



Helge Meinhard / 08 May 2018

Status: draft 1 for participants' comments

Fourth Scientific Computing Forum CERN, Geneva, 23 February 2018 Minutes

Present at CERN

Maite Barroso Lopez (CERN, IT Department)

Ian Bird (CERN, IT Department)

Tommaso Boccali (INFN Pisa, Italy, representing CMS)

Concezio Bozzi (INFN Ferrara, Italy, representing LHCb)

David Britton (University of Glasgow, UK)

Peter Clarke (University of Edinburgh, UK)

Eckhard Elsen (CERN, Director for Research and Scientific Computing, Chair)

Josep Flix Molina (PIC/CIEMAT, Spain)

Frédéric Hemmer (CERN, IT Department)

Brij Kishor Jashal (TATA Institute, India)

Robert Jones (CERN, IT Department)

Thorsten Kollegger (GSI, Germany)

Mario Martinez-Perez (Autonomous University of Barcelona, Spain)

Helge Meinhard (CERN, IT Department, Scientific Secretary)

Andreas Petzold (KIT, Germany)

Stefan Roiser (CERN, IT Department, representing LHCb)

Thomas Schulthess (ETHZ-CSCS, Switzerland)

Oxana Smirnova (Lund University, Sweden)

Torre Wenaus (BNL, USA, representing ATLAS)

Antonio Zoccoli (INFN Bologna, Italy)

Present via teleconferencing

Sadaf Alam (ETHZ-CSCS, Switzerland)

Volker Beckmann (CNRS/IN2P3, France)

Volker Gülzow (DESY, Germany)

Michel Jouvin (CNRS/LAL, France)

Rob Roser (FNAL, USA)

Elisabeth Sexton-Kennedy (FNAL, USA, representing CMS)

Achim Streit (KIT, Germany)

Jeff Templon (Nikhef, The Netherlands)

Apologies

Eric Christian Lancon (BNL)

Christoph Grab (ETHZ)



Kajari Mazumdar (TATA Institute, India)
Karl Jacobs (ATLAS Spokesperson)
Giovanni Passaleva (LHCb Spokesperson)

Introduction and approval of the minutes (E Elsen)

E Elsen welcomed participants to the fourth Scientific Computing Forum (SCF). The meeting is informal and open. Relevant experts are invited to attend for exchange of opinions on topical scientific computing subjects. All information about the forum is publicly available from the SCF Indico page <https://indico.cern.ch/category/9249/>.

There were no remarks on the agenda.

E Elsen stated that as SCF does not take decisions, the minutes do not need formal approval. He invited factual corrections to and comments on the draft minutes of the third SCF. As there were none, these minutes were considered final.

Brief WLCG status report (I Bird)

I Bird gave a short summary of the status of WLCG. 2017 was a very good year for the LHC machine and the experiments, adding to the data volume by October 2017 of now 230 PB on tape at CERN; The amount of data per integrated luminosity is shrinking because all experiments successfully compressed their event information. The WLCG system overall behaved very well, delivering a new peak of 210 million HS06 days, which corresponds to 700 000 cores used simultaneously. It is however clear that with constant budgets, the requirements for HL-LHC cannot be satisfied; I Bird referred to the Community White Paper (CWP) by the HEP Software Foundation, to be addressed in more detail in a later presentation during the meeting. WLCG are preparing a strategy document for spring 2018 that is essentially a prioritisation of the list of topics mentioned in the CWP, as a precursor to a TDR to be delivered in 2020. With the rapid progress in computing technology, the 2020 report cannot be a design report as for detector hardware, but will rather be a snapshot of the understanding at the time of how to meet the requirements of the experiments at the HL-LHC. To that end, WLCG are setting up a programme of R&D and prototyping aimed at an in-depth understanding of system performance across CPU, storage and networks in view of optimising the system overall; the programme includes a major activity on the concept of data lakes. For 2018 the needs are compatible, at least for CPU and disk storage, with the flat-budget assumptions if technology progress continues to result in 20% more CPU and storage per year for the same investment; note however that for 2018 this evolution is far from being clear. The LHC machine is expected to deliver high ($2.0 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$) luminosity at 25 ns bunch spacing with 2 556 bunches. The bunch luminosity will be levelled such that the pile-up will be 55 similar to the situation end 2017; despite of the levelling, the average pile-up will increase from 35 to 45, leading to some 20...25% higher requirements on CPU for reconstruction.

