

# Report of the Computing Resources Scrutiny Group

## CRSG current composition

C.Bozzi (Italy), T.Cass (CERN), **G. Lamanna** (France), D.Espriu (Spain, *Chairman*),  
J.Flynn (UK), M.Gasthuber (Germany), D.Groep (The Netherlands),  
**T. Schalk** (USA), W.Trischuk (Canada), B.Vinter (Nordic Grid),  
**H.Meinhard** (CERN/IT, *Scientific Secretary*)

# Report of the Computing Resources Scrutiny Group

Every year the CRSG should scrutinize

- The resource accounting figures for the preceding year
- The use the experiments made of these resources
- The overall request for resources for every experiment for the following year and forecasts for the subsequent two years

The CRSG shall also examine the match between the refereed requests and the pledges from the Institutions.

The CRSG shall make recommendations concerning apparent under-funding

# Report of the Computing Resources Scrutiny Group

Contents of this report:

- Overall usage of the WLCG resources during 2010 (January to December). [Part A of the written report]
- The use the experiments made of the committed resources [Part B of the written report]
- Scrutiny of the experiments' requests for 2012 and 2013 [Part C of the report]
- Mitigation of the growth in resources [p. 4]
- Recommendations and requests to the experiments [p. 5]

A detailed usage report of the Tier 2 is appended to the written report [p.38]

The scrutiny of the experiments' requests is **preliminary**, subject to revision during the October C-RRB. Some revision/ additional information is required from several collaborations before a firm recommendation can be given..

*In preparation of the October C-RRB we ask the experimental collaborations to provide their documents by **1st SEPTEMBER 2011.***

# Report of the Computing Resources Scrutiny Group

Live time: 30 days/month = **720** hours

Folding in efficiencies  $720 \times 0.7 \times 0.4 = 201.6$  effective hours/month = **725760** s/month

RRB year	RRB year start	RRB year end	Months (max) Data taking	Total live time (in Ms)	Pp	HI
2010	June '10	March '11	8	5.8	5.1 (~3)	0.7
2011	April '11	March '12	8	5.9	5.2	0.7
2012	April '12	March '13	8	5.9	5.2	0.7

During 2010 the luminosity has increased by 5 orders of magnitude and it is expected to go up by yet another order of magnitude in the present run.

During 2010 the computing models have passed the challenge with good marks. Most of the data was obtained during the last weeks of the run, showing the adaptability of the models.

On average the experiments have recorded **30%** less events than expected

# Report of the Computing Resources Scrutiny Group

## PART A

### Scrutiny of the WLCG resources utilization in 2010

# Report of the Computing Resources Scrutiny Group

WLCG accounting report for 2010

<http://lcg.web.cern.ch/LCG/accounting/Tier1/2010>

EGI accounting portal at CESGA.ES

[http://www3.egee.cesga.es/gridsite/accounting/CESGA/tier1\\_view.html](http://www3.egee.cesga.es/gridsite/accounting/CESGA/tier1_view.html) ,

Reports provided by the four experiments to the CRSG. T2 usage compiled by Ian Fisk (with thanks)

## Overall usage 2010

Resource	Site(s)	Used/Available [mean occupancy] (Oct 2010)
CPU	CERN	32 % (22 %)
	T1	62 % (50 %)
	T2	122 % (92 %)
Disk	CERN	110 [75] % (71 %)
	T1	110 [89] % (92 %)
	T2	Not available
Tape	CERN	69 [52] % (52 %)
	T1	60 [52] % (49 %)

# Report of the Computing Resources Scrutiny Group

Efficiency of the utilization of the CPU at Tier 2s for the whole 2010 (left column) compared to the October 2010 C-RRB (right column)

ALICE	73%	74%
ATLAS	85%	84%
CMS	66%	62%
LHCb	88%	88%

In the current state of implementation of the computing models a substantial amount of chaotic/individual analysis is performed at the Tier 2 (where applicable), coexisting with organized tasks.

*In view of these figures we recommend a revision of the official figure of **60%** assumed for Tier2 up to **66%** (2/3). This represents an optimization of 10% in the Tier 2 CPU request.*

*The CRSG recommends that a final decision is taken in the October C-RRB. The tables shown in this presentation do not include this optimization factor yet.*

# Report of the Computing Resources Scrutiny Group

Percentage of use of the resources by experiment in 2010 (CERN+Tier 1s)

Collaboration	% of tape in T1+CERN used at end of period	% of disk in T1+CERN used at end of period	% of CPU in T1+CERN used	% of which at CERN (Oct 2010)
ALICE	8 %	6 %	16 %	33 % (26%)
ATLAS	35 %	57 %	59 %	14 % (10%)
CMS	52 %	29 %	17 %	20 % (22%)
LHCb	5 %	9 %	9 %	46 % (42%)



# Report of the Computing Resources Scrutiny Group

## Percentage of use of the resources by experiment in 2010 (Tier 2s)

Collaboration	% of total disk in T2 used at end of period	% of total CPU in T2 used in 2010 ( <i>October 2010</i> )	
ALICE	N/A	7 %	(7 %)
ATLAS	N/A	59 %	(59 %)
CMS	N/A	30 %	(29 %)
LHCb	N/A	4 %	(4 %)

Statistics show a marked stability and quite definite patterns.

