

The ATLAS Data Carousel Project Status

Mikhail Borodin², Alessandro Di Girolamo¹, Edward Karavakis¹, Alexei Klimentov^{ID}^{4,*}, Tatiana Korchuganova^{ID}^{3,5,6}, Mario Lassnig^{ID}¹, Tadashi Maeno⁴, Siarhei Padolski⁴, and Xin Zhao⁴ on behalf of ATLAS Software and Computing

¹European Organization for Nuclear Research (CERN), Geneva, Switzerland

²The University of Iowa, USA

³Institute of System Programming, Russian Academy of Science, Moscow, Russia

⁴Brookhaven National Laboratory, Upton, NY, USA

⁵Plekhanov Economy University, Moscow, Russia

⁶University Andres Bello, Santiago, Chile

Abstract. The High Luminosity upgrade to the LHC, which aims for a ten-fold increase in the luminosity of proton-proton collisions at an energy of 14 TeV, is expected to start operation in 2028/29 and will deliver an unprecedented volume of scientific data at the multi-exabyte scale. This amount of data has to be stored, and the corresponding storage system must ensure fast and reliable data delivery for processing by scientific groups distributed all over the world. The present LHC computing and data management model will not be able to provide the required infrastructure growth, even taking into account the expected hardware technology evolution. To address this challenge, the Data Carousel R&D project was launched by the ATLAS experiment in the fall of 2018. State-of-the-art data and workflow management technologies are under active development, and their current status is presented here.

1 Introduction

The overarching common challenge for particle physics experiments is data handling. The evolution of the computing facilities and the way storage will be organized and consolidated, will play a key role in how the possible shortage of resources will be addressed by the LHC [1] experiments. Technologies that will address the High Luminosity LHC (HL-LHC) computing challenges may be applicable to other scientific communities, such as SKAO, DUNE, Vera Rubin Observatory, BELLE II, and JUNO for the management of large-scale data volumes. To address the HL-LHC distributed data handling challenge ATLAS [2] has launched the *Data Carousel* R&D project to study the feasibility of getting input data from tape directly for various ATLAS workflows.

The Data Carousel is the orchestration between the workflow management systems ProdSys2 and PanDA [3][4], the distributed data management (DDM) system Rucio [5], and the tape services. It enables a bulk production campaign, with input data resident on tape, to be executed by staging and promptly processing a sliding window of a fraction of the input onto buffer disk such that only a percentage of the data are pinned on disk at any one time. From the very beginning we defined three phases of the project:

*e-mail: aak@bnl.gov

- Phase I: Tape system performance evaluation at CERN and the WLCG [6] Tier-1 centers.
- Phase II: Workflow management and data management systems integration, and integration of the ATLAS specific distributed software components with middleware, such as the File Transfer Service (FTS) [7], and with Grid sites. During Phase II we identified missing distributed software components needed for an effective Data Carousel operation (it was mostly related to Rucio/ProdSys2 integration and a special service to release data processing before the whole data sample is staged in).
- Phase III: Run Data Carousel at scale in production for the selected workflows with an ultimate goal to have it operational before LHC Run 3 in 2022.

Phase I and Phase II results have been presented at the CHEP2019 conference [8]. In this manuscript we describe the recent software developments, Phase III accomplishments, and our future plans. All three phases have now been completed. ATLAS Distributed Computing ran the full LHC Run 2 data reprocessing and has been continuously running Monte-Carlo simulation in Data Carousel mode since 2020. Derivation production was demonstrated at small scale, and it will be in full production in 2021.

2 ATLAS Run 2 data reprocessing in Data Carousel mode

One of the Phase III goals was to run Data Carousel in production for bulk data processing. It was decided to perform a very challenging task to demonstrate the efficiency and advantages of Data Carousel for the reprocessing of all data collected by ATLAS in 2015-2018. The total data volume was close to 18.5 PB. The data were processed in reverse order, with the largest data volume from year 2018 processed first and the smallest data volume from year 2015 processed at the end.

