



Computing Resources Review Board

11th November 2008

Minutes of the 14th Resources Review Board Meeting Held at CERN on 11th November 2008

Present:

Europe:

J. Lemonne (FWO, Belgium); J. Sacton (FNRS, Belgium);
 M. Lokajicek, J. Niederle (Institute of Physics AS CR, Czech Republic);
 J.D. Hansen (Niels Bohr Institute, Denmark);
 D-O. Riska, J. Tuominiemi Riska (Helsinki Institute of Physics, Finland);
 E. Auge (CNRS/IN2P3, France); F. Malek (CNRS/LCG-France, France);
 S. Bethke (MPI, Germany); C. Diaconu (CPM/DESY, Germany); K. Ehret, M. Fleischer, V Guelzow (DESY, Germany); T. Hebbeker (RWTH, Germany); P. Malzacher (GSI, Germany); K-P. Mickel (Karlsruhe Inst. of Technology, Germany);
 G. Vesztergombi (KFKI-RMKI, Hungary);
 L. Levinson (Weizmann Institute of Science, Israel); E. Rabinovici (Hebrew University, Israel);
 F. Ferroni (INFN, Italy);
 A. Bernotas (The Lithuanian Academy of Sciences, Lithuania);
 A. Van Rijn (Nikhef, Netherlands);
 B. Jacobsen (The Research Council of Norway);
 M. Turala (IFJ PAN, Poland);
 G. Barreira (LIP, Portugal);
 F-D. Buzatu (Institute of Atomic Physics, Romania); L. Puscaragiu (Permt. Mission of Romania);
 Y. Kozlov (Russian Federal Agency of Science and Innovation); R. Lednicky, A. Vodopyanov (JINR, Russia); A. Petrov (Permt. Mission of Russia); V. Savrin (Institute of Nuclear Physics, Russia);
 N. Colino (CIEMAT, Spain);
 D. Espriu (Universitat de Barcelona, Spain);
 T. Ekelöf (Uppsala University, Sweden); P. Karlsson (Swedish Research Council);
 T. Nakada (CHIPP and EPFL, Switzerland);
 I. Turk Cakir (TAEA, Turkey);
 J. Seed (STFC, United Kingdom).

N. America:

I. Blain (NSERC, Canada); W. Davidson (NRC, Canada); R. Mcpherson (University of Victoria, Canada);
 L. Bauerdick, J. Butler (FNAL, U. S. A.); A. Boehnlein, S. Gonzalez, M. Procario (DOE, U. S. A.);
 R. Cousins (UCLA, U. S. A.); H. Gordon (BNL, U. S. A.); M. Pripstein (NSF, U. S. A.);
 M. Tuts (Columbia University, U. S. A.); J. Shank (Boston University, U. S. A.).

Asia:

P. Ji, Y. Zhang (NSFC, China);
 Y-I. Choi (Sungkyunkwan University, Republic Of Korea); J. E. Ha (Minister of Education, Science & Technology, Republic Of Korea); D.M. Lee (KICOS, Republic Of Korea);
 A. Gurtu (Tata Inst of Fundamental Research, India); P. Mukherjee (DAE, India);
 T. Kawamoto (University of Tokyo, Japan);
 S.C. Lin (ASGC, Taiwan).

CERN:

J. Engelen (Chairman), H. Renshall (Scientific Secretary), G. Lafferty (Univ. Manchester, Scrutiny Group),
 S. M. Schmeling, D. Jacobs, T. Lagrange, E. Van Hove, W. Von Rueden.

LCG:

I. Bird, S. Foffano;

ALICE:

F. Carminati, J. Schukraft, Y. Schutz;

ATLAS:

D. Barberis, F. Gianotti, P. Jenni;

CMS:

M. Kasemann, J. Virdee;

LHCb:

A. Golutvin, A. Schopper.

1. Introduction

J. Engelen, Chief Scientific Officer

J. Engelen welcomed delegates to this fourteenth meeting of the Computing Resources Review Board. He asked delegates to please identify their name and institution when making an intervention.

