LCG 3D Project Status

and proposed steps to enter maintenance mode

Dirk Düllmann,
LCG Management Board,
15. April 2008

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

LCG 3D Project Goals

CERN IT Department

- The Distributed Database Deployment project was started by the MB in July 2004
  - http://lcg.web.cern.ch/LCG/PEB/Minutes/Minutes20040720.html
- **Why a LCG Database Deployment Project?**

  - What's missing?
    - LCG provides an infrastructure for distributed access to file based data and data replication
    - Physics applications (and Grid services) require a similar infrastructure for data stored in relational databases
      - Several applications and services already use RDBMS
      - Several sites have already experience in providing RDBMS services
    - Need for some standardisation as part of LCG
      - To allow applications to access data in a consistent, location independent way
      - To allow to connect existing db services via data replication mechanisms
      - To simplify a shared deployment and administration of this infrastructure during 24*7 operation
      - To increase the availability and scalability of the total LCG system
  - Need to bring service providers (site technology experts) closer to database users/developers to define a LCG database service for the upcoming data challenges in 2005

Data Management

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

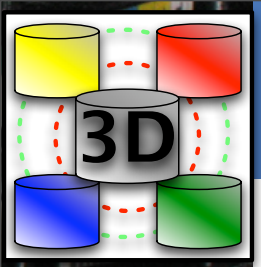

Project Goals

- Define distributed database services and application access allowing LCG applications and services to find relevant database back-ends, authenticate and use the provided data in a location independent way.
- Help to avoid the costly parallel development of data distribution, backup and high availability mechanisms in each experiment or grid site in order to limit the support costs.
- Enable a distributed deployment of an LCG database infrastructure with a minimal number of LCG database administration personnel.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

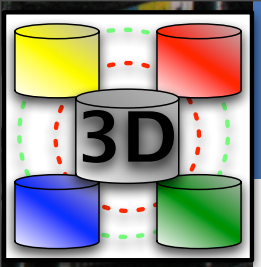

Status Today

- Large scale database installations are in place at CERN Tier 0 and ten Tier 1 sites and deployed in production mode since April 2007
  - according to the resource requests by the experiments
- Suitable replication technologies have been identified and integrated into the site database services
  - experiments have tested the installations during several larger production activities
  - site administrators and experiments are regularly using the LCG wide DB procedures and monitoring tools that been developed in the 3D project
- ATLAS, CMS and LHCb use streams replication for consistent replication between online and offline
  - ATLAS and LHCb also for Tier 1 replication
  - CMS uses FroNTier for Tier 1 and 2

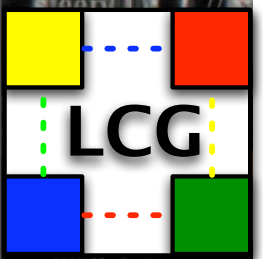

Data Management

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

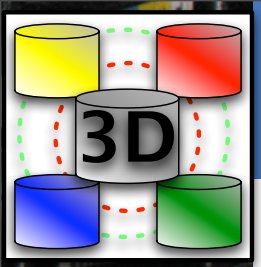

Software Integration

- Site replica look-up and grid authentication (including VOMS roles) has been integrated in the LCG Persistency Framework
- FroNTier protocol has been integrated into the LCG Persistency Framework
  - FroNTier servers have been setup as part of the 3D project and are now operated directly by CMS
- Some remaining scalability and security issues being addressed by the CORAL server development in the Persistency Framework

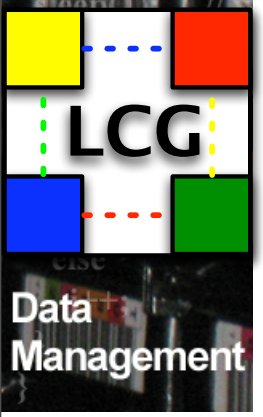

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Proposal

- The main service development goals of the project are now achieved
- The resulting services, procedures and tools are already well integrated into the daily database operations at CERN and Tier 1
- We propose to conclude LCG 3D as a service development project
  - possibly by defining any last experiment validation milestones, if not (yet) part of the general CCRC'08 activity
  - to summarise the work by issuing a final project report

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Move to Operations Mode

- During the project a well functioning group of database experts at experiments and sites has been formed
  - need to maintain this communication channel, which has been very productive
- Propose to continue with the established regular meetings and workshops
  - database version upgrade schedule
  - regular resource, hardware and license reviews
  - technical discussion between database experts from experiment and LCG sites
- Natural choice would be Maria Girone from the Tier0 physics database service as new chair for these operational meetings and workshops

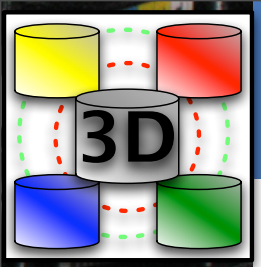

Acknowledgement

- I would like to thank the many experiment and site database experts and collaborators from CERN openlab and the s/w development projects who allowed to run the 3D project successfully without significant dedicated resources.

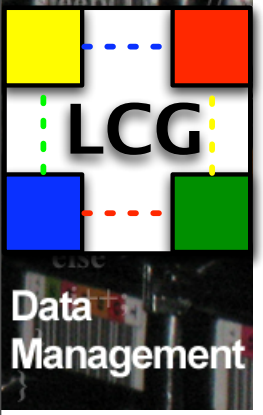

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it