The ATLAS Production System

J. Kennedy, O. Biebel, T. Christiansen, J. Dubbert, G. Duckeck, J. Elmsheuser, M. Erlebach, F. Fiedler, R. Hertenberger, R. Mameghani, D. Merkl, T. Nunnemann, M. Obermaier, F. Rauscher, D. Schaile, M. Schott, A. Staude, H. Steffens, R. Ströhmer, A. Varga, F. Vollmer, and C. Zupancic

1. ATLAS Data Processing

The ATLAS experiment at CERN will begin data taking in 2007 when the LHC is commissioned. From then on data is expected to be recorded at a rate of 2 PetaBytes per year with an additional 1 PetaByte of simulated data produced per year.

This huge data volume means that it is impossible to perform all the processing and analysis at CERN. ATLAS aims to utilise distributed computing resources in collaborating countries from all around the world.

To aid with the development of this distributed computing infrastructure a series of data challenges were started in 2002. The data challenges allow us to evaluate the ATLAS computing model, the full software suite, the data model and also to ensure the correctness of technical choices made for ATLAS computing.

The second data challenge, DC2, began in summer 2004. DC2 focuses on the use of an automated production system and GRID software to produced 10^7 simulated physics events. In addition DC2 provides an opportunity to test the simulation of physics events using the ATLAS software framework(ATHENA). The physics events produced during DC2 may be utilised by the physicists for large scale physics analysis and feedback from both the event production and physics analysis will be incorporated into the ATLAS computing model to allow for a realistic estimation of our computing needs.

2. The ATLAS Production System

The LMU group participated in the design of the ATLAS production system which provides a common framework in which any grid flavour or legacy system may be integrated [1,2,3,4]. The production system is formed from several individual elements which when plugged together provide the required functionality for the submission, tracking, recovery and validation of jobs. The individual elements of the production system may be summarised as follows.

- Common database for production jobs
- Common Supervisor run by all facilities/managers
- Executors developed by middleware experts
- Data Management system to allow inter-grid data transfer and file cataloging

Figure 1 shows an overview of the ATLAS production system. The core of the ATLAS production system is formed from the coupling of a supervisor with an executor. The supervisor provides an interface to the job definition data and metadata associated with the jobs while the executor provides an interface to the computing resources. The communication between the two elements is performed using either Jabber or SOAP, this allows some flexibility in the design and implementation of the system. The use of XML messaging to communicate between the supervisor and executor provides a means of allowing the individual elements of the production system to be designed without being strictly tied to one technology. The supervisor interfaces to a central production database and retrieves job specific information. The supervisor-executor system then allows these jobs to be passed to one of the grid flavours or standard batch system. Through its interface to the executor the supervisor may continuously monitor the state of the submitted jobs and finally retrieve detailed information about the job once the job has finished.

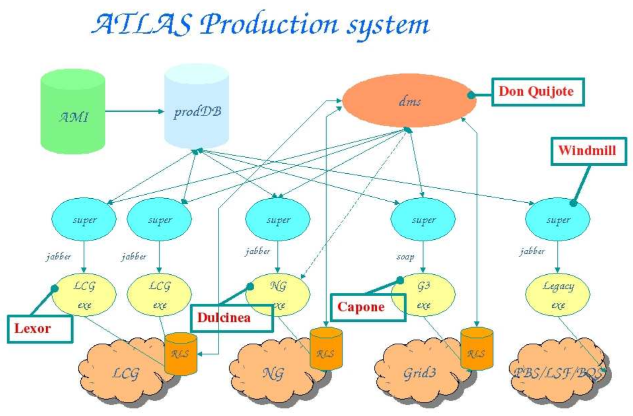

Fig. 1: A schematic of the ATLAS Production System

3. The Legacy Executor

In addition to work on the production system design the LMU group produced a legacy executor interface to batch systems [1]. The Legacy executor may be used as a backup solution should unforeseen problems arise with the development of grid executors, provide a testing facility for use when developing the production system and to provide access to computing resources at sites which have not yet deployed any grid middleware. The legacy executor was developed using the PBS batch system at the GridKa facility in Karlsruhe and later ported to the LSF and BQS batch systems. Extensive tests of the executor were performed with several hundred jobs being submitted to the GridKa facility.

References

[1] J. Kennedy, CHEP04 proceedings
[2] L. Goossens *et al.*, CHEP04 proceedings
[3] `http://heppc12.uta.edu/windmill/`
[4] M. Branco, CHEP04 proceedings