


NVIDIA Reinvents Computer Graphics with Turing Architecture

Hybrid Rendering Fundamentally Changes Computer Graphics by Fusing Real-Time Ray Tracing, AI, Simulation and Rasterization

August 13, 2018



SIGGRAPH — NVIDIA today reinvented computer graphics with the launch of the [NVIDIA Turing™ GPU architecture](#).

The greatest leap since the invention of the CUDA GPU in 2006, Turing features new RT Cores to accelerate ray tracing and new Tensor Cores for AI inferencing which, together for the first time, make real-time ray tracing possible.

These two engines — along with more powerful compute for simulation and enhanced rasterization — usher in a new generation of hybrid rendering to address the \$250 billion visual effects industry. Hybrid rendering enables cinematic-quality interactive experiences, amazing new effects powered by neural networks and fluid interactivity on highly complex models.

The company also unveiled its [initial Turing-based products](#) — the NVIDIA® Quadro® RTX™ 8000, Quadro RTX 6000 and Quadro RTX 5000 GPUs — which will revolutionize the work of some 50 million designers and artists across multiple industries.

“Turing is NVIDIA’s most important innovation in computer graphics in more than a decade,” said Jensen Huang, founder and CEO of NVIDIA, speaking at the start of the annual SIGGRAPH conference. “Hybrid rendering will change the industry, opening up amazing possibilities that enhance our lives with more beautiful designs, richer entertainment and more interactive experiences. The arrival of real-time ray tracing is the holy grail of our industry.”

NVIDIA’s eighth-generation GPU architecture, Turing enables the world’s first ray-tracing GPU and is the result of more than 10,000 engineering-years of effort. By using Turing’s hybrid rendering capabilities, applications can simulate the physical world at 6x the speed of the previous Pascal™ generation.

To help developers take full advantage of these capabilities, NVIDIA has enhanced its [RTX development platform](#) with new AI, ray-tracing and simulation SDKs. It also announced that key graphics applications addressing millions of designers, artists and scientists are planning to take advantage of Turing features through the RTX development platform.

“This is a significant moment in the history of computer graphics,” said Jon Peddie, CEO of analyst firm JPR. “NVIDIA is delivering real-time ray tracing five years before we had thought possible.”

Real-Time Ray Tracing Accelerated by RT Cores

The Turing architecture is armed with dedicated ray-tracing processors called RT Cores, which accelerate the computation of how light and sound travel in 3D environments at up to 10 GigaRays a second. Turing accelerates real-time ray tracing operations by up to 25x that of the previous Pascal generation, and GPU nodes can be used for final-frame rendering for film effects at more than 30x the speed of CPU nodes.

“Cinesite is proud to partner with Autodesk and NVIDIA to bring Arnold to the GPU, but we never expected to see results this dramatic,” said Michele Sciolette, CTO of Cinesite. “This means we can iterate faster, more frequently and with higher quality settings. This will completely change how our artists work.”

AI Accelerated by Powerful Tensor Cores

The Turing architecture also features Tensor Cores, processors that accelerate deep learning training and inferencing, providing up to 500 trillion tensor operations a second.

This level of performance powers AI-enhanced features for creating applications with powerful new capabilities. These include DLAA — deep learning anti-aliasing, which is a breakthrough in high-quality motion image generation — denoising, resolution scaling and video re-timing.

These features are part of the NVIDIA NGX™ software development kit, a new deep learning-powered technology stack that enables developers to easily integrate accelerated,

enhanced graphics, photo imaging and video processing into applications with pre-trained networks.

Faster Simulation and Rasterization with New Turing Streaming Multiprocessor

Turing-based GPUs feature a new streaming multiprocessor (SM) architecture that adds an integer execution unit executing in parallel with the floating point datapath, and a new unified cache architecture with double the bandwidth of the previous generation.

Combined with new graphics technologies such as variable rate shading, the Turing SM achieves unprecedented levels of performance per core. With up to 4,608 CUDA cores, Turing supports up to 16 trillion floating point operations in parallel with 16 trillion integer operations per second.

Developers can take advantage of NVIDIA's CUDA 10, FleX and PhysX SDKs to create complex simulations, such as particles or fluid dynamics for scientific visualization, virtual environments and special effects.

Availability

Quadro GPUs based on Turing will be initially available in the fourth quarter.

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](#).

Media Contacts

Gail Laguna

Professional Visualization

NVIDIA Corp.


+1-408-386-2435

glaguna@nvidia.com

Downloads

1. [Download Press Release](#)
2. [Download Attachments](#)

More Images



More News




NVIDIA Stockholder Meeting Set for June 25; Individuals Can Participate Online

June 11, 2025



NVIDIA Partners With Novo Nordisk and DCAI to Advance Drug Discovery

June 11, 2025



NVIDIA Builds World's First Industrial AI Cloud to Advance European Manufacturing

June 11, 2025



Siemens and NVIDIA Expand Partnership to Accelerate AI Capabilities in Manufacturing

June 11, 2025



NVIDIA DGX Cloud Lepton Connects Europe's Developers to Global NVIDIA Compute Ecosystem

June 11, 2025

About NVIDIA

[NVIDIA](http://nvidianews.nvidia.com/)'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: NVIDIA reinventing and changing computer graphics with the launch of the Turing GPU architecture; the benefits, abilities, impact and performance of the Turing architecture, RT Cores, the RTX platform and Tensor Cores, including the ability to make and accelerate ray-tracing possible, enabling cinematic quality rendering, new creative tools and fluid interactivity, being NVIDIA's most important innovation, Turing resetting the industry and opening up creative possibilities, Turing making possible a new generation of applications that will simulate the physical world at unprecedented speeds, enabling partners to iterate faster and more frequently, completely changing how artists work, creating applications with powerful new capabilities, ability to create complex simulations; NVIDIA Quadro RTX products revolutionizing the work of designers and artists; key graphics applications planning to take advantage of Turing features through the RTX development platform; the

combination of technologies resulting in new results and the impact of the technology in every industry that does visualization with computers; the benefits, performance and abilities of the NVIDIA NGX software development kit; and the availability of Quadro GPUs 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.

© 2018 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, CUDA, NVIDIA NGX, NVIDIA RTX, NVIDIA Turing, Pascal, Quadro and Quadro RTX 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.