

Efficient Deployment of Oracle Real Application Clusters (RAC) in a Virtualized Environment Using Enterprise Manager Provisioning Pack

A Dell Technical White Paper

Authors:

Kai Yu
Dell Database Solutions Engineering, Dell Product Group

Akanksha Sheoran Kaler
Enterprise Manager Product Management, Oracle Corporation

Contributing Author:

Rajat Nigam
Enterprise Manager Product Management, Oracle Corporation

January 2010



Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2010 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *Dell* logo, *OpenManage*, *PowerConnect*, *EqualLogic* and *PowerEdge* are trademarks of Dell Inc; *Linux* is the registered trademark of Linus Torvalds in the United States and other countries. *Intel* and *Xeon* are registered trademarks of Intel Corporation in the U.S. and other countries; Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

Contents

Introduction to Oracle Enterprise Manager and Oracle VM	3
Overview of Oracle Enterprise Manager Provisioning and Patch Automation Pack	4
Benefits of Using Provisioning and Patching Features:	6
Overview of VM Technology and Virtualization Management Pack	6
Oracle VM Management Pack Features:	7
Lifecycle Automation for Guest Virtual Machines	7
Adding Oracle VM Technology as the Latest Evolution of Grid Computing	8
Provisioning and Cloning of Guest Virtual Machines Using VM Management Pack	10
Oracle Virtualization Infrastructure Implementation	11
Dell Blade Server as VM Server Platform	11
EqualLogic Storage Provides the Shared Storage for VM Servers	12
Oracle VM Server Installation	13
Oracle VM server manual installation tasks	13
Oracle VM Server Network Configuration	13
Oracle VM Server Storage Configuration on EqualLogic iSCSI Storage	18
Configuring the dom0 Host Access to the EqualLogic Storage Volumes	18
Convert the OVM Repository from Local Storage to the Shared Storage	19
Connect the VM Servers to Enterprise Manager Grid Control	21
Provisioning Guest Virtual Machines Using Oracle Enterprise Manager on Virtual Server	25
Oracle VM Templates	25
Figure 14: Import VM template	28
Create Guest VMs Using a VM Template	28
Customizing Oracle VM Templates to Create “Gold Images”	33
Configure Guest VMs for Oracle 11g RAC Database	37
Deploying Real Application Cluster using Enterprise Manager Provisioning and Patch Automation Pack	37
Pre-Requisite for RAC Provisioning	38
RAC Provisioning using Deployment procedures	40
VM template for RAC	43

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

Extending RAC..... 43

Conclusion..... 46

Appendix 47

References: 48

Introduction to Oracle Enterprise Manager and Oracle VM

Oracle Enterprise Manager provides the most extensive, comprehensive, deep, and best-of-breed management for Oracle products. It manages the entire IT stack—from packaged applications (EBS, Siebel, PSFT, etc.) to middleware and to the databases in both physical and virtual environments.

Oracle Enterprise Manager has deep knowledge of Oracle technologies, which allows it to provide intelligent and unique solutions designed to maximize quality of service and minimize cost and complexity.

While Oracle Enterprise Manager's core strength is management of the Oracle ecosystem, it is not a closed solution. A rich set of plug-ins allows for monitoring and management of non-Oracle products as well including MS SQL Server, IBM Websphere, EMC Symmetrix, F5 BIG-IP, etc.

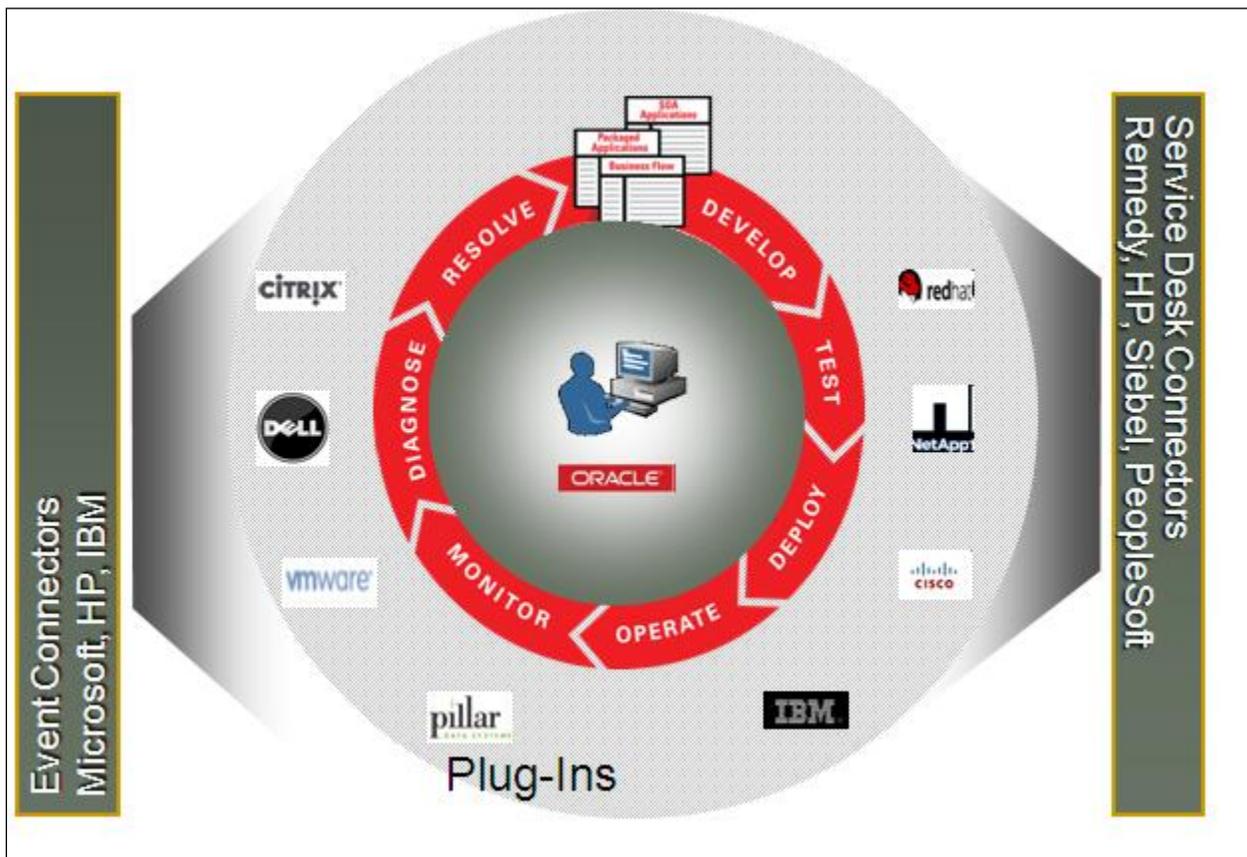


Figure 1: Oracle Enterprise Manager Overview

Oracle Enterprise Manager also features a rich set of cross-tier management solutions spanning application, middleware, database and infrastructure layers. This enables a unique "top-down"

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

management approach which allows Information Technology (IT) departments to focus on what matters to businesses-- greater agility, better service quality, and lower operational costs. Using Oracle Enterprise Manager, you can manage your applications from a top-down perspective-- from monitoring service levels to proactively isolating business exceptions before they escalate, and remediating issues at any level of the IT stack.

Its breadth of capabilities as a management solution is another hallmark of Oracle Enterprise Manager as it provides a complete and integrated lifecycle management solution, including testing, provisioning, patching, monitoring, day-to-day administration, performance, and configuration management.

This paper illustrates the best practices for provisioning an Oracle 11g Real Application Clusters (RAC) database in a virtualized environment using Enterprise Manager VM Management Pack and Enterprise Manager Provisioning Pack. Starting with the introduction to Enterprise Manager, the paper covers the best practices for Oracle virtualization infrastructure implementation including the storage and network configuration. Additionally, the entire process is described on how to provision Oracle Virtual servers, Virtual machines, Oracle 11g RAC on the guest virtual machines, and adding nodes to the RAC using the Enterprise Manager Provisioning Pack.

Overview of Oracle Enterprise Manager Provisioning and Patch Automation Pack

Oracle Enterprise Manager Provisioning and Patch Automation Pack provides an end-to-end solution for automated provisioning, patching, and lifecycle management for entire system stacks including physical and virtual infrastructure, servers, databases, middleware, and applications.

The provisioning features allow deployment of new servers and Oracle software using either standard gold images which can be created and stored in the Enterprise Manager Software Library, or a Live Reference Installation. These features allow mass deployment of software and servers in a standard, repeatable, and unattended manner. This mechanism aids implementing a skill-agnostic process for deployment.

The gold image-based deployment approach ensures that all deployments conform to standards and that there are no resource integrity threats while performing mass deployment of software. This approach, complemented by industry certified ITIL release management processes, allow you to create and manage multiple versions of gold images for Oracle software right from virtual servers, operating systems, databases, middleware, and applications.

The reference-based provisioning deployment approach allows user to provision databases from an existing pre-patched running database, while ensuring that the provisioned database is up to required standards.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

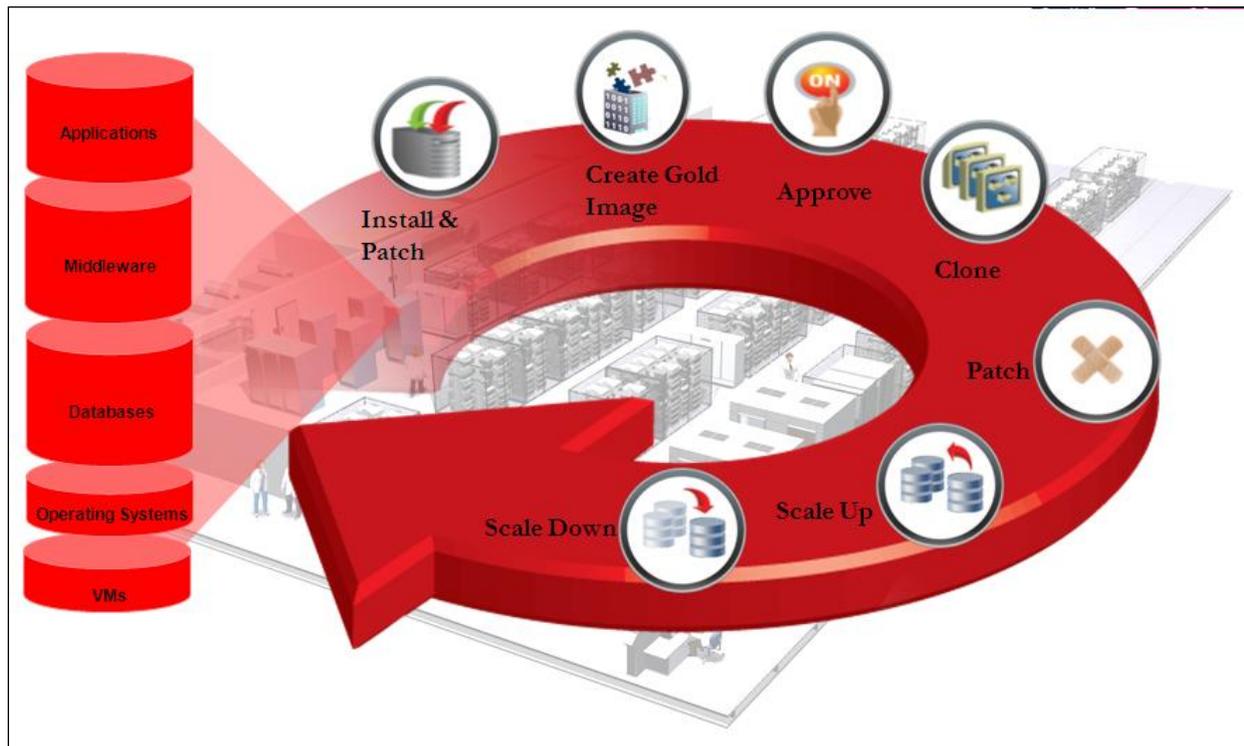


Figure 2: Software lifecycle management using Provisioning and Patch Automation Pack

Figure 2 illustrates how Oracle Enterprise Manager Grid Control orchestrates both the initial reference sandbox deployment, and then the mass unattended deployment of gold images from reference deployments.

Besides deployment of software, Enterprise Manager Grid Control also automates the on-going patch lifecycle management of these deployments-proactively informs you of critical patches and vulnerabilities in deployments, automatically downloads required patches, applies them to deployments, and verifies whether patches were applied successfully.

Oracle Enterprise Manager Grid Control offers lifecycle management solutions in the form of Deployment Procedures that help to manage the entire lifecycle of software, applications, and servers. Deployment Procedures, not only allow for automated deployment of software, applications, and servers, but also the maintenance of these deployments. This makes critical data center operations easy, efficient, and scalable resulting in lower operational risk and cost of ownership.

Provisioning and patching solutions are offered across the stack for:

- operating systems, with Bare Metal Provisioning of Linux and operating system patching

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

- databases, with both single instances and Real Application Clusters (RAC) provisioning, extension, deletion and flexible patching for Oracle database and Oracle Real Application Clusters
- middleware, with application server J2EE, BPEL, SOA provisioning, and patching

Benefits of Using Provisioning and Patching Features:

- Provides a repeatable, reliable, and automated solution for performing mass, unattended, and schedulable deployment of
 - Software and servers based on gold images created using reference deployment or installation media
 - Software and operating system updates
 - Complex and multi-tier software like Oracle Real Application Clusters (RAC) and Application Server Clusters
- Handles not only provisioning of software, but completely automate configuration of software and ensure zero-time for patching of mission critical systems by managing rolling patching for complex multi-tier installations like Real Application Clusters (RAC) databases and Application Server clusters.
- New resources are provisioned at short notice based on compliant and tested gold images.
- Accommodate multiple operations in a single change window.
- Provides command-line interface support to all out-of-box provisioning and patching deployment procedures which can then be invoked by custom scripts.
- Software provisioning and patching features support SUDO and PAM authentication.
- Single interface for multiple players. For example, component designers responsible for creating gold images based on corporate standards and the operators all use the same Oracle Enterprise Manager Grid Control console.
- Automation of repeatable installation and patching operations across the stack leads to substantial cost savings in terms of costs and man hours.

