Technical Validation

WekaFS: Storage for Modern Exascale Workloads in the Enterprise

Enhancing Performance and Productivity with WekaIO

By Tony Palmer, Senior Validation Analyst

October 2020

This ESG Technical Validation was commissioned by Weka and is distributed under license from ESG.

Contents

ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

Introduction

This ESG Technical Validation documents hands-on testing of the WekaIO Parallel File System (WekaFS). We evaluated how WekaFS achieves extremely scalable and consistently high performance for AI/ML and technical computing workloads anywhere with enterprise-focused functionality and management. We examined Weka's ability to extend a file system namespace across flash and object storage, across on-premises data centers and the cloud, transparently to applications.

Background

ESG recently uncovered that more than 40% of organizations have at least 1 PB of primary storage capacity. It is not surprising, since 56% of these organizations anticipated their on-premises data growth to be at least 21%, with one in four expecting greater than 50% year-over-year growth over the following three years. With organizations facing this level of growth, software-defined storage (SDS) technologies are being considered as a viable way to grow storage capacity in a simple and flexible manner. In fact, 55% of surveyed organizations are already using SDS technology in production, with another 14% formally evaluating it (see Figure 1). Forty-six percent of those using or evaluating SDS technology already achieve or expect to achieve increased performance.[1]

Figure 1. Software-defined Storage Usage

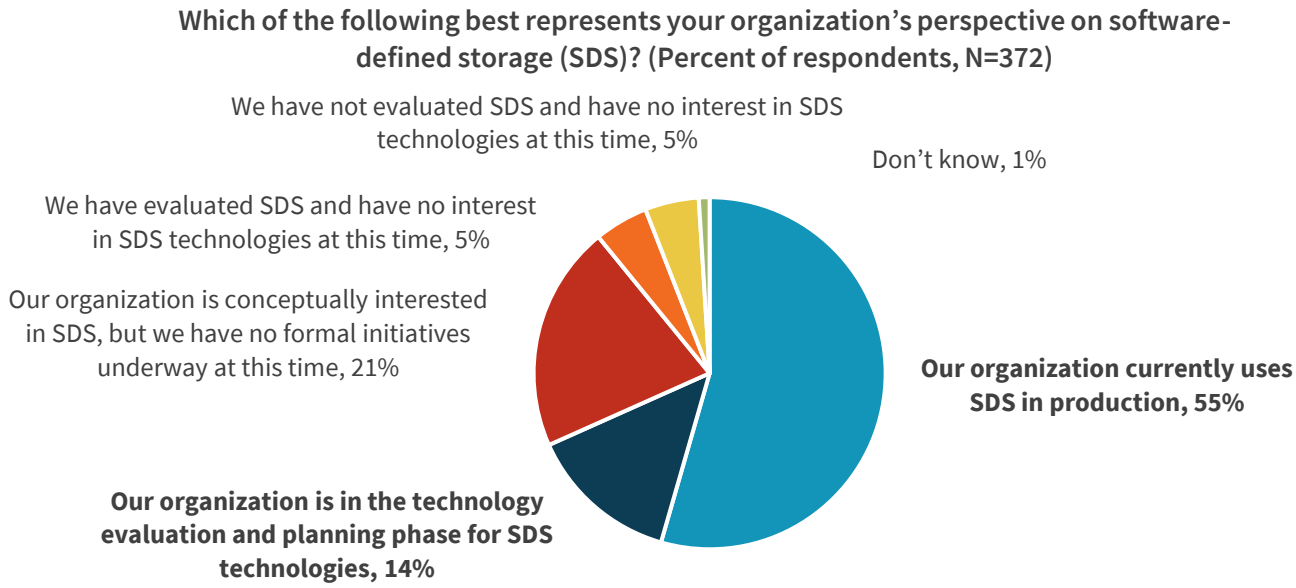

Source: Enterprise Strategy Group

Further, 64% of respondents to a separate survey indicated that spending on AI/ML in 2020 would increase over the prior year.[2] Larger amounts of both structured and unstructured data are generated and collected, and organizations have turned to SDS to build out their storage infrastructure more efficiently; one primary goal is to scale both processing power and capacity independently of each other to meet their current needs without overprovisioning either element. In addition to AI/ML, organizations are employing data-intensive applications such as those used for life sciences, financial modeling, data analytics, post-production editing, and the internet of things (IoT), which all utilize unstructured data. Organizations seek out SDS solutions that will provide fast and consistent storage performance for reads, writes, and metadata operations—which are key to application performance at scale—as their storage grows and supports the large amount of application processing required so they can process data and extract value without unnecessary delay.

[1] Source: ESG Master Survey Results, *2019 Data Storage Trends*, November 2019.
[2] Source: ESG Research Report, *2020 Technology Spending Intentions Survey*, February 2020.

The Solution: WekaFS

WekaFS is a POSIX-compliant file system that uses industry-standard NVMe-based servers on-premises and/or cloud instances to cost-effectively provide high-performance, high-capacity, and resilient storage at scale. WekaFS software uses industry-standard servers and network infrastructure to create a shared pool of file storage in the Weka global namespace. With Weka, storage can share the compute infrastructure with other applications, or it can run on dedicated servers. Weka's global namespace can run conventionally on bare metal servers; as a virtual machine (VM) managed by a hypervisor; as a Kubernetes containerized application; and in the cloud on Amazon Web Services (AWS). Every application server instance, physical or virtual, running Weka contributes storage resources and performance. Additionally, Weka can use an AWS S3 compatible object store to expand the namespace to accommodate exascale data in a highly economical way, on-premises or in the public cloud.

Weka is especially well-suited to the most performance-sensitive and resource-intensive applications—AI/ML, genomics and life sciences, financial analytics, media and entertainment, and manufacturing/engineering. Weka delivers flash storage performance using NVMe SSDs that are direct-attached to the application servers in a converged mode or installed on dedicated Weka storage servers. The software can take advantage of NVMe 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 NVMe SSDs to a server, by adding more servers with NVMe devices, or by allocating additional S3 capacity, and the performance can scale by allocating additional CPU cores or adding server instances to the global namespace. Weka's file storage can be consumed by applications in a conventional NAS topology, using NFS, SMB, a WekaFS native POSIX client. Weka also supports the NVIDIA GPUDirect Storage protocol and S3 clients via a gateway to access data on a cluster. In addition, clients and servers can reside together on the same hosts to take maximum advantage of server hardware and server virtualization.

