

Fueling Growth and Discovery with the Speed of Flash and the Efficiency of Private Cloud **WEKA.IO**

Highlights

- **Breakthrough Economics:** Weka's hardware independent storage software coupled with Western Digital's object storage system consolidates multiple tiers of storage, saving power, cooling, and rack space
- **Unparalleled Simplicity:** Rack space, power and network connections, and your existing servers with SSDs are all you need
- **Easy Management:** Policy-based tiering automatically migrates data to cost-optimized object storage, all managed from a single screen so one admin can easily manage petabytes of data
- **Unmatched agility and scalability:** Capacity scales to petabytes within a single namespace, and performance scales dynamically and independently, allowing the system to adapt to varying workloads

Customer Challenges

- **Diverse workloads:** Today's workloads are a complex mix of small and large files, random and sequential access, structured and unstructured data, and require both high throughput and low latency
- **Complex and difficult to manage storage architectures:** Legacy storage systems were not designed for cloud scale and result in silos of data that are difficult and costly to manage
- **Growing data sets and longer retention periods:** Corporations and research institutions generate enormous amounts of data that must remain available for longer periods

Solution

- WekaIO's Virtual File Server™ software converts your existing server clusters into a massively parallel scale-out storage system with linear performance scalability and sub-millisecond latency
- ActiveScale object storage system is a complete, easy-to-install, easy-to-manage active archive solution delivering cloud economics to enterprises

Modern Applications Pose Unique Storage Challenges

In today's rapidly changing world, corporations and research institutions alike are creating data at an unprecedented rate. These data sets are digital assets that hold great value as a way for corporations to gain strategic competitive advantage and for researchers to make important new scientific discoveries that improve our quality of life.

Modern scientific research applications frequently generate terabyte-sized files while real-time IoT and trading system applications stream millions of tiny files. The sheer volume and variety of data pose significant challenges because traditional storage systems are based on legacy technology that was never designed for the performance levels and scale found in today's data centers. As a result, storage systems have become a major bottleneck to progress. CIOs are increasingly looking to advanced technology such as cloud and hyper converged systems to deliver the business agility, scale and manageability required to regain lost momentum and optimize their IT investment.

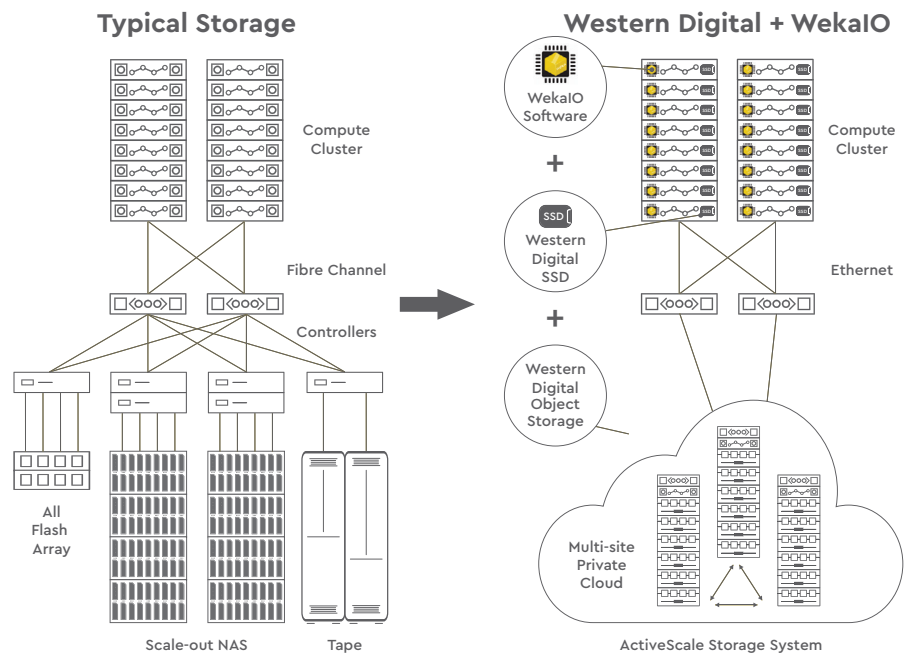


Figure 1: The Western Digital+WekaIO solution combines the performance of flash with the durability and efficiency of object storage.

