



Five Reasons You Need a Data Platform to Solve Your Complex HPC Storage Problems

High-Performance Computing (HPC)

relies on a critical “infrastructure triangle” consisting of network, compute, and storage to supply and process essential data for addressing the world’s most complex problems. Yet, over time, networks and compute have experienced exponential performance improvement with 2x-5x CPU/GPU density, up to 8x network capacity, and 3x increases in memory. However,

storage has lagged, only making limited iterative improvements, and so the storage environments, requiring constant care and performance tuning, have become a nightmare to manage at scale. What further exposes these limitations is the adoption of GPU-powered server farms, which is quickly becoming the new supercomputer of today and consolidating traditional HPC and AI workloads, creating unpredictable

I/O patterns and mixed data sizes. The fact that storage has not kept up with the advances on the other legs of the “triangle” results in the underutilization of resources, increased costs to deliver performance, and erratic time to results.

To try and overcome these complexities, All-Flash-NAS appliances were developed to ease management and improve performance, but it quickly showed that while adding NVMe flash does improve some performance of legacy protocols like NFS, it does not keep up with the needs of performance-dense GPU-powered server farms. As a result, while easy to deploy, these NAS solutions have fallen short of expectations as they were not designed for the performance and scale required by modern workloads.

To truly keep up with the significant performance improvements made on the compute and networking fronts, a unilaterally different design is needed to deliver dynamic and independent performance and capacity scaling.

In 2013, WEKA’s founders set out with a blank sheet of paper and a vision to create a fundamentally different approach to storing, managing, and moving data that could eradicate the compromises of the past AND power the possibilities of the future. The WEKA Data Platform is purpose-built for large-scale data processing and high performance. It delivers the best large and small file performance at ultra-low latencies while not compromising enterprise features that ensure data security and availability.

01 **Get Faster Research Outcomes for Next-Generation Workloads**

The WEKA Data Platform has repeatedly proven to hundreds of customers the massive performance delta, from 10-100x, it has over similar competitive solutions. Customers

have measured major improvements in application performance, resulting in faster insights, quicker time to market, and huge improvements in infrastructure efficiency.

02 Zero Tuning Required Based on Workload

WEKA is ideally suited for the challenges of mixed workloads—large and small files, random and sequential access, and structured and unstructured data. The software

has been engineered to require zero tuning for one workload or another. All workloads will improve, and IT administrators can focus on managing the application workloads instead.

03 Freedom of Choice: Have a Preferred Hardware Vendor?

There is no need to sign on new vendor relationships. Keep buying from your preferred source. WEKA has been rigorously tested on multiple server platforms from the top server vendors, and the software can be purchased as an integrated solution. WEKA actively tests interoperability with

the processor, SSD, networking, and object storage solutions to ensure the best performance and compatibility. In addition, WEKA is thoroughly tested with CPU and GPU-based servers and applications for continuous performance improvements.

04 No PhD or Ponytails Needed to Run a High-Performance Environment

WEKA eliminates the complexity associated with traditional parallel filesystems. It is available fully pre-configured and tested from some of the most respected server manufacturers as a turnkey reference architecture file system component,

making deployment as easy as any appliance. Once deployed, the WEKA Data Platform can manage exabytes of data from the GUI or integrate with widely adopted automation, monitoring, and management tools.

05 **Weka: Simple and Seamless**

- ✓ WEKA can scale up or down easily. Scale up to exabytes of cloud storage while eliminating infrastructure complexity by extending the file system namespace to include any S3-compatible object store (public or private). Scale down when priorities change without affecting the entire cluster.
- ✓ Access application data anywhere you want: in the cloud, on-premises, or hybrid
- ✓ WEKA's global namespace seamlessly combines the performance tier of NVMe for the lowest latency with the capacity tier of S3 object storage for massive scalability for the best economics at Petascale.

Looking for breakthrough performance, elasticity, scalability, and flexibility of your data for your applications?

Contact a WEKA Solutions Consultant for a complementary review today:

- Contact Us: www.weka.io/contact-us
- Learn More on Our Website: www.weka.io/cloud



weka.io

| 844.392.0665

