# ATLAS S&C toward the HL-LHC challenges

HEP-CCE All Hands Meeting

18th December 2023

Edward Moyse & Verena Martinez Outschoorn (UMass)
with much appreciated inputs from Marilena Bandieramonte, Attila Krasznahorkay, Doug Benjamin, Peter van Gemmeren and several others

ATLAS EXPERIMENT

# Introduction

Thank you for sharing the draft plans with us

In the next slides we will try to cover (our interpretation of) the ATLAS response to the four technical areas:

- Portable Applications and Workflows (PAW)
- Storage Optimization (SOP)
- Accelerating HEP Simulation (AHS)
- Scalable AI/ML on HPC systems (SML)

We will summarise the stated goals, and our thoughts…

*Disclaimer* - we have some ATLAS experts in leadership positions and so ATLAS priorities were considered “by design”…

# Portable Applications and Workflows (PAW)

ATLAS EXPERIMENT

- **Five year goals**
  - Delivering to ASCR facilities and the HEP community mini-apps and benchmarks representative of accelerated HEP workflows
  - Enabling more HEP workflows to run on multiple HPC systems
  - Training early career researchers to understand HEP computing on HPC systems
  - Reaching out to a broad university community (including MSIs) through seminars, lectures, hackathons, etc.
- **First year goals**
  - Develop a cookbook for application portability layers based on Phase 1 findings
  - Support experiments to implement portable accelerated applications through workshops and hackathons followed by regular “office hours”
  - Survey existing HEP experiment solutions to run workflows on HPC systems
  - Explore HEP needs in terms of ML workflows and microservices (synergistic with SML)
  - Create two mini-apps based on the Phase I portable parallelization testbeds that can be executed at ALCF, NERSC, and OLCF, preferably with the same software environment
  - Use mini-apps to extract figures of merit for ASCR facilities and LCFs to use as baselines
  - Investigate common layers and interfaces (batch scheduler, policies, pilots, etc.) to facilitate portability and interoperability across ASCR facilities
  - Create two representative HEP experiment workflows portable across two different HPC systems

# PAW: Our comments (1)

- **Generally on Training / outreach / workshops / hackathons:**
  - ATLAS is based in Europe; for useful training would need to cover several time zones/accommodate remote learning
  - It would be helpful to partner this with existing ATLAS/US ATLAS as well as HSF programmes
- **Cookbook**
  - We can try to identify personpower to test and give feedback
- **Survey existing HEP experiment solutions to run workflows on HPC systems**
  - Information can be found in the report on USATLAS/USCMS HPC/Cloud (see link here)
  - ATLAS has focused primarily on simulation workflows in the past
- **Explore HEP needs in terms of ML workflows and microservices (synergistic with SML)**
  - Increasing interest in use-cases, especially for object identification (e.g. b-tagging) and for analysis methods (e.g. simulation based inference, etc)
  - We need to connect these effortds
- **Create two (year 1) mini-apps based on the Phase I portable parallelization testbeds that can be executed at ALCF, NERSC, and OLCF, preferably with the same software environment**
  - Interested to learn about the experience and benefit from it internally

# PAW: Our comments (2)

ATLAS EXPERIMENT

- **Investigate common layers and interfaces (batch scheduler, policies, pilots, etc.) to facilitate portability and interoperability across ASCR facilities**
  - Here we think it would to useful to connect our operational experts
- **Create two representative HEP experiment workflows portable across two different HPC systems**
  - Can try to find effort to contribute to provide representative HEP experiment workflows.
  - For example Traccc might be possible to try on different HPC systems (we have existing CUDA + SYCL implementations, and partial Alpaka/Kokkos)

# Storage Optimization (SOP)

- **Five year goals**
  - Implement direct mapping of accelerator-friendly data models to disk storage
  - Provide a suite of experiment-driven tools for mimicking and characterizing I/O of HEP workflows
  - Provide solutions to store full experimental data (including simulation outputs) using ROOT RNTuple technology
  - Apply at least two “intelligent” compression algorithms to data from at least two HEP experiments
  - Determine, in collaboration with our stakeholders, the applicability and impact of data/workflow orchestration approaches like on-demand data delivery, filtering, augmentation, etc.
- **1st year goals**
  - Add to Darshan features needed to monitor HEP workflows
  - Demonstrate direct mapping of accelerator-friendly data models to disk storage
  - Identify the feature set (data model, tools) needed to store experiments reconstruction and simulation output data in RNTuple
  - Develop infrastructure to apply at least one intelligent compression algorithm to HEP production data streams
  - Prototype experiment agnostic samples of ROOT data to be used to demonstrate object store technologies
  - Build toy prototype allowing HDF5 data streaming via xrootd
  - Explore and evaluate existing data orchestration solutions in HEP, streaming data between at least one experiment data store and DOE HPC systems

# SOP: Our comments

- **Monitoring HEP workflows:**
  - ATLAS is developing plans to integrate Darshan into their SPOT measurements and would benefit from continued improvements of this tool.
- **ROOT RNTuple**
  - ATLAS expects to write their offline event data to ROOT RNTuple, so this area is of crucial importance.
  - Combining efforts to make improvements here can have great benefits.
- **Intelligent compressions:**
  - Intelligent or lossy compression have not yet been widely or generically adapted into ATLAS workflows. However, with the increase of HL-LHC data volumes potential benefits are increasing and combined research could proof valuable.
  - For example, for a long time we have used a lossy compression scheme for our calorimeter cell collections - but perhaps something more sophisticated could help us? Can we extend compression to more domains? Perhaps better validation tools would help?
  - Stability is important - we need to be able to read old datasets, no matter what compression was used
- **Data stores:**
  - New technology, such as object stores, could benefit from combined R&D to make them applicable to HEP experiments. Object stores could enable deduplication of data and save storage.

**Short term goals:**

- Exploring parallel, GPU-friendly implementations of Next-to-Leading-Order (NLO) 1-loop event generator calculations
- Exploring GPU porting of infrared subtraction methods needed for NLO
- Porting of optical photon physics models to GPU in Celeritas
- Development of GPU-enabled surface-based shape models needed for LHC, and other experimental detector models

# AHS: Our comments

- **Porting of optical photon physics models to GPU in Celeritas**
  - We already started testing Celeritas (and also AdePT) with FullSimLight and test beam applications
    - We are looking forward testing them with more complex geometry setup and integrating them in Athena.
  - Recent review by GEANT4
- **Development of GPU-enabled surface-based shape models needed for LHC, and other experimental detector models**
  - Sounds similar to approaches taken by Traccc / ACTS (which ATLAS is heavily invested in).
  - *Common solutions are always nice!*
  - Orange by Celeritas and Adept also going in the same direction by developing the same kind of model in the next generation of VecGeom
  - We look forward to having the full set of volumes to run further benchmarks

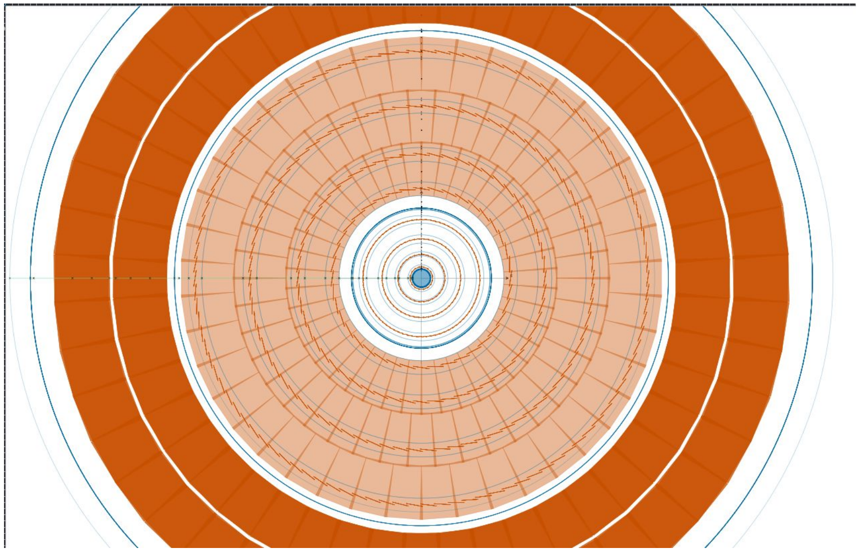

Ray scan in Detray/ODD using pure surface based description

# Scalable AI/ML on HPC systems (SML)

ATLAS EXPERIMENT

**Five year goals:**

- Publish performance studies of ML pipelines (both for training and inference) on heterogeneous resources
- Shorten the development cycle for two large models weeks→hours
- Develop a CCE distributed ML platform adopted by at least two experiments

**First year goals:**

- Identify target ML models in collaboration with experiments
- Port, train, and run at least two target models on two different HPC systems
- Compare two data parallel training solutions for at least one target model
- Compare two hyperparameter optimization tools on at least one target model
- Setting up a prototype Inference as-a-service platform on at least one DOE HPC system

# SML: Our comments

- **Develop a CCE distributed ML platform adopted by at least two experiments**
  - What about using existing OpenSource platforms?
- **Identify target ML models in collaboration with experiments**
  - Can provide examples
- **Port, train, and run at least two target models on two different HPC systems**
  - We can contribute here.
  - Aside: does “run” here include integrating with reconstruction code?
  - FastCaloSim has already been ported on GPU and we have a preliminary work on containerizing FastCaloGan with a FastCaloGantainer prototype and some first results on HPC systems:
    - Report (slides 14 and 15)
  - Simulation based inference methods for analysis
    - Report
  - Object identification methods, including tracking (GNN4ITk), flavor tagging, tau identification, lepton isolation, etc

