



# Analysis Grand Challenge

Alex Held (NYU)  
Oksana Shadura (UNL)

IRIS-HEP / Ops Program Analysis Grand Challenge Planning  
March 1, 2022: <https://indico.cern.ch/event/1131750/>

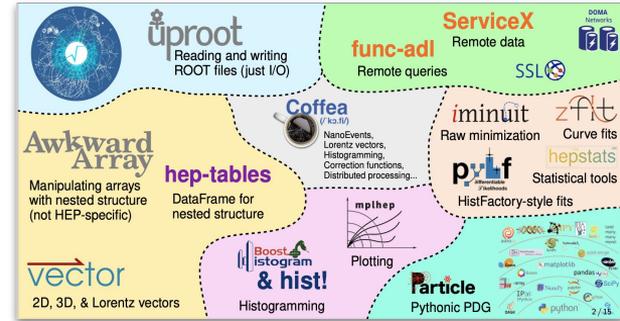
This work was supported by the U.S. National Science Foundation (NSF) Cooperative Agreement OAC-1836650 (IRIS-HEP).



# Analysis Grand Challenge

## Motivation:

- Allow coping with HL-LHC data sizes by rethinking data pipeline
  - Evaluating the new Python analysis ecosystem and integrating a differentiable analysis 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

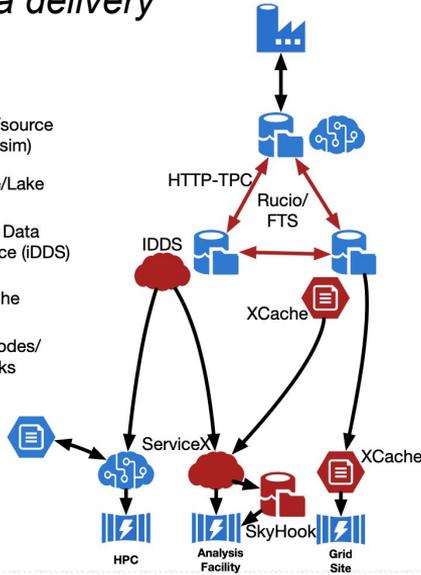
[Coffea-casa vCHEP 2021 plenary talk](#)

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.*

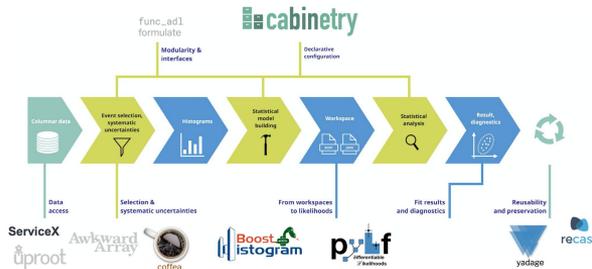
# AGC is connecting different IRIS-HEP focus areas

## DOMA: Data delivery

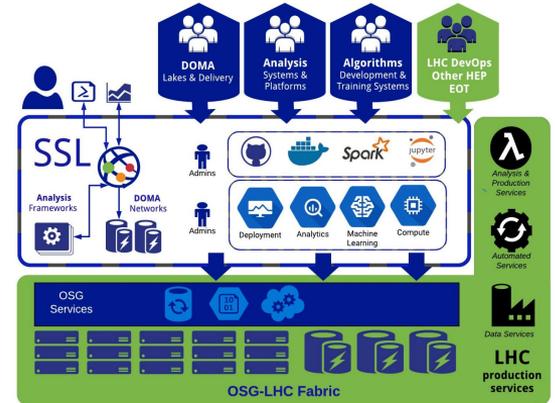
-  Data Factory/source (e.g. T0 or sim)
-  Data Store/Lake
-  Intelligent Data Delivery Service (iDDS)
-  Data Cache
-  Compute Nodes/Data Sinks



## AS: tools



## SSL: deployment techniques and resources



---

# Towards a benchmark analysis

# Towards a benchmark analysis: datasets

---

- Main AGC analysis example will be based on **Run-2 CMS Open Data**
  - CMS released ~400 TB of 2015 **miniAOD** at the end of 2021
  - Prefer to work with **nanoAOD** inputs
    - Results in realistic workflows as envisioned in ATLAS / CMS with PHYSLITE / nanoAOD
  - Eventually include upcoming second batch of 2015 Open Data (will be miniAOD + nanoAOD)
- No **official miniAOD -> nanoAOD conversion utility** available
  - In contact with CMS to understand longer-term options
- Using **PhysObjectExtractorTool** to build columnar output directly from miniAODs
  - Next step: running this via ServiceX transformer, then benchmarking performance
  - If conversion is a major bottleneck, may want to pre-convert to have nanoAOD-like input available

