

**MINIO**

A Data Leader's Guide

# Building an S3-Compatible Enterprise

How to deliver hybrid and private clouds with MinIO AIStor

---



# Executive Overview

## The Data Challenge

Enterprises are contending with explosive data growth across on-premises, public clouds, and edge environments. AI, analytics, and application modernization initiatives demand storage that performs, scales, and integrates seamlessly across these domains. At the same time, sensitive data and evolving regulations are forcing many organizations and governments to address data sovereignty, keeping data within specific geographic or jurisdictional boundaries while maintaining full operational agility.

Fragmented infrastructure, inconsistent APIs, legacy storage stacks, and escalating public cloud costs make it increasingly difficult to build the unified data foundation that modern data requires. Whether it's powering AI pipelines, modernizing data lakehouses, or building cloud-native, object-native applications, enterprises need a consistent and performant way to store, access, and move data everywhere.

## MinIO AIStor Eliminates Silos with a Universal Data Plane

The Amazon S3 API has become the de facto standard for object storage and data access, adopted by every major cloud provider, platform, and application. This standardization enables true interoperability across hybrid environments, allowing enterprises to move data seamlessly between clouds, colocation facilities, and on-premises data centers. S3 compatibility is the key to breaking cloud lock-in, simplifying architecture, and enabling consistent developer experiences across every environment. S3 API compatibility is the core of the universal data plane.

MinIO AIStor delivers ultra-scalable, high-performance, software-defined, S3-compatible object storage built for hybrid and private cloud architectures and object-native applications, unifying data infrastructure while maximizing performance, scalability, and TCO efficiency.

With industry-leading performance, exascale scalability, and the lowest total cost of ownership in its class, AIStor provides a cloud-native foundation for data-driven enterprises that need to operate at global scale, under sovereign control, and without compromise.

# The New Reality of Enterprise Data

## Unprecedented Data Growth

Global data volumes continue to surge. This growth is driven by connected devices, IoT sensors, machine learning, and digital transformation initiatives across every industry. The diverse sources of data include sprawling environments where data is created, processed, and stored.

Traditional approaches built on monolithic storage arrays, proprietary APIs, or rigid cloud architectures cannot keep pace with the dynamic, distributed nature of enterprise data today.

## Hybrid by Design

Enterprises no longer choose between cloud and on-premises. They operate across both. Data lives in hyperscale clouds, regional data centers, manufacturing plants, hospitals, financial networks, and sovereign infrastructure. These hybrid realities create operational complexity, but also a strategic opportunity: to standardize data management around a single, cloud-native API.

## S3 as the Universal Data Plane

The Amazon S3 API has become the global standard for object storage. Every major cloud provider, modern application platform, AI framework, and analytics engine integrates with S3, making it the most widely supported interface for storing and accessing data.

This ubiquity matters. A common API reduces integration friction, simplifies application development, and enables teams to work with data the same way across different environments. With S3, developers and data pipelines can read, write, and manage data consistently across on-premises infrastructure, public clouds, colocation facilities, and edge or far-edge deployments.

The broad adoption of the S3 API does not automatically create a unified data architecture, but it does make one possible. Because S3 is so deeply embedded across the modern data ecosystem, it provides the shared foundation required for a universal data plane, one where data and applications remain portable even as infrastructure becomes more distributed and diverse.

Many storage systems claim S3 compatibility, but deliver only partial implementations or gateway-based

By standardizing on a complete and consistent version of the S3 API for object storage, enterprises gain a common interface that spans environments and supports every workload type.

translations, missing key API operations, introducing latency from additional metadata layers, or limiting consistency guarantees. These limitations make it difficult to move applications and data across clouds, on-premises environments, and the edge. A complete, consistent implementation of the S3 API is essential for achieving true portability and building a data architecture that can evolve as AI and analytics demands accelerate.

This is why object storage has become the architectural foundation for modern data at scale.

