

Preparing ATLAS Distributed Computing for LHC Run 2

Eric Lançon on behalf of the ATLAS collaboration

CEA-Saclay

Irfu, 91191 Gif/Yvette Cedex, France

E-mail: eric.lancon@cea.fr

ATLAS software and computing is in a period of intensive evolution. The current long shutdown presents an opportunity to assimilate lessons from the very successful Run 1 (2009-2013) and to prepare for the substantially increased computing requirements for Run 2 (from spring 2015). Run 2 will bring a near doubling of the energy and the data rate, high event pile-up levels, and higher event complexity from detector upgrades, meaning the number and complexity of events to be analyzed will increase dramatically. At the same time operational loads must be reduced through greater automation, a wider array of opportunistic resources must be supported, costly storage must be used with greater efficiency, a sophisticated new analysis model must be integrated, and concurrency features of new processors must be exploited. This presentation will survey the distributed computing aspects of the upgrade program and the plans for 2014 to exercise the new capabilities in a large scale Data Challenge.

*International Symposium on Grids and Clouds (ISGC) 2014
Academia Sinica, Taipei, Taiwan
23-28 March, 2014*

1. Introduction

With the anticipated increases in data rates foreseen for Run 2, the ATLAS experiment is updating its computing model to reflect the advances in understanding of the most effective ways to use the distributed computing and storage resources, based upon the experience in LHC Run 1.

The anticipated increase in data rate is expected to be a factor of ~ 2 higher rate in Run 2. The increasing trigger rates are motivated by a desire to maintain similar trigger thresholds and sensitivity to Higgs physics and to potential new physics as in Run 1. With the anticipated increase in instantaneous luminosity and the increase in cross sections due to the higher beam energy the output rates in ATLAS will rise to close to 1kHz.

In 2012, the average trigger rate during proton-proton collisions in stable beam conditions was ~ 550 Hz. This includes all physics streams (prompt +delayed), which, with an average compressed RAW event size of 0.8 MB/event, gave an average 440 MB/s streaming rate of RAW events.

For 2015-2018 ATLAS expects an average trigger rate of 1 kHz as event input to the Tier-0 for prompt reconstruction. For 25 ns bunch spacing, the average event size will be similar to the Run-1 (2012) values, at least in the first year when the pile-up average value at 25 ns is unlikely to exceed $\mu \sim 25$ and would go up to about 1 MB/event at the average pile-up of $\mu \sim 40$. The anticipated average streaming throughput of RAW data in Run-2 is thus between 800 MB/s and 1000 MB/s.

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 the majority of the ATLAS analyses used these triggers. 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.

This increase is close to a factor of two in 2015 compared with 2012 in terms of the needs for CPU and disk. However, it should be noted that in estimating this increased computing and storage need, assumptions have been made on several ongoing developments being able to significantly improve the current performances of the reconstruction, simulation, and analysis software, compared with that used in 2012. This work has resulted in a number of improvements, including: a reduction in the number of data replicas that are made, in order to minimise the storage requirement; reducing the number of (re-) processing passes, and often better organising when they are scheduled; making significant efforts to reduce event sizes, minimize the reconstruction times, and to control the memory consumption of the software.

The area of simulation is important, since significant CPU resources are devoted to this activity. Work is under way to improve the efficiency of full simulation, but also to use

parameterized or very fast simulations where possible, avoiding the full Geant 4 process where that is feasible.

2. The ATLAS Run-2 Computing Model

The ATLAS Run-2 Computing Model will introduce several innovations in ATLAS resource management with the aim of keeping the resource requirements in Run-2 within reasonable constraints. The presented changes are evolutionary updates, based on the Run-1 operational experience, which has evolved from the original ATLAS Computing Model [1].

