The 200 Gbps Challenge: Imagining HL-LHC analysis facilities

Sam Albin [1], Garhan Attebury [1], Ken Bloom [1], Brian Paul Bockelman [2], Lincoln Bryant [3],
Kyungeon Choi [4], Kyle Cranmer [5], Peter Elmer [6], Matthew Feickert [5], Rob Gardner [3], Lindsey Gray [7],
Alexander Held [5], Fengping Hu [3], David Lange [6], Carl Lundstedt [1], Peter Onyisi [4], Jim Pivarski [6],
Oksana Shadura [1], Nick Smith [7], John Thiltges [1], Ben Tovar [8], Ilija Vukotic [3], Gordon Watts [9],
Derek Weitzel [1], Andrew Wightman [1]

[1] University of Nebraska–Lincoln, [2] Morgridge Institute for Research, [3] University of Chicago,
[4] University of Texas at Austin, [5] University of Wisconsin-Madison, [6] Princeton University,
[7] Fermilab, [8] University of Notre Dame, [9] University of Washington

CHEP 2024

https://indico.cern.ch/event/1338689/

Oct 21, 2024

This work was supported by the U.S. National Science Foundation (NSF) under Cooperative Agreements OAC-1836650 and PHY-2323298.

CHEP 2004 and 2024

- Until the mid-1980s HEP's "computing problem" was often thought to be about obtaining enough processor power
- Then we worried about storage capacity
- The real problem has always been, in my opinion, getting people to collaborate on a solution

[David Williams: "50 years of Computing at CERN ", CHEP 2004]

CHEP'04

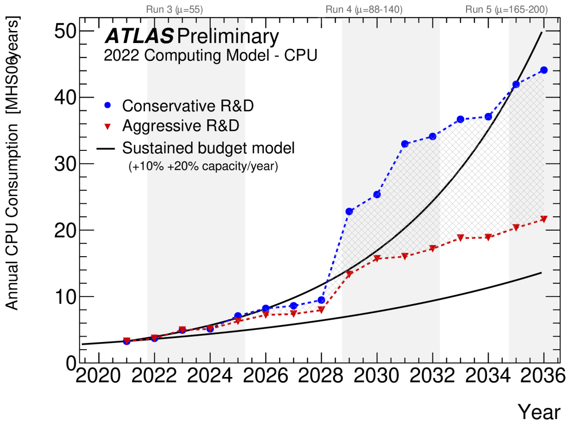

CERN-LHCC-2022-005

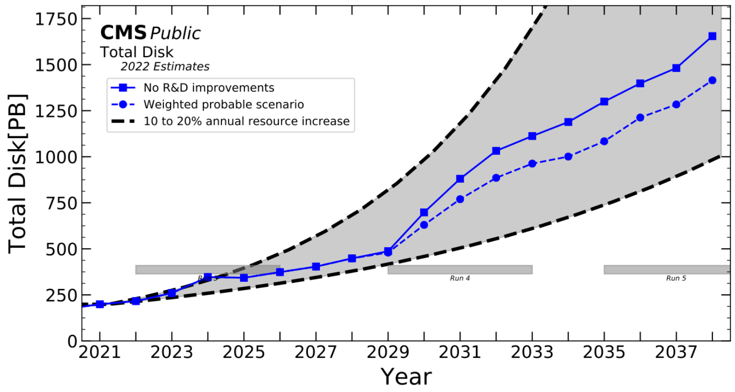

CERN-CMS-NOTE-2022-008

IRIS-HEP and a HL-LHC vision

The IRIS-HEP software institute

- **IRIS-HEP**: *"Institute for Research and Innovation in Software for High Energy Physics"*
  - a software institute funded by the US National Science Foundation, running 2018–2028
  - working in close collaboration with LHC experiments and facilities

U.S. National Science Foundation

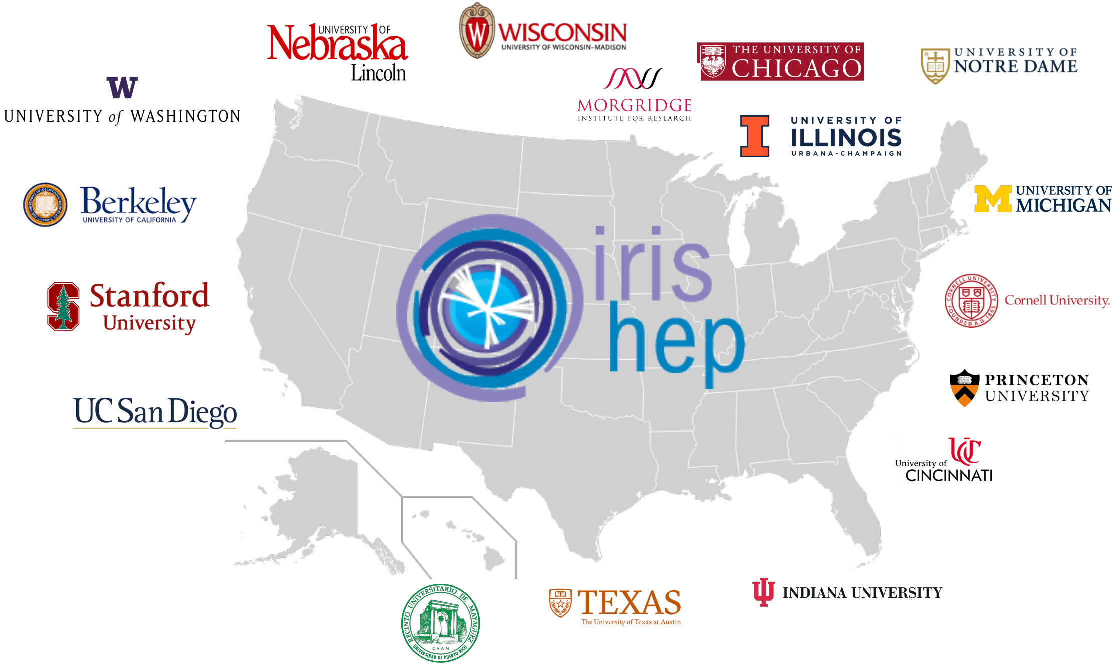

[2024 IRIS-HEP retreat]

R&D for the HL-LHC

- IRIS-HEP is working on **computing and software R&D for the HL-LHC**
  - a software upgrade accompanying detector hardware upgrades
  - focusing on a subset of activity areas today, connected through “challenge” format

DOMA
data organisation and management

ServiceX

XCache

Coffea-Casa

WLCG Data Challenge

AS
analysis systems and tools

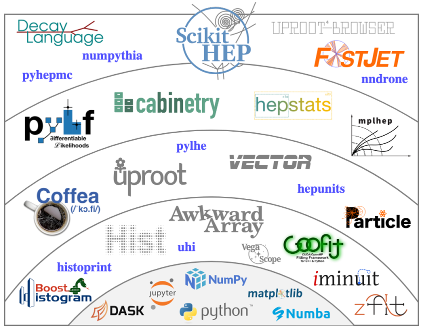

SSL and OSG-LHC
deployment techniques and resources

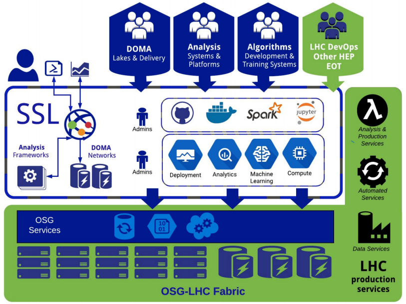

SSC
training

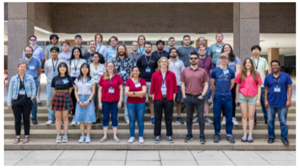

Empowering physicists, today and tomorrow

- This work is driven by the desire to **minimize time-to-insight** and **maximize the HL-LHC physics reach**
  - let **physicists** spend more time doing physics, less time debugging, bookkeeping, waiting, …
  - tighten feedback & support cycles by connecting communities together
- **Physics is the end goal**: strive to find ways to overcome computing challenges