# Backup

# HEP-CCE Project Questions / Comments (I)

From ATLAS in July

ATLAS EXPERIMENT

- We haven't seen a lot of visible impact (yet?) from some HEP-CCE projects
  - Maybe that's natural — HL-LHC isn't here yet, and these things are hard to deploy
  - OTOH, several ATLAS members deeply involved in HEP-CCE and doing great things
- Our Roadmap includes a decision about (portability) languages in 2024
  - This is very much because we expect to learn from HEP-CCE's recommendations in 2023
  - ATLAS members are deeply involved in this effort; we don't think you're missing anything
- We are aware of the IOS activities (and ATLAS members are involved)
  - We haven't seen a killer use-case for e.g. the HDF5 writing yet
  - When we talk about "I/O limiting" in our applications, it is *rare that we really mean the I/O layer itself*. Usually we mean I/O+decompression+ROOT overheads+access overheads, and recent measurements indicate that these last two are the most important to us.
- Offloading (e.g. to GPU) data efficiently is an ongoing project in our Roadmap
  - Perhaps looking for further overlap/collaboration here would be useful
- Intelligent data compression is certainly interesting to us
- Without thinking deeply about it:
  - Metadata is an ongoing fight for ATLAS in a variety of ways. Perhaps there's some opportunity here for collaboration / harmonization / common tooling?

# Traccc ITK demonstrator almost finished.

https://indico.cern.ch/event/1295479/contributions/5616016/attachments/2747644/4781460/2023-11-07-ACTS-WS-Introduction.pdf

| | Q4/2023 | Q1/2024 | Q2/2024 | Q3/2024 |
|---|---|---|---|---|
| Open Data Detector Demonstrator | finish demonstrator | Performance, optimisation<br>traccc paper writing | traccc paper writing<br>& finishing | Traccc standalone wrap-up. |
| ITk Demonstrator | start with ITk<br>geometry integration | finish demonstrator | Performance studies<br>ATLAS documentation | |
| ACTS Integration | | traccc as ACTS/thirdparty | ACTS -> traccc offloading<br>(Partly, full chain) | ACTS reintegration discussion |
| ATLAS/ATHENA Integration | | Review components<br>for direct integration in ATHENA | Traccc deomonstrator<br>In ATHENA | |

# LHC schedule

- We are just entering the second year of Run 3
  - Run 3 will end in 2025
- HL-LHC will start in 2029
  - Run 4: 2029–2032, Run 5: 2035–2038. Formal plan extends beyond 2040.

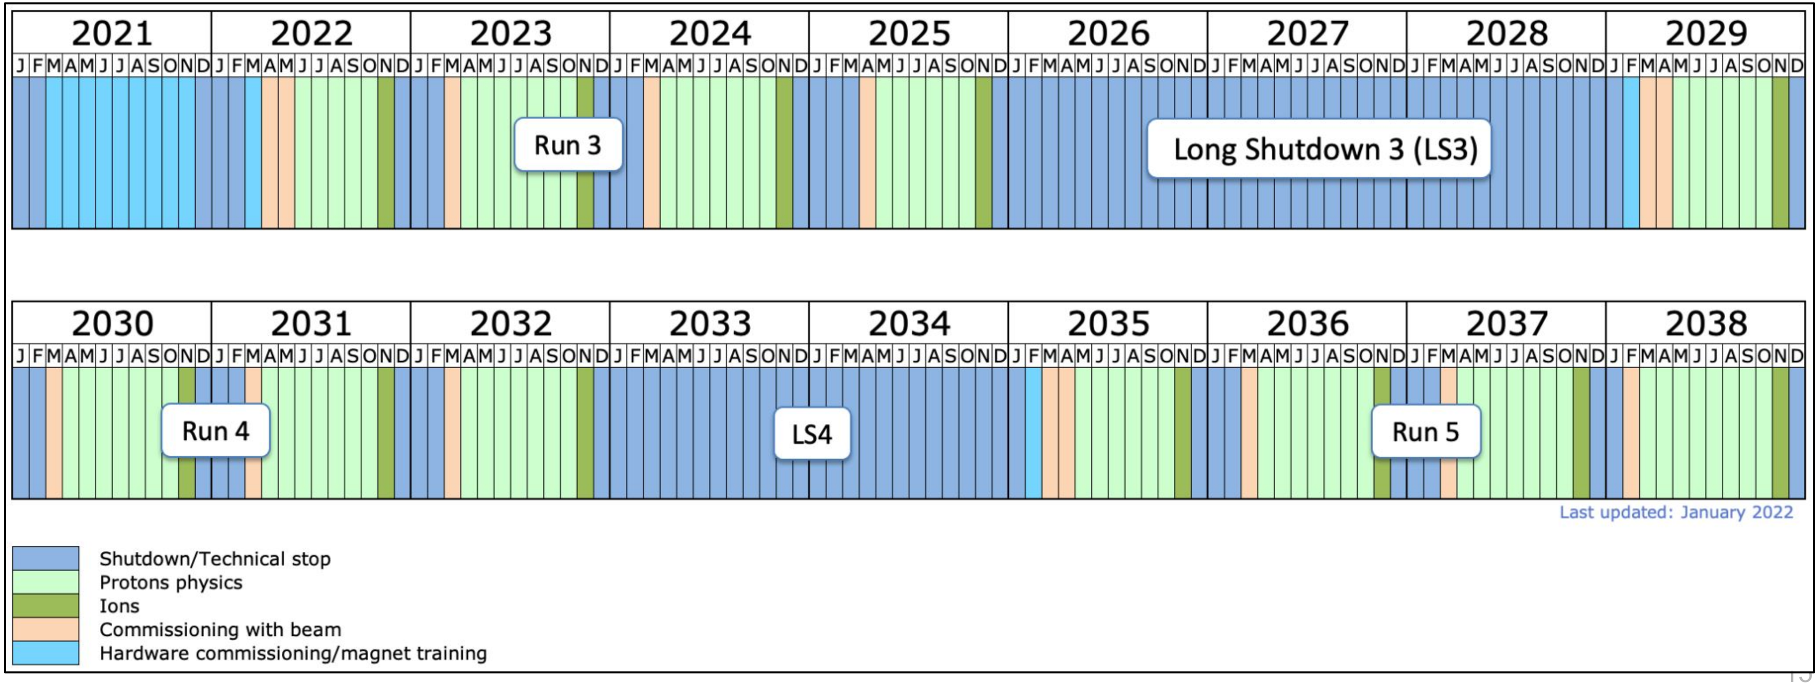

Our “interesting” data set size is measured in $fb^{-1}$

- Trigger rate (incl. delayed) ~ 3.5 kHz
- RAW size ~ 1.3 MB/ev
- Pile up ~ 60
- ~ 30B MC events/year

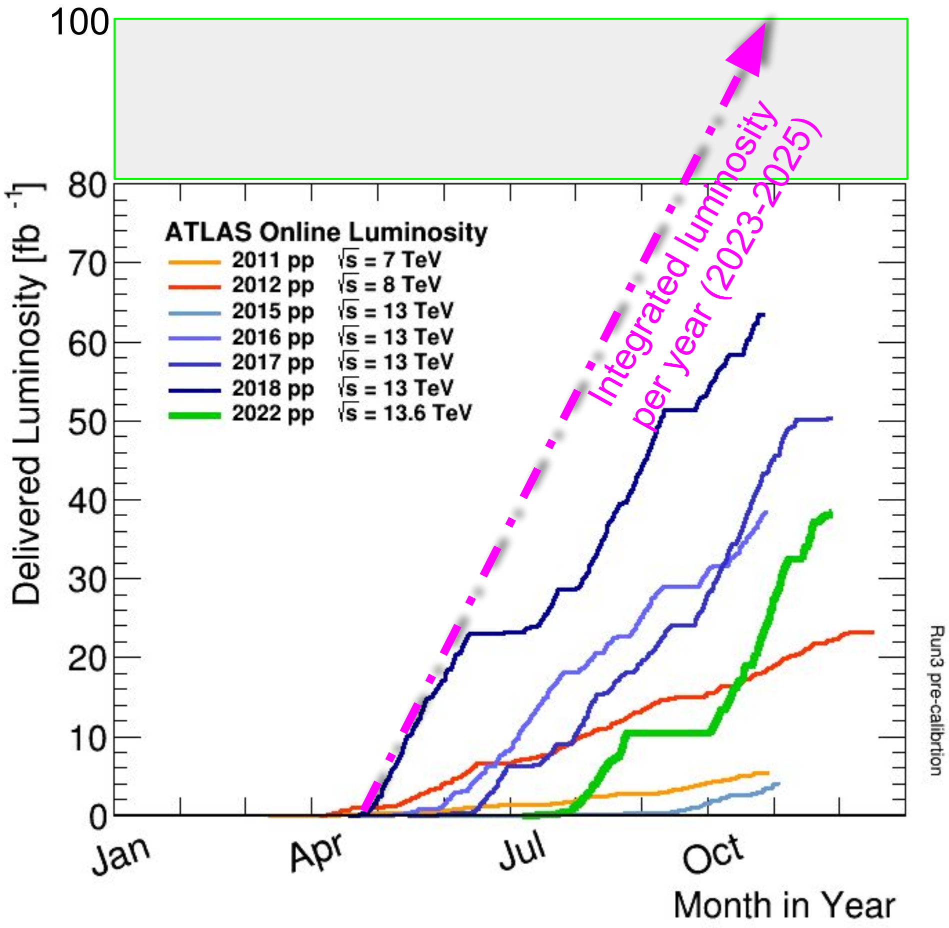

# The Physics challenge: HL-LHC (Run 4)

Our "interesting" data set size is measured in $fb^{-1}$

- Trigger rate ~ 10 kHz
- RAW size ~ 4 MB/ev
- Pile up ~ 140+
- ~ 150B MC events/year

