CERN–RRB–2017–056
April 2017

Computing Resources Scrutiny Group

C Allton (UK), V Breton (France), G Cancio Melia (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 2016, the pledges for 2017 and the requests for 2018 from the four main LHC experiments ALICE, ATLAS, CMS and LHCb. The year 2016 corresponds to unexpectedly high performance of the LHC machine that are expected to be at the same level also in 2017 and 2018. The increased amount of data collected is requiring additional resources from CERN and the Funding Agencies (FA), and a significant effort by the collaborations. The fact that the 2017 pledges do not cover completely the experiments requests and the consequent implications on the requests for 2018 are discussed. The experiments were asked also to provide a preliminary evaluation of 2019 resources, but it was too early to provide a precise estimation due to the current uncertainties.

The report summarizes the C-RSG comments for the 2016 usage, 2017 pledges and the recommendations for 2018 resources procurement.

2 CRSG membership

Membership of the CRSG has changed for this scrutiny. There is a new representative for CERN, German Cancio Melia who replaces Tony Cass. The group thanks Tony for his important contributions over the years of his membership.

The chairperson thanks the 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 1^{st}. The CRSG thanks the experiments for the timely submission of their detailed documents [1–6]. 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 2017 RRB we ask the experiments to submit their documents by September 11^{th} 2017.

Fulfillment of pledges as of February 2017.

CPU		Disk		Tape	
CERN	101%	CERN	97%	CERN	99%
Tier-1	108%	Tier-1	105%	Tier-1	95%

Table 1 Fulfillment of pledges as of February 2017. Data from the Rebus WLCG repository [10].

Pledges Balance

	ALICE	ATLAS	CMS	LHCb
CERN CPU	0%	0%	0%	0%
CERN disk	0%	0%	0%	0%
CERN tape	0%	0%	0%	0%
T1 CPU	-8%	-12%	-14%	-4%
T1 disk	-14%	1%	-21%	-6%
T1 tape	-1%	-7%	-24%	-3%
T2 CPU	-24%	-13%	-7%	27%
T2 disk	-28%	-7%	-22%	-30%

Table 2 Fractional pledges balance, (total offered pledges - experiment's required pledges)/(experiment's required pledges). Data from the REBUS WLCG repository [11].

4 Overall assessment

The experiments in 2016 have made a very intensive use of the WLCG resources. They were forced to do a lot of operations based on human intervention in order to keep pace with the exceptional LHC performance. In order to reduce the 2018 requests they have optimized resource usage by further diminishing the derived data formats and by reducing the number of data replicas in the tiers relying now on the availability of fast networks connecting the majority of the sites.

The computing models are continuing their evolution towards a configuration in which there will be large data centers and those with limited disk capacity focus on provisioning CPU resources. The discussion of the infrastructure configuration has started among the experiments, WLCG, CERN management and Funding Agencies.

4.1 Fulfillment of Pledges

Table 1 summarizes the pledge fulfillment for RRB year 2017. These represent the resources available, actually delivered to the experiments respct to those pledged by the Funding Agencies. But 2017 was treated as special year. Given the exceptional performances of the machine, the experiments asked for more resources than the *flat budget* expectations and during the October RRB the agencies were asked to contribute on a best effort basis. In table 2 we report the balance defined as (total offered pledges - experiment's required pledges)/(experiment's required pledges), for each experiment for each tier. CERN has provided the experiments with the requested resources while at T1 and T2 level not all the required requests have been actually offered. We note that ATLAS can count almost on the requested resources, the missing CPU power can be easily compensated by the overpledeges that the experiment always has. LHCb is in a similar situation, the missing disk space at T2 is not worrying given the small size of the T2. ALICE is lacking disk space both T1 and T2 level. CMS is lacking disk space at T1 and T2 centers, and tape at T1 level.

Percentage increase in pledges 2013-2017

	ALICE	ATLAS	CMS	LHCb
CERN CPU	37%	45%	38%	23%
CERN disk	31%	29%	39%	35%
CERN tape	23%	40%	33%	43%
T1 CPU	25%	27%	39%	19%
T1 disk	27%	19%	20%	24%
T1 tape	32%	50%	36%	39%
T2 CPU	14%	29%	19%	38%
T2 disk	17%	14%	17%	24%

Table 3 Annual increase in the pledges from fitting 2013-2017 REBUS data [11] as shown in Figs 1-4.

4.2 Pledged Resources in the recent years

Triggered by the pattern of promised resources in 2017 shown in Table 2, C-RSG has studied the pledges as function of year. We used the numbers as recorded in the REBUS database and plotted them as function of the year starting from 2013 when the RRB decided to provide resources to the LHC experiments on the *flat budget* basis. Then, data is fitted to find the average increase. Figures 1-4 show this analysis with the solid line representing the fit result and the dotted line the extrapolation to 2018. The percentage increase is summarized in table 3 for each experiment and each tier.

We remind here that the Computing Models presented by the experiments [7] assumed an increase of 20% of CPU and 15% of disk and tape space each year for fixed budget.

As it can be seen from the figures, the pledges have been rising faster than expected for a flat budget. In particular we see an increase almost two times the expected one for tape.

4.3 Analysis of Access-Frequency data

Data popularity plots have been provided by all experiments. In figure 5 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. Recently ATLAS, CMS and LHCb started to change data access model by attaching a lifetime to datasets, that are deleted or moved to tape when not used for a given period of time. By comparing the plots with those of previous year we noticed that the amount of data not accessed keeps reducing. CMS has sizeable amount of data in the second bin, not accessed in the period produced, that was not present before. CMS has explained that the most recently produced data that has not been accessed in the period is due to a temporary effect caused by a massive Monte Carlo production in view of Moriond 2017 and simulated samples that were not yet accessed at the time of the plot creation. We have to recall also that CMS plot is CRAB-based and does not show all the disk space accessed and therefore it underestimates CMS usage. ALICE is working to implement a procedure to reduce old not accessed data that appears in the first bin.

