



COLLABORATE14

TECHNOLOGY AND APPLICATIONS FORUM
FOR THE ORACLE COMMUNITY

#C14LV

High Availability DB on the Cheap using Oracle VM, Oracle Linux, ASM, and DataGuard

*Oracle Virtualization for High
Availability and Cost Containment*

Session # 331

Prepared by:
Bill Petro
Database Architect



REMINDER

Check in on the
COLLABORATE mobile app

Bill Petro - Database Architect

wpetro@gmail.com

- Over 20 years working with Database Server Technology
 - Mostly Oracle including versions 5 through 12
 - Linux, Solaris, Windows, UNIX
 - VMWare, Oracle VM
- Small and large company experience
 - Environmental Testing, Transportation/Logistics, Insurance, Technology, Pharmaceutical, Banking, and Media
 - UPS, Oracle, Merck & Co, Nielsen Media, American Express, + off hours consulting
 - Consulting experience working for Oracle spanning industries, including Telecom and Government
- Board member of SOUG



Disclaimer:
Any comments made or implied during this presentation regarding Oracle licensing are solely my opinion. Please consult your Oracle Sales team for actual license terms.

Business Environment

- Small company
- Banking industry related
- Primary business is moving money between bank accounts for client companies
 - No calculation errors acceptable
 - Near real-time bank transfers
 - Data loss is unacceptable
- Bank transfer data accepted 24 x 7 x 7
 - Bi-monthly scheduled downtime available 9 PM Friday – 9 AM Saturday only



Business Environment (Continued)

- Any unscheduled downtime reduces confidence of clients and banks
 - Cooperative banks can be hard to come by
 - Limited banking hours for fund transfers
 - Unscheduled downtime greater than 20 minutes jeopardizes batch deliverables
 - Clients can be fined by state governments if payments are not processed same day
- Database and application optimal performance less important than availability and consistency
 - Most customer reports finish in seconds
- OLTP and Batch
 - Analytics and Warehouse separate



Business Requirement Targets

- Full hardware redundancy at HQ (almost)
- Full hardware redundancy at DR (almost)
- Near real-time data synchronization to DR
- Maintain Oracle Database CPU based current license level with hardware refresh
- Limit downtime to less than 20 minutes from hardware failure
- Upgrade to supported Oracle version
- Migrate database servers from Windows to Linux



Database Architecture – Original

- Oracle 9i SE with 2nd database in Recovery mode and 3rd database in DR
 - Archive log shipping and apply via Visual Basic scripts with compression
 - Daily manual log manipulation required
- Limited monitoring of Database availability and performance
- Window NT Server OS – 32 Bit
- Hardware Redundancy
 - Extra server similar to Production – manual failover
- DR Server same as Production
 - limited bandwidth DB logs up to 1.5 days lag



Database Architecture – Pre-VM

- Oracle 10g EE - DataGuard
- Oracle 10g OEM – VMware VM -> Oracle VM 2.2
- Tuning Pack
- Diagnostics Pack
- Window 2003 Servers – 64Bit
- Internal hardware redundancy
- Test Server same as Production – manual failover
 - Production on SAN
 - Test on local disk
- DR Server same as Production
 - DataGuard Physical Standby database
 - Limited bandwidth DB logs up to 3 hours lag, usually 20 minutes



Virtualization Decision Factors

- OEM in Windows VMWare VM
 - Worked well until it didn't
- OEM in Oracle VM Server (2.x) VM
 - Fully supported and software was “free”
- Oracle database CPU licensing (named user for Dev/Test)
 - No new commodity fault tolerant hardware with 4 cores
- Oracle VM Server makes hardware effectively “smaller”
 - CPU Pinning
- Oracle VM Server worked well for OEM
- Oracle 10g support ended
- Hardware is near end of life



Why not VMWare?

- VMWare Shop
 - 3 Way Production cluster and 2 Way Test
 - All internal infrastructure was virtualized
 - Email, accounting, custom applications
 - All custom client facing web applications virtualized
- Original OEM installation ran under VMWare
 - It broke once and Oracle Support said call us back when on a supported platform
- Database \ Web Apps \ email resource contention
 - Separate databases
- VMWare licenses = \$ x 6 servers
- CPU pinning enabling commodity hardware with internal redundancy – not with VMWare




Oracle VM and VMWare Costs

(<http://www.oracle.com/us/media/calculator/vm/vm-home-2132015.html>)

- 1 Calculate your Hypervisor and Management Costs:
Select Guest OS you plan to deploy.
For Oracle VM
For VMware
- 2 Enter your system configuration.
Number of servers with 1 or 2 processors sockets
Number of servers with more than 2 processors sockets
Total number of processor sockets across all physical servers
- 3 Enter total number of virtual guests.
- 4 Choose a support term.

VMware vSphere with Red Hat Enterprise Linux costs between **\$35k** and **\$72k** more than Oracle VM and Oracle Linux.

That's **5x** more!



