

CernVM-FS powered container hub

Wednesday 19 May 2021 17:53 (13 minutes)

Containers became the de-facto standard to package and distribute modern applications and their dependencies. The HEP community demonstrates an increasing interest in such technology, with scientists encapsulating their analysis workflow and code inside a container image. The analysis is first validated on a small dataset and minimal hardware resources to then run at scale on the massive computing capacity provided by the grid. The typical approach for distributing containers consists of pulling their image from a remote registry and extracting it on the node where the container runtime (e.g., Docker, Singularity) runs. This approach, however, does not easily scale to large images and thousands of nodes. CVMFS has long been used for the efficient distribution of software directory trees at global scale. In order to extend its optimized caching and network utilization to the distribution of containers, CVMFS recently implemented a dedicated container image ingestion service together with container runtime integrations. CVMFS ingestion is based on per-file deduplication, instead of the per-layer deduplication adopted by traditional container registries. On the client-side, CVMFS implements on-demand fetching of the chunks required for the execution of the container instead of the whole image.

Primary authors: BOCCHI, Enrico (CERN); BLOMER, Jakob (CERN); MOSCIATTI, Simone (CERN); VALENZUELA RAMIREZ, Andrea (Universitat Oberta de Catalunya (ES))

Presenter: BOCCHI, Enrico (CERN)

Session Classification: Storage

Track Classification: Distributed Computing, Data Management and Facilities