Accounting and monitoring tools enhancement for Run 3 in the ATLAS distributed computing

Aleksandr Alekseev[1,*], *Dario* Barberis[2], and *Michal* Svatos[3]
on behalf of the ATLAS Computing Activity

[1]University of Texas at Arlington, TX, USA
[2]Università e INFN Genova, Italy
[3]Czech Academy of Sciences, Czech Republic

Abstract. The ATLAS experiment at the LHC utilizes complex multicomponent distributed systems for processing (PanDA WMS) and managing (Rucio) data. The complexity of the relationships between components, the amount of data being processed and the continuous development of new functionalities of the critical systems are the main challenges to consider when creating monitoring and accounting tools able to adapt to this dynamic environment in a short time. To overcome these challenges, ATLAS uses the unified monitoring infrastructure (UMA) provided by CERN-IT since 2018, which accumulates information from distributed data sources and then makes it available for different ATLAS distributed computing user groups. The information is displayed using Grafana dashboards. Based on the information provided, they can be grouped as "data transfers", "site accounting", "jobs accounting" and so on. These monitoring tools are used daily by ATLAS members to spot and fix issues. In addition, LHC Run 3 required the implementation of significant changes in the monitoring and accounting infrastructure to collect and process data collected by ATLAS during the LHC run. This paper describes the recent enhancements to the UMA-based monitoring and accounting dashboards.

1 Introduction

One of the key challenges for distributed computing in the ATLAS experiment [1] at the Large Hadron Collider (LHC) [2] is the effective accounting and monitoring of computational resources usage. At the current stage of the experiment, a significant improvement of the existing monitoring tools is necessary to provide comprehensive access to data generated during the operation of distributed systems such as PanDA [3] (workflow management system) and Rucio [4] (data management system), as well as their components. One such tool that needs improvement is a monitoring application based on the Unified Monitoring Infrastructure [5] (UMA). This infrastructure, provided by CERN IT and based on modern NoSQL software solutions including [6], Kafka [7], ElasticSearch [8], InfluxDB [9], Grafana [10], etc., is used by various CERN experiments and departments for collecting information from distributed sources, storing and visualizing a large volume of diverse data.

*e-mail: aleksandr.alekseev@cern.ch

The first version of this monitoring tool for ATLAS was developed and implemented in 2018 [11] and has been continually upgraded since then to meet the evolving requirements of the experiment. Today, this tool represents a universal monitoring environment for all possible dashboards, which are used by different user groups to detect and spot issues in the computational infrastructure, obtain retrospective data, create reports, and so on.

For LHC Run 3 the main focus in the development of the monitoring environment is on updating existing dashboards and creating new ones that provide complete information regarding data popularity for the distributed systems mentioned above. This paper provides an overview of existing dashboards and the latest updates for them, as well as new dashboards dedicated to data popularity.

2 Overview of the ATLAS Monitoring Environment in UMA

To access ATLAS data that is collected by UMA, the Grafana visualization system is used. This system implements a generic environment consisting of a set of dashboards divided into logical groups depending on the information they provide. There are several large groups of dashboards:

- Jobs Monitoring. The group contains dashboards designed to monitor the execution of computing jobs on computing infrastructure by the PanDA system.
- Data Monitoring is used to monitor data operations and storage for data. All this information is provided by the Rucio system.
- Data Popularity. This is a special group of dashboards that allows users to get statistics on the most frequently used datasets using information from Rucio and PanDA systems.
- Site Monitoring, consists of dashboards for monitoring the computing infrastructure at all sites participating in ATLAS distributed computing operations.
- Other monitoring tools. This group contains various dashboards that cannot be classified as above and are developed to monitor various other services and some components of the PanDA system, for example: ART [12], EventIndex [13] Harvester [14], iDDS [15], dataset information, etc.

2.1 Jobs Monitoring and Accounting

Dashboards from this group allow users to obtain comprehensive information about computational jobs generated by the PanDA system. There are two types of dashboards, depending on the type of information used: aggregated and non-aggregated. Aggregated information is available since 2010 and updated every hour; it is used for tabular reports and plots in the Jobs Accounting and HS23 (CPU consumption) Reports dashboards.

Using aggregated information has a number of pros and cons. The pros include the speed of plot building regardless of the selected time period, and also the gradual volume increase of new data in the data source for the plots. The cons of this representation of information are the inability to obtain complete information about a specific job. An example is a plot visualizing information about the number of running job slots for each activity over the last ten years, shown in Figure 1.

