ATLAS S&C toward the HL-LHC challenges

US ATLAS Computing Facilities 2022

1 Dec 2022

Alessandro Di Girolamo (CERN) and Zach Marshall (LBNL)

LHC schedule

- We have just now finished the first year of Run 3
  - Run 3 will end in 2025
- HL-LHC will start in 2029:
  - Run 4 2029-2032, Run 5 2035->2038

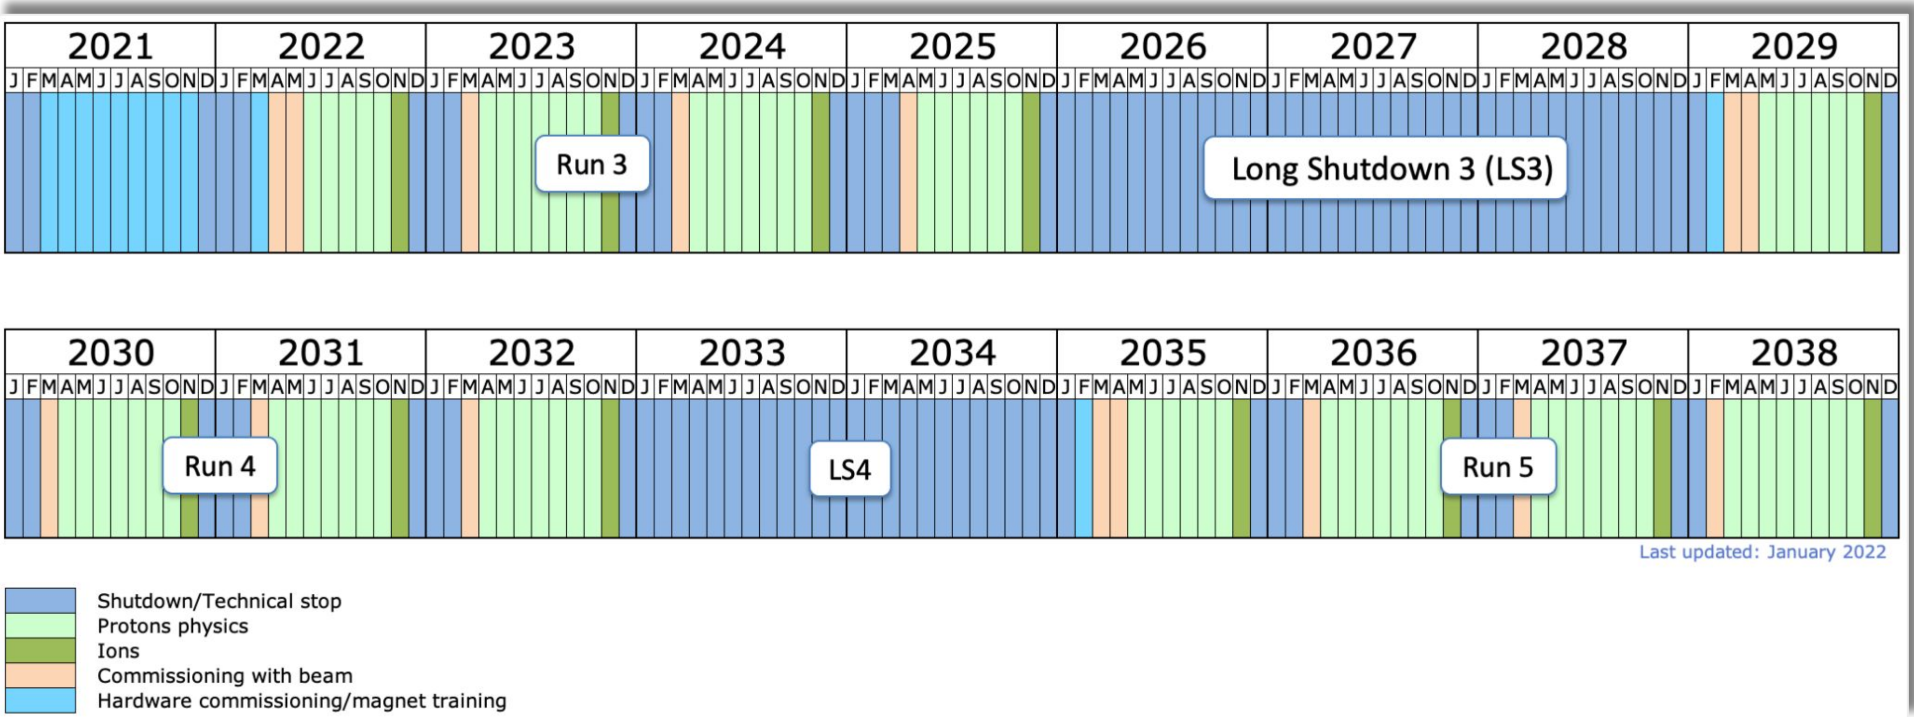

The Physics challenge: Run 3

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

- Trigger rate (physics_Main + delayed) ~ 3.5 kHz
- RAW size ~ 1.2 MB/ev
- Pile up ~ 50
- ~ 30B MC events/year

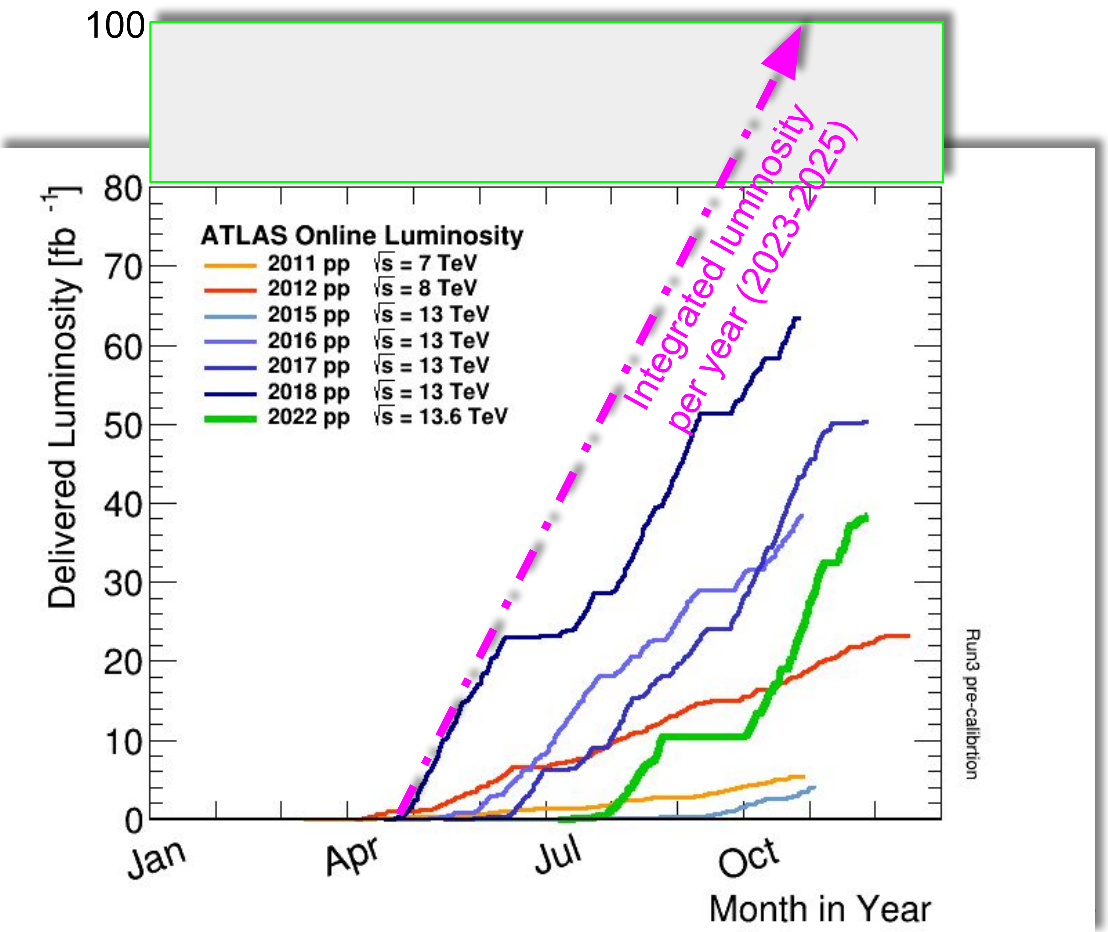

The Physics challenge: HL-LHC (Run 4)

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

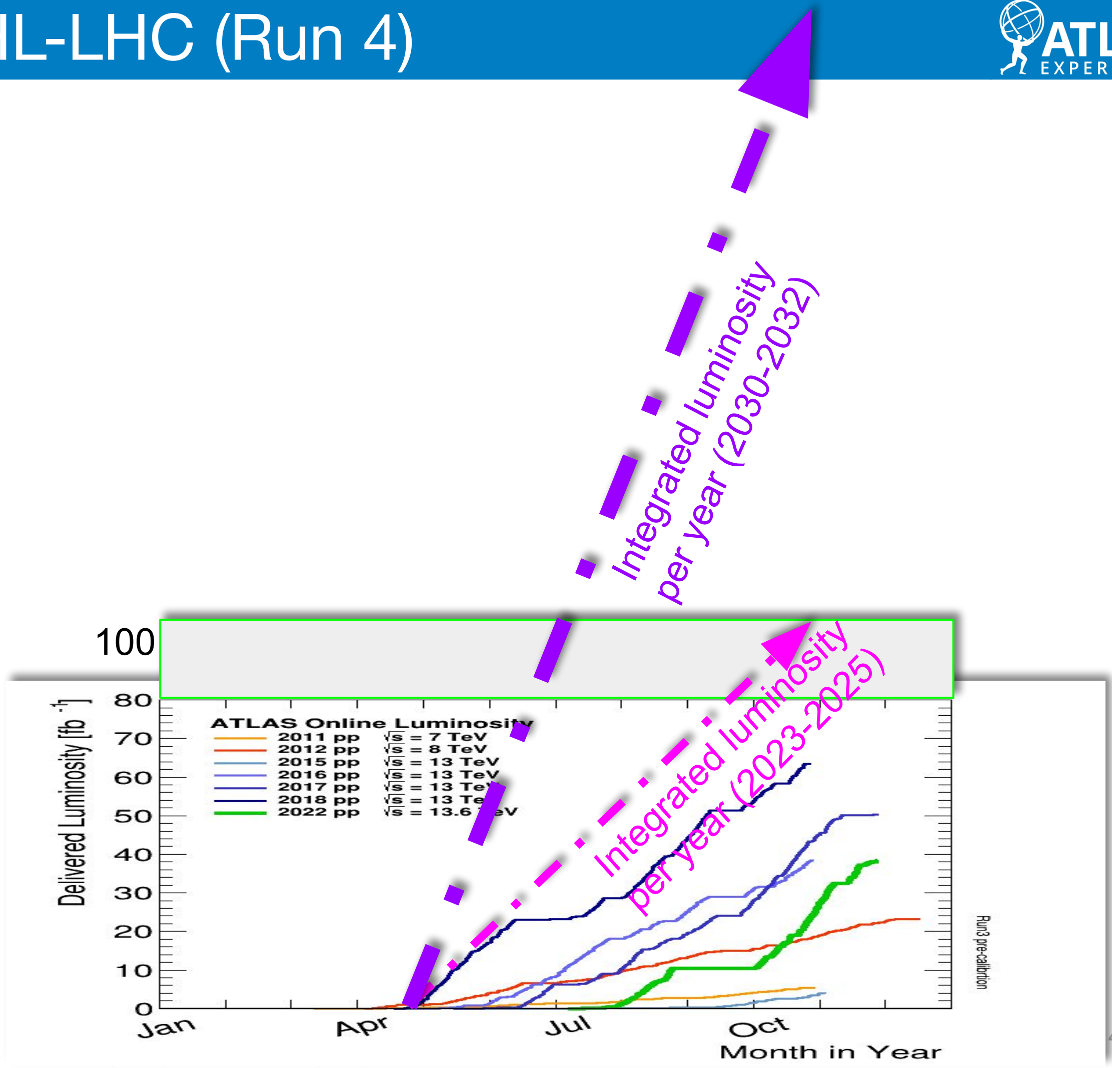

Mixing the ingredients...