2. Approval of the Minutes of the 13th Meeting (CERN-RRB-2008-079)

J.Engelen asked delegates to give their approval or make any comments on the minutes of the last meeting, CERN-RRB-2008-079. M.Turala (Poland) referred to his intervention concerning resource usage following the 'Status of Resources and Financial Plan' presentation. He had asked that experiments should also look into the resource usage numbers and comment and, if he remembered rightly, this was followed by experiment statements but this is not reflected in the minutes. J.Engelen agreed this should be corrected. (Later replay of the audio recording showed there were, in fact, no experiment statements at that time so these must have been made in another forum - H.Renshall, secretary).

3. Status of the LCG Project

I.Bird, Project Leader

Paper CERN-RRB-2008-101

Presentation CERN- RRB-2008-102

J.Engelen proposed to now move to the report on the status of the LCG project.

I. Bird began by saying that the report covered the period from the last C-RRB meeting in April 2008 and, referring to the agenda on slide 2, that improving the WLCG services and executing the CCRC'08 challenge were the major activities during this period. He said that in general the WLCG service has been quite successful.

His slide 3 summarised these activities. The April target for all 2008 resources to be in place had not been met but this had not stopped the CCRC'08 activities. The main things not tested were reprocessing with tape recall at Tier 1 sites for more than one experiment concurrently and analysis at a scale of hundreds of users. He thought that how to handle such analysis would be an unknown over the next couple of years. CCRC'08 was run in a sustained way during May and since then there has been continually increasing workloads of simulations, cosmic data processing and functional tests of all aspects of the service. Slide 4 showed graphically the performance of CCRC'08 and beyond. A graph of jobs run per month in WLCG had a peak in May but has since continued to grow. A plot of the CMS inter-site data traffic during CCRC'08 showed that this was very successful with up to 200 TB of data moved per day and the plot of average data throughput from CERN to Tier 1 sites during CCRC'08 showed that the target rate of 1.3 GB/s was easily exceeded. The ongoing nature of these activities was illustrated by the 4-months plot of throughput from CERN to Tier 1 which included a dip in load preceding the LHC first beam day followed by a resumption to it's plateau level. His 5th slide, on usage patterns, showed the difference in cpu usage where Tier 2 contributed 45% during the May CCRC'08 while in September they were contributing 48%. I.Bird commented that neither of these was illustrative of the workload we expect with real data.

He then moved on to outstanding service issues commenting that during incidents we clearly saw cases where people were on vacation and unable to respond. He added that the recovery process after power or cooling failures was an area we could improve by correct configuration. Slide 7 showed service incidents since May demonstrating that we can expect at least one major incident per site per year and that there is very little that can be pinpointed in the middleware or software to avoid this. Slide 8, on reliabilities, showed a plot of the results of test jobs that probe services at a site. There has been a gradual improvement over the years accelerating since May this year especially at Tier 2 sites and this was encouraging. He added that we are also starting to publish the results of VO-specific tests. The next slide showed Tier 2 reliabilities with numbers given per federation. The aim is to get all Tier 2's to be above 90% efficient with the target being 95%. Some sites have services that are reliable when up but are only up for half the time. A few federations indicated are still not reporting and I.Bird encouraged them to do so. The final slide on

service issues was on VO-specific testing where validation will be in progress over the next few months. As an example of the results RAL had storage problems in August where the site was overall available but the ATLAS services were not.

I.Bird then moved on to the consequences of the LHC shutdown, starting at slide 11, saying that the LCG management board has considered what to do during the shutdown in the three areas of capacity and procurements, upgrade plans and re-validation of the service as shown on the following three slides. He emphasised that in the current plans the amount of data to be gathered in 2009 might be close to or exceed that originally anticipated for 2008 and 2009 together. On the upgrade plans the agreed list of SRM changes will go ahead but there will be no more significant changes until we have real data. The glxec/SCAS middleware is to be used by pilot jobs that pull down workloads from different users and is wanted for next spring and new clients and a WMS submission interface are needed to make the CREAM compute element available. The program of work behind these upgrades needs to be done but in a way that maintains the running services. Finally he emphasised that we need to think about re-validating the service next year.

Slides 15 and 16 discussed the applications area mentioning that two other R&D projects in the areas of operating system virtualisation and multi-core architectures are maintained by the architects forum.

The presentation then moved on to procurement issues in slide 17 showing that not all 2008 resources were in place by April and that some are still missing. This is not a particular problem for this year but it is worrying. Two sites currently had problems – ASGC which is installing now and the NL-T1 which needs additional power and will not be complete till mid-2009. There are many issues behind these delays, such as faulty equipment. Unfortunately there is no indication that this will improve and we are hitting the limit of what people can deal with. I.Bird mentioned the proposal to split disk acquisitions into two per year emphasising that a second round must not be too late with deployment targeted before the end of year.

