CERN–RRB–2023–012
10th April 2023

Computing Resources Scrutiny Group

C Allton (UK), P Christakoglou (The Netherlands) A Connolly (USA), J Hernandez (Spain), J Kleist (Nordic countries), T Mkrtchyan (Germany), N Neyroud (France), P Sinervo (Canada, chair person), V Vagnoni (Italy), A Valassi (CERN, scientific secretary), J van Eldik (CERN)

1 Introduction

The Computing Resources Scrutiny Group (C-RSG) is responsible for evaluating the computing needs for the Large Hadron Collider (LHC) experiments, ALICE, ATLAS, CMS and LHCb, and for making recommendations to the Computing Resources Review Board (C-RRB) regarding the World-Wide LHC Computing Grid (WLCG) computing resources required by the collaborations to address their approved physics programme. The C-RSG has completed its spring 2023 review, which considered the computing usage in 2022, anticipated usage in 2023 and requests for 2024 of each of the four LHC experiments. This report summarises the results of this scrutiny. Unless otherwise noted, the start of each reporting period is 1st April.

The C-RSG solicited reports on the utilization of computing resources in 2022, the anticipated utilization in 2023 and proposed utilization in 2024. The C-RSG reviewed these reports, which also addressed the C-RSG recommendations from the most recent scrutinies in spring 2022 and autumn 2022. The C-RSG sent questions to each collaboration before dedicated meetings were held with collaboration representatives for clarifications. The face-to-face meetings with the collaboration representatives allowed the C-RSG to develop a more informed understanding of the computing needs of each collaboration.

The C-RSG also had a meeting with the LHCC (LHC Experiments Committee) WLCG referees, who focus on the role of the WLCG in supporting the approved physics programmes of the experiments. This allowed the LHCC and the C-RSG to confer on areas of common interest and concern prior to the scrutiny, and helped inform both groups regarding the issues related to computer resources and utilization. Although the mandates of the two groups are quite different, the overlap in the evaluation and oversight of computing resources made the discussion helpful and informative. The C-RSG thanks the LHCC referees for their participation in this discussion.

The C-RSG recommendations for 2024 resource procurement by the WLCG are provided in this report. They are based on the current LHC Run 3 schedule for 2024 [1]. The C-RSG recommendations assume that the 2024 operation of the LHC will increase the Run 3 datasets for all four collaborations, with a rapid increase in volume of data.

2 C-RSG Membership

There have been no changes to the C-RSG membership since the last report.

However, both the French funding agency and CERN have informed the C-RSG Scientific Secretary that they will be appointing new representatives to the C-RSG following the expiration of the terms of their representatives. In particular, these organizations have proposed that Eric Fede (IN2P3) will replace Nadine Neyroud and that Markus Schulz (CERN) will replace Jan van Eldik. The C-RRB is requested to approve this new membership.

The chair thanks the C-RSG members for their commitment and expert advice. The committee is particularly grateful for the work of the experiment representatives to address the questions and the previous recommendations by the C-RSG.

Thanks are also due to the CERN management for its support and to the C-RSG scientific secretary, A Valassi (CERN), for ensuring the smooth operation of the C-RSG and for his expertise in organising and conducting this and preparing subsequent rounds of scrutiny.

3 Interactions with the Experiment Collaborations

The experiment collaborations submitted their reports 21st February 2023. The C-RSG thanks the collaborations for the timely submission of their detailed documents [2–6], which also contained responses to the findings and recommendations from the autumn 2022 scrutiny round [7]. The group would like to thank the computing representatives of the experiments 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.

Specific teams of C-RSG referees were assigned to review the ALICE and LHCb requests. As usual, by agreement with the ATLAS and CMS managements, a single team of C-RSG referees scrutinised the ATLAS and CMS reports and requests to ensure a consistent approach. The referees subsequently reported to the full C-RSG, which then developed the recommendations in this report.

In anticipation of the autumn 2023 scrutiny, the C-RSG asks the experiments to submit their documents by Tuesday 29th August 2023. The Scrutiny Group is aware that this timeframe is a challenge given summer schedules, but is unfortunately constrained by the need to submit its report to the C-RRB by the 9th October 2023. The C-RSG requests that as part of their submission, the experiments respond to its general recommendations as well as those specific to the respective experiment in this report.

4 Resource Usage in Calendar Year 2022

4.1 Overall Usage

The C-RSG requested the collaborations to summarise their use of both pledged and non-pledged computing resources for the period January 2022 to December 2022. This approach to collecting these data has been found to be more accurate than relying on an independent assessment performed by the scrutiny group using the accounting resources available online. The collaboration statistics take into account a number of relatively small corrections and provide a measure of the opportunistic computing resources used by each experiment.

The collaborations have been focussing their analysis efforts on the first Run 3 data and in completing the Run 2 analyses. The WLCG resources are used for a mixture of data reconstruction, physics and detector simulation, and physics analysis. The production of large simulated event samples dominates the CPU use, while disk and tape storage largely scales with the recorded data volume. This usage report will follow the usual classification of CERN (T0), Tier-1 (T1) sites and Tier-2 (T2) sites.

The resources pledged by the funding agencies were approved in March 2021 when a full-year of LHC data-taking was anticipated in 2022. The events in Europe since that decision was taken had significant effects on these plans: The LHC running period in 2022 was shortened, the heavy-ion run

		Used/pledged resources						
		2022	2021	2020	2019	2018	2017	2016
CPU	T0	136%	131%	187%	119%	107%	105%	122%
	T1	128%	116%	118%	116%	92%	97%	119%
	T2	177%	172%	151%	154%	131%	143%	151%
Disk	T0	80%	83%	83%	87%	80%	72%	97%
	T1	93%	89%	88%	94%	79%	88%	72%
	T2	90%	83%	79%	90%	94%	—	—
Tape	T0	74%	78%	91%	92%	75%	64%	98%
	T1	76%	89%	91%	89%	63%	53%	67%

Table 1 Usage summary by tier and WLCG year for 2016 through 2022. Data is taken from the experimental reports, where some corrections have been made to account for inaccuracies in the WLCG reporting. Pledged resources are the available resources reported by the experiments at the end of the 2022 calendar year.

was postponed and run plans for 2023 have been revised. In addition, the plans for LHCb data-taking were revised given the delay in fully commissioning the new VELO detector. These factors are reflected in the reported utilization of the available WLCG resources. The computing plans for the ATLAS and CMS collaborations were least affected as the changes in 2022 data-taking resulted in a reduction of the total data volume collected but didn't effect their simulation production plans or analysis efforts, while data-taking by ALICE and LHCb was limited.

A significant risk arose in 2022 from the potential loss of computing resources provided by Russian and Belarusian institutions. The collaborations ensured that the affected sites did not hold any data that was not replicated elsewhere, and treated the CPU available at these sites in a manner similar to opportunistic "non-pledged" resources. As reported elsewhere, the collaborations are working with WLCG management to identify potential new T1 sites to provide longer-term mitigation of this risk.

