



Computing Resources Review Board

20th April 2010

Minutes of the 17th Resources Review Board Meeting Held at CERN on 20th April 2010

Present: Delegates connected by Videoconference are marked as (EVO)

Europe:

J. Lemonne (EVO)(FWO, Belgium);
 J.D. Hansen (Niels Bohr Institute, Denmark);
 D-O. Riska (EVO), J. Tuominiemi (Helsinki Institute of Physics, Finland);
 E. Auge (CNRS/IN2P3, France); F. Malek (CNRS/LCG-France, France); E.Lancon (CEA/DSM IRFU/SPP);
 P. Chomaz (EVO)(CEA/DSM/IRFU/DIR, France);
 S. Bethke (EVO)(MPI, Germany); V. Guelzow (EVO)(DESY, Germany); K-P. Mickel (Karlsruhe Inst. of
 Technology, Germany); D. Muller (GSI, Germany);
 G. Vesztergombi (KFKI-RMKI, Hungary); T. Csorgo (EVO)(KFKI, Hungary);
 L. Levinson (EVO)(Weizmann Institute of Science, Israel);
 D. Bonacorsi (Deputy CMS Computing coordinator, Italy); U.Dosselli (EVO)(INFN, Italy); F. Ferroni
 (EVO)(Universita' Sapienza e INFN Roma, Italy);
 M. Turala (IFJ PAN, Poland);
 V. Ilyin (Moscow State University replacing F. Koslov of Federal Agency of Science and Innovations); V.
 Karjavan, V. Korenkov (JINR, Dubna);
 N. Colino (EVO)(CIEMAT, Spain); J. Fuster (EVO)(MICINN, Spain);
 T. Ekelöf (Uppsala University, Sweden);
 T. Nakada (CHIPP, Switzerland);
 D. Charlton, A. Medland(EVO) (STFC, United Kingdom);

Americas:

W. Davidson (EVO)(NRC, Canada); R. Mcpherson (EVO)(University of Victoria, Canada);
 I. Fisk, (FNAL, U.S.A.); A. Lankford (UCI, U.S.A.) R. Cousins (UCLA replacing J. Butler of FNAL,
 U.S.A.); H. Gordon (EVO)(BNL, U.S.A.); J. Shank (Boston University, U.S.A.).

Asia:

T. Kawamoto (University of Tokyo, Japan); H. Kamiyama (Permanent Mission of Japan);
 G. Taylor (University of Melbourne, Australia);

CERN:

S. Bertolucci (Chairman), H. Renshall (Scientific Secretary), R. Heuer (DG), E. Van Hove, T. Lagrange, S.
 Lettow, E. Tsesmelis;

LCG: I. Bird, S. Foffano;

C-RSG: D. Espriu (University of Barcelona)

ALICE: J. Schukraft;

ATLAS: F. Gianotti, K.Bos;

CMS: J. Incandela, G. Tonelli;

LHCb: A. Golutvin, A. Schopper;

Excused (due to air travel restrictions):

J. Sacton (FNRS, Belgium);
 S. Novaes (UNESP, Brazil); R. C. Shellard (CPBF, Brazil);
 Y. Zhang (National Natural Science Foundation of China);
 M. Lokajicek, J. Ridky, (Institute of Physics AS CR, Czech Republic); M. Sumbera (Nuclear Physics
 Institute ASCR, Czech Republic);
 B. Erasmus (IN2P3/CNRS, France);
 K. Ehret (DESY, Germany);
 A. Gurtu (Tata Inst of Fundamental Research, India);
 E. Rabinovici (Hebrew University, Israel);

M. Diemoz (INFN Roma, Italy);
U. Marconi, L. Silvestris, A. Vacchi (INFN, Italy);
A. Van Rijn (NIKHEF, Netherlands);
B. Jacobsen (The Research Council of Norway, Norway);
G. Polok (H.Niewodniczanski Institute of Nuclear Physics, Poland);
G. Barreira (LIP, Portugal);
Y-I. Choi (Sungkyunkwan University, Republic Of Korea); I. K. Yoo (Pusan National University);
F-D. Buzatu (Institute of Atomic Physics, Romania);
S. Pajovic (Ministry of Science of Serbia, Serbia);
P. Karlsson (Swedish Research Council);
G. Zinovjev (National Academy of Sciences of Ukraine);
J. Porter (LBNL, U.S.A.); T. Wenaus (BNL, USA); S. Gonzalez, A. Boehnlein, M. Procaro (DOE, U.S.A.);
M. Pripstein (NSF, U.S.A.);

