CERN–RRB–2017–125
September 2017

# Computing Resources Scrutiny Group

C Allton (UK), V Breton (France), G Cancio (CERN), A Connolly (USA), M Delfino (Spain), F Gaede (Germany), J Kleist (Nordic countries), D Lucchesi (Italy, Chairman), H Meinhard (CERN, scientific secretary), J Templon (Netherlands)

## 1 Introduction

The fall 2017 report analyses the resource requests for 2019 for the four main LHC experiments ALICE, ATLAS, CMS and LHCb. It evaluates the ALICE 2018 requests that were not recommended for adoption by the April RRB and confirms the ATLAS, CMS, and LHCb resources approved by RRB. The 2017 usage is not discussed experiment-by-experiment; only a general overview is given.

## 2 CRSG membership

Membership of the C-RSG has changed for this scrutiny. Dugan O′Neil, the Canada representative, stepped down and, at the moment, no replacement has been nominated. The group thanks Dugan for his dedication and important contributions over the years of his membership.

The chairperson thanks the C-RSG members for their dedication and commitment to this activity and the experiments representatives for their availability and their cooperation. Thanks are also due to the CERN management for their support and to the scientific secretary, H Meinhard (CERN), for ensuring the smooth running of the group.

## 3 Interactions with the experiments

The ATLAS, CMS and LHCB were asked to submit their resource requests by September $4^{th}$, ALICE by August $28^{th}$ in order to have sufficient time to evaluate the 2018 request. The C-RSG thanks the experiments for the timely submission of their detailed documents [2–5]. The group also thanks the computing representatives of the experiments for their availability, their responses to our questions, subsequent requests for further information, and for their helpful discussions.

By agreement with ATLAS and CMS management, a single team of C-RSG referees scrutinized the ATLAS and CMS reports and requests to ensure a consistent approach.

For the April 2018 RRB we ask the experiments to submit their documents by March $5^{th}$ 2018.

## 4 Resource Usage

The C-RSG reports on the general usage of resources up to September 2017, a complete analysis experiment-by-experiment will be provided in the Spring report.

At the writing of this report the machine's live time has been lower than 2016. The computing resources were used efficiently by the experiments as briefly reported in each specific sub-section. In

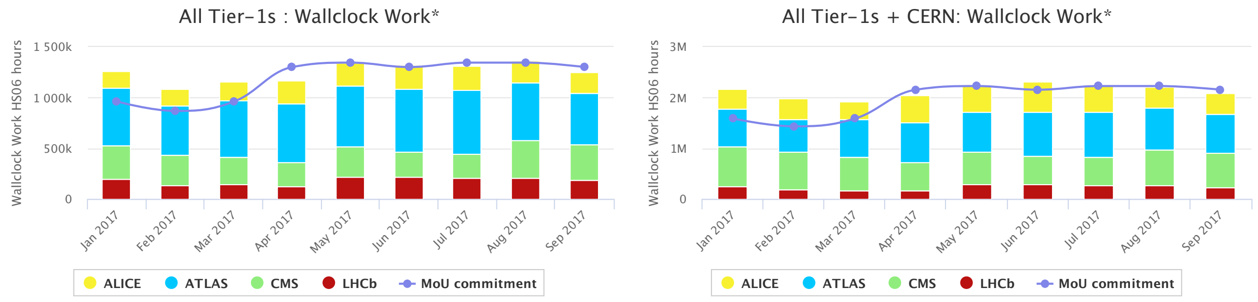


**Figure 1** Wallclock time as reported by the batch system normalized by benchmarked HEPSPEC06 power of a given CPU resource and multiplied by the number of processors as function of time. On the left the Tier-1 and on the right Tier-1+Tier-0 usage [6].

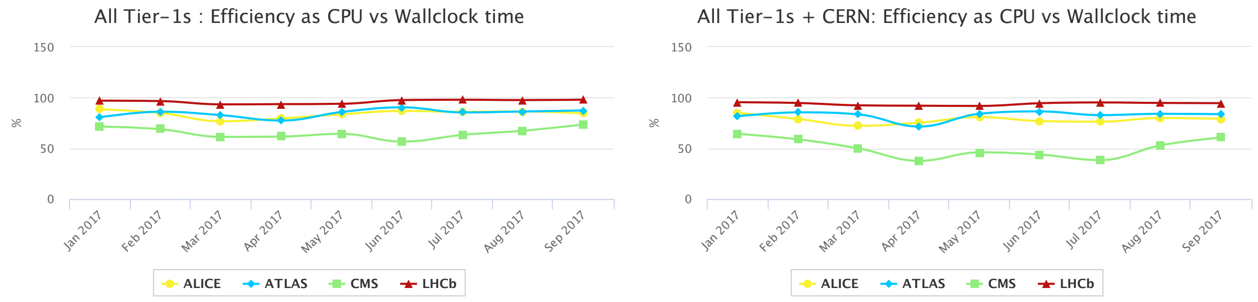


**Figure 2** CPU efficiency as function of time on the left for the Tier-1 and on the right Tier-1+Tier-0 for the four experiments [6].

