

<center>Montenegro, Budva, Becici, 28 september - 02 october 2015</center>



Contribution ID: 40

Type: **not specified**

Evolution of the use of relational and NoSQL databases in the ATLAS experiment.

Wednesday, 30 September 2015 11:10 (30 minutes)

The ATLAS experiment used for many years a large database infrastructure based on Oracle to store several different types of non-event data: time-dependent detector configuration and conditions data, calibrations and alignments, configurations of Grid sites, catalogues for data management tools, job records for distributed workload management tools, run and event metadata. The rapid development of “NoSQL” databases (structured storage services) in the last five years allowed an extended and complementary usage of traditional relational databases and new structured storage tools in order to improve the performance of existing applications and to extend their functionalities using the possibilities offered by the modern storage systems. The trend is towards using the best tool for each kind of data, separating for example the intrinsically relational metadata from payload storage, and records that are frequently updated and benefit from transactions from archived information. Access to all components has to be orchestrated by specialised services that run on front-end machines and shield the user from the complexity of data storage infrastructure.

This talk will describe this technology evolution in the ATLAS database infrastructure and present a few examples of large database applications that benefit from it.

Primary author: Prof. BARBERIS, Dario (University and INFN Genova (Italy))

Presenter: Prof. BARBERIS, Dario (University and INFN Genova (Italy))

Session Classification: Non-relational databases and heterogeneous repositories