PROCEEDINGS OF SCIENCE

Providing the computing and data to the physicists: Overview of the ATLAS distributed computing system

Michal Svatoš[a,*] **on behalf of the ATLAS Collaboration**

[a] *Institute of Physics, v.v.i., Academy of Sciences of the Czech Republic, Na Slovance 1999/2, Prague, Czech Republic*

E-mail: Michal.Svatos@cern.ch

The ATLAS experiment at CERN uses more than 150 sites in the WLCG to process and analyze data recorded by the LHC. The grid workflow system PanDA routinely utilizes more than 400 thousand CPU cores of those sites. The data management system Rucio manages about half an exabyte of detector and simulation data distributed among these sites. With the ever-improving performance of the LHC, more data is expected to come and the ATLAS computing needs to evolve and adapt to that. Disk space will become more scarce which should be alleviated by more active usage of tapes and caches and new smaller data formats. Grid jobs can run not just on the WLCG sites but also on opportunistic resources, i.e. clouds and HPCs. A new grafana-based monitoring system facilitates operation of the ATLAS computing. This presentation will review and explain the improvements put in place for the upcoming Run 3 and will provide an outlook to the many improvements needed for the HL-LHC.

40th International Conference on High Energy physics - ICHEP2020
July 28 - August 6, 2020
Prague, Czech Republic (virtual meeting)

*Speaker

PoS(ICHEP2020)926

1. Introduction

The ATLAS detector [1] is a multi-purpose particle physics detector with a forward-backward symmetric cylindrical geometry and nearly 4π coverage in solid angle. It is located at the Large Hadron Collider (LHC) [2] at CERN near Geneva.

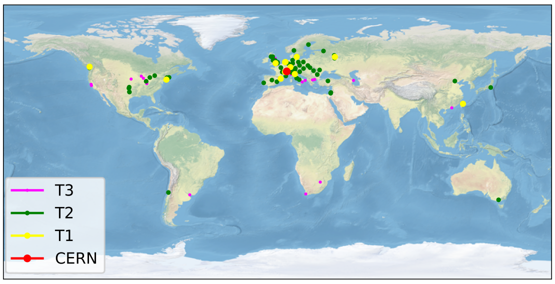

Figure 1: Locations of computing centers which provide resources to ATLAS. T0 (i.e. CERN) is the largest computing resource which also has detector data archived on tape. T1s are the largest computing centres with second copy of the detector data on tape. T2s are computing centres (usually smaller than T1 sites). T3s are small sites (sometimes, only for local users).

The collaboration comprises about 3000 scientific authors from 183 institutions, representing 38 countries. The collaboration of such size also needs significant amount of computing resources. These are managed by the ATLAS Distributed Computing (ADC) system.

The ADC manages about half an exabyte of detector and simulation data and more than 400 thousand CPU cores (Figure 2) at more than 150 sites located around the world (Figure 1). The system is running non-stop.

2. Current status

The ADC system uses various CPU resources to run jobs (Figure 2). Namely Worldwide LHC Computing Grid (WLCG) sites, cloud resources (including opportunistic usage of ~90k cores of High Level Trigger farm), HPCs, Volunteer computing (BOINC), etc.

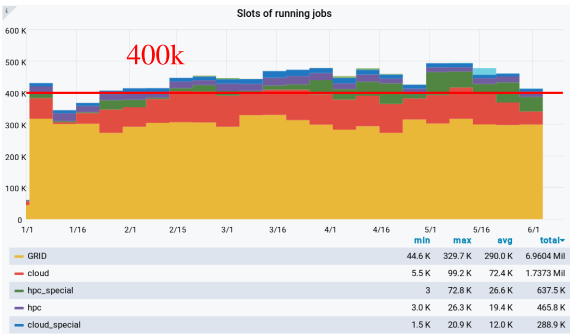

Figure 2: CPU resources available to the ATLAS experiment. With all the pledged and opportunistic resources, the ADC system can run on more than 400k cores sustained.

The approximately half exabyte of data managed by the ADC system is split almost evenly between disk and tape (Figure 3). Because the disk storage is always almost full, various mechanisms were developed to clean it. For example, campaigns to remove older versions of analysis-format data or data which were not used for certain period of time, are run several times each year. Even with these mechanisms in place, amount of space which can be used for additional data caching purposes (yellow band on the left plot of Figure 3) is low.

3. Run3 / HL-LHC outlook

With the expected increase in luminosity at the LHC and the corresponding increase in event size and rate, the data volume will increase by an order of magnitude. This outlook presents exciting

PoS(ICHEP2020)926

Figure 3: Storage used by the ATLAS experiment. On the left, there is used disk space (green color representing primary data which need to stay on disk and yellow color representing secondary data which can be deleted if necessary). On the right, there is used tape space.

opportunities for physicists but serious challenges for computing. Figure 4 shows the projection of the expected CPU and disk resources needed to manage, process, and analyze data of such size. Run-3 should be manageable with the current system but aggressive R&D will be necessary to manage the data-rates expected from the HL-LHC in Run-4 from 2027.

