
Host Based Adapters	8	2 per physical server

2.2 Software Environment

Component	Version	Purpose
vSphere	5.0	Hypervisor hosting all virtual machines
VMware vCenter™	5.0	Management of vSphere
Oracle Enterprise Linux x86_64	5.5	Guest operating system for virtual machines running Oracle Database server
Oracle Database 11g (with Oracle RAC and Oracle Grid Infrastructure) Enterprise Edition	11.2.0.1	Oracle cluster database software for grid computing
Microsoft Windows Server	2008 R2	Operating system to host vCenter Server

2.3 VMware Virtual Machines

Virtual Machine Role	Quantity	Configuration
Oracle RAC nodes	4	12 vCPU, 160GB RAM, Oracle Enterprise Linux x86_64 5.5
vCenter	1	2 vCPUs, 8GB RAM, Windows Server 2008 R2

2.4 Virtual Disk Layout

Datstore Name	Virtual Disk on ESX	Guest OS Device Name	Virtual Device	Virtual SCSI Driver	VMDK Size (GB)	Purpose
VMDATASTORE	VMDK – Hard disk 1	/dev/sda	SCSI 0:0	LSI Logic	50	Oracle Enterprise Linux 5.5 OS and Oracle binaries
CRS1	VMDK – Hard disk 2	Shared disk /dev/sdb1	SCSI 1:0	Paravirtual	20	CRS and voting disk
CRS2	VMDK – Hard disk 3	Shared disk /dev/sdc1	SCSI 1:1	Paravirtual	20	CRS and voting disk
CRS3	VMDK – Hard disk 4	Shared disk /dev/sdd1	SCSI 1:2	Paravirtual	20	CRS and voting disk
VMFSDATA01	VMDK – Hard disk 5	Shared disk /dev/sde	SCSI 1:3	Paravirtual	300	RAC database DATA
VMFSDATA02	VMDK – Hard disk 6	Shared disk /dev/sdf1	SCSI 1:4	Paravirtual	300	RAC database DATA
VMFSDATA03	VMDK – Hard disk 7	Shared disk /dev/sdg1	SCSI 1:5	Paravirtual	300	RAC database DATA
VMFSDATA04	VMDK – Hard disk 8	Shared disk /dev/sdh1	SCSI 1:6	Paravirtual	300	RAC database DATA
VMFSDATA05	VMDK – Hard disk 9	Shared disk /dev/sdi1	SCSI 1:8	Paravirtual	300	RAC database DATA
VMFSDATA06	VMDK – Hard disk 10	Shared disk /dev/sdj1	SCSI 1:9	Paravirtual	300	RAC database DATA
VMFSDATA07	VMDK – Hard disk 11	Shared disk /dev/sdk1	SCSI 1:10	Paravirtual	300	RAC database DATA
VMFSDATA08	VMDK – Hard disk 12	Shared disk /dev/sdl1	SCSI 1:11	Paravirtual	300	RAC database DATA
VMFSDATA09	VMDK – Hard disk 13	Shared disk /dev/sdm1	SCSI 1:12	Paravirtual	300	RAC database DATA
VMFSDATA10	VMDK – Hard disk 14	Shared disk /dev/sdn1	SCSI 1:13	Paravirtual	300	RAC database DATA

Datastore Name	Virtual Disk on ESX	Guest OS Device Name	Virtual Device	Virtual SCSI Driver	VMDK Size (GB)	Purpose
VMFSREDO01	VMDK – Hard disk 15	Shared disk /dev/sdo1	SCSI 2:0	Paravirtual	64	RAC database REDO
VMFSREDO02	VMDK – Hard disk 16	Shared disk /dev/sdp1	SCSI 2:1	Paravirtual	64	RAC database REDO
VMFSREDO03	VMDK – Hard disk 17	Shared disk /dev/sdq1	SCSI 2:2	Paravirtual	64	RAC database REDO
VMFSREDO04	VMDK – Hard disk 18	Shared disk /dev/sdr1	SCSI 2:3	Paravirtual	64	RAC database REDO

2.5 Oracle RAC Database Environment

Node Name	ASM Instance	RAC Instance	Database	ASM Storage
VMORARAC1.vmware.com	+ASM1	VMORARAC1	VMORARAC	ASM file system
VMORARAC2.vmware.com	+ASM2	VMORARAC2	VMORARAC	ASM file system
VMORARAC3.vmware.com	+ASM3	VMORARAC3	VMORARAC	ASM file system
VMORARAC4.vmware.com	+ASM4	VMORARAC4	VMORARAC	ASM file system

2.6 Network Configuration

All four nodes in the cluster must be able to communicate with each other and with external clients using TCP/IP. Communication between clients and the nodes in the cluster is across the public network. All nodes need a network adapter configured for the public network.

To enable high availability and failover, a virtual IP (VIP) address is also required for each node. A VIP address can be moved between nodes in case of a failure. CRS manages the VIP addresses. To support a virtual IP address, both nodes require an unused IP address that is compatible with the public network's subnet and subnet mask.

For communications between instances running on the four nodes, a private network is required. This private network connects only the nodes in the cluster and cannot be accessed from outside the cluster. All nodes need a separate network adapter configured for this private network.

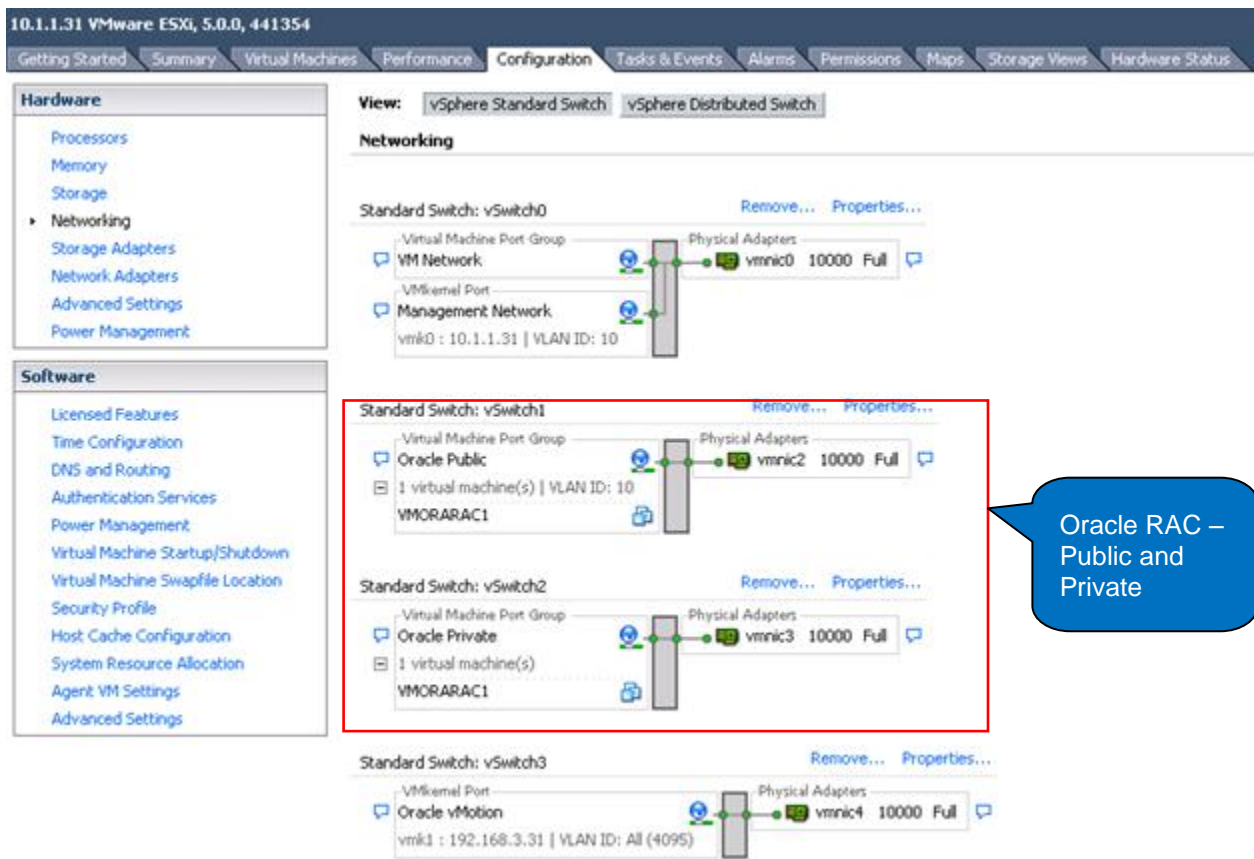
This deployment uses the following network configuration:

- Four physical network interface cards (NIC) – vmnic1 and vmnic5 for Oracle Public and Oracle Private, vmnic0 for ESX Console and vmnic2 for VMware vSphere® vMotion® operations.
- Two virtual switches – Oracle Public (vSwitch1) and Oracle Private (vSwitch2).
- Three static IP addresses per server – An example for node 1 (VMORARAC1) is as follows:
 - One public static IP address (eth0): 10.1.1.30
 - One private static IP address (eth1): 192.168.2.30
 - One virtual static IP address(VIP): 10.1.1.40

The physical NIC vmnic0 (Gigabit recommended) is connected to the public network which is then connected to public virtual switch vSwitch1. The Service Console is also connected to vSwitch0, and vMotion operations use vSwitch3. The physical NIC vmnic1 is connected to the private network which is connected to the private virtual switch vSwitch2. The RAC network is also connected to vSwitch1. To create virtual switches vSwitch1 and vSwitch2, refer to Section 2 of *vSphere Networking* (<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-networking-guide.pdf>). VMware recommends vSphere's Distributed Switch (vDS) as it spans many vSphere hosts and aggregates networking to a centralized cluster level administration and monitoring through VMware vCenter.

Each server in the cluster requires three IP addresses. One IP address is for the public network connected to the public network virtual switch (Oracle Public), and a second IP address is for the private cluster interconnects connected to the private network virtual switch (Oracle Private). The third IP address, which is a virtual IP address, is configured by Oracle when Oracle Grid Infrastructure is installed.


Figure 1. Network Configuration



2.7 Physical Architecture

The following architecture diagram depicts the physical architecture of the Oracle RAC deployment environment on vSphere 5.0.

