EMBARGOED for release until August 13 at 6:00pm PT or when Jensen announces news

NVIDIA Unveils Quadro RTX, World's First Ray-Tracing GPU

New Turing-Based Design Revolutionizes Workflow of Millions of Designers and Artists on the Desktop and in the Datacenter

VANCOUVER, British Columbia—SIGGRAPH—Aug. 13, 2018—NVIDIA today announced its first Turing™ architecture-based GPUs, revolutionizing the work of 50 million designers and artists by enabling them to render photorealistic scenes in real time, add new AI-based capabilities to their workflows, and enjoy fluid interactivity with complex models and scenes.

Unveiled by NVIDIA founder and CEO Jensen Huang at the annual SIGGRAPH conference, the NVIDIA® Quadro® RTX™ 8000, Quadro RTX 6000 and Quadro RTX 5000 bring hardware-accelerated ray tracing, AI, advanced shading and simulation to creative professionals. Also announced was the Quadro RTX Server, a reference architecture for highly configurable, on-demand rendering and virtual workstation solutions from the datacenter.

"Quadro RTX marks the launch of a new era for the global computer graphics industry," said Bob Pette, vice president of Professional Visualization at NVIDIA. "Users can now enjoy powerful capabilities that weren't expected to be available for at least five more years. Designers and artists can interact in real time with their complex designs and visual effects in ray-traced photo-realistic detail. And film studios and production houses can now realize increased throughput with their rendering workloads, leading to significant time and cost savings."

Quadro RTX Professional GPUs

Quadro RTX GPUs are designed for the most demanding visual computing workloads, such as those used in film and video content creation; automotive and architectural design; and

scientific visualization. They far surpass the previous generation with groundbreaking technologies, including:

- New RT Cores to enable real-time ray tracing of objects and environments with physically accurate shadows, reflections, refractions and global illumination.
- Turing Tensor Cores to accelerate deep neural network training and inference, which are critical to powering AI-enhanced rendering, products and services.
- New Turing Streaming Multiprocessor architecture, featuring up to 4,608 CUDA® cores, delivers up to 16 trillion floating point operations in parallel with 16 trillion integer operations per second to accelerate complex simulation of real-world physics.
- Advanced programmable shading technologies to improve the performance of complex visual effects and graphics-intensive experiences.
- First implementation of ultra-fast Samsung 16Gb GDDR6 memory to support more complex designs, massive architectural datasets, 8K movie content and more.
- NVIDIA NVLink® to combine two GPUs with a high-speed link to scale memory capacity up to 96GB and drive higher performance with up to 100GB/s of data transfer.
- Hardware support for USB Type-C™ and VirtualLink™(1), a new open industry standard being developed to meet the power, display and bandwidth demands of next-generation VR headsets through a single USB-C™ connector.
- New and enhanced technologies to improve performance of VR applications, including Variable Rate Shading, Multi-View Rendering and VRWorks Audio.

Key attributes of the new GPUs:

GPU	Memory	Memory with NVLink	Ray Tracing	CUDA Cores	Tensor Cores
Quadro RTX 8000	48GB	96GB	10 GigaRays/sec	4,608	576
Quadro RTX 6000	24GB	48GB	10 GigaRays/sec	4,608	576
Quadro RTX 5000	16GB	32GB	6 GigaRays/sec	3,072	384

Quadro RTX Server

The Quadro RTX Server defines a new standard for on-demand rendering in the datacenter, enabling easy configuration of on-demand render nodes for batch and interactive rendering.

It combines Quadro RTX GPUs with new Quadro Infinity software (available in the first quarter of 2019) to deliver a powerful and flexible architecture to meet the demands of creative professionals. Quadro Infinity will enable multiple users to access a single GPU through virtual workstations, dramatically increasing the density of the datacenter. End-users can also easily provision render nodes and workstations based on their specific needs.

With industry-leading content creation and render software pre-installed, the Quadro RTX Server provides a powerful and easy-to-deploy rendering solution that can scale from small installations to the largest data centers, at one quarter of the cost of CPU-only render farms.

Industry Support for Quadro RTX

The makers of the world's most widely used design and creative applications are already working closely with NVIDIA to bring the power of Quadro RTX to customers.

Developers can access these Quadro RTX features through the new NVIDIA RTX™, a graphics platform that includes APIs for ray tracing, AI, rasterization, and simulation plus support for NVIDIA MDL materials and Pixar USD asset interchange to transform the creative process. Initially supported by 30 ISV applications and addressing more than 50 million users, RTX can be easily accessed by professionals across industries through a range of creative applications and tools.

"AI and real-time ray tracing are transforming the way designers and artists work," said Amy Bunszel, senior vice president of Design & Creation Products at Autodesk. "Autodesk and NVIDIA are working together to deliver these advancements to the market and we look forward to the new tools and capabilities Quadro RTX will deliver to our customers."

