



nTeligence Model T Edge Computing Node

The Model T is the world's most powerful* and advanced** artificial intelligence appliance. It was designed to be used at the edge of the network, and to fill a gap in the hardware market, which exists between general-purpose desktop workstations, rack mounted servers, and gaming PCs. The machine serves as a development environment for machine learning, operational decision support, electronic worker (eWorker),

and conversational AI applications. In addition, it also serves as the actual deployment platform for production ready applications. This high-powered computational device is optimized to run **AIOS**, our Artificial Intelligence Operating System (AIOS). A license for AIOS is included along with every Model T. The software stack enables organizations to build and deploy a completely new generation of intelligent business, medical, government, industrial, and military applications. We call it "AI 2.0"

The Model T was engineered to provide the computational power necessary to support Hybrid Intelligence™, which seamlessly combines mathematical approaches to machine learning along with human like thought processes. In addition, it also supports True Meaning™ for natural language understanding. The AI based workstation is powered by AMD's 3rd generation 3990x Threadripper microprocessor. The 3990X has 64 CPU cores, supports 128 threads, has a maximum clock speed of 4.3 Ghz, up to 288 MB of combined cache (L1, L2 and L3), and 88 PCIe Gen4 lanes to meet demanding data transfer requirements. It also includes 192 GB of quad channel 3200 MHz DDR4 memory, a blazingly fast 1 TB PCIe 4.0 NVMe solid state drive, and a liquid cooling system complete with a radiator and variable speed pump. In addition, to support building and running deep and reinforcement learning models, it incorporates a Geforce RTX 3090 graphics adapter. This video card has 24GB of GDDR6 memory and is based upon the latest NVIDIA Ampere GA102 GPU. It has a total of 10,496 CUDA cores, 3,500 more than its closest cousin the A100 GPU, which is targeted primarily towards servers located within data centers. Like other chips in the Ampere product line the GA102 can be sub-divided into a varying number of smaller logical GPUs if needed. The RTX 3090 is fully compatible with today's high resolution 4K video monitors.

Expected delivery time is somewhere between 2-3 weeks. For additional information, current pricing, or to place an order, please contact our sales department at info@nteligence.com or call 609-651-0070.

*Based Upon a Single CPU/GPU Configuration

** When bundled with AIOS