ESG Technical Review

WekaIO Matrix

Date: July 2017 **Author:** Mike Leone, Senior Validation Analyst

Abstract

This ESG Lab Review documents the results of hands-on testing of the WekaIO Matrix software-defined storage solution with a focus on management simplicity, data protection, and performance at scale.

The Challenges

IT is in a continual battle to simplify and "do more with the same, or less," but the continued growth of data has drastically increased the complexity and cost of deploying, managing, and processing data using traditional IT infrastructure components. One key component, enterprise storage—and more importantly, the underlying software that controls the storage—is fortunately evolving nearly as fast as the data is growing. With a goal of simplifying the overall storage experience, meeting strict IT budget requirements, and transitioning to the cloud, software-defined storage solutions are being rapidly adopted. In fact, recent ESG research shows that the top two reasons for evaluating software-defined storage are simplifying management and reducing operational costs.[1]

Figure 1. Factors Driving the Consideration of Software-defined Storage

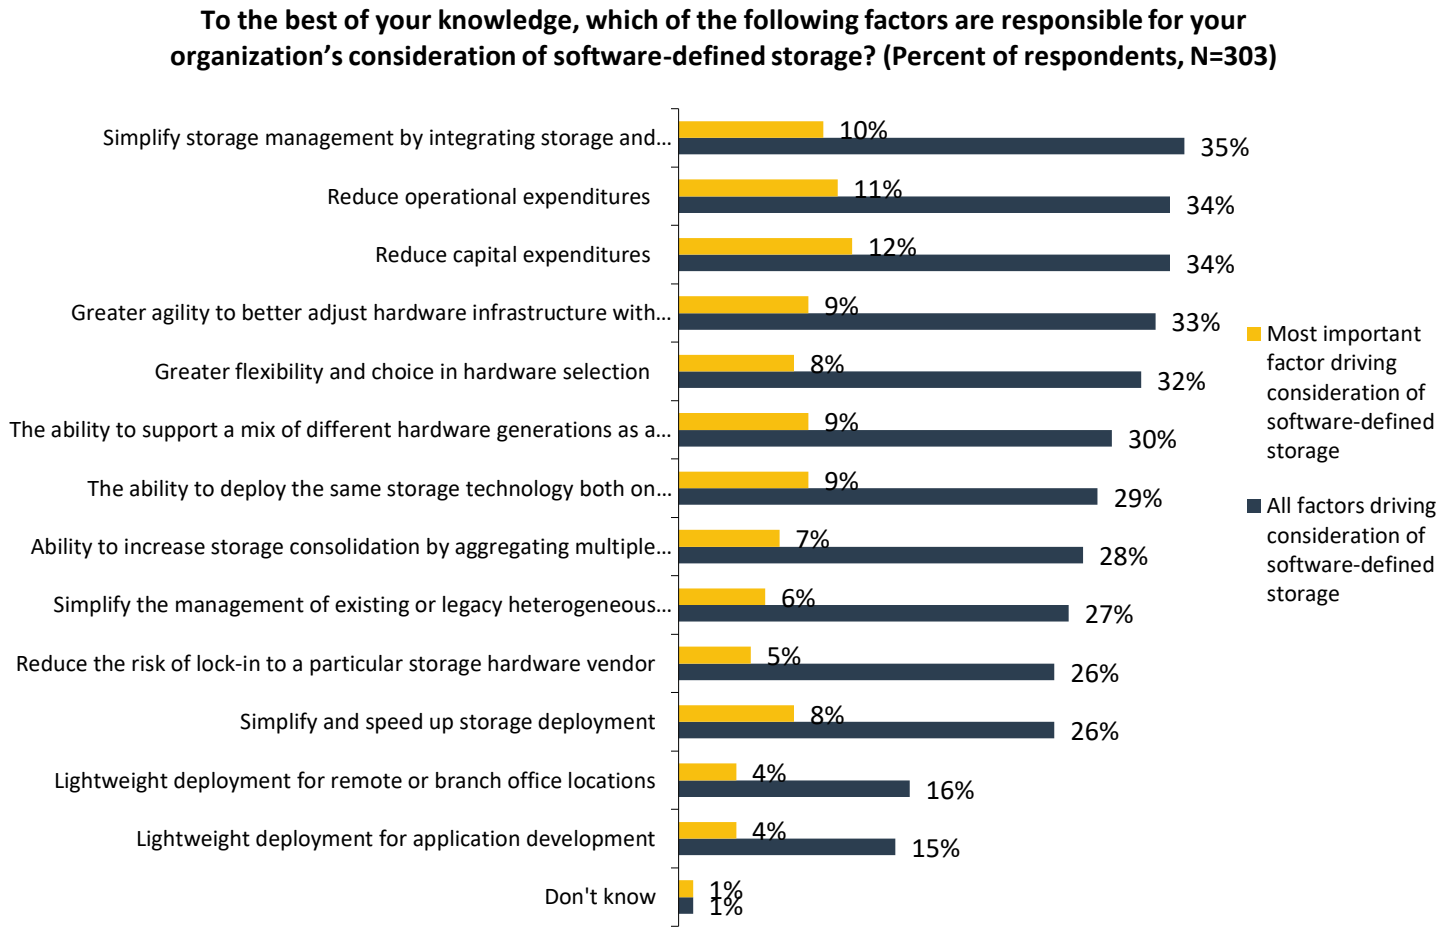

Source: Enterprise Strategy Group, 2017

[1] Source: ESG Research Report, *Software Defined Storage (SDS) Market Trends*, February 2017

The Solution: WekaIO Matrix

WekaIO Matrix is POSIX compliant virtual filesystem software that uses off-the-shelf servers and SSDs to cost-effectively provide high performance, high capacity, and resilient storage at scale, on-premises and in the cloud. The WekaIO Matrix software uses existing host server and network infrastructure to create a shared pool of file storage for applications running on servers in the WekaIO global namespace. With WekaIO, storage is just another application in the compute infrastructure. WekaIO's global namespace can run conventionally on bare metal servers; as a virtual machine (VM) managed by a KVM or VMware hypervisor; as a Docker containerized application; and in the cloud on Amazon Web Services (AWS). Every application server instance, physical or virtual, running WekaIO contributes storage resources and performance. Additionally, WekaIO can use Swift and Amazon S3 cloud storage as a repository in which to store 'cold' data.

As high-performance file storage, WekaIO works well in analytics-heavy applications—data mining, video editing, financial and seismic analysis, genomics and software dev/test—all of which require high performance at high capacity. WekaIO delivers flash storage performance using SSDs that are direct-attached to the application servers. The software can take advantage of SSDs of various sizes and technologies, including NVMe devices available in local servers or in Amazon EC2 I3 instances. The