SOLUTION BRIEF

# SUPERMICRO® SERVERS EXCEL AS A WEKAIO REFERENCE ARCHITECTURE

*The Next Generation Storage Solution with 3rd Gen AMD EPYC™ 7003 Processors*

TABLE OF CONTENTS



## Executive Summary


Unstructured data is a customer's most valuable asset, which must be preserved and protected forever. However, relentless data growth and retention trends drive demands for more efficient, resilient, and secure Exabyte-scale storage solutions. These demands continue to pressure IT budgets and administrators. Simultaneously, organizations are looking to unlock the value in their data, which makes the task even more challenging. The correct storage architecture can allow organizations to leverage more of their data and make facilitating 'data forever' a realistic prospect. Jointly with our market-leading, strategic partners, Supermicro can easily transition from legacy storage to next-generation storage platforms.

WekaIO™was founded on the idea that current storage solutions have forced IT organizations to choose complex solutions to address their highest storage need at the expense of other desirable capabilities. The three dominant architectures are block, file, and object, each servicing a different need: speed, shareability, and scalability in that order. In today's "data-as-a-service" market, organizations need a flexible infrastructure that addresses the many business needs within a single framework. The design philosophy behind the Weka file system – WekaFS - was to create a single storage architecture that runs on-premises or in the public cloud with the performance of all-flash arrays, the simplicity and feature set of network-attached storage (NAS), and the scalability and economics of object storage.

## WekaIO Solution Architecture Overview

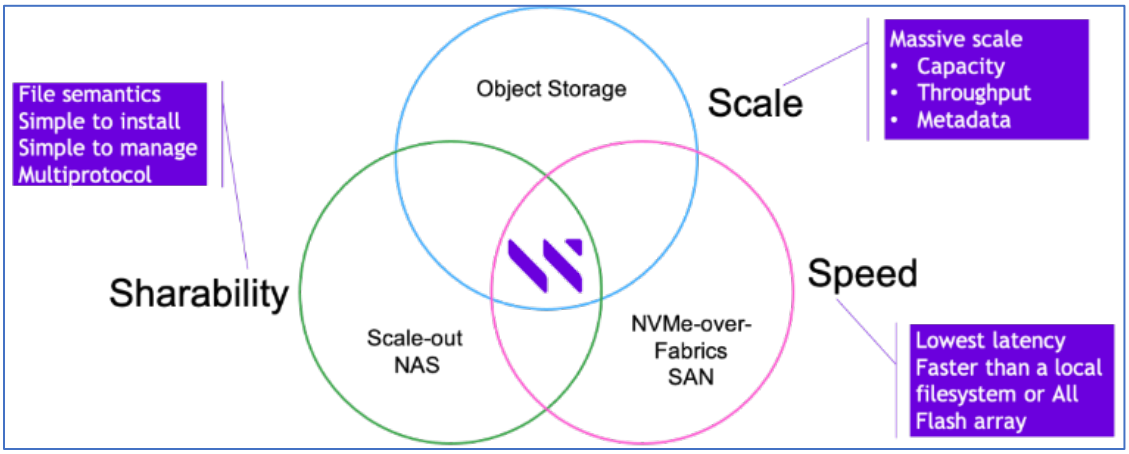


*Figure 1 - Weka File System Structure*

WekaIO's file system (WekaFS) is a fully distributed, parallel file system that was written entirely from scratch to deliver the highest performance file and object services by leveraging NVMe flash as its primary storage for persistent data across a wide range of applications. WekaFS will also, and transparent to the application layer, seamlessly expand the filesystem namespace to include an extended layer built on any S3 compliant object storage system. There is no need for a particular data migration software or complex scripts; all data resides in a single global namespace for easy access and management while maintaining the best performance. The intuitive graphical user interface allows a single administrator to quickly and easily manage hundreds of petabytes of data without any specialized storage training.

Leveraging existing technologies in new ways and augmenting them with engineering innovations, Weka's software delivers a more powerful and more straightforward solution that would have traditionally required several disparate storage systems.

### SUPERMICRO

Supermicro (Nasdaq: SMCI), the leading innovator in high-performance, high-efficiency server and storage technology is a premier provider of advanced server Building Block Solutions® for Enterprise Data Center, Cloud Computing, Artificial Intelligence, and Edge Computing Systems worldwide. Supermicro is committed to protecting the environment through its "We Keep IT Green®" initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.

