MINIO

Industrial and Operational Analytics

A Modern Data Foundation for Full-Fidelity Operational Intelligence

---

"AIStor has enabled us to scale our smart metering infrastructure faster and more efficiently than we imagined. The time savings, simplicity, and performance have been game-changing."

— National electric utility representative

Every signal retained. Every investigation accelerated. Every AI model fed.

- **2–3×** more operational history without increasing budgets.
- **Minutes** to investigate failures, not weeks
- **Native** integration with analytics, AI, and OT tooling
- **PB to EB** Full-fidelity retention, write once, read everywhere

The Challenge: Industrial Data Infrastructure Was Not Built for This

Traditional operational analytics architectures were designed for a world that no longer exists. Sensor data was sampled once per minute, not per second or faster. Data types were numeric only, not video, images, or acoustic signals. Retention was measured in days or weeks, not years or decades. And analytics powered dashboards and thresholds, not ML models and AI inference.

That world is gone. Enterprises must plan for two to three times data growth over the next four years (Gartner). Edge and IIoT data is growing at 30%+ per year, outpacing other enterprise data categories (IDC/McKinsey). Over 90% of sensor data is never stored or analyzed. Insufficient historical data is a top barrier to industrial AI adoption (Gartner). And regulated industries must retain some operational and quality records for seven to 10 years or more.

Three Limiting Factors

Scale. Legacy platforms force sampling and discarding. When your historian or Hadoop cluster hits capacity limits, the default response is to downsample or purge. The data you throw away is the data your ML models need.

Cost. Hadoop imposes a 3x storage tax through replication, limiting what you can afford to retain. For organizations exploring hybrid or cloud-adjacent architectures, per-operation and egress fees compound further at industrial data volumes.

Analytics ceilings. Proprietary data formats and closed ecosystems block the path from descriptive dashboards to predictive ML. When your data is locked in a historian or siloed across plant-specific systems, teams spend 30-40% of their analytics time stitching data across silos instead of learning from it.

The Operational Cost of the Status Quo

The consequences show up on the plant floor. Investigations hit "data not available" when engineers need to trace a quality excursion back through months of process history. Root cause analysis takes weeks instead of hours because cross-system data has to be manually assembled. Process drift is detected after yield loss, not before. ML initiatives stall because there is insufficient historical data at full fidelity to train models that generalize. And capital keeps flowing into infrastructure that cannot scale with the data.

AIStor: A Data Foundation Built for Industrial Scale

MinIO AIStor is a high-performance, S3-native data store purpose-built for AI and analytics workloads. It provides a unified data foundation beneath your existing historians, analytics engines, and operational tools, consolidating the fragmented layers of Hadoop clusters, siloed archives, and cloud object stores into a single platform that can ingest, retain, and serve operational data from petabytes to exabytes, on-premises, at the edge, or in hybrid configurations.

AIStor enables industrial organizations to centralize and analyze two to three times more data for the same cost they are paying today. It reduces downtime and maintenance waste, lifts equipment effectiveness and throughput, improves yield and right-first-time quality, and provides the data foundation to scale industrial AI from a single pilot to an entire plant network.

Operational Analytics Workloads

Process History Retention and Historian Modernization

The foundation workload. AIStor serves as a durable system of record for operational process data, preserving raw, high-frequency sensor and control data for long-term replay and reuse without downsampling or runaway cost. Full-fidelity retention means every signal is available for historical analysis, regulatory audits, and ML training, not just the summary statistics your legacy historian keeps.

AIStor's full S3 compatibility provides native integration with analytics platforms like Databricks, Spark, ClickHouse, and Splunk. Data is written once and accessed by any tool in the ecosystem, eliminating the copies, exports, and ETL pipelines that fragment operational data today.

Reliability Engineering and Root Cause Forensics

When an asset fails or a production line drifts out of specification, the investigation depends entirely on data availability. AIStor collapses investigation timelines from weeks to minutes by making multi-year, cross-system process data queryable through standard analytics tooling. Engineers can correlate vibration data with temperature profiles, control setpoints, and maintenance records in a single environment rather than requesting extracts from three different teams.

Full-fidelity data access enables the shift from reactive forensics to proactive detection. Reliability teams can build drift detection models in Spark, Databricks, or Python-based analytics that identify degradation patterns before they reach the threshold that triggers a failure or yield loss.

Quality Analytics and Traceability

Quality workloads are shifting from manual inspection and statistical sampling to real-time, ML-powered defect detection. This transition demands both high-throughput data ingestion (camera feeds, measurement streams, test results) and long-term retention for traceability and regulatory compliance.

AIStor handles both sides of this equation. Line-speed ingestion keeps pace with high-frequency inspection systems. And multi-year retention with versioning and object locking provides the storage-layer controls required to support compliance with FDA 21 CFR Part 11, EU Annex 11, and ISO 9001 quality records requirements for 7 to 10 years or more, including in air-gapped and sovereign environments where data residency is non-negotiable.

Asset Health and Predictive Maintenance

Predictive maintenance models need long-horizon sensor data across entire asset fleets to detect failure patterns that only emerge over months or years. AIStor provides the scale and economics to retain this data at full fidelity without the storage tax imposed by Hadoop replication or cloud egress fees.

Because AIStor scales linearly from tens of petabytes to exabytes in a single namespace, organizations can expand from a single-plant pilot to fleet-wide predictive maintenance without re-architecting their data infrastructure. The same S3 API that serves descriptive analytics today serves AI and ML training pipelines tomorrow.

Computer Vision and Visual Inspection

AI-driven visual inspection replaces manual QA with consistent, high-throughput defect detection across production lines. These workloads generate massive volumes of image and video data that must be ingested at line speed, retained for model retraining, and served to inference endpoints with low latency.

