



# Performance Benchmarks on Intel® x86 Platforms

Shared Accelerated Storage (SAS) is a new category of software centric storage that decouples storage services from the underlying physical hardware and uses solid-state drive (SSD) technology to deliver extremely high performance and low latency. The result is a hardware independent storage architecture based on industry standard servers, memory, networking, and storage devices that is both agile and highly responsive to the applications it supports. This technology brief illustrates how the WekaIO Matrix™ storage platform provides unprecedented file-based storage performance and scalability in private and hybrid cloud environments, leveraging Intel’s high-performance processors, networking, and solid-state drive technology, delivered on Dell servers. Enterprises will benefit from improved infrastructure flexibility, reduced storage costs, and better performance and scalability.

## OVERVIEW

Today, many large enterprise IT organizations support thousands of users and several hundred applications. The diversity of supported workloads places significant demands on storage infrastructure, and real-time applications such as security log management and line-of-business applications require high performance computing (HPC) levels of performance and scalability. Some applications demand small-file performance and are I/O and latency sensitive, while others require large-file performance, high bandwidth, and throughput. IT organizations struggle to deliver the right mix of storage services on-demand due to the structural limitations of legacy external storage architectures. Traditionally, the solution has been to buy more dedicated storage resources, resulting in data silos and complex infrastructure that is un-sharable, difficult to manage, and does not scale. (Figure 1)

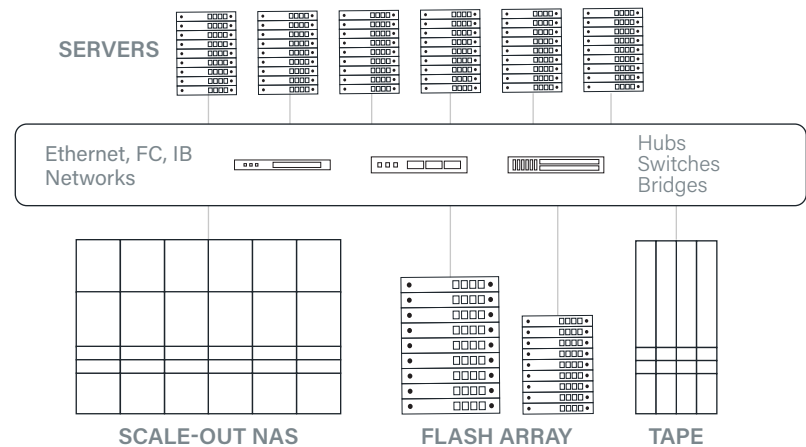


Figure 1 – Traditional External Storage Complexity

WekaIO Matrix is a next-generation software solution designed to solve the key challenges of on-demand shared storage— performance, cost, and scalability. Rather than utilizing existing technology with pre-established design limitations, Matrix was written from scratch to leverage the benefits of standard Intel x86 architecture combined with flash technology. The result is an easy to deploy, easy to manage storage architecture that is a radical departure from traditional purpose built appliances. By optimizing Matrix for flash, the storage solution is ideal for mixed use workloads that consist of large and small files with random or sequential access patterns.

WekaIO Matrix platform utilizes Intel’s DPDK technology, found on high-performance Intel network interface cards (NICs), for best latency and performance. Through optimization in the network stack, Matrix is capable of servicing I/O requests with

application latencies as low as 200 microseconds. Figure 2 illustrates the flexible deployment options—hyper-converged mode where Matrix runs as a process on the compute cluster like any other application or on dedicated storage servers. The result is a software based storage solution that delivers the performance and latency of an all flash array, the simplicity and manageability of a NAS system, and the scalability of cloud storage.

