BNL-226158-2024-JAAM

The ATLAS Workflow Management System Evolution in the LHC Run3 and towards the High-Luminosity LHC era

A. Pacheco Pages, A. Klimentov

To be published in "EPJ Web of Conferences"

May 2024

Physics Department

Brookhaven National Laboratory

U.S. Department of Energy

USDOE Office of Science (SC), High Energy Physics (HEP)

Notice: This manuscript has been authored by employees of Brookhaven Science Associates, LLC under Contract No.DE-SC0012704 with the U.S. Department of Energy. The publisher by accepting the manuscript for publication acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

EPJ Web of Conferences **295**, 04043 (2024) https://doi.org/10.1051/epjconf/202429504043
CHEP 2023

The ATLAS Workflow Management System Evolution in the LHC Run3 and towards the High-Luminosity LHC era

Andrés Pacheco Pages[1,2],* and *Alexei* Klimentov[3]

[1]Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra (Barcelona), Spain
[2]Port d'Informació Científica (PIC), Campus UAB, 08913 Bellaterra (Cerdanyola del Vallès), Spain
[3]Brookhaven National Laboratory, 98 Rochester St, Upton, NY 11973, USA

Abstract. The ATLAS experiment has 18+ years of experience using workload management systems to deploy and develop workflows to process and to simulate data on the distributed computing infrastructure. Simulation, processing and analysis of LHC experiment data require the coordinated work of heterogeneous computing resources. In particular, the ATLAS experiment utilizes the resources of 250 computing centers worldwide, the power of supercomputing centres, and national, academic and commercial cloud computing resources. In this contribution, we present new techniques for cost-effectively improving efficiency introduced in workflow management system software. The evolution from a mesh framework to new types of computing facilities such as cloud and HPCs is described, as well as new types of production and analysis workflows.

1 Introduction

The Workflow Management System (WFMS) is an essential tool for the ATLAS experiment [1]. It allows physicists to manage their computing jobs and access the computing resources they need.

In the past, the WFMS was limited to grid resources. However, in recent years, the WFMS has been adapted for the transparent use of HPC and cloud resources. This means that scientists can now use the computing resources of supercomputing centers and cloud platforms, together with grid resources.

This adaptation has been a great success. It has allowed scientists to access more computing resources, which has led to greater scientific output. For example, the use of the cloud has allowed scientists to accelerate data analysis [2].

The ATLAS WFMS is an example of ATLAS's commitment to scientific excellence. The experiment continues to work to improve its computing infrastructure.

2 The ATLAS WFMS: A Reliable and Scalable Platform for Data Processing

The ATLAS Workflow Management System (WFMS) is responsible for processing detector and Monte Carlo simulated data on a computing infrastructure made available for the ex-

*e-mail: pacheco@ifae.es
Copyright 2023 CERN for the benefit of the ATLAS Collaboration. CC-BY-4.0 license.

© The Authors, published by EDP Sciences. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

periment. The WFMS has been constantly evolving since the experiment began operation at CERN's LHC accelerator [3]. The main functionalities are based in the PanDA ecosystem [4]. The objectives of the WFMS are:

- To process the detector and Monte Carlo simulated data efficiently and effectively.
- To provide an easy-to-use interface for scientists working with the data.
- To handle large amounts of data reliably.
- To be scalable as the experimental data grow.

3 Grid Computing: The Key to Success for the ATLAS Experiment

The ATLAS experiment initially used the Worldwide LHC Grid (WLCG) [5] as its computing infrastructure. The WLCG remains the primary source of computing resources for ATLAS today.

Grid computing is a distributed computing model that unites computing resources scattered around the world under a single infrastructure. This is achieved by homogenizing the resources at the level of operating system compatibility, middleware, and authentication and security system.

The benefits of grid computing include:

- Access to a large amount of computing resources: Grid computing can provide access to a much larger pool of computing resources than any single institution could provide on its own. This allows scientists to run large-scale simulations and analyses that would not be possible otherwise.
- Scalability: Grid computing is scalable, meaning that it can be easily expanded to accommodate more users and more computing resources. This is important for experiments like ATLAS, which are constantly generating new data.
- Reliability: Grid computing is reliable, meaning that it is designed to continue operating even if some of the resources become unavailable. This is important for experiments like ATLAS, which need to be able to process data continuously.
- Security: Grid computing is secure, meaning that it is designed to protect data from unauthorized access. This is important for experiments like ATLAS, which are working with hundreds of institutions around the globe.

The WLCG is a large and complex system, but it has been essential for the success of the ATLAS experiment.

