Journal of Physics: Conference Series

Assessment of data quality in ATLAS

To cite this article: M G Wilson 2008 *J. Phys.: Conf. Ser.* **119** 042034

Journal of Physics: Conference Series **119** (2008) 042034 doi:10.1088/1742-6596/119/4/042034

Assessment of data quality in ATLAS

M G Wilson, for the ATLAS collaboration

European Organization for Nuclear Research (CERN), Geneva, Switzerland

E-mail: `Michael.G.Wilson@cern.ch`

Abstract. Assessing the quality of data recorded with the ATLAS detector is crucial for commissioning and operating the detector to achieve sound physics measurements. In particular, the fast assessment of complex quantities obtained during event reconstruction and the ability to easily track them over time are especially important given the large data throughput and the distributed nature of the analysis environment. The data are processed once on a computer farm comprising $\mathcal{O}(1,000)$ nodes before being distributed on the Grid, and reliable, centralized methods must be used to organize, merge, present, and archive data-quality metrics for performance experts and analysts. A review of the tools and approaches employed by the detector and physics groups in this environment and a summary of their performances during commissioning are presented.

1. Overview of the data-assessment computing environment and processes

The computer programs and infrastructure providing the means for assessing data quality at ATLAS must be relatively lightweight, reliable and thorough; in general, we aim for a monitoring system that is more reliable than the processes being monitored. Therefore, we require relatively few services and provide redundancy where possible in the development of software, application chains and refined-data storage. Common frameworks and tools are developed and shared among all trigger, detector and physics-performance monitoring groups in order to provide a consistent and complete coverage of data-quality metrics.

Among the challenges faced are those of scalability and speed. The data are processed locally at CERN within one to two days after being recorded, after which they are distributed to other computing centers via the Grid. This data processing requires a computer farm, Tier 0, with $\mathcal{O}(1,000)$ nodes, which processes $\mathcal{O}(10,000)$ files per run[1]. Since the data are processed in this manner at CERN only once before being distributed and possibly used for detector and physics analyses, a reasonably accurate assessment of the data quality must also have a timescale of one to two days. The monitoring frameworks and tools used for assessment must account for and be reliable within the parallel processing on the Tier 0 farm, and they must be flexible enough to accommodate the diverse needs of trigger, detector and physics-performance monitoring groups. In particular, monitoring methods must maintain compatibility, to the greatest extent possible, with existing tools used for commissioning, yet they must also be scalable and reliable for LHC operation.

The main computing components of the data-assessment process are shown schematically in Figure 1. During the data-acquisition process, detector-status data are stored in the conditions database, and selected events are recorded in files. The files are organized into streams:

[1] Assuming a run last approximately 10 hours.

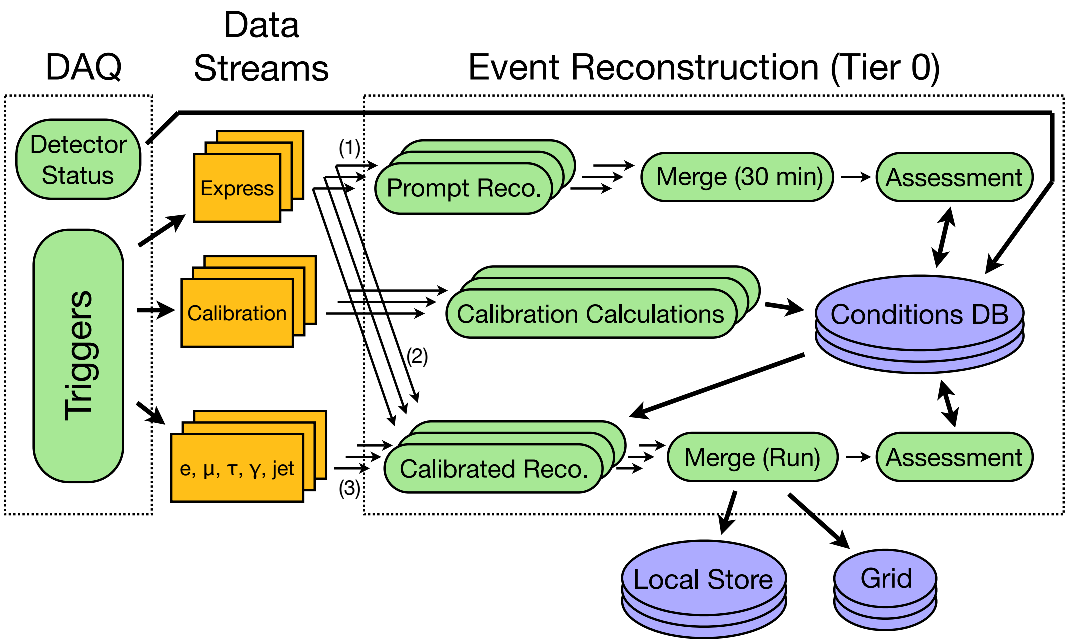

Figure 1. A schematic overview of the main components in the data-assessment process. The data flow is from left to right. The data-acquisition (DAQ) system selects events according to trigger criteria and it records the status of detector components in the conditions database (CondDB). The data are written by the DAQ into streams according to the trigger composition. Calibration-stream data are specially selected by subdetectors to ensure that calibration constants in the CondDB are always current when reconstructing analysis data. Analysis data comprises the express stream and physics-object (ex., e, μ, τ, γ, and jet) streams. The analysis events are reconstructed in parallel at Tier 0, the output files are merged, and the merged files are assessed and distributed. The express stream, comprising a prescaled selection of many types of triggers, is processed promptly, possibly without the most accurate calibrations, to provide rapid feedback. After a run is completed and calibration constants have been updated, the express stream is processed and assessed again to ensure that the calibration is correct. Finally, the physics-object streams are processed and assessed. The data-quality metrics used for assessment are stored locally and on the Grid, for redundancy and ease of retrieval.

