



Analysis Grand Challenge Steering Board #11

Alex Held (NYU)
Oksana Shadura (UNL)

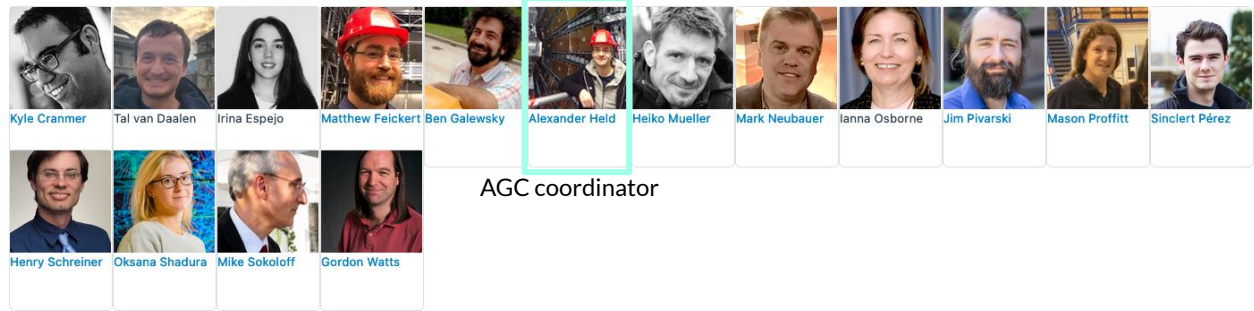
Sept 7, 2021
<https://indico.cern.ch/event/985525/>

This work was supported by the U.S. National Science Foundation (NSF) Cooperative Agreement OAC-1836650 (IRIS-HEP).




Team

AS



DOMA



SSL



Analysis Grand Challenge

Motivation:

- Allow coping with HL-LHC data sizes by rethinking data pipeline
- Provide flexible, easy-to-use, low latency analysis facilities



Looking for new ideas for Analysis Facilities

- **New pythonic ecosystem**
- Discovering the benefits of **column-oriented (columnar) data analysis**
- **Interactivity** for user data analysis
- Deliver the needed data to the processing workflow in a fine-grained approach (**data delivery services**) and **efficient storage technologies** (e.g. object stores)
- **Kubernetes (k8s)** and new concept of "**infrastructure as code**"
- **Portability** and flexibility across different environments
- Integration with existing resources: current infrastructure is not going to be replaced in one day








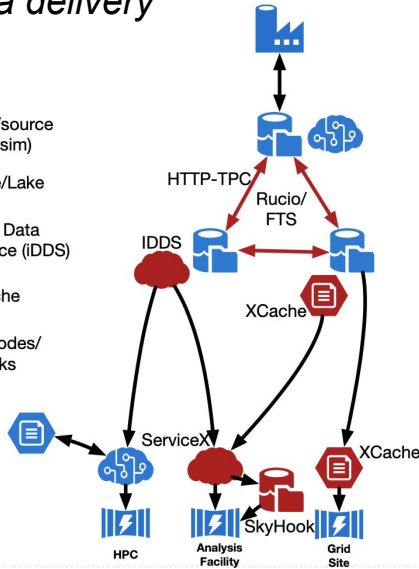
Coffea-casa vchep 2021 plenary talk

Analysis Grand Challenge will be conducted during **2021–2023**, leaving enough time for tuning software tools and services developed as a part of the IRIS-HEP ecosystem before the start-up of the HL-LHC and *organized together with the US LHC Operations programs, the LHC experiments and other partners.*

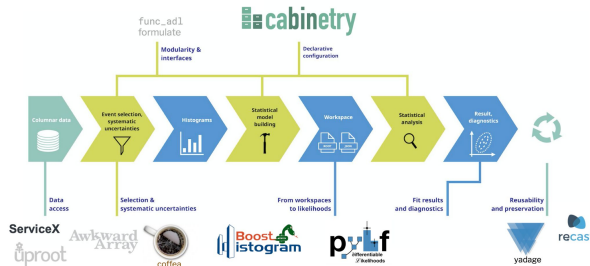
AGC is connecting IRIS-HEP focus areas

DOMA: Data delivery


-  Data Factory/source (e.g. T0 or sim)
-  Data Store/Lake
-  Intelligent Data Delivery Service (iDDS)
-  Data Cache
-  Compute Nodes/Data Sinks



