

Functional and Large-Scale Testing of the ATLAS Distributed Analysis Facilities with Ganga

Thursday, 26 March 2009 15:20 (20 minutes)

Effective distributed user analysis requires a system which meets the demands of running arbitrary user applications on sites with varied configurations and availabilities. The challenge of tracking such a system requires a tool to monitor not only the functional statuses of each grid site, but also to perform large-scale analysis challenges on the ATLAS grids. This work presents one such tool, the ATLAS GangaRobot, and the results of its use in tests and challenges. For functional testing, the GangaRobot performs daily tests of all sites; specifically, a set of exemplary applications are submitted to all sites and then monitored for success and failure conditions. These results are fed back into Ganga to improve job placements by avoiding currently problematic sites. For analysis challenges, a cloud is first prepared by replicating a number of desired DQ2 datasets across all the sites. Next, the GangaRobot is used to submit and manage a large number of jobs targeting these datasets. The high-loads resulting from multiple parallel instances of the GangaRobot exposes shortcomings in storage and network configurations. The results from a series of cloud-by-cloud analysis challenges starting in fall 2008 are presented.

Primary author: VAN DER STER, Daniel Colin (Conseil European Recherche Nucl. (CERN))

Co-authors: SERFON, Cedric (Ludwig-Maximilians-Universität München); GALEAZZI, Fulvio (INFN - Roma 3); ELMSHEUSER, Johannes (Ludwig-Maximilians-Universität München); SLATER, Mark (University of Birmingham); BIGLIETTI, Michela (INFN - Napoli)

Presenter: VAN DER STER, Daniel Colin (Conseil European Recherche Nucl. (CERN))

Session Classification: Distributed Processing and Analysis

Track Classification: Distributed Processing and Analysis