IRIS-HEP is active in all these areas

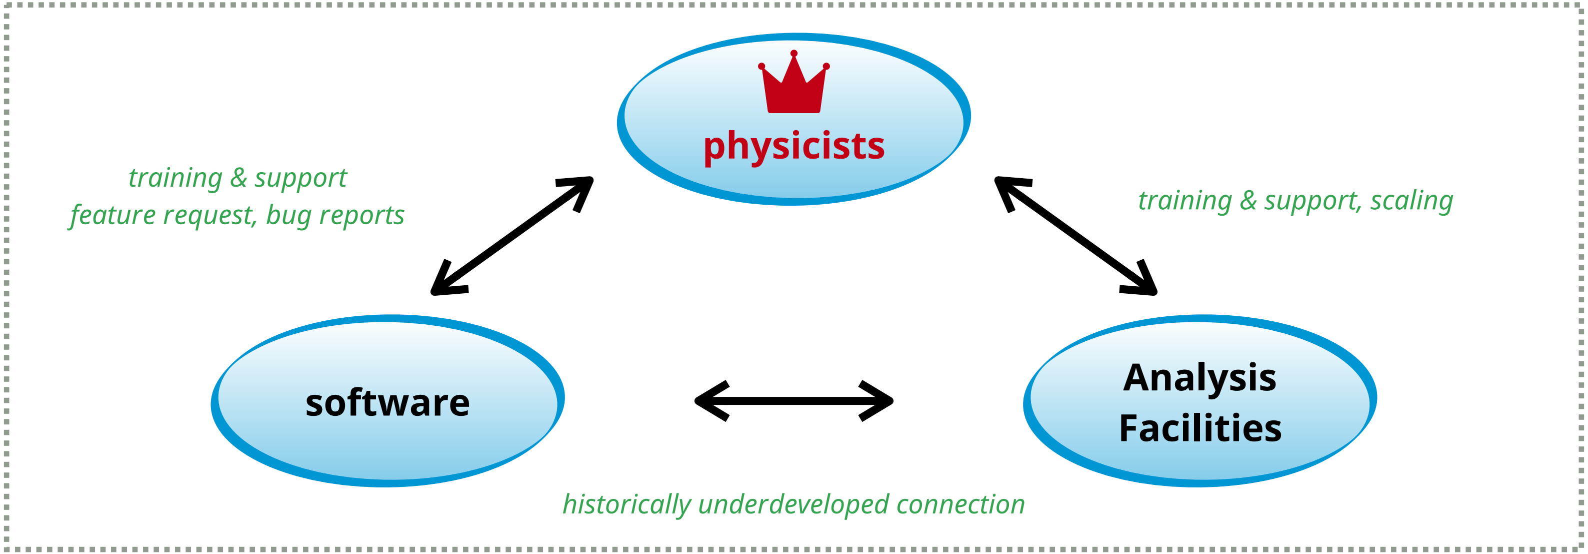

Our end-user analysis vision

- **Analyze O(1000) TB of data within a few hours**
- Interactive analysis turnaround: “coffee break” timescale
- Fully integrated Analysis Facilities (AFs)
- UX to empower big & small teams
- Easy access to state-of-the-art ML + techniques
- Reproducibility, preservation, reuse

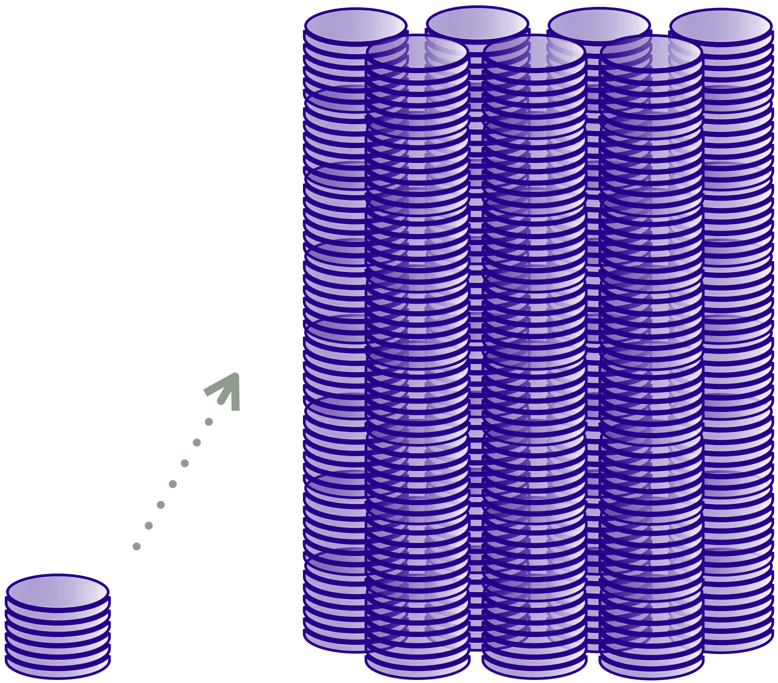

today:

create $\mathcal{O}(1-10)$ TB ntuples
on the grid
in $\mathcal{O}$(days – weeks),
analyze on Tier-3 in $\mathcal{O}(h$ – days)

HL-LHC:

analyze $\mathcal{O}(1000)$ TB of data
straight out of central
PHYSLITE / NanoAOD files in $\mathcal{O}(h)$

Our end-user analysis vision

- Analyze O(1000) TB of data within a few hours
- **Interactive analysis turnaround: “coffee break” timescale**
- Fully integrated Analysis Facilities (AFs)
- UX to empower big & small teams
- Easy access to state-of-the-art ML + techniques
- Reproducibility, preservation, reuse

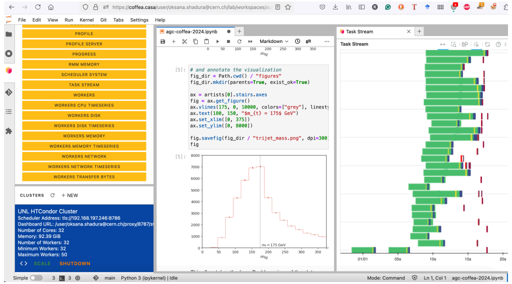

meaningful analysis iterations on timescale of a coffee break, interactive analysis design

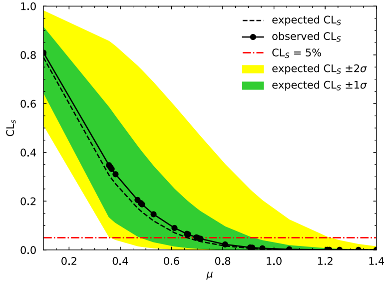

Our end-user analysis vision

- Analyze O(1000) TB of data within a few hours
- Interactive analysis turnaround: “coffee break” timescale
- **Fully integrated Analysis Facilities (AFs)**
- UX to empower big & small teams
- Easy access to state-of-the-art ML + techniques
- Reproducibility, preservation, reuse

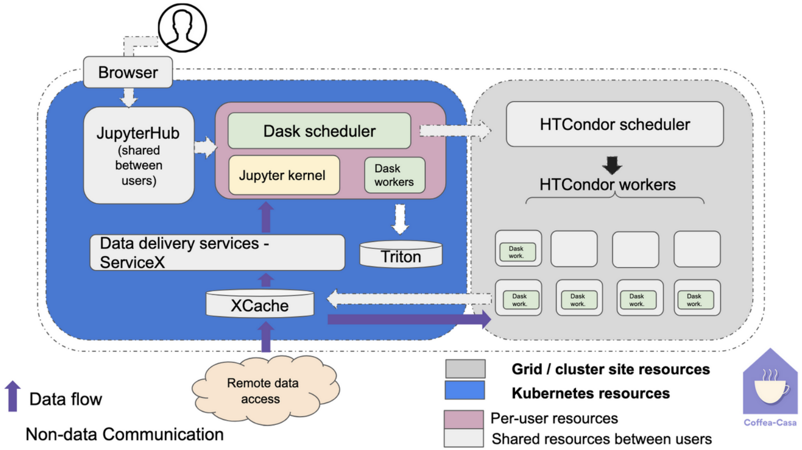

required services available,
convenient interfaces,
access to powerful resources

Our end-user analysis vision

- Analyze O(1000) TB of data within a few hours
- Interactive analysis turnaround: “coffee break” timescale
- Fully integrated Analysis Facilities (AFs)
- **UX to empower big & small teams**
- Easy access to state-of-the-art ML + techniques
- Reproducibility, preservation, reuse

efficient collaboration within analysis team

data access & output sharing

software environment handling

fast turnaround

seamless laptop-to-AF/grid transition

see also: ***HSF AF White Paper***
https://arxiv.org/abs/2404.02100

Our end-user analysis vision

- Analyze O(1000) TB of data within a few hours
- Interactive analysis turnaround: “coffee break” timescale
- Fully integrated Analysis Facilities (AFs)
- UX to empower big & small teams
- **Easy access to state-of-the-art ML + techniques**
- Reproducibility, preservation, reuse

