

Computing Resources Scrutiny Group Report

Donatella Lucchesi
University and INFN of Padova

For the Computing Scrutiny Group

C-RSG membership

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, scient. secr.)

P Sinervo(Canada)


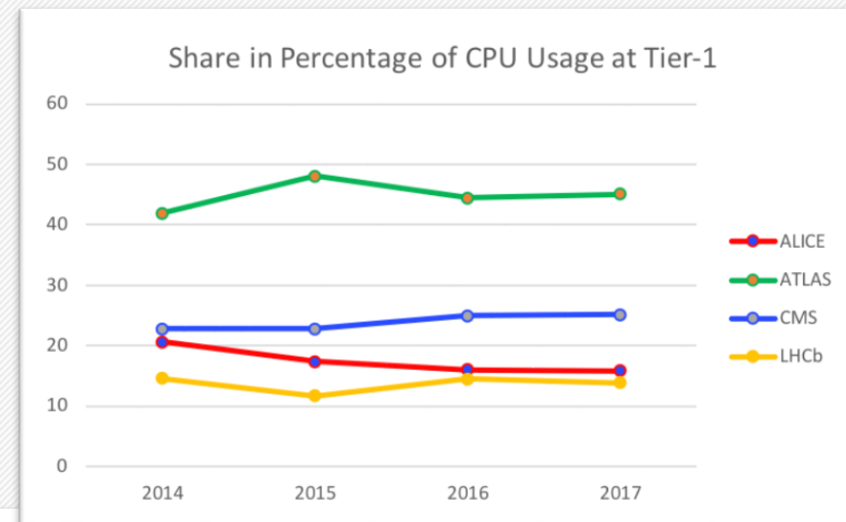

J Templon (Netherlands)

- **P Sinervo**, is the new Canada representative. We welcome Sinervo, who actively participated to this scrutiny round.
- The group thanks Jeff Templon for his valuable contribution during the years and ask the Netherland to nominate a new representative.
- Thanks to the experiment representatives for their collaboration and to CERN management for the support.
- I want to thank the C-RSG members for their commitment, it is a great team!

General Usage - CPU

For each Tier, Percentage, $y = (CPUUsed)_{experiment} / \sum_{All-experiments} (CPUUsed)$



3



General Usage – Disk Space



For each Tier, Percentage, $y = (DISKused)_{experiment} / \sum_{All-experiments} (DISKused)$

4




General Usage – Tape Space

For each Tier, Percentage, $y = (TAPeUsed)_{experiment} / \sum_{All-experiments} (TAPeUsed)$



