



# Computing Resources Review Board

13<sup>th</sup> October 2009

## Minutes of the 16<sup>th</sup> Resources Review Board Meeting Held at CERN on 13<sup>th</sup> October 2009

### Present:

#### *Europe:*

J. Lemonne (FWO, Belgium); J. Sacton (FNRS, Belgium);  
M. Lokajicek, J. Ridky, (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); F. Malek (CNRS/LCG-France, France); P. Rebougeard, U. Bassler (DSM IRFU)  
S. Bethke (MPI, Germany); K. Ehret, V. Guelzow (DESY, Germany); K-P. Mickel (Karlsruhe Inst. of Technology, Germany); M. Pantea (Ministry of Education & Research, Germany);  
G. Vesztegombi (KFKI-RMKI, Hungary);  
L. Levinson (Weizmann Institute of Science, Israel);  
M. Basile, C. Bozzi (INFN, Italy);  
B. Jacobsen (The Research Council of Norway);  
A. Van Rijn (NIKHEF);  
M. Turala (IFJ PAN, Poland);  
G. Barreira (LIP, Portugal);  
F-D. Buzatu (Institute of Atomic Physics, Romania);  
V. Savrin (Institute of Nuclear Physics, Russia); V. Ilyin; A. Sissakian (JINR, Dubna)  
N. Colino (CIEMAT, Spain);  
T. Ekelöf (Uppsala University, Sweden); P. Karlsson (Swedish Research Council);  
T. Nakada (CHIPP, Switzerland);  
G. Zinovjev (National Academy of Sciences of Ukraine);  
J. Butterworth, D. Charlton, A. Medland (STFC, United Kingdom).

#### *Americas:*

S. Novaes (UNESP, Brazil); R. C. Shellard (CPBF, Brazil).  
W. Davidson (NRC, Canada); R. Mcpherson (University of Victoria, Canada);  
P. McBride, I. Fisk, J. Butler (FNAL, U.S.A.); S. Gonzalez, A. Boehnlein, M. Procaro (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:*

Z. Wang (Institute of High Energy Physics, CAS, China); Y. Zhang (National Natural Science Foundation of China);  
Y-I. Choi (Sungkyunkwan University, Republic Of Korea); J. E. Ha (Minister of Education, Science & Technology, Republic Of Korea); H. Kim (National Research Foundation of Korea); I. I. Lee (MEST); I. K. Yoo (Pusan National University) I. H. Kim;  
A. Gurtu (Tata Inst of Fundamental Research, India);  
T. Kawamoto (University of Tokyo, Japan); M. Yamauchi (KEK);  
S.C. Lin (ASGC, Taipei).

#### *CERN:*

S. Bertolucci (Chairman), H. Renshall (Scientific Secretary), F. Hemmer, G. Tonelli, P. Bloch, E. Van Hove, J. Salicio Diez, H. De Groot, T. Lagrange, S. Lettow.

*LCG:* I. Bird, S. Foffano;

*C-RSG:* D. Espriu

*ALICE:* J. Schukraft, Y. Schutz;

*ATLAS:* D. Barberis, F. Gianotti;

*CMS:* M. Kasemann, 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 inviting them to start what he expected to be a busy day. The computing this year for the LCG had meant a lot of exercises together with the experiments processing millions of events and gaining a lot of experience from the fact that, for the first time, experiments have been working all together this year, as required by the scientific committees, to stress the system by addressing all of its capabilities at the same time. This would be reported on by Ian Bird.

## 2. Approval of the Minutes of the Last Meeting (CERN-RRB-2009-076)

S. Bertolucci announced he had received one small correction to the previous minutes namely that the second polish delegate had been G. Polok and not P. Grzegors, which we would amend, and asked delegates if they had any other amendments to the minutes of the last meeting, CERN-RRB-2009-076. None were forthcoming so he considered the minutes as approved.

