

Analysis Facilities discussion

US ATLAS - US CMS joint meeting at HTC24
July 10, 2024

Analysis Facilities for the HL-LHC

In the past few years much attention has been given to so-called *Analysis Facilities*.

- HEP Analysis Ecosystem 2017 [workshop](#)
- IRIS-HEP [Analysis Systems R&D on Scalable Platforms](#) 2019
- WLCG pre-CHEP 2019 [workshop](#)
- IRIS-HEP Future Analysis Systems 2020 blueprint [workshop](#)
- HSF Analysis Ecosystems II 2022 [workshop](#)
- WLCG pre-CHEP 2023 [workshop](#)
- WLCG/HSF May 2024 [workshop](#)
- HSF AF forums: <https://indico.cern.ch/category/14976/>

HSF [Analysis Facilities Whitepaper](#) was published in April covering:

- User perspectives, compute and data access, consistency across infrastructures, continuous integration deployment and other features of current AFs

Yet significant questions remain:


- What are the use cases, analysis model differences Run 3 to Run 4, organization, benchmarks, and dedicated hardware needed



Only a subset of what are called AFs out there...


CERN and NAF examples

SWAN + HTCondor for interactive analysis



- Swan fits very well my needs for:
 - prototyping code and algorithms
 - plotting final results
 - working on ML models interactively
- It fills the gap between:
 - full-scale analysis (condor jobs)
 - interactive play with the results (difficult to do by running scripts on lxplus) == definition of the jupyter notebook.)


- NAF is considered an analysis facility for ATLAS.
- One of the main benefits of NAF is large and accessible storage.
 - Ease of sharing of the data between analysers inside DESY and in Germany.
 - Many workflows supported so everything can be done in one location.
- NAF is vital for German CMS analysers
 - for many, grid jobs are not even necessary



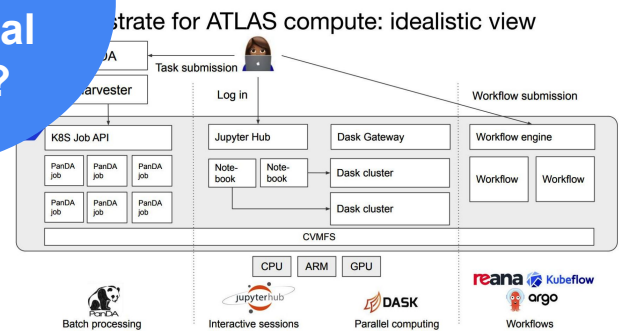
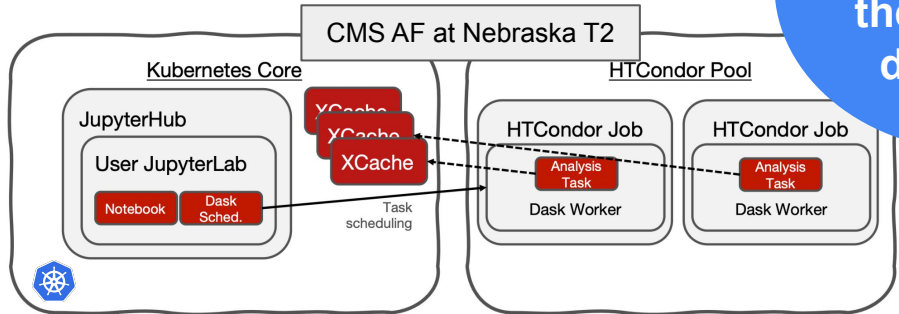
NAF: <https://indico.cern.ch/event/1214418>

SWAN: <https://indico.cern.ch/event/1180396/>

[D. Ciangottini, A. Forti Pre-CHEP 2023](#)
[E. Tejedor, Analysis Facility at CERN](#)



How to benefit from the artisanal diversity?