On capacity and benchmarking, slide 18, I.Bird hoped that the automated gathering of installed capacity for Tier 1 should be in place by the next RRB though we do not have an automated mechanism for Tier 2 sites. The benchmark proposal includes that sites will be required to run an agreed exact recipe. A web site of results will be set up to help Tier 2 sites.

The subject then moved on to planning covering the end of EGEE in 2010, planning for EGI and the CERN infrastructure (slides 19 to 22). A particular issue for LCG is the current lack of representation of the user communities implied in the EGI blueprint. Also the lack of clarity in the evolution of the present operational infrastructure of the Tier 1 and Tier 2 in Europe that support the WLCG is a worry. I.Bird emphasised the importance that WLCG members raise these issues to their NGI and national funding agencies in order to ensure they get the required operational support. At the same time the WLCG Oversight Board has agreed we should now document a concrete plan on how WLCG could operate without relying on the European level support, either temporarily or indefinitely. Over the next few months we will document how each of today's core functions will be managed in future if there is insufficient support from EGI. As regards the CERN Tier 0 planning I.Bird said there is probably now sufficient capacity to last until the end of 2010 instead of the beginning and that four solutions are being followed including expansion in the current building by adding water-cooled racks and aggressive replacement of older power-hungry equipment. In addition CERN is trying to understand what the needs would be in a new second site at Preessin and in possible stop-gap solutions outside of CERN.

I.Bird then presented a summary slide, slide 23, reminding that in the 2008 data challenges some experiments did some reprocessing but not at the level of stressing the Tier 1. His final point was that he did not see resource procurement improving in future years so we must plan around this.

Discussion

J. Engelen thanked I. Bird for his presentation and invited comments.

The first intervention was by F.Ferroni (INFN, Italy) who asked I.Bird if he could show a summary of CERN and Tier 1 resource usage. The reply was that this will be shown in the Status of Resources presentation. F.Ferroni then made the comment that the WLCG is asking for a lot of resources for 2009 and, given that the resources were committed in 2008 on the idea that LHC would start in July that year and that the resources available (in INFN) have only been used to 30%, those resources were good enough for 2008 while now 2009 will be like 2008 was supposed to be so what is the need for any new resources at all? I.Bird replied that he expected to come back to this after the next talk but that it was difficult to make a judgement on current usage with the accelerator off. F.Ferroni repeated his view that the resources that were good enough for 2008 should equally be good enough for 2009. I.Bird answered this by saying that by the end of 2009 there should now be as much data as had been previously planned for 2008 and 2009 combined. F.Ferroni thought this was a dream and I.Bird replied that this was based on the information they currently had in which case they have to plan to have the resources available. J.Engelen added that he saw this as a plea for an optimisation of resources that was legitimate but difficult being offset by resources that were not available on time. He assured F.Ferroni that his point was well taken and that the WLCG must work with the experiments to optimise resources but that it should not be done in an overly conservative way.

The next intervention was from T.Ekelöf (Uppsala University, Sweden) who asked if the worries exposed about the future of EGI were mainly because of organisational aspects or those of funding? I.Bird replied that by now we had hoped to have a blueprint with a transition plan from EGEE to EGI but he did not see this emerging over the next few years nor see the end point. As a result we must plan to be independent of this process as much as possible and ensure that all software that we absolutely rely on be supported in some way.

A. Gurtu (DAE, India) then remarked that for Tier 2 centres in particular that even if resources were in place it appeared to be difficult to properly measure their availability and reliability nor the number of jobs passed. He suggested that a good way to bring all the centres up to speed would be for the experiments to focus on a few weak centres at a time and try and bring them up to a 90% reliability level. This would be a good systematic approach. I.Bird thought this was mixing two things because the reliability of a site does not depend on the number of jobs sent there but is an independent measure of the service done by sending probes. The experiments also have their own sets of such probes and are doing systematic studies. It is true that when you look at the accounting of a site you cannot tell if low values are because the site is unusable or the experiments are just not sending work there.

