

WLCG Overview Board

Minutes of the thirty-first meeting held on Friday 25 May 2018

Attendance

Function	Name	Presence
Director of Research and Scientific Computing	E Elsen	Present via teleconference
Scientific secretary, WLCG resource manager	H Meinhard	Present at CERN
IT Department head	F Hemmer (chair)	Present at CERN
EP Department representative	R Forty	Present at CERN
WLCG project leader	I Bird	Present at CERN
ALICE	F Antinori	Absent
ATLAS	K Jacobs, replaced by T Wenaus	Present at CERN
CMS	J Butler, replaced by T Boccali	Present at CERN
LHCb	G Passaleva, replaced by M Cattaneo	Present at CERN
Canada	M Vetterli	Absent
France	L Duflot	Present via teleconference
Germany	J Bluemer, replaced by A Heiss	Present via teleconference
Italy	A Zoccholi	Absent
Japan	J Tanaka	Present at CERN
The Netherlands	J Templon	Absent
Nordic countries	I Bearden	Absent (apologies received)
Republic of Korea	S Y Noh, replaced by S U Ahn	Present via teleconference
Spain	M Delfino Reznicek, replaced by J Flix Molina	Present via teleconference
Russia	No representative	
Taiwan	S Lin	Absent
United Kingdom	D Britton	Present at CERN
United States ATLAS	E Lançon	Present via teleconference
United States CMS	K Bloom	Present via teleconference
Tier-2 representative	M Jouvin	Present at CERN

Function	Name	Presence
Representative of non-EU/US Tier-2 sites	K Mazumdar	Present via teleconference
GDB chair	I Collier	Present via teleconference
Collaboration Board chair	R Jones	Present at CERN
Additional CMS representative	L Sexton-Kennedy	Present via teleconference
Deputy WLCG project leader	S Campana	Present at CERN
Deputy United States ATLAS representative	H Severini	Present via teleconference

Material: <https://indico.cern.ch/event/677596/>

1. Introduction and approval of minutes (F Hemmer)

F Hemmer, replacing E Elsen, who was attending remotely, welcomed delegates. The only comment to the draft minutes of the 30th meeting on 27 October 2017 was that there should not be any highlighting in yellow. With the removal of the highlighting, leaving the text intact, the minutes were **approved**.

2. WLCG status report and discussion (I Bird)

I Bird presented the report on the WLCG status, mentioning that concerning the collaboration, the Korean Tier-2 centre has been moved to KISTI, and that S Campana is now acting as his deputy as WLCG project leader.

The LHC machine continues to run very well in 2018, the integrated luminosity so far exceeds expectations; the turnaround between fills is very small (of the order of two hours). The WLCG resources are used at a very high level above pledges. In November 2017 the CNAF Tier-1 data centre in Bologna was suffering from an accident: a water pipe broke outside the centre, and the centre was flooded. The impact on WLCG could be limited thanks to other Tier-1 centres stepping in with additional resources; LHCb was most affected, as 20% of the raw data were destroyed, but second copies existed, hence no significant data were lost.

Most Tier-1 centres mostly run dual-stack (IPv4 and IPv6) services now; the deployment of IPv6 at Tier-2 sites now needs to be pushed forward. The Spectre and Meltdown security vulnerabilities announced towards end 2017 caused some disruption due to the patching and reboot campaigns required; initial worries about loss of performance due to the patches were fortunately not confirmed.

ALICE expects to take a significant amount of Heavy Ion (HI) data in 2018 (corresponding to 60% of the HI statistics in Run 2); LHCC has recommended ratios of raw over simulated data of 1:1 for proton-proton and 1:0.3 for HI data. ATLAS has collected 47 fb⁻¹ in 2017; the Tier-0 resources hardly kept up with the load. They introduced a fast chain of simulation, which not only saves CPU cycles, but is also advantageous in terms of storage. ATLAS is successfully using unpledged resources at a significant level. CMS have recorded very high data volumes as well, creating high load in particular on the Tier-0, the capacity of which was increased by adding resources primarily intended for analysis, and resources from Tier-1 centres. The Tier-0 resources are no longer provided as virtual

machines in a private cloud, but as capacity administered by CERN's batch services based on HTCondor. LHCb is working on improving algorithms for Run 3 in order to contain their resource needs.