storage capacity can scale by adding SSDs to a server and the performance can scale by allocating additional CPU cores or adding server instances to the global namespace. End user applications can consume WekaIO's file storage in two ways. First, applications can use it in a conventional NAS topology, using an NFS client or a WekaIO-provided native client to access data on a cluster. Second, clients and servers can reside together on the same host servers, to take maximum advantage of multicore server hardware and server virtualization.

WekaIO Matrix has a strong storage heritage and represents significant advances in storage software technology. Currently, WekaIO has applied for 20 patents, seven of which have been issued. WekaIO has identified about 20 more patentable aspects of the WekaIO Matrix technology. WekaIO innovations include the distributed filesystem and distributed data protection (DDP) algorithms, and leverages advanced technologies such as DPDK and SRIOV in conjunction with a customized network stack.

WekaIO Matrix

- Software-defined
- Flexible deployment: hosts & clients
- Advanced data protection
- Large scale capacity & performance
- Cloud-aware compute & storage

WekaIO Matrix Distributed Data Protection (MatrixDDP) is designed to run in a cloud environment and to deliver high data resiliency with minimal performance impact. MatrixDDP distributes data across failure domains, which can be physical host servers or cloud availability zones. MatrixDDP can employ either N+2 or N+4 data redundancy. N+4 data redundancy is especially appropriate for customers who want to scale up SSD capacity in each node or for whom the compute/networking environment is not considered stable. Data protection is inherent in the WekaIO Matrix filesystem with journaling, so recovery from failure is fast, not requiring a full filesystem check.

Cloud computing techniques are used to offer options to WekaIO's on-premises capabilities. First, clouds can be used as a low-cost storage tier by automatically moving seldom-used data to either a Swift or an Amazon S3 object store. WekaIO can expand its global namespace to WekaIO Matrix instances running in the AWS public cloud, for 'cloud bursting' load sharing and to deploy disaster recovery in the cloud. Snapshot capability, within its global namespace spanning on-premises and cloud resources, enables storage to take advantage of public cloud availability zones and multi-site private clouds for disaster tolerance.

ESG Lab Tested

ESG Lab tested the WekaIO Matrix storage solution on WekaIO systems on premises in San Jose, California, and on WekaIO application and storage deployments in the cloud on Amazon Web Services. ESG Lab tested deployment and management simplicity, data resiliency, and performance at scale, on-premises and in the cloud.

