



Computing Resources Review Board

10th November 2007

Minutes of the 12th Resources Review Board Meeting Held at CERN on 23rd October 2007

Present:

Europe:

C. Wulz (bm:bwk, Austria);
J. Lemonne (FWO, Belgium); G. Wilquet (FNRS, Belgium);
M. Lokajicek (MSMT CR, Czech Rep.), J. Niederle, M. Sumbera;
J. Dines Hansen (National Science Research Council, Denmark);
M. Kadastik (NICPB, Estonia) (*Observer*);
D.O. Riska (HIP, Finland), J. Rak, T. Tuominiemi;
J. Zinn-Justin, J.-P. Meyer (CEA/DSM/DAPNIA, France);
F. Le Diberder (CNRS/IN2P3, France);
K. Ehret (DESY, Germany), M. Fleischer; H. Marten (FZK, Germany); P. Malzacher (GSI, Germany), D. Müller;
S. Bethke (MPG, Germany); T. Hebbeker;
G. Vesztergombi (KFKI/RMKI, Hungary) (*Observer*);
E. Rabinovici (Hebrew University Jerusalem, Israel), L. Levinson (*Observer*);
F. Ferroni (INFN, Italy);
F. Linde (NIKHEF, The Netherlands), A. J. Van Rijn;
B. Jacobsen (Norwegian Research Council, Norway), F. Ould-Saada;
M. Turala (Ministry of Science & Education, Poland), G. Polok;
G. Barreira (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. Sissakian (JINR, Dubna);
J. Fuster (MEC, Spain), N. Colino;
T. Ekelof (Research Council, Sweden), O. Thylander;
A. Rubbia (SER/SNF/ETH/CSCS, Switzerland);
G. Zinovjev (National Academy of Sciences of Ukraine);
R. Wade (STFC, United Kingdom), J. Seed, G. Lafferty.

N. America:

W. Davidson (CFI, Canada), R. McPherson;
M. Pripstein (NSF, U.S.A.), J. Whitmore;
S. Gonzalez (DOE, U.S.A.), T. Ferbel, E. Henry, G. Rai;
L. Bauerdick, J. Butler; H. Gordon, J. Shank, M. Tuts; T. Cormier, J. Symons.

Asia:

P. Ji (MoST/NSFC, China), Y. Zhang;
P.S. Dhekne (DAE, India), A. Gurtu;
K. Saito (University of Tokyo, Japan), T. Kawamoto;
K. Song (KICOS, Korea), J. Choi, Y.-I. Choi, S. Hwang, D.-W. Kim (*Observer*);
S.C. Lin (Academia Sinica, Taipei), C. Chien-Hong.

Australia:

S. Tovey (AusHEP, Australia).

CERN:

J. Engelen (Chairman), C. Jones (Secretary), S. Lettow,
J.-J. Blaising, P. Geeraert, D. Jacobs, B. Salami, E. Tsesmelis, E. Van Hove, W. Von Rueden, J. Salicio Diez;

LCG:

L. Robertson, C. Eck, S. Foffano;

ALICE:

J. Schukraft, Y. Schutz;

ATLAS:

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

CMS:

J. Virdee, M. Kasemann;

LHCb:

T. Nakada.

1. Introduction

J. Engelen, Chief Scientific Officer

J. Engelen welcomed RRB delegates to this 12th meeting of the Computing Resources Review Board.

2. Approval of the Minutes of the 11th Meeting (CERN-RRB-2006-006)

The minutes of the 11th meeting, CERN-C-RRB-2007-006, 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-2007-101

Presentation CERN-RRB-2007-102

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:

- Applications support was in good shape
- For the WLCG service:
 - Baseline services were in production with the exception of SRM 2.2
 - Continuously increasing capacity and workload
 - General site reliability was improving
 - Data and storage remained the weak points
- Experiment testing was progressing –
 - involving most sites, approaching full dress rehearsals
- Sites & experiments were working well together to tackle the problems
- Major Combined Computing Readiness Challenge next year before the machine starts
- There was a steep ramp-up ahead to delivering the capacity needed for 2008 run

Discussion

J. Engelen thanked L. Robertson for his presentation. In inviting questions he requested delegates to take into account the LHCC Deliberations, which had been provided in a paper by the scientific secretary of the LHCC, E. Tsesmelis.

