STRIKE AI GOLD WITH SCALABLE, EFFICIENT DATA LAKES & LAKEHOUSES

Introducing 4 storage essentials for enterprise AI success

When AI models are finely tuned with all your enterprise data, the potential for better decisions and innovation goes through the roof. But without the right infrastructure, intelligence is lost.

Explore how to right-size data lakes (for aggregating enterprise data) and lakehouses (for running analytics) to power success in the AI era.

THE CHALLENGES OF ENTERPRISE AI WORKLOADS

Enterprise AI is all about unlocking real-time insights from massive amounts of data. But it can be complex to put prototypes into production. To enable high performance at scale, you need:

- **Robust data integration** — Enterprise AI requires fast access to diverse data (transactional data, chat logs, images, sensor data).
- **Modern storage & governance** — High data quality, consistency, and reliability are crucial for model fine-tuning.
- **Flexible architectures** — Modular designs let you scale storage and compute independently, making AI more affordable and sustainable.

HOW FUTURE-READY INFRASTRUCTURE CAN HELP

AMD has joined Supermicro, MinIO, and EDB to build resilient storage that can power enterprise AI workloads at scale. The pre-validated solutions combine:

- Rack-scale storage systems optimized for enterprise AI
- High-performance object storage for massive data lakes
- Next-generation systems and CPUs
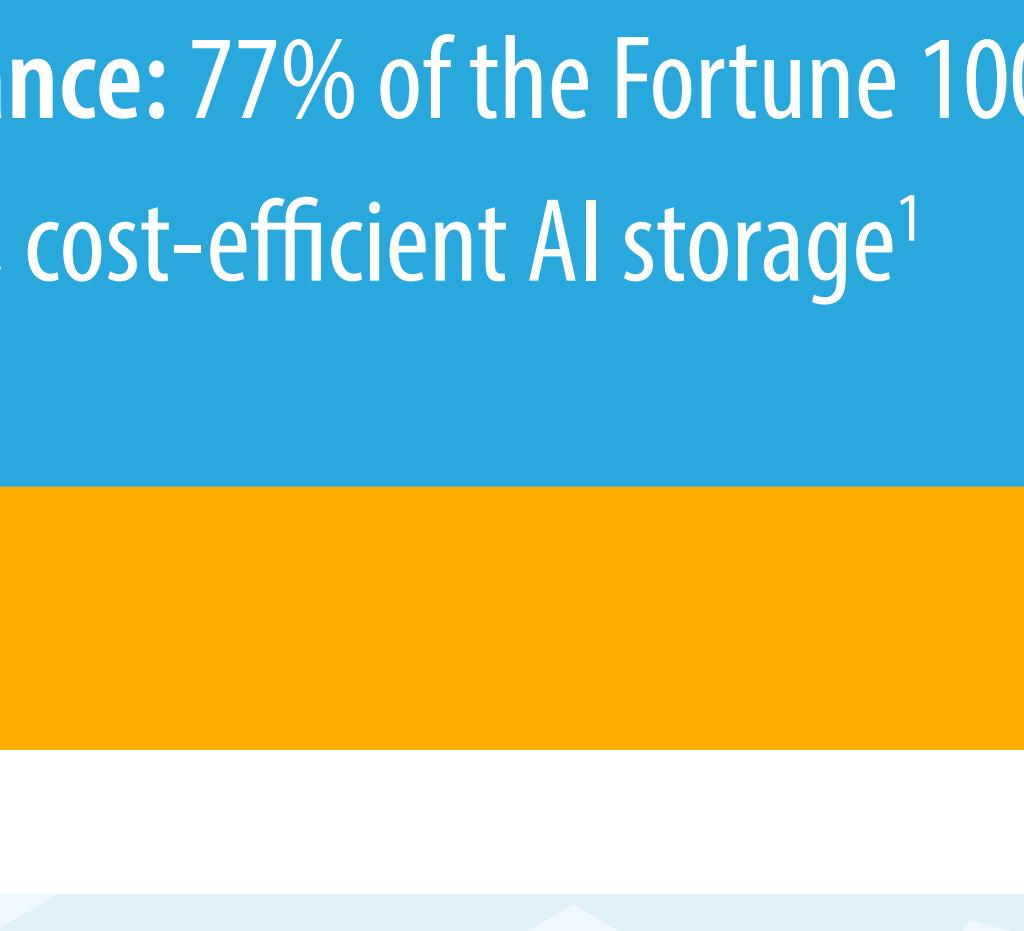- Software-defined storage management (for smarter data lakehouses)

THE RESULT IS FAST, SCALABLE, AND COST-EFFICIENT INFRASTRUCTURE, ESPECIALLY MADE FOR THE FUTURE OF ENTERPRISE AI WORKLOADS.

Here are the top 4 essentials for successful AI.

1 AMD: An end-to-end platform for AI innovation

Enterprise AI workloads need fast access to high-quality, relevant data to meet business goals. AMD provides the full-stack AI solution to prevent bottlenecks.

- **Optimized for storage with large PCIe lane count**
- **AMD EPYC™ CPUs:** Power enterprise AI workloads with leadership per-core performance
- **AMD Instinct™ MI355X GPUs:** Supercharge AI with direct data paths for training & **inference**
- **AMD Pensando™/DPUs:** Reduce storage latency with ultra-fast data movement (400Gbps)

2 MinIO: Object storage for AI brilliance

MinIO delivers high-performance, secure, and massively scalable storage. It can feed the data for enterprise AI workloads, from gigabytes to exabytes and beyond.

- **Object-native storage:** Generative AI, analytics & lakehouses directly access object-native storage
- **Proven performance:** 77% of the Fortune 100 use MinIO for fast, cost-efficient AI storage[1]

3 EDB Postgres: High performance database for faster AI pipelines

EDB Postgres AI is a new data lakehouse stack that utilizes object storage, an open table format, and query accelerators to fuel AI workloads, with high performance and scalability.

- **Unified management** of entire Postgres estate with always-on availability (99.999%)[2]
- **90% better value** running EDB Postgres AI on Supermicro than running EDB on Amazon EC2[3]
- **Automatic tiered storage** moves older data to object storage to improve AI performance

4 Supermicro: Rack-scale solutions for enterprise AI growth

Powered by AMD processors, Supermicro servers with MinIO storage are built to accelerate enterprise AI performance, while saving energy in the data center.

- **Right-size deployments:** Work with AI experts for the right CPUs, GPUs, networking & storage
- **Advanced cooling options:** Lower costs with liquid-cooling and energy-efficient performance
- **Building blocks for AI ROI:** Speed deployment with modular systems, built for easy upgrades

ASG-2015S-E1CR24H

SET THE FOUNDATION FOR SCALE, SPEED & SUCCESS

Discover how future-ready storage infrastructure can meet your enterprise AI needs.

WATCH THE WEBINAR

[1] https://www.min.io
[2] https://www.enterprisedb.com/sites/default/files/pdf/Sovereign_Data_and_AI_Factory_Overview.pdf
[3] https://www.enterprisedb.com/sites/default/files/2024-10/EDB_on_SMC_Benchmark_Report.pdf