

Computing activities at the Spanish Tier-1 and Tier-2s for the ATLAS experiment towards the LHC Run3 and High-Luminosity periods

Santiago González de la Hoz^{1,*}, *Carles* Acosta-Silva^{2,3}, *Javier* Aparisi Pozo¹, *Jose* del Peso⁴, *Álvaro* Fernández Casani¹, *José* Flix Molina^{3,5}, *Esteban* Fullana Torregrosa¹, *Carlos* García Montoro¹, *Julio* Lozano Bahilo¹, *Almudena* Montiel⁴, *Andrés* Pacheco Pages^{2,3}, *Javier* Sánchez Martínez¹, *José* Salt Cairols¹, and *Aresh* Vedae^{2,3}

¹Instituto de Física Corpuscular (IFIC), centro mixto CSIC-Universitat de València, Paterna, Spain

²Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra (Barcelona), Spain

³Port d'Informació Científica (PIC), Campus UAB, 08913 Bellaterra (Cerdanyola del Vallès), Spain

⁴Departamento de Física Teórica y CIAFF, Universidad Autónoma de Madrid, Madrid, Spain

⁵Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain

Abstract. The ATLAS Spanish Tier-1 and Tier-2s have more than 15 years of experience in the deployment and development of LHC computing components and their successful operations. The sites are already actively participating in, and even coordinating, emerging R&D computing activities and developing new computing models needed for the Run3 and High-Luminosity LHC periods. In this contribution, we present details on the integration of new components, such as High Performance Computing resources to execute ATLAS simulation workflows. The development of new techniques to improve efficiency in a cost-effective way, such as storage and CPU federations is shown in this document. Improvements in data organization, management and access through storage consolidations (“data-lakes”), the use of data caches, and improving experiment data catalogs, like Event Index, are explained in this proceeding. The design and deployment of new analysis facilities using GPUs together with CPUs and techniques like Machine Learning will also be presented. Tier-1 and Tier-2 sites, are, and will be, contributing to significant R&D in computing, evaluating different models for improving performance of computing and data storage capacity in the High-Luminosity LHC era.

* Corresponding author: santiago.gonzalez@ific.uv.es

1 ATLAS Computing Model (Mesh Model)

In the original computing model of the ATLAS experiment [1] at the LHC accelerator at CERN, Tier-1 centers had associated Tier-2 centers, which were close-by in terms of network connectivity, and they formed a cloud, which in ATLAS means a regional setup of one Tier-1 and its Tier-2s in a certain geographical area. All data flowed to and from Tier-2s via its Tier-1. But more and more Tier-2s had excellent worldwide network connections and could exchange data directly between themselves. This motivated and led to the new ATLAS computing model (Mesh) where there are Nucleus and Satellites sites [2].

Tier-2s with a significant storage and excellent network connections are designated as “Nucleus” sites, passing job production on to smaller Tier-2s (Satellites) in any cloud, exchanging data directly. Currently, all the Spanish ATLAS Tier-2 sites are of the Nucleus type.

2 Spanish Tier-1 and Tier-2s inside ATLAS

The Spanish cloud (ES) [3] inside ATLAS has one Tier-1 at PIC in Barcelona providing 5% of Tier-1s data processing of CERN’s LHC detectors ATLAS, CMS and LHCb. Also, it consists of one federated Spanish Tier-2 (ES-ATLAS-T2) at IFIC in Valencia, at IFAE (co-located in PIC) in Barcelona, and at UAM in Madrid with 60%, 25% and 15% of the resources respectively, and one Tier-2 in Lisbon, Portugal. ES-ATLAS-T2 represents 4% of the total Tier-2s resources and together with the Tier-1, they are integrated in the World Wide LHC Computing GRID (WLCG) project [4] and strictly follow the ATLAS computing model.

In Table 1 the Spanish resources in October 2019, together with their availability and reliability, are shown. The ES cloud is at the top of availability and reliability ranks.

Table 1. Resources and performance of the Spanish cloud facilities for the ATLAS experiment.

Site	CPU (HEP-SPEC06)	Disk (TB)	TAPE (PB)	Availability (2019)	Reliability (2019)
PIC-Tier-1	42300	3500	8.8	100%	100%
IFIC-Valencia	26751	2600		99.75%	99%
IFAE-Barcelona	10420	980		100%	100%
UAM-Madrid	10358	1220		97%	96%

3 Development Activities for ATLAS Distributed Computing

The Spanish sites are actively participating in, and even coordinating, emerging R&D computing activities aimed at developing the new computing models needed for the LHC Run3 and HL-LHC periods.