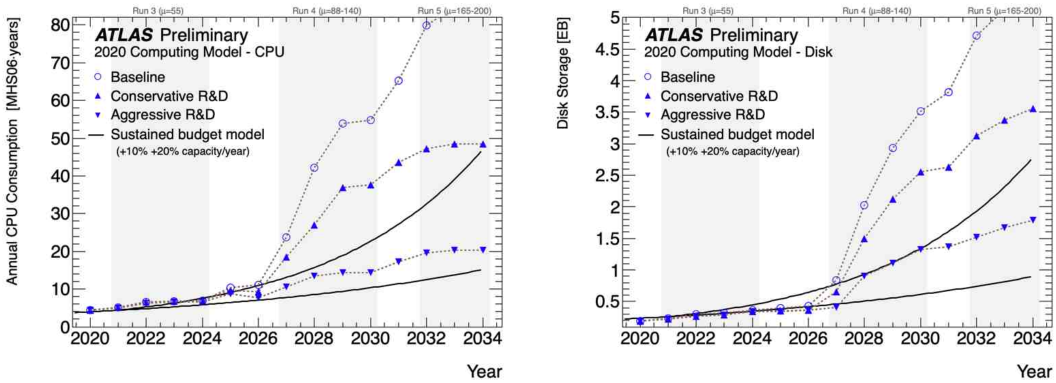

Figure 4: Projected CPU (left) and disk (right) requirements of the ATLAS experiment between 2020 and 2034 based on 2020 computing design report assessments [3]. Both plots show that, while the ATLAS computing system will be able to sustain the Run3 data rates, aggressive R&D is necessary to fit resource constraints of HL-LHC era.

In addition, the computing budget is assumed to stay flat and performance gains from technology advancements are decreasing. One example would be price/performance evolution of installed CPU servers at CERN [4] (Figure 5). In the last decade and half, there was a trend of decreasing price/performance, i.e. CERN was able to get more performance for the same amount of money. In the last few years, the price/performance variable stagnates or even increases, i.e. CERN needed to pay more money for the same performance. The outlook for the next few years is only a slow decrease.

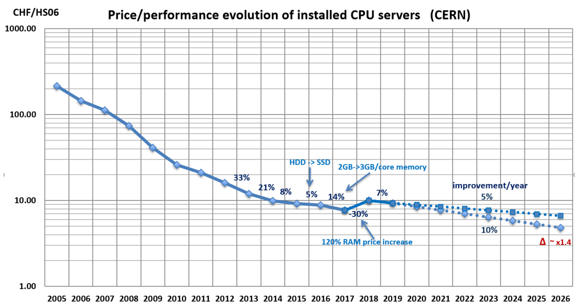

Figure 5: Price/performance evolution of installed CPU servers at CERN [4]. In the past, the trend of price/performance variable was decreasing. In the recent years, the trend is stagnating or even increasing.

3.1 Software

Various improvements are needed in software which is used to process recorded data (or generated events) and store them in data formats suitable for analysis.

3.1.1 Event Generation

The Monte Carlo generated events are expected to be needed at next-to-leading (NLO) and next-to-next-to-leading (NNLO) level of precision. As opposed to other processing steps, event generators are not a product of the collaboration, i.e. there is a limited influence on development and optimizations.

There are several options which can help to save resources. First, optimization of physics choices in event generators can bring significant savings. For example, in Sherpa NLO-merged V+jets events, the usage of a different clustering scale allowed for a speed-up by about a factor of two with no visible impact in modelling. Second, biasing the event generation (as a function of a kinematic quantity of interest) allows significant reduction in the number of events that need to be generated. Otherwise, large resources are spent on populating extreme regions of phase space with inclusive samples. Also, computation of uncertainties from scales and PDFs through a re-weighting technique removes the need to produce large numbers of Monte Carlo samples to evaluate systematic variations.

The LHC experiments can also share their generated events (mainly relevant for ATLAS and CMS), as those are experiment-independent.

3.1.2 Simulation

Simulation means modelling the interaction of particles with the detector. There are many R&D projects in ATLAS and GEANT4 dedicated to resource usage reduction. The resource usage can be currently decreased by usage of so called fast simulation (primarily parametrized calorimeter response) instead of FullSim (based on GEANT4).

Currently, the full detector simulation is used in the majority of simulations but there are plans to increasingly replace this with the parametrized fast simulation.

3.1.3 Digitization

The digitization in MC simulation models means modelling of the output of the detector readout. For this step, it is planned to use pre-mixed pile-up datasets, i.e. hard-scatter events will be digitized and then "overlaid" on top of a pre-mixed pile-up event. This approach has several advantages. It is considerably faster, has reduced disk I/O requirements, and scales much less steeply with pile-up luminosity. On the other hand, the pre-mixed pile-up events will occupy some additional disk space.

