Analysis Grand Challenge

Alexander Held (UW-Madison)
Oksana Shadura (UNL)

IRIS-HEP / Ops Program Analysis Grand Challenge Planning

Aug 23, 2022: https://indico.cern.ch/event/1173741/

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

Analysis Grand Challenge

Motivation:

- Allow coping with HL-LHC data sizes by **rethinking data pipeline**
  - Evaluating the new Python "datascience" analysis tools, integrating differentiable analysis pipeline
- Provide flexible, easy-to-use, low latency **analysis facilities**

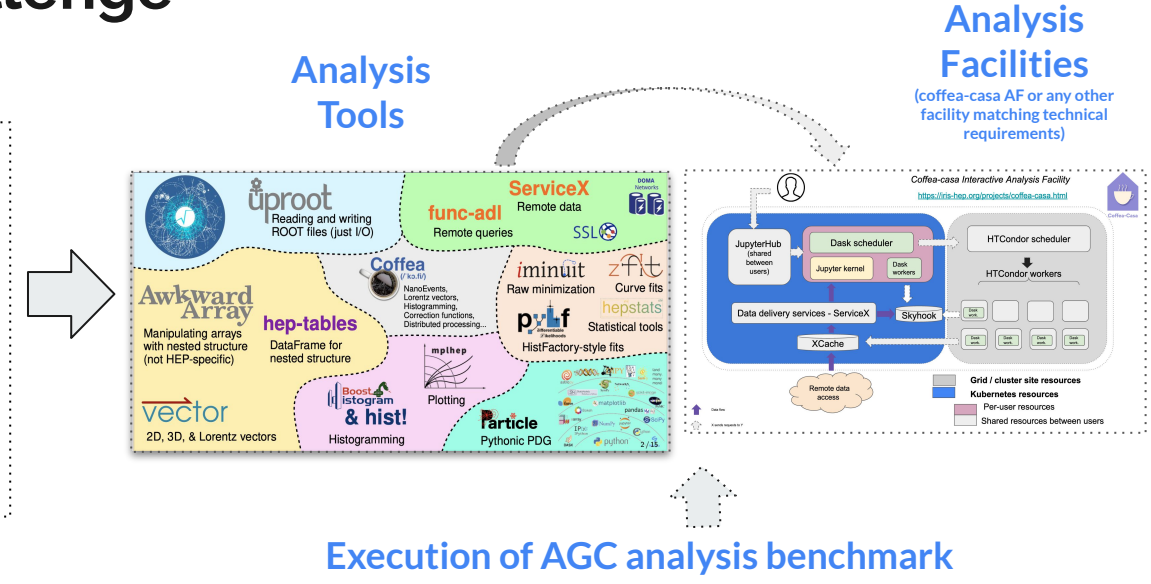

The AGC is being conducted during **2021–2023**, leaving enough time for tuning software tools and services developed as a part of the IRIS-HEP ecosystem before the start-up of the HL-LHC and *organized together with the US LHC Operations programs, the LHC experiments and other partners.*

Recent news:

Fellow projects, CMS Open Data Workshop

AGC related IRIS-HEP Fellow projects

- **Durbar Chakraborty:** Metrics to define user activities and engagement on the various coffea-casa Analysis Facility deployments (https://iris-hep.org/fellows/durbar2003.html) - finished
- **Haoran Sun:** Enabling support for MiniAOD Transformer for ServiceX Data Delivery Service (https://iris-hep.org/fellows/hswhhe.html) - in progress
- **Andrii Falko**: documentation of AGC analysis task & implementation in ROOT's RDataFrame (https://iris-hep.org/fellows/andriiknu.html) - in progress
- **Jake Li:** Jupyterhub plugin that will provide users access to a ServiceX dashboard directly from AF (https://iris-hep.org/fellows/jakel2014.html) - finished
    - Collecting list of all ServiceX instances: please see Slack (or send us URLs)

