

WLCG

Minutes of the 5th Collaboration Board Meeting

(Held at CERN on 13 November 2009)

Present:

CERN IT Head
 CERN PH Head
 Scientific Secretary
 LCG Project Leader

J. Shiers (for F. Hemmer)
 L. Mapelli (for Ph. Bloch)
 J. de Groot
 I. Bird

LHC Experiment Spokespersons:

ALICE Spokesperson
 ATLAS Spokesperson
 CMS Spokesperson

J. Schukraft, L. Betev
 F. Gianotti, D. Barberis
 G. Tonelli (for J. Virdee)
 Ph. Charpentier (for A. Golutvin),
 A. Schopper

LHCb Spokesperson

International Membership:

Canada, Simon Fraser University / TRIUMF, CB Chair
 Czech Rep., FZU AS, Prague
 Finland, NDGF/ HIP Tier-2
 France, CC-IN2P3
 France, GRIF, Paris

M. Vetterli
 M. Lokajicek
 D. Riska
 R. Rumler
 M. Jouvin, J. Meyer
 V. Guelzow (by phone),
 P. Fuhrmann (by phone)
 K. Mickel
 L. Levinson (by phone)
 L. Dell'Agnello (by phone for
 M. Maserà)
 G. Carlino
 H. Sakamoto
 M. Bouhuis
 L. Nixon
 R. Gokieli
 G. Stoica
 J. Salt
 M. Delfino Reznicek (by phone)
 C. Grab
 Q. Gang, J. W. Huang
 R. Jones (by phone)
 N. Geddes
 J. Gordon
 M. Ernst (by phone)

Germany, DESY, Hamburg
 Germany, FZK-GridKa
 Israel, ICHEP

Italy, INFN ALICE Federation
 Italy, INFN ATLAS Federation
 Japan, ICEPP, Tokyo
 Netherlands, LHC/ Tier 1
 Nordic Data Grid Facility (NDGF)
 Poland, Polish Tier-2 Federation
 Romania, Romanian Tier-2 Federation
 Spain, ATLAS Federation
 Spain, PIC
 Switzerland, CHIP
 Taipei, ASGC
 UK, NorthGrid
 UK, RAL
 UK, Southgrid T2 Federation
 USA, BNL

Invited:

EMI Project Director
OPN Convener

A. Di Meglio
W. Salter

Absent:

Australia, University of Melbourne	T. Dyce
Austria, Austrian Tier-2 Federation	D. Kuhn
Austria, Austrian Tier-2 Federation	C. Wulz
Belgium, Belgian Tier-2 Federation	G. Bruno
Belgium, Belgian Tier-2 Federation (UA, Antwerpen)	O. Devroede
Belgium, Belgian Tier-2 Federation (UCL, Louvain-la-Neuve)	P. Vanlaer
Canada, Eastern Tier-2 Federation	P. Savard
CERN DRSC	S. Bertolucci
China - Tier2, IHEP Beijing	G. Chen
Estonia, NICPB	M. Kadastik
France, LAPP, Annecy	S. Jezequel
France, LPC, Clermont-Ferrand	D. Pallin
France, SUBATECH, Nantes	L. Aphenetche
Germany, ATLAS Federation, FR/ W	T. Hareberg
Germany, ATLAS Federation, FR/ W	J. Sundermann
Germany, ATLAS Federation, Munich	S. Bethke
Germany, GSI, Darmstadt	P. Malzacher
Hungary, HGCC Federation	G. Vesztergombi
Hungary, HGCC Federation (KFKI-RMKI, Budapest)	D. Horvath
Hungary, HGCC Federation (SzTAKI, Budapest)	C. Hajdu
India, TIFR, Mumbai	A. Gurtu
India, VECC/ SINP, Kolkata	Y. Viyogi
Italy, CNAF	M. Mazzucato
Italy, INFN CMS Federation	M. Paganoni
Italy, INFN LHCb Federation	U. Marconi
Norway, UNINETT SIGMA Tier-2	J. Koster
Pakistan, Pakistan Tier-2 Federation	H. Hoorani
Republic of Korea, KISTI, Daejeon	S. Hwang
Russian Fed., Russian Data-Intensive GRID (RDIG)	V. Ilyin
Slovenia, SIGNET	B. Kersevan
Spain, CMS Federation	F. Matorras
Spain, LHCb Federation	R. Graciani Diaz
Sweden, SNIC Tier-2	S. Holmgren
Turkey, Turkish Tier-2 Federation (TAEK)	L. Baskus
Turkey, Turkish Tier-2 Federation (TAEK)	İ. Cakir
Turkey, Turkish Tier-2 Federation (ULAKBIM)	B. Akcan
Turkey, Turkish Tier-2 Federation (ULAKBIM)	B. Ortakaya
UK, London Tier-2	D. Colling
UK, Scotgrid	N. Glover
UK, SouthGrid	P. Watkins
Ukraine, Ukrainian Tier-2 Federation	G. Zinovjev
USA, Caltech CMS T2	H. Newman
USA, Florida CMS T2	P. Avery
USA, FNAL	V. White
USA, Great Lakes ATLAS T2 Federation	S. Mckee
USA, Midwest ATLAS T2 Federation	R. Gardner
USA, MIT CMS T2	C. Paus
USA, Nebraska CMS T2	K. Bloom
USA, Northeast ATLAS T2 Federation	J. Shank
USA, Purdue CMS T2	N. Neumeister

