

FAQ for WekaIO and WekaFS™:

The Weka File System

GENERAL

Q. Who is WekaIO?

WekaIO is a private company building high-performance storage solutions for compute and data-intensive workloads such as machine learning, GPU acceleration, genomics, life sciences, financial analytics, log management, and technical computing. The leadership team has a long legacy of storage expertise from industry-leading companies such as IBM, XIV, Intel, NetApp, Panasas, and EMC. The core engineering team came from XIV, an acclaimed scale-out storage startup acquired by IBM in 2007. With the rapid adoption of cloud computing, the team saw an opportunity to revolutionize how storage is provisioned in the datacenter. The company has a significant portfolio of intellectual property with over 40 patents filed and many already granted.

Q. Who are the company's investors?

The company was founded in 2013 and has raised over \$67M in funding from leading venture capital firms, including Walden Riverwood Ventures, Norwest Ventures, Qualcomm Ventures, CID Group, and Gemini Israel Ventures. It has strategic investments from Western Digital, Seagate, Mellanox, and NVIDIA, and is a HPE Pathfinder company.

PRODUCT

Q. What does WekaIO make?

WekaIO has developed an easy-to-deploy, software-based, scale-out storage solution that provides all-flash level performance, NAS simplicity and manageability, cloud scalability, and breakthrough economics. The software runs on standard server architectures and delivers 10x the performance of traditional NAS-based architectures.

The Weka file system (WekaFS) has some very important core attributes. It is a:

- fully distributed and scalable file system
- POSIX-compliant file system
- software-defined and hardware agnostic solution
- solution with integrated cloud tiering

Q. How can WekaFS be hardware agnostic?

A key system design principle is that hardware technology changes over time; a software-defined storage solution should accommodate such changes, which means that it must be able to run on any standard hardware platform. WekaFS is software that was designed to run on any standard Intel x86-based server hardware and commodity SSDs or run natively in the public cloud. This eliminates the cost overhead of expensive specialized hardware and allows you to benefit from advances in technology without suffering the pain of forklift upgrades to next-generation architectures.

Q. What are the typical use cases that WekaFS supports?

- AI, Machine Learning, and any GPU-accelerated application
- Big Data Analytics: fraud detection, retail optimization, call center analytics, and IoT data analytics
- Life Sciences: NGS, bio-imaging, structural biology, and bioinformatics
- Financial Services: high-frequency trading, risk management, fraud detection, and market simulation
- HPC: climate change simulation, computational physics, earthquake studies, space research, simulation, and intelligence
- Media and Entertainment: nonlinear editing, VFX rendering, transcoding, and content delivery

Q. How can WekaFS be deployed?

The Weka file system is an easy-to-configure, easy-to-deploy storage solution that fluidly adapts to your environment, giving you complete deployment flexibility.

- Hyperconverged deployments leverage your existing compute infrastructure while eliminating your storage footprint and reducing power and cooling costs.
- Dedicated storage server deployment is ideal when you want to maintain separate storage and compute infrastructure for application isolation, performance, or scalability.
- Public cloud deployments allow you to realize the promise of truly elastic computing and limitless scale by running WekaFS on public cloud server instances.

Q. What environments does WekaFS support?

WekaFS provides flexibility, ease of deployment, and resiliency, whether on-premises, in a hybrid configuration, or entirely in the cloud for on-demand scalability. WekaFS is a single, no-compromise, storage solution that provides the freedom to choose the environment best suited for your application based on performance, scale, and economics.

Q. How many file systems will WekaFS support?

The Weka file system can support up to 1,024 file systems, each of which has access to all cluster resources. Each file system may have its own data management settings.

Q. Does WekaFS support data tiering?

WekaFS stores data on two tiers. Hot data is stored on flash SSDs within the storage server cluster, and warm data is tiered to any S3 compatible object storage system with hard disks for cost-optimized capacity. Tiering is transparent to the application, so no modifications are required to access data on the secondary tier.

Q. Is WekaFS susceptible to performance hot spots?

WekaFS uses a patented data and metadata distribution mechanism to avoid hot spots. The software transparently monitors cluster storage resources and automatically redistributes data to balance the load across all servers whenever an out-of-balance condition is detected.

