SOLUTIONS BRIEF

Distributed File Storage Deployment with Optimized Supermicro All-Flash NVMe Systems

Featuring WekaFS™ Software

TABLE OF CONTENTS

EXECUTIVE SUMMARY 1

WEKAIO MATRIX™ 3

AN ENGINEERED SOLUTION 4

CHOOSE A PRE-CONFIGURED SOLUTION THAT DELIVERS PERFORMANCE, FLEXIBILITY, AND VALUE 5

DERIVE BENEFITS FOR YOUR BUSINESS ..

SUPERMICRO STORAGE APPLIANCES ...

STORAGE SERVER CONFIGURATION .....

SUPERMICRO

Supermicro is a global leader in high performance, green computing server technology and innovation. We provide our global customers with application-optimized servers and workstations customized with blade, storage, and GPU solutions. Our products offer proven reliability, superior design, and one of the industry's broadest array of product configurations, to fit all computational need.

EXECUTIVE SUMMARY

The boom of Big Data in today's world has created amazing new opportunities for innovation but unforeseen problems as well, particularly in the areas of artificial intelligence (AI), machine learning (ML), deep learning (DL), and high-performance computing (HPC) workloads. These workloads hold the promise of finding solutions for significant challenges across industries, including healthcare and financial services. To accelerate the process of obtaining insight from mountains of data, these workloads require a modern infrastructure with extreme performance in compute, storage, and networking resources. Supermicro has teamed with WEKAIO to create a solution that addresses these needs and can help you gain the insight that you need to gain a competitive advantage, deliver customer value, and grow your business.

The Supermicro Storage Appliances, featuring WekaFS, the world's fastest and most scalable file system, are preconfigured and optimized for maximum acceleration of AI workloads and reduced training times, delivering unmatched performance at scale. This family of appliances is ideal for today's data-intensive applications in artificial intelligence and technical computing. It has proven scalable performance, delivering over 10x more performance than blade-based, all-flash, scale-out NAS, and 3x more than locally attached NVMe SSDs. This solution supports deep learning workloads that are akin to HPC but with unique storage requirements. Because these workloads may consist of billions of small files, the storage system must deliver high throughput, low latency, and excellent metadata performance at extreme scale to keep graphics processing units (GPUs) and applications saturated with data. This requirement is unattainable with legacy storage systems but made possible with WekaIO's modern NVMe-optimized, distributed, and parallel file system running on Supermicro servers.

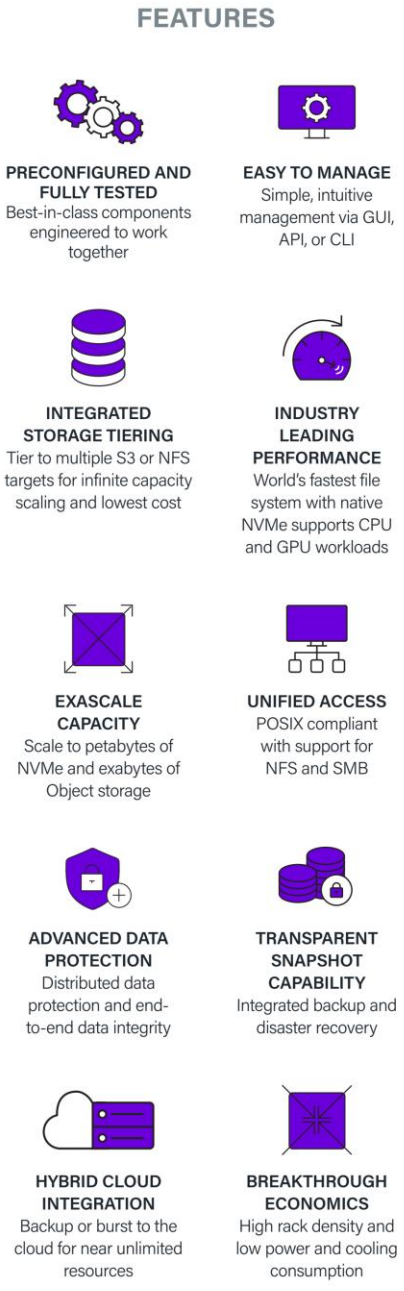

WekaIO MATRIX™

WekaIO's Matrix is the world's fastest and most scalable file system, ideal for today's data-intensive applications in artificial intelligence (AI) and technical computing. It has proven scalable performance, delivering over 6x more performance than blade-based all flash scale-out NAS and 2x more than locally attached NVMe SSDs.

