

WekaFS™ Architecture White Paper

JULY 2020



ABSTRACT

This white paper provides technical overview of the features and benefits of WekaFS file system. The white paper provides details on theory of operation, features and management, and independent performance validation.

REVISIONS

DATE	DESCRIPTION
JULY 2017	Initial Document, authored by Liran Zvibel and David Hiatt
JUNE 2019	Feature updates, authored by Barbara Murphy
JULY 2020	Major update with expanded content to include new features, performance, sizing, best practises, authored by Barbara Murphy

CONTENTS

Introduction	1
Common Storage Challenges in the Cloud Era	1
Designing a Unified, Cloud Era Storage Solution	2
WekaFS – Storage Without Compromise	2
WekaFS Benefits	3
WekaFS Architecture	4
File System Design	5
Supported Protocols	5
Weka Storage Nodes	5
Integrated Flash and Disk Layers for Hybrid Storage	6
Networking	6
Network High Availability (HA)	7
Management GUI	7
3rd Party Monitoring Tools	8
Adaptive Caching	8
Global Namespace and Expansion	9
Tiering	9
Snapshots and Clones	10
Snap-to-Object	10
Data Protection	11
Data Distribution	12
WekaFS Rebuilds	12
Power-Fail and End-to-End Data Protection	12
Automated Data Rebalancing	13
Container Storage Integration (CSI)	13
Multi-Tenant Organizations	13
Authentication and Access Control	13
In-Flight and At-Rest Encryption	14
Key Rotation and Key Management	14
Flexible Deployment Options (on-premises and cloud)	14
Converged	15
Dedicated Storage Server	15
Native Public Cloud	16
WekaFS Performance	16
Standard Performance Evaluation Corporation: SPEC.org	17
Performance Scaling	18
Performance to GPU Storage	19
Summary	21

INTRODUCTION

WekaIO™ (Weka) was founded on the idea that current storage solutions have only provided incremental improvements to legacy designs, widening the gap between compute performance and the underlying data storage. Storage remains a bottleneck to application performance, and the gap has widened with the continued densification of compute in areas such as GPU-based applications. In today's hyper-competitive market, organizations need flexible infrastructure; application workloads are becoming increasingly complex and data sets are continuing to grow unchecked, forcing enterprises to architect overly complicated and costly systems that reduce IT agility. As a result, important business insights remain locked away, out of reach of decision makers.

IT organizations are adopting cloud technology for its fluid, on-demand scalability that supports diverse workloads at scale. However, while network and compute can be virtualized to operate at scale very effectively, storage remains largely isolated in silos based on system performance profiles. Consequently, organizations are forced to architect a storage system that is highly customized for their environment and workloads from building blocks that do not scale. The result is a storage solution that is complex, temperamental, expensive, and slow.

WekaIO has built a software-only, high-performance file-based storage solution that is highly scalable and easy to deploy, configure, manage, and expand. The design philosophy behind the Weka file system (WekaFS) was to create a single storage architecture that runs on-premises or in the public cloud with the performance of all-flash arrays, the simplicity and feature set of network-attached storage (NAS), and the scalability and economics of the cloud.




Figure 1: WekaFS Benefits Summary

WekaFS Highlights (Figure 1)

- World's fastest shared file system as validated on SPEC SFS 2014¹, IO-500² and STAC³ benchmarks
- Supports bare-metal, containerized, virtual, and cloud (on-prem, hybrid, public) environments
- Deployable as a converged platform, a dedicated storage appliance, or native in the cloud
- Flexible application storage access, including POSIX, NFS, SMB and NVIDIA® GPUDirect® storage
- Zero performance tuning to support small and large files simultaneously, with both mixed random and sequential I/O patterns
- Application-level 4K I/O, consistent sub-250 microsecond latency on high-speed networks, unlimited random IOPs performance – linear scaling in line with the size of the cluster
- Automated built-in tiering to expand the namespace from fast flash to hard disk storage via on-premises or cloud object storage
- Fully distributed data and metadata to ensure that there are no hotspots in the storage cluster
- Distributed resilience that eliminates the bottlenecks of traditional data protection
- Full cloud integration with cloud-bursting for a hybrid cloud model or 100% public cloud
- Snapshots and snapshot to the cloud for backup, archive and disaster recovery

COMMON STORAGE CHALLENGES IN THE CLOUD ERA

Modern-day applications have a wide variety of performance requirements (IOPs, bandwidth, latency), and when combined with the diversity of application file formats, access protocols, and data structures, it can all lead to increased IT complexity.

¹ <https://www.weka.io/press-releases/wekaio-Weka-cements-storage-leadership-groundbreaking-performance-latency-results-spec-sfs-2014/results/>

² <https://www.weka.io/press-releases/wekaio-places-first-on-io-500-challenge/>

³ <https://www.weka.io/blog/weka-hpe-and-kx-systems-set-17-new-stac-m3-records/>

Storage architects try to work around these limitations by employing multiple storage architectures for specific applications. All-flash storage-area networks, commonly called all-flash arrays (AFAs) are optimized for performance but they provide neither the scale of the cloud nor the simplicity and shareability of network-attached storage (NAS). Furthermore, AFAs provide block storage protocols that are not shareable across servers, thus they are unsuitable for shared storage use cases. The result of these workarounds has been a costly and endless cycle of rearchitecting infrastructure to keep pace with changing application requirements.

It is a daunting task to determine which storage solution is best suited to a particular environment or application workload because of the variety of options available. Some solutions are optimized for performance while others are optimized for scale. Workloads in the technical compute space, such as Artificial Intelligence (AI) and Machine Learning (ML), genomic research, and financial analytics, all of which generate both large file sequential access and small file random access on very large data sets, are especially problematic. No single traditional storage design has been able to address all of these workload patterns. The workaround has always been to use multiple storage systems and complex data management platforms.

DESIGNING A UNIFIED, CLOUD ERA STORAGE SOLUTION

When designing a modern storage solution, a key consideration is to account for the continuing evolution and improvement of technology. A truly software-defined storage solution should accommodate such changes, which means that it must be able to run on commodity server hardware, adapt to customer environments, and add cloud-like agility, scalability, and on-demand performance. It should also be simple to deploy and expand fluidly without incurring the typical procurement delays associated with traditional external storage appliances.

The limitations created by legacy design constraints led the founders of Weka to develop a brand new file system that delivers the performance of all-flash arrays, the simplicity of scale-out NAS, and the scalability of the cloud in a single architecture. Figure 2 provides an overview of the Weka file system value proposition.