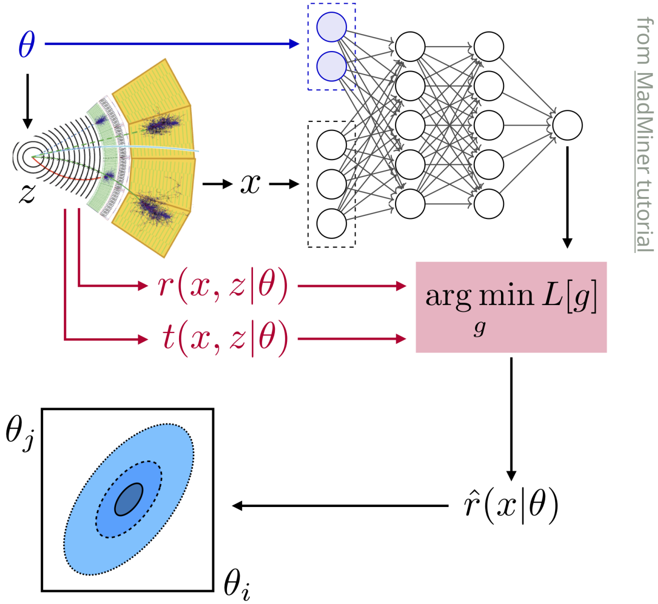

simulation-based inference techniques use different workflows from traditional histogram-based approaches

Our end-user analysis vision

- Analyze O(1000) TB of data within a few hours
- Interactive analysis turnaround: “coffee break” timescale
- Fully integrated Analysis Facilities (AFs)
- UX to empower big & small teams
- Easy access to state-of-the-art ML + techniques
- **Reproducibility, preservation, reuse**

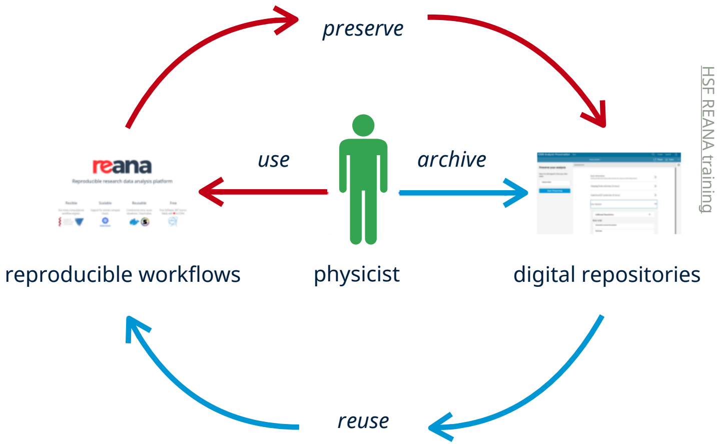

sustainable research
maximizing long-term impact and legacy

see also:
Nature 533, 452–454 (2016)
arXiv:2203.10057 [hep-ph]

The Analysis Grand Challenge (AGC)

A test case for HL-LHC analysis

- The **Analysis Grand Challenge (AGC)** defines a **physics analysis task** with **Open Data** to test **HL-LHC workflows**
  - columnar data extraction from large datasets & data processing into histograms
  - statistical model construction and statistical inference, relevant visualizations
  - ML training & inference

CERN OPEN DATA PORTAL

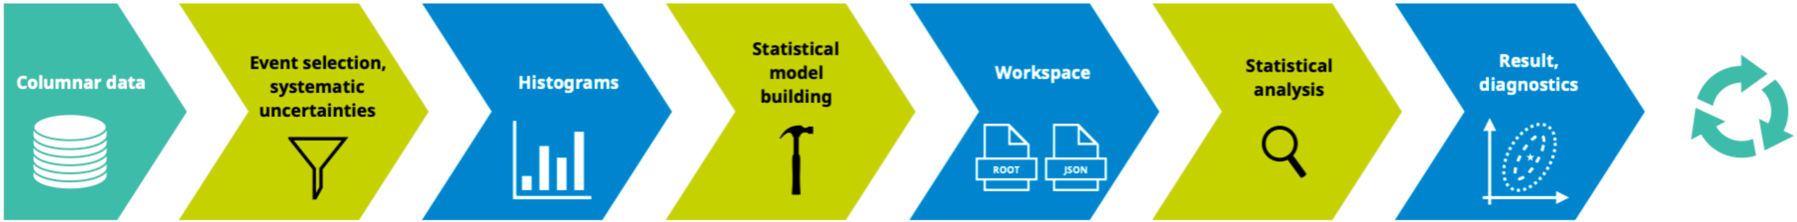

Analysis with task graphs and dask

- We employ **task graphs** to **express & execute** data analysis operations
- This relies on dask, a **Python library** providing
  - an interface to describe manipulations of data via task graphs
  - a task scheduler to execute task graphs
- **Deep integration of Dask** and existing **Python HEP toolset** with minimal API changes
  - arrays via dask-awkward, histograms, coffea etc.
  - Dask emerging as common feature in Analysis Facilities

example of a task graph

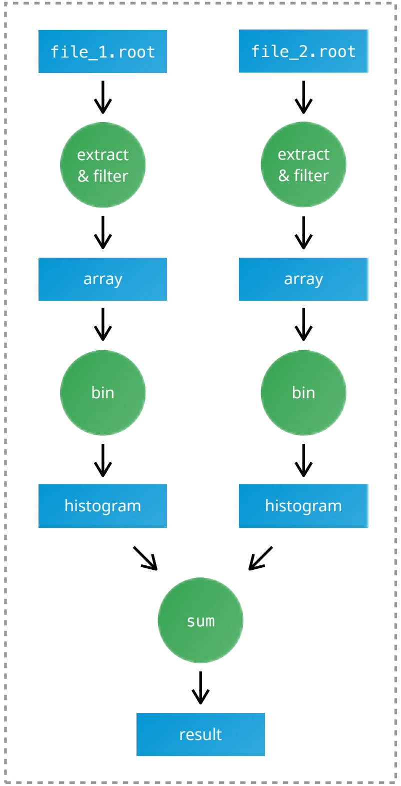

The IRIS-HEP AGC reference implementation

The **IRIS-HEP reference implementation** employs the **Scikit-HEP/ PyHEP ecosystem** and serves as **ideal environment** to test our **latest R&D**.