# Analysis selection



- Previously we thought about extending the [Run-1 Higgs->tautau example analysis](#)
  - Argument was that it may be simpler to get this approved by building on a public example
- New [Open Data release provides a lot of flexibility](#)
  - All major MC samples available for many analyses
  - Will do a [generic search in a ttbar phase space](#), likely modeled after existing public analysis
    - More familiarity with relevant objects / phase space / systematics / techniques
    - Possible synergies and collaboration with [Swift-HEP](#) / University of Manchester
- Developing an analysis from scratch gives us [flexibility](#)
  - Can e.g. easily showcase columnar kinematic reconstruction or MVAs, and all other relevant aspects
  - It also allows us to proceed [step by step](#): some aspects of this kind of analysis are quite generic, so can implement overall structure now and follow up with details later when they matter

---

# Getting involved

# How to participate



- Hoping to run **technical tests** at all interested sites to **evaluate compatibility with AGC plans**
  - Various examples testing pieces of the ecosystem are available:  
<https://github.com/CoffeaTeam/coffea-casa-tutorials>
- Proposed first step: **evaluate ServiceX setup** via the func\_adl + ServiceX example notebooks provided
  - Currently uses (public) CMS files, ATLAS to be added as well
  - Happy to help resolve issues, any feedback related to ease of setup would be great as well
    - Any particular things required that would simplify setup? Happy to iterate with you + ServiceX team
- Next step: **coffea without and with ServiceX**, and the various executors (including dask)
  - Examples also provided in repository
  - Expect that after this stage, remaining required pieces are going to be comparatively simple

# Feedback from SB meeting

---

- Very useful feedback at SB meeting two weeks ago (<https://indico.cern.ch/event/985528/>), thank you!
  - AGC efforts seem generally **aligned with CMS interests**
  - Envisioned workflow is **slightly more disconnected** from reality of current **ATLAS analysis patterns**
  - Also looking forward to LHCb / HSF feedback in next SB meeting

## ❖ Analysis Grand Challenge [Kaushik De's slides](#)

- Need to work with experiments to bring real users into AGC
  - **With well debugged end-to-end tools**
  - **Goals and plans defined in collaboration with ATLAS**

## ❖ Need algorithms and tools (and facilities) to work with common ATLAS data formats - PHYS and PHYSLITE

- Need to engage physicists doing analysis with hot data - while ATLAS putting a lot of effort on “Open Data,” it has lower physics priority
- Run 3 data is good for testing “aggressive” scenarios

# How to better engage ATLAS?

---

- **Happy to hear any suggestions here!**
- Closely following ATLAS efforts to push forward **PHYSLITE / columnar analysis**
  - Wide adoption of nanoAOD in CMS presumably is correlated with experiment adoption of columnar tools
- **PHYS vs PHYSLITE**
  - PHYSLITE may find some use in Run-3, but is mainly targeted at HL-LHC (as are many IRIS-HEP developments)
  - Trying to understand with ServiceX team what a PHYS transformer would entail
  - PHYS workflows are going to be different from PHYSLITE
    - We expect that PHYSLITE workflows will be very similar (though not identical) to nanoAOD setups
    - Expect that AGC analysis will be representative for nanoAOD / PHYSLITE-type analyses

---

# Related activities and upcoming events



## HSF AF forum

- Following suggestions pointing out the community need for an AF forum (in this meeting & similar ones), we proposed an **AF forum as a HSF activity area**
  - <https://hepsoftwarefoundation.org/activities/analysisfacilitiesforum.html> (see for mailing list, mattermost)
  - HSF provides the ideal neutral forum to bring together a rich community working on this topic
  - Four coordinators:
    - Diego Ciangottini (INFN, Perugia U, CMS)
    - Alessandra Forti (Manchester, WLCG/ATLAS)
    - Lukas Heinrich (TUM, ATLAS)
    - Nicole Skidmore (Manchester, LHCb)

