CERN–RRB–2021-107
6 October 2021

Autumn 2021 Report of the Computing Resources Scrutiny Group

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

1 Introduction

The Computing Resources Scrutiny Group (C-RSG) is charged with reviewing the computing requests for the four Large Hadron Collider experiments, ALICE, ATLAS, CMS and LHCb, and making recommendations on allocations of resources given their approved physics programs. This report summarizes the autumn 2021 C-RSG review of the computing resources being used for 2021, agreed-upon pledged resources for 2022 and preliminary projections for the 2023 fiscal year. Unless otherwise noted, the start of each reporting period is 1 April.

The Autumn C-RSG review focused on the computing activities for the 2021 and 2022, and preliminary computing plans for 2023. At this stage, the resources for the 2022 year are well-defined and there is a firm schedule for LHC data-taking. This allowed for a well-informed discussion between the scrutiny group and each collaboration. The C-RSG comments and makes recommendations regarding the 2022 and 2023 plans, anticipating the formal requests for 2023 resources that will be made in Spring 2022.

The inputs to this review are information from previous years, estimates of resource needs for 2022 and 2023 from each collaboration, subsequent questions from the C-RSG to each collaboration and remote face-to-face discussions with the computing representatives for clarifications.

2 C-RSG membership

The chair thanks the C-RSG members for their commitment and expert advice, and the collaborations' computing representatives for their cooperation with the C-RSG and this review process. The group membership was updated in Spring 2021, as were the referee teams for each experiment. Thanks are due to the CERN management for its support and to the C-RSG scientific secretary, H Meinhard (CERN), for ensuring the smooth running of the group.

3 Interactions with the experiments

The experiments were asked to submit their reports by August 24$^{\text{th}}$, 2021. The C-RSG thanks the experiments for the submission of their detailed documents [1–4]. The dedicated face-to-face meetings with experiment representatives were particularly helpful and appreciated by the C-RSG, even though they were all done remotely.

Specific sets of C-RSG referees were assigned to review the ALICE and LHCb requests. As usual, by agreement with ATLAS and CMS management, a single team of C-RSG referees scrutinized the

ATLAS and CMS reports to ensure a consistent approach. The referees subsequently reported to the full C-RSG, which then developed the recommendations contained in this report.

In anticipation of the Spring 2022 scrutiny, the C-RSG asks the experiments to submit their documents by February 22th, 2022. The C-RSG reminds the experiments to respond in their submissions to the general and experiment-specific recommendations in this report, as has now become practice.

4 Overview of 2022 plans and 2023 resource requirements

The 2022 computing resources recommended in Spring 2021 by the C-RSG for the four collaborations appear to be appropriate to meet their scientific needs, as there were virtually no adjustments in the resources needed by the collaborations to achieve their approved physics programs in the 2022 year.

As has been the case in the past, the experiments continue to secure significant opportunistic CPU resources. Although these are not pledged, the C-RSG has encouraged this practice, though recognizing the risk that reliance on these resources creates.

Each experiment is making significant effort to innovate and improve on the utlization of the resources they have available, although in some cases this is limited by personnel resources. The C-RSG continues to monitor future developments, and where appropriate encourages the experiments to keep these in mind as they plan their development activities for the next few years. An example of such a development is the increasing deployment of GPUs for more general-purpose computing.

The projected computing resources for 2023 are summarized in the following sections. The C-RSG finds that they are consistent with the longer-term physics goals for each experiment, especially in the case of ALICE and LHCb where significant changes in the computing models have been made for Run 3.

5 Resource requests: ALICE

The C-RSG report is based on the usage and resource requests provided by the ALICE experiment, a written set of responses to C-RSG scrutiny questions, and a remote meeting with the ALICE computing coordinators. The C-RSG thanks the ALICE team members for providing detailed responses to the C-RSG recommendations from the Spring 2021 scrutiny report and for their timely responses to questions from the C-RSG in this scrutiny.