The situation concerning 2018 pledges looks healthy. The pledges almost meet the requests; no particular issue is requiring special attention at this time.

The Community White Paper (CWP) by the HEP Software Foundation (HSF) was published at the end of 2017, and was the topic of a dedicated article in the CERN Courier. It was followed up by a joint WLCG and HSF workshop in March 2018 with some 200 attendees, during which R&D activities were started. Based on the CWP, a WLCG strategy document has been drafted (see dedicated presentation), which prioritises the R&D activities mentioned in the CWP. LHCC will consider the draft in its forthcoming meeting (week of May 28th); it has been circulated to the WLCG Management Board, and will be published after the LHCC meeting. LHCC are proposing to review the strategy document and the ensuing R&D activities during early 2019; choosing reviewers and selecting a mandate for the review will require some care, as software as well as infrastructure aspects need to be covered by relevant expertise. Efforts to watch how technology and markets evolve are ongoing, hinting at average gains of 15% and 20% per year for CPU and disk capacity, respectively. However, there are significant risks and uncertainties, for example because very few companies dominate the markets. In addition, some market trends, for example around GPUs and accelerators, are driven by machine learning requirements, and may not be immediately useful for HEP data processing.

Open access to data (Open Data) and data preservation are currently organised in an ad-hoc manner; the experiments have different scales and use-policies. CERN currently provides an Open Data portal based on EOS disk storage funded by CERN's WLCG's resources; a comprehensive policy, assignment of responsibilities and in particular, allocation of resources will be required.

The ESCAPE proposal within the framework of the Horizon 2020 programme has been submitted to the European Commission; funding for HEP is requested within WP2 on data infrastructure for open science. No news has been received yet, but we expect feedback in September or October 2018 for funding to start early 2019 if approved.

D Britton asked whether the purpose of the LHCC review of the strategy document is only to engage the funding agencies, or to check whether the strategy is adequate. **I Bird** replied that it is the latter, and that the review may allow for pushing back the computing TDR initially requested by the LHCC for 2020. **R Jones** commented that pushing back the TDR may create tensions between computing and software, to which **I Bird** replied that the storage R&D activities ("data lakes") are indeed a software contribution. **L Sexton-Kennedy** added that CMS are currently reviewing their approach to workload and data management; reconstruction requires particularly intense collaboration with the detector experts. She expressed agreement with the time scales proposed in the presentation. Replying to a question by **I Bird**, **T Wenaus** explained that in addition to the WLCG computing TDR, ATLAS intends to write an experiment-specific one, for which numerous activities have started within ATLAS.

M Cattaneo asked whether the situation concerning Open Data and data preservation has changed recently, and commented that if the funding agencies request these services to be provided, they should finance them. **I Bird** replied that no additional funding was made available. **D Britton** added

that this topic requires an in-depth discussion; the funding agencies are likely to be positive towards these activities, but additional funding cannot necessarily be expected. **S Campana** said that a proper estimate of the resource needs requires clarifying the purpose of the activities. **E Elsen** commented that the funding agencies are expecting HEP to follow other sciences, who have been providing Open Data facilities for quite some time already, mentioning astroparticle physics as a particularly outstanding example, and agreed with the need to properly discuss needs and strategy. **I Bird** said that defining the WLCG policy would be a task for the Overview Board. **L Sexton-Kennedy** said that CMS has been pioneering Open Data for LHC experiments, which has led to some physics output already. In the US the National Science Foundation is very interested, and is considering providing additional funds for this purpose. **M Jouvin** commented that collecting input on Open Data from outside the experiments is problematic, since apart from the funding agencies, there are no structures to contact. **F Hemmer** suggested discussing the policy, once defined by the WLCG Overview Board, in the CERN Scientific Policy Committee or even in Council. **H Meinhard** mentioned that the computing resources scrutiny group (C-RSG) had been confronted with requests for Open Data resources already, and wondered whether the group should consider them, given that they are not needed for the activities C-RSG was established to review, but they are funded from the same limited budgets. **I Bird** and **E Elsen** replied that C-RSG should not review these requests, which should be funded separately. **E Lançon** wondered whether these activities would not be a CERN-only concern; **I Bird** replied that according to the WLCG MoU, the responsibility is shared with the Tier-1 centres. **D Britton** suggested considering reducing costs by using tapes rather than disks, noting that using tapes actively may increase costs compared with a pure archive usage. **F Hemmer** stated that there was obviously need for further discussion, but suggested in view of the limited time to defer such discussion to a later occasion.

