

Quantum TXR410-144R

Intel® Xeon® / NVIDIA® Tesla™ C2050/C2070 GPU Server



Features

- The fastest GPU server available, with up to 3,584 GPU cores, 8.24 Teraflops of single precision floating point, and 4.12 Teraflops of double precision floating point performance in a 4U chassis
- Powered by the massively parallel CUDA architecture, transforming a workstation to perform like a small cluster at 1/20th the power consumption and 1/10th the cost
- CUDA programming environment with broad support of programming languages and APIs, including C, C++, OpenCL, DirectCompute, or Fortran
- Offers protection of data in memory to enhance data integrity and reliability for applications. Register files, L1/L2 caches, shared memory, and DRAM all are ECC protected
- Supports NVIDIA® Parallel DataCache™ Technology, accelerating algorithms such as physics solvers, ray-tracing, and sparse matrix multiplication where data addresses are not known beforehand
- Supports NVIDIA® GigaThread™ Engine, maximizing throughput with faster context switching, concurrent kernel execution, and improved thread block scheduling
- Seamlessly fits into existing HPC environments
- Ideal for life sciences, geosciences, engineering & sciences, molecular biology, medical diagnostics, electronic design automation (EDA), government and defense, visualization, financial modeling, oil & gas, and electromagnetic (EM) simulation applications

SPECIFICATIONS

Processor

- Supports two 32nm Intel® Xeon® processor Westmere-EP 5500/5600 series, with up to six cores & 12 threads, and 12MB L2 cache per processor
- Supports Trusted Execution, Advanced Encryption Security - New Instructions (AESNI), Turbo Boost, Intel® Virtualization, Intel® QuickPath Interconnect, and Intel® Hyper-Threading technologies

Motherboard Chipset

- Intel® 5520 chipset
- Integrated memory controllers for each processor
- Intel® QuickPath architecture with 6.40 GT/s / 5.86 GT/s / 4.80 GT/s (Gigatransfers per second)

Memory

- Up to 144GB DDR3 Reg. ECC 1333/1066 memory
- 18 x DDR3 memory DIMMs

Graphics Processing Unit (GPU):

- Supports eight Tesla C2050/C2070 GPU cards (448 computing cores per card, 3,584 cores total)
- Delivers 8.24 Teraflops of single precision floating point and 4.12 Teraflops of double precision floating point performance in a 4U chassis
- IEEE 754 single & double floating point precision
- Up to 24 or 48 GB dedicated DDR5 memory (organized as 3GB per Tesla C2050 GPU and 6GB per Tesla C2070 GPU)
- Up to 8x 384-bit GDDR5 memory interface (organized as a 384-bit interface per GPU)
- Up to 1.152 TB/sec memory bandwidth (144 GB/s per GPU to local memory)

Expansion Slots

- 2 x PCI-E x16 slots (w/ x4 link)
- 8 x PCI-E Gen.2 x16 slots
- 2 x PCI-E Gen.2 x1 slots
- 1 x PCI 32-bit slot

Graphics & LAN

- 4 x Gigabit Ethernet connections w/ Intel 82574L chipset
- Onboard graphics controller

Server Management

IPMI 2.0 compliant baseboard management controller (BMC) / USB 2.0 virtual hub / BIOS update

Chassis

- 4U rack-optimized chassis
- Power supply: 2,800 watt power supply
- 1 x 3.5" fixed hard disk drive bay

Chassis Dimensions

6.93" x 17.245" x 27.96" / 176 x 438 x 710 mm (H x W x D)

Warranty

Three years