Contribution ID: **6** Type: **not specified**

Evolution of ATLAS ADC relational databases

Monday 29 May 2017 11:50 (20 minutes)

The ATLAS Distributed Computing (ADC) project delivers production tools and services for ATLAS offline activities such as data placement and data processing on the Grid. The system has been capable of sustaining with high efficiency the needed computing activities during the first and in the ongoing second run of LHC data taking.
Databases are a vital part of the whole ADC system. The Oracle Relational Database Management System (RDBMS) has been addressing a majority of the ADC database requirements for many years. Much expertise was gained which could be used as a good foundation for next generations PanDA (Production ANd Distributed Analysis) and DDM (Distributed Data Management) systems.
By extrapolating of the current data volume in the roadmap to Run4, one could expect Grid operations on exabytes of data (factor ten from now). The corresponding catalog entries in the database and the rate of DB data change and read operations will greatly depend on the used object granularity: containers, datasets, files or events plus having event-level analysis workflow instead of file-level. It is expected that such evolution would require larger database capabilities and capacity in the medium-term future.

Author: DIMITROV, Gancho (CERN)

Presenter: DIMITROV, Gancho (CERN)

Session Classification: Requirements for run3&4

Track Classification: Requirements for run3&4