

# Evolution of ATLAS analysis workflows and tools for the HL-LHC era

David Cameron<sup>1</sup>, Alessandra Forti<sup>2</sup>\*, Alexei Klimentov<sup>3</sup>, Andrés Pacheco Pages<sup>4,5</sup>, and David South<sup>6</sup>

<sup>1</sup>University of Oslo, P.b. 1048 Blindern, 0316 Oslo, Norway

<sup>2</sup>School of Physics and Astronomy, University of Manchester, Oxford Rd, Manchester, M13 9PL, UK

<sup>3</sup>Brookhaven National Laboratory, 98 Rochester St, Upton, NY 11973, USA

<sup>4</sup>Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra (Barcelona), Spain

<sup>5</sup>Port d'Informació Científica (PIC), Campus UAB, 08913 Bellaterra (Cerdanyola del Vallès), Spain

<sup>6</sup>Deutsches Elektronen-Synchrotron (DESY), Notkestr. 85, 22607 Hamburg, Germany

**Abstract.** The High Luminosity LHC project at CERN, which is expected to deliver a ten-fold increase in the luminosity of proton-proton collisions over LHC, will start operation towards the end of this decade and will deliver an unprecedented scientific data volume of multi-exabyte scale. This vast amount of data has to be processed and analysed, and the corresponding computing facilities must ensure fast and reliable data processing for physics analyses by scientific groups distributed all over the world. The present LHC computing model will not be able to provide the required infrastructure growth, even taking into account the expected evolution in hardware technology. To address this challenge, several novel methods of how end-users analysis will be conducted are under evaluation by the ATLAS Collaboration. State-of-the-art workflow management technologies and tools to handle these methods within the existing distributed computing system are now being evaluated and developed. In addition the evolution of computing facilities and how this impacts ATLAS analysis workflows is being closely followed.

## 1 Introduction


The experiments at the Large Hadron Collider [1] use a worldwide complex and distributed computing infrastructure with almost 1 million computing cores and an exabyte of storage, interconnected through high-speed networks. The bulk of data processing to produce end-user analysis objects is done through the Worldwide LHC Computing Grid, the WLCG [2], which consists of hundreds of individual sites worldwide at universities and national laboratories. However, the extreme computing needs of the experiments running from 2027 in the High Luminosity LHC (HL-LHC) era, primarily for data processing and analysis that are crucial for physics results, will not be satisfied by the current infrastructure, even allowing for the expected decrease in hardware costs (see Figure 1). The ATLAS experiment [3] is

---

\*e-mail: [Alessandra.Forti@cern.ch](mailto:Alessandra.Forti@cern.ch)

@Copyright [2020] CERN for the benefit of the ATLAS Collaboration. Reproduction of this article or parts of it is allowed as specified in the CC-BY-4.0 license

therefore exploring the use of new technologies and new computing facilities. The evolution of how computing facilities will be organised and consolidated will play a key role in how any possible shortage of resources will be addressed. Technologies that will address the HL-LHC computing challenges may be applicable for other scientific communities in high-energy physics, astronomy and beyond to analyse large-scale data volumes.



**Figure 1.** Projected CPU requirements of ATLAS between 2020 and 2034 based on 2020 assessment [4]. Three scenarios are shown, corresponding to an ambitious (“aggressive”), modest (“conservative”) and minimal (“baseline”) development programme. The black lines indicate annual improvements of 10% and 20% in the computational capacity of new hardware for a given cost, assuming a sustained level of annual investment. The blue markers with the dotted lines represent the 3 ATLAS scenarios following the present LHC schedule.

One of the computing challenges for particle physics experiments is the evolution of data analysis. For ATLAS, bulk data processing and reduction down to analysis objects is traditionally performed on WLCG resources, with final results produced on resources local to the analyser. In recent years however there has been an explosion of ideas and technologies from the wider data science community, some of which can be and have been applied to analyses of ATLAS data. These include Machine Learning and Deep Learning techniques, use of alternative hardware such as GPU and FPGAs, and a Python-based ecosystem of numerical libraries for vectorised array computation. It is expected that by the time of HL-LHC, data analysis through these technologies will become mainstream, and therefore it is a requirement of the distributed computing systems that they evolve accordingly.