E Elsen thanked I Bird for his presentation and added information from the recent “Chamonix” meeting of the accelerator sector: For 2018, an integrated luminosity of 60 fb^{-1} is expected (in 2017 it was 50 fb^{-1}). If the data volume increases by 20...30%, it will be an issue for the data storage despite all efforts to optimise the latter; it will probably mostly affect tape media. The detector performance decreases as expected due to irradiation. I Bird commented that depending on the sites’ specificities, the higher data volume may also imply that investments on the tape infrastructure (drives and



libraries) are needed, and that the implications of the heavy-ion run at the end of 2018 are not yet fully clarified.

Replying to a request for additional explanation of the data lake concept, I Bird clarified that a data lake is a storage federation with automatic replication; the concept will be prototyped by CERN and some Tier-1 and Tier-2 centres (STFC-RAL and DESY; CNAF could be added later, Nikhef was considering).

Replying to a comment that the assumptions about the 2018 running conditions may not be aggressive enough, E Elsen explained that the estimates are very solid, even though perhaps a bit on the conservative side due to the assumed radiation damage on the triplet magnets. If radiation levels turn out to be lower than expected, the luminosity could be ramped up further, leading foremost to longer fills with luminosity levelling.

Lessons from HEP Software Foundation Community White Paper (M Jouvin)

E Elsen briefly introduced the topic emphasising that the CWP process was very much a bottom-up approach. M Jouvin apologised for presenting remotely due to family constraints and referred to the presentation given by G Stewart in the previous meeting, pointing out that his presentation will be focused on what has finally been delivered, and what steps to take now.

As already mentioned by I Bird, HL-LHC may deliver pile-ups exceeding 200, leading to requirements on computing that very significantly exceed what can be expected from Moore's law and be dealt with in a flat-budget scenario. While in general due to processor architecture advances, CPUs become more powerful over time, this does not apply for all parameters; for example, the performance for floating-point calculations is even going down. DUNE and Belle II have similar requirements and issues as the LHC experiments; it is hence appropriate to join forces and work together. HSF was established in 2015 and has hence held workshops and launched working groups on topics of common interest. In 2016 HSF decided to launch the CWP process; dedicated workshops were held in January and June 2017. Some 250 persons contributed actively to the process. The document was finalised by an editorial board, of which M Jouvin is a member; it is available in arXiv, publication in the "Computing and Software for Big Science" journal is foreseen. For members of the community, signing the document will still be possible until the HSF-WLCG workshop in March in Naples. The document is being complemented by topical working group documents giving significant additional details.

M Jouvin then covered the main points of the CWP. Physics event generators, by means of including next-to-leading-order calculations, are becoming a significant consumer of CPU resources; the codes are usually maintained by the theory community and hence outside the direct influence of the experiments. Current and future work on detector simulation includes improved physics models, vectorisation, fast simulation and geometry modelling; the CWP process has played an important role in ensuring consistency between these activities. The trigger and reconstruction session is a bit biased towards HL-LHC and focuses on the resource consumption for particle tracking, new computing architecture and validation. For data analysis and interpretation, the key metric is "time to insight"; work is needed on interactive data analysis clusters and on making data formats interoperate seamlessly. This is an area with strong potential for collaboration with the non-HEP world. Data processing frameworks will need to be adapted to new hardware, including external (co-)processing