Matrix uses a clean-sheet design with an optimized network stack that runs Ethernet (100Gbit and above) or InfiniBand, so data locality is no longer a necessary factor for performance. The software presents a POSIX file system that distributes both data and metadata across the entire storage cluster to ensure massively parallel access. Using NVMe flash technology ensures the highest performance and lowest latency.

AN ENGINEERED SOLUTION

Based on Supermicro's A+ Server, BigTwin™, and Ultra servers, WekaIO's Matrix software, and Mellanox's network adapters, the engineered solution provides industry-leading performance and value. The BigTwin™ is the first and only 2U multi-node system supporting the highest performance processor, memory, storage, I/O, and an incredible 30% better thermal capacity. Mellanox ConnectX-5/6 adapters provide the choice of either Infiniband or Ethernet to deliver 100Gb/s bandwidth in a single port, the lowest available latency, 150 million messages per second, and application hardware offloads, satisfying even the most demanding application requirements.

CHOOSE A PRE-CONFIGURED SOLUTION THAT DELIVERS PERFORMANCE, FLEXIBILITY, AND VALUE

Based on the Supermicro A+ Server, Ultra SuperServer® and BigTwin™ server platforms, WekaFS software, and Mellanox network adapters, you can take advantage of a plug-and-play engineered solution that helps extract greater value from data. This solution is available in a set of appliances that allows you to easily select a choice that best suits your performance, capacity, and footprint requirements. You can make your decision based on your need for:

- Entry-level capacity with granular scaling of nodes (1U model)

- The best value per GB with granular scaling of nodes (1U model)
- Best performance density per RU (2U model)

The 1U models offer flexibility, an attractive entry price, and the highest number of drives and performance per server node; the 2U model provides the highest performance density per rack unit and great power and space economy. The Supermicro BigTwin 2U Storage Appliance is the industry's first 2U multi-node system with high-performance processors, memory, storage, and I/O. Furthermore, it can lower energy consumption in the data center with an incredible 30 percent better thermal capacity. Performance and capacity can be easily increased with expansion storage for each appliance and drives for partially populated appliances.

WekaFS is a POSIX-compliant file system that distributes both data and metadata evenly across the entire storage cluster to ensure massively parallel access. WekaFS reduces time to innovation by delivering more data to the applications that need it, faster than any other storage system. With a single namespace that can offer on-premises storage and cloud connectivity, the software delivers simplified storage management and data protection. Its performance is 3x that of local file systems and 10x that of traditional NAS. Mellanox adapters provide the choice of either InfiniBand or Ethernet to deliver 100Gb/sec bandwidth in a single network port, the lowest available latency, 150 million messages per second, and application hardware offloads, satisfying even the most demanding application requirements. The SPEC SFS®2014[1] results show that WekaFS has performance that is far superior to all competitors in Overall Response Time (ORT) for database, Electronic Design Automation (EDA), Video Data Acquisition (VDA), Virtual Desktop Infrastructure, and software build workloads.

Derive Benefits For Your Business

These high-performance storage solutions not only address your specific needs based on workload types and performance and capacity requirements but go beyond that to provide more significant overall benefits for your business:

- Optimization of your IT environments with a modern storage infrastructure and simplified storage management
- Datacenter agility with faster data access and effective resource utilization
- Data transformation for machine learning and analytics with faster time to value and insight.

In partnership with Weka, Supermicro and Weka deliver a differentiated solution beyond current market and performance standards for storage and the best solutions for your IT and business challenges.

