

Oracle RAC on Vsphere 5

Solution Presentation - Don Sullivan

– Senior Systems Engineer - Database Specialist

sullivanand@vmware.com

Don Sullivan – sullivan@vmware.com

- Oracle Certified Master, Server Products Trainer for Oracle University and consultant with Oracle Advanced Technology Services - 1998-05.
- Oracle SA for Polyserve/HP – 05-10
- Vmware SE DB specialist 2010 – Present

Agenda

- **RAC on VMware – The Vsphere Platform**
 - Why use vSphere 5 as the platform for Oracle RAC
 - How to setup RAC on Vsphere - The installation process
 - Vsphere 5 features that make Vsphere complimentary to RAC
- **Performance capabilities**
- **Customer Success Stories**
- **The VRAC-DBA**
- **Summary**
- **Q/A**

Why RAC on Vmware (ESX/vSphere 5)

- **Uptime**

- Greater uptime than RAC on physical
 - Vsphere HA and RAC are complimentary technologies

- **Provisioning**

- New RAC nodes provisioned in minutes rather than hours

- **Comprehensive Resource Management**

- The VM/Server is created to the exact specifications desired
- The VM/Server resources can be adjusted as requirements change
- Zero Downtime Hardware maintenance

- **The changing role of the DBA – The *VRAC-DBA***

- The modern Oracle DBA has influence over the entire stack
- Database Administrators do not have to re-learn their skill set
- The *VRAC-DBA* occupies the preeminent position in modern IT
 - From the Virtualized Infrastructure to the RAC instances and Apps
 - From Storage through the network Architectures

vSphere High Availability Features

VMware HA

- Detects operating system and hardware failures
- Automatically restarts failed database virtual machine
- Provides a simple and reliable first line of defense for all databases
- Can be used in conjunction with Symantec App HA to provide application aware protections

VMware vMotion

- Enables live migration of database virtual machines from one physical server to another without service interruption
- Can reduce virtual machine planned downtime
- Perform host maintenance any time of the day

VMware DRS

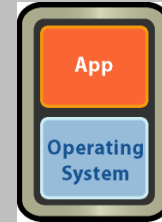
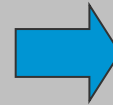
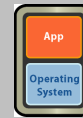
- Monitors state of virtual machine resource usage
- Can automatically and intelligently locate virtual machine
- Directs compute resources where needed
- Maintains database response time and SLAs

Scalability on Demand

Dynamic Scaling on VMware

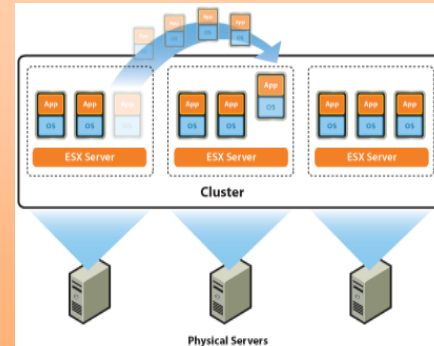
Hot-Add Capacity

1 vCPU
2 GB

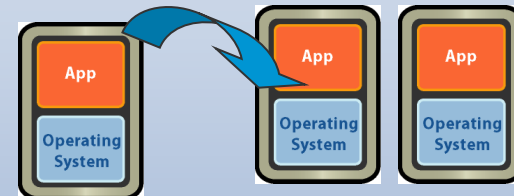


4 vCPU
64 GB

VMotion to More Powerful Host



Provision Additional App Instance in Minutes



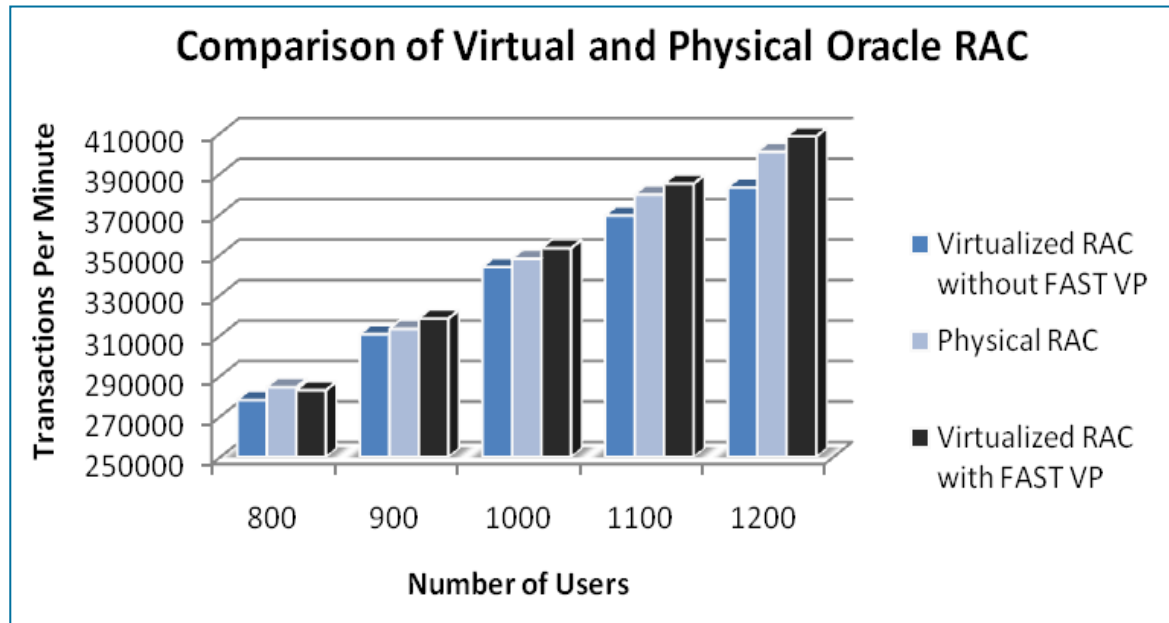
Maximum Scalability and Performance With vSphere 5

Application's Performance Requirements

% of Applications

95% of Apps Require		ESX 1	ESX 2	VMware Inf. 3.0/3.5	VMware vSphere 4	VMware vSphere 5
CPU	1 to 2 CPUs	1 VCPUs	2 VCPUs	4 VCPUs	8 VCPUs	32 VCPUs
Memory	< 4 GB at peak	2 GB per VM	3.6 GB per VM	16/64 GB per VM	256 GB per VM	1,000 GB per VM
Network	<2.4 Mb/s	<.5Gb/s	.9 Gb/s	9 Gb/s	30 Gb/s	>36Gb/s
IOPS	< 10,000	<5,000	7,000	100,000	300,000	1,000,000

RAC on Vsphere - Physical & Virtual Oracle RAC comparison



- Virtual environment delivered performance consistently within 4% of physical environment (without FAST VP)

- Virtual environment processed 2% more transactions after FAST policy applied

Agenda

- **RAC on VMware – The Vsphere Platform**
 - Why use vSphere 5 as the platform for Oracle RAC
 - **How to setup RAC on Vsphere - The installation process**
 - Vsphere 5 features that make Vsphere complimentary to RAC
- **Performance capabilities**
- **Customer Success Stories**
- **The VRAC-DBA**
- **Summary**
- **Q/A**