Spanish sites are participating in the DOMA-TPC test and storage system performance studies for the implementation of the tape carousel [5]. We are developing a catalog of data of all events in all processing stages, called Event Index [6], which is needed to meet multiple use cases and search criteria. The ES cloud is helping the Event Service project where the main goal is to allow for a more flexible and efficient usage of CPUs available when running ATLAS simulation jobs. Finally, we are designing and deploying novel analysis facilities on GPUs and adopting Machine Learning (ML) methods. We also contribute to monitoring tasks, like monitoring frontier-servers, site and cloud support tools and live pages.

3.1 Data access studies and storage performance

Data access and storage performance studies are important to further optimize data organization and management. In the ES cloud these studies have been carried out by researchers of the CMS experiment working at Spanish sites, that run dCache gathering data from the billingDB [5]. The goal of these studies is to answer several questions. How is the storage system utilized? Is it working in the most optimal way? How to identify data that could usefully be cached and what benefits would this bring? If the sites are close enough (10 ms in terms of latency), should we aim for a data federation and/or consolidation of the deployed storage in the region?

Detailed studies are being performed on how the files are accessed in time, with the aim of detecting cases in which storage system usage might be optimized, as well as identifying all of the possible cases in which data could be cached, turning into optimized deployed storage at the site. Simulations of caches and how to properly handle deletion of files from the cache, based on real usage or applying ML methods, are being investigated as well, prior to deploying caches in production in the region. The aim is to perform the same studies for the ATLAS experiment at some point soon, to have a coherent and almost complete view from a regional perspective (except LHCb).

3.2 Monitoring frontier-servers

Frontier-servers [7] optimize access to the so-called “conditions” database. It contains the operational parameters of all ATLAS subdetectors that are needed to run simulation or production jobs. A squid server provides caching of data and a Tomcat servlet connects to the Oracle database when required.

The monitoring system collects information about the queries from “log” files. It operates in a steady and stable way, processing 12M queries daily on average during data taking periods across all the ATLAS sites and allows the visualization of meaningful information by means of a dashboard. Figure 1 shows one of its pages with several relevant summary tables and histograms.

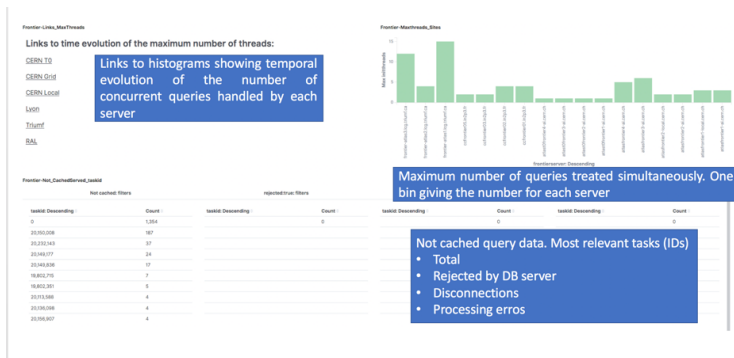


Fig. 1. Monitoring frontier-servers dashboard developed by the ES cloud.

The system incorporates an Alarms&Alerts subsystem to send e-mail warnings when a site performance deteriorates. CMS servers are also monitored.

3.3 Event Index project

The Event Index project [6] aims to provide a catalog of real and simulated data of all events in all processing stages, which is needed to meet multiple use cases and search

criteria. Billions of events (PetaBytes) have been indexed so far since 2015. Some of the use cases are: event picking, duplicate event checking, overlap, trigger check, event skimming, or trigger counter.

The latest developments are aimed to optimize storage and operational resources, in order to accommodate the higher amount of data produced by ATLAS, which is expected to increase in the future with a prediction of 35 billion new real events per year in Run3, and 100 billion at the HL-LHC. At IFIC we have improved the data collection system [8], and we are currently developing new storage schemas using HBase/Apache Phoenix for the final data backend with promising results.

3.4 Physics case and machine learning

The Spanish sites are involved in several physics analyses. In one of them we are using a sample of simulated events where we have both $t\bar{t}$ resonances, which represent our signal, and Standard Model events as would be observed by an experiment running in a proton-proton collider. Our aim is to eventually apply ML algorithms to ATLAS data in order to improve the identification of the signal events and their separation from background processes [9]. We have evaluated the performance of several standard ML algorithms such as Neural Networks and Random Forest techniques provided by well-known Python libraries (Keras, Scikit-learn, ...), thereby investigating parameters such as the number of layers, number of neurons, activation function, etc, as a first approach to understand the benefit of these methods. Figure 2 shows the AUC (Area Under Curve) as obtained for each resonance mass value from the percentages of signal and background identified as signal by the algorithm after its training. The samples used in Figure 2 come from a repository [10] with public ATLAS toy simulated events.




Fig. 2. AUC vs mass of the signal ($t\bar{t}$ resonances) [10]. NN: Neural Network; MLP (Keras): MultiLayer Perceptron.