Getting Started

The first phase of testing consisted of walking through the installation process for a WekaIO Matrix cluster. This phase also included touring the product's manageability and reporting capabilities, as well as reviewing the creation process for local and tiered (cloud) filesystems.

Installation

ESG Lab tested the creation of the WekaIO cluster onto a ten-server environment. The WekaIO software packages had already been loaded onto the servers and a WekaIO signature had already been applied to the SSDs to reserve them for WekaIO Matrix. ESG considers WekaIO Matrix to be an accommodating software system, in that there are few specific hardware dependencies. Of course, there are still prerequisites, including: 64-bit x86 Intel CPU, Ivy Bridge or later; CPU cores to be dedicated to WekaIO software nodes; 10Gbit, or faster, Ethernet NICs supporting DPDK and SR-IOV to be dedicated as one physical or virtual NIC per core; Layer 3 connectivity between all the WekaIO cores (nodes); SSDs to be dedicated to WekaIO nodes; a 64-bit RHEL/CentOS/Scientific Linux 6.5 or later operating system; and 5GB RAM per WekaIO core plus 6GB RAM per host server. The RedHat Package Manager (RPM) or WekaIO Yum repository can be used to install the four required WekaIO Matrix software packages.

The complete installation process is shown in Figure 2. We began the installation process by logging into the WekaIO GUI on one of the servers. Note that the installation process can start from any of the servers in the cluster. We selected the installation type, which was the formation of a new hyperconverged cluster.

Next, the configuration wizard guided ESG Lab through the rest of the setup process. ESG Lab selected the 10 host servers to form the cluster from the list of servers that WekaIO had discovered, and assigned one core per server to be used by WekaIO. The software then automatically selected one NIC per core to create the data plane. ESG Lab advanced to the ***SSDs*** tab and entered two SSDs per node, and WekaIO automatically selected the devices. Next, we entered cluster configuration information, including cluster name, MatrixDDP stripe size (redundancy level), and a range of valid IP addresses for intra-node communication. We bypassed the ***UDP Host*** tab since it was a hyperconverged installation. Finally, the ***Preview*** tab showed the cluster ready to form, with ten hosts, plus one core and two SSDs per node. This total installation process took less than eight minutes. We clicked on the ***Form*** button and WekaIO created an IO-ready cluster in less than 45 seconds.

Figure 2. WekaIO Installation Process

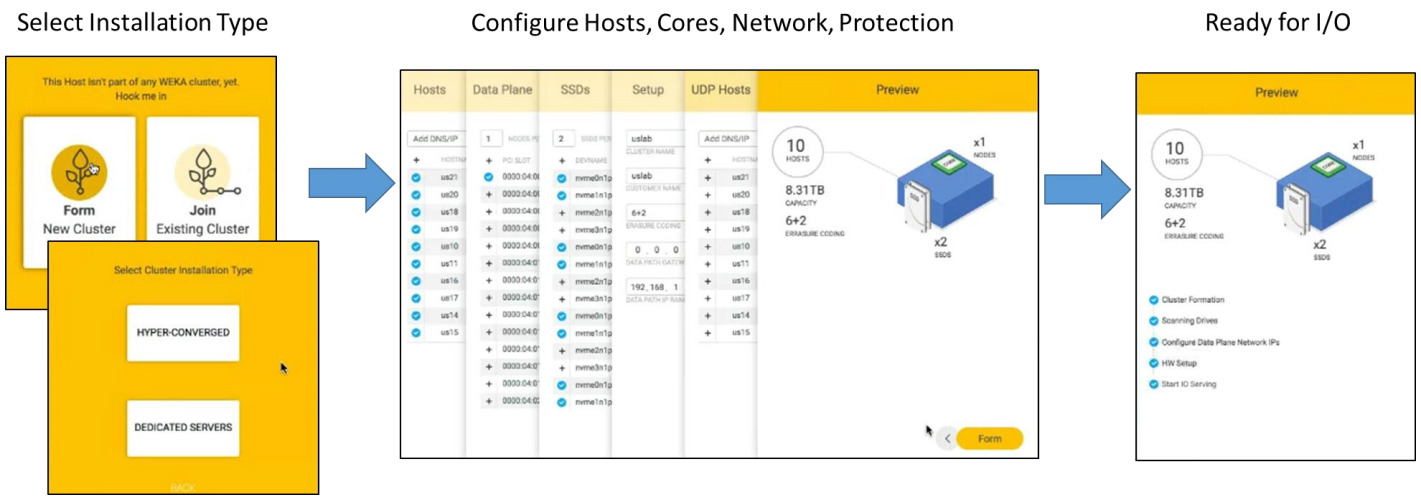

Source: Enterprise Strategy Group, 2017

Management

The WekaIO Matrix management GUI is a well-organized tool for managing and reporting on a WekaIO Matrix cluster. The management application itself is clustered for high availability, with its own quorum process for maintaining control and consistency of management data in the cluster.