RAC on Vsphere - History

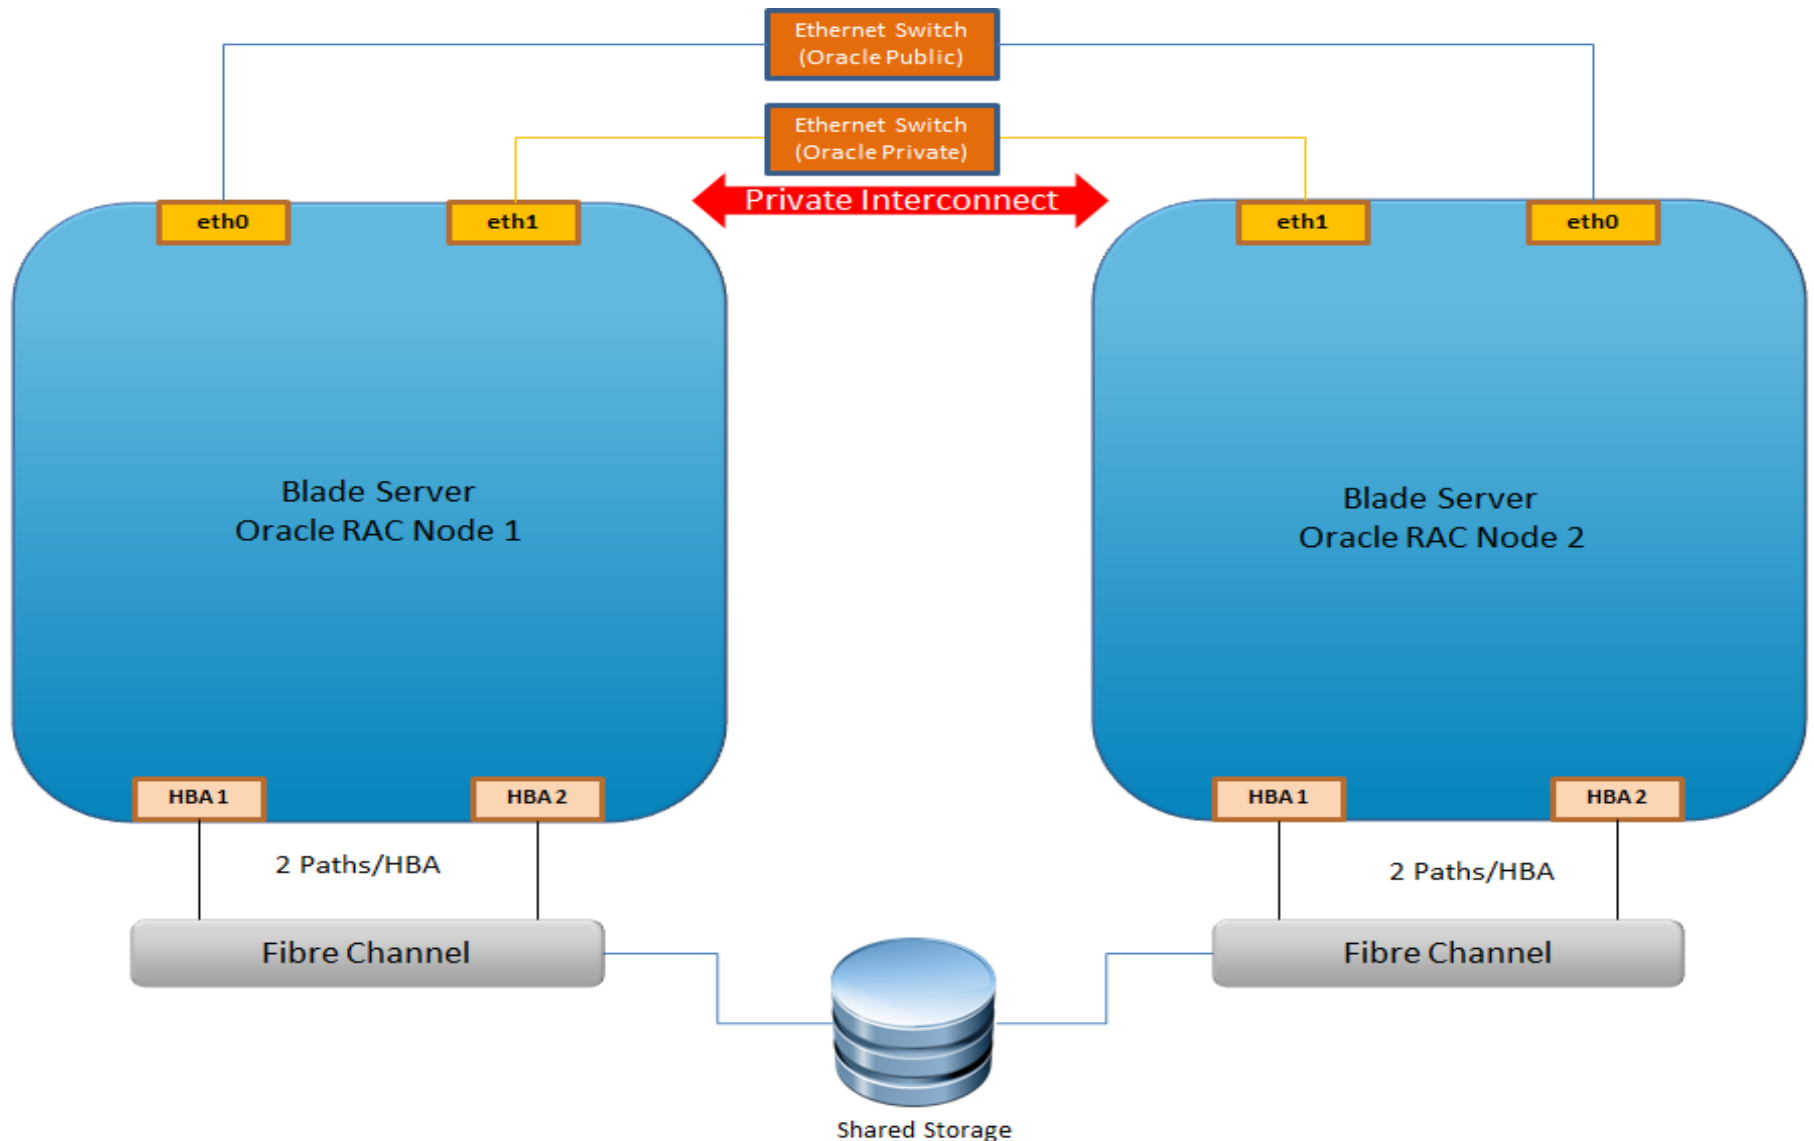
- **Milestones**

- Oracle support statement extension – 11.2.0.2 – Fall 2010
 - MyOracleSupport.com Note 249212.1
 - “For Oracle RAC, Oracle will only accept Service Requests as described in this note on Oracle RAC 11.2.0.2 and later releases.”
- Release of the “Oracle Databases on VMware - RAC Deployment Guide” – Jan 2011

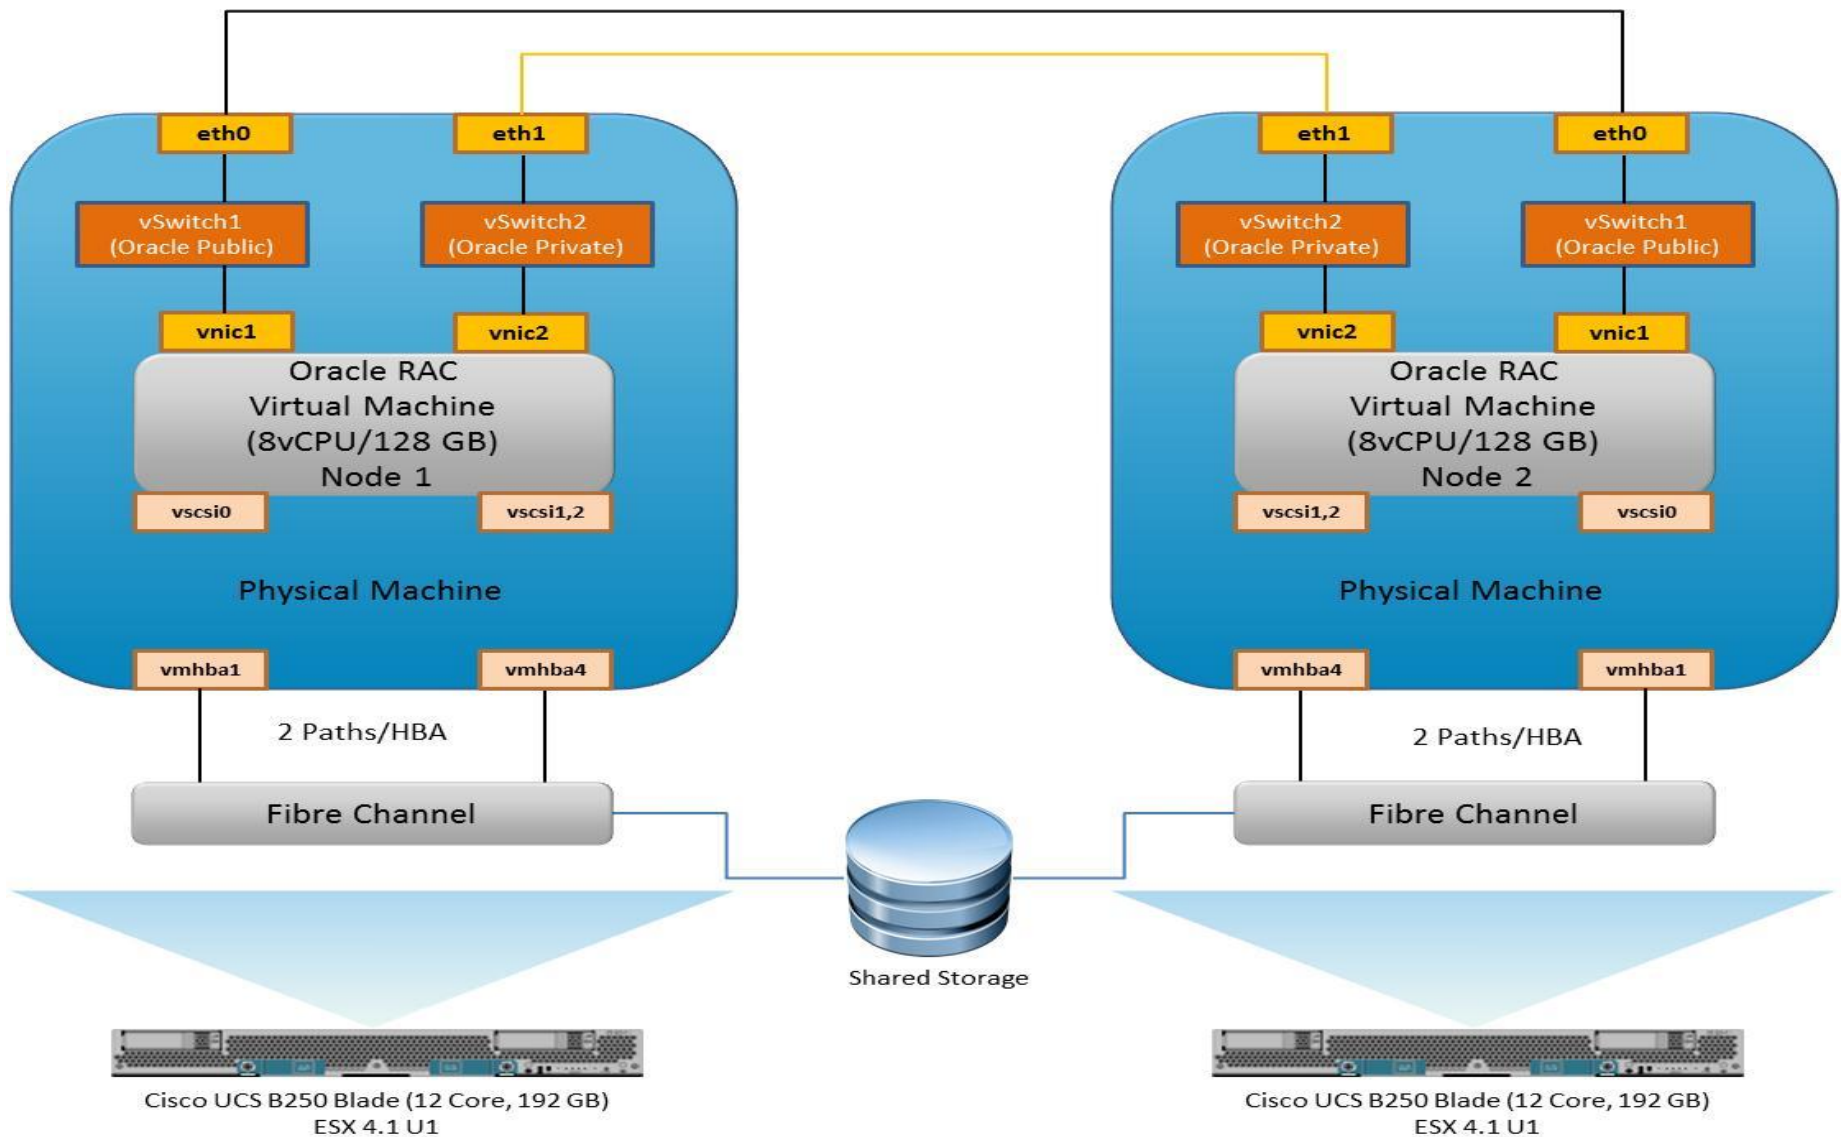
- **Methods**

- RDMs
- iSCSI Gateway – Cognizant Corporation – Chris Williams
- VMFS/VMDK

Oracle Real Applications Clusters (RAC) – Physical Server

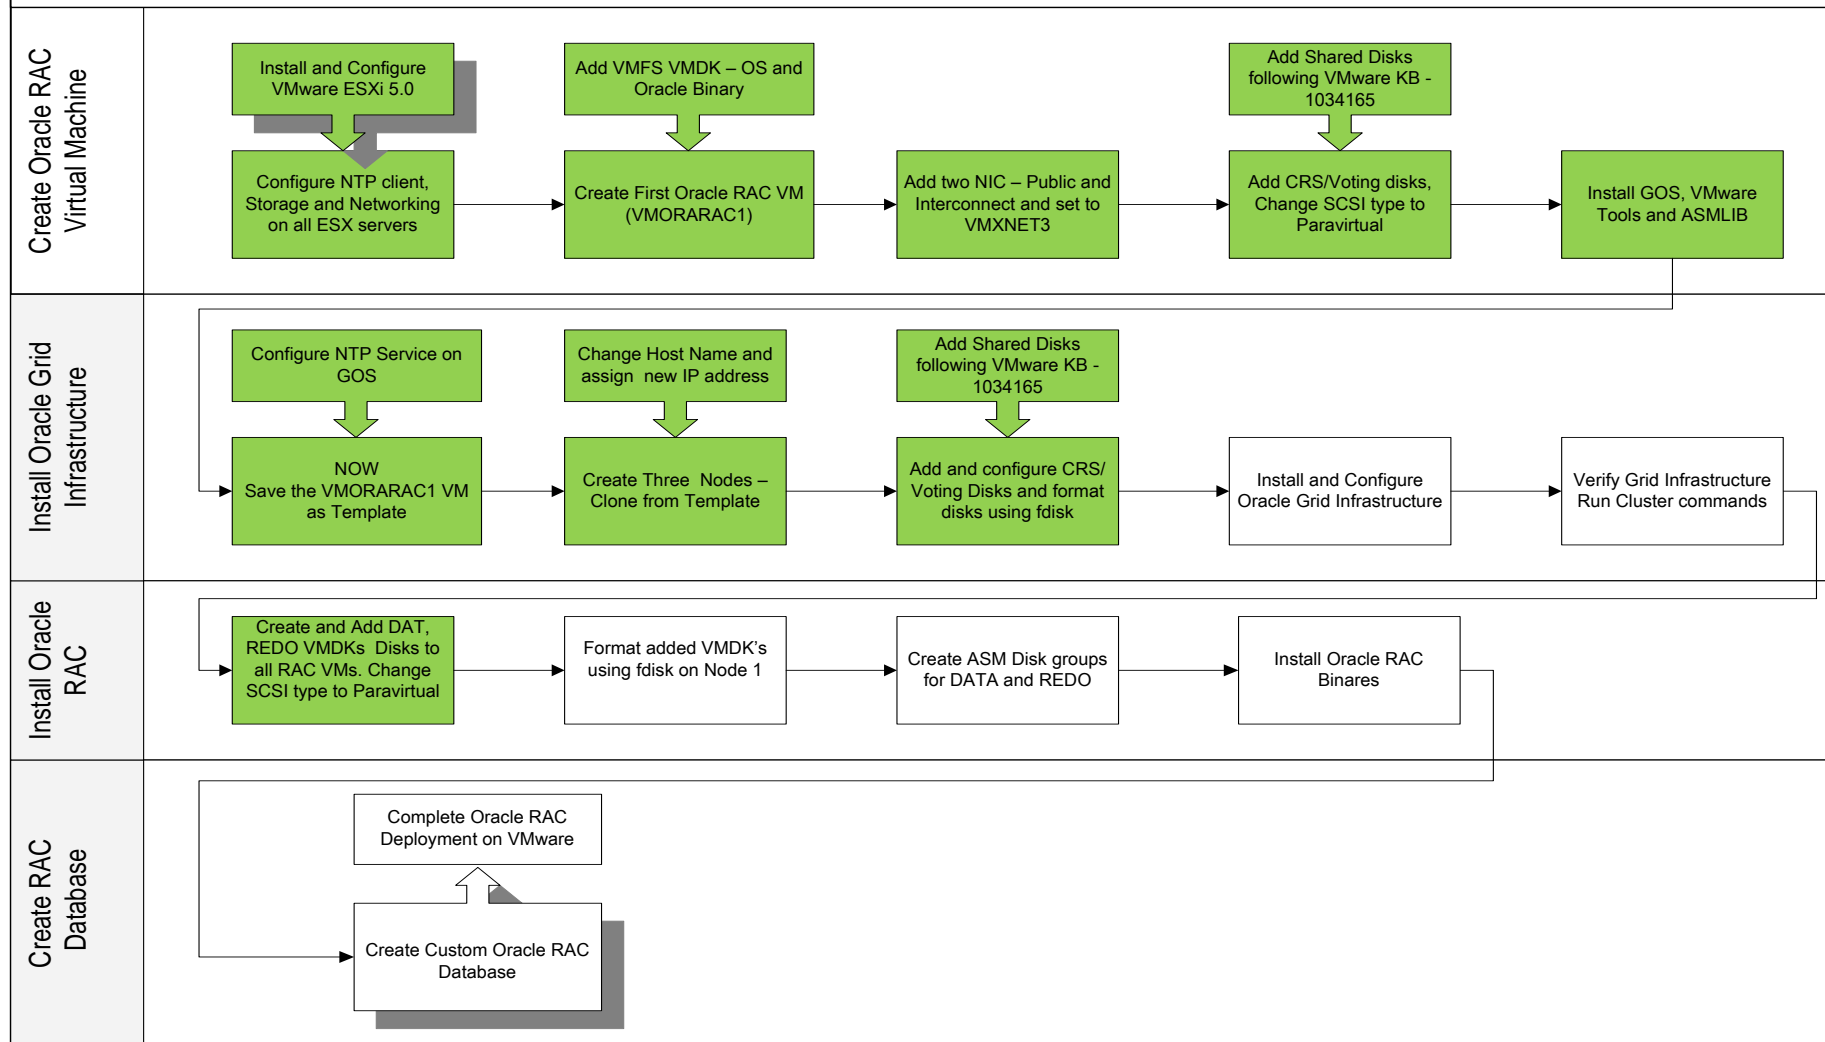


Oracle Real Applications Clusters (RAC) – VMware VMs



Oracle RAC on vSphere – Deployment Process Chart

Oracle RAC Deployment Process on VMware vSphere - VMFS



RAC on Vsphere – Networking

10.1.1.31 VMware ESXi, 5.0.0, 441354

Getting Started Summary Virtual Machines Performance Configuration Tasks & Events Alarms Permissions Maps Storage Views Hardware Status

Hardware

- Processors
- Memory
- Storage
- ▶ Networking
 - Storage Adapters
 - Network Adapters
 - Advanced Settings
 - Power Management

Software

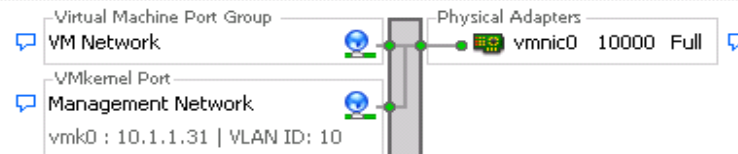
- Licensed Features
- Time Configuration
- DNS and Routing
- Authentication Services
- Power Management
- Virtual Machine Startup/Shutdown
- Virtual Machine Swapfile Location
- Security Profile
- Host Cache Configuration
- System Resource Allocation
- Agent VM Settings
- Advanced Settings

View: vSphere Standard Switch vSphere Distributed Switch

Networking

Standard Switch: vSwitch0

Remove... Properties...



Management

Standard Switch: vSwitch1

Remove... Properties...



Oracle RAC –
Public and
Private

Standard Switch: vSwitch2

Remove... Properties...



vMotion

Standard Switch: vSwitch3

Remove... Properties...