Despite all these planning challenges, the WLCG resources in 2022 were used effectively by each experiment, with all experiments using CPU capacity made available on an opportunistic basis by the WLCG facilities, high-performance computing centres (HPCs) and specialised computing centres that have agreed to provide access to CPU cycles (such as several national laboratories). In total, these sources provided over 5000 kHS06-years of computing (4200 and 1100 kHS06-years for WLCG and non-WLCG sites, respectively), which is ~60% more CPU than pledged by T0, T1 and T2 sites. The large opportunistic WLCG contribution in particular reflects increases in these resources at all tiers. The C-RSG congratulates all experiments for utilising these resources.

Disks and tapes for data storage continue to be the most constrained resource, with 662 PB of disk being used by the experiments (114 PB at T0, 276 PB at T1 sites and 272 PB at T2 sites). The experiments have also used 880 PB of tape storage (333 PB at T0 and 547 PB at T1 sites) typically for archiving of raw and processed data that is not immediately required. Both T0 and T1 tape storage increased by over 10% over 2021, reflecting the archiving of raw data from the first year of the LHC Run 3. Together, the four experiments utilised over ~1.5 EB of storage in 2022.

4.2 CERN, Tier-1 and Tier-2 Usage

The usage relative to the pledged resources for CERN, T1 sites and T2 sites is shown in Table 1 for the last seven years. Values are summed over the four experiments. Disk and tape numbers give the occupancy of available resources at the end of the calendar years.

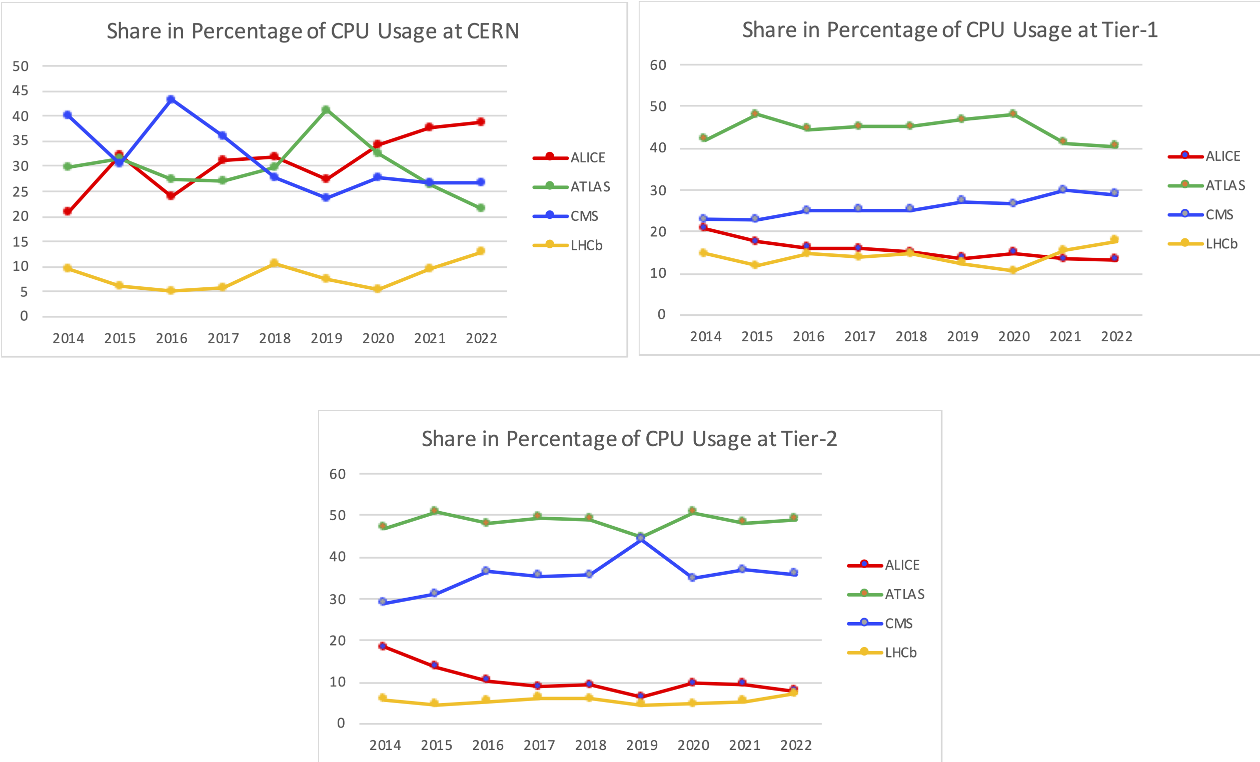

Figure 1 The figure on the top-left shows the percentage of CERN CPU that is used by each experiment. The figure on the top-right shows how the CPU resources at T1 sites are shared by experiment. The bottom figure shows the same sharing of T2 CPU resources. The shares for each year sum to 100%.

As in past years, the CPU usage exceeds the amount formally pledged by WLCG sites and so reflects opportunistic use of resources. These data do not include non-pledged CPU resources arising from the use of HLT farms, HPC centres and other non-WLCG resources. Most of these opportunistic resources are used for production of simulated events. In case these resources do not become available, the production of physics results will be delayed, but this is not expected to compromise the experiments' physics goals. As in past years, the C-RSG has discussed with the experiments the effect on their physics programmes of losing a significant fraction of these resources. The collaborations have been confident that the physics goals would be met but the delivery of that programme would be likely protracted.

The disk space is completely utilised when one takes into account the "headroom" needed to manage large data sets and the distributed analysis efforts underway by each collaboration. Although tape usage is below pledges, the expenditures for tape media by WLCG sites are made as experiments require increased storage within the limits of the pledges.

Figure 1 shows the yearly evolution of the share of CPU usage by experiment at CERN (top left), T1 sites (top right) and T2 sites (middle bottom). In each plot, the percentage used by each experiment normalised to the total CPU cycles used is plotted, therefore they sum up to 100% year by year. The major users at CERN are ALICE, ATLAS and CMS, with LHCb being a minor user below the 10% level. T1 CPU usage is dominated by the ATLAS experiment at 40%, with the CMS experiment at around 30% and the ALICE and LHCb experiments around 15% each. The ATLAS and CMS experiments dominate the use of T2 computing resources, maintaining a consistent pattern even as T2 resources have increased.

