Present and future of PanDA WMS integration with Titan supercomputer at OLCF

A. Angius, F. Barreiro Megino, K. De, S. Jha, A. Klimentov, T. Maeno, R. Mashinistov,
P. Nilsson, D. Oleynik, S. Oral, S. Padolski, S. Panitkin, P. Svirin, M. Turilli, J. Wells, T. Wenaus

BNL, ORNL, Rutgers, UTA

Abstract

Modern physics experiments collect peta-scale volumes of data and utilize vast, geographically distributed computing infrastructure that serves thousands of scientists around the world.
Requirements for rapid, near real time data processing, fast analysis cycles and need to run massive detector simulations to support data analysis pose special premium on efficient use of available computational resources. A sophisticated Workload Management System (WMS) is needed to coordinate the distribution and processing of data and jobs in such environment. We will discuss PanDA WMS developed by the ATLAS experiment at the LHC. Even though PanDA was originally designed for workload management in Grid environment, it was successfully extended to include cloud resources and supercomputers. In particular we'll described current state of PanDA integration with Titan supercomputer at Oak Ridge Leadership Computing Facility (OLCF). Our approach utilizes a modified PanDA pilot framework for job submission to Titan's batch queues and for data transfers to and from OLCF . The system employs lightweight MPI wrappers to run in parallel multiple, independent, single node payloads on Titan's multi-core worker nodes. It also gives PanDA a 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. The initial implementation of this system already allowed to collect in 2016 more than 70M core hours of otherwise left unused resources on Titan and execute tens of millions of PanDA jobs. Based on experience gained on Titan the PanDA development team is exploring designs of next generation components and services for workload management on HPC, Cloud and Grid resources.

PanDA integration with Titan

PanDA WMS was developed for the ATLAS Experiment at LHC for job scheduling on the distributed computational infrastructure. The system has been designed to meet ATLAS production and analysis requirements for a data-driven workload management system capable of operating at LHC data processing scale. Currently, as of 2017, PanDA WMS manages processing of over one million jobs per day, serving thousands of ATLAS users worldwide. It is capable of executing jobs on heterogeneous distributed resources which include WLCG, supercomputers, and public and private clouds. As part of the DOE ASCR supported BigPanDA project aimed at extension of PanDA beyond ATLAS Grid PanDA WMS was integrated with Titan supercomputer.

The Titan supercomputer is located at the Oak Ridge Leadership Computing Facility in the Oak Ridge National Laboratory, USA. It has theoretical peak performance of 27 PetaFLOPS. Titan has 18,688 worker nodes with GPUs and a total of 299,008 CPU cores.

Current approach to integration of PanDA with Titan 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. In this setup PanDA pilot effectively plays a role not of a pilot but of a Broker. Data delivery to and from Titan was fully automated and synchronized with PanDA jobs. Figure 1 shows a schematic view of PanDA interface to Titan.

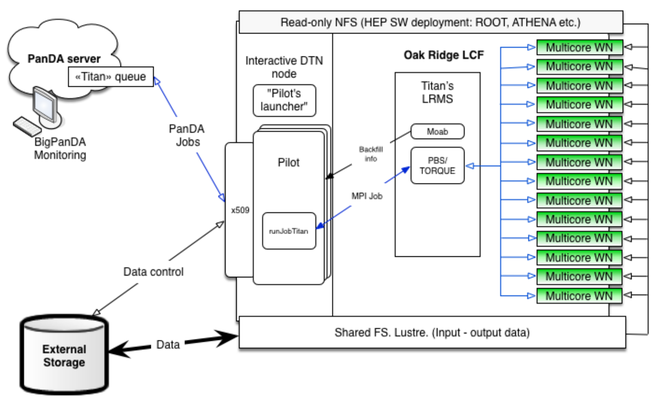

Fig 1: Schematic view of PanDA interface to Titan

This setup was integrated with ATLAS production system and used to ran large scale simulations campagns, using only backfill . Figure 2 shows Titan core hours used by ATLAS detector simulation jobs from July 2016 to July 2017. During that period the system collected ~146M Titan core hours, processed ~4.2M jobs, with ~309M events. Peak resource consumption reached ~20M core hours per month.

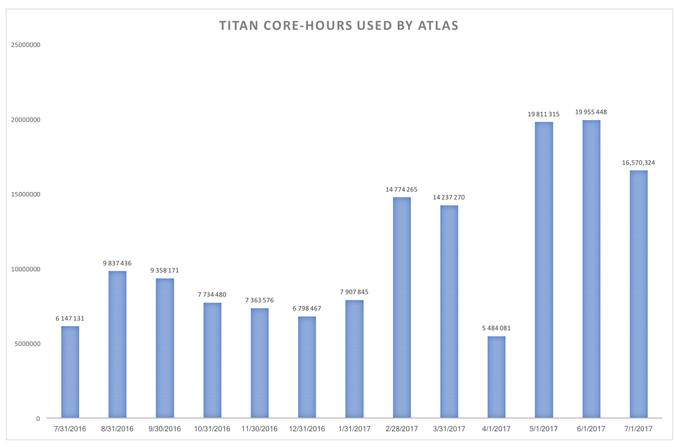

Fig. 2: Titan core hours used by ATLAS detector simulation jobs from July 2016 to July 2017

Based on experience on running large scale simulations on Titan and other HPCs the BigPanDA team is developing next generation set of tools and services.

Harvester

Harvester is a resource-facing service between PanDA server and collection of workers. Worker is a generalization of a pilot concept and can be, depending on a resource and workflow, a pilot, an MPI job, or a virtual machine. Harvester has modular multithreaded design to support heterogeneous resources and to accommodate for special workflows requirements. Harvester will provides for flexible scheduling of job execution and asynchronous data transfer to and from the controlled resource.

