



Computing Resources Review Board

28th April 2009

Minutes of the 15th Resources Review Board Meeting Held at CERN on 28th April 2009

Present:

Europe:

C.-E. Wulz (HEPHY, Austria);
 J. Lemonne (FWO, Belgium); J. Sacton (FNRS, Belgium);
 D. Adamova, M. Sumbera, J. Niederle (Institute of Physics AS CR, Czech Republic);
 J.D. Hansen (Niels Bohr Institute, Denmark);
 D-O. Riska, J. Tuominiemi (Helsinki Institute of Physics, Finland);
 E. Auge (CNRS/IN2P3, France); P. Chomaz (IRFU, France); F. Malek (CNRS/LCG-France, France);
 S. Bethke (MPI, Germany); K. Ehret, M. Fleischer; T. Hebbeker (RWTH, Germany); K-P. Mickel
 (Karlsruhe Inst. of Technology, Germany); M. Pantea (Ministry of Education & Research, Germany);
 T. Csörgo (MTA KFKI-RIMKI); G. Vesztergombi (KFKI-RMKI, Hungary);
 E. Rabinovici (Hebrew University, Israel);
 F. Ferroni, G. Tonelli, U. Marconi, A. Vacchi (INFN, Italy);
 B. Jacobsen (The Research Council of Norway);
 M. Turala (IFJ PAN, Poland); P. Grzegorz (Polish Academy of Sciences);
 G. Barreira (LIP, Portugal);
 F-D. Buzatu (Institute of Atomic Physics, Romania);
 V. Savrin (Institute of Nuclear Physics, Russia);
 N. Colino, P. Ladron de Guevara (CIEMAT, Spain); J. Fuster (MICINN, Spain); E. Higon (IFIC, Spain)
 T. Ekelöf (Uppsala University, Sweden); P. Karlsson (Swedish Research Council);
 I. Turk Cakir (TAEA, Turkey);
 G. Zinovjev (National Academy of Sciences of Ukraine);
 J. Seed (STFC, United Kingdom).

N. America:

W. Davidson (NRC, Canada); R. Mcpherson (University of Victoria, Canada);
 L. Bauerdick, J. Butler (FNAL, U.S.A.); S. Gonzalez, H. Marsiske (DOE, U.S.A.); A. Lankford (UCI,
 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:

Y-I. Choi (Sungkyunkwan University, Republic Of Korea); J. E. Ha (Minister of Education, Science &
 Technology, Republic Of Korea); K. H. Choi, 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. Lee (ASGC, Taiwan);
 G. Taylor, University of Melbourne, Australia.

CERN:

S. Bertolucci (Chairman), H. Renshall (Scientific Secretary), F. Hemmer, B. Loehr (Scrutiny Group
 Chairman), D. Espriu (Computing Scrutiny Group), P. Bloch, G. Cosmo, E. Van Hove, J. Salicio Diez,
 D. Jacobs, T. Lagrange, S. Lettow, R. Heuer.

LCG:

I. Bird, S. Foffano;

ALICE:

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

ATLAS:

D. Barberis, F. Gianotti, R. Jones;

CMS:

M. Kasemann, P. McBride, J. Virdee;

LHCb:

A. Golutvin, A. Schopper.

1. Introduction

S.Bertolucci, Director for Research and Computing

S.Bertolucci, as chairman, started the meeting wishing the delegates a good morning and adding that he expected, as always for computing, a lively discussion.

2. Approval of the Minutes of the 14th Meeting (CERN-RRB-2009-007)

S.Bertolucci then asked delegates if they had any amendments to the minutes of the last meeting, CERN-RRB-2009-007. None were forthcoming so he considered the minutes as approved.

3. Status of the LCG Project

I.Bird, LCG Project Leader

Paper CERN-RRB-2009-039

Presentation CERN-RRB-2009-040

S.Bertolucci proposed to move to the report on the status of the LCG project.

I.Bird began with his agenda which would cover a general overview of the WLCG service status, the Roadmap for the Scale Testing for the Experimental Program 2009 (STEP'09), the planning for 2010 and finally the re-assessment of experiment requirements in the light of the LHC schedule and the experience gained last year during the common computing readiness challenge and by taking cosmics data.

