



# Scalable Oracle 10g for the Physics Database Services

---

Luca Canali, CERN IT  
January, 2006

# Outline




- Oracle 10g RAC main architectural components
- Installation overview
- ASM examples
- Utilities and Benchmark Tools
- Q/A

# Scalable Oracle Architecture



**Goal:** A database infrastructure that provides the required system resources to the end-users and applications.

**How:** A modular architecture that can scale up to a large number of components



# Architectural Elements



- **Oracle RAC for HA and Performance**
  - Multiple nodes
  - Services to balance and segment workload
  - Clients configured with TAF
- **Low cost scalable storage**
  - SAN infrastructure
  - Oracle's ASM (volume manager and RAID)
- **Network infrastructure**
  - Cluster interconnects
  - Public network

# Deployment Elements



- Linux
  - Oracle certified Linux distribution RHEL 3 or 4
  - Processor architecture
    - At CERN so far 32 bit and Intel Xeon
    - Planned: 64 bit architectures
- Monitoring and Manageability
  - EM for monitoring
  - Additional custom scripting/alerts
  - Backups with RMAN
    - to tape
    - to disk (flash backups)

# Pre-Installation



- SAN configuration
  - Attach storage arrays to the Linux servers (SAN network)
  - Configure the HBA drivers (possibly using multipathing)
- Configure TCP/IP networks
  - Cluster interconnects
  - Public networks (possibly using teaming)
- Setup ssh equivalence (logon without pass) for the cluster nodes
- Create the necessary disk partitions
- Setup raw devices
  - For RHEL 3 use devlabel to map rawdevices to obtain persistency
- Setup asmlib to map disk partitions for ASM
- Check OS prerequisites on Oracle's installation guide (kernel parameters, etc)

# Installation



- Install CRS
- Install RDBMS
- Install latest Patchset
- Install 'CPU' (latest security update)
- Install EM agents
- Setup ASM (dbca can do it, but double check the results)
  - ASM instances
  - ASM diskgroups
- Create DB

# Useful Tools



- **Performance tools:**
  - IO tests: [ORION](http://otn.oracle.com) from [otn.oracle.com](http://otn.oracle.com)
  - Hammerora ([hammerora.sourceforge.net](http://hammerora.sourceforge.net))
  - Bytemark (nbench port by U. Mayer )
- **Oracle RAC on VMWARE from [OTN.oracle.com](http://OTN.oracle.com)**
- **Utilities and time savers**
  - pconsole: allows to send keyboard input to multiple cluster nodes shell.
  - rlwrap: command history to sqlplus for Linux
  - Oracle's CVU (cluster verification utility)




# ASM's Configuration – Examples 1




- ASM is a volume manager, its output are disk groups (DG) that Oracle databases can mount to allocate their files

