Minutes of the 3rd Scientific Computing Forum
(CERN, Geneva, 27 October 2017)

Present:
M. Barroso Lopez (CERN, IT Department)
I. Bird (CERN, IT Department)
T. Boccali (representing CMS)
P. Buncic (representing ALICE)
D. Britton (University of Glasgow, UK)
J. Butler (representing CMS)
P. Clarke (University of Edinburgh, UK)
J. Cochran (BNL, USA)
I. Collier (STFC RAL, UK)
E. Elsen (CERN, Director for Research and Scientific Computing, Chair)
J. Flix Molina (PIC/Ciemat, Spain)
S. Foffano (CERN, IT Department, Scientific Secretary)
F. Hemmer (CERN, IT Department)
R. Jones (University of Lancaster, UK)
M. Jouvin (LAL Orsay, France)
T. Kollegger (GSI, Germany)
E. Lancon (BNL, USA)
M. Martinez-Perez (University of Barcelona, Spain)
P. Mato Vila (CERN, EP Department)
K. Mazumdar (Tata Institute of Fundamental Research, India)
B. Roser (FNAL, USA)
H. Sakamoto (University of Tokyo, Japan)
T. Schulthess (CSCS, Switzerland)
E. Sexton-Kennedy (representing CMS)
O. Smirnova (Lund University, Sweden)
G. Stewart (CERN, EP Department)
V. Vagnoni (INFN Bologna, Italy)
T. Wenaus (representing ATLAS)

Remote:
S. Alam (CSCS-ETHZ, Switzerland)
V. Beckmann (CNRS/IN2P3, France)
K. Bloom (FNAL, USA)
C. Broekema (ASTRON, Netherlands)
P. Calafiura (BNL, US)
K. De (University of Texas, US)
P. Elmer (Princeton University (US)
C. Grab (ETH Zurich, Switzerland)
V. Guelzow (DESY, Germany)
A. McNab (University of Manchester, UK)
J. P. Meyer (CEA Saclay, France)
F. Ould-Saada (University of Oslo, Norway)
A. Petzold (KIT, Germany)
J. Shank (Boston University, USA)
J. Templon (NIKHEF, Netherlands)
A. Valassi (CERN, IT Department)

Introduction and approval of the minutes (E. Elsen)

E. Elsen welcomes participants to the third Scientific Computing Forum (SCF) reminding them the meeting is informal and open to encourage the relevant experts to attend for discussion on topical scientific computing subjects. The size of the meeting must remain appropriate to ensure discussion at the right level. All information about the forum is publically available from the SCF Indico page: https://indico.cern.ch/category/9249/.

There are no comments on the minutes of the second meeting which are therefore considered approved.

EC Interactions (E. Elsen)

For the large infrastructures there is one pending call INFRA-EOSC-04-2018 available for ESFRI–listed projects or landmarks for which HL-LHC, FAIR, KM3NeT and other physics projects are eligible. The main question is how to form alliances in this context.

Other interactions, restricted to particle physics, are the annual meetings between the European Commission and CERN where topics such as computing developments and the Open Science Cloud are discussed. As mentioned at the last SCF meeting, a position paper "The Federated Scientific Data Hub" was submitted to the Commission summarizing plans for the EIROForum organisations. The intention is to develop a coherent framework for computing in the future, however much discussion is still required.

On-going projects such as HNSciCloud and the EOSC-Pilot initiative are progressing.

Brief WLCG status report for non WLCG Overview Board members (I. Bird)

I. Bird very briefly summarises WLCG and the experiments are working well, the resources needed for 2017 have coped well with the high LHC performance in the past couple of months, particularly for pileup and luminosity. Concerning the Community White Paper there has been significant effort from the community to be used for alignment of the future WLCG strategy for HL-LHC. Since the last SCF meeting a formal Collaboration Agreement was signed with SKA, and the first Collaboration meeting has taken place.

E. Elsen thanks I. Bird adding the short summary is to avoid repeating the more detailed information presented to many of the same participants in the earlier WLCG Overview Board (OB) meeting, and invites participants to refer to the OB Minutes for more information. There are no comments or questions.

US perspective on Scientific Computing over the next decade (R. Roser, E. Lancon)

The US HEP computing environment representing 20% of the HEP budget and contributing to about 20% of the LHC CPU and Disk requirements is presented. The Office of Advanced Scientific Computing Research (ASCR), DOE High Energy Physics and National Science Foundation are the three agencies funding computing in the US for Particle Physics.

The P5 Strategy document for US recognized the importance of computing and strengthening global collaboration between laboratories and universities, and initiated the HEP Centre for Computational Excellence to develop cross-cutting solutions. Leveraging HPC facilities, cloud computing, evolution and maintenance of the software base are important aspects of the future computing environment.

Following an exascale requirements review in 2016, a strategic computing initiative sponsored by the US Government aims to have one exascale machine on the floor by 2021, and two more by 2023 increasing the future amount of compute capacity very significantly.