Figure 2 shows the year-by-year usage of disk storage by the experiments at CERN (left) and T1 sites (right). The percentage is obtained from the space used by each experiment divided by the total disk space used at CERN and T1 sites, so that by definition they sum up to 100% year-by-year. At

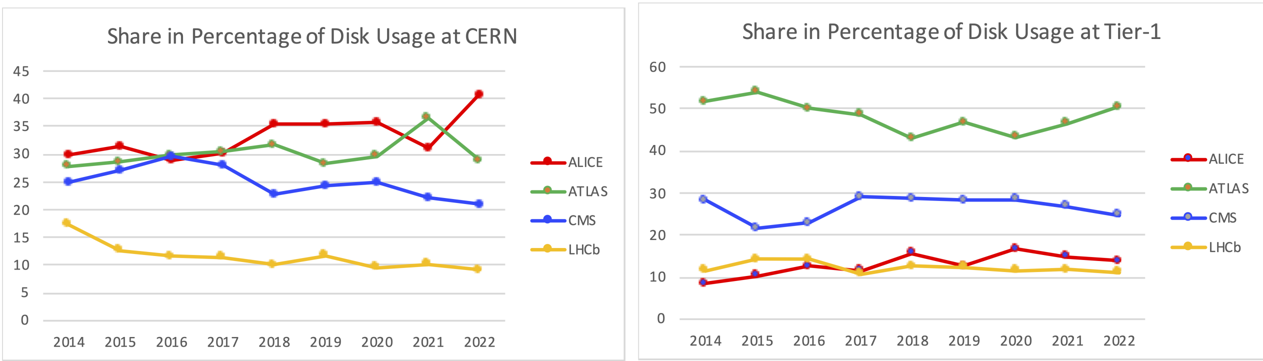

Figure 2 Year-by-year share of disk space by experiment used at CERN (left) and T1 sites (right). Data were obtained from CRIC [8] using the WLCG year.

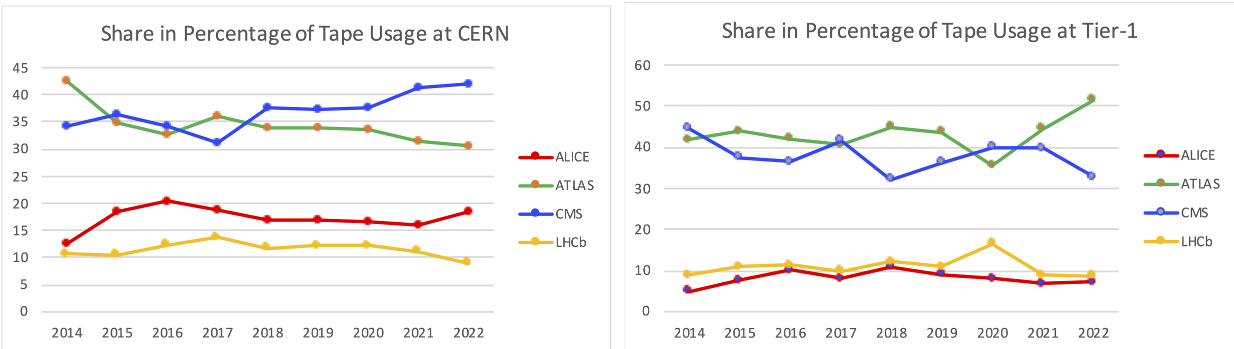

Figure 3 Year-by-year share of tape storage space used by experiments provided by CERN (left) and the T1 facilities (right). Data were obtained from the experiments using the calendar year.

CERN, ALICE, ATLAS and CMS continue to use comparable amounts of disk space, though CMS has significantly reduced its usage in the last two years, reflecting changes that the collaboration has made in its analysis model. LHCb's share decreased modestly in the last year, which reflects the reduced data-taking during 2022.

The largest disk storage increases have taken place at T1, with the total volume having increased over the last five years from 174 PB in 2018 to 276 PB in 2022. As shown in Figure 2, this increase has been distributed across all experiments, though the relative share for ATLAS has increased by over 5% in the last two years with a commensurate drop in the shares of T1 disk space for the other collaborations.

Usage of tape storage at T0 and T1 sites is shown in Figure 3, illustrating the fraction of space used by each collaboration. Tape usage has increased by 89 PB, with tape storage at CERN increasing by 64 PB and at T1 sites by 25 PB, which is a reversal of the growth in tape capacity compared with the 2021 year. The share of tape used by each collaboration has not changed significantly, with ATLAS and CMS using approximately 40% each and ALICE and LHCb using between 10 and 15% of the total tape storage.

Finally, Figure 4 explores the year-by-year evolution of the fraction of CERN CPU resources and disk storage used by a given collaboration compared with its total utilisation. CPU usage and disk space utilisation are shown on the left and right, respectively. The CERN CPU contributions to the ATLAS, CMS and LHCb experiments' needs are around 10-20% , while ALICE relies on CERN for over 45% of its CPU resources.

The proportion of CERN disk space used by the collaborations has shown a secular decline for all experiments until 2022, when the need for more CERN disk by the ALICE and CMS experiments

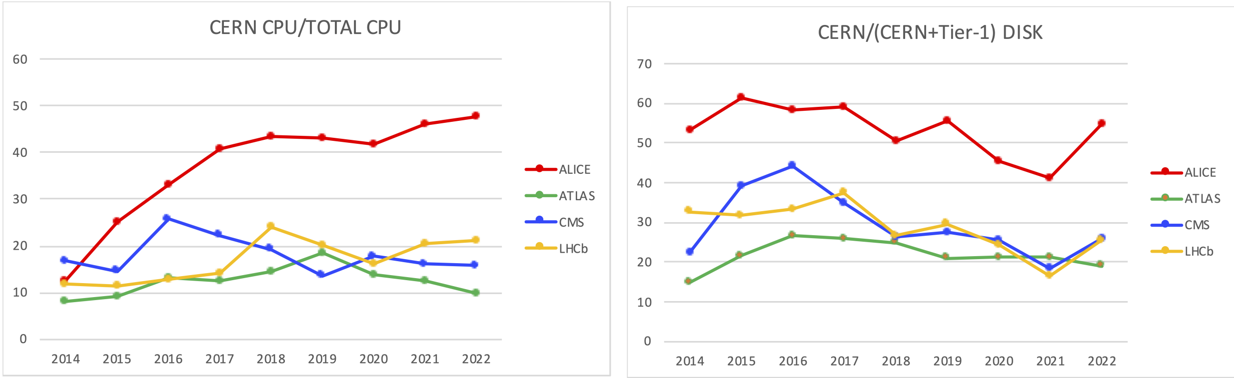

Figure 4 CPU usage at CERN divided by the total usage (left), defined as the sum of CERN, T1 and T2 CPU. Disk space usage at CERN divided by the sum of CERN and T1 usage (right). T2 disk space utilization is not available for the entire period.

increased. The overall trend reflects an increasing amount of T1 disk space available to all collaborations. The ALICE collaboration has the greatest reliance on CERN disk resources with over 50% of its disk space provided by CERN. The CMS, ATLAS and LHCb collaborations have between 20 and 30% of their disk space provided by CERN. These data illustrate the important role played by CERN in the WLCG computing model, with the T1 and T2 sites still being critical computing resources. Note that these plots do not take into account T2 disk resources, as these were not consistently tracked until 2018.

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 [2], a written set of responses to C-RSG scrutiny questions, and a face-to-face meeting with the ALICE computing coordinators. The ALICE report focuses on the computing resources used from January 2022 to December 2022, planned usage starting in April 2023 and the resource requirements for 2024. The ALICE experiment's 2022 resource utilization is summarised in Table 2.

In September 2022, the CERN Council approved measures to reduce power consumption significantly as a mark of social responsibility in view of the ongoing energy crisis. Two main decisions were taken by the Council: stopping the LHC operation in 2022 two weeks earlier than initially planned and reducing by 20% the data-taking in 2023. In agreement with the experiments, it was decided to postpone the heavy-ion (HI) running period from 2022 to 2023 and extend the 2023 run from 4 to 5 weeks. In addition, two days in November 2022 were dedicated to Pb-Pb collisions at low interaction rate for test purposes.

