CERN–RRB–2025–009
1st April 2025

Computing Resources Scrutiny Group

R Aaij (The Netherlands), C Allton (UK), J Amundson (US), D Elia (Italy), E Fede (France), D Gingrich (Canada), J M Hernández (Spain), J Kleist (Nordic countries), H Meinhard (CERN, acting scientific secretary), T Mkrtchyan (Germany), M Schulz (CERN)

1 Introduction

The Computing Resources Scrutiny Group (C-RSG) is responsible for evaluating the computing needs of the Large Hadron Collider (LHC) experiments, ALICE, ATLAS, CMS and LHCb, and providing recommendations to the Computing Resources Review Board (C-RRB) on the Worldwide LHC Computing Grid (WLCG) resources required by the Collaborations to support their approved physics programmes. The C-RSG has completed its Spring 2025 review, which evaluated computing resource usage in the calendar year 2024, projected needs for 2025, and the computing resource requests for 2026. This report summarizes the findings of this review and the recommendations for the experiments.

The C-RSG requested reports from each Collaboration detailing their computing resource utilization in 2024, anticipated usage in 2025, and proposed resource needs for 2026. These reports also addressed the recommendations issued following the most recent scrutiny in Autumn 2024 [1].

After reviewing the reports, the C-RSG submitted questions to each Collaboration to clarify specific points. Dedicated meetings were then held with Collaboration representatives, providing an opportunity for further discussion. These face-to-face meetings allowed the C-RSG to gain a well-informed understanding of the experiments' computing requirements.

The C-RSG also had an informal meeting with the LHCC WLCG referees on January 29th, 2025. This meeting allowed the LHCC and the C-RSG to confer on areas of common interest related to computer resources and utilization. Although the mandates of the two groups are different, the overlap in the evaluation and oversight of computing resources made for a helpful and informative discussion. The C-RSG thanks the LHCC referees for their participation in this discussion.

The C-RSG recommendations for 2026 resource provisioning by WLCG are provided in this report. They are based on the updated LHC Run 3 schedule for 2025 and 2026 [2].

2 C-RSG Membership

The Spring 2026 scrutiny marked the first review led by J M Hernández (Spain) as chair of the C-RSG. The members extend their sincere appreciation to the previous chair, P Sinervo (Canada), for years of dedicated leadership and successful guidance of the C-RSG. They also thank J M Hernández for stepping into this role and leading the group forward.

Just before this scrutiny round, the U.S. Department of Energy (DoE) nominated J Amundson (FNAL) as the successor to A. Connolly (University of Washington) as the U.S.-proposed member of the C-RSG. We extend our gratitude to the DoE for the nomination, to J Amundson for agreeing to serve, and to A Connolly for years of dedicated and highly valued contributions to the C-RSG and WLCG.

Following established practices, the group has invited J Amundson to actively participate in this scrutiny. We now invite the RRB to formally confirm his nomination.

The chair thanks the C-RSG members for their commitment and expert advice.

The group thanks H Meinhard, who served as the scientific secretary of the C-RSG from 2011 to 2022, for stepping in to replace A Valassi during this scrutiny round due to his unavailability. His support was instrumental in ensuring the smooth operation of the group's activities. The group also thanks CERN management for its continued support of the C-RSG's work.

The group sincerely appreciates the efforts of the Collaboration representatives in supporting the scrutiny process and addressing its previous recommendations.

3 Interactions with the Experiment Collaborations

All Collaborations submitted their reports by 18$^{\text{th}}$ February 2025. The C-RSG thanks the Collaborations for the timely submission of their detailed documents [3–7], which also contained responses to the findings and recommendations from the Autumn 2024 scrutiny round [1]. The group thanks the computing representatives of the Collaborations for their availability, their constructive responses to the questions raised by the C-RSG and responses to subsequent requests for further information. The dedicated meetings with experiment representatives were particularly helpful and greatly appreciated by the C-RSG.

Dedicated teams of C-RSG referees were assigned to review the ALICE and LHCb requests. As in previous years, and in agreement with the ATLAS and CMS managements, a single team of referees scrutinized the ATLAS and CMS reports and requests to ensure a consistent evaluation. The referees then presented their findings to the full C-RSG, which collectively developed the observations and recommendations outlined in this report.

In preparation for the Autumn 2025 scrutiny, the C-RSG requests that the experiments submit their documents by Tuesday, August 26, 2025. While this deadline poses a challenge due to summer schedules, it is necessary to ensure the timely submission of the C-RSG report to the C-RRB by October 13, 2025.

The C-RSG asks that the reports include the preliminary computing resource requirements for 2027 and that the Collaborations address both the general recommendations and those specific to their experiments as part of their submission.

4 Resource Usage in Calendar Year 2024

The C-RSG requested that the Collaborations provide a summary of their computing resource usage from January to December 2024. The 2024 data-taking year was highly successful for all experiments, with the LHC delivering record-breaking numbers of proton-proton (p-p) and lead-lead (Pb-Pb) collisions in a single year. The p-p running period was extended by four weeks, which were brought forward from the 2025 schedule. Notably, 2024 marked the first Run 3 data-taking period in which LHCb operated under the expected experimental conditions and fully deployed its Run 3 computing model. In 2024, ALICE collected a lead-lead dataset slightly larger than that of 2023, but within a significantly smaller number of allocated data-taking days. The detector and data acquisition system operated at sustained interaction rates of 50 kHz for over an hour across multiple fills, demonstrating remarkable stability. The ATLAS and CMS experiments greatly expanded the 13 TeV dataset collecting data with excellent efficiency and quality.

Table 1 summarizes the RRB's approved baseline resource levels for 2024, categorized by tier level and resource type. The approved CPU power is approximately 11 MHS23, while disk capacity exceeds one exabyte, and tape capacity approaches two exabytes.

	CPU [kHS23]	Disk [PB]	Tape [PB]	CPU [%]	Disk [%]	Tape [%]
Tier0	2690	201	825	25%	19%	43%
Tier1	3648	414	1072	34%	40%	57%
Tier2	4421	431		41%	41%	
Total	10759	1045	1897	100%	100%	100%

Table 1 Baseline computing resource capacities approved by the RRB for WLCG sites in 2024. The last three columns represent the share of resources across the Tier levels.

Table 2 shows the balance (in percentage) between the RRB's approved CPU, disk and tape resources and those pledged by various funding agencies for each experiment, broken down by tier level.

With the exception of the Tier0, which provided exactly the approved resource levels, experiments faced resource shortfalls at the Tier1 and Tier2 levels, particularly affecting storage. Additionally, the 2024 LHC schedule was extended by four extra weeks of data-taking, moved forward from 2025. The resulting increase in data volume, which was not accounted for in the computing resource requests established in April 2023, placed further strain on disk and tape storage.

		ATLAS	CMS	ALICE	LHCb
Tier0	**CPU**	0	0	0	0
	Disk	0	0	0	0
	Tape	0	0	0	0
Tier1	**CPU**	0	10	-14	21
	Disk	0	-5	-13	-6
	Tape	2	-7	-4	0
Tier2	**CPU**	12	-5	-3	11
	Disk	-3	-8	2	-33

Table 2 Balance (in percentage) between pledged and RRB-approved computing resources for 2024. Numbers in red highlight a shortfall in pledged resources, while numbers in green indicate a surplus.

To mitigate this impact, experiments implemented various measures, including deletion campaigns to remove infrequently accessed data from disk and obsolete data from tape. Additionally, ALICE and LHCb negotiated an early allocation of a fraction of the 2025 disk resources.

The Collaborations efficiently utilized their allocated computing resources for processing, simulation and analysis of both Run 2 and Run 3 data. In 2024, total CPU usage reached 19.6 MHS23 years, marking a 16% increase compared to 2023. Experiments effectively leveraged CPU resources beyond pledged capacities, including idle time on High Level Trigger (HLT) farms, additional availability at WLCG sites and allocations at High Performance Computing (HPC) facilities. WLCG sites accounted for 78% of the total CPU consumption, with HLT farms contributing 7% and HPC facilities providing the remaining 15%. Overall, the CPU usage exceeded the baseline capacity approved by the RRB at WLCG sites by 82%, with WLCG sites alone providing 42% more CPU than the baseline level. It is worth noting that the expected CPU contribution from HLT farms was incorporated into the resource planning, allowing for a corresponding reduction in the CPU requested from WLCG sites.

Table 3 shows the CPU used by each of the experiments at the different tier levels, compared with the RRB's approved baseline values. All experiments enjoyed substantial over-pledge CPU capacity at the WLCG sites.