A short summary of the changes with the highest impact on ATLAS computing resource requirement is as follows:

- ATLAS has set strict benchmark targets for the CPU/event consumption and AOD event sizes of data reconstruction, Monte Carlo production chain and analysis, and is currently investing a lot of manpower to achieve this.
- ATLAS is developing several flavours of fast simulation within the new Integrated Simulation Framework (ISF). The aim is to achieve an optimal balance between required precision for analyses and resource requirements.
- ATLAS has reworked its analysis model in order to optimize and reduce the resource requirements. The model involves innovations impacting resource requirements as follows:
 - an update to the AOD format, to make the analysis objects ROOT-readable, thus eliminating the need for ‘intermediate’ big ROOT-readable group analysis ROOT N-tuples used in Run-1.
 - Introducing AOD2AOD re-processings to reduce the CPU consumption of group analysis, which currently introduces reconstruction improvements at the group analysis level.
 - Introduction of a centralized ‘Reduction Framework’ to derive the group analysis specialized N-tuples with reduced effort and processing, thus reducing the CPU consumption.
- ATLAS has reduced, at the end of Run-1, the pre-placement policy of AOD replicas to one replica on Tier-1 disks and one replica on Tier-2 disks. ATLAS will consequently rely even more on the existing popularity-based dynamic data replication.
- ATLAS has in the new model introduced dynamic deletion scheduling of pre-placed copies based on the usage patterns. Over time this algorithm will gradually reduce the disk copies of unused pre-placed data sets on disk to zero and only custodial tape replicas will remain. The algorithm parameters will be fine-tuned with Run-2 operational experience.

- This also means ATLAS will rely more on tape access for AOD processing and analysis for older data sets.
- ATLAS has introduced the possibility of ‘over-spilling’ the prompt reconstruction from Tier-0 to Tier-1 centres in case of resource congestion on Tier-0 resources.
- ATLAS will use the HLT farm during sufficiently long time windows. The re-configuration of the HLT farm from TDAQ operations to Grid activities has been demonstrated to take less than one hour. The available time windows will thus be set by typical job durations and TDAQ operational constraints.

Table 1: Summary of ATLAS Data Formats

RAW	<i>Raw data</i> : a persistent representation of the event data produced at the ATLAS online cluster (involving High Level Trigger) in byte-stream format.
ESD	<i>Event Summary Data</i> : a C++ object representation contains the detailed output of the detector reconstruction and is produced from the RAW data. It contains sufficient information to allow particle identification, track re-fitting, jet calibration etc. thus allowing for the rapid tuning of reconstruction algorithms and calibrations. Transient for data, i.e. in a rotating buffer as described in the text.
AOD	<i>Analysis Object Data</i> : C++ object representation, contains a summary of the reconstructed event, and contains sufficient information for common analyses.
DPD	<i>Derived Physics Data</i> : Skims, Slims and Thins of AOD, prepared in Group and User Analysis and specific to one or a few analysis groups; written out in ROOT N-tuple format (see section 4).
NTUP	<i>Data in ROOT N-tuples</i> : In ATLAS synonymous to the DPD.
EVNT	<i>Event Data</i> : C++ object representation, contains the truth event information as produced by Monte Carlo generators (Alpgen, Sherpa, MC@NLO etc...).
HITS	<i>Hits data from full or fast simulation</i> : C++ object representation contains simulated energy deposits in active detector volumes and related particle information.
RDO	<i>Raw Data Object</i> : a C++ object representation of the byte-stream information (i.e. RAW data format), used predominantly in simulation.
DESD	<i>Derived Event Summary Data</i> : C++ object representation the DESD are derived from the ESD. Data reduction is used to select targeted events and store only the necessary information, taking into account the performance goals and rate estimates.
D(format)	<i>Further derived formats (DRAW, DAOD...)</i> : C++ object representation of data, where reduction is used to select targeted events and store only the necessary information.

3. The ATLAS data flow

The ATLAS prompt data reconstruction is performed at the Tier-0. The RAW data are processed in two steps within one job, producing first the ESDs and then the derived AODs and DESDs in the second step. The RAW data and the reconstruction outputs are exported to the ATLAS Grid storage according to the replication policy described later in the document. There is also a small set of ‘monitoring data’ in specialized file formats (ROOT files) which are produced in data (re)processing for specialized studies (e.g. data quality assessment). The data re-processing from RAW is performed at Tier-1s using the same workflow.