- ☐ OS license and support cost
- ☐ VMware Hypervisor license and support cost
- ☐ VMware Management license and support cost
- ☐ Oracle VM Cost



COLLABORATE14

TECHNOLOGY AND APPLICATIONS FORUM
FOR THE ORACLE COMMUNITY

Virtualized Database Architecture Target

- Oracle 11g EE, ASM - Clusterware - Dataguard
- Oracle 11g OEM – Oracle VM Server VM
 - Tuning Pack
 - Diagnostics Pack
- Oracle Enterprise Linux
- Oracle VM Server 3.1
- Two clustered Production Servers
 - Shared SAN
- DR Servers same as Production
 - Separate SAN
 - Limited archive log transfer bandwidth



Migration Path

- All at once
 - Test custom applications once
 - Limit migration downtimes
- All new hardware
 - 2 socket dual core (1x) -> 2 socket 4 core (2x)
 - Xiotech -> NetApp SAN
- Oracle 10 -> 11
 - Database Files - OS -> ASM
 - Clusterware Oracle restart
- Windows -> Linux
- Bare-metal single server -> clustered VM Servers
- HQ and DR in one weekend
 - Test Running on Production hardware— moved back to Test after



OVM Additional Benefits

- High Availability VM architecture
 - VM Live Migration with running databases
 - Think about it though
- Standardized Linux template installation
 - Portable across differing hardware
 - Quick deployment
- Minimal performance impact
 - Use physical devices, not virtual files for databases!
- CPU Pinning!

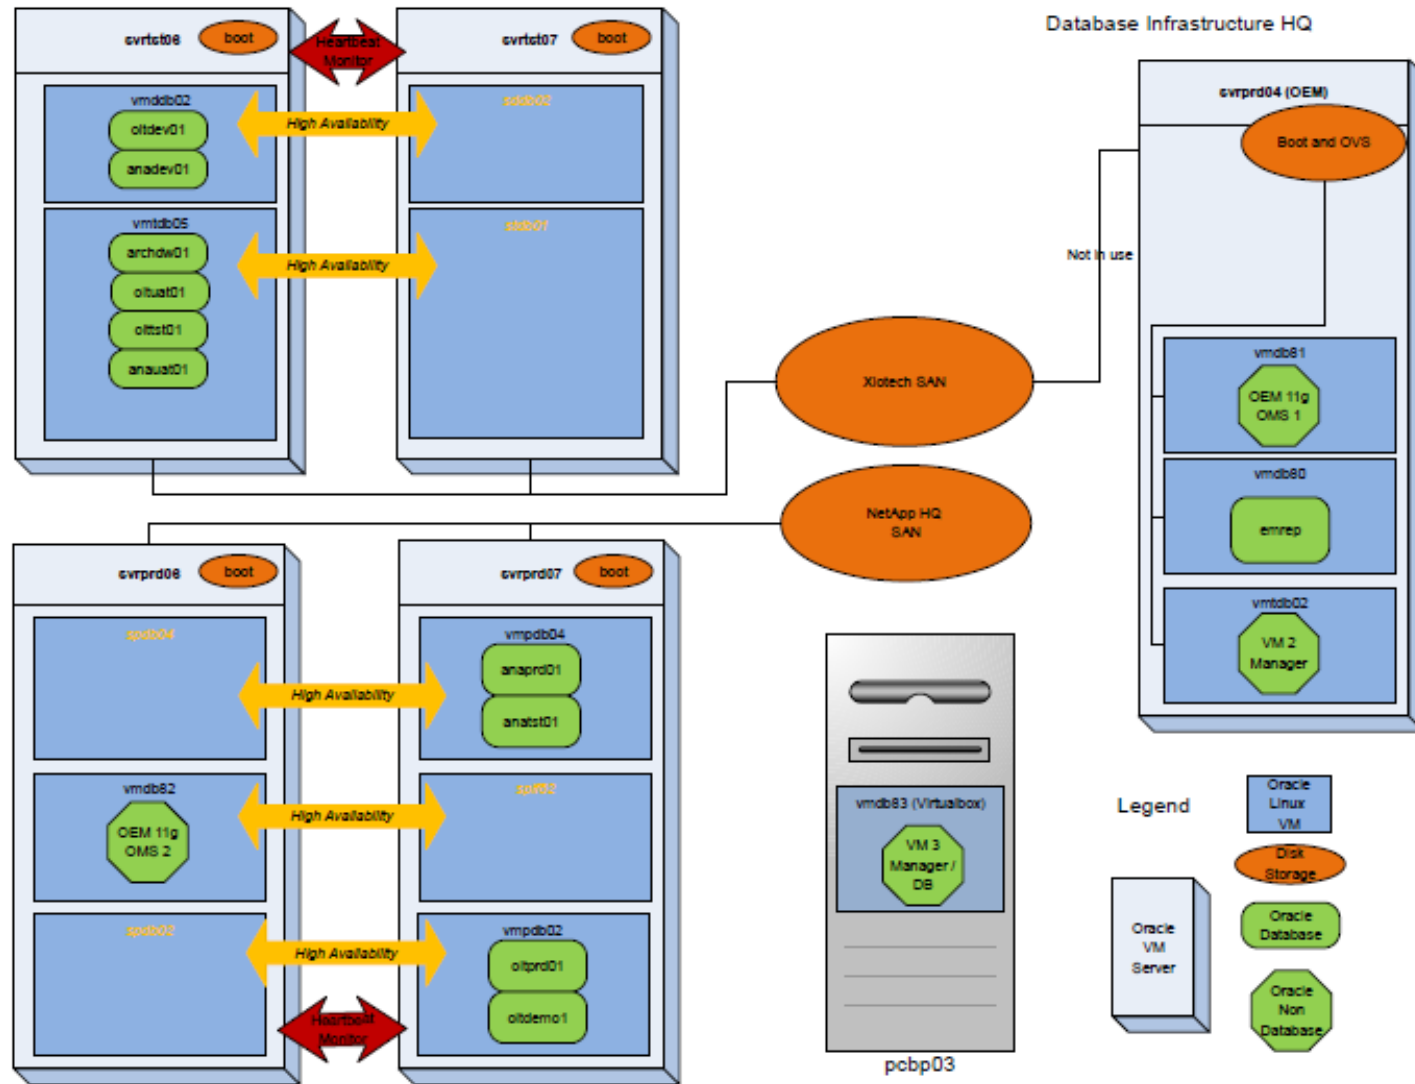


Target Database Architecture Overview

- OEM - 2 OMS
- Dev
- Test
 - Warehouse
- Production
 - OLTP
 - Analytics
- DR Site
 - OLTP
 - Analytics
 - OEM Installation (passive)