Figure 4 illustrates the annual CPU utilization of the experiments since 2017, highlighting individual contributions from WLCG sites, HLT farms and HPC facilities. The dashed line represents the RRB's

	Tier0			Tier1			Tier2		
	Used	RRB	Used/RRB	Used	RRB	Used/RRB	Used	RRB	Used/RRB
ATLAS	858	936	0.92	1642	1516	1.08	3555	1852	1.92
CMS	1510	980	1.54	1406	930	1.51	2247	1600	1.40
ALICE	1109	600	1.85	452	630	0.72	571	650	0.88
LHCb	468	174	2.69	952	572	1.66	514	319	1.61
Total	3945	2690	1.47	4452	3648	1.22	6887	4421	1.56

Table 3 CPU usage (in units of kHS23 years) by experiments across different tier levels in 2024, compared with the RRB's approved baseline values.

approved baseline CPU capacity for WLCG sites. The significant contributions of CPU from the experiments' online farms and HPC facilities are evident, along with the over-pledged CPU provided by the WLCG sites.

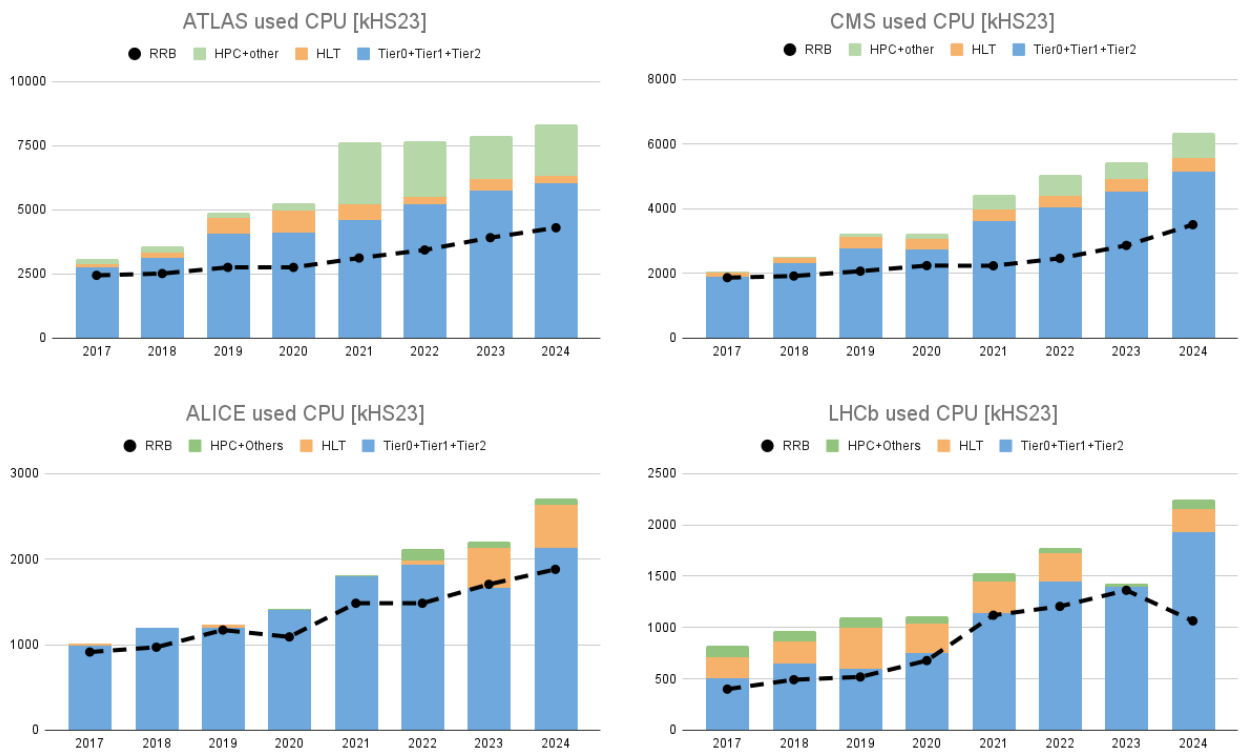

Figure 1 Annual CPU utilization of the experiments (in units of kHS23 years) with contributions from WLCG sites, HLT farms and HPC facilities stacked. The dashed line represents the RRB's approved CPU baseline capacities for WLCG sites.

Experiments utilized most of the non-pledged CPU resources to extend their Monte Carlo (MC) simulations, enhancing statistical precision and reducing systematic uncertainties in various physics analyses. The CMS experiment used part of these resources to promptly reconstruct parking data streams before archiving them to tape. The extension of simulation datasets might have a significant impact on pledged storage resources since opportunistic storage resources do not exist. To assess this impact, we provide several recommendations in Section 14.

Disk storage capacity grew by 18% from 886 PB in 2023 to 1045 PB in 2024, while tape storage increased from 1532 PB to 1897 PB during the same period. Both Tier0 and Tier1 tape storage increased by 24%, reflecting the increase in Run-3 recorded raw data. Disk and tape utilization are adequate when considering the headroom required to effectively manage large data volumes. At the end of 2024, ALICE's tape utilization remained low (68% at Tier0 and 50% at Tier1). This was primarily due to the significantly smaller-than-expected data collection in 2023 and the scheduling of the 2024 Pb-Pb data-taking period at the end of the year, resulting in the collected data not yet being archived.

Disk and tape storage were limited to the pledged resources at WLCG sites. As previously mentioned, the four-week extension of the 2024 data-taking period, the storage under-pledges faced by the experiments and the increased MC production driven by non-pledged CPU resources all placed significant strain on the storage capacity and the management efforts of the Collaborations. Efficient use of disk resources is even more critical in the context of constrained availability. Removing long-unaccessed data from disk and optimizing data replication policies will help improve disk utilization. We provide several recommendations in Section 14 to address these challenges.

The ATLAS and CMS experiments record additional data streams, known as parked or delayed data, which are not immediately reconstructed offline by the Tier0 system but are processed at a later stage. These streams are justified as a means to maximize the physics potential of analyses that benefit from increased data collection rates. The storage, delayed processing and simulation of these data consume a significant share of available resources, particularly storage, which is a highly constrained resource. In this regard, we have issued a recommendation for CMS and ATLAS to quantify the resources allocated to parked data, including both current usage and projected requirements for 2025 and 2026.

The Collaborations have made significant efforts to port their software stacks to the ARM architecture and are now essentially ready to run production workflows on ARM resources as they become available. They have also begun accepting CPU pledges based on ARM processors. A common request from the experiments is an improved ARM infrastructure at CERN to facilitate faster software building, validation and debugging for the ARM architecture.

The Collaborations continue their efforts to adapt their experimental applications for execution on GPUs, which offer a better performance-to-cost ratio compared to standard CPUs. Additionally, GPUs are becoming increasingly prevalent in large HPC facilities, making their adoption essential for leveraging these potentially vast resources.

However, not all algorithms can be efficiently executed on GPUs, making a hybrid approach (applications run on CPUs while offloading specific workloads to GPUs) the best match. Refactoring software for GPUs is labour-intensive, necessitating careful cost-benefit analyses.

As GPU usage becomes more widespread, proper benchmarking will be required to assess GPU performance in terms of the standard HS23 metric, ensuring accurate accounting of GPU resource utilization.

In recent years, experiments have successfully leveraged large CPU allocations from HPC facilities thanks to significant efforts in evolving workload management systems to integrate these resources.

At the end of January 2024, the HEP and HPC communities convened at CERN for a workshop to discuss strategic and technical aspects of HPC resource utilization. Key topics included the need for strategic access beyond traditional allocation requests and the development of common interfaces to facilitate seamless and uniform integration of HPCs into the WLCG distributed computing infrastructure. The recognition of high energy physics as a strategic science in EuroHPC will be pursued. In parallel, key projects aimed at developing common interfaces for seamless HPC integration have been launched in both the U.S. and Europe.

The C-RSG strongly supports these initiatives, recognizing the substantial potential of accessing vast CPU resources to enhance computing capabilities.

A more detailed description of the computing resource utilization by the experiments in 2024 is provided in the following sections.

ALICE		2024				
		RRB approved	Pledged	Pledged / RRB	Used	Used / RRB
CPU	Tier0	600	600	100%	1109	185%
	Tier1	630	540	86%	452	72%
	Tier2	650	630	97%	571	88%
	Total	**1880**	**1770**	**94%**	**2132**	**113%**
	*EPN**				*501*	*27%*
	*Others**				*76*	*4%*
Disk	Tier0	67.5	67.5	100%	62.4	92%
	Tier1	71.5	61.9	87%	53.1	74%
	Tier2	66.5	68.0	102%	60.1	90%
	Total	**205.5**	**197.4**	**96%**	**175.6**	**85%**
Tape	Tier0	181.0	181.0	100%	123.0	68%
	Tier1	107.0	102.4	96%	54.0	50%
	Total	**288.0**	**283.4**	**98%**	**177.0**	**61%**

*: Percentage taken with respect to the total RRB approved CPU recommendation

Table 4 Summary of ALICE resource usage in 2024 and the comparison with RRB-approved and pledged resources. The CPU and storage resources are in units of kHS23 years and PB, respectively.

5 Resource Usage: ALICE

The report on the ALICE Collaboration's use of computing resources is based on the usage data and resource request provided by the ALICE experiment [3], a written set of responses to the C-RSG scrutiny questions, and a face-to-face meeting with the ALICE computing and physics coordinators. The ALICE report details the computing resources used from January to December 2024 and is summarised in Table 4.

