

Report of the Computing Resources Scrutiny Group

CRSG current composition

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

Report of the Computing Resources Scrutiny Group

Contents of this report:

- Overall usage of the WLCG resources during 2010 (January to September).
- The use the experiments made of the committed resources (summary on p. 7)
- The experiments' reports are appended to our assessment, (p. 20)
- Detailed usage report of the Tier 2 (p. 54)
- CRSG work plan for 2011.
- Composition of the CRSG.

- It has not been possible to review the forecast for 2012 and 2013 as the experiments (with one exception) have not submitted their tentative requests.
- A revision of the April 2010 scrutiny for the ALICE experiment is included (p. 4).
- A full new scrutiny for 2011 has not been deemed necessary.

*In preparation of the respective C-RRBs we ask the experimental collaborations to provide their reports and by **1st MARCH** and **1st SEPTEMBER***

Greater precision concerning the future LHC schedule (predicted running time, especially for 2013) is needed from the CERN management in order to provide reasonable estimates

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
2009	LHC start	May '10	3	2.2	2.2	0
2010	June '10	March '11	8	5.8	5.1	0.7

Year	Start	End	Months data taking	Total live time (in Ms)	pp	HI
2010	<i>April '10</i>	<i>August '10</i>	5	2.3	2.3	0

460000 seconds per month compared to the **725000** seconds theoretically expected (**63%**)

On average the experiments have recorded **40%** less events than expected, but for the first time they have had to deal with very sizeable volumes of real data.

In the last 10 weeks experiments are accumulating data at a rate nearly equal to the expected

Report of the Computing Resources Scrutiny Group

WLCG accounting report for 2010

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

EGEE accounting portal at CESGA.ES

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
CPU	CERN	22 % (26 %)
	T1	50% (69 %)
	T2	92 % (>100 %)
Disk	CERN	71 % (71 %)
	T1	92 % (81 %)
	T2	Not available
Tape	CERN	52 % (41 %)
	T1	49 % (64 %)

Report of the Computing Resources Scrutiny Group

Efficiency of the utilization of the CPU at Tier 2s per experiment in 2010 (left column) compared to 2009 (right column)

ALICE	74%	65%
ATLAS	84%	85%
CMS	62%	66%
LHCb	92%	90%

In the current implementation of the computing models only a fraction of the Tier 2 is used for user & group analysis and this surely reflects in the fact that the percentage of utilization in the Tier2 even in this early stage is well above the theoretical 60%. (**76%** when averaged over the Tier2 - It was 75% in the April 2010 report.)

*In view of these figures we recommend a continuous monitoring and if the tendency to large efficiencies is confirmed we would recommend a revision of the official figure of **60%**.*

Report of the Computing Resources Scrutiny Group

Installed CPU and disk **installed/used** at the Tier 2 is not properly centrally accounted yet.

The CRSG requests the WLCG collaboration to finish the implementation of a complete accounting of the Tier 2 resources.

Tier 1 and Tier 2 disk: A metric that takes into account a disk access pattern would be most appropriate, allowing to identify "hot spots", namely data that are frequently accessed. We are aware of the technical difficulties involved and of the timescale needed to define and implement such a metric.

For the time being, we require that the disk utilization in the next reports from the experiments should be given in terms of the various data types involved and how frequently they were changed/replaced on disk, whenever possible. We acknowledge receipt of a detailed document from ATLAS along these lines.

We encourage the experimental collaborations to adopt an active policy of optimization of disk resources.