The resulting software solution provides high performance for all workloads (big and small files, reads and writes, random, sequential, and metadata heavy). Furthermore, it is designed to run on a server infrastructure that does not rely on any specialized hardware-assist. As future hardware innovations come to market, WekaFS is well-positioned to leverage emerging technologies for the continued delivery of best cost and performance. The system can be expanded online to handle more demanding performance or store more capacity with no service interruption.

## Configuration

- Our WekaIO reference configuration offers the lowest cost, most flexible, outstanding performance on a single processor platform leveraging AMD EPYC processors with 128 lanes of PCIe Gen4 that supports up to 20 PCIe Gen 4 NVME storage devices, greater than 256 GB of memory, with 24 cores per CPU.
- The reference system provides raw capacity of up to 6.5 PB/rack or 306 TB/server using 15.3 TB NVME PCIe Gen 4 (21x20x15.3=6.5 PB per 42U Rack)
- High-performance PCIe4 Kioxia NVME drives with measured performance of 6.9 GB/s vs. 3.3 GB/s and 1.6 MIOPs vs. 800 KIOPs for PCIe 3 drives.
- One of the key advantages of this platform is the online upgrade of additional capacity without adding IO/switches/network port/or nodes by inserting NVMe drives to empty slots within the cluster systems. This allows users to simply install additional drives for extra capacity, turn on the software, and the systems are being configured automatically.Users can also easily scale out the number of storage cluster hoststo improve performance and capacity as needed.

An example of a validated cluster solution using AS -2114S-WN24RT single-node servers:

| Type | Description | Per System | 6 System config |
|---|---|---|---|
| **System** | AS -2114S-WN24RT AMD WIO 2U/24 NVMe A+ Rack Server | 1 | 6 |
| **CPU** | 3rd Gen AMD EPYC 74F3 24C/48T 3.2G 256M 240W | 1 | 6 |
| **Memory** | 32GB DDR4-3200 2Rx4 ECC Registered DIMM | 8 | 48 |
| **Boot Drive** | Micron 2300 2TB NVMe PCIe Gen4 M.2 SSD | 1 | 6 |
| **Storage Drive** | Kioxia CM6 7.68TB NVMe PCIe Gen4 2.5" U.2 SSD | 20 | 120 |
| **NIC** | Mellanox ConnectX-6, Low-Profile Dual-port VPI HDR 200GbE, QSFP56, PCIe Gen4 | 2 | 12 |

