

Job Centric Monitoring for ATLAS jobs in the LHC Computing

Thursday, 26 March 2009 08:00 (20 minutes)

As the Large Hadron Collider (LHC) at CERN, Geneva, has begun operation in September, the large scale computing grid LCG (LHC Computing Grid) is meant to process and store the large amount of data created in simulating, measuring and analyzing of particle physics experimental data. Data acquired by ATLAS, one of the four big experiments at the LHC, are analyzed using compute jobs running on the grid and utilizing the ATLAS software framework 'Athena'. The analysis algorithms themselves are written in C++ by the physicists using Athena and the ROOT toolkit. Identifying the reason for a job failure (or even the occurrence of the failure itself) in this context is a tedious, repetitive and - more often than not - unsuccessful task. The debugging of such problems was not foreseen and tracing back problems is even more difficult by the fact that the output-sandbox, which contains the jobs' output and error logs, is discarded by the grid middleware if the job failed. So, valuable information that could aid in finding the failure reason is lost. These issues result in high job failure rates and less than optimal resource usage.

As part of the High Energy Particle Physics Community Grid project (HEPCG) of the German D-Grid Initiative, the University of Wuppertal has developed the Job Execution Monitor (JEM). JEM helps finding job failure reasons by two means: It periodically provides vital worker node system data and collects jobrun-time monitoring data. To gather this data, a supervised line-by-line execution of the user job is performed. JEM is providing new possibilities to find problems in largely distributed computing grids and to analyze these problems in nearly real-time.

All monitored information is presented to the user almost instantaneously and additionally stored in the jobs' output sandbox for further analysis. As a first step, JEM has been seamlessly integrated into ATLAS' and LHCb's grid user interface 'ganga'. In this way, submitted jobs are monitored transparently, requiring no additional effort by the user.

Primary author: KALININ, Sergey (Universite Catholique de Louvain)

Presenter: KALININ, Sergey (Universite Catholique de Louvain)

Session Classification: Poster session

Track Classification: Grid Middleware and Networking Technologies