



Computing Resources Scrutiny Group Report

Donatella Lucchesi
for the CRSG
CERN, October 24th 2017

Report:
CERN-RRB-2017-216

C-RSG membership

2

Membership of the CRSG changed for this scrutiny:

C Allton (UK)

J Kleist (Nordic countries)

V Breton (France)

D Lucchesi (Italy, Chairman)

G Cancio Melia (CERN)

H Meinhard (CERN, scient. secr.)

A Connolly (USA)

D O'Neil (Canada)

M Delfino (Spain)

J Templon (Netherlands)

F Gaede (Germany)

D O'Neil, Canada representative ended his period, thank you Dugan for your important contribution over the years. We ask the Canadian Funding Agency to nominate a replacement.

Thanks to CRSG members for their commitment, to the experiment representatives for their collaboration and to CERN management for the support.

Pledged Resources for 2018

3

The experiments has made a large effort to keep the resource needs within *flat budget*.

$$Pledges\ balance = \frac{total\ offered - experiment\ required}{experiment\ required}$$

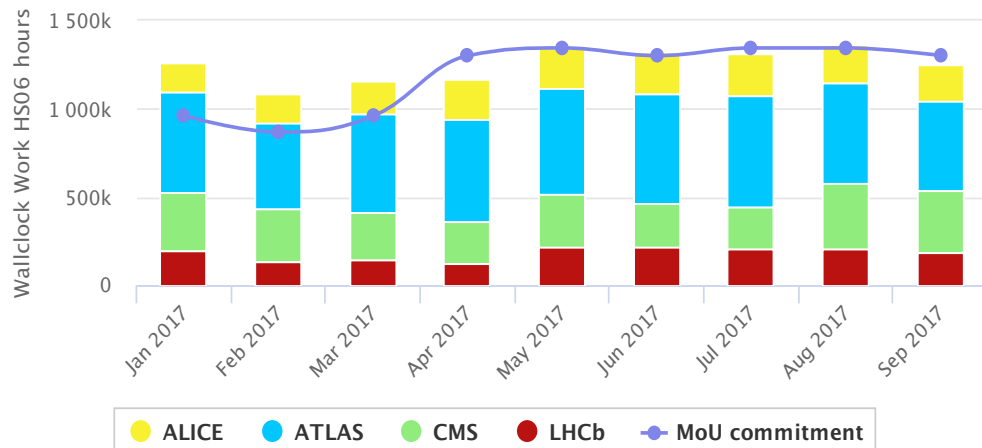
2018	ALICE	ATLAS	CMS	LHCb
CERN CPU	0%	0%	0%	0%
CERN disk	0%	0%	0%	0%
CERN tape	0%	0%	0%	0%
Tier-1 CPU	-9%	2%	-6%	-1%
Tier-1 disk	0%	11%	-8%	7%
Tier-1 tape	3%	1%	-12%	12%
Tier-2 CPU	-21%	-2%	3%	16%
Tier-2 disk	-18%	-2%	-6%	-36%

