MINIO

Building AI-Ready Infrastructure: Storage, Performance & Resilience at Scale

Murat Karslioglu
Head of Global Field Architects / Field CTOs
@muratkarslioglu
November 2025

AI Storage is Object Storage

MINIO

Every AI Leader Uses Object Storage as its Primary Storage for Building LLMs

MINIO

AI storage is the new bottleneck.

Not GPUs. Not compute. Storage.

Object Native Advantages are Required at AI Scale

Single-Namespace Scalability

Store **raw data**, **tokenized corpora**, and **model checkpoints**, easily reaching 10s to 1,000s of petabytes, all in a single global namespace.

Performance for Unstructured Data

Efficiently train using thousands of GPUs with **massive parallel throughput across both large *and* small files** such as txt, JSON, parquet, audio, image, video etc.

Cost Effectiveness

Object-native storage persists data more efficiently than file storage due to its **flat namespace and software-defined** use of non-proprietary hardware.

Deep Ecosystem Support

All leading AI/ML tools such as **Hugging Face**, **PyTorch**, **TensorFlow**, **Ray**, **Kubeflow** natively support S3/object storage APIs.

Versioning & Immutability

Easily **reproduce training runs and rollback models or data**, via safe, fine-grained continuous data protection powered by object immutability and versioning.

Operational Efficiency

Object-native storage has **no file system hierarchy or metadata bottlenecks** making it faster and easier to manage, automate, and monitor at AI scale.

The CSP Challenge: Cost & control barriers too high for most enterprises

Architecture Matters

The On-Prem Choice: Legacy Retrofit vs. Object-Native

Retrofit Architecture (Legacy Enterprise)

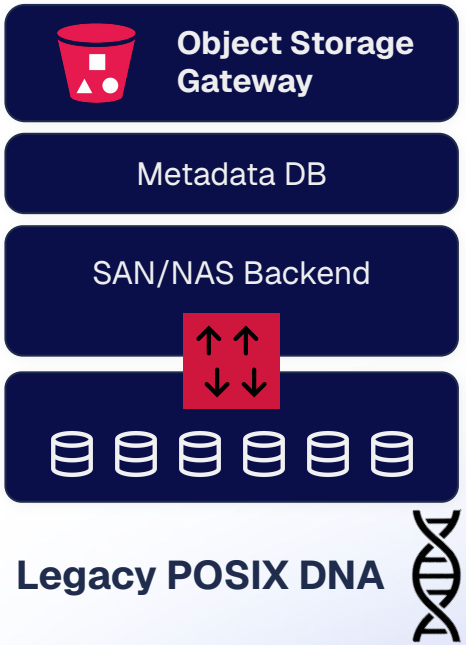

VS.

Object-Native Architecture (AI Leaders)

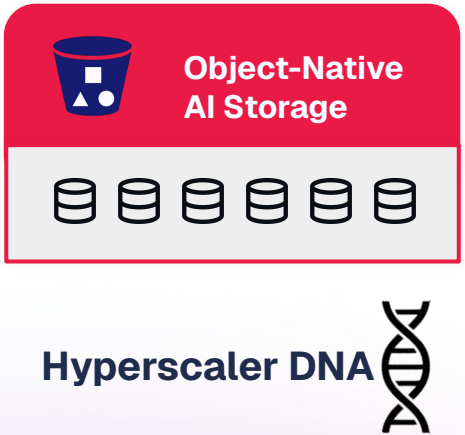

MinIO is Software-Defined Object Storage Leader

MinIO delivers the world's fastest and most cost-efficient AI storage at any scale

The World's Enterprises Run on MinIO

52% of the **Fortune 500**

77% of the **Fortune 100**

9 of the 10 Largest **automotive**

All 10 of the Largest **US banks**

8 of the 10 Largest **US retailers**

The World's Developers Build with MinIO

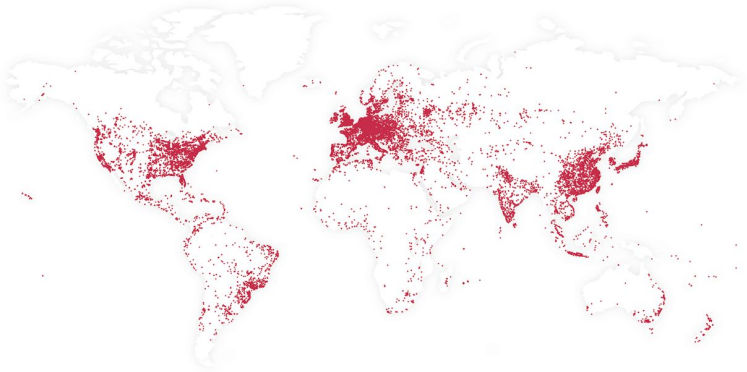

2B+ Downloads

51.4K+ GitHub Stars

32.5K+ Community Members

1.4K+ Developers

MINIO AISTOR

AI Storage for the Data-Driven Enterprise

MINIO AISTOR

Platform Foundations

- Administration
- Data Management
- Data Protection
- Replication & Security

API & Protocols

- S3
- GPUDirect
- S3 Express
- S3 over RDMA
- SFTP
- MCP

High Performance Object Storage

Intelligent Ecosystem

- Observability
- AI & Extensibility
- Automation
- Integrations

MINIO SUBNET — Direct to engineer support across the full stack, Applications to H/W

