

Unlocking Cost Efficiency

With the WEKA® Data Platform.

Executive Summary

In an era of rapid digital transformation, organizations face mounting challenges to optimize their data infrastructures amidst unprecedented growth in data volumes and complexity. Balancing the need for high performance, scalability, and cost efficiency has become a defining priority for business executives. The WEKA Data Platform offers a groundbreaking solution, leveraging cutting-edge technology to drive transformative cost reductions while maintaining unmatched performance for AI, HPC, and data-intensive workloads. This document explores how WEKA's unique capabilities deliver value across financial, operational, and strategic dimensions.

Key Cost-Saving Capabilities of WEKA

1 Autoscaling for Elastic Cloud Storage and Compute

Autoscaling is a cornerstone of cloud economics, and WEKA extends this principle to the storage layer, a feat few systems can match. Unlike legacy storage systems that require pre-provisioned capacity—often leading to overpayment for unused resources—WEKA dynamically adjusts both storage and compute capacity to meet real-time workload demands.

- **How it Works:** WEKA leverages cloud-native capabilities, such as AWS Auto Scaling and Azure VM Scale Sets, to dynamically expand or contract the storage environment. When workloads peak, additional resources are spun up to maintain performance. When demand drops, resources are scaled back down automatically, ensuring that businesses only pay for what they use.
- **Cost Impact:** This capability eliminates the need for over-provisioning, which can save organizations millions annually in operational and capital expenses. Traditional storage systems often tie performance to fixed capacity, leading to resource wastage. In contrast, WEKA's approach achieves optimal resource utilization.
- **Outcome:** Up to 95% reduction in infrastructure costs during low-demand periods.
- **Example:** Stability AI, a pioneer in generative AI, used WEKA's autoscaling to achieve 15x more storage capacity while cutting costs by 20% compared to previous solutions. Their infrastructure scaled seamlessly to handle rapid data.

2 Intelligent Data Tiering for Lower TCO

Intelligent data tiering is critical for organizations managing large datasets, where only a fraction of the data requires high-performance storage. WEKA addresses this challenge with an innovative tiering mechanism that integrates high-speed NVMe flash and cost-efficient object storage into a single namespace.

- **How it Works:** WEKA automatically identifies “hot” data that needs fast access and keeps it on NVMe flash storage. Less frequently accessed (“cold”) data is seamlessly moved to cost-efficient object storage such as AWS S3 or Azure Blob. This process occurs without manual intervention, reducing operational complexity and ensuring optimal placement of data across performance and cost tiers.
- **Strategic Advantage:** Traditional systems rely on manual policies or external software for data migration, often leading to delays and increased costs. WEKA's built-in tiering eliminates the need for separate software and ensures continuous, real-time optimization of data placement.
- **Outcome:** Customers report up to 75% reduction in storage costs while maintaining high application performance.
- **Example:** A genomics firm utilized WEKA's data tiering to store massive genomic datasets efficiently, cutting per-genome costs by over 70% and accelerating research cycles.

3 Simplified Management and Reduced Operational Overhead

The cost of managing traditional storage systems extends beyond hardware and includes significant personnel expenses. These systems require specialized teams to handle tasks like provisioning, performance tuning, and troubleshooting. WEKA's streamlined architecture minimizes these burdens, allowing a single administrator to manage even the largest environments.

- **How it Works:** WEKA consolidates disparate storage silos into a unified system with intuitive management tools. Automation of routine tasks, such as scaling, data placement, and system health monitoring, significantly reduces the need for manual intervention.
- **Strategic Advantage:** By reducing management complexity, organizations can redirect IT talent toward strategic initiatives that drive innovation and revenue growth.
- **Outcome:** Up to 64% reduction in data operations staff, equating to millions in annual savings for large enterprises.
- **Example:** A global media company using WEKA saw staffing requirements drop from almost 20 staff down to just 1, freeing IT teams to focus on developing data-driven customer experiences instead of maintaining infrastructure.

4 Optimized Infrastructure Utilization

For enterprises running performance-intensive workloads, underutilization of high-value compute resources, such as GPUs, can significantly inflate costs. WEKA maximizes infrastructure efficiency by ensuring seamless, high-speed data delivery to compute nodes, preventing bottlenecks that lead to idle time.