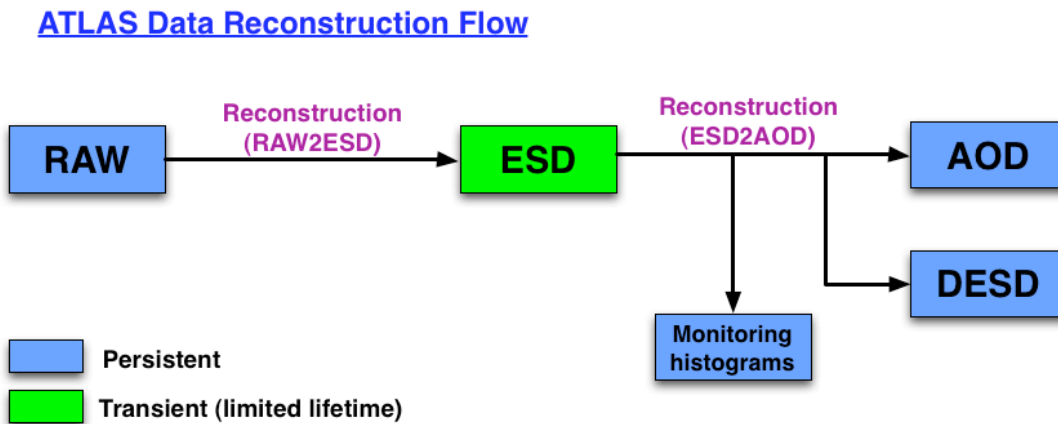


Figure 1: ATLAS Reconstruction Data Flow

The ATLAS production of the Monte Carlo (i.e. simulated) samples starts with the generation of hard-process collisions using the Monte Carlo generator programs (Sherpa, Alpgen, Powheg, AcerMC...), producing EVNT files. In some cases pre-generated inputs for the processes are produced off-Grid and registered on Grid storage for MC event generation and EVNT production. The EVNT files are then processed in the detector simulation step, producing HITS files. The modelling of pile-up is added in the next processing stage and the detector response (digitization) is simulated at the same time, producing RDO files. As a separate step, the trigger response simulation is performed again producing RDO files with the simulated trigger information added. The rest of the reconstruction chain is the same as the prompt data reconstruction. The pile-up, digitization, trigger simulation and reconstruction are usually performed together in one Grid job.

ATLAS envisages one full data re-processing from RAW and a corresponding Monte Carlo re-processing per year as well as two AOD2AOD re-processings of both real and Monte Carlo data. Fine-tuning the existing dynamic data replication to increase the availability of popular data sets and a dynamic cleaning policy for less accessed data sets will optimize the disk space usage.