Figure 2. WekaFS in a Production Environment

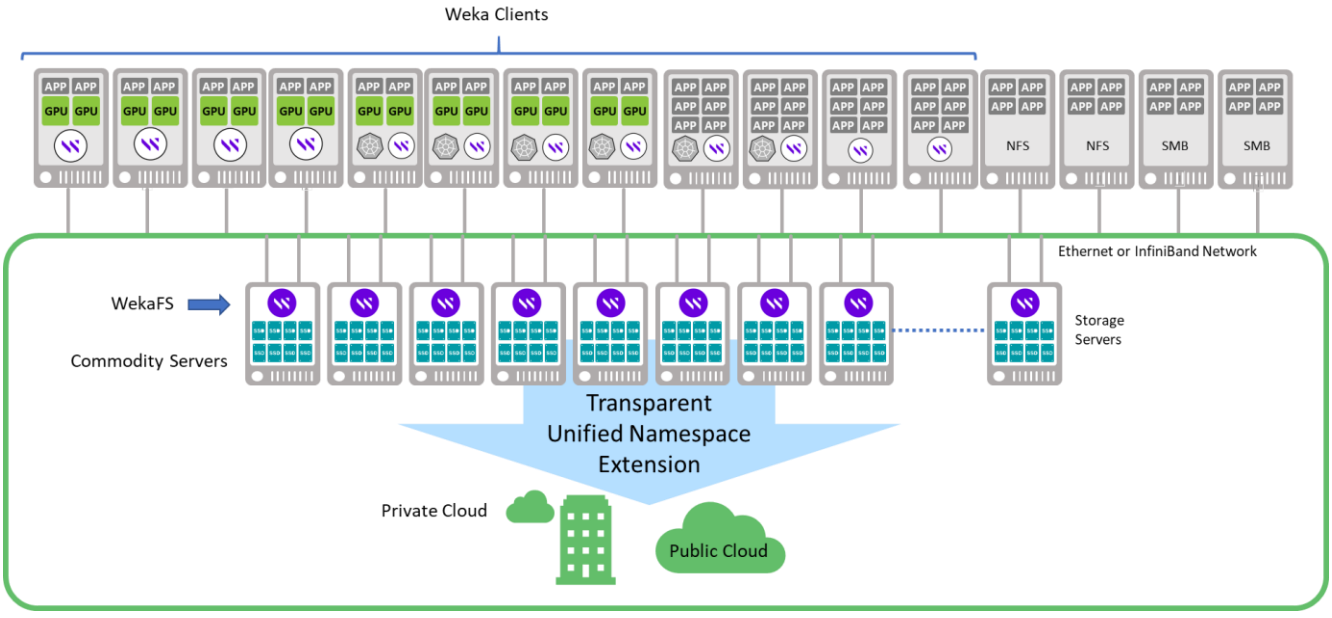

Source: Enterprise Strategy Group

Weka's distributed data protection is designed to run in a hybrid cloud environment and to deliver high resilience with minimal impact. WekaFS distributes data across failure domains, which can be physical host servers or cloud availability zones. The file system can employ either N+2, N+3, or N+4 data redundancy. Journaling is integral to WekaFS, so recovery from a server failure or from a complete system failure is fast, never requiring a full file system check (FSCK) process. Organizations can use object storage as a low-cost capacity storage tier by having WekaFS automatically move seldom-used data to an S3 object store in AWS and/or on-premises.

Organizations report that the volume of their primary data is massive and growing fast. More than half (57%) of respondent organizations report currently managing at least 1 PB of primary data storage capacity.[3] Protecting this volume of data is costly and time consuming using traditional methods, as organizations have been telling ESG for a long time: Managing high volumes of data copies and the need to reduce backup time were both cited as challenges with data protection processes and technologies by 26% of respondents to an ESG survey.[4]

Weka has built a storage solution designed for the enterprise. Weka provides data services like multiple sub-organizations on multiple file systems, point-in-time snapshots per file system, and writeable snapshots (clones) that are accessed through the namespace. Weka snap-to-object functionality allows snapshots to be saved to an S3 object storage bucket so that any other Weka cluster system—even from different locations and of differing size—can attach to that snap-to-object snapshot and keep running from that point in time. Snap-to-object snapshots are saved on the object storage in differential manner designed for efficiency, and snap-to-object can be used as a backup, a cloud-bursting mechanism, or for disaster recovery (DR).

GPU-accelerated computing has long been associated with high performance computing (HPC) and AI/ML; those workloads require massive I/O bandwidth to satisfy GPUs' ability to process data. Modern software libraries are opening GPU acceleration to workloads across the spectrum, from life sciences, to financial analysis, to media and entertainment and beyond, providing end-to-end data science and analytics pipelines run entirely on GPUs. Weka has developed a framework—Weka AI—to help organizations accelerate storage for GPUs using NVIDIA GPUDirect storage with its NVMe file storage. Weka AI is a framework composed of customizable reference architectures and software development kits, leveraging NVIDIA GPUs and WekaFS parallel file system software.

Figure 3. Weka with NVIDIA GPUDirect Storage (GDS) Protocol

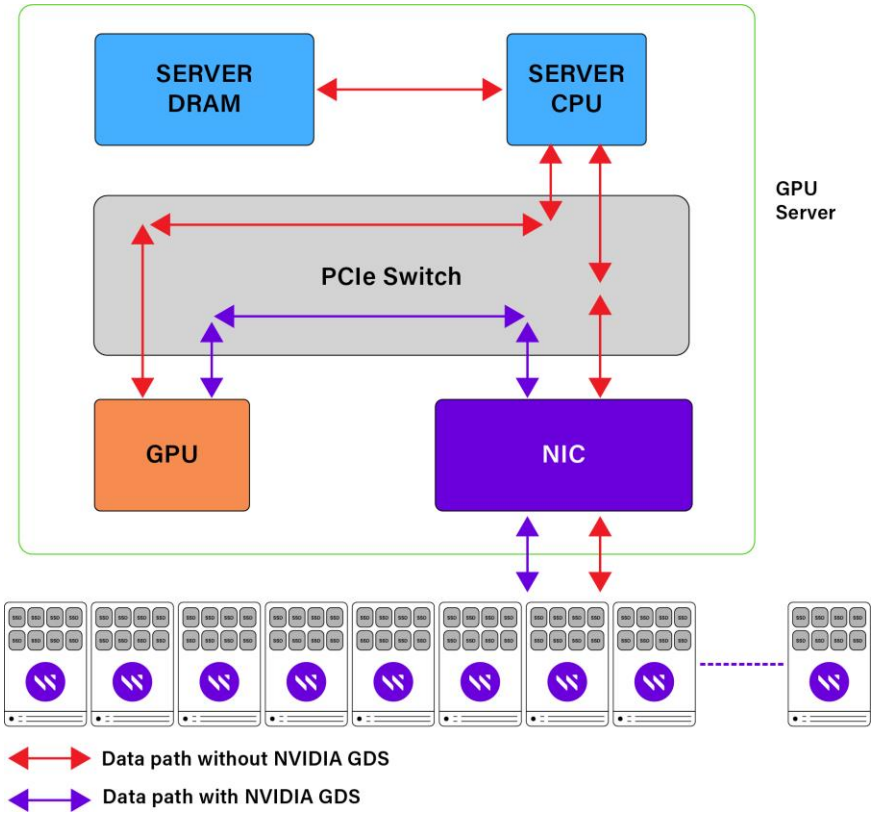

Source: Enterprise Strategy Group

Many of the features and functionality of WekaFS are summarized in Table 1; for more detailed descriptions and the most up to date list, contact Weka.

[3] Source: ESG Research Report, *Data Storage Trends in an Increasingly Hybrid Cloud World* , March 2020.
[4] Source: ESG Master Survey Results, *2018 Data Protection Landscape Survey*, November 2018.

Table 1. WekaFS Features

Weka Feature	Description	Weka Feature	Description
Snapshots	Instantaneous	Active Directory	Full integration
Snapshot to object (S3)	Local or cloud	Client protocols	NFS, SMB, WekaFS POSIX client, LDAP, S3 (via gateway)
Clones	Writable	Multiple file systems	1024 per namespace
Storage tiering	To S3 Storage, on-premises or in the cloud	Utilization management	Quotas, File system organizations
Dynamic scalability to thousands of nodes	Near-linear performance scaling	Networking	HA bonding, 100 gigabit Ethernet, HDR InfiniBand
Data security	AES-XTS-256 encryption in transit/at rest, SMB encryption, end-to-end checksum	Local file system caching	Provides coherency with local FS performance
Data protection	N+2 to N+4 (16+4)	Bandwidth	80 GB/sec per node, 160 GB/sec starter cluster on-premises
Hardware independence	Software only, no requirement for special hardware	IOPS (starter configuration)	5 million
Cloud-native	Private/public cloud and containers	GPU acceleration	NVIDIA GPUDirect Storage
Dynamic namespace	Zero-impact expansion and contraction	S3 access	Via gateway
Access security	POSIX, ACLs, Multi-organization authenticated mounts	Containers	Kubernetes CSI Plugin
Access management	POSIX and Windows ACL sync, authenticated mounts, Key management integration (KMIP)	Write cache	Coherent
Exclusive Mount	Yes		

Source: Enterprise Strategy Group

ESG Technical Validation

ESG tested the Weka storage solution on WekaFS-enabled systems on-premises and on WekaFS application and storage deployments in AWS. ESG validated Weka's performance as compared to the market using three rigorous industry-standard benchmarks. Time to deploy and extend a file system to AWS, deployment and management simplicity, and data durability and recoverability were also tested and examined.

