Report of the Computing Resources Scrutiny Group

2012 has been an exciting year for particle physics. Computing has been a crucial ingredient of the successful LHC run leading to the discovery of the Higgs-like particle announced on July 4th.

The CRSG wishes to express its praise for the outstanding performance of the LHC, of the experiments under review at this RRB and especially of the WLCG.

Contents of this report:

- Overall usage of the WLCG resources during 2012 (Jan to Aug).
  [Introductory part of the report]
- The use the experiments made of the committed resources
  [Part A of the written report]
- Scrutiny of the experiments' requests for 2013 & 2014
  [Part B of the report]
- Preliminary remarks concerning 2015 and beyond
  [Part C of the report]

Report of the Computing Resources Scrutiny Group

Live time: 30 days/month = **720** hours
Folding in efficiencies 720 x 0.7 x 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
2012	**April '12**	**March '13**	**8**	**5.9**	**5.2**	**0.7**
2012	***April '12***	***March '13***	***10***	***7.2***	***6.5***	***0.7***
2013	**April '13**	**March '14**	**-**	**-**	**-**	**-**
2014	**April '14**	**March '15**	**-**	**-**	**-**	**-**

- Excellent accelerator performance
- The time available for physics in 2012 has been 3.8 Ms ~ 100% of the theoretical maximum.
- Experiments have recorded data at increased rates:
  - ALICE: 400 for pp, 560 for pPb /100 H
  - ATLAS: 340 prompt + 150 delayed /200 Hz
  - CMS: 375 prompt + 300 delayed /300 Hz (includes 25% data set overlap)
  - LHCb: 4000 prompt 1000 deferred/ 2000 Hz

Report of the Computing Resources Scrutiny Group

- Pile-up: has averaged to ~20 interactions per crossing (as large as 40 just after a fill).
  At the LHCb IP pile up has been 1.7 (below the 2.5 peak value for 2010 but above the value of 1.5 averaged in 2011).
  Pile-up now seems less of a threat to the efficient running of the experiments than it appeared in previous scrutinies but it has a definite impact on resources.

- During the early months of 2012 the collaborations reconstructed and analysed the events recorded during the second PbPb run.

- A satisfactory trial pPb run took place in September in preparation of the HI run in early 2013. Unlike in 2011 LHCb plans to take data during the pPb period in 2013.

After the excellent performance of the LHC, the limits of the existing resources appear to have been reached and some collaborations have limited further increases in data taking rate that were previously envisaged.

Report of the Computing Resources Scrutiny Group

- There has been a massive use of all the available WLCG resources. The GRID fabric works well; data distribution and network performance are excellent.
- All aspects of the computing models have been put to a test and the challenge has been passed very successfully.
- The pressure on grid resources has been considerable, particularly in the weeks preceding the July 4th announcement but the system has responded adequately.
- The collaborations have implemented more realistic and more organized data distribution policies. The number of reprocessing passes has decreased and processing times optimized.
- We detect a more efficient use of the resources in those collaborations where the computing model tends to favour organized analysis and a more hierarchical structure, but the efficiency is high overall.
- Some experiments provide more information on their disk usage including details on the popularity of the various files and datasets and are implementing dynamical data placement policies, taking advantage of the good connectivity.
- The collaborations have been very active in redistributing tasks among CERN, Tier 1 and Tier2 to optimize the usage of resources.

Report of the Computing Resources Scrutiny Group

Optimization has allowed the experiments not only to cope with the increasing amounts of data generated by the excellent LHC performance and the more complex events due to pile-up, but also allowed them to record at increased rates

There appears to be room for improvement in Montecarlo production.

Experiments are able to simulate many more events than envisaged in their computing models, but the impact of simulation is growing rapidly and may become unsustainable. More manpower effort needs to be dedicated to optimize Montecarlo production.

Some experiments plan to use their HLT farms (or parts thereof) for reprocessing or MonteCarlo production during 2013 and 2014. We encourage this line of development not only during long shutdowns but in all periods when the beam is off.

We welcome the action taken by some experiments to take into account their non-WLCG resources making the scrutiny process more transparent.

