IDC PERSPECTIVE

WekaIO: Redefining What Digitally Transforming Enterprises Should Expect from Unstructured Storage Platforms

Eric Burgener

EXECUTIVE SNAPSHOT

FIGURE 1

WekaIO Re-Defines Enterprise-Class, High Performance Unstructured Storage

For digitally transforming enterprises, the evolution to more data-centric business models is re-defining how they want to use data analytics, and this is changing the requirements for the underlying storage infrastructure. In addition to supporting high throughput and massive capacity for large file environments, this storage must also now support low latency access for small file, metadata-intensive operations that require low latency, high availability and rapid recovery. Cloud integration is also critical as most enterprises move to hybrid cloud.

Key Takeaways

- Artificial intelligence, machine learning (AI/ML)-driven big data analytics workloads exhibit requirements that are not well met with traditional general-purpose file sharing and parallel file system platforms, limiting their scalability and making them harder to use.
- New technologies, innovations and storage designs are helping to overcome the classic concerns that have driven enterprises towards inefficient and expensive siloed unstructured storage strategies.
- Foremost among these factors are distributing metadata, enabling a global, media-optimized namespace for both file and object storage, and a file-based storage layer that does not require a separate volume manager.

Recommended Actions

- Understand your specific unstructured storage workload requirements in terms of latency, throughput, band-width, availability, scalability, data protection, random vs sequential I/O, and small vs large files, mapping these requirements to the functional capabilities of different unstructured storage platforms.
- Don't forget to consider backup, disaster recovery and archive capabilities, and don't underplay the importance of good cloud integration capabilities in unstructured storage platforms.
- Consider the opportunity to consolidate multiple unstructured workloads onto a single platform – newer designs that leverage newer technologies and features make this much more viable than it has been in the past, and successful consolidation can significantly reduce total cost of ownership.

Source: IDC, 2020

THE EVOLVING UNSTRUCTURED STORAGE MARKET

As enterprises undergo digital transformation (DX) – the evolution towards much more data-centric and data-driven business models – the unstructured storage market is beginning to change in significant ways. Enterprises are capturing, storing, protecting, analyzing and retaining more data than ever before, and many of them are already dealing with petabytes of data. The high rate of data growth is expected to continue, and in 2024 IDC expects that 143 zettabytes of data will be created (that's roughly 19TB of data for each person alive that year, assuming a world population of 7.5 billion). 70-80% of that data will be unstructured data. What's more – almost 25% of the data created in 2024 will be real-time data, a factor which has significant implications for the type of storage infrastructure that will be required to handle this data growth (particularly in terms of low latencies).

Next generation applications (NGAs) leveraging artificial intelligence, machine and deep learning-driven algorithms are being deployed to leverage the data collected by enterprises to drive better business decisions, and in many cases drive the business itself (e.g. social media, internet search, autonomous vehicles, etc.). These developments are all making information technology (IT) infrastructure more central to business success, driving a need for high availability for the types of systems supporting next generation applications. To meet their requirements for massive data scalability, enterprises often look to locate these NGAs on scale-out unstructured storage platforms that have not historically been able to deliver low latency or high availability for their traditional workloads.

In response, enterprises are modernizing their IT infrastructures to be able to better handle the new workloads and requirements of an increasingly digitized business environment. IDC's modernized IT infrastructure survey (n = 900, March 2020) indicated that roughly 70% of enterprises will be modernizing their server, storage and/or data protection infrastructure within the next two years, and that over 91% of organizations expect infrastructure modernization to be a critical determinant of DX success. As part of this effort, organizations expect to enhance their infrastructures with NVMe, solid-state media, software-defined, scale-out architectures, and cloud technologies. Spending on unstructured storage will grow at a compound annual growth rate (CAGR) of 7.7% to hit $33.9B by 2024. 89% of that spending will be on scale-out solutions.

