




# NVIDIA RTX 4000 Ada Generation

Performance for endless possibilities.



## Powering the Next Era of Innovation


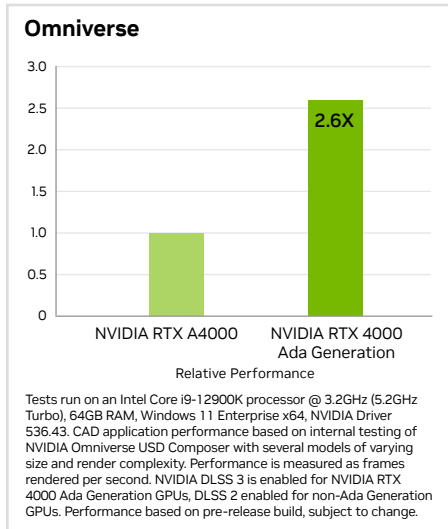
Industries are embracing accelerated computing and AI to tackle powerful dynamics and unlock transformative possibilities. Generative AI is reshaping the way professionals create and innovate across various domains, from design and engineering to entertainment and healthcare.

The NVIDIA RTX™ 4000 Ada Generation is the most powerful single-slot GPU for professionals, delivering remarkable acceleration for AI, real-time rendering, graphics, and compute workloads to your desktop. Built on the NVIDIA Ada Lovelace architecture, RTX 4000 combines 48 third-generation RT Cores, 192 fourth-generation Tensor Cores, and 6,144 CUDA® cores with 20GB of graphics memory to effortlessly handle large datasets and complex visual workloads.


NVIDIA RTX professional graphics cards are certified with a broad range of professional applications, tested by leading independent software vendors (ISVs) and workstation manufacturers, and backed by a global team of support specialists. Get the peace of mind to focus on what matters with the premier visual computing solution for mission-critical business.

## Key Features

- > PCIe Gen4
- > Four DisplayPort 1.4a connectors
- > AV1 encode and decode support
- > DisplayPort with audio
- > 3D stereo support with stereo connector
- > NVIDIA® GPUDirect® for Video support
- > NVIDIA GPUDirect remote direct memory access (RDMA) support
- > NVIDIA Quadro® Sync II<sup>1</sup> compatibility
- > NVIDIA RTX Experience
- > NVIDIA RTX Desktop Manager software
- > NVIDIA RTX IO support
- > HDCP 2.2 support
- > NVIDIA Mosaic<sup>2</sup> technology




## Graphics




Tests run on an Intel Core i9-12900K processor @ 3.2GHz (5.2GHz Turbo), 64GB RAM, Windows 11 Enterprise x64, NVIDIA Driver 528.49. Relative speedup for 4K geomean score. Performance based on pre-release build, subject to change.

## Generative AI



Tests run on an Intel Core i9-12900K processor @ 3.2GHz (5.2GHz Turbo), 64GB RAM, Windows 11 Enterprise x64, Stable Diffusion WebUI v1.3.1, NVIDIA Driver 536.15. Relative speedup for 512x512 image generation. Performance based on pre-released build, subject to change.

## Inference



Tests run on an Intel Core i9-12900K processor @ 3.2GHz (5.2GHz Turbo), 64GB RAM, Windows 11 Enterprise x64, TensorRT v8.5.1, NVIDIA Driver 525.85. Relative speedup for ResNet-50 V1.5 Inference, mixed precision. Performance based on pre-release build, subject to change.

## Specifications

<b>GPU Memory</b>	20GB GDDR6
<b>Memory Interface</b>	160 bit
<b>Memory Bandwidth</b>	360GB/s
<b>Error Correcting Code (ECC)</b>	Yes
<b>NVIDIA Ada Lovelace Architecture-Based CUDA Cores</b>	6,144
<b>NVIDIA Fourth-Generation Tensor Cores</b>	192
<b>NVIDIA Third-Generation RT Cores</b>	48
<b>Single-Precision Performance</b>	26.7 TFLOPS <sup>3</sup>
<b>RT Core Performance</b>	61.8 TFLOPS <sup>3</sup>
<b>Tensor Performance</b>	327.6 TFLOPS <sup>4</sup>
<b>System Interface</b>	PCIe 4.0 x16
<b>Power Consumption</b>	Total board power: 130W
<b>Thermal Solution</b>	Active
<b>Form Factor</b>	4.4" H x 9.5" L, single slot
<b>Display Connectors</b>	4x DisplayPort 1.4a <sup>5</sup>
<b>Max Simultaneous Displays</b>	4x 4096 x 2160 @ 120Hz 4x 5120 x 2880 @ 60Hz 2x 7680 x 4320 @ 60Hz
<b>Encode/Decode Engines</b>	2x encode, 2x decode (+AV1 encode and decode)
<b>VR Ready</b>	Yes
<b>Graphics APIs</b>	DirectX 12, Shader Model 6.7, OpenGL 4.6 <sup>6</sup> , Vulkan 1.3 <sup>6</sup>
<b>Compute APIs</b>	CUDA 12.2, OpenCL 3.0, DirectCompute
<b>NVIDIA NVLink<sup>®</sup></b>	No

## Ready to Get Started?

To learn more about NVIDIA RTX 4000, visit:  
[www.nvidia.com/rtx-4000](http://www.nvidia.com/rtx-4000)

1 Quadro Sync II card sold separately. | 2 Windows 10 and Linux. | 3 Peak rates based on GPU boost clock. | 4 Effective FP8 teraFLOPS (TFLOPS) using sparsity. | 5 Display ports are on by default for RTX 4000. | 6 Product is based on a published Khronos specification and is expected to pass the Khronos conformance testing process when available. Current conformance status can be found at [www.khronos.org/conformance](http://www.khronos.org/conformance)

© 2023 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, CUDA, GPUDirect, NVLink, Quadro, and 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. All other trademarks are the property of their respective owners. 2895141. SEP23

