# Computing Resources Scrutiny Group Report

**José M. Hernández** (CIEMAT)

for the Computing Resources Scrutiny Group

April 28-29, 2025

CERN-RRB-2025-10

# C-RSG membership

| | | | |
|---|---|---|---|
| C Allton (UK) | ALICE, lead | R AAij (Netherlands) | ATLAS, CMS |
| M Schulz (CERN) | ALICE | T Mkrtchyan (Germany) | LHCb, lead |
| J Amundson (USA) | ALICE | D Gingrich (Canada) | LHCb |
| J Kleist (Nordic C.) | ATLAS, CMS, lead | H Meinhard (CERN) | scient. secr. |
| E Fede (France) | ATLAS, CMS | J Hernández (Spain) | chair, LHCb |
| D Elia (Italy) | ATLAS, CMS | A Valassi (CERN) | scient. secr. |

- First review led by J Hernández (Spain) as chair of the C-RSG
- C-RSG thanks the previous chair, P Sinervo (Canada), for years of dedicated leadership and successful guidance of the C-RSG
- C-RSG thanks the C-RSG Scientific Secretary, H Meinhard (covering for A Valassi), the collaboration computing representatives and CERN management for their support
- J Amundson (FNAL) replaces A. Connolly (University of Washington) as the U.S.-proposed member of the C-RSG. We invite the RRB to formally confirm his nomination

# Spring 2025 Scrutiny Process

- C-RSG met with LHCC WLCG referees and chair in advance
  - Identified a number of issues of common concern
- LHC Collaborations submitted report on
  - 2024 computing resource usage
  - Computing activities and plans for 2025
  - Final resource requests for 2026
- C-RSG responded with written questions
  - Met with Collaboration computing representatives
  - Provided Collaborations with draft report for any corrections of facts
- Made recommendations regarding requests and plans

# Computing resources in 2024

Resources **approved** by RRB for 2024

| | 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% |

- ~11M HS23 CPU, ~1 EB disk, 2 EB tape
- **Under-pledges at Tier1 and Tier2 levels**
  - Some strain on storage

Balance (**Pledged**-RRB)/RRB [%]

| | | 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 |

# CPU utilization in 2024

CPU in kHS23 units

| | 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 |

- **Substantial beyond-pledge CPU leveraged:** $Used_{total}/RRB = 182\%$
  $Used_{WLCG}/RRB = 142\%$
- Extra CPU from WLCG sites, HLT farms and HPC facilities
- Extension of MC simulation samples
- The HLT CPU is “dependable”, used in planning and reducing CPU requested to WLCG sites

