

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 152

Type: oral presentation

Integration of PanDA workload management system with Titan supercomputer at OLCF.

Thursday, April 16, 2015 11:15 AM (15 minutes)

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.

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.

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 queues 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 precisely define the size and duration of jobs submitted to Titan according to available free resources.

This capability significantly reduces PanDA job wait time while improving Titan's utilization efficiency.

This implementation was tested with a variety of Monte-Carlo workloads on Titan and is being tested on several other supercomputing platforms.

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

Co-authors: VANIACHINE, Alexandre (ATLAS); Dr KLIMENTOV, Alexei (Brookhaven National Laboratory (US)); PETROSYAN, Artem (Joint Inst. for Nuclear Research (RU)); OLEYNIK, Danila (Joint Inst. for Nuclear Research (RU)); SCHOVANCOVA, Jaroslava (CERN); DE, Kaushik (University of Texas at Arlington (US)); Dr WENAUS, Torre (Brookhaven National Laboratory (US))

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

Session Classification: Track 8 Session

Track Classification: Track8: Performance increase and optimization exploiting hardware features