

# UXC RED ROCK CONSULTING

---

ALWAYS SEEK TO DELIGHT OUR CUSTOMERS

TRUSTED TO DO THE RIGHT THING

LISTEN BEFORE TALKING

GOOD ENOUGH IS NOT ENOUGH

WE LIKE TO HAVE FUN...SERIOUSLY!

# UXC RED ROCK CONSULTING

---

Presentation for NZOUG

Oracle Linux for DB and APPS

March 2013

## About the Speaker

- Alvaro Miranda from Chile
- Senior Consultant at UXC RedRock
- 15+ years of Linux experience
  - RedHat Instructor 2003 – 2011
- 9 Years DBA
- 5 Years Oracle e-Business Suite APPS DBA
  - Oracle Support ATG, Sev1
  - Experian US/UK, 11i -> R12 MAA global instance 15k users

## What we will do today?

- On this presentation, we will explore features present in Oracle Linux, that will enable us to get the maximum benefit for Databases and Applications

## Agenda

- Kernel UEK2
- Oracle Linux 6
- OS Update
- Clusterware and ASM
- Filesystems

## Kernel UEK2

## Why use Oracle UEK2 Kernel?

- Oracle UEK2 Kernel reduce the complexity of module dependencies for several components like ASM, BTRFS or OCFS2
- Used today in benchmark records and engineered systems
- One kernel, for bare metal or OracleVM virtual guest

## From oracle documentation

- Unbreakable Enterprise Kernel Release 2 includes performance and scalability enhancements that allowed Oracle to achieve record database benchmark results.
- Based on the 3.x mainline kernel, Unbreakable Enterprise Kernel Release 2 also includes improved memory and resource management, and is optimized to be deployed as a virtual guest.
- The Btrfs file system is now production-ready with this release.

## Kernel UEK2 version

- Kernel-uek-2.6.39
  - Number keep as 2.6 to avoid potential issues
  - Based on 3.x Kernel
- Actual version as today:  
kernel-uek-2.6.39-400.17.1.el6uek

## Actual Kernel UEK2 available

Name : kernel-uek                      Relocations: (not relocatable)  
Version : 2.6.39                        Vendor: Oracle America  
Release : 400.17.1.el6uek            **Build Date: Sat 23 Feb 2013 01:19:00 AM**  
Install Date: (not installed)        Build Host: ca-build44.us.oracle.com  
Group : System Environment/Kernel   Source RPM: (none)  
Size : 84878054                        License: GPLv2  
Signature : RSA/8, Sat 23 Feb 2013 04:10:12 AM UTC  
URL : <http://www.kernel.org/>  
Summary : The Linux kernel

## Switching to Kernel-UEK2

- Enable the repo UEK\_latest
- Yum install kernel-uek
- Reboot
- Review software that is kernel version dependent, ie hp support pack, emc powerpath, etc

## Kernel errata fixes

- Wouldn't be nice to be able to patch the kernel in memory for the errata fixes
- Online
- Without downtime

## What is Ksplice?

Ksplice Uptrack lets you apply **100% of the important kernel security updates without rebooting.**

Ksplice Uptrack is available for Oracle Linux, free of charge, for Oracle Linux customers with a Premier support subscription

Reboots avoided on EL 5 

Reboots avoided on EL 6 

## ol6\_playground

- If you will like to test a newer kernel, for
  - Some feature
  - Hardware compatibility
  - Close to upstream kernel

There's a channel for that: **ol6\_playground**

```
# rpm -qip ./latest/x86_64/kernel-3.7.10-3.7.y.20130305.ol6.x86_64.rpm
```

```
Name      : kernel                Relocations: (not relocatable)
Version   : 3.7.10                Vendor: Oracle America
Release   : 3.7.y.20130305.ol6  Build Date: Wed 06 Mar 2013 02:44:45 AM UTC
```

## Oracle Linux 6

## Release dates

- Oracle Linux 6 update 4:
  - released to the Unbreakable Linux Network on 2013-2-27
- Oracle Linux 6 update 3:
  - released to the Unbreakable Linux Network on 2012-6-27
- Oracle Linux 6 update 2:
  - released to the Unbreakable Linux Network on 2011-12-15
- Oracle Linux 6 update 1:
  - released to the Unbreakable Linux Network on 2011-05-29
- Oracle Linux 6 GA :
  - release to the Unbreakable Linux Network on 2011-02-03



## Why use Oracle Linux 6 ?

- Binary compatible with RedHat
- Single vendor for OS and Application  
<http://support.oracle.com>
- Pay support where you need support.
- Updates and errata available for paid and non-paid linux machines

All development, testing and certification occurs on Oracle Enterprise Linux

For Basic and Premier support customers :

Oracle Management Pack for Linux

Complete Linux server lifecycle management

Provisioning - Patching – Monitoring – Administration - Configuration - Management



- 24x7 global coverage, 145 countries, 27 local languages
- Dedicated engineering and delivery team
- Backed by world's largest support team
- Enterprise-quality, Lower cost
- Service Excellence in Integrated Services (TSIA)
- 2008 Service Excellence in Mission Critical Support Award (SSPA)
- 2006 J.D. Powers and Associates Global Certification for Outstanding Customer Support

## Lower time to deploy

OL4 or OL5	OL6
Start installer	Start installer
Setup installation	Setup installation + internal yum repo
Start installation	Start installation
First Boot	First Boot
Setup yum repository	Done
Update	
Reboot	
Second boot	
Done	

## Installation

- Test what works better
- On a 1 Gb network, net install is fast!
- Linux installer can be started from DVD, Network, pendrive/hard drive or at OracleVM guest creation
- Installation media can be on DVD, Network, pendrive/hard drive

## How much Swap we need ?

- Today memory standard is way bigger than few years ago
- Servers are being deployed with 64GB or more
- 144GB is becoming normal
- The 2x RAM or 1.5 times is not accurate today
- Thumb rule of swap is, if you use the swap, then you have a memory issue that need review.

## What Oracle RDBMS say?

### Oracle 11G R2

<http://www.oracle.com/pls/db112/homepage>

Database Quick Installation Guide for Linux x86-64

<b>Available RAM</b>	<b>Swap Space Required</b>
Between 1 GB and 2 GB	1.5 times the size of the RAM
Between 2 GB and 16 GB	Equal to the size of the RAM
More than 16 GB	16 GB

## Partition

- Keep the things simple

– OS DISK

/boot	/usr
/	/opt
/home	/var

– Oracle DISK

/u01
/u02

## device path

- Path are different
  - /dev/xvda for OracleVM
  - /dev/sda for normal scsi driver
  - /dev/cciss/c0d0 for hp local disk
  - /dev/powera for EMC Powerpath
  - /dev/mpath/<alias> for dm multipath
- For filesystem, use UUID
- For database, use asmlib

```
# grep UUID /etc/fstab
UUID=c30da00f-2591-4f64-8cb7-8edf938d7f8f / ext4 defaults 1 1
UUID=88331ff9-5c4c-4264-bbb4-77cdd4de2bed /boot ext4 defaults 1 2
UUID=259f7359-f79d-4468-bff3-2e5a4a949860 swap swap defaults 0 0

# oracleasm listdisks
DATA1
DATA2
# oracleasm status
Checking if ASM is loaded: yes
Checking if /dev/oracleasm is mounted: yes
# ls -al /dev/oracleasm/disks/
brw-rw---- 1 grid oinstall 8, 4 Mar 14 01:42 DATA1
brw-rw---- 1 grid oinstall 8, 20 Mar 14 01:42 DATA2
```

## User and groups

- If you use specific UID / GID
  - Create role separation user and groups before using
  - oracle-validated on ol5
  - or
  - oracle-rdbms-server-11gr2-preinstall on ol6

## RPM Packages

- **ID 401167.1** Defining a "default RPMs" installation of the Oracle Enterprise Linux (OEL) OS
- **ID 376183.1** Defining a "default RPMs" installation of the RedHat EL OS

“ .. de-selecting any "default RPM" groupings or individual RPMs can result in failed RDBMS installation attempts. Oracle recommends that you install the Linux operating system with the default software packages (RPMs). Do not customize the RPMs during installation.”

## RPM Packages

- **OL5**  
# yum install oracle-validated
- **OL6**  
# yum install oracle-rdbms-server-11gR2-preinstall

## OS Update

### We will update the OS?

- By experience, if you skip OS update..  
Will you update later?
- Be in control
- Exclude from auto update what might  
cause troubles
- Test !

