SOLUTION BRIEF

Intel® Storage Builders

WekaIO Matrix™
Intel® Xeon® Processor E5-2600 Product Family
Intel® Ethernet Converged Network Adapter X520
Intel® SSD Data Center Family
Data Plane Development Kit

Create a Flexible, Scalable High-Performance Storage Cluster with WekaIO Matrix™

Matrix and Intel® SSD Data Center Family Improve Storage Performance and Scalability

Executive Summary

Across industries, data generation is growing exponentially, creating in enterprise IT both an incredible demand for storage capacity as well as the capability to access it quickly—often in real-time.

WekaIO software creates an extremely high-performance, flexible, and scalable, parallel file system, deployed either hyper-converged onto an existing compute cluster or as a dedicated storage platform. The WekaIO Matrix™ storage platform delivers a POSIX-compliant, single namespace data repository that scales in both performance and capacity by leveraging Intel® architecture, Intel® SSDs, and Intel® technologies, running on industry-standard, high-volume servers in private and hybrid cloud environments. The software is optimized for the high-performance Intel® Xeon® processor E5-2600 Product Family, Intel® Ethernet Converged Network Adapters, and Intel® SSD Data Center Family of storage devices.

WekaIO's software-based storage technology is transforming enterprise and cloud storage by radically simplifying the way storage is deployed in the data center. Running Matrix, enterprises will benefit from improved infrastructure flexibility, reduced storage costs, and better performance and scalability for their data demands.

The Need to Rethink Storage

Today, many large enterprise IT organizations support thousands of users and several hundred applications with an ever-increasing hunger for data acquisition and consumption. The diversity of supported workloads and the exponentially expanding generation of data places significant demands on the storage infrastructure.

Real-time applications, such as security log management and line-of-business informatics, require high-performance computing (HPC) levels of performance and scalability—both in compute and storage. These diverse workloads consume data differently. Some process small files and are I/O—and latency—sensitive, while others require high bandwidth and throughput for large files.

IT organizations struggle to deliver the right mix of storage services on demand, due to the inflexibility, complexity, and lack of federation with a mix of legacy storage architectures, making them difficult to manage and easily scale (Figure 1). Traditionally, the solution has been to add more hard drives to a Network Attached Storage (NAS) system, resulting in data silos that are un-sharable, difficult to manage, and do not scale. Alternatively, available high-performance parallel files systems are either difficult to deploy and operate or overly complex and expensive.

WekaIO—Shared Accelerated Storage

WekaIO Matrix is designed to solve the key challenges of on-demand shared storage—performance, cost, and scalability (Figure 1). Matrix is a zero-footprint, hardware-independent, Shared Accelerated Storage (SAS) software solution. SAS is a new category of software-centric storage that decouples storage services from the underlying physical hardware and le-verages solid state drive (SSD) technology to deliver extremely high-performance and low latency access to data. The result is a hardware-independent, parallel file system architecture based on industry-standard servers, memory, networking, and storage devices that is agile, scalable, and highly responsive to the applications it supports.

Rather than utilizing existing technology with pre-established design limitations, Matrix was written from the ground up—with its own operating system and network stack. Matrix lever-ages the benefits of industry-standard Intel architecture found in Intel Xeon processors, combined with the flash technology of Intel SSD DC Family devices and network performance of Intel® Ethernet Converged Network Adapters.

Matrix can run hyper-converged within a cluster on any number of servers, using CPU cores assigned to it. Or, it can run as a stand-alone, dedicated storage system. The software assembles a storage pool from available SSDs, creating a man-ageable, high-performance, scalable storage repository that is POSIX-compliant and with a single-namespace.

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 and multiple work-loads that consist of large and small files with random and sequential access patterns.