Q. How large can WekaFS scale?

WekaFS scales independently on two dimensions—capacity and performance. WekaFS is designed to scale to hundreds of petabytes, thousands of compute instances, and billions of files. WekaFS has production deployments with hundreds of nodes and tens of petabytes of storage under management. Performance scales linearly as more nodes are added to the storage cluster. The system is capable of reaching 10s of millions of IOPS or >2.5TB/sec of bandwidth at less than 300-microsecond latency.

COMPARISON TO CLOUD SERVICES AND OTHER FILE SYSTEMS

Q. How does WekaFS compare to Amazon's EBS?

An EBS volume is simply block storage. WekaFS is a high-performance, shareable file system that enables participating instances to share the same data coherently with POSIX semantics.

Q. How does WekaFS compare to Amazon's S3?

WekaFS is a high-performance, distributed file system that allows file sharing by applications that require a file interface. Weka extends these capabilities by providing an interface for these applications to access AWS S3 storage at sub-millisecond latencies as part of a massively scalable global namespace.

Q. How does WekaFS compare to other parallel file systems?

- WekaFS greatly simplifies deployment; cluster configuration typically takes only a few minutes.
- WekaFS provides exceptionally high IOPS performance for small files (<4K) while delivering excellent bandwidth. It has achieved top ranking on SPEC 2014 and the IO-500 performance benchmark tests.
- WekaFS does not have separate metadata services that can limit performance; metadata is distributed throughout the cluster via patented mechanisms that prevent hot spots. WekaFS enables a scale-out architecture that continues to add performance as more nodes are added to the cluster.
- WekaFS uses innovations in networking to achieve the lowest latencies regardless of cluster size.
- WekaFS uses patented data protection mechanisms that provide the most resilient large cluster deployment with the shortest rebuild time in the industry.

Q. How does WekaFS compare to purpose-built storage appliances?

- WekaFS delivers 10X the performance of legacy NFS or Windows-based storage appliances.
- WekaFS allows performance and capacity to scale up or down dynamically and independently,
- WekaFS provides users with more granular control over performance for a given server or instance type.
- WekaFS performance scales linearly based on the number of cores allocated to the file system and the number of server nodes in the cluster.
- WekaFS is hardware-agnostic and enables cloud bursting to increase system agility and deliver hybrid cloud capability.

DATA PROTECTION AND AVAILABILITY

Q. How is data protected?

Weka's distributed data protection scheme consists of multiple layers of protection to ensure the highest levels of availability, performance, and data resiliency.

These layers include:

- data distribution within the cluster
- a proprietary error correction scheme that is faster, more efficient, and more resilient than RAID or other forms of erasure coding
- metadata journaling that is used to protect against file system corruption in the event of a power failure, eliminating the need to run FSK against the entire file system (a process that can take days) before placing it back in service
- an end-to-end checksum that ensures that data is written exactly as it was intended, preventing bit-flips and other forms of silent data corruption.

WekaFS can easily handle multiple failures and is configured with +2 or +4 levels of failure domain reliability.

Q. What is the impact of data protection on performance?

WekaFS delivers excellent performance and minimizes overhead by using patented data protection mechanisms coupled with the computing power of multiple participating instances. The result is almost no impact on I/O performance.

DATA ACCESS AND DEPLOYMENT

Q. What protocols does Weka Support?

WekaFS supports access via POSIX, SMB, NFS, and S3 via gateway.

Q. How do I deploy Weka with NFS Access on AWS?

It is extremely easy and only takes a few minutes with the following steps:

- configure a local or tiered filesystem group
- configure a filesystem
- make sure that you have Valid AWS IPs for the NFS connections from AWS
- configure the Weka interface groups - choose the servers that will export NFS as well as provide the AWS range of NFS IPs
- configure the Weka client groups as well as the filesystems to be exported
- mount the file systems as NFS on your clients via the Linux mount command (e.g. `mount -t nfs x.x.x.x:/filesystemname/ /mnt/weka`)

PERFORMANCE

Q. What latency can I expect when using WekaFS?

Read and write latency for file operations against active data can be as low as 200 microseconds. Factors affecting latency include the load on the cluster, tiering, and choice of access methods. NVMe will provide the lowest latency.

Q. How many IOPS can WekaFS provide?

Depending on the server and SSD used, Weka currently delivers over 3.5 million IOPS in a starter cluster. IOPS performance scales as the cluster grows.

Q. How much bandwidth can WekaFS provide?

WekaFS can utilize all available bandwidth in the cluster. Total performance depends on the cluster size and the amount of networking bandwidth provided. Users can expect to see more than 45GBytes/second on a starter cluster, and performance per node will scale as more nodes are added to the cluster.