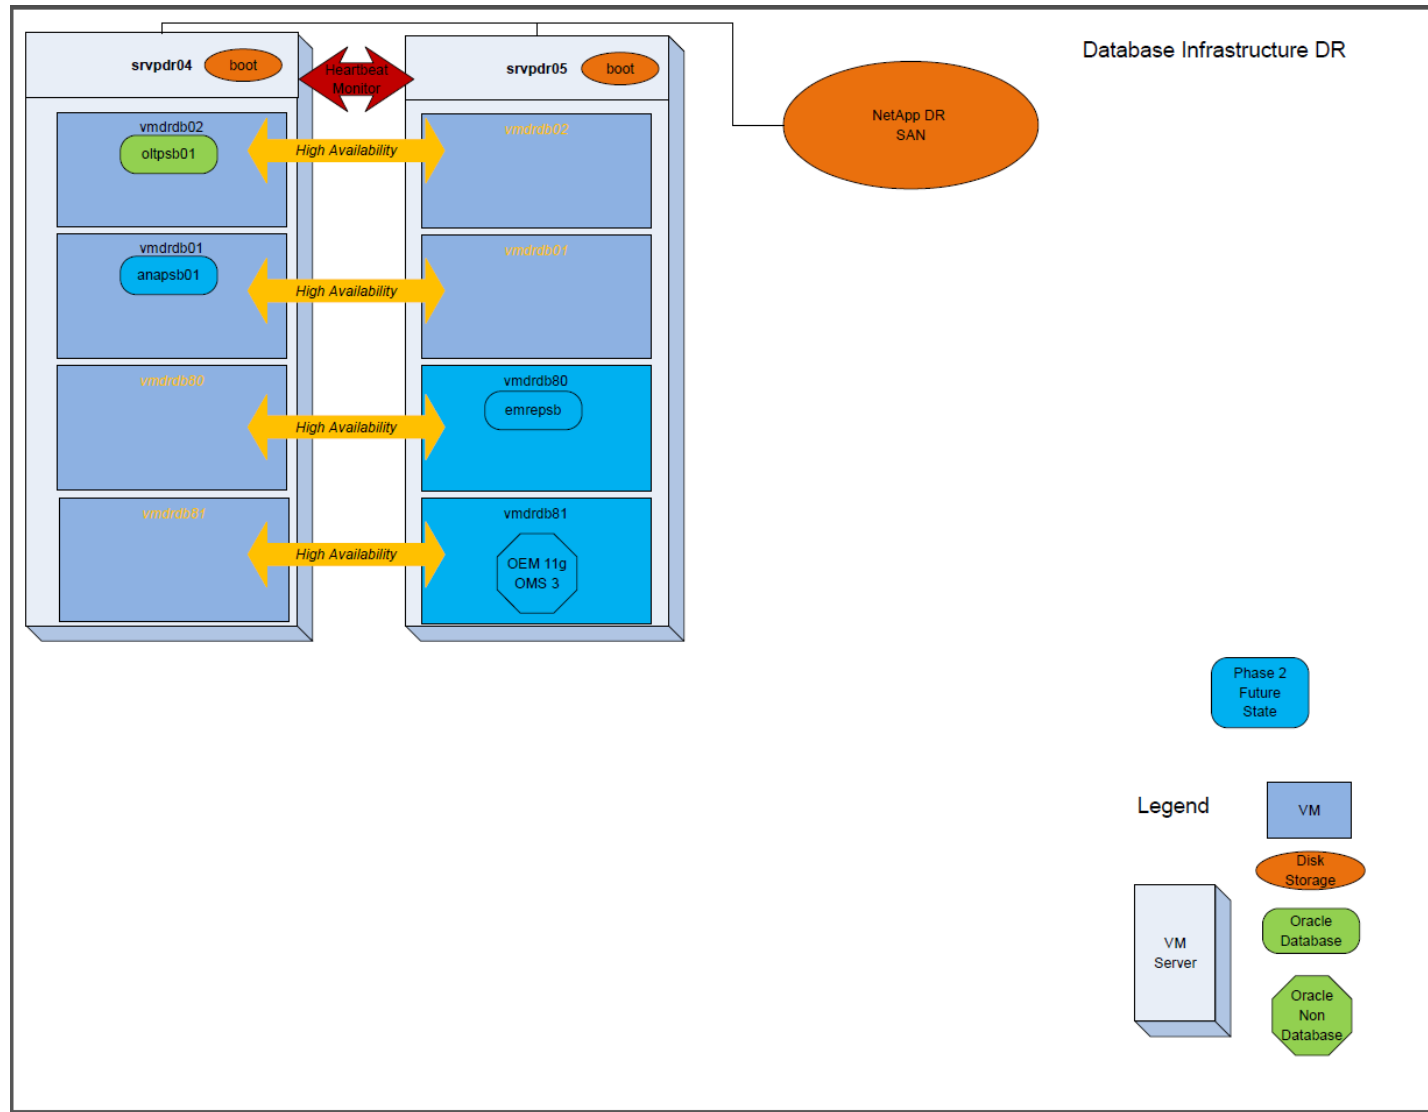
HQ Database Architecture



COLLABORATE14

TECHNOLOGY AND APPLICATIONS FORUM
FOR THE ORACLE COMMUNITY

DR Site Database Architecture



COLLABORATE14

TECHNOLOGY AND APPLICATIONS FORUM
FOR THE ORACLE COMMUNITY

HA VM Server Minimum Requirements

- 2 identical servers
 - Memory to suit intended purpose
 - Internal boot storage
 - Hard Drive or USB Drive
- Shared Storage – SAN
 - iSCSI
 - Fiber Channel – HBA
 - NAS
- Networks
 - Public
 - Storage
 - Private - Heartbeat

