CERN–RRB–2016–049
April 2016

Computing Resources Scrutiny Group

C Allton (UK), V Breton (France), T Cass (CERN), A Connolly(USA), M Delfino (Spain), F Gaede (Germany), J Kleist (Nordic countries), D Lucchesi (Italy, Chairman), H Meinhard (CERN, scientific secretary), D O'Neil (Canada), J Templon (Netherlands)

1 Introduction

This report analyses the resource usage in 2015 and the requests for 2017 from the four main LHC experiments ALICE, ATLAS, CMS and LHCb. Usage covers the first experience with Run 2 data, i.e. with pp collisions at $\sqrt{s} = 13$ TeV, and the tests of the modified computing model performances. Based on these results and on their computing models [1], the four experiments requested resources for 2017 and gave initial estimates of their resource needs for 2018.

The report summarises the C-RSG comments for the 2015 usage and the recommendations for 2017 resources procurement, 2018 resources requests were not scrutinized.

2 CRSG membership

Membership of the CRSG has changed for this scrutiny. There are new representatives for:

- Spain, Manuel Delfino
- UK, Chris Allton
- USA, Andrew Connolly

The chairperson is now Donatella Lucchesi. The group thanks the previous members, Jonathan Flynn as the previous chairperson, Terry Schalk and Jesus Marco de Lucas, for their important contributions over the years of their membership.

The chairperson thanks CRSG members for their commitment and the experiments' representatives for their collaboration with us. Thanks are also due to the CERN management for their support and to our 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 and resource requests by February 15. The CRSG thanks the experiments for the timely submission of their detailed documents [5–9, 11]. 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 with us.

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

For the October 2016 RRB we ask the experiments to submit their documents by August 30, 2016.

