



Contribution ID: 147

Type: **Sectional**

Optimizing new components of PanDA for ATLAS production on HPC resources

Thursday 28 September 2017 12:15 (15 minutes)

The Production and Distributed Analysis system (PanDA) has been used for workload management in the ATLAS Experiment for over a decade. It uses pilots to retrieve jobs from the PanDA server and execute them on worker nodes. While PanDA has been mostly used on Worldwide LHC Computing Grid (WLCG) resources for production operations, R&D work has been ongoing on cloud and HPC resources for many years. These efforts have led to the significant usage of large scale HPC resources in the past couple of years. In this talk we will describe the changes to the pilot which enabled the use of HPC sites by PanDA, specifically the Titan supercomputer at Oakridge National Laboratory. Furthermore, it was decided in 2016 to start a fresh redesign of the Pilot with a more modern approach to better serve present and future needs from ATLAS and other collaborations that are interested in using the PanDA System. Another new project for development of a resource oriented service, PanDA Harvester, was also launched in 2016. The main goal of the Harvester is flexible distribution of payloads for opportunistic resources like HPC and clouds. Both applications are now in full development after a year of studying use cases, trying different designs and deciding on the shared components model. This talk will give an overview of the evolution of the HPC pilot into Pilot 2 and Harvester projects for better utilization of HPC resources.

Author: Mr OLEYNIK, Danila (JINR LIT)

Presenter: Mr OLEYNIK, Danila (JINR LIT)

Session Classification: Distributed Computing. GRID & Cloud computing

Track Classification: Distributed Computing. GRID & Cloud Computing