Monitoring the ATLAS Production System

J. Kennedy, O. Biebel, G. Duckeck, J. Elmsheuser, G. Krobath, B. Ruckert, D. Schaile, and C. Serfon

1. Introduction

Throughout 2006 ATLAS has operated a system to continuously produce simulated Monte-Carlo data for its physics community. The Monte-Carlo production is driven by the so called production system. The production system is formed from several elements which allow jobs to be distributed to the 3 grids used within ATLAS (EGEE,OSG,NorduGrid).

Several teams of "shifters" are responsible for distributing the jobs, monitoring their progress and reporting problems at both the site and the job level.

During 2006 the production system reached peak rates of O(20,000) successful jobs per day. This rate is expected to double each quarter in 2007. With such a high number of inhomogeneous jobs being distributed across O(150) sites worldwide an efficient monitoring system is indispensable.

2. Monitoring

The LMU group has taken a leading role in developing the production system monitoring, providing tools for people running production and physics end users who wish to be informed of the progress on their MC datasets.

The Monitoring system uses a set of database tables which compress the relevant information from the global database into smaller more accessible data. This approach ensures that the monitoring itself has a limited impact on the production, avoiding many long queries of the production database. In addition it guarantees quicker information retrieval from the monitoring queries.

The project provides a set of websites which present views of the jobs with their success/failures as well as accounting views. Users can view the last 24hrs based on the job type, the centers at which they were executed and the system(team) through which they were distributed.

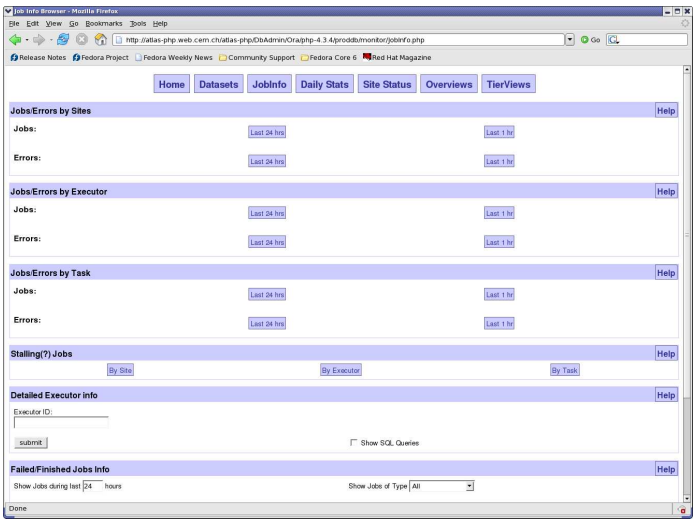

Fig. 1: Example Monitoring page

Fig. 1 shows an example of one of the web pages provided by the monitoring system, simple click buttons allow users to gather relevant information quickly to identify problems.

Fig. 2 shows a view of errors experienced on one grid over a certain date range, in this case just 1 day. These types of views allow the shifting teams to spot poor performance during a given time window and also to pinpoint the cause.

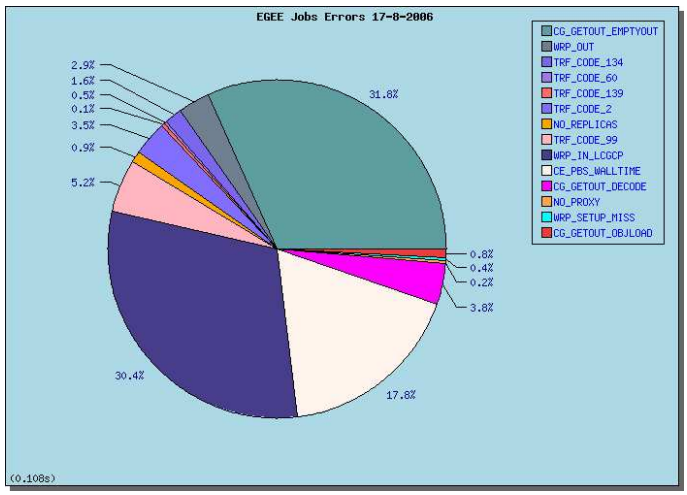

Fig. 2: Errors in the EGEE grid flavor on 17th August 2006

As the project evolved it became apparent that the shift teams responsible for the production itself would be the main users of the monitoring system. Thus, to aid their work, additional pages were developed to help display information such as, sites with missing software, correlated file transfer errors in groups of sites as well as views of the jobs waiting to be run and jobs which appear to be in stuck/unchanging states.

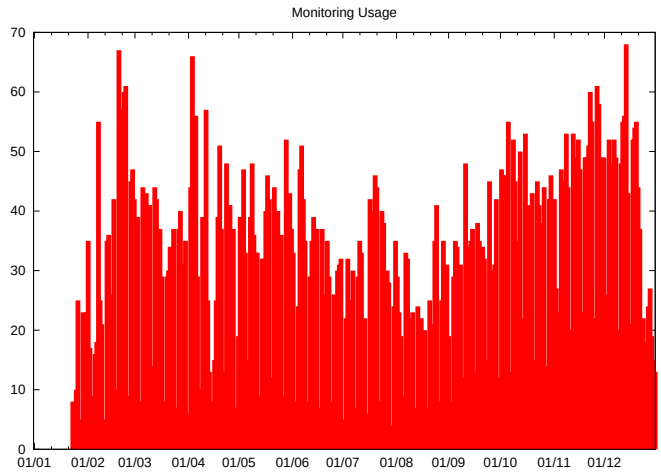

Fig. 3: Monitoring Usage during 2006

Throughout 2006 the monitoring pages have been used steadily by both production managers and physics end users. Fig. 3 shows the usage of the monitoring pages where users per day are accounted.

The monitoring project will continue to develop during 2007. The main goals are to improve our ability to identify and solve problems and to allow us to study the use of our resources to best utilize them.