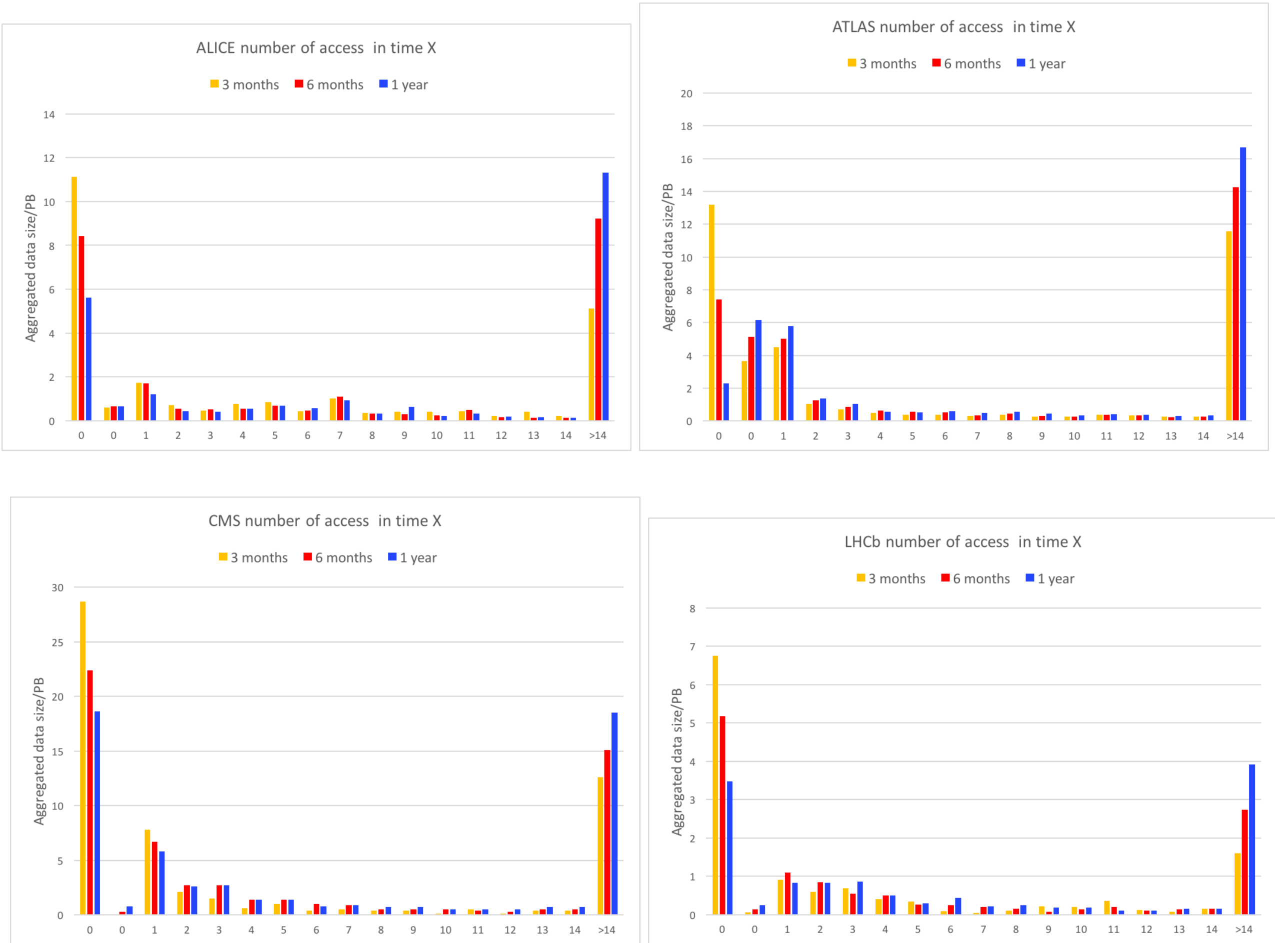

Figure 1 Volumes of data versus number of accesses in 3-, 6- and 12-month periods for ALICE (top left), ATLAS (top right). CMS (bottom left) and LHCb (bottom right). For each period X, data created in that period but not accessed is in the second bin. The first bin is for data created before the period began and not accessed during that period.

4 Overall assessment

The experiments continue to make intensive use of the WLCG resources. They are all working to optimize resource usage by diminishing the derived data formats and by reducing the number of data replicas in the tiers thanks to the availablity of fast networks connecting the majority of the sites.

The computing models are evolving towards a configuration in which the original hierarchical distinctions between tasks performed at different tier levels are eliminated, or at least very much diminished. A discussion of a new possible infrastructure configuration has started between the experiments and WLCG. Some operational savings may be realised by reducing the number of centers with limited disk space while expanding their CPU capacity.

Data popularity plots have been provided by all experiments for this scrutiny. In figure 1 we show the number of accesses in a given time period. We are interested in particular in the first two bins that contain data with zero access. The behavior is similar for the four experiments. The data access model has been changed and is changing in order to address the concerns raised by these plots. A lifetime has been attached to a dataset requiring its deletion or moving to tape when expired. We are pleased that these data popularity plots have become a standard feature of the experiments' usage reports. Experiments continue to make progress in the reduction of the CPU used by reconstruction and simulation programs. Unfortunately, however, this is not directly reflected in the CPU resource requests, as the reduction is outweighed by increases due to other factors, principally the increased event complexity. Simulation remains the biggest consumer of CPU capacity. There is common use of Geant4, although ALICE still relies almost entirely on Geant3. ATLAS and CMS included in their requests the resources needed for the simulation of Run 4 studies.

A sizable amount of CPU power is coming from the beyond-pledge and opportunistic use of resources from outside the WLCG. While we applaud this, and encourage the experiments to use any available resources, we are concerned by the dependence the experiments have on such resources which can become a problem with the increasing of the simulation and analysis complexity. ATLAS and CMS (two US sites) have experimented with the usage of commercial clouds, but they do not yet have conclusive results on the costs compared to the WLCG resources cost. ALICE, ATLAS, CMS and LHCb used in Run 2 the online computing farm (HLT) as expected during the shutdown and technical stops.

5 Comments and recommendations

1. We rely on resource monitoring and accounting from the WLCG [2] and the EGI accounting portal [3]. WLCG follows up reported problems and updates its accounting reports once the problems are fixed. The EGI accounting just recently fixed the multicore jobs CPU consumption. Data has been used to evaluate used CPU power and the efficiency. When we compared with numbers obtained by the experiments we see sometimes significant differences. We recommend experiments, EGI and WLCG to work together to ensure consistency between accounting systems.

2. The CRSG strongly supports software engineering development and recommends that sufficient effort be funded to support this. Improving the efficiency of software, including making optimal use of new hardware designs and improving the code optimization is essential to mitigate the growth in resource use. There have been substantial improvements for the current data taking, but in the longer term, with orders-of-magnitude increases in the expected computing needs, this work is even more essential.

3. Resources pledged to CMS at T1 are below the approved requests in 2014, 2015 and 2016 looking at the pledge summary. For CPU and disk this has been compensated by over-pledging at T2s, but this is *not* so for tape. CMS reports this as an important problem and we recommend the funding agencies and the RRB to provide a solution.

4. ALICE and ATLAS have reported an unexpected necessity of using computing nodes with more than 2GB/core of memory in the event reconstruction. The experiments have adopted different strategies to mitigate the problem. As a consequence of that, the CPU efficiency in particular at T0 is lower than expected. ALICE is now confident that the memory demand in RAW data reconstruction is a solved issue. The required CPU/core for event reconstruction has to be monitored and carefully evaluated in the coming data taking. It appears that the experiments would benefit from having CPU with more that 2GB/core of memory since some tasks exceed this limit.

