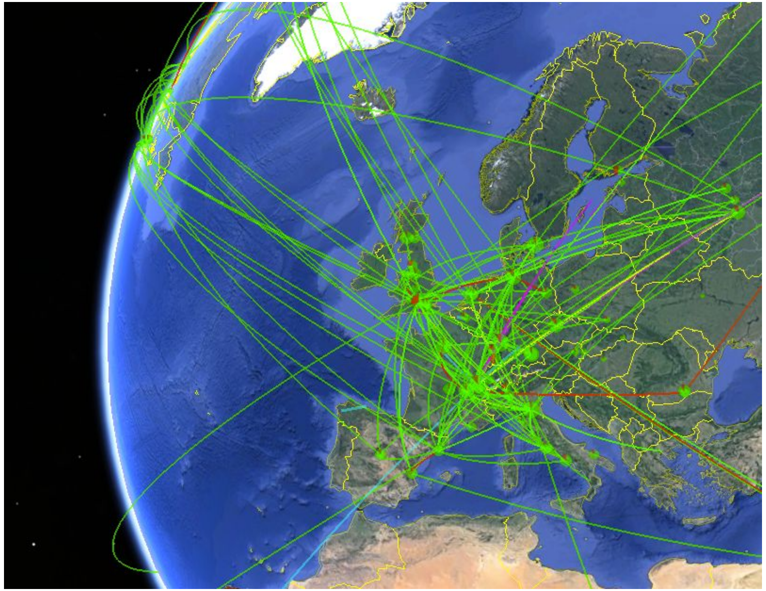

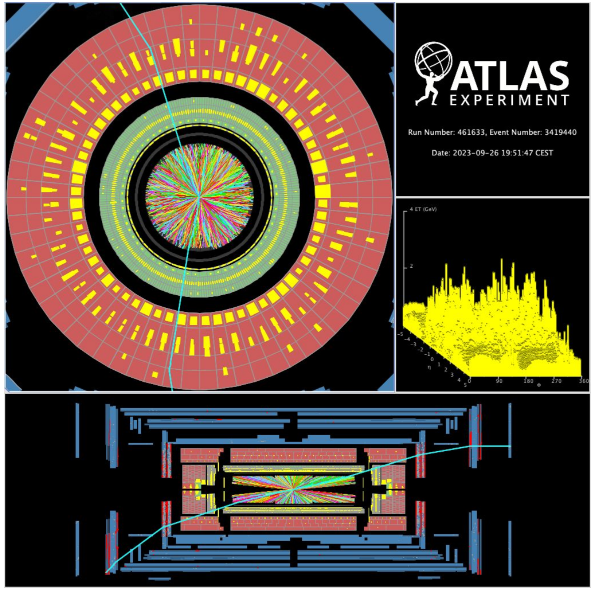

ATLAS Software & Computing

Goals, Challenges, and Opportunities for Collaboration with BNL

Workshop on ATLAS Computing and Software Activities at BNL
Navigating Distributed Computing, Storage, Compute, and Beyond
Brookhaven, March 18th, 2024

Mario Lassnig (CERN), Zach Marshall (LBNL),
Andreu Pacheco Pages (PIC) and David South (DESY)

ATLAS EXPERIMENT

LHC, ATLAS, and the Data Harvest

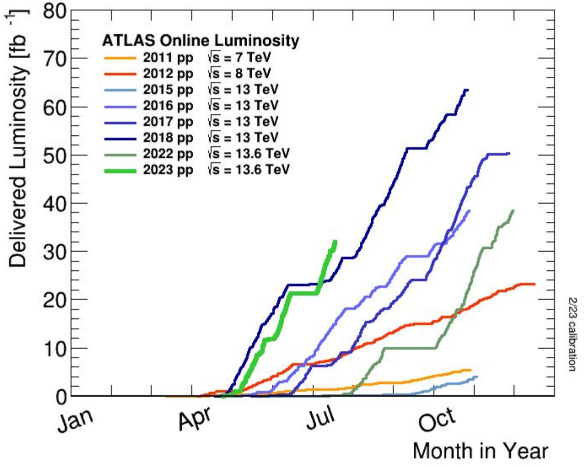

Run 1 data (2011-2013)
Centre of mass energy **7-8 TeV**

Run 2 data (2015-2018)
Centre of mass energy **13 TeV**

Run 3 data (since 2022)
Centre of mass energy **13.6 TeV**

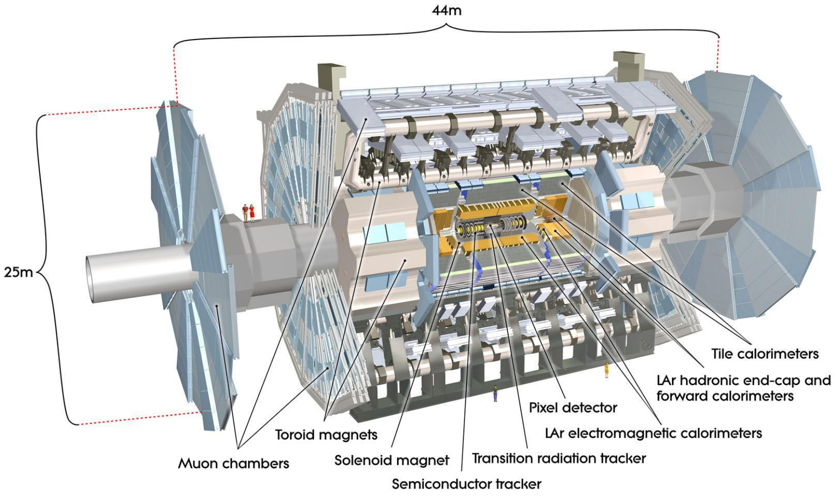

Main physics stream

Year	Raw events	SFO total volume	SFO event volume
2015	1,694,555,330	1.4 PB	828.2 KB
2016	5,387,420,813	4.9 PB	1004.8 KB
2017	5,649,311,254	5.5 PB	1 MB
2018	6,400,342,575	6.2 PB	1 MB