S. Tovey asked, given that the best 50% of the Tier-2s were performing as expected, whether the worst 50% were performing rather badly. L. Robertson reported that they had not been reporting these numbers back to the sites. These sites were being tested because they were part of the EGEE grid, and the sites might not know they were being tested. In fact the tests were rather harsh, and all of them had to work, otherwise the site was considered to be down. As from October they would be sending the results of these tests back to the sites.

M. Turala wished to express his appreciation to L. Robertson for having brought this project into a working state, given that ten years ago the whole issue of LHC computing had seemed an enormous challenge. However he would like ask experiment coordinators to comment on the status. Did they feel secure that in a year from now they would have the tools to analyse the data?

D. Barberis for ATLAS explained that a few years ago they had been very ambitious and wanted to build a very special system. In the meantime they had become more realistic. They could see that progress was being made in the essential features. They hoped that, with the last few features that would come together during this winter and spring, that they would be able to use the system, albeit perhaps not with the very high expectations that they had in 2000/2001. The system was usable and was being used for production, data distribution and for many tests. There were people who still complained quite rightly that not everything worked perfectly and they were trying to improve the system. It was at least usable.

M. Kasemann for CMS noted that they knew that they had a year of hard work in front of them. As soon as they tested the system they saw that many things worked very well, but they were not yet where they needed to be next year. The tests and challenges were very important to them. In particular they needed all the experiments to test simultaneously and this they had to do before data taking, otherwise they could not be confident that they could use the system. The basic services would allow them to analyse the data, but not yet in a way that every scientist would have easy access to the data. He considered that there had been good progress but certain ways needed to be improved in order that the scientist had easy access to their data. This needed work both by the experiment and also on the grid services side. They had notably to work on site reliability.

Y. Schulz for ALICE thought that taking the pragmatic approach had been useful in order to get an operational system. They were now in a position to have commonly running production and reconstruction using the whole infrastructure, with more and more sites as they came up. The bottleneck for the moment was the analysis and getting access to distributed data. The common data challenge of all four experiments together was very important.

T. Nakada reported that most of the technical issues had been sorted out, apart from the SRM which had to be made reliable. Things were now moving into a more operational phase. For example, upgrades needed to be propagated quicker, and similarly new software. Their observation was that things needed attention on the operational side.

J. Engelen thanked M. Turala for this question and the experiment representatives for their comments. He asked the delegates for their reactions to the not-unexpected large ramp-up that was needed up to say April 2008. Even that capacity was only sufficient for 2008 which was not yet a full running year. Was this consistent with the plans of delegates? Was the silence of the delegates a sign that everything was in order?

D.O. Riska replied that silence meant that everyone around the table knew that there were resource limitations back home and it was very hard to say that they could promise exactly suggested the ramp-up, for example the factor 3.7 needed for CPU. Maybe there would be different numbers. J. Engelen noted that might not be satisfactory.

G. Barreira noted that Portugal was planning to ramp-up by a factor of four by the end of the year and a factor of seven by the target of April 2008.

Y. Schultz wished to add to his comments above that ALICE had a chronic problem in obtaining all of the resources required. They were slowly closing the gap but the problem remained. They had to work hard with the delegates to get close to the required resources.

T. Ferbel wondered whether Moore's law was a reason for people holding off with the acquisition of capacity. L. Roberston agreed this was indeed a factor but that they were coming to a period close to the accelerator operation and there was this important test of all experiments simultaneously to be conducted before that moment. There had been problems with capacity being installed late for this reason in the past, and he thought that at this moment it would be better to err on the side of making sure the capacity was installed. They hoped that the LHC machine would start well and the first months of operation would be an enormous challenge. W. Von Rueden added that when they saw the first really big wave on installations several computer centres got into trouble with electricity and cooling. He thought that this could well happen again, and such things took relatively long to correct. Some problems were only discovered when everything was installed. J. Engelen thanked him for this warning. G. Barreira considered that these were indeed critical parameters that were sometimes insufficiently considered and he felt that some effort should be spent asking the collaborators about their plans.

