EPJ Web of Conferences **295**, 07043 (2024)
CHEP 2023
https://doi.org/10.1051/epjconf/202429507043

The Creation and Evolution of the US ATLAS Shared Analysis Facilities

Ofer Rind[1,*] *Douglas* Benjamin[1], *Lincoln* Bryant[2], *Costin* Caramarcu[1], *Robert* Gardner[2], *Farnaz* Golnaraghi[2], *Christopher* Hollowell[1], *Fengping* Hu[2], *David* Jordan[2], *Judith* Stephen[2], *Ilija* Vukotic[2], and *Wei* Yang[3] on behalf of the ATLAS Computing Activity

[1]Brookhaven National Laboratory, Upton, NY, USA
[2]Enrico Fermi Institute, University of Chicago, Chicago, IL, USA
[3]SLAC National Laboratory, Palo Alto, CA, USA

Abstract. Prior to the start of the LHC Run 3, the US ATLAS Software and Computing operations program established three shared Tier 3 Analysis Facilities (AFs). The newest AF was established at the University of Chicago in the past year, joining the existing AFs at Brookhaven National Lab and SLAC National Accelerator Lab. In this paper, we will describe both the common and unique aspects of these three AFs, and the resulting distributed facility from the user's perspective, including how we monitor and measure the AFs. The common elements include enabling easy access via Federated ID, file sharing via EOS, provisioning of similar Jupyter environments using common Jupyter kernels and containerization, and efforts to centralize documentation and user support channels. The unique components we will cover are driven in turn by the requirements, expertise and resources at each individual site. Finally, we will highlight how the US AFs are collaborating with other ATLAS and LHC wide (IRIS-HEP and HSF) user analysis support activities, evaluating tools like ServiceX and Coffea-Casa.

1 Overview of Analysis Facilities

The US ATLAS operations program provides three distinct shared Tier3 analysis facilities (or AFs), which play a pivotal role in bridging the gap between grid jobs and localized interactive analysis on personal computers. The overarching purpose of these facilities is to furnish physicists with resources that expedite the extraction of insights from the increasing volumes of data generated by the ATLAS detector [1] at the CERN LHC. By granting access to interactive SSH logins, local batch processing, storage provisions, GPU resources, and other advanced computational services, the AFs empower physicists to navigate complex data sets more efficiently.

The trio of ATLAS analysis facilities in the US comprises the Brookhaven National Laboratory AF (BNL AF), the University of Chicago AF (UC AF), and the Stanford Linear Accelerator Center AF (SLAC AF). These facilities offer a diverse array of hardware resources, and take varying approaches in their architectural designs. In order to combine them into a

*e-mail: rind@bnl.gov

©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/).

useful complex of distributed analysis facilities, a number of essential considerations had to be taken into account, including cross-facility access, user file sharing and common environments. Details of the individual facilities are provided in the following sections, highlighting the solutions used to address these distributed design considerations.

2 Brookhaven National Lab Analysis Facility

The BNL AF is co-located with the US ATLAS Tier-1 and was designed from the bottom up to leverage the local resources and expertise at that facility. BNL "Tier-3" users can access interactive login nodes through an SSH gateway and submit jobs from there to a shared HTCondor pool providing $\sim 2\,000$ dedicated cores with opportunistic access to up to $\sim 40\,000$ cores. A Jupyterhub portal provides access via web browser, through a custom web proxy frontend, that handles user authentication via Keycloak, and uses HTCondor to spawn notebook servers across a dedicated set of virtual machines. The custom interface allows users to select from a set of standard environments, or to upload and instantiate their own custom containers. The virtual machines allow standard batch access to the shared HTCondor pool, as well as access via Dask-jobqueue. Each user is provided a home area with 500 GB of GPFS storage along with a large scratch area on Lustre storage for storage of data files. Over the past year, 106 unique users have logged in to the BNL AF and run over 1.6M batch jobs consuming 1.3M CPU hours. Over that same period, 29 have accessed the AF through the Jupyter portal.