## Why Object Storage Wins

Object storage has become the preferred architecture for modern data because it was designed for scale, durability, and API-driven access. S3 changed how developers think and how they build applications and data infrastructure. It freed them to solve problems more creatively.

Unlike traditional file or block storage systems, object storage decouples data from the underlying infrastructure and exposes it through a simple, RESTful interface that aligns with how modern applications and data pipelines are built.

Several core characteristics make object storage the foundation for today's AI, analytics, and cloud-native workloads:

- **Simplicity for Developers:** Access through a clean API makes it easy to integrate storage directly into applications, services, and data pipelines.
- **Designed for Scale:** A flat namespace and distributed design allow object storage to simply scale from terabytes to exabytes without the hierarchical limitations or bottlenecks of traditional file systems.
- **Flexible Data Model:** The schema-free, metadata-rich object format supports everything from logs and event streams to AI training sets, unstructured media, and application artifacts.
- **Cloud-Native by Design:** Object storage was built for distributed systems and performs predictably across public cloud, on-premises infrastructure, and edge deployments.
- **Cost Effective:** Industry-standard hardware, simpler operational models, and policy-driven data management make object storage significantly more cost-effective than legacy SAN/NAS approaches.
- **Object-level Encryption:** Apply different encryption keys or methods based on the data type, user, or security requirements—not system or volume levels—to deliver fine-grained control over the data.

These attributes explain why nearly every modern analytics engine, AI/ML framework, and cloud-native application integrates with object storage first. While implementations vary widely, the architectural advantages of object storage itself have made it the data foundation for the next decade of enterprise computing.

## Challenges in Building Hybrid and Private Clouds

In order to realize the promise of a unified, S3-compatible data foundation, enterprises face a number of significant challenges:

### 1. Fragmented Storage Architectures

Legacy storage systems (SAN/NAS) and retrofit object gateways introduce multiple layers of translation, metadata bottlenecks, and silos of data. These systems were not built for modern workloads and distributed architectures.

## 2. Cloud Lock-in & Cost Pressure

Public cloud storage seems flexible until egress fees, retrieval penalties, and runaway scaling costs start to bite. Data gravity and vendor lock-in undermine agility and amplify TCO.

## 3. Performance Gaps

Modern data pipelines for AI and analytics demand high throughput, low latency, and massive parallelism. Traditional file/block or gateway-based object stores struggle to feed GPU clusters or large-scale analytic engines.

## 4. Data Sovereignty & Compliance

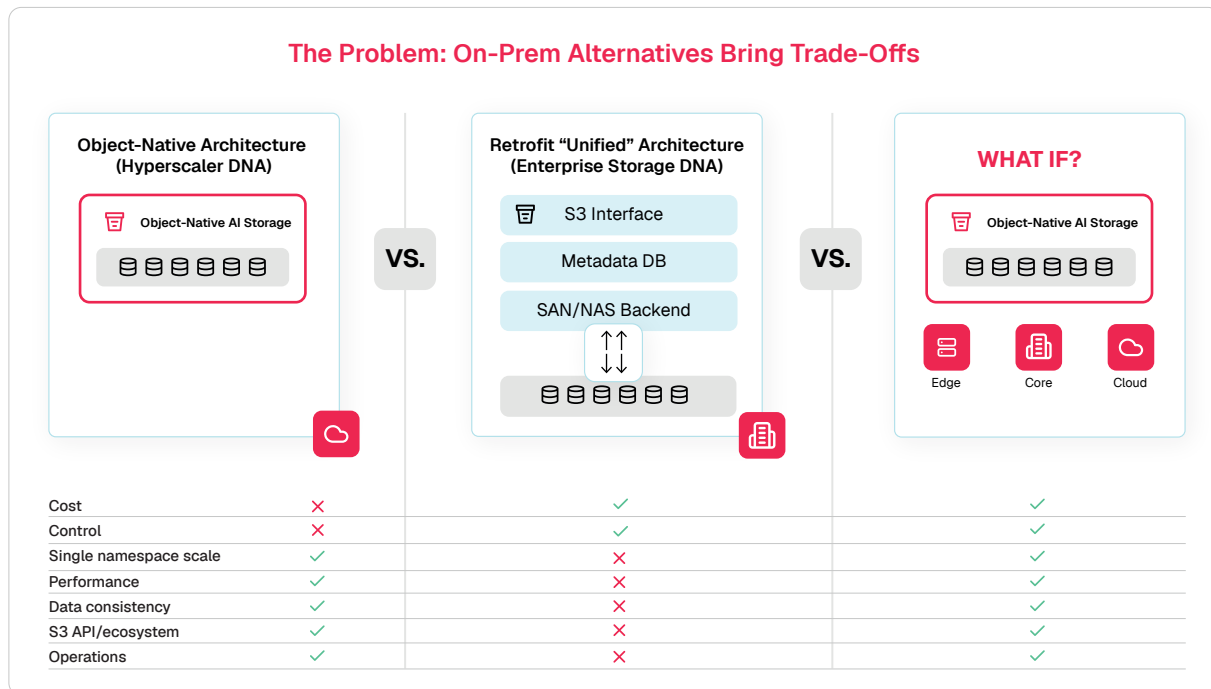
Governments and enterprises must comply with regulatory mandates (GDPR, Schrems II, HIPAA, national security zones) that require in-region, on-premises, or sovereign infrastructure. Meeting these requirements without compromising performance or user experience is a significant architecture challenge.