His slides 3 and 4 showed the usual site reliability plots and the increasing numbers of jobs run per month showing that the ramp up of the load continues and that the service is already running at the level expected for data taking this year. The average Tier 1 reliabilities seem good but the individual site plots show that month by month individual sites have many problems. Slide 5 of the table of the reliability of Tier 2 federations shows an encouraging improvement since the last RRB and only one federation is not yet reporting. Slide 6 showed the time variation over the Tier 1 sites of experiment specific tests which are now run in addition to those of WLCG and are giving a consistent picture. These tests are in the process of being validated but are already used by experiments to decide whether to run work at a site.

He then moved on in slide 7 to comparing pledged, delivered and used cpu and disk resources at Tier 1 sites commenting that at the end of 2008 most pledges had been installed but individual site plots showed there are many problems to ramp up resources, taking up to a year for some sites, and that sites should really start a year in advance.

In slides 8 and 9 he showed serious site incidents of the last 6 months and gave general advice on how to improve reliability. Serious incidents are classed as ones where we ask for formal follow-up with a written document. As examples he mentioned the long term CASTOR oracle problems at ASGC and RAL due to using an unrecommended Oracle programming interface (that was not documented as such), the several day network downtime at PIC due to a service provider failure and the fire in an ASGC Uninterruptible Power Supply which has obliged them to move to a temporary hosting centre. He then proposed simple actions to improve reliability, which he stated was our outstanding problem, and also encouraged the holding of workshops and site visits.

His next topic was resources, slides 10 and 11, where he reported on the agreement of a new cpu benchmark, the HEP-SPEC06, which will be abbreviated to HS06, and where 1 HS06 = 4 KSi2K. A simple to run benchmark suite has been built as an accurate measure of each site is required and conversion of requirements, pledges and accounting to the new units has begun. In addition the process for automated gathering of installed resources, which will improve accuracy and normalisation between sites, has been agreed and deployment has started.

I.Bird then moved on in slide 12 to the expected timeline of WLCG activities in 2009-2010 which starts with the STEP'09 tests this June and leading to almost a year of data taking from October. In terms of computing resources the WLCG delayed the commissioning of the full 2009 capacity till October 2009 but maintains the 2010 commissioning to be by April 2010. I.Bird pointed out that

this 2010 commissioning will be right in the middle of data taking. He then gave details of the proposed STEP'09 tests in slide 13 pointing out that in the CCRC'08 tests tape recall for reprocessing at Tier 1s was done for several experiments but independently and it had been agreed at the WLCG workshop in Prague that in STEP'09 such testing schedules should be aligned and should also include analysis scenarios. Also on the WLCG timeline was the ending of EGEE in April 2010 and the transition to EGI and this was discussed on slides 14 and 15. The final EGI blueprint was published at the end of December and we are now awaiting the transition plan from the EGI-DS project. The EGI organisation should have been setup by September 2010 and calls for funding for EGI projects, such as an HEP Specific Support Cluster, will close at the end of November. I.Bird remarked that the timeline to achieve the full transition to EGI was very optimistic but that we have to go ahead with this process. He stressed that the parts of EGEE that WLCG relies on have to be in place.

The presentation then moved on to the updated experiment resource requirements where for planning purposes WLCG is breaking the long run into two resource periods of October 2009 to March 2010, which is termed the 2009 period, and April 2010 till March 2011, which is termed the 2010 period. He started with general comments, slide 17, emphasising that the LHCC mini-review held in February concluded we must ensure that computing is not a limiting factor when data comes. He then presented slides for each experiment and a summary (slides 19 to 23) where he compared their updated requirements for the 2009 and 2010 resource periods with the existing 2009 pledges and the old 2010 requirements. For ATLAS he remarked that a lot of effort is going into understanding the detector using cosmic data and that it is now realised to be important to store raw data on disk at Tier-1s for a few weeks to avoid tape recall delays. For CMS there have been some changes in their model concerning the number and turnaround timing of their reconstruction cycles which require more resources and for ALICE a big change is their intention to take p-p data at the maximum rate possible, thought to be about 300 Hz. LHCb changes included allowing for additional event pileup due to a new uncertainty in the running mode. The summary showed that cpu requirements at CERN had increased quite a lot and also the Tier 2 disk space but most other requirements had fallen.

