ATLAS status for Run 3

Frédéric Derue, LPNHE Paris

LCG France meeting, CC-IN2P3 Lyon

30th november 2023

- **Run 3 program in 2024-2025**
- **Run PbPb**
- **LS3 program**
- **Token migration**
- **Heterogeneous resource evolution**
- **Site evolution**

Run 3 program (2024-2025)

- **<u>2025 preliminary resource requests</u> discussed with the C-RSG team (RRB)**

Stable beam fraction	Average <μ>	Stable beam seconds (pp)	Integrated luminosity (fb^{-1})	Main physics trigger rate (kHz)	Main physics stream events	Delayed trigger rate (kHz)	Delayed stream events
50%	62 (peak 65)	6.3×10^6	<120	2.3	14.5 billion	1.5	9.5 billion

Table 2: summary of the LHC operating programme in 2025 and the number of events that ATLAS expects to record.

- processing and reprocessing the Run 3 data taken in 2024 and 2025 using the newest release;
- Simulation of MC samples for 2025 (campaign starting in 2024);
- production of derived formats for data and MC (DAOD_PHYS and its skims as well as DAODs made directly from AOD)
- User analysis on Run 2 and Run 3 data and MC samples
- continuation of physics studies for the HL-LHC phase; and
- processing heavy-ion data and MC

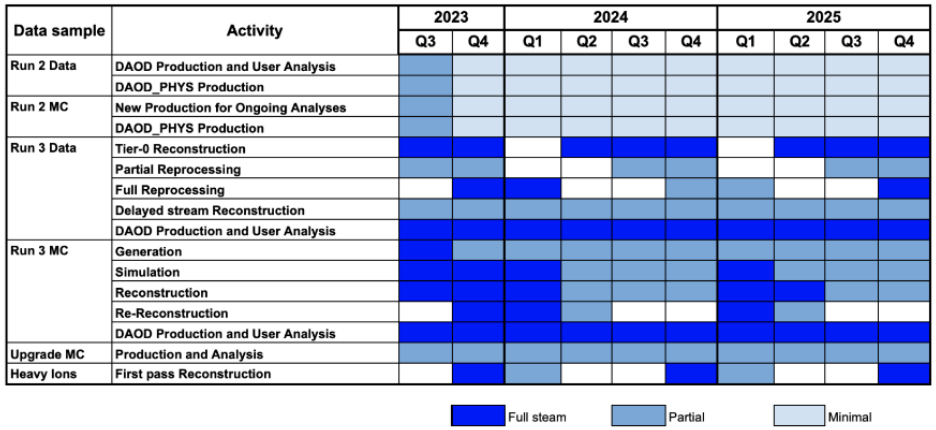

Figure 1: Plan for activities in 2023–2025. A darker shading indicates periods when it is expected that the activity will be most concentrated.

Run 3 program (2024-2025)

- **2025 preliminary resource requests discussed with the C-RSG team (RRB)**

Processing Times in 2025 (HS23s/Event)	
Full Simulation	2600
Fast Simulation	550
Generators	600
Data reconstruction	350
MC digitization+reconstruction	205+735
DAOD production (remnants of DAOD plus DAOD_PHYS and PHYSLITE)	30

Table 3: The main parameters concerning event sizes and processing time per event used to estimate the 2025 computing needs.

Event size (MB/Event) in 2025	
RAW data	1.8
MC HITS	0.8
Data AOD	0.35
MC AOD	0.55
MC RDO	2.0
DAOD_PHYS data	0.035
DAOD_PHYS MC	0.055
DAOD_PHYSLITE data	0.010
DAOD_PHYSLITE MC	0.012
DAOD fraction	0.40

Table 4: The main parameters concerning event sizes and processing time per event used to estimate the 2025 computing needs. The "DAOD fraction" is the total size of the DAODs (including formats supporting performance studies and unusual analyses, but not including PHYS or PHYSLITE) relative to the AOD. In Run 2 this fraction was above 1; in Run 3 it should be much lower thanks to the new analysis model and the use of DAOD_PHYS by most analyses.

The replication policy for 2025 is expected to follow that of previous years, as detailed in previous C-RSG reports.

