HPC and its Role in CERN’S ATLAS Experiment

Wesley Kwiecinski

Overview

- Large Hadron Collider
- ATLAS Experiment
- HPC in ATLAS
  - Considerations of Using HPC With ATLAS
  - PanDA
  - Harvester
  - ATLAS Workflows with HPC
- Conclusion

Large Hadron Collider (LHC)

- Largest particle accelerator in the world
- 27km tunnel located at CERN in Switzerland
- Accelerates subatomic particles, creates proton-proton and heavy ion-heavy ion collisions, when active
- Supports a number of different projects at CERN

Large Hadron Collider (LHC)

Aerial photo of the LHC, from CERN

Large Hadron Collider (LHC)

Aerial photo of the LHC, from CERN

ATLAS Experiment

- An experiment designed to use the full discovery potential of the LHC
- Aims to gain insight into building blocks of matter, dark matter & fundamental forces, will continue for decades
- Focuses on proton-proton (p-p) collisions from the LHC
- 1 billion p-p collisions per second
- Filtering of event data results in 1000 events (approx. 300 MB) per second

ATLAS Detector

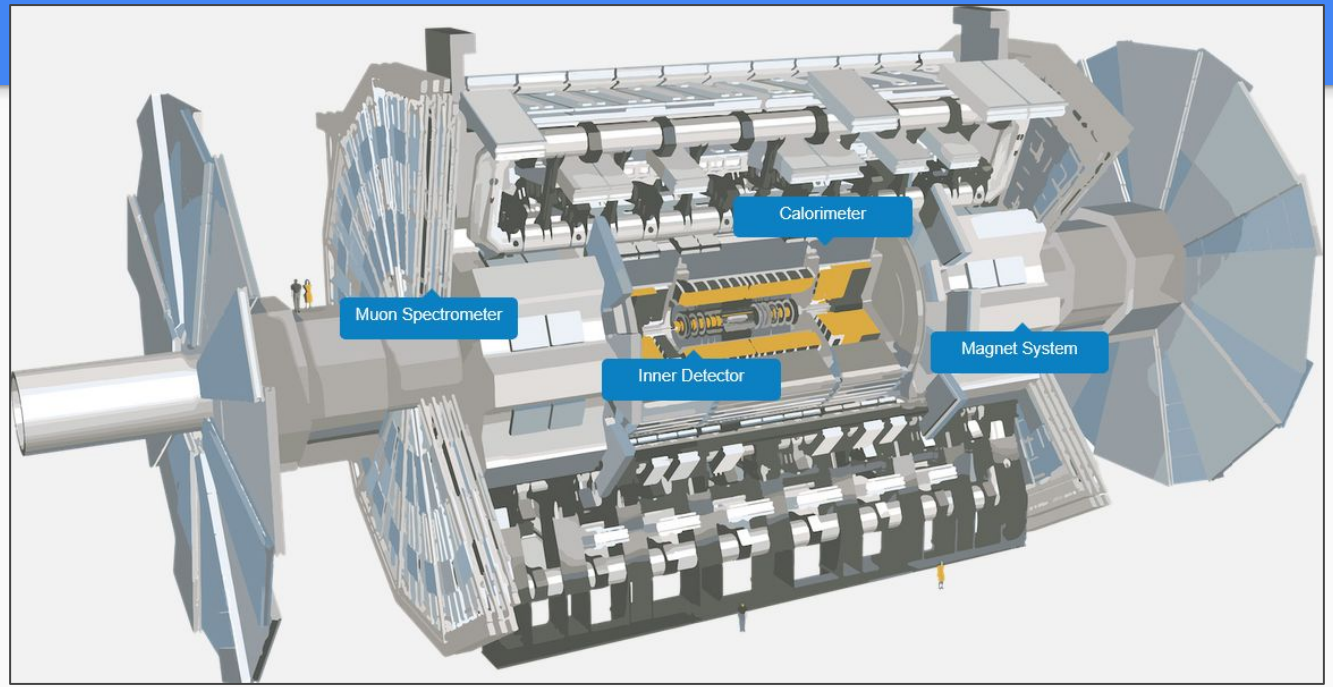

The ATLAS Detector, from CERN

Considerations of Using HPC Systems in ATLAS

- HPCs are able to process data that has overflowed from WLCG computing resources
- Many HPCs are heterogeneous systems, so software can achieve more performance through the use of accelerators (e.g. GPUs or FPGAs)
    - Though to take advantage of heterogeneous architectures, software may need to be modified
- There are many different HPC systems all owned by different vendors.
    - All have their own unique access steps, job schedulers, hardware, and software stacks
- ATLAS Software often does not use MPI, makes it harder to use multi-node systems
- HPCs are designed to process large jobs. If not enough work is scheduled, nodes are wasting time by sitting idle

