




NVIDIA RTX 4000 SFF Ada Generation

Power for endless possibilities.



The NVIDIA RTX™ 4000 SFF Ada Generation packs a powerful punch, delivering full-size performance in a compact form factor. Designed for professionals who demand performance but without the footprint of full-sized workstations to achieve outstanding results across industries. The RTX 4000 SFF provides excellent performance and capabilities essential for design, real-time rendering, AI, and high-performance compute workflows in a small form factor.


Built on the NVIDIA Ada Lovelace architecture, the RTX 4000 SFF combines 48 third-generation RT Cores, 192 fourth-generation Tensor Cores, and 6,144 CUDA® cores with 20GB of error correction code (ECC) graphics memory. The RTX 4000 SFF delivers incredible acceleration for 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


- > Four mini DisplayPort 1.4a
- > AV1 encode and decode support
- > DisplayPort with audio
- > NVIDIA Quadro® Sync II¹ compatibility
- > 3D stereo support with stereo connector
- > NVIDIA GPUDirect® for Video support
- > NVIDIA GPUDirect Remote Direct Memory Access (RDMA) support
- > NVIDIA RTX Experience™
- > NVIDIA RTX Desktop Manager software
- > NVIDIA RTX IO support
- > HDCP 2.2 support
- > NVIDIA Mosaic² technology

Graphics




Tests run on an Intel® Core™ i9-12900K Processor @ 3.20GHz (5.20GHz Turbo), 64GB RAM, Windows 11 Enterprise x64, SPECviewperf 2020, NVIDIA Driver 528.35. Relative speedup for 4K energy composite score.

CAD




Tests run on an Intel® Core™ i9-12900K Processor @ 3.20GHz (5.20GHz Turbo), 64GB RAM, Windows 11 Enterprise x64, SPECviewperf 2020, CATIA v6, NVIDIA Driver 528.35. Relative speedup for 4K Siemens NX composite score.

Rendering




Tests run on an Intel® Core™ i9-12900K Processor @ 3.20GHz (5.20GHz Turbo), 64GB RAM, Windows 11 Enterprise x64, NVIDIA Driver 528.35. Relative speedup of Arnold render test score.

AI Inference



Tests run on an Intel® Core™ i9-12900K Processor @ 3.20GHz (5.20GHz Turbo), 64GB RAM, Windows 11 Enterprise x64, TensorRT 8.5, NVIDIA Driver 525.85. Relative speedup of ResNet-50 Inference mixed precision test score.

High-Performance Computing



Tests run on an Intel® Core™ i9-12900K Processor @ 3.20GHz (5.20GHz Turbo), 64GB RAM, Windows 11 Enterprise x64, LAMMPS, NVIDIA Driver 525.85. Relative speedup of Atomic Fluid Lennard-Jones 2.5 FP64 precision test score.

Specifications

GPU memory	20GB GDDR6
Memory interface	160-bit
Memory bandwidth	280 GB/s
Error correcting code	Yes
NVIDIA Ada Lovelace architecture-based CUDA Cores	6,144
NVIDIA fourth-generation Tensor Cores	192
NVIDIA third-generation RT Cores	48
Single-precision performance	19.2 TFLOPS ³
RT Core performance	44.3 TFLOPS ³
Tensor performance	306.8 TFLOPS ⁴
System interface	PCIe 4.0 x 16
Power consumption	Total board power: 70 W
Thermal solution	Active
Form factor	2.7" H x 6.6" L, dual slot
Display connectors	4x Mini DisplayPort 1.4a
Maximum 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
Encode/decode engines	2x encode, 2x decode (+AV1 encode and decode)
VR ready	Yes
vGPU software support	No
Graphics APIs	DirectX 12, Shader Model 6.6, OpenGL 4.6 ⁵ , Vulkan 1.3 ⁵
Compute APIs	CUDA 11.6, OpenCL 3.0, DirectCompute
NVIDIA NVLink®	No

Ready to Get Started?

To learn more about the RTX 4000 SFF, visit:

[nvidia.com/rtx-4000-sff](https://www.nvidia.com/rtx-4000-sff)

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 sparsity feature. | 5 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 and affiliates. All rights reserved. NVIDIA, the NVIDIA logo, CUDA, NVIDIA GPUDirect, NVIDIA NVLink, NVIDIA Quadro, NVIDIA RTX and NVIDIA RTX Experience are trademarks and/or registered trademarks of NVIDIA Corporation and affiliates in the U.S. and other countries. All other trademarks and copyrights are the property of their respective owners. 2742450. APR23