ServiceX JupyterHub plugin

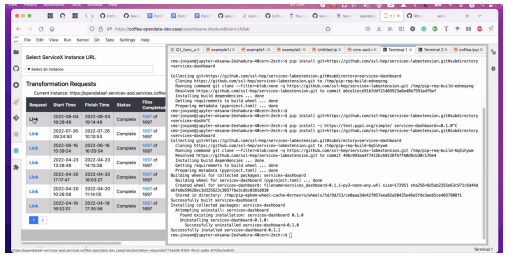

AGC related US.CMS Fellow projects

- **Holly Wingren:** AGC component benchmarking - finished
- **Sneha Dixit:** Data generation (NanoAOD) from CMS Opendata 2015 (recent release), https://github.com/cms-sw/cmssw/pull/39040 - finished
  - Will replace AGC input data (flat ntuples) by NanoAOD files produced in this official way
  - Expect that these NanoAOD files will also be provided officially/centrally by CMS

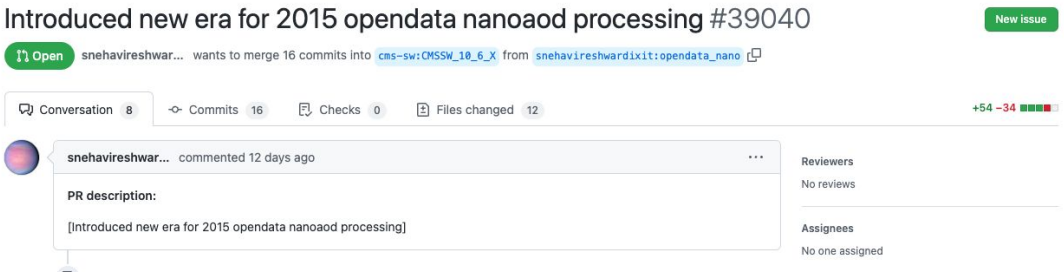

ICHEP 2022

- Presented **AGC at ICHEP** (slides)
- There is interest in adding **ML inference** to the workflow
  - Clearly frequently used in practice
  - Less clear how to best integrate it: efficient use of accelerators, likely many site-dependent differences
  - We are looking into including this in the pipeline
- Less in scope, but also interesting: **ML training workflow**
  - Could become a standalone project?

CMS Open Data workshop

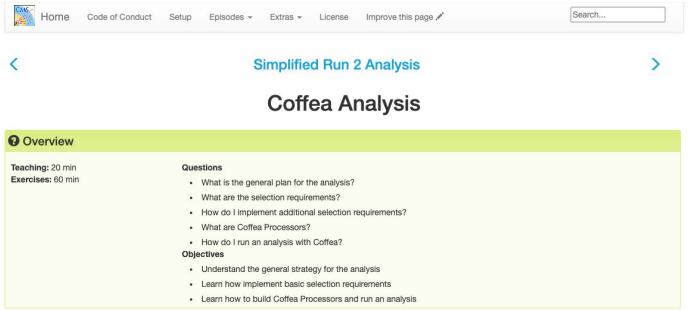

- Aug 1-4: **CMS Open Data workshop** (https://indico.cern.ch/event/1139022/)
  - Hands-on tutorials of how to use CMS Open Data
  - Included demonstration of using Run-2 Open Data, heavily based on AGC (tutorial page)
    - the AGC analysis is turning into an “official” CMS Open Data demonstration
- CMS **re-processed input files** for analysis (more precise calibration, similar ntuple structure)
  - Will investigate switching to this, or directly to NanoAOD format

AGC updates:

PyHEP + ACAT, benchmarking plans, site status

Next targets: PyHEP + ACAT

- **PyHEP 2022** (https://indico.cern.ch/e/PyHEP2022)
  - Presenting a tutorial, focused on the analysis pipeline & interactions between the libraries used
  - Expect participants to be able to follow along via Binder (using small set of input files)

- **ACAT 2022** (https://indico.cern.ch/event/1106990/)
  - Talk focused on first performance benchmarking results
  - Want to measure multiple setups:
    - Comparison of processing remote files read through XRootD (with / without XCache) and local files
    - At UNL: scaling to Tier-2 HTCondor queue vs HTCondor in k8s (running on new hardware at UNL)
    - Scaling with fraction of data accessed in files
    - How I/O limited is the setup (by removing event processing & comparing to theoretical limits of what the hardware could provide)
    - Happy to hear other suggestions!
  - ServiceX use in this context relies on next release (thanks to Burt Holzman for finding / resolving minio bug!)

