

Advanced Technologies for Scalable ATLAS Conditions Database Access on the Grid

Monday 23 March 2009 14:20 (20 minutes)

During massive data reprocessing operations an ATLAS Conditions Database application must support concurrent access from numerous ATLAS data processing jobs running on the Grid. By simulating realistic workflow, ATLAS database scalability tests provided feedback for Conditions DB software optimization and allowed precise determination of required distributed database resources. In distributed data processing one must take into account the chaotic nature of Grid computing characterized by peak loads, which can be much higher than average access rates. To validate database performance at peak loads, we tested database scalability at very high concurrent jobs rates. This has been achieved through coordinated database stress tests performed in series of ATLAS reprocessing exercises at the Tier-1 sites. The goal of database stress tests is to detect scalability limits of the hardware deployed at the Tier-1 sites, so that the server overload conditions can be safely avoided in a production environment. Our analysis of server performance under stress tests indicates that Conditions DB data access is limited by the disk I/O throughput. An unacceptable side-effect of the disk I/O saturation is a degradation of the WLCG 3D Services that update Conditions DB data at all ten ATLAS Tier-1 sites using the technology of Oracle Streams. To avoid such bottlenecks we prototyped and tested novel approach for database peak load avoidance in Grid computing. Our approach is based upon the proven idea of “pilot” job submission on the Grid: instead of the actual query ATLAS utility library sends to the database server a “pilot” query first.

Presentation type (oral | poster)

oral

Summary

We present detail results of our database access scalability studies as well as technologies developed to eliminate database access bottlenecks in a Grid computing environment.

Authors: VANIACHINE, Alexandre (Argonne); VALASSI, Andrea (CERN); WONG, Andrew (TRIUMF); VIEGAS, Florbela (CERN); DIMITROV, Gancho (LBNL); CANALI, Luca (CERN); GIRONE, Maria (CERN); NEVSKI, Pavel (BNL); HAWKINGS, Richard (CERN); WALKER, Rodney (LMU Munich); BASSET, Romain (CERN)

Presenters: VANIACHINE, Alexandre (Argonne); WALKER, Rodney (LMU Munich)

Session Classification: Software Components, Tools and Databases

Track Classification: Software Components, Tools and Databases