## 3. Status of the LCG Project

### I.Bird, LCG Project Leader

Paper CERN-RRB-2009-126

Presentation CERN- RRB-2009-127

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

I. Bird began with his agenda which would cover a status report on STEP'09, which had been the main activity in the last few months, then give the usual reviews of service performance, middleware, storage and the status of Tier-0 planning and conclude with the status of the EGEE to EGI transition.

His slide 3 showed the WLCG timeline for 2009/10 showing STEP'09 was completed in May and June of 2009. He indicated that the timeline for LHC running had changed slightly thus supporting the C-RSG recommendation to delay the 2010 acquisitions until June 2010. Slide 4 showed key performance plots from STEP'09. CERN tape writing was well above the required rate and the LHCOPN was able to run well above the required rates even surviving a double fibre cut to one site, where the backup fibre was in the same trench, by falling back to the general purpose network. Most Tier-1 showed impressive operations and where there had been some issues these sites were retested over the summer. Slide 5 showed that after STEP'09 workloads have been continuous with the WLCG service running according to the defined procedures and where the focus had been on stability and resolution of problems in preparation for data taking. He showed the level of open problem tickets indicating a steady level even though there had been an increase in workload.

I. Bird then gave specific information on STEP'09 from each experiment starting with CMS on slides 6 to 8. He highlighted that CMS processing ran smoothly at the pledge levels of the numbers of job slots. Where sites had downtimes these were either known of ahead or the sites had since been retested. In data transfers there were some Tier-0 to Tier-1 latencies with tails of up to days impacted by the tape system state. Analysis was tested at a higher level than before and CMS found that moving data to and from Tier-2s was not optimal especially for inexperienced sites and Bird remarked that other experiments had also seen this. CMS closely monitor site readiness where they test links by real CMS jobs and Bird showed a plot that worries them where there was a decrease in 'ready' sites over the summer holidays and he wondered what might happen should there be a flu pandemic. His last slide for CMS comparing Tier-2 site readiness between March and October 2009 did, however, show far more sites to be ready and performing a significant amount of analysis.

I. Bird then moved on to ATLAS and STEP'09 (slides 9 and 10). STEP'09 ran at full rate for 11 days comparing well with 2008 when only a few days were sustained at full rate. In terms of data transfer performance 5 Tier-1 sites reached 5 times the nominal data taking rate while 4 reached

90% of this metric and one reached only 50% with some of these reductions being due to scheduled outages. ATLAS added 4 PB of data in the 11 days to the ATLAS total of 12 PB on the grid and real physics simulation production continued during the whole period. A world wide peak of greater than 32000 running jobs was reached with about 1 million jobs submitted and 26.3 billion events processed. The mean cpu to walltime ratio was 0.39 which was not as good as it could have been.

Moving on to LHCb (slides 11 to 13) Bird showed that their data transfer tests were successful but their data reconstruction was delayed by problems with their conditions database. Over the summer LHCb built up to running over 45000 jobs a day and, importantly, involving all of the 116 sites that will support LHCb.

Finally Bird reported on ALICE who sustained a storage rate of 1.25 GBytes/second into CASTOR from their pit and performed data replication to their 6 Tier-1 sites. Between 11<sup>th</sup> and 18<sup>th</sup> of June they exercised all of their workflows but did not save the data sent to the Tier-1 in order to save on storage. He then showed some later information from ALICE showing a lead beam splash event and plots showing that they ran an average of 7500 concurrent jobs from June to October involving more than 200 different users.

I. Bird concluded his survey of STEP'09 and its follow up by saying that the LHCC review held in July had suggested a second STEP'09 exercise but this had been excluded by the existing individual experiment schedules.

I. Bird then moved on to WLCG reliability and availability in slides 17 to 20 reminding that reliability is measured by running jobs over the system. CERN and the Tier-1 ran well over this year with a few Tier-1 problems reducing their reliability over the summer months. For the Tier-2 sites only 60% have a reliability at the top end of the spectrum and WLCG has set an internal milestone to raise this up to the 90% level. At the last C-RRB Bird had promised to show reliabilities per experiment and these he showed on slide 18 adding that the data prior to March 2009 was not always reliable. It was important to note that the operations tests were not always a good indication of what the experiments see. Slide 19 showed the matching experiment site availability pictures where each square is a day and the colour shows where experiment SAM tests failed to various levels. Bird said it was sometimes hard to understand what is going on underneath in these measures but they are reviewed weekly and attempted to be correlated with known events. Slide 20 then showed recent incident reports which are now systematically produced for any outages resulting in service unavailability for a few hours or more. Many of the issues are due to everyday service problems and not necessarily with grid middleware and they are probably at an irreducible level that we must live with and manage around.