During 2022, the ALICE detector and data processing have been fully commissioned with *pp* collisions. Since the main 2022 Pb-Pb data-taking has been postponed, almost all *pp* physics has been accommodated in the O2 disk buffer allowing optimization of the *pp* processing plan. Two large processing calibration campaigns on the full statistics have been carried out, with the aim of improving the *pp* physics selection and validating data quality. During data-taking, the event-processing node (EPN) farm was used primarily for the online synchronous processing, and most of the asynchronous processing was performed on T0 and T1 sites. A few GPUs on the EPNs were also used to test the asynchronous reconstruction during the data taking, and the full EPN farm was used for asynchronous reconstruction of pp events at the end of data taking. Two large Run 3 simulations for *pp* and Pb-Pb collisions were carried out during 2022 and some specific Run 2 MC productions have been produced

ALICE		2022				
		C-RSG recomm.	Pledged	Pledged /C-RSG	Used	Used /C-RSG
CPU	Tier-0	471	471	100%	921	196%
	Tier-1	498	448	90%	505	101%
	Tier-2	515	517	100%	507	98%
	HLT	n/a	n/a	n/a	n/a	n/a
	Total	**1484**	**1436**	**97%**	**1933**	**130%**
	Others				*139*	*9%**
Disk	Tier-0	50.0	50.0	100%	46.6	93%
	Tier-1	55.0	49.7	90%	38.3	70%
	Tier-2	49.0	55.2	113%	40.3	82%
	Total	**154**	**154.9**	**101%**	**125.2**	**81%**
Tape	Tier-0	95.0	95.0	100%	61.4	65%
	Tier-1	63.0	71.8	114%	39.5	63%
	Total	**158.0**	**166.8**	**106%**	**100.9**	**64%**

*: Percentage taken with respect to the total C-RSG CPU recommendation

Table 2 Summary of ALICE resource usage in 2022 and the comparison with recommended and pledged resources. The CPU and storage resources are in units of kHS06-years and PB, respectively.

at low CPU capacity and low priority to finalise specific analyses for publication.

The average size of the compressed time-frame (CTF) event with with "Strategy A" (a non-lossy compression method) was found to be 30% greater the estimates based on MC. This increase comes mostly from a higher number of clusters in the time-projectiong chamber (TPC) not predicted by simulations. In addition, the validation of the more aggressive strategy for CTF compression ("Strategy B") was postponed until 2024 due to the lack of reliable sample of HI data to perform the physics validation.

The C-RSG congratulates ALICE for its progress in validating the Run 3 processing, simulation and analysis systems.

The T0 and T1 tape usage has increased by 10.5 PB from the last report in July 2022. In December 2022, tape usage was 61.4 PB (64% of the pledged capacity) at T0 and 39.5 PB (55% of the pledged capacity) at the T1 sites. The tape utilization is expected to remain significantly below the pledged capacity until the end of the Pb-Pb data-taking period in 2023.

A total of 155.8 PB of disk space has been deployed (including the disk buffers in front of the tape systems at T0 and the T1 sites). Of this, 125.2 PB are used (46.6 PB at T0, 38.3 PB at the T1 sites and 40.3 PB at the Tier-2 sites) corresponding to 80% of the total deployed capacity. The increase is mainly due to the *pp* data asynchronous passes. The 5.3 PB T1 disk deficit is compensated by a 6.2 PB surplus at Tier-2. To keep the disk occupancy stable on the T0 and T1, some of the data of the asynchronous passes was moved to the Tier-2 storage, allowing faster analysis.

The C-RSG congratulates ALICE for its effective use of opportunistic CPU resources. ALICE CPU usage during the reporting period January 2022 – December 2022 amounts to 921 kHS06-years at T0 (96% more than the pledged resources), 505 kHS06-years at the T1 sites (13% more than the pledged resources), and 507 kHS06-years (2% less than the pledged resources) at Tier-2 sites. Other opportunistic resources provide a further 139 kHS06-years including 20 kHS06-years from the recently added Wigner facility.

6 Resource Usage: ATLAS

The report on the ATLAS collaboration's computing usage is based on the information provided by the ATLAS experiment [3], written responses to questions by the C-RSG, and a face-to-face meeting with the ATLAS computing coordinators.

The ATLAS experiment has provided summaries of resource usage for the calendar year 2022. The average total CPU usage has been 7680 kHS06-years, well above the pledged 3509 kHS06-years, mostly from extensive use of unpledged HPC resources. In particular the Vega EuroHPC has provided more than 90% of HPC CPU resources. Though the majority of HPC resources are unpledged, some T1 and T2 centres have provided a portion of their WLCG pledge using HPC systems.

The overall ATLAS CPU efficiency on Grid sites has been consistently above 85%. The HPC systems provided lower efficiency, particularly during commissioning. Ongoing investigations are in place to address this issue. The fact that each HPC resource is unique prevents there being a universal solution that can be easily applied at all HPC centres. The successful utilization of HPC resources appears to correlate with the presence of ATLAS-knowledgeable personnel at the sites.

Disk and tape utilizations remain high with 313 PB used out of 304 PB pledged and 383 PB out of 400 PB, respectively. The data on disk is divided into three categories: persistent data, temporary data, which is stored disk while it is being processed but scheduled for deletion in the future, and cached data, eligible for deletion when space is needed. Such separation aims to optimise job scheduling and minimise data re-call from tape.

The new data formats, DAOD_PHYS and DAOD_PHYSLITE, are already in use for the derivation of calibrations and systematic uncertainties for both the reprocessed Run 2 data and the new Run 3 data in the current ATLAS software release (Release 22). It is anticipated that wider adoption of DAOD_PHYS and DAOD_PHYSLITE will reduce the overall number of different file formats used by ATLAS.

The ATLAS software porting to ARM CPUs has continued. The physics validation of the simulation software has been performed in 2022. The validation of event generation, reconstruction and derivation production is planned for 2023.

The ATLAS Google Cloud Project, started in mid-2022, has been providing 5–10k CPU cores for several months running a variety of workloads. The project goals include: R&D work to explore what cloud resources are best used for, the setup and operation of a "standard" ATLAS site, and the evaluation of the total cost of ownership (TCO) of the cloud-based site. The report on this project is expected later in 2023.

The recent developments in ATLAS software enable the use of available GPU resources, which allows utilization not only of NVidia CUDA-based GPUs but as well other technologies, such as Intel oneAPI-based graphics accelerators.

7 Resource Usage: CMS

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

During 2022, the CMS experiment made effective use of the resources provided. Table 4 shows an overview of the CMS resource usage in 2022 along with the comparison with the amount of resources recommended by the C-RSG and the pledged resources by the sites.

