


PNY RTX 4000 Ada

Artikel	132624
Herstellernummer	VCNRTX4000ADA-PB
EAN	3536403393960
PNY	



NVIDIA RTX 4000 ADA GENERATION **Performance for Endless Possibilities**

The NVIDIA RTX 4000 Ada Generation is the most powerful single-slot GPU for professionals, providing massive breakthroughs in speed and power efficiency over the previous generation to tackle demanding creative, design, and engineering workflows from the desktop. Harnessing the latest-generation RT Cores, Tensor Cores, and CUDA® cores alongside 20GB of graphics memory, RTX 4000 empowers professionals to create intricate product engineering, visionary cityscapes, and immersive entertainment experiences. With its power-efficient, single-slot design, RTX 4000 seamlessly fits into various workstation chassis, offering limitless potential for professionals.

NVIDIA Ada Lovelace Architecture

Experience fast, interactive performance—powered by the latest NVIDIA Ada Lovelace architecture-based GPU—with ultra-fast, onboard graphics memory technology and optimized software drivers for professional applications. The Shader Execution Reordering (SER) system allows on-the-fly organization and reordering of workloads, grouping similar performing threads so the streaming multiprocessor (SM) and RT Core can operate more efficiently.

CUDA Cores

The NVIDIA Ada Lovelace architecture-based CUDA cores deliver up to 2X the single-precision floating-point (FP32) throughput compared to the previous generation, providing significant performance improvements for graphics workflows such as 3D model development and compute for workloads like desktop simulation for CAE.

Third-Generation RT Cores

Third-generation RT Cores provide up to 2X the throughput of the previous generation and the ability to run ray tracing with either shading or denoising capabilities concurrently. This accelerates renders for M&E content creation, AECO design evaluations, and manufacturing virtual prototyping. Third-generation RT Cores deliver up to 2X the ray-tracing performance over the previous generation, delivering groundbreaking performance for photorealistic rendering. Enhanced RT Cores combined with new SER technology dynamically reorder inefficient workloads, dramatically improving shader performance to accelerate end-to-end ray-traced image rendering performance.

Fourth Generation Tensor Cores

Fourth-generation Tensor Cores provide up to 4X faster AI training performance than the previous generation with FP16 precision. Support for the new FP8 data format for inference provides more than 4X faster performance than the previous generation and reduces data memory usage by half (compared to the FP16 data format).

PCIe Gen 4

RTX 4000 supports PCIe Gen4, which doubles the bandwidth of PCIe Gen3 from 15.75GB/sec to 31.5 GB/sec for x16 connections, improving data transfer speeds from CPU memory for data-intensive tasks such as AI, data science, and creating 3D models from large datasets. Faster PCIe performance also accelerates GPU DMA transfers, providing faster video data transfers from GPUDirect for Video-enabled devices and faster IO with GPUDirect Storage.

Higher Speed GDDR6 Memory

RTX 4000 features 20GB of GDDR6 memory, providing the memory needed for rendering, data science, engineering simulation, and other GPU memory-intensive applications. Support for ECC memory delivers uncompromised computing accuracy and reliability for mission-critical applications.

Multi-Display

Four mini-DisplayPort 1.4a connectors, support for multiple 8K monitors, NVIDIA Quadro Sync, Mosaic, and Warp and Blend enable massive immersive environments for CAVE, video walls, and location-based entertainment deployments.

Extended Reality

Support for the latest high-resolution HMD devices, high-performance graphics, and 20GB of GPU memory enable the creation of AR and VR experiences for training, product validations, building walkthroughs, and compelling entertainment.

Form Factor and Power Efficiency

With a single-slot, low-profile form factor, and power-efficient design, RTX 4000 fits into a wide range of workstation chassis, providing professionals with a generous selection of compatible workstations from worldwide OEM vendors and delivering a groundbreaking GPU architecture for AI and AI-based neural graphics in desktop workstations.

Motion BVH

Hardware-accelerated rendering of motion blur—a common cinematic effect that's difficult to render—means artists no longer need to rely on traditional methods of using motion vectors to achieve motion blur. Motion vectors allow the artist to adjust motion blur in post but require visual fixes for reflections and translucency.

NVIDIA DLSS

The Ada Lovelace GPU architecture features a new Optical Flow Accelerator and AI-based DLSS Super Resolution with DL denoiser that boost DLSS 3.0's frame rates up to 4X compared to the previous version while maintaining or exceeding native image quality.