The main dashboard of the management GUI, shown in Figure 3, is organized by tabs that appear across the left-hand border. The tabs include:

- SYSTEM OVERVIEW – Cluster dashboard with summary performance details and drill-down capability for nodes. This is the default view upon starting the GUI.
- SYSTEM EVENTS – Event log.
- STATISTICS – Reporting and presentation of cluster information.
- FILESYSTEMS – Configuration and management of groups and filesystems, local and tiered (cloud).
- OBJECT STORE – Configuration and management of cloud storage.
- IP INTERFACES – Manage cluster servers to export NFS.
- NFS CLIENT PERMISSIONS – Specify how NFS exports of WekaIO filesystems.

ESG Lab opened the tab that displayed the recently created ten-node cluster. By clicking on one of the nodes, details about that node were displayed, including IP address, WekaIO version, and the number of drives, cores, and NICs.

Figure 3. WekaIO Matrix Management and Statistics

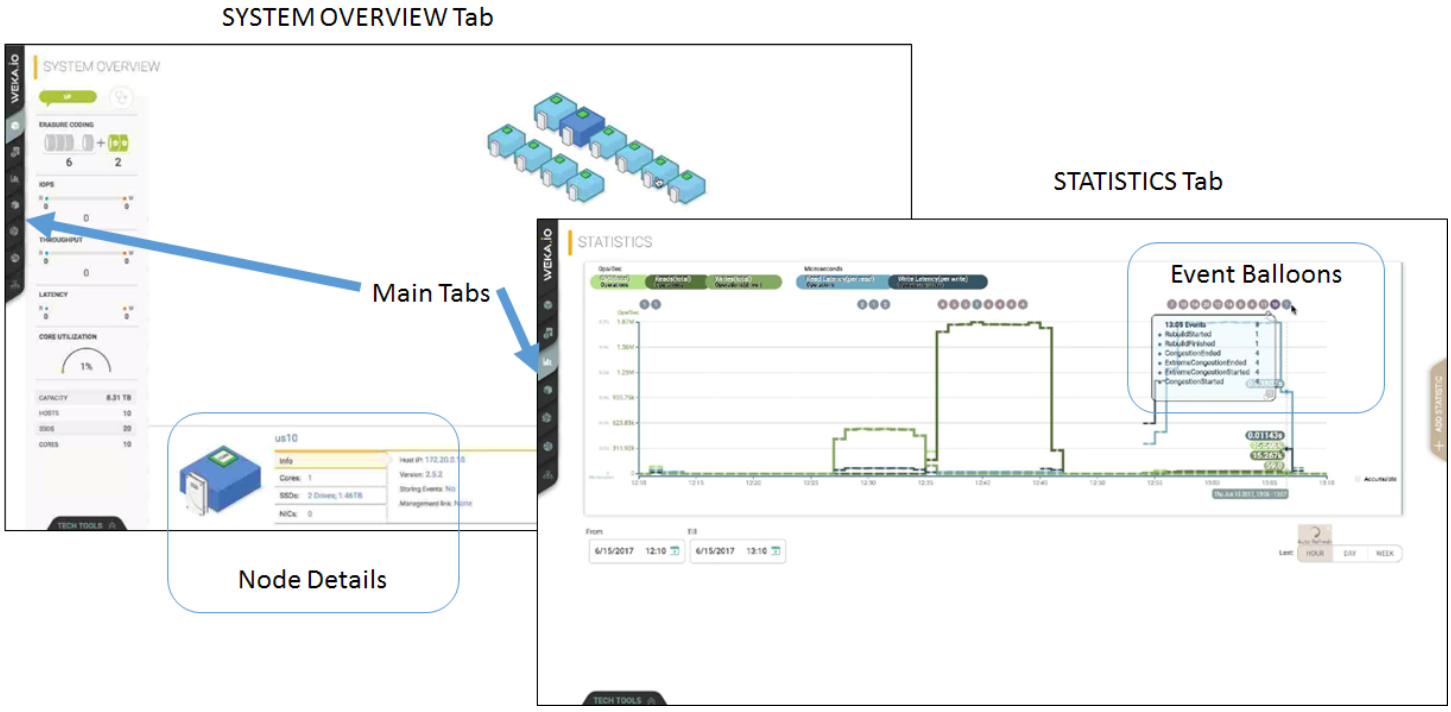

Source: Enterprise Strategy Group, 2017

Next, we logged into a separate cluster that had been running for an extended period of time to explore the ***Statistics*** tab, which displayed a histogram of events that could be drilled down on for more details. This time series display enabled us to add any statistics that we thought were relevant, and by looking at event balloons, we could visually correlate cluster-wide events to specified data measurements.

