

Polaris: Paving the Way to Exascale

Argonne National Laboratory's new supercomputing resource will accelerate efforts to prepare applications and workloads for Aurora.



Polaris, a hybrid CPU-GPU system built by HPE, will help ready scientists for the arrival of the lab's Aurora exascale supercomputer.

With the arrival of the Polaris supercomputer at the Argonne Leadership Computing Facility (ALCF), researchers now have a powerful new tool to prepare for science in the exascale era.


Polaris, a leading-edge system built by Hewlett Packard Enterprise (HPE) in collaboration with NVIDIA and AMD, gives scientists and application developers a platform to test and optimize codes for Aurora, Argonne's upcoming Intel-HPE exascale supercomputer. Like Aurora, Polaris is a hybrid system equipped with both graphics processing units (GPUs) and central processing units (CPUs).

Designed with industry-leading high-performance computing (HPC) and artificial intelligence (AI) technologies, Polaris will help advance investigations into society's most complex and pressing issues, from understanding the biology of viruses to revealing the secrets of the universe. It will also augment Argonne's ongoing research efforts in areas such as clean energy, climate resilience, and manufacturing.

The HPE Apollo 6500 Gen10+ based supercomputer will be powered by 3rd Gen AMD EPYC processors and NVIDIA A100 Tensor Core GPUs. The system will deliver approximately 44 petaflops of peak double precision performance and nearly 1.4 exaflops of theoretical AI performance based on mixed-precision compute capabilities.

Initially, Polaris will be dedicated to research teams participating in DOE's Exascale Computing Project, ALCF's Aurora Early Science Program, and the ALCF Data Science Program. In 2022, Polaris will be made available to the broader HPC community for a wide range of science and engineering projects.

With architectural features that support AI and data-centric workloads, Polaris is particularly well suited to handle the massive amounts of data being produced by large-scale simulations, experimental facilities, and vast sensor networks. Initial efforts will be focused on integrating HPC and AI workloads with experimental facilities located at Argonne, including the Advanced Photon Source and the Center for Nanoscale Materials.



Polaris will be equipped with 2,240 NVIDIA A100 Tensor Core GPUs.
(Image courtesy of NVIDIA)

SYSTEM SPECS

Peak Performance

44 petaflops (double precision)

Chassis

HPE Apollo 6500 Gen10+ (two nodes)

System Size

560 Nodes

Node Architecture

1 3rd Gen AMD EPYC processor; 4 NVIDIA A100 Tensor Core GPUs; Unified Memory Architecture; NVMe SSDs

Node Memory

160 GB HBM + 512 GB DDR4

System Fabric

HPE Slingshot 11; 2 endpoints per node; Dragonfly topology