USA, SLAC ATLAS T2
USA, Southwest ATLAS T2 Federation
USA, U. Wisconsin CMS T2
USA, UC San Diego CMS T2

W. Yang
K. De
S. Dasu
F. Wuerthwein

Introduction

M. Vetterli welcomes the participants to the meeting.

Minutes of the last Meeting

The minutes of the last meeting are approved.

Agenda

The Collaboration Board agrees on the following Agenda:

- Status of WLCG (I. Bird)
 - Status of EGI and the Specialized Support Centre for HEP (J. Shiers)
 - The European Middleware Initiative (EMI) (A. Di Meglio)
 - Status of WLCG Networking - OPN (W. Salter)
 - Transition to EGI (or future Grid infrastructure)
 - 5 minute statements by Tier-1s on how they plan the transition to EGI and how they will incorporate themselves into their NGIs.
 - Readiness of Experiments for Data
 - General Discussion & Closing
-

Status of WLCG (I. Bird)

I. Bird presents the status of the WLCG Project.

He shows the timeline of the major events affecting the WLCG project in 2009 and 2010. WLCG resources were reviewed in July. The 2009 capacity should be installed in October 2009 and 2010 capacity in June 2010.

The STEP'09 exercise was completed successfully. Overall performance was very satisfactory with tape writing above the required level. Almost all sites participated. After STEP'09 the emphasis has been on stability and resolution of problems in preparation for data taking.

Site reliabilities are measured by Operations using generic tests. Because of their differing requirements the experiments test somewhat different aspects. Also, the experiments have their own tests. Incidents resulting in significant service unavailability are now systematically documented. There is at present no evidence that the general level of problems is decreasing.

The discovery during the summer of serious vulnerabilities in the Linux kernel required rapid updates. The security coordination worked as expected for notification, but many sites did not implement the appropriate updates. This is posing a serious risk to the entire infrastructure. I. Bird emphasizes that this is a serious issue and should be treated accordingly. During a short discussion it is suggested to suspend sites automatically in case of serious security shortcomings. Several members ask for the statistics concerning their area.

Mass Storage has generally seen good performance. No real development is presently being undertaken; emphasis is on providing stable versions for data taking.

Middleware is subject to a continuous process of patching and updating. There are some upgrades that improve performance or functionality.

Resources requirements were re-assessed by the experiments and reviewed by the LHCC and C-RSG. The new requirements were presented to and endorsed by the C-RRB in October. Tier-2 pledges for ALICE remain significantly below requirements. I. Bird summarizes the status of the different Tier-1's.

The EGIeu organization is being set up in Amsterdam. About 30 NGI's have signed the MoU and the EGI Council has been formed and has met. Several EU FP7 projects proposed will support EGI. Among the latter are proposals for specialized support centres as well as the European Middleware Initiative.

I. Bird concludes his presentation:

- STEP'09 was carried out as planned – showing that we are ready to take data
- Business as usual since STEP'09
- The Tier-0 strategy is evolving
- EGEE to EGI transition: the situation is encouraging

Status of EGI and the Specialized Support Centre for HEP (J. Shiers)

J. Shiers summarizes two proposals for the EU 7th. Framework Program presently being prepared. The deadline for submission is very near (November 24). The two proposals, ROSCOE and EGI InSPIRE include requests for services and support and involve different communities.