The aggregated CPU usage at T0, T1 sites, T2 sites and the HLT facility amounts to 4394 kHS06-years, which is about 69% higher than pledged. A significant part of the additional CPU resources comes

ATLAS		2022				
		C-RSG recomm.	Pledged	Pledged /C-RSG	Used	Used /C-RSG
CPU	Tier-0	740	544	74%	512	69%
	Tier-1	1300	1349	104%	1545	119%
	Tier-2	1588	1616	102%	3151	198%
	HLT	n/a	n/a	n/a	311	n/a
	Total	**3628**	**3509**	**97%**	**5519**	**152%**
	Others				*2161*	*60%**
Disk	Tier-0	32.0	32.0	100%	33.0	103%
	Tier-1	116.0	130.0	112%	139.0	120%
	Tier-2	142.0	142.0	100%	141.0	99%
	Total	**290.0**	**304.0**	**105%**	**313.0**	**108%**
Tape	Tier-0	120.0	120.0	100%	102.0	85%
	Tier-1	272.0	280.0	103%	281.0	103%
	Total	**392.0**	**400.0**	**102%**	**383.0**	**98%**

*: Percentage taken with respect to the total C-RSG CPU recommendation

Table 3 Summary of planned and used resources for ATLAS in the calendar year 2023. The CPU and storage resources are in units of kHS06-years and PB, respectively.

from beyond pledge contributions from T2 sites. CMS managed to increased the CPU efficiencies at the different tiers to 76% (from 70% in 2021).

CMS reports a significant saving of 10 kHS06-years provided by so-called VO-boxes at T0, with a full review together with CERN IT staff expected to result in additional savings. In addition, the CMS experiment used an additional 446 kHS06-years at HPC sites and other opportunistic resources, which is comparable to the opportunistics resources used during 2021.

Disk usage peaked at 179 PB amounting to 84% of pledged capacity and tape usage peaked at 319 PB amounting to 78% of pledged capacity. During 2022, CMS gained experience with the Rucio service, which was configured conservatively without overcommitting disk storage quotas. Work is ongoing to increase the disk space usage, with CMS aiming for 90% utilization.

The CMS collaboration has continued its efforts to merge datasets and reducing overlaps. The collaboration reports that the nano-AOD and mini-AOD datasets are in increasing use in physics analyses and expect this trend to continue. As the analyses using nano-AOD datasets are performed on local (non-Grid) clusters, the collaboration is unable to include the computational resources used in these analyses. The C-RSG would be interested in a more detailed report to understand how the increased popularity of the nano-AOD datasets will affect future Grid resource requests, in particular the need for growth in disk space.

8 Resource Usage: LHCb

The LHCb collaboration made efficient use of offline computing resources throughout 2022 for the production of simulated events, user jobs execution and for centralised production of analysis data formats for the various physics working groups. MC production, both full and fast simulations, dominated the CPU use (97%), with data analysis consuming the remaining CPU cycles. The CPU usage for MC production is, in turns, dominated by full simulations (84%), followed by fast (9%) simulations, while MC reconstruction accounts for the remaining 4%.

CMS		2022				
		C-RSG recomm.	Pledged	Pledged /C-RSG	Used	Used /C-RSG
CPU	Tier-0	540	540	100%	638	118%
	Tier-1	730	730	100%	1107	152%
	Tier-2	1200	1200	100%	2302	192%
	HLT	n/a	n/a	n/a	347	n/a
	Total	**2470**	**2470**	**100%**	**4394**	**178%**
	Others				*648.5*	*26%**
Disk	Tier-0	35.0	35.0	100%	24.0	69%
	Tier-1	83.0	83.0	100%	68.0	82%
	Tier-2	98.0	98.0	100%	87.0	89%
	Total	**216**	**216**	**100%**	**179.0**	**83%**
Tape	Tier-0	155.0	155.0	100%	140.0	90%
	Tier-1	260.0	260.0	100%	179.0	69%
	Total	**415.0**	**415.0**	**100%**	**319.0**	**77%**

*: Percentage taken with respect to the total C-RSG CPU recommendation

Table 4 Summary of planned and used resources for CMS in the calendar year 2022. The CPU and storage resources are in units of kHS06-years and PB, respectively.

Table 5 shows an overview of the LHCb resource usage in 2022 along with a comparison to the amount of resources approved by C-RRB and the resources pledged by the sites. The usage information is based on the WLCG accounting [9] records from the LHCb 2022 resource usage report [5]. The delivered CPU work for sites not included in the WLCG accounting (e.g., the HLT farm) is calculated from the LHCb Dirac accounting [10]. Pledged resources are extracted from CRIC [8].

Overall, the pledged resources to LHCb were significantly below the C-RRB approved values in CPU (by around 17% at T1), disk (by around 10% at T1 and 32% at T2) and tape (by around 17% at T1).

The WLCG accounting reports a CPU consumption during 2022 of 1721 kHS06-years at the T0, T1 and T2 sites, about 43% higher than the pledged capacity. In addition, the HLT farm provided about 271 kHS06-years. Opportunistic resources outside WLCG (clouds, HPCs, volunteer computing and institutional clusters) resulted in a contribution of 53 kHS06-years, or around 2.6% of the total CPU usage. LHCb continues its efforts to integrate and exploit HPC resources.

The collaboration used 50% of disk and 35% of the tape space allocated to it. This low occupancy is attributed to the change in the running schedule of LHCb. The size of the storage resources for the WLCG 2022 was set based on the assumption that 2022 would be a full data-taking year for the experiment. Most of the time, however, was allocated to commissioning of the various sub-detectors that resulted in a rather small fraction of collision data recorded in stable beam conditions. Given this reduction in storage utilization, the shortfall in pledged resources below the C-RRB recommendations had no impact on the LHCb Collaboration's physics programme.

9 Anticipated 2023 Resource Utilization

The collaborations are planning for data-collection according to the revised LHC running schedule for 2023. This is a somewhat smaller dataset for ATLAS and CMS than originally planned, while the ALICE dataset will be somewhat larger given the longer HI run. The LHCb dataset will be smaller due to the status of their detector during this running period.

LHCb		2022				
		C-RSG recomm.	Pledged	Pledged /C-RSG	Used	Used /C-RSG
CPU	Tier-0	189	189	100%	305	161%
	Tier-1	622	515	83%	676	109%
	Tier-2	345	333	97%	470	136%
	HLT	50	50	n/a	271	n/a
	Total	**1206**	**1087**	**90%**	**1722**	**143%**
	Others				*53*	*4%**
Disk	Tier-0	26.5	26.5	100%	10.5	40%
	Tier-1	52.9	47.8	90%	30.6	58%
	Tier-2	10.2	6.9	68%	4.0	39%
	Total	**89.6**	**81.2**	**91%**	**45.1**	**50%**
Tape	Tier-0	81.0	81.0	100%	29.8	37%
	Tier-1	139.0	116.0	83%	47.1	34%
	Total	**220.0**	**197.0**	**90%**	**76.9**	**35%**

*: Percentage taken with respect to the total C-RSG CPU recommendation

Table 5 Summary of LHCb resource usage in 2022 and the comparison with recommended 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 2022. The CPU and storage resources are in units of kHS06-years and PB, respectively.