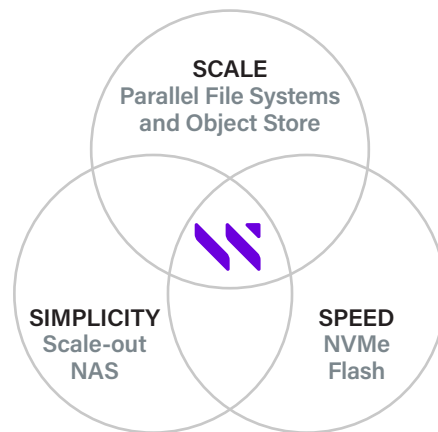


Figure 2: The Weka File System Design Approach

The Weka file system WekaFS™ is a software-only storage solution with a clean sheet design that solves the problems associated with traditional storage systems. It runs on any standard AMD or Intel x86-based server hardware with commodity NVMe solid-state disks (SSDs), eliminating the need for custom specialized hardware. This approach allows you to take advantage of improvements in technology without the pain of complete forklift upgrades to next-generation architectures, including public cloud deployments.

WekaFS – STORAGE WITHOUT COMPROMISE

WekaFS solves the common storage challenges previously mentioned by eliminating the chokepoints that impact application performance. It is well-suited for demanding environments that need shareable storage with low-latency, high-performance, and cloud scalability.

Example use cases include:

- Artificial intelligence and machine learning

- Life sciences including genomics and Cryo-EM
- Financial trading, including backtesting, time-series analysis, and risk management
- Engineering DevOps
- Electronic design and automation
- Media rendering and visual effects (VFX)
- High-performance computing (HPC)

By leveraging existing technologies in new ways and augmenting them with engineering innovations, Weka's software delivers a more powerful and simpler solution that would have traditionally required several disparate storage systems. The resulting software solution delivers high performance for all workloads (big and small files, reads and writes, random, sequential, and metadata heavy). Furthermore, because it is designed to run on commodity server infrastructure, it does not rely on any specialized hardware.

WekaFS is a fully-distributed, parallel file system that was written entirely from scratch to deliver the highest-performance file services by leveraging NVMe flash. The software also includes integrated tiering that seamlessly expands the namespace to and from hard disk drive (HDD) object storage, without the need for special data migration software or complex scripts; all data resides in a single namespace for easy access and management. The intuitive graphical user interface allows a single administrator to quickly and easily manage hundreds of petabytes of data without any specialized storage training.




Figure 3 – WekaFS combines NVMe flash with cloud object storage in a single global namespace

WekaIO's unique architecture, as shown in Figure 3, is radically different from legacy storage systems, appliances, and hypervisor-based software-defined storage solutions because it not only overcomes traditional storage scaling and file sharing limitations but also allows parallel file access via POSIX, NFS, SMB and GPUDirect Storage. It provides a rich enterprise feature set, including local and remote snapshots to the cloud, clones, automated tiering, cloud-bursting, dynamic cluster rebalancing, private cloud multi-tenancy, backup, encryption, authentication, key management, user groups, quotas and much more.

WekaFS Benefits

- Highest performance – ideal for mixed small and large file workloads
- Scalable capacity – start as small as 30TB and scale to hundreds of petabytes in a single namespace
- Strong security – keep data safe from threat or rogue actors with encryption and authentication
- Hybrid Cloud – burst to AWS for compute agility or run natively in the cloud
- Backup – push backups straight to private or public cloud for long term retention
- Best economics – combine flash and disk for best cost at scale

WekaFS ARCHITECTURE

Weka's parallel file system is designed to provide a cloud-like experience, whether you run your applications on-premises or plan to move them to the cloud. WekaFS provides a seamless transition to the cloud and back.

Most legacy parallel file systems overlay file management software on top of block storage, creating a layered architecture that impacts performance. WekaFS is a distributed, parallel file system that eliminates the traditional block-volume layer managing underlying storage resources. This vertically integrated architecture does not suffer the limitations of other shared storage solutions and delivers both scalability and performance effectiveness (productivity).

Figure 4 below provides an overview of the software architecture from the application layer all the way to the physical persistent media layer. The Weka core components, including the WekaFS unified namespace and other functions such as virtual metadata servers (MDSs), execute in user space in a Linux container (LXC), effectively eliminating time-sharing and other kernel-specific dependencies. The notable exception is the Weka Virtual File System (VFS) kernel driver, which provides the POSIX filesystem interface to applications. Using the kernel driver provides significantly higher performance than what can be achieved using a FUSE user-space driver, and it allows applications that require full POSIX compatibility to run on a shared storage system.