resources; the interface with the workload management can be much improved and consolidated. Machine learning remains a pretty hot topic, the field has advanced significantly with new techniques. Further work is needed on speeding up the calculations, enhancing the physics reach, improving the data compression, and better detecting anomalies. Good links with the broader (non-HEP) communities are required. Other technical areas covered by the CWP include conditions data and visualisation, where the process has had a catalysing effect. Concerning data management and organisation, access to training data sets needs to be sped up, and in general data access must allow for higher granularity and better throughput. In view of better understanding facilities and distributed computing, dedicated effort is needed for the complex task of understanding the effective costs involved in delivering computing; links are being established with other big-data sciences (e.g. SKA) as well as the Computer Science community. In the area of security, cooperation with other HEP and non-HEP communities is key; adapting to evolving regulatory frameworks is challenging. Further work is needed on trust and policy, operational security, and authentication and authorisation. Data, software and analysis preservation is at least as much an issue of preserving knowledge as preserving bits; DPHEP has looked into both, this area appears to be well under control. The experiments have modernised their software development models and processes a lot; in future, commonality across experiments is key, as already prototyped successfully by ALICE and LHCb. The software development environment is becoming ever more complex, requiring physics expertise as well as professional programming skills.

Summarising the technical areas of the CWP, M Jouvin concluded that the master direction is to avoid HEP-specific solutions wherever possible, noting that for some areas, replacements from outside for previously used HEP-specific solutions may not readily available.

Covering the important topic of training and careers, M Jouvin stated that in the HEP community, expertise on software and computing must be raised and improved; training must become a first-class activity, and successful work in software and computing must be fully recognised in view of the career evolution of the person concerned. Ideally, much of the software work would be done by people combining physics and computing expertise, which risks to be a difficult career path, as the person would not excel in either area.

M Jouvin concluded that the CWP was an important milestone engaging many people. It is not a final step, but a solid basis for future work. Each experiment must now build its own prioritised R&D programme out of the CWP, but for each topic relevant to several experiments, it is essential that the work is done together. The software upgrades for HL-LHC require significant investments. While HL-LHC is a natural driver, the process must include the whole HEP community. Support by organisations and funding agencies is required for marshalling and re-focusing the R&D efforts and helping to attract new investments into critical areas. The evolution of the careers of relevant experts is of critical importance.

E Elsen thanked M Jouvin for this comprehensive overview and for opening the discussion towards priorities and common work.

Discussion on lessons from HSF Community White Paper

E Elsen invited questions on M Jouvin's presentation as well as contributions on what lessons to draw from the CWP.



P Clarke said that the CWP statement on not being able to afford multiple solutions is perhaps not strong enough. HEP needs to prove now that work on common solutions can successfully be done together. Funding for distinct activities by experiments is entirely unrealistic. M Jouvin fully agreed, but emphasised that it's all about finding the right effort. E Elsen commented that continuing with home-grown HEP solutions would cut us off further from the mainstream, and is hence not viable. What would happen if we spent one MCHF on software rather than on hardware? Would this make the most effect on simulation (QCD, vectorisation), data management and storage, or some other area?

D Britton commented that one of the base assumptions of the CWP appears to be that budgets will remain flat which is an artificial and arbitrary constraint that may not be 100% realistic; in addition it is not clear whether this concerns only material and hardware, or covers software as well. The work on software, as has become evident from the CWP, is extremely important and must be funded adequately. The community will need to address the constant budget constraint and, being maximally transparent, challenge that constraint. I Bird said that HEP must not lose sight of the overarching aim to optimise physics output given a certain amount of resources. An example of an already on-going activity in that sense is the data lake prototype in view of data management, moving commonality in data management into the infrastructure. This leads to a much needed operational discussion as well, as we want to optimise data management once rather than for each experiment separately. His understanding is that flat budget mostly means hardware only. P Clarke agreed that commonality on data management is really needed, and mentioned CVMFS as an excellent example. Non-HEP communities (e.g. astronomy) need to be involved from the bottom up, otherwise they will not embark on our solutions. I Bird referred to his forthcoming presentation during the meeting, and said that there is already a collaboration agreement on data management with SKA.

