



# Virtualizing Oracle 11g/R2 RAC Database on Oracle VM: Methods/Tips

**Saar Maoz, RACPack – RAC Development, Oracle**  
**Kai Yu, Oracle Solutions Engineering, Dell Inc**

# About Authors

- **Saar Maoz**

- Consulting Software Engineer, with Oracle since 1993
- Worked with many customers & partners
- Specializing Linux, worked on most other OS
- RAC Development on RAC OVM Templates

- **Kai Yu**



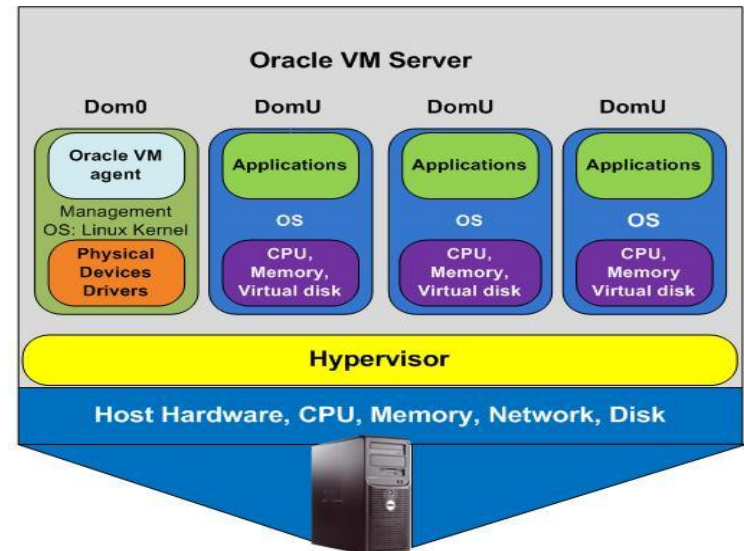
- Dell Oracle Solutions Engineering Lab
- 16 years with Oracle Technology: RAC, OVM, EBS
- IOUG Oracle RAC SIG President (2009-2010)
- IOUG Virtualization Committee Member
- 2011 OAUG Innovator of Year Award Winner
- Oracle Blog: <http://kyuoracleblog.wordpress.com/>

# Agenda

- ❖ Introduction to Oracle VM and Oracle RAC
- ❖ Establishing Oracle VM Infrastructure
- ❖ Virtual Machine Creation
- ❖ <Add more here>
- ❖ <Add more here>
- ❖ <Add more here>


# Introduction to Oracle VM and Oracle RAC

- Introduction of Oracle VM Virtualization Solution
  - Software based server Virtualization solutions
  - Server partitions and consolidation
  - High availability and scalability
  - Fast deployment and standardization
  - Provide the virtualization support for Grid /Cloud
- Oracle VM Architecture
  - Oracle VM server: Xen
    - Hypervisor, Dom 0 management Linux kernel with support of devices, IO, networking, etc
  - VM: Dom U, guest OS with applications running on domU



# Introduction to Oracle VM and Oracle RAC

- Virtual Server Pool: an autonomous region of VM servers
- Oracle VM Manager: GUI Management solution for Oracle VM
- Oracle VM Agent: installed with VM server and communicating with Oracle VM manager

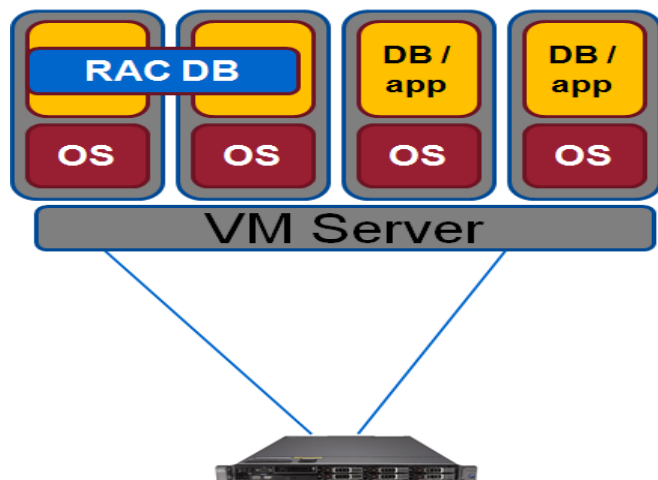


# Introduction to Oracle VM and Oracle RAC

## Why running RAC On Oracle VM


- Server partition and consolidation
- Sub-capacity licensing
- Virtual cluster on single physical server
- Template base deployment