Several fundamental changes were implemented before bulk production was started. ATLAS developers agreed with the system administrators of the Tier-1 sites and the dCache and CTA software teams on the joint monitoring of storage system performance to identify potential bottlenecks at an early stage. Each Tier-1 site and the CERN Tier-0 site were asked to provide their preferred data staging profile. This is a formal description of the data staging timeline, as burst data staging was used during Phase II. Profiles are stored in the ATLAS information system (CRIC), and they are used by the Production System. Production System submits staging requests to Rucio following site's defined staging profile policy. The Production System doesn't send new requests to Rucio to stage a new data chunk until the processing of the previous one has reached a predefined level, usually 50+%. In addition, the ATLAS Physics Coordination defined priorities and shares for data reprocessing. All these factors defined the size of the bulk requests.




Figure 1: Total Tier-1 staging throughput (GB/s) during the reprocessing campaign.




Figure 2: Staging throughput (GB/s) from a single Tier-1 (RAL) during the reprocessing campaign, which clearly shows the wave-like staging pattern.

The figures 1 and 2 demonstrate data staging throughput, and Figure 3 shows the RAW data volume on disk at any given time. Data was marked for removal from disks immediately after the reconstruction process was completed. One can see that the rule not to exceed 3 PB (day average for primary data) was respected. Table 1 shows the tape throughput recorded for each Tier-0 and Tier-1 site during the campaign. The peak data staging performance was approximately 16 GB/s, which is a great improvement over the throughput we reached in 2018, as shown in the table. Figure 2 is the staging throughput from a single Tier-1 site, which shows clearly a wave-like pattern with artificial delays in between bulk staging requests. This is an efficient use of the tape resources that the results from the Production System following site staging profiles. The Data Carousel monitoring was improved to address site and operational requirements, in particular data staging and task execution/delays are promptly monitored in the new version. In order to promptly process staged data and to improve turnaround time, a new software module, the Intelligent Data Delivery Service (iDDS) has been developed and integrated with the existing system. It is described in more details at the next section. iDDS leverages the existing data orchestration. The collaboration between the Data Carousel and iDDS R&D projects is an excellent example of early HL-LHC R&D delivery and commissioning for LHC Run 3.




Figure 3: RAW data volume on disk. Primary RAW data sample in green, secondary data in yellow. Secondary data are ready to be deleted when disk space is needed. It shows full Run 2 reprocessing did not cause significantly larger disk usage by RAW data and 3PB (day-average) buffer was respected

Sites	2018 Phase I Test (MB/s)	2020 Reprocessing (MB/s)
CERN (CTA Test)	2000	4300
BNL	866	3400
FZK	300	1600
INFN	300	1100
PIC	380	540
TRIUMF	1000	1600
CC-IN2P3	3000	3000
SARA-NIKHEF	640	1100
RAL	2000	2000
NDGF	500	600

Table 1: Stable Rucio tape throughput for the ATLAS Tier-1 sites and CERN, measured from the 2020 reprocessing campaign, with comparison to the Phase I results.

2.1 New distributed software components: Intelligent Data Delivery Service

The Intelligent Data Delivery Service (iDDS) has been developed to orchestrate the Workflow Management and Distributed Data Management systems in order to optimize resource usage in various workflows. It dynamically transforms and delivers data to let computing resources process data on time, i.e., it decouples data pre-processing, delivery, and main processing in each workflow and allows them to run asynchronously.

iDDS has been introduced into ATLAS Distributed Computing to improve inefficiencies in the old Data Carousel scheme that worked with coarser data granularity due to constraints in the workflow. In the old system, tasks were released only when most of the input data were staged-in from tape storage, which led to a significant delay before processing could start and required huge disk caches for the entire processing period. iDDS propagates the detailed information on the input data status from Rucio to the Job Execution and Definition Interface (JEDI), as shown in Figure 4, and allows JEDI to incrementally release tasks so that tasks can start processing even if input data are only partially staged-in. iDDS has been in production since the middle of 2020, and it has solved the issue with the delayed processing in bulk reprocessing campaigns.