# CPU utilization history

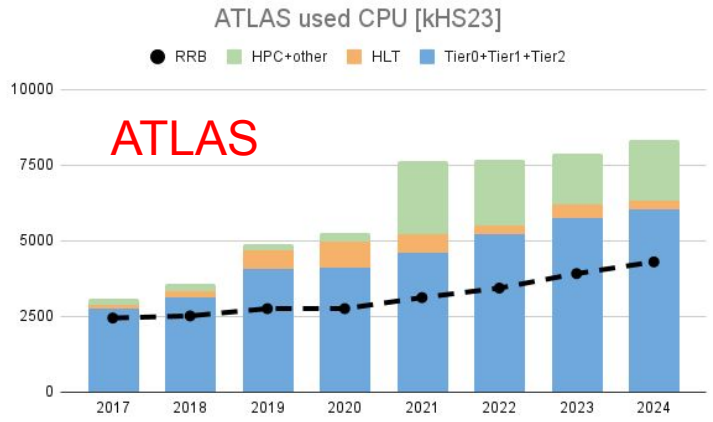


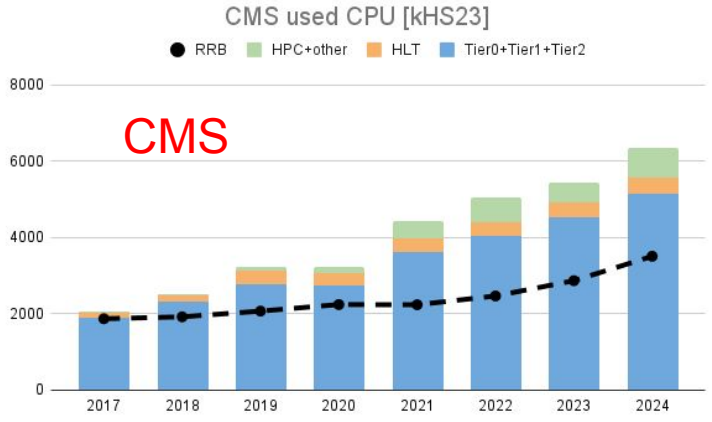


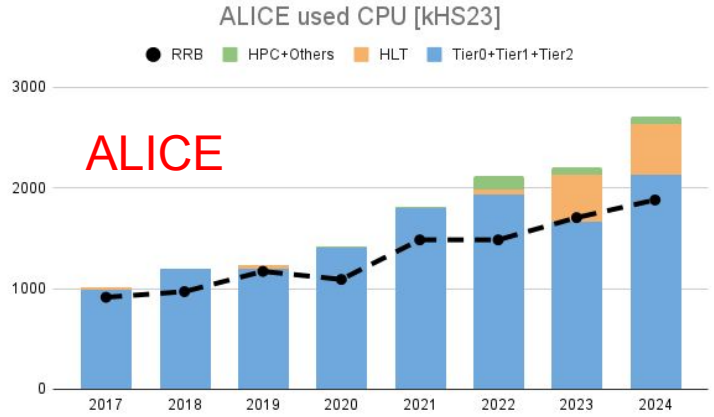


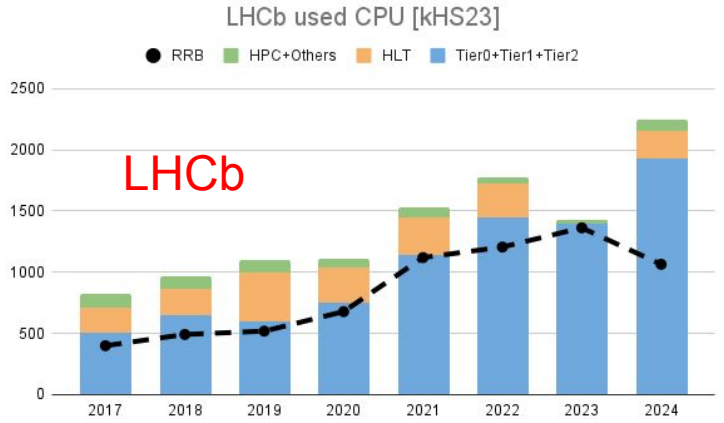

# Heterogeneous compute resources

- Increasing use of High Performance Computing (**HPC**) facilities
  - Workshop at CERN in Jan. 2025 discussing strategic and technical aspects of HPC utilization
    - **Strategic access** beyond traditional allocation
    - Development of common interfaces to facilitate **seamless integration**
- Usage of **energy efficient, cost-effective** processor architectures
  - **ARM**
    - Experiments have ported their software stacks, finishing physics validations, and even accepted a fraction of the pledged CPU as ARM
    - Experiments would benefit from an extended infrastructure at CERN to facilitate faster software building, validation and debugging for the ARM architecture
  - **GPU**
    - Continuous efforts to adapt experimental applications for execution on GPUs
    - Extended use for online event triggering
    - Increasingly prevalent in large HPC facilities
    - Benchmarking of GPUs will be needed for accounting and pledging

# Storage utilization in 2024

- Disk and tape limited to pledged resources at WLCG sites
  - No beyond-pledge or “opportunistic” storage
- Some strain on storage capacity and management in 2024
  - Data-taking four weeks longer than planned
  - Storage under-pledges 5-10%
  - Increased MC production driven by beyond-pledged CPU resources
  - ALICE and LHCb negotiated an early allocation of a fraction of the 2025 disk resources
  - **Efficient use of disk resources is** even more **critical** in the context of constrained availability
- ATLAS and CMS record additional data streams (parked / delayed data)
  - Not immediately reconstructed offline by the Tier0 system but are processed at a later stage
  - ATLAS: 2.2 kHz (4 GB/s) main, 1 kHz (2 GB/s) delayed
  - CMS: 2.4 kHz (2 GB/s) main, 4.8 kHz (4 GB/s) parked
  - Aimed at maximizing physics potential of analyses benefit from increased data collection rates
  - **Parked/delayed streams consume a significant fraction of storage resources**