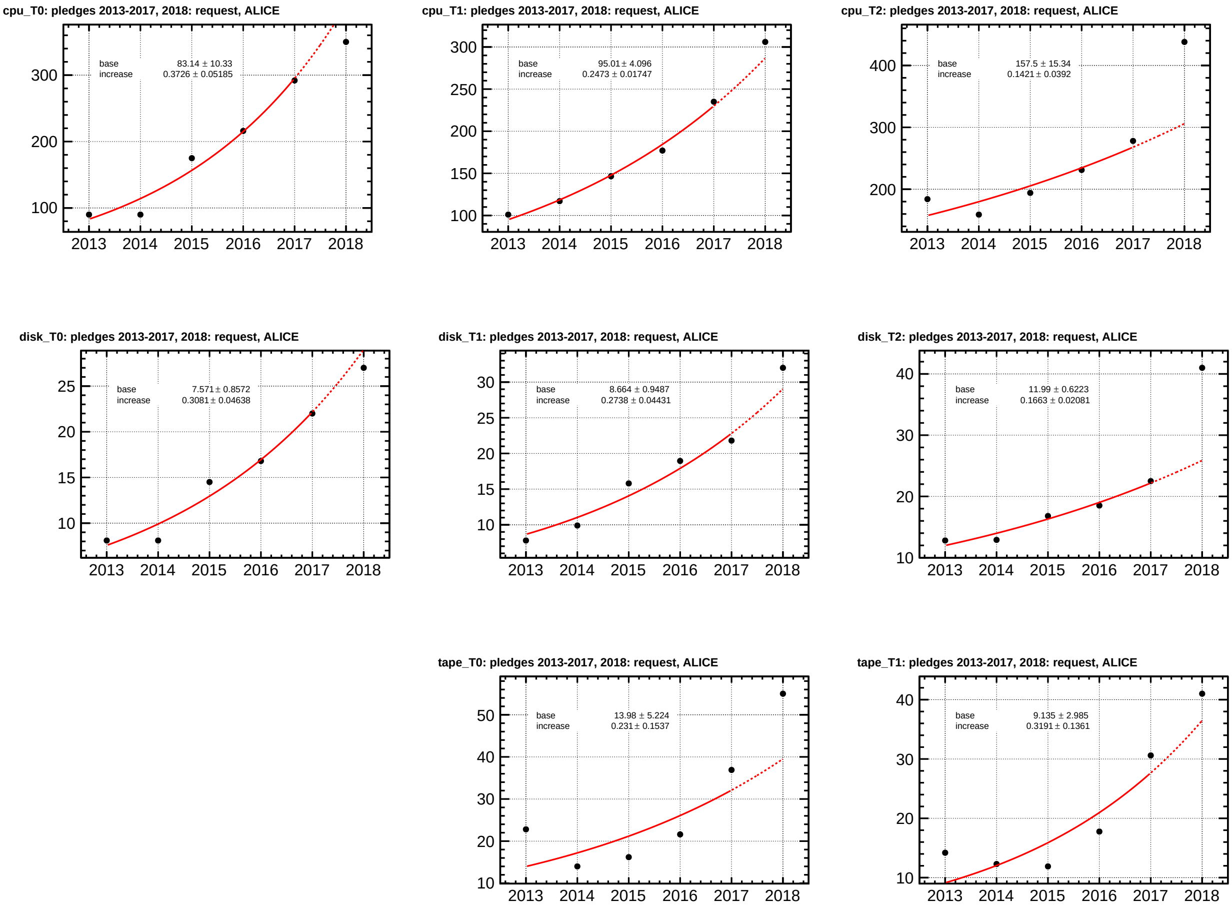

Figure 1 ALICE: plots in the first row show the CPU increase at T0, T1 and T2 while the second one has the disk space increase for the same tiers and the last row the tape space increase respectively at T0 and T1. The last point represents the 2018 requests, not considered in the fit.

