

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

Members:

J. Knobloch (CERN)

W. Trischuk (Canada)

H. Marten (Germany)

J. Flynn (UK)

A. Lazzarini (USA)

H. Renshall (CERN/IT, *Scientific Secretary*)

B. Vinter (Nordic Grid)

C. Diaconu (France)

C. Bozzi (Italy)

D. Groep (the Netherlands)

D. Espriu (Spain, *Chairman*)

The CRSG wishes to thank the four experiments ALICE, ATLAS, CMS and LHCb, and in particular their respective computing managements, for the collaboration offered and their remarkable openness.

The CRSG also wishes to thank the CERN management for the help provided to the CRSG.

Particular thanks are due to the LHCC for making available to us their previous material on the WLCG.

As chairman I would like to express my gratitude to all CRSG members for their dedication and availability.

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.

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 started their activities on 10 December 2007 with an inaugural meeting chaired by the CSO where experiments presented their computing plans.

In early 2008 the chairman was appointed and the CRSG approved a work plan and calendar for 2008, and priorities. The CRSG decided to concentrate on the 2009 resource request, and considered 2008 as a test exercise for the understanding of the computing models.

Referees were appointed (2-3 per experiment) and proceeded to acquire the required critical understanding of the respective computing models.

The deliberations of the spring C-RRB meeting were taken into its work.

Agreed common but flexible template to review each experiment by the same metrics as far as possible

It had originally hoped to finish their review by September 1st to give advance notice to the funding agencies of its findings

Had to take into account the effect of the new LHC schedule on 2009 resources.

Report of the Computing Resources Scrutiny Group 2008

Several rounds of talks with the experiments took place. We noted that experiments did not wish to change their parameters until they have real data, sticking to their Sept 2007 request. ATLAS however eventually submitted a revision of their request in August.

The actual methodology changed from experiment to experiment, but uniformity in the scrutiny was a priority. The CRSG hold a total of six plenary meetings and numerous phone conferences and mail exchanges.

Decided to prepare a long stand-alone paper (~15 pages) per experiment, not part of this report but available to the experiments. A summary, recommendations and summary tables extracted (~5 pages) and joined into single report with a preamble of methodology, general recommendations and matters to be referred to the LHCC for further review form the document being submitted to the C-RRB now.

On October 6th the draft final report was submitted to the C-RRB chairman and later released to the experiments for comments.

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

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.

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

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.

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.