



Lambda Vector One

4 GPU deep learning workstation with AMD Ryzen™ Threadripper™ processor



Configure your Vector [↗](#)



TECHNICAL SPECIFICATIONS

GPU DETAILS

1x RTX 4090 24 GB Liquid-cooled

PROCESSOR

AMD Ryzen™ 9 7950X 16-core, 32-thread

SYSTEM RAM

64 GB or 128 GB DDR5

STORAGE

OS - Up to 3.84 TB M.2 (NVMe)
Data - Up to 3 x 3.84 TB M.2 (NVMe) + 3 x 7.68 TB (SATA)

NETWORK INTERFACE

1x 10Gb Ethernet and 1x 2.5Gb Ethernet

Sleek power that doesn't disturb

The Vector One has been meticulously designed with both liquid cooling for the CPU and GPU, ensuring optimal performance without the noise. Even under typical high workloads, it only emits a mere 39 dB of noise, making it perfect for maintaining a quiet workspace.

Next-gen graphics for advanced AI/ML tasks

Equipped with the cutting-edge RTX 4090 boasting 24 GB of VRAM, the Vector One stands ready to tackle demanding tasks. From fine-tuning Stable Diffusion to handling the complexities of LLAMA 2 7B, this machine ensures high-level computations are a breeze.

Experience the power of future-ready architecture

At the heart of Vector One lies the state-of-the-art AMD Ryzen 9 7950X CPU, hosted on the advanced X670E chipset. This powerhouse not only supports PCIe Gen 5 and DDR5 but also offers up to twice the bandwidth of its predecessors. Dive into the future of computing with unrivaled speed and efficiency.

Delivering the optimal experience for AI/ML

Through rigorous research and experience, our engineers have crafted the ultimate system configuration tailored for AI/ML tasks. No more guesswork or configurations needed; the Vector One is fine-tuned to deliver unparalleled performance right out of the box. Additionally, every Vector One comes with a one-year warranty on hardware, with an option to extend to three years. For added peace of mind, choose to include dedicated technical support for Ubuntu and all ML frameworks and drivers that come pre-installed with your machine.

Pre-installed with the software you need

Each Vector One is pre-installed with Lambda Stack, which includes everything you need to get started training neural networks.





Lambda Vector One

GPU

1x RTX 4090 Liquid Cooled with 240mm radiator

CPU

AMD Ryzen™ 9 7950X 16-core, 32-thread (4.5 GHz Base, 5.7 GHz Boost)

MEMORY

64 GB 4800 MHz DDR5 or 128 GB 3600 MHz DDR5
4 DIMM slots

MOTHERBOARD

ASUS ProArt X670E-CREATOR WIFI

STORAGE

OS storage: Up to 3.84 TB M.2 (NVMe)
Extra storage: Up to 3 x 3.84 TB M.2 (NVMe SSD) + 3 x 7.68 TB (SATA SSD)

POWER SUPPLY

1300W 80 Plus Gold power supply

NETWORKING

1 x Marvell® AQtion 10Gb Ethernet
1 x Intel® 2.5Gb Ethernet

FORM FACTOR

Midtower desktop

COOLING

Liquid-cooled for both CPU and GPU:

AIO with 360mm Radiator + 3x 120mm fans for CPU
AIO with 240mm Radiator + 2x 120mm fans for GPU

DIMENSIONS

H 20.28" x W 9.45" x D 19.69" (H 515mm x W 240mm x D 500mm)

INPUT/OUTPUT

Front panel:

- 2x USB 3 Type-A
- 1x USB 3 Type-C
- 1x Microphone/Headphone Jack

Rear panel:

- 2x USB 4 Type-C ports
- 7x USB 3.2 Gen2 Type-A ports
- 1x USB 3.2 Gen 2x2 Type-C ports
- 1x USB 2.0 Type-A ports
- 1x DisplayPort (Input only)
- 1x HDMI port
- 1x 10Gb Ethernet port
- 1x 2.5Gb Ethernet port
- 7.1 audio out

NOISE*

39 dB (Typical)
44 dB (Maximum)

*Measured 30" diagonally from the front left side of the system. Typical represents noise during [Lambda's Deep Learning benchmarks](#). Maximum represents noise with all fans spinning at 100%.

LIGHTING

D-RGB lighting for chassis front, side, front fans and CPU block with dedicated hardware controls

SOFTWARE

Ubuntu Linux
Lambda Stack with CUDA, cuDNN, TensorFlow, PyTorch, Keras

COMPLIANCE

RoHS Compliant
TAA Compliant
PSU: 80 Plus Gold, CE mark

WARRANTY AND SUPPORT

Up to 3 years of hardware coverage, plus technical support from Lambda engineer