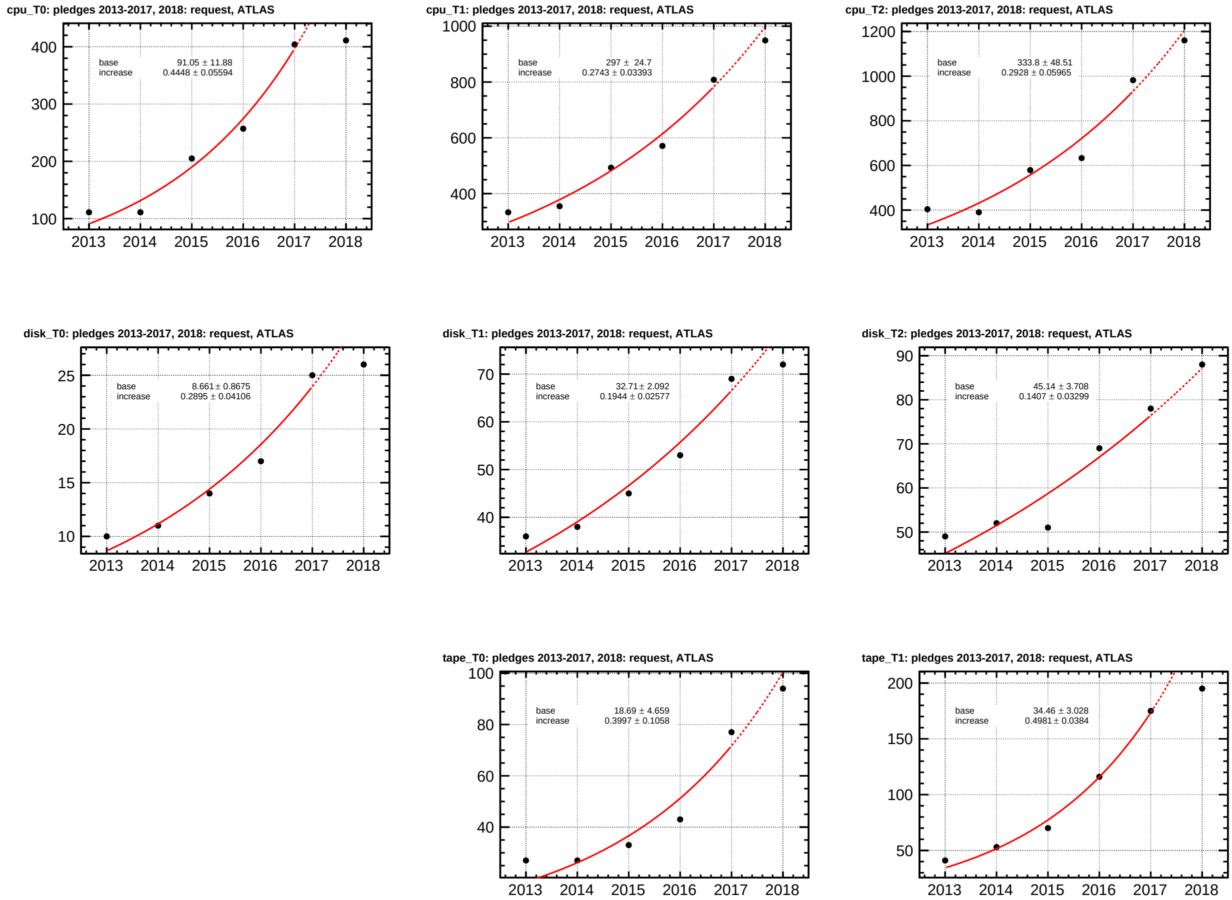

Figure 2 ATLAS: plots in the first row show the CPU increase for T0, T1 and T2 while the second one has the disk space increase for the same tiers and the last one the tape space increase respectively at T0 and T1. The last point represents the 2018 requests, not considered in the fit.

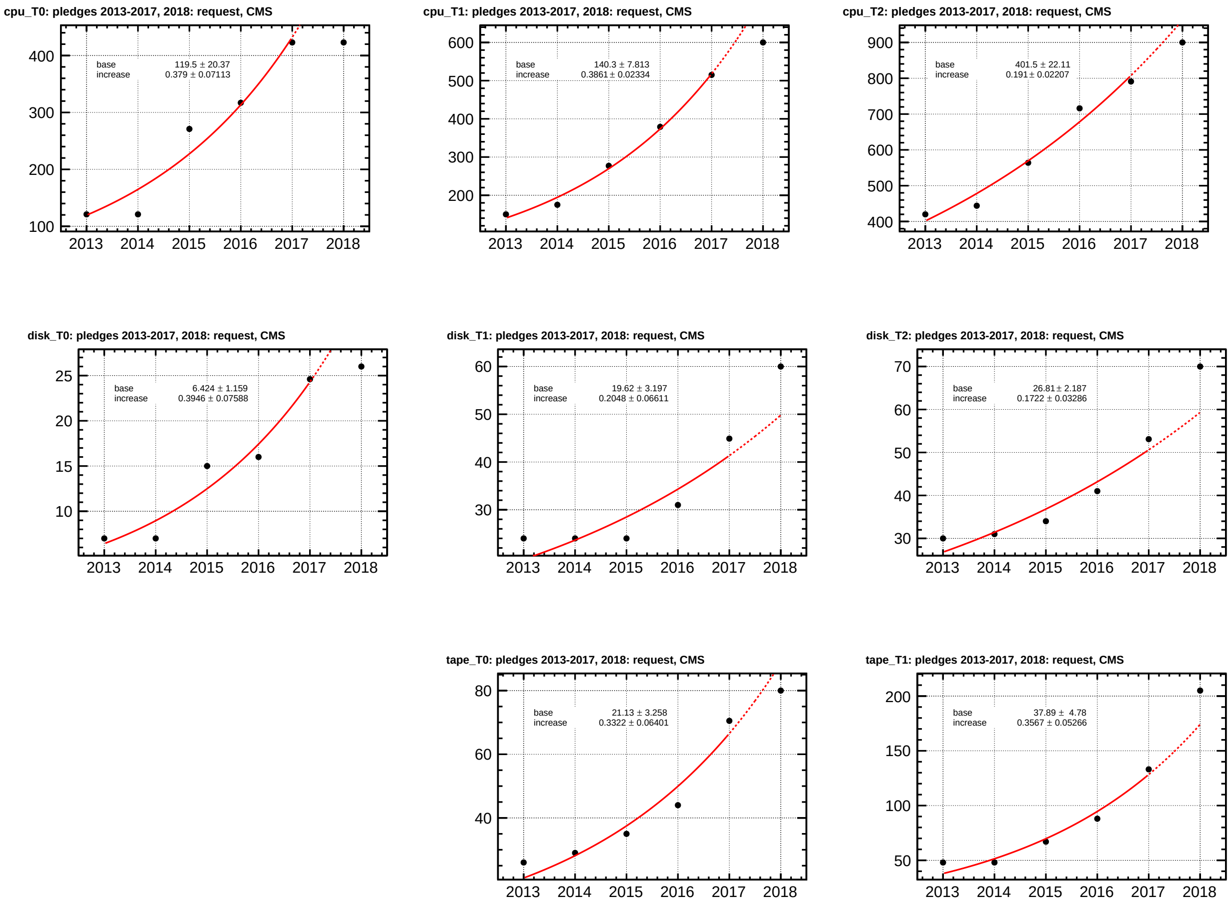

Figure 3 CMS: plots in the first row show the CPU increase for T0, T1 and T2 while the second one has the disk space increase for the same tiers and the last one the tape space increase respectively at T0 and T1. The last point represents the 2018 requests, not considered in the fit.