The new baseline for Run 3, approved in June 21, assumes three years of data collection (2022-2024) and the closure of the caverns 1st February, 2022, the option of early cavern closure in November 2021 being no longer possible. As ALICE computing resources for Run 3 are driven by the Pb-Pb data collection, the first Pb-Pb run is expected to be undertaken in November 2022 and to last 24 days (with an additional 143 days of pp data collection in 2022).

The development and provisioning of the new 02 system continues in preparation for Run 3 (with 20 PB of the 80 PB disk buffer now installed). The C-RSG notes that the O2 system will have significant compute resources (16k CPU cores and 2k GPUs) and welcomes the report that asynchronous reconstruction can be performed on GPUs as well as CPUs. Reporting of the usage of the O2 will be important as ALICE enters commissioning and then Run 3 operations. Working to integrate monitoring tools that are accessible through the standing accounting portals will help in tracking the usage of this opportunistic resource and we encourage ALICE to work with CERN to implement such an accounting system.

The C-RSG recognises that, unlike x86 systems, there is no standardized benchmark to account for this GPU usage. This situation is unlikely to change in the near term given the priority for updating the CPU benchmarks. It is, however, important that this information is captured and reported to the C-RSG as Run 3 operations begin. From discussions with ALICE, and in view of the fact that these O2 GPU resources are homogeneous, the C-RSG notes that events per second for workflows with and without GPUs could be used as an initial scale factor to place GPU usage in the same units as current CPU reporting.

For 2022, ALICE will start its first data-taking with redesigned and redeployed main detectors along with a completely new readout and data processing system. The objective of the activities during the initial start-up of LHC beam operations, such as the beam commissioning and ramp-up, will be to validate the entire data collection and processing chain with data and to prepare for efficient operation during the first pp and Pb-Pb campaigns of Run 3 (later in November 2022). The full pp period is planned to commission and gain experience with the upgraded detector, as well as to collect data samples that enable a complete validation of the calibration, processing and analysis chain.

The updated running conditions, with a higher expected pp pile-up profile, does not affect the ALICE experiment which is planning to collect pp at a fixed 500kHz interaction rate. The computing and storage resources requests remains unchanged since the last report with an increase from “priority” needs of 18% for CPU (divided 17%, 19%, 19% across T0, T1 and T2, respectively), 21% for disk (divided 38%, 14%, 14% across T0, T1 and T2, respectively) and 73% for Tape (divided 89% and 53% across T0 and T1, respectively). The C-RSG notes that the 9.3 PB (mainly CTF data) tape storage, requested in 2021 for contingency against an early closure of the cavern, is still needed for commissioning (pp ramp-up and 1 MHz pp data), but only in 2022 at the start of Run 3. The associated disk storage will be removed as soon as the commissioning analyses are completed, likely in 2023.

For 2023, ALICE estimates remain unchanged since the Spring 2021 report, with an increase of 15% for CPU at all tiers, 15% for disk (divided 16%, 15%, 14% across T0, T1 and T2, respectively), 30% for tape (divided 33% and 25% across T0 and T1, respectively). The beginning of 2023 will be dedicated to two data reconstruction passes and simulations following the 2022 Pb-Pb data-taking (through June 2023) in parallel with new pp data-taking (mid-May 2023) and a new Pb-Pb campaign (mid-October 2023). This will result in CTF production of 44 PB (Strategy B compression) stored on disk and archived on tape (divided 2/3 and 1/3 at T0 and T1, respectively). All CTF data collected in the previous year will be removed from all disk buffers before the new Pb-Pb data-taking period starts.

Nevertheless, the current T1 under-pledged disk (85% of requested), even if compensated by the T2 over pledge, is a source of concern for ALICE for its asynchronous data processing, where T1s will store 1/3 of the CTF data, process it and keep a copy of the resulting AODs (2 passes). The critical point will be reached in April/May 2023 when the second year of pp data-taking starts. Other proposed solutions such as a partial redirection of asynchronous processing to T2 have been rejected due to the need to balance CPU and storage resources during processing as well as a concern about potential network bottlenecks (i.e. there are not sufficent T0 to T2 network bandwidth guarantees).