AS: tools



SSL: deployment techniques and resources





Activities

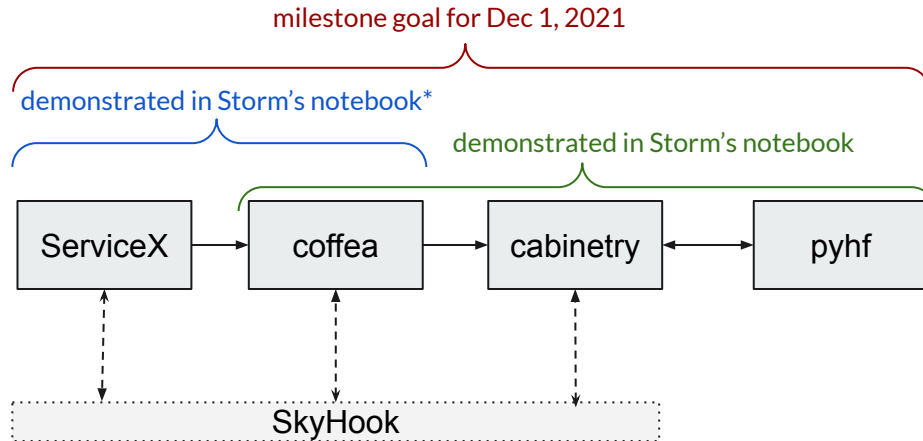
- Define **target analysis and dataset** to be used in the Analysis Grand Challenge
 - Building towards full analysis, now working with smaller prototypes
- Define and investigate **baseline programming interfaces** between components
 - Investigate differentiable pipeline as a part of ACG
- Prototyping and deploying **Analysis Facilities** for executing Analysis Grand Challenge
- **Coordinate** with AS, DOMA, SSL, and operations programs to benchmark performance of prototype system components to be used for Analysis Grand Challenge and to execute the Analysis Grand Challenge

Analysis Grand Challenge Analysis requirements

- **Analysis needs to include features encountered in realistic ATLAS/CMS analyses**
 - Handling of **large data volumes** (rough number we have in mind: ~200 TB)
 - Handling of **different types of systematic uncertainties**
 - Ideally use of **modern formats** *NanoAOD / PHYS / PHYSLITE*, which would make it easier for ATLAS/CMS analyzers to adopt to their use case
- **Not intent to send physics message with analysis**
 - Want to show **realistic workflow**, not make physics claims
 - No need for real data, **simulation fully sufficient** (ideally many samples to simulate book-keeping)
- **Want to demonstrate enhanced functionality**
 - Possibility to **end-to-end optimize physics analysis**, potentially via automatic differentiation
 - Analysis needs to run on analysis facility
- **Analysis needs to be sufficiently specified for others to re-implement**
 - **Ideally: data is open and available to everyone (or scheduled to become public in the near future)**
 - Hoping to learn from comparing to implementations developed by others outside IRIS-HEP
 - Want to turn parts of analysis into mini-benchmarks for facility and tool benchmarking

Recent activities from Analysis Systems

- IRIS-HEP fellow Storm ported **ATLAS H>ZZ* open data** example to **IRIS-HEP tools**
 - Repository: [stormsomething/CoffeaHZZAnalysis](https://github.com/stormsomething/CoffeaHZZAnalysis), also contributed back to ATLAS ([PR](#))
 - Outcome of project: set of notebooks for training and functionality testing, usability & API/interface feedback, first tests of OpenData coffea-casa



- Related AS activities: [discussion of differentiable pipeline](#) & interfaces (see also Dec 1 milestone)

