Computing Resources Review Board

10th May 2006

Minutes of the 9th Resources Review Board Meeting
Held at CERN on 25th April 2006

Present:

Europe:
C. Wulz (HEPHY, Austria) *(Observer)*;
J. Lemonne (FWO, Belgium) ; J. Sacton (FNRS, Belgium);
M. Lokajicek (MSMT CR, Czech Rep.), M. Sumbrera, V. Vrba;
J. Dines Hansen (National Science Research Council, Denmark);
D.O. Riska (HIP, Finland);
J. Zinn-Justin (CEA/DSM/DAPNIA, France), P. Rebourgeard;
F. Le Diberder (CNRS/IN2P3, France), F. Etienne, D. Boutigny;
J. Richter (BMBF, Germany) ; S. Bethke (MPG, Germany) ; K.P Mickel (FZK, Germany);
V. Gülzow (DESY,Germany) ; D. Müller (GSI, Germany);
E. N. Gazis (National Technical University of Athens, Greece) *(Observer)*;
G. Vesztergombi (KFKI/RMKI, Hungary) *(Observer)*;
L. Levinson (Weizman Institute, Israel) *(Observer)*;
U. Dosselli (INFN, Italy), F. Ferroni;
F. Linde (NIKHEF, The Netherlands), A. J. Van Rijn;
B. Jacobsen (Norwegian Research Council, Norway);
G. Polok (Ministry of Science & Education, Poland), M. Turala;
F. Bello (GRICES/FCT/UMIC, Portugal);
F-D. Buzatu (Nat. Authority for Scientific Research, Romania), L. Puscaragiu;
Y. Kozlov (Federal Agency of Science and Innovation, Russia), V. Savrin;
A.N. Sissakian (JINR, Dubna), A. Vodopianov;
D. Espriu (MEC, Spain), N. Colino;
A.C. Lagerkvist (Research Council, Sweden);
A. Rubbia (SER/SNF/ETH/CSCS, Switzerland), C. Grab, U. Straumann;
G. Zinovjev (National Academy of Sciences of Ukraine) *(Observer)*;
R. Wade (PPARC, United Kingdom), R. Jones;

N. America:
W. Davidson (NSERC, Canada), R. Orr; C. Oram;
M. Pripstein (NSF, U.S.A.), H. Gordon ;
J. O'Fallon (DOE, U.S.A.), T. Ferbel, D. Green, S. Gonzalez;
R. Cousins, J. Shank, M. Tuts;

Asia:
P. Ji, Y. Peng (MoST/NSFC, China), Y. Zhang, Y. Huang, C. Jiang;
C. Ananda Bose (DAE, India); P. Dhekne;
T. Kawamoto (University of Tokyo, Japan), H. Iwasaki;
J-H. Kim (Ministry of Science and Technology, Korea), D. Son *(Observer)*;
S.C. Lin (Academia Sinica, Taipei), S.C. Lee;

Australia:
S. Tovey (AusHEP, Australia);

CERN:
R. Aymar (Director General), J. Engelen (Chairman), C. Jones (Secretary),
P. Geeraert, D. Jacobs, A. J. Naudi, E. Tsesmelis, W. Von Rueden;
LCG:
L. Robertson, C. Eck;
ALICE:
J. Schukraft, Y. Schutz, F. Carminati;
ATLAS:
P. Jenni, D. Barberis, F. Gianotti;
CMS:
M. Della Negra, L. Taylor, J. Virdee;
LHCb:
T. Nakada, N. Brook.

9th Meeting of the Computing Resources Review Board RRB, 25th April 2006

1. Introduction — J. Engelen, Chief Scientific Officer

J. Engelen welcomed RRB delegates to this 9th meeting of the Computing Resources Review Board. The aim was to gain an understanding of the current status of the project, and to discuss possible shortfalls and what to do about them.

2. Approval of the Minutes of the 8th Meeting (CERN-C-RRB-2005-111)

The minutes of the 8th meeting, CERN-C-RRB-2005-111, were **approved** with no corrections. J. Engelen thanked C. Jones for having taken these minutes. There were no matters arising.

