



Contribution ID: 242

Type: Oral

Deployment of containers on the diverse ATLAS infrastructure

Tuesday, November 5, 2019 4:45 PM (15 minutes)

We will describe the deployment of containers on the ATLAS infrastructure. There are several ways to run containers: as part of the batch system infrastructure, as part of the pilot, or called directly. ATLAS is exploiting them depending on which facility its jobs are sent to. Containers have been a vital part of the HPC infrastructure for the past year, and using fat images - images containing several releases and other data - created using the `cvmfs_shrinkwrap` utility has allowed ATLAS to run production jobs as part of the HPC infrastructure at places like NeRSC and produce 100s millions events. At other sites using non-redhat-based linux distributions, ATLAS could run thanks to containers embedded in the batch system infrastructure and using the ADC images in CVMFS, which are also used at BOINC sites. To run more extensively at all grid sites, we have devised and integrated in the new ATLAS pilot two ways to deploy containers: one wrapping them in ALRB (AtlasLocalRootBase), which is how ATLAS set up the environment, and the other is using standalone containers that can run code independently from the environment they are executed in. Both methods respond to different use cases; the first one will be completely transparent to the user and production teams, while the second will put the user in a position to choose the software to run and how to run it. The two methods also meet the requirement of running on a diverse range of sites from standard grid sites, to cloud sites, to HPC sites with no network and different architectures; in particular, it opened the possibility to use hardware accelerated workflows at grid sites and in the future at HPC sites. Access to the images for both methods is handled differently since ALRB containerization relies on images being distributed via CVMFS, while standalone containers can be more integrated in a CI workflow, which needs a faster turnaround of the images on registries such as docker and gitlab during development, or can be used as a software distribution method for networkless sites. Methods to robustly transfer and manage a large number of popular user images from the registries to CVMFS are under evaluation. We will describe different setups required at diverse sites to achieve payload isolation from the pilot environment required to comply with the WLCG traceability and isolation policy. To be able to adapt to the evolving containers ecosystem and the sites requests, we will also describe how the flexibility to use different runtimes, other than singularity, has been incorporated in the infrastructure.

Consider for promotion

No

Primary authors: FORTI, Alessandra (University of Manchester (GB)); FILIPCIC, Andrej (Jozef Stefan Institute (SI)); HEINRICH, Lukas Alexander (CERN); DE SILVA, Asoka (TRIUMF (CA)); NILSSON, Paul (Brookhaven National Laboratory (US)); DE SALVO, Alessandro (Sapienza Universita e INFN, Roma I (IT)); BOGDANCHIKOV, Alexander (Budker Institute of Nuclear Physics (RU)); LOVE, Peter (Lancaster University (GB)); PANITKIN, Sergey (Brookhaven National Laboratory (US)); BENJAMIN, Doug (Argonne National Laboratory (US)); YANG, Wei (SLAC National Accelerator Laboratory (US))

Presenter: FORTI, Alessandra (University of Manchester (GB))

Session Classification: Track 7 – Facilities, Clouds and Containers

Track Classification: Track 7 – Facilities, Clouds and Containers