

W L C G

Minutes of the 2nd COLLABORATION BOARD Meeting

(Held at CERN on 25th January 2007)

Present:

CERN CSO
 CERN IT Dept. Head
 CERN PH Deputy Dept. Head
 Scientific Secretary
 LCG Project Leader

J. Engelen
 W. Von Rueden
 M. Doser
 D. Jacobs
 L. Robertson

LHC Experiment Spokespersons

ATLAS
 CMS
 LHCb

D. Barberis (for P. Jenni)
 T. Virdee
 P. Charpentier (for T. Nakada)

International Membership

Austria, Austrian Tier-2 Federation
 Belgium, Belgian Tier-2 Federation
 Canada, TRIUMF
 Czech Rep., FZU AS, Prague
 France, CC-IN2P3
 France, CC-IN2P3 AF
 France, GRIF, Paris
 France, SUBATECH, Nantes
 Germany, ATLAS Federation, Munich
 Germany, DESY, Hamburg
 Germany, CMS Federation
 Germany, GSI, Darmstadt
 Germany, GridKa
 Italy, INFN Tier-2 Federation
 Japan, ICEPP, Tokyo
 Netherlands, NIKHEF/SARA
 Pakistan, Pakistan Tier-2 Federation
 Poland, Polish Tier-2 Federation
 Russian Fed., Russian Data-Intensive GRID (RDIG)
 Spain, ATLAS Federation
 Spain, CMS Federation
 Spain, LHCb Federation
 Spain, PIC
 Switzerland, CSCS
 Taipei, ASGC
 Taipei, Taiwan AF Federation
 UK, London Tier-2
 UK, NorthGrid
 UK, RAL
 UK, SouthGrid
 USA, BNL
 USA, Boston/Harvard ATLAS T2 Federation
 USA, FNAL

C.E. Wulz
 G. Bruno
 M. Vetterli (*phone*)
 M. Lokajicek
 F. Hernandez (for F. Malek)
 J.P. Meyer
 L. Aphecetche
 G. Duckeck (for S. Bethke)
 P. Fuhrmann (for V. Gülzow)
 K. Schwarz (for P. Malzacher)
 K.-P. Mickel
 L. Perini & M. Paganoni (*phone*) for
 P. Capiluppi
 H. Sakamoto
 K. Bos, J. Templon
 H. Hoorani
 R. Gokieli.
 V. Ilyin
 J. Salt
 F. Matorras (*phone*)
 R. Graziani Diaz (*phone*)
 G. Merino (for M. Delfino)
 D. Feichtinger (for M-C. Sawley)
 Hurng-Chun Lee (for S. Lin)
 D. Colling
 R. Jones
 N. Geddes
 J. Tseng
 B. Gibbard (*phone*)
 J. Shank
 V. White, P. McBride

Speakers

LCG Resource Coordinator	C. Eck
CERN-IT Communication Systems Group Leader	D. Foster
Joint Security Policy Group	D. Kelsey

Excused

India, TIFR, Mumbai	A. Gurtu
USA, Caltech CMS T2	H. Newman

Absent

ALICE	J. Schukraft
Australia, Univ. of Melbourne	
China, IHEP, Beijing	G. Chen
France, LPC, Clermont-Ferrand	D. Pallin
Italy, CNAF	M. Mazzucato
India, VECC/SINP, Kolkata	Y.P. Viyogi
Nordic Data Grid Facility (NDGF)	B. Vinter
Portugal, LIP T2	J. Gomes
Romania, Romanian Tier-2 Federation	M. Dulea
UK, Scotgrid	N. Glover
USA, Midwest ATLAS T2 Federation	R. Gardner
USA, Southwest ATLAS T2 Federation	K. De
USA, Florida CMS T2	P. Avery
USA, MIT CMS T2	C. Paus
USA, UNL Nebraska CMS T2	D. Swanson
USA, Purdue CMS T2	N. Neumeister
USA, UC San Diego CMS T2	F. Wuerthwein
USA, U. Wisconsin CMS T2	S. Dasu

2nd Meeting of the WLCG Collaboration Board (CB), 25 January 2007

Documents can be found via: [WLCG CB Home Page](#)

1. Introduction

CB Chair N. Geddes

N. Geddes welcomes those present, recalling that although this Board is the senior body of the WLCG Collaboration, most of the executive roles have been delegated to other committees such as the Overview Board (OB), Management Board (MB), etc.. Short reports from these bodies are on the agenda of this meeting, as is a review of the known capacity plans. In the context of the last point, he counts particularly on the Tier-2 representatives present to point out any funding difficulties they have.

