

# Integrating HPC into an agile and cloud-focused environment at CERN

*Thursday, 12 July 2018 14:30 (15 minutes)*

CERN's batch and grid services are mainly focused on High Throughput computing (HTC) for LHC data processing. However, part of the user community requires High Performance Computing (HPC) for massively parallel applications across many cores on MPI-enabled infrastructure. This contribution addresses the implementation of HPC infrastructure at CERN for Lattice QCD application development, as well as for different types of simulations for the accelerator and technology sector at CERN. Our approach has been to integrate the HPC facilities as far as possible with the HTC services in our data centre, and to take advantage of an agile infrastructure for updates, configuration and deployment. The HPC cluster has been orchestrated with the Openstack Ironic component, and is hence managed with the same tools as the CERN internal cloud. Experience and benchmarks of MPI applications across Infiniband with shared storage on CephFS is discussed, as well the setup of the SLURM scheduler for HPC jobs with a provision for backfill of HTC workloads.

**Primary authors:** LINDQVIST, Carolina (University of Helsinki (FI)); LLOPIS SANMILLAN, Pablo (CERN); Mr GANZ, Philippe (CERN); HOIMYR, Nils (CERN); VAN DER STER, Dan (CERN)

**Presenter:** LLOPIS SANMILLAN, Pablo (CERN)

**Session Classification:** T7 - Clouds, virtualization and containers

**Track Classification:** Track 7 –Clouds, virtualization and containers