

# FERMILAB

The 200 Gbps Challenge: Imagining HL-LHC analysis facilities

FERMILAB-CONF-25-0158-CSAID

arXiv:2503.15464

This manuscript has been authored by Fermi Forward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

## The 200 Gbps Challenge: Imagining HL-LHC analysis facilities

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

<sup>1</sup>University of Wisconsin–Madison, United States

<sup>2</sup>University of Nebraska–Lincoln, United States

<sup>3</sup>Morgridge Institute for Research, United States

<sup>4</sup>University of Chicago, United States

<sup>5</sup>University of Texas at Austin, United States

<sup>6</sup>Princeton University, United States

<sup>7</sup>Fermilab, United States

<sup>8</sup>University of Notre Dame, United States

<sup>9</sup>University of Washington, United States

**Abstract.** The IRIS-HEP software institute, as a contributor to the broader HEP Python ecosystem, is developing scalable analysis infrastructure and software tools to address the upcoming HL-LHC computing challenges with new approaches and paradigms, driven by our vision of what HL-LHC analysis will require. The institute uses a “Grand Challenge” format, constructing a series of increasingly large, complex, and realistic exercises to show the vision of HL-LHC analysis. Recently, the focus has been demonstrating the IRIS-HEP analysis infrastructure at scale and evaluating technology readiness for production.

As a part of the Analysis Grand Challenge activities, the institute executed a “200 Gbps Challenge”, aiming to show sustained data rates into the event processing of multiple analysis pipelines. The challenge integrated teams internal and external to the institute, including operations and facilities, analysis software tools, innovative data delivery and management services, and scalable analysis infrastructure. The challenge showcases the prototypes — including software, services, and facilities — built to process around 200 TB of data in both the CMS NanoAOD and ATLAS PHYSLITE data formats with test pipelines.

The teams were able to sustain the 200 Gbps target across multiple pipelines. The pipelines focusing on event rate were able to process at over 30 MHz. These target rates are demanding; the activity revealed considerations for future testing at this scale and changes necessary for physicists to work at this scale in the future. The 200 Gbps Challenge has established a baseline on today’s facilities, setting the stage for the next exercise at twice the scale.

---

\*e-mail: alexander.held@cern.ch

# 1 Introduction

The upcoming High Luminosity Large Hadron Collider (HL-LHC) will deliver a dataset of collisions with a total size significantly exceeding that resulting from the Run-2 and ongoing Run-3 operation of the LHC. Both the CMS [1] and ATLAS [2] collaborations have studied the computing challenges associated with handling this data and outlined the need for R&D to meet the computational requirements of physics analyses with HL-LHC data.

The focus of the work reported here lies in the last stages of a High Energy Physics (HEP) analysis pipeline, in the realm of *end user physics analysis*. Earlier stages in the pipeline are commonly centrally organized and the end user part begins with centrally provided input files. The subsequent workflow can vary a lot between analyses, but typically has to be run frequently from the design stage of a physics analysis all the way through to the publication of the results. The turnaround time for this end user workflow increase with dataset size. This poses the risk of large increases in the time it takes to perform a physics analysis at the HL-LHC. Mitigating this risk is one of the missions of the IRIS-HEP software institute.

## 1.1 IRIS-HEP and the “Grand Challenge” format

IRIS-HEP [3] is the Institute for Research and Innovation in Software for High Energy Physics. Since 2018, IRIS-HEP has performed computing and software R&D for the HL-LHC: it targets a *software upgrade* accompanying the detector hardware upgrades that are being pursued for the CMS and ATLAS experiments. This work is being done in close collaboration with the experiments and computing facilities. IRIS-HEP members are a mix of physicists, computer scientists, and engineers, distributed across many institutes in the United States. They work on a broad range of topics, many of which are connected via “Grand Challenges”. These challenges outline a series of exercises of increasing complexity towards realistic HL-LHC physics analysis scale. A recent focus has been end user physics analysis: demonstrating IRIS-HEP analysis infrastructure at scale and evaluating technology readiness. Two recent key projects are situated in this field:

- **Analysis Grand Challenge (AGC)** [4–7]: an end-to-end analysis pipeline that serves as an integration exercise,
- **200 Gbps Challenge**: demonstrating data throughput of 200 gigabit per second for physics analysis applications.