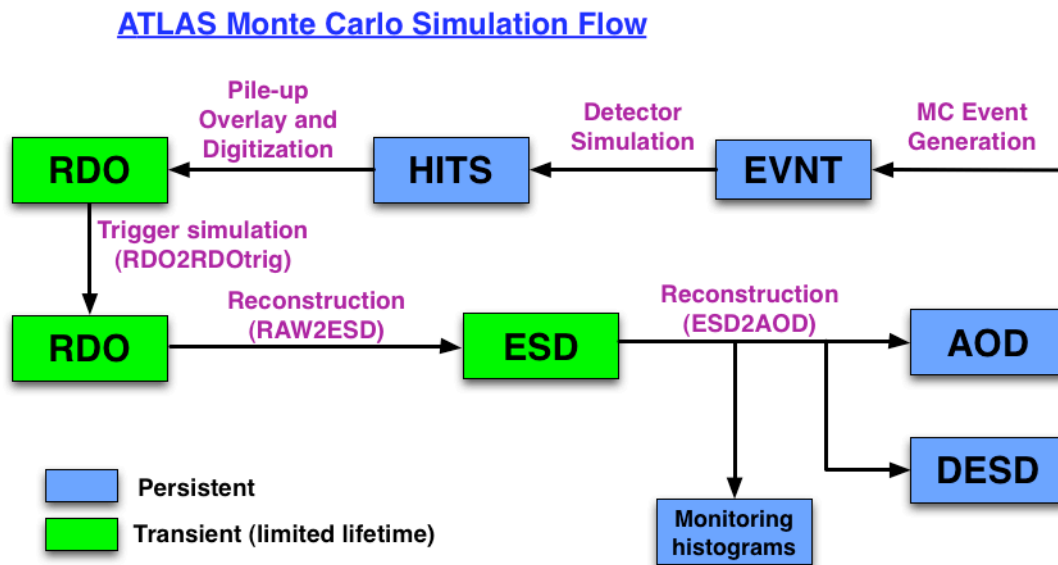


Figure 2: ATLAS Monte Carlo Simulation Data flow

4. New Analysis Environment

In Run 1, ATLAS followed a “frozen Tier-0” policy to ensure stability and reproducibility. In practice, this meant that the T0 AODs did not reflect the best available understanding; so each physics or performance group’s AOD-based analysis would first apply all the latest fixes to create a large private dataset for further analysis. In addition, since the AOD created by Athena jobs was not ROOT-readable, a majority of groups created their own large ROOT-readable datasets for group members to use. It became clear that this overall procedure was wasteful of CPU, storage and effort – hence the New Analysis Environment outlined in Figure 3.

The ATLAS Run-2 analysis starts from AODs (with an exception of specific analyses using dedicated DESDs as the input format). The data analysis can be performed within the ATLAS Athena framework or using ROOT directly to produce the final ROOT N-tuples used in individual analyses. In Run-2 the ATLAS AOD format will be updated to ensure simple ROOT readability of high-level analysis objects and thus simplify the analysis process. One can perform the analysis on the AOD files or use the reduction framework on the Grid to derive dedicated data samples in a common analysis.

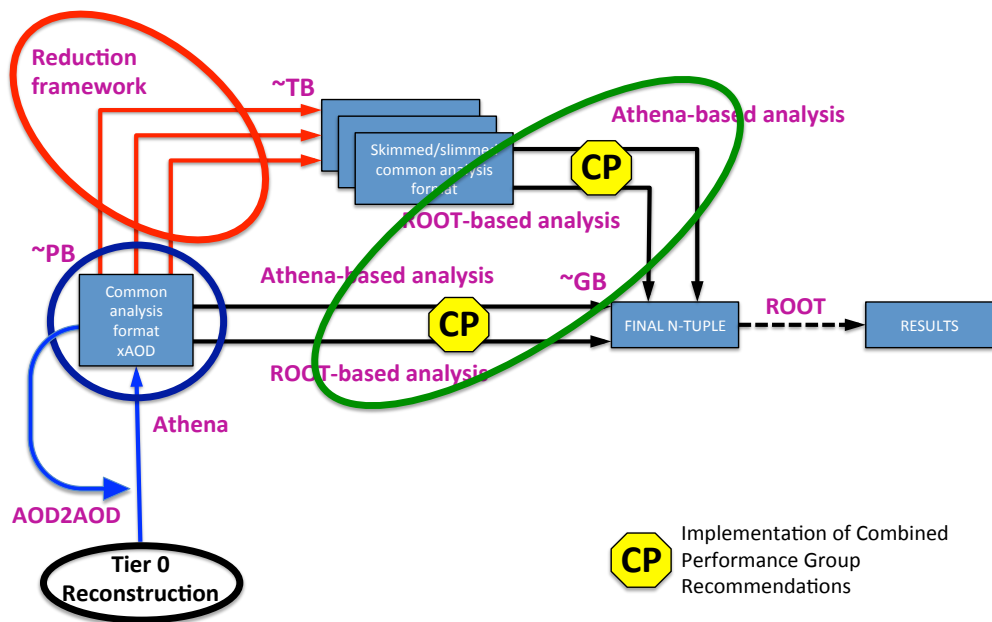


Figure 3: The new ATLAS analysis workflow showing the main areas of software activity required for its implementation