This results in modest changes to the computing plans and no significant change in the required resources, except for LHCb Collaboration. It will use the resources available to continue its simulation production campaign to complete the collaboration's Run 2 analyses and use the data collected in 2022 and 2023 to understand the performance of the detector upgrades. The collaboration estimates that it will require 680 kHS06-years of CPU, 48 PB of disk and 129 PB of tape storage. These are well-within the resources that have already been made available to the collaboration for the 2023 year. The C-RSG recommends that the LHCb Collaboration confer with WLCG management on how to best arrange the provision of these resources.

Overall, the C-RSG believes the resources available to the collaborations in 2023 are well-matched to their physics programmes.

10 Background for 2024 Resource Requests

The computing resource requests for the 2024 WLCG year by the experiments are based on the current schedule for LHC operations for 2023 and 2024. For the ATLAS and CMS collaborations, this means that their required computing resources will increase significantly at T0 but with smaller incremental growth at the T1 and T2 sites. The ALICE experiment expects to require a relatively modest increase in computing resources given the deployment of its O2 facility and the resources already available to it. The LHCb experiment will not require additional resources in 2024 above what has already been approved by the C-RRB for the 2023 year given their expected running plans.

The justification for these requests and the recommendations of the C-RSG are discussed in the next sections.

ALICE		2022			2023		2024		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	471	471	921	541	541	600	111%	600
	Tier-1	498	448	505	572	506	630	110%	630
	Tier-2	515	517	507	592	567	650	110%	650
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**1484**	**1436**	**1933**	**1705**	**1614**	**1880**	**110%**	**1880**
	Others			139					
Disk	Tier-0	50.0	50.0	46.6	58.5	58.5	67.5	115%	67.5
	Tier-1	55.0	49.7	38.3	63.5	57.6	71.5	113%	71.5
	Tier-2	49.0	55.2	40.3	57.5	60.4	66.5	116%	66.5
	Total	**154.0**	**154.9**	**125.2**	**179.5**	**176.5**	**205.5**	**114%**	**205.5**
Tape	Tier-0	95.0	95.0	61.4	131.0	131.0	181.0	138%	181.0
	Tier-1	63.0	71.8	39.5	82.0	87.7	107.0	130%	107.0
	Total	**158.0**	**166.8**	**100.9**	**213.0**	**218.7**	**288.0**	**135%**	**288.0**

Table 6 ALICE resource requests for 2024 and C-RSG recommendations. For reference, the 2022 utilization and 2023 pledged resources are also given.

11 Resource Requests: ALICE

As noted previously, the 2022 Pb-Pb data taking was canceled due to the year-end technical stop starting two weeks earlier than initially planned. For the 2024 resource requirements the experiment expects to collect 140 days of *pp* collisions (including 9 days at low magnetic field) and either five weeks of Pb-Pb data in 2024 and 3 weeks of *p*-Pb data in 2025, or 4 weeks in each of the two years for Pb-Pb and *p*-Pb collisions. For its resource estimates ALICE assumes 28 days of Pb-Pb and 5 days of *pp* reference runs with a stable beam time of around 50%. Two additional short O-O and *p*-O runs have also been considered.

The 2024 runs will be the second time ALICE will have collected Pb-Pb data and the experiment will adopt a more aggressive compression strategy (Strategy B) for the “compressed time frames” (CTF) with an expected 33% reduction in event size over Strategy A. The C-RSG notes that the increase in the CTF event size (30%) due to the increase in the number of the clusters in the TPC volume has increased the overall size of the CTF by 30% as well. This offsets the savings in storage predicted for Strategy B by ALICE in previous scrutinies. The 33% reduction in CTF size arising from the use of Strategy B remains an estimate from the experiment and the C-RSG encourages ALICE to actively work on reducing the event sizes as this drives the resource increases requested by the experiment for 2024 and the need for additional resources for the O2 disk.

The 2024 *pp* runs will collect about 3.3×10^{12} collisions at full-field with a CTF output of 90 PB. After event selection 1.5% (increased from 1.2% in the last scrutiny) will be written to disk (1.4 PB). A further 8.1 PB of low field *pp* data will be generated and stored on disk. The 2024 Pb-Pb runs will produce 2.5×10^{10} events that correspond to 62 PB of CTF data. Two-thirds of the CTFs will be archived to tape at T0 facilities with the remaining one-third archived to tape at T1 sites. After the first asynchronous pass of the Pb-Pb data, 11 PB of AOD data (2 replicas) will be generated and stored to disk. The *p*-O and O-O runs will generate about 0.6 PB of CTF data and 0.7 PB of AODs (after two passes and assuming two replicas).

ALICE has achieved a four-fold increase in the efficiency of the MC generation relative to the CPU estimates based on Run 2 performance. The C-RSG congratulates ALICE on this achievement. This was offset by an increase in the number of MC events that could be generated with approximately the same amount of CPU allocation (by a factor of four). This increase in MC events also increases the disk and tape requirements for the experiment. ALICE decided to make use of increased code efficiencies to increase the number of minimum bias equivalent MC events relative to the number of collisions (i.e. an increase in the MC/event ratio compared to previous scrutinies).

The resources requested to support the 2024 data-taking and processing are given in Table 6. The 2023 resources in this table have been calculated after subtracting the Russian and JINR contributions. Allowing for carryover of unused 2023 resources (11.0 PB of disk and 3.8 PB of tape), the ALICE Collaboration's request represents a 22% increase for CPU, a 22% increase for disk and a 35% increase for tape (compared to 2023 pledges without Russian and JINR contributions).

Conclusions

The C-RSG finds the ALICE Collaboration 2024 computing requests to be appropriate to achieve its physics programme and recommends that they be approved.

The C-RSG makes the following observations and recommendations to the ALICE Collaboration:

ALICE-1 The C-RSG recommends that the ALICE Collaboration's resource requests for 2024 be approved.

ALICE-2 The C-RSG notes that the CTF event size has increased by 30% and that this has affected the tape resource and O2 buffer size during Pb-Pb data taking. THe ALICE Collaboration is encouraged to focus on CTF compression strategies.

ALICE-3 The C-RSG notes that the number of MC events that will be generated by ALICE has increased by a factor of four (without a corresponding increase in physics events). This four-fold increase in MC events does not have much effect on the CPU request because of significant improvements in the efficiency of the simulation code though it may drive storage growth. The C-RSG would appreciate an explanation of the need to increase the MC/event fraction relative to previous scrutinies and how this decision was made. The increase in MC events impacts not just CPU usage but also disk and tape resources and should not be solely driven by the availability of CPU resources.

ALICE-4 The C-RSG requests that ALICE further encourages their users to migrate to the new AOD format to take advantage of their faster throughput.

ALICE-5 The C-RSG notes that ALICE has successfully used spare T0 opportunistic resources, but encourages the collaboration to also increase its opportunistic resources elsewhere, i.e. on the other Tiers and on other HPC.

12 Resource Requests: ATLAS

Table 7 shows the resource requests submitted by the ATLAS Collaboration for 2024, the resources used in 2022 and the pledged resources for 2022 and 2023 along with the C-RSG recommendations. The experiment is asking for increases of 9.9% for CPU, 19.8% for disk and 25.1% for tape, in line with the expected increased output from the detector.