Gartner Peer Insights Customers' Choice 2023

DATA Breakthrough Award 2023 — Object Storage Solution of the Year

Datanami Readers Choice Award 2022 2023 2024

TechTarget Storage Product of the Year 2022 2023

GigaOm Radar Object Storage 2019 2020 2021 2022 2023

GigaOm Radar High Performance Object Storage 2021 2023

Redpoint Infra Red 2023 2024

Storage Technology Award 2021 2022 2024

DATAwire Reader's Choice Award 2021 2022 2024

Why AIStor + Red Hat OpenShift = AI-Ready Platform

- Certified on OpenShift operators
- Works natively with Ray, Kubeflow, Spark on OpenShift
- Fully software-defined > runs on any OpenShift bare-metal hardware
- AIStor = S3 backbone for OpenShift Data Science (ODH)

MinIO Kubernetes Deployment (OpenShift)

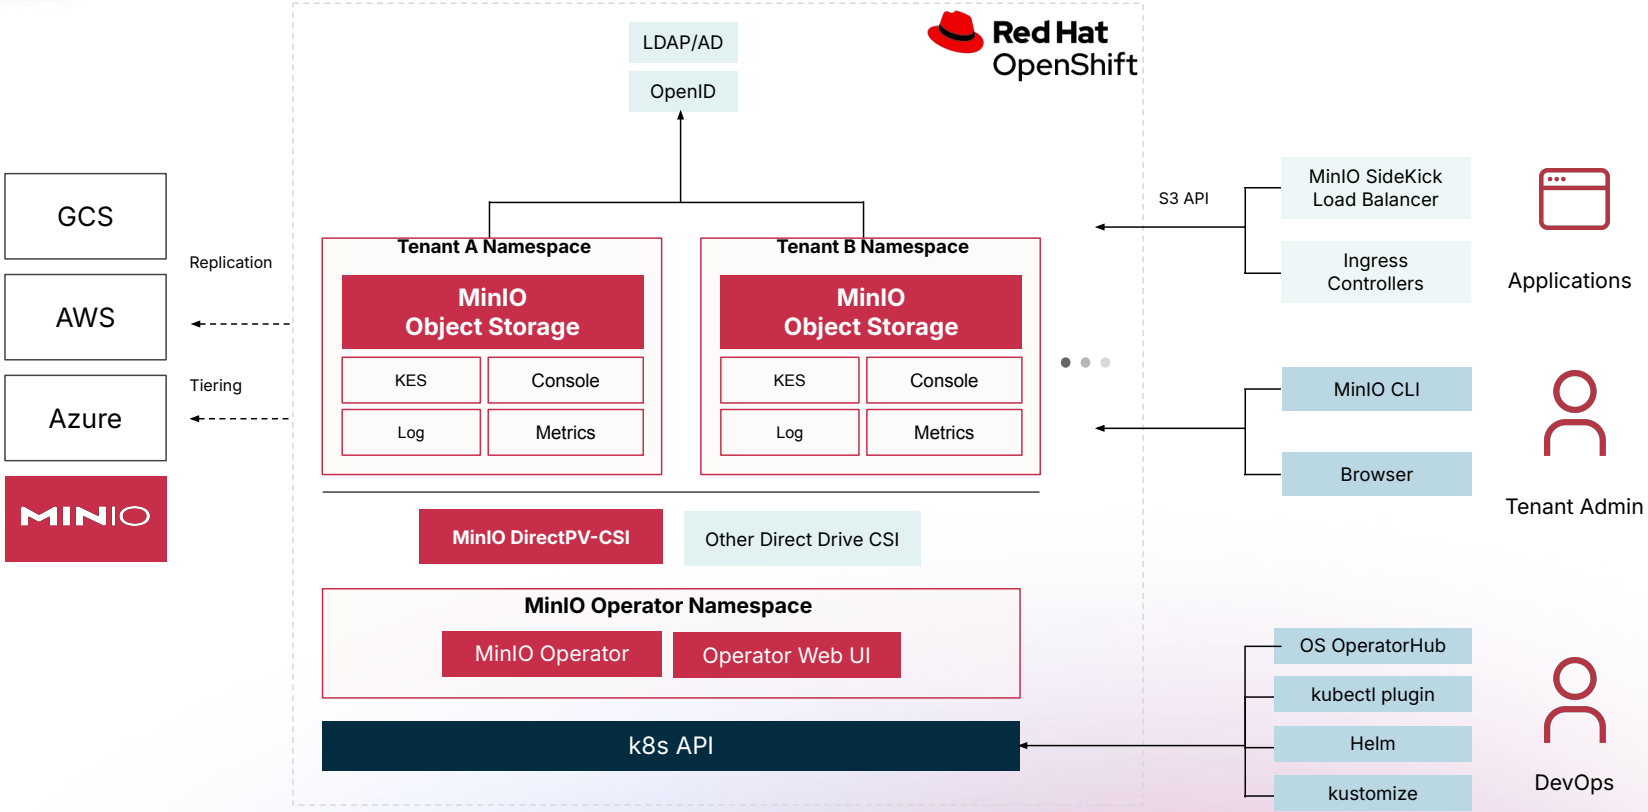

Why AIStor + Red Hat OpenShift = AI-Ready Platform

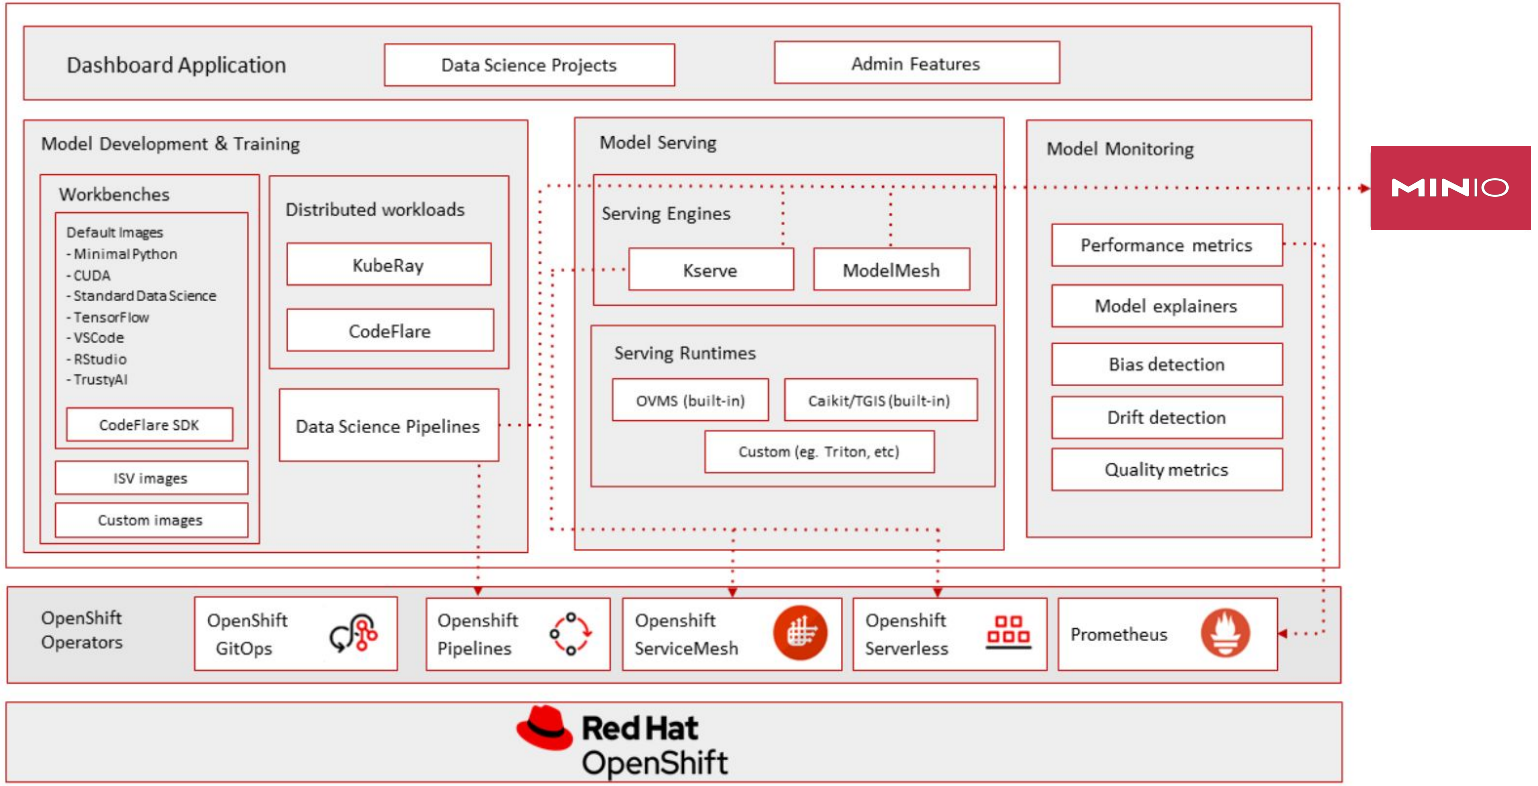

10

AIStor Architecture is Object-Native

Strictly Consistent

Guaranteed read-after-write and list-after-write correctness.
DAS O-DIRECT (direct I/O).
All metadata updates are atomic and strictly serialized.
No delayed propagation cache.
Erasure coding with strict quorum enforcement.