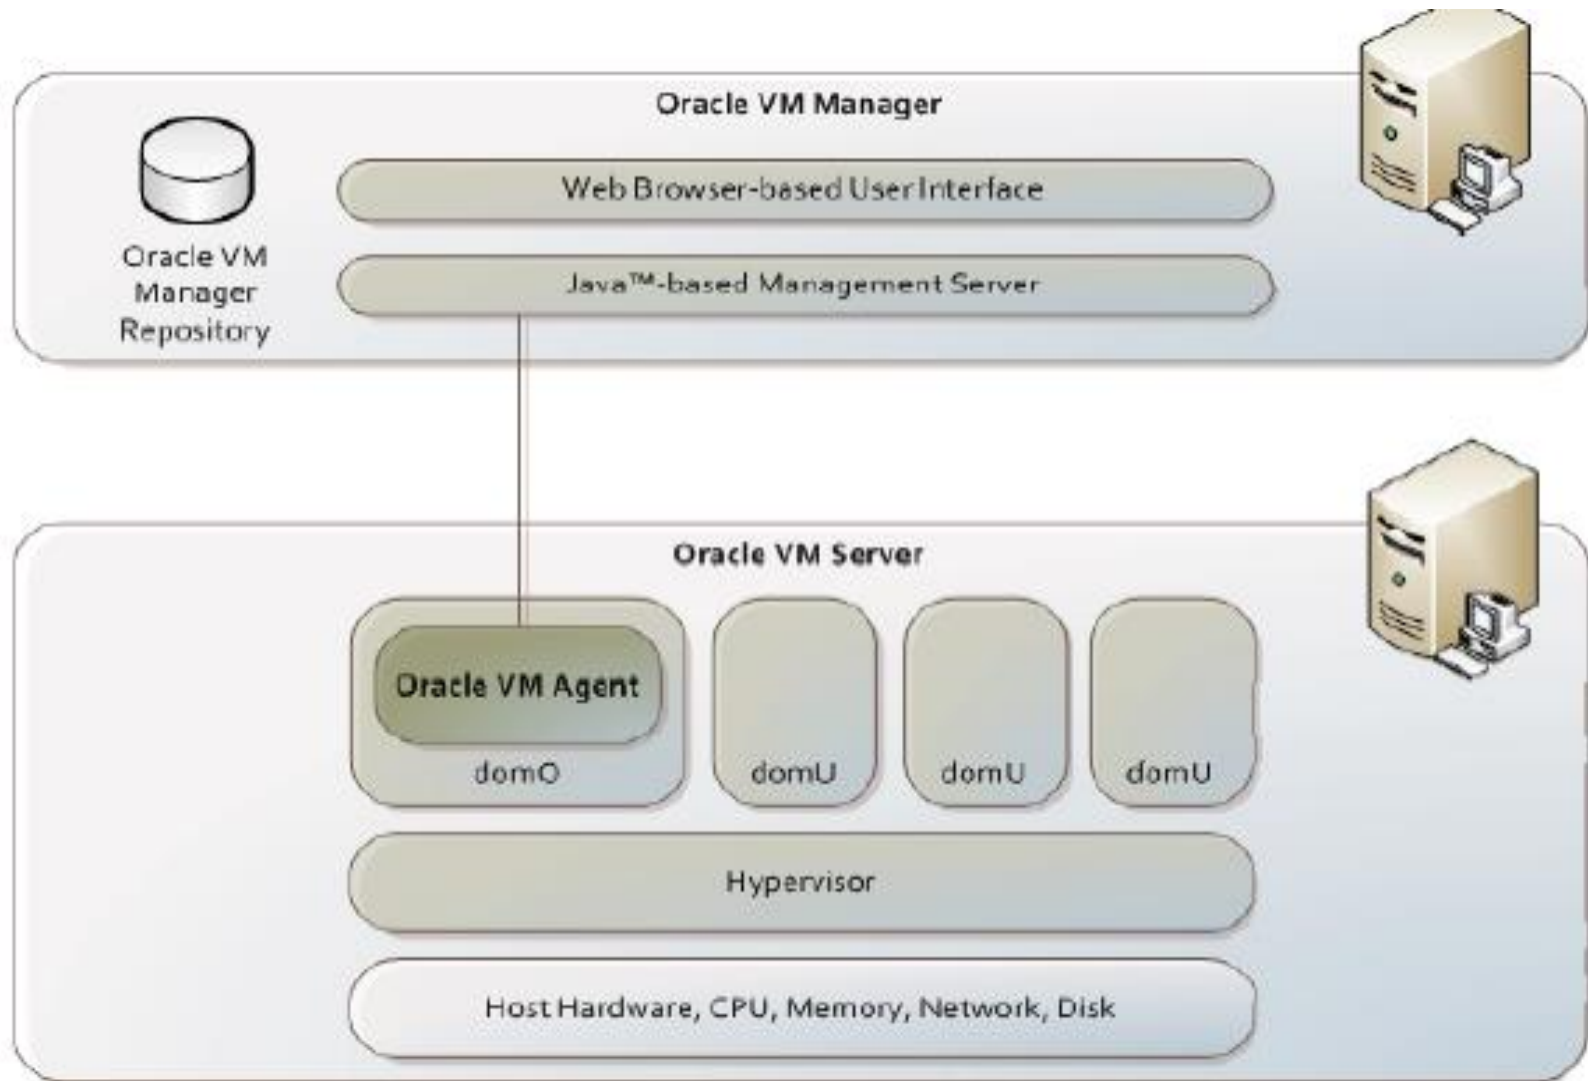


Database Server VM configuration

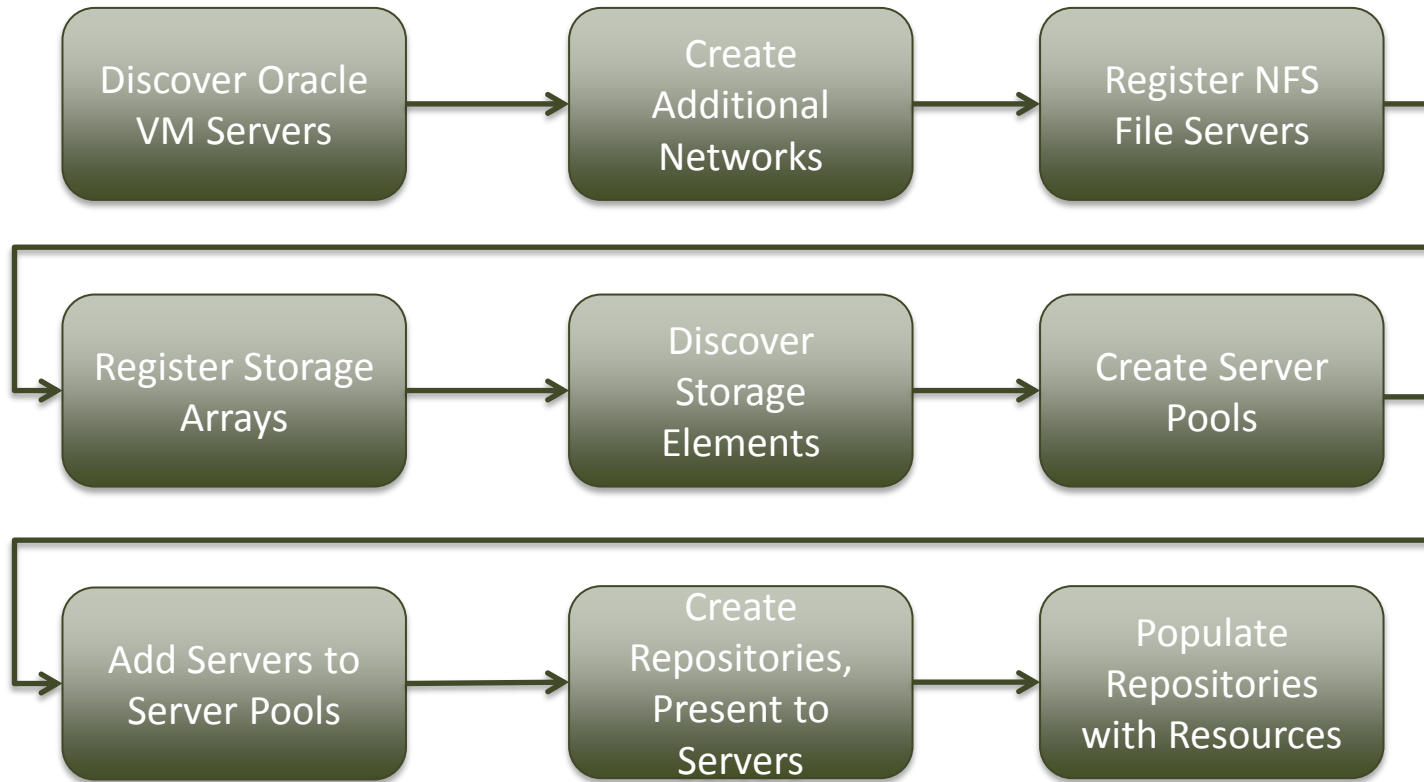
- Oracle Linux 5 release 6 64-Bit
- Oracle 11.2.0.3
 - Clusterware
 - ASM
 - Database
 - **Use “raw” volumes as physical drives, not virtual files for database files. 50%+ performance degradation with virtual files**
- Create “golden image” template of Linux with Oracle software staged for installation and cloning



OVM Architecture Components



OVM Configuration Overview



Oracle Virtualbox

- This is not Oracle VM
 - Requires full underlying OS GUI install
 - No High Availability features
 - High OS overhead impact on performance
- VM technology for the desktop
 - Linux or Windows Host
 - Linux, Windows, Solaris, ... Guests
- Host OS 64 Bit recommended, but not required
- CPU and BIOS virtualization features recommended



