

# Move Any Workload To Any Cloud, Even The “Impossible” Ones

## Challenges

- Mission-critical applications often experience performance challenges after migrating to the cloud.
- Low-latency or IO-intensive applications or workloads that share large amounts of data have been considered impossible to run affordably in the cloud.

## Solution

- WEKA Data Platform provides a single, unified product with consistent high performance and robust data services regardless of where your data and applications are located.

## Benefits

- Deliver a cloud experience that feels local.
- Quickly leverage on-demand compute without a significant CAPEX investment.
- Minimize data copies and move data more easily.
- Get unlimited scaling at the best cost with automatic tiering.
- Eliminate the need to continually optimize systems.

## WEKA Delivers the Multicloud Data Platform for the Future of Federal Workloads

Federal agencies need flexibility in the way their applications are deployed and consumed. Depending on the mission and program goals, agencies need to deploy mission-critical applications in public and private clouds as well as on-premises—and, enabled by the rise of containers, often in multiple locations simultaneously. According to IDC, public cloud storage spending surpassed on-premises storage spending in 2019 and grows as the cloud becomes the primary deployment option to meet business goals.

Most agencies have at least one cloud initiative that leverages the economy, scale, or elastic nature of the cloud. The cloud's increased speed and agility are particularly attractive with today's supply chain disruptions. Missions and programs cannot afford to have their critical business and research initiatives hindered by snarled supply chains, and the cloud has become the de facto solution.

### The Challenge: Large, Resource-Intensive Workloads in the Cloud

There are many well-known benefits of moving workloads to the cloud in terms of scalability, flexibility, security, and elasticity. What is not talked about much is that agencies have struggled with moving many applications to the cloud. Running large or resource-intensive workloads in the cloud is extremely challenging, particularly at the scale of modern deployments.

Architecting workloads in the cloud requires different approaches than on-premises. Cloud infrastructure works on loosely-connected systems, and workloads designed for traditional, dedicated infrastructure usually can consume excessive resources or fail to meet performance requirements when run in the cloud. A common workaround is to run applications locally on a single instance, which prevents you from taking advantage of many of the benefits of the cloud.

Low-latency or IO-intensive applications that share large amounts of data have often been considered impossible to run affordably in the cloud and have either remained on-premises or been painstakingly brought back after failed cloud deployments.

## The Promise of Multicloud and Cloud Openness

According to both 451 Research and Gartner, 76% of companies are using two or more public clouds, with the average having 2.3 clouds in use. Organizations with more than a billion dollars in revenue are twice as likely to be using three or more clouds than smaller businesses.

You want the freedom of choice and access to vendor-specific capabilities of different clouds. However, the promise of openness and transparency of the cloud has been stymied by proprietary lock-in of data sets and prohibitive economics associated with data egress. The inability to easily move data from cloud to cloud or share data across multiple clouds in an economically viable way has slowed down missions and programs from reaching their goals to adopt a multicloud strategy.

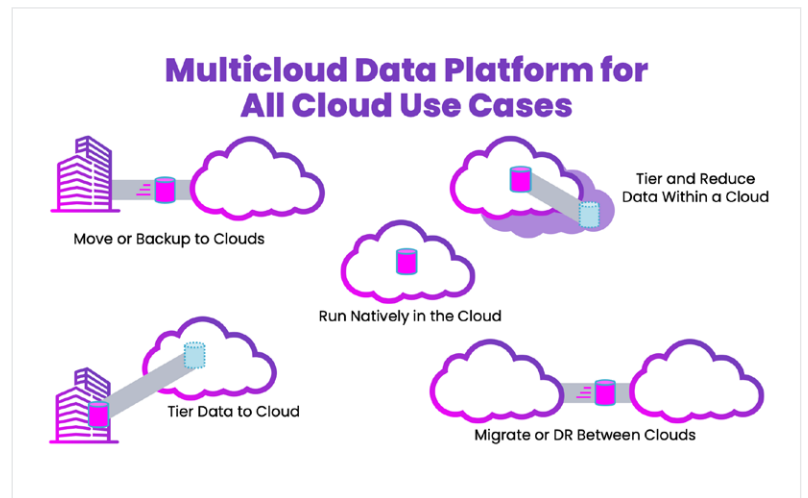
## The Solution: The Industry's First Multicloud Data Platform

The WEKA Data Platform is the industry's first multicloud data platform for AI and next-generation workloads. It provides a single, unified product with consistent high performance and robust data services that deliver a seamless, simplified data management experience with best-in-class economics—regardless of where your data and applications are located.

