

The ATLAS Distributed Computing project for LHC Run-2 and beyond

Alessandro Di Girolamo¹ on behalf of the ATLAS Collaboration

CERN

Geneva, Switzerland

E-mail: Alessandro.Di.Girolamo@cern.ch

The ATLAS Distributed Computing infrastructure has evolved after the first period of LHC data taking in order to cope with the challenges of the upcoming LHC Run 2. An increased data rate and computing demands of the Monte-Carlo simulation, as well as new approaches to ATLAS analysis, dictated a more dynamic workload management system (ProdSys2) and data management system (Rucio), overcoming the boundaries imposed by the design of the old computing model.

In particular, the commissioning of new central computing system components was the core part of the migration toward the new, more flexible computing model. The flexible computing utilization exploring the opportunistic resources such as HPC, cloud, and volunteer computing is embedded in the new computing model, the data access mechanisms have been enhanced with the remote access, and the network topology and performance is deeply integrated into the core of the system. Moreover, a new data management strategy, based on defined lifetime for each dataset, has been defined to better manage the lifecycle of the data. In this note the overview of the operational experience of the new system and its evolution is presented.

The European Physical Society Conference on High Energy Physics

22–29 July 2015

Vienna, Austria

1 Speaker

1. The LHC Run 2 ATLAS computing challenges

The Large Hadron Collider (LHC) has been upgraded during 2013 and 2014 (the so-called Long Shutdown 1, LS1) to increase the energy of the colliding beam up to 6.5 TeV per proton beam and increase the luminosity at the collision points. These improvements brought new computing challenges for the ATLAS experiment: the trigger rate, for Run 2 (2015 onwards) foreseen to be at an average of 1 kHz, has more than doubled with respect to the one during Run 1 (400Hz); the average pile-up (the amount of proton-proton collisions per each bunch crossing), for the beginning of Run 2 at an average of 40, almost double with respect to the one during Run 1, increases the complexity of the events and their processing time of more than a factor three.

In addition, the improvements of the ATLAS sub-detectors performed during LS1, which affect mostly the tracking and muon systems, need to be fully understood also from the software and computing point of view.

ATLAS has to address these challenges, taking into account that it cannot rely on an increase of the budget dedicated to the computing resources, as the agreements with the funding agencies are that ATLAS computing will be on a “flat-budget” regime for Run 2: a growth in the computing power of approximately 20% per year is expected following Moore’s law, but this increase will only partially satisfy the new needs. Several measures were taken in the software and computing areas to tackle these challenges.

2. The ATLAS Distributed Computing systems for LHC Run 2

The ATLAS Distributed Computing (ADC) [1] systems have substantially evolved in the past few years: a completely new Distributed Data Management system (DDM), called Rucio [2], a new Production System, called ProdSys2 [3], and a new Job Management system, called JEDI, have been developed and commissioned. Also, completely new frameworks, namely the EventIndex and the Event Service, were developed and are being integrated in the ADC in these first months of LHC Run 2 data taking.

2.1 The new ATLAS Distributed Data Management System: Rucio

Rucio, the Run 2 ATLAS DDM system, has replaced the one used during Run 1 called DQ2. Rucio has been intensely tested during LS1 and as been put in production at the end of 2014. As of July 2015 Rucio is managing approximately 700 million files, which occupy approximately 170 PetaBytes of space. The files are organized in datasets, of which there are several millions. Rucio catalogs both the files and the datasets: this is a substantial difference with respect to Run 1 where the file catalog was done through the LFC (LHC computing grid File Catalog) and the dataset catalog was a part of DQ2,

with this split relying on intense operational efforts to avoid any de-synchronization between the two catalogs.

One other important characteristic of Rucio is the possibility to integrate different protocols for data transfers. As of today, all the computing centers that contribute to ATLAS computing resources offer at least the GridFTP transfer protocol, often behind an SRM (Storage Resource Manager) service. Through Rucio it is now possible to use, for specific activities as upload/download between worker nodes and the storage or for log file inspections, also sites which expose xrootd and/or HTTP protocols, and are more efficient than SRM for such operations.

In terms of performance, Rucio has demonstrated during these first months of data taking to be able to transfer files at a rate up to 35 Hz with an overall throughput of about 10GB/s (Figure 1), without showing any architectural limitations.

