ATLAS view on Analysis Facilities

IRIS-HEP SB, 16 January 2024

Zach Marshall (LBNL), David South (DESY)

Introduction

- ATLAS was asked to contribute to establishing a “high level strategy” within IRIS-HEP, including where we see opportunity for “larger community participation”
  - Well, first of all we’d better get our own strategy clear

- But do we really need to define this now? And if we do, do we risk being too rigid?
  - The key question: Do we know enough do do this?
  - Note that this is not the same as, “*Do we presume to know what we need?”*

- Simply put, our strategy should be:
  - *To ensure our (analysis) resources evolve in a way that is aligned with the requirements of the experiments (users)*
  - That’s the most important aspect, otherwise we end up developing something hypothetical that may not being used in the end — which is a road oft travelled
  - Getting input from the users is critical

So, can we try to define it?

- ATLAS definition at vCHEP:

  "Resources that can provide such integrated solutions [featuring not only reliable batch systems but also the latest, relevant interactive tools] are referred to as 'analysis facilities'. They are defined more by the set of applications they offer, rather than the resources on which they run. The ideal analysis facility is one or more dedicated resources with user support and federated access for all ATLAS users."

- More recently, the HSF white paper in preparation, defines an AF:

  "We loosely define an Analysis Facility as 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"

Analysis Facilities and the HSF

- There is a nice HSF talk from the LHCC focus session last June, which has many great ideas and descriptions of work going on in this field
  - Two slides from that talk:

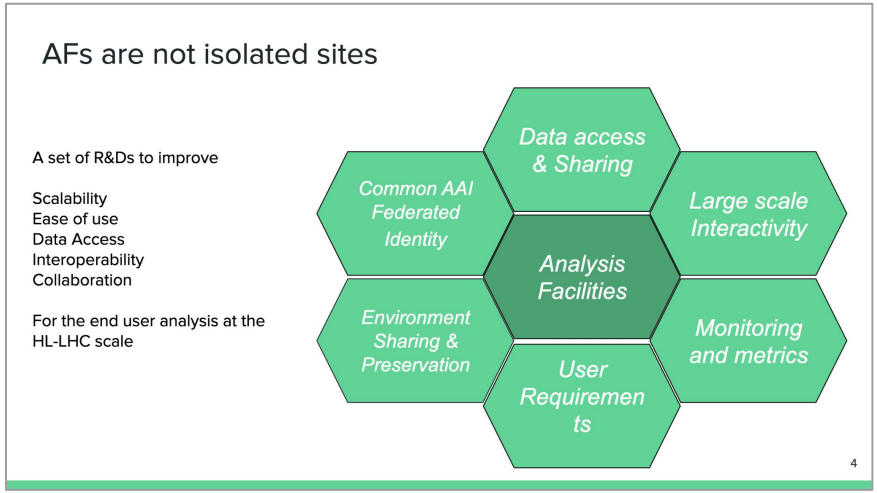

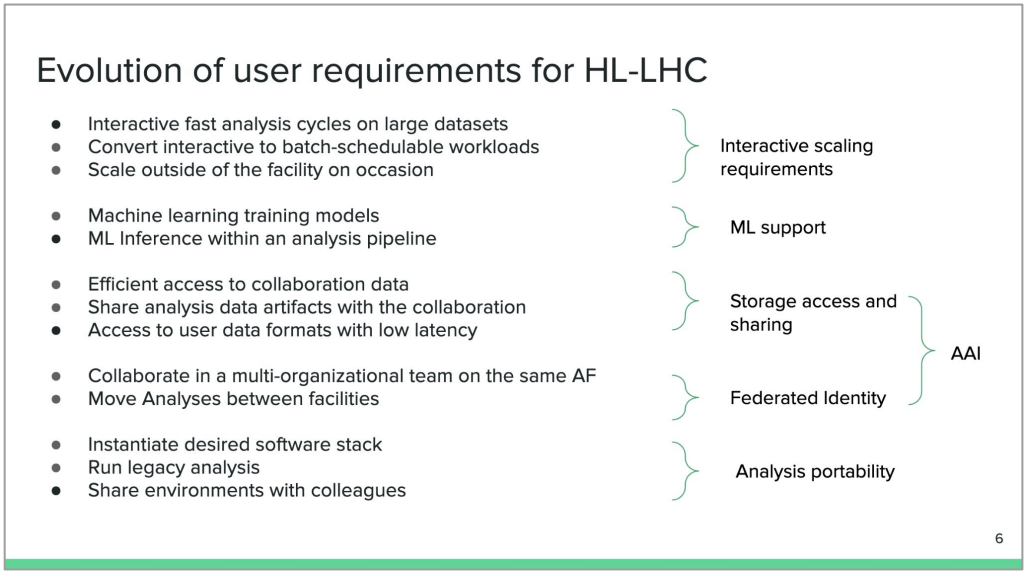

- These requirements are of a general nature, starting points for the experiments
  - ATLAS is examining them already without the need to concretely define an analysis facility
  - In fact the outcome of any associated R&D may help to form a strategy in the future

