Oracle for Physics

Services and Support Levels

physics-database.support@cern.ch

Maria Girone, IT-ADC

24 January 2005

Outline

- Database Services for Physics at CERN
- Current problems and limitations
- Scalable service for the LHC
- Planning the deployment of physics applications
  - Defining applications requirements
  - Optimization
  - Resource allocation
- Conclusions

Database Services for Physics

- Mandate
  - Coordinate the deployment of physics database applications
  - Administers the physics databases in co-operation with the experiments or grid deployment teams
  - Consultancy for application design, development and tuning
- Provide database services for LHC and non-LHC experiments
  - Grid File Catalogs
  - Sun Cluster (PDB01)
    - public 2-node for applications related to detector construction, calibration and physics production processing
  - ~30 dedicated Linux disk servers, including COMPASS and HARP

Service People and Responsibilities

- Physics Database Services and CERN T0 in 3D: Maria Girone
- Link people to experiments/projects:
  - Grid Deployment & Middleware: Miguel Anjo
  - ALICE: Jacek Wojcieszuk
  - ATLAS: Dirk Geppert / Ioannis Papadopoulos (POOL)
  - CMS & COMPASS: Giacomo Govi
  - LHCb & HARP: Andrea Valassi
- Openlab Oracle Fellow: Marta Jakubowska-Sobczak

Current Problems and Limitations

- Service requirements from experiments will increase for LHC start-up preparation
  - Requirements still uncertain in some areas
  - Need to plan for a scalable service
- Current infrastructure based on Sun cluster is already not sufficient to achieve required performance and application isolation
  - Linux based solution with ORACLE Real Application Cluster looks promising
  - Need to provide guaranteed resources to key applications
    - Need to identify them with experiments
    - Deploy them only after a validation phase

Towards a scalable service for LHC

- 10g RAC/Linux based service may provide:
  - 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
- Hardware in place, under acceptance tests for individual components..
  - We set-up a common work-plan across several IT groups
- ..but target date for pre-production is summer 2005
  - Need a stop-gap solution until then

RAC Work plan

- Main items view
  - RAC assembly & config tests: Q1 2005
  - RAC optimization & stability test: Q2 2005
  - Ramping up to production phase: Q3 2005
  - Migration of all applications: Q4 2005