HPC platforms will offer unique computational opportunities however investment is needed as such machines are not designed for HEP compute. A combination of HPC resources, HEP owned GRID-like resources and commercial clouds are thought to be the approach for HEP computing in the next decade, with a lot of investment needed to modernise the software stack.

E. Elsen thanks R. Roser for his presentation and mentions that the DOE recently agreed to officially request computing support from ASCR.

Points arising from the round table discussion

There are comments about the significant number of Full-time Equivalents (FTEs) required to work on the adaptation of the middle-layer software for HPC estimated at 150/experiment/site.

The approaches in the UK, Italy, Switzerland and France with respect to use of HPC machines are outlined: politics differ from country to country and there is a danger that decision makers think a few HPC machines will solve all the computing problems. HPC certainly seems to be part of the future computing solution, the real question is, how large a part? In order for the HPC strategy to work correctly, at least for Europe, CERN should take a strong coordination role.

Examples were shared of HEP jobs running on supercomputers, or in Switzerland the move of the Tier 2 to a supercomputer which is seen to be an efficient way to perform data analysis for HEP. There is some debate about this, HEP software was not written for HPC architectures, so how realistic is it that around 12 million lines of code can run efficiently in supercomputing architectures? Supercomputing architectures at the processor level are the same as the commodity computing architectures of the grid, therefore software that does not run efficiently on a supercomputer at the node level, will not run efficiently anywhere else.

From the economy of scale point of view, supercomputers are advantageous as the hardware procurement is cheaper. Making more efficient use of the architectures as they develop is becoming increasingly important, so when the aim is to compute more for the same, or less, cost in future, the challenge is to identify where the performance comes from and see the 12 million lines of code as an opportunity for an efficiency gain structured around service-oriented architectures, and goal-oriented computing.

In the UK, 7 or 8 HPC machines are optimized for different uses such as shared memory and high-speed interconnects. It is not seen to be cheaper to use an expensive machine architected to solve complex and tightly coupled problems. The observation so far is that the effort taken to get HEP workflows running on HPC machines is substantial, and each change of HPC machine represents a new set of problems. There has been a lot of effort in the US to get the code to work on other machines. It is not felt to be a code portability problem, rather how to use the HPC as a plug-in to the infrastructure for data management, data transfer and metadata. The situation in France is very similar to that of the UK.

E. Elsen suggests listening to the next two presentations before continuing the debate.

UK perspective on Scientific Computing over the next decade (D. Britton, P. Clarke)

UK Computing for Particle Physics, GridPP, which represents 10% of WLCG, is presented comprising 18 sites: the Tier 1 at RAL, and four distributed Tier 2 centres linked via the Janet core network rather than LHCONE. 10% CPU and 5% storage is available to non-LHC VOs: LHC VO CPU usage is currently at 93%.

Funding comes from the Science and Technology Facilities Council (STFC) supporting programmes in HEP, Astronomy, Astro-Particle and Nuclear, and national facilities including Diamond Light Source, ISIS for neutrons and the central laser facility. Following a strategic computing review published in December 2015, in view of the significant future increase in data volumes and funding limitations, scientific disciplines supported by the STFC were encouraged to work more closely together and share the computing infrastructures. Some examples of this include SKA, LSST, Advanced LIGO and EUCLID.

In April 2018, the funding agency structure will change, all UK Research Councils, including STFC, will belong to a single organisation; UK Research and Innovation (UKRI). Future directions include shared computing across UKRI, working more closely with large astronomy projects, progress towards a national e-infrastructure for research, and a push towards the Cloud where appropriate.

Evolution of the GridPP Tier 2 sites is described which includes consolidating disk at specific sites; four for ATLAS and one for CMS, dedicating smaller sites to CPU with an appropriately sized disk cache, and for HEP creating the UK data lake from the Tier 1 and five large Tier 2 sites.

E. Elsen thanks the speakers and suggests deferring discussion until after the presentation from Spain.

Spanish perspective on Scientific Computing over the next decade (J. Flix, M. Martinez)

Spain participates heavily in ATLAS, CMS, LHCb, also astro-particle and cosmological projects such as MAGIC, EUCLID and CTA. As one of the initial contributors, Spain has offered strong and reliable contributions to WLCG representing around 5% via the multi-experiment Tier 1 and experiment-specific Tier 2's. Following an international review of the performance of the Spanish tier sites in 2015 and resulting recommendations, and in line with the general evolution of WLCG and trends to integrate HPC and commercial clouds, the strategy is to provide a modern, powerful and sustainable data and compute platform linking the most reliable and cost-effective sites in Spain, including a portal to integrate external resources for CPU intensive workflows.

Supercomputing is very strategic in Spain, with good financial investment and support from the government. Proposals have been made to exploit the MareNostrum4 HPC facility at BSC, Barcelona with interest expressed for participation of the supercomputer in LHC; potentially a very useful opportunity for the future. Contacts are already ongoing using ATLAS and CMS input.