figure 1 the used wallclock time is shown as function of the time for all the Tier-1 (left) and all the Tier-1+CERN Tier-0 (right). Different colors are used for different experiments and the solid line represents the resources pledged by the Funding Agencies assumed available as of April $1^{st}$.

During the April 2017 RRB a drop in CPU efficiency, defined as CPU time divided by wallclock time, was reported for CMS. Since Jul 2017, CMS CPU efficiency has increased substantially. In a many cores environment, the definition of CPU efficiency is not something that is straightforward to estimate. Both experiments, ATLAS and CMS, are studying in detail the causes of any inefficiencies in the workflow and analyses (as briefly reported below). Figure 2 shows the CPU efficiency for each experiments for all Tier-1 (left) and all Tier-1+CERN Tier-0 (right) as function of time. LHCb shows a very high efficiency since it proceeds in single core mode. CMS issues are discussed 4.2.

## 4.1 ATLAS CPU efficiency

CPU efficiency for ATLAS is around 80% over Run-2 and therefore around 10% lower than in Run-1, following the introduction of multicore jobs. ATLAS has started investigating the workflows behind the efficiency losses and has initiated activities to address them.

A few examples are reported here. Efficiency losses in due to the single-threaded initialisation of multi-core workflows will be addressed in the short term by using pre-configured checkpointed ATLAS software instances in virtual machines, and in the long term by rewriting the Athena configuration system. Another improvement will address using shared writing modules allowing processes to write concurrently to the same file thus avoiding expensive merging operations. ATLAS has also set up a task force for enabling high pile-up pre-mixing avoiding serialised thus expensive run time mixing of minimum bias Monte Carlo events. Derivation generation, being a new workflow in Run-2 is another

| RRB year | pp/$10^6$ s | HI/$10^6$ s | pp pileup |
|---|---|---|---|
| 2017 | 7.8 | - | 35 |
| 2018 | 7.8 | 1.2 | 35 |
| 2019 | - | - | - |

**Table 1** Assumptions on live time in seconds for LHC machine in Run 2, 2017, 2018 and 2019. The final column gives the anticipated average pileup for ATLAS and CMS during pp running for each year.

area where tuning is still required, given its high requirements on I/O. Each of these changes may yield efficiency improvements in the order of 1.5%-2%.

ATLAS is prioritising implementation of these improvements driven by the effort required as well as the impact and effort regarding physics validation.

### 4.2 CMS CPU efficiency

A dedicated CMS collaboration task-force is undertaking an in depth review of CPU and workflow efficiencies. From figure 2 it is evident that CPU efficiency started recovering from July 2017. As for ATLAS, the CPU efficiency drop is a consequence of multicore jobs. Multiprocessing codes now utilize eight cores to better match available compute systems. Development continues on a complex batch processing system that utilizes pilot jobs to assign work across a heterogenous set of resources. Aggressive sniping of pilot jobs that do not correctly initialize or start has improved the utilization of the package payloads. In addition CMS is working to optimize the throughput, for example they are aggressively back porting GEANTV improvements to GEANT4 to increase the throughput of the simulations. This will be the first vectorized application they will use.

There remains a trade between CPU efficiency and storage (i.e. the computational efficiency cost of moving data between systems vs the financial cost of additional copies of the data). Current workflows do not always map to the heterogeneous capabilities of the Tier-1 and Tier-2 facilities.

We note that CMS uses Agile at CERN Tier-0, which has definition of CPU efficiency (normalized to the uptime) that will result in an underestimate of CPU efficiency when compared to systems that use wall clock time as the normalization factor.

