



# NVIDIA RTX PRO Graphics and AI Solutions


You need to do great things—create and collaborate from anywhere, on any device, without the distractions of slow performance, poor stability, or application incompatibility. With NVIDIA RTX™ technology, you can unleash your vision and enjoy ultimate creative freedom.

NVIDIA RTX PRO™ products power a wide range of laptop, desktop, and data center solutions. Leverage the latest advancements in real-time ray tracing, AI, virtual reality (VR), and interactive, photorealistic rendering to develop revolutionary products, tell vivid visual stories, and design groundbreaking architecture like never before. Support for advanced features, frameworks, and SDKs across all of our products gives you the power to tackle the most challenging visual computing tasks, no matter the scale.




## NVIDIA RTX PRO Laptop GPUs

Professionals today increasingly need to work on complex visual computing and AI workflows on the go. NVIDIA RTX PRO laptop GPUs deliver world-class performance in a portable form factor combining the latest advancements in real-time ray tracing, advanced shading, and AI-based capabilities, so professionals can tackle demanding workflows from anywhere.



## NVIDIA RTX PRO Desktop GPUs

Power the most innovative workflows in design, engineering, and beyond with NVIDIA RTX PRO Desktop GPUs. Featuring breakthroughs in AI, ray tracing, and neural rendering technology, they accelerate photorealistic rendering, AI inference, and complex visualization tasks. Built with massive memory, enterprise-grade reliability, and certified drivers for 100+ professional applications, RTX PRO desktops deliver unmatched performance—seamlessly integrated into global partner ecosystems to power the next era of work.



## NVIDIA Data Center GPUs

Demand for visualization, rendering, data science, and simulation continues to grow as businesses tackle larger, more complex workloads. Scale up your visual compute infrastructure and tackle graphics-intensive workloads, complex designs, photorealistic renders, and augmented and virtual environments at the edge with NVIDIA GPUs. Optimized for reliability in enterprise data centers, NVIDIA GPUs feature both active and passive thermal solutions to fit into a variety of servers.