ATLAS Distributed Computing System

- ATLAS's distributed computing system acts as a layer between the user and the computing resources
- Interacts with every computing resource used by ATLAS
- 2 major subsystems used when considering HPCs, PanDA, Harvester
- We'll see examples of HPCs that use PanDA and Harvester to process data from ATLAS

ATLAS Distributed Computing System

- Hierarchy of ATLAS's distributed computing system
- PanDA interacts with HPCs through Harvester
- Harvester interacts directly with HPCs

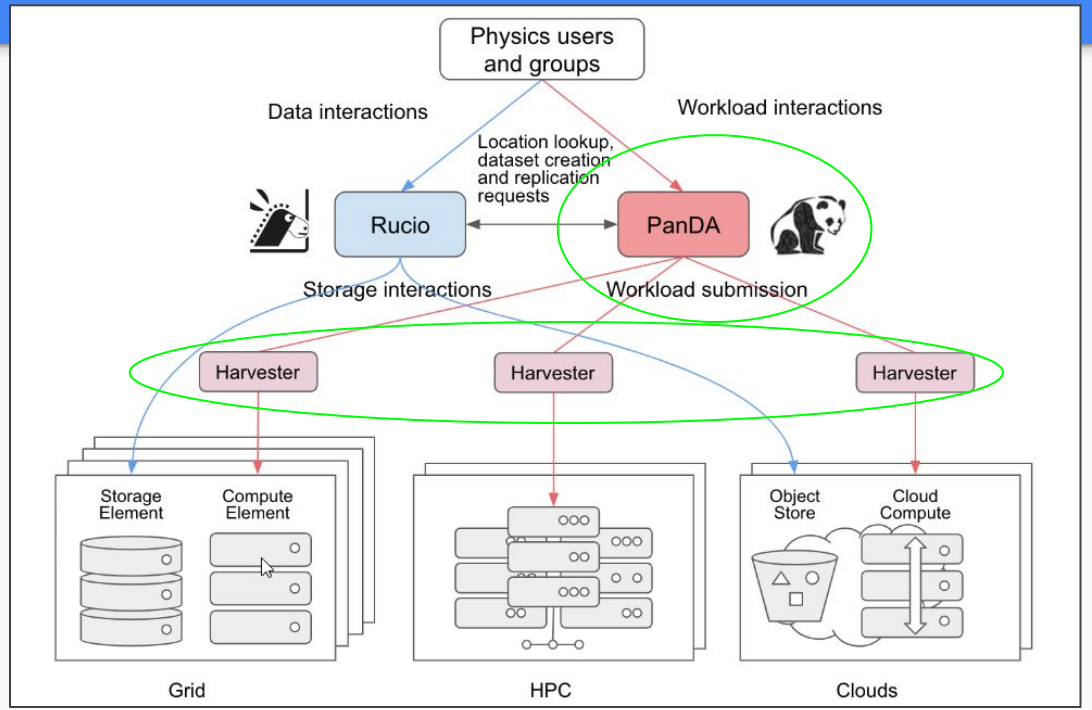

Overview of the ATLAS Distributed Computing System, from J. Elmsheuser.

Production and Distributed Analysis System (PanDA)

- Workload management system, supports massively parallel jobs
- Abstracts complexity of using computing resources
- Pilot jobs are ran on worker nodes and receive end-user jobs from a central DB
  - Pilot jobs organize workflow processing
  - Pilots improve resource utilization & job reliability

PanDA With Titan at OLCF

- ATLAS & OLCF Collaboration integrated PanDA onto the Titan supercomputer in 2015
- To take full advantage of HPC hardware:
  - Pilots use MPI wrappers to send jobs to Titan
  - PanDA interacts with Titan's scheduler to improve underused resources
  - ATLAS software was ported to Titan, focusing on physics analysis that could not be done well on WLCG (Event generation, simulations)
- Titan ran over 100,000 MC simulations in a few months
- Titan was ultimately decommissioned in 2019