DATA-DG

RECOVERY-DG



**Config 1:** Disk groups created with dedicated disks

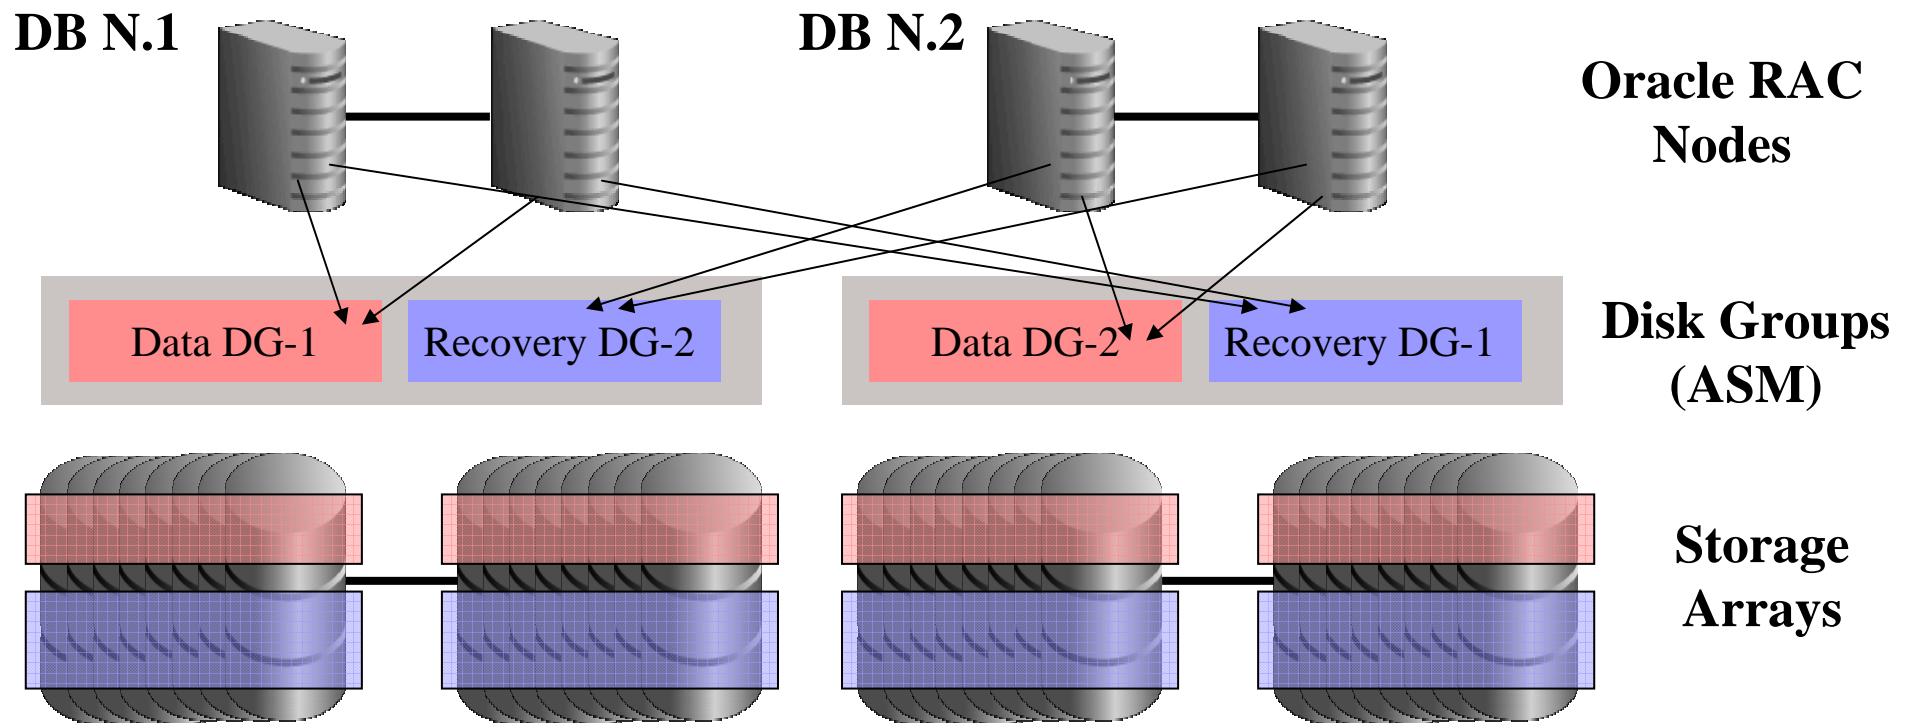


**Config 2:** Disk groups created by 'horizontal' slicing

# ASM Configuration –Examples 2



- ‘Coupled-RACs’ storage configuration:
  - High availability
  - High performance
  - DBs have dedicated resources
  - Allows backups to disk
  - Allows clusterware mirroring (10.2)



# Conclusions



- See also our wiki page
  - <https://uimon.cern.ch/twiki/bin/view/PSSGroup/DbArea>
    - Installation procedures
    - Post install actions
  
- Q/A