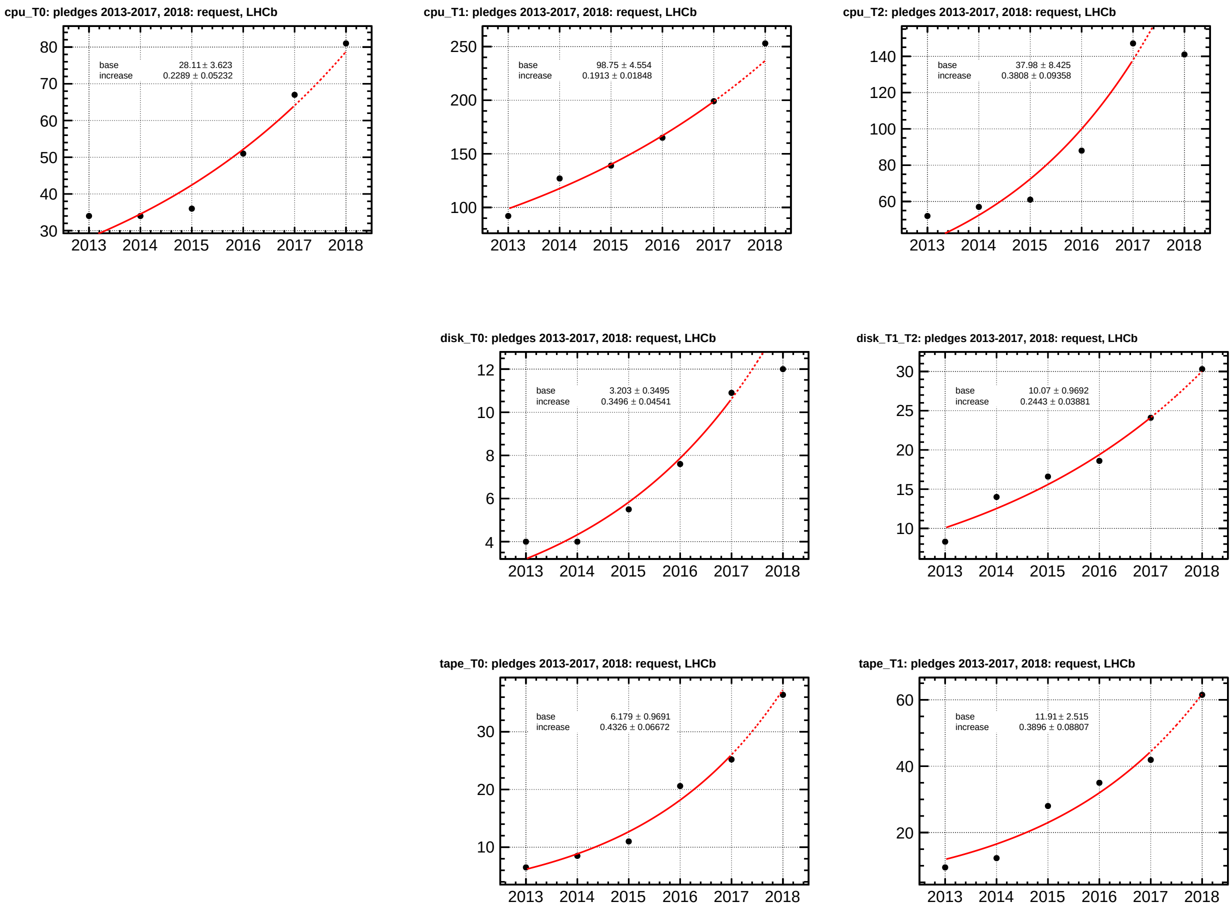

Figure 4 LHCb: plots in the first row show the CPU increase at T0, T1 and T2 while the second one has the disk space increase where T1 and T2 are merged and the last one the tape space increase respectively at T0 and T1. The last point represents the 2018 requests, not considered in the fit.

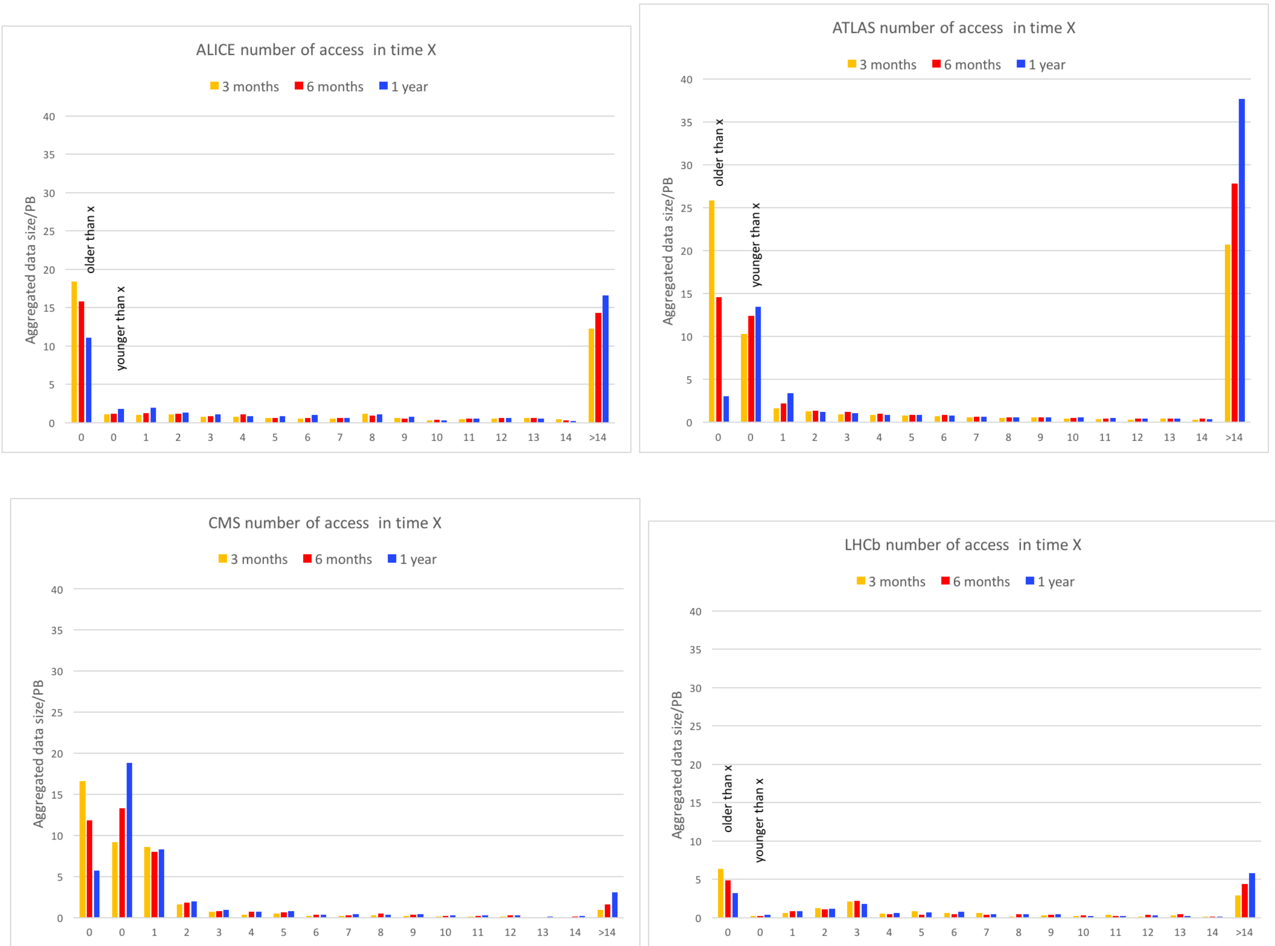

Figure 5 Volumes of data versus number of accesses in 3-, 6- and 12-month periods for ALICE (top left), ATLAS (top right). CMS (bottom left) and LHCb (bottom right). For each period X, data created in that period but not accessed is in the second bin. The first bin is for data created before the period began and not accessed during that period.