Scrutiny of the WLCG resources utilization in 2012

Report of the Computing Resources Scrutiny Group

Overall usage 2012 (Jan - Aug)

Resource	Site(s)	Used/Available [mean occupancy] *(2011)*
CPU	CERN	58 % *(55 %)*
	T1	95 % *(93 %)*
	T2	167 % *(166 %)*
Disk	CERN	99 [102] % *(119 %)*
	T1	129 [131] % *(137 %)*
	T2	Not available
Tape	CERN	86 [84] % *(97 %)*
	T1	72 [63] % *(51 %)*

Report of the Computing Resources Scrutiny Group

Delivered versus pledged

Resource	Site(s)	Available / pledged
CPU	CERN	100 %
	T1	105 %
	T2	108 %
Disk	CERN	85 %
	T1	97 %
	T2	Not available
Tape	CERN	100 %
	T1	87 %

- The large turnout in CPU at the Tier 2 indicates that the percentage installed is actually well above 100%.
- NL-LHC-T1 and KIT are below their pledges in CPU.
- CNAF, FNAL , NL-LHC-T1 and TRIUMF are below 85% of their pledges in disk.

Report of the Computing Resources Scrutiny Group

Percentage of use of the resources by experiment in 2012 (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	11 %	14 %	11 %	48 % (52 %)
ATLAS	38 %	45 %	56 %	15 % (17 %)
CMS	42 %	33 %	26 %	29 % (21 %)
LHCb	9 %	9 %	8 %	24 % (26 %)

- Figures are more or less stable and generally reflect the basic tenets of the respective computing models.
- CMS has increased their share @ CERN after correcting some past difficulties. ATLAS has experienced some slowdown due to the CERN batch system.
- ALICE has reduced their (still large) dependence on CERN resources.

Report of the Computing Resources Scrutiny Group

Efficiency of the utilization of the CPU at Tier 2s in 2012 (left column) compared to 2011 (right column)

ALICE	**60 %**	**54 %**
ATLAS	**88 %**	**88 %**
CMS	**84 %**	**83 %**
LHCb	**95 %**	**93 %**

CPU efficiency is very high with the exception of ALICE whose efficiency of CPU use is still affected by large I/O requirements and a rather heterogeneous user community and it lags the other LHC experiments. CPU consumption per event has stabilized for RAW data and MC processing, but is still growing for analysis trains.

The collaboration is aware of these difficulties and is trying to favour organized activities. The CRSG urges ALICE to take appropriate actions to improve the efficiency.

Report of the Computing Resources Scrutiny Group

Percentage of use of the resources by experiment in 2012 @ Tier 2s

Collaboration	% of total CPU in T2 used in 2012 *(2011)*
ALICE	**7 %** *(11 %)*
ATLAS	**54 %** *(54 %)*
CMS	**36 %** *(33 %)*
LHCb	**3 %** *(3 %)*

Statistics show a marked stability and quite definite patterns.

Disk @ Tier 2 not centrally accounted.

PART A
Usage by the experimental collaborations

Report of the Computing Resources Scrutiny Group

ALICE

Resource	Site(s)	2012 request	2012 pledge	2012 usage	Efficiency
CPU/kHS06	T0+CAF	125	90	66	58 %
	T1	95	95	73	56 %
	T2	207	115(194)	136	60 %
Disk/PB	T0+CAF	7.8	8.1	8.4	
	T1	7.0	7.2	6.8	
	T2	12.4	9.1(12.9)	9.6	
Tape/PB	T0+CAF	17.1	20.0	9.7	
	T1	11.3	11.5	3.9	

In parenthesis: non-WLCG resources

Report of the Computing Resources Scrutiny Group

ATLAS

Resource	Site(s)	Pledged	Used	Used/ Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	111	111	100 %	89 %
	T1	285	436	153 %	91 %
	T2	328	612	187 %	87 %
Disk (PB)	T0+CAF	9	6	67 %	
	T1	30	27	90 %	
	T2	45	30	67 %	
Tape (PB)	T0+CAF	18	20	111 %	
	T1	38	24	63 %	

Report of the Computing Resources Scrutiny Group

CMS

Resource	Site(s)	Pledged	Used	Used/ Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	121	75	62%	86%
	T1	137	138	101%	89%
	T2	320	368	115%	84%
Disk (PB)	T0+CAF	7	5.9	84%	
	T1	21	20.2	96%	
	T2	27	24	89%	
Tape (PB)	T0+CAF	23	18.9	82%	
	T1	47	36.4	77%	

Report of the Computing Resources Scrutiny Group

LHCb

	Pledges			Usage		
2012	**CPU (kHS06)**	**Disk (PB)**	**Tape (PB)**	**CPU (kHS06)** **Efficiency (%)**	**Disk (PB)**	**Tape (PB)**
Tier 0	34	3.5	6.4	12.0 92 %	2.46	3.9
Tier 1	91	7.3	5.5	43.2 90 %	6.21	5.53
Tier 2	48			27.9 92 %		

PART B
Scrutiny of the requests for 2013 and 2014 (preliminary)

Report of the Computing Resources Scrutiny Group

Recommendations

- We recommend intensive use of the HLT farms during 2013 and 2014 for reprocessing and simulated data production. The possibility of using the HLT farms for MC production or reprocessing during periods with no beam should be thoroughly investigated.

