# HSF AF White Paper Overview

**Oksana Shadura (University Nebraska–Lincoln)**

**The 8th Asian Tier Center Forum**
https://indico.cern.ch/event/1411901

# HSF AF Forum

The Analysis Facilities (AF) Forum provides a community platform for those interested in contributing to the development of analysis facilities for use within HEP experiments (e.g. inter experiment)

- ***Goal:*** the contribution of **ideas from potential end users** for functionality to support the analysis of HEP data, **specification and planning of the facilities** themselves, and **technical developments** needed to implement AFs
- HSF AF White Paper as one of deliverables
  - Each section has a set of open questions (check backup slides)

# HSF AF Forum

- *25 March 2022* - first kick-off meeting of the HSF "Analysis Facilities Forum" organised in collaboration with IRIS-HEP
- Multiple interesting contributions over last years: Indico category
  - Talks from facilities representative, analysis tools and services developers
- The AF Forum also collaborates with related HSF Working Groups, such as the Data Analysis Working Group, the PyHEP Working Group, and with WLCG/DOMA.

Conveners:

- Diego Ciangottini (INFN, Perugia U, CMS)
- Alessandra Forti (Manchester, WLCG/ATLAS)
- Lukas Heinrich (TUM, ATLAS)
- Nicole Skidmore (Manchester, LHCb)

# What is an Analysis Facility (AF)

- **Analysis Facility definition**: "*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.*"
- **What we consider as an Analysis Facility?**
  - CERN (all WLCG users), NAF (only German users), FNAL (CMS), BNL (ATLAS),...
    - But also T3s, grid, commercial clouds….

# Key areas to be addressed by HSF AF whitepaper

- The user perspective
- Compute resources: access and provisioning
- DOMA (Data Organisation, Management, Access)
- Federated Identity Management and Authentication, Authorization Infrastructure (AAI)
- Accelerator Resources
- Analysis portability and preservation
- Monitoring and Metrics
- End user documentation and support

# User requirements for HL-LHC AFs

Assumptions from HSF AF whitepaper

## Evolution of user requirements for HL-LHC

- Interactive fast analysis cycles on large datasets
- Convert interactive to batch-schedulable workloads
- Scale outside of the facility on occasion

→ Interactive scaling requirements

- Machine learning training models
- ML Inference within an analysis pipeline

→ ML support

- Efficient access to collaboration data
- Share analysis data artifacts with the collaboration
- Access to user data formats with low latency

→ Storage access and sharing

- Collaborate in a multi-organizational team on the same AF
- Move Analyses between facilities

→ Federated Identity

→ AAI

- Instantiate desired software stack
- Run legacy analysis
- Share environments with colleagues

→ Analysis portability

"Analysis Facilities Community R&D Status", LHCC report, A.Forti

# Computing resource provisioning

Assumptions  from HSF AF whitepaper

**Workflow management:** while switching to larger and more complex workflows, AF should offer the seamless transition between local “interactive” analysis development to multi-scale full production analysis.
One of the solution is to **offloading to external resources for “interactive” analysis workflows**

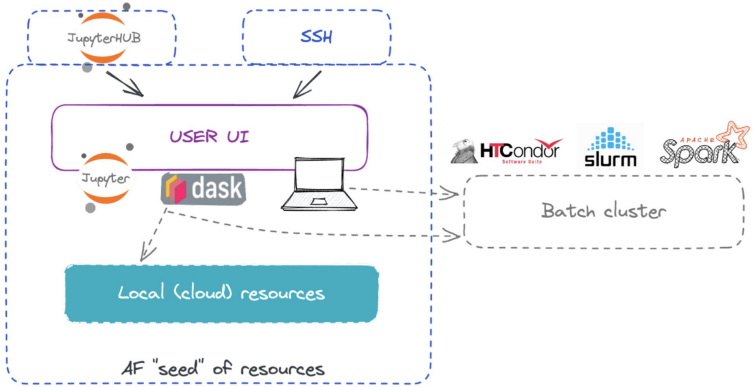


Examples: Coffea-casa AF, EAF, INFN AF, Purdue AF and many others

# Computing resource provisioning

Assumptions from HSF AF whitepaper

**UX/UI**: Access to AF resources could be provided through WebUI access interfaces, such as Jupyterhub, in addition to traditional SSH access

**AAI**: wider adoption tokens in HEP simplifies access to data from diverse resources and services ( e.g. including cloud services) including available facilities

**Containers**: containerised payloads will become more common approach at AF (e.g. a partial solution for the analysis portability and reproducibility)