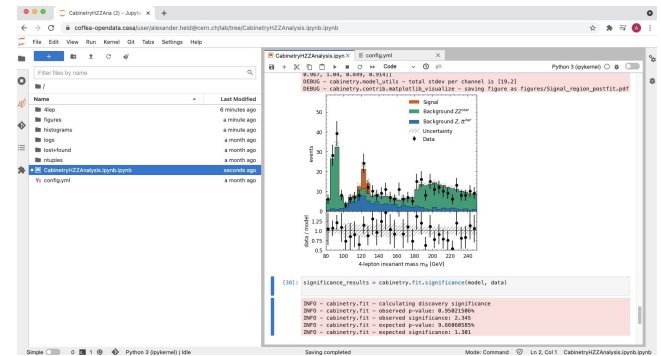
Benchmarks



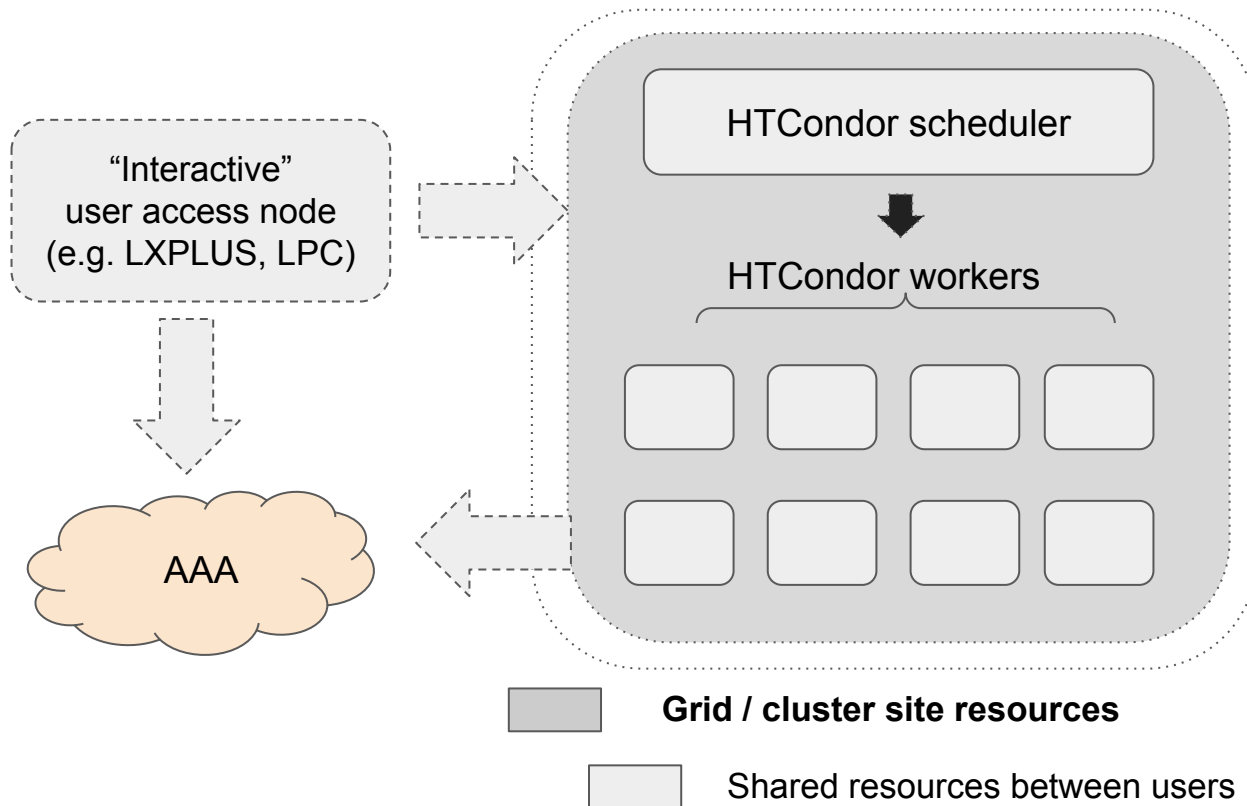
- **HSF DAWG** interested in expanding existing [ADL benchmarks](#)
- **HSF DAWG and ACG** identified several potential **directions for extensions**
 - *Testing interfaces between different tools in analysis pipeline*
 - *Handling of systematic uncertainties*
- **Idea:** to specify **AGC** sufficiently well so that it can be used as **very large benchmark**
 - Also want to split into **sub-tasks** that can be used for benchmarks
 - Detailed specification may attract other users to write new implementations
- Specification of N new benchmarks potential new **milestone for ACG**