- Web site: http://cern.ch/it-adc/phydb/rac/

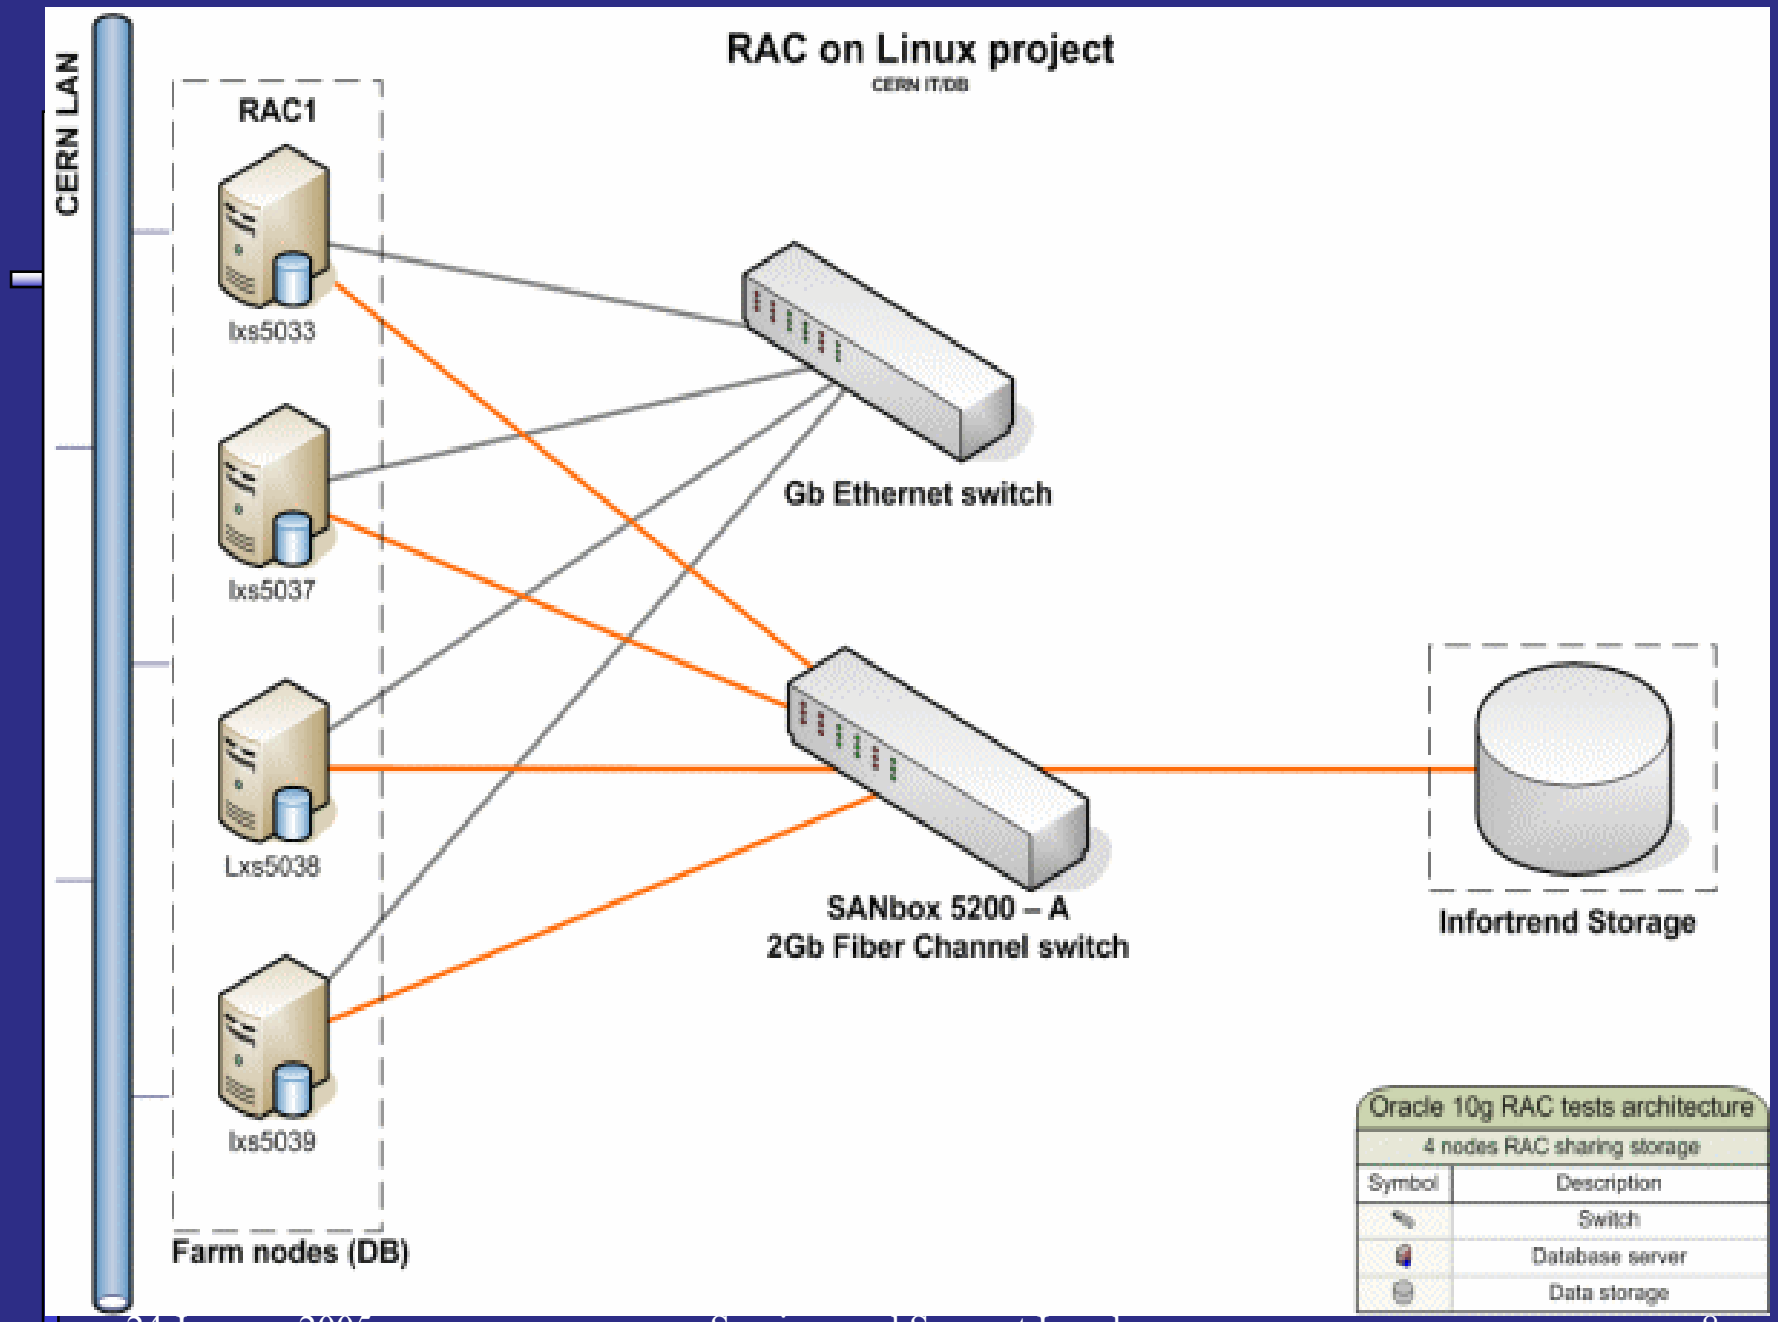
RAC on Linux project
CERN IT/DB
CERN LAN
RAC1
lxs5033
lxs5037
Lxs5038
lxs5039
Farm nodes (DB)
Gb Ethernet switch
SANbox 5200 – A
2Gb Fiber Channel switch
Infortrend Storage
Oracle 10g RAC tests architecture
4 nodes RAC sharing storage
Symbol
Description
Switch
Database server
Data storage