Overview of VM Technology and Virtualization Management Pack

Oracle VM Management Pack provides a comprehensive management solution for end-to-end monitoring, configuration management, and lifecycle automation of virtualized infrastructure to capture and maximize the benefits of virtualization. The Oracle VM Management Pack helps customers accelerate the adoption of virtualization, enabling them to optimize IT resources, improve hardware utilization, streamline IT processes, and reduce costs without adding the complexity and cost associated with multiple management tools.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

Oracle VM Management Pack Features:

- Manage virtual environments with an application perspective: Allows customers quickly to diagnose whether the root cause of a problem is in an application component, in a virtual resource, or in a physical resource, and enables efficient consolidation and optimization of data center resources through virtualization.
- Built-in configuration management: Enables administrators easily to track relationships between virtual infrastructure components and software to ensure quick problem resolution and configuration compliance.
- Lifecycle automation: Allows automation spanning test, deployment, and patching and maintenance capabilities, including provisioning of Oracle VM servers and guest virtual machines as well as live migration of guest virtual machines to other servers during server maintenance windows.
- Faster server and software rollout: Automates software deployment through Oracle VM templates for packaged applications, middleware, database, and Oracle Enterprise Linux.
- High availability for virtual infrastructure: Provides enterprise-class capabilities such as server pooling, automatic load balancing, and server failover, which helps to ensure business continuity in the virtual environment traditionally associated with high-cost physical servers.

Lifecycle Automation for Guest Virtual Machines

Oracle VM Management Pack enables Oracle Enterprise Manager to automate the guest virtual machine lifecycle. This reduces the time required to provision new operating system/application environments, helps standardize application deployments by cloning pre-certified configurations, and minimizes planned downtime by allowing guest virtual machines to be live-migrated to other servers during server maintenance windows. Oracle Enterprise Manager also allows automated patching of operating systems and Oracle software running inside the guest virtual machines.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

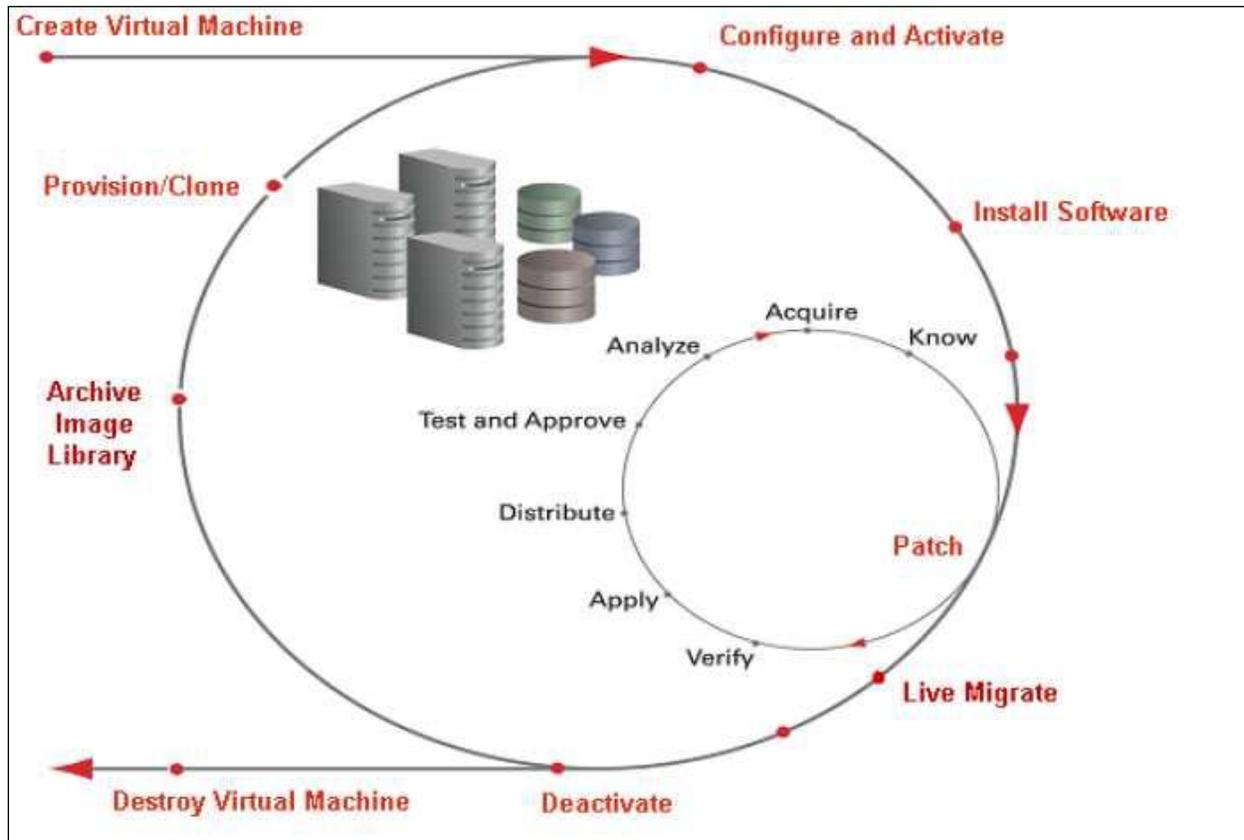


Figure 3: Oracle VM Management Pack facilitates the automation of the guest virtual machine lifecycle

Adding Oracle VM Technology as the Latest Evolution of Grid Computing

Grid computing enables groups of systems to be pooled and provisioned to accommodate the changing needs of businesses. Rather than using dedicated servers for specific tasks, each application may make use of shared computing resources. This approach results in better scalability and high availability, enabling different applications to consolidate infrastructure, save on cost, and enhance power efficiency. Introducing virtualization to the Grid infrastructure environment allows for greater flexibility, enabling consolidation beyond the database grid by adding virtualization to the Grid infrastructure and databases. This allows multiple application servers or database servers to run together on one physical server, expanding the database grid to the virtual servers, and providing an ideal location for test and development for the Oracle RAC infrastructure.

The Virtual Grid discussed here is based on Oracle VM Technology and is composed of the virtual machines (VM) that are enabled by Oracle VM. The Oracle VM environment includes the following components:

- Oracle VM Server: A virtualization environment that provides a lightweight, server-based platform for running virtual machines. It is based on an updated version of the underlying Xen Hypervisor technology. The Oracle VM software is installed on a bare metal x86 server a Dell

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

blade server. It includes a Linux kernel running as dom0 with support to manage one or more domU virtual machines. Dom0 is the management domain that handles the physical device's IO, networking, and etc. Dom0 also runs the VM agents that connect to the VM manager.

- Virtual Server Pool: An autonomous region of VM servers that collects all of the resources of VM servers. All the VM servers with one virtual server pool need to access the shared storage.
- VM (Virtual Machine): Guest Operating System with applications running on domU.
- Oracle VM Manager: provides a user interface that can be used to manage Oracle VM servers, virtual machines, and resources. VM manager can be configured independently, or used as a VM management pack plug-in as a part of Oracle Enterprise Manager Grid Control. For this POC project, we used Oracle Enterprise Manager Grid Control with the VM Management Pack plug-in as the unified management solution for the physical grid as well as the virtual grid. Figure 4 illustrates the Oracle VM server components and Oracle virtual

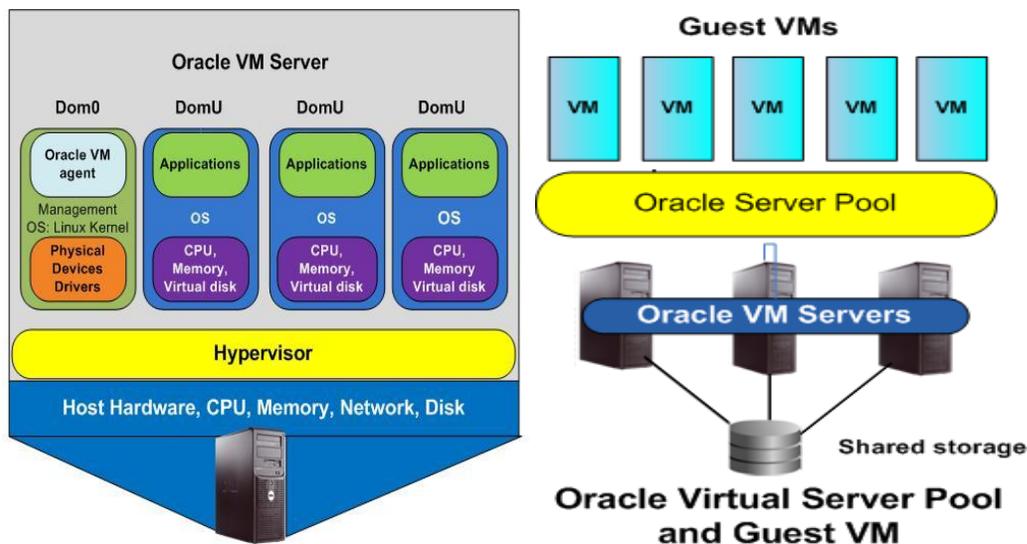


Figure 4: Oracle VM environment components and VM server pool

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

- Virtual Grid Components:
 - Virtual Servers: Dell blade servers
 - Shared Storage: Dell EqualLogic iSCSI SAN
 - Unified Management: Oracle Enterprise Manager Grid Control with VM Management

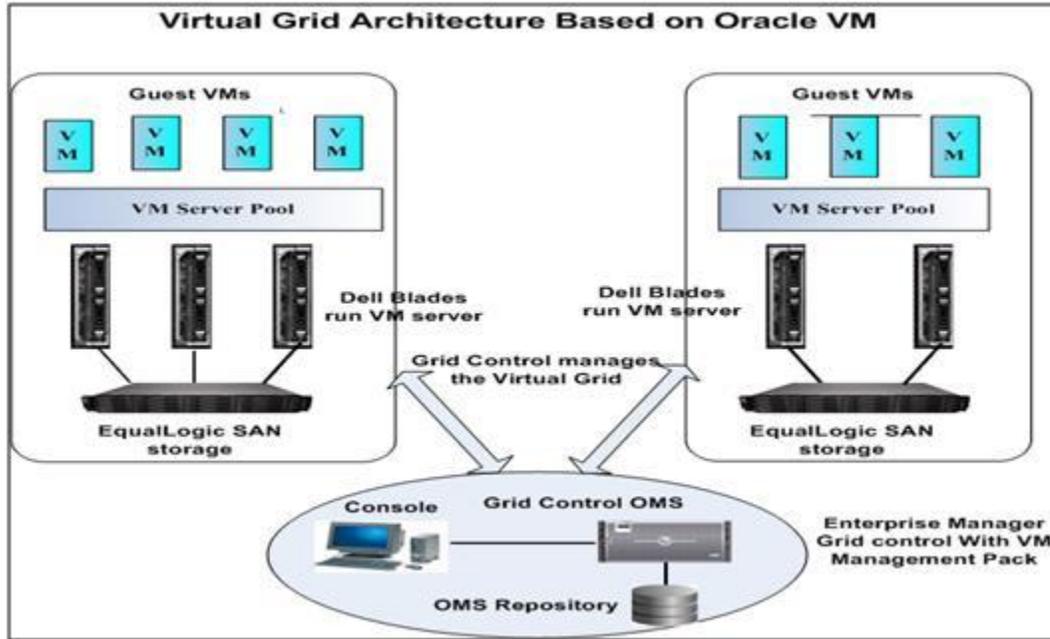


Figure 5: VM grid architecture and its management solution

Provisioning and Cloning of Guest Virtual Machines Using VM Management Pack

Oracle Enterprise Manager eliminates the challenge and complexity of provisioning and cloning new guest virtual machines. Oracle Enterprise Manager allows administrators to save an existing guest virtual machine as an Oracle VM template, and then use this template to create new guest virtual machines. Alternatively, users can point to an existing guest virtual machine to create a software gold image.

This standardized gold image can be cloned to provision software on other machines. For example, gold images can be created for database and middleware components, as well as Linux environments. Organizations can utilize these cloning features to standardize software deployments throughout a data center.

Oracle Virtualization Infrastructure Implementation

Dell Blade Server as VM Server Platform

With the ever growing requirements of IT datacenters, space is always a valuable commodity. Dell PowerEdge blade servers offer a dense solution to optimize space constraints. Dell's M1000e blade enclosure offers the ability to enclose 16 blades in a 10U-sized enclosure. To take advantage of these benefits while implementing VM infrastructure, Dell used the blade server as the VM server platform. Figure 6 shows the Dell M1000e blade chassis with the M610 blade servers .

