

NVIDIA RTX 6000 ADA GENERATION

SPECIFICATIONS

Product	NVIDIA® RTX™ 6000 Ada Generation
Architecture	NVIDIA Ada Lovelace Architecture
Foundry	TSMC
Process Size	4 nm NVIDIA Custom Process
Transistors	76.3 billion
Die Size	608.4 mm ²
CUDA Parallel Processing Cores	18,176
RT Cores	142 Gen 3
NVIDIA Tensor Cores	568
NVIDIA RT Cores	142
Single-Precision Performance ¹	91.1 TFLOPS
RT Core Performance ¹	210.6 TFLOPS
Tensor Performance ¹	1457.0 TFLOPS ²
GPU Memory	48 GB GDDR6 with ECC
Memory Interface	384-bit
Memory Bandwidth	960 GB/s
Max Power Consumption	300W
Graphics Bus	PCI Express 4.0 x16
Display Connectors	DP 1.4a (4) ³
Form Factor	4.4" H x 10.5" L Dual Slot
Product Weight	1.180 kg
Thermal Solution	Blower Active Fan
vGPU Software Support ³	NVIDIA vPC/vApps, NVIDIA RTX Virtual Workstation (vWS)
vGPU Profiles Supported	1 GB, 2 GB, 3 GB, 4 GB, 6 GB, 8 GB, 12 GB, 16 GB, 24 GB, 48 GB
NVIDIA® 3D Vision® and 3D Vision Pro	Support via 3-pin Mini DIN

Frame Lock	Compatible (with Quadro Sync II)
NVLink	Not Supported
NVLink Interconnect	Not Supported
Power Connector	1x PCIe CEM5 16-pin
NVENC NVDEC	3x 3x (+AV1 Encode & Decode)

Performance numbers may be subject to change until product availability.

¹Peak rates are based on GPU boost clock.

²Effective FP8 TFLOPS using the new sparsity feature.

³Display ports are on by default for the RTX 6000 Ada Generation. Display ports are not active when using vGPU software.

⁴Virtualization support for the RTX 6000 Ada Generation will be available in an upcoming NVIDIA virtual GPU (vGPU) release, anticipated in Q1, 2023.

SUPPORTED PLATFORMS

- Microsoft Windows 11 (64-bit)
- Microsoft Windows 10 (64-bit)
- Red Hat Enterprise Linux 7.x
- SUSE Linux Enterprise Desktop 15.x
- OpenSuse 15
- Fedora 31
- Ubuntu 18.04
- FreeBSD 11.x
- Solaris 11