Demo Environment

OVM Server on Virtualbox

- Desktop or Laptop
 - 150 GB storage
 - 16 GB memory (12 will do)
 - Processor/BIOS with virtualization support
- 64 Bit OS Linux or Windows
- Oracle Virtualbox Software and guest additions
- VM Templates
 - OVM Manager
 - OVM Server
 - 12g RAC
 - Cluster Configuration Utility

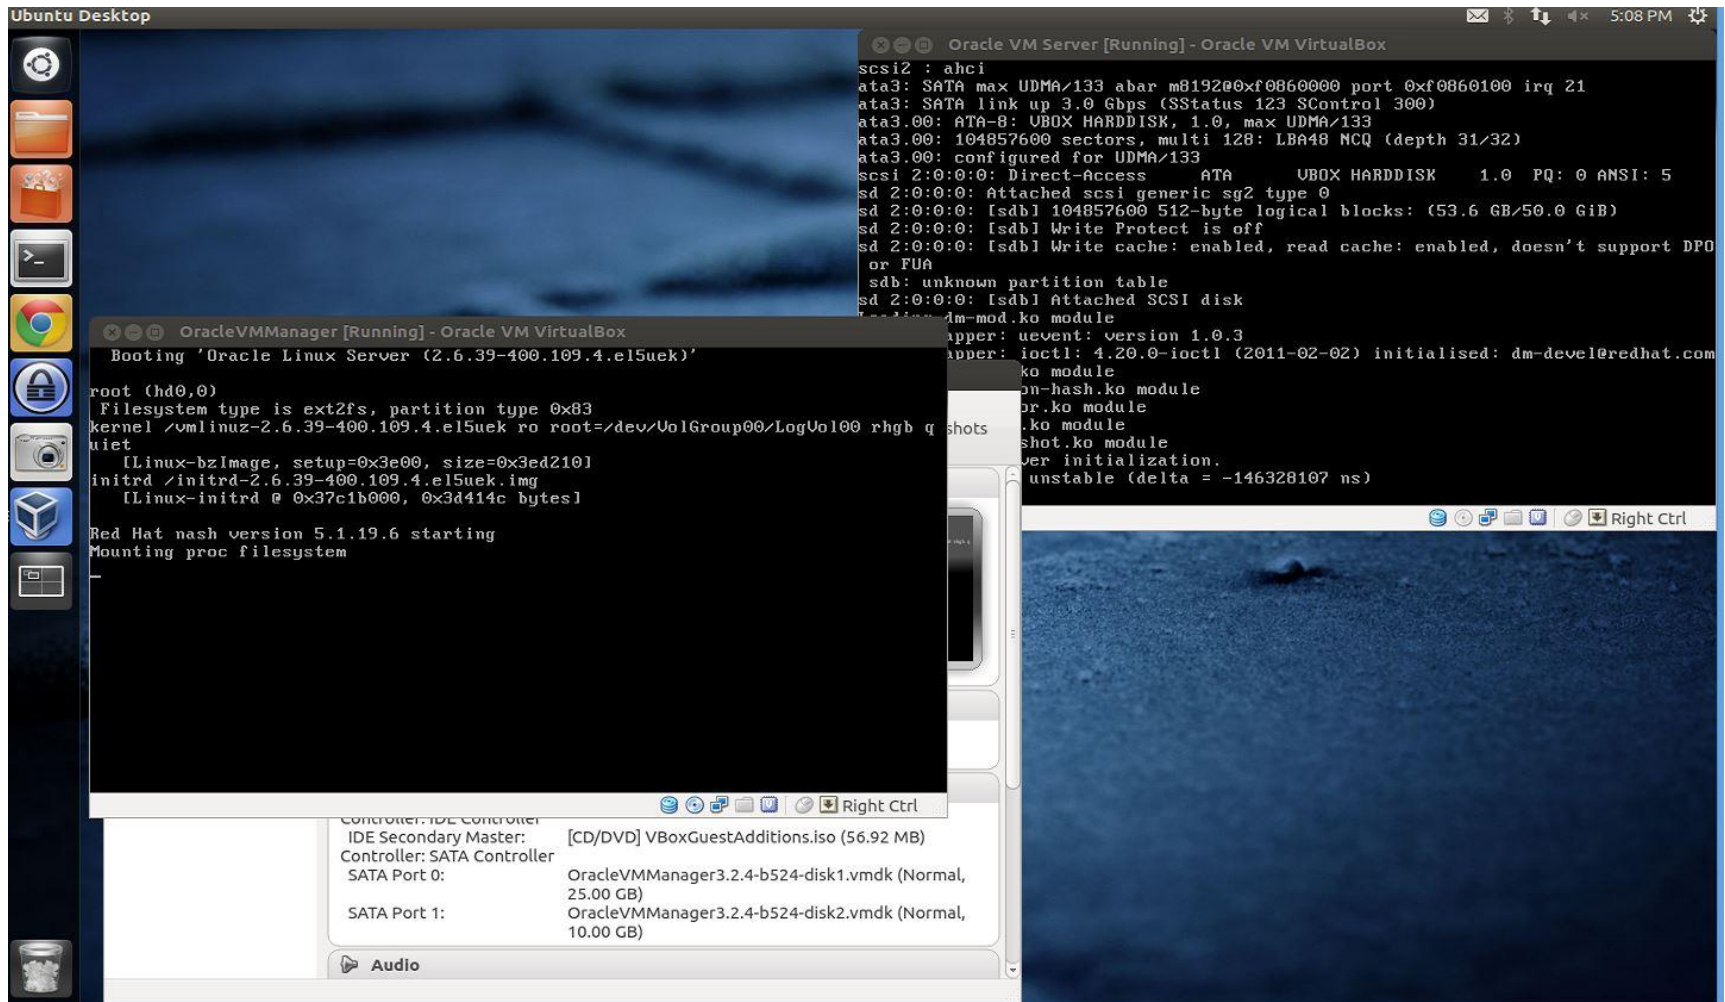


Demo OVM environment configure

- VM Manager
- Discover Servers
- Configure VM Network
- Create Server Pool
- Create Storage Repository
- Import templates with Oracle RAC
- Create VMs from templates
- Create shared storage for ASM
- Use OVM Deploycluster Tool to configure Oracle RAC Nodes