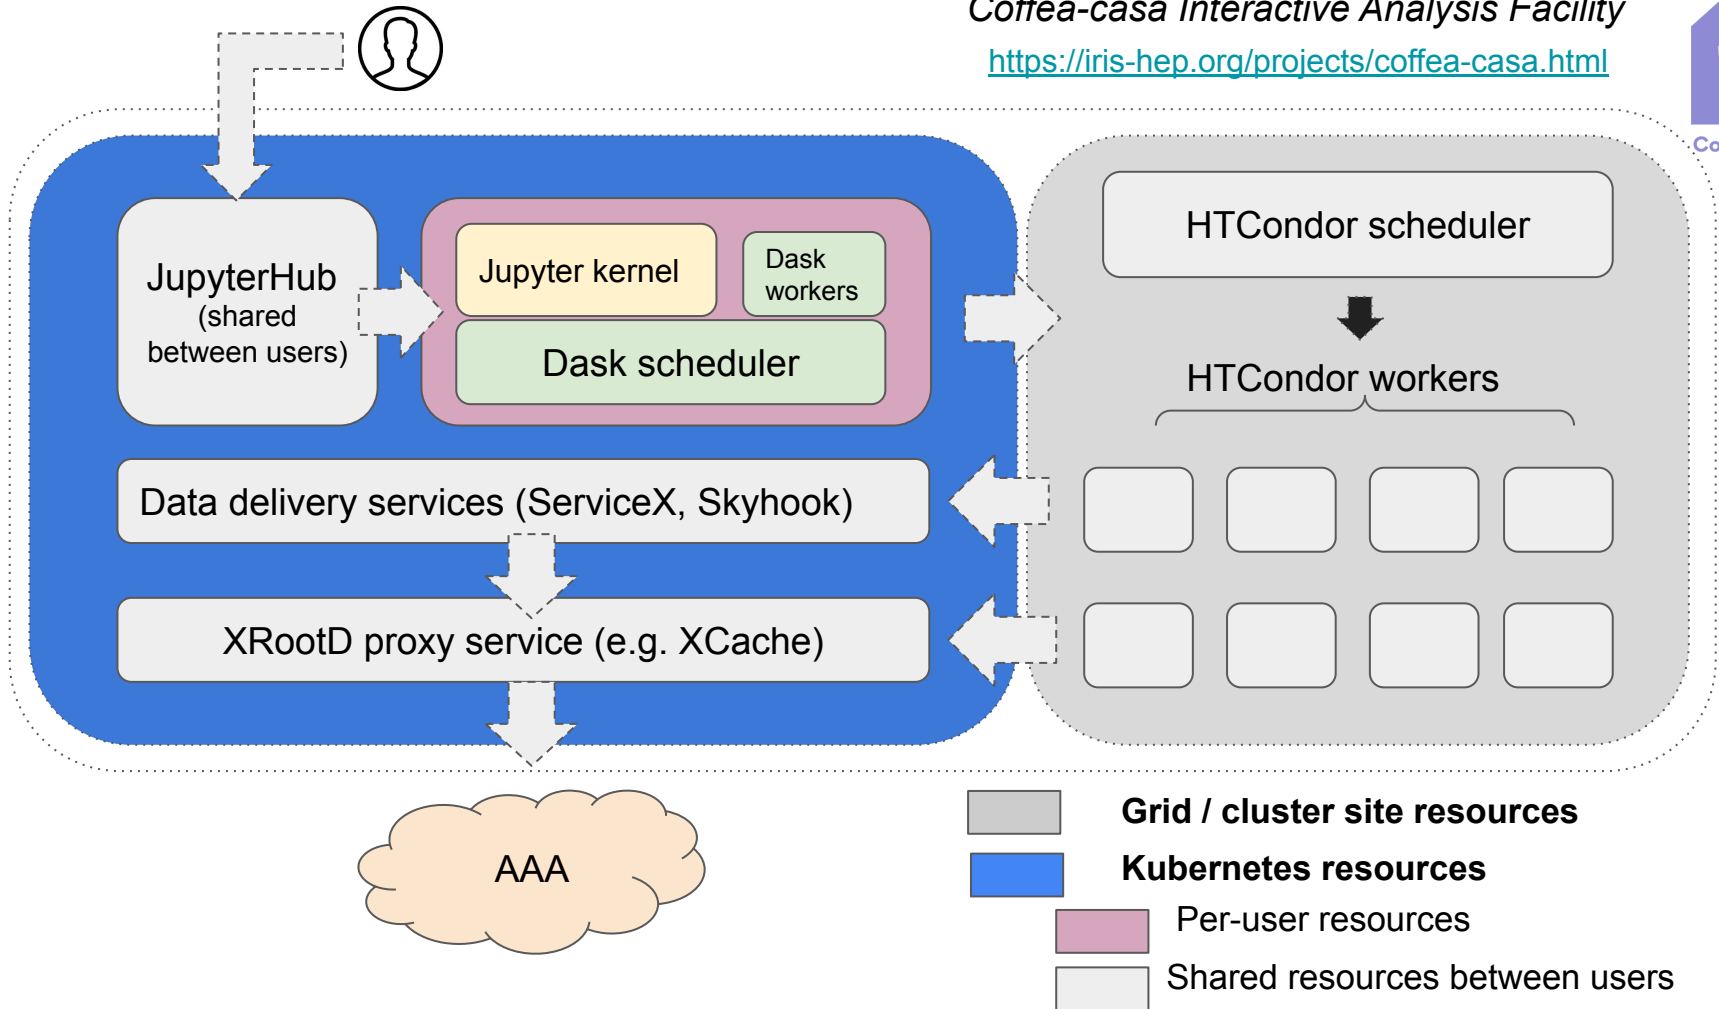
Recent activities on Analysis Facilities

