



Contribution ID: 64

Type: oral presentation

Managing ATLAS data on a petabyte-scale with DQ2

Monday, September 3, 2007 5:30 PM (20 minutes)

The ATLAS detector at CERN's Large Hadron Collider presents data handling requirements on an unprecedented scale. From 2008 on the ATLAS distributed data management system (DQ2) must manage tens of petabytes of event data per year, distributed globally via the LCG, OSG and NDGF computing grids, now known as the WLCG.

Since its inception in 2005 DQ2 has continuously managed all datasets for the ATLAS collaboration, which now comprises over 3000 scientists participating from more than 150 universities and laboratories in more than 30 countries.

Fulfilling its primary requirement of providing a highly distributed, fault-tolerant as well as scalable architecture

DQ2 has now been successfully upgraded from managing data on a terabyte-scale to data on a petabyte-scale.

We present improvements and enhancements to DQ2 based on the increasing demands for ATLAS data management. We describe performance issues, architectural changes and implementation decisions, the current

state of deployment in test and production as well as anticipated future improvements. Test results presented here show that DQ2 is capable of handling data up to and beyond the requirements of full-scale data-taking.

Primary authors: Dr GAIDIOZ, Benjamin (CERN); Dr KOBLITZ, Birger (CERN); Mr LASSNIG, Mario (CERN & University of Innsbruck, Austria); Dr LAMANNA, Massimo (CERN); Mr BRANCO, Miguel (CERN); Mr SALGADO, Pedro (University of Texas at Arlington, USA); Mr ROCHA, Ricardo (CERN); Dr GARONNE, Vincent (CERN)

Presenter: Mr LASSNIG, Mario (CERN & University of Innsbruck, Austria)

Session Classification: Grid middleware and tools

Track Classification: Grid middleware and tools