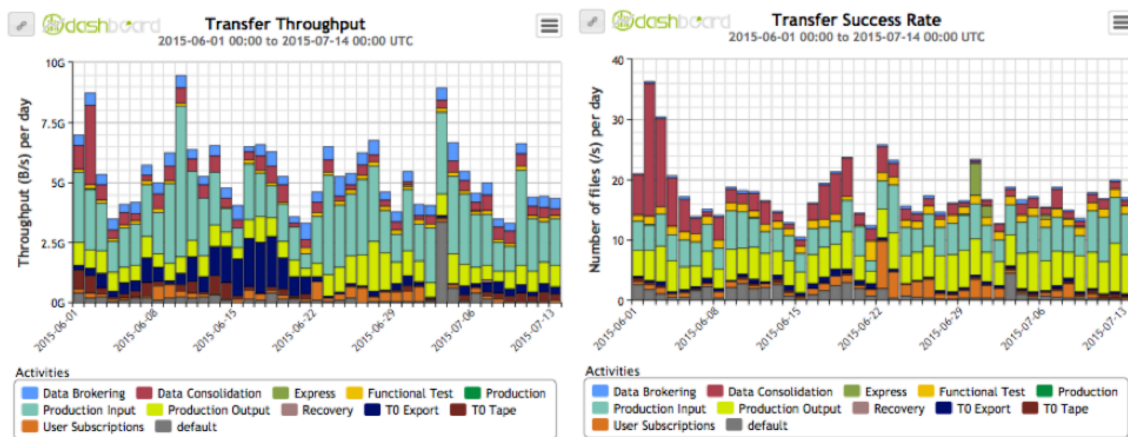


Figure 1, taken from [4]- On the left: daily volume of data transfer; on the right: file transfer frequency

2.2 The workload management system: JEDI and ProdSys2

The core of the ATLAS workload management system is the PanDA (Production AND Distributed Analysis) framework, and during the LS1 efforts have been put to automate the workload management system. For the production activities a new request interface (DEFt: Database Engine for Task) has been developed to manage the requests and translate them into chains of tasks. Furthermore, a new system, called JEDI (Job Execution and Definition Interface) has been developed to dynamically create the jobs and automatically manage the failures and the retries on the distributed infrastructure. Summing production and analysis activities, the workload management system is managing jobs that use over 200k CPU slots, as shown in Figure 2.

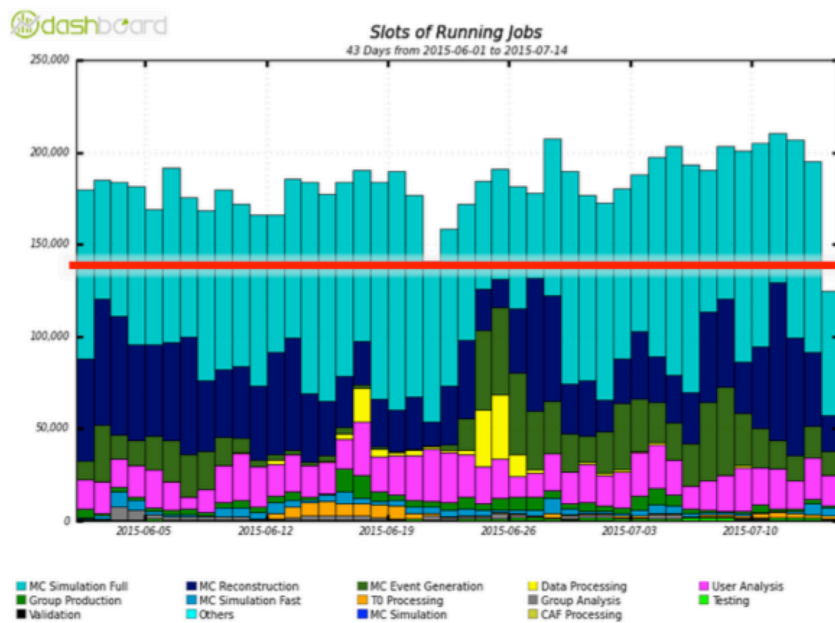


Figure 2, taken from [4]- ATLAS concurrent running job slots. The red line represents the pledged resources.