Dell EqualLogic iSCSI SAN as the shared storage

- PS6000 XV

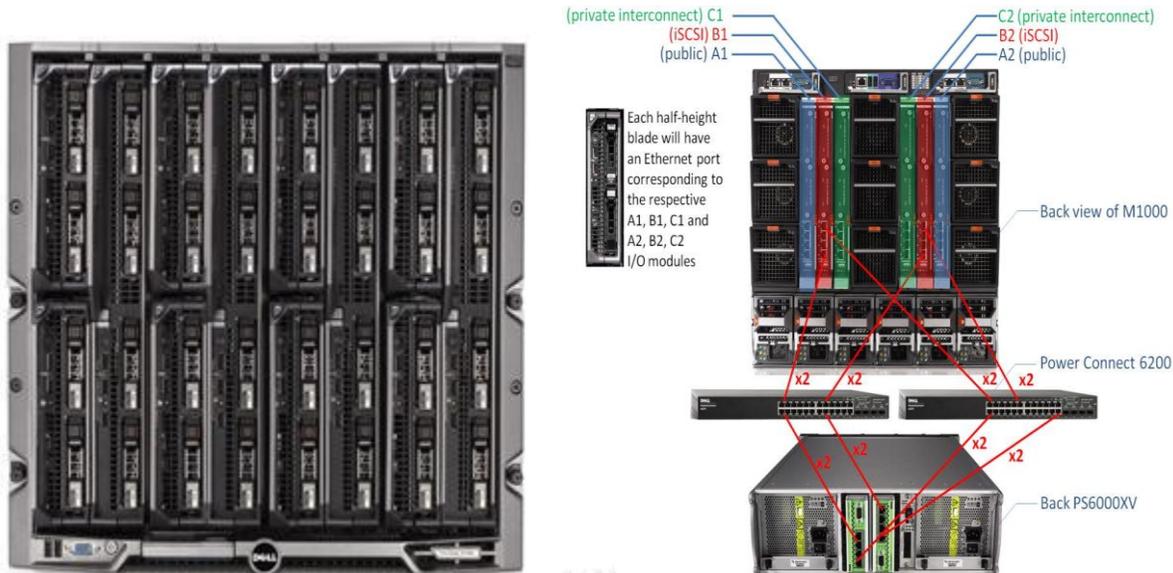


Figure 6: Dell M610 blade servers as the server platform for Oracle VM, EqualLogic storage provides the shared storage for VM servers

Each Dell M610 server runs the Oracle VM servers to contribute resources including memory, virtual CPUs, network devices and disk devices to the VM server pool. A guest virtual machine has the operating system running on a virtual server with a set of these configurable resources. As a best practice from Dell, the total number of virtual CPUs on a virtual server should be less than or equal to two times the total number of CPU cores that exist with the VM Server. In our example reference design, each OVM server has 8 CPU cores, limiting us to a total of 16 virtual CPUs for the total number of guest VMs running on each VM server. The total amount of virtual memory available to all of the Virtual Machines running on a physical server should be less than or equal to the total amount of RAM installed. Each server in our reference design includes 24GB of memory, so the aggregate memory of our guests running on an OVM sever should not be larger than 24GB.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

Application software may be configured and run on the virtual machines as with physical machines. A guest virtual machine is associated with a VM server pool from which the resources are assigned to the guest virtual machine. A VM server pool contains one or more VM servers running on physical servers and, possibly, shared storage. Shared storage is compulsory unless the virtual pool contains a single VM server. With shared storage configured in the VM server pool, a guest virtual machine associated with the VM server pool may be started and run on any VM server within the pool that is available and has the most resources free. With high availability option (HA) being enabled on the server pool level and the guest virtual machine level, the virtual machine is migrated or restarted on another available VM server if the VM server that runs the guest virtual machine is shutdown or fails.

To complement the VM Manager, the native management solution for Oracle VM environment, Oracle has released the Oracle Enterprise Manager 10g R5 with Oracle VM Management Pack which provides a complementary management solution for the virtual server infrastructure.

This solution includes the management and provisioning of the components of the virtual Grid such as virtual server pool, VM servers, guest VMs, and the resources allocations to the guest VMs such as CPUs, disks, networks, and memory. Figure 5 shows the virtual grid infrastructure managed by the Enterprise Manager Grid control 10g R5 with the Oracle VM Management Pack. Refer to the My Oracle Support Note 781879.1 for more information on enabling the VM Management Pack.

EqualLogic Storage Provides the Shared Storage for VM Servers

As shown in Figure 6, the Blade servers connect to Dell EqualLogic storage PS6000 XV through a Gigabit Ethernet IP Storage Area Network (SAN). The Ethernet configuration includes the two Blade switches on the IO modules, B1 and B2, and two external Dell PowerConnect 6200 switches. This fully redundant configuration provides high availability and scalability and maximum IO bandwidth for the storage IOs.

One or more EqualLogic storage arrays, called members, can be organized into a storage group which is assigned a DNS name or IP address. Each storage group can be segregated into pools and, at any given time, a member may be assigned to one pool. The member may be moved between pools while keeping the storage online. For this implementation, two members, oracle-member01 and oracle-member02, are assigned to pool Raid10 with the Raid10 disk array configuration. In order to provide the usage storage for Grid, physical disks are grouped by a usable component called volume. A volume represents a portion of the storage pool that can be spread across multiple disks and group members and is seen as an iSCSI target visible to storage clients. To the client OS, it is a logic disk.

Oracle VM Server Installation

Oracle VM Server can be installed using following methods:

1. Oracle Enterprise Manager for Bare Metal Provisioning of Oracle VM servers

Oracle Enterprise Manager uses industry standard PXE boot technology to automate provisioning of Oracle VM Server software on Bare Metal machines. For more information, refer to [Oracle Enterprise Manager Administrator's Guide for Software and Server Provisioning and Patching](#)

2. Manual installation of Oracle VM Server software

Download the Oracle VM Software from: <http://www.oracle.com/virtualization>. Burn the Oracle VM Server ISO file to a bootable CDROM. Refer to the [VM Server Installation Guide](#) for more details.

The recommended method for installing the VM server is the first option, “Oracle Enterprise Manager for Bare Metal Provisioning of Oracle VM servers,” but for the purpose of this exercise we will use the second option, “Manual Installation of Oracle VM Server software”.

Oracle VM server manual installation tasks

- Enable virtualization for CPUs on BIOS
- Install Oracle VM Server 2.1.2
- Configure dom0 memory by editing /boot/grub/menu.lst
- Secure OVM agent is running: `service ovs-agent status`

Oracle VM Server Network Configuration

With the updated version of the underlying Xen hypervisor technology, the Oracle VM server also includes a Linux kernel that runs as dom0 to manage one or more domU guest virtual machines. Oracle VM uses Xen bridge to provide the network for guest VMs running on domU. This bridge functions as a virtual switch presented to the guest VMs. When the VM server is installed in a Dell blade server M610, six network interfaces are shown in the dom0 by default. Each interface is associated with a Xen bridge that can be presented to guest VM in domU. The default configuration can be modified to meet your Grid configuration needs. Figure 7 shows the customization of the network configuration used in this proof of concept project.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

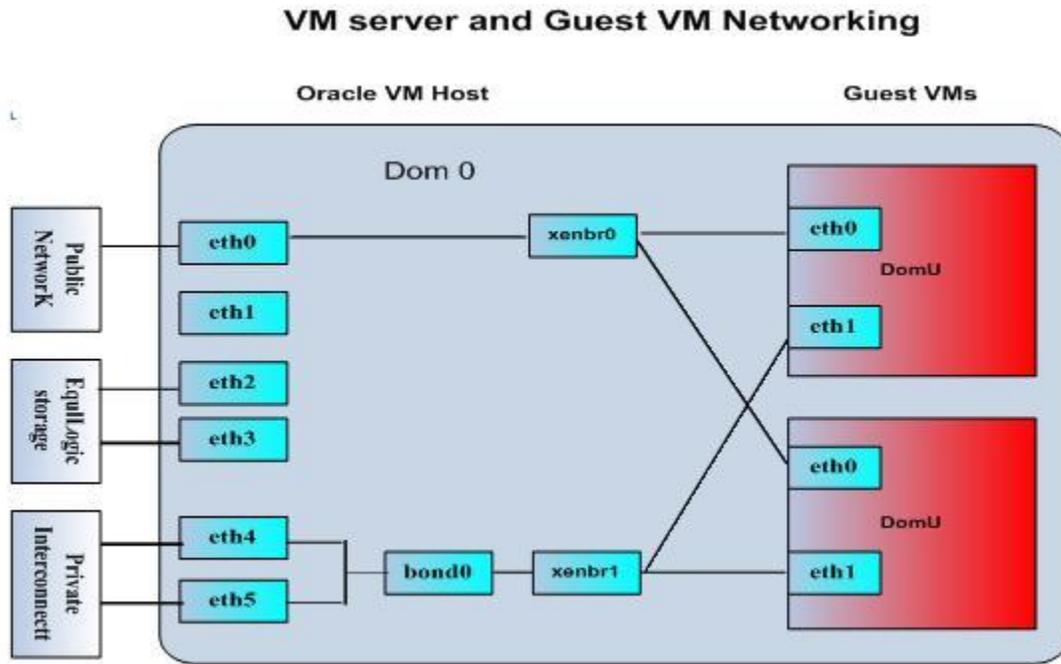


Figure 7: VM servers and guest VMs networking configuration

In dom0, six of the network interfaces play similar roles to those in the servers of the physical grid. The difference is that some of these network interfaces need to be presented to the guest VMs through the Xen bridges. Figure 7 shows the network configuration for the VM servers and how they are presented to the guest VMs.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

- eth0 is for a public network connecting to the IO module A1 (eth1 connecting with A2 can be an optional second public network interface).
- eth2 and eth3 connecting to IO modules B1 and B2 are for the iSCSI connections. As all the guest VM IO operations are handled by dom0, there is no need to expose eth2 and eth3 to guest VMs, and therefore no Xen bridge is needed for eth2 and eth3.
- eth4 and eth5 are bonded as bond0 for the private interconnect. This bonded private interconnection is also used for Oracle Cluster File System 2 (OCFS2) heartbeat traffic between VM servers in the VM server pool. This OCFS2 cluster file system provides storage for the VM repositories. On other hand, bond0 is also presented to the guest VMs as a virtual network interface eth2 through Xen bridge Xenbr1. This virtual network interface eth2 can be used for the private interconnection between guest VMs which is required for deploying Oracle 11g RAC database on the guest VMs.

Table 1 illustrates the example network configuration in VM server dom0 and how these network interfaces are bonded, bridged, and presented to the guest VM as the virtual network interfaces.

Network Interface in Dom0	IO modules	Connections	IP address	Xen Bridge	Dom#1	Dom#2	Dom#N
eth0	A1	Public network	155.16.9.82-85	Xenbr0	eth0: 155.16.99.101	eth0: 155.16.99.102	eth0: 155.16.99.10N
eth2	B1	iSCSI connection	10.16.7.228-234(even)				
eth3	B2	iSCSI connection	10.16.7.229-235(odd)				
eth4	C1	eth4 and eth5 bonded to form bond0	192.168.9.82-85	Xenbr1	eth1:192.168.9.101	eth1:192.168.9.102	eth1: 155.16.99.10N
eth5	C2						

Table 1: Network configuration in Oracle VM Server and guest VMs

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

Following are the steps used to implement these network configurations by customizing the default Xen bridge configuration and network configurations from the VM server installation:

1. Shut down the default Xen bridge configuration:

```
/etc/xen/scripts/network-bridges stop
```

2. Make sure no Xen bridge exists:

```
[root@kblade9 scripts]# brctl show
bridge name      bridge id          STP enabled      interfaces
```

3. Modify `/etc/xen/xend-config.sxp`
change the line:

```
(network-script network-bridges)
```

To:

```
(network-script network-bridges-dummy)
```

4. Edit `/etc/xen/scripts/network-bridges-dummy` to include only the following two lines:

```
#!/bin/sh
/bin/true
```

5. Configure the network interfaces and bonding and Xen bridges `xenbr0` and `xenbr1` by editing the following network scripts in `/etc/sysconfig/network-scripts`:

```
ifcfg-eth0:
DEVICE=eth0
HWADDR=00:1D:09:FC:B8:10
ONBOOT=yes
BOOTPROTO=none
BRIDGE=xenbr0
```

```
ifcfg-eth2:
DEVICE=eth2
BOOTPROTO=none
HWADDR=00:1D:09:72:29:F4
ONBOOT=yes
IPADDR=10.16.7.229
NETMASK=255.255.255.0
USERCTL=no
MTU="9000"
```

```
ifcfg-eth3:
DEVICE=eth3
BOOTPROTO=none
HWADDR=00:1D:09:72:29:F6
ONBOOT=yes
IPADDR=10.16.7.230
```

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

```
NETMASK=255.255.255.0  
USERCTL=no  
MTU="9000"
```

```
ifcfg-eth4:  
DEVICE=eth4  
BOOTPROTO=none  
HWADDR=00:1D:09:FC:B8:12  
TYPE=Ethernet  
MASTER=bond0  
SLAVE=yes  
USERCTL=no  
ONBOOT=yes
```

```
ifcfg-eth5:  
DEVICE=eth5  
ONBOOT=yes  
BOOTPROTO=none  
HWADDR=00:10:18:3b:37:e0  
TYPE=Ethernet  
MASTER=bond0  
SLAVE=yes  
USERCTL=no
```

```
ifcfg-bond0:  
DEVICE=bond0  
ONBOOT=yes  
BOOTPROTO=none  
BRIDGE=xenbr1
```

```
ifcfg-xenbr0:  
DEVICE=xenbr0  
ONBOOT=yes  
TYPE=Bridge  
DELAY=0  
STP=off  
BOOTPROTO=none  
IPADDR=155.16.9.82  
NETMASK=255.255.0.0
```

```
ifcfg-xenbr1  
DEVICE=xenbr1  
ONBOOT=yes  
TYPE=Bridge  
DELAY=0  
STP=off  
BOOTPROTO=none  
IPADDR=192.168.9.82  
NETMASK=255.255.255.0
```

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

- Restart the network service:

```
[root@kblade9 scripts]# Service network restart
```

- Check the Xen bridge configuration:

```
[root@kblade9 scripts]# brctl show
bridge name      bridge id                STP enabled        interfaces
xenbr0           8000.002219d1ded0       no                  eth0
xenbr1           8000.002219d1ded2       no                  bond0
```

Oracle VM Server Storage Configuration on EqualLogic iSCSI Storage

All Oracle VM resources, including guest VMs, Virtual Machine templates (guest seed images), ISO images, shared/non-shared Virtual Disks, etc. are stored in Oracle VM repositories. The Oracle VM server installation creates a single non-shared, OCFS2 repository (/OVS) on the local storage. As this default configuration does not work with a VM server pool with the multiple VM servers required to operate on the shared resources in the repository on which the virtual Grid is based, the repositories infrastructure of the VM servers in the virtual Grid need to be moved to the shared storage. EqualLogic PS6000XV provides the shared storage for this repository infrastructure.

The following storage volumes are created for the virtual Grid:

- blade-data4 (400GB) and blade-data6(500GB) for OVM repositories
Four 1GB volumes vmorc1, vmorc2, vmorc3, vmorc4 are created as the disk partitions that will be presented to the guest VMs as the virtual disks for OCR and votingdisks of the 11g RAC that are running on the guest VMs.
- vmrac1 and vmrac2 are the partitions that will be used as the virtual disks for ASM diskgroups of the 11g RAC database running on guest VMs.