Pipelines for performance testing

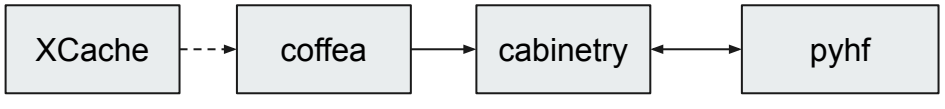

Columns from NanoAOD or flat ntuples processed by coffea

Starting out with plain coffea, reading inputs through XCache from remote

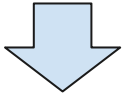

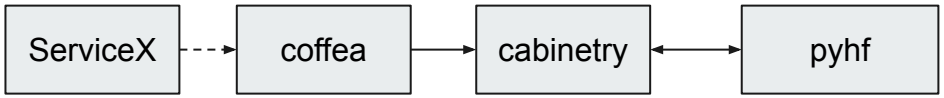

Columns from NanoAOD or flat ntuples from ServiceX

Want to also compare effect of serving data through ServiceX

Towards AGC at various sites across the US

- Investigate **XCache setup and configuration** for AGC execution
  - Deployed already / planned? Restrictions to which kinds of files can be cached?
- **Deploy and test `coffea`** at your facility
  - Please feel free to report issues to us, we can help relay them
- Investigate **scaling for `coffea`** (various backends available)
  - e.g. latency, performance with number of workers, workers a single user can get, …
- **Deploy and test `ServiceX`**
  - AGC only uses uproot transformers currently
  - Expect significant performance improvements from upcoming release
- How to **support additional dependencies** (on interactive nodes + workers)
  - Should users bring images?

- **Helps us help you**: **which issues are you running into** that we should prioritize / communicate to the respective teams?

What currently runs where? (please help us fill in the gaps!)

	BNL	FNAL	SLAC	UNL	UChicago
basic `coffea` (e.g. `IterativeExecutor`) -> notebook with `USE_DASK = False`	✓	✓		✓	✓
`coffea` scaling (e.g. with Dask) -> notebook with default settings*		✓		✓ (using HTCondor @ Tier2, planning to switch to k8s)	✓ (occasional segfaults at scale)
standalone `ServiceX` -> notebook (no configuration)	✓ (?)	✓		waiting for new release	✓
`ServiceX`+`coffea`+scaling -> notebook with `PIPELINE = "servicex_processor"`				waiting for new release	✓
`XCache` support	✓	✓ (some performance caveats, to be understood)	✓	✓	✓

* may need site-dependent Dask cluster configuration, see implementation, please get in touch in case of questions

AGC: timeline & plans

- **Nov 2021:** first demonstration of toolchain at AGC workshop (agenda)
- **April 2022:** second iteration of AGC workshop, new ttbar analysis pipeline (agenda)
- **Summer 2022:** benchmarking of system components in the AGC context
    - Results to be shown at ACAT 2022
- **Spring 2023:** execution of AGC at full scale

- Proposal for **next meeting**: **Sep 27, 2022** (9:00 PT / 11:00 CT / 12:00 ET / 18:00 CERN)

Stay in touch via analysis-grand-challenge@iris-hep.org (sign up: google group link), and please also feel free to contact us if you'd like to get involved or have any questions!

Backup slides

What do we mean by “analysis”?

- In view of the HL-LHC: “analysis” **starts** from centrally produced **DAOD_PHYSLITE**
  - Includes **all subsequent steps** to produce results needed for publication
    - (Re-) calibrate objects & calculate columns for systematic variations
    - Filter events & calculate observables
    - (Optional) histogramming
    - Construction of statistical model + statistical inference
  - Do all these steps in a **reproducible** way