NVIDIA RTX Broadcast Engine

The NVIDIA RTX Broadcast Engine transforms offices into broadcast studios, upgrading standard webcams and microphones into premium smart devices with the power of AI. Improve the video and audio quality of live streams with AI capabilities such as virtual backgrounds, webcam auto frame, and microphone noise removal. With dedicated AI processors called Tensor Cores on NVIDIA RTX GPUs, the AI networks can run in real-time alongside your applications.

NVIDIA RTX Experience

NVIDIA RTX Experience delivers a suite of productivity tools to your desktop workstation, including desktop recording in up to 8K, automatic alerts for the latest NVIDIA RTX Enterprise driver updates, and access gaming features.

NVIDIA CUDA Parallel Computing Platform

Natively execute standard programming languages like C/C++ and Fortran, and APIs such as OpenCL, OpenACC, and Direct Compute to accelerate techniques such as ray tracing, video and image processing, and computation fluid dynamics.

Unified Memory

A single, seamless 49-bit virtual address space allows for the transparent migration of data between the full allocation of CPU and GPU memory.

NVIDIA GPU Direct for Video

GPUDirect for Video speeds communication between the GPU and video I/O devices by avoiding unnecessary system memory copies and CPU overhead.

NVIDIA Enterprise Management Tools

Maximize system uptime, seamlessly manage wide-scale deployments, and remotely control graphics and display settings for efficient operations.

Zusammenfassung

NVIDIA RTX 4000 ADA GENERATION

Performance for Endless Possibilities

The NVIDIA RTX 4000 Ada Generation is the most powerful single-slot GPU for professionals, providing massive breakthroughs in speed and power efficiency over the previous generation to tackle demanding creative, design, and engineering workflows from the desktop. Harnessing the latest-generation RT Cores, Tensor Cores, and CUDA@ cores alongside 20GB of graphics memory, RTX 4000 empowers professionals to create intricate product engineering, visionary cityscapes, and immersive entertainment experiences. With its power-efficient, single-slot design, RTX 4000 seamlessly fits into various workstation chassis, offering limitless potential for professionals.

NVIDIA Ada Lovelace Architecture

Experience fast, interactive performance—powered by the latest NVIDIA Ada Lovelace architecture-based GPU—with ultra-fast, onboard graphics memory technology and optimized software drivers for professional applications. The Shader Execution Reordering (SER) system allows on-the-fly organization and reordering of workloads, grouping similar performing threads so the streaming multiprocessor (SM) and RT Core can operate more efficiently.

CUDA Cores

The NVIDIA Ada Lovelace architecture-based CUDA cores deliver up to 2X the single-precision floating-point (FP32) throughput compared to the previous generation, providing significant performance improvements for graphics workflows such as 3D model development and compute for workloads like desktop simulation for CAE.

Third-Generation RT Cores

Third-generation RT Cores provide up to 2X the throughput of the previous generation and the ability to run ray tracing with either shading or denoising capabilities concurrently. This accelerates renders for M&E content creation, AECO design evaluations, and manufacturing virtual prototyping. Third-generation RT Cores deliver up to 2X the ray-tracing performance over the previous generation, delivering groundbreaking performance for photorealistic rendering. Enhanced RT Cores combined with new SER technology dynamically reorder inefficient workloads, dramatically improving shader performance to accelerate end-to-end ray-traced image rendering performance.

Fourth Generation Tensor Cores

Fourth-generation Tensor Cores provide up to 4X faster AI training performance than the previous generation with FP16 precision. Support for the new FP8 data format for inference provides more than 4X faster performance than the previous generation and reduces data memory usage by half (compared to the FP16 data format).

PCIe Gen 4

RTX 4000 supports PCIe Gen4, which doubles the bandwidth of PCIe Gen3 from 15.75GB/sec to 31.5 GB/sec for x16 connections, improving data transfer speeds from CPU memory for data-intensive tasks such as AI, data science, and creating 3D models from large datasets. Faster PCIe performance also accelerates GPU DMA transfers, providing faster video data transfers from GPUDirect for Video-enabled devices and faster IO with GPUDirect Storage.