2017	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%

Overall Resource Use

4

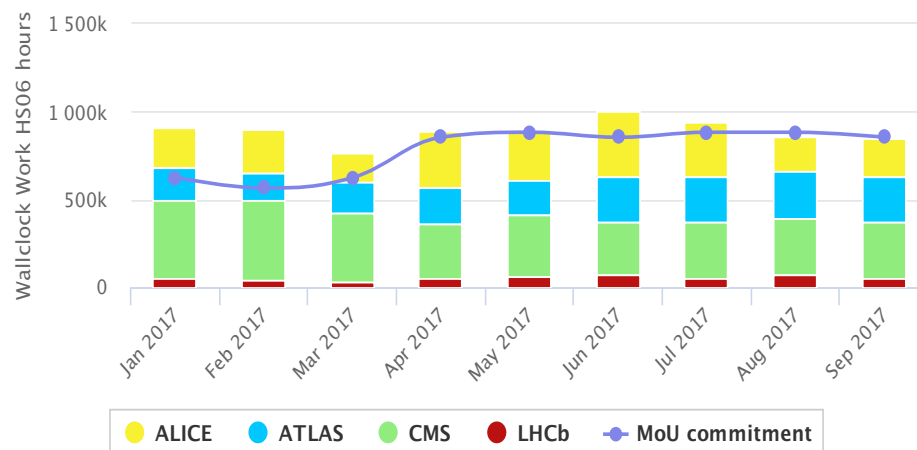
All Tier-1s : Wallclock Work*



January – September 2017

CPU resources fully used

CH-CERN: Wallclock Work*

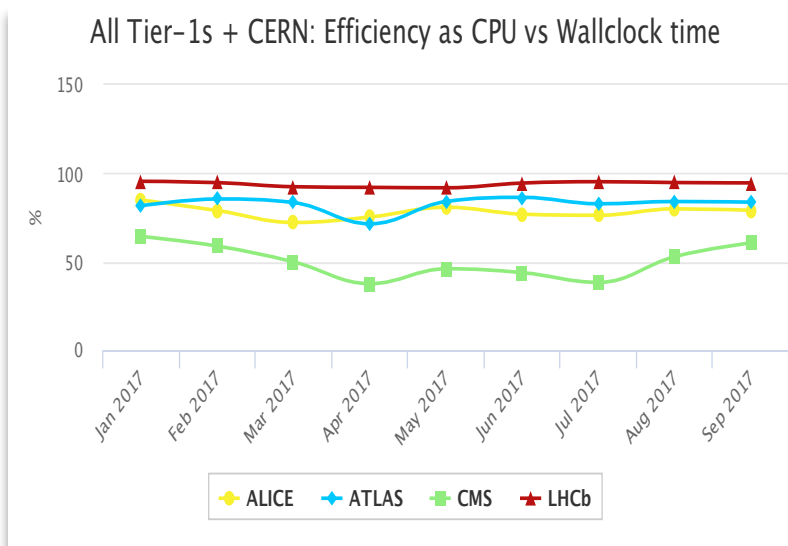
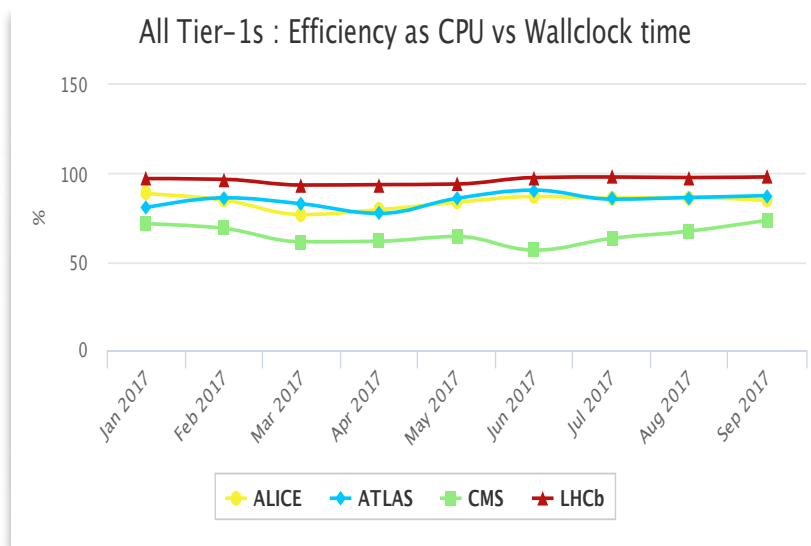


CPU Efficiency

5

$$CPU\ Efficiency = \frac{CPU\ Time}{Wallclock\ Time}$$

Question: is this the right parameter to optimize? Would be better to look at the “physics output” considering also the disk space usage?



CMS is addressing issues related to inefficiency.

ATLAS and CMS are working to improve the system performances including data movement.

