Argonne Leadership Computing Facility

Blue Gene/Q Systems and Supporting Resources

The Argonne Leadership Computing Facility is part of the U.S. Department of Energy's Advanced Scientific Computing Research (ASCR) leadership computing program to enable significant advances in computational science.

The ALCF provides substantial awards of supercomputing time and user support services to scientific and engineering teams worldwide to conduct large-scale modeling and simulation research aimed at solving the world's largest and most complex problems related to energy and the environment.

Virtually any process or problem can be advanced with precision and speed using Mira, the ALCF's IBM Blue Gene/Q supercomputer.

Our user community conducts R&D spanning a wide range of disciplines. Active projects include:

- Materials science research to extend the range and power of electric vehicles
- Integration and modeling of global climate science data
- Prototyping technologies to improve fuel efficiency and to reduce aerodynamic noise
- Optimizing the next-generation power grid
- Modeling thermal hydraulics in next-generation nuclear reactors

MIRA

Mira, the ALCF's 10-petaflops supercomputer, opens the door for researchers and industry to analyze data more efficiently, design products more quickly, and address some of society's biggest problems in ways that would otherwise be impossible. An engineering marvel, the system is capable of carrying out 10 quadrillion calculations per second. To put this speed into perspective, Mira can do in one hour what it would take an average personal computer 20 years to achieve.

In addition to being one of the fastest computers in the world, Mira is also among the most energy efficient. The supercomputer saves considerable energy through innovative chip designs and a unique water-cooling system.

By fitting more cores onto a single chip, Mira speeds the communication between cores and saves the energy lost when transporting data across long distances. Mira's water-cooling system uses copper tubes to pipe cold water directly alongside the chips, saving power by eliminating an extra cooling step.

Beyond enabling scientific discoveries in a sustainable way, Mira itself is a stepping stone toward the next great goal of supercomputing: exascale speed, where computers will operate a thousand times faster than today's top machines.

Mira's Blue Gene/Q system consists of:

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

Argonne NATIONAL LABORATORY

CETUS

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

Cetus has four racks, 4,096 nodes, 64 TB RAM, 5D torus interconnect, 32 I/O nodes, and a peak performance of 838 teraflops.

VESTA

Vesta is the ALCF's test and development platform, serving as a launching pad for researchers planning to use Mira. Vesta has the same architecture as Mira, but on a much smaller scale (two computer racks compared to Mira's 48 racks). This system enables researchers to debug and scale up codes for the Blue Gene/Q architecture in preparation for Mira.

Vesta is also a resource for users applying for a Director's Discretionary allocation and researchers preparing proposals for INCITE and ALCC. By allocating time on Vesta for testing and development, Mira is kept clear for capability jobs.

Vesta has two racks, 2,048 nodes, 32 TB RAM, 5D torus interconnect, 32 I/O nodes, and a peak performance of 419 teraflops.

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

The full system has 126 nodes, 1,512 cores, a FDR InfiniBand interconnect, 47 TB of RAM, 3 TB of GPU RAM, and a GPU peak performance (aggregate) of over 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 GPFS 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.

U.S. DEPARTMENT OF ENERGY

Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC