PART NUMBER:
TCSM40M-24GB-PB

NVIDIA TESLA M40 24GB
THE WORLD'S FASTEST DEEP LEARNING TRAINING ACCELERATOR

The **NVIDIA® Tesla® M40** is the world's fastest deep learning training accelerator, purpose-built to train larger, more sophisticated neural networks within hours—versus days on CPU-only systems. It features an **NVIDIA Maxwell GPU** and **24 GB** of large-capacity memory, letting users increase the detection and prediction accuracy of their deep learning models and accelerate time-to-deployment.

Deep learning is redefining what's possible. From early-stage startups to large web service providers, deep learning has become the fundamental building block in delivering amazing solutions for end users.

Today's leading deep learning models typically take days to weeks to train, forcing data scientists to make compromises between accuracy and time to deployment. The NVIDIA Tesla M40 GPU accelerator is the world's fastest accelerator for deep learning training, purpose-built to dramatically reduce training time.

Running Caffe and Torch on the Tesla M40 delivers the same model within days versus weeks on CPU based compute systems.

Some of the world's largest data centers take advantage of Tesla accelerators to deliver unprecedented system throughput. The Tesla Platform supports industry-standard applications and system management tools, making it easier than ever before for IT managers to maximize uptime and system performance.

Deliver faster discoveries and scientific insights to your users by deploying GPU accelerators in your data center. With broad support for HPC developer tools like MPI, scientific libraries, and OpenACC, most applications offer performance boost with GPUs today.

TESLA M40 24GB - PRODUCT SPECIFICATION

MEMORY SIZE (PER BOARD)	24 GB GDDR5
MEMORY INTERFACE	384-bit
MEMORY BANDWIDTH	288 GB/s
CUDA CORES	3072
PEAK SINGLE PRECISION FLOATING POINT PERFORMANCE	~7 Tflops (GPU Boost Clocks)
SYSTEM INTERFACE	PCI Express 3.0 x16
MAX POWER CONSUMPTION	250 W
THERMAL SOLUTION	passive Heatsink
FORM FACTOR	111,15 mm (H) x 267,7 mm (L) Dual Slot, Full Height
DISPLAY CONECTORS	None
POWER CONNECTORS	8-pin CPU power connector 2-pin power brake header
WEIGHT (W/O EXTENDER)	898g
PACKAGE CONTENT	1x Power adapter (2 x PCIe 8-pin to CPU 8-pin)
PART NUMBER AND EAN	TCSM40M-24GB-PB 3536403348922

NVIDIA. | PNY®