ROSCOE

ROSCOE (Robust Scientific Communities for EGI) targets different disciplines including mature Virtual Research Communities such as High-Energy Physics, but also nascent communities. The objectives are to enable and increase the number of users and communities and provide for more effective collaboration by removing constraints and barriers.

ROSCOE covers a number of Network and Service Activity work packages across 7 different communities. The target for the project is to ask for 8.5 MEUR from the European Commission.

J. Shiers briefly reviews the different work packages. The HEP manpower adds up to 18.5 FTE contributed by 8 different institutes.

EGI InSPIRE

InSPIRE (Integrated Sustainable Pan-European Infrastructure for Researchers in Europe) focuses on support for Heavy Users Communities (HUCs). HEP Service Activity tasks include dashboards and applications such as GANGA. HEP manpower adds up to 486 person-months.

If successful, the sum of the two proposals represents extremely valuable manpower more than compensating the on-going losses, to cover the early years of LHC data taking. Manpower reductions have already started and there will be a dip before funding arrives. This requires efforts from both CERN and other partners.

J. Shiers notes that, in his view, collaboration with other communities in these two projects will have positive effects beyond the immediate goals of these proposals. This will come at the cost of participating in activities like meetings, training and schools. He proposes to further develop this theme at the first EDI conference.

In conclusion:

- The deadline for submitting these proposals is November 24th, 2009 at 17:00 Brussels time
- The proposals include requests for funding for:
 - Services for Heavy Users, including Ganga, Dashboards & LHC VO-specific services
 - Support for (Experiment) Integration, Operations, Distributed Analysis etc.
- The sum of the two proposals should provide extremely valuable effort to cover the early years of LHC data-taking
- Using these projects, we can prepare not only for the medium-long term but also demonstrate significant socio- economic impact that goes way beyond HEP
- This may open the door to future funding.

Discussion:

F. Gianotti emphasizes the unfortunate coincidence of the funding gap and the analysis of the first LHC data and says that CERN should assure proper support for the experiments. F. Hemmer notes that efforts are being made by CERN: three indefinite appointments have been agreed and two new positions will be opened. F. Gianotti adds that the shortfall concerns three staff for ATLAS alone.

J. Schukraft mentions that the fellowship selection committee is being prepared now but that IT Department has used its entire quota. If extra money could be found this could be used for additional fellowships.

The formal decision on the proposals will only be announced in June. The EU has been informed of CERN's scheduling difficulties.

The European Middleware Initiative (EMI) (A. Di Meglio)

A. Di Meglio, Project Director of the European Middleware Initiative (EMI) starts by explaining that EMI represents a close collaboration between the three middleware providers ARC, gLite and UNICORE, together with other software providers. The goal is to establish a sustainable model to support, harmonise and evolve the grid middleware.

The project is funded in the context of the EU FP7, with a provisional starting date of 1 May 2010 and a period of three years. The total budget is 26 MEUR, shared equally between the EU and the 24 EMI Partners.

The objectives of EMI are to consolidate the existing middleware distribution, evolve the middleware services and functionality and maintain the middleware distribution.

EMI is organized in the form of Product Teams responsible for all technical tasks from design to release and support. A Project Technical Board oversees the Product Teams. Di Meglio emphasizes the importance of Quality Assurance in the form of continuous monitoring of projects. Release policies include major and more frequent minor releases as well as revisions to fix bugs.

A. Di Meglio then presents the selection criteria for the services to be included in EMI. New services will be included based on new user requirements when needed.

A. Di Meglio ends his presentation by listing the middleware evolution targets for the different services.

Discussion:

Replying to a question from M. Vetterli A. Di Meglio says that the first priorities among the different services are the security model and messaging and improvements in monitoring, accounting and data management will follow soon after. M. Vetterli further notes that OSG is not listed as collaborator. A. Di Meglio replies that there are discussions with OSG.

P. Charpentier points out that Pilot Jobs are not mentioned even though they are being used by all experiments.

OPN (W. Salter)

W. Salter presents the status of the Large Hadron Collider Optical Private Network (LHCOPN). The network links CERN with all Tier-1 sites by means of 10 Gbps links except TRIUMF (5 Gbps) and BNL (10 +7 Gbps). Standard internet backup is available for all sites except RAL. He notes that some traffic between Tier-1 sites uses LHCOPN while there are additional Tier-1 – Tier-1 links outside LHCOPN. He shows some details of the USLHC upgrade presently under way.