This meeting comes at the end of a WLCG workshop that attracted some 270 people from 27 countries – a remarkably strong attendance.

2. Summary of OB meetings

OB Secretary D. Jacobs

D. Jacobs begins this summary of the first year of OB activities by recalling that all of the documents related to OB meetings are available on the [OB Web Site](#). The OB is a standing sub-committee of the CB and as such is one of the strategic committees of WLCG, as against the executive committees such as the MB, GDB etc. Its membership comprises representatives of the managements of the LHC experiments, LCG and the Tier-1 centres, and its role as defined in the WLCG MoU is to undertake technical discussion on strategy, policy and conflict resolution. At its first meeting, the OB additionally decided that points for its consideration and decision should be presented well ahead of the meetings. This has proven hard to implement. In view of its membership, the OB also acknowledged that its decisions could not be binding on Tier-2's without prior consultation.

In 2006 the OB met on 20 March, 12 June and 11 September. A meeting was planned for 18 December but has been delayed to 29 January. All meetings heard project status reports illustrating the steady progress being made but also allowing many issues and problems to be flagged. Beyond this, some particular themes were followed through the meetings:

- Use of VO-boxes

It was rapidly agreed that VO-boxes were to be deprecated but tolerated at least at the Tier-1's. It was noted that it would be difficult for many Tier-2's to install them due to the range of activities at these centres. Subsequently, a series of dedicated meetings led by C. Loomis allowed the identification of two classes of VO-box service: (1) those that can be tolerated and (2) those that must be eliminated as soon as possible. Two Class 2 services were identified: Package management (ALICE) and SRM functionality (CMS). Since then the matter has been left to the LCG executive structure to handle.

- Accounting

It was recognised at the first meeting that it was urgent to implement accounting at the VO level, extending later to the user level. Monthly figures were becoming available by June and in September the OB approved the figures to be shown to the October C-RRB, by then systematically covering the Tier-0 and Tier-1's up to August, in terms of installed capacity.

- Tier-1 – Tier-2 dependencies

The OB took the view that arranging these dependencies was mainly homework for the experiments and in June all four experiments reported, also explaining their data models. Their views ranged however from a wish for no defined dependencies (CMS) to specific associations forming a radial hierarchy (ATLAS). The project underlined the importance of taking proper account of the realities of network connectivity. At the same time, the need to separate service issues from those of data storage and transfer was recognised. In the discussion, the Tier-1's emphasised that, in making their capacity plans, they were not counting on any timesharing amongst the experiments. The OB asked the MB to make a workable technical implementation of associations for its next meeting. At the third meeting, the experiments reported on the data transfer and storage issues, revealing the load on the Tier-1's to be very uneven, along with a real capacity shortage for ALICE. The task of further follow-up was passed to the MB.

- Capacity ramp-up

The OB has been following the revision of the LHC schedule, while noting that centres could only realise economies from shifting purchase schedules by 6 months or more. By September the experiments were updating their requirements in the light of expected running in 2007 and 2008, with the result that the CERN capacity over this period appears now to be just 10% short of needs. It is thus assumed that the experiment requirements at CERN can be compressed to fit with the presently available funding.

In conclusion, D. Jacobs lists the agenda points for the forthcoming OB meeting.

He confirms to J. Tseng that, despite the implied economies in capacity if experiments timeshare, there is no pressure for ATLAS and CMS to do this.

3. Report from the MB

Project Leader L. Robertson

L. Robertson emphasises that he is presenting a summary of MB activities rather than the details. For those who wish more information there are the quarterly reports and the status reports to the OB. Showing an LCG organizational diagram, he points out the executive committees – the MB and, reporting to it, the Architects Forum (development and maintenance of common applications software) and the GDB (service aspects – the coordination of Grid operation). The GDB interacts with all sites, not just the Tier-1's.

