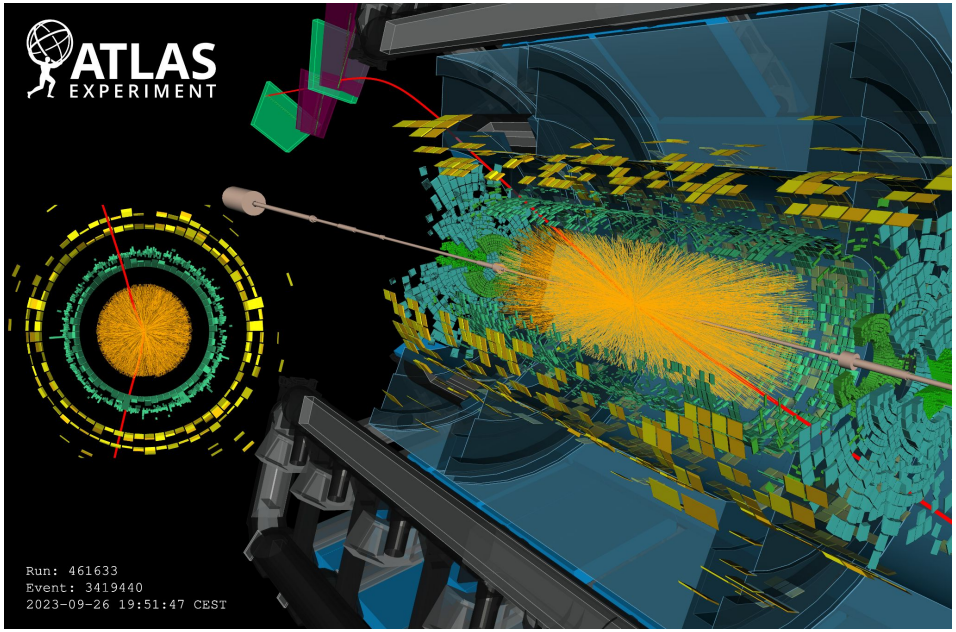

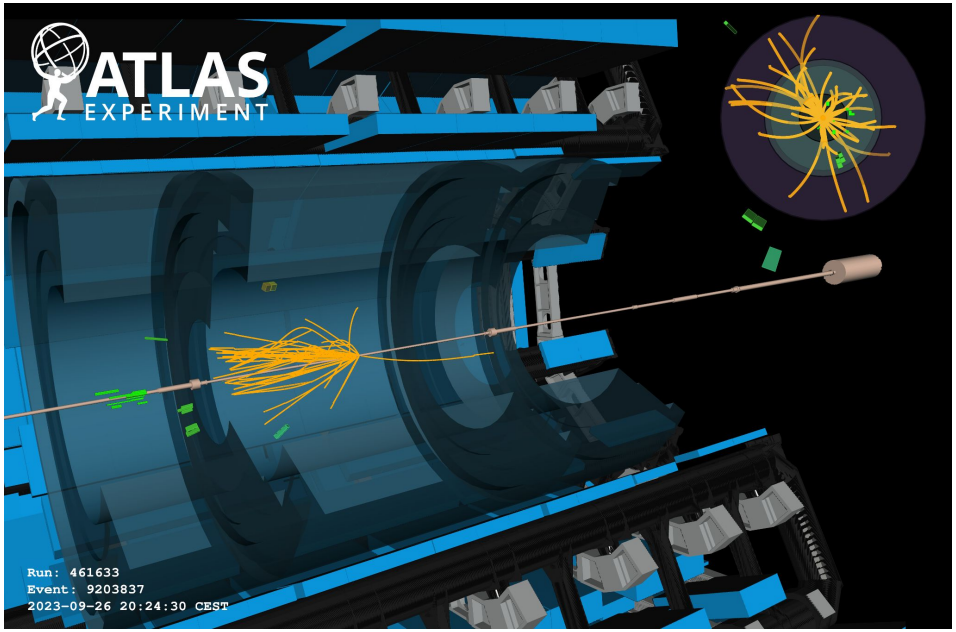

Overview and future perspectives of ATLAS Computing

Meeting annuale ATLAS Italia Computing

Genova, 28 November 2023

Zach Marshall (LBNL) and David South (DESY)

On behalf of the ATLAS S&C community

Where We Are, Where We Are Going

- We have just now finished 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; Current plans extend to 2041

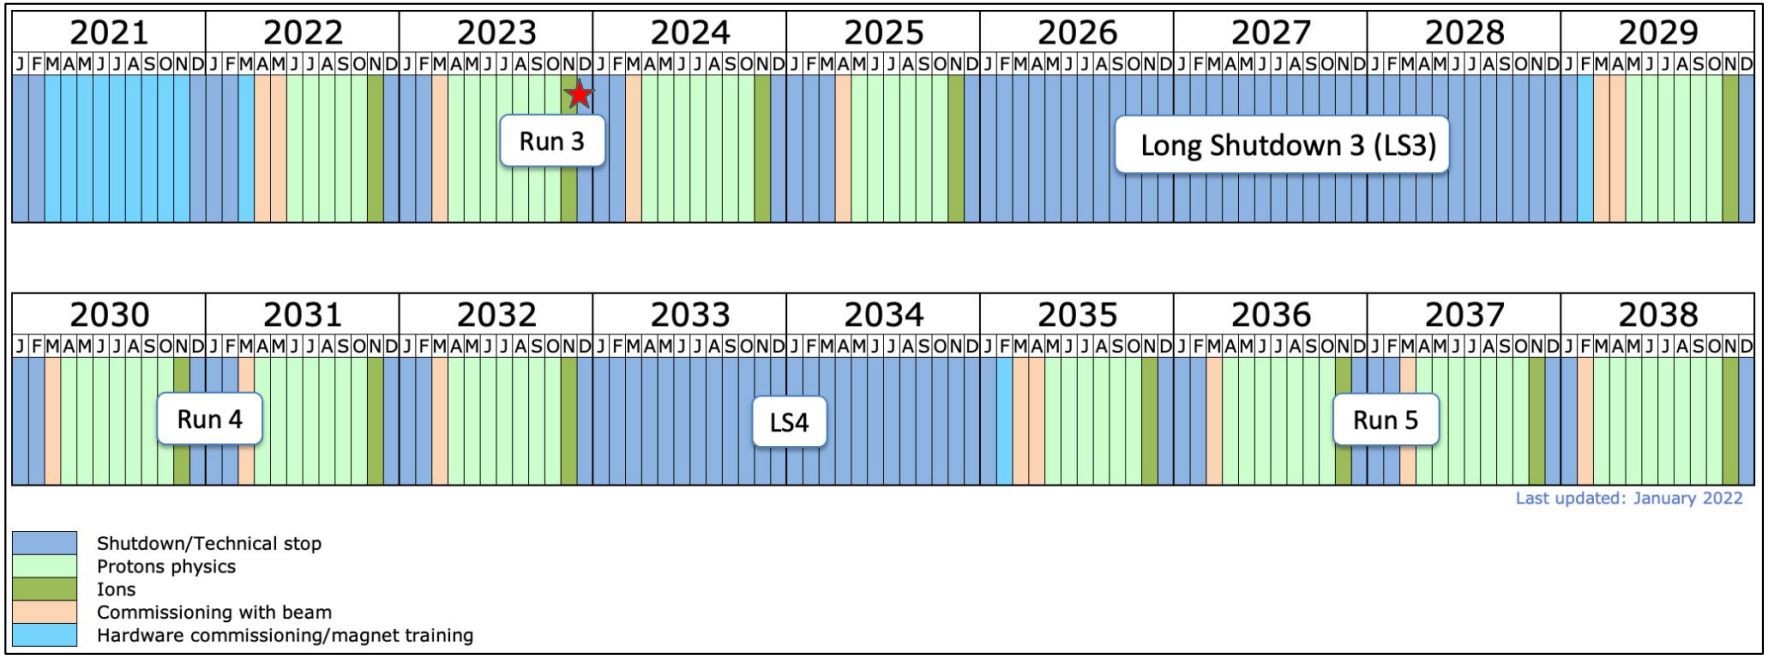

It was an Up and Down Year

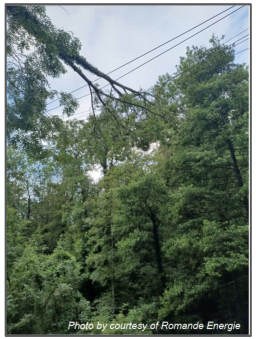

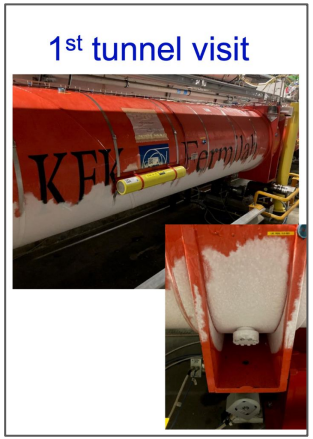

- We had some bumps along the way this year
- Didn't get quite as much data as we'd hoped for
- The ATLAS control room is still a pretty great place
- Don't forget to occasionally come visit the experiment, take a shift, see things in person, meet your colleagues