Filesystem Creation

We then switched to a 60-node WekaIO Matrix cluster running on Amazon Web Services to create filesystems (see Figure 4). The process for creating local, SSD-based filesystems or tiered, cloud-backed filesystems is the same. From the ***Filesystems*** tab in the GUI, a new filesystem could be created on a Weka filesystem group. ESG Lab navigated to the ***Local Filesystems*** row and clicked on ***Add Group***. We traversed the menu to add a group named *esg,* and a 3TB filesystem named *'esg1'*. After clicking the ***Create*** button, the filesystem was created – the whole process took about a minute.

ESG used a similar process to create a tiered filesystem using Amazon S3 storage, but first, we had to create the S3 object store for the tiered filesystem. ESG Lab used the WekaIO ***Object Stores*** tab on the left side of the GUI to create the cloud storage, which identified it to WekaIO. We then navigated to the ***Tiered Filesystems*** row in the ***Filesystems*** tab and clicked ***Add Group***. ESG traversed the menu to add a group named *archive* and a 10PB filesystem named *cold*. Using the Linux command line on a server in the cluster, ESG confirmed that WekaIO was in fact presenting the newly created tiered filesystem in the AWS cloud with 10PB of capacity.

Figure 4. Creating WekaIO Matrix Local and Tiered Filesystems

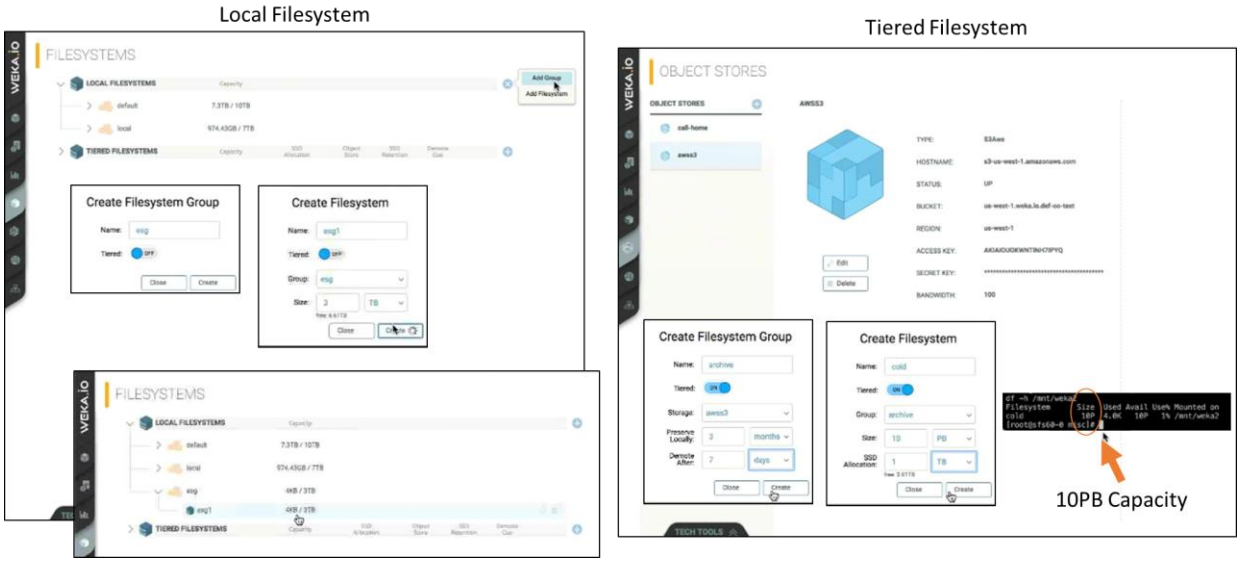

Source: Enterprise Strategy Group, 2017

Why This Matters

While the journey to amass data begins with a single byte, in our digital world those bytes collect at an exponential rate, quickly growing from gigabytes to terabytes to petabytes and beyond. ESG research shows that over a two-year period, the number of respondents reporting typical analytics jobs as being below 5 TB has dropped significantly by 25%, while 10% more are now using between 11-25 TB, and there is a significant increase in the number of organizations now exceeding 25 TB. [2] With this level of growth, a storage system must not only be easy to manage, but also easy to grow. Further, with the most precious resource for managing IT deployments being person-hours, a system that lets one person manage constantly growing data sets in the cloud makes efficient use of peoples' time.