## 5. Operational Complexity

Coordinating consistency, replication, site failover, metadata coherence, and API uniformity across clouds, edge sites, and on-premises data centers is hard. Without a unified object layer and standard interface, operational burden grows dramatically.

## 6. Data Silos

Modern data is naturally disaggregated. It is developed and processed at the edge, on-premises, and in private clouds. For large enterprises, there are often a number of shadow IT organizations doing all of this in their own silos, so data ends up being distributed around a vast number of data storage systems, many of which do not communicate well with each other.

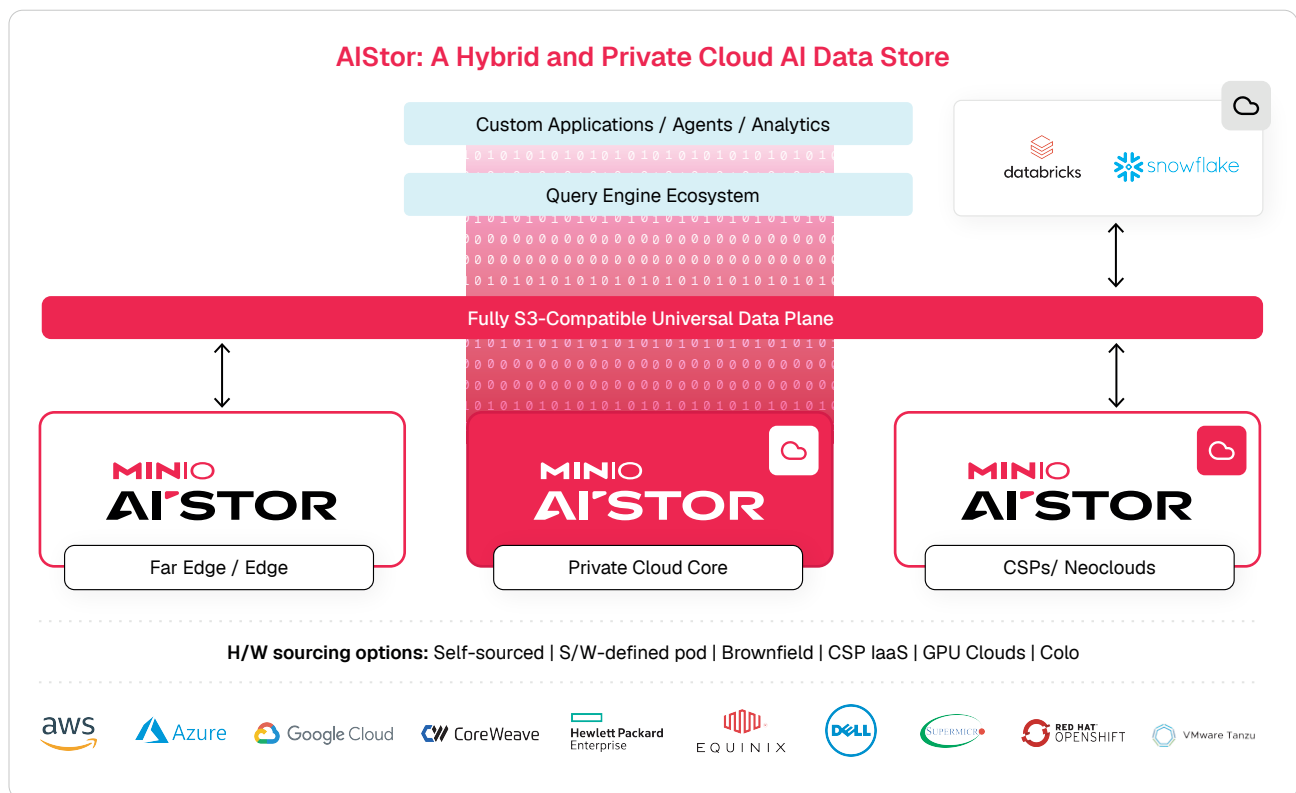


The on-premises choice in building private and hybrid clouds: legacy retrofit vs. object-native.

# The MinIO AIStor Solution

The vision outlined above is addressed by MinIO AIStor through a software-defined, cloud-native, S3-compatible object storage architecture that is designed for hybrid and private clouds and built on the most widely tested S3 alternative to AWS.

S3 is not a bolt on protocol in AIStor. It is the native API, implemented end-to-end with strict adherence to AWS semantics. AIStor's implementation of S3 is [publicly documented](#), validated at scale and updated quickly as AWS evolves. That is why AIStor slots into existing S3 workflows without brittle shims or rewrites.



AIStor delivers S3 compatibility for hybrid and private clouds to create a universal data plane. It can be deployed as a private cloud, a hybrid cloud, a sovereign cloud, and to support custom object-native applications out of the box.

## Cloud-Native Design

AIStor is built to run indistinguishably on bare metal, virtualized infrastructure, or within Kubernetes. Its stateless, object-native architecture eliminates gateway layers and central metadata bottlenecks, ensuring strict read-after-write and list-after-write consistency even at massive scale.