Oracle RAC on vSphere – Deployment Steps

- Install and Configure VMware ESXi 5.0
 - Configure Storage, Networking and NTP for all ESXi servers
- Create Oracle RAC VM Template
 - Add VMDK for Guest OS and binary
 - Add two vNICs for Oracle Public and Private – Change the adapter to VMXNET3
 - Add voting and CRS disks – Follow VMware KB – [1034165](#) for sharing the VMDKs
 - Change the SCSI type to Paravirtual
 - Install GOS, VMware tools, ASMLIB
 - Configure NTP service in GOS
 - Save the VM as Oracle RAC VM Template
- Install Oracle Grid infrastructure
 - Create Clones from RAC VM Template
 - Change the Cloned VM Host name and IP address
 - Add existing voting and CRS disks – Follow VMware KB – [1034165](#) for sharing the VMDKs
 - Install, configure and verify Oracle Grid infrastructure

Oracle RAC on vSphere – Deployment Steps ...Continued

■ Install Oracle RAC

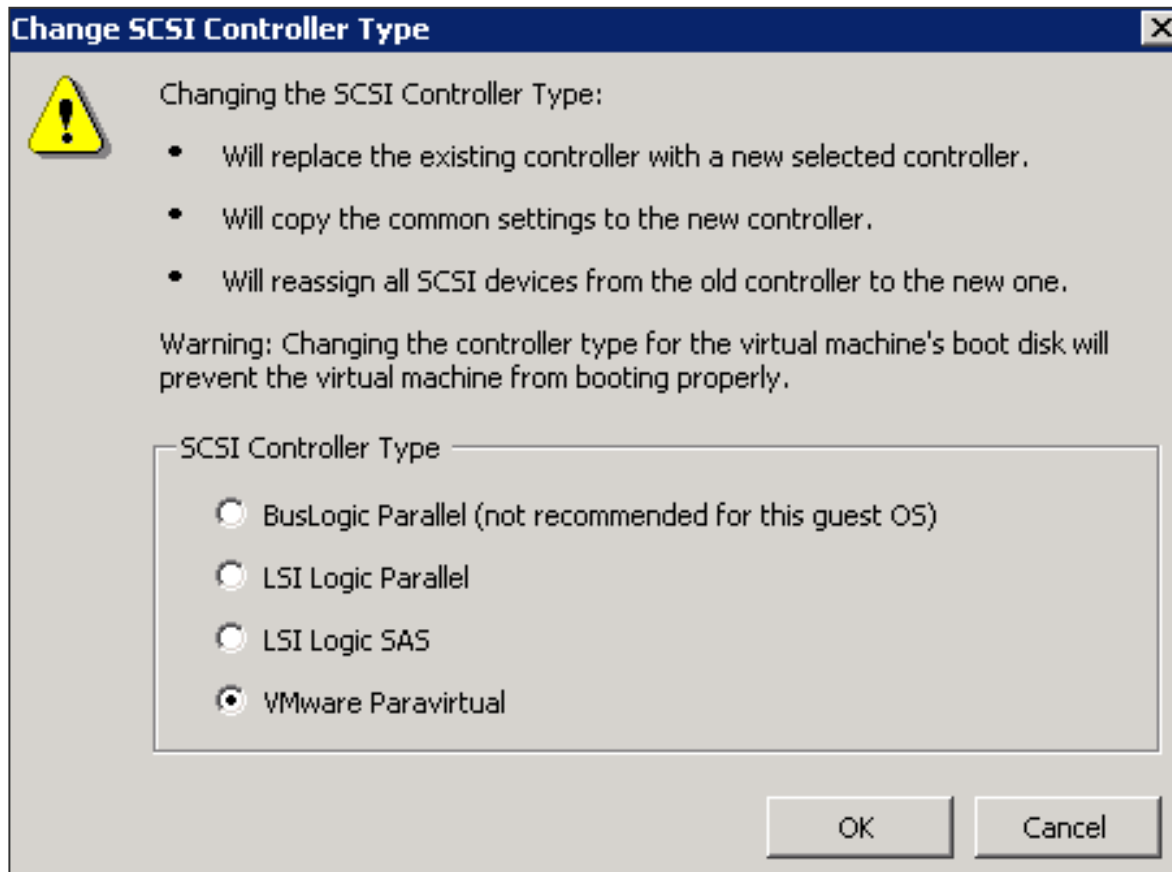
- Create Data, Redo VMDKs and add to all RAC VMs - Follow VMware KB – [1034165](https://kb.vmware.com/s/article/1034165) for sharing the VMDKs. Change the SCSI type to paravirtual.
- Format VMDKs using fdisk and create ASM disk groups for DATA and REDO
- Install Oracle RAC software

■ Create Oracle RAC Database

- Create custom oracle database
- Configure Oracle Enterprise Manager

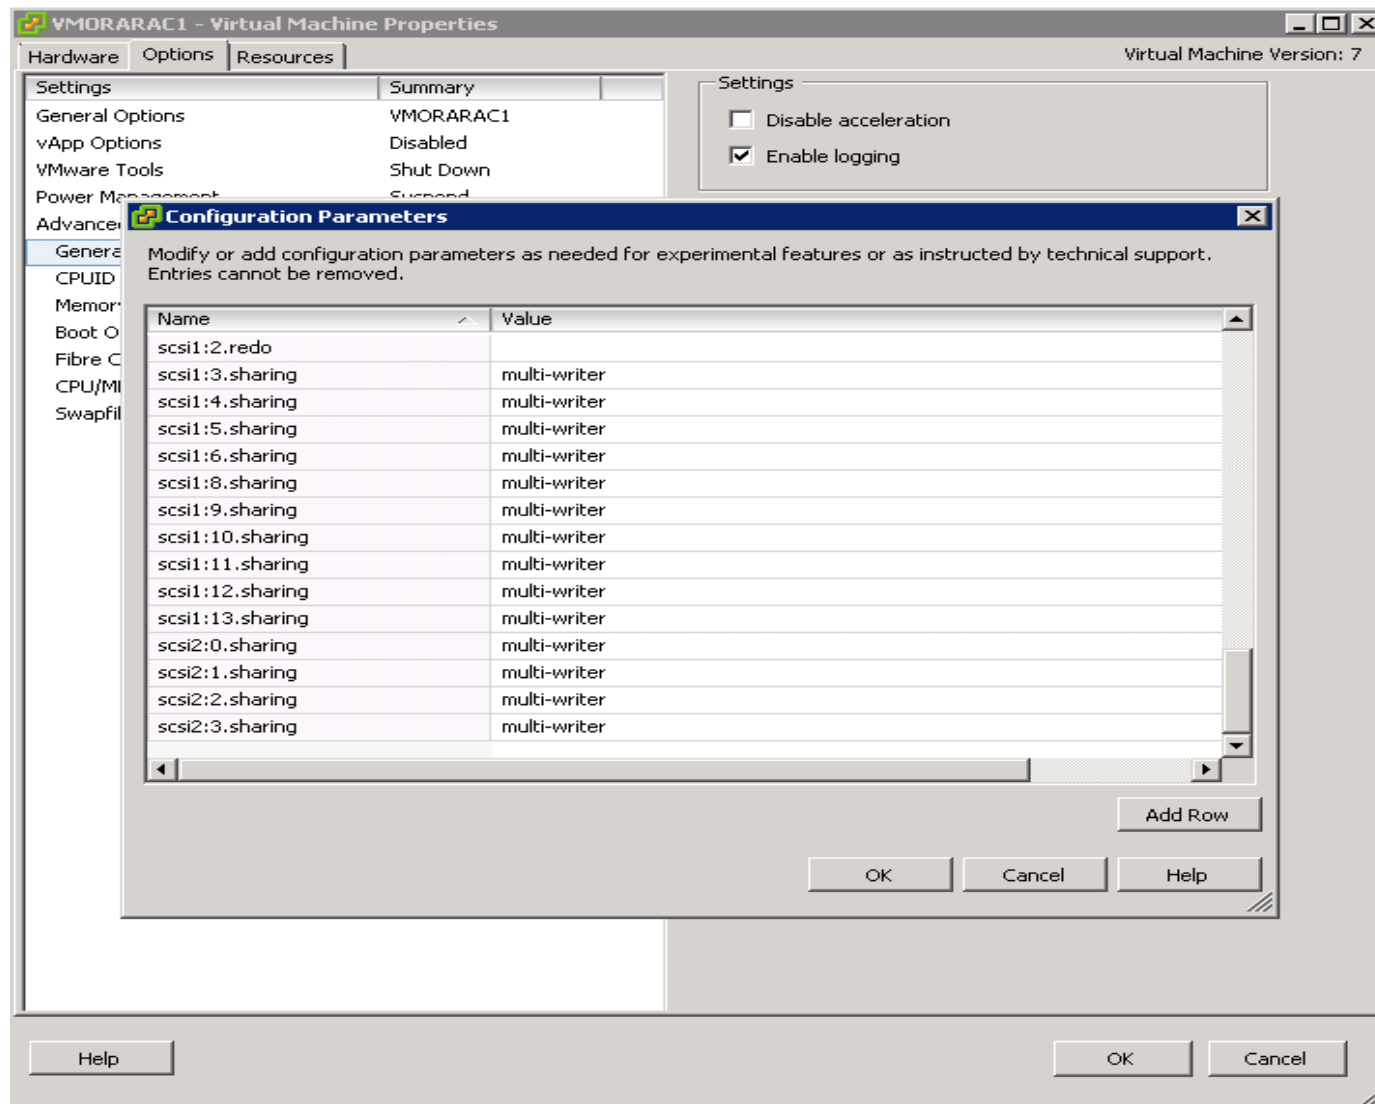
RAC on vSphere - Storage

- **Provision the storage – for the database VMDKs**
 - Change the SCSI controller to PVSCSI (VMware Paravirtual) for Data, Redo and Temp



RAC on Vsphere – Storage – Multi-Writer Flag

- Setup SCSI multi-writer (VMware KB – [1034165](https://kb.vmware.com/s/article/1034165))



RAC on vSphere – RAC VM Properties

VMORARAC1 - Virtual Machine Properties

Hardware | Options | Resources | Profiles | vServices | Virtual Machine Version: 8