ATLAS Data Volume

- ATLAS is rapidly approaching an exabyte of data on disk+tape, but not done yet!
- We have collected about 5–8% of our total data so far (in fb^{-1})
- In terms of PB, we have a lot of growth ahead…

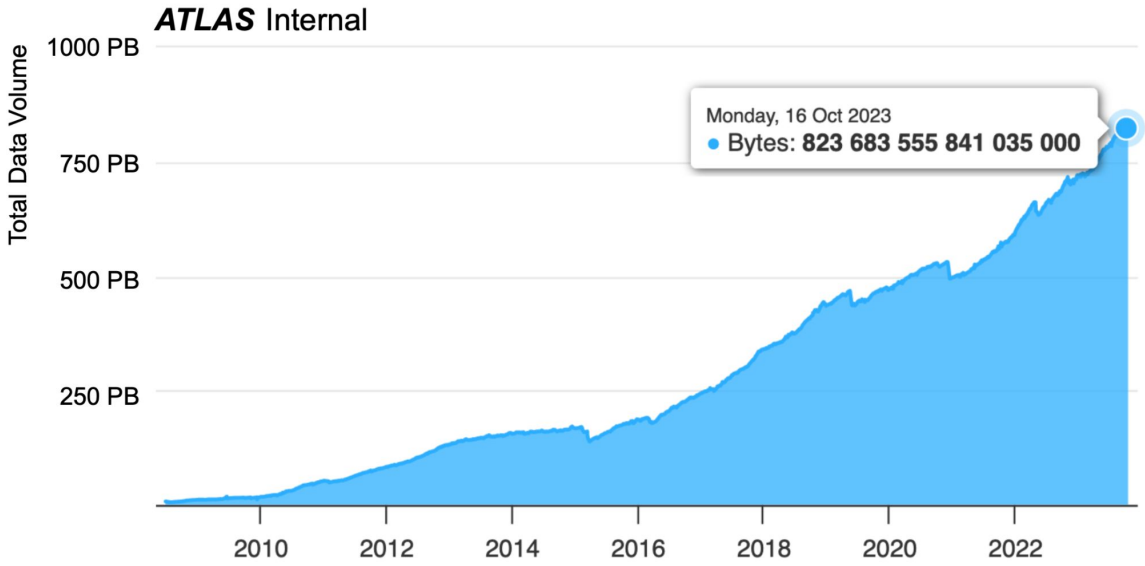

ATLAS Data Volume

- ATLAS is rapidly approaching an exabyte of data on disk+tape, but not done yet!
- We have collected about 5–8% of our total data so far (in fb^{-1})
- In terms of PB, we have a lot of growth ahead…

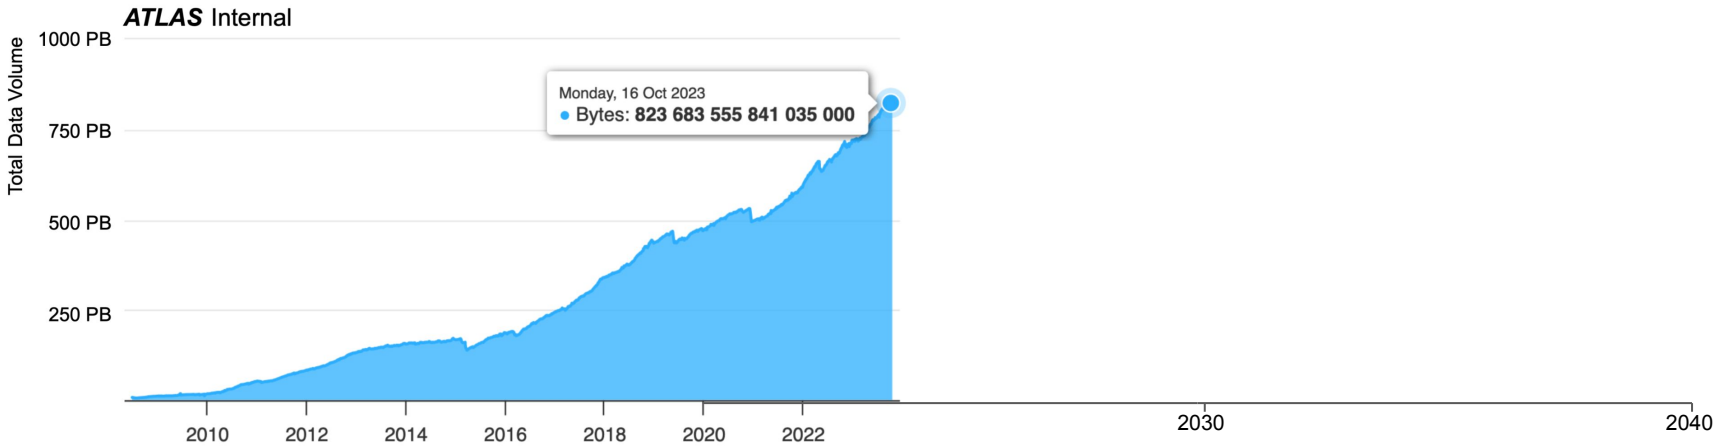

ATLAS Data Volume

- ATLAS is rapidly approaching an exabyte of data on disk+tape, but not done yet!
- We have collected about 5–8% of our total data so far (in fb^{-1})
- In terms of PB, we have a lot of growth ahead…

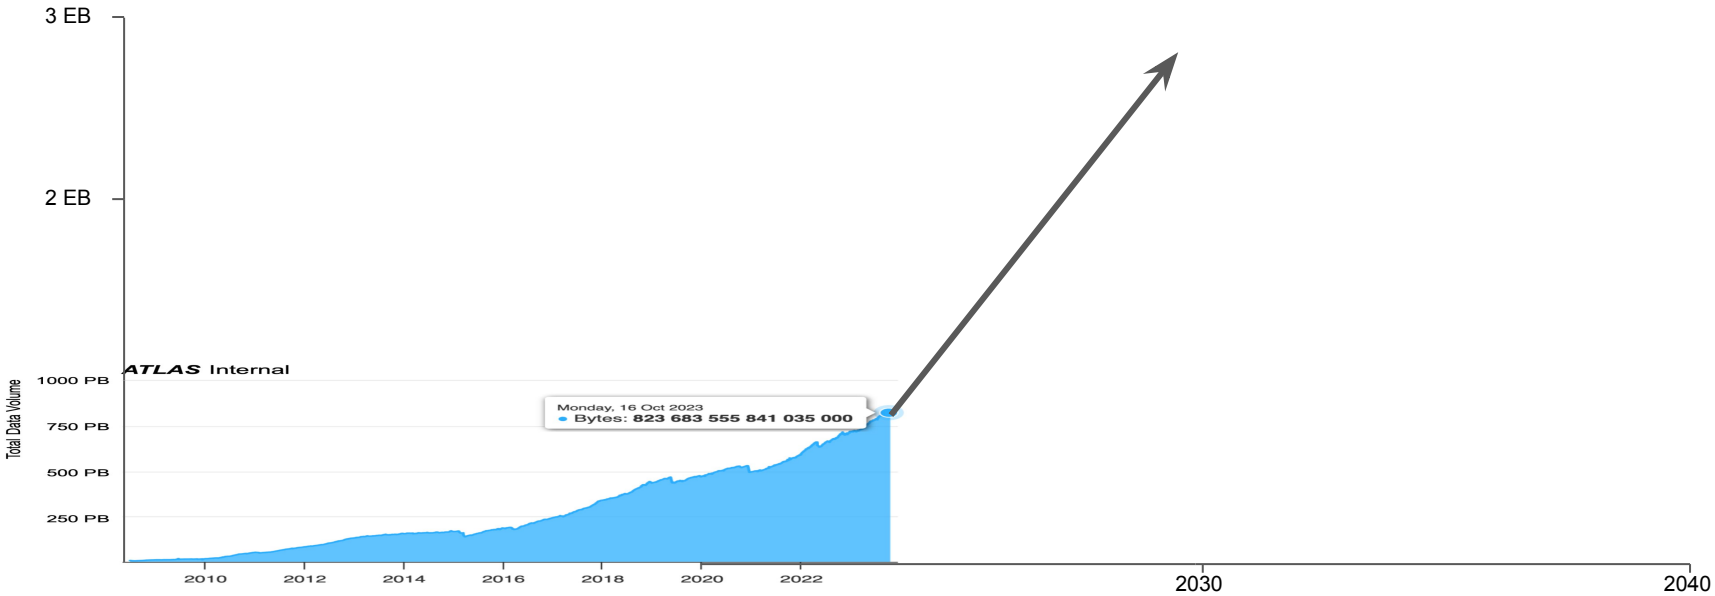

ATLAS Data Volume

- ATLAS is rapidly approaching an exabyte of data on disk+tape, but not done yet!
- We have collected about 5–8% of our total data so far (in fb^{-1})
- In terms of PB, we have a lot of growth ahead…