ESG Lab validated that the WekaIO Matrix installation and management process was quick and easy. The browser-based and menu-driven user interface enabled us to deploy a new cluster in minutes, guided by an intuitive wizard. Cluster management was easy, with interactive diagrams that enabled us to point-and-click to visualize granular details of the cluster and correlate events to cluster performance. Creating a new filesystem on a cluster, whether on-premises or in the cloud, took just a minute with a few clicks.

[2] Source: ESG Research Report, *Enterprise Big Data, Business Intelligence, and Analytics Trends: Redux*, July 2016

Data Protection/Resiliency

WekaIO Matrix was designed to be a software solution that has the performance of an all-flash storage array with the economics and scalability of the cloud. MatrixDDP, implemented as part of the filesystem, implements patented (or patent-applied-for) error correction algorithms to optimize performance and minimize cost, by managing how data is distributed across a cluster for resiliency and fast recovery in case of a failure. The WekaIO Matrix uses a journaled filesystem, so writes that are acknowledged to applications are committed to non-volatile storage in SSDs. Journaling enables WekaIO Matrix to run on conventional hardware, with no requirement for an uninterruptible power supply (UPS) or NVRAM to protect data.

The WekaIO filesystem eliminates the need for the full filesystem check (fsck) common to traditional file-based storage. A full filesystem check increases failure recovery time and becomes intolerable as filesystems grow. During I/O, the WekaIO Matrix only opens a thin chunk of a file at a time, so there is only a small amount of data affected during a failure.

Algorithm and architecture techniques add to WekaIO Matrix's failure tolerance. WekaIO error correction algorithms provide either N+2 or N+4 resiliency. N+2 is the same resiliency level as triple mirroring or RAID 6. N+4 provides protection for up to four failures, and is well-suited to large cloud deployments or high capacity servers. WekaIO distributes data based on its concept of failure domains. A failure domain can be a physical server or a particular location, such as a public cloud availability zone. MatrixDDP manages data placement so there can never be the possibility of a multiple logical failures.

Local and cloud clustering improve I/O and rebuild performance since WekaIO maintains performance and availability as the cluster size increases. Simply put, the more nodes, the more data that can be spread across them. Furthermore, the larger the cluster, the faster the rebuild because more processors participate in the rebuild. In the event of multiple failures, Matrix DDP prioritizes rebuilds starting with the data that is least protected, to return the cluster to the next higher level of resiliency as soon as possible. Data tiering also helps WekaIO minimize rebuilds and improve resiliency. Because WekaIO Matrix protects data at a file level, the only data that needs to be rebuilt is the data stored on the failed component. There is no need to rebuild the data already tiered to the cloud, because it is safely stored there. ESG Lab viewed a data rebuild in process (see Figure 5); it took less than 15 seconds for a 60-node cluster to go from N+1 resiliency back to full N+2 data protection.

Figure 5. Failure Domain Rebuild

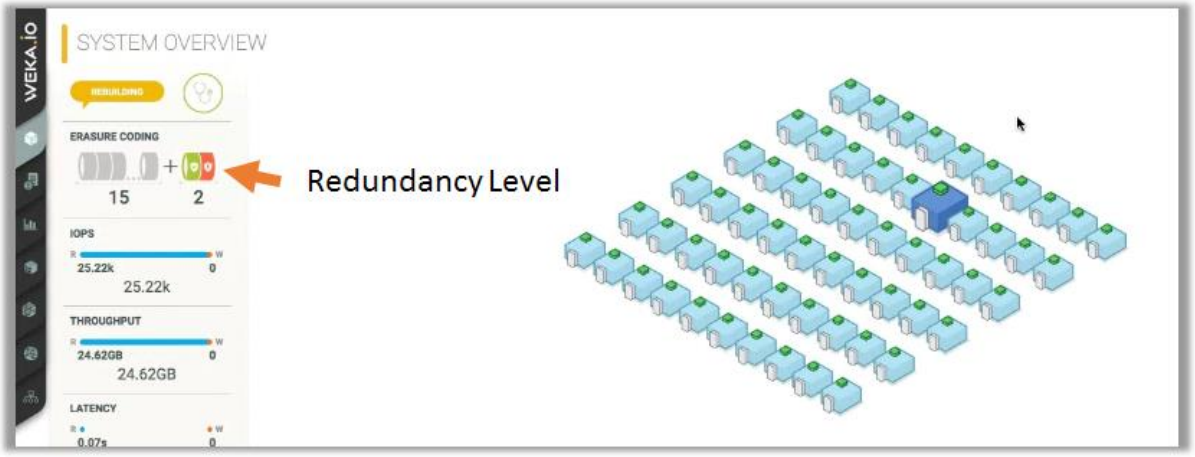

Source: Enterprise Strategy Group, 2017

Why This Matters

No matter where data resides, protecting it is always a top IT concern. ESG research shows that more than half (51%) of organizations report having a downtime tolerance of less than an hour for their "high priority" applications. We reviewed and validated WekaIO's resiliency architecture on premises, and witnessed a rebuild on a 60-node cluster taking less than 15 seconds to recover—going from a cluster failure to full availability.