J.Engelen (chairman) then pointed out that I.Bird has now given the projection that CERN will have sufficient infrastructure (power and cooling) to satisfy requirements till 2010 and asked if the Management Board also discussed similar problems at the Tier 1 centres adding that at the Dutch Tier 1 they appear to have money for hardware resources but nowhere to put them. I.Bird replied that the MB has indeed discussed power issues many times and that any Tier 1 that comes up with this as a problem now has no excuse though this has been a source of surprise in the IT industry as a whole. He hoped there would be no more such surprises in the future. E.Auge (CNRS/IN2P3, France) remarked that the situation in some Tier 1s is similar to that of CERN. Problems with power and cooling have been identified but still the delays are long and lead to stopgap solutions which are expensive. For example in Lyon the need will be for 2009 not 2010 and this is why they are sensitive to the previous remarks of F.Ferroni on the need for expansion in 2009 as in order to fulfil their pledges and be really useful, as they wished to be, they will have to adopt expensive stopgap solutions. J.Engelen then reminded the meeting that they would later be hearing a detailed report from the Computer Resources Scrutiny Group to assure us or not if the requests are genuine and that the pledges would need to follow. In conclusion of this point A. Van Rijn (Nikhef,

Netherlands) underlined that planning for computer centre power and cooling requirements is a tedious effort that should not be underestimated.

J.Engelen then gave his conclusions to this item that we should not forget that this fantastically complex system (of the WLCG) is delivering, it works and it is there to stay. This is in spite of the fact that EGI is not forthcoming, which he found disappointing and for which he did not fully understand the reasons, but it was reassuring that in the Management Board they were preparing fallback solutions in order to not depend on an infrastructure that may not be there. As far as the capacity planning is concerned the interventions by Ferroni and Auge are on record and he expected to come back to this later in the meeting.

G. Barreira (LIP, Portugal), who is also chairman of the EGI Policy Board, then made a statement in reply about EGI. He supposed members are all looking at the EGI report but said that it is difficult to understand. Up till now grid efforts had been based on two legs – one being CERN with WLCG and the other being the European Union supporting projects from data grids to EGEE with some site projects like crossgrid. EGI is a different player made up of NGIs which depend on funding agencies and on governments. These are new players that were not involved before in the grid effort so the effort to establish EGI is enormously more difficult. He thought that till now they had been successful with now 38 European NGI represented at the policy board level. A difficulty is that the NGI do not all have the same status and while the EU is demanding they be legal entities this is not the case for all of the NGI. To make them such, to be government bodies, is enormously difficult as each country has to do this. Even in well organized countries like Spain and Portugal the NGI are not legal entities so this will take time. CERN is part of the EGI-DS consortium but to his surprise CERN never applied to have associate status on the policy making board and he hoped this would be taken into account. He did not know how successful EGI would be in the next 12 months and added that bidding to host the EGI site was now open, with some 5-6 countries already bidding, and they hoped to decide around March. He did not, however, anticipate that all the 38 NGI would be fully ready by the end of next year estimating that perhaps one third would have achieved the necessary legal status. EGI will continue discussions at the beginning of next year with some changes to increase flexibility being proposed. He concluded by saying that we (High Energy Physics) should not discourage the future of EGI but anticipate that it will be a difficult task as it was for the installation of the European Communication network which is now a great success. Governments in many countries are financing their grid efforts under the assumption that they will not serve only the interests of the particle physics community but of a wider community and he hoped that CERN would become more involved with the EGI effort. J.Engelen thanked G.Barreira for his clarifications and wise words and agreed that CERN management will come back to this issue outside of this meeting and stated that the goodwill of CERN was not in doubt. Finally on this point I.Bird repeated that whether he was optimistic or pessimistic it was necessary to have a plan to deal with the eventuality that EGI may not happen and J.Engelen agreed that we need to be practical as well as trying to be constructive.

4. LHCC Deliberations (paper only)

Paper CERN-RRB-2008-103

J.Engelen stated that these deliberations are consistent with what the delegates had just heard, if not a little obsolete, and took it that members had taken note of them.

E. Tsismelis, LHCC Scientific Secretary

5. Status of Common Project Accounts

Paper CERN-RRB-2008-067

T. Lagrange, CERN Finance Dept.

J.Engelen welcomed the new head of Finance department, Thierry Lagrange. He then reminded delegates that the Common Projects were absolutely vital in getting the LCG effort going. At this moment they were still important but relatively small compared to the resources that the delegates represented.