Oracle VM



- Oracle VM
- Partitioning single server
- Consolidation
- Development/Test RAC

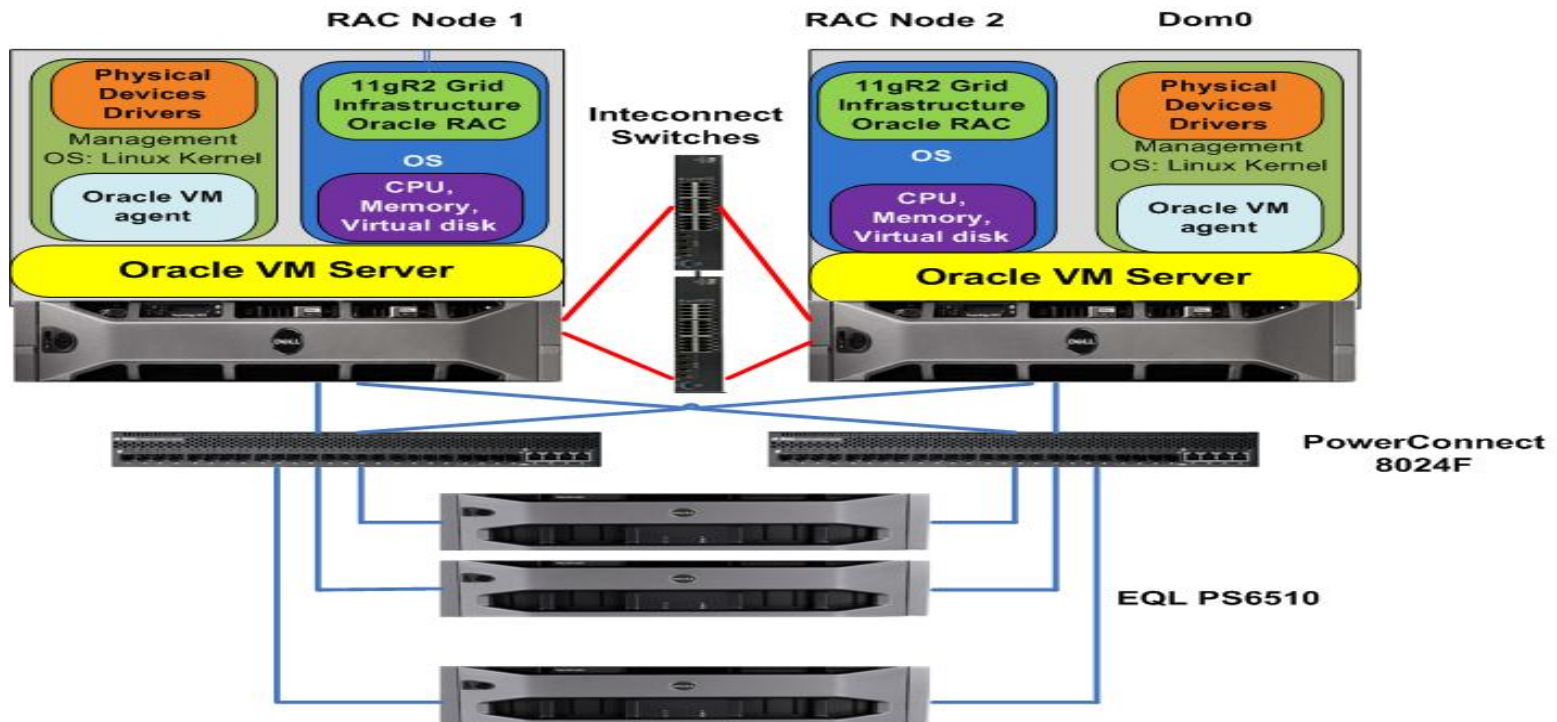
Oracle RAC on Oracle VM



- Server partitioning and clustering
- Scale-up and scale-out
- Production RAC

# Introduction to Oracle VM and Oracle RAC

- Deployment Methods of RAC On Oracle VMs
  - Guest VMs as Oracle RAC Nodes
  - Virtual Infrastructure for CPUs, Memory, Network, Storage
  - RAC VM templates or Manual RAC Configuration





# Introduction to Oracle VM and Oracle RAC

- Oracle VM 3.0 New Features
  - Application Driven Virtualization
  - Scalability: up to 160 physical CPUs, 2 TB of memory  
128 virtual CPUs and 1TB memory per guest VM
  - Oracle VM Manager provides centralized management
    - Fusion middleware Apps on Weblogic and Database
    - Provide Browser-based GUI Interface
    - Provide completed network and storage management
    - Cloning, sharing, configuring and migrating VMs.
    - Dynamic Resource Management for VMs
    - All VM operations are submitted as jobs
  - Centralized network management
    - Separate logic networks; management through GUI
  - Centralized Storage Configuration
    - Storage Connect, Automaticity storage repository configuration

# Establishing Oracle VM Infrastructure

- Implement Tasks Overview
  - Virtual Servers and VM Manager Installation
  - Virtual server network and storage configuration.
  - Guest VM creation and Network/storage configuration
  - Configure network and shared storage on VMs
- VM Server Installation:
  - Bare metal Installation, turn on the virtualization in BIOS
  - Change Dom 0 memory configuration to 1024GB in VM2.2
  - Oracle VM 3.0: no manual work on VM server after install

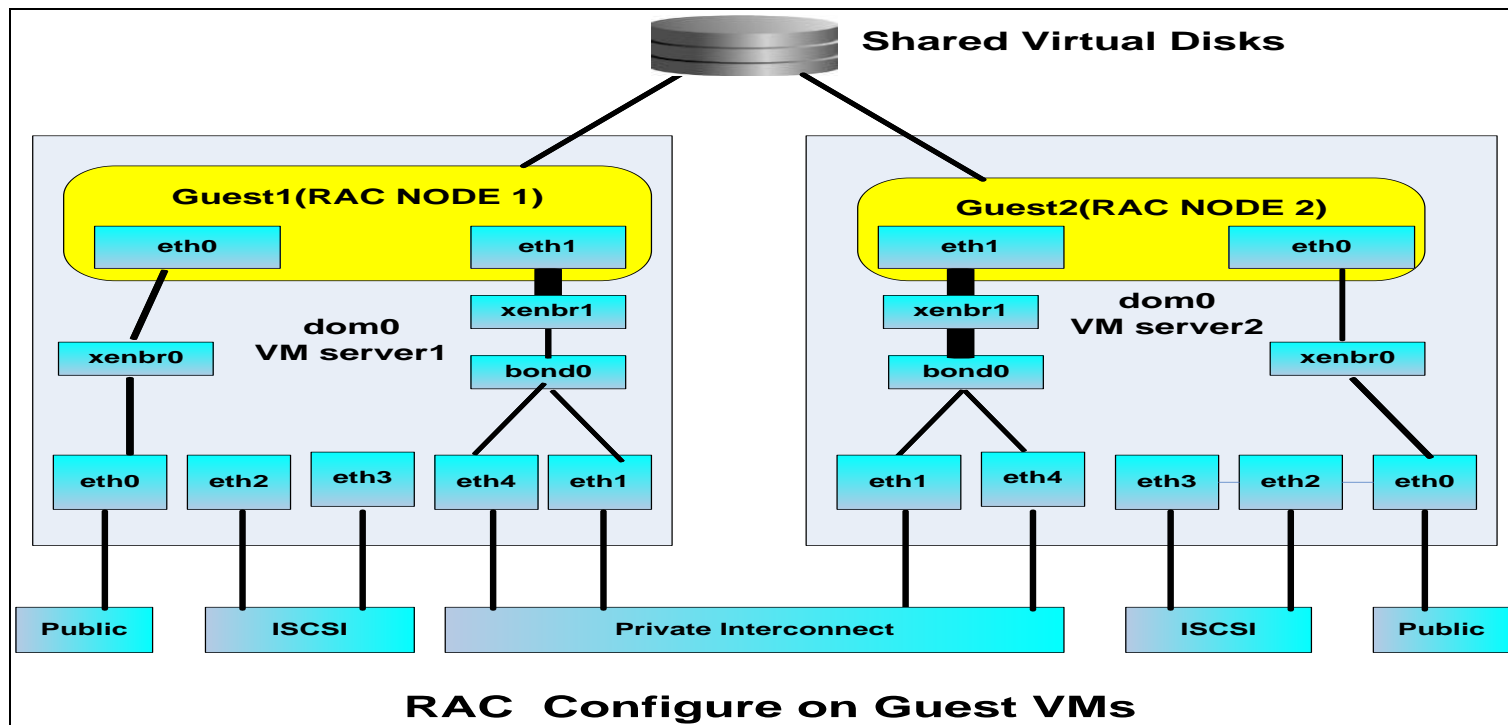
```
[root@kovmmgr2 ~]# ssh root@k4r710ovs2.dblab.com
root@k4r710ovs2.dblab.com's password:
Last login: Tue Sep 13 11:51:38 2011 from 192.168.9.81
Warning: making manual modifications in the management domain
might cause inconsistencies between Oracle VM Manager and the server.
[root@k4r710ovs2 ~]#
```

- VM Manager Installation:

–

# Establishing Oracle VM Infrastructure

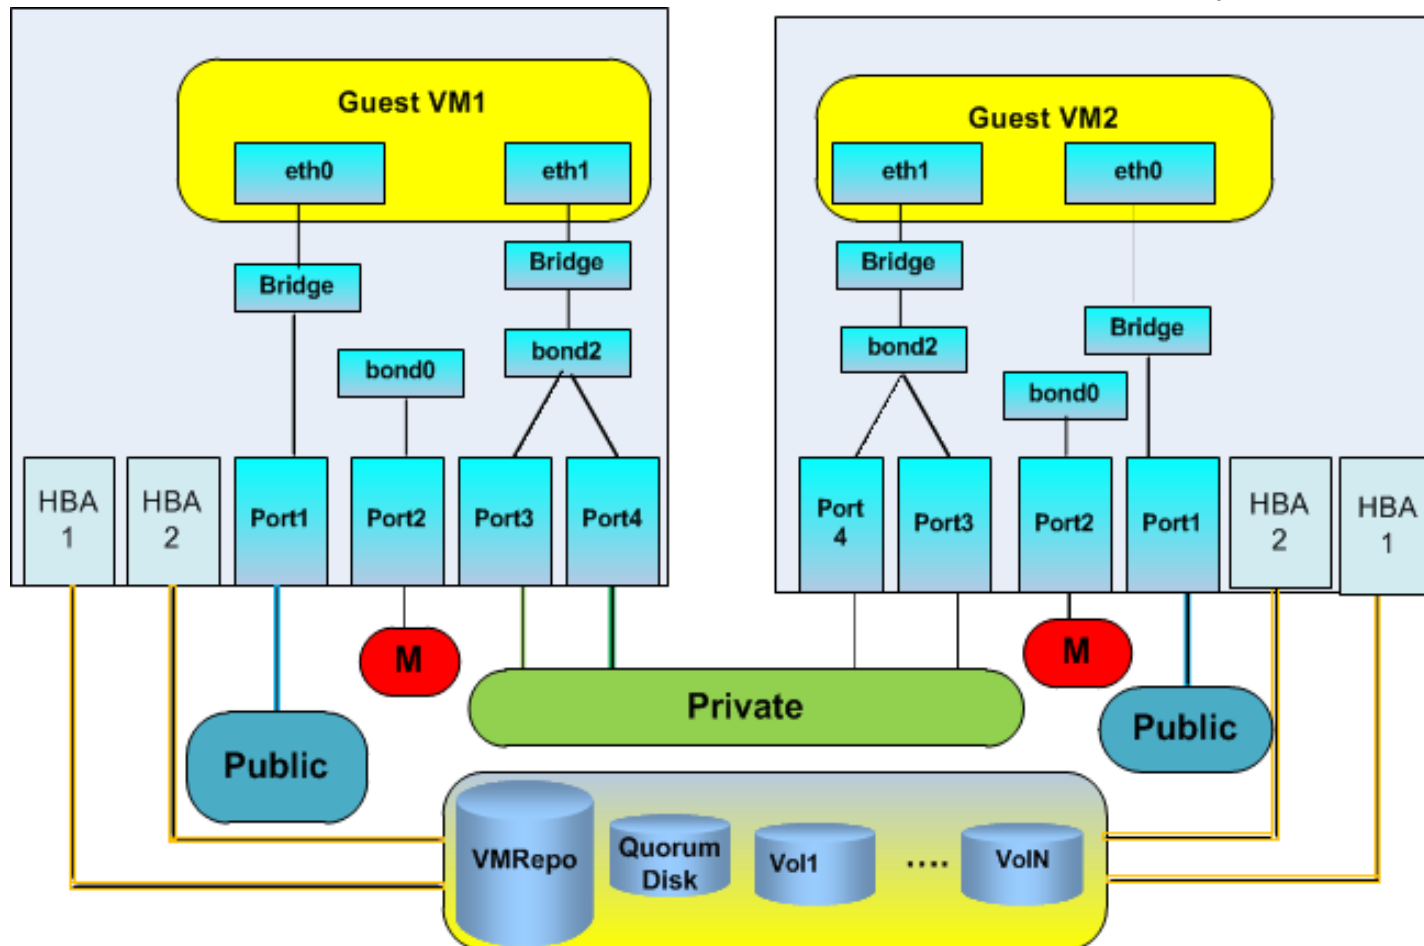
- Discover/connect VM servers vs OVS agent
- Virtual Infrastructure Design
  - Virtual Infrastructure based Oracle VM 2.2 (iSCSI storage)