19 billion events collected by ATLAS

18 PB of raw data

From Detector to Publication

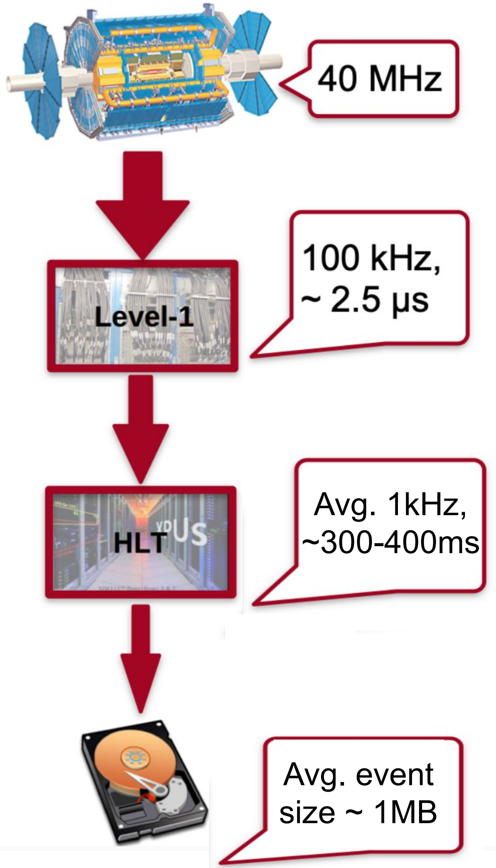

Collisions at a rate of 40 MHz

Level 1 Hardware Trigger: 100 kHz

First selection based on calorimeter and muon systems

Limit from detector and trigger hardware

High Level Software Trigger: 1 kHz

Processing time of 300-400ms

Size of HLT farm comprising ~100k cores

Final output rate ~ 1kHz

In Run-3 this increased substantially

Acceptance rate at 3 kHz

Event size increased to 1.8MB

The Data Processing Chain

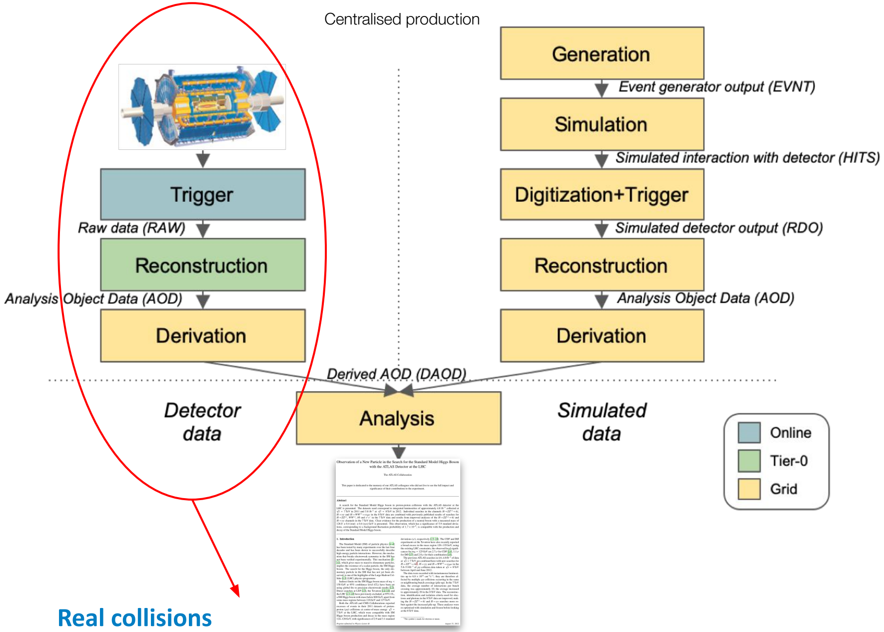

Real collisions

RAW Data : Sensor hits, energy deposits, timing information

Analysis Object Data (AOD): 4-vector momentum of tracks

Energy in jet clusters, particle identification, first calibrations

Derived AODs (DAODs): Starting point for analysis

Selected analysis level information with full calibration

From Detector to Publication

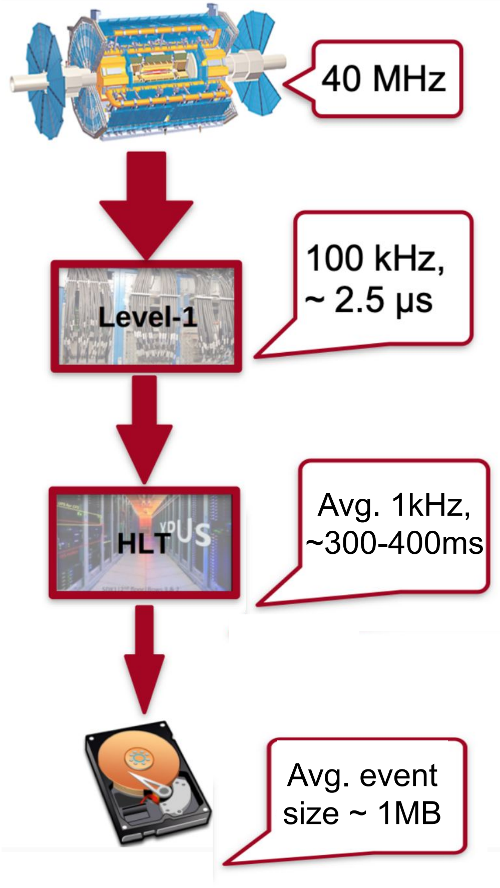

Collisions at a rate of 40 MHz

Level 1 Hardware Trigger: 100 kHz

First selection based on calorimeter and muon systems

Limit from detector and trigger hardware

High Level Software Trigger: 1 kHz

Processing time of 300-400ms

Size of HLT farm comprising ~100k cores

Final output rate ~ 1kHz

In Run-3 this increased substantially

Acceptance rate at 3 kHz

Event size increased to 1.8MB

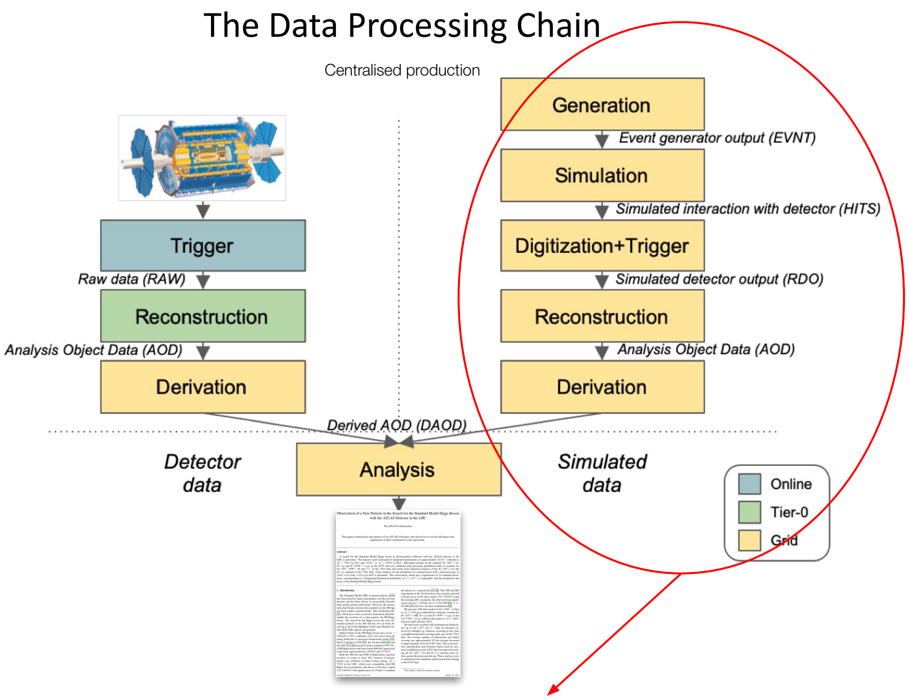

Monte Carlo Simulation

Event generation	EVNT	Calculated particle interactions
Simulation	HITS	Interactions with detector material
Digitization	RDO	Simulated energy transformed to detector response
Reconstruction	AOD	Performed the same way as for data

Resource Usage — Compute — Last 12 months

- Smooth running of our resources with a variety of workflows
- Example campaigns in the last few months
  - Re-reconstruction of all 2023 MC with new pile-up profile *mc23d*
  - New evgen in preparation for 2024 MC *mc23e*

	max	avg
MC Simulation Full	627 K	243 K
MC Reconstruction	406 K	128 K
Group Production	241 K	75.1 K
User Analysis	103 K	58.9 K
MC Event Generation	209 K	46.4 K
MC Simulation Fast	457 K	39.2 K
Data Processing	156 K	10.6 K
Group Analysis	35.2 K	7.97 K
t0_processing	37.1 K	7.71 K
MC Resimulation	69.6 K	4.35 K
MC Merge	12.5 K	2.38 K

Slots of Running jobs by ADC activity

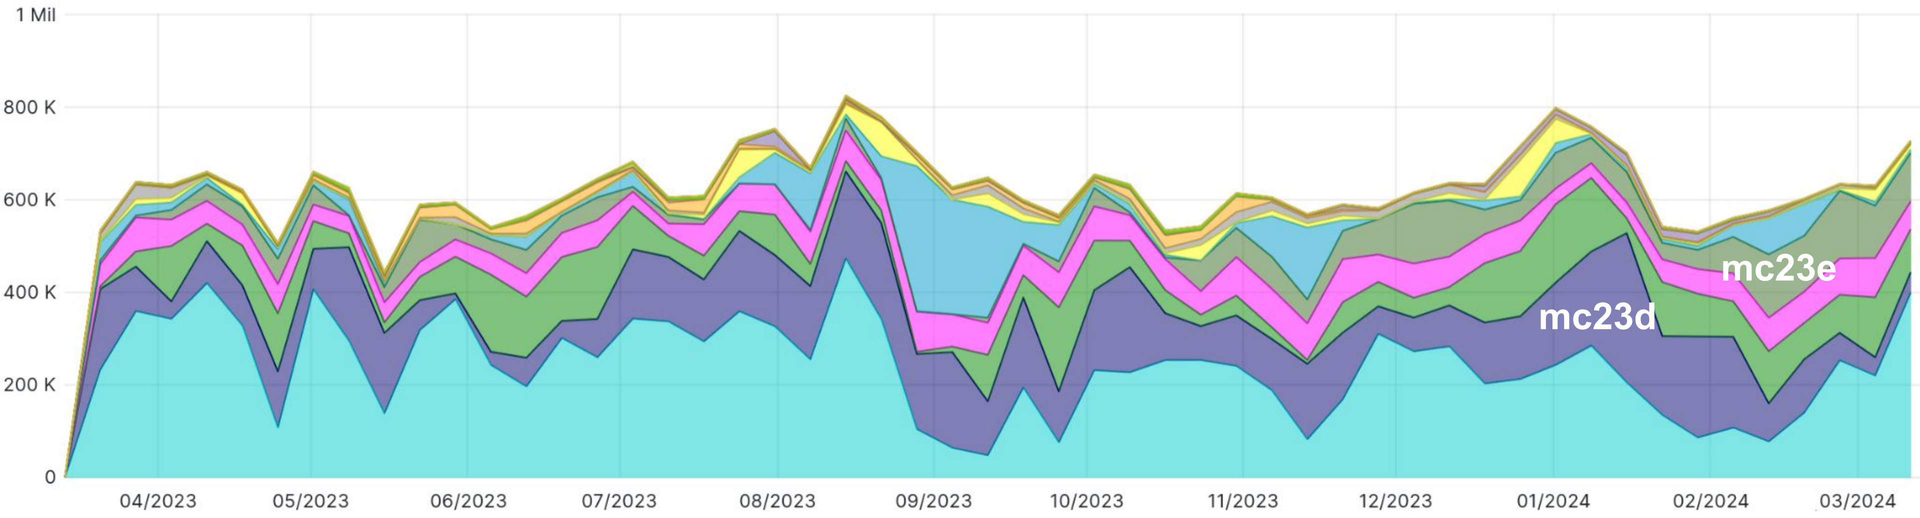

Resource Usage — Compute — Last 12 months

- Grid sites, especially Tier-2s, providing good and beyond-pledge CPU
- HPC (Vega, Perlmutter) and cloud (mainly from the HLT at P1) continue to make significant contributions

	max	avg
GRID	6.96 Mil	5.47 Mil
Pledges	4.10 Mil	4.08 Mil
hpc	5.78 Mil	1.62 Mil
cloud	1.35 Mil	621 K
Tier-0	477 K	101 K

