

WekaFS™ for AWS

Fastest, Most Scalable File System



BEST UTILIZATION OF EC2 INSTANCES

Ensure applications never have to wait for data



FAST PERFORMANCE FOR YOUR HPC APPLICATIONS

Leverage flash-native file system architecture for the highest throughput and lowest latency



BEST ECONOMICS

Use EC2 instances for performance and S3 for capacity scaling. Only use EC2 instances when applications need them

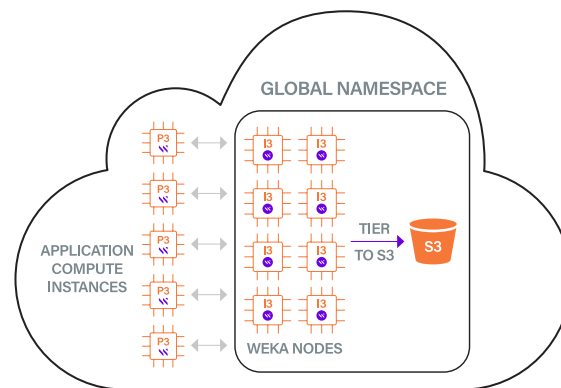


GROW YOUR INFRASTRUCTURE ON DEMAND

Expand on-premises infrastructure with cloud-bursting to AWS

HIGH-PERFORMANCE COMPUTING WITH THE WEKA FILE SYSTEM 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 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 available as a self-provisionable storage service on Amazon Marketplace.



WekaFS on EC2 Instances with Integrated Tiering to S3

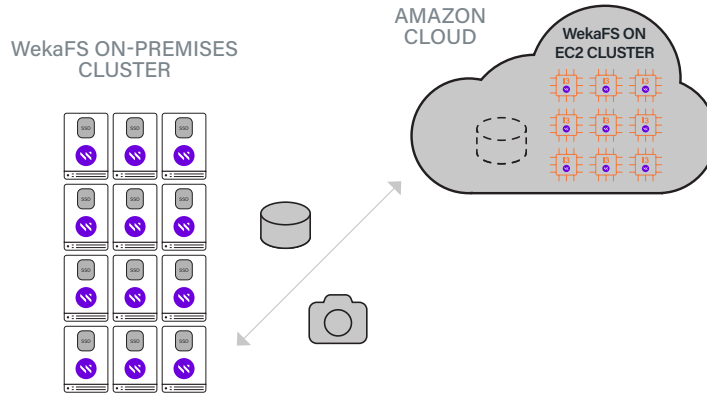
BEST ECONOMICS FOR DATA-INTENSIVE APPLICATIONS

EBS (Elastic Block Services) delivers excellent latency but is not built for the large-scale data sets found in HPC applications. WekaFS seamlessly combines the performance tier of I3 EC2 instances with the capacity tier of S3 in a single namespace to provide best economics at scale. An integrated tiering mechanism automatically and seamlessly moves data between the hot tier and the S3 data lake.

ONE SOLUTION FOR BACKUP AND HYBRID CLOUD DEMANDS

WekaFS includes snapshots that can be used for continuous backup, creation of clones, or to pause and resume applications on EC2. A snapshot of the entire file system and its metadata can be created 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. 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 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.



Snapshot your data to the cloud for continuous backup, creation of clones, or to pause and resume applications on EC2.

HIGHEST PERFORMANCE FOR FASTER RESULTS

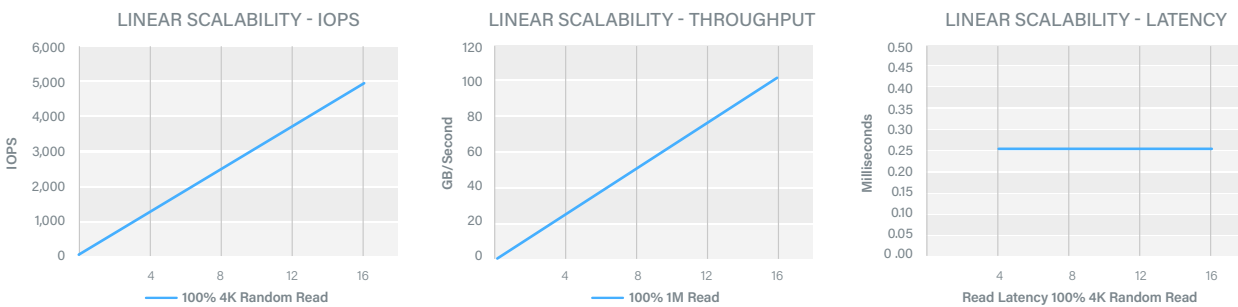
Many AI and technical compute applications have demanding IO patterns that require highly parallel file access to large data sets, testing the limits of cloud storage services. WekaFS provides an easy way to provision a single file system that delivers millions of IOPS and hundreds of gigabytes of bandwidth at very low latencies across Amazon EC2 instances.

WekaFS is a fully parallel and distributed file system that runs natively in Amazon EC2 instances with local NVMe or SATA SSD storage. Both data and metadata are distributed across the entire storage infrastructure to ensure the fastest data access without any hot-spots. WekaFS software presents a POSIX file interface to the applications and is optimized to leverage the performance of flash technology to support large and small files for both random and sequential access.

The following charts demonstrate Weka performance and scalability in the Amazon Cloud. Large file (1MB) read performance reached over 100GB/sec and small file (4KB) read performance reached over 5 million IOPS across 16 EC2 instances. Throughout the small file test, latency consistently measured below 250 microseconds across the entire cluster.

FILE SYSTEM SCALES LINEARLY WITH CLUSTER SIZE

WekaFS exhibits linear scaling of performance as the cluster size increases — and the latency at scale is unmatched.



Performance was measured using I3en.24xlarge EC2 instances in AWS

To find out more or to arrange for a free trial, contact us at info@weka.io.

To get started right away on your own self provisioned cluster go to <https://start.weka.io>.