Resources usage in 2015

The resources usage is analyzed first for the whole WLCG infrastructure then experiment by experiment. In both cases the numbers are obtained from the WLCG accounting reports [2] and from the EGI portal [3].

6 WLCG resources usage in 2015

The usage of CERN, T1 and T2 resources is shown in table 1 for years 2015 and the two previous years for comparison. The numbers represent the average calculated using time-integrated CPU power or storage capacity over the calendar year. The T2 disk usage information are not available. The CRSG would welcome WLCG/EGI reporting of T2 disk use. In the WLCG reports the disk-use efficiency factor was 70% until December 2013, but 100% thereafter. The CPU efficiency factor was 85% for T0 and T1 and 70% for T2 until May 2015 (included) and 100% thereafter. Tape efficiency is 100%. After removing these factors, we see the artificial decrease in CPU usage, but still greater than 100% with an important contribution coming from resources beyond the pledges of which the experiments benefit a lot.

Table 2 reports the situation of the pledges fulfillment at the end of January 2016. The pledged resources are reduced by the aforementioned efficiency factors before comparison to delivered resources.

Used/pledged resources

		2015	2014	2013
CPU	CERN	39%	53%	66%
	T1	102%	123%	114%
	T2	111%	152%	105%
Disk	CERN	80%	81%	116%
	T1	82%	95%	140%
	T2	—	–	–
Tape	CERN	76%	96%	106%
	T1	69%	89%	82%

Table 1 Usage summary for calendar year 2015 (left), 2014 (center) and 2013 (right). Data are from T1 and T2 accounting summaries in the CERN WLCG document repository [2].

Fulfillment of pledges as of January 2016.

CPU		Disk		Tape	
CERN	100%	CERN	99%	CERN	97%
T1	157%	T1	108%	T1	104%
T2	134%†	T2	—		

Table 2 Fulfillment of pledges as of January 2016. Data from the master accounting summary in the WLCG document repository [2]. †: the T2 CPU percentage is delivered/pledged from the T2 reports in [2].

Resource use at CERN end of 2015

	CPU	Disk	Tape
ALICE	30%	31%	17%
ATLAS	28%	29%	37%
CMS	32%	27%	36%
LHCb	10%	14%	10%

Resource use at CERN ens of 2014

	CPU	Disk	Tape
ALICE	26%	30%	12%
ATLAS	27%	28%	42%
CMS	35%	25%	35%
LHCb	12%	17%	11%

Resource use at CERN end of 2013

	CPU	Disk	Tape
ALICE	30%	29%	12%
ATLAS	28%	28%	43%
CMS	32%	28%	35%
LHCb	10%	15%	10%

Table 3 Use of resources at CERN by the experiments at the end of 2015 (top left), 2014 (top right) and 2013 (bottom). Data is from the master accounting summaries in the WLCG document repository [2].

Resource use at T1 end of 2015

	CPU	Disk	Tape
ALICE	19%	10%	7%
ATLAS	49%	54%	43%
CMS	17%	23%	41%
LHCb	14%	14%	10%

Resource use at T1 end of 2014

	CPU	Disk	Tape
ALICE	17%	8%	5%
ATLAS	48%	52%	41%
CMS	19%	30%	45%
LHCb	16%	11%	9%

Resource use at T1 end 0f 2013

	CPU	Disk	Tape
ALICE	12%	8%	6%
ATLAS	61%	52%	39%
CMS	15%	31%	46%
LHCb	12%	10%	10%

Table 4 Use of resources at T1 by the experiments at the end of 2015 (top left), 2014 (top right) and 2013 (bottom). Data is from the master accounting summaries in the WLCG document repository [2].

Tables 3 and 4 show the division of the resources at, respectively, CERN and at all T1 during the last three years. At CERN, CPU and disk space are almost equally used by ALICE, ATLAS and CMS with a smaller fraction for LHCb while tape space usage is dominated by ATLAS and CMS. At T1 around 50% of CPU and disk resources are used by ATLAS, with tape usage similar to that at the T0.

In table 5 we show the distribution of CPU usage at T2 in the last four years. Almost half of the CPU power goes to ATLAS, followed by CMS at around 30% then ALICE and lastly LHCb (which has very few T2s). Table 6 reports CERN's fraction of the total CPU use by an experiment (hence the column values need not sum to 100%). Historically, ALICE have been the experiment for which CERN has provided the largest fraction of total CPU, but the reduction observed in 2014 is confirmed in 2015; CMS concentrated an increased share of its total CPU at CERN in 2013, but this has reduced to the average previous value.

CPU consumption by experiment at T2

	2015	2014	2013	2012
ALICE	13%	15%	11%	8%
ATLAS	56%	51%	56%	54%
CMS	26%	29%	26%	34%
LHCb	5%	6%	7%	5%

Table 5 Distribution of the time-integrated normalized CPU time by the experiments at T2 centers for calendar year 2015 and three preceding years. Data from EGI Accounting Portal [3] with update time 2016-03-24 03:06 UTC.