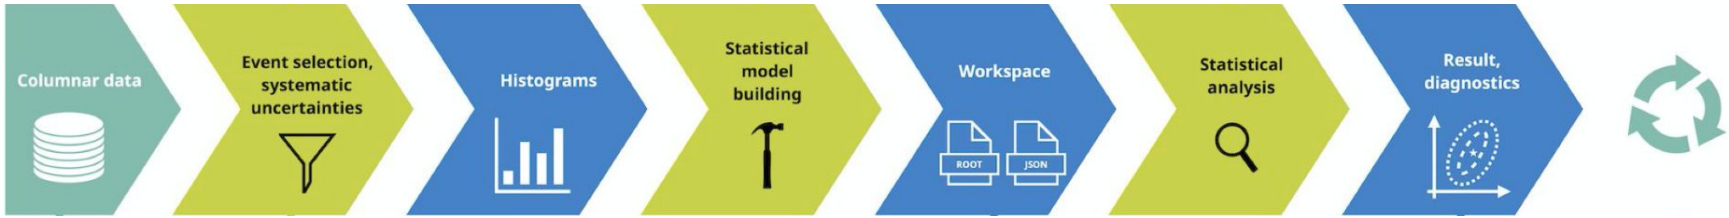

The AGC analysis setup

- Main AGC analysis example is based on **Run-2 CMS Open Data** (~400 TB of miniAOD available)
  - **Open Data is crucial**: allows everyone to participate!
  - Input data currently ntuples derived from miniAODs, will switch to nanoAOD once available
    - nanoAOD setup will more closely correspond to PHYSLITE workflows

- Generic single-lepton **ttbar** selection, simple top reconstruction, x-sec measurement
  - Including on-the-fly evaluation of systematic uncertainties
  - Flexibility to also turn setup into a BSM search
  - Currently using ~4 TB of ntuple inputs (~1B events before cuts) → through Rucio and xrootd(https) @ UNL

- Goal of setup is showing **functionality**, **not** discovering new physics
  - Want to capture workflow, but can made-up tools for calibrations & systematic uncertainties

Goals of the AGC

- Demonstrate method for **handling HL-LHC data pipeline requirements**
  - Large data volumes of O(100 TB) + bookkeeping
  - Handling of different types of systematic uncertainties
  - Use of PHYSLITE / nanoAOD formats, aligned with LHC experiments

- **"Interactive analysis"**: turnaround time of ~minutes or less
  - Made possible by highly parallel execution in short bursts, low latency & heavy use of caching

- **Specify all analysis details** to allow for **re-implementations** and re-use for benchmarking

- Stretch goal: **differentiable analysis pipeline** (see e.g. https://indico.cern.ch/event/1096431/)
  - Investigate end-to-end analysis optimization, evaluate usefulness vs cost of gradient information

Status

- **Demonstrated analysis pipeline** at last AGC workshop: https://indico.cern.ch/e/agc-tools-2
  - Workshop also features details of all tools & services used in pipeline
  - **Pipeline setup**
    - `ServiceX` delivers columns following declarative `func_adl` request
    - `coffea` orchestrates distributed event processing & histogram production
      - Using `uproot`, `awkward-array`, `vector`
    - Visualization with `hist` & `mplhep`
    - Statistical model construction with `cabinetry` & inference with `pyhf`

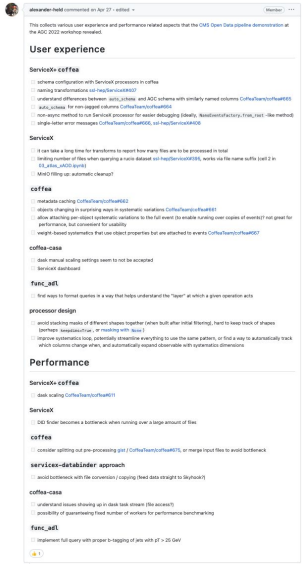

- Pipeline demonstrator highlighted possible **usability & performance improvements**
  - Everything tracked in agc#64
  - Main performance bottlenecks are being addressed

- **Everything is openly developed** (AGC repository)
  - Including categorization of datasets in terms of role in AGC demonstrator (AGC repository)

Expanding the AGC, connections with ATLAS & CMS

- Looking to further **expand AGC demonstrator** & **connect more closely to ATLAS and CMS**
  - "ATLAS version": PHYSLITE input + (simplified?) CP tools (related: PHYSLITE schema for coffea)
  - "CMS version": more CMS specific NANOAOD input
  - R00T **`RDataFrame`-based implementation** of pipeline
  - Would be interesting to talk to analyzers currently using **distributed RDF / `coffea`** in ATLAS/CMS
    - Please get in touch!
  - Actively **looking for more person power** to expand scope

