



COLLABORATE12

TECHNOLOGY AND APPLICATIONS FORUM
FOR THE ORACLE COMMUNITY



Alternative Options to Enhance the High Availability of your Oracle Database

IOUG Collaborate 12, Session#840,
Monday, April 23rd, 2012, 12:15pm-12:45pm

Kai Yu



Oracle Solutions Engineering Lab
Dell Inc.



- **Kai Yu, *kai_yu@dell.com***
 - 17 years with Oracle Technology: DBA, Apps DBA, Architect
 - Hands-on experience on IT solution stack: hardware to software
 - Focus on Oracle RAC, Oracle VM and Oracle EBS
 - Oracle ACE Director, author and frequent presenter
 - IOUG Oracle RAC SIG President (2009-2010)
 - IOUG Virtualization SIG Webinar Chair
 - 2011 OAUG Innovator of Year Award Winner
 - Oracle Blog: <http://kyuoracleblog.wordpress.com/>
- **Dell Oracle Solutions Engineering: www.dell.com/oracle**
 - Oracle Technology Solutions on Dell systems/storages
 - Dell | Oracle Solutions Components
 - Solutions stack: servers, storage, network, OS, Oracle VM, Oracle RAC, Oracle Applications



Agenda

- Introduction of Oracle Database High Availability
- Database HA with Oracle RAC One Node
- Database HA with Oracle VM
- QA



Oracle Database High Availability

Goal of HA: provide uninterrupted database services

- Meet Service -Level Agreement (SLA)
- Avoid or minimize downtime during the essential time periods
- Planned downtime: system maintenance and upgrade
- Unplanned downtime: system failure, data corruption, hummer errors

Oracle HA solutions to reduce planned downtime

- Oracle Real Applications Cluster for OS/Hardware upgrades
- Online patching
- Database upgrade
- Transportable technology for across platform migration
- Online Application maintenance and upgrades



Oracle Database High Availability

Oracle HA solutions for unplanned downtime

- Server HA : Real Application Clusters (RAC)
- Data HA: Raid, ASM , Data Guard, Stream, Recovery manager
- Data Corruption: Recovery manager,
- Site HA: Data Guard, Stream, Recovery manager,
- MAA Architecture: an integrated HA solutions

