Contribution ID: 45

Experience running Engineering Applications at CERN on HPC clusters in HTCondor and SLURM

Thursday, 27 August 2020 14:30 (30 minutes)

CERN IT department has been running two Linux based computing infrastructures in HTCondor and SLURM for many years. HTCondor resources are used for general purpose parallel but single-node type jobs, providing computing power to the CERN experiments and departments for tasks such as physics event reconstruction, data analysis and simulation. For HPC workloads that require multi-node parallel environments for MPI programs, there is a dedicated HPC service with MPI clusters running under the SLURM batch system and dedicated hardware with fast interconnects.

Engineering users at CERN need to run critical simulations in very different domains. They use applications like CST, COMSOL and Ansys. These simulations are very demanding in terms of computing power and storage. In the past, there used to be a dedicated Windows based HPC cluster that was running for five years. However, the Windows HPC infrastructure was eventually decommissioned in 2019 to consolidate all computing resources under Linux. Since mid 2018, engineering users at CERN were migrated to run their simulations in the HTCondor and SLURM clusters. The change of infrastructure implied some technical and human challenges like the lack of Linux expertise among engineering teams, the lack of application specific knowledge on the IT side, and the fact that HTCondor and SLURM were not supported by CST, COMSOL or Ansys.

After a successful migration of CST, COMSOL and Ansys to Linux, the challenge has changed focus to running the simulations in the most optimized way, to make the most of the available computing resources. Some of the tasks where the IT team has worked in close collaboration with the engineers are: fine tuning the applications to reduce I/O access, understanding how to calculate the maximum number of cores to gain on processing time, integrating the Windows GUI interface to submit jobs to Linux and learning how to debug problems. In this contribution we will describe how we have dealt with all these challenges to offer a production computing infrastructure that meets the engineering users needs.

Primary authors: ALANDES PRADILLO, Maria (CERN); Dr LLOPIS SANMILLAN, Pablo (CERN)

Presenter: Dr LLOPIS SANMILLAN, Pablo (CERN)

Session Classification: Converging High Performance infrastructures: Supercomputers, clouds, accelerators Session

Track Classification: Physics (including HEP) and Engineering Applications