- Interested to hear **feedback**: how to best **align AGC effort with ATLAS**
  - Balance between retaining R&D nature and potential usefulness on Run-3 scale

Running the AGC: coffea-casa AF example

https://iris-hep.org/projects/coffea-casa.html

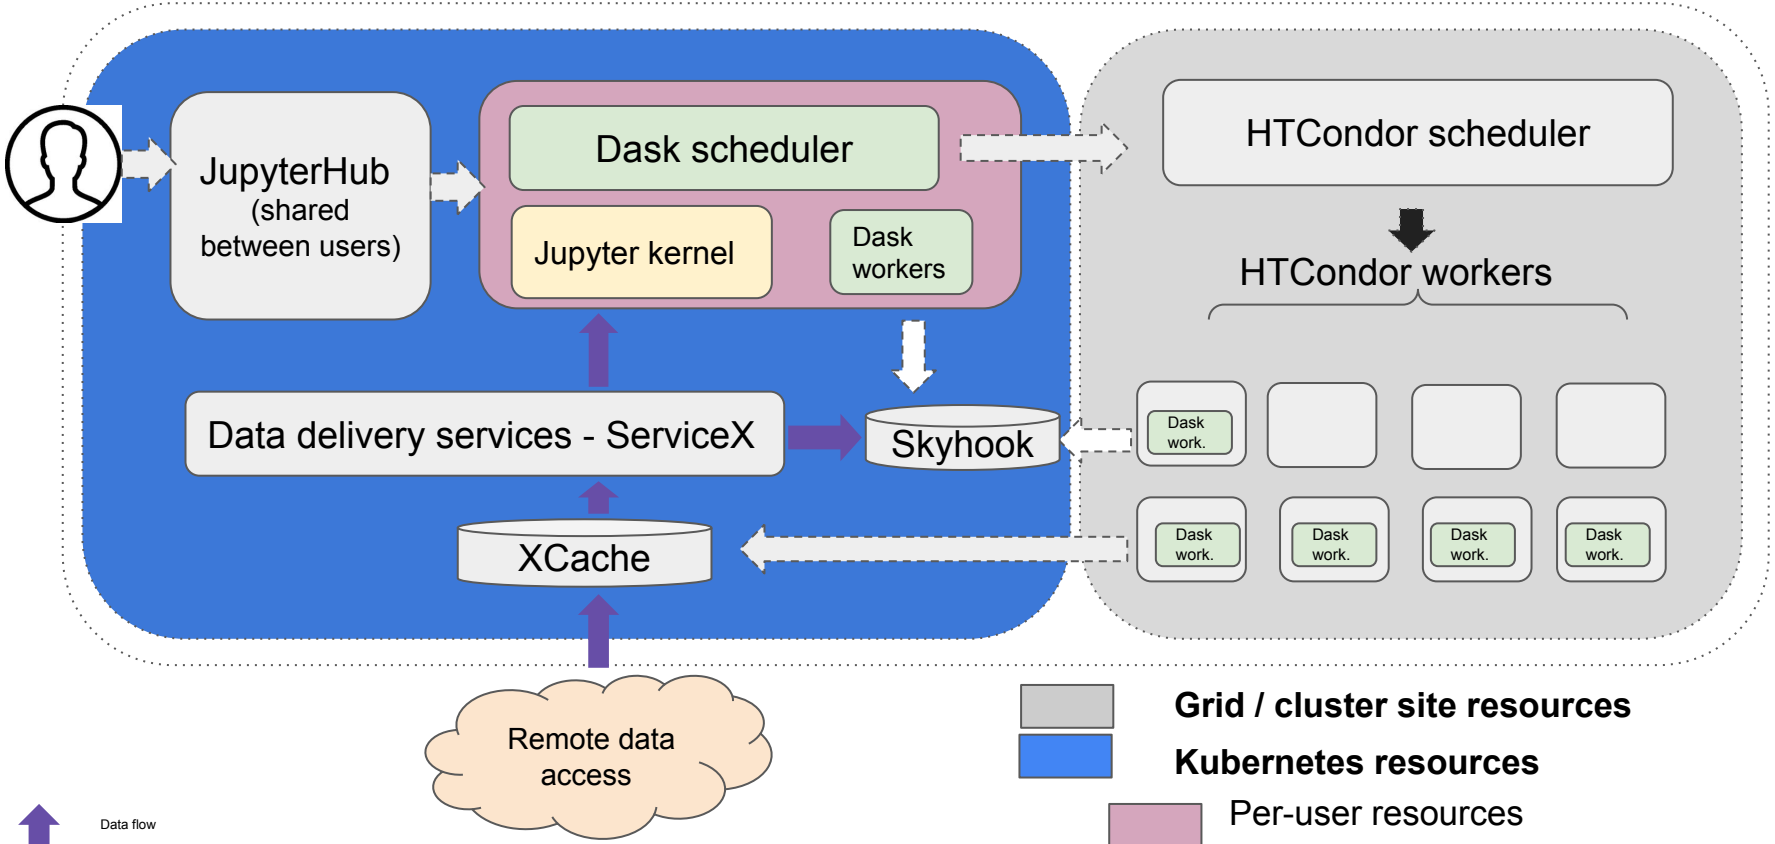

Matous Adamec, Garhan Attebury, Brian Bockelman, Ken Bloom, Carl Lundstedt, John Thiltges, Oksana Shadura, Andrew Wightman

