

6th Scientific Computing Forum

CERN, 25 February 2019, 14:00 h

Draft minutes

Indico event page: <https://indico.cern.ch/e/Scientific-Computing-Forum-2019-02-25>

Introduction and approval of the minutes (E Elsen)

E Elsen welcomed participants and reminded them that the purpose of the meeting is to provide for a forum for informal information exchange about topics of relevance in the context of scientific computing at CERN and the Member States.

Although the meeting is not formal and hence does not require formal minutes, E Elsen invited comments to the minutes of the previous meeting. As none had been received beforehand, and none were raised during the meeting, the minutes were considered final.

HPCs and Large Scale Scientific Instruments - FZ Julich's view (T Eickermann)

E Elsen welcomed T Eickermann from the Julich Supercomputer Centre, a unit of Forschungszentrum Julich (FZJ).

T Eickermann gave an overview of services and strategies of FZJ, highlighting the opportunities and obstacles for synergies between High-Performance Computing (HPC) and computing installations for large-scale scientific instruments. In an era when single-core performance does not increase any further, and accelerators as well as parallel execution promise to provide performance enhancements, new steps have to be taken. HPC resources are not optimised for HTC applications, requiring the latter to be adapted; available resources are usually fully committed; HPC applications usually are granted computing resources for one year following a peer review cycle; worker nodes usually do not have external Internet connectivity. He proposed to start with a pilot project similar to the one already launched with ESRF.

E Elsen thanked T Eickermann and commented that modular, heterogeneous facilities match HEP workloads well, which are dominated by simulation, but also have a large component of data-intensive calculations (e.g. reconstruction). T Eickermann replied that for further performance improvements, GPUs are currently the best choice, which implies Nvidia programmed via CUDA; there are promising efforts towards high-level programming, but they are not yet sufficiently mature.

Replying to a question by E Elsen, T Eickermann gave details of user support that consists of multiple levels covering all required aspects.

Replying to a question by M Girone, T Eickermann confirmed that work on a demonstrator of HPC services for HEP is going on in the PRACE context; GEANT will be invited to participate, too.

A discussion on software adaptation followed, during which it was mentioned that very large efforts are required, possibly similarly as 20 years ago, when codes needed to be ported from vector machines to parallel computers. Hopefully more low-level details will be encapsulated in higher-level libraries, encouraging communities to migrate that have been reluctant so far; currently the high-level libraries do not provide the required functionality, stability and/or performance. Developing higher-level open abstract interfaces may not be in the interest of the

two major commercial players on the market; hence it is important that major user communities from all concerned areas of scientific computing express their interests towards the manufacturers and public bodies such as the European Commission, taking advantage of the current interest in the “digital agenda”, even though it is not entirely clear how technical obstacles can be overcome, which exist in addition to the issue of the strategic direction driven by economic factors.

Replying to a question by M Girone, T Eickermann confirmed that current plans of modular computing go well beyond the DEEPEST project. Independently from the project, existing resources will continue to be used efficiently.

Replying to a question by E Elsen, T Eickermann said that FZJ see themselves both as experts on choosing hardware setups as well as a pool of software expertise.

Replying to a question by D Britton, P Buncic said that ALICE have indeed proven that investing into software adaptations pays off. For example, they managed to speed up the TPC reconstruction by a factor of 40, but not all applications could be accelerated that much.

“Scientific Software Institute” (S Roiser)

S Roiser presented the current status of ideas around a “Scientific Software Institute”, the proposed mandate of which is to establish a group of researchers placed at universities and research institutions to lead and engage in R&D and software engineering activities. Enabling training and qualification leveraging in view of developing career paths is key. HEP labs, universities and research labs and experiments should be stakeholders. A close collaboration should be envisaged with IRIS-HEP, CERN openlab and a number of EU projects. How to engage with other data-intensive sciences still needs to be clarified.

I Bird commented that the scope must be broader than HEP, which is already covered by HSF. E Elsen agreed and pointed out that new topics need to be introduced at universities into the academic context, which requires a few ‘early adopter’ universities to start establishing this together with other sciences, which must hence be involved from the start. A Streit emphasised the importance of clarifying our aims between better software, career paths and potentially others, which will help find the right contact persons. E Elsen added that the idea is a distributed organisation, not a “software house”, which would be at risk to see people go to industry rapidly. J Templon commented that care must be taken to distinguish national initiatives from experiment ones, and said that the overhead of the institute must essentially be zero.

P Clarke suggested the name “software engineering institute” rather than “software institute”, and agreed with previous contributions that the value generated by the institute must be clear. Previous lack of communication between physicists and computer scientists has not been due to missing willingness, but rather fully busy schedules; a funded activity is needed to focus the attention. A Zoccoli said that to a large extent, the knowledge is available already, but coordination and focus on a specific set of goals is required; each country must find it worth while to contribute actively.

E Elsen thanked all participants for the discussion and concluded that the topic will be brought again to a future SCF.

Cost evolution of compute and storage servers (B Panzer-Steindel)

B Panzer-Steindel presented the slides attached to the agenda in Indico.

Replying to a question by P Clarke, B Panzer-Steindel explained that the differences of growth rates observed in different countries are due to different architectural constraints and choices, mostly determined by the set of experiments and activities a given site supports. M Jouvin commented that a number or at least a range should be defined, as sites need it for planning purposes. C Bozzi emphasised the importance of reducing RAM requirements of the experiment applications in order to contain CPU costs. I Bird concluded that the discussion

will continue at the WLCG/HSF/OSG workshop in March, and that the situation is more complex now than the single number we have had previously; regular follow-ups are required.

Results from HNSciCloud (R Jones)

R Jones presented the slides attached to the agenda in Indico.

E Elsen suggested to postpone the data preservation topic to a future meeting, and noted that it is of relevance for a large spectrum of sciences rather than limited to HEP.

Replying to questions by D Britton and A Streit, R Jones said that a comprehensive report on TCO had been written, to which all members of the HNSciCloud Buyers' group have access, and that the comparison of costs of in-house solutions was incomplete due to sites' constraints to release information. F Hemmer added that the pricing of commercial offerings is highly non-transparent and unpredictable. R Jones commented that real use cases are required to deliver useful information.

Round table

E Elsen noted that at this meeting, software development and software engineering was a common theme even with HPC, which shows the need to move forward with the software engineering or scientific software institute. P Clarke added that the institute must be positioned as a place where new ideas from computer science can be put into practice. A Streit informed the meeting about a European initiative across computer science faculties (<https://www.informatics-europe.org/>) that we may wish to approach. P Elmer added that the NSF funding for IRIS-HEP also included money for common workshops with computer science; the experience was positive, but their scope was too broad, they should have focused more on specific topics. They found it difficult to explain to non-HEP people what our problems are; the collaboration was generally easier if the universities involved had both a computer science and a HEP department. He is interested to be involved and give advice.

Any other business and date of next meeting

The date of the next meeting was discussed. *Note added after the meeting: The next meeting will be held on 02 October 2019 at 10:00 h.*

Annex: Attendance

At CERN:

Maite Barroso Lopez
Ian Bird
Concezio Bozzi
Predrag Buncic
Simone Campana
Mihnea Dulea
Thomas Eickermann
Eckhard Elsen (chair)
Maria Girone
Frederic Hemmer
Andreas Hoecker
Robert Jones
Markus Klute
Thorsten Kollegger
Patricia McBride
Helge Meinhard (scientific secretary)
Bernd Panzer-Steindel
Stefan Roiser
Achim Streit

Via teleconferencing:

Tommaso Boccali
David Britton
Peter Clarke
Peter Elmer
Jose Flix Molina
Patrick Fuhrmann
Miguel Gila
Christophorus Grab
Volker Guelzow
Michel Jouvin
Matthias Kasemann
Mark Klein
Dirk Krucker
Eric Lancon
Pierre Etienne Macchi
Gonzalo Merino
Jean-Pierre Meyer
Pepe Salt (replacing Maria Garcia Borge)
Jeff Templon
Christoph Wissing
Antonio Zoccoli

Apologies:

Maria Garcia Borge
Kajari Mazumdar