The agenda of the RRB meetings are to be found on the web at <http://indico.cern.ch/categoryDisplay.py?categId=852>

1. Introduction

S.Bertolucci, Director for Research and Computing

S. Bertolucci, as chairman, started the meeting wishing the delegates a good morning and inviting them to start the long session of today beginning with computing. Tests of readiness with real data have been done in a massive way both in November and now and he thought the system had reacted very nicely to date as he expected I.Bird would tell us.

2. Approval of the Minutes of the Last Meeting (CERN-RRB-2010-003)

S. Bertolucci requested any comments on the minutes of the last meeting, CERN-RRB-2010-003. None were forthcoming so he considered the minutes as approved.

3. Status of the LCG Project

I.Bird, LCG Project Leader

Paper CERN-RRB-2010-032

Presentation CERN-RRB-2010-033

S. Bertolucci then asked Ian Bird to present the status of the LCG project.

I. Bird began with his agenda which would cover the current status including experience with real data, planning and milestones, planning for the new Tier 0, a summary of the EGEE to EGI transition and finish by looking at the resource planning up to 2012.

His slide 3 showed some interesting events from the current 7 TeV run showing that the accelerator is operating and that the computing system is working as expected. Bird said that you could not tell from looking at the WLCG status that there was an impact from real data. He showed on slides 4 and 5 that the WLCG is running higher and higher workloads, up to 650K jobs per day today, and with data transfers between sites at unprecedented rates. The internal traffic inside CERN CASTOR of 4 GB/sec from the Experiments is also at a rate we have never seen before and this has been achieved without any particular notice. Slide 6 made the point that WLCG also uses the EGEE and OSG grid infrastructures and that non-LHC work is also at a significant level and has not really been affected by the LHC start-up. He further illustrated the readiness of the computing on slides 7, 8 and 9 showing that data is processed and available at Tier 2 sites within hours, there are more and more users and the first physics papers have appeared. On slide 10 he showed a schematic of the LHC optical private network pointing out that the last missing Tier 1 backup path, that of RAL, was now in place so that the OPN is now fully redundant.

Bird then moved on to reliabilities and availabilities (slides 11 and 12) showing that although the average is high and stable not all Tier 1 deliver a stable service all of the time and this must be addressed. In the weeks before data taking sites did make significant improvements and reliability did improve as data taking started.

Bird then addressed planning and milestones on slides 13 to 17. WLCG is on track for the 2010 hardware commissioning that has been delayed from the usual April till June. Most milestones have been completed but there is still work to do on publishing site capacities, a slow and

laborious process. Having moved from setting up to regular operations there are now very few formal milestones and we must move more to tracking metrics for performance, reliability and scalability.

Bird's next subject was planning for the Tier 0 (slides 18 to 21). Current estimates, based on the experiment requirements and technology evolution, are now that the CERN Computer Centre will run out of power around 2013. IT department has started to prepare several stop gap solutions – an additional 400KW in building 513, remote hosting in Geneva and containers on the Preveessin site – as well as alternative options such as far away remote hosting. Decisions for the medium term should be taken in 2010 in the light of experience of data taking and once alternative options can be fully evaluated.

Bird then moved on to the EGEE to EGI transition. On slide 22 he described three projects - the EGI-Inspire project, the core of EGI focussing on support for existing large communities, and which is likely to be funded; the EMI middleware project which is also likely to be funded though with a cut of 900K euro and finally the support for virtual research communities, including HEP, which was not funded. He said that in spite of these cuts the situation does not represent a major risk for WLCG. The main change will be that work with other application communities will significantly reduce at CERN. He also clarified the point from the last C-RRB on the status of EGI in some non-EC states in that there is now an EGI associate member status and that non-EC countries could become full partners of the EGI-Inspire project.