☐ Show All Devices

Hardware	Summary
Memory	163840 MB
CPUs	12
Video card	Video card
VMCI device	Restricted
SCSI controller 0	LSI Logic Parallel
SCSI controller 1	Paravirtual
SCSI controller 2	Paravirtual
Hard disk 1	Virtual Disk
Hard disk 2	Virtual Disk
Hard disk 3	Virtual Disk
Hard disk 4	Virtual Disk
Hard disk 5	Virtual Disk
Hard disk 6	Virtual Disk
Hard disk 7	Virtual Disk
Hard disk 8	Virtual Disk
Hard disk 9	Virtual Disk
Hard disk 10	Virtual Disk
Hard disk 11	Virtual Disk
Hard disk 12	Virtual Disk
Hard disk 13	Virtual Disk
Hard disk 14	Virtual Disk
Hard disk 15	Virtual Disk
Hard disk 16	Virtual Disk
Hard disk 17	Virtual Disk
Hard disk 18	Virtual Disk
Network adapter 1	Oracle Public
Network adapter 2	Oracle Private

Memory Configuration

Memory Size: GB

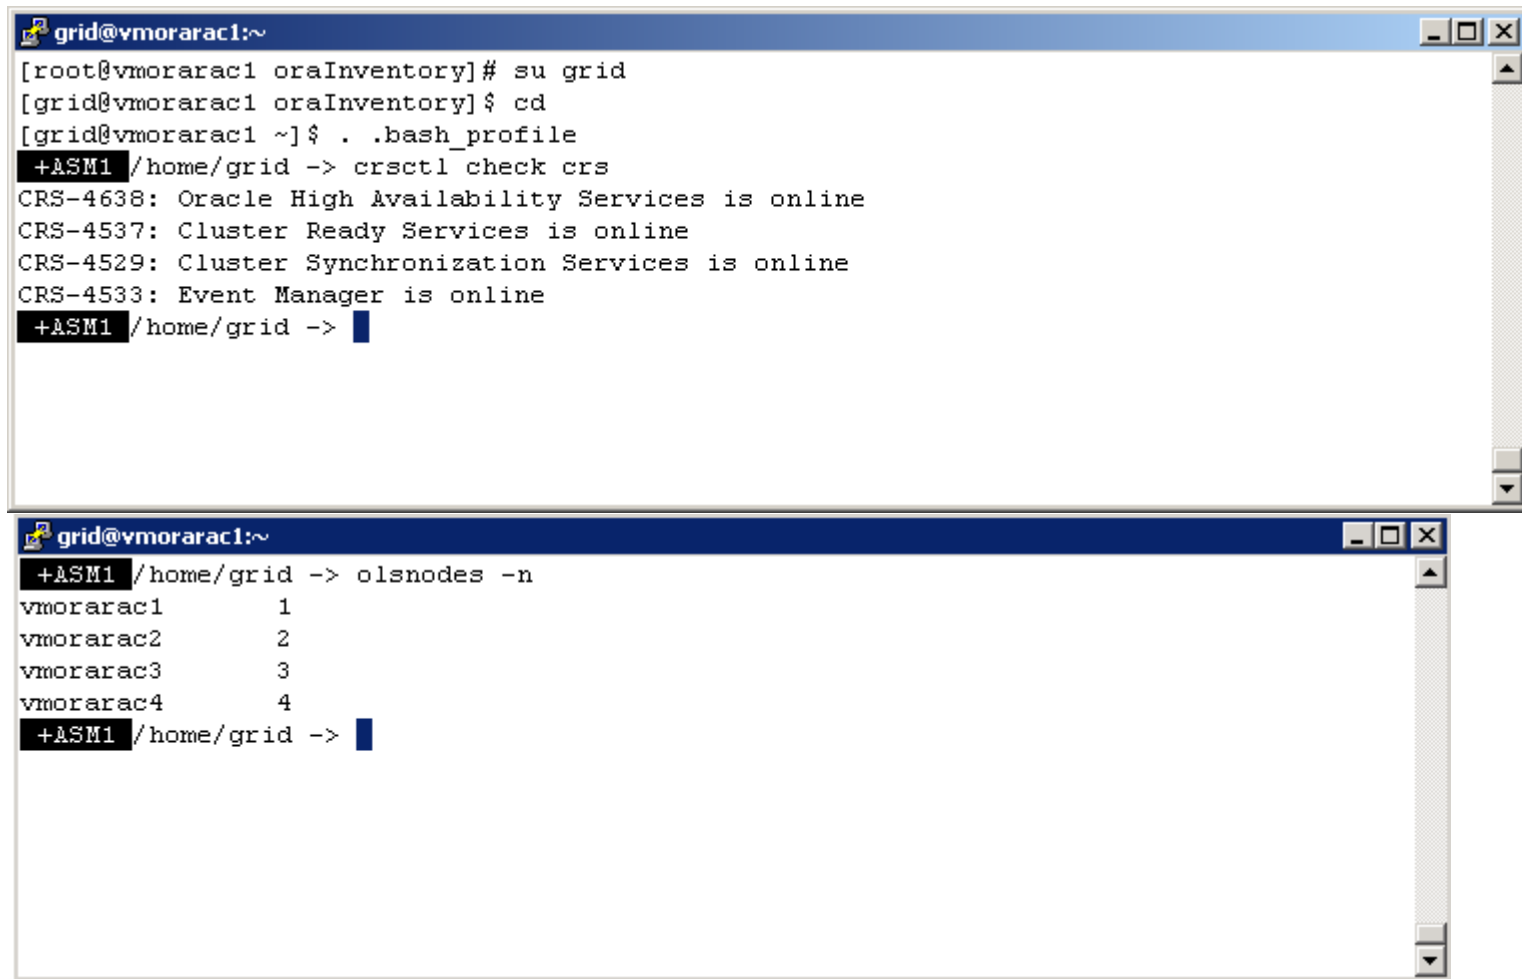
Maximum recommended for this guest OS: 1011 GB.
Maximum recommended for best performance: 196540 MB.
Default recommended for this guest OS: 2 GB.
Minimum recommended for this guest OS: 256 MB.

Callouts:

- For Root Disk (points to SCSI controller 1)
- For Data and Redo Disks (points to SCSI controller 2)
- OS and Binary (points to Hard disk 1)
- CRS Disks (points to Hard disk 2)
- Data Disks (points to Hard disk 3)
- Redo Disks (points to Hard disk 4)
- vNICs (points to Network adapter 1 and 2)

RAC on vSphere – crsctl

- Observe the rac install
 - srvctl, crsctl & olsnodes



The image displays two terminal windows from a VMware environment. The top window shows the user 'grid' logging in as 'root' and running 'crsctl check crs', which reports that all Oracle High Availability Services are online. The bottom window shows the user 'grid' running 'olsnodes -n', which lists the four nodes in the RAC cluster: vmorarak1, vmorarak2, vmorarak3, and vmorarak4.

```
grid@vmorarak1:~  
[root@vmorarak1 oraInventory]# su grid  
[grid@vmorarak1 oraInventory]$ cd  
[grid@vmorarak1 ~]$ . .bash_profile  
+ASM1 /home/grid -> crsctl check crs  
CRS-4638: Oracle High Availability Services is online  
CRS-4537: Cluster Ready Services is online  
CRS-4529: Cluster Synchronization Services is online  
CRS-4533: Event Manager is online  
+ASM1 /home/grid ->  
  
grid@vmorarak1:~  
+ASM1 /home/grid -> olsnodes -n  
vmorarak1      1  
vmorarak2      2  
vmorarak3      3  
vmorarak4      4  
+ASM1 /home/grid ->
```