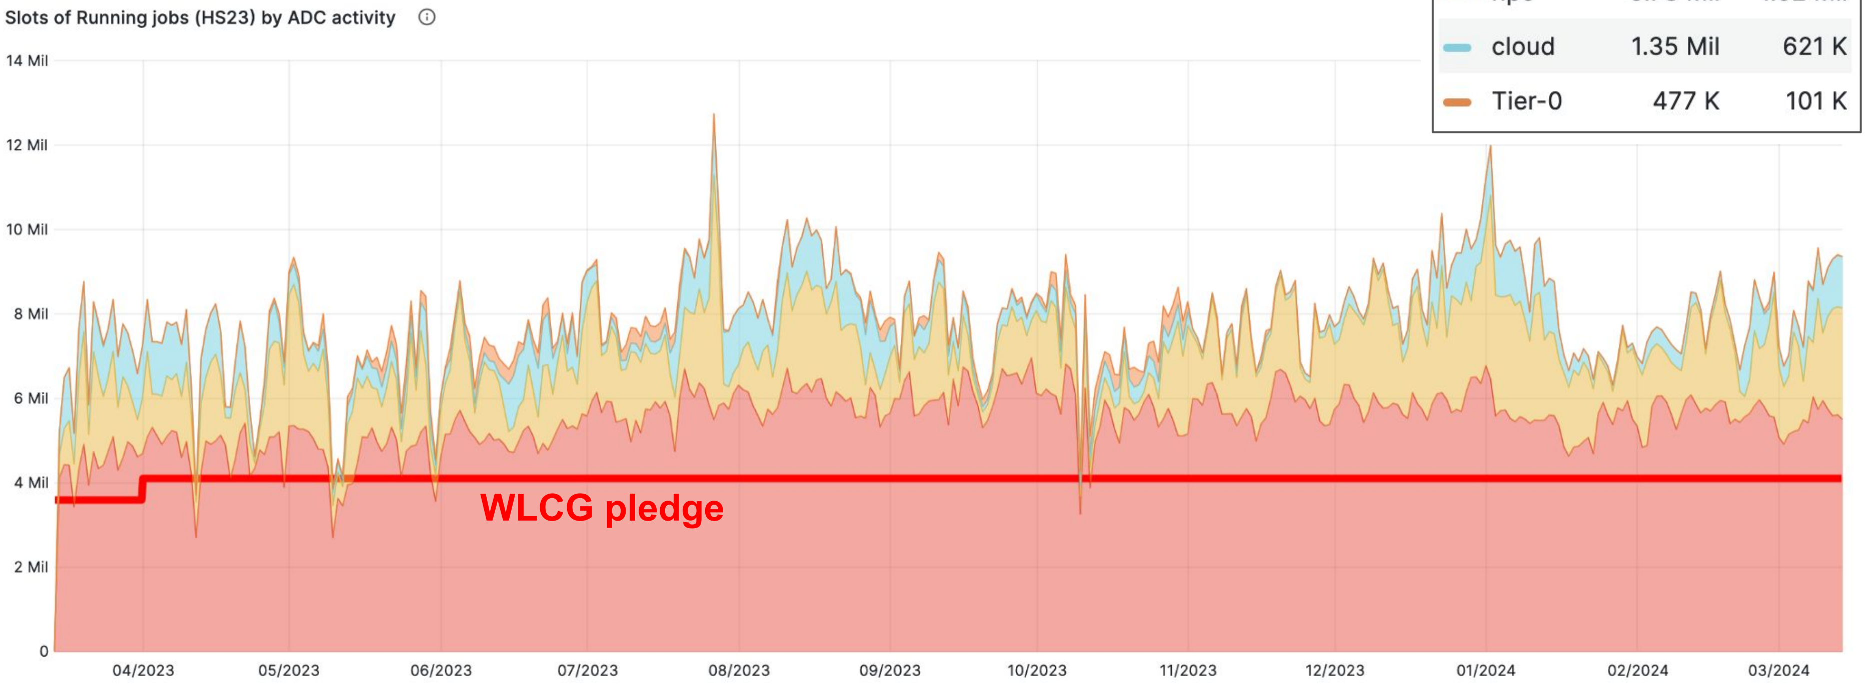

Resource Usage – Tape – Since Start of Run 3

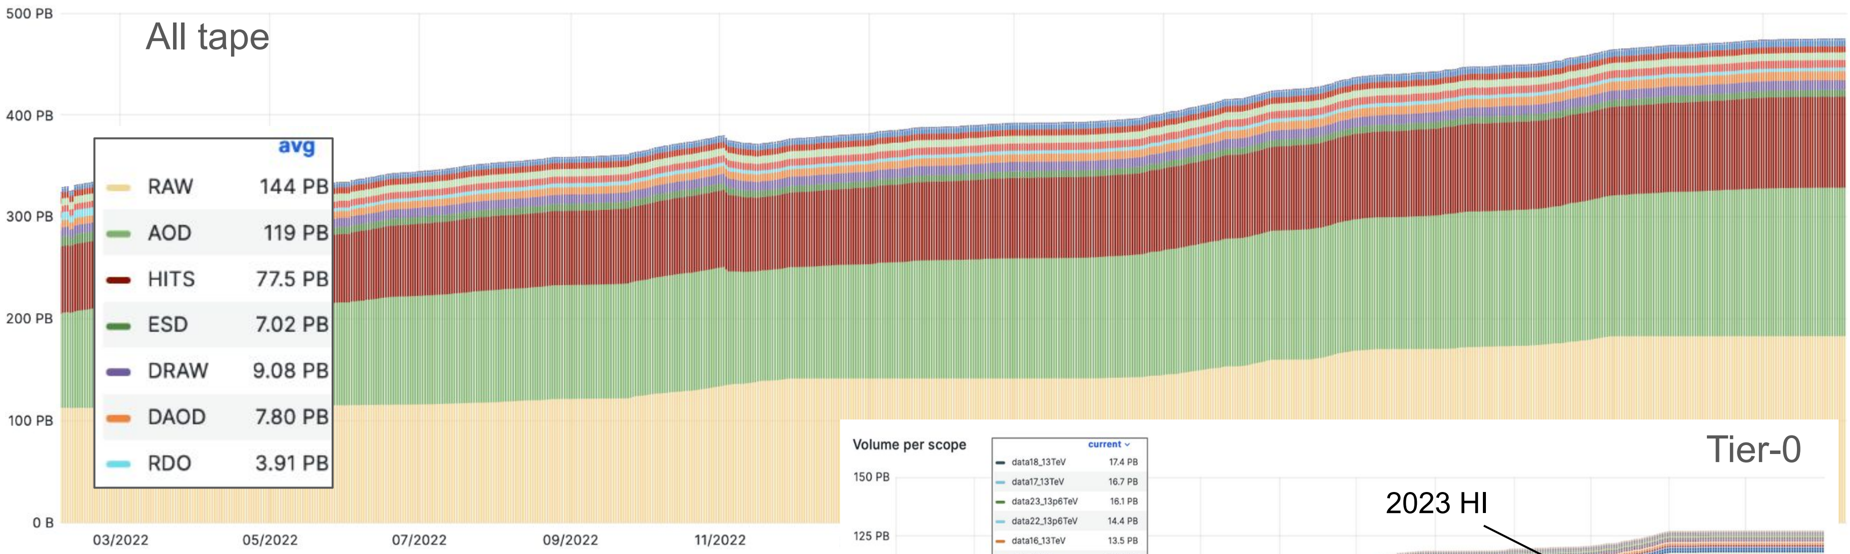

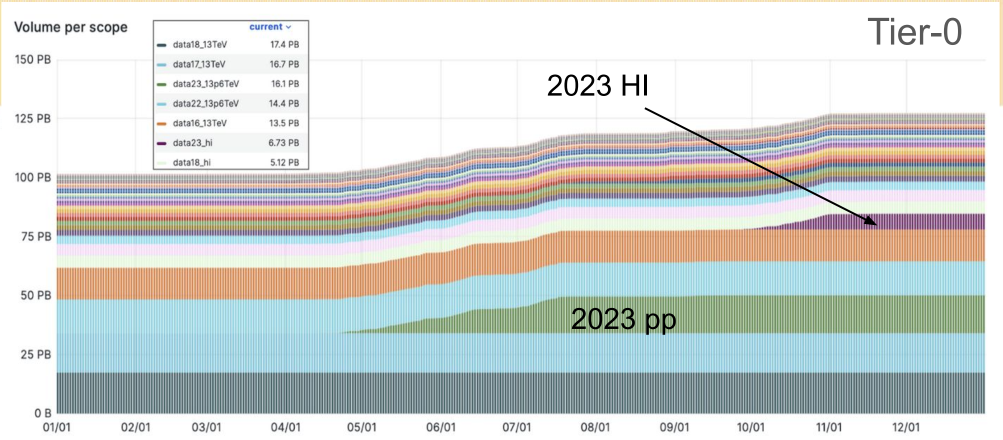

- Another year of data taking done

- 2023 was shorter than planned
- Hoping for a stronger 2024 and 2025

Resource Usage – Disk – Last 12 months

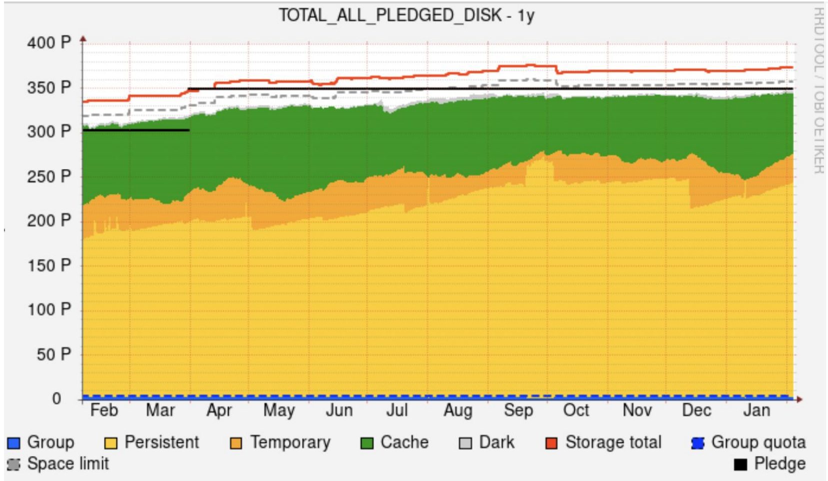

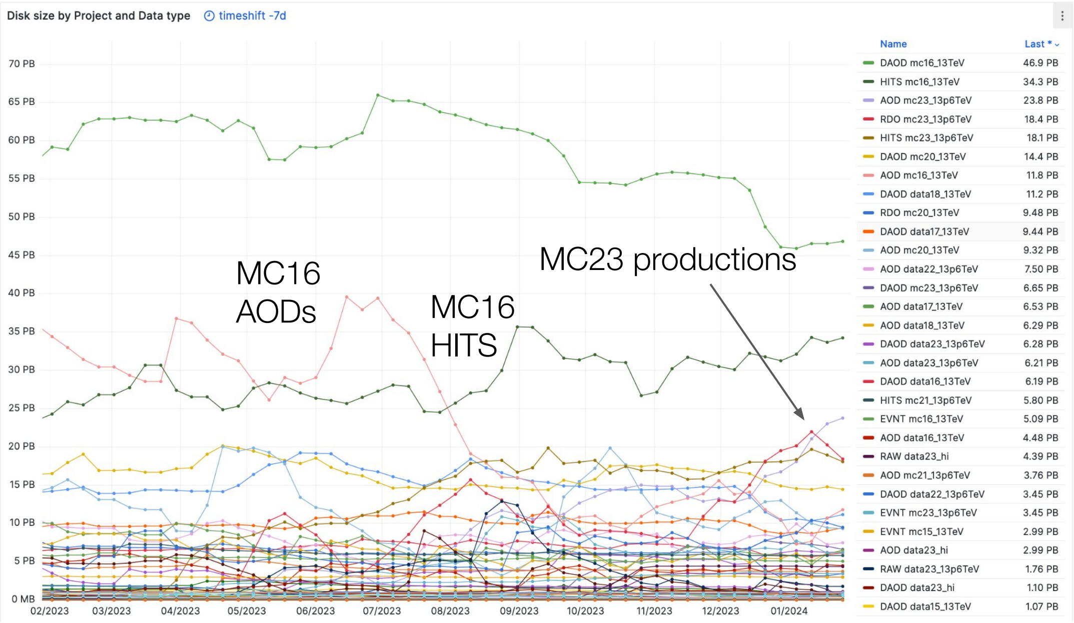

- Keeping storage under control

- Regular applications of the lifetime model
- Creative work by DDM Operations
- Group production obsoletions

- Cache is still comfortable
- 2024 pledges are starting to come in

- Space still dominated by R21 MC16/20
- Slowly moving off disk as MC23 comes in

Experiment job mix

Globally configured shares are employed to allocate the available resources among the activities

- Done by **agreement** between the various physics groups
- **Hierarchical** implementation of the configuration parameters
- Related activities have the opportunity to **inherit idle resources**

Essentially two major categories of jobs

Production
- Data processing and reprocessing
- Event Generation / Simulation / Reconstruction
- Derivation

Analysis
- User Analysis
- Group Analysis

The main activity at a given time can depend on many things

- Data **reprocessing** or Monte Carlo **production** campaigns
- **Conference** deadlines, need for an increase for user analysis
- Global pandemics…