The new environment is centred on the “xAOD”. The xAOD has a completely redesigned event data model and is readable by both ROOT and Athena. The xAOD data are directly usable in ROOT without a persistent-transient translation. The xAODs produced by the T0 will be reprocessed within Athena (AOD2AOD) approximately twice each year to incorporate the latest “fixes” to reconstruction. A standardized reduction framework is being created which will run “trains” with “carriages” for each group needing a reduced dataset. The groups can put group-specific reconstruction, skimming (selecting events), slimming (selecting objects), and thinning (selecting data) into their carriages. Any more general creation of new objects will be run in the AOD2AOD processing. A standardised analysis framework is being created supporting generic tools, for example, to apply performance group recommendations to data before performing physics analysis.

In view of these improvements in Run-2 the volume of the derived group analysis formats will become more manageable and is envisaged to occupy twice the total volume of one AOD replica distributed among 50-100 derived collections (formats), all written with the same new ROOT-readable AOD technology. In combination with the AOD2AOD re-processing, the changes should also considerably reduce the CPU consumption and processing duplication. A detailed study has shown that space savings up to a factor four should be achieved by the introduction of these improvements compared to this area of the Run-1 computing model.

5. Role of the Tiers

In Run-1, the Tier-0 facility at CERN was utilized for prompt data reconstruction along with the supporting activities like the prompt data calibration, data quality determination etc. The Tier-0 was in Run-1 also used as the primary archival (tape) repository of ATLAS RAW data as well as the promptly reconstructed data in AOD and DESD format. The role of the Tier-0 facility is expected to remain unchanged in Run-2.

In Run-1 ATLAS used part of the available resources at CERN as the CAF (CERN facility for data calibration and alignment). The CAF function included providing servers, non-automated data calibration, partial reconstruction, debugging, monitoring and other expert activities as well as high priority group and user analysis activities within the Grid environment. During Run-1 ATLAS developed procedures to simply and dynamically allocate the idle resources in Tier-0 to the CAF Grid environment (e.g. during LHC technical stops) and to reassign CAF resources to the Tier-0 in case of processing congestions. This functionality is expected to remain also in Run-2. In addition, ATLAS commissioned the High Level Trigger (HLT) farm to work in the ATLAS Grid environment for simulation production, using the Open Stack IaaS interface. In Run-2, ATLAS is also planning to use the HLT farm, when idle, in the Grid environment. The HLT farm can be reconfigured within an hour. ATLAS expects to use the fraction of the HLT farm not taken by the ATLAS Trigger/DAQ activities for Monte Carlo production.

Developments in the ATLAS software and the distributed computing environment, as well as the technological progress manifest in the ATLAS Grid sites, have enabled ATLAS to perform workflows formerly restricted to Tier-1s (data re-processing, group analysis, Monte-Carlo reconstruction) also at most of the Tier-2s at the end of Run-1. This advantage is expected to be expanded upon in Run-2, enabling ATLAS to perform all the required workflows across the ATLAS Tier-1s and Tier-2s and thus to optimize the load and throughput on the sites by providing each an optimal mixture of jobs in terms of high and low throughput, memory consumption and CPU load. The Tier-1s will still have the distinguished role as the data archival centres, providing the tape archival service for the second copy of the RAW data exported from CERN as well as the archival of the ATLAS data deemed crucial for data preservation. In addition, the Tier-1s will remain the centres for the most memory and I/O intensive computing tasks.

In Run-1, the ATLAS Tier-2s were used to perform the bulk of Monte-Carlo event simulation and user analysis. Based on operational experience, the Tier-2s were also categorized into several sub-categories using the connectivity and reliability metrics automatically collected by the ATLAS distributed computing monitoring. This categorization enabled ATLAS to optimize the data placement with respect to its criticality and use-cases. This optimization is expected to remain in place and be expanded further in Run-2.