Performance

ESG began by auditing performance results from three industry-standard benchmarks run by Weka to gauge complex workload performance capabilities for a variety of performance-critical workloads. First, we examined Weka's published SPEC SFS 2014 results. Weka earned the top spot for multiple workloads. The results are summarized in Figure 4.

Figure 4. WekaFS SPEC SFS 2014 Performance Compared to Other Solutions

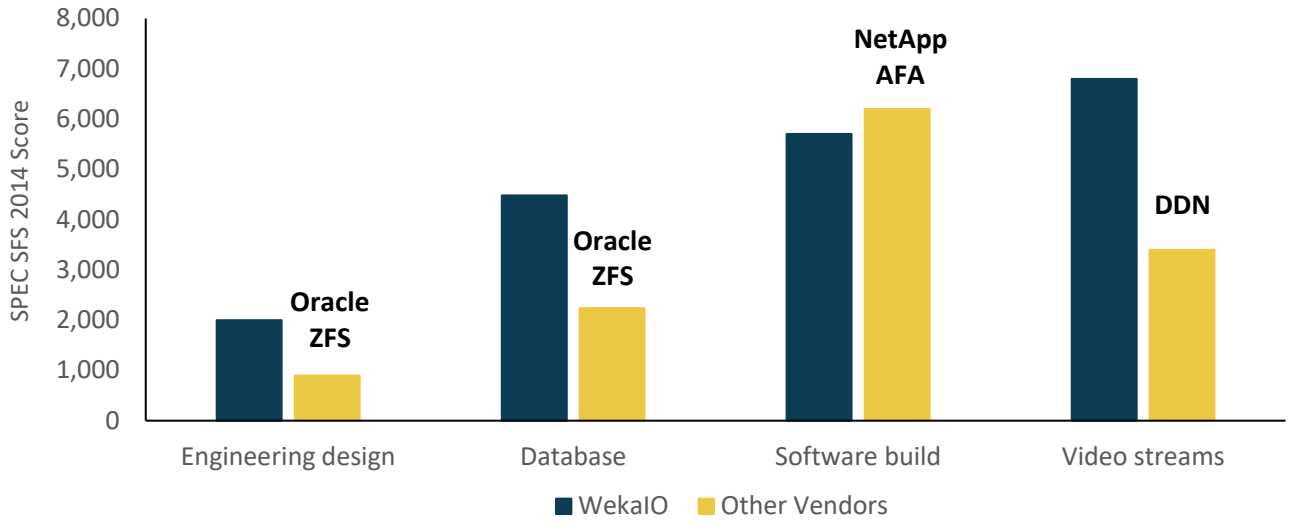

Source: Enterprise Strategy Group

Of particular interest is the software build benchmark. The software build benchmark is a file system benchmark that mimics the behavior of software builds and traces. 70% of the operations performed by the benchmark are a stat call returning file attributes (metadata). As such, the software build workload has a high amount of small block reads with interspersed larger block file I/O. While the overall number of concurrent builds was a bit higher than Weka, response times tell a more complete story.

Figure 5. SPEC SFS 2014 Software Build Response Time Curve

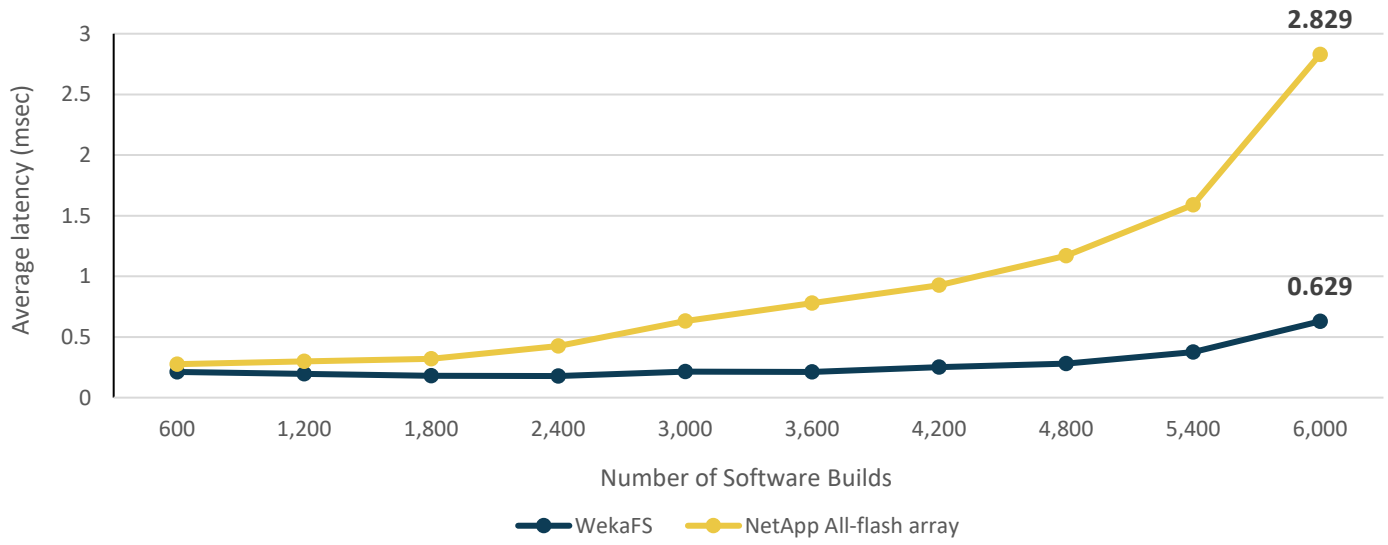

Source: Enterprise Strategy Group

As seen in Figure 5, response time stayed low as the number of concurrent software builds was increased from 570 to 5,700 with an overall response time for the entire test run of 480 µs or .48 ms. The All-flash array tested by NetApp, scaled to 6,200 concurrent builds, but with 4.5x the response time of Weka. Weka posted the highest performance in three of the four categories we examined, and as seen in Table 2, performance was double the next closest result for three out of the four tests, with lower overall response time across the board.

Table 2. WekaFS SPEC SFS 2014 Performance Across Multiple Workloads

Benchmark	Top Scoring Solution	SPEC SFS 2014 Score	Overall Response Time (ms)	Number Two Solution	SPEC SFS 2014 Score	Overall Response Time (ms)
Engineering Design	WekaFS	2,000	0.26	Oracle	900	0.61
Database	WekaFS	4,480	0.34	Oracle	2,240	0.78
Software Build	WekaFS	5,700	0.48	NetApp	6,200	0.83
Video Streams	WekaFS	6,800	1.56	DDN	3,400	50.07

Source: Enterprise Strategy Group

Next, ESG analyzed the results of the IO-500 benchmark and compared WekaFS performance on AWS to other file system solutions. The IO-500 benchmark suite was designed to enable comparison of high-performance storage systems with full disclosure of how the tests were performed and a methodology and ranking system designed to maximize credibility and usefulness. IO-500 scoring is derived from the square root of the product of metadata IOPS and throughput.