Single Layer, Gateway-Free

Simple, efficient topography helps deliver the lowest TCO.
Linear, per-node scalability.
No protocol gateways, translations, or overhead.
No hierarchy, flat namespace.
No application/OS tuning.

Metadata DB-Free

Metadata embedded with objects for true scale-out.
Consistent low latency regardless of file size & object count.
High obj/s across billions of samples, billions of embeddings.

SIMD Accelerated

GO + assembly optimized SIMD extensions, e.g. AVX512, NEON, VSX.
Speeds MD5, SHA256, Highway Hash bitrot protection, erasure coding for x86 and Arm.
Enables throughput and IOPS saturation of any storage H/W, even through H/W failure.

Data Mobile

Batch replication for multi-cloud data movement.
Strictly consistent, synchronous active-active replication.
Objects, metadata, tags & locks.
Common source and destination bucket names.

Inline Only

All data services are inline and high performance to eliminate complexity and variability.
Object-level encryption, compression, erasure coding, replication and more.

Beyond S3

Full S3 compatibility (baseline), including S3 Express.
PromptObject to talk to objects as if they were an LLM.
Extension to list, stat and download contents of ZIP files in any bucket.

Object-Native AI Storage

The Data Foundation for the AI-Ready Stack

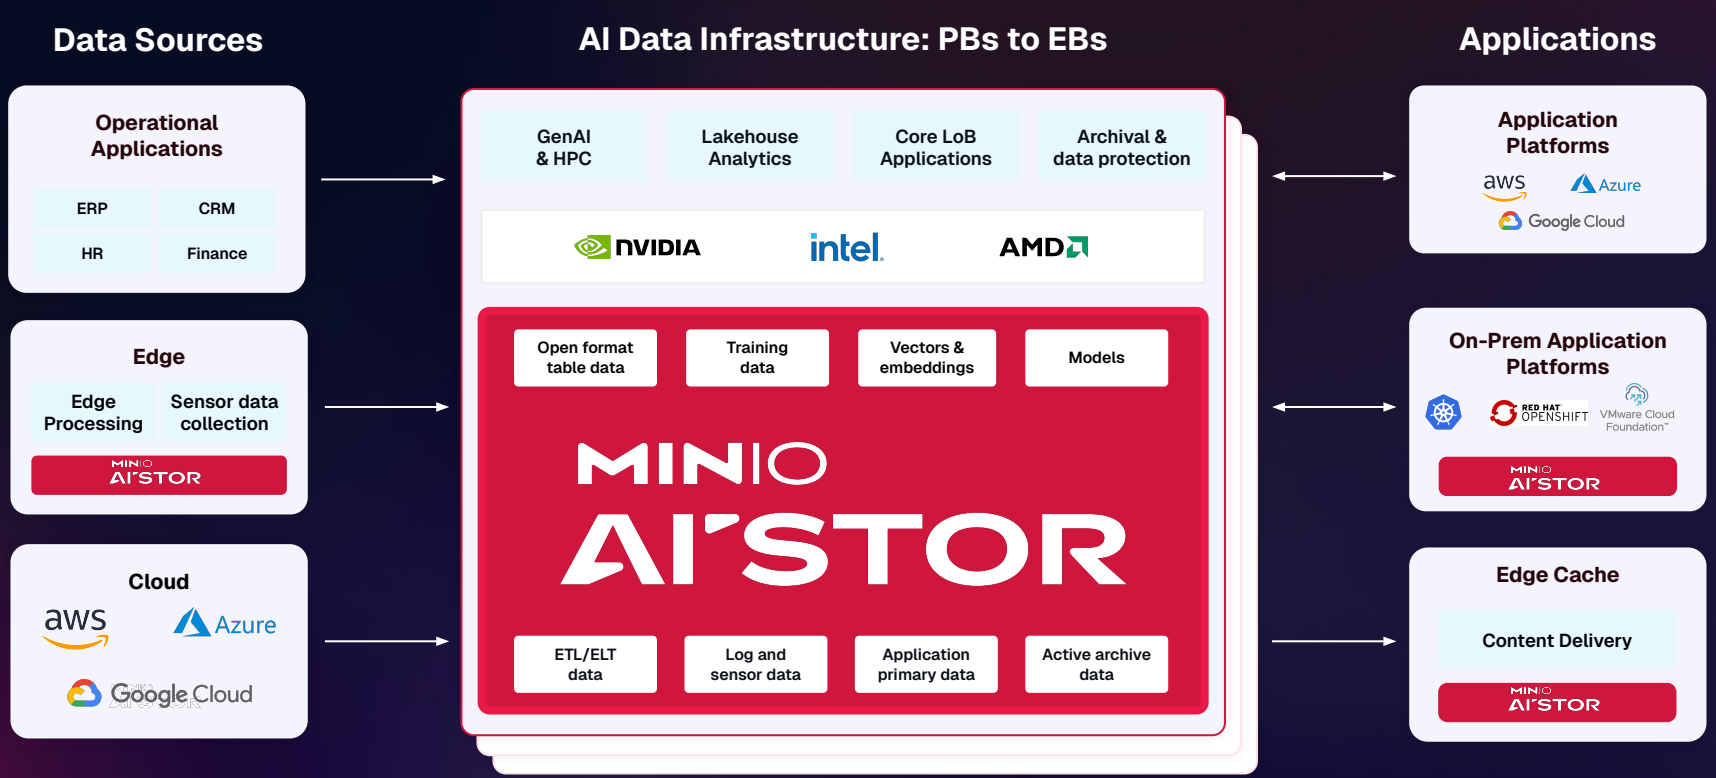

AIStor Multi-Site Replication: Active/Active, Sync, Async

Production Customer Success at AI Scale

MINIO

Generative AI and HPC	Data Lakehouse Management and Analytics	Object-Native Core Business Applications
TESLA	CROWDSTRIKE	PayPal
Recursion	AMERICAN EXPRESS	Walt Disney

10s to 1,000s of petabytes of data, tables, vector embeddings, and models.

Production Customer Success at AI Scale

Generative AI and HPC

Data Lakehouse Management and Analytics

Object-Native Core Business Applications

10s to 1,000s of petabytes of data, tables, vector embeddings, and models.