The next subject of Bird was mass storage, slide 21, where performance had been generally very good but there were dache and CASTOR upgrades that should be deployed to reach stable versions for data taking. Migration of dcache sites to a new namespace (Chimera) for better scalability was being encouraged while for CASTOR there was a consolidation version that would give better support for analysis. Finally he pointed out that the open SRM functionality issues had been revisited and only a few need to be addressed.

On slide 22 Bird discussed middleware reporting that there remain relatively few upgrades in the pipeline. He concluded that middleware is not really an issue now admitting he had never thought he would be able to say that.

Tier-0 status, slide 23, was his next subject. Decision on construction of a new computer centre in Preveessin has now been suspended while costing of container solutions is looked at and an attempt is being made to get a clearer understanding of the long-term computing requirements. Reinforcing the critical power supply to building 513 remains urgent independently of any new centre and investigations are also being made how to operate and manage remote hosting solutions.

I. Bird's final subject was the EGEE to EGI transition where he showed 4 slides. The two main aspects of this transition namely setting up the organisation, to be hosted in Amsterdam, and

responding to EU FP7 calls of which 3 are relevant to WLCG are both moving along. Within EGI there will be a transition task of service deployment for heavy users where it is hoped to get 2-3 FTE for each experiment for specific tasks though, he warned, these numbers may be cut. Under a specialised support centres call we were hoping for an additional level of effort for analysis support. Under a middleware support call we see for the first time a collaboration by all three European middleware providers. Bird's judgement was that all these activities can ensure what we need and we must make sure that it stays that way. On slides 26 and 27 he showed tables of the components and services required by WLCG and their attributions among sites and as EGI tasks.

His final slide summarised his conclusions: STEP'09 showed that we are ready to take data surprising some people with how well it went; business has been as usual since then with resources being heavily used; the Tier-0 strategy is evolving and the EGEE to EGI transition is encouraging but we must be prepared for possible disruption.

## Discussion

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

The first intervention was by Prof. Barreira (LIP, Portugal and chairman of the EGI Policy Board) who, referring to Bird's final remark on the potential disruption of the EGEE to EGI transition, recalled that it is the job of the High Energy Physics community to mobilise their national grid initiatives to create the conditions for non-disruption and this is our responsibility. Bird replied that he could not agree more adding that it was very important that the national initiatives really are understanding the need for them to continue to support the existing operations in the way that they are currently running.

Prof. Ekelof (Uppsala, Sweden) said that he had a worry with the separation of the European middleware initiative from the computer centres where there is a tendency for the centres to say that middleware is not their concern and that they would pick the middleware that suits them best. To his mind the centres should make a commitment to having the middleware sustained in a way connected to their machines and operations and not just separated out. Bird responded that he thought it was equally worrisome when the middleware was part of the operations project, for a different reason, and that we must draw the line somewhere and live with the consequences. He then responded to Prof. Ekelof querying where this line was by saying he thought it better to have the projects separated because the operations have to be free to choose the middleware they require and not be forced to take one that is part of a project. Prof. Ekelof then added that he was worried about the funding of the middleware project. Bird replied that this was an issue for the commission who had designed the calls this way and this was out of our hands. Given a free hand Bird would have put some of the middleware required by LCG in the service activities of the heavy users call (SA4) but there was a fear that if middleware appeared other than in a middleware project that it might be reviewed away by the commission.

