

A catalog, not a SKU.

Five reasons neoclouds choose AMD Instinct™ MI350P with Liquid. One model per GPU means one sellable unit per GPU, on one server.

144 GB

HBM3E per
MI350P

4.32 TB

HBM3E on one
server

30 units

sellable, one
model per GPU

0 fees

per-GPU software
licensing



LIQID GPU chassis
4 to 30 MI350P over
PCIe Gen 5

One host
Dual-socket AMD
EPYC™ 9005

In your jurisdiction
Tenant data stays put

Five Reasons for Neoclouds

You sell GPU-hours, and GPU-hours are a commodity. Airlines don't sell seats by the pound; they sell fare classes on one aircraft.

1

Memory per card your competitors can't match in a PCIe slot

144 GB of HBM3E per GPU with about 4 TB/s of bandwidth. Memory-hungry inference customers pay a premium, and memory is what you sell them.

Up to 30 GPUs on one host: 4.32 TB of HBM3E

2

One model per GPU is one sellable unit per GPU

Each GPU runs its own model: a catalog tier, or a tenant's own fine-tuned model. Isolation is easy to explain: one tenant, one card.

20 GPUs = up to 20 dedicated tenant models

3

No per-GPU software licensing against your margin

The AMD enterprise AI stack is open source. Licensing fees leave the bill of materials instead of scaling with your GPU count.

ROCm™ · AMD Inference Microservices · no license per GPU

4

Better tokenomics is your floor; tiers are your ceiling

Fewer servers means fewer licenses, less networking and less rack space, so your cost per token falls. What you charge on top is a product decision, not a price war.

Tokens per dollar is gross margin

5

A second source is a supply strategy

Operators adding a second GPU source treat it as a differentiator, not a hedge: capacity you can add now, on an architecture you control.

Available now, from 4 to 30 GPUs

Not the fit for large-scale distributed training: this is a scale-up architecture for inference. Bring us your tenant pipeline and tier mix, and we will model your cost floor.

Built for

Neoclouds · Cloud service providers · Sovereign cloud · Managed service providers

Proven on

AMD Instinct™ MI350P · AMD EPYC™ 9005 · ROCm™ · 200+ Liquid deployments

Size Your Catalog

sales@liquid.com · +1 303.500.1551