3. Status of the LCG Project — L. Robertson, Project Leader

Paper CERN-RRB-2006-014 Presentation CERN- RRB-2006-061

L. Robertson reported on the status of the LCG Project. The details are reported both in his paper and his presentation, which are both referenced above, and are not further abbreviated in these minutes.

He summarized the current status as follows:

LCG depended on long term funding of two major science grid infrastructures to provide the basic operational environment, namely EGEE - Enabling Grids for E-Science, and OSG - US Open Science Grid.

Grids were now operational with around 200 sites involved in EGEE and OSG, and Grid operations centres had been running for well over a year. There were more than 20K LCG jobs per day accounted. However there was a long way to go in terms of reliability.

Data recording and Tier-0 testing at CERN was on target. Data distribution from CERN to Tier-1s was close to the nominal rates required for LHC operation for short periods. They still needed to push rates up to reach the full target, to achieve continuous operation at these levels, and to include recording to tape at the Tier-1s.

In the next six months Service Challenge 4 would be crucial. They needed a stable environment for experiments and this should include the Tier-2 centres. This would represent the final test phase for sites to validate their operation.

Discussion

J. Engelen thanked L. Robertson for his presentation. In inviting questions he requested delegates to take into account as well the LHCC Deliberations, which had been provided in a paper by the scientific secretary of the LHCC, E. Tsesmelis. In addition he asked any of the LCG users present also to ask questions.

J. Virdee noted that it was very important for CMS that LCG kept to its planned milestone of entering operation state at the end of SC4. CMS was counting on starting its CSA at around that time. L. Robertson replied that this was well understood, that dates were now the important thing, and that functionality had become secondary to providing a service with whatever could be made available.

D.O. Riska noted that the impression had been given that the whole LCG system was up and running, yet they had heard that there was a shortfall of as much as 50% in some areas. F. Carminati wished to emphasize that, of the four experiments, it was ALICE that was particularly short of resources. L. Robertson replied that he had described how they were setting up the service. He felt that the task of putting capacity into that service was rather a concern for

this committee. J. Engelen thanked D.O. Riska for this very relevant question and pointed out that they would return to the issue after the presentation by the Resource Coordinator.

F. Linde was concerned by the requests for extra manpower made by the collaborations, and he urged LCG and the experiments to keep this to a minimum, bearing in mind that some Tier-1 sites were part of a National GRID, which also served other communities. He had seen that this was an issue in more countries than just the Netherlands. L. Robertson said this was indeed a concern and that they had been working hard to find a reasonable compromise. Different experiments had different needs. The most recent compromise as to how to proceed this year was agreed a month ago. He felt that this could become a bigger issue in the longer term. As the services built up, the people needed to run these services would become overloaded with the basic day to day operation, and would not be able to deal with several different varieties. He thought this was understood by everybody, but it was also important that the experiments received the services they needed. It was a difficult area and one to which they were paying quite a lot of attention. J. Engelen noted that this had been discussed by the WLCG Overview Board and that the practical attitude taken was that, for a Tier-1 centre, in the end, there should be no experiment-specific tasks. However at the beginning one needed some flexibility. L Robertson felt that they also needed some experience as to how much effort was needed to run these services. They had to find compromises here that would enable them to move forward.

4. LHCC Deliberations — E. Tsesmelis, LHCC Scientific Secretary

Paper CERN-RRB-2006-015

Delegates had no further comments to make and the RRB took note of the report of E. Tsemelis.

5. Status of Common Project Accounts — P. Geeraert, Head, CERN Finance Dept.

Paper CERN- RRB-2006-016

P. Geeraert noted that he would be brief. In Phase 1, expected contributions amounted to 21.432 MCHF and they had received so far 20.949 MCHF.

He summarized the commitments so far in Phase 2 which amounted to 15.614 MCHF out of a budget of 45 MCHF.

Discussion

U. Dosselli commented that there was a negative number reported in the table for Italy for Phase 1 and he expected this to be paid very soon. J. Engelen noted that this would leave the Phase 1 balance as very close to zero. They would return to LCG Phase 2 after the next talk.

There being no comments on these numbers, the RRB **took note** of this financial report.

