Dense, Accelerated Supercomputing

The Cray® CS-Storm™ system is purpose-built for AI

Cray can help you move artificial intelligence (AI) from pilot to production. Our dense, accelerated Cray® CS-Storm™ supercomputer is an integrated system ready to tackle machine learning and deep learning problems at production scale. As machine learning and deep learning become integral to production workloads, Cray's CS-Storm system is your fastest path to value and the foundation for future innovation through new discoveries.

DESIGNED FOR SPEED

Powered by NVIDIA, Intel and Nallatech accelerators and processors, each CS-Storm server node and rack system is integrated by Cray to deliver maximum performance across the broadest range of machine learning and deep learning environments.

DESIGNED FOR SCALE

As machine learning and deep learning become core requirements for business and scientific discovery, the need for consistent and timely processing is driving the use of supercomputer-scale systems, which are designed to handle the largest data sets and perform the most complex calculations. Cray has the depth of experience to guide organizations on the journey from AI trials and pilots to production workloads.

DESIGNED FOR SIMPLICITY

Designing, implementing and using an AI system doesn't have to be difficult. You can rely on Cray, the expert in production supercomputing, to simplify your environment. Our systems are built on open, industry-standard technologies. We deliver a complete solution for AI including high-performance storage, system management, developer tools and validated deep learning frameworks.

A Look Inside

	CS-Storm 500GT	CS-Storm 500NX
Processors & Accelerators	Up to 8 450W or 10 400W accelerators NVIDIA® Tesla® P40 or P100 PCIe GPU accelerators Nallatech FPGA accelerators Two Intel® Xeon® "Skylake" processors	Up to 10 NVIDIA Tesla P100 SXM2 GPU accelerators Two Intel Xeon E5-2600 v4 "Broadwell" processors
Interconnects	PCIe with single-root and balanced configuration option InfiniBand® or Intel® OmniPath architecture	NVIDIA® NVLink™ for 8-way GPU-to-GPU communications InfiniBand or Intel OmniPath architecture
Architecture	19" 3U or 4U rackmount chassis 16 hot-swap 2.5" drives (up to 8 NVMe)	19" 4U rackmount chassis 12 2.5" drives (up to 4 NVMe)

Cray Inc. ▪ 901 Fifth Avenue, Suite 1000 ▪ Seattle, WA 98164 ▪ Tel: 206.701.2000 ▪ Fax: 206.701.2500 ▪ www.cray.com