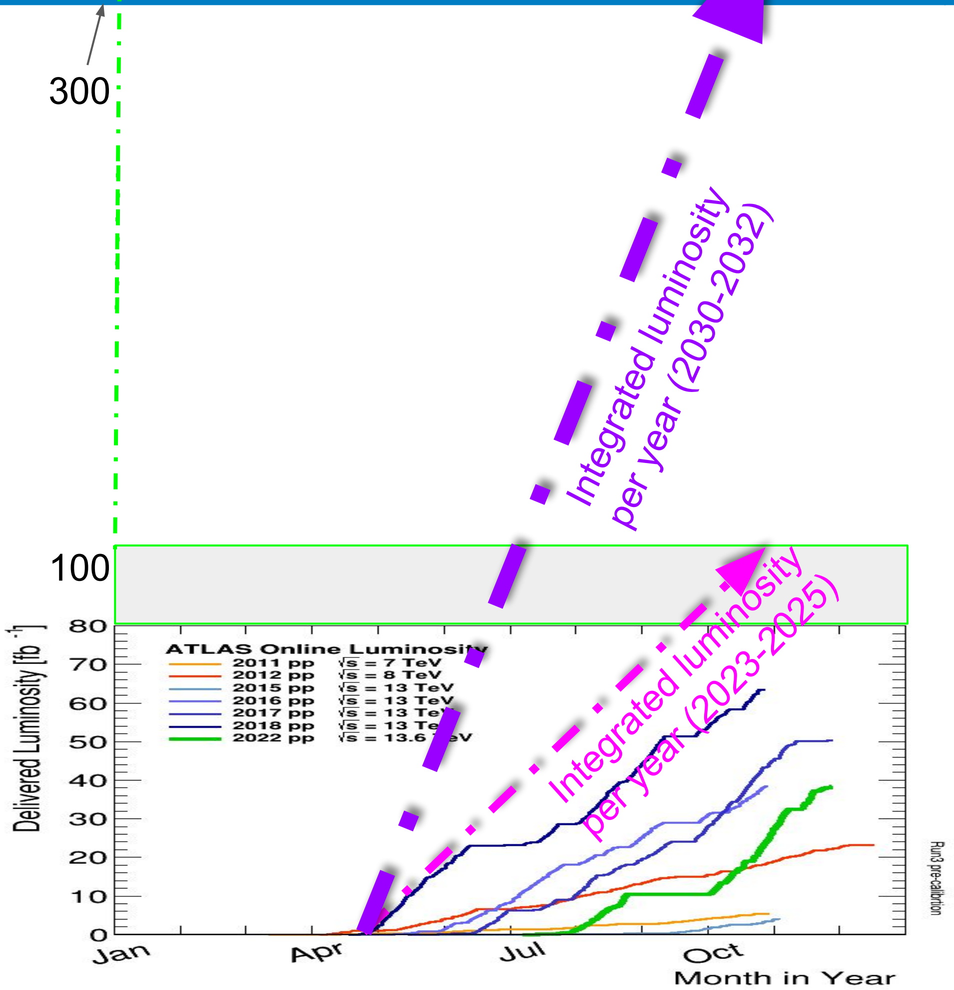

# The Computing Challenges

- We hope you're all familiar with these plots by now!
  - Model updated in Spring 2023 — no significant change to these projections
- Need for major R&D (or budgetary) effort to achieve HL-LHC physics potential
  - We have defined *Conservative R&D* and *Aggressive R&D* scenarios
    - N.b. some projects for which we were/are not able to estimate the concrete impact are not (yet) included (e.g. GPU usage, FastChain simulation)
  - The black lines indicate the "flat budget" of 10% (lower line) and 20% (upper line)

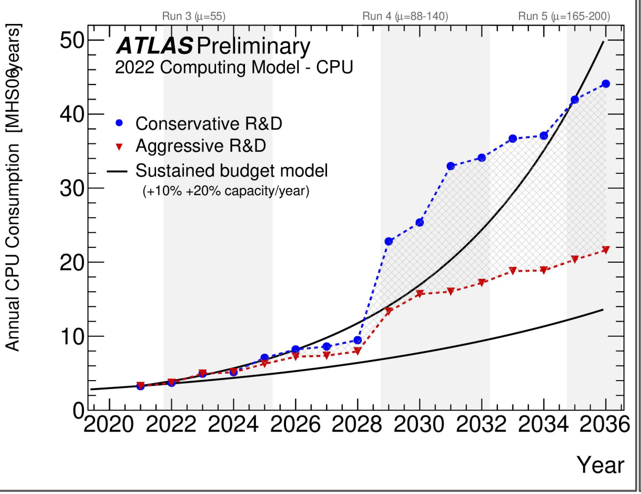


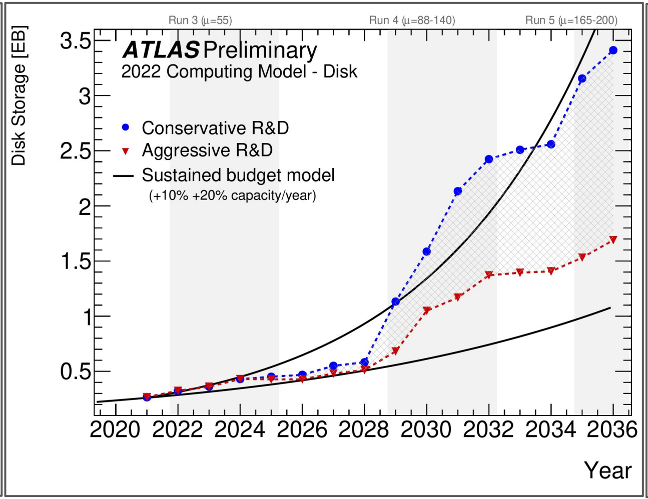


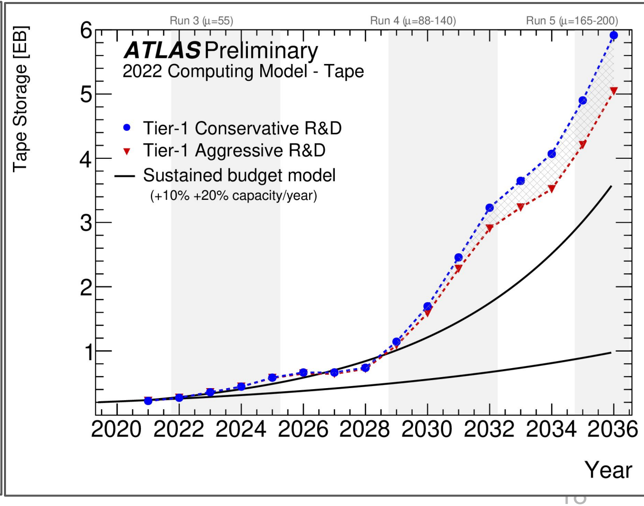

# Getting There from Here

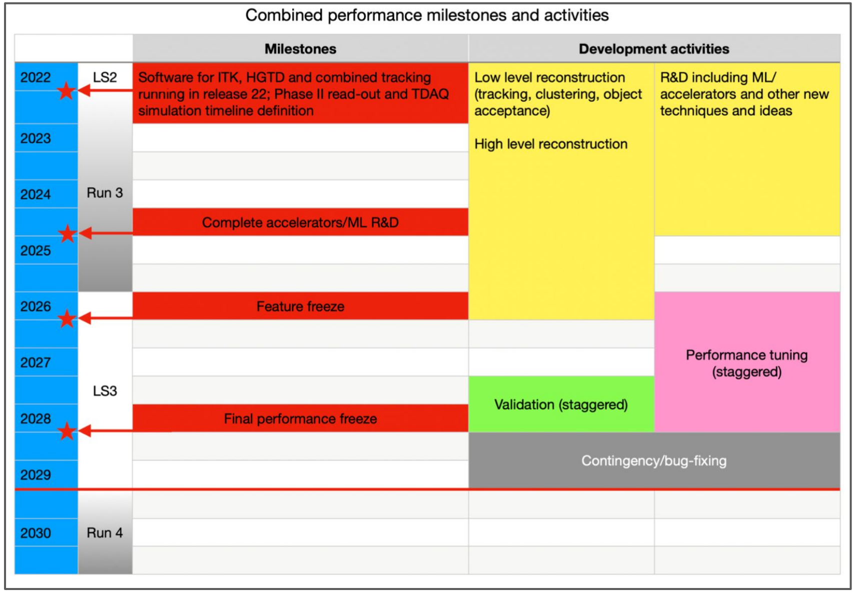


- ATLAS S&C HL-LHC Roadmap
- Defines milestones and deliverables to get to the HL-LHC successfully
- R&D is ongoing
  - We've just spent some time discussing and reviewing "demonstrator" prototype projects for the HL-LHC
  - We see already lot of engagement!
- Integration and validation will require **time**
  - Late arriving R&D is risky
- In 2025 we expect to have in the TDR a detailed path to HL-LHC data taking
  - For example, accelerators: Yes or No
  - This is also the timeline for a decision from our trigger group on accelerators