6. Status of Resources and Financial Plan — C. Eck, Resources Coordinator

Paper CERN- RRB-2006-017 Presentation CERN- RRB-2006-062

C. Eck began by noting that the details of this presentation could be found in the written report (CERN-RRB-2006-017), in the WLCG MoU (CERN-C-RRB-2005-001/Rev.), and on the LCG planning web pages: http://lcg.web.cern.ch/LCG/planning/planning.html, under Current WLCG MoU Documents.

6.1 Concluding remarks on LCG Phase 1

The voluntary contributions to personnel and materials during Phase 1 reached exactly the amounts predicted in October 2005. It was thus not necessary to present the tables again, but he wished to take the opportunity to thank once more all the contributors.

The LCG materials expenditure at CERN in 2005 appeared to be 1.8 MCHF less than that presented last October. This was largely an accounting difference. Networking equipment worth 1.4 MCHF, originally shown under Phase 1, had now been accounted under Phase 2, which corresponded more correctly with its usage. Air conditioning equipment for 0.4 MCHF could not yet be billed in 2005. In this revised accounting the 0.5 MCHF balance would grow to 0.8 MCHF when the final contribution was received. He emphasized that most of this positive balance was just the result of shifting payments into Phase 2 and not a real saving.

6.2 Regional Centre Resources

Since the previous RRB, Portugal had entered the MoU tables as a Tier-2 centre for ATLAS and CMS. The Finnish contribution to CMS was now shown as a new Tier-2 and not as part of the NDGF Tier-1. Multiple small changes had led to many versions of the MoU tables from October until now, but the sum of Tier-1 and Tier-2 capacities had changed by less than 2‰. Did this represent a wonderful level of stability or a terrible lack of growth?

He showed the external Tier-1 and Tier-2 Summary Table, giving the pledged capacity, the TDR requirements and the balance for each of the three resources, CPU, disk and tape. Clearly some experiments were better positioned than others. Forty Tier-2 centres had contributed to the data in these tables, and 9 more Tier-2 centres planned to join as soon as possible. C. Eck showed a table of these additional Tier-2 centres and explained in some detail the current status (see slide 7 of the presentation).

Discussion

J. Engelen suggested at this point that they should pause and discuss the consequences of these current tables, and consider which dynamics could lead to changing these numbers. He noted that the situation presented was correctly that for 2008, not the start up year, and hence the requirements were not going to be reduced by less running. The requirements numbers were based on models, but they had been reviewed by a specialist committee of the LHCC, so one had to conclude that the overall picture of the requirements was not going to change dramatically. Therefore one had to foresee that the pledges needed to be changed. He asked the experiments whether they considered the requirements to have changed.

F. Carminati, on behalf of ALICE, considered that these stated requirements were still valid and whilst he did not exclude some fine tuning, he feared this could only go in the wrong direction. Clearly ALICE had a problem.

D. Barberis stated that for ATLAS the requirements were still the same because they were based on two or three years of discussion of their computing model and on the common assumptions on the periods of running on 2007, 2008 and further on. Hence the only variable that could be changed would be to reduce the running time which it was agreed would not be an attractive option.

J. Virdee for CMS echoed the previous two speakers as far as requirements were concerned. He added that several of their Tier-1 centres, especially in Europe, were sub-nominal, especially for tape, and they were discussing how to rectify this. M. Della Negra noted that the requirements were based on a number of assumptions that were somewhat arbitrary, and he felt that it would be much better to replace the word “requirement” with “request”. The experiment could request as much computing as they could afford. The only limitation was that they could not write too much data if they could not process them off-line. So they were cutting on physics in order to stay reasonable with the number of useful events they could write. However, there was no absolutely well defined threshold on how much data they could write. In the end they were limited by resources, by the money to do the computing. The question was not so much how much they needed but how they would share the available resources. This was what was on the table for this committee and other relevant committees. They needed to have a scheme in order to share the

available resources such that all the experiments were competitive and could do their physics programme.

T. Nakada for LHCb felt that they had already adjusted their demands to reality such that they therefore seemed to be reasonable in their request. They believed that some of their shortfall could be covered by Tier-3 centres.

J. Engelen noted that the comments of M. Della Negra required a prepared discussion, perhaps at a later time, when they had maximized the resources and understood a little better the consequences. He proposed to continue the present discussion on the basis of the currently presented numbers.

