



# Oracle/Frontier and CondDB Consolidation

Dario Barberis

Genoa University/INFN



# Database distribution and access (1)

- Only *Geometry*, *Trigger* and *Conditions* data are accessed by ATLAS jobs running on the *Grid*
  - The other databases are used
    - Only by jobs running at CERN
    - Through special web interfaces (*AMI*, *TAGs*)
- *Trigger* and *Conditions* DB are distributed from CERN to the Tier-1s with Oracle Streams for analysis access
- The *Geometry* database (including only the base geometry, not alignments and calibrations that are in the *Conditions* DB) is distributed to each *Grid* site in the DB release
- *Conditions* data files are distributed to each *Grid* site using DDM
- Special *Conditions* Database Releases are created for reprocessing, containing:
  - An *SQLite* extraction of the *COOL* data for the run range to reprocess
  - The *conditions* data files for that run range
- DB and CDB releases are distributed on the *Grid*, downloaded on the worker node by each *Grid* job and accessed from the local disk



# Database distribution and access (2)

- We want to prevent any interference between scheduled production activities and chaotic user analysis
  - For the benefit of both sides
- Simulation jobs and analysis jobs running on simulated data only need the DB release with which the simulation was produced (or any more recent one)
- Reprocessing jobs should need only the Conditions DB release
  - This statement was found to be false in the current reprocessing campaign
    - Found several direct accesses to Oracle from Python code during job configuration (magnetic field and some detector status parameters) and one direct access to Oracle from within C++ code
    - These problems are solved (or about to be solved) with code modifications in forthcoming releases
- Analysis jobs access conditions data through local Squids and Frontier servers placed at Tier-1s and CERN
  - Analysis runs mostly at Tier-2s and Tier-3s (but also at Tier-1s)
    - Local caching of data is useful to reduce access times
  - Load on the Oracle servers is reduced also for jobs running at Tier-1s



# Database distribution and access (3)

- ATLAS has 10 Tier-1s, therefore Oracle streams were set up initially for all of them
- Early in 2009 we asked the sites that wanted to build experience with Frontier servers to go ahead and install it
  - At that time we had no idea of the load on Oracle and Frontier servers due to analysis jobs but we wanted to be on the safe side
- We now have 7 Frontier servers (CERN + 6 Tier-1s) and measure light loads
  - As a side-effect of the Frontier deployment, it turned out that Oracle databases without Frontier server should not be accessed by any job
    - Apart from the problem explained in the previous slide
- Discussion active in ATLAS Computing Model and Resource groups (CREM+ICB) on the consolidation of distributed database services:
  - How many Oracle/Frontier servers do we need, and which size, for analysis tasks?
  - Do we want to use Oracle/Frontier instead of CDB releases also for reprocessing?
    - And in this case how much more load shall we have?



# Discussion in ICB and WLCG 3D Workshop

- The ICB took note of the two options presented for Frontier deployment in ATLAS :
  - a) Continue with adiabatic approach used so far. Let sites willing to support Oracle/Frontier servers continue to do so but do not force them to have a very fail-proof service, as the global ATLAS system is robust through redundancy.
  - b) Have 2 large and well-maintained sites (CERN and BNL would be the natural choices for ATLAS) and discontinue the other sites.
- Sites that have already invested in setting up Oracle capacity, Frontier servers and expertise don't want to waste their investment (quite correctly!)
- Dario's proposal:
  - Keep the existing servers and if necessary consolidate their set-up
    - Improve monitoring of service availability and server loads
  - Use (with site agreement) the Oracle capacity at the other Tier-1s to increase the available storage capacity and server power for the TAG DB
    - Technical requirement specifications are in preparation by the TAG group
    - Work load for sites would not be appreciably different from now but they would provide a more useful service to the Collaboration
    - Discussions just started with CNAF, NDGF and SARA