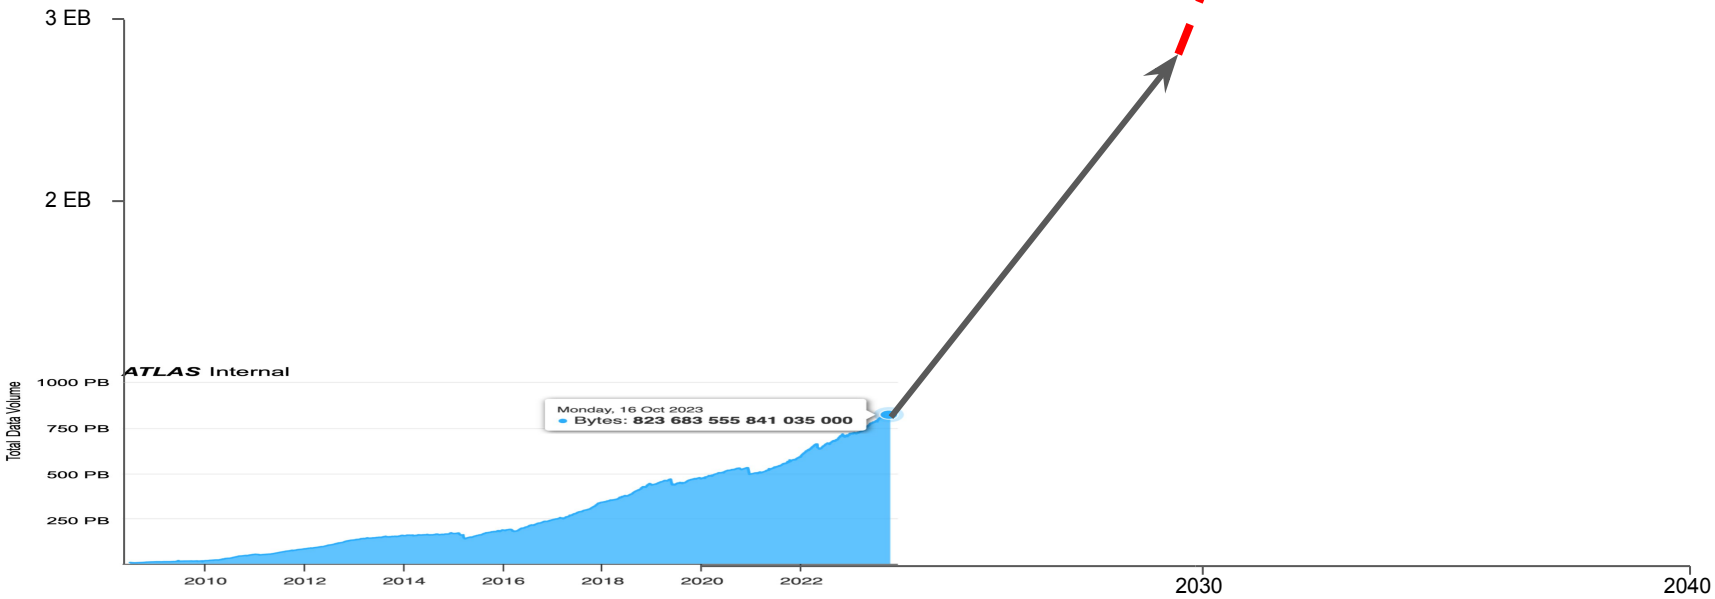

Getting There from Here

- You've all read our HL-LHC Roadmap, of course
- Lots of material about milestones and deliverables to arrive at the HL-LHC in reasonable shape
- We are planning a **TDR for early 2025** to lay out the detailed path
- Lots of ongoing **demonstrators** to test technologies and ideas
  - We still have a long lever-arm: decisions today have years of development to ensure things work well
- Expecting **two years** for integration, validation, and final performance tuning
  - If we want a simulation campaign, that means all features are included by 2026!
- **Your engagement** with the process, milestones, demonstrators, etc **is very important!**

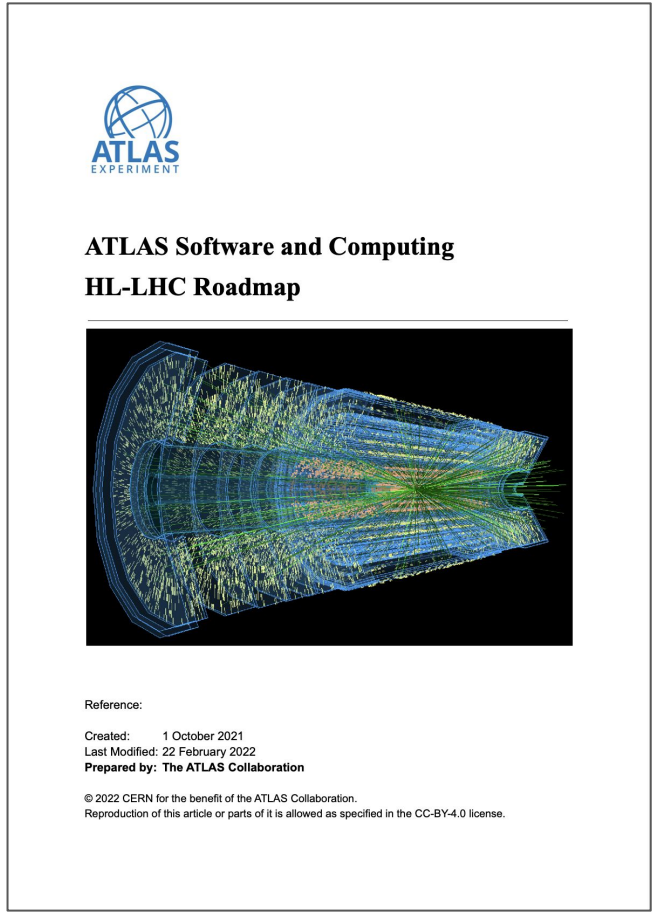

ATLAS EXPERIMENT

ATLAS Software and Computing

HL-LHC Roadmap

Reference:

Created: 1 October 2021
Last Modified: 22 February 2022
Prepared by: The ATLAS Collaboration

Lots of Work Ahead!