While file- and object-based storage will still host their more traditional workloads separately, there is a clear trend towards the need for a global namespace that unifies both. Workloads in the financial services, life sciences, massively parallel computing and other application areas that leverage artificial intelligence and machine learning (AI/ML) are increasingly being architected around a secure, high-performance file-based front end that can simultaneously deliver low latency and high throughput and a massively scalable, object-based "data lake" back end that delivers non-disruptive expansion and low $/GB cost. The entire environment needs to be resilient to server and storage failures, rapidly recoverable, and offer flexible deployment options (cloud, non-cloud and hybrid cloud). With a single unified namespace that spans both file- and object-based storage and automated data mobility that places data on the appropriate tier according to user-defined policies in a manner fully transparent to applications, enterprises get the performance they need when data is being used and the cost-effective long-term storage of data in a flexible data lake that simplifies the use of data across a variety of different applications with different access methods.

Traditional vendors of scale-out file-based storage platforms have attempted to enhance their platforms to accommodate these changing requirements. It is difficult, however, to modify the parallel file system products that were popular in massively parallel and other technical workloads that primarily needed high throughput for sequential I/O to operate efficiently in the new environments. Most of

these systems were architected before solid-state storage became a mainstream media option, and they tend to be optimized for use with hard disk (HDDs) rather than solid-state drives (SSDs). The way they handle metadata limits both their ability to efficiently handle millions of small files and their overall scalability. They typically do not support high availability, a global namespace across multiple unstructured data types, or the data services needed to provide efficient data protection, simple backup and recovery, or cybersecurity. Given that many of the NGAs that were designed to perform AI/ML-driven big data analytics are small file and metadata intensive while at the same time requiring low latency, high availability and rapid recovery, traditional scale-out file system vendors have been challenged.

Making the File System Taxonomy Very Clear

In *IDC's Enterprise Storage Systems Taxonomy 2020* (IDC #US45949020, February 2020) we draw the distinction between the three different types of file systems. Scale-up file systems are designed to run on a single server system image. File system performance is mostly limited by controller performance, and many of these types of systems are limited to a single controller pair (for high availability). While system capacity can be expanded to a certain extent, moving to next generation controllers to reap significant additional performance generally requires a disruptive and time-consuming forklift upgrade. Over time, these types of systems have become known for their ability to handle general purpose file sharing (i.e. environments with millions of small files that needed lots of random, low latency I/O). Enterprise versions of these systems provide high availability and extensive storage management functionality (data reduction, snapshots, encryption, replication, etc.).

As file system workloads began to evolve, particularly in the massively parallel computing space, customers tried to run their workloads on these scale-up systems. These types of environments worked with very large files that were often accessed sequentially, and the ability to deliver these files to application servers fast enough became a real challenge. These environments also typically required the ability to scale up to the multi-petabyte range while the sweet spot for general purpose file serving often required an order of magnitude less storage capacity. In the traditional scale-up model, performance metrics such as latency, throughput and capacity are limited to the maximum performance deliverable from a single storage system.

In an attempt to address some of these limitations, some vendors of scale-up file systems fielded a software-based abstraction layer that allowed them to "cluster" multiple scale-up file servers together in a single namespace. From a management point of view, this allowed unified management for multiple file servers under a single namespace, extending the capacity scalability to more resemble a scale-out model. It did not, however, really address the need for extremely high throughput to very large files. In these clustered environments, the performance with which any single file could be accessed could still not exceed the maximum performance capability of a given file server. Newer technologies did extend the performance capabilities of single servers, but these high-performance scale-up file servers could be very expensive and were still limited by their dual controller designs. In the IDC taxonomy these types of file system environments are referred to as "federated scale-out file systems".

In the late 1990s, vendors began to develop off-the-shelf parallel file systems that could distribute a single file across multiple nodes in a clustered environment. This broke the single system performance "limit", allowing the throughput and bandwidth with which a single large file could be accessed to be scaled as more cluster nodes were added. While these systems did not perform very well on workloads that required lots of random, low latency I/O across many small files, they performed much better in technical environments with very large files that needed highly scalable throughput and

cost-effective, massive capacity. Distributing the data in this manner helps to avoid hot spots in accessing it since many more network connections, servers, and devices are brought into play to serve up the data. These types of scale-out, parallel file systems began replacing the scale-up designs for many of the more performance-sensitive large file workloads and have come to be widely used in environments across many industries that required high throughput for very large file sizes.