Leveraging the Performance of the Data Plane Development Kit (DPDK)

Matrix's network stack leverages Intel's open source Data Plane Development Kit technology (dpdk.org). DPDK is a set of libraries that takes advantage of Intel technology found on high-performance Intel Ethernet Converged Network Adapters and controllers. Leveraging DPDK, Matrix accelerates and improves networking performance for data access. Through optimizations to the network stack, Matrix is capable of servicing I/O requests with application latencies as low as 200 microseconds.
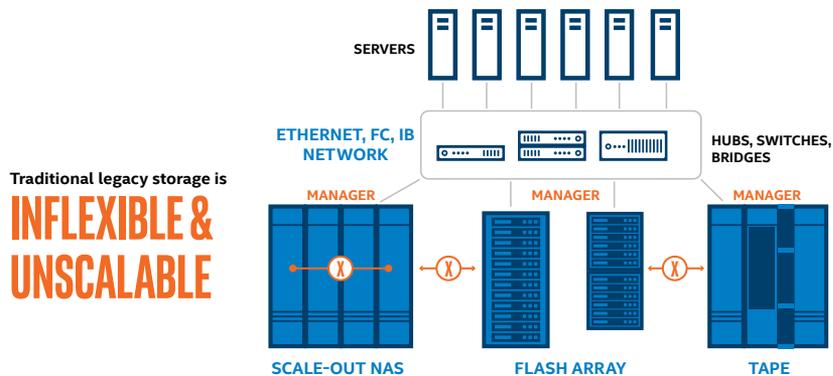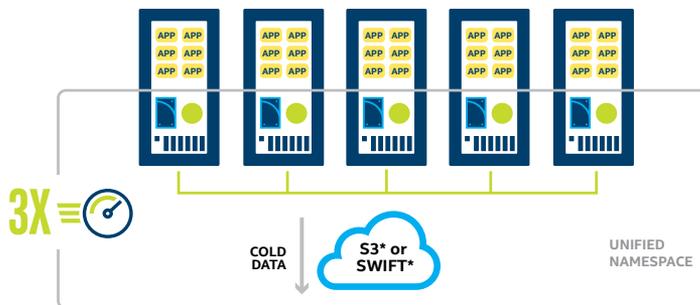

Figure 1. Traditional legacy storage is inflexible and difficult to manage.

Flexible Deployment

Figure 2 illustrates the flexible deployment options of Matrix.

• Hyper-converged mode—Matrix runs as a process on the compute cluster like any other application.

• Dedicated storage—Matrix runs on fixed storage servers, independent of compute.

The result is a SAS solution that delivers the performance and latency of an all-flash array, the simplicity of a NAS system, and the scalability of cloud storage.

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. Matrix utilizes SSDs with SATA or PCIe* interfaces to provide extreme high performance, irrespective of file size or access patterns.

Matrix performance depends on the number of cores assigned to it. It can be dynamically scaled up or down based on application needs at any time. When more storage perfor-

mance is needed—for example, during final verification of a large circuit design—more cores are assigned to process storage requests.

Capacity scales out as nodes are added to the cluster or storage devices are added to a node. As capacity needs change, users simply add more SSDs and/or more nodes. The file system automatically detects and assimilates new devices.

Advanced Data Protection for Data Resiliency

Matrix's data protection scheme consists of multiple layers to help ensure high levels of availability, performance, and data resiliency. These layers include checkpointing, distributing data across the entire cluster, and a proprietary error correction scheme that is fast, efficient, and resilient without using RAID or erasure coding. WekaIO's patented data protection scheme can easily handle multiple failures. It is configured automatically based on the cluster size and other parameters. In the event of a failure, Matrix-enabled nodes participate equally in 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 Amazon 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.