## 5 Assumptions for resource requests

The assumptions used by the experiments to determine the resources needs are based on the LHC running conditions [7] and on the updated approved schedule [1]. Table 1 reports the anticipated LHC beam live times updated to the latest official schedule [1]. Looking at 2016 data taking, the machine efficiency in 2017 and 2018 for pp runs is assumed to be 60%. The final column gives the average pileup (average number of collisions in each beam-crossing) for ATLAS and CMS pp collisions. The LHC luminosity is expected to be $1.9 \times 10^{34}\ \mathrm{cm}^{-2}\ \mathrm{s}^{-1}$ for 2017 and 2018. The current LHC schedule foresees no data in 2019 and 2020. The Heavy Ion (HI) run at the end of 2018 is of 24 days with an expected efficiency around 60%

## 6 ALICE

This report is based on the original submission by ALICE on August $28^{th}$ 2017, a face-to-face meeting on October $4^{th}$ 2017 with the computing coordinators that was attended remotely by the physics

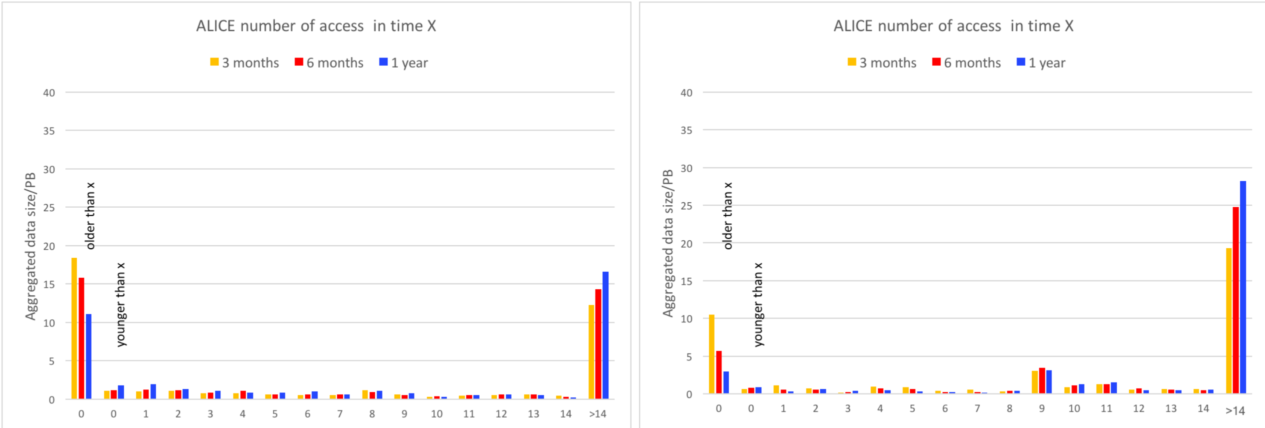


**Figure 3** Volumes of data versus number of accesses in 3-, 6- and 12-month periods for ALICE in March (left), and in July (right) for comparison. For each period X, data created in that period but not accessed is in the second bin. The first bin is for data created before the period began and not accessed during that period.

coordinator and additional information provided by mail until October $19^{th}$ in response to reviewers questions and comments.

The August report contains request for 2018 and 2019. ALICE has performed a lot of work since April resulting in a reduction of the resources requested. The total amount of tape needed in 2018 is lowered by 7 PB due to

- change of gas mixture in TPC (from Ar to Ne) that reduces spurious clusters
- improved HLT compression of TPC data from 5.2 to 7.2

The new working conditions cause also a small reduction (0.8 PB) of the disk space needed at Tier-0. C-RSG congratulates the collaboration for this important result and accepts the request of reducing the amount of tape and disk space for 2018.

Regarding the requests for 2018 which were not approved in April 2017 (CPU at Tier-2, disk at Tier-1 and Tier-2), C-RSG appreciates the effort made by the collaboration to reduce the request. Since April the ALICE computing group has performed an aggressive clean-up activity of the disk space at Tier-1 and Tier-2 as shown in figure 3, recovering 7 PB. Thanks to that, the April request have been reduced: from 32 PB to 30.5 PB at Tier-1 and from 41 PB to 35.2 PB at Tier-2. C-RSG recommends ALICE to keep the disk space monitored to avoid unused data on disk.

The CPU at Tier-2 went from 438 kHS06 to 398 kHS06. This is the result of a re-evaluation of the computing power needed to reconstruct one event, taking into account the actual processing time observed in 2016. The reduction is applied to Tier-2 budget since there ALICE has the largest deficit.

Before and during the face-to-face meeting several questions were raised by the C-RSG to understand the actual needs of the experiment triggered also by the large request for 2019.