# Pain Points (Processing / CPU)

- No single application dominates CPU in 2031
- That's good news and bad news for us
  - No silver bullet to "solving" our resource crunch
  - Also not fatal if one workflow isn't improved
  - Can diversify our R&D — lots of interesting projects!
- Lots of effort and ideas around these problems
  - Very happy with the diverse portfolio HEP-CCE is investing in; many of the key problems are covered
- Biggest (by some metric) "CPU" efforts currently in:
  - Faster simulation (Geant4 optimization / on GPU, better fast simulation, FastChain…)
  - ML/accelerator-based charged particle tracking
  - New approaches to analysis
- Effort spread around reasonably well
  - Other ML/accelerator approaches to reconstruction, event generation, etc
  - Cleaning up "waste" (e.g. unused / failed production)

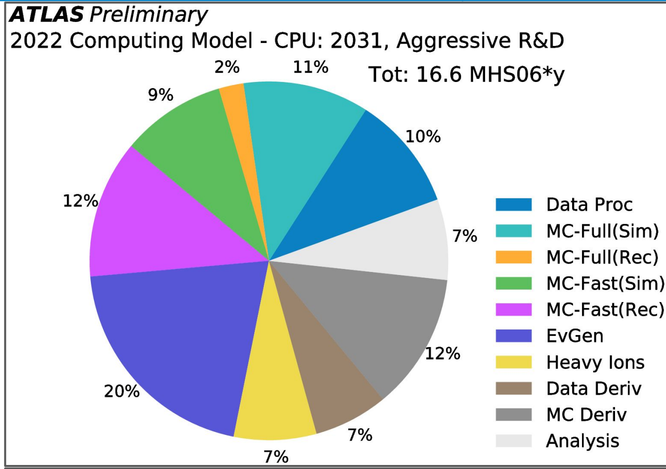


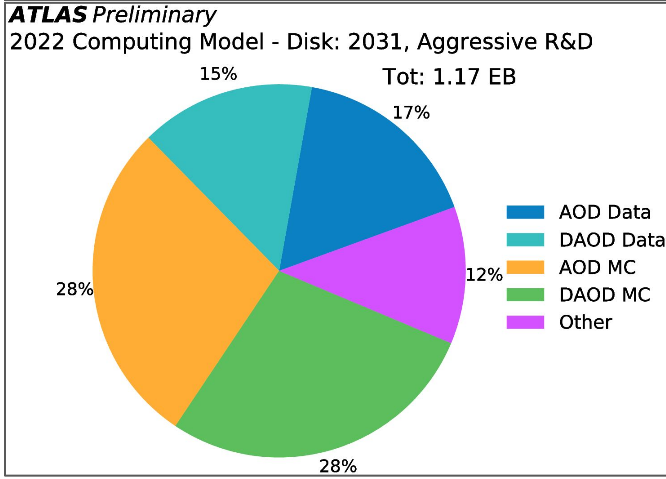

# Pain Points (Storage)

- Several ongoing disk efforts as well
  - RNTuple (of course)
  - Lossy compression: difficulty is not the infrastructure but the physics validation
  - Augmentation to support sparse additional data
  - Constant revision / review of file contents
  - Alternative compression/settings — delicate balance
  - More aggressive deletion / recreation
- Anticipating our PHYSLITE format will serve a wide variety of analyses in Run 4
  - The disk model is driven by “remnants”: how many analyses *don’t* use PHYSLITE, and what they use instead. This is where the hard work goes!
- Very successful model for data distribution
  - Our “data carousel” uses tape effectively as a warm storage medium, reducing disk needs
  - Already have a mechanism for replication of popular datasets, and expecting to continue this way

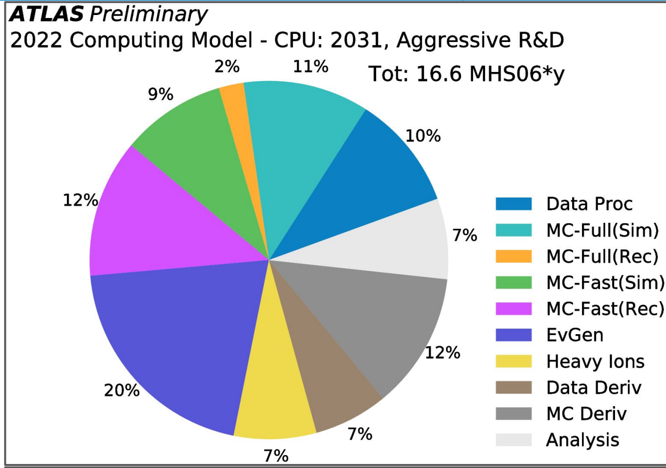


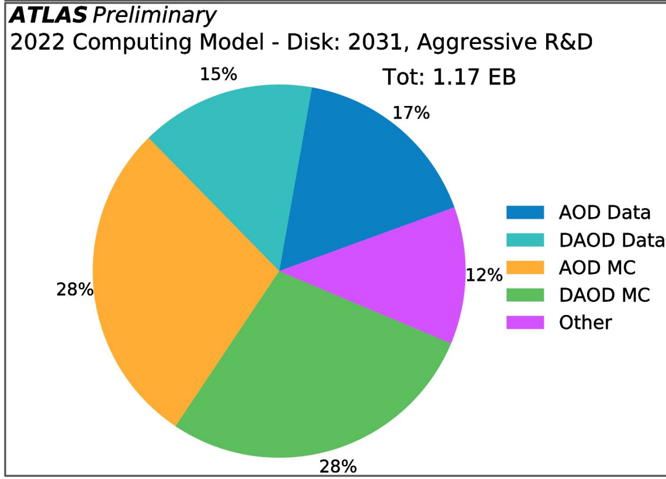

# High-Performance Computing

- We have benefited enormously from HPC systems (especially since 2021)

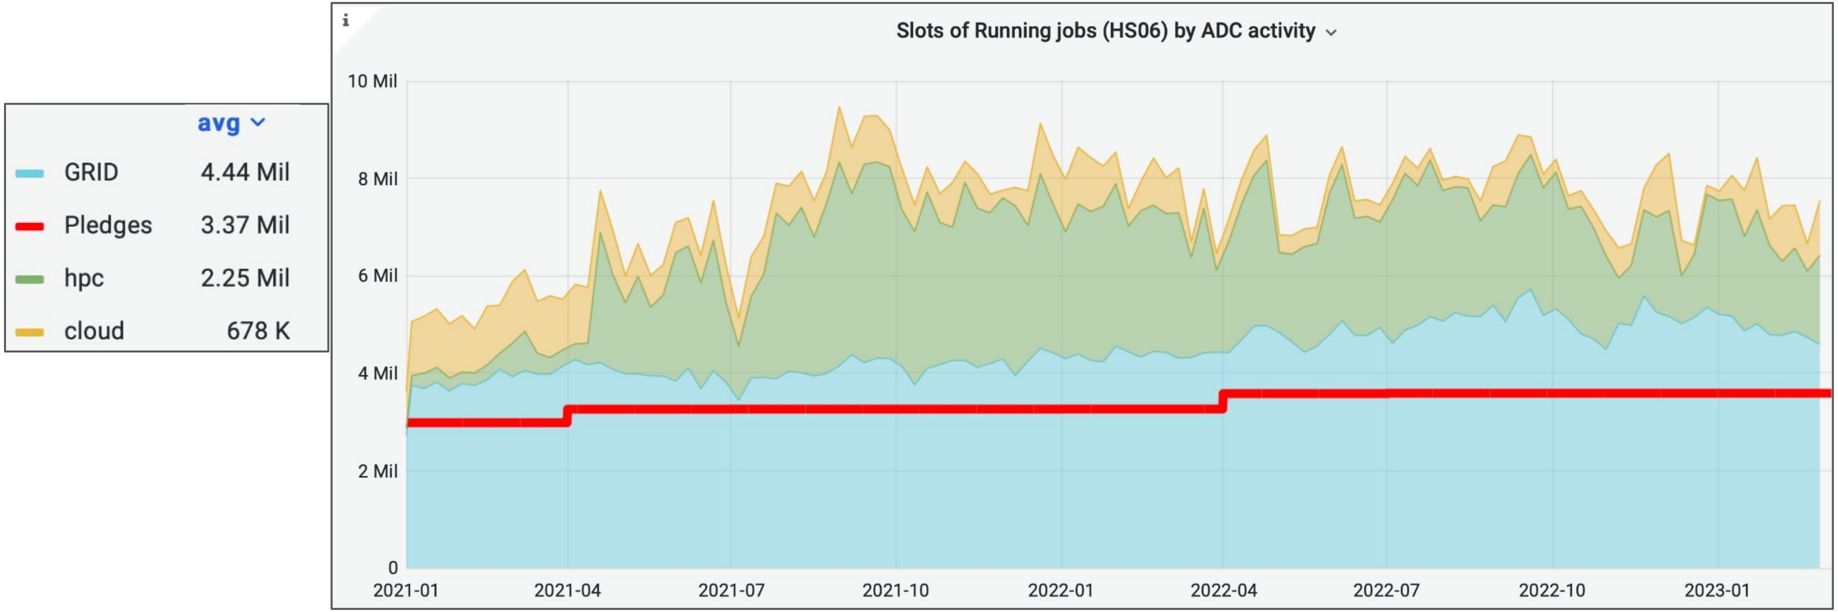

# High-Performance Computing

- We have benefited enormously from HPC systems (especially since 2021)
- Recently allowed us to cross 1M simultaneous cores for the first time