To address the HL-LHC distributed data handling challenge, ATLAS has launched several R&D projects to study the feasibility of setting up dedicated computing facilities for end-user analysis, to evaluate new analysis workflows (many using Machine Learning and Artificial Intelligence), and to identify new tools to be developed to describe more complex analysis workflows. This paper describes the next generation of analysis tools in ATLAS, ideas of the roles of analysis facilities and the changes required in distributed computing software.

## 2 Physics analysis in ATLAS

In this section we review the progressive reduction of the data formats for analysis. Reconstructed real and simulated data are stored in Analysis Object Data (AOD) files. These are too large data for each analyser to go through independently. To facilitate timely analysis of the ATLAS data and reduce the amount of CPU cycles used, smaller formats are produced centrally. In Run-2 (2015-2018) the AOD datasets were processed in the ATLAS derivation framework [5], producing about 80 different Derived AOD (DAOD) formats that contain a subset of events and reduced reconstruction information tailored for specific physics analysis and performance groups. These DAOD types are then processed by individual users on the grid to produce final ntuples which are then analysed on non grid resources.

Concerns about the storage resources required for the DAOD format led to an R&D program that will introduce progressively smaller data sizes. For Run-3 (2022-2024), a universal derived format has been introduced, DAOD\_PHYS [6]. This is a reorganisation of the same format to avoid duplication of events in different derivations. It is foreseen the vast majority of analysis will use this format, which has already been tested. For the HL-LHC, an even smaller derived dataset, DAOD\_PHYSLITE, containing already calibrated physics objects, will likely be the only frequently updated format available to physicists due to the greatly increased sample size. The aim of a format like DAOD\_PHYSLITE is to reduce the overall size of the HL-LHC analysis data from 100 PB per year it would have with the current format to 2 PB. This is important to be able to distribute the data to analysis facilities. Table 1 presents the typical sizes per event.

**Table 1.** Event sizes of different formats of Analysis Object Data (AOD) and Derived AOD (DAOD).

xAOD Type	Size per event
AOD	600 kB
DAOD	40 – 450 kB
DAOD_PHYS	50 kB
DAOD_PHYSLITE	10 kB

## 3 Analysis facilities

### 3.1 The current computing infrastructure

The ATLAS experiment employs a sophisticated distributed computing system (ADC), which comprises hundreds of clusters and associated storage at the WLCG sites. These resources are managed by interconnected workflow management (PanDA [7] and ProdSys [8]) and data management (Rucio [9]) systems. The sites are connected via a global data network and organised as a series of Tiers, with an increasing number of sites per Tier with correspondingly decreasing size. The large and unique Tier-0 resource at CERN is tasked with recording and promptly reconstructing the collision data from the LHC, as well as providing the custodial archive. Ten geographically separated Tier-1 sites store a tape copy of detector data, and along with around 60 Tier-2 sites they provide the bulk of the computing resources for ATLAS data processing and Monte Carlo simulation, including the central production of the analysis formats described above. Whilst the original ATLAS computing model associated the various Tier-2 sites around a given Tier-1, based typically on national or geographical groupings with common funding agencies, this rigid hierarchical distinction has long since disappeared. The evolution of WLCG networks over the last two decades has meant that transfers are no longer limited to within a single grouping of sites, and that any stable and

performant Tier-1 or Tier-2 site, a so called nucleus, can aggregate the output of a task, i.e. a group of related jobs that can be executed at any site.

Several High Performance Computing (HPC), Cloud Computing and other opportunistic resources have been integrated into ADC where possible, in order to augment the capacity of the ATLAS managed Tier-1 and Tier-2 sites. As many of these resources are volatile they are not considered suitable for custodial roles or to provide time-critical processing and storage services. Thus far ATLAS has used such resources almost exclusively for Monte Carlo simulation purposes, as to avoid any necessary add-ons or payload complications that may be associated with analysis or other more complex workflows.