At IFIC-Valencia we are using the ARTEMISA (ARTificial Environment for Machine learning and Innovation in Scientific Advanced computing) [11] facility where the core computing resources are based on GPUs (one server with four GPUs NVIDIA Tesla Volta and two servers with only one GPU each). The worker nodes are composed of several Intel Xeon Platinum CPUs and Tesla Volta GPUs. ARTEMISA is being used to find the optimal configuration of each ML algorithm. For this purpose, the code implemented in Python has been tested and first results with this facility in batch mode have been obtained, showing improvements on the performance running on CPUs in the IFIC batch system.

4 Challenges for Run3 and High-Luminosity LHC Periods

Higher luminosity is equivalent to higher data flow, so there is an increase on the horizon and this represents a challenge on storage, network bandwidth and processing power [12]. As can be seen in Figure 3, HL-LHC CPU estimations show a factor of three shortfall to the

flat budget model and a factor of six shortfall to today's estimate in the storage on disk. The HL-LHC will require to increase the network bandwidth by a factor 10.




Fig. 3. CPU (left) and disk storage (right) requirements prediction for Run3 and High-Luminosity LHC periods [12]. Approaches to solve CPU and storage shortfall are shown.

There are a few ways to face the CPU challenge: to use HPCs, cloud computing and High Level Trigger farms; to use fast instead of full simulation and speed up the Monte Carlo generators by a factor two; to run on GPUs but one needs significant time and effort to adapt the software to the new architecture.

Approaches for solving the storage shortfall are: increased investment in computing; new file formats; reduction of data volumes and more use of tape storage. This last option slows down the workflow. This will be the major impact on how we are going to deploy the infrastructure: data-lakes with centers connected to sites with large storage capacity, use of caches, smaller data, more use of the network, etc ...

4.1 Integration of High Performance Computing resources

Opportunistic resources turn out to be a meaningful way to face the future HL-LHC challenges in terms of CPU requirements. Spanish HPC centers have been exploited recently, by deploying all of the necessary edge services that allowed us to integrate the resources into the ATLAS workflow management system.

Lusitania and MareNostrum HPC resource integration in the ATLAS production system started in April 2018 in a collaboration between IFIC and PIC. Since then, we have been granted periods where we have been able to exploit several million hours of HPC resources, including 4.8 million hours in 2019 used by IFIC and PIC on the MareNostrum4 (MN4) facility.

Figure 4 shows how IFIC and IFAE-PIC led ATLAS simulation Monte Carlo production using opportunistic HPC resources. The impact of the start of MN4 in the ES cloud was to jump from 7k to 12k slots. Sixty million events were simulated in 2019 and more than 90% of jobs ended successfully. The Spanish community is working with the HTCondor developers to define a similar approach for CMS, which will allow the execution of CMS pilot jobs in MareNostrum [13,14].

4.2 Increasing bandwidth

Barcelona, Madrid and Valencia are connected by a dedicated network with a bandwidth of 20 Gbits/s that is expected to be upgraded up to 100 Gbits/s in 2020 thanks to REDIRIS (Spanish academic network provider – RedIris-Nova). This will allow the Spanish cloud to meet the network challenges for the Run3 and HL-LHC.

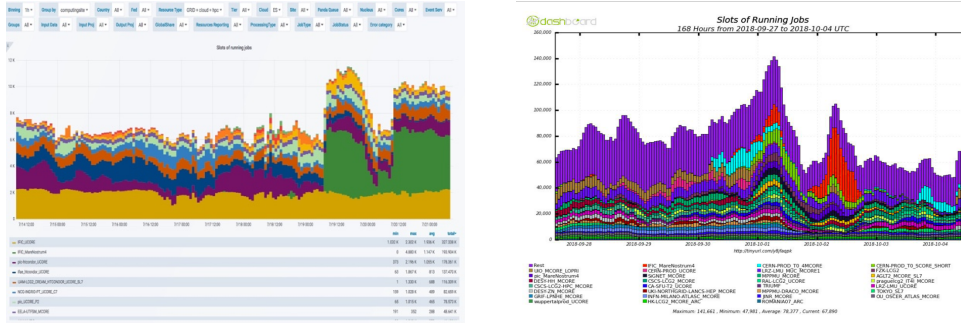


Fig. 4. Left: Impact of the start of MN4 (green) in the Spanish cloud. Right: MN4 HPC (red) resource utilization running ATLAS simulation jobs in October 2018.

4.3 Toward a regional federation (Tier-1-Tier-2s)

With the goal of federating the resources at a national level, first studies have been done in order to understand how latency affects the job workloads when reading data remotely. Overflow of CMS analysis jobs from PIC (Barcelona) to CIEMAT (Madrid) and vice versa was enabled by deploying a regional XRootD re-director with high availability. Since May 2019, pilot jobs from PIC to CIEMAT and vice versa are flocked, since HTCondor BS is deployed in both sites on dedicated machines (80 CPU-cores available at each site). Sites are connected with low latency, measured to be around 10 ms. In these tests, regional input file reads are always preserved, since a regional XRootD redirector was deployed. This allows job degradation to be studied when some jobs are running data remotely in the region.