The initial document and the experiment requirements were significantly modified following the exchanges between ALICE management and C-RSG reviewers. Between the document submitted on August $28^{th}$ and the last document submitted on October $11^{th}$ a number of relevant numbers have changed in particular the ratio of Monte Carlo to raw events has been increased from 0.18 to 0.3 for Pb-Pb and from 1. to 1.4 for p-p.

The determination of the appropriate number of Monte Carlo events per real collision goes beyond the mandate of the C-RSG and has to be discussed with the LHCC.

Table 2 provides the last figures communicated by ALICE management in the e-mail dated October $11^{th}$, together with C-RSG recommendations for 2018. Regarding the requests for 2018 which were not approved in April 2017, C-RSG accepts ALICE request for disk at Tier-1s, already pledged by Funding Agencies. CPU and disk space at Tier-2 are heavly affected by the Monte Carlo production, and C-RSG does not have enough information to make a recommandation.

For 2019, our discussions with ALICE have resulted in their revision downwards in their CPU request. This was due to an error they found in some of the assumptions made in their resource calculations. However, C-RSG remains unable to make a recommendation for 2019 until we have had guidance from LHCC on the number of Monte Carlo events per experimental event because, obviously, this parameter plays a crucial role in the overall resource requirements. We postpone any evaluation of the 2019 resources until our Spring scrutiny since this will allow the interaction with the LHCC.

In order to make the scrutiny process more efficient, the C-RSG requests that all relevant parameters and formulae used to justify the resource requests are provided in all future requirements documents.

| Resource | Site | 2017 CRSG | 2017 Pledge | 2018 Spring request | 2018 Fall request | 2018 CRSG | 2018 Growth over 2017 | 2019 Fall request | 2019 Growth over 2018 |
|---|---|---|---|---|---|---|---|---|---|
| CPU (kHS06) | T0 | 292 | 292 | 350 | 350 | 350 | 20% | 430 | 23% |
| | T1 | 256 | 235 | 306 | 307 | 307 | 30% | 375 | 22% |
| | T2 | 366 | 280 | 438 | 398 | - | - | 475 | - |
| Disk (PB) | T0 | 22.4 | 22.4 | 27 | 26.2 | 26.2 | 17% | 30.7 | 17% |
| | T1 | 25.4 | 21.8 | 32 | 30.5 | 30.5 | 40% | 35.8 | 14% |
| | T2 | 31.4 | 22.7 | 41 | 35.1 | - | - | 39.7 | - |
| Tape (PB) | T0 | 36.9 | 36.9 | 55 | 49.1 | 49.1 | 33% | 49.1 | 0% |
| | T1 | 30.9 | 30.6 | 41 | 40.9 | 40.9 | 34% | 40.9 | 0% |

**Table 2** ALICE 2018 and 2019 requests. T0 means T0+CAF. Growth over 2017 is calculated as (2018 CRSG - 2017 Pledge)/2017 Pledge. Similarly, growth over 2018 is calculated as (2019 request - 2018 CRSG)/2018 CRSG

# 7 ATLAS

The ATLAS report is based on the information provided by ATLAS [3], written responses to initial and follow-up questions by C-RSG, and an in-person meeting with the ATLAS computing coordination team.

While no usage report has been requested in this scrutiny round, ATLAS usage numbers for January to July 2017 are inline with the requests approved by RRB for disk and tape. The average CPU utilisation exceeds the pledge values by over 50% (2353 kHS06 vs. 1539 kHS06), which shows the extent of the availability of beyond pledge resources (from HPC, Grid and Cloud) to ATLAS. As already highlighted by C-RSG, this continued reliance remains a risk for the experiment.

The ATLAS 2018 resource needs are based on an average pileup of 33 and 7.8M seconds of p-p data taking. Since the 2017 RRB meeting, there have been no significant changes communicated for LHC schedule and running conditions, therefore the ATLAS computing resource requests for 2018 (Table 3) are unchanged from the ones agreed at that meeting.

In 2019, all Run-2 data will be reprocessed with a final set of updated condition data. A large Monte Carlo production activity involving three major campaigns will span into 2019, with 3.9B events to be produced in the first 4 months of 2019. While this activity will provide a minimal scenario allowing for early searches for new physics, an increase in statistics is needed for precision measurements to reduce statistical fluctuations hence allowing for the full exploitation of the Run-2 physics data, requiring around 16.8B MC events to be produced between 2019 and 2020. Another 6B events are expected to be produced for allowing taking advantage of new MC generator developments and parameter tunings based on first Run-2 measurements. In addition, a reprocessing of 2018 heavy ion data may take place in 2019. Resources will also be required for the Phase-II detector TDR design studies, with around 3.5B events to be produced in 2020. In total, ATLAS expects to generate 30.2B events in 2019/2020, with a target of producing at least 50% in 2019. Around 40% (8.5B) of the events will be generated using fast simulations. These allow for reducing MC generation time by an order of magnitude, and ATLAS is looking to increasing the percentage of fast simulations in order to optimise CPU resources.