Figure 6. WekaFS IO-500 Performance on AWS Compared to Other File System Solutions

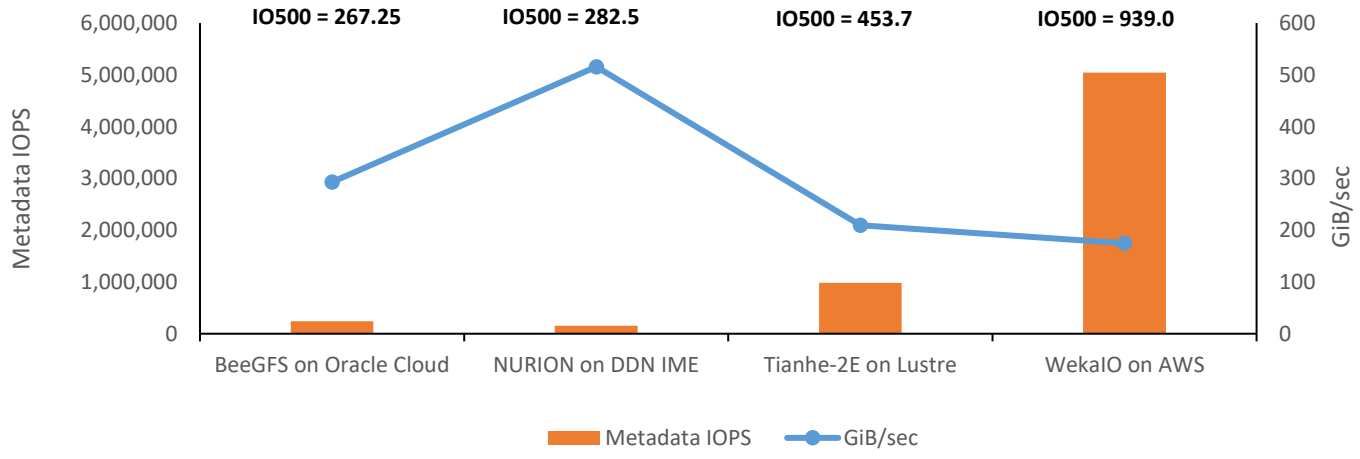

Source: Enterprise Strategy Group

As seen in Figure 6, Weka's IO-500 results on AWS are significantly higher than other file system solutions whether on-premises or running in the cloud. It's important to note the throughput and metadata IOPS are weighted equally in this benchmark, so a solution will be scored fairly. It's particularly impressive that Weka supported 5x the metadata IOPS of the next highest file system solution while sustaining 175 Gibibytes (GiB)per second of throughput. Note: IO 500 reports throughput in Gibibytes/sec. Gibibytes (GiB) differ from Gigabytes (GB)in that they are calculated in multiples of 1,024, rather than 1,000. 175 *Gibibytes*/sec is equal to 187.9 *Gigabytes*/sec. It's important to note that while the NURION on DDN IME bandwidth was 515 GiB/sec, that result was achieved with 2048 clients—about .25 GiB/sec per client. Weka's 175 GiB/sec was achieved with just 345 clients—about .5 GiB per client.

Next, ESG audited the results of the STAC-M3 benchmark. STAC-M3 is a benchmark published by the Securities Technology Analysis Center (STAC), with the goal of helping end-user firms relate the performance of new technologies to that of their existing systems. The STAC Benchmark Council is an organization of leading financial institutions and technology vendors that specifies standard ways to assess technologies used in finance.

STAC-M3 is a set of industry-standard enterprise tick-analytics benchmarks for database stacks that manage large time series of market data—aka "tick data." We examined results from the baseline benchmark suite (Antuco) and the optional

scaling suite (Kanaga). STAC-M3 focuses on the complete system under test—compute, storage, and networking elements tested together. The Antuco benchmark suite simulates gathering tick data from 12 specific (not consecutive) days from one year's worth of tick data. Clients are scaled from one to 50 to 100. Searches are performed concurrently. The Kanaga benchmark suite includes multi-year high bid analytics that reads terabytes of data to answer a query, placing significant load on storage I/O, while other tests such as Theoretical P&L in the Antuco suite are computationally intense with less impact from the storage system.

This multidimensional model incorporating both the Antuco and Kanaga benchmark suites highlights the limits of solutions while articulating where another solution starts to make sense for the anticipated workloads. It also helps users understand the resources—compute, storage I/O, and networking—required to reach the performance demands for an organization.

Figure 7. WekaFS STAC-M3 Tick Data Query Response Time (Shorter Is Better)

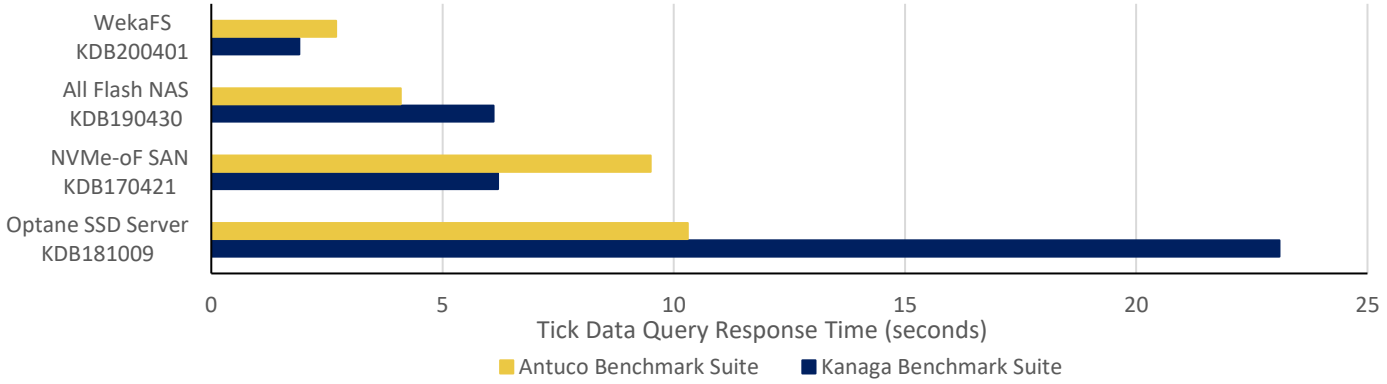

Source: Enterprise Strategy Group

The systems compared include a server with internal SSDs, an NVME over Fabric SAN, and a high-end all-flash NAS array. As seen Figure 7, Weka posted the shortest query response time in both benchmarks by a significant margin.

Finally, ESG tested performance and scaling of Weka software in AWS. We deployed an eight-node WekaFS cluster on AWS i3en.24xlarge all-flash instances. We used the fio load generation tool to test IOPS, throughput, and latency. Each test was run for five minutes and repeated three times and results were averaged. After we completed each series of tests, we added two nodes to the cluster and tested again. The results are summarized in the next three charts. First, we'll look at throughput. Fio was configured to drive 1MB I/O with each of the 16 clients running 32 threads. As shown in Figure 8, throughput scales linearly as nodes are added to the cluster.