Higher Speed GDDR6 Memory

RTX 4000 features 20GB of GDDR6 memory, providing the memory needed for rendering, data science, engineering simulation, and other GPU memory-intensive applications. Support for ECC memory delivers uncompromised computing accuracy and reliability for mission-critical applications.

Multi-Display

Four mini-DisplayPort 1.4a connectors, support for multiple 8K monitors, NVIDIA Quadro Sync, Mosaic, and Warp and Blend enable massive immersive environments for CAVE, video walls, and location-based entertainment deployments.

Extended Reality

Support for the latest high-resolution HMD devices, high-performance graphics, and 20GB of GPU memory enable the creation of AR and VR experiences for training, product validations, building walkthroughs, and compelling entertainment.

Form Factor and Power Efficiency

With a single-slot, low-profile form factor, and power-efficient design, RTX 4000 fits into a wide range of workstation chassis, providing professionals with a generous selection of compatible workstations from worldwide OEM vendors and delivering a groundbreaking GPU architecture for AI and AI-based neural graphics in desktop workstations.

Motion BVH

Hardware-accelerated rendering of motion blur—a common cinematic effect that's difficult to render—means artists no longer need to rely on traditional methods of using motion vectors to achieve motion blur. Motion vectors allow the artist to adjust motion blur in post but require visual fixes for reflections and translucency.

NVIDIA DLSS

The Ada Lovelace GPU architecture features a new Optical Flow Accelerator and AI-based DLSS Super Resolution with DL denoiser that boost DLSS 3.0's frame rates up to 4X compared to the previous version while maintaining or exceeding native image quality.

NVIDIA RTX Broadcast Engine

The NVIDIA RTX Broadcast Engine transforms offices into broadcast studios, upgrading standard webcams and microphones into premium smart devices with the power of AI. Improve the video and audio quality of live streams with AI capabilities such as virtual backgrounds, webcam auto frame, and microphone noise removal. With dedicated AI processors called Tensor Cores on NVIDIA RTX GPUs, the AI networks can run in real-time alongside your applications.

NVIDIA RTX Experience

NVIDIA RTX Experience delivers a suite of productivity tools to your desktop workstation, including desktop recording in up to 8K,

automatic alerts for the latest NVIDIA RTX Enterprise driver updates, and access gaming features.

NVIDIA CUDA Parallel Computing Platform

Natively execute standard programming languages like C/C++ and Fortran, and APIs such as OpenCL, OpenACC, and Direct Compute to accelerate techniques such as ray tracing, video and image processing, and computation fluid dynamics.

Unified Memory

A single, seamless 49-bit virtual address space allows for the transparent migration of data between the full allocation of CPU and GPU memory.

NVIDIA GPU Direct for Video

GPUDirect for Video speeds communication between the GPU and video I/O devices by avoiding unnecessary system memory copies and CPU overhead.

NVIDIA Enterprise Management Tools

Maximize system uptime, seamlessly manage wide-scale deployments, and remotely control graphics and display settings for efficient operations.

PNY RTX 4000 Ada, RTX A4000, 20 GB, GDDR6, 160 bit, PCI Express x16 4.0, 1 fan(s)

PNY RTX 4000 Ada. Graphics processor family: NVIDIA, Graphics processor: RTX A4000. Discrete graphics card memory: 20 GB, Graphics card memory type: GDDR6, Memory bus: 160 bit. Interface type: PCI Express x16 4.0. Cooling type: Active, Number of fans: 1 fan(s)

Merkmale

		Memory	
Power		Discrete graphics card memory	20 GB
Power consumption (max)		Graphics card memory type	GDDR6
		Memory bus	160 bit
		Memory bandwidth (max)	360 GB/s
Ports & interfaces		Performance	
Interface type	PCI Express x16 4.0	TV tuner integrated	No
DisplayPorts quantity	4	Virtual Reality (VR) ready	Yes
DisplayPort version	1.4a	Dual Link DVI	No
Weight & dimensions		NVIDIA 3D Vision	Yes
Weight	528 g	Processor	
Depth	241.3 mm	CUDA	Yes
Height	111.8 mm	CUDA cores	6144
Design		Graphics processor family	NVIDIA
Cooling type	Active	Graphics processor	RTX A4000
Number of fans	1 fan(s)	Parallel processing technology support	Not supported
Number of slots	1	Lithography	4 nm
Product colour	Black		

Preisänderungen und Irrtümer vorbehalten. Alle Produkte solange der Vorrat reicht.