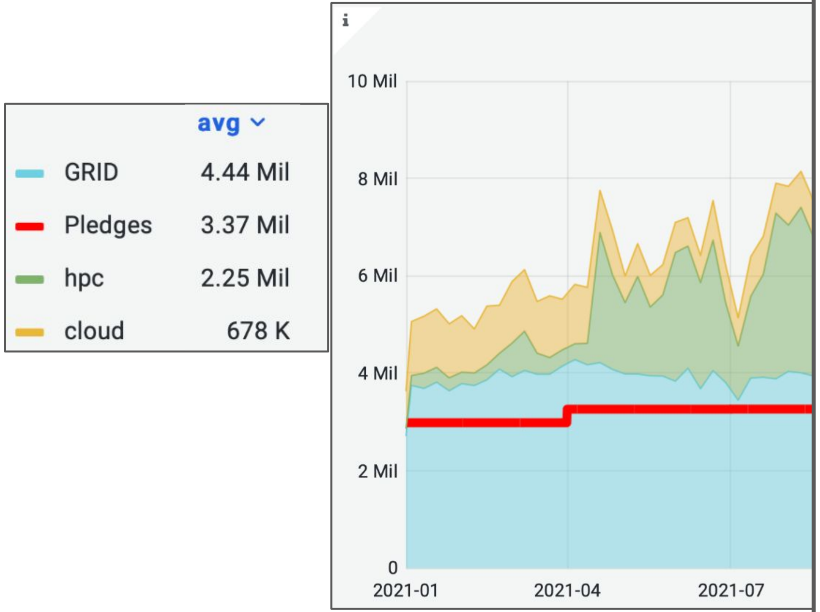


**ATLAS Experiment** @ATLASexperiment

New record! 💻 For the first time, over 1 million CPU cores simultaneously contributed to ATLAS computing.

ATLAS uses a global network of data centres to perform data processing and analysis, including HPC (supercomputers) in the US & Europe and the Worldwide LHC Computing Grid.

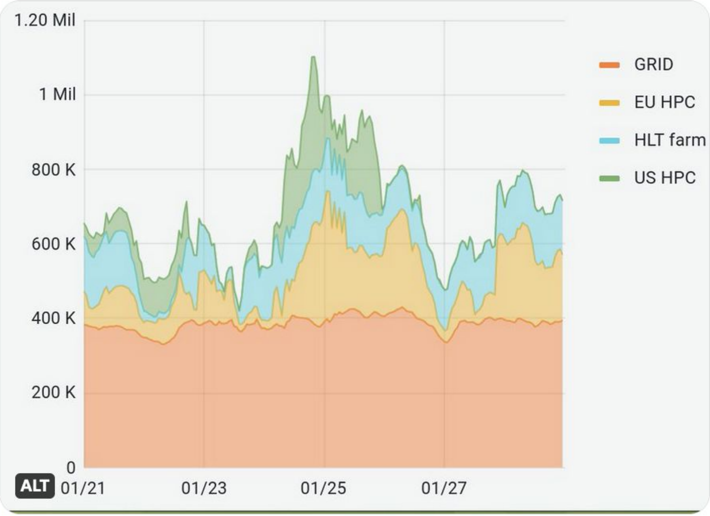


4:55 PM · Feb 7, 2023 · **5,219** Views

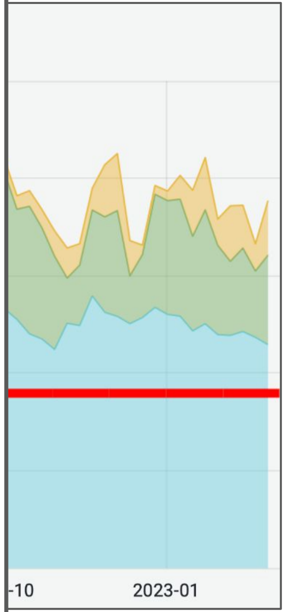

# High-Performance Computing

ATLAS EXPERIMENT

- We have benefited enormously from HPC systems (especially since 2021)
- Recently allowed us to cross 1M simultaneous cores for the first time
- We have a nice mixture of “transparent” and “complex” HPCs today
  - Transparent: Mare Nostrum, Norway, Vega, Karolina, CSCS, (Leonardo), …
  - Complex: Cori/Perlmutter, Toubkal, …

| | max | avg |
|---|---|---|
| Pledges | 3.59 Mil | 3.37 Mil |
| Vega_MCORE | 2.61 Mil | 1.02 Mil |
| Vega | 1.96 Mil | 600 K |
| praguelcg2_Karolina_MCORE | 899 K | 134 K |
| Vega_FULLNODE | 1.21 Mil | 86.8 K |
| NERSC_Cori_p2_mcore | 582 K | 70.8 K |
| CSCS-LCG2 | 109 K | 67.3 K |
| RIVR.UM | 127 K | 58.4 K |
| NERSC_Perlmutter | 4.19 Mil | 51.0 K |
| Vega_largemem | 238 K | 36.6 K |
| pic_MareNostrum4 | 197 K | 35.3 K |
| LRZ-LMU_MUC | 226 K | 28.7 K |
| NSC | 36.5 K | 25.9 K |
| IFIC_MareNostrum4 | 55.3 K | 21.6 K |
| CSCS-LCG2-ALPS | 108 K | 18.1 K |
| TACC-FRONTERA-UCORE | 203 K | 16.4 K |
| HPC2N | 43.5 K | 16.3 K |
| CERN-HPC | 37.3 K | 13.1 K |
| praguelcg2_Barbora_MCORE | 114 K | 12.2 K |
| pic_MareNostrum4_Test | 158 K | 12.0 K |
| TOKYO_HPC | 394 K | 10.4 K |
| UAM_MareNostrum4 | 50.0 K | 10.1 K |
| praguelcg2_Salomon_MCORE | 79.8 K | 9.95 K |
| LRZ-LMU_MUC_ES | 113 K | 9.07 K |
| CERN-HPC_ES | 35.8 K | 6.83 K |
| ALCF_Theta | 138 K | 6.37 K |
| UM6P | 34.7 K | 3.69 K |
| DESY-HH_HPC | 14.3 K | 3.60 K |
| MPPMU-DRACO_MCORE | 43.2 K | 3.19 K |