4 The WFMS Team: Essential for the Success of the ATLAS Experiment

The Workflow Management System (WFMS) team is a group within the ATLAS Distributed Computing Area (ADC). It is composed of 12 people, most of whom are software developers. WFMS developers directly contribute to the infrastructure operation effort as second-level support, in addition to developing new features for existing software.

The WFMS team is responsible for:

- Developing and maintaining the workflow management software for ATLAS.
- Operating and maintaining the computing infrastructure for workflow management services.

- Providing second-level technical support to scientists working with ATLAS data.

The WFMS team is an essential team for the success of the ATLAS experiment. They have developed workflow management software that is efficient and effective, and they have operated and maintained the computing infrastructure for ATLAS reliably. The WFMS team has also provided excellent technical support to scientists working with ATLAS data.

The workflow management software that the WFMS team has developed for ATLAS is being used for other experiments. This software has been well-received by other experiments, and it has helped to improve the efficiency and effectiveness of these experiments.

The WFMS team is a world-class team that has made a significant contribution to the ATLAS experiment.

5 The Grid: Still a Key Technology for High-Energy Physics

The grid was initially developed thanks to the efforts of the high-energy research centers. At that time, there was no other solution capable of meeting the LHC's computing needs. However, the current situation is different. Now there are other relevant sources of computing resources, such as cloud computing and quantum computing.

The grid also helps to ensure that the data from high-energy physics experiments are available to a wider audience. Currently the grid provides to the ATLAS experiment around 62% of the computing resources as seen in fig. 1.

In the future, the grid will continue to play an important role in high-energy physics research. It will be combined with other technologies, such as cloud computing, to create a more powerful and flexible computing infrastructure. This infrastructure will allow scientists to conduct more sophisticated research and make more important particle physics discoveries.

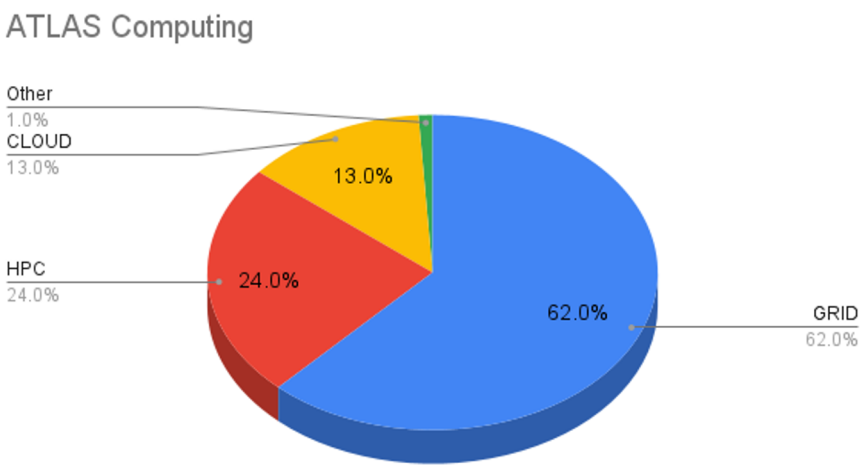

Figure 1. Distribution of ATLAS jobs' wallclock time consumption in Grid, HPC and Cloud resources from August 2022 to July 2023. The total wallclock time corresponds to 20 trillion HS06-seconds.

6 Supercomputer Centers

Supercomputer centers are an alternative source of computing resources. They are emerging all over the world to meet the needs not only of high-energy physics, but of any area of science.

Unlike grid resources, supercomputer centers do not adapt to the experiment. Instead, it is the experiment that adapts to the configuration of each supercomputer center. Compared to grid resources, supercomputer centers could deliver more computing capacity within shorter time frames. This increased power allows researchers to conduct more complex analyses in their experiments.

Supercomputer centers will become increasingly important for scientific research. They provide a scalable and reliable way to access large amounts of computing power. This allows scientists to conduct more sophisticated research and make more important discoveries.

Supercomputer centers do not provide resources on a continuous basis. Instead, users must request and justify the resources they need for limited periods of time. In some centers, if the resources are exhausted, users can use the resources opportunistically, taking advantage of the computing gaps in other projects.

In the case of the ATLAS experiment, from August 2022 to July 2023 24% of the resources came from supercomputer centers as seen in fig. 1. These centers are perfect for massive simulation workloads, which can take advantage of the computing capacity until the quota is full.

Here are some specific examples of how ATLAS uses the resources of supercomputer centers:

- The experiment uses supercomputer centers to run LHC simulations, which help scientists understand how the accelerator works and how it produces data.
- The experiment uses supercomputer centers to analyze LHC data.
- The experiment uses supercomputer centers to develop new tools and methods for data analysis, which help scientists work more efficiently and effectively.