This brought Bird to his last subject namely resource planning. On slide 26 he showed the LHC schedule to run through till 2011 and accumulate one inverse femto-barn of collisions followed by a long LHC shutdown in 2012 to increase the LHC energy limit to 14 TeV. He pointed out that this did not imply a shutdown of computing in 2012. His guidance for the next three years was that the 2010 pledges and installation schedules cannot be changed but for 2011 we should revert to the April resource installation deadline. He showed on slide 28 the summary of experiment requirements up to 2012 where there is a flattening of storage in 2012 due to there being no new raw data. He expressed concern in the planning over an announced 40% budget cut in France for 2010, that the EMI middleware support will be close to the limit of what we need leaving little for development and that data access for analysis is a concern. In this last area there will be a workshop in June which is likely to lead to significant changes over the next two years.

Bird summarised his report on his last slide, number 30, saying that first experience with data has been positive for WLCG, resource planning for the coming years are a concern, having many more non-expert users could have an effect, that the transition from EGEE to EGI is underway now and, finally, that we must start to address the long term sustainability of the system we have.

Discussion

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

The first intervention was by E. Auge (CNRS/IN2P3, France) referring to their budget cut mentioned by I.Bird. This did not come by chance and he wanted to make two points. Firstly the grid is now supposed to be in place so in terms of project management we have moved from a construction phase to one of upgrade and evolution and the speed of evolution is less critical once the infrastructure is in place. Secondly we must be very careful to be consistent in what we show in the usage of our resources and this was a weak point on our side where, for the year 2009, we had more than what was needed so part of it was not used hence putting us in a weak situation to ask for more money for the next years. His concern now was to make every effort that this does not happen a second time, he will be fighting hard in 2011 but this will depend on the 2010 resources really being used. We must show to our authorities that we are not over-equipped. I.Bird replied to this saying that WLCG had ramped up resources to a certain level expecting the accelerator to start but then we had delays so it is hard to say all those resources could have been used and we must be careful that we do not start to affect future years by making cuts too early. E. Auge then added that specifically for the Tier 1 in Lyon the pledged cpu will be there but they had

pledged about 5 PB of disk to be added to the existing 3 or 4 but will now provide only 3.5 PB extra.

Prof. Ekelof (Uppsala, Sweden) then queried I.Birds remark that effort within the European Middleware Initiative project is worrisome as there is a funding cut of 900 K Euro. Bird replied that he thought the cut would not have too much effect in that the partners have agreed how to absorb the cut without really reducing the program of work. He continued that the project proposal itself was very specific on its ongoing support for a certain set of components and attempting to integrate existing middleware stacks but that there is no investment in there for significant new developments. Any thinking we have on doing data management differently in the future will not be funded by EMI but must be found in our community. We can rely on EMI for the ongoing support of existing components we use but it will not be a source of funding for major new developments. Prof. Ekelof then asked if the EMI leadership should have included such development in its project to which Bird replied that the guidance from the commission was that this would not have been acceptable. Prof. Ekelof then commented that he agreed there is a general difficulty in getting support from funding authorities for development on middleware and that he found this to be strange. Bird replied that the thrust now was not on grids but rather on data and that we as a community should understand what we need to do and identify those developments where it makes sense to go to the commission or other funding agencies. We must be sure this goes in directions where we need to go and not just jump on a bandwagon to get money. Prof. Ekelof then asked where was the initiative to sustain long-term development of new middleware to come from. Bird confirmed what he had said earlier that some of this has started – there have been first discussions and there will be a workshop in June to flesh this out and he clarified that this initiative is coming from the experiments and WLCG.

Prof. Ekelof then asked on the planning for 2013 which he had not seen in Bird's tables. Bird agreed that there was a problem there as we had little idea of what is going to happen in 2013. Our best guess would be a linear extrapolation based on 2011 but nobody has the experience with data yet on which to base further estimates. It would be risky of us to now make strong statements on 2013 though if this is a problem we could decide in the management board what we could say but there would not be a lot of science in any numbers.

