

BigPanDA Workflow Management on Titan for HENP and extreme scale applications

Monday 9 July 2018 11:45 (15 minutes)

The Titan supercomputer at Oak Ridge National Laboratory prioritizes the scheduling of large leadership class jobs, but even when the supercomputer is fully loaded and large jobs are standing in the queue to run, 10 percent of the machine remains available for a mix of smaller jobs, essentially ‘filling in the cracks’ between the very large jobs. Such utilisation of the computer resources is called “backfill”.

Smaller scientific groups and data science experiments are able to conduct their computations on Titan in backfill mode. For example, simulations in high-energy physics don’t require large MPI scale jobs on a supercomputer. As a universal jobs scheduler to optimize Titan utilization it was successfully used to schedule and to submit ATLAS MC simulation jobs earlier. The current R&D project is aimed to enabling the convergence of HPC and HTC paradigms for a range of application types and communities in particle physics and beyond.

In March 2017, we implemented a new PanDA server instance within ORNL operating under Red Hat OpenShift Origin - a powerful container cluster management and orchestration system in order to serve various experiments at Titan supercomputer. We implemented a set of demonstrations serving diverse scientific workflows including LQCD and IceCube, biology studies of the genes and human brain, and molecular dynamics studies.

Authors: Dr MASHINISTOV, Ruslan (University of Texas at Arlington (US)); DE, Kaushik (University of Texas at Arlington (US)); KINCL, Jason C. (Oak Ridge National Laboratory (US)); KLIMENTOV, Alexei (Brookhaven National Laboratory (US)); SVIRIN, Pavlo (Brookhaven National Laboratory (US)); Dr WELLS, Jack C. (Oak Ridge National Laboratory (US))

Presenter: KLIMENTOV, Alexei (Brookhaven National Laboratory (US))

Session Classification: T3 - Distributed computing

Track Classification: Track 3 –Distributed computing