2.3 EventIndex and Event Service

Completely new frameworks, introducing new functionalities, have been developed during LS1. The EventIndex [5] is designed to be the global catalog of the ATLAS events: it contains one record for each event for each processing stage. For each of its records the EventIndex is storing the event identification information like run number, luminosity, data stream and type, the trigger information and the pointers to the files containing the event is and its internal location. This new system is useful in several cases, such as sample selections on rare trigger combinations, allowing retrieving specific events, it can also help in minimizing the overlap of the events on the various outputs of the derivation frameworks. The EventIndex has been filled with all the Run 1 data and it is now running in parallel to the normal data taking and production activities, being filled as the data are produced.

The Event Service [4] has been developed to optimize the usage of short-living resources, i.e. resources available at highly discounted prices, if agreed that they can be taken back by the provider with notice as short as a few minutes, or High Performance Computing resources, which can be used in backfill mode. The Event Service streams single input events into the worker nodes and collects the outputs as soon as produced by each event, minimizing the time lost in case of a sudden vanishing of the resources. The commissioning and the physics validation of the Event Service were ongoing at the end of summer 2015.

3. Leveraging new type of resources

Intense research, development and integration work has been done to be able to exploit several types of computing resources that have become more available over the past years, like High Performance Computing (HPC) machines, Cloud resources [7] and Volunteering computing through BOINC [8]. Figure 3 shows the usage of the MIRA supercomputer at Argonne (USA) over the past year: 62 Million hours were used to produce computationally intense events for ATLAS. For what concerns HPCs, there are still several challenges to address, because as of today each HPC is different and requires dedicated efforts to fully integrate it within ProdSys2. In addition, porting the ATLAS software to scale up to million of threads to exploit the full power of the HPC is an activity that will require great efforts in software development.

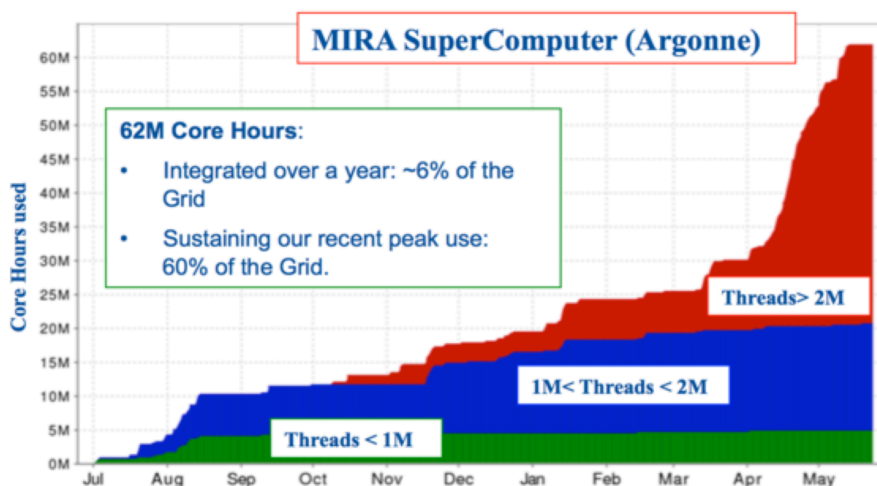


Figure 3 - Usage of the MIRA supercomputer in Argonne between July 2014 and May 2015

4. Resource Usage optimizations

To optimize the usage of the Distributed Computing resources, ATLAS made several changes, notably to the software and also in the analysis model.

The ATLAS reconstructions jobs [9] are particularly memory hungry, needing more than 3GB of physical RAM per core: to keep the memory below the 2GB physical RAM per core, which is the typical configuration of the resources in the Grid, ATLAS commissioned and deployed AthenaMP, a framework that allows to run the jobs in multi-process mode: as it is shown in the right plot of Figure 4, with an 8 core AthenaMP job the physical memory used is well below the 16 GB limit.

The reconstruction software has also been heavily improved during LS1, bringing down the reconstruction time per event by a factor 3 (from approximately 60 to 20 seconds), as it is visible in the right plot of Figure 4.