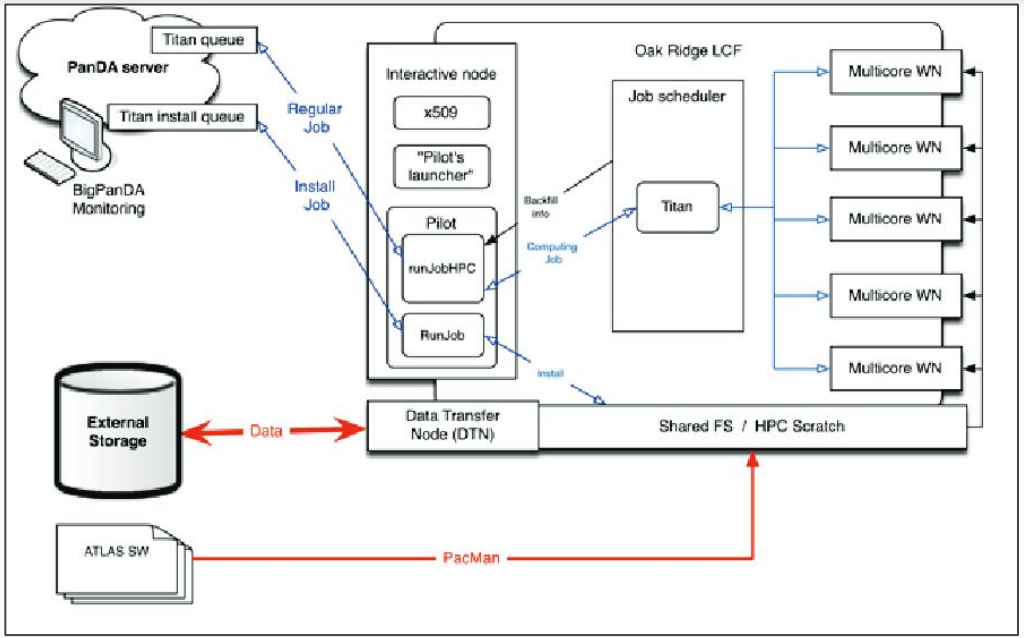

PanDA integration with Titan schema, K. De et al.

Harvester

- Another abstract layer to interfacing with HPC systems
- Designed to improve resource usage without manual intervention
- Uses databases to track resource and job use, data is communicated back to PanDA to manager Harvester instances
- Harvester instances on HPC nodes will communicate directly with the batch scheduler on the system

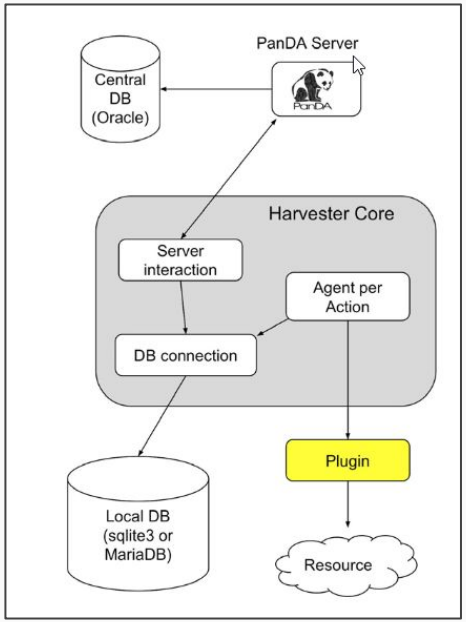

Harvester schema, T. Maeno et al.

Harvester on HPCs

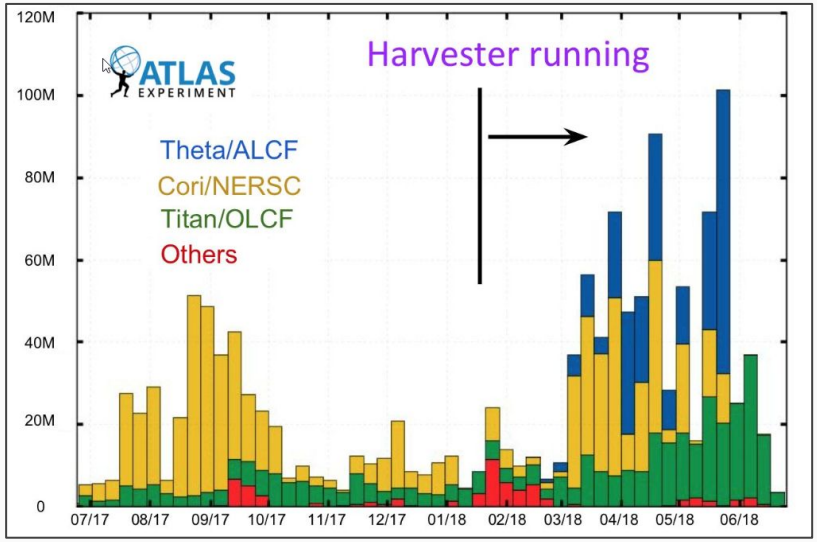

Events generated per week with Harvester, T. Maeno et al.

