



Lambda Hyperplane 8-H200

8 NVIDIA® H200 Tensor Core GPUs with NVLink® & NVSwitch™

Configure your Hyperplane [↗](#)



TECHNICAL SPECIFICATIONS		
GPU DETAILS 8x NVIDIA H200 SXM5 141 GB GPUs with NVLink and NVSwitch	PROCESSOR 2x AMD EPYC™ or Intel Xeon® processors	SYSTEM RAM Up to 2304 GB DDR5
STORAGE Up to 16x 30.72 TB NVMe	NETWORK INTERFACE Up to 9x NVIDIA InfiniBand NDR 400 Gb/s NICs	

8 H200 SXM5 GPUs deliver 32 petaFLOPS of AI performance

Eight NVIDIA H200 GPUs interconnected with an NVLink and NVSwitch fabric allow the Hyperplane 8-H200 to serve up to 32 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.



Lambda Hyperplane 8-H200

GPU

8x NVIDIA H200 SXM5 GPUs with 141GB VRAM each NVLink with NVSwitch GPU-GPU interconnect

CPU

2x AMD EPYC or Intel Xeon Processors

- AMD EPYC 9654 (2.4 GHz, 96-core, 192 thread)
- Intel Xeon Platinum 8570 (2.1 GHz, 56-core, 112-thread)

*server specifications are the same for both AMD and Intel processors

MEMORY

AMD:

- 2304 GB (24 × 96 GB) DDR5

Intel:

- 2048 GB (32 × 64 GB) DDR5

POWER SUPPLY

Up to 6x n+n redundant 3000W 80 Plus Platinum power supplies

STORAGE

OS storage:

- AMD: 1× 3.84 TB NVMe onboard M.2
- Intel: 2× 3.84 TB NVMe onboard M.2

Extra storage: up to 16× 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

8U 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 or Red Hat Enterprise Linux

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

DIMENSIONS

H 14" x W 17.2" x D 31.5" (H 356mm x W 437mm x D 800mm)

COMPLIANCE

RoHS Compliant

TAA Compliant

PSU: 80 Plus Titanium, CE mark

SYSTEM POWER

200V and above required

Cables: 6x PWCD, US/EU/Canada/China/Australia, IEC60320 C14 TO C13, 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