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
ALICE	5 %	6 %	16 %	26 %
ATLAS	35 %	57 %	59 %	10 %
CMS	55 %	29 %	17 %	22 %
LHCb	5 %	9 %	9 %	42%

Percentage of use of the resources by experiment in 2009 (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
ALICE	6%	4%	10%	29%
ATLAS	30%	57%	55%	9%
CMS	61%	34%	26%	35%
LHCb	3%	6%	9%	20%

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 (2009)
ALICE	N/A	7 % (7 %)
ATLAS	N/A	59 % (48 %)
CMS	N/A	29 % (35 %)
LHCb	N/A	4 % (10 %)

Report of the Computing Resources Scrutiny Group

ALICE

Resource	Site(s)	Pledge [1]	Alice usage [2]	Alice usage CRSG [3,4,5]
CPU/kHS06	T0+CAF	46.8	15 (30)	17.9 (37.8)
	T1	45.6	23.4 (62.6)	30.8 (73.3)
	T2	52.6	21.0 (54.4)	31.5 (64.9)
Disk/TB	T0+CAF	5500	610 + 520 buffer	781
	T1	6122	986 + buffer	1317
	T2	4326	1090	?
Tape/TB	T0+CAF	6300	1551	1754
	T1	8485	620	418

[1] Pledges from April 2010 RRB.

[2] From Alice Computing Resources Usage in 2010. Values in parentheses are averages over July and August.

[3] CERN and Tier 1 data are from the WLCG Tier 1 Accounting Summary.

Tier 2 CPU is from the Tier 2 spreadsheet supplied to the CRSG by Ian Fisk.

[4] CPU is total normalised wall time in kHS06 days from Jan-Aug, divided by the number of days, 243.

Values in parentheses are usage in Aug divided by 31 days.

[5] Disk and tape values are at the end of the period..

Report of the Computing Resources Scrutiny Group

ATLAS

Resource	Site(s)	Required [1]	Pledged [1]	Used [2]	Used/ Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	67	67	24 (31)	36% (46%)	47% (55%)
	T1	192	211	117 (106)	55% (50%)	83% (80%)
	T2	240	215	174 (167)	81% (78%)	84% (79%)
Disk (TB)	T0+CAF	3900	3900	3145	81%	-
	T1	21900	22018	18007	82%	-
	T2	20900	21238	12000	57%	-
Tape (TB)	T0+CAF	8900	8900	6426	72%	-
	T1	14200	15372	7814	51%	-

Report of the Computing Resources Scrutiny Group

CMS

Resource	Site(s)	Requested [1]	Pledged [1]	Used [2]	Used/ Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	96.6	96.6	16.6	17% (19%)	37% (35%)
	T1	100.5	103.5	33.5 (36.7)	32% (35%)	71% (70%)
	T2	195	196	149	76% (80%)	62% (57%)
Disk (TB)	T0+CAF	4100	4100	2456	60%	-
	T1	13400	12183	8393	69%	-
	T2	9200	13627	7500	55%	-
Tape (TB)	T0+CAF	14600	14600	7133	49%	-
	T1	23300	23677	15439	65%	-

Report of the Computing Resources Scrutiny Group

LHCb

Site	CPU (kHS06) Request	CPU (kHS06) (last 3 months)	Disk (TB) Request	Disk (TB) Used	Tape (TB) Request	Tape (TB) Used
CERN	21	3.2 (4.9)	1220	850	1500	788
Tier-1	41	10.3 (18.2)	2870	2064	2800	835
Tier-2	36	12.7 (19.7)	20	1	0	0

Report of the Computing Resources Scrutiny Group

Delivered versus pledged

Resource	Site(s)	Available / pledged
CPU	CERN	104 %
	T1	100 %
	T2	At least 92 %
Disk	CERN	100 %
	T1	97 %
	T2	N/A
Tape	CERN	100 %
	T1	100 %

Report of the Computing Resources Scrutiny Group

These percentages of fulfilment are quite satisfactory. However there are local anomalies that are sometimes compensated by other centers' over-commitment. For instance:

- CNAF delivered only 73% of the pledged CPU
- CNAF and ASGC delivered less than 75% of the pledged disk
- ASGC delivered 71% of the pledged tape.

The large turnout in CPU at the Tier 2 indicated that the percentage installed is actually superior to 100%, however automated accounting of this is still not in place.