The provision of both Tier-1 and Tier-2 resources is based on bilateral agreements, or pledges, between the WLCG and contributing funding agencies that cover the entire duration of the full LHC program, including the HL-LHC. In addition to production workflows, the WLCG Tier-1 and Tier-2 sites provide distributed processing capability for the user analysis of ATLAS data, at the scale of up to a million analysis jobs executed per day. However, these resources are not well suited for the final interactive stages of user analysis, such as visualisations, plot generation, re-weighting, systematic studies, limit testing, and so on. To fulfil this role dedicated local batch clusters or facilities are typically utilised, such as those deployed at CERN, DESY in Germany or BNL and SLAC in the USA. These local resources (often referred to as Tier-3s), are funded and managed outside the WLCG framework and pledges, and provide the necessary, additional resources for end-user analysis. However the resource requirements for user analysis have evolved, and it has become desirable to have more integrated solutions featuring not only reliable batch systems but also the latest, relevant interactive tools. This provides an opportunity to bridge both the gap between batch on the grid and on shared facilities as well as between non-interactive and interactive analysis by expanding on the grid and local batch/Tier-3 concepts. Resources that can provide such integrated solutions are referred to as ‘analysis facilities’, and focus on interactivity, usability, as well as a strong user support. They are defined more by the set of applications they offer, rather than the resources on which they run.

### 3.2 Potential implementations of analysis facilities

The ideal analysis facility is one or more dedicated resources with user support and federated access for all ATLAS users. This of course is going to be difficult to get funded so one has to envisage a number of potential implementations. They may be dedicated resources, integrated into large existing infrastructures such as Tier-1s or be fully virtualised and deployed on private or commercial cloud or HPC providers. Each have pros and cons in terms of funding, availability, user accessibility and support, and proximity to the data. It is likely that most of the implementations described below will co-exist in some form or other, and ATLAS will ensure that it can take advantage of all of them wherever possible.

**Dedicated Analysis Facilities:** A computing resource that is dedicated to data analysis and is not used for centralised data processing tasks such as simulation or reconstruction. It provides the necessary software and hardware (e.g. GPUs) to enable a diverse range of data analysis. It could be a national centre for the use of a particular region or a pledged resource for the whole ATLAS community and will typically provide interactive access through notebooks for example.

**Tier-3 evolution:** Many institutes and laboratories provide shared computing facilities to local users. In ATLAS these are typically used for the end-stages of user analysis after all bulk processing has been done on the Grid. For convenience these resources are sometimes linked to the Grid, for example to make it easier to move data between the facility and Grid storage. Leveraging and expanding existing resources is technically a low hanging fruit but

since these resources are locally-owned and not pledged they are not normally accessible to the wider ATLAS community and the challenge in opening up these resources is therefore more political than technical.

**Co-location with Tier-1 and Tier-2 centres:** Tier-1 and Tier-2 centres form the backbone of the Grid and perform the majority of data processing on resources pledged to ATLAS. Some of these centres currently provide analysis facility-like resources such as GPUs, but these are generally opportunistic and shared with other communities. In order to provide dedicated resources to ATLAS, a framework needs to be in place to account for the analysis facility usage as part of the pledge delivered to ATLAS by these centres. This framework must also ensure a fair balance between resources provided as analysis facility and those currently used for Grid data processing.

**HPCs as Analysis Facilities:** HPCs present both a challenge and an opportunity for HL-LHC. Technology and policy differences compared to typical Grid sites require significant changes in software and computing models. However, the rewards can be large if physicists are able to take advantage of the most powerful machines in the world boasting cutting-edge technology. Interactive data analysis poses a particular challenge in that it is the exact opposite of the extremely large parallel sandboxed batch workflows preferred by HPCs. Therefore it is unlikely that most HPC centres will run data analysis, but there may be some opportunity with upcoming HPC infrastructures with more interactive services.

**Commercial Clouds:** Cloud computing was designed to scale up and down elastically to meet user demand, and therefore seems ideally suited to unpredictable and bursty user analysis. In addition it is far more cost-effective to “rent” a specialised hardware resource which is only required for a short period of time than purchasing and installing it in a data centre. There have been ongoing projects in ATLAS to use Amazon and Google clouds [10] which have successfully integrated these resources into ATLAS’ workflow and data management systems. The main challenge of cloud computing is data access and management, due to the costs involved, and therefore it is likely that mainly computationally-intensive rather than data-intensive workloads will be run there.

### 3.3 Further considerations

In addition to the implementation issues described in the previous section, which are primarily concerned with the location and access of a potential analysis facility, several other details on its configuration also need to be considered.

Firstly, a real-life scenario in which several analysis facilities are employed is more likely than a single entity for all users. Whilst the R&D phase will most likely begin with a single analysis facility instance, this would be expanded over time to be several facilities, or a federation of facilities for practical reasons, as it would be impractical, inconvenient and of high risk for all users to rely on one discrete set of resources.