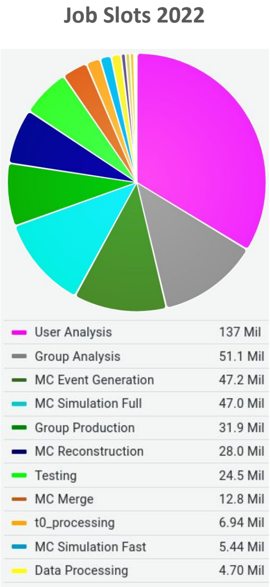

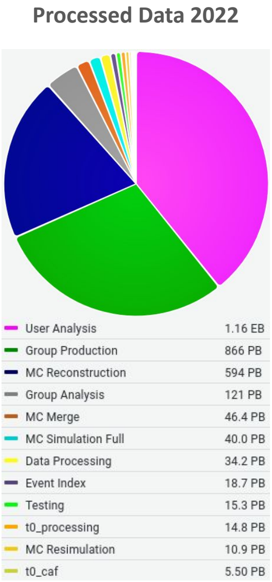

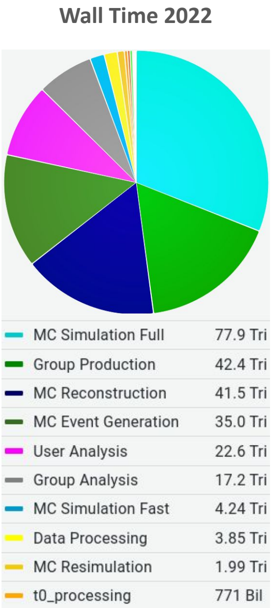

ATLAS Data Management Scale

- A few numbers about the ATLAS scale
  - 1B+ files, 850+ PB of data, 400+ Hz interaction rate
  - 120 data centres, 5 HPCs, 3 clouds, 1000+ users
  - 1.5 Exabytes/year transferred
  - 3 Exabytes/year uploaded & downloaded

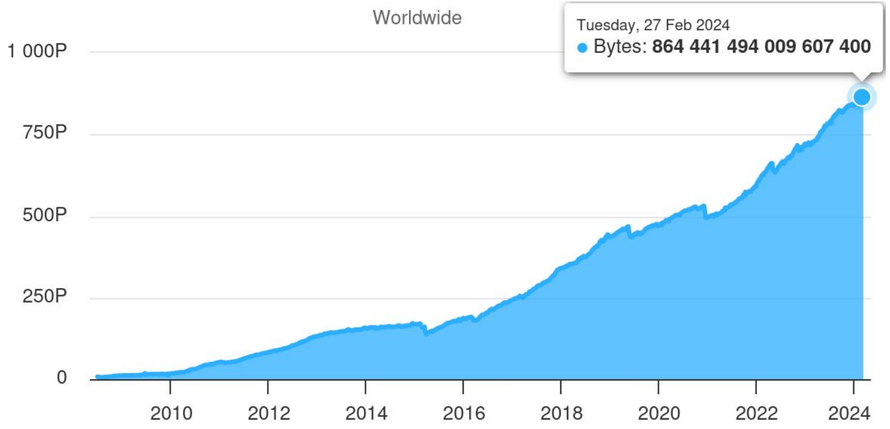

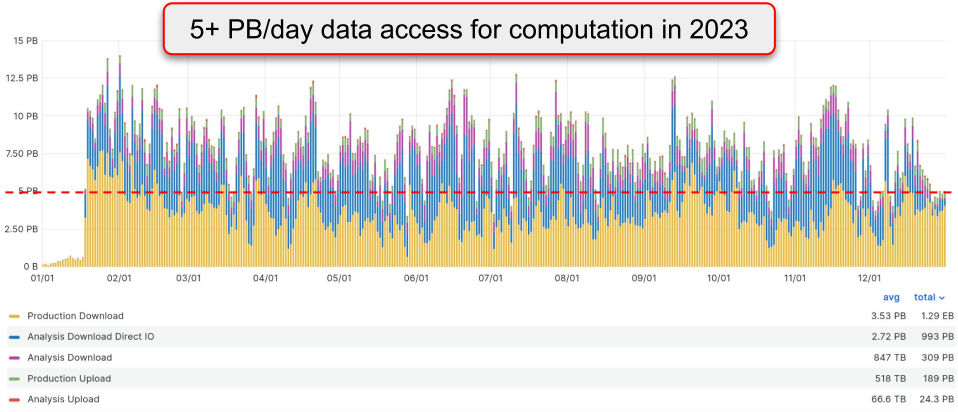

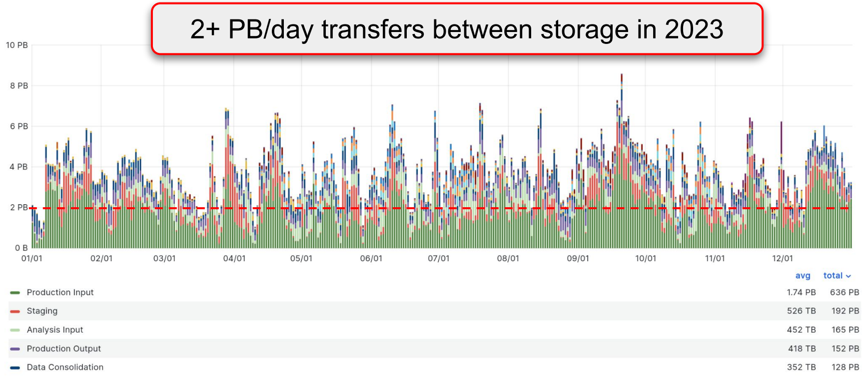

BNL resources — Last Ten Years

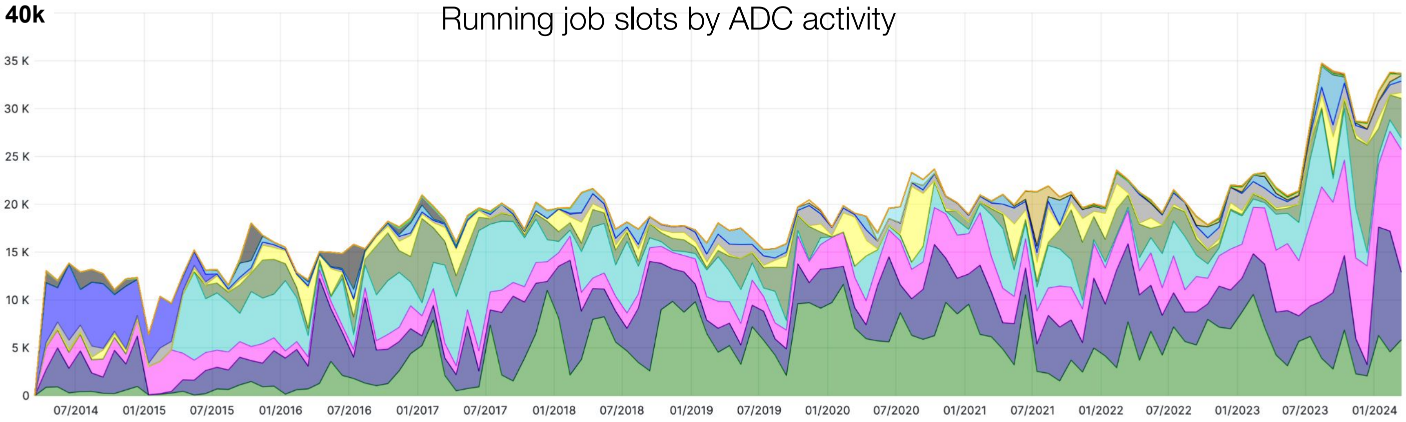

	max	avg
Group Production	11.6 K	4.41 K
MC Reconstruction	12.6 K	4.12 K
User Analysis	12.8 K	3.07 K
MC Simulation Full	12.5 K	2.78 K
MC Event Generation	11.3 K	1.49 K
Data Processing	7.88 K	855
Group Analysis	2.99 K	709
MC Simulation	7.89 K	540
MC Simulation Fast	5.23 K	412
Validation	4.69 K	198
MC Resimulation	5.61 K	127

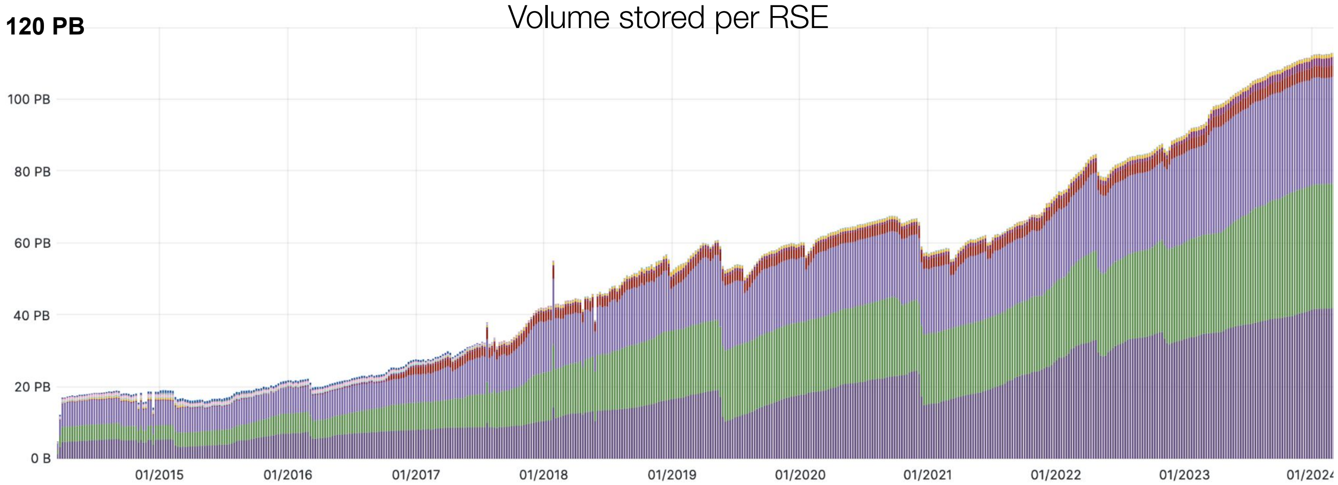

	max	avg
BNL-OSG2_MCTAPE	42.0 PB	17.2 PB
BNL-OSG2_DATATAPE	34.8 PB	16.2 PB
BNL-OSG2_DATADISK	29.8 PB	15.3 PB
BNL-OSG2_GROUPTAPE	3.14 PB	1.98 PB
BNL-OSG2_LOCALGROUPDISK	2.45 PB	772 TB
BNL-OSG2_SCRATCHDISK	1.45 PB	521 TB
BNL-OSG2_PHYS-HI	236 TB	174 TB

ATLAS Distributed Computing

- ATLAS Distributed Computing (ADC) comprises the hardware, software, and operations to
  - Support **distributed computing activities** of the experiments
  - Support the **evolving needs** of the experiment

- Running 24 / 7 / 365
  - **Computing never stops**
  - 80+ people contributing centrally
  - 50+ people across the WLCG

- Four major areas
  - Physics activities requiring computing
  - Infrastructure & operations
  - Data management
  - Workload & workflow management