E Elsen said that he would like to see the community formulate a request to funding agencies for targeted help with the CWP process, which does not necessarily mean additional money. Currently, we are not sufficiently prepared to make the statement that one MCHF would make a measurable difference; how can we get into such a position? D Britton commented that in the UK (probably like elsewhere), asking money for physics generators is not likely to succeed, but for data management developments of potential benefit to other communities as well, it may be successful.

J Templon said that that experiment-specific solutions are sometimes due to grassroots activities in the experiments; strong guidance in the experiments is hence required to avoid island solutions. E Elsen suggested that now the successful bottom-up CWP process needs to be complemented by a strong top-down incentive. The CWP has proved the willingness of the community, and has delivered a very comprehensive list of potential activities that could help improve the situation. We now need to choose and prioritise these activities. P Clarke wondered whether there are any possibly helpful statements we can make, given that if the statements and the requests are made correctly, there may be hope to get funding. E Elsen replied that the Resource Review Board (RRB) meetings are opportunities to address the funding agencies, but we need to have agreed what the statements and requests should be, which requires prior selection and prioritisation.

A Streit advised HEP to very seriously consider working with computer science and maths departments as well; at KIT they are already moving into this direction, as is CC-IN2P3, looking into the physics experiments' code from an algorithmic point of view. T Schulthess agreed that it is very important to formulate a project and a scope, and said that in view of past investments, an amount of one MCHF appears low; the right figure may perhaps be more in the vicinity of five to ten MCHF.



E Elsen replied that one MCHF was an arbitrary choice, guided perhaps by an overall expense of some 100 MCHF per year all over WLCG; the real amount depends on whether we can define the right projects. T Schulthess added that lab-directed R&D often amounts to about 5..6% of typical budgets, hence one MCHF really appears too small, but agreed that concrete proposals need to be defined. M Jouvin commented that HEP need to be pragmatic; with the CWP we can hope to move out of a period where asking for person-power was hopeless, as we now understand much better what needs to be done about software; it is however still a long process to define the projects well. E Elsen said that another complication is that the funding comes from many different sources and cannot simply be moved around. D Britton agreed that we need to choose and prioritise projects, and requested that the experiments are tightly involved with this. E Elsen agreed and mentioned that the LHCC is already trying to promote this, but we would need to avoid the risk of choosing HEP-specific solutions again. P Clarke expressed reservations against a top-down approach, as the stakeholders may not feel committed sufficiently. I Bird added that the strategy paper already mentioned in his earlier presentation is supposed to address exactly the point of prioritising software activities; collaborating with SKA means that the projects with exabyte-scale data volumes are already addressing data management in common. Software remains an issue; the world has moved on, but the software, not necessarily optimised for performance when written, has not. We need to identify very specific projects, just saying to invest one MCHF on software will not help at all. M Jouvin suggested that the focus should probably be on simulation, as it takes a very significant fraction of the computing resources; any efficiency improvement will help the whole community. A Petzold said that the major improvements we need must come from the experiment code and the common simulation packages, hence the experiments must be fully involved and supportive of the developments, even if they use tools also used by other communities. E Elsen replied that the experiments are indeed taking efficiency improvements a lot more seriously than in the past.

INFRAEOSC-04 (Open Science Cloud for ESFRI projects/landmarks) (I Bird)

I Bird presented ESCAPE, a proposal in reply to a call on data stewardship and open science in the context of the European Open Science Cloud (EOSC). The key objectives of the call are to address stewardship of data handled by research infrastructures according to the FAIR (findable, accessible, interoperable and re-usable) principle, and to ensure integration and interoperability of data and tools within the EOSC. Clusters of ESFRI infrastructures are eligible, other world-class research infrastructures may be part of a cluster as well. ESFRI landmark infrastructures such as HL-LHC, FAIR and SKA are expected to request about two MEUR each. Projects are expected to improve access to data and related tools in view of enabling a broad interdisciplinary research, and to suggest support activities, most of which our community is already involved with. The intended proposal is called ESCAPE (European Science Cluster for Astronomy and Particle physics ESFRIs). The cluster leverages the ASTERICS one comprising astronomy, astrophysics and astroparticle physics experiments, to which particle physics (HL-LHC and FAIR) have been added; the cluster hence represents a significant weight of the science communities concerned, linking them strongly as many WLCG Tier-1 sites provide services to the other science communities as well. It is hence in a very strong position to define the future EOSC data infrastructure. The proposed project, coordinated by Giovanni Lamanna from LAPP Annecy, is split into five work packages, of which WP2 on “data lake” infrastructure concepts is particularly relevant. The work package is jointly managed by Ian Bird and Simone Campana from CERN. Other work packages concern a repository of open-source