Oracle Clusterware

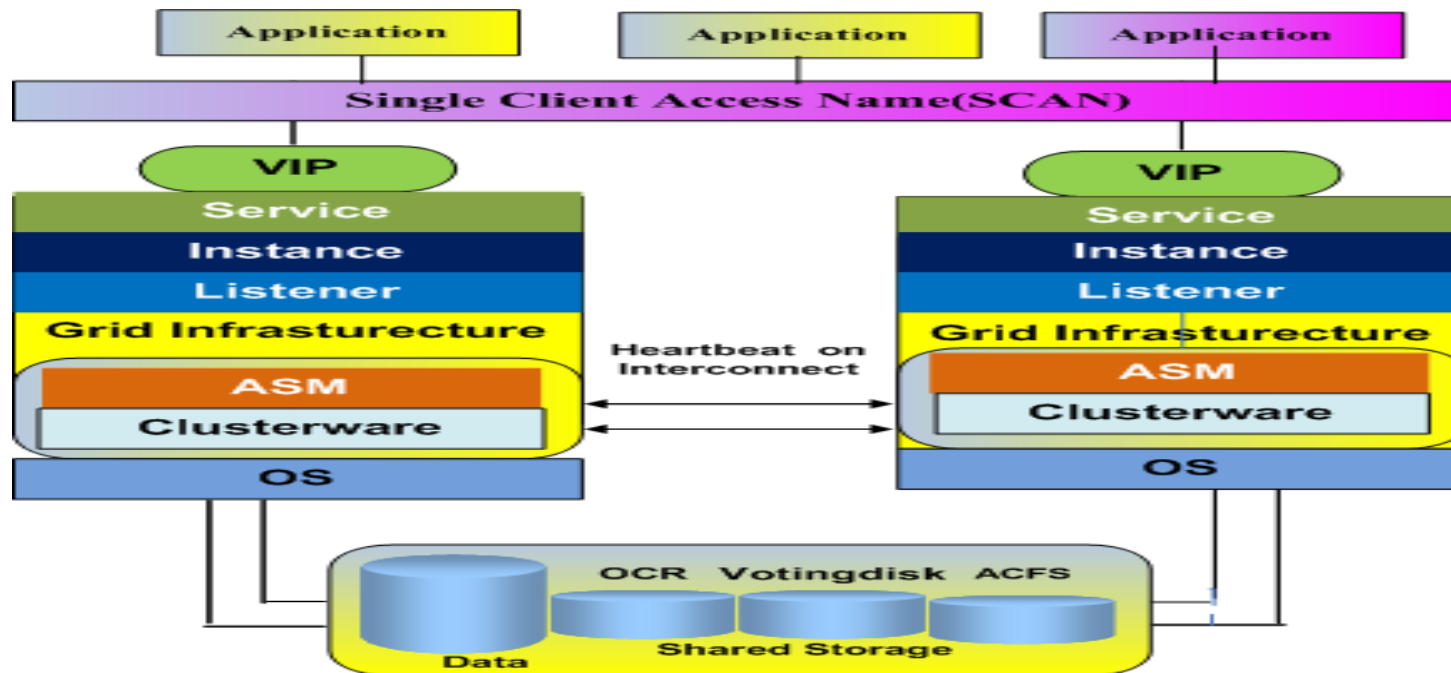
- Make multiple server work as one single server
- Manage the availability of the database and application
- Foundation for RAC database and HA features
- Manages RAC database + cluster resources
- Manage failover of Virtual IP to other node



Oracle Database High Availability

Oracle RAC Architecture

Multiple instance linked together to access the shared DB based on Oracle Clusterware





Oracle Database High Availability

Oracle RAC for Database:

Service Availability to reduce unplanned time:

- Redundancy and high availability with n-1 node failure
- Virtual IP automatically failover to remaining node in the server pool
- Fast Application Notification (FAN) for application connection fast failover

Service Availability to reduce the planned downtime:

- Avoid downtime for Server and OS upgrade
- Rolling upgrade to avoid downtime
- Oracle RAC provide load balance among the multiple nodes.

Oracle RAC License schema:

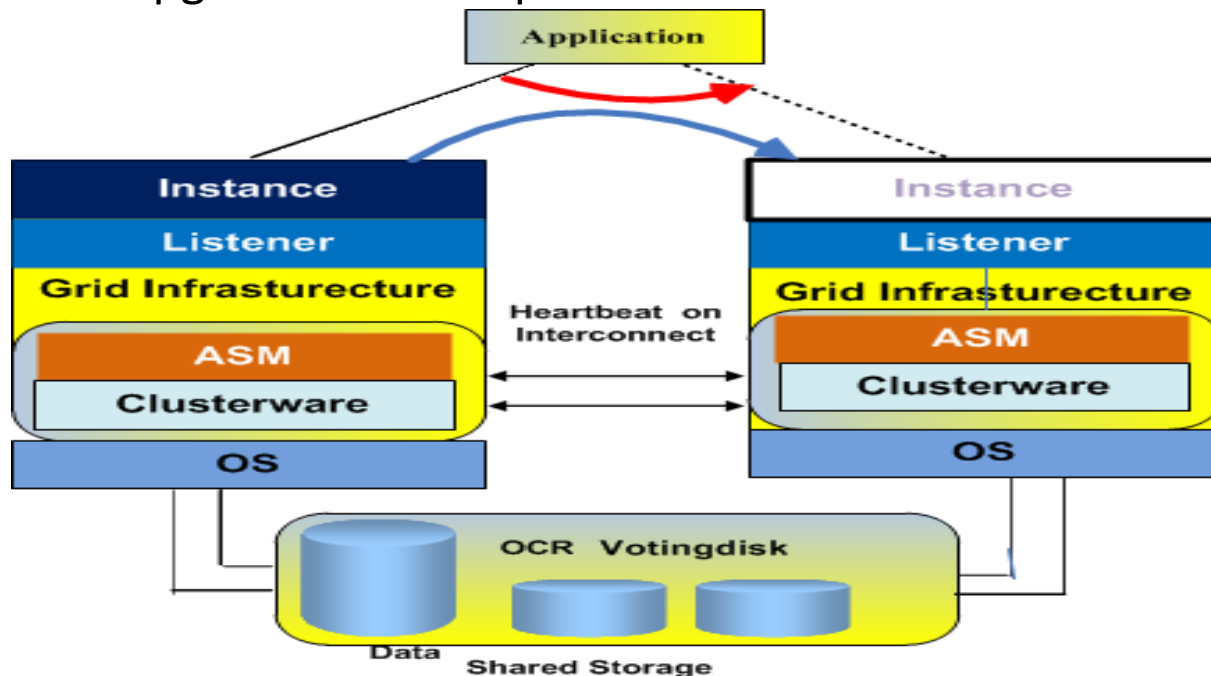
- Separate license for RAC on top of Oracle Database
- EE Edition: \$23k/processor , SE Edition: \$17k/processor , max 4 sockets
- For two nodes , two socket 8 core /server, $2 \times 2 \times 8 \times 0.5 = 16$ processors
- RAC license: \$512k (EE), \$272 K(SE), any way to reduce this cost?



Oracle RAC One Node Database

What is Oracle RAC One Node Database

- Single instance database running on Oracle cluster infrastructure
- Can be online relocated to difference node of the cluster
- Can be upgraded to multiple node RAC database





Oracle RAC One Node Database

Why do we want Oracle RAC One Node Database

- HA against unplanned downtime: automatically failover
- HA against planned downtime: online relocation to other node
- License advantage over RAC
- Allow to consolidate multiple databases to private cloud
- Easily online upgradeable to multiple nodes RAC database
- Work with server virtualization like Oracle VM

Oracle RAC One Node license:

- \$10K/processor vs RAC \$23k/processor (EE) , 17k/processor (SE)
- Only pay for one node and allow 10 days/year to failover another node.
- Two nodes 2 sockets 8 core server: $1 \times 2 \times 8 \times 0.5 = 8$ processors, \$80K
- Compared for RAC: \$512k (EE), \$272 K(SE),