5 Resources usage in 2016

The resources usage is analyzed for the whole WLCG infrastructure and then experiment by experiment. In both cases the numbers are obtained from the EGI portal and validated against the experiments. The analyzed period goes from April 1 2016 up to March 31 2017.

5.1 WLCG resources usage in 2016

The usage of CERN, T1 and T2 resources is shown in table 4 for years 2016 and the three previous years for comparison. The numbers represent the average calculated using time-integrated CPU power or storage capacity over the RRB year 2016 and the calendar year for the previous ones. As in the previous years the T2 disk usage information are not available and C-RSG would welcome WLCG/EGI reporting of T2 disk use. For 2016, CERN CPU usage includes the local jobs, now providing a complete accounting. CPU usage includes an important contribution coming from resources beyond the pledges, from which the experiments benefit significantly in 2016. Disk and tape space have been largely used at T0 due to the data logging. Tables 5 and 6 summarize the division of the resources at CERN and

Used/pledged resources

		2016	2015	2014	2013
CPU	CERN	122%	39%	53%	66%
	T1	119%	102%	123%	114%
	T2	151%	111%	152%	105%
Disk	CERN	97%	80%	81%	116%
	T1	72%	82%	95%	140%
	T2	—	—	–	–
Tape	CERN	98%	76%	96%	106%
	T1	67%	69%	89%	82%

Table 4 Usage summary for 2016 RRB year and for calendar year 2015, 2014 and 2013. Data is from Tier-1 and Tier-2 accounting summaries for WLCG obtained from EGI [8]. CERN percentage is now taken into account properly including the local jobs and this explain the difference with the previous years.

at all T1s respectively during the last four years. At CERN, CPU is used almost equally by ATLAS and ALICE with CMS being the major user in 2016 while LHCb is limited to a few percent. Disk space is almost equally used by ALICE, ATLAS and CMS with a smaller fraction for LHCb, while tape space usage is dominated by ATLAS and CMS with a significative increase in ALICE usage. At T1 the main user is ATLAS, which has a share of more than 40% for CPU, disk and tape space. CMS is around 20% and finally ALICE and LHCb are almost at the same level. In table 7 we show the distribution of CPU usage at T2 in the last four years. As in the previous year, in 2016 almost half of the CPU power goes to ATLAS, followed by CMS, then ALICE and lastly LHCb, which has very few Tier-2s and can not be expected to compete with the others.

Table 8 reports CERN's fraction of the total CPU use by each experiment, hence the column values are not expected to sum to 100%. Historically, ALICE has been the experiment for which CERN has provided the largest fraction of total CPU. There was a reduction in 2014 and 2015 but in 2016, with the increased requests, CERN is again the major contributor. The trend is constant for the other experiments.

Table 9 summarizes the efficiency for T1 plus CERN and for T2 for the last five years. Efficiency is defined as the sum of normalized CPU time divided by sum of normalized wall clock time times the

Resource use at CERN
Year 2016

	CPU	Disk	Tape
ALICE	24%	29%	21%
ATLAS	28%	30%	33%
CMS	43%	30%	34%
LHCb	5%	11%	12%

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
End 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 5 Use of resources at CERN by the experiments for RRB year 2016 (top left) and at the end of 2015 (top right), 2014 (bottom left) and 2013 (bottom right). Data is generated from the EGI accounting web page [8].

Resource use at Tier-1
year 2016

	CPU	Disk	Tape
ALICE	16%	12%	10%
ATLAS	44%	52%	43%
CMS	25%	22%	35%
LHCb	15%	14%	12%

Resource use at Tier-1
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 Tier-1
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 Tier-1
End of 2013

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

Table 6 Use of resources at Tier1 by the experiments for RRB year 2016 (top left) and at the end of 2015 (top right), 2014 (bottom left) and 2013 (bottom right). Data is generated from the EGI web page [8].

CPU consumption by experiment at Tier-2

	2016	2015	2014	2013
ALICE	10%	13%	15%	11%
ATLAS	49%	56%	51%	56%
CMS	36%	26%	29%	26%
LHCb	5%	5%	6%	7%

Table 7 Distribution of the time-integrated normalized CPU time by the experiments at T2 centres for RRB year 2016 and calendar years 2015, 2014 and 2013. Data is from EGI Accounting page [8].

CPU usage at CERN

	2016	2015	2014	2013
ALICE	32%	29%	24%	44%
ATLAS	12%	13%	10%	13%
CMS	24%	27%	27%	40%
LHCb	13%	13%	13%	21%

Table 8 Distribution of the fraction of the experiment's total CPU consumption which has been at CERN. Data is taken from EGI web page [8].

CERN plus Tier-1 CPU efficiency

	2016	2015	2014	2013	2012
ALICE	78%	81%	81%	84%	62%
ATLAS	83%	84%	85%	92%	91%
CMS	74%	74%	74%	84%	85%
LHCb	95%	94%	94%	95%	91%

Tier-2 CPU efficiency

	2016	2015	2014	2013	2012
ALICE	85%	83%	78%	76%	64%
ATLAS	81%	86%	87%	89%	88%
CMS	65%	70%	75%	81%	86%
LHCb	96%	96%	97%	96%	95%

Table 9 CPU efficiency for CERN plus T1 sites and for T2 sites by experiment for RRB year 2016 and for four preceding years. Data from EGI Accounting Portal [9].

number of processors. The values are obtained from the official portal [9] and take into account the multicore jobs. CMS efficiency is lower than that of any other experiment and of the previos years due to the different way CMS accounts for the scheduling overhead plus other effects as will be discussed in the CMS section.

5.2 Usage by the experiments

5.2.1 ALICE

We summarize the computing resource usage by the ALICE experiment for the period January 1st to December 31st 2016, based on the report provided by ALICE [1]. Pledged resources are extracted from REBUS, used resources are extracted from EGI accounting portal [9].

