Harvesting HPC Resources on Titan with BigPanDA

S. Wilkinson, P. Svirin, J. Kincl, R. Mashinistov, A. Merzky, D. Oleynik, S. Oral, S. Panitkin, M. Turilli,
K. De, S. Jha, A. Klimentov, J. Wells, T. Wenaus

Overview

- Our project, BigPanDA, has adapted the PanDA workload management system from the ATLAS collaboration to run on the Titan supercomputer.
- PanDA manages High Throughput Computing (HTC) workloads on the Worldwide LHC Compute Grid (WLCG), and BigPanDA is the first project to manage Big Data workloads from High Energy Physics (HEP) by using High Performance Computing (HPC) resources.
- BigPanDA contributed almost 500 million Titan core hours to ATLAS physics results between January 2016 and September 2018, thanks to 128 million hours from an ALCC allocation and 363 million hours from harvesting unused resources on Titan via “backfill mode”.
- Ongoing work continues to generalize and adapt PanDA for use by other projects and workloads in HEP, Genomics, and Molecular Dynamics.

Harvester

- Harvester is a resource-facing service between the PanDA server and the HPC batch system for optimal resource provisioning and workload shaping.
- It is a lightweight stateless Python service running on an edge node of an HPC center to provide a uniform interface for HPC resources.

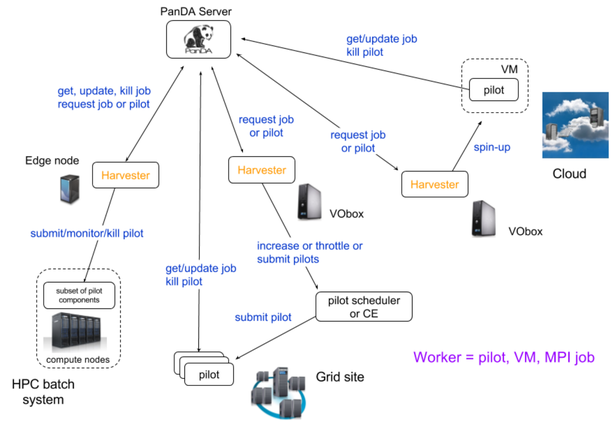

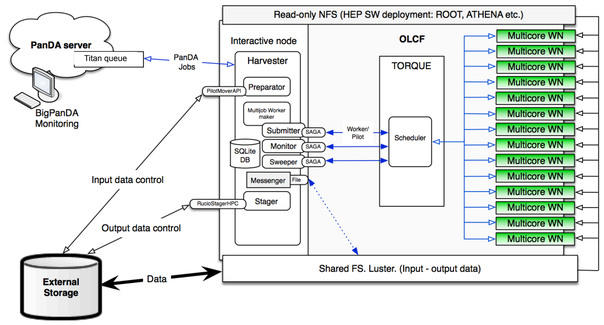

- In February 2018, the Harvester instance at OLCF went into production for the HEP110 project (ALCC for ATLAS). It is deployed on a Data Transfer Node with its own set of plugins, SQLite for storage, and WAN access to the PanDA server (not LAN).
- Work is underway to use Harvester to manage and reshape workloads for applications beyond ATLAS, such as Molecular Dynamics.

Performance and Metrics

- Titan was the first large-scale HPC integrated into ATLAS production, and it is currently used exclusively for simulation.
- From January 1, 2018 through September 1, ATLAS used 7.6% of all core hours consumed on Titan, and this represented 5.4% of the total computing consumed by ATLAS and 12% of its consumption for simulations.

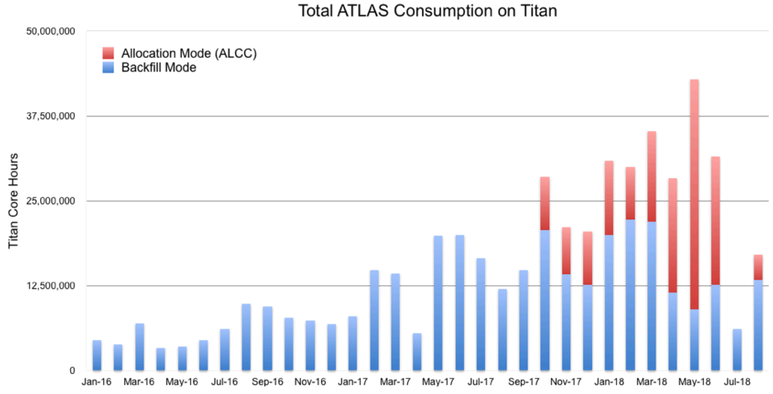

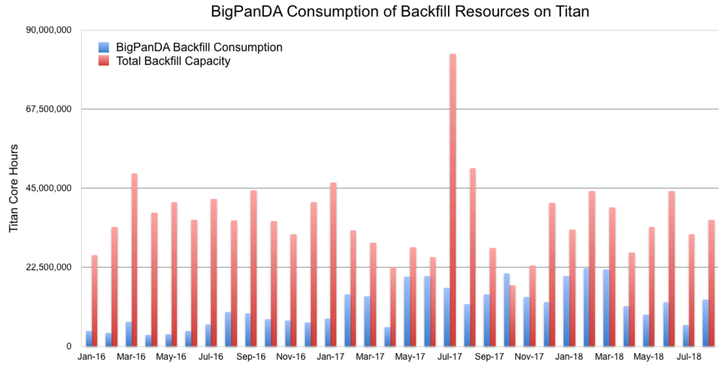

- BigPanDA’s ability to use backfill resources efficiently also increases Titan’s overall utilization. From January 2016 to September 2018, ATLAS consumed 363 million core hours via backfill mode.
- Assuming that BigPanDA’s use of backfill mode consumes only resources which would have gone unused, ATLAS is responsible for a ~3% increase in overall utilization on Titan.

- BigPanDA is granted an exception to run more concurrent jobs than OLCF policy normally allows, but these jobs run under the lowest possible priority.
- To quantify the effects of BigPanDA on the overall utilization of Titan, work has begun to define and compute metrics from data sampled from the MOAB scheduler.