- **RAW data size set to 1.8 MB (discussion on SuperCells still ongoing)**
  - several questions from C-RSG on the RAW size topic

Run 3 program (2024-2025)

- **2025 preliminary resource requests discussed with the C-RSG team (RRB)**

CPU requirements in 2025 Tier-1 and Tier-2		Disk requirements in 2025 Tier-1, Tier-2 and Tier-0		Tape requirements in 2025 Tier-0, Tier-1	
Activity	**kHS23*yr**	**Samples**	**PB**	**Samples**	**PB**
Event gen	666	HepMC events	6	T0 Run 1/2 Raw	57
MC full sim	1085		14	T0 Run 3 Raw	110
MC fast sim	407	MC Hits	1	T0 ESD/AOD/ DRAW/DESD	64.6
MC reco	1749	MC AOD	74	T0 HeavyIon Raw/AOD	31.1
MC derivation	548	MC DAOD	112	T1 Run 1/2 Raw	56
Data reconstruction	833	Data AOD	33	T1 Run 3 Raw	110
Data derivation	465	Data DESD	7	T1 ESD/DESD	3.8
User analysis	518	Data DAOD	58	T1 data AOD	57.2
Heavy ions	401	Heavy ions	13	T1 MC AOD	134.1
		Raw data (ops)	4	T1 RDO	19.8
		Others (logs, HISTS, etc)	0	T1 Hits	89.6
		Secondary copies	103	T1 Others	42.2
		Temporary copies	30	T1 HeavyIon Raw/AOD	54.3
		Tier-0 buffer	15		
TOTAL	**6672**	**TOTAL**	**469**	**TOTAL**	**829**

Table 5: A summary of the CPU and disk requirements at Tier-1 and Tier-2 sites, and tape at the Tier-0 and Tier-1 sites, for 2025, with a breakdown per activity or data sample.

Run 3 program (2024-2025)

- **2025 preliminary resource requests discussed with the C-RSG team (RRB)**

ATLAS Report to the C-RSG, September 2023

		2024 agreed @ April 2023 RRB	2025 Prel. Request @ September 2023 RRB	Balance 2025 wrt 2024
CPU	T0 (kHS23)	936	1100	17.5%
CPU	T1 (kHS23)	1516	1661	9.6%
CPU	T2 (kHS23)	1852	2030	9.6%
CPU	SUM (kHS23)	4304	4791	11.3%
Disk	T0 (PB)	49	56	14.3%
Disk	T1 (PB)	163	186	14.1%
Disk	T2 (PB)	200	227	13.6%
Disk	SUM (PB)	412	469	13.9%
Tape	T0 (PB)	207	264	27.4%
Tape	T1 (PB)	452	567	25.4%
Tape	SUM (PB)	659	831	26.1%

Table 6: Summary of the preliminary requests for computing resources in 2025.

"Old" RR for 2024 (approved)

ATLAS Report to the C-RSG, April 2023

		2023 agreed @ April 2022 RRB	2024 Request @ September 2022 RRB	2024 Request @ April 2023 RRB	Balance 2024 wrt 2023
CPU	T0 (kHS06)	740	850	936	26.5%
CPU	T1 (kHS06)	1430	1501	1516	6.0%
CPU	T2 (kHS06)	1747	1834	1852	6.0%
CPU	SUM (kHS06)	3917	4185	4304	9.9%
Disk	T0 (PB)	40	46	49	22.5%
Disk	T1 (PB)	136	162	163	20.2%
Disk	T2 (PB)	168	198	200	18.9%
Disk	SUM (PB)	344	405	412	19.8%
Tape	T0 (PB)	174	205	207	19.0%
Tape	T1 (PB)	353	448	452	28.1%
Tape	SUM (PB)	527	653	659	25.1%

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

Run 3 PbPb program

- **<u>2025 preliminary resource requests</u> discussed with the C-RSG team (RRB)**

3.5 Run-3 Heavy Ion data processing

According to the LPC plans, a heavy ion run will take place at the end of the 2025 data taking period. The throughput to storage is anticipated to be comparable to that of the p-p running, around 7–8 GB/s, leading to a total (RAW and derived products) anticipated data volume of 10 PB.

