
Plans for Exploitation of the ORNL Titan Machine

Richard P. Mount

*ATLAS Distributed Computing Technical Interchange Meeting
May 17, 2013*



Advancing the Era of Accelerated Computing

Titan Overview

The Oak Ridge Leadership Computing Facility (OLCF) has completed the first phase of an upgrade of the Jaguar system that will result in a hybrid-architecture Cray XK7 system named Titan, with a peak theoretical performance of more than 20 petaflops. Titan will be the first major supercomputing system to utilize a hybrid architecture, or one that utilizes both conventional 16-core AMD Opteron CPUs and unconventional NVIDIA Kepler graphics processing units (GPUs). The combination of CPUs and GPUs will allow Titan and future systems to overcome power and space limitations inherent in previous generations of high-performance computers.

Titan Specifications

Titan System Configuration	
Architecture:	Cray XK7
Processor:	16-Core AMD
Cabinets:	200
Nodes:	18,688 AMD Opterons
Cores/node:	16
Total cores:	299,008 Opteron Cores
Memory/node:	32GB + 6GB
Memory/core:	2GB
Interconnect:	Gemini
GPUs:	18,688 K20 Keplers
Speed:	20+ PF
Square Footage	4,352 sq feet

Current Status

Announcement

Beginning Friday, April 12th at 8:00 am, Titan was removed from production while the OLCF conducts the final tests to ensure that the system is ready for your science experiments. We expect Titan to be unavailable to users for final checkout and testing through mid- to late May. Further updates will be sent in the weekly messages and posted on the

Titan Overview

Operating System

- Titan employs the Cray Linux Environment as its OS. This consists of a full-featured version of Linux on the login nodes, and a Compute Node Linux microkernel on compute nodes. The microkernel is designed to minimize partition overhead allowing scalable, low-latency global communications.

File Systems

- The OLCF's center-wide Lustre file system, named [Spider](#), is available on Titan for computational work. With over 52,000 clients and (10) PB of disk space, it is the largest-scale Lustre file system in the world.

(from the) Titan User Guide

Cray Supercomputers are complex collections of different types of physical nodes/machines. For simplicity, we can think of Titan nodes as existing in one of three categories: *login* nodes, *service* nodes, or *compute* nodes.

Login nodes are shared by all users of a system, and should only be used for basic tasks such as file editing, code compilation, data backup, and job submission.

When a job is submitted to the batch system, the job submission script is first executed on a **service node**. Any job submitted to the batch system is handled in this way, including interactive batch jobs (e.g. via `qsub -l`).

On Cray machines, when the `aprun` command is issued within a job script (or on the command line within an interactive batch job), the binary passed to `aprun` is copied to and executed in parallel on a set of **compute nodes**.

Compute nodes run a Linux microkernel for reduced overhead and improved performance. Only User Work (Lustre) and Project Work (Lustre) storage areas are available to compute nodes.

Login and service nodes use AMD's Istanbul-based processor, while compute nodes use the newer Interlagos-based processor. Interlagos-based processors include instructions not found on Istanbul-based processors, so ...

Timeline of ATLAS Involvement with OLCF

- December 7, 2012
 - » RM is called by Ken Read (long ago an L3 grad student, now an ALICE physicist at ORNL and UTK)
 - » Ken wants to explore ATLAS/ALICE/nEDM etc./Geant4/OLCF common interests.
 - » nEDM = neutron Electric Dipole Moment experiment, one of several ORNL activities interested in heavy Geant4 usage
- February 1, 2013
 - » RM submits “Implementing and Accelerating Geant4-Based Simulations on Titan” proposal to the ALCC (ASCR Leadership Computing Challenge) program
 - » Request 10M core hours on Titan, plus 1 PB of disk storage, primarily focused on exploring use for ATLAS simulation. (About 1/3 of a US T2/year).

Timeline of ATLAS Involvement with OLCF

- February 11, 2012
 - » Meeting at ORNL: RM, Alexei, Sergey, Kaushik, SLAC G4 People, OLCF People, nEDM people
 - » The OLCF People (e.g. Jack Wells, Director of Science at the ORNL National Center for Computational Sciences) strongly encourage ATLAS efforts to exploit Titan
- March to May, 2013
 - » ATLAS and Geant4 people succeed in getting accounts on OLCF
- May 2013
 - » Result of ALCC application expected
- July 2013
 - » Start of 1-year allocation on Titan if ALCC is approved

Plans

- Leverage the existing (ASCR+HEP funded) “Big PanDA” project, plus ALCF experience to interface PanDA to Titan.
- Get ATLAS simulation running (ignoring GPUs):
 - » SLAC Geant4 group
 - » Wei Yang plus SLAC student(s)
 - » Need ATLAS simulation expertise
- Start producing useful simulation ($H \rightarrow b \bar{b}$?)
- Explore GPU use:
 - » Communications overheads mean that multithreading on the CPU side is required to mask the latency of the GPUs
 - » SLAC G4 group is now driving G4 multithreading
 - » Not clear when multithreaded G4 will be validated by ATLAS