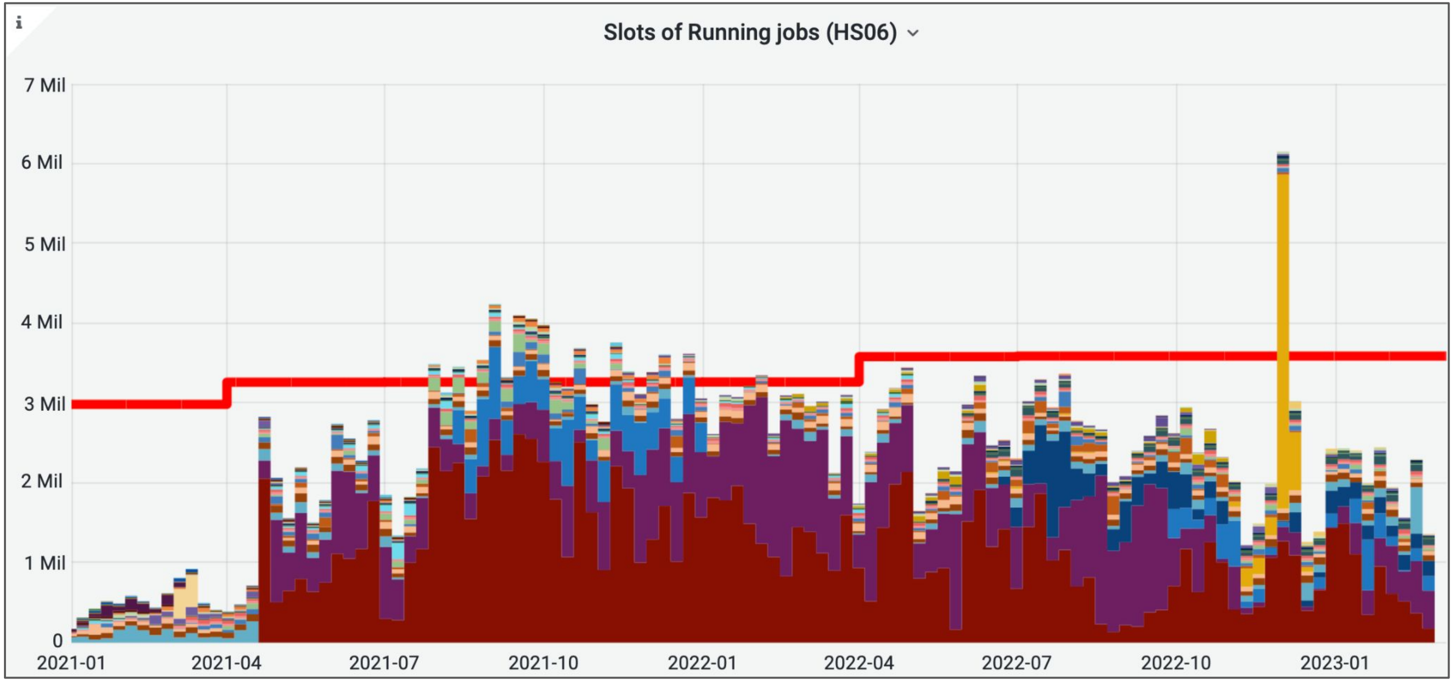

# HPC and Heterogeneity Philosophy

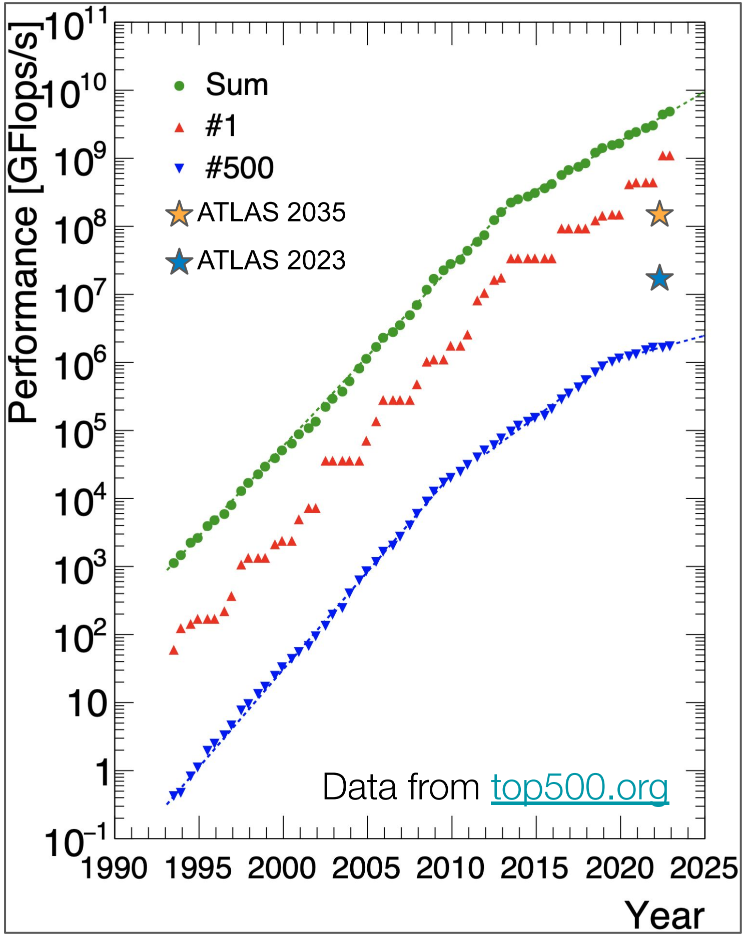


- We need more compute for the HL-LHC
- We need to keep our resources diverse
  - Expect HPCs to be an important component
- We do **not** need to run everywhere!!
  - Even a small HPC today could deliver a huge fraction of our required cycles
  - Vega is barely top 100 and is easily our #1 site — and we only use the CPU partition
  - This means we can find *the most friendly* HPC machines to use
- Corollary: we don't need to run on all hardware
  - We are validating ARM now and have some GPU developments in the pipeline
  - It appears likely that ARM+CPU *might* already be sufficient for us
  - Portability languages will be key for us to port to other hardware

# HEP-CCE Project Questions / Comments (I)

- We haven't seen a lot of visible impact (yet?) from some HEP-CCE projects
  - Maybe that's natural — HL-LHC isn't here yet, and these things are hard to deploy
  - OTOH, several ATLAS members deeply involved in HEP-CCE and doing great things
- Our Roadmap includes a decision about (portability) languages in 2024
  - This is very much because we expect to learn from HEP-CCE's recommendations in 2023
  - ATLAS members are deeply involved in this effort; we don't think you're missing anything
- We are aware of the IOS activities (and ATLAS members are involved)
  - We haven't seen a killer use-case for e.g. the HDF5 writing yet
  - When we talk about "I/O limiting" in our applications, it is *rare that we really mean the I/O layer itself*. Usually we mean I/O+decompression+ROOT overheads+access overheads, and recent measurements indicate that these last two are the most important to us.
- Offloading (e.g. to GPU) data efficiently is an ongoing project in our Roadmap
  - Perhaps looking for further overlap/collaboration here would be useful
- Intelligent data compression is certainly interesting to us
- Without thinking deeply about it:
  - Metadata is an ongoing fight for ATLAS in a variety of ways. Perhaps there's some opportunity here for collaboration / harmonization / common tooling?

# HEP-CCE Project Questions / Comments (II)

- It isn't clear to us what benefit object store R&D can bring
    - We're happy to learn more, of course
- We expect to run significant LO, NLO, and (hopefully) NNLO generation in HL-LHC
    - Accelerating any of these is welcome; we will be happy to test prototypes when ready
    - Our biggest (in CPU) sample is NLO Sherpa V+jets, followed by Powheg NLO ttbar
    - We generate many different samples with MadGraph (more configurations than Sherpa)
    - It would not be shocking (to me at least) if this were to remain the case in the HL-LHC
- GPU-based Geant4 would be a wonderful development
    - Recall that GeantV encountered *serious* difficulties when it got to hadronic physics, and neither Celeritas nor Adept has gotten there yet. Attacking this issue is important.
    - We are working on various steps to prepare ourselves (e.g. geometry offloading)
- We are aware of some of the FuncX/Parsl work
    - We don't see a big potential impact for this yet but are happy to learn more and discuss it

# Bonus Point(s)

- Significant interest in power / carbon issues worldwide
  - Is there interest in a HEP-CCE project in that space?
  - Still some time to recommend best practices before big purchases for HL-LHC

# Summary

- ATLAS is facing interesting, difficult, but solvable software and computing challenges for the HL-LHC
  - One of the biggest challenges not mentioned here is supporting and retaining skilled developers — your help is always welcome!
- Now is a great time for R&D, demonstrators, prototypes, and pilot projects!
  - From experience we know how long and painful integration in our frameworks and full physics validation are: we should take this into consideration to manage our expectations!
- Focusing our efforts on common, shared objectives is paramount
  - The *way* in which we work can make the difference between success and failure!
  - Fragmented efforts are lethal — and ineffective
- This was a “quick” overview of some of the challenges
  - Much more in our HL-LHC Roadmap, and we are happy to discuss further with anyone interested in contributing!

# Extras

# The ATLAS detector now

## Phase-I upgrade

New LAr L1 trigger

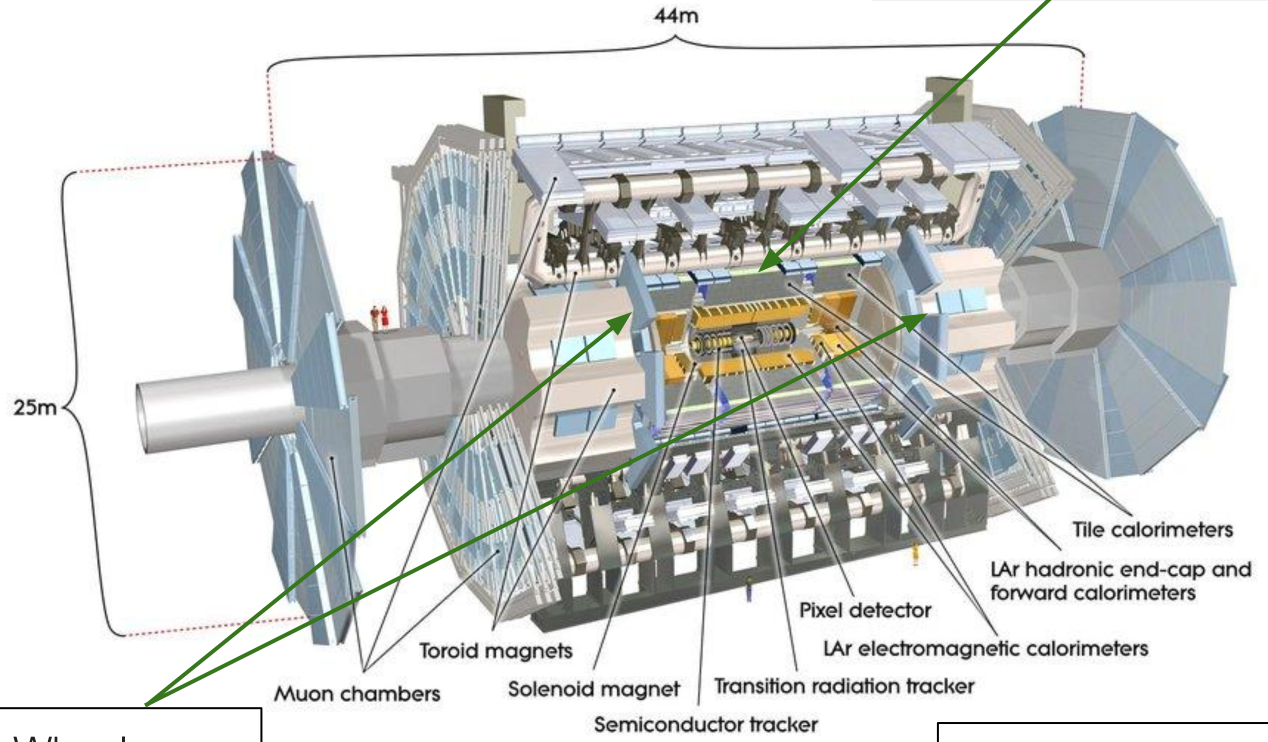


New Small Wheels

New SW: Athena Multi-Threaded

# The ATLAS detector in HL-LHC

## Phase-II upgrade

High Granularity Timing Detector (HGTD)