The storage capacity needs for 2023, assuming commissioning data removal (9.3 PB) in 2023 and taking into account additional removal of the Run 1 and Run 2 AODs and MCAODs after data format conversion together with low popularity data deletion, is compatible with the ALICE disk storage estimates. The increase in tape is consistent with the required need for archive resources due to the large data collection from the Pb-Pb run.

Conclusions and Recommendations

The CRSG compliments the ALICE experiment on its continued development of the O2 system in preparation for Run 3.

ALICE-1 The C-RSG requests that for future scrutinies ALICE report the usage of the O2 facility as an opportunistic resource. The C-RSG asks that the experiment work with CERN to define

ALICE		2021		2022			2023	
		C-RSG recomm.	Pledged	Request	2022 req. /2021 C-RSG	C-RSG recomm.	Preliminary Request	2023 req. /2022 C-RSG
CPU	Tier-0	471	471	471	100%	471	541	115%
	Tier-1	498	412	498	100%	498	572	115%
	Tier-2	515	481	515	100%	515	592	115%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**1484**	**1364**	**1484**	**100%**	**1484**	**1705**	**115%**
	Others							
Disk	Tier-0	45.5	45.5	50	110%	50.0	58.0	116%
	Tier-1	53.3	45.4	55	103%	55.0	63.0	115%
	Tier-2	44.8	50.4	49	109%	49.0	56.0	114%
	Total	**143.6**	**141.3**	**154**	**107%**	**154.0**	**177.0**	**115%**
Tape	Tier-0	86.0	86.0	95.0	110%	95.0	126.0	133%
	Tier-1	57.0	59.8	63.0	111%	63.0	79.0	125%
	Total	**143.0**	**145.8**	**158.0**	**110%**	**158.0**	**205.0**	**130%**

Figure 1 Alice resource requests and C-RSG recommendations for 2021 and 2022, and resource estimates for 2023.

a mechanism by which the 02 facility can be monitored using standard tools such as the CRIC and EGI accounting portals.

ALICE-2 With benchmarks for GPU usage unlikely to be available for the next scrutiny, the C-RSG asks that when reporting on the O2 GPU usage ALICE scale the GPU-to-CPU resources using the ratio of the number of reconstructed events per second with CPU only to workflows with the joint CPU+GPU (for asynchronous reconstruction). This will ensure that the GPU and CPU usage are placed on a comparable scale.

6 Resource requests: ATLAS

The ATLAS report is based on the information provided by ATLAS [2], written responses to questions by C-RSG, and a remote face-to-face meeting with the ATLAS computing coordinators.

While no detailed usage report was requested from the experiment, ATLAS provided a brief overview on 2021 resource usage. The average total CPU usage has been 4126 kHS06-years, well above the pledged 3265 kHS06-years, mostly from T1 and T2 beyond-pledge resources. In addition, ATLAS reported a CPU usage of 1432 kHS06-years of opportunistic resources (HLT farm, HPC centres and BOINC).

ATLAS noted that in some cases procuction jobs run slower than the published HS06 number indicates. Investigations are ongoing, in collaboration with the relevant sites, and the C-RSG will be interested to learn of their finding in Spring 2022.

Disk utilisation remains high with 249 PB used out of 272 PB pledged. The tape usage approaches the pledge value on the T0 site (with 84 PB out of a total of 95 PB occupied), while on the Tier-1s a planned deletion campaign has freed space (179 PB out of 242 PB occupied) in preparation for the Run 2 reprocessing and Run 3 data-taking.

ATLAS makes the following assumptions on the running scenario for 2023: one full year of operation with an integrated luminosity of $L = 100\ \text{fb}^{-1}$, pile-up of 50 collisions/bunch crossing and a total running time of $t_{run} = 6 \times 10^6$ s for pp and 1.2×10^6 s for heavy ion running. With these assumptions ATLAS expects to record 10.2×10^9 events in the main trigger stream and an additional 9.6×10^9 events in a delayed stream.

