
Oracle RAC 11gR2 with Red Hat Enterprise Linux 6



Roger Lopez, Dell Inc.
Systems Engineer

Objectives

- Notable changes between RHEL5 and RHEL6 as they relate to Oracle 11gR2
- Deployment and Configuration of Oracle 11gR2 - Two Node RAC on Red Hat Enterprise Linux 6
- Best practices running Oracle 11gR2 RAC on Red Hat Enterprise Linux 6

Points of Interest – Notable Changes

- My Oracle Support ID 1089399.1
 - Official statement that Oracle will not provide ASMLib packages for kernels distributed by Red Hat as part of RHEL 6 or the Red Hat compatible kernel in Oracle Linux 6
- SysV Init replaced by Upstart
 - Oracle 11gR2 (as of 11.2.0.3) Installer still relies on legacy inittab
- Dell Oracle Deployment Utilities
 - GNU GPL Deployment Scripts and RPMs
- Automatic System Tuning for Database Storage
 - Tuned package recommended for auto tuning the system for common workloads



No more ASMLib - What Next?

- The Legacy Package Set
 - oracleasm-support
 - oracleasmlib
 - Oracleasm
- No ASMLib? No Problem!
 - Native RHEL mechanisms such as udev and device-mapper allow us to handle device maps and permissions with ease.
- It is important to note Oracle ASM and Oracle ASMLib are not the same.
- Oracle ASM is a volume manager for your Oracle data.
- Oracle ASMLib is a set of extensions used for device management.



Prerequisites

- Standard Package Set
 - cloog-ppl compat-db42 compat-gcc-34 compat-gcc-34-c++ compat-libcap1 compat-libstdc++-33 compat-libstdc++-33 cpp elfutils-libelf-devel gcc gcc-c++ glibc glibc-devel glibc-devel glibc-headers kernel-headers ksh libICE libSM libX11 libXau libXext libXi libXp libXt libXtst libaio libaio-devel libaio-devel libgcc libstdc++ libstdc++-devel libtool-ltdl libuuid libxcb mpfr ncurses-libs nss-softokn-freebl ppl readline unixODBC unixODBC unixODBC-devel unixODBC-devel

Prerequisites

- Kernel Tuning Parameters
 - Set our /etc/sysctl.conf parameters
 - › fs.file-max net.ipv4.ip_local_port_range
 - net.core.rmem_default net.core.rmem_max
 - net.core.wmem_default net.core.wmem_max kernel.shmall
 - kernel.shmmax kernel.shmmni kernel.sem
- Create Users and Permissions needed
 - Oracle and Grid
 - › Best Practices: Need UID/GID higher than 1000
 - › <add about user groups>
- This seems tedious, lets automate!

Automation – Dell Deployment TAR

- What's inside?
 - Dell Oracle Deployment Utilities RPM
 - Dell Validated RPM
 - Directory of drivers
 - › Drivers based on a quarterly release
 - › Validated in collaboration with extended peripherals teams
 - › Only replace native Enterprise Linux drivers where peripheral teams see fit
 - README
 - › Documentation for quick start use of contents of TAR
 - › References pointing to Dell's collaterals
 - › Wiki How-Tos and articles
 - › Links to Dell Whitepapers
 - › Information about the Dell Oracle SDL



Dell Validated RPM – Overview

- Dell's GNU GPL compliant fork of Oracle's Validated RPM
 - Motivation:
 - › Provide customers with an easy to deploy configuration based upon Oracle's Best Practices
 - › Give customers the option of deploying RHEL or Oracle UEK
 - › Closely based on Oracle Validated RPM
 - › Keeps low long term maintenance and overhead
 - › Incorporates Dell & Oracle's Best Practices
 - › Provide compatibility with Dell Tested and Validated Configurations



Dell Oracle Deploy – dodeploy

- Easy install via RPM
 - RPM Package built using Fedora Packager Guidelines
 - Depends on Dell-Validated RPM to confirm setup of system prerequisites
- Command line utility called 'dodeploy'
 - Designed to have look and feel of traditional GNU/Linux utilities
 - Set sysctl configuration parameters for the system based on Oracle Release
 - Set PAM Limits configuration for Oracle user if not previously set
 - Provide setup and configuration of Grid user
 - Grid software home setup
 - Group memberships
 - Permissions for those above
 - Optionally configure SELinux settings



Pre-Install Setup – Use the Automated Utilities

- Install the RPMs needed.
 - yum -y install dell-validated* --nogpgcheck
 - yum -y install dell-oracle-utilities* --nogpgcheck
 - yum -y install device-mapper-multipath
- Running dodeploy:
 - dodeploy -r 11gR2 -g