ATLAS DISTRIBUTED COMPUTING
March 2024

ADC COORDINATION
Mario Lassnig, Andreu Pacheco Pages

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
Coordination
V. Garonne

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

Operations
- Computing Run Coordination
- 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
R. Walker, F. Barreiro Megino

System
- Workflow Definition
- Workload Management
- Workload Execution

Operations
- System Deployment
- Monitoring

Research
- Data Analytics
- Analysis Facilities
- Cloud
- HPC

- Plus many task forces and working groups, e.g., HPC or monitoring

Organisational observations

- 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
    - 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

- Therefore, we did the ADC reorganisation
    - **Consolidate** work areas… but still a long way to go
    - More **streamlined** and **flexible** organisation
    - **Improved communication**
    - Bringing **more people into leadership roles**

ADC Operations

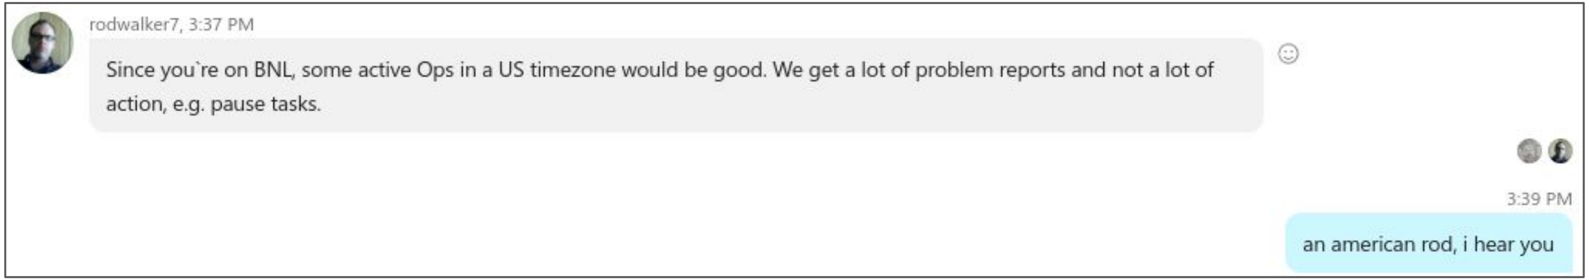

- Distributed Processing and Analysis
  - Managing pledged compute resources
  - Production campaign support, User analysis support
- Data Management
  - Managing pledged storage resources
  - Data replication policies and lifecycle processes
- Shared contributions
  - Immediate and long-term help with user and site issues
  - Supporting R&D initiatives
  - Developments for more efficient operations

- Both areas are **critically understaffed**!
  - Difficult to sustain because these areas require long-term investment to be useful
  - Time Zones, holidays, accidents, etc… are real and actual issues

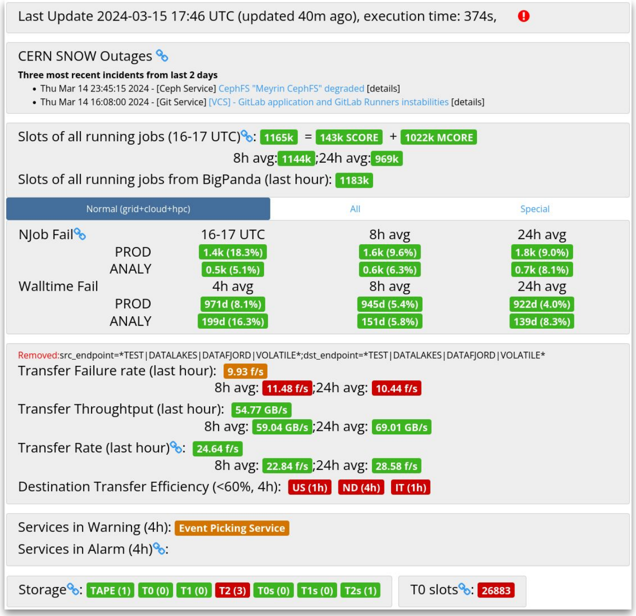

Our Data Challenges and Opportunities

The High Luminosity upgrade to the LHC

10 times increase in **accelerator performance**
Leads to more and bigger, complex events
10 times increase in **data volume/usage**
In a very tight computing capacity envelope

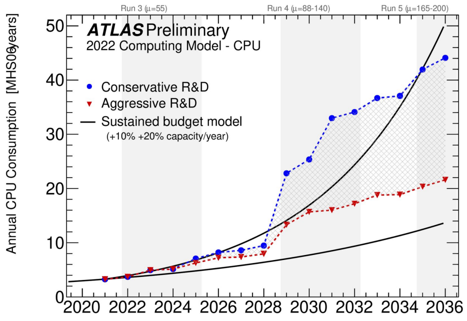

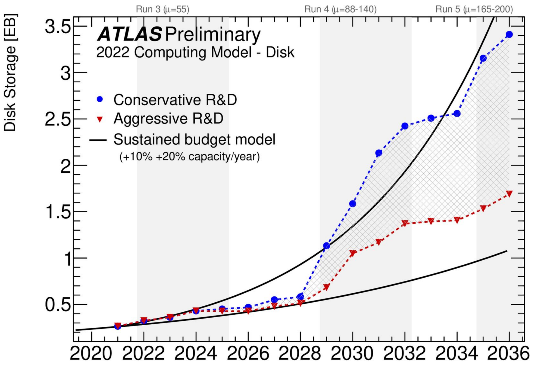

We cannot compromise physics performance

Long-term R&D programme to address the gap

To support the **European Strategy for Particle Physics**

Community-driven computing R&Ds
- Advanced **software-defined networks** (SDNs)
- **Smart content delivery** and caches
- New analysis **data formats** and models
- Integration of new **external developments**
- **Industry collaborations** for new technologies

Collaborations with other sciences
- Shared infrastructure with other big communities
- Prototype of a common European Data Infrastructure

MSc and PhD studies to train our future computing engineers

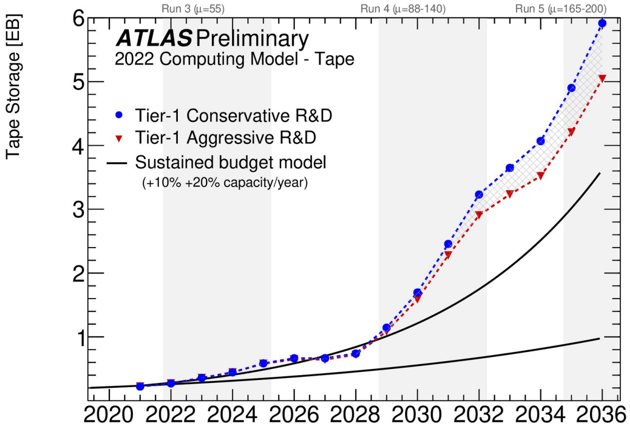

We are working on keeping between the lines!

- Software improvements
  - Visible impact
  - Finalized requests in Feb
- Slight decrease in CPU requirements
  - Full simulation getting faster
- Slight decrease in Tape requirements
  - RAW data getting smaller

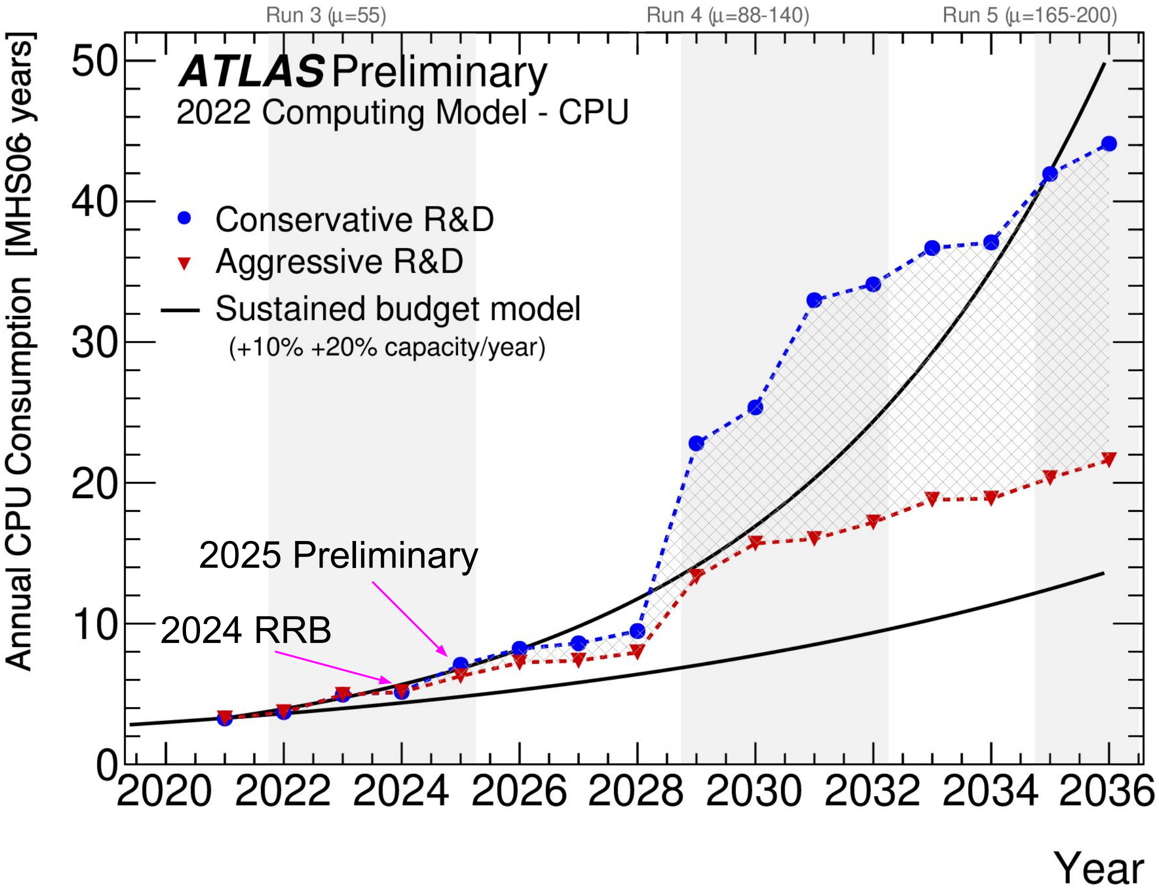

We are working on keeping between the lines!

- Software improvements
  - Visible impact
  - Finalized requests in Feb
- Slight decrease in CPU requirements
  - Full simulation getting faster
- Slight decrease in Tape requirements
  - RAW data getting smaller

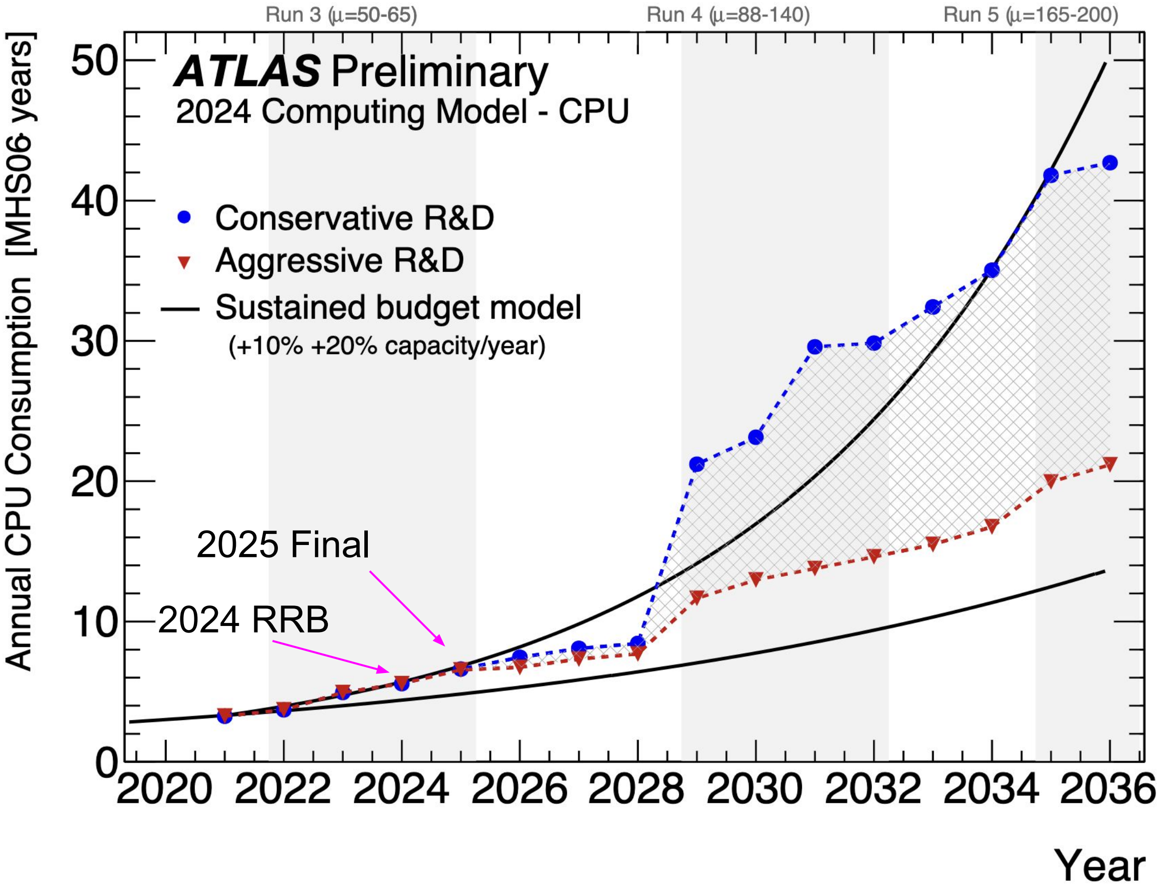

HL-LHC Data Challenges

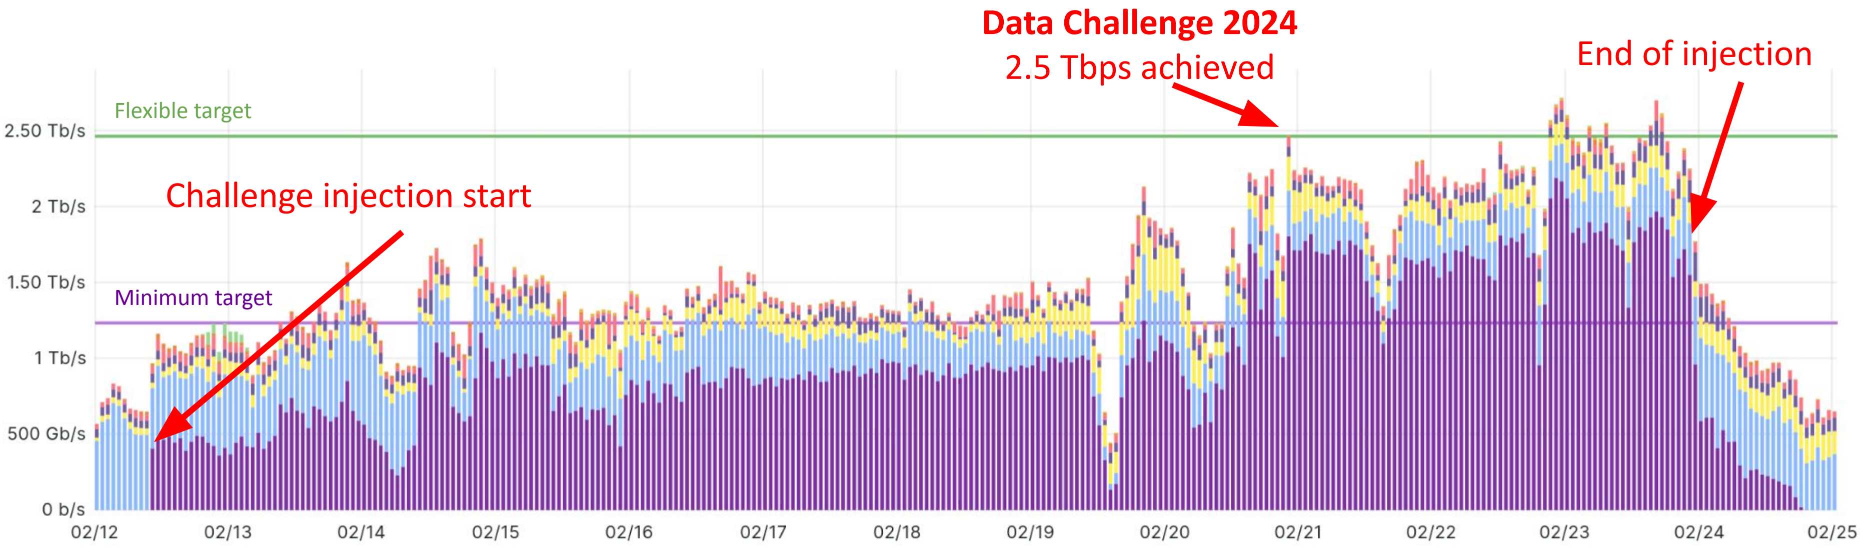

	max	avg
Data Challenge	2.19 Tb/s	987 Gb/s
atlas	706 Gb/s	316 Gb/s
alice xrootd	349 Gb/s	114 Gb/s
cms	271 Gb/s	56.8 Gb/s
cms xrootd	191 Gb/s	67.7 Gb/s
lhcb	83.1 Gb/s	2.35 Gb/s
belle	38.9 Gb/s	9.33 Gb/s
dune	28.6 Gb/s	5.47 Gb/s

2020 estimation of HL-LHC needs

- **4.8 Tbps** of total network capacity for the Run-2 computing model
- **9.6 Tbps** for the Run-3 (and beyond) computing model

Data Challenges until HL-LHC startup

- **Bi-annual steps** of 25% expected capacity
- With an **accompanying R&D** programme for software and hardware