With new storage technologies and designs becoming available, a new opportunity arose to create a single parallel file system product that could simultaneously deliver low latency for random small file I/O and high throughput and cost-effective scalability for sequential I/O against extremely large files. In addition, the need to access data via a global namespace, regardless of where or on what tier the data resides, is an emerging requirement. A storage startup, WekaIO, delivered a new class of parallel file system in 2017 that met this consolidated requirement and extended the size of usable unstructured data environments into the multi-exabyte range (an exabyte is 1000 petabytes). For customers that have been running up against scalability limitations with more traditional parallel file system designs, require a unified namespace for millions of smaller files, and/or want to consolidate multiple different types of file-based workloads onto a single platform, IDC recommends a close look at the vendor's Weka File System (WekaFS) platform. This is true regardless of whether the data sets are on-premises, in the cloud, or in a hybrid cloud environment.

WekaIO: Performant Unstructured Storage for Next Generation Workloads

WekaIO is a privately held, modern data platform vendor headquartered in Campbell, California with engineering located in Tel Aviv, Israel. Founded in 2013, the vendor's core product, WekaFS, became generally available in 2017. WekaFS is a true parallel file system platform, architected around a scale-out, software-defined design, that in three short years has arguably proved itself as the most scalable unstructured storage solution on the planet for today's AI/ML-driven big data analytics workloads. They primarily target life sciences, financial services, massively parallel computing (MPC), and other AI/ML-driven big data analytics workloads and count among their customers 9 of the Fortune 64, the world's 10 largest data farms, and many production customers with multiple tens of petabytes of storage capacity under management. They routinely compete against the two file-based storage giants (Dell EMC PowerScale and NetApp ONTAP) and currently hold world records in the IO-500 Challenge, STAC-M3 and SPEC benchmarks.

WekaIO sells 100% through the channel and has a strong channel program (called the Weka Innovation Network (WIN)) with more than100 channel partners in North America specifically selected for their vertical market and parallel file system expertise. But interestingly (and somewhat uniquely), customers (and resellers) can also buy servers directly from the server OEMs with WekaFS software factory-loaded using a single SKU that appears on the server OEM's price list (a program the vendor refers to as "Weka Within"). The Weka Within validation and certification is a designation of interoperability and compatibility, functionality, and performance that meets Weka's exacting standards for data platforms designed for I/O-intensive workloads. Partners in the program include Amazon Web Services (AWS), Cisco, Dell, Hitachi Vantara, HPE, Lenovo, NVIDIA, Penguin Computing, and Supermicro. This addresses some key support and procurement issues that exist with other software-only go-to-market strategies in a way that customers really like.

Key WekaFS Differentiators

Most vendors of file system products tout the performance and scalability of their environments, leaving astute customers to sort out the suitability of different offerings for different types of file system workloads. WekaIO also headlines the performance and scalability of their offering but is clear about

the fact that these claims simultaneously apply to both low latency, random small file I/O and high throughput, sequential large file I/O workloads. WekaIO bases their performance and scalability claims on six sustainable technical differentiators that set WekaFS apart from other players:

- **Data and metadata are distributed.** While all parallel file systems distribute data to provide high throughput to large files that span multiple nodes, most place all the metadata in dedicated servers that ultimately become performance bottlenecks as environments scale. WekaFS distributes the metadata as well, retaining it at all times on NVMe media to promote high performance for random small file I/O as well as sequential large file I/O.
- **Global, media-optimized namespace that spans both file- and object-based storage.** With the increasing popularity of cloud technology, many parallel file systems can tier data to external, object-based storage using S3 to extend the scalability of the overall environment, but they cannot do so in a manner transparent to applications accessing the file system. The WekaFS environment is composed of two separate components (a parallel file system and an object store), linked by intelligent information life cycle management processes that automatically move data back and forth between the two tiers without requiring any application changes.