… (while adding a few more)

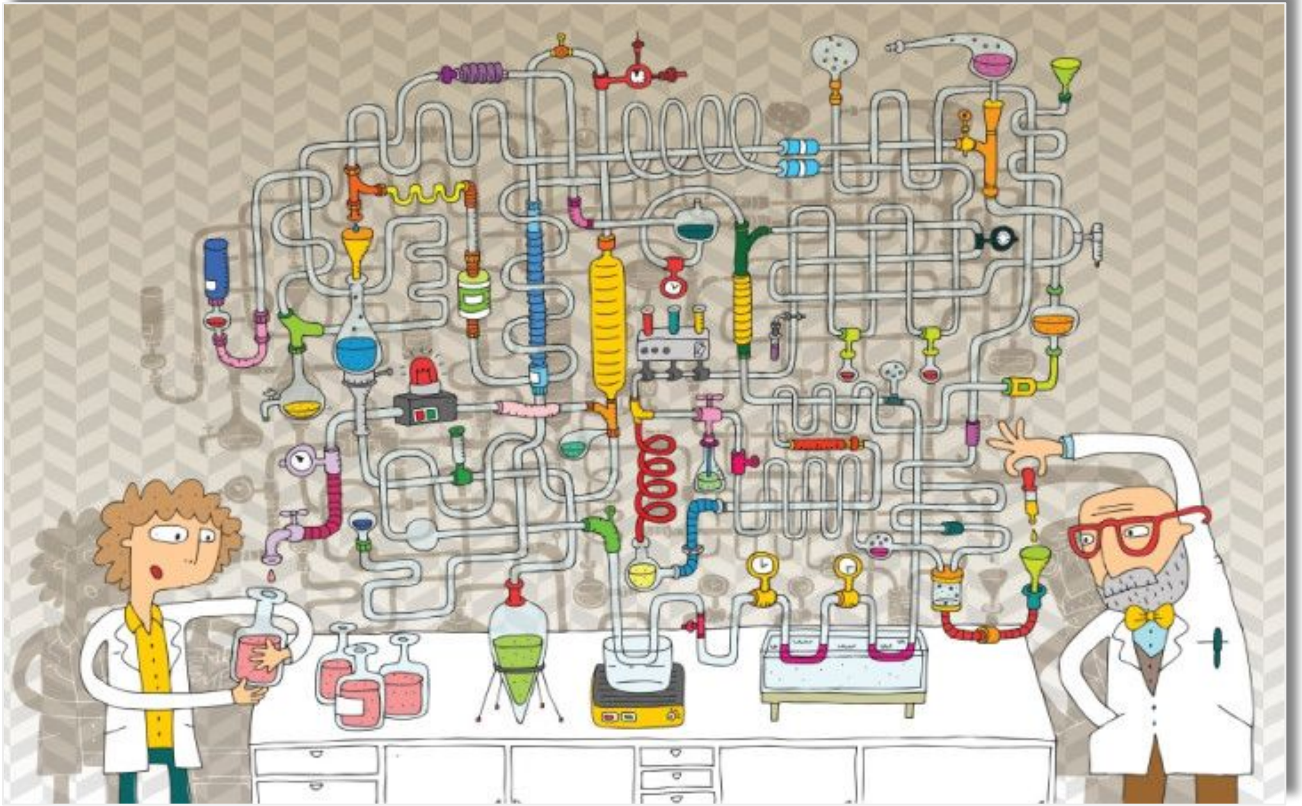

The Computing Challenges

- Computing Resource Estimates show 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
    - Conservative should be achieved with today's effort and people
    - Aggressive requires more people, more R&D projects bringing resource savings
    - N.b. some projects for which we were/are not able to estimate the concrete impact are not (yet) included (e.g. GPU usage)
  - The black lines indicate the "flat budget" of 10% (lower line) and 20% (upper line).

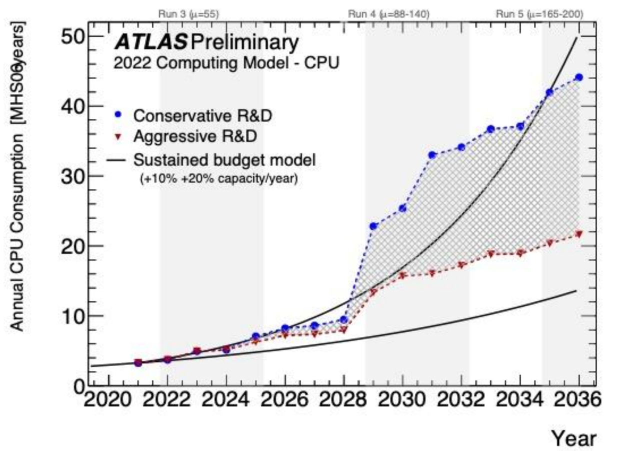

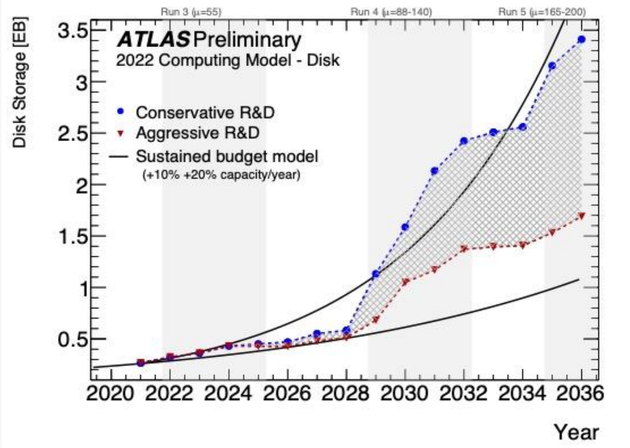

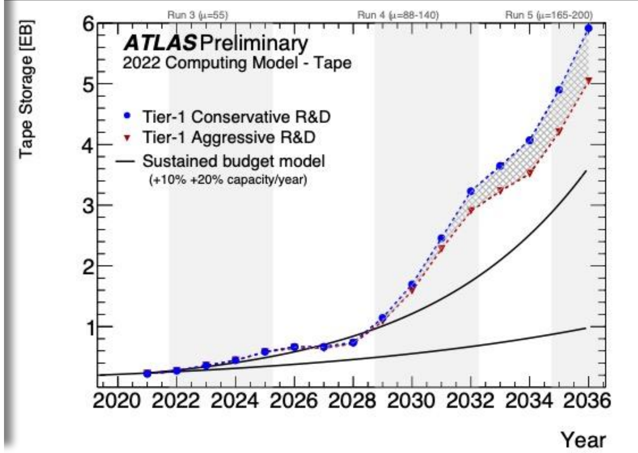

Time to act - now!

Combined performance milestones and activities

		Milestones	Development activities	
2022	LS2	Software for ITK, HGTD and combined tracking running in release 22; Phase II read-out and TDAQ simulation timeline definition	Low level reconstruction (tracking, clustering, object acceptance) High level reconstruction	R&D including ML/ accelerators and other new techniques and ideas
2023				
2024	Run 3	Complete accelerators/ML R&D		
2025				
2026		Feature freeze		Performance tuning (staggered)
2027				
2028	LS3	Final performance freeze	Validation (staggered)	
2029			Contingency/bug-fixing	
2030	Run 4			

ATLAS S&C HL-LHC Roadmap

- Based on the HL-LHC schedule, it is clear that R&D need to be going now
  - from R&D to demonstrators, with measurable impact
  - We see already lot of engagement!
- Integration and validation will require lot of time
  - Late arriving R&D is risky

- In 2025 we expect to have in the TDR a detailed path to HL-LHC data taking
  - (e.g. w or w/o accelerators)

Some caveats: lots of “business as usual”

- We have to keep on running the experiment while we are planning for major upgrades
  - As Karl J liked to say “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”

HL-LHC Organizational Milestones

- By experience, we know we need to prototype, evaluate, integrate, validate and iterate back!
- *PR*oject organization milestones:
  - High level description of expectations from each area and activity in terms of demonstrators.
  - (Re)evaluation of impact
  - Evaluation of effort needed
- Milestones evaluation twice a year
  - First checkpoint done in June S&C week
  - Lightweight reviews: the goal is to keep clear expectations on each area – will learn while going through

Project Organization

MID	DID	Description	Due
PR-1		First evaluation of effort needed to deliver on HL-LHC milestones	Q2 2022
PR-2		HL-LHC Computing TDR	Q3 2024
	2.1	R&D projects targeting Run 4 ("Run 4 projects) define scope and potential impact of their demonstrators, and a program of work with effort and risk estimates to the end of Phase 2.	Q4 2022
	2.2	Define release, datasets and platforms to be used to evaluate Run 4 performance impact of demonstrators	Q1 2024
	2.3	Run 4 projects release their demonstrators	Q2 2024
	2.4	Run 4 projects evaluate the performance impact of their R&D demonstrators and estimate the effort needed to develop fully functional prototypes	Q2 2024
PR-3		Run 4 Release	Q2 2027
	3.1	Run 4 projects release fully functional prototypes, estimate risks and effort needed to bring to production quality	Q2 2026
	3.2	Run 4 developers tutorial	Q3 2026
	3.3	Run 4 Feature Freeze	Q2 2027
PR-4		Ready for Run 4 Data Taking	Q2 2029
	4.1	Run 4 projects demonstrate required functionality in release	Q3 2026
	4.2	Run 4 release validated	Q1 2029