## Yum repo

- With paid support:
  - Linux.oracle.com
- Without paid support:
  - public-yum.oracle.com
- Once configured  
# yum update

## Clusterware and ASM



## Considerations

- Clusterware (Grid, ASM) is a separate product from DB
- The same powerful motor behind RAC
- The same HA tool behind Oracle Restart

## U01 layout

Mount Point	USE
/etc/orainst.loc	global inventory location
/u01/app/orainventory	oracle global inventory
/u01/app/grid	GRID_BASE
/u01/app/11.2.0.3/grid	Grid Home
/u01/app/oracle	ORACLE_BASE
/u01/app/oracle/product/11.2.0.3/<dbhome_1>	Database Home

## Role separation

```
157189 0 drwxr-xr-x 1 grid oinstall 50 Feb 19 00:58 /u01/app/grid/Clusterware
176199 0 drwxr-xr-x 1 grid oinstall 180 Feb 19 01:17 /u01/app/grid/cfgtoollogs/netca
176200 4 -rw-r--r-- 1 grid oinstall 228 Feb 19 01:20 /u01/app/grid/cfgtoollogs/netca/netca_Ora11g_gridinfracome1-.log
176201 136 -rw-r--r-- 1 grid oinstall 135370 Feb 19 01:20 /u01/app/grid/cfgtoollogs/netca/trace_Ora11g_gridinfracome1.log
208203 34244 -rw-r----- 1 grid oinstall 35057792 Mar 14 03:05 /u01/app/grid/diag/tnslsnr/mirandaa00/listener/trace/listener.log
```

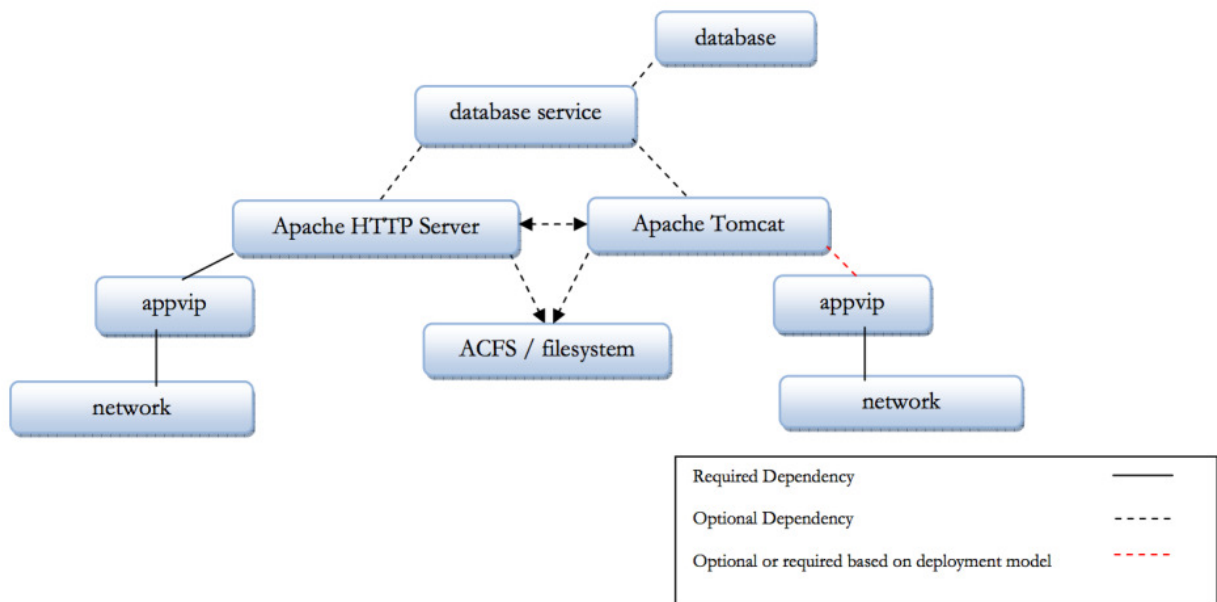
```
157195 4 -rw-r--r-- 1 root root 3074 Feb 19 00:59 /u01/app/11.2.0.3/grid/log/mirandaa00/client/clscfg1.log
157196 4 -rw-r----- 1 root root 914 Feb 19 01:02 /u01/app/11.2.0.3/grid/log/mirandaa00/client/crsctl_root.log
```

```
157334 4 -rw-r--r-- 1 root root 1805 Feb 19 01:02 /u01/app/11.2.0.3/grid/cfgtoollogs/patchauto2013-02-19_report.log
161775 20 -rwxrwx--- 1 grid oinstall 16505 Feb 19 01:02 /u01/app/11.2.0.3/grid/cfgtoollogs/patchauto2013-02-19_log
```

## Clusterware

- Customers with Linux support, can use Clusterware (Grid, CRS) to protect any application (restarting or failing over the application in the event of a failure)
- Same software, you already are familiar with:
  - VIP
  - ASM/ACFS
- Out of the box agents for Goldengate, Peoplesoft, Apache webserver and Tomcat

Resource Dependency Options



# Filesystems

## What we will check here?

- Oracle Linux includes several filesystem
  - Cluster filesystem, free to use
  - Filesystem with advanced features like Copy on Write, Snapshots