Why AIStor: SCALE

AIStor Capacity Scales From TBs to EBs In a Single Namespace

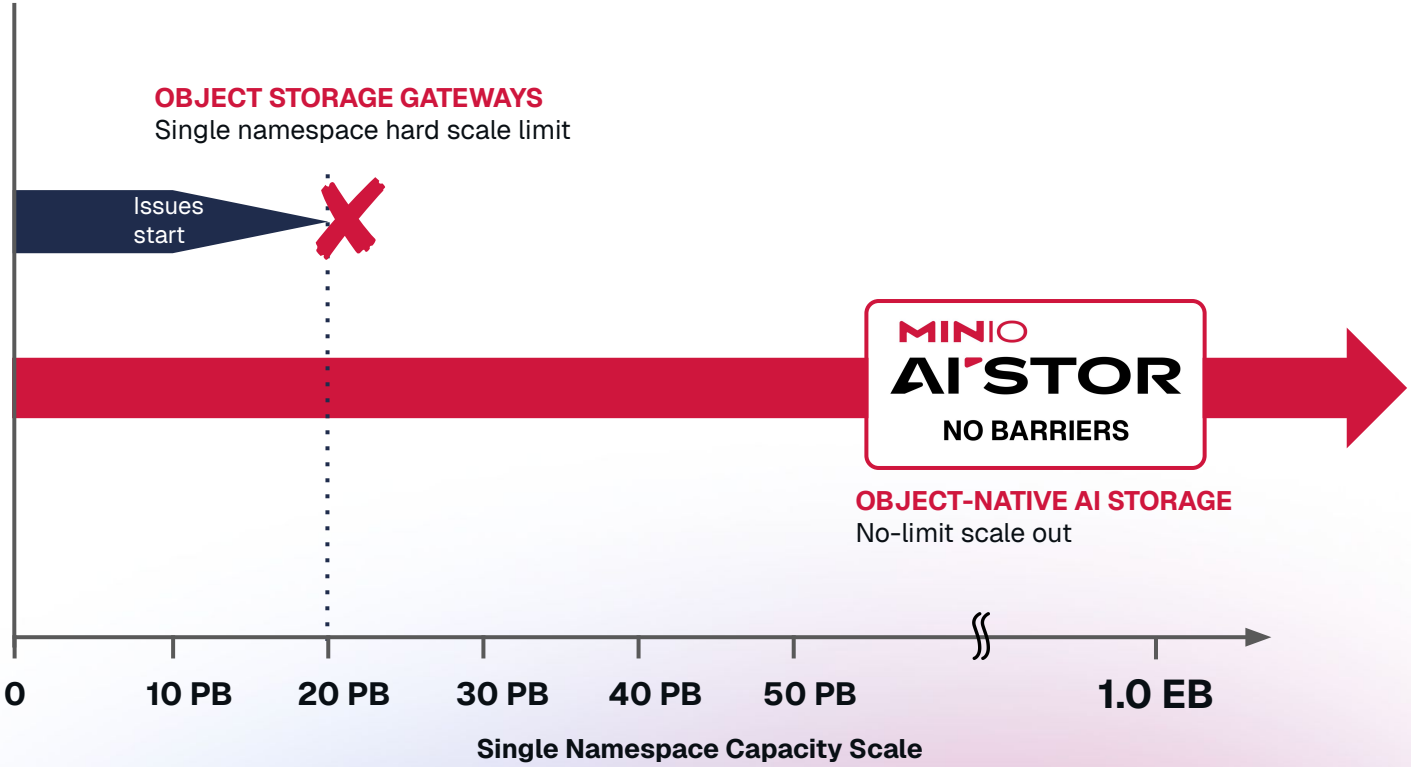

Why AIStor: PERFORMANCE

AIStor Performance Scales Linearly With Capacity, In a Single Namespace