ATLAS computing requests for 2023 (Figure 2) are driven by the following planned computing activities:

- processing and reprocessing of the Run 3 data taken in 2012,
- user analysis of Run 2 and Run 3 data and production of the necessary Monte Carlo (MC) samples,
- simulation of MC samples for 2023, and starting the MC sample production for 2024,
- production of derived data formats of data and MC samples for physics analysis,
- physics studies for HL-LHC, and
- processing of heavy-ion data and MC simulation.

The collaboration aims to produce 30 billion MC events with the 2023 running configuration, with 40% of the events produced in full simulation, and the remaining 60% as fast simulation. These numbers were both roughly 50% for 2022, and ATLAS expects the fraction of fully-simulated events to decrease to 25% by the end of Run 3. Significantly, ATLAS has been able to reduce the CPU burden associated with simulating radiation damage in its silicon tracker.

ATLAS expects to use 10 PB to record heavy-ion data at the end of the 2023 data-taking period. Although the definite plans for this run are not yet available, ATLAS believes that this volume will not change with the various scenarios.

The estimated increases for CPU, disk and tape capacity at the T0 are considerable (at 35%, 25% and 45%, respectively), but are in line with the expected increase in integrated luminosity.

CPU is estimated to increase by 18% on T1/T2 sites and 35% at T0 over the resources agreed in 2022. The CPU requests for T1/T2 sites have been lowered by 700 kHS06-years, taking into account the expected resources from beyond-pledge contributions at these sites, as well as HPC and BOINC usage.

The C-RSG notes, as in previous scrutinies, that the growth in requested CPU resources by ATLAS has consistently been greater than that of CMS, even though the two experiments have comparable physics programs. Though the detailed computing models differ for the experiments, the C-RSG views it appropriate for ATLAS to consider how to manage the growth in CPU resources more aggressively, especially as the estimated increases for 2023 exceed the "flat budget scenario" of approximately 10-15% annual increases that has been used by the WLCG for longer-term planning.

ATLAS estimates for 2023 disk capacity represent an overall increase of 18% over those for 2022. The request for 40 PB of storage at the T0 includes 28 PB of Grid storage, which has been subtracted from the storage requests for the T1 and T2 sites.

In order to store Run 3 pp and heavy-ion raw data, Run-3 MC production and AODs from Run 3 and Run 2 reprocessing, ATLAS estimates an increase by 30% on T1 sites and 45% at T0 over the tape resources pledged in 2022.

ATLAS		2021		2022			2023	
		CRSG recomm.	Pledged	Request	2022 req. /2021 C-RSG	C-RSG recomm.	Preliminary Request	2023 req. /2022 C-RSG
CPU	Tier-0	525	525	550	105%	550	740	135%
	Tier-1	1170	1243	1356	116%	1300	1536	118%
	Tier-2	1430	1497	1656	116%	1588	1877	118%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**3125**	**3265**	**3562**	**114%**	**3438**	**4153**	**121%**
	Others							
Disk	Tier-0	29.0	29.0	32.0	110%	32.0	40.0	125%
	Tier-1	105.0	116.3	121.0	115%	116.0	136.0	117%
	Tier-2	130.0	127.2	148.0	114%	142.0	167.0	118%
	Total	**264.0**	**272.5**	**301.0**	**114%**	**290.0**	**343.0**	**118%**
Tape	Tier-0	95.0	95.0	120.0	126%	120.0	174.0	145%
	Tier-1	235.0	241.2	272.0	116%	272.0	353.0	130%
	Total	**330.0**	**336.2**	**392.0**	**119%**	**392.0**	**527.0**	**134%**

Figure 2 ATLAS resource requests and C-RSG recommendations for 2021 and 2022, and resource estimates for 2023.

Conclusions and Recommendations

ATLAS-1 The C-RSG finds the ATLAS resource projections justified, based on the currently known 2023 parameters. The C-RSG expects that these parameters will become more firm by Spring 2022, resulting in more refined resource requests for 2023.

ATLAS-2 The C-RSG applauds ATLAS for the effort on decreasing the penalty occurred from the simulation of detector radiation damage and expects to see further reports on progress in this area.

