



Contribution ID: 567

Type: **Poster**

The ATLAS database application enhancements using Oracle 11g

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

The ATLAS experiment at LHC relies on databases for detector online data-taking, storage and retrieval of configurations, calibrations and alignments, post data-taking analysis, file management over the grid, job submission and management, data replications to other computing centers, etc. The Oracle Relational Database Management System has been addressing the ATLAS database requirements to a great extent for many years. Several database clusters were deployed for the needs of the different applications. The data volume, complexity and demands from the users are increasing steadily with time. Nowadays about 20 TB of data are stored in the ATLAS Oracle databases at CERN (not including the index overhead), but the most impressive number is the hosted 260 database schemas (in the common case each schema is related to a dedicated client application with its own requirements). At the beginning of 2012 all ATLAS databases at CERN are upgraded to the newest Oracle version 11g Release 2. In order to make the ATLAS DB applications more reliable and performant we explored and evaluated the new 11g database features. In this work we present some of the Oracle 11g enhancements and typical ATLAS application use cases which suit best and the gain from the implemented changes.

Author: DIMITROV, Gancho (Brookhaven National Laboratory (US))

Co-authors: CANALI, Luca (CERN); BLASZCZYK, Marcin (CERN); SOROKOLETOV, Roman (University of Texas at Arlington (US))

Presenter: DIMITROV, Gancho (Brookhaven National Laboratory (US))

Session Classification: Poster Session

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