V. Ilyin (JINR, Dubna) then raised the issue of a serious disruption between WLCG and EGI where the EGI consortium had decided that their MoU could only be signed by eligible European countries. Several countries participating in WLCG and represented here, Ukraine, Russia and some others cannot sign the EGI MoU and this a serious potential danger for our activity. Bird replied that this was correct and he apologised for not having mentioned it. It is something he was concerned about and in fact the EGEE project management board had sent a message to the EGI council that we need that the countries that the LCG relies upon should be part of this structure.

E. Auge (CNRS/IN2P3, France) then asked if he understood correctly that I. Bird was not much worried by the manpower problems next year in particular in connection with the end of the EGEE contract. Bird's reply was that if we really got funded at the level that was being proposed in the EGI projects then he would stop worrying. There would still be a transition of people with expertise to new people which is part of the concern over disruption but the overall level is probably going to be alright.

The next points were from V. Guelzow (DESY, Germany) who made the remark that one of the problems that is facing us with the EGI funding is that it has to be spread over four years instead of three besides not having enough money anyway. Also we have this nice OPN connecting all the Tier-1 but looking at the Tier-2 in his opinion a better connectivity was needed at least at the major Tier-2's. Bird's reply was that this had recently been discussed at an OPN meeting and also at the Grid Deployment Board where the OPN project leader came and asked for advice. He continued that the result was to ask the OPN community to look at general connectivity between all of the tiers. This has certainly been discussed in various networking communities as something they would like to do but now we have given a mandate to the OPN to address this. He agreed with Guelzow that one of the observations from STEP'09 was that connectivity to certain of the Tier-2s was not optimal. He added that for some of the Tier-2 the internal data rates that they have to support for analysis are probably higher than anticipated and we need to foresee some improvements internally in the networking as well.

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-2009-128

#### **E. Tsemelis, LHCC Scientific Secretary**

S. Bertolucci stated that, Tsemelis not being present, he would inform the delegates that the LHCC had reviewed, in cooperation with the Computing Scrutiny Group, the computing status and requests and as a result confirms the report given by I. Bird. The tendency he is now seeing is that the convergence between the requests of the experiments and what the LHCC and scrutiny group assess are the needs is quite good. There were no principal concerns except that the LHCC underlines the fact that the real test of the models will be when real data comes. The general impression from the LHCC is that the currently proposed level of hardware requests is correct as is the staging of them which has been slightly moved in time from the original schedule.

#### **5. Status of Common Project Accounts**

Paper CERN-RRB-2009-099

#### **T. Lagrange, CERN Finance Dept.**

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 on his side there was nothing to add to the document that had been submitted. S. Bertolucci closed this item and invited S. Foffano to present the next item.

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

Paper CERN-RRB-2009-129

Presentation CERN-RRB-2009-130

S. Foffano began by reviewing her agenda (slide 2) saying she would bring delegates up to date on the WLCG MoU status, look at funding and expenditure at CERN for the 2009-2013 period then continue with the regular resource usage and accounting reports. Her final subject would be resource requirements and pledges which is where she would spend most time. She had hoped to present final tables today but this turned out to not be possible. She would also be soliciting delegates input on how to improve the requirement/pledge process. She also encouraged delegates to read the written report and visit the LCG web pages which contain detailed tables but which she expected to change.

##### **6.1 WLCG MoU Signature Status**

S. Foffano reviewed (slides 3 and 4) the status of signatures of the WLCG MoU. Since the April meeting the state of Sao Paulo, Brazil, has signed the WLCG MoU supporting the CMS experiment as a Tier-2 site. A Brazil wide committee will coordinate participation of all other tier centres as a federation so she is expecting another MoU signature from that body. Some other countries have expressed interest in becoming Tier-2 centres on which she hopes to report in future and she will work in collaboration with the external relations unit of the DGs office on this follow-

up. On slide 4 she showed the current list of countries, both member and non-member state, who have signed the WLCG MoU with either Tier-1 or Tier-2 sites.