# HSF AF kick-off event

- **Kick-off event** on March 25 <https://indico.cern.ch/event/1132360/>
- Idea: introduce activity area & context, very briefly show aspects of developments occurring
  - This cannot possibly capture all the ongoing work, but would hope to have dedicated follow-up meetings that can go into more detail
- Hoping that this activity area can become the place for AF-related discussions where the whole community can be involved
  - Expecting strong participation from IRIS-HEP

|       |  |  |
|-------|--|--|
| 15:00 | Welcome and Overview (10')                                   | Mark Neubauer<br>15:00 - 15:10                   |
|       | Introduction to the AF Forum (20+10')                        | Alessandra Forti et al.<br>15:10 - 15:40         |
|       | AFs in the context of the IRIS-HEP AGC (10+5')               | Alexander Held<br>15:40 - 15:55                  |
| 16:00 | DESY NAF (15+5') (TBC)                                       | 15:55 - 16:15                                    |
|       | Break  | 16:15 - 16:25                                    |
|       | SWAN over Spark and HTCondor at CERN (15+5')                 | Enric Tejedor Saavedra<br>16:25 - 16:45          |
|       | Coffea-Casa Facility at U Nebraska Tier-2 (15+5')            |  |
| 17:00 | Distributed Dask-based national facility at INFN (15+5')     | 16:45 - 17:05                                    |
|       | AF activities in LHCb (15+5') (tbc)                          | 17:05 - 17:25<br>17:25 - 17:45                   |
|       | Break  | 17:45 - 17:55                                    |
| 18:00 | AF activities in DOE multi-purpose computing centers (15+5') | 17:55 - 18:15                                    |
|       | Analysis on Cloud Facilities (15+5')                         | Fernando Harald Barreiro Megino<br>18:15 - 18:35 |
|       | A kubernetes-based AF at UChicago (10+5')                    | Lincoln Bryant<br>18:35 - 18:55                  |
| 19:00 | AF Activities on OSG (15+5')                                 | Brian Hua Lin<br>18:55 - 19:15                   |
|       | Discussion (30')   | 19:15 - 19:45                                    |

# Upcoming events



- HSF AF forum kick-off <https://indico.cern.ch/event/1132360/>
  - March 25
- IRIS-HEP AGC workshop: <https://indico.cern.ch/event/1126109/>
  - April 25-26, two afternoons CERN time with similar format as last time
  - Official announcement to follow
  - Will position this as something that produces input for HSF workshop below
- HSF analysis ecosystem workshop: <https://indico.cern.ch/event/1125222/>
  - May 23-25 in Paris
  - Oksana & Alex co-convening AF track and analysis user experience / declarative language track



## Next IRIS-HEP / Ops program AGC meeting

- Current time slot worked reasonably well according to poll results
- Proposal: next meeting on **April 5 at the same time** (9:00 PT / 11:00 CT / 12:00 ET / 18:00 CERN)
  - Would that generally work? Can otherwise run another poll

# Related IRIS-HEP Fellow proposals ([see more here](#))

- Application deadline: March 8, 2022

• **Enabling support for MiniAOD Transformer for ServiceX Data Delivery Service:** ServiceX is a distributed, cloud-native application that extracts columnar data from HEP event data and delivers it to an analyst. The func\_adl data query language is used to tell ServiceX how to extract the data (the columns, simple cuts, etc.). The func\_adl data query language has two backends that are currently part of ServiceX - one based on C++ for ATLAS data and CMS data, and one based on columnar processing using uproot and awkward arrays. The C++ backend currently runs only on the ATLAS binary format (xAOD) and CMS binary format (CMS AOD). This project will modify the C++ backend to also run on CMS MiniAOD binary files (available publicly as a part of [Run 2 CMS Opendata release](#)). The MiniAOD transformer is an important ingredient for a physics analysis workflow envisioned in the [Analysis Grand Challenge](#). (Contact(s): [Gordon Watts Ben Galewsky Oksana Shadura Alexander Held](#) )

• **Benchmarking of prototype analysis system components:** The [Analysis Grand Challenge](#) of IRIS-HEP focuses on performing a high energy physics analysis at scale, including all relevant features encountered by analyzers in this context. It is performed using tools and technologies developed within both IRIS-HEP and the broader community, making use of the Python ecosystem and the required cyberinfrastructure to run at scale. This project will happen after a first preliminary benchmarking has been performed, and it will build on that: the prospective fellow will use pieces of an example physics analysis to study the performance of different system components in more detail. Fellows are expected to have prior Python experience and interest in working with a diverse stack of analysis tools available in the ecosystem. (Contact(s): [Oksana Shadura Alexander Held](#) )