Configuring the dom0 Host Access to the EqualLogic Storage Volumes

To access these EqualLogic volumes from dom0, eth2 and eth3 two network interfaces, shown in Table 1, are configured in dom0 using the open-iSCSI administration utility in the same way as they were for a physical server.

Secondly, the device mapper multipath is configured for all the iSCSI volumes in the same way as it was in a physical server. Refer to the corresponding steps in the physical Grid for details.

The following devices should be created and available in dom0:

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

```
[root@kblade10 ~]# ls /dev/mapper/*
/dev/mapper/blade-data6      /dev/mapper/mpath5
/dev/mapper/ovs_data4p1
/dev/mapper/blade-data6p1    /dev/mapper/vmocr-css1
/dev/mapper/vmocr-css1p1
/dev/mapper/vmocr-css1p2    /dev/mapper/vmocr-css1p3 /dev/mapper/vmocr-
css2
/dev/mapper/vmocr-css3      /dev/mapper/vmocr-css4  /dev/mapper/vmracdb
/dev/mapper/vmracdb1      /dev/mapper/vmracdb1p1 /dev/mapper/ovs_data4
```

Convert the OVM Repository from Local Storage to the Shared Storage

Configure OCFS2 file system for VM repositories:

1. Create `/etc/ocfs2/cluster.conf` file:

```
node:
    ip_port = 7777
    ip_address = 192.168.9.82
    number = 0
    name = kblade9
    cluster = ocfs2

node:
    ip_port = 7777
    ip_address = 192.168.9.83
    number = 1
    name = kblade10
    cluster = ocfs2

node:
    ip_port = 7777
    ip_address = 192.168.9.85
    number = 2
    name = kblade11
    cluster = ocfs2

cluster:
    node_count = 3
    name = ocfs2
```

- a. Stop ocfs2 service: `service o2cb stop.`
- b. Configure oc2b service:

```
[root@kblade9 ocfs2]# service o2cb configure
```

- c. 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 "2 boot." The current values will be shown in brackets ('1').

Hitting ENTER without typing an answer will keep that current value.

Ctrl-C will abort the procedure.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

- Load O2CB driver on boot (y/n) [n]: y
 - Cluster to start on boot (Enter "none" to clear) [ocfs2]:
 - Specify heartbeat dead threshold (>=7) [31]: 100
 - Specify network idle timeout in ms (>=5000) [30000]: 60000
 - Specify network keepalive delay in ms (>=1000) [2000]: 1000
 - Specify network reconnect delay in ms (>=2000) [2000]: 2000
 - Writing O2CB configuration: OK
 - Mounting configfs filesystem at /sys/kernel/config: OK
 - Loading module "ocfs2_dlmfs": OK
 - Creating directory '/dlm': OK
 - Mounting ocfs2_dlmfs filesystem at /dlm: OK
- d. Make ocfs2 file system on /dev/mapper/ovs_data4p1
- ```
[root@kblade9 /]# mkfs.ocfs2 -b 4k -C 64k -L ovs /dev/mapper/ovs_data4p1
```
- e. Unmount the existing /OVS on local disk:
- ```
umount /OVS
```
- f. Change /etc/fstab to have the shared volume mounted at boot:
- ```
#/dev/sda3 /OVS
ocfs2 defaults 1 0
/dev/mapper/ovs_data4p1 /OVS
ocfs2 _netdev,datavolume,nointr 0 0
```
- g. Mount ocfs2 partitions: mount -a -t ocfs2
- h. Repeat the steps above on all the VM servers in the VM server pool.
2. Add a new volume to the repositories:
- To increase the size of OVM the repositories, you can add more additional disk partitions to the VM repositories:
- a. Make OCFS2 file system on the new disk partition /dev/mapper/blade-data6p1
- ```
mkfs.ocfs2 -b 4k -C 64k -L ovs /dev/mapper/blade-data6p1
```
- b. Add the new partition to the repositories:
- ```
/usr/lib/ovs/ovs-makerepo /dev/mapper/blade-data6p1 1
OVS_repository
```
- c. Check the new repositories:
- ```
df -k | grep /OVS
/dev/mapper/ovs_data4p1 1 419432960 327838144 91594816 79%
/OVS
/dev/mapper/blade-data6p1 524297280 1071936 523225344 1%
/OVS/9A87460A7EDE43EE92201B8B7989DBA6
```

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

The 500GB new repository `/OVS/9A87460A7EDE43EE92201B8B7989DBA6` is added using the new partition `/dev/mapper/blade-data6p1`

Connect the VM Servers to Enterprise Manager Grid Control

1. Install the latest version of Oracle Enterprise Manager Grid Control and enable the Virtual Management Pack.

The following steps are performed to connect VM servers to the VM Management Pack of the Oracle Enterprise Manager Grid Control:

- a. Create Oracle user with group name `oinstall` in the VM server `dom0`.
- b. Create `/OVS/proxy` directory and set the ownership `chown oracle.oinstall /OVS/proxy`
- c. Oracle Enterprise Manager supports only Oracle VM server version 2.1.2 or above. All Oracle VM servers version 2.1.2 or above managed by Oracle Enterprise Manager should be updated to have OVS agent of version 2.2-70 or higher.
- d. Set the sudo privilege for Oracle user in the VM server

```
Vi sudo -f /etc/sudoers
Add line: oracle ALL=(ALL) NOPASSWD: ALL

Comment out line: Defaults    requiretty
```

2. Create a Virtual Server Pool

It is essential that you create a server pool to manage all Oracle VM servers. A server pool must have a minimum of one virtual server registered with it.

- a. Login to Oracle Enterprise Manager Grid Control console as `sysman`
- b. In Grid Control, click **Targets**, then click **Virtual Servers**
- c. In the Virtualization Central page, select **Action as Virtual Server Pool Create Virtual Server Pool** and click **Go**.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack



Figure 8: Start creating virtual server pool

- d. Specify the parameters for creating the virtual pool including the first VM server host name, Enterprise manager agent that is used to communicate with the VM server, Oracle user passwords on the VM server, and the Grid control host, etc.

Note: To understand the values in this page, refer to Table 2 in the Appendix.

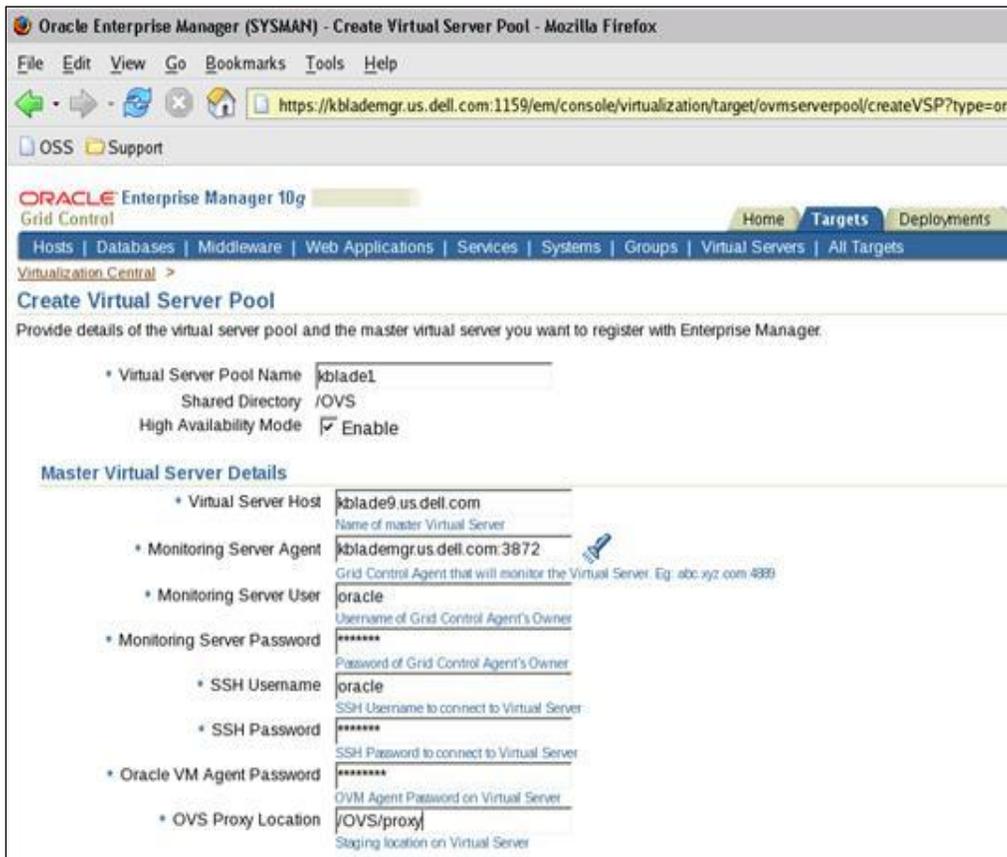


Figure 9: Specify the first VM server kblade9 on the Virtual server pool kblade1

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

VM server pool kblade1 successfully created with the first VM sever kblade9 registered as:

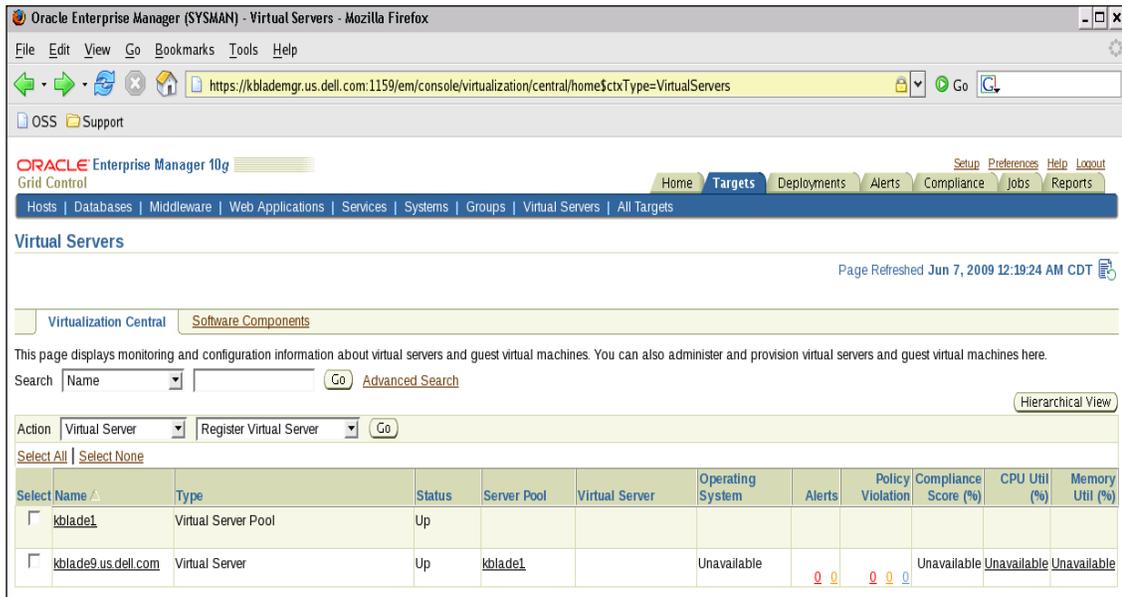


Figure 10: Virtual server pool kblade1 with the first VM server kblade9 is created

3. Registering Virtual Servers

For Enterprise Manager Grid Control to monitor and manage a virtual server, you must register it with a server pool. Once you create a master virtual server, you can also register other virtual servers to the server pool.

- a. Add additional virtual servers by repeating Step 1 on kblade10 and register it to the VM server pool kblade1:

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

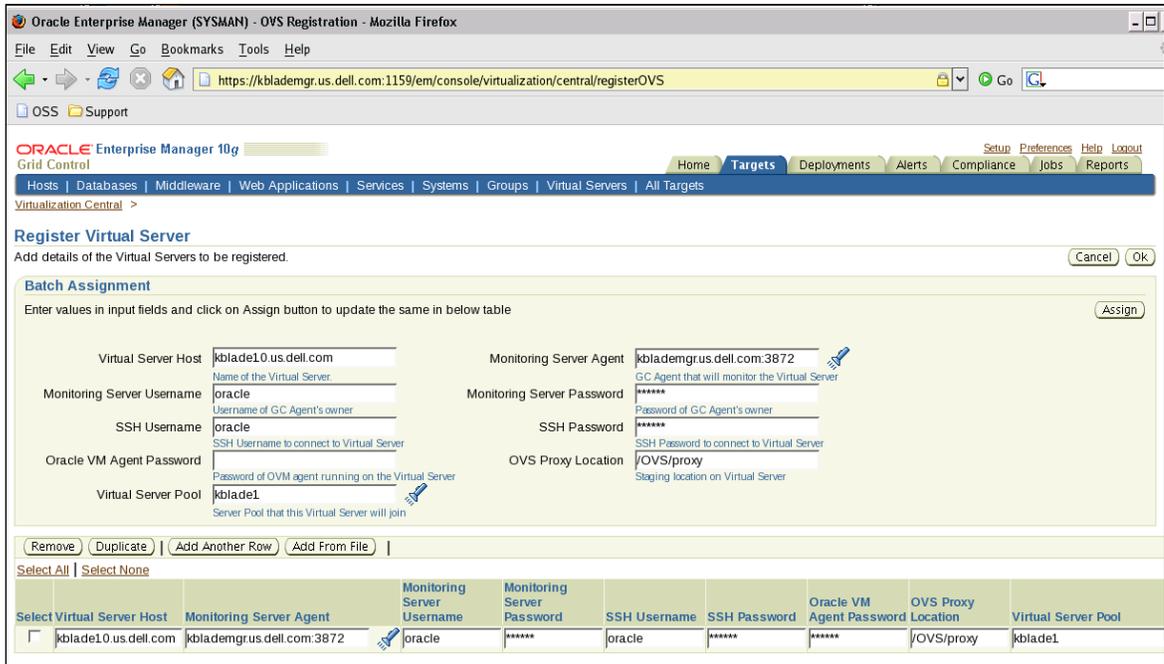


Figure 11: Register the second VM server kblade10 to the VM server pool kblade1

VM server Kblade10 is added to VM server pool kblade1 successfully :