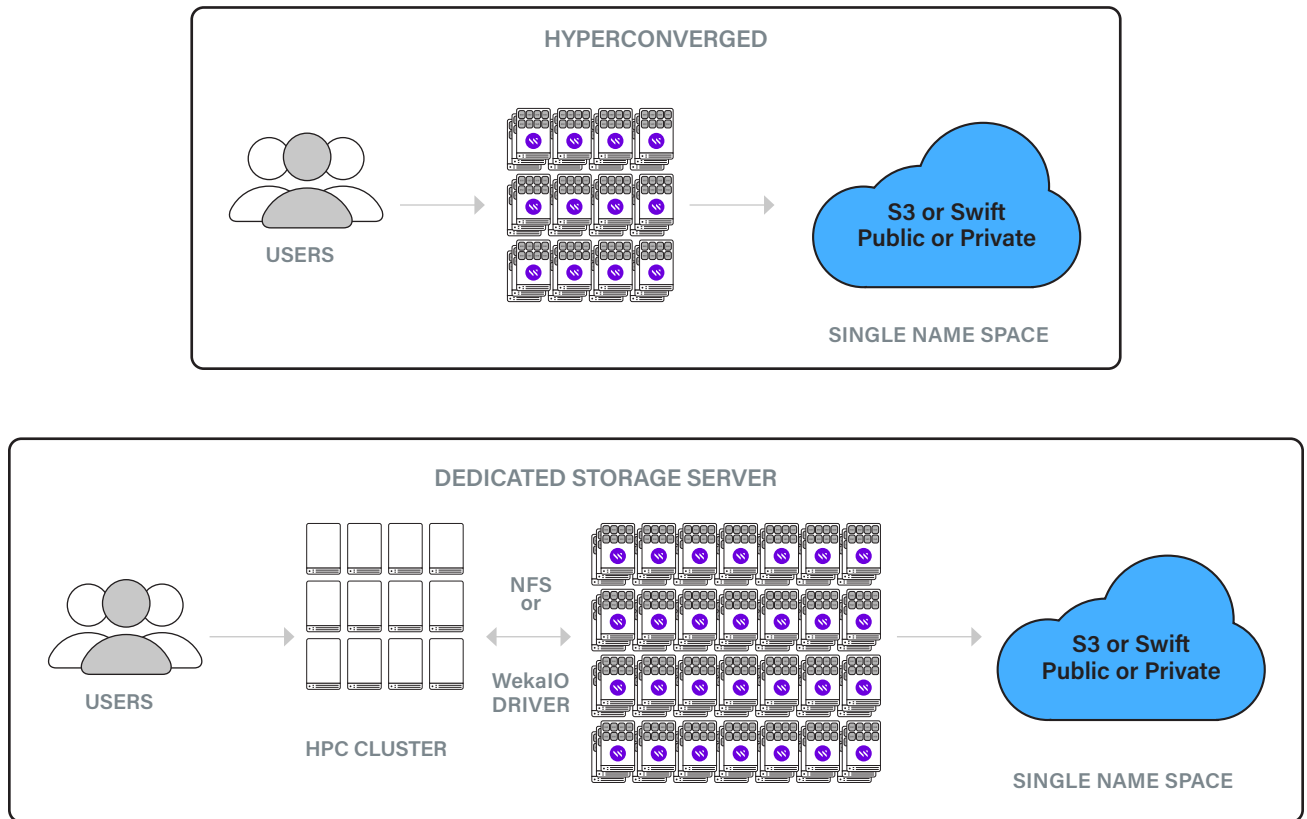


Figure 2 – Matrix Deployment Modes

## DYNAMIC PERFORMANCE AND CAPACITY SCALING

Most traditional storage solutions scale performance and capacity in tightly coupled, pre-defined increments. With Matrix, performance and capacity scale independently and dynamically in user defined increments. Performance can be scaled up or down based on application needs and utilizes SSD or NVMe technology to provide extremely high performance regardless of file size or access pattern. Matrix eliminates the additional power, cooling, and floor space required for a traditional storage array by utilizing the storage slots inside each Intel x86 standard server. As capacity needs change, add more SSDs to any available storage slot in a Matrix enabled server. The file system automatically detects and assimilates the new device.

