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*)

Report of the Computing Resources Scrutiny Group

Contents of this report:

- Overall usage of the WLCG resources during 2011.
  [Introductory part of the report]
- The use the experiments made of the committed resources
  [Part A of the written report]
- Preliminary scrutiny of the experiments' requests for 2013 & recommendations
  [Part B of the report]
- Usage of the Tier 2 by country
  [Part C of the report]

*In preparation of the October 2012 C-RRB we ask the experimental collaborations to provide their documents by **1st SEPTEMBER 2012.***

At any moment the CRSG is prepared to discuss with the experimental collaborations specific issues or recommendations, as deemed necessary .

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	PbPb
2011	**April '11**	**March '12**	**8**	**5.9**	**5.2**	**0.7**
2012	**April '12**	**March '13**	**8**	**5.9**	**5.2**	**0.7**

- The time avalilable for physics was 4.7Ms = 90% of the theoretical maximum. Experiments recorded ~1.5 B events (~10B in LHCb)
- This large number of recorded events has been possible thanks to experiments using all the available bandwidth and effectively recording events at rates larger than the nominal ones.
- Pile up started lower than expected but it reached 12-16 on average per crossing after summer. This represented longer processing times and larger data sets.
- The PbPb run at the end of the year was equally successful. During the early months of 2011 the collaborations reconstructed and analyzed the events recorded during the first PbPb run.

Report of the Computing Resources Scrutiny Group

- Real/nominal rates in 2011 and expectations for 2012:

ALICE: 380/100 Hz (200 Hz for PbPb, expected in 2012: 400Hz, for pPb 560 Hz)
ATLAS: 340/200 Hz (expected in 2012: 400 Hz)
CMS: 375/300 Hz (includes 25% overlap, expected in 2012: 400 Hz and up to 600Hz)
LHCb: 3000/2000 Hz (expected in 2012: 4500 Hz)

- Running time during 2012 expected to be very similar to 2011. No running in 2013

Pile-up is expected to increase up to 25-30 events per crossing.

Experiments plan to reprocess all data since 2010 in 2013 as well as analyze low priority streams (`parked data')

The collaborations have been forced to revise some of their assumptions regarding data placement policies, number of copies, etc. They have optimized their reprocessing times, reduced raw data sizes and moved towards derived formats for analysis.

Optimization has allowed 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 to record at the increased rates.

Report of the Computing Resources Scrutiny Group

- The experiments' computing models and the WLCG have demonstrated in a remarkably smooth way their capability to record, distribute and analyse the substantial amounts of data delivered to them by the very successful run of the LHC during last year.
- There has been a massive use of the available resources. Some aspects of the computing models such as large individual non-organized computing usage, format and distribution of the data sets, the flexibility to cope with increasingly challenging running conditions, and the urgency to reprocess and analyse large amounts of data in a short time have represented a real challenge for the computing models and for the WLCG as a whole. This challenge has been passed very successfully.
- Efficiency has been high generally.
- CPU resources are generally exceeding the experiments' needs at this point (by a factor that has clearly decreased with respect to previous years) and the experimental collaborations have had substantial headroom that they have employed to increase simulation production.
- Although less visible, there is still some headroom for disk due to installed resources surpassing pledges in some instances. Computing has never been a limiting factor in any case.

Report of the Computing Resources Scrutiny Group

- The collaborations have implemented more realistic and more organized data distribution policies. Boundaries between tiers have disapperaed to some extent (not necessarily meaning less structured analysis but often more)
- The reprocessing policy is quickly converging to the one indicated in the computing models as the number of events disfavours frequent reprocessing.
- The GRID fabric works well, data distribution and network performance are excellent.
- No significant problems concerning middleware have been reported to the CRSG.
- Accounting information remains imperfect (it does not show much improvement with respect previous reports)

Report of the Computing Resources Scrutiny Group

Scrutiny of the WLCG resources utilization in 2011

Report of the Computing Resources Scrutiny Group

WLCG resources and accounting for 2011:

http://lcg.web.cern.ch/

EGEE accounting portal at CESGA.ES:

http://www3.egee.cesga.es/

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

Overall usage 2011 (Jan-July)

Resource	Site(s)	Used/Available [mean occupancy] *(October 2011)*
CPU	**CERN**	**55 %** *(52 %)*
	T1	**93 %** *(83 %)*
	T2	**166 %** *(117 %)*
Disk	**CERN**	**119 [105] %** *(99 %)*
	T1	**137 [121] %** *(116 %)*
	T2	**Not available**
Tape	**CERN**	**97 [75] %** *(75 %)*
	T1	**51 [47] %** *(43 %)*

Report of the Computing Resources Scrutiny Group

Delivered versus pledged

Resource	Site(s)	Available / pledged
CPU	CERN	100 %
	T1	99 %
	T2	136 % [was 117%]
Disk	CERN	100 %
	T1	109 % [was 101%]
	T2	Not available
Tape	CERN	100 %
	T1	89 %

- The large turnout in CPU at the Tier 2 indicates that the percentage installed is actually above 100%, and that the efficiency is large.

