NVIDIA.

NVIDIA Turbocharges Extreme-Scale AI for Argonne National Laboratory's Polaris Supercomputer

Largest GPU-Powered Supercomputer for U.S. Department of Energy's Argonne Lab Will Enable Scientific Breakthroughs in Era of Exascale AI

The largest GPU-based supercomputer at the U.S. Department of Energy's Argonne National Laboratory will run on NVIDIA's accelerated computing platform, the company announced today.

The Polaris supercomputer, hosted at the laboratory's Argonne Leadership Computing Facility (ALCF), will supercharge research and discovery with extreme scale for users' algorithms and science. Accelerated by 2,240 NVIDIA® A100 Tensor Core GPUs, the system can achieve almost 1.4 exaflops of theoretical AI performance and approximately 44 petaflops of peak double-precision performance.

Polaris, to be built by Hewlett Packard Enterprise, will combine simulation and machine learning by tackling data-intensive and AI high performance computing workloads, powered by 560 total nodes, each with four NVIDIA A100 GPUs.

"The era of exascale AI will enable scientific breakthroughs with massive scale to bring incredible benefits for society," said Ian Buck, vice president and general manager of Accelerated Computing at NVIDIA. "NVIDIA's GPU-accelerated computing platform provides pioneers like the ALCF breakthrough performance for next-generation supercomputers such as Polaris that let researchers push the boundaries of scientific exploration."

"Polaris is a powerful platform that will allow our users to enter the era of exascale AI," said ALCF Director Michael E. Papka. "Harnessing the huge number of NVIDIA A100 GPUs will have an immediate impact on our data-intensive and AI HPC workloads, allowing Polaris to tackle some of the world's most complex scientific problems."

The system will accelerate transformative scientific exploration, such as advancing cancer treatments, exploring clean energy and propelling particle collision research to discover new approaches to physics. And it will transport the ALCF into the era of exascale AI by enabling researchers to update their scientific workloads for Aurora, Argonne's forthcoming exascale system.

Polaris will also be available to researchers from academia, government agencies and industry through the ALCF's peer-reviewed allocation and application programs. These programs provide the scientific community with access to the nation's fastest supercomputers to address "grand challenges" in science and engineering.

About NVIDIA
NVIDIA's (NASDAQ: NVDA) invention of the GPU in 1999 sparked the growth of the PC gaming market and has redefined modern computer graphics, high performance computing and AI. The company's pioneering work in accelerated computing and artificial intelligence is reshaping trillion-dollar industries, such as transportation, healthcare and manufacturing, and fueling the growth of many others. More information at https://nvidianews.nvidia.com/.

Certain statements in this press release including, but not limited to, statements as to: the Polaris supercomputer running on NVIDIA's accelerated computing platform, being built by Hewlett Packard Enterprise, and its impact, performance and availability; the benefits, impact, and performance of our products and services, including NVIDIA's accelerated computing platform and NVIDIA A100 Tensor Core GPUs; the Polaris supercomputer supercharging research and discovery with extreme scale for users' algorithms and science; and the era of exascale AI enabling scientific breakthroughs with massive scale to bring benefits for society are forward-looking statements that are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic conditions; our reliance on third parties to manufacture, assemble, package and test our products; the impact of technological development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; as well as other factors detailed from time to time in the most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its annual report on Form 10-K and quarterly reports on Form 10-Q. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

notice.

Alex Shapiro
Enterprise Networking
1-415-608-5044
ashapiro@nvidia.com