## 1.2 Overall vision

A central goal for IRIS-HEP is to empower physicists by minimizing the time-to-insight and thereby maximizing the HL-LHC physics reach. The time-to-insight is the turnaround time of physics analysis and incorporates time spent debugging, bookkeeping, and waiting for computing to finish. The institute aims to tighten feedback and support cycles in particular among three groups of people:

- physicists,
- software developers,
- analysis facility experts.

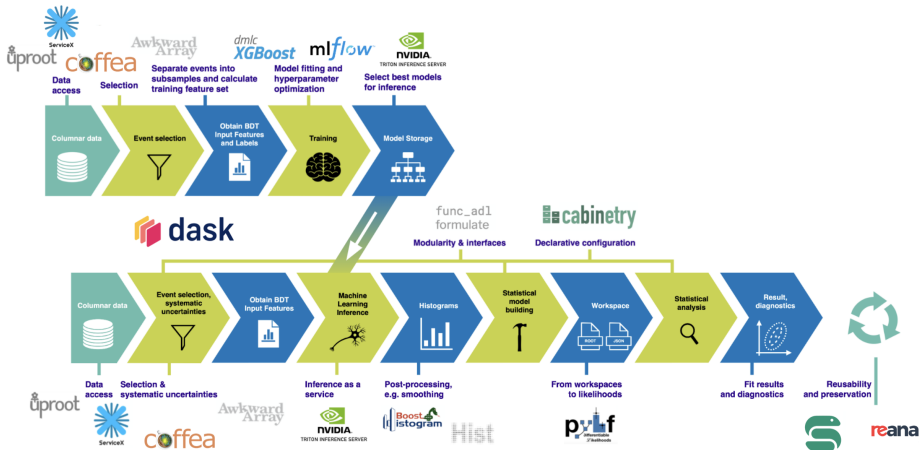
An efficient and scalable implementation of a physics analysis pipeline requires expertise in these three areas. By having people from all three work together within the IRIS-HEP challenge format, the institute targets the development of solutions that can serve as a blueprint for the future.

The unifying IRIS-HEP vision for end-user physics analysis encompasses the following:

- **Analyze  $O(1000\text{ TB})$  of data within a few hours:** This scale captures the envisioned HL-LHC physics analysis needs and multiple hours of turnaround time for a full pipeline run allow for efficient progress in a physics analysis.
- **Interactive analysis turnaround with a “coffee break” timescale:** The possibility of iterating on ideas on the timescale of a coffee break with a meaningful subset of data means that development can proceed rapidly and waiting time does not become a bottleneck.
- **Fully integrated Analysis Facilities (AFs):** These facilities provide a convenient interfaces to access required services and computational resources.
- **User experience to empower big & small teams:** A recent dedicated white paper [8] describes desirable features from a user perspective, including the handling of software environments and collaboration with colleagues.
- **Easy access to state-of-the-art Machine Learning (ML) techniques:** ML techniques are widely used throughout physics analysis and new ideas have the potential to shape analysis workflows, so sufficient flexibility to incorporate novel developments is desirable.
- **Reproducibility, preservation, reuse:** Developing sustainable physics analyses will help maximize their long-term impact and legacy [9].

## 2 The Analysis Grand Challenge

The AGC defines a physics analysis task starting from publicly accessible CMS Open Data [10]. It features the extraction and processing of data into histograms, ML training and evaluation, as well as statistical inference. The reference implementation of this task is provided by IRIS-HEP [11] and heavily uses Python libraries from the Scikit-HEP [12] project. This pipeline is depicted in Figure 1.



**Figure 1.** The IRIS-HEP reference implementation for the AGC analysis task, figure adapted from [6].


The AGC project provides a useful environment for integration tests and a variety of implementations beyond the IRIS-HEP reference have been developed. For IRIS-HEP, the implementation employs task graphs to express and execute data analysis operations. This is done with Dask [13], a Python library providing both an interface to describe manipulations

of data through task graphs, as well as ways to efficiently schedule the distributed execution of such graphs. Deep integration of Dask with other HEP-specific Python libraries, such as coffea [14], has been achieved over time and Dask is additionally emerging as a common feature in many AFs.

The AGC has been successful in its integration role but is limited in scale as the size of the input data is only around 2 TB. This motivated the start of a new Grand Challenge project to specifically target scaling behavior.

### 3 The 200 Gbps Challenge

End user physics analysis pipelines are not centrally prescribed and there is limited agreement on how physics analyses will evolve towards the HL-LHC. It is therefore currently unclear what a “representative” HL-LHC analysis might look like. This poses a challenge for software and computing R&D: in the absence of agreed-upon benchmarks, the evaluation of technical readiness is more difficult. Figure 2 proposes a way to factorize computational aspects of physics analysis pipelines into independent challenges. One axis is pure data throughput, focused on how fast data can be read, decompressed and be made available for any subsequent physics analysis operations. Another axis is the total computational cost of the pipeline: this can capture aspects such as the cost to train and evaluate ML models, but also the processing of systematic variations for the evaluation of associated uncertainties in the physics result. The third axis is the analysis complexity, which can relate to the number of different steps required in the workflow or the types of operations and external services that are required to be supported. A given analysis example, such as the AGC shown in red, can be located somewhere within this space.



**Figure 2.** Factorizing computational aspects of physics analyses into independent challenges: the 200 Gbps project is focused on data throughput.

A benchmark analysis example like the AGC probes a specific point in this space. Defining technical capabilities today, as well as moving forward towards the HL-LHC, requires addressing what region of this space is technologically accessible. The 200 Gbps Challenge project focuses purely on data throughput to evaluate what can be achieved with the approaches envisioned by IRIS-HEP and to subsequently push these boundaries further outwards with targeted development.

### 3.1 Defining the challenge

The 200 Gbps Challenge is named after its key goal: demonstrate data throughput in a physics analysis context at a rate of 200 Gbps, sustained over half an hour. These numbers target physics analysis at scale to probe HL-LHC needs. A given physics analysis will frequently only read a subset of the content in the centrally provided input files. Given a target 25% of the file content needing to be read, a rate of 200 Gbps corresponds to processing a 180 TB dataset within half an hour.

The relevant data formats in the HL-LHC context for the CMS and ATLAS experiments are NanoAOD [15] and PHYSLITE [16]. Table 1 provides more context about what this challenge means with event sizes typical for these formats. For the NanoAOD example, reaching 200 Gbps with 1000 cores means processing events with a rate of 50 kHz and a throughput of 25 MB/s per core. For PHYSLITE, a slightly more complicated data structure, the use of 2000 cores would mean event rates of 5 kHz and throughput of 12.5 MB/s per core.

**Table 1.** Event numbers and rates for 200 Gbps data throughput sustained over half an hour for CMS NanoAOD and ATLAS PHYSLITE examples.

	CMS NanoAOD example	ATLAS PHYSLITE example
size per event	2 kB	10 kB
number of events	90 billion	18 billion
target total event rate	50 MHz	10 MHz

As the challenge targets data throughput in the context of subsequent physics analysis, it includes all the steps leading up to this: the initial read, network transfer, as well as decompression. It captures the steps up until the point where arrays of data are available in memory, ready for any case-specific additional physics operations.


Two different setups were studied in the context of the 200 Gbps Challenge. The CMS-targeted implementation [17] uses Run-3 CMS NanoAOD data and was operated at an AF located at the University of Nebraska–Lincoln. The ATLAS version [18] uses Run-2 ATLAS PHYSLITE data and ran at the University of Chicago AF. While both implementations have a similar target of demonstrating sustained throughput of 200 Gbps, they do differ in various ways. Most importantly, the input files vary in event size, object types, and compression used; the AFs differ in their setup and hardware.

The initial timeline for this project was tight: the goal was to present the status at the 2024 WLCG/HSF workshop, which took place 8 weeks after the creation of the challenge. Despite this, the presentation [19] at the workshop demonstrated that the goal was achieved and it was further improved upon subsequently.

### 3.2 Key Analysis Facility elements

Figure 3 schematically depicts the data flow and relevant AF components used in the challenge context. The input data starts out distributed across the Worldwide LHC Computing Grid (WLCG). Reading over the wide-area network puts limits on throughput and latency which are incompatible with the goals of this challenge. XCache [20, 21] instances are deployed within the AFs to address this. Remote data is cached during the first read and all subsequent reads happen out of the cache and are contained to the local network for fast and stable data access. For the purpose of the following measurements all caches are assumed to be warm, with the caching having already taken place previously. This corresponds to the

typical situation during the development of physics analysis where physicists iterate on a design and will repeatedly run similar data processing pipelines. The AFs also provide ways to distribute the workflow across available resources. In the cases discussed here the individual worker nodes provide CPU resources for data handling and decompression.



**Figure 3.** Key AF components used to address the 200 Gbps Challenge.

## 4 Measurements with the CMS setup at Nebraska


Benchmarking with the CMS setup [17] took place on the coffea-casa [22] AF located at the University of Nebraska–Lincoln. Coffea-casa is built as a prototype for a HL-LHC AF with a hybrid setup employing both Kubernetes resources and resources via the local CMS Tier-2 and Tier-3 site. Eight XCache instances were deployed, each with a 2x100 Gbps uplink, to scale to the intended throughput level. Details about the setup can be found in a dedicated contribution to this conference [23].

The measurements run a task graph built with Dask. Nodes in this graph encapsulate the work to be done for a single file: reading the data with uproot [24] into awkward [25] arrays and returning metadata for benchmarking purposes. Both Dask and TaskVine [26] were used to perform the distributed execution using resources provided by both HTCondor [27] and Kubernetes.

Figure 4 shows a measurement example with Dask on HTCondor resources. Using a constant allocation of around 1300 workers, a data rate exceeding 200 Gbps was sustained over a period of roughly 15 minutes. This example processed 40 billion events distributed across 64k NanoAOD files for an event rate of 32 MHz. A total of 30 TB of compressed data was read, corresponding to an uncompressed size of 71 TB. This measurement was performed on NanoAOD inputs that were re-compressed from the default LZMA to instead use the Zstandard algorithm. The change in compression algorithm resulted in an event rate increase per core of roughly a factor 2.5 and made reaching the 200 Gbps target possible.

## 5 Measurements with the ATLAS setup at Chicago

The ATLAS setup [18] used the University of Chicago AF. This AF is used in production for ATLAS, is fully Kubernetes-based and was partially reconfigured to address the computational needs of the 200 Gbps Challenge. A detailed description about the facility configuration is provided in a separate contribution to this conference [28]. Eight XCache instances were deployed at this AF as well, distributed across the network to maximize the available total bandwidth to the computing resources.




**Figure 4.** Data rate measurement at the coffee-casa AF located at the University of Nebraska–Lincoln. A rate exceeding the target of 200 Gbps was sustained over 15 minutes in a distributed setup using around 1300 workers.

Two different data pipeline configurations were pursued: a Dask-distributed reading of data through an XCache with uproot and a setup employing ServiceX [29] as a data delivery service.

### 5.1 Results with Dask and uproot

Using Dask to distribute reading of data into awkward arrays with uproot, the target 200 Gbps data rate was sustained over a period of 20 minutes in the example shown in Figure 5. The worker allocation was dynamic and controlled by the Dask scheduler. It quickly ramps up to a peak of around 1700 workers and scales down again towards the end of the execution as the number of remaining tasks decreases. The example processed 32 billion events from 218k PHYSYLITE files with a total size of 190 TB for an event rate of 15 MHz, with per-core rates of 5–20 kHz. The uncompressed size of the data read was 32 TB, corresponding to 80 TB uncompressed.




**Figure 5.** Data rate measurement at the University of Chicago AF. A rate exceeding the target of 200 Gbps was sustained over 20 minutes in a distributed setup using up to around 1700 workers.



## 5.2 Results with ServiceX

A second data processing pipeline that was tested for the challenge employed ServiceX as a data delivery service. Data is read through the caches, processed by ServiceX, and then written out to object storage. The resulting output can then be processed further in a next stage, for example with another Dask-distributed setup. This multi-stage schema can be beneficial especially when the data size significantly decreases in the first processing step, for example due to event filtering being applied. Repeated executions of the following processing step can then be particularly efficient if they only infrequently necessitate re-running the earlier step.

Figure 6 shows the data rates achieved in an example run, measured through network monitoring. This example resulted in a rates of 170 Gbps over a 25 minute runtime. It processed 19 billion events corresponding to 146 TB of data, using up to 1000 pods. The processing time of subsequent analysis steps, which read data out of the object storage, strongly depends on the filtering performed during this ServiceX stage.