Figure 8. WekaFS Linear Performance Scaling on AWS—Throughput

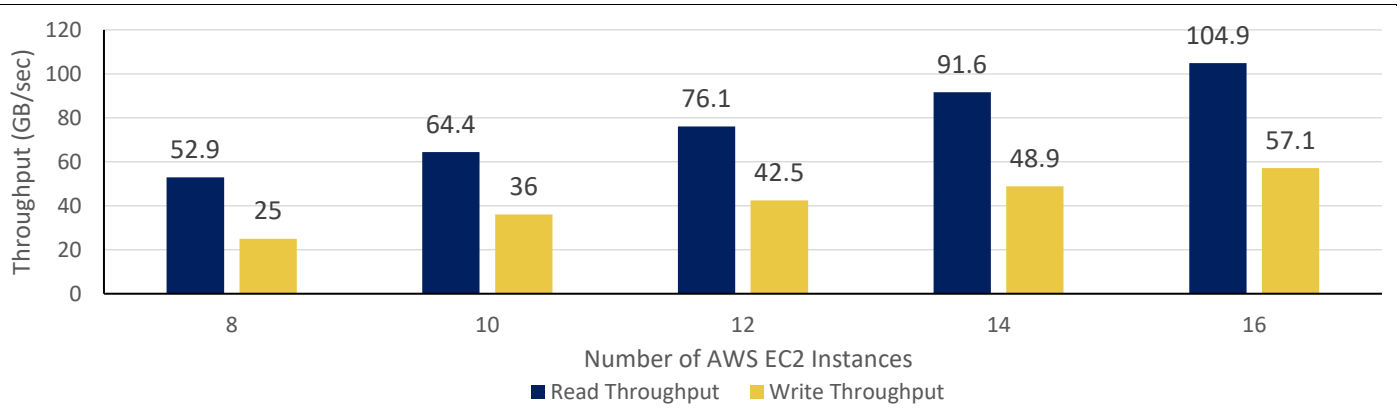

Source: Enterprise Strategy Group

Next, we'll examine IOPS. Fio was configured to drive 4KB I/O with each of the 16 clients running 192 threads. As seen in Figure 9, Read IOPS scaled with near-perfect linearity, while write IOPS increased consistently as nodes were added, adding an average of 73,000 IOPS as each pair of nodes were added. During these tests, latency averaged 748 microseconds across all IOPS tests for all cluster sizes.

Figure 9. WekaIO Performance Scaling on AWS—IOPS

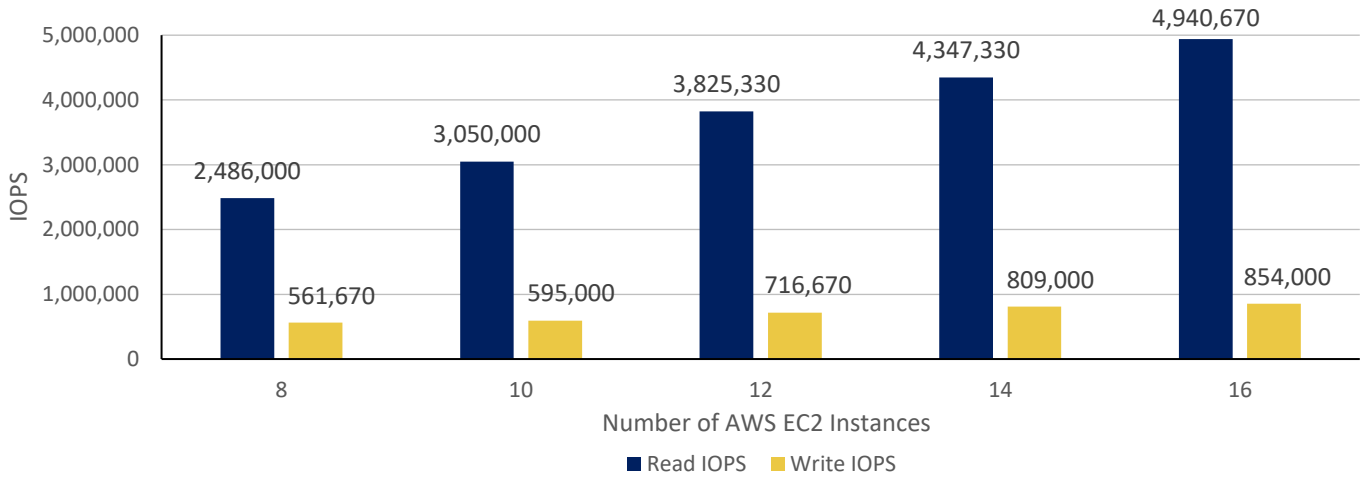

Source: Enterprise Strategy Group

Finally, we looked at latency. In this test, Fio was configured to drive 4KB I/O with each of the 16 clients running just one thread. As Figure 10 shows, latency stayed steady as nodes were added to the cluster, never exceeding 250 µsec.

Figure 10. WekaIO Performance on AWS—Latency

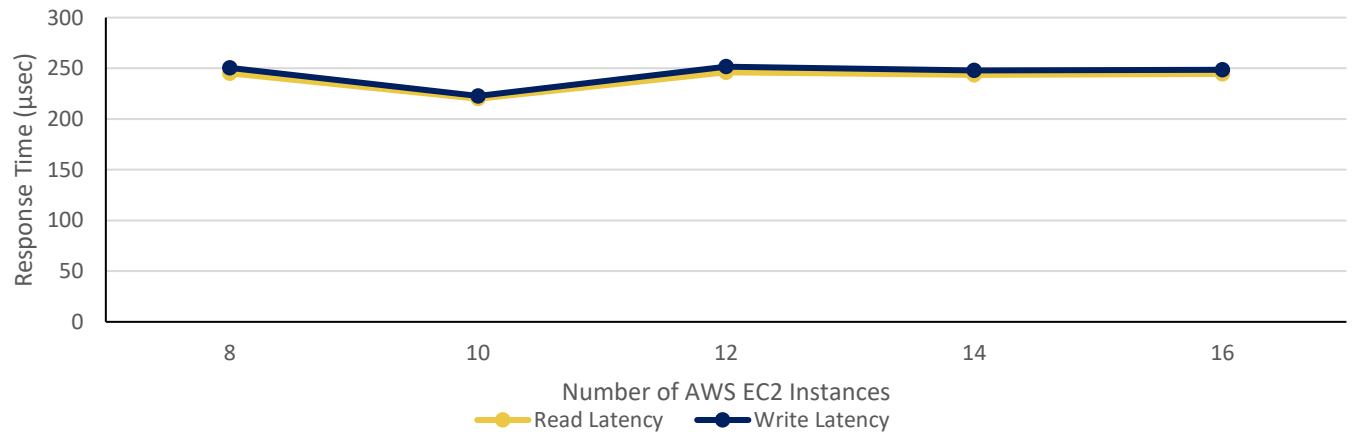

Source: Enterprise Strategy Group

It's important to note that the latency WekaFS can achieve in on-premises environments running on commodity hardware is considerably lower. ESG observed average response times of 450 µsec under similar IOPS load as seen in Figure 9 and less than 100 µsec in single-threaded tests.

Why This Matters

Though software-defined storage has historically provided benefits such as deployment simplicity, agility, and cost savings, organizations running modern and traditional HPC workloads need performance, too. It's not surprising that 71% of organizations are already using or plan to use NVMe-based solid-state storage technology, and more than half of organizations say they are using it to improve performance for new (56%) or existing (55%) applications.[5]

ESG verified WekaIO's ability to outperform all-flash storage arrays and to use their distributed, parallel file system to expand performance even further. WekaFS on AWS delivered nearly 5.8 million 4KB IOPS with sub-millisecond response times and more than 160 GB/sec of throughput. An on-premises cluster delivered the same performance with response times averaging under 500 µsec. Weka has posted top results in multiple industry-standard benchmarks including SPEC SFS 2014, emulating multiple application workloads, and STAC-M3 running intensive financial analysis. In the IO-500 benchmark, Weka posted the top result for solutions running on a file system.