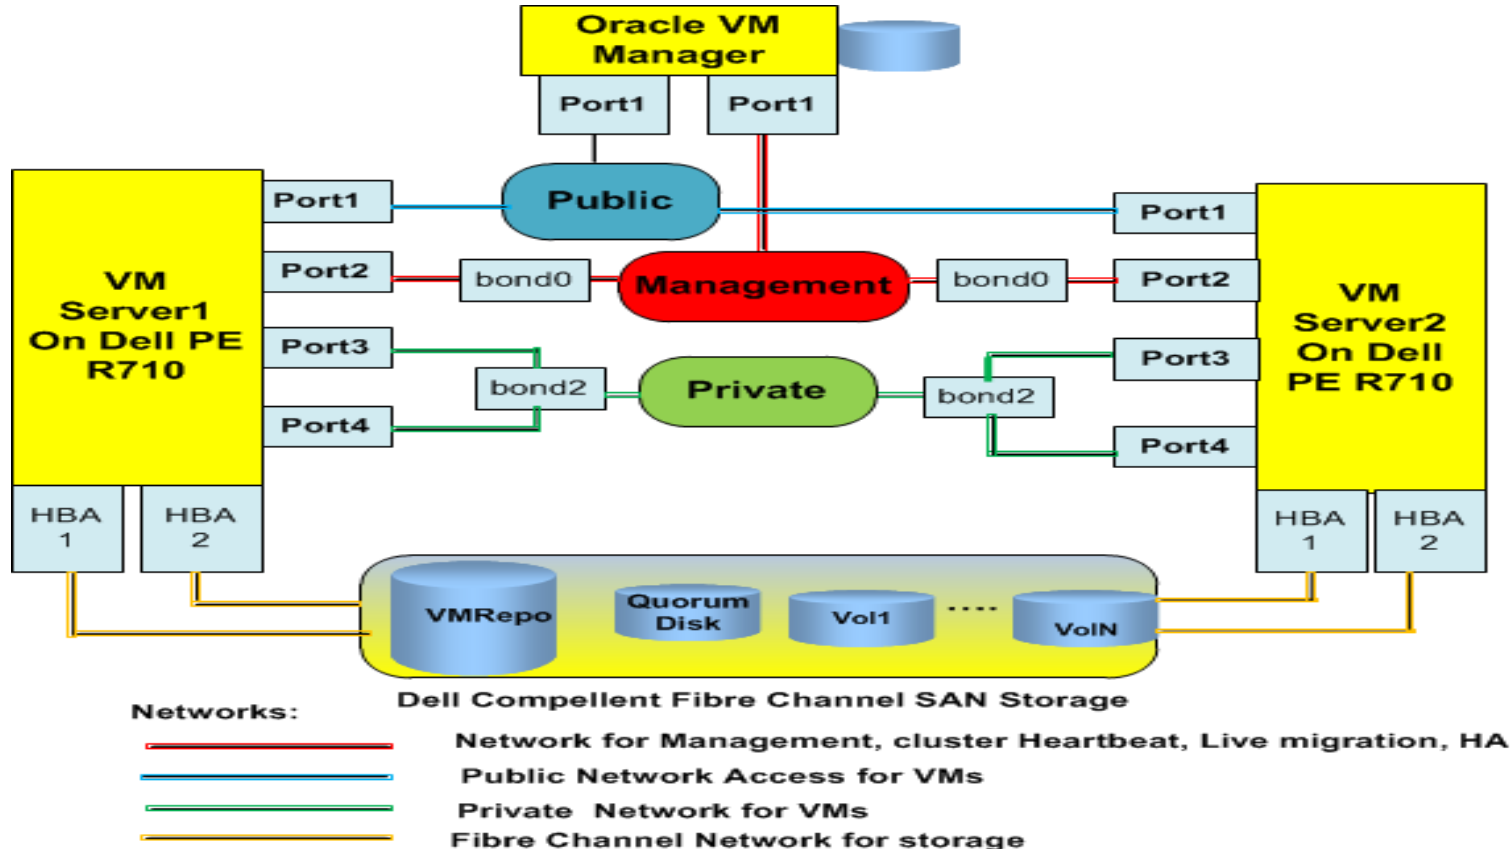
# Establishing Oracle VM Infrastructure

- Virtual Infrastructure based on Oracle VM 3.0 (FC storage)



# Establishing Oracle VM Infrastructure

- VM server network configure :
  - . Oracle VM 2.2 Manually create bonding/Xen bridges
  - . Oracle VM 3.0: separate networks, managed through VM Manager



# Establishing Oracle VM Infrastructure

## Separate Networks

The screenshot shows the Oracle VM Manager interface with the 'Networks' tab selected. A table displays the configuration for three networks: 'public', 'management', and 'private'. The 'private' network is highlighted in yellow.

Name	Intra Server	Network Use					VLAN Seg
		Server Management	Cluster Heartbeat	Live Migrate	Storage	Virtual Machine	
public						√	
management		√	√	√			
private						√	

## The bonding and ports for the private network

This screenshot provides a detailed view of the 'private' network configuration. It shows the network name, description, and a list of bonded ports with their IP addresses and masks.

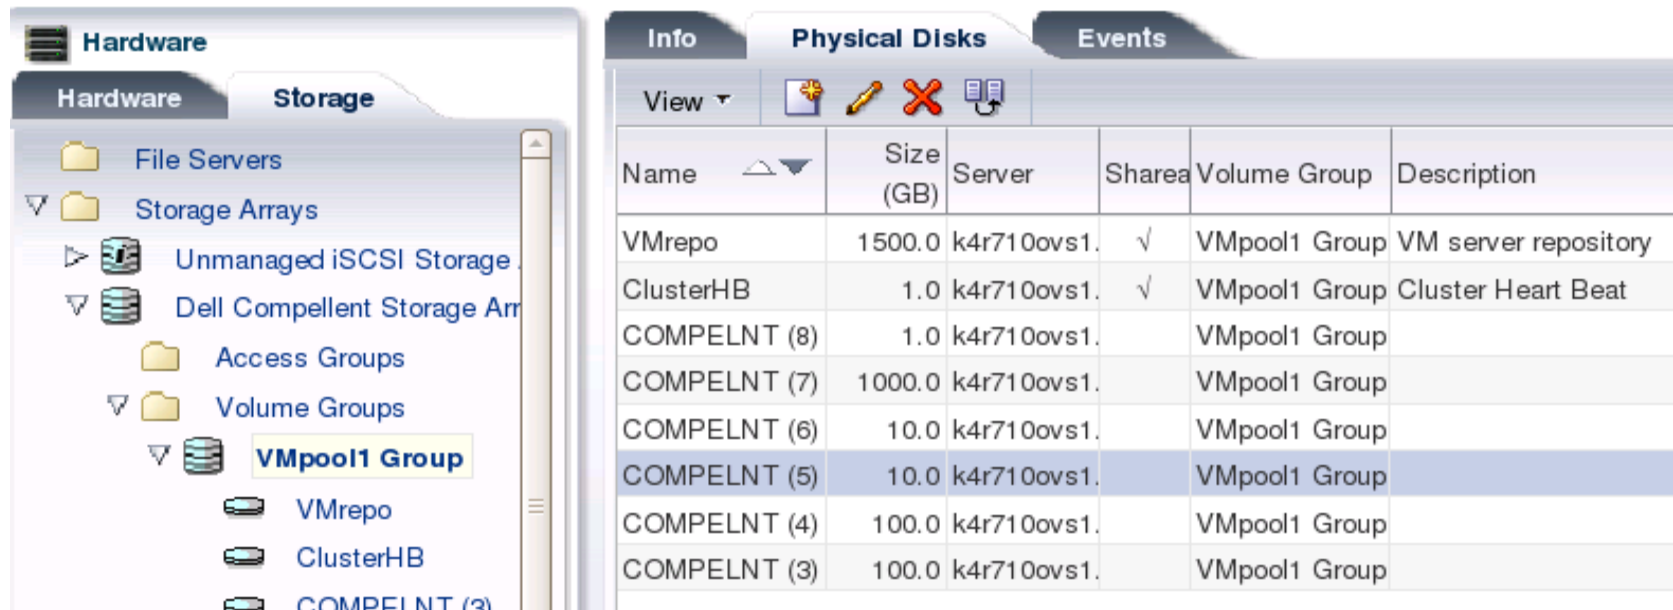
Port Name	Address T	IP Address	Mask	Port Status	Bonding
k4r710ovs1.dblab.com Bond Port (3)	Static	192.168.8.63	255.255.255.0	Up	Active Backup
k4r710ovs2.dblab.com Bond Port (3)	Static	192.168.8.65	255.255.255.0	Up	Active Backup