[1] SPEC® SFS2014 Results: https://bit.ly/35UHLDk

SUPERMICRO STORAGE APPLIANCES (8-NODE MINIMUM)

	A+ Server	BigTwin		Ultra	
SMC SKU	8x AS-1114S-WN10RT	2x SYS-2029BT-HNR		8x SYS-1029U-TN10RT	
Total Rack Units	8U	4U	4U	8U	8U
Software License (per Usable TB)	323TB	193TB	64TB	161TB	64TB
Storage Media NVMe[2]	80 (10x 7.68TB/node)	48 (6x 7.68TB/node)	16 (2x 7.68TB/node)	80 (10x 3.84TB/node)	32 (4x 3.84TB/node)
CPU	8x AMD Rome 7402P	16x Intel 4214R	16x Intel 4214R	16x Intel 4214R	16x Intel 4214R
Memory	2048GB	1536G	768GB	1536G	768GB
Network	16x single-port Mellanox CX-6 200G VPI	16x single-port Mellanox CX-5 100G VPI	8x single-port Mellanox CX-5 100G VPI	16x single-port Mellanox CX-5 100G VPI	8x single-port Mellanox CX-5 100G VPI
4K Read IOPs	6.8M[3]	4.9M	1.6M	6.8M	2.7M
Bandwidth	320GB[4]/s	119GB/s	40GB/s	160GB/s	80GB/s
4K Read IO Latency	Sub-400µs	Sub-400µs	Sub-400µs	Sub-400µs	Sub-400µs
Data Protection	• Distributed Data Protection (N+2 or N+4) • Drive Hot Sparing • Error Detection: End-to-end Data Protection • In-flight and at-rest Data Encryption				
Protocols	POSIX, NFS, SMB, S3 gateway				
Snapshots and Clones	File System Level, Up to 1024 Snapshots				
Tiering	S3 Compatible Cloud Object Store (Public or Private), and Ceph				
System Monitoring	Cloud-based Monitoring and Analytics for Application Tuning and Remote Support				
Software Subscription	1-year or 3-year				

[2] *NVMe capacity options include 1.92TB, 3.84TB, 7.68TB, and 15.36TB. For different capacities, drives request, please contact weka-pm@supermicro.com.
[3] Estimated.
[4] Estimated.

STORAGE SERVER CONFIGURATION[5] – MINIMUM 8 NODES

A+ Server – 1 node

COMPONENTS	SPECIFICATIONS	QUANTITY
Supermicro Storage Enclosure	AS-1114S-WN10RT	1
CPU(s)	AMD Rome 7402P, 2.8GHz	1
RAM	32GB DDR4-3200	8
Networking	Mellanox ConnectX-6 VPI Single Port PCIe	2
NVMe(s)	3.84TB U.2	10
Boot Drive	960GB or larger M.2 SATA	1

BigTwin – 4 nodes

COMPONENTS	SPECIFICATIONS	QUANTITY
Supermicro Storage Enclosure	SYS-2029BT-HNR	1
CPU(s)	Intel® Xeon® 4214R 2.4GHz	8
RAM	16GB DDR4	48
Networking	Mellanox ConnectX-5 or later 100Gbit/s VPI Single Port PCIe	8
NVMe(s)	3.84TB U.2	24
Boot Drive	480GB or larger M.2 SATA	4
Supermicro I/O Module	AOC-MTG-i2TM-O	4

[5] Configuration details subject to change without notice. Please contact weka-pm@supermicro.com for custom configuration details.

Ultra Server – 1 node

COMPONENTS	SPECIFICATIONS	QUANTITY
Supermicro Storage Enclosure	SYS-1029U-TN10RT	1
CPU(s)	Intel® Xeon® 4214R 2.4GHz	2
RAM	16GB DDR4	12
Networking	Mellanox ConnectX-5 or later 100Gbit/s VPI Single Port PCIe	2
NVMe(s)	3.84TB U.2	10
Boot Drive	480GB or larger M.2 SATA	1