Report of the Computing Resources Scrutiny Group 2008

The purpose of the CRSG is to inform the decisions of the Computing Resources Review Board (C-RRB) for the LHC experiments.

According to the WLCG, MoU the starting point of the scrutiny is the request information presented at the spring C-RRB meeting and any guidance that the C-RRB cares to give. From that moment the CRSG enters into a sustained dialogue with each experiment seeking to understand to what extent the computing resource requests are well motivated. The recommendations are presented in the C-RRB autumn meeting.

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

This is the first time that such an independent and detailed scrutiny of the computing yearly request is carried out.

Report of the Computing Resources Scrutiny Group 2008

In carrying out the present scrutiny the scope of this group is largely limited to the implementation of the respective computing models whose TDRs date back to 2005. The review of the computing models themselves has been in the hands of the LHCC. However:

- A gray zone where the respective competences of the LHCC and the CRSG overlap exists.
- The natural evolution of the commissioning of the experiments as well as the implementation of the computing models in successive tests and a better understanding have motivated a number of changes, sometimes representing limitations in the original model or assumptions.

When the CRSG feels it is not competent to judge the validity or convenience of changes, particularly on the physics side, we bring them to the attention of the LHCC.

While the methodology used has varied, uniformity in the scrutiny has been a priority.

For a proper scrutiny we have yet to see real collisions and real data with the computing models going through a reality check. The CRSG prefers not to commit itself to any specific forecast for 2010 and beyond. Only specific recommendations for 2008 and 2009 are provided.

With a few exceptions, no gross discrepancies have been found so the existing envelope could be used as a guidance. Yet, some of the discrepancies found between the scrutinized needs and the 'historic' request may be worrisome.

Report of the Computing Resources Scrutiny Group 2008

The CRSG proposed a standard set of assumptions on beam time. These assumptions have been used for scrutinizing all experiments. They differ considerably from previous scenarios in the case of 2008, and only slightly in 2009. An ‘efficiency’ of 50% has been assumed in order to extract useful beam time from the total amount that the accelerator will be running. This is an optimistic assumption (recent public presentations suggest that 40% is closer to reality and this is perhaps still too optimistic for the first months of running).

These beam times would correspond to 3 months of data-taking in 2008 and 7 months of data-taking in 2009 for proton-proton (pp) operations, and 0 months in 2008 and 1 month in 2009 for heavy ion (AA) operations. These were rather optimistic, but attainable, expectations.

However the 19 September events forced a last-minute change of the scrutiny. The new ones assume that 2009 will be a nearly normal year as far as running conditions.

Year	pp	AA
	Beam time (seconds/year)	Beam time (seconds/year)
2008	0.3×10^7	0
2009	0.9×10^7	10^6
2010	10^7	10^6

Year	pp	AA
	Beam time (seconds/year)	Beam time (seconds/year)
2008	0	0
2009	0.9×10^7	10^6
2010	10^7	10^6

Report of the Computing Resources Scrutiny Group 2008

ALICE

The ALICE computing model is well tested for MC simulation and event reconstruction. Chaotic (end user) analysis requirements for real data are the hardest to estimate in advance.

One of the elements of our assessment is an independent and simplified CM produced by the CRSG.

The requests for 2009 look reasonable overall. The CRSG thinks the disk requirement at T1s and particularly at T2s is underestimate, but slightly overestimate at CERN.

The mass storage request looks overestimated, at least partly because the ALICE model is accumulating data from an assumed 2007 start-up.

MC simulation is very demanding. ALICE addresses this by generating underlying AA events which are merged several times with a signal. ALICE aims for a 1:1 ratio of real to MC events: reducing MC production can produce savings but risks compromising physics.

While the total CPU request agrees with the CRSG scrutiny, the distribution differs from the request.

The ALICE model distributes fractions of raw and reconstructed data to T1 and T2 disk storage, allowing duplication of data in high demand. Experience with early running should allow the assumed fractions to be checked and perhaps revised.

It is clear that ALICE's computing requirements are unlikely to be met in practice. The CRSG recommends that ALICE makes a clear statement to the LHCC how their physics programme will be affected.

Report of the Computing Resources Scrutiny Group 2008

ALICE

Resource	**CERN**	**T1 ext**	**T2 ext**	**Total**	**change**	**Sep 07 request**	**change**
CPU/MSI2k	10	12	21	**43**	0	43.8	-2%
Disk/PB	2.5	9.9	9.6*	**22.1**	0	15.3	44%
MS/PB	7.0	10.1	0	**17.1**	-7%	19.7	-13%

Report of the Computing Resources Scrutiny Group 2008

ATLAS

Recent experience has given ATLAS a good understanding of the strong and weak points of their computing model. The CRSG believes that the ATLAS CM will successfully pass the test of first contact with LHC data.

