

Hit the Ground Running

Oracle 10gR2 RAC on Linux

Chris Page (cpage@cptech.com)

Corporate Technologies
(<http://www.cptech.com/>)

December 12, 2006



CORPORATE TECHNOLOGIES
Tough IT. Solved.

Agenda

- Concepts and Definitions
- Hardware and Software Requirements
- Overview of the Installation Process
- Some Useful Links
- Best Practices
- Recommended Reading Topics



Oracle 10gR2 RAC Concepts and Definitions

- RAC – Real Application Clusters
 - One Database, Many Instances
 - Shared Storage Architecture
- Database
 - What persists when everything shuts down
- Instance
 - The Processes and Memory on a node accessing the database
- Clusterware
 - Oracle's cluster management services
- Listener
 - Client network access to instances
- Flash Recovery Area
 - The place storing all files needed for database recovery



Oracle 10gR2 RAC Concepts and Definitions (2)

- Automatic Storage Management
 - AKA Oracle's Logical Volume Manager
 - Stripes across raw devices
 - Requires a dedicated ASM Instance per node
- ASMLib (optional with ASM)
 - Support Libraries for ASM devices
- Oracle SID
 - Identifier that uniquely defines an instance
- Oracle Services
 - Defines groups of instances
 - Used for workload management
- Cache Fusion
 - Oracle RAC's mechanism for cache coherency



Planning: Storage Requirements

- Shared Storage
 - SCSI-3 Reservations or NFS Certified Appliance
 - Device Naming Persistence
 - Configure Raw devices if needed
- Binaries, Inventory, and Trace Files
 - local or shared (Cluster File System)
- Oracle Cluster Devices
 - Cluster Registry and Voting disks [raw or NFS]
- Database Files
 - Data files, control files, spfile, online redo log files
 - Standard Edition: Must use ASM (Automatic Storage Management)
 - Enterprise Edition: NFS, Raw, CFS, or ASM
- Backup Files, Flash Recovery Files, Archive Log Files, etc
 - NFS, CFS, ASM, or local (not recommended)



Node Requirements

- All Node configurations must have the same...
 - Same Architecture and OS
 - Same Network Interface Names
 - Same Disk Device Names
 - Same uid and gids for Oracle user
 - Directory Structure
- Except the following is permitted...
 - Different Number and Speed of CPUs
 - Different Memory Sizes



Network Requirements

- Public and Private Networks
 - Private Interconnect
 - GigE is popular choice
 - No cross-over cables
 - UDP (Cache Fusion) and TCP (CRS)
 - Can be bonded
 - Same subnet throughout cluster
 - Public Network
 - Same subnet throughout cluster
- Virtual IP (VIP) Addresses
 - Listeners listen on and redirect to VIPs
 - Each node has its own VIP
- Three Addresses for each Node
 - Node : host address
 - Node-priv: host interconnect address
 - Node-vip: host virtual IP address



Software Requirements

- Certification Matrix
- The Linux Choices
 - SuSe or RHEL
 - “Unbreakable” Enterprise Linux
- Oracle Clusterware
- Oracle Standard or Enterprise Edition with RAC option
 - SE: Free but limited RAC License and must use ASM
 - EE: No limit to number of CPUs, but need to purchase RAC licenses



Useful Links

- Oracle Support (<http://metalink.oracle.com>)
 - Certification Matrix
- Oracle Technology Network (<http://otn.oracle.com>)
 - Software Downloads
 - Guides
 - Quick Start Installation
 - Installing Oracle RAC
 - Installing Oracle Database
 - Reference material
 - Read the release notes



Some More Links (<http://www.cptech.com>)

- Certification Matrix
 - <http://metalink.oracle.com/> (Click on Certification tab)
- Oracle Database Software
 - <http://www.oracle.com/technology/software/products/database/oracle10g/htdocs/10201linuxsoft.html>
- Oracle Enterprise Linux
 - <http://www.oracle.com/technologies/linux/index.html>
- Oracle 10gR2 Documentation
 - <http://www.oracle.com/technology/documentation/database10gr2.html>
- Oracle 10gR2 RAC Installation Guide
 - http://download-east.oracle.com/docs/cd/B19306_01/install.102/b14203/toc.htm
- Oracle 10gR2 on Linux Installation Guide
 - http://download-east.oracle.com/docs/cd/B19306_01/install.102/b15660/toc.htm
- Oracle 10gR2 on Linux Release Notes
 - http://download-east.oracle.com/docs/cd/B19306_01/relnotes.102/b15659/toc.htm
- Oracle Cluster Verification FAQ
 - http://www.oracle.com/technology/products/database/clustering/cvu/faq/cvu_faq.pdf
- Oracle RAC Deployment Guide
 - http://download-east.oracle.com/docs/cd/B19306_01/rac.102/b14197/toc.htm
- Installing Oracle on Linux Walk-Through (non-RAC) - background
 - http://www.oracle.com/technology/pub/articles/smiley_10gdb_install.html
- Oracle 10gR2 RAC installation on Firewire Walkthrough - background
 - http://www.oracle.com/technology/pub/articles/hunter_rac10gr2.html
- Oracle ASM Intro
 - http://www.oracle.com/technology/oramag/webcolumns/2003/techarticles/scalzo_asm.html
- Oracle ASMLib
 - <http://www.oracle.com/technology/tech/linux/asmlib/index.html>



