

Captions

AI Video Platform Speeds Training and Cuts Cloud Costs



O

GPUs Waiting
for Data

Situation

Captions is pioneering foundation models for AI video. With rapid growth and increasingly large model runs, the company needed to stream massive training datasets from object storage to globally distributed compute environments. High egress fees and inflexible region constraints from a prior provider created inefficiencies for a team moving toward training with H100 clusters.

95%

Estimated Cost
Savings

Solution

After a competitive bakeoff with other providers, Backblaze B2 Overdrive won out for its superior throughput, cost predictability, and flexibility. Key drivers included direct peering with GPU compute partners for fast, zero-cost transfers; free cloud replication for multi-region caching of training sets for geographically distributed compute, and white glove support with real-time help and fast response times during testing and production.

\$0

Egress

Result

In production, Captions achieved a technical win: high-speed data transfer, seamless integration via the S3 compatible API, cost projections showing up to 95% savings compared to their previous solution, and performance that exceeded expectations with better pricing predictability and support.

“Backblaze B2 Overdrive delivered the speed and flexibility we needed—with unlimited egress and real human support. The performance and cost savings were clear.”

— **Drew Jaegle**, Head of AI, Captions