## **6.2 Funding and Expenditure for WLCG at CERN**

S. Foffano then showed in slides 5 and 6 WLCG budget estimates for 2009-2013 reminding that the book closing for Phase 2, which ended in 2008, had been given at the last CRRB. The personnel funding had changed little since the April report while the materials budget had changed largely due to the change in profiling of the Computer Centre project now compressing the previous 10-year funding into the three years 2010-2012. Once the Computer Centre strategy is confirmed the materials budget will be re-profiled accordingly and reported back to a future meeting.

## **6.3 CERN, Tier 1 and Tier 2 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 from January 2009 to August 2009 obtained by summing up CERN and all external Tier-1s. She remarked that the plots showed, on a monthly basis, how experiment use of the installed resources was creeping up to the 2009/10 pledge levels and that there was currently a lot of installation activity going on. She reminded of the agreed shift in the pledge installation for 2009 from April to September.

S. Foffano turned to accounting for Tier-2s (slides 9 to 11) showing a snapshot of the August status ordered by 2008 pledges and split into two plots for sites with less than and more than 750 KSi2K pledged. In many cases the results were good though some sites were not reporting while others had quiet months while they were installing new resources. She had been dialoguing with sites that were delivering few resources and would like to meet them face to face. She remarked that overall a lot of sites had been doing very well in compensation for the under-performing sites. Finally on this subject she reminded that the accounting is circulated each month with sites being invited to comment allowing one week before publishing to the WLCG website and she encouraged sites to regularly check their accounting reports.

## **6.4 Computing resource requirements and pledges**

Slide 13 summarised the recent pledge collection exercise pointing out that at the last CRRB meeting there had not been a convergence between the experiment requirements and the Scrutiny Group which resulted in a commitment to continue discussions over the summer, concentrating on 2010, and waiting for more precision on the LHC schedule. Sites were asked to give a status report on their 2009 and 2010 resources and confirm their 2010 pledges by the end of September 2009. At the end of this deadline input from two Tier-1, eight Tier-2 and formal confirmation from some official bodies was still outstanding. On 5 October ATLAS signalled a change to their 2010 Tier-1 requirements after many confirmed pledges for 2010 had been received and on 7 October the preliminary resource tables for this meeting were published. Slide 14 showed the 2009 pledge installation status classified into 3 main categories – installation complete or planned to be by the end of October; installation progressing well with a revised schedule on a time scale agreed with the experiments; installation delayed for procurement reasons but planned to complete by the end of 2009. S. Foffano said that she regarded 2009 as an exceptional year with a number of sites wanting to change their 2009 pledges following changed experiment requirements. She pointed out that the WLCG needed to be consistent and treat sites equitably following the spirit of the MoU adding that in fairness to sites that have installed their 2009 pledges changes could not be made to the already confirmed 2009 pledge data however explanatory footnotes could be added to the published pledge tables for those sites requesting changes. Slide 15 showed the preliminary percentage balance between 2010 pledges and the experiment requirements as known on the 7 October 2009. New experiment requirements were given at the beginning of September with some having a large increase in CERN resources. For some of the Tier-1 and Tier-2 sites she was awaiting official confirmation and while Russia has confirmed its global 2010 pledge it cannot confirm the experiment split until January. Globally Tier-1 cpu has an excess over requirements of 5%, Tier-1 disk is 5% below requirements, Tier-1 tape is 3% below and Tier-2 cpu and disk are

12% and 22% below requirements respectively. Finally under this heading S. Foffano discussed, on slide 16, how to improve the requirements and pledge process saying that in an ideal world experiment requirements would be discussed and approved by the LHCC and C-RSG bodies in February/March of each year then announced to funding agencies at the April C-RRB meeting with a set of fixed experiment requirements. For the 2010 exercise there will be a push to get experience of data taking for which the C-RSG is asking for experiment reports by March.