# Computing resource provisioning

Assumptions from HSF AF whitepaper

**Computing resource provisioning responsiveness:** the possible solution is enable the dynamic provisioning resources for low latency interactive analysis alongside traditional batch systems

**Efficiency and performance**: analysis frameworks should be validated, targeting the efficiency of resource utilisation and overall physics throughput

**Site provisioning model**: an open question how AF could be seamlessly integrate HPC and commercial clouds at AFs and how to benefit the most from them

**Resource provisioning model**: an open question how to organise the server level resource provisioning, e.g. multiple servers vs “fat” node dedicated for fast interactive analysis turn-around

# Accelerator resource provisioning

Assumptions from HSF AF whitepaper

**Specialised hardware at AFs**: investigation of GPUs and possible FPGA access and what should be the resource provisioning patterns from the infrastructure perspective

**Provisioning complex ML workflows:** as a community, we need to understand how to handle complex use cases (e.g. how to integrate GPUs in interactive ML workflows or how organise the efficient multi-GPU sharing for large ML models or provide inference-as-a-service at AFs)

**GPU-on-demand:** the open question about how HEP community could benefit from commercial cloud GPU resources and how to integrate such type of resources in facility

# DOMA: Input data organisation and access

Assumptions from HSF AF whitepaper

- Data flows for HL-LHC analysis are **still being defined**
- **Latency** of intense workflows usually reduced by a fast local storage serving the interactive resources
- **Caches** gives a notable advantages
  - Automatically transfers only accessed data, it doesn't copy the entire datasets
  - Many analyses have a highly repetitive stage, particularly in the development stages when ideas are tested

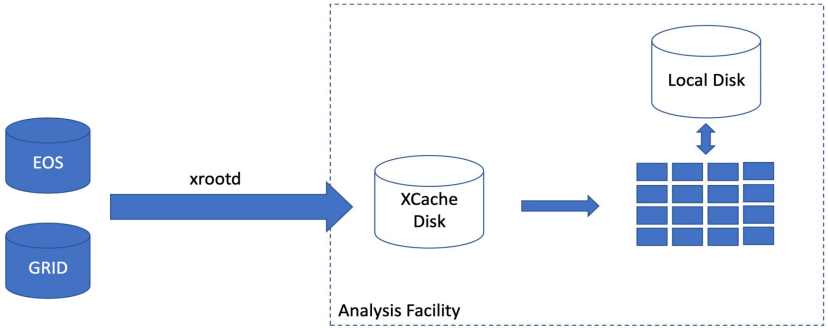

# DOMA: shared storage provisioning

Assumptions from HSF AF whitepaper

- *Integration data management tools* (e.g. RUCIO) in AF user interface
- *Federated storage* could also be used to share data between facilities
- *The open question about object stores availability at AF* and its comparison vs POSIX access
- *Data transformation services availability at AF:* data extraction and columnar data delivery services (e.g. ServiceX) could be crucial for complex user workflows (e.g. convert to ML data formats) and to enable data transformation on the fly for interactive analysis

# Analysis Portability & Preservation

Assumptions from HSF AF whitepaper

- *Users want to share with colleagues their setup, code, configuration, small amount of input data with in facility and other facilities*
- **CVMFS, conda environments and containers** are key solution
  - CVMFS distribution and containers can be easily combined
    - Currently /cvmfs/unpacked.cern.ch has ~3000 images
- **Containers** could help users to migrate between AFs and **are a key ingredient for interoperability**

# How do we know all this is useful, works and is it enough

Assumptions from HSF AF whitepaper

- Enhanced **monitoring, tracking, benchmarking**
- *Encourage advanced user analytics at AF*:
  - How to avoid “dark” analysis or when users go off to their local resources and we lose track of them
- *Each facility should provide user-friendly documentation* how to start to use it
- Benchmark AF setup with examples of typical analysis workflows with experiment-specific or Open Data datasets  (e.g. Analysis Grand Challenge)

