

WEKA on Google Cloud Platform

A Leap in Scalability, Flexibility, and Performance for Cloud

Challenges

- Realizing the benefits of the cloud for modern applications
- Running large or resource-intensive workloads in the cloud at scale
- Leveraging the benefits of the cloud to make progress in High-Performance Computing (HPC)

Solution

- WEKA on GCP delivers a leap in scalability, flexibility, and performance with ease-of-use and affordable enterprise features for life sciences, financial analytics, VFX, AI/ML and HPC

Benefits

- Autoscaling to add or remove VM instances for linear performance, flexibility, and cost savings
- Expansion of the namespace through interoperability with S3 object storage
- High performance for both data and metadata with no tuning required
- Multiprotocol support: POSIX, NFS, SMB, S3, and CSI access to a single data set

The Challenge: Effective Use of the Cloud for Modern Workloads

There are many potential benefits to moving workloads to the cloud in terms of scalability, flexibility, and performance, and they have accelerated migration to the cloud. However, there is a struggle to realize the benefits of the cloud for modern applications. Running large or resource-intensive workloads in the cloud at modern enterprise scale is extremely challenging. This hinders progress in Life Sciences, Financial Analytics, Media Creation, AI/ML, and High-Performance Computing (HPC).

WEKA on Google Cloud Platform (GCP) addresses these challenges and delivers a leap in scalability, flexibility, performance, and ease-of-use for modern cloud-based workloads.

WEKA Data Platform

The WEKA Data Platform is the first multicloud data platform for AI and next-generation workloads. It provides a single, unified product with consistent high performance, robust data services, a seamless, simple data management experience, and best-in-class economics—regardless of the location of your data and applications.

WEKA combines the performance tier of NVMe-enabled compute instances in GCP with the capacity tier of native object storage in a single namespace. WEKA provides the lowest latency and massive scalability for best economics from just a few terabytes up to exascale. Integrated data management automatically moves data between the hot tier and the object storage data lake, so actively accessed data is always available at the highest performance. Because you can scale performance and capacity independently with WEKA, its cost is usually half that of native storage services at comparable performance levels.


WEKA enables modern high-performance applications to migrate to, and run on, GCP. WEKA excels at both latency-sensitive and large file IOs with no tuning required. Its cloud-native design and software also runs on-premises, so WEKA's scalable architecture can be deployed as a cloud-native or hybrid solution. WEKA provides a single global namespace and a robust enterprise feature set, including data protection, snapshots, and replication, all with the highest performance in the cloud for any application.

Cloud Scalability, Flexibility, Performance, and Ease-of-Use

The synergy between WEKA and GCP delivers distinct advantages. WEKA leverages GCP capability with autoscaling groups to automatically add or remove VM instances based on workload changes to provide scalability, flexibility, and cost savings. For applications with varying demands on compute resources, WEKA autoscaling ensures consistent support at optimal cost. There is no need to size for

peak demand as performance can scale with application needs. Unlike other solutions that only allow for growth, WEKA autoscaling groups support shrinking to reduce customer spend after spikes and achieve true cloud elasticity. All of these scale-out and scale-in capabilities are provided online without any impact to the system and/or applications.

Broad multiprotocol support for POSIX, NFS, SMB, S3, and the Kubernetes CSI plugin allows many applications to access a single data set. This especially benefits AI/ML pipelines as multiple users and applications can simultaneously share and collaborate with full data consistency across protocols. All data is available through all protocols at the same time, eliminating the need for internal and/or external movement to utilize data with different protocols in various stages of the pipeline.



WEKA leverages Terraform and autoscaling groups for seamless automation and integration with your GCP ecosystem for ease of orchestration and automation.

WEKA on GCP delivers maximum acceleration of your business outcomes by matching the speed of your compute and data storage resources. WEKA provides integrated tiering to S3 object storage for maximum scalability to exascale capacity and best economics. Its patented data layout and protection scheme distributes data and metadata evenly across the entire file system for the highest level of protection and eliminates hot spots and performance bottlenecks. This applies not only to small file and metadata-intensive applications but also to workloads requiring massive throughput scale-out. The WEKA data platform needs no tuning for small or large files and is ideal for mixed workloads with unpredictable data patterns. The storage cluster can be configured to survive four concurrent node failures with no loss of service.

Affordable Enterprise Features

WEKA on GCP also delivers enterprise features that customers expect, including data reduction, user quotas, high availability, encryption, multitenancy, data protection with instantaneous snapshots, clones, remote snapshots to secondary object storage for disaster recovery, backup, cross-region replication, and cloud-bursting from on-premises. The unique snap-to-object feature allows users to easily create a replica of production data and instantly push it to a second object store—enabling snapshot-based replication. The second cloud copy can facilitate workload migration to another application cluster or provide a fast recovery point objective (RPO) service guarantee.

WEKA on GCP delivers outstanding economics, return-on-investment, and application consolidation. Users get the most out of their GPU, TPU, and CPU resources by saturating them with data and maximizing their efficiency. WEKA on GCP also helps to optimize compute resource usage by growing or shrinking the global namespace, all without downtime. Users can add or reduce compute and storage capacity independently as needed, keeping hot data on NVMe and warm/cooler data on S3 Object Storage.

The WEKA Data Platform addresses all the challenges of running modern workloads in the cloud by delivering scalability, flexibility, and performance along with ease-of-use and affordable enterprise features on GCP.

For More Information or to Arrange a Free Trial

Visit us at <https://www.weka.io/get-started> or email us at info@weka.io.



[weka.io](https://www.weka.io)

844.392.0665