Another important consideration is whether analysis facilities are to be considered pledged resources and included in the current infrastructure described in section 3.1. In this case they would be provided for the whole collaboration by the same funding agencies and with the same funding mechanisms as the Tier-1 and Tier-2. Furthermore, if they are to be pledged resources, it must be decided if analysis facilities are part of the current pledge, thus reducing the budget available for the current model, or something in addition requiring new, dedicated funding.

Whilst the Tier-1 and Tier-2 sites employ a variety of compute, storage and transfer technologies, the overall interface presented to the end user is rather uniform. However, in the case of analysis facilities, this may not necessarily be the case, and the nature, size and range

of tools and levels of functionality provided by each instance may differ depending on the location and host infrastructure.

Access controls to analysis facilities must also be considered, whether they should follow a similar model to existing WLCG resources, or continue to be more regional, dependent on the local funding, analogous to the Tier-3s. A proper accounting methodology must also be factored into any analysis facility implementation, which maybe prove difficult when more exotic resources such as GPUs or interactive notebooks are involved. There are also user support questions that need to be resolved. Currently ATLAS has more than 1500 active users that need to be supported. This alone makes the idea of expanding the pledged resources to cover the analysis facilities more important to support the services.

## **4 New technologies and tools**

One of the key aspects of the Analysis Facilities R&D described in this document is to explore and integrate new technologies which allow new types of workflows to run. Experiments have benefited from the homogeneity agreed in terms of CPU architecture, Operating System (OS) and grid middleware for the WLCG resources. Towards HL-LHC, this uniformity is progressively less guaranteed and a much greater diversity of resources with a range of OS, hardware architectures and configurations will need to be integrated.

As described earlier, the analysis of ATLAS data will be based on smaller formats such as DAOD\_PHYSLITE which will allow users to carry out analysis both with and without the full ATLAS software and environment. This ability to decouple the analysis from the ATLAS framework opens the door to data analysis with different approaches and different software stacks. This section examines how different new technologies might help integrating different types of analysis workflows in a more heterogeneous environment and yet still offering the user a uniform experience.

### **4.1 Containers**

ATLAS has successfully introduced generic OS containers for most workflows as a way to abstract the OS layer and separate the payload environment from the host environment. This has worked well on the grid to support the current Run-2 type of analysis using the ATLAS framework but it is still tied to centrally-managed distribution of official software releases. To fully exploit the power of containers for the evolution of analysis and also for production, ATLAS has also integrated the possibility to run standalone containers as if they were a normal payload. This work started with the simpler user analysis submission [11], which has been further developed to run more complicated (Machine Learning) ML workflows as we will illustrate later and has now been expanded to streamline the building and testing of containers to run simulation and possibly reconstruction at HPCs [12]. Streamlining image production for all releases and the full adoption of standalone containers to replace standard software distribution would help with the integration of diversified resources. In the analysis case it would allow users to take advantage of any kind of resource (analysis facilities, grid, cloud, HPC) with a single container. The expansion of container streamlined production beyond the HPC use case is under review.

### **4.2 Kubernetes and containerisation**

Kubernetes [13] is a container orchestrator developed to automate application and services deployment, scaling and management. As far as analysis facilities are concerned, it can be



used to deploy a batch system such as HTCondor [14] as a service [15], it can be used as a batch system directly, or it can run as a back-end for services like JupyterHub [16] or services that are being developed to make analysis more efficient as described in section 4.3. This flexibility and the ability to expand and shrink resources for different services makes it an important technology to explore in its capability. At a typical grid site it could be the glue between grid services and the interactive services characteristic of an analysis facility simplifying the deployment of multiple services required at HL-LHC scale as would be the case in an Analysis Facility co-located to a grid site. ATLAS has already explored integrating Kubernetes as a batch system in its Workflow Management System (WFMS) [17]. Harvester [18], the ATLAS WFMS component that submits jobs has been adapted to interact with Kubernetes. The simplest method is for it to treat the pods of a Kubernetes cluster as standard batch system worker nodes using nested containers, that is the pods run the pilots which in turn run the containerized payloads. The system could be simplified by running the pilot functionality within individual pods corresponding to each production step. Submission to Kubernetes works both on grid sites and without extra effort on commercial clouds [10]. Kubernetes activities are still R&D and the use of Kubernetes in ATLAS for service deployment, as a batch system and as a backend of user services is under review.

