

Nvidia Intros Quadro RTX 4000

Written by Frederick Douglas
16 November 2018

Nvidia presents the Quadro RTX 4000 at the annual Autodesk Conference-- a first Turing-based midrange professional GPU from the company, packing the appropriate amount of CUDA, ray-tracing (RT) and Tensor cores.



The Quadro RTX 4000 is essentially a workstation version of the GeForce RTX 2070 gaming GPU reworked for professional rendering and deep learning workloads. As such, it carries 2304 CUDA cores, 36 RT cores and 288 Tensor cores, plus 8GB GDDR6 graphics memory. It demands 150W of power, and Nvidia says the fully enabled TU106 GPU pushes up to 7.1 TFLOPS of FP32 performance, 6 Giga Rays/Sec and 43T RTX-OPS.

The GPU has a single-slot design connecting to a PCIe 3.0 x16 slot, with max power consumption reaching 160W. Onboard outputs include x3 DP 1.4 and x1 VirtualLink-enabled USB-C, and the card features 8K video decoding and encoding capability for multiple professional formats.

Nvidia says the card ships from December 2018.

Go [Nvidia Quadro RTX 4000](#)