



22 May 2014

**Minutes of the 25th LHC Computing Resources Review Board Meeting
(CERN, Geneva, 29th April 2014)**

Present:

G. Taylor (University of Melbourne, Australia)
J. Lemonne (FWO, Belgium)
G. Wilquet (FNRS, Belgium)
S. Novaes (UNESP, Sao Paulo, Brazil)
D. Wayner (National Research Council, NSERC, Canada)
M. Vincker (Carleton University, Canada)
Y. Zhang, S. Zhang (National Natural Science Foundation, China)
N. Fejksova (Ministry of Education, Youth and Sports, , Czech Republic)
M. Lokajicek (Institute of Physics AS CR, Czech Republic)
P. Hansen (Niels Bohr Institute, Denmark)
J. Aysto (Helsinki Institute of Physics, University of Helsinki, Finland)
A.I. Etienvre, E. Lançon (CEA/IRFU, France)
U. Bassler, F. Malek, G. Lamanna (CNRS/IN2P3, France)
R. Feldmann (Federal Ministry of Education and Research, BMBF, Germany)
M. Groll (BMBF/PT-DESY, Germany)
M. Fleischer, V. Guelzow (PT-DESY, Germany)
T. Kolleger (GSI Helmholtzzentrum, Germany)
A. Streit (BMBF/Karlsruhe Institute of Technology, Germany)
S. Bethke (Max Planck Institut fuer Physik, MPI, Germany)
G. Vesztegombi (Wigner RCP, RMKI, Hungary)
L. Levinson (Weizmann Institute of Science, Israel)
A. Zoccoli, F. Bedeschi, N. Pastrone (INFN, Italy)
A. Di Ciaccio (University of Roma Tor Vergata, Italy)
T. Kawamoto (University of Tokyo, ICEPP, Japan)
J. Son (Ministry of Science, ICT and Future Planning, Korea)
K.W. Lee, T.Y. Lee (National Research Foundation, NRF, Korea)
F. Linde (NIKHEF, Netherlands)
B. Jacobsen (Research Council, Norway)
F. Ould-Saada (University of Oslo, Norway)
F.D. Buzatu (Institute of Atomic Physics, Romania)
M. Dulea (IFIN-HH, Romania)
V. Matveev (Joint Institute for Nuclear Research, Russia)
V. Savrin (Institute of Nuclear Physics, Moscow State University, Russia)
V. Shevchenko (National Research Centre, Kurchatov Institute, Russia)
Z. Hlavacikova (Ministry of Education, Science, Research and Sports, Slovakia)
L. Sandor (Institute of Experimental Physics, Kosice, Slovakia)
D. Bruncko (Institute of Experimental Physics, SAS, Slovakia)
F. Del Aguila (Ministry of Economy and Competitiveness, Spain)
N. Colino (CIEMAT, Spain)

P. Karlsson (Swedish Research Council, Sweden)
M. Nylen (Umea University, Sweden)
D. Meister (ETH Zurich, Switzerland) (replacing C. Grab)
S.C. Lee (Academia Sinica, Taiwan)
G. Zinovjev (Bogolyubov Institute for Theoretical Physics, National Academy of Sciences, Ukraine)
G. Blair, A. Medland (STFC, United Kingdom)
C. Parkes (University of Manchester, United Kingdom)
S. Rolli (Department of Energy, United States of America)
J. Cochran (Iowa State University, United States of America)
P. McBride (Fermilab, United States of America)
S. Rajagopalan, T. Wenaus (Brookhaven National Laboratory, United States of America)

WLCG

I. Bird, S. Foffano, F. Hemmer

CERN

S. Bertolucci, C. Decosse, T. Lagrange, L. Mapelli, J. Salicio-Diez, E. Tsesmelis, E. van Herwijnen, E. Van Hove

Computing Resources Scrutiny Group

J. Flynn, D. Lucchesi

Resources Scrutiny Group

C. Touramanis

ALICE: P. Giubellino, Y. Schutz

ATLAS: D. Charlton, R. Mount, T. Wengler

CMS: T. Camporesi, M. Girone, P. Sphicas

LHCb: P. Campana, M. Cattaneo

Excused

R. Heuer (CERN), S. Lettow (CERN), K.A.P. Sinha (Department of Atomic Energy, DAE, India), J. Królikowski (Ministry of Science and Higher Education, Poland), S. Gonzalez (National Science Foundation, United States of America)

Documents can be found in the RRB indico pages; accessible via the LHC-RRB home page
<http://cern.ch/committees/all/welcomeLHCRRB.html>

1. Introduction S. Bertolucci, Director of Research and Scientific Computing.

S. Bertolucci welcomes delegates to the meeting and invites I. Bird to present the status of the WLCG Project.

2. Approval of the minutes of the last meeting S. Bertolucci, Director of Research and Scientific Computing.

The minutes of the last Computing Resources Review Board meeting CERN-RRB-2013-135 are approved without comments.

3. Status of the WLCG Project and Financial Status Report I. Bird, WLCG Project Leader CERN-RRB-2014-044 (Report), CERN-RRB-2014-045 (Slides).