## ADVANCED DATA PROTECTION FOR DATA RESILIENCY

WekaIO Matrix data protection scheme consists of multiple layers of protection to ensure the highest levels of availability, performance, and data resiliency, with only 25% additional overhead. These layers include checkpointing, distributing data across the server cluster, and a proprietary error correction scheme that is faster, more efficient, and more resilient than replication, RAID or other forms of erasure coding. Our patented data protection scheme can easily handle multiple failures and is configured automatically based on the cluster size and other parameters. In the event of a failure, all of the Matrix enabled nodes participate equally in the recovery to ensure data is accessible and performance is maintained.

## SCALE STORAGE INFINITELY

MatrixFS, the underlying distributed, parallel file system, delivers cloud scalability by aggregating server based SSDs and optional S3 or Swift compatible cloud storage into a single namespace for easy management. Data and metadata are dynamically

and transparently distributed across the file system, ensuring that no performance hotspots occur. The integrated data tiering feature automatically and transparently monitors data access patterns and migrates data to lower cost cloud storage based on user defined rules. The result is that enterprises get the best utilization of the expensive high performance tier, without ever losing long term access to their data, and tiering to public or private cloud storage effectively provides limitless capacity at the lowest TCO.

## PERFORMANCE BENCHMARKS

WekaIO Matrix is designed to run on standard Intel x86 servers. Selection criterion for the server platform used in this test included server popularity, attractive price-performance, and suitability to file based workloads. We chose the popular Dell PowerEdge R430 entry level rack server with Intel processors, NICs, and SSD storage components to demonstrate the performance that can be achieved with Matrix. This highly affordable, basic server is an ideal building block for storage environments that need a scale-out architecture to support high performance computing (HPC), web tech, DevOps, EDA, genomics, media rendering and similar storage intensive workloads.

### DELL R430 TEST PLATFORM

COMPONENT	SPECIFICATIONS	QUANTITY
Enclosure	1U enclosure with 8 2.5" hot plug drive bays	1
Network	Intel x520 DP 10Gb DA/SFP+	2
Boot Drive	120GB SSD SATA 6Gbps hot plug	1
Processor	Intel Xeon E5-2620 V3 2.4GHz	1
SSD Storage	Intel S3610 1.6TB SATA SSD, MLC 6Gb/sec, 2.5" hot plug	5
Memory	8GB RDIMM 2400MT/s	8
Power	Hot plug, power supply 550W	2
Operating System	Red Hat Enterprise Linux	1

6x Dell R430 rack servers

Matrix Version 2.1


1 core per server assigned to MatrixFS (in hyper-converged mode)