Integrated data tiering technology automatically and transparently monitors data access patterns. It migrates data to lower-cost cloud storage based on user-defined rules, and it automatically moves cold data as local capacity fills. Enterprises get the best utilization of the expensive high-performance tier, without ever losing long-term access to their data. Tiering to public or private cloud storage effectively provides limitless capacity at the lowest TCO.
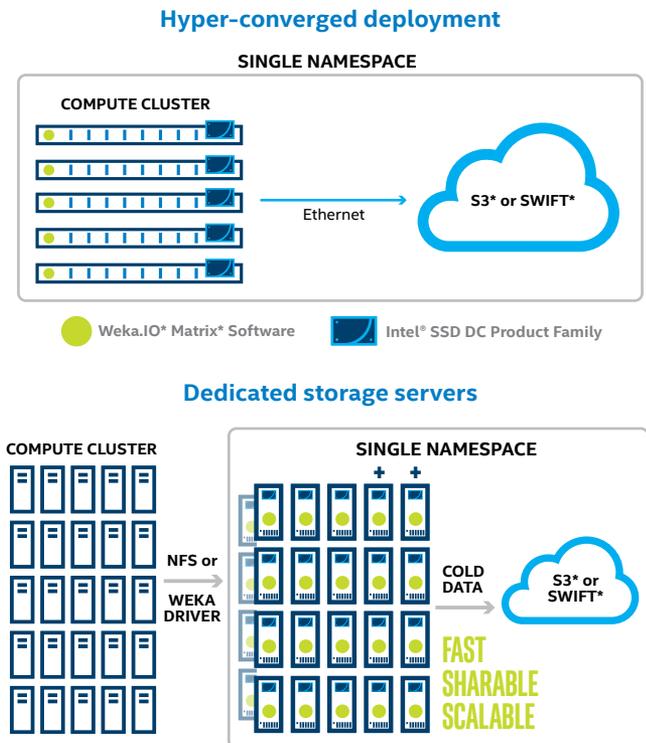
Hyper-converged deployment

Dedicated storage servers

Figure 2. WekaIO Matrix deployment modes.

Performance Benchmarks

The following benchmark testing illustrates the performance of a WekaIO Matrix platform solution. Testing revealed Matrix delivered over 32,000 IOPS per core, over 450 MB/s per core throughput, with latency as low as 260 microseconds across 30 parallel processes.
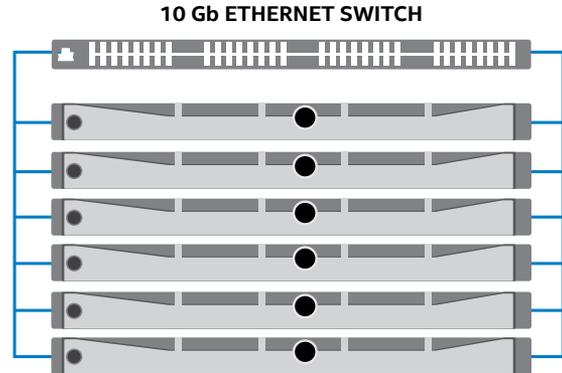
Test Configuration

Selection criteria for the server platform used in the following tests included server popularity, attractive price-performance, and suitability to file-based workloads. Dell PowerEdge* R430 entry-level rack server with Intel® Xeon® processor E5-2620 v3 and Intel® Ethernet Converged Network Adapter X520 were chosen to demonstrate the performance that can be achieved with Matrix on commercially available, industry-standard systems. The servers were populated with Intel® SSD DC S3610 series SATA drives (Table 1).

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 services, DevOps, and similar storage-intensive workloads.

Table 1. Dell R430 Test Platform

COMPONENT	SPECIFICATION	QTY.
Enclosure	1U enclosure with 8 2.5" hot plug drive bays	1
Network	Intel® Ethernet Converged Network Adapter X520, Dual-Port 10 Gbps, DA/SFP+	2
Boot Drive	120 GB SSD SATA 6 Gbps hot plug	1
Processor	Intel® Xeon® processor E5-2620 V3 2.4 GHz	1
SSD Storage	Intel® SSD DC S3610 series, 1.6 TB SATA SSD, MLC, 6 Gb/sec, 2.5" hot plug	5
Memory	8 GB RDIMM 2400 MT/s	8
Power	Hot plug, power supply 550W	2
Operating System	Red Hat Enterprise Linux*	

Testing was completed using six Dell PowerEdge R430 servers clustered as shown in Figure 3 with test parameters listed in Table 2.

10 Gb ETHERNET SWITCH

Dell R430* Server Cluster Running Red Hat Enterprise Linux* 6.x

Figure 3. Test configuration.