These tests tell us that within the region, PIC could run jobs either at PIC or CIEMAT (reading files from PIC), except merge, which should run always locally without significant degradation (as expected, since it is the most I/O activity). This working mode would cause an increase in PIC exports, stressing both the network and storage systems. At higher latencies (40 ms), analyses start to be degraded, as seen with some tests made when running ATLAS jobs in the Amazon cloud (AWS) [15]. We plan to perform similar tests with ATLAS in ES sites.

5 Summary

The Spanish cloud contributes around 5% of the total resources deployed in Tier-1 and around 4% in Tier-2 sites. The Spanish Tier-2 has the so-called Nucleus state in ATLAS, which implies significant responsibilities and larger work volume. Not only are we contributing to the deployment of CPU and storage resources for the ATLAS experiment but also several research activities are carried out by the teams in the Spanish cloud.

The current ATLAS computing model will not be able to cope with the Run3 and HL-LHC challenges. CPU resources, disk storage and bandwidth shortfalls have to be assessed and faced, for instance in CPU by using opportunistic resources like HPC and in storage by studying new data format and data processing. Concerning the bandwidth, requirements seem to be satisfied at this time.

MICINN in Spain partially supported this work under grants FPA2016-75141-C2-1-R, FPA2016-75141-C2-2-R and FPA2016-80994-C2-2-R, which include FEDER funds from the European Union. PIC is maintained through a collaboration of CIEMAT and IFAE, with additional support from UAB and ERDF. The authors thankfully acknowledge the computer resources at Mare Nostrum and

Lusitania, and the technical support provided by BSC (FI-2019-3-0017, FI-2019-2-0030, FI-2019-1-0035).

References

1. ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST*, **3**, S08003 (2008)
2. S. Campana, *Distributed Computing in LHC Run2*, *J. Phys: Conf. Ser.* **664**, 032005 (2015)
3. S. Gonzalez de la Hoz, C. Acosta, J. Aparisi, M. Delfino, J. del Peso, A. Fernández, J. Flix, E. Fullana, C. García, J. Lozano, A. Montiel, A. Pacheco, J. Sanchez, J. Salt, A. Vedace, *Spanish ATLAS Tier-1 & Tier-2 perspective on computing over the next years*, *EPJ Web of Conference*, **214**, (2019) 03013
4. Worldwide LHC Computing Grid project: <http://wlcg.web.cern.ch>
5. J. Flix, *CMS data access and usage studies at PIC Tier-1 and Ciemat Tier-2*. In these proceedings
6. D. Barberis, E. Alexandrov, I. Alexandrov, Z. Baranoswski, G. Dimitrov, A. Fernández, E. Gallas, C. García, S. González de la Hoz, H. Hrivnav, A. Kazymov, M. Mineev, F. Prokoshin, G. Rybkin, J. Sanchez, J. Salt, M. Villaplana, *The ATLAS Event Index for LHC Run3*. In these proceedings
7. N. Ozturk, M. Si Amer, A. Formica, J. Lozano, E. Gallas, I. Vukotic, *An Information Aggregation and Analytics System for ATLAS Frontier*. In these proceedings
8. A. Fernández, D. Barberis, J. Sánchez, C. García, S. González de la Hoz, *Distributed Data Collection for the Next Generation ATLAS EventIndex Project*. *EPJ Web of Conference*, **214**, (2019) 04010
9. J. Salt, S. González de la Hoz, M. Villaplana, S. Campos, *Comparison of different ML methods applied to the classification of events with $t\bar{t}$ in the final state at the ATLAS experiment*. In Proceedings of Connecting The Dots International Conference, (CTD/WIT 2019, Valencia)
10. D. Dua, C. Graff. UCI Machine Learning Repository [<http://archive.ics.uci.edu/ml>]. Irvine, CA: University of California, School of Information and Computer Science (2019)
11. ARTificial Environment for ML and Innovation in Scientific Advanced Computing: <http://artemisa.ific.uv.es/web/>
12. T. Wenaus, D. Costanzo, A. Di Girolamo, *How HL-LHC challenges inform workload management R&D: An ATLAS view*, (WLCG&HSF Workshop, Naples, 2018)
13. A. Perez-Calero Yzquierdo, *CMS Strategy for HPC resource exploitation*. In these proceedings
14. J. Flix, *Exploiting network restricted compute resources with HTCondor: a CMS experiment experience*. In these proceedings
15. A. Delgado Peris, *Lighweight site federation for CMS support*. In these proceedings