7 Resource requests: CMS

As in previous years, the C-RSG scrutinised the information provided by the CMS collaboration, elaborating a set of questions and asking for written responses. Furthermore, more in-depth information has been gathered by a remote face-to-face meeting with the CMS computing coordinators. The discussion has focused on the preliminary resource request for 2023, but also has covered ongoing activities, along with transition to the upcoming data-taking in 2022.

Figure 3 summarises the CMS computing resources for 2021 and 2022, and the corresponding recommendations for 2022 by the C-RSG. The same figure also shows the preliminary estimates for 2023.

CMS computing efforts are presently focused on preparation for the Run 3 data-taking. One critical operational aspect during 2021 has been the transition to Rucio as the core data management system, which CMS reports as a success. CMS has also been successful in accessing tape resources using CTA as a replacement for CASTOR.

CMS anticipates that the T0 2023 resources will be primarily used for prompt reconstruction of events. The C-RSG notes that CPU needs are higher than those of 2022 due to the increased level of pileup

CMS		2021 C-RSG recomm.	2021 Pledged	2022 Request	2022 2022 req. /2021 C-RSG	2022 C-RSG recomm.	2023 Preliminary Request	2023 2023 req. /2022 C-RSG
CPU	Tier-0	500	500	540	108%	540	720	133%
	Tier-1	670	764	730	109%	730	800	110%
	Tier-2	1070	1151	1200	112%	1200	1350	113%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**2240**	**2415**	**2470**	**110%**	**2470**	**2870**	**116%**
	Others							
Disk	Tier-0	30.0	30.0	35.0	117%	35.0	45.0	129%
	Tier-1	77.0	76.0	83.0	108%	83.0	98.0	118%
	Tier-2	92.0	96	98.0	107%	98.0	117.0	119%
	Total	**199.0**	**202**	**216.0**	**109%**	**216.0**	**260.0**	**120%**
Tape	Tier-0	120.0	120.0	155.0	129%	155.0	228.0	147%
	Tier-1	230.0	219.0	260.0	113%	260.0	316.0	122%
	Total	**350.0**	**339**	**415.0**	**119%**	**415.0**	**544.0**	**131%**

Figure 3 CMS resource requests for 2021 and 2022, C-RSG recommendations for 2022, and estimates of resources needed for 2023.

and to the larger integrated luminosity expected in 2023. That also motivates increases in disk and tape needs. Similar considerations hold for the seven CMS T1 centres, where event reprocessing and MC production will take place with high priority, and T2 sites, mostly devoted to user analysis, MC production, and production of updated analysis data formats.

In summary, the C-RSG notes that CMS requests for 2023 an overall increase of 16% in CPU, 20% in disk and 26% in tape with respect to 2022, with most of the relative increase taking place at the T0. These appear to be justified based on the expected 2022 and 2023 data-taking and physics analysis requirements.

Conclusions and Recommendations

CMS-1 The C-RSG finds the CMS resource projections justified, as they are based on the currently known 2023 parameter seems justified. The C-RSG expects that these parameters will become more firm by Spring 2022, resulting in more refined resource requests for 2023.

CMS-2 The C-RSG recognizes CMS for its continued effort to support non-X86 architectures, as this is expected to increase the robustness of its software as well as prepares for the future hardware landscape.

CMS-3 The C-RSG supports CMS plans to adapt lossy compression algorithms for heavy-ion data, algorithms that promises a high level of compression without sacrificing the accuracy of the physics data.

CMS-4 The C-RSG applaud CMS for its continued effort to decrease the size of the analysis data and especially its plans to have 50% of all analysis to be on nano-AODs by the end of Run 3.