DC24 exposed real issues in our infrastructure, esp. in the transfer scheduling

BNL throughput during DC24

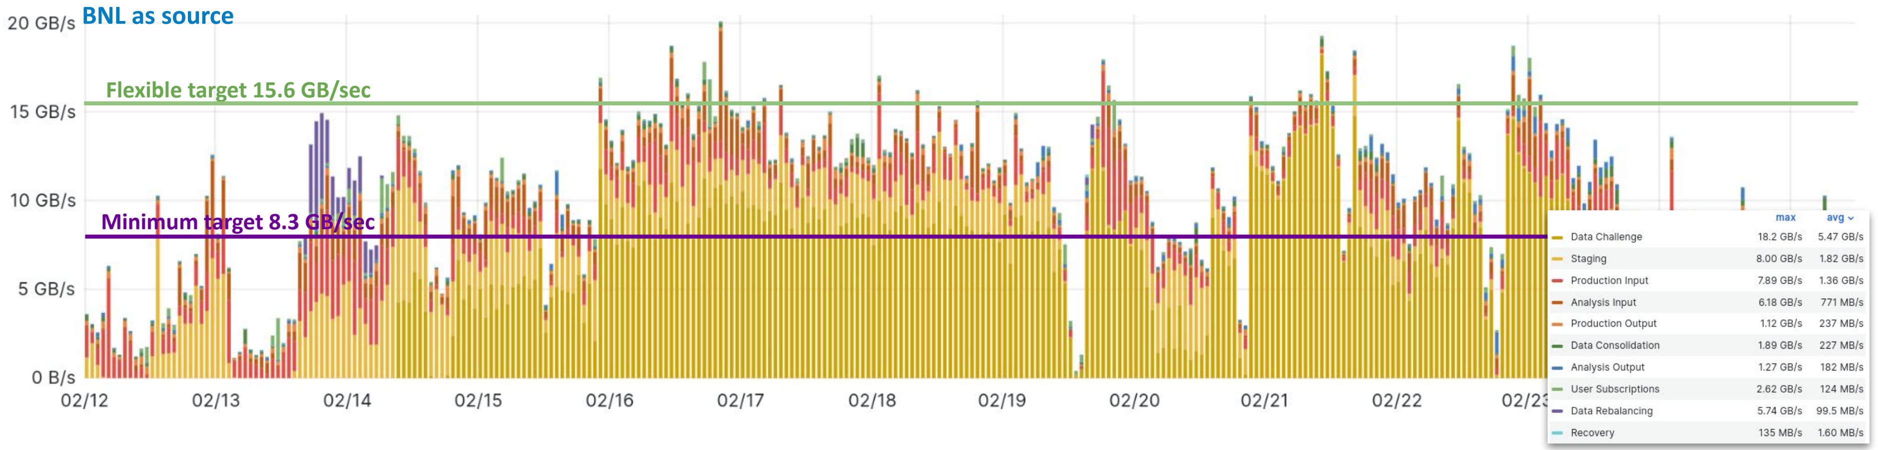

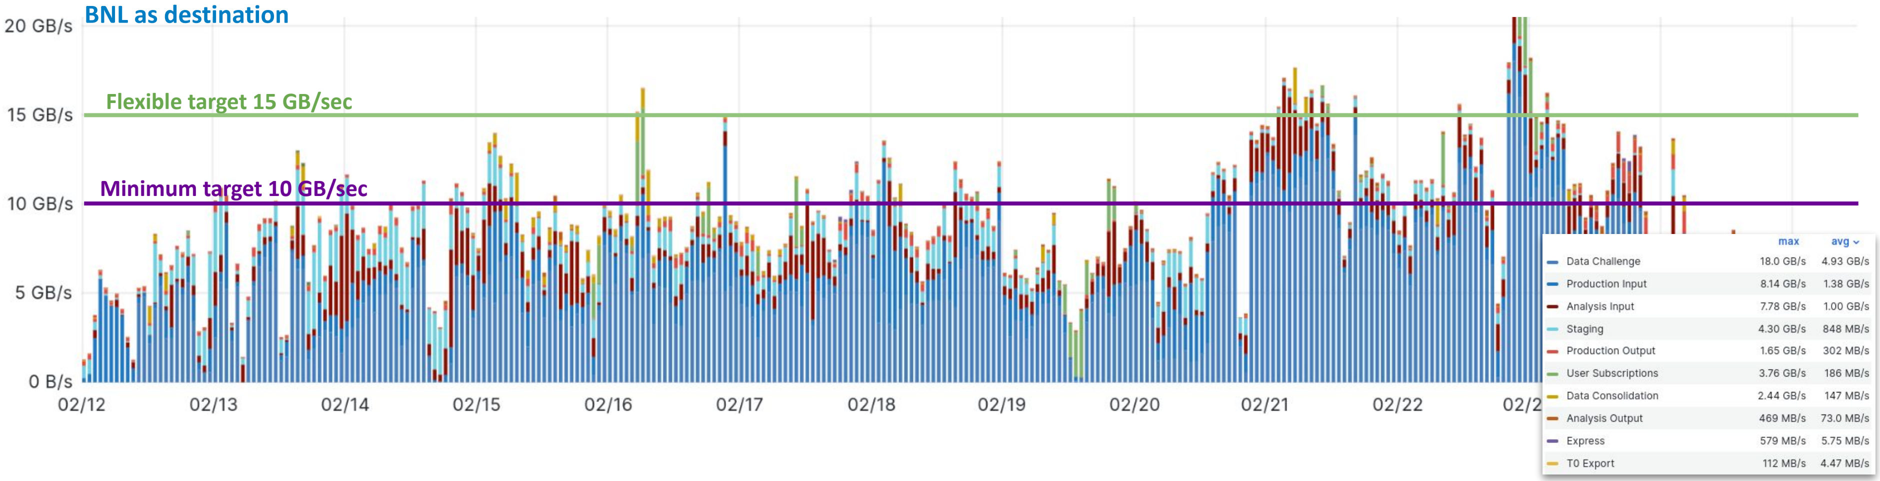

ADC Milestones

Distributed Computing							
B	DC	1		Transition to tokens	12/2025	12/2025	Q4 2025
B	DC		1.1	Submission from Harvester to all HTCondor CEs with tokens	3/2022	3/2022	Q1 2022
B	DC		1.2	All users move from VOMS to IAM for X509	12/2022	6/2023	Q2 2023
B	DC		1.3	All job submission and data transfers use tokens	12/2025	12/2025	Q4 2025
B	DC		1.4	Evaluate ATLAS progress as well as the WLCG full token transition progress	12/2023	12/2023	Q4 2023
B	DC	2		Storage evolution	12/2025	12/2025	Q4 2025
B	DC		2.1	No GridFTP transfers at any site	3/2022	3/2022	Q1 2022
B	DC		2.2	SRM-less access to tape	12/2025	12/2025	Q4 2025
B	DC		2.3	Recommended transition plan from DPM completed	12/2021	12/2021	Q4 2021
B	DC		2.4	Transition plan from all DPM sites	12/2022	12/2022	Q4 2022
B	DC		2.5	All sites moved away from DPM	6/2024	6/2024	Q2 2024
B	DC	3		Next operating system version	6/2024	6/2024	Q2 2024
B	DC		3.1	Ability to run on "future OS" on grid sites	12/2022	10/2023	Q4 2023
B	DC		3.2	Central services moved to "future OS"	12/2023	6/2024	Q2 2024
B	DC		3.3	(CentOS 7/8 EOL)	6/2024	6/2024	Q2 2024
B	DC	4		Network infrastructure ready for Run 4	11/2027	11/2027	Q4 2027
B	DC		4.1	Network challenge at 10% expected rate	11/2021	11/2021	Q4 2021
B	DC		4.2	Network challenge at 25% expected rate	11/2023	3/2024	Q1 2024
B	DC		4.3	Network challenge at 60% expected rate	11/2025	11/2025	Q4 2025
B	DC		4.4	Network challenge at 100% expected rate	11/2027	11/2027	Q4 2027
B	DC		4.5	Check on changes to network usage patterns	3/2024	3/2024	Q1 2024
C	DC	5		Integrating next generation of HPCs	12/2024	12/2024	Q4 2024
C	DC		5.1	Integration of at least 2 EuroHPC sites	12/2022	12/2022	Q4 2022
C	DC		5.2	Integration of next generation US HPCs for production	6/2023	6/2023	Q2 2023
C	DC		5.3	Integration of a further 2 EuroHPC sites	12/2024	12/2024	Q4 2024
A	DC	6		Exploratory R&D on non-x86 resources and next generation HPC	12/2023	12/2025	Q4 2025
A			6.1	Employing and/or deployment of ARM resources for CPU	12/2022	9/2023	Q3 2023
A			6.2	Support for workflows using a significant number of GPUs	12/2025	12/2025	Q4 2025
			6.3	Test FastCaloSimGPU at reasonably large scale	6/2024	6/2024	Q2 2024

A	DC	7		HL-LHC datasets replicas and versions management	6/2024	6/2024	Q2 2024
A	DC		7.1	Replicas and versions detailed accounting	12/2022	12/2023	Q4 2023
A	DC		7.2	DAOD replicas reduction	12/2023	12/2023	Q4 2023
A	DC		7.3	DAOD versions reduction	6/2024	6/2024	Q2 2024
A	DC		7.4	Recommendations from dynamic data working group	12/2023	12/2023	Q4 2023
A	DC		7.5	DAOD recreation by Prodsys/Panda checkpoint	12/2023	12/2023	Q4 2023
A	DC		7.6	DAOD recreation by Prodsys/Panda production	12/2024	12/2024	Q4 2024
B	DC	8		Data Carousel for storage optimization	12/2025	12/2025	Q4 2025
B	DC		8.1	Investigate with sites the cost of Tape infrastructure and the estimated cost in case of sensible increase of read/write throughput	12/2022	12/2025	Q4 2025
A	DC		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.	12/2023	12/2023	Q4 2023
A	DC		8.3	Smart writing investigations and evaluated cost benefits to ATLAS	6/2024	6/2024	Q2 2024
A	DC	9		Disk management: secondary(cached) dataset	6/2023	12/2023	Q4 2023
A	DC		9.1	Evaluate the impact on job brokering and task duration if disk space for secondary data is varied	6/2023	6/2023	Q2 2023

A	DC	10		Evaluation of commercial cloud resources	12/2023	12/2023	Q4 2023
A	DC		10.1	Commission Google as a Grid site	12/2022	12/2022	Q4 2022
A	DC		10.2	Trial use as a “bursty” resource	3/2023	3/2023	Q1 2023
A	DC		10.3	Total Cost of Ownership evaluation	12/2023	12/2023	Q4 2023
B	DC	11		Optimising the user analysis experience	12/2023	12/2023	Q4 2023
B	DC		11.1	Evaluation of which sites should be used to run analysis jobs	3/2023	12/2022	Q4 2022
B	DC		11.2	Implement new mechanisms to monitor and reduce the tails of analysis tasks, to reduce waiting time to complete all jobs	12/2023	12/2023	Q4 2023
A	DC	12		Sustainability in ATLAS Computing	12/2025	12/2025	Q4 2025
A	DC		12.1	What is current picture, where can we improve?			
A	DC		12.2	Reduce current failed wall time			
A	DC		12.3	Investigate potential adaptations needed by ATLAS to exploit savings			

First column — Original dates

Second column — Latest dates, to be updated next week

Latest presentation from October 2023 S&C Week

..And Demonstrators

- Review next week
  - Input from November ADC TIM to be incorporated
- Further progress reports at Oslo S&C in June
- Demonstrators should “have demonstrated” by this summer
  - If they are to be included in the HL-LHC TDR !
  - Ideally also written up as CHEP 2024 contributions

	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

Name / Brief Topic
Remote read of selected data vs copy of whole file
Stress-test xcache by running many (or all) the US analysis jobs at a single site and have all inputs done through xcache.
Adaptive data placement/job scheduling using Intel’s Loihi2 neuromorphic computing platform.
Use ATLAS-Google site as “bursty” resources (demonstrate ATLAS bulk campaign [reprocessing or MC production] on this site)
ATLAS-Google site for analysis workflows and new technologies for physics analysis
ARM based PanDA queue setup and operation for Athena porting and Physics Validation
Recreate DAOD datasets on demand by PanDA using Data Carousel (Delete DAOD datasets from the lifetime model exception list and reproduce them on demand if needed)
Using Xrootd to seamlessly integrate S3 storage with ATLAS DDM system, including Google and Amazon cloud S3 storage.
Large scale AI/ML services across diverse resources, complex analysis workflows
General ARM OSG queue at BNL
Tiered storage with different storage classes (Object Store, HPC, ..)
Storage for data intensive application
Optimized writing to tape for efficient reading
Measurement of ATLAS applications energy consumption and data center optimization
Tape smart writing in Data Carousel
Sustainability in ATLAS Computing

Analysis Facilities

- Important topic, recent talk on ATLAS position. To be discussed at WLCG/HSF w/shop at DESY
- Establish feedback channel between user and computing communities
  - Allow us to understand what's really needed
    - We need to define what the users should be able to do
    - They also need to tell us what they need
  - **What do *all* of our users need?**
    - Look to the bigger picture rather than super-user / edge case solutions

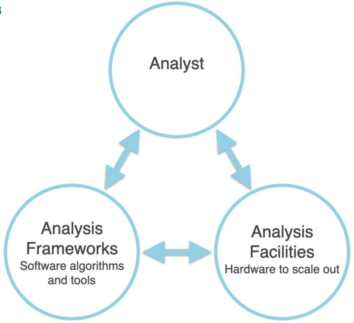

- Employ use-case driven R&D to focus the definition and contribute to the overall strategy
  - Concentrate strategic core infrastructure evolution
    - AAI, evaluation of object stores, containers, caches
    - Integration of AFs with the grid, particularly to facilitate storage access
- Decrease the isolation between already existing and upcoming AFs
  - Such as those at CERN (lxplus), SLAC, UChicago, BNL
  - As well as the more limited access variants at DESY (NAF), Valencia…
  - Aligning the current AFs can have a positive influence on the future AF development and design
  - Extra benefit to users who see smooth transitions between AFs

Opportunistic Resources: HPC

- ATLAS continues to exploit significant HPC resources
  - EuroHPC (Vega, Karolina) and US HPC (Perlmutter, TACC, …)
  - We do see a downward trend in the amount of HPC compute we are able to harvest
  - Mainly driven by getting fewer slots on Vega (which was expected)