B. Bockelman, [IRIS-HEP 200 Gbps challenge](#)

Fernando Barreiro Megino, Lukas Heinrich, [KubeCon Oct 2022](#)

One approach: focusing on building platform factories

- Ultimately we have a long future ahead with many possible workflows and services under development, many more to be created
- In the near term it makes sense to focus on tools of the day:
 - Moulding the AF to support users where they are today - frameworks and formats
 - Build around promising initiatives with Coffea, Awkward, Dask, RDataFrame
 - stake out where we are with demonstrators of varying scales and types
- But also take a step back and consider the "how" we are creating these
 - Start from a **software defined infrastructure** => deployment => Ops for efficient usability
 - Imagine making the platform easy to build (Coffea Casa's approach)
 - Can we think of "products", or even "parts"?
 - If so, with a "factory" could we "**mass**" produce them? (At least so components/parts can be re-used around different resources -- where it makes sense.)
 - Today this might look like Kubespray & FluxCD & special sauce by various sous chefs.. is there room for abstractions on top, to speed up devOps?
 - eg. take advantage of CI creation tools from the cloud native community, e.g. dagger.io, stateful.com (DevOps notebooks)?
 - Can we create space for dev teams to do full devOps and test out "parts" themselves?
 - Making it easy to test and benchmark new tools - essential to mark progress - inform decisions on hardware

Keep a flexible base to deploy and try out the best ideas

CERN and NAF examples

SWAN + HTCondor for interactive analysis

- Swan fits very well my needs for:
 - prototyping code and algorithms
 - plotting final results
 - working on ML models interactively
- It fills the gap between:
 - full-scale analysis (condor jobs)

SSH and Jupyter-based access to private containers and shared batch login with shared SDATA.





"Behind every great product is a great factory"

Solomon Hykes, Docker founder (ref)

Continuously improve the factory and the product together

.. and leverage the platform building community

<https://landscape.cncf.io/>



Defining Analysis Facilities for the HL-LHC

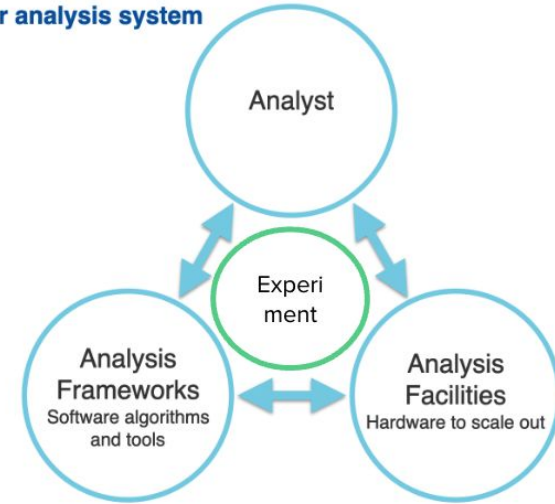
From [Follow-up on the Focus Session on Analysis Facilities held at the 154th LHCC in June 2023](#)

- The HSF Analysis Facilities Forum has focused on the technical building blocks that can make the user experience more productive. Novel concepts of Analysis Facilities, integrating interactivity, scalability (off-load) and Machine Learning tools are being built.
- 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 the HL-LHC Computing and SW review panel. The document is expected to be regularly updated in the process towards HL-LHC.

Charge

To follow-up on the LHCC recommendation cited above, the experiments, the sites and the HSF Analysis Facilities Forum must engage in a dialogue towards defining use cases and establishing realistic benchmarks for Analysis Facilities. The LHCC is charging WLCG to oversee this continuous process and to report regularly to the LHCC on the progress and steps taken as described below.

A better analysis system



Nick Smith

1- Establishing a list of questions

→ questions reviewed & finalized at HSF/WLCG [workshop](#) in Hamburg [[Excellent summary of the AF session](#) by A. Forti]

LHCC charge, cont

[Follow-up on the Focus Session on Analysis Facilities held at the 154th LHCC in June 2023](#) (cont)

- Analysis Facilities is a broad topic. The use cases and the expectation of the experiments, and of the analysers may evolve over time. The scope and the expected content of the document to be provided by the experiments must be clearly defined. To this end, a list of questions must be defined first that seek to define the expectations from experiments for Analysis Facilities. The questions, to be answered by experiments, must be picked such that the answers are useful for sites and are representative of a broad spectrum of analyses and analysers. The list of questions might evolve in the future. **1 The LHCC charges WLCG with establishing such a list.** This process will be iterative and must involve the HSF Analysis Facilities Forum, **the sites and the experiments.** The monthly GDB (Grid Deployment Board) meetings of WLCG could be a place for the relevant discussions to occur, but it is up to WLCG to establish the appropriate mechanism. A first list will be established by the June 2024 LHCC meeting and presented by a WLCG representative at the LHCC to the LHCC for comments.

