PART NUMBER:
TCSM4M-PB

NVIDIA TESLA M4 by PNY

THE WORLD'S FIRST ACCELERATOR FOR
THE HYPERSCALE DATA CENTER

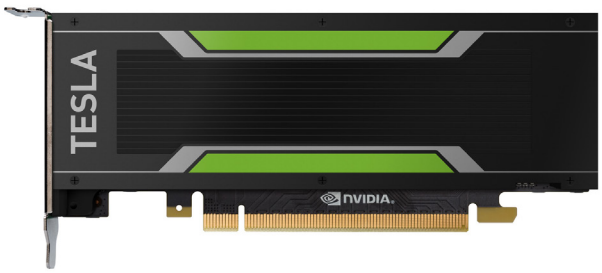

The Tesla M4 GPU brings the power of the Tesla Accelerated Computing Platform to the hyperscale data center. This is a low-power, small form factor GPU accelerator optimized for video transcoding, image processing, and machine learning inference that efficiently offloads demanding applications and boosts data center throughput.

Exploding volumes of user-generated data are redefining what's required for hyperscale data centers. Today's cloud applications harness valuable data to deliver smarter, real-time experiences using modern video and image processing and deep learning techniques. These applications can benefit greatly from GPU acceleration in the data center.

The NVIDIA® Tesla® M4 is the world's first accelerator designed for hyperscale servers, enabling customers to keep up with ever-growing amount of data. It's engineered to accelerate application throughput in a small, low-power design, slashing data center costs by half and deliver up to 7x more power-efficient processing than CPUs for deep learning at 20 images/sec/watt, inference, machine learning prediction, and video workloads.

The combination of the Tesla accelerated computing platform and NVIDIA hyperscale software suite provides an end-to-end solution to build and deploy modern hyperscale applications.

TESLA M4 - PRODUCT SPECIFICATION

MEMORY SIZE (PER BOARD)	4 GB GDDR5	
MEMORY INTERFACE	128-bit	
MEMORY BANDWIDTH	88 GB/s	
CUDA CORES	1024	
PEAK SINGLE PRECISION FLOATING POINT	2.2 Tflops (GPU Boost Clocks)	
SYSTEM INTERFACE	PCI Express 3.0 x16	
MAX POWER CONSUMPTION	Configurable: 50W minimum, 75 W default	
THERMAL SOLUTION	passive Heatsink	
FORM FACTOR	68.58 mm (H) x 167.64 mm (L)	
DISPLAY CONECTORS	None	
POWER CONNECTORS	None	
PACKAGE CONTENT	1x ATX Bracket	
WEIGHT (W/O EXTENDER)	234.9g	
PART NUMBER UND EAN	TCSM4M-PB	3536403348878

PNY Technologies Europe
sales@pny.eu | T: +33 (0)5 56 13 75 75