CPU usage at CERN

	2015	2014	2013
ALICE	29%	24%	44%
ATLAS	13%	10%	13%
CMS	27%	27%	40%
LHCb	13%	13%	21%

Table 6 Distribution of the fraction of the experiment's total CPU consumption which has been at CERN. Data is from the master accounting summaries in the WLCG document repository [2].

The evaluation of the CPU efficiency is a currently much complicated by the move to jobs that use multiple cores. We interacted directly with EGI support with the help of WLGC in order to have the numbers in the official portal [3]. Table 7 summarizes the efficiency for T1 plus CERN and for T2 for the last four years. Efficiency is defined as sum of normalized CPU time divided by sum of normalized wall clock time times the number of processors. Efficiency values for 2014, 2013 and 2012 have been recalculated accordingly. The results are not in full agreement with those in previous reports due to the change in the handling of accounting for multicore jobs.

CERN plus T1 CPU efficiency

	2015	2014	2013	2012
ALICE	81%	81%	84%	62%
ATLAS	84%	85%	92%	91%
CMS	74%	74%	84%	85%
LHCb	94%	94%	95%	91%

T2 CPU efficiency

	2015	2014	2013	2012
ALICE	83%	78%	76%	64%
ATLAS	86%	87%	89%	88%
CMS	70%	75%	81%	86%
LHCb	96%	97%	96%	95%

Table 7 CPU efficiency for CERN plus T1 sites and for T2 sites by experiment for calendar year 2015 and for three preceding years. Data from EGI Accounting Portal [3] with update time 2016-04-01 02:06 UTC.

7 Usage by the experiments

7.1 ALICE

We summarize the computing resource usage by the ALICE experiment for the period January 1st 2015 to January 6th 2016, based on the report provided by ALICE [5]. In order to avoid potential discrepancies between ALICE MonALISA system and WLCG accounting, only the values reported by WLCG are reported.

ALICE has fulfilled its data taking program as established for 2015 collecting pp data for diffractive physics when the LHC was operating with isolated bunches, data triggered by single and double muon events (rare trigger) during running with 50 ns bunch spacing and, the bulk of the data, with a trigger mix including minimum bias events, single and double muon events and high multiplicity events during running with 25 ns bunch spacing. For all trigger categories the objectives in terms of statistics have been reached and even slightly exceeded. ALICE also collected pp and Pb-Pb events at 5.02 TeV according to its goals.

CPU, disk and tape usages are documented in the following table 8. Pledged resources are extracted from REBUS [4].

Tape usage exceeded the original requirements and the pledged resources as a consequence of an unexpected high pile-up factor during the pp run at 13 TeV in the 25 ns bunch spacing mode. As of January 2016, the cumulative amount of data stored was 18.4 PB at T0 (for a 2015 pledge of 16.2 PB) and 11.3 PBs in T1s (for a 2015 pledge of 10.2 PB). New values of the pp and Pb-Pb event sizes have been used to estimate the resources required in the period 2016-2018.

Disk usage sums up to 31 PBs (9,4 PB at T0, 10.1 PB at T1s and 11.5 PB at T2s) excluding 3.4 PB of disk buffer in front of the taping system at T0 and 2.8 PB of disk buffer in front of the taping system at T1s. ALICE request for 2015 amounted to 55 PB (14.5 PB at T0, 17.8 PB at T1s and 22.7 PB at T2s), including the required disk buffer sizes at T0 and T1 REBUS [4]. The usage of disk resources is below the required resources because the fully calibrated reconstruction of 2015 data was delayed due to unexpected distortions in the TPC. A procedure to correct for these distortions was developed and a full reconstruction of 2015 Pb-Pb and pp data is foreseen to take place during the spring of 2016. To minimize the interference between data processing and 2016 data taking, a strategy has been set-up with CERN-IT to run two tape buffers, one for data taking and one for data reconstruction. A significant fraction of the resources pledged for ALICE have not yet been installed, especially at T2s, but it had no impact on the experiment because of the delay in the reconstruction of 2015 Pb-Pb and pp data.

CPU usage amounts to 517 kHEP-SPEC06: 127 were used at T0 (including CAF), 190 at T1s and 200 at T2s. Compared to 2015 pledges, T0 resources appear underused (127 out of 175) but this was due to a misevaluation of the kHEP-SPEC06 values that was corrected from July 2015 on.

The share among tasks was 9% for raw data reconstruction, 69% for Monte-Carlo production, 14% for train analysis and 7% for end user analysis.

The efficiency of jobs calculated as the ratio of the used CPU time and the wall time was on average 73% at T0, 83% at T1s and 82% at T2s.