## Performance & Scale

AIStor supports exabyte-scale deployments and is optimized to deliver high throughput and low latency. For example, deployments have scaled to more than 1 EiB of data. The architecture eliminates the typical 20-30 petabyte wall that many legacy systems encounter.

## TCO Optimization

Because AIStor is software-defined and hardware-agnostic, enterprises avoid appliance lock-in and can leverage standard industry hardware footprints. The elimination of gateway layers and metadata overhead also reduces operational cost. A cybersecurity firm migrated off a public cloud object store to achieve significant cost reduction by bringing workloads on-premises with AIStor.

## Unified Data Infrastructure

With full S3 API compatibility, AIStor enables a consistent developer experience across environments: on-premises, hybrid, or cloud. Developers don't need to rewrite or refactor applications; the same S3 calls work everywhere. This interoperability is the cornerstone of hybrid mobility and sovereign architectures, and this unified foundation also supports open table formats such as Apache Iceberg, giving enterprises a consistent storage layer for both objects and structured analytics data.

## Enterprise-Grade Support

MinIO AIStor offers direct-to-engineer enterprise support models, reducing escalation delays and accelerating deployment and issue resolution. This level of support is critical for enterprises deploying mission-critical data infrastructure across hybrid and sovereign zones.

# Four Key S3 Use Cases

Here we highlight four distinct use cases of S3-compatible hybrid and private cloud storage, each tailored to a specific enterprise need and clearly differentiated. While there is overlap in challenge/solution narrative, each use case stands alone for business-oriented clarity.