- Onboarding a new generation of experts
  - Ensuring ‘central’ ADC is well-equipped to help worldwide is very important!
  - Moving from VOMS to IAM for authentication
- Onboarding new sites
  - Both HPCs (more shortly) and ‘standard’ sites
  - Thank you for your help!
- **Big gains** are still possible from ‘standard’ work
  - Geant4 optimization has delivered 2x so far!
- Pushing *fast simulation* to the limit will be key
  - The “Fast Chain” could completely change our resource outlook if successful
- Machine learning has lots of potential
  - Most critical is bringing projects from “demonstration” into **real use** — this isn’t happening enough today
- Accelerators (e.g. GPUs) are on everyone’s mind

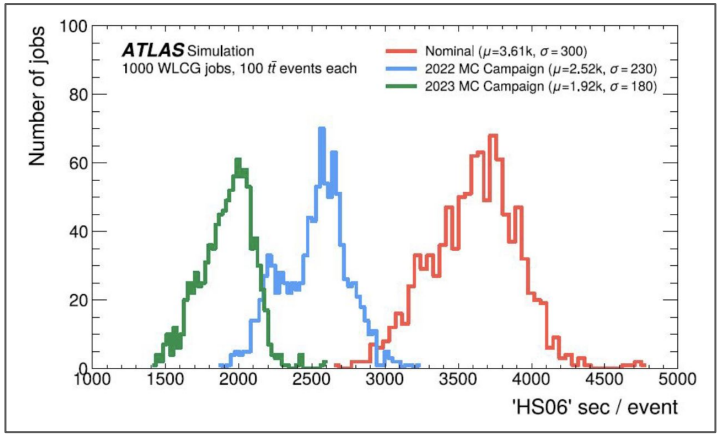

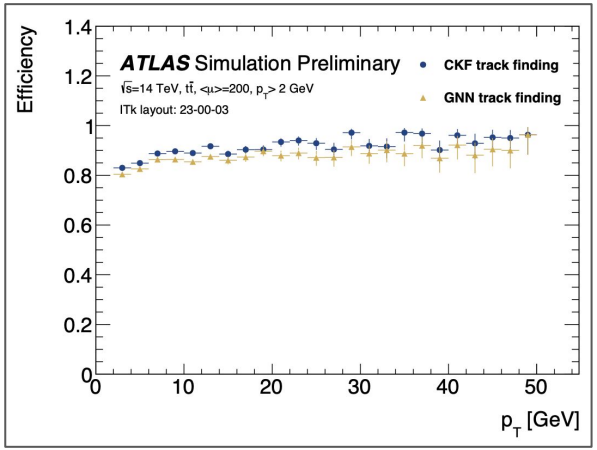

Alternative Computing: HPCs and Cloud Computing

- We have some excellent HPC partners within ATLAS!
- Looking forward to taking as much advantage of Leonardo as possible!
- "Good" HPCs are able to run *all workloads* — very important to minimize ADC work!
- Having insiders to help with edge services, setup, policies, etc has been very helpful
  - We have a pretty good toolkit now that should "just work"
- Having a multi-year CPU allocation and solid disk allocation would be fantastic
  - The long-term reliability of these CPUs is a major concern for us

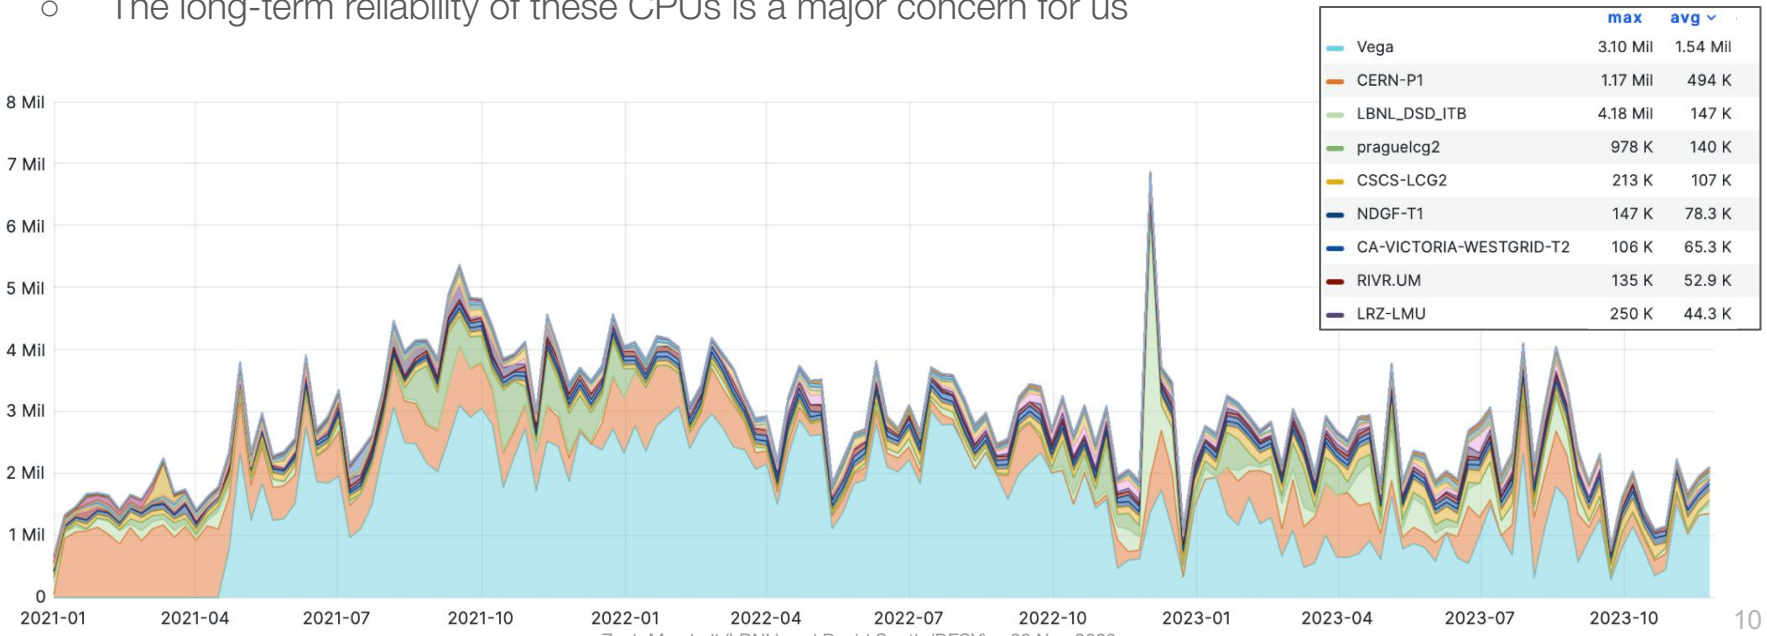

Alternative Computing: ATLAS Google Project

- 15 month project utilising Google Cloud resources
- Service Agreement (“subscription”), flat rate
- Run ATLAS workflows on site with elastic compute queue via k8s and Object Store RSE
- Many successful studies, to be published soon

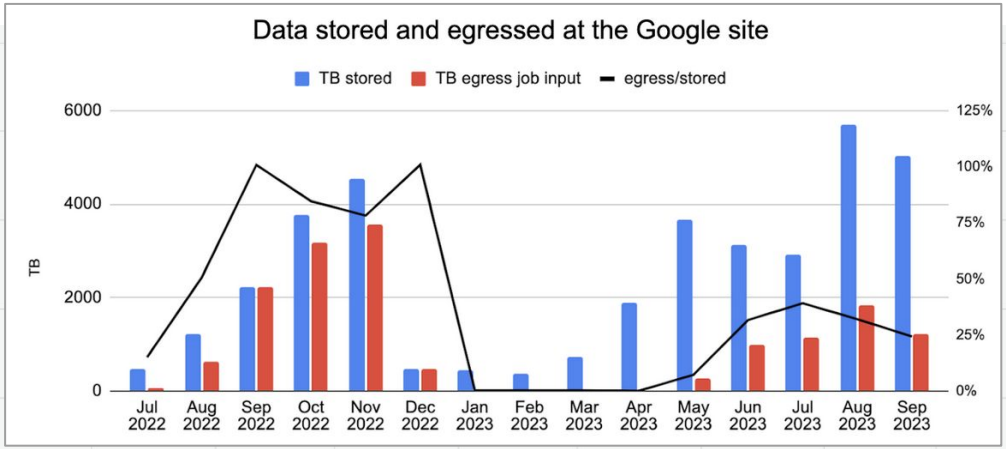

ATLAS Google Project: Activities

- R&D projects including bursting and use of GPU and ARM resources in the cloud
  - Very interesting possibility for provisioning / testing expensive or non-standard resources
- Trying to operate “like a WLCG site”
- Tests with each ATLAS workload to understand associated costs / variation with workload

	Simulation *06/01 - 19/01*	**Reconstruction** *05/02 - 09/02*	**Group Production** *11/02 - 19/02*	**Event Generation** *07/04 - 12/04*	**Data Reprocessing** *12/07 - 16/07*
Compute	76%	69%	56%	49%	26%
Storage	9%	9%	8%	30%	12%
Network egress	5%	8%	21%	2%	58%
Other	10%	14%	15%	19%	4%

