




The Creation and Evolution of the US ATLAS Shared Analysis Facilities

Ofer Rind on behalf of the ATLAS Collaboration

International Conference on Computing in High Energy Physics
May 8th, 2023



Analysis Facilities

Broad [Snowmass Report](#) definition:

“The infrastructure and services that provide integrated data, software and computational resources to execute one or more elements of an analysis workflow. These resources are shared among members of a virtual organization and supported by that organization.”

In the HL-LHC era, scale becomes an issue, impacting the line between interactive/batch analysis and access to datasets. Increasingly complex workflows and heterogeneous architectures will also play a role.

“What elevates a resource to the level of an analysis facility is official support as a shared resource within an organization of people with shared interests”

Overview of US Analysis Facilities

US ATLAS has three shared Tier 3 analysis facilities providing software & computing

- Resources that fill gaps between grid jobs and interactive analysis on local computers
- All leveraging substantial local “Tier 3” batch and storage resources
- Interactive ssh login, local batch, non-grid storage, LOCALGROUPDISK, PanDA
- GPU resources




BNL Facility

~2000 cores, part of a larger shared pool,
opportunistic access up to 40k cores

User quota: 500GB GPFS plus 10TB Lustre

~200 users



SLAC National Lab Shared
Scientific Data Facility (SDF)

SLAC Facility

~1200 cores, part of larger shared pool,
opportunistic access up to 15k cores

User quota: 100GB home, 2-10TB for data

~100 users



ANALYSIS
FACILITY
at
UCHICAGO

UChicago Facility

~3000 cores, co-located with MWT2,
opportunistic access up to 16k cores

User quota: 100GB home, 10TB for data

~210 users

Launched Oct 2021

Access via Federated Login

UChicago: leveraging CIConnect and InCommon federation via in-house user management system; ATLAS IAM for Coffea Casa service



Log in to use af-portal

Use your existing organizational login

e.g., university, national lab, facility, project

CERN

By selecting Continue, you agree to Globus terms of service and privacy policy.

Continue



Globus uses CILogon to enable you to Log In from this organization. By clicking Continue, you agree to the CILogon privacy policy and you agree to share your username, email address, and affiliation with CILogon and Globus. You also agree for CILogon to issue a certificate that allows Globus to act on your behalf.

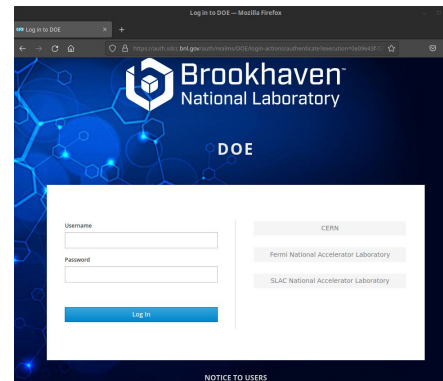
OR



Didn't find your organization? Then use Globus ID to sign in. (What's this?)

BNL: Federated account

- Created via CERN/SLAC/FNAL MFA login, satisfying DOE requirements (No SSH access)
- Users must fill out a form, requiring ORCID; approval after confirmation of ATLAS affiliation
- Account activation turn-around time goal of 1-2 business days
- EPPN OIDC token attribute mapping used in LDAP for authorization and to tie federated ID to a local UNIX account



SLAC: In development (based on NERSC IRIS system)

File Sharing

- Local storage for each user in addition to LOCALGROUPDISK
- CERN EOS R/W access enabled on interactive hosts at BNL and UChicago via fuse mount with auth via CERN kerberos
- EOS data also accessible directly and via local xcache at each site
- Globus GCS endpoint supporting HPC resources at BNL and SLAC, and in a pre-production containerized service at UChicago

Shared Environments and Containerization

Motivated by the desire to have a common user environment across all three US Analysis Facilities – seeking a uniformity of user experience despite slightly different facility designs

Initial solution (BNL, SLAC): Jupyterhub virtual environment from shared volume on CVMFS Stratum-0 operated at BNL

Current Effort: Common containerized (Apptainer) environments

- Base layer for specific facility integrations (e.g. Jupyterhub, HTCondor/Slurm, bindmounts)
- Workflow-based package requirements (e.g. tensorflow for ML)
- Mechanism for user customizations at build or runtime
- Integration with ATLAS container build/validation infrastructure