RAC on vSphere – OCR and voting disk

- Observe the Grid Files

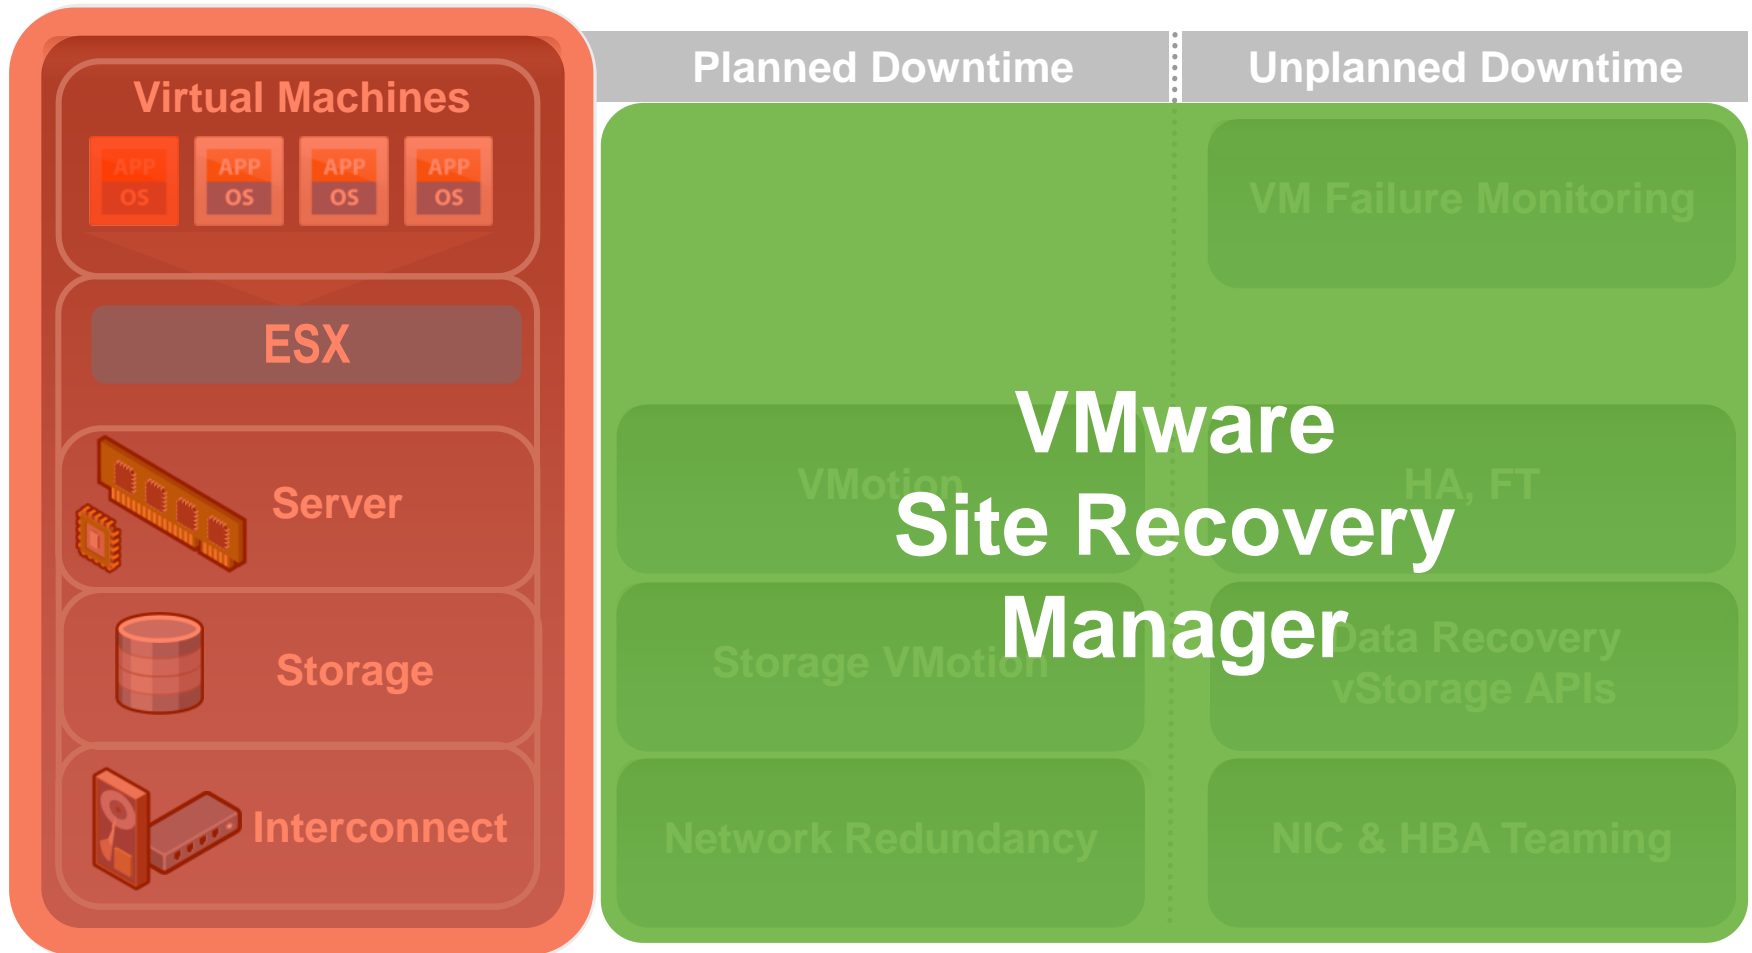
```
grid@vmoraracl1:~  
+ASM1 /home/grid -> ocrcheck  
Status of Oracle Cluster Registry is as follows :  
      Version                :                3  
      Total space (kbytes)    :            262120  
      Used space (kbytes)     :             2636  
      Available space (kbytes):            259484  
      ID                     :    1585277307  
      Device/File Name       :                +CRS  
                               Device/File integrity check succeeded  
  
                               Device/File not configured  
  
                               Device/File not configured  
  
                               Device/File not configured  
  
                               Device/File not configured  
  
      Cluster registry integrity check succeeded  
  
      Logical corruption check bypassed due to non-privileged user  
  
+ASM1 /home/grid ->  
  
grid@vmoraracl1:~  
+ASM1 /home/grid -> crsctl query css votedisk  
##      STATE      File Universal Id                File Name Disk group  
--      -  
1.  ONLINE     91667ba2625c4fefbf7e2c8938b8ebf8 (ORCL:CRSVOL1) [CRS]  
2.  ONLINE     29277d62fa324f5dbfad8b987c6b7332 (ORCL:CRSVOL2) [CRS]  
3.  ONLINE     490468754f174f0fbff297cd027925e8 (ORCL:CRSVOL3) [CRS]  
Located 3 voting disk(s).  
+ASM1 /home/grid ->
```

- **Commentary on NFS and DNFS**
 - **NFS to ESX**
 - NFS Datastores
 - **In-Guest NFS with DNFS**
 - Use Direct NFS (dNFS) with NetApp
 - Optimized for Oracle databases direct I/O access
 - Easy to configure - mount tab settings; change disk mount library
 - Still requires Kernel NFS for RAC quorum disk

Agenda

- **RAC on VMware – The Vsphere Platform**
 - Why use vSphere 5 as the platform for Oracle RAC
 - How to setup RAC on Vsphere - The installation process
 - Vsphere 5 features that make Vsphere complimentary to RAC
- **Performance capabilities**
- **Customer Success Stories**
- **The VRAC-DBA**
- **Summary**
- **Q/A**

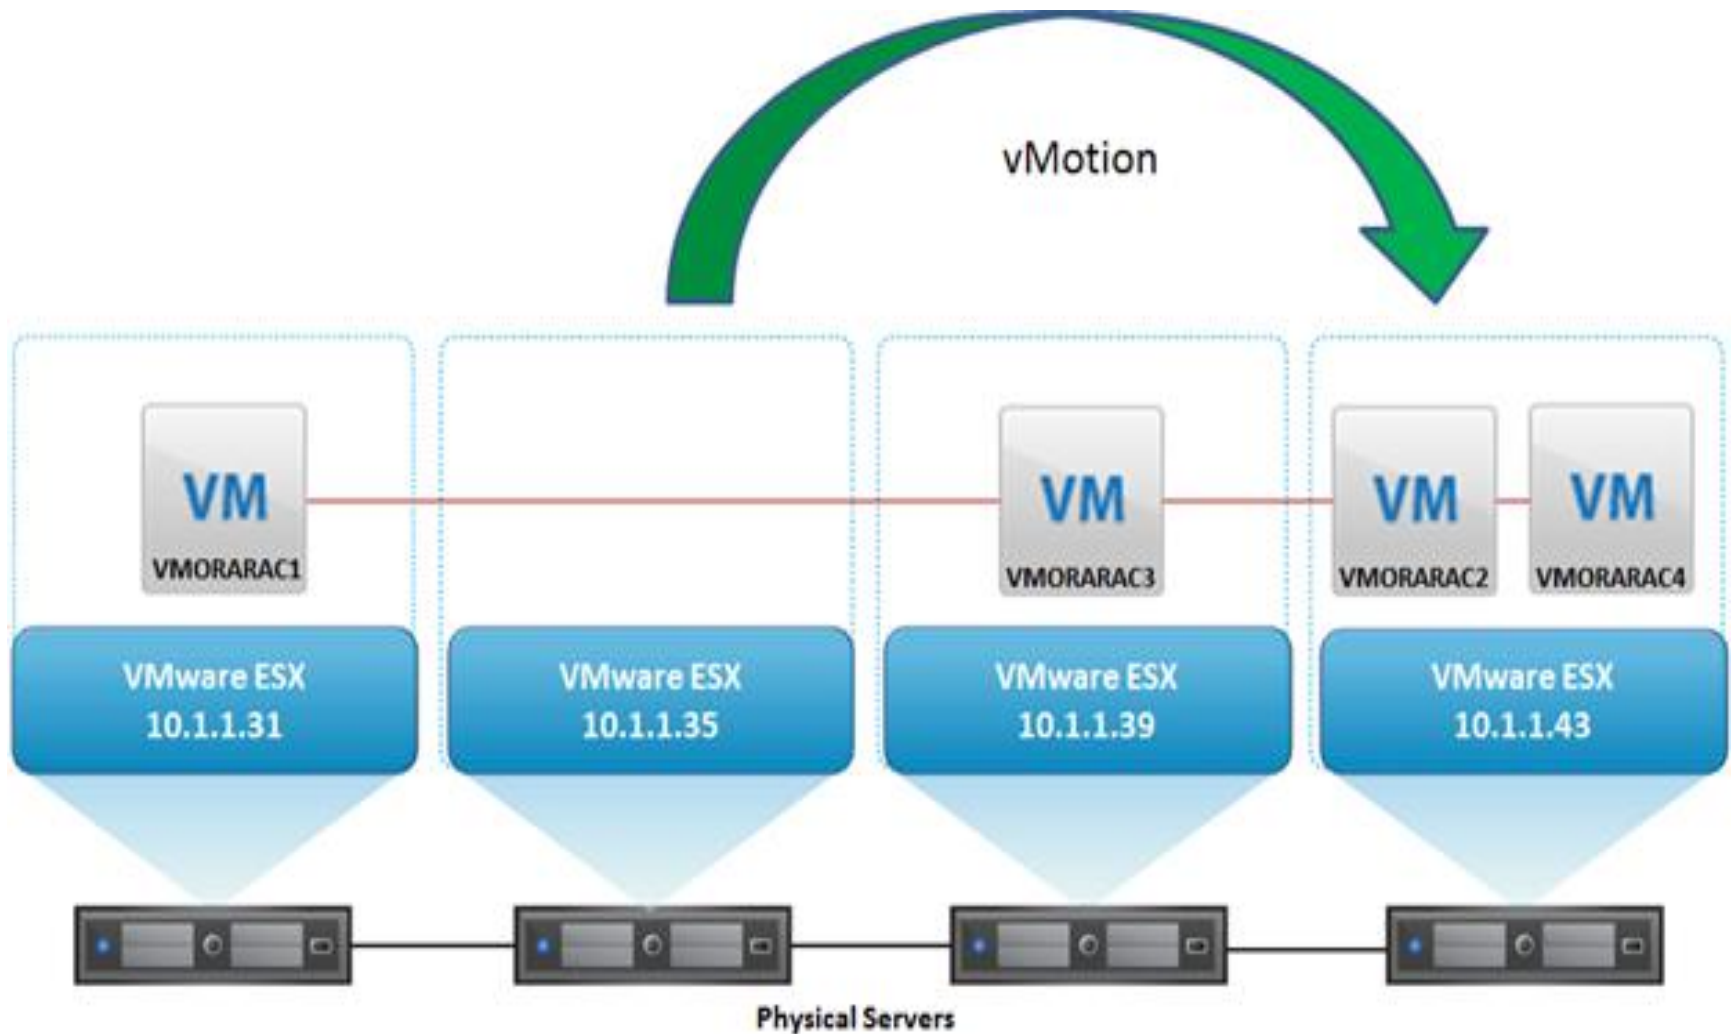
VMware Features to Eliminate Planned/Unplanned Downtime