L. Levinson pointed out that for the Tier-2 centres, a major fraction of the computing resources was taken by simulation. For ATLAS the simulation rate in the model was 20%, which was to say that 20% of the number of events recorded was considered to be the number of events that would have to be simulated. There had been comments that this was a low number. If this were to move to 30% or 40%, then there would be a serious imbalance.

U. Dosselli explained that Italy was participating in major experiments and the numbers were correspondingly large. Nonetheless they had committed to credible numbers that they could guarantee. Perhaps these were the minimum, and that they would strive to go beyond that, but they could not commit to more at this stage. They had committed to major technical works in the Tier 1 centre in order to achieve the desired configuration and this was also expensive. Time was a critical element was in the planning, and purchasing just in time.

F. Linde noted that their absolute numbers for the coming four years were fixed, and they had tried to favour the experiments in which they had, percentage-wise, a large component: namely LHCb, and then ATLAS and then ALICE. Nonetheless, the needs of ATLAS and ALICE were big and hence, in absolute terms, they had much more than LHCb.

C. Eck asked about the Nordic Data Grid Facility. Their numbers in the table were the oldest and he wondered if there had been any changes. J. Dines Hansen replied that the Danish numbers were increasing. C. Eck noted that he had heard that this was true in Sweden as well, and that therefore there was some hope of improvement.

T. Ferbel had heard that the LHCC was looking into discrepancies in the requests from ATLAS and CMS. He wondered if there were any developments here. J. Engelen replied that there were no results as yet and that the analysis was not a simple one. There would be a study of the computing models in more detail than had been done so far. The results would be public but had not yet been finalized.

J. Engelen summarized that they would **take note** of the tables and the discussion. The tables were not old and it was probable that an update of the numbers would produce a picture that was qualitatively the same. At a certain moment the resources that were there would have to do. Therefore, there was no reason to believe that, in a resource sense, all problems had been solved, and delegates should leave this meeting concerned rather than relieved.

6.3 Progress in Signing the WLCG MoU

C. Eck returned to his presentation. The number of signatures to be collected from Funding Agencies with pledges in the current MoU was 30. The number of signed MoUs so far received was 18, including Italy who were currently in the process of signing. Therefore they had, at this time, signatures for 80% of the pledged capacities in 2008. There seemed to be no indication that it would not be possible to collect the remaining signatures before October 2006.

6.4 LCG Phase 2 Budget at CERN

The table on slide 10 showed the cost and funding estimates for LCG Phase 2 at CERN. With respect to the table shown at the previous RRB, funds had been re-profiled over the years to fit better to the needs. The personnel budget now showed fully its planned balance. The overall balance had improved by 780 kCHF. Of this 500 kCHF had been received from cuts in the IT Department budget, and PH Department had been able to pay for their additional LCG personnel from its 2005 budget. Nonetheless, he felt that the situation was far from satisfactory.

The materials budget shortfall was essentially 13.9 MCHF. The Infrastructure and Tier-0 spending corresponded necessarily and fully to the original profile. The CERN Analysis facility, CAF, was budgeted to provide 80% of the requested capacity, but lacked 83% of its funding.

Discussion

J. Engelen wished to remind Funding Agencies that the structure of LCG in two phases was agreed some time ago, when it was already clear that CERN could not provide everything. The result of Phase 1 was very successful, with major contributions from a number of Funding Agencies. However the situation with Phase 2 was not the same. He wished to make it quite clear that several Funding Agencies, notably INFN, had lived up to what was expected from them, and they were not at fault for this current situation. Nonetheless there was a big negative number and this needed discussion.

The decision had been made deliberately, more than a year ago, to move the CERN problem to materials, because it was clear that momentum in LCG should not be lost by allowing key personnel to leave. They had tried, by various gymnastics, to make the -14 MCHF smaller, and they could see how perhaps to achieve something like 3 MCHF, but not 14 MCHF. In the end, if this persisted, they would have to start to reduce performance, in a manner as responsibly as possible. The Tier-0 facility, on which the whole GRID was based, had been carefully protected in all of this, and he thought this was clearly the right approach. They were doing their best to rearrange matters internally and minimize the impact of the shortfall. Nonetheless it was clear that the concerns of the Funding Agencies were focussed on the capacity at home. He invited delegates to make their comments.