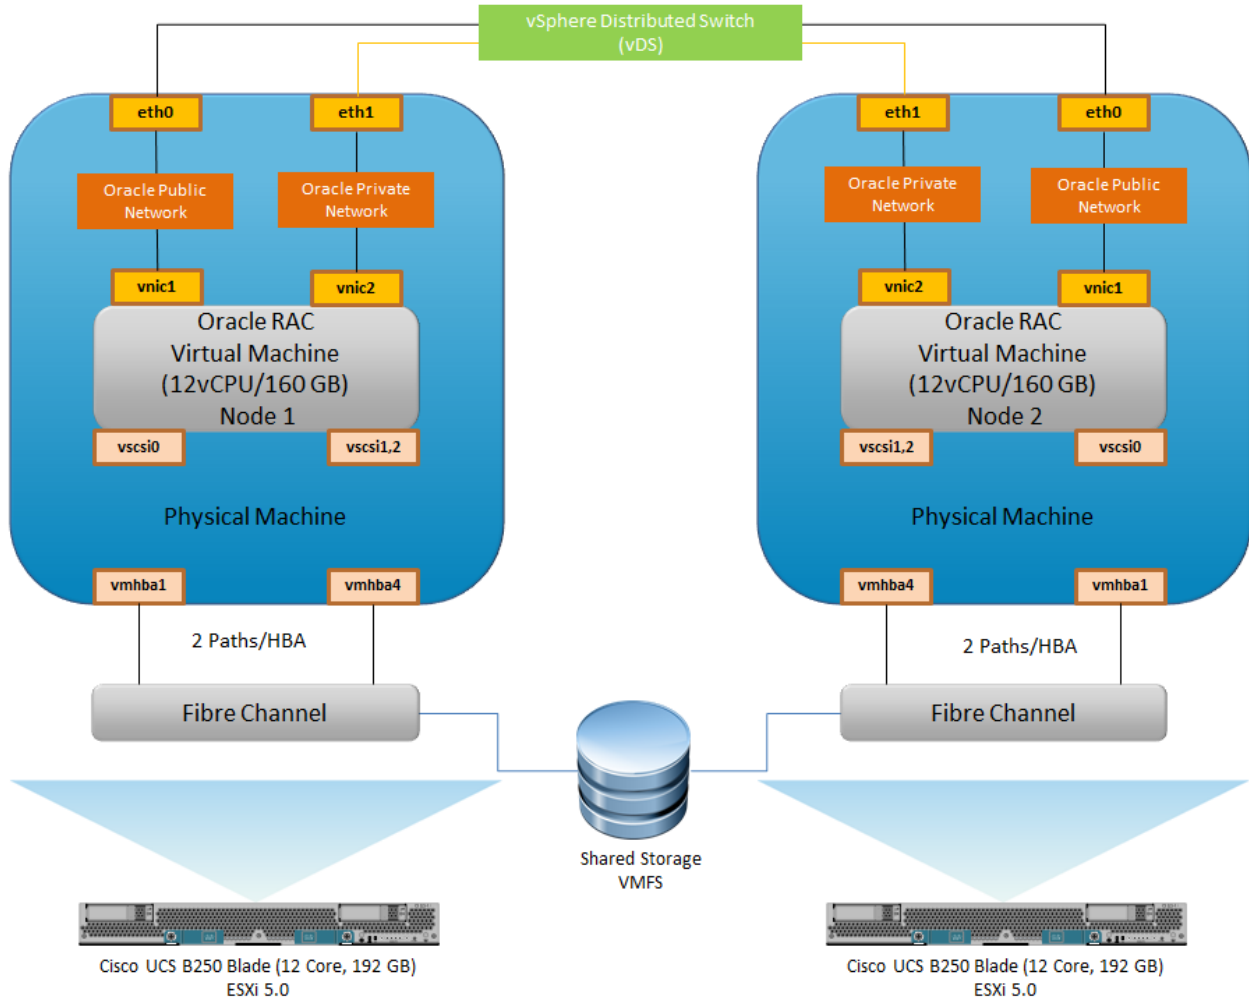
Figure 2. Physical Architecture of Oracle RAC on vSphere



2.8 Logical Architecture

The following diagram describes the logical architecture of a two-node RAC on vSphere 5.0. The actual deployment uses four nodes.

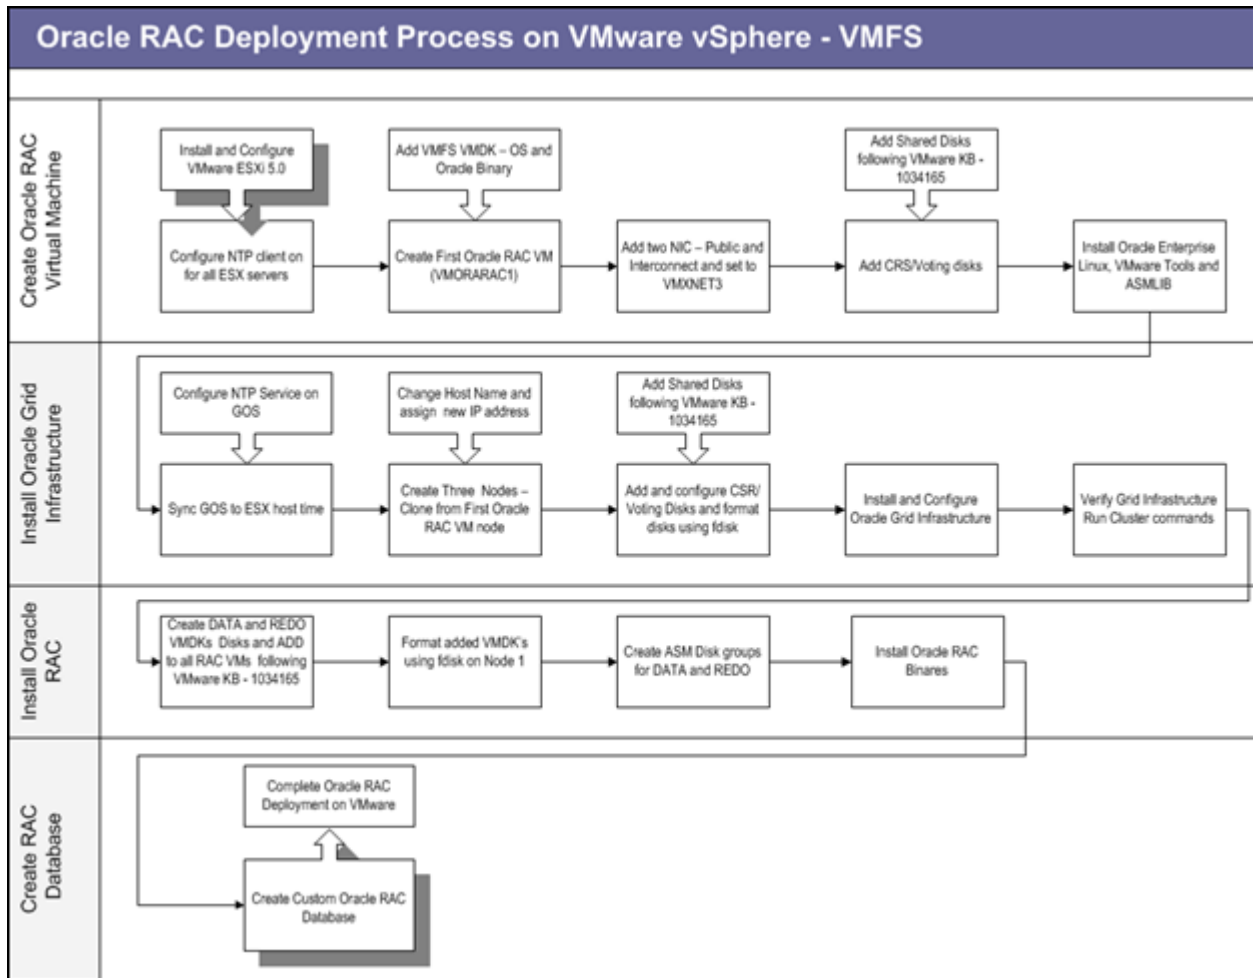
Figure 3. Logical Architecture of Two-Node Oracle RAC on vSphere



3. Overview of Deployment Steps

The following process diagram illustrates the deployment steps for deploying a four-node Oracle RAC on vSphere 5.0 with VMFS. For deploying Oracle RAC with RDM refer Appendix A: Deployment of Oracle RAC with RDM.

Figure 4. Deployment Steps for Four-Node Oracle RAC on vSphere with VMFS



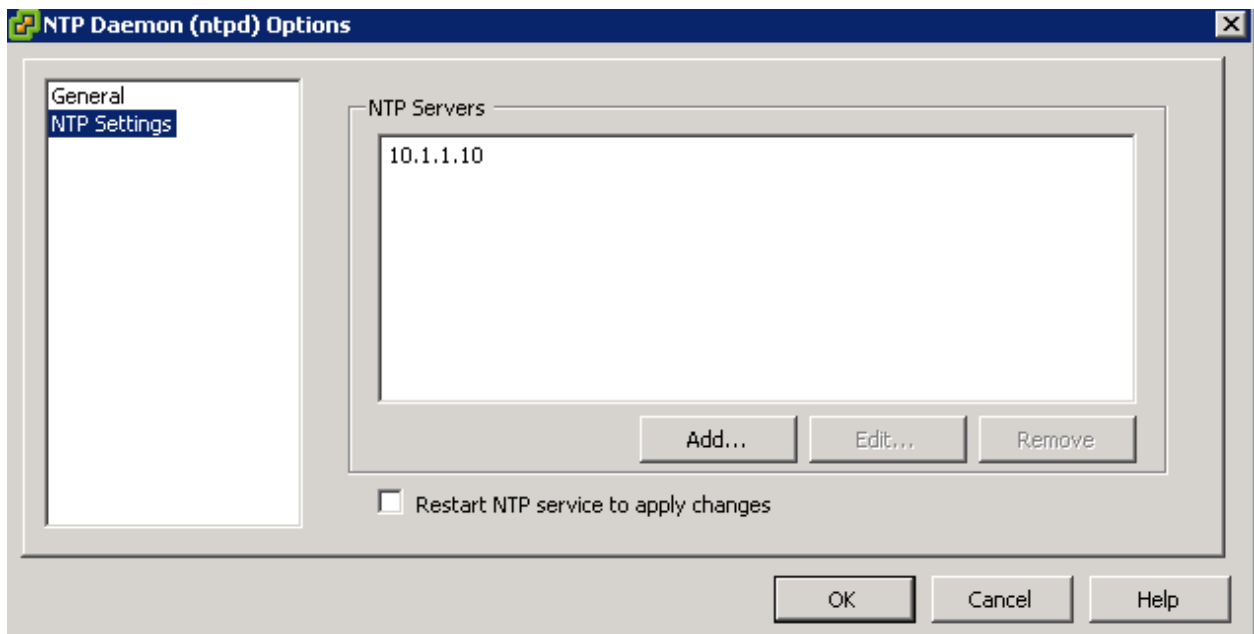
4. Configure NTP Client on All ESX Hosts

Oracle Database requires the correct system time. VMware Tools has an option to allow virtual machines to synchronize the time with the ESX host they are running on. Synchronize the system clocks of all ESX hosts in a resource cluster.