• **Metrics to define user activities and engagement on the various coffea-casa Analysis Facility deployments:** coffea-casa is a prototype of analysis facility (AF), which provides services for “low latency columnar analysis”, enabling rapid processing of data in a column-wise fashion. These services, based on Dask and Jupyter notebooks, aim to dramatically lower time for analysis and provide an easily-scalable and user-friendly computational environment that will simplify, facilitate, and accelerate the delivery of HEP results. The goal of the project is to define a set of various user engagement metrics, collected from Jupyterhub and other AF services, as well from underlying infrastructure (e.g. Kubernetes) and available through Elasticsearch. Expected results are the development of the various metrics, a data collection infrastructure for them, and possibly visualization dashboards. (Contact(s): [Brian Bockelman Ken Bloom Oksana Shadura](#) )



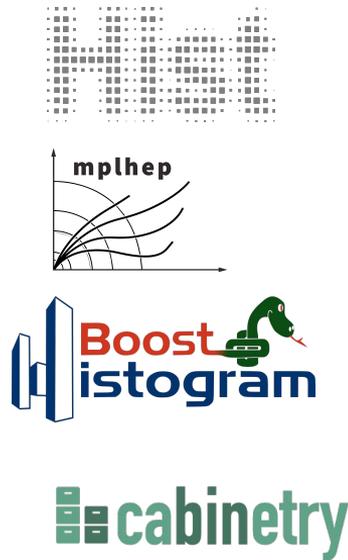
# Summary

- Made progress with technical aspects of **handling CMS Open Data**
  - Have a way forward towards benchmarking milestone this summer
  - Now working towards shaping this into a benchmark analysis
- Hope to be able to run **technical tests at all available sites**
  - Described first step to evaluate compatibility of setups with AGC requirements
- Range of **interesting upcoming events**
  - HSF AF forum kick-off
  - Next AGC workshop & HSF analysis ecosystem workshop

---

# Backup slides

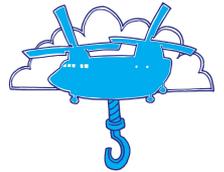
# Expanding analysis pipeline: software components



