


ATLAS Data Quality Monitoring: Experience with First Collision Data

Peter Waller, on behalf of the ATLAS collaboration
ICHEP 2010 Paris



Introduction

It is essential to get quick feedback at all stages of the collection and reconstruction of particle physics data, and to correctly record quality decisions to ensure that only good data are used to obtain physics results. The ATLAS Data Quality (DQ) system provides prompt investigation of collected data, initial calibrations, and later reconstruction, and propagates the corresponding quality decisions to analysis users.

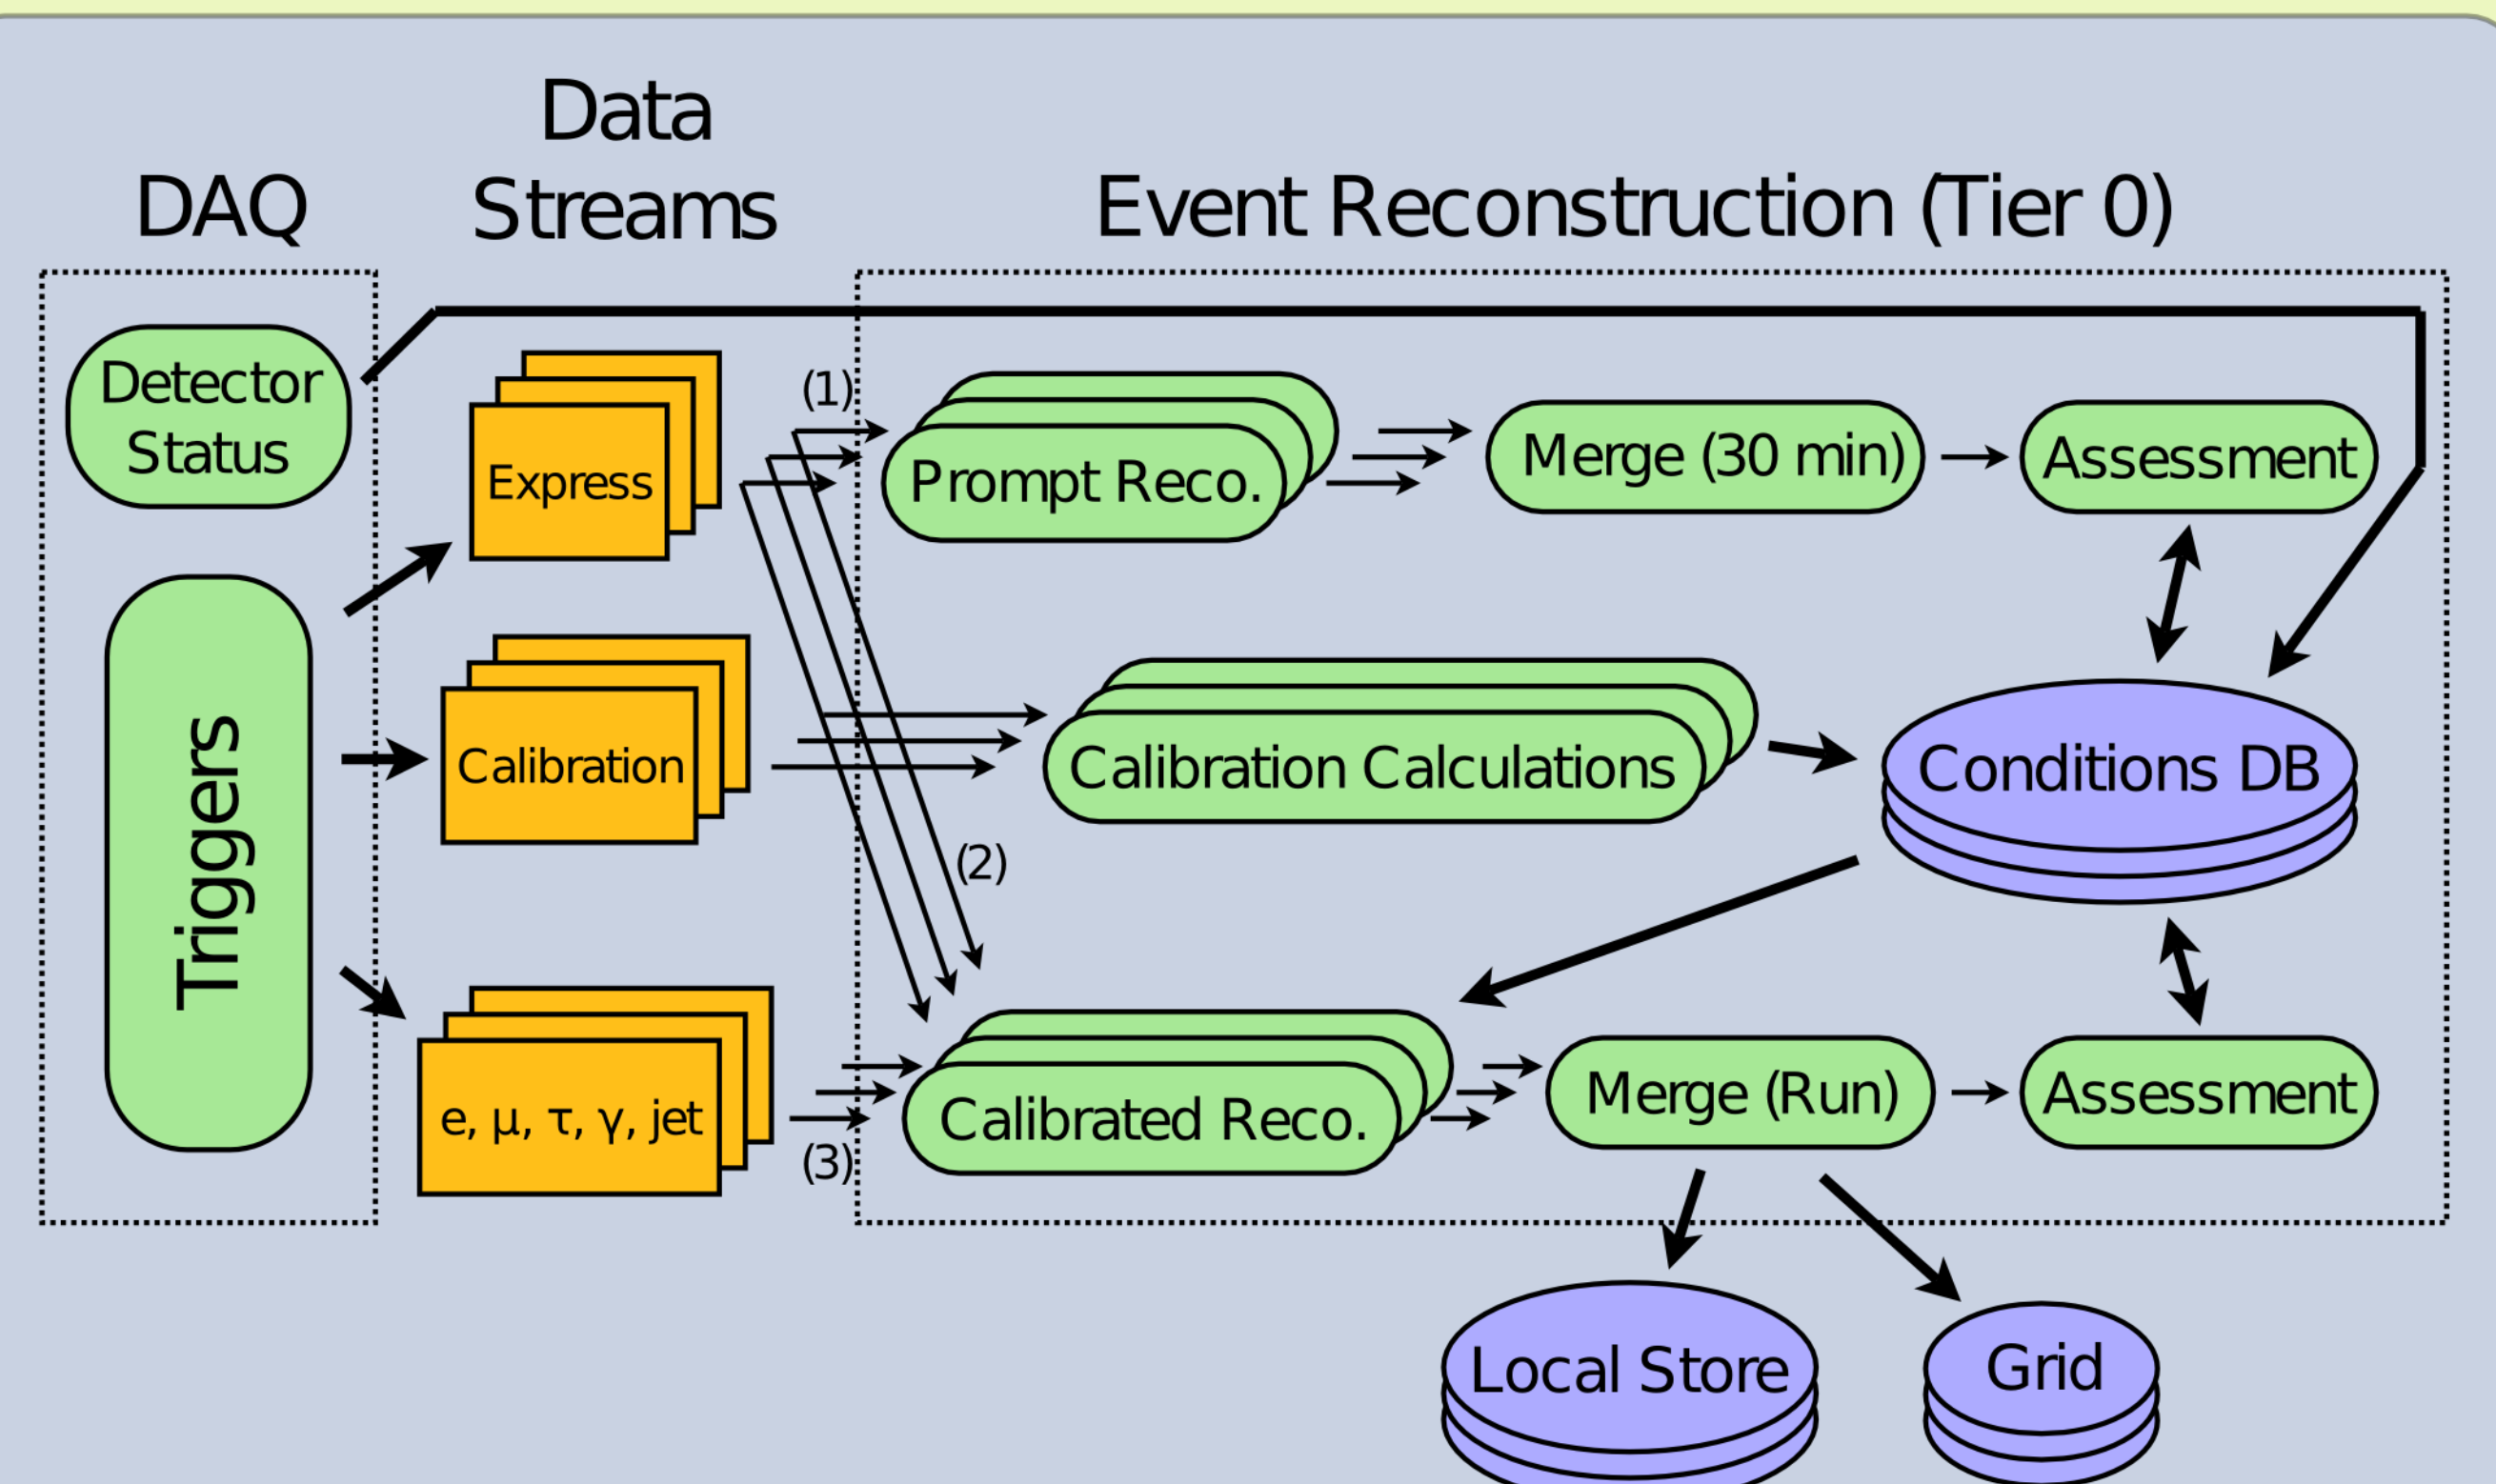


Figure 1: A schematic overview of the Data Quality structure. Data enters the system from the Data Acquisition (DAQ) and is separated into different event streams depending on the trigger status. The express and calibration streams are used to update calibration information in the conditions database. The calibration information is used in the second reconstruction pass, which is ultimately used for physics analysis. Further passes can be made if necessary. An example showing the effect of calibration can be seen in figure 2. The assessment stages are further described in figure 3.