find it all on GitHub and https://agc.readthedocs.io/

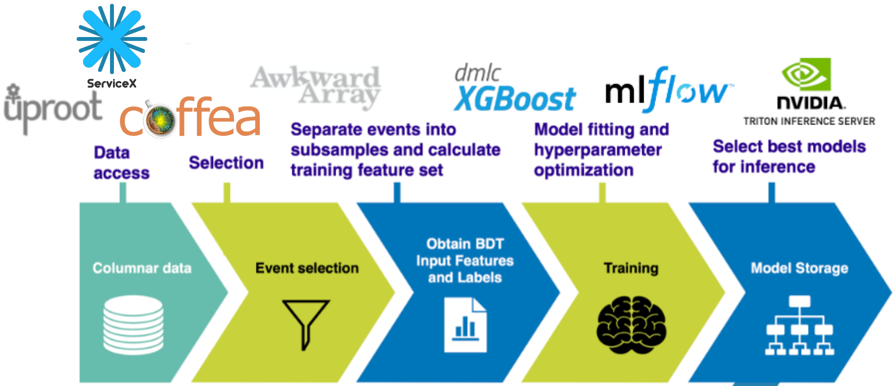

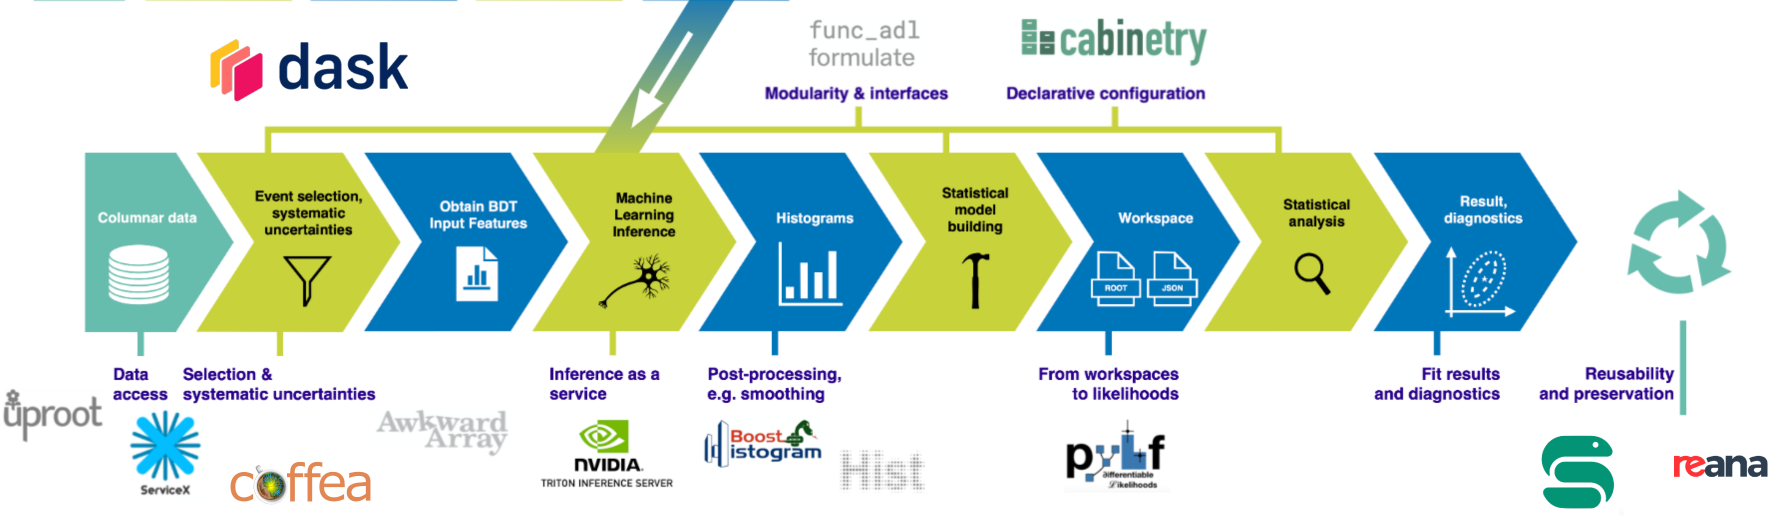

A community project

- **Variety of AGC implementations** have been developed, more are welcome!

- **Regularly hosting related events**
  - AGC workshops (#1, #2, #3, #4)
  - "Demo day" event series

AGC workshop 2023

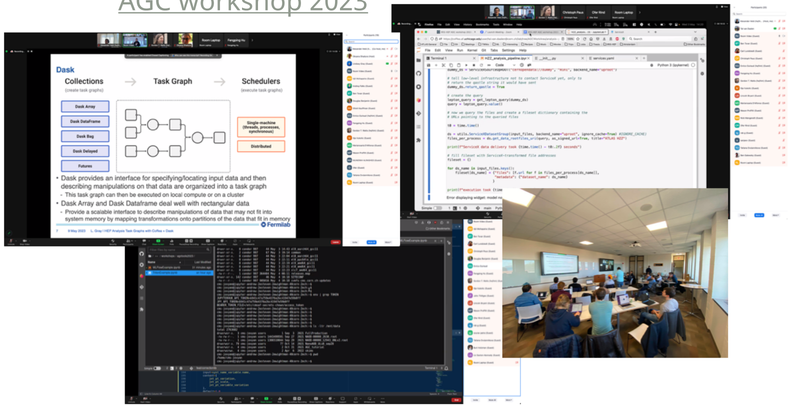

- **AGC keeps evolving:** wishlist for future additions

The 200 Gbps Challenge

Breaking down physics analyses

- Currently **limited agreement** on how **"HL-LHC physics analyses"** will look
  - factorize into independent challenges, push the boundaries in all directions

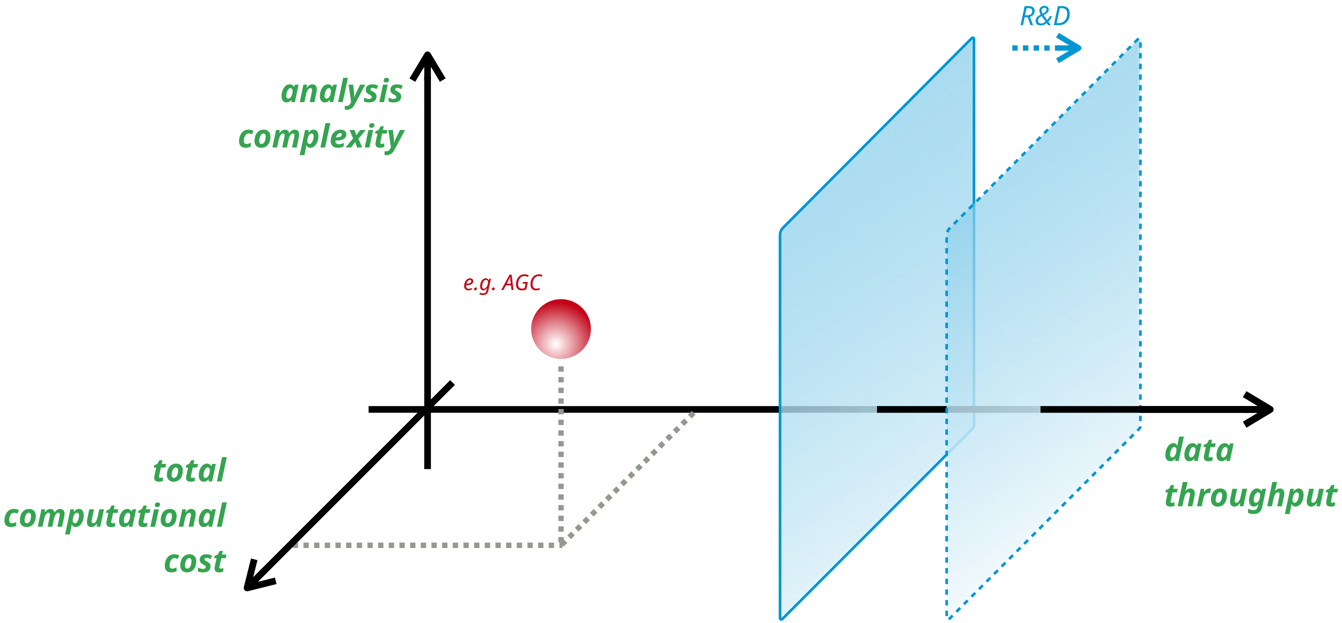

Defining the 200 Gbps Challenge

200 Gbps →

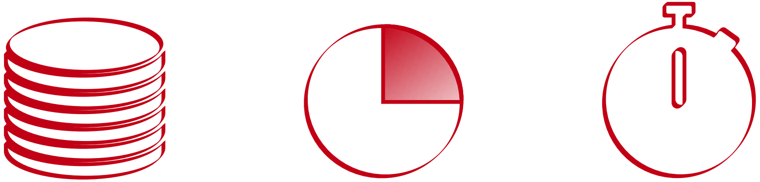

180 TB dataset | *25% of file content* | *30 minutes*

Reaching these scales poses significant challenges. We set ourselves an ambitious goal.

… and had only 8 weeks to reach it.

CMS NanoAOD example

With **2 kB / event**, this means

- **90 B events,**
- **50 MHz event rate,**

or 1k cores with 50 kHz and 25 MB/s each.

ATLAS PHYSLITE example

With **10 kB / event**, this means

- **18 B events,**
- **10 MHz event rate,**

or 2k cores with 5 kHz and 12.5 MB/s each.