What the RAC one node doesn't provide

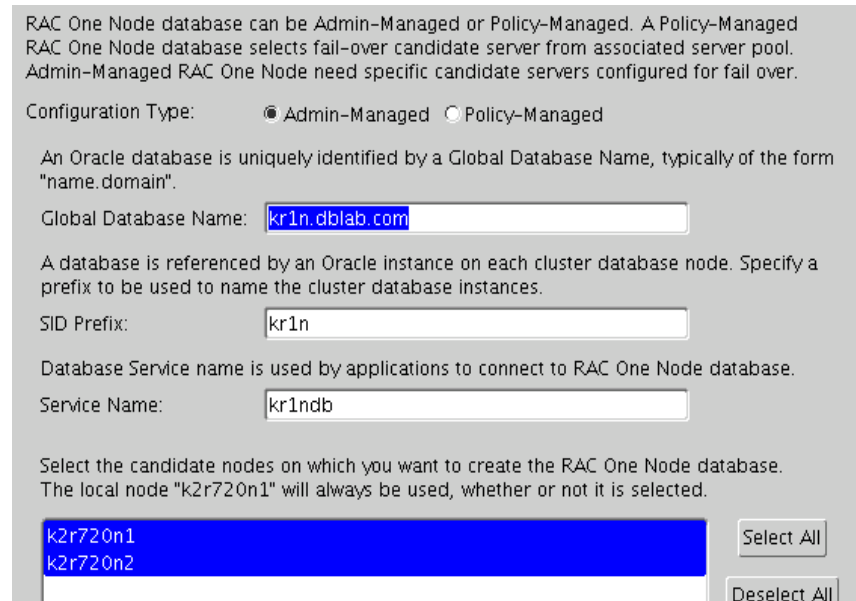
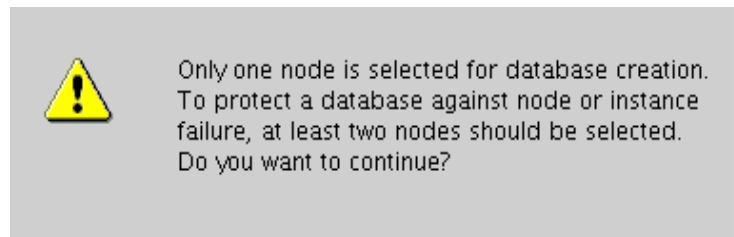
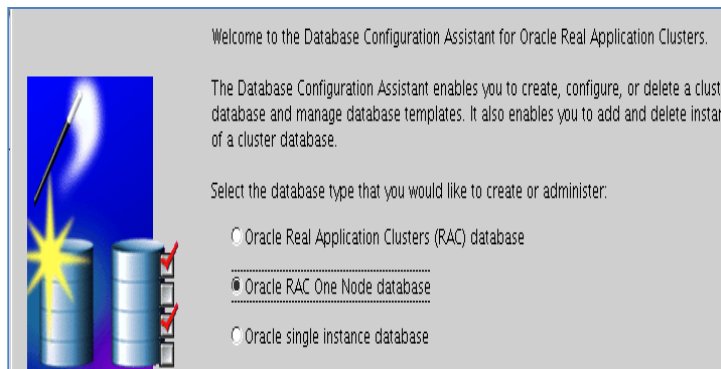
- Load balance among nodes , all the nodes are utilized.



Oracle RAC One Node Database

Deploy RAC One Node Database

- Prepare Hardware: servers, shared storage, and network
- Install Oracle 11gR2 Grid Infrastructure on all nodes
- Install Oracle 11gR2 RAC software on all nodes
- Create RAC one Node Database





Oracle RAC One Node Database

```
[oracle@k2r720n1 ~]$ srvctl config database -d kr1n
```

Database unique name: kr1n

Database name: kr1n

Oracle home: /u01/app/oracle/product/11.2.0/dbhome_1

Oracle user: oracle

Spfile: +DATA/kr1n/spfilekr1n.ora

Domain: dblab.com

Start options: open

Stop options: immediate

Database role: PRIMARY

Management policy: AUTOMATIC

Server pools: kr1n

Database instances:

Disk Groups: DATA,BKDG

Mount point paths:

Services: kr1ndb

Type: **RACOneNode**

Online relocation timeout: 30

Instance name prefix: kr1n



Oracle RAC One Node Database

Avoid the planned downtime: Rolling online patching

Rolling upgradable RAC patches are can be applied to RAC One node as rolling upgrade patch

- Steps:
- online relocate the database to secondary instance
 - Apply the patch on the primary instance
 - move the database back to the primary instance
 - Apply the patch on the secondary instance

Online relocation from one node to other node : srvctl utility

```
[oracle@k2r720n1 ~]$ srvctl relocate database -d kr1n -n k2r720n2 -w 15 -v
Configuration updated to two instances

Instance kr1n_2 started
Services relocated
Waiting for up to 15 minutes for instance kr1n_1 to stop ...
Instance kr1n_1 stopped
Configuration updated to one instance
```