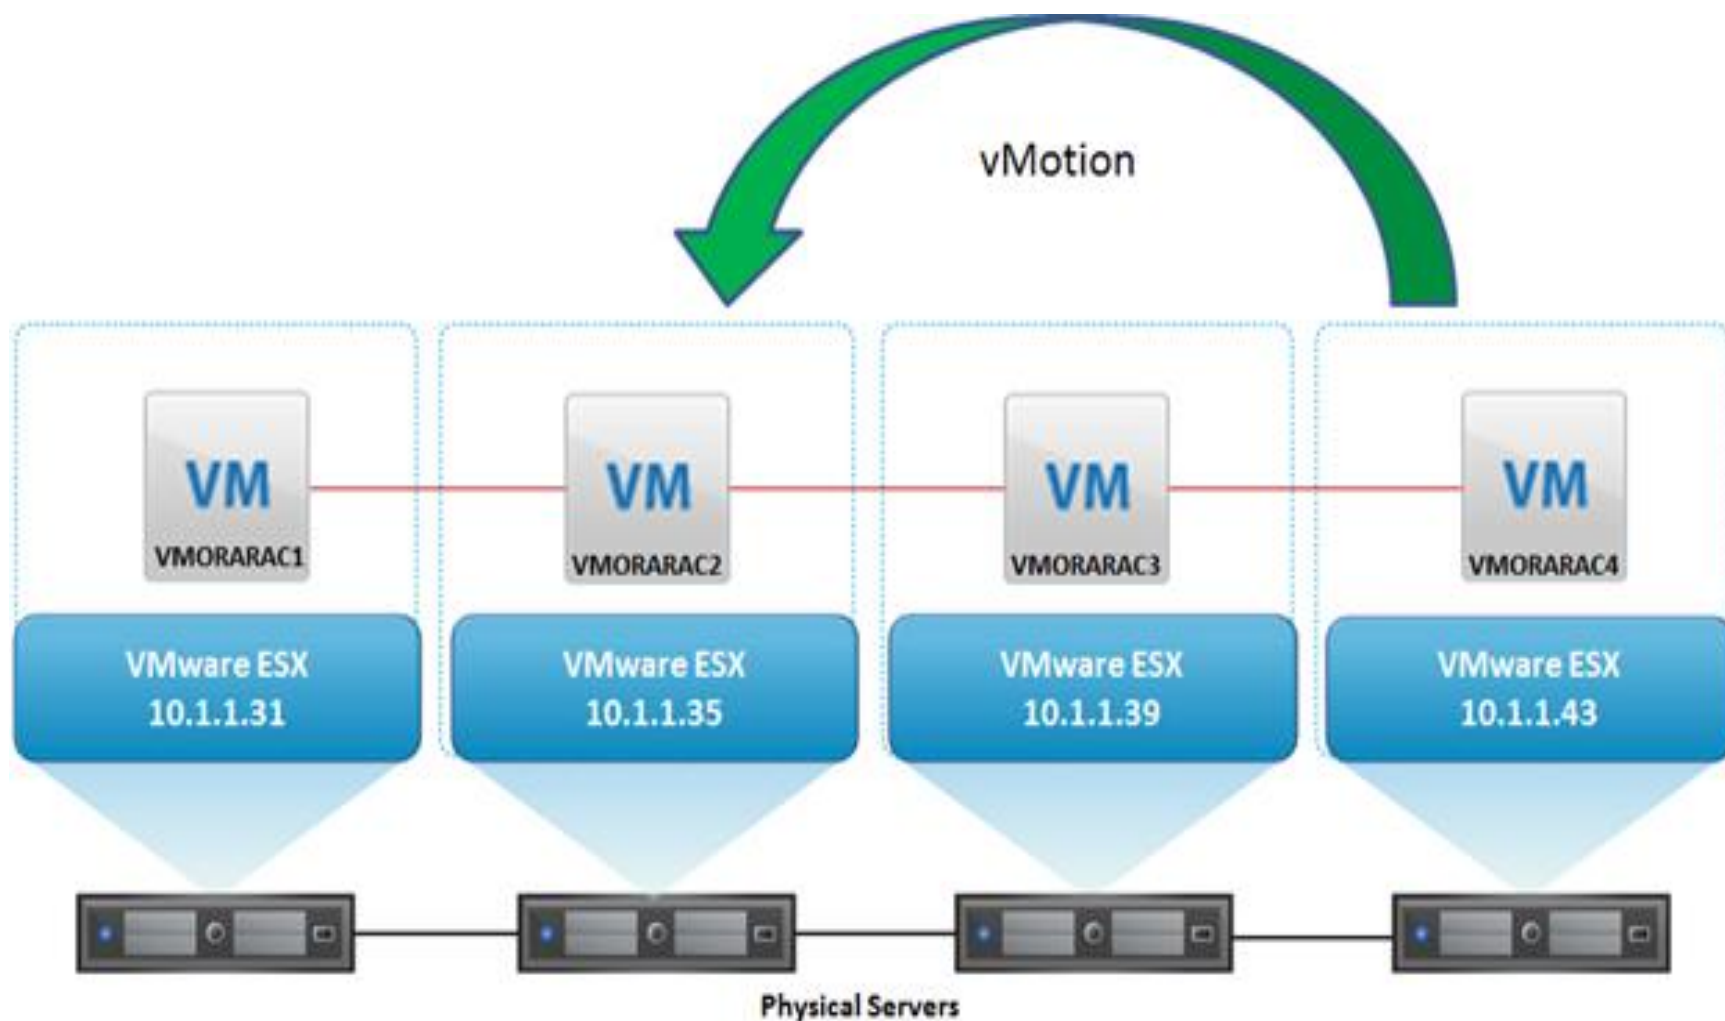
RAC on Vsphere – Vsphere 5 Features

- **Useful features for RAC**
 - Vmotion any Logical RAC node to an available ESX host
 - HotAdd one vCPU to any node
 - VMware HA – RAC Node restart on available ESX Hosts.

