

Integration Of PanDA Workload Management System With Supercomputers for ATLAS and Data Intensive Science

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2016 J. Phys.: Conf. Ser. 762 012021

(<http://iopscience.iop.org/1742-6596/762/1/012021>)

View the [table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 128.219.49.13

This content was downloaded on 23/11/2016 at 13:03

Please note that [terms and conditions apply](#).

You may also be interested in:

[Integration of PanDA workload management system with Titan supercomputer at OLCF](#)
K. De, A. Klimentov, D. Oleynik et al.

[Non-conventional mesons at PANDA](#)
Francesco Giacosa

[The future of PanDA in ATLAS distributed computing](#)
K De, A Klimentov, T Maeno et al.

[Thermal analysis of the cold mass of the 2T solenoid for the PANDA detector at FAIR](#)
G Rolando, H H J ten Kate, A Dudarev et al.

[Integration of Globus Online with the ATLAS PanDA Workload Management System](#)
C Contreras, W Deng, T Maeno et al.

[A Study of ATLAS Grid Performance for Distributed Analysis](#)
Sergey Panitkin, Valery Fine and Torre Wenaus

[Distributed Data Analysis in the ATLAS Experiment: Challenges and Solutions](#)
Johannes Elmsheuser and Daniel van der Ster

Integration Of PanDA Workload Management System With Supercomputers for ATLAS and Data Intensive Science

A Klimentov¹, K De², S Ha³, T Maeno¹, P Nilsson¹, D Oleynik^{2,5}, S Panitkin¹, J Wells⁴ and T Wenaus¹

¹ Physics Department, Brookhaven National Laboratory, Upton, NY, US

² Physics Department, University of Texas at Arlington, Arlington, TX 76019, USA

³ Department of Electrical and Computer Engineering Rutgers University, NJ 08854-8058, USA

⁴ Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory, Oak Ridge, TN, 37831-6161

⁵Laboratory of information Technologies, Joint Institute for Nuclear Research, Dubna 141980, Russia

E-mail : alexei.klimentov@cern.ch

Abstract: *LHC, operating at CERN, is leading Big Data driven scientific explorations. Experiments at the LHC explore the fundamental nature of matter and the basic forces that shape our universe.*

ATLAS, one of the largest collaborations ever assembled in the sciences, is at the forefront of research at the LHC. To address an unprecedented multi-petabyte data processing challenge, the ATLAS experiment is relying on a heterogeneous distributed computational infrastructure. The ATLAS experiment uses PanDA (Production and Data Analysis) Workload Management System for managing the workflow for all data processing on over 150 data centers. Through PanDA, ATLAS physicists see a single computing facility that enables rapid scientific breakthroughs for the experiment, even though the data centers are physically scattered all over the world. While PanDA currently uses more than 250,000 cores with a peak performance of 0.3 petaFLOPS, LHC data taking runs require more resources than grid can possibly provide. To alleviate these challenges, LHC experiments are 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 supercomputers in United States, in particular with Titan supercomputer at Oak Ridge Leadership Computing Facility. Current approach utilizes modified PanDA pilot framework for job submission to the supercomputers batch queues and local data management, with light-weight MPI wrappers to run single threaded workloads in parallel on LCFs multi-core worker nodes. This implementation was tested with a variety of Monte-Carlo workloads on several supercomputing platforms for ALICE and ATLAS experiments and it is in full production for the ATLAS since September 2015.

We will present our current accomplishments with running PanDA at supercomputers and demonstrate our ability to use PanDA as a portal independent of the computing facilities infrastructure for High Energy and Nuclear Physics as well as other data-intensive science applications, such as bioinformatics and astro-particle physics.



Content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1. Introduction

The largest scientific instrument in the world – the Large Hadron Collider [1] – operates at the CERN Laboratory in Geneva, Switzerland. The ATLAS [2] experiment at the LHC explore the fundamental nature of matter and the basic forces that shape our universe. To address an unprecedented multi-petabyte data processing challenge, experiments are relying on the computational grids infrastructure deployed within the framework of the Worldwide LHC Computing Grid (WLCG) [3].

The 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 including the opportunistic use of supercomputers and high-performance computing clusters (HPC's).

2. PanDA Workload Management System

PanDA (Production and Distributed Analysis) is the workload management system developed in 2005 for the ATLAS experiment. Its move to production represented a paradigm shift in ATLAS computing by federating the O(100) heterogeneous, independent computing centers of the WLCG into a unique job submission system. Since then, PanDA manages both the analysis jobs from around 3000 users and the production jobs from a few groups of power-users, and brokers them to the best available resource. Furthermore, PanDA is more than a job submission system, and is capable of managing complex tasks, defining different steps and dependencies between the jobs inside a task.

PanDA is built as a central, well-tuned service that knows the status of the distributed resources and manages these intelligently by matching jobs and sites to ensure quick turnaround and full resource usage (figure 1). In addition PanDA implements global fair-shares, job priorities and policies. PanDA communicates with so called pilot jobs [4,5], which are running on the worker nodes and request the real payload once a worker node health-check has been completed.

