


PNY NVIDIA RTX 4000 ADA

Artikel
Herstellernummer
EAN
PNY

132621
VCNRTX4000ADALP-SB
0751492776989



NVIDIA RTX 4000 SFF ADA GENERATION

Power for Endless Performance

In recent years, the demand for smaller footprint workstations has been skyrocketing, fueled by an array of factors, including cutting-edge hardware advancements, the extreme miniaturization of components, an increase in remote and mobile workforces, and the need for more energy efficient solutions suited for space available in modern offices and data centers. To meet this growing demand, NVIDIA has introduced the NVIDIA RTX™ 4000 SFF Ada Generation, the most powerful GPU on the market for Small Form Factor (SFF) workstations. Built on the revolutionary NVIDIA Ada Lovelace GPU architecture, the RTX 4000 SFF Ada brings new levels of performance, features, and capabilities to compact workstations, continuing to drive the evolution of visual computing for professionals. Designed for the needs of artists, architects, designers, and engineers across industries, the RTX 4000 SFF Ada is the ultimate compact computing solution to help you take your work to the next level

NVIDIA Ada Lovelace Architecture

NVIDIA RTX 4000 SFF Ada is the most powerful dual slot, low profile workstation GPU offering high-performance, real-time ray tracing, AI-accelerated compute, and professional graphics rendering, all of this in a compact design with unmatched performance. Building upon the major SM enhancements from the Ada Lovelace GPU, the NVIDIA Ada Lovelace architecture provides more cores, higher clocks, and a larger L2 cache for more performance to enhance ray tracing operations, tensor matrix operations, and frame rates with DLSS 3.0.

CUDA Cores

The NVIDIA Ada Lovelace architecture-based CUDA Cores offer more than 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 such as desktop simulation for computer-aided engineering (CAE). The RTX 4000 SFF Ada enables two FP32 primary data paths, doubling the peak FP32 operations.

Third Generation RT Cores

Incorporating 3rd generation ray tracing engines, NVIDIA Ada Lovelace architecture-based GPUs provide incredible ray-traced rendering performance. A single NVIDIA RTX 4000 SFF Ada board can render complex professional models with physically accurate shadows, reflections, and refractions to empower users with instant insight. Working in concert with applications leveraging APIs such as NVIDIA OptiX, Microsoft DXR, and Vulkan ray tracing, systems based on the NVIDIA RTX 4000 SFF Ada will power truly interactive design workflows to provide immediate feedback for unprecedented levels of productivity. The NVIDIA RTX 4000 SFF Ada features up to 2X faster ray-triangle intersection throughput compared to the previous generation.

Fourth Generation Tensor Cores

Specialized for deep learning matrix multiply and accumulate math operations at the heart of neural network training and inferencing functions, the NVIDIA RTX 4000 SFF Ada includes enhanced Tensor Cores that accelerate more data types and still support the Fine-

Grained Structured Sparsity feature that delivers more than 2X throughput for tensor matrix operations compared to the previous generation. New Tensor Cores will accelerate new FP8 precision modes. Independent floating-point and integer data paths allow more efficient execution of workloads using a mix of computation and addressing calculations.

PCIe Gen 4

The NVIDIA RTX 4000 SFF Ada supports PCI Express Gen 4, which provides double the bandwidth of PCIe Gen 3, improving data-transfer speeds from CPU memory for data-intensive tasks like AI and data science.

Higher Speed GDDR6 Memory

Built with 20GB GDDR6 memory, the NVIDIA RTX 4000 SFF Ada provides an ideal memory footprint to address datasets and models in latency-sensitive professional applications and at volume.

Error Correcting Code (ECC) on Graphics Memory

Meet strict data integrity requirements for mission-critical applications with uncompromised computing accuracy and reliability for workstations.

Fifth Generation NVDEC Engine

NVDEC is well suited for transcoding and video playback applications for real-time decoding. The following video codecs are supported for hardware-accelerated decoding: MPEG-2, VC-1, H.264 (AVCHD), H.265 (HEVC), VP8, VP9, and AV1 video formats. Video encoding at 8K/60 will be achievable for professional video editing.

Eighth Generation NVENC Engine

NVENC can take on the most demanding 4K or 8K video encoding tasks to free up the graphics engine and the CPU for other operations. The NVIDIA RTX 4000 SFF Ada provides better encoding quality than software-based x264 encoders. The NVIDIA RTX 4000 SFF Ada incorporates AV1 video encoding which is 40% more efficient than H.264 encoding for 4K HDR video. AV1 will provide better quality at the same bitrate bandwidth.

Graphics Preemption

Pixel-level preemption provides more granular control to better support time-sensitive tasks such as VR motion tracking.

Compute Preemption

Preemption at the instruction level provides finer-grain control over compute tasks to prevent long-running applications from either monopolizing system resources or timing out.