New muon chambers (new RPC and sMDT)

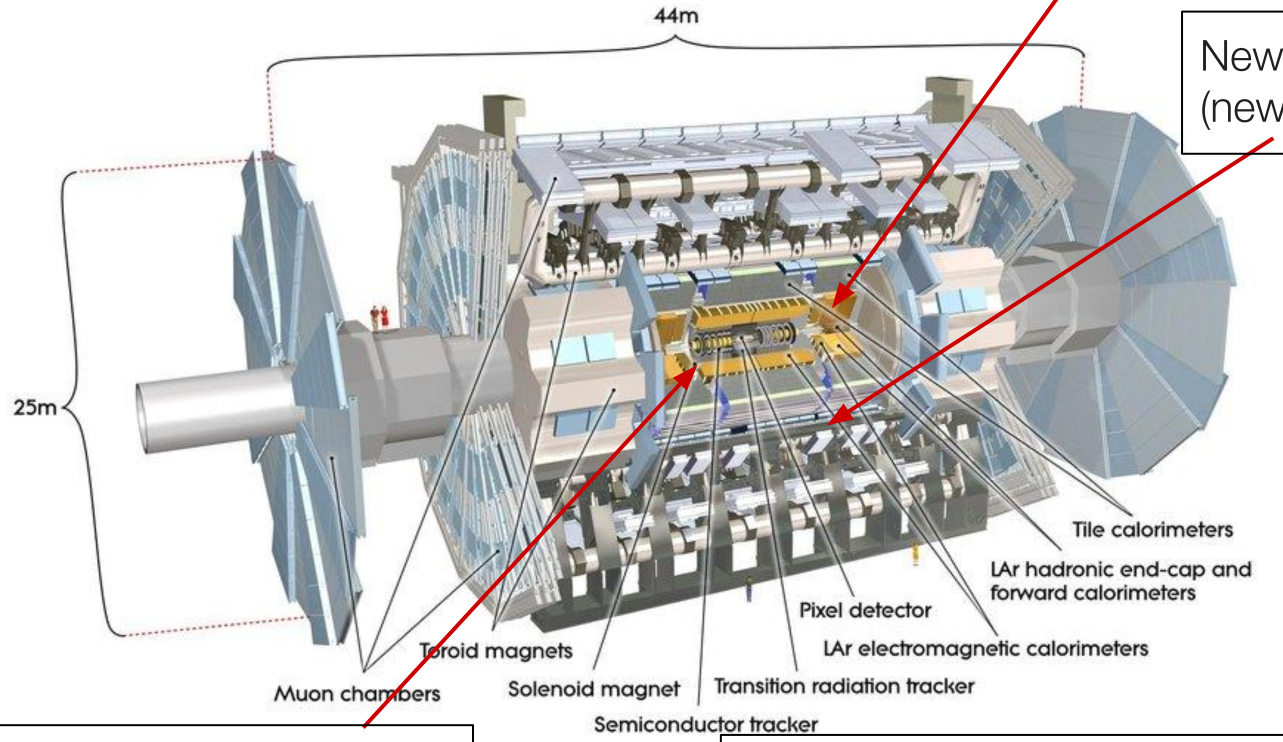


New Inner Tracking detector (ITk)

Upgraded Trigger and Data Acquisition system

# The ATLAS Collaboration


# Some caveats: lots of "business as usual"

ATLAS EXPERIMENT

- We have to keep on running the experiment while we are planning for major upgrades
  - As Karl Jakobs said, "We are building a new ATLAS while we are running ATLAS"
  - Failing is not an option
- Lot of efforts that need to go into non-R&D work (or at the boundaries)
  - Maintaining our current software
  - Updates to database infrastructure
  - Improvements in metadata handling
  - Upgrade geometry and digitization
  - SW performance improvements
  - Re-tuning of Fast and G4 simulation for Run 4 new detector, and re-tuning reco
  - Distributed computing: lots of fundamental stuff, building blocks (tokens, OS, network, storage technologies)
- R&D projects are on top of this:
  - Balance (between R&D and "business as usual") is key
  - And we need to have a strong focus on "impact"

# Power and Energy

- Interesting & useful discussion of [energy consumption during WLCG Workshop]
    - Discussions dominated up to now by ATLAS members + sites
- Happy to engage further on energy consumption, power, C02, etc
- Various positive steps in terms of energy reduction
    - ATLAS full (Geant4) simulation fully validated on ARM (and now working on evgen+reco)
    - Clearly defined list of priorities for sites in case of power-shedding needs (switching off disk should be the last resort)