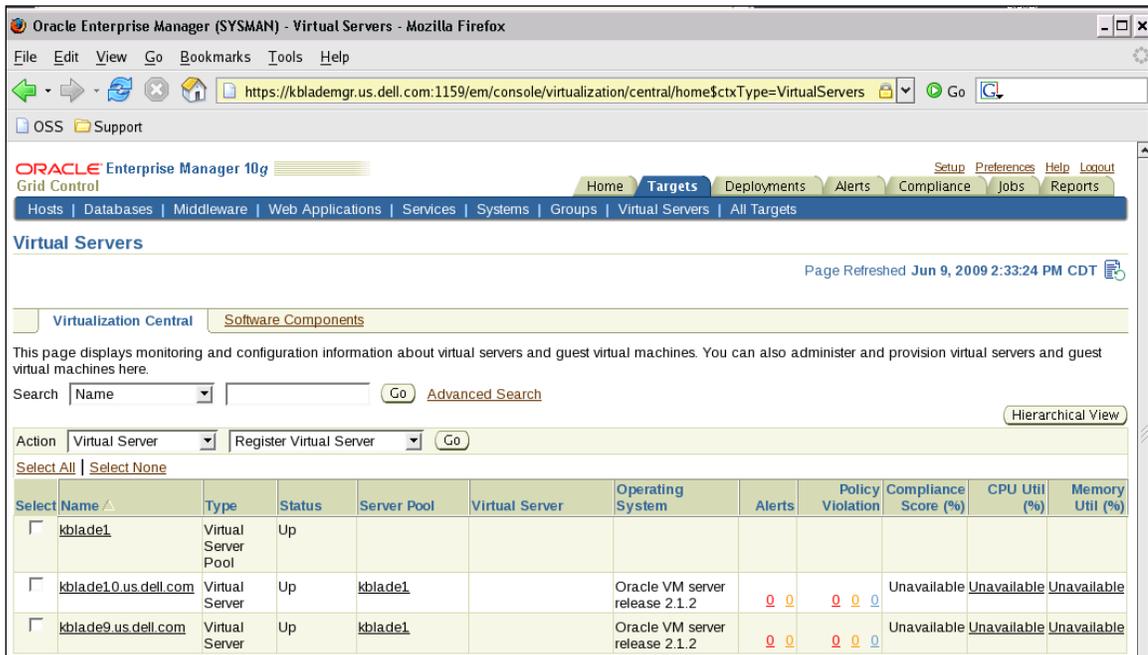


Figure 12: VM server pool kblade is composed two VM servers kblade9 and kblade10

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

- b. Repeat step 3-a to add the third VM server kblade11 to VM server pool kblade1 in the same way.

Provisioning Guest Virtual Machines Using Oracle Enterprise Manager on Virtual Server

The VM server pool and VM servers associated with the VM server pool provide the resources for the virtual Grid. These resources include CPUs, memory, disk storage, network resources and VM templates. As the basic element of the virtual Grid, guest VMs can be created and run on one of the virtual servers with a specific number of resources such as virtual CPUs, memory and disk space. The guest VM will contain its own operating system and associated application software.

Depending on the type of virtualization method used, a virtual machine can be full virtualized (also known as hardware virtualized (HVM)) or paravirtualized. In a full virtualized configuration (hardware virtualized), the unmodified operating system runs on the virtual machine, while in a paravirtualized configuration, the guest operating system is recompiled. The paravirtualized virtual machine runs on near native speed. In the proof of concept discussed below, only paravirtualized VMs are created to form the virtual Grid.

Guest virtual machines may be created using any of three sources:

- Oracle VM Templates
- ISO images of Linux and Windows operating systems
- Bare Metal provisioning of Operating System on a network bootable (PXE boot) virtual machine

In this paper we will be using Oracle VM Templates to create the guest VM.

Oracle VM Templates

An Oracle VM Template is a VM, or group of VMs, containing Oracle or other software that is pre-built, pre-installed, pre-configured, and ready to use—no installation required. Simply download the Templates from Oracle E-Delivery, import to an Oracle VM instance, deploy, and start-up the VM(s) to begin using the encapsulated product(s).

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

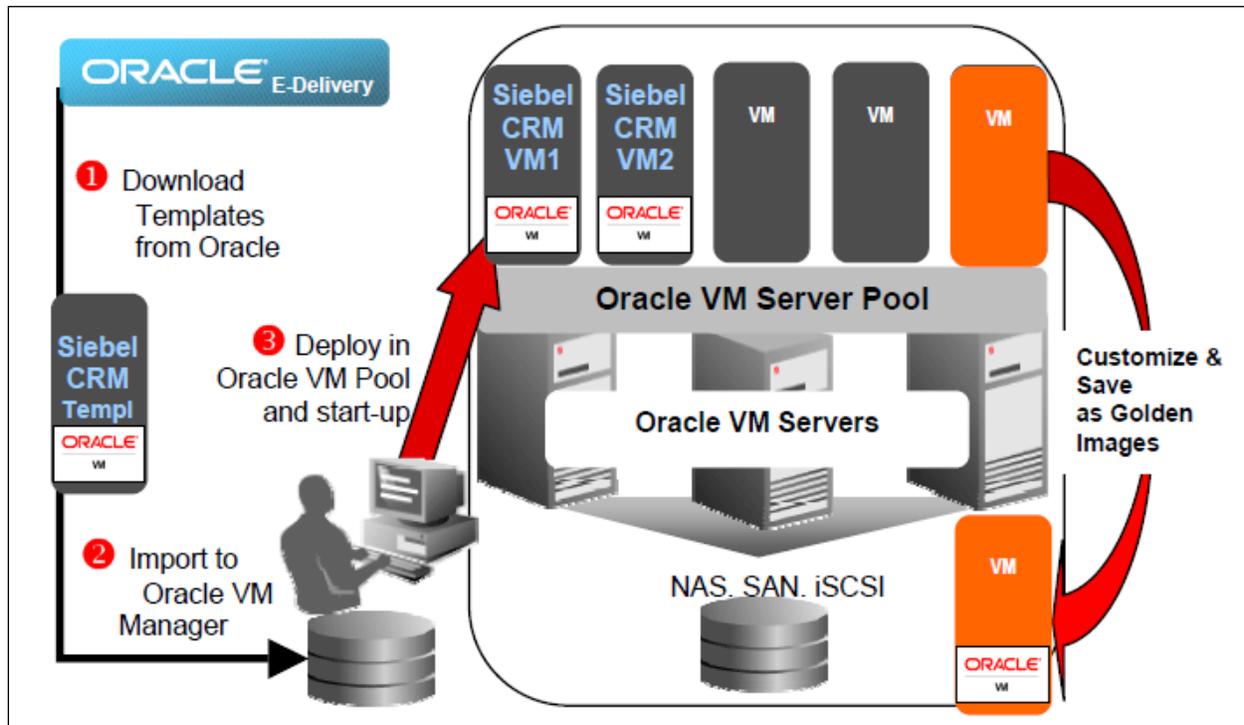


Figure 13: Oracle VM template creation process

Oracle VM Templates produced by Oracle are available from Oracle E-Delivery, ready to download and start-up in an Oracle VM infrastructure. Already licensed for production use, Oracle VM Templates can save users days or weeks learning to install and configure. Instead, users can focus on downloading, starting the VM(s), and beginning to use the product right away with the confidence that they are deploying an exact, validated copy of their VM(s) and not a VM hand-built from scratch where the risk of human error such as missing a patch dependency, or incorrectly configuring a parameter is far greater.

Within these templates, Oracle software is laid out in the same manner as the software would be if it were installed and patched manually “the old fashioned way”, meaning the exact same directories and Oracle “homes” are used, and the package and patch inventories are completely standard and up-to-date so that no changes to your normal Oracle procedures are required to maintain the instances over time.

Downloading Oracle VM Templates

Oracle VM Templates are downloaded from the Oracle E-Delivery site

<http://edelivery.oracle.com/oraclevm>.

Ideally, Oracle VM Templates should be downloaded directly to the shared `/OVS/seed_pool` directory used by every Oracle VM Manager instance. This directory serves as the central repository directory for all Oracle VM Template files associated with an Oracle VM. Of course, alternatively, they could be

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

downloaded anywhere, but ultimately they will need to be moved to the `/OVS/seed_pool` directory to be imported, so a direct download to this directory will save a step.

The following steps were followed for downloading the Oracle VM templates:

1. Download the Oracle Linux OEL5.2 template zip file `OVM_EL5U2_X86_64_11GRAC_PVM.gz` from the Oracle VM template download website: <http://edelivery.oracle.com/oraclevm>
2. Unzip this file to `/OVS/seed_pool` to form the `OVM_EL5U2_X86_64_PVM_10GB` directory which includes three files: `README`, `System.img`, `vm.cfg`

The next step is to import these Templates into Oracle VM Manager so they can be used to create VMs.

Importing Templates

For the Templates to be used, each of the Templates associated with the included VMs must be imported into the desired instance of Oracle VM Manager by going to the **Resources** tab in the UI and clicking on the **Virtual Machine Templates** subtab, and then **Import**.

Within the Import wizard screens, you can choose to import from an external location, including directly from Oracle, via ftp or http if the Template is not in the `/OVS/seed_pool` directory, or you can **Discover and Register** Templates that are already in the `/OVS/seed_pool` directory as described above.

Import Template

* Name

Component **AGENT_TEMPLATE_RC**

Description **123**

Virtualization Type **Paravirtualized**

OS Type **Linux**

Add Server Pools

|

Select Name

<input checked="" type="radio"/>	ca-emhostsPool
----------------------------------	----------------

Monitoring Server Credentials

Specify the credentials to be used for provisioning. You can either choose to use the preferred credentials

Monitoring Server **stbdh19.us.oracle.com**

Use Preferred Credentials Override Preferred Credentials

Virtual Server Host Credentials

Specify the credentials to be used for provisioning. You can either choose to use the preferred credentials that the user you specify here must be allowed to use sudo on that virtual server host.

Virtual Server Host **ca-emhost3.us.oracle.com**

Use Preferred Credentials Override Preferred Credentials

Figure 14: Import VM template

Create Guest VMs Using a VM Template

Figure 15 shows a workflow to create guest VMs from the Oracle VM Templates.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

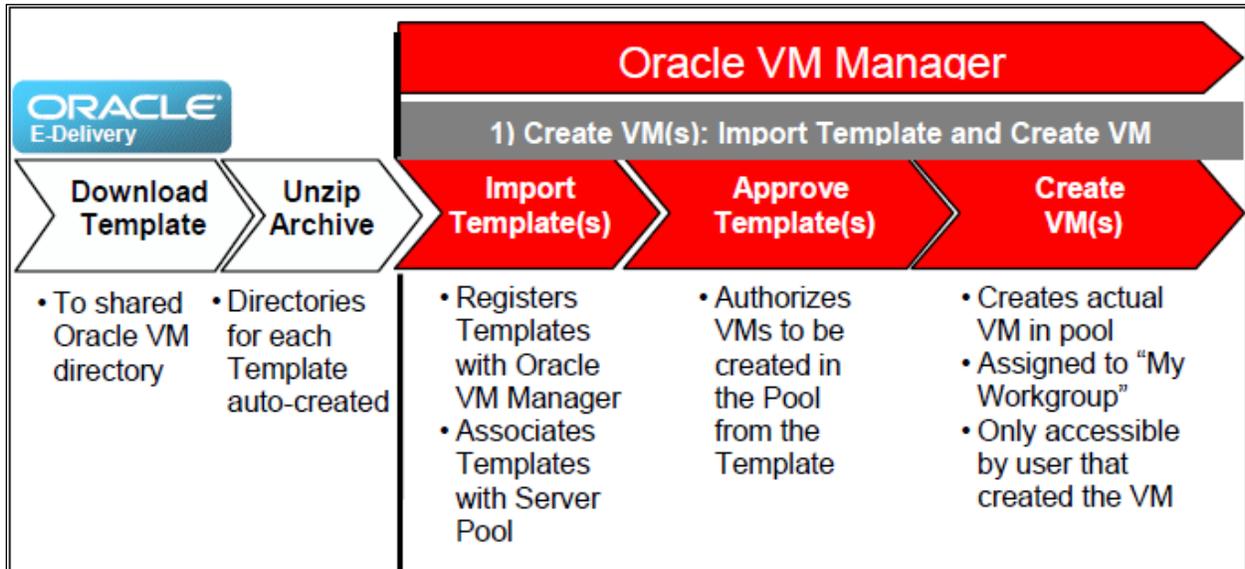


Figure 915 Overview of process for creating VMs from Oracle VM templates

1. Discover/approve the VM template
 - a. In the VM Management Pack **Virtualization Central** tab, select **VM server pool kblade1** and the **Action: Virtual server Pool, Discover Guest VM**.

