Contribution ID: **425** Type: **Poster**

Accounting the ATLAS DDM system – A case study with Oracle, MongoDB and HBase

Thursday 24 May 2012 13:30 (4h 45m)

The ATLAS Distributed Data Management system requires accounting of its contents at the metadata layer. This presents a hard problem
due to the large scale of the system and the high rate of concurrent modifications of data. The system must efficiently account more than 80PB of disk and tape that store upwards of
500 million files across 100 sites globally.
In this work a generic accounting system is presented, which is able to scale to the requirements of ATLAS. The design and architecture is presented, and three implementations are discussed, the reference implementation in Oracle RAC, and two alternative implementations in MongoDB and HBase. A strong emphasis is placed on the necessary design choices such that the underlying data models are generally applicable to many kinds of accounting, reporting and monitoring. The evaluation then focuses on principal architectural differences,
read-insert-update-delete performance, support for concurrent operations, deployment and operational effort, and possible means to
calculate the actual accounting values based on metadata critera. Finally, a recommendation is presented for the applicability of each
implementation under different accounting use cases, as well as an overall recommendation for useful and required data models.

Primary author: ATLAS, Collaboration (Atlas)

Co-authors: DIMITROV, Gancho (Brookhaven National Laboratory (US)); CHINZER, Lisa Azzurra (Universita e INFN (IT)); CANALI, Luca (CERN); LASSNIG, Mario (CERN); GARONNE, Vincent (CERN)

Presenter: LASSNIG, Mario (CERN)

Session Classification: Poster Session

Track Classification: Software Engineering, Data Stores and Databases (track 5)