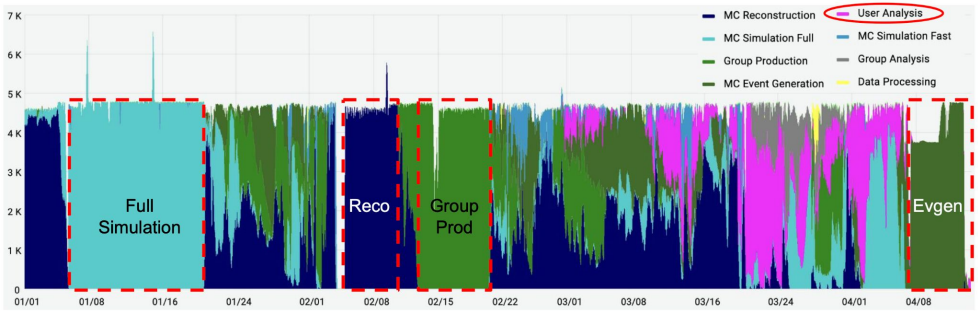

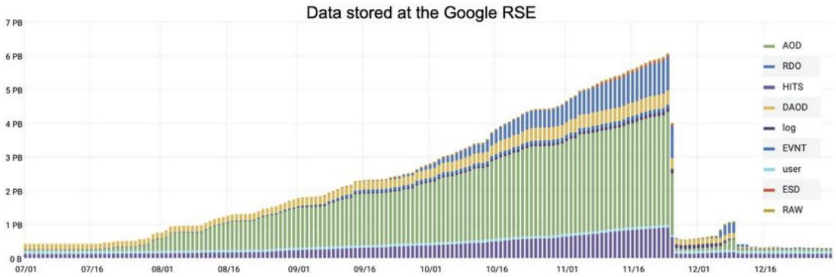

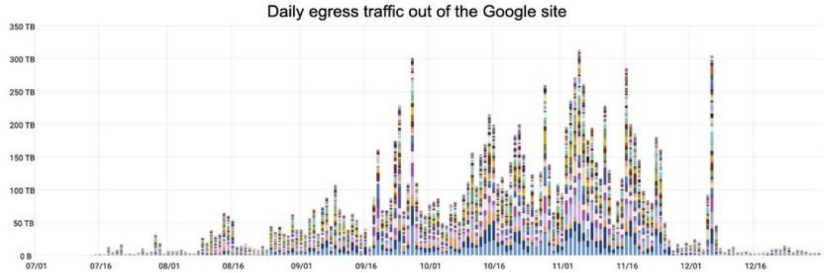

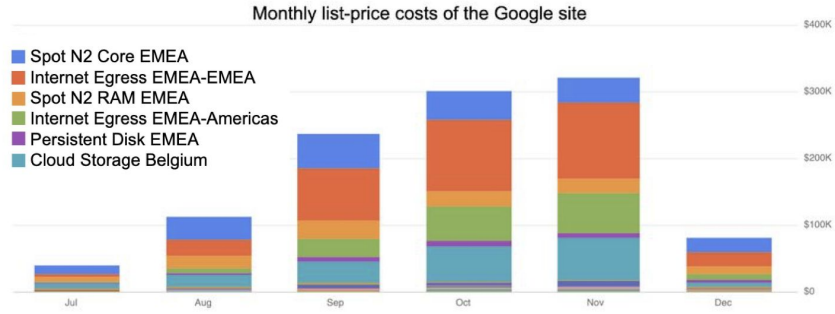

ATLAS Google Project: Outcomes

- Exploring peering options for Google egress to move data over NRENs (GEANT, ESNET, etc)
- Eviction rate (1–2%) was far lower than previously seen / feared (up to 15%)
- Stability and failure rate were quite good (5%), comparable to a good WLCG site
- Effort was made to keep things provider-agnostic, so no Google lock-in
- **Very clear and important** message that **we cannot afford to lose expertise** in ADC
  - If a site moves resources to the cloud, the effort *should not* go to zero!
- *No one* pays the list price. This means whether it's "worth it" depends on the details of the deal you can get and the internal site economics.
  - Some private comparisons suggest that in some cases it's close; and the math changes annually

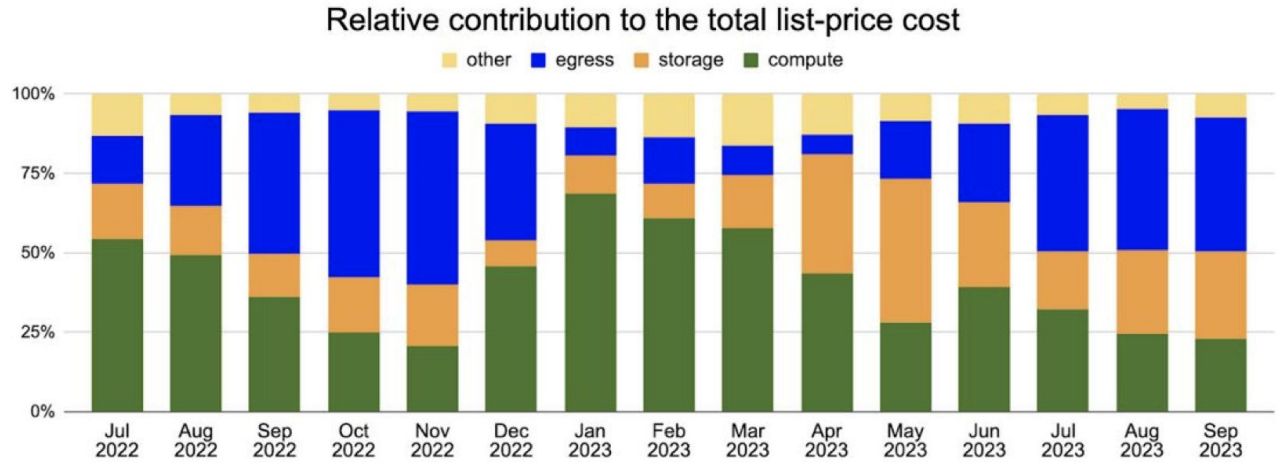

A Deployment R&D Project

- Lots of cloud / HPC implementations are already based on Kubernetes; trigger farm may be as well in 5 years
- Trigger farm in HL-LHC projected to be 𝒪(10M) HS23
- During shutdowns, we use the HLT farm for offline computing
  - Assume ⅓ of the year — 3M HS23
- HL-LHC fills are 5–6 hours, with a 1–2 hour break
  - That's for a pretty fast break (interfill) between
  - That means for a reasonable fraction of the day, the trigger farm is not in use for ATLAS
- Using 20% of the HLT farm would give us ~1.5M HS23

- All together, that could satisfy 15–30% of our CPU need
  - Has a *direct* impact on our WLCG needs
- This relies on our ability to *quickly* move onto and off of the resources
  - Such an ability could also be useful for backfill, quick ramp-up, and other features on HPC systems

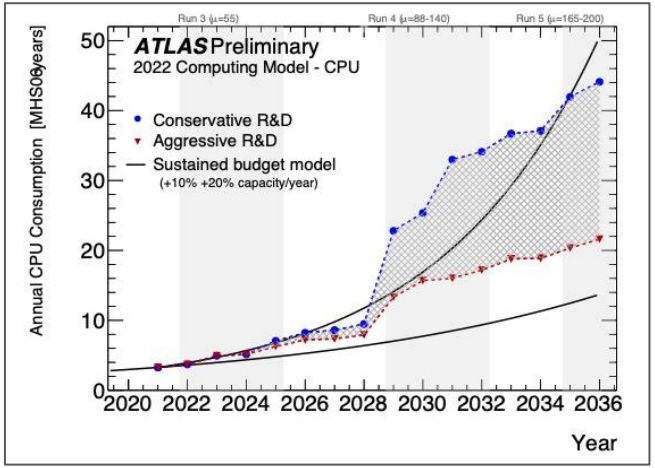

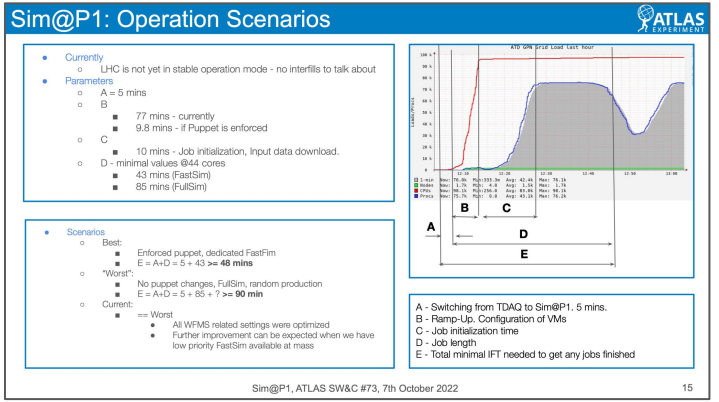

Sustainability: Growing Interest in a Great Challenge

- Sustainability at ATLAS S&C week, bringing awareness by adding carbon footprint estimates for computational jobs:
  - BigPanDA monitor task page
  - Notification emails sent to users
  - Explanation how this value is calculated can be found here

