

# FAQ for WekaIO Matrix™

## GENERAL

### Q. Who is WekaIO?

WekaIO is a private company building innovative cloud storage solutions for on-premises or cloud deployment. The leadership team has a long legacy of storage expertise from industry leading companies such as IBM, XIV, Intel, NetApp, Sandforce 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 that includes 7 patents issued, 16 patents filed, and 40 more identified.

### Q. Who are the company's investors?

The company was founded in 2013 and has raised over \$32M in funding from leading venture capital firms including Walden Riverwood Ventures, Norwest Ventures, Qualcomm Ventures, and Gemini Israel Ventures.

## PRODUCT

### Q. What is WekaIO Matrix™?

WekaIO Matrix is 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.

WekaIO Matrix consists of:

- MatrixFS POSIX file system
- Integrated cloud tiering system
- Trinity management software

### Q. What is meant by Zero-footprint Storage?

By leveraging existing application compute resources or public cloud based compute instances and storage, Matrix eliminates the need to deploy a dedicated storage infrastructure. This results in zero added footprint, as well as huge savings in power, cooling, floor space, infrastructure, and management costs.

### Q. How can Matrix be hardware agnostic?

A key system design principle is that technology changes over time; a true software defined storage solution should accommodate such changes, which means that it must be able to run on any hardware. Matrix is software that was designed from the beginning to run on any standard Intel x86 based server hardware and commodity SSDs. This eliminates the cost overhead of expensive specialized hardware and allows you to take advantage of improvements in technology without the pain of a forklift upgrade to next generation architectures.

### Q. What are example use cases that Matrix supports?

- Web 2.0: Cloud-hosted services (on-premises or public cloud), cloud bursting, web applications, Dev/Test, and content sharing
- Big Data Analytics: Fraud detection, retail optimization, call center analytics, and IoT data analytics
- Life Sciences: NGS, bio-imaging, structural biology, bioinformatics
- Media and Entertainment: Nonlinear editing, VFX rendering, transcoding, and content delivery
- Financial Services: High frequency trading, risk management, fraud detection, and market simulation
- Academic/Government: Climate change simulation, computational physics, earthquake, space research, and intelligence

### Q. How can Matrix be deployed?

WekaIO Matrix is a single, easy-to-configure, easy-to-deploy storage solution that fluidly adapts to your environment giving you complete deployment flexibility.

- Hyperconverged deployments leverage your existing 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. It also works for environments where the blade-compute infrastructure that does not allow for local SSD.
- Public cloud deployments allow you to realize the promise of truly elastic computing and limitless scale by running Matrix on public cloud server instances.

### Q. What environments does Matrix support?

Whether your applications run on-premises for performance, in a virtual environment for ease of deployment and resiliency, or entirely in the cloud for on-demand scalability, Matrix is a single, no-compromise, storage solution providing the freedom to choose the environment best suited for your application based on performance, scale, and economics.

**Q. How many file systems will Matrix support?**

The Matrix global namespace can contain up to 1,024 file systems, where each file system has access to all the cluster's resources. Each file system may have its own quality of service and data management settings.

**Q. Does Matrix support data tiering?**

Matrix stores data on two tiers. Hot data is stored on SSD within the compute cluster or in a dedicated storage server cluster. Cold data is tiered to any S3, Swift, or REST compatible object storage device, including S3 enabled NAS filers for cost optimized capacity. Tiering remains transparent to the application, so no application modifications are required to access data from the cold tier.

**Q. Is Matrix susceptible to performance hot spots?**

Matrix 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 situation is detected.

**Q. What is MatrixFS?**

MatrixFS is a distributed, scalable, high performance file system that transforms a collection of standard servers with locally attached SSDs into a POSIX compliant global namespace. MatrixFS runs on either AWS instances or an on-premises server cluster and is a key component of WekaIO's Matrix storage solution.

**Q. How large can MatrixFS scale?**

MatrixFS scales independently on two dimensions—capacity and performance. The design is scalable to hundreds of petabytes, thousands of compute instances and billions of files. MatrixFS has been tested in configurations with over a 1000 server nodes and hundreds of PBs.