## 6.5 Conclusions

In conclusion (slide 17) S. Foffano summarised that:

- WLCG MoU: Brazil Tier-2 signed since the April CRRB and another federated Tier-2 is expected in future. Follow-up is on-going with other interested countries and the delegates will be kept informed.
- WLCG Funding & Expenditure: some future over-spending is estimated on personnel. Once the replacement Computer Centre strategy is finalised, which should be by the next meeting, materials funding and expenditure will be re-aligned accordingly.
- Monitoring of Tier-1 and Tier-2 accounting continues. 2009 pledge values start after September. Follow-up is on-going with Tier-2 sites having little or no accounting activity. Delegates should look carefully at their sites accounting.
- Recent pledge collection exercise is still not concluded. Rather than send a new version after each change this will be done when the exercise is formally concluded expecting that the definitive situation can only be published in December 2009 or January 2010.
- The process must change for the future to allow formal approval of pledge tables at the Autumn CRRB meeting.

S.Bertolucci thanked S. Foffano and suggested to hear the associated report from the Computing Resources Scrutiny Group before taking questions on both presentations.

## 6.6 Report from the C-RSG

### D.Espriu

Paper CERN- RRB-2009-079

D. Espriu began by reminding the delegates, on slide 1, of the mandate of the C-RSG to inform the decisions of the C-RRB for the LHC experiments adding that he would only be discussing the 2010 period and that a complete version of the scrutiny report was attached to the agenda. On slide 2 he reviewed the usage of installed resources showing that during the first six months of 2009 this had been at 56% for cpu resources, 75% in disk and 62% in tape for CERN and the Tier-1s. He added that the situation has since improved in that the middleware is more mature and more users are submitting jobs but more resources are still installed than are used. He also thought it was clear from S. Foffano's presentation that the 2009 resources were sufficient. He continued on slide 3 remarking that the few underperforming Tier-1's that had been identified during the STEP09 exercise were a source of worry to the scrutiny group as was the general underuse of resources. Analysis of resource usage was an important part of the mandate of the group but one that they had not been able to get into yet due to the changing accelerator schedules but Espriu hoped that this would change now that operations were starting. The scrutiny group was now asking experiments to provide a summary of their use of the computing resources made available to them by March 1<sup>st</sup> 2010 in preparation for the April CRRB report. He added that in granting resources for the following year it will be important to see how the experiments got on with real data.

On slide 4 Espriu reminded of the LHC schedule which had led the scrutiny group to recommend that the 2010 resources should exceptionally be in place by June 1<sup>st</sup> rather than April 1<sup>st</sup>. Espriu then showed on slide 5 the expected 2010 beam times of  $4.3 \times 10^6$  seconds/year of pp running and  $5 \times 10^5$  seconds/year of AA running. On slide 6 he reviewed the revised experiment requirements that were made public on April 7<sup>th</sup> 2009 pointing out that a number of changes in computing models or their parameters were proposed in some cases justifying the need for increased resources but that such changes had not been validated. After discussions with the experiments a number of changes were accepted by the C-RSG in view of the special characteristics of the run. Advice from the LHCC regarding modifications the group thought were

beyond its mandate was sought. He showed on slide 7 the questions submitted to the LHCC in April 2009, that the delegates had already seen, then on slide 8 the additional questions submitted in June 2009. The first question was that one of the main reasons for increased resource needs in the revised requests was for much more MonteCarlo simulation and the LHCC accepted this justification. They also approved an initial high trigger rate for ALICE. The third question related to the planned ATLAS and CMS event reprocessing rates which were not foreseen in the computing TDRs and which had a big impact on disk requirements and these rates were also approved. On other points the LHCC recommended that this was not the moment to attempt to do some savings. The final point was the different estimates of pileup between ATLAS and CMS for the same running conditions where the LHCC urged the collaborations to converge.

D.Espriu then moved on in slide 9 to say that the C-RSG was happy with the LHCC recommendations and, although the nature of the scrutiny carried out by the C-RSG is different to the one implemented by the LHCC reviews, the suggestion now was to hold joint meetings twice a year for experiments to present and justify their requests.

On slide 10 Espriu explained that the C-RSG was generally satisfied with the amount and quality of the information provided by the experiments this time adding that there was still room for improvement in communications. He then gave a list of requirements for future reviews including that experiment requirements and models be frozen during the review, all documents be provided sufficiently early, simplified spreadsheets accurate to within 10-20% should be provided and that a reasonably common format, especially in calendar years, be used for the experiment requests.

