Argonne NATIONAL LABORATORY

ALCF COMPUTING RESOURCES

MIRA

Mira, the ALCF's 10-petaflops IBM Blue Gene/Q supercomputer, opens the door for scientists to analyze data more efficiently, design products more quickly, and address some of society's biggest problems in ways that would otherwise be impossible. The system is capable of carrying out 10 quadrillion calculations per second. Mira is also among the most energy-efficient supercomputers, saving considerable energy through innovative chip designs and a unique water-cooling system.

- 48 racks
- 16 1600 MHz PowerPC A2 cores per node
- 49,152 nodes
- 786,432 cores
- 768 TiB RAM
- 5D torus interconnect
- 384 I/O nodes
- Peak performance of 10 petaflops

CETUS

The primary role of Cetus is to run small jobs to debug problems that occur on Mira. Cetus shares the same software environment and file systems as Mira.

- 4 racks
- 4,096 nodes
- 64 TB RAM
- 5D torus interconnect
- 32 I/O nodes
- Peak performance of 838 teraflops

VESTA

As the ALCF's test and development platform, Vesta serves as a launching pad for researchers preparing to use Mira. Vesta has the same architecture as Mira, but on a much smaller scale.

- 2 racks
- 2,048 nodes
- 32 TB RAM
- 5D torus interconnect
- 32 I/O nodes
- Peak performance of 419 teraflops

COMING TO THE ALCF: AURORA AND THETA

In April 2015, the U.S. Department of Energy announced its investment to deliver Aurora and Theta, the ALCF's next-generation supercomputers.

Designed in collaboration with industry leaders Intel and Cray, Aurora is scheduled for delivery in 2018. Aurora's revolutionary architecture features Intel's HPC scalable system framework and second-generation Intel® Omni-Path Fabric, providing a peak performance of 180 petaflops.

A second system, named Theta, will be delivered in 2016 to serve as a bridge between Mira and Aurora. Based on Intel's second-generation Xeon Phi processor, Theta will enable breakthrough science and engineering research while providing an early production system to help ALCF users transition their applications to the new technology.

ALCF users have access to several computing resources, including one of the most powerful supercomputers in the world.

COOLEY

Cooley is the ALCF's analysis and visualization cluster. Equipped with graphics processing units (GPUs), Cooley converts computational data from Mira into high-resolution visual representations. The resulting images and videos help users to better analyze and understand the data generated by Mira. Cooley can also be used for statistical analysis, helping to pinpoint trends in the simulation data. Additionally, the system is capable of preprocessing efforts, such as meshing, to assist users preparing for Mira simulations. Cooley shares file systems with Mira, enabling direct access to Mira-generated results.

Each Cooley node has:

- 2 2.4 GHz Intel Haswell E5-2620 v3 6-core processors
- NVIDIA Tesla K80 GPU accelerator containing two Kepler GK210 GPUs
- 384 GB RAM
- 24 GB GPU RAM

The full Cooley system has:

- 126 nodes
- 1,512 cores
- FDR Infiniband interconnect
- 47 TB RAM
- 3 TB GPU RAM
- Peak performance of 293 teraflops

DATA STORAGE

At the ALCF, disk storage provides intermediate-term storage for active projects, offering a means to access, analyze, and share simulation results. Tape storage is used to archive data from completed projects.

DISK STORAGE The Blue Gene/Q data systems consist of 384 I/O nodes that connect to 22 storage arrays that control 13,000 disk drives with a total useable capacity of 27 PB and a maximum aggregate transfer speed of 330 GB/s over two file systems. The ALCF uses the General Parallel File System to access the storage.

TAPE STORAGE The ALCF has two 10,000-slot libraries using LTO 6 tape technology. The LTO tape drives have built-in hardware compression with compression ratios typically between 1.25:1 and 2:1, depending on the data, giving an effective capacity of 26-40 PB.

NETWORKING

The Blue Gene/Q systems have an internal proprietary network for communicating between nodes. InfiniBand enables communication between the I/O nodes and the storage system. Ethernet is used for external user access, and for maintenance and management of the systems.

The ALCF connects to other research institutions using up to 100 Gb/s of network connectivity. Scientists can transfer datasets to and from other institutions over fast research networks, such as the Energy Science Network (ESNet) and Internet2.