

# Monitoring the ATLAS Production System

J. Kennedy, G. Duckeck, J. Elmsheuser, and D. Schaile

## 1. Introduction

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.

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.

## 2. Production System

The production system is designed to provide a common framework in which any grid flavour may be integrated. The 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 grid middleware experts
- Data Management system to allow intergrid data transfer and file cataloging

Jobs are defined in the central production database and are picked up by the supervisor element. The supervisor then uses an executor as an interface to the chosen grid. During the previous Data Challenge 2 and associated Rome production a peak rate of about 10,000 verified jobs per day was observed. The current phase of ATLAS MC production, Computing Systems Commissioning (CSC), is expected to surpass this performance by an order of magnitude within 2006. Considering the scale of the planned production it is clear that monitoring tools are an essential part of the production system framework. The monitoring framework described here and the snapshots shown in the figures was developed during the pre-CSC production system development/testing phase.

## 3. Monitoring

Monitoring the ongoing ATLAS production is an important and diverse task. The idea behind the current monitoring framework is to implement the functionality found to be of most use in previous production and to act as a testbed for monitoring/accounting developments. Furthermore the current system attempts to abstract the monitoring away from the central production database in an aim

to lower the impact that monitoring production has on the actual production itself.

The monitoring can be split into two separate aspects. Firstly monitoring the production via interfaces to the production database and secondly monitoring the grid resources by using pre existing grid monitoring tools.

The majority of the functionality provided thus far is based on the former method.

Some areas of interest may be summarised as follows,

- Datasets - Identify which datasets are being processed and the success rate of previous sets
- Jobs - Identify where they are running and correlate any errors
- Accounting statistics - Ensure some simple global statistics are available at the click of a button
- Grid Status - Live time status of grid sites, available resources, jobs running, storage space, installed SW, etc.

Monitoring the ATLAS production via the central production database allows us to provide generic monitoring views which cover the several grid flavours which support ATLAS. The use of individual tools which are specific to a given grid flavour is being explored and it is hoped that the information from these tools can be extracted and converted to a generic format.

Figure 1 shows the results for production per grid flavour in the first month of 2006. The Figure shows results for four flavours of Grid, NorduGrid - The Nordic Grid (Green), OSG - The US Open Science Grid (Red), LCG - The LHC Computing Grid Project (Pink) and the LCG-CG - A CondorG interface to the LCG resources (Blue). It should be noted that the LCG and LCG-CG flavours represent the same LCG resources which are simply accessed in a different manner.




Fig. 1: Production jobs per grid flavour during Jan 2006