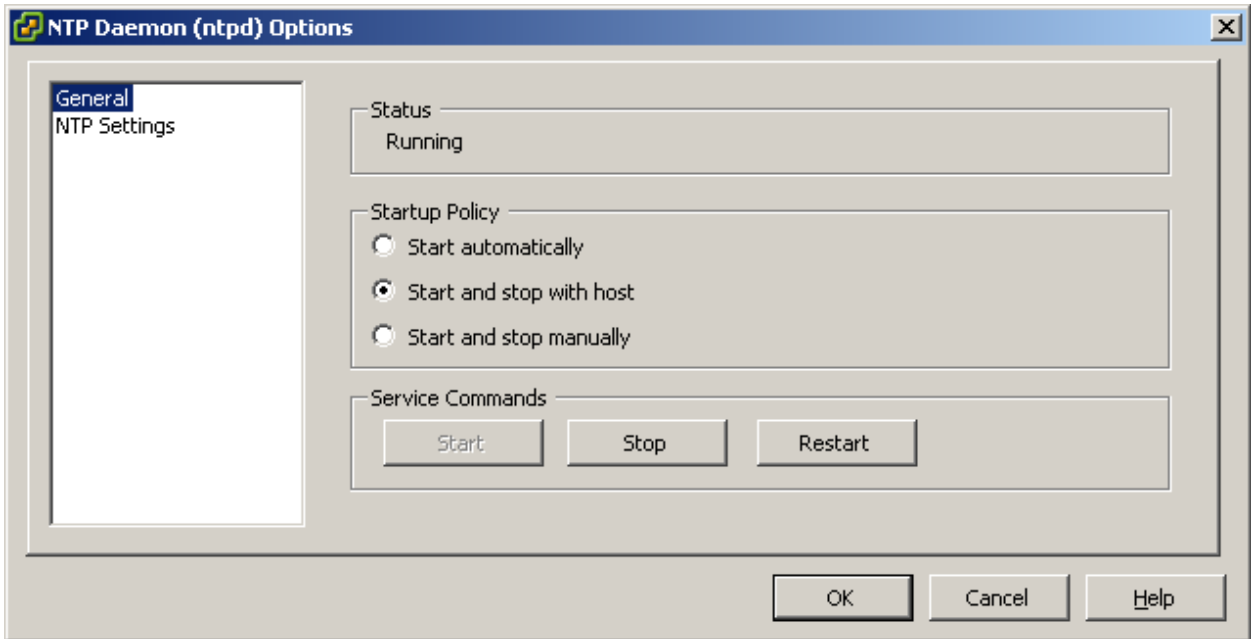
ESX hosts include NTP software and start the ntpd service by default. ESX hosts are not configured to make use of NTP upon installation, nor is there anything in the standard installation process which gives the administrator the option to set this. In the four-node example, time synchronization is done for all four ESX hosts.

To configure the NTP client settings

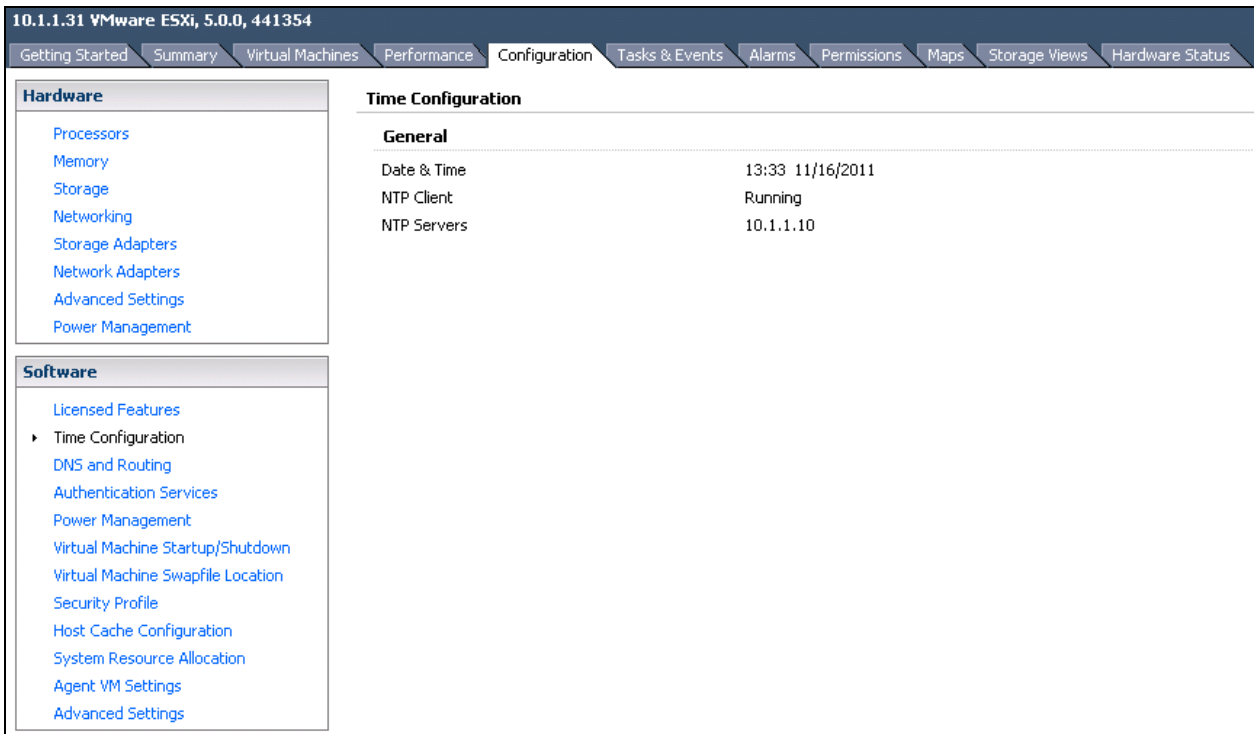
1. Click **Properties** to display the **NTP Daemon (ntpd) Options** dialog box.
2. Click **NTP Settings**.
3. Click **Add** to add NTP servers.



4. Click **General**.
5. For the **Startup Policy**, select **Start and stop with host**.
6. Click **OK**.



7. After the configuration is complete, the configuration screen looks similar to the following illustration.



5. Prepare a Virtual Machine for the First RAC Node

This section describes the following steps:

- Creation of the first virtual machine, which is the first RAC node, using vSphere Client. This is created on VMFS and corresponds to the root drive. Two NICs are assigned for the public and private networks. Three datastores, CRS1, CRS2, and CRS3, are created from the storage LUNs and three VMDKs are created and assigned to the virtual machine for CRS and voting. Only one VMDK is required but three are used for redundancy.
- Installation of the Linux operating system in the virtual machine followed by installation of VMware Tools.
- Synchronization of the time within the guest OS with the ESX host.

5.1 Create a Virtual Machine

Four ESX hosts have been created and are visible in vCenter.

To create a new virtual machine

1. Log in to vCenter using vSphere Client.
2. Select a host and click **Create a new virtual machine** on the **Getting Started** tab.
3. In the **Create New Virtual Machine** wizard, select **Typical** for the **Configuration**, and click **Next**.
4. On the **Name and Location** page, enter **VMORARAC1** for the **Name**, and click **Next**.
5. On the **Datastore** page, select **VMDATASTORE**, and click **Next**.
6. On the **Guest Operating System** page, select **Linux** from the **Guest Operating System** list. For the **Version**, select **Oracle Linux 4/5 (64-bit)** from the drop-down menu. Click **Next**.
7. On the **Create a Disk** page, set the **Virtual Disk Size** to **50GB**, and click **Next**.
8. On the **Ready to Complete** page, select **Edit the virtual machine settings before completion**. Click **Continue**.
9. On the **Virtual Machine Properties** page for the VMORARAC1 virtual machine, select **8 CPUs** and **128GB Memory**.
10. Click **Finish**.

5.2 Add Two NICs

In this section, two NICs are added. One NIC is for the public network, and the other is for the private network. The procedure for each is similar, differing only in the selection of **RAC Public** or **Private** for the network label.

To add a NIC to the virtual machine for the public network

1. Select a virtual machine.
2. Click Edit settings to display the **Virtual Machine Properties** pane.
3. Click **Add**.
4. Select **Ethernet Adaptor** and click **Next** to display the **Add Hardware** wizard.
5. On the **Network connection** page, select **VMXNET 3** from the **Adapter Type** drop-down menu.
6. In the **Network Connection** section of the **Network connection** page, select **Named network with specified label**.

7. In the **Named network with specified label** drop-down menu, select **RAC Public**.
8. In the **Device Status** section of the **Network connection** page, select **Connect at power on**.
9. On the **Ready to Complete** page, click **Finish**, and then click **OK**.

To add a NIC to the virtual machine for the private network

1. Repeat steps 1 through 6 from the preceding procedure.
2. In the **Named network with specified label** drop-down menu, select **Private**.
3. Continue with steps 8 and 9 from the preceding procedure.

5.3 Add CRS and Voting Disk

Three 20GB disks are created and VMDKs are created and added to the virtual machine for RAC CRS and voting. Although only one is required, three are used for redundancy. The datastores are created from the assigned LUNs which are attached to a SCSI controller configured as VMware Paravirtual. The SCSI bus sharing policy of the SCSI controller is set to none.