Defining the 200 Gbps Challenge

200 Gbps →

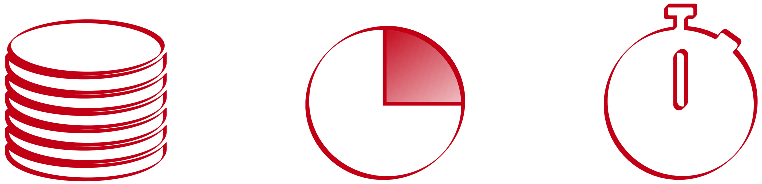

180 TB dataset *25% of file content* *30 minutes*

- Targeting **"HL-LHC scale" analysis**, including decompression & data in memory as arrays

- **Two different setups, targeting realism**, all code on GitHub
  - Nebraska: analyze Run-3 NanoAOD CMS data (iris-hep/idap-200gbps)
  - UChicago: analyze Run-2 PHYSLITE ATLAS data (iris-hep/idap-200gbps-atlas)
  - similar tasks broadly, important differences: facilities, event sizes, object types, compression, …

Ingredients for 200 Gbps throughout

team of experts from IRIS-HEP and beyond,
rallied behind a shared vision

planning, structure, schedule

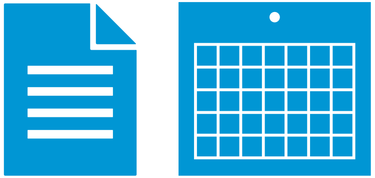

100s of messages per day,
dedicated communication channels

slack

[image by DALL·E 3]

challenging timeline: 8 weeks
from first idea to WLCG/HSF workshop

dedicated meetings in multiple formats

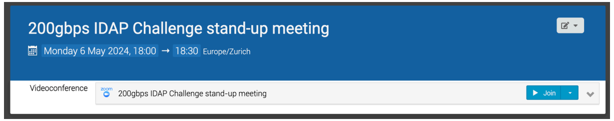

Key Analysis Facility elements for 200 Gbps

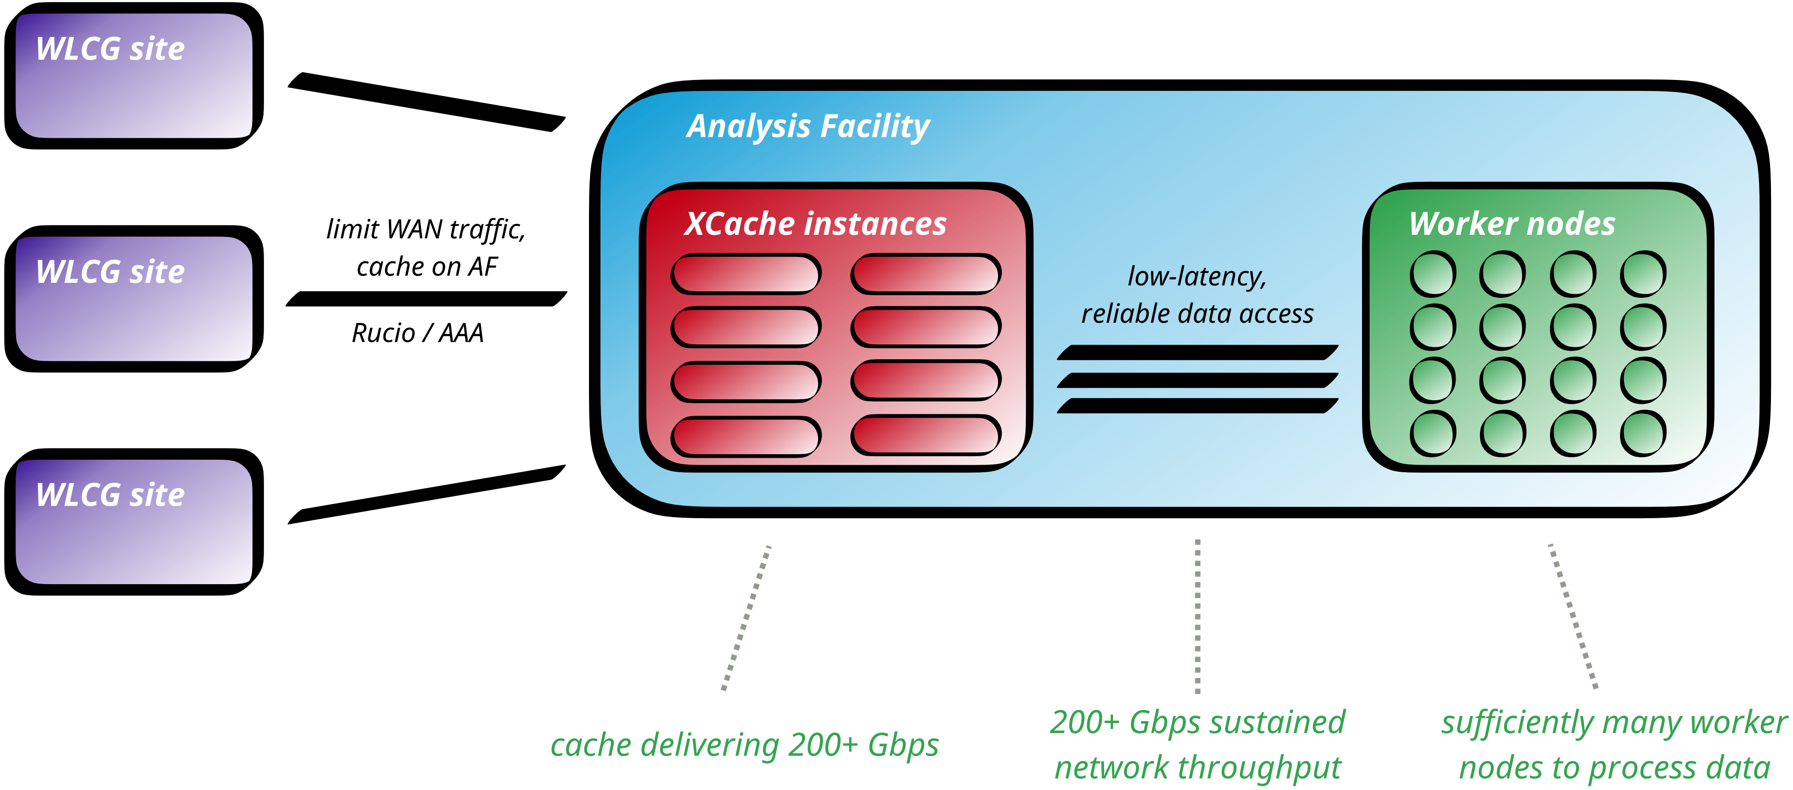

see Oksana Shadura's talk

Coffea-casa at Nebraska: 200 Gbps setup

developed in close collaboration

- **R&D prototype of a future Analysis Facility**
  - designed as hybrid setup including Kubernetes and Nebraska CMS Tier-2 / Tier-3 resources

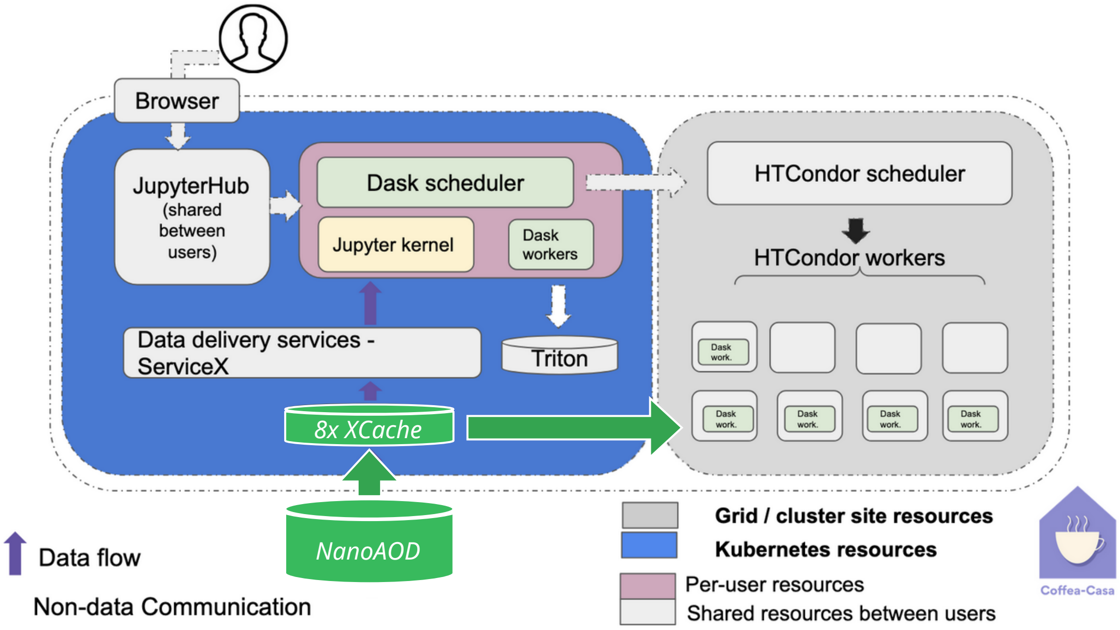

using 8 XCache instances behind 2x100 Gbps uplink each

Coffea-casa at Nebraska: measurements

- **200 Gbps sustained throughput of data for physics**
  - scheduling with Dask & TaskVine, scaling with HTCondor & Kubernetes
  - re-compressed NanoAOD (LZMA → ZSTD) for 2.5x event rate increase

details for this example:
- *40 B events, 64k files*
- *1304 workers*
- *32 MHz event rate*
- *data processed (compressed): 30 TB*
- *data processed (uncompressed): 71 TB*