Disk @ Tier 2 not centrally accounted yet.

# Report of the Computing Resources Scrutiny Group

## Delivered versus pledged

Resource	Site(s)	Available / pledged
CPU	CERN	129 %
	T1	109 %
	T2	122 %
Disk	CERN	112 %
	T1	100 %
	T2	Not available
Tape	CERN	100 %
	T1	101 %

### Very satisfactory !

- NL-LHC/T1 delivered only 79% of the pledged disk
- ASGC delivered 71% of the pledged disk and 82% of the pledged tape only..

The large turnout in CPU at the Tier 2 indicates that the percentage installed is actually above 100%, however automated accounting is not fully implemented yet.

# Report of the Computing Resources Scrutiny Group

## PART B

### Usage by the experimental collaborations

# Report of the Computing Resources Scrutiny Group

ALICE

Resource	Site(s)	Pledge	Usage	Used/Pledge %
CPU/kHS06	T0+CAF	46.8	25.8 (6.9, 55.2)	55% (38%)
	T1	45.6	37.5 (9.1, 73.3)	82% (68%)
	T2	52.6	32.3 (5.9, 63.6)	61% (60%)
Disk/TB	T0+CAF	5500	2582	47% (14%)
	T1	6122	1958	32% (22%)
	T2	4326	Not available	--
Tape/TB	T0+CAF	6300	3014	48% (28%)
	T1	8485	1229	14% (5%)

CPU usage over the first two months of 2011 is 68.1 kHS06 for T0+CAF, 52.2 kHS06 for T1 and 49.4 kHS06 for T2.

Storage was below pledges in 2010, but is now growing fast. Disk and tape storage at T0+CAF had both grown to 3.1 PB at the end of February 2011, while disk storage at T1s had grown to 2.3 PB

# Report of the Computing Resources Scrutiny Group

## ATLAS

Resource	Site(s)	Pledged	Used	Used/ Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	67	36	54% (36%)	
	T1	211	171	81% (55%)	81% (83%)
	T2	215	183	85% (81%)	84% (84%)
Disk (TB)	T0+CAF	3900	3900	100% (81%)	-
	T1	22018	22000	100% (82%)	-
	T2	21238	12000	56% (57%)	-
Tape (TB)	T0+CAF	8900	9500	107% (72%)	-
	T1	15372	6800	44% (51%)	-

# Report of the Computing Resources Scrutiny Group

CMS

Resource	Site(s)	Pledged	Used	Used/ Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	96.6	21	22% (17%)	
	T1	103.5	46	44% (32%)	77% (71%)
	T2	196	152	78% (76%)	64% (62%)
Disk (TB)	T0+CAF	4100	4396	107% (60%)	-
	T1	12183	9223	76% (69%)	-
	T2	13627	10000	74% (55%)	-
Tape (TB)	T0+CAF	14600	9728	67% (49%)	-
	T1	23677	17676	75% (65%)	-

# Report of the Computing Resources Scrutiny Group

LHCb

Site	CPU [kHS06]	CPU [kHS06] used	Disk (TB)	Disk (TB) used	Tape (TB)	Tape (TB) used
CERN	23	14.4 (63%)	1290	922 (71%)	1800	844 (47%)
Tier-1	43	14.7 (34%)	3254	2096 (64%)	3036	903 (30%)
Tier-2	42	13.6 (32%)	434	n/a	0	0

# Report of the Computing Resources Scrutiny Group

## PART C

Scrutiny of the requests for 2012 and 2013 (preliminary)



# Report of the Computing Resources Scrutiny Group

The CRSG is generally satisfied with the amount of information provided to the referees to assess the past usage and the request. The CRSG has been granted access to rather detailed spreadsheets or equivalent documentation, and it has been able to check the accuracy of their estimates as well as the adequacy of the requests to the current status of the computing models. However,

- We recommend the ALICE collaboration to submit more detailed reports to the CRSG.
- The CRSG would also like to see a better time granularity in the requests of the ATLAS collaboration.
- The quantitative impact of pile up on the needed resources is still not fully documented.

# Report of the Computing Resources Scrutiny Group

Some experiments expressed interest in taking data at a higher rate than contemplated in the current computing models

*The LHCC endorses the inclusion of charm as a valuable extension to the LHCb physics programme and encourages the collaboration to explore ways to record these extra data. In case the computing resources are restricted then the committee recommends delaying the processing and analysis of these data, or reducing the amount of simulation, as necessary.*

*The LHCC encourages CMS and ATLAS to explore ways to record the data needed to maintain high sensitivity to new physics at low thresholds as the luminosity increases. Such an initiative is supported to fulfil the physics goals of the 2011-12 run. In case the available computing resources are restricted, the committee recommends delaying the processing, analysis or simulation, as necessary.*

The CRSG does not see how a substantial increases of the data taking rate can be accommodated with the existing computing resources. To do so would require further, substantial modifications to the computing models we were presented with.