Performance scales linearly with the number of cloud instances or the number of CPU cores allocated to the Matrix processes. The system is capable of reaching 10s of millions of IOPS or >2.5TB/sec of bandwidth at sub-millisecond latencies.

## COMPARISON TO CLOUD SERVICES AND OTHER FILE SYSTEMS

**Q. How does MatrixFS compare to Amazon's EBS?**

An EBS volume is simply block storage. MatrixFS is a high-performance sharable file system such that participating instances share the same data coherently with POSIX semantics.

**Q. How does MatrixFS compare to Amazon's S3?**

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

**Q. How does MatrixFS compare to other distributed files systems?**

- MatrixFS greatly simplifies deployment; cluster configuration typically takes only a few minutes
- MatrixFS provides exceptionally high IOPS performance for small files
- MatrixFS has no dedicated metadata server that can limit performance, and metadata is distributed throughout the cluster via patented mechanisms that prevent hot spots and enable a scale-out architecture
- MatrixFS uses innovations in networking to enable the lowest latencies regardless of cluster size
- MatrixFS uses patented data protection mechanisms that provide the most resilient large cluster deployment with the shortest rebuild time in the industry

**Q. How does MatrixFS compare to purpose built storage appliances?**

- MatrixFS allows performance and capacity to scale up or down dynamically and independently
- MatrixFS provides users with more granular control over performance for a given server or instance type
- MatrixFS performance scales linearly based on the number of cores allocated to the file system and the number of server nodes in the cluster
- MatrixFS is hardware agnostic and enables cloud bursting to increase system agility, reliability, and performance

## DATA PROTECTION AND AVAILABILITY

**Q. How is data protected?**

The Matrix distributed data protection scheme (MatrixDDP™) 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 is used to protect against file system corruption in the event of a power failure. This prevents the need to run FSCK against the entire file system before placing it back in service, a process that can take days
- An end-to-end checksum ensures that the data written is exactly as it was intended. This prevents bit-flips and other forms of silent data corruption

MatrixDDP can easily handle multiple failures and is configured automatically based on the cluster size and other parameters.

**Q. What is the impact of MatrixDDP on performance?**

MatrixFS delivers excellent performance and minimizes data protection overhead by using patented data protection mechanisms coupled with the computing power of multiple participating instances.

**Q. Is MatrixDDP resilient to cluster node failures?**

The MatrixDDP algorithm can be configured so that even upon the failure of an entire availability zone, all data is safe and access remains uninterrupted.

## DATA ACCESS AND DEPLOYMENT

### Q. What protocols does Matrix Support?

Matrix supports access via SMB, NFS, REST, HDFS, and via a kernel driver.

### Q. How do I deploy Matrix with NFS Access?

It is extremely easy and only takes a few minutes:

- 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`)

### Q. How do I deploy Matrix with WekaIO direct access?

It is extremely easy and only takes a few minutes:

- Configure a local or tiered filesystem group
- Configure a filesystem
- Mount the weka filesystem on your weka hosts in one of the supported mount modes using the linux mount command (e.g. `mount -t wekafs filesystemname /mnt/weka`)

## PERFORMANCE

### Q. What latency can I expect using Matrix?

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 Matrix provide?

Depending upon the server and SSD used, Matrix currently provides 20K-50K IOPS per core. Total performance depends on the cluster size.

### Q. How much bandwidth can Matrix provide?

Matrix can utilize all available bandwidth in the cluster. Total performance depends on the cluster size and the amount of networking provided. Users can expect to see more than 300MB/sec per core in a typically configured 10Gbit network attached cluster, performance per node will scale as a larger network pipe is added.

## TIERING

### Q. How does data tiering work?

Tiering is specified at the file system level. When configuring tiering, users can specify the percentage of data to be stored on the cluster versus on S3 (e.g. 10% on the cluster, 90% on S3). Tiering can be enabled on any file system at any time.