5.1 Data Taking and Processing

ALICE data taking during 2024 can be divided into three distinct periods: two p-p runs at 13.6 TeV and 5.36 TeV and one Pb-Pb run at 5.36 TeV. During the p-p run at 13.6 TeV, from 5 April on the experiment started data taking at 650 kHz. On 19 April, the runs for physics started and ALICE recorded physics data with all detectors. Up to 16 October more than 53 pb^{-1} have been recorded. ALICE efficiency has been at 92%. The modified schedule added four weeks of p-p running in 2024. The C-RSG notes that this was not taken into account in the resource request in 2023. Resources pledged for 2025 had to be allocated already in 2024. This affected especially disk space at the Tier0 and Tier1 level. At each tier, 2 PB of additional disk space had been required. In addition, 22 PB of tape had to be moved forward. In total, 204 PB have been collected. In line with the computing model, the data have been skimmed by a factor of 4.5‰, and the disk space has been freed.

During the p-p reference run at 5.36 TeV in late October, ALICE reached the Run 3 integrated luminosity goals and no further reference data need to be collected for the Run 3. 14 PB have been recorded on the O2 buffer. This data is needed as a reference for the Pb-Pb run and will be kept without skimming.

On 12 November, ALICE started recording Pb-Pb data at 50 kHz. Due to the improvements of the machine, ALICE accumulated 1.54 nb^{-1} of physics data before the run ended prematurely on November 23. More data than in 2023 could be recorded, despite the shortness of the run. Before calibration, the efficiency was better than 90%. A volume of 39.2 PB have been recorded, corresponding to more than 12 million events.

RAW data reconstruction, skimming and validation of the 2022 and 2023 pp data have been fully completed. The original p-p compressed trigger format (CTF) files for 2022 and 2023 have been removed from the disk buffer and are now archived on tape, with one copy at the Tier0 and one at the Tier1s. The luminosity recorded in 2024 was also processed and the 204 PB of data have been reduced to 9.2 PB. The selection and skimming were done at the Tier0 accessing the O2 disks.

For the 2023 Pb-Pb data, 20% of the data have been processed three times for optimisation, followed by a pass on all data. During the run in November 2024, 10% of the data have been used for a designated calibration pass. This led to the first production pass in which periods of stronger TPC distortions have been excluded, to be reconstructed in the future when sufficiently good calibration is available. This first pass came to completion on 13 January 2025.

In parallel to the reconstruction and calibration passes, MC productions anchored to 2022 and 2023 conditions took place during the year; those based on 2023 are still ongoing.

5.2 Resource Utilisation

The reported usage of tape, disk and CPU resources does not include the Russian sites that were excluded from WLCG in November 2024 with the exception of JINR. However, ALICE chose to use these resources as opportunistic ones until November 2024; all data have been replicated. This corresponds to 1.5 PB of disk space and 25 kHS23 of CPU. In addition to the resources at JINR, ALICE lost 5.7 PB of tape, 4.5 PB of disk and 32.8 kHS23 CPU capacity at the Tier1, and 1.5 PB disk and 12 kHS23 CPU at the Russian Tier2s.

The C-RSG noticed that the report did not contain a roadmap towards the implementation of Strategy B for data compression, which would result in a reduction of storage needs. ALICE explained that the move to taking data at 50 kHz changed the conditions drastically. The level of understanding of the TPC calibration is currently not good enough to move to lossy compression. The experiment is tracking the calibration progress closely, but can not give a concrete date at which this issue will be resolved.

Tape Usage The tape usage at Tier0 and Tier1, with 123 PB and 54 PB, respectively, have been at the 70% and 50% level of the RRB-approved and pledged resources. A data volume of 40.8 PB from 2023 has been removed at the Tier0, and the archival of the 2022 and 2023 data has been completed. The 2024 p-p and Pb-Pb data files still need to be migrated to tape.

Disk Usage At the end of December 2024, a capacity of 208 PB has been deployed. This includes 2 PB of non-pledged resources at Wigner and the LBNL Analysis Facilities. While the Tier1s are 14% below the pledges, however, this gap has been reduced to 9% during the year by Tier1s providing storage beyond their pledge. In ALICE's report, the development of pledges, installed capacity and usage throughout the year at the different tiers are illustrated. At all levels, the installed storage exceeds the pledges and the pledges exceed the usage of disk storage. The gap is especially pronounced at the Tier2s. In total 84% of the RRB-approved disk space has been used.

Concerning disk usage, 35% of the capacity has been used for simulation, 56% for reconstruction and 9% for analysis related work.

The information provided on the data volume resident on disk vs the number of data accesses has been difficult to interpret. In the discussion with the C-RSG, ALICE explained that the data entering the statistics are those formats that dominate the used disk space and are mostly related to analysis and MC productions. ALICE further clarified that only a small part of the data provided carry meaningful information. Those are describing data that have not been accessed for 3, 6 or 12 months, either produced before this year or during the current period. The amount of these data has been on the order of 10 PB. In addition, it was noted that ALICE removes all data from disk that have not been accessed within one year. ALICE suggested that data popularity should be redefined by the C-RSG.

CPU Usage The table 4 contains the information related to the pledged CPU resources. At the Tier0, 185% of the pledged resources have been utilised, while at the Tier1 and Tier2 level, the utilisation has been 72% and 88%. While requested and pledged resources at the Tier1 and Tier2 level haven't been fully utilised, ALICE used significant CPU capacity at the Tier0 beyond the pledge and non-pledged resources, with the event processing nodes (EPN) contributing 501 kHS23 years and other centres 76 kHS23 years. The EPN has been used primarily for reconstruction tasks during non-data-taking periods and utilised also the GPUs. With the Lawrencium supercomputer at LBNL, ALICE demonstrated the remote reconstruction of Pb-Pb data with good results.

It was noted that CPU efficiency was below the values reported by other experiments. For the Tier0, it was 83%, for the Tier1 sites 69%, and for the Tier2s 60%. ALICE is aware of the problem, and upon request, provided data showing that over the last year, the efficiency of all workloads has improved, increasing from an average of 60% to 75%. The analysis workflow saw the greatest progress. The switch of the framework for organised analysis jobs from the legacy alitrain to the new alihyperloop improved the efficiency from about 55% to close to 80%.

Regarding the distribution of CPU usage among the various activities, 44% of the resources have been used for simulation, 36% for reconstruction and 20% for analysis.

Over the year, the usage of resources has not been continuous. Apart of changes due to the accelerator programme and the evolution of the analysis framework, there have been also periods, especially in April and December 2024, with low utilisation of resources. ALICE explained that these were periods when detector calibration and code were optimised for reprocessing campaigns. During significant calibration improvement campaigns, the production of MC data was also halted until these had been completed.

6 Resource Usage: ATLAS

The report of the ATLAS Collaboration's computing usage is based on the information provided by the ATLAS experiment [4], written responses to questions by the C-RSG, and a face-to-face meeting with the ATLAS computing coordinators.

6.1 Overview of Activities

In 2024, the ATLAS experiment successfully utilized its computing resources to process a record 117 fb^{-1} of p-p collision data delivered by the LHC. The year saw an increase in data processing efficiency and an expansion of MC simulations for both Run 2 and Run 3 analyses. The collaboration submitted 127 physics papers, reinforcing ATLAS's diverse physics programme.

Computing resources were engaged in supporting MC simulations, data reconstruction and user analyses. These resources included pledged allocations from WLCG sites and significant contributions from unpledged sources such as HPC facilities, the HLT farm and volunteer computing (BOINC).

CPU Usage The computing infrastructure of ATLAS was fully utilized to support its ambitious programme. MC simulation played a central role in resource usage, accounting for 35% of total CPU consumption, primarily through full Geant4-based detector simulations. MC event reconstruction followed, consuming 18% of CPU resources, while derivation production—used to generate analysis-ready data—used another 13%. Event generation workflows, especially those focused on B-physics samples, accounted for 14% of the overall usage. Although fewer total MC events were generated in 2024 compared to 2023, the simulations required more CPU time due to the higher complexity of the samples. User analysis also made up a significant portion of resource consumption, using about 9% of total CPU resources.

Total CPU usage reached 8.3 million HS23, with a CPU efficiency of 85.5%. The highest CPU efficiency (89%) was achieved at Tier1 and Tier2 sites. In addition to pledged CPU resources, unpledged sources contributed significantly to computing capacity. HPC centres such as the Vega EuroHPC system provided 26% of total MC simulation CPU resources. The HLT farm and commercial clouds contributed 11%, and volunteer computing through BOINC accounted for another 1%.

ATLAS successfully used all pledged CPU resources while effectively leveraging unpledged sources at Tier2 sites, which provided nearly twice the pledged CPU capacity.

At CERN's Tier0, 83% of CPU resources were dedicated to prompt and bulk data reconstruction, merging and file handling. The remaining 17% was allocated to detector calibration and alignment.

Data transfer rates from Point 1 to Tier0 were higher than in previous years, averaging well above 6 GB/s. The LHC's exceptional performance resulted in reduced downtime, limiting the availability of the HLT farm for offline processing compared to previous years.