In addition to the "Jobs Accounting", a "Jobs Monitoring" dashboard is available to users. It uses enriched non-aggregated data from the PanDA database for plotting. This dashboard provides comprehensive information about completed jobs for the last 3 months, with the possibility to obtain the parameters of a specific job. A new heatmap plot, as seen in Figure 2 on this dashboard, illustrates the distribution of jobs depending on the duration of their processing.

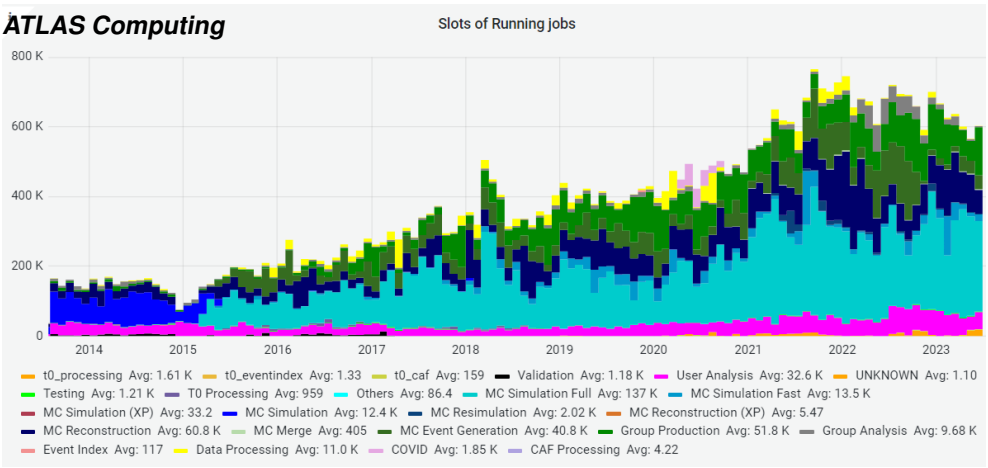

Figure 1. Slots of running jobs by activity for the last 10 years. Each bin shows the number of CPU cores used on average by each distributed computing activity during one month.

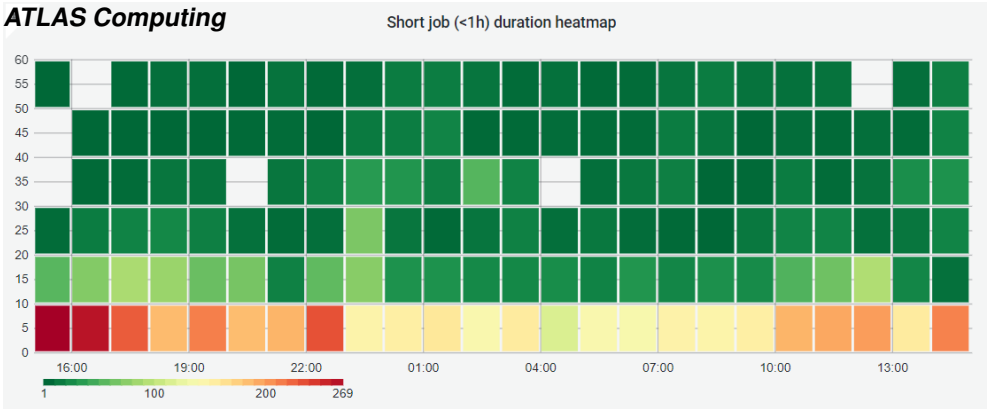

Figure 2. Heatmap of the job duration for jobs shorter than one hour. The x axis displays the completion time (in 1-hour bins), the y axis the job duration (in 10-minute bins). Each bin contains the number of completed jobs with the given range of durations (recorded every 10 seconds). This is useful for quick debugging of low CPU efficiency which can be caused by very short jobs.

2.2 Data Management Monitoring and Accounting

Another important group in the ATLAS environment is devoted to monitoring the data and the state of data storage elements. There are multiple Rucio data sources from which the unified monitoring infrastructure receives and processes information for end users. In total, the following dashboards can be distinguished depending on the information they provide:

- DDM Global Accounting provides information about data evolution in ATLAS over time as shown in Figure 3. Data for the dashboards is updated weekly from a snapshot of the Rucio database.
- DDM Transfers is mainly used by shifters and experts for global transfer monitoring and problem detection on the Grid.

