




NVIDIA RTX 6000 Ada Generation

Performance for endless possibilities.



Powering the Next Era of Innovation

The NVIDIA RTX™ 6000 Ada Generation is the ultimate workstation graphics card designed for professionals who demand maximum performance and reliability to deliver their best work and breakthrough innovations across industries. The RTX 6000 provides the unmatched performance and capabilities essential for high-end design, real-time rendering, AI, and high-performance compute workflows.

Built on the NVIDIA Ada Lovelace architecture, the RTX 6000 combines 142 third-generation RT Cores, 568 fourth-generation Tensor Cores, and 18,176 CUDA® cores with 48GB of error correction code (ECC) graphics memory. This all helps deliver the next generation of AI graphics and petaflop inferencing performance for unprecedented speed-up in rendering, AI, graphics, and compute 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 Gen 4
- > 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 virtual GPU (vGPU) software support
- > NVIDIA Quadro® Sync II¹ compatibility
- > NVIDIA RTX Experience™
- > NVIDIA RTX Desktop Manager software
- > NVIDIA RTX IO support
- > HDCP 2.2 support
- > NVIDIA Mosaic² technology

Technical Specifications

GPU memory	48GB GDDR6
Memory interface	384-bit
Memory bandwidth	960 GB/s
Error correction code (ECC)	Yes
NVIDIA Ada Lovelace architecture-based CUDA Cores	18,176
NVIDIA fourth-generation Tensor Cores	568
NVIDIA third-generation RT Cores	142
Single-precision performance	91.1 TFLOPS ³
RT Core performance	210.6 TFLOPS ³
Tensor performance	1457.0 TFLOPS ⁴
System interface	PCIe 4.0 x16
Power consumption	Total board power: 300 W
Thermal solution	Active
Form factor	4.4" H x 10.5" L, dual slot, full height
Display connectors	4x DisplayPort 1.4a ⁵
Max simultaneous displays	<ul style="list-style-type: none">> 4x 4096 x 2160 @ 120 Hz> 4x 5120 x 2880 @ 60 Hz> 2x 7680 x 4320 @ 60 Hz
Power connector	1x PCIe CEM5 16-pin
Encode/decode engines	3x encode, 3x decode (+AV1 encode and decode)
VR ready	Yes
vGPU software support ⁵	<ul style="list-style-type: none">> NVIDIA vPC/vApps> NVIDIA RTX Virtual Workstation
vGPU profiles supported	See the Virtual GPU Licensing Guide
Graphics APIs	DirectX 12, Shader Model 6.6, OpenGL 4.6 ⁶ , Vulkan 1.3 ⁶
Compute APIs	CUDA 11.6, OpenCL 3.0, DirectCompute
NVIDIA NVLink®	No

Ready to Get Started?

To learn more RTX 6000, visit www.nvidia.com/rtx-6000

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 the new sparsity feature. ⁵ Display ports are on by default for RTX 6000. Display ports are not active when using vGPU software. 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

© 2023 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, CUDA, GPUDirect, NVLink, Quadro, RTX Experience, 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. 2647623. FEB23