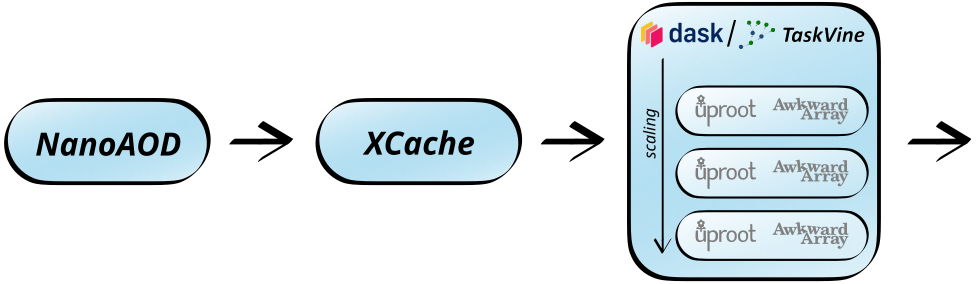

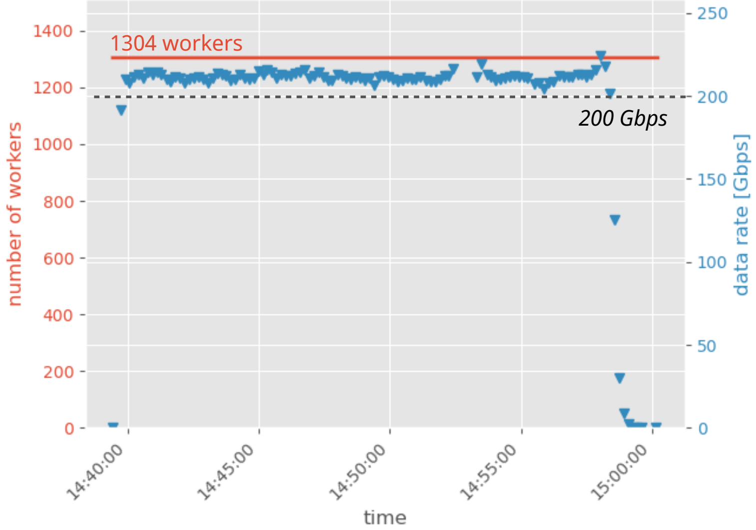

200+ Gbps with Dask + HTCondor

UChicago AF: 200 Gbps setup

see Rob Gardner's talk

developed in close collaboration

US ATLAS iris hep

- **Production Analysis Facility for ATLAS**
  - built on Kubernetes, partially reconfigured for needs of challenge

- **Two configurations explored** with Kubernetes scaling (HTCondor available)
  - `uproot` scaled with Dask reading from XCache
  - ServiceX as data delivery service writing to object storage