## S3-Compatible Private Cloud

### Challenge:

Enterprises must modernize legacy infrastructure—block/SAN/NAS or HDFS-based architectures—that cannot scale, operate cost-effectively, or support cloud-native AI and analytics workloads. They also need to keep critical data on-premises for control, latency, or compliance reasons.

### Solution:

MinIO AIStor provides a high-performance, cloud-native foundation for private cloud environments. It delivers S3 compatibility, enterprise performance, linear scalability, and operational simplicity.

Because it runs on-premises in a private cloud environment with the same full S3 API compatibility as the public cloud, it enables developers and data teams to operate consistently. In essence, AIStor is “Amazon S3 in a box.”

#### Case Studies:

A [global digital payments provider](#) migrated from legacy object-gateway systems to MinIO AIStor, scaling from petabytes into exabyte-scale data volumes, while maintaining billions of small-file workloads and meeting SLAs they could not manage previously.

A leading streaming media company accelerated time-to-first-byte (TTFB) to improve end user experience, while enabling them to scale from 100 PB to a projected 300 PB at half the TCO of the public cloud they previously operated on, all while maintaining S3 compatibility.

An IT solutions company was able to accelerate their network log retrieval from more than an hour to less than five minutes—while taking advantage of active-active replication across data centers—by replacing legacy file-based storage to a modern object store with AIStor.

## S3-Compatible Hybrid Cloud Data Mobility

#### Challenge:

As data becomes distributed across edge locations, on-premises facilities, public clouds, and remote data centers, enterprises need seamless data mobility—active-active replication, tiering, and fail-over—without sacrificing performance, consistency or developer experience.

#### Solution:

MinIO AIStor supports active-active replication across sites, automated tiering between on-premises and cloud, and uses the S3 API as the universal interface. This enables consistent data pipelines across edge, on-premises core, and cloud environments, unlocking unified mobility and global data access.

#### Case Studies:

An [automotive manufacturer](#) used MinIO AIStor along with network infrastructure to ingest and store petabytes of autonomous-vehicle data daily across factory edge, regional data centers, and centralized AI training hubs, enabling high throughput and low latency inter-site flows.

A large pharmaceutical company uses AI/ML to automate and accelerate drug discovery and development, and are storing 20+ PB of biological image data in AIStor, driven by millions of experiments every month. They are using AIStor as their high performance AI storage to ensure their GPUs are fully utilized. They are also using AIStor’s full S3 compatibility to manage data across a hybrid cloud, including their data center, colocation facilities, and cold data in Google Cloud.

An aerospace and defense company needed high-performance object storage as a foundation for AI/ML and a modern data lakehouse to reduce the public cloud spend. With AIStor, they achieved significant cost savings while gaining compatibility across cloud-native environments, improving operational efficiency while maintaining compatibility with S3.

## S3-Compatible Sovereign Cloud

### Challenge:

Enterprises operating in regulated industries or government contexts must meet data-sovereignty mandates. Data must reside in specific jurisdictions, under certain control regimes, yet still be accessible and federated across the organization. They cannot rely solely on public cloud regions that might not meet regulatory or chain-of-custody requirements.

### Solution:

MinIO AIStor enables sovereign cloud architectures by supporting on-premises or regionally isolated deployments with full S3 compatibility and federation. Enterprises can deploy AIStor in-region, maintain full control over data location and replication, yet still enable cross-site data access through the same S3 interface.

### Case Study:

Organizations can meet compliance mandates (e.g., GDPR, Schrems II, FedRAMP) while preserving application portability, avoiding vendor lock-in, and retaining high performance. A federal health services organization was able to scale from zero to 750TBs with near zero-touch operations with their sovereign cloud built on AIStor.

## S3-Compatible Object-Native Core Business Applications

### Challenge:

Developers need storage that is fully S3-native—not an adapter or gateway—to build modern applications, data pipelines, analytics, and AI/ML workflows. Many legacy systems act as “object gateways” over block or file storage, introducing latency, complexity, and vendor lock-in.