Resource	Site(s)	2016 request	2016 pledged	2016 used	2016 used/pledged	2016 CPU efficiency
CPU (kHS06)	T0 + CAF	215	216	218	101 %	73%
	T1	157	177	253	143 %	80%
	T2	237	231	255	110 %	82%
Disk (PB)	T0	16.80	16.80	13.3	79 %	
	T1	21.0	18.95	17.4	92 %	
	T2	26.1	18.51	14.0	76 %	
Tape (PB)	T0	21.6	21.6	25.5	118 %	
	T1	15.6	17.77	18.5	104 %	

Table 10 Summary of resource usage by ALICE in 2016 (January 2016 to December 2016), with 2016 pledges and 2015 requests. Requested and pledged resources are from REBUS [11]. T0 and T1 disk data includes 6 PB of disk buffer in front of the tape systems.

ALICE has fulfilled its data taking program as established for 2016. All of the 2016 objectives in terms of statistics have been reached.

The tape usage is dominated by RAW data recording. The data collected at T0 in 2016 amounts to 7.5 PB for a cumulated total of 25.5 PB exceeding by 18% the 2016 pledged resources. The accumulated data at the T1s storages amounts to 18.5 PB exceeding by 4% the 2016 pledged resources.

The disk storage resources are distributed as follows: 30% in T0, 39% in T1s and 31% in T2s. Simulation represents 40% of disk resources used, reconstruction 45% and data analysis 15%. According to ALICE usage report, 70% of the 2016 requested resources for disk have been used. In later discussion ALICE pointed out that 90% of the installed capacity is used. About 40% of the proton-proton raw data for 2016 have not been reconstructed yet.

Simulation represents 70% of CPU resources used. The adoption of Geant4 is ongoing with several MC productions done exclusively with it. The additional TPC distortion corrections double the CPU cost of reconstruction but do not significantly impact overall CPU budget as the reconstruction continues to represent only 11% of used CPU. Data analysis represents 19% of CPU resources used.

From the data popularity plot, it appears that about 11PB of data have not been accessed in the last year. ALICE is encouraged to maintain an active disk cleanup policy in order to reduce the volume of infrequently used data to maximum possible extent.

The HLT farm of ALICE has been successfully integrated into the Grid and is providing about 5% of the total CPU resources and the contribution from non WLCG sites amounts to 4% of the pledged resources.

5.2.2 ATLAS

Usage of the computing resources by the ATLAS experiment is summarized in Table 11 and is based on the report from ATLAS [3], together with the pledged resources extracted from REBUS and the used resources extracted from the EGI accounting portal.

Overall ATLAS is making high use of the pledged resources with the notable exception that only 56% of tape pledged at T1 sites was used. T1 tape usage was lower due to a delay in the MC reconstruction planned for the fall of 2016. We also note that ATLAS has been able to use twice the amount of CPU resources pledged by T2 sites and that they also managed to make efficient use of the HLT farm. In addition to WLCG sites, ATLAS are making use of resources provided by non-WLCG sites such as HPC and cloud resources, amounting to approximately 15% of the total amount of available CPU wall-clock time.

Computation continues to be dominated by Monte Carlo simulation and reconstruction. The introduction of "derivation trains" has managed to keep the amount of CPU resources spent on ad-hoc user analysis at a constant level despite an increase in the amount of data. The issues previously encountered with the CPU efficiency at the CERN T0 have been resolved after CERN IT installed hardware without hyper-threading, increasing the available memory per core.

Efficiency of the ATLAS software on T1 and T2 facilities remains high with approximately 80% utilization. Significant work has been undertaken to reduce the number of copies of DAODs based on the popularity charts and to reduce the amount of Run-1 data and MCs to 6 PB.

Use of fast simulations remains below predictions (by 50%) due to a focus on commissioning Geant4 V10 with an expectation that increasing the number of fast simulations would be a focus for 2017 and developing fast chains (the full simulation through to derivation) for Run-3.

Table 11 shows an overview of ATLAS' resource usage for 2016.

Resource	Site	Pledged	Used	Used/Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	235	241	103%	87%
	T1	538	642	129%	82%
	T2	610	1235	202%	80%
	HLT	22	56	255%	76%
Disk (PB)	T0+CAF	17	14	82%	
	T1	53	48	91%	
	T2	69	72	104%	
Tape (PB)	T0	42	41	98%	
	T1	119	67	56%	

Table 11 Fulfillment of pledges. The table reports the ATLAS situation at the end of 2016. Data from the master accounting summary in the WLCG document repository [8].

5.2.3 CMS

During 2016 CMS had to operate in a challenging situation due to the large data delivery and the deficit of resources at T1. The software and computing group took significative actions that mitigated the deficit but still the resources available are not at the required level. CMS performed a massive tape deletion campaign at T1 that made available ~30 PB of space. In addition an aggressive disk clean-up was done at T1 where the space was below the safe point.

CMS has created the "ECoM-17 Evolution of the Computing Model" committee, a group with experts from Physics, Trigger, R&D and Software and Computing to evaluate the computing resource needed due to the increased data volume and come up with an evolved strategy before Run3.

CMS continues to evolve the computing model by increasing the flexibility of the workflows executable at each tier. The HTCondor Global Pool handles multi-core payloads that are run at all T1. Multi-core processing is possibile also at almost all T2. We noted that the CPU efficiency is lower with respect to the past years. This is due to the scheduling overheads, switching from single-core to multi-core and to the fact that idle pilot time may be accounted differently among various experiments. In fact the payloads CPU efficiency is around 80%.

The HLT farm has been commissioned as an opportunistic resource as well as during the interfill, being capable of providing 15k cores during these periods. On average in 2016 HLT provided 50 kHS06 averaged contributing for 14% of the total T1 CPU.

CMS used the Amazon Web Services commercial cloud to expand the available resources by 25% for two weeks to produce Monte Carlo events for conferences. In addition the Google Cloud Platform was exploited to produce Monte Carlo samples for Moriond 2017. The use of this resource was a proof of concept, demonstrating the elastic nature of the Cloud, and was available for a short period of time.

