

BigPanDA Experience on Titan for the ATLAS Experiment at the LHC

Sean Wilkinson, Alexei Klimentov, Kaushik De, Jack Wells, Sergey Panitkin, Danila Oleynik

University of Texas - Arlington

sean.wilkinson@uta.edu

Abstract

The PanDA software is used for workload management on distributed grid resources by the ATLAS experiment at the LHC. The BigPanDA effort was launched to extend PanDA to access HPC resources, funded by the US Department of Energy (DOE-ASCR). Through this successful effort, ATLAS now uses more than 25 million hours monthly on the Titan supercomputer at Oak Ridge National Laboratory as part of its production workflow. Many challenges were met and overcome in using HPCs for ATLAS simulations. ATLAS uses two different operational modes at Titan. The traditional mode uses allocations, which require software innovations to fit the low latency requirements of experimental science. New techniques were implemented to shape large jobs using allocations on a leadership-class machine. In the second mode, work is streamed steadily to Titan to backfill high priority leadership-class jobs. This has resulted in impressive gains in overall utilization of Titan, while benefiting the physics objectives of ATLAS. For both modes, BigPanDA has integrated traditional grid computing with HPC architecture for ATLAS production workflows, and ATLAS has replicated this accomplishment to additional HPC centers around the world. This presents a summary of the innovations to use Titan successfully for LHC physics goals.