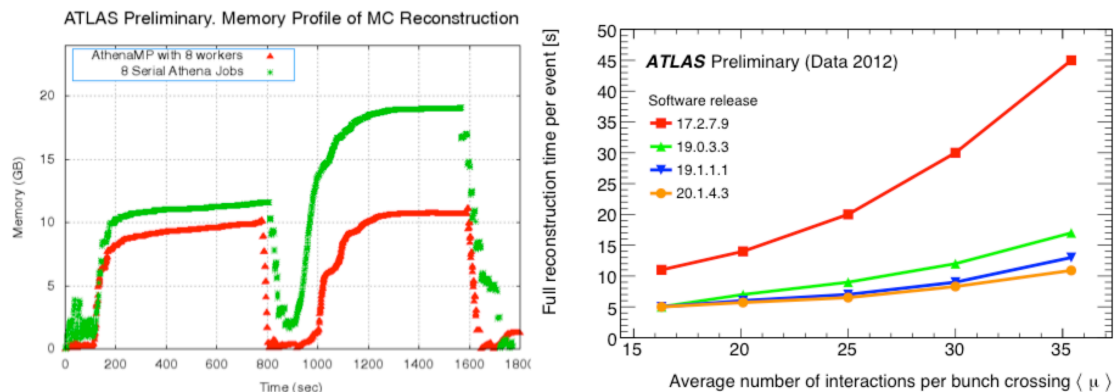


Figure 4 - On the left: memory usage of 8 single-core jobs (green) and one 8-core job (red). On the right: reconstruction time per event as a function of the pile-up for the various ATLAS software releases: in red a Run 1 release, the others are improvements done during the LSI.

Another important change introduced in Run 2 is the new ATLAS analysis model [8], which is based on xAOD (Analysis Object Data) that are readable through the ATLAS software and also directly by ROOT: during Run 1 various physics groups were using a further step to transform the AOD into ROOT-readable ntuples, so this change saved precious processing time. Coupled with the new analysis model, there is the implementation of the Derivation Framework, which enables the central processing through ProdSys2 of the petabytes of xAOD produced: the derived outputs are for each group about 1% of the original input size, thus allowing for the users more usability and quick access to the data. Thanks to this new model we are able to steadily provide data to physicists within six days from data taking and we have also demonstrated to be able to finish the full workflow within three days in case of urgent deadlines.

5. The new Data Lifetime Model

The ATLAS Distributed Computing infrastructure consists of about 200 PB of storage, shared between disk and tape. During Run 1 the data distribution was implemented mainly by two complementary methods: a pre-placement of the data newly created with a certain amount of minimal copies, 2 copies for most of the data, and a dynamic placement based on the popularity of the data, which triggered the replication of new data in case of heavily accessed datasets, and the reduction of the number of replicas down to the minimal 2 once these dynamically created replicas were not accessed anymore. This method worked well during Run 1 but showed its limits: while the pre-placed number of copies were meant to be not more than 60-70% of the total available space to leave enough space to the dynamic placement, at the end of Run 1 most of the space was used by these pre-placed copies, leaving not so much space for the dynamic placement; in addition, we observed that there was a non-negligible amount of produced data which were never or very rarely accessed.

For Run 2, to optimize the usage of the storage resources, ATLAS defined a Data Lifetime policy, learning from the experience acquired during Run 1. This data lifetime policy consists in applying to each dataset a lifetime at creation time, and extending this

lifetime in case the data are accessed. Data not accessed and with an expired lifetime can automatically be deleted from disk or tape in case the storage space is needed to store other data. For each type of data a specific lifetime is defined, based on the difficulties in reproducing them and on the importance of the data themselves: for example, RAW data have infinite lifetime. Within the boundaries of the lifetime of the data it is up to ATLAS Distributed Computing to manage the accessibility and the resiliency of the replicas: most of the data are stored directly on tape for resiliency, from where they are deleted only in case of expired lifetime and exceeded tape quota used, and the replicas on disk are dynamically increased or decreased depending on data popularity.

One of the implications of this model is the increase of usage of the tape systems: the amount of space available on tape has not increased significantly, thus data will be deleted also from tape once their lifetime will expire. This will induce more load on Tier1s to repack the tapes from which data are deleted to be able to reuse the space. In addition tape systems are now used to keep replicas not heavily used, and thus not anymore on disk, but that can be required by specific analyses, and this can increase the number of file recalls with respect to Run 1. The observation over the first six months of Run 2 is that the increase of the tape recall is of the order of 10%.

Since the end of 2014 the data lifetime model is enforced in a semi-automatic way each few months, leading to a substantial reduction of unused data, which dropped from 10PB before the first implementation of the model to a number that now is approximately 5PB. More than 15 PB of data were also deleted from the tape systems. Figure 5 shows the evolution of the used storage on disk and on tape.