M. Turala noted that this capacity issue did not just refer to new resources but also to getting the required capacity from existing capacity. He wanted to pass the message to the collaborations to make an effort to work closely with their computer centres and to press on them to get the resources they wanted.

4. LHCC Deliberations

Paper CERN-RRB-2007-109

E. Tsesmelis, LHCC Scientific Secretary

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

5. Status of Common Project Accounts

Paper CERN-RRB-2007-107

B. Salami, CERN Finance Dept.

Presentation CERN-RRB-2007-108

B. Salami noted that he would be brief. The CERN expenditure for the three years 2005-2007 stood, at the end of February 2007, at 30 MCHF for personnel and 33 MCHF for materials, making a total of 63 MCHF.

Concerning the accounts of the Funding Agencies, the report presented the accounts where there was a flow of money (Germany, Italy, Israel and India). The balance on these accounts was positive by 200 kCHF.

There being no comments on these numbers, the RRB **took note** of this financial report. J. Engelen thanked P. Geeraert for his report.

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

Paper CERN-RRB-2007-103

Presentation CERN-RRB-2007-104

S. Foffano, who was taking over from C. Eck as Resource Coordinator, began by noting that the details of this presentation could be found in the written report (CERN-RRB-2007-103) and on the LCG planning web pages.

6.1 Signatures of the WLCG MoU

S. Foffano reviewed the status of signatures of the WLCG MoU. Since April 2007 the following federations had signed:

- Spain (T1+T2)
- Switzerland (T2)
- Australia (T2)
- Israel (T2)
- JINR, Dubna (T2)
- Russia (T2)
- Slovenia (T2)
- Germany -ATLAS Federation FR/W (T2)

Of the 11 Tier-1 Centres, all had signed except the Nordic Data Grid Facility (NDGF) which was still to be signed by Finland, Norway and Sweden, Denmark having already signed.

Signatures were expected soon for the following Tier-2s:

- Austria –Institute for HEP, Vienna & University of Innsbruck
- Germany –Ludwig Maximilian Universität (LMU) for the ATLAS Federation, Munich
- Germany –Albert-Ludwigs-Universität (ALU), Freiburg for the new ATLAS Federation FR/W (Universities of Freiburg & Wuppertal)

Following information from the WLCG Collaboration Board member, the signature from the Czech Republic was pending budget approval. It was currently unclear when this would be received.

Slide 5 showed a list of additional Tier-2 centres planning on joining the WLCG collaboration. S Foffano then noted some very recent news:

- A new Tier-2 had been added on 12/10/07: France, LAPP, Annecy serving ATLAS and LHCb with priority
- 2 MoU signatures were planned to take place that day:
 - Korea for an ALICE Tier-2
 - Finland for the NDGF Tier-1 and their CMS Tier-2
- 4 countries had provided pledge information and requested preparation of the MoU for signature:
 - Norway on 11/10/07 for the NDGF Tier-1 and their ATLAS Tier-2
 - Hungary on 16/10/07 for their ALICE and CMS Tier-2
 - Canada on 17/10/07 for their East and West ATLAS Tier-2s
 - Estonia on 18/10/07 for their CMS Tier-2
- This increased the current number of Tier2 Federations from 53 (as documented in CERN-RRB-2007-103) to 55 by the day of the RRB and 59 in the near future.

The conclusion was that the WLCG continued to grow.

6.2 LCG Phase 2 Budget at CERN

S. Foffano then showed in slides 7 & 8 the CERN funding and planned expenditure for 2005-2008 and 2009-2012 of the LCG project at CERN in MCHF. It should be noted that the LCG personnel planning at CERN assumed that a successor project to EGEE2 would deliver 14 FTE to Grid Deployment activities from April 2008-March 2010. The future estimated materials expenditure included the latest requirements data from the experiments up to 2012, and an allowance for adapting the CERN infrastructure to the increased power and cooling requirements.