The ATLAS CM emphasizes accessibility of the data.

The ATLAS computing model presented in the TDR was optimistic with respect to event sizes, event data formats, the distribution model and the required resource capacity. There were no uncertainties assigned to the input parameters. At this time some elements of the model should probably be reassessed.

Figures for the requested storage are consistent with the combination of event size, expected rate and number of copies. However event sizes have grown and the use of different formats has proliferated. ATLAS should take a hard look at possibly redundant utilization of different formats by different groups for similar purposes. A decision to rely on either AOD or on the ensemble of physics DPDs is recommended to optimize resources.

Unless ATLAS makes a number of choices, the resource estimates by the CRSG for ATLAS in 2009 indicate that the computing needs for 2010 and beyond may be hard to materialize.

ATLAS has submitted a request for an increment to their T0 and CAF resources while keeping their T1 and T2 resources requirements roughly constant. The CRSG recommends that a minimum level of essential tasks that must be done at CERN should be determined and limitations enforced. The de-emphasis of the T1 and T2 roles relative to the CERN role is cause for some concern. The CRSG understands that the original ATLAS CM relied heavily on the T1s but the bulk of the first pass analyses must be done at CERN and the request is largely justified.

The updated fraction of full MC generation of only 15% of the real-event fraction is very low, and may inhibit effective calibration in 2009. Ways to improve MC generation capacity should be pursued.

Report of the Computing Resources Scrutiny Group 2008

ATLAS

RESOURCE	T0	CAF	T1 ext	T2 ext	TOTAL	Sept 07 request	Change
CPU/kSI2K	7587	5783	29391	30321	**73082**	62020	+18%
Disk/TB	650	3304	21394	14645	**39993**	36300	+10%
MS/TB	8557	1523	15050	-	**25130**	22000	+14%

Report of the Computing Resources Scrutiny Group 2008

CMS

CMS has made good progress in understanding the realities associated with their computing model. The CMS computing model appears to be viable and has successfully weathered the various simulated test campaigns.

Subsequent to discussions with CMS the CRSG produced an independent simplified spreadsheet that mirrors the CMS computing model. We find good agreement between this independent assessment and the CMS computing resource requests, with a couple of exceptions.

The Tier0 CPU requirements we compute are significantly higher than foreseen by CMS. This is due to the different data taking profile assumed. We believe that one should be more conservative as far as the Tier0 is concerned, being able to keep up with increased data rates due to the LHC duty cycle, adverse running conditions and so forth.

The Tier2 CPU requirements we compute for 2009 is 40% lower than the CMS request. While we are able to account for half of the Tier2 CPU required in terms of MC production, our scrutiny of the CMS scheduled skimming/analysis activities only amount to 15% of their requested Tier2 CPU. User driven analysis activities should account for the remainder. CMS assumes that the sum scheduled and chaotic analysis activities is comparable to the resources required for MC production.

Report of the Computing Resources Scrutiny Group 2008

CMS

Resource	Location	Original	Revised	Change	Sep 07 request	Change
CPU (MSI2k)	Tier0/CAF	16.4	**16.4**	-	58.1	-5%
	Tier1	16.4	**12.8**	-22%		
	Tier2	17.2	**15.5**	-9%		
Disk (PB)	Tier0/CAF	2.9	**3.4**	18%	17.9	-9%
	Tier1	7.4	**7.5**	1%		
	Tier2	5.3	**5.3**	-		
Tape (PB)	Tier0/CAF	8.4	**7.5**	-10%	24.3	-6%
	Tier1	16.5	**15.2**	-8%		

Report of the Computing Resources Scrutiny Group 2008

LHCb

The CRSG has analysed the CM implemented by the LHCb experiment and concluded that the model is viable and solid. The successful recent tests have demonstrated the stability of the implemented solutions and suggest that the model will enable a successful data taking at the LHC start-up and beyond.

The CRSG recommends careful reconsideration of the strategy for a full replication of the DST data at all T1 centres and to review this model assumption in the present real conditions expected for 2009.

The reduced LHC running in 2008 lead to a total CPU requirement in 2009 reduced by about 6% of the original request. The percentage impact is small because the CPU requirement is dominated by the Monte Carlo simulation as well as reconstruction and analysis of data from the current running year.