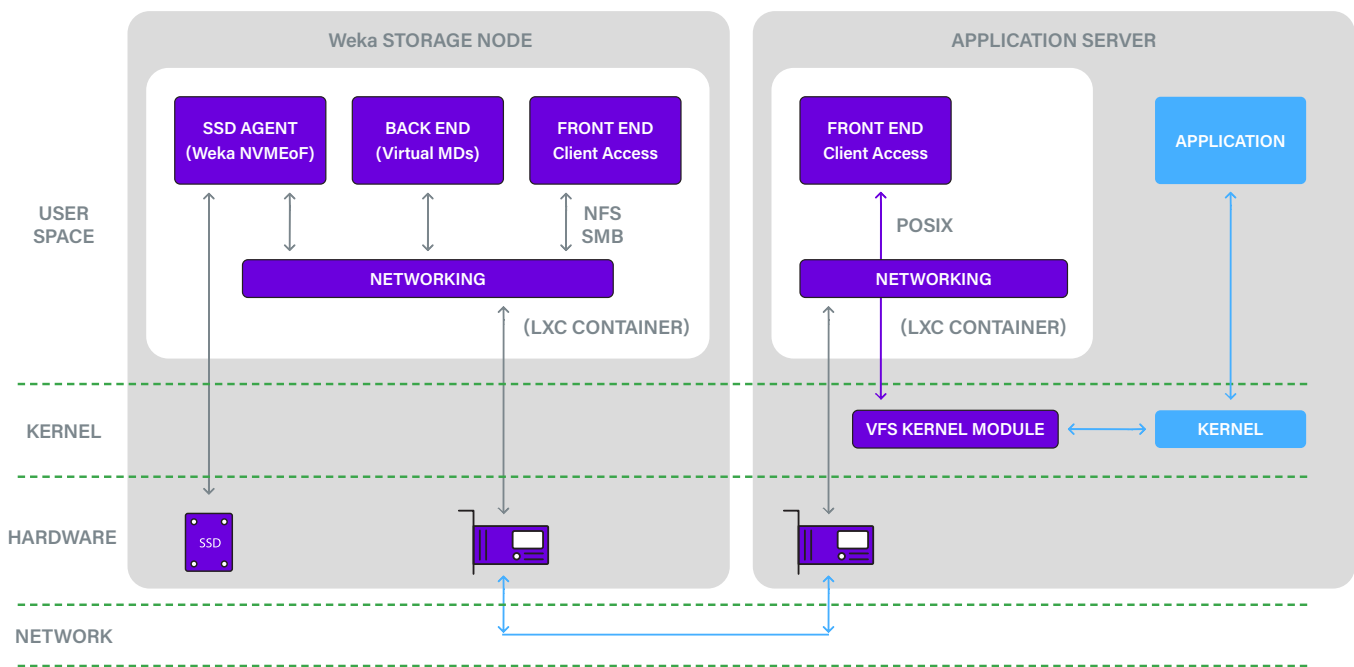


Figure 4 - WekaFS Software-Based Storage Architecture

Weka supports all major Linux distributions and leverages virtualization and low-level Linux container techniques to run its own RTOS (Real-Time Operating System) in user space, alongside the original Linux kernel. Weka manages its assigned resources (CPU cores, memory regions, network interface cards, and SSDs) to provide process scheduling, memory management, and to control the I/O and networking stacks. By not relying on the Linux kernel, WekaFS effectively utilizes a zero-copy architecture with much more predictable latencies.

WekaFS functionality running in its RTOS (figure 4) is comprised of the following software components:

- File Services (Front End) – manages multi-protocol connectivity
- File System Clustering (Back End) – manages data distribution, data protection, and file system metadata services
- SSD Access Agent – transforms the SSD into an efficient networked device
- Management Node – manages events, CLI, statistics, and call-home capability
- Object Connector – read and write to the object store

Weka core software runs inside an LXC container that has the benefit of improved isolation from other server processes. A Weka VFS driver enables WekaFS to support full POSIX semantics leveraging lockless queues for I/O to achieve the best performance while enhancing interoperability. The WekaFS POSIX file system has the same runtime semantics as a local Linux file system (e.g., Ext4, XFS, and others), enabling applications that previously could not run on NFS shared storage because of POSIX locking requirements, MMAP files, performance limitations, or other reasons. These applications will enjoy massively improved performance compared to the local file system (figure 14).

Bypassing the kernel means that Weka's software stack is not only faster with lower latency, but is also portable across different bare-metal, VM, containerized, and cloud instance environments.

Resource consumption is often a problem with traditional software-based storage designs because these solutions either take over the entire server or share common resources with applications. This extra software overhead introduces latency and steals precious CPU cycles. By comparison, Weka only uses the resources that are allocated to it inside its LXC container, which means it can consume as little as one server core and a small amount of RAM in a shared environment (converged architecture) or as much as all of the resources of the server (a dedicated appliance). The same software stack is utilized in either case.

File System Design

From the outset, WekaFS was designed to solve many of the problems inherent with legacy scale-out NAS solutions. One of the key design considerations was to build a software platform that could address the requirements of different groups within an organization at scale or a multi-tenant environment. The most popular scale-out NAS file systems support a construct of a single file system and a single namespace, utilizing directories and a quota system to allocate resources and manage permissions. While this solution worked for smaller scale, it has made management complex when the number of users and/or directories scale. Full isolation of groups requires the creation of new file systems and namespaces, which then creates islands of physical storage to manage. Additionally, directory scaling is a problem and typically requires creating multiple directories to maintain performance, further exacerbating the complexity. As such WekaFS differs from other scale-out NAS solutions in that it embraces the concept of many file systems within the global namespace that share the same physical resources. Each file system has its own "persona" and can be configured to provide its own snapshot policies, tiering to object store, organizations, role based access control, quotas and much more. A Weka file system is a logical construct and unlike other solutions, the file system capacity can be changed on the fly. Clients that are mounted can observe the change in file system size right away without any need to pause I/O. As already mentioned, each file system has a choice to tier to an object store, and if it is a tiered file system, the ratio of hot (NVMe) tier and object (HDD) tier can also be changed on the fly. A file system can be split into multiple organizations managed by their own administrator. A single file system can support billions of directories and trillions of files, delivering a scalability model more akin to object stores than NAS systems, and directories scale with no loss in performance. More details can be found later. Currently WekaFS supports up to 1024 file systems, and up to 4096 snapshots in a single global namespace. This is an artificial limit which can be expanded.

WekaFS Limits:

- Up to 6.4 trillion files or directories
- Up to 8 Exabytes
- Up to 6.4 billion files in a directory
- Up to 4 petabytes for a single file

Supported Protocols

Clients with the appropriate credentials and privileges can create, modify, and read data using one of the following protocols:

- POSIX
- NFS (Network File System) v3
- SMB (Server Message Block) v2 and v3
- S3 via gateway
- NVIDIA® GPUDirect® Storage (GDS)⁴

Note: WekaFS does not support the HDFS protocol (Hadoop Distributed File System); however, Weka's POSIX connector can directly mount to a Hadoop node to provide much higher performance.

Data written to the file system from one protocol can be read via another one, so the data is fully shareable among applications.

Weka Storage Nodes

Weka storage nodes are created by installing WekaFS on any standard AMD EPYC™ or Intel Xeon™ Scalable Processor-based server

⁴ NVIDIA GPUDirect Storage is a protocol developed by NVIDIA to improve bandwidth and reduce latency between the NIC and GPU memory. It is currently in beta and available on select NVIDIA GPU-based systems.

with the appropriate memory, CPU processor, networking, and NVMe solid-state drives. A configuration of 8 storage nodes is required to create a cluster that can survive a two node failure. To create an appliance-like experience, Weka has worked with leading server vendors to create a single part number that can be used to order a complete storage system inclusive of software license. More details on specific configurations from Cisco, Dell, HPE, Hitachi Vantara, Lenovo, Penguin, and Supermicro as well as AWS are available on Weka's website.