## How we will grow the filesystem? (on that Friday night patch/upgrade)

- Online or Offline ?
- Adding a new device?
- Expanding an existing device?
- Plan ahead

## U01 layout

Mount Point	USE
/etc/orainst.loc	global inventory location
/u01/app/oraInventory	oracle global inventory
/u01/app/11.2.0.3/grid	Grid Home
/u01/app/oracle	ORACLE_BASE
/u01/app/grid	GRID_BASE
/u01/app/oracle/product/11.2.0.3/<dbhome_1>	Database Home 1
/u01/app/oracle/product/11.2.0.3/<dbhome_2>	Database Home 2

## Create the base of U01

- As root:

```
#mkdir -p /u01/app/grid
#mkdir -p /u01/app/oracle
#mkdir -p /u01/app/11.2.0.3/grid
#mkdir -p /u01/app/oraInventory
#chown -R grid:oinstall /u01
#chown -R oracle:oinstall /u01/app/oracle
```

## Filesystems

Filesystem	Online Grow	Reflink	Snapshot	Cluster	Certified for datafiles
ext3	Yes	no	No	No	
ext4	Yes	no	No	No	
ocfs2	Yes	Yes	No	Yes	Yes
btrfs	Yes	Yes	Yes	No	

## OCFS2

- Shared Cluster Filesystem
- Free to use
- Support included with OL
- Used on OracleVM Server
- Can be used as local filesystem
- Support reflink

## reflink

- Similar to snapshot, but for files.
- As with snapshots. Quicker than a full copy, and save space.
- To put it simple, a reflink is a copy in time to a file (same content as the time we created the reflink)

## Ocfs2 reflink

- Inode based writable Snapshots
- Used by OracleVM as an option to clone Virtual Guests
- Used by OracleVM as an option to deploy template to guest
- `# reflink file1 file2`

Note: doesn't follow directory or wildcards

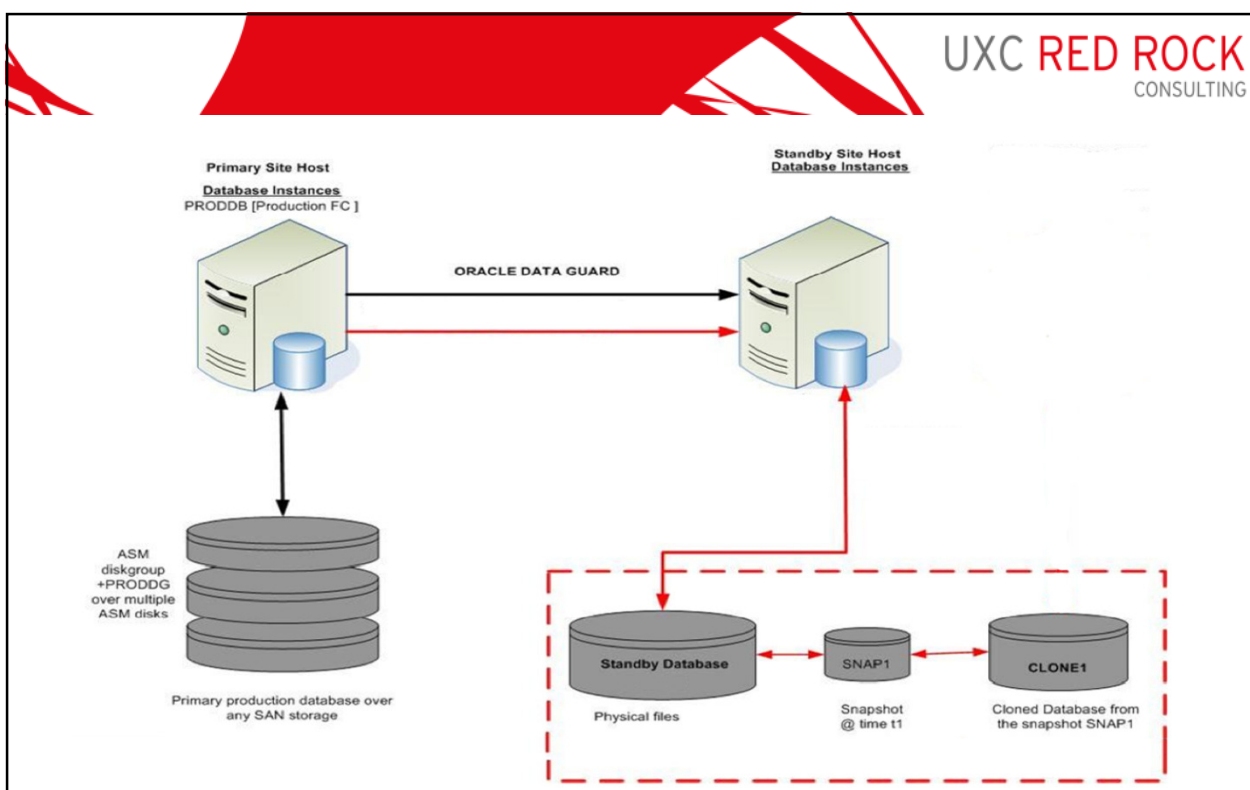
## To have a local OCFS2 FS + reflink

```
# mkfs.ocfs2 -F -T datafiles -M local --fs-features=refcount /dev/<device>
# vi /etc/fstab
    /dev/<device> /<mount_point> ocfs2 datavolume,nointr,noatime 0 0
```

