

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 461

Type: **oral presentation**

Lightweight scheduling of elastic analysis containers in a competitive cloud environment: a Docked Analysis Facility for ALICE

Thursday, 16 April 2015 09:30 (15 minutes)

During the last years, several Grid computing centers chose virtualization as a better way to manage diverse use cases with self-consistent environments on the same bare infrastructure. The maturity of control interfaces (such as OpenNebula and OpenStack) opened the possibility to easily change the amount of resources assigned to each use case by simply turning on and off virtual machines. Some of those private clouds use, in production, copies of the Virtual Analysis Facility, a fully virtualized and self-contained batch analysis cluster capable of expanding and shrinking automatically upon need: however, resources starvation occurs frequently as expansion has to compete with other virtual machines running long-living batch jobs. Such batch nodes cannot relinquish their resources in a timely fashion: the more jobs they run, the longer it takes to drain them and shut off, and making one-job virtual machines introduces a non-negligible virtualization overhead. By improving several components of the Virtual Analysis Facility we have realized an experimental “Docked” Analysis Facility for ALICE, which leverages containers instead of virtual machines for providing performance and security isolation. We will present the techniques we have used to address practical problems, such as software provisioning through CVMFS, as well as our considerations on the maturity of containers for High Performance Computing. As the abstraction layer is thinner, our Docked Analysis Facilities may feature a more fine-grained sizing, down to single-job node containers: we will show how this approach will positively impact automatic cluster resizing by deploying lightweight pilot containers instead of replacing central queue polls.

Primary author: BERZANO, Dario (CERN)

Co-authors: GANIS, Gerardo (CERN); CHARALAMPIDIS, Ioannis (CERN); BLOMER, Jakob (CERN); BUNCIC, Predrag (CERN); MEUSEL, Rene (CERN)

Presenter: BERZANO, Dario (CERN)

Session Classification: Track 7 Session

Track Classification: Track7: Clouds and virtualization