The Calibration Loop

There are many parameters which can vary between runs and need to be taken into account in the reconstruction before the data can be used for analysis. Some of these parameters can only be easily measured during a reconstruction pass. Figure 2 shows two examples: before calibration, a nominal beamspot is assumed in the reconstruction of tracks, causing an incorrect measurement of the distance of closest approach. In the second reconstruction pass, the actual beamspot information is available.

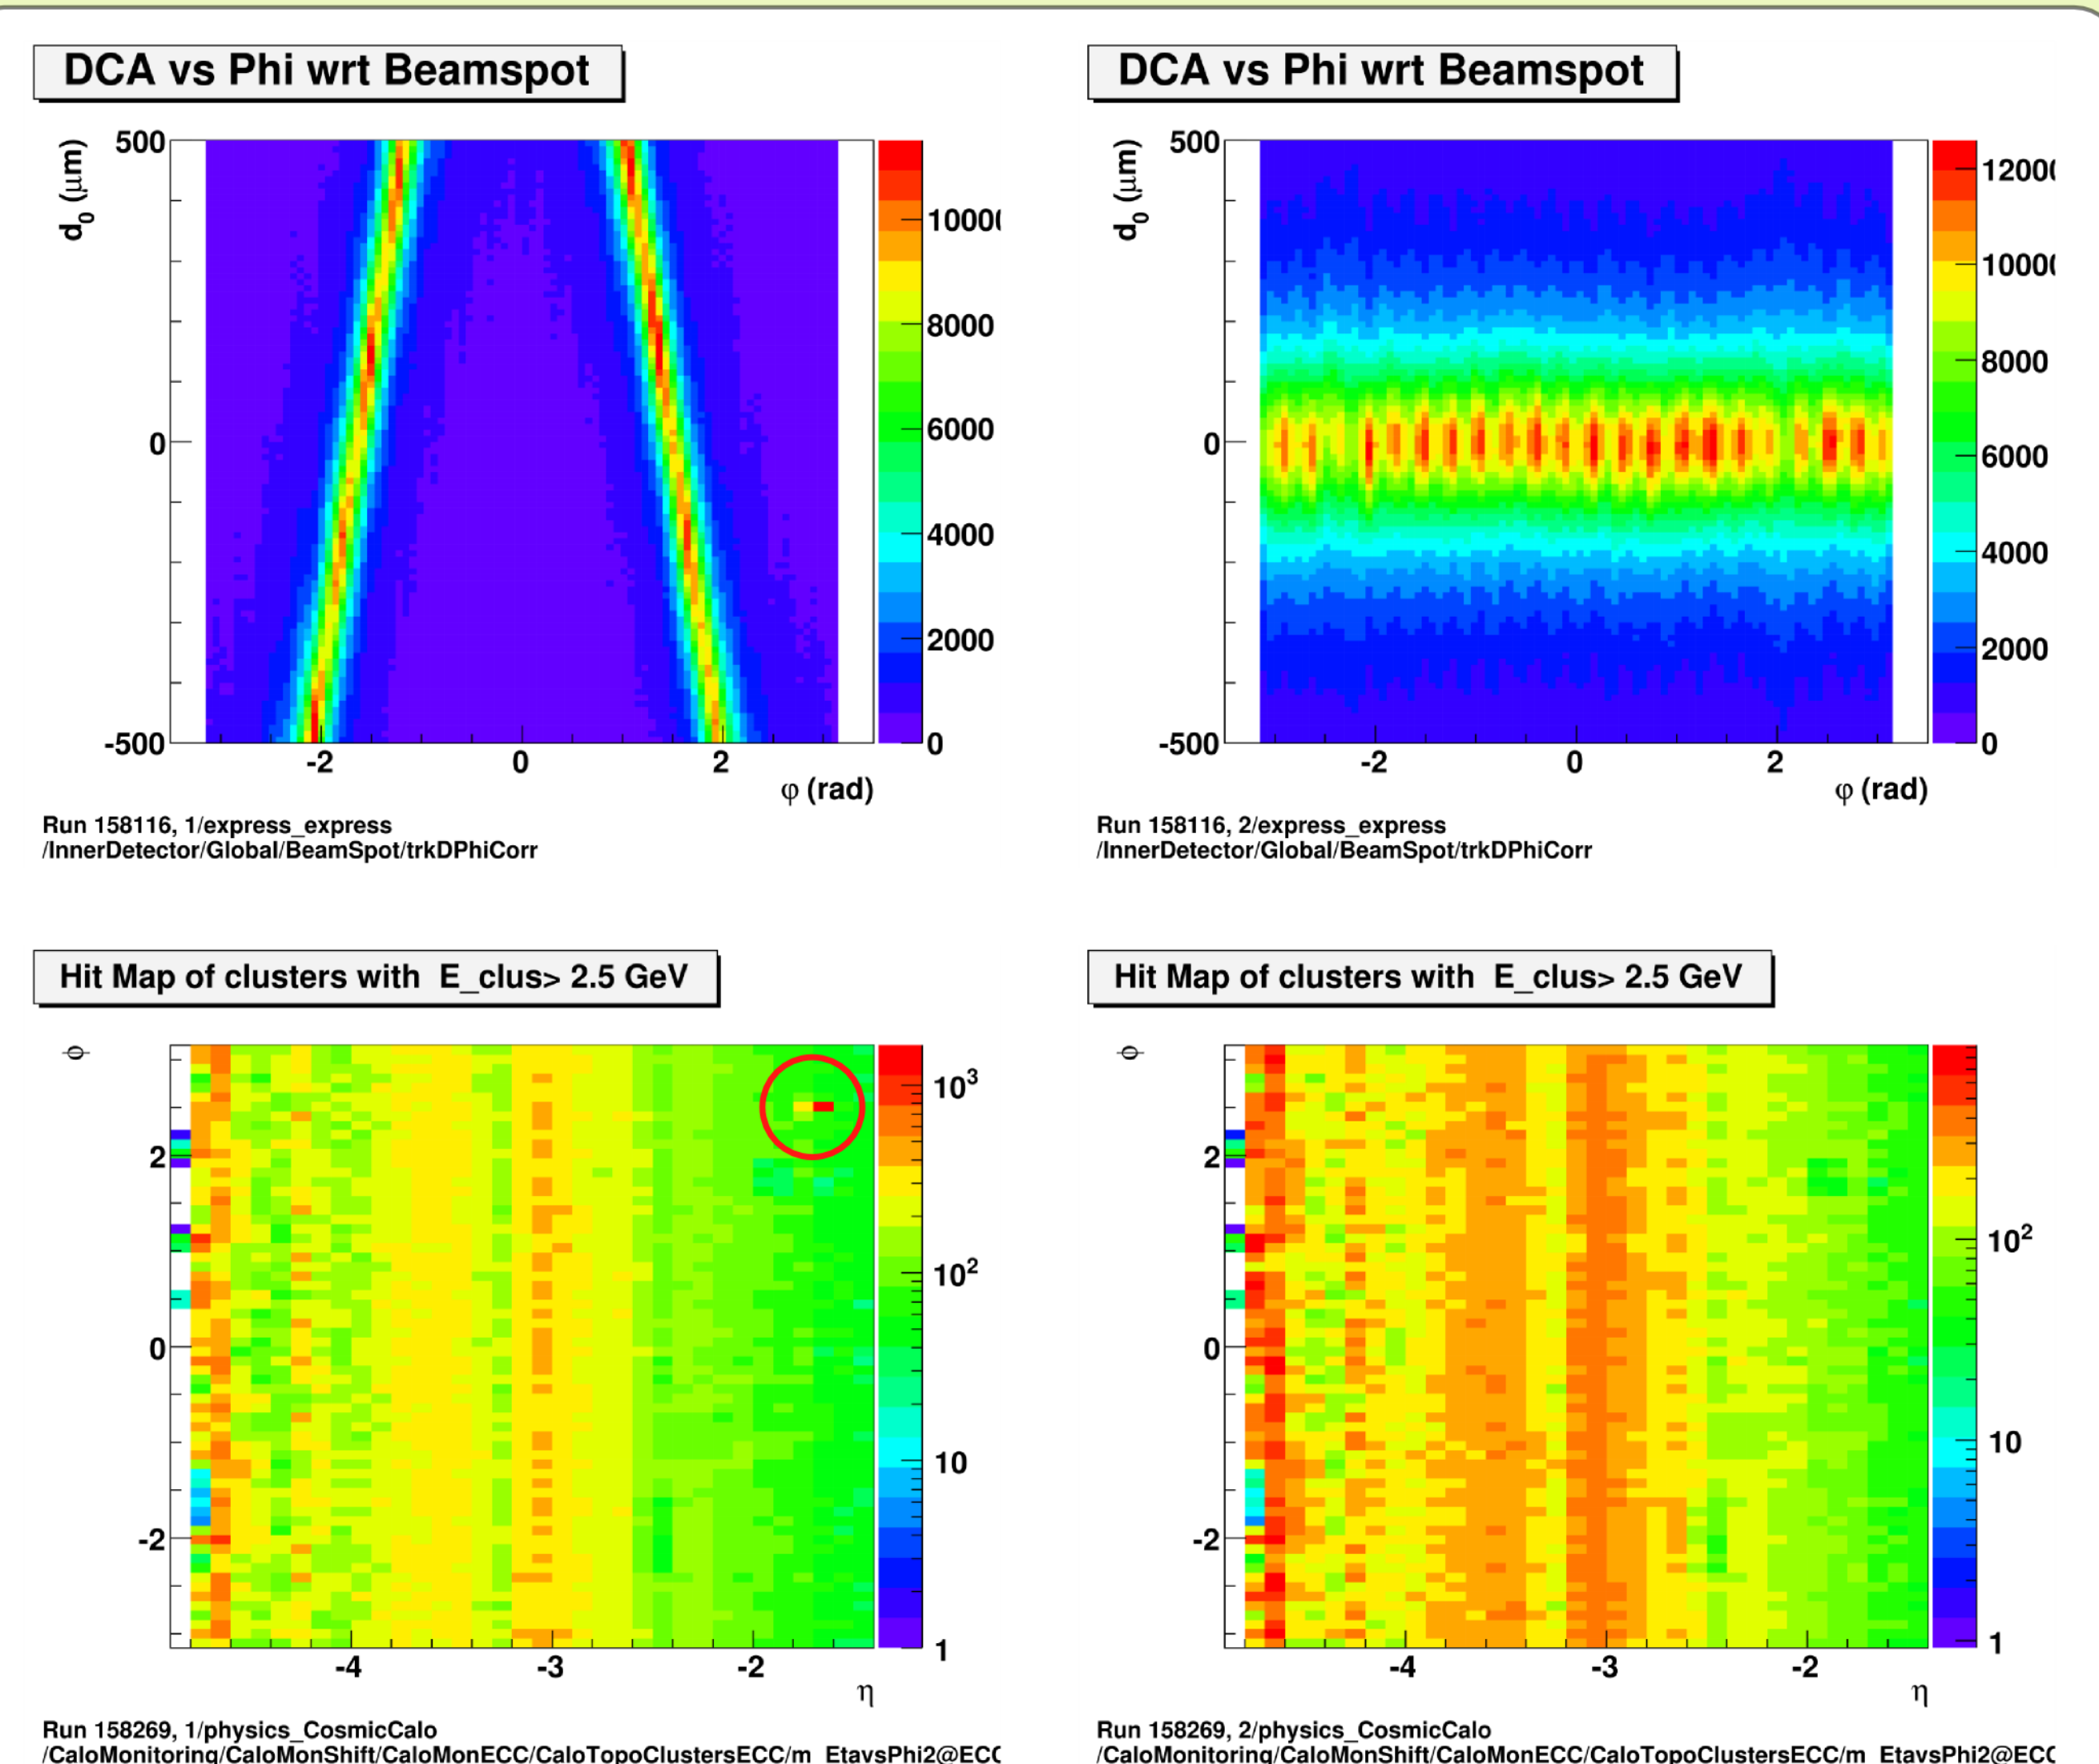


Figure 2: The above plots are example histograms from the offline DQ monitoring webpages. *Left:* From the express stream. *Right:* The same stream after an iteration of the calibration loop. *Upper:* This shows the distance of closest approach of tracks to the beamspot versus phi in run 158116. After calibration the beamspot is set correctly. *Lower:* Calorimeter topocluster occupancy for an endcap side in run 158269. Before the calibration loop there is a hot spot (circled); afterwards it is removed.