I Bird mentioned that C-RSG had requested that unpledged resources be listed, and expressed concern, wondering whether opportunistic resources are part of the pledge. Currently WLCG accounts for using non-pledged resources posteriori, but only for sites/federations that pledge themselves. Only the experiments have the full information what resources they use. **K Bloom** commented that it is indeed a question of definition, which requires clarification. **D Britton** suggested sending a statement to C-RSG agreeing with providing posteriori accounting figures and pointing out that prior commitments are unrealistic and possibly counter-productive. **I Bird** agreed to provide such statement (*action*).

K Bloom expressed surprise about S Campana's nomination as deputy project leader, as he was not aware of any formal process, adding that this was not to imply any personal criticism. **I Bird** replied that this nomination, not having any bearing on a future decision of his successor as project leader, was entirely at his discretion and did not require any formal process.

3. Run 3 requirements and budgets (I Bird)

I Bird presented a summary of the current understanding of requirements and budgets for Run 3. In 2018 the average pile-up increased from 35 to 45, reaching peaks of 55 limited by luminosity levelling, which means larger fractions of fills at high pile-up. Similar levels are to be expected for Run 3, even though there are still many unknowns. The main limiting factor for the LHC machine performance is expected to be the heat load created in the cryogenics. To first approximation, similar conditions as in 2018 can be expected. However, the predictions of running conditions by various models show very significant spread; there could be much more data and longer levelling periods

leading to higher resource needs, up to the extreme of applying luminosity levelling for the entire fill. In total, some 30% more resources can be expected to be needed for Run 3 than for 2018 (a safer estimate is 50%); however the data volumes and resource needs for 2018 are still to evolve. The data-taking period in 2021 could be shorter than for the following years, which may allow for late provisioning.

If 50% more resources need to be made available, the situation appears to be rather under control. However, ALICE and LHCb will change their computing models entirely during LS2; no firm numbers are known yet, but in particular, LHCb will probably need a much steeper increase than the canonical 15...20% per year that can be expected to be covered by flat budgets. The ALICE TDR has been approved by the LHCC; it foresees increases of 50% for CPU and a little less for disk.

CERN plans to conduct procurements for 2021 as late as possible, re-profiling resources from 2018 and 2019, even though there is demand for resource increases during LS2. One of CERN's options to prepare for Run 3 is to re-configure the space in the B513 data centre, which would imply disruptions and temporary drops of resources during LS2. The funding agencies have hinted strongly that nothing more than flat funding can be expected; as 2022 and 2023 are likely to be very demanding on resources, it is quite probable that WLCG will hardly be able to provide the required resources, in particular at CERN.

R Forty and **M Cattaneo** commented that the LHCb software TDR has been released; its implication on resources are being discussed with the LHCC. The document is expected to become public in September 2018.

L Duflot said that in France, resource allocations are proportional to the number of physicists; additional resources for ALICE and LHCb are unlikely.

E Elsen emphasised the risk of very high pile-ups (up to an average of 55) if beta-star squeezing is put into operation early during the fills. **I Bird** concurred, reminding that the assumption for 2018 is an average of 35 only.