Oracle RAC One Node Database

Check the online relocation status:

```
[oracle@k2r720n2 pfiles]$ srvctl status database -d krin
Instance krin_1 is running on node k2r720n1
Online relocation: ACTIVE
Source instance: krin_1 on k2r720n1
Destination instance: krin_2 on k2r720n2
```

Both instances up during the online relocation

```
[oracle@k2r720n2 pfiles]$ ps -ef | grep pmon
grid      26415      1  0 Mar09 ?          00:00:56 asm_pmon_+ASM2
oracle    29216      1  0 17:42 ?          00:00:00 ora_pmon_krin_2
```

```
[oracle@k2r720n1 ~]$ ps -ef | grep pmon
grid      13751      1  0 Mar12 ?          00:00:24 asm_pmon_+ASM1
oracle    25440      1  0 17:18 ?          00:00:00 ora_pmon_krin_1
```



Oracle RAC One Node Database

Reduce the unplanned downtime:

- Server fail → Database failover to other node by clusterware

Node1:

```
[root@k2r720n1 ~]# reboot ← shutdown node1
```

```
Broadcast message from root (pts/3) (Fri Mar 16 16:33:54 2012):
```

```
The system is going down for reboot NOW!
```

Node2:

```
[oracle@k2r720n2 ~]$ ps -ef | grep pmon
```

```
oracle 25151 24935 0 16:34 pts/2 00:00:00 grep pmon
```

```
grid 26415 1 0 Mar09 ? 00:00:58 asm_pmon_+ASM2
```

```
[oracle@k2r720n2 ~]$ ps -ef | grep pmon
```

```
oracle 25177 1 0 16:34 ? 00:00:00 ora_pmon_kr1n_1 ← instance failover
```

```
oracle 25930 24935 0 16:36 pts/2 00:00:00 grep pmon
```

```
grid 26415 1 0 Mar09 ? 00:00:58 asm_pmon_+ASM2
```



Oracle RAC One Node Database

Convert RAC one node to RAC Database

- **Convert to RAC** `$ srvctl convert database -d kr1n -c RAC -n k2r720n1`

```
[oracle@k2r720n1 ~]$ srvctl config database -d kr1n
```

```
Database unique name: kr1n
```

```
.....
```

```
Database instances: kr1n_1
```

```
....
```

```
Type: RAC
```

```
Check the database status: $ srvctl status database -d kr1n
```

```
Instance kr1n_1 is running on node k2r720n1
```

- **Add an instance:** `$srvctl add instance -d kr1n -i kr1n_2 -n k2r720n2`

```
Check the status: $srvctl status database -d kr1n
```

```
Instance kr1n_1 is running on node k2r720n1
```

```
Instance kr1n_2 is not running on node k2r720n2
```



Oracle RAC One Node Database

Convert RAC Database to RAC One Node:

- **Stop/remove 2nd instance** \$ srvctl stop instance -d kr1n -i kr1n_2
\$srvctl remove instance -d kr1n -i kr1n_2

. Convert to RAC one Node:

\$srvctl convert database -d kr1n -c RACONENODE -w 30 -i kr1n

Check status: ~]\$ ~]\$ **srvctl config database -d kr1n**

Database unique name: kr1n

...

Type: RACOneNode

...

Candidate servers: k2r720n1

...



