


Scaling AI Data Ingestion with MinIO AIStor, Red Hat OpenShift, and F5 BIG-IP

Integrating F5® BIG-IP® Local Traffic Manager™ (LTM) with Red Hat® OpenShift® and MinIO AIStor (S3 Compatible Storage) transforms AI training data ingestion by elevating performance, scalability, and security. The combined solution streamlines the flow and protection of data, ensuring that AI development processes run efficiently and remain resilient. This approach helps organizations ingest large volumes of data without compromising security, speed, or integrity, ultimately accelerating time-to-insight for AI initiatives.



Key Benefits

High Performance at Scale

Effortlessly manage large-scale AI data ingestion with F5® BIG-IP® purpose-built hardware-accelerated architecture, ensuring high throughput and sustained performance for increasingly demanding workloads.

End-to-end security for sensitive AI data

Protect AI training datasets both in transit and at rest with robust security policies and automated traffic management, safeguarding intellectual property and ensuring compliance across hybrid deployments.

Scalable, unified AI storage pipeline

By integrating MinIO AIStor high-performance, S3-compatible object storage with Red Hat OpenShift, the solution ensures dynamic resource allocation, simplifying scalability for fluctuating workloads and eliminating bottlenecks in the AI process.

Seamless Integration

Red Hat OpenShift integrates seamlessly with MinIO AIStor object storage, providing scalable and high-performance storage solutions for AI/ML workflows.

Keeping data consistent, up to date, and secure in distributed environments becomes exponentially more challenging when faced with surging workloads in the exabyte range—far exceeding the capacity of manual processes to keep pace.

Overcome Large-Scale AI Data Ingestion Hurdles

Ingesting AI training data into S3-compatible storage often comes with obstacles tied to capacity, speed, and security. As volumes of data surge into the exabyte range, traditional architectures struggle to handle the influx without inducing bottlenecks or risking data loss. Managing so much incoming information also demands a level of performance that many legacy systems cannot sustain, introducing inefficiencies that can delay AI development timelines.

Moving data from cloud storage to on-premises environments adds further complexity. Organizations must ensure their on-premises solutions match or surpass the scalability and reliability of the cloud. This shift may introduce significant hardware and software investments, as well as the need for comprehensive security measures that protect data in transit and at rest. Any gaps in these areas can compromise both performance and trust in the AI pipeline.

The interplay of hybrid cloud environments raises additional concerns around data governance and accessibility. Keeping data consistent, up to date, and secure in distributed environments becomes exponentially more challenging when faced with surging workloads in the exabyte range—far exceeding the capacity of manual processes to keep pace. Without the right strategy, organizations risk fragmentation, difficulties in resource allocation, and unforeseen operational costs that can stall AI innovation.

Build a Scalable, Secure AI Data Pipeline with F5, Red Hat OpenShift, and MinIO

F5 BIG-IP LTM forms the backbone of a holistic solution that streamlines AI data ingestion. F5 purpose-built hardware, F5® rSeries and VELOS, are engineered to handle heavy data flows, balancing traffic efficiently and maintaining robust security policies to prevent unauthorized access and data corruption.

MinIO AIStor provides S3-compatible object storage that seamlessly integrates with Red Hat OpenShift, enabling a high-performance, scalable foundation for AI training data. This integration automates resource allocation and accelerates data operations, giving teams the agility to handle fluctuating demand. By pairing AIStor with OpenShift, organizations can more easily maintain data consistency, optimize performance, and eliminate bottlenecks that hinder AI development cycles.

Bringing these components together addresses the performance, scalability, and security challenges inherent to large-scale AI data ingestion. The solution's unified approach safeguards data in motion and at rest, supports hybrid cloud or on-premises deployments, and minimizes the overhead of manually coordinating complex workflows. With this framework in place, users achieve faster, more reliable data processing and can focus on refining AI applications rather than troubleshooting infrastructure.

Key Benefits

Accelerate Data Repatriation and Replication

As demand stabilizes and planning improves, data can be moved back to its original location for better performance and cost efficiency, leveraging regional replication and backup to provide secure and accessible data across different regions while minimizing the premium costs associated with cloud elasticity.

S3-Compatible Storage for Hybrid Cloud Deployments

MinIO AIStor offers high-performance, scalable solutions that seamlessly integrate with Red Hat OpenShift, optimizing resource utilization and accelerating AI development processes.

Scalable to Your Business

MinIO can be scaled to meet and exceed your business needs, from public clouds and on premises to Kubernetes and bare metal.

Figure 1: Optimized application delivery for distributed data ingestion infrastructure, designed to handle high-volume training data with efficiency. This solution ensures availability, resiliency, and load balancing to support large-scale training datasets crucial for machine learning and neural network training. Additionally, it facilitates high-speed ingestion and seamless movement of data objects for AI training.


Drive AI Innovation with an Integrated, High-Performance Architecture

The combined solution of BIG-IP, MinIO AIStor, and Red Hat OpenShift delivers a cohesive way to ingest and manage extensive AI training data streams. By aligning advanced load balancing with container orchestration and high-performance storage, organizations overcome the primary obstacles that slow AI initiatives—namely performance bottlenecks, security gaps, and difficulties scaling across hybrid environments.

Through hardware-accelerated traffic management, consistent policy enforcement, and streamlined container workflows, overall data handling becomes more agile and predictable. This architecture checks the boxes for robust security without sacrificing speed or functionality, allowing massive amounts of data to be moved without introducing risk. It also ensures that workloads scale seamlessly as AI models evolve and require deeper or more repetitive training.

With these capabilities in place, teams benefit from a faster path to iteration and development, unencumbered by infrastructure constraints. Massive volumes of data remain continuously accessible and secure, allowing analytical insights to be derived efficiently. Ultimately, this integrated approach powers the next generation of AI applications, supporting innovation while keeping operational complexity in check.

F5 helps safeguard data in motion and at rest, supporting hybrid cloud and on-premises deployments, while minimizing the overhead of manually coordinated and complex workflows.



Next Steps

F5 and Red Hat connect and secure AI apps and APIs powering enterprises with validated joint solutions.

F5 and MinIO enhance AI workloads with high-performance object storage and distributed application services.

Find out how F5 products and solutions can enable you to achieve your goals.

