

High-Performance Data Storage for Observability and Telemetry

“AIStor helped us turn what was once a fragile, monolithic system into a very forward looking Data Lakehouse that supports a true hybrid cloud.”

— Conor Brennan, Managing Director, Risk IT, Nomura

Reduce TCO. Query data faster. Manage petabyte-scale observability data seamlessly.

- **Store 2–3× more data** on the same hardware with erasure coding vs. replication
- **Retrieve historical data 12× faster** than traditional archive-tier storage
- **Save up to 60%** in total observability storage costs
- **Native S3 tiering** with Splunk SmartStore, Elastic ILM, and Grafana Mimir/Loki
- **Deploy anywhere:** on-prem, edge, air-gapped, Kubernetes-native

The Challenge: Observability Data Is Outpacing the Storage That Supports It

Modern enterprises generate staggering volumes of telemetry. A single Kubernetes cluster emits millions of log lines per minute. Distributed tracing across microservices produces petabytes of span data per quarter. One of the world's largest digital payments platforms, an AIStor customer, generates over 3PB per day in observability data alone. AI inference pipelines add model performance signals, token usage metrics, and prompt audit logs to the mix. And every one of these data streams is growing.

The observability platforms that consume this data (Splunk, Elastic, Grafana) are built for analysis, not for storage at scale. Their storage backends rely on coupled compute-storage architectures or cloud object stores with usage-based pricing that compounds quickly. Datadog, for example, charges \$0.10/GB for log ingestion plus \$1.70/GB/month for indexing at its default 15-day retention.

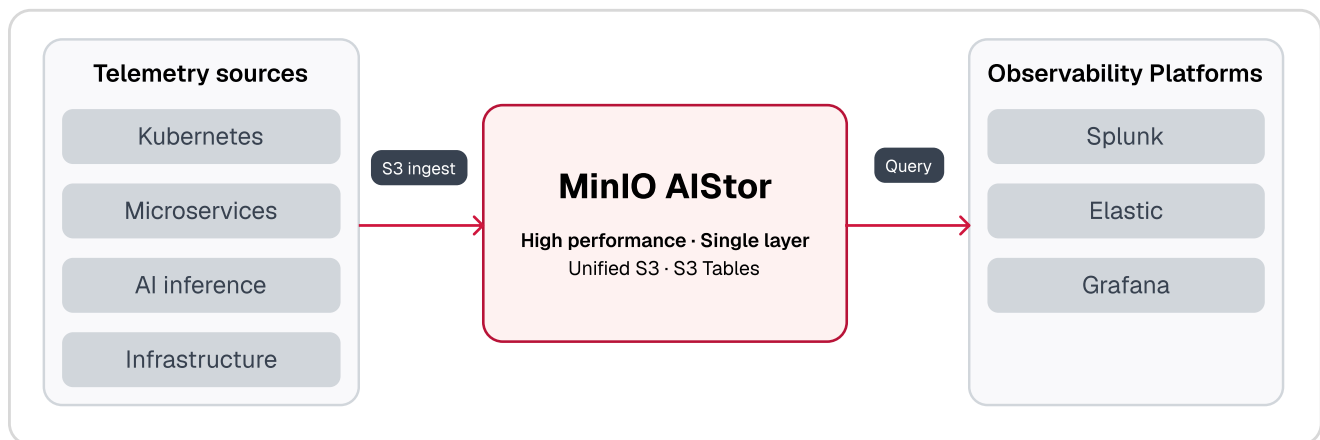
At 1 TB/day with 30-day retention, log costs alone exceed \$1.2 million per year. These numbers are not outliers. Gartner reports that 44% of enterprise clients now spend over \$1 million per year on observability, with more than half of that going to logs. Faced with these economics, the compromise is always the same: shorter retention windows, sampled data, and blind spots during incident response.

