Why a LCG Database Deployment Project?

- LCG today provides an infrastructure for distributed access to file based data and file replication
- Physics applications (and grid services) require a similar services for data stored in relational databases
  - Several applications and services already use RDBMS
  - Several sites have already experience in providing RDBMS services
- Goals for common project as part of LCG
  - increase the availability and scalability of LCG and experiment components
  - allow applications to access data in a consistent, location independent way
  - allow to connect existing db services via data replication mechanisms
  - simplify a shared deployment and administration of this infrastructure during 24 x 7 operation
- Need to bring service providers (site technology experts) closer to database users/developers to define a LCG database service
  - Time frame: First deployment in 2005 data challenges (autumn '05)

Project Non-Goals

- Store all database data
  - Experiments are free to deploy databases and distribute data under their responsibility
- Setup a single monolithic distributed database system
  - Given constraints like WAN connections one can not assume that a single synchronously updated database work and provide sufficient availability.
- Setup a single vendor system
  - Technology independence and a multi-vendor implementation will be required to minimize the long term risks and to adapt to the different requirements/constraints on different tiers.
- Impose a CERN centric infrastructure to participating sites
  - CERN is one equal partner of other LCG sites on each tier
- Decide on an architecture, implementation, new services, policies
  - Produce a technical proposal for all of those to LCG PEB/GDB

Tiers, Resources and Level of Service

- Different requirements and service capabilities for different tiers
  - Tier1 Database Backbone
    - High volume, often complete replication of RDBMS data
    - Can expect good network connection to other T1 sites
    - Asynchronous, possibly multi-master replication
    - Large scale central database service, local dba team
  - Tier2
    - Medium volume, often only sliced extraction of data
    - Asymmetric, possibly only uni-directional replication
    - Part time administration (shared with fabric administration)
  - Tier3/4 (eg Laptop extraction)
    - Support fully disconnected operation
    - Low volume, sliced extraction from T1/T2

- Need to deploy several replication/distribution technologies
  - Each addressing specific parts of the distribution problem
  - But all together forming a consistent distribution model

Starting Point for a Service Architecture?