I. Bird informs delegates about the status of WLCG Collaboration since the last meeting with the approval of KISTI as a full Tier-1 for ALICE by the WLCG Overview Board in November 2013. The upgrade to the LHCOPN 10 Gbps is still outstanding with respect to the milestones however it is planned and funded; I. Bird stresses the importance of this connectivity for the new Tier-1. Several new sites have been added to the Latin America Tier-2 serving all 4 experiments, and a new Tier-2 site in Pakistan is in preparation. The revised financial situation takes into account late purchasing in 2014 to meet the latest requested resources and should be viewed as advanced planning rather than a budget surplus. WLCG activity is reported on with resources fully used at anticipated levels and heavy use of the HLT farms. Plots are shown reflecting the current activity with continuous large-scale workload above or at pledge levels.

The Wigner Centre status is presented as in production with over 1000 worker nodes. There have been several network incidents recently due to firmware incompatibilities, now resolved. Generally there is no significant difference in job efficiency between CERN and Wigner.

Changes to the reliability and availability reporting, approved by the WLCG Management Board, are outlined. Four reports with more meaningful metrics are produced monthly including experiment-specific tests to better reflect the experiments use of the sites.

The computing model document requested from the LHCC was recently published; results obtained in Run 1 were built on implying significant changes in the computing models of each experiment. In addition the WLCG community is investing a lot of effort in software evolution with a frequent difficulty to hire and keep the necessary skills for software and computing development.

A HEP software Collaboration workshop was held at CERN in early April 2014 to build on the work of the Concurrency Forum and attract additional help and expertise, as a perceived lack of computing resources could limit the physics to be accomplished in the coming years. The goals of the initiative were to better meet the rapidly growing needs for simulation, reconstruction and analysis of current and future HEP experiments i.e. not limited to LHC. In addition promote maintenance and development of common software projects and components, adapt to new technologies, improve performance, enable potential new collaborators to become involved, promote collaboration and other scientific and software domains. There was agreement to create a Foundation and hold a follow-up workshop in Autumn 2014.

The HEP and Astroparticle Funding Agencies have put together a new initiative to develop data and computing infrastructures building on the CERN/EIROForum vision document called EU-T0 as a set of services for scientists and science programmes built on common underlying infrastructures. A vision is to be built over the coming months following a discussion with EC representatives and a recent workshop to discuss the H2020 engagement.

I. Bird concludes activities continue at a high level during LS1 with on-going preparations for Run 2 and various on-going initiatives around the evolution of the computing models and the distributed computing infrastructure. A major concern is the level of funding likely to be available for computing in Run 2 with a currently unclear situation putting the future LHC physics output at risk.

U. Bassler questions the rationale behind the choice of “Foundation” rather than “Collaboration” for the future HEP software initiative. I. Bird explains the concerns from those involved in the discussion that a balance is needed between encouraging innovation and considering new things, leading to the choice of “Foundation”.

G. Taylor enquires about the status of discussions on the use of non-dedicated resources which has an impact for small countries and is important to be able to show the use of shared facilities. I. Bird replies much has happened to improve this with less grid infrastructure now needed to set up a site. The long term goal is to be able to use resources without having to install much or any grid infrastructure.

T. Medland comments in preparation for Run 2 the machine is likely to take a cautious approach to starting up and asks how this is factored into the planning. I. Bird explains it is based on the run time and the experience from Run 1 in 2010 as approved by the Scrutiny Group. S. Bertolucci adds start up should be quite fast despite the fact it is a new machine, the dominant factor will be the number of seconds runtime.

F. Linde asks when the concerns on I. Bird’s last slide [the level of funding likely to be available for computing in Run 2 with a currently unclear situation putting the future LHC physics output at risk] will be addressed, adding that physics should not be compromised due to lack of computing.

A. Zoccoli remarks while Funding Agencies are ready to support the analysis of the experiments, there are problems for the future to meet the experiment’s computing requirements. The situation needs to be discussed in detail for the experiments to fully understand what the Funding Agencies can offer.

S. Bertolucci comments on the need to balance many things: improving the computing models, the continuing effort devoted to software development, the transition phase between FP7 and H2020, the need for extra resources as obtained in the past, and the need to consolidate the level of commitment from the Funding Agencies. He suggests listening to the Scrutiny Group before discussing further.

4. LHCC Deliberations E. Tsismelis, LHCC Scientific Secretary. CERN-RRB-2014-046 (Report)

On behalf of the LHCC E. Tsismelis congratulates WLCG and the experiments on the efficient data processing underlining that funding for the pledged WLCG resources remains under pressure.