NVIDIA RTX PRO Graphics and AI Solutions

New		GPU Specifications							Performance			Display Technology					Virtual Reality		Options					
		NVIDIA CUDA® Processing Cores¹	NVIDIA RT Cores	Tensor Cores	GPU Memory²	Error-Correcting Code (ECC) Memory	Peak Memory Bandwidth	NVIDIA® NVLink®	Floating-Point Performance, Single Precision (TFLOPS, Peak)³	Accelerated Double Precision (TFLOPS, Peak)	AI TOPS⁴	Maximum Active Displays	DisplayPort⁵	HDMI via Adaptors, HDMI	NVIDIA SLI®⁶	High-Dynamic Range (HDR)⁷	NVIDIA Mosaic Technology	VR Ready⁸	Variable Rate Shading	GPUDirect® for Video	Graphics Synchronization with NVIDIA RTX PRO Sync	3D Stereo	Encode/Decode⁹	Multi-Instance GPU
Laptop GPUs																								
•	NVIDIA RTX PRO 5000 Blackwell	10,496	4th Gen	5th Gen	24 GB	•¹⁰	896 GB/s		49.8		1,824	4*	Yes*	Yes*		•	•	•	•	•		•	•	
•	NVIDIA RTX PRO 4000 Blackwell	7,680	4th Gen	5th Gen	16 GB	•¹⁰	896 GB/s		38.7		1,334	4*	Yes*	Yes*		•	•	•	•	•		•	•	
•	NVIDIA RTX PRO 3000 Blackwell	5,888	4th Gen	5th Gen	12 GB	•¹⁰	672 GB/s		29.1		992	4*	Yes*	Yes*		•	•	•	•	•		•	•	
•	NVIDIA RTX PRO 2000 Blackwell	3,328	4th Gen	5th Gen	8 GB		384 GB/s		17.7		798	4*	Yes*	Yes*		•	•	•	•	•		•	•	
•	NVIDIA RTX PRO 1000 Blackwell	2,560	4th Gen	5th Gen	8 GB		384 GB/s				572	4*	Yes*	Yes*		•	•	•	•	•		•	•	
•	NVIDIA RTX PRO 500 Blackwell	1,792	4th Gen	5th Gen	6 GB		288 GB/s				294	4*	Yes*	Yes*		•	•	•	•	•		•	•	
	NVIDIA RTX 5000 Ada Generation	9,728	3rd Gen	4th Gen	16 GB	•¹⁰	576 GB/s		42.6		682	4*	Yes*	Yes*		•	•	•	•	•	•	•	•	
	NVIDIA RTX 4000 Ada Generation	7,424	3rd Gen	4th Gen	12 GB	•¹⁰	432 GB/s		33.6		538	4*	Yes*	Yes*		•	•	•	•	•	•	•	•	
	NVIDIA RTX 3500 Ada Generation	5,120	3rd Gen	4th Gen	12 GB	•¹⁰	432 GB/s		23.0		369	4*	Yes*	Yes*		•	•	•	•	•		•	•	
	NVIDIA RTX 3000 Ada Generation	4,608	3rd Gen	4th Gen	8 GB	•¹⁰	256 GB/s		19.9		319	4*	Yes*	Yes*		•	•	•	•	•		•	•	
	NVIDIA RTX 2000 Ada Generation	3,072	3rd Gen	4th Gen	8 GB		256 GB/s		14.5		232	4*	Yes*	Yes*		•	•	•	•	•		•	•	
	NVIDIA RTX 1000 Ada Generation	2,560	3rd Gen	4th Gen	6 GB		192 GB/s		12.1		193	4*	Yes*	Yes*		•	•	•	•	•		•	•	
	NVIDIA RTX 500 Ada Generation	2,048	3rd Gen	4th Gen	4 GB		128 GB/s		9.2		154	4*	Yes*	Yes*		•	•	•	•	•		•	•	
Desktop GPUs																								
•	NVIDIA RTX PRO 6000 Blackwell Workstation Edition	24,064	4th Gen	5th Gen	96 GB	•	1,792 GB/s		125		4,000	4	4	4	•	•	•	•	•	•	•	•	•	•
•	NVIDIA RTX PRO 6000 Blackwell Max-Q Workstation Edition	24,064	4th Gen	5th Gen	96 GB	•	1,792 GB/s		110		3,511	4	4	4	•	•	•	•	•	•	•	•	•	•
•	NVIDIA RTX PRO 5000 Blackwell	14,080	4th Gen	5th Gen	48 GB	•	1,344 GB/s		TBC		TBC	4	4	4	•	•	•	•	•	•	•	•	•	•
•	NVIDIA RTX PRO 4500 Blackwell	10,496	4th Gen	5th Gen	32 GB	•	896 GB/s		TBC		TBC	4	4	4	•	•	•	•	•	•	•	•	•	•
•	NVIDIA RTX PRO 4000 Blackwell	8,960	4th Gen	5th Gen	24 GB	•	672 GB/s		TBC		TBC	4	4	4	•	•	•	•	•	•	•	•	•	•
	NVIDIA RTX 6000 Ada Generation	18,176	3rd Gen	4th Gen	48 GB	•¹⁰	960 GB/s		91		1,457	4	4	4	•	•	•	•	•	•	•	•	•	•
	NVIDIA RTX 5000 Ada Generation	12,800	3rd Gen	4th Gen	32 GB	•¹⁰	576 GB/s		65		1,044	4	4	4	•	•	•	•	•	•	•	•	•	•
	NVIDIA RTX 4500 Ada Generation	7,680	3rd Gen	4th Gen	24 GB	•¹⁰	432 GB/s		40		634	4	4	4	•	•	•	•	•	•	•	•	•	•
	NVIDIA RTX 4000 Ada Generation	6,144	3rd Gen	4th Gen	20 GB	•¹⁰	360 GB/s		27		427	4	4	4	•	•	•	•	•	•	•	•	•	•
	NVIDIA RTX 4000 SFF Ada Generation	6,144	3rd Gen	4th Gen	20 GB	•¹⁰	280 GB/s		19		307	4	4	4	•	•	•	•	•	•	•	•	•	•
	NVIDIA RTX 2000 Ada Generation	2,816	3rd Gen	4th Gen	16 GB	•¹⁰	224 GB/s		12		191	4	4	4	•	•	•¹²	•	•			•	•	
	NVIDIA A800 40GB Active	6,912		432 (3rd Gen)	40 GB	•¹¹	1,555 GB/s	•	20	9.7	1,247				•		•¹²		•				•	
	NVIDIA RTX A1000	2,304	18 (2nd Gen)	72 (3rd Gen)	8 GB		192 GB/s		7		107	4	4	4		•	•¹²		•			•	•	
	NVIDIA RTX A400	768	6 (2nd Gen)	24 (3rd Gen)	4 GB		96 GB/s		3		43	4	4	4		•	•¹²		•			•	•	
Data Center GPUs																								
•	NVIDIA RTX PRO 6000 Blackwell Server Edition	24,064	188 (4th Gen)	752 (5th Gen)	96 GB	•¹⁰	1,613 GB/s		117.3		3,753	4	4		•	•	•	•	•	•	•	•	•	•
	NVIDIA L40S	18,176	142 (3rd Gen)	568 (4th Gen)	48 GB	•¹⁰	864 GB/s		91.6		1,466	4	4	4	•	•	•	•	•	•	•	•	•	•

**\* Check with OEM manufacturer for specific display topology.**

1. CUDA parallel processing cores cannot be compared between GPU generations due to several important architectural differences that exist between streaming multiprocessor designs.

2. NVIDIA Blackwell architecture uses GDDR7 memory. NVIDIA Ada Lovelace and Ampere architecture use GDDR6 memory. A800 40GB Active uses HBM2 as its primary memory type.

3. Peak rates are based on GPU Boost clock.

4. NVIDIA Blackwell GPUs are based on FP4 AI TOPS with sparsity. For Ada-based GPUs, AI TOPS is based on FP8 with sparsity. AI TOPS for Ampere and Turing-based GPUs are based on INT8 with sparsity.

5. Feature support varies by system-level implementation. Check with your workstation OEM vendor for system specific configurations. NVIDIA Blackwell supports DisplayPort 2.1. Ada and Ampere support DisplayPort 1.4.

6. SLI functionality is provided via NVLink.

7. Supported adaptors are required for HDMI.

8. Supports multi-view rendering (MVR) feature.

9. For more details on GPU-specific video encode/decode format support, refer to: <https://developer.nvidia.com/video-encode-and-decode-gpu-support-matrix-new>

10. Ensures data integrity and reliability by eliminating soft errors on direct random-access memory (DRAM) only.

11. Ensures data integrity and reliability by eliminating soft errors on both GPU cache and on-board DRAM.

12. This Mosaic setup does not offer framelock synchronization or display overlap functionality.

For more information on NVIDIA professional graphics solutions, visit: [nvidia.com/rtx-pro/](https://nvidia.com/rtx-pro/)