Data Usage ATLAS fully utilized the available disk space, ensuring active datasets were accessible while managing a healthy cache for optimizing data access. Disk capacity at Tier1 and Tier2 sites remained nearly full throughout the year, with active datasets stored for ongoing processing and analysis. Regular deletion campaigns were performed to remove obsolete datasets.

Disk storage was dominated by analysis data (DAOD) and reconstructed events (AOD), reflecting the transition to the new streamlined analysis model. A breakdown of disk use shows that persistent data was balanced with temporary and cache data, with about 30% of disk space allocated for caching to facilitate data movement and access.

Tape Storage The usage of tape storage at the Tier0 increased significantly due to the large volume of 2024 collision data (64.2 PB), including heavy-ion collisions (5.8 PB) and a special low-energy p-p reference run (3.2 PB). Tier0 and Tier1 tape occupancy increased steadily, though the rapid growth, combined with an LHC schedule shift that moved some 2025 data-taking into 2024, strained available tape capacity. Tier1 tape storage increased by nearly 120 PB over the year, and several Tier1 sites approached or exceeded their pledged limits. In response, ATLAS initiated the use of non-pledged tape space at some Tier2 centres, such as DESY and NET2/NESE, to store MC samples, easing pressure on Tier1 sites.

Data Transfer Activities Data transfers across the grid remained robust and efficient. Production input transfers—from tape storage to compute sites—accounted for approximately half of all grid data movements. Additional transfers supported user analyses and the consolidation of datasets across the network. In February, the ATLAS computing system successfully participated in the 2024 WLCG Data Challenge, which stress-tested the infrastructure by simulating 25% of the network throughput expected at the HL-LHC. This challenge validated ATLAS's preparedness and exposed areas for further optimization.

ATLAS		2024				
		RRB approved	Pledged	Pledged / RRB	Used	Used / RRB
CPU	Tier0	936	936	100%	858	92%
	Tier1	1516	1514	100%	1642	108%
	Tier2	1852	2074	112%	3555	192%
	Total	**4304**	**4524**	**105%**	**6055**	**141%**
	*HLT**				*290*	*7%*
	Others*				*1982*	*46%*
Disk	Tier0	49.0	49.0	100%	43.4	89%
	Tier1	163.0	163.1	100%	172.0	106%
	Tier2	200.0	194.0	97%	179.0	90%
	Total	**412.0**	**406.1**	**99%**	**394.4**	**96%**
Tape	Tier0	207.0	207.0	100%	194.0	94%
	Tier1	452.0	460.0	102%	467.0	103%
	Total	**659.0**	**667.0**	**101%**	**661.0**	**100%**

*: Percentage taken with respect to the total RRB approved CPU recommendation

Table 5 Summary of planned and used resources for ATLAS in the calendar year 2024. The CPU and storage resources are in units of kHS23 years and PB, respectively.

Distributed Computing & HPC Integration ATLAS effectively managed its worldwide computing grid while integrating new architectures and resources. A highlight is that ARM architecture pledges are now accepted (with up to 50%) – expanding ATLAS computing flexibility. Additional Tier2 sites are planned in Greece, Morocco, Armenia, and UAE to expand computational capacity. Commercial cloud resources tested via Google Cloud, improving cloud-based workload execution.

Summary In 2024, ATLAS computing resources handled a record year of data taking, ensuring operation of Run 2 final analyses and Run 3 expansions. The year also reinforced the importance of alternative computing resources (HPC, ARM, and commercial cloud) to support evolving experimental needs.

7 Resource Usage: CMS

The report on the CMS Collaboration's usage is based on the information provided by the CMS experiment [5], written responses to questions by the C-RSG and a face-to-face meeting with the CMS computing coordinators.

The CMS collaboration reported successful data-taking operations in 2024, achieving record-high p-p and Pb-Pb collision rates at the LHC. CMS recorded 112.7 fb^{-1} of p-p collisions (92% efficiency) and 1668.26 μb^{-1} of Pb-Pb collisions (87.6% efficiency). The total RAW data volumes were 49.2 PB for p-p collisions and 9.4 PB for Pb-Pb collisions.

Significant advancements were achieved in computing performance and efficiency. CMS utilized GPUs for around 40% of the reconstruction sequence at HLT nodes, significantly optimizing resource use.

The Tier0 infrastructure efficiently handled higher-than-anticipated data rates, processing collision data promptly at input and output rates of up to 8.5 GB/s and 11 GB/s, respectively.

The MC simulation produced over 73 billion events, primarily for supporting Run 2 and Run 3 analyses and preliminary studies for the Phase-2 upgrade. Geant4-based simulations provided highly accurate data descriptions, with performance and physics modeling significantly enhanced by the Geant4 11.1 series integrated into the CMS software framework CMSSW.

Resource utilization at CERN averaged approximately at 68 745 CPU cores with a Tier0 average CPU efficiency increasing from 67% in 2023 to 77% in 2024. CERN's total disk capacity increased to 54 PB, peaking at 51.6 PB utilization by year-end. Tape storage usage reached 94% of the pledged capacity (320 PB).

At Tier1 sites, CPU resource utilization averaged 141% of the pledged capacity with a CPU efficiency of 78%, while Tier2 sites reported 153% CPU utilization with an efficiency of 67%. Both tiers reported strong performance, though Tier1 faced persistent storage shortfalls, highlighting the need for additional storage resources, for example those potentially provided by future prospective Tier1 sites.

CMS has resumed utilizing the Tier-1 storage resources at JINR—the second largest in terms of pledged capacity and utilization—after a period when these resources were considered opportunistic due to the political situation.

CMS is continuing the search for additional Tier1 sites. Serbia signed an MoU with WLCG to establish a CMS Tier1 site. The Serbian facility is expected to provide around 10% of the total CPU, disk and tape capacity, with its contribution ramping up from 2026.

In 2025, Poland plans to provide 5 PB of tape storage for CMS, with the potential to double this capacity. However, the amount of CPU and disk resources that can be pledged remains uncertain.

CMS continues to leverage opportunistic CPU resources significantly, including contributions from HPC centres and volunteer computing resources.

The Run 2 HLT farm was continuously used for offline processing, while the Run 3 HLT farm was used during periods without LHC beams. The availability of the HLT farms was instrumental to allow the prompt reconstruction of the parking data streams.

We note that a significant fraction of the disk resources is occupied by rarely accessed data, primarily analysis data in the AOD format. It appears there is potential to optimize data placement, enhancing disk utilization efficiency.

8 Resource Usage: LHCb

The usage of computing resources in the 2024 calendar year has been in general smooth for LHCb, with the notable exception of disk shortage due to lower pledged resources than requested ($\approx$ 10%) and data volumes higher than initially foreseen due to the one-month extension of the LHC p-p run. The disk shortage was covered by additional storage deployed in a number of Tier1 sites and lower-than-expected data rates out of the detector, namely 8.8 GB per LHC second instead of the estimated 10 GB/s. This lower data rates also implied that less tape than estimated was used.

As in previous years, the offline computing activities were dominated by MC production, close to 96% of the total CPU usage, followed by data analysis activities. The average CPU work required per simulated event is 2.06 kHS23 s, which is almost twice as much as in 2023. This increase is partially explained by the higher fraction of detailed simulated events (for every event, the overlaid pile-up collisions are simulated) for Run 1 and Run 2, but needs further investigation. It is expected to decrease in 2025 and 2026.

CMS		2024				
		RRB approved	Pledged	Pledged / RRB	Used	Used / RRB
CPU	Tier0	980	980	100%	1510	154%
	Tier1	930	1020	110%	1406	151%
	Tier2	1600	1526	95%	2247	140%
	Total	**3510**	**3526**	**100%**	**5162**	**147%**
	*HLT**				*409*	*12%*
	*Others**				*772*	*22%*
Disk	Tier0	54.0	54.0	100%	47.1	87%
	Tier1	122.0	115.7	95%	91.1	75%
	Tier2	149.0	136.7	92%	110.6	74%
	Total	**325**	**306.4**	**94%**	**248.8**	**77%**
Tape	Tier0	320.0	320.0	100%	300.7	94%
	Tier1	380.0	353.9	93%	279.9	74%
	Total	**700.0**	**673.9**	**96%**	**580.6**	**83%**

*: Percentage taken with respect to the total RRB approved CPU recommendation

Table 6 Summary of planned and used resources for CMS in the calendar year 2023. The CPU and storage resources are in units of kHS23 years and PB, respectively.

Table 7 shows an overview of the LHCb resource usage in 2024, along with a comparison of the amount of resources approved by the RRB and the resources pledged by the sites. The usage information is based on EGI accounting [8] provided by the WLCG Accounting Utility (WAU) portal [9].

WLCG sites provided 1934 kHS23 years of CPU work, 60% higher than the pledged CPU for LHCb at those sites. The contribution from non-WLCG sites, including HLT and HPC centres were 310 kHS23 years.

