




Prompt Processing of LHC Collision Data with the ATLAS Reconstruction Software

Nick Barlow, for the ATLAS Collaboration

CHEP 2010, Taipei

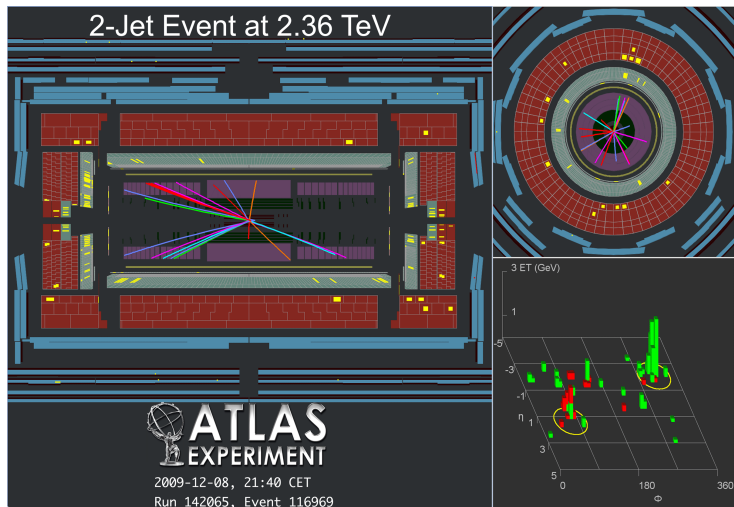
LHC milestones (1)

- LHC delivered first collisions, at 900GeV CM energy, on 23rd November 2009.




LHC milestones (2)

- December 8th 2009, became the world's highest energy collider, with 2.36 TeV collisions.




LHC milestones (3)

- March 30th 2010, first collisions at 7TeV CM energy.



LHC milestones (4)

- On each of those occasions, data were processed, scanned, and the result (with event displays) made available, within *minutes* of the collisions taking place!
- Luminosity delivered by LHC increasing rapidly...



- Despite this, data continue to be reconstructed, monitored, and distributed ready for analysis in very short timescale.

Data format nomenclature


- **RAW** (a.k.a. ByteStream) files - output of ATLAS DAQ.
- **ESD** (Event Summary Data) - POOL format, used for detailed analysis.
- **AOD** (Analysis Object Data) - POOL format, intended for most user analysis.
- **dESD (dAOD)** - *derived* ESD (AOD), formerly known as DPD (Derived Physics Data).
 - Interesting events skimmed from ESD (AOD) files, and/or uninteresting containers within events removed.
 - Several variants, used by specific performance (physics) groups.
- **TAG** (formerly also **TAG_COMM**) - ROOT TTree, for fast browsing of data.
 - Also uploaded to TAG database - can be queried with different event selections and used to provide “skimmed” ESD or AOD files.
- **HIST (NTUP)** - ROOT histograms (TTrees) used for data quality monitoring (and calibration).

- Data are organized into inclusive *streams*, based on trigger chains.
- **Calibration** streams:
 - Partial events, used by specific ATLAS subdetectors.
- **Express** stream:
 - Used to check the data quality, calculate calibration constants..
 - Subset of the physics data, corresponding to 10Hz
- For early data, **physics** streams were based on L1 trigger:
 - **MinBias, MuonswBeam, L1Calo, L1CaloEM, CosmicCalo, CosmicCaloEM**
- As luminosity increased, moved to final streaming model:
 - **MinBias, EGamma, Muons, JetTauEtMiss, CosmicCalo**

- Focus on robustness, rather than latest features.
- *Production cache* built on top of full ATLAS release, can be built and deployed within matter of hours if necessary.
- “Frozen Tier0” policy - no changes allowed that affect reconstruction.
 - Ability to combine first-processed data with reprocessed data for analysis.
 - Also have MC re-reconstructed with same release.
- Last reprocessing was in May, all data since then has been processed with a compatible release.
- Autumn reprocessing campaign getting under way with new release, with many code improvements.
 - Will be used at Tier-0 for first 2011 pp collisions.

- Extensive tests run on nightly builds of production cache:
 - **ATN** (ATLAS Testing Nightly) tests - run on very small number of events, mainly to test that jobs do not crash in various configurations.
 - **RTT** (RunTime Tester) tests - run on a few more events, more possibility to analyze output.
 - Both are checked daily by dedicated shifters, who file bug reports to appropriate trackers.
- Also more detailed tests:
 - **TCT** (Tier-0 Chain Test) - runs full Tier-0 workflow (RAW→ESD,AOD,TAG etc.).
 - Tools to compare output files from successive nightlies also help enforce “frozen Tier-0” policy.
 - **BCT** (Big Chain Test) - several million events, run on the grid.

ATLAS computing model



- Tier-0 tasks:
 - First pass ESD, AOD, dESD, TAG file production.
 - First pass calibration and alignment processing.
 - Upload TAG files into TAG databases
 - Merging of RAW files.
 - Registration of data files (RAW and derived) with ATLAS DDM (Distributed Data Management).