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

Simply specify the tiering policy and tiering destination upon configuration. No changes to applications are required. Applications will access one file system, regardless of any tiering operations. However, data tiering may 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 based on what criteria?

File content is tiered, but metadata is always stored on the cluster. Files are tiered by user defined policies based on the 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 tiering of a specific file?

Tiering can be forced simply by changing the file's modification time attribute.

### 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 either at the file or file system level.

### Q. What is the performance impact of a snapshot?

Practically none, a sub-second (and for small clusters even sub 100ms) service interruption is all that would be experienced across the entire file system.

### Q. How many snapshots can be taken?

1,024

**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 snapshot?**

- Physical backup, where a daily snapshot is copied to another storage service
- Logical backup, where copies are kept periodically and can be used in case of logical data corruption
- Archive, 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

## PRICING AND LICENSING

**Q. How do I get started?**

Getting started with Matrix is easy. The software is available through our network of qualified reseller partners. Contact us for a list of resellers near you or to schedule a free Proof of Concept (POC) test.

For increased IT agility and scalability, Matrix is also available in Amazon's Marketplace, or you can Bring-Your-Own-License. From Marketplace, simply deploy a Matrix cluster for on-demand flexibility. Don't have an Amazon Marketplace account? You can purchase a Matrix for Amazon license from a qualified reseller and bring it with you.

**Q. How is Matrix software packaged?**

Matrix software includes the foundation MatrixFS file system and Trinity management system, everything you need for high performance, scale-out primary storage.

**Q. How is Matrix licensed?**

WekaIO Matrix is licensed as an annual subscription that includes our white glove support. Flexible software licensing accommodates hyperconverged or dedicated storage server deployment models.

*Hyperconverged Licensing* - Each server (physical or virtual) that takes part in the Matrix-enabled cluster would require a license that entitles you to use up to two CPU cores per server. For even more performance, you can purchase the optional MatrixEP™ (extreme performance) license that entitles you to use an unlimited number of cores, ideal for on-demand performance scaling for bursty workloads.

In addition, the Matrix server 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).

Licensing example: a 30-server cluster using 2 cores per server requires 30 Matrix server licenses. Because licensing is server-based, you can have as many or as few SSDs in a server and still pay only one price.

Need more than 60 cores? Purchase the additional MatrixEP license and scale to as many cores as the server can support.

*Dedicated Storage Server Licensing* - This deployment model is ideal for large capacity deployments when you choose to separate compute and storage resources. WekaIO's Matrix is licensed on a Per Raw Terabyte basis. There are no per-core or server based charges. You can use as many servers as you need to support your capacity requirements. Capacity charges per terabyte vary based on the storage media. Object storage capacity is charged at a significantly lower cost per terabyte compared to SSD.

For specific pricing details or help with custom configurations, please contact your authorized reseller.

**Q. What is WekaIO's software upgrade policy?**

Matrix license holders are entitled to all software fixes and minor product updates 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 be entitled to support, maintenance, or bug fixes but you will have free and clear access to your data.

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

You can reinstate support at any time after expiration by simply contacting your authorized WekaIO reseller. WekaIO will charge a recertification fee equal to the amount of support costs accrued from the time of license expiration to the time of reinstatement.



2001 Gateway Place, Suite 400W, San Jose, CA 95110 USA T 408.335.0085 E [info@weka.io](mailto:info@weka.io) [www.weka.io](http://www.weka.io)

©2018 All rights reserved. WekaIO Matrix, Trinity, MatrixDDP, MatrixFS, the WekaIO logo and Radically Simple Storage are trademarks of WekaIO, Inc. and its affiliates in the United States and/or other countries. Other trademarks are the property of their respective companies. References in this publication to WekaIO's products, programs, or services do not imply that WekaIO intends to make these available in all countries in which it operates. Product specifications provided are sample specifications and do not constitute a warranty. Information is true as of the date of publication and is subject to change. Actual specifications for unique part numbers may vary.

W04r1DS201805