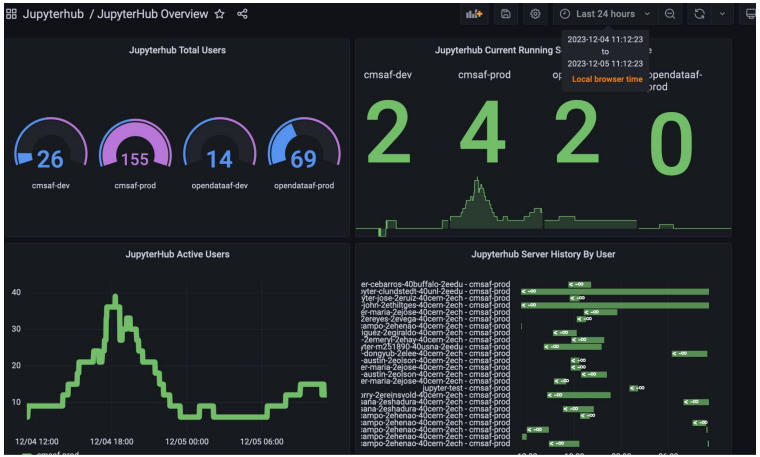


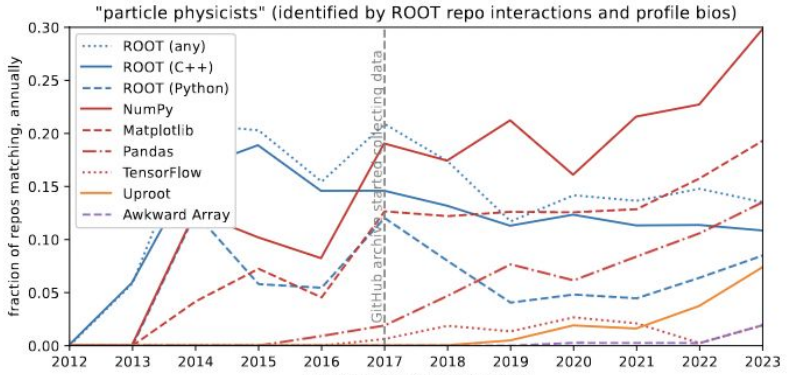


Tracking of CMS software trends

# Conclusions

- *HSF AF Forum activity provided a platform for sharing the experience between experiments* and helped to define the key areas needed for future AF R&D
  - *One of the goals of HSF AF white paper* was to identify a *list of AF "building blocks"* necessary to run HL-LHC workflows *as well the open questions for community*
  - *Columnar analysis approach should cover the majority of analysis in the future*
    - Experiments should look more in depth at how the analysis at the HL-LHC will look and how to adopt it at AFs
    - It is important to make sure that remaining analysis are supported without any issues
- **Integrating new technologies takes several years and AFs are ideal place to experiment/adopt on new approaches needed for HL-LHC analysis**

# Backup

## Open Questions

**On interactivity of analysis sessions**

What is the user expectation of an *interactive* session? Interactive analysis sessions on large datasets implies the ability to scale out to distributed resources and users will expect *instant* access to such an analysis session. What are the required disk resources, computing cores (order 100s) and network capabilities to realize this? The answers to these questions are analysis/workflow dependent, vary widely, and will need to be quantified as such. (The definition of interactive remains an open question for the community). The question is very pertinent to analysis in the HL-LHC era, where datasets will be much larger (~x10) for ATLAS and CMS from Run 4 onwards and will also grow again for LHCb and ALICE from Run 5.

This does not *automatically* mean that interactive needs jump by x10, but an increase will happen and should be quantified.

**On the Interfaces to resources provided**

Analysis work is multi-faceted and while a browser-based interface can be attractive for some use-cases such as interactive computing with notebooks, it should not be the only possible interface to the resource. Should users expect to be able to connect via their preferred text-based terminals or IDE such as VSCode?

## Open Questions

**On resource interoperability?** If we cannot guarantee uniformity of technologies how do we manage “interoperability” between resource providers?

**On a common interface**
Given the large number of new technologies should users be exposed to all of these or should there be a common entry point interface such as JupyterLab?

## Open Questions

**On data “locality” at facilities**
Should analysts expect the data, particularly reduced formats, to be local to any facility they wish to use (thus providing low latency access)? Often analysts work with derived datasets (with extra cuts, derived variables), does the same apply?

**On POSIX file access?**
Is POSIX required? Maybe just interacting with an object store via e..g xrootd / https / analysis software is sufficient? How much work is required to fully support object stores? Users like filesystem-like semantics, but what part of POSIX is really needed and can we decouple mass storage access from more interactive, smaller scale activities?

**On user interaction with Distributed Data Management systems**
Can all Distributed Data Management queries be hidden from the user and is this desirable? Should users expect that this is managed for them?

**On provenance of intermediate data products**

Which intermediate files need to be promoted from local to global storage so that users can run at different sites and how is this declared by the user?

**On common file sharing services**

Do we need a common file sharing service to help users share their files? In which case can any existing services fulfill this purpose?