scientific software, interoperable archives and data preservation, and an open-science platform for providing open access to research data. The mission of the data lake work package, to prototype a reliable and scalable data infrastructure meets SKAs requirement very well, and is entirely in line with WLCG prototype work that was intended anyway. In addition, our community can contribute to other work packages wherever relevant, for example with DPHEP, service catalogues, volunteer computing, open data portals/open access, Zenodo, Invenio etc. Projects can run for up to 36 months, with a few additional months for getting started and winding down.

E Elsen thanked I Bird for the short presentation and emphasised that it is very natural to collaborate with astronomy, astrophysics and particle astrophysics, as we know them already and understand the common requirements well. There is just another month before the deadline for project proposals (22 March); the amount of person power we could hope to get is not entirely clear yet, but it should not be spread out too thinly. I Bird agreed, suggesting that funding agencies should not be active in too many areas.

Replying to a number of questions, I Bird clarified that the ESFRI landmark projects will receive the funding from the EU, and will forward it to the relevant funding agencies for project contributions. The experiments have been kept informed; further involvement was not necessary, as they were favourable towards data lake-related work anyway. They fully support the proposal as an attempt to ensure that EOSC delivers something useful. HPC centres and research network providers will be involved in the data lake workpackage via automatic data transfers within the centres, between them as well as with external centres including commercial clouds. Involving HPC centres is important in order to ensure that they are able to handle the high data volumes and bandwidths. However EOSC should be driven by science rather than infrastructure providers. Links with other related communities, for example CTA, are being established.

Round table

As many aspects had been mentioned already under “Lessons from the HSF CWP”, and time was running short, this point was dropped.

AoB and date of next meeting

E Elsen mentioned issues with licensed software and in particular Microsoft’s announcement to withdraw academic status from CERN, strongly impacting yearly licence costs. He explained that offering Microsoft-based services to all CERN personnel irrespective of their home institute may be questioned, and asked to what extent other laboratories depended on Microsoft and/or considered alternatives, in particular open-source ones. Following a first fact-finding round at CERN, it appears that the issue is linked much more with backoffice services such as Active Directory, Exchange, Sharepoint and Skype For Business rather than with Windows, Office and other products on personal workstations. He wondered whether such change would have an impact on users from participants’ countries, and whether representatives would consider migrating off Microsoft products together with CERN.

Several speakers explained that their institutes had chosen different solutions in a number of areas already, that they did not see a strong link with CERN’s migration strategy, and that they did not expect their users to be significantly affected by potential CERN server-side changes. In a number of cases, contacts with technical experts in CERN’s IT Department had been established already.



E Elsen thanked speakers for their contributions, pointed out that while negotiations with Microsoft had not been completed yet, CERN wished to consider alternatives in parallel, invited further input on the question off-line, and suggested to stay in contact about this issue.

E Elsen asked about the date for the next meeting. Several speakers expressed their feeling that the meeting was very useful (much more so than expected from the agenda), and that the frequency should at least be kept at two meetings per year, if not increased; a potential increase was however seen as problematic due to travel constraints. The fifth Scientific Computing Forum was scheduled for Thursday 20 September from 10:30 h to 13:00 h CEST (08:30 – 11:00 h UTC).

E Elsen thanked all participants for their attendance and contributions, and closed the meeting.