Scalability, Flexibility, and Availability

The next phase of testing included extending a Weka file system using object storage in AWS, creating a snapshot, uploading it to the cloud, and using the snap to create a duplicate file system. We also walked through the process of scaling a cluster up and down non-disruptively, while workloads were running.

Extending a File System with Object Storage

We began with an eight-node on-premises cluster with 76 TB of capacity. We logged into the GUI and created a 10-terabyte file system, then mounted it from a Linux client and created a 10-megabyte file with the *dd* utility. Next, we attached an object store to it by selecting the file system in the GUI, clicking *Configure*, then clicking *Object Store.* We selected from the pull-down list, then configured the file system, as seen in Figure 11. We then expanded the file system on the fly to 5 petabytes without any need to stop I/O or take the file system offline.

Figure 11. WekaIO—Extending a File System with AWS S3 Storage

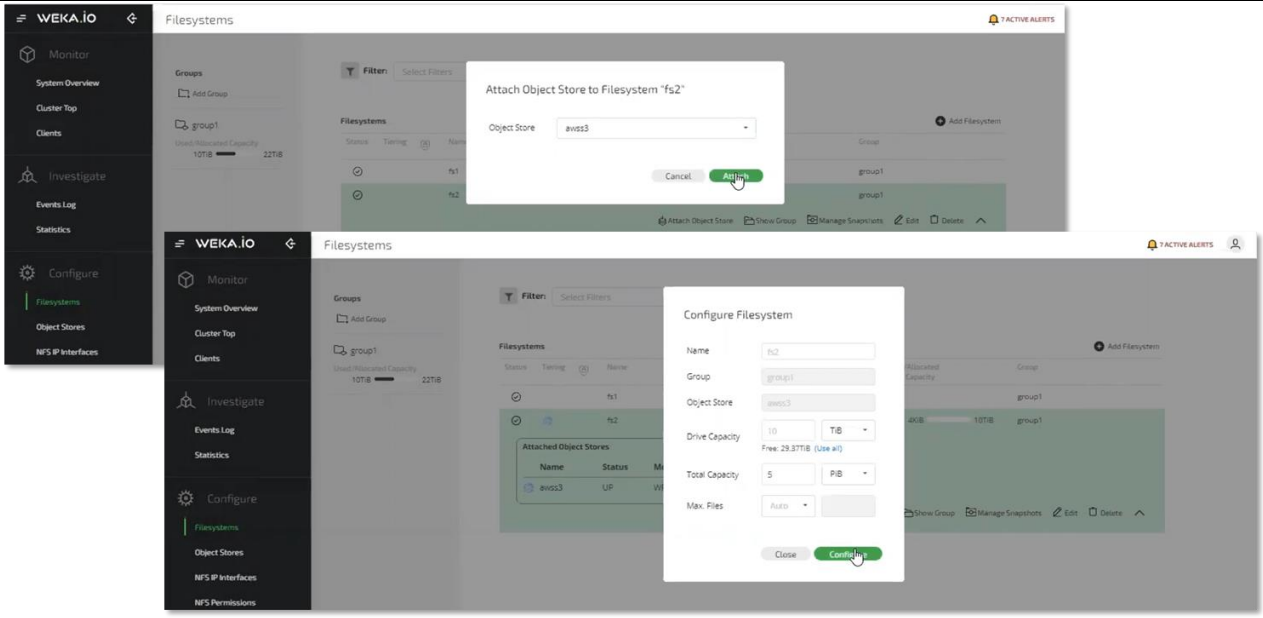

Source: Enterprise Strategy Group

[5] Source: ESG Research Report, *Data Storage Trends in an Increasingly Hybrid Cloud World*, March 2020.

We then verified that the Linux client could see the entire 5-petabyte file system. The entire process, including mounting from the client and creating the file, took less than a minute. Next, we expanded the file system to 10 petabytes with one click. The Linux client was able to see the additional capacity immediately.

Snap to Cloud

At this point, we took a snapshot of the file system. With one click, we uploaded the snap to object storage. When you snap to the cloud, data and metadata are moved together. At this point, any Weka system with the correct security credentials can connect to the same bucket and mount the file system using the *Object Store Locator*, shown in Figure 12.

Figure 12. WekaIO—Snapping a File System to Object Storage in AWS

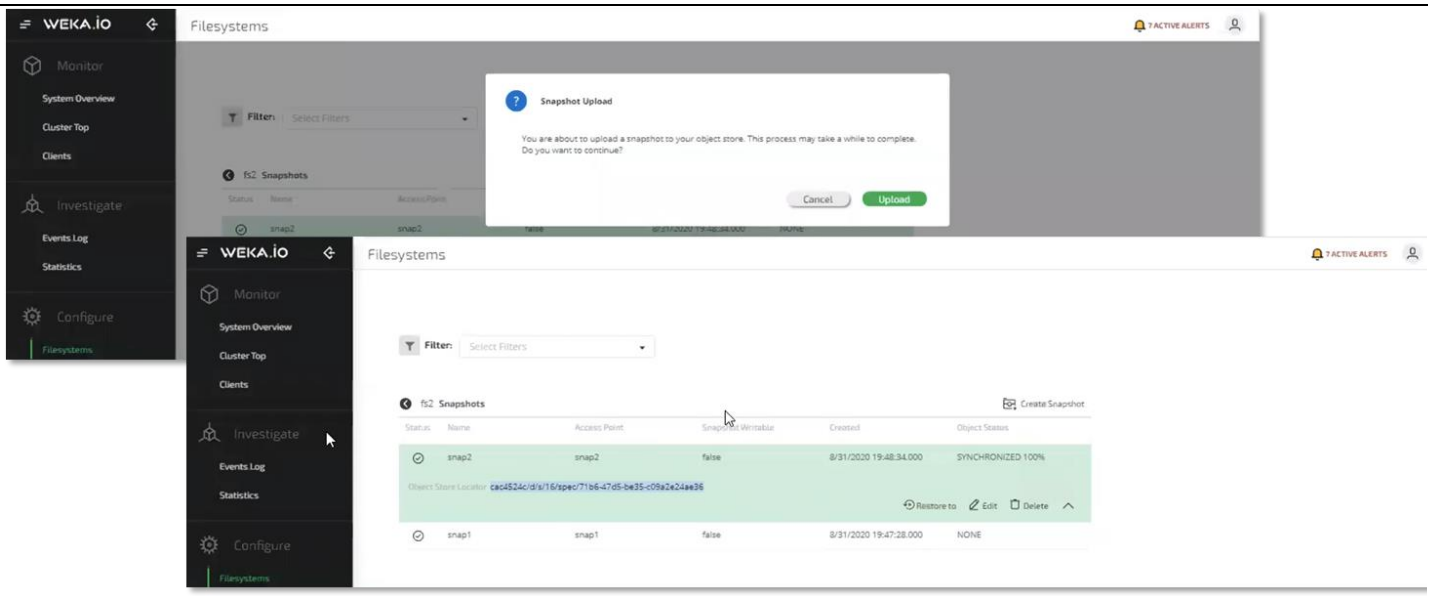

Source: Enterprise Strategy Group

Burst to Cloud

Next, we connected to a Weka cluster in AWS and created a new file system. We selected the *From Uploaded Snapshot* option and pasted in the *Object Store Locator.*