Espriu then reviewed, on slide 11, the most important points of disagreement with the experiments or items the group were not able to scrutinise. For ALICE the scrutiny recommends more tape at Tier-1 than they request. There is a good overall agreement with ATLAS after a number of revisions. There was a large discrepancy with the scrutiny in their Tier-1 cpu request for reprocessing that is now agreed and should lead to a further revision of the ATLAS request. With CMS and LHCb there is now a good agreement after implementing the recommendations of the LHCC.

Various general remarks were then made on slides 12 to 14. The group worried that the tendency to use available resources might lead to irreversible changes in the models leading to inefficient use of resources. The group felt strongly that the experience gained from the first data taking should be promptly fed into next years scrutiny process. This would give a reality check where the group was particularly worried about some cases of disk usage. The group recommended the experiments to make maximal use of the distributed resources in the grid.

On his final slide, slide 15, Espriu summarised that the scrutiny group had critically examined all possible aspects of the different computing models and their implementations. The group recognised the special characteristics of the combined 2009/2010 run and while they were still worried with some long term implications for the sustainability of the models they had reached a reasonable degree of convergence with the experiments. The group felt that updates to the models will be needed and that the scrutiny after the first round of real data will be of great relevance. The scrutiny group did not doubt that the experiment models will survive their first contact with real data in 2009 and wished them good luck.

## **Discussion**

S. Bertolucci then invited an open discussion on the last two reports. The first intervention came from M. Turala (IFJ PAN, Poland) who asked a question to S. Foffano on her table of the preliminary pledge balances for 2010. He wanted to know the uncertainties on these numbers since only two Tier-1 and 8 Tier-2 have not committed completely and she must have preliminary numbers from the previous year so what were the changes expected when these missing sites provide their information, adding that he expected they would not be too big. S. Foffano replied that she had now received from the two Tier-1 their numbers and for her the biggest changes risked to come from the bodies that do not go along with the proposal. She did not think the changes would be big but could not be certain. She added that the numbers she had received from

the French Tier-1 and Tier-2 at IN2P3 are not included in the balance that she showed and their numbers have reduced so there will be a change to the overall balance. Her overall message was that she would like to be able to say there was no problem for 2010 but she did not feel she had enough information to be able to confirm this. M. Turala then asked that, while appreciating that some official confirmations were needed, when he looked at the balance table he thought the numbers were almost right. S. Foffano replied that these numbers were as close as she could get and that as soon as she had the confirmed numbers they would be sent. She expected this to take several months as it depended on external committees but thought that the risk of resource problems for 2010 was small.

The next intervention was from F. Gianotti, ATLAS spokesperson, who apologized for the late change in the ATLAS request for Tier-1 cpu resources. This was due to a mistake in the calculation of the amount of cpu needed for data reprocessing and which was spotted by the scrutiny group and she thanked them for finding it. The change actually reduces the ATLAS request for cpu at Tier-1s which, she thought, would make the delegates happy. Jim Shank, the US ATLAS computing manager, then added his comment on this issue firstly apologizing to S. Foffano for giving her information so late and thanking her for doing a very nice job in difficult circumstances but he would like the minutes to reflect that the US Tier-2 pledges that are mentioned on the slide, referring to the preliminary pledge balance 2010, are actually an increase in the pledges that cover exactly the requirements for ATLAS in 2010.

There were no further points raised so S. Bertolucci proposed to pass on to his summary.

## **7. Summary**

### **S.Bertolucci**

S.Bertolucci began by saying that he thought the delegates would agree with him that the system to follow the very challenging and dynamic evolution we were experiencing is in place. He thanked very much the scrutiny group for their excellent work with more being prepared for the future. He thought that in general what was being done now was trying to avoid that experiments be damaged by any lack of computing power in the current delicate phase and at the same time his impression was that resources were not being wasted. These issues are ongoing and will always be with us but he thought all the tools were in place to follow them appropriately.

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 <b>Monday 19 to Wednesday 21 April 2010</b></p>
---

H.Renshall  
28 January 2010