CERN-IT Oracle Database Physics Services

Maria Girone, IT-DB

13 December 2004

Outline

- Oracle Physics Services
  - Service Overview
  - Current Limitations
- Evolution of the Physics Services – 2005 and beyond
  - Moving to RAC on linux
- Proposal for lifecycle of key applications
- Conclusions

Service Overview

- Provide database and data management services for
  - >10 dedicated disk servers (COMPASS, HARP, Cristal II, CMS Tracker, …)
  - Replica Location Service (RLS), 6 VOs
    - One WN as AS per VO (6 VO) for production service
    - Two DS, hosting the DB, coupled via hot standby
  - Sun Cluster (public 2-node for applications related to detector construction, calibration and physics production processing)
- Total volume ~500 GB
- Offer production quality services (24 x 7 operation)
  - crucial for the experiments production and grid related applications

Support team

- 24 x 7 support, two levels of expertise
  - 1st line on-call team (10 people)
  - 2nd line on-call team for expert DB and AS administrators (3 people)
  - Weekly shifts, two people
  - On duty procedures on the web page
- Mandate
  - Implement procedures for high availability database and application server setups
    - Oracle technologies used: database clusters (RAC), data protection (Data Guard), Oracle streams
  - Standardise procedures and techniques for monitoring, backup and configurations
  - Consultancy for application design, development and tuning

Evolution of the Physics Services

- Current systems not scalable to initial exploitation phase of LHC
  - Disk server poor match for DB needs; Sun Cluster under-configured
- ☺ Tests of Oracle 10g RAC on Linux, as proposed to PEB, promising

- Uncertain physics services requirements
  - storage volume
  - resource split among different services
  - access I/O patterns

- Main goals of replacement service:
  - **Isolation** – 10g 'services' and / or physical separation
  - **Scalability** - in both database processing power and storage
  - **Reliability** – automatic failover in case of problems
  - **Manageability** – significantly easier to administer than now

- Target date for RAC based service: Summer 2005

RAC work-plan

- Hardware has arrived and is under acceptance testing (IT/FIO)
- Main RAC work plan items
  - RAC functionality & config tests: Q1 2005
  - RAC stability and work load: Q2 2005
  - Establish services on key applications: Q3 2005
  - Migration of all applications: Q4 2005
- Expected to be completed by June 2005
  - Regular progress reports to the PEB

Current problems and limitations

- High load on the Sun cluster has caused performance degradation of key production applications, caused by
  - increased use of existing applications
  - new not fully optimized applications

- Current cluster setup couples applications of different experiments and priority

- Proposed Stop-gap solution until RAC becomes available
  - Identify key applications which either are high priority or large resource consumers
  - Allocate dedicated resources (disk servers for now) to service these applications
  - Introduce a better defined policy to insure the proper planning and optimisation for key application

Proposed Lifecycle for Key Applications

Following Fermilab model, we propose

- Development Service
  - Use the development service offered by IT-DB (limited service and data volumes)
- Integration and Optimization Service
  - Deployment includes a realistic work load for definition of the performance metric of the new service and optimization in collaboration with IT/ADC and IT/DB
- Production Service
  - full production quality service, including backup, monitoring services, on call intervention procedures

- Requests for high priority key applications also coordinated with COCOTIME request

Conclusions

- Oracle deployment needs at CERN are rapidly ramping up
  - Raising need for consultancy, deployment resources
  - Investigate RAC/linux to achieve scalability
  - Need deployment policy to ensure proper planning of hardware and consultancy
- CERN Tier 0 service will be an important participant in the LCG distributed database infrastructure
  - But current resources don't leave much space for additional tasks
- Current performance limitations will require stop-gap solution to ensure guaranteed resources to production applications