General Usage

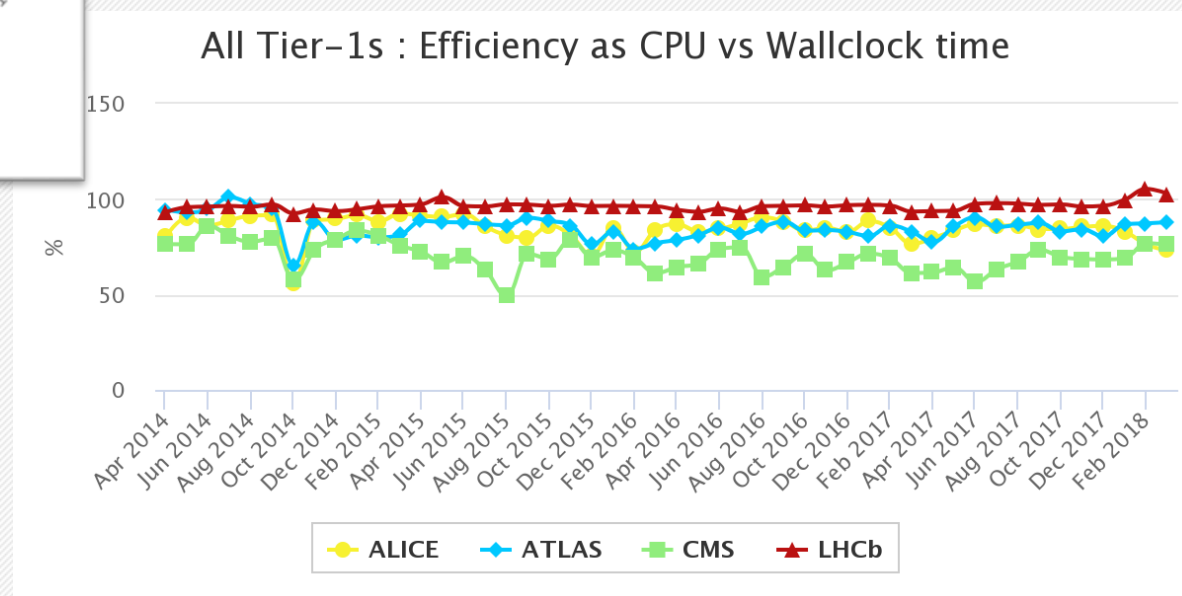

$$y = \frac{(CPUUsed@CERN)_{experiment}}{\sum_{All-experiments, All-Tiers}(CPUUsed)}$$



		Used/pledged resources			
		2017	2016	2015	2014
CPU	CERN	105%	122%	39%	53%
	T1	97%	119%	102%	123%
	T2	143%	151%	111%	152%
Disk	CERN	72%	97%	80%	81%
	T1	88%	72%	82%	95%
	T2	—	—	—	—
Tape	CERN	64%	98%	76%	96%
	T1	53%	67%	69%	89%

CPU Efficiency

7




Usage experiment by experiment

8

ALICE Usage

ALICE		2017				
		CRSG recomm.	Pledged	Pledged /CRSG	Used	Used /CRSG
CPU	Tier-0	292	292	100%	389	133%
	Tier-1	256	235.5	92%	295	115%
	Tier-2	366	279.6	76%	299	82%
	HLT	n/a	n/a	n/a	26	n/a
	Total	914	807.1	88%	1010	110%
	Others			n/a	n/a	
Disk	Tier-0	22.4	22.4	100%	19.3	86%
	Tier-1	25.4	21.8	86%	18.245	72%
	Tier-2	31.4	22.7	72%	20.06	64%
	Total	79.2	66.9	84%	57.6	73%
Tape	Tier-0	36.9	36.9	100%	29.7	80%
	Tier-1	30.9	30.6	99%	22.3	72%
	Total	67.8	67.5	100%	52	77%

- Good use of opportunistic CPU resources.
- More cycles could be exploited on HPC systems, they are not of easy access by ALICE.
- Disk space still under-pledged but much less than last year, ALICE thinks this does not impact performances.




- Raw event size reduced thanks to improved HLT compression ⇒ less tape space needed
- 2017 unused tape will be accounted for in 2018 and 2019 requests.

ATLAS Usage

ATLAS		2017				
		CRSG ² recomm.	Pledged	Pledged ² /CRSG	Used	Used ² /CRSG
CPU	Tier-0	404	367	91%	421	104%
	Tier-1	921	786	85%	826	90%
	Tier-2	1125	917	82%	1505	134%
	HLT	n/a	22	n/a	120	n/a
	Total	2450	2092	85%	2872	117%
	<i>Others</i>			230	9%*	
Disk	Tier-0	25	25	100%	25,0	100%
	Tier-1	68	70	103%	67	99%
	Tier-2	83	78	94%	78	94%
	Total	176	173	98%	170,0	97%
Tape	Tier-0	77	77	100%	58	75%
	Tier-1	188	175	93%	102	54%
	Total	265	252	95%	160	60%

- Very high use of opportunistic CPU resources.
- Sizeable amount of CPU cycles from HPC farms, contributions also from volunteer computing.
- Disk space ~100% full.
- Increase of *not accessed data in last 6 months* due to large dataset not fully ready for analysis.




- Tape usage less than expected due to: smaller cosmic and calibration dataset, deleted Run 1 non-raw events and better than expected lifetime model at Tier-1.

CMS Usage

CMS		2017				
		CRSG ² recomm.	Pledged	Pledged ² /CRSG	Used	Used ² /CRSG
CPU	Tier-0	423	397	94%	326	77%
	Tier-1	600	470	78%	425	71%
	Tier-2	850	772	91%	1133	133%
	HLT	n/a	n/a	n/a	148	n/a
	Total	1873	1639	88%	2032	108%
	Others			30	2%*	
Disk	Tier-0	24,6	25	102%	21,0	85%
	Tier-1	57	45	79%	39	68%
	Tier-2	68	54	79%	48,6	71%
	Total	149,6	124	83%	108,6	73%
Tape	Tier-0	70,5	70	99%	49	70%
	Tier-1	175	133	76%	111	63%
	Total	245,5	203	83%	160	65%

