

Database Access Patterns in ATLAS Computing Model

Monday, 13 February 2006 15:00 (20 minutes)

In the ATLAS Computing Model widely distributed applications require access to terabytes of data stored in relational databases. In preparation for data taking, the ATLAS experiment at the LHC has run a series of large-scale computational exercises to test and validate multi-tier distributed data grid solutions under development.

We present operational experience in ATLAS database services acquired in large-scale computations run on a federation of grids harnessing the power of more than twenty thousand processors. Among the lessons learned is the increase in fluctuations in database server workloads due to the chaotic nature of grid computations. The observed fluctuations in database access patterns are of a general nature and must be addressed through services enabling dynamic and flexibly managed provisioning of database resources. ATLAS is collaborating with the LCG 3D project and the OSG Edge Services Framework activity in the development of such services.

ATLAS database services experience relevant to local CERN data taking operations is also presented including the conditions data flow of ATLAS Combined Test Beam operations, prototype Tier 0 scalability tests and event tag database operations.

Summary

ATLAS database services experience and lessons learned in large-scale hierarchical multi-tier distributed computations is presented.

Primary authors: NAIRZ, A. (CERN); VANIACHINE, A. (ANL); MALON, D. (ANL); ZEMA, F. (CERN); GIER-ALTOWSKI, G. (ANL); POULARD, G. (CERN); CRANSHAW, J. (ANL); KARR, K. (ANL); GOOSSENS, L. (CERN); BARROS, N. (CERN); NEVSKI, P. (BNL); HAWKINGS, R. (CERN); WENAU, T. (BNL); SHAPIRO, Yu. (CERN); SMIRNOV, Yu. (BNL)

Presenter: VANIACHINE, A. (ANL)

Session Classification: Distributed Event production and Processing

Track Classification: Distributed Event production and processing