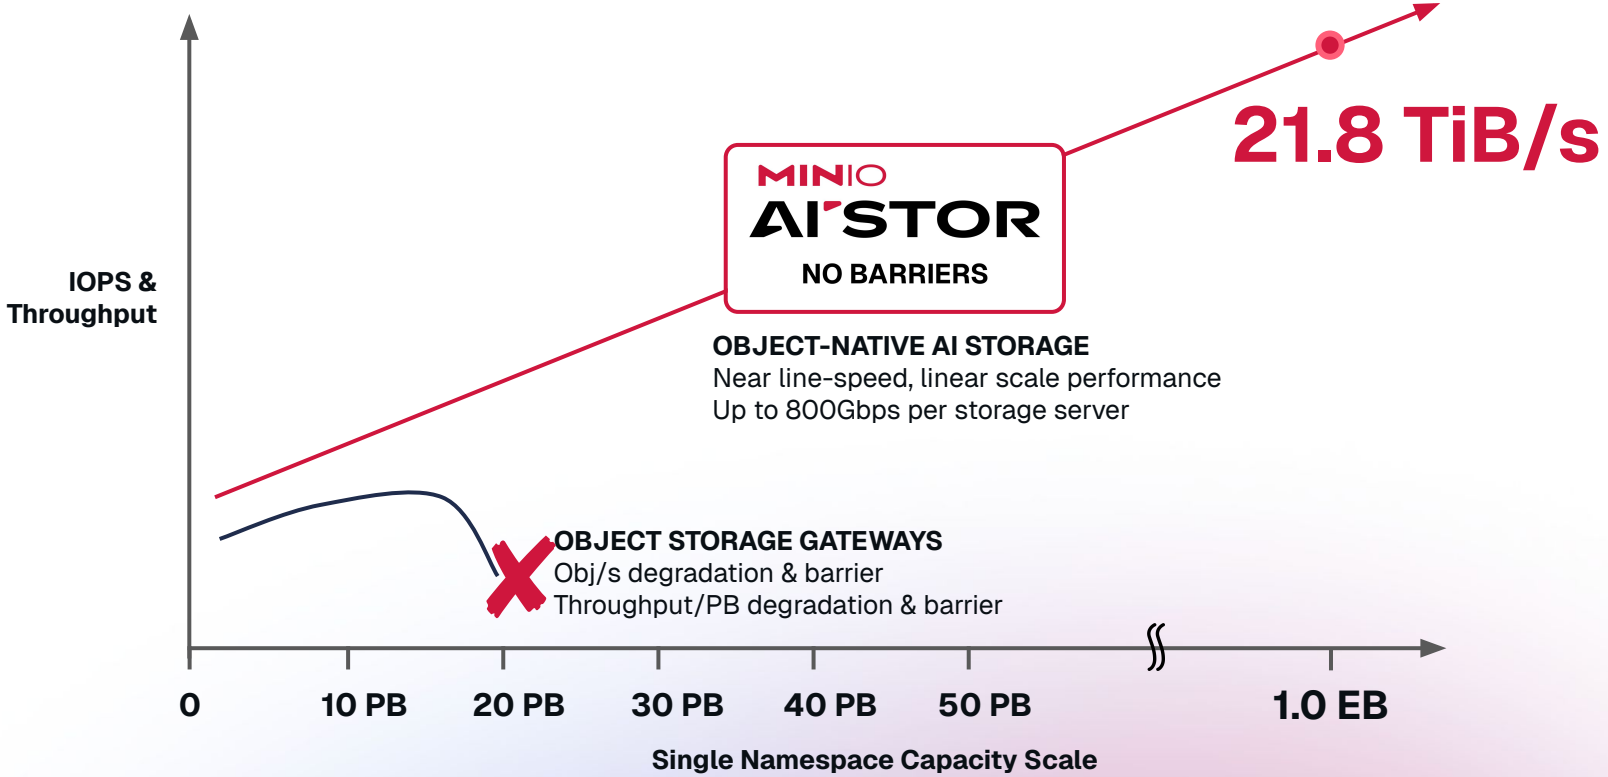

Why AIStor: ECONOMICS

AIStor Delivers Industry-Leading Unit Economics at Every Scale

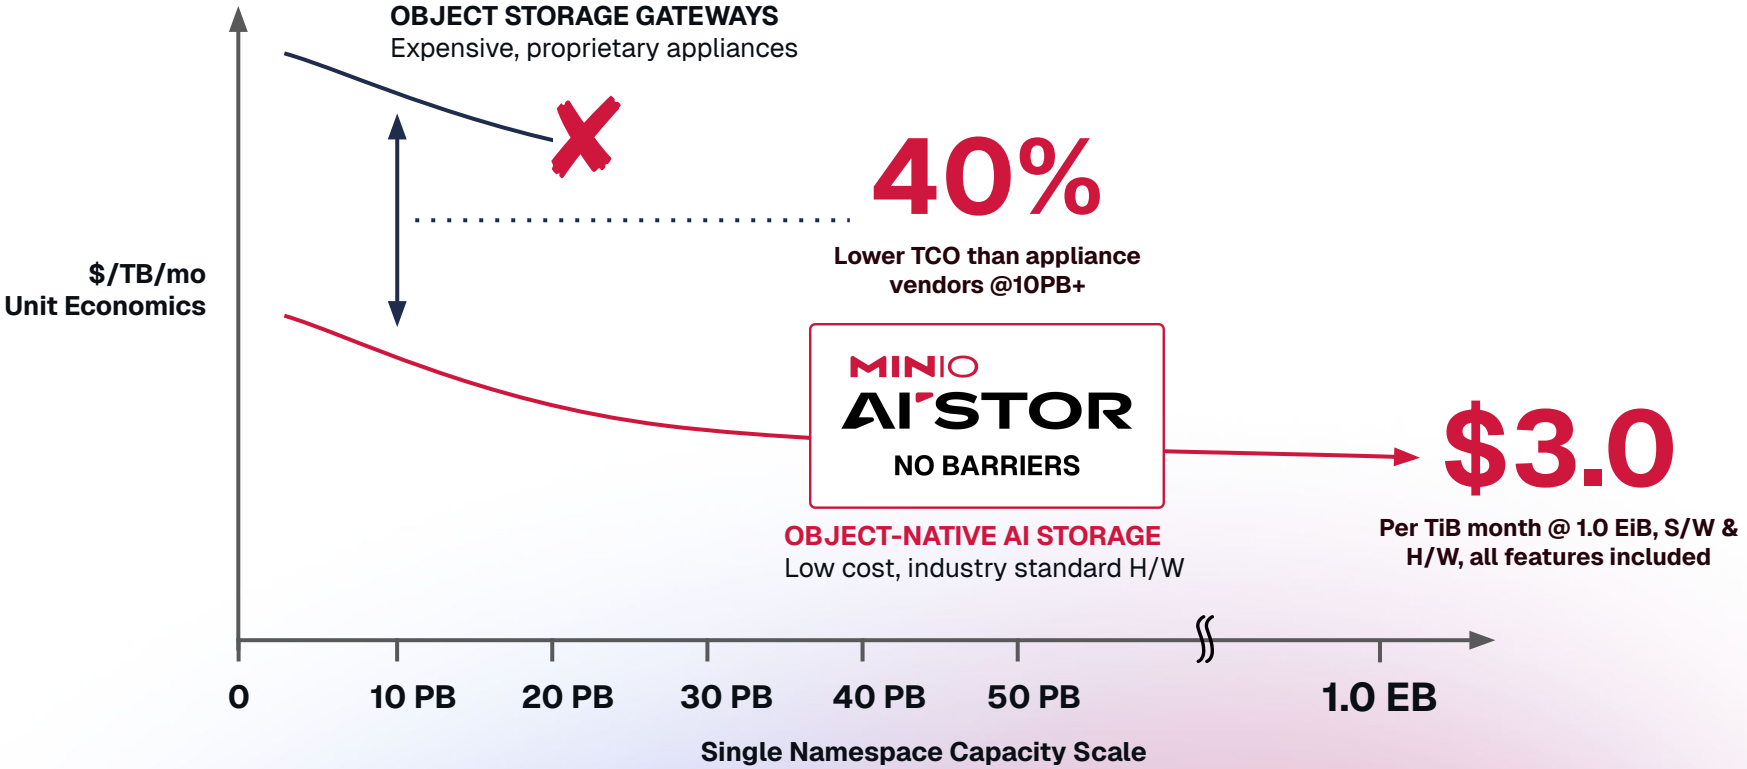

Why AIStor: ECOSYSTEM

AIStor's S3 API Compatibility Eliminates Friction and Risk to Speed Time to Production

OBJECT STORAGE GATEWAYS

Partial S3 API compatibility with low maturity

% of S3 API Supported

"It might work"

vs.