2.1 Federated Login

One challenge in developing an analysis facility at a DOE national laboratory is providing effective access to interactive resources for collaboration members without guest appointments. This issue has been addressed at BNL by using federated login for access to restricted accounts. The frontend UI has been approved and instrumented to allow multi-factor login using CERN, SLAC and FNAL credentials. New user accounts are requested via a simple form that requires the user to provide an ORCID, and activated after a manual verification process that typically takes less than 1-2 business days. The user's EPPN OIDC token attribute is used to tie their federated identity to a local LDAP account. This account is restricted: it can be used on the Jupyterhub portal, but does not allow interactive facility access through the aforementioned SSH gateways.

2.2 File Sharing

It's a common observation that users will choose to work where they have access to their data. As noted above, an essential consideration in designing a complex of distributed analysis facilities is enabling distributed access to user files, both small files, such as Jupyter notebooooks, and large ones, such as group datasets. At BNL, users have long had access to Rucio-managed storage on ATLAS LOCALGROUPDISK, while data stored on CERN EOS is also accessible directly or via local XCache. For smaller files, BNL (and UChicago as well), after consultation with CERN IT, has now enabled read/write access on interactive hosts to CERN EOS home directories through fuse mounts authenticated via CERN kerberos. In addition, a Globus Connect Server (GCS) endpoint has been provided to support workflows utilizing HPC resources at BNL (and SLAC as well).

2.3 Common Environments

Despite differing facility designs, the desire to provide a common user environment across all the US Analysis Facilities has been an important consideration from the outset. Prior to the startup of the UC AF, the initial solution was to deploy the BNL and SLAC Jupyterhub virtual environments from a shared volume on the CVMFS Stratum-0 server operated at BNL. With the expansion of the facility and its usage, this approach has shifted to developing a set of common containerized (Apptainer) environments that provide specific ATLAS workflow-based package requirements, are validated to ensure compatibility with the different facility infrastructures, and afford a mechanism for specific facility integrations as well as user runtime customizations. A first set of these containers is under development for deployment in the coming year.

3 UChicago Analysis Facility

In terms of user quotas, all three AFs offer similar allotments: a 100 GB capacity for home and data storage, with data storage spanning the expansive 10 TB range. Despite the similarities in raw hardware resources among the facilities, the UC AF (Figure 1.) distinguishes itself by embracing a forward-looking approach. This approach rests upon a Kubernetes-based foundation, enabling the integration of cutting-edge services such as ServiceX [2] and Coffea-Casa [3], alongside other container-based applications.

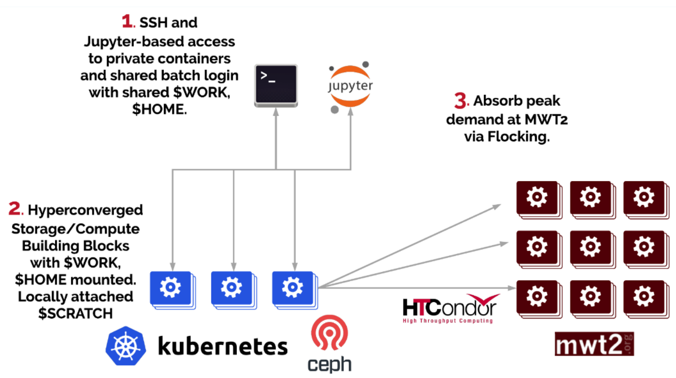

Figure 1: An overview of the Analysis Facility co-located with the Midwest Tier2 Center (MWT2) at the University of Chicago

The UC AF is firmly rooted in Kubernetes [4] architecture. All services, including computational tasks, storage management, interactive Jupyter environments, machine learning services, and Coffea-casa operations, are orchestrated within the Kubernetes framework. Notably, the UC AF has pioneered the integration of a production-grade HTCondor batch system into Kubernetes – an innovation that affords streamlined management of over 3 000 cores, 3.6 PB of storage, and diverse GPU configurations. Interconnectivity within the UC AF infrastructure is facilitated via 25G connections for high-speed compute nodes and the XCache [5] node, whereas other nodes, encompassing GPUs and login nodes, operate on 10G connections.

