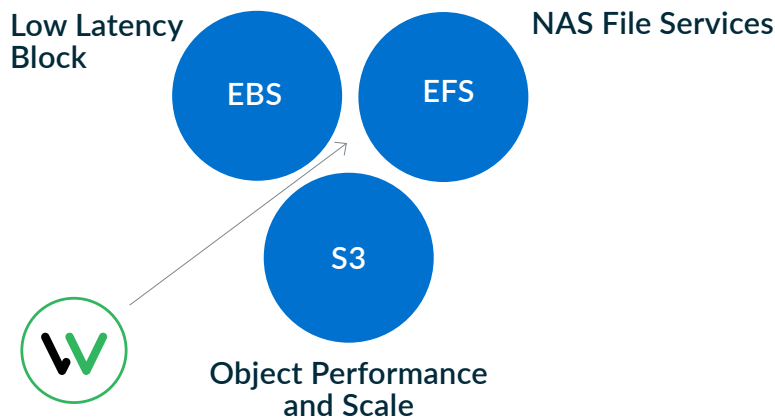


WekaIO Matrix™ for AWS

HIGH PERFORMANCE FILE STORAGE SOFTWARE FOR 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 enterprises that are considering adopting high performance computing (HPC) methods to accelerate design and discovery, the AWS cloud is the ideal platform to provide an agile compute environment. Applications that have bursty or variable compute demands will benefit from the ability to spin-up and spin-down services based on usage.

WekaIO Matrix is a unique solution in AWS that fills a void for application profiles that do not easily fit into the existing available services. S3 delivers great performance and capacity scaling but relies on eventual consistency for data sets, EBS (Elastic block services) delivers great latency but the volume size is limited to 16TB, and EFS provides shared file services with strict consistency, but has limited bandwidth. For extreme performance and scalability found in HPC applications, WekaIO Matrix can fill the void.



WekaIO Matrix is a fully distributed, fully parallel file system that was built in the Amazon cloud to deliver highest performance file services natively within the application compute infrastructure. It is hardware agnostic and can be deployed 100% in the AWS cloud or can support a hybrid IT model from on-premises to AWS cloud for peak workloads or on-demand compute. A single WekaIO cluster can scale from 6 to 300 instances in a single namespace.

WekaIO Matrix on Amazon EC2 Instances with Local SSD




MASSIVE SCALE AT SUB-MILLISECOND LATENCY

Many high performance applications have demanding I/O patterns that require highly parallel file access to large data sets, testing the limits of cloud storage services. WekaIO provides an easy way to provision a single file system that delivers millions of IOPS or hundreds of gigabytes of bandwidth at very low latencies across Amazon EC2 instances.

WekaIO Matrix is a high performance file service that runs natively in Amazon EC2 instances with local SSD storage. Matrix software presents a POSIX file interface to the HPC applications and is optimized to leverage the performance of flash technology to support both large and small file access, either randomly or sequentially.

The following charts demonstrate WekaIO Matrix performance and scalability in the Amazon cloud. Small file (4KByte size) read performance reached over 7 million IOPS across 240 r3.8xLarge instances while large file (1MByte size) read performance reached over 90GBytes/second across the same 240 node cluster. Throughout the test, latency was consistently maintained below 500 microseconds across the entire cluster.

~30K OPS/AWS Instance
~375MB/sec/AWS Instance
<400 microsecond latency




INTEGRATED TIERING AND SNAPSHOTS TO S3 OBJECT STORE

Research from leading hyper-scale web companies has shown that over 80% of data stored in the cloud is rarely or never accessed. SSD capacity is a premium on EC2 instances and WekaIO has developed a transparent tiering capability to AWS S3 to ensure customers can take advantage of lower cost storage for their cold data. The Matrix file system keeps all file metadata on the SSD tier so data is instantly available to the applications even if it is actually tiered off to S3.

Matrix global namespace is dynamic so users do not have to provision the S3 tiering capability from day one. A Matrix global namespace can begin with a minimum set of EC2 instances based on application needs and grow as the storage demands increase or as applications scale to more compute instances. The S3 object storage tier can be added to the file system at a later date as capacity requirements grow beyond the SSD capacity limit. EC2 clusters can also shrink down in SSD capacity as compute demands drop, with the excess capacity pushed to the S3 tier.

Users no longer have to maintain separate file systems and manually copy data from one tier to another, WekaIO seamlessly manages the data tiering service.

Automated Tiering to S3 for Unlimited Scalability



SNAPSHOTS FOR DEV/TEST AND DISASTER RECOVERY

WekaIO Matrix supports snapshots, which provide a point-in-time copy of application data for disaster recovery or for engineering testing. The snapshots can be cloned to take advantage of testing against production data or to create a new branch of development. The file system software can support up to 1024 snapshots per file system.


In addition, WekaIO has developed a unique snapshot capability that integrates EC2 and S3 architectures. Users can snapshot the entire filesystem plus its metadata to a remote S3 object store allowing for disaster recovery in the event of a catastrophic availability zone failure. In addition, it allows users to treat EC2 instances as ephemeral storage with the ability to shut down instances when compute resources are no longer required. The entire file system can be re-hydrated at a later date when applications are active and there is no requirement to re-hydrate a file system to the exact same instance footprint as was originally saved. For example, if the original cluster had 100 EC2 instances, the re-hydrated cluster could have only 50 instances of compute and the file system will adjust the hot tier based on the new available SSD capacity, with the remaining data on the S3 tier.

AVAILABILITY ZONE LEVEL FAULT TOLERANCE FOR UNINTERRUPTED SERVICE

Multiple availability zones provide site level fault tolerance for applications running in the Amazon cloud. WekaIO Matrix provides a high performance storage tier across multiple availability zones maintaining strict consistency. For applications that are critical to business continuity WekaIO Matrix will maintain service even in the event of a complete data center outage.

SUPPORTED CONFIGURATIONS

WekaIO Matrix is a very flexible software solution that runs on r3 and i3 based EC2 instances with locally attached SSD providing storage services and performance. In addition, applications running on compute-only instances with no local storage (such as c4 instances) can access a cluster of Matrix enabled EC2 instances via Matrix native client software for highest performance, or via the NFS protocol. This configuration provides the ultimate flexibility and value as lower cost compute instances can be leveraged for CPU intensive applications.



WEKA.IO

2001 Gateway Place, Suite 400W, San Jose, CA 95110 USA T 408.335.0085 E info@weka.io www.weka.io

©2017 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.

W06DS201707