LHCb		2021 C-RSG recomm.	2021 Pledged	2022 Request	2022 req. /2021 C-RSG	2022 C-RSG recomm.	2023 Preliminary Request	2023 req. /2022 C-RSG
CPU	Tier-0	175	175	189	108%	189	361	191%
	Tier-1	574	470	622	108%	622	1185	191%
	Tier-2	321	292	345	107%	345	657	190%
	HLT	50	50	50	100%	50	50	100%
	Total	1120	987	1206	108%	1206	**2253**	**187%**
	Others		50	50			50	
Disk	Tier-0	18,8	18,8	26,5	141%	26,5	42,2	159%
	Tier-1	37,6	33,9	52,9	141%	52,9	84,4	160%
	Tier-2	7,3	6,1	10,2	140%	10,2	16,2	159%
	Total	63,7	58,8	89,6	141%	89,6	**142,8**	**159%**
Tape	Tier-0	43,8	43,8	81	185%	81,0	131,6	162%
	Tier-1	75,9	64,7	139	183%	139,0	227,7	164%
	Total	119,7	108,5	220	184%	220,0	**359,3**	**163%**

Figure 4 LHCb C-RSG recommended and pledged resources for 2021, requested resources and C-RSG recommendations for 2022, and preliminary estimates of resources for 2023.

8 Resource requests: LHCb

The assessment of the preliminary LHCb computing resource requests for 2023 is based on the document provided by the experiment [4] and discussions between the C-RSG and LHCb computing management.

Figure 4 shows the LHCb resource projections for 2023, as well as the C-RSG recommendations for 2022 and the pledged resources in 2021. At the Tier-1 and Tier-2 levels, a 10-20% deficit on pledged resources in 2021 with respect to the C-RSG recommendations is observed. Due to the delayed start-up of the LHC Run 3 to 2022, this shortage has no impact on the experiment provided the requested resources for 2022 are delivered on time. In 2023, LHCb projects a significant increase in computing resources, in relation to those recommended for 2022, by about 90% for CPU and 60% for storage. A large amount of additional computing resources is requested to store and process the data to be recorded in 2023 and to fulfill the growing data simulation needs of the experiment.

CPU requirements are dominated by MC simulation, consuming close to 80% of the total LHCb CPU budget. Simulations for Run 3 data accounts for more than 90% of the total simulation needs, while the rest will be devoted to simulating Run 1 and Run 2 conditions for the remaining analyses. CPU needs for Run 3 simulations have been calculated following the parameters established in the LHCb Upgrade Computing Model Technical Design Report [5], back in 2018. The fraction of events produced with full/fast/parametric simulations (40%/40%/20%, respectively) is assumed to be the same as in 2022, and no speedup of the simulation is expected for 2023. Compared to the Run 2 simulation, a similar event processing time is anticipated for the Run 3 full simulation (due to improvements in the simulation software), while the fast simulation is expected to be four times slower. The C-RSG recommends updating the estimates using actual measurements from realistic Run 3 simulations.

The C-RSG notes that the CPU work required for Run 3 simulation in 2023 is three times larger than that requested for 2022. This is a consequence of the LHCb simulation production model where, following the experience in Run 2, the production of simulated events needed to represent a given data-taking year is expected to be performed over six years. During the first year, the number of events produced is set to 50% of the nominal value (4.8×10^9 events per fb^{-1} of collision data collected), while in the five subsequent years the nominal value is expected. Hence in 2023, the production of the nominal number of simulated events for 2022 conditions and 50% for 2023 conditions is foreseen, while in 2022 only 50% of the nominal number of simulated events are planned to be produced. Given this pile-up effect in the simulation needs, the C-RSG encourages LHCb to increase the fraction of events produced with fast simulation.

The projected increase of data storage is driven by the integrated data volume expected to be recorded in 2023. An additional 120 PB of tape storage will be needed to store two replicas of the raw data, and close to 50 PB of disk storage will be required to keep on disk two replicas of the latest data reconstruction pass for physics analysis. No significant reduction of these requirements is expected, neither by data format reductions nor reductions in the number of replicas. Two replicas on tape are needed for data redundancy.

The estimates of the required computing resources for the heavy-ion run, planned for the end of 2023, is affected by large uncertainties on the data-taking conditions, detector performance and reconstruction throughput. Nevertheless, the projected resources are much smaller than those required for the pp run. References for the calculations are taken from the 2016 p-Pb run. The reconstruction of the raw events is expected to take place on the online farm.