Installation Overview

- Pre-Installation Prep
- Clusterware Installation
- DB Software Installation
- Database Creation



Pre-Installation Steps

- Configure Network
 - Private network bonding
 - DNS Entries
 - /etc/hosts
- Prep the OS
 - Required RPMs and System Configuration File Changes
 - Create user, groups, and environment variables
 - SSH user equivalence across all nodes in the cluster
 - Check 'ssh {nodeN} date' and 'ssh {nodeN.domain} date' as oracle user
 - Configure hangcheck timer
- Configure Disks (see earlier storage slide)
 - Create ASM Disks via ASMLib if using ASMLib
- Get Software
 - Download and Unpack Clusterware and Database Software
 - Important: Two Oracle Homes!



Cluster Installation

- Create a directory for the Cluster software (\$ORA_CRIS_HOME)
- Install Clusterware binaries into \$ORA_CRIS_HOME
 - {cluster-sw}/rootpre/rootpre.sh as root
 - {cluster-sw}/runInstaller –record –destinationFile /tmp/clus-sw-install.rsp
 - Creates startup files (/etc/init.d/init,crs etc)
 - Run root scripts when prompted on specified nodes in order
 - Execution on first node will initialize the voting and CRS devices
 - Execution on last node will configure startup the node applications via 'vipca'
- Installer copies files via scp
 - Verify as this step will fail silently



DB Software Installation

- Create a directory for the database software (\$ORACLE_HOME)
- Install Database binaries into \$ORACLE_HOME
 - `{database-sw}/runInstaller -record -destinationFile /tmp/db-sw-install.rsp`
 - Run root scripts when prompted on specified nodes in order
- Installer with replicate binaries to remote nodes for 'local' \$ORACLE_HOMEs
- Configure Oracle Network Using Network Configuration Assistant
 - Run `$ORACLE_HOME/bin/netca`



Database Creation

- Configure ASM
 - Run `$ORACLE_HOME/bin/dbca`
 - select Configure ASM
 - Provide path to directory with raw devices when prompted
- Create One or More Databases using Database Configuration Assistant
 - Run `$ORACLE_HOME/bin/dbca` and select Create Database
 - Set `DBCA_RAW_CONFIG` if using Raw Devices
 - Indicate that this is a cluster installation
- Configure Client
 - Make sure client is connecting to VIP and specifies a Service rather than a SID



Oracle's Cluster Verification Utility (cluvfy)

- Oracle utility that tries to do just what it says
- Attempts to identify problems early and throughout the installation process
 - Verifies node connectivity
 - Verifies
- Find it at `{cluster-sw}/cluvfy/cluvfy.sh`



Best Practices

- Synchronize time (ntp)
- Cluster Verification Utility
 - Use it each step of the way
- During Cluster installation: Check for failure of 'vipca' during 'root.sh'
 - Will fail if using non-routeable addresses
 - Run 'vipca' by hand as root before continuing
- NFS
 - Use recommended mount settings for data
 - Use "normal" mount settings for Oracle binaries
- ASM
 - Use ASM to stripe across RAID 1 devices
- CFS
 - Don't use Oracle CFS for data files



Best Practices (2)

- Host Name Resolution Peculiarities
 - /etc/hosts
 - Don't have the nodes name on localhost line entry
 - List FQDN before 'shortname'
 - Verify 'hostname' returns FQDN
 - Have all addresses in DNS
- When using dbca...
 - Don't create sample schemas
 - Don't configure for shared server (MTS)
- Monitor Flash Recovery space warnings in alert file
 - 2G by default and can quickly fill with archive logs and halt system



Recommended Topics

- Backup and Recovery
 - Recovery Manager (rman)
 - Cluster Device Backup (ocrconfig)
- Automatic Storage Management (ASM)
 - ASMLib (not required for ASM)
- Interacting with RAC
 - 'srvctl' – database control
 - 'emctl' – enterprise manager control
 - 'crsctl' – cluster management
- Get to know your log files
 - \$ORACLE_BASE/admin/{DB}/bdump/
 - \$ORA_CRS_HOME/log/
- Oracle Enterprise Manager (aka dbconsole)



Questions

- Contact:

- Chris Page (cpage@cptech.com)

- <http://www.cptech.com/>