For all other platforms, each storage node must conform to the following minimum specification outlined below in Tables 1 and 2:

COMPONENT	DESCRIPTION	MINIMUM QUANTITY
Processor (recommended)*	Intel Xeon Silver 4214R processor	2
Memory	DDR4-2666, 8GB DIMM	12
SSD (for boot device)	M.2 or U.2 SATA, 480GB or larger	1
SSD (for WekaFS)	U.2 NVMe, 960GB or larger	4
Network Interface Card (Ethernet or InfiniBand)	10Gb or above with DPDK support. 100Gb required for the referenced performance	1

Table 1: Minimum Hardware Specification for Intel Scalable Processor-Based Server

* WekaFS requires an x86-64 architecture, with a minimum of one processor with 2.2GHz and 12 cores.

COMPONENT	DESCRIPTION	MINIMUM QUANTITY
Processor	AMD EPYC 7402P processor	1
Memory	DDR4-2666, 16GB DIMM	8
SSD (for boot device)	M.2 or U.2 SATA, 480GB or larger	1
SSD (for WekaFS)	U.2 NVMe, 960GB or larger	4
Network Interface Card (Ethernet or InfiniBand)	10Gb or above with DPDK support. 100Gb required for the referenced performance	1

Table 2: Minimum Hardware Specification for AMD EPYC™ Scalable Processor-Based Server

INTEGRATED FLASH AND DISK LAYERS FOR HYBRID STORAGE

WekaFS consists of two separate layers, an NVMe SSD-based flash tier that provides high-performance file services to the application layer, and an S3-compatible object storage tier that manages the long-term data lake (outlined in Figure 3). The two tiers are physically separate but logically serve as one comprehensive system to the application layer. WekaFS expands the namespace from the flash tier to the object tier, presenting a single global namespace that scales with the object store. The object store can be from any S3-compliant vendor for either on-premises or public cloud deployment. WekaFS only requires the presence of an S3 bucket, so an existing object store can be used for the namespace extension as well as for supporting other applications. As we will see later, WekaFS leverages components of the object store capability to enable cloud-bursting, backup to the cloud, DR to another WekaFS-enabled storage cluster, or file system cloning.

Networking

The Weka system supports the following types of networking technologies:

- InfiniBand (IB) HDR and EDR
- Ethernet - recommended 100Gbit and above

The customer's available networking infrastructure dictates the choice between the two as WekaFS delivers comparable performance on either one. For networking, the Weka system does not use standard kernel-based TCP/IP services, but a proprietary infrastructure based on the following:

- Use of DPDK to map the network device in the user space and make use of the network device without any context switches and

with zero-copy access. This bypassing of the kernel stack eliminates the consumption of kernel resources for networking operations and can be scaled to run on multiple hosts. It applies to both backend and client hosts and enables the Weka system to fully saturate up to 200Gbit Ethernet or InfiniBand links.

- Implementation of a proprietary Weka protocol over UDP, i.e., the underlying network may involve routing between subnets or any other networking infrastructure that supports UDP. Clients can be on different subnets as long as they are routable to reach the storage servers. The use of DPDK delivers operations with high throughput and extremely low latency. Low latency is achieved by bypassing the kernel and sending and receiving packages directly from the NIC. High throughput is achieved because multiple cores in the same host can work in parallel, eliminating any common bottleneck.

On legacy platforms lacking support for virtualization (SR-IOV) and kernel offloading technologies (DPDK), the Weka system software relies on in-kernel processing and UDP as the transport protocol. This mode of operation is commonly referred to as the 'UDP mode' and is typically used with older hardware such as the Mellanox CX3 family of NICs.

In addition to being compatible with older platforms, the UDP mode yields CPU resources to other applications. This can be useful when the extra CPU cores are needed for other purposes. On InfiniBand, the protocol leverages RDMA-enabled environments to supply high performance without the need to dedicate cores to the Weka front-end processes.

Application clients connect to the Weka storage cluster via Ethernet or InfiniBand connections. The Weka software supports 10GbE, 25GbE, 40GbE, 50GbE, 100GbE Ethernet networks, and EDR and 200Gb HDR InfiniBand networks. For the best performance outlined in this document, Weka recommends using at least 100Gbit network links.

Note: Unlike RoCE implementations that require Priority-based Flow Control (PFC) to be configured in the switch fabric, Weka does not require a lossless network setting to support its NVMe-over-fabrics implementation, and can even deliver this level of low latency performance in the public cloud networks.

Network High Availability (HA)

The Weka file system supports high availability (HA) networking to ensure continued operation should a network interface card (NIC) or network switch fail. HA performs failover and failback for reliability and load balancing on both interfaces and is operational for both Ethernet and InfiniBand. For HA support, the Weka system must be configured with no single component representing a single point of failure. Multiple switches are required, and hosts must have a connection on each switch. HA for clients is achieved through the implementation of two network interfaces on the same client. WekaFS also supports the Link Aggregation Control Protocol (LACP) on the compute clients for a single dual-ported NIC.

Management GUI

Weka provides two quick and easy ways to manage the Weka file system, either through a Graphical User Interface (GUI) or a Command Line Interface (CLI). Reporting, visualization, and overall system management functions are accessible using the CLI or the intuitive GUI-driven management console (see Figure 5 below).




Figure 5 – WekaFS Management Software User Interface

Point-and-click simplicity allows users to rapidly provision new storage; create and expand file systems within a global namespace, establish tiering policy, data protection, encryption, authentication, permissions, NFS and SMB configuration, snapshots, snapshot-to-objects, and quality of service policies, as well as monitor overall system health. Detailed event logging provides users the ability to view system events and status over time or drill down into event details with point-in-time precision via the time-series graphing function (Figure 6).



Figure 6 – Time-Series Charts for Event Monitoring

A System Events tab lists the events that have occurred in a Weka environment. The events displayed in this window are also transmitted to the Weka Support Cloud (Weka Home) so that they can be used by WekaIO support to actively assist you when necessary or to proactively notify you when action needs to be taken. The WekaFS GUI is entirely cloud-based, eliminating the need to physically install and maintain any software resources, and you always have access to the latest management console features. For customers who wish to use their own monitoring systems.

3rd Party Monitoring Tools

Weka supports integration with other monitoring and management tools for customers who wish to use their own monitoring systems. Weka currently supports integration with Prometheus and Grafana.