HL-LHC - a glance into the Resource Estimates for 2031

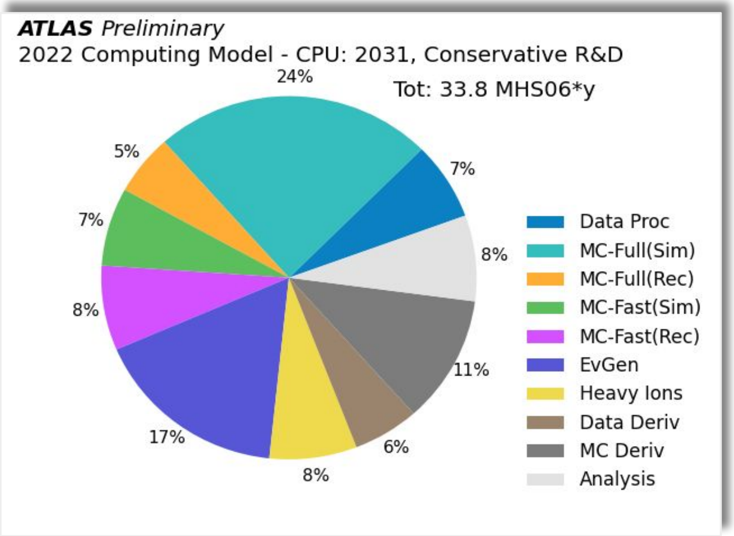

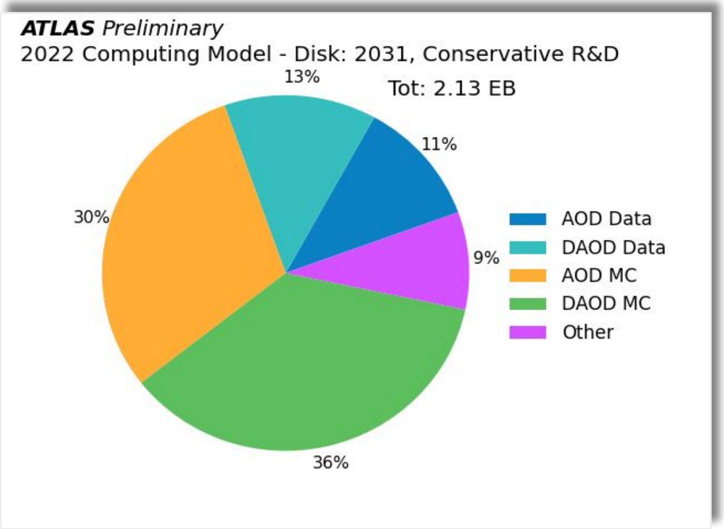

- There is no “silver bullet”:
  - it’s not that we do have 1 (one) single problem and if we solve that “we are done”.
- we need to work on each and every part of our system
- The opportunities are “everywhere”: we need to work hard to scavenge a few % here and there
  - An evergreen motivational speech, [Any Given Sunday](#) : true we do not play *against* anyone, but we still want to bring out the best from ourselves - while having fun!

Tackling the challenges

- There is nothing bad in being “two” in tackling difficult challenges!
  - Important we work together and we organize ourselves
  - We don’t compete
- There are no shortcuts and not so many “low hanging fruits”
  - It requires disciplines from all of us
  - Others might suffer from our “being late”

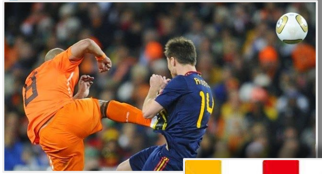

- Expectations management is paramount
  - We do not dictate! We do bring people onboard, we (try to) motivate them (and people motivate us!)
  - Crucial to clearly define our commitments taking into account all the “other” stuff we have to do
  - <u>We do not need to know what people do, but it is critical that we all estimate *by when* “that something” will be done.</u>

Tackling the challenges: infrastructure evolution

- HPC and Clouds (and Industrial partners)
  - More and more often we are seeing them discussed together
  - Not really much overlap, except that they “both” are not standard Grid sites.
- They really deserve two different “set of slides” → next

High performance computing refers to computing systems with extremely high computational power that are able to solve hugely complex and demanding problems.

European Commission - Exscalate4Cov project : supercomputing to identify new therapies for COVID-19

In the digital decade, high performance computing (HPC) is at the core of major advances and innovation, and a strategic resource for Europe's future.

Cloud computing[1] is the on-demand availability of computer system resources, especially data storage (cloud storage) and computing power, without direct active management by the user.[2] Large clouds often have functions distributed over multiple locations, each of which is a data center. Cloud computing relies on sharing of resources to achieve coherence and typically uses a "pay as you go" model, which can help in reducing capital expenses but may also lead to unexpected operating expenses for users.[3]

Contents [hide]

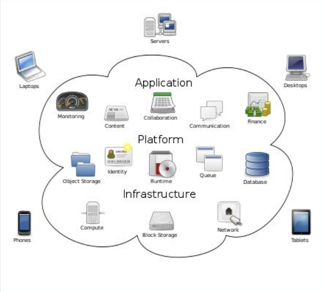

Cloud computing metaphor: the group of networked elements providing services need not be individually addressed or managed by users; instead, the entire provider-managed suite of hardware and software can be thought of as an amorphous cloud.

HPC and ATLAS - last 12 months - HPC only

- A few HPCs make a big difference
  - But still it's important to highlight that many HPCs are integrated!
  - Looking forward to Leonardo and Perlmutter (what a name they chose here ;-))

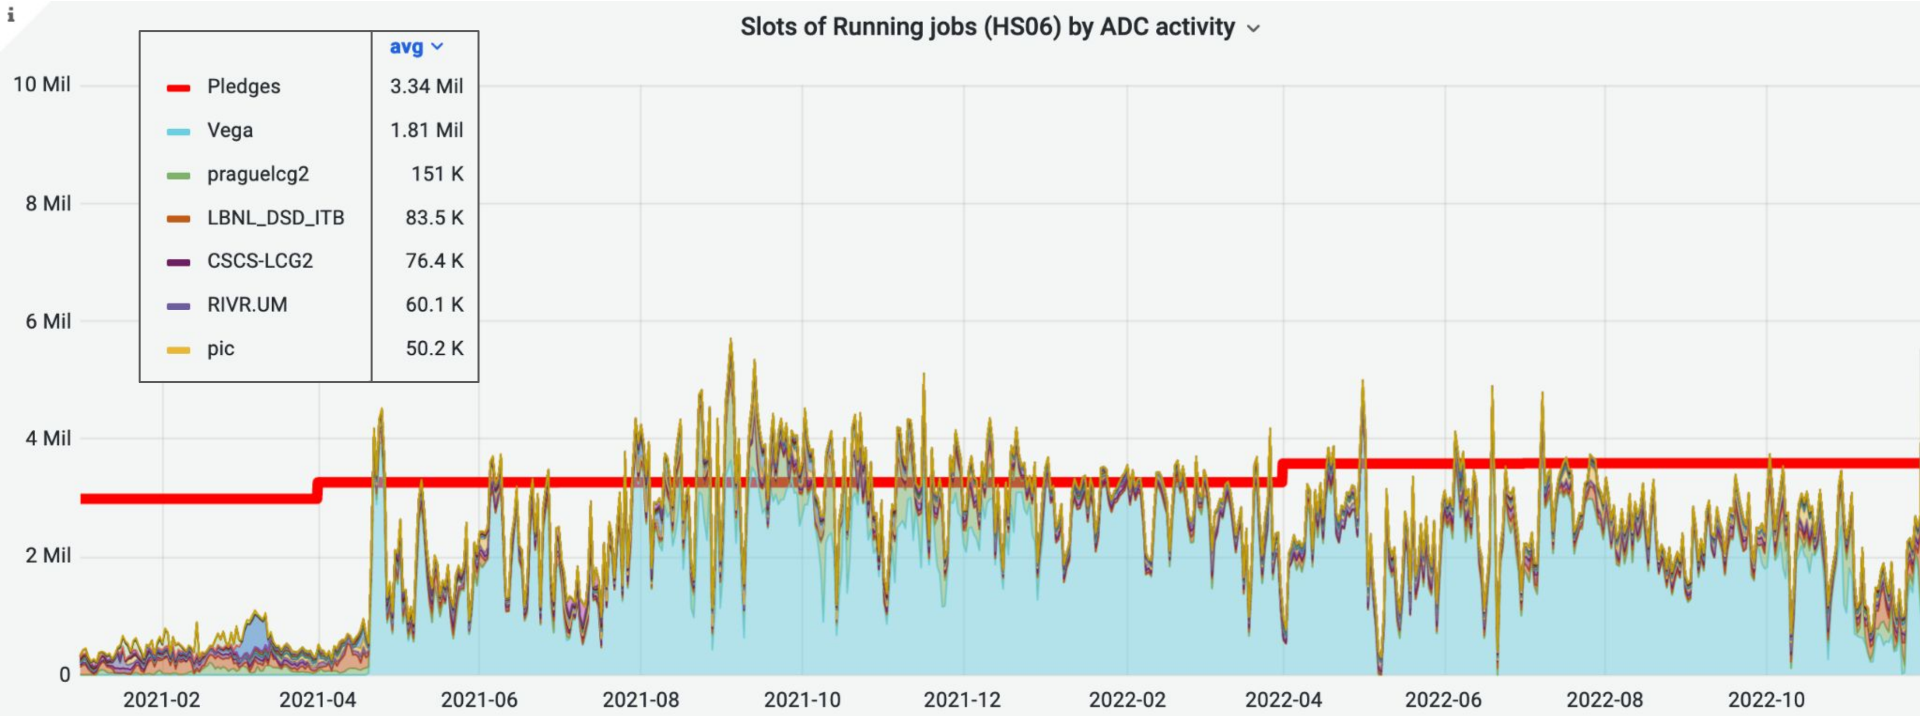

HPC and ATLAS - considerations

- Most resources provided by:
  - Vega (large majority) and Karolina, both EuroHPC systems
  - Both run ~all workloads (sometimes also analysis)
- Why were these so successful?
  - Basically x86 machines (no major code adaptation; similar to a std grid site)
  - Insiders helping with edge services (and even machine setup)
  - Early adoption helping give large allocations (in before others)
- Biggest problem (challenge?) with HPCs in general:
  - No associated disk allocations!
- Second biggest problem:
  - Single-year allocations

The ATLAS strategy towards opportunistic HPC

- Investment in proportion to opportunity
  - At the moment, e.g., we see ARM as a higher priority than Power
  - Not trying to get in on every machine!!
    - We don't need to – getting the easy/with internal good support ones gives ample CPU.
  - All the "easy" HPCs transparently integrated, running all workflows, are much welcome!
- We have a good (and growing) toolkit for edge services
  - With several available solutions, one is likely to fit a new machine
- Work towards several big HPCs is ongoing
  - Perlmutter in the US has just passed PhysiscsValidation; hope to capture some of Fugaku(2)/CSCS with ARM development; Getting close to using Toubkal (UM6P, largest HPC in Africa).
- The details can make the difference:
  - HPC are not easy, need expertise on I/O, shared file systems, scheduling, AAI…
- Without forgetting:
  - The risk that the opportunistic CPU resources will disappear at any time is very real
  - The main physics program of the experiment should be supported through pledged resources
  - A constant yearly increment for both CPU and Disk is reasonable to minimize the risks that, if these resources disappear, the computing centers providing pledged resources won't be able to provide what is needed