The C-RSG is concerned about the gap which exists between the requested resources and the actually pledged resources. Particularly, the amount of tape space does not cover the request of the experiment. Table 12 shows an overview of CMS resource usage for 2016. C-RSG appreciates the reports made by CMS, which show a good understanding of the interplay of resources of various types (CPU, Disk, Tape) deployed at the various Tiers.

Resource	Site	Pledged	Used	Used/Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	306	260	85%	81%
	T1	348	358	103%	68%
	T2	677	872	129%	65%
	HLT	-	51	-	80%
Disk (PB)	T0+CAF	16	15	90%	
	T1	30	30	100%	
	T2	41	37	90%	
Tape (PB)	T0+CAF	44	40	91%	
	T1	88	73	83%	

Table 12 Fulfillment of pledges. The table reports the CMS situation at the end of 2016. Data from the master accounting summary in the WLCG document repository [8]. Note that CMS manually keeps T2 disk usage at 90% utilization, so the T2 disk utilization does not indicate lack or surplus of T2 disk resources

5.2.4 LHCb

The report covers all of 2016 (January to December) and is based on a report from LHCb [5]. Activities in 2016 included simulation, user analysis, an incremental stripping of Run-1 data, validation cycles of 2015 TURBO data, reconstruction of data taken in 2015 in proton-ion collisions and processing of data taken in 2016 proton-proton collisions.

Several unanticipated factors impacted the use of storage and compute resources by LHCb in 2016: the LHC live time in 2016 was considerably higher than originally anticipated and put strain on resources, the stripping output used a bandwidth of 120MB/s rather than the design of 100MB/s, and the size of the TURBO stream was 50kB/event instead of 10kB/event. Dataset parking, reduction of the number of stored copies and the removal of unused dataset mitigated the impact of these factors.

Resource	Site(s)	2016 pledge	2016 used	Used/pledged	Average CPU efficiency
CPU (kHS06)	T0	51	31.4	62%	91%
	T1	165	201.4	122%	95%
	T2	88.6	115.6	130%	96%
Disk (PB)	T0	7.6	4.59	60%	
	T1	15.9	13.0	82%	
	T2	2.7	2.8	103%	
Tape (PB)	T0	20.6	16.6	81%	
	T1	35.0	22.8	65%	

Table 13 2016 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.

RRB year	pp/10^6 s	HI/10^6 s	pp pileup
2015	3	0.7	25
2016	5	0.7	35
2017	7.8	-	35
2018	7.8	0.7	35

Table 14 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.

6 Resource requirements for 2018

6.1 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 [12]. Table 14 reports the anticipated LHC beam live times updated to the latest official schedule [12]. 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 1.7×10^{34} cm^{-2} s^{-1} for 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.

7 Resource requirements from the experiments

7.1 ALICE

Table 15 summarizes the requests from ALICE for 2018 and the CRSG recommendations, where an agreement with the experiment could be reached. CERN CAF resources are included in the T0 requests.

We noticed that pledges for ALICE are not at the level of their requests, in particular at T2 since few years pledges are below the requests which appears now as large requests. Taking as a reference, the pledged resources for 2017 in table 15, ALICE requests for 2018 go far beyond the goal of keeping a

Resource	Site	2017 pledge	2018 ALICE	Growth	2018 CRSG	Growth
CPU (kHS06)	T0+CAF	292	350	20%	350	20%
	T1	235	306	30%	306	30%
	T2	278	438	58%	-	-
Disk (PB)	T0+CAF	22	27	23%	27	23%
	T1	21.8	32	47%	-	-
	T2	22.5	41	82%	-	-
Tape (PB)	T0+CAF	36.9	55	49%	55	49%
	T1	30.6	41	34%	41	34%

Table 15 ALICE resources requests and CRSG recommendations. 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.

flat budget. However, the ALICE computing coordinator made it very clear that the increased requests are needed to achieve the physics program endorsed by LHCC.

From numerous discussions, it appears there are limited possibilities to further optimize the resources needed for ALICE to achieve its scientific goals.

Compared to the resources pledged in 2017 by the funding agencies, tape requests for 2018 represent an increase by 49% for T0 and 34% for T1. The CRSG perception is that these requests for tape are acceptable because tape are mandatory for storing raw data and the unit cost of tape is relatively cheap.

However, the other requirements for 2018, specifically: T1-disk (47% increase), T2-disk (82%) and T2-cpu (58%) are well beyond "flat budget" and we noticed that T2 pledges are more than 20% less than the request. CRSG is not in a position to accept these requests and asks the experiment, the LHCC and the FA to take the actions necessary to reduce the deficit, which imply to increase the pledges and reduce the resource needs. We propose to scrutinize these requests in October. CRSG is willing to follow and to help in this process if asked.

ALICE usage statistics show that only about 80% of the disk pledged for 2016 has been used. In later discussion ALICE pointed out that 90% of the installed capacity is used. A flat budget scenario for the disk in 2018 is not perceived by CRSG as a risk for the physics program.

Regarding CPU, confusion was raised by a discrepancy between the numbers provided by ALICE and CERN for T0 resource usage in 2016. A plausible explanation for the discrepancy was found on April 13th 2017 by the ALICE computing coordinators that results in a significant reevaluation of the CPU usage compared to the numbers provided by ALICE in their initial usage report. These new numbers are documented in table 10.

7.2 ATLAS

Table 16 shows the pledged resources for ATLAS for 2017, the request for 2018, and the recommendations of the CRSG. Using the 2017 pledges as a baseline the ATLAS requests for 2018 represent a 15% increase in CPU (over all tiers), an 8% increase in disk (over all tiers), and a 15% increase in tape (for T0 and T1).

These requests are consistent with the expectations of a flat budget and lower respect to the projections of growth in ATLAS resources over the last four years (see Figure 2). We note that the request for T0 CPU resources and T0 and T1 disk resources are substantially below the expectations for 2018 assuming a flat budget.

The CRSG perspective on these increases is that the requested resources are acceptable and we recommend that they be fully supported. We note, however, a number of concerns and comments about the request below.

ATLAS expects to continue to receive substantial amount of over pledge CPU from HPC, Grid, and Cloud resources. These beyond pledge resources currently amount to about 30% of the pledge resources for 2017. The continued reliance on these resources (offsetting CPU requests for increases in disk and tape) remains a risk for the experiments.