  WekaFS has written their parallel file system specifically (and only) for NVMe-based storage, a design which reads and writes data so as to optimize free space management, promote more consistent read and write latencies, and maximize the endurance of solid-state media. Prospective buyers should note that WekaIO is the only unstructured storage vendor that has taken this approach, a design decision that looks prescient given the industry's large-scale movement towards solid-state media for both primary and secondary storage.
- **Single click "snap-to-object" feature for simple, reliable and comprehensive backup.** Because most parallel file systems are deployed on top of block-based volume managers, it can be very difficult to create a recoverable copy of the entire file system environment (if it can be done at all). WekaFS includes a single click "snap to object" function that instantly creates a complete, recoverable copy of the front-end file-based data that includes all its own metadata (enabling recovery to and/or use from anywhere, not just its original source). The snapshots can then be recovered to a completely different system, even in a different data center or the public cloud. The recovering system can be smaller or larger than the original, another departure from the common approach in the industry that requires recovering to the exact same configuration in many cases.
- **Multiprotocol access to the same data.** Unstructured data stacks now are diverse and contain different applications accessing the same data from different platforms. WekaFS allows accessing the same data using WekaIO's POSIX client, NFS, Hadoop, SMB for Windows access, NVIDIA's GPUDirect Storage (GDS) and S3. This provides increased flexibility for organizations to be able to support all their needs from a single system.
- **Patented, distributed data protection protocol specifically designed for performant, efficient operation in large scale environments.** Traditional data protection mechanisms suffer from limited scalability, long rebuild times, heavy CPU utilization, high latencies, and/or high costs, all of which impose increasing burdens as environments scale. Weka Distributed Data Protection (DDP) adapts its data and parity distribution algorithms as nodes are added and optimizes data movement based on the nature of the failure for rapid, efficient recovery. DDP can protect against (and automatically recover from) 2 simultaneous node and/or device failures with only 11% redundant capacity while maintaining latencies in the sub-millisecond range.
- **Ease of management.** Initial deployment and ongoing management leverage system defaults to simplify common tasks like storage provisioning, file system administration, and expansion.

Storage device and/or node addition can be done in minutes without impacting application performance. WekaFS was designed as a self-optimizing system that has few parameters that require manual tuning – data protection and data distribution dynamically adjust to changing conditions without operator intervention to enable predictably consistent performance as the system scales. The fact that WekaFS installs directly on top of NVMe storage (no separate volume manager has to be installed or managed) makes it much simpler and easier to configure, backup and upgrade. There is a graphical user interface for intuitive manual administration as well as a command line interface which provides full access to all system functionality when creating automated workflows.

It's hard to argue with the success WekaIO has achieved in terms of bringing over 100 very large customers, many of whom have many years of experience with more traditional scale-out file system designs, into the fold. For the last full year of financial results available at this writing (2019), WekaIO grew revenues by 600%.

The WekaFS Architecture

WekaFS is a software-defined, parallel file system environment that deploys on top of commodity Intel x86-based or AMD EPYC gen2 CPU server hardware with internal storage. The compute tier can be composed of either Linux or Windows clients. A variety of different Linux servers are supported as long as they have internal NVMe SSDs and are running a Weka client (which runs an optimized proprietary file system protocol). Windows clients can communicate with the file system tier over NFS or SMB but will not be running a Weka client. No specialized hardware of any kind (storage class memory, etc.) is required, a feature which allows WekaFS to be run cost-effectively in the public cloud as well. Compute servers can be composed of general-purpose CPUs or accelerated compute GPUs, and the different types of CPUs can be mixed and matched across various vendors and technology generations. Compute servers can be connected to the WekaFS storage tier over either Ethernet or InfiniBand networks (or actually on both infrastructures concurrently – which is very unique).

The WekaFS storage tier is comprised of two components: a front-end scale-out cluster of WekaFS-based storage servers with internal NVMe SSDs and an optional back end tier based on 3rd party object-based storage. The back-end object tier (when desired) is used as a data lake, providing much more cost-effective longer term storage for colder data. Supported object storage platforms and vendors include Amazon S3, Ceph, Cloudian, Hitachi Vantara HCP, IBM (Red Hat) COS, Quantum ActiveScale, Scality and SwiftStack. The entire storage tier is managed as a unified, global namespace with WekaFS managing data placement based on usage patterns. As colder data is moved to the object storage using S3, a "stub" is maintained on the file-based storage that includes the relevant metadata, allowing the data to be transparently accessed as needed without requiring any application modifications.

A container housing WekaFS (shown in Figure 2 as the light green circle) runs on each of the storage cluster nodes. WekaIO's containerized implementation offers a unique capability not found in many other unstructured storage offerings: customers can deploy it in a "dedicated" mode where all the configured resources on the storage node are used by WekaFS or in what the vendor calls a "converged" mode where the file system is limited to using only certain pre-defined resources on that node, keeping other resources available to run applications (effectively letting customers deploy it in a hyperconverged infrastructure (HCI) model). This deployment flexibility gives customers additional configuration options, particularly in edge deployments. When considering edge deployments, however, the reader should note that the smallest viable configuration is a five-node cluster.