Supercomputer centers are an essential part of the ATLAS infrastructure. They provide scientists with the resources they need to conduct their research.

7 The Rise of Cloud Computing

Cloud resources are another alternative source of computing resources that is very relevant today [6]. Cloud resources are paid for, but they have an innate advantage that they are very elastic, which allows for an increase or decrease in computing capacity in a very short time. Cloud resource providers offer discounts if resources are used opportunistically, meaning that jobs can be killed if the resources are needed by full-price paying customers.

In the case of ATLAS, currently 13% of the resources come from commercial cloud providers as seen in fig. 1. This type of resource is perfect for tasks that must be completed in a very short time, such as reconstructing all the physics data from experiments in less than a week.

Here are some specific examples of how ATLAS uses cloud resources:

- The experiment uses cloud resources to store and process large amounts of data.
- The experiment uses cloud resources to develop and run scientific applications.
- The experiment uses cloud resources to train machine learning models.

Cloud resources are an essential part of the ATLAS infrastructure. They provide scientists with the resources they need to conduct their research.

The progressive adaptation of the ATLAS workflow management system for the transparent use of grid, HPC, and cloud resources has been one of the most significant successes in recent years.

8 The Rise of Alternative Computing Architectures

The computing market has undergone a number of changes in recent years, in terms of the design of computing processors. Initially, Intel dominated the market, but in recent years, many alternative architectures have emerged [7], such as ARM [8] and PowerPC. These architectures are more efficient and offer a number of advantages over Intel, such as lower power consumption, smaller size, and higher performance per watt. Additionally, there has been an increase in the use of computing accelerators, such as GPUs. GPUs are ideal for tasks that require a large amount of parallel computing, such as graphics processing and machine learning.

These changes have been driven by a number of factors, such as the growing demand for computing and the need to process large amounts of data. As a result of these changes, the computing market is becoming more competitive and diversified.

This is a positive change for the scientific community, as it means that there are more options available to choose from. It is also positive for the computing industry, as it means that there is more innovation and competition.

The computing market is constantly evolving, with new processor designs emerging to meet the needs of new applications and devices. It is likely that we will also see the development of new computing architectures and accelerators that are even more efficient and powerful than the current ones.

9 The Evolution of Workflow Management Software to Meet the Challenges of Heterogeneous Computing

The emergence of different types of CPUs and GPUs has forced workflow management software to become more transparent [9]. Currently, each resource in the ATLAS production system is identified with the CPU architecture type, CPU chipset version, and GPU model, if any. The type of hardware on which a job is running can have a significant impact on its efficiency and effectiveness. To ensure that jobs run more efficiently, workflow management software must be able to automatically identify the type of hardware and adjust the job parameters accordingly.

The ATLAS workflow management software is designed to meet this requirement. It can automatically identify the type of hardware and adjust the job parameters accordingly, ensuring that jobs run efficiently and effectively on the available hardware.

The ability of the ATLAS workflow management software to use different types of CPUs and GPUs transparently is one of the reasons why the experiment has been so successful. The software has allowed scientists to access a wide range of computing resources, which has allowed them to perform compute-intensive analyses.

10 Helping to Create a More Sustainable Future

One of the current social concerns is the carbon footprint and the economic cost of resources that are used depending on the time, which is indirectly related to the use of more polluting

energies. In this sense, there is a wide diversity of centers that obtain electricity exclusively from renewable sources, while other centers do not.

With the aim of tracing carbon footprint information, all centers participating in computing tasks of the experiment are being asked to provide the origin of the electricity and the consumption of the equipment. Although this information is not currently taken into account when sending the payload, in the future this could be a reality for the good of our planet [10].

These measures are important to reduce the environmental impact of the ATLAS experiment. The experiment uses a large amount of resources, and it is important that these resources are used in a sustainable way. By tracking the carbon footprint of the experiment, scientists can take steps to reduce their environmental impact.

For example, scientists can choose to use data centers that obtain electricity from renewable sources. They can also choose to use more energy-efficient equipment.

11 Closing the Gap Between Official and Physicist-Driven Workflows in ATLAS

In ATLAS, there has always been a divide between the official workflows that process the experiment's real and simulated data for analysis, which are highly automated processes, and the physicists' data analysis workflows, which are highly manual. Since 2022, efforts have been made to close this gap between the two workflows by integrating the analysis workflows that can be executed automatically as part of the experiment's data processing chain. This allows for centralized and more efficient management of analyses, as data can be analyzed in the same center where it is generated.

The benefits of this integration are numerous. First, it allows scientists to focus on data analysis rather than workflow management. Second, it allows scientists to share their findings more easily with other scientists. Third, it allows scientists to conduct more complex analyses and obtain more precise results.

