



ATLAS: Computing Model Document

Richard P Mount
SLAC National Accelerator Laboratory





Status of the ATLAS sections

Cristi's comments:

- Addressed most already – so far without deleting the comments themselves.
- Now making a final pass to see where further changes would help – we see no particular issues in addressing them

Amber's comments:

- Specific responses on subsequent slides.

Eckhard's comments and edits:

- Edits accepted – thanks!
- Specific responses to comments on subsequent slides.

Specific responses

2.2.2 Anticipated event rates and data streams

For 2015-2018 ATLAS uses as the baseline an average trigger rate of 1 kHz as event input to the Tier-0 for prompt reconstruction. For 25 ns bunch spacing, the average

The menu will be adjusted during a fill, setting/removing pre-scales, bringing in/out triggers with different thresholds, to keep the rate at around 1 kHz max. The system could cope with more rate temporarily, however we don't think we will go to more than about 1.1-1.2 kHz for short periods, mostly to cope with operational needs. Especially in the beginning, we expect to be significantly below the 1 kHz, moving up as the run progresses.

Amber Boehnl..., 1/29/14 7:20 PM

Comment [2]: So this is a factor of 2 average-what is the peak?

Specific responses

2.2.2 Anticipated event rates and data streams

The primary motivation for the doubling of the HLT output rate is the desire to keep the effective single electron and muon p_T threshold at 25-30 GeV and therefore well below the Jacobian peak from W boson decays. This strategy was very successfully employed for Run-1 and these triggers were used by the majority of the ATLAS analyses. The single lepton triggers allows ATLAS to keep an inclusive and unbiased trigger for most electro-weak processes, for example associated production of the Higgs boson in the WH channel. Options for relying on more exclusive selections are under study, but generally incur efficiency losses for many physics channels and carry an increased risk of missing new physics signals. At $\mu=40$, the rate of the single lepton triggers alone are expected to reach 4-500 Hz with the rate dominated by irreducible W and Z boson decays. This can be compared to a typical average in 2012 of approximately 120 Hz. Other triggers will also see increased rates, but this increase is expected to be more modest as thresholds will need to be raised to keep the L1 rate below the maximum possible.

Agreed. This text also appears in the introduction as applying to ATLAS and CMS. It has been replaced by a reference to the introduction.

Eckhard Elsen 2/7/14 8:08 PM

Comment [3]: Basically this is a repetition of the introduction. Suggest less detail in the intro.

Specific responses

2.2.5 Data Categories

analysis needs. The same mechanisms are expected to remain in Run-2, with further improvements and refinements, e.g. by fine-tuning the dynamic data replication triggers (data popularity measurement).

We lean towards creating a table/section as requested and adjust the text accordingly. It will take us some more time (~ days) to make all the required edits.

Amber Boehnl..., 1/29/14 7:28 PM

Comment [7]: There quite a few development projects that are listed. It would useful to have them in a table with an explanation, an estimate of the effort needed, and a prioritization.

Eckhard Elsen 2/11/14 9:49 PM

Comment [8]: Estimate of impact. Right place to mention?

Cristinel Di..., 11/22/13 10:52 AM

Comment [9]: Examples?

Specific responses

2.2.8 Description of Workflows

AOD2AOD Data re-processing

In Run-2 the state of the art reconstruction and fixes will be applied in dedicated AOD2AOD re-processings, which will happen a couple of times per year. The AOD2AOD reprocessing will be commissioned during the 2014 ATLAS Data Challenges. This new

The two reprocessing are expected based on experience in Run 1. Changes to software, calibrations and algorithms build pressure for reprocessing.

Amber Boehnl..., 1/29/14 7:34 PM

Comment [13]: What motivates the number of times per year?

Specific responses

2.2.10 Non-event data

volumes may grow by a factor of 2 to 3 for Run-2). ADCR is much more difficult to predict, given the uncertainty in the offline event rate and size, number of files, as well as the current major evolutions of the DQ2/Rucio and ProdSys-II systems. The LFC service and DB will not exist any longer but Rucio will include its functionality.

In Run-2, ORACLE support will not be needed at all T1s. Only a few T1s will continue to provide ORACLE and Frontier services. The reduction in number of required sites is a consequence of an improved overall stability of the distributed system.

[23] The uncertainty over the ADCR database size is not a problem for Rucio. Tests of Rucio already indicate an ability to scale beyond likely ATLAS needs.

[24] Other sites will also use the Frontier client communicating via Squid caches with the small number of Frontier servers.

Amber Boehn..., 1/29/14 7:43 PM
Comment [23]: Is this a source of risk to the Rucio deployment?

Amber Boehn..., 1/29/14 7:44 PM
Comment [24]: How do the other sites get the relevant info? A central service?

Specific responses

2.6 Data Preservation and Open Access

embargo period, be used by non-collaboration members. ATLAS is currently evolving a policy on these matters. The policy must inform the other software and computing efforts, and will require effort over and above the normal exploitation of the data.

We are of course interested on how the other experiments are handling these issues. At this point we are still developing our strategy, but a common approach in the future may indeed be an advantage.

Amber Boehnlein 1/30/14 7:01 PM
Comment [18]: Is there any benefit in CMS and ATLAS collaborating?