OBJECT-NATIVE AI STORAGE

Full S3 API is hardened and broadly validated

"It just works"

Bare Metal Reference Architecture

MINIO AISTOR NO BARRIERS

Storage Server Reference

	Dell	HPE	Supermicro
1U	 PowerEdge R7615	 ProLiant DL325	 ASG-1115S-NE316R
2U	 PowerEdge R7715	 ProLiant DL345	 ASG-2115S-NE332R

Single Node Configuration:

- **CPU** - Single AMD EPYC™ 9655P 96 cores or higher
- **RAM** - DDR5 RDIMM-6400
  394GB or higher (12×32 GB) (1U Medium) or
  786GB or higher (12×64 GB) (2U Large)
- **NIC** - 2×400GbE or 2×100GbE
- **STORAGE** - E3.S NVMe SSD
  16x E3.S NVMe SSD 15.36TiB (1U Medium) or
  32x E3.S NVMe SSD 61.44TiB (2U Large)

Example: 2U Large Config: (~9.6 PiB useable)

- **8×2U nodes (16U)**
- **Erasure Code Stripe Size (K+M) - 8**
- **Erasure Code Parity (M) - 3**
  - Drive Failure Tolerance: 96
  - Server Failure(s) Tolerance: 3

NOTE: For sizing and throughput estimates, please see the MinIO sizing calculator

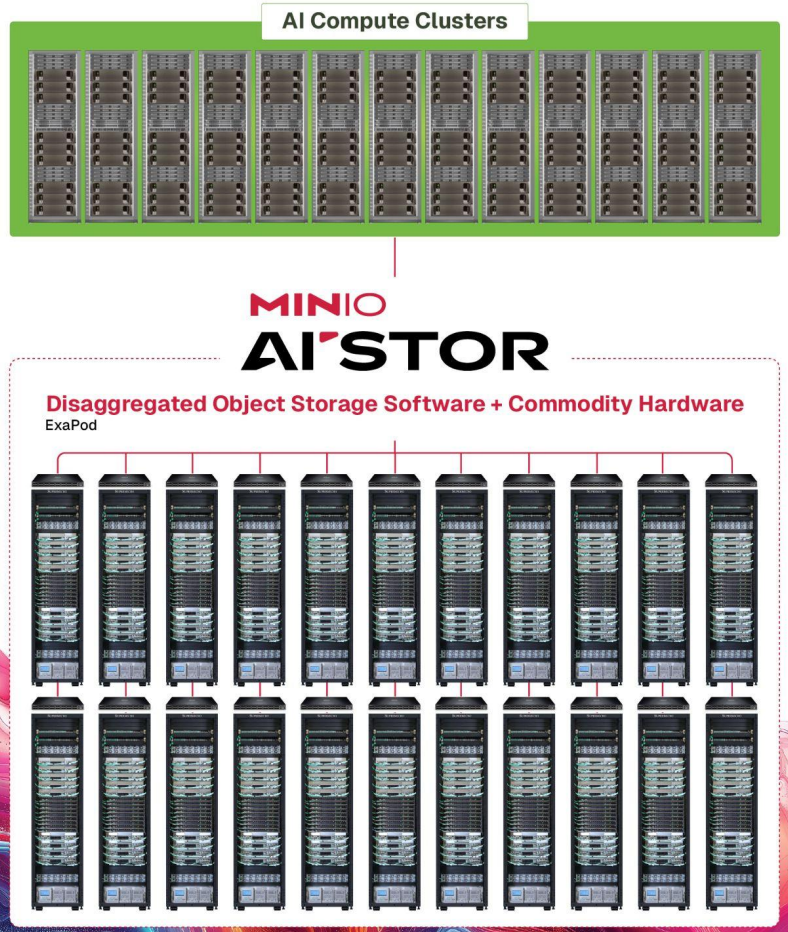

The Only Solution for Exascale AI: ExaPod

MINIO AI'STOR NO BARRIERS

1.0 EiB
Single namespace capacity

640
Storage servers

19.2TiB/s
Read throughput

$4.6/TiB
Per month, S/W & H/W, all features included

Hardware Trends

Dense NVMe

- E3.L/E3.S/U.2 QLC > E2 (122.88-245.76+ TB) > 1PiB+ drives
- Front-service, better thermals, fewer chassis

Network First

- 2×100-800 GbE per node standard
- Non-blocking leaf-spine, jumbo frames

Right-Sizing

- Spend on NICs/NVMe before extra CPUs
- EC 12:4/16:4 vs 3× replication = 2× capacity

What to Measure

- $/TiB-usable, watts/TiB, throughput per rack
- GPU util >90%, p99 GET latency

Thank you

@minio | min.io
@muratkarslioglu