Identify Storage – ASM disks

- Simple Example: 1 Database Disk, 1 Vote Disk
 - Two LUNs, each with Two Paths
 - Subset of 'cat /proc/partitions' output:

8	16	524288000	sdb
8	48	524288000	sdc
8	32	524288000	sdd
8	64	524288000	sde



Identify Paths – Mapping Devices

- In this example the LUNs are equal in size, however this may vary upon your environment.

```
[root@rhel6 ~]# for i in sdb sdc sdd sde; do  
printf "%s %s\n" "$i" ;  
"$(scsi_id --page=0x83 --whitelisted --device=/dev/$i)" ;  
done  
  
sdb 360026b900061855e000008a54ea5356a  
sdc 360026b9000618571000008b54ea5360b  
sdd 360026b900061855e000008a54ea5356a  
sde 360026b9000618571000008b54ea5360b
```

udev Rules – Quick Notes

- Quick refresher:
 - udev supplies the system software with device events, manages permissions of device nodes and may create additional symlinks in the /dev directory, or renames network interfaces.
- udev rules are executed in enumerated order
 - 10-* rules are run before 20-*
 - device-mapper rules are 10-* rules therefore ours will be 20-*



udev Rules – Writing Ours

- Create an udev rule file
`/etc/udev/rules.d/20-redhat_dell_oracle.rules`
- Populate with the following:
 - `KERNEL=="dm-*" , PROGRAM="scsi_id --page=0x83 --whitelisted
--
device=/dev/%k" , RESULT=="360026b9000618571000008b54ea5360b`
 - `KERNEL=="dm-*" , PROGRAM="scsi_id --page=0x83 --whitelisted
--
device=/dev/%k" , RESULT=="360026b900061855e000008a54ea5356a`



Multipath – Path fail over + consistent naming

- Device-mapper-multipath - /etc/multipath.conf
 - Set defaults, blacklist, and multipaths
 - Need a multipath.conf
 - › man 8 mpathconf
 - › man 5 multipath.conf
- Copy the template and set friendly names
 - cp /usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf.defaults /etc/multipath.conf
 - user_friendly_names yes



Multipath – Defaults and Blacklists

```
defaults {  
    udev_dir          /dev  
    polling_interval  5  
    path_selector     "round-robin 0"  
    path_grouping_policy failover  
    getuid_callout  
    "/lib/udev/scsi_id --whitelisted --  
    device=/dev/%n"  
  
    prio              const  
    path_checker      directio  
    rr_min_io         1000  
    rr_weight         uniform  
    fallback          manual  
    no_path_retry     fail  
    user_friendly_names yes  
}  
  
blacklist {  
    wwid <local_disk_sda>  
    devnode "^(ram|raw|loop|fd|md|dm-  
    |sr|scd|st)[0-9]*"  
    devnode "^hd[a-z]"  
    devnode "^dcssblk[0-9]*"  
}
```



Multipath – Consistent Naming

- Configure the WWIDs for our paths with friendly names

```
multipaths {  
    multipath {  
        wwid      360026b9000618571000008b54ea5360b  
        alias     db  
    }  
    multipath {  
        wwid      360026b900061855e000008a54ea5356a  
        alias     vote  
    }  
}
```



Verification – Permissions and Aliases

- Confirm dm- devices have correct user:group perms

```
[root@rhel6 ~]# ls -l /dev/dm-*
```

brw-rw----.	1	root	disk	253, 0	Dec 6 16:05	/dev/dm-0
brw-rw----.	1	root	disk	253, 1	Dec 6 16:05	/dev/dm-1
brw-rw----.	1	grid	asmadmin	253, 2	Dec 7 14:16	/dev/dm-2
brw-rw----.	1	grid	asmadmin	253, 3	Dec 7 14:16	/dev/dm-3
brw-rw----.	1	root	disk	253, 4	Dec 6 16:06	/dev/dm-4

- Confirm multipath aliases

```
[root@rhel6 ~]# ls -l /dev/mapper/db /dev/mapper/vote
```

lrwxrwxrwx.	1	root	root	7 Dec 6 16:06	/dev/mapper/db ->/dm-3
lrwxrwxrwx.	1	root	root	7 Dec 6 16:06	/dev/mapper/vote ->/dm-2

Verification – Paths

- Confirm that the device-mapper-multipath daemon does in fact recognizes the multiple paths

```
[root@rhel6 ~]# multipath -ll

db (360026b9000618571000008b54ea5360b) dm-3 DELL,MD32xx
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 rdac' wp=rw
| +-+ policy='round-robin 0' prio=6 status=active
|   `-- 3:0:1:1 sde 8:64 active ready running
`--+ policy='round-robin 0' prio=1 status=enabled
   `-- 3:0:0:1 sdc 8:32 active ghost running

vote (360026b900061855e000008a54ea5356a) dm-2 DELL,MD32xx
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 rdac' wp=rw
| +-+ policy='round-robin 0' prio=6 status=active
|   `-- 3:0:0:0 sdb 8:16 active ready running
`--+ policy='round-robin 0' prio=1 status=enabled
   `-- 3:0:1:0 sdd 8:48 active ghost running
```