Below the table, the 'Ethernet Ports' section shows the status of the individual ports:

Port Name	Port Statu
k4r710ovs1.dblab.com Port (3)	Up
k4r710ovs1.dblab.com Port (4)	Up

# Establishing Oracle VM Infrastructure

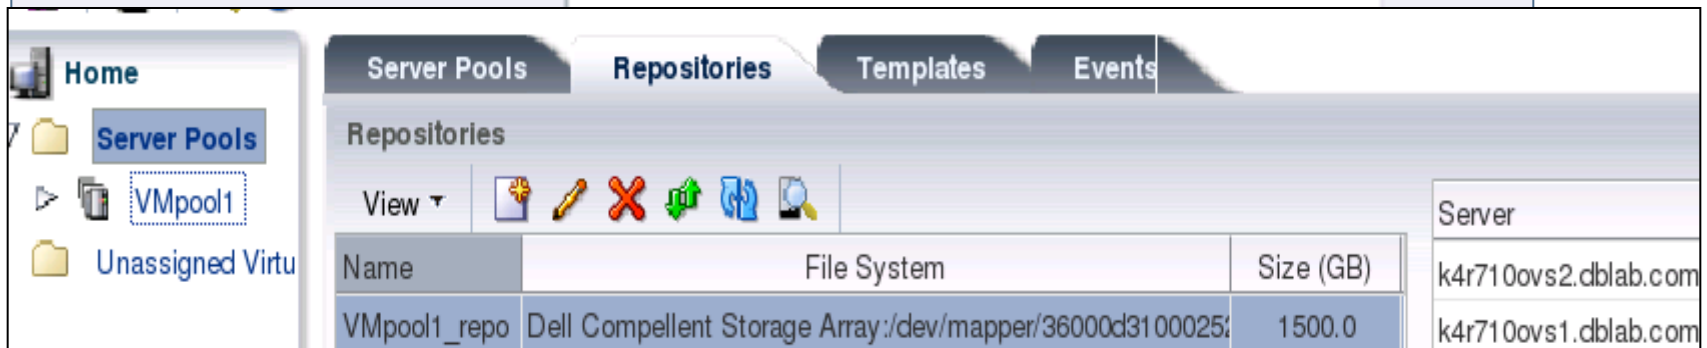
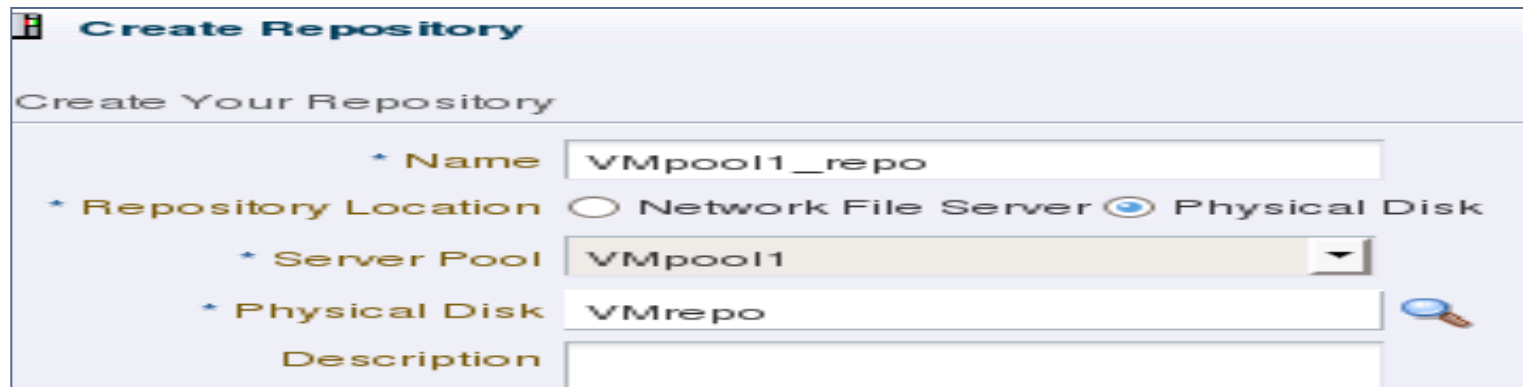
- Storage Configuration
  - Connect to shared storage and setup multipath devices
    - Oracle VM 2.2: manual configuration, eg: iSCSI commands
    - Oracle VM 3.0: Oracle VM manager connects to the storage through a set of plug-ins (storage connect)
      - Oracle generic iSCSI plug-in and SCSI plug-in



