

ALCF AI Testbed

The ALCF AI Testbed is a collection of the world’s most advanced AI accelerators available for open science.



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

Contact

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

ALCF AI Testbed

The AI accelerator systems are available to the research community with data-intensive problems to solve. Researchers can submit project proposals for the systems via the ALCF’s Director’s Discretionary program. For more information, visit: alcf.anl.gov/alcf-ai-testbed

SYSTEM NAME	CEREBRAS CS-2	SAMBANOVA 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
Estimated Performance of a Single Accelerator (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