

RED HAT :: CHICAGO :: 2009

SUMMIT

Beginner's Guide to Running Oracle on RHEL

Brian Likosar (liko@redhat.com)
Solutions Architect, Red Hat
03-Sep-2009

presented by



Before we begin...

- Remember – this presentation is titled a “Beginner's Guide”
- Disclaimer – I am **not** an Oracle DBA
- Quick Audience Survey
- Bad news – no demo today
- “Cultural Issues” - let's get them out of the way

Where to start?

- Installation considerations
 - Database workload (heavy transactional versus data warehouse)
 - To virtualize – or not...?
 - Always use 64-bit
 - Database size – storage requirements and layout
 - Try to avoid LVM for storage for items besides the binaries
 - NFS rant (this can be done properly, but...)

Heavy Transaction vs Data Warehouse

- Is the application workload expected to be heavy read & write (AKA “transactional”)?
- Is this backing a web application?
- For transactional workloads – be sure to use the highest performing disk possible
- Is this going to be backing a data warehouse?
- For read-mostly workloads, lower performing disk might be acceptable

Virtualization

- The big question: is it a good idea to virtualize a database?
- All virtualization incurs overhead, period.
- Overhead is typically in the I/O space
- Databases are always I/O intensive
- That all being said – your judgment is best
- My recommendation is to consolidate more databases into existing environments

32-bit vs 64-bit

- Always use 64-bit
- 32-bit kernels have a concept of low memory and high memory, which can introduce artificial limits
- 32-bit kernels also have process limits for memory (of about 3GB)
- Most 32-bit kernels have been modified to support larger memory amounts (albeit in non-optimal fashion)
- More detail here: <http://kerneltrap.org/node/2450> and <http://www.redhat.com/rhel/compare>

Database Sizing & Storage

- The “old fashioned way” is to place everything inside /u01
- The old adage “as many spindles as possible” still applies, even with SAN disk
- You want to spread the data out as much as possible across controllers, disks, arrays, switches, etc.

The LVM Compromise

- As a systems administrator, you might want to set up Oracle filesystems under LVM
- As a database administrator, you will dread LVM for your archive, redo log, and data filesystems
- LVM causes overhead but does provide easier management
- How difficult is it (really) to have incremental filesystems?
- Leverage shell regular expressions `disk[0-3]` and the “seq” command

NFS Rant

- I **hate** NFS – it can be unreliable, unpredictable, and difficult to troubleshoot
- If you have a network hiccup (or a DNS one, or a lost route), the mount may hang
- If you insist on using NFS, please follow these recommendations:
 - Separate NFS network
 - Specialized NFS hardware (NetApp, etc.)
 - Use TCP, and tune it (rsize, wsize)

Preparing for Installation

- Ensure that storage is provisioned and has the correct permissions
- Determine if prerequisites are installed
(http://www.oracle.com/pls/db111/to_toc?pathname=install.111/b32285/toc.htm)
- Package listing for RHEL includes compilers – this might be against your company's security policy
- I recommend including the xauth RPM as a requirement so that you can forward your X sessions over SSH instead of using xhost as described in the documentation

Oracle Prerequisite Details

- Installation document talks about “kernel parameters”
- These are **key** to performance!
- Tread lightly until you research their potential effects
- See Scott Croft's presentation “Tuning RHEL for Oracle”
- Also see www.puschitz.com

Tools to ease preparation

- Oracle wiki has a kickstart configuration file that is useful as a starting point: <http://wiki.oracle.com/page/Linux+installation+kickstart+for+Oracle+database>
- To rapidly partition storage, leverage sfdisk
- Oracle RAC requires passphrase-less ssh connectivity between all nodes in a cluster, this snippet is useful:

```
cat /etc/hosts | grep -v '^#' | grep -v '^127\.0\.0\.1' | sed -  
re 's/[[[:space:]]+/,/g' | ssh-keyscan -t rsa,dsa -f - >  
~/.ssh/known_hosts
```

Tools to ease preparation (continued)