The main role of the MB, which L. Robertson chairs, is to make sure that the project is running. As such it follows the fluctuating priorities, monitors service levels and use of resources, and undertakes corrective action where needed. It meets by phone weekly and face-to-face every month. EGEE and OSG representatives participate in the meetings as well as the project Management, the experiment Computing Coordinators and one person from the Tier-0 and each of the Tier-1's. Tier-2's are represented in the project through their membership of the CB, participation in the GDB, operationally through the grid infrastructures and more directly to the MB via their associated Tier-1's, the GDB chair or himself. L. Robertson underlines the importance of Tier-2's expressing issues as forcefully as required. They should not feel inhibited by the formal structure.

MoU commitments have been made not just in terms of capacity but also service levels. The TDR's define how the resources are used and the inter-site relationships. This information is extracted into the "MegaTable" maintained by C. Eck.

When all of this has been implemented, the MB must oversee the operation and so suitable metrics are being set-up. The Site Availability Monitoring (SAM) framework

measures reliability at the EGEE sites, along with BNL and FNAL. If a defined subset of the tests all work then the site is marked as “available”. This is clearly simplistic and more sophisticated experiment-dependent tests are desirable. One must recognise that these tests measure grid availability – a site may still be running (e.g.) Monte-Carlo work even if it appears down to the grid. There must also be a definition of what is needed to declare a site as “down”. OSG and NDGF will report into this framework. The monitoring plots illustrate the wide spread in measured availability of sites. Other metrics are being studied, e.g. resource provision & usage accounting, network quality & data throughput, and reliability.

Monthly accounting figures are now available for Tier-0 and Tier-1 storage (disk and tape), cpu capacity and usage, along with the relationship of installed and pledged figures. The framework is as yet manual and must now be automated. CERN’s plan is to use the APEL accounting repository at RAL. This will allow all EGEE sites as well as OSG and NDGF to insert data. An experiment-level view is also required and should be available in a few months. During 2007 the scheme will be extended to Tier-2’s and it is hoped to add user-level accounting at many sites during the year.

GridView is the tool being developed to measure network throughput and quality. It will also report job status and service availability.

Job reliability has received much attention from the ARDA team, with the main aim of identifying the reasons for job failure. It is not easy to decide the performance level that is “good” and hard to analyse the job logs and determine failure reasons. The work has nevertheless resulted in many good operations tools for sites to investigate failures.

Turning to an outline of the WLCG commissioning schedule, L. Robertson points to the residual services presently being introduced – full FTS services, the 3D distributed database for ATLAS and LHCb, SRM 2.2 (enables management of the various required storage classes), VOMS roles for site scheduling. In addition, SL(C)4 (Scientific Linux 4) must be installed everywhere, along with the latest gLite. All of this is, however, just routine business. The goal is to commission by July the capacity needed to handle first collisions in November. This is partially to accommodate holiday plans and also to allow the experiments time to get their software running. L. Robertson underlines that, if first collisions are delayed, there is a real possibility of extending the run through the end of the year. Sites will thus need to plan flexible coverage over the end of the year both for this eventuality and for the possibility that the experiments are busy processing first data that they have already collected.

In conclusion, L. Robertson reminds members of the various sources of information on the [LCG Web](#), including the planning page and those for the MB and GDB. In addition, there is a bi-weekly news bulletin (mail A. Aimar to subscribe to this).

N. Geddes underlines the message that accounting and monitoring information for the Tier-2’s will become necessary over the next year.

4. Report from the GDB and outstanding issues *GDB Chair K. Bos*

Showing the GDB mandate, K. Bos stresses its role as a platform for Tier centres to make their views felt. All members have defined deputies but the formal structure is rarely used and the meetings (monthly) are essentially open. Any decisions are passed to the MB for endorsement. He shows the dates of the 2007 meetings, mostly at CERN. The Vancouver meeting coincides with CHEP. Each meeting is normally preceded by a pre-GDB meeting. Recalling the GDB’s position in the LCG organisational chart, K. Bos points out that all of the points mentioned by L. Robertson were already discussed in the GDB.

In view of the large size of the GDB itself, use is made of several working-groups and task forces. Security was handled in one such working-group, which has now grown into the independent Joint Security Policy Group (nevertheless still reports to the GDB). The situation is similar for networking. The efforts of the working-group on installation and configuration tools (Quattor) are perhaps less visible. More recently, a working-group for storage deployment has been formed. This will focus on SRM 2.2.

In conclusion, K. Bos shows a list of examples of the topics handled by the GDB and notes that it will be necessary in February to select a new GDB chair. During his time as chair, he considers that the GDB has become more open and has focused more on discussion and on site issues. Nevertheless, he would have liked more Tier-2 representation and better synchronisation with the MB than has so far been possible.