The total disk requirement in 2009 is reduced by 14%. The reduction in disk requirement comes from the fact that LHCb plans to keep on disk AOD, TAG and analysis data of the previous year (i.e. 2008). Tier-2 disk storage is not affected since Tier-2 centres are foreseen to produce only Monte Carlo data.

The reduction of tapes requirements reaches 34% and is larger than on disk because of missing RAW and rDST data from 2008 that would have been stored on tape under normal planned beam conditions.

It should be noted that the above estimates were done under the assumption of zero data from the detector and hence no analysis activities at all. However, even without a beam LHCb has already done and will continue to do analysis of cosmic ray events etc. We believe that these corresponding necessary resources, although not explicitely considered in the LHCb CM, can be handled with the resources installed in 2008.

Report of the Computing Resources Scrutiny Group 2008

LHCb

Estimation of total resource requirements for LHCb (zero beam time in 2008 and 0.9×10^7 s in 2009)		**Change**	**Sep 07 request**	**Change**
CPU (MSi2k)	**16.4***	-3%	17.4*	-6%
Disk (TB)	**3238**	-8%	3773	-14%
Tape (TB)	**3516**	-24%	5340	-34%
(*) Note that the online farm which is only used within 2 months has been removed in this summary table in order to convert the CPU requirement into the usual MoU units of the installed capacity given in MSi2k.				

Report of the Computing Resources Scrutiny Group 2008

In the process of scrutinizing the 2008 and 2009 requests of the four LHC experiments we have critically examined all possible aspects of the different computing models and their implementation.

While we find some points of discrepancy and a few potentially troublesome issues, we conclude that the overall demand of resources for 2009 largely remains within the envisaged envelope.

A very limited degree of redistribution of resources may be advisable in 2009, however care has to be taken not to harm experiments with a more consolidated CM in favour of those whose CM is less defined or consolidated at this stage.

To remain in the future within this envelope will require some updates and revisions of the computing models, perhaps of some substance in some cases. The scrutiny after the first round of real data will be of great relevance.

The CRSG believes that the different computing models have largely proven their validity and we have no doubt that they will survive their first contact with real data in 2009.

BACK-UP SLIDES

Report of the Computing Resources Scrutiny Group 2008

ALICE

2009

		CERN	T1 ext	T2 ext	Total
CPU/MSI2k	**Scrutiny**	**10**	**12**	**21**	**43**
	Request	*9.7*	*19.9*	*14.3*	*43.8*
Disk/PB	**Scrutiny**	**2.5**	**9.9**	**9.6***	**22.1***
	Request	*4.4*	*6.8*	*4.0*	*15.3*
MS/PB	**Scrutiny**	**7**	**10.1**	**0**	**17.1**
	Request	*7.4*	*12.4*	*0*	*19.7*

Report of the Computing Resources Scrutiny Group 2008

ATLAS

2009

		T0	CAF	T1 ext	T2 ext	Total
CPU/MSI2k	**Scrutiny**	**7.6**	**5.8**	**29.4**	**30.3**	**73.1**
	Request	*4.1*	*2.6*	*28.4*	*27.0*	*62.0*
Disk/PB	**Scrutiny**	**0.7**	**3.3**	**21.4**	**14.6**	**40.0**
	Request	*0.3*	*1.8*	*20.9*	*13.3*	*36.3*
MS/PB	**Scrutiny**	**8.6**	**1.5**	**15.1**	**0**	**25.1**
	Request	*5.6*	*0.7*	*15.8*	*0*	*22.0*

Report of the Computing Resources Scrutiny Group 2008

CMS

2009

		CERN	T1 ext	T2 ext	Total
CPU/MSI2k	**Scrutiny**	**16.4**	**12.8**	**15.5**	**44.7**
	Request	*13.7*	*16.3*	*28.1*	*58.1*
Disk/PB	**Scrutiny**	**3.4**	**7.5**	**5.3**	**16.2**
	Request	*2.5*	*9.7*	*5.7*	*17.9*
MS/PB	**Scrutiny**	**7.5**	**15.2**	**0**	**22.7**
	Request	*9.3*	*15.0*	*0*	*24.3*