3.1 Project Evolution

UC AF's evolution unfolded across three implementation phases. The initial phase, executed in 2021, focused on establishing foundational services such as SSH login access, the /home NFS file system, and the HTCondor batch system. Phase 2 concentrated on enabling Jupyter Notebook environments and the /data file system. Notably, a homegrown web application supported the Jupyter service, while the /data distributed file system was empowered by the

Rook-Ceph storage operator. In the third phase, advanced services, including Coffea-casa, machine learning tools like Triton Inference Server and MLflow, and the Harbor container image caching service, were integrated. Additionally, optimization measures were applied to the batch system, including the segregation of HTCondor job queues into short and long queues for improved efficiency.

The hardware ecosystem of the AF comprises hyperconverged nodes, fast compute nodes, GPU nodes, login nodes, an XCache node, an NFS server, and a head node. Hyperconverged nodes, accounting for 19 units, incorporate AMD CPUs and 12×16 TB spinning disks. Intel processors and 10×3.2 TB NVMe drives equip the 16 fast compute nodes, while 8 interactive nodes, of which 2 are specific to the UC AF, offer accessibility. The XCache node, pivotal for accelerating remote data retrieval, is equipped with 2×25 Gbps connections and 24×1.5 TB NVMe drives. Diverse GPU nodes, including those with Nvidia A100, V100, 2080, and 1080 GPUs, enhance computational capabilities.

3.2 Batch, Data Access and Storage

UC AF's hallmark is the fusion of conventional high-throughput computing (HTC) with cloud-native technology through the deployment of HTCondor as a Kubernetes-based service. Here we have opted for software containers and cloud-native [6] application management methods (Helm charts, GitOps and continuous delivery processes [7]) to standardize deployments. This capitalizes on the advantages of both paradigms, minimizing provisioning efforts while optimizing resource utilization. Horizontal Pod Autoscaling (HPA) augments resource management by dynamically adjusting resource allocation, while real-time monitoring via Prometheus and HTCondor queue status monitoring through Elasticsearch and Kibana ensures efficient operations. To address diverse user needs, UC AF has implemented a dual queue system, bifurcating tasks into short and long queues. The short queue, with a 4-hour wall-time limit, prioritizes rapid job execution, catering to interactive requirements. The queue monitoring interface affords users insights into queue utilization and workload trends, facilitating task planning. UC AF offers an array of data access services, with storage provisions at the forefront. A rapid solid-state /home NFS file system aids software development, offering users 100 GB per account. The distributed /data file system, underpinned by the Rook (distributed Ceph object store [8]) Kubernetes operator, spans 400 disks across 35 servers, delivering up to 5 Gbps read/write performance. An upcoming addition, a high-speed all-flash disk pool, will supplement the existing /data file system. Furthermore, per-node /scratch file systems, each furnished with multiple terabytes of solid-state drives, complement data storage strategies.