5. Status of Common Projects accounts T. Lagrange, CERN Finance and Procurement Department. CERN-RRB-2014-047 (Report)

T. Lagrange confirms he has nothing to add to the written report.

6. Report from the Computing Resources Scrutiny Group. J- Flynn, C-RSG Chair.
CERN-RRB-2014-048 (Report), CERN-RRB-2014-049 (Slides)

J. Flynn begins by presenting the assumptions for Run 2 including the revised efficiency assumptions: CPU organized 85%, CPU analysis 70%, disk 100% and tape 85%. The overview of the sum of all Run 2 requests starting from 2014 pledges with respect to meeting flat budgets (geometric growth from hardware improvements only) shows that CPU and disk at the T0 rise above the flat budget in 2015 then subsequently grow more slowly. With the exception of T2 CPU other resources grow above flat budget earlier or later in Run 2.

Resource use in 2013 is reported on at each Tier level and by experiment. All experiments have been asked to provide data popularity information in the hope that revealing and monitoring this information will lead to more efficient use of disk space by minimizing storage of data rarely or never used.

J. Flynn reports on the scrutiny for 2015 with individual assessments and comments for each of the experiments. The following global comments and recommendations are then made:

- Flat budget assumption: Run 2 requests are made with a flat funding profile i.e. not adjusted for inflation
- Data preservation: Distinguish the ability to read/reanalyze old data from requirements for open/public access. Allowing for public access to data has implications for future resource requirements, requires substantial human effort to implement, and risks taking resources from LHC physics output if not funded separately.
- Use of HLT farms: The CRSG acknowledges the use of online farms during LS1 and the plans to use them during technical stops and shutdowns in Run 2. The CRSG does not consider this use to be opportunistic; it is within the control of the experiments.
- Support software engineering: Improving software efficiency (ultimately physics per euro) is essential to constrain the growth in requests. The resulting gains are already assumed. CRSG strongly supports this for Run 2 and as preparation for Run 3, and recommends that sufficient software engineering effort is funded.
- Disk efficiency: Effectiveness of disk use is only partly captured by occupancy. The CRSG welcomes the experiments' efforts to purge obsolete or unused data and their agreement to provide data popularity information for future scrutiny rounds.
- Importance of networking: Good networking has been exploited to reduce disk use and move processing between tiers. There is a danger that poorly-networked sites will be underused and there are potential implications on the cost of providing network capacity.
- Scrutiny schedule: Following the new schedule approved at the October 2013 RRB, the first scrutiny of 2016 requests will be reported on in October 2014, revisiting 2015 only if necessary. It is not intended to report resource usage to the RRB in October 2014.

The CRSG membership is reviewed. G. Lamanna from France and M. Gasthuber from Germany will stand down following the current scrutiny round and are thanked for their

contributions, the CRSG looks forward to nominations of their replacements. J. Flynn concludes by thanking the experiments for their constructive dialogue and the CERN management for its support.

A. Zoccoli thanks the experiments for the work done to improve their software and build up a new computing model to meet the needs of Run 2. As the computing budget comes from the same source as the experiment upgrade, he stresses the importance of preserving the computing budget and requests clarification from the CRSG on the point from when the flat budget is defined, and clarification on the equipment replacement rate. In addition he remarks on the impressive plot of data access vs time adding the experiments should use this to optimise their disk space usage. Finally he asks if contributions from the other Tiers can be reduced as new Tier-1's contribute new resources, and why LHCb requests an increase in disk resources. On the last point M. Cattaneo explains this is to make better use of T2 resources and address a shortfall in disk. All the analysis is done at T1's however large T2's were contacted who commissioned disk and were used like T1's. The future aim is to increase this. Countries with a T1 and a T2 can choose where they attribute the disk.

F. Linde remarks ATLAS used twice as much CPU as CMS in 2013 adding usage includes beyond pledged resources from which ATLAS had greatly benefitted. D. Charlton explains the large use of CPU was needed for calorimeter simulation. This and improved reconstruction times have been worked on during the shutdown as there is no future guarantee of opportunistic resources. D. Charlton explains that the complex geometry of the ATLAS calorimeter is reflected in the large CPU use for simulation. Improving the flexibility of the simulation to allow more use of faster parametrisations, as well as improving reconstruction times have been a priority during the shutdown, as there is no guarantee of future opportunistic resource availability.

V. Guelzow stresses the importance of the work of the Scrutiny Group and confirms the process is on-going to ensure replacement of the German member before the next meeting.

U. Bassler confirms assessment of the most suitable of two candidates to replace G. Lamanna is currently on-going in France. A replacement will be nominated before the next meeting.

T. Medland shares the concerns expressed about the flat baseline funding and the need to understand the requests now to ensure optimal physics and use of resources. The LHC machine and detector performance at start-up is still a big unknown, so Run 1 may not be a reliable source for estimates.

S. Bertolucci confirms experience from Run 1 indicates the LHC is over-performing. For Run 2 the experiments are expected to have more data and if the machine has to run at 50ns it will be at increased luminosity. The driving factor will be the number of seconds of running time.

7. Summary. S. Bertolucci, Director of Research and Scientific Computing.

In order to define the flat baseline more precisely, S. Bertolucci proposes the CRSG and experiments should consider more closely the evolution of requests over Run 2 and the consequences for future pledges which would be needed to meet them. The proposal is accepted.

S. Bertolucci thanks the delegates and closes the meeting.