As of last week, S&C coordinators were expecting the questions from LHCC very soon


How can current **facilities** (sites) contribute to this process?

- development of platforms (e.g. multi-site AFs)
- benchmarks & demonstrators
- trying new technology
- use of cloud, bursting to resources
- questions of storage, caching in-network
- cross-site integrations
- we don't have to work in silos, even if we do different things!

breakdown by A. Forti



Questions structure



- 6 question with a number of sub-questions each divided in 3 categories
 - Analysis model(s) to understand the evolution
 - Run3 and then Run4 (Run5) to highlight the differences
 - What would these AFs look like?
 - AF organisation resources, support, technology and hardware
 - How is the experiment going to get there and measure the benefits
 - Analysis, use cases, benchmarks and R&D

Caveats: Answers for CMS/ATLAS as a whole (not just US parts and not just the 'new style AF'). Also, answers for CMS and ATLAS can and will differ.

LHCC questions

1. Describe the current Run-3 analysis model:
 - a. Briefly describe the Run-3 analysis model and main analysis workflows, noting the required data reductions steps and how closely chained they need to be.
 - b. Describe the data formats used in the analysis chains, including their relative size and the fraction of analyses that can use reduced-sized formats (current and future Run-3 expectation).
 - c. How much compute, storage, and network resource are used for Run-3 analysis? Indicate broadly, what fraction of this is pledged resource, and how much is accessed in an 'interactive' (as opposed to 'batch') like manner.
 - d. What features in the analysis chain do current users find most useful and what aspects are the most difficult?
 - e. What are the successes and problems of the current analysis system as far as the providers (Experiment software & computing teams, and Sites) are concerned?

LHCC questions

2. Explain what is driving the need to develop dedicated Analysis Facilities for Run-4 and Run-5 data.
 - a. What data volumes (per year and total per run; both data and MC) are expected?
 - b. How will the data model change in terms of the number of versions, the number of replicas, and the degree to which the data is centrally managed?
 - c. What access to extra information from other resources do analysis workflows need, and will this change?
 - d. What aspects of the current analysis system do not simply scale to Run-4 and Run-5 or are already a limitation?
 - e. Over the next few years, what are the expected evolutionary, or possibly revolutionary, factors that support a decision to develop Analysis Facilities (do you anticipate new resources, new technologies, new analysis techniques, capabilities, or paradigms) ?
 - f. Why are analysis facilities a better option than evolving the current paradigm that uses the general Grid infrastructure?

LHCC questions

3. How might a dedicated Analysis Facility be organised?
 - a. How many Analysis Facilities would be needed? How similar would they need to be? What would be the access policy? How centrally organised would the Analysis Facilities need to be?
 - b. ~~Would Analysis Facilities use just local resources, mostly distributed resources, or a mixture? Which of these resources would be pledged?~~
 - c. What combination of batch and interactive usage might be envisaged?
 - d. Would an Analysis Facility be tailored for specific analysis workflows and/or specific working groups?
 - e. How closely would an Analysis Facility need to be integrated with an existing Grid site?
 - f. Are there any new security issues concerns specifically associated with an Analysis Facility?

LHCC questions

4. What hardware and support would a dedicated Analysis Facility require?
 - a. What specific hardware configuration would an Analysis Facility require, such as: CPU-GPU combinations; high memory nodes; high performance disk; caches; storage classes; bandwidth; etc?
 - b. What level of support would an Analysis Facility require for deployment, operation, software and services; user-support; documentation etc?
 - c. What level of commonality could or should there be between Analysis Facilities for different experiments in the situation where hosting sites support multiple experiments?

5. What are the plans to develop specific use cases that could be used to motivate and benchmark Analysis Facilities so that a comparison can be made with the existing infrastructure and identify the key technology challenges identified?

6. Describe what currently exists that could be described as an Analysis Facility or component there-of. What is the status of any R&D development work that is currently underway towards an Analysis Facility?