Monitoring Infrastructure

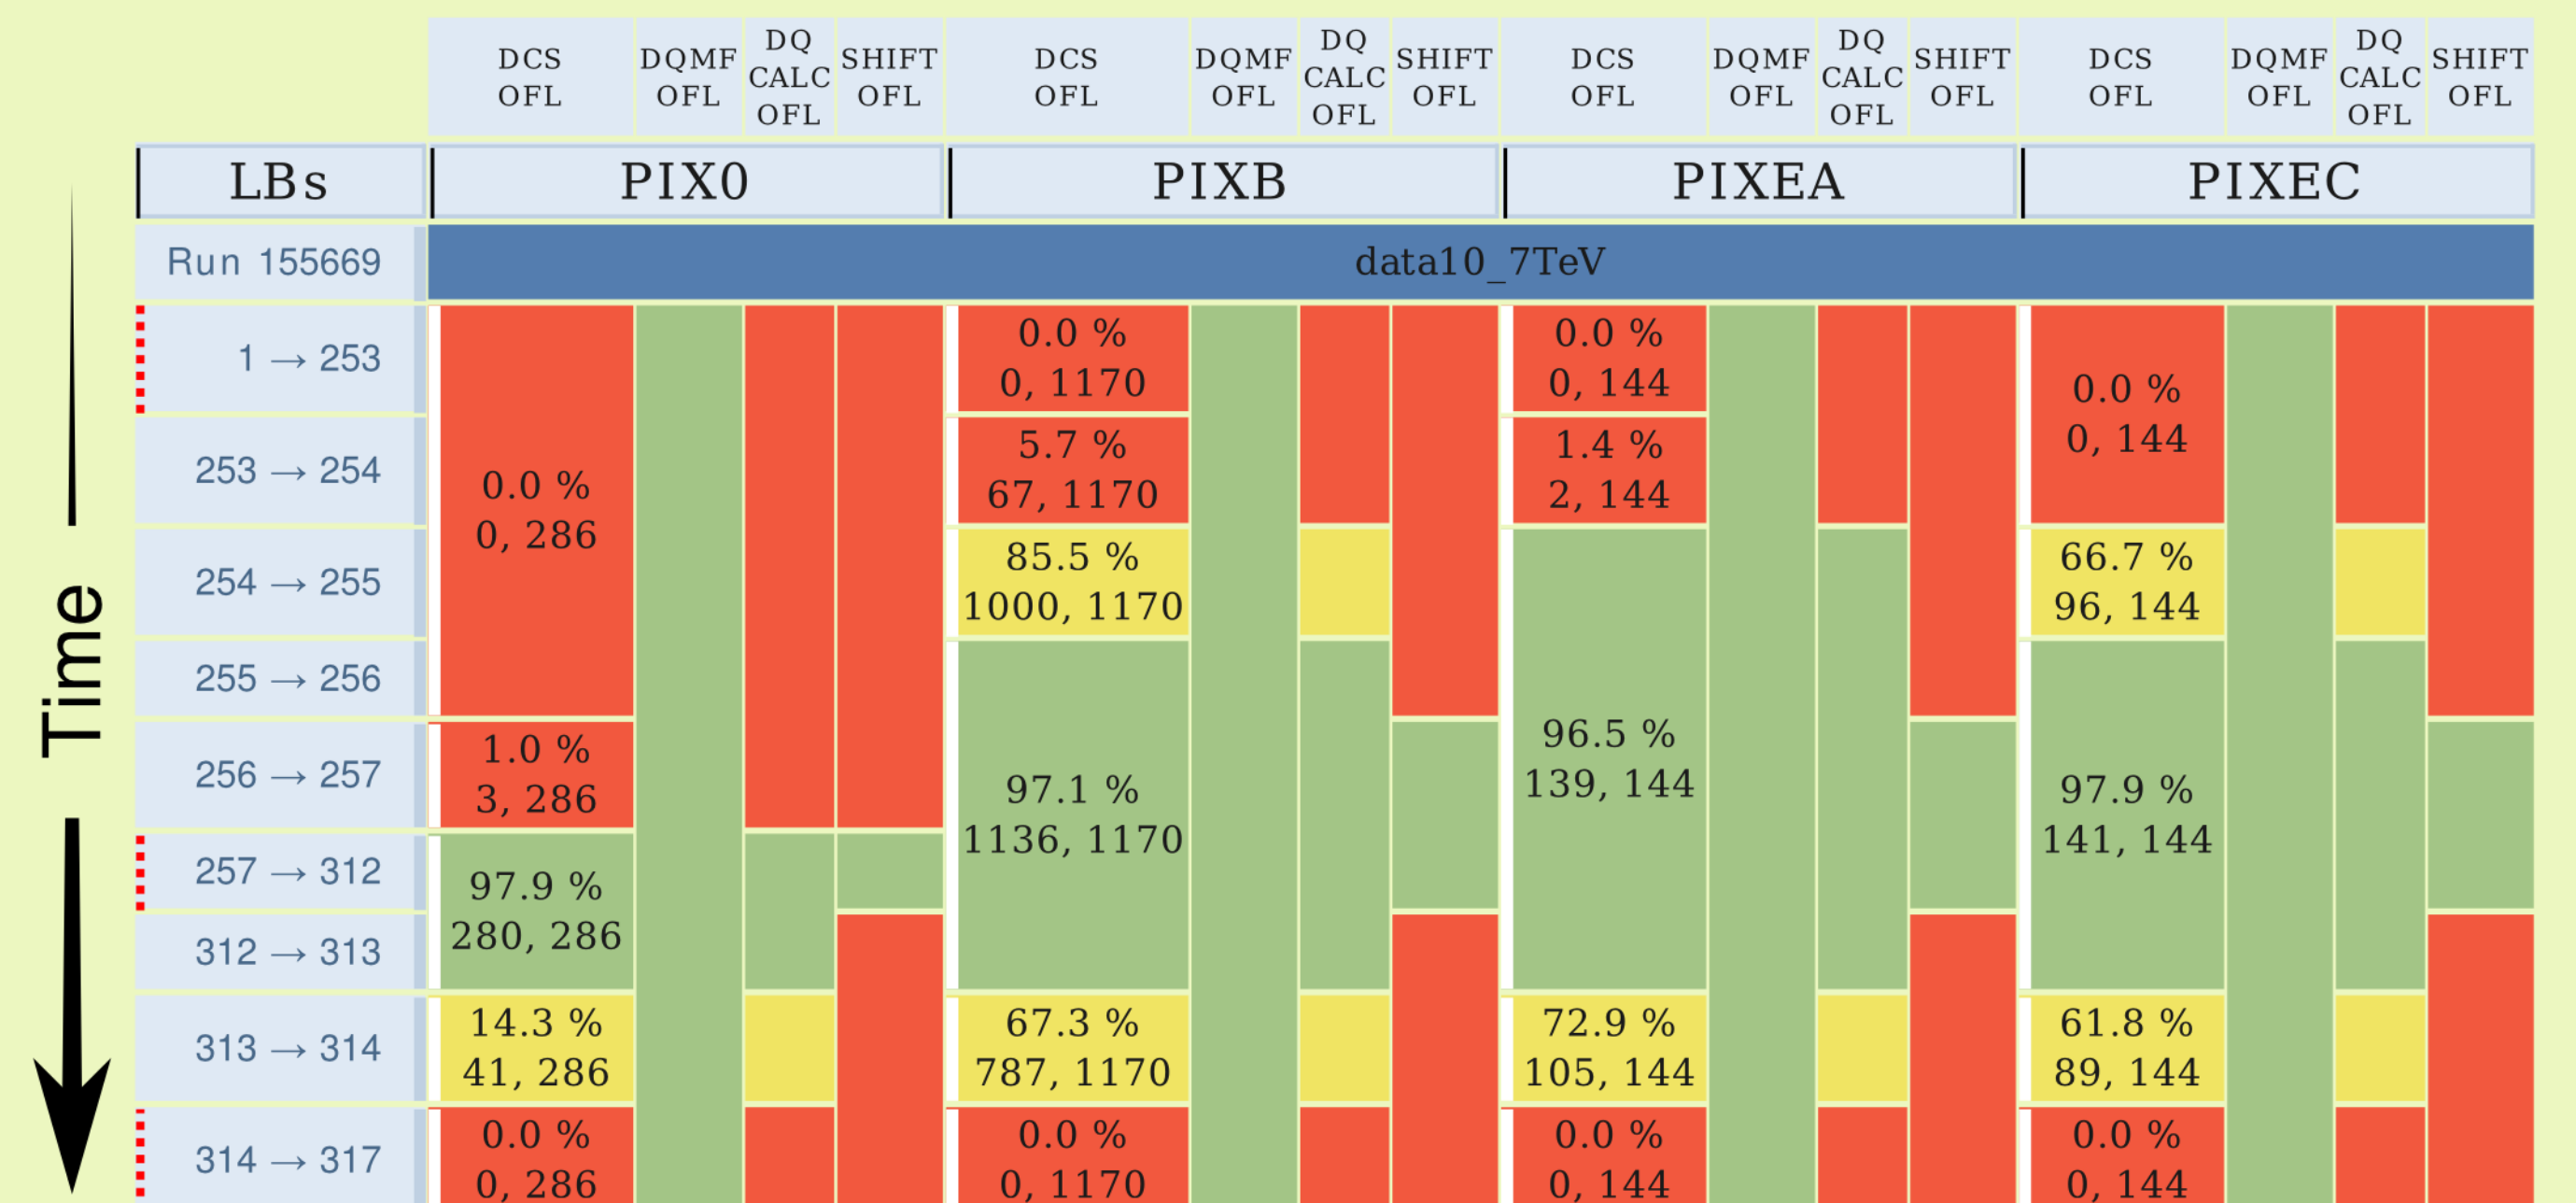


Figure 3: A display showing the DQ flags for the four pixel detector DQ regions (PIX0, PIXB, PIXEA and PIXEC) as determined by different methods. The warm start is seen where the detector modules are gradually brought into a working state as the run starts, and the reverse occurs when the run ends. The numbers shown for the DCS are the percentage of modules considered working, the absolute value and the number in configuration. Also shown are the automated histogram checks (DQMF), summary (DQCALC) and offline shifter decisions.

Data quality decisions are made by people as well as several automated systems. Figure 3 shows an excerpt of these flags for the pixel detector. At the moment, a shifter makes the ultimate decision. In the future, they may passively agree with the automated decision.

The Data Quality Flags

- : good for analysis
- : may be good after some corrections
- : the data are unrecoverable
- : unable to reach a decision
- : detector is switched off

The histogram analysis framework produces, per active stream, 20,000 histograms (100 MB) per run, along with 700 histograms (6 MB) per ten luminosity blocks. Originally, these were all dumped as images so that they could be served over the web as static content, but the disk space and CPU usage per run were not sustainable. Given that the vast majority of histograms are not looked at for each run, it is better to generate the image files as they are requested and to cache them for a short time.

