

Exploitation of HPC Resources for data intensive sciences

Wednesday, 19 May 2021 18:06 (13 minutes)

The Large Hadron Collider (LHC) will enter a new phase beginning in 2027 with the upgrade to the High Luminosity LHC (HL-LHC). The increase in the number of simultaneous collisions coupled with a more complex structure of a single event will result in each LHC experiment collecting, storing, and processing exabytes of data per year. The amount of generated and/or collected data greatly outweighs the expected available computing resources. In this paper, we discuss efficient usage of HPC resources as a prerequisite for data-intensive science at exascale. We discuss the work performed within the contexts of three EU-funded projects, DEEP-EST, EGI-ACE and CoE RAISE, with primary focus on three topics that emphasize the areas of work required to run production LHC workloads at the scale of HPC facilities. First, we discuss the experience of porting CMS Hadron and Electromagnetic calorimeters to utilize Nvidia GPUs; second, we look at the tools and their adoption in order to perform benchmarking of a variety of resources available at HPC centers. Finally, we touch on one of the most important aspects of the future of HEP - how to handle the flow of PBs of data to and from computing facilities, be it clouds or HPCs, for exascale data processing in a flexible, scalable and performant manner. These investigations are a key contribution to technical work within the HPC collaboration among CERN, SKA, GEANT and PRACE.

Primary authors: GIRONE, Maria (CERN); SOUTHWICK, David (CERN); KHRISTENKO, Viktor (CERN)

Co-authors: FONTES MEDEIROS, Miguel (CERN); GIORDANO, Domenico (CERN); HOEGSTOEYL, Ingvild Brevik (Norwegian University of Science and Technology (NTNU) (NO)); ATZORI, Luca (CERN)

Presenter: SOUTHWICK, David (CERN)

Session Classification: Facilities and Networks

Track Classification: Distributed Computing, Data Management and Facilities