

Enabling production HEP workflows on Supercomputers at NERSC

Monday, July 9, 2018 11:30 AM (15 minutes)

Many HEP experiments are moving beyond experimental studies to making large-scale production use of HPC resources at NERSC including the knights landing architectures on the Cori supercomputer. These include ATLAS, Alice, Belle2, CMS, LSST-DESC, and STAR among others. Achieving this has involved several different approaches and has required innovations both on NERSC and the experiments' sides. We detail the approaches taken, comparing and contrasting the benefits and challenges. We also describe the innovations and improvements needed particularly in the areas of data transfer (via DTNs), containerization (via Shifter), I/O (via burst buffer, Lustre, or Shifter per-node-cache), scheduling (via developments in SLURM), workflow (via grid services or on-site engines), databases, external networking from compute nodes (via a new approach to networking on Cray systems), and software delivery (via a new approach to CVMFS on Cray systems). We also outline plans, and initial development, for future support of experimental science workloads at NERSC, via a 'Superfacility API' that will provide a more common, plug-and-play base for such workflows, building on best practises to provide a lower bar of entry to HPC for new experiments as well as consistency and performance.

Primary authors: GERHARDT, Lisa (LBNL); Dr MUSTAFA, Mustafa (Lawrence Berkeley National Laboratory); Mr LEE, Rei (LBNL); CANON, Shane (NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER); BHIMJI, Wahid (Lawrence Berkeley National Lab. (US))

Presenter: BHIMJI, Wahid (Lawrence Berkeley National Lab. (US))

Session Classification: T3 - Distributed computing

Track Classification: Track 3 – Distributed computing