

S3 Object Storage for AI and Data-Intensive Workloads

Transform your data pipeline with WEKA's purpose-built S3 interface

Challenges

- Traditional S3 struggles with high-performance, low-latency demands.
- Inefficient handling of small objects creates processing bottlenecks.
- Scaling workloads introduces latency and performance degradation.
- Lack of unified access leads to fragmented operations and data silos.

Solution

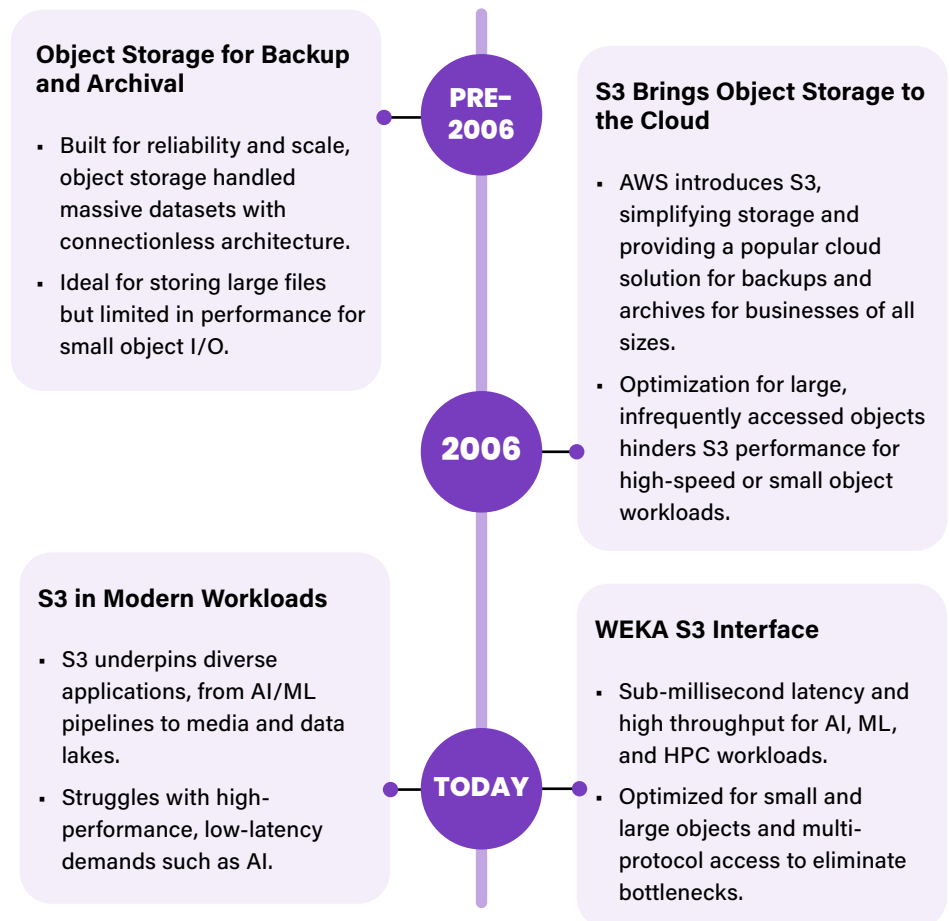
- WEKA's S3 interface ensures ultra-low latency and high throughput for all workloads.
- Optimized for small-object handling, AI pipelines, and real-time data demands.
- Multi-protocol support bridges traditional and cloud-native applications.
- Built for seamless scalability and consistent performance.

Benefits

- Accelerate AI and HPC workloads with 20x faster throughput.
- Maintain sub-millisecond latency for real-time applications.
- Eliminate data silos with unified hybrid cloud support.
- Reduce API overhead for better resource utilization.

The WEKA® Data Platform offers fully compliant native S3 protocol access, seamlessly integrated with its multi-protocol capabilities, including POSIX, NFS, SMB, NVIDIA GPUDirect Storage, and CSI plug-in support. Designed for high-performance workloads, WEKA delivers line-rate S3 performance for small objects and files with linear scalability, enabling unparalleled speed and efficiency for modern data pipelines.

A Brief History of Object Storage



WEKA vs. Traditional S3: Closing the Gap in Performance for Modern Workloads

Standard S3 storage systems provide reliability and scalability but lack the performance needed for modern workloads. High latency, inefficient small-object handling, and bottlenecks in metadata processing often lead to delayed insights and operational inefficiencies.

WEKA's S3 interface addresses these limitations by fully implementing the S3 protocol using a purpose-built architecture designed for performance-sensitive applications. WEKA's distributed metadata engine removes traditional congestion points, enabling the processing of millions of small objects without compromising speed or efficiency. WEKA's S3 implementation is specifically optimized for high-performance workloads, offering ultra-low latency and high throughput for millions of small objects. WEKA's architecture avoids problems inherent in traditional S3 solutions designed primarily for large objects and archival use cases that get overwhelmed by the protocol's chattiness. By leveraging NVMe SSDs, kernel bypass, and distributed metadata handling, WEKA delivers ultra-low latency and high throughput, even at scale. These innovations ensure consistent performance across large datasets, unlocking the potential of AI and HPC workloads with unparalleled speed and reliability.