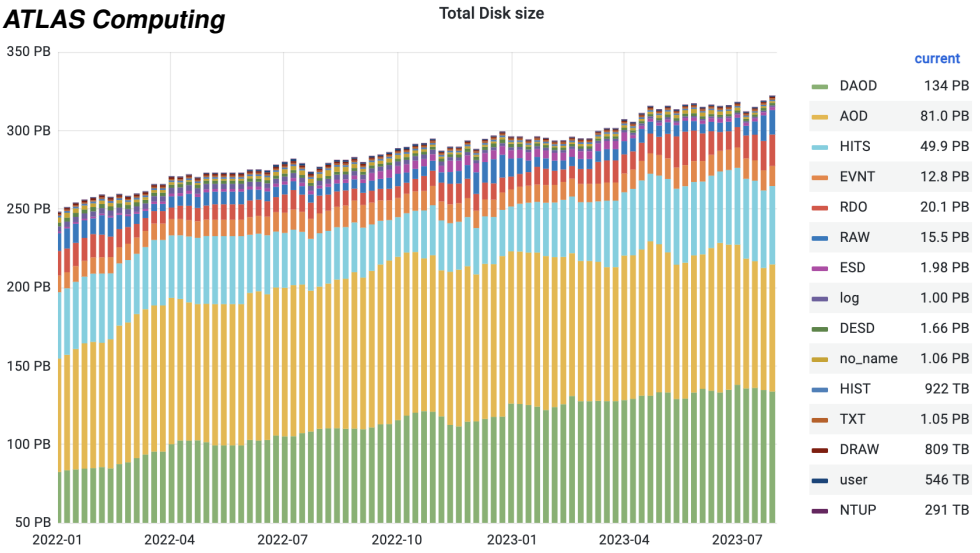

Figure 3. Histogram of total data size evolution on disk in 2022-2023, split by data type.

- DDM Site Accounting shows the change in space usage over time in storage elements on the sites. The information from this dashboard is mainly used by operators and site administrators.

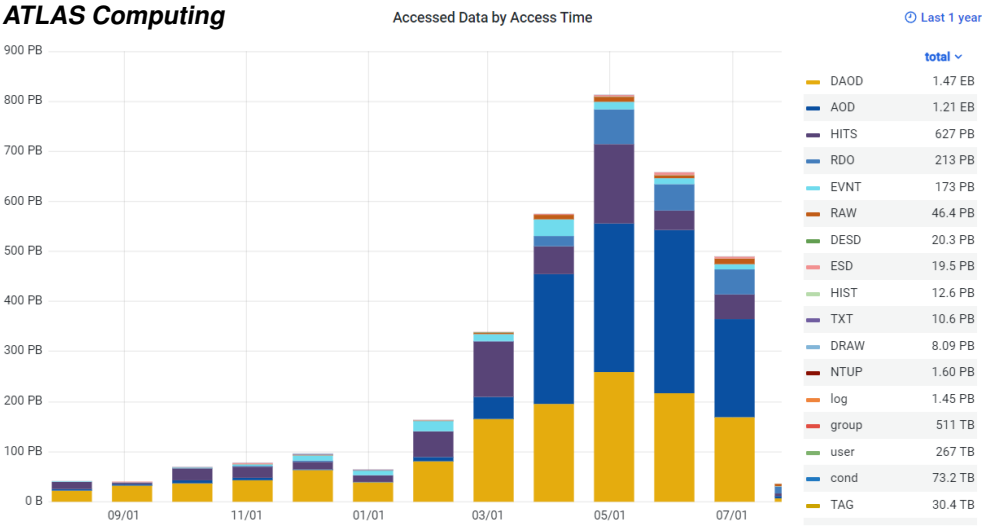

Figure 4. Accessed data in petabytes by dataset type for the past year (2022-23). Each bin corresponds to the month containing the last access date for each dataset; the height of the bin is the sum of the sizes of all datasets that were last accessed during that month.

2.3 Data Popularity

The increasing volume of data to process during Run 3 requires the improvement of special monitoring and analytic tools to identify the most popular datasets. The development of such tools can significantly increase the efficiency of data access by placing more highly-requested data on fast storages, as well as deleting unused data from them. The Data Popularity [16] dashboard, recently developed as a part of the DDM global accounting framework, is a tool that allows users to get information about used and unused datasets using weekly snapshots of tables in the Rucio database. An example of a plot from this dashboard in Figure 4 gives a complete picture of the data volume (in PB) residing on disk that has been used during the last year, by last access date and data type.