Adaptive Caching

Applications, particularly those with small files and lots of metadata calls, benefit greatly from local caches. The data is available with very low latency and it reduces the load on the shared network as well as on the back-end storage itself. The Weka file system provides a unique advanced caching capability, called adaptive caching, that allows users to fully leverage the performance advantages of Linux data caching (page cache) and metadata caching (dentry cache) while ensuring full coherency across the shared storage cluster. NFS v3 does not support coherency so utilizing Linux caching can lead to data inconsistency for read cache and potential data corruption for write cache. WekaFS supports leveraging Linux page cache – typically reserved for direct attached storage (DAS) or file services run over block storage – on a shared networked file system while maintaining full data consistency. The intelligent adaptive caching feature will proactively inform any client that was an exclusive user of a file (and hence running in local cache mode) that another client now has access to the data set. Once this flag is set, the client can continue running in local cache mode until the file is modified by another client. WekaFS will now invalidate local cache ensuring that both clients are only accessing the most recent iteration of the data. This ensures the highest performance from local cache when appropriate and full coherency on data at all times. WekaFS does not require specific mount options to leverage local page cache as the system dynamically manages caching, making the provisioning of the Weka environment very simple to manage with no danger of an administrative error causing data corruption.

WekaFS provides the same capability for metadata caching, also known as Linux dentry cache. A client can leverage local metadata cache for a directory, reducing latency significantly. However, once another client has access to the same directory, Weka will ensure that any directory changes from one client will invalidate the cached metadata for all other clients accessing that directory. The caching capability also includes extended attributes and access control lists (ACLs).

While some shared file storage vendors allow local caching, no other file system provides the adaptive caching capability. Caching is typically disabled by default and requires an administrator to change the mount option. That is because write coherency typically depends on some form of battery backup protection on the client to ensure data consistency on a committed write. Weka's caching implementation will work out-of-the-box without any administrator intervention as it does not depend on battery protection. As a result, the same software that runs on-premises can be seamlessly deployed in the public cloud with no software changes. This feature is ideal

⁵ Note: An on-premises flash tier, with a cloud-based object tier would suffer the performance penalty of a WAN connection and the ingress/egress costs of cloud.

for use cases such as file “Untar” which will run significantly faster as a local process vs. across a shared file system.

Global Namespace and Expansion

WekaFS manages all data within the storage system as part of a global namespace and supports two persistent storage price tiers in a single hybrid architecture – NVMe SSD for active data and HDD-based object storage for a data lake. Expanding the namespace to object store is an optional addition to the global namespace and can be configured during installation, or any time after, with a few clicks from the management console. A file resides on flash while it is active or until it is tiered off to object storage based on preset or user-defined policies. When a file is tiered to HDD, the original file is kept on the flash layer until the physical space is required by new data, and hence acts as a cached file until overwritten. While file data is demoted to HDD, file metadata always remains locally on the flash tier, so all files are available to the applications, irrespective of physical placement, even if the bucket was in the public cloud⁵. As NVMe flash system capacity runs low and usage reaches a high watermark, data is dynamically pushed to the object tier, which means you never have to worry about running out of capacity on the flash tier. This is particularly useful for write-intensive applications as no administrator intervention is required. The flash tier and the HDD-based object tier can scale independently depending on the required usage capacities.

The global namespace can be sub-divided into 1024 file systems and file system capacity can be expanded at any time simply by allocating more space to it. By segmenting the namespace, storage capacity can be allocated to individual users, projects, customers, or any other parameter, yet be easily and centrally managed. Data within a file system is fully isolated from every other file system to prevent noisy neighbor issues.

Tiering

WekaFS has a built-in, policy-based automated tiering feature that transparently moves data across storage tiers. Weka supports moving data from the NVMe flash storage tier to on-premises or cloud-based object storage (Figure 7). Tiering is set at the file system level and is optional. To ensure the highest performance at all times, users may want to keep certain file systems exclusively on NVMe SSD, while other file systems implement tiering to S3 object storage for the best cost dynamics. Metadata is always stored on the flash tier, and a snapshot of the file system can be stored on the object storage tier to protect against a failure on the flash tier. The application clients see all the files in a given file system regardless of their placement, thus no change in application is needed to leverage cost-optimized object storage-based solutions.

This integrated capability eliminates the need for additional Hierarchical Storage Management (HSM) or tiering software that adds complexity and cost. The by-product of the integrated tiering is an elastic unified namespace that can scale to the size of your cloud vendor’s capacity limits.

Maximum tiering granularity is in 64MB range (chunks) of data, and smaller files are “packed” into larger chunks to ensure the best performance between the Flash tier and object tier. WekaFS supports byte-range reads so small file reads do not require re-hydration of the entire 64MB chunk, rather just of the data that is being called on by the file system. By chunking data into 64MB sizes, WekaFS provides parallel and concurrent access to large files between the object tier and the flash tier. So in essence, WekaFS has two parallel file systems, an all-flash solution to the clients and a hard disk-based solution between the flash tier and the object tier. A secondary benefit is that only the data range that is required by an application has to be retrieved from the object store, which significantly reduces data network traffic for large files and improves read performance for partial files.

There is no hard and fast rule on how much data should be on NVMe flash versus HDD, but a survey of Weka’s customers shows that the typical distribution is now around 20% on flash. The rule of thumb is to measure the flash tier such that it will hold the normal working data sets for current workloads and enough capacity to pre-stage new workloads for maximum performance.

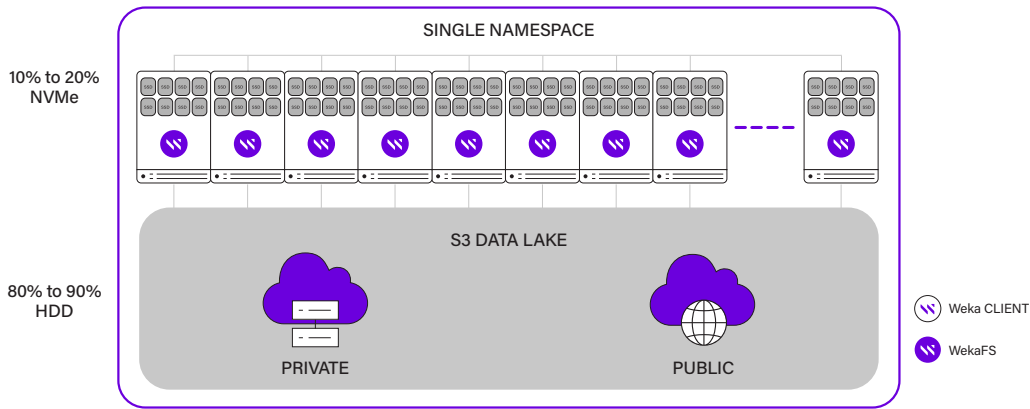


Figure 7 – Integrated Tiering to any S3 or Swift-Compatible Object Store

Snapshots and Clones

Weka supports user-definable snapshots for routine data protection including backup as well as for cloud data migration and cloud bursting. For example, WekaFS snapshots can be used to back up files locally on the flash tier as well as making copies to cloud storage tiers for backup or disaster recovery. Also, Weka snapshots can be utilized for saving (parking) files and data to lower-cost cold storage such as public cloud and on-premises object storage. In addition to point-in-time snapshots, WekaFS can create full clones (snapshots that can be converted into writable snapshots) with pointers back to the originating data. WekaFS snapshots and clones occur instantaneously and are differential after the first instance, which dramatically reduces the time and storage required for protection. Furthermore, overall system performance is unaffected by the snapshot process or when writing to a clone. Each file system has its own snapshots and a single WekaFS cluster supports up to 4096 snapshots.

Snap-to-Object

Once Tiering is enabled in a file system, WekaFS supports a unique feature which we call snap-to-object. This feature enables the committing of all the data of a specific snapshot, including metadata, to an object store. Unlike data lifecycle management processes, this feature involves all the contents of the snapshot, including all files and metadata. After the first snapshot-to-object has been completed, subsequent snapshots are stored in a differential manner so backup time is limited to the changes and should be very fast.

The outcome of using the snap-to-object feature is that the object store contains a full copy of the snapshot of the data, which can be used to restore the data on the original Weka cluster or on another cluster. The secondary cluster that mounts the WekaFS snap-to-object snapshot does not need to be a mirror of the primary system. In fact the primary system could have 20 storage nodes in the cluster while the second system could have 8 or even 100 nodes, this makes it ideal for cloud bursting. Consequently, the snap-to-object feature is useful for a range of use cases, as follows:

1. Generic Use Cases (on-premises and cloud)

- Remote backup of data to an on-premises or cloud-based object store: If too many hardware components in a Weka NVMe cluster fail beyond recovery because of a failure of the system or an external event such as a fire, the snapshot saved to the remote object store can be used to re-create the same data on another Weka cluster, or re-hydrate on the original cluster.
- Data archival: The periodic creation of data snapshots, followed by uploading the snapshot to an object store or the cloud to create an archive copy of the data.
- Asynchronous mirroring of data: Combining a Weka NVMe cluster with a replicated object store in another data center will create a mirror of the data that can be mounted on a second Weka cluster.

2. Cloud-Only Use Cases

- Public Cloud pause and restart: In the AWS cloud, Weka utilizes instances with local SSDs to create a cluster. Sometimes for bursty project-specific work, users may want to shut down the cluster to save costs. The snapshot can be saved on the AWS S3 object store and re-hydrated when needed again at a later time.
- AWS protection against single availability zone failure: Utilizing the snap-to-object feature allows users in AWS to recover from an availability zone (AZ) failure. Should the first AZ fail, the Weka snapshot can be replicated to a second AZ and re-hydrated in minutes by a cluster in the secondary AZ.

3. Hybrid Cloud Use Case

⁶ RAID = Redundant Array of Independent Disks

- Cloud bursting: An on-premises customer can benefit from cloud elasticity to consume additional computational power for short periods. By uploading a snapshot to S3, the file system can be run in the cloud. The data can be deleted on the shutdown of the instances.

Data Protection

Data protection is a critical function of any storage system, and challenges are amplified at scale. Without an appropriate protection schema, file systems would need to be limited in size to accommodate the rebuild time window and minimize the risk of data exposure. Popular data protection schemes such as RAID⁶, replication, and erasure coding are a compromise between scalability, protection, capacity, and performance.

With Weka, there is no concept of data or metadata locality as all data and metadata are distributed evenly across the storage nodes, which improves the scalability, aggregate performance and resiliency. On high-speed networks, data locality actually contributes to performance and reliability issues by creating data hot spots and system scalability issues. By directly managing data placement on the SSD layer, WekaFS can shard the data and distribute it across the storage cluster for optimal placement based on user-configurable stripe sizes. WekaFS uses advanced algorithms to determine data layout; sharded data perfectly matches the block sizes used by the underlying flash memory to improve performance and extend SSD service life. Stripe sizes can be set to any value from 4 to 16, while parity can be set to either +2 or +4. Figure 8 illustrates data placement across SSDs in a 4+2 configuration. The minimum recommended cluster size is 8, which allows for two full virtual spares for a rebuild from a 4+2 configuration. The bigger the Weka cluster, the bigger the stripe size that it can support, and the greater the storage efficiency and write performance.

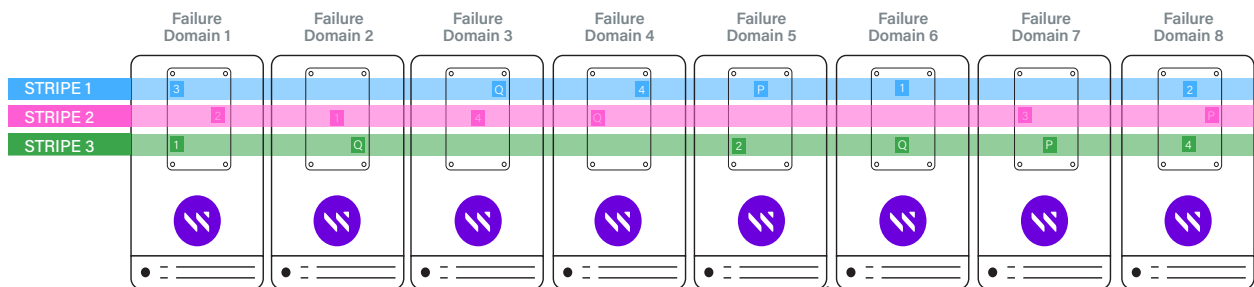


Figure 8: WekaFS data distribution

WekaFS Data Protection Schema

WekaFS manages protection so data is always safe and accessible:

- Configurable data protection levels from 4+2 to 16+4
- Patented distributed data protection schema
- Configurable failure domains
- End-to-end data protection checksums for data integrity
- Metadata journaling
- Network redundancy support for two top-of-the-rack switches over two network ports
- Local or cloud-based snapshots and clones
- Snapshot-to-object for backup and DR
- Automated tiering to the cloud

Weka uses failure domains to define data protection levels. Failure domains are fully configurable starting at the server node level, which provides single or multiple SSD level granularity. Data protection levels are flexible depending on the size and scale of the server cluster—the larger the cluster, the larger the recommended data stripe size for the best utilization of SSD capacity, improved performance, and higher resiliency. For granular protection, the data protection level is set at the cluster level and parity can be set to two or four, meaning that the system can survive up to two or four simultaneous node failures without impacting data availability.