Name	Size (GB)	Server	Share	Volume Group	Description
VMrepo	1500.0	k4r710ovs1.	√	VMpool1 Group	VM server repository
ClusterHB	1.0	k4r710ovs1.	√	VMpool1 Group	Cluster Heart Beat
COMPELNT (8)	1.0	k4r710ovs1.		VMpool1 Group	
COMPELNT (7)	1000.0	k4r710ovs1.		VMpool1 Group	
COMPELNT (6)	10.0	k4r710ovs1.		VMpool1 Group	
COMPELNT (5)	10.0	k4r710ovs1.		VMpool1 Group	
COMPELNT (4)	100.0	k4r710ovs1.		VMpool1 Group	
COMPELNT (3)	100.0	k4r710ovs1.		VMpool1 Group	

# Establishing Oracle VM Infrastructure

- Create VM Repository (Storage Repository in OVM3.0)
  - Create OCFS2 cluster file system on the storage volume  
Oracle VM 2.2: Manual step. Oracle VM3.0: Automated by OVM
  - Create VM repository:  
Oracle VM 2.2: repos.py Oracle VM 3.0: VM Manager GUI:

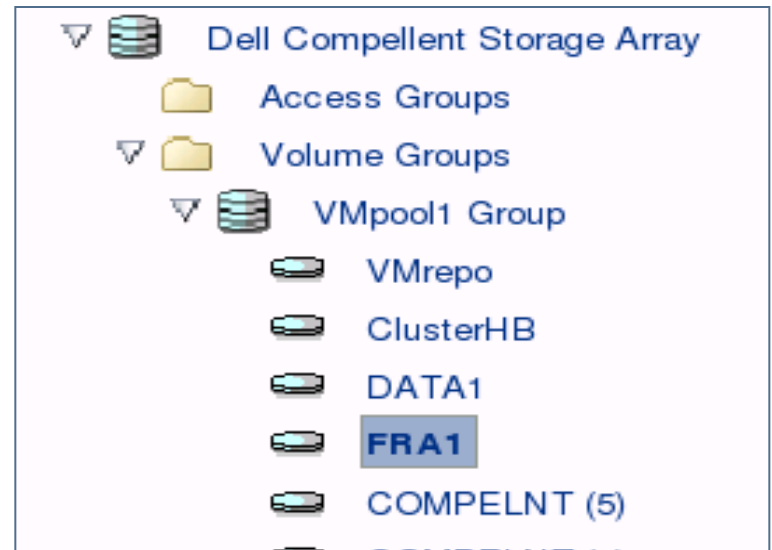
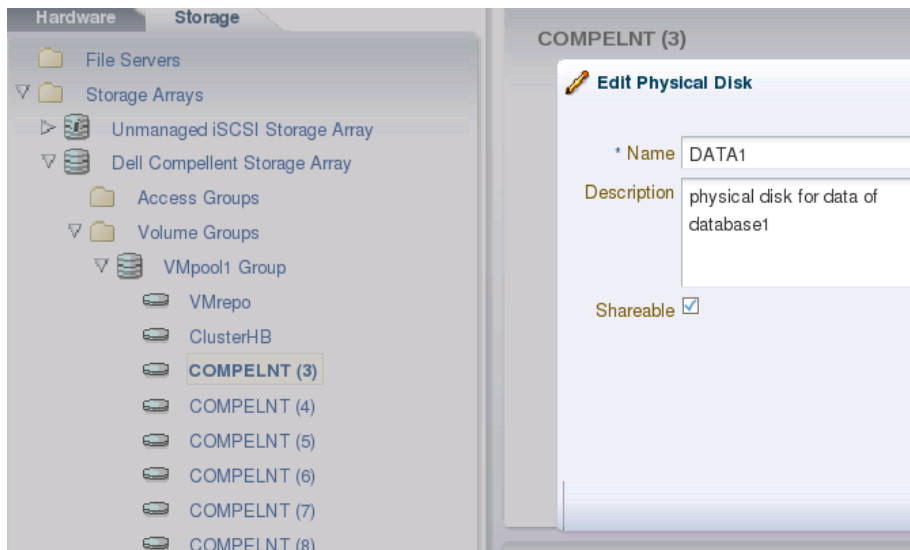