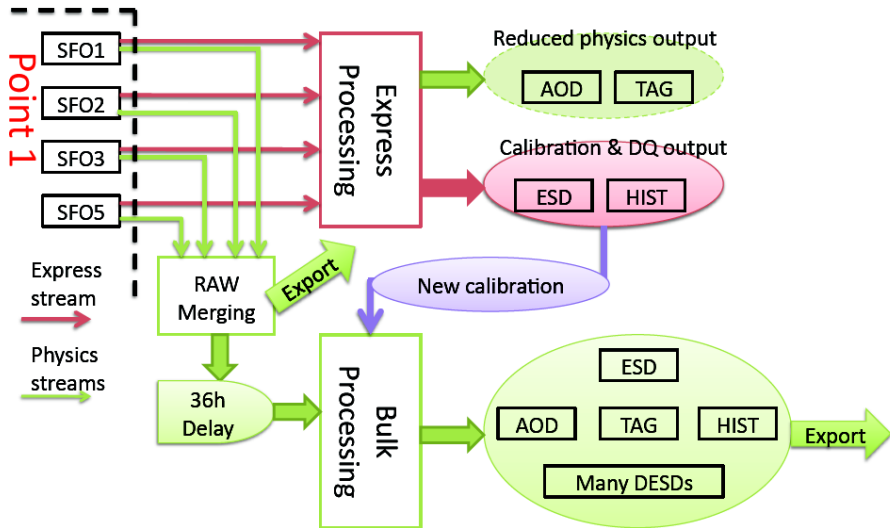
Tier-0 throughput and performance

- In 2010, Tier-0 has run $> 4\text{M}$ jobs. ¹
 - Includes express-stream reco, bulk reco, merging of RAW, ESD, AOD, TAG files, TAG upload.
- Extremely high success rate!
 - Of 163588 reconstruction jobs in past 30 days, 19 jobs (0.012%) crashed due to software problems.
 - 7 of those failures were due to corruption in RAW files.
- Average reconstruction time per event varies according to software release, stream pile-up, but is typically 11s.
- Can cope with 200Hz output of ATLAS at 100% LHC lifetime.
 - LHC beam time less than 100% this year, as expected for new collider!
 - ATLAS recording rate recently between 200Hz and 300Hz, but Tier-0 still keeping up..

¹more details in talk by I.Ueda.

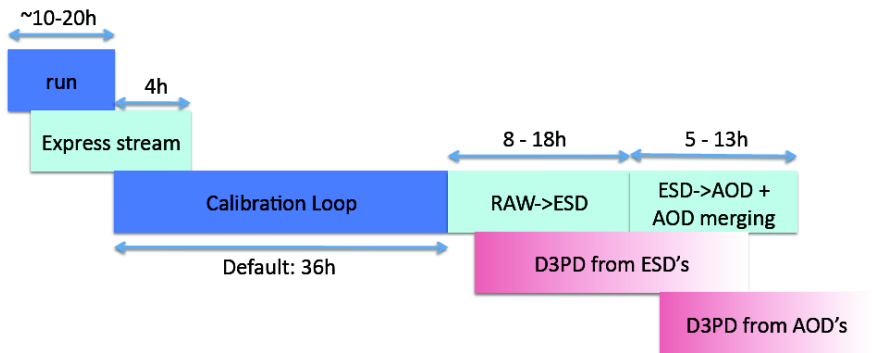
- Two-pass reconstruction:
 - **Express stream** reconstructed at Tier-0, and outputs sent to CERN Analysis Facility (CAF) for analysis.
 - e.g. Beamspot determination (run at Tier-0), dead and noisy channel corrections.
 - Calibrated run is assessed at the daily Data Quality meeting before bulk processing starts.
 - Calibration constants are uploaded to COOL database, then used in bulk processing of the rest of the data, which starts automatically, 36 hours after the end-of-run time.
 - If calibrations are not ready in time, start of bulk processing can be delayed by manual intervention.
 - Database upload tools disallow calibration updates for runs that have already been bulk-processed, to guarantee reproducibility.

Prompt processing flowchart



Data availability

- Many users perform final analysis on “D3PD” ntuple format.
- Can be produced from either (d)ESD or (d)AOD, depending on required level of detail.



Accelerated strategy for conferences

- With luminosity increasing so fast, one extra day of data can make a big difference to physics plots and results.
- In exceptional circumstances, can reduce latency between data being recorded and being available for analysis:
 - Fast-track calibration loop - when finished begin bulk reconstruction jobs manually rather than waiting for 36 hours to elapse.
 - Manual killing of “tails”, last 0.1% jobs can hold up processing of a run.
- For ICHEP this summer, data taken on a Monday was reconstructed, analyzed, and results were approved, in time to be presented on Friday the same week!

- Prompt assessment of data quality (DQ)² is vital for feedback to detector operations.
- Online DQ for each subdetector is loaded into COOL conditions database, for each lumi block (2 minutes).
- Tier-0 populates histograms every 10 minutes for further DQ checks.
- More than 20,000 histograms generated on demand per run per stream, and cached for future use.

²For more details see talk by J. Frost.

- Prompt reconstruction of ATLAS data is working well!
 - Result of extensive preparations in cosmic ray and single beam runs.
- Short latency between raw data being recorded, and reconstructed, calibrated, data being signed-off by DQ and becoming available to the collaboration.
 - Can be reduced even more in extreme circumstances...
- Systems will improve, become more automated.
- Ready for the challenge of more and more data from LHC!