

# ALCF Computing Resources

Supporting large-scale, computationally intensive scientific projects.



## Supercomputer Systems

The ALCF is available to any researcher in the world with a large-scale computing problem. Researchers gain access to ALCF systems through competitive, peer-reviewed allocation programs supported by DOE and Argonne National Laboratory, and publish their findings in high-impact journals and publications.

SYSTEM NAME	AURORA	POLARIS	SOPHIA
Architecture	HPE Cray EX	HPE Apollo 6500 Gen10+	NVIDIA DGX A100
Peak Performance	2 EF	34 PF	3.9 PF
Processors per Node	2 Intel Xeon CPU Max Series processors	3rd Gen AMD EPYC	2 AMD EPYC 7742
GPU per Node	6 Intel Data Center GPU Max Series	4 NVIDIA A100 Tensor Core	8 NVIDIA A100 Tensor Core
Nodes	10,624	560	24
Cores	9,264,128	17,920	3,072
Memory	20.4 PB	280 TB (DDR4); 87.5 TB (HBM)	26 TB (DDR4); 8.32 TB (GPU)
Interconnect	HPE Slingshot 11 with Dragonfly Configuration	HPE Slingshot 11 with Dragonfly Configuration	NVIDIA HDR InfiniBand
Racks	166	40	7



## ALCF AI Testbed

The ALCF AI Testbed provides an infrastructure of next-generation AI-accelerator machines for research campaigns at the intersection of AI and science.

SYSTEM NAME	CEREBRAS CS-2	SAMBANOVA CARDINAL SN30	GROQRACK	GRAPHCORE BOW POD-64	HABANA GAUDI-1
System Size	2 Nodes (Each with a Wafer-Scale Engine) Including MemoryX and SwarmX	64 Accelerators (8 Nodes and 8 Accelerators per Node)	72 Accelerators (9 Nodes and 8 Accelerators per Node)	64 Accelerators (4 Nodes and 16 Accelerators per Node)	16 Accelerators (2 Nodes and 8 Accelerators per Node)
Compute Units per Accelerator	850,000 Cores	1,280 Programmable Compute Units	5,120 Vector ALUs	1,472 Independent Processing Units	8 TPC + GEMM Engine
Single Accelerator Performance (TFlops)	> 5,780 (FP16)	>660 (BF16)	>188 (FP16) >750 (INT8)	>250 (FP16)	>150 (FP16)
Software Stack Support	Cerebras SDK, TensorFlow, PyTorch	SambaFlow, PyTorch	GroqWare SDK, ONNX	PopART, TensorFlow, PyTorch, ONNX	SynapseAI, TensorFlow, PyTorch
Interconnect	Ethernet-based	Ethernet-based	RealScale™	IPU Link	Ethernet-based

## Data Storage Systems

ALCF disk storage systems provide intermediate-term storage for users to access, analyze, and share computational and experimental data. Tape storage is used to archive data from completed projects.

SYSTEM NAME	EAGLE	GRAND	SWIFT	TAPE STORAGE
File System	Lustre	Lustre	Lustre	—
Storage System	HPE ClusterStor E1000	HPE ClusterStor E1000	All NVMe Flash Storage Array	LTO6 and LTO8 Tape Technology
Usable Capacity	100 PB	100 PB	123 TB	300 PB
Sustained Data Transfer Rate	650 GB/s	650 GB/s	48 GB/s	—
Disk Drives	8,480	8,480	24	—

## Joint Laboratory for System Evaluation

Argonne’s Joint Laboratory for System Evaluation (JLSE) provides access to leading-edge testbeds for research aimed at evaluating future extreme-scale computing systems, technologies, and capabilities.

**Florentia:** Test and development system equipped with early versions of the Intel Max Series CPUs and Intel Max Series GPUs that power Aurora

**NVIDIA and AMD GPUs:** Clusters of NVIDIA V100, A100, and A40 GPUs, and AMD MI50 and MI100 GPUs for preparing applications for heterogeneous computing architectures

**Aurora Software Development Kit:** Frequently updated version of the publicly available Intel oneAPI toolkit for Aurora development

**Arm Ecosystem:** Apollo 80 Fujitsu A64FX Arm system, NVIDIA Ampere Arm and A100 test kits, and an HPE Comanche with Marvell ARM64 CPU platform provide an ecosystem for porting applications and measuring performance on next-generation systems

**Presque:** Intel DAOS nodes for testing the Aurora storage system

**Edge Testbed:** NVIDIA Jetson Xavier and Jetson Nano platforms provide a resource for testing and developing edge computing applications

**Arcticus, DevEP, Iris:** Intel discrete and integrated GPU testbeds to support the development, optimization, and scaling of applications and software for Aurora

**NVIDIA Bluefield-2 DPU SmartNICs:** Platform used for confidential computing, MPICH offloading, and APS data transfer acceleration

**NextSilicon Maverick:** First-generation product being tested by Argonne researchers

**Atos Quantum Learning Machine:** Platform for testing and developing quantum algorithms and applications