R&D Projects with Industry

- ATLAS is engaging with a number of industrial partners on R&D projects, e.g.:
  - SEAL, a decentralized cloud storage start-up
  - Google Compute Engine and Amazon Web Services

- These are valuable opportunities to evaluate technologies and to define a resilient long term strategy
  - Use of cloud compute for *all* workflows, and a comparison to traditional resources of the total cost of ownership (TCO)
    - Google already used for interactive analysis
  - Evaluation of test technologies (e.g. TPU, ARM) using cloud resources without requiring large-scale purchases
    - AWS extensively used in setup and validation of ARM Athena nightlies builds
  - Evaluation of remote cold storage as a complement / supplement to tape systems

- These projects also present great connections with strong partners
  - Amazon and Google are major players that many groups already collaborate with
  - Being able to speak the same language and learn how they operate is invaluable

ATLAS and Cloud Resources - TCO exercise

ATLAS EXPERIMENT

- ATLAS Google Project (AGP)
  - Established a “site” on Google (cloud) resources with a total cost of ownership (TCO) exercise. Full report by end of 2023
- Cloud TCO Phase 1 (“What functionality is needed?”) concluding soon.
- Understanding complexity – many moving parts to track
  - Hardware, operation expenses, long-term expenses, risks
- Engaging a broad pool of “consultants” to ensure all issues & metrics are understood
- Interesting discussions about cost reporting: no one pays “sticker price”
  - A lot depends on the Service Agreements (3-5y) between Google and the stakeholder; not easy to scale the costs.

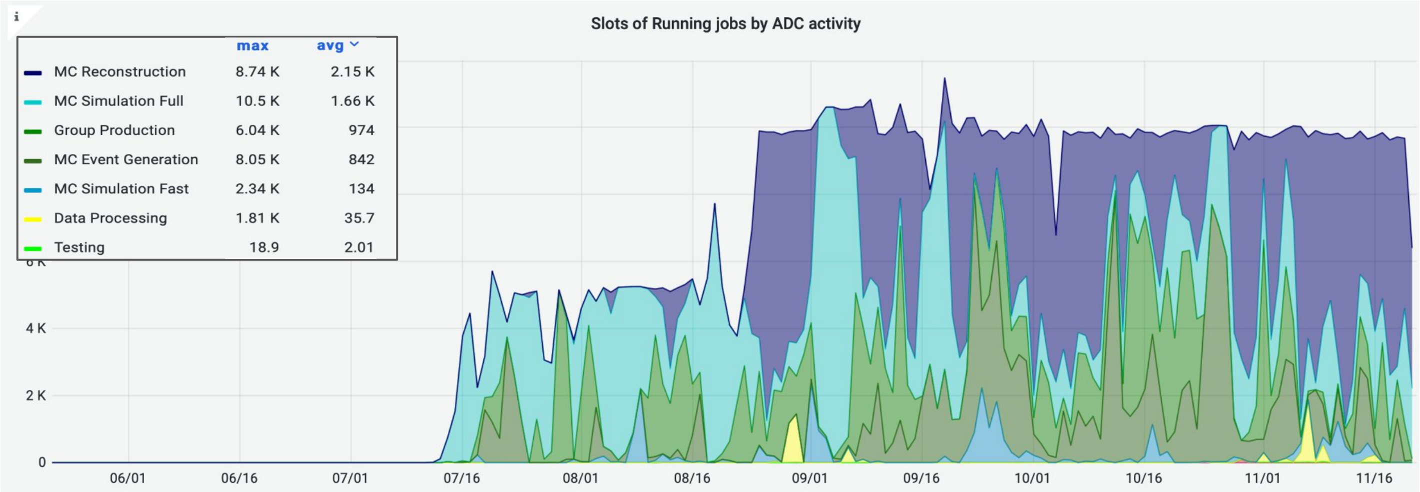

Storage challenges

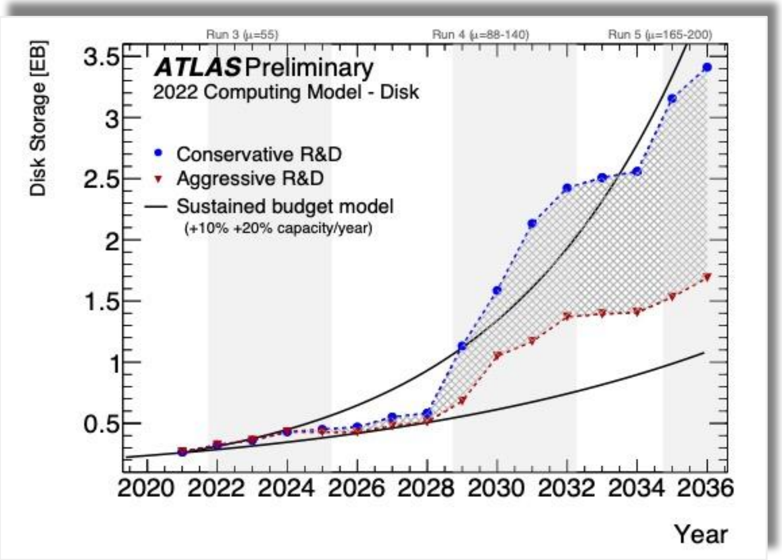

- Our data sample is going to be 10x bigger!
  - today trigger rate is 3.5kHz, will be 10kHz in HL-LHC,
  - RAW size from 1.2MB/event to ~4MB in HL-LHC
  - MC simulated events from 30B/year to 150B/year
- We are working on several R&D (and lots of "business as usual") to make sure we make the best use of our Disk and Tape:

- New analysis model: from a plethora of DAOD_[TOP,HIGGS..] to DAOD_PHYSLITE → savings up to 50% of DAOD disk space
- Data Carousel: trading Tape for Disk
- Constant review of the content and the objects stored in each datatype

Storage challenges: more

- Lot of analytics:
  - (Too much, never enough - and always complicated to keep all the chain working! - help welcome!)
  - understand what we use, how often, how many replicas we need…

Some more far R&D are fascinating, e.g.:

- fastchain: trading CPU for Disk
  - A fast end-2-end simulation, we could choose to save less (or zero) intermediate data - potentially important impact on disk and tape savings
- More aggressive deletion - and recreation on demand
  - Again trading CPU for Disk, one of the complexity is on the timely reproduction (with proper metadata)

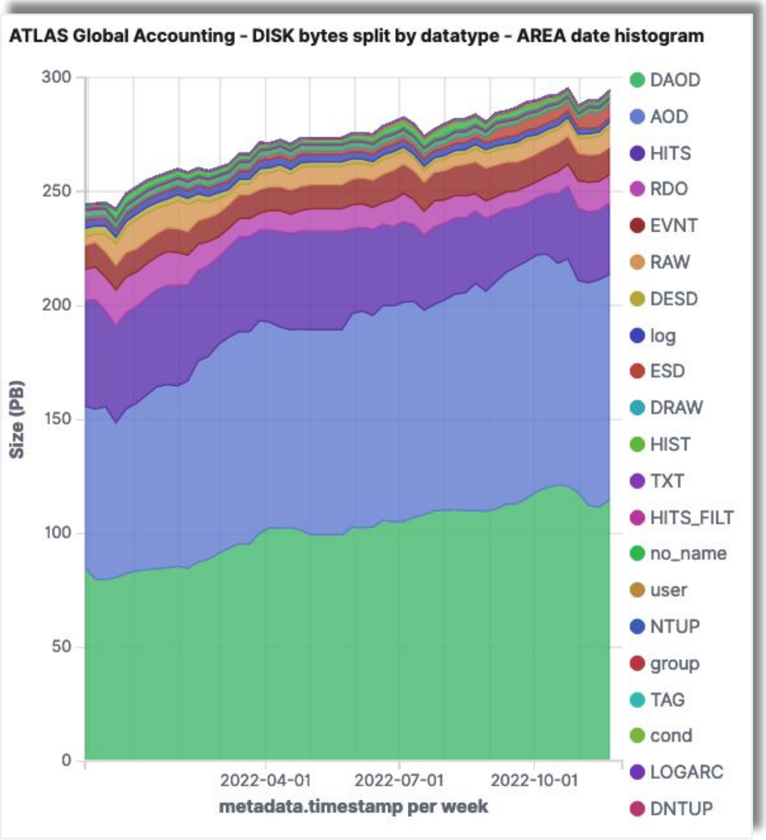

Tackling challenges: considerations

ATLAS EXPERIMENT

- HPCs are very useful to mitigate the risk of being short in CPU resource
  - With some caveats - as we discussed
- Clouds and in general project with industry help us in improving our frameworks, make it more sustainable and “community standard”
  - Some interesting “opportunistic storage” R&D
  - And TCO evaluation is important
- Many more R&D to mitigate the HL-LHC not-enough-CPU risks:
  - We are following with great interest the Geant4 collaboration (+Adept and Celeritas) and Event Generator groups’ R&D efforts
- Many R&D to mitigate the storage challenges’ risks:
  - Our systems are complex: need lots of brains and analytics expertise to take the right informed decisions
- Note that not all our R&D will reduce our resource consumption!
  - We want to do *great* physics, some of the R&D will help us on this too
  - And some others, e.g. ACTS, will push us to rewrite our (reco and tracking) code, to be able to be accelerators friendly
    - → independently on the usage or not of accelerators, we will for sure gain having better written code: sustainability and capability to engage more people

Accelerators: why, where, when?

- In terms of Total Cost of Ownership, accelerators often don't pay
  - Modern GPUs are extremely expensive
  - Good FPGA/GPU-based software engineering is very expensive and requires long-term expert support (v. CPU-based code)
- Accelerators might "pay" when you have a *fixed space* or *power budget* and still you need to expand your total compute
  - This is the case for e.g. our trigger farms…
- The big decision point for accelerators will be in ~2025/6
  - Several demonstrators, and prototypes, planned to prepare for the decision point
  - Likely both online and offline will decide within a year of each other

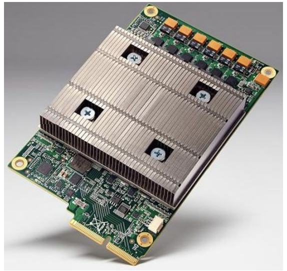

Analysis (& Facilities)

- It's not just a matter of tools!
  - Actually, that's the "easy" part!
- Understanding what users really want, what are their pain-points, how crazy some of their(our!) requests are - this is the real challenge!
  - Requirement engineering is critical - "what", and not "how"!
  - We have to be careful in not building up a solution for just a handful of users.