The procedures in this section describe the steps needed to create disks for one virtual machine. Repeat the steps in this section for each of the four virtual machines named VMORARAC1, VMORARAC2, VMORARAC3, and VMORARAC4.

To create one 20GB datastore for a virtual machine

1. In vSphere Client, select a virtual machine.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. Click **Add** to display the **Add Hardware** wizard.
4. In the **Add Hardware** wizard select **Hard Disk** for the **Device Type**, and click **Next**.
5. In the **Disk** section of the **Select a Disk** page, select **Create a new virtual disk**. Set the disk size to **20GB**, select the **Support clustering features such as Fault Tolerance** check box, and select CRS1 to specify a datastore. Also, for **vSphere 5**, in the **Disk Provisioning** section select **Thick Provision Eager Zeroed**. Click **Next**.
6. On the **Compatibility Mode** page, click **Next**.
7. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:0)**. In the **Mode** section, select **Independent**. Click **Next**.
8. On the **Ready to Complete** page, click **Finish**.
9. Click **OK**.

To create a second 20GB datastore for the virtual machine

1. Select the virtual machine that was selected in the preceding step 1.
2. Repeat steps 2 through 6 in the preceding procedure to create one 20GB datastore.
3. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:1)**. In the **Mode** section, select **Independent**. Click **Next**.
4. Repeat steps 8 and 9 of the preceding procedure to create one 20GB datastore.

To create a third 20GB datastore for the virtual machine

1. Select the virtual machine that was selected in the preceding step 1.
2. Repeat steps 2 through 6 in the preceding procedure to create one 20GB datastore.
3. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:2)**. In the **Mode** section, select **Independent**. Click **Next**.
4. Repeat steps 8 and 9 of the preceding procedure to create one 20GB datastore.

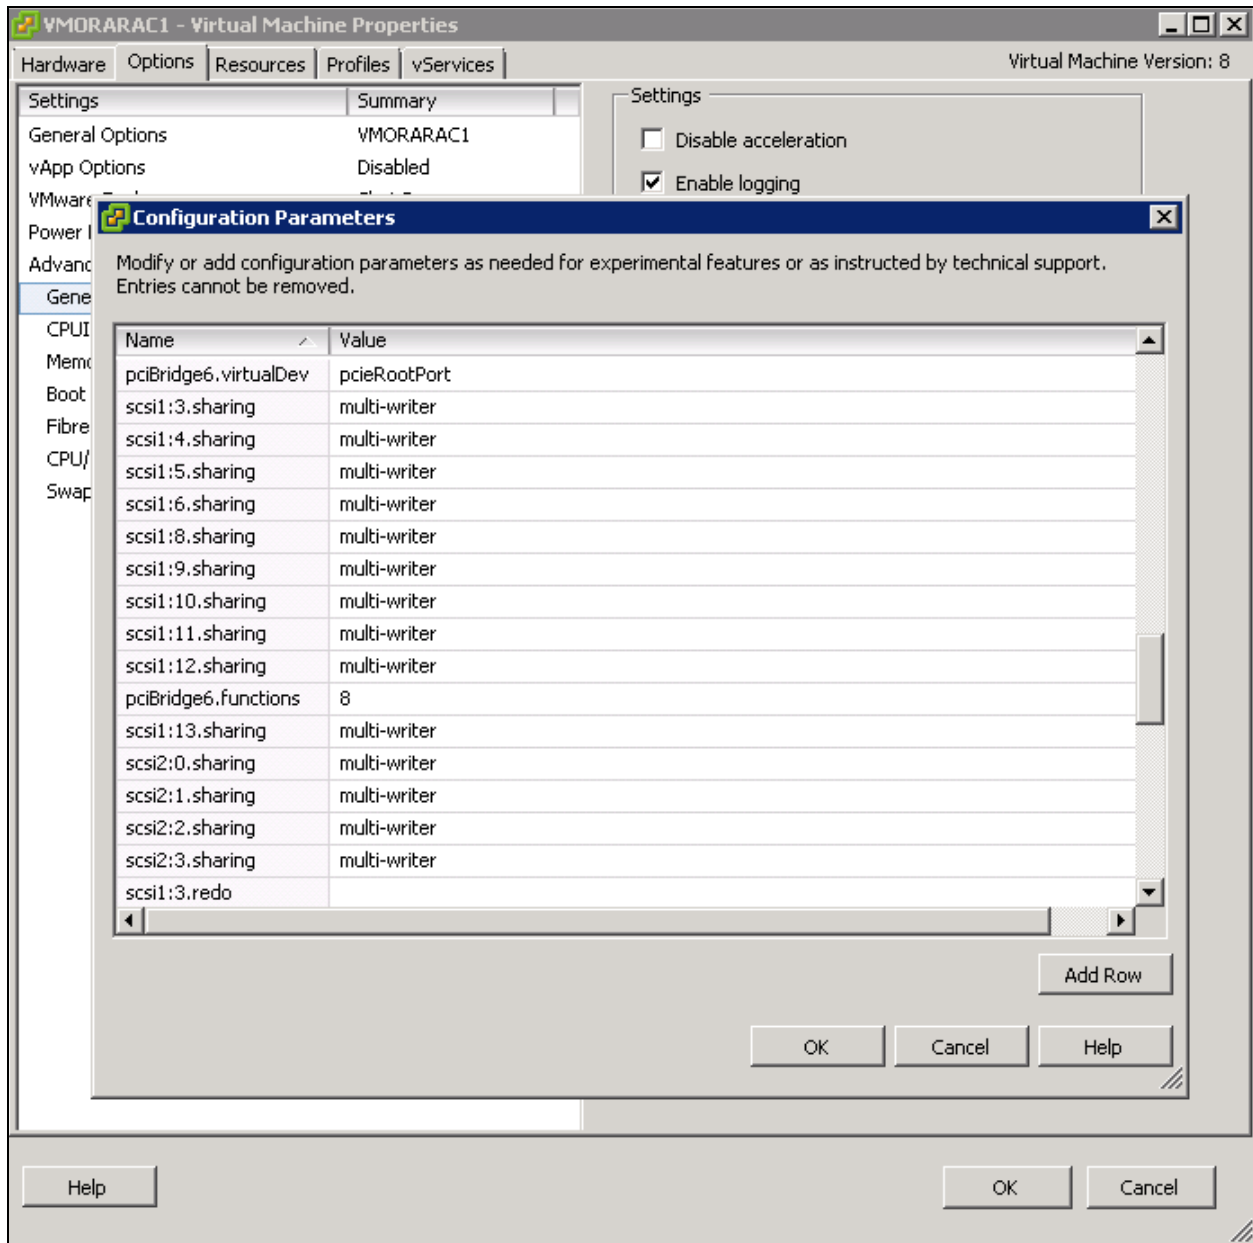
To set the SCSI controller type to paravirtual

1. Select the virtual machine that was selected in the preceding step 1.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. In the device list, select **SCSI controller 1**.
4. In the **SCSI Bus Sharing** section, select **None**, keep the default selection.
5. In the **SCSI Controller Type** section, click **Change Type**.
6. Select **VMware Paravirtual**.
7. Click **OK**, and click **OK** again.


Add configuration parameters for each Oracle RAC virtual machine as described in *Disabling simultaneous write protection provided by VMFS using the multi-writer flag* (<http://kb.vmware.com/kb/1034165>).

To add configuration parameters to the virtual machines

1. In vSphere Client, select a virtual machine.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. Click the **Options** tab, click **Advanced**, and click **General**.
4. Click **Configuration Parameters**.
5. On the **Configuration Parameters** page, add rows and entries for the shared disks in the virtual machine's configuration parameters as shown in the following screenshot.
6. Repeat steps 1 through 5 for each of the remaining virtual machines.



You can view the virtual machine properties for the VMORARAC2 virtual machine in vCenter as in the following screenshot.



5.4 Install the Guest Operating System

The Oracle Enterprise Linux 5.5 x64 operating system is installed from an ISO image that has been loaded into a VMFS datastore.

Note: This document is specific to Linux guest operating system. Oracle RAC 11.2.0.2 and above versions can be deployed and supported on the list of guest operation systems using VMware multi-writer flag (<http://kb.vmware.com/kb/1008027>).

To install Linux as the guest operating system in a virtual machine

1. In vSphere Client, select a virtual machine.
2. Click **Edit settings**.
3. In the **Hardware** tab of the **Virtual Machine Properties** pane, select **CD/DVD Drive 1**.
4. In the **Device Status** section of this pane, select the **Connect at power on** check box.
5. In the **Device Type** section of this pane, select **Datastore ISO File** and enter the path to your Linux ISO image file.
6. In the **Virtual Device Node** section of this pane, select **IDE (1:0) CD/DVD Drive 1** from the drop-down menu.
7. Click **OK**.
8. Power on the virtual machine.
9. Install Linux with the following settings.

Custom partitioning scheme:

```
/dev/sda1 ext3 15GB mounted on /  
/dev/sda2 ext3 25GB mounted on /u01  
/dev/sda3 swap 10GB
```

Installed packages – use the default package selection, plus the following:

- oracle-validated-1.0.0-22.el5.x86_64
 - oracleasm-2.6.18-194.el5-2.0.5-1.el5.x86_64
 - oracleasm-support-2.1.3-1.el5.x86_64
 - systat-7.0.2-3.el5.x86_64
10. Download and install the ASM library from the Oracle Web site:
<http://www.oracle.com/technetwork/topics/linux/asmlib/index-101839.html>.

11. View the installed packages with this command:

```
run command "rpm -q <package name>
```

12. Confirm that the following packages have been installed:

- binutils-2.17.50.0.6
- compat-libstdc++-33-3.2.3
- elfutils-libelf-0.125
- elfutils-libelf-devel-0.125
- elfutils-libelf-devel-static-0.125
- gcc-4.1.2
- gcc-c++-4.1.2
- glibc-2.5-24
- glibc-common-2.5
- glibc-devel-2.5
- glibc-headers-2.5
- ksh-20060214
- libaio-0.3.106
- libaio-devel-0.3.106
- libgcc-4.1.2
- libstdc++-4.1.2
- libstdc++-devel 4.1.2
- make-3.81
- sysstat-7.0.2
- unixODBC-2.2.11
- unixODBC-devel-2.2.11

5.5 Install VMware Tools

VMware Tools is a suite of utilities that enhance the performance of the virtual machine's guest operating system and improve management of the virtual machine. VMware Tools must be installed inside the guest OS. The installation steps are described in *Installing VMware Tools in a Linux virtual machine using Red Hat Package Manager (RPM)* (<http://kb.vmware.com/kb/1018392>).