- Compare 109% installed disk@T1 with the 121% / 137% usage : more disk than pledged is available and efficiency is higher than theoretical 70%. Some sites have large excess of disk.

- NL-LHC-T1 and NDGF are below their pledges in disk and CPU.

- Tier 1 have adapted to the low usage of tape.

Report of the Computing Resources Scrutiny Group

Percentage of use of the resources by experiment in 2011 (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	12 %	14 %	15 %	52 % (59 %)
ATLAS	39 %	46 %	51 %	17 % (18 %)
CMS	41 %	33 %	23 %	21 % (18 %)
LHCb	8 %	7 %	11 %	26 % (28 %)

- Figures are stable and generally reflect now the basic tenets of the respective computing models.
- The large collaborations seem to converge in the relative fraction of CERN resources they use.
- LHCb does now a reasonable fraction of their total computing at CERN. 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 2011 (left column) compared to 2010(right column)

ALICE	**54 %**	**50 %**
ATLAS	**88 %**	**89 %**
CMS	**83 %**	**80 %**
LHCb	**93 %**	**90 %**

70% *nominal efficiency for CPU@T2 in 2013 adopted.*

CPU efficiency is very high with the exception of ALICE.
While the overall CPU efficiency of ALICE at Tier 2 is low but still acceptable, the one associated to chaotic analysis drops to a worrisome 16% with a huge dispersion among users. Due to the implementation of the ALICE computing model the average efficiency for Tier 1 and Tier 2 is very similar.

LHCb uses T2 mostly for MC production – far more efficient than user analysis.

Report of the Computing Resources Scrutiny Group

Percentage of use of the resources by experiment in 2011 (Tier 2s)

Collaboration	% of total CPU in T2 used in 2010 *(October 2010)*
ALICE	**11 %** *(7 %)*
ATLAS	**54 %** *(59 %)*
CMS	**33 %** *(30 %)*
LHCb	**3 %** *(4 %)*

Statistics show a marked stability and quite definite patterns.

Disk @ Tier 2 not centrally accounted yet.

Report of the Computing Resources Scrutiny Group

PART A

Usage by the experimental collaborations

Report of the Computing Resources Scrutiny Group

ALICE

Resource	Site(s)	2011 request	2011 pledge	2011 usage	Efficiency
CPU/kHS06	T0+CAF	**62**	**62**	**56**	**67%**
	T1	**117**	**71**	**60**	**59%**
	T2	**121**	**81**	**107**	**54%**
Disk/PB	T0+CAF	**6.1**	**6.1**	**5.0**	**--**
	T1	**7.9**	**5.5**	**6**	**--**
	T2	**6.6**	**7.3**	**9.9**	**--**
Tape/PB	T0+CAF	**6.8**	**6.8**	**7.9**	**--**
	T1	**13.0**	**8.0**	**3.5**	**--**

Report of the Computing Resources Scrutiny Group

ATLAS

Resource	Site(s)	Pledged	Used	Used/ Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	**75**	**82**	**109 %**	**90 %**
	T1	**248**	**244**	**99 %**	**87 %**
	T2	**285**	**405**	**142 %**	**88 %**
Disk (PB)	T0+CAF	**7**	**5**	**70 %**	**-**
	T1	**26**	**27**	**103 %**	**-**
	T2	**35**	**22**	**62 %**	**-**
Tape (PB)	T0+CAF	**12**	**14**	**117 %**	**-**
	T1	**32**	**16**	**50 %**	**-**

Report of the Computing Resources Scrutiny Group

CMS

Resource	Site(s)	Pledged	Used	Used/ Pledged	Average CPU efficiency
CPU (kHS06)	T0+CAF	**106**	**39**	**37%**	**59%**
	T1	**130**	**114**	**88%**	**85%**
	T2	**305**	**265**	**87%**	**80%**
Disk (PB)	T0+CAF	**5.4**	**3.7**	**68%**	**-**
	T1	**16.2**	**15.5**	**97%**	**-**
	T2	**18.1**	**12.7**	**70%**	**-**
Tape (PB)	T0+CAF	**21.6**	**14**	**65%**	**-**
	T1	**45**	**30**	**67%**	**-**