Leading workstation and systems providers have also voiced their support for new Turing-based Quadro RTX GPUs:

- "Dell Precision workstations strive to bring the latest in technology to creators, engineers and data scientists to enable workflows previously thought impossible. We're proud to support NVIDIA Quadro RTX 6000 and 5000 GPUs in Dell Precision tower workstations to enable more immersive workflows, inferencing, training and hyperrealistic visualization by leveraging the power of physically based rendering and real-time intelligence using Tensor Cores. We will offer these options on the Dell Precision 5820, 7820 and 7920 workstations over the coming months," said Rahul Tikoo, vice president and general manager of Commercial Specialty Products at Dell.

- "With the high-performing graphics architecture of Turing-based Quadro RTX GPUs for enterprise datacenter workloads, we are unlocking extraordinary capabilities for our customers to power new creative experiences," said Bill Mannel, vice president and general manager of HPC and AI Solutions at Hewlett Packard Enterprise.

- "In our technology proof of concept, Turing-based Quadro RTX GPUs, coupled with the HP Z8 Workstation – the most powerful desktop workstation in the world – showed that this solution can unlock new levels of performance. AI-based capabilities can allow creators and developers to simulate and interact with their creations in ways not currently possible," said Xavier Garcia, vice president and general manager of HP Z Workstations at HP Inc.

- "Turing is about to turn the graphics industry on its head. The groundbreaking AI and real-time ray-tracing capabilities of Turing-based Quadro RTX GPUs, combined with the sheer horsepower of Lenovo's ThinkStation P920 or P720, will advance the creative process faster than anything our customers have experienced ever before," said Rob Herman, general manager of the Lenovo Workstation & Client AI Group at Lenovo.

- "The ability to support NVIDIA Turing-based Quadro RTX GPUs in select models of Lenovo ThinkSystem servers brings the advancement of RTX technology to those users needing to scale their rendering in leading-edge datacenter implementations. RTX provides orders of magnitude faster rendering over a traditional render farm

deployment, unleashing the creative minds of media and entertainment professionals," said Scott Tease, executive director of HPC and AI at Lenovo Data Center Group.

Pricing

Quadro RTX 8000 with 48GB memory: $10,000 estimated street price

Quadro RTX 6000 with 24GB memory: $6,300 ESP

Quadro RTX 5000 with 16GB memory: $2,300 ESP

Availability

Quadro RTX GPUs will be available starting in the fourth quarter on nvidia.com. For shipment dates, contact the world's leading OEM workstation manufacturers, including Dell EMC, HPE, HPI and Lenovo, and system builders and authorized distribution partners, including PNY Technologies in North America and Europe, ELSA/Ryoyo in Japan and Leadtek and Ingram in Asia Pacific.

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

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/.

# # #

For further information, contact:
Gail Laguna
Sr. PR Manager, Professional Visualization
NVIDIA Corporation
+1-408-386-2435
glaguna@nvidia.com

(1) In preparation for the emerging VirtualLink standard, Turing GPUs have implemented hardware support according to the "VirtualLink Advance Overview." To learn more about VirtualLink, see www.virtuallink.org.

Certain statements in this press release including, but not limited to, statements as to: the benefits, performance, impact, abilities and availability of Quadro RTX GPUs, including their ability to revolutionize the work of millions of designers and artists, RT Cores enabling real-time ray tracing and environments, Turing Tensor Cores accelerating deep neural network training and inference, advanced programmable shading technologies improving the performance of complex visual effects and graphics-intensive workloads, compute performance acceleration, ability to support more complex designs, massive architectural datasets, 8K movie content, the technology improving performance of VR applications, unlocking extraordinary capabilities to power new creative experiences, enabling new levels of GPU performance, advancing the creative process faster than customers have experienced before, bringing RTX technology to those users needing to scale their rendering, and RTX providing faster rendering and unleashing the creative minds of media and entertainment professionals; Quadro RTX marking the launch of a new era for the global computer graphics industry and enabling users to enjoy capabilities that weren't expected for more than five years; Quadro RTX enabling designers and artists to interact in real time with their designs and visual effects and film studios and production houses' ability to realize increased throughput with rendering workloads, leading to time and cost savings; the benefits and abilities of NVIDIA NVLink; the development and benefits of USB Type-C and VirtualLink; the availability of Quadro Infinity software; the benefits, performance, impact and abilities of the Quadro RTX Server, Quadro Infinity software and NVIDIA RTX; the makers of design and creative applications working with NVIDIA to bring Quadro RTX to customers; NVIDIA RTX's accessibility and use by professionals; AI and real-time ray tracing transforming the way designers and artists work; Autodesk and NVIDIA working together to deliver advancements to the market and the benefits Quadro RTX will deliver to customers; Dell Precision workstations striving to bring the latest technology and enable workflows previously thought impossible; the benefits, abilities and availability of Dell Precision tower workstations using NVIDIA Quadro GPUs; AI-based capabilities allowing creators and developers to simulate and interact with their creations in ways not currently possible; and Turing turning the graphics industry on its head 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.