This year, ALICE provided a data popularity plot using the same template as the other experiments. The plot shows that the large majority of the data are accessed, mainly through train analyses. The trains are a collection of user tasks (wagons), which are running over the same input data set. The input can be a single or a multiple data sets. The wagons can analyze different portions of the data set. Each wagon produces its own output set, which is merged in a common container at the end of the train run. The trains are configured per Physics Working Group and per analysis type. About 2200 analysis trains have been executed over 2015. "Accessed" means in the context of analysis trains that the data have been processed by at least one wagon.

Resource	Site(s)	2015 request	2015 pledged	2015 used	2015 used/pledged	2015 CPU efficiency
CPU (kHS06)	T0 + CAF	175	175	127 (80+47)[73]	72,6 %	73%
	T1	120	146.5	190 [188]	130 %	83%
	T2	200	194.1	200 [214]	103 %	82%
Disk (PB)	T0	14.5	14.5	9.4 + 3.4 [10.8]	88 %	
	T1	17.8	15.8	10.1 + 2.8 [7.2]	82 %	
	T2	22.7	16.8	11.5	71 %	
Tape (PB)	T0	16.2	16.2	18.4 [11.9]	114 %	
	T1	10.2	11.9	11.3 [7.2]	95 %	

Table 8 Summary of resource usage by ALICE in 2015 (January 2015 to December 2015), with 2015 pledges and 2015 requests. T0 disk requirement includes 3.4 PB of disk buffer in front of the taping system. T1 disk requirement includes 2.8 PB of disk buffer in front of the taping system. Requested and pledged resources are from REBUS [4]. T1 pledged tape resources have been adjusted after interaction with ALICE.

7.2 ATLAS

ATLAS 2015 resource usage continues to be dominated by simulation and reconstruction of Monte Carlo samples, but a considerable amount of CPU capacity has also been spent on data analysis.

ATLAS achieves a CPU efficiency of approximately 90% at all Tiers but T0, which only displays a value of 80%. The low efficiency at T0 is caused by the first pass of the bulk reconstruction that requires more than the assumed 2GB/core of memory and therefore leads to under utilization. Multicore jobs also spend a large fraction of their time in inherently sequential phases.

Since the last scrutiny ATLAS has worked heavily on improving the efficiency of the HLT farm and it has achieved an efficiency of 93% at the end of 2015 compared to 70% at the start. ATLAS has also made an effort to investigate the use of the HTL farm during offline periods. Unfortunately, the current conclusions are that the HTL farm will probably not be usable for ATLAS between fills during 2016 due to the time needed to switch back and forth between online and simulation mode and the expected time available between fills. We encourage ATLAS to continue working on lowering the switch time.

ATLAS continues the heavy usage of beyond pledge CPU resources, at T2 75% is beyond pledge.

ATLAS has made good progress at making the best use of available disk, improvements comes from usage of data lifetime to allow for deletion of data with expired lifetime, a reduction in the number of different space tokens and a flattening of the T1/2 hierarchy.

Compared to the last scrutiny, the access frequency plots show a significant decrease in untouched data, showing the improvements in ATLAS' data management.

ATLAS 2015 data on tape amounts to approximately 28PB, compared to 33PB expected in April 2015. ATLAS has also started applying lifetime driven deletion from tape and deletion campaigns are clearly visible in the occupancy data.

Table 9 shows an overview of ATLAS' resource usage for 2015. The pledged are from rebus [4]. Used CPU are taken Atlas report, in square brackets the values obtained from the master accounting summary in the WLCG document repository [2]. ATLAS in their report pointed out that the numbers in the EGI accounting portal are skewed by incorrect reporting of multicore usage, we report also these numbers.

Resource	Site	Pledged	Used	Used/Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	205	223 [74]	109%	84%
	T1	493	648 [482]	131%	86%
	T2	579	1018 [896]	176%	87%
Disk (PB)	T0+CAF	14	11[10.2]	76%	
	T1	45	38 [40]	84%	
	T2	51	42	82%	
Tape (PB)	T0	33	27 [26]	85%	
	T1	70	53 [48]	76%	

Table 9 Fulfillment of pledges. The table reports the ATLAS situation at the end of 2015.

7.3 CMS

CMS CPU usage has been in line with pledges, except for some additional use of T2 CPU made possible through available capacity at some sites. CMS has made good use of the HTL farm during 2015, which when available provides an amount of CPU comparable to that T1 centers together.

The low CPU efficiency at T1 seen is caused by known effects, the situation improved during 2015 and CMS expect this to continue during 2016.

CMS has implemented a dynamic data management system for driving automatic deletion and replication of data. The results are very promising and we encourage further development of the system.