F-D. Buzatu asked why he had not heard that Romania had joined the LCG, since this had now been signed. C. Eck replied that the inclusion of Romania as a full Tier-2 member had been announced and agreed in the previous RRB, and that clearly this was seen as a very positive move.

S. Tovey had some concern about the numbers. Australia had a large GRID effort in terms of CPU and manpower, but it was not dedicated to high energy physics, it was also being used by the astrophysicists and other groups. Even within HEP it was being split between ATLAS and the Belle experiment in Japan. It was very difficult to provide hard numbers as to what was going to CERN, and he could imagine that other countries had the same problem.

R. Wade noted that this gap had been known for quite some time. Although they had sat round this table in this rather full room, and had several times made requests for people to provide more resources, the gap did not seem to have closed. It seemed to him as a result of this that the basis for planning now should be that it was not going to close before the start up of operations.

As he had said in the CMS RRB yesterday, there were a number of gaps in resources needed, both in the experiments for the cost to completion and for some particular deliverables, as well as in the computing. He was concerned that, whilst this was a Resources Review Board which could look at the resources available, managing the way to the successful operation of the LHC required something else. He did not have a particular proposal to make, although he was trying to find one in order to be helpful. Perhaps the way to a solution was not anymore through these RRBs but through some wider management action. He did not know which was the most suitable forum was

for this. He thought they needed to look at resources across the board for the LHC and not at individual areas. They needed to understand, in the timeliest manner they possibly could, how to get the most physics out of the whole endeavour. He had doubts that this particular forum would be able to solve the problem.

A, Naudi noted that there were a number of delegates around the table from Finance Committee and from Council, who knew very well the financial situation of CERN. There was a very difficult financial situation because CERN was in deficit. He could hardly see any way of producing more funds from the CERN budget without going further into deficit.

D.O. Riska, as a Council member, expressed the view that he had heard, circulating around the Council, that this particular deficit ought to be covered from the CERN budget. The reason for this view was that it was also seen by these people that it might be possible to cover the deficit within the experiments by having an additional RRB specifically for each experiment. This was the way that agreements to pay for the M&O funds for the experiments were brought to successful conclusion. The CERN Director responsible at that time, R. Cashmore, called an extraordinary RRB in September 2001. This was, in his memory, an unpleasant meeting, and rather brutal demands and comments were made, but in the end it led to the successful agreement on a level of M&O funds. He believed that this was the method by which the deficits in the experiments might be hammered out. It would fall upon the Chairman to call such meetings, and clearly they would not be pleasant. However, he believed that this particular problem of the LCG had to be solved through the CERN budget. J. Engelen thanked him for this remark with its proposal of how some of the missing money might be agreed.

R. Wade, also as a Council member, added that, to his knowledge, this level of deficit had not been discussed in Council. Furthermore this concerned not just the Member States but also the non-Member States. There was a need for a forum and a mechanism for this situation to be discussed with both sets.

J. Engelen noted that this situation had been presented quantitatively to Council and Finance Committee on more than one occasion, starting already more than a year ago. A. Naudi confirmed this statement. R. Aymar added that it was not appropriate, at first sight, to put this computing shortfall into the deficit, thus making the assumption that the budget would pay for it, before there was a real appreciation as to what could be made as a global effort by everyone. J. Engelen noted that this problem was owned collectively by all members of the LCG Collaboration, Member States and non-Member States alike, and both were represented around this table.

7. Summary, Future Activities & A. O. B. J. Engelen

J. Engelen noted that, despite the above situation which they had high-lighted, one should not forget that very large technical progress was being made and that a large GRID was coming to life. On the other hand, in order to be optimally prepared to use the LCG, there were a few outstanding issues, not necessarily of a technical nature, which needed to be dealt with fairly urgently because one was already in 2006.

The next RRB meetings in 2006 will take place at CERN on
Monday 23rd, Tuesday 24th and Wednesday 25th October 2006

There being no questions and no further business, the Chairman thanked the participants and closed the meeting.

C. Jones
June 2006