Table 2. Performance Test Parameters

6x Dell R430* rack servers
Matrix Version 2.1
1 core per server assigned to MatrixFS (in hyper-converged mode)
Data protection level
SpecSFS 2014 SP2 Test Suite

- 4K READ
- 1M READ

Total IOPS, Latency, and Throughput Results

The MatrixFS-enabled six-server cluster delivered 192,551 IOPS (Figure 4). That equates to 32,092 OPS per MatrixFS-enabled core, utilizing Intel SSD DC Family SATA drives. Latency of 260 microseconds was recorded for 30 processes; as the number of processes grew to 300, latency only increased to 1.56 milliseconds.
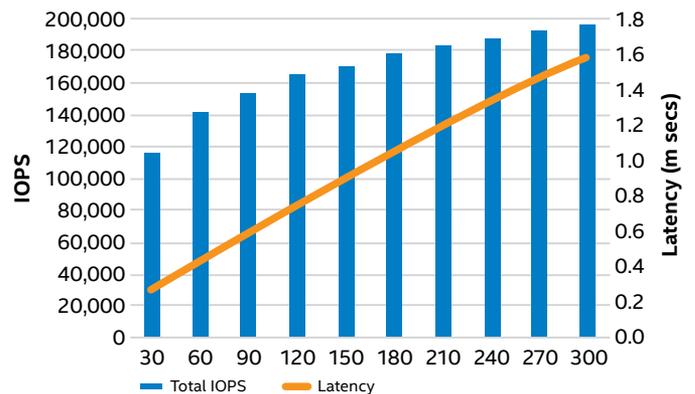

Figure 4. Total IOPS and latency results—4kB Read.

The cluster delivered a peak throughput of 2.7 gigabytes/second across the six-server cluster (Figure 5). That equates to a peak throughput of 452 megabytes per second per MatrixFS enabled core.
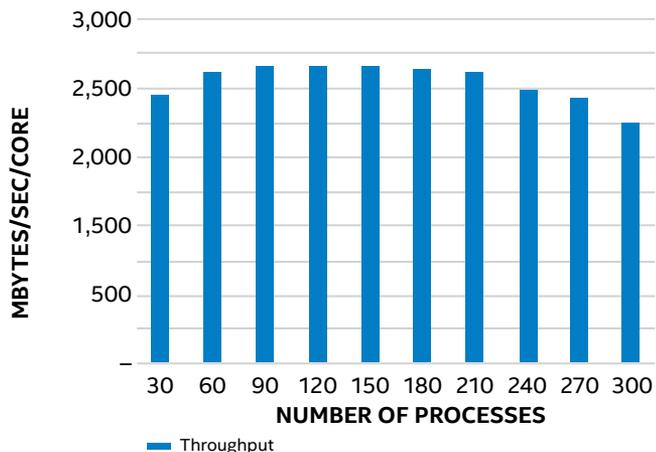

Figure 5. Total single-core throughput results—1 MB Read.

The system can scale to hundreds of servers, where performance scales up linearly with each new Matrix-enabled node. 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 IOPS, all with extremely low latency. More cores in the cluster can be assigned to MatrixFS during peak performance periods, and they can be spun down when the demand surge has ended.

Summary

WekaIO Matrix Shared Accelerated Storage is designed to create a flexible, scalable, high-performance storage system that leverages existing infrastructure of industry-standard servers for low TCO. Matrix can be deployed as part of a hyper-converged cluster or a dedicated high-performance storage resource.

The above benchmarks demonstrate that Matrix can deliver very high performance running on Intel Xeon processor E5-2620 v3-based servers with Intel SSD DC Family SATA drives and Intel Ethernet Converged Network Adapters. Utilizing a single core from the Intel Xeon processor E5-2620 v3, the DPDK libraries enabled the network adapter to maximize throughput. With Intel SSD DC Family of SATA drives in 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.

About WekaIO

WekaIO is a provider of innovative, high-performance Shared Accelerated Storage solutions that turn any standard Intel-based server cluster into a high-performance, scale-out storage system. WekaIO was founded in 2013 and funded by leading venture firms and strategic investors. For more information, visit www.weka.io or follow @wekaIO on Twitter.*

Learn more at http://weka.io

To find out more about Intel storage solutions, visit **www.intel.com/storage**.

For more details about Intel Xeon processors, visit **intel.com/xeon**.

Visit the Intel® Builders solution library: **https://builders.intel.com/solutionslibrary**.

Find us on Twitter with **#IntelBuilders**.