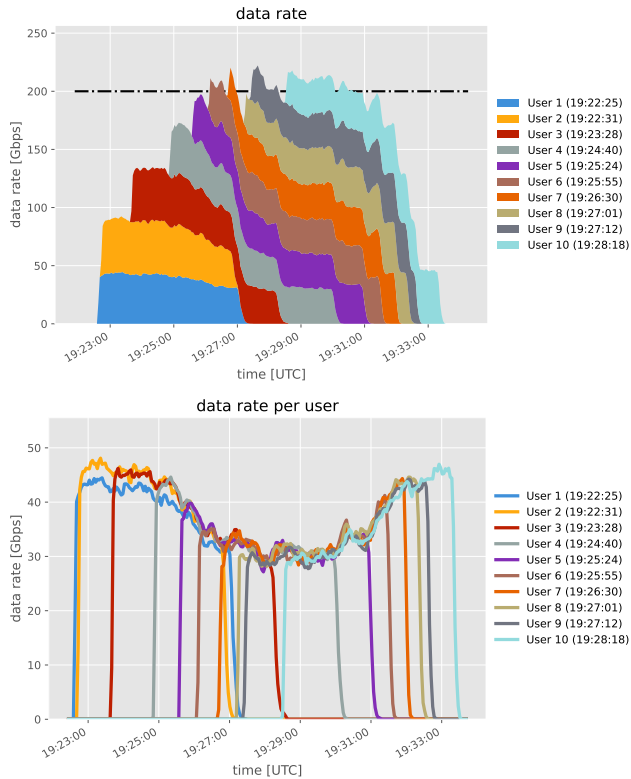
**Figure 6.** Data rate measurement at the University of Chicago AF. Using ServiceX, the network monitoring shows rates up to 170 Gbps in this 25 minute example using up to 1000 pods.

## 6 Multi-user interactive analysis

Future AFs are aiming to host many physicists who all would like to simultaneously scale to sufficiently many resources to enable interactive analysis turnaround. It is therefore important to demonstrate that good performance can also be achieved with multiple users running in parallel. The two exercises reported here looked at different scenarios in this context at the AFs at the University of Nebraska–Lincoln and the University of Chicago.

The first exercise featured ten ATLAS users at the University of Chicago AF launching data processing pipelines adapted from the 200 Gbps setup. Launch times were randomized and each user was restricted to be assigned at most 200 cores, with Dask dynamically scaling the allocation. The total data rate shown in Figure 7 (top) reached 200 Gbps, distributed between multiple different users. Figure 7 (bottom) depicts the data rate per user as a function of time. The rates dropped in the middle of the exercise as the network resources were saturated at this point in time.

The second exercise used the coffea-casa AF at the University of Nebraska–Lincoln. Figure 8 (top) shows the worker allocation per participant. Five CMS users launched a data processing pipeline at the same point in time and were allocated 150 workers each. Dask dynamically scaled the allocation up to the 150 worker limit. The sixth user received a more limited allocation due to the total availability of resources. After the first tasks finished up,



**Figure 7.** Multi-user test at the University of Chicago AF. With ten users executing data processing pipelines, the available bandwidth was distributed between users and reached 200 Gbps (top). A saturation effect is visible in the per-user data rates (bottom).

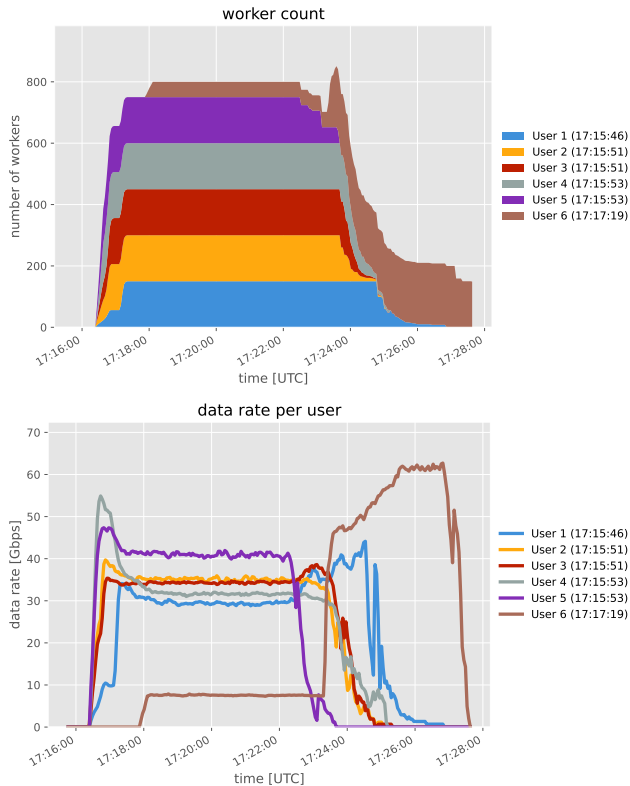
the sixth user was dynamically allocated more of the available resources through Dask. The data rate for this user subsequently increased significantly as shown in Figure 8 (bottom).