Stop-gap Solution

- Several applications which either are high priority or large resource consumers have been identified
  - Priorities are set by the experiments
- Allocate dedicated resources to service these applications
  - One box per experiment for now - propose to move to Oracle 10g
  - Once ready move to a defined slice of the new service cluster
- Application owners have been contacted to prepare the move

Planning for Application Deployment

- Introduce a better defined deployment process to insure
  - Proper planning of database capacity (volume & CPU)
  - Insure the optimisation of key applications before production starts
- Propose to separate between two basic application types
  - Resource consuming applications
    - Eg >20% of CPU, >20% of server memory, >20 concurrent sessions
    - Provide application isolation
  - Standard applications
    - Smaller database applications which can run in a shared service

Service Levels

- Following Fermilab model, we propose
  - Development Service
    - Code development, no large data volumes, no backup
  - Integration and Optimization Service (for key apps)
    - Sufficient resources for larger tests and optimisation
    - Allocated together with consultancy manpower in advance: need 3 month notice
  - Production Service
    - full production quality service, including backup, monitoring services, on call intervention procedures
    - Monitoring to detect new resource consuming applications or changes in access patterns
- All Physics Database Service will be based on Oracle 10g

A Typical Application Deployment Cycle

- Development and definition of application requirements
  - Application development will take place as usual using the database development service
  - Once sufficiently stable application code moves to a validation phase
- Schema validation and optimization phase
  - Performance requirements are needed for key apps
    - How many operations? Which query mix? From how many concurrent clients?
    - Experiments should provide test work load
  - Work load is run in a specialised validation service until required performance is achieved
- Resource Allocation & Planning
  - Optimised application with resource requirements is used to estimate h/w resources required for the service
  - Total database volume requirements are handled via COCOtime

Conclusions

- Oracle deployment needs for physics are rapidly ramping up
  - Raising need for consultancy, deployment resources
- Developing a new service to achieve scalability, availability, isolation
- Current performance limitations will require stop-gap solution to ensure guaranteed resources to production applications
- Resource consuming applications need to be identified early in order to schedule validation and optimisation and to allocate guaranteed resources