The next point was from V. Guelzow (DESY, Germany) who posed a question concerning the VCards information from ATLAS where had received an email from ATLAS indicating a change in the resource information contained therein in particular the job memory requirement has increased from 2GB per core to 3GB per core and he wondered how this could be financed or compensated for by perhaps reducing cpu power. Bird said there was no formal request behind this – there had been formal agreements on the amount of memory per core in the past and any changes in this would need to be agreed. He then invited K.Bos of ATLAS to comment. Bos expressed his agreement that this was not a formal request for an increase in real memory but this was in fact concerning virtual memory. For short periods a job might take more than the available 2GB of memory so would need to use swapping to disk but this does not mean an increase in cost nor in real memory. V. Guelzow added that there was also a request for more scratch space per ATLAS job but his issue was that such increases imply extra money which is not available at Desy. Bird confirmed that there is no requirement here to buy extra physical memory.

E. Auge (CNRS/IN2P3, France) then asked Bird about his Tier 0 expansion planning where there are still open questions so when did he expect decisions would be taken in particular is the question of a new building as compared to containers decided. Bird replied that this decision had not yet been made – it will depend on more concrete proposals being made such as from Norway – but he agreed that we really must make a decision on what the strategy should be in the first half of this year.

There being no further questions the chairman proceeded to the next item on the agenda.

4. LHCC Deliberations (paper only)
Paper CERN-RRB-2010-036

E. Tsesmelis, LHCC Scientific Secretary

S.Bertolucci invited E. Tsesmelis to comment on his paper. Tsesmelis explained that since the last C-RRB there had been one LHCC meeting in February this year with its conclusions written in the attached paper. He summarised them by saying that the LHCC considers that the WLCG has made excellent progress in all aspects and congratulates both the WLCG and the experiments on the performance of the grid until that point with the first collision data and analysis while recognising that the system still needed to be stress tested. That has been happening in the intervening period so some very good and encouraging words from the committee.

There were no comments forthcoming to this report.

5. Status of Common Project Accounts T. Lagrange, CERN Finance Dept.

Paper CERN-RRB-2010-037

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

T.Lagrange thanked the chairman and announced that there was nothing to add to the paper that had been submitted. S.Bertolucci invited questions then, there being none, closed this item and invited D.Espriu to present the next item.

6. Report from the CRSG

D.Espriu, CRSG Chairman

Paper CERN- RRB-2010-052

Presentation Slides attached to the agenda

D. Espriu began by reminding the delegates, on slide 1, of the members and the mandate of the Computing Resources Scrutiny Group. On behalf of the C-RRB the group scrutinises the resource accounting figures for the preceding year, the use the experiments have made of those resources, the match between the refereed requests and the pledges from the institutes and the overall requests for resources for the following year and forecasts for the subsequent two years. In addition the group can make recommendations concerning apparent under-funding though this has not been necessary so far. For the first time a substantial part of the report is dedicated to resource usage.

In slides 2 to 9 Espriu reviewed the resource accounting and usage by the experiments in 2009 with information coming from the WLCG accounting reports, the EGEE accounting portal and the reports that the experiments are asked to submit to the CRSG by 1 March every year. He added that unfortunately there is no full accounting from the Tier 2 sites yet. Referring to the detailed tables on slide 4 Espriu said that some of the experiments had made large use of their resources while others had not. This was clearly shown on slide 6, the percentage of use of the Tier 1 and CERN resources by experiment in 2009. Espriu summarised the group's conclusions that the experiments have not used all the WLCG resources available to them but the usage has been both extensive and intensive. One of the experiments is clearly below the average in using the requested resources but the group thought it too early to draw consequences.

Espru then reported on delivered versus pledged resources in slides 10 to 13 showing that some sites had over-delivered notably cpu at CERN. At the end of the reporting period only two Tier 1 sites, TW-ASGC and IT-CNAF were below having delivered 90% of their pledged cpu. The mismatches were more visible in disk capacity which is to be expected as it is the most expensive commodity. Only four Tier 0/1 sites were above the 90% level delivered while five were actually below the 80% delivered line. Espriu commented that none of these mismatches jeopardised physics analysis or simulation in view of the performance of the LHC and in some cases were triggered by lack of use of existing resources. However the CRSG thought there was a danger in this attitude and recommended that in the future pledged resources are made available within the agreed WLCG deadlines.

Espru then moved on to general issues to be monitored by the LHCC and CRSG explaining that over the last few months a fluid dialogue, including joint meetings, has been established between the CRSG and the LHCC. During 2010 no specific issues appeared that were thought necessary to be referred to the LHCC and in slide 16 Espriu repeated the standing recommendations to the

experiments also presented in the last meeting (reducing event sizes, proliferation of data formats and so on).