The ATLAS Collaboration makes the following assumptions on the running scenario for 2024: one full year of operation with an integrated luminosity of $L = 110\ \text{fb}^{-1}$, pile-up of 62 collisions/bunch crossing and a total running time of 5.2×10^6 s for pp runs. With these running conditions ATLAS

ATLAS		2022 C-RSG recomm.	2022 Pledged	2022 Used	2023 C-RSG recomm.	2023 Pledged	2024 Request	2024 req. /2023 C-RSG	2024 C-RSG recomm.
CPU	Tier-0	740	544	512	740	740	936	126%	936
	Tier-1	1300	1349	1545	1430	1520	1516	106%	1516
	Tier-2	1588	1616	3151	1747	1841	1852	106%	1852
	HLT	n/a	n/a	311	n/a	n/a	n/a	n/a	n/a
	Total	**3628**	**3509**	**5519**	**3917**	**4101**	**4304**	**110%**	**4304**
	Others			2161					
Disk	Tier-0	32.0	32.0	33.0	40.0	40.0	49.0	123%	49.0
	Tier-1	116.0	130.0	139.0	136.0	150.5	163.0	120%	163.0
	Tier-2	142.0	142.0	141.0	168.0	160.0	200.0	119%	200.0
	Total	**290.0**	**304.0**	**313.0**	**344.0**	**350.5**	**412.0**	**120%**	**412.0**
Tape	Tier-0	120.0	120.0	102.0	174.0	174.0	207.0	119%	207.0
	Tier-1	272.0	280.0	281.0	353.0	360.3	452.0	128%	452.0
	Total	**392.0**	**400.0**	**383.0**	**527.0**	**534.3**	**659.0**	**125%**	**659.0**

Table 7 ATLAS Collaboration resource requests for 2024 and C-RSG recommendations. For reference, 2022 and 2023 resources are also given.

expects to record 10.2×10^9 events in the main trigger stream and an additional 9.4×10^9 events in a delayed stream. The event size has been increased by 7% (i.e., by 100kB/event) to account for the 15% higher average pile-up.

The increase in average pile-up from 50 to 62 simultaneous collisions results in a simultaneous increase of the event reconstruction time by about 60% and data-taking trigger rate, bringing the main physics stream rate from 1.7 kHz to 2.0 kHz. To ensure prompt reconstruct of at least the main physics stream within 24 hours, the T0 CPU resources need to be increased by 26% with respect to the preliminary report.

ATLAS full simulation remains a major CPU consumer, but significant improvements of up to a 20% decrease in simulation time per event for the newest production configuration have been achieved for Run 3-like MC simulation compared to the software release employed in 2022.

Some uncertainty around the 2024 running conditions remains. If the readout limitations of the inner detector can be overcome, there is potential for ATLAS to reach an even higher average pile-up rate before the LHC machine limitations are reached. The HI schedule has also not yet been finalised, which could lead to an additional 5 PB of HI data at both the T0 and T1 sites. These possibilities have not been explicitly accounted for in the resource estimations.

Conclusions

The C-RSG considers that the ATLAS experiment's resource requests for 2024 shown in Table 7 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 the ATLAS Collaboration's resource requests for 2024 be approved.

CMS		2022			2023		2024		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	540	540	638	720	720	980	136%	980
	Tier-1	730	730	1107	800	916	930	116%	930
	Tier-2	1200	1200	2302	1350	1313	1600	119%	1600
	HLT	n/a	n/a	347	n/a	n/a	n/a	n/a	n/a
	Total	**2470**	**2470**	4394	**2870**	**2949**	**3510**	**122%**	**3510**
	Others			649					
Disk	Tier-0	35.0	35.0	24.0	45.0	45.0	54.0	120%	54.0
	Tier-1	83.0	83.0	68.0	98.0	96.8	122.0	124%	122.0
	Tier-2	98.0	98.0	87.0	117.0	109.7	149.0	127%	149.0
	Total	**216.0**	**216.0**	**179.0**	**260.0**	**251.5**	**325.0**	**125%**	**325.0**
Tape	Tier-0	155.0	155.0	140.0	228.0	228.0	320.0	140%	320.0
	Tier-1	260.0	260.0	179.0	316.0	303.7	380.0	120%	380.0
	Total	**415.0**	**415.0**	**319.0**	**544.0**	**531.7**	**700.0**	**129%**	**700.0**

Table 8 CMS Collaboration resource requests for 2024 and C-RSG recommendations. For reference, 2022 and 2023 resources are also given.

ATLAS-2 The C-RSG would like to understand whether the escalation in the adoption of the DAOD_PHYS and DAOD_PHYSLITE data formats would influence furture Grid resource requests, specifically in the context of disk space.

13 Resource Requests: CMS

The CMS Collaboration's resource requests for 2024 together with the C-RSG recommendations are shown in Table 8. The requested resources have increased with respect to the preliminary request provided to the C-RSG in autumn 2022 as reported in [7], reflecting differing assumptions in data-taking conditions, simulation and analysis requirements. CMS is asking for increases of 22% for CPU, 25% for disk and 29% for tape, in line with the expected increased output from the detector. This is based on the same running assumptions as employed by the ATLAS experiment.

The computing activities planned by CMS for 2024 are:

- Run-3 data-taking, with limited commissioning activities, with reprocessing of of the data taken in the beginning of 2024 and a full reprocessing of all the 2024 data when the final detector calibrations are derived;
- Run-3 MC event generation to support detector studies and development of reconstruction software, followed by MC simulation for physics analyses;
- High-Luminosity LHC (HL-LHC) MC generation to support the work associated with detector definition and software algorithm evolution; and
- data analysis, focussing on the data acquired during 2022, 2023 and 2024, as well as completion of a number of analyses using Run 2 data.

The significant increase of requested CPU resources at the T0 is required to cope with the higher pile-up conditions, and for the optimized compression of the raw data. The predicted increase of

LHCb		2022 C-RSG recomm.	2022 Pledged	2022 Used	2023 C-RSG recomm.	2023 Pledged	2024 Request	2024 req. /2023 C-RSG	2024 C-RSG recomm.
CPU	Tier-0	189	189	305	215	215	174	81%	174
	Tier-1	622	515	676	707	598	572	81%	572
	Tier-2	345	333	470	391	434	319	82%	319
	HLT	50	50	271	50	50	50	n/a	50
	Total	**1206**	**1087**	**1722**	**1363**	**1297**	**1115**	**82%**	**1115**
	Others			53					
Disk	Tier-0	26.5	26.5	10.5	30.3	30.3	30.6	101%	30.6
	Tier-1	52.9	47.8	30.6	60.5	54.7	61.2	101%	61.2
	Tier-2	10.2	6.9	4.0	11.6	7.9	11.8	102%	11.8
	Total	**89.6**	**81.2**	**45.1**	**102.4**	**92.9**	**103.6**	**101%**	**103.6**
Tape	Tier-0	81.0	81.0	29.8	91.0	91.0	117.1	129%	117.1
	Tier-1	139.0	116.0	47.1	157.0	133.7	133.3	85%	133.3
	Total	**220.0**	**197.0**	**76.9**	**248.0**	**224.7**	**250.4**	**101%**	**250.4**

Table 9 LHCb Collaboration resource requests for 2024 and C-RSG recommendations. For reference, 2022 and 2023 resources are also given.

integrated luminosity explains the increased tape and disk space requests at the T0. The increased disk space will also be used as buffer in case of unexpected data transfer issues with T1 sites.