T.Lagrange then announced that there had been no changes to the document issued by the management at the end of August so he had no further comments to make unless there be any questions. None were forthcoming so the chairman closed this item and invited the LCG Resources Coordinator, Sue Foffano, to present the next item.

6. Status of Resources and Financial Plan S. Foffano, CERN

Paper CERN- RRB-2008-104

Presentation CERN- RRB-2008-105

S. Foffano began by mentioning that in her paper report she now included an annex containing the LCG Tier 0, 1 and 2 pledges up to 2013. She reviewed the agenda of her presentation (slide 2) pointing out that we were about to reach the end of phase 1 of the WLCG project (2005-2008) and enter phase 2 (2009-2012). She showed the WLCG web pages address and reported that these pages have been updated to have a more consistent way of finding information and added that following input from RRB members she would later be proposing a shorter forward planning cycle.

6.1 WLCG MoU Signature Status

S. Foffano reviewed (slide 3) the status of signatures of the WLCG MoU. The Austrian Tier 2 signed yesterday so all Tier 1 and 2 have now signed. Korea have already signed as an ALICE Tier 2 and will now sign as a CMS Tier 2 while Brazil, though anxious to sign as a Tier 2 federation supporting all 4 experiments, are not yet in a position to do so.

6.2 Funding and Expenditure for WLCG at CERN

S. Foffano then showed in slides 5, 6 & 7 the CERN funding and planned expenditure for 2005-2008 and 2009-2012 of the LCG project at CERN in MCHF. The books have been closed on the years 2005, 6 and 7 so nothing has changed while there have been small revisions for 2008 where the book closing will be next April. The current estimate for 2008 is that there will be a carry forward of 1.4 MCHF which includes the external contributions. The estimates for 2009-2012 include all the items reported by Ian Bird, for example the new CERN infrastructure. The estimation is that there will be a small positive balance of some 0.2 MCHF at the end of 2012.

6.3 Tier 1 and Tier 2 Resource accounting

S. Foffano then showed graphs (slides 9,10) giving a summary of the CPU Time, Disk and Tape Storage accounting from January 2008 to September 2008 obtained by summing up CERN and all external Tier-1s. The graphs showed the installed capacity, the pledged capacity and the resources usage. She noted that the sharp increase in MoU pledged capacity in April was the normal annual increase and that though cpu installation followed well the pledges those of disk and tape suffered from delivery delays and faulty equipment. She remarked that we must ensure better convergence of all resources in future.

S, Foffano turned to accounting for Tier-2s (slides 11 to 13). She noted that collection of these data is not as automated as for the Tier 1 but this is being worked on and she encouraged sites to review and verify their data. Slide 12 showed the September accounting for the 10 largest Tier 2 sites as being easier to read than slide 13 showing the other 49 Tier 2.

6.4 Computing resource pledges

Slide 15 reminded of the resource pledge responsibilities where, currently, the Autumn C-RRB should be given planned pledge values for the next 4 years as she had been soliciting over the last few months. In the spring C-RRB there is a check to see if sites met their commitments and she encouraged sites to check, in the annex of the paper, that their information is correct.

Slide 16 showed the pledge balance for 2009 showing the split across all the experiments. As regards the missing pledges from three German federations she would be meeting with them that afternoon and she added that resources from the Korean MoU were not yet included. When the missing information was obtained she would republish the figures on the LCG web site but the improvements would only be small.

Slide 17 showed the global picture for 2008-2013 where the experiment requirements were from 2007 and are expected to be revised after the 2009 data taking. Experiments are saying they cannot give requirements as far ahead as 2016 while some sites are pointing out they only have budgets approved to 2011.

This lead on to slide 18 showing the new proposal to reduce the site pledge timeframe to be 1 year ahead of confirmed pledges then a further 2 years ahead of planned pledges.

6.5 Conclusions

In conclusion (slide 19) S. Foffano added that:

- Only Brazil needs an MoU signature follow-up.
- The global funding situation till 2012 suggests the project will not run into major financial difficulties.
- Most Tier 1 and Tier 2 sites are now reporting their accounting.
- New experiment requirements and scrutiny group recommendations are to be fed into the next round of pledges where the balance today still shows a significant shortfall.
- It is proposed to reduce the pledge timeframe from 5 to 3 years.