- Challenging 2017 due to incident at CNAF, issues and less pledged resources than recommended.
- Mitigations came from opportunistic CPU resource and tape deletion campaign @Tier-1
- Very good use of disk space: dynamic storage management and reduced data formats



Mitigations were possible thanks to additional operational efforts


LHCb Usage

LHCb		2017				
		CRSG recomm.	Pledged	Pledged /CRSG	Used	Used /CRSG
CPU	Tier-0	67	67	100%	63	94%
	Tier-1	207	199	96%	245	118%
	Tier-2	116	147	127%	196	169%
	HLT	10	0	0%	208	2080%
	Total	400	413	103%	712	178%
	Others			114	29%*	
Disk	Tier-0	10.9	10.9	100%	6.6	61%
	Tier-1	22.1	20.9	95%	16.9	76%
	Tier-2	4.7	3.3	70%	3.6	77%
	Total	37.7	35.1	93%	27.1	72%
Tape	Tier-0	25.2	25.2	100%	22.5	89%
	Tier-1	43.3	42.0	97%	31.3	72%
	Total	68.5	67.19	98%	53.8	79%

- Very good use of not pledged CPU resources, and HLT
- Substantial contribution from non WLCG CPU resources used for non critical tasks as additional Mont Carlo production
- Refinement of the TURBO data format helps optimize storage use, which results not fully utilized here but peak utilization happened during the year

*Percentage taken with respect to Total CRSG CPU recommendation

CNAF incident had a substantial impact on LHCb, ~18% of data inaccessible for many weeks.



Resource Requirements for 2019

13

- 2019 – shutdown year.
- Tier-0 management expects CERN resources pledged in 2019 identical to 2018 with the intention to re-profile a maximum of funding to the beginning of Run 3.
- Tier-0: ATLAS and CMS no resource increase, LHCb modest and ALICE substantial increase.

Alice Requests for 2019

ALICE		2017			2018		2019		
		CRSG recomm.	Pledged	Used	CRSG recomm.	Pledged	Request	2019 req. /2018 CSG	C-RSG recomm.
CPU	Tier-0	292	292	389	350	350	430	123%	430
	Tier-1	256	235.5	295	307	279.5	365	119%	365
	Tier-2	366	279.6	299	312.9	312.9	376	120%	376
	HLT	n/a	n/a	26	n/a	n/a	n/a	n/a	0
	Total	914	807.1	1010	969.9	942.4	1171	121%	1171
	Others								39
Disk	Tier-0	22.4	22.4	19.3	26.2	26.2	34.3	131%	34.3
	Tier-1	25.4	21.8	18.245	30.5	30.4	37.9	124%	37.9
	Tier-2	31.4	22.7	20.06	29	29	33.9	117%	33.9
	Total	79.2	66.9	57.6	85.7	85.6	106.1	124%	106.1
Tape	Tier-0	36.9	36.9	29.7	49.1	49.1	44.2	90%	44.2
	Tier-1	30.9	30.6	22.3	40.9	42.2	37.7	92%	37.7
	Total	67.8	67.5	52	90	91.3	81.9	91%	81.9

- No C-RSG recommendations for 2018 Tier-2 CPU and disk. The pledged values are assumed as recommended.
- Requests changed substantially since October 2017: -25% for CPU, and -12% for disk space.
- Tier-0 CPU and disk space requests are due to Pb - Pb data taking with event data size increased due to new trigger mix.

ALICE Recommendations

15

- C-RSG strongly encourages ALICE to keep working to reduce the size of Pb – Pb data and associated Monte Carlo
- C-RSG welcomes ALICE's work on Geant4 multithreading and strongly encourage further developments that will better match the opportunistic resources, in particular HPC