These measurements demonstrate how resources can be dynamically distributed across users. Finding the right resource provisioning and fair-share patterns in the context of interactive analysis will likely require more work in this direction.

## 7 Conclusion and next steps

The 200 Gbps Challenge successfully demonstrated technology readiness and provided a checkpoint towards the HL-LHC. It served as a valuable mechanism to generate feedback and identify potential bottlenecks with large-scale data throughput analysis use cases. The project relied on close collaboration between physicists, analysis software developers, and analysis facility experts in order to reach the 200 Gbps throughput target.

Looking ahead, IRIS-HEP intends to further study the space of physics analyses outlined in Figure 2. Two directions are of particular interest: adding complexity to the existing 200 Gbps setup to capture more of the event-by-event processing of a physics analysis, as well as extending the throughput towards a 400 Gbps Challenge. The future of HEP end user physics analysis remains difficult to predict, but these targeted efforts aim to ensure readiness for a broad range of HL-LHC use cases.



**Figure 8.** Multi-user test on the coffea-casa AF at the University of Nebraska–Lincoln. When the resource demand for all users exceeded the available resources, users received more resources as other tasks finish up (top). The per-user data rates show the effect of resource redistribution and subsequent change in data rates (bottom).

This work was supported by the U.S. National Science Foundation (NSF) cooperative agreements OAC-1836650, OAC-2029176, OAC-2115148, PHY-2120747, PHY-2121686, and PHY-2323298.

## References

- [1] CMS Offline Software and Computing, CMS Phase-2 Computing Model: Update Document (2022), <https://cds.cern.ch/record/2815292>
- [2] ATLAS Collaboration, ATLAS Software and Computing HL-LHC Roadmap (2022), <https://cds.cern.ch/record/2802918>
- [3] P. Elmer, M. Neubauer, M.D. Sokoloff, Strategic Plan for a Scientific Software Innovation Institute (S2I2) for High Energy Physics (2017), [arXiv:1712.06592](https://arxiv.org/abs/1712.06592).
- [4] A. Held, O. Shadura, The IRIS-HEP Analysis Grand Challenge, PoS **ICHEP2022**, 235 (2022). [10.22323/1.414.0235](https://arxiv.org/abs/10.22323/1.414.0235)
- [5] O. Shadura, A. Held, First performance measurements with the Analysis Grand Challenge (2023), [arXiv:2304.05214](https://arxiv.org/abs/2304.05214)