Slide 18 showed the agreed beam time expectation up to 2012 with 19 months of data taking delivering 12.4 Million seconds of pp live time and 1.4 million seconds of heavy ion live time. This lead on to the scrutiny tables presented in slides 19 to 25. The experiment requests contained only a limited number of changes with respect to previous iterations of the scrutiny procedure, an indication of the maturity of the computing models. The CRSG generally endorsed changes due to larger event sizes and processing times but asks the experiments to try and establish reduction profiles. The CRSG regards the situation for 2010 as somewhat exceptional, driven by the transient character of the start-up months, and commits itself to revising the present scrutiny ahead of the October 2010 C-RRB in the light of the usage of the resources up to that moment.

Reviewing the individual experiments Espriu said that for ALICE the group had some reservations due to the low resource usage they had made. For ATLAS the CRSG estimates were typically 10% lower than those of the experiment but those of the CRSG should be regarded as a bare minimum compared with the experiments healthy estimates. For LHCb there was no difference between the CRSG and experiment estimates. After several interactions the scrutiny and experiments requests now show a good degree of convergence and the provision of simplified spreadsheets by the experiments was very helpful. The two sides will work together to clarify those cases where experiment requests are systematically higher than those of the CRSG.

Espriu then made several general recommendations from the CRSG to the collaborations (slide 27). The experiments should incorporate the running conditions, in particular the implications of pileup, into their models in a uniform way. Experiments should make maximal use of the distributed resources in the grid. At CERN a clear separation should be made between Tier 0 function resources and those used to perform physics analysis by the CERN based physicists. Different data distribution strategies are being used and the CRSG recommends the experiments use the upcoming data taking period to determine which strategies optimise physics output while keeping resource requirements at a reasonable and sustainable level.

The final subject was on the CRSG membership where M. Gasthuber (DESY) is proposed to replace H. Martin and B. Panzer (CERN) is proposed to replace J. Knobloch. Also in the coming year three members should be rotated as specified in the WLCG MoU.

The last slide, number 29, summarised the report. The 2011 requests have been scrutinised after the first data taking in 2009 as has the resource usage in 2009. Though there are still some potentially troublesome issues the CRSG and experiments have reached a reasonable degree of convergence and the CRSG is confident the computing models will pass the test of the 2010/11 run. The CRSG is nevertheless concerned about the long term sustainability of the computing models.

Finally the CRSG congratulates the four experiments and the machine people for their excellent performance in the first days of the LHC and for their outstanding work.

Discussion

S. Bertolucci thanked the speaker and invited comments and questions.

The first comment came from the ALICE spokesperson, J. Schukraft, who reminded that in 2009 the LHC delivered only a very small fraction of the data which was originally foreseen, even including the collected cosmic ray data, so he thought the low usage was to be expected as it is clearly related to the amount of data we take. ALICE does have a problem with lack of computing resources so have adopted a rather diligent policy of using resources, particularly of tape and disk which are cumulative, so having only a small fraction of the data used only a small fraction of the resources. The fact that ALICE used less than the average share of resources is due to our very strict policy to be careful with the resources particularly when they are cumulative. Tapes that we fill today cannot be used tomorrow which is why, for example, we did not replicate our cosmic

data outside, We are trying to keep our usage of resources in proportion with the amount of data we have otherwise we will be in trouble once we get the big amount of data we expect this year.

