



Accelerate your most demanding HPC and hyperscale data center workloads with NVIDIA Data Center GPUs.

NVIDIA's accelerators also deliver the horsepower needed to run bigger simulations faster than ever before. Plus, NVIDIA GPUs deliver the highest performance and user density for virtual desktops, applications, and workstations.

TRAINING

Training increasingly complex models faster is key to improving productivity for data scientists and delivering AI services more quickly. Servers powered by NVIDIA GPUs use the performance of accelerated computing to cut deep learning training time from months to hours or minutes.

INFERENCE

Inference is where a trained neural network really goes to work. As new data points come in such as images, speech, visual and video search, inference is what gives the answers and recommendations at the heart of many AI services. A server with a single GPU can deliver 27X higher inference throughput than a single-socket CPU-only server resulting in dramatic cost savings.

AN ENTERPRISE-READY PLATFORM FOR PRODUCTION AI

NVIDIA AI Enterprise, an end-to-end, secure, cloud-native suite of AI software, accelerates the data science pipeline and streamlines the development and deployment of production AI.

	PART #	GPU SPECIFICATIONS			PERFORMANCE			DISPLAY TECHNOLOGY			OPTIONS									
		Part Number	CUDA Cores	GPU Memory	Memory Bandwidth	Tensor Cores	Ray Tracing Cores	Floating Point Performance, Single Precision (TFLOPS)	DisplayPort 2/3	Maximum Active Displays	Maximum Resolution Support @ 60 Hz	Dimensions	Mosaic Technology	Maximum Power Consumption	H.264 and HEVC Encode/Decode Engines	Vulkan Support	GPU Direct [™] for Video	Graphics Synchronization	NVLink Bridge Support	NVIDIA AI Enterprise Support
NVIDIA FOR DATA CENTERS																				
NVIDIA H200 NVL	NVH200NVLTGPU-KIT	-	141 GB	4.8 TB/s	-	-	-	-	-	-	4.4 H" x 10.5" FH DS	-	600W	-	-	-	-	-	✓	✓
NVIDIA RTX PRO 6000 Blackwell Server Edition	NVRTXPRO6000TCGPU-KIT	24,064	96 GB	1.6 TB/s	752	188	-	4	4	-	4.4 H" x 10.5" FH DS	-	600W	✓	-	-	-	-	-	✓
NVIDIA L40S	NVL40STCGPU-KIT	18,176	48 GB	864 GB/s	568	142	91.6 SP	4	4	7680x4320	4.4 H" x 10.5" FH DS	✓	350W	✓	✓	✓	✓	-	-	
NVIDIA A16	NVA16TCGPUNC-KIT	5,120	64 GB	800 GB/s	160	40	-	-	-	-	4.4 H" x 10.5" FH DS	-	250W	-	-	-	-	-	-	
NVIDIA RTX PRO 4500 Blackwell Server Edition	NVRTXPRO4500TCGPU-KIT	10,496	32 GB	800 GB/s	-	82	-	-	-	-	4.4 H" x 10.5" L FHFL SS	-	165W	✓	-	-	-	-	-	
NVIDIA L4	NVL4TCGPU-KIT	7,680	24 GB	300 GB/s	240	60	-	-	-	-	HH LP	-	72W	-	-	-	-	-	-	

For more information, contact a [PNY Account Manager](#) or email: GOPNY@PNY.COM or visit WWW.PNY.COM/NVIDIA-DATA-CENTER-SOLUTIONS

PNY Technologies, Inc. 100 Jefferson Road, Parsippany, NJ 07054 | Tel 973-515-9700 | Fax 973-560-5590 | WWW.PNY.COM

Features and specifications subject to change without notice. The PNY logo is a registered trademark of PNY Technologies, Inc. All other trademarks are the property of their respective owners. © 2026 PNY Technologies, Inc. All rights reserved.

March 2026 08:42 AM