N. Geddes asks those present to comment on the lower than expected attendance of Tier-2's at GDB meetings. D. Colling responds that this is the first time he hears that the GDB is open to Tier-2's. He will now endeavour to attend. For US-ATLAS, J. Shank comments that he feels they already get all the input necessary via their formal GDB representation.

5. Privacy and policies

D. Kelsey – Joint Security Policy Group

Addressing first the security policies, D. Kelsey emphasises that the policies adopted are important for all users. The JSPG started life in LCG but has now grown into a joint body that also involves EGEE, OSG and NDGF. It is intended that the policy documents should apply across all grid infrastructures, without overriding local site policies (although this is sometimes necessary). Many people are involved in the JSPG, which produces and maintains documents for the approval of the management bodies of the various projects concerned. The document set is structured with the Security and Availability Policy at the top level. Under this are links to various Acceptable Use Policies etc., which D. Kelsey lists. The current versions are held at <http://cern.ch/proj-lcg-security/documents.html>.

N. Geddes asks about the consultation process for these documents, since they set out policies at the top level of the Collaboration. D. Kelsey responds that at its meetings the JSPG consults with the site managers and then reports to the GDB. It can therefore be hoped that most issues are sorted out by the time the matter gets to the MB. R. Jones comments that this procedure is satisfactory if no significant issues arise. If they did, they would have to be brought to the CB. The Tier-2's expect to be kept informed.

D. Kelsey draws attention to the present policy revision, contained in a new document entitled "[Agreement on Grid Site Operations](#)". This started life in EGEE-II as an operational policy document and is intended to apply to all Grid Sites (they will have to sign-up to it). The work started in August 2006 and the document is hopefully now near-final. He emphasises, however, that all sites should read the draft and comment rapidly if they wish.

N. Geddes gets confirmation that only a few of the sites present did not know about this development. The question remains for those absent. K. Bos adds, however, that there is a mechanism in place to distribute the information via the GDB. N. Geddes emphasises that it is important for sites to realise what they are signing up to.

D. Kelsey points to the revision of another important document originally called "[Grid Security and Availability](#)". This has reached version 5.4 and has been re-named "Grid Security Policy". It has become less LCG-specific and is more general and simple. It should be ready for approval in 2-3 months and meanwhile all feedback is strongly encouraged.

T. Virdee asks about verification that sites are implementing these policies. D. Kelsey replies that site audits are not foreseen but self-audits are required. Future evolution will be guided by experience. T. Virdee worries about possible damage to the Grid but D. Kelsey responds that the documents do address sanctions etc.. N. Geddes questions the value of signing to accept expulsion for misbehaviour if in fact there is no policing but L. Robertson points out that there are security professionals at all of the sites. He is confident that this body is doing the right thing. D. Kelsey points out that the documents are not written to be legally binding, being more in the style of Collaboration MoU's, but the JSPG is taking useful legal advice on the drafting. J. Templon asks about the Security Service Challenges and D. Kelsey agrees that these do constitute auditing of a kind.

Moving on to privacy matters, D. Kelsey notes that both operations and VO managers require accounting, auditing, logging and monitoring – at the individual user level. So far the JSPG has been concentrating on accounting. There are legal issues with this, starting with EU directives and national laws, and so informed user consent is needed. There are also cross-border issues. A user must have the right to see and correct his/her data. All of this creates a very complex situation and so many sites are unwilling to supply data. The question of informed consent is addressed by the Grid Acceptable Use Policy document, which users must sign. Hopefully this, along with a policy document on personal data management, will be enough to convince sites to grant access to the appropriate logs. While the WLCG management requires accounting for all participating sites for all jobs run by members of the LHC VO's, the LHC VO managers want more detailed and timely information down to the individual user level. The JSPG has the basis of a policy document to address all of this and at least a technical solution for the encryption of accounting data is available. The policy document must describe all aspects of the storage and access of data, and is now being treated with top priority.

In the future, two further areas must be addressed for which no documents exist – VO agreements and agreement on Grid Services operations (perhaps as a single document). The next face-to-face JSPG meeting on 13/14 March will work on these topics.

Lastly D. Kelsey draws the attention of CB members to the JSPG Web site at <http://proj-lcg-security.web.cern.ch/> and the agendas and material for meetings at <http://indico.cern.ch/categoryDisplay.py?categId=68>.

It was agreed that the CB mailing list be used to disseminate information and requests for input on these policy documents to ensure that all sites are contacted.

6. OPN Status

D. Foster – CERN-IT/CS Group Leader

Showing the complete network layout at the CERN computer centre, D. Foster notes that the WAN uses many different infrastructures: NREN's worldwide, the interconnection of NREN's in Europe (GEANT2), USLHCnet for transatlantic links and the LHC Optical Private Network (LHCOPN). Neither the management nor the funding is centralised.

LHCOPN was started in 2004 for the purpose of data distribution from CERN and GEANT2 has evolved with OPN needs in mind. The results can be viewed as a considerable success. The 2004 LHCOPN topology was simple, comprising a star of 10Gbps links from CERN to the Tier-1's in order to get the data out. It is now developing more as a mesh, partially with the aim of providing redundancy. In 2004 it was not known how the Tier-2 connectivity would be implemented, although the feeling was that they would connect via the general-purpose research networks, i.e. independently of the circuit-based OPN.

Showing the geographic layout of GEANT2, D. Foster points to the heavy black lines marking the 10Gbps trunks and the position of Switzerland as a key hub. The layout allows additional connectivity for low extra cost. The project depends on EC support, with DANTE acting as the commercial contracting body. Policy is decided in the NREN Policy Group.

The end-end picture of LHCOPN shows the connectivity of infrastructures and management domains. The local loops to FNAL and BNL are provided by ESnet. The mission of USLHCnet, which provides the transatlantic links, is broad, extending much beyond the immediate needs of LHC. Under the planned management layout, there will be additional transatlantic connectivity and a management scheme to divide-up links according to the Tier-1 traffic. Adding Tier-2's will bring additional changes. There will also be peering between ESnet and the NREN's.

It is expected that on the 2007-8 timescale, LHCOPN will be able to support all of the requirements for Tier-0 - Tier-1 traffic and probably all of the Tier-1 - Tier-1 traffic (if Europe and the US are looked at separately). There is already ~25Gbps of peering available between ESnet/I2 and GEANT2 but so far this is little used. D. Foster thus sees no network performance issues for Tier-2 - Tier-1 connectivity. At the same time, he cautions that the existence of a 1Gbps connection locally does not guarantee 1Gbps end-end, due to possible intervening general IP infrastructures, and only does so for dedicated links. The network requirements in the LCG "MegaTable" must also be treated with caution since they seem to be a mixture of peak and average performance figures, it is not clear if there is any headroom for failures etc. and the assumed patterns of data movement may be simplistic. Thus although the figures are good order-of-magnitude requirement estimates, future network provision will need more detailed prediction of usage. That said, all of the Tier-0 - Tier-1 and Tier-1 - Tier-1 traffic seems to fit well within the 10Gbps available each way. For Tier-1 - Tier-2 traffic no real problems are visible, although some sites need more than 1Gbps. There appear thus to be no striking problems, although Tier-1 sites should look at general-purpose connectivity to make sure that they can sustain the aggregate rates foreseen.

Concerning longer term evolution, the provisioning of additional circuits should be possible at marginal cost both in Europe and the US, but overall transatlantic capacity is getting tight and is likely to remain expensive (~300kUSD/year at present). The same risks to be true for other remote links (e.g. circuits to India are currently full and so 1Gbps would cost ~200kUSD/month). Caution is also needed regarding GEANT cost-sharing and acceptable use policy - the actual costs are cross-subsidised from the membership fee, an arrangement that gives rise to sensitivity about traffic going outside the GEANT "club". LHCOPN already suffers from this but is so far tolerated by the GEANT community. Additional circuits amongst the Tier-2's and Tier-1's will require real additional funding.

LHCOPN organisation is steered by quarterly meetings, organised as a sub-activity of the GDB. Below this come a number of working groups, whose roles evolve with time. Nearly all the required infrastructure for LHCOPN is now in place and the same should be true by mid-year for all of the planned Cross-Border Fibre (CBF) - being put in place by the European NREN's independently of GEANT. Although several links are already in production, effort will still be needed to create coherent operational management across organisational domains. All operational information is documented at the [LHCOPN Twiki](#).