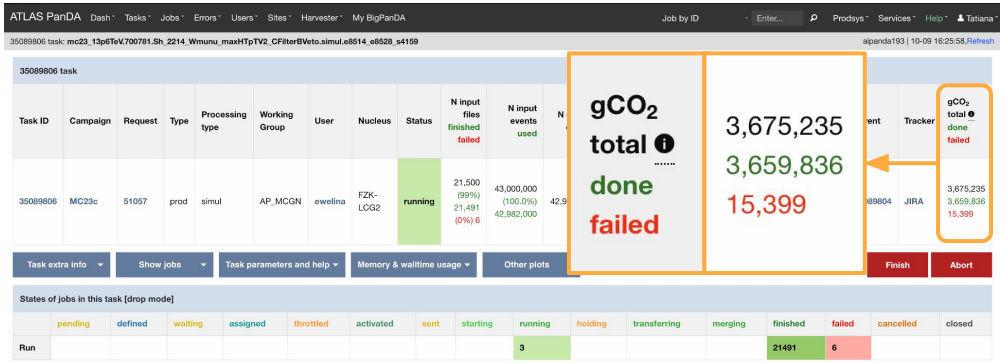

https://bigpanda.cern.ch/task/35089806/

```
Subject: JEDI notification for TaskID:34620652 (10/10 All Succeeded)
From: atlas-adc-panda-no-reply@cern.ch
To: xxx@cern.ch

Summary of TaskID:34620652

Created : 2023-08-30 14:33:06 (UTC)
Ended   : 2023-08-30 15:24:30.687287 (UTC)

Estimated carbon footprint for the task
     Succeeded   : 6.73 gCO2
     Failed      : 0 gCO2
     Cancelled   : 0 gCO2
     ____________
     Total       : 6.73 gCO2
(More details on estimation: https://panda-wms.readthedocs.io/en/latest/advanced/carbon_footprint.html)
```

- Still developing estimate of carbon footprint for storage (disks, tapes)
- Related ongoing R&D projects:
  - Power Usage Efficiency variations
  - CPU power modulations
  - Low energy platforms (ARM)
- The goal is transparency and to make clear recommendations to sites and developers in future

Approaching Storage Analytics as a Scientist

- Our computing *uses* significant carbon
- The first thing we want to eliminate first is *waste*
  - Understanding unused data and job failures
- Bridging the computing and physics communities is critical to understand and solve this problem
  - We have seen that active site admins can make a big difference in helping users facing difficulties
  - We've seen users doing very surprising things — but it's only surprising if you know what "normal" is
- Lots of analytics
  - We have the data! We need more people looking
- Big impacts on the future
  - Changing the numbers of replicas or versions, cache size, our ability to reproduce data on demand, etc can have *huge impacts* on our HL-LHC footprint

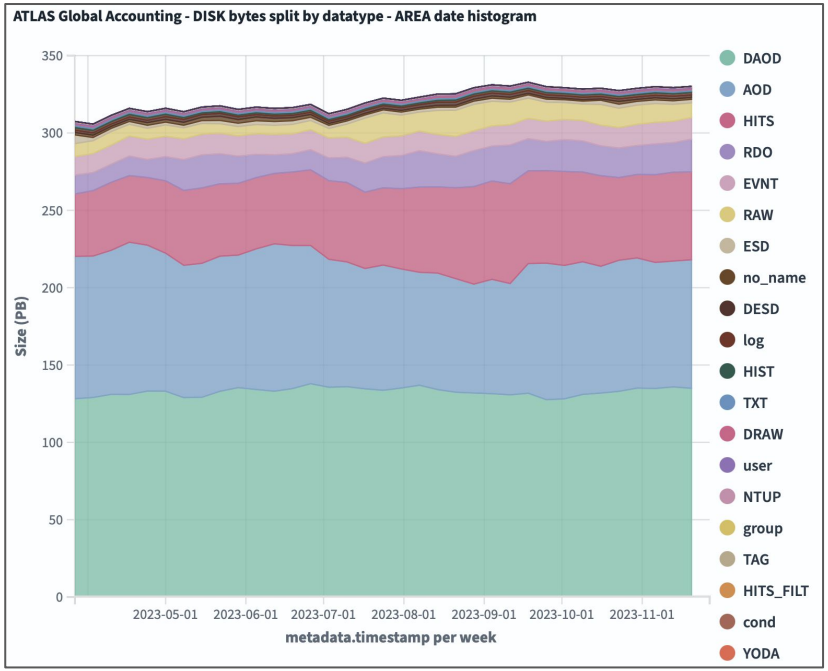

Analysis Facilities: What Does It Mean?

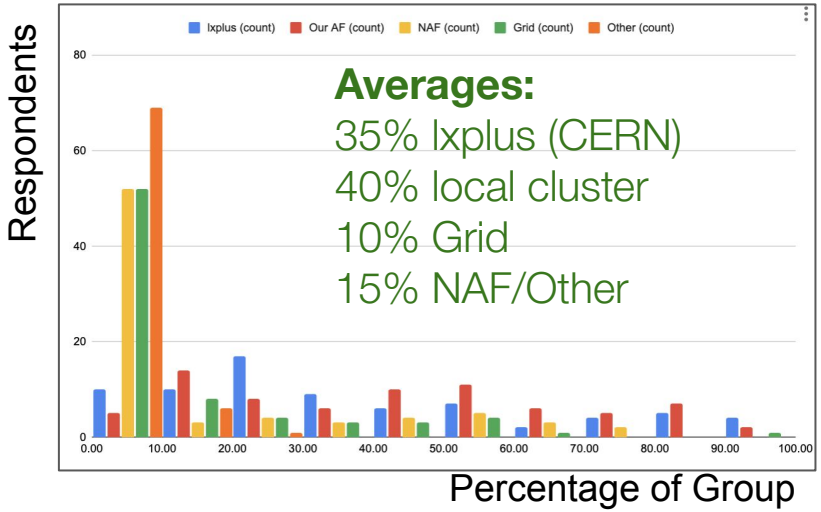

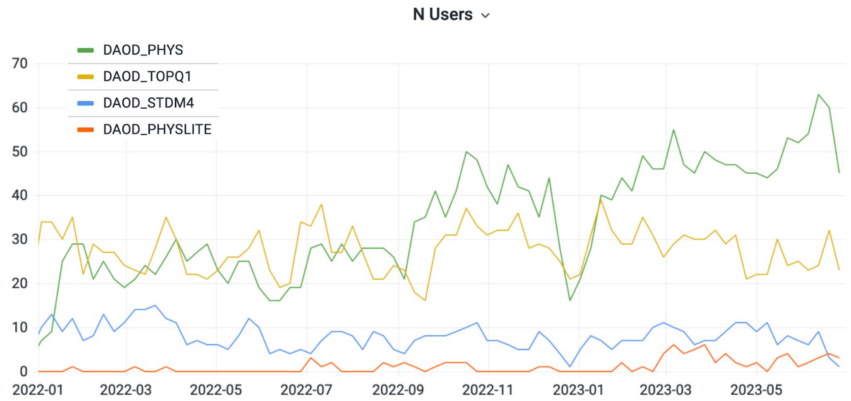

Figure 4: Number of unique users in a given week accessing the three most popular physics derivation formats, and PHYSLITE.

- Lots of recent enthusiasm around *analysis facilities*
  - The phrase, of course, can mean anything
- Focus is on how we'll do analysis in five years
  - This is **extremely** hard to predict, most of all because it requires **understanding the evolution of our physics programme**
- Lots of people still rely on lxplus (CERN) for analysis
  - Understanding ***why*** is key to the success of a NAF
- There are a lot of resources in play worldwide
  - ~50 PB of disk, 1M HS23, 100s of GPUs
- **Working with your users** is key to success here
  - Most importantly, not just working with the super-expert users, but with *all* the users
  - We see again and again that super-experts have cool solutions, but most users rely on basics
- Still, a big part of the model is **DAOD_PHYSLITE**
  - Lossy compression, efficient tools, columnar analysis are all important parts of the effort
  - Also key: understanding what users need when it isn't in PHYSLITE (augmentation, alternatives, etc)

Summary / Conclusions

- ATLAS doesn't have *just one* problem to solve
  - That's good — it means we have lots of interesting problems, and no single huge risk!
- There are lots of exciting opportunities and projects to tackle for the HL-LHC
  - This might be the biggest computing challenge our field will face for 50 years
- Most importantly, please help ensure that **your** projects and interests are communicated to ATLAS, so that we all can help each other!
  - **Faster** alone, **further** together!