Data taking assumptions

6

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

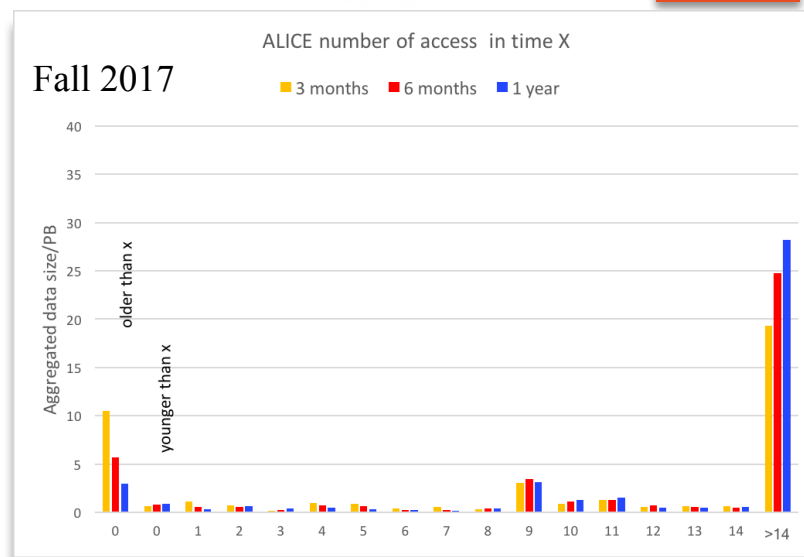
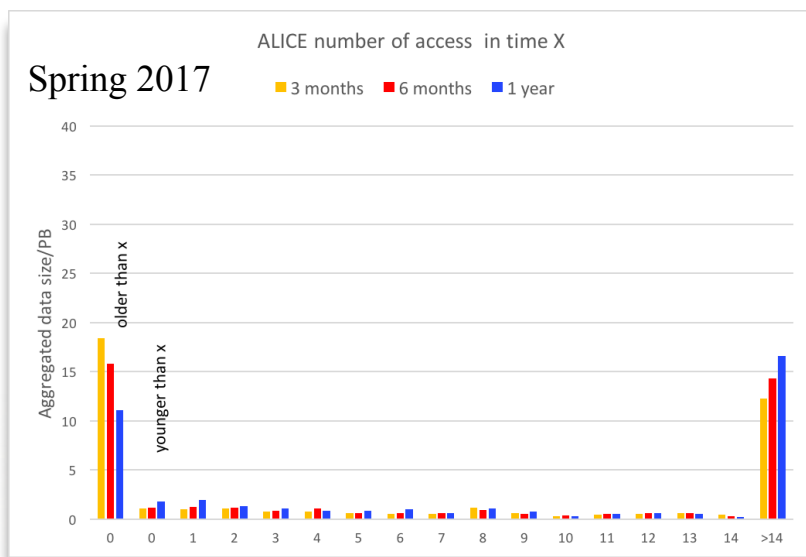
Peak luminosity $1.9 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. It may be pessimistic... while the pileup may be too optimistic.

Higher values of pile-up will increase the event complexity, in principle pile-up and CPU scale linearly but both ATLAS and CMS are working to mitigate it.

No data taking is foreseen in 2019 in the following assumption.

ALICE: Disk Space Use

7



Volume of data vs. n° of access in 3-, 6-, 12-months.

1st bin: data created before the period and not accessed in the period.

2nd bin: data created in that period and not accessed in the period.

Disk space clean-up at Tier-1 and Tier-2 allowed to recover 7PB of space

➡ 2018 resource request reduced

❑ C-RSG recommends ALICE to keep the disk space monitored to avoid unused data on disk.

- ❑ ALICE recovered 7 PB of tape space by reducing the pp event size from 2.87 MB to 1.7MB thanks:
 - Change of the gas mixture in the TPC \Rightarrow spurious clusters
 - HLT compression of TPC
- ❑ The CPU at Tier-2 went from 438 kHS06 to 398 kHS06 for improvements in event reconstruction, budgeted at Tier-2 to mitigate the deficit.
- ❑ The initial scrutiny document and experiment requirements were modified following the exchange of C-RSG reviewers and ALICE management:
 - Bugs fixes
 - Monte Carlo events for real collision for pp and Pb-Pb
 - Error in the 2019 assumptions
- ❑ The determination of the appropriate Monte Carlo to data ratio goes beyond the mandate of the C-RSG and has to be discussed with the LHCC.

Alice Scrutiny 2018

9

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%

❑ C-RSG for 2018:

- Accepts tape reduction and Tier-0 disk reduction
- Accepts Tier-1 disk space already pledged by Funding Agency
- Does not have enough information to make recommendation on CPU and disk space at Tier-2

Alice Scrutiny 2019

10

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%

- ❑ C-RSG postpones any evaluation of 2019 resources until Spring scrutiny. C-RSG will interact with the LHCC for the determination of Monte Carlo to data value.
- ❑ 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.

ATLAS Scrutiny 2018

11

Resource	Site	2017 CSRG	2017 Pledge	2018 Re- quest	2018 CSRG	2018 Growth over 2017	2019 Re- quest	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%

- ❑ From January to July 2017 ATLAS average CPU utilization exceeds the pledge values by over 50%
- ❑ No changes from April RRB for 2018
- ❑ ATLAS investigating further workflows optimization (ex. Geant4 compilation optimization) and extension of AOD size reduction (30%) to DAOD

ATLAS Scrutiny 2019

12

Resource	Site	2017 CSRG	2017 Pledge	2018 Re- quest	2018 CSRG	2018 Growth over 2017	2019 Re- quest	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%

- ❑ Major activities foreseen in 2019:
 - ❑ Run-2 data reprocessing
 - ❑ Large Monte Carlo production that will span to 2020
 - ❑ Monte Carlo for Phase-II detector TDR
- ❑ Request based on same processing times and event sizes, 700 kHS06 from opportunistic resources not included
- ❑ Tier-0 and HLT (50%) used for MC production and data re-processing

CMS Scrutiny 2018

13

Resource	Site	2017 CRSG	2017 Pledge	2018 Spring request	2018 CRSG	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%

- ❑ No changes from April RRB for 2018 other than +2PB of tape at Tier-0 for Pb-Pb data not considered before
- ❑ CMS working on additional improvements:
 - ❑ Use MiniAOD for high level studies
 - ❑ Reduction of AOD copies from 1 to 0.7 at Tier-2 with 0.3 at Tier-1

CMS Scrutiny 2019

14

Resource	Site	2017 CRSG	2017 Pledge	2018 Spring request	2018 CRSG	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%

- ❑ Major activities foreseen in 2019:
 - ❑ 2016-2018 data reprocessing, 2015 under discussion
 - ❑ Large Monte Carlo production that will span to 2020
 - ❑ Monte Carlo for Phase-II detector TDR
- ❑ Tier-0 and HLT (80%) used for MC production and data re-processing
- ❑ CMS express concerns on reliability/capability of EOS and P5 networking to support the Tier-0 and HLT data reconstruction

In view of Phase-II computing C-RSG would like to monitor the extra-pledged resources therefore requires that ATLAS and CMS experiments document and report separately their use of overpledge and opportunistic resources including the source of these resources (e.g. grid farms, and HPC centers).

LHCb Scrutiny 2018

16

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%

- ❑ Due to reduced LHC efficiency, less data is recorded and LHCb anticipated some operations originally postponed to 2019.
- ❑ LHCb completed the micro-DST format development that will allow stored data reduction.
- ❑ Consequences: reduction of tape requests of 8% and 26% at Tier-0 and Tier-1s and small decrease in disk space at Tier-0 and Tier-2

LHCb Scrutiny 2019

17

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%

❑ Main activities will be:

❑ Simulation → dominates CPU request

❑ Data re-stripping → disk space request

❑ LHCb plans to develop a single platform to access any kind of dataset for analysis preservation. CERN Tier-0 is considered a very good choice by C-RSG. This justify the small CPU request in 2018 and 2019 at Tier-0

- ❑ C-RSG, LHCC and WLCG, propose to update procedures for reviewing experiments' requests:
 - Prior to the year N Autumn RRB, experiments make initial resource requests for year N+2 and, if possible, a preview for year N+3;
 - A six month review starts, involving C-RSG, LHCC and experiments, with a possibility to review physics scope;
 - 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
- ❑ 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.