- Today defining Analysis Facility is at the same time simple and hard:
  - Some of our facilities are already analysis facilities, e.g. lxplus at CERN, and some in US, Italy, Germany, etc.
  - For us - for ATLAS: federated access, offering interactive and batch resources, with why not state of the art technology available (e.g. notebooks, GPUs, etc)
- But what are the critical points? What can/should we do better? User friendliness? Usability? Fast turnaround?
  - … it's going to be fun!

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!
- Time to act is *now*
    - By 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 to common shared objectives is paramount
    - The *way* in which we work can make the difference between success and failure!
    - Fragmenting efforts is going to be lethal - not effective
    - An interesting read: Socio economic impact of big science center
- We have provided here a “quick” overview of some of the challenges
    - Much more in our HL-LHC Roadmap
    - … and we will be happy to discuss furthermore with any of you interested in contributing!

Plan B

Success is no accident.
It is hard work, perseverance,
learning, studying, *sacrifice*,
and MOST of all,
love of what you are doing.
-Pele

Extras

The ATLAS detector now

Phase-I upgrade

New LAr L1 trigger

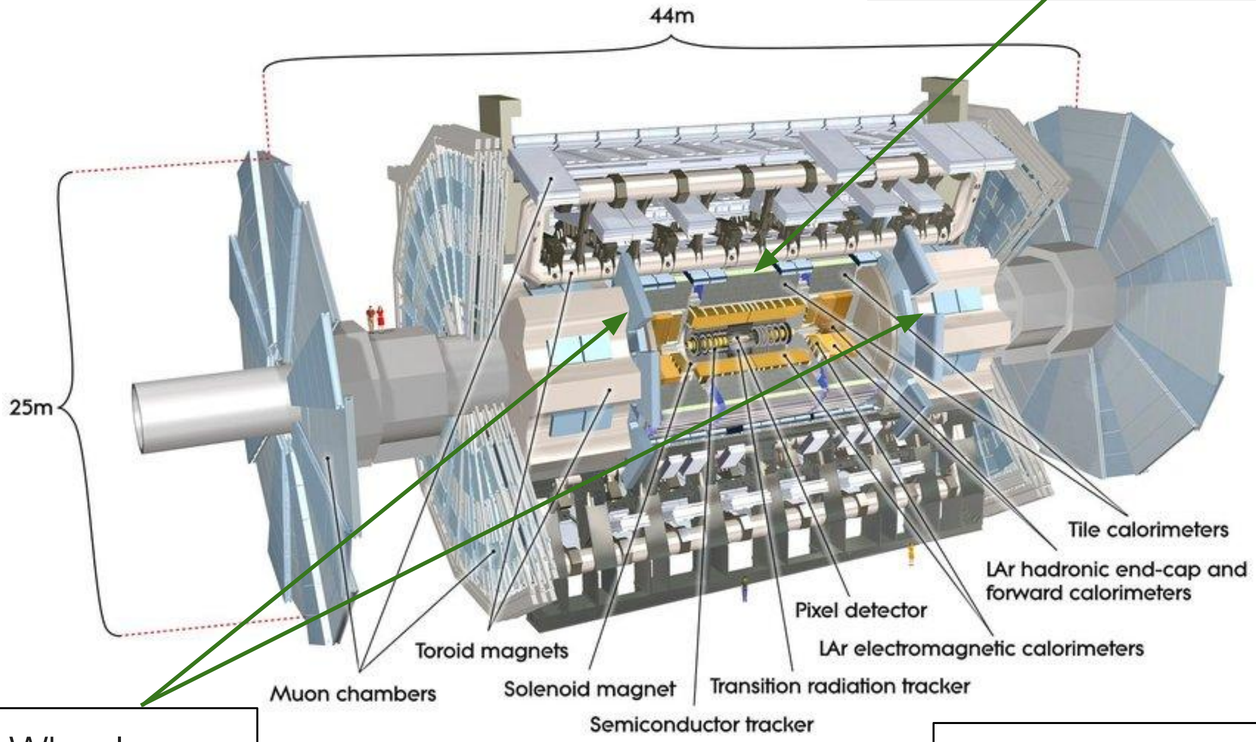

New Small Wheels

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)

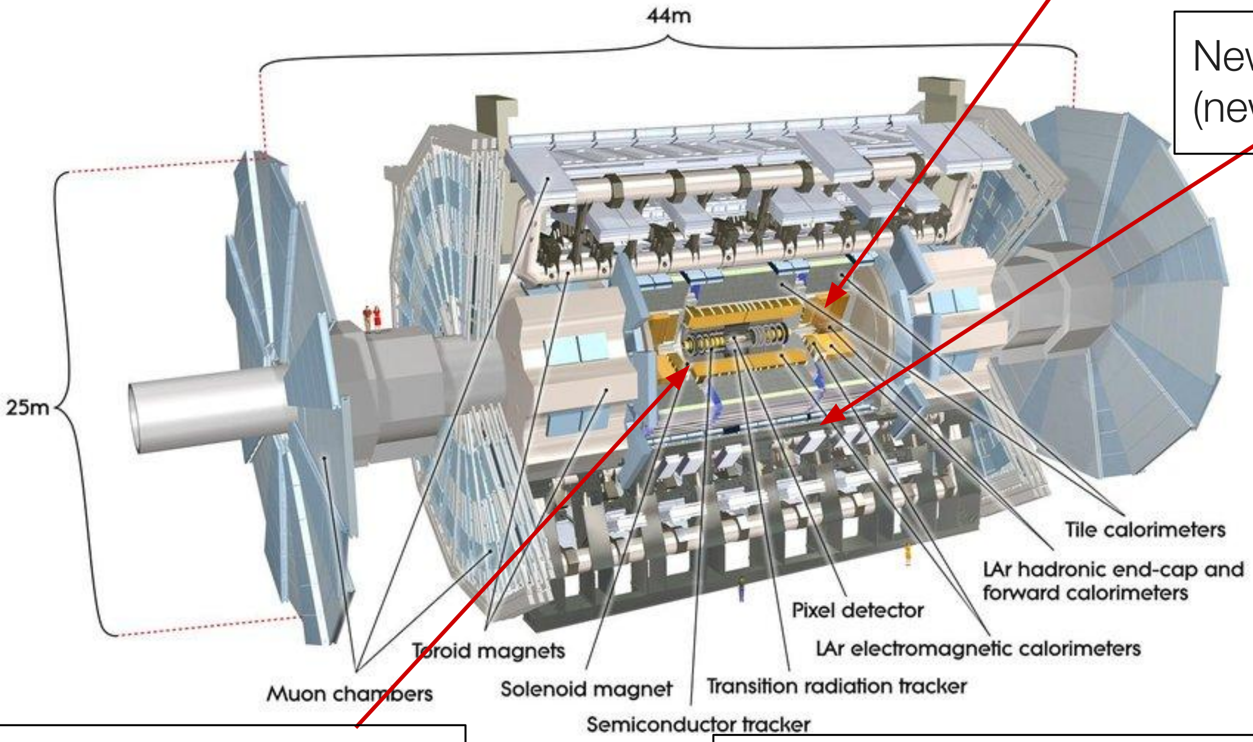

New Inner Tracking detector (ITk)

Upgraded Trigger and Data Acquisition system

The ATLAS Collaboration

HPC and ATLAS - last 12 months

- We have been extremely successful in integrating and exploiting the CPU from opportunistic HPC resources: doubling our pledge!
  - And some Grid Sites deliver (part of the) pledge transparently through HPC

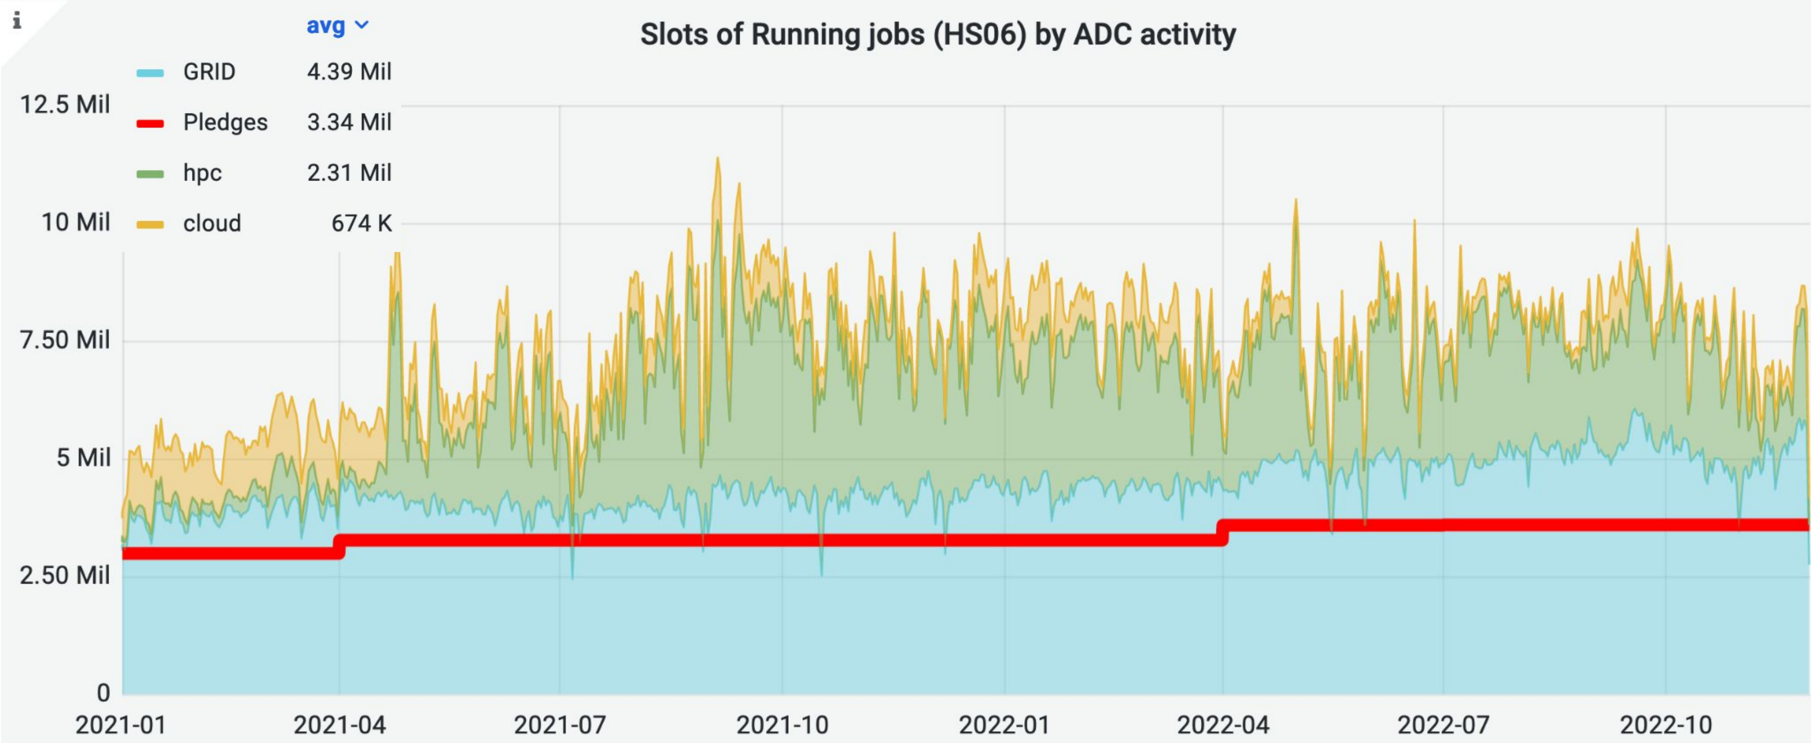

Power and Energy

- Interesting & useful discussion of energy consumption during WLCG Workshop
  - Discussions dominated up to now by ATLAS members + sites
- Happy to engage with WLCG 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)