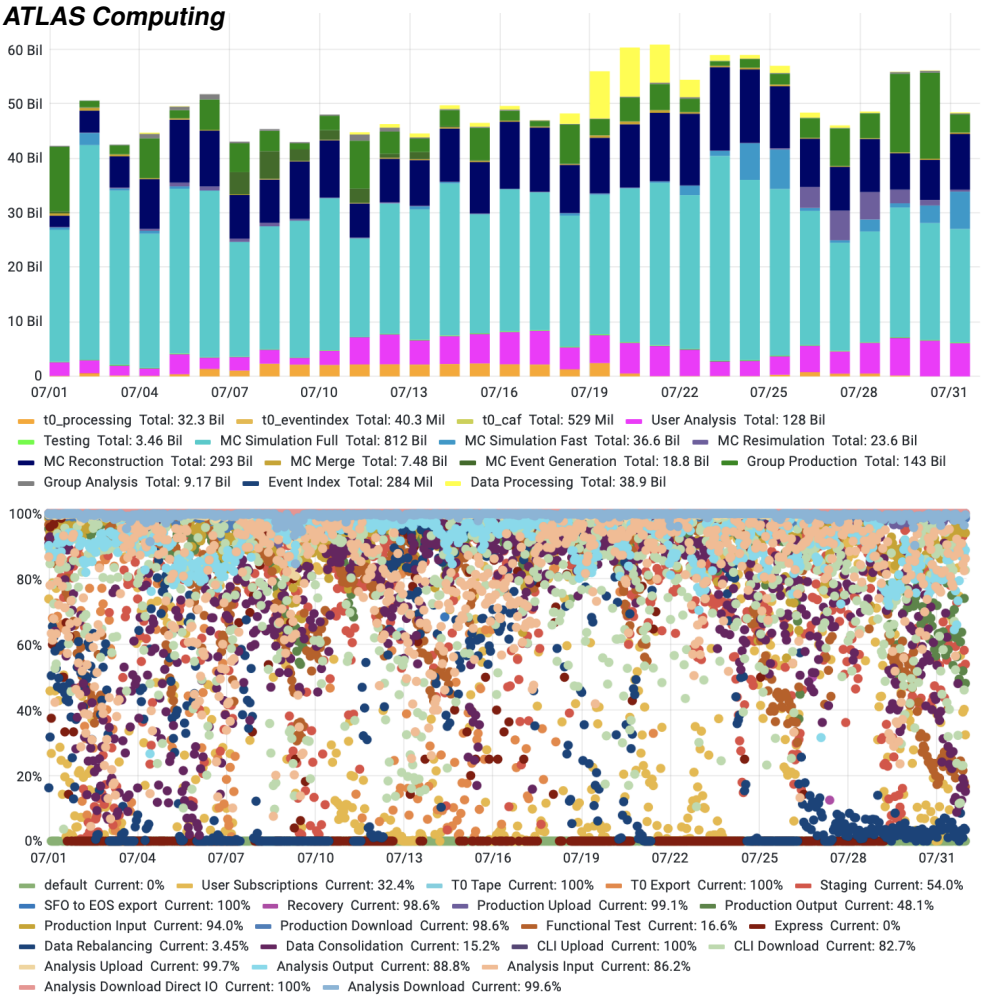

Figure 5. CPU consumption from PanDA (top) and transfer efficiency from Rucio (bottom) plots for one month from the Site-oriented dashboard, separately for each distributed computing activity.

2.4 Site Monitoring

Another large group of dashboards is dedicated to site monitoring and is mainly used by site administrators to identify potential problems and fix them in their computing infrastructure. The Site-oriented dashboard is the main one in this group; its key feature is that it combines plots based on information from Rucio and PanDA. This combination of plots makes it much easier to compare information about data transfers and processing for further analysis. Figure 5 shows an example of a plot combination.

2.5 Other monitoring tools

The dashboards in this group provide information about the operation of various PanDA components such as Harvester and iDDS, about the datasets used by Panda jobs, and other ATLAS services like Tier-0, ART and EventIndex..

Dataset Info is a new dashboard that was created using information about datasets utilized by the compute tasks generated by the PanDA system. The information for the dashboard is updated every hour, using a cron job to extract from the PanDA database relevant data about tasks and datasets used. One of the graphs from this dashboard is shown in Figure 6 and can be compared to the data popularity graphs in DDM accounting to get more detailed information on datasets usage.