Performance

ESG Lab audited performance tests run by WekaIO to gauge peak performance capabilities on-premises, as well as scalability with a public cloud (AWS) implementation.

ESG Lab used SPECSFS 2014 as a workload generator to create small block (4KB) workloads to measure I/O performance (IOPS) and large block (1MB) workloads to measure throughput (GB/s). ESG used a WekaIO Matrix hyperconverged configuration in which the workload generator clients were running on the same nodes as the WekaIO Matrix filesystem. The first phase consisted of testing a 10-node cluster on premises. ESG found results comparable to currently available, state-of-the-art, all-flash arrays (see Table 1).

Table 1. WekaIO Matrix Performance with Ten On-premises Servers

Workload	Ten-server Cluster Performance
4KB random read	983,638 IOPs; 0.34 ms latency
4KB random write	132,190 IOPS; 0.25 ms latency
1MB random read	23.8 GB/s
1MB random write	7.9 GB/s

Source: Enterprise Strategy Group, 2017

For testing in the AWS cloud, ESG sized the cluster at 30, 60, and 120 nodes, again in a hyperconverged configuration. Figure 6 highlights the linear performance scalability achieved by WekaIO for a 4KB random I/O size. Peak performance, achieved at the 120-node count, reached 3.76 million IOPS for reads and 399K IOPS for writes, while response times for both reads and writes remained under 1ms for all test cases.

Figure 6. WekaIO Matrix Performance Scalability - IOPS

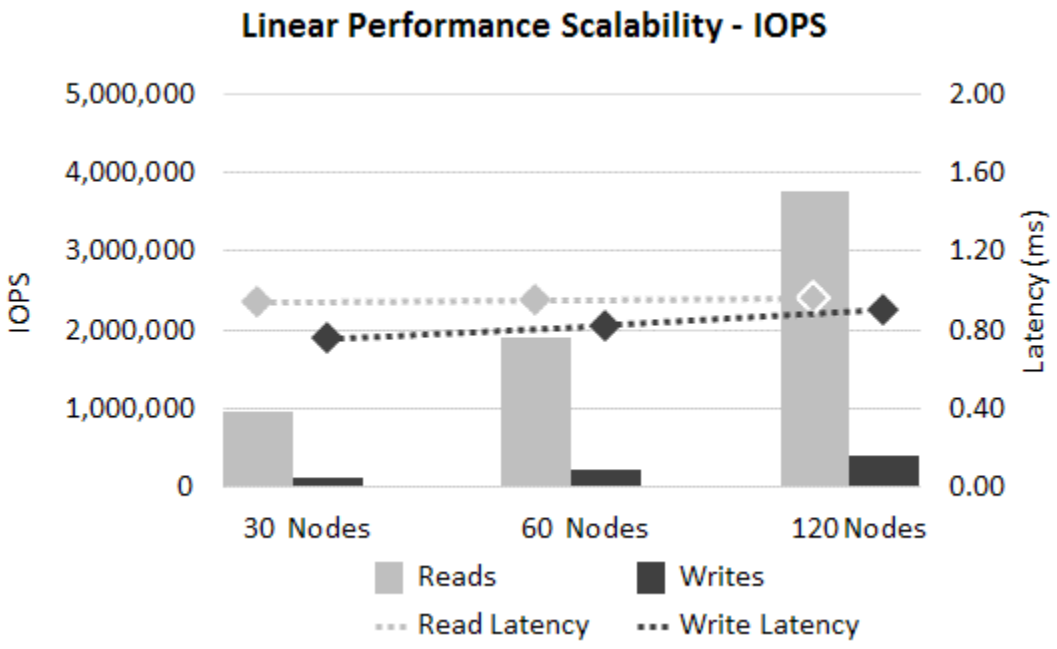

Source: Enterprise Strategy Group, 2017

Next, ESG Lab audited the throughput capabilities of WekaIO in AWS. 1MB random reads and writes were run and Figure 7 highlights the results. Again, we saw linear performance scalability as the test bed doubled in size, from 30 to 60 to 120 nodes. At the 120-node count, ESG witnessed 59.1 GB/s of throughput for reads, while writes yielded an impressive 34.7 GB/s.