To install VMware Tools

1. Confirm that the Linux virtual machine is powered on.
2. In vSphere Client, right-click the virtual machine and select **Guest > Install VMware Tools**.
3. In the Linux guest OS, create a mount point and run the following command:

```
mkdir /mnt/cdrom
```

4. Mount the CD-ROM with the following command:

```
mount /dev/cdrom /mnt/cdrom
```

5. Install VMware Tools using RPM, with the following command:

```
rpm -ivh /mnt/cdrom/VMwareTools-8.3.2-257589.rpm
```

As an example, "8.3.2-257589" is the version of VMware Tools used here.

6. Configure VMware Tools, by running the following command:

```
/usr/bin/vmware-config-tools.pl
```

7. Unmount the CD-ROM with the following command:

```
umount /mnt/cdrom
```

8. Click **VM** in the virtual machine menu, and click **Guest > End VMware Tools Install**.

5.6 Timekeeping with the Linux Guest Operating System

Follow the recommendations in *Timekeeping best practices for Linux guests* (<http://kb.vmware.com/kb/1006427>) to configure NTP, and when using NTP in the guest, disable VMware Tools *periodic* time synchronization.

To disable VMware Tools periodic time synchronization, perform one of these options

- Set `tools.syncTime = "False"` in the configuration file (`.vmx` file) of the virtual machine.
- Deselect **Time synchronization between the virtual machine and the host operating system** in the VMware Tools toolbox GUI of the guest operating system.
- Run the `vmware-guestd --cmd "vmx.set_option synctime 1 0"` command in the guest operating system.

For ESX 4.x, use the following Linux commands:

To display the current status of the service

```
vmware-toolbox-cmd timesync status
```

To disable periodic time synchronization

```
vmware-toolbox-cmd timesync disable
```

6. Prepare Subsequent Virtual Machines/RAC Nodes

Now that the first RAC node virtual machine (VMORARAC1) has been created, this virtual machine can be cloned to create the remaining three cluster node virtual machines. The high level procedure is:

- In vCenter, clone VMORARAC1.
- On the newly created virtual machine remove the three RDM virtual disks that came from the source virtual machine.
- Attach the same three VMFS disks (VMDKs) for CSR and voting as used by the first virtual machine, VMORARAC1. These three disks were configured to be shared by multiple virtual machines using the procedures outlined in *Disabling simultaneous write protection provided by MVFS using the multiwriter flag* (<http://kb.vmware.com/kb/1034165>).
- Change the hostname and IP address in the guest OS.

6.1 Clone the First Virtual Machine/RAC Node

To clone the first virtual machine for the second, third, and fourth nodes

1. Confirm that the first virtual machine VMORARAC1 is offline.
2. In vSphere Client, select the virtual machine, VMORARAC1.
3. Right-click the name of the virtual machine and select **Clone**.
4. In the **Clone Virtual Machine** wizard on the **Name and Location** page, enter **VMORARAC2** for the **Name**. Click **Next**.
5. On the **Host / Cluster** page, select a host or cluster. Click **Next**.
6. On the **Specific Host** page, select a unique ESX host (that is, one not running a RAC node virtual machine). Click **Next**.
7. On the **Datastore** page, select a VMFS datastore. Click **Next**.
8. On the **Disk Format** page, select **Same format as source**. Click **Next**.
9. On the **Guest Customization** page, select **Do not customize**.
10. On the **Ready to Complete** page, click **Finish**.
11. Repeat steps 1 through 10 for nodes 3 and 4 using **VMORARAC3** and **VMORARAC4** in step 4.

6.2 Change the Hostname and Assign an IP Address

To assign unique names and IP addresses to each RAC Node virtual machine

1. Using the console, log in as root to the VMORARAC2 cloned virtual machine.
2. Go to **System > Administration > Network**.
3. On the **Devices** tab, change the IP addresses for eth0 and eth1, with the required subnet masks, as follows:
 - a. For eth0, change the IP address from 10.1.1.30 to **10.1.1.31**.
 - b. For eth1, change the IP address from 192.168.2.30 to **192.168.2.31**.
4. On the **DNS** tab, change the hostname from VMORARAC1 to **VMORARAC2**.
5. Repeat steps 1 through 4 for the remaining two virtual machines (VMORARAC3 and VMORARAC4), using appropriate IP addresses and subnet masks in steps 3a and 3b, and the corresponding hostname in step 4.

6.3 Configure CSR and Voting Disks

The new virtual machine has been cloned, but the links to hard disks 2, 3, and 4 are incorrect. These are removed so that one disk (VMFS) remains, corresponding to the root drive. The CSR and voting disks are then added. Finally, the SCSI controller type is changed to paravirtual.

The procedures in this section describe the steps needed for one virtual machine. Repeat the steps in this section for each of the four virtual machines named VMORARAC1, VMORARAC2, VMORARAC3, and VMORARAC4.

To remove incorrect links from the clone

1. In vSphere Client, select a virtual machine.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. Select hard disk 2 and click **Remove**. In the **Removal Options** section, select **Remove from virtual machine** and click OK.
4. Repeat step 3 for hard disk 3 and hard disk 4.

To add the first CSR and voting disk for a virtual machine

1. In vSphere Client, select a virtual machine.
2. Right-click the name of a virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. Click **Add** to display the **Add Hardware** wizard.
4. In the **Add Hardware** wizard select **Hard Disk** for the **Device Type**, and click **Next**.
5. In the **Disk** section of the **Select a Disk** page, select **Use an existing virtual disk** and click **Next**.
6. On the **Select Existing Disk** page, enter the path to the first VMDK of the virtual machine (selected in step 1) used for the CSR and voting datastores. Click **Next**.
7. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:0)**. In the **Mode** section, select **Independent** and **Persistent**. Click **Next**.
8. On the **Ready to Complete** page, click **Finish**.
9. Click **OK**.

To add the second CSR and voting disk for a virtual machine

1. Repeat steps 1 through 5 in the preceding procedure to add the second CSR and voting disk.
2. On the **Select Existing Disk** page, enter the path to the second VMDK of the virtual machine (selected in step 1) used for the CSR and voting datastores. Click **Next**.
3. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:1)**. In the **Mode** section, select **Independent** and **Persistent**. Click **Next**.
4. Repeat steps 8 and 9 in the preceding procedure to add the second CSR and voting disk.

To add the third CSR and voting disk for a virtual machine

1. Repeat steps 1 through 5 in the preceding Configure CSR and Voting Disks procedure to add the third CSR and voting disk.
2. On the **Select Existing Disk** page, enter the path to the third VMDK of the virtual machine (selected in step 1) used for the CSR and voting datastores. Click **Next**.
3. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:2)**. In the **Mode** section, select **Independent** and **Persistent**. Click **Next**.
4. Repeat steps 8 and 9 in the preceding procedure to add the third CSR and voting disk.

To set the SCSI controller type to paravirtual

1. In vSphere Client, select the virtual machine that was selected in the preceding step 1.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. In the device list, select **SCSI controller 1**.
4. In the **SCSI Controller Type** section, click **Change Type**.
5. Select **VMware Paravirtual**.
6. Click **OK**, and click **OK** again.

6.4 Format the Added Disks Using fdisk

On node 1, you must partition and format the virtual disks. You can use the Linux `fdisk` utility to accomplish this. Refer to the following illustration showing the output of the `fdisk` utility.

To partition and format three virtual disks

1. Log in as root on node1.
2. Enter the command:
`fdisk /dev/sdb`
3. Enter **n** to create a new partition.
4. Enter **p** to create a primary partition.
5. Enter **1** to create the first partition.
6. Press **Enter** to take the default value of 1 for the first cylinder.
7. Press **Enter** to take the default value of 20480 for the last cylinder.
8. Enter **w** to write the new partition table to disk and exit `fdisk`.
9. Repeat steps 2 through 8, using the following command for the second virtual disk:
`fdisk /dev/sdc`
10. Repeat steps 2 through 8, using the following command for the third virtual disk:
`fdisk /dev/sdd`

The output from the `fdisk` command should look somewhat like the following:

```
root@vmoraracl ~# fdisk /dev/sdb
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF
disklabel
Building a new DOS disklabel. Changes will remain in memory only,
until you decide to write them. After that, of course, the previous
content won't be recoverable.

Warning: invalid flag 0x0000 of partition table 4 will be corrected by
w(rite)

Command (m for help): n
Command action
e extended
p primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-20480, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-20480, default 20480):
Using default value 20480
Command (m for help): w

The partition table has been altered!
Calling ioctl() to re-read partition table.
Syncing disks.
```

7. Install Oracle Grid Infrastructure

The four virtual machines are cloned and ready for the installation of Oracle Grid Infrastructure.

