



NVIDIA Delivers Quantum Leap in Performance, Introduces New Era of Neural Rendering With GeForce RTX 40 Series

Powered by Ada Lovelace Architecture and DLSS 3; Third-Gen RTX up to 4x Faster Than NVIDIA Ampere Architecture GPUs

NVIDIA today unveiled the [GeForce RTX® 40 Series of GPUs](#), designed to deliver revolutionary performance for gamers and creators, led by its new flagship, the RTX 4090 GPU, with up to 4x the performance of its predecessor.

The world's first GPUs based on the new NVIDIA® Ada Lovelace architecture, the RTX 40 Series delivers massive generational leaps in performance and efficiency, and represents [a new era of real-time ray tracing and neural rendering](#), which uses AI to generate pixels.

"The age of RTX ray tracing and neural rendering is in full steam, and our new Ada Lovelace architecture takes it to the next level," said Jensen Huang, NVIDIA's founder and CEO, at the GeForce® Beyond: Special Broadcast at GTC.

"Ada provides a quantum leap for gamers and paves the way for creators of fully simulated worlds. With up to 4x the performance of the previous generation, Ada is setting a new standard for the industry," he said.

DLSS 3 Generates Entire Frames for Faster Game Play

Huang also announced [NVIDIA DLSS 3](#) — the next revolution in the company's Deep Learning Super Sampling neural-graphics technology for games and creative apps. The AI-powered technology can generate entire frames for massively faster game play. It can overcome CPU performance limitations in games by allowing the GPU to generate entire frames independently.

The technology is coming to the world's most popular game engines, such as Unity and Unreal Engine, and has received support from many of the world's leading game developers, with [more than 35 games and apps coming soon](#).

Additionally, the RTX 40 Series GPUs feature a range of new technological innovations, including:

- **Streaming multiprocessors** with up to 83 teraflops of shader power — 2x over the previous generation.
- **Third-generation RT Cores** with up to 191 effective ray-tracing teraflops — 2.8x over the previous generation.
- **Fourth-generation Tensor Cores** with up to 1.32 Tensor petaflops — 5x over the previous generation using FP8 acceleration.
- **Shader Execution Reordering (SER)** that improves execution efficiency by rescheduling shading workloads on the fly to better utilize the GPU's resources. As significant an innovation as out-of-order execution was for CPUs, SER improves ray-tracing performance up to 3x and in-game frame rates by up to 25%.
- **Ada Optical Flow Accelerator** with 2x faster performance allows DLSS 3 to predict movement in a scene, enabling the neural network to boost frame rates while maintaining image quality.
- **Architectural improvements** tightly coupled with custom TSMC 4N process technology results in an up to 2x leap in power efficiency.
- **Dual NVIDIA Encoders (NVENC)** cut export times by up to half and feature AV1 support. The NVENC AV1 encode is being adopted by OBS, Blackmagic Design DaVinci Resolve, Discord and more.

New Ray-Tracing Tech for Even More Immersive Games

For decades, rendering ray-traced scenes with physically correct lighting in real time has been considered the holy grail of graphics. At the same time, geometric complexity of environments and objects has continued to increase as 3D games and graphics strive to provide the most accurate representations of the real world.

Achieving physically accurate graphics requires tremendous computational horsepower. Modern ray-traced games like [Cyberpunk 2077](#) run over 600 ray-tracing calculations for each pixel just to determine lighting — a 16x increase from the first ray-traced games introduced four years ago.

The new third-generation RT Cores have been enhanced to deliver 2x faster ray-triangle intersection testing and include two important new hardware units. An Opacity Micromap Engine speeds up ray tracing of alpha-test geometry by a factor of 2x, and a Micro-Mesh Engine generates micro-meshes on the fly to generate additional geometry. The Micro-Mesh Engine provides the benefits of increased geometric complexity without the traditional performance and storage costs of complex geometries.

Creativity Redefined With RTX Remix, New AV1 Encoders

The RTX 40 Series GPUs and DLSS 3 deliver advancements for [NVIDIA Studio](#) creators. 3D artists can render fully ray-traced environments with accurate physics and realistic materials, and view the changes in real time, without proxies.

Video editing and live streaming also get a boost from improved GPU performance and the inclusion of new dual, eighth-generation AV1 encoders. The [NVIDIA Broadcast](#) software development kit has three updates, now available for partners, including Face Expression Estimation, Eye Contact and quality improvements to Virtual Background.