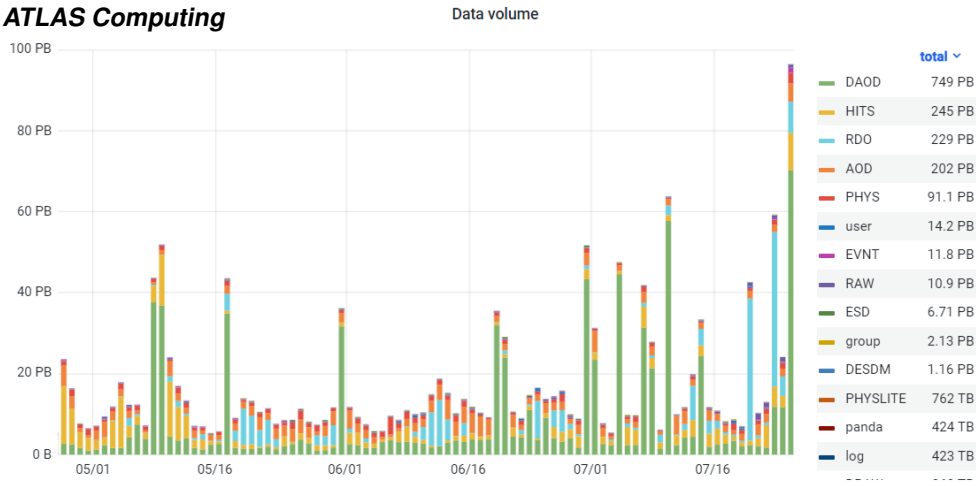

Figure 6. The amount of data consumed and generated by the PanDA system grouped by dataset type for 90 days.

3 Summary

Today, the unified monitoring infrastructure (UMA) provided by CERN IT is a flexible and powerful tool for creating monitoring dashboards that are used daily by various user groups to obtain real-time and historical information about the operation of many systems and subsystems used in the ATLAS collaboration, for any time period. These dashboards continue to be actively developed, so new slices of monitoring in the ATLAS monitoring environment have been created to provide information about ATLAS data popularity. It helps to optimize all data stores in the computing infrastructure, thereby responding to one of the modern challenges of LHC Run 3.

References

[1] ATLAS Collaboration, Journal of Instrumentation **3**, S08003 (2008), https://dx.doi.org/10.1088/1748-0221/3/08/S08003

[2] L. Evans, P. Bryant, Journal of Instrumentation **3**, S08001 (2008), https://dx.doi.org/10.1088/1748-0221/3/08/S08001

[3] T. Maeno, K. De, T. Wenaus, P. Nilsson, R. Walker, A. Stradling, V. Fine, M. Potekhin, S. Panitkin, G. Compostella, Journal of Physics: Conference Series **396**, 032071 (2012), https://dx.doi.org/10.1088/1742-6596/396/3/032071

[4] M. Barisits et al., Comput. Softw. Big Sci. **3**, 11 (2019), https://doi.org/10.1007/s41781-019-0026-3, 1902.09857

[5] A. Aimar, A.A. Corman, P. Andrade, S. Belov, J.D. Fernandez, B.G. Bear, M. Georgiou, E. Karavakis, L. Magnoni, R.R. Ballesteros et al., Journal of Physics: Conference Series **898**, 092033 (2017), https://dx.doi.org/10.1088/1742-6596/898/9/092033

[6] *Apache flume*, https://flume.apache.org/

[7] *Apache kafka*, https://kafka.apache.org/

[8] *ElasticSearch: An open search and analytics solution*, https://www.elastic.co/

[9] *Influxdata: Influxdb times series data platform*, https://www.influxdata.com/

[10] *Grafana: The open observability platform | grafana labs*, https://grafana.com/

[11] Beermann, Thomas, Alekseev, Aleksandr, Baberis, Dario, Crépé-Renaudin, Sabine, Elmsheuser, Johannes, Glushkov, Ivan, Svatos, Michal, Vartapetian, Armen, Vokac, Petr, Wolters, Helmut, EPJ Web Conf. **245**, 03031 (2020), https://doi.org/10.1051/epjconf/202024503031

[12] T. Cuhadar Donszelmann et al., EPJ Web Conf. **245**, 05015 (2020), https://doi.org/10.1051/epjconf/202024505015

[13] D. Barberis et al., Comput Softw Big Sci **7**, 00096 (2023), https://doi.org/10.1007/s41781-023-00096-8

[14] Maeno, Tadashi, Barreiro Megino, Fernando Harald, Benjamin, Doug, Cameron, David, Childers, John Taylor, De, Kaushik, De Salvo, Alessandro, Filipcic, Andrej, Hover, John, Lin, FaHui et al., EPJ Web Conf. **214**, 03030 (2019), https://doi.org/10.1051/epjconf/201921403030

[15] W. Guan, T. Maeno, B.P. Bockelman, T.J. Wenaus, F. Lin, S. Padolski, R. Zhang, A. Alekseev, CoRR **abs/2103.00523** (2021), https://doi.org/10.1051/epjconf/202125102007, 2103.00523

[16] M. Grigoryeva, A. Klimentov, M. Schulz, N. Popova, M. Shubin, *Operational Analytics Studies for ATLAS Distributed Computing: Data Popularity Forecast and Utilization of the WLCG Centers* (2023), in these Proceedings