The integration of analysis workflows is an important step in improving the efficiency and effectiveness of the ATLAS experiment.

12 Data Access Efficiency with New Strategies

In line with the previous section, all data processing optimizations are limited by their availability and the possibility of being accessed more or less quickly. Often this availability involves massive movements of data between centers, and storage and network resources are always limited [11].

Several strategies are being implemented to make more efficient use of disk storage. One of these strategies is to move less accessed data to slower storage, such as tape robots [12]. Another strategy is to recreate the data at the time it is requested as input, without having a pre-existing disk copy waiting to be used.

In the latter case, it is important to weigh the storage cost against the cost of generating the file. It is also important to have an infrastructure elastic enough to give maximum priority to file recreation so that they can be served as if they were on disk.

These strategies are helping to improve the efficiency and effectiveness of the ATLAS experiment. They allow scientists to access data more quickly and efficiently.

13 Summary

The text discusses the ATLAS workflow management system (WFMS), which is responsible for processing the experiment's data, both real and simulated, within a computing infras-

tructure made available to the experiment. The WFMS has been constantly evolving since the beginning of the experiment at CERN's Large Hadron Collider, and now uses a variety of sources of computing resources, including supercomputing centers, cloud resources, and distributed computing resources.

In recent years, the WFMS has been adapted to be more efficient and effective. For example, it can now use different types of processors and accelerators, which allows it to process data more quickly and efficiently. The WFMS has also been adapted to be more sustainable, and now uses renewable energy sources and other technologies to reduce its environmental impact.

The WFMS is a valuable tool for the ATLAS experiment, and is an essential part of the experiment's computing infrastructure. The WFMS will continue to be an essential tool for the experiment in the years to come. Here are some specific conclusions from the text:

- The WFMS has been constantly adapting to be more efficient and effective, and now uses a variety of computing resources.
- The WFMS is a sustainable tool, and uses renewable energy sources and other technologies to reduce its environmental impact.
- The WFMS is an essential part of the ATLAS experiment's computing infrastructure, and will continue to be an essential tool for the experiment in the years to come.

Acknowledgements

We gratefully acknowledge the following funding and support:

- The MICINN in Spain under grants PID2019-110942RB-C22, including FEDER funds from the European Union.
- The U.S. Department of Energy, Office of Science, High Energy Physics contract No. DE-SC0012704.
- The participants and contributors to the ADC Technical Interchange Meeting workshop, which was held virtual in December 2022.
- The members of the ADC Workflow Management System group for their crucial role in the project: T. Alekseev, F. Barreiro-Megino, M. Borodin, K. De, D. Golubkov, M. Grigoryeva, T. Korchuganova, W. Guan, E. Karavakis, V. Kotliar, F.H. Lin, T. Maeno, P. Nilson, T. Wenaus, X. Zhao.
- The computing centers and personnel of the Worldwide LHC Computing Grid and other centers for delivering so effectively the computing infrastructure essential to our analyses.

We are grateful for the support of all of these individuals and organizations. Their contributions have made this work possible.

References

[1] ATLAS Collaboration, JINST **3** S08003 (2008) 10.1088/1748-0221/3/08/S08003

[2] M. Dunford et al., *Active Learning reinterpretation of an ATLAS Dark Matter search constraining a model of a dark Higgs boson decaying to two b-quarks*, ATL-PHYS-PUB2022-045

[3] L. Evans and P. Bryant (editors), JINST **3** S08001 (2008) 10.1088/1748-0221/3/08/S08001

[4] The PanDA Production and Distributed Analysis System. panda-wms.readthedocs.io/en/latest
[5] J. Shiers, Comput. Phys. Commun. **177** 219–223 (2007) 10.1016/j.cpc.2007.02.021
[6] F.H. Barreiro Megino, *Accelerating science: the usage of commercial clouds in ATLAS distributed computing*, these proceedings
[7] IDC (2023): Reports that in 2022, ARM shipments grew 8.8% while Intel shipments declined 3.3%. (Source: idc.com:prUS51383823
[8] J. Elmsheuser, *The ATLAS experiment software on ARM*, these proceedings
[9] T. Maeno et al., *Utilizing Distributed Heterogeneous Computing with PanDA in ATLAS*, these proceedings
[10] R. Walker, *Sustainability in HEP Computing*, these proceedings.
[11] A. Klimentov et al., *Operational Analytics Studies for ATLAS Distributed Computing: Data Popularity Forecast and Ranking of the WLCG Centers*, these proceedings
[12] X. Zhao et al., *Updates to the ATLAS Data Carousel Project*, these proceedings