As in previous years, a fraction of heavy-ion data will be reconstructed at the Tier-0. Given the complexity of the events, it will not be possible to reconstruct all data within 24 hours at the Tier-0. Some data will be processed on the grid in spill-over mode, and others will be processed with a delay. In 2018 most heavy-ion data were processed for the first time within two weeks of data-taking; only test runs were taken in 2022.

HL-LHC network needs

- **DC21 planning**

T1 Sites (T0 export / T1→T2 reco)	HL-LHC Minimal Scenario [Gbps]	HL-LHC Flexible Scenario [Gbps]	DC27 (100%) [Gbps]	DC26 ~~(60→50%)~~ [Gbps]	DC24 (25%) [Gbps]	DC24 ATLAS [Gbps]	DC24 CMS [Gbps]	DC24 Alice [Gbps]	DC24 LHCb [Gbps]	~~DC23 (30%) [Gbps]~~	DC21 (10%) [Gbps]
CA-TRIUMF	200	400	100	60	30	30	0	0	0	~~30~~	10
DE-KIT	600	1200	300	180	80	32	26	11	11	~~90~~	30
ES-PIC	200	400	100	60	30	13	13	0	3	~~30~~	10
FR-CCIN2P3	570	1140	290	170	70	33	21	7	9	~~90~~	30
IT-INFN-CNAF	690	1380	350	210	90	24	35	14	16	~~100~~	30
KR-KISTI-GSDC	50	100	30	20	10	0	0	10	0	~~10~~	0
NDGF	140	280	70	40	20	16	0	4	0	~~20~~	10
NL-T1	180	360	90	50	20	15	0	1	4	~~30~~	10
NRC-KI-T1	120	240	60	40	20	8	0	8	4	~~20~~	10
UK-T1-RAL	610	1220	310	180	80	39	21	1	18	~~90~~	30
RU-JINR-T1	200	400	100	60	30	0	30	0	0	~~30~~	10
US-T1-BNL	450	900	230	140	60	60	0	0	0	~~70~~	20
US-FNAL-CMS	800	1600	400	240	100	0	100	0	0	~~120~~	40
(transatlantic link)	1250	2500	630	380	160	60	100	0	0	~~190~~	60
Sum	**4810**	**9620**	**2430**	**1450**	**640**	**270**	**246**	**56**	**65**	~~**730**~~	**240**

Planning for 2026 and beyond

- **LHCC met in September; report only recently available**
  - starting to think about 2026 and beyond

> • The year 2026, the first of the LS3, is expected to introduce sizable changes in the proportions each medium is needed by each experiment. The **LHCC appreciates** having a coarse projection of the computing resource needs of each experiment over the next years.