TIERING

Q. How does data tiering work?

Tiering is specified at the file system level. When configuring tiering users specify the sizes of the flash tier and the object tier. Data will be migrated from one tier to the other based on tiering policy.

Q. What actions do I need to take for my applications to utilize tiering?

From the GUI, simply specify the tiering policy and tiering destination. No changes to applications are required. Applications will access one file system, regardless of any tiering operations. However, data tiering can affect application performance. Users can choose either the default tiering policy or define their own, and policies can be changed at any time.

Q. What data is tiered and what criteria are used?

File content is tiered, but metadata is always stored on the cluster. Files are tiered with user-defined policies based on access and modification times.

Q. When a file is tiered to the object store, is it deleted from the cluster?

The authoritative file resides on the object store while the original file remains as a cached copy for accelerated access until the physical SSD space is required for new file system writes.

Q. Can I force the tiering of a specific file?

Tiering can be forced simply by changing the file's modification time attribute. Files can also be immediately demoted to object store to accommodate data migration from a legacy system to Weka.

Q. Can I limit the bandwidth consumed by tiering?

Users can limit the bandwidth available to the tiering process (e.g. 200 MB/sec). This parameter can be changed at any time.

SNAPSHOTS

Q. How granular are snapshots?

Snapshots can be taken at the file system level.

Q. What is the performance impact of a snapshot?

Practically none as a sub-second service interruption is all that would be experienced across the entire file system.

Q. How many snapshots can be taken?

4,096

Q. How much space does a snapshot consume?

Space consumption is differential and depends on the difference between the snapshot and the master file system.

Q. Can a snapshot be a full clone?

Yes. A snapshot can be promoted to be a full clone such that applications can write to the clone. Space consumption is still differential.

Q. What are the use cases for snapshots?

- backup, where a daily or periodic snapshot is copied to another storage service for data recovery in case of data corruption
- archival, where old copies of the data are kept for compliance and applicative needs
- DevOps and testing, where new application releases can be tested against production data by making clones of the data that can be destroyed later

DATA SECURITY

Q. What kind of data security does WekaFS support?

WekaFS has full encryption all the way from the application clients to the storage system, supporting encryption on the fly and at rest. It is tightly integrated with key management systems that are KMIPS compliant.

Q. What kind of authentication does WekaFS support?

WekaFS has support for client-server authentication to prevent rogue clients from accessing the storage cluster. In addition, it has full integration with directory services for user authentication and permissions.

PRICING AND LICENSING

Q. How do I get started?

Getting started with Weka is easy. The software is available through our network of qualified reseller partners and can be purchased fully integrated with hardware platforms from select OEMs. Contact us for a list of resellers near you to schedule a free Proof of Concept (POC) test.

For increased IT agility and scalability, WekaFS is also available in Amazon's Marketplace, or you can bring your own license. From Marketplace, simply deploy a Weka cluster for on-demand flexibility. If you don't have an Amazon Marketplace account, you can purchase a license for Amazon from a qualified reseller and bring it with you.

Q. How is the software packaged?

Weka software includes everything you need for high-performance, scale-out primary storage. There are no additional license fees for standard enterprise features.

Q. How is WekaFS licensed?

WekaFS is licensed as an annual subscription that includes our white-glove support. The software price is based on usable capacity under management. You only pay for the capacity used by the file system and can increase your entitlement to manage more capacity with WekaFS as your needs grow. The license enables tiering to an object store (on-premises or in the cloud). Pricing for the tiered storage is on a consumption basis (per Terabyte). For specific pricing details or help with custom configurations, please contact your authorized reseller.

Q. What is WekaIO's software upgrade policy?

Weka license holders are entitled to all software fixes and product enhancements as part of the base product license. From time to time, WekaIO may offer high-value features as optional product upgrades available for purchase. Such features would not be included in the base product license entitlement.

Q. What if I let my license expire? Can I still access my data?

WekaIO will never prevent you from accessing your data. If your license expires, you will no longer receive support, maintenance, or bug fixes. You will have free and clear access to your data on the system and active use of the data.

Q. How will I know if my license has expired?

WekaFS has an entitlement assigned to every storage cluster that matches the license term.

Q. How can I reinstate support after it has lapsed?

You can reinstate the software license at any time after expiration by simply contacting your authorized WekaIO reseller.