NVIDIA Mosaic Technology

Transparently scale the desktop and applications across up to 4 GPUs and 16 displays from a single workstation while delivering full performance and image quality.

Mini DisplayPort 1.4a

Support up to four 5K monitors at 60Hz, or dual 8K displays at 60Hz per card. The NVIDIA RTX 4000 SFF Ada supports HDR color for 4K at 60Hz for 10/12b HEVC decode and up to 4K at 60Hz for 10b HEVC encode. Each Mini DisplayPort connector can drive ultra-high resolutions of 4096 x 2160 at 120 Hz with 30-bit color.

NVIDIA RTX Desktop Manager

Gain unprecedented end-user control of the desktop experience for increased productivity in a single large display or multi-display environments, especially in the current age of large, widescreen displays.

Included NVIDIA Quadro Sync II

Synchronize the display and image output of up to 32 displays from 8 GPUs (connected through two Sync II boards) in a single system, reducing the number of machines needed to create an advanced video visualization environment.

Frame Lock Connector Latch

Each frame lock connector is designed with a self-locking retention mechanism to secure its connection with the frame lock cable to provide robust connectivity and maximum productivity.

OpenGL Quad Buffered Stereo Support

Provide a smooth and immersive 3D Stereo experience for professional applications.

Ultra-High-Resolution Desktop Support

Get more Mosaic topology choices with high-resolution display devices with a 32K Max desktop size.

Professional 3D Stereo Synchronization

Robust control of stereo effects through a dedicated connection to directly synchronize 3D stereo hardware to an NVIDIA RTX professional graphics card.

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. The application is available for download [here](#).

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 GPUDirect 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 SFF ADA GENERATION

Power for Endless Performance

In recent years, the demand for smaller footprint workstations has been skyrocketing, fueled by an array of factors, including cutting-edge hardware advancements, the extreme miniaturization of components, an increase in remote and mobile workforces, and the need for more energy efficient solutions suited for space available in modern offices and data centers. To meet this growing demand, NVIDIA has introduced the NVIDIA RTX™ 4000 SFF Ada Generation, the most powerful GPU on the market for Small Form Factor (SFF) workstations. Built on the revolutionary NVIDIA Ada Lovelace GPU architecture, the RTX 4000 SFF Ada brings new levels of performance, features, and capabilities to compact workstations, continuing to drive the evolution of visual computing for professionals. Designed for the needs of artists, architects, designers, and engineers across industries, the RTX 4000 SFF Ada is the ultimate compact computing solution to help you take your work to the next level

NVIDIA Ada Lovelace Architecture

NVIDIA RTX 4000 SFF Ada is the most powerful dual slot, low profile workstation GPU offering high-performance, real-time ray tracing, AI-accelerated compute, and professional graphics rendering, all of this in a compact design with unmatched performance. Building upon the major SM enhancements from the Ada Lovelace GPU, the NVIDIA Ada Lovelace architecture provides more cores, higher clocks, and a larger L2 cache for more performance to enhance ray tracing operations, tensor matrix operations, and frame rates with DLSS 3.0.

CUDA Cores

The NVIDIA Ada Lovelace architecture-based CUDA Cores offer more than 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 such as desktop simulation for computer-aided engineering (CAE). The RTX 4000 SFF Ada enables two FP32 primary data paths, doubling the peak FP32 operations.

Third Generation RT Cores

Incorporating 3rd generation ray tracing engines, NVIDIA Ada Lovelace architecture-based GPUs provide incredible ray-traced rendering performance. A single NVIDIA RTX 4000 SFF Ada board can render complex professional models with physically accurate shadows, reflections, and refractions to empower users with instant insight. Working in concert with applications leveraging APIs such as NVIDIA OptiX, Microsoft DXR, and Vulkan ray tracing, systems based on the NVIDIA RTX 4000 SFF Ada will power truly interactive design workflows to provide immediate feedback for unprecedented levels of productivity. The NVIDIA RTX 4000 SFF Ada features up to 2X faster ray-triangle intersection throughput compared to the previous generation.

Fourth Generation Tensor Cores

Specialized for deep learning matrix multiply and accumulate math operations at the heart of neural network training and inferencing functions, the NVIDIA RTX 4000 SFF Ada includes enhanced Tensor Cores that accelerate more data types and still support the Fine-Grained Structured Sparsity feature that delivers more than 2X throughput for tensor matrix operations compared to the previous generation. New Tensor Cores will accelerate new FP8 precision modes. Independent floating-point and integer data paths allow more efficient execution of workloads using a mix of computation and addressing calculations.

PCIe Gen 4

The NVIDIA RTX 4000 SFF Ada supports PCI Express Gen 4, which provides double the bandwidth of PCIe Gen 3, improving data-transfer speeds from CPU memory for data-intensive tasks like AI and data science.