The projected CPU requirements for 2023 rely on the assumption of extensive software performance improvements, such as a more efficient event processing framework and data model, a substantial speed-up of the simulation and moving most of the user analysis activity to organized workflows. The C-RSG recommends careful monitoring of the evolution of the software performance upgrades and regular reporting in future scrutinies.

The C-RSG notes that no significant advances in software performance have been reported. There has been no significant progress on other C-RSG recommendations, such as increasing the access and exploitation of HPC opportunistic resources or optimizing the disk usage by implementing a service that purges unused data from disk. The C-RSG is concerned about the reported lack of personel resources for computing development and optimization activities.

Conclusions and Recommendations

LHCb-1 The MC simulation for Run 3 data dominates the CPU requirements of the experiment. CPU resource calculations for Run 3 simulation have been performed using preliminary estimates taken from the upgrade LHCb computing proposal released in 2018. The C-RSG recommends refining these estimates using more realistic event simulation times from actual Run 3 conditions. Likewise, the C-RSG encourages LHCb to explore increasing the fraction of events produced with fast simulation.

LHCb-2 The CPU cycles obtained from opportunistic HPC resources continue to be marginal. The C-RSG restates its recommendation of devoting appropriate effort to access and exploit HPC facilities. The availability of HPC resources for HEP computing is growing and its contribution to LHCb data analysis would mitigate the pressures the new computing model places on WLCG CPU resources.

LHCb-3 LHCb faces tremendous computing challenges in Run 3. While recognizing the very significant reductions the collaboration has made to reduce the CPU requirements, the C-RSG is concerned about an apparent shortage of personnel available for computing activities. Some of the C-RSG's past recommendations could not be addressed apparently due to personnel constraints. In particular, software improvements required to further contain the growth of computing resources requirements involve a substantial expert workforce. The C-RSG encourages the experiment and the funding agencies to increase the support for LHCb computing development.

9 Comments and General Recommendations

The C-RSG makes the following overall observations and recommendations regarding the status of computing resources for the 2022 and 2023 year.

ALL-1 The C-RSG thanks all four experiments for the responses to the Spring 2021 recommendations. It appreciates the constructive discussions the group had with the computing coordinators of each experiment.

ALL-2 The C-RSG notes that computing resource allocations are usually made at the beginning of the year, providing a well-defined and committed resource for the subsequent 12 months. The group is concerned that reductions in the WLCG MOU commitment mid-year compromise the experiments' ability to complete their physics programs.

ALL-3 The use of GPUs as a source of T1 and T2 computing is increasing, but the benchmarks necessary to account for the GPU computing resources have not been established. The C-RSG recommends that an interim approach to GPU accounting be established to allow a consistent assessment of the total resources available to each collaboration.

ALL-4 Given that the C-RSG sees a future where processor architectures may evolve, the group suggests the experiments continue to assess the long-term outlook for architectures alternative to the current X86 platform (e.g., ARM, Power, RISC-V) and anticipate porting implications.

ALL-5 The C-RSG notes the increasing importance of networking in view of the changing computing models for the LHC experiments.. The group requests that experiments report on any risks to their computing plans from possible network constraints.

References

[1] ALICE Collaboration. ALICE Computing Resources Usage February to July 2021, Resource Requirements for 2022 and Estimates for 2023, submitted 10-Aug-2021.

[2] ATLAS Collaboration. ATLAS Computing Status and Plans – Report to the C-RSG September 2021, submitted 25-Aug-2021.

[3] CMS Collaboration. Fall 2021 Scrutiny – Preliminary Resource Request for 2023 - CMS Offline Software and Computing, submitted 25-Aug-2021.

[4] LHCb Collaboration. LHCb Computing Resources: preliminary 2023 requests. *LHCb-PUB-2021-008*, submitted 24-Aug-2021.

[5] LHCb Collaboration. LHCb Upgrade Computing Model Technical Design Report. *CERN/LHCC 2018-14*, 26 November 2018.