




MINIO

SOLUTION OVERVIEW

F5 and MinIO for secure, performant AI infrastructure



Key benefits

Reduce public cloud costs

Address data repatriation needs and lower cloud storage costs.

Eliminate AI pipeline bottlenecks

Deliver consistent, high-speed data movement for AI workflows.

Enable comprehensive data security

Ensure data storage and access remain protected from cyber threats.

Optimize cloud

Deploy and synchronize data across on-premises, hybrid, and cloud environments seamlessly.

AI applications and modern data-hungry workloads are quickly shifting from experiment to production at scale, putting inadequate infrastructure into sharp focus as the next hurdle to successful execution. Organizations are discovering that AI workloads that once made sense in the cloud are now facing high public cloud costs, latency bottlenecks in AI workflows, and a growing demand for secure, performant, and reliable data access.

Organizations need an alternative to their existing public cloud infrastructure that delivers AI securely at scale while reducing cost and ensuring high-speed data movement.

Power high-performance, secure AI data delivery

By deploying secure, high-performance data management and delivery for AI and enterprise workloads. By deploying F5® Application Delivery and Security Platform (ADSP) alongside MinIO's scalable, high-performance S3-compatible object storage solution, AIStor, organizations are empowered with a seamless and powerful solution for AI data pipelines, data repatriation, regional replication, and data storage. Together, these solutions deliver the high throughput, unified security, and optimized performance needed to maximize GPU utilization and enable scalable, resilient pipelines for data-driven applications, including AI and ML, across on-premises, cloud-native, and hybrid environments.

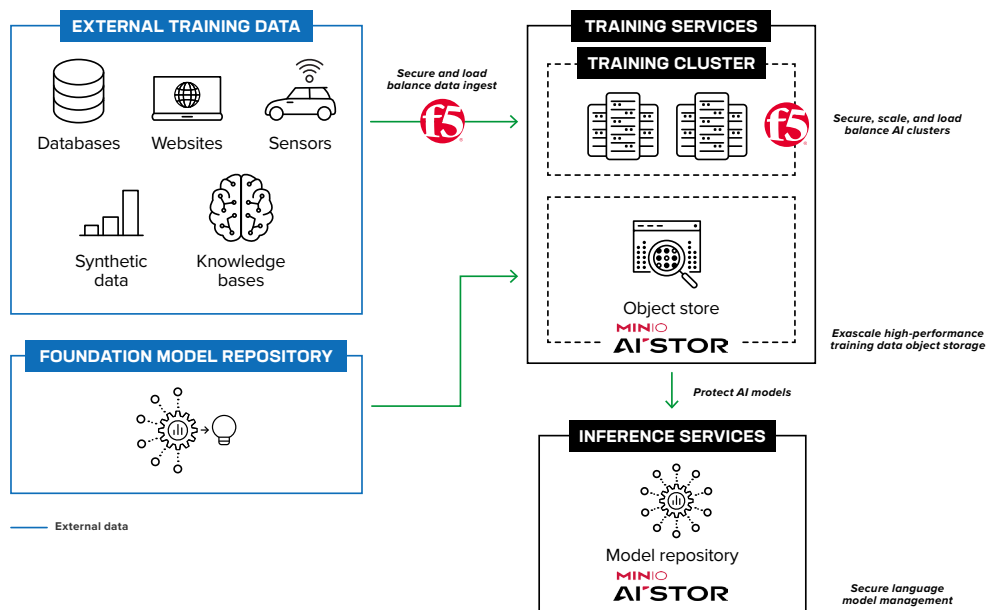


Figure 1: F5 Application Delivery and Security Platform (ADSP) and AIStor enable high-performance data delivery and security for AI workloads.

Key features

Seamless scalability and performance

Exascale data operations and consistent, low-latency performance for demanding workloads.

Unified security and resilience

Secure, multi-tenant architectures that mitigate risks and enable compliance across diverse environments.

Operational efficiency

Intelligent load balancing and centralized orchestration to reduce operational overhead and maximize resource utilization.

Streamlined deployment

Minimized deployment complexity and expedited production readiness.

Performant traffic management for AI data delivery

The combination of F5 ADSP, F5 BIG-IP LTM and MinIO AIStor provides high-performance load balancing, high-volume throughput, and secure ingestion pipelines for data training and inference. This setup allows for scaling hundreds of Gbps to Tbps of bandwidth, ensuring that even the most data-intensive operations can be handled smoothly. Solution capabilities include:

- **Intelligent traffic distribution:** BIG-IP optimizes data flow across AIStor clusters to ensure efficient resource allocation, even under exascale demand, while eliminating bottlenecks and hotspots.
- **High throughput object storage:** AIStor delivers S3-compatible performance that can scale to accommodate burst data loads for intensive AI workflows.
- **Low-latency pipeline optimization:** The combined solution minimizes data transmission delays, enhancing the speed of AI training and real-time inference.