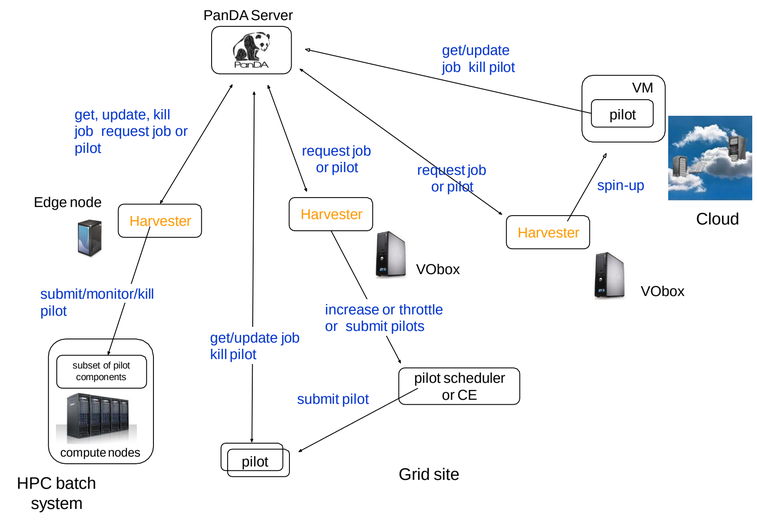

Fig. 3: Schematic view of Harvester integration with PanDA managed resources

Next Generation Executor (NGE)

Next Generation Executor (NGE) is a prototype of a pilot system capable of executing on Titan and other HPC mashines. NGE is a runtime system to execute heterogeneous and dynamically determined workloads. Fig. 3 illustrates its current architecture as deployed on Titan: the two management modules (Pilot and Unit) represent a simplified version of the PanDA Broker while the agent module is the pilot submitted to Titan and executed on its worker nodes.

NGE exposes an API to describe workloads and pilots, and to instantiate a PilotManager and a UnitManager. The PilotManager submits pilots to Titan's PBS batch system via SAGA API. The NGE Agent uses the Open Run-Time Environment (ORTE) for communication and coordination of the execution of units. NGE allows to run multiple generations of heterogeneous workloads during agent's time allocation which helps to increase resource utilization efficiency.

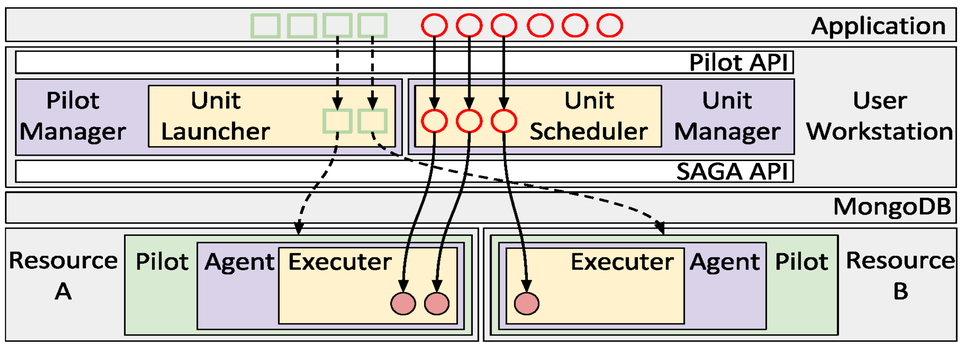

Figure 4: NGE Architecture: The PilotManager and UnitManager reside on a Titan data transfer nodes while the Agent is executed on its worker nodes. Color coding: gray for entities external to NGE, white for APIs, purple for NGE's modules, green for pilots, yellow for module's components.

We ran a set of tests with detector simulations workload in order to determine scalability of NGE on Titan. Both weak and strong scalability were tested. As an example Fig. 5 shows the results of the strong scalability tests where pilots with less than 2048 nodes ran multiple generations of workloads.

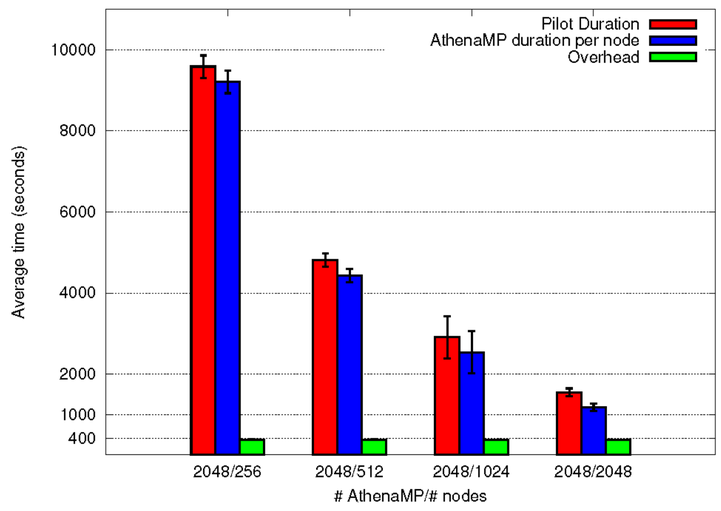

Figure 5: Average pilot duration, average duration of sequential jobs, and pilot's overhead for Pilots with 256, 512, 1024 and 2018 nodes

The next generation services for PanDA are currently under development and undergoing testing. They aim to improve resource utilization and operational efficiency and will be ready soon for the production on Titan.