ATLAS Requests for 2019

ATLAS		2017			2018		2019		
		CRSG ² recomm.	Pledged	Used	CRSG ² recomm.	Pledged	Request	2019 ³ req. ² /2018 ² CRSG	C-RSG ² recomm.
CPU	Tier-0	404	367	421	411	411	411	100%	411
	Tier-1	921	786	826	949	969	1057	111%	1057
	Tier-2	1125	917	1505	1160	1136	1292	111%	1292
	HLT	n/a	22	120	n/a	0	0	n/a	n/a
	Total	2450	2092	2872	2520	2516	2760	110%	2760
	Others			230					
Disk	Tier-0	25	25	25,0	26	27,0	26	100%	26,0
	Tier-1	68	70	67	72	80,0	88	122%	88,0
	Tier-2	83	78	78	88	86	108	123%	108,0
	Total	176	173	170,0	186	193	222	119%	222,0
Tape	Tier-0	77	77	58	94	105,0	94,0	100%	94,0
	Tier-1	188	175	102	195	196,0	221	113%	221,0
	Total	265	252	160	289	301	315	109%	315,0

- Sizeable contribution from beyond-pledge CPU resources is expected
- Disk space request at Tier-1 and Tier-2 are slightly high, but given ATLAS computing model this appears a critical request
- Even if tape occupancy is below pledged, ATLAS expects tape usage to catch up with pledges over 2018, 2019. Additional space is for contingency

ATLAS Recommendations

17

- C-RSG is concerned about the growth in the ATLAS disk usage and reliance on uncommitted CPU resources.
- C-RSG notes that ATLAS has significantly more disk resources than CMS on the Tier-1 and Tier-2 sites.
- The difference continues to grow in the 2019 requests: +29% at Tier-1 and + 38% at Tier-2.
- ATLAS explains this by differences in the computing model such as the data format and event sizes.
- C-RSG encourages ATLAS to consider investigating smaller data formats, higher compression rates and/or virtual data for fast simulation streams, similar to what CMS has achieved.

CMS Requests for 2019

CMS		2017			2018		2019		
		CRSG recomm.	Pledged	Used	CRSG recomm.	Pledged	Request	2019 Req. /2018 CRSG	C-RSG recomm.
CPU	Tier-0	423	397	326	423	423	423	100%	423
	Tier-1	600	470	425	600	562	650	108%	650
	Tier-2	850	772	1133	900	940	1000	111%	1000
	HLT	n/a	n/a	148	n/a	n/a	n/a	n/a	n/a
	Total	1873	1639	2032	1923	1925	2073	108%	2073
	Others			30					
Disk	Tier-0	24.6	25	21.0	26	26.1	26.1	100%	26.1
	Tier-1	57	45	39	60	55.4	68	113%	68.0
	Tier-2	68	54	48.6	70	66.7	78	111%	78.0
	Total	149.6	124	108.6	156	148.2	172.1	110%	172.1
Tape	Tier-0	70.5	70	49	99	97.0	99.0	100%	99.0
	Tier-1	175	133	111	188	166.0	220	117%	220.0
	Total	245.5	203	160	287	263	319	111%	319.0

- CMS requests are minimal.
- Since October RRB C-RSG notes a reduction in tape space request at Tier-1.

CMS Recommendations

- C-RSG notes that the 2018 pledge is now less than 5% below for disk and 8% for tape compared to C-RSG recommendations, which lowers previous C-RSG concerns about the lack for pledged resources to CMS.
- C-RSG applauds CMS's work on using reduced size data-formats such as miniAOD and their investigations into even more compact nanoAOD.
- C-RSG recommends CMS to continue to seek improvements to the relatively low CPU efficiency, and appreciates the work done to understand the background to this.

LHCb Requests for 2019

20

LHCb		2017			2018		2019		
		CRSG recomm.	Pledged	Used	CRSG recomm.	Pledged	Request	2019 req. /2018 CRSG	C-RSG recomm.
CPU	Tier-0	67	67	63	88	88	86	98%	86
	Tier-1	207	199	245	253	250	271	107%	271
	Tier-2	116	147	196	141	164	152	108%	152
	HLT	10	0	208	10	0	10	100%	10
	Total	400	413	712	492	502	519	105%	519
	<i>Others</i>			114					
Disk	Tier-0	10.9	10.9	6.6	11.4	11.4	14.1	124%	14.1
	Tier-1	22.1	20.9	16.9	24.5	26.3	27.9	114%	27.9
	Tier-2	4.7	3.3	3.6	5.7	3.7	6.8	119%	6.8
	Total	37.7	35.1	27.1	41.6	41.4	48.8	117%	48.8
Tape	Tier-0	25.2	25.2	22.5	33.6	33.6	35.0	104%	35.0
	Tier-1	43.3	42.0	31.3	45.6	56.9	50.9	112%	50.9
	Total	68.5	67.19	53.8	79.2	90.5	85.9	108%	85.9

