Site Services and Policies Summary

Dirk Düllmann, CERN IT

More details at http://lcg3d.cern.ch

Tier 0/1 Database Service

- Significant Oracle services and experience exist at several (but not all) LCG Tier 1 sites
  - Proposed (homogenous) database services for Tier-0 and Tier-1 could be implemented
    - Assuming adequate resources
- Agreement on software setup seems not a major issue
  - Eg Oracle 10g for 3D prototyping and deployment next year
  - Even though application of patches and upgrades needs an greed procedure
- Hardware setups and backend storage are different
  - … and will likely continue to be different
  - Impact 3D “shared” DBA services needs to be taken into account

Tier 2 - Database Service

- No major application side conflict with the proposed very limited Tier 2 database service
  - No media recovery or backup
  - Local queriable cache (of slices) of Tier0/1 data
    - Writes are possible but propagation of changes to Tier 1/0 are not provided by the service
    - Restricting to FroNtier capabilities may be a good idea
  - Consequence of available manpower
- Completely automated installation and pre-configuration
  - ATLAS deploys MySQL rpms successfully
  - What is required in addition to that
- U Chicago agreed to act as prototype Tier 2 in 3D
  - Need a Tier 2 application test plan
  - …after agreement on a plan on the Tier 1 level
- Proposal to Tier 2 community after successful tests

Tier 1 Service Split

- Discussion only just starting
  - Draft proposal based on input received from FNAL (Anil Kumar)
- Local Services
  - Server installation, bug & security fixes
  - OS patches/upgrades
  - Backup/recovery support
  - Data migration (between db servers)
- Shared Services
  - Db and OS accounts & privileges
  - Storage support (adding more space)
  - Monitoring
    - DB alerts, “killer queries cron job output
    - Host system load, space thresholds
  - Performance problems & optimization
- Site Communication
  - Proposal to setup a shared (web based) Log-Book, mailing lists
  - Need to establish regular DBA meeting
  - eg as part of weekly/bi-weekly 3D meetings

3D Resources..

- 3D will soon need to ask for resources to schedule task and provide hardware resources
  - My personal impression:
    - Hardware costs are not too significant for 2005
    - Available manpower resources _are_ especially during the early deployment phase of new database applications!
- Need to review manpower plan in overall LCG database area to plan for a deployable LCG database service in 2005
  - Not sure that the development effort on the experiment and project sides is matched with required consultancy resources for controlled deployment
- Realistic planning which applications go into production by when will be key to avoid overloading the database services teams at all tiers
  - Deploying new applications without proper integration and optimization is not a solution to this
  - Asking for more flexibility from the application side (eg more DB flavors deployed in parallel, more access protocols, more complex services) won't really help either..

Expected Services Requests in 2005

- Expect significant workload load on the database service teams
  - DB applications for physics community are rapidly ramping up in volume and performance requirements
  - Ongoing R&D to achieve required database scalability and availability at CERN and several T1 sites
  - LHC application development support at FermiLab and CERN
    - Many new applications are planned to move from development into Tier0 deployment with the help of DBA teams
- 3D services depend on a reliable Tier 0 service
  - Resource allocation at CERN needs to insure that priorities for Tier 0 are preserved…
  - …but provide sufficient coverage for coordination of distribution technology tests in parallel to T0 integration

Database Service Policies

- Several sites have deployment policies in place
  - E.g. FNAL:
    - Staged service levels
      - Development -> Integration -> Production systems
    - Well defined move of new code / DB schema during development process
      - Apps developers and DB experts review and optimize schema before production deployment
- Similar policy proposal prepared for CERN physics database services
  - To avoid recent interference between key production applications of different experiments on shared resources
    - Caused by missing indices, inefficient queries, inadequate hardware resources
  - Storage volume alone is not a sufficient metric to define a database service
  - Need dummy workload for each key application to define and optimize the service
    - How many requests from how many clients on how much data are required?
- Especially for distributed DB service some agreement like this will be essential to avoid surprises on either side
- This may need resources from the database services which need to be planned at all/some sites

Questions

- Which sites are ready to join the test bed now?
  - Initial focus will be technology tests deployment setup of Streams and FroNtier
    - Tools for shared administration and diagnostic
  - Need to be ready when applications are certified for tier 0 deployment
- Need a draft tier 0 integration and resource plan as soon as possible
  - Based on application priorities and availability as defined by experiments and projects
- Which sites are prepared provide application development & integration consultancy?
  - To speed up the T0 integration which is pre-requisite for higher tiers