Higher Speed GDDR6 Memory

Built with 20GB GDDR6 memory, the NVIDIA RTX 4000 SFF Ada provides an ideal memory footprint to address datasets and models in latency-sensitive professional applications and at volume.

Error Correcting Code (ECC) on Graphics Memory

Meet strict data integrity requirements for mission-critical applications with uncompromised computing accuracy and reliability for workstations.

Fifth Generation NVDEC Engine

NVDEC is well suited for transcoding and video playback applications for real-time decoding. The following video codecs are supported for hardware-accelerated decoding: MPEG-2, VC-1, H.264 (AVCHD), H.265 (HEVC), VP8, VP9, and AV1 video formats. Video encoding at 8K/60 will be achievable for professional video editing.

Eighth Generation NVENC Engine

NVENC can take on the most demanding 4K or 8K video encoding tasks to free up the graphics engine and the CPU for other operations. The NVIDIA RTX 4000 SFF Ada provides better encoding quality than software-based x264 encoders. The NVIDIA RTX 4000 SFF Ada incorporates AV1 video encoding which is 40% more efficient than H.264 encoding for 4K HDR video. AV1 will provide better quality at the same bitrate bandwidth.

Graphics Preemption

Pixel-level preemption provides more granular control to better support time-sensitive tasks such as VR motion tracking.

Compute Preemption

Preemption at the instruction level provides finer-grain control over compute tasks to prevent long-running applications from either monopolizing system resources or timing out.

NVIDIA Mosaic Technology

Transparently scale the desktop and applications across up to 4 GPUs and 16 displays from a single workstation while delivering full performance and image quality.

Mini DisplayPort 1.4a

Support up to four 5K monitors at 60Hz, or dual 8K displays at 60Hz per card. The NVIDIA RTX 4000 SFF Ada supports HDR color for 4K at 60Hz for 10/12b HEVC decode and up to 4K at 60Hz for 10b HEVC encode. Each Mini DisplayPort connector can drive ultra-high resolutions of 4096 x 2160 at 120 Hz with 30-bit color.

NVIDIA RTX Desktop Manager

Gain unprecedented end-user control of the desktop experience for increased productivity in a single large display or multi-display environments, especially in the current age of large, widescreen displays.

Included NVIDIA Quadro Sync II

Synchronize the display and image output of up to 32 displays from 8 GPUs (connected through two Sync II boards) in a single system, reducing the number of machines needed to create an advanced video visualization environment.

Frame Lock Connector Latch

Each frame lock connector is designed with a self-locking retention mechanism to secure its connection with the frame lock cable to provide robust connectivity and maximum productivity.

OpenGL Quad Buffered Stereo Support

Provide a smooth and immersive 3D Stereo experience for professional applications.

Ultra-High-Resolution Desktop Support

Get more Mosaic topology choices with high-resolution display devices with a 32K Max desktop size.

Professional 3D Stereo Synchronization

Robust control of stereo effects through a dedicated connection to directly synchronize 3D stereo hardware to an NVIDIA RTX professional graphics card.

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. The application is available for download [here](#).

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 GPUDirect 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 NVIDIA RTX 4000 ADA, Quadro RTX 4000, 20 GB, GDDR6, 160 bit, 7680 x 4320 pixels, PCI Express x16 4.0

PNY NVIDIA RTX 4000 ADA. Graphics processor family: NVIDIA, Graphics processor: Quadro RTX 4000. Discrete graphics card memory: 20 GB, Graphics card memory type: GDDR6, Memory bus: 160 bit. Maximum resolution: 7680 x 4320 pixels. DirectX version: 12.0, OpenGL version: 4.6. Interface type: PCI Express x16 4.0. Cooling type: Active, Number of fans: 1 fan(s)

Merkmale

Power

Power consumption (typical)	70 W
-----------------------------	------

Technical details

Sustainability certificates	CE, RoHS
-----------------------------	----------

Weight & dimensions

Depth	167.6 mm
Height	68.5 mm

Design

Cooling type	Active
Number of fans	1 fan(s)
Number of slots	2

Ports & interfaces

Interface type	PCI Express x16 4.0
DisplayPorts quantity	4
DisplayPort version	1.4a

Memory

Discrete graphics card memory	20 GB
Graphics card memory type	GDDR6
Memory bus	160 bit
Memory bandwidth (max)	320 GB/s

Processor

CUDA	Yes
CUDA cores	6144
Graphics processor family	NVIDIA
Graphics processor	Quadro RTX 4000
Maximum resolution	7680 x 4320 pixels
Maximum displays per videocard	4

Performance

DirectX version	12.0
Shader model version	6.6
OpenGL version	4.6
Virtual Reality (VR) ready	Yes
HDCP	Yes
HDCP version	2.2
Dual Link DVI	No

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