6.3 Resource usage and accounting

S. Foffano then showed graphs giving a summary of the CPU Time, Disk and Tape Storage accounting from September 2006 to August 2007 obtained by summing up all external Tier-1s and CERN. The graphs showed the installed capacity, the pledged capacity and the resources usage. In particular for disk storage used there was a large gap which was a source of concern between the commitment and the installed capacity. It had become urgent that 2007 commitments were fully honoured and that 2008 commitments were installed and fully in production by April 2008

S, Foffano turned to accounting for Tier-2s. In June an information campaign had started amongst the Tier-2 sites to inform them that formal reporting of CPU usage would begin from September 2007 onwards on a monthly basis. Test reporting took place in July and August. Of the 53 Federations asked to identify sites to participate in the reporting, 5 had not yet provided this information, namely Finland, Germany (ATLAS Federation, Munich), Norway, Sweden and Ukraine. On the other hand 48 Federations had provided the information, with 108 sites identified, of which 94 reported accounting data in September.

The first formal report was produced mid October for September data and was available on the LCG web site giving details of the breakdown by Federation. Globally 48% of the pledged CPU was delivered in September, with many sites far from their pledge value. Work was still required to identify remaining sites, and ensure that all sites reported data. Federations which had not yet identified sites were urgently requested to do so. A more complete report for the first 6 month period would be given in the April C-RRB meeting.

6.4 Revised computing capacity pledges

The 4 LHC experiments were asked to revise and update their computing requirements for 2008-2012 in July 2007. The information was received from ATLAS in July, and ALICE, CMS and LHCb later in the summer. All Tier-1 and Tier-2 federations were asked to confirm their pledge values for 2008 and provide values for 2009-2012 inclusive in accordance with the MoU. A majority of federations had not yet responded:

- Tier-1s: 5/11 had responded
- Tier-2s: 24/54 had responded

From the data received so far, particularly from Tier-2s, many pledge values for 2011 and 2012 remain as for 2010. This was a major cause for concern, as if this trend continued there would be major discrepancies between rising requirements and pledges which did not rise as expected.

Slides 15 and 16 showed the revised capacity pledge situation as it was known at the present time. The overall picture was described as very pessimistic. Therefore it became urgent that all remaining federations should reply as soon as possible in order to confirm or update their pledge for 2008 and to provide their planned pledge for 2009-2012. Federations should plan to renew their computing resources in the future to enable increased pledges to be guaranteed. Slides 18 and 19 showed graphically the same information and what would happen if nothing changed. S. Foffano emphasized that the participation and commitment of the delegates were essential to improve these trends.

In conclusion S. Foffano noted that:

- MoU: The signing of the 2 remaining NDGF countries (Norway and Sweden), Austria, Germany (LMU and ALU) were expected very soon.
- Conclusion and signing of the MoU for the additional announced T2 centres before the next C-RRB meeting.
- MoU pledge levels for 2007 must be fully honoured and for 2008 installed and in production by April 2008 to ensure sufficient time for full-scale testing.
- Tier-2 Federations must ensure that their sites are identified and reporting data for accounting purposes, and their pledge level is reached.
- Tier-1 and Tier-2 Federations MUST confirm their 2008 pledges and provide pledge values from 2009-2012:

Discussion

J. Engelen thanked S. Foffano for this clear presentation. He asked whether there were any comments or questions.

S. Tovey reported that they had guaranteed funding until 2010 but that they could not make a pledge beyond that date as it would have no value.

J. Seed wished to raise a process issue. Assuming all the numbers in the tables were correct, then one could be facing a situation where some experiments had more capacity pledged that was required and some less. What was the process for looking at this at a top level and deciding whether there was a need to move computing resources between the experiments or for optimising the use of those resources that were available? J. Engelen replied that, personally, he would hope that this would be considered by the Resources Scrutiny Group that would be discussed later on in this agenda, and that they would advise how to deal with such a situation. L Robertson noted that it was well understood that funding was not secured until the point that the funding process was complete. However all computer centres were now facing problems of power and energy and these took time to resolve. A long term was essential. J. Engelen thought that the Overview Board should start to consider such a long term view.

T. Nakada wished to comment that LHCb seemed reasonably well provided for in the table. He explained that the 2007 requirements numbers had been revised downwards to reflect the short run, and this did not mean that they would be well-off in 2009.

G. Barreira commented that in Portugal they had found such long term planning was very hard. Computing was quite different from the detectors. The MoUs provided a long-term foundation for the Funding Agencies to address the funding of the detectors and similarly for the M&O. They had been unable to do the same for computing where things were bound up with the national computing initiatives. For the detectors they had some form of quantification related mainly to the number of people involved, the PhDs that signed the documents. The same matrix did not exist for the computing. At a certain point the Funding Agencies could ask whether they were contributing more than their fair share. This was quite an issue and he considered that WLCG had to face it. Probably they had to find some matrix to represent the fair contributions. Without this matrix and without a more stable commitment from the Funding Agencies for the long term it was really very difficult to provide numbers beyond 2010. He knew that these commitments were necessary but he was simply not in a position to provide the numbers. J. Engelen thanked him for these comments and agreed that they would work on trying to improve the situation and to provide such a matrix. A quite general reaction for the sustained support of these projects would be that the involvement in national computing grid initiatives beyond our own field brought some more stability and security, as seen in several countries, and he believed that to be a good development for the WLCG.

W. Von Rueden brought some words of warning about the expectations on national grid initiatives. He was sure that these initiatives would help for the general operation and infrastructure, and to ensure the connectivity between the different countries. However the capacity needed in terms of storage and disks for particular disciplines was usually very different. This might not be so true if countries tried globally to make available centres that were shared by all sciences. In most cases he believed this was not the case. Hence one should not count on someone else to solve the problem.

7. Establishment of a Resources Scrutiny Group

J. Engelen

J. Engelen noted that a Resources Scrutiny Group had existed for a long time within the context of the experiment RRBs whose role was in particular to scrutinize the M&O budgets. This was written also into the MoUs for the computing grid, namely that it would be useful to have a similar resources scrutiny group looking at the computing needs felt by the experiments and at the way in which they should be satisfied. They had written, in accordance with what was agreed, to the Tier-1 countries to ask them to delegate a person for this scrutiny group. They had received replies and they were going to nominate a person from each of Denmark, USA, the Netherlands, Spain, Italy, Germany, France and Canada. This meant that they still needed delegates from the UK and from Taipei, and there should also be a representative from CERN. It was hoped that this body could meet already in 2007, not to start scrutinizing, but in order to start defining a process, to defining what types of input were required. He hoped this discussion could take place in interaction both with the members of the project and also with the experiments in order to see how the resources could be scrutinized in a useful manner. He would report at the next RRB on the first results and also on a more precise mandate and working procedures.

J. Seed noted that a UK representative would be named soon. She asked whether there were some terms of reference. D. Jacobs replied that they were written in the MoU and they had not changed. J. Engelen wished to invite this new scrutiny group itself to consider whether they considered this mandate to be adequate or whether they would like some modifications which he would then offer to this board for approval.

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

J. Engelen noted that this was the last RRB for L. Robertson as the project leader. A new project leader had been selected and appointed to take office on the first of January 2008, namely I. Bird. L. Robertson would not be leaving immediately but would be around to offer his support for a further six months distributed over the twelve months of next year. They would organize a proper ceremony at the appropriate time but he asked the meeting to join in a round of applause for L. Robertson as a prelude.

J. Engelen noted that one had an enormous and dynamic project with a very large positive derivative. One had to hope that it would be large enough when real data arrived. It was impressive how the project had evolved technically and organisationally over the years and he felt that high energy physics had shown the way how computing was to be done in the future. He thought one could look positively towards the culmination of the project into a service sometime at the beginning of 2008.

<p>The next RRB meetings in 2008 are provisionally scheduled to take place at CERN on Monday 14th, Tuesday 15th and Wednesday 16th April 2008 and on Monday 10th, Tuesday 11th and Wednesday 12th November 2008</p>

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

C. Jones
November 2007