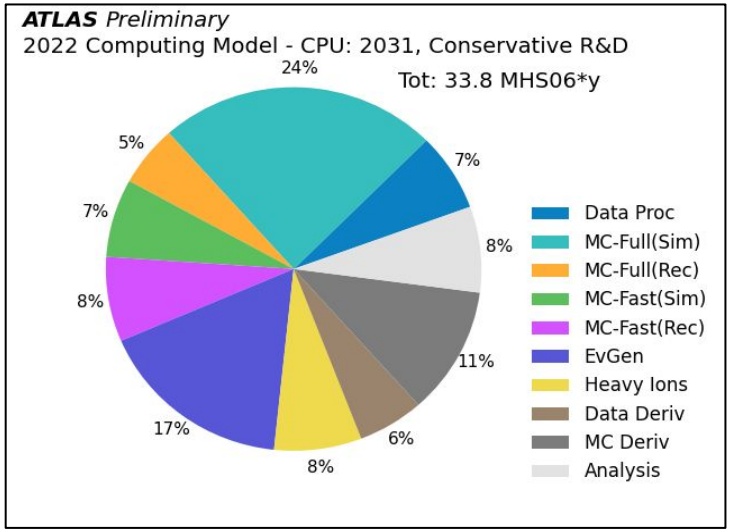

Extras

ARM resources

- Lots of discussion on this last week
  - Is it a game changer?
- Study ARM resources with Glasgow
  - Dave talk, Dwayne talk
  - Around 1600 cores available
  - 50M event validation sample almost finished
- Studies also at BNL - Imran's talk

Systems were evaluated during HEPscore23 runs

ARM 128 core	ARM 64 core	X86_64 Xeon
HEPscore – 2030	HEPscore – 1042	HEPscore – 1278
Max power – 370W	Max power – 252W	Max power – 601W
Idle – 125W	Idle – 107W	Idle – 179W
HEPscore / Watt – 5.49	**HS / W – 4.13**	**HS / W – 2.13**

What's the best way of going about this in the future

At some point we will want to get rid of the test CE

Option 1

- Modify exiting Queue

UKI-SCOTGRID-GLASGOW_CEPH
ce01.gla.scotgrid.ac.uk (x86)
ce02.gla.scotgrid.ac.uk (x86)
ce03.gla.scotgrid.ac.uk (x86)
ce04.gla.scotgrid.ac.uk (x86)
ce05.gla.scotgrid.ac.uk (aarch64)

- Dangerous, will impact the workflow of many VO's. Will have all the ARM traffic on one CE - not scalable in the future?

Option 2

- Add a queue pointing to the same CE's

UKI-SCOTGRID-GLASGOW_CEPH
ce01.gla.scotgrid.ac.uk (x86)
ce02.gla.scotgrid.ac.uk (x86)
ce03.gla.scotgrid.ac.uk (x86)
ce04.gla.scotgrid.ac.uk (x86)

UKI-SCOTGRID-GLASGOW_ARM
ce01.gla.scotgrid.ac.uk (aarch64)
ce02.gla.scotgrid.ac.uk (aarch64)
ce03.gla.scotgrid.ac.uk (aarch64)
ce04.gla.scotgrid.ac.uk (aarch64)

- If Job requirements can be successfully injected this seems the safest option

Option 3

- Read architecture from the jobs themselves.

UKI-SCOTGRID-GLASGOW_CEPH
ce01.gla.scotgrid.ac.uk (x86, aarch64)
ce02.gla.scotgrid.ac.uk (x86, aarch64)
ce03.gla.scotgrid.ac.uk (x86, aarch64)
ce04.gla.scotgrid.ac.uk (x86, aarch64)

- Condor_submit has architecture flags. Could try to pulling the architecture flag from condor_submit into the ARCsub, maybe modify in job description language (JDL)?
- Would potentially need to set up/inject default architecture so that "standard" x86 jobs don't get sent to ARM cores.

Option 4

- Pilots report to VO what architecture it's running on.
- That VO sends jobs of that type. Would require every VO to add this functionality to their pilots, not really default-able.
- Potentially wasteful if a site gets a pilot running on an ARM server and has no ARM work, long term solution?

Dwayne Spiteri, University of Glasgow — Glasgow ARM Cluster: Experience and Findings — 19

GridPP UK Computing for Particle Physics

HEPScore/Watt

Machine	CPU	Threads	HS/Watt	
Test AMD	1x EPYC 7643 HT	96	3.7	
Std AMD WN	2x EPYC 7513 HT	128	4.2	
AMD Bergamo	2x EPYC 9754 HT	512	6.3	50% improvement from 7513 WN
Test ARM	1x ALTRA 80	80	5.6	
Farm ARM	2x ALTRA 80	160	5.2 ??	
Ampere Max	1x ALTRA Max	128	7.0	25% improvement from 80c

Note added: We use Average Watts to calculate HS/Watt. Other people also use Max Watts, so care needs to be taken comparing numbers.

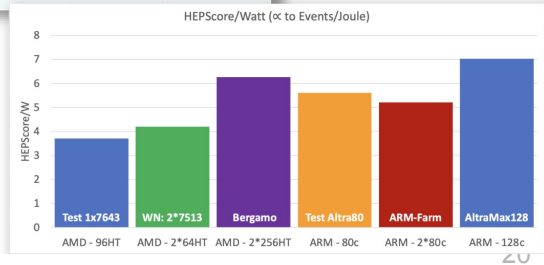

Looking forward: The road to HL-LHC

- LHCC performs a series of reviews of the Software and Computing plans of the LHC experiments towards HL-LHC
  - The ATLAS HL-LHC Computing Conceptual Design Report was published in May 2020
- A follow up ATLAS Software and Computing HL-LHC Roadmap was published in March 2022 with clearly defined *milestones*