[NVIDIA Omniverse™](#) — included in the NVIDIA Studio suite of software — will soon add [NVIDIA RTX Remix](#), a modding platform to create stunning RTX remasters of classic games. RTX Remix allows modders to easily capture game assets, automatically enhance materials with powerful AI tools, and quickly enable RTX with ray tracing and DLSS.

Portal Is RTX ON!

RTX Remix has been used by NVIDIA Lightspeed Studios to reimagine Valve's iconic video game [Portal](#), regarded as one of the best video games of all time. Advanced graphics features such as full ray tracing and DLSS 3 give the game a striking new look and feel. *Portal with RTX* will be released as free, official downloadable content for the classic platformer with RTX graphics in November, just in time for *Portal's* 15th anniversary.

The GeForce RTX 4090 and 4080: The New Ultimate GPUs

The RTX 4090 is the world's fastest gaming GPU with astonishing power, acoustics and temperature characteristics. In full ray-traced games, the RTX 4090 with DLSS 3 is up to 4x faster compared to last generation's RTX 3090 Ti with DLSS 2. It is also up to 2x faster in today's games while maintaining the same 450W power consumption. It features 76 billion transistors, 16,384 CUDA® cores and 24GB of high-speed Micron GDDR6X memory, and consistently delivers over 100 frames per second at 4K-resolution gaming. The RTX 4090 will be available on Wednesday, Oct. 12, starting at \$1,599.

The company also announced the RTX 4080, launching in two configurations. The RTX 4080 16GB has 9,728 CUDA cores and 16GB of high-speed Micron GDDR6X memory, and with DLSS 3 is 2x as fast in today's games as the GeForce RTX 3080 Ti and more powerful than the GeForce RTX 3090 Ti at lower power. The RTX 4080 12GB has 7,680 CUDA cores and 12GB of Micron GDDR6X memory, and with DLSS 3 is faster than the RTX 3090 Ti, the previous-generation flagship GPU.

Both RTX 4080 configurations will be available in November, with prices starting at \$1,199 and \$899, respectively.

Where to Buy

The GeForce RTX 4090 and 4080 GPUs will be available as custom boards, including stock-clocked and factory-overclocked models, from top add-in card providers such as ASUS, Colorful, Gainward, Galaxy, GIGABYTE, Inno3D, MSI, Palit, PNY and Zotac.

The RTX 4090 and RTX 4080 (16GB) are also produced directly by NVIDIA in limited Founders Editions for fans wanting the NVIDIA in-house design.

Look for the GeForce RTX 40 Series GPUs in gaming systems built by Acer, Alienware, ASUS, Dell, HP, Lenovo and MSI, leading system builders worldwide, and many more.

About NVIDIA

Since its founding in 1993, [NVIDIA](#) (NASDAQ: NVDA) has been a pioneer in accelerated computing. The company's invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined computer graphics and ignited the era of modern AI. NVIDIA is now a full-stack computing company with data-center-scale offerings that are reshaping industry. More information at <https://nvidianews.nvidia.com/>.

Certain statements in this press release including, but not limited to, statements as to: the benefits, impact, specifications, performance, features and availability of our products and technologies, including GeForce RTX 40 Series, RTX 4090 GPU, NVIDIA Ada Lovelace architecture, NVIDIA DLSS 3, RT Cores, Opacity Micromap Engine, Micro-Mesh Engine, NVIDIA Studio suite of software, NVIDIA Broadcast software development kit, NVIDIA Omniverse, NVIDIA RTX Remix and RTX 4080 GPU; Ada Lovelace architecture providing a quantum leap for gamers and paving the way for creators of fully simulated worlds; the game-changing innovations of Ada Lovelace architecture setting a powerful new standard for the industry; a range of the world's popular game engines incorporating DLSS 3 and leading game developers supporting DLSS 3; and the availability of *Portal with RTX* 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.

© 2022 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, CUDA, GeForce, GeForce RTX, NVIDIA Omniverse, and NVIDIA RTX are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. All other trademarks and copyrights are the property of their respective owners. Features, pricing, availability and specifications are subject to change without notice.

Bryan Del Rizzo

GeForce Desktops and Notebooks, eSports

NVIDIA Corp.

+1-408-486-2772

bdelrizzo@nvidia.com