Weka's data protection scheme follows the convention of data (N) + parity (2 to 4), and the N + 4 data protection level is unique in cloud

⁷ United States Patent 9448887 – <http://www.freepatentsonline.com/9448887.html> to learn more

storage. Weka's data protection scheme delivers exponentially better resiliency than triple replication – which only protects to 2 failures – without the expensive storage and throughput impact. An $N + 2$ protection level is sufficient for most production environments, whether with converged clusters or dedicated appliances. An $N+4$ protection level is recommended for clusters with a large number (hundreds) of converged cluster nodes because applications can impact server availability.

Data Distribution

WekaIO utilizes a patented⁷ distributed data protection coding scheme that actually increases resiliency as the number of the nodes in the cluster scale. It delivers the scalability and durability of erasure coding but without the performance penalty. Unlike legacy hardware and software RAID and other data protection schemes, Weka's rebuild time actually gets faster and more resilient as the system scales because every node in the cluster participates in the rebuild process.

WekaFS equally distributes data and metadata across logical nodes that span failure domains (FD). A failure domain can be a server node, a rack, or even a data center. In cloud environments, WekaFS can span availability zones (AZ).

Unlike traditional hardware and software data protection schemes, Weka only places a single data stripe inside any one server (or FD), so in the event of two drive failures within a single server it will still be considered a single failure. Even if the node has multiple SSDs, data stripes are always spread across different server nodes, racks, or AZs depending on the resiliency chosen. Weka's resiliency is user-configurable to define the number of faults within a Weka cluster to meet the application workload service level requirements. When a failure occurs, the system considers the FD a single failure, regardless of how large the domain is defined. In addition to distributing stripes at the FD level, Weka also ensures a highly randomized data placement for improved performance and resiliency. As the cluster size grows the probability of a hardware failure goes up proportionally, but WekaFS overcomes this issue by distributing the stripes in a randomized manner. The more nodes the higher the amount of random stripe combinations, making the probability of a double failure lower. Example: for a stripe size of 18 (16+2) and a cluster size of 20 the number of possible stripe combinations is 190, however as the cluster size grows to 25, the number of possible stripe combinations is now 480,700.

WekaFS Rebuilds

Weka uses several innovative strategies to return the system to a fully protected state as quickly as possible and be ready to handle a subsequent failure. This ensures that applications are not impacted by long rebuild processes.

Weka protects data at the file level, so it only needs to rebuild the data that is actively stored on the failed server or SSD. This means that the rebuild times are faster compared to a traditional RAID solution or file server that protects data at the block layer. RAID controller-based systems typically rebuild all blocks on an affected device, including empty blocks, prolonging rebuilds and the time of exposure. Weka only needs to rebuild the data that has been affected by the failure. A further benefit of the Weka tiering policy is that data that has already been tiered off to an object store is never impacted by a node failure because it is protected on the object store. In addition, any cached data (data that was tiered to object but still remains on the flash tier functioning as a cache) does not need to be rebuilt either, limiting the rebuild priority to data that only resides on the flash tier.

Since all of the server nodes (failure domains) share all the stripes, a failed component has stripes shared on the remaining good nodes; hence all the healthy nodes participate in the recovery process. This means that the larger the cluster size, the faster the rebuild and the more reliable the system becomes because more compute resources are available to participate in the rebuild process and the more randomized the stripes. In the event of multiple failures, the system prioritizes data rebuilds starting with data stripes that are least protected. Weka looks for data stripes that are common to the failed nodes and rebuilds these data stripes first so the system can be returned to the next level higher of resiliency as fast as possible. This prioritized rebuild process continues until the system is returned to full redundancy. By contrast, in a replicated system only the mirrored servers participate in the recovery process, impacting application performance significantly. Erasure coding suffers from a similar problem where only a small subset of the servers participates in the recovery. With WekaFS, the recovery rate is user-configurable and the amount of network traffic dedicated to rebuild can be changed, can be changed at any time, so administrators have complete control to determine the best tradeoff between continued application performance and time-to-recovery.

Power-Fail and End-to-End Data Protection

Using a checksum process to ensure data consistency, WekaFS provides end-to-end data protection for both reads and writes. Weka always stores data and checksum information separately from each other on different physical media for improved protection.

WekaFS provides additional data integrity capabilities by protecting against data loss due to power failures. When a write is acknowledged

back to the application, it is safely protected from server failures or data-center-wide power failure through a journaling process. Weka's innovative data layout and algorithms enable it to recover from a data-center-wide power failure in minutes because there is no need to do a complete file system consistency check (FSCK). For most other file systems, the FSCK process recovery time is proportional to the size of the recovered file system. In large scale deployments, this recovery can take days.

Automated Data Rebalancing

WekaFS proactively monitors and manages the performance, availability, and capacity health status of a Weka cluster. This allows the system to calculate the utilization levels (performance and capacity) of nodes to automatically and transparently redistribute data across the cluster to prevent hot spots.

The benefit is that Weka can maintain well-balanced cluster performance and data protection as capacity and usage change. Another advantage is that as additional SSDs are added to existing server nodes or the cluster is expanded with more nodes, Weka automatically rebalances to enhance performance, resiliency, and capacity without requiring costly downtime for data migration. Matched capacity SSDs are not required, which allows you to leverage new technology and save money as SSD prices decline.

Container Storage Integration (CSI)

CSI is a standard that has been developed to provision and manage shared file storage for containerized workloads. WekaFS CSI plugin for Kubernetes provides an interface between the logical volumes in a Kubernetes environment (known as Persistent Volumes) and the storage. The CSI plugin provisions a Kubernetes pod volume either via Persistent Volume (by an administrator) or it can be dynamically provisioned via a Persistent Volume Claim (PVC). This feature simplifies the process of moving containerized workloads to the cloud or sharing across multiple clusters.