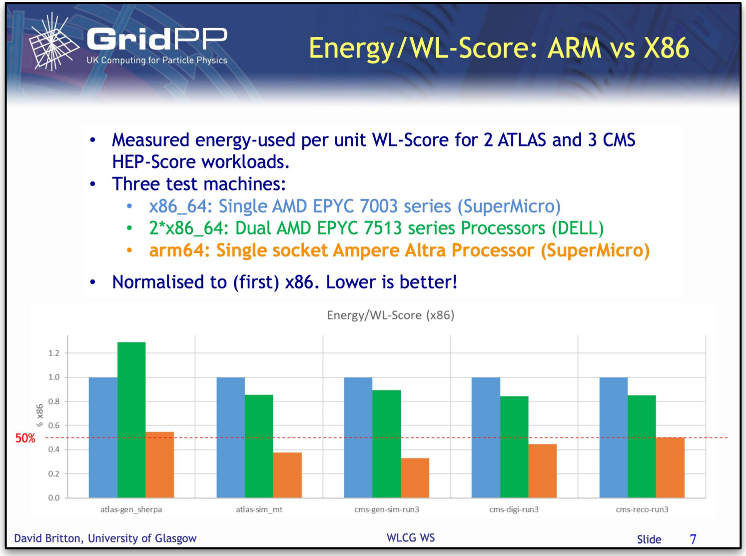


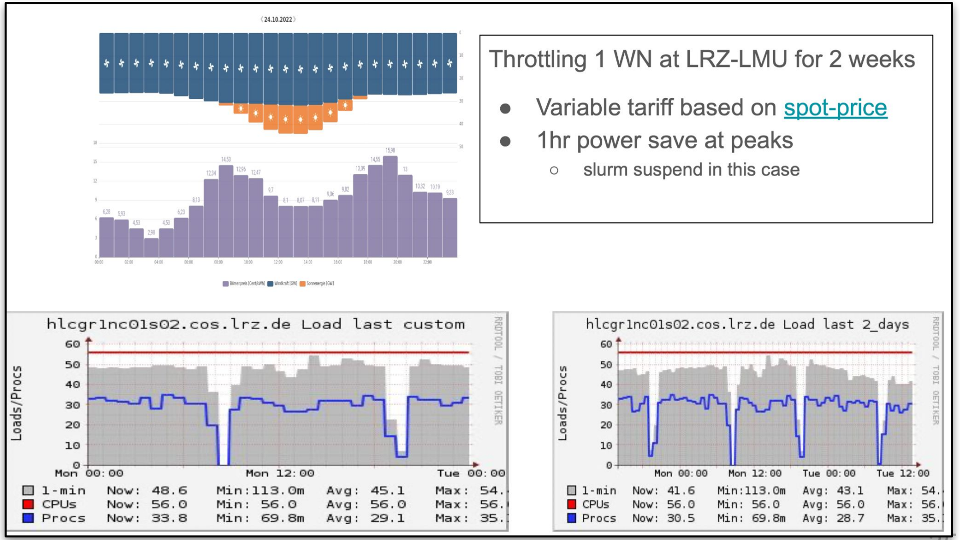

# CERN Hardware cost

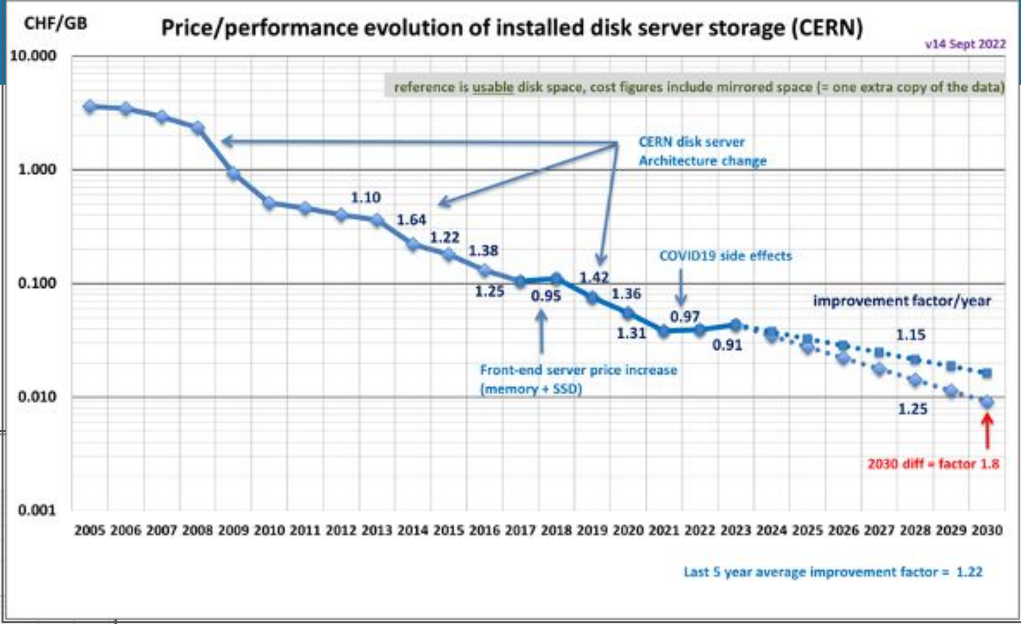


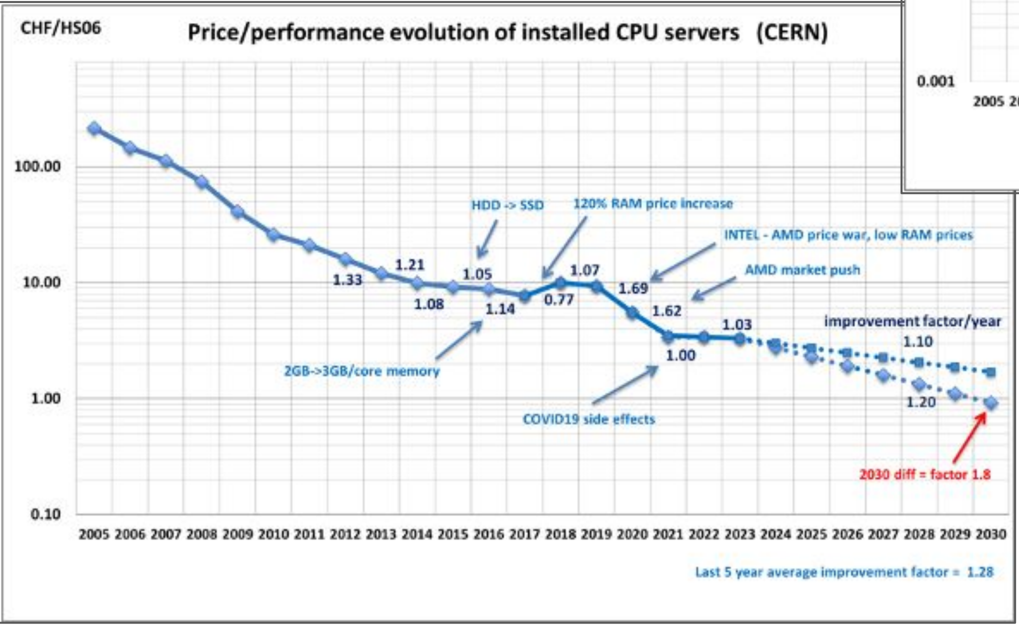

# Challenges: Tracking

- Charged particle tracking is a great example of (constantly) hot R&D
  - It is slow in the HL-LHC – it's always been one of the heavier parts of the reconstruction
- Lots of work was done in preparation for Run 3
  - The vast majority of the speed up came from physics-driven optimizations of the algorithms
  - We should be investing in some of these optimizations now to avoid wasted code optimizations
  - N.B. physicists will find ways to use the CPU again!!

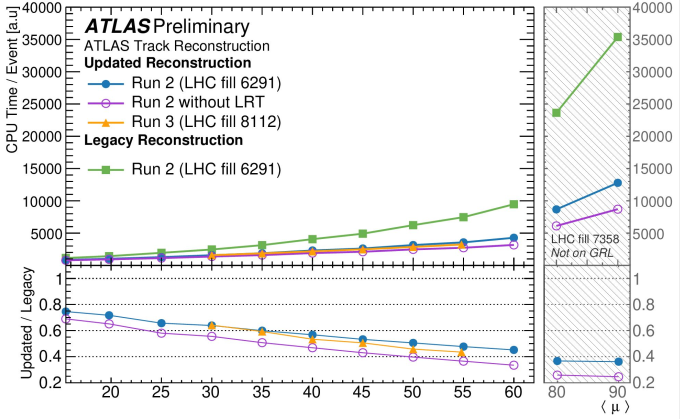


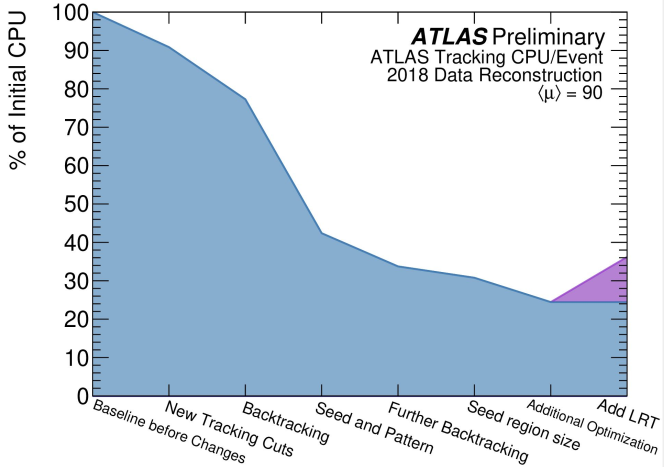

# Presentation Notes (to be deleted)

## 30 minutes talk+questions; this is already about the right volume

https://docs.google.com/document/d/1af2ndWz4MzgG2WMkeg4juk9hL53qzGAT_6Ry8wHRd-o/edit

Questions about your experiment's computing:
1) Have you updated your resource projections for the next decade? Have you identified "pain points" that require significant R&D?
2) What is the role of HPC systems in your experiment today? Are you planning to increase this role in the future? Do you envision significant changes in your computing strategy that do not involve HPC?
3) If this particular issue is a concern, what is your strategy for increasing heterogeneity (ARM vs. x86, NVIDIA vs. AMD, FPGAs, ML accelerators)?
4) What are the experiment priorities for algorithmic R&D (e.g., pattern recognition, simulation, generators, data management, and analysis pipelines), and what are the associated strategies (e.g., physics optimization, parallelization, ML, dedicated resources)?
5) What are the experiment strategies for I/O and storage optimization from physics-driven data reduction (e.g., nanoAODs) to lossy compression and network and storage-driven workflow optimization (e.g., dynamic replicas, tape vs. disk, intelligent networks)?
6) What else is important to know about your experiment computing needs?

# Presentation Notes (to be deleted)

Questions about current and future HEP-CCE work:

1) Are you familiar with the CCE evaluation of portability layers for parallel applications? If so, are the layers we are working on (Kokkos, SYCL, alpaka, OpenMP/OpenACC, std::par) the right ones to evaluate? Did we forget something?

2) We test portable parallelization layers on pattern recognition applications (patatrack and p2r from CMS, Wire-Cell Toolkit from DUNE/LArTPC) and simulation (ATLAS FastCaloSim, LArTPC Wire-Cell). Are these sufficiently representative of your resource-intensive workflows? Should we add different applications in the future (suggestions are welcome)?

3) Are you familiar with HEP-CCE's I/O and storage (IOS) activities? These include (1) an experiment-agnostic mini-app that mimics HEP application I/O patterns, (2) measuring and understanding I/O patterns of HEP applications with Darshan, and (3) an HPC-friendly HDF5 mechanism to write HEP data originally serialized using ROOT I/O. Are these relevant to your I/O issues? Can you suggest other directions that will be useful for you?

4) The HEP-CCE IOS effort started two study groups dedicated to (1) HEP event data models that are efficient for both disk I/O and accelerator offloading and (2) "intelligent" data compression (domain-specific, guaranteed precision, and recoverable precision). Are these directions relevant to your experiment R&D plans? Are there other I/O projects that CCE should focus on?

5) There are R&D activities in the HPC and HEP communities on Object Stores (e.g., HEPnOS, RNTuple/DAOS). HEP-CCE IOS has relevant expertise in this area. Would you be interested in an "event object store" R&D activity?

6) CCE contributes to two GPU-accelerated event generators, madgraph4gpu and BlockGen/Sherpa. While neither of these is ready for the experiments, it would be helpful to know if and how your experiment would benefit from an accelerated LO and NLO generator. NLO on GPU will require a significant investment to develop. To this end, on a timescale of 5 years, can you tell us which physics processes you could live generating with LO accuracy? Among the processes which will need NLO accuracy, which are the highest priorities for your experiment?

7) How important would a GPU-accelerated G4-style simulation be for your experiment? Two projects (Adept and Celeritas) are dedicated to parallelizing electromagnetic physics targeting GPUs. Should HEP-CCE play a role in bringing either or both of these to production quality and making them portable across multiple platforms?

8) The CCE complex workflows group is exploring the applicability of Parsl/FuncX to HEP workflows. FuncX is already used as a resource abstraction layer by Coffea, and ATLAS is testing it as an endpoint for their complex ML workflows. Are you familiar with this work? Could your experiment profit from a FaaS or a complete workflow execution engine that abstracts away the resources (clouds, grids, HPCs) it is running on?

9) Any other inputs about CCE current work and future plans?