Tools & services we are focusing on for benchmarking

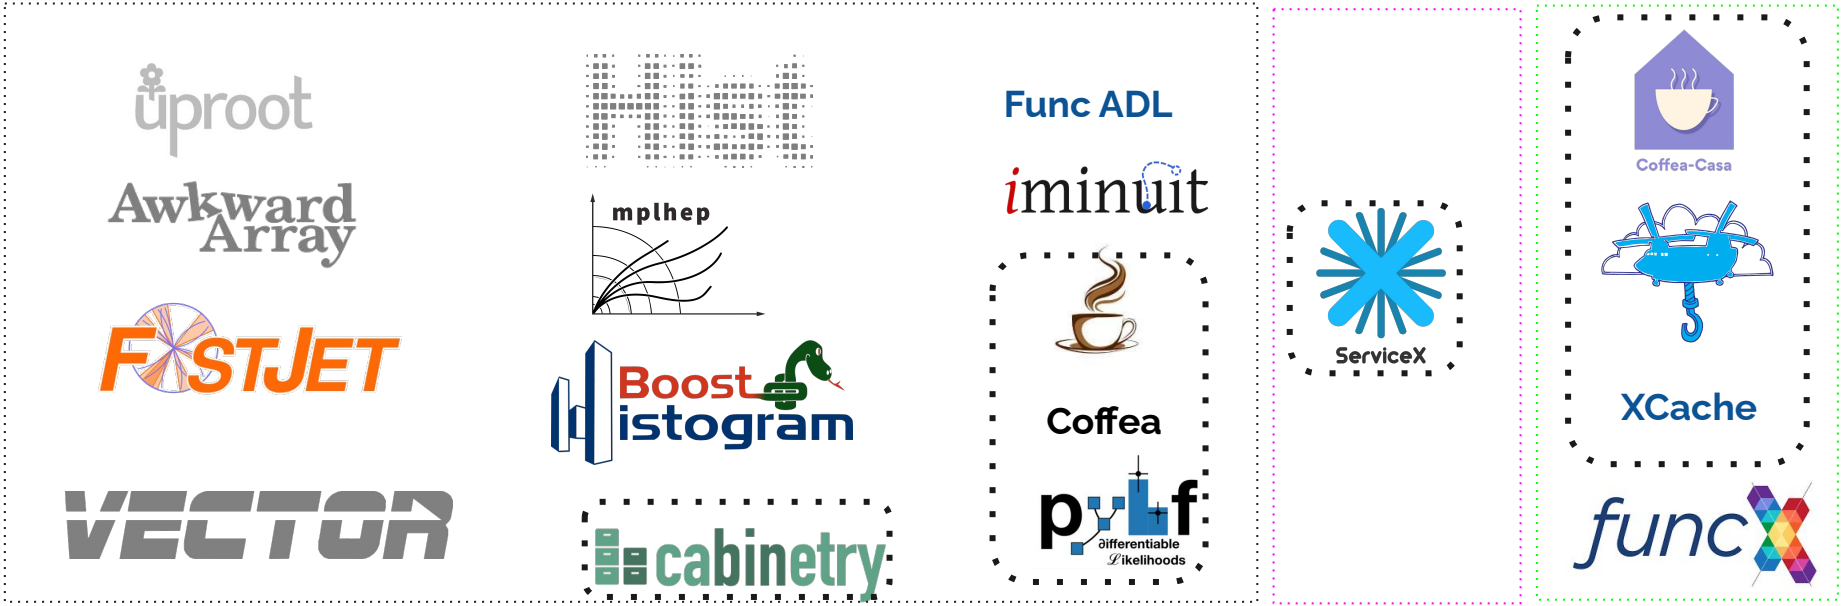

Analysis specific frameworks and packages (available in Docker container)

Data delivery service (k8s)

Optional services (k8s)

Requirements for AFs

Modern authentication (AIM/OIDC), tokens, macaroons, scitokens	Ongoing R&D on moving to use scitokens natively for AF (write/read)
Efficient data delivery and data management technologies	Ongoing work on integration ServiceX/Skyhook data delivery services
Columnar analysis and support new pythonic ecosystem	
Modern deployment and integration techniques	
Support for object storage	
Efficient data caching solutions	Integrating XCache in analysis pipeline
Easy integration with existing HPC resources	Looking to add support for other batch systems and task managements frameworks

[Coffea-casa team]

Analysis Grand challenge pipelines

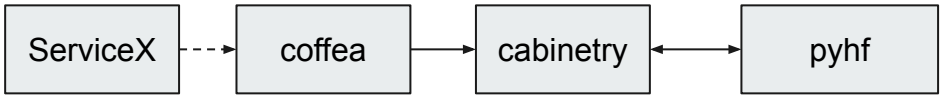

Columns from NanoAOD or flat ntuples from ServiceX or directly through coffea

Generic analysis pipeline based on *Open Data dataset* allowing to easily port AGC to other analysis frameworks

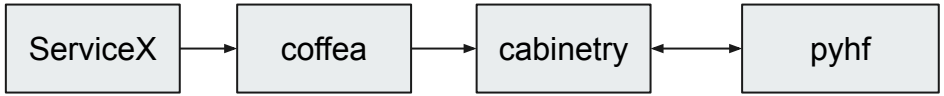

Columns from NanoAOD
and request column
from MiniAOD (only ServiceX)

CMS specific analysis pipeline based on *Open Data datasets* and *CMS datasets*

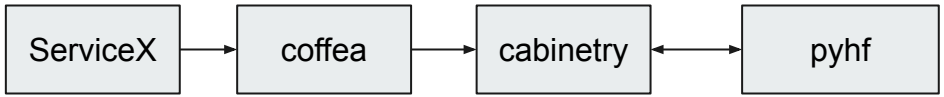

Columns from PHYSLITE
and request column
from PHYS (only ServiceX)