## Business requirement

- From PROD DB
  - We need a clone for development
  - We need a clone for new development
  - We need a clone to troubleshoot a functional issue






UXC RED ROCK  
CONSULTING

## Ocfs2 reflink - cont

- Example use:
  - rman backup as copy database format '/u02/data/%b'

```
$ du -sh data clone
189G data      0 clone
$ time for x in $(ls data/*dbf); do
    reflink $x clone;/; done
real    0m0.622s   user 0m0.072s  sys 0m0.412s
[oracle@mirandaa00 apps_st]$ du -sh data clone
189G data      189G clone
```



## BTRFS

- **Btrfs** — New file system featuring error detection, correction, and tolerance which will provide enhanced scalability and simplified management for large storage configurations
  - ✓ Copy on Write – COW (reflink)
  - ✓ Subvolumes
  - ✓ Snapshot
  - ✓ Compression

## Btrfs reflink

```
$ du -sh dbhome_1*
6.3G    dbhome_1
$ time cp -ax dbhome_1/. dbhome_1_bkp_normalcopy
0.71 user 26.82 system 2:40.23 elapsed
$ time cp -ax --reflink=always dbhome_1/. dbhome_1_bkp_reflink
0.20 user 3.47 system 0:03.87 elapsed
```

## Btrfs reflink

```
$ du -sh cloudcontrol/
9.4G    cloudcontrol/
time cp -ax --reflink=always cloudcontrol/. cloudcontrol_bkp
real    0m50.225s
```

50seconds for a 9.4GB backup, with 13.000 directories, 77.000 files.

Anything went wrong with that patch?

```
$ mv cloudcontrol cloudcontrol.bad
$ mv cloudcontrol_bkp cloudcontrol
```

## BTRFS - Subvolume

- When mkfs.btrfs initialize a device. It create a default subvolume
- A subvolume will show up at OS as a folder
- Inside a subvolume, we can store files, folders and other subvolumes
- Integrated with os. Ie chown chmod mv

## BTRFS - Snapshot

- Is a copy of a subvolume
- As the original data is referenced, the snapshot is quick, and save space  
(poor man dedup)
- Once the snapshot is created, source and clone, are independent subvolumes.

## BTRFS - Examples

- Subvolume

```
[root@mirandaaeb ~]# btrfs subvolume list -p /u01 | grep r1213prd
ID 470 parent 319 top level 5 path app/oracle/r1213prd
ID 471 parent 470 top level 5 path app/oracle/r1213prd/apps
ID 473 parent 470 top level 5 path app/oracle/r1213prd/inst
ID 479 parent 470 top level 5 path app/oracle/r1213prd/db
# du -sh *
28G apps          6.3G db          18M inst
```

## BTRFS - Examples

- Snapshot

```
# time ./bkp_r1213prd.sh rockers
Create subvolume '/u01/bkp/r1213prd_rockers'
Create a snapshot of '/u01/app/oracle/r1213prd/apps' in '/u01/bkp/r1213prd_rockers/apps'
Create a snapshot of '/u01/app/oracle/r1213prd/db' in '/u01/bkp/r1213prd_rockers/db'
Create a snapshot of '/u01/app/oracle/r1213prd/inst' in '/u01/bkp/r1213prd_rockers/inst'

real    0m1.169s
user    0m0.001s
sys     0m0.036s
```

## As today:

- Btrfs module is included in kernel-uek from oracle.
- Btrfs-tools is available in oracle yum repo
- OEL 6.3 and 6.4 can be installed on BTRFS for root file system



**RED ROCK** CONSULTING  
A UXC Company

YOUR SUCCESS IS OUR SUCCESS