ATLAS bases its computing resource requirements on similar event sizes and processing times as in 2018. While AOD size has been successfully reduced by 30% at the end of 2016, this has not yet been extended to DAOD's, where similar gains could be achieved, although a task force has now started working on this area. Given its large impact on disk resources, C-RSG encourages continuing this work.

ATLAS is investigating how to further optimise their workflows. For simulation, preliminary tests indicate that Geant4 compilation tuning may yield performance gains in the order of 10-15%, without affecting physics results.

ATLAS expect to keep one full copy of the AOD's from the latest processing on disk as processing from tape alone has been showing limitations at several Tier-1's. These limitations are due to the tape retrieval rate which falls below expectations in particular when sites are serving multiple VO's and thus multiple concurrent activities that compete for disk buffer and tape drives. Given that processing from tape is a new workflow that may result in significant disk space savings, C-RSG encourages to continue working closely with the Tier-1's in order to identify, understand and address the underlying shortcomings.

During LS2, resources freed up on ATLAS Tier-0 and HLT will be used for offline processing and simulation generation, thus no Tier-0 resource increases are requested for 2019. Due to infrastructure work, the HLT availability is expected to be around 50% in 2019 and 20% in 2020. ATLAS still needs to discuss the configuration of CERN CPU resources as enabling hyper-threading and moving from 4GB to 2GB of memory per core may yield an increase of 25% more HS06 for Monte-Carlo production.

| Resource | Site | 2017 RRB | 2017 Pledge | 2018 Request | 2018 RRB | 2018 Growth over 2017 | 2019 Request | 2019 Growth over 2018 |
|---|---|---|---|---|---|---|---|---|
| CPU (kHS06) | T0+CAF | 404 | 404 | 411 | 411 | 2% | 411 | 0% |
| | T1 | 921 | 808 | 949 | 949 | 17% | 1057 | 11% |
| | T2 | 1125 | 982 | 1160 | 1160 | 18% | 1292 | 11% |
| Disk (PB) | T0+CAF | 25 | 25 | 26 | 26 | 4% | 27 | 0% |
| | T1 | 68 | 69 | 72 | 72 | 4% | 88 | 23% |
| | T2 | 83 | 78 | 88 | 88 | 13% | 108 | 23% |
| Tape (PB) | T0+CAF | 77 | 77 | 94 | 94 | 23% | 105 | 0% |
| | T1 | 188 | 174 | 195 | 195 | 12% | 221 | 13% |

**Table 3** ATLAS 2017 and 2018 resources, and 2019 requests

Resource requests for 2019 (Table 3) show an overall increase of 11% for CPU at Tier-1s and Tier-2s, 23% for disk at Tier-1s and Tier-2s, and 13% for tape on Tier-1s over the 2018 approved requests. These requests are driven by the resources required for the main activities during LS2 that will be centered around MC event production for finalising Run-2 analyses, preparing for Run-3 data taking as well as studies for the HL-LHC upgrade. Similar to previous years, the resource requests do not include around 700 kHS06 that ATLAS expects to obtain from opportunistic resources.

Overall, based on 2018 requests approved by RRB, the 2019 requests from ATLAS are overall consistent with nominal growth limits as requested by the funding agencies: While requests for Tier-1 and Tier-2 disk capacity exceed the maximum 15% annual growth rate foreseen for disk resources, CPU increases are well below the maximum 20% annual growth limit. For ATLAS, a reduction of disk space would impact MC production for precision measurements required for the full physics exploitation of Run-2 data.

For 2020, ATLAS cannot yet quantify concretely their resource needs. While the Monte Carlo needs are considered to be well understood, a decision needs to be taken whether Run-2 and MC data is to be reprocessed during that year, and in addition the availability of Tier-0 and HLT CPU resources is yet not full known. ATLAS anticipates larger CPU needs in 2020 than in 2019.

# 8 CMS

The CMS report is based on the resource requests provided by the CMS experiment [4], written responses to a set of scrutiny questions, and an in-person meeting with the CMS computing coordinators.