Report of the Computing Resources Scrutiny Group

LHCb

Resource	Site(s)	Pledged	Used	Used/ Pledged
CPU (kHS06)	T0 (CERN)	**21.0**	**7.2**	**34 %**
	T1	**69.2**	**40.0**	**58 %**
	T2 + others	**40.5**	**47.0 + 26.2**	**180 %**
Disk (PB)	T0	**1.5**	**1.2**	**80 %**
	T1	**3.7**	**2.7**	**73 %**
	T2	**--**	**--**	**--**
Tape (PB)	T0+CAF	**2.5**	**2.1**	**84 %**
	T1	**3.9**	**3.3**	**85 %**

Report of the Computing Resources Scrutiny Group

Steps to mitigate the growth in resources

- Experiments have made an effort to reduce the raw event size (and the size of all subsequent derived formats) . These efforts have mitigated the serious challenge that pile-up represents and allowed experiments to record events at a higher rate, indicating some margin of safety and redundancy in the resources.
- Experiments have set up task forces to reduce processing times and they have generally improved, partly under the pressure to deal with increasing values of pile-up.
- The collaborations have made substantial changes in their data distribution policies, reducing the number of copies stored in Tier 1 or Tier 2 and moved to more compact datasets for analysis. They have been very active in redistributing tasks among CERN, Tier 1 and Tier2.
- The collaborations have continued implementing aggressive data cleaning policies
- Substantial progress in the implementation of fast MonteCarlo simulations has been made in some experiments.
- The user efficiency is better than planned and continuously improving.
- Some experiments plan to use their HLT farms (or parts thereof) for reprocessing or MonteCarlo production during 2013.
- The collaborations attempt to smoothen out their peak CPU demands throughout the year.

Report of the Computing Resources Scrutiny Group

PART B

Scrutiny of the requests for 2013 (preliminary)

Report of the Computing Resources Scrutiny Group

Recommendations for 2013

- We recommend the use of the on-line farms during 2013 for reprocessing and simulated data production. This would entail “parking” some fraction of the data for later processing. The T0 resources should also be available in 2013. These two resources should permit an approximately flat profile for CPU requests in 2013 compared to 2012.

- We recommend to smoothen out the CPU needs throughout the year and consider the possibility of using external resources for very localized demands, particularly for MonteCarlo production.

- We remind the collaborations that, while they are welcome to write data at increased rates they cannot expect that resources automatically increase to match these rates. Therefore they should be selective in the kind of `dark’ or `parked’ data they plan to collect. Maintaining a reasonably flat profile is essential for the sustainability of the WLCG

- The CRSG would like to keep a balanced usage of the different Tiers. Ensuring such a balance will maintain a healthy WLCG collaboration and so ensure long term success.

Report of the Computing Resources Scrutiny Group

Recommendations for 2013 (continued)

- We recommend keeping the request for new disk under close scrutiny. Some collaborations have enlarged their physics scope and this may justify some increases but others have not fully justified the usage of existing disk resources yet.
  If possible, the collaborations should present data access statistics to better understand and demonstrate that the data placement policies are meaningful and effective.

- 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. In particular the implications on network bandwidth for best-use of resources should be considered. We consider this issue very relevant for the operation of the LHC experiments after 2014.

- We encourage the experimental collaborations to continue working on realistic estimates for the computing needs in 2014 and beyond, keeping the budgetary constraints in mind and working with the CRSG as necessary.

Report of the Computing Resources Scrutiny Group

Comments on the scrutiny for the 2013 requests

ALICE

- Presents a request in line with the expected resources and describes in detail possible new contributors in the short term.
- Some unpledged resources are listed and acounted for, helping to bridge the gap with requests.
- Stays within the `natural envelope' of resources.
- Low CPU efficiency is the major concern. We ask the collaboration to use the 2013+ period to reformulate some of the computing strategies aiming to reach efficiencies comparable to the other experiments as much as possible.
- Some unknowns concerning the pPb run at the end of 2012.