Func ADL

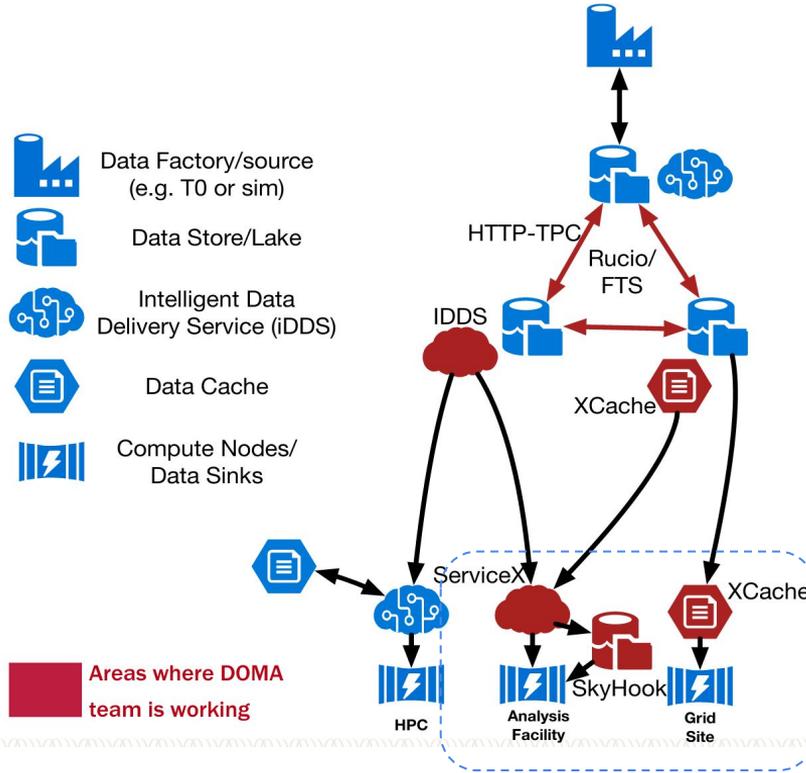


iminuit

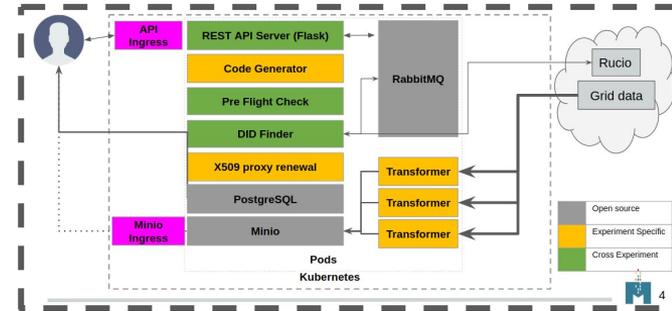


func

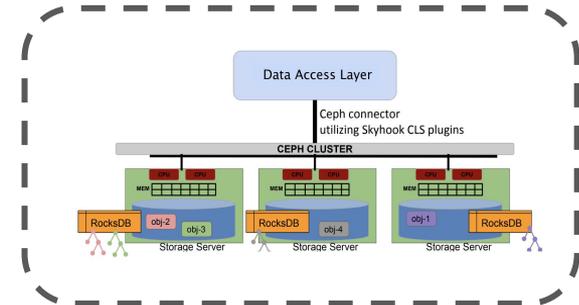
# Analysis Facility and Distributed Ecosystem (Data Lakes)



Coffea-casa AF



ServiceX

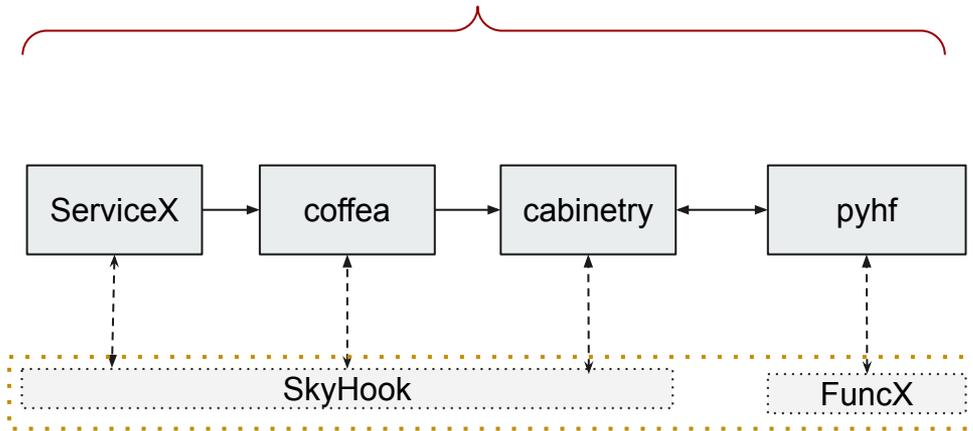


Skyhook

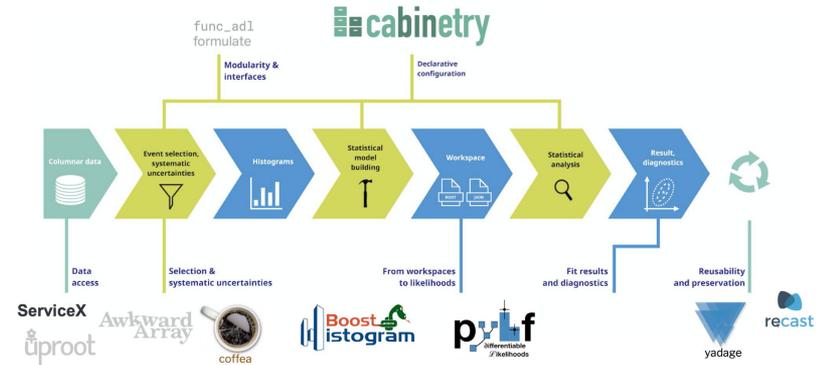
# Expanding existing analysis pipeline

- Demonstration of **ServiceX -> Skyhook -> coffea -> cabinetry -> pyhf** pipeline on Open Data