# 2026 computing resource requests

- Higher p-p luminosity (~25%) expected in 2025
  - Resources for 2025 fixed in 2024 before the change of schedule
  - ATLAS and CMS include additional resources for 2025 in the 2026 request
  - Expecting early allocations
- Early start and shorter data-taking in 2026
  - Experiments have recalculated the required resources
  - Notably, LHCb reduces the requested resource increase from ~60% to 20%
  - For ATLAS and CMS, the extra resources for 2025 offset the decrease for 2026
  - **Earlier deployment of 2026 resources is important**

# 2026 computing resource requests

- Experiments have established 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
- 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 funding agencies provide **the requested CPU, disk and tape resources**

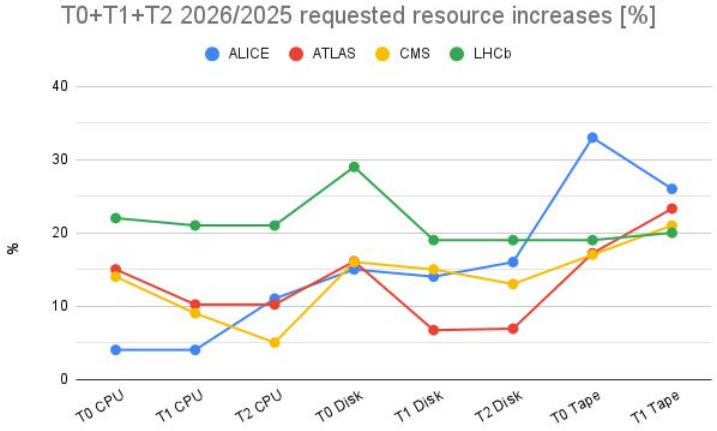


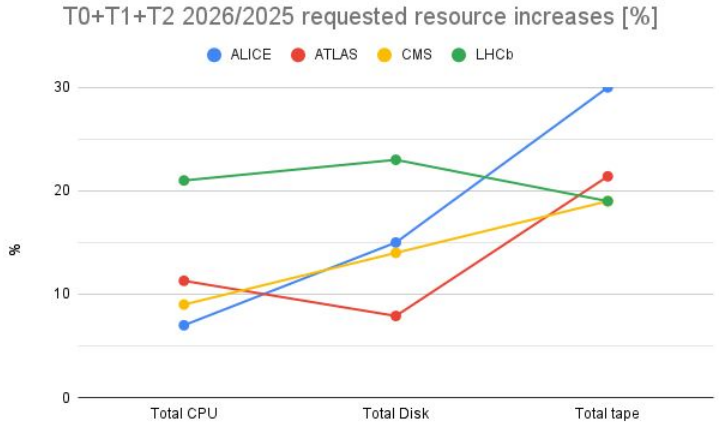

# Computing resource evolution

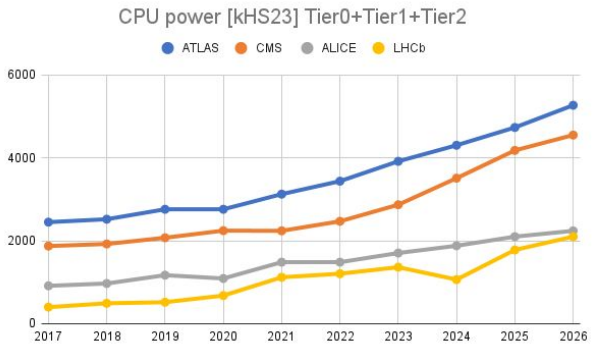


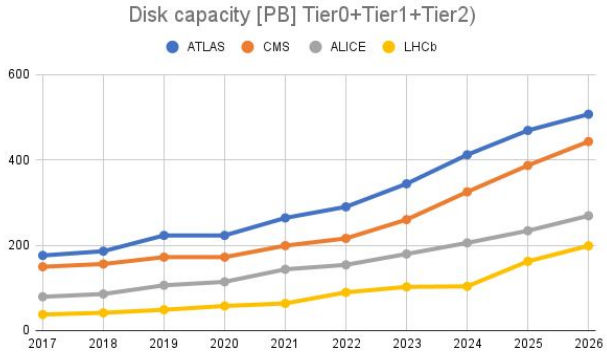


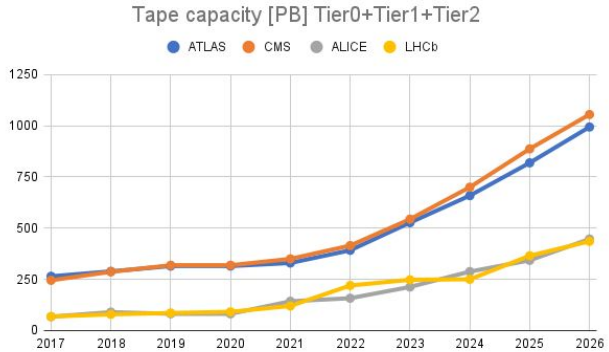


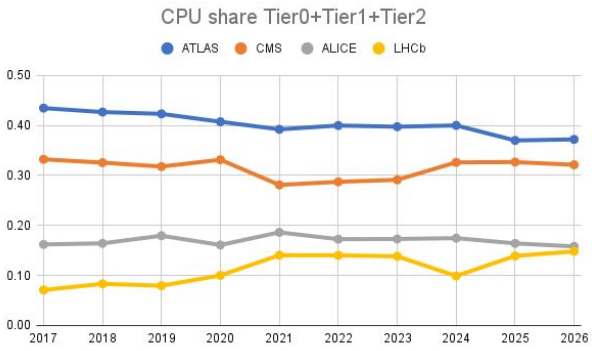


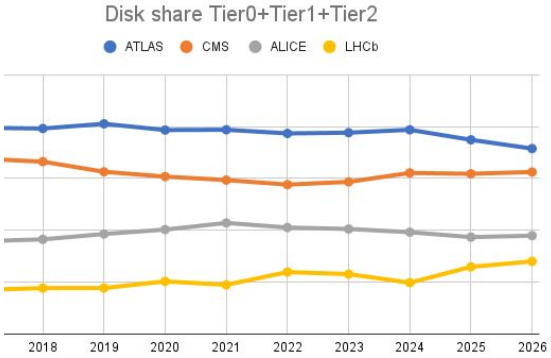


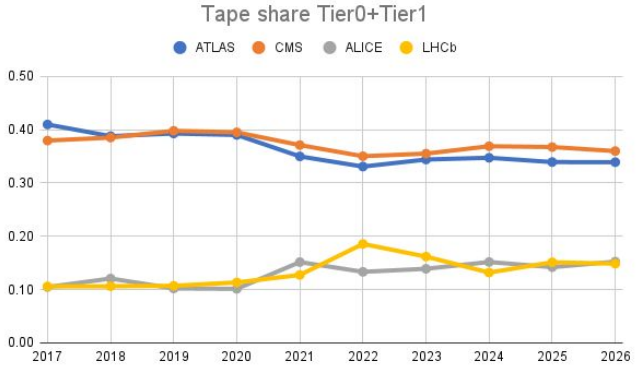

# Computing resource annual increases

- CPU and disk capacity growing at a rate of **~15%**, tape **~25%**
  - Growth achieved with essentially flat budget, since hardware got cheaper at the same rate
    - Not anymore? In 2023 and 2024 disk storage got cheaper by 8% and 3% respectively
- In the period 2021-2026, CPU and disk have **~doubled**, tape **~tripled**
  - Run3 data ~ **4** x (Run1 + Run2) data ➡ Large optimizations in data processing and storage