The requests for resources at the T1 and T2 sites are largely driven by the larger Run-3 dataset, and the increased MC event samples required for the physics analyses of these data.

Conclusions

The C-RSG considers that the CMS Collaboration's resource requests for 2024 shown in Table 8 are necessary to achieve its approved physics programme, and recommends that the C-RRB arranges for pledges correspondingly.

The C-RSG makes the following request:

CMS-1 The C-RSG recommends that the CMS Collaboration's resource requests for 2024 be granted.

CMS-2 The C-RSG would like to understand if the increased popularity of the nano-AOD datasets will impacts the Grid resources requests, in particular for what concerns diskspace.

CMS-3 The C-RSG recognises the effort to increase disk utilization on the T0 and T1 sites, and encourages CMS to raise this to the 90% limit considered optimal.

CMS-4 The C-RSG requests that the CMS Collaboration continue reporting on efforts to pinpoint and reduce the causes behind the relatively low 2022 CPU efficiency.

14 Resource Requests: LHCb

The computing resource requests submitted by the LHCb Collaboration for 2024 are shown in Table 9. The resources used in 2022, the pledged resources for 2022 and 2023 along with the C-RSG recommendations are also listed in the table.

The LHCb Collaboration requires no increase in computing resources for 2024, compared to the resources already approved by the C-RRB for 2023. During 2023 the experiment will collect fewer physics data than anticipated. The LHC running time in 2023 has been shortened by 30% to reduce electricity costs, and an incident in the LHC vacuum system has affected the interface with the LHCb VELO detector volume and will prevent the VELO from operating fully closed in 2023. All in all, the recommended resources for 2023 are sufficient to provide the capacity required for 2024. In fact, the CPU requirements for 2024 are smaller than the pledged capacity for 2023. The C-RSG notes, however, that the level of pledged disk resources by the funding agencies for 2023 at the T1 and T2 levels is below 90% of the required capacity for 2024. The missing storage resources should be provided in time for the 2024 data-taking.

LHCb has provided the C-RSG with long-term projections for computing resource requirements through 2032. A large increase is expected for 2025 when resources will have to be doubled compared to 2024 requirements. A constant level of resources will then suffice until 2028, followed by an annual increase of 15-20%. These long-term estimates can be useful for funding agencies to plan the provision of resources.

LHCb is negotiating with Poland (NCBJ) and China (IHEP) for the incorporation of new T1 sites. In recent years, there has been a 10-20% deficit in T1 resources for the LHCb Collaboration that may worsen if the Russian T1 centre's resources cannot be utilised. The WLCG Overview Board has endorsed in December 2022 the bids of NCBJ and IHEP to become T1 sites for LHCb, setting an 18-month transition period to address all the needed steps. Both centers are already T2-D sites for LHCb, offering significant CPU work for simulation and physics analysis. Milestones in data management, storage, and network capabilities have been defined to be achieved in the time scale of one year.

In view of the large increase in resources forecasted for 2025, LHCb is investigating ways of reducing CPU consumption and data volume. The simulation work accounts for most of the CPU needs, while disk storage is dictated by the size of the analysis data format, and tape storage by the size of the detector raw data. Improvements in the handling of the out-of-time and in-time pile-up in the simulation, largely increased for Run 3, have been identified as an aspect where CPU consumption could be significantly reduced. Improving data compression can potentially reduce significantly the storage footprint. The C-RSG considers it important to devote appropriate effort to these optimization activities.

Conclusions

The C-RSG considers that the LHCb Collaboration's resource requests for 2024 conform to the experiment's physics programme. The following findings and recommendations are made:

LHCb-1 The LHCb Collaboration requests no increase in resources for 2024, compared to the resources already approved by the C-RRB for 2023. This decision is based on the fact that the resources allotted for 2023 will not be fully utilised due to changes in the physics programme following the incident affecting the VELO detector operations. However, less than 90% of the disk resources at the T1 and T2 levels approved by the C-RRB for 2023 have been pledged. The C-RSG encourages the funding agencies to provide the LHCb collaboration with the 2023 approved resources in time for the 2024 data-taking.

LHCb-2 The C-RSG notes that the zero growth of LHCb computing resources in 2024 will be followed by a significant increase (of almost 100%) in 2025, as per the long-term projections shared by the LHCb Collaboration. The experiment considers it acceptable to stagger the growth over 2024 if it makes the procurement process for the sites easier to manage.

LHCb-3 The C-RSG recommends that the LHCb Collaboration allocates the necessary effort to carry out the activities identified to decrease CPU consumption in simulations (such as optimizing pile-up handling) and reduce storage footprint (such as implementing more aggressive data compression).

15 Comments and Recommendations

The C-RSG makes the following overall observations and recommendations.

ALL-1 The C-RSG requests that the collaborations report in subsequent scrutinies high-level summaries of the manner in which their disk space is utilised, how they have optimised the allocations and how they propose to allocate the requested space. This summary should identify the space used for: i) persistent datasets (and differentiating between primary copies and additional replicas), ii) cached storage that is used to hold datasets for short periods of time, and iii) buffering space used, for example, to transfer data from one storage media to another. This information will assist the C-RSG in better understanding the pressures on disk utilization and help identify best practice.

ALL-2 The C-RSG recognises the significant efforts of the collaborations and WLCG to identify additional T1 sites. As noted in earlier scrutiny rounds, the addition of such sites is essential to mitigate the effects on the physics programmes of the loss of any existing T1 facility.

References

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

[2] ALICE Collaboration. ALICE computing resources usage January to December 2022, resource requirements for 2023 and 2024. Submitted 21 February 2023.

[3] ATLAS Collaboration. ATLAS computing status and plans – Report to the C-RSG – April 2023. Submitted 21 February 2023.

[4] CMS Collaboration. Spring 2023 scrutiny – CMS utilization report 2022 and resource request for 2024. Submitted 21 February 2023.

[5] LHCb Collaboration. LHCb computing resource usage in 2022. *LHC-PUB-2023-002*, submitted 21 February 2023.

[6] LHCb Collaboration. LHCb computing resources: 2024 requests. *LHC-PUB-2023-001*, submitted 21 February 2023.

[7] Computing Resources Scrutiny Group. C-RSG autumn 2022 report. *CERN-RRB-2022-078*, 10 October 2022.

[8] WLCG CRIC. Computing Resource Information Catalogue. http://wlcg-cric.cern.ch.

[9] EGI Foundation. European Grid Initiative. https://www.egi.eu/federation.

[10] LHCb. LHCb DIRAC Accounting Portal. http://lhcb-portal-dirac.cern.ch.