Figure 13. WekaIO—Snapping a File System to Object Storage in AWS

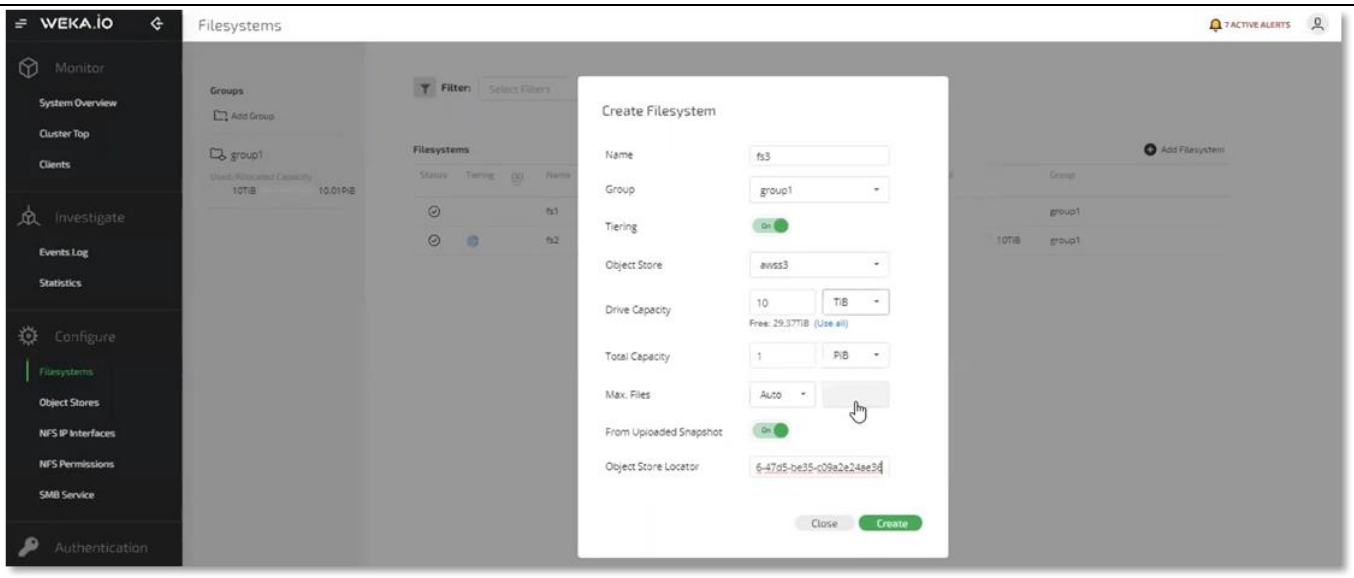

Source: Enterprise Strategy Group

It's important to note that the new file system does not need to be the same size as the original; it just needs to be large enough to hold the meta data that exists in the uploaded snapshot.

We mounted the new file system with a Linux client and verified that the 10MB file we created in the on-premises file system was there.

Encryption

ESG also walked through configuring encryption. Encryption is performed at the client and data is encrypted in flight and at rest. Because of this, there is no need for self-encrypting drives. It only took a few seconds to connect a Hashicorp key management system (KMS) and create a new file system with encryption. We tested performance with and without encryption on two file systems on the same cluster, using the same client. We used fio to run a 100% random 4KB workload against two identical Weka clusters. Without encryption, we saw approximately 330,000 IOPS at 96 microseconds, and with encryption, we saw almost identical results: 330,000 IOPS at 87 microseconds.

Scaling

In this test, we expanded a cluster from eight nodes to nine, growing from 76 TiB to 88 TiB while the cluster was running an 11 gigabyte per second workload. The entire process was completed with just a few clicks and took less than three minutes from start to finish. The process was completely transparent to clients and non-disruptive to applications. Shrinking the cluster was even easier and took just two minutes to complete, again, non-disruptively. Neither adding nor evicting a node had any effect on the workload.

Availability

WekaIO was designed to be a software solution that can provide better performance than an all-flash storage array with the economics and scalability of the cloud. Weka's distributed data protection, implemented as part of the file system, applies patented 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. WekaFS is a journaled file system, so writes that are acknowledged to applications once the data is safely stored on the NVMe devices. Journaling enables WekaFS to run on conventional hardware, with no requirement for an uninterruptible power supply (UPS) or NVRAM to protect data.

The Weka file system eliminates the need for the full file system check (fsck) common to traditional file-based storage. A full file system check increases failure recovery time and has become intolerable as HPC file systems have grown to hundreds of terabytes and even petabytes. During I/O, WekaFS writes data across several drives in small chunks so that during a failure only a small amount of data is affected, significantly reducing recovery time—to seconds in ESG's testing.

Maximizing Performance and Availability

Weka keeps as much data as possible on fast, primary storage to maximize performance. When capacity reaches the high-water mark (95%), Weka expires the least-used data from flash to free up capacity until the system reaches the low-water mark (90%). If object storage capacity is available, clients will never run out of space. To prevent a performance hit, Weka uploads data from flash continuously in the background, so that when the high-water mark is reached, all Weka has to do is expire data from the flash tier without having to upload the data at that time. If Weka has no data to expire (all 95% is active), a feature called back pressure is invoked and I/O to clients is throttled to the max throughput of the upload to object storage under the assumption that organizations would rather suffer a temporary performance hit than experience an out of space error. An event is triggered that alerts administrators that local capacity is at the threshold and needs to be expanded. Admins can move files to the flash tier on demand and pull back to the flash tier on demand as well using Weka APIs.

Algorithm and architecture techniques add to Weka's failure tolerance. WekaFS error correction algorithms provide either N+2, N+3, or N+4 resiliency. N+2 is the same resiliency level as triple mirroring without the 300% capacity requirement or

RAID 6, but across multiple servers and not just within a single server. N+4 provides protection for up to four failures and is well-suited to large cloud deployments or high-capacity servers. WekaFS distributes data based on failure domains; a failure domain can be a physical server, a data center rack, or a particular location, such as a public cloud availability zone. WekaFS manages data placement so there can never be the possibility of multiple logical failures.

Figure 14. Failure Domain Rebuild

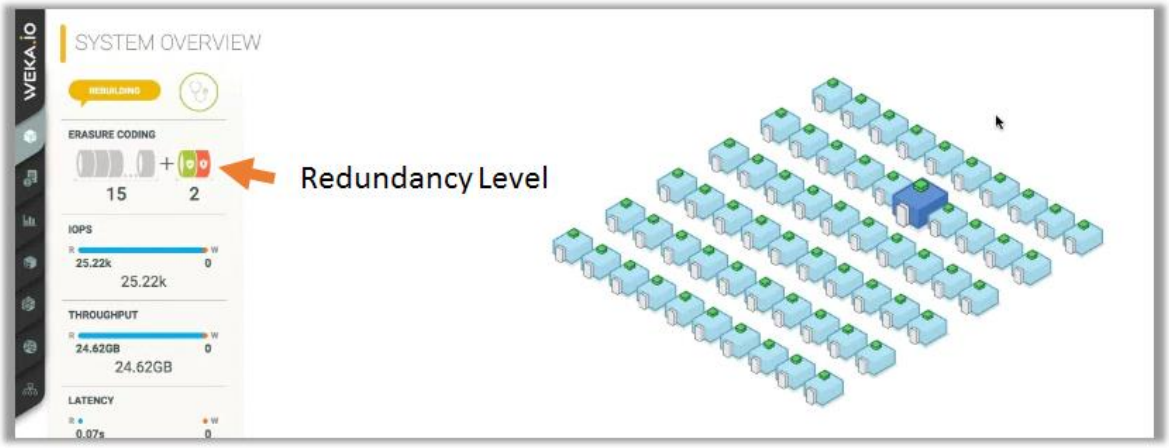

Source: Enterprise Strategy Group

WekaFS clustering— both on-premises and in the cloud improves I/O and rebuild performance as the cluster size increases. Simply put, the more nodes, the more data that can be safely spread across them. Furthermore, the larger the cluster, the faster the rebuild because more processors and NVMe drives participate in it. In the event of multiple failures, WekaFS 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 Weka minimize rebuilds and improve resiliency. Because WekaFS protects data at a file level, the only data that needs to be rebuilt is the data stored on the failed component and not simply rebuild blocks that might not even be in use. There is no need to rebuild the data already tiered to an object store or the cloud because it is already safely stored there. ESG observed a data rebuild in process (see Figure 14); it took less than 15 seconds for a 60-node cluster to go from N+1 resiliency back to full N+2 data protection.

