

WEKA for AWS

Fastest Time to Results

- Cut applications run-times from hours or days to minutes
- Ensure latency-sensitive applications run smooth
- Scale performance in line with application needs

Minimize Infrastructure Costs

- Store your data on S3, enjoy local-SSD performance
- Save even more with S3 Glacier Instant retrieval
- Easily Shut down and spin up projects over the weekend
- Eliminate multiple copies of the same data

Serve All Your Applications from Your Data Lake

- Run Windows, Linux, S3 and POSIX applications to a single data set
- Share data between applications, maintaining file consistency
- Never copy your data again

High-Performance Computing With WekaFS on Amazon Web Services (AWS)

The Amazon cloud infrastructure is an ideal platform for rapid deployment of new services and applications. Enterprises can leverage the latest high-performing, on-demand compute infrastructure as a service without the need for huge investment in CAPEX-intensive computer equipment. For artificial intelligence (AI) and technical compute workloads that require high-performance computing (HPC) methods to accelerate design and discovery, the AWS Cloud is the ideal platform to provide an agile compute environment. WekaIO™ (WEKA) delivers the WEKA File System (WekaFS™), the fastest, most scalable file system for AWS, to ensure applications on EC2 instances never have to wait for data.

WekaFS is the only storage solution that supports auto scaling groups in AWS. For demanding performance applications, or peak demands, the cluster can be scaled out to meet demand. On completion of the workload, the cluster can be scaled back down to the previous scale.

Best Economics for Data-Intensive Applications

EBS (Elastic Block Services) delivers good latency but does not provide file semantics required for HPC applications. WekaFS seamlessly combines the performance tier of I3 EC2 instances with local NVMe for lowest latency with the capacity tier of S3 for massive scalability, all presented in a single namespace for best economics at Petascale. An integrated data management mechanism automatically and seamlessly moves data between the hot tier and the S3 data lake, so the relevant data is always available at highest performance. WekaFS on AWS is over 2x more cost effective than native AWS storage services at comparable performance levels, as performance is not hard bound to the capacity deployed.

One Solution for Latency Sensitive and Unpredictable Workloads

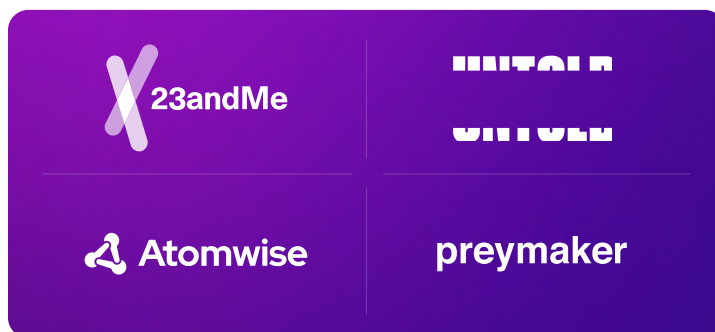
WekaFS is an ideal solution for applications that have varied I/O patterns with unpredictable performance. It excels at both latency sensitive and large file I/O with no tuning required. Additionally, WekaFS has broad multi-protocol support, allowing multiple applications (Windows, Linux, Mac, POSIX, S3 native) to share and collaborate on a single data set.

Backup and Hybrid Cloud Demands

WekaFS includes snapshots that can be used for continuous backup to S3 storage, creation of clones, or to pause and resume applications on EC2. A snapshot of the entire file system and its metadata can be saved on S3 storage so that EC2 instances can be shut down when compute resources are not needed. When applications become active again, the file system can be re-hydrated to a new EC2 cluster with a size that is different from the prior WEKA storage cluster. This feature is ideal for applications that shut down over weekends and holidays as expensive compute resources can be torn down when not in use. Snapshots can also be utilized for cloud-bursting between on-premises and AWS for elastic infrastructure scaling during peak compute demand periods. 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 and 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. The snapshot can also be saved as an immutable copy to meet compliance and regulatory requirements.

Full Integration With AWS Cloud

- WekaFS is available as a self-provisioned storage service on Amazon Marketplace.
- WekaFS can be easily deployed in the AWS cloud through cloud formation templates. To find the best fit for your application demands use the WEKA configurator found at start.weka.io
- WekaFS can be purchased through the [AWS Marketplace](https://aws.amazon.com/marketplace).



Sample Instance Configurations for WekaFS

	Starter Instance	Mixed I/O Patterns	Bandwidth Oriented	Highest Performance
Instance	I3en.2xlarge	i3en.6xlarge	I3en.12xlarge	i3en.24xlarge
Minimum Number	6 – scales to Hundreds	6 – scales to Hundreds	8 - scales to Hundreds	8 – scales to Hundreds
Raw Capacity	30TB	90TB	240TB	480
Usable Capacity	15TB	45TB	142TB	284TB
S3 Storage	Yes	Yes	Yes	Yes
Networking	Up to 25Gbit/second burst	25Gbit/second guaranteed	50Gbit/second guaranteed	100Gbit/second guaranteed
Software	WekaIO WekaFS™ File System			
Data Protection	Distributed Data Protection (N+2 to N+\$), Encryption, End-to-End Data protection,			
Protocols	POSIX, NFS, SMB, S3			
Max Filesystems	Up to 1024 per Cluster			
Snapshots	File System Level, Up to 4096 Snapshots			
Tiering	S3, support for multiple buckets, all available in the same namespace with dynamic movement from flash to object tier			
Security and Management	Active Directory, LDAP, Private VPC, Quotas			
System Monitoring	Cloud-based Monitoring and Analytics for Application Tuning and Remote Support			
Minimum Performance (Scales up with more instances)	Sequential read performance up to 4.4GB/s	Sequential read performance up to 7.2GB/s	Sequential read performance up to 19.5GB/s	Sequential read performance up to 25.1GB/s
	Up to 435K IOPs	Up to 607K IOPs	Up to 1.6M IOPs	Up to 2M IOPs

For More Information

- Speak to a representative or get more information [contact us](#)
- Check out our [Website](#)


weka.io

844.392.0665