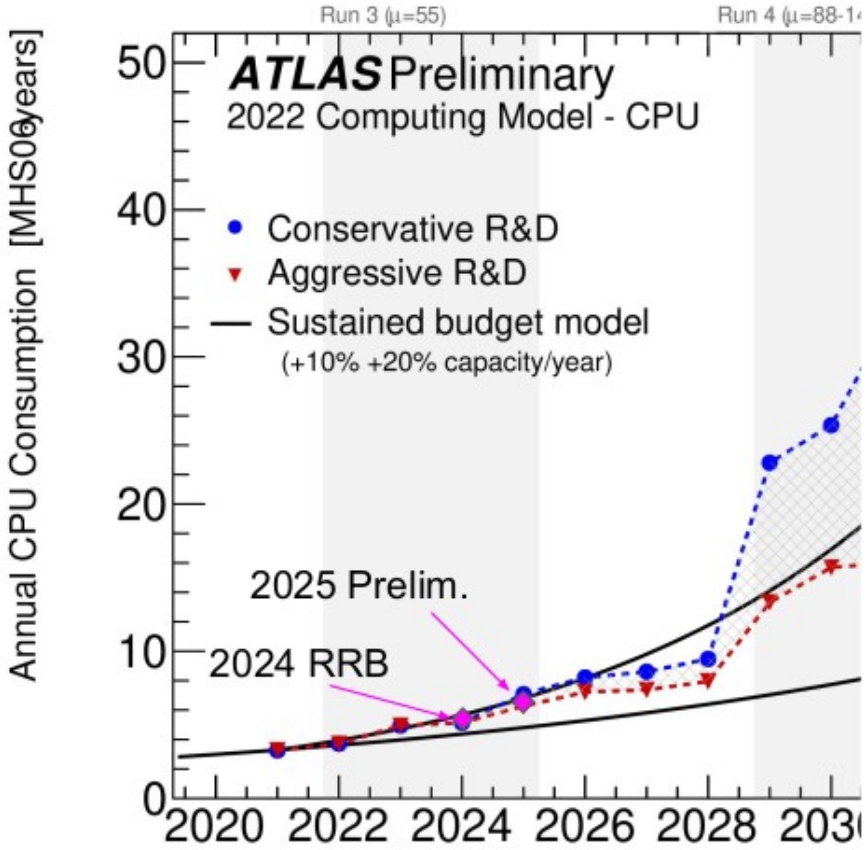

2025 Prelim.: **6672 kHS23**
HLT contrib.: 600 kHS23
HPC etc.: 1300 kHS23
WLCG sites: 3691 kHS23
(Rest is Tier0 contributions)

2026:
HLT contribution → Zero (unplugged)
→ 16% increase in WLCG site requirements assuming *no* change to MC, *no* additional data processing, etc

(The plan included increases, which is why you see the point go up in 2026)

Planning for 2026 and beyond

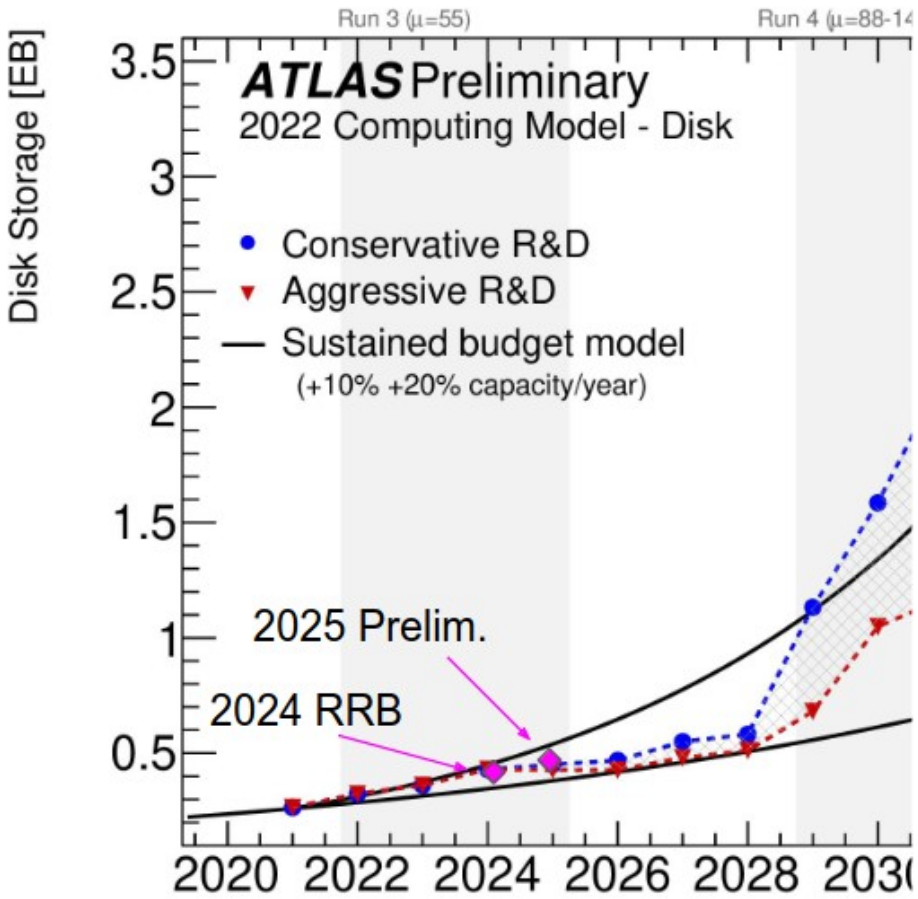

2025 Prelim.: 469 PB
Tier0 contrib.: 56 PB
WLCG contrib.: 410 PB

2028: 550–610 PB
2029: 680–1220 PB
2030: 1050–1630 PB

23–100% increase in 2029
55–34% increase in 2030

Move to tokens

Dimitrios Christidis at ATLAS S&C meeting Oct. 2023 (link)