The LHCOPN Operations Model has been developed and agreed and has been in production since May 2009.

Following his presentation, W. Salter raises the following questions:

1. Should Tier1-Tier1 links (not foreseen for Tier0-Tier1 traffic) be part of LHCOPN?
2. Are there any new networking requirements, e.g. for improving Tier1-Tier2 connectivity?
3. Should we define SLA for links?

Concerning the first question, M. Jouvin mentions that data transfer from Lyon is problematic and would benefit from LHCOPN.

I. Bird pleads for an extension of the OPN mandate that is currently limited to Tier-0 – Tier-1 traffic. He will take this as a request to the Management Board.

On a question from M. Ernst concerning the timeline for OPN monitoring, W. Salter replies that a reasonable solution could be available by the second quarter of 2010.

In summary, W. Salter says that LHCOPN is fulfilling the 'known' needs of WLCG with a good level of resilience and that the operational model is now in production.

Update on Status of Tier-1s

- **UK (N. Geddes)**

N. Geddes briefly discusses the transition to EGI. The impact on the Tier-1 itself should be small, but there is a risk of losing key staff because of lower funding in EGI than EGEE. GridPP and NGS funding will continue at current levels for another year but up to 50% of non Tier1 specific support staff may be lost from 2011 onwards.

- **France (R. Rumler):**

The French NGI is called "France Grilles" and is based in the 'Institut de Grilles' of CNRS. R. Rumler shows the organizational structure. Key persons have been designated for all subtasks and functional entities; some are delegated from their home institute. R. Rumler is optimistic about the switch to a EGI structure, but some concern remains about temporary contracts held by some of the key people.

- **Germany (K.-P. Mickel)**

KIT (Karlsruhe Institute of Technology; the former FZK) hosts the German Tier-1 GridKa and represents Germany in the different WLCG bodies. The supporting organization for NGI-DE will be the Gauss Allianz, an association of large German scientific computing centres. KIT will lead the NGI-DE consortium of four members (KIT, LRZ (Munich), FZJ (Jülich) and DESY. The consortium agreement has been signed. In summary, the German Tier-1 GridKa is well embedded in the new structure.

- **The Netherlands (M. Bouwhuis)**

The Funding Agency NCF serves as de facto legal NGI with SARA and NIKHEF as operational partners. Activities are funded through the Dutch e-science grid BIG GRID. A national plan for the consolidation of Dutch e-science was submitted in 2009; action is expected in 2010.

- **Italy (L. Dell' Agnello)**

IGI (Italian Grid Initiative, the Italian NGI) is supported by the Ministry of Research and Education. Pending the formation of a new legal entity, the present JRU will act as NGI. Services presently provided by INFN Grid, mostly hosted at CNAF, will be taken up by IGI. In the future, INFN and other IGI partners will support the current level of effort to ensure a smooth transition.

- **Spain (M. Delfino)**

The Spanish Network for e-Science is funded by the Ministry of Science and Education and has the mandate to help organize / set up the Spanish NGI. Spain has signed the EGI LoI and MoU. M. Delfino reviews the status of the different operational tasks and shows the roadmap proposed by the Spanish NGI coordinators. The personnel continuity in critical operations tasks is a challenge.

- **Nordic DataGrid Facility (L. Nixon)**

NDGF is a collaboration between four NGIs: Denmark, Finland, Norway and Sweden. The collaboration fits well in the EGI/ NGI setup. Funding is secured until the end of 2010 and a plan for future funding is being made. The present operations collaboration between Nordic Tier-1 and Nordic/ Baltic Tier-2 and EGEE sites is planned to continue also after EGEE-III. In conclusion, NDGF fits well in the EGI model and WLCG obligations are not endangered.

- **Taipei, ASGC (Q. Gang)**

ASGC will continue running the Asia Pacific Regional Operation Center (APROC) in support of e-Science in Asia. ASGC leads the JRU of the Asia Pacific partners to join EGI-InSPIRE. ASGC will join ROSCOE to support HEP applications.

- **Canada (M. Vetterli)**

Canada will not be part of EGI.

Funding for High Performance Computing sites and hardware is now centralized in Compute-Canada. The Canadian Tier-1 has dedicated funding until ~2012. Funding for network infrastructure is up for renewal.