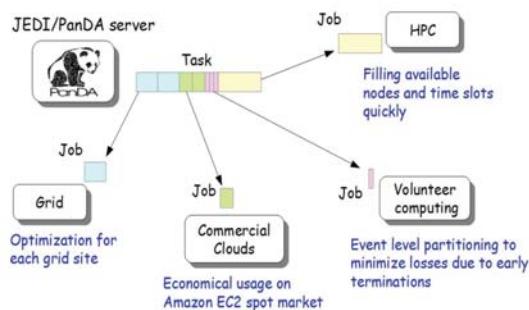


Figure 1: Dynamic job definition adapted to the computing resource

Using of pilot-based workflows helps to improve job reliability, optimize resource utilization, allows for opportunistic resources usage, and mitigates many of the problems associated with the inhomogeneities found on the Grid. Extending PanDA beyond the Grid will further expand the potential user community and the resources available to them.

3. Extending PanDA to Titan supercomputer

Titan is the most powerful supercomputer for open science (number 2 in Top 500 list) hosted at the Oak Ridge Leadership Computing Facility in Oak Ridge National Laboratory, USA [5].

Titan is a hybrid-architecture Cray XK7 system with a theoretical peak performance exceeding 27 petaflops. Titan features 18,688 compute nodes, (each with one 16-core AMD Opteron CPU and 1 NVIDIA Kepler K20X GPU), 299,008 x86 cores, a total system memory of 710 terabytes, and a high-performance proprietary network [6]. Worker nodes use Cray's Gemini interconnect for inter-node MPI messaging but have no network connection to the outside world. Titan is served by the shared Lustre filesystem that has 32 PB of disk storage and by the HPSS tape storage that has 29 PB capacity. Titan's worker nodes run Compute Node Linux which is a run time environment based on the Linux kernel derived from SUSE Linux Enterprise Server.

Taking advantage of modular and extensible design of PanDA, the PanDA pilot code and logic has been enhanced with tools and methods relevant for HPC. The pilot runs on Titan's front-end nodes which allows it to communicate with the PanDA server, since front end nodes have connectivity to the Internet. The interactive front-end machines and the worker nodes use a shared file system which makes it possible for the pilot to stage-in input data that is required by the payload and stage-out the produced output at the end of the job. The ATLAS Tier 1 computing center at Brookhaven National Lab is currently used for data transfer to and from Titan, but in principle that can be any Grid site. The pilot submits ATLAS payloads to the worker nodes using the local batch system (PBS) via the SAGA (Simple API for Grid Applications) interface [7][8]. Figure 2 shows the schematic diagram of PanDA components on Titan.

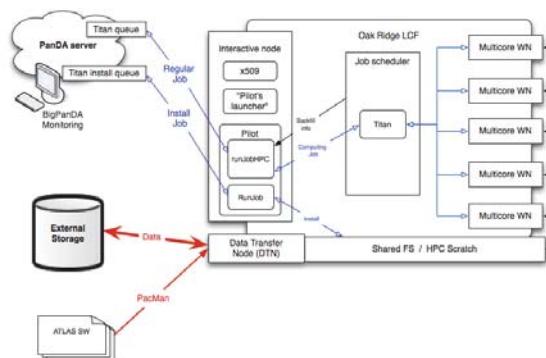


Figure 2 Titan integration with PanDA

The majority of experimental high energy physics workloads do not use Message Passing Interface (MPI). They are designed around event level parallelism and thus are executed on the Grid independently. Typically, detector simulation workloads can run on a single compute node using multiprocessing.

For running such workloads on Titan we developed an MPI wrapper that allows us to launch multiple instances of single node workloads simultaneously. MPI wrappers are typically workload specific since they are responsible for setup of workload specific environment, organization of per-rank worker directories, rank specific data management, input parameters modification when necessary and cleanup on exit. The wrapper scripts is what pilot actually submits to a batch queue to run on Titan. The pilot reserves the necessary number of worker nodes at submission time and at run time a corresponding number of copies of the wrapper script will be activated on Titan. Each copy will know its MPI rank (an index that runs from zero to a maximum number of nodes or script copies) as

well as the total number of ranks in the current submission. When activated on worker nodes each copy of the wrapper script, after completing the necessary preparations, will start the actual payload as a subprocess and will wait until its completion. In other words the MPI wrapper serves as a "container" for non-MPI workloads and allows us to efficiently run unmodified Grid-centric workloads on a parallel computational platforms, like Titan.

Leadership Computing Facilities (LCF), like Titan, are geared towards large scale jobs by design. Time allocation on an LCF machines is very competitive and large scale projects are often preferred. This is especially true for Titan at OLCF which was designed to be the most powerful machine in the world, capable of running extreme scale computational projects. As a consequence, on average, about 10% of capacity on Titan is unused due to mismatches between job sizes and available resources. The worker nodes sit idle because there are not enough of them to handle a large scale computing job. On Titan, these 10% correspond to estimated 300M core hours per year. Hence, a system that can occupy those temporarily free nodes would be very valuable. It would allow delivery of more CPU cycles for scientific research, while simultaneously improving resource utilization efficiency on Titan. This offers a great possibility for PanDA to harvest these opportunistic resources on Titan.