Multi-Tenant Organizations

WekaFS supports the construct of organizations, an element of multi-tenancy. Groups can be separated such that data is managed at the organizational level and only visible to that group's members. The Weka file system can support up to 64 organizations, and within an organization, logical entities participating in obtaining control of that data are managed by the Organization Admin, not the Cluster Admin. The Cluster Admin can create organizations, define the Organizational Admin, delete organizations, and monitor capacity usage by the organization file system. Cluster Admins have root access to the file system, but they are obscured from the organizational data in the Weka file system.

Authentication and Access Control

WekaFS provides authentication services at the user level and the client-server level to validate that the user or client has the security levels to view and access data. WekaFS allows different authenticated mount modes such as read-only, or read-write and is defined at the file system level. WekaFS supports the following:

- LDAP (Lightweight Directory Access Protocol), a networking protocol that provides directory services across many different platforms.
- Active Directory, a Microsoft implementation of LDAP, a directory service that can store information about the network resources. It is primarily used to authenticate users and groups who want to join the cluster.
- WekaFS also offers Role-Based Access Control (RBAC), delivering different levels of privileges to users and administrators. Some users can be granted full access rights while others have read-only rights.

Every Weka system user has one of the following defined roles:

Cluster Admin: A user with additional privileges over regular users. These include the ability to:

- Create new users
- Delete existing users
- Change user passwords
- Set user roles
- Manage LDAP configurations
- Manage organizations

Additionally, the following restrictions are implemented for Cluster Admin users to avoid situations where a Cluster Admin loses access to a Weka system cluster:

- Cluster Admins cannot delete themselves
- Cluster Admins cannot change their role to a regular user role

Organization Admin: A user who has similar privileges to cluster admins, except that these privileges are limited to the organization level. They can perform the following within their organization:

- Create new users
- Delete existing users
- Change user passwords
- Set user roles
- Manage the organization LDAP configuration

Furthermore, to avoid situations where an organization admin loses access to a Weka system cluster, the following restrictions are implemented for organization admins:

- Organizational Admins cannot delete themselves
- Organizational Admins cannot change their role to a regular user role

Regular: A user with read and write privileges

Read-only: A user with read-only privileges

In-Flight and At-Rest Encryption

WekaFS provides full end-to-end encryption from the client all the way to the object storage solution, making it the most robust encrypted file system commercially available. Encryption is set at the file system level upon creation of the file system, so some file systems that are deemed critical can be encrypted while others are not. Weka's encryption solution protects against physical media theft, low-level firmware hacking on the SSD, and packet eavesdropping on the network. File data is encrypted with XTS-AES using 512-bit keys, which provides 256 bits of effective security.

Weka has demonstrated that encrypted file systems have minimal impact on application performance when using the Weka client.

Key Rotation and Key Management

WekaFS supports any key management system (KMS) that is compliant with KMIP (the Key Management Interoperability Protocol), as well as Hashicorp Vault proprietary API. Cluster keys are rotated by the KMS, file system keys can be rotated via the KMS and re-encrypted with the new KMS master key, and file keys can be rotated by copying the file.

File data on the object store is also encrypted. When uploading a snap-to-object to the object store, among other file system parameters, the file system key is included and encrypted with a special "backup-specific" cluster key that is available via the KMS and is used for all snap-to-object backups and restores. When WekaFS pushes a snapshot to the AWS cloud, the data is fully protected and can be authenticated only through the on-premises KMS system.

FLEXIBLE DEPLOYMENT OPTIONS (ON-PREMISES AND CLOUD)

Whether your applications run on bare-metal for performance, in a virtual environment for ease of deployment and resiliency, or entirely in the public cloud for on-demand scalability and elasticity, Weka is a single, no compromise, storage solution providing the freedom to choose the environment best suited for your application based on performance, scale, and economics. WekaFS supports running on bare metal as well as in virtualized, containerized, and cloud environments (Figure 9).

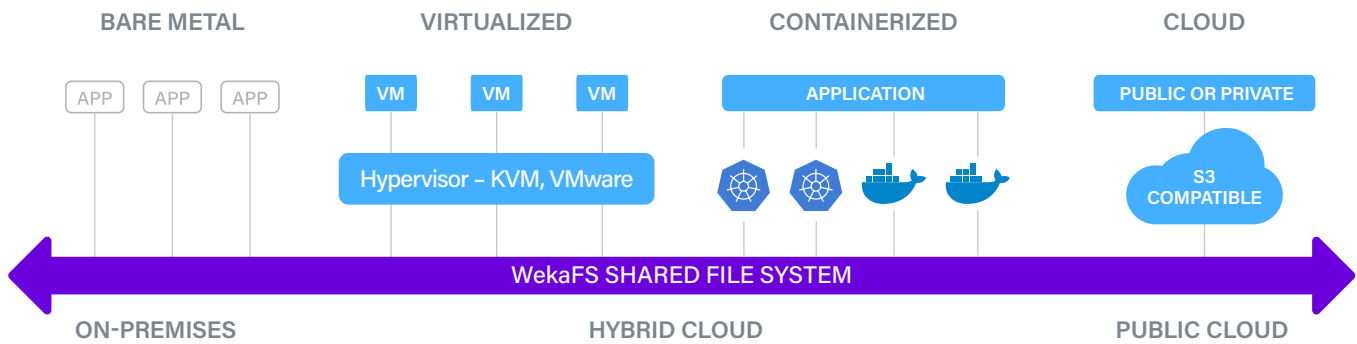


Figure 9 – Industry-Leading Flexibility to Support Your Operating Environment

Weka provides industry-leading flexibility by supporting an extremely broad array of operating environments and deployment models. It runs on standard x86-based servers using Ethernet or InfiniBand network adapters and off-the-shelf NVMe SSDs.

Starting at only eight servers, Weka can scale to many thousands of servers. When infrastructure matures, Weka enables expanding the cluster with new servers, then retiring older generation servers, thus enabling continuity over hardware refresh generations without the need to perform complete lift-and-shift data migration to different storage.

Converged

In a converged deployment (Figure 10), Weka is integrated into standard AMD or Intel x86-based application servers, or public cloud instances. Combining storage and compute resources with applications and data into a single building block delivers a highly optimized data center solution across a mix of workloads. This deployment model provides the ability to create integrated application-based solutions at better economics, by eliminating the need for external storage, reducing hardware footprint and power while boosting performance, availability, and capacity. WekaFS carves out storage resources and only utilizes the resources in its container, with the remaining resources available to the applications. This deployment model is popular for clusters of GPU servers which have local NVMe devices.

The following (Figure 10) shows WekaFS in converged mode.