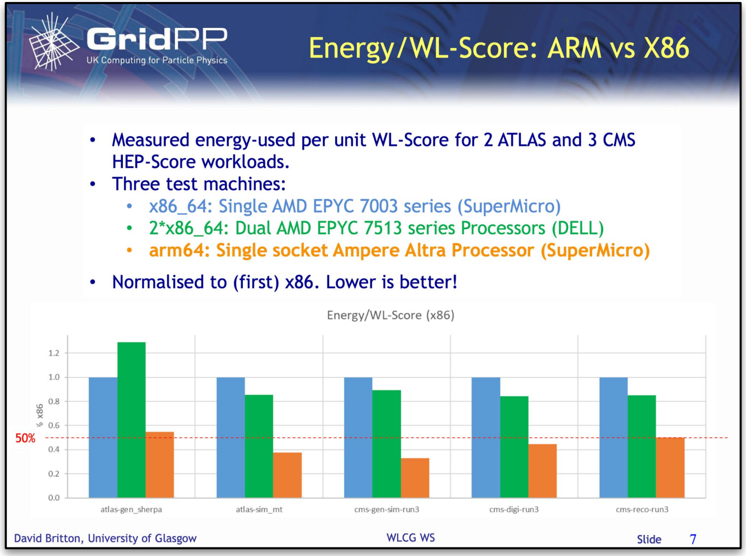

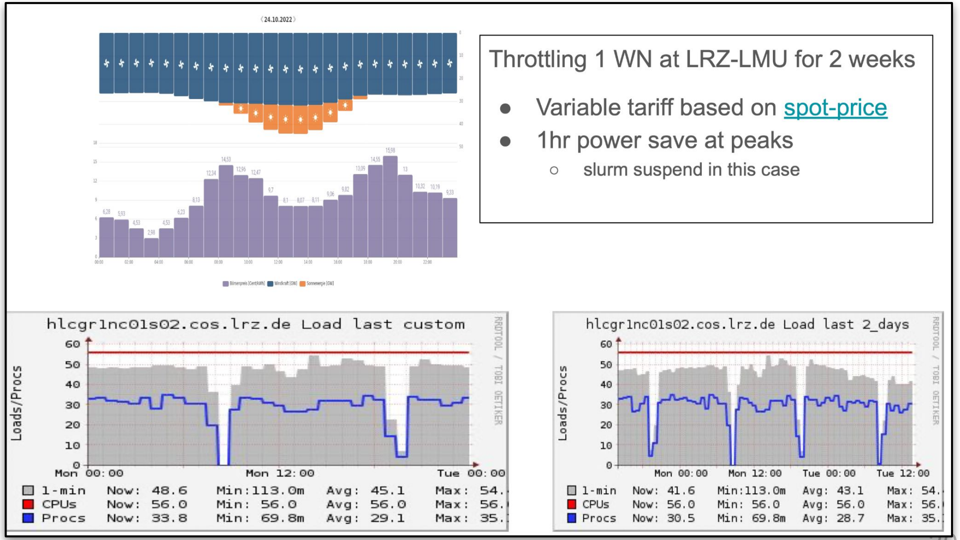

C-RSG Request for 2024

ATLAS EXPERIMENT

- Discussed our preliminary request for 2024 computing resources with C-RSG.
  - Some concerns about the tail of release 21 analysis (we're concerned as well)
  - Questions about LHC machine uncertainties and power crisis
- In the preliminary report the indications are that they will not object to our preliminary request
  - Spring is when this gets finalized in any case

		2023 agreed @ April 2022 RRB	2024 Request @ September 2022 RRB	Balance 2024 wrt 2023
CPU	T0 (kHS06)	740	850	14.9%
CPU	T1 (kHS06)	1430	1501	5.0%
CPU	T2 (kHS06)	1747	1834	5.0%
CPU	SUM (kHS06)	3917	4185	6.8%
Disk	T0 (PB)	40	46	15.0%
Disk	T1 (PB)	136	162	18.9%
Disk	T2 (PB)	168	198	17.6%
Disk	SUM (PB)	344	405	17.8%
Tape	T0 (PB)	174	205	17.9%
Tape	T1 (PB)	353	448	26.8%
SUM (PB)	527	518.0	653	23.9%

Table 6: Summary of the preliminary requests for computing resources in 2024.

- Essentially no change to Tier 0 resources w.r.t. 2023 (very similar run parameters)
- Modest increase to CERN resources, particularly for lxbatch (user feedback) and central services
- Very small increase in CPU (standard HW replacement usually provides more than 5%)
- Overall: sum of CPU, disk and tape requests should be consistent with "flat budget" scenario

CERN Hardware cost

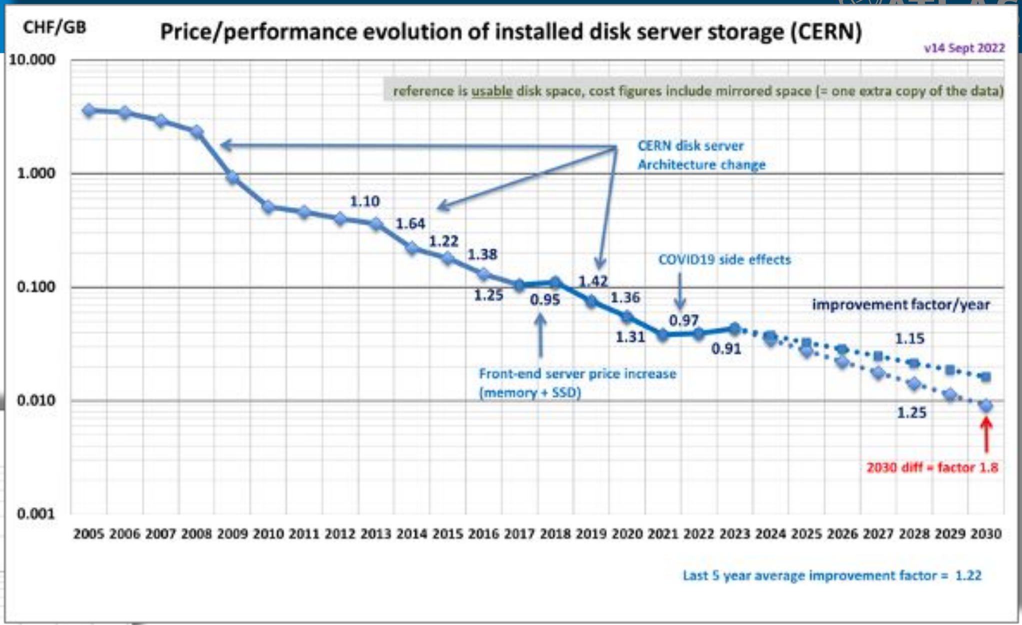

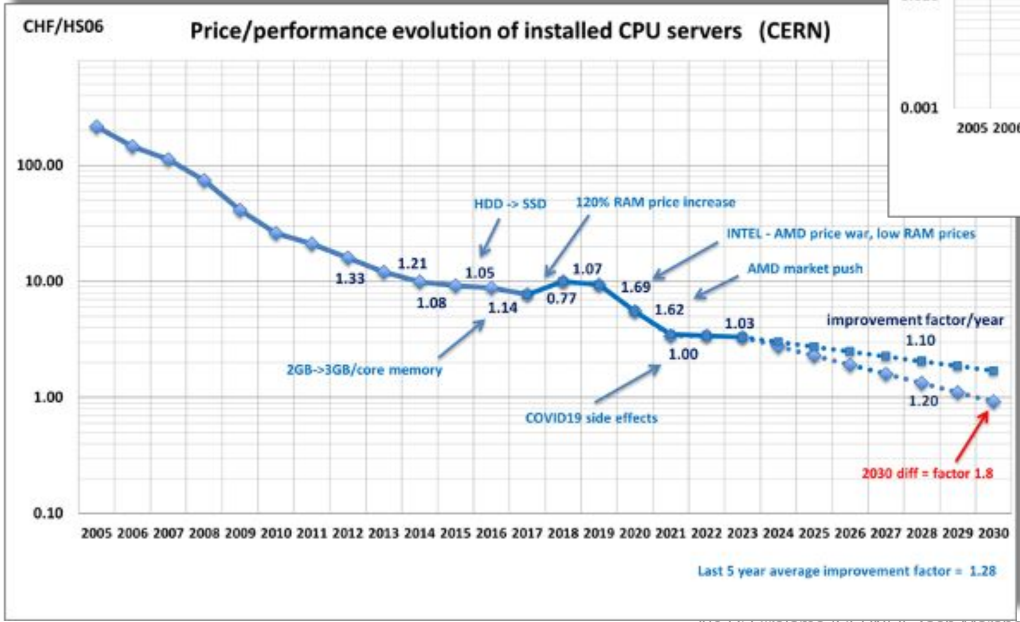

The Software and Computing HL-LHC Roadmap

Planning for the HL-LHC era

- First major step was the Conceptual Design Report in 2020
  - Broad plans for improvements in several key areas
- Second major step is the HL-LHC Roadmap
  - More detailed milestones and deliverables
- Plan to produce a TDR in ~2024

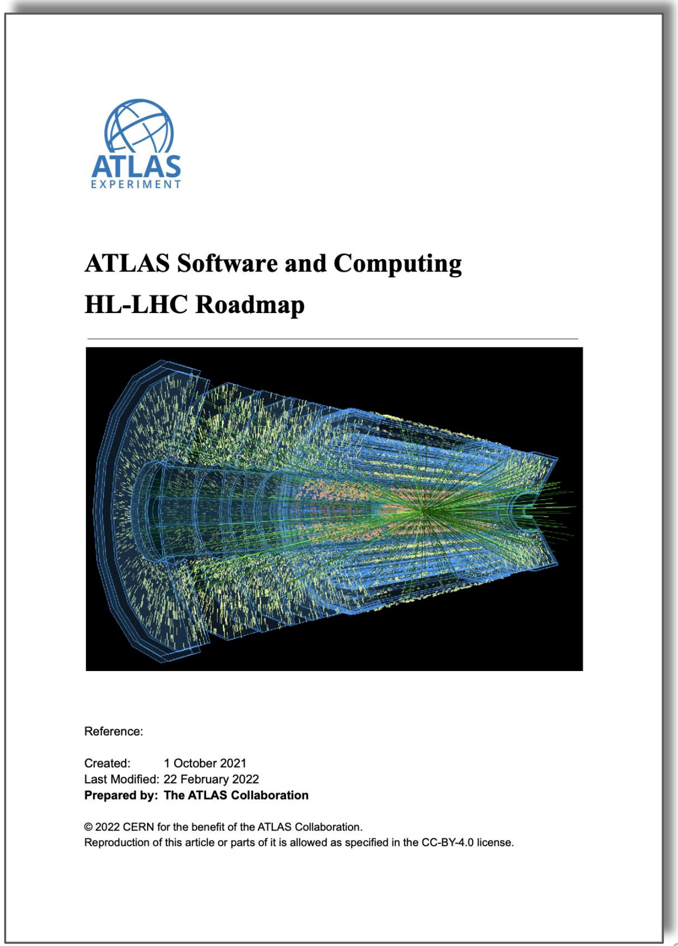

ATLAS EXPERIMENT

ATLAS Software and Computing HL-LHC Roadmap

Reference:

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

© 2022 CERN for the benefit of the ATLAS Collaboration.
Reproduction of this article or parts of it is allowed as specified in the CC-BY-4.0 license.