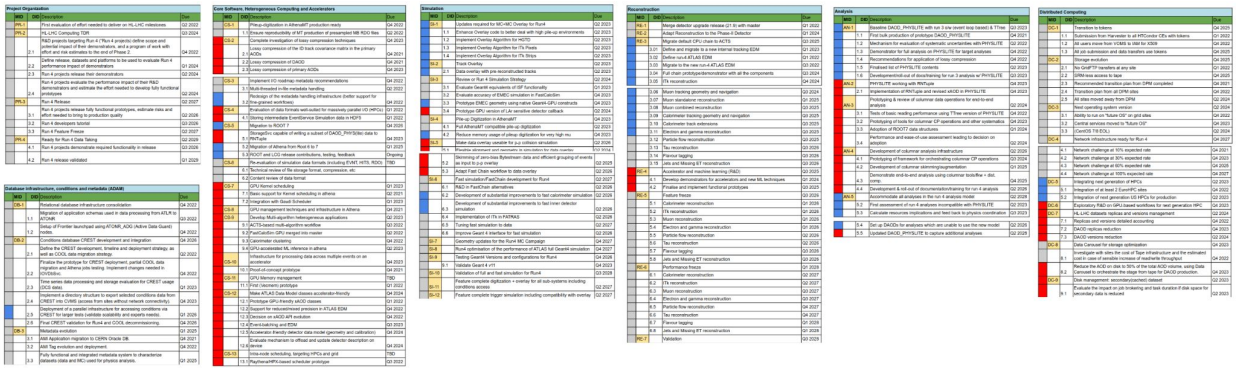

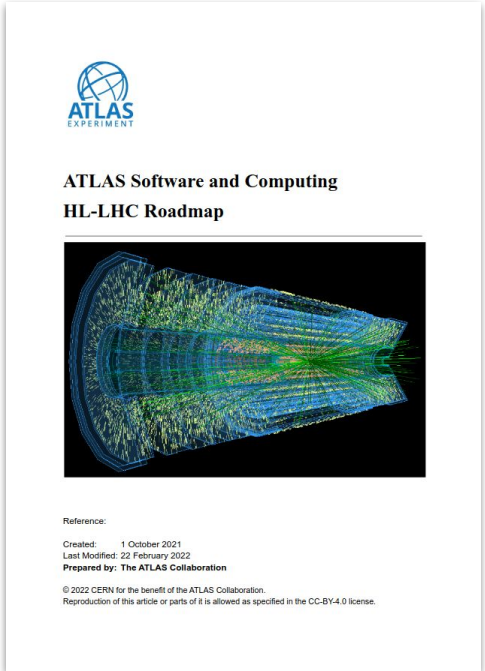

- An ATLAS HL-LHC Computing TDR is planned for 2024

ATLAS Distributed Computing milestones

Distributed Computing				
	MID	DID	Description	Due
	DC-1		Transition to tokens	Q4 2025
		1.1	Submission from Harvester to all HTCondor CEs with tokens	Q1 2022
		1.2	All users move from VOMS to IAM for X509	Q4 2022
		1.3	All job submission and data transfers use tokens	Q4 2025
	DC-2		Storage evolution	Q4 2025
		2.1	No GridFTP transfers at any site	Q1 2022
		2.2	SRM-less access to tape	Q4 2025
		2.3	Recommended transition plan from DPM completed	Q4 2021
		2.4	Transition plan from all DPM sites	Q4 2022
		2.5	All sites moved away from DPM	Q2 2024
	DC-3		Next operating system version	Q2 2024
		3.1	Ability to run on "future OS" on grid sites	Q4 2022
		3.2	Central services moved to "future OS"	Q4 2023
		3.3	(CentOS 7/8 EOL)	Q2 2024
	DC-4		Network infrastructure ready for Run 4	Q4 2027
		4.1	Network challenge at 10% expected rate	Q4 2021
		4.2	Network challenge at 30% expected rate	Q4 2023
		4.3	Network challenge at 60% expected rate	Q4 2025
		4.4	Network challenge at 100% expected rate	Q4 2027
	DC-5		Integrating next generation of HPCs	Q2 2023
		5.1	Integration of at least 2 EuroHPC sites	Q4 2022
		5.2	Integration of next generation US HPCs for production	Q2 2023
	DC-6		Exploratory R&D on GPU-based workflows for next generation HPC	Q4 2023
	DC-7		HL-LHC datasets replicas and versions management	Q2 2024
		7.1	Replicas and versions detailed accounting	Q4 2022
		7.2	DAOD replicas reduction	Q4 2023
		7.3	DAOD versions reduction	Q2 2024
	DC-8		Data Carousel for storage optimization	Q4 2023
		8.1	Investigate with sites the cost of Tape infrastructure and the estimated cost in case of sensible increase of read/write throughput	Q4 2022
		8.2	Reduce the AOD on disk to 50% of the total AOD volume, using Data Carousel to orchestrate the stage from tape for DAOD production.	Q4 2023
	DC-9		Disk management: secondary(cached) dataset	Q2 2023
		9.1	Evaluate the impact on job brokering and task duration if disk space for secondary data is reduced	Q2 2023

Maintenance & Operations | Conservative R&D | Aggressive R&D

- Several HL-LHC milestones related to distributed computing
  - Not a static list! Regularly reviewed, updated and/or expanded
  - New milestones defined since the roadmap publication
  - Essentially two types: Maintenance and Operations, R&D
- ADC Maintenance and Operations Milestones
  - These are essential changes, needed “just to get by”
    - *Tokens, evolution of storage technologies/access protocols, OS changes, network/data challenges, ..*
  - More details on these milestones in the ATLAS report at the WLCG Workshop in Lancaster last November and in other talks this week
- ADC R&D Milestones
  - Conservative R&D: New developments achievable with current effort
  - Aggressive R&D: New developments requiring extra effort
  - These are translated into “Demonstrators”, described in the following

The road to Run-4 (R2R4)

- ADC has 14 items on the R2R4 work plan

DC-1	Transition to tokens	*Job submission, queues, IAM, transfers*
DC-2	Storage evolution	*GridFTP, DPM, SRM-less tape*
DC-3	Next operating system version	*CentOS EOL, future OS*
DC-4	Network infrastructure ready for Run 4	*Data Challenges, SDNs*
DC-5	Integrating next generation of HPCs	*EuroHPC, US HPC*
DC-6	Exploratory R&D on non-x86 resources and next generation HPC	*ARM, GPU*
DC-7	HL-LHC datasets replicas and versions management	*DAOD versions/reduction/recreation*
DC-8	Data Carousel for storage optimisation	*Tape cost, AOD reduction, smart writing*
DC-9	Disk management: secondary(cached) dataset	*Impact on job brokering*
DC-10	Evaluation of commercial cloud resources	*Google, Burst scheduling, TCO*
DC-11	Optimising the user analysis experience	*Site evaluation, tails, waiting times*
DC-12	Sustainability in ATLAS Computing - session	*Failed wall times, potential adaptations*

Details: Mario's talk from last week

Sustainable ADC

- ADC is delivering great results!
  - Thanks to a lot of very dedicated individuals
- Some observations
  - Overall available personpower is slowly but surely shrinking
    - Temporary placements (6m-1y) to fill the most critical gaps
    - Adds significant overhead to find funding, replacements, work reshuffling and prioritisation
  - Senior technical people are increasingly spending time in coordination
    - Replenishment by new & young people stagnated
    - Important projects and tasks have become dormant
  - Many areas are dependent on single individuals
  - But people also fragmented over many different areas and tasks: Lots of 5% - 10% OTP work classifications
- Can we do better?
  - Gain more time for development and operations
  - Clearer communication flow
  - Set up a more sustainable coordination framework
  - Relax the historical boundaries within ADC
  - Make it easier for site contributions - we're thinking of CRC here
  - Attract new people (and excite the oldtimers!)

New ADC Organigram

ATLAS DISTRIBUTED COMPUTING
OCTOBER 2023

ADC COORDINATION

Mario Lassnig, Alexei Klimentov

PHYSICS

Production Coordination
M. Borodin
Analysis Coordination
A. Forti

Centralised Production
- Monte Carlo Production
- Group Production
- Data Reprocessing
- Physics Validation
- HLT Reprocessing

Physics Analysis
- User Analysis Tools
- Analysis Model Group
- DAST

FABRICS

Computing Run Coordination
Rotating appointment

Infrastructure
- Tier-0
- Grid
- HPC
- Cloud
- BOINC
- Analysis Facilities

Operations
- DA Operations
- DPA Operations
- Central Services
- CRIC
- HammerCloud
- Monitoring
- ADCoS

DATA MANAGEMENT

Coordination
S. McKee, P. Vokac

System
- Rucio

Operations
- System Deployment
- DDM Central Operations
- Monitoring

Research
- Networks
- Caches
- Storage
- Cloud

WORKFLOW MANAGEMENT

Coordination
A. Pacheco, R. Walker

System
- Workflow Definition
- Workload Management
- Workload Execution

Operations
- System Deployment
- Monitoring

Research
- Data Analytics
- Analysis Facilities
- Cloud
- HPC

Organisation and communication

- ADC Coordination Board
  - Consists of the area coordinators and meets weekly
  - The CRC has the option to join if they want to have a bigger say in things!
- Rewritten mandates
  - Previously they were very detailed, almost like job adverts ("do task X", "look after service Y", …)
  - Completely rewritten to focus on the expected outcomes of the activity areas
  - The coordinators are your point of communication, no need to funnel through ADC Coordination
  - Give the coordinators the freedom to act without micromanagement from the top
  - Give the engineers the freedom to act in a more grassroots fashion
- Recognition of Task Forces
  - No more static “boxology”, but a more flexible way to deal with limited duration, focused topics
  - (And it's the way we do things anyway…)
- ADC Weekly
  - Very positive response — people seem to enjoy that it's short and to the point!
  - Possibilities to improve the meeting? Let's discuss!

Notes

1) in our (small) community we have persons basically working only on running the computing resources. I would like to make them feeling they are doing something important for the experiment life. If possible maybe they could make a little step towards the experiment if they find something ‘small' that can be done on top of what they are already doing.

2) on the other hand I’m trying to push since quite some time already to enlarge our community by involving new (younger) persons to work on interesting R&D projects. In this case I think it would be nice to give people a reasonable number of possibile R&D projects or QT to solicit possible future involvements (both on computing and software sides). I know there’s a lot going on here : eterogeneous computing, analysis facilities, cloud computing. It would be interesting to see if connections can be established and strengthened with the ongoing italian government computing program (PNRR/ICSC).

Dario already indicated some enthusiasm for additional information about the Google Project.

Maybe Open Data
~~Maybe Analysis Facilities & Challenges~~
~~Analytics~~
~~Mention Fast Sim (incl. Fast Chain)~~
~~Mention e.g. Lossy compression~~
~~Support for non-Italian sites~~
~~HPC~~