- [6] A. Held, E. Kauffman, O. Shadura, A. Wightman, Physics analysis for the HL-LHC: Concepts and pipelines in practice with the Analysis Grand Challenge, EPJ Web Conf. **295**, 06016 (2024). [10.1051/epjconf/202429506016](https://doi.org/10.1051/epjconf/202429506016)
- [7] E. Kauffman, A. Held, O. Shadura, Machine Learning for Columnar High Energy Physics Analysis, EPJ Web Conf. **295**, 08011 (2024). [10.1051/epjconf/202429508011](https://doi.org/10.1051/epjconf/202429508011)
- [8] D. Ciangottini et al., Analysis Facilities White Paper (2024), [arXiv:2404.02100](https://arxiv.org/abs/2404.02100).
- [9] S. Bailey et al., Data and Analysis Preservation, Recasting, and Reinterpretation (2022), [arXiv:2203.10057](https://arxiv.org/abs/2203.10057).
- [10] CMS Data preservation and open access group, Getting Started with CMS 2015 Open Data, <https://opendata.cern.ch/docs/cms-getting-started-2015> (2022)
- [11] A. Held, E. Kauffman, O. Shadura, E. Guiraud, M. Feickert, J. Chakraborty, M. Proffitt, A. Wightman, K. Choi, A. Novak et al., Analysis Grand Challenge, <https://doi.org/10.5281/zenodo.7274936>
- [12] E. Rodrigues et al., The Scikit HEP Project – overview and prospects, EPJ Web Conf. **245**, 06028 (2020). [10.1051/epjconf/202024506028](https://doi.org/10.1051/epjconf/202024506028)
- [13] Dask Development Team, Dask: Library for dynamic task scheduling, <https://dask.org> (2016)
- [14] L. Gray, N. Smith, B. Tovar, Y.M.E. Chen, I. Krommydas, A. Novak, J. Chakraborty, N. Krug, P. Fackeldey, D. Valsecchi et al., coffea, <https://doi.org/10.5281/zenodo.3266454>
- [15] A. Rizzi, G. Petrucciani, M. Peruzzi, A further reduction in CMS event data for analysis: the NANOAOB format, EPJ Web Conf. **214**, 06021 (2019). [10.1051/epjconf/201921406021](https://doi.org/10.1051/epjconf/201921406021)
- [16] J. Elmsheuser et al., Evolution of the ATLAS analysis model for Run-3 and prospects for HL-LHC, EPJ Web Conf. **245**, 06014 (2020). [10.1051/epjconf/202024506014](https://doi.org/10.1051/epjconf/202024506014)
- [17] IRIS-HEP, IDAP 200 Gbps with CMS NanoAOD, <https://github.com/iris-hep/idap-200gbps> (2024)
- [18] IRIS-HEP, IDAP 200 Gbps with ATLAS PHYSYLITE, <https://github.com/iris-hep/idap-200gbps-atlas> (2024)
- [19] B. Bockelman, IRIS-HEP 200 Gbps Challenge, WLCG/HSF Workshop 2024, <https://indico.cern.ch/event/1369601/contributions/5924000/> (2024)
- [20] L.A.T. Bauerdick et al. (CMS), XRootd, disk-based, caching proxy for optimization of data access, data placement and data replication, J. Phys. Conf. Ser. **513**, 042044 (2014). [10.1088/1742-6596/513/4/042044](https://doi.org/10.1088/1742-6596/513/4/042044)
- [21] A. Hanushevsky et al. (ATLAS), Xcache in the ATLAS Distributed Computing Environment, EPJ Web Conf. **214**, 04008 (2019). [10.1051/epjconf/201921404008](https://doi.org/10.1051/epjconf/201921404008)
- [22] M. Adamec, G. Attebury, K. Bloom, B. Bockelman, C. Lundstedt, O. Shadura, J. Thiltges, Coffea-casa: an analysis facility prototype, EPJ Web Conf. **251**, 02061 (2021). [10.1051/epjconf/202125102061](https://doi.org/10.1051/epjconf/202125102061)
- [23] O. Shadura et al., Tuning the CMS Coffea-casa facility for 200 Gbps Challenge, <https://indico.cern.ch/event/1338689/contributions/6010711/> (2024)
- [24] J. Pivarski, H. Schreiner, A. Hollands, P. Das, K. Kothari, A. Roy, J. Ling, N. Smith, C. Burr, G. Stark, Uproot, <https://doi.org/10.5281/zenodo.4340632>
- [25] J. Pivarski, I. Osborne, I. Ifrim, H. Schreiner, A. Hollands, A. Biswas, P. Das, S. Roy Choudhury, N. Smith, M. Goyal, Awkward Array, <https://doi.org/10.5281/zenodo.4341376>
- [26] B. Sly-Delgado, B. Tovar, J. Zhou, D. Thain, Reshaping High Energy Physics Applications for Near-Interactive Execution Using TaskVine (2024), <https://doi.org/10.5281/zenodo.1144444>

[1109/SC41406.2024.00068](#)

- [27] D. Thain, T. Tannenbaum, M. Livny, Distributed computing in practice: the Condor experience, *Concurr. Comput.: Pract. Exper.* **17**, 323 (2005). [10.1002/cpe.938](#)
- [28] R.W. Gardner Jr, D. Benjamin, L. Bryant, M. Feickert, F. Golnaraghi, A. Held, F. Hu, D.A. Jordan, J.L. Stephen, I. Vukotic et al. (ATLAS), Operating the 200 Gbps IRIS-HEP Demonstrator for ATLAS (2025), <https://cds.cern.ch/record/2923996>
- [29] B. Galewsky, R. Gardner, L. Gray, M. Neubauer, J. Pivarski, M. Proffitt, I. Vukotic, G. Watts, M. Weinberg, ServiceX A Distributed, Caching, Columnar Data Delivery Service, *EPJ Web Conf.* **245**, 04043 (2020). [10.1051/epjconf/202024504043](#)