3 Distributed data management

The ATLAS experiment has developed the Rucio data management system to keep track of its files and datasets across the worldwide distributed data centers throughout their whole lifetime. Rucio also manages the orchestration of transfers to enable coordinated physics processing as well as user analysis and long-term archival. The actual data movements are executed by the File Transfer Service (FTS) [7], which is one of the *transfer tools* with which Rucio communicates.

Rucio is directly integrated with the ATLAS workflow management systems, and thus orchestrates the necessary input and output files for the brokered jobs. To support the Data Carousel mode, several characteristics had to be understood and addressed. The first was that we noticed that all datasets were going to a single site in each regional cluster of sites, when

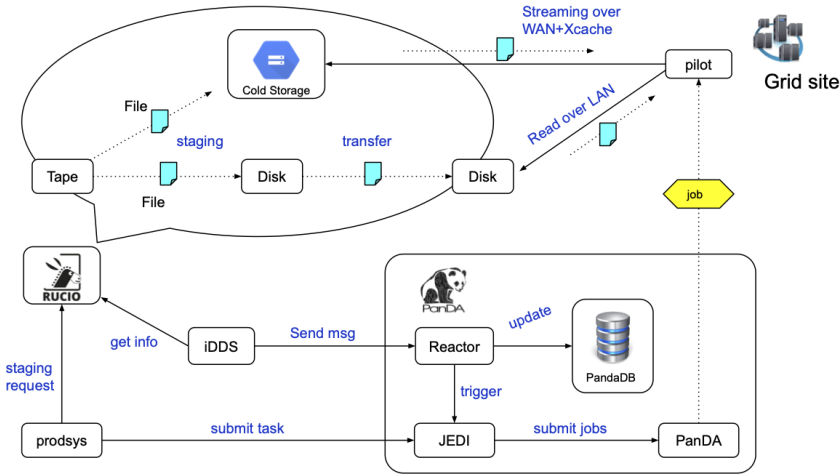


Figure 4: Data Carousel with iDDS

they should have been distributed by the weight of free space at the sites. Rucio uses as concept called replication rules to enable users and external applications to replicate data. After investigation, it was discovered that the weighting option of a replication rule, which influences the selection of destination storage, is ignored if any files in the dataset have replicas on any storage in the replication rule’s RSE expression, since the algorithm aims to minimize data movement. The bias in the Data Carousel therefore was normal, because some sites had RAW files already available, and so those sites were preferred even though they had less free space. The second issue is related to the pinning of files in the buffer of the tape system after recall. Due to many different scenarios, it is possible that a file is evicted from the buffer before it has a chance to get transferred. This problem leads to a cycle of recall and deletion without any transfer progress. This was addressed through an exhaustive study of FTS pinning mechanism, Rucio pinning mechanisms, and selective adaptations of the necessary timeouts.

Thirdly, the orchestration engine in Rucio has a throttling component which releases transfers in FIFO mode. For the Data Carousel, it is necessary to release the transfer in a smarter grouped FIFO mode. This means that if a transfer is being released, it should also release all the transfers of the same dataset, so they are submitted to FTS in the same time-window to allow subsequent grouping on FTS side, and thus, on the tape system. This mostly covered throttling per activity and destination pair. An additional mode was added to allow the Data Carousel to be able to throttle per destination over all potential activities. The last, and arguably most important extension was the addition of the rule progress meter. Since a single replication rule in Rucio can potentially affect thousands of files, it is necessary for the ATLAS workflow management system to understand how far the replication rule has progressed, and not only wait for the final completion. A new percentage based messaging mechanism was developed, which is directly consumed by the Production System via ActiveMQ.

4 Improvements on tape systems from sites

As explained in our previous paper [8], one important metric in this R&D project is tape efficiency, which is defined by the ratio of the throughput delivered to end users over the

vendor-specified nominal throughput of the tape system. It is essential that ATLAS Distributed Computing teams work together with our facilities' tape experts, in order to achieve the optimal usage of our tape resources.