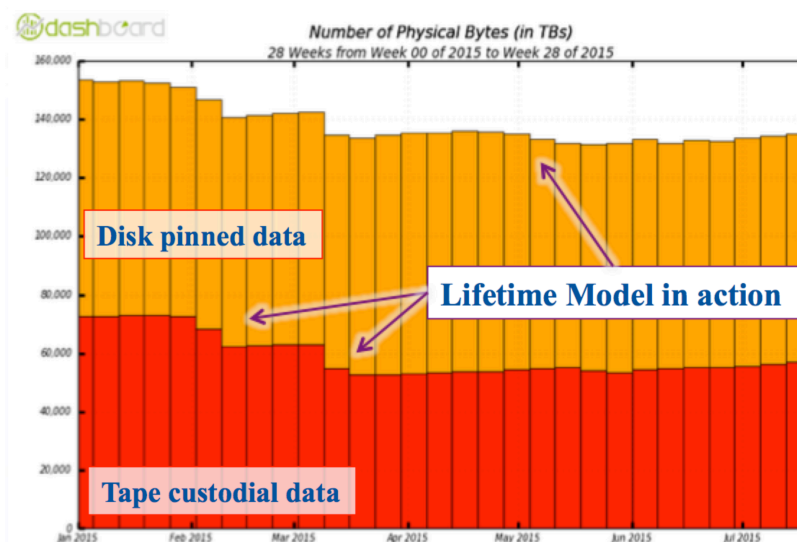


Figure 5, taken from [4]– evolution of disk and tape used storage space with the Data Lifetime model

6. Conclusions

The ATLAS Distributed Computing tools and procedures have gone through an intense upgrade program during LS1. All systems have been fully commissioned and are now working efficiently, providing ATLAS with a solid and reliable infrastructure to accomplish its rich physics program. Major evolutions are now in place with respect to Run 1: Rucio, the new Distributed Data Management system, ProdSys2 and JEDI, new services such as the EventIndex and the Event Service. Thanks to these improved services and to the improvements done in the software and in the analysis model, the ATLAS computing infrastructure is able to fully exploit all the pledged and opportunistic resources available, providing quickly results to physicists.

References

- [1] S. Campana et al. “ATLAS Distributed Computing in LHC Run 2”, *2015 Int. Conf. on Computing in High Energy and Nuclear Physics 2015, Okinawa*
- [2] V. Garonne et al. “The ATLAS Data Management system - Rucio: commissioning, migration and operational experiences” *2015 Int. Conf. on Computing in High Energy and Nuclear Physics 2015, Okinawa*
- [3] A. Klimentov et al. “Scaling up ATLAS production system for the LHC Run 2 and beyond: project ProdSys2”, *2015 Int. Conf. on Computing in High Energy and Nuclear Physics 2015, Okinawa*
- [4] J. Andreeva et al., “Experiment Dashboard - a generic, scalable solution for monitoring of the LHC computing activities, distributed sites and services”, *2012, J. Phys.: Conf. Ser.* **396** 032093 [doi:10.1088/1742-6596/396/3/032093](https://doi.org/10.1088/1742-6596/396/3/032093)
- [5] T. Wenaus et al. “The ATLAS Event Service: A new approach to event processing”, *2015 Int. Conf. on Computing in High Energy and Nuclear Physics 2015, Okinawa*
- [6] D.Barberis et al. “The ATLAS EventIndex: architecture, design choices, deployment and first operation experience” *2015 Int. Conf. on Computing in High Energy and Nuclear Physics 2015, Okinawa*
- [7] R. Taylor et al. “Evolution of Cloud Computing in ATLAS”, *2015 Int. Conf. on Computing in High Energy and Nuclear Physics 2015, Okinawa*
- [8] D. Cameron et al. “ATLAS@Home: Harnessing Volunteer Computing for HEP”, *2015 Int. Conf. on Computing in High Energy and Nuclear Physics 2015, Okinawa*
- [9] J. Mitrevski et al. “Preparing ATLAS Reconstruction for LHC Run 2”, *2015 Int. Conf. on Computing in High Energy and Nuclear Physics 2015, Okinawa*
- [10] J. Catmore et al. “New Petabyte-scale Data Derivation Framework for ATLAS”, *2015 Int. Conf. on Computing in High Energy and Nuclear Physics 2015, Okinawa*