



**How to complete
your GPU cloud
without building
storage from scratch.**


Protect Your Margins. Increase Your Revenue.

Your customers need more than GPU capacity.



Keep data close to compute

Deploy persistent object storage alongside your GPU infrastructure. Customer AI data stays on your platform, not a hyperscaler.



Offer storage under your brand

S3-compatible storage through your own domain and customer endpoints. No "powered by" badge. No third-party branding.



Increased storage revenue

Storage that previously flowed to Amazon is now billed under your brand. Turn a cost pass-through into a revenue line.



Protect your margins

Flat pricing, zero egress fees, zero storage headcount. Earn healthy margins on a product your customers already need.

You rent the GPUs.

Someone else is billing for the storage next to them.

Storage lives on someone else’s cloud. Compute lives on yours.

CURRENT HYPERSCALER STACK

Customer data

Stored on a hyperscaler

aws

Egress charge every read

↓

Your GPU cluster

Reads run against remote storage

Cost.

Moving 100 TB to a GPU cluster runs roughly \$8,000–9,200 in egress per pass. That bill lands on the customer, not on you.

NEW EDGE STACK

Customer data

Stored on your platform

AKAVE

CLOUD

Zero egress

↓


Your GPU cluster

Reads stay inside your platform

Revenue.

Storage becomes a native line on your invoice, billed under your brand. Reads and rewrites do not compound the bill.

Your customers need persistent storage for datasets, model weights, checkpoints, logs, and outputs that remain available across workloads. Pick the layer where you can build a durable advantage. Don’t dilute engineering resources trying to compete across two layers of the infrastructure stack. Akave gives you the storage technology and operational support to add that layer to your platform.



1. GPU Compute

Training, inference, and AI workloads

↓



2. Hot Working Layer

Local NVMe, shared files, and active workload data

↓



3. Akave Object Storage

Persistent datasets, model weights, checkpoints, logs, and outputs

FOR NEOCLOUDS • CYBER RESILIENT • SOVEREIGN STORAGE • COMPUTE AGNOSTIC FOR AI WORKLOADS

Sovereign object storage built for a secure AI infrastructure.

Sovereign object store, built for secure AI.
Enabling data lakes to keep AI datasets on-platform, close to the NeoCloud.

FEATURE FOCUS

Akave Managed Storage
Add a branded storage service to your own cloud environment.

- Custom domain and SSL
- Tenant-isolated endpoints
- Shared, dedicated, or partner-hosted infrastructure
- Managed deployment and operations
- Wholesale, bespoke pricing, healthy margins

01

Isolated Endpoints Per Customer
No shared namespaces. No shared data paths. Isolation is independently auditable per tenant, so your customers can confirm it without relying on Akave's word.

02

Custom White-Label Options
Separate endpoints, security keys, and policies by customer, with tenant operations recorded through Akave's logs.

03

Flexible Deployment
Shared, dedicated, or partner hosted infrastructure placed around regional, performance, and sovereignty requirements. Scalable from terabytes to petabytes.

04

Audit-Native Data Protection
Erasure coding, Object Lock, cryptographic integrity proofs, and an independently auditable record of storage operations.

When storage lives on your platform, so does your customer.

Platform-native storage is one of the strongest retention signals in enterprise AI contracts.



Separate endpoints, keys, policies, and audit records by customer.

Storage that used to cost you money now earns it.

Akave gives NeoClouds a managed, S3-compatible object-storage layers for persistent AI data. Deploy it alongside your compute infrastructure, offer it under your own brand, and keep more customer data on-platform.

WHAT OPERATORS USING AKAVE CAPTURE

- ✦ **Storage margin, under your brand.**
Bill storage as a native platform service, not a pass-through.
 - ✦ **Isolated endpoints per tenant.**
One isolated S3-compatible endpoint per customer.
 - ✦ **Ideal for sovereign industries.**
Customer-held keys (BYOK/HYOK), regulated-sector ready.
- ✦ **White-label deployment.**
Custom domain, custom SSL certificate, customers see your brand.
 - ✦ **Zero egress.**
Move data between storage and compute as often as your workloads demand, at no charge.
 - ✦ **Managed by Akave.**
No storage ops, no engineering overhead on your side.

Enabling data lakes that keep AI datasets on-platform.

The most valuable AI workloads create data that must be retained, reused, protected, and accessed across compute jobs.

- ✦ **VM AND CONTAINER SNAPSHOTS**
Preserve environments and restore workloads without rebuilding from scratch.
- ✦ **MODEL WEIGHTS AND CHECKPOINTS**
Keep versioned training artifacts available across jobs, clusters, and teams.
- ✦ **GEN AI DATA**
Store outputs and logs from inferencing and fine-tuning pipelines.
- ✦ **REGULATED SOURCE DATA**
Maintain auditable data for sovereign and compliance-sensitive environments.



Neoclouds



Pricing



Talk to Sales



Free Trial