### Solution:

MinIO AIStor provides a truly object-native storage foundation. Application developers can build directly against the S3 API with no translation layers, delivering predictable performance, linear scalability, and a consistent developer experience across on-premises, private cloud, and hybrid environments. Whether powering analytics engines, AI/ML training pipelines, microservices, or large-scale application data stores, AIStor behaves exactly like public-cloud S3 while running anywhere the enterprise needs it.

This object-native approach also extends upward into the data analytics stack. AIStor Tables integrates directly with Apache Iceberg, providing an open, performant foundation for modern data lakehouse architectures. This allows enterprises to unify structured and unstructured data under a single S3 interface, simplify lakehouse operations, and deliver consistent high-performance analytics across hybrid environments.

### Case Studies:

A [real-time data pipeline](#) using Java, Kafka, and MinIO AIStor shows how enterprises can ingest tens of millions of events, store raw data for analytics while generating aggregated dashboards, all seamlessly across deployment locations.

[Kobiton](#), a leading mobile device testing platform, generates enormous volumes of unstructured data, including screenshots, logs, video, and sensor data from real devices. Their teams needed a storage layer that could ingest this data at high speed, retain it cost-effectively, and feed AI models used for automated testing and anomaly detection. By adopting an S3-compatible object-native architecture, Kobiton simplified data capture across distributed environments, eliminated the scaling and consistency issues of traditional file systems, and enabled their engineering and AI teams to operate on a unified S3 interface. This allowed them to accelerate development cycles, improve test automation accuracy, and support large-scale growth without re-architecting applications.

A Fortune 100 consumer goods company was able to reduce manual QA costs on its factory floor for one of its divisions by 20% by building its visual AI inference application on top of AIStor.

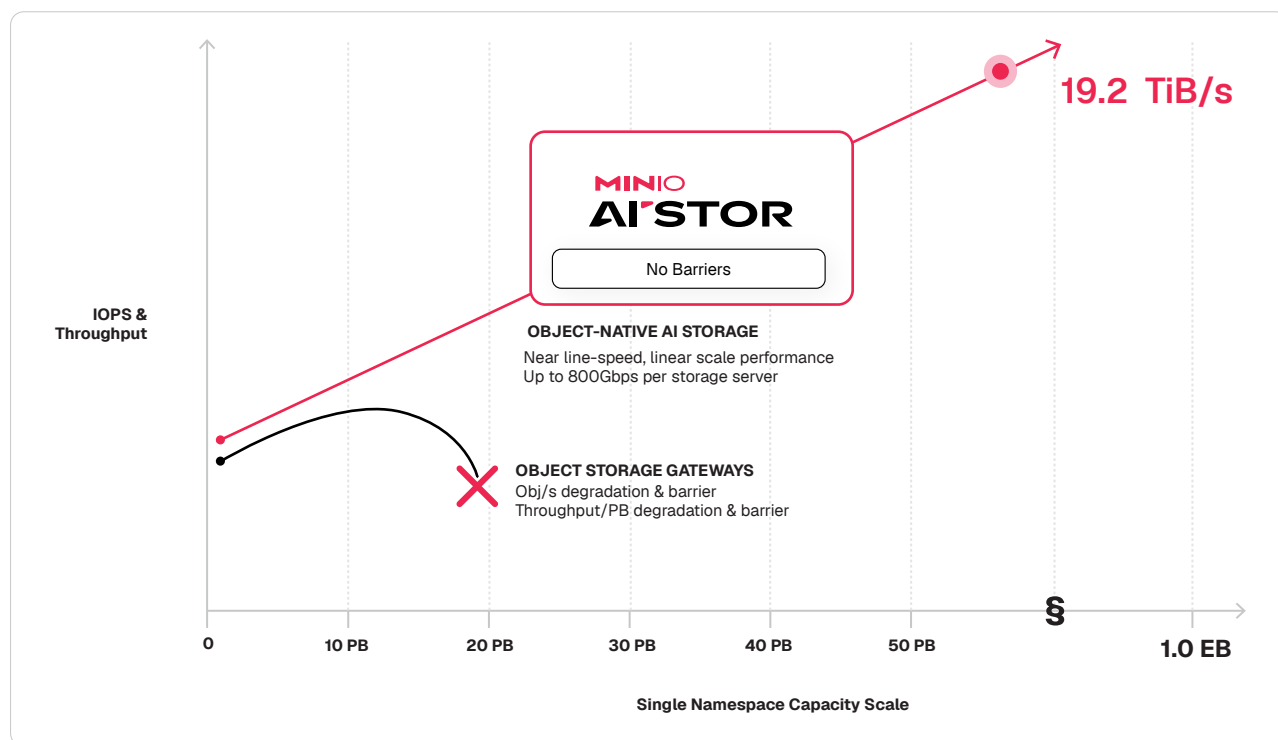
A leading payments company modernized their legacy application on AIStor to reduce month-end reporting by 7 days, while ensuring ample capacity to scale from 100 million to 430 million customers without performance degradation. Across multiple initiatives, AIStor delivered significant improvements in scalability, cost efficiency, and developer agility, while strengthening the company's overall data infrastructure for future growth.