- Harvester was introduced to various HPC systems in 2018
- Most systems saw a large increase in events generated per week with the introduction of Harvester
- Cori supercomputer at NERSC saw little benefit since it already had good resource utilization

Harvester on HPCs

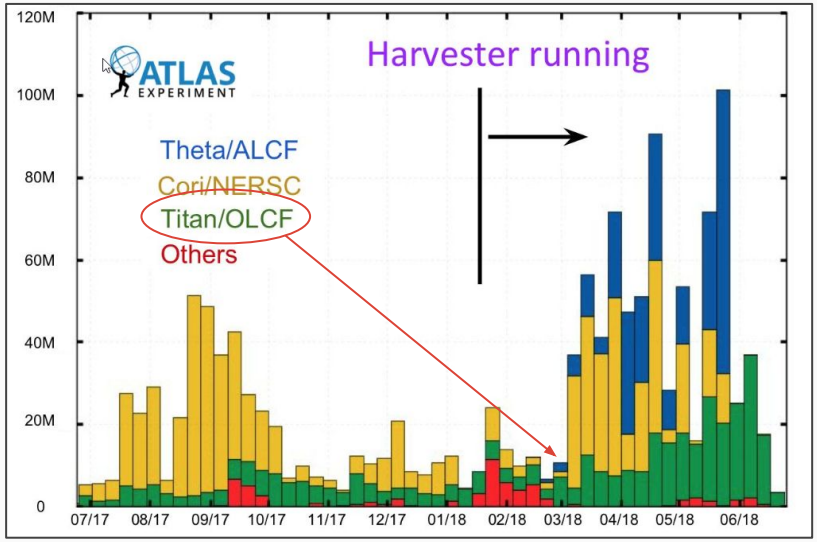

Events generated per week with Harvester, T. Maeno et al.

- Harvester was introduced to various HPC systems in 2018
- Most systems saw a large increase in events generated per week with the introduction of Harvester
- Cori supercomputer at NERSC saw little benefit since it already had good resource utilization

ATLAS Workflows with HPC

- Often HPCs do event generation & Monte Carlo (MC) simulations
  - Event Generation - generation of particle collision event data from the detector, 1-5000 CPUs/event
  - MC Simulation - mocks particle collisions from the detector data, 200-1000 CPUs/event
- We saw Titan running MC Simulations with PanDA
- Various HPC systems ran event generation jobs with Harvester
- ATLAS collaborators also ran Monte Carlo simulations at the Swiss National Supercomputing Center (CSCS), ported ATLAS software code to run on GPUs
  - Used the Piz Daint Cray HPC
  - GPU code saw minor performance benefit, shows work can be done to optimize ATLAS software for specific HPC systems

ATLAS Workflows with HPC

- Other ATLAS workflows with HPC:
  - Reconstruction - Reconstruction of data from trigger in ATLAS
  - Derivations - Copy and write to data files from reconstruction
  - Analysis - user analysis of physics data
    - Some newer analysis workflows use machine learning

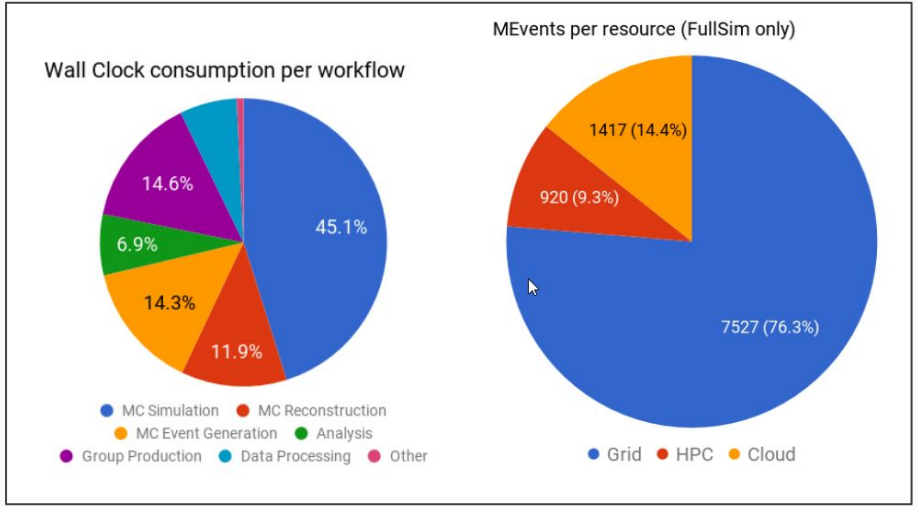

Resource usage per million events, 2017, J. Elmsheuser

Summary

- HPCs were used to help process extra data to stop resource shortages on the WLCG
- PanDA and Harvester were used on HPCs to run many ATLAS workflows, many of which are MC Simulations and Event Generation
  - Both are also used for effective resource utilization on HPC systems
- Some utilization of heterogeneous architectures by using GPUs for ATLAS workflows
- 10% of total events processed on HPC systems

Thank You!