Lambda

Lambda Hyperplane 8-B200

8 NVIDIA® B200 GPUs with NVLink® & NVSwitch™

Configure your Hyperplane

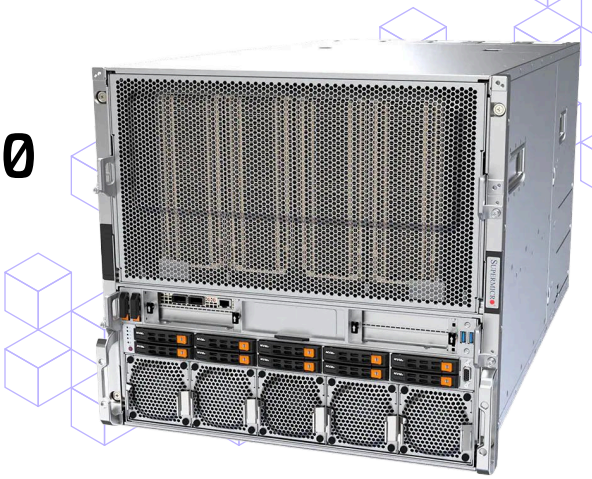

Microsoft · Los Alamos National Laboratory · Carnegie Mellon · Anthem · MIT

TECHNICAL SPECIFICATIONS

GPU DETAILS

8x NVIDIA B200 SXM 141 GB GPUs with NVLink and NVSwitch

PROCESSOR

2x AMD EPYC™ or Intel Xeon® processors

SYSTEM RAM

Up to 3072 GB DDR5

STORAGE

Up to 8× 30.72 TB NVMe (AMD)

Up to 10× 30.72 TB NVMe (Intel)

NETWORK INTERFACE

Up to 9x NVIDIA InfiniBand NDR 400 Gb/s NICs

8 B200 SXM GPUs deliver 72 petaFLOPS of AI performance

Eight NVIDIA B200 GPUs interconnected with an NVLink and NVSwitch fabric allow the Hyperplane 8-B200 to serve up to 72 petaFLOPS of FP8 performance for unprecedented acceleration of today's most demanding AI tasks.

Enterprise-class support

Focus on research and development, not Linux system administration and hardware troubleshooting. Lambda takes care of the details, providing optional parts depots in your data center as well as on-site parts replacement services to minimize downtime.

Engineered for your workflow and workloads

With Lambda servers, you get the world's best compute hardware backed by the expertise of experienced AI engineers. You get a total system designed, optimized, and ready to use for your specific deep learning workloads.

Pre-installed with the software you need

Each Lambda machine is pre-installed with Lambda Stack, which includes everything you need to start training neural networks.

PyTorch · TensorFlow · Keras · NVIDIA

Lambda

Lambda Hyperplane 8-B200

GPU

8x NVIDIA B200 SXM GPUs with 180GB VRAM each NVLink with NVSwitch GPU-GPU interconnect

CPU

2x AMD EPYC or Intel Xeon Processors

- AMD EPYC 9555 (3.2 GHz, 64-core, 128-thread)
- Intel Xeon Platinum 8570 (2.1 GHz, 56-core, 112-thread)

MEMORY

AMD:

- 3072 GB (24 × 128 GB) DDR5

Intel:

- 3072 GB (32 × 96 GB) DDR5

POWER SUPPLY

Up to 6x n+n redundant 5250W 80 Plus Titanium power supplies

STORAGE

OS storage:

- AMD: 2× 3.84 TB NVMe onboard M.2
- Intel: 2× 3.84 TB NVMe onboard M.2

Extra storage:

- AMD: up to 8× 30.72 TB NVMe hot-swap U.2 bays
- Intel: up to 10× 30.72 TB NVMe hot-swap U.2 bays

NETWORKING

Default: 1x Dual Port 10GbE RJ45 (modular)
Storage: Up to 1x NVIDIA ConnectX-7 InfiniBand/Ethernet 400 Gb/s PCIe NIC (dual-port)
GPU clustering: Up to 8x NVIDIA ConnectX-7 InfiniBand 400 Gb/s PCIe NICs (single-port)

IPMI

IPMI 2.0 with virtual media over LAN and KVM-over-LAN support

FORM FACTOR

10U Rackmount with Rackmount Kit (assembly required)

INPUT/OUTPUT

1x RJ45 Dedicated IPMI LAN port
2x USB 3.0 ports
1x VGA connector
1x TPM 2.0 (optional - upon request)

SOFTWARE

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

DIMENSIONS

H 17.2" x W 17.6" D 33.2" (H 438.8mm x W 449mm x D 843.28mm)

COMPLIANCE

RoHS Compliant

TAA Compliant

PSU: 80 Plus Platinum, CE mark

SYSTEM POWER

200V and above required

Cables: 12x PWCD, US/EU/Canada/China/Australia, IEC60320 C20 to C19, 4FT

OPERATING CONDITIONS

10°C ~ 35°C (50°F ~ 95°F) and 8% to 90% humidity
Non-operating: -40°C to 60°C (-40°F to 140°F) and 5% to 95% humidity

WARRANTY & SUPPORT

Up to 5 years of hardware coverage, plus technical support from Lambda engineers