The focus of this scrutiny round is the resource allocations for 2018, and 2019, and for initial projections of resources for 2020. The assumptions of CMS for 2018 running conditions are 7.8 Ms live seconds and an average pileup of 35. Given that these conditions are consistent with the values used in the 2017 RRB and that there have been no official updates to the LHC schedule, the 2018 CMS resource requests have not changed since the last CRSG meeting (see Table 4) other than an increase of Pb-Pb live seconds (from 0.7 Ms to 1.2 Ms) resulting in additional 2PB of tape space requested at the Tier-0.

In 2019, data for the years 2016-2018 will be reprocessed to produce a final set of Run-2 analysis results papers. While there are no current plans to reprocess the 2015 data this remains under discussion by the experiment. Parked data will not be included within the legacy reprocessing unless analyses on the scouting datasets show features that require further exploration. Reprocessing is expected to start in May 2019 (after establishing the calibration factors for the data taking periods) and is expected to be completed by the end of the year.

A large Monte-Carlo (MC) production run will be undertaken in conjunction with the reprocessing with a goal of simulating between 25B and 35B events. The initial component of the MC run, 25B events, is expected to be completed by the end of 2019 (in time for the science conferences) with 2020 used to finish the production tails of the MC run for the final Run-2 papers.

The HLT farm and Tier-1 resources will be used to support the MC generation (from Geant4 simulations to reconstruction). Under the assumption of 80% availability of the HLT farm, it is estimated that an additional 50 kHS06 will be required in the Tier-1 to complete MC generation. Tier-0 and Tier-2 resources will perform the reprocessing of the data and undertake the science analyses. The use of Tier-0 resources for data reprocessing substantially reduces the need for Tier-2 CPU resources.

CMS has shifted most of its high level studies to utilize the MiniAOD data to reduce disk usage. Disk resource requirements are minimized at the Tier-2 facilities by reducing the number of AOD copies from 1 to 0.7 (with a further 0.3 on Tier-1 disks).

Tier-1 tape requests for 2019 represent a 22% increase, which is above a flat budget, due to the need to complete the legacy reprocessing prior to the start of Run-3. The increase in needs for new tape

| Resource | Site | 2017 RRB | 2017 Pledge | 2018 Spring request | 2018 RRB | 2018 Fall request | 2018 Growth over 2017 | 2019 Fall request | 2019 Growth over 2018 |
|---|---|---|---|---|---|---|---|---|---|
| CPU (kHS06) | T0 | 423 | 423 | 423 | 423 | 423 | 0% | 423 | 0% |
| | T1 | 600 | 515 | 600 | 600 | 600 | 17% | 650 | 8% |
| | T2 | 850 | 791 | 900 | 900 | 900 | 14% | 1000 | 11% |
| Disk (PB) | T0 | 24.6 | 25 | 26 | 26 | 26 | 4% | 26 | 0% |
| | T1 | 57 | 45 | 60 | 60 | 60 | 33% | 68 | 13% |
| | T2 | 68 | 53 | 70 | 70 | 70 | 32% | 78 | 11% |
| Tape (PB) | T0 | 70.5 | 71 | 97 | 97 | 99 | 39% | 99 | 0% |
| | T1 | 175 | 133 | 205 | 188 | 188 | 41% | 230 | 22% |

**Table 4** CMS 2018 resources as approved by RRB, and 2019 requests. T0 means T0+CAF

beyond 2019 is expected to be significantly lower once a tape cleaning campaign has been completed in 2020 after completion of the Run-2 analysis.

Resources will also be required for Upgrade Studies for Phase 2, for a Trigger TDR in 2019, and possibly a Physics TDR at a later date. The resource requirements for the TDR work are modest, 50M events, but will require a large fraction of the events to be in full format. Heavy Ion (HI) analysis is not accounted for in the current resource requests as it undertaken using resources outside of those provisioned through the RRB. Reprocessing of 2018 heavy ion data may take place in 2019.

There remain a number of areas of uncertainty in these resources estimates that will require reconsideration and updates in the next scrutiny round. There is some concern about the reliability and capability of EOS and P5-to-networking to support the Tier-0 and HLT data reconstruction. This is being evaluated and will be revisited in the spring 2018 scrutiny. The scheduled availability of the HLT farm in 2019 has not been finalized and the assumption of 80% availability may not be possible to achieve. The current assumption of a Pile-Up value of 35 may be an underestimate. With a conservative estimate of a 20% increase in Pile-Up this would translate to an uncertainty of between 10 and 20 % in the CPU resources requests. There is significant uncertainty in the computational performance of the reconstruction algorithms for the high granularity calorimeter (HGCAL) (though the impact of this is likely to be at the 1-2% level for CPU resources). Delays in the completion of the calibration may extend the reprocessing and simulation schedule beyond the baseline described here.

