




NVIDIA RTX PRO 4500 Blackwell Workstation Edition

Powering the next era of AI.



Transform Workflows With Next-Level Workstation Performance

As AI continues to advance at an incredible pace, industries face mounting pressure to harness its transformative power and adopt tools capable of handling generative AI, real-time simulation, and hyper-realistic rendering. Enterprises need solutions that combine breakthrough performance, scalability, and versatility to tackle the rise of increasingly complex workloads—from training domain-specific AI models to rendering billion-polygon engineering designs or simulating real-world physics with higher fidelity and precision.

The NVIDIA RTX PRO™ 4500 Blackwell Workstation Edition redefines professional workflows with groundbreaking AI acceleration, neural rendering, and 32GB of ultra-fast memory. Built on the revolutionary NVIDIA Blackwell architecture, it features 5th Gen Tensor Cores, 4th Gen RT Cores, and advanced CUDA® cores to power next-generation creative, engineering, and scientific applications. Professionals can securely deploy autonomous, agentic AI systems using NVIDIA NemoClaw, run physics simulations, and render cinematic-quality 3D worlds—all while handling massive datasets and complex workloads with ease. Whether designing sustainable infrastructure or driving generative AI pipelines, the RTX PRO 4500 delivers enterprise-grade reliability and future-proof performance for tackling the most demanding challenges.

Breakthrough Innovations

The NVIDIA Blackwell architecture combines breakthrough AI, ray tracing, and neural rendering technology, with massive performance and memory improvements to drive cutting-edge professional creative, design, and engineering workflows and power the next decade of innovation.

NVIDIA Blackwell Streaming Multiprocessor: The new SM features increased processing throughput, and new neural shaders that integrate neural networks inside of programmable shaders to drive the next decade of AI-augmented graphics innovations.

5th Gen Tensor Cores: Deliver up to 3x the performance of the previous generation and support for FP4 precision for faster AI model processing times with reduced memory usage, enabling local fine-tuning of LLMs and generative AI.

Key Features

- Enhanced streaming multiprocessors (SMs) built for neural shaders
- 5th Gen Tensor Cores support FP4 precision, DLSS 4 Multi Frame Generation
- 4th Gen Ray Tracing Cores built for detailed geometry
- 32 GB of GDDR7 memory
- 896 GB/s of memory bandwidth
- 9th Gen NVENC and 6th Gen NVDEC with 4:2:2 support
- PCIe Gen 5
- Four Display Port 2.1b connectors
- AI management processor

4th Gen Ray Tracing Cores: Double the ray-triangle intersection rate of the previous generation to create photoreal, physically accurate scenes and immersive 3D designs with RTX Mega Geometry, which enables up to 100X more ray-traced triangles.

Next-Gen Video Engines: Enhance video conferencing, video production, and streaming workflows with real-time AI processing. Ninth-generation NVENC and sixth-generation NVDEC engines provide support for 4:2:2 encoding and decoding to explore a new realm of high-resolution video workflows.

GDDR7 Memory: New and improved GDDR7 memory significantly boosts bandwidth and capacity, empowering your applications to run faster and work with larger, more complex datasets. With 32 GB of GPU memory and 896 GB/s bandwidth, tackle massive 3D and AI projects, fine-tune AI models locally, explore large-scale VR environments, and drive larger multi-app workflows.

DLSS 4: Multi Frame Generation ensures ultra-smooth frame pacing for lifelike simulations. Experience up to 3x faster frame rates and stunning image quality in supported game engines and 3D rendering applications for smoother, more responsive performance.

PCIe Gen 5: Support for PCIe Gen 5 provides double the bandwidth of PCIe Gen 4, improving data-transfer speeds from CPU memory and unlocking faster performance for data-intensive tasks like AI, data science, and 3D modeling.

DisplayPort 2.1b: Achieve unparalleled visual clarity and performance, driving high-resolution displays at up to 8K at 240 Hz and 16K at 60 Hz. Increased bandwidth enables seamless multi-monitor setups, ideal for multitasking and collaboration, while HDR and higher color depth support ensures superior color accuracy for precision work, such as video editing, 3D design, and live media.

Enterprise Reliability

Designed for professionals who demand the best, NVIDIA RTX PRO solutions deliver unparalleled performance, reliability, and support. Every GPU is rigorously tested for a wide range of design, engineering, and AI workflows and continually optimized through enterprise drivers. With extensive ISV certifications, robust IT management tools, and enterprise-grade support, RTX PRO workstations are the trusted choice for enterprise and mission-critical work. Support for NVIDIA NemoClaw allows organizations to enforce strict privacy and security guardrails when running local, always-on AI agents, ensuring data remains protected during complex AI workflows.

Specifications

GPU architecture	NVIDIA Blackwell
NVIDIA® CUDA® Cores	10,496
Tensor Cores	5th Generation
Ray Tracing Cores	4th Generation
AI TOPS	1,617 AI TOPS ^{1,2}
Single-precision performance	51 TFLOPS ¹
RT Core performance	153 TFLOPS ¹
GPU memory	32 GB GDDR7 with ECC
Memory interface	256-bit
Memory bandwidth	896 GB/s
System interface	PCIe 5.0 x16
Display connectors	4x DisplayPort 2.1b
Max simultaneous displays	>4x 3840 x 2160 @ 165 Hz >2x 7680 x 4320 @ 100 Hz
Video engines	>2x NVENC (9th Gen) >2x NVDEC (6th Gen)
Power consumption	Total board power: 200 W
Power connector	1x PCIe CEM5 16-pin
Thermal solution	Active
Form factor	4.4" x 10.5" L, dual slot, full height
Graphics APIs	DirectX 12, Shader Model 6.6, OpenGL 4.6 ³ , Vulkan 1.4 ³
Compute APIs	CUDA 12.8, OpenCL 3.0, DirectCompute

Ready to Get Started?

To learn more, visit: [nvidia.com/rtx-pro-4500](https://www.nvidia.com/rtx-pro-4500)

1. Peak rates based on GPU Boost Clock.
2. Effective FP4 TOPS with sparsity.
3. 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

© 2026 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, CUDA, RTX, and RTX PRO 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. 5108623. APR26