FIGURE 2

The WekaFS Environment

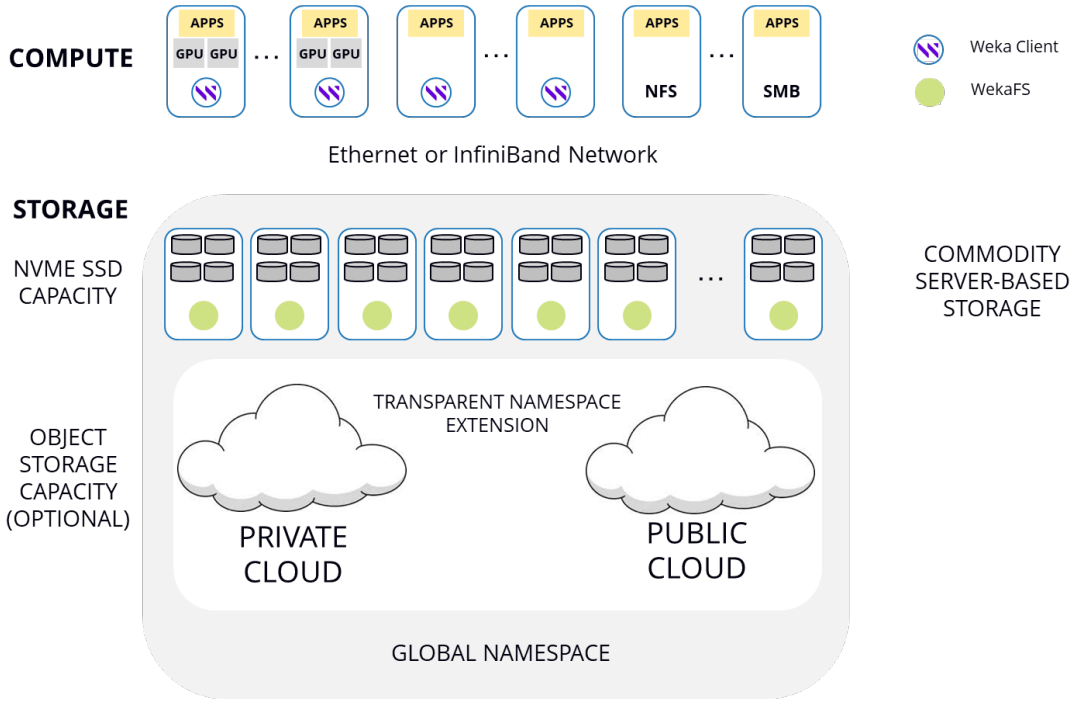

Source: IDC, 2021

WekaFS is specifically optimized for NVMe-based solid-state storage using I/O algorithms that maximize the performance, reliability and endurance of flash media. The 3rd party object storage back ends have generally been written to specifically optimize the performance, reliability and endurance of the spinning disk media used in the object tier. The two components work seamlessly together, appearing to applications as a single, massively scalable, high performance file tier. This design provides NVMe-like performance for all data at HDD-like prices (with the actual blended $/GB cost being determined by the actual capacity of each of the two tiers).

WekaFS is highly scalable with the ability to manage up to 6.4 trillion files or directories and up to 1024 unique virtual file systems in one cluster as well as supporting a global namespace that can be up to 8 exabytes in size. A single file can be as large as 4 petabytes, there can be up to 6.4 billion files in a single directory, and the file-based storage cluster can be scaled to 3275 nodes. Working with NVIDIA, WekaIO has demonstrated the ability to deliver 162GB/sec of bandwidth for a single NVIDIA DGX A100 client using NVIDIA's GPUDirect Storage protocol. And this is with a single file system mount point! Traditional scale-out NAS needs many mount points to deliver that amount of bandwidth, which translates directly to a more complex environment. This means that it takes a smaller WekaIO system that is easier to manage to meet any given throughput or bandwidth requirement. A maximally

configured WekaFS storage cluster is capable of handling multiple tens of millions of IOPS and more than two terabytes per second of bandwidth at less than 300 microsecond latencies.

