

DON'T LET STORAGE BE YOUR AI BLINDSPOT

S3-compatible object storage engineered for the speed and scale of AI training.



Why teams choose AI storage from Impossible Cloud

Most cloud storage was never designed for the metadata-intensive, checkpoint-heavy reality of AI training. Slow operations, unpredictable egress costs, and API gaps push storage onto the critical path.

Impossible Cloud changes that. S3-compatible object storage engineered specifically for AI training workloads. Built to keep pipelines fast and predictable — regardless of model architecture or modality. With native LanceDB support, industry-leading metadata performance, and flat pricing.

What makes us different

01 Built for parallel training pipelines

Multi-writer support without external lock services or hyperscaler dependencies. Native LanceDB compatibility included.

02 High-volume metadata performance

Up to 1,000 bulk DELETES per millisecond. Built for the object operations behind checkpoints, datasets, and training runs.

03 Flat pricing. No egress surprises.

No per-request fees on API calls. No egress fees up to your monthly upload volume.

Everything you need to make the switch



S3-compatible API

No per-request fees on API calls. No egress fees up to your monthly upload volume.



Multilayer security

Encryption at rest and in transit. Object lock and ransomware protection.



Sovereign by design

Geofenced across seven European regions. GDPR-compliant, ISO27001 certified



Proven performance

Matches AWS bulk DELETE throughput, outperforms Wasabi by 314 ms.



Predictable pricing

No tiering, no per-request fees, no egress fees up to your monthly upload volume.



Easy migration

Full migration support included. Most teams are live within two weeks.

AI changed what infrastructure teams need from storage. That's why we built Impossible Cloud for what comes next.

Christian Kaul, PhD — CEO & Co-Founder, Impossible Cloud



WHO THIS IS FOR

Built for the workloads driving modern AI

From foundation model training to vector-native pipelines, Impossible Cloud delivers the performance and predictability modern AI infrastructure requires.



LLM pre-training and fine -tuning

Petabyte-scale storage for training and fine-tuning frontier models on proprietary datasets.



Image and video diffusion models

Stable, high-throughput storage for large-scale image and video generation workloads.



Multimodal vision-language models

Single-bucket storage for mixed modalities: document AI, medical imaging, and vision tasks.



Embedding training and vector indexes

Native LanceDB support for RAG pipelines and vector indexes. No DynamoDB required.



Speech, audio and music models

High-throughput sequential reads for audio, speech and music training at any scale.



Bring-your-own-model architectures

Standard S3 API. Works with any model architecture, no proprietary tooling required.

Our pricing model

PREDICTABLE, FLAT-RATE PRICING

Storage	Pay-as-you-go or lower on 1, 3 or 5-year contracts	€7.99/TB/MONTH
API operations	HEAD, LIST, DELETE, conditional writes	INCLUDED
Egress	Up to your monthly upload volume	INCLUDED
	Egress exceeding monthly upload volume	€7.99/TB
Tiered pricing		NONE
Minimum retention time		24 HOURS

LET'S GET STARTED

Book a meeting



Contact us

impossiblecloud.com
info@impossiblecloud.com



Impossible Cloud GmbH
Friesenweg 12, Haus 5 · 22763 Hamburg (Germany)