Discussion

J. Engelen then opened the floor for any comments or questions.

F.Ferroni (INFN, Italy) first made a statement that INFN had submitted their pledges for 2009 but with the caveat of not knowing yet the accelerator schedule for 2009. If they learn something different from what they have been told, so the boundary conditions have been changed, they would feel free to modify their pledges accordingly. He then referred to the transparency of user jobs accounting for 2008 where he said he was not regretting to have committed the cpu resources for 2008, done on the assumption that the LHC was going to produce data in July 2008, but that these were provided to satisfy the needs of data taking in 2008. In his mind the schedule had just been shifted by one year so the data taking of 2009 should be equivalent to the data taking of 2008. If those cpu were good enough for 2008 they should still be good enough for 2009 unless something drastic has changed which he queried. S.Foffano replied that, as already mentioned by I.Bird, the ramping up process always envisaged a doubling of resources in 2009 over 2008 and she would be interested to hear other comments. As regards the first statement she pointed out that such reservations were what she had tried to bring out in footnotes (For INFN see Note 1 of the LCG Tier 0 and 1 Resources annex) in her paper and offered to change them if so required. She hoped, however, not to have to change any of the numbers in her tables. I.Bird then added that in terms of what their plan is they had asked the experiments to go away and look again at their requirements. They had all replied that the original agreed resource planning should be maintained as what they foresaw doing in 2009 would require all of those resources. The scientific merits of these statements now have to be judged by the LHCC.

J.Engelen proposed at this point to listen to the report by the scrutiny group chair, who also had comments on their interactions with the LHCC, before coming back to this discussion.

6.6 Report from the C-RSG

Paper CERN-RRB-2008-106

D.Espriu

Presentation Slides

D.Espriu began by saying he hoped to clarify some of the issues that had just been raised. He added that this report was more quantitative than the one he gave yesterday to the plenary RRB and that the C-RSG believed their recommendations to be a minimum of what the experiments need.

His first two slides summarised the role and mandate of the C-RSG. He pointed out that he was only reporting today on the resources of 2009 and that the group had been hoping to use the

realities of data in 2008 as a proof of their concepts. As an executive summary he stated that, with a few exceptions, no gross discrepancies had been found in the resource requests but that some of the discrepancies found between the scrutinised needs and the historic requests were a source of worry. In slide 3 he showed the groups working assumptions on beam time for 2008, 2009 and 2010 up until the LHC failure then the revised values of zero beam time in 2008 but assuming 2009 and 2010 would be normal years with $9 \cdot 10^{*6}$ and 10^{*7} seconds of pp beam time respectively and each with 10^{*6} seconds of heavy ion collisions. He stated the group was being very optimistic in this but could not allow data to be lost.

He then presented (slides 4 to 11) two slides for each LHC experiment summarising the groups findings and recommendations of CPU, disk and mass storage for 2009 broken down into CERN, external Tier 1 and external Tier 2 levels and indicating the percentage change the group is recommending compared to the official experiment requests from September 2007.

For ALICE he stated that there are still some unknowns in their computing model. The total CPU requested agreed with the scrutiny but the scrutiny found a different distribution with more required at the Tier 1 level and less at the Tier 2 level. The scrutiny found that ALICE had underestimated their disk requirements and recommended a 44% increase including more than doubling of the quantity requested at the Tier 2 level. They recommended a 13% reduction in mass storage since the ALICE model used in their request accumulated data from an assumed 2007 startup. Finally he observed that it was clear that ALICE's computing requirements are unlikely to be met in practise and the C-RSG was recommending that ALICE make a clear statement to the LHCC of how their physics program will be affected.

The ATLAS requests of September 2007 assumed only two-thirds of the 2009 beam time that the C-RSG was assuming but the effect of this was largely compensated for by the lack of data in 2008. The C-RSG found that the storage requested was consistent with the ATLAS computing model but remarked that event sizes have grown and recommended that ATLAS take a fresh look at possibly redundant usage of different event formats by different groups for similar purposes. ATLAS had recently submitted a request for an increment to their Tier 0 and CAF resources while keeping their Tier 1 and Tier 2 requirements roughly constant, largely on the grounds that the bulk of the first pass analyses must be done at CERN, and the C-RSG found this to be largely justified. The net result was a recommended increase over the 2007 requests of 18% in CPU, 10% in disk and 14% in mass storage, essentially due to the increased resource demand at CERN. The C-RSG also noted that ATLAS propose a full Monte-Carlo generation equal to only 15% of their real events which might inhibit effective calibration in 2009 and recommended that ways to improve their Monte-Carlo generation capacity be pursued.