### 4.3 New services and tools development

The landscape of an HL-LHC facility will be somewhat different with an ecosystem of services under development to partly minimise the amount of data that needs to be served to the users and partly to make the analysis development more agile. There are three ways to integrate them: one is for the Workflow Management Software to interact with the service on the user or application behalf, another is incorporating the functionality directly and finally to redesign the clients to offer a more uniform access to grid services and Analysis Facilities services.

**JupyterHub** Since the language for high-level analysis is shifting to Python, users want to use more interactive tools to develop their analysis, and Jupyter Notebooks are becoming the standard. JupyterHub is a way to run Notebooks in a scalable way. It can use different backends but is natively built on kubernetes and can make effective use of its horizontal scaling capabilities. A JupyterHub at an Analysis Facility could be seen as the evolution of the more traditional batch system, with additional interactive capabilities.

**REANA Hub** [19] is a REANA (Reusable Analyses) platform enabling users to structure their research data analysis in view of enabling future reuse. Like JupyterHub, REANA Hub is natively integrated with Kubernetes and Docker [20]. It offers the users the ability to monitor what the jobs are doing in real time. It is different to JupyterHub in that it uses a workflow engine and a declarative approach to submitting the analysis. REANA Hub can run sequential and more complex workflows with thousands of steps (such as Directed Acyclic Graph (DAG)).

**iDDS** (Intelligent data delivery service) [21] is a service developed to transform and deliver only the data required to a given application. It has several capabilities but the main features being used by the WFMS are the dynamic data management capability and an integrated decision engine and task chain management. The former is used to serve data to the application from tape as they are requested thus reducing the jobs queuing time. It is currently heavily used in production [22] and will be evaluated for analysis to provide a faster job turn around on the grid particularly for the type of analysis that will still require larger datasets. The latter is used to run multi-step DAG workflows and chain tasks on the grid. Such a workflow in a single cluster would rely on the ability of jobs to communicate with

each other. In a distributed environment like the grid a dedicated service is needed to take the decisions and this function is carried out by iDDS.

**Command Line Clients** are technically not a service, but are being expanded to add functionalities. ATLAS users typically submit jobs to the grid using quite powerful command lines clients which often require long and complicated strings of options and are completely different from, for example, the declarative approach of a REANA facility or the Graphical User Interface (GUI) of Jupyter notebooks. The development carried out on the clients is quite important. The functionality to express the options in a declarative way using a JSON file has been added, but this is still for grid only clients. Effort is being invested in an API that can be integrated into JupyterLab giving users the possibility to use uniform tools when using the grid and when using other Analysis Facilities. Prototype JupyterLab GUIs are currently being implemented to interface with a standard JupyterHub, submit jobs to the grid with integrated grid commands, interact with iDDS to refine the decision engine parameters and query a rucio catalogue.

#### 4.4 Authentication and authorisation implementations

The grid authentication and authorisation implementations (AAI) were built on an X.509 certificate infrastructure. This has worked well for almost two decades but is starting to show its limitations particularly when trying to integrate cloud resources and technologies like JupyterHub and kubernetes described in sections 4.2 and 4.3, which use native token based authentication. Rudimentary support for WLCG tokens is already available in PanDA and some Grid services, but there is currently no official ATLAS service for end-users to obtain tokens and no consensus on federated access to facilities. Frameworks and policies must be developed to provide access and properly account for usage of these facilities, as well as the technical implementation in the various services and tools. Moving away from X.509 is not something experiments can do in isolation, and requires a concerted effort. Work to adopt a token AAI based on the industry standard protocols like OAuth2.0 and OpenID Connect is ongoing at the WLCG level [23].

The new workflows and clients will be the first to be adapted to use the new AAI enabling the users to engage with tokens in an already new environment.