With the technological progress of wide area networks and consequently improved network connectivity ATLAS has already in Run-1 gradually moved away from the hierarchical

association of Tier-2s to one Tier-1 (MONARC model) and is now able to associate workflows between a well-connected Tier-2 and several Tier-1s, based on monitored connectivity metrics. This functionality will be further pursued in Run-2 and is being incorporated into the production system and distributed data management developments.

In 2012 and now during Long Shutdown 1, ATLAS Computing has also further explored and is investigating the use of opportunistic resources, either ‘Cloud’/IaaS resources or High Performance Computing (HPC) centres and the creation of a transparent and straightforward interface between these and the ATLAS distributed computing environment. The current development work is implementing the connection either directly through the PanDA infrastructure or, in case of HPCs, also through the ARC middleware interface to PanDA. In 2012 and 2013, ATLAS Computing has already demonstrated successful use of the Amazon EC2, Google Computing Engine ‘Cloud/IaaS’ and Helix Nebula resources for Monte-Carlo event simulation.

In Run-2 ATLAS will continue to exploit and enhance this versatility of using opportunistic resources to complement our existing resources to off-load CPU intensive and low I/O workloads which under-use our I/O optimized grid resources. A typical example would be ‘Monte Carlo generation of hard-process events’ (AlpGen, Sherpa Monte Carlo event generators). Another target use is the CPU intensive Geant4 simulation. The Grid resources can then be better used for I/O intensive work: (simulation, reconstruction, group production, analysis), important especially during peak demand.

In addition, if ATLAS Grid sites want to move to use cloud middleware (e.g. OpenStack) or provide dual-use resources (‘batch’+HPC) the ATLAS versatility gained by the ongoing R&D provides this option.

Based on the Run1 experience, ATLAS is upgrading and re-modelling both its distributed workload and data management systems.

6. Distributed Computing Environment

In 2014 the new data management system (Rucio) will replace the existing ATLAS data management system (DQ2) and introduce more granular handling of ATLAS data (file replicas rather than dataset replicas) together with a more intelligent implementation of data ownership for users and groups. This will allow optimization of the space usage while offering more functionalities in terms of local and global quotas, thus e.g. relaxing the need for dedicated group storage (SRM space tokens) at sites. The modular structure of the new system will allow better interfacing with newly developed middleware clients and services such as FTS3 allowing the utilization of different protocols than SRM for file transfers/access/deletion and storage management (e.g. HTTP for LAN and WAN access, WebDAV for storage management, xRootd for WAN and LAN access, plain gridFTP for data transfers). Additionally, ATLAS has been commissioning, and intends to use in Run-2, a xRootd based federated storage infrastructure, which will offer more resilience in data access (increasing job efficiency) and

allow transparent file access through the WAN, reducing the need of massive data replication and therefore optimizing disk space usage. An HTTP based federation model is also being evaluated for the longer term, because of the benefits in utilizing a standard protocol such as HTTP rather than an HEP-specific one.

The new workload management system will continue to rely on the pilot-based PanDA service as a core component. It will be complemented by an extension offering the functionalities of job definition and management (JEDI) and a higher modular component for workflow and task management (DEFT). Briefly, JEDI will allow dynamic job definition for better matching of site capabilities with payload requirements. It will also allow a better handling of existing use cases such as file merging, reducing the requirements for network resources and storage I/O bandwidth. DEFT will allow users and production managers to more flexibly define high level workflows and improve the automation and reliability of the system, speeding up the job processing sequence and minimizing the time needed to complete production and analysis tasks. The combination of JEDI and DEFT (prodsys2) will soon replace the existing ATLAS production system and will offer a unique platform for distributed production and analysis, reducing the cost for maintenance and support, while offering new functionalities to end users.