# Report of the Computing Resources Scrutiny Group

## ATLAS

CPU [kHS06]	2012	2013
CERN	73	<b>73</b>
Tier-1	259	<b>273</b>
Tier-2	295	<b>315</b>
Disk [PB]		
CERN	9	10
Tier-1	27	30
Tier-2	47	53
Tape [PB]		
CERN	18	18
Tier-1	36	40

# Report of the Computing Resources Scrutiny Group

## CMS

<b>CPU [kHS06]</b>	<b>2012</b>	<b>2013</b>
<b>CERN</b>	<b>121</b>	<b>119</b>
<b>Tier-1</b>	<b>145</b>	<b>145</b>
<b>Tier-2</b>	<b>350</b>	<b>350</b>
<b>Disk [PB]</b>		
<b>CERN</b>	<b>5</b>	<b>4</b>
<b>Tier-1</b>	<b>22</b>	<b>27</b>
<b>Tier-2</b>	<b>26</b>	<b>26</b>
<b>Tape [PB]</b>		
<b>CERN (including HI)</b>	<b>23</b>	<b>23</b>
<b>Tier-1</b>	<b>51</b>	<b>59</b>

# Report of the Computing Resources Scrutiny Group

## LHCb

Date	Site	kHS06 peak-power	Disk (PB)	Tape (PB)
2012 period	CERN	34	3.5	6.4
	Tier-1	113	9.5	6.2
	Tier-2	48	0	0
2013 period	CERN	33	4.0	7.7
	Tier-1	110	11.1	8.0
	Tier-2	48	0	0

# Report of the Computing Resources Scrutiny Group

LHCb has substantially increased their request due to the new charm channels and also due to the larger-than-expected pile-up. Securing enough computing resources may be a problem.

The referees express their concerns about the usage of CERN based resources, even if available and idle. The Job scheduling (pilot scheduling controlled by LHCb) should always try to use all of the T1 resources before scheduling starts to CERN resources (especially for the user analysis and MC jobs).

The current model foresees a 25% processing at CERN (down from 46% at present).

LHCb should have a 'plan B' if the supplied resources do not allow the peak power requested for certain times. For this purpose the long-promised commissioning of the on-line farm for physics analysis seems quite relevant. We urge the collaboration to move forward in this direction.

# Report of the Computing Resources Scrutiny Group

## ALICE

CPU/kHS06	2012	2013
CERN T0	85	85
CERN CAF	31	31
Tier 1 Ext	160	157
Tier 2	161	157
Disk/PB		
CERN T0+CAF	14	14
Tier 1	11	9
Tier 2	8	8
Tape/PB		
CERN	20	25
Tier1	21	28

# Report of the Computing Resources Scrutiny Group

The CRSG has confirmed the resource calculations within the stated assumptions made by ALICE

The CRSG acknowledges that computing resources in the amount and proportion requested by ALICE are not available and are unlikely to be offered in the future. The theoretical needs are likely to exceed resources in 2011. Now that reconstruction and analysis of real HI events are under way, the ALICE computing model is being fully exercised and this prediction will be tested.

We ask ALICE to revise their computing model in order that it can fit in the existing resources envelope within a 10% error bar. The ALICE collaboration is requested to submit to the October 2011 C-RRB this revision for final endorsement.

Some suggestions as to how to proceed are included in our written document.



# Report of the Computing Resources Scrutiny Group

## Mitigation of the growth in resources

Experiments were urged to use the experience gained in the first year of running to modify the implementation of the computing models in ways that would make them sustainable in the long run and mitigate the growth in resources.

The emergence of larger-than-planned pile up has meant a real need for optimization to temper the need for additional resources.

- Experiments have made an effort to reduce the raw event size and event processing times.
- Reduction in the number of copies stored in Tier 1 or Tier 2.
- Continuous optimization of the Monte Carlo simulation generators.
- Some experiments are introducing dynamical data placing policies.
- The role of different formats revised.
- Experiments have been active in redistributing tasks among Tier 1 and Tier2.
- The experimental collaborations have implemented aggressive data cleaning policies.

# Report of the Computing Resources Scrutiny Group

## Recommendations and requests

- WLCG resources should be used as much as possible as there is a tendency by some collaborations to place heavier demands on CERN resources or suggest that a larger than planned part of their analysis should be done at CERN. In 2010 Tier 1 resources were still underused.
- On efficiency grounds the CRSG recommends sharing CERN resources when allocations are not fully used, perhaps using low priority queues. This should not be used to increase even more the CERN-based share of the analysis.
- The implications for best-use of resources of the interplay between improvements in network bandwidth and dynamical data placement policies should be evaluated.
- Experiments should carefully quantify the impact of pile-up and incorporate these values uniformly and realistically. We expect a careful evaluation of this impact by the October C-RRB meeting.
- The CRSG recommends revising the assumed Tier 2 efficiency up to  $2/3$  (currently is 60%). This represents a saving of 10% in Tier 2 CPU to funding agencies. A more careful evaluation of the relative efficiencies of chaotic vs organized activities would be convenient. A final decision should be taken in the October C-RRB.
- The WLCG accounting of Tier 2 resources is improving steadily but it is still insufficient.