



Contribution ID: 254

Type: Talk

## Integrating the Perlmutter HPC system in the ALICE Grid

Thursday 24 October 2024 16:33 (18 minutes)

The Perlmutter HPC system is the 9th generation supercomputer deployed at the National Energy Research Scientific Computing Center (NERSC). It provides both CPU and GPU resources, offering 393216 AMD EPYC Milan cores with 4 GB of memory per core, for CPU-oriented jobs and 7168 NVIDIA A100 GPUs. The machine allows connections from the worker nodes to the outside and already mounts CVMFS for users who need to access software from it. These two options make Perlmutter an ideal candidate for integrating into Grid infrastructures.

Due to the specific highly parallel and massive CPU and memory requirements of the native payloads running on supercomputers, there is always an idle part of the computing capacity. Conversely, Grid payloads require few CPU cores for a single task and can take advantage of the idle resources. This 'backfill' is advantageous both for the supercomputer operators, increasing the overall use efficiency of the machine and for the Grid users, allowing them to opportunistically use a substantial amount of CPUs. ALICE takes advantage of these conditions, the architecture of the Perlmutter supercomputer, and facilities offered by NERSC by deploying a standard Grid interface to Perlmutter through the NERSC SuperFacility API scheduling tool to submit and monitor normal Grid payloads. Perlmutter has been integrated into the ALICE Grid, running Monte Carlo simulation, with measurements and tests having been made to also integrate analysis jobs connecting to an EOS instance hosted at LBNL shared with the main Tier 2 site. The resulting HPC-based Grid site has proven to be a reliable resource contributor to the ALICE Grid, providing 8000 cores on average, with its only constraints being the short lifetime of jobs and the current time allocation from NERSC.

This paper describes the path taken to integrate Perlmutter in the ALICE Grid and the usual modifications needed to integrate HPC facilities into the standard Grid infrastructure.

**Authors:** WEISZ, Sergiu (National University of Science and Technology POLITEHNICA Bucharest (RO)); WEISZ, Sergiu (Lawrence Berkeley National Lab)

**Co-authors:** GRIGORAS, Costin (CERN); CHAKABERIA, Irakli (Lawrence Berkeley National Lab. (US)); BETEV, Latchezar (CERN)

**Presenters:** WEISZ, Sergiu (National University of Science and Technology POLITEHNICA Bucharest (RO)); WEISZ, Sergiu (Lawrence Berkeley National Lab)

**Session Classification:** Parallel (Track 4)

**Track Classification:** Track 4 - Distributed Computing