Resource requests for 2019 (Table 4) show an overall increase of 10% for CPU at Tier-1s and Tier-2s, 12% for disk at Tier-1s and Tier-2s, and 22% for tape on Tier-1s over the 2018 approved requests. These requests are driven by the resources required for the reprocessing of the Run-2 data, MC event production for finalising Run-2 analyses, preparations for Run-3 data taking as well as studies for the HL-LHC upgrade.

CMS has proposed a low risk approach for the completion of the Run-2 reprocessing and associated Monte-Carlo campaign by utilizing available CPU, disk, and tape resources across the Tiers. Their timeline is such that delays at a level of 10-20% in completing the calibration, or the availability of resources for the reprocessing of data or the MC generation should not impact preparations for Run-3.

Overall, based on 2018 requests approved by RRB, the 2019 requests from CMS are consistent with a flat budget, as requested by the funding agencies. While the Tier-1 requests for tape represent a 22% increase, which exceeds the nominal 15% growth anticipated for a flat budget, the disk and CPU requests are both below this nominal growth. We note that a reduction in the tape request could impact the CMS experiments ability to complete the reprocessing and analyses prior to the end of LS-2. If

| Resource | Tier | 2017 CRSG | 2017 Pledge | 2018 Spring request | 2018 Fall request | 2018 Fall CRSG | 2018 Growth over 2017 | 2019 Fall request | 2019 Growth over 2018 |
|---|---|---|---|---|---|---|---|---|---|
| CPU | T0 | 67 | 67 | 81 | 88 | 88 | 31% | 93 | 6% |
| (kHS06) | T1 | 207 | 199 | 253 | 253 | 253 | 27% | 271 | 7% |
| | T2 | 116 | 147 | 141 | 141 | 141 | -4% | 152 | 8% |
| Disk | T0 | 10.9 | 10.9 | 12.0 | 11.4 | 11.4 | 5% | 14.2 | 25% |
| (PB) | T1 | 22.1 | 20.9 | 24.5 | 24.5 | 24.5 | 11% | 27.9 | 14% |
| | T2 | 4.7 | 3.3 | 5.8 | 5.7 | 5.7 | 21% | 6.8 | 18% |
| Tape | T0 | 25.2 | 25.2 | 36.4 | 33.6 | 33.6 | 33% | 35.0 | 4% |
| (PB) | T1 | 43.3 | 42.0 | 61.5 | 45.6 | 45.6 | 9% | 50.9 | 4% |

**Table 5** LHCb resource requests and CRSG recommendations. 2018 growth is calculated with respect to 2017 pledges. 2019 growth is calculated with respect to 2018 Fall CRSG recommendation.

necessary, the 22% increase in tape for the Tier-1 (in 2019) could be split across the Tier-0 and Tier-1 facilities.

For 2020, CMS did not specify its computational needs other than noting that the MC generation may extend into 2020 and that they will undertake a large tape cleaning exercise. They note, however, that they do not expect a large increases in resources.

While CMS continues to see the provisioning of fewer resources compared to the recommendations of the CRSG this under-provisioning has been substantially reduced compared to the situation in 2017 (to less than 10% across CPU, disk, and tape) due to the delivery of additional resources through 2017 and the availability of overpledge and opportunistic resources.

As with the ATLAS experiment we request that CMS document and report separately their use of overpledge and opportunistic resources including the source of these resources (e.g. grid farms, and HPC centers).

# 9 LHCb

The LHCb report is based on the resource requests provided by the LHCb experiment [5], written responses to a set of scrutiny questions, and an in-person meeting with the LHCb computing coordinators.

The focus of this scrutiny session was to evaluate adjustments to the resource requests for 2018 (previously evaluated in Spring 2017), to start evaluation of the 2019 requests (to be completed in Spring 2018), and to receive initial projections of resources for 2020. This last point had to be postponed, as LHCb was not able to provide 2020 projections at this time.

The assumptions of LHCb for 2018 running conditions are based on the current accelerator schedule and 60% LHC efficiency, giving 7.8 Ms of proton physics, 1.2 Ms of heavy ions and 0.2 Ms of fixed target (SMOG) data taking. LHCb operates with a luminosity leveling scheme, hence their data accumulation rate is primarily determined by this live time.