The CRSG wishes to state that the recommendations contained in our scrutinies are to the best of our knowledge rigorous. They correspond to the real needs of the experiments for a given LHC live time in the present stage of their computing model implementation. There is no contingency for late delivery or failure to meet the pledges included in our estimates or for less than 100% availability of these resources.

Report of the Computing Resources Scrutiny Group

Summary of the scrutiny (I)

- The past months of 2010 have witnessed a substantial change in the WLCG usage with respect to 2009 both in qualitative and in quantitative terms. 2010 has brought about sizeable amounts of real data valid for physics analysis. The experiments have made extensive and intensive use of the WLCG resources.
- Generally speaking the experiments' computing models and the WLCG have demonstrated in a remarkably smooth way their capability to record, distribute and analyze the so far moderate, but rapidly increasing, amounts of data delivered to them by the LHC.
- The performance of the WLCG throughout the year has been generally regular and without any noticeable difficulties, with periods where usage has been quite intensive corresponding to simulation and reprocessing campaigns and physics analysis activities jumping to high levels just before the summer conferences.
- The GRID fabric works reasonably well, data distribution and network performance are excellent, much better that could perhaps have been envisaged some time ago. A similar comment applies to the middleware.

Report of the Computing Resources Scrutiny Group

Summary of the scrutiny (II)

- The data placement policy and the detailed computing activities have been different from what was envisaged in the computing models but this is regarded by the CRSG as being reasonable at this moment.
- Resources are exceeding the experiments' needs at this point and the experimental collaborations have had substantial headroom to increase simulation production, making reprocessing passes more often and making more copies in Tier 1 and Tier 2 to increase accessibility.
- Tier 2s large usage is partly due to an increasing demand by individuals or group for non-organized physics analysis, but the overall use is still dominated by Monte Carlo production. Tier 1s are still underused at this stage (50%) but this situation is changing rather quickly.
- If the present trend is to persist it may be the case that Tier 2 resources are stretched in a not too distant future. Redistribution of the organized activities may be advisable.

Report of the Computing Resources Scrutiny Group

Summary of the scrutiny (III)

- Data analysis should gradually move from the more complete data formats (RECO, ESD) to derived formats such as AOD and DPD. Proliferation of data formats should be carefully watched. Given the good performance of the networking, and the possibility of staging copies from tape, unnecessary multiplication of copies can be avoided thus saving on disk.
- The noticeable under usage of CERN CPU merits a careful evaluation. In periods where there is no data-taking the CRSG thinks that sharing of resources may be advisable.

Report of the Computing Resources Scrutiny Group

CRSG work plan for 2011

1.- For 2011 we ask the experimental collaborations to improve the modelling and the amount of information so that the estimates of the CRSG can be assumed to be accurate at the **5%** level. The CRSG acknowledges the receipt of ATLAS documents dealing with this,

2.- It seems natural that some of the assumptions of the computing models should be revised after the first months of data taking. There are several assumptions that have a direct impact on computing requirements and it seems advisable to work on an improvement plan to optimize processing times and data sizes and other parameters so as to limit the expected increase in the resource request in coming years.

We ask the experiments to work on such an optimization and resource-contention plan in the coming months and to submit their plans to us during 2011 so that, after revision by the CRSG, it can be presented to the C-RRB in October 2011. We take note that since the HI runs have not started and analysis of that data will take place during 2011 this item applies only partly to ALICE.

Report of the Computing Resources Scrutiny Group

CRSG membership

Bernd Panzer-Steindel from CERN who replaced Jürgen Knobloch, has asked to step down from the CRSG citing conflict of interests. Tony Cass, also from CERN has been proposed as a new member of the CRSG. The appointment of CRSG members is made by the respective funding agencies, according to the terms of the WLCG MoU contingent on the approval of the C-RRB.

Three new members should join the CRSG in order to comply with the mandate of the WLCG MoU and Funding Agencies have been contacted in the past weeks. New appointments should become effective after this C-RRB.