Contribution ID: **43**   Type: **Poster**

Integration of PanDA workload management system with Titan supercomputer at OLCF

Tuesday, 2 September 2014 08:00 (1 hour)

Experiments at the Large Hadron Collider (LHC) face unprecedented computing challenges.
Heterogeneous resources are distributed worldwide, thousands of physicists analyzing the data need remote access to hundreds of computing sites,the volume of processed data is beyond the exabyte scale, and data processing requires more than billions of hours of computing usage per year.
The PanDA (Production and Distributed Analysis) workload management system (WMS) was developed to meet the scale and complexity of LHC distributed computing for the ATLAS experiment.
While PanDA currently uses more than 100,000 cores at well over 100 Grid sites with a peak performance of 0.3 petaFLOPS, next LHC data taking run will require more resources than Grid computing can possibly provide.
The Worldwide LHC Computing Grid (WLCG) infrastructure will be sufficient for the planned analysis and data processing, but it will be insufficient for Monte Carlo (MC) production and any extra activities.
Additional computing and storage resources are therefore required.
To alleviate these challenges, 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 turn this activity drives evolution of the PanDA WMS.
We will describe a project aimed at integration of PanDA WMS with
Titan supercomputer at Oak Ridge Leadership Computing Facility (OLCF).
Current approach utilizes modified PanDA pilot framework for job submission to
Titan's batch queue and local data management, with light-weight MPI wrappers to run single threaded workloads in parallel on Titan's multi-core worker nodes.
It also gives PanDA new capability to collect, in real time, information about unused worker nodes on Titan, which allows
to precisely define the size and duration of jobs submitted to Titan according to available free resources.
This capability can reduce job wait time and improve Titan's utilization efficiency.
This implementation was tested with Monte Carlo simulation jobs and is suitable for deployment with many other supercomputing platforms.

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

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

Session Classification: Poster session

Track Classification: Computing Technology for Physics Research