- **Big concern is the move to tokens as replacement for X.509 authentication**
  - touches everything, and there is a very small pool of experts
  - in parallel there is the change to IAM (Identity and Access Management) as a replacement for VOMS
  - concrete initiatives now ("Rucio SIG") with objectives and milestones
    - expected end-date March 2026, in line with the WLCG Token Transition timeline
    - Mattermost channel: #tokens
  - for communities depending on X.509, the transition must be allowed to happen progressively
    - an RSE attribute for most workflows
      - for TPC, must be all RSEs related to a transfer job
      - won't support a situation where an RSE supports only tokens while others use X.509
    - account identities for user authentication
  - will be running in hybrid mode for Data Challenge in February

VOMS to IAM migration by June 2024

Petr Vokáč at ATLAS IT meeting Oct. 2023 (link)

- **CERN provided date for legacy VOMS retirement** (WLCGTOKENS-117)
  - originally thought it would be possible to move already in 2022
  - questionable CERN IAM support in the past, long IAM release/deployment cycle
- **Concerns: still no clear (published / presented) timeline by CERN IT**
  - still few open questions and work that needs to be done
  - migration CERN IAM to Kuberentes cluster & new LSC deployment
    - stress tests of new infrastructure for IAM VOMS AA
    - deployment of LSC for second / fallback hostname for IAM VOMS AA service (RAL IAM Hackathon notes)
  - waiting for IAM update that allows us to synchronize also AUP from legacy VOMS
    - scheduled 1.8.2 OTG0144943 should be followed soon by 1.8.3 that finally supports AUP updates via API
  - move user management from legacy VOMS to IAM interface
    - use SCIM IAM API instead of VOMS Admin API in production
  - turning off legacy VOMS servers – several steps (remove vomses, LSC, turning off VOMS, …)
  - discover & fix missing features in IAM
    - we need faster develop & test & deploy cycle (currently ~ ½ year)
  - Data Challenges 2024 in February, CentOS7 EOL in June (same deadline as VOMS to IAM migration)
    - experts might be busy... clear timeline for VOMS to IAM migration would help to avoid conflicts with other projects

Heterogeneous resource evolution

Mario Lassnig at ATLAS S&C meeting Oct. 2023 (link)

- **Resource evolution**
  - ARM looking good (PhysVal for both sim and reco) - Progress at scale validation request
    - if successful, can/will embrace ARM for production - What about user analysis?
    - we expect a reasonable amount of sites to deploy ARM due to expected energy saving
  - GPU resource need currently limited to 10s of GPUs: training, user analysis - not used in production
    - however, if we progress like CMS (10% of reco task on GPU) may not have access to enough
    - FastCaloSimGPU large scale PhysVal (code is ready and queues exist), expected until Q2 2024
  - HPC integration - There's always challenges and new developments needed (e.g. Goettingen)
    - still believe we will have x86, however, we fully expect future HPCs to have ARM support (e.g., TACC and EuroHPC)
    - onboarding of new HPCs will become part of regular ADC operation
    - also, RISC-V is coming
  - ISA deprecation
    - moved to x64-v2
    - dropping 5% of resources for improvements from v3 doesn't seem to be coming soon, ATEAM-905

Site evolution

Mario Lassnig at ATLAS S&C meeting Oct. 2023 (link)

- **Site evolution**
  - recent examples of evolution like the Taiwan T1→T2 or NET2 reorganisation
    - necessary, but can cause disruptions on a small pool of experts
  - several ATLAS clouds considering paradigm shift on how they provide resources
    - consolidation of Tier-2s, emergence of more HPCs
      - Is the corresponding support there ?
    - may incur updates on the computing model
      - integration of commercial clouds, not necessarily grid-like behaviour
      - Tier-2s with non-pledged tape
  - outcome of Google Total Cost of Ownershipe (TCO) may influence site pledging/purchasing
  - emphasis on user activities
    - Analysis Facilities are for sure an opportunity, but not at the expense of production
    - do we know what we truly need for an AF ?
      We are definitely further on than before
      - user and facility input is crucial
      - new ADC Physics Analysis coordinator
    - most likely will provide some sort of guidance/document to Physics Coordination

Eduardo Bach at ATLAS S&C meeting Oct. 2023 (link)