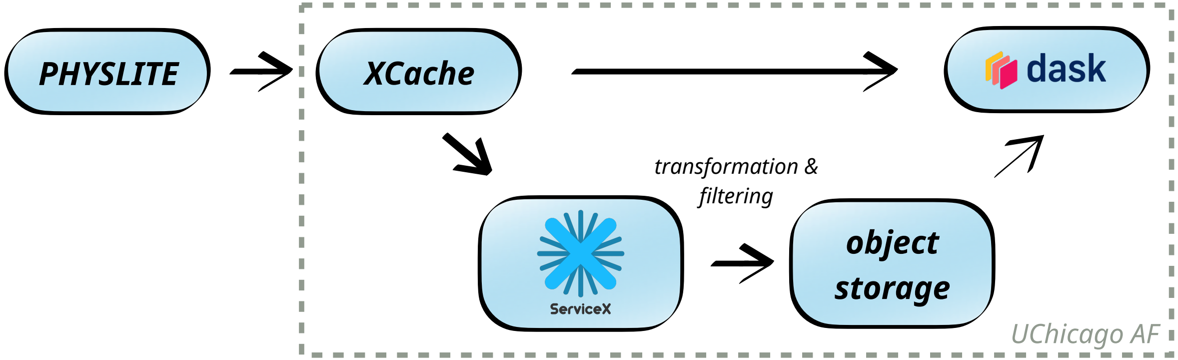

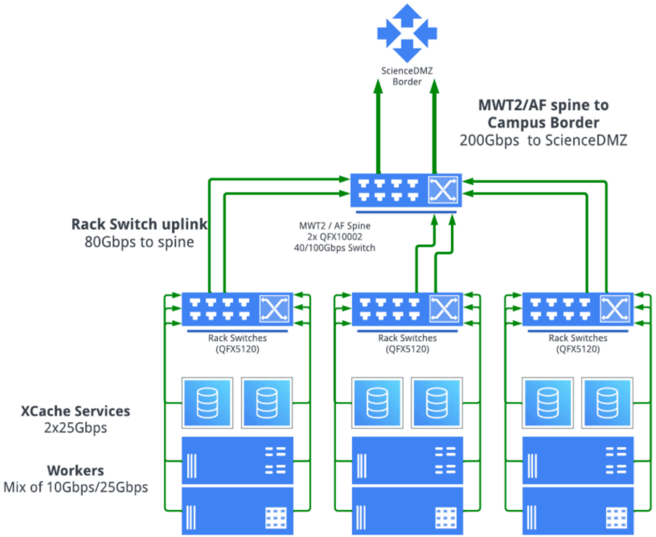

8 XCache instances total, distributed to maximize bandwidth

UChicago AF: measurements

- **200 Gbps sustained throughput of data for physics**

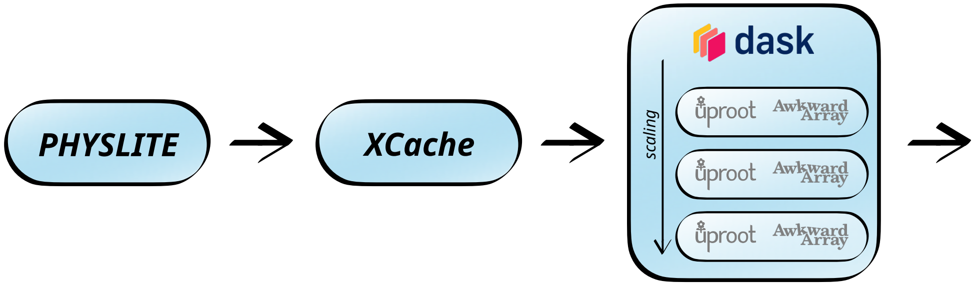

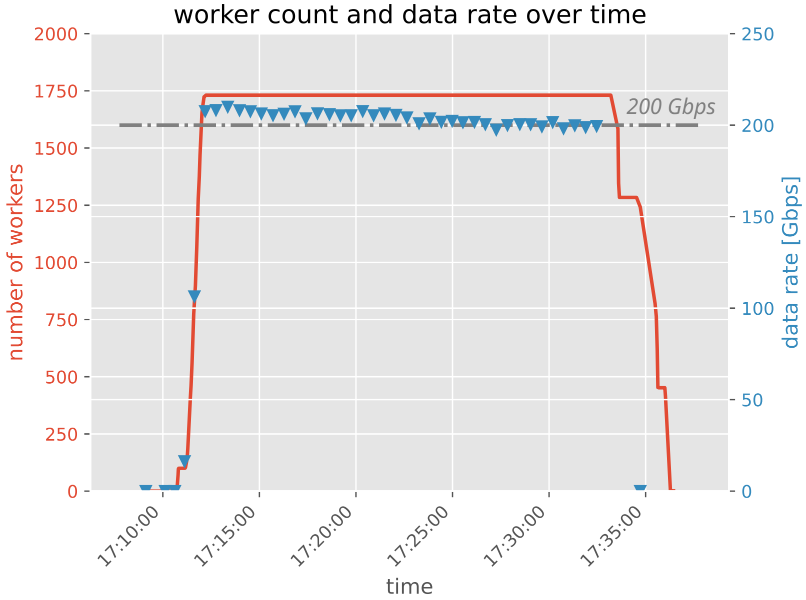

network monitoring

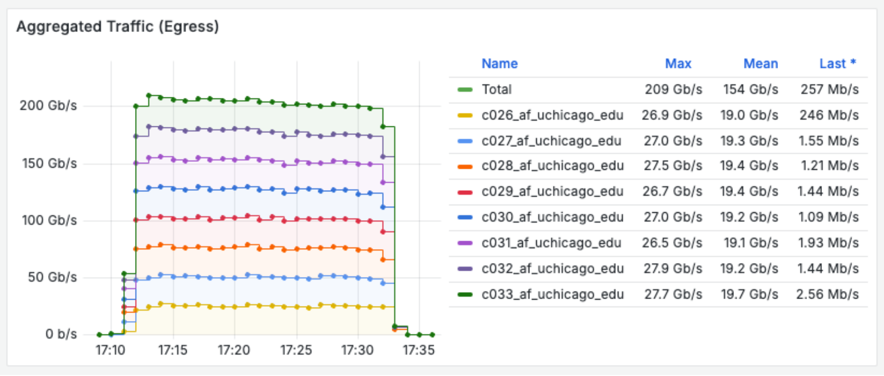

Name	Max	Mean	Last *
Total	209 Gb/s	154 Gb/s	257 Mb/s
c026_af_uchicago_edu	26.9 Gb/s	19.0 Gb/s	246 Mb/s
c027_af_uchicago_edu	27.0 Gb/s	19.3 Gb/s	1.55 Mb/s
c028_af_uchicago_edu	27.5 Gb/s	19.4 Gb/s	1.21 Mb/s
c029_af_uchicago_edu	26.7 Gb/s	19.4 Gb/s	1.44 Mb/s
c030_af_uchicago_edu	27.0 Gb/s	19.2 Gb/s	1.09 Mb/s
c031_af_uchicago_edu	26.5 Gb/s	19.1 Gb/s	1.93 Mb/s
c032_af_uchicago_edu	27.9 Gb/s	19.2 Gb/s	1.44 Mb/s
c033_af_uchicago_edu	27.7 Gb/s	19.7 Gb/s	2.56 Mb/s

more details:

- *32 B events, 190 TB data, 218k files*
- *1739 workers peak*
- *15 MHz event rate, 5–20 kHz per core*
- *200 Gbps throughput sustained*
- *data processed (compressed): 32 TB*
- *data processed (uncompressed): 80 TB*

ServiceX as data delivery service

- **Idea: filter data with ServiceX**, then further process output with Dask
  - rapid turnaround from cached ServiceX output

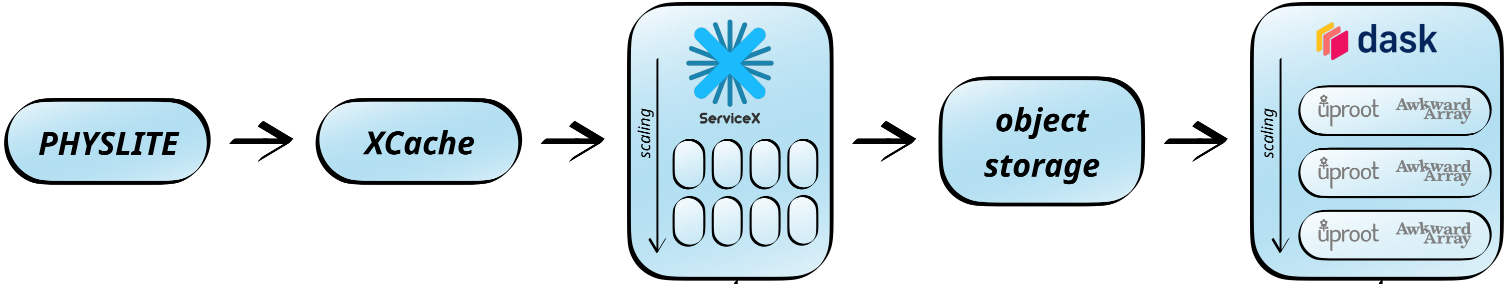

- *19 B events, 146 TB data*
- *up to 1k pods*
- *10 MHz event rate*

- *1 TB data (skimmed)*
- *500 worker nodes*
- *10 MHz event rate*

170 Gbps parallel data processing with ServiceX

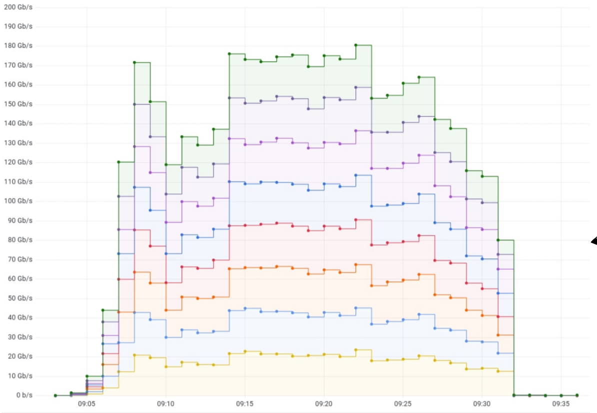

multi-stage processing schema, transparent to users

Dask tasks

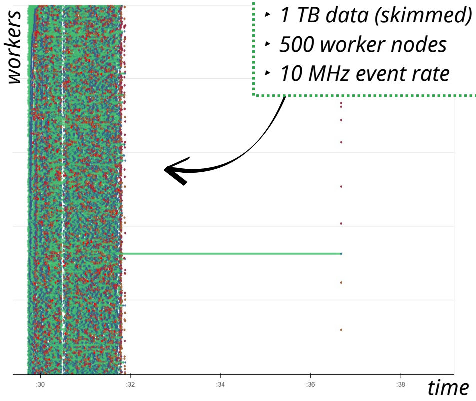

Towards multi-user interactivity

- Analysis Facilities will host **many users** looking to achieve **interactive turnaround simultaneously**
  - ensure scale testing includes number of users
- **Live exercise** at 2024 IRIS-HEP retreat: ***200 Gbps setup with 16 participants***
  - automatic CPU scaling with Dask
  - limited maximum number of cores per user
- Intended as part of a bigger **discussion about fair-share & interactivity**

live demonstration with retreat participants

Bandwidth shared between multiple users

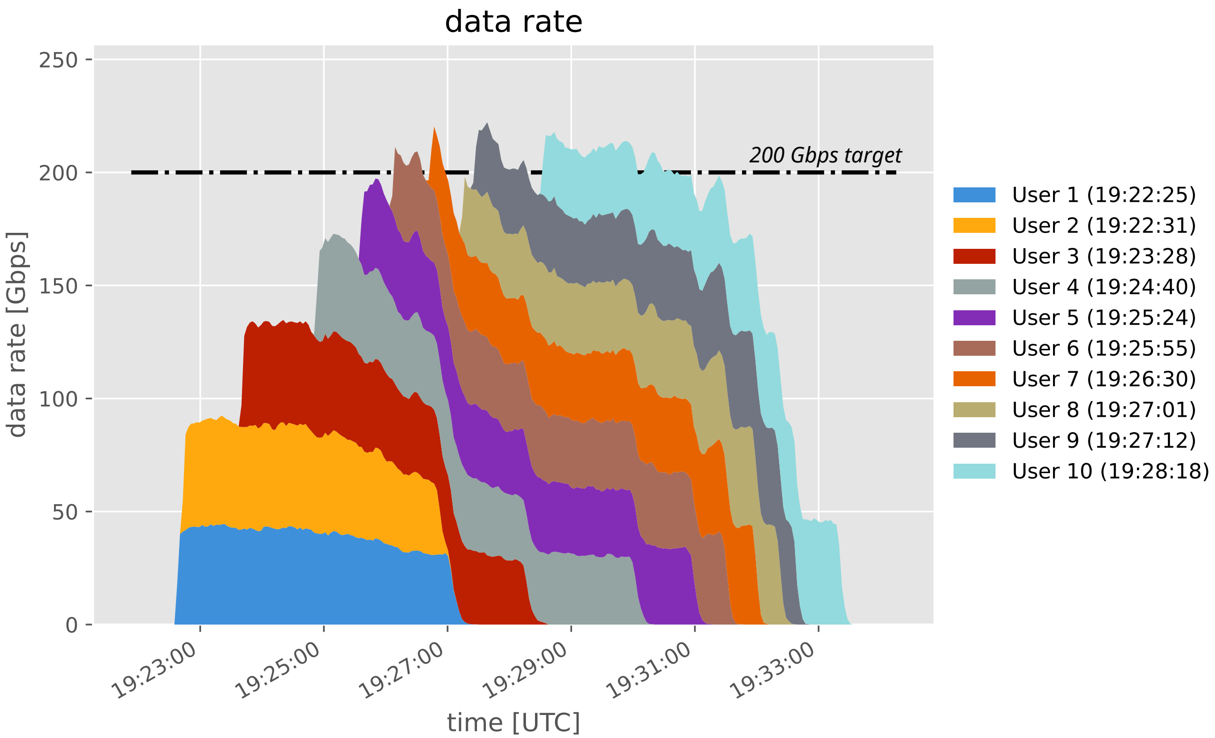

task launch times were randomly distributed to simulate reality of random submissions