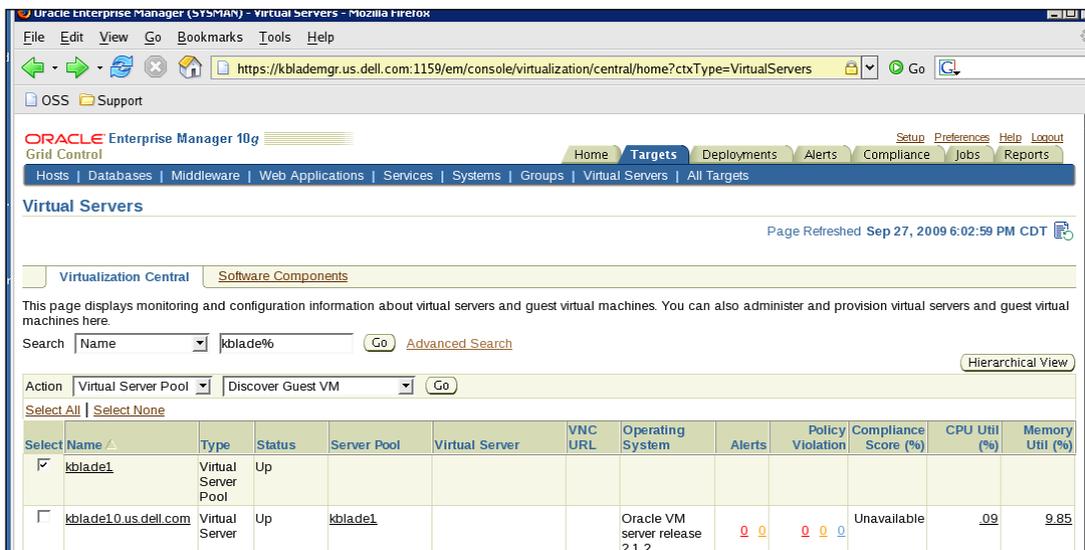


Figure 16: Select a VM server pool to discover a VM template

- b. Select undiscovered template.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

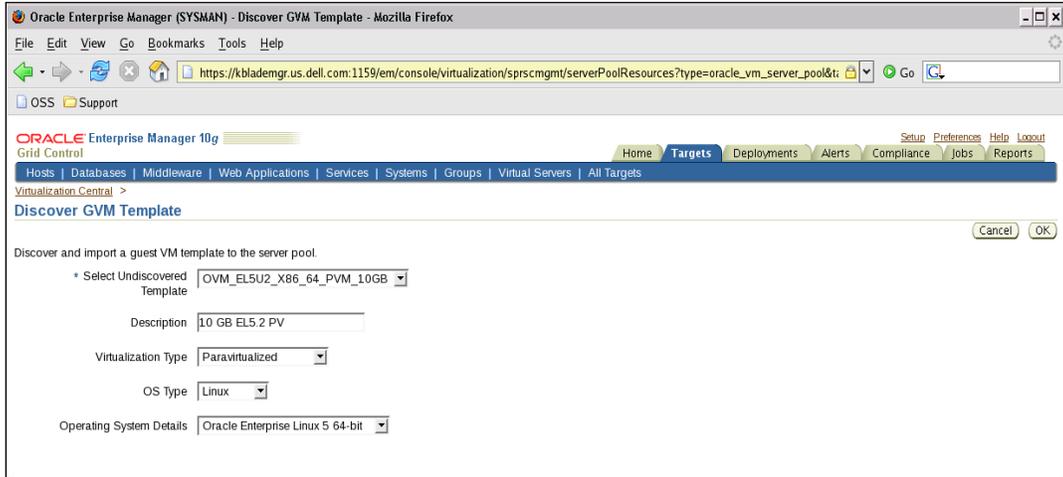


Figure 17: Discover a VM template

2. OVM_EL5U2_X86_64_PVM_10GB template is discovered, listed as a VM template resource in VM server pool kblade1, and is ready to be used to create a guest VM.

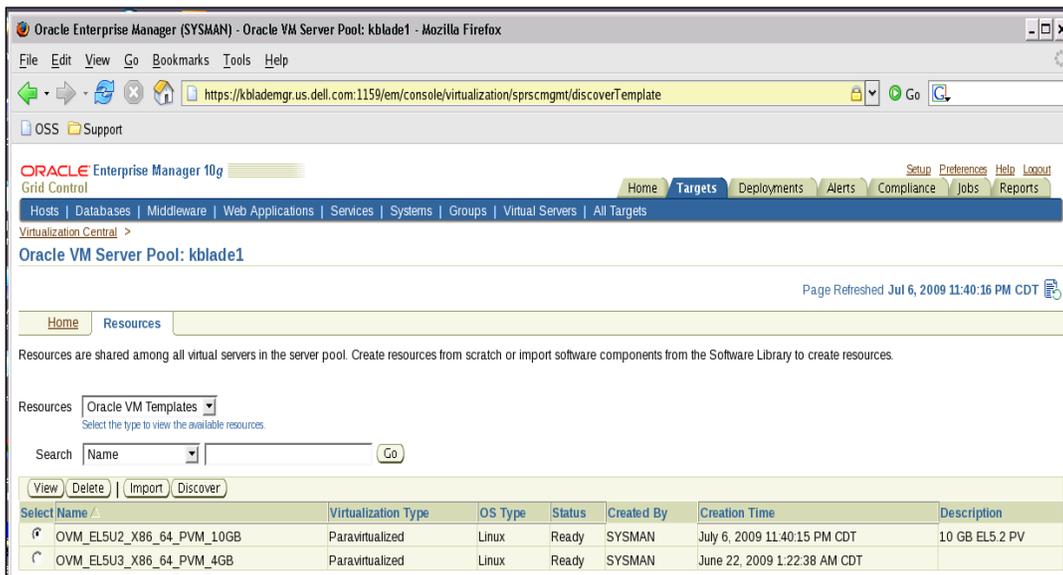


Figure 18: Select a VM template to add to the VM server pool

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

3. Create a guest VM using the Template :
 - a. In **Virtualization Central**, select **Action as Guest VM, Create Guest VM** and click **Go**.
 - b. In the **Create Guest Virtual Machines: Server Pool** page, specify the number of guest virtual machines to be created and the server pool with which you want to associate the guest virtual machines.
 - c. Specify the preferred virtual servers where the guest virtual machines will run. Choose **Automatic** to select the virtual server automatically. Click **Next**.
 - d. In the **Create Guest Virtual Machines: Source** page, select **Oracle VM Template**.

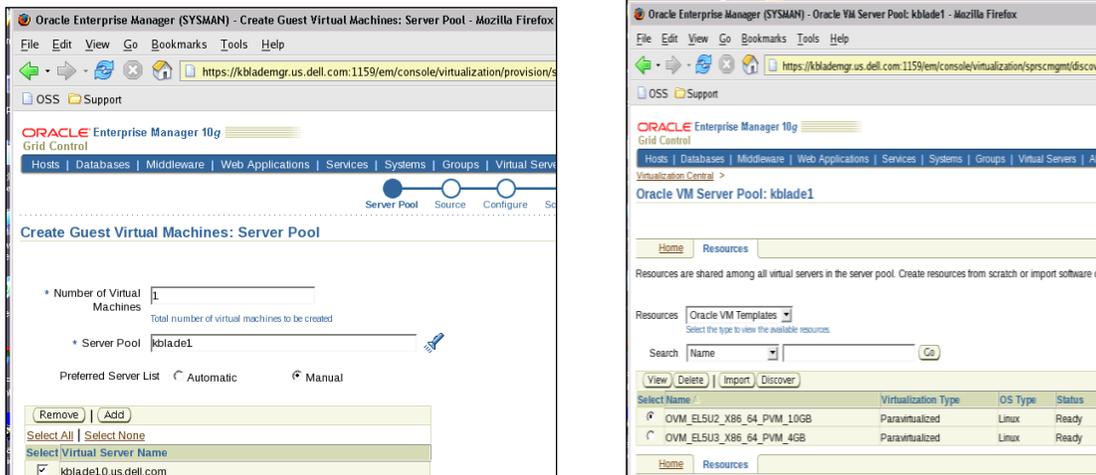


Figure 19 Start creating a guest VM from a VM template

- e. Specify the parameters for resource assignment for this guest VM such as virtual CPUs and memory and network configuration, local disks and shared disk.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

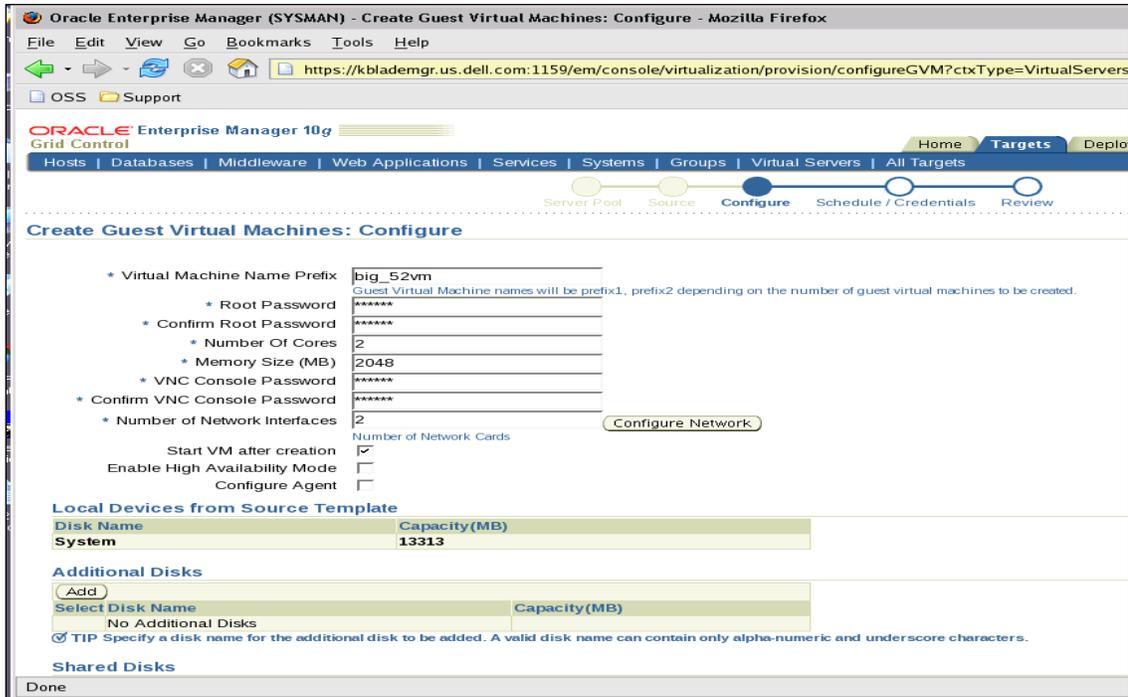


Figure 20: Specify the parameter values for the guest VM creation

f. Guest VM is created and running on VM server kblade 10:

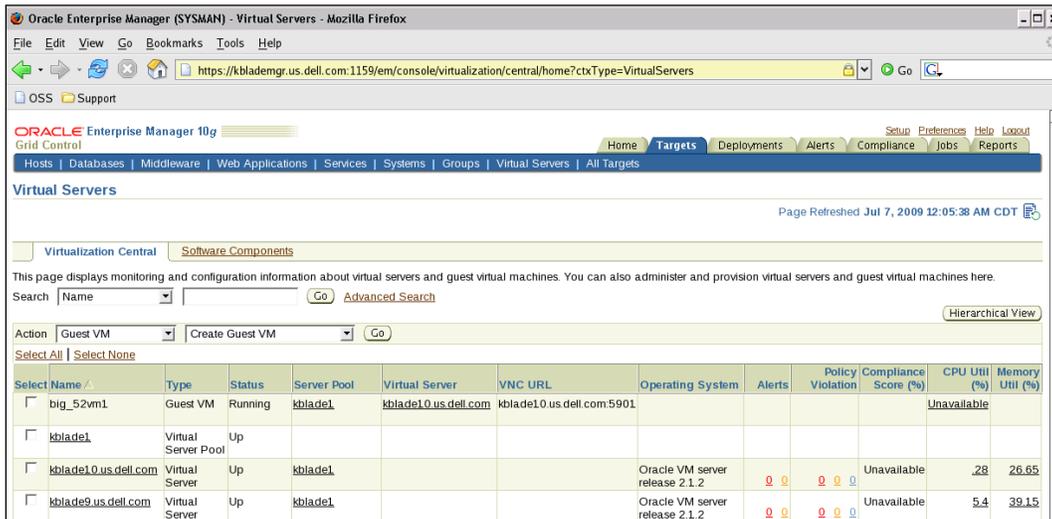


Figure 21: A new guest VM big_52vm1 was created and is running on VM server kblade10.us.dell.com

Customizing Oracle VM Templates to Create “Gold Images”

By their nature, Templates typically reflect the general best practices for the products they contain, but any one configuration cannot cover all scenarios, especially in complex enterprise environments. As a result, you may wish to update the Templates and then save those changes in your Template VMs so that any future VMs created from that Template would automatically reflect those customizations. This is easily done by making the changes and then selecting the “Save As Template” option while the VM is in a Powered Off state.

By using this option, the VM is then saved in the central `/OVS/seed_pool` directory as a template that can be used to create VMs in the future. Note that the amount of time required to complete the Save as Template operation depends on the VM size and the performance of your shared storage since the VM is effectively cloned and saved to the `/OVS/seed_pool` directory.

For example, if the VM template only has one image file, `system.img`, additional disks may need to be added as the local disk partition, as well as the shared disks in the guest VM, to install the application. If the guest VM is created for Oracle RAC database, an additional local disk needs to be added for Oracle clusterware and Oracle RAC software homes. Shared disks may also need to be added for clusterware OCR and Votingdisk and RAC database. The local disk and shared disk can be specified during the guest VM creation process or can be added after the VM creation as follows:

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

1. Add local disks:
 - a. In the **Virtual Central** page of VM Management Pack, select the **Guest VM**.
 - b. Select Edit Guest VM from the **Action** menu.
 - c. Go to the **Edit Guest Virtual Machine** page.
 - d. Add local disk: Click **Add Disk Dtab** and specify the disk name and size of the local disks.

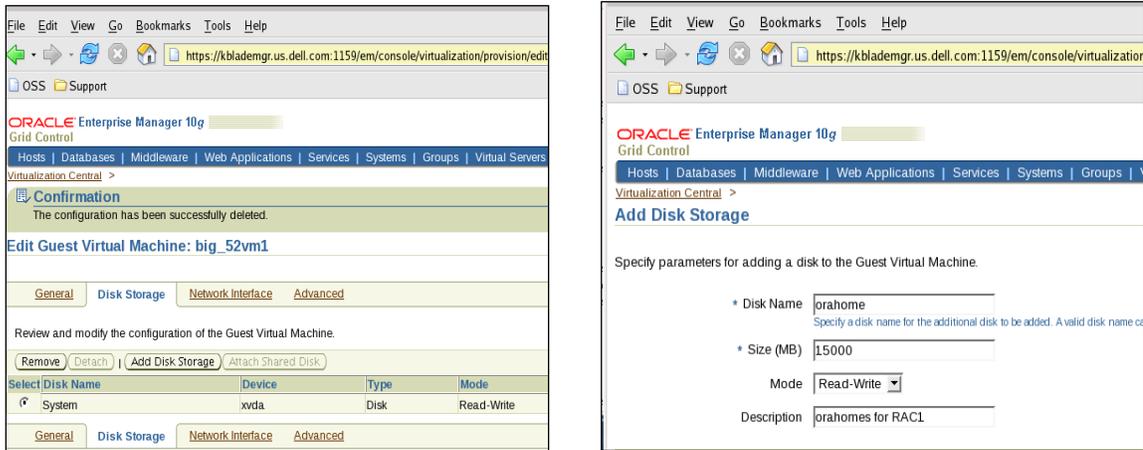


Figure 22: Add a local disk to the guest VM

2. Create shared disk racdb on server pool and attach to the share disk racdb to the guest VM.
3. Create shared disk and attach the shared disk to VM.

