



NVIDIA RTX 6000 ADA

PERFORMANCE FOR ENDLESS POSSIBILITIES

Amplified Performance for Professionals

The NVIDIA RTX™ 6000 Ada Generation is designed to meet the challenges of today's professional workflows. Built on the NVIDIA Ada Lovelace architecture, the RTX 6000 combines 142 third-generation RT Cores, 568 fourth-generation Tensor Cores, and 18,176 CUDA® cores with 48GB of graphics memory to deliver the next generation of AI graphics and petaflop inferencing performance for unprecedented speed-up of rendering, AI, graphics, and compute workloads. RTX 6000-powered workstations provide what you need to succeed in today's ultra-challenging business environment.

NVIDIA CUDA CORES

With up to 2X the throughput over the previous generation, third-generation RT Cores deliver massive speedups for workloads like photorealistic rendering of movie content, architectural design evaluations, and virtual prototyping of product designs. This technology also accelerates the rendering of ray-traced motion blur with greater visual accuracy.

Fourth-Generation Tensor Cores

Double-speed processing for single-precision floating point (FP32) operations provides significant performance improvements for graphics and simulation workflows, such as complex 3D computer-aided design (CAD) and computer-aided engineering (CAE), on the desktop.

Third-Generation RT Cores

Fourth-generation Tensor Cores provide faster AI compute performance, delivering more than 2X the performance of the previous generation. These new Tensor Cores support acceleration of the FP8 precision data type and provide independent floating-point and integer data paths to speed up execution of mixed floating point and integer calculations.

48 Gibabytes (GB) of GPU Memory

With 48GB GDDR6 memory, the RTX 6000 gives data scientists, engineers, and creative professionals the large memory necessary to work with massive datasets and workloads like rendering, data science, and simulation.

SPECIFICATIONS

Part Number	VCNRTX6000ADA-PB
EAN Number	3536403392635
GPU memory	48 GB GDDR6 with ECC
Memory interface	384-bit
Memory bandwidth	768 GB/s
Error-correcting code (ECC)	Yes
CUDA Parallel Processing Cores	18176
NVIDIA fourth-generation Tensor Cores	568
NVIDIA third-generation RT Cores	142
Single-precision performance	Up to 87.5 TFLOPS
Rendering performance	Up to 202.3 TFLOPS
Tensor performance	Up to 700 TFLOPS*
NVIDIA NVLink	No
System interface	PCI Express 4.0 x16
Power consumption	Total board power: 300 W
Thermal solution	Active
Form factor	111.76 cm H x 266.7 cm L, dual slot, full height
Display connectors	4x DisplayPort 1.4a
Max simultaneous displays	4x 4096 x 2160 @ 120 Hz, 4x 5120 x 2880 @ 60 Hz, 2x 7680 x 4320 @ 60 Hz
Power connector	1x CEM5 16-pin CPU
Encode/decode engines	3x encode, 3x decode (+ AV1) 4 JPEG Decode Engines
VR ready	Yes
vGPU software support	NVIDIA vPC/vApps, NVIDIA RTX Virtual Workstation
vGPU profiles supported	1 GB, 2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 16 GB, 24 GB, 48 GB
Graphics APIs	DirectX 12.07, Shader Model 6.5, OpenGL 4.68, Vulkan 1.18
Compute APIs	CUDA, DirectCompute, OpenCL™

* using FP8 data format with structural sparsity enabled.

