



Contribution ID: 376

Type: Talk

Efficient and fast container execution using image snapshotters

Tuesday 22 October 2024 14:06 (18 minutes)

A large fraction of computing workloads in high-energy and nuclear physics is executed using software containers. For physics analysis use, such container images often have sizes of several gigabytes. Executing a large number of such jobs in parallel on different compute nodes efficiently, demands the availability and use of caching mechanisms and image loading techniques to prevent network saturation and significantly reduce startup time. Using the industry-standard containerd container runtime for pulling and running containers, enables the use of various so-called snapshotter plugins that “lazily” load container images. We present a quantitative comparison of the performance of the CVMFS, SOCI, and Stargz snapshotter plugins. Furthermore, we also evaluate the user-friendliness of such approaches and discuss how such seamlessly containerised workloads contribute to the reusability and reproducibility of physics analyses.

Primary authors: THUNDIYIL, Amal Santosh (Sardar Patel Institute Technology of Technology (SPIT) (IN)); LANGE, Clemens (Paul Scherrer Institute (CH)); FEICHTINGER, Derek (Paul Scherrer Institute (CH)); BLOMER, Jakob (CERN); FATOUROS, Max (Paul Scherrer Institute (CH)); VOLKL, Valentin (CERN)

Presenter: LANGE, Clemens (Paul Scherrer Institute (CH))

Session Classification: Parallel (Track 7)

Track Classification: Track 7 - Computing Infrastructure