With its software-defined design, WekaFS can run on in-house server hardware infrastructure or it can be run in the public cloud (on Amazon S3). The same exact binaries are deployed in each case. Since this is not necessarily true of competitive unstructured storage products that can be run on either in-house or public cloud infrastructure, it is worth pointing out that regardless of deployment location WekaIO customers enjoy exactly the same functionality. This makes it easier to run WekaFS in hybrid cloud environments.

Other Relevant WekaFS Features

- **Granular snapshots.** Snapshots are taken at the file system level, can be immutable or writeable (as a full clone), and are space-efficient (a redirect-on-write architecture so space consumption is based on write deltas from the originating snapshot or clone). Taking a snapshot is simple (a single click through the GUI), fast (there is a sub-second service interruption that does not require quiescing any applications), and the system supports up to 4096 snapshots. Snapshots are most often used for backup, archive, compliance, test and development purposes.
- **Tiering.** Administrators specify the size of the NVMe and object tiers separately, and tiering is specified at the file system level. Data is tiered using either a default or a user-defined policy based on access and modification times, and applications will only ever access the front-end file system regardless of where data is actually located. No changes to applications are required. Metadata is always retained on the NVMe storage on the front-end file cluster to ensure low latency access to both small and large files.
- **End-to-end encryption.** Through tight integration with key management systems that are KMIPS compliant, WekaFS supports AES 256-bit encryption both on-the-fly (from application clients to the storage system) and at-rest. The system also provides end-to-end data protection enabling high-end data integrity checks usually leveraged in block storage systems ("SAN"), making sure each 4KB of data is verified against checksum before it is accepted by the client.
- **Security.** In addition to end-to-end encryption, WekaFS supports role-based access control (through its GUI) as well as Lightweight Directory Access Protocol (LDAP) integration for user authentication and permissions. The snap-to-object feature supports an "air gap protected" copy on remote object storage which provides protection against malware or ransomware attacks. The WekaFS security supports multi-tenant management: a single WekaFS environment can support multiple organizations with each organization's administrator(s) being able to mount, access and manage only its own subset of the filesystems and resources.
- **Multi-protocol support.** WekaFS supports access via POSIX, NFS, SMB, GDS, and S3. Customers should note that WekaFS provides the full POSIX semantics on Linux (with support for cache coherency with the standard Linux Page Cache, atomic open, atomic append, etc.) but does not suffer the performance degradation traditionally associated with making a shared file system POSIX compliant. This allows WekaFS customers to enjoy the ease of use of shared file systems with much better performance than that supported by competitors.
- **Container Storage Interface (CSI) support.** WekaIO supports a Kubernetes CSI plugin to allow its customers to deliver Container-as-a-Service functionality on top of WekaFS. This is an important feature since many of the applications being used in AI/ML-driven big data analytics

like MySQL, Cassandra and MongoDB, require the stateful, easily portable enterprise-class storage that WekaIO's CSI plugin makes available.

- **Efficient updating operations.** When modifying files, WekaFS does not need to retrieve the entire file, only the small portion that must be updated. This saves time and makes better use of IT infrastructure resources, particularly when data being updated resides in the object tier. And WekaFS does not perform read/modify/write operations against pre-existing objects – changes will just be uploaded as a new object – so the performance of the back-end object store will not impact the performance of the front-end file system.
- **Subscription-based licensing**. WekaFS is licensed based on usable capacity under management and includes the functionality described in this document, tiering to an object store and 7x24 support. From time-to-time WekaIO may release major new features that are licensed separately.

Given the strong growth of AI/ML-driven big data analytics workloads in digitally transforming enterprises, accelerated compute has become one of the most sought-after new technologies during infrastructure modernization. NVIDIA is the clear leader in the accelerated compute space and has struck partnerships with several of the parallel file system vendors to enable the use of their GPUs with these software vendors' products. It is notable, however, that NVIDIA has only invested in one parallel file system vendor, and that is WekaIO. In May 2019, both NVIDIA and their recently purchased subsidiary Mellanox announced investments in WekaIO focused around AI/ML-driven workloads, lending strong credence to WekaIO's strategy and platform.