Analysis preservation is a side benefit

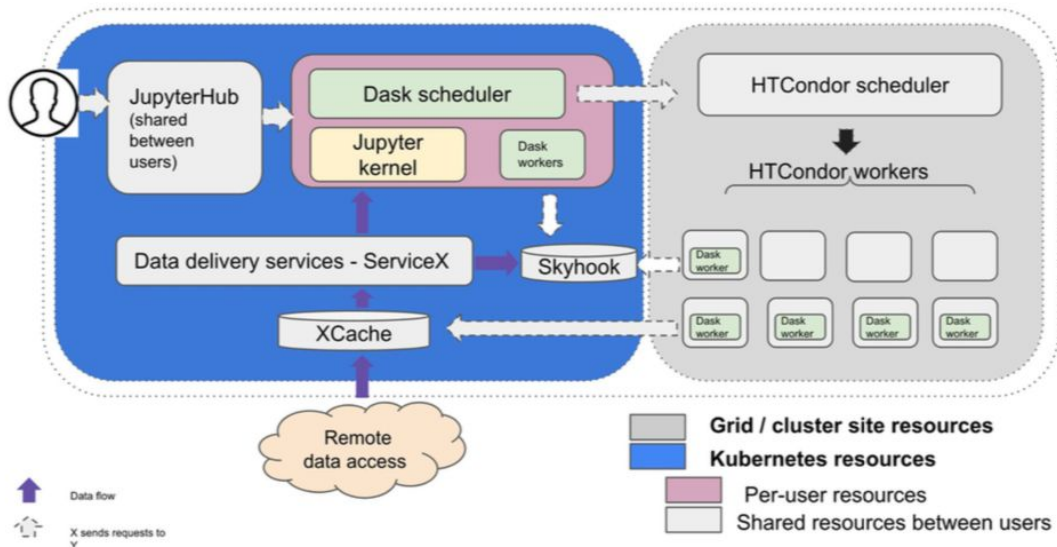
US Analysis Facilities Accounting

Leveraging ATLAS analytics infrastructure for centralized and unified collection of user batch and interactive usage trends across all three facilities.



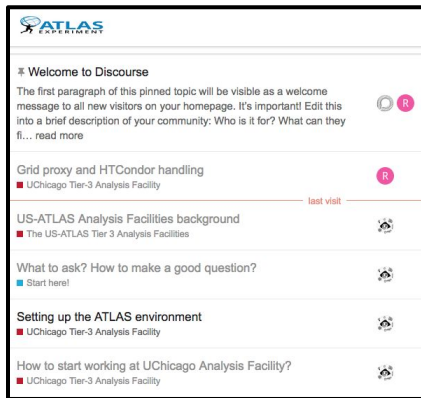
Platforms for Modern Analysis Tools

- Participating in Analysis Grand Challenge with IRIS-HEP
- Coffea-Casa analysis service at UChicago
 - FluxCD deployment; images from OSG Harbor
 - HTCondor on Kubernetes
- Dask scale-out to local HTCondor and Slurm pools under development
- ServiceX instances at UChicago for columnar data analysis (ATLAS xAOD and Uproot)
- BNL testing instances of ServiceX and REANA on OKD cluster
- GPU resources at all sites

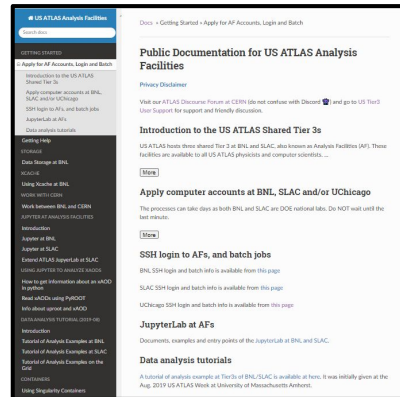


AF User Support, Documentation & Training

- Focused effort on improving user support, documentation & training
 - <https://usatlas.readthedocs.io/projects/af-docs/en/latest/>
- Increasing support for users (onboarding, documentation, container maintenance,...)



- Employed new tools ([Discourse](#), [Mattermost](#)) for user support to foster and build an active, engaged and inclusive online user community



Plans

- Developing a set of uniform containerized work environments with facility specific integrations and mechanism for user customizations
- Demonstrate complex, portable distributed workflows implementing IRIS-HEP tools at all sites
 - [Analysis Grand Challenge](#)
- Exploring integration with CERNBox and ScienceMesh for sharing user working directories
 - Storage token implementation
- Expanded usage metrics and accounting

Thank you

Douglas Benjamin, Costin Caramarcu, Christopher Hollowell, Thomas Smith, Patrick Pascual, Wei Yang, Lincoln Bryant, Robert Gardner, Farnaz Golnaraghi, Fengping Hu, David Jordan, Judith Stephen, Andrew Taylor, Suchandra Thapa, Ilija Vukotic