The deficit of pledged tape at T1 centers continues and only through aggressive deletion campaigns CMS has been able to accommodate data. The 82% of used/pledged is a snaptshot of the tape consuption at the end of 2015 after deletion campaigns. This situation is a continued concern for CRSG.

Table 10 shows an overview of CMS resources usage for 2015.

Resource	Site	Pledged	Used	Used/Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	271	77 [62]	28%	82%
	T1	277	256 [169]	92%	56%
	T2	564	518 [416]	92%	68%
Disk (PB)	T0+CAF	15	11 [10]	73%	
	T1	24	22 [18]	92%	
	T2	34	34	100%	
Tape (PB)	T0+CAF	35	31 [25]	88%	
	T1	67	55 [45]	82%	

Table 10 Fulfillment of pledges. The table reports the CMS situation at the end of 2015. Data from the master accounting summary in the WLCG document repository [2].

7.4 LHCb

This report covers all of 2015 (January to December) and is based on a report from LHCb [9], and follow-up with the experiment via email. The pledged are from rebus [4]. Used resources are from the LHCb DIRAC portal [10], in square brackets the values obtained from the master accounting summary in the WLCG document repository [2].

Operations in 2015 included continued simulation, a legacy stripping of Run1 data, the running of user jobs and the processing of new Run2 data. Overall, the usage of compute, disk and tape was very close to original expectations. The used CPU resources distribution was slightly different than anticipated with less compute resources provided at T0+T1 and more provided by T2s as shown in the table Table 11.

Due to the lower than anticipated LHC live time in 2015, the experiment was able to take higher rates for both their TURBO and FULL streams and to exercise the split HLT concept for the first time, while still remaining within their pledges.

Resource	Site(s)	2015 pledge	2015 used	Average CPU efficiency
CPU (kHS06)	T0	36	18 [19.5]	91%
	T1	139	136 [136]	96%
	T2	61	84 [76.8]	97%
Disk (PB)	T0	5.5	5 [4.8]	
	T1	14	14.7 [10.3]	
	T2	2.6	1.5	
Tape (PB)	T0	11	8.6 [7.4]	
	T1	28	13 [10.7]	

Table 11 2015 LHCb usage table. Disk used at T0 and T1 is T0D1 class plus cache space for tape storage, it does not include stage and read pools for dCache. CPU is CPU power in kHS06 averaged over one year. The T2 cpu usage also includes non pledged WLCG sites, while T1 CPU usage also includes the HLT resource.

Resource requirements for 2017

8 Assumptions for resource requests

The assumptions used by the experiments to determine the resources needs are based on the LHC running conditions [1] and on the updated approved schedule [12]. Table 12 reports the anticipated LHC beam live times updated to the latest official schedule [12]. The machine efficiency in 2017 and 2018 for pp runs is assumed of 40% which may be considered optimistic but possible. The final column gives the average pileup (average number of collisions in each beam-crossing) for ATLAS and CMS pp collisions. The event sizes and reconstruction times increase as pileup increases bringing to a larger demand of resources in particular for CPU power. The LHC luminosity is expected to be $1.5 \times 10^{34}\,cm^{-2}\,s^{-1}$ for 2016, 2017 and 2018.

Since June 2015 the assumed efficiencies on CPU, disk and tape usage is 100%, it was 85% and 70% for organized and analysis CPU usage.

RRB year	pp/10^6 s	HI/10^6 s	pp pileup
2015	3	0.7	25
2016	5	0.7	35
2017	5.5	-	35
2018	5.5	0.7	35

Table 12 Assumptions on live time for LHC running in Run 2, 2015 to 2018. The final column gives the anticipated average pileup for ATLAS and CMS during pp running for each year.

9 Resource requirements from the experiments

9.1 ALICE

Table 13 summarizes the requests from ALICE for 2016, 2017 and 2018. For a sake of clarity, the following information is provided:

- For years 2016 and 2017, ALICE requests are documented for October 2015 and April 2016 RRB meetings.
- CERN CAF resources are included in the T0 requests.
- T0 disk requirement includes 3.4 PB of disk buffer in front of the tape system. T1 disk requirement includes 2.8 PB of disk buffer in front of the tape system.