HL-LHC - a glance into the Resource Estimates for 2031

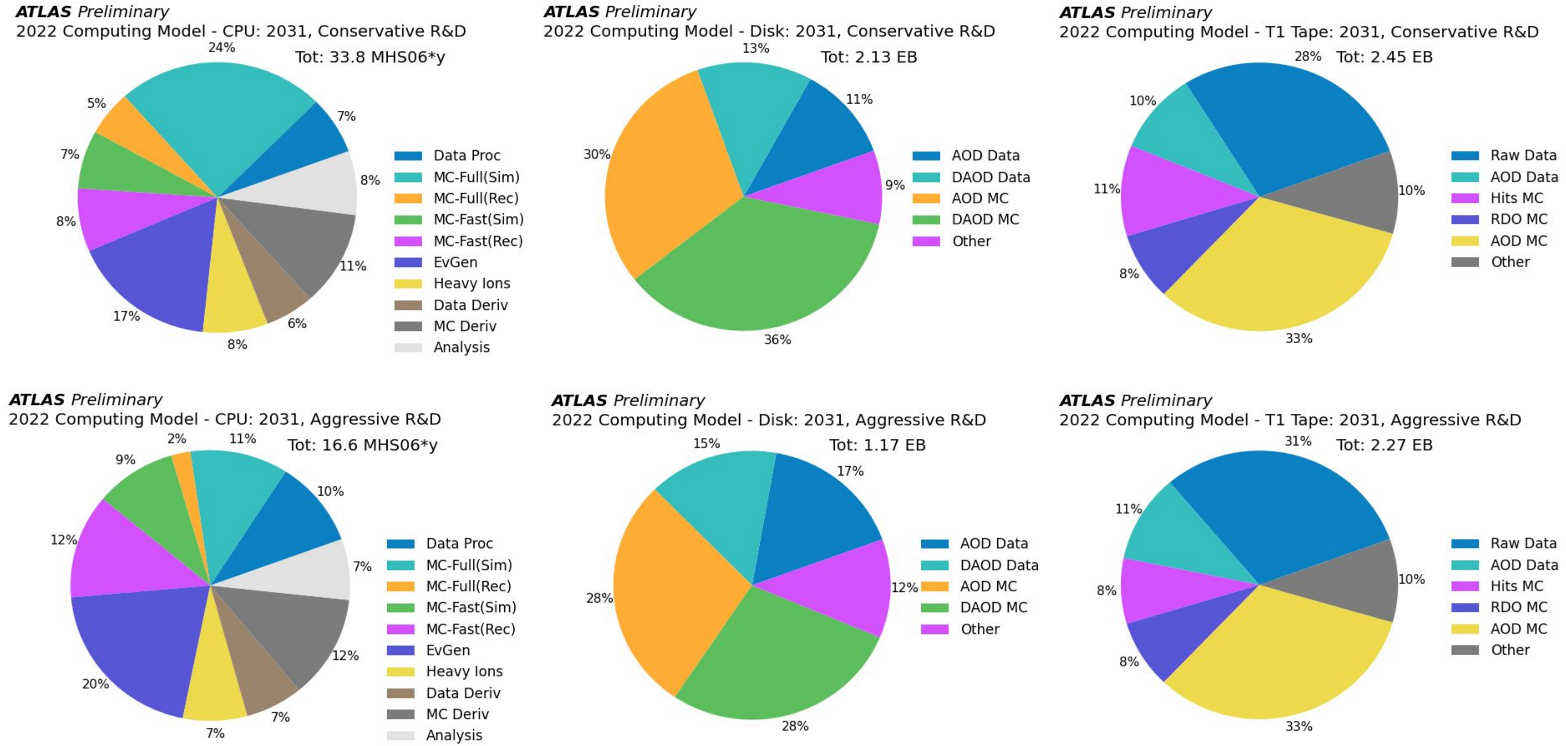

- Observations:
  - We do not have areas which are too much resource hungry
  - There is no “silver bullet”, we need to work on each and every part of our system!

HL-LHC Milestones

ATLAS EXPERIMENT

- In the document we have several different kinds of milestones (examples from simulation):
  - Combinations of deliverables to a main goal with an expectation of delivery date (e.g. “Track Overlay in FastChain”, #16)
  - Decision points (e.g. “Decide if we will run Geant4 on a GPU for Run 4 start-up”, #17)
  - Deadlines by which work must be delivered, or where outside work must be integrated (e.g. “Geometry updates for the Run 4 MC Campaign”, #21; “Validation of full and fast simulation for Run 4”, #24)

- Milestones are here to *help us understand* what work is in front of us for Run 4 and *communicate expectations*

Simulation

	MID	DID	Description	Due
	15		Updates required for MC+MC Overlay for Run4	Q4 2022
		15.1	Enhance Overlay code to better deal with high pile-up environments	Q4 2022
		15.2	Implement Overlay Algorithm for HGTD	Q4 2022
		15.3	Implement Overlay Algorithm for ITk Pixels	Q4 2022
		15.4	Implement Overlay Algorithm for ITk Strips	Q4 2022
	16		Track Overlay in FastChain	Q4 2022
		16.1	Data overlay with pre-reconstructed tracks	Q4 2022
	17		Review or Run 4 Simulation Strategy	Q2 2023
		17.1	Evaluate G4 equivalents of ISF functionality	Q1 2023
		17.2	Evaluate accuracy of EMEC simulation in FastCaloSim	Q1 2023
		17.3	Prototype EMEC geometry using native G4-GPU constructs	Q1 2023
		17.4	Prototype GPU version of LAr sensitive detector callback	Q1 2023
	18		Pile-up Digitization in AthenaMT	Q4 2023
		18.1	Full AthenaMT compatible pile-up digitization	Q4 2022
		18.2	Reduce memory usage of pileup digitization for very high mu	Q4 2023
	19		Make data overlay useable for p-p collision simulation	Q4 2024
		19.1	Flexible alignment and geometry in simulation for data overlay	Q4 2022
		19.2	Skimming of zero-bias Bytestream data and efficient grouping of events as input to p-p overlay	Q4 2023
		19.3	Adapt Fast Chain workflow to data overlay	Q4 2024
	20		Fast simulation/FastChain development for Run4	Q3 2026
		20.1	R&D in FastChain alternatives	Q3 2026
		20.2	Development of substantial improvements to fast calorimeter simulation	Q3 2026
		20.3	Development of substantial improvements to fast inner detector simulation	Q3 2026
		20.4	Implementation of ITk in Fatras and Fatras tuning for Run4	Q3 2026
		20.5	Tuning fast simulation to data for Run4	Q3 2026
		20.6	Improve Geant 4 interface for fast simulation	
	21		Geometry updates for the Run4 MC Campaign	Q3 2026
	22		Run4 optimisation of the performance of ATLAS full Geant4 simulation	Q3 2026
	23		Testing Geant4 Versions and configurations for Run4	Q4 2026
	24		Validation of full and fast simulation for Run4	Q1 2027

Addressing the computing challenges of the HL-LHC

COMPUTATION
Make full use of all available resources, including new hardware architectures and facilities
Make physics choices that save computation without sacrificing experimental reach
Produce fewer simulated events with very high fidelity and more events with reduced fidelity
Take advantage of 10 years of operating experience and redesign workflows to be more efficient without sacrificing quality
Adapt the analysis model to maximise efficiency, especially by re-thinking user analysis

STORAGE
Adapt computing model to maximise flexibility and to be able to operate with a wide range of storage types
Reduce the number of analysis formats and their size
Make more use of cheaper storage technologies for data products that are not frequently accessed
Write smaller/fewer events

The more aggressive the development,
the more efficient the use of resources
→ aggressive development plans = investment in people

James Catmore

Adapt to latest trends, more heterogeneity
Use what we have more efficiently

Reduce the disk footprint
Use what we have more efficiently

S&C and ATLAS

the ATLAS community (... and not just the ATLAS S&C community!)

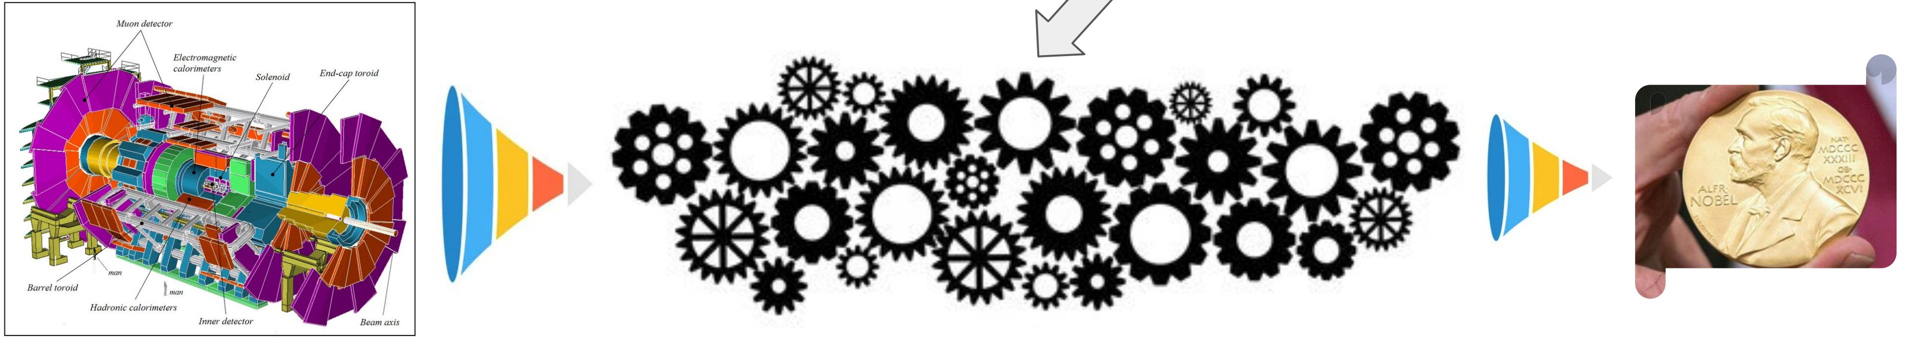

- The Software and Computing community rely on the engagement of each sub detector and Combined Performance/Physics Group, especially in areas like Simulation, Digitization and Reconstruction:
  - We cannot define milestones and deliverables for them, but we need to agree with the collaboration in such a way that the various pieces can come together to solve ATLAS’s HL-LHC challenge