milestone goal for May 1, 2022

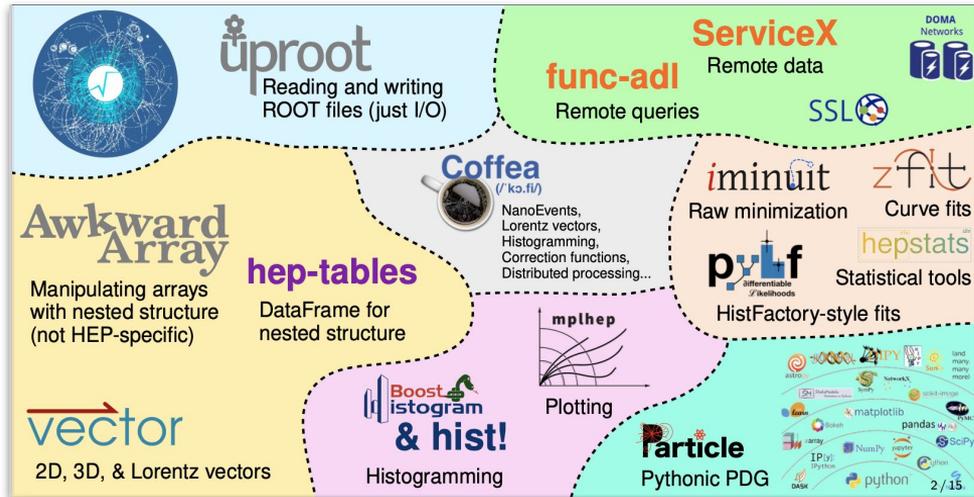


not included in 2021 workshop demonstration

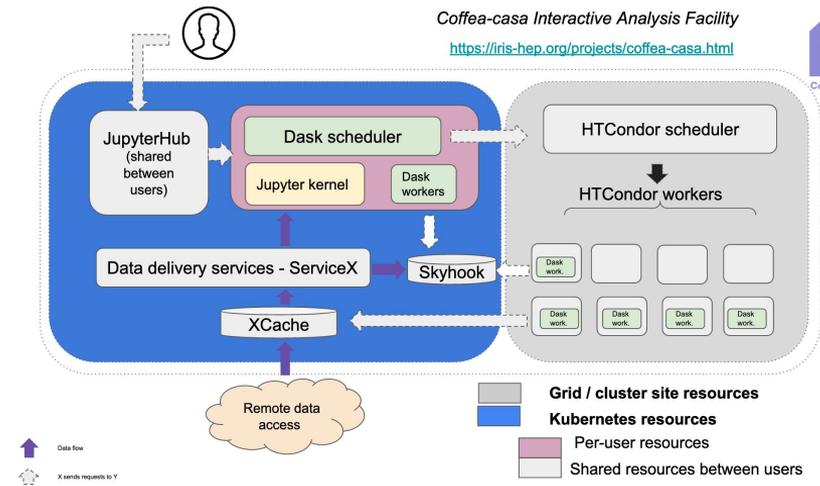


# Building blocks used for designing AFs

## Analysis Tools



## Analysis Facilities