The symptoms are predictable. Retention windows shortened to fit budget constraints rather than compliance or operational needs. Forensic investigations that stall because the relevant data was deleted or moved to slow archive tiers weeks ago. Alert thresholds tuned on sampled metrics that miss edge cases. For organizations adding AI inference to their stack, these same constraints extend further: model drift detection, inference latency monitoring, and token cost attribution all demand long-lived, high-fidelity telemetry that existing storage economics make prohibitively expensive.

The AIStor Solution: High-Performance Observability Storage

AIStor decouples storage from compute in your observability stack and delivers consistent, high-throughput performance that does not degrade as data volumes grow, allowing each layer to scale independently based on actual demand. Your observability platform handles indexing, querying, and alerting. AIStor handles storing every byte of telemetry at scale, with consistent performance whether the data is one hour old or one year old.

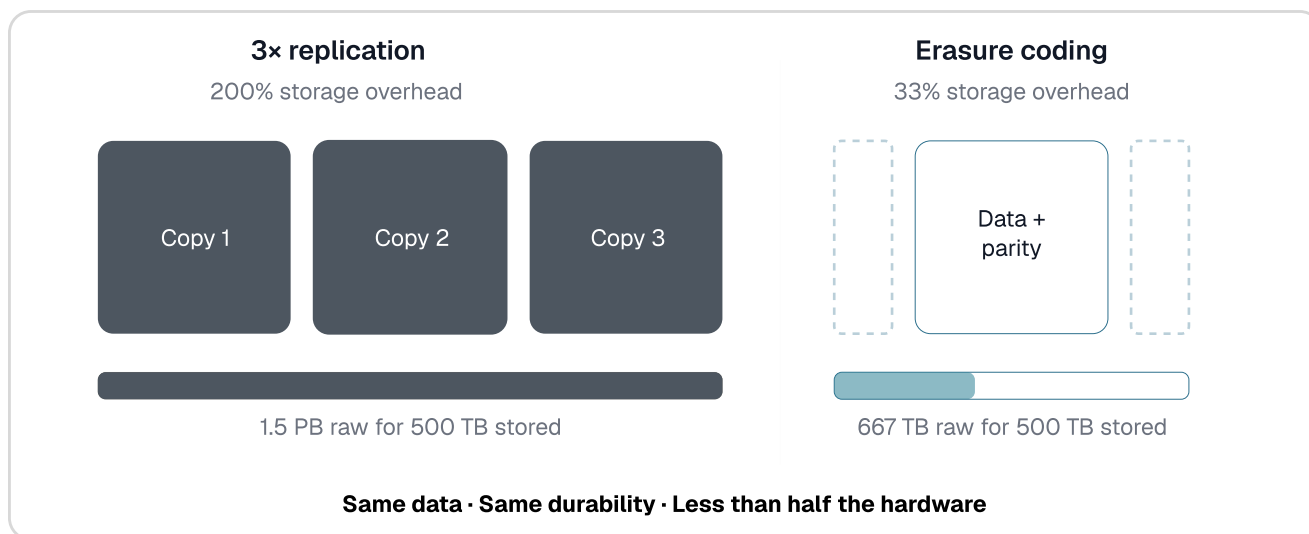
The integration is native. AIStor connects directly to observability platforms via S3 with no custom connectors, no middleware, and no gateway translation. Splunk SmartStore, Elastic Index Lifecycle Management, and Grafana Mimir and Loki all support S3-compatible remote storage, making AIStor a direct drop-in once the S3 endpoint is configured. AIStor serves as the high-performance primary store for observability data, delivering consistent throughput for both real-time alerting and historical queries. With AIStor Tables, structured telemetry data can be queried directly using standard analytics engines, eliminating the need for separate hot, warm, and cold tiers. You search last quarter's logs as fast as yesterday's.



Flexible Data Protection: Erasure Coding and Replication, Matched to Your Needs

Traditional observability backends protect data through replication, typically maintaining three full copies of every log, metric, and trace. That means for every petabyte of telemetry, you provision three petabytes of raw storage. For single-site durability, AIStor eliminates this overhead with flexible, configurable erasure coding. In a typical 12:4 configuration, AIStor distributes 12 data shards and 4 parity shards across drives in an erasure set, achieving 11 nines of durability at just 33% storage overhead instead of 200%. That translates directly into storing 2–3× more telemetry data on the same hardware.