RAC on Vsphere – Features – Vmotion for Maintenance



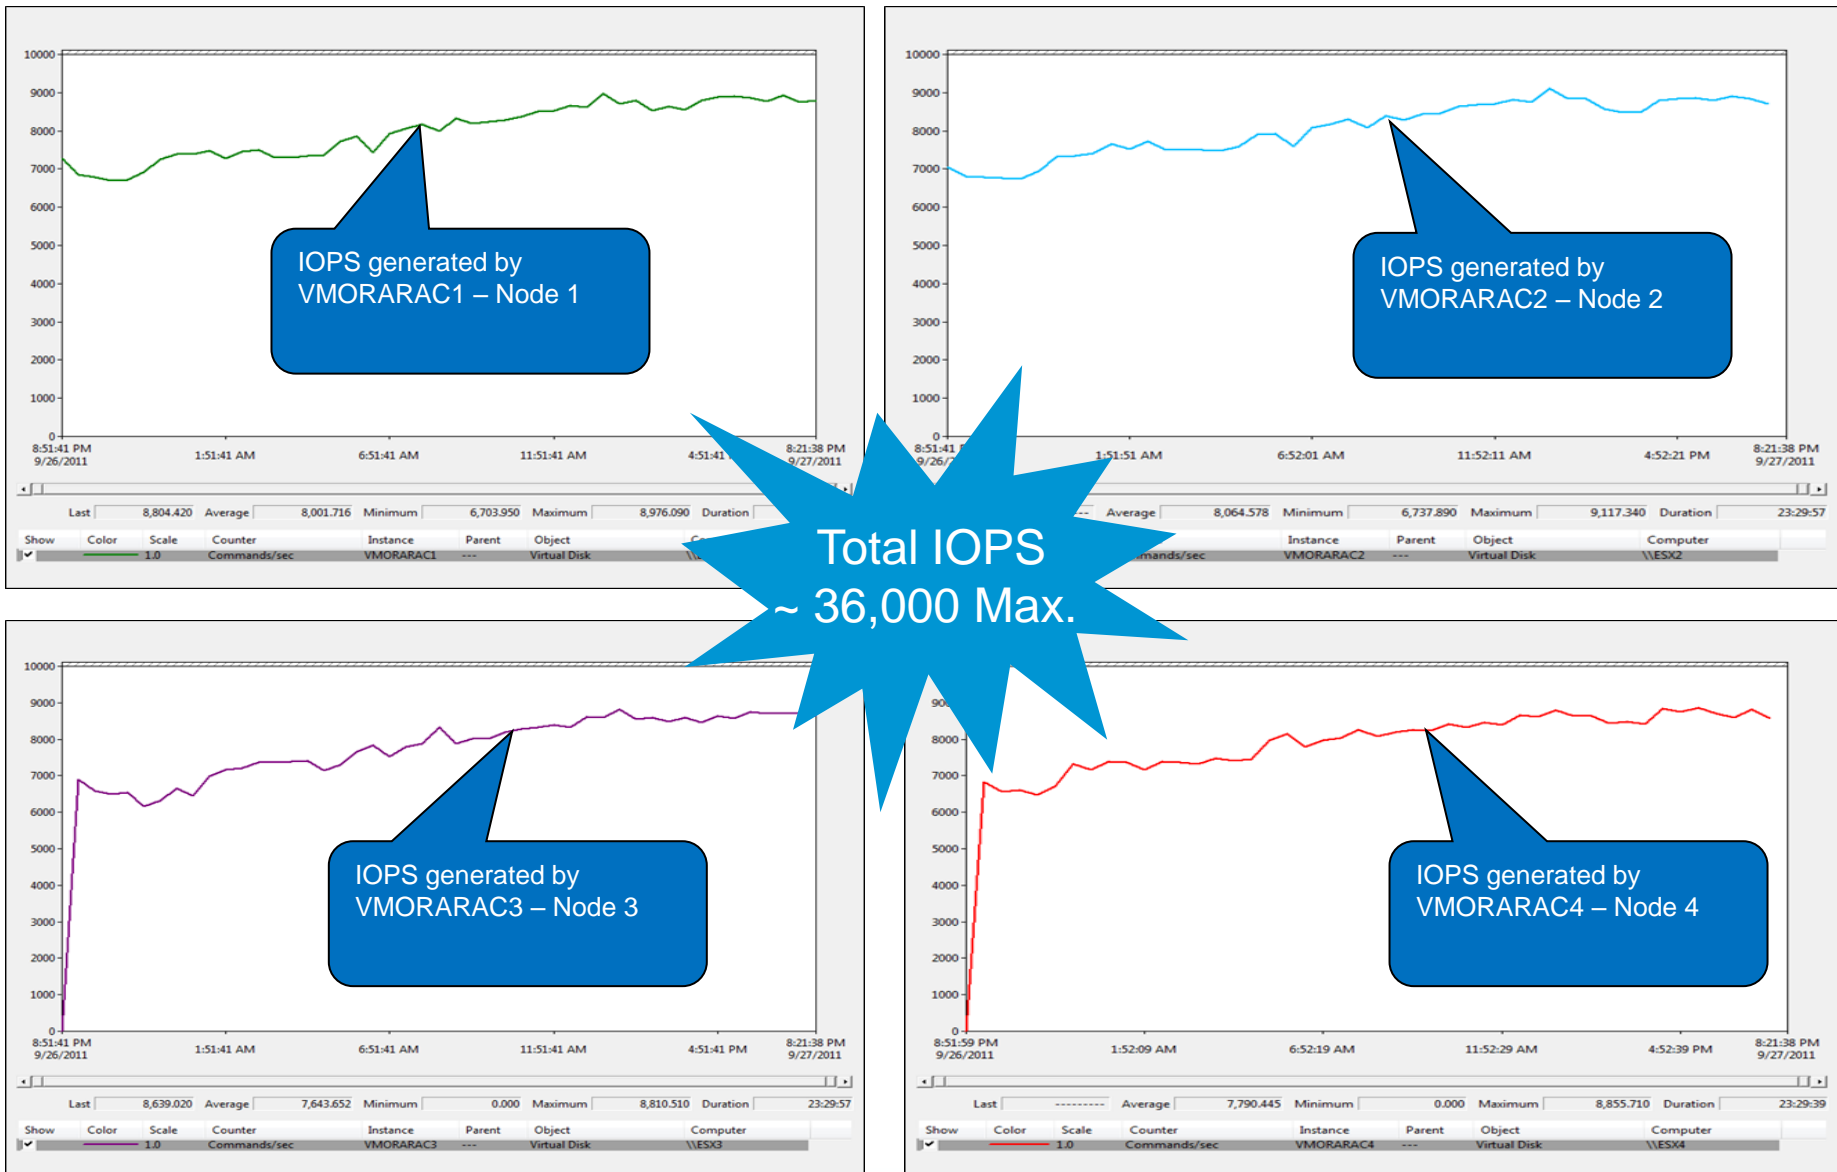
RAC on Vsphere – Features – After Maintenance Completion



Agenda

- **RAC on VMware – The Vsphere Platform**
 - Why use vSphere 5 as the platform for Oracle RAC
 - How to setup RAC on Vsphere - The installation process
 - Vsphere 5 features that make Vsphere complimentary to RAC
- **Performance capabilities**
- **Customer Success Stories**
- **The VRAC-DBA**
- **Summary**
- **Q/A**

Performance - IOPS for all the Four Node Oracle RAC VMs



Storage: Key Indicators

■ Kernel Latency Average (KAVG)

- This counter tracks the latencies of IO passing thru the Kernel
- Investigation Threshold: 1ms

■ Device Latency Average (DAVG)

- This is the latency seen at the device driver level. It includes the roundtrip time between the HBA and the storage.
- Investigation Threshold: 15-20ms, lower is better, some spikes okay

■ Aborts (ABRT/s)

- The number of commands aborted per second.
- Investigation Threshold: 1

Host Level Monitoring

- **vSphere Client:**
 - GUI interface, primary tool for observing performance and configuration data for one or more ESX/ESXi hosts
 - Does not require high levels of privilege to access the data
- **Resxtop/Esxtop**
 - Gives access to detailed performance data of a single ESX/ESXi host
 - Provides fast access to a large number of performance metrics
 - Requires root-level access
 - Runs in interactive, batch, or replay mode



The screenshot shows the Resxtop/Esxtop command-line interface running on a host. The header indicates the user is root@bk09-h380-11, connected via SSH on a 118x21 terminal. The system status shows the time as 8:18:46am, up for 4 days, 4:27, with 143 worlds and a CPU load average of 0.01, 0.01, 0.01. Below this is a table of performance metrics for various VMs.

ADAPTR PATH	NPTH	CMDS/s	READS/s	WRITES/s	MBREAD/s	MBWRTN/s	DAVG/cmd	KAVG/cmd	GAVG/cmd	QAVG/cmd
vmhba0 -	1	0.59	0.00	0.59	0.00	0.00	0.00	0.01	0.09	0.00
vmhba1 -	6	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba2 -	6	1741.29	1740.29	0.99	54.38	0.01	17.31	1.00	10.39	0.00
vmhba3 -	1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba32 -	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba33 -	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba34 -	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba35 -	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
vmhba36 -	0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

RAC on vSphere – Performance Monitoring

CPU Utilization during vMotion Test – vCenter Chart



Agenda

- **RAC on VMware – The Vsphere Platform**
 - Why use vSphere 5 as the platform for Oracle RAC
 - How to setup RAC on Vsphere - The installation process
 - Vsphere 5 features that make Vsphere complimentary to RAC
- **Performance capabilities**
- **Customer Success Stories**
- **The VRAC-DBA**
- **Summary**
- **Q/A**

Oracle RAC on vSphere – Joint Effort with Partners

- EMC

- EMC IT's "On-Ramp" to the Journey to the Private Cloud - <http://www.emc.com/collateral/hardware/white-papers/h8170-emc-it-on-ramp-cloud-wp.pdf>
- EMC IT's Virtual Oracle Deployment Framework, EMC IT is one the largest Oracle deployment which is running on VMware vSphere - <http://www.emc.com/collateral/hardware/white-papers/h8989-emc-it-virtual-oracle-deploy-framework.pdf>

- NetApp

- NetApp's Oracle Database 11g Release 2 Performance Using Data ONTAP 8.1 Operating in Cluster-Mode (4- Node Oracle 11gr2 RAC on vSphere compared with Bare Metal) - <http://media.netapp.com/documents/tr-3961.pdf>

Agenda

- **RAC on VMware – The Vsphere Platform**
 - Why use vSphere 5 as the platform for Oracle RAC
 - How to setup RAC on Vsphere - The installation process
 - Vsphere 5 features that make Vsphere complimentary to RAC
- **Performance capabilities**
- **Customer Success Stories**
- **The VRAC-DBA**
- **Summary**
- **Q/A**

- **The Future of the Oracle DBA**

- Virtualization endorsement
- Where will Virtualization be in 3-5 years?
- Ubiquitous and Comprehensive Tier1 Resource management
- Provisioning turnover expectations greatly increased
- Involvement in the creation and adjustment of the VM/Server to the exact specifications required
- Zero Downtime Hardware maintenance execution

- **The changing role of the DBA – The *VRAC-DBA***

- The modern Oracle DBA has influence over the entire stack
- Oracle DBAs do not have to re-learn their skill set
- The *VRAC-DBA* occupies the preeminent position in modern IT
 - From the Virtualized Infrastructure to the RAC instances and Apps
 - From Storage through the network Architectures

Agenda

- **RAC on VMware – The Vsphere Platform**
 - Why use vSphere 5 as the platform for Oracle RAC
 - How to setup RAC on Vsphere - The installation process
 - Vsphere 5 features that make Vsphere complimentary to RAC
- **Performance capabilities**
- **Customer Success Stories**
- **The VRAC-DBA**
- **Summary**
- **Q/A**

RAC on Vsphere - Resources

- <http://www.vmware.com/solutions/partners/alliances/oracle-database.html>
- <http://www.emc.com/solutions/application-environment/oracle/oracle-virtualization-vmware.htm>
- www.houseofbrick.com/oracle-on-vmware
- <http://blogs.vmware.com/apps/oracle/>

- Great Videos
 - <http://www.youtube.com/watch?v=V51en55eVnM>
 - <http://bit.ly/uy7o7W>
 - <http://bit.ly/t1Kskb>

Thank You

Solution Presentation - Don Sullivan

– Senior Systems Engineer - Database Specialist

sullivanand@vmware.com