For the longer term, ATLAS is investigating the possibility to further refine the granularity of the production and analysis payload from a collection of events to a single event. A prototype of such an “Event Server” is under development and expected to be ready for testing by 2015. Such an innovative system would allow to enable and/or optimize the usage of (mostly) opportunistic resources where job preemption cannot be avoided, such as cloud resources, opportunistic HPC resources and even user desktops through volunteer computing. The workflows where the gain of this approach in processing efficiency is highest are high in CPU and low in I/O, such as Geant4 simulation.

The increased processing times and memory footprints of the ATLAS software due to higher energy and pile-up in Run-2 processing suggest a different approach to efficiently utilize the resources provided by the WLCG. A simple model with one job executing on one computing core is not sufficient any more due to increased memory requirements. AthenaMP provides a solution to efficiently run multi-processing jobs, using less memory, which will be exploited further. The full deployment of multi-core setup on all sites will follow the schedule of AthenaMP deployment with the commissioning of the ATLAS Run-2 software environment (i.e. the target ATLAS software release with the reconstruction, simulation and analysis tools and software) in parallel with the planned ATLAS Data Challenges in 2014.

7. Replication and deletion Policies

ATLAS, in Run-1, controlled the static replication of the pre-placed data replicas. In the updated model foreseen for Run-2, the number of pre-placed AOD replicas kept in Tier-1 and Tier-2 disks for both real and simulated data will be reduced as much as possible, placing more

reliance on the dynamic data replication. In Run-2, as in Run-1, the final user analysis files, produced by individual users, are expected not to be stored on pledged ATLAS resources. Temporary storage (scratch) space will be provided for the users to retrieve the files produced on the ATLAS Grid resources.

The number of primary (pre-placed) copies of AODs, corresponding to different (re)processings and simulation productions in Tier-1s and Tier-2s, will be decreased from two to zero according to their relevance and popularity for physics analysis. While in Run-1 this was done only via a human decision process, in Run-2 auxiliary automated mechanisms based on popularity will be introduced to provide an additional dynamic component in reducing the number of pre-placed primary AOD and DESD replicas in Tier-1 and Tier-2 disks.

In addition, as a further correction algorithm, an automated procedure to re-distribute the primary copies of very popular datasets to disk resources with enough free space, sufficient CPU resources and good network accessibility will also be put in place.

As a consequence of this model update for Run-2, ATLAS will more actively use the tape resources for data retrieval than in Run-1.

It should be stressed that during Run-1, with the introduction of the ATLAS-developed dynamic data replication mechanism (PD2P) and the automated deletion mechanisms, (Victor), ATLAS was able to reduce the number of pre-placed AOD copies considerably, especially in Tier-2s, where 20 copies of AODs were used in 2010 (10 current+10 previous), 10 copies in 2011 (8 current+2 previous) to 4 (2 current+2 previous) at the end of Run-1.

8. Description of Workflows

8.1 Prompt reconstruction

In Run-2 the new ATLAS Grid production system and the Tier-0 processing environment are being aligned to be able to use the same reconstruction job configurations. In case of congestion of CERN resources this would enable ATLAS to reconstruct fractions of data also at Tier-1s in a straightforward way if needed.

8.1.1 Data re-processing from RAW

In Run-1 the RAW data was reprocessed at Tier-1s (and a set of Tier-2s in 2012 as a proof of concept) on average once per year with an improved reconstruction software release. At the same time the same software release was introduced in the Tier-0 data processing to ensure the consistency of the whole data set. In Run-2, about one re-processing from RAW data per year is envisaged and it will be complemented by several AOD2AOD re-processings. Simulated samples are reprocessed (re-digitized and re-reconstructed) in synchronization with the data re-processing.

8.1.2 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 model has several advantages, in particular more frequent incorporation of relevant improvements in the analysis software. Also, the series of reconstruction improvements which were in Run-1 applied at the group analysis level and thus repeated several times for each group analysis format, will in Run-2 be applied only once at the central AOD2AOD re-processing level.