Report of the Computing Resources Scrutiny Group

ALICE

Resource	Site(s)	2013
CPU/kHS06	T0+CAF	***125***
	T1	***95***
	T2	***195***
Disk/PB	T0+CAF	***13.4***
	T1	***10.9***
	T2	***19.4***
Tape/PB	T0+CAF	***23.5***
	T1	***19.1***

Report of the Computing Resources Scrutiny Group

ATLAS

- Plans to record data @ 400 Hz and `park' the less relevant part for later analysis.
- Plans to make intensive use of the DAQ farm and T0 resources.
- Makes an intensive use of all resources available. They were able to make much more simulation than originally envisaged and can use MC production as a lever.
- Submitted a `revised' 2012 estimate with increased requests. Requests for 2013 are even larger and appear to the CRSG unrealistic in view of the existing spending profile and the availability of free resources in 2013+.
- 2012 will be very much similar to 2011 as data taking is concerned, except that pile-up will increase.
- Taking into account the LHCC recommendations and having the previous considerations in mind we conclude that the committed resources should match the **revised** 2012 ones.
- This is a tentative scrutiny; the final one will be provided in the October 2012 C-RRB where the present estimates can be revised if deemed necessary.

Report of the Computing Resources Scrutiny Group

ATLAS

CPU [kHS06]	2013 (this scrutiny)	2013 (previous estimate)
CERN	***111***	*111*
Tier-1	***297***	*273*
Tier-2	***319***	*281*
Disk [PB]		
CERN	***10***	*10*
Tier-1	***29***	*30*
Tier-2	***49***	*53*
Tape [PB]		
CERN	***19***	*18*
Tier-1	***34***	*33*

Report of the Computing Resources Scrutiny Group

CMS

- Plans to record up to 600 Hz, 400 Hz on average, and `park' the less relevant data for later analysis.

- The use of HTL farm unclear. Its use is strongly encouraged by the CRSG.

- Makes an intensive use of all resources available. They were also able to make much more simulation than originally envisaged. During 2011 experienced problems with the memory footprint that reduced their CERN usage, hopefully partly solved.

- Also submitted a `revised' 2012 estimate with increased requests. Requests for 2013 are even larger, particularly on CPU @T1 and T2. The CRSG cannot endorse this large request.

- 2012 will be very much similar to 2011 as data taking is concerned, except that pile-up will increase. Taking into account the LHCC recommendations and having the previous considerations in mind we also concluded in this case that the committed resources should match the revised 2012 ones. Some additional disk @T1 appears justified.

- This is a tentative scrutiny; the final one will be provided in the October 2012 C-RRB where the present estimates can be revised if deemed necessary.

-

Report of the Computing Resources Scrutiny Group

CMS

CPU [kHS06]	2013 (this scrutiny)	2013 (previous estimate)
CERN	***121***	*120*
Tier-1	***145***	*145*
Tier-2	***350***	*306*
Disk [PB]		
CERN	***7***	*7*
Tier-1	***26***	*27*
Tier-2	***26***	*26*
Tape [PB]		
CERN (including HI)	***23***	*23*
Tier-1	***45***	*59*

Report of the Computing Resources Scrutiny Group

LHCb

- Plans to record data @ 4500 Hz, justified on basis of the revised charm physics program
- Plans to make intensive use of the on-line farm.
- LHCb computing is very mature but a clear underuse of their CERN usage has been observed which had to be compensated by redistribution of tasks, particularle in the T2. The model has shown good flexibility in adapting to tighter resources.
- While the total computing power is OK, the CRSG is of the opinion that some rethinking of the model may be necessary.
- The 2013 request is flat with respect to previous requests.
- A substantial amount of unpledged resources will help LHCb to fulfill their new ambitious physics program.

Report of the Computing Resources Scrutiny Group

LHCb

Site	kHS06	Disk (PB)	Tape (PB)
CERN	***21***	***3.5***	***6.2***
Tier-1	***55***	***7.6***	***6.1***
Tier-2	***47***	***0***	***0***
Unpledged	***(54)***	***--***	***--***
Total	***123 (177)***	***11.1***	***12.3***

Report of the Computing Resources Scrutiny Group

Composition of the CRSG

Concezio Bozzi, representing INFN, has been replaced by Donatella Lucchesi. A replacement for William Trischuk (Canada) is now pending. During 2012 it will be necessary to renew or replace those members of the CRSG (including the chairman) that were not replaced during 2010 and 2011.

Report of the Computing Resources Scrutiny Group

2011 has been an excellent year for the LHC and 2012 will be even better. Computing has been an essential ingredient to the LHC success.

Computing is now very mature. The collaborations have to try to make the most of the large resources continuously invested by the participant institutes and agencies.

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.

We recommend the funding agencies to endorse this review and continue providing support to the LHC computing.

Report of the Computing Resources Scrutiny Group

THANK YOU