A multi-threaded version of the LHCb simulation is currently under development. This is expected to reduce RAM usage per core and significantly enhance the scalability of MareNostrum5 HPC utilization. At present, the scale of exploitation is constrained by the number of single-core simulation jobs that the Dirac servers can manage concurrently.

The disk storage usage of 99.2 PB in 2024 was a bit higher than foreseen by the pledge of 96.4 PB. Additional storage was deployed by a number of Tier1 sites as an advance on 2025 resources. The total tape usage of 206.4 PB was 18% lower than the 250.4 PB requested (and 20% lower than the 259.3 PB pledged) for 2024.

LHCb usage is based on EGI and DIRAC accounting reports, providing complementary resource usage perspectives. WLCG CPU resource accounting is obtained from the WAU portal, while additional non-WLCG CPU resources are accounted for by DIRAC. Storage accounting relies solely on information provided by DIRAC.

The additional 7.5 PB Tier0 tape buffer and 5.5 PB of user and detector studies and commissioning space are not visible in either DIRAC or WLCG storage accounting.

LHCb		2024				
		RRB approved	Pledged	Pledged / RRB	Used	Used / RRB
CPU	Tier0	174	174	100%	468	269%
	Tier1	572	692	121%	952	166%
	Tier2	319	356	112%	514	161%
	Total	**1065**	**1222**	**115%**	**1934**	**182%**
	*HLT**				*216*	*20%*
	*Others**				*94*	*9%*
Disk	Tier0	30.6	30.6	100%	27.2	89%
	Tier1	61.2	57.9	95%	66.7	109%
	Tier2	11.8	7.9	67%	5.3	45%
	Total	**103.6**	**96.4**	**93%**	**99.2**	**96%**
Tape	Tier0	117.1	117.1	100%	94.6	81%
	Tier1	133.3	142.2	107%	111.8	84%
	Total	**250.4**	**259.3**	**104%**	**206.4**	**82%**

*: Percentage taken with respect to the total RRB approved CPU recommendation

Table 7 Summary of LHCb resource usage in 2024 and the comparison with RRB-approved and pledged resources. CPU numbers are taken from the EGI and Dirac accounting portals. Disk and tape occupancy are taken from the WLCG storage space accounting and refer to 31st December 2024. The CPU and storage resources are in units of kHS23 years and PB, respectively.

9 Anticipated 2025 Resource Utilization and Background for 2026 Resource Requests

The final two years of Run 3, the 2025 and 2026 data-taking periods, hold the potential to substantially expand the accumulated datasets, further advancing the physics programmes toward even greater achievements.

The Collaborations have outlined comprehensive plans for computing activities in 2025 and have requested the required computing resources for 2026, taking into account the final LHC schedule (see Fig. 2), which notably extends the delivered luminosity to the experiments in 2025 and introduces an additional data-taking period running until the end of June 2026.

In 2025, the LHC is expected to provide p-p collisions for physics between May and October. The heavy-ion data-taking period will occur in November, lasting for three weeks. In 2026, p-p collisions will be provided between March and May, followed by three weeks of heavy-ion data taking in June. In 2025, most of the heavy-ion production days will be dedicated to Pb-Pb collisions, with some runs of O-O and Ne-Ne being planned. The heavy-ion plan in 2026 remains under discussion. No p-p reference runs are planned for 2025 or 2026, as the dataset collected in 2024 is considered sufficient for Run 3 analyses.

Given the extension of the 2025 data-taking period beyond the originally planned schedule at the time of approval of the computing resources for 2025, experiments have accounted for additional 2025 resource needs in their 2026 requests. Advancing a portion of the 2026 storage resources to late 2025 would help experiments manage the expected increase in data volume more effectively.

Similarly, with the early start of the 2026 data-taking period in March, ahead of the established 1 April

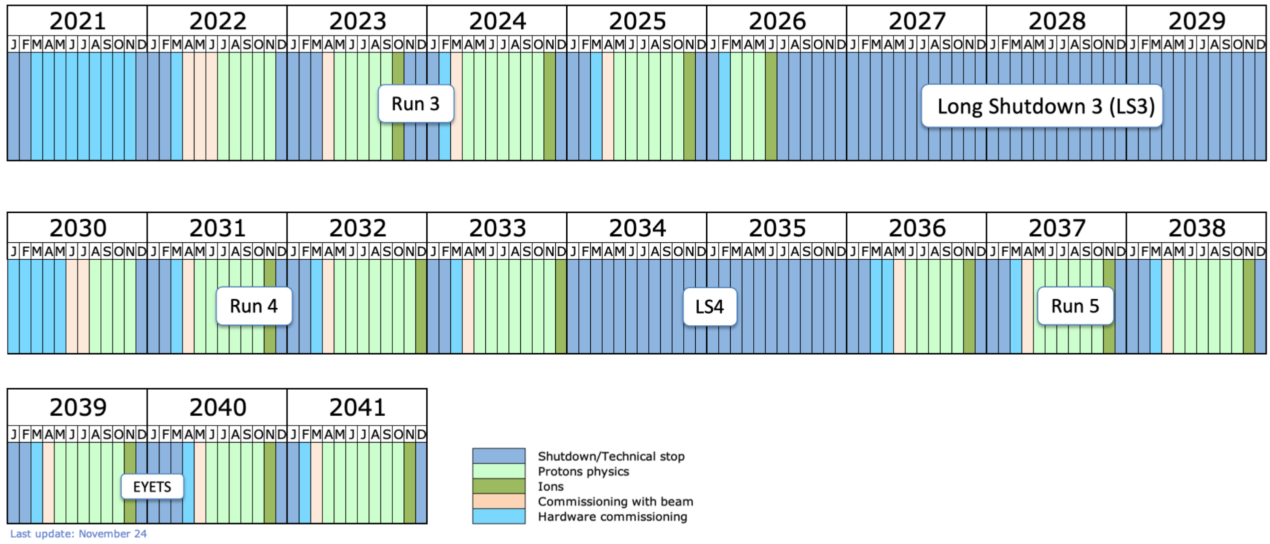

Figure 2 Long-term LHC schedule as revised in September 2024

date for the provisioning of 2026 computing resources, ensuring the ahead-of-time delivery of some storage resources by February will be crucial for the experiments.

When formulating their 2026 resource requests, experiments have accounted for the expected CPU availability from HLT farms when not used for online processing. These resources, which reduce the CPU requested from WLCG sites, have a significant impact on overall resource needs. Additionally, experiments plan to repurpose and, in the case of ATLAS, relocate their Run 3 HLT farms for offline processing during Long Shutdown 3.

The following sections summarize the justification for the requests and the recommendations of the C-RSG.

10 Resource Requests: ALICE

The ALICE request for 2026 is shown in Table 8.

The extension of Run 3 into 2026 will result in a heavy-ion run of 21 days commencing on 30 May. The mix of Pb-Pb versus p-Pb in this period has yet to be decided. The estimated integrated luminosity collected during this run will be 3.24 nb^{-1} (200 nb^{-1}) for Pb-Pb (p-Pb) collisions. There are no plans to collect further pp reference run data for the remainder of Run 3 since the integrated luminosity goal has already been reached.

During the 2024 run, a new algorithm was implemented for the CTF (compressed time frame) data which reduced the data size per event by 14% for Pb-Pb collisions. Together with other improvements, ALICE have attained a total of ~25% decrease in the AO2D data size per event. A volume of 102 PB CTF data will be collected in 2026 p-p runs, which will be skimmed by a factor of 5% corresponding to 5 PB. In the Pb-Pb heavy-ion scenario, the storage requirements are 76 PB for the CTF files, two-thirds and one-third of which will be archived at the Tier0 and the Tier1 sites, respectively, and 13.8 PB of AO2D files. If p-Pb runs are chosen, 51 PB of CTF files and 13.4 PB AO2D data will be produced.

In the heavy-ion case, the number of MC events to be produced will be 8% of the physical events considered, unchanged from previous years.

The C-RSG notes the recurring compression issue. The current compression algorithm "Strategy A" is lossless and has been used for many years. There have been plans to move to a more aggressive compression, "Strategy B", which would result in 30% lower storage requirements. However since

ALICE		2024 RRB approved	2024 Pledged	2024 Used	2025 RRB approved	2025 Pledged	2026 Request	2026 req. / 2025 RRB	2026 C-RSG recomm.
CPU	Tier0	600	600	1109	680	680	710	104%	710
	Tier1	630	540	452	690	596	720	104%	720
	Tier2	650	630	571	730	750	810	111%	810
	Total	**1880**	**1770**	**2132**	**2100**	**2026**	**2240**	**107%**	**2240**
	HLT			*501*					
	Others			*76*					
Disk	Tier0	67.5	67.5	62.4	78.0	78.0	90.0	115%	90.0
	Tier1	71.5	61.9	53.1	79.0	69.1	90.0	114%	90.0
	Tier2	66.5	68.0	60.1	77.0	82.5	89.0	116%	89.0
	Total	**205.5**	**197.4**	**175.6**	**234.0**	**229.6**	**269.0**	**115%**	**269.0**
Tape	Tier0	181.0	181.0	123.0	220.0	220.0	292.0	133%	292.0
	Tier1	107.0	102.4	54.0	123.0	117.4	155.0	126%	155.0
	Total	**288.0**	**283.4**	**177.0**	**343.0**	**337.4**	**447.0**	**130%**	**447.0**

Table 8 ALICE resource requests for 2025 and C-RSG recommendations. For reference, the 2023 utilization and 2024 pledged resources are also given.

Strategy B is a lossy compression, transitioning to this algorithm is conditional on understanding the calibration. While significant progress has been made in this area, moving to a 50 kHz interaction rate in 2024 uncovered further issues with the buildup of space charge in the TPC which need to be understood before Strategy B can be implemented. Since the compression is a significant factor in the resource requirements, the C-RSG requests to be kept regularly updated on the progress towards implementing Strategy B.

Conclusions

The C-RSG finds the ALICE Collaboration's 2026 computing requests to be appropriate to achieve its physics programme and recommends that they be approved.

The C-RSG makes the following recommendation to the ALICE Collaboration:

ALICE-1 The C-RSG notes the ongoing issues surrounding the implementation of Strategy B compression. We understand that moving to a 50 kHz interaction rate uncovered new issues regarding the build up of space charge and that this is a subject of active discussion within ALICE. Since this affects the calibration and other parts of the processing, this is a significant factor in the resource requirements. Following discussions with the LHCC, we request that ALICE provide an update on the progress of this endeavor at the upcoming LHCC meeting in June, in the computing request report to be submitted to the C-CRSG by the end of August, and whenever there are significant developments.

11 Resource Requests: ATLAS

Table 9 shows the resource requests submitted by the ATLAS Collaboration for 2026, the resources used in 2024 and the pledged resources for 2024 and 2025 along with the C-RSG recommendations.

ATLAS		2024 RRB approved	2024 Pledged	2024 Used	2025 RRB approved	2025 Pledged	2026 Request	2026 req. / 2025 RRB	2026 C-RSG recomm.
CPU	Tier0	936	936	858	1100	1100	1265	115%	1265
	Tier1	1516	1514	1642	1635	1639	1802	110%	1802
	Tier2	1852	2074	3555	1998	2297	2202	110%	2202
	Total	**4304**	**4524**	**6055**	**4733**	**5036**	**5269**	**111%**	**5269**
	HLT			*290*					
	Others			*1982*					
Disk	Tier0	49.0	49.0	43.4	56.0	56.0	65.0	116%	65.0
	Tier1	163.0	163.1	172.0	186.0	186.7	199.0	107%	199.0
	Tier2	200.0	194.0	179.0	227.0	218.9	243.0	107%	243.0
	Total	**412.0**	**406.1**	**394.4**	**469.0**	**461.6**	**507.0**	**108%**	**507.0**
Tape	Tier0	207.0	207.0	194.0	258.0	258.0	302.0	117%	302.0
	Tier1	452.0	460.0	467.0	561.0	567.6	692.0	123%	692.0
	Total	**659.0**	**667.0**	**661.0**	**819.0**	**825.6**	**994.0**	**121%**	**994.0**

Table 9 ATLAS Collaboration resource requests for 2025 and C-RSG recommendations. For reference, 2023 and 2024 pledged resources are also given.

The computing resource request for 2026 reflects the updated LHC schedule, in which data taking will occur through June before the start of LS3. This request accounts for the reduced duration of the 2026 run while incorporating additional resource needs related to previous years' shifts in luminosity and schedule.

The experiment is asking for increases of 11% for CPU, 8% for disk and 21% for tape, in line with the expected increases in recorded data.

The ATLAS Collaboration makes the following assumptions on the running scenario for 2026: 81 days of p-p collisions with an integrated luminosity of $L = 80$ fb^{-1}, pile-up of 65 collisions/bunch crossing (up from 62) and a total running time of 3.5×10^6 s. With these running conditions, ATLAS expects to record 8×10^9 events in the main trigger stream and additional 5.3×10^9 events in a delayed stream.

ATLAS full simulation remains a major CPU consumer. Significant improvements in Run 3 simulation have increased its speed by 50% compared to Run 2. In 2024, Run 2 simulation accounted for 35% of the total simulation produced, a fraction that is expected to decline further.

In 2025, ATLAS has fully implemented a decrease of 12% of the RAW event size to 1.6 MB/event as a result of changes in the liquid-argon calorimeter (LAr) trigger data. ATLAS is exploring further reductions by investigating a potential change in the compression algorithm, which could yield additional savings. However, this change is unlikely to be implemented before the end of Run 3..

Several factors influence the requests. The HLT farm is expected to contribute less than previous years due to its relocation during the second half of 2026. Beyond-pledge CPU from WLCG sites and HPC resources are expected to contribute 1600 kHS23 for 2026. Disk needs remain steady due to the expected increase in MC production and the intensive derivation of data formats required for calibrations and uncertainty studies in the second half of the year. Additionally, the Tier0 CPU request remains consistent with full data-taking scenarios, ensuring adequate support for user analyses and MC production during the downtime following data collection.

Conclusions

The C-RSG considers that the ATLAS experiment's resource requests for 2026 shown in Table 9 are necessary to achieve the experiment's approved physics programme and recommends that the C-RRB arranges for pledges correspondingly.

The C-RSG makes the following observations and recommendations:

ATLAS-1 The C-RSG recommends that sufficient resources are allocated to fast simulation to ensure that the planned transition stays on track. The C-RSG also suggests to provide quantitative milestones in order to more easily track progress.

ATLAS-2 The C-RSG recommends that the ATLAS Collaboration continue its effort on the reduction of the raw event size.

ATLAS-3 The C-RSG requests that the ATLAS Collaboration investigate why a noticeably large fraction of data stored on disk has not been accessed.

ATLAS-4 The C-RSG requests that the ATLAS Collaboration report quantitative information on the resources allocated to parked data, both in absolute terms and as a fraction of total storage, including current usage and projected requirements for 2025 and 2026.

12 Resource Requests: CMS

CMS anticipates a p-p luminosity of 80 fb^{-1} in 2026 with an average pileup of 65. CMS plans to increase the data taking rate in 2025 and 2026 to 2.8 kHz (up from 2.6 kHz) to be promptly reconstructed and 6.0 kHz (up from 4.9 kHz) to be parked. The resource request for 2026 also includes provisions for the extended p-p running in 2024 and the increased luminosity expected to be collected in 2025. These additional resource needs offset the reduced requirements for 2026, which is not a full data-taking year.

To manage this increased data production and processing, CMS requests detailed resource increases across Tier0, Tier1 and Tier2 facilities. For CPU resources, CMS requests a 14% increase at Tier0 to manage higher pileup and event rates, essential for prompt reconstruction. At Tier1 sites, a 9% CPU increase is sought to accommodate reprocessing campaigns, MC production and user analyses. A 5% increase at Tier2 sites is requested, primarily to facilitate MC simulation, user analysis jobs and centrally managed data processing activities.

In terms of disk storage, CMS requests a 16% increase at Tier0 to expand buffer capacities and working space, facilitating management of higher data volumes and analysis datasets. A 15% disk storage increase at Tier1 sites is necessary to support the storage of processed data, MC production outputs and to enhance analysis capacity. At the Tier2 level, a 13% increase in disk storage is requested to accommodate the growing datasets required for user analyses and centrally managed data productions.

Tape storage requirements include a 17% increase at Tier0, reflecting anticipated RAW data volumes from increased collision rates and integrated luminosity. At Tier1 sites, a substantial 21% tape storage increase is essential to provide adequate archival capacity for RAW and processed data, aligning with the increased storage demands arising from boosted data production and MC simulation volumes.

CMS plans to leverage parked data to significantly boost physics acceptance and scouting data rates to explore potential discoveries, particularly in exotic new physics and heavy-flavour analyses.

CMS also plans to further utilize opportunistic resources, such as HPC and ARM architectures, following successful validation campaigns. Strategic expansions include integrating new Tier1 sites

CMS		2024			2025		2026		
		RRB approved	Pledged	Used	RRB approved	Pledged	Request	2026 req. / 2025 RRB	C-RSG recomm.
CPU	Tier0	980	980	1510	1180	1180	1350	114%	1350
	Tier1	930	1020	1406	1100	1166	1200	109%	1200
	Tier2	1600	1526	2247	1900	1830	2000	105%	2000
	Total	**3510**	**3526**	**5162**	**4180**	**4176**	**4550**	**109%**	**4550**
	HLT			*409*					
	Others			*772*					
Disk	Tier0	54.0	54.0	47.1	70.0	70.0	81.0	116%	81.0
	Tier1	122.0	115.7	91.1	142.0	133.8	164.0	115%	164.0
	Tier2	149.0	136.7	110.6	175.0	159.6	198.0	113%	198.0
	Total	**325.0**	**306.4**	**248.8**	**387.0**	**363.4**	**443.0**	**114%**	**443.0**
Tape	Tier0	320.0	320.0	300.7	442.0	442.0	515.0	117%	515.0
	Tier1	380.0	353.9	279.9	445.0	411.5	540.0	121%	540.0
	Total	**700.0**	**673.9**	**580.6**	**887.0**	**853.5**	**1055.0**	**119%**	**1055.0**

Table 10 CMS Collaboration resource requests for 2026 and C-RSG recommendations. For reference, 2024 and 2025 figures are also given.