For CMS the C-RSG found the experiment had made good progress in understanding the realities associated with their computing model. The C-RSG produced an independent simplified spreadsheet that mirrored the CMS computing model and which found some exceptions to their requests namely the C-RSG found higher Tier 0 CPU requirements but lower Tier 2 CPU requirements. The net result, folding in the absence of beam data in 2008, was to recommend overall a 5% drop in CPU, a 9% drop in disk and a 6% drop in mass storage as compared to the CMS requests of September 2007.

The C-RSG concluded that the computing model implemented by the LHCb experiment was viable and solid. The reduced LHC running in 2008 lead to a 6% reduction in the total CPU requirement in 2009, a small impact because their CPU requirement is dominated by Monte Carlo simulation. The total disk requirement was found to be reduced by 14% coming from the fact that LHCb plan to keep on disk all of a previous years AOD, TAG and analysis data. Tier 2 disk requirements do not change as their only function is to produce Monte Carlo data. In mass storage the reduction recommended is by 34% and is larger than for disk since there will be no raw data from 2008. The group noted that LHCb will continue to do analysis of cosmic ray and Monte Carlo events but concluded that this could be handled with the resources installed in 2008.

Finally on slide 12 D.Espriu summarised the conclusions of the C-RSG in their scrutiny of 2009. In the process of scrutinizing the 2008 and 2009 requests of the four LHC experiments the group

had critically examined all possible aspects of the different computing models and their implementation. While they found some points of discrepancy and a few potentially troublesome issues, they concluded 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 Computing Model in favor of those whose computing model 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 C-RSG believes that the different computing models have largely proven their validity and has no doubt that they will survive their first contact with real data in 2009.

J.Engelen thanked D.Espriu for this precise report on the conclusions of his group and invited questions and comments.

M. Turala (IFJ PAN, Poland) said he had a comment to make after the last two presentations namely that he felt uneasy seeing the tables of S.Foffano showing that the resources balance is usually red and then seeing the plots which show the usage of resources which are pledged is usually at the level of 60 to 70%. This is not a satisfactory situation. He understood that the scrutiny group did not want to play a role of negotiator but we have to do something about this. We have resources that are not used and experiments usually request much more than they need and they use. We never hear directly from experiments at this meeting and at the previous meeting he had asked if experiments could confirm if they were happy even though they were not getting all the resources they were requesting. He thought we should, in the future, attack this problem of resources that are not used adding that in his country he followed this issue and when an experiment could not use its resources he encouraged other experiments to do so. F.Carminati (ALICE) volunteered an answer from ALICE. He said that ALICE have tried to use all their resources but this is not trivial to achieve. One thing is that there were long periods when the pledges were not met and also some centres do not run their jobs even if the resources are there. It is difficult to disentangle the reasons why there is a low global efficiency but they would not be happy with less resources. They really try hard to use them all but come up with only 60 to 70% efficiency so have work to do there.

F.Gianotti (ATLAS) then made a statement in reply to those of M.Turala and the earlier remarks of F.Ferroni. On the issues of storage resources ATLAS have taken, and are still taking, cosmic data and have taken more cosmic data than originally foreseen as these data are the only ones the experiment has to debug and calibrate the detector before collisions start. Although the rate is lower than the 200 Hz for beam the events are larger as more of the detector is read out than for collisions so the cosmic data volume is five times greater than what ATLAS will have with collisions data. Cpu resources are another story and, as F.Ferroni has said, depend on what happens in 2009. The original management input was there would be an LHC resumption early in 2009 where the C-RSG assumed $9 \times 10^{**6}$ seconds of beam time in 2009 whereas in the past ATLAS had been assuming $6 \times 10^{**6}$ seconds of beam time in 2009 and the experiment needs to deal with that. Before there is a clear schedule ATLAS is unable to say if the cpu resources will be enough or not.