Report of the Computing Resources Scrutiny Group 2008

LHCb

2009

		CERN	T1 ext	T2 ext	Total
CPU/MSI2k	**Scrutiny**	**0.9**	**4.1**	**11.4**	**16.4**
	Request	*1.1*	*5.0*	*11.4*	*17.4*
Disk/PB	**Scrutiny**	**0.8**	**2.4**	**0.0**	**3.2**
	Request	*1.0*	*2.8*	*0.0*	*3.8*
MS/PB	**Scrutiny**	**1.5**	**2.0**	**0**	**3.5**
	Request	*2.3*	*3.1*	*0*	*5.3*

Report of the Computing Resources Scrutiny Group 2008

The CRSG makes the following general comments recommendations:

- It seems prudent to scrutinise the experiments' use of resources after a few months of data taking in 2009. It is also important, given the resource acquisition cycle, to inform the Tier1 and Tier2 computing centres of the resource acquisition plans for calendar year 2010 as soon as possible. The CRSG commits itself to provide a scrutiny at the earliest feasible date and would recommend an earlier CRRB meeting.

  While it may be difficult in this startup period to suggest definite dates and a substantial advancement may not be feasible in 2009, we think that in future years it would be very helpful to the funding agencies and the different institutes to have a scrutiny ready by the end of summer, thus giving more time to the Tier1 and Tier2 to complete the procurement process.
- The WLCG represents a computing effort of an unprecedented scale. In spite of increasingly demanding tests being passed uncertainties remain. We recommend that the different collaborations undertake a proper risk analysis and take stock of their results in future requests in order to cope with the most likely failures or shortfalls. We feel that this assessment is particularly worthwhile for some experiments.
- In some cases the information provided to us about the AA program of the collaborations has been rather sketchy. While this may not be the main physics goal they are pursuing, and it will impact their 2009 needs in a very limited manner, it will surely have an impact on their future computing needs. We would be thankful to them for more detailed information in successive scrutinies.
- As running conditions may vary in the future (with the presence of 75ns bunch crossings) the collaborations should be aware that this has to be accommodated within the existing envelope by decreasing the event rate or similar measures.

Report of the Computing Resources Scrutiny Group 2008

General comments and recommendations (continuation)

- The experiments are asked to actively pursue the policy of reducing the size of their raw events, and other derived formats, in future years as much as possible as detectors become better understood.
- A very strict policy of removing all ‘dark’ or ‘orphaned’ data should be enforced by the collaborations.
- The CRSG recommends to the experiments to keep their computing models and needs under constant revision. We came across in this first scrutiny a somewhat conservative approach according to which some requests had not been officially modified even if it was clear that they were not realistic anymore.
- We recommend the experiments make maximal use of the distributed resources in the GRID avoiding as much as possible the use of CERN facilities.
- In the case of CERN resources, we advocate for a very clear separation between the contributions used for calibration and first pass reconstruction and central analysis (‘express stream’ or similar), and those used to perform physics analysis by the CERN based physicists.
- The CRSG wishes to state that the recommendations contained in the scrutiny 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 the commissioning and of their computing model implementation. Shortfalls of any kind would seriously jeopardize the success of the experiments. We therefore recommend that the funding agencies ensure the effective and timely delivery of the pledged resources.

Report of the Computing Resources Scrutiny Group 2008

Our scrutiny has identified several aspects that need to be brought to the attention of the LHCC.

- Most experiments propose using increased trigger rates as compared to the ones stated in the TDR reviewed by the LHCC. We feel we are not sufficiently competent to review the need or convenience to do so.
- ALICE wants to increase substantially their amount of pp data; in particular they stress the benefit of acquiring data at 10TeV. We have not assessed these needs from the physics point of view and we do not know whether such lower energies will be available in the 2009 run or anytime in the future.
- One of our conclusions is to recommend that ALICE undertakes a full assessment of how their physics reach might be affected by requested computing resources not materializing.
- The event size has a very direct impact on the computing requirements. Some experiments, such as CMS, have made an effort to reduce the event sizes by establishing a reduction profile after startup. We believe that this example should be followed by all experiments.
- We take note of potential modifications of the computing models due to the use of different data formats serving the same purposes, not always well justified.
- The realization of the computing model for ATLAS seems to differ slightly from the implementation originally envisaged in the TDR for reasons discussed in the report. This implies, in particular, heavier demands on CERN resources. We believe these demands are largely justified, however.
- Cosmic data taking is now much emphasized by experiments; while it is clear that cosmics are extremely useful in commissioning for calibration, this data is by nature transient and it seems somewhat questionable to us to support substantial requests based on cosmic runs, but we do feel we have not sufficient insight to make a definite scientific judgement on this.