Along these same lines, it's also worth noting that WekaIO is a "first mover" with support for NVIDIA's GPUDirect Storage (for improving bandwidth and latency by enabling a direct path between local or remote storage and GPU memory). These enhancements are extremely important for the types of AI/ML-driven NGAs that digitally transforming enterprises are deploying and WekaIO's first mover status in supporting them has clearly helped to close new customers who need this type of performance and efficiency in their unstructured storage environments.

Go-To-Market Strategy

WekaIO has crafted a go-to-market strategy that maximizes flexibility both for their channel partners as well as for their customers. While many parallel file system vendors do sell both an appliance and a software-only version of their platform, only the appliance (which is available on usually just one kind of server hardware) provides a true single point of support contact. Indirect channel partners that load the software on their server of choice still have to split support between the server and parallel file system vendors. This strategy also provides little comfort to procurement organizations in potential customers, who must still overcome their hesitancy in dealing with startups.

WekaIO has done things a bit differently. The vendor has established partnerships with all the major x86-based servers OEMs, including Cisco, Dell, Hitachi Vantara, HPE, Lenovo, Supermicro and Western Digital, whereby the server OEM puts WekaFS on their pricing schedule. The server partners create SKUs that include their server hardware pre-loaded with WekaIO software and can sell those either to value added resellers (VARs) or directly to end user customers. This is the "Weka Within" program referred to earlier and has several important implications for customers. It makes the server OEM a single point of support contact, a simplifying feature for administrators looking for a "single throat to choke", and appeases procurement, who prefers to deal with larger, more established vendors. It also makes it easier to include a WekaIO purchase as part of an Enterprise License Agreement (ELA) with one of the server OEMs. Any VARs that can resell the server OEM's products

also like this since it gives them a direct ability to sell (and get compensated by the server vendor for) WekaFS deployments.

It is notable that many of these server vendor relationships are backed with investments by those companies in WekaIO. Cisco, Hitachi Vantara, HPE, NVIDIA, Seagate and Western Digital have all made financial investments in the vendor. Additionally, Hitachi Vantara recently struck an OEM deal whereby they will be white-labeling WekaFS under their own brand (an arrangement that includes sales commitments on Hitachi's part that gives WekaIO a significant amount of guaranteed revenue). NVIDIA has done a number of Reference Architectures with WekaIO and Seagate embeds WekaFS into their object store (CORTX). WekaIO is on the AWS MarketPlace, and Micron has also invested in the unstructured storage vendor. These relationships underline WekaIO's status as a leader in unstructured storage and give the vendor a market reach that far exceeds the typical company of their size.

The existence of server SKUs with embedded WekaFS available directly from the server OEMs has clearly strengthened WekaIO's sales momentum. This strong backing from well-established and proven enterprise infrastructure providers incents potential customers to look at WekaFS as a platform that transcends the server hardware with which it is initially sold. Customers often sign contracts to purchase WekaFS that are longer than the typical server life cycle, planning to stay with the platform even as they upgrade to next generation servers from the same server OEM underneath it. This has given WekaIO the ability to close many more longer-term contracts that most startups, an ability which strengthens their financial status. And the fact that Weka Within applies to multiple server OEMs gives customers another deployment option that can be of interest: running different WekaFS storage clusters that are 100% supported on different server hardware that all leverage the same object tier. For example, a WekaFS storage cluster in a primary location could be running on Dell PowerEdge hardware, a disaster recovery configuration in a remote secondary location could be running on Supermicro hardware, and they could both be communicating with a data lake object tier running on Hitachi hardware.

Roadmap Items for Prospective Buyers to Ask WekaIO About

WekaIO has both a very strong technical case and strong benchmark proof points to validate their claim to be a high performance, massively scalable file system that can simultaneously and cost-effectively be used to consolidate both latency-sensitive, random small file I/O and throughput and bandwidth-sensitive sequential large file I/O on the same platform. They provide functionality that meets common customer needs like snapshots for data mobility, snap-to-object for easy recovery, tiering with a unified global namespace to deliver massive scalability, multi-protocol support for flexibility in usage, role-based access control and encryption for security, subscription-based pricing to smooth storage acquisition costs over time, and different deployment models (appliances, software-only, public cloud). There are several key enhancements, however, that would make WekaFS an even more compelling unstructured storage option:

- **Native S3 support.** Today, WekaFS requires a gateway for S3 front-end functionality, resulting in some management overhead. A native S3 interface would increase performance for S3 operations and make managing a multi-protocol environment easier. The vendor has already stated their intent to deliver this capability in 2021.
- **Data reduction.** Data reduction uses technologies like compression, deduplication and others to increase capacity utilization. Unstructured data often is not as supportive of the high data reduction ratios administrators may be familiar with from backup environments, but every little

bit helps. Given that most WekaFS environments are quite large, even a low data reduction ratio could help to drive considerable savings though.