- **How it Works:** WEKA's zero-copy and zero-tuning architecture ensures that data is always delivered at the speed required by compute-intensive applications. Its distributed metadata management system eliminates single points of failure, while its integration with GPU Direct Storage maximizes data throughput.
- **Cost Impact:** By improving data access efficiency, WEKA enables organizations to achieve higher utilization rates for expensive compute resources, such as GPUs. This reduces the need for additional hardware, driving down capital expenditures.
- **Outcome:** Customers have reported a 93% GPU utilization rate, compared to industry averages below 60%, and a 35% reduction in training times for AI models.
- **Example:** Stability AI used WEKA to increase GPU utilization efficiency, reducing training cycles by three weeks and accelerating their product roadmap.

5 Reduced Hardware and Energy Costs

Physical infrastructure costs, including servers, networking equipment, and energy consumption, represent a significant portion of IT budgets. WEKA's high-performance density and software-defined architecture reduce these costs while supporting sustainability goals.

- **How it Works:** WEKA consolidates workloads onto fewer nodes by delivering unmatched performance per server. Its architecture minimizes the need for additional hardware, reducing data center footprint, cabling, and energy consumption.
- **Environmental Impact:** WEKA supports sustainability initiatives by lowering power usage and cooling requirements. Its efficient scaling ensures that businesses meet performance goals without increasing their carbon footprint.
- **Outcome:** Customers have reduced physical infrastructure requirements by up to 20x and power consumption by 50%, significantly cutting operational expenses.
- **Example:** A financial services company using WEKA saw dramatic reductions in their energy bills and data center complexity, enabling them to reinvest savings into customer-centric innovations.

1

Optimized Infrastructure Utilization

For enterprises running performance-intensive workloads, underutilization of high-value compute resources, such as GPUs, can significantly inflate costs. WEKA maximizes infrastructure efficiency by ensuring seamless, high-speed data delivery to compute nodes, preventing bottlenecks that lead to idle time.

- **How it Works:** WEKA's zero-copy and zero-tuning architecture ensures that data is always delivered at the speed required by compute-intensive applications. Its distributed metadata management system eliminates single points of failure, while its integration with GPU Direct Storage maximizes data throughput.
- **Cost Impact:** By improving data access efficiency, WEKA enables organizations to achieve higher utilization rates for expensive compute resources, such as GPUs. This reduces the need for additional hardware, driving down capital expenditures.
- **Outcome:** Customers have reported a 93% GPU utilization rate, compared to industry averages below 60%, and a 35% reduction in training times for AI models.
- **Example:** Stability AI used WEKA to increase GPU utilization efficiency, reducing training cycles by three weeks and accelerating their product roadmap.

2

Reduced Hardware and Energy Costs

Physical infrastructure costs, including servers, networking equipment, and energy consumption, represent a significant portion of IT budgets. WEKA's high-performance density and software-defined architecture reduce these costs while supporting sustainability goals.

- **How it Works:** WEKA consolidates workloads onto fewer nodes by delivering unmatched performance per server. Its architecture minimizes the need for additional hardware, reducing data center footprint, cabling, and energy consumption.
- **Environmental Impact:** WEKA supports sustainability initiatives by lowering power usage and cooling requirements. Its efficient scaling ensures that businesses meet performance goals without increasing their carbon footprint.
- **Outcome:** Customers have reduced physical infrastructure requirements by up to 20x and power consumption by 50%, significantly cutting operational expenses.
- **Example:** A financial services company using WEKA saw dramatic reductions in their energy bills and data center complexity, enabling them to reinvest savings into customer-centric innovations.

How WEKA Enhances Your Bottom Line

Enabling Cloud Migration Without Cost Barriers

Migrating data-intensive workloads to the cloud is often hindered by concerns about cost and performance. WEKA bridges this gap with its ability to handle even the most demanding workloads efficiently in cloud environments, enabling cost-effective migration.

Supporting Modern Workloads

Whether it's AI/ML, genomic research, or real-time analytics, WEKA eliminates bottlenecks, accelerates insights, and reduces operational costs. Its unified platform ensures that all applications—regardless of scale—run seamlessly.

Future-Proofing Your Infrastructure

WEKA's software-defined design supports both cloud-native and hybrid deployments. Its scalability ensures that enterprises can adapt to future needs without disruptive infrastructure upgrades.

Quantifiable Benefits of WEKA's Cost Efficiency

35%

reduction in
cloud storage costs

64%

reduction in
operational staff

95%

reduction in AI/ML Model
Training Costs

20%

reduction in
Infrastructure Footprint

Why Business Leaders Choose WEKA

WEKA empowers enterprises to achieve their strategic goals by delivering unmatched performance and efficiency. Its ability to reduce costs, improve resource utilization, and accelerate time-to-value makes it a critical enabler of digital transformation. By choosing WEKA, organizations can unlock the full potential of their data while achieving sustainable growth.



weka.io

844.392.0665

