Datasheet

NVIDIA RTX 2000E Ada Generation

Performance for endless possibilities.



Powering the Next Era of Innovation

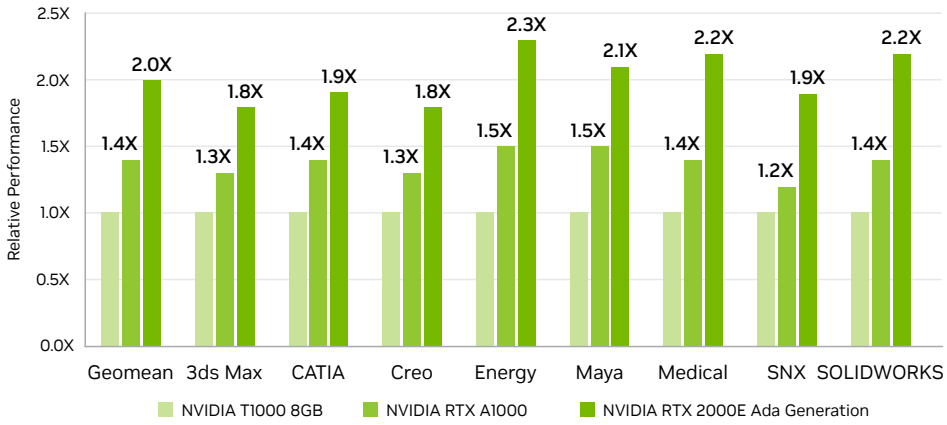
The NVIDIA RTX™ 2000E Ada Generation is a powerful, low-profile GPU purpose-built for demanding embedded and edge computing applications. Powered by the NVIDIA Ada Lovelace GPU architecture, it packs 2,816 CUDA® cores, 88 fourth-generation Tensor Cores, 22 third-generation RT cores, and 16GB of GDDR6 memory with ECC support into a compact, single-slot form factor. The RTX 2000E delivers incredible performance and power efficiency for accelerating graphics, AI and compute workloads, while meeting the highest quality and reliability standards. From powering medical devices and advanced digital displays to enabling cutting-edge robotics, the RTX 2000E brings the power of real-time ray tracing, accelerated computing, and accelerated AI capabilities to a wide range of embedded and edge systems.

NVIDIA RTX professional graphics cards are certified for 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 RTX Experience™
- NVIDIA RTX Desktop Manager software
- NVIDIA RTX IO support
- HDCP 2.2 support
- NVIDIA Mosaic[1] technology

Graphics



Tests run on an Intel Core i9-12900K Processor @ 3.2GHz (5.2GHz Turbo), 64GB RAM, Windows 11 Enterprise x64, SPECviewperf 2020, NVIDIA Driver 551.57. Relative speedup for 4K geomean score, 3ds Max, CATIA, Creo, energy, Maya, medical, SNX, and SOLIDWORKS composite scores. Performance may vary based on scene, application, and system configuration

Specifications

GPU memory	16GB GDDR6
Memory interface	128-bit
Memory bandwidth	224 GB/s
Error-correcting code (ECC)	Yes
NVIDIA Ada Lovelace architecture-based CUDA Cores	2,816
NVIDIA fourth-generation Tensor Cores	88
NVIDIA third-generation RT Cores	22
Single-precision performance	8.9 TFLOPS[2]
Ray Tracing performance	20.5 TFLOPS[2]
Tensor (AI) performance	142 AI TOPS | 142 TFLOPS[3]
System interface	PCIe 4.0 x 8[4]
Power consumption	Total board power: 50 W
Thermal solution	Active
Form factor	2.7” H x 6.6” L, single slot
Display connectors	4x Mini DisplayPort 1.4a
Max simultaneous displays	4x 4096 x 2160 @ 120 Hz 4x 5120 x 2880 @ 60 Hz 2x 7680 x 4320 @ 60 Hz
Encode/decode engines	1x encode, 1x decode (+AV1 encode and decode)
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

Ready to Get Started?

To find a partner, visit
www.nvidia.com/rtx-partners

1 Windows 10 and Linux. I 2 Peak rates based on GPU Boost Clock. I 3 Theoretical FP8 TOPS and TFLOPS using the sparsity feature. I 4 RTX 2000E Ada Generation utilizes a full-length PCIe Gen 4 x8 interface. | 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

© 2025 NVIDIA Corporation and affiliates. All rights reserved. NVIDIA, the NVIDIA logo, CUDA, RTX, and 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. 4328932. OCT25

NVIDIA