NVIDIA TITAN V Transforms the PC into AI Supercomputer

Volta-Powered GPU Delivers 110 Teraflops of Deep Learning Horsepower — 9x Its Predecessor — to Researchers and Scientists

Conference on Neural Information Processing Systems -- NVIDIA today introduced TITAN V, the world's most powerful GPU for the PC, driven by the world's most advanced GPU architecture, NVIDIA Volta.

Announced by NVIDIA founder and CEO Jensen Huang at the annual NIPS conference, TITAN V excels at computational processing for scientific simulation. Its 21.1 billion transistors deliver 110 teraflops of raw horsepower, 9x that of its predecessor, and extreme energy efficiency.

“Our vision for Volta was to push the outer limits of high performance computing and AI. We broke new ground with its new processor architecture, instructions, numerical formats, memory architecture and processor links," said Huang. “With TITAN V, we are putting Volta into the hands of researchers and scientists all over the world. I can't wait to see their breakthrough discoveries."

NVIDIA Supercomputing GPU Architecture, Now for the PC
TITAN V's Volta architecture features a major redesign of the streaming multiprocessor that is at the center of the GPU. It doubles the energy efficiency of the previous generation Pascal™ design, enabling dramatic boosts in performance in the same power envelope.

New Tensor Cores designed specifically for deep learning deliver up to 9x higher peak teraflops. With independent parallel integer and floating-point data paths, Volta is also much more efficient on workloads with a mix of computation and addressing calculations. Its new combined L1 data cache and shared memory unit significantly improve performance while also simplifying programming.

Fabricated on a new TSMC 12-nanometer FFN high-performance manufacturing process customized for NVIDIA, TITAN V also incorporates Volta's highly tuned 12GB HBM2 memory subsystem for advanced memory bandwidth utilization.

Free AI Software on NVIDIA GPU Cloud
TITAN V's incredible power is ideal for developers who want to use their PCs to do work in AI, deep learning and high performance computing.

Users of TITAN V can gain immediate access to the latest GPU-optimized AI, deep learning and HPC software by signing up at no charge for an NVIDIA GPU Cloud account. This container registry includes NVIDIA-optimized deep learning frameworks, third-party managed HPC applications, NVIDIA HPC visualization tools and the NVIDIA TensorRT™ inferencing optimizer.

Immediate Availability
TITAN V is available to purchase today for $2,999 from the NVIDIA store in participating countries.

Keep Current on NVIDIA
Subscribe to the NVIDIA blog, follow us on Facebook, Google+, Twitter, LinkedIn and Instagram, and view NVIDIA videos on YouTube.

About NVIDIA
NVIDIA’s (NASDAQ:NVDA) invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined modern computer graphics and revolutionized parallel computing. More recently, GPU deep learning ignited modern AI — the next era of computing — with the GPU acting as the brain of computers, robots and self-driving cars that can perceive and understand the world. More information at http://nvidianews.nvidia.com/.

Certain statements in this press release including, but not limited to, statements as to: the impact, performance, efficiencies, benefits and availability of TITAN V GPUs; TITAN V setting new performance standards for what is possible in a PC; and the value and impact TITAN V will deliver to deep learning researchers and its use to help change the world 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 reports NVIDIA files with the Securities and Exchange Commission, or SEC, including its Form 10-Q for the fiscal period ended October 29, 2017. 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.

Media Contacts

Bryan Del Rizzo
+1 408 486 2772
bdelrizzo@nvidia.com