To install Oracle Grid Infrastructure

1. Configure SSH for all the nodes to allow SSH to all nodes without a password.
2. Configure Oracle ASM and create three disks CRSVOL1, CRSVOL2, and CRSVOL3 from /dev/sdb1, /dev/sdc1, and /dev/sdd1 respectively.
3. From VM1 (VMORARAC1) 64-bit (x86_64) installations download and unzip the following software packages (File 1 and File 2) from http://www.oracle.com/technology/software/products/database/oracle11g/112010_linx86_64soft.html.
 - Oracle Database 11g Release 2 Grid Infrastructure (11.2.0.1.0) for Linux x86_64
 - Oracle Database 11g Release 2 (11.2.0.1.0) for Linux x86_64
4. Launch the Oracle Grid Infrastructure installation wizard.
5. On the **Installation Option** page, in the **Select any of the following installation options**, select **Install and Configure Infrastructure for a Cluster**. Click **Next**.
6. On the **Installation Type** page, select **Advanced Installation** and click **Next**.
7. On the **Product Languages** page, select **English** and click **Next**.
8. On the **Grid Plug and Play** page, enter **vmorarac** for the **Cluster Name**, **vmorarac-scan.vmware.com** for the **SCAN Name** and **1521** for the **SCAN Port**. Click **Next**.
9. On the **Cluster Node Information** page, click **Add** to enter a **Hostname** and **Virtual IP Name**. Repeat the process for each of the four hosts, as follows:

Hostname	Virtual IP Name
vmorarac1.vmware.com	vmorarac1-vip.vmware.com
vmorarac2.vmware.com	vmorarac2-vip.vmware.com
vmorarac3.vmware.com	vmorarac3-vip.vmware.com
vmorarac4.vmware.com	vmorarac4-vip.vmware.com

10. Click **Next**.
11. On the **Network Interface Usage** page, enter the following information for the **Interface Name**, **Subnet**, and **Interface Type**:

Interface Name	Subnet	Interface Type
eth0	10.1.1.0	Public
eth1	192.168.2.0	Private

12. Click **Next**.

13. On the **Storage Option** page, select **Automatic Storage Management (ASM)**, and click **Next**.
14. On the **Create ASM Disk Group** page, enter **CRS** for the **Disk Group Name**, and select **Normal** for **Redundancy**.
15. In the **Add Disks** section of the **Create ASM Disk Group** page, select **Candidate Disks**. Make the following entries:

Disk Path	Size (in MB)	Status
ORCL:CRSVOL1	20479	Candidate
ORCL:CRSVOL2	20479	Candidate
ORCL:CRSVOL3	20479	Candidate

16. Select the check box to the left of each of these entries, and click **Next**.
17. On the **ASM Password** page, select **Use same passwords for these accounts** and enter the password **"oracle1"** in the **Specify Password** and **Confirm Password** fields. Click **Next**.
18. On the **Failure Isolation** page, select **Do not use Intelligent Platform Management Interface (IPMI)**. Click **Next**.
19. On the **Operating System Groups** page, make the following entries:

Field	Selection
ASM Database Administrator (OSDBA) Group	asmdba
ASM Instance Administration Operator (OSOPER) Group	asmoper
ASM Instance Administrator (OSASM) Group	asmadmin

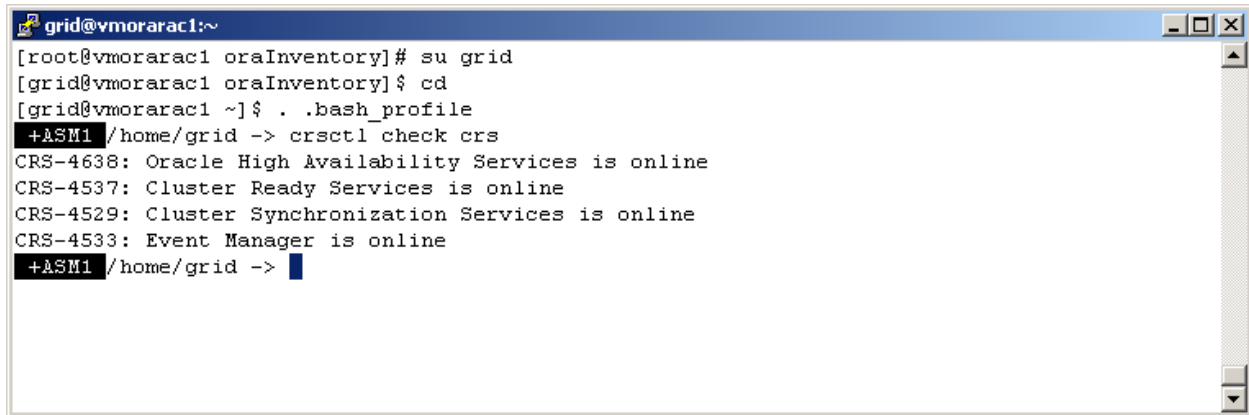
20. Click **Next**.
21. On the **Installation Location** page, enter `/u01/app/grid` for the Oracle Base, and `/u01/app/11.2.0/grid` for the **Software Location**. Click **Next**.
22. On the **Create Inventory** page, enter `/u01/app/orainventory` for the **Inventory Directory**. Click **Next**.
23. Wait while the setup application conducts the prerequisite checks and installs Grid on each node. After setup is complete, execute the following configuration scripts on the four nodes VMORARAC1, VMORARAC2, VMORARAC3, and VMORARAC4. These scripts must be executed in a terminal window as root.


```
/u01/app/orainventory/orinstRoot.sh
/u01/app/11.2.0/grid/root.sh
```
24. Click **OK**, and click **Finish**.

8. Verify the Grid Infrastructure Installation

The following figures show how to run various RAC commands in the guest OS to check the status of the cluster.

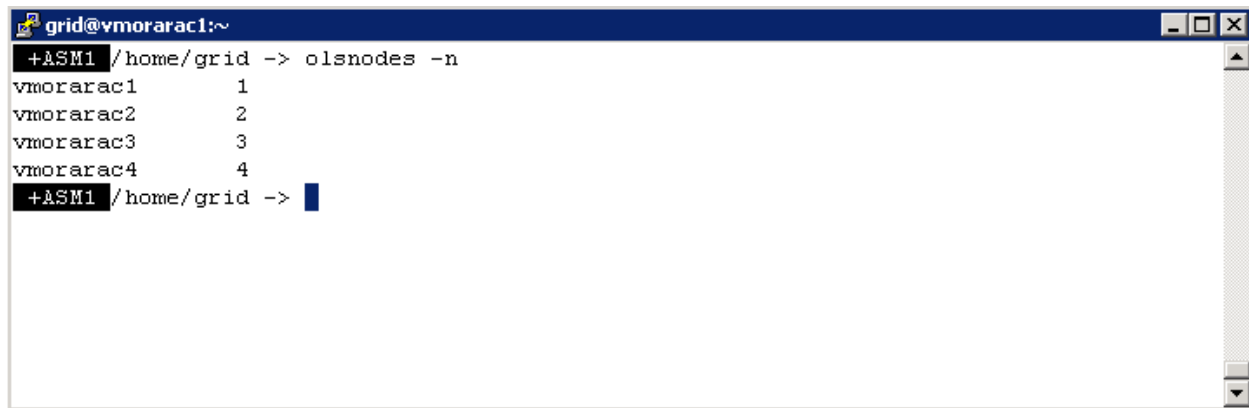
Figure 5. Check the CRS Status on the First Node as Grid User



```

grid@vmorara1:~
[root@vmorara1 oraInventory]# su grid
[grid@vmorara1 oraInventory]$ cd
[grid@vmorara1 ~]$ . .bash_profile
+ASM1 /home/grid -> crsctl check crs
CRS-4638: Oracle High Availability Services is online
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
+ASM1 /home/grid ->
  
```

Figure 6. Check for Cluster Nodes



```

+ASM1 /home/grid -> olsnodes -n
vmorara1      1
vmorara2      2
vmorara3      3
vmorara4      4
+ASM1 /home/grid ->
  
```

Figure 7. Check for Cluster Registry (OCR)

```

grid@vmorara1:~
+ASM1 /home/grid -> ocrcheck
Status of Oracle Cluster Registry is as follows :
  Version          :          3
  Total space (kbytes) :    262120
  Used space (kbytes)  :     2636
  Available space (kbytes) :  259484
  ID                : 1585277307
  Device/File Name   :      +CRS
                    : Device/File integrity check succeeded
                    :
                    : Device/File not configured
                    :
                    : Device/File not configured
                    :
                    : Device/File not configured
                    :
                    : Device/File not configured

Cluster registry integrity check succeeded

Logical corruption check bypassed due to non-privileged user

+ASM1 /home/grid ->
  
```

Figure 8. Check Voting Disk

```

grid@vmorara1:~
+ASM1 /home/grid -> crsctl query css votedisk
## STATE      File Universal Id                File Name Disk group
---  -
 1. ONLINE    91667ba2625c4fefbf7e2c8938b8ebf8 (ORCL:CRSVOL1) [CRS]
 2. ONLINE    29277d62fa324f5dbfad8b987c6b7332 (ORCL:CRSVOL2) [CRS]
 3. ONLINE    490468754f174f0fbff297cd027925e8 (ORCL:CRSVOL3) [CRS]
Located 3 voting disk(s).
+ASM1 /home/grid ->
  
```

9. Add DATA and REDO Disks to Virtual Machine RAC Nodes

In this section, the database disks for data and redo logs are added to the virtual machine RAC nodes in the following way:

- Data and redo datastores are created with the assigned LUNs, and corresponding VMDKs are added to the virtual machine VMORARAC1. The SCSI controller type is set to VMware Paravirtual.
- The same Data and Redo VMDKs are added to the remaining three nodes (VMORARAC2, VMORARAC3, and VMORARAC4) by selecting the VMDKs that were created for VMORARAC1. Again, the SCSI controller type is set to VMware Paravirtual.
- The added disks are formatted using the Linux `fdisk` utility.
- Configuration parameters are added for each virtual machine.

The procedures in this section describe the steps needed for one virtual machine. Repeat the steps in this section for each of the four virtual machines named VMORARAC1, VMORARAC2, VMORARAC3, and VMORARAC4.

To add DATA or REDO disks to the first virtual machine

1. In vSphere Client, select the virtual machine.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. Click **Add** to display the **Add Hardware** wizard.
4. In the **Add Hardware** wizard select **Hard Disk** for the **Device Type**, and click **Next**.
5. In the **Disk** section of the **Select a Disk** page, select **Create a new virtual disk**. Set the disk size to **300GB** for DATA or **64GB** for REDO, and select the **Support clustering features such as Fault Tolerance** check box. Also, for **vSphere 5**, in the **Disk Provisioning** section, select **Thick Provision Eager Zeroed**. Select VMFSDATA01 for DATA, or VMFSREDO01 for REDO to specify a datastore. Click **Next**.
6. On the **Compatibility Mode** page, click **Next**.
7. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:3)** for DATA, or **SCSI (2:0)** for REDO. Set the **Mode** to **Independent**. Click **Next**.
8. On the **Ready to Complete** page, click **Finish**.
9. Click **OK**.
10. Repeat steps 1 through 9 for all the DATA and REDO datastores. Change step 7 to use **SCSI (1:4)** for the next DATA datastore, or **SCSI (2:1)** for the next REDO datastore. For additional DATA or REDO datastores, repeat steps 1 through 9 and continue the numbering sequence in step 7. For example, the third DATA datastore uses the numbering **SCSI (1:5)**, and the third REDO datastore uses the numbering **SCSI (2:2)**.