Startup of VM Manager and VM Server Virtualbox



VM Manager and VM Server Virtualbox

#C14LV



COLLABORATE14

TECHNOLOGY AND APPLICATIONS FORUM
FOR THE ORACLE COMMUNITY

Four VMs running Virtualbox

The screenshot displays the Oracle VM VirtualBox interface. The main window shows the 'Oracle VM Manager' with a sidebar on the left containing icons for applications, places, and system. The central pane is titled 'Oracle VM Home - Mozilla Firefox' and shows the 'ORACLE VM Manager' interface. The 'Servers and VMs' tab is active, displaying a table of running VMs.

The 'Oracle VM Server [Running] - Oracle VM VirtualBox' window is open, showing the 'xentop' command output. The output indicates 5 domains: 2 running, 0 blocked, 0 paused, 0 crashed, 0 dying, 0 shutdown. The memory usage is 7863868k total, 6987220k used, and 876648k free. The CPU usage is 2 @ 2485MHz.

NAME	STATE	CPU(sec)	CPU(%)	MEM(k)	MEM(%)	MAXMEM(k)	MAXMEM(%)	UCPU	SS
0004fb0000	-----	3	30.5	1572864	20.0	1572864	20.0	1	
0	0	0	2	1459	45	28198	512		
0004fb0000	-----	25	32.1	1572864	20.0	1572864	20.0	1	
0	31	1	4	7272	282	181326	5202		
0004fb0000	-----	43	43.7	1572864	20.0	1572864	20.0	1	
0	29	9	4	10671	669	391502	12122		
0004fb0000	-----r	12	43.1	1572864	20.0	1572864	20.0	1	
0	0	0	4	2951	51	51766	576		
Domain-0	-----r	707	58.8	606208	7.7	606208	7.7	2	
0	0	0	0	0	0	0	0		

The 'Oracle VM Manager' window shows a table of running VMs:

Name	Status	Tag(s)	Event Severity	Server	Max. Memory (MB)	Memory (MB)
rac.0	Running		Normal	ovm-srv.oow.com	1536	1536
rac.1	Running		Normal	ovm-srv.oow.com	1536	1536
rac.2	Running		Normal	ovm-srv.oow.com	1536	1536
rac.3	Running		Normal	ovm-srv.oow.com	1536	1536



OVM Server Virtual Disk Repository

OracleVMManager [Running] - Oracle VM VirtualBox

Applications Places System 2:30 PM

Oracle VM Home - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Oracle VM Home

https://localhost:7002/ovm/console/faces/resource/resourceView.jspx?_afrcWindowMode=0&_afrcLoop=1003930622919&_adf.ctrl- Google

ORACLE VM Manager Logged in as: admin

Health Servers and VMs Repositories Networking **Storage** Tools and Resources Jobs

File Servers
SAN Servers
Unmanaged FibreChannel Storage
FibreChannel Volume Group
Unmanaged iSCSI Storage Array
iSCSI Volume Group
Local File Systems

View

Name	Storage Device	Event Severity	Size (GiB)			Location
			Free	Used	Total	
fs_OVMRepo	SATA_VBOX_HARDDISK_	Normal	29.74	20.26	50.00	Repository: OVMRepo



OVM Server Virtual Disks

OracleVMManager [Running] - Oracle VM VirtualBox

Applications Places System 2:28 PM

Oracle VM Home - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Oracle VM Home

https://localhost:7002/ovm/console/faces/resource/resourceView.jspx?_afWindowMode=0&_afLoop=1003930622919&_adf.ctrl- Google

ORACLE VM Manager Logged in as: admin Logout Set

Health Servers and VMs Repositories Networking Storage Tools and Resources Jobs

Show My Repositories
Show All Repositories

Create New Repository...

- VM Templates
- Assemblies
- ISOs
- Virtual Disks

Name	Used (GiB)	Max (GiB)	Shareable	Description
Oracle12101DBRAC_x86_64-xvdb.img	9.93	30.0	No	
Oracle12101DBRAC_x86_64-xvdb.img (2)	9.93	30.0	No	
Oracle12101DBRAC_x86_64-xvdb.img (3)	9.93	30.0	No	
Oracle12101DBRAC_x86_64-xvdb.img (4)	9.93	30.0	No	
Oracle12101DBRAC_x86_64-xvdb.img (5)	9.93	30.0	No	
System.img	1.66	12.0	No	

Rows Selected 1



OVM Server Virtual and Physical Disks – Single VM View

Oracle VM Home - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Oracle VM Home

https://localhost:7002/ovm/console/faces/resource/resourceView.jspx?_afrcWindowMode=0&_afrcLoop=5943918767800&_adf.ctrl-state=ryr3e: Google

ORACLE VM Manager

Logged in as: admin Logout Settings Help

Health Servers and VMs Repositories Networking Storage Tools and Resources Jobs

Server Pools
mypool
ovm-srv.oow.com
Unassigned Servers
Unassigned Virtual Machines