ATLAS specific analysis pipeline based on *ATLAS datasets*

Analysis Facilities participating in AGC (for now)

CMSAF @ T2 Nebraska
“Coffea-casa”
https://coffea.casa

OpenData AF @ T2 Nebraska
“Coffea-casa”
https://coffea-opendata.casa

ATLAS AF @ Scalable System Lab (UChicago)
“Coffea-casa”

New facility with ATLAS IAM, setting this up generated valuable feedback for future coffea-casa developments.

Elastic AF @ Fermilab (U.S. CMS)

FNAL team is participating in AGC, helping to test coffea-casa AF and re-using its components at the FNAL facility

U.S. ATLAS facilities
BNL / SLAC are participating in AGC
(we are evaluating the possibility to use coffea-casa experience)

Benchmarking & performance

- **How to benchmark AF performance** to demonstrate HL-LHC readiness?
  - **Multi-user environments**, where each user has a different workflow & requirements
    - Collection of “representative” analyses to simulate many users?
  - **`Data volume / (time * numCPUs)`**
    - Does not inherently account for I/O bound jobs where more CPUs != higher throughput
    - Assumption: `# branches read` matters, while `total # branches` does not, does this generally hold?
  - Fair to assume that **all input data is cached** (xCache)?
  - **Latency** is crucial for interactive workflows, how to best account for that?

- **Related questions**
  - How much control does a user have over caching (e.g. when re-calculation from scratch is expensive)?