CSRG had several meetings with ALICE computing coordinators to understand a number of issues coming out of table 13:

- A first issue concerns the significant increase in 2016 and 2017 requests from the previous scrutiny in October 2015. Tape requests for 2016 increased by 21.7% at T0 (26.3 instead of 21.6) and 30.1% at T1 (20.3 instead of 15.6) while CPU request for 2016 was increased by 5% at T0 and T1s. The reason for the increase in CPU request is that the processing power for pp events has been tentatively increased by a factor 2 to take into account the increased processing

Resource	Site(s)	2016 RRB	2016 ALICE rev.	2017 ALICE	2017 CRSG	2018 ALICE
CPU (kHS06)	T0	215	224	255(245)	255	318
	T1	157	168	221(207)	221	284
	T2	237	235	275(267)	275	320
Disk (PB)	T0	16.8	16.8	21.4(19.2)	21.4	27.7
	T1	21.0	21.4	25.0(24.2)	25.0	31.0
	T2	26.1	26.1	31.3(30.7)	31.3	39.7
Tape (PB)	T0	21.6	26.3	34.4(25.7)	34.4	45.9
	T1	15.6	20.3	28.4(19.7)	28.4	39.9

Table 13 Summary of ALICE updated resource requests for 2016, 2017 and 2018. For years 2016 and 2017, ALICE requests are documented for October 2015 and April 2016 RRB meetings. CERN CAF resources are included in T0 CPU requests. T0 disk requirement includes 3.4 PB of disk buffer in front of the tape system. T1 disk requirement includes 2.8 PB of disk buffer in front of the tape system. Numbers in brackets for 2017 show the October 2015 CRSG recommandations.

time needed for high pileup events. Similarly, the pp raw data size has been increased by a factor 3.5 to match the real average size in 2015. When asked for a plan B in case the funding agencies were not able to respond to these unexpected requests, ALICE computing coordinators explained that reducing the amount of minimum bias pp event would impact directly the ALICE physics program and/or mean that only one copy of the RAW data files was kept in some instances leading to potential data loss.

- Two scenarios were documented in the ALICE requirements document [6] for April 2016 RRB on T2s in 2016 and 2017: one scenario corresponds to the direct output of the computing model that generates a very significant peak in CPU needs in 2016 in order to simulate Pb-Pb events collected at 5.02 TeV in 2015. To avoid this peak and adjust for a smoothed growth of CPU requests on T2, a T2-smoothed scenario has been proposed that postpones some Pb-Pb Monte-Carlo simulation to 2017. CSRG supports the use of the so-called T2-smoothed scenario that is documented in table 13.

The CRSG group endorses ALICE requests for 2016, 2017 and 2018 with the following caveats:

- The updated requests for 2016 and 2017 go far beyond the initial goal of keeping a flat budget. However, ALICE computing coordinators made it clear that changing the triggers and especially reducing the amount of high pile-up minimum bias pp event would negatively affect the physics program. In the case that the new resource requests from ALICE could not be matched by the funding agencies for 2016 and 2017, a redefinition of the ALICE physics program would be required, which is clearly outside the scope of the CRSG.
- Once the procedure to correct for the distortions in TPC is validated, the full reconstruction of 2015 Pb-Pb and pp data is foreseen to take place during the spring of 2016. CRSG will revise 2017 requests at next scrutiny in October 2016.

Moreover, the CRSG makes the following recommendations to the ALICE computing coordinators:

- We encourage and support the migration from Geant3 to Geant4 for Monte-Carlo production to take advantage of multicore architectures. More generally, we strongly encourage software developments to optimize CPU usage and take advantage of parallel and multicore architectures.

- We appreciate that the different scenarios for T2 requests were documented in the requirements document provided by ALICE. We are happy to have only one smoothed scenario for the future provided the smoothing process is detailed.
- we appreciate that ALICE provided a data popularity plot following the same template as the other experiments.

9.2 ATLAS

Notable in the ATLAS requests there are increases in the requests for CPU at T0 and T2s and a decrease in the request for disk at T2s. ATLAS explains that the CPU increase at T0 is caused by better knowledge about Run-2 data, the performance of the software and its behavior on multiprocessor architectures, previous estimates have been too optimistic. The CPU growth at T2s is explained with the need for additional power for the upgrade simulation as well as an unanticipated high use of full simulation and event generation. The reduction of T2 disk request is based on two observations: a lower number of events produced by fast simulation and a wish to consolidate storage only in big sites.

The wish to consolidate data resulted in a request that T2-sites below 400TB of disk invest in CPU rather than disk. We have expressed concerns about this strategy as this might not fit with the funding schemes of the T2 sites. ATLAS note that the scheme will be implemented gradually and in concert with the funding agencies.

The T0 CPU resources have been reduced in agreement with ATLAS. The original request was motivated by the large memory (>2GB/core) necessary for event reconstruction. The extra 83 kHS06 will, hopefully, not needed thanks to better performances of the delivered hardware to CERN and the optimisations in the ATLAS workflows. The HLT farm, that contributed with 42 kHS06 to CPU resources has not been included in the CPU pledged resources therefore, we propose to lower the CPU requests at T1.

ATLAS expects to continue their large use of beyond-pledge CPU resources for full simulation and though ATLAS is quite certain that such resources continue to be available, CRSG continues to be concerned about the reliance on such resources.

Table 14 shows the latest ATLAS computing requests for 2017 and 2018 along with the current resource pledges and the CRSG recommendations.

