


Aurora

The Argonne Leadership Computing Facility's future exascale system will be used to dramatically advance scientific discovery and innovation.



Aurora's compute nodes will be equipped with two Intel Xeon Scalable processors and six general-purpose GPUs based on Intel's X^e architecture.
Image: Intel Corporation

Argonne National Laboratory's first exascale computer is coming soon, and will exclusively serve the research community. Scientists will use the new machine, named Aurora, to pursue some of the farthest-reaching science and engineering breakthroughs ever achieved with supercomputing.

A Brand-New Class of System

Each machine generation provides a fresh challenge to U.S. computer manufacturers—from the racks to the processors to the networking to the I/O system. Similarly, fulfilling the science potential of each new computing architecture requires significant changes to today's software. The initiative is, and will continue to be, guided by pioneering visionaries in the mathematics and computational science community, stewarded by the DOE's Office of Science, and operated at the cutting edge.

Aurora will be based on Intel's Xeon Scalable processors and high-performance Intel X^e GPU compute accelerators. The system will rely 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 will leverage historical advances in software investments along with increased application portability via Intel's oneAPI. The supercomputer will also introduce a new I/O system called Distributed Asynchronous Object Storage (DAOS) to meet the needs of new exascale workloads.

Aurora Early Science Program

The Aurora Early Science Program is preparing key applications for Aurora's scale and architecture, and will solidify libraries and infrastructure to pave the way for other production applications to run on the system. The program has selected 15 projects, proposed by investigator-led teams from universities and national labs and covering a wide range of scientific areas and numerical methods.

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

Exascale Training

The ALCF and the Exascale Computing Project (ECP) are offering several training opportunities, including workshops, webinars, and hackathons, to help researchers prepare for Aurora and other DOE exascale systems. Stay tuned to the ALCF and ECP websites for upcoming opportunities.



The ALCF's Intel-HPE exascale system is scheduled to arrive in 2022.

SYSTEM SPECS

Delivery

CY 2022

Compute Node

2 Intel Xeon scalable "Sapphire Rapids" processors; 6 Xe arch- based GPUs; Unified Memory Architecture; 8 fabric; RAMBO

GPU Architecture

Xe arch-based "Ponte Vecchio"
GPUPile-based, chiplets, HBM stack, Foveros 3D integration

On-Node Interconnect

CPU-GPU: PCIe
GPU-GPU: Xe Link

Aggregate System Memory

>10 PB

System Interconnect

HPE Slingshot 11; Dragonfly topology with adaptive routing

Sustained Performance

$\geq 1\text{EF DP}$

Network Switch

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

High-Performance Storage

$\geq 230\text{ PB}$, $\geq 25\text{ TB/s}$ (DAOS)

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 Size

>9,000 Nodes