

**XENON**
High Performance Computing

NVIDIA RTX PRO 6000 Blackwell Workstation Edition

Powering the next era of AI.



Transform Workflows With the Ultimate Single-GPU Powerhouse

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.

Experience unparalleled compute power with the NVIDIA RTX PRO™ 6000 Blackwell Workstation Edition—made for professionals who demand peak performance for the most complex workloads. Built on the revolutionary NVIDIA Blackwell architecture, the RTX PRO 6000 delivers up to 4000 TOPS of AI processing power, enabling lightning-fast LLM fine-tuning, real-time rendering, and precision-driven simulations. Equipped with 96 GB of GDDR7 memory it effortlessly handles massive datasets and multi-billion-parameter models while sustaining seamless multi-app workflows, from developing AI assistants and photorealistic visualization to AI-powered engineering simulations.

The RTX PRO 6000 harnesses a 600W power design to fuel the most demanding professional workloads—from fine-tuning trillion-parameter AI models to hyperrealistic cinematic rendering and simulating computational fluid dynamics. Its double-flow-through cooling system channels airflow through dual axial pathways, dissipating heat more efficiently to maintain stability while operating compute-intensive workloads. This thermal innovation ensures engineers can simulate physics models uninterrupted, AI researchers can iterate on neural networks faster, and VFX artists can render complex scenes without performance dips.

Ideal for AI developers, researchers, engineers, and creative professionals, the RTX PRO 6000 unlocks breakthroughs in real-time ray tracing, low-latency AI inference, and scientific modeling. Whether optimizing LLMs, developing agentic AI applications, rendering cinematic-grade assets, or simulating hyperdetailed 3D environments, the RTX PRO 6000 sets a new level of desktop workstation performance.

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
- 96 GB of GDDR7 memory
- 1.8 TB/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
- Multi-instance GPU (MIG) support
- AI Management Processor
- Double flow through thermal design

Breakthrough Innovations

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.

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 96 GB of GPU memory and 1.8 TB/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 2.1: 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 broadcasting.

Universal MIG: Divide a single RTX PRO 6000 Blackwell into multiple isolated instances, each with dedicated resources, allowing for concurrent execution of multiple workloads, optimized GPU utilization, and secure isolation of different applications or users.

Double-Flow-Through Design: The RTX PRO 6000 Blackwell features an innovative double-flow-through cooling design, optimizing thermal efficiency and airflow to sustain peak performance under 600W power loads. It operates quietly even during the most demanding tasks, eliminating thermal throttling and maximizing productivity.

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 is 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.

Specifications

GPU architecture	NVIDIA Blackwell
NVIDIA® CUDA® Cores	24,064
Tensor Cores	5th Generation
Ray Tracing Cores	4th Generation
AI TOPS	4000 AI TOPS ^{1,2}
Single-precision performance	125 TFLOPS ¹
RT Core performance	380 TFLOPS ¹
GPU memory	96 GB GDDR7 with ECC
Memory interface	512-bit
Memory bandwidth	1792 GB/s
System interface	PCIe 5.0 x16
Display connectors	4x DisplayPort 2.1b
Max simultaneous displays	>4x 4096 x 2160 @ 120 Hz >4x 5120 x 2880 @ 60 Hz >2x 7680 x 4320 @ 60 Hz
Video Engines	4x NVENC (9th Gen) 4x NVDEC (6th Gen)
MIG Support	>Up to 4x 24 GB >Up to 2x 48 GB >Up to 1x 96 GB
Power consumption	Total board power: 600 W
Power connector	1x PCIe CEM5 16-pin
Thermal solution	Double-flow-through
Form factor	5.4" H x 12" L, dual slot, extended height
Graphics APIs	DirectX 12, Shader Model 6.6, OpenGL 4.6 ³ , Vulkan 1.3 ³
Compute APIs	CUDA 12.8, OpenCL 3.0, DirectCompute

Ready to Get Started?

To learn more about NVIDIA RTX 2000, visit

<https://xenon.com.au/product/nvidia-rtx-pro-6000-workstation-edition/nvidia.com/rtx-pro-6000/>

Contact XENON today!

www.xenon.com.au | info@xenon.com.au | 1300 888 030

1. Peak rates based on GPU Boost Clock. 12. Effective FP4 TOPS with sparsity. 13. 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. All rights reserved. NVIDIA, CUDA, NVIDIA RTX PRO and the NVIDIA logo 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.