- We recommend an aggressive implementation of dynamical data placement policies and a close scrutiny of disk usage.
  The collaborations should provide data access statistics to demonstrate that the data placement policies are meaningful and effective. The information provided to date not yet satisfactory.
  We ask the WLCG to compile statistics from the Tier1 and principal Tier 2 stating what fraction of the disk volume is seldom or never accessed, detailed by VO, and establish suitable metrics

- The CRSG is considering a revision of the disk efficiency for subsequent scrutinies. Some collaborations have suggested assuming a 100% efficiency if buffering for data is included in the disk budget. This concept needs precise quantification.

- We ask the WLCG to automate the retrieval of the information that the CRSG needs.

Report of the Computing Resources Scrutiny Group

Recommendations for 2013 (continued)

- We anticipate the need for substantial progress in the implementation of fast MonteCarlo simulations to reduce significantly the relative weight of simulation in the computing budget.
  Experimental collaborations should be prepared to dedicate enough manpower to this task and report progress to the CRSG in subsequent scrutinies.

- We underline the importance of smoothing out CPU needs throughout the year and consider the possibility of using external resources for very localized demands. The desire to present new results in winter and summer conferences has a substantial cost in terms of CPU peak demand.

- The CRSG encourages close collaboration of the different centres with the experiments to continue the implementation of intelligent storage management policies to allow efficient and cost-effective access to data. We consider this issue very relevant for the operation of the LHC experiments after 2014.

- We recommend that hardware upgrades at the T1 and T2 are in synchrony with the newly upgraded software to take full advantage of the technology available.

Report of the Computing Resources Scrutiny Group

ALICE

Resource	Site(s)	2013 ALICE	2013 CRSG	2014 ALICE	2014 CRSG
CPU/kHS06	T0+CAF	*126*	**126**	*135*	**135**
	T1	*160*	*(95)* **120**	*160*	**130**
	T2	*145*	*(195)* **145**	*211*	**200**
Disk/PB	T0+CAF	*11.0*	*(13)* **11.0**	*11.0*	**11.0**
	T1	*10.8*	**10.8**	*10.8*	**10.8**
	T2	*15.8*	*(19)* **15.8**	*15.8*	**15.8**
Tape/PB	T0+CAF	*22.8*	*(24)* **22.8**	*26.1*	**26.1**
	T1	*21.0*	*(19)* **21.0**	*23.9*	**23.9**

Report of the Computing Resources Scrutiny Group

ALICE

ALICE has historically faced a lack of computing resources. Partly because of insistence by the CRSG, ALICE lowered its requests in order to better reflect the anticipated resources.

However, in 2012 the reduced request led to an unanticipated corresponding reduction in the pledge at GridKa. This situation may arise at further T1 sites for the 2013 requests and has led the collaboration, in particular, to raise the T1 CPU request in response.

The CRSG deplores a situation where the collaboration derives a request from the computing model, reduces it to accommodate expected overall resource availability, but subsequently feels forced to modify it in anticipation of the way pledges are calculated.

We hope this can be addressed satisfactorily in future and make a plea for the funding agencies to revise their contributions to ALICE resources so that 100% of the scrutinized needs can be met.

For now we propose a change of the T1 CPU request which maintains the sum of T1 and T2 CPU power at the level ALICE would be requesting without the pledge calculations taken into account.

Report of the Computing Resources Scrutiny Group

ATLAS

Resource	Site(s)	2013 ATLAS	2013 CRSG	2014 ATLAS	2014 CRSG
CPU/kHS06	T0+CAF	*111*	**111**	*111*	**111**
	T1	*319*	*(297)* **319**	*373*	**355**
	T2	*355*	*(319)* **350**	*408*	**350**
Disk/PB	T0+CAF	*11*	*(10)* **11**	*11*	**11**
	T1	*35*	*(29)* **33**	*36*	**33**
	T2	*53*	**49**	*56*	**49**
Tape/PB	T0+CAF	*27*	*(19)* **23**	*31*	**23**
	T1	*43*	*(34)* **40**	*53*	**44**

Report of the Computing Resources Scrutiny Group

CMS

Resource	Site(s)	2013 CMS	2013 CRSG	2014 CMS	2014 CRSG
CPU/kHS06	T0+CAF	*121*	**121**	*121*	**121**
	T1	*175*	*(145)* **165**	*175*	**175**
	T2	*350*	**350**	*350*	**350**
Disk/PB	T0+CAF	*7*	**7**	*7*	**7**
	T1	*26*	**26**	*26*	**26**
	T2	*28*	**26**	*29*	**27**
Tape/PB	T0+CAF	*26*	*(23)* **26**	*26*	**26**
	T1	*50*	*(45)* **50**	*60*	**55**

