

WEKA On Microsoft Azure

The Most Scalable, Highest Performing File System for Hybrid Cloud

About Us

WEKA is accelerating innovation, research and discovery for the world's leading enterprises and research organizations with the industry's first multicloud data platform for AI, machine learning, and other next-generation workloads.

Product

The WEKA® Data Platform delivers the radical simplicity, epic performance, and infinite scale required to support enterprise AI, machine learning, and other next-generation workloads in virtually any location.

Customers

WEKA is accelerating innovation, research, and discovery for the world's leading enterprises and research organizations—including eight of the Fortune 50—and helping them to achieve first-to-market results.

Organizations of all types are turning to Microsoft Azure to make their hybrid cloud strategy a reality. Next-generation workloads that rely on accelerated compute and machine learning are introducing new challenges for customers as they bring them to the cloud.

These performance-intensive workloads add new pressure on infrastructure across both data center and cloud environments. Organizations with data center investments are trying to augment their on-premises applications by using the cloud to increase scalability and improve their economics. As customers increasingly add cloud infrastructure into the mix, their users demand the same performance and user experience they were used to with the on-premises deployment. Across an increasingly diverse set of industries from media and entertainment to financial services to life sciences to manufacturing, technical teams are being pressed to make smart decisions about the right infrastructure that can bridge from their on-premises deployment to the cloud. The WEKA™ Data Platform® provides the fastest, most scalable file system for Microsoft Azure, delivering the performance users expect and the scalability and simplicity the cloud promises - for every workload, even the most demanding ones.

Why Choose WEKA on Microsoft Azure

Run your most performance-intensive applications 7x faster in Microsoft Azure

Deliver unbeatable file storage performance for your most demanding applications running in Azure, supporting high I/O, low latency, small files, and mixed workloads with zero tuning with automatic storage rebalancing.

Dynamically Scale Up and Down to Save Costs and Time

Desegregate storage and compute and scale each independently. Scale up to meet peak demands and scale down when no longer needed. Eliminate painful data resynching and resharing to make upgrades and maintenance simpler and free staff to work on more valuable projects.

Accelerate Hybrid Cloud Workflows running on Microsoft Azure

Maintain a usable copy of your on-premises data in your Azure environment, where you can use elastic compute resources to run calculations and analyses and gain new insights. Use easy, low-impact data portability to enable back-up and disaster recovery to the cloud, cloud bursting, and cloud retention and archive.

Improve time to results by up to 90% on Microsoft Azure

Your researchers, data scientists, creative teams, or other knowledge workers can analyze data faster, and collaborate seamlessly using a single copy of data optimized to meet the performance needs of every step in your workflow, and accessible from any client (Windows, Linux, Containers, or POSIX-compliant applications).

Migrate more applications to Microsoft Azure

Send snapshots of your WEKA filesystem to any Azure Blob Storage object store for backup and disaster recovery. Use your disaster recovery copy in Azure for either recovery or cloud bursting. Full and incremental snapshots include metadata to enable seamless data portability between on-prem and Azure.

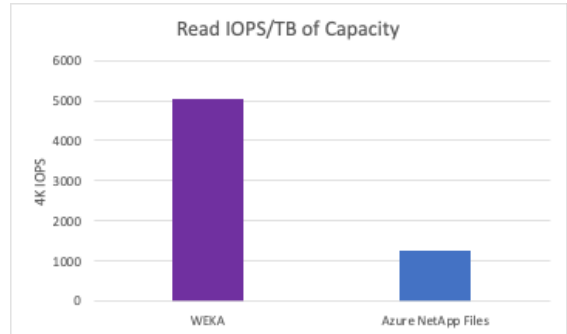
Optimize utilization of accelerated compute workloads and reduce infrastructure costs by up to 65%

WEKA enables you to run your compute-intensive workloads, including HPC and ML applications, faster and more efficiently by ensuring they are never waiting on the storage infrastructure for data.

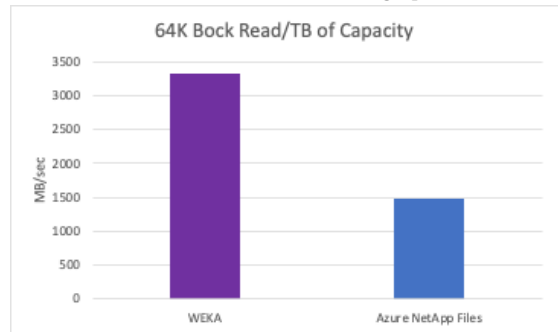
WEKA on Azure Outperforms the Alternatives

*Versus Azure Netapp Files

4x Faster IOPs



2.25X Greater Throughput



WEKA on Microsoft Azure Use Cases

Run high-performance applications on Azure

Deliver performance for your most demanding applications running in Azure with the world's fastest cloud-native data platform supporting high I/O, low latency, small files, and mixed workloads with zero tuning.

Seamlessly tier data in Azure

Intelligent tiering automatically moves data between high-performance flash-based storage on Azure Virtual Machines and low-cost, massively scalable object storage in Azure Blob Storage, all in a single namespace for the best performance, scale, and economics.

Migrate data to Azure

Send snapshots of your on-premises filesystem data and metadata to any Azure Blob Storage object store for backup and disaster recovery. Full and incremental snapshots include metadata to enable seamless data portability between your premises and Azure.

Burst data analysis to Azure

Maintain a usable copy of your on-premises data in your Azure environment, where you can use elastic compute resources to run on-demand calculations and analyses and gain new insights.

Build agile data pipelines in Azure

Your researchers, data scientists, creative teams, and more can collaborate faster by using a single copy of data optimized to meet the performance needs of every step in your workflow. Customers who use WEKA from their data pipelines deliver up to 90% faster time to insights from their data.

Sample Instance Configurations for WEKAFS

	Mixed I/O Pattern	Highest Performance
Virtual Machine	Standard_L8s_v3	Standard_L16s_v3
Cluster Size (nodes)	6 to Hundreds	6 to Hundreds
Raw Capacity	11.52 TB	23 TB
Usable Capacity*	5.9 TB	11.7 TB
Blob Storage	Yes	Yes
Networking	12.5 Gb/second	12.5 Gb/second
Software	WEKA(R) Data Platform WekaFS™ File System	
Data Protection	Distributed Data Protection (N+2 to N+\$), Encryption, End-to-End Data protection	
Protocols	POSIX, NFS, SMB, S3	
Max File Systems	Up to 1024 per Cluster	
Snapshots	File System Level, Up to 24,576 Readable Snapshots (14,336 if all writable)	
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	
Performance	Up to 400 MBps / TB	

*Note: Clusters use 2 node parity, regardless of size so the net capacity will scale to >98% of raw capacity as the cluster size count grows.

For More Information

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

To speak to a representative or get more information [contact us](#)



weka.io

844.392.0665