LHCb has an excellent internal resource usage monitoring and accounting system, which is cross-checked with the WLCG accounting. Information on resource usage is then fed into their computing model.

LHCb reports the successful completion of their micro-DST data format development. This allows an eventual reduction in stored data, as ancillary information used for micro-DST physics verification will be removed.

Given that the efficiency of LHC operation has been lower than expected during 2017, less new data will be accumulated. As a consequence, using the resources already pledged and installed for 2017, some operations which had been postponed to 2019 can be accommodated during 2017 and 2018, benefiting the physics output of the experiment. In addition, the accumulation of less raw data during 2017 feeds forward into the estimates of resources needed in later years.

For these reasons, LHCb has reassessed their 2018 resource request, as shown in Table 5.

The overall impact of improvements in the computing model (including completion of the micro-DST) and the smaller data accumulation in 2017 has a significant impact on 2018 Tape resources. The updated Tape requests are 8% and 26% less for Tier-0 and Tier-1s, respectively, compared to the C-RSG recommendation in Spring 2017. Small decreases in the Disk requests at Tier-0 and Tier-2s have also been made. CPU resources are dominated by simulation needs to support physics analyses, with growth rates similar to previous years.

LHCb is planning to reinforce their development activity on analysis preservation, in order to be ready for the many analyses which will be completed in the next years. In order to do this efficiently, a single development platform will be used which needs to have agile access to many different types of datasets. CERN is a very good choice for this, and therefore LHCb has increased their 2018 Tier-0 CPU request by 7 kHS06, that is 9% more compared to the C-RSG recommendation in Spring 2017.

The 2019 initial resource requests can also be seen in Table 5. Small growth in CPU is dominated by simulation needs, with smaller contributions from re-stripping and user analysis. Growth in Disk resources is requested in order to accommodate the outputs of the analysis activities and the full re-stripping of the Run2 data. A very small growth in Tape resources is due to the shutdown of the LHC.

C-RSG thanks LHCb for the reassessment of the 2018 resources and recommends granting them as requested. The 2019 resource estimates seem well founded, and should be taken as guidance by the funding agencies. A final review of the 2019 resource request will be made by C-RSG in Spring 2018.

C-RSG congratulates LHCb on the excellent operation of the computing model/resource monitoring/resource estimation feedback loop.

LHCb is strongly requested to submit their estimates for 2020 resources in early 2018, to allow C-RSG to review them and provide guidance at the Spring 2018 RRB meeting.

# 10 Comments and recommendations

- The C-RSG, together with LHCC and WLCG, proposes to update the procedures for reviewing experiments' requests, in order to allow more time for interactions with the LHCC. The proposed procedures are the following:
  - Prior to the year N Autumn RRB, experiments make initial resource requests for year N+2 and, if possible, a preview of needed resources in year N+3;
  - A six month review of the requests starts, involving C-RSG, LHCC and experiments, with a possibility to review physics scope with the LHCC;
  - At the year N+1 Spring RRB, C-RSG makes final recommendations for year N+2 resources and, if possible, gives guidance regarding the year N+3 resources;
  - At the year N+1 Autumn RRB, Funding Agencies confirm pledges for year N+2.

- The C-RSG appreciates the continued work by the experiments on increasing the computational efficiency of their workflows and simulations, and on reducing the CPU and disk resources required to addressing the increase in the luminosity of the LHC.
- The C-RSG continues to strongly support software engineering development and recommended that sufficient effort is funded to support this activity in the collaborations in particular now that the demand for resources is exceeding the expectations.
- The C-RSG would request that as part of future resources assessments that the experiments provide a proposed mitigation strategy to address changes in the assumed running conditions for the experiment (e.g. pile-up or luminosity) at the level of a 20% increase. These assessments should assume a scenario of no additional compute or storage beyond Tier-0 tape.

## References


[1] CERN. LHC Schedule.

[2] ALICE Collaboration. ALICE 2016 computing resources usage, 2017.

[3] ATLAS Collaboration. ATLAS Computing Status and Plans. Report to the C-RSG, October 2017. September 2017.

[4] CMS Collaboration. CMS Software/Computing Resource Utilization & Request 2017-2019, Spring 2017. February 2017.

[5] LHCb Collaboration. LHCb Computing Resources: 2019 requests and reassessment of 2018 requests. *LHC–PUB–2017–019*.

[6] EGI. EGI Accounting Portal.

[7] I Bird et al. Update of the Computing Models of the WLCG and the LHC Experiments. *CERN-LHCC-2014-014*.