Data protection level 4+2

SpecSFS 2014 SP2 Test Suite

- 4K READ
- 1M READ

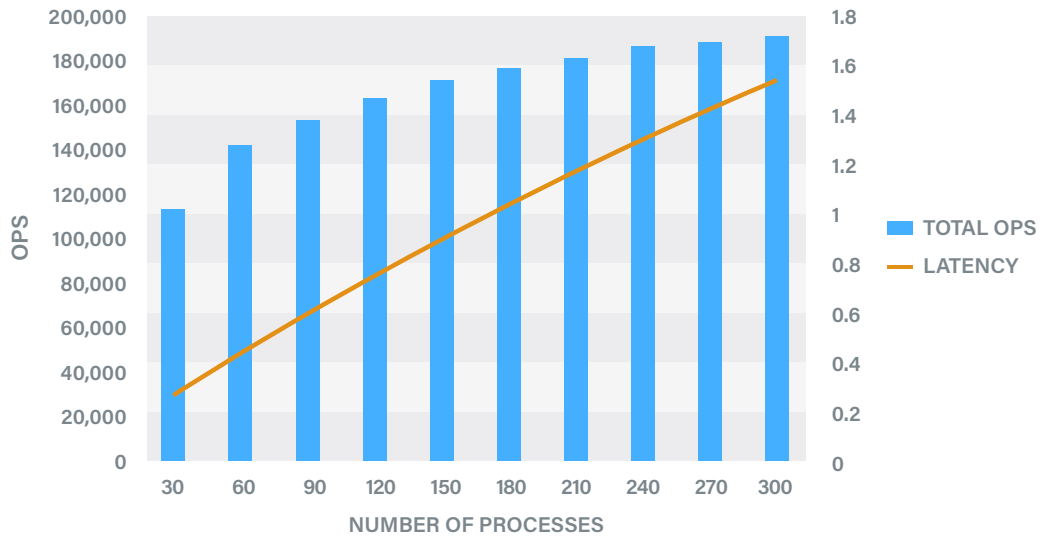
### PERFORMANCE TEST PARAMETERS



### TEST RESULTS:

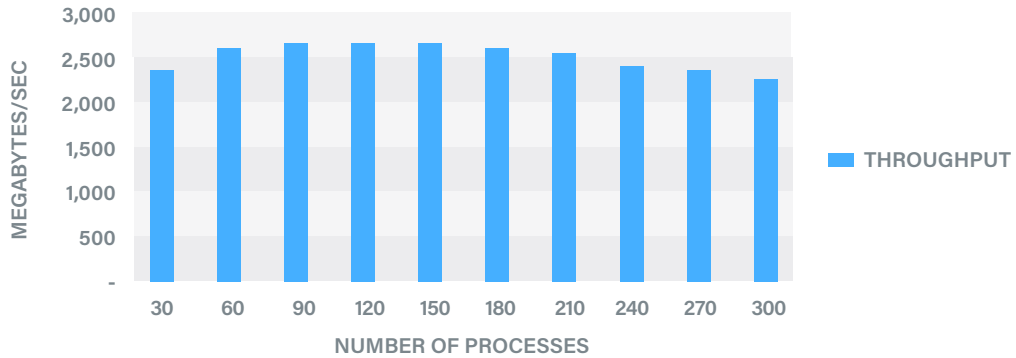
The MatrixFS enabled six-server cluster delivered 192,551 OPS. That equates to 32,092 OPS per MatrixFS enabled core utilizing Intel SATA SSDs. A latency of 260 microseconds was recorded for 30 processes, as the number of processes grew to 300, latency only increased to 1.56 milliseconds.

**Total OPS and Latency – 4K Read - 6 Dell R430 Servers Each With a Single Core**



The MatrixFS enabled cluster delivered a peak throughput of 2.7 gigabytes/second throughput across the six-server cluster. That equates to a peak throughput of 452 megabytes per second per MatrixFS enabled core. Bandwidth is limited by the NIC throughput and can be doubled with 2 NICs per server.

**Total Throughput – 1M Read - 6 Dell R430 Servers Each With a Single Core**



**SUMMARY**

The benchmarks demonstrate that very high performance can be achieved with Matrix storage software running on standard Dell rack servers with Intel x86 processors. Utilizing a single core from the Intel Xeon E5-2620 v3 processor, Intel 10Gbit Ethernet NICs, and Intel high-performance SATA SSD’s delivered through the Dell R430 series rack server, the MatrixFS software delivered a very high performance integrated storage tier within the server cluster. This combined solution eliminates the need for costly or specialized storage appliances because applications and storage share a single high performance platform.

The system can scale to hundreds of servers where performance scales up linearly with each new MatrixFS enabled core. A 100-node server with a single core dedicated to MatrixFS can deliver up to 45 gigabytes per second of throughput and over 3 million OPS, all with extremely low latency. More cores in the cluster can be assigned to MatrixFS during peak performance periods and can be spun down when the demand surge has ended. No other storage platform provides this level of performance and agility.