- **Explore WekaIO's cloud integration capabilities.** While the vendor already has a story here with CSI support and the ability to run WekaFS in the public cloud, the future is multi hybrid cloud. As more enterprises implement DevOps, add cloud-native applications, and want to leverage multiple public clouds, make sure to understand how WekaFS enhancements will further enable more agile hybrid cloud infrastructure.

ADVICE FOR THE TECHNOLOGY BUYER

- **Understand your specific workload requirements.** Make very sure you understand exactly what type of I/O your unstructured workload(s) require. Know the difference between low latency, random small file I/O and high throughput/bandwidth sequential large file I/O and take into account any newer AI/ML-driven big data analytics workloads that may exhibit a mix of these different I/O profiles. Take requirements around high availability, scalability, data protection and recovery into account. Also consider the different facets of high availability – how easy is it to backup and recover both individual files and objects and the entire system, etc. Also consider archive requirements and a system's ability to support those – an effective archive strategy can help to lower costs significantly.
- **Map your I/O requirements to the functional capabilities of different offerings.** Historically, different file systems have been optimized for different types of environments. Make sure to understand the strengths and weaknesses of alternatives you are considering when it comes to handling file I/O so that you aren't "surprised", particularly if yours is a high data growth environment.
- **Consider the opportunity to consolidate multiple unstructured workloads onto a single platform.** With newer technologies and designs addressing many of the reasons enterprises have pursued siloed storage environments in the past, IDC recommends that customers actively consider opportunities to consolidate different types of workloads onto fewer storage platforms as they modernize infrastructure. When consolidation can be successfully achieved, it can significantly reduce storage acquisition and ongoing management costs, as well as streamline data re-use and recovery processes. Make sure as part of this process to consider the value of distributed metadata, unified global namespaces and multiple access methods.
- **Don't undervalue cloud integration capabilities.** Although not all enterprises may require multi public cloud support today, it is highly likely that digitally transforming organizations will ultimately move to a hybrid cloud environment in the future (a high percentage of them are already there today). As IT infrastructure gets modernized, more and more enterprises will be looking to optimally place workloads in one of three different models - traditional in-house, private cloud and public cloud infrastructure – and the use of server-based, software-defined in-house infrastructure with excellent cloud integration points will make that process easier.

LEARN MORE

Related Research

- *Flash Usage in Unstructured Data Environments is Surprisingly High and Only Expected to Increase Over Time* (IDC #US47007120, November 2020)
- *Worldwide Scale-Out NAS and Object Storage Appliance Market Shares, 2019: Modest Growth Before COVID-19 Impact* (IDC #US46555520, June 2020)

- *Worldwide File- and Object-Based Storage Forecast, 2020-2024* (IDC #US45992520, June 2020)

Synopsis

The unstructured storage market is undergoing significant change right now. In addition to maintaining more traditional general-purpose file sharing and scale-out file system workloads, enterprises are also deploying next generation, artificial intelligence and machine learning-driven big data analytics applications. These newer workloads are exhibiting requirements that pose a real challenge for more traditional unstructured storage platforms. WekaIO, a storage software vendor, has incorporated a number of newer technologies and design innovations into their parallel file system that deliver compellingly differentiating value for large scale customers struggling with the evolving storage requirements of digital transformation. This Vendor Profile takes a look at WekaIO and its WekaFS parallel file system platform, discussing both its technical differentiators and go-to-market strategy.

"Unstructured workloads are changing, and traditional scale-out file designs are having difficulty in scaling to the levels of performance, availability, and capacity that customers are increasingly demanding," said Eric Burgener, research vice president, Infrastructure Systems, Platforms and Technologies Group, IDC. "WekaIO has raised the bar for what enterprises should expect from their scale-out file systems with a number of innovations that sustainably differentiate its WekaFS platform."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com