

HL-LHC blueprint & 200 Gbps: the next iteration

Alexander Held¹, Oksana Shadura²

¹ University of Wisconsin-Madison

² University of Nebraska-Lincoln

IRIS-HEP Steering Board meeting

<https://indico.cern.ch/event/1507606/>

July 29, 2025

“Physics analysis at the HL-LHC” workshop

Workshop took place May 15/16: <https://indico.cern.ch/event/1515493/>


- Understand **computational cost** of physics analyses today
- Provide a set of “**representative**” **physics analyses** & characterize **possible exceptions**
- Define **extrapolation scenarios** for analyses today towards HL-LHC

- **Summary document** with discussion outcomes -> **revisit this in the future** as needed to adjust

- **Overall aim**: align R&D with future needs, ensure **resource availability** for your future HL-LHC analyses!
 - required: **projection of computational and storage needs for facilities**

Breaking down physics analyses

- Currently **limited agreement** on how “**HL-LHC physics analyses**” will look: aiming to change this!



Organization

Day 1: CMS focus

14:00	Intro 6/R-012 - conference room, CERN	Alexander Held et al. 14:00 - 14:05
	Run-3 CMS compute model 6/R-012 - conference room, CERN	Nick Smith 14:05 - 14:20
	CMS survey results 6/R-012 - conference room, CERN	Oksana Shadura 14:20 - 14:45
	CMS analysis example: dark QCD search (15'+10') 6/R-012 - conference room, CERN	Kevin Pedro 14:45 - 15:10
15:00	CMS analysis example: long-lived particles search (15'+10') 6/R-012 - conference room, CERN	Jeremi Niedziela 15:10 - 15:35
	Coffee break	15:35 - 16:00
16:00	CMS analysis example: W boson mass measurement (15'+10') 6/R-012 - conference room, CERN	David Walter 16:00 - 16:25
	CMS analysis example: Higgs to mu mu search (15'+10') 6/R-012 - conference room, CERN	Hyeon-Seo Yun 16:25 - 16:50
	Brainstorming: what did we not capture here? 6/R-012 - conference room, CERN	16:50 - 17:00
17:00	Discussion about HL-LHC analysis extrapolation 6/R-012 - conference room, CERN	17:00 - 18:00
18:00		

intro, context
& survey results

selected analysis
examples

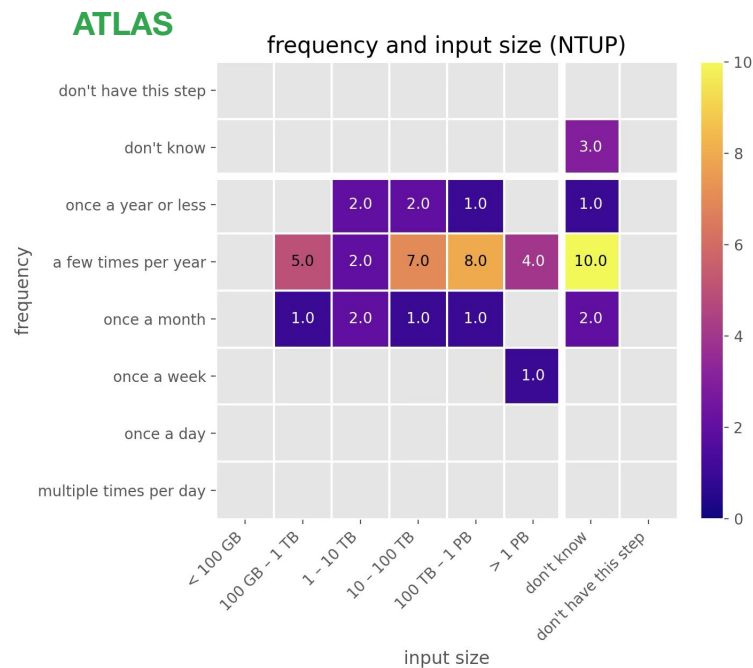
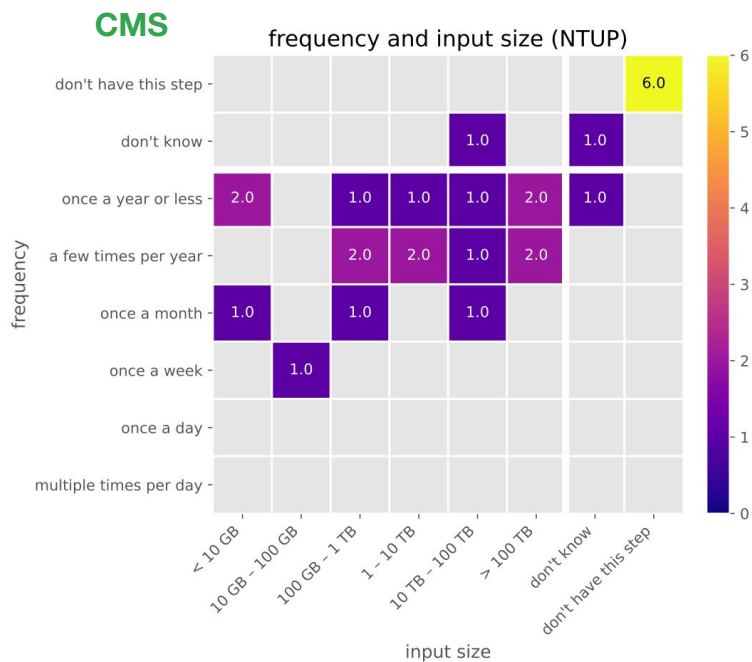
discussion & extra-
polation scenarios

Day 2: ATLAS focus

14:00	Intro 40/S2-D01 - Salle Dirac, CERN	Alexander Held et al. 14:00 - 14:05
	Run-3 ATLAS compute model 40/S2-D01 - Salle Dirac, CERN	Emma Torro Pastor et al. 14:05 - 14:20
	ATLAS survey results 6/R-012 - conference room, CERN	Alexander Held 14:20 - 14:45
	ATLAS analysis example: di-Higgs search (15'+10') 40/S2-D01 - Salle Dirac, CERN	Sebastien Rettie 14:45 - 15:10
15:00	ATLAS analysis example: jet energy calibration with ttbar (15'+10') 40/S2-D01 - Salle Dirac, CERN	Tomas Dado 15:10 - 15:35
	Coffee break	15:35 - 16:00
16:00	ATLAS analysis example: SUSY in multi-lepton final states (15'+10') 40/S2-D01 - Salle Dirac, CERN	Tina Potter 16:00 - 16:25
	ATLAS analysis example: R(K*) measurement (15'+10') 40/S2-D01 - Salle Dirac, CERN	Dvij Mankad 16:25 - 16:50
	Brainstorming: what did we not capture here? 40/S2-D01 - Salle Dirac, CERN	16:50 - 17:00
17:00	Discussion about HL-LHC analysis extrapolation 40/S2-D01 - Salle Dirac, CERN	17:00 - 18:00
18:00		

Survey results

- See the [workshop slides](#) for many more details (including invited talks about example use cases!)
- **Scales vary by multiple of magnitude** between analyses—we now have numbers to quantify this
- **Diverse landscape of approaches**, multi-stage pipelines to efficiently filter down



200 Gbps: the next iteration
Integration Challenge

Defining the 200 Gbps Challenge (2024)



We identified a **200 Gbps data rate** as relevant target in 2024 and pursued this goal:

see <https://arxiv.org/abs/2503.15464>

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.

Aspects we want to capture in an **Integration Challenge**

- **Dataset of HL-LHC-relevant scale** (with possibility to duplicate), **TTree + RNTuple**
- **Physics** complexity, e.g. systematic uncertainty handling
 - increased compute **cost per event** (capturing full physics analysis needs)
- **Possibilities to interface to additional services** (data delivery, ML inference, ...)

-> **Creating a pipeline for realistic HL-LHC analyses to study functionality, interfaces, performance.**



workflow example [<https://agc.readthedocs.io/>]

Overall plan

- **Analysis setup**
 - ▶ **ttbar phase space**: flexible to do e.g. x-sec measurement like in AGC or BSM interpretation (e.g. Z')
 - ▶ full set of **systematic uncertainties** (experimental + theory)
- **Input datasets** (aiming for TTree + RNTuple)
 - ▶ Run-2 + Run-3 **NanoAOD** for CMS
 - ▶ Run-2 + Run-3 **PHYSLITE** for ATLAS
- **Output**
 - ▶ **O(10k) histograms** + diagnostics for performance measurements and debugging
- Details of analysis now being worked on!
- Similar to AGC: open for **additional implementations** (just like for AGC)

ATLAS-specific details

- **IRIS-HEP pipeline approach**
 - ▶ **TopCPToolkit** producing ntuples from PHYSLITE (with systematic variations evaluated)
 - ▶ **coffea** with **virtual arrays**
 - ▶ optional: **ServiceX** as intermediate step (delivering ntuples to coffea)
- Aiming to **align with US-ATLAS** goals

CMS-specific details

- **IRIS-HEP pipeline approach**
 - ▶ On-the-fly systematic variations with correctionlib
 - ▶ coffea with virtual arrays
 - ▶ optional: ServiceX for data delivery

Preparing next iteration of CMS Analysis Facilities

Running analysis frameworks and tools at scale

- **Adopt diverse computing executors** to support execution of complex task graphs
 - **Dask, TaskVine, VirtualArrays (new coffea development)**
- Flexible computing **resource provisioning model** optimised for given facility
 - **Kubernetes, Tier-2 resources, HPC cluster**
- Fairshare / Batch
 - **HTCondor, Slurm, Kubernetes**
- Workflow management
 - **Luigi**

USCMS Ops program has dedicated milestone to execute challenge on all USCMS facilities:
Facility can prioritise the interesting setup they are going to explore (based on this list and not only)

XCache - service provides caching of data accessed using xrootd protocol

CMS ServiceX - data extraction and delivery service

1000+ cores - expect that no meaningful results can be obtained with significantly less resources

Preliminary timeline

Timeline is a work in progress (synchronizing with parties involved)

- **Sep 2025**: pipeline implementations available (Python-based, optionally as notebook)
- **Sep–Dec 2025**: scale testing, usability & performance improvements
- **Dec 2025**: possibility of small event to discuss results (open to all interested collaborators)
- **CHEP 2026**: presentation of final outcomes

Beyond this timeline: can use the implementations for other future studies: functionality, usability, benchmarking, ...


Questions / comments? Happy to hear your input!

Backup

Analysis Grand Challenge (AGC)

- 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**



- Focused on **workflow & interfaces**, not computational cost: **1.8 TB inputs, 1 B events**