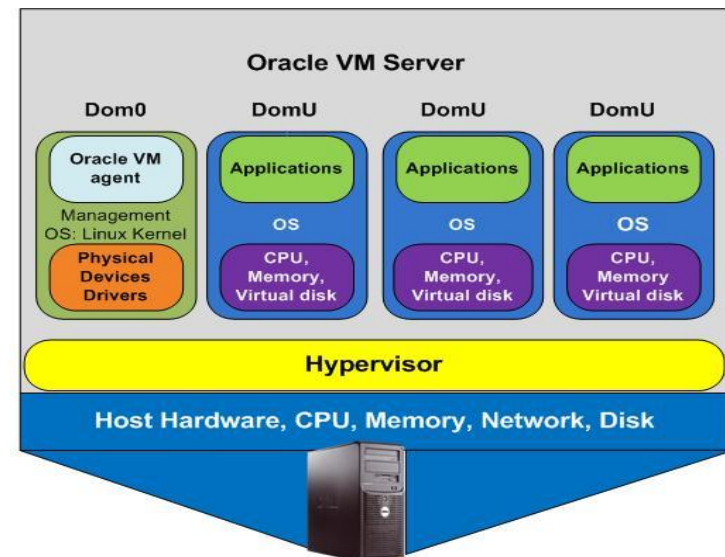
Database HA with Oracle VM

Oracle VM Virtualization Solution.

- Software based server virtualization solutions for X86
- Server partitions and consolidation
- High availability and scalability

Oracle VM Architecture

- Oracle VM server: Xen Hypervisor, management Linux kernel with support of devices, IO, networking, etc.
- VM: guest OS with applications running on domU
- Virtual Server Pool:
an autonomous region of VM servers




For more details of running Oracle Database on Oracle VM, Please attend my **Virtualizing Oracle 11g/R2 RAC Database on Oracle VM: Methods/Tips** Session#328, **Monday, 3:45-4:45pm, Palm D**



Database HA with Oracle VM

Live Migration of Database Virtual Machine (VM):

online migration of VM from one VM server to another.



Virtual Machines

Select and

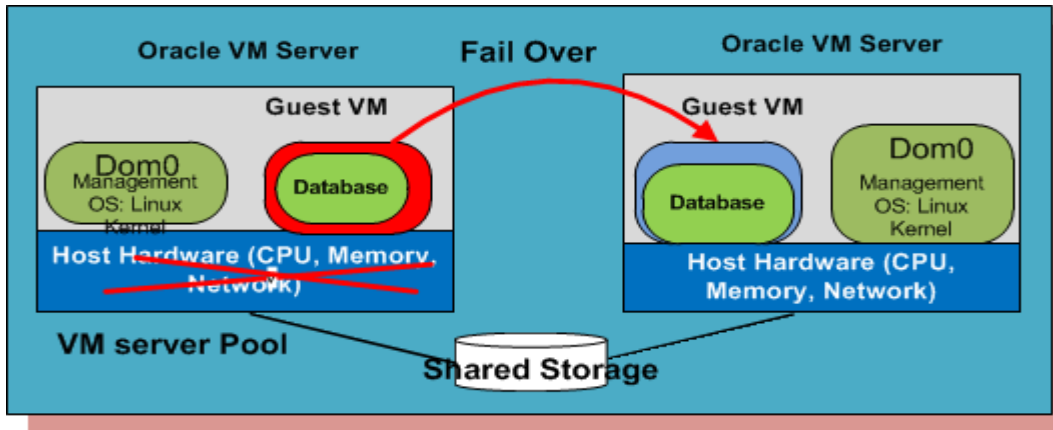
Power On Console Power Off Configure More Actions: --Select--

| Select | Details | Virtual Machine Name | Memory Size (MB) | Status | Owner | Group Name | Server Name |
|----------------------------------|----------------------|----------------------|------------------|-----------|-------|--------------|-----------------|
| <input checked="" type="radio"/> | Show | ebs 1211 db | 2,048 | Migrating | admin | My Workspace | owivs1.us.dell. |
| <input type="radio"/> | Show | ebs 12 1 1 apps | 1,048 | Running | admin | My Workspace | owivs1.us.dell. |



Database HA with Oracle VM

Failover of single node database VM from one physical server to another:



Select and


Power On Console Power Off Configure More Actions: --Select-- Go

| Select | Details | Virtual Machine Name | Memory Size (MB) | Status | Owner | Group Name | Server Name | Server Pool Name |
|--------------------------|--------------------------|----------------------|------------------|---------|-------|--------------|-------------------|------------------|
| <input type="checkbox"/> | <input type="checkbox"/> | Show ebs 1211 db | 2,048 | Running | admin | My Workspace | owvs2.us.dell.com | owi_pool |

Select and

Power On Console Power Off Configure More Actions: --Select-- Go

| Select | Details | Virtual Machine Name | Memory Size (MB) | Status | Owner | Group Name | Server Name | Server Pool Name |
|--------------------------|--------------------------|----------------------|------------------|---------|-------|--------------|--------------------|------------------|
| <input type="checkbox"/> | <input type="checkbox"/> | Show ebs 1211 db | 2,048 | Running | admin | My Workspace | owivs1.us.dell.com | owi_pool |



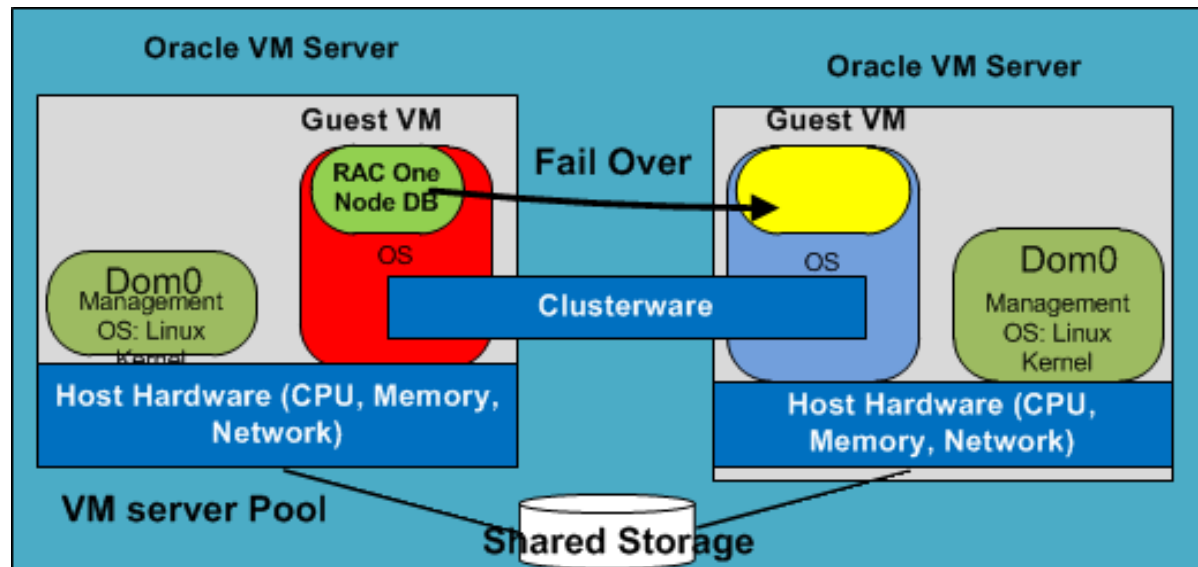
Failed over to OWIVS1 in 1-2 minutes



Database HA with Oracle VM

RAC One Node Database in Oracle VM environment.

- RAC One Node fully supported in Oracle VM environment.
- Run RAC One Node on Virtual Machine:
- Ensure the high availability of RAC One Node database in VM
The database will failover to another VM if the this VM node has issue





