

# The ATLAS Distributed Data Management Central Catalogues and steps towards scalability and high availability

*Monday 23 March 2009 08:00 (20 minutes)*

The ATLAS Distributed Data Management system, Don Quijote2 (DQ2), has been in use since 2004.

Its goal is to manage tens of petabytes of data per year, distributed among the WLCG.

One of the most critical components of DQ2 is the central catalogues which comprises a set of web services with a database back-end and a distributed memory object caching system.

This component has proven to be very reliable and to fulfill ATLAS requirements regarding performance and scalability.

In this paper we present the architecture of the DQ2 central catalogues component and implementation decisions regarding performance, scalability, replication and memory usage. The exploitation of techniques and features of the Oracle database which hosts the application is described together with an overview of the disaster recovery strategy that needs to be in place to address the requirement of high availability.

**Authors:** VIEGAS, Florbela (CERN); DIMITROV, Gancho (Lawrence Berkeley National Laboratory); SALGADO, Pedro (CERN)

**Co-authors:** CAMERON, David (Department of Physics, University of Oslo); LASSNIG, Mario (Distributed and Parallel Systems, University of Innsbruck); BRANCO, Miguel (CERN); WENAUS, Torre (Brookhaven National Laboratory); GARONNE, Vincent (CERN)

**Presenter:** SALGADO, Pedro (CERN)

**Session Classification:** Poster session

**Track Classification:** Software Components, Tools and Databases