



WekaFS™ :

The Weka File System




EPIC PERFORMANCE AND LIMITLESS SCALE

Delivers > 10 GB/sec to a single client and performance scales linearly as system grows




RUNS ANYWHERE

Runs anywhere your data lives, on any standard server hardware, whether on-premises, in the cloud, or shared across both



IDEAL FOR ANY WORKLOAD

Instantly adapt to any new workload with no knobs to tune, no expertise required




UNIFIES YOUR DATA

One global namespace for your entire data lake and easily access and manage billions of files from one directory

FASTEST FILE STORAGE FOR AI AND TECHNICAL COMPUTING

New applications in technical computing, analytics, and deep learning have placed unprecedented demands on storage I/O, throughput, latency, and scalability. WekaFS is the world's fastest shared parallel file system and delivers unmatched performance at ANY scale while offering the same enterprise features and benefits of traditional storage. It meets all storage challenges, delivering 10x the performance of legacy network attached storage (NAS) systems and 3x the performance of local server storage. WekaFS is an innovative file system, built for those who solve big problems, and runs anywhere on any standard server hardware, whether on-premises, in the cloud, or shared across both. It gives you control of hardware choices and achieves dramatic cost savings for IT infrastructure.



1MB Read Performance to a Single Client

WekaFS software is optimized to leverage the speed and low latency of NVMe technology and supports both small and large file access — either randomly or sequentially — at the lowest latency. It can be deployed on any standard X86 infrastructure as a dedicated storage server (appliance model) or integrated into the application cluster (converged). The software also supports a hybrid cloud model, allowing enterprises to leverage on-demand public compute resources for cloud-bursting, remote backup, and disaster recovery.

WekaFS is ideally suited for performance-intensive workloads such as artificial intelligence, deep learning, data analytics, life sciences research, financial modeling, log management, engineering development, media rendering, and government/university research.

“WekaFS was the clear choice for our on-premises deep neural network training... a NAS solution would not be able to scale to the extent we would need it to...and WekaFS was the most performant of all the parallel file systems we evaluated.”

Dr. Xiaodi Hou, Co-founder and CTO



BETTER THAN LOCAL FILE SYSTEM PERFORMANCE

Poor NAS performance has forced many organizations to serve data from local flash drives for improved application performance. But copying data to a server from a networked storage system defeats the purpose of a shared file system and adds considerable wall clock time to production runs. WekaFS delivers faster performance than a local file system, and can fully saturate a 100 Gbit network link, delivering over 10 GB/sec to a single application client. Predictable and seamless performance scaling is also a challenge for traditional NAS filers and legacy parallel file systems, resulting in poor utilization of expensive compute resources. With WekaFS, performance scales linearly as the infrastructure grows.

EXASCALE CAPACITY WITH INTEGRATED TIERING

WekaFS delivers the features that enterprise customers have come to expect. It supports instantaneous snapshots, clones, remote snapshots to object storage for disaster recovery, user quotas, high availability, direct-to-cloud backup, private cloud multi-tenancy, and cloud-bursting. The patented data layout and protection scheme distributes both data and metadata evenly across the entire file system to ensure the highest level of protection and eliminate hot spots and I/O bottlenecks. The storage cluster can survive up to four concurrent failures without loss of service. The WekaFS software has advanced authentication and encryption suitable for the most demanding enterprise environments. Integrated tiering to any S3 object store delivers the highest performance and maximum scalability, all in a single namespace.

Weka's unique snap-to-object feature allows users to easily create a replica of the production data and instantly push it to any S3 object store – on-premises or in the cloud – enabling snapshot-based replication. The cloud copy can be used for workload migration to another application cluster or to provide a fast recovery point objective (RPO) service guarantee.

USE CASES

AI and Analytics • Machine learning, fraud detection, retail optimization, voice analytics, and IoT data analytics

Life Sciences • Next-generation genomics sequencing, microscopy, bioimaging, structural biology, and informatics

Financial Services • Algorithmic trading, time series analysis, risk management, and market simulation

Manufacturing • EDA simulation and verification, software builds, CFD, and thermal imaging

Media and Entertainment • Nonlinear editing, VFX rendering, transcoding, and content delivery

Academic/Government • Climate change simulation, computational physics, earthquake studies, space research, and intelligence

SPECIFICATIONS	
Software	WekaFS
Virtual File Systems	Up to 1024 unique configurable virtual file systems per storage cluster
Data Protection	Distributed data protection (N+2 or N+4)
Supported Protocols	POSIX, NFS, SMB, S3 via gateway
Management Interface	GUI, CLI
Snapshots and Clones	Up to 4096 snapshots
Error Detection	End-to-end data protection
System Monitoring	Proactive cloud-based monitoring and analytics for application tuning
Alerts	Configurable to Email, SMS, and Slack
Access Control	User Authentication, LDAP, extended ACLs
Security	Encryption in flight and at rest, Key Management System integration
Tiering	S3 compatible cloud (public or private)
Hardware Vendors	Dell, HPE, Lenovo, Penguin Computing, Supermicro
Certified Object Stores	Amazon S3, Cloudian, IBM COS, Quantum (ActiveScale), Scality, SwiftStack S3