I.Bird then dealt in slides 24 to 26 with why so many resources are needed when the beam live time will only be one third of that anticipated for 2009/10 and the implications of this. He pointed out that this will be the first year for data taking so is a critical opportunity and that important lessons have been learned from exercising the experiment models and from cosmic data which have led to the models being refined. He emphasised, for example, that analysis will have to be completely from disk rather than involving any tape mounting. Additional reconstruction passes will also be needed as the software evolves and the cosmic ray data is now part of the history of the detectors and cannot be simply discarded. He added that having a quarterly resource profile from the experiments could be helpful in scheduling installations during the foreseen long data taking period.

I.Bird then presented a summary slide, slide 27, reminding that we must use the analyses of service incidents to improve reliability between now and data taking and reemphasising that we must ensure that the computing resources are not a limiting factor compared with the huge investment in the accelerator and detectors.

Discussion

S.Bertolucci thanked I. Bird for his presentation and invited comments.

The first intervention was by Dr. Pripstein (USA/NSF) who referred to the proposed transition from EGEE to EGI next April or May and pointed out that he could not understand the impact should EGI be significantly delayed since Bird had said the essential services would still be there and he thus queried what EGI would be bringing to the table. I.Bird answered that EGI will bring two things – it is not just for HEP but will provide a general long term infrastructure to support us

and other disciplines and secondly it will be a source of funding. Today our model is fairly heavy on resources in operations and there are tools missing to reduce this. There is still some development effort needed in the next few years to produce a sustainable infrastructure and we need to be part of this effort to reduce our overall level of operational involvement. S.Bertolucci then added that there is an ongoing process between CERN and the European Commission to try to ease this difficult transition. He continued that we are discussing with them on a regular basis and will be meeting again soon. It is clear the transition to EGI will be long and we must make provision that this happens smoothly – we are after all the largest users of this infrastructure – and it is understood that we must not put ourselves at risk.

Prof. Barreira (LIP, Portugal and chairman of the EGI Policy Board) then intervened to say that though he thought the EGI transition timescale presented was optimistic he was convinced that they, EGI, would not be too far from that timescale. The experience now is that problems they had a year ago, mainly in relations between EGI, EGEE and CERN, have been completely overcome. Channels are working with constant communication between the departments and relations with the European Commission have been strengthened. EGI is strategic for CERN but the grid, in competing for support from the European Community, is not seen as the property of the HEP community. The National Grid Initiatives (NGI's) governance imposes upon them to be of service to not only particle physics but to the entire community and this is the only way to keep the support of the European Community. The call that will be open till November reserves a substantial amount of resources for the next two years for the launching of EGI. The numbers now approved are of the order of 25M Euros. He drew the attention of the delegates to the fact that EGI will be established as a non-profit private foundation where the partners and owners would be the NGI's which will be legal entities under their funding agencies with sustainable support from their governments and that EGI will not be able to do the national work. He reminded that EGI is trying to follow the same steps as the successful Geant network model which works not only because it is a European supported initiative but because it is based on NREN's (National Research and Education Networks) which are legal entities at national level having clear support from their governments. We, as the owners of the largest part of the operations of the grid, should establish in our countries strong relations between our material resources with national entities and guarantee that the lines of funding, as they exist now for NREN's, should also exist for the grid. This is the only condition for EGI to be a success and for the continued support of the European Community. What we are trying to do in the next few weeks is to try and identify among the 42 European NGI that exist at different levels of implementation the 15-20 that will be able now to sign the act of foundation of EGI in Amsterdam. It was very impressive that we were able to select in a short time the location, to be installed in the Amsterdam Science centre, and we hope that everything will be ready to take the opportunity of the call that will be open till November. If EGI will not be ready to sign the project by then measures have already been taken for having an existing entity to take in charge this project, which could be CERN, but we will not lose this first window to obtain support from the European Commission.

Prof. Ekelof (Uppsala, Sweden) asked that when I.Bird reported on requested and pledged requirements if the pledged requirements were what has come forward since the last RRB. I.Bird replied that the pledges are what was previously presented in the last RRB. Prof. Ekelof then asked how recent are the experiment requirements to which I.Bird replied that these are from the last few weeks, documents that were sent in a few weeks ago, adding that for ATLAS they had already increased their request at CERN last August, when they had reanalyzed their model, and this is not reflected in the pledge numbers. Their new request now is not so different from that of last August.