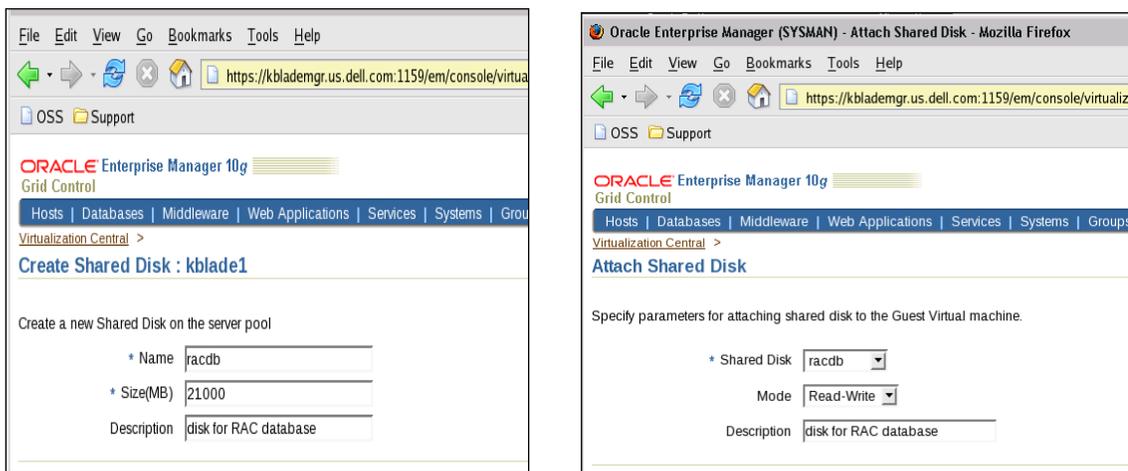


Figure 23: Add a shared disk racdb to the guest VM

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

4. Add shared disk. Click the **Attach Share** disk button and specify the name and size of the shared disk to attach.

Now the guest VM has two disks, `system` for OS, `orahome` for Oracle software, and one share disk, `oradb`, for Oracle database.

Two local disks correspond to image files in the OVM repository:

```
ls -l /OVS/running_pool/150_big_52vm2/* .img
/OVS/running_pool/150_big_52vm2/orahome.img
/OVS/running_pool/150_big_52vm2/System.img
```

And a shared disk corresponds to `racdb.img` in the repository:

```
ls /OVS/sharedDisk/racdb.img
/OVS/sharedDisk/racdb.img
```

They are also shown as the resources for the guest VM.

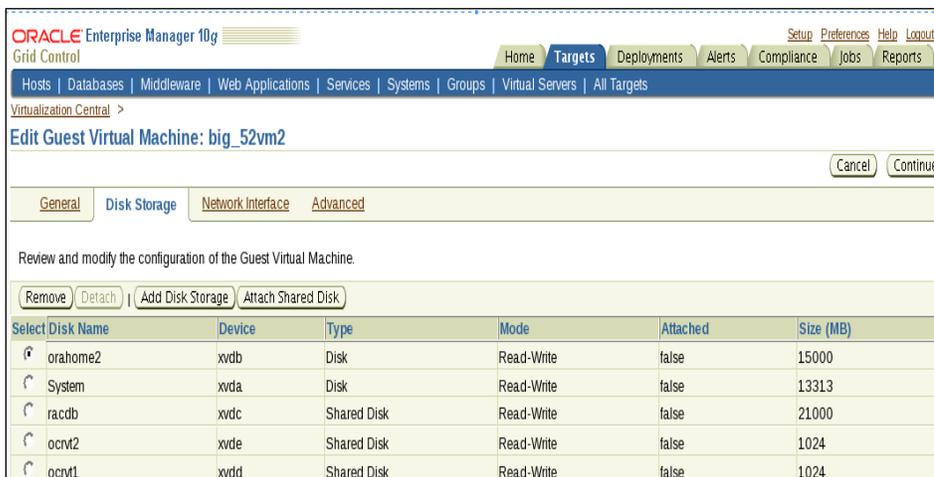
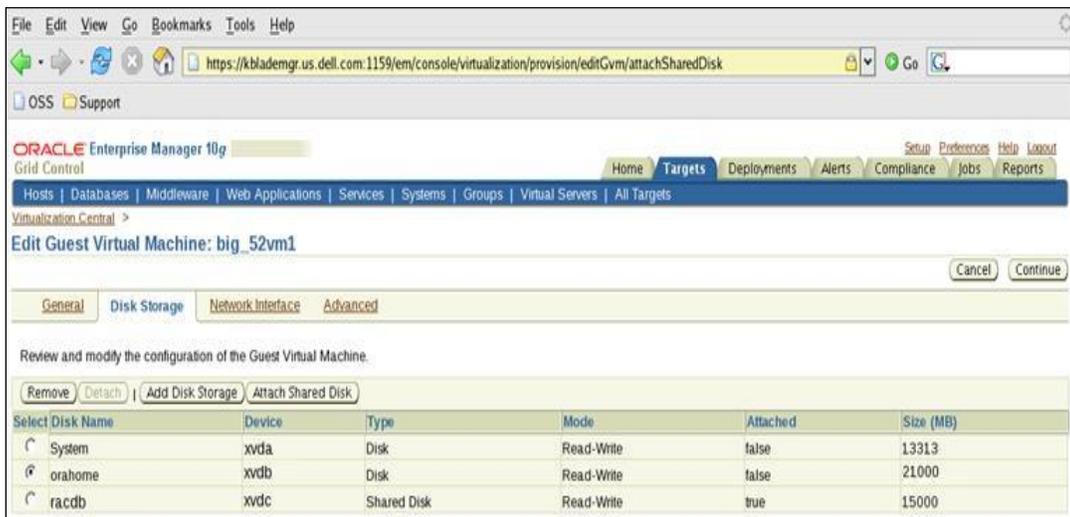


Figure 24: The guest VM has one local disk and three shared disks

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

In the guest VM OS, the virtual disk partitions can be seen.

```
more /proc/partitions
major minor #blocks name

202      0 13631488 xvda
202      1   104391 xvda1
202      2 13526730 xvda2
16             15360000 xvdb
17             15342043 xvdb1
202     32 21504000 xvdc
202     33 21486906 xvdc1
```

The virtual disks and the associate image files in the OVM repositories

Virtual device Name p in VM	disk name in VM	image file in /OVS
xvda	System	System.img
xvdb	orahome	orahome.img
xvdc	Racdb	oradb.img

Table 2: Virtual disks and the associated image files

5. Attach disk partition in dom0 to the guest VM:

The disk partition, such as the multipath device, can be attached to the guest VM by specifying the physical disk mapping in the vm.cfg file. The following configuration attached the disk partition in dom0 /dev/mapper/vmracdbp1 to the guest VM as the virtual disk /dev/xvdc

```
vm.cfg:
disk = [ 'phy: /dev/mapper/vmracdbp1, xvdc,w!'
]
```

This configuration currently cannot be configured using Grid Control Virtual Management Pack or VM Manager. Manually editing the vm.cfg file of the guest VM is required.

For the mapping from the image file in the OVM repository to the virtual device in the guest VM, the following mapping in vm.cfg file was generated by the VM Management Pack:

```
vm.cfg:
disk = ['file:/OVS/running_pool/150_big_52vm2/System.img,xvda,w',
'file:/OVS/running_pool/150_big_52vm2/orahome.img,xvdb,w',
'file:/OVS/sharedDisk/racdb.img,xvdc,w!'],
```

Configure Guest VMs for Oracle 11g RAC Database

Oracle 11g RAC database is certified to run on the guest VMs. To configure the guest VMs as the database nodes for Oracle 11g RAC configuration, the guest VMs are configured with network and shared storage as shown in the following diagram:

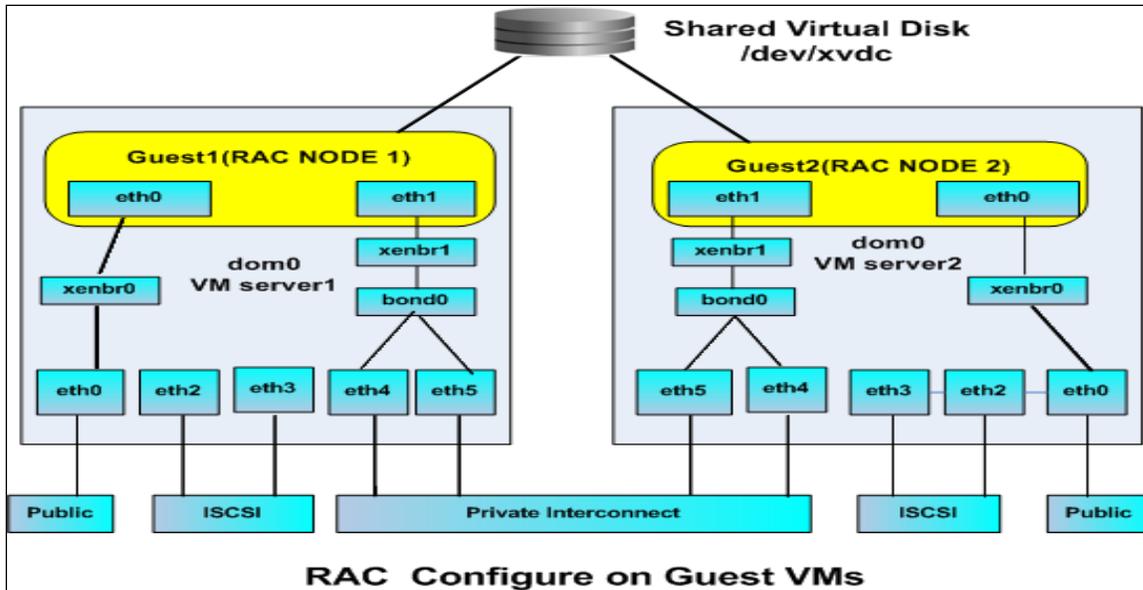


Figure 25: Guest VMs and OVM Server Configuration for 11g RAC on VMs

The shared virtual storage should be configured for 11g RAC clusterware as well as the Oracle RAC database. The shared virtual storage can be either the shared virtual disks from the OVM repository or the disk partition on the OVM server dom0, as illustrated in the previous session. For a production database or an IO intensive database, it is recommended that you use the disk partition on dom0 for the shared virtual disk on guest VM.

Deploying Real Application Cluster using Enterprise Manager Provisioning and Patch Automation Pack

Enterprise Manager comes with out-of-box Deployment Procedures to install the Oracle RAC Database, Oracle Clusterware and Oracle Automatic Storage Management from gold images following the best practices for maximum availability.

When you provision Oracle RAC the Deployment Procedures deploy the following core components:

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

- Oracle Clusterware
- Oracle RAC Database
- Automatic Storage Management (ASM) (optional)

You can deploy ASM either in the same Oracle home as the one for Oracle RAC Database, or in a completely different Oracle home (recommended).

Pre-Requisite for RAC Provisioning

Before we start with RAC provisioning on Guest VM servers, make sure that the following Pre-requisites are met:

1. Clustered Management Agent is installed on Guest VM machines. The recommended method for installing Clustered Management Agent is by using the Agent push application from the Oracle Enterprise Manager. Refer to [Enterprise Manager Grid Control Installation and Configuration Guide](#) for more details.
2. Oracle Software Library is set up. *Oracle Software Library* (Software Library) is a repository that stores certified software images (for example, Oracle Database, operating system, Oracle Real Application Clusters, third party software, and so on) and other related entities. These can then be automatically mass-deployed to provision software, software updates, and servers using Enterprise Manager Grid Control in a reliable and repeatable manner. These provisioning operations, which are unattended and can be scheduled, lead to substantial cost savings.
 - a. Access the Oracle Enterprise Manager Provisioning Application by navigating to the **Deployments** tab.
 - b. Under the **Deployments** tab, go to the **Provisioning** tab.
 - c. Access the **Administration** tab. This requires super administrator privileges similar to SYSMAN user.
 - d. In the Software Library Configuration section of the **Administration** tab, click **Add**.
 - e. On the Add Software Library Location page, enter the directory location and click **OK**. When the Software Library is configured, out-of-box Provisioning Archive files (PAR files) will be deployed. These files contain pre-build entities such as components, directives, and so on for various applications such as Bare Metal Provisioning and Patching.
3. Software Library components have been created. *Components* are entities in the Software Library that represent the primary building blocks of Deployment Procedures. Components are individually maintained within the Software Library, and you can associate versions, states, and maturity levels with each component.
 - a. In Grid Control, click **Deployments**, and select the **Provisioning** tab.
 - b. In the **Components** tab, expand **Oracle Component** → **RAC Provisioning**
 - c. Select the required version (10g or 11g), and further expand it to select the required platform.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

- d. Select the component **Oracle Database shiphome** and click **Edit** to upload the 11g Database Installation binaries.
- e. Select the component Oracle **Clusterware shiphome** and click **Edit** to upload the CRS Installation binaries.

Once the binaries are uploaded successfully the status of the Components will change from **Incomplete** to **Ready**.