The next question came from V.Ilyin (Moscow State University, Russia) who made the point that, looking at the 2010 numbers, there are differences between the computing model resource requests and the pledged, available and used resources on which the CRSG is reporting. For example in the Tier 2 the total disk space offered for 2010 was much higher than that required by the computing models except in the case of ALICE where it was lower. He thought we should try and understand why there are these differences and what they mean. Espriu began his reply by saying that to first approximation offered and available resources were the same. The CRSG monitors usage using the available information which does not yet include disk space at Tier 2. It will be more important in the coming years that the experiments show that they are really using correctly what they requested. Ilyin then added that from the table to which he was referring (this was later clarified by J. Schukraft to be table 5 in the CERN-RRB-2010-034 report of S. Foffano) the Tier 2 offer of disk space to CMS for 2010 was about 14 PB while the experiment computing model requirement was for 9 PB. For ALICE he remembered that the offer was only about a third of what was required. Espriu answered that in 2010 CMS used 40% of the Tier 2 disk that was available but that this number came from the experiment and cannot be checked. Ilyin reiterated that his point was somewhat different – it is that what countries are offering for Tier 2 resources is not correctly connected with the computing models. Espriu thought that in the case of Tier 2 there was an over-commitment of both cpu and disk resources in 2010 for two reasons – the efficiency for the Tier 2 was larger than our standard assumption then the deadline for the installation of 2010 resources was moved in the middle of the year from April to September and many Tier 2 had already commenced procurement so ended up installing more resources than was actually requested by the experiments. Espriu did not think this situation, due to the shifting accelerator timescale, would repeat itself in the future. The scrutiny group has always planned to be as rigorous as possible but did not want to leave the experiments without the resources they needed and taking that into account they have tried to do their best and finally the decisions on resources will be taken around this table. At this point J. Schukraft intervened to clarify that the numbers referenced by Ilyin came out of the report of S. Foffano, yet to be presented, showing in table 5 the resource offers, requirements and balances for 2010 where there are quite large differences with one experiment having 2000% more disk than required, another 60% more and another 60% less. He added that the offered resources for 2010 were not part of what the scrutiny group looked into.

The next point was made by A. Medland (STFC, UK) who asked about the scrutiny process saying that we had heard earlier about the importance of being aware of the budget pressures and concerns facing some of the funding agencies and their need for some assurance on the optimisation of resource usage particularly as experience with data taking is gained. This is referred to in the scrutiny group report which says they will revisit some of the requirements before September so he was wondering how the scrutiny group felt about the timescales and would they be able to get a better feel for the balance between resource needs and resource provision. Espriu replied that he thought the CRSG would be able, with the help of the experiments, to meet this timescale. By now they had a rather complete understanding of the models and had the tools prepared. They will have to see how the data taking and analysis summer season develops, which will stress the models, and it will be very important the experiments report to them on time. By mid-October he expected the group would be able to say if their initial scrutiny presented here is maintained or required some modification. He also hoped in October to say more on the resource expectations for 2012. Medland then continued that the CRSG report recommends that experiments reach a compromise in the requested resources and should try to establish a data-reduction profile and asked Espriu if he could elaborate on that. Espriu replied that this depended a lot on the experiment – with some the interactions have been more efficient than others which is to some extent related to the experiments internal organisation and structure. The experiments are aware in general that there should be an optimisation of the data and MonteCarlo processing times and we have seen substantial improvements for example in the time ATLAS need for simulation which we have welcomed. ATLAS is also trying to optimise their data distribution policy – this is very ambitious intending to make the data as available as possible to the individual physicists and groups. This consumes a lot of resources and they are thinking in the mid-term to modify this policy which we welcome. CMS also has been quite helpful and after conversations with us have

understood that in 2011 they will need less to replace data on disk and other modifications that are explained in our written report. The experiments understand that this is an area they need to optimise but this needs some time and 2010 will be essential to get a more complete idea of what is going on. Above all the experiments and the CRSG do not want data to 'fall on the floor'.

There were no further points raised so S. Bertolucci again thanked D. Espriu for his in-depth work and requested S.Foffano to report on the status of resources and financial plan.

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

Paper CERN- RRB-2010-034

Presentation CERN- RRB-2010-035

S. Foffano began by wishing a good morning to those in the room and connected remotely. Her presentation would follow-on from some of the things I. Bird said earlier. Her agenda, shown on slide 2, would cover the WLCG and MoU status, funding and expenditure for WLCG at CERN over the next 5 years and a little further, resource usage accounting, an area that is being looked at more and more, and resource planning where the 2010 pledges are now final and the next pledge exercise will begin soon.

7.1 WLCG Collaboration and MoU Status

S. Foffano showed on slide 2 a colourful photographic representation of the WLCG today. There is one Tier 0 site, 11 Tier 1's and 64 Tier 2 federations comprising 125 sites. Thirty-four countries are represented with a total of 49 MoU signatories. There have been no new signatories since the last C-RRB but there are a couple of countries interested in becoming Tier 2 but with no formal commitment. Following some recent requests, a complete electronic archive of all MoU documents has now been created and is available via the WLCG website or by request to lcg.office@cern.ch

7.2 Funding and Expenditure for WLCG at CERN

S. Foffano then showed in slides 4 and 5 WLCG budget estimates for WLCG at CERN 2009-2015. The 2009 book closing showed a slight overspend in the personnel budget that had been absorbed by the IT personnel budget. The materials budget had a carry forward of 4.4 MCHF into 2010 since not all equipment had been accepted by the time of the 2009 book closing. The personnel planning up to 2015 includes current commitments, planned replacements and some new requests. The materials planning is more difficult being based on costs to satisfy experiment requirements up to 2012 then extrapolated beyond and also on the future of the Tier 0 resources. Last year she had shown figures including re-profiling for a new computer centre building where the agreed 60MCHF featured in 2011-2013 inclusive. In the meantime the future strategy has been clarified and, even if not finalised, it comprises getting the maximum out of the existing centre, profiting from some technical advances, planning for a container solution and considering offers for external hosting. The bottom-line is that the costs are expected to be the same as for a new building but spread more over a longer time period. She hoped to give more final figures on the materials budget at the next meeting.

7.3 Resource usage accounting

S. Foffano then showed graphs (slides 7 and 8) giving a summary of the pledged, installed and used CPU Time, Disk and Tape Storage accounting separately for 2009 with cpu in SI2K units and the first two months of 2010 with cpu in HS06 units both obtained by summing up CERN and all external Tier-1s. There had been technical problems with the accounting portal but these have now been resolved.

S, Foffano turned to accounting for Tier-2s (slides 9 and 10) showing a snapshot of the March 2010 status ordered by 2009 pledges and split into two plots for sites with more than and less than 4500 HS06 pledged. For the larger sites many are already advanced on their 2010 pledges while the plot of smaller sites showed many not yet reporting their accounting. She urged the sites not yet reporting to do so and encouraged all the sites to consult the accounting portal in particular to check their 2010 status.

7.4 Resource Planning

Slide 13 summarised the pledge table changes since the last meeting at which there had been numerous footnotes to be resolved. The Russian pledge allocation per experiment has now been received as had the NDGF input clarifying their Tier 1/2 shares. This allowed tables to be published on 24 February but then on 12 April came news of the French budget cut. The table on the WLCG web site dated 12/04/10 contains all these changes. Slide 12 showed an extract summarising the external Tier 1 and Tier 2 per experiment pledges together with the experiment requirements and the resulting balances and a sidebar showing the status as shown at the Autumn 2009 C-RRB. On slide 13 she reviewed the 2011 and 2012 status where the experiments have recently reviewed their requirements and those for 2011 have been through a thorough scrutiny. Foffano requested that funding agencies and sites prepare as usual, for the Autumn C-RRB, their confirmed pledges for 2011 including the distribution between experiments for sites supporting more than one and their planned pledges for 2012 which can be more flexible. She would like this input before the end of September 2010 and would show the results at the next meeting. She will discuss with the experiments and the C-RSG the experiment requirements numbers then will circulate this request.

7.5 Conclusions

In conclusion (slide 14) S. Foffano summarised that:

- WLCG is now a production service with 49 MoU signatories representing 34 countries.
- Funding and expenditure planning has been adapted to the latest Tier 0 strategy but will continue to evolve.
- Tier 1 and 2 accounting is now being watched much more than before.
- Pledge data for 2010 is hopefully now stabilized and resource commissioning at all sites should be in its final stages.
- Planning should now begin on confirmed 2011 pledges and planned 2012 pledges to be sent before the end of September 2010 to prepare for the Autumn C-RRB meeting.

S. Bertolucci thanked S. Foffano and invited questions.

V. Ilyin (Moscow State University, Russia) began by making a comment on the delay in splitting the resources for Russia. The summary figures for the Russian Tier 2 were available already in October but there was then a four month period of complicated and heavy discussions, he would say a struggle between the experiments, on how to split these resources and he would explain why. Firstly they collected requests from the experiments on their 2010 expectations from Russia and these were at least two times higher than they could provide by their budget, even three times higher for some experiments. Finally a consensus was reached in January to base the splitting on the number of official authors in the experiments. This calculation was done so that in February a table was submitted. This exercise showed a very serious discrepancy between the requests based on the computing models and what is a countries vision to provide resources, a serious problem probably also for other countries. Now we have this model in place there will not be such a delay next year.