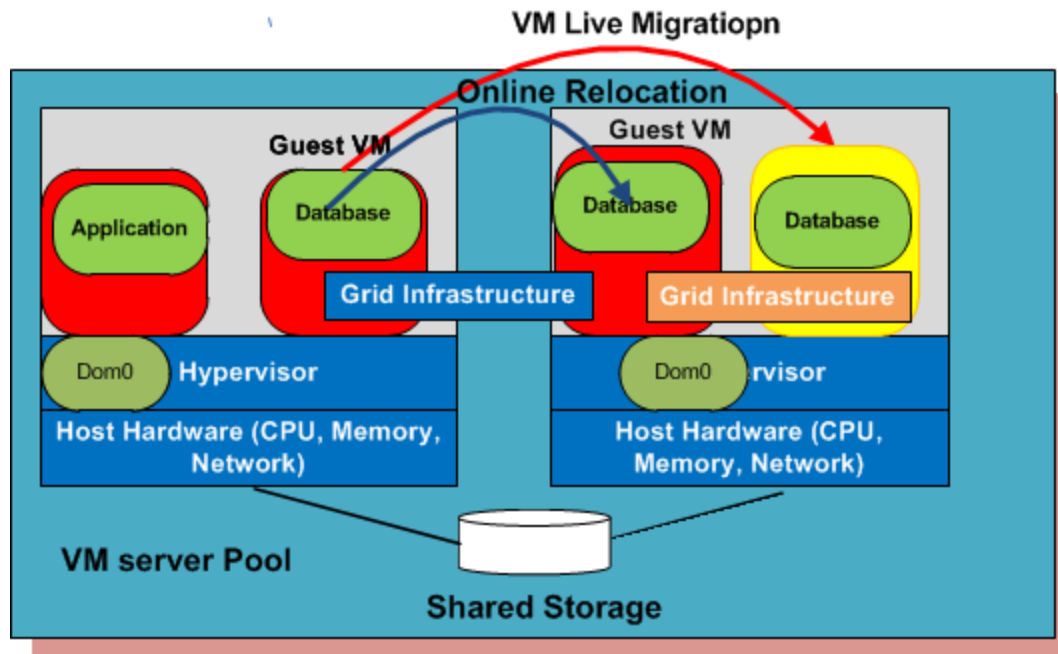
Database HA with Oracle VM

Oracle VM Live Migration:

Migrate VM from one physical server to 2nd physical server

RAC One Node Online Relocation

Relocate the database from 1st VM node to 2nd node (VM) instance





Thank You and QA

Visit my Oracle Blog at <http://kyuoracleblog.wordpress.com>

IOUG Virtualization SIG: www.ioug.org/virtualization

My Collaborate 12 Sessions and Panel

| | | |
|------------------------------------------------------------------------------------------|-----------------|------------|
| #920, Enterprise Manager 12c Deep Dive | Sunday 9:am-3pm | Mandalay K |
| #840, Alternative Options to Enhance the High Availability of your Oracle Database | Monday, 12:15pm | Palm C |
| #328, Virtualization Boot Camp: Virtualizing Oracle 11g/R2 RAC Database on Oracle VM 3.0 | Monday, 3:45pm | Palm D |
| #312, Configuring and Managing a Private Cloud with Oracle Enterprise Manager 12c | Tuesday, 4:30pm | Mandalay K |
| Oracle on Oracle VM - Expert Panel, Oracle VM Expert Panel | Wed, 9:30am | Mandalay L |



Thank You and QA

Visit my Oracle Blog at <http://kyuoracleblog.wordpress.com>

My Collaborate 12 Sessions and Panel

| | | |
|------------------------------------------------------------------------------------------|-----------------|------------|
| #920, Enterprise Manager 12c Deep Dive | Sunday 9:am-3pm | Mandalay K |
| #840, Alternative Options to Enhance the High Availability of your Oracle Database | Monday, 12:15pm | Palm C |
| #328, Virtualization Boot Camp: Virtualizing Oracle 11g/R2 RAC Database on Oracle VM 3.0 | Monday, 3:45pm | Palm D |
| #312, Configuring and Managing a Private Cloud with Oracle Enterprise Manager 12c | Tuesday, 4:30pm | Mandalay K |
| Oracle on Oracle VM - Expert Panel, Oracle VM Expert Panel | Wed, 9:30am | Mandalay L |