During the early steps of this project aimed to integration of Titan with PanDA, it was demonstrated that PanDA is capable of improving the overall utilization of Titan by running in backfill mode. In this mode, PanDA submits and executes jobs that fit within unused cycles between large leadership class jobs. New functionality was added to the PanDA Pilot to collect information in real time about available unused worker nodes on Titan. This allows PanDA to predict precisely the size and duration of jobs submitted to Titan based on available free resources. Dynamic occupation of backfill resources of Titan shown on Figure 3.

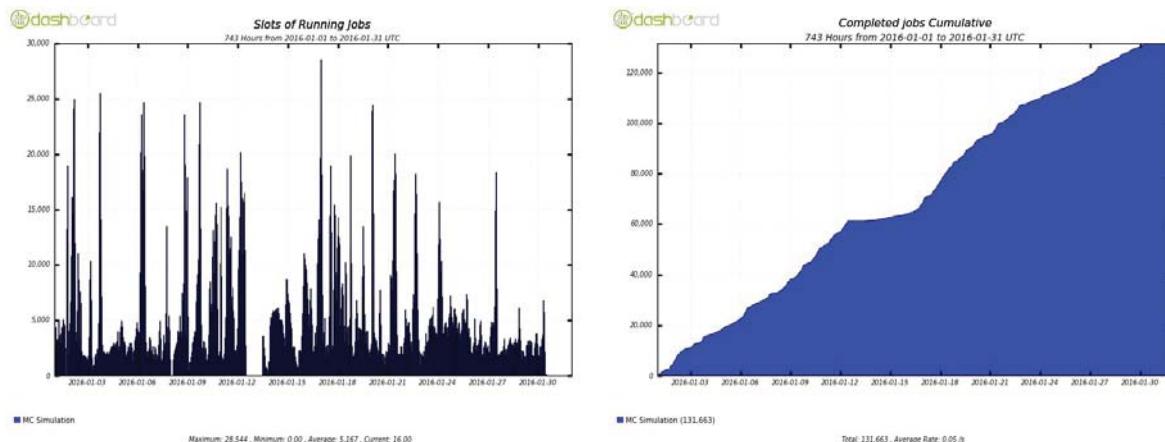


Figure 3. Job slots (CPU cores) handled by PanDA on Titan during January 2016

Figure 4. Number of completed jobs on Titan in January 2016

Titan was fully integrated with the ATLAS PanDA based Production and Analysis system and now the ATLAS experiment routinely runs Monte-Carlo simulation tasks there. All operations, including data transfers to and from Titan, are transparent to the ATLAS Computing Operations team and physicists.

4. Conclusion

PanDA WMS was successfully integrated with one of most powerful supercomputer in the world, and this enables Titan to process data from the ATLAS experiment as a valuable computing resource.

Just in one month in January 2016, more than 130,000 PanDA jobs were completed on Titan (Figure 4)

One PAnDA pilot was able to manage up to 300 simultaneously launched PanDA jobs filling up to 4800 cores. With seven simultaneously launched pilots this solution allows to consume >2M CPU/hours per month of backfill resources on Titan.

Initial results approve that advised architecture works, but also shows which parts should be improved for increasing of efficiency:

- Improving of warming-up procedures, environment setup and initial checks before launching of payload may take significant time for big chunks of jobs;
- Data Management procedures, especially Stage-Out.
- IO operations as in Pilot, so in payloads
- Communication with PanDA server, to update information about jobs in bulk.

References

- [1] Maeno T et al 2007 PanDA: Distributed Production and Distributed Analysis System for ATLAS *International Conference on Computing in High Energy and Nuclear Physics* (Victoria, Canada)
- [2] ATLAS Collaboration 2008 *J. Inst.* **3** S08003.
- [3] 2005 LHC Computing Grid: Technical Design Report. document LCG-TDR-001, *CERN-LHCC-2005-024 (The LCG TDR Editorial Board)* ISBN 92-9083-253-3
- [4] Nilsson P et al 2007 Experience from a pilot based system for ATLAS *International Conference on Computing in High Energy and Nuclear Physics* (Victoria, Canada)
- [5] Top500 List <http://www.top500.org/>
- [6] Titan at OLCF web page <https://www.olcf.ornl.gov/titan/>
- [7] The SAGA Framework web site <http://saga-project.github.io>
- [8] Merzky A Weidner O Jha S 2015 SAGA: A Standardized Access Layer to Heterogeneous Distributed Computing Infrastructure *Software-X* DOI: 10.1016/j.softx.2015.03.001

Acknowledgements

This project was funded in part by U. S. DOE, Office of Science, High Energy Physics and Advanced Scientific Computing Research (ASCR) under Contract No. DE-AC02-98CH10886. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.