J.Virdee (CMS) then remarked that the CMS requirements stated in August 2007 had assumed $3 \times 10^{**6}$ seconds of beam in 2008 with 300 events/second and $9 \times 10^{**6}$ seconds beam in 2009 with 200 events/second and the guidance CMS needs is the number of seconds of beam. For them the amount of disk, however, does not change very much – at CERN it goes from 2.2 to 2.5 PB, hardly a change, while Tier 1 disk goes from 7.2 to 9.7 PB, also not a big change. What matters is really the number of seconds as has been said.

A.Golutvin (LHCb) added an information that was specific for LHCb that they would be starting a big computing exercise called FEST (Full Experiment System Test) which is trying to replace data which they hope to use to commission their high level trigger, calibration and alignment and otherwise they endorsed the C-RSG report.

J.Engelen thanked the members for their comments and summarized that the message is that the experiments are not just sitting and waiting but they are using their computing resources to be even better prepared for next year when real data comes. He thought the simple translation of resources that F.Ferroni has proposed would be an oversimplification. J.Schukraft then pointed out the particular case of ALICE. All have agreed the resources needed are scaled by the beam time we will have so ALICE is waiting for these numbers to firm up but has the particular case of missing a lot in 2008 so if they are getting in 2009 the resources they needed in 2008 they would actually be happy. S.Lin (ASGC, Taiwan) stated that his team was responsible for the ASGC Tier 1 and that he took the pragmatic view that what is required should be the ultimate target and speculated that resources could be deployed early if they had the funding or this could be done much later with the possibility that unexpected things might happen. He thought that a 6 month delay would not create operationally such a big difference for them so they did not object to such a schedule. J.Engelen continued that he still persisted in his conclusion that a simple translation would be an oversimplification. He then returned to M.Turala's comments saying that he believed that sometime next year when the experiments have the first operational experience with colliding beams that the two aspects of interest, the implementation of the models and the adequacy of the models themselves and whether physics is being lost or not either due to the models or lack of resources, then a longer report on behalf of the LHCC rather than the paper reports submitted now would be in order. This would allow both pictures to be visible in this meeting giving a more complete view and this would address the issue of having a good communication on the limitations the experiments might experience once real data is there. He suggested this be put on the C-RRB agenda for a year from now.

6.7 Time Frame for Computing resource pledges

J.Engelen then moved discussion on to the proposal from S.Foffano, the LCG resource coordinator, to change the time frame of LCG resource planning. S.Foffano repeated that the proposal before the C-RRB was that for the future pledge exercises members would be asked to commit to the year ahead as is the case today but rather than providing four extra years of planned information to provide two extra years. In the autumn of 2009 she would be asking members to commit to 2010 and give plans for 2011 and 2012. This would be based on the same time frame for the experiments who are also interested in moving to this planning. J.Engelen asked the members if they could all agree to this practical proposal. S.Bethke (MPI, Germany) replied that it makes sense not to go too far ahead for the funding agencies but that these needed the longer term requests by the experiments as they needed to plan for the power consumption and things like that so if they had a proposal of what would be needed for the years to come that would make more sense. J.Engelen agreed this was a good idea and proposed that for the practical planning we do as proposed by S.Foffano but for the planning ahead, preparing for the less formal part, that the scrutiny group looks into the requirements on a longer time scale.

7. Summary

J. Engelen

There being no further comments J. Engelen announced the imminent conclusion of this meeting of the Resources Review Board. He did not think modifications were foreseen for the Scrutiny Group for next year so there was no need to cover this point. In summary he reemphasised that the system is working and that it is ramping up according to expectations. The board had had a discussion, understandable and justified, on the change to the required resources as a consequence of the delayed startup of the LHC and that for himself he had concluded that we have strong arguments from the experiments, corroborated by a report from the scrutiny group, that the simple translation by one year of requirements is an oversimplification that would be dangerous and that we should be more ambitious than that.

He then invited the members to make any last remarks. M.Pripstein (NSF, USA) made the comment that computing is the common denominator of all the experiments and as a result this computing RRB has more attendance than the other working sessions so an observation worth making here is that this is the last set of RRBs over which Jos will be presiding and he would like

to use this occasion for the members to express their gratitude for his tremendous stewardship of the program. This was greeted with warm applause from the members to which J.Engelen expressed his thanks and the meeting concluded.

The next RRB meeting in 2009 is scheduled to take place at CERN on
Monday 27 to Wednesday 29 April 2009

H.Renshall
20 February 2009