Argonne NATIONAL LABORATORY

Argonne Leadership Computing Facility

Aurora

Argonne's exascale system will be used to dramatically advance scientific discovery and innovation.

Argonne Leadership Computing Facility's first exascale computer integrates world-class compute capabilities with powerful memory, storage, and networking to advance the frontiers of science.

A Brand-New Class of System

Built in partnership with Intel and Hewlett Packard Enterprise (HPE), Aurora is one of the fastest supercomputers in the world and one of the nation's first exascale systems. It is also one of the world's largest supercomputers, occupying 10,000 square feet and weighing 600 tons.

The system is powered by 21,248 Intel Xeon CPU Max Series processors and 63,744 Intel Data Center GPU Max Series processors. Notably, Aurora features more GPUs and more network endpoints in its interconnect technology than any system to date.

The system relies on HPE Cray EX supercomputer exascale-class architecture and HPE Slingshot technology, which can provide concurrent support for advanced simulation and modeling, AI, and analytics workflows.

Aurora leverages historical advances in software investments along with increased application portability via Intel's oneAPI. The supercomputer also introduces a new I/O system called Distributed Asynchronous Object Storage (DAOS) to meet the needs of new exascale workloads.

Aurora is built to tackle a wide range of scientific problems including designing more efficient airplanes, investigating the mysteries of the cosmos, modeling the impacts of climate change, and accelerating the discovery of new materials.

System Overview

Intel-Hewlett Packard Enterprise (HPE) System

Exascale System

10,624 Compute Blades

21,248 Intel Xeon CPU Max Series

63,744 Intel Data Center GPU Max Series

166 Racks

U.S. DEPARTMENT OF ENERGY

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

July 2024

The Compute Blade

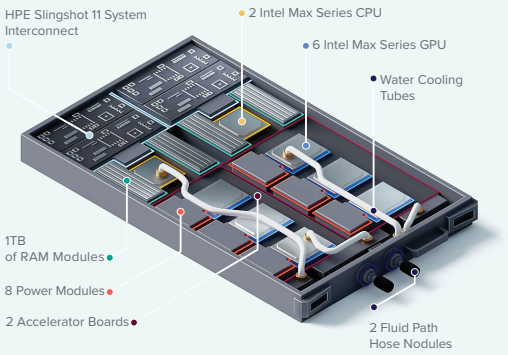

The Compute Racks

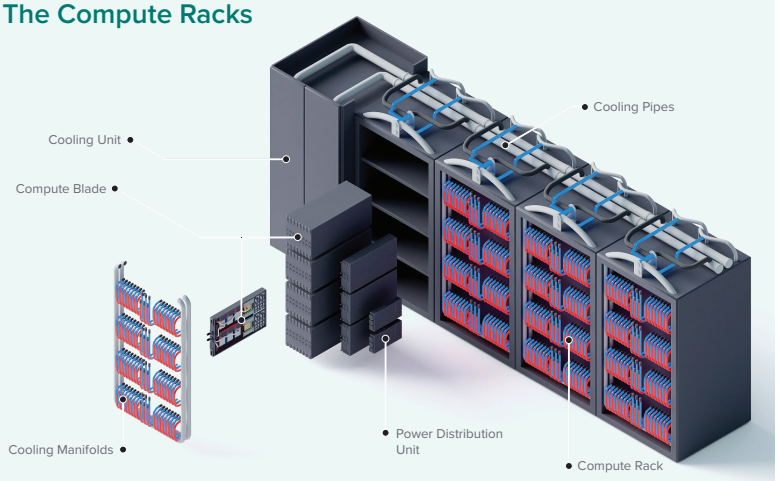

Science on Day One

The Aurora Early Science Program prepares key applications for Aurora's scale and architecture, and helps solidify libraries and infrastructure for other production applications to run on the system. The program supports 15 projects that cover a wide range of scientific areas and computational methods, including research involving data science and machine learning approaches as well as traditional modeling and simulation-based campaigns.

In addition to fostering application readiness for the supercomputer, the Early Science Program allows researchers to pursue innovative computational science campaigns not possible on today's leadership-class supercomputers.

Early Performance

The Aurora Early Science teams have demonstrated strong performance gains while scaling and optimizing simulation, data, and AI applications for the system. The initial Aurora projects include efforts to advance research in cosmology, fusion energy science, drug discovery and the design of new materials for clean energy technologies.

Compute Node

2 Intel Xeon CPU Max Series processors; 6 Intel Data Center Max GPUs; Unified Memory Architecture; 8 fabric endpoints; RAMBO

GPU Architecture

Intel Data Center GPU Max Series; Tile-based, chiplets, HBM stack, Foveros 3D integration

On-Node Interconnect

CPU-GPU: PCIe
GPU-GPU: Xe Link

Aggregate System Memory

20.4 PB

High-Performance Storage

230 PB, 31 TB/s (DAOS)

System Interconnect

HPE Slingshot 11; Dragonfly topology with adaptive routing

Network Switch

25.6 Tb/s per switch, from 64–200 Gbs ports (25 GB/s per direction)

Programming Environment

Intel oneAPI, MPI, OpenMP, C/C++, Fortran, SYCL/DPC++

Software Stack

HPE Cray EX software stack +Intel Enhancements + Data and Learning

Platform

HPE Cray EX supercomputer

System Performance

Exascale

System Size

10,624 nodes

Contact

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