#### 4.5 Accelerators and distributed computing

ATLAS is exploring the use of ML workflows that are suited to run on GPUs. GPUs are available at different scale and with different degrees of accessibility. This ranges from a small number of GPUs at some grid sites, a significantly larger number on commercial clouds, and several hundreds at large scale installations such as HPC centres, which are not particularly suitable for small users or R&D models. ML and GPUs also require specialised software that is not part of the standard ATLAS software distribution and is dependent on the particular GPU model and manufacturer. WFMS jobs and task brokering to GPUs is a subject of R&D and will evolve greatly in the future. At the moment it has a very simplistic model adapted from the more uniform CPUs use case with a single architecture. This was good enough to support a handful of users and running COVID applications [24] but as soon as the GPU models were expanded it demonstrated insufficient.

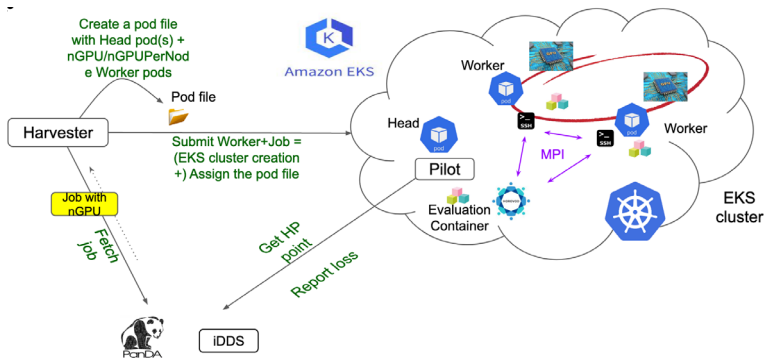
#### 4.6 Example of a new workflow

An example of R&D that brings several of these technologies together is the Hyper Parameter Optimisation (HPO) service [25]. This is a set of WFMS functions designed to run multi-step



ML containerised workflows using the iDDS decision engine to select which tasks to submit depending on the output of previous tasks as described in the iDDS section of this paper. The HPO service was developed using GPUs at grid sites and was expanded to submit work to a Kubernetes cluster on the Amazon cloud accessing multiple GPUs. The JupyterLab clients can now interact also with the HPO service. An important workflow being adapted to use it is FastCaloGAN [26] which will be used to parametrize the next generation of ATLAS fast simulation.

A diagram of this workflow is shown in Figure 2 with Harvester creating the Kubernetes resource description yaml file with the appropriate parameters to use multiple GPUs and "submitting" it. The evaluation pods then talk between each other and with the head pod which in turn talks to iDDS to send the HP points and the loss, and get a decision on how to continue.



**Figure 2.** A Multi-GPU Hyper Parameter Optimisation workflow on a Kubernetes cluster on the Amazon cloud.

## 5 Evolution of ATLAS distributed computing services

The workflow and data management systems of ATLAS distributed computing have been built to handle large-scale batch processing and have so far succeeded in providing the resources required for the wide breadth of ATLAS physics results. With the paradigm shifts in techniques and technologies described in the previous sections, these systems must adapt to be ready to handle the diversity of data analysis in the HL-LHC era.

Integration with new services and tools is key to providing a smooth analysis experience. This begins with ensuring a single point of entry and access control, i.e. a federated identity system allowing access to all resources with a single credential and the same tools for both the grid and the analysis facilities as described in section 4.3 and 4.4.

A "Data Analysis as a Service" (DAaaS) front-end would provide users access to specialised resources for data analysis through the same interfaces as those used to access grid resources. The work described in section 4.2 to seamlessly integrate containerised workloads with ATLAS distributed computing services will provide a means both to encapsulate workloads and exploit a variety of heterogeneous resources through a common technology. A DAaaS service could also offer interactive access to notebooks and act as a common interface to data transformation services such as iDDS.

The ATLAS Production System (ProdSys) handles centrally-managed data processing tasks such as simulation and reconstruction campaigns, but it could be extended to also handle data analysis tasks from end-users where appropriate. For example, chained workflows

with multiple interdependent steps would be better managed by ProdSys than the users themselves. Along with technological changes, improvements can be gained through more intelligence in existing tools and policies. For example, end-user analysis data objects (DAOD) are currently distributed among grid sites following strict ATLAS policies on the number of copies and their lifetime. However some of this data is more popular than others and there may be different access patterns over the life of the data. A smarter data management policy is required which proactively increases or decreases the number of copies of data based on its popularity between users and real or predicted access patterns. ATLAS may also consider a copy of less popular DAOD samples on tapes, keeping the most popular samples on disks with increased number of replicas. Data may also be dynamically distributed to where it can be optimally processed, according to the kind of analysis facilities used at the time. It is the subject of an R&D to develop new algorithms for dynamic management of DAOD replicas and selection of the best site to place additional replicas to improve physics analysis performance and turnaround time.

## 6 Conclusions

This paper describes physics analysis workflows in ATLAS and the current computing infrastructure, and discusses various types of Analysis Facilities and their potential implications. New technologies, services and tools relevant to ATLAS, both under active development and the subject of future research is also examined. Changes and workflows needed by ATLAS distributed computing services to fulfill the new requirements for improving the analysis experience for the next years are also discussed. Over the coming years, the fruits of these research and development activities will provide a solid base to address the impending challenge of HL-LHC computing.

## Acknowledgements

We would like to acknowledge support from the MICINN in Spain under grants PID2019-110942RB-C22 and FPA2016-80994-C2-2-R, including FEDER funds from the European Union. 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. We also thank the participants and contributors to the ADC Technical Interchange Meeting workshop, which was held virtual in December 2020, which has brought together many ideas for this paper. With particular emphasis we thank Tadashi Maeno, Fernando Harald Barreiro Megino, Lukas Heinrich, and Johannes Helmsheuser for their crucial role in the R&D program. In addition, we gratefully acknowledge the computing centres and personnel of the Worldwide LHC Computing Grid and other centres for delivering so effectively the computing infrastructure essential to our analyses.

## References

- [1] L. Evans and P. Bryant (editors), JINST **3** S08001 (2008)
- [2] J. Shiers, Comput. Phys. Commun. **177** 219–223 (2007)
- [3] ATLAS Collaboration, JINST **3** S08003 (2008)
- [4] ATLAS Collaboration, *ATLAS HL-LHC Computing Conceptual Design Report*, doi:10.3204/PUBDB-2020-04972, (2020)
- [5] J. Catmore et al., J. Phys.: Conf. Ser. **664** 072007 (2015)
- [6] J. Elmsheuser et al., EPJ Web Conf. **245** 06014 (2020)

- [7] T. Maeno et al., J. Phys.: Conf. Ser. **396** 032071 (2012)
- [8] F.H. Barreiro et al., J. Phys.: Conf. Ser. **898** 052016 (2017)
- [9] M. Barisits et al., Comput. Softw. Big Sci. **3** 11 (2019)
- [10] J. Elmsheuser et al. *Seamless integration of commercial Clouds with ATLAS Distributed Computing*, these proceedings
- [11] A.C. Forti et al., J. Phys.: Conf. Ser. **1525** 012059 (2020)
- [12] N. Ozturk et al., *Containerization in ATLAS Software Development and Data Production*, these proceedings
- [13] *Kubernetes*, <https://kubernetes.io> (2021), accessed: 2021-06-07
- [14] *HTCondor* <https://research.cs.wisc.edu/htcondor/> (2021), accessed: 2021-06-07
- [15] L.F. Alvarez et al., EPJ Web Conf. **245**, 07048 (2020)
- [16] *JupyterHub*, <https://jupyter.org/hub> (2021), accessed: 2021-06-07
- [17] F.H. Barreiro et al., EPJ Web Conf. **245** 07025 (2020)
- [18] F.H. Barreiro et al., EPJ Web Conf. **245** 03010 (2020)
- [19] T. Simko et al., EPJ Web Conf. **214** 06034 (2019)
- [20] *Docker*, <https://www.docker.com> (2021), accessed: 2021-06-07
- [21] W. Guan et al., *An intelligent Data Delivery Service for and beyond the ATLAS experiment*, these proceedings
- [22] X. Zhao et al., EPJ Web Conf. **245**, 04035 (2020)
- [23] B. Bockelman et al., EPJ Web Conf. **245** 03001 (2020)
- [24] A.C. Forti et al., *The fight against COVID-19: Running Folding@Home simulations on ATLAS resources*, these proceedings
- [25] R. Zhang et al, *Hyperparameter optimisation for Machine Learning using iDDS*, IRIS-HEP workshop, (2020) <https://indico.cern.ch/event/960587/contributions/4070339/attachments/2130656/3588122/IRIS20201027.pdf>, accessed: 2021-06-07
- [26] H. Ahmed et al., *AtlFast3: Next Generation of Fast Simulation in ATLAS*, submitted to Comput.Softw.Big Sci.