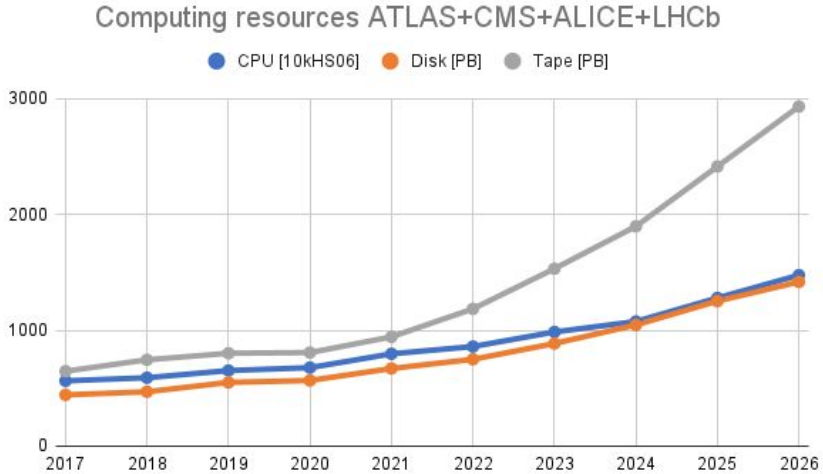


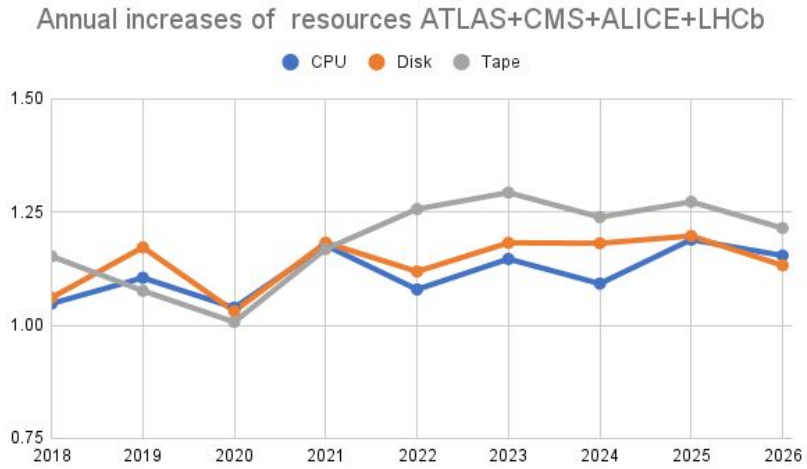

# ALICE Usage in 2024 and Request for 2026

| ALICE | | 2024 | | | 2025 | | 2026 | | |
|---|---|---|---|---|---|---|---|---|---|
| | | RRB approved | Pledged | Used | RRB approved | Pledged | Request | 2026 req. / 2025 RRB | 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** |

2026/2025 requested resource increases:

- CPU: 7% (+140 kHS23)
- Disk: 15% (+35 PB)
- Tape: 30% (+104 PB)

## 2024 utilization

- p-p @ 5.36 TeV
  - Reference data collected, with no further data needed in Run 3
- Pb-Pb @ 5.36 TeV
  - 1.54 $nb^{-1}$ physics data; more than in 2023
  - Data taking at 50 kHz
  - Issues with TPC calibration prevents moving to lossy compression of raw data
    - **Impact on resources is a concern**
- Advanced allocation of disk and tape at Tier0 and Tier1s was required

## 2026 plan is 21 days of heavy-ion run

- 3.24 $nb^{-1}$ collected data expected

# ALICE recommendations

The C-RSG considers ALICE's computing requests for 2026 to be appropriate for achieving its physics program and **recommends their approval**

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

# ATLAS Usage in 2024 and Request for 2026

| ATLAS | | 2024 | | | 2025 | | 2026 | | |
|---|---|---|---|---|---|---|---|---|---|
| | | RRB approved | Pledged | Used | RRB approved | Pledged | Request | 2026 req. / 2025 RRB | 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** |

2026/2025 requested resource increases:

- CPU: 11% (+536 kHS23)
- Disk: 8% (+38 PB)
- Tape: 21% (+175 PB)

## 2024 utilization: p-p @ 13 TeV

- 117 $fb^{-1}$ data collected
- Large beyond-pledge CPU at Tier2 (192%)
- Large opportunistic CPU from VEGA HPC
- Single copy of analysis datasets on disk complemented with dynamically created replicas on disk cache
- Tape fully utilized
  - Exploring use of tape at T2s (DESY, NET2/NESE)

## 2026 resource request

- 12% reduction in RAW event size for 2025
- Relocation of HLT farm at the end of Run 3 for offline usage

# ATLAS recommendations

The C-RSG considers ATLAS's computing requests for 2026 to be appropriate for achieving its physics program and **recommends their approval**

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
2. The C-RSG recommends that the ATLAS Collaboration continue its effort on the **reduction of the raw event size**
3. The C-RSG requests that the ATLAS Collaboration investigate why a noticeably large fraction of **data stored on disk has not been accessed**.
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

# CMS Usage in 2024 and Request for 2026

| 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** |

2026/2025 requested resource increases:

- CPU: 9% (+370 kHS23)
- Disk: 14% (+56 PB)
- Tape: 19% (+168 PB)

## 2024 utilization: p-p @ 13 TeV

- 113 $fb^{-1}$ data collected
- Large beyond-pledge CPU at all Tier levels (~150%)
- Low CPU efficiency at Tier2 (67%)
- Good use of Run 2 and Run 3 HLT farms
- Storage underpledges
- Large fraction of disk space used by data barely used for analysis (AOD)
- Resumed using storage resources at JINR
- New Tier1s coming (Poland and Serbia)

## 2026 resource request

- Increase of data rate (main and parking) ~15%

# CMS recommendations

The C-RSG considers CMS's computing requests for 2026 to be appropriate for achieving its physics program and **recommends their approval**

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.
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
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
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

# LHCb Usage in 2024 and Request for 2026

| LHCb | | 2024 | | | 2025 | | 2026 | | |
|---|---|---|---|---|---|---|---|---|---|
| | | RRB approved | Pledged | Used | RRB approved | Pledged | Request | 2026 req. / 2025 RRB | 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** |

2026/2025 requested resource increases:

- CPU: 21% (+321 kHS23)
- Disk: 23% (+37 PB)
- Tape: 19% (+71 PB)

## 2024 utilization: p-p @ 13 TeV

- First complete data-taking year in Run 3!
- Doubling collected dataset so far
- Lower data rates than foreseen (8.8 GB/s)
- Large beyond-pledge CPU at all Tier levels (~180%)
  - Increase CPU for MC production
- Good use of Run 2 and Run 3 HLT farms
- Storage under-pledges
  - Advanced allocations of disk at Tier1

## 2026 resource request

- Resource requests increases reduced from ~60% to ~20% due to shorter 2026

# LHCb recommendations

The C-RSG considers LHCb's computing requests for 2026 to be appropriate for achieving its physics program and **recommends their approval**

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.
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.
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.
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.

# Overall recommendations summary

- Impact of beyond-pledge CPU resources
  - Historically, experiments have leveraged significant CPU from HLT farms, HPC and WLCG
    - Typically used for significantly extending MC simulation samples
  - C-RSG  requests experiments to quantify the **impact on storage resources**
  - C-RSG  has asked experiments to quantify the beyond-pledge **“dependable” CPU**
- Efficient use of disk resources
  - C-RSG proposes to simplify the metric quantifying **unaccessed data** on disk
  - C-RSG recommends establishing automated procedures to **purge** unaccessed data
  - C-RSG requests experiments to describe and justify their disk **replication policies**
- Trading storage with CPU
  - C-RSG intends to discuss with the experiments during the Autumn 2025 scrutiny the possibility of **trading storage for CPU**
    - Regenerating reconstructed or simulated data as needed, rather than storing them on disk or archiving them on tape

# Overall recommendations

1. The CRSG asks the experiments to report the **outcomes** achieved in 2024 through **beyond-pledge and opportunistic CPU usage**
2. The CRSG requests that experiments classify non-pledged CPU resources into two categories:
    a. **Dependable CPU**, resources that can be reasonably trusted to be available and thus incorporated into planning, thereby reducing the CPU requested from WLCG sites
    b. **Opportunistic CPU**, resources with unpredictable availability that cannot be reliably accounted for in planning. This 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.

# Overall recommendations

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

4. The CRSG requests that experiments replace the popularity plot with a report detailing the total volume of **disk-resident data** that have **not** been **accessed** for three, six and 12 months, categorized by tier. The report should also include the corresponding fractions relative to the total data volume stored on disk

# Overall recommendations

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

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

# Summary

- The final two years of LHC Run 3 will substantially expand the accumulated datasets, advancing physics programmes toward even greater achievements
- Corresponding computing resources will be essential to support the expected increases in data volume, processing, simulation and analysis
- The experiments have demonstrated an outstanding ability to leverage a large volume of beyond-pledge CPU resources, nearly doubling the baseline capacity
  - The C-RSG would like to assess the impact of the enlarged MC simulation samples on storage resources
- The C-RSG recommends explororing further optimization of storage resources
  - Review disk replication policies, reduced levels of unaccessed data, trading storage for CPU
- The C-RSG will collaborate with the WLCG LHCC to review experiment-specific parameters that drive computing resource requirements