Boost data repatriation

As data-driven workloads stabilize, cost and optimization become a priority over elasticity of environments—your workloads may need to change location as they move to production at scale and start hitting the performance limits or exceeding cloud budgets. But data repatriation presents hurdles such as inconsistent network performance, security gaps during data transfers, and rigid dependencies on S3 APIs, which limit flexibility and scalability. Traditional approaches can fall short due to inefficient data transfer processes, weak traffic optimization, and inadequate security controls, which can hinder performance and inflate costs.

F5 BIG-IP and MinIO AIStor simplify this transition. AIStor's consistent S3-compatible API enables applications to migrate smoothly, while BIG-IP intelligently orchestrates incremental traffic shifts through DNS routing, allowing for phased, disruption-free migrations with minimal adjustments.

Unified security and resilience

Integrated security across networking and storage fortifies hybrid and multicloud infrastructures. F5 delivers robust protection through advanced encryption, secure load balancing, and resilient multi-tenant architectures, ensuring data integrity and regulatory compliance while safeguarding mission-critical AI operations. Three key benefits include:

1. **Comprehensive threat protection:** F5's advanced security features actively safeguard against network and application attacks while securing data access across distributed systems.

- 2. End-to-end encryption:** Integrated SSL/TLS offloading ensures that data is secured in transit.
- 3. Resilient multi-tenant architecture:** Active-active clustering and regional replication deliver robust data availability and rapid disaster recovery for critical workloads.

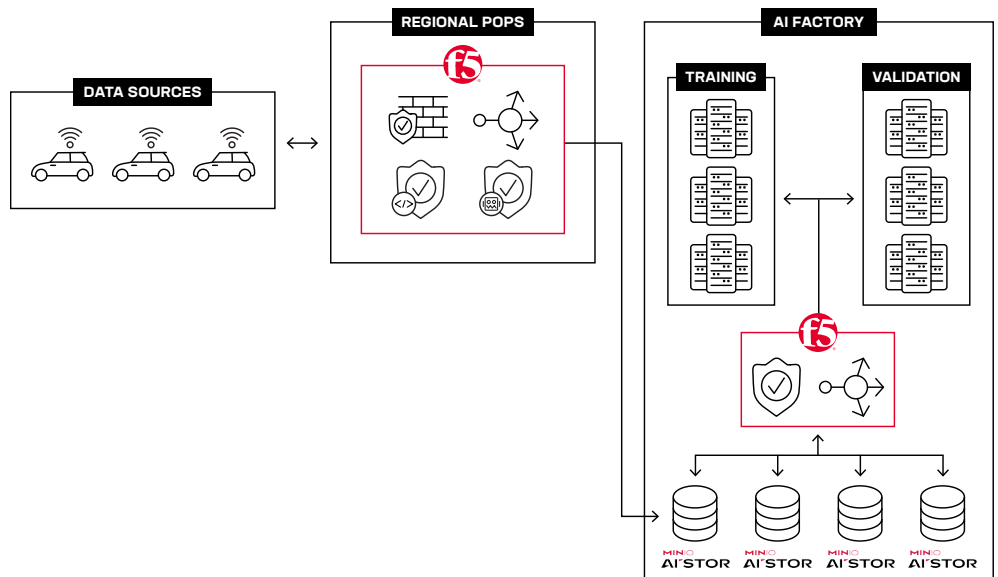
An automotive success story

From autonomous vehicle training data management to healthcare analytics and financial compliance, the flexibility offered by F5 and MinIO translates into a powerful competitive advantage.

For example, a major global automotive manufacturer is using the joint solution provided by F5 and MinIO to securely and reliably deliver and manage proprietary data from its vehicle fleets worldwide to [AI Factories](#). This enables continuous improvements to their AI models, providing significant business advantages and has enabled millions in savings annually by repatriating data from the public cloud. The manufacturer has also reduced downtime, enhanced reliability, maximized ROI on AI infrastructure, and ensured secure management of exascale data.

Leveraging F5's secure traffic management capabilities, vehicle-generated data is securely collected and efficiently routed through F5-powered regional points of presence (PoPs), ultimately arriving at a centralized data lake running AIStor inside the AI Factories. MinIO's advanced object storage solution forms the essential foundation for training AI models. Updated models are then delivered securely back to vehicle systems, with the entire cycle managed and protected by F5 infrastructure.

Figure 2: A leading automotive manufacturer uses F5 and MinIO for secure AI data delivery.



Value across industries

This large-scale deployment illustrates the significant advantage organizations in other industries can gain by adopting the combined F5 and MinIO solution. For instance:

- Healthcare providers can use secure, high-performance storage to enhance predictive analytics and medical imaging, improving patient outcomes through rapid, reliable data access.
- Financial institutions leverage BIG-IP's security and MinIO's scalability for compliant storage of sensitive data, enabling AI-driven fraud detection and risk management.
- Companies with edge computing deployments can count on F5 and MinIO for secure, low-latency data processing and storage—critical for IoT and smart industry use cases.

Accelerate AI data pipelines, boost data repatriation, and secure AI workloads at scale with F5 and MinIO.

Learn more about F5 and MinIO or send an email to minio@f5.com.