4. WLCG strategy towards HL-LHC (S Campana)

S Campana introduced the WLCG strategy document, explaining that the starting point was the CWP, which is a crucial step for WLCG in preparing for HL-LHC. CWP also addresses the computing needs of other research communities with data and processing requirements at a similar scale; the strategy document is a specific view, with prioritisations, of the CWP content. The R&D work addressed in the strategy document will lead to prototyping solutions, which will serve as the basis for the computing TDR.

The strategy document focuses on five themes: software performance, algorithmic improvements and changes, reduction of data volumes, management and containment of operational costs, and optimisation of hardware investments. It is assumed that for HL-LHC, ATLAS and CMS will need about 20 times more resources than now; based on canonical growth, which is an optimistic assumption, the resources will still fall short of what will be required by a factor of five. A key element of the process is modelling the impact certain changes will have on the cost, which is being studied by a dedicated working group.

The question will need to be addressed how much simulated data with respect to detector data is needed. Aggressive reduction of data formats will be required, with more extensive formats being stored on less expensive media (such as magnetic tape). The potential gains from a more organised production and analysis need to be evaluated. The event generators take an ever-increasing fraction of the CPU time spent, and will need much more attention concerning optimising their performance. The simulation and digitisation programs will need to be modernised, introducing multi-threading, vectorisation, and pre-mixing for pile-up; dedicated investments are required into fast simulation. The reconstruction is the main challenge at high pile-up; vectorisation as well as concurrency (the latter to reduce the memory requirements) will need to be studied in more detail. The data layout will need to be reviewed in view of performance; investments are needed into a more automatic physics validation. Storage will be consolidated, but will remain distributed, with a layer of content delivery and caching; prototyping distributed storage services is part of the plan, which also includes modernising data transfer and access protocols. The overhead of authentication needs to be reduced. Further-going storage federations (possibly a single “data lake” for all WLCG) need to be studied; the boundary between infrastructure and high-level experiment-specific solutions needs to be defined. Concerning networks, the usefulness of software-defined networks and protocols other than TCP needs to be investigated; it needs to be understood how commercial cloud resources can be attached to the WLCG network. Brokerage of jobs in order to optimise the co-location of data and CPU needs to be studied.

A project on data management and data organisation has started already; there is a serious person-power issue around software and current core tasks. Leveraging commonality at its maximum will be key.

F Hemmer commented that it is not clear how the plan can be resourced. **S Campana** replied that at the Naples workshop in March 2018, a good level of interest has been shown. We will need to further prioritise and to involve the funding agencies, some of which are waiting for this input. **I Bird** added that the community needs to recognise software-related work in a much better way than previously. **E Lançon** agreed with the assessment that further prioritisation is needed; the list is still too large and too broad. We must better understand how to use (expensive) low-latency storage and (less expensive) high-latency one. **S Campana** re-iterated that compression and tiering of data will be required. **T Wenaus** reported that ATLAS is looking at reducing the size and the number of data formats as well as at using more high-latency storage.

E Elsen stated that there is a lot of room for experiments to adopt common solutions, and suggested involving the CERN SPC or even Council to promote exploitation of public e-infrastructures. **I Bird** added that common solutions were seriously mentioned often times at the Naples workshop, and that some concrete projects had started already; the situation looks better than previously.

5. AOB

Concerning the date of the next meeting, participants agreed that the preferred date would be Friday November 30th from 13:00 h to 16:00 h. (*Editor’s note added after the meeting: The date was confirmed.*)

F Hemmer summarised the actions emerging from the meeting (see table below), thanked participants and closed the meeting.

List of actions

Added	Action on	Action	Comments/status
2018-05-25	All OB members	Propose reviewers for LHCC strategy review (contact I Bird)	
2018-05-25	E Elsen	Bring Open Data and data preservation to Council's (or SPC's) attention	
2018-05-25	I Bird	Write message to C-RSG about handling non-pledged resources	
2018-05-25	All funding agencies	Consider Run 3 and course of action if resources are insufficient	
2018-05-25	H Meinhard	Revise membership list, ensure mailing list is aligned and up-to-date	