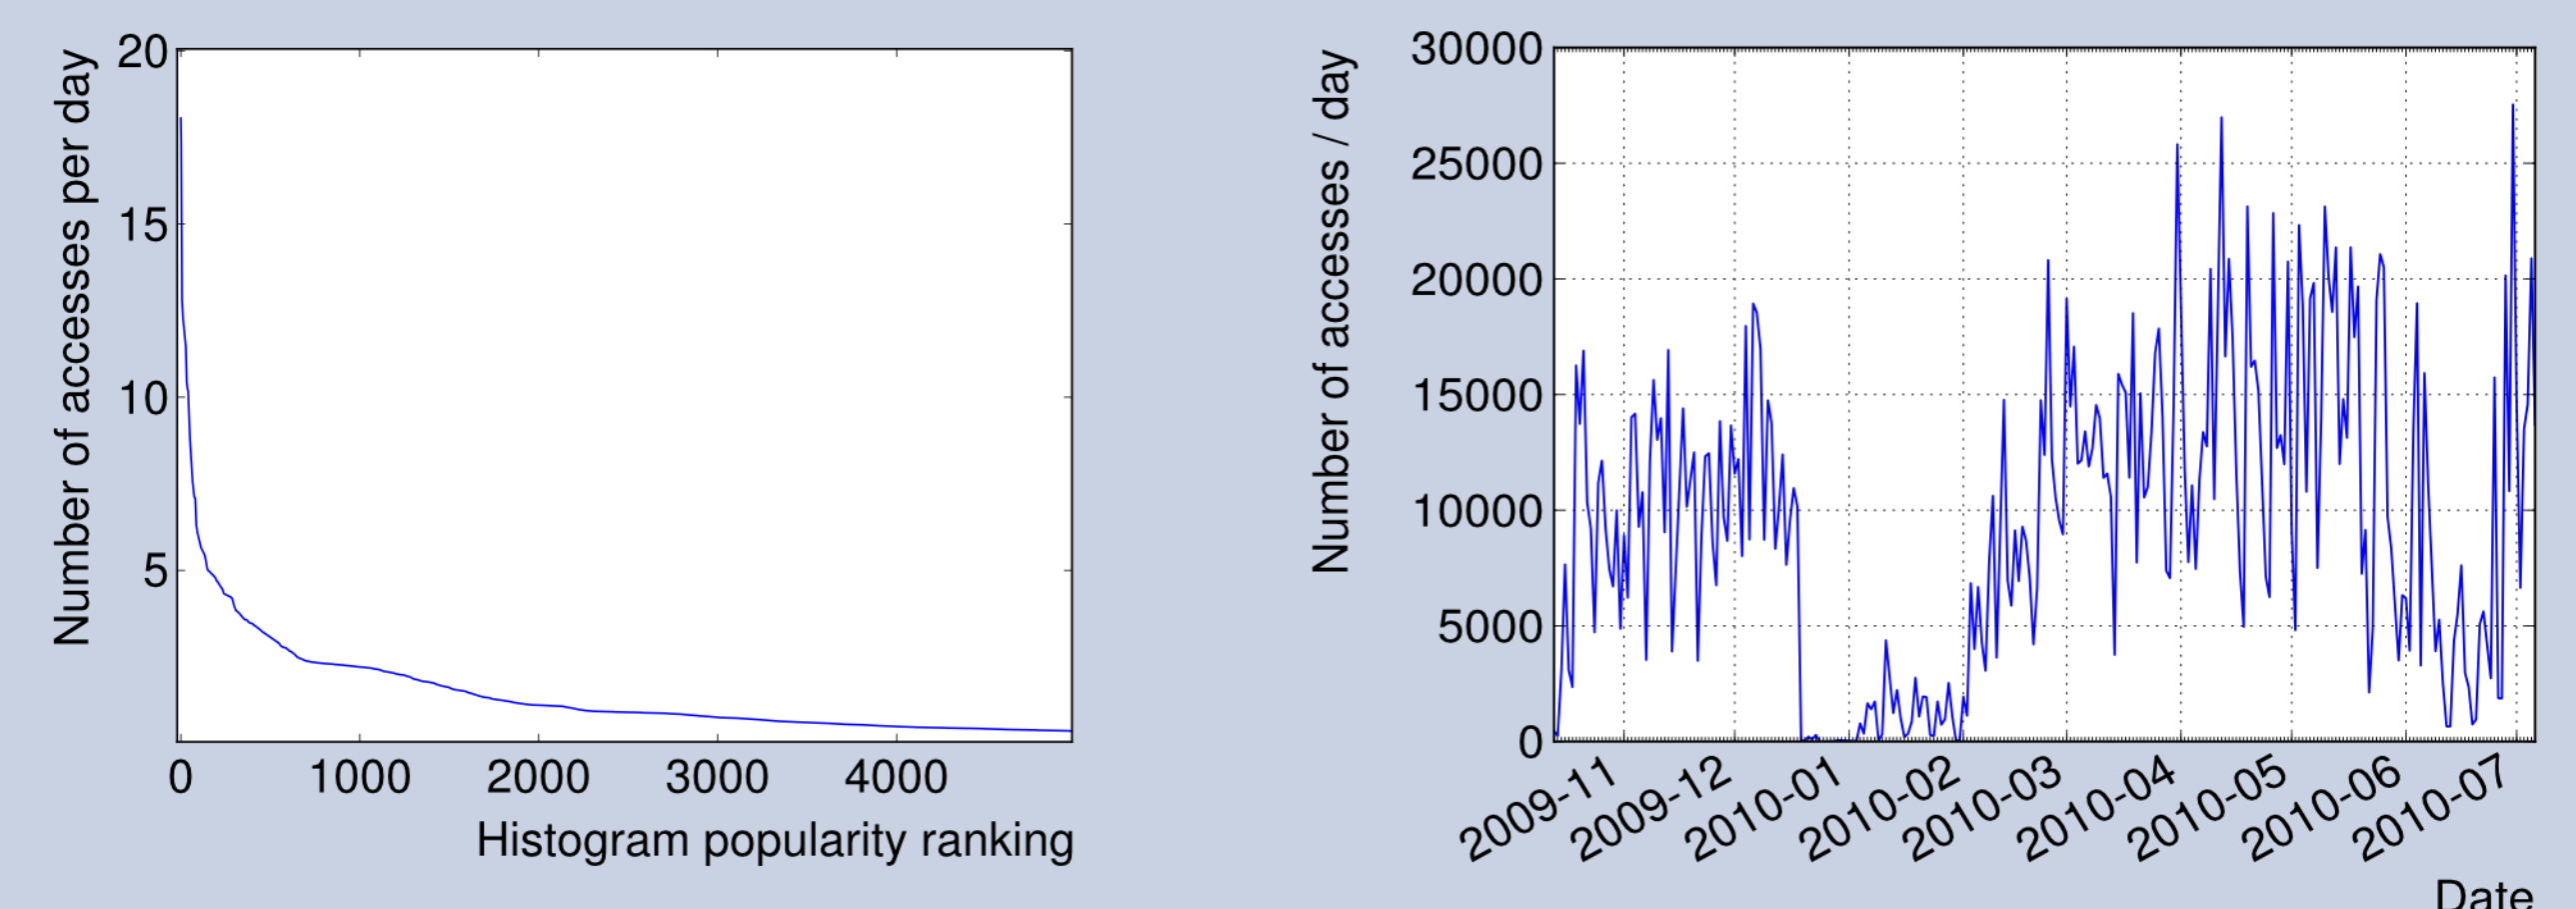


Figure 4: Traffic to the histogram web server. *Left:* Number of hits to an individual histogram, showing the most popular histogram on the left and decreasing with a long tail of histograms viewed. *Right:* Number of accesses to the histogram server per day. Dips can be seen in periods with no beam.

