

Lightweight integration of Kubernetes clusters for ATLAS batch processing

Tuesday 1 December 2020 10:10 (20 minutes)

The PanDA team has evaluated the possibility of native Kubernetes job submission in order to process ATLAS workloads and offer the possibility of immediate integration of major cloud computing providers. This model also offers a novel way to set up lightweight compute sites, without the need of setting up a Grid stack.

During the last year we have been running several queues at clusters setup by institutes associated to ATLAS (ASGC, CERN, University of Chicago, University of Victoria) and on cloud providers (Amazon and Google), and have focused on increasing the stability and efficiency.

This contribution will discuss the advantages and challenges we have faced during our experience and also briefly introduce ongoing work to integrate less trivial (non pleasantly parallel) workloads.

Authors: BARREIRO MEGINO, Fernando Harald (University of Texas at Arlington); LIN, Fa-Hui (University of Texas at Arlington (US))

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

Session Classification: block 1 - presentations