Spain is very interested to maintain an important role in WLCG computing preparing for HL-LHC. The network is thought to be adequate, prototyping is beginning to federate WLCG resources and to use and integrate HPC and Cloud resources, with substantial funding from the Spanish Research Plan. To shape the future, synergies with other scientific projects and trans-national projects are needed, in addition to a sustainable WLCG HL-LHC model. It is hoped that the Scientific Computing Forum can play an important role in making this happen.

E. Elsen thanks the presenters.

Points arising from the round table discussion

E. Elsen summarises the last 2 presentations as an aggregation towards larger centres rather than scattered centres, reducing personnel and hardware costs. This is a repeating pattern with respect to the presentations made at the 2nd SCF meeting. With the development towards BSC and the earlier discussions on use of supercomputers to download the binary on the supercomputer, how can supercomputers be best used for HEP purposes ?

The model being followed for BSC is for them to realize just how complex the HEP use case is; as a lot of CPU will be used they have expressed an interest in data management. BSC will learn that managing LHC data in a sustainable way is quite different from data management for biology or astronomy projects.

The experiences of moving the Swiss Tier 2 to the flagship supercomputer and the experience at Oakridge prove that despite the different technical details it is possible to move middleware onto a modern supercomputer, make use of modern technology for queuing and use of the resources, and as a result run ATLAS and CMS analysis. Containers seem to solve the scalability issue of delivering the software to the computing nodes. The workflow requires rethinking, as an example the ATLAS event service uses the concept of trying to be stateless to be able to run on different HPCs. It is not so much the technology that is the main issue, rather the openness and speed of the different centers to deploy new software differently; some centers move much faster then others and the overall cost depends on how the HPC procurements are made.

The US exascale community is very interested in data intensive science, therefore if the HEP community is proactive, engaging and willing to provide HEP benchmarks, it can influence the future architectures. There is agreement that collaboration is needed in the future to ensure the software will work hand-in-hand with the hardware, and a request that the SCF pushes in this direction.

There is more debate on the cost-effectiveness of HPC. In most countries HPC is more expensive than HTC, however moving a small Tier 2 into an existing large HPC center is understandably cost effective, as there are savings from the infrastructure costs and operational manpower, however investing in HPCs may not be the most cost effective long-term solution.

E. Elsen comments on the PRACE trans-national initiative bringing several of the larger installations together. Although the architectures are different, how can the infrastructure be used to minimize the effort required to run the software of those structures?

PRACE has gone through a crisis for a few years due to the need to transition towards supercomputing operations. In Europe the EC may invest heavily in certain HPC systems therefore it is very important that the science and HEP communities make their needs clear to ensure science driven machines are among the top HPC machines in future. In September a meeting was organized between CERN and PRACE (https://indico.cern.ch/event/666963/) the outcome of which was to organize a workshop for people to share experience on their use of HPC machines. The date is to be confirmed, however this should be reported on at a future SCF.

In summary E. Elsen points out that heterogeneous hardware is a fact of life although HPC seems of interest, is developing quickly, is highly-regarded therefore attractive for funding, and may be useful in particular for deep learning. The hurdles are the number of FTE's and the investment in the software of the middleware. More in depth discussion is needed on some of the issues from this meeting and the discussion with PRACE needs to be followed up on – plenty of food for thought and material for future discussions in the next SCF meetings.

HEP Software Foundation Community White Paper (G. Stewart)

The Community White Paper (CWP) process is presented which has engaged over 250 people producing over 300 pages of detailed analysis in the following areas:

- Conditions Databases
- Data analysis and interpretation
- Data and software preservation to enable reuse
- Data organisation, management and access
- Data-flow processing framework
- Detector simulation
- Facilities and distributed computing
- Machine learning
- Physics generators
- Software development, deployment, validation and verification
- Software trigger and event reconstruction
- Visualisation

The summary CWP Roadmap 60-page document projects a path forward and identifies the main areas needing investment for the future to support the HL-LHC Computing TDRs. The current first draft is undergoing refinement by the CWP Editorial Board and will conclude within a few months. In addition to defining the future strategy and direction, the CWP Roadmap has enabled community consensus, which in itself is an achievement.

The next steps involve focusing the R&D efforts in the community, shaping the current effort, and helping to attract new investment in critical areas. The HEP Software Foundation will play a vital role in spreading knowledge of new initiatives, encouraging collaboration and monitoring progress. Workshops are planned for 2018 including sessions before CHEP.

E. Elsen thanks G. Stewart for the overview and comments on the very readable and comprehensive account of what needs to be addressed in future, encouraging SCF participants to read the document and make comments. Some of the areas relevant for the SCF should be selected for continued discussion in a future meeting.

AoB and date of next meeting

It is agreed to hold the next meeting of the SCF in February 2018 to discuss certain CWP areas, further in-depth analysis of HPC, and other topics as necessary. [After the meeting 23rd February 2018 at 10:30 was confirmed for the 4th SCF]

E. Elsen thanks participants for their attendance and contributions, and closes the meeting.