*Table 1 - 6 System Config Specifics*

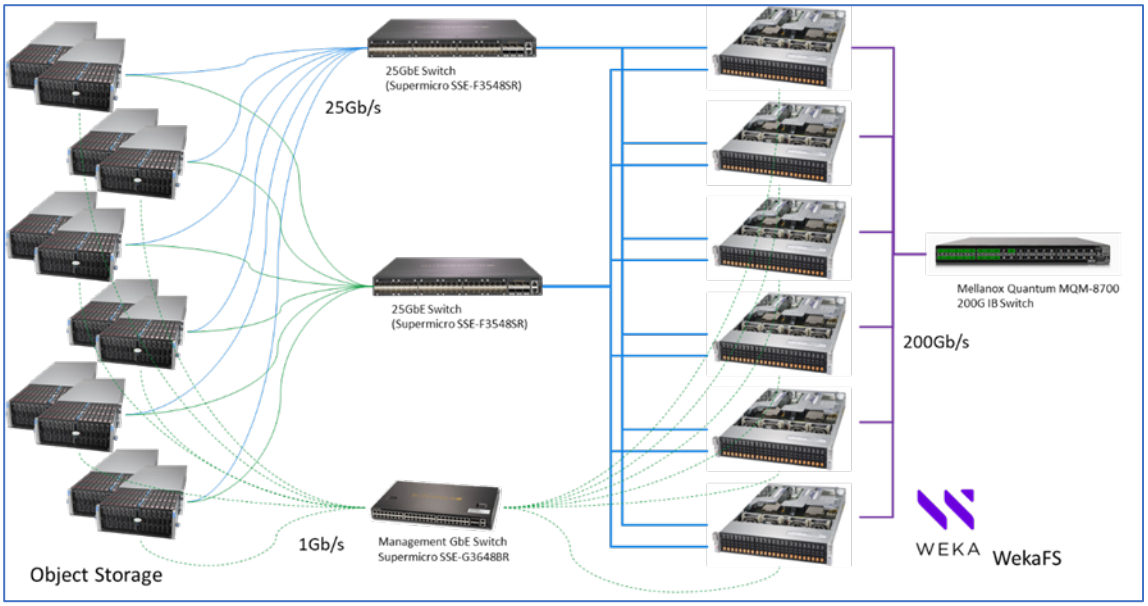


*Figure 2 - Cluster Network Topology*

## WekaIO Performance Overview

Using FIO IO generators on 12 AMD clients, massive performance of 217 GB/s (36GB/s per node) was measured on just 6-nodes of Supermicro WIO platform with 3rd Gen AMD EPYC 74F3 processor, 2x200 Gb/s CX6 IO, 20 Kioxia PCIe4 NVME, and WekaIO (10.x.x).

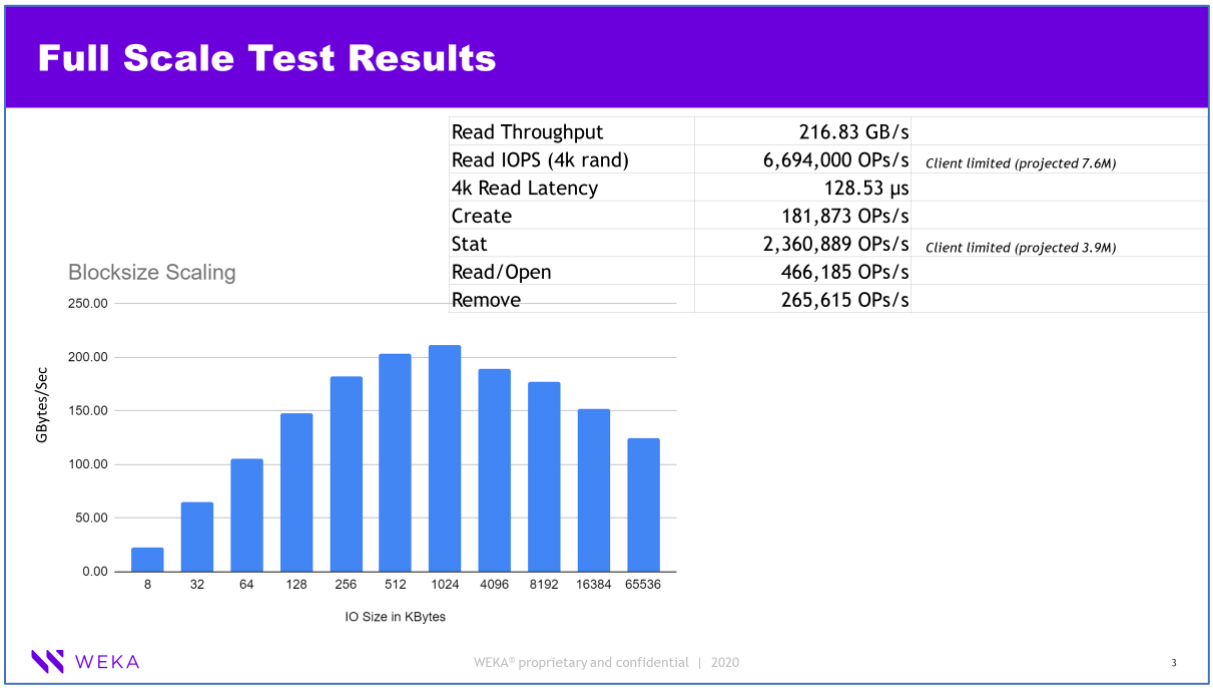


*Figure 3 - Weka Performance Overview*

## Supermicro Server Overview

### SUPERMICRO SERVER AS -2114S-WN24RT

*Figure 4 - 2U UP Single Node Server*

# AS -2114S-WN24RT

| |
|---|
| **Processor Support**<br>Single EPYC 7002/7003 Series up to **280W TDP** |
| **Memory Capacity**<br>**16 DIMM**, up to 4TB DDR4 3200 MHz Reg. ECC |
| **Expansion Slots**<br>**1 PCI-E 4.0 x16** (FH) slot |
| **I/O ports**<br>2 10Gbe RJ45 ports<br>1 Built-in Video<br>6 USB 3.1 ports (4 rear; 2 front) |
| **System management**<br>Built-in Server management tool (IPMI 2.0, KVM/media over LAN) with dedicated LAN port |
| **Storage**<br>2 M.2 Connector (SATA/NVMe)<br>**24 hot-swap U.2 NVMe** drive bays (optional SATA) |
| **System Cooling**<br>3 heavy duty fans w/ Optimal Fan Speed Control |
| **Power Supply**<br>1200W (Titanium Level) Redundant* PWS |