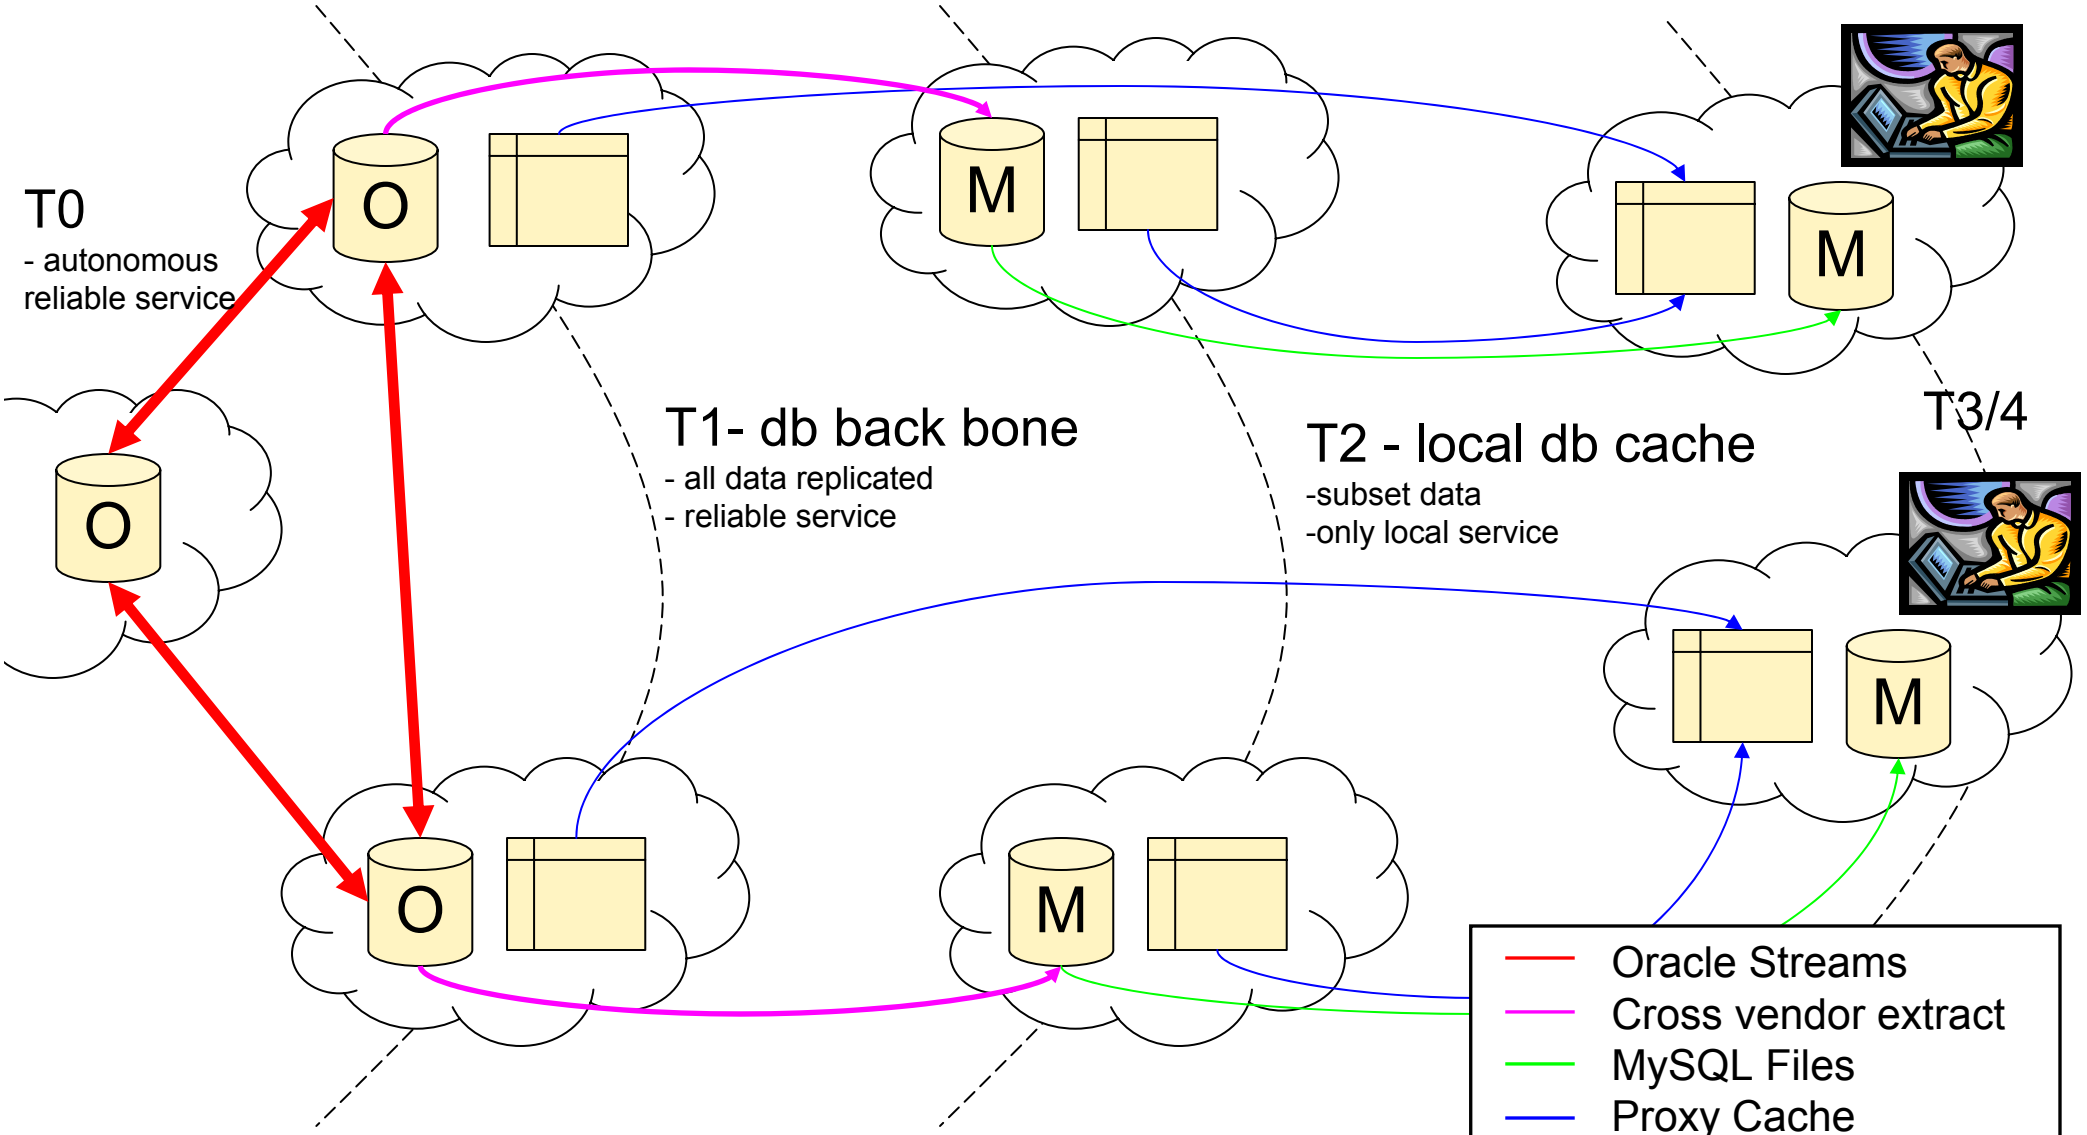

LCG 3D Project

WP1 -Data Inventory and Distribution Requirements

- Members are s/w providers from experiments and grid services that use RDBMS data
- Gather data properties (volume, ownership) requirements and integrate the provided service into their software

WP2 - Database Service Definition and Implementation

- Members are site technology and deployment experts
- Propose a deployment implementation and common deployment procedures

WP3 - Evaluation Tasks

- Short, well defined technology evaluations against the requirements delivered by WP1
- Evaluation are proposed by WP2 (evaluation plan) and typically executed by the people proposing a technology for the service implementation and result in a short evaluation report

Data Inventory

- Collect and maintain a catalog of main RDBMS data types
  - Select from catalog of well defined replication options
    - which can be supported as part of the service
  - Conditions and Collection/Bookkeeping data are likely candidates
- Experiments and grid s/w providers fill a table for each data type which is candidate for storage and replication via the 3D service
  - Basic storage properties
    - Data description, expected volume on T0/1/2 in 2005 (and evolution)
    - Ownership model: read-only, single user update, single site update, concurrent update
  - Replication/Caching properties
    - Replication model: site local, all t1, sliced t1, all t2, sliced t2 …
    - Consistency/Latency: how quickly do changes need to reach other sites/tiers
    - Application constraints: DB vendor and DB version constraints
  - Reliability and Availability requirements
    - Essential for whole grid operation, for site operation, for experiment production,
    - Backup and Recovery policy
      - acceptable time to recover, location of backup(s)

Service Definition and Implementation

- DB Service Discovery
  - How does a job find a close by replica of the database it needs?
  - Need transparent (re)location of services - eg via a database replica catalog
- Connectivity, firewalls and connection constraints
- Access Control - authentication and authorization
  - Integration between DB vendor and LCG security models
- Installation and Configuration
  - Database server and client installation kits
    - Which database client bindings are required (C, C++, Java(JDBC), Perl, ..) ?
  - Server and client version upgrades (eg security patches)
    - Are transparent upgrades required for critical services?
  - Server administration procedures and tools
    - Need basic agreements to simplify shared administration
    - Monitoring and statistics gathering
- Backup and Recovery
  - Backup policy templates, responsible site(s) for a particular data type
  - Acceptable latency for recovery

- **Bottom line: service effort should not be underestimated!**
  - We are rather close to LHC startup and can only afford to propose models that have a good chance of working!
  - Do not just hope for good luck; These services will be a critical part of the experiments' infrastructure and should be handled accordingly!