calibration streams contain events selected to calibrate subdetectors, and analysis streams contain events for detector and physics analyses. A special analysis stream is the express stream, containing a prescaled mixture of different types of physics signals; it is primarily intended to assess data quality quickly. The basic steps in assessing the data in analysis streams are histogram production, histogram-file merging, assessment and histogram archival. Histograms for assessing data quality are produced when event-reconstruction programs are run on the Tier-0 nodes. These histograms are saved into files, and the files from each node are collected and merged into a single file. During this merge step, some histograms are combined to produce new histograms with higher statistics, while other histograms are preserved at a finer granularity. The merged file is assessed by a combination of automatic algorithms and physicists, via a display program, and the merged file is saved for possible retrieval later. The express stream is processed twice. First, it is processed promptly in 30-minute blocks to provide a fast, but not fully calibrated, feedback to the detector shift crew. Second, after a run is completed and the detector has been calibrated, the express stream is processed for the entire run with the new calibration applied. When the data from this step are validated, the physics-object streams, containing all of the recorded data, are processed and assessed.

2. Software tools for conducting assessments

Software frameworks and tools have been developed for performing the steps to assess the data: histogram production, histogram-file merge, histogram-file archival and retrieval, histogram display and assessment-status recording. The steps are performed on the Tier 0 in parallel with event-reconstruction and data-transformation programs, and a final assessment status is determined by physicists for each run.

Histogram production is done within and immediately after the event-reconstruction program. Developers create and fill their own histograms, but a small framework is provided for registering histograms so that they can be collected and organized into a single output file. The main interface to this framework is the pseudocode method

```
regHist( h, logicalPath, levelOfDetail, interval ).
```
(1)

In this method, `h` is a histogram object, `logicalPath` is a logical name within the output file, `levelOfDetail` is a flag describing the suggested relevance of the histogram (*i.e.*, 'shift', 'expert', or 'debug') and `interval` is a flag describing the time interval over which data for the histogram should be collected. The actual name of the histogram within the output file is determined dynamically from the given logical name, and it may include information about the run (such as the run number). The `levelOfDetail` and `interval` flags are saved in the output file to avoid the necessity of complex configurations in later stages. The `levelOfDetail` is used to build automatic displays and to assist users in retrieving subsets of histograms; the `interval` is used when merging histogram files to ensure that final histograms cover the intended time intervals.

Since each output file contains histograms filled from only a fraction of a run, histogram files are merged into larger blocks by collecting and combining the histograms from the initial output files. For fast validation, the express stream is reconstructed promptly, and histogram files are merged into blocks of data of $\mathcal{O}(1)$ hour. After a run is completed, it is evaluated using histogram files merged from the entire run, producing one file per run per stream.

After merged files are created, they are stored locally at CERN and distributed on the Grid, and information needed to retrieve the files is stored in a database. Because the merged files contain histograms from all systems for an entire run, the files may be quite large, $\mathcal{O}(10)$ GB. A utility has been developed to allow retrieval of subsets of histograms based on run, stream, and the `levelOfDetail` flag. In this way, histograms are stored in large files appropriate for mass-storage devices, yet users may search for and retrieve individual histograms as if they were stored separately in a database.

Histograms are displayed centrally on the web for fast and easy feedback. The `levelOfDetail` flag is used to automatically select a representative selection of histograms from each system for web display. In addition to the histograms themselves, derived quantities and checks on those quantities may be available for display. Derived quantities may be stored in the conditions database as a function of run, so that their time evolution may be monitored.

Finally, quality-assessment status is stored in the conditions database for access by analysts. A configurable calculator makes an automatic assessment based on information obtained during data acquisition and checks performed on reconstructed data. This assessment is reviewed and adjusted by physicists who have evaluated the data quality for that run based on the histogram displays. A web application is used to navigate and update the quality-assessment status.

3. Experience from detector and software commissioning and conclusions

The software that has been developed for assessing data quality is in active use in detector and software commissioning activities. The detector is being commissioned using cosmic-ray data, and because the subdetector groups have long had monitoring and validation utilities, a comparison between existing standalone tools and the new central tools is quite helpful in debugging and further development. By using central tools now, we are able to understand features which must be added before recording LHC data. In addition, by providing a common monitoring framework for all systems, we are able to more easily identify unmonitored quantities, scalability limitations, and integration and synchronization problems.

Many features of the anticipated LHC-event reconstruction cannot be tested with cosmic-ray data, and must be tested instead with simulations. For example, the streaming model, luminosity calculations, and use of the trigger menu may be tested with simulated data. As these are an essential part of future physics analyses, we must be able to monitor and validate them from the start of LHC data taking.

We believe that with the preparations we have made, ATLAS will be ready to assess LHC data quality as soon as the data are available. The infrastructure has been carefully considered and planned with participation from all systems, including trigger, detector and physics-performance monitoring groups. We are currently learning about how to use the full system so that it can be tuned and extended for optimal performance and reliability.