- Oracle's Metalink can be your best friend (or worst enemy) <http://metalink.oracle.com>
- If building an Oracle RAC cluster, the CVU can be helpful:
http://www.oracle.com/technology/products/database/clustering/cvu/cvu_download_homepage.html
- Oracle's installer warns you if have missed a prerequisite

Repeating Oracle Installations

- When preparing to automate Oracle software installations, bear in mind that different applications have different requirements
- Disk is “cheap” - install all options
- Don't have the installer create the default database
- When starting the installation, pass these flags:
 - `./runInstaller -record -destination /tmp/oracleinstall.rsp`
- To use that response file:
 - `./runInstaller -silent -responseFile /tmp/oracleinstall.rsp`
- For more information, see http://download.oracle.com/docs/cd/B14117_01/em.101/b12140/toc.htm

Comparison of Filesystems for Oracle Use

- Raw vs. “Cooked”
- OCFS/OCFS2
- GFS
- ASM

Filesystems

- Raw devices
 - “Deprecated” as of RHEL 5
 - Customer demand (literally due to Oracle) brought it back
 - Set up is somewhat cumbersome – documentation that exists contradicts itself
 - Working way is to leverage udev; however, this behavior is not as straightforward as you'd expect
 - If system crashes and Oracle cannot apply redo logs properly, nothing can be done – beware of fsck expectation (no such thing with raw)!

Filesystems

- “Cooked” filesystems (can we stop using this term?)
- Ext3
 - Excellent general purpose filesystem
 - Not intended or recommended for clustering
 - Scales well enough for Oracle filesystems
 - Beware of default “superuser” space – uses 5% by default for lost+found
 - Recommended for quick/simple setups (development, labs, etc.)
 - Can be checked for consistency in case of system crash

Filesystems

- OCFS
 - Decent first attempt at a clusterable filesystem
 - Not POSIX compliant, not upstream
 - Not recommended whatsoever
- OCFS2
 - Fixes flaws of OCFS
 - Generally used for Cluster Registry (OCR) and Voting disks with Oracle 10g
 - Recommended **only** for OCR and vote disks

Filesystems

- GFS (including GFS2)
 - Works excellently as a scalable clustered filesystem
 - Generally intended for large installations (more than 8 nodes)
 - Can be used for any files
 - Requires Red Hat Cluster Suite
 - Somewhat complicated, and leaves control of the system in the System Administrator's hands

Filesystems

- ASM
 - An additional RPM which can be downloaded from Oracle's web site: <http://www.oracle.com/technology/tech/linux/asmlib/install.html>
 - Manages disk slated for Oracle use as a database instance
 - Recommended for all setups – easiest for DBAs and Systems Administrators
 - Ensure persistent naming by leveraging `/etc/multipath.conf` parameters

System Startup and Shutdown

- Oracle has historically not included scripts to properly start and stop the databases in a system standard fashion
- The following slide includes a startup and shutdown script that can be rather easily integrated into your environment
- Enable them with chkconfig:
 - `chkconfig add oracledb`
 - `chkconfig oracledb on`

/etc/init.d/oracledb

```
#!/bin/sh
# chkconfig: 345 99 10
# description: Oracle auto start-stop script.
#
# Change the value of ORACLE_HOME to specify the correct Oracle home
# directory for your installation.

ORACLE_HOME=/u01/app/oracle/product/10.2.0/db_1
#
# Change the value of ORACLE to the userid of the
# oracle owner at your site.
#
ORACLE=oracle

PATH=${PATH}:$ORACLE_HOME/bin
export ORACLE_HOME PATH
#
case $1 in
'start')
    su - $ORACLE -c "$ORACLE_HOME/bin/dbstart $ORACLE_HOME"
    ;;
'stop')
    su - $ORACLE -c "$ORACLE_HOME/bin/dbshut $ORACLE_HOME"
    ;;
*)
    echo "usage: $0 {start|stop}"
    exit
    ;;
esac
exit
```

QUESTIONS?

**TELL US WHAT YOU THINK:
[REDHAT.COM/SUMMIT-SURVEY](https://redhat.com/summit-survey)**