

Using Docker in HEP

Thursday 13 October 2016 16:30 (15 minutes)

Docker is a container technology that provides a way to “wrap up a piece of software in a complete filesystem that contains everything it needs to run” [1]. We have experimented with Docker to investigate its utility in three broad realms: (1) allowing existing complex software to run in very different environments from that in which the software was built (such as Cori, NERSC’s newest supercomputer), (2) as a means of delivering the same development environment to multiple operating systems (including laptops), and allowing the use of tools from both the host and container system to their best advantage, and (3) as a way of encapsulating entire software suites (in particular, a popular cosmology-based MCMC parameter estimation system), allowing them to be supported for use on multiple operating systems without additional effort. We report on the strengths and weaknesses of Docker for the HEP community, and show results (including performance) from our experiments.

[1] “What is Docker?”, <https://www.docker.com/what-docker>.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Data processing workflows and frameworks/pipelines

Primary Keyword (Mandatory)

Virtualization

Author: PATERNO, Marc (Fermilab)

Co-authors: LYON, Adam (Fermilab); KOWALKOWSKI, Jim (Fermilab)

Presenter: PATERNO, Marc (Fermilab)

Session Classification: Posters B / Break

Track Classification: Track 7: Middleware, Monitoring and Accounting