ATLAS-Canada has received a 3-year grant to fund 50% of a person at each of the Tier-2 sites. The CERN ROC will shut down end 2009. TRIUMF is establishing a new ROC to service sites previously serviced by CERN (Canada, China ...). A TRIUMF staff member has been designated ROC manager.

- **United States (R. Pordes)**

OSG provides the infrastructure for the US participation in the WLCG. The OSG is presently funded until October 2011. A proposal for a further 5-year term of funding is in the process of being prepared; a draft proposal should be ready by March 2010 for endorsement by the OSG Council.

OSG is ready for the transition to EGI, which is based on the same federated model as OSG. Letters of intent to collaborate with EGI and ROSCOE have been written.

CMS Readiness (P. Kreuzer)

P. Kreuzer shows an overview of the CMS computing model. In all, more than 60 sites are involved. A second copy of the raw data is stored on the Tier-1s where also re-reconstruction is carried out. User analysis and simulations are carried out on the Tier-2s. Tier-3s are entirely dedicated to analysis.

P. Kreuzer reviews the readiness of the different layers. The main conclusions are:

- The Tier-0 has been performing well with cosmics data with good stability and performance of the CMS software. The CERN Analysis Facility (CAF) is used for commissioning and low latency work.
- Tier-1 availability has improved but some issues need to be resolved. Tier-1 readiness for analysis was tested in October with a job rate comparable to expectations in the Computing Model.
- Data samples were successfully replicated to Tier-2 sites. Latency problems need to be taken up with developers.
- Distributed MC production was very successful in 2009.
- Operations still need substantial effort.
- There are 45 registered Tier-3 centres of which 25 have received data during the last quarter.

Overall, the Computing Project seems to be ready for first collision data.

Reacting to a question from M. Vetterli, P. Kreuzer says that the biggest problem is storage. This involves many different technologies. There appears to be no magic solution.

ATLAS Readiness (D. Barberis)

D. Barberis explains the ATLAS software release strategy featuring a stable release for Tier-0 production, HLT and online monitoring and another stable release for simulation production, reprocessing and analysis. Code is built for SLC4 and SLC5. Much work has been done to improve memory and CPU usage.

Throughput tests are performed periodically to test the links between the Tier-0, the Tier-1's and the Tier-2s. Last month's test was partially successful. More tests and developments are needed.

Some statistics of production jobs on the Grid are shown. Distributed analysis has been tested using HammerCloud, leading to improvements in performance.

Concluding, D. Barberis states that ATLAS has a robust Software and Computing system ready for LHC collision data.

On a question from M. Bouwhuis if enough Tier-2 resources are available, D. Barberis replies that ATLAS is indeed short of some Tier-2 resources.

ALICE Readiness (L. Betev)

L. Betev presents the strategy for software releases. He starts by describing the role of the core offline team for the AliRoot framework, the detector groups for their modules and physics working groups for the analysis packages. AliRoot releases are done twice a year. The latest release dates from July 2009 and is what will be used for first LHC physics.

Problems are tracked in AliRoot Savannah.

In summary:

- The AliRoot release policy ensures code stability through major revisions twice a year.
- Portability (platforms, compilers) is emphasized
- Tracking and weekly reviews with the different participants assure code readiness for LHC startup.

On a question from J. Gordon, L. Betev replies that ALICE is ready for LHC startup. Replying to N. Geddes, L. Betev says that the question of missing resources is still pending, but that this should not have a significant impact on the first physics.

LHCb Readiness (P. Charpentier)

P. Charpentier shows an overview of the LHCb computing model. It differs from other experiments to the extent that Tier-2s are used for simulation only.

The resources required and agreed with the CRSG have been pledged and LHCb does not at present have any particular resource problems.

An overview of the different software components and their interrelations is shown. Software releases follow a bottom-up strategy: releases are performed on demand from subsystems, core software or production teams.

LHCb has used 140 sites in 2009, mostly for simulations. CPU usage statistics are shown.

Despite the fact that there are many problems, LHCb applications are ready for first physics. and pledged resources just match requirements. The main risks for LHCb are site stability and the scarcity of manpower.

On a question from M. Delfino about T1 stability, P. Charpentier replies that PIC is not a major site for LHCb. LHCb distributes tasks according to pledges and wants to do the same for simulation.

Next Meeting

The next meeting will take place in the first half of 2010, but the date is not decided.