Over the past two years, the tape storage experts at various sites have made many studies and improvements on both the tape system themselves as well as the frontend of the storage, e.g., at FZK and TRIUMF Tier-1s. This is one of the important reasons for the significant increase of our staging throughput, shown in Table 1. However, there is still big difference between sites in terms of tape efficiency. At sites which group files by datasets on tapes, we observed much higher tape recall efficiency (sometimes close to stream reading speed) compared to other sites. Currently our average tape recall efficiency is approximately 30 to 40%. Tape test results indicate that, with optimal file grouping (a.k.a. smart writing), tape recall efficiency can be doubled.

In ATLAS Distributed Computing, we have started to increase file sizes to tapes, targeting 5 to 10 GB files. This will help us with tape efficiency in Run 3. For the longer term, better file placement on tapes, achieved through more sophisticated writing mechanisms, colloquially called *smart writing*, will be our goal. All these efforts call for continued collaboration between tape facilities and the experiments; not only ATLAS, because Tier-0 and many Tier-1 sites support multiple experiments.

5 Summary and future plans

We successfully and quickly completed the ATLAS Data Carousel R&D project phases, involving collaboration between ATLAS, FTS, dCache, CTA and the WLCG centers. During the first year we obtained metrics vital for the project's success. Many problems were encountered and solved along the way. Known challenges, such as smart data writing, data grouping on tape and meta-information passing between ATLAS distributed computing and storage management software still remain. During full Run 2 data reprocessing, i.e., 18.5 PB of RAW data, ATLAS demonstrated the real Data Carousel mode in action, in a production environment with many other concurrent activities such as data writing, data rebalancing, or data consolidation between ATLAS Grid sites. New software modules are being evaluated and integrated with the ATLAS Production System, which mitigates the latency issue of staging inputs from tape directly. Deep integration and communication protocols between data and workflow management systems were defined and implemented. We have evaluated the optimal file size to enable more efficient tape I/O and, based on this, file size has been increased for data produced by prompt reprocessing, i.e., Tier-0 data processing and by the Production System. We will continue to move forward on the various areas of Data Carousel project. In particular, we will work closely with various service providers, such as the dCache, CTA and FTS teams, to improve scalability of services, and also explore ways to increase tape recall efficiency, starting with smart writing mechanisms. We also plan to evaluate how end-user analysis can be run in Data Carousel mode with data staging from tape. Several cross-experiment Data Carousel exercises are under discussion as well, including a large scale reprocessing campaign conducted simultaneously by multiple experiments. In Run 3, we expect that major campaigns requesting data from tape will run in Data Carousel mode while we continue to improve tape recall efficiency and grow tape capacity towards the needs of the HL-LHC.

Acknowledgments

The work at Plekhanov University is funded by the Russian Science Foundation grant (project No.19-71-30008). The work at Brookhaven National Laboratory is funded by the U.S. Department of Energy, Office of Science, High Energy Physics contract No. DE-SC0012704

References

- [1] LHC – The Large Hadron Collider, <http://lhc.web.cern.ch/lhc/>
- [2] ATLAS Collaboration, The ATLAS Experiment at the CERN Large Hadron Collider, *JINST* **3** S08003 (2008).
- [3] F. H. Barreiro et al, The ATLAS Production System Evolution: New Data Processing and Analysis Paradigm for the LHC Run2 and High-Luminosity, *J. Phys.: Conf. Ser.* **898**, 052016 (2017)
- [4] F H Barreiro et al, PanDA for ATLAS distributed computing in the next decade, *J. Phys.: Conf. Ser.* **898** 052002 (2017)
- [5] M. Barisits et al, Rucio: Scientific Data Management, *Comput. Softw. Big Sci.* (2019) **3**: 11
- [6] J. Shiers, World LHC Computing Grid, *Computer Physics Communications* **177**, 219–223, (2007)
- [7] E. Karavakis et al, FTS improvements for LHC Run-3 and beyond, *EPJ Web of Conferences* **245**, 04016 (2020), DOI : <https://doi.org/10.1051/epjconf/202024504016>
- [8] M. Barisits et al, ATLAS Data Carousel, *EPJ Web of Conferences* **245**, 04035 (2020), DOI : <https://doi.org/10.1051/epjconf/202024504035>