Why This Matters

While the journey to amass data begins with a single byte, those bytes collect at an exponential rate in our digital world, quickly growing from terabytes to petabytes, soon to reach exabytes and beyond. ESG research shows that more than four in ten organizations have at least one petabyte of primary storage capacity and data is growing at a rapid pace; 56% of organizations report more than 20% on-premises storage capacity annual growth, with 25% seeing more than 50% annual growth.[6] With this level of growth, a storage system must not only be easy to manage, but also easy to scale. Further, with the most precious resource for managing IT deployments being person-hours, businesses need a system that enables an admin to manage petabyte-sized data sets on-premises or in the cloud.

ESG validated that the WekaIO 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, expanding, and shrinking a file system on a cluster, whether on-premises or in the cloud, took less than a minute with a few clicks. WekaIO made snapping a file system to the cloud fast and easy as well. Once uploaded, reconstituting that snap into a live file system on a new cluster was even faster.

No matter where data resides, protecting it is always a top IT concern. We reviewed and validated WekaIO's resiliency architecture on-premises and witnessed 300 TiB cluster rebuild from a double node failure in less than 15 seconds.

[6] Source: ESG Master Survey Results, *2019 Data Storage Trends*, November 2019.

WekaFS in the Real World

Weka reports that eight of the top 64 Fortune companies run Weka for their AI initiative to make them more competitive. ESG spoke with several customers to learn how Weka has impacted their businesses. Interviewees included executives and IT administrators across a spectrum of industries. An international data services integrator switched to Weka when their existing scale-out NAS cluster reached its capacity and performance limits. Their 22 PB environment has since grown to 60 PB and their storage admin requirements have shrunk by 80%. They told ESG that Weka has enabled them to win clients they could not even talk to before.

Grew clients and revenue; reduced admin costs by 80%.

Accelerated workloads; reduced admin costs by 90%.

A US-based biotechnology company focused on research, development, and delivery of treatments globally has selected Weka to replace a 20 PB tiered NAS environment. Their environment is growing at nearly 2% per month and a team of 10 storage admins is required to manage the environment. Their HPC team lead told ESG that while the system they are replacing is a fantastic NAS platform, it is not up to the demands of a modern HPC environment. Snap-to-cloud and burst-to-cloud capabilities are differentiators for them. Their pilot deployment is two petabytes—1 PB on premises and 1 PB in AWS—and it is being managed by one admin part-time. They are confident that they will not need more as the system grows to tens of petabytes.

We spoke with the CIO of a genomics company based in the US that is working to leverage applied genetic research to develop smart diagnostics and targeted therapeutics to improve disease outcomes. They also switched to Weka from an NFS-based system that they describe as "not well suited" to HPC at the level they need. Their total environment is 10 petabytes—7 PB active and 3 PB cold archive that gets rehydrated frequently. Their Weka deployment is 1 PB of on-premises active storage today. Object integration with the cloud is key for them, so they can move workflows instead of data, according to the CIO. Weka is used today for concurrent workloads that need extreme throughput. Again, they have a fraction of a full-time employee managing the system.

Four petabytes, 75% object storage with NVMe performance.

Grew revenue; reduced admin costs by 66%; Four petabytes, 80% object storage with NVMe performance.

Finally, we spoke with the lead architect for an AI/ML-based product being produced by a high-tech manufacturer in the US. The 30-petabyte environment comprised a 6 PB on-premises WekaFS cluster, extended with 24 PB in object store. The environment was used to develop a software system to solve extremely complex problems at independent remote nodes in real time. Again, the administrative overhead is negligible, with three admins managing the entire environment. "Weka lets us design the cluster the way it works best for us, which makes for a much more efficient ML environment," said the architect. Further, they said they appreciated not being forced to adjust their processes and architecture to adapt to platform constraints.

There were a few common threads in our discussions. Customers told ESG that the value of using a single global namespace makes object stores more manageable, which drives lower administration requirements. Weka's response to issues and high-touch support were mentioned by all the customers we spoke with; working closely with their customers and other vendors made the troubleshooting and resolution of challenges fast and painless. Also, customers typically refresh infrastructure software at the same time as servers for convenience, traditionally on a three-year cadence. Customers often choose to purchase Weka via a five-year license (at a 40% savings over three-year licenses) because it will run on any industry-standard server and so they can optimize hardware and Weka can take advantage of it immediately.

The Bigger Truth

Organizations are continuing to generate and store exceptionally large amounts of unstructured data. ESG uncovered that more than half of organizations expect their on-premises data to grow by at least 21% annually over the next three years.[7] With the increasing adoption and use of data-intensive applications—life sciences, financial analysis, artificial intelligence, and machine learning, to name just a few—organizations require a solution that can efficiently store and process data with consistently high performance. The solution should also scale in a manner that enables organizations to increase processing power and capacity independently, on-premise or in the cloud.

Weka is a software-defined storage solution that provides a distributed file system on-premises or in the cloud to cost-effectively provide extremely high-performance file storage using hyperscale techniques to provide high-performance, high-capacity, and resilient storage at scale. WekaFS is especially well-suited to the most performance-sensitive and resource-intensive applications—AI/ML, life sciences (genomics and cryogenic electron microscopy (Cryo-EM) analysis), financial analytics, media and entertainment, and manufacturing/engineering. Weka delivers flash storage performance using a tiered methodology that enables organizations to extend their unified namespace to private and public clouds.

ESG validated the simplicity, manageability, and performance of Weka's storage software. 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 file system took another minute and a few mouse clicks. The simplicity and speed of accomplishing these tasks was impressive. For performance, we witnessed results that placed Weka head and shoulders above traditional all-flash storage arrays; WekaFS achieved millions of IOPS and hundreds of GBs of throughput for common HPC workloads, while delivering linear scalability in an AWS cloud deployment. Providing such high throughput and IOPS at the same time at extremely low latency is a differentiator and enables organizations to use the Weka file system in place of servers with local drives while improving the overall time to completion of customers' workloads, saving time and money.

With IT modernization in full swing, organizations are searching for new technologies that deliver faster deployment times, improved manageability, greater agility, better resiliency, and improved scalability and performance. To address the complexities that come with modern, constantly growing, data sets and their strict performance, recovery, and cost SLAs, enterprises are rapidly adopting software-defined storage to replace traditional, complex, and costly storage solutions. ESG interviewed four organizations with deployments up to 62 petabytes that view Weka as providing them a competitive edge for AI/ML and other demanding workloads.

Extreme performance requirements, once considered solidly in the HPC domain, are becoming a common challenge for enterprises, especially with the increased use of GPUs with AI/ML for business- and mission-critical analytics. ESG recommends a serious look at WekaFS as a mature, robust, software-defined storage solution that is much more than an HPC storage system. Weka combines cloud-scale NAS and extreme performance with flexibility, agility, and protection designed for the enterprise.

[7] Source: ESG Master Survey Results, *2019 Data Storage Trends*, November 2019.

Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

 www.esg-global.com

 contact@esg-global.com

 508.482.0188