Figure 7. WekaIO Matrix Performance Scalability - Throughput

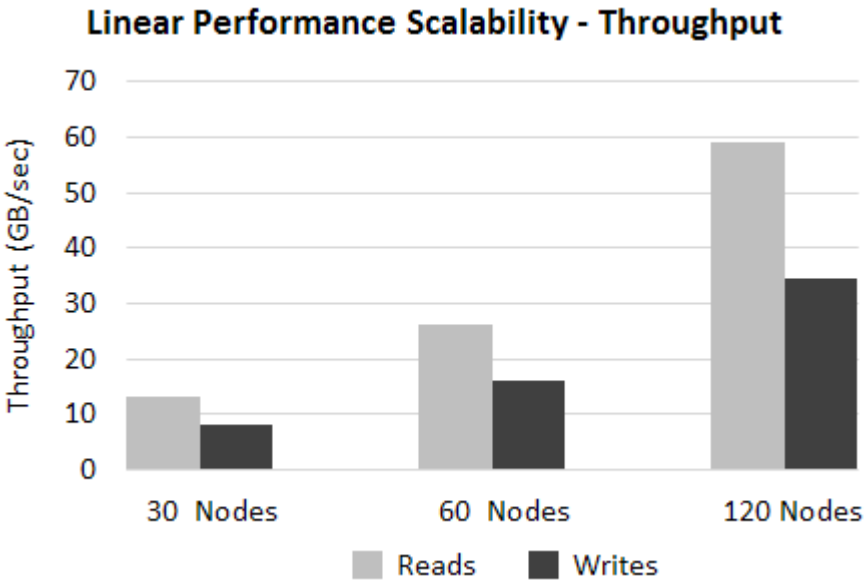

Source: Enterprise Strategy Group, 2017

Why This Matters

Though software-defined storage is gaining market traction by providing benefits such as deployment simplicity, agility, and cost-savings, some organizations continue to stick with their traditional storage implementations. One reason for this is a fear of poor performance. In fact, ESG research shows that 15% of organizations feel one of their biggest challenges or concerns if they were to adopt software-defined storage solutions would be poor performance.[3]

ESG Lab verified WekaIO Matrix's ability to enable a cluster of conventional servers to perform like an all-flash array, and to use a distributed filesystem and parallel processing to expand performance even further, whether on-premises or in the cloud. A local 10-server cluster delivered 984,000 4KB read IOPs, competitive with a standalone all flash array, while a 30-, 60-, and 120-server AWS cluster delivered predictable, linear performance scalability, with a peak of 3.8M read IOPs and nearly 60 GB/s of read throughput. As expected, latency measurements remained low throughout all testing, never exceeding 1ms.

[3] Source: ESG Research Report, *Software Defined Storage (SDS) Market Trends*, February 2017

The Bigger Truth

For IT administrators to meet the demands of constantly growing data sets, infrastructure modernization is inevitable. Software-defined storage technology is being adopted to simplify storage management, reduce costs, and provide greater agility for modern IT infrastructures, whether on-premises or in the cloud.

WekaIO Matrix is a software-defined storage solution that provides a distributed, virtual filesystem using off-the-shelf servers and SSDs to provide high performance file storage, and cloud techniques to scale performance and capacity on-premises or in the cloud. Flexible deployment options include hyperconverged, traditional client-server, NFS, or WekaIO clients, bare metal, virtualized, containers, and cloud.

ESG Lab validated the simplicity, manageability, and performance of WekaIO. The user interface for installation and management was well-organized and provided helpful automation that an IT administrator can customize. Cluster installation was quick and easy, taking less than ten minutes to be data-ready; creating a local or cloud-tiered filesystem took another minute and a few mouse clicks. The simplicity and speed to accomplish these tasks was impressive. For performance, we witnessed results that placed WekaIO in the same category as traditional all-flash storage arrays; WekaIO achieved millions of IOPS and GBs of throughput for common I/O sizes, while delivering linear scalability in a cloud deployment consisting up to 120 nodes.

With IT modernization in full swing, organizations are exploring and adopting new technologies that deliver faster deployment times, improved manageability, greater agility, better resiliency, improved scalability and performance, or cost reduction. To address the complexities that come with constantly growing data sets and their strict performance, recovery, and cost SLAs, software-defined storage is being rapidly adopted to replace traditional, complex, and costly storage solutions. ESG recommends WekaIO as a cost-optimized, software-defined storage solution that combines cloud-like scale, agility, and protection with on-premises, all-flash storage performance.