Dr. Heuer (CERN, DG) also posed a question to ATLAS pointing out that they were increasing their CPU request at CERN but decreasing it at their Tier1 and asking what is the logic behind that. Dr.F.Gianotti (CERN, ATLAS spokesperson) replied that the increased requests for CPU and disk at CERN were put forward already last year before the LHC incident and were due to the realisation by analyzing cosmic data and performing dress rehearsals and data challenges that the ATLAS requests to CERN were too low. To give a comparison they were at the level of one-third of the CMS requests. ATLAS then doubled their request because they will need more cpu and disk to do first pass calibration and alignment and this was independent of the LHC incident. Concerning the new Tier 1 and Tier 2 requests, in most cases ATLAS reduced these reflecting the

change in the LHC schedule. I.Bird added that there had been a perception within the LCG that, even from the days of the TDR, that ATLAS did not seem to be asking for enough at CERN. F.Gianotti confirmed that ATLAS had decided to rectify this in 2008 based on their experience with real data. I.Bird then pointed out that the budget was approved based on the TDR. Dr. Heuer then, in a lighthearted way and asking if he could be provocative, asked if could expect ATLAS to double their 2010 request because they were below CMS ? F.Gianotti replied this would not be the case at which point Dr. Heuer, to laughter from the delegates, asked for this reply to be noted in the minutes.

The next question on this item was from F.Ferroni (INFN, Italy) who queried the increased cpu request at CERN from LHCb. This was replied to by A.Golutvin (LHCb) making two comments – firstly that LHCb cannot use cosmic runs like the other experiments and that in order to prepare the detector for physics they will have to use LHC data to calibrate and align the detector and secondly, though they hoped to reach the nominal LHCb luminosity, the way this luminosity will be collected is different and in order to accumulate a substantial number of pico-barns they will probably have to run with increased pileup as implied by one of the LHC running scenarios. LHCb have to test their tracking and triggering in a completely different environment and for that they have now initiated full system tests and when they start data taking they have to be ready to tune their triggering.

There being no further questions the chairman thanked I.Bird again and moved on to the next item.

4. LHCC Deliberations (paper only)

Paper CERN-RRB-2009-041

S.Bertolucci stated that there was no material that had not been covered by I.Bird's presentation and that essentially what the LHCC recommends is that a simultaneous test of the experiments be done this year and then everyone will be waiting for data to make the next reality check. They also recognise that the experiments requests for resources are consistent with the plans presented before. R.Heuer (CERN, DG) queried the issue of ALICE resources but the chairman suggested to leave this till the Resource Scrutiny Group report.

E. Tsesselis, LHCC Scientific Secretary

5. Status of Common Project Accounts

Paper CERN-RRB-2009-042

S.Bertolucci then asked the head of Finance department, T. Lagrange, to inform the meeting of the status of the common project accounts.

T. Lagrange, CERN Finance Dept.

T.Lagrange announced that there had been no changes to the document so he had no further comments to make. No questions were forthcoming so the chairman closed this item and invited the chairman of the Computing Resource Scrutiny Group, D.Espriu, to present the next item.

6. Report from the C-RSG

Paper CERN-RRB-2009-067

D.Espriu

Presentation CERN-RRB-2009-068

D. Espriu reminded the delegates that this was the second time he had reported to the C-RRB and in slide 1 reviewed the mandate of the C-RSG. In slide 2 he reviewed the 2008 usage of the grid resources pointing out that the level of usage by the experiments does not seem totally optimal to the group which was a cause for concern. In slides 3 and 4 he gave the changes in the beam time assumptions being used by the group since the LHC incident. The 2009 beam time assumption (which covers the period to the end of April 2010) is now 1.7×10^6 seconds of p-p running while for the 2010 period 4.3×10^6 seconds of p-p running and 5×10^5 seconds of A-A running are assumed. These numbers include an accelerator efficiency beginning at 10% in 2009 and ramping up to 32% by the end of 2010.

Moving on to slide 5 D.Espriu emphasised the short period of time the group had had to review the new experiment requests and that this scrutiny has to be considered preliminary. The group had used the same methodology as in their 2008 review limiting themselves to the implementation of the respective computing models and considering any review of them to be a matter for the LHCC. Slide 6 compared the revised 2009/10 beam times to be about half those of the previous plans for

2008/9 and noted that some of the experiment computing models had changed in ways that the CRSG did not expect leaving them the impression that some of the modifications proposed do not have an understandable bottom-up justification. This led to the conclusions of the scrutiny shown on slide 7 where, taking into account the new efficiency and availability conditions, the convenience of some modifications for the first year of running and the necessity of additional Monte Carlo data in some cases, the group concluded that the experiment needs for 2009 and 2010 should be roughly satisfied by granting the resources originally planned for 2008 and 2009 respectively. In addition there should be savings on mass storage (tape) as the original plans envisaged data taking already in 2007.