ATLAS Actionables

- The ATLAS approach is to perform R&D projects centered on ideas relevant to AFs, in order to scope the required strategy and program of work

- Several areas where this approach is applicable - more details here
  - Work on scaling up GPUs and dask integration, offloading to batch system, providing entry to various resources
  - Initiatives within DOMA:
    - Choice of shared global storage: posix or object store?
    - Input data organisation and access, expanding the use of caches
  - Common AAI, integration of tokens for easier grid integration as well as to ease cloud access

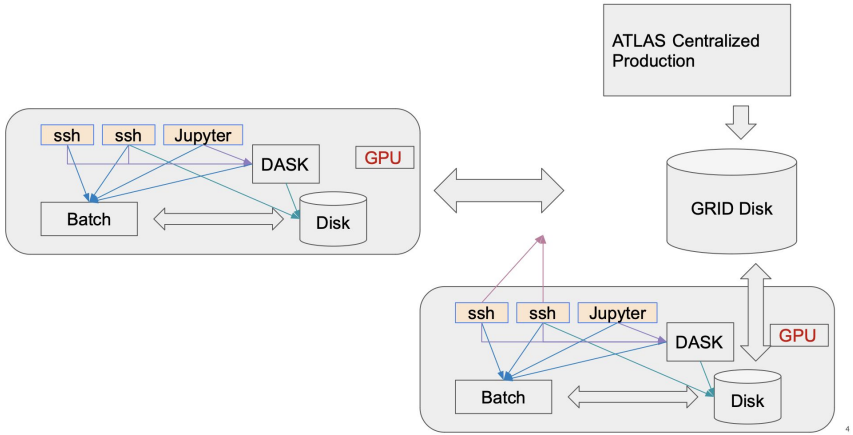

Connectivity between grid infrastructure and new technologies at various AFs

ATLAS Actionables

- The ATLAS approach is to perform R&D projects centered on ideas relevant to AFs, in order to scope the required strategy and program of work

- Several areas where this approach is applicable - more details here
  - Improve environment sharing and “analysis portability”
    - Distribution, provisioning of containerised workflows
  - The Run 4 analysis level format DAOD_PHYSLITE has a big role to play
    - Examining more centralised, automated production
  - Better understanding of the needs of users by the monitoring of current jobs, disk usage, practices, types of resources accessed and the associated trends

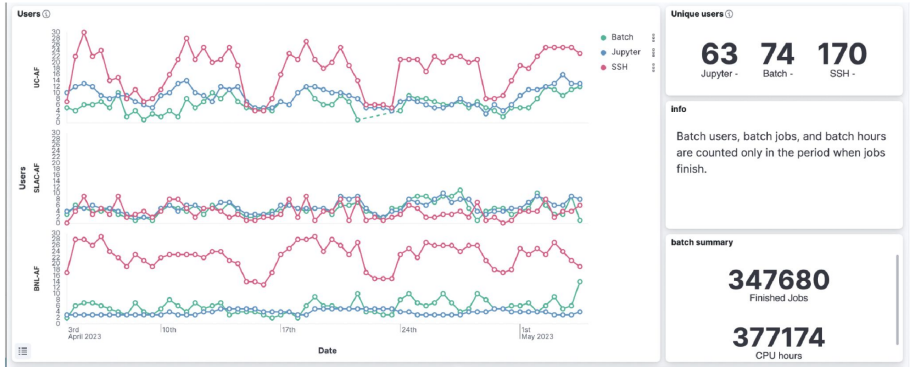

Monitoring of ssh, batch, Jupyter usage at US-ATLAS federated T3

LHCC now asking how we'll do analysis in five years

"The LHCC recommends that experiments engage in the process of developing and defining the structure of the future Analysis Facilities and requests they produce a document which defines the use cases in order to establish realistic benchmarks. This process should be coordinated with and SW review panel. The document is expected to be regularly updated in the process towards HL-LHC."