3.1.4 Reconstruction

Reconstruction means creation of high-level physics objects. This step produces AODs (Analysis Object Data). ATLAS initiated the ACTS (A Common Tracking Software) [5] open source project to develop the next generation tracking software in a common cross-experiment project.

3.2 Analysis model

The current analysis model starts with the centralized data reduction system which uses the output of the reconstruction (AODs). It produces data in the DAOD (i.e. Derived AODs) format. The AODs and DAODs together occupy three quarters of disk space (Figure 6).

Process of transformation of AODs into DAODs can include slimming (removing unnecessary variables), thinning (removing unnecessary objects), skimming (removing unwanted whole events), or adding new variables or objects. Analysis teams can define formats tailored for their specific analysis. The problem is that there is a significant overlap in the output formats produced by the various analysis groups (currently there are 84 different formats) which is causing a heavy disk footprint.

A new analysis model is being prepared which will address many of these issues. One of the main ingredients will be the new common un-skimmed data formats - DAOD_PHYS (with about 50kB/event) and DAOD_PHYSLITE (with about 10 kB/event). The goal is to cover the needs of up to 80% of ATLAS analyses. It should also have several other advantages for analysers. As those new formats will be much smaller, ATLAS can keep more copies (i.e. availability of data for analysers will improve). The event data model will be much flatter to allow for better integration with the growing Python-based analysis ecosystem. More disk space saving could be achieved by appropriate application of lossy compression.

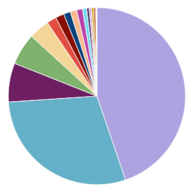

	current	percentage
DAOD	101.6 PB	45%
AOD	66.3 PB	29%
HITS	16.1 PB	7%
user	13.09 PB	6%
EVNT	8.58 PB	4%
log	3.97 PB	2%
RAW	3.55 PB	2%
NTUP	2.822 PB	1%
DRAW	2.676 PB	1%
group	2.435 PB	1%

Figure 6: Sizes of data formats located on disk. Three quarters of disk space are used by DAODs (analysis format) or AODs (output of reconstruction used to make DAODs).

3.3 Improvements in compute

Jobs in ADC can be split in two categories - production and analysis. Production jobs process a chain starting at raw data from the experiment (or product of event generation) and ending at the common formats used for analysis. They are submitted by a dedicated team of experts. Analysis jobs are jobs from individual analysers or analysis groups. The analysis takes only a fraction of CPU resources but dominates in number of files it reads.

3.3.1 Grand-Unification

Production and analysis jobs used to be sent to separated dedicated queues at each site. The separation existed both on the side of ADC as well as in the batch systems of sites. The process of grand-unification, i.e. unification of access to resources for production and analysis, should remove the separation. With grand-unification, the number of analysis jobs and their location is based on global shares defined by the ADC and not, for example, some hard-coded split in a batch system. This allows analysis jobs to be sent where the data are (instead of moving the data around) and therefore speed up the data processing by analysers.

PoS(ICHEP2020)926

3.3.2 Containerization

Almost all jobs run inside a generic (singularity) container. It contains a bare minimum of software and dependencies to run ATLAS job. But the container can also be used to run jobs in user specific containers (if specific software is needed) and for data preservation.

3.3.3 Interactive analysis

The usual way of developing analysis code is by interactively running on small sample of data. In recent years, the approach using Jupyter Notebooks connected to horizontally scalable compute clusters such Spark, Dask or Ray or batch systems, appeared. This way, analysers can develop the analysis code using more performant resources than a personal computer and therefore use more data or finish the processing sooner.

3.4 Data Carousel

The Data Carousel is a sliding window approach to orchestrate data processing with the majority of data resident on tape storage. The processing is executed by staging (i.e. copying from tape to disk) only part of the data which can be immediately processed by available CPUs. This means only the minimum required input data are located on disk at any time. The Data Carousel was successfully tested on full Run2 RAW data reprocessing (18 PB staged over several weeks rather than all at once).

4. Summary and conclusion

The distributed computing of the ATLAS experiment at LHC is preparing for the challenges of "HL-LHC" era. Its computing activities need to evolve as the current model will not meet the HL-LHC requirements. Many improvements are necessary in the ADC to be able to make LHC data available to ATLAS physicists and to provide them with means to analyse the data. Several examples like the new analysis model introducing new small data formats, internal and external software improvements, improvements in compute access for analysers, improvements in disk/tape usage were mentioned.

References

[1] ATLAS Collaboration, JINST **3**, S08003 (2008).

[2] L. Evans and P. Bryant, JINST **3**, S08001 (2008)

[3] ATLAS Experiment Computing and Software - public results, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ComputingandSoftwarePublicResults, [accessed 2020-05-19]

[4] 6th Scientific Computing Forum, https://indico.cern.ch/event/759361/contributions/3327688/attachments/1801413/2938427/cpu_disk_server_price_evolution_CERN_Feb2019_v3.pdf, [accessed 2020-06-10]

[5] Acts Common Tracking Software, http://acts.web.cern.ch/ACTS/