Slides 8 to 11 then showed tables for ALICE, LHCb and CMS giving the new recommendations of cpu, disk and tape resources for 2009 and 2010 reflecting the groups conclusions given above. Slide 12 reviewed in more detail the ATLAS case stating that the group was not able to understand their latest requests to the same level as was possible last year but felt that the general trend of shifting resources by one year in view of the substantial reduction in beam time should also apply to ATLAS.

Some general remarks were then presented on slide 13, some of which were already given in the last review, notably that the C-RSG committed to scrutinising the experiments use of resources after the first months of data taking in 2009 and that they felt it important to inform on resource requirements for the calendar year 2011 as soon as possible suggesting the C-RRB modified the timing of their meetings to facilitate the procurement process.

Slide 14 reviewed the issues to be brought to the attention of the LHCC with a new one being that the issue of data analysis strategies be soon reviewed by the LHCC in order to ensure a reliable start-up and a coherent long term strategy. In particular the official reprocessing strategies should be better specified then reviewed by the LHCC.

The final slide reviewed the group's conclusions emphasising that this was a preliminary scrutiny but that they have noted unexpected changes in this years requests and they recommend taking into account only part of them. The final point was that the C-RRB should have the flexibility to react to unexpected changes and be able to commit more resources if they are really needed.

Discussion

S.Bertolucci invited questions and comments.

The first intervention was by F.Ferroni (INFN, Italy) who began by suggesting delegates look into the matching C-RSG paper as it contains many details. He saw now a general problem in that the scrutiny group for computing is looking at a vital part of the experiments, one that is alive (whereas the Maintenance and Operations reviews are not) and has pointed out there is some ambiguity between the LHCC and themselves and this has to be solved otherwise the timing will always be wrong. The very short time in which the C-RSG had to work is in some sense a waste of opportunity. He then asked that in the presentation of I.Bird there were tables with experiment requests and pledges while the C-RSG presents data in a different format so he did not at the moment understand which fraction of resources were they suggesting to be committed. D.Espriu replied to this that they had not wanted to overload the presentation with tables adding that last years review basically agreed to the experiment requests within the 10% level. For this year, 2009, they disagree with the experiments at this point though they are open for further interactions and he confirmed the judgement of the C-RSG that shifting resources by just one year should be sufficient and that this was their clear message.

A.Gurtu (Tata Institute of Fundamental Research, India) commented that we now have two different views and he wondered how we were going to reconcile them. He asked if the experiments had had some time to see the report and said he would like to hear what they had to say.

J.Virdee (CMS spokesperson) responded first saying that he would try to be diplomatic. He thought that it was unfortunate they saw the report only after it was posted, namely yesterday, and he thought it simplistic to shift resources by one year. As he had stressed in yesterdays R-RRB the implications of the long first run are wide ranging and also only decided some three months ago. It was clear that to see the competition off a large amount of data had to be accumulated and it was likely that the bulk of the good and analysable data would be taken in the last four months and was likely to be under conditions of significant pileup, conditions we were not expecting till a few years down the line. We must focus on this aspect and ramp-up the resources with this fact in mind – a possible luminosity that is only three times lower than the design luminosity. We would need to respond fast to changing conditions and have a rapid turnaround and feed this into our computing model like all the experiments. CMS plan to run at 300 Hz which, together with the complexity, determines the cpu need which is, as I.Bird said, an instantaneous need. This cpu determines the disk usage and, as has been said, the tape to disk recall speed is not as good as it should be which puts emphasis on the disk and this is especially true for CMS since their tape to disk ratio is much higher and analysis will have to be done off disk. He pointed out that CMS already have 4 Petabytes of tape data. Another aspect was the C-RSG report of an event overlap for CMS of 40% while that for CDF it is only 10% but this is after they have run for 20 years while this will be CMS first year. They are hence many issues in the report with which CMS did not agree and he added that it was unfortunate the C-RSG had set up its own simplified model while theirs was full and detailed. There had been disagreements last year and he thought these should be understood. It was important to engage in a dialogue now and sort this out because the factors were big and the impact could be not so nice.

S.Bertolucci then added that everyone had been expecting a higher level of disagreement than last year with a major difference being that this year there are some real data. He also expected the resource requirements to evolve very rapidly as soon as there was data from the LHC. He regarded the current C-RSG work as the start of a process because there is, unfortunately, as yet no beam data feedback from the experiments. The scrutiny group only got the new requests a few weeks ago and the scrutiny is not a process that could be finished in three weeks. He recommended that the RRB takes this as a starting point of discussions that should proceed during the summer and come to a first point of decision in October this year to be immediately followed when there is first data to assess the evolution of the requirements before the bulk of the luminosity comes, though this may be already too late since procurement takes time, but we will have to react as fast as we can in this moving situation.

F.Gianotti (ATLAS spokesperson) then spoke for ATLAS starting by saying that she would try to be more diplomatic. She thought that, as D.Espriu had said, it is not possible to understand in detail the quite complex computing models of the LHC experiments in a timescale of two weeks including the long Easter weekend. In her opinion the analysis could only be superficial without having had enough feedback and discussion. The C-RSG has recommended that since the LHC has shifted by one year the computing requirements could also be shifted by one year but she thought that this was wrong. Since last year the machine luminosity profile and the running scenario profile have changed from that of one relatively short run followed by a similar shutdown to a mode with very long runs followed by very long shutdowns and which has implications for the computing resources. She added that also the machine energy profile would possibly have more energy points on the way to 5 TEV, starting at injection energy as planned last year but then spending some time at 4 TEV, which would have to be taken into account in the ATLAS simulations. In addition another bunch structure has been introduced with a new 50 nanosecond point intermediate to the previously planned 75 and 25 nanosecond spacings. In conclusion, she said, there is a difference between 2008 and 2009 - the LHC schedule and machine profile have changed and this must be taken into account in the computing resources.

J.Schukraft (ALICE spokesperson) then contributed that ALICE has not changed its computing model and that concerning the scrutiny report they needed time to study it but he presumed that because they had not changed their model the discrepancies in requirements between ALICE and the scrutiny would be much smaller than for the other experiments. They would study the report then enter into a dialogue with the C-RSG.

Dr. Heuer (CERN, DG) then intervened to say that he thought it was quite obvious that there had not been enough time to get feedback between the experiments and the scrutiny group. He proposed that we should now use today and tomorrow to start a dialogue between the experiments and the scrutiny group and organize intermediate meetings between the group and the experiments before the summer break in order to come to a conclusion for the 2009 run that could be distributed to the RRBs. He was concerned about getting enough resources installed so wanted, for 2009, a conclusion before the summer break. He added that he did not intend to delay the LHC startup due to lack of computing resources. He also expressed the wish that a conclusion for 2010 could already be reached and, if not, the 2010 budget would come early autumn while we must keep in mind that it takes time to get computing resources delivered. He concluded by reemphasising his wish that the experts get together today or tomorrow to see how to continue so as to reach a first conclusion before the summer break and asked if this was acceptable to the scrutiny group and the experiments to which there was general agreement. S.Bertolucci announced that with this provision this closed this item and at this point Dr. Heuer left the meeting. S.Bertolucci then invited S.Foffano to present her report.

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

Paper CERN- RRB-2009-045

Presentation CERN- RRB-2009-046

S. Foffano began by reviewing her agenda (slide 2) where she would bring delegates up to date on the WLCG MoU status, briefly look at funding and expenditure at CERN then continue with the regular resource usage and accounting reports which, she pointed out, are becoming increasingly automated and more reliable. Her final subject would be resource requirements and pledges. She also encouraged delegates to read the written report and visit the LCG web pages which contain a wealth of information including detailed accounting reports.

7.1 WLCG MoU Signature Status

S. Foffano reviewed (slides 3 and 4) the status of signatures of the WLCG MoU. Since the last C-RRB Austria has signed the MoU and also Tier-2 sites in Germany, the University of Goettingen supporting ATLAS, and in Korea at the CHEP site supporting CMS. The signature of the MoU for the Brazilian Tier-2 supporting CMS is imminent.

7.2 Funding and Expenditure at CERN

S. Foffano then showed in slide 5, for what she said would probably be the last time, the final status of the WLCG Phase 2 (2205-2008) budget which ended with a zero balance of personnel and a materials balance of 0.6 MCHF. She pointed out that funding for this phase had been mainly from CERN but with contributions from certain countries for personnel namely Portugal, Israel, India, Italy, Germany, Taipei and Russia. Many of these were planned to stop at the end of 2008 but there is continuity with Russia, Taipei, Italy and India who are still providing people. The current plan for Italy, for example, takes us up to 2011 or 2012, currently being discussed, and with India we go up to 2010. She then took the opportunity to thank these countries for their contributions. Slides 6 and 7 showed the WLCG budget estimates for the 2009-2013 period which concluded that, based on the current plans, the project does not look likely to run into major financial difficulties.

7.3 CERN, Tier 1 and Tier 2 Resource accounting

S. Foffano then showed graphs (slide 8) giving a summary of the CPU Time, Disk and Tape Storage accounting from January 2008 to February 2009 obtained by summing up CERN and all external Tier-1s. She remarked that individual Tier-1 show different profiles for various reasons, such as the impact of the fire at ASGC Taiwan, or the impact of the LHC incident on the procurement plans of certain Tier-1 sites, and that follow up is regularly done via the WLCG meetings to understand these differences. Globally in this period pledge levels were met for CPU and largely met for disk and tape.

S, Foffano turned to accounting for Tier-2s (slides 10 to 13). She noted that globally CPU usage compared with the pledges has improved from the 40-50% level previously to over 100% now. The work on automatically obtaining installed capacities is now close to fruition and WLCG can also address the way it presents these data and would welcome the delegates suggestions.

7.4 Computing resource pledges

Slide 15 summarised the changes to the pledge tables since the last C-RRB meeting which mostly concern additional or revised values for Tier-2 sites and where the cpu is now expressed in the new HEP-SPEC06 unit. The tables are available on the LCG web site (<http://lcg.web.cern.ch/LCG>) and are dated 05/02/09. They cover the timeframe up to 2013 but the next pledge collection, which will be before the summer, will use the new agreed timeframe of +1 year for confirmed pledges and then +2 years for planned pledges.

On slides 16 and 17 S.Foffano reminded that there had been changes to the experiment requirements since the last C-RRB meeting and following the new LHC schedule. She pointed out that she needed the conclusions from the LHCC and C-RSG bodies before she could launch a new pledge collecting. On slide 18 she compared the confirmed 2009 pledge data with these new (unconfirmed) requirements showing that the main problems are in cpu at CERN and disk at the Tier-1s and reminded that exceptionally 2009 pledged resources are due to be commissioned before October 2009. Finally, on slide 18, she showed the current global pledge balance for 2009-2010 and pointed out that confirmed 2010 pledged data and planned 2011-2012 data would be needed by the end of September for the Autumn C-RRB.

7.5 Conclusions

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

- Brazil needs an MoU signature follow-up but there is nothing else in the pipeline.
- WLCG Phase 2 is now officially closed with an 0.6 MCHF balance.
- The global funding situation till 2013 suggests the project will not run into major financial difficulties.
- Most Tier 1 and Tier 2 sites are now reporting their accounting.
- Revised experiment requirements should be soon confirmed and will trigger the pledge collection exercise of confirmed pledges for 2010 and planned pledges for 2011-2012 to be presented at the Autumn C-RRB meeting.

Discussion

S.Bertolucci then invited questions or comments but none were forthcoming so he concluded there was general satisfaction with this status report and moved on to his meeting summary.

7. Summary

S.Bertolucci

He began by saying that the meeting had decided today that follow-up of the resources situation is of primary importance for many reasons. The experiments and the scrutiny group must work together to produce a more agreed way and this is needed by September otherwise we will not be able to meet any decisions for 2010. It was also pointed out several times that we need to increase our attention on two aspects – one is the transition from EGEE to EGI which may be very disruptive not only because we are losing personnel here at CERN and from the EGEE project just in the middle of 2010 when we will be running but also because the well established EGEE will be replaced by something which is growing and not yet established so we must follow that very closely. The other aspect he thought important was that the monitoring of the resources has to be

made more and more automatic and present in the future not only to see how much hardware we have available but the efficiency with which it is used so as to understand where the models might have their bottlenecks in order to be able to react. For this we need more tools and more automated tools.

These points are what he had learned today and he then opened the floor for any discussion with a view to making these conclusions more agreed. There were no remarks forthcoming at this point so S.Bertolucci declared the meeting closed.

<p>The next RRB meeting in 2009 is scheduled to take place at CERN on Monday 12 to Wednesday 14 October 2009</p>

H.Renshall
15 July 2009