- Whilst ATLAS is engaging in this activity, it requires an understanding of the evolution of our physics programme, which is hard to accurately predict at this point
  - Many technical questions also remain, such as the impact of columnar analysis, scale of adoption of data formats, level of adoption of ML, …
- In order to build a the requested use cases requires close cooperation between the users and computing communities
  - In fact there needs to be a constant feedback loop between the two parties
- In late summer 2023 a survey was performed to help understand local and national analysis facility resources and their usage in ATLAS

Analysis Facilities: Key results from the survey

- There are a lot of resources in play worldwide
  - For ATLAS, national and local clusters make up ~50 PB of disk, 1M HS23, 100s of GPUs
- Lots of people still rely on lxplus for analysis
  - Understanding why is key to the success of an AF
  - Conduct survey to ask users what they like / don't like
- Supporting all users is key
  - Not just working with the super-experts, who often have dedicated, specific solutions, whilst most users still rely on the basics - large scale batch submission
    - Important to understand: Is this all that's needed, or a question of adequate support for new workflows..
- ATLAS Run 4 analysis model relies on DAOD_PHYSLITE
  - Lossy compression, efficient tools, columnar analysis are all important parts of this effort
  - Also need to understand what users need when PHYSLITE is not suitable (augmentation, alternatives, etc)
- Should ensure that AFs support the full range of major use cases, not only those adopted by a few users

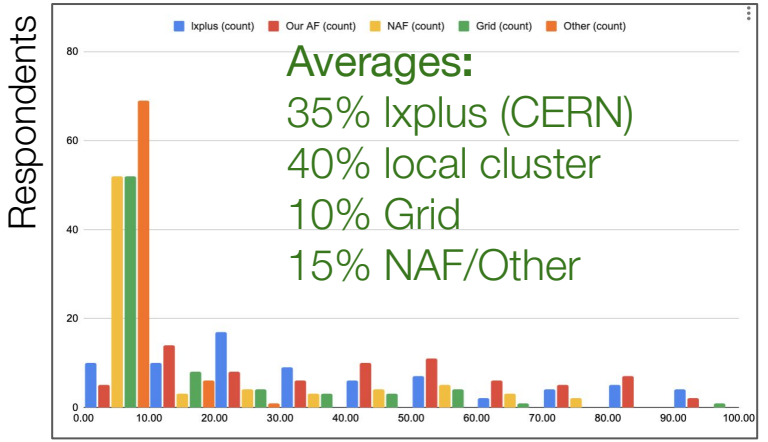

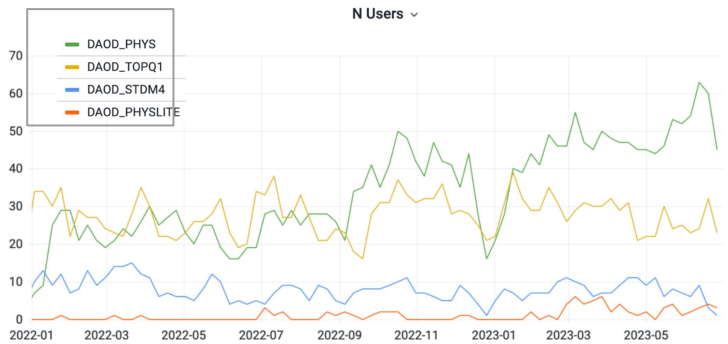

No. of user Grid jobs accessing various ATLAS derivation formats

Conclusions - AF design and implemention

- Back to the strategy: *To ensure our (analysis) resources evolve in such a way that is aligned with the requirements of the experiments (users)*
  - Simply: A large Tier 3-like facility that supports all ATLAS users and allows both interactive and batch jobs
- Is a universal AF definition possible? Is it necessary that AFs all provide the same features?
  - Probably can't enforce a universal AF design, but we can identify certain features that may be desirable
  - Each AF does not not necessarily need things like GPUs, ML tools, Jupyter…
  - Hardware resources are expensive: better if we don't *require* everyone to provide everything
- Make distinction between identical (they don't need to be) and coherent development/access
  - For example, how global storage works, how entry/authorisation is implemented - *the ATLAS Actionables*
  - i.e. AFs all looks the same to the users, but some may provision certain resources while others do not
- AFs have existed for a while, but have been established rather isolated from each other
  - Should be brought under one roof, developed with a common set of goals
  - Look at central documentation, recommendations to sites, a common support team..
  - When AFs co-located with grid resources, can maximise support/sustainability, to some extent flexibility

Conclusions - How can we collaborate

- Establish a feedback channel between the user and computing communities, to allow us to understand what's really needed
  - We need to define what the users need, but they also need to tell us what they need
  - Cautious about adopting super-user / edge case solutions: what do *all* of our users need?

- Employ use-case driven R&D to focus the definition and contribute to the overall strategy
  - Concentrate strategic core infrastructure evolution like AAI, evaluation of object stores, containers, caches and integration of AFs with the grid particularly to facilitate storage access

- Decrease the isolation between existing, universal access AFs such as those at CERN (lxplus), SLAC, UChicago, BNL, as well as the more limited access variants at DESY (NAF), Valencia…
  - Aligning the current implementations can also have a positive influence on the future AF development and design