We also note that the reason ATLAS is not requesting significantly larger increases in resources in 2018 is because the 2017 pledges met the needs of the experiment (in contrast to CMS).

Resource	Site	2017 Pledge	2018 ATLAS	Growth	2018 CRSG	Growth
CPU (kHS06)	T0+CAF	404	411	2%	411	2%
	T1	808	949	17%	949	17%
	T2	982	1160	18%	1160	18%
Disk (PB)	T0+CAF	25	26	4%	26	4%
	T1	69	72	4%	72	4%
	T2	78	88	13%	88	13%
Tape (PB)	T0+CAF	77	94	22%	94	22%
	T1	174	195	12%	195	12%

Table 16 ATLAS resources request and CRSG recommendations.

7.3 CMS

This report is based on the original submission by CMS [4]. Table 17 shows the latest CMS computing requests for 2017 and 2018 along with the corresponding CRSG recommendations.

CMS has taken several actions to reduce the needs for 2018. The number of replicas of AOD samples has been reduced as well as the amount of RECO and RAW on disk. In addition part of the Run-1 Monte Carlo has been removed from disk. However, despite those steps, CMS still requests disk in excess of what is expected from a flat budget. The major reason for this is a deficit resulting from the fact that CMS was not pledged the disk resources they requested for 2017 (150 PB requested and 123 pledged). The C-RSG, therefore, still endorses the request for disk in order to bring CMS back to a level where they can cope within the needs. We encourage CMS to take the actions with T1 and T2 necessary to fill the gap.

As for tape CMS has for quite a number of years, as pointed out in several reports from C-RSG, suffered from a deficit between requested and pledged tape resources. This continues for the 2018 request. To mitigate the rather large request for T1 tape, C-RSG suggests to move some of the requested tape resources from T1 to T0. It has be noted that at T0 CMS is not asking for any increase in CPU and the requested disk space is kept as low as possible.

7.4 LHCb

This report is based on the original submission by LHCb [6]. Table 18 shows the latest LHCb computing pledge for 2017 and 2018 resources requests along with the corresponding CRSG recommendations.

Since the last report, LHCb has changed its anticipated stripping size from 120MB/s to 165MB/s. The stripping strategy has been adjusted to compensate. The TURBO format continues to be larger

Resource	Site	2017 Pledge	2018 CMS	Growth	2018 CRSG	Growth
CPU (kHS06)	T0+CAF	423	423	0%	423	0%
	T1	515	600	17%	600	17%
	T2	791	900	14%	900	14%
Disk (PB)	T0+CAF	25	26	6%	26	6%
	T1	45	60	34%	60	34%
	T2	53	70	32%	70	32%
Tape (PB)	T0+CAF	71	80	13%	97	36%
	T1	133	205	54%	188	41%

Table 17 CMS resources request and the CRSG recommendations.

per live second than originally planned. To compensate, LHCb has decided to "park" 35% of their TURBO stream for the duration of Run-2.

Based on 2017 pledge , the 2018 requests from LHCb are roughly consistent with a flat budget, as requested by the funding agencies. An exception is the tape request. While the CPU request increases by approximately 15% averaged over all Tiers compared to 2017 pledge and disk increases 21%, the tape request increases by more than 40% in 2018 compared to 2017 pledge. Significant increases in tape requirement were anticipated for 2017 and 2018 in previous reports as a result of the increase in LHC live time. From this new base, the projected increase in tape for 2019 is only 7%. The use of tape in this way allows some of the Run-2 data to be "parked" for later processing and mitigates the growth in CPU requirements.

Resource	Sites(s)	2017 Pledge	2018 LHCb	Growth	2018 CRSG
CPU(kHS06)	T0	67	81	21%	81
	T1	199	253	27%	253
	T2	147	141	-4%	141
	HLT + Yandex	20	20	-	20
Disk (PB)	T0	10.9	12	12%	12
	T1	20.8	24.5	20%	24.5
	T2	3.3	5.8	-	5.8
Tape (PB)	T0	25.2	36.4	44%	36.4
	T1	41.9	61.5	47%	61.5

Table 18 LHCb resources pledge for 2017, requests for 2018 and CRSG recommendations. T1 and T2 disk space increase is considered together due to the small amount disk available at T2.

8 Comments and recommendations

- 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.
- From an analysis of the last four years of operations, the assumption of a flat budget with a yearly 20%, 15%, and 15% increase in CPU, disk, and tape respectively is not consistent with the historical pledge resources and we recommend a reevaluation of the assumptions of what a flat budget entails.
- The experiments continue to increase the efficiency of their MC generation and analyses as well as reducing the amount of data stored on disk (based on lifetimes and usage). It is not clear that there is substantially more efficiency that can be gain without extensive reworking of the simulations and analysis frameworks (e.g. the introduction of fast chains for the simulations).
- The CRSG notes that some experiments have been particularly successful in securing non-WLCG cpu resources and we encourage all experiments to pursue this. To help monitor this, we recommend that all experiments quantify more fully the non-WLCG resources that they have obtained in their future reports. Furthermore we welcome the fact that every experiment has made use of their HLT farms to augment their cpu resources.

References

[1] ALICE Collaboration, ALICE 2016 computing resources usage. Spring 2017

[2] ALICE Collaboration, ALICE Update on computing resources requirements: 2017-2019. Spring 2017

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

[4] CMS Collaboration, CMS Software/Computing Resource Utilization & Request 2017-2019, Spring 17. 1 February 2017

[5] LHCb Collaboration, C. Bozzi(Ed), LHCb Computing Resource Usage in 2016(II), LHCb–PUB–2017–010. 19 February 2017

[6] LHCb Collaboration,C. Bozzi (Ed), LHCb Computing Resources: 2018 requests and preview of 2019 requests LHC–PUB–2017–009. 23 February 2017

[7] 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

[8] EGI Accounting Portal, https://accounting.egi.eu/wlcg/report/tier1/

[9] EGI Accounting Portal, http://accounting.egi.eu/

[10] REBUS Pledge Fulfillment, CERN, https://wlcg-rebus.cern.ch/apps/acc_report/customised_report/?from=April%202016&months=11&tables=cpu&graphs=cputime&federations=CH-CERN&experiments=ALICE,ATLAS,CMS,LHCb,total&accsum=1&expsum=1

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

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