Limitless Performance and Capacity Scalability

Business needs are constantly changing, and your storage infrastructure should adapt and scale right along with them. WekaIO's Virtual File Server (VFS) software paired with ActiveScale™ object storage system deliver unparalleled performance, scale, and affordability.

VFS has been architected to allow performance and capacity to scale automatically and independently. Performance scales based on the number of CPU cores allocated to VFS; cores can be added and removed dynamically to handle peak workload needs such as month-end processing or a complex research project.

At the storage layer, VFS provides a scalable tier of solid-state disk (SSD) storage to support performance sensitive workloads while the ActiveScale object storage system provides highly durable and cost-effective capacity. A fully configured ActiveScale system provides one of the highest density, lowest power footprints in the industry, allowing you to optimize expensive data center floor space and dramatically lower power and cooling expenses. The system architecture scales from hundreds of terabytes to exabytes.

Solution details

Parallel File Access Accelerates Your Business

Performance starts with the file system. VFS is a distributed, scale-out, POSIX-compliant file system that runs on your existing compute cluster and uses off-the-shelf SSDs, eliminating vendor lock-in. Because the data resides on flash-based storage inside the server and part of the single global namespace, data access is near instantaneous. Valuable data is protected with a combination of erasure coding and our patented data distribution algorithm, and can sustain up to four simultaneous node or SSD failures without losing access to your data.

Data Durability and Integrity at Scale Safeguards Digital Assets

ActiveScale's system software ensures that valuable data is protected and always available with up to 19 nines durability, including site-level fault tolerance in a multi-site configuration. Robust data integrity checks occur automatically and transparently protecting your data from silent data corruption known as bit-rot. Each object can tolerate 1000 bit-errors without the loss of data, which can eliminate the risks, costs, and media management activities associated with tape-based archives.

Easy Deployment and Data Management Lower TCO

The transformative simplicity of the combined solution compared to traditional storage architectures is shown in the figure. Multiple tiers of storage are consolidated into a performance tier for hot data and a capacity tier for cold data. Data is automatically migrated to cost-effective ActiveScale object storage for long-term retention either on-demand or automatically based on policies.

The entire solution can be up and running in minutes. Simply run the automated VFS installation procedure on the selected servers, and VFS appears as a single NFS mount point to the Linux cluster. Roll the ActiveScale system in place, connect the power, configure the network connections, and it presents an Amazon S3™ compliant object interface and global namespace to VFS. The system self protects and heals allowing a single storage admin to easily manage multiple petabytes of data. With ActiveScale Cloud Manager, system analytics can offer insights for predictive and preemptive support and issue resolution.

Conclusion

Designed for the most demanding data intensive technical workloads, the combined solution from WekaIO and Western Digital can dramatically change the economics of high performance storage. Leveraging off-the-shelf servers, SSDs, virtualization and object storage, the solution delivers the performance of an all-flash-array with the durability and efficiency of on-premises cloud object storage.

To learn more, visit www.Weka.IO and www.wdc.com/dc-systems

Western Digital

5601 Great Oaks Parkway
San Jose, CA 95119, USA
US (Toll-Free): 800.801.4618
International: 408.717.6000

www.westerndigital.com

© 2017-2018 Western Digital Corporation or its affiliates. Produced 03/17, Rev. 06/18. All rights reserved. Western Digital, the Western Digital logo and ActiveScale are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the U.S. and/or other countries. WekaIO and WekaIO Virtual File Server are trademarks or registered trademarks of WekaIO, Inc. or its affiliates in the U.S. and other countries. Amazon S3 is a trademark of Amazon.com, Inc. or its affiliates. All other marks are the property of their respective owners.