3.3 Advanced Services

UC AF extends its services to encompass advanced analytics capabilities. **XCache.** The XCache server expedites remote data access through optimized XRootD-based configurations. Additionally, a container image caching service accelerates image pulls by locally caching remote images, ensuring compatibility and transparency while enhancing service efficiency. **Notebook Access.** UC AF's JupyterLab-based service empowers users with customizable resource requests and a selection of predefined container images. These images are equipped with essential tools such as NumPy, pandas, scikit-learn, and matplotlib. Users can tailor the number of cores, memory, and GPU resources according to their analysis needs. The platform features multiple image choices, including those incorporating TensorFlow, PyTorch, and Dask. **Columnar Analysis and Data Delivery**. Coffea-casa, a Dask and JupyterLab-powered framework, supports low-latency columnar analysis. It capitalizes on the Coffea framework, facilitating vectorized event analysis in parallel, boosting efficiency. Moreover, the ServiceX data delivery service streamlines data access, allowing users to specify datasets and criteria for extraction. ServiceX employs parallel workers to extract and store data in columnar formats, ensuring compatibility with existing analysis tools. **Machine Learning Services.** UC AF extends its capabilities to encompass machine learning services. The Nvidia Triton Inference Server, equipped with GPU nodes, scales based on inference request queue times. Machine learning models reside in a Ceph S3 store, with user access credentials provided. MLflow, an open-source ML lifecycle management platform, is also available, enabling experimentation, reproducibility, deployment, and model registry management. External access requires institutional identity authentication, and model registry data is stored in Ceph S3. **Server-Side Data Delivery.** In addition to the AF instance, we are also exploring running a ServiceX instance at CERN to leverage the FABRIC Across Borders (FAB), an extension of the FABRIC testbed [9] connecting the core North America infrastructure to international sites. We call this *server side data delivery*. This idea is to co-locate the data transformation service to the data origin and transfer the filtered data via a network path that's managed and high performance.

3.4 User Growth and Achievements

UC AF's user base has witnessed substantial growth, surging from 25 users at launch in July 2021 to over 250 at present. In the past 90 days alone, the facility has serviced 77 Jupyter users, 105 SSH users, and 48 batch users, culminating in over 900 000 completed jobs and nearly 700 000 CPU hours. This trajectory underscores UC AF's success in meeting the dynamic demands of the ATLAS analysis community.

4 SLAC Analysis Facility

Like the other AFs previously mentioned, the SLAC AF offers 100 GB of backed-up home space and a maximum of 10 TB for private data storage. The SLAC AF is not an independent facility; instead, it is integrated into the broader lab-wide SLAC Data Facility (SDF) and is slated for relocation to the second phase of SDF, known as the SLAC Shared Scientific Data Facility (S3DF).

4.1 SLAC AF Resource

SLAC AF started in 2017 with hardware inherited from the former US ATLAS Western Tier 2. The current ATLAS-owned hardware resources in service include:

- Four CPU nodes, each has 128 AMD EPYC 7702 (Rome) cores in two CPU sockets, 512GB RAM, a 10Gbps NIC and a Mellanox Infiniband controller, as well as 500GB local scratch.
- 830TB usable space in a pool of 17PB Lustre file system.
- One data transfer node (64-core, 128GB RAM) with both 100Gbps NIC and a Mellanox infiniband controller
- One GPU node with 4× Nivida A100 GPUs connected by NVlink, 64 CPU cores and 1TB RAM. It also has a 10Gbps NIC and an Infiniband controller.

In addition to ATLAS-owned resources, SLAC AF users have opportunistic access to other CPU and GPU resources available at SDF/S3DF.

The SLAC S3DF will deploy AMD Milano-based CPU nodes, each equipped with a 100Gbps NIC and a much larger local scratch space, but without an Infiniband controller.

The current Lustre file system will not be expanded. S3DF is building a Weka file system based storage instead. This Weka storage has a large flash storage frontend, tiered with a Ceph HDD storage backend. Users will only see a POSIX file system (the tiering is invisible to users). SLAC AF plans to eventually migrate to use Weka and retire Lustre.

4.2 SLAC AF Software Environments

SDF currently operates on CentOS 7, while S3DF will use EL 8. All batch nodes are equipped with CVMFS for accessing ATLAS software and the Grid environment.

The batch system is administered using Slurm. Within SDF, each Slurm account is assigned a set of CPU and GPU resources associated with the account. Users have the choice between a "usatlas" account or a "shared" account. Jobs executed under the "shared" account can be preempted. In S3DF, resource allocations are determined based on the purchased resource for each Slurm account, eliminating the need for users to switch accounts, although excessive resource use will be subject to preemption.

SDF/S3DF uses OpenOnDemand [10] to offer a web interface for the widely used Jupyter service. SLAC AF and other SLAC communities offer various preconfigured Jupyter environments. Users can also initiate their own Jupyter sessions from the OpenOnDemand web interface, including those within Singularity containers or Conda environments. These Jupyter sessions are executed as Slurm jobs on batch nodes. Examples are provided to help users create custom Jupyter environments capable of submitting Slurm jobs or running Dask on the Slurm farm from within Jupyter.

