


From Bits to Particles

Each part of the ATLAS detector produces an electronic signal every time a particle crosses it. The combination of detector position, time and amplitude of the signal, provide information on the particle trajectory and energy deposition in each detector element. The signals recorded by all detector elements at a given time constitute an “event”, and provide information on all particles produced by a single proton-proton interaction in the centre of the ATLAS detector.




Special software, developed by ATLAS physicists and computer scientists, is used to process each event. Processing means using the detector signals to reconstruct the trajectories and energies of all particles, and identify them as electron, muons, pions, protons etc. The output of such reconstruction process for each event consists in lists of particles (and their properties), which are then further analysed by physicists to determine the nature of the interaction that produced them.

Calibration and Alignment

In order to process the electronic signals produced by each detector component and reconstruct particles, calibration and alignment constants are needed. Calibration constants are the conversion factors between the recorded signal amplitudes and the energy deposited by particles in each detector element; alignment constants are the position in space of each detector element. All these constants vary with time, depending on the local temperature of the detectors and atmospheric pressure, but also on LHC beam conditions (as higher interaction rates produce more activity in the electronics and therefore more heat) and on external factors like the humidity of the rocks surrounding the ATLAS cavern.

Calibration and alignment data, as well as all recorded “conditions data” (temperatures, gas pressures, voltages and currents in power supplies etc.), are recorded in a database as a function of time. The database information is then replicated in real time from CERN to the 10 larger (“Tier-1”) computing centre for ease of access. Event reconstruction and analysis code uses the conditions data to improve the precision of physics measurements.



The multi-tier structure of ATLAS world-wide computing centres

World-wide Computing Infrastructure

ATLAS has set up a multi-tiered world-wide computing infrastructure. The CERN computer centre hosts the Tier-0 and the “CERN Analysis Facility” (CAF) clusters. They are used for immediate data processing and detector data quality monitoring; the Tier-0 also stores on tape all data that are recorded by the ATLAS detector. Raw and processed data are distributed from CERN to the Tier-1 and Tier-2 centres.

Ten large laboratories that support ATLAS provide Tier-1 computing centres. They provide data storage, including keeping a second copy of the detector raw data on tape (for safety). They also keep on disk the processed data from Tier-0 for fast analysis access. Their computing capacity is used to reprocess the data, if better calibration or alignment constants, or better reconstruction code, are available.

About fifty ATLAS institutes have so far set up a Tier-2 computing facility. They hold copies of processed data, in summary form that can be directly used for most physics analyses; they also host user derived data, which are produced by the analysis procedures and have to be shared with other members of a given analysis working group. Tier-2 sites in addition produce simulated data for the needs of the whole collaboration. All ATLAS institutes act as Tier-3 sites. These are gateways into the Grid system; users can submit jobs to the Grid, query their status, retrieve the outputs, or copy data for local processing.



ATLAS on the Grid

- Three Grid middleware infrastructures are used by the ATLAS distributed computing project:
- EGEE (Enabling Grids for E-science) in most of Europe and the rest of the world.
 - NorduGrid in the Scandinavian countries.
 - OSG (Open Science Grid) In the United States of America.
- The ATLAS Grid tools interface to all middleware types and provide uniform access to the Grid environment.
- The VOMS (Virtual Organisation Management System) database contains the computing information and the privileges of all ATLAS members; it is used to allow ATLAS jobs to run on ATLAS resources and store their output files on ATLAS disks.
 - The DDM (Distributed Data Management) system catalogues all ATLAS data and manages the data transfers in the ATLAS Grid.
 - The ProdSys/PanDA production system schedules all organised data processing and simulation activities, distributes jobs to the sites holding corresponding data, and organise the safe storage of output data.
 - The Ganga and Pathena interfaces allow any ATLAS member to submit analysis jobs to all ATLAS computing facilities; jobs go to the site(s) holding input data, and output data can be stored locally or sent back to the submitting site.
 - The “Dashboard” monitoring system supports central monitoring of all operations, including site status, data transfer and job execution rates.

Resources & Statistics

Total CERN capacity (2008): CPU 5.8 MSI2k, disk 1.3 PB, tape 2.8 PB
Sum of all Tier-1s (2008): CPU 18.1 MSI2k, disk 10.7 PB, tape 8.1 PB
Sum of all Tier-2s (2008): CPU 17.5 MSI2k, disk 7.7 PB, no tapes
1 MSI2k = 500 CPU cores = 65 modern professional machines.
1 PB of disk = 1,000 TB = 1 million GB; 1 PB of tape = 2,500 tapes
Data transfer rate from the ATLAS detector to the CERN computing centre: 320 MB/s during LHC operation
Data export rate from CERN to the Tier-1 centres, including raw and processed data: 1.2 GB/s
Network links between CERN and the Tier-1 centres (LHC Optical Private Network): minimum 10 Gb/s for all pairs of computer centres
Network links between Tier-2 and Tier-1 centres: at least 1 Gb/s with the nearest Tier-1 site
The total ATLAS disk storage per year of data is the equivalent of 200,000 top of the range iPod classics, and they would cover 85 tennis courts.