**AIStor empowers enterprises to build once and deploy anywhere, across private, hybrid, or sovereign environments.**

# Impact & Outcomes

The business value of adopting MinIO AIStor spans multiple dimensions and is measurable across enterprise deployments.

- **Performance:** Enterprise-grade throughput, low latency, and linear scale from petabytes to exabytes. Customer deployments already span >1 EiB.
- **Scale:** Proven exascale deployments, eliminating the “petabyte ceiling” characteristic of legacy systems.
- **TCO:** By leveraging industry-standard hardware, avoiding gateway overhead, and reducing vendor-lock-in, enterprises achieve significantly lower total cost of ownership by up to 70% compared with competing solutions. Organizations can reduce [cloud-storage spend](#) substantially by repatriating to on-premises AIStor and tiering cold data to lower cost hardware or the public cloud. Simplified operations minimize the number of FTEs dedicated to managing the system.
- **Development Velocity:** Achieve faster time to market for new applications with software-defined infrastructure-as-code and a familiar operating environment for developers and application teams.
- **Always-on Availability:** Upgrade and replace hardware without downtime. Erasure coding provides unmatched durability. Applications keep running and teams can continue developing.
- **Compatibility:** Full S3 API compliance ensures application portability, developer consistency, and no rewrites across hybrid/private/cloud environments.
- **Support:** Enterprise-grade direct-to-engineer support accelerates deployment, reduces operational risk, and ensures infrastructure reliability.



MinIO AIStor performance scales linearly in a single namespace.

# Conclusion & Next Steps

## The Future Is S3-Native

Enterprise data infrastructure must evolve to meet unprecedented growth, hybrid deployment realities, and regulatory constraints. The future is hybrid, sovereign, and S3-compatible.

## Why MinIO?

MinIO AIStor delivers the universal data plane required for modern enterprises, offering unmatched performance, scale, simplicity, and cost efficiency. Whether your deployment is on-premises, in a private cloud, or across hybrid infrastructure, MinIO AIStor provides the consistent foundation your developers and data teams expect.

## Take the Next Steps

- [Request a Demo](#)
- [Talk to a Hybrid Cloud Expert](#)
- [Try MinIO AIStor](#)

About

# MINIO

MinIO is the data foundation for enterprise AI. Built for exascale performance and limitless scale, MinIO AIStor delivers a secure, sovereign, and AI-ready data store that spans from edge to core to cloud. With rampant adoption across the Fortune 100 and 500, MinIO is redefining how organizations and government agencies store, manage, and mobilize all of their data in the AI era. MinIO is backed by Jerry Yang's AME Cloud Ventures, Dell Technologies, General Catalyst, Index Ventures, Intel Capital, Softbank Vision Fund 2 and others.