SLAC AF also supplies multiple Rucio Storage Elements, enabling users to transfer ATLAS data products and access them from batch jobs. Additionally, an XCache service is available.

Although the SLAC AF does not currently possess ATLAS-owned Kubernetes resources, S3DF hosts an extensive pooled Kubernetes resource for all SLAC experiments. The SLAC AF will have the option to access these resources.

4.3 SLAC AF Users

Since the inception of the AF service at SLAC, the facility has logged over 260 users. The composition of active users varies over time due to factors such as student graduations; however, it consistently shifts from a user base of 30–60 individuals. Notably, in recent years, SLAC AF has observed a higher number of Jupyter users compared to batch users, despite substantial CPU and GPU hours being utilized by batch users.

S3DF is currently in the process of implementing a Federated Access model. This approach will enable users to authenticate using their home institutional accounts, alleviating the need to remember yet another password. It's important to note that the user vetting process, in compliance with DOE regulations, and the cybersecurity training requirements will remain unchanged.

5 Overall Usage

We have endeavored to compile usage metrics encompassing the three analysis facilities, with a particular focus on the mode of access, see Figure 2. This entails discerning whether users employ direct logins through the secure shell or engage with Jupyter notebooks. Additionally, we are keenly interested in tracking the number of distinct users who interact with the batch systems.

We have observed a noteworthy presence of both conventional batch methods, characterized by HTCondor submissions, and the utilization of Jupyter notebook instances across all three facilities. Over the specified timeframe, we have recorded 146 unique Jupyter users, 177 individuals accessing our batch systems, and 297 users employing secure shell access. Collectively, these users have contributed to a substantial workload, resulting in the completion of 3 056 009 jobs and the consumption of 2 790 320 CPU hours.

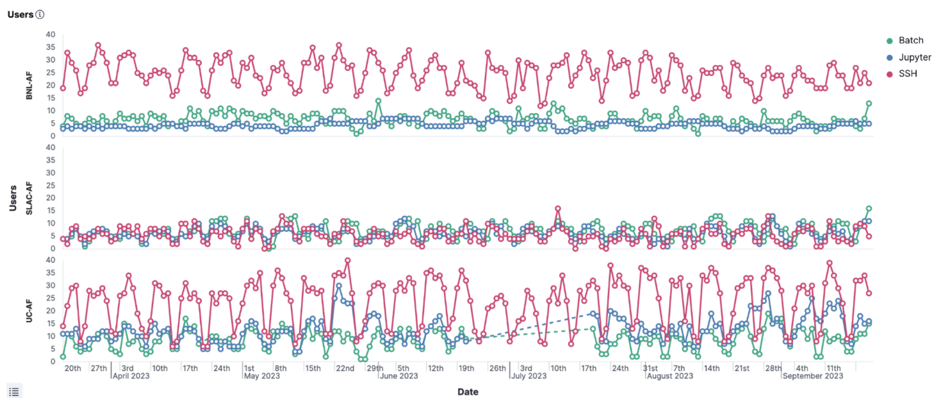

Figure 2: Overall usage metrics for the three US ATLAS analysis facilities over a 6 month period. Shown are three primary access methods: SSH-login, batch submission, and Jupyter notebooks launched.