- ADC HPC task force
  - Understand current landscape
  - Where we want to focus
  - ATLAS policy is not to use *every* HPC on the market
  - Rather exploit those easiest to interface

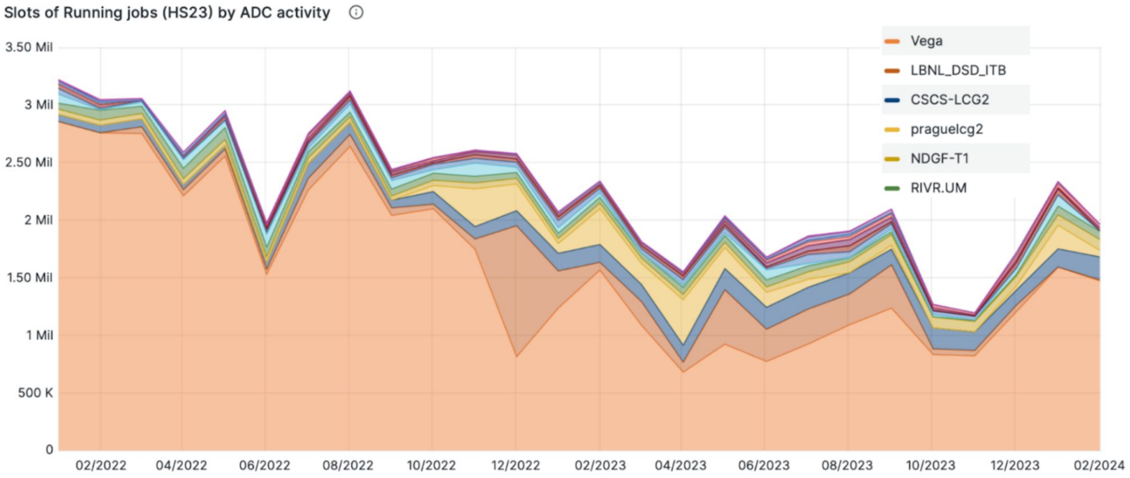

Opportunistic Resources: Clouds

- ATLAS Google Project
  - 15 months, agreement with a flat cost rate
  - ATLAS workflows with elastic compute queue
  - Uses Kubernetes and S3-style storage
  - Many successful studies with TCO close to publication
  - Google Phase 3 project starting up
- SEAL Web3 Cloud Storage demonstrator currently ongoing
- FRESNO Amazon AWS project concluded

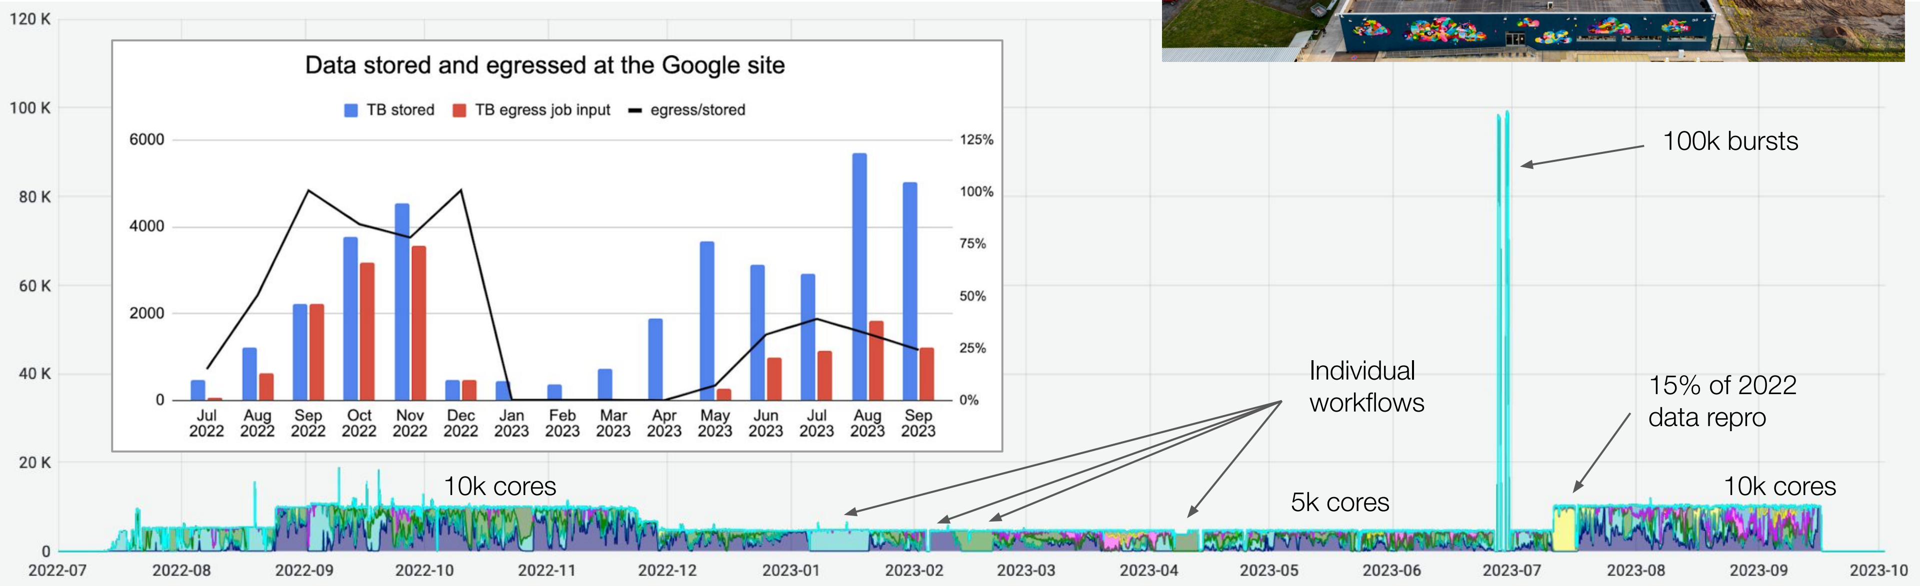

Opportunistic Resources: HLT Farm, Sim@P1

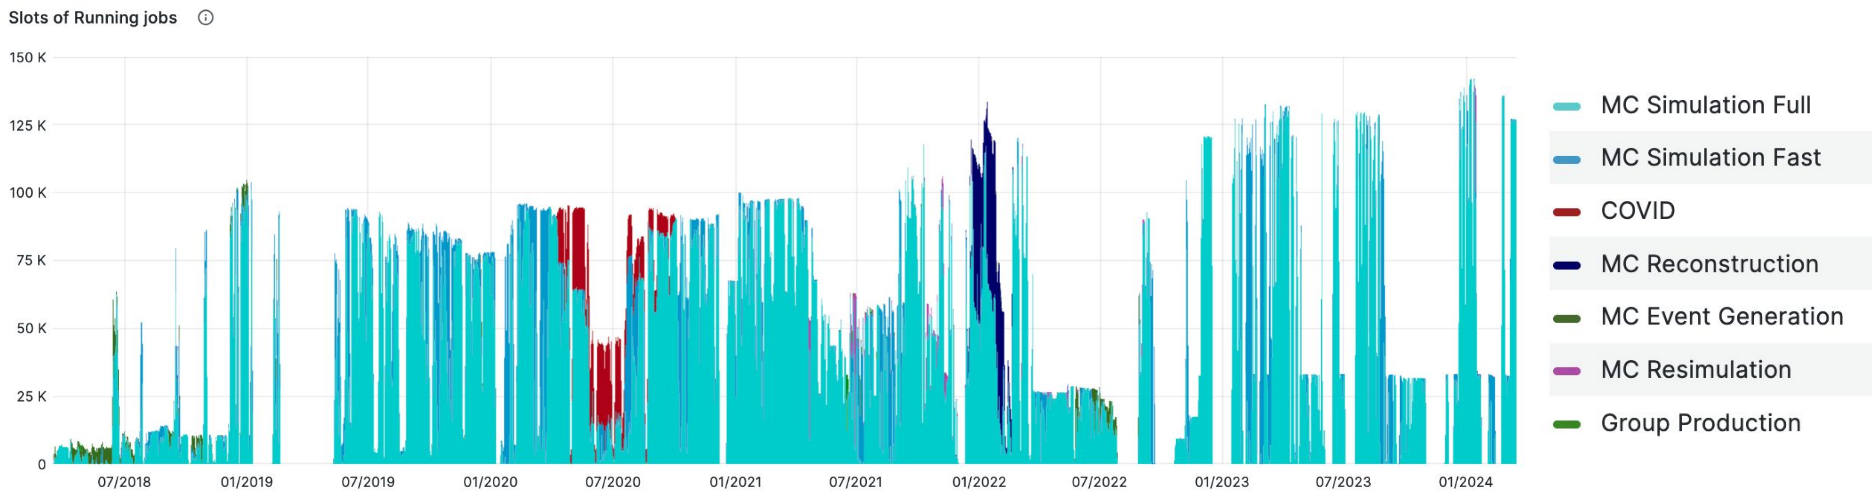

- Huge resource, 2.3 MHS23, successfully employed for production outside data taking
  - Initially only for FullSim and FastSim
  - Can also run Reconstruction as well as Folding@Home for COVID research
- Must be unplugged and removed at the end of Run 3
  - If we can find a new home, where we could move the hardware?
  - **Would be huge benefit to ATLAS**
  - Of course, will also reduce the pledge(burden) on what the Tier-2s need to provide

BNL, US-ATLAS, IRIS-HEP potential contributions

- Operations
  - Support for central DPA and DDM Operations.
    - Both suffer from the bus number / single point of failure fragility. And both are crucial
  - DPA seems easier because here we are actively looking for more timezone coverage
  - DDM Operations basically requires CERN presence
    - We are happy to see the plan for such a position coming to fruition

- Development
  - Injection of core development effort
    - General rejuvenation and modernisation across all our systems
  - DC24 has demonstrated transfer scheduling deficiencies
    - CS help very much appreciated
  - Smaller, but very important parts of ADC infrastructure
    - CRIC, HammerCloud, SAM/ETF, RPG, Catmore, Lifetime, ART, …

BNL, US-ATLAS, IRIS-HEP potential contributions

- Step R&D *"We have something and need to improve it"*
  - Back-pressure mechanisms to protect sites from overwhelming requests
  - Automating more operations of the data lifecycle management
  - Better HPC integration, especially with data ingress/egress

- Exploratory R&D *"Bluesky ideas for the next big thing"*
  - In agreement with HL-LHC roadmap and it's review process
  - Usually easier to get people excited (and funded) for this
  - However, high chance of being disruptive on the few remaining operators
  - This type of R&D should not be the main focus of the limited available effort
    - Priority must go to operations, development and Step R&D

Summary