Figure 5 shows how flags are passed down a chain and merged at each stage. For each region, the Detector Control System (DCS), online shifter and DQMF (Data Quality Monitoring Framework) decisions are merged into a single flag which can be overridden later by an offline shifter. This final decision is merged across systems into Virtual Flags. These are used to make decisions for physics objects. Finally, the green virtual flags are taken along with information such as the state of the LHC to make a "good runs list", a standard ATLAS XML file containing a list of run numbers and luminosity blocks which are considered good for a particular analysis.

Lastly, there is currently no system to record whether a person has signed off on an automated flag. At the moment all flags are always checked by a human, but in the long term a record will be necessary as more trust is placed in the automated systems.

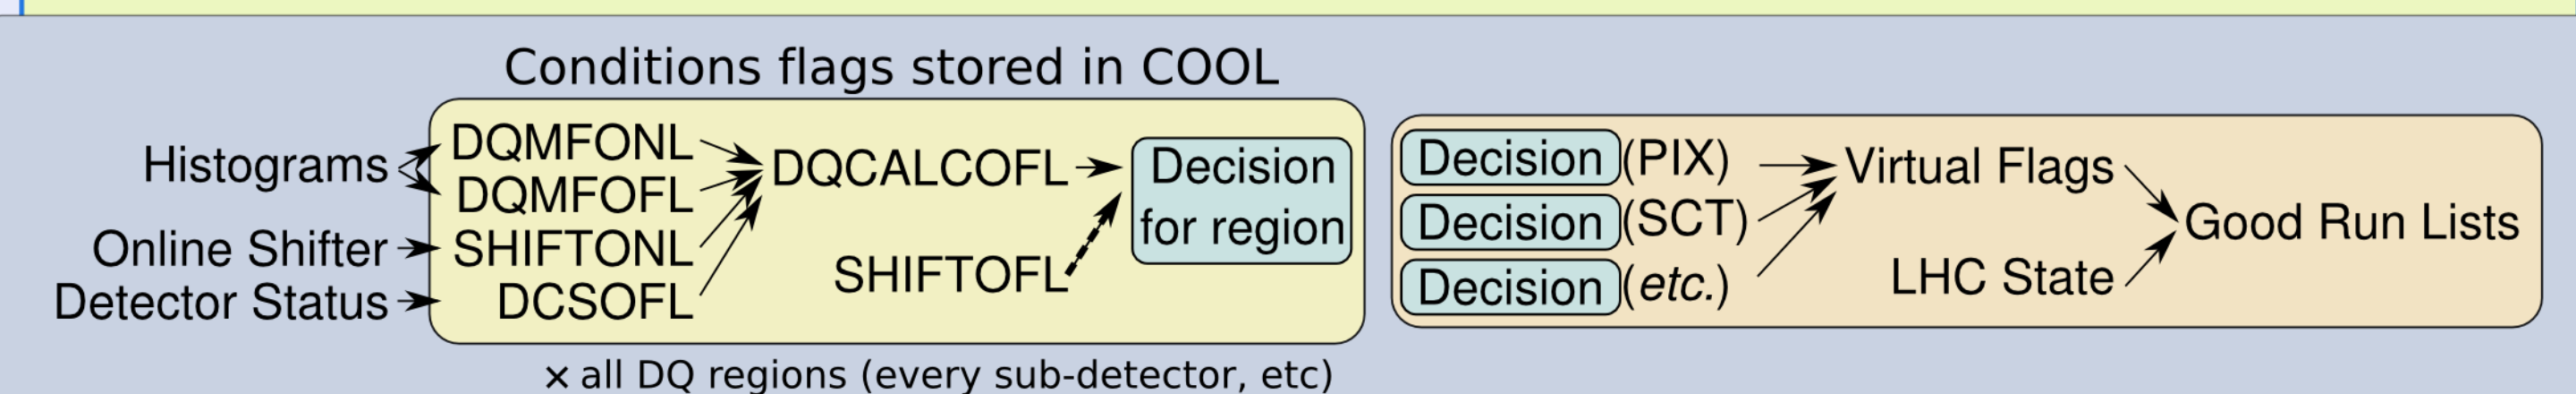


Figure 5: An illustration of the DQ chain, showing how raw status information makes it into a good run list XML file and affects analysis. A virtual flag can require different parts of the detector depending on the type of analysis.

Human factors

The ATLAS DQ group has been successful at promptly assessing problems with the data and marking which data is good for analysis. Nevertheless, there is some room for improvement. In particular, given the large scope of the DQ chain, it can be hard to communicate subtleties downstream to the people performing physics analyses. Powerful tools are provided globally for generating good runs lists, but often it is easy to misunderstand a feature of a flag or a system.

| Inner Tracking Detectors | | | Calorimeters | | | | Muon Detectors | | | |
|--------------------------|------|-----|--------------|---------|---------|------|----------------|------|------|------|
| Pixel | SCT | TRT | LAr EM | LAr HAD | LAr FWD | Tile | MDT | RPC | TGC | CSC |
| 94.9 | 98.3 | 100 | 95.4 | 96.3 | 99.8 | 100 | 97.6 | 98.3 | 97.8 | 98.3 |

Luminosity weighted relative detector uptime and good quality data delivery during 2010 stable beams at $\sqrt{s}=7$ TeV between March 30th and June 5th (in %)

Figure 6: This table shows the overall luminosity weighted fraction of the data which was marked good for data analysis. Most of the inefficiencies occur during "warm start", when the detectors are switched on after the LHC declares stable beams.