(e.g., Serbia, Poland, and Saudi Arabia) to address storage deficits and computational demands effectively.

Conclusions

The C-RSG considers the CMS Collaboration's resource requests for 2026 shown in Table 10 necessary to achieve its approved physics programme, and recommends that the C-RRB arranges for pledges correspondingly.

The C-RSG makes the following recommendations:

CMS-1 The C-RSG is concerned about the persistent under-pledging of storage resources, particularly for tape. To better understand the root cause of this issue, the C-RSG requests that CMS provide details on how the shares of storage between specific Tier1 and Tier2 sites are determined.

CMS-2 The C-RSG requests that CMS investigate the reason behind the large amount of disk space occupied by rarely accessed data, as indicated by the data popularity plot.

CMS-3 The C-RSG requests that the CMS Collaboration report quantitative information on the resources allocated to parked data, both in absolute terms and as a fraction of total storage, including current usage and projected requirements for 2025 and 2026.

CMS-4 The C-RSG requests that the CMS Collaboration investigate and quantify the factors contributing to the low CPU efficiency of 68% at Tier2 sites.

13 Resource Requests: LHCb

The computing resource requests submitted by the LHCb Collaboration for 2025 are shown in Table 11. The resources used in 2024, the pledged resources for 2025 and 2026, along with the C-RSG recommendations, are also listed in that Table.

LHCb		2024 RRB approved	2024 Pledged	2024 Used	2025 RRB approved	2025 Pledged	2026 Request	2026 req. / 2025 RRB	2026 C-RSG recomm.
CPU	Tier0	174	174	468	283	283	344	122%	344
	Tier1	572	692	952	928	856	1127	121%	1127
	Tier2	319	356	514	518	535	629	121%	629
	Total	**1065**	**1222**	**1934**	**1729**	**1674**	**2100**	**121%**	**2100**
	HLT			*216*					
	Others			*94*					
Disk	Tier0	30.6	30.6	27.2	54.9	54.9	70.9	129%	70.9
	Tier1	61.2	57.9	66.7	89.9	82.7	107.1	119%	107.1
	Tier2	11.8	7.9	5.3	17.4	15.2	20.7	119%	20.7
	Total	**103.6**	**96.4**	**99.2**	**162.2**	**152.8**	**198.7**	**123%**	**198.7**
Tape	Tier0	117.1	117.1	94.6	170.4	170.4	202.2	119%	202.2
	Tier1	133.3	142.2	111.8	194.8	164.2	233.7	120%	233.7
	Total	**250.4**	**259.3**	**206.4**	**365.2**	**334.6**	**435.9**	**119%**	**435.9**

Table 11 LHCb Collaboration resource requests for 2026 and C-RSG recommendations. For reference, 2024 and 2025 resources are also shown.

LHCb requests for 2026 an increase of about 20% in computing resources for CPU power, disk storage and tape archival capacity relative to the values approved by the RRB for 2025. This reduced growth with respect to preliminary requests from Autumn 2024 is defined by the shorter-than-expected LHC run in 2026 and follows the LHCb computing model. After the start of LHC Long Shutdown 3 at the end of data taking in summer 2026, the current HLT farm will become available for offline computing and provide an additional 600 kHS23 of CPU power.

The CPU requirements are dominated by the MC simulation for Run 3 data, consuming 74% (2094 kHS23) of the total LHCb CPU budget. User analysis activities account for 23% (647 kHS23) of the requested CPU, while data processing and computing infrastructure consume the remaining 2%. In total, LHCb requires for 2026 2800 kHS23 CPU power, of which 2100 kHS23 will be provided by WLCG-pledged resources and 700 kHS23 will be provided by non-pledged resources.

The projected increase in data storage capacity is driven by the integrated data volume expected to be recorded in 2026. With a projected p-p run time of 3.5×10^6 s and a nominal throughput from the trigger farm to the offline system of 8.8 Gb/s, the expected data volumes on tape are 80.9 PB for Run 1 and Run 2 real and simulated (archived) data and 355 PB for Run 3 collision data, heavy-ion and fixed target data, minimum bias/no-bias stream and simulated data.

Conclusions

The C-RSG considers that the LHCb Collaboration's resource requests for 2026 are necessary to achieve the experiment's physics programme.

The following recommendations are made:

LHCb-1 The C-RSG requests LHCb to investigate the high CPU usage per full simulated event and quantify the proportions of detailed simulated events produced by LHCb.

LHCb-2 The C-RSG requests LHCb to understand and address discrepancies between the CPU utilization reported by EGI and DIRAC accounting and consider using only EGI accounting for future reports to avoid confusion.

LHCb-3 The C-RSG endorses LHCb to continue working on software improvements, namely the adoption of ARM-based CPUs and the porting of the LHCb codebase to support multi-threading for more efficient memory utilization, especially in relation to HPC resources.

LHCb-4 The C-RSG requests LHCb to base its CPU resource projections on event simulation times measured from real Run 3 MC simulations rather than design parameters.

14 Comments and Overall Recommendations

The C-RSG completed this scrutiny round impressed by the vast amount of data collected by the experiments in 2024 and the remarkable scale of data processing, simulation and analysis activities undertaken. The experiments have demonstrated an outstanding ability to leverage a large volume of opportunistic CPU resources, nearly doubling the baseline capacity for data processing, simulation and analysis.

Comprehensive plans for computing resource utilization in 2025 and 2026, along with the estimated computing resources required for the planned processing, simulation and analysis of the collected data, have been established. The C-RSG deems the computing resource requests from the experiments for 2026 to be well-aligned with the advancement of their physics programmes and recommends that the funding agencies provide the requested CPU, disk and tape resources.

Figure 3 illustrates the annual evolution of the computing resources approved by the RRB and the distribution of these resources among the experiments. ATLAS and CMS require significantly more resources than ALICE and LHCb, although the gap between these two groups is gradually narrowing driven by the high-rate upgrades implemented by ALICE and LHCb in Run 3. Additionally, the resources required by ATLAS and CMS are also converging over time.

Figure 4 (left) shows the evolution of WLCG computing resource capacities, aggregating the capacities of the four experiments and breaking them down by resource type (CPU, disk, and tape). Figure 4 (right) displays the annual increases in these resources. Since the start of LHC Run 3 in 2022, CPU and disk resources have grown at an average rate of approximately 15%, while tape capacity has increased at an average rate of 25%.

CPU and disk capacities are projected to double between 2021 and 2026, while tape capacity is set to triple, all with essentially the same annual budgets, thanks to the continuous decrease in hardware costs. Notably, during this period, the experiments are expected to increase by a factor four the amount of data collected during Run 1 and Run 2, highlighting the significant ongoing optimizations in data processing and storage.

Experiments have extensively utilized beyond-pledge CPU capacity available at WLCG sites, exceeding the formally requested amounts. Additionally, they have leveraged idle time in online farms for offline processing and accessed significant opportunistic and other non-pledged CPU resources, primarily from HPC centres. Leveraging these resources has significantly expanded MC production, leading to additional simulations that could cause a substantial impact on storage requirements. To evaluate the effects of these resources, the CRSG issues the following Recommendations:

ALL-1 The CRSG asks the experiments to report the outcomes achieved in 2024 through beyond-pledge and opportunistic CPU usage.

ALL-2 The CRSG requests that experiments classify non-pledged CPU resources into two categories: i) Dependable CPU, resources that can be reasonably trusted to be available and thus incorporated into planning, thereby reducing the CPU requested from WLCG sites; ii) Opportunistic CPU, resources with unpredictable availability that cannot be reliably accounted for in planning. This

