

Integration of the Titan supercomputer at OLCF with the ATLAS Production System

Thursday 13 October 2016 15:15 (15 minutes)

The PanDA (Production and Distributed Analysis) workload management system was developed to meet the scale and complexity of distributed computing for the ATLAS experiment.

PanDA managed resources are distributed worldwide, on hundreds of computing sites, with thousands of physicists accessing hundreds of Petabytes of data and the rate of data processing already exceeds Exabyte per year.

While PanDA currently uses more than 200,000 cores at well over 100 Grid sites, future LHC data taking runs will require more resources than Grid computing can possibly provide.

Additional computing and storage resources are required.

Therefore ATLAS is engaged in an ambitious program to expand the current computing model to include additional resources such as the opportunistic use of supercomputers.

In this talk we will describe a project aimed at integration of ATLAS Production System with Titan supercomputer at Oak Ridge Leadership Computing Facility (OLCF).

Current approach utilizes modified PanDA Pilot framework for job submission to Titan's batch queues and local data management, with lightweight MPI wrappers to run single node workloads in parallel on Titan's multi-core worker nodes. It provides for running of standard ATLAS production jobs on unused resources (backfill) on Titan.

The system already allowed ATLAS to collect on Titan millions of core-hours per month, execute hundreds of thousands jobs, while simultaneously improving Titan's utilization efficiency.

We will discuss details of the implementation, current experience with running the system, as well as future plans aimed at improvements in scalability and efficiency.

Secondary Keyword (Optional)

Distributed workload management

Primary Keyword (Mandatory)

High performance computing

Tertiary Keyword (Optional)

Primary author: PANITKIN, Sergey (Brookhaven National Laboratory (US))

Presenter: PANITKIN, Sergey (Brookhaven National Laboratory (US))

Session Classification: Track 7: Middleware, Monitoring and Accounting

Track Classification: Track 7: Middleware, Monitoring and Accounting