- Test with **ten simultaneous users at UChicago**
  - users limited to max 200 cores
- Reached **200 Gbps collectively**
  - network saturation effect visible

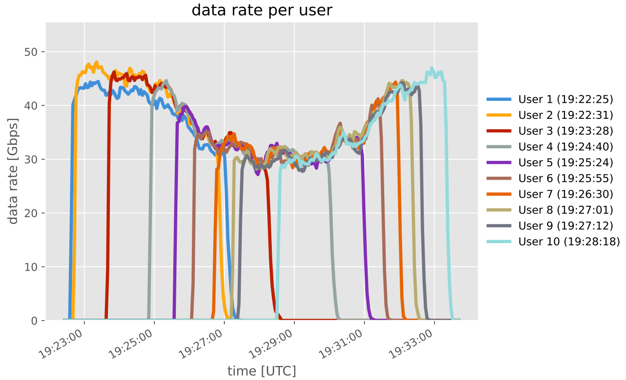

Automatic worker scaling

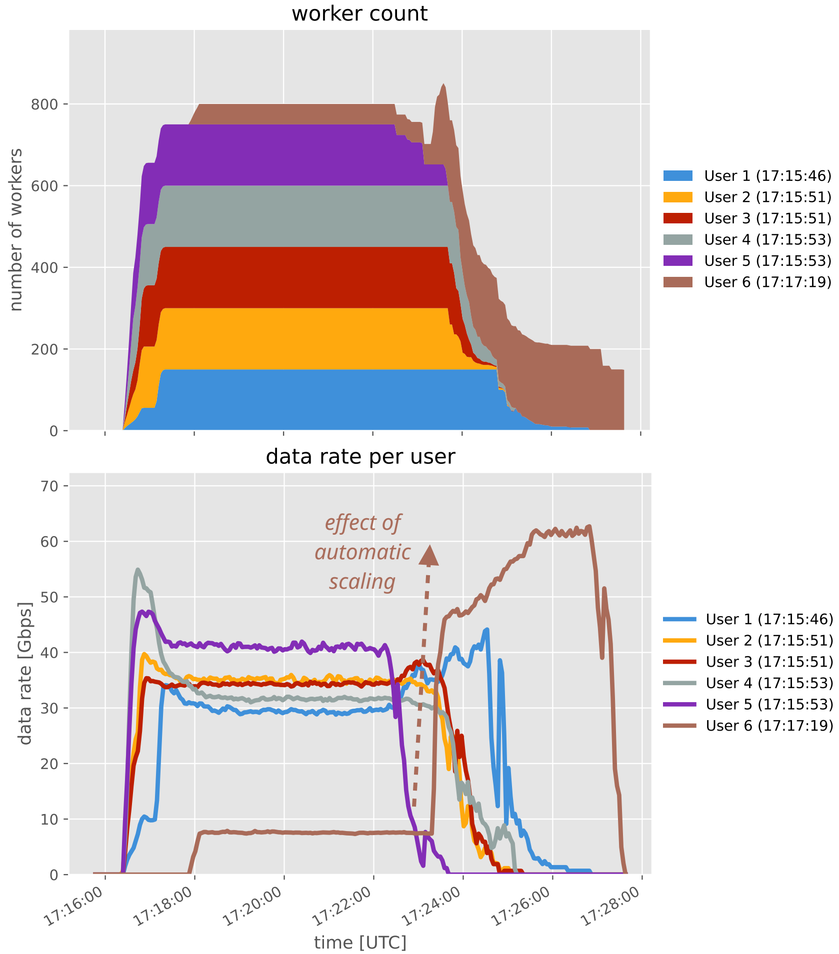

- Test with **six simultaneous users at Nebraska**
- First five users launch at the same time
  - stable parallel processing
- **User 6** receives last available cores, slower initially
  - rapid automatic scaling once more resources become available

Where to go from here?

Next steps

- **Further explore the parameter space** of HL-LHC analyses
  - extend 200 Gbps setup towards full AGC-type workflow
  - medium term: 400 Gbps exercise
- Further **collaboration with community & knowledge sharing**
  - lessons learned help analyses *already today*
- Help **inform Analysis Facility evolution**

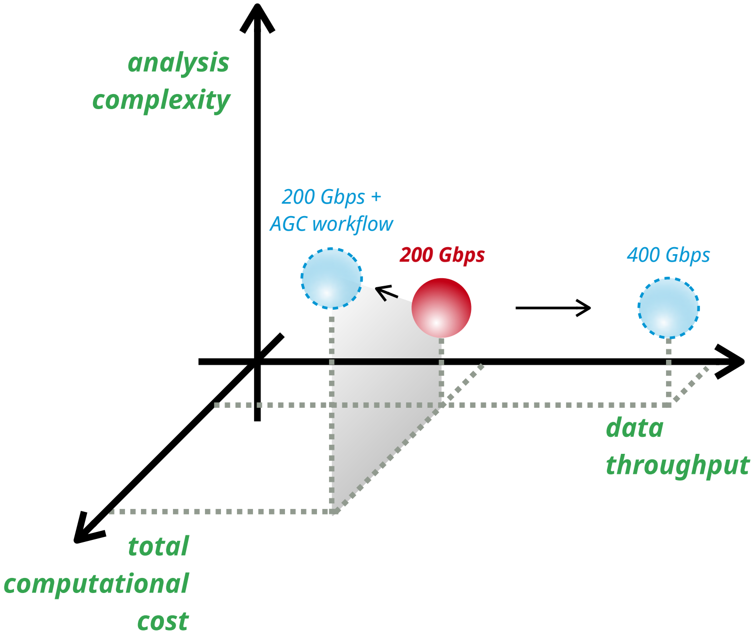

axis closely connected to environmental sustainability → David Britton's talk

Conclusion

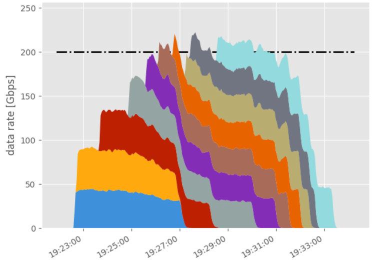

- **Successful 200 Gbps Challenge** shows technology readiness, checkpoint towards HL-LHC
  - extremely valuable project to generate feedback and identify potential bottlenecks
  - planned extensions for more realism (Analysis Grand Challenge)

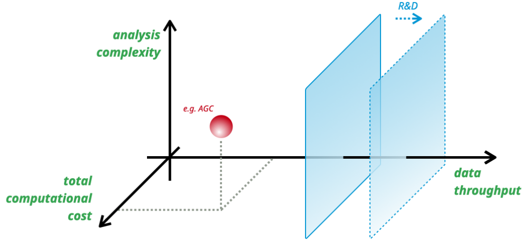

- **Difficult to predict future** of HEP end-user data analysis
  - factorize challenges & push boundaries, remain open to new ideas
- **Close collaboration remains crucial:** take advantage of CHEP to connect to colleagues from different areas!
  - physics is the end goal

Also check out these *related contributions* this week!

Investigating Data Access Models for ATLAS: A Case Study with FABRIC Across Borders and ServiceX [Oct 22, 15:18]
GIL-free scaling of Uproot in Python 3.13 [Oct 23, 14:42]
Building Scalable Analysis Infrastructure for ATLAS [Oct 23, 17:09]
Operating the 200 Gbps IRIS-HEP Demonstrator for ATLAS [Oct 24, 16:33]
Tuning the CMS Coffea-casa facility for 200 Gbps Challenge [Oct 24, 16:51]
Benchmarking massively-parallel Analysis Grand Challenge workflows using Snakemake and REANA [Oct 24, 17:45]

- **Contact:** join us at analysis-grand-challenge@iris-hep.org (sign up: google group link)