4. Prepare shared devices: 2 X OCR and 3 X votingdisk, ASM disk. As raw device service has been depreciated in OEL5.x, we used the two rules files in /udev/rules.d to establish the raw devices mapping:

/udev/rules.d/65-raw.rules to map the block devices to raw device:

```
root@bvmrac1 rules.d]# more 65_raw.rules
ACTION=="add", KERNEL=="xvdc1", RUN+="/bin/raw /dev/raw/raw1 %N"
ACTION=="add", KERNEL=="xvdd1", RUN+="/bin/raw /dev/raw/raw2 %N"
ACTION=="add", KERNEL=="xvdd2", RUN+="/bin/raw /dev/raw/raw3 %N"
ACTION=="add", KERNEL=="xvdd3", RUN+="/bin/raw /dev/raw/raw4 %N"
ACTION=="add", KERNEL=="xvde1", RUN+="/bin/raw /dev/raw/raw5 %N"
ACTION=="add", KERNEL=="xvde2", RUN+="/bin/raw /dev/raw/raw6 %N"
ACTION=="add", KERNEL=="xvde3", RUN+="/bin/raw /dev/raw/raw7 %N"
```

/udev/rules.d/89-raw.rules to set the proper permission and ownership of the raw devices:

```
[root@bvmrac1 rules.d]# more 89-raw.rules
#OCRs
KERNEL=="raw2",OWNER="root",GROUP="oinstall",MODE="640"
KERNEL=="raw5",OWNER="root",GROUP="oinstall",MODE="640"
#Votingdisk
KERNEL=="raw3",OWNER="oracle",GROUP="oinstall",MODE="640"
KERNEL=="raw6",OWNER="oracle",GROUP="oinstall",MODE="640"
KERNEL=="raw7",OWNER="oracle",GROUP="oinstall",MODE="640"
#ASM disks
KERNEL=="raw1",OWNER="oracle",GROUP="dba",MODE="640"
#ASM spfile
KERNEL=="raw4",OWNER="oracle",GROUP="dba",MODE="640"
```

5. And run the udev script to make the rules take effect immediately: /sbin/start_udev

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

1. Establish the network interfaces:
As illustrated in Figure 10, each Guest VM should have public network/private network configuration:
 - a) Public network interfaces eth0 connects to the physical network interface eth0 on the VM server through the xenbr0 Xen bridge.
 - b) Private network interface eth1 through xen bridge xenbr1 connects to the bond0 bonding by eth4 and eth4 physical network interfaces of VM server.
2. Prepare /etc/hosts to include public, private, VIPs: hostnames/IPs .

RAC Provisioning using Deployment procedures

Follow the following steps to perform RAC provisioning in an unattended manner:

1. In Grid Control, click the **Deployments** tab.
2. On the **Deployments** page, in the **Deployment Procedure Manager** section, click **RAC Provisioning Procedures**.
3. On the **Deployment Procedure Manager** page, in the **Procedure** subtab, from the table, select **Oracle Clusterware / RAC Provisioning For UNIX**.
4. In the **Select Source** section, select **Select from Software Library**.
5. In the **Source for Clusterware** section, click the torch icon and select the generic component that has the software binaries of Oracle Clusterware. Ensure that you select only components that are in "Ready" status. Once you select the component name, the application automatically displays the component location.

Source for Clusterware	
Name	Oracle Clusterware shiphome
Location	Components/Oracle Components/RAC Provisioning/11.1.0.6.0/linux32/Oracle Clusterware shiphome

Figure 26: Select 11g clusterware shiphome component from software library as the source for clusterware provisioning

6. In the **Source for RAC** section, click the torch icon and select the generic component that has the software binaries of Oracle Database. Ensure that you select only components that are in "Ready" status. Once you select the component name, the application automatically displays the component location.

Source for RAC	
Name	Oracle Database shiphome
Location	Components/Oracle Components/RAC Provisioning/11.1.0.6.0/linux32/Oracle Database shiphome

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

Figure 27: Select 11g RAC database shiphome component from software library as source for the RAC database provisioning

7. On the **Select Hosts** page, click Add and select the target hosts that should form the cluster. By default, **Private Host Name** and **Virtual Host Name** are automatically prefilled with values. Edit them and specify values that match your environment. Optionally, you can also specify their IP addresses.
8. Provide credentials for the Oracle home.
9. On the **Configure Cluster** page, provide the **Cluster Name**, **Location** etc.

Cluster Name and Location
Provide a name for the cluster. Specify the installation path. If you choose to create a new Oracle Base for RAC or ASM, you must specify the Oracle Base for RAC or ASM.

* Cluster Name	<input type="text" value="stnsp11-12"/>	Cluster Name cannot contain characters such as #,\$%!.
* Clusterware Home Location	<input type="text" value="/u01/cluster"/>	It is not recommended to specify the Clusterware Oracle Base as a subdirectory of the Oracle Base for RAC or ASM.
* Oracle Base for RAC	<input type="text" value="/u01/app/asdf123"/>	
* Database Home Location	<input type="text" value="/u01/app/asdf123/product/11.1.0/Oracle"/>	It is recommended to specify the RAC Oracle Home as Oracle Base for RAC.
* Oracle Base for ASM	<input type="text" value="/u01/app/asdf123"/>	
* ASM Home Location	<input type="text" value="/u01/app/asdf123/product/11.1.0/Oracle"/>	It is recommended to specify the ASM Oracle Home as Oracle Base for ASM.
* Scratch Location	<input type="text" value="/tmp/rac/"/>	
Additional Parameters	<input type="text"/>	Eg: -debug

Figure 28: Specify clusterware configuration

10. In the **Database Details** section, retain the default selection for creating a starter database.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

Figure 29: Specify RAC database configuration

11. In the **Shared Storage Configuration** section, provide details about the storage devices and click **Next**.

Partition Name	Mount Location	Mount Format	Size	OCR	OCR Mirror	VDSK	VDSK1	VDSK2	Data Files	Remove
/dev/xvdd1	/dev/raw/raw2	raw		<input type="checkbox"/>						
/dev/xvde1	/dev/raw/raw5	raw		<input type="checkbox"/>						
/dev/xvdd2	/dev/raw/raw3	raw		<input type="checkbox"/>						
/dev/xvde2	/dev/raw/raw6	raw		<input type="checkbox"/>						
/dev/xvde3	/dev/raw/raw7	raw		<input type="checkbox"/>						
/dev/raw/raw1	/dev/raw/raw*	asm		<input type="checkbox"/>						
Do not provision storage										

Figure 30: Specify the shared storage for the RAC database provisioning

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

12. On the **Review** page, review the details you have provided for provisioning Oracle RAC, and click **Submit**.
13. Check the status of the provisioning process and ensure it completes successfully.

VM template for RAC

After you have completed RAC provisioning on the VM Guest machines, you can create a Template from the guest VM so that it will capture the RAC installation details. The next time you use this VM template to create new guest VMs, the template will put the RAC binaries on the new VMs and you will only have to configure the new guest VM. This may also be done by modifying the out-of-box deployment procedures.

Extending RAC

The previously installed RAC cluster can be extended to “N” number of nodes by using any of the following methods:

1. Single Click Extend Deployment procedure of the Provisioning and Patch Automation Pack.
2. Clone a VM: Create a template out of the existing RAC on a VM and use this Template for further VM cloning.

In the next section of this paper we will be discussing the first method. To extend an existing Oracle RAC, follow these steps:

1. Prepare a new guest VM as the new RAC node by configuring it in the same way as the existing VM RAC node was configured, including VM OLE OS kernel parameters and configuration, shared virtual storage and raw devices on the shared storage, public and private networks and VIP, Oracle enterprise manager agent 10.2.0.5, etc.
2. In Grid Control, click the **Deployments** tab.
3. On the **Deployments** page, in the **Deployment Procedure Manager** section, click **RAC Provisioning Procedures**.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

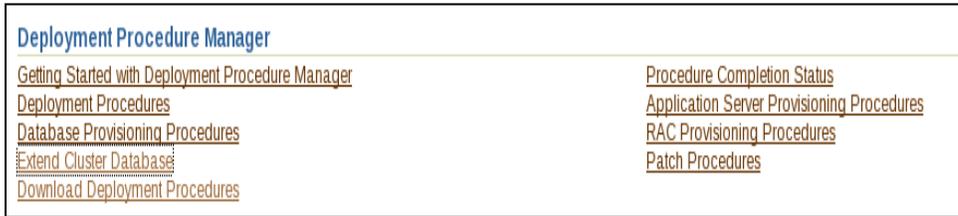


Figure 31: Navigate to RAC provisioning procedure

4. On the **Deployment Procedure Manager** page, in the **Procedures** subtab, select **One Click Extend Cluster Database** from the table. Click **Schedule Deployment**.
3. In the **Select Real Application Clusters (RAC)** section, select the Oracle RAC you want to extend. The associated clusterware and Automatic Storage Management (ASM) also get extended if they are not already existing.

Expand All Collapse All		Member Nodes	Oracle Home	Platform	Product
Available Cluster Databases					
<input checked="" type="radio"/>	▼ pqttest (3)	strrac13.us.oracle.com, strrac15.us.oracle.com, strrac14.us.oracle.com	/opt/oracle/oracle11g/db	Enterprise Linux Enterprise Linux AS release 4 (October Update 4)	Oracle Database 11.1.0.7.0
	str13_cluster (3)	strrac13.us.oracle.com, strrac14.us.oracle.com, strrac15.us.oracle.com	/opt/crs11g	Enterprise Linux Enterprise Linux AS release 4 (October Update 4)	Oracle Clusterware 11.1.0.7.0
	+ASM1_strrac13.us.oracle.com (3)	strrac13.us.oracle.com, strrac15.us.oracle.com, strrac14.us.oracle.com	/opt/oracle/oracle11g/asm	Enterprise Linux Enterprise Linux AS release 4 (October Update 4)	Oracle ASM 11.1.0.7.0

Figure 32: Select a RAC database to extend

4. In the **Reference Host Options** section, from the **Reference Host** list, select a host that you want to use as the primary host for performing this operation. Reference Host is the host that is selected for creation of clone archives and then transferred to the new target nodes being added.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

Home Location	Files To Exclude
/opt/crs11g	log/*,crs/init,racg/dump,svrm/log,cdata,EMStagedP
/opt/oracle/oracle11g/db	log/*,EMStagedPatches,oratab,*.trc,*.dbf,cdump,lis
/opt/oracle/oracle11g/asm	log/*,EMStagedPatches,oratab,*.trc,*.dbf,cdump,lis

Figure 33: Specify the reference host

5. In the **Select New Nodes** section, click **Add** to add new nodes that you want to include to the selected Oracle RAC.

Host	Virtual Node Name	Remove
strrac16.us.oracle.com	strrac16-vip.us.oracle.com	

Figure 34: Select the new node to add to the RAC

6. In the **Schedule** section, schedule the **Deployment Procedure** to run either immediately or later.
7. In the **Prerequisites (Run Prerequisites and Fix-Ups)** section, by default, **Skip prerequisites and fix-ups** is not selected and therefore, the deployment procedure runs the prerequisite checks and fix-ups on the selected nodes.

The prerequisite checks are required to ensure that the nodes meet all the requirements of this operation and are ready to be added to the cluster. The option is not selected assuming that you have not already run the prerequisite checks on the selected nodes beforehand.

8. On the **Review** page, review the details you have provided for extending Oracle RAC, and click **Submit**.

Conclusion

In this paper, we explored some of the best practices that are related to provision and extend Oracle 11g RAC database on Oracle virtual machines:

1. Create Oracle VM infrastructure such as VM server and connect them to Oracle Enterprise Manager Virtual Management Pack.
2. Create guest VMs using VM templates.
3. Prepare guest VMs configuration such as shared storage and network for Oracle RAC provisioning.
4. Provision Oracle 11g RAC on guest VMs using Oracle Enterprise Manager Provisioning Pack.
5. Extend Oracle 11g RAC to new guest VMs using Oracle Enterprise Manager Provisioning Pack.

This has been a joint project between the Dell Oracle Solutions Engineering Team and the Oracle Enterprise Manager Product Team. The related materials have also been presented in *Oracle OpenWorld 2009*:

- 1) Session: #S308185, "Building an Oracle Grid with Oracle VM on Blade Servers and iSCSI Storage," Kai Yu and David Mar.
- 2) Session: #S312109, "Provisioning Oracle RAC in a Virtualized Environment, Using Oracle Enterprise Manager," Kai Yu and Rajat Nigam.

Appendix

For more information concerning the Provisioning and Patch Automation Pack refer to the Oracle site [here](#).

To learn about solutions to common problems and scenarios that you might encounter while provisioning and patching your multiple databases in the datacenter refer to [Oracle® Enterprise Manager Grid Control Installation and Configuration Guide](#)

You can also refer to the recorded viewlets available on the [Oracle website](#) under the “Getting Started” section.

Efficient Deployment of Oracle Real Application Clusters in a Virtualized Environment Using Enterprise Manager Provisioning Pack

References:

- 1 Oracle Press Release: Oracle® Enterprise Manager Extends Management to Oracle VM Server Virtualization:http://www.oracle.com/us/corporate/press/017955_EN
- 2 Oracle® Enterprise Manager Concepts 10g Release 5 (10.2.0.5)
Part Number B31949-10
- 3 Oracle® Enterprise Manager Grid Control ReadMe for Linux x86-64
10g Release 5 (10.2.0.5) April 2009
- 4 How to Enable Oracle VM Management Pack in EM Grid Control 10.2.0.5, Metalink Note #: 781879.1
- 5 Oracle VM: Converting from a Local to Shared OVS Repository, Metalink note # 756968.1
- 6 How to Add Shared Repositories to Oracle VM Pools with Multiple Servers, Metalink Note #869430.1
- 7 [Deploying Oracle VM Release 2.1 on Dell PowerEdge servers and Dell/EMC storage](#)
Dell white paper
- 8 [Building an Oracle Grid with Oracle VM on Dell Blade Servers and EqualLogic iSCSI storage](#), Kai Yu, David Mar, Oracle OpenWorld 2009 Presentation Session #S308185 and Dell Whitepaper
- 9 Provisioning Oracle RAC in a Virtualized Environment, Using Oracle Enterprise Manager, Kai Yu, Rajat Nigam, Oracle OpenWorld 2009 Presentation, Session #S312109
- 10 Technical Best Practices for Virtualization & RAC—Oracle RAC SIG Webseminar slides: Michael Timpanaro-Perota & Daniel Dibbets
- 11 Oracle® Enterprise Manager Administrators' Guide for Software and Server Provisioning and Patching, 10g Release 5 (10.2.0.5.0), E14500-01, May 2009