HL-LHC Computing Challenge - LHC/HEP community

- Challenge shared with the entire LHC/HEP community
  - ATLAS involved with HSF (HEP Software Foundation forum) from its inception (2015)
  - HSF community produced roadmap identifying common software and shared R&D priorities
- Our HL-LHC C&S planning relies on successful R&D in "external" software
  - e.g. Geant4, Generators, Root, DOMA.
- We will succeed (or not…) as a community

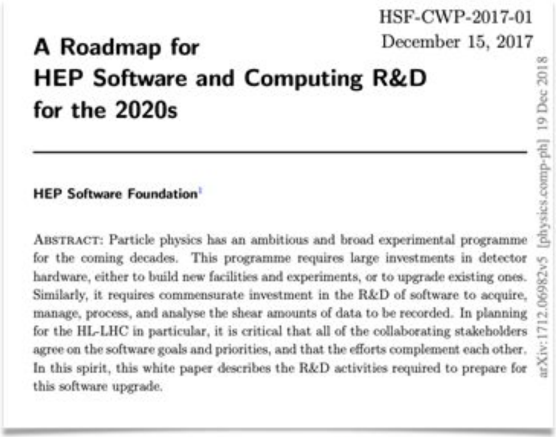

HSF-CWP-2017-01
December 15, 2017

A Roadmap for HEP Software and Computing R&D for the 2020s

HEP Software Foundation[1]

ABSTRACT: Particle physics has an ambitious and broad experimental programme for the coming decades. This programme requires large investments in detector hardware, either to build new facilities and experiments, or to upgrade existing ones. Similarly, it requires commensurate investment in the R&D of software to acquire, manage, process, and analyse the shear amounts of data to be recorded. In planning for the HL-LHC in particular, it is critical that all of the collaborating stakeholders agree on the software goals and priorities, and that the efforts complement each other. In this spirit, this white paper describes the R&D activities required to prepare for this software upgrade.

arXiv:1712.06982v5 [physics.comp-ph] 19 Dec 2018

https://arxiv.org/abs/1712.06982

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!!

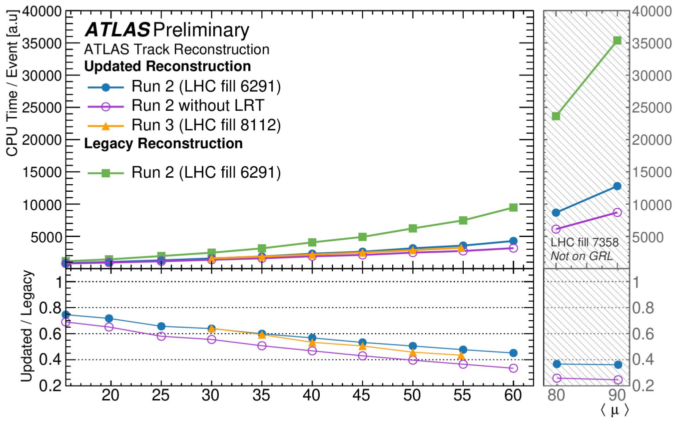

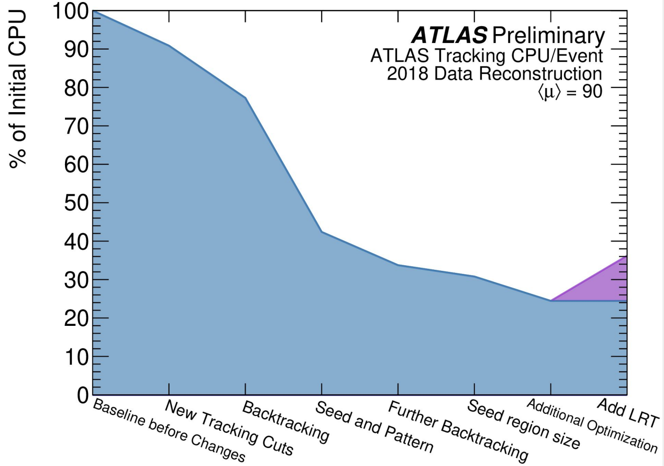

The ATLAS data processing chain

- Raw data:
  - Sensor hits, energy deposits, timing information
- Reconstructed data (AODs):
  - Momentum of tracks (4-vectors), energy in clusters (jets), particle identification, calibration
- Derived data sets (DAODs):
  - Selected analysis level information, some of which is calibrated. Starting point for analysis

- Monte Carlo: Simulated data, comparison to theory
- **Event generation** (EVNT): Calculated particle interactions
- **Simulation** (HITS): Particle interactions with detector material
- **Digitization** (RDO): Transforms simulated energy into a detector response that looks like the raw collision data
- **Reconstruction**: Performed the same way as for real data: produce AODs, DAODs, with some additional information

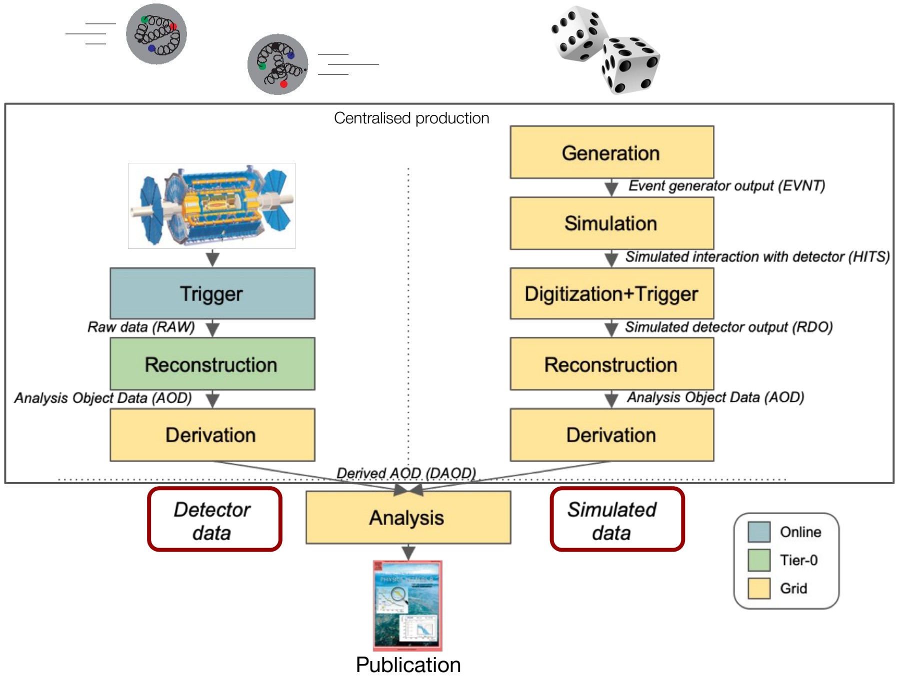

Centralised production: In essence, several steps of data processing followed by data reduction

ATLAS Software and Computing organization

Software and Computing

Computing co-Coordinator 1 [H] *Alessandro Di Girolamo (CERN)*
Computing co-Coordinator 2 [H] *Zach Marshall (Berkeley LBNL)*

Appointments

International Computing Board Chair [H] *Oxana Smirnova*

Show inactive

ADC

ADC coordinator [H]	*David Cameron (Oslo)* *David South (DESY)*

Appointments

Derivation production manager [H]	*Radek Novotny, Emma Torro Pastor*
Central Services coordinator [H]	*Alessandro De Salvo, Chris Lee*
Distributed Production and Analysis coordinator [H]	*Ivan Glushkov, Mayuko.Kataoka Maeno, Miguel Villaplana*
MC Production Coordinator [H]	*John Kenneth Anders, Ewelina Maria Lobodzinska*
Derivation Coordinator [H]	*Cristiano Alpigiani*
Monitoring and Analytics Coordinator [H]	*Aleksandr Alekseev, Dario Barberis, Markus Schulz*
Distributed Data Management Coordinator [H]	*Mario Lassnig, Shawn Mc Kee*
ADC Technical Coordinator	*Misha Borodin, Alessandra Forti, Rodney Walker*
Data Management Operations coordinator [H]	*Dimitrios Christidis*
Workflow Management Coordinator [H]	*Alexei Klimentov, Andreu Pacheco Pages*
CRIC Coordinator [H]	*Alexey Anisenkov, Alessandro Di Girolamo*
Derivation release coordinator [H]	*Oana Vickey Boeriu*
T0 coordinator	*Luc Goossens, Armin Nairz*

ATLAS Database and Metadata (ADAM)

ADAM Coordinator	*Andrea Formica (Saclay CEA)* *Davide Costanzo (Sheffield)*

Appointments

Physics Metadata Coordinator	*Pierre Antoine Delsart, Jose Enrique Garcia Navarro*
Data Engineering Coordinator	*Luca Canali*
Database Operations Coordinator [H]	*Elizabeth Gallas, Nurcan Ozturk*
Conditions Evolution Coordinator	*Andrea Formica*

Computing Resources

Computing Resource Management (CREM) Chair [H]	*Alexei Klimentov (Brookhaven BNL)*

General

Appointments

Software & Computing Speakers Committee Member [H]	*Roger Jones, Ryu Sawada, Isabel Trigger*
ICB Scrutiny Group Member [H]	*Jiri Chudoba, Doug Gingrich, Armin Nairz*
Machine Learning Forum Conveners [H]	*Walter Hopkins, Daniel Whiteson*

Software

Software Coordinator [H]	*James Catmore (Oslo)* *Nicholas Styles (DESY)*

Appointments

Physics Validation Coordinator [H]	*Giuseppe Callea, Daniel Hayden*
Core Software Coordinator [H]	*Attila Krasznahorkay, Scott Snyder*
Analysis Model Group Coordinator (AMG) [H]	*Verena Ingrid Martinez Outschoorn, Heather Russell*
Software and Computing Infrastructure Group (ASCIG) Coordinator [H]	*Dario Barberis, Chris Lee*
Upgrade Software Coordinator [H]	*Jiri Masik, Nora Emilia Pettersson*
Reconstruction Coordinator [H]	*Christos Anastopoulos, Mark Hodgkinson*
Heavy Ions software coordinator	*Sebastian Tapia Araya*
SPOT coordinator [H]	*Alaettin Serhan Mete*