



NVIDIA GPUs Power Facebook's New Deep Learning Machine

New Purpose-Built Computing Platform for Developing Deep Neural Networks

NVIDIA today announced that Facebook will power its next-generation computing system with the NVIDIA® Tesla® Accelerated Computing Platform, enabling it to drive a broad range of machine learning applications.

While training complex deep neural networks to conduct machine learning can take days or weeks on even the fastest computers, the Tesla platform can slash this by 10-20x. As a result, developers can innovate more quickly and train networks that are more sophisticated, delivering improved capabilities to consumers.

Facebook is the first company to adopt [NVIDIA Tesla M40 GPU accelerators](#), introduced last month, to train deep neural networks. They will play a key role in the new ["Big Sur" computing platform](#), Facebook AI Research's (FAIR) purpose-built system designed specifically for neural network training.

"Deep learning has started a new era in computing," said Ian Buck, vice president of accelerated computing at NVIDIA. "Enabled by big data and powerful GPUs, deep learning algorithms can solve problems never possible before. Huge industries from web services and retail to healthcare and cars will be revolutionized. We are thrilled that NVIDIA GPUs have been adopted as the engine of deep learning. Our goal is to provide researchers and companies with the most productive platform to advance this exciting work."

In addition to reducing neural network training time, GPUs offer a number of other advantages. Their architectural compatibility from generation to generation provides seamless speed-ups for future GPU upgrades. And the Tesla platform's growing global adoption facilitates open collaboration with researchers around the world, fueling new waves of discovery and innovation in the machine learning field.

Big Sur Optimized for Machine Learning

NVIDIA worked with Facebook engineers on the design of Big Sur, optimizing it to deliver maximum performance for machine learning workloads, including the training of large neural networks across multiple Tesla GPUs. Two times faster than Facebook's existing system, Big Sur will enable the company to train twice as many neural networks -- and to create neural networks that are twice as large -- which will help develop more accurate models and new classes of advanced applications.

"The key to unlocking the knowledge necessary to develop more intelligent machines lies in the capability of our computing systems," said Serkan Piantino, engineering director for FAIR. "Most of the major advances in machine learning and AI in the past few years have been contingent on tapping into powerful GPUs and huge data sets to build and train advanced models."

The addition of Tesla M40 GPUs will help Facebook make new advancements in machine learning research and enable teams across its organization to use deep neural networks in a variety of products and services.

First Open Sourced AI Computing Architecture

Big Sur represents the first time a computing system specifically designed for machine learning and artificial intelligence (AI) research will be released as an open source solution.

Committed to doing its AI work in the open and sharing its findings with the community, Facebook intends to work with its partners to open source Big Sur specifications via the [Open Compute Project](#). This unique approach will make it easier for AI researchers worldwide to share and improve techniques, enabling future innovation in machine learning by harnessing the power of GPU accelerated computing.

Keep Current on NVIDIA

Subscribe to the [NVIDIA blog](#), follow us on [Facebook](#), [Google+](#), [Twitter](#), [LinkedIn](#) and [Instagram](#), and view NVIDIA videos on [YouTube](#) and images on [Flickr](#).

About NVIDIA

Since 1993, [NVIDIA](#) (NASDAQ: [NVDA](#)) has pioneered the art and science of [visual computing](#). The company's technologies are transforming a world of displays into a world of interactive discovery -- for everyone from gamers to scientists, and consumers to enterprise customers. More information at <http://nvidianews.nvidia.com/> and <http://blogs.nvidia.com/>.

Certain statements in this press release including, but not limited to, statements as to: Facebook's use of the NVIDIA Tesla Accelerated Computing Platform; and the impact and benefits of NVIDIA Tesla M40 GPUs and the Big Sur computing platform 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 25, 2015. 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.

© 2015 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo and Tesla are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Features, pricing, availability and specifications are subject to change without notice.

Ken Brown
Corporate Communications
+1-408-486-2626
kebrown@nvidia.com