Speaking lastly of present issues and activities, D. Foster notes that the provision of backup in case of failure is still a major issue, since many fibres often occupy the same physical trench. A risk analysis is under way, which DANTE is leading. Operational procedures are still being refined and further capacity planning will clearly depend on actual use patterns. The requirements for the US-ALICE Tier-1 must also be addressed.

J. Engelen asks what WLCG is paying in total for networking. D. Foster replies that he has individual numbers and understands all for the Tier-1's. A European 10Gbps link costs 40kEUR/y. He stresses, however, that WLCG is relying on developments in the countries and is not making direct investments. The NREN's pay about 2MEUR/y to be in the GEANT club. The way in which costs are reflected back to the institutes is very variable.

7. Revised capacity plans

C. Eck

The capacity plan revision was triggered by the [new estimates](#) made in July 2006 for the LHC start-up schedule and subsequent running in 2007-2010. In the light of this, the experiments came forward with [revised computing requirements](#) in October and the C-RRB that month called for new capacity plans to be presented in April 2007. Overall capacity pledges seem now to fit global needs out to 2009 but the time profile, type of equipment and allocation to individual experiments does not fit so well. Revising the time profile with incurring loss of budget will be delicate, as will the optimisation of allocations amongst the experiments. It is now urgent to resolve these issues before the April C-RRB meeting.

Regarding types of equipment, the main issues are with the Tier-2's, as the situation at the Tier-1's is better understood. The key issue is that there must be enough disk at the Tier-2's to allow efficient use of the pledged CPU capacity. This is not a question of spending more money, rather just a change of emphasis. In terms of TB/kSI2k, the Tier-2 disk/CPU requirements of the experiments are: ALICE 0.40; ATLAS 0.44; CMS 0.22; LHCb 0. The anomalous LHCb figure is due to their plan to do only simulations at Tier-2's. At many sites the current figures are very low compared with the ALICE, ATLAS and CMS needs.

Before inviting discussion, C. Eck points to the location of the "MegaTable" at <http://lcg.web.cern.ch/LCG/planning/planning.html> under "Technical Design Reports and Experiment Resource Requirements". The table has a sheet for each Tier-1, which shows (as well as their own pledges) the known plans and future hopes of all the Tier-2's. Alongside is shown the network traffic and storage at the Tier-1 for each particular Tier-2 and experiment. The table is regularly updated.

N. Geddes comments that the table was originally put together by the experiment representatives. He underlines that the question now is to identify the constraints on Tier-2 planning to meet the needed ratios for disk/CPU, at least for the experiments that they are committed to support. He opens the floor to discussion, noting that he would like to get a feeling for possible issues and for what is needed by way of follow-up.

R. Jones, supported by J. Tseng, points out that the UK funding cycle is such that funding applications are already now being made in which the revised experiment requirements have been taken into account. Funding will not, however, be allocated annually but rather as a fixed amount for four years, somewhat limiting flexibility.

L. Perini reports about the INFN Tier-2's : they have an annual funding cycle. 2007 will be a difficult year but the Tier-2's will be funded, taking into account the disk/CPU ratio needed by the experiments (the Tier-2 funding will be allocated on a per-experiment basis). The funding for 2006 was received in September 2006 and has all been spent.

J. Shank notes that the US-ATLAS Tier-2 Federation is following the ATLAS model out to 2010, after which time there may be problems.

For the UK London Tier-2, D. Colling re-iterates that situation regarding UK funding and emphasises the need to see good use made of the capacity available.

R. Gokieli points out that Poland is about to sign the WLCG MoU. A detailed list of needs has been prepared and there are preliminary indications that these will be funded by

the Ministry. The situation will hopefully be satisfactory for 2007-8, in which period it should be possible to exceed the figures presently in the MegaTable. Thereafter it is less clear and will depend on what has been achieved up to that time.

G. Merino says that PIC Barcelona will have no difficulty meeting the required ratios.

For the DESY Tier-2, P. Fuhrmann also is confident that the needed ratios can be met.

The same is true for the ALICE Tier-2 at GSI Darmstadt (K. Schwarz) and the Belgian Tier-2 Federation (G. Bruno).

V. Ilyin recalls that the disk/CPU ratio for RDIG is still just 0.1. It is planned from this year to improve the ratio to what is needed. The budget request for 2007 is progressing well and there are good hopes for 2008.

F. Hernandez is not in a position to make a statement for all of the French Tier-2's. For CC-IN2P3 (Tier-1 and AF), given that experiment representatives were involved in the planning process for the resource deployment at the site, the requested ratios should be satisfactory. Nevertheless, he will check to make sure that they are.

