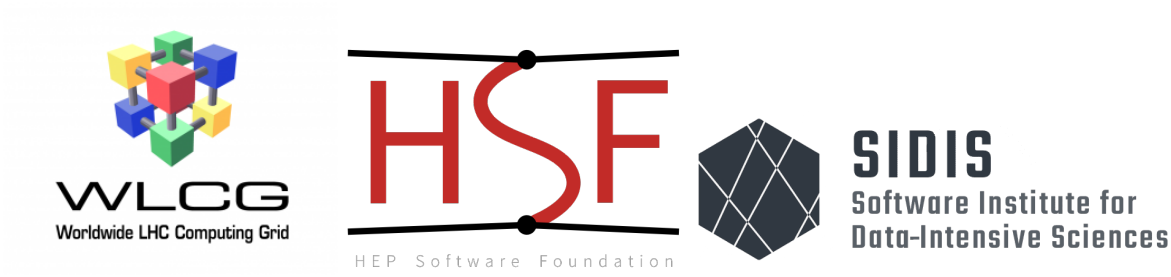

Dear Davide, Luigi

This letter is in support of the ExaTEPP proposal being submitted by the UK Particle Physics community to the ExCALIBUR programme.

In the coming decade, major new opportunities to advance our understanding of the fundamental building blocks of nature and their interactions will come online. These include the upgraded High Luminosity Large Hadron Collider (HL-LHC) at CERN, the European Laboratory for Particle Physics, in Geneva, Switzerland, and the Deep Underground Neutrino Experiment (DUNE) at Fermilab, in the US. The Higgs Boson discovery in 2012 was the capstone of a decades long effort to experimentally confirm the so-called Standard Model (SM) of particle physics. The SM describes, often with breathtaking precision, nearly all experimental observations to date. But it is known to be incomplete, and these new facilities are key to opening windows on what lies beyond the Standard Model and an even deeper understanding of the physical world.

One of the key challenges in exploiting the investments in these cutting-edge facilities will be managing the deluge of data which they will produce. Computational and data science has long been a critical enabler of scientific progress, especially at particle physics facilities. However, the planned scale of the jump in data volume and complexity is beyond anything previously seen and beyond what could be achieved simply by Moore's Law gains in capability per unit cost. Even this assumes that Moore's Law will hold, and it is already clear that exploiting what industry can provide in terms of new processor architectures (including “accelerators”) and storage technologies will require significant R&D in the software and algorithms used. An order of magnitude improvement in cost-performance was required, both for computation and data.

As plans for the facilities of the 2020s became clear, the particle physics community recognised the scale of the problem it was facing. One of the strengths of the global particle physics community is its capability to build international collaborations and make them effective at pursuing grand challenges. As a first step, the community worked together during 2017 to produce a white paper "A Roadmap for HEP Software and Computing R&D for the 2020s" (https://doi.org/10.1007/s41781-018-0018-8) to describe the efforts that are required to maximise the physics reach of the new facilities. Several things became clear during the process that led to the R&D roadmap. First, the scale of the problem requires significant investments of resources as a real "software upgrade" to accompany the planned hardware upgrades. Second, as is often the case in particle physics, the effort is larger than any single institution or country. Collaboration is key and multiple national partners will need to step up and work together to accomplish the full set of R&D activities in the roadmap, both in terms of raw resources and intellectual input.

The ExaTEPP proposal is addressing exactly this space. The proposal will tackle several of the key areas identified in the roadmap, all of which are critical to the future of HEP and in particular the LHC Run-3. These include: Intelligent data and workflow management focussing on the development and use of data management systems and data lakes that forms a key part of the strategy; Simulation, and in particular the deployment of Geant4 on accelerators which is again a key area to allow the LHC to produce the simulated data it requires within the available computing resources in Run-3; and improvements to reconstruction using accelerators and novel algorithms, which is highly relevant not only to the LHC but the future neutrino experiments that have a different pattern recognition problem.

Collaboration with the wider scientific community and with industry was identified as crucial by the recently released European Strategy for particle physics document. The participation of our UK colleagues to the wider ExCALIBUR initiative is addressing exactly this. It is also very good to note that the proposal intends to put substantial weight on training and education with a planned participation in the SIDIS initiative and into work with its partners.

This letter of support is sent on behalf of three major International Particle Physics Computing bodies:

- The Worldwide LHC Computing Grid: The WLCG is a global collaboration that coordinates and manages the infrastructure and services for the data processing and analysis of the data from the LHC. It is a collaboration between CERN, the four LHC experiments, and national computing contributions from more than 40 countries, providing a distributed computing infrastructure at unprecedented scale. Looking forward to the HL-LHC upgrade in a few years, the challenges of multi-Exabyte data volumes, software performance, and technology evolution require significant investment in R&D within the global particle physics community, and the WLCG relies on relevant investments and initiatives at national level.

- The HEP Software Foundation: The HSF is a community organisation of researchers in High-Energy Physics across many experiments (LHC, Belle II, DUNE and beyond) who are focused on improving software for HEP beyond single experiment solutions. The HSF authored the community white paper roadmap cited above and now works with partners to deliver the software improvements needed for HL-LHC, sharing knowledge, best practice and tools, and striving for the most general solutions for multiple experiments, including going beyond traditional CPU based algorithms.

- The SIDIS virtual institute aims to establish a supporting environment for interactions between natural science and computer science, which will foster closer and better structured collaboration between the two fields. This will help natural sciences to make best use of current computing hardware and allow computer science to tackle concrete problems and to refine and develop concepts on real physical science data.

We fully support the ExaTEPP proposal as an essential component of the global Particle Physics ecosystem, and one which will bring the particular strengths of the UK to bear on this problem. It is completely aligned with the global roadmap, and the proponents will work coherently with the existing global efforts represented here. Should it be funded it will make

an important contribution to the evolution of software for exascale distributed scientific computing, and we are sure this will have implications well outside of Particle Physics.

Stefan Roiser
SIDIS Coordinator

Simone Campana
Project Leader of the Worldwide LHC Computing Grid (WLCG)

Graeme Stewart
For the HEP Software Foundation (HSF)