# Requirements for AFs



Modern authentication (AIM/OIDC), tokens, macaroons, scitokens

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

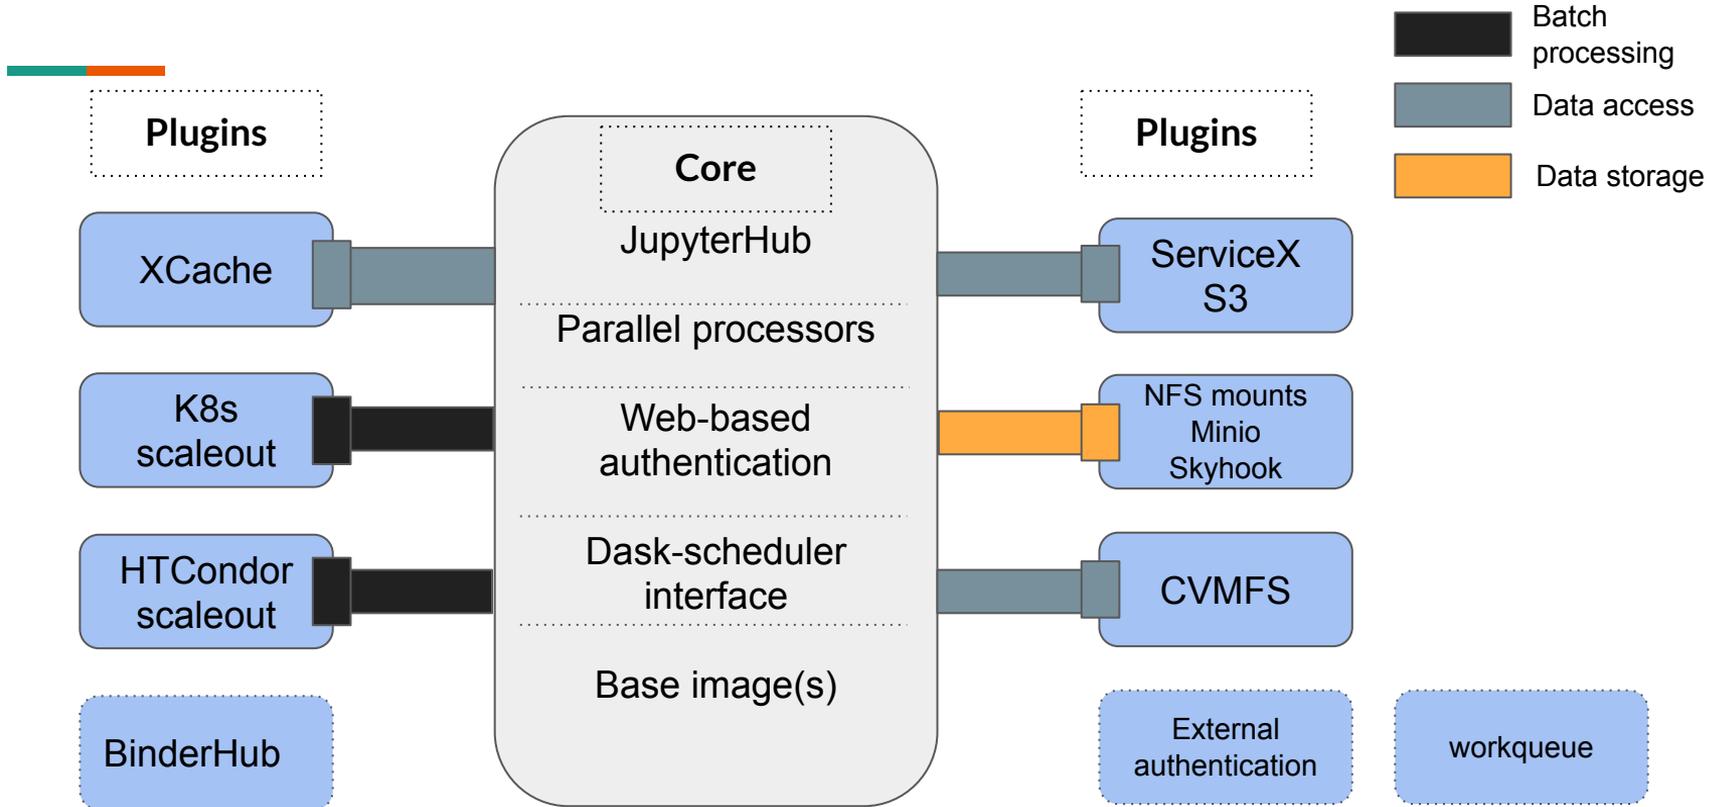
Ongoing R&D on moving to use scitokens natively for AF (write/read)