Report of the Computing Resources Scrutiny Group

ATLAS

The CRSG thinks that the request in CPU at T1's in 2013 is justified based on the extended running period. The amount of CPU recommended in 2014 is based on a 20% increase with respect to our recommendation for 2013 in April 2012. In the CPU at Tier 2 centers we recommend no increase in 2014 over 2013. We note that ATLAS appears to have no problem in obtaining ample resources in Tier 2's.

Regarding disk, we can only recommend a small increase in 2014 with respect to the April 2012 approved request for 2013 . We assume an improved usage of disk by aggressively implementing monitoring techniques based on data access and popularity. This recommendation could be revised in spring 2013 based on new evidence. We recommend an early installation of the 2014 resources.

CMS

An increase in CPU is based on the extended running period this year. The reduction in CPU at T1's with respect to CMS requests is justified by assuming a partial use of the HTL farm. The recommendation for disk at T2's is based on assuming an improved usage of disk using monitoring techniques based on data access and popularity. We recommend an early installation of the 2014 resources.

Report of the Computing Resources Scrutiny Group

LHCb

Resource	Site(s)	2013 Needed	2014 Needed
CPU (kHS06)	T0 (CERN)	**34**	**34**
	T1	**110** *(91)*	**110**
	T2 + others	**46**	**46**
	HLT farm	**20**	**20**
Disk (PB)	T0	**4.4** *(3.5)*	**5.5**
	T1	**8.6** *(7.6)*	**10.4**
Tape (PB)	T0	**6.5** *(6.2)*	**7.3**
	T1	**10.8** *(6.1)*	**11.9**

We recommend early deployment of the 2014 resources.

PART C
Preliminary remarks concerning 2015 and beyond

Report of the Computing Resources Scrutiny Group

For this scrutiny, the experiments have made available to us and the LHCC first estimates of the potential requirements for the period 2015-2017 when data taking resumes in full at a planned energy of 13 TeV.

Both ATLAS and CMS think that they are capable of increasing their trigger rate to 1 kHz (the TDR value was 200Hz and 300 Hz, respectively), and LHCb has the potential to double its current rate to 10 kHz. Only ALICE would expect a modest increase of data volumes by about 20%. This would translate into a significant step up in CPU, disk and tape requirements that the CRSG has not quantified with precision at this point.

Up to now, it has been possible to accommodate the increasing computing demands within a roughly flat budget thanks to Moore's law and to continuous improvements: the speed of the reconstruction code has been improved by a factor of 8 since 2010, and memory reduced by about 40%. Although still lagging behind in efficient use, disk is now much better used.

It is doubtful whether the required rise in 2015 or 2016 will be possible within a roughly flat budget. This growth may require significant annual investments, difficult to get given the ongoing financial constraints.

Report of the Computing Resources Scrutiny Group

- The experiments have to invest heavily in software improvements to be able to accommodate the needs within reasonable resources. Monte Carlo simulation is a major resource consumer. Experiments are urged to invest heavily in making large scale simulations sustainable by e.g. devoting the necessary manpower to optimize the processing and make ample use of fast simulations. Likewise, in reprocessing times there is surely still room for improvement too.
- Of particular importance is to adapt the existing software and analysis tools to the new generation of CPU hardware. This may bring in large gains that will not come for free in any case as it will require adapting all existing processes. This will only pay off if this new hardware is made readily available at all of the tier centers
- Disk volumes are harder to improve - event size and further reduction in the number of effective copies thanks to improved network access and data caching can help.
- A substantial investment in software manpower will be needed to achieve all the previous objectives.
- We encourage the experiments to make their plans for the 2015-2017 period in close consultation with the LHCC and the CRSG.

Report of the Computing Resources Scrutiny Group

SUMMARY

2012 has been an excellent year for the LHC. Computing has been an essential ingredient to the LHC success.

The CRSG acknowledges the excellent work done by the computing teams of the LHC experiments.

The CRSG congratulates the Tier1 and Tier2 and all institutions participating in the WLCG for the overall success of the LHC computing.

The collaborations have to try to make the most of the large resources continuously invested by the participant institutes and agencies.

We recommend the funding agencies to endorse our recommendations and continue providing support to the LHC computing.

Report of the Computing Resources Scrutiny Group

CRSG current composition

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

THANK YOU