



Contribution ID: 217

Type: **Poster**

Standalone containers with ATLAS offline software

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

This talk describes the deployment of ATLAS offline software in containers for use in production workflows such as simulation and reconstruction. For this purpose we are using Docker and Singularity, which are both lightweight virtualization technologies that can encapsulate software packages inside complete file systems. The deployment of offline releases via containers removes the interdependence between the runtime environment needed for job execution and the configuration of the computing nodes at the sites. Docker or Singularity would provide a uniform runtime environment for the grid, HPCs and variety of opportunistic resources. Additionally, releases may be supplemented with a detector's conditions data, thus removing the need for network connectivity at computing nodes, which is normally quite restricted at HPCs.

In preparation to achieve this goal, we have built Docker and Singularity images containing single full releases of ATLAS software for running event simulation jobs in runtime environments without network connection. These images have been successfully tested at the Theta supercomputer (ALCF) and at MareNostrum (BSC). Unlike similar parallel efforts to produce containers by packing all possible dependencies of every possible workflow into heavy images (~200GB), our approach is to include only what is needed for specific workflows and to manage dependencies efficiently via software package managers. This leads to more stable packaged releases where the dependencies are clear and the resulting images have more portable sizes (~16GB). In an effort to include a wider variety of workflows, we are deploying images that can be used in raw data reconstruction. This is particularly challenging due to the high database resource consumption during the access to the experiment's conditions payload. We describe here a prototype pipeline in which images are provisioned only with the conditions payload necessary to satisfy the jobs' requirements. This database-on-demand approach would keep images slim, portable and capable of supporting various workflows in a standalone fashion in environments with no network connectivity.

Consider for promotion

No

Primary authors: VOGEL, Marcelo (Bergische Universitaet Wuppertal (DE)); FORTI, Alessandra (University of Manchester (GB)); HEINRICH, Lukas Alexander (CERN); BORODIN, Misha (University of Iowa (US))

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

Session Classification: Posters

Track Classification: Track 7 – Facilities, Clouds and Containers