20x Faster Data Throughput

Accelerate small and large object workloads with WEKA's optimized architecture.

S3 API Support

S3 Artifact	Methods Supported
Buckets	HEAD/GET/PUT/DEL
Bucket Lifecycle	GET/PUT/DEL
Bucket Policy	GET/PUT/DEL
Bucket Tagging	GET/PUT/DEL
Object	GET/PUT/DEL
Object Tagging	GET/PUT/DEL
Object Multiparts	POST Create/Complete/GET/DEL/PUT/GET Parts

Built for AI and Data-Intensive Workloads

AI and data-intensive workloads require more than just storage. They demand uninterrupted integration with high-speed data pipelines that can match the pace of modern computational needs. Traditional S3 systems often fall short, creating obstructions that slow progress and waste valuable GPU cycles.

WEKA's S3 interface transforms AI workflows with a coherent high-performance platform designed to overcome these limitations. By ensuring data flows efficiently between compute and storage, WEKA reduces delays in training and inference cycles, enabling organizations to handle even the most demanding workloads. Its architecture supports real-time responsiveness, ensuring faster time-to-insight and maximizing the value of every compute resource.

The WEKA S3 Interface:

- Streamlines AI pipelines to eliminate data flow inefficiencies.
- Accelerates training cycles with real-time responsiveness for high-demand tasks.
- Optimizes GPU utilization, delivering consistent performance across workloads.

Bridge the Gap Between Traditional and Cloud-Native Applications

Modern hybrid environments demand unrestricted data access across diverse applications. Yet, traditional storage solutions often create silos, requiring duplicative systems that increase complexity and operational costs. These limitations slow workflows and hinder innovation.

WEKA's Multi-Protocol Support eliminates these challenges with a unified platform that provides access across S3, POSIX, NFS, SMB, and GPUDirect® Storage. A single dataset can now support traditional and S3-enabled applications without duplication or migration. By enabling interoperability across on-premises, cloud, and hybrid environments, WEKA simplifies operations. It enhances flexibility, allowing organizations to adapt quickly to evolving workload demands.

**Accelerate
AI model
training times
by up to
20x**

WEKA's high-speed architecture eliminates bottlenecks for faster insights.

**Millions
of IOPS**

Scale effortlessly with WEKA's high-performance architecture.

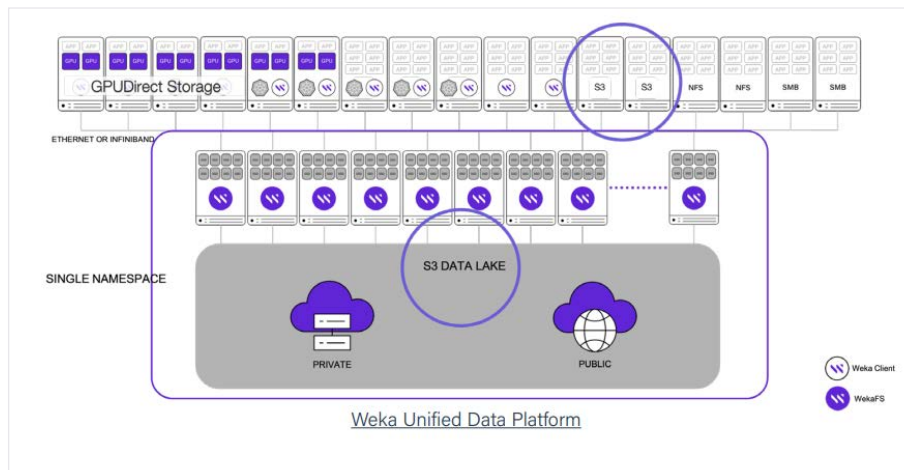
**One
Platform
for Every
Protocol**

Simplify operations and eliminate silos with WEKA.

Driving Innovation Across Industries

By accelerating data processing and removing storage bottlenecks, WEKA enables businesses to innovate faster.

- **Life Sciences:** Accelerate genomics workflows, enabling faster sequencing and analysis.
- **Financial Services:** Deliver low-latency data access for real-time trading and risk calculations.
- **Media and Entertainment:** Optimize high-speed rendering and editing for 4K/8K video production.
- **Research and Development:** Run simulations and analyze massive datasets without delays.



Why WEKA for S3?

WEKA's S3 interface delivers what traditional object storage cannot: the speed, scalability, and efficiency demanded by today's AI, HPC, and data-intensive applications. Its unique architecture eliminates bottlenecks, supports hybrid workflows with multi-protocol access, and ensures consistent performance at scale. With WEKA, organizations can achieve faster time-to-insight, reduce operational complexity, and unlock the full potential of their data.

- ✓ **Accelerate AI and HPC workflows with ultra-low latency and high throughput.**
- ✓ **Simplify hybrid cloud operations with effortless interoperability across protocols.**
- ✓ **Scale effortlessly while maintaining consistent performance for complex workloads.**



weka.io

844.392.0665