Note: additional PCIe4 slots can be exchanged for 4 NVME connect (2 max)

## Case Study: Biomedical and Life Sciences Research

The data center infrastructure team's key focus at biomedical research centers is supporting high-performance computing (HPC) for general bioinformatics work. One very common workflow is Next-Gen Sequencing (NGS) analysis using the GATK pipeline for sequence alignment and variant calling. However, the HPC cluster has also to support numerous research jobs running simultaneously with unique toolsets. The data center infrastructure team's challenge is to architect a system for scientists with growing informatics needs for their research: more compute, more storage, faster storage, and "bigger" data.

The team typically manages the data from high throughput lab instruments, such as NGS and Electron Microscopy. (EM) They support mixed workloads that can vary greatly, ranging from jobs with a few very large files (100s GB), to jobs with many very large files, to jobs with thousands of tiny (< 1MB) files, and jobs doing lots of meta operations. Existing storage architectures based on decades-old technologies used to meet the needs of capacity and were never a significant bottleneck to productivity when software stacks were not well developed. But with the increase in workload size, complexity, and throughput and the maturing of analysis workflows (highly-tuned CPUs and the emergence of GPUs), the legacy storage system can no longer keep pace with the growing performance demands. For the "hot" active data tier, researchers demand faster storage to handle mixed workloads and a cost-efficient object storage back-end solution to manage larger capacity. The ideal storage solution would:

- Provide better throughput to remove the storage I/O bottleneck and speed data access to the applications
- Enable concurrent research jobs
- Tier seamlessly to object storage archive

- Be cost-efficient
- Future-ready the data center that has begun exploring and using GPUs to accelerate compute

By implementing the WekaIO solution, researchers have achieved better throughput and run more research jobs concurrently without negatively impacting other jobs or workloads. In addition, the turn-around time is better because the jobs finish faster and the results get to the scientist quicker, which accelerates the next stage of their research. The research workflows are greatly simplified because using Weka eliminates the complexity of staging data in and out of a compute node's local SSD. Research outcomes are no longer limited by how much data can be stored on a compute node's local SSD, and with WekaFS acting as a front-end, they have faster and easier access to their object archive tier, ensuring the applications have access to all the data. Ultimately, researchers now have so much performance and expandable capacity available to all nodes that nobody has to think about storage any longer, enabling a greater focus on outcomes of their research.

For one customer, they had the following results:

- Faster Time to Answer: research jobs were reduced 10X; one job was reduced from 70 days to 7 days; another typical analysis workflow was reduced from 12 hours to 2 hours.
- Multiple Concurrent Projects: researchers supported 3x the number of simultaneous new research initiatives, all while having faster turn-around time and job completion times. This gets results to researchers faster.
- Cost-Efficiency: The object tiering feature of Weka doubled the scratch space's available capacity for a lower overall storage cost. Alternative solutions based on All-flash combined with object storage for additional capacity were found to be 1.9-2.4X the price of Weka per usable TB. And while other SATA-based hybrid solutions being considered were slightly less $/TB, the massive performance improvement and time-to-value delivered by Weka offset the slight difference in acquisition cost. By comparing $/IOP or $/RW throughput, the Weka solution came out 8-10X ahead of both the all-flash and hybrid solutions.

## Summary

Dramatic improvements in computational power and exascale needs for storage in today's digital mediums have meant that typical file systems traditionally used to address complex workloads are often impractical or inadequate to the task. WekaIO combined with Supermicro servers provides a stunning performance, protection, and data management story for Deep Learning, High-Performance Compute, and high-throughput low-latency storage workloads. WekaIO removes your computational storage bottlenecks by leveraging the power of NVMe and task-optimized servers with software designed for performance, scalability, and flexibility.

The combination of Supermicro with the WekaIO application solution provides customers with solutions that can leverage our collective building-block architecture design to provide for the most cost-optimized capital expenditure while maximizing operational efficiencies. With Supermicro's professional services, the Supermicro Rack Integration can fully integrate, pre-tested, tuned, racked, and be operational in less than 30 minutes after receiving.

## Additional Resources

WekaIO - https://www.weka.io/how-it-works/
Supermicro Servers https://www.supermicro.com/en/products/rackmount
Contact: Weka-Milan@supermicro.com