

## The ATLAS Computing Model

---

### D Barberis

*INFN Genova and Università di Genova*

*Dipartimento di Fisica, Via Dodecaneso 33, IT-16146 Genova, Italy*

*E-mail: [Dario.Barberis@cern.ch](mailto:Dario.Barberis@cern.ch)*

### RWL Jones

*Lancaster University*

*Physics Department, Lancaster LA1 4YB, United Kingdom*

*E-mail: [Roger.Jones@cern.ch](mailto:Roger.Jones@cern.ch)*

### On behalf of the ATLAS Collaboration

The ATLAS computing model was constructed to exploit the opportunities of the Grid in handling the large volumes of data from the Large Hadron Collider and to allow easy and relatively local access to data for all of its collaborators worldwide. Despite delays with collision data, the model has now been tested with beam-related and cosmic ray data, and much has been learned. The model has retained its overall design, but with adjustments for the actual functionality of the delivered Grid middleware and services, and for the realities of data access. While much has still to be learned, the model has worked effectively. This paper covers the roles of the various Tiers, the resources required and the expected workflows.

*European Physical Society Europhysics Conference on High Energy Physics EPS-HEP 2009,  
Krakow , Poland  
July 16-22, 2009*

## 1. Introduction

The ATLAS computing model has been developed over more than ten years to meet the challenges of the LHC era [1]. These challenges are many:

- Very many petabytes of raw and processed data per year
- Diverse data formats, with both large and small files
- A world-wide analysis community

The first modelling exercises were very hierarchical. With the advent of the Grid paradigm, a refined version of the model emerged, with the Grid providing ‘clouds’ of resources that still had an underlying hierarchical structure. The model was initially tested through exercises of increasing complexity with simulated data (data challenges, functional tests, full-dress rehearsal), and since the autumn of 2007 with cosmic ray triggers recorded with the full detector in the underground cavern and for a brief period in 2008 with beam-related data. Through this process, the computing model evolved to meet the problems encountered in the exercises and to match the actual level of functionality delivered by the various Grid middleware.

## 2. The Tier Structure and Roles

The model retains a tiered structure, with the Tier-0 at CERN performing first pass processing of both the express and calibration stream, and subsequently of the bulk dataset. The ATLAS Tier-0 runs completely automatic processes:

- Runs reconstruction on event samples to extract data for new calibrations and alignments
- Runs automatic calibration and alignment processes and validates the results
- Runs full reconstruction on all events
- Runs monitoring tasks to assess the data quality, saving the information into web servers and databases
- Saves RAW and reconstructed data to tape
- Exports RAW and reconstructed data to Tier-1s (one copy of RAW data distributed amongst all Tier-1s, two copies of the full reconstruction output data (ESD) distributed amongst all Tier-1s, and one copy of the physics analysis data (AOD) to each Tier-1, to be further distributed to Tier-2s in each cloud)
- Exports samples of RAW data, calibration data, AODs and histograms to the CAF

All data movements take place from/to Castor disk pools; tapes are treated as data archives. The ATLAS CAF runs primarily high-priority non-automatic processes:

- Receives samples of RAW data, calibration data, AODs and monitoring histograms from the Tier-0
- Members of relevant activity groups run jobs to debug the trigger and the detectors, as well as reconstruction code (if needed) to give fast feedback to online operations and to Tier-0 processing
- Prompt calibration processes that are not yet completely automatic and need iterations and human checking are also run on the CAF

There are 10 Tier-1 facilities for ATLAS, hosted at major computing centres. These receive a portion of the raw data, and are responsible for its long-term curation and reprocessing to derived formats. In the steady state, this is expected to occur twice a year: once after a few months, when better calibrations have been obtained; and once at the end of data taking, when algorithmic improvements can be applied. The Tier-1s also provide a facility for scheduled access by physics and detector performance groups to large quantities of data. As the requirements for simulation have increased, the Tier-1s also provide part of the experiment simulation capacity.

Beyond this, the Tier-1s provide services for a cloud of associated Tier-2 facilities, which are usually based in universities. The Tier-1s host file catalogues and file transfer services, and act as a staging point in the transfer of data around the world. They also host in their Oracle servers a complete replica of the Conditions Database, streamed to them from CERN in real time.

Currently 37 Tier-2 federations (~70 computing centres) support ATLAS computing. They are organizationally connected into "Tier-1 clouds", each consisting of a Tier-1 and its associated Tier-2s. Associations were defined based on common funding, quality of networking or direct agreements. ATLAS Tier-2s have a few main functions, providing CPU capacity for simulation tasks, disk storage for their data fractions, CPU capacity for group and user access to the data they keep on disk and disk storage for group and user derived data.

### 3. Functional Tests

Several rounds of tests of increasing complexity have been performed over many years. Since Autumn 2008, ATLAS had two major periods of cosmic-ray data taking; these data were treated as LHC collision data and the same processing, and reprocessing, procedures were applied. Distributed simulation production also continued in the background and filled all currently available CPU slots (in excess of 36k). In parallel, intensive tests of distributed analysis tasks were performed on all ATLAS Tier-1 and Tier-2 facilities. A low-level load of data transfers and production and analysis test jobs keeps flowing through the system in order to provide constant information on the health of the sites and Grid services.

### 4. Conclusions

The ATLAS computing model has stood up well to tests with real data at large scale and rate. These tests continue, and the analysis load will now be added to the exercises. The model detail has had to change to meet the problems of the real world and the delivered middleware functionality, but we believe we are in a good position to process and analyse the collision data from autumn 2009.

### References

- [1] The ATLAS Computing Technical Design Report, ATLAS-TDR-017; CERN-LHCC-2005-022, June 2005.