8.2 Monte Carlo production

The volume of simulated data required for analysis has historically been hard to predict. The excellent LHC and ATLAS performance in Run-1 resulted in the ability to address a much wider range of physics analyses, with a higher level of precision, surpassing the most optimistic expectations. In addition, detailed physics studies established that the simulation is of unprecedented quality compared to previous generations of experiments, describing the data very well in most analyses; this quality opened-up more ‘safe’ applications for the simulated data. These two facts together significantly enhanced ATLAS physics output in 2011 and 2012, and they motivated production of larger than foreseen simulation statistics. Consequently, in Run-1 the ratio of required simulated to real data was of the order of two to one, and the required volume of simulated data might be even higher in Run-2 with higher required precision for background estimation and an even larger phase space for New Physics searches.

For the very fast simulation types, the digitization and reconstruction CPU time becomes dominant. The developments during the Long Shutdown 1 are targeting the development of fast-digitization and fast-reconstruction algorithms. More compact output formats are being explored, to allow taking advantage of the high Monte Carlo simulation and reconstruction speed to produce very large Monte Carlo samples.

An unexpected development at the end of Run-1 was that a non-negligible CPU consumption was taken by the Monte-Carlo ‘event generation’ of hard processes with the introduction of large-scale use of the advanced (multi-leg, Next-to-Leading-Order..) Monte Carlo generators (Sherpa, Alpgen, Madgraph..) which in 2012 and 2013 represented 15% of the total ATLAS pledged CPU resources. ATLAS is aiming to use opportunistic resources (HLT farm, HPC centres) for these high CPU and low I/O tasks whenever possible and the aim is to significantly reduce the amount of standard grid CPU resources used for event generation though the use of these opportunistic resources, especially HPCs, in spite of the growing need for more sophisticated (and therefore slower) MC generator configurations in Run2.

9. Software improvements

Long Shutdown 1 has provided an opportunity to step back and examine ATLAS reconstruction code. This examination revealed numerous opportunities for simplification and

optimization, all of which are without compromise to physics precision. Particular examples include the new magnetic field service and numerous cleanups. One particular reconstruction improvement is the replacement of CLHEP vector and matrix manipulation by the Eigen open-source library. Eigen was selected after a study that included also CLHEP, SMATRIX and Intel MKL. Overall, it is expected that the reconstruction improvements will meet the goal of at least a factor of two speedup measured by running the new and old code on Run-1 data.

Improvements are underway to both full and fast simulation, plus the creation of a framework in which both can be combined. The Geant4-based full simulation is essential for analysis requiring the lowest systematic uncertainties for electromagnetic calorimeter measurement, and to validate fast simulation. The code is under examination by a combination of Geant4 experts and detector experts that are already delivering speedups without any compromise to physics.

ATLAS fast simulation is a combination of the FATRAS tracker simulation and the FastCaloSim calorimeter simulation. Used together, they can simulate events with two orders of magnitude less CPU time than the full simulation. FastCaloSim has been improved for Run-2. For example it now simulates hadronic tau decays in the calorimeters with a fidelity that should be useful for physics. All Run-2 simulation production will be run in the Integrated Simulation Framework (ISF) that allows any a choice of full simulation, FATRAS or FastCaloSim on a particle-by-particle basis. An example of the approach would be to use full simulation for particles in a cone around a high-energy electron, and for any muons, but fast simulation for the remaining particles. When using fast simulation, the CPU needs of “digitization” (calculating the detector response) and reconstruction dominate unless these too are made faster. ATLAS is developing fast digitization and truth-seeded reconstruction to be able to take advantage of fast simulation. These faster approaches do, of course, increase systematic uncertainties somewhat.

10. Summary

ATLAS is developing a new computing model to meet the operational challenges of Run-2. On-going efforts are taking place in both the areas of distributed computing and software. The new model and its various components will be tested at full scale during the Data Challenge that will take place at the end of 2014.

References

- [1] The ATLAS Computing TDR CERN/LHCC/2005-022