



Contribution ID: 23

Type: **not specified**

ATLAS Cloud R&D, Kubernetes and CVMFS

Tuesday 13 September 2022 16:00 (20 minutes)

Different groups, sites and experiments in the WLCG community have started using Kubernetes to manage services, implement novel analysis facilities or run batch services. Despite being in a native containerised environment, many of these use cases depend on CVMFS to stay compatible with the existing Grid model or to benefit from a well established software distribution model. Multiple of these groups therefore implemented their own helm charts and images to install the CVMFS client on Kubernetes.

In the case of the ATLAS Cloud R&D we use Kubernetes native batch capabilities to integrate multiple research and commercial clouds with the PanDA ecosystem. The clusters can be quickly scaled up and down, and the nodes are usually fully exploited. The CVMFS client needs to run like a clock to avoid expensive job failures due to hanging clients. We spent a significant amount of time comparing and customising some of the existing CVMFS clients until we reached a satisfactory situation.

This contribution will give an overview of some CVMFS clients, describe the most common issues and propose where the expertise of the CVMFS team and an official Kubernetes client would make a difference for the WLCG community.

Presenter: BARREIRO MEGINO, Fernando Harald (University of Texas at Arlington)