Facilitating Collaboration with HPC to Support Data Intensive Science
A Snowmass Letter of Interest for the Computational Frontier

Maria Girone, CERN openlab CTO

August 7, 2020

The field of High Performance Computing (HPC) is undergoing a transition to the next major phase of its development, namely that of exascale computing. The first systems with such capabilities are slated for the early 2020s with numerous development efforts underway in Europe, the United States, China and Japan. In anticipation of the delivery of these systems, the global HPC community is broadening its horizons: as well as providing a step change in capability for its traditional user base (computational fluid dynamics, quantum chemistry etc.) exascale systems will need to provide the e-Infrastructure required by large experimental facilities that are due to generate unprecedented volumes of data as new capabilities come on line in the next decade. In addition, exascale platforms will be required to run artificial intelligence/machine learning (AI/ML)-related workloads as the fields of HPC and machine/deep learning continue to converge. ML/DL techniques will be widely adopted to enhance both scientific analysis and improve the performance, usability and reliability of the underlying platforms.

Ensuring effective collaboration between HPC experts, scientists and engineers working on world leading 'big science' experiments, and the national and international bodies that facilitate the underlying e-Infrastructure to move data across the world is therefore critical for the successful delivery of frontier projects. The challenges and opportunities of the exascale computing era and its convergence with AI/ML are the motivation for enhancing collaboration between the HPC and data intensive science communities.

CERN has recently signed a collaboration agreement between CERN, SKA, PRACE and GÉANT, with the goal of building collaborations between entire science projects and infrastructure provider communities. There have been a number of very successful projects to connect individual experiments with individual HPC sites. We expect that by engaging at the community level we will achieve more general solutions to common challenges in HPC. The current collaborative agreement has a European focus, bringing together PRACE and GÉANT with two large international science projects. We would like to prove that this model of collaboration is beneficial across continental boundaries and applicable to other scientific disciplines.

The collaboration will execute a series of technical demonstrators to make progress on the identified challenges of using HPC facilities for data intensive science. Software improvements for heterogenous architectures, benchmarking and accounting, data management and access are high priority and high impact development projects. The demonstrators are intended to show the feasibility of a technical capability and to define the R&D program. A benchmarking proof-of-concept will be established to measure the performance of different computing resources for workflows of SKA and HEP. A data access demonstrator is in development to show that we can effectively use HPC sites for data intensive science. The pilot would include the ability to deliver and validate a multi-petabyte dataset to local HPC storage significantly faster than the expected processing time. To show that the workflows performed by

members of the scientific communities can be securely supported by HPC, a demonstrator on authenticated and authorised access to resources, including job submission, is being defined.

The parties signed the formal agreement this summer and a dedicated technical workshop is organized for September. There will be regular updates and reports as the activities progress. We are interested in expanding the collaboration and sharing across additional HPC communities globally and additional science projects.