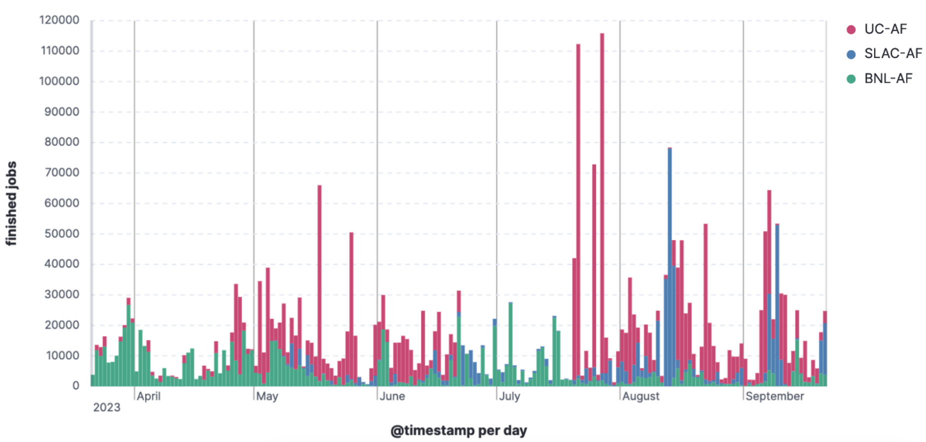

Figure 3: Overall jobs completed for the three US ATLAS analysis facilities over the past six months.

We also collected statistics of job metrics from the three analysis facilities, and have aggregated these for relative comparison as well, see Figure 3. What stands out is the discernible pattern of fluctuating demands within analysis facilities concerning job execution. These fluctuations are marked by periods of heightened demand, which often align with the schedules of individual users and analysis groups. These peaks in demand can be attributed, in part, to factors such as conference deadlines or the submission of theses for analysis.

With a growing user base, there is an increased need to provide coherent and effective user support, which can be a challenge across a diverse set of facilities. To that end, the facilities have worked together to provide a unified user's guide on a single documentation platform [11]. A shared Discourse server [12] has also been deployed at BNL to provide cross-facility user support from the facilities and the analysis community alike.

6 Summary

The US ATLAS Analysis Facilities (AFs) have been instrumental in accelerating data analysis at the CERN LHC. Collaboration among these facilities has allowed for seamless cross-facility access, user file sharing, development of common containerized environments and deployment of unified user support platforms and documentation. Each facility has its strengths: Brookhaven leverages local resources efficiently, Chicago pioneers Kubernetes-based solutions, and SLAC integrates ATLAS-owned and shared resources. Usage metrics are collected on a common platform and demonstrate the importance of these AFs, with users employing various access methods. Infrastructure development at these sites continues with a focus on portability of the complex user analysis workflows expected at HL-LHC.

This work was supported in part by the National Science Foundation awards PHY-2120747, OAC-2115148, OAC-2029176, OAC-1836650, and OAC-1724821.

References

[1] ATLAS Collaboration, JINST **3**, S08003 (2008)
[2] *Towards Real-World Applications of ServiceX, an Analysis Data Transformation System*, Vol. 251 (2021), 2107.01789, https://doi.org/10.1051/epjconf/202125102053
[3] *Coffea-casa: an analysis facility prototype*, Vol. 251 (2021), 2103.01871, https://doi.org/10.1051/epjconf/202125102061
[4] B. Burns, B. Grant, D. Oppenheimer, E. Brewer, J. Wilkes, ACM Queue **14**, 70 (2016)
[5] *Xcache documentation* (2018), http://slateci.io/XCache/
[6] *Cloud Native Computing Foundation*, https://www.cncf.io/
[7] *Flux application deployment*, https://fluxcd.io/
[8] S.A. Weil, S.A. Brandt, E.L. Miller, D.D.E. Long, C. Maltzahn, *Ceph: A Scalable, High-performance Distributed File System*, in *Proceedings of the 7th Symposium on Operating Systems Design and Implementation* (USENIX Association, Berkeley, CA, USA, 2006), OSDI '06, pp. 307–320, ISBN 1-931971-47-1, http://dl.acm.org/citation.cfm?id=1298455.1298485
[9] *FABRIC Testbed*, https://fabric-testbed.net/
[10] *Open On Demand*, https://openondemand.org
[11] *US ATLAS Analysis Facilities*, https://usatlas.readthedocs.io/projects/af-docs/en/latest/
[12] *ATLAS Talk*, https://atlas-talk.sdcc.bnl.gov/