# Establishing Oracle VM Infrastructure

- Storage Repositories structure

```
[root@k4r710ovs2 0004fb00000300002ca24cefa9bc29e5]# pwd
/OVS/Repositories/0004fb00000300002ca24cefa9bc29e5
[root@k4r710ovs2 0004fb00000300002ca24cefa9bc29e5]# ls -l
total 0
drwx----- 2 root root 3896 Sep 11 15:52 Assemblies
drwx----- 2 root root 3896 Sep 11 15:52 ISOs
drwxr-xr-x 2 root root 3896 Sep 11 15:52 lost+found
drwx----- 2 root root 3896 Sep 11 15:52 Templates
drwx----- 2 root root 3896 Sep 11 15:52 VirtualDisks
drwx----- 2 root root 3896 Sep 11 15:52 VirtualMachines
```

- Prepare the physical disks for VMs:



# Virtual Machine(VM) creation-OVM3.0

- Create VMs by cloning VM templates

**Virtual Network Interface Card Manager**

Virtual Network Interface Card Manager

**MAC Selector**

Specify a unique starting MAC address or collisions may occur in applications. Press the Next button to see what starting range

Specify an Initial MAC Address:

00 : 21 : f6 :  :  :

Number of MACs to generate:

---

**Manage MAC Addresses**

View

MAC Address	Ethernet Network	Virtual Machine
00:21:f6:00:00:00	public	VM1_oe15u6
00:21:f6:00:00:01	private	VM1_oe15u6
00:21:f6:00:00:02	public	VM2_OEL5U6
00:21:f6:00:00:03	private	VM2_OEL5U6
00:21:f6:00:00:04		
00:21:f6:00:00:05		

**Clone Virtual Machine or Template**

Select Clone Type  Simple Clone  Advanced Clone

Simple Clone: V26302-01\_oe15u6\_pv\_64.zip

---

Target Clone Type  Template  Virtual Machine

\* Name

**Edit Virtual Machine**

**Networking Options**

Edit Network Configuration

Network Adapter (Vnic)	Selected Value(s)
	<input type="checkbox"/> 00:21:f6:00:02-00:21:f6:00:02-public
	<input type="checkbox"/> 00:21:f6:00:03-00:21:f6:00:03-private

# Virtual Machine Creation- OVM 3.0

Add the shared physical disks and local disk to VM:

Edit Storage Configuration

ISOs	Virtual Disks	Physical Disks
Name	Size(GB)	Storage Array
<input type="checkbox"/> COMPELNT (7)	1000.0	Dell Compellent S
<input type="checkbox"/> COMPELNT (6)	10.0	Dell Compellent S
<input type="checkbox"/> COMPELNT (5)	10.0	Dell Compellent S
<input type="checkbox"/> COMPELNT (8)	1.0	Dell Compellent S
<input type="checkbox"/> FRA1	100.0	Dell Compellent S
<input type="checkbox"/> DATA1	100.0	Dell Compellent S

ISOs	Virtual Disks	Physical Disks
Name	Size(GB)	Shareable
<input type="checkbox"/> 0004fb000012000052ef31ecb14fcc60.img	12.0	No
<input checked="" type="checkbox"/> localdisk_vm2	10.0	No
<input type="checkbox"/> 0004fb000012000083da8813d56100fb.img	12.0	No
<input checked="" type="checkbox"/> 0004fb000012000012fabc0e06fd11cf.img	12.0	No

**VM2\_OEL5U6**  
Assigned Virtual Machines

**Networks & Storage**

Vnic	Ethernet Network	Storage	Repository
00:21:f6:00:00:02	public	0004fb000012000012fabc0e06fd11cf.img	VMpool1_repo
00:21:f6:00:00:03	private	FRA1	
		DATA1	
		localdisk_vm2	VMpool1_repo



# Virtual Machine Creation- OVM 3.0

```
VNC: VM2_OEL5U6
UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
RX packets:261 errors:0 dropped:0 overruns:0 frame:0
TX packets:59 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:17508 (17.0 KiB)  TX bytes:8067 (7.8 KiB)
Interrupt:13

eth1  Link encap:Ethernet  HWaddr 00:21:F6:00:00:03
      inet addr:192.168.9.50  Bcast:192.168.9.255  Mask:255.255.255.0
      inet6 addr: fe80::221:f6ff:fe00:3/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
      RX packets:49 errors:0 dropped:0 overruns:0 frame:0
      TX packets:46 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:2536 (2.4 KiB)  TX bytes:6580 (6.4 KiB)
      Interrupt:12

lo    Link encap:Local Loopback
      inet addr:127.0.0.1  Mask:255.0.0.0
      inet6 addr: ::1/128 Scope:Host
      UP LOOPBACK RUNNING  MTU:16436  Metric:1
      RX packets:8 errors:0 dropped:0 overruns:0 frame:0
      TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:0
      RX bytes:560 (560.0 b)  TX bytes:560 (560.0 b)

[root@v2oel5u6 ~]# more /proc/partitions
major minor #blocks name
202      0 12582912 xvda
202      1  104391  xvda1
202      2 10040625 xvda2
202      3  2433847 xvda3
202     16 104857600 xvdb
202     32 104857600 xvdc
202     48  10485760 xvdd
[root@v2oel5u6 ~]#
```