View Perspective: Virtual Machines

Name	Status	Tag(s)	Event Severity	Server	Max. Memory (MB)	Memory (MB)	Max. Processors	Processors	Keymap	Operating System
rac.0	Stopped		Normal	ovm-srv.oow.com	1536	1536	1	1	en-us	None

Configuration Networks Disks

Slot	Disk Type	Content	Size (GiB)	Repository	Location
0	Virtual Disk	System.img (2)	12.0	OVMRepo	/dev/mapper/SATA_VBOX_HARDDISK_VB79
1	Virtual Disk	Oracle12101DBRAC_x86_64-xvdb.img (2)	30.0	OVMRepo	/dev/mapper/SATA_VBOX_HARDDISK_VB79
2	Virtual Disk	racasm1	5.0	OVMRepo	/dev/mapper/SATA_VBOX_HARDDISK_VB79
3	Physical Disk	SATA_VBOX_HARDDISK_VBb1499d16-a263a9d0_	2.0		Generic Local Storage Array @ ovm-srv.oow.c



CPU Pinning OVM 3.1

- Binds a VM to physical CPUs (actually threads)
 - Supports CPU based licensing
 - Isolates VMs from each another
- Must be configured for each VM
 - Use the Oracle VM Utilities (ovm_vmcontrol) to set hard partitioning
 - Must be downloaded separately Oracle Support ID 13602094
 - Or modify vm.cfg manually on the VM Server
 - `vi /<repository file system path>/VirtualMachines/<virtual machine ID>/vm.cfg`
- Disabled during VM migrations between hosts
 - Must be re-enabled via the VM Utilities



VMServer CPU Detail

Oracle VM Server [Running] - Oracle VM VirtualBox

```
nr_nodes : 1
cores_per_socket : 2
threads_per_core : 1
cpu_mhz : 2482
hw_caps : 178bf bff :28100800:00000000:00000140:00000201:00000000:0
00000001:00000000
virt_caps :
total_memory : 7679
free_memory : 2392
free_cpus : 0
xen_major : 4
xen_minor : 1
xen_extra : .30VM
xen_caps : xen-3.0-x86_64 xen-3.0-x86_32p
xen_scheduler : credit
xen_pagesize : 4096
platform_params : virt_start=0xffff800000000000
xen_changeset : unavailable
xen_commandline : dom0_mem=592M
cc_compiler : gcc version 4.1.2 20080704 (Red Hat 4.1.2-48)
cc_compile_by : mockbuild
cc_compile_domain : us.oracle.com
cc_compile_date : Wed May 29 12:16:42 PDT 2013
xend_config_format : 4
[root@ovm-srv ~]# _
```



VM Server CPU Topology

```
Oracle VM Server [Running] - Oracle VM VirtualBox
[root@ovm-srv ~]# xenpm get-cpu-topology
CPU      core      socket  node
CPU0      0          0        0
CPU1      1          0        0
[root@ovm-srv ~]# xm vcpu-list
Name                                           ID    VCPU    CPU State    Time(s) CPU Affinity
0004fb00000600002bc62c3cc7a8535b             5      0      0  -b-      148.9 any cpu
0004fb000006000035d84c39c5ce807e             2      0      1  r--      803.4 1
0004fb0000060000479b642099243ca8             4      0      0  ---      192.6 0
Domain-0                                       0      0      1  r--      257.1 any cpu
Domain-0                                       0      1      1  -b-      197.7 any cpu
[root@ovm-srv ~]#
[root@ovm-srv ~]#
[root@ovm-srv ~]#
```



VM Server CPU Pinned

```
Oracle VM Server [Running] - Oracle VM VirtualBox
[root@ovm-srv ~]# xenpm get-cpu-topology
CPU      core      socket  node
CPU0      0          0        0
CPU1      1          0        0
[root@ovm-srv ~]# xm vcpu-list
Name                                     ID   VCPU   CPU State   Time(s) CPU Affinity
0004fb00000600002bc62c3cc7a8535b       5     0     0  -b-      148.9 any cpu
0004fb000006000035d84c39c5ce807e       2     0     1  r--      803.4 1
0004fb0000060000479b642099243ca8       4     0     0  ---      192.6 0
Domain-0                                0     0     1  r--      257.1 any cpu
Domain-0                                0     1     1  -b-      197.7 any cpu
[root@ovm-srv ~]#
[root@ovm-srv ~]#
[root@ovm-srv ~]#
```



Oracle's VM hardware

- Oracle Database Appliance
 - Separate Dev/Test/Prod
 - Run you Applications with Database
- Oracle Exalytics
 - Separate Dev/Test/Prod
 - Separate Products
 - Separate Versions
- Oracle Virtual Compute Appliance
 - Brand new
 - Presentation from Wednesday at 4:30



Key feature review

- Key features enabling project success
 - CPU Pinning to limit hardware usage for purposes of database licensing
 - Physical disk configuration to support high performance
 - Clustering to guard against hardware failure
 - Excess CPU capacity was available to support OEM primary, secondary, and DR installations
 - VM Manager was actually later run using extra capacity in the Test cluster



Please complete the session evaluation on the mobile app

We appreciate your feedback and insight

Session # 331

This box will have simplified instructions about how to complete the session evaluation online



COLLABORATE14

TECHNOLOGY AND APPLICATIONS FORUM
FOR THE ORACLE COMMUNITY