To set the SCSI controller type to paravirtual

1. Select the virtual machine that was selected in the preceding step 1.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. In the device list, select **SCSI controller 1**.
4. In the **SCSI Bus Sharing** section, select **None**, keep the default selection.
5. In the **SCSI Controller Type** section, click **Change Type**.
6. Select **VMware Paravirtual**.
7. Click **OK**, and click **OK** again.

To format disks using the Linux `fdisk` utility


1. Follow the steps in Section 6.4 to format one disk.
2. Repeat, using `fdisk` to format each disk.

To add configuration parameters for each Oracle RAC virtual machine


Add Configuration parameters for each Oracle RAC virtual machine as per *Disabling simultaneous write protection provided by VMFS using the multi-writer flag* (<http://kb.vmware.com/kb/1034165>).

To add configuration parameters to the virtual machines

1. In vSphere Client, select a virtual machine.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. Click the **Options** tab, click **Advanced**, and click **General**.
4. Click **Configuration Parameters**.
5. On the **Configuration Parameters** page, add rows and entries for the shared disks in the virtual machine's configuration parameters as shown in the following screenshot.
6. Repeat steps 1 through 5 for each of the remaining virtual machines.



The devices for the virtual machine should look like the following screenshot.



10. Create DATA and REDO ASM Disk Groups as Grid User

To prepare for the installation of Oracle Database binary and database

1. Using Oracle ASM, create disks DATAVOL1 through DATAVOL10 using /dev/sde1 through /dev/sdn1, and disks REDOVOL1 through REDOVOL4 using /dev/sdo1 through /dev/sdr1.
2. Using ASM Configuration Assistant logged in as Grid user at node1 (VMORARAC), create two disk groups RACDATA (selecting all ten disks DATAVOL1 through DATAVOL10) and RACREDO (selecting all four disks REDOVOL1 through REDOVOL4).
3. In the ASM Configuration Assistant, on the **Create Disk Group** page, enter **RACDATA** for the **Disk Group Name**. In the **Redundancy** section, select **External (None)**. Click **OK**, and click **OK** again.
4. On the **Create Disk Group** page, in the **Select Member Disks** section, select ORCL:DATAVOL1, ORCL:DATAVOL2, ORCL:DATAVOL3, and ORCL:DATAVOL4. Set the size of each to 300GB. Click **OK**.
5. On the **Create Disk Group** page, enter **RACREDO** for the **Disk Group Name**. In the **Redundancy** section, select **External (None)**. Click **OK**, and click **OK** again.
6. On the **Create Disk Group** page, in the **Select Member Disks** section, select ORCL:REDOVOL1, ORCL:REDOVOL2, ORCL:REDOVOL3, and ORCL:REDOVOL4. Set the size of each to 64GB. Click **OK**.

11. Install and Create the Oracle Database 11g R2 RAC

To install and create the Oracle Database 11g R2 RAC

1. Log in into VMORARAC1 as oracle user and launch `runInstaller` from the Oracle software location to install the Oracle binaries.
2. In the Oracle Installer wizard, on the **Configure Security Updates** page, leave the **Email** field blank, and do not select **I wish to receive security updates via My Oracle Support**. Click **Next**.
3. On the **Installation Option** page, select **Install database software only**. Click **Next**.
4. On the **Grid Options** page, select **Real Application Clusters database installation**. Select the four nodes VMORARAC1, VMORARAC2, VMORARAC3, and VMORARAC4. Click **Next**.
5. On the **Product Languages** page, select **English**. Click **Next**.
6. On the **Database Edition** page, select **Enterprise Edition (4.29GB)**. Click **Next**.
7. On the **Installation Location** page, enter `/u01/app/oracle` for **Oracle Base**, and `/u01/app/oracle/product/11.2.0/dbhome_1` for **Software Location**. Click **Next**.
8. On the **Operating System Groups** page, select **dba** for the **Database Administrator (OSDBA) Group** and **oinstall** for the **Database Operator (OSOPER) Group**. Click **Next**.
9. Review the configuration on the **Summary** page, and click **Finish**.
10. After the installation is complete, open a terminal window, log in as root, and run the following script:

```
/u01/app/oracle/product/11.2.0/dbhome_1/root.sh
```

Repeat for each of the four nodes, and click **OK**.
11. On the **Finish** page, click **Close**.
12. Create the Oracle RAC database using Oracle Database Configuration Assistant (DBCA).

12. References

The following are resources and references for Oracle and VMware vSphere.

vSphere Installation and Setup

<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-installation-setup-guide.pdf>

vSphere Networking

<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-networking-guide.pdf>

vSphere Storage

<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-storage-guide.pdf>

Guide to configure NTP on ESX servers

<http://kb.vmware.com/kb/1003063>

Timekeeping in VMware Virtual Machines

<http://www.vmware.com/files/pdf/Timekeeping-In-VirtualMachines.pdf>

Installing VMware Tools in a Linux virtual machine using Red Hat Package Manager (RPM)

<http://kb.vmware.com/kb/1018392>

Oracle ASMLib

<http://www.oracle.com/technetwork/topics/linux/asmlib/index-101839.html>

Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide 10g Release 2 (10.2) for Linux

http://download.oracle.com/docs/cd/B19306_01/install.102/b14203/toc.htm

Oracle Clusterware Installation Guide 11g Release 1 (11.1) for Linux

http://www.oracle.com/pls/db111/to_toc?pathname=install.111/b28263/toc.htm

Oracle Real Application Clusters Installation Guide 11g Release 1 (11.1) for Linux and UNIX

http://www.oracle.com/pls/db111/to_toc?pathname=install.111/b28264/toc.htm

Appendix A: Deployment of Oracle RAC with RDM


This appendix describes the following:

- Deployment steps for Four-Node Oracle RAC with RDM.
- Virtual disk layout for RDMs.
- Adding CRS, DATA, and REDO RDMs to an Oracle RAC virtual machine.

Deployment Steps for Oracle RAC on vSphere with RDM

The following process diagram illustrates the deployment steps for deploying a four-node Oracle RAC on vSphere 5.0 with RDM.

Figure 9. Deployment Steps for a Four-Node Oracle RAC on vSphere with RDM



Virtual Disk Layout

Table 1. Virtual Disk Layout

Virtual Disk on ESX	Guest OS Device Name	Virtual Device	Virtual SCSI Driver	Size (GB)	Purpose
VMDK – Hard disk 1	/dev/sda	SCSI 0:0	LSI Logic	50	Oracle Enterprise Linux 5.5 OS and Oracle binaries
RDM – Hard disk 2	Shared disk /dev/sdb1	SCSI 1:0	Paravirtual	20	CRS and voting disk
RDM – Hard disk 3	Shared disk /dev/sdc1	SCSI 1:1	Paravirtual	20	CRS and voting disk
RDM – Hard disk 4	Shared disk /dev/sdd1	SCSI 1:2	Paravirtual	20	CRS and voting disk
RDM – Hard disk 5	Shared disk /dev/sde1	SCSI 1:3	Paravirtual	300	RAC database DATA
RDM – Hard disk 6	Shared disk /dev/sdf1	SCSI 1:4	Paravirtual	300	RAC database DATA
RDM – Hard disk 7	Shared disk /dev/sdg1	SCSI 1:5	Paravirtual	300	RAC database DATA
RDM – Hard disk 8	Shared disk /dev/sdh1	SCSI 1:6	Paravirtual	300	RAC database DATA
RDM – Hard disk 9	Shared disk /dev/sdi1	SCSI 1:8	Paravirtual	300	RAC database DATA
RDM – Hard disk 10	Shared disk /dev/sdj1	SCSI 1:9	Paravirtual	300	RAC database DATA
RDM – Hard disk 11	Shared disk /dev/sdk1	SCSI 1:10	Paravirtual	300	RAC database DATA
RDM – Hard disk 12	Shared disk /dev/sdl1	SCSI 1:11	Paravirtual	300	RAC database DATA
RDM – Hard disk 13	Shared disk /dev/sdm1	SCSI 1:12	Paravirtual	300	RAC database DATA
RDM – Hard disk 14	Shared disk /dev/sdn1	SCSI 1:13	Paravirtual	300	RAC database DATA

Virtual Disk on ESX	Guest OS Device Name	Virtual Device	Virtual SCSI Driver	Size (GB)	Purpose
RDM – Hard disk 15	Shared disk /dev/sdo1	SCSI 2:0	Paravirtual	64	RAC database REDO
RDM – Hard disk 16	Shared disk /dev/sdp1	SCSI 2:1	Paravirtual	64	RAC database REDO
RDM – Hard disk 17	Shared disk /dev/sdq1	SCSI 2:2	Paravirtual	64	RAC database REDO
RDM – Hard disk 18	Shared disk /dev/sdr1	SCSI 2:3	Paravirtual	64	RAC database REDO

12.1

Create a Virtual Machine

The section describes creation of the first virtual machine (that is, the first RAC node) using vSphere Client. This is created on VMFS and corresponds to the root drive. Two NICs are assigned for the public and private networks. Three RDM LUNs are assigned to the virtual machine for CRS and voting. Only one LUN is required but three are used for redundancy.

Four ESX hosts have been created and are visible in vCenter.

To create a new virtual machine

1. Log in to vCenter using vSphere Client.
2. Select a host and click **Create a new virtual machine** on the **Getting Started** tab.
3. In the **Create New Virtual Machine** wizard, select **Typical** for the **Configuration**, and click **Next**.
4. On the **Name and Location** page, enter **VMORARAC1** for the **Name**, and click **Next**.
5. On the **Datastore** page, select **VMDATASTORE**, and click **Next**.
6. On the **Guest Operating System** page, select **Linux** from the **Guest Operating System** list. For the **Version**, select **Oracle Linux 4/5 (64-bit)** from the drop-down menu. Click **Next**.
7. On the **Create a Disk** page, set the **Virtual Disk Size** to **50GB**, and click **Next**.
8. On the **Ready to Complete** page, select the **Edit the virtual machine settings before completion** check box. Click **Continue**.
9. On the **Virtual Machine Properties** page for the VMORARAC1 virtual machine, select **8 CPUs** and **128GB Memory**.
10. Click **Finish**.

Add CRS and Voting Disk

Three 20GB LUNs are attached to the virtual machine for RAC CRS and voting. Although only one is required, three are used for redundancy. The following characteristics apply to these LUNs:

- The LUNs are assigned as RDMs.
- The LUNs are attached to a SCSI controller configured as VMware Paravirtual.

- For the SCSI controller the SCSI BUS sharing policy is set to physical to allow these LUNs to be shared by multiple virtual machines on any ESX host.

The procedures in this section describe the steps needed to add RDMs for one virtual machine. Repeat the steps in this section for each of the four virtual machines named VMORARAC1, VMORARAC2, VMORARAC3, and VMORARAC4.

To add one 20GB RDM for a virtual machine

1. In vSphere Client, select a virtual machine.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. Click **Add** to display the **Add Hardware** wizard.
4. In the **Add Hardware** wizard select **Hard Disk** for the **Device Type**, and click **Next**.
5. In the **Disk** section of the **Select a Disk** page, select **Raw Device Mappings**. Set the LUN size to **20GB**. Click **Next**.
6. On the **Select Datastore** page, in the **Select datastore on which to store LUN mapping** section, select **Store with virtual machine**. Click **Next**.
7. On the **Compatibility Mode** page, in the **Compatibility** section, select **Physical**. Click **Next**.
8. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:0)**. Click **Next**.
9. On the **Ready to Complete** page, click **Finish**.
10. Click **OK**.

To add a second 20GB RDM for the virtual machine

1. Select the virtual machine that was selected in the preceding step 1.
2. Repeat steps 2 through 7 in the preceding procedure to add one 20GB datastore.
3. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:1)**. Click **Next**.
4. Repeat steps 9 and 10 of the preceding procedure to add one 20GB datastore.

To add a third 20GB RDM for the virtual machine

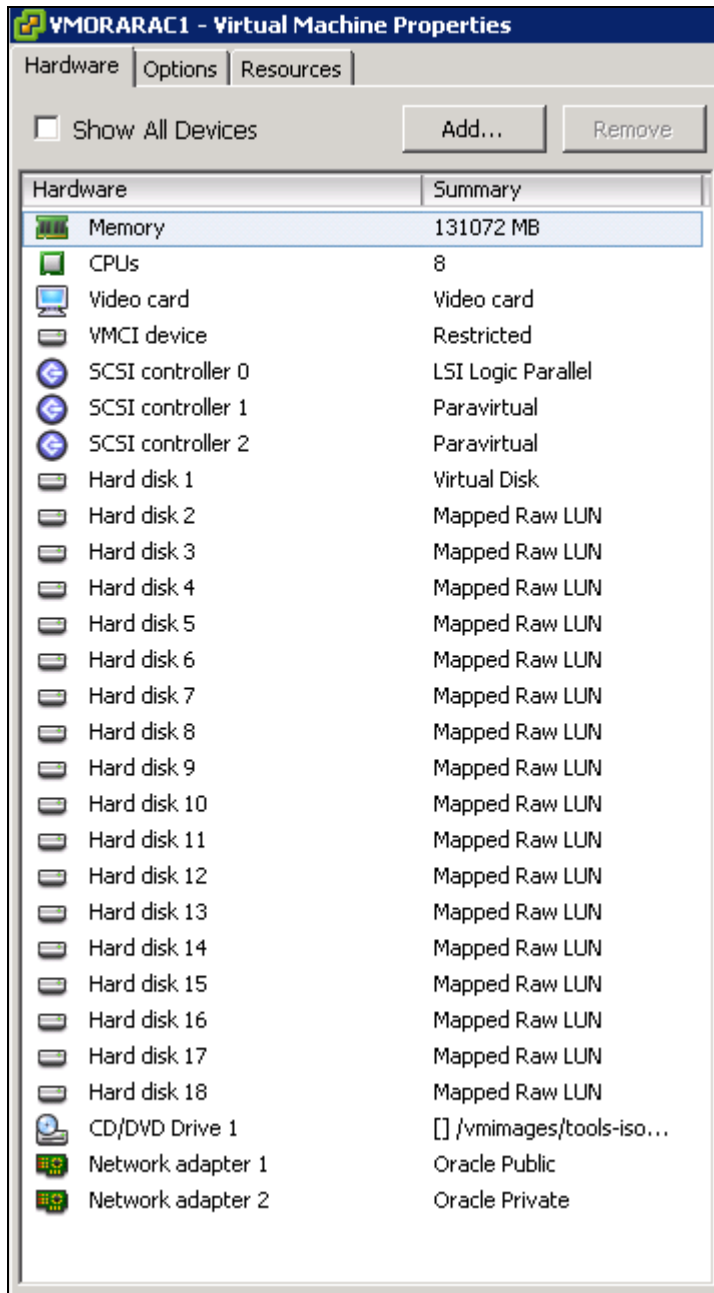
1. Select the virtual machine that was selected in the preceding step 1.
2. Repeat steps 2 through 7 in the preceding procedure to add one 20GB datastore.
3. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:2)**. Click **Next**.
4. Repeat steps 9 and 10 of the preceding procedure to add one 20GB datastore.

Change the properties of the SCSI controller for the RDMs so they can be shared and configured as paravirtual.

To set the SCSI controller type to paravirtual

1. Select the virtual machine that was selected in the preceding step 1.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. In the device list, select **SCSI controller 1**.
4. In the **SCSI Bus Sharing** section, select **Physical**.
5. In the **SCSI Controller Type** section, click **Change Type**.
6. Select **VMware Paravirtual**.
7. Click **OK**, and click **OK** again.

The virtual machine properties can be viewed in vCenter, as seen in the following screenshot.



Add DATA and REDO Disks to Virtual Machine RAC Nodes

In this section, the database disks for data and redo logs are added to the virtual machine RAC nodes. The general procedure is:

- Data and redo LUNs are added to virtual machine VMORARAC1 as RDMs. This creates RDM mapping files in the VMFS datastore. The SCSI controller type is set to VMware Paravirtual.
- The same data and redo LUNs are added to the remaining three nodes (VMORARAC2, VMORARAC2, and VMORARAC3) by selecting the same RDM mappings created in the previous step. Again, the SCSI controller type is set to VMware Paravirtual.

To add DATA or REDO LUNs to the first virtual machine, VMORARAC1

1. In vSphere Client, select the virtual machine, VMORARAC1.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. Click **Add** to display the **Add Hardware** wizard.
4. In the **Add Hardware** wizard select **Raw Device Mappings** for the **Device Type**, and click **Next**.
5. In the **Select a target LUN** section of the **Select a Disk** page, select a **300GB** LUN for DATA or a **64GB** LUN for REDO. Click **Next**.
6. On the **Select Datastore** page, in the **Select datastore on which to store LUN mapping** section, select **Store with virtual machine**. Click **Next**.
7. On the **Compatibility Mode** page, in the **Compatibility** section, select **Physical**. Click **Next**.
8. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:3)** for DATA, or **SCSI (2:0)** for REDO. Click **Next**.
9. On the **Ready to Complete** page, click **Finish**.
10. Click **OK**.
11. Repeat steps 1 through 7 for all the DATA and REDO LUNs. Change step 8 to use **SCSI (1:4)** for the next DATA LUN, or **SCSI (2:1)** for the next REDO LUN. For additional DATA or REDO LUNs, repeat steps 1 through 7 and continue the numbering sequence in step 8. For example, the third DATA LUN uses the numbering **SCSI (1:5)**, and the third REDO LUN uses the numbering **SCSI (2:2)**.

To set the SCSI controller type for the RDM LUNs to paravirtual

1. Select the virtual machine VMORARAC1.
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. In the device list, select **SCSI controller 1**.
4. In the **SCSI Controller Type** section, click **Change Type**.
5. Select **VMware Paravirtual**.
6. Click **OK**, and click **OK** again.

To add DATA or REDO LUNs to the remaining virtual machines

1. In vSphere Client, select a virtual machine (VMORARAC2, VMORARAC3, or VMORARAC4).
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. Click **Add** to display the **Add Hardware** wizard.

4. In the **Add Hardware** wizard select **Hard Disk** for the **Device Type**, and click **Next**.
5. In the **Select the type of disk to use** section of the **Select a Disk** page, select **Use an existing virtual disk**. Click **Next**.
6. On the **Select Existing Disk** page, in the **Disk File Path** section, select the path to the vmdk pointer file for the LUN. Click **Next**.
7. On the **Compatibility Mode** page, click **Next**.
8. On the **Advanced Options** page, in the **Virtual Device Node** section, select **SCSI (1:3)** for DATA, or **SCSI (2:0)** for REDO. Set the **Mode** to **Independent and Persistent**. Click **Next**.
9. On the **Ready to Complete** page, click **Finish**.
10. Click **OK**.
11. Repeat steps 1 through 7 for all the DATA and REDO LUNs. Change step 8 to use **SCSI (1:4)** for the next DATA LUN, or **SCSI (2:1)** for the next REDO LUN. For additional DATA or REDO LUNs, repeat steps 1 through 7 and continue the numbering sequence in step 8. For example, the third DATA LUN uses the numbering **SCSI (1:5)**, and the third REDO LUN uses the numbering **SCSI (2:2)**.
12. Repeat steps 1 through 11 for each of the other virtual machines (VMORARAC2, VMORARAC3, or VMORARAC4).

To set the SCSI controller type for the RDM LUNs to paravirtual

1. Select a virtual machine (VMORARAC2, VMORARAC3, or VMORARAC4).
2. Right-click the name of the virtual machine and select **Edit Settings** to display the **Virtual Machine Properties** pane.
3. In the device list, select **SCSI controller 1**.
4. In the **SCSI Controller Type** section, click **Change Type**.
5. Select **VMware Paravirtual**.
6. Click **OK**, and click **OK** again.

To format disks using the Linux `fdisk` utility

1. Follow the steps in Section 6.4 to format one disk.
2. Repeat, using `fdisk` to format each disk.

The devices for the virtual machine should look like the following screenshot.