- LHCb requests are minimal
- Increase is driven mainly by re-stripping

- Additional CERN requests for data and analysis preservation Open Data project:
 - +30 kHS06-years CPU
 - +0.5 PB +0.5 PB of disk for the two projects

LHCb Recommendations

21

- C-RSG congratulates LHCb on the very successful management of their computing model and the links to their data processing operations, including usage monitoring and future resource prediction.
- C-RSG considers that LHCb has an appropriate contingency plan in case more data is accumulated in 2018 than planned.
- C-RSG takes note of the needs reported by LHCb for analysis and data preservation as well as for participation in CERN's Open Data project.

Contingency Plan

C-RSG asked the experiments, as part of future resources assessments, to provide a mitigation strategy to address changes in the assumed running conditions at the level of a 20% increase.

Answers:

ALICE:

- will keep data on tape. Data reconstruction, MC, and analysis delayed until sufficient amount of CPU and disk resources are available. The required resources are expected to scale linearly with the amount of data.

ATLAS:

- CPU Tier-0, off-load to Tier-1's via spill-over operations with very low increase of resources.
- Increase of +0.5 FTE to guarantee the operations.
- Storage Tier-0: +8PB of disk and more tapes depending on how much disk will be available.

Contingency Plan – cont'd

C-RSG asked the experiments, as part of future resources assessments, to provide a mitigation strategy to address changes in the assumed running conditions at the level of a 20% increase.

Answers:

CMS:

- CPU Tier-0 use Tier-1s or CERN_T2.
- CPU Tier-1 and Tier-2 delays in analysis and reprocessing.
- Disk space Tier-0 move operation to Tier-1.
- Disk space Tier-1, Tier-2 further reduction of AOD(SIM) copies that will go below 1.
- All of the above will requiring additional operations.

LHCb:

- use the powerful online selection capabilities to park on tape specific datasets and postpone analysis. This has a multiplying effect \Rightarrow simulation will also be postponed.

2020 Outlook

24

- **ALICE:** no CPU and tape increase. Needs for disk space: around 20% at T0 and T1 and 15% at T2. Activities: third pass for pp data and the final two passes for the Pb-Pb data and the associated MC production.
- **ATLAS:** ATLAS has not been able to quantify their needs for 2020 as their computing needs for LS2 still need to be understood.
- **CMS:** full model not in place yet, as it depends on the 2021 LHC machine parameters. At the moment, no increase is foreseen for 2020 while it is expected a substantial increase in 2021.
- **LHCb:** well advanced in estimating resource needs. There will be a need for 15-20% growth in CPU and disk resources, whilst tape growth is expected to be negligible. The resources will be needed to overlap the consolidation of all processed data from Run2 to prepare for Run3

Comments and Recommendations

- ❑ Only partial disk space accounting is available, C-RSG would require to monitor also the Tier-2 disk space usage starting from the next scrutiny.
- ❑ In the shutdown years the purchasing profile can be different with respect to data taking periods. C-RSG would request that the experiments for the next scrutiny provide requests assuming:
 - the purchasing will be delayed until the begin of Run 3;
 - the resources are provided year by year during the shutdown.
- ❑ The C-RSG requests that CERN's expectations for the resource as function of time for Run 3 be documented and communicated to the experiments prior to the start of Run 3.
- ❑ The C-RSG commends the experiments for having started the activity on long term data preservation and their involvement on making data available for education through Open Data Portal. C-RSG recommends that these projects are discussed within WLCG and the experiments in order to have a structured projects with the appropriate funding.

Comments and Recommendations cont'd

- ❑ The C-RSG appreciates the continued work on increasing the computational efficiency and on reducing the CPU and disk resources requirements. The C-RSG hopes that all the experiments reach the same level of involvement in improving the computational efficiency and reducing the data size format.
- ❑ The C-RSG would request that as part of future resources assessments the experiments provide a proposed mitigation strategy to address changes in the assumed running conditions for the experiment at the level of a 20% increase.
- ❑ Experiments are encouraged to gather accounting information on the use of non-WLCG resources in order to be able to report them in future to the C-RSG. Furthermore we welcome the fact that every experiment has made use of their HLT farms to augment their CPU resources.