For GRIF, Paris, J.P. Meyer reports that the disk/CPU ratio is satisfactory. The overall financial situation is satisfactory for 2007 but 1MEUR is still missing for 2008.

L. Aphecetche notes that the disk/CPU ratio for SUBATECH, Nantes (an ALICE Tier-2) is presently just 0.1. He thinks it will be difficult to reach 0.4 even if only disks are purchased from now on. The situation for funding after 2008 is unclear.

H. Sakamoto observes that funding for the Tokyo Tier-2 is allocated on a 3 year cycle, within which there is no possibility of modification. There is now more disk than expected but the disk/CPU ratio is presently only 0.2 and it is not clear if 0.4 can now be achieved in this cycle.

For GridKa, K.-P. Mickel reports that the 2007-8 funding is assured and that for 2009-2012 is being discussed. There is no problem with the required disk/CPU ratio.

M. Lokajicek stresses that the funding for the Czech Tier-2's is still not allocated. If it becomes available, there will be no difficulty meeting the required disk/CPU ratio.

For the Swiss Tier-2, D. Feichtinger says that the 2007-8 funding is satisfactory and the required disk/CPU ratio can be met.

K. Bos recalls that NIKHEF/SARA meets the desired disk/CPU ratio by design. He reminds the meeting, however, that a purchase cycle takes about 6 months and that effort is needed to convince funding agencies to allow money to be spent as required.

M. Vetterli reports that the situation for Canada is relatively easy since only ATLAS is served. The ATLAS requirements are being followed and both the Tier-1 and the Tier-2's are now funded through to 2010.

Concerning the ATLAS Tier-2 Federation (Munich + other candidates), G. Duckeck says that the Munich funding is satisfactory and the required disk/CPU ratio is being met. For the other candidate centres the basic installations have been made for 2007 and there are good prospects of longer-term funding.

For the Spanish LHCb Tier-2 Federation, R. Graziani Diez reports that there are no constraints on disk/CPU ratio.

At the conclusion of this round of the table, J. Engelen remarks that he feels he is getting the message from centres that they are taking a constructive and flexible attitude to the requirements, and that the situation is less dramatic than was presented by C. Eck. The

latter encourages sites to update their MegaTable entries accordingly but R. Jones warns of the need to distinguish between pledges (money available and acquisition details settled) and intentions that are still conditional on funding.

N. Geddes agrees with J. Engelen's view but warns the experiments about multi-year funding constraints. He hopes therefore that the mismatches pointed to by C. Eck are mainly due to a timing issue with the update of the MegaTable. At the same time he notes that some Tier-2's are not represented at the meeting, although there is good representation of the overall capacity.

C. Eck asks sites to check with the experiments before submitting numbers to the MegaTable. The meeting advises C. Eck to use the CB mailing list for all related correspondence.

8. Summary and Future Activities

N. Geddes

N. Geddes considers that the meeting has reflected the effectiveness of the executive structure of the project. From the discussions he has gained the impression that things are going well. No serious problems have been identified. He therefore concludes that the CB does not need to meet more frequently than at present.

R. Jones points out that more has been learned about the accessibility of the GDB. He is concerned by the poor attendance of Tier-2's at GDB meetings. With K. Bos now coming to the end of his mandate, he also wonders if the CB might consider appointing the next GDB chair. N. Geddes is not in favour of this suggestion, pointing out that, although WLCG is set up like an experiment, its CB differs from their's since the project has the unique function of delivering computing capacity. The CB thus only interferes in case of problems. The GDB has historically elected its own chair and he proposes that this should continue to be so. It is so agreed.

N. Geddes underlines the importance of the security and privacy policy documents, which should be circulated to the CB.

8.1. Election of next chair

N. Geddes assumes that the next CB meeting will be the last that he will chair. He will therefore set up a search committee of 3-4 people to identify candidates to take over this role and invites volunteers prepared to serve on a search committee to contact him.

8.2. Date of Next Meeting

N. Geddes recalls that opinion was split last year on the timing of the CB meeting, resulting in it finally being attached to the WLCG workshop. There were also split views about associating it with a conference such as CHEP. He presently does not see a need to plan for a further full meeting before January 2008.

8.3. Any Other Business

There being no other business, the Chairman thanks participants once more and closes the meeting.