WEKA provides seamless integration with the compute and object storage resources of Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP), and Oracle Cloud Infrastructure (OCI) to give agencies a single, unified deployment option that performs consistently across the edge, core, hybrid, and multicloud environments—regardless of your preferred cloud vendor or location.

While many vendors offer some form of cloud integration or support, none of them provides a single unified platform that supports all the major cloud use cases:

- Running Natively on the Cloud
- Tiering and Reducing Data within a Cloud
- Moving or Backing Up Data from On-Premises to Clouds
- Use the Cloud for Data Tiering
- Migrating or Replicating Data Between Clouds or Zones



And, unlike other offerings, the WEKA Data Platform implements the exact same code base across all deployments, standardizing management and capabilities across all clouds.

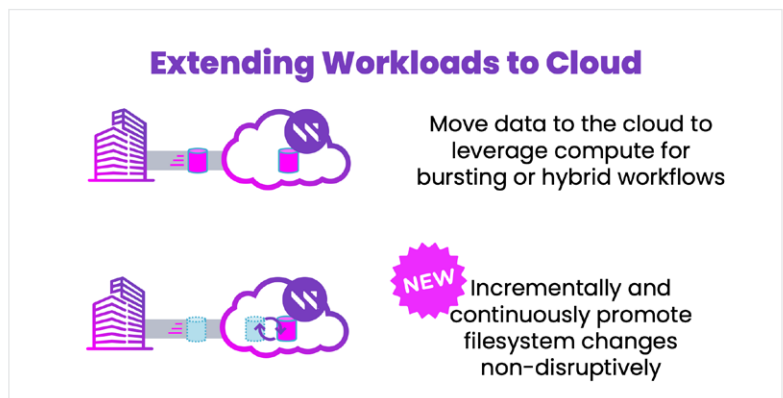
## Best Economics and Performance for Data-Intensive Applications

WEKA seamlessly combines the performance tier of NVMe-enabled cloud compute instances for the lowest latency with the capacity tier of native object storage for massive scalability, all presented in a single namespace for best economics at petascale. Integrated data management automatically and seamlessly 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, it is usually more than twice as cost-effective as native storage services at comparable performance levels.

The WEKA Data Platform is an ideal solution for missions and programs with applications that have varied IO patterns with unpredictable performance needs. It excels at both latency-sensitive and large file IO without tuning. WEKA has broad multi-protocol support, allowing multiple applications (Windows, Linux, Mac, POSIX, S3 native) to simultaneously share and collaborate on a single data set with full data consistency across protocols.

## Make Hybrid Cloud Work for You

WEKA makes it easy to have an entirely usable copy of any quantity of data in the cloud, so missions and programs can start analyzing that data without building out new infrastructure. Use elastic compute resources to get new insights by running calculations and analyses in the cloud on data shared from on-premises. And then release—and stop paying for—the compute resources when analysis is complete.



Or build sophisticated hybrid workflows where you start processing the data on-premises but then use WEKA to move it to the cloud for final analysis—or vice versa. And you can incrementally and continuously promote changes to the cloud—updating the file system non-disruptively. This minimizes the bandwidth needed and keeps the hybrid workflows running non-stop.

## Mission and Program Benefits

Whether on-premises, in between, or headed to an all-cloud approach, WEKA helps make your cloud journey seamless and painless. WEKA makes the transition to cloud-based workflows easier. It enables you to get easy access to extra compute resources in the cloud without shifting large quantities of data. WEKA allows you to leverage the best of what the cloud offers without compromising your data or security requirements. And if you are just starting on your cloud journey, WEKA takes the pressure off, allowing you to move at a measured pace without the stress of having to make all-or-nothing decisions.

- Deliver a cloud experience with the ultra-low-latency and high bandwidth of a local server for all applications.
- Easily leverage the latest high-performing, on-demand compute infrastructure as a service without the need for massive CAPEX investment.
- Eliminate the need to continually optimize systems for each new workload—zero performance tuning.
- Enable users or applications to access the same data set simultaneously—even from on-premises—to reduce copies, movement, and complexity of data as well as the number of people required to manage the environment.
- Get unlimited scaling at the best cost by automatically tiering between the NVMe native compute and object storage resources of the clouds, all in a single namespace.
- Easily manage and automate your environments, no matter the scale, with full GUI, CLI, and REST API management.

## To Find Out More or Arrange for a Free Trial

- **Contact us at:** [US.Federal@weka.io](mailto:US.Federal@weka.io)
- **Find us online:** [www.weka.io/get-started](http://www.weka.io/get-started)



[weka.io](http://weka.io)

844.392.0665