Installation Guides – Recommended Resources

- Dell Oracle Quickstart:
 - http://en.community.dell.com/techcenter/enterprise-solutions/w/oracle_solutions/1425.4-1-1-2-0-how-do-i-install-oracle-rac-11g-r2.aspx
- Red Hat Deployment Guide:
 - <https://access.redhat.com/knowledge/refarch/2012-deploying-oracle-11gr2-rhel-6>



Install – Install Oracle Grid and Database

- Skip ASMLib steps from installation documentation
- During Grid Installation section "Create ASM Disk Group", set the Discover Path to /dev/mapper/*

The screenshot shows the Oracle 11g Grid Infrastructure setup tool. The title bar reads "Create ASM Disk Group" and "ORACLE DATABASE 11g". On the left, a navigation pane lists steps: Download Software Updates, Installation Option, Installation Type, Product Languages, Grid Plug and Play, Cluster Node Information, Network Interface Usage, Storage Option, Create ASM Disk Group (which is selected), ASM Password, Operating System Groups, Installation Location, Prerequisite Checks, Summary, Install Product, and Finish. The main panel is titled "Select Disk Group Characteristics and select disks". It shows a "Disk Group Name" field with "OCR_VOTE", a "Redundancy" radio button group (Normal is selected), and a "Candidate Disks" radio button group (selected). A table lists disk paths, sizes, and statuses:

	Disk Path	Size (in MB)	Status
<input type="checkbox"/>	ORCL:DATA1	5119	Candidate
<input type="checkbox"/>	ORCL:DATA2	5119	Candidate
<input checked="" type="checkbox"/>	ORCL:DATA3	5119	Candidate
<input type="checkbox"/>	ORCL:DATA4	5119	Candidate
<input checked="" type="checkbox"/>	ORCL:OCR_VOTE01	1023	Candidate
<input checked="" type="checkbox"/>	ORCL:OCR_VOTE02	1023	Candidate
<input checked="" type="checkbox"/>	ORCL:OCR_VOTE03	1023	Candidate

A red box highlights the "Change Discovery Path" button at the bottom right of the table.

Post Installation – SysV Init to Upstart Migration

- Create a file called /etc/init.d/ohasd.conf with the following:

```
# As of Oracle 11gR2 there is the following line placed in /etc/inittab but is
# no longer in use as of RHEL6 because of the switch from SysV init to Upstart
### h1:35:respawn:/etc/init.d/init.ohasd run >/dev/null 2>&1 </dev/null

# The following configuration is simply the migration of that line to the new
# syntax

start on stopped rc RUNLEVEL=[35]

stop on runlevel [S01246]

respawn

script
    /etc/init.d/init.ohasd run >/dev/null 2>&1 </dev/null
end script
```



Post Installation – Setting HugePages

- RHEL 6 introduced the concept of Transparent Huge Pages (THP)
 - THP currently only maps anonymous memory regions (AnonHugePages) which only effects heap and PGA stack.
 - SGA needs shared memory, it does not use the anonymous memory region, therefore, it does not take advantage of the THP feature
- Recommendation
 - One must still set Huge Pages manually for SGA sizes bigger than 4GB
 - Make use of Oracle's HugePages script found in My Oracle Support ID **401749.1** - Shell Script to Calculate Values of Recommended Linux HugePages



Resources – Useful Online Documentation

- Dell Database Solutions Wiki
 - <http://delltechcenter.com/enterprise>
 - Engineering Pre-release of RHEL6 Oracle 11gR2 guide:
http://en.community.dell.com/techcenter/enterprise-solutions/w/oracle_solutions/3336.aspx
 - Current Releases page:
 - http://en.community.dell.com/techcenter/enterprise-solutions/w/oracle_solutions/1429.current-release.aspx
- Red Hat Documentation and Customer Portal
 - <http://docs.redhat.com/docs/en-US/index.html>
 - <https://access.redhat.com/knowledge/refarch/2012-deploying-oracle-11gr2-rhel-6>
 - <https://access.redhat.com/knowledge/techbriefs/storage-management-oracle-database-red-hat-enterprise-linux-6-using-asm-or-with>

Questions and Answers

Questions
or
Feedback?



Acknowledgements

- Adam Miller, Red Hat Inc



Thank you