For organizations that need multi-site durability, whether for active-active observability dashboards, disaster recovery, or regulatory requirements, AIStor supports asynchronous replication across sites. Because replication is asynchronous, the performance of the primary site is not compromised by cross-site writes, which makes it well-suited for high-volume log and telemetry workloads where ingest throughput cannot be sacrificed. Customers choose the protection model that fits their architecture: erasure coding for cost-efficient single-site durability, replication for cross-site resilience, or both.



As an example: if your current Splunk or Elastic deployment stores 500TB of telemetry using 3x replication for single-site protection, that consumes 1.5PB of raw capacity. AIStor with erasure coding stores the same 500TB in approximately 667TB. The savings translate directly into longer retention, broader data coverage, or both, without purchasing additional hardware. One of the world's largest content platforms leveraged this exact approach to deploy 38.4PB of usable capacity from 51.2PB of raw storage across just 16 AIStor nodes.

Multi-TB/s Ingest with Linear Scaling

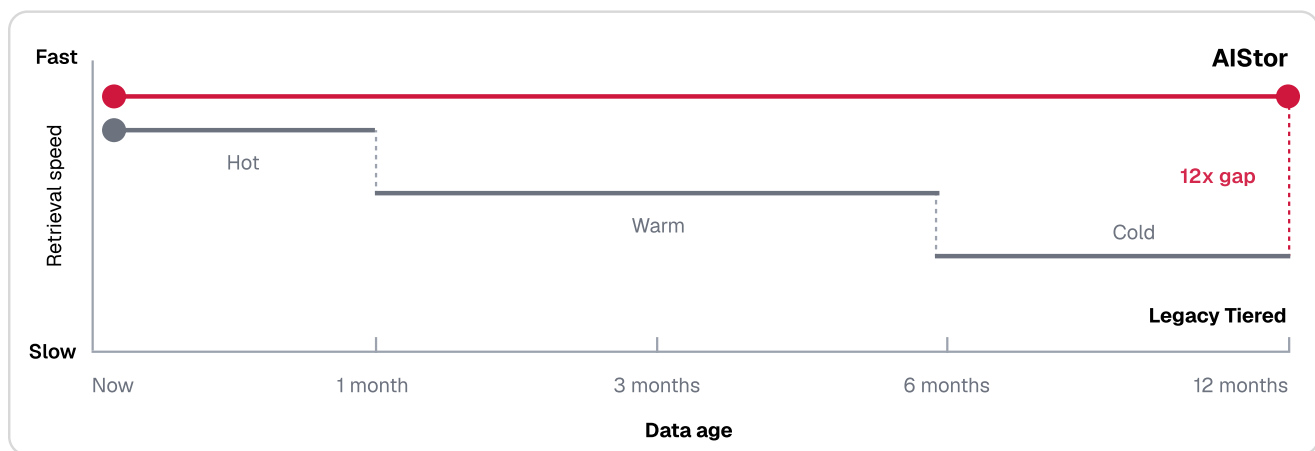
Observability workloads are write-heavy by nature. Thousands of agents, exporters, and collectors push data continuously, and the storage layer must absorb that ingest without throttling the pipeline. AIStor's stateless architecture has no centralized metadata server to bottleneck ingest or reads. Every node in the cluster contributes to aggregate throughput at near line-speed. Object placement is determined by deterministic hashing. A hash of the full object path selects the erasure set, so any node can receive a request, compute the hash, and route directly to the correct storage location without querying a central authority.

This is why performance remains consistent as the cluster grows. Adding capacity means adding server pools, which expands the number of available erasure sets. There is no rebalancing, no migration jobs, no metadata resharding. A cluster with one pool and a cluster with ten pools operate the same way. Each node added increases aggregate throughput linearly. For the global content platform running one of the largest Splunk SmartStore deployments in production, this architecture delivered high throughput using AVX-512 accelerated erasure coding while maintaining the performance their observability pipeline demanded.

