ALCF Computing Resources

The Argonne Leadership Computing Facility operates supercomputing resources that support large-scale computing projects aimed at solving some of the world's most complex problems in science and engineering.

Aurora

Designed in collaboration with Intel and Cray, Aurora will be one of the nation's first exascale systems when it arrives in 2021. The supercomputer will be based on Intel's Xeon Scalable processors and high-performance Intel X^e GPU compute accelerators. It will rely on Cray's Shasta exascale-class architecture and Slingshot interconnect technology, which can provide concurrent support for advanced simulation and modeling, AI, and analytics workflows. Aurora will leverage historical advances in software investments along with increased application portability via Intel's oneAPI. The system will also introduce a new I/O system called Distributed Asynchronous Object Storage (DAOS) to meet the needs of new exascale workloads.

- CY 2021 delivery
- 2 Intel Xeon scalable "Sapphire Rapids" processors; 6 X^e arch-based GPUs; Unified Memory Architecture; 8 fabric endpoints
- X^e arch-based "Ponte Vecchio" GPU archticture; Tile-based, chiplets, HBM stack, Foveros 3D integration, 7nm
- On-Node Interconnect CPU-GPU interconnect: PCIe on-node; GPU-GPU interconnect: X^e Link
- > 10 PB aggregate system memory
- Cray Slingshot; Dragonfly topology with adaptive routing
- ≥ 1EF DP sustained performance
- 25.6 Tb/s per network switch, from 64–200 Gbs ports (25 GB/s per direction)
- ≥ 230 PB, ≥ 25 TB/s (DAOS) high-performance storage
- Intel oneAPI, MPI, OpenMP, C/C++, Fortran, SYCL/DPC++
- Cray Shasta software stack + Intel Enhancements + Data and Learning
- Cray Shasta
- > 100 cabinets

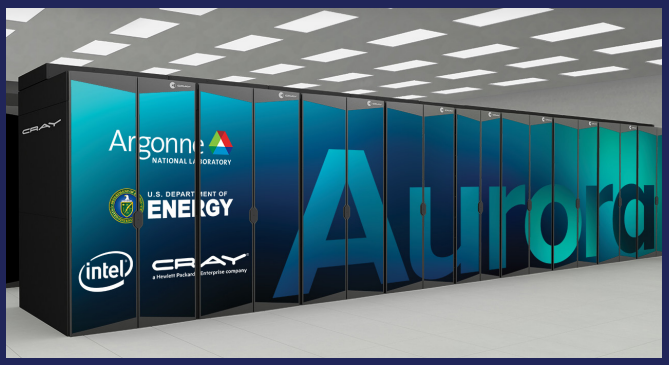

Theta

Theta, an 11.69 petaflops Intel-Cray system, is the engine that drives breakthrough computational science and engineering research at the ALCF. Theta is a massively parallel, many-core system based on Intel processors and interconnect technology, an advanced memory architecture, and a Lustre-based parallel file system, all integrated by Cray's HPC software stack.

- Intel-Cray XC40 architecture
- 11.69 petaflops
- 64-core, 1.3-GHz Intel Xeon Phi 7230 processor per node
- 4,392 nodes
- 281,088 cores
- 843 TB of memory
- 70 TB of high-bandwidth memory
- Aries interconnect with Dragonfly configuration
- 24 racks

U.S. DEPARTMENT OF ENERGY

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

CONTACT
media@alcf.anl.gov
alcf.anl.gov

Cooley

Cooley is the ALCF's data analysis and visualization cluster.

- Intel Haswell architecture
- 293 teraflops
- Two 6-core, 2.4-GHz Intel E5–2620 processors per node
- 1 NVIDIA Tesla K80 GPU per node
- 126 nodes
- 1,512 cores
- 47 TB of memory
- 3 TB of GPU memory
- FDR InfiniBand interconnect
- 6 racks

Iota

Iota serves as the ALCF's Intel-Cray test and development platform.

- Intel-Cray XC40 architecture
- 117 teraflops
- 64-core, 1.3-GHz Intel Xeon Phi 7230 processor per node
- 44 nodes
- 2,816 cores
- 12.3 TB of memory
- 1 TB of high-bandwidth memory
- Aries interconnect with Dragonfly configuration
- 1 rack

Testbeds

Through Argonne's Joint Laboratory for System Evaluation (JLSE), the ALCF provides access to next-generation hardware and software to explore low-level experimental computer and computational science, including operating systems, messaging, compilers, benchmarking, power measurements, I/O, and new file systems. JLSE testbeds include:

- Intel DAOS nodes (DCPMM and NVMe storage) with Intel DAOS file system v0.9
- Intel Xeon Phi Knights Landing Cluster
- NVIDIA DGX–1
- Intel Xeon Platinum Skylake Cluster
- Intel Iris Gen9 nodes (integrated GPUs) with Intel oneAPI (public beta)
- HPE Comanche Prototype ARM64 Cluster
- IBM Power System AC922
- Supermicro NVIDIA V100 and P100 cluster
- Atos Quantum Learning Machine
- Kubernetes Cluster with Rancher
- Petrel v3 (Ceph 3.2 PB filesystem)

Data Storage

The ALCF's data storage system is used to retain the data generated by simulations and visualizations. 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 Theta system consists of 30 I/O nodes that connect to a storage array that controls 2,300 disk drives with a total useable capacity of 9 PB and a maximum aggregate transfer speed of 240 GB/s. Theta uses Lustre to access this storage. The ALCF also utilizes a 10 PB file system based on an IBM Elastic Storage Server (ESS) to host data for science running on the Theta and Cooley systems. The ESS system is a software defined storage system based on IBM's GPFS file system and consists of 60 I/O nodes controlling 7,260 disk drives.

TAPE STORAGE

The ALCF has three 10,000-slot libraries using LTO-6 and LTO-8 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 approximately 65 PB.

Networking

Theta has 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 ESnet and Internet2.

Theta, an 11.69 petaflops Intel-Cray supercomputer, is the engine that drives scientific discoveries at the ALCF.

CONTACT
media@alcf.anl.gov
alcf.anl.gov