AIStor's performance profile, up to 800 Gbps per storage server with linear scaling across nodes, keeps pace with multi-camera inspection systems. Its S3 compatibility means the same storage layer that holds training data also serves the inference pipeline, eliminating the data movement overhead that slows deployment of vision systems at scale.

What Makes AIStor Different

Performance. Line-speed ingestion and high-concurrency access for telemetry, time-series, and video. Supports concurrent analytics, AI, and operational workloads without contention.

Scale. From billions of small objects to petabytes and exabytes in a single namespace. Multi-year retention without architecture changes.

Economics. 40%+ lower TCO with industry-standard hardware and no proprietary lock-in. Predictable economics for data that must be retained for years, not days.

Ecosystem. Full S3 compatibility across analytics, AI, and industrial platforms. On-premises, edge, and hybrid deployments without app rewrites. Native integration with Databricks, Spark, ClickHouse, Splunk, and more.

Support. Mission-critical reliability with direct engineer access 24x7x365. Built for environments where downtime is not acceptable.

Traditional Industrial Data Infrastructure vs. MinIO AIStor

	Traditional Infrastructure	MinIO AIStor
Data Retention	Downsampled, often days or weeks	Full fidelity, years to decades
Storage Tax	3x replication (Hadoop) or cloud egress fees	Erasure coding, no egress fees
Analytics Ecosystem	Proprietary formats, closed tools	S3-native, open ecosystem
Scale Path	Forklift upgrades or cloud cost escalation	Linear scale, industry-standard hardware
AI/ML Readiness	Requires separate data pipelines	Serves analytics and ML from one layer
Deployment Flexibility	Cloud-locked or appliance-locked	On-premises, edge, hybrid, air-gapped
Operational Overhead	Multi-FTE, multi-vendor	Sub-1 FTE, single platform
Time to Production	Months	Weeks

Proven in Production

"AIStor was the perfect solution to fix a challenging operational issue around customer reporting, while also enabling us to transform our architecture to be more agile and position us to maximize our ability to leverage AI."
— Global digital payments provider

National Smart Grid Modernization

A national electric utility needed to support a national smart metering initiative on a data platform that could handle AI-driven grid management at scale. Expanding their legacy Hadoop environment would have required 240 additional servers, 14 months to implement, and 3 full-time FTEs to operate.

AIStor deployed on 90 servers in 10 weeks, delivering **>50% lower TCO**, a **62.5% reduction in hardware**, and operational overhead of **less than 1 FTE**. AI models for anomaly detection and load disaggregation reached >90% accuracy on the new platform.

Semiconductor Manufacturing Quality

A major semiconductor manufacturer integrated on-premises AIStor with Databricks across two manufacturing zones to generate near real-time manufacturing insights. The deployment provided 250 TB of securely replicated storage with full S3 compatibility and native Delta Sharing integration, improving yield and quality analytics while reducing loss and waste.

AI-Driven Visual Inspection at Scale

A leading consumer packaged goods company shifted a manufacturing product line from manual QA to AI-based visual inspection, targeting a **20% reduction in inspection labor, or $2-4M per month in savings**. AIStor serves as the storage layer for both training data and inference pipelines, providing the throughput and scale needed to extend AI inspection across multiple production lines.

Exabyte-Scale AI

One of the world's largest manufacturers of advanced data storage components built a **30 PB system with a roadmap to 1 exabyte** to power industry-specific GenAI operations. AIStor provides the performance and single-namespace scale required to serve specialized AI workloads on manufacturing data at a scale that would be economically infeasible on legacy or cloud infrastructure.

Pharmaceutical Compliance in Air-Gapped Environments

A leading pharmaceutical manufacturer stores 10+ years of batch images across air-gapped pharmaceutical production sites. AIStor delivers S3-compatible private cloud storage that meets regulatory retention mandates without sacrificing performance or requiring network connectivity to external services.

Sovereign and Secure Deployments

The US Army Corps of Engineers requires secure, on-premises, S3-compatible storage with FIPS compliance and air-gapped capabilities to meet DoD retention requirements. AIStor delivers the security, scalability, and compliance posture required for sovereign environments without the cost of proprietary alternatives.

Business Value Summary

Quantifiable Outcomes

Retain 2-3x more operational data on the same budget. Erasure coding eliminates the 3x storage overhead of Hadoop replication. Inline compression reduces physical footprint prior to erasure coding. No per-operation or egress fees.

Collapse investigation timelines from weeks to minutes. Full-fidelity, cross-system data is queryable through standard analytics tooling. Engineers find answers instead of requesting data extracts.

Reduce infrastructure footprint and operational overhead. A national electric utility cut hardware by 62.5% and operational staff from 3 FTEs to less than 1. Time to production dropped from 14 months to 10 weeks.

Scale AI from pilot to plant network. AIStor's linear scaling and consistent S3 API mean the infrastructure that serves a single-plant PoC is the same architecture that serves an enterprise rollout.

Strategic Outcomes

Eliminate "data not available." When every signal is retained at full fidelity, investigations, audits, and ML training are never blocked by missing data.

Unified governance across operational data. One platform, one IAM model, one audit log. Versioning and object locking protect data integrity. Lifecycle automation enforces retention policy continuously.

No vendor lock-in. Software-defined, S3-native architecture runs on industry-standard hardware. Deploy on-premises, at the edge, or in hybrid configurations without rewriting applications or migrating data.

Next Steps

Extend operational insights without extending budgets. Improve operations faster. Lower TCO with software-defined economics.

Executive Deep Dive. Strategy session with your leadership team

Technical Deep Dive. Architecture review with your platform engineers

Live Demo. See AIStor in action with your OT stack.

Proof of Value. Deploy in your environment with your data.

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.