- **HATS training** ([CoffeaTeam/coffea-hats](https://github.com/CoffeaTeam/coffea-hats)) successfully executed at coffea-casa at Nebraska
 - More than 50 users (total number of participants over two days training)
 - Test for scaling coffea-casa resources with the number of users
- Released a first version of **coffea-casa Helm charts**
 - Collecting feedback from SSL and FNAL teams
- **Deployment of the OpenData coffea-casa instance**
Nebraska: <https://coffea-opendata.casa>
 - Open to everyone (after approval process)
 - Preparing to deploy OpenData and ATLAS *coffea-casa* at University of Chicago (SSL)



Simplified diagram of hypothetical Analysis Facility currently used by users





Building blocks used for designing AFs



Modern authentication (AIM/OIDC), tokens, macarons

Efficient data delivery and data management technologies

Columnar analysis and support new pythonic ecosystem

Modern deployment and integration techniques

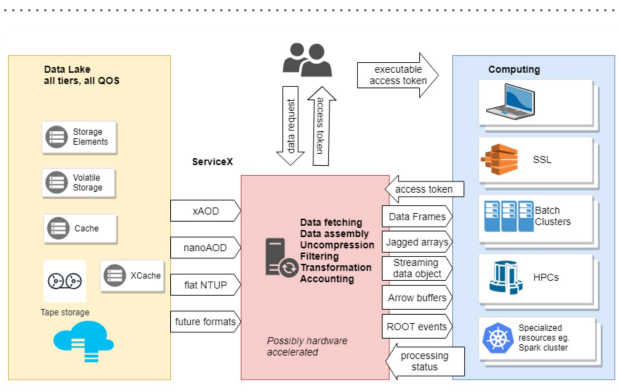
Support for object storage

Efficient data caching solutions

Easy integration with existing HPC resources

Data delivery services


ServiceX



ServiceX provides user level ntuple production

- Converts experiment-specific datasets to columns (e.g. NanoAOD, DAOD)
- Enable simple cuts or simple derived columns and fields (*heavy-weight analysis will still happen via some separate processing toolchain*)


Skyhook DM



The Skyhook DM is converting event data from ROOT files to the internal object-store format

- Mechanism to access data kept in CephFS through the popular Arrow libraries
- Enables pushing down filters, projections, compute operations directly to the storage backend to minimise network overhead
- Allows writing files to a POSIX filesystem

Columnservice



Columnservice (FNAL) is a multi-tenant service for caching columnar data that removes the need to curate skims and re-run expensive algorithms.

Next steps (1)

- Training event showing **IRIS-HEP toolchain at coffea-casa**, aimed at PhD / postdoc level
 - Format: brief introduction to individual packages, notebook talks focusing on interfaces between tools
 - Initially using OpenData examples, then splitting into ATLAS / CMS - specific tracks
 - Homework + Q&A session
 - 2 afternoons CERN time (~15:00 - 19:00) between **Oct 25 and Nov 5** (after US ATLAS computing bootcamp scheduled for Oct 18-22)
- Demonstration of **ServiceX -> coffea -> cabinetry -> pyhf pipeline on ATLAS OpenData**
 - Extending existing ATLAS OpenData H>ZZ* examples
 - Experience from this exercise will inform further API developments in AS
 - **Target: Dec 1, 2021**

Next steps (2)



- Test integration of **SkyHook in coffea-casa@UNL**: end of September
- Develop, test and deploy **different 'flavours' of coffea-casa instance** (at least one instance each):
 - *CMS*
 - *Opendata*
 - *ATLAS* (need to collect technical requirements)

Target: **should be ready before IRIS-HEP training event**

Question: Can we use modern OpenData for AGC scale?

- Existing large datasets in OpenData restricted to CMS Run-1
- Would ideally prefer **modern ATLAS PHYS/PHYSLITE or CMS NanoAOD formats**
 - Makes AGC implementation more relevant to current/future analyses & ideally re-usable
- Following up with ATLAS & CMS to understand whether we may be allowed to use (a) new dataset(s) for technical demonstration



Backup slides

AGC analysis definition: proposed approach

In 2021

- stick to **ATLAS H>ZZ*** example for demonstrations, allows testing interfaces and is ready
- identify possibility of using datasets with new ATLAS / CMS formats in parallel
 - May involve re-formatting existing OpenData if no new datasets become available

June 2022 milestone (“benchmark performance of prototype system components for AGC”)

- If **modern ATLAS / CMS formats** are available: design analysis around what samples we get
- Otherwise: extend **CMS H>tautau OpenData** analysis ([OpenData record](#)) with systematic uncertainties covering all uncertainty types identified in taxonomy (different types requiring different approaches)

March 2023 milestone (execute AGC)

- If **new ATLAS / CMS data** is unavailable: (reformatted) **CMS Run-1 OpenData**-based analysis

Past relevant milestones (selection)



- **Mar 1, 2021 (AS): G2.6** Scoping and specification of the target analysis to be used in the Analysis Grand Challenge
- **Mar 1, 2021 (AS): G2.7** Baseline programming Interfaces between components like ServiceX, func_ADL, HEP_tables, Coffea, cabinetry, and pyhf

Next relevant milestones related to AGC

- **Dec 1, 2021 (AS): G2.9 Differentiable programming roadmap** across services needed for analysis challenge
 - *Thinking about the possibility of a blueprint meeting in November to address this*
- **June 1, 2022 (AS, DOMA, SSL): G2.15** Coordinate with DOMA, SSL, and operations programs to benchmark **performance of prototype system components** to be used for Analysis Grand Challenge
 - *Define and test interoperability of software components (services and packages) to be used for ACG and show the demonstrator(s)*
 - *Select dataset(s) and examples (CMS Run1 AOD, NanoAOD or DAOD) applicable to be used for AGC (from G2.6)*
 - *Deploy multiple ATLAS, CMS and OpenData coffea-casa instances ready for benchmarking*
- **Sep 1, 2022 (management): G1.13 Blueprint workshops** (Grand Challenges Planning)
 - *Falls in between prototype phase (due June 1, 2022) and execution (March 1, 2023), could be used to identify issues based on experience with the prototype and figure out a plan to address them*

Next relevant milestones related to AGC



- **Dec 1, 2022 (DOMA): G3.11** Demonstrate ability to **filter / process data** at rate necessary for analysis challenge using **SkyHook**
 - *Prepare coffee-casa setup ready for benchmarking (related to G2.15)*
- **Mar 1, 2023 (AS, DOMA, SSL) G2.16** Coordinate with DOMA, SSL, and operations programs to **execute the Analysis Grand Challenge**
- **Dec 1, 2023 (management) G1.22** **Blueprint workshops** (Grand Challenges Progress)

Communication channels



- <https://github.com/iris-hep/analysis-grand-challenge> to accommodate milestones, related materials and documentation
- IRIS-HEP slack: [#analysis-grand-challenge](#)
- Possible topical meeting slot with AS, DOMA, SSL groups to discuss relevant ACG topics (can schedule *on demand* during IRIS-HEP Mon/Wed meeting slot)