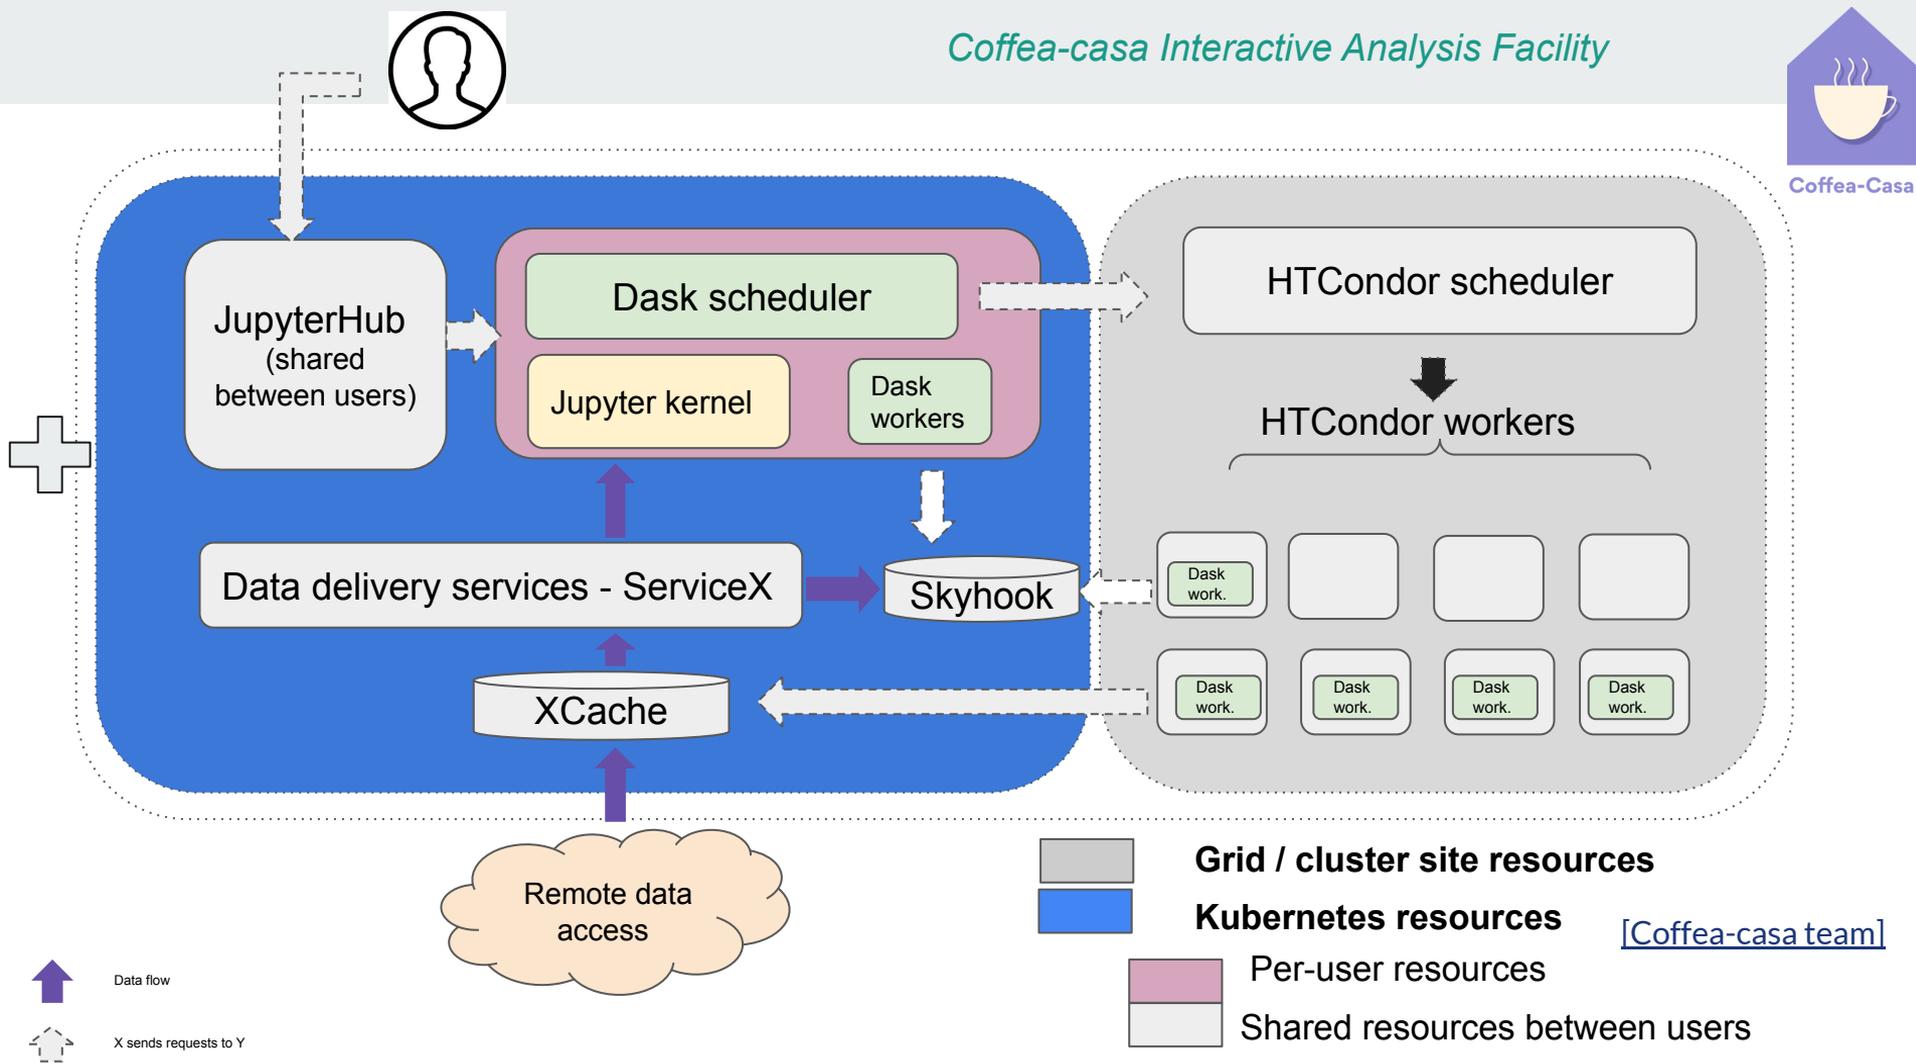
Ongoing work on integration ServiceX/Skyhook data delivery services

Integrating XCache in analysis pipeline

Looking to add support for other batch systems and task managements frameworks

# Designing AF: components of Coffea-casa Analysis Facility





↑ Data flow  
⤵ X sends requests to Y