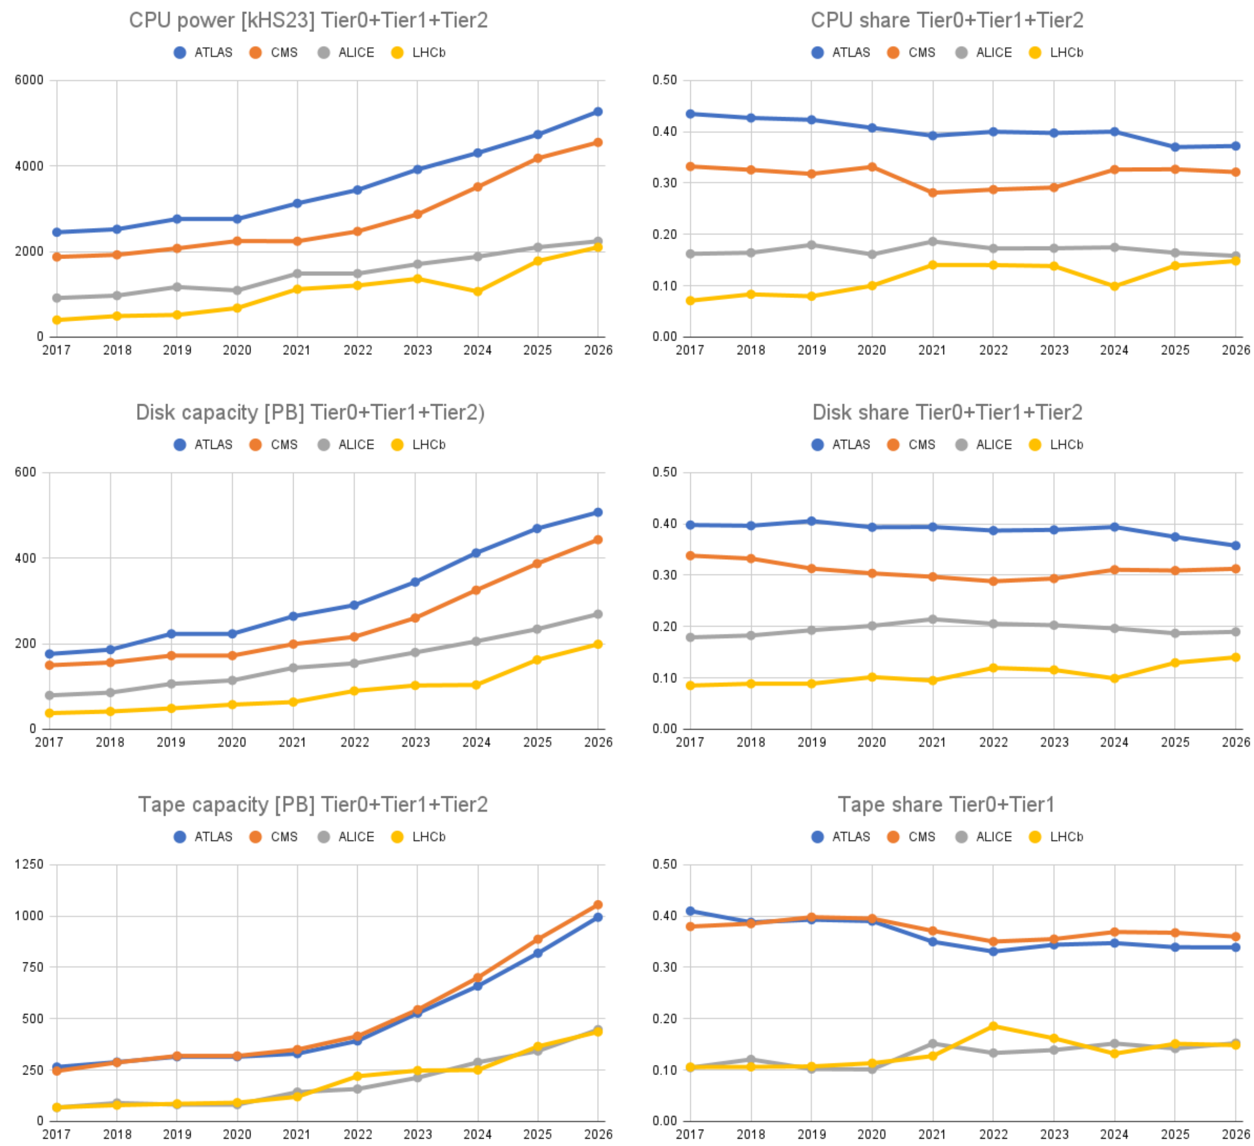

Figure 3 Left: Annual evolution of the RRB's approved baseline capacities for CPU, disk, and tape at the WLCG sites for the LHC experiments; the CPU and storage resources are in units of kHS23 and PB, respectively. Right: Distribution of resources among the LHC experiments. The 2026 values reflect the recommendations of the C-RSG.

distinction will help optimize resource allocation and ensure more accurate CPU requests from WLCG sites. The usage of dependable and opportunistic CPU resources should be reported separately in the annual usage report. When formulating resource requests, the amount of dependable CPU should be explicitly specified.

ALL-3 The CRSG requests that experiments quantify the data volume generated by the additional MC production in 2024, both on disk and tape. Additionally, experiments should estimate the storage volume required for similar extra MC productions anticipated in 2025 and 2026 and clarify whether these needs have been accounted for in the requested storage resources for 2026.

Given the substantial and consistent use of beyond-pledged CPU resources by the experiments over the years, coupled with the shortage of storage resources, the C-RSG intends to discuss with the experiments during the Autumn 2025 scrutiny the possibility of trading storage for CPU. This would involve regenerating reconstructed or simulated data as needed, rather than storing them on disk or archiving them on tape.

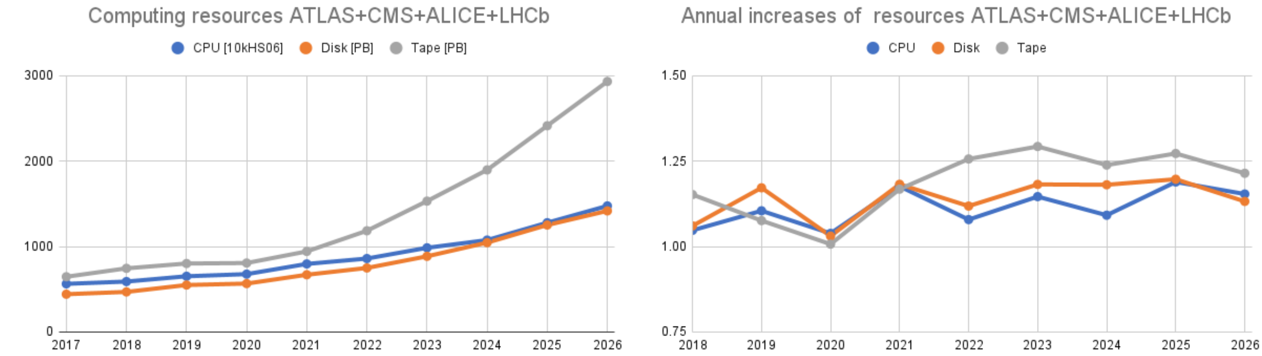

Figure 4 Left: Evolution of the RRB's approved overall baseline capacities for the WLCG computing resources (CPU in units of 10 kHS23, storage in PB). Right: annual resource increases. The 2026 values reflect the recommendations of the C-RSG.

Disk storage is the most expensive infrastructure at WLCG sites. To optimize its use, experiments aim to keep only frequently accessed data on disk. Accurate monitoring of disk usage and effective procedures to remove cold data are essential for efficient resource management. Producing the usual popularity plot, which shows data volumes on disk as a function of access frequency, is both difficult and time-consuming for experiments. To streamline this process and standardize the calculation of this important metric, the CRSG proposes a simplified approach, issuing Recommendation ALL-4. In addition, Recommendation ALL-5 is issued to ensure cold data are systematically evicted from disk.

ALL-4 The C-RSG requests that experiments replace the popularity plot with a report detailing the total volume of disk-resident data that has not been accessed for three, six, and twelve months, categorized by data tier. The reference timestamp for determining the last data access should be the end of the reported calendar year. The calculation should treat different instances of the same dataset separately—if a dataset has multiple copies on disk and only one is accessed, the remaining instances should still be considered unaccessed. The report should also include the corresponding fractions relative to the total data volume stored on disk.

ALL-5 The CRSG recommends that experiments establish automated procedures for unpinning data on disk that have remained unaccessed for the past 12 months.

Experiments typically maintain multiple instances of data on disk to enhance access performance or ensure data safety when the data are not backed up on tape. However, data replication policies significantly impact total disk space usage. To help assess the efficiency of disk usage across experiments, the CRSG delivers Recommendation ALL-6:

ALL-6 The CRSG requests that experiments describe and justify their disk replication policies, including the total volume occupied by second and additional copies of the data, as well as the fraction of total disk space these copies represent. The reported figures should be broken down by data tier.

References

[1] Computing Resources Scrutiny Group. C-RSG autumn 2024 report. *CERN-RRB-2024-079*, 14 October 2024.

[2] CERN. LHC Schedule. http://lhc-commissioning.web.cern.ch/lhc-commissioning/schedule/LHC-long-term.htm.

[3] ALICE Collaboration. ALICE Computing Resources Usage January through December 2024, Resource Requirements for 2026. 1 February 2025, submitted 18 February 2025.

[4] ATLAS Collaboration. ATLAS computing status and plans – Report to the C-RSG – April 2025. 18 February 2025.

[5] CMS Collaboration. Spring 2025 scrutiny – CMS utilisation report 2024 and resource request for 2026. 18 February 2025.

[6] LHCb Collaboration. Utilisation of LHCb offline computing resources in 2024. *LHC-PUB-2025-007*, 18 February 2025.

[7] LHCb Collaboration. LHCb computing resources: 2026 requests. *LHC-PUB-2025-006*, 18 February 2025.

[8] EGI Foundation. Accounting portal. https://accounting.egi.eu/.

[9] WAU. WLCG Accounting Utility portal. https://monit-grafana.cern.ch/dashboards/f/Jux9_-1ik/accounting.