

Wasabi for AI Storage

Affordable, secure, predictable cloud object storage for AI

Wasabi delivers hot, S3-compatible object storage at roughly 80% less than the hyperscalers, with zero egress fees and zero application programming interface (API) request charges—so AI teams can focus on project success instead of managing mounting storage costs.

AI processing is read- and iteration-heavy. Training datasets, checkpoints, embeddings, model artifacts, and inference logs are pulled, moved, and reprocessed constantly. On hyperscaler storage, every move is metered.

Wasabi gives AI teams a consistent storage layer across the entire data lifecycle, from ingest and training data through checkpoints, model artifacts, and long-term retention, with built-in security and immutability, and flat predictable pricing.

In [Wasabi's 2026 Global Cloud Storage Index](#) (a Vanson Bourne survey of roughly 1,700 IT decision makers), 32% of organizations report positive AI ROI today, while 51% expect to get there within twelve months; closing that gap is exactly where storage economics comes in.

Affordable: storage economics built for AI scale

The two costs that quietly drain AI storage budgets (paying to read your data and paying to move it) simply aren't on the Wasabi storage bill. As datasets, checkpoints, and model versions multiply, your cost scales with capacity, not with how hard your team works the data. That savings can go straight into what moves AI forward: graphics processing units (GPUs), talent, and innovation.

Predictable: one flat rate, no surprise bills

Wasabi uses flat, published per-TB pricing. Because there are no egress or API call charges, the number you model at the start of a project is the number you can budget for, even when retraining, frequency spikes or inference traffic grows. Teams read, move, and reprocess data as often as the work demands without having to meter anxiety over the cost, and finance gets a line item it can plan against across every environment and at any scale.

Features

- Flat, predictable pricing with no data access or API call fees
- Built-in immutability (Object Lock) and 11 9's of object durability
- Data encryption and access controls with MFA and Multi-User Authorization
- Amazon S3-compatible REST API

Benefits

- Cyber-resilient storage for data and models; reduce risk from ransomware and accidental deletion
- Predictable costs across environments and as data volumes scale
- Fast onboarding with S3-compatible tools; reuse existing tools and workflows
- Enables flexible computing and GPU selection (cost, availability, locality) based on workload needs
- Supports long retention of data locality, sovereignty, governance and audit workflows
- Future-ready architecture without vendor lock-in

Secure: cyber-resilient by design

Object Lock immutability, encryption, Multi-Factor Authentication (MFA), and Multi-User Authorization protect your training data and model weights from ransomware, corruption, and accidental deletion by default. Eleven nines (99.999999999%) durability and an expanding regional footprint support the data residency, sovereignty, retention, and audit requirements that are increasingly first-screen criteria for AI built on regulated data.

Open and compute-neutral: integrates with your entire AI stack without lock-in

Wasabi speaks the same S3 API that every machine learning (ML) framework, lakehouse format, and MLOps tool already uses, so adoption is an endpoint swap rather than a re-architecture.

Wasabi keeps your data open and portable across any compute; it is the one layer of your stack that doesn't compete with anything you run.

Zero egress means you can train wherever the GPUs are—on a hyperscaler, a neocloud, or your own cluster—so storage stops dictating where you build and your data stays portable as the AI landscape shifts.

The result: the potentially most volatile line item in your AI budget becomes the one you never have to worry about, because eliminating fees is what makes Wasabi's storage economics work at scale.

AI teams already building on Wasabi include:

AI RoA (AI robotics foundation models): Wasabi powers a 2.2 petabyte (PB) shared data hub for robot model development, projected to save AI RoA more than 100 million Japanese yen (JPY) versus alternative vendors—and with no data transfer or API fees.

Performance Hub / UBX (AI video surveillance, global gym network): with Wasabi's flat-rate pricing, Performance Hub / UBX cut cloud storage costs by 4 to 6 times and customer-billing calculation costs by 20 times versus its previous provider.

Black.AI (computer vision on CCTV video): Wasabi gives Black.AI a platform for PB-scale video and AI training data, with predictable pricing, no data loss, and the regional coverage to meet its customers' data-sovereignty rules.