Query Historical Data at Full Speed

The value proposition of long retention is useless if historical data is slow to access. Legacy tiered storage treats older data as second-class, moving it to archive tiers with retrieval latencies measured in hours. Because AIStor stores all data on a single high-performance tier with no internal archival layers, it maintains near line-speed reads across the entire retention window. There are no cold-tier penalties and no rehydration delays.

When an incident occurs and the investigation requires correlating today's alerts with patterns from six months ago, the query against AIStor runs at the same speed as one from yesterday because no data migration between tiers is involved. This enables faster root-cause analysis, reduced mean time to resolution, and improved SLA compliance. Organizations using AIStor for observability storage report retrieving historical data up to 12x faster than traditional archive-tier solutions.



Single Namespace from Terabytes to Exabytes

Legacy object storage platforms hit hard ceilings, typically 20 to 30PB, that force re-architecture, cluster splits, and fragmented search history. AIStor provides one logical namespace from your first terabyte to your first exabyte. Logs, metrics, and traces live in one unified store. No cluster splits that fragment searchable history. No re-architecture as volumes grow.

For observability workloads, this matters operationally. A single namespace means a single endpoint for your observability platform, a single access control model, and a single set of retention policies. There is no shard management, no cross-cluster query routing, and no namespace migration when capacity thresholds are reached. AIStor handles all routing within pools and erasure sets transparently while the application continues to use the same endpoint.

Software-Defined and Kubernetes-Native: Deploy Where Your Stack Runs

AIStor is software-defined and runs on commodity hardware in your data center, at the edge, or in air-gapped environments. There are no proprietary appliances and no vendor lock-in. For organizations with strict data sovereignty or compliance requirements, including financial services, healthcare, government, and defense, AIStor provides sovereign deployment options that keep every byte of telemetry within your controlled perimeter.

AIStor is also Kubernetes-native, deploying alongside the same container orchestration that runs your observability stack. This eliminates the operational burden of managing a separate storage infrastructure and enables infrastructure-as-code workflows for provisioning, scaling, and monitoring storage alongside compute.

Traditional Observability Storage vs. MinIO AIStor

	Traditional Approach	MinIO AIStor
Data Protection	3x replication, 200% overhead	Erasure coding, 33% overhead, 11 nines durability
Historical Query Speed	Archive latency, hours to rehydrate	Near line-speed reads, no cold-tier penalties
Scaling	Coupled compute-storage, rebalancing required	Disaggregated, add pools with no rebalancing
Metadata	External database, bottleneck at scale	Inline with object data, no external dependencies
Platform Integration	Custom connectors, middleware layers	Native S3, works with Splunk, Elastic, Grafana out of the box
Pricing Model	Per-operation, egress, and ingest fees	Capacity-based, no per-operation or egress charges
Deployment	Cloud-only or proprietary appliances	Software-defined, Kubernetes-native, deploy anywhere

MinIO AIStor for Observability and Telemetry

AIStor's stateless architecture with deterministic hashing eliminates the centralized metadata bottlenecks that constrain legacy object storage at scale, while maintaining near line-speed reads across the entire retention window regardless of data age. As a unified data store, AIStor provides native S3 and Iceberg interfaces alongside AIStor Tables and Delta Sharing, enabling observability platforms and analytics engines to query telemetry data directly without middleware, gateway translation, or data movement between systems. Capacity-based pricing with no per-operation or egress fees keeps storage costs predictable as data volumes grow. In production, AIStor deployments span tens of petabytes of usable capacity, with customers reporting 40% or greater reductions in total cost of ownership and migration timelines measured in weeks rather than months.

Ready to see it in action?

Visit min.io to learn more. [Download AIStor](#) and test it yourself. [Request a demo](#) to see how observability data work in your environment.