Resource	Site	2016 RRB	2017 ATLAS	2017 CRSG	2018 ATLAS
CPU (kHS06)	T0+CAF	257	383(270)	300	389
	T1	520	703(662/637)	682	763
	T2	566	846(702/675)	846	946
Disk (PB)	T0+CAF	17	20(18)	20	20
	T1	47	57(54)	57	60
	T2	72	78(91)	78	84
Tape (PB)	T0+CAF	42	53(51)	53	67
	T1	116	173(185)	173	257

Table 14 ATLAS resources requests and CRSG recommendations. The numbers in parentheses () are those previously requested by ATLAS / proposed by CRSG (if different) in October 2015.

9.3 CMS

Overall CMS requests for 2017 are in line with the numbers suggested by CRSG in October 2015 scrutiny, with the major change being a decrease in the request for CPU and increase in the request for disk at T2s. CMS explains the decrease in the need for T2 CPU due to the extensive use of MiniAOD for analysis (approx. 75% of the analysis jobs), whereas the increase in disk space at T2s is to accommodate opportunistic use of T2 CPU for doing GEN-SIM and RAW events processing. CRSG would like to return to this item in September/October and see if the additional disk indeed leads to better usage of T2 CPU.

CRSG is worried about the continued deficit in tape which the current request is not lowering and we encourage CMS and sites to consider ways to improve on the situation in the future.

Table 15 shows the latest CMS computing requests for 2017 and 2018 along with the current resource pledges and the CRSG recommendations.

Resource	Site	2016 RRB	2017 CMS	2017 CRSG	2018 CMS
CPU (kHS06)	T0+CAF	319	378(377)	378	383
	T1	400	550(550)	550	600
	T2	700	750(800)	750	900
Disk (PB)	T0+CAF	16.3	17.8(17)	17.8	19.8
	T1	33	49(49)	49	58
	T2	38	53(47)	53	63
Tape (PB)	T0+CAF	44	58(58)	58	72
	T1	100	135(135)	135	175

Table 15 CMS resources requests and the CRSG recommendations. The numbers in parentheses () are those suggested by CRSG in October 2015.

9.4 LHCb

This report is based on the original submission by LHCb [11] and follow-up via email.

Table 16 shows the latest LHCb computing requests for 2016 and 2017 along with the CRSG recommendations. The submitted requests are the same as the requests submitted in fall 2015, except for a change in projected LHC live time from 5.1 to 5.5Ms. The preliminary CRSG recommendations from fall 2015 remain unchanged, but have been adjusted for this lifetime projection.

The CRSG applauds previous work on software efficiency and optimization which has already resulted in substantial savings. We encourage continued work on fast MC simulation and the “particle gun” approach to further reduce CPU and disk consumption.

Resource	Site(s)	2016 LHCb	2016 CRSG	2017 LHCb	2017 CRSG
CPU (kHS06)	T0	51	51	61	61
	T1	156	156	189	189
	T2	88	88	106	106
	HLT + Yandex	20	20	20	20
Disk (PB)	T0	7.6	7.6	8.7	8.7
	T1	13.5	13.5	17.7	17.7
	T2	4.0	4.0	3.8	3.8
Tape (PB)	T0	20.6	20.6	22.0	22.0
	T1	42.1	42.1	38.8	38.8

Table 16 LHCb resources requests for 2016 and 2017 and CRSG recommendations.

References

[1] I Bird, P Buncic, F Carminati, M Cattaneo, P Clarke, I Fisk, J Harvey, B Kersevan, P Mato, R Mount and B Panzer-Steindel (Eds), Update of the Computing Models of the WLCG and the LHC Experiments

[2] WLCG Document Repository, CERN, https://espace.cern.ch/WLCG-document-repository/

[3] EGI Accounting Portal, CESGA, Spain, http://accounting.egi.eu/

[4] REBUS Pledge Summary, CERN, http://wlcg-rebus.cern.ch/apps/pledges/summary/

[5] ALICE Collaboration, ALICE 2015 Computing Resources Usage. 21 January 2016

[6] ALICE Collaboration, Update on Computing Resources Requirements 2016–2018. 25 January 2016

[7] ATLAS Collaboration, Computing Status and Plans Report to the C-RSG, April 2016. 15 February 2016

[8] CMS Collaboration, CMS Software/Computing Resource Utilization & Request 2015-2018, Spring 16. 15 February 2016

[9] LHCb Collaboration, C. Bozzi(Ed), LHCb Computing Resource Usage in 2015(II), LHCb–PUB–2016–004. 15 February 2016

[10] http://lhcb-portal-dirac.cern.ch.

[11] LHCb Collaboration,C. Bozzi (Ed), LHCb Computing Resources: 2017 request, LHC–PUB–2016–003. 11 February 2016

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