S. Bertolucci thanked Ilyin for this comment adding that his perception was that for most countries this large discrepancy between what is perceived by the agency and what is requested by the experiments is not so big. Experiments tend to put themselves on the safe side while agencies tend to be slightly on the other side but he thought that in most European countries the fluctuation is 10 or 20% and not factors of two. He would encourage this approach with local groups for the sake of fairness and not to give the wrong impression to a collaboration. He invited Foffano's view and she added that it comes back to what she said that it is now important that between the scrutiny group and the experiments we come up with an agreed request number for 2011 for the sites to base their pledges on. She thought we were almost there but one more iteration would be needed and asked Espriu his opinion. Espriu replied that the CRSG had been pretty close to agreement with the experiments. The worst discrepancies were at most 10% and the group do not claim the

precision of their numbers is 1%. If there is no variation larger than 10% he suggested to stick to the experiment request. If there were larger variations than 10%, though he thought not, then he would recommend to add 10% on top of the scrutiny group judgement.

V.Ilyin (Moscow State University, Russia) then made another comment related to efficiency which for them was a problem and where he exposed two cases. One was that the software of the experiments was not stable with new versions coming very fast requiring some time for their centres to recover to stable mode. His hope was that the 2010-11 period would have more stability in the experiments software. Another example started at the end of last summer when some experiments strongly requested sites should install the CREAM Computing Element. However this was not included in the regular LCG packaging for installation so their system administrators in many centres had to do much manual work. This work was then wasted when a new version of LCG was released. There was also a problem that some centres refused to install the CREAM-CE because it was not in the regular LCG package and an experiment then blacklisted the site. His request was hence that there should be much stronger coordination between experiments and the LCG when starting to use newer software components. I. Bird then added that he thought it important that the Russian system administrators, or indeed any concerned, join the Grid Deployment Board discussions where exactly these issues are discussed.

E. Auge (CNRS/IN2P3, France) then intervened to make a small precision that he was talking in the name of IN2P3 as regards their funding cut and that the other funding agency in France, which is CEA, has been paying its share.

There were no further points raised so S. Bertolucci proposed to make a very brief summary.

8. Summary

S.Bertolucci

S.Bertolucci began by saying that it was clear that in general the perception of how well we are using the resources is increasing. It is also clear that we still want to improve on that – the worst thing that could happen is that we give the impression that we are not using optimally the resources that we have. He was sure that we would never be able to agree on the perfect utilisation as in computing, like in most other things in life, one must allow some headroom. He would not expect that in the future we would be able to express our needs to within 1 per cent and he would not like that the experiments be impeded by lack of computing after so much investment. He thought that in general what happens will be dictated by two things – the experience we will gain in the next months with real data and the number of seconds the LHC will be running in the next years. Before the next meeting most of this information will be available and he thought that measured decisions could then be taken.

M. Turala (IFJ PAN, Poland) then intervened saying that he would like to state that WLCG is a great success. Fifteen years ago we did not know how to do computing, ten years ago we were not sure we could manage but in fact seeing the results of the experiments it looks like things are working very well. He suggested it would be a good idea if the CERN management could issue a statement on this to those who had signed the Memorandum of Understanding which could later be passed to the contributing institutions. His institute was using three different institutes as Tier 2 sites and he would like to pass them this message on a more formal basis. He thought there were more than one hundred institutes that supported them altogether and it would be good to pass to our colleagues that this was appreciated and, of course, we need more in the future – that is also important to show. Bertolucci thanked him for this good idea saying he would be happy to issue such a statement. The CERN management were also very pleased that the WLCG works and is a good example of a successful project.

There were no further questions or statements forthcoming at this point so S.Bertolucci declared the meeting closed.

<p>The next RRB meeting in 2010 is scheduled to take place at CERN on Monday 11 to Wednesday 13 October 2010</p>



H.Renshall
25 May 2010