

Using Shifter to Bring Containerized CVMFS to HPC

Wednesday, October 12, 2016 12:30 PM (15 minutes)

Bringing HEP computing to HPC can be difficult. Software stacks are often very complicated with numerous dependencies that are difficult to get installed on an HPC system. To address this issue, amongst others, NERSC has created Shifter, a framework that delivers Docker-like functionality to HPC. It works by extracting images from native formats (such as a Docker image) and converting them to a common format that is optimally tuned for the HPC environment. We have used Shifter to deliver the CVMFS software stack for ALICE, ATLAS, and CMS on the Edison and Cori supercomputers at NERSC. As well as enabling the distribution of TBs of software to HPC, this approach also offers performance advantages. We show that software startup times are significantly reduced (up a factor of 4 relative to the Lustre file system in the ATLAS case) as well as scaling with minimal variation to 1000s of nodes. We will discuss how this was accomplished as well as future outlook for Shifter and HEP at NERSC.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Virtualization

Primary Keyword (Mandatory)

High performance computing

Primary author: GERHARDT, Lisa (LBNL)

Co-authors: JACOBSEN, Douglas (Lawrence Berkeley National Laboratory); PORTER, Jeff (Lawrence Berkeley National Lab. (US)); CANON, Shane (NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER); TSULAIA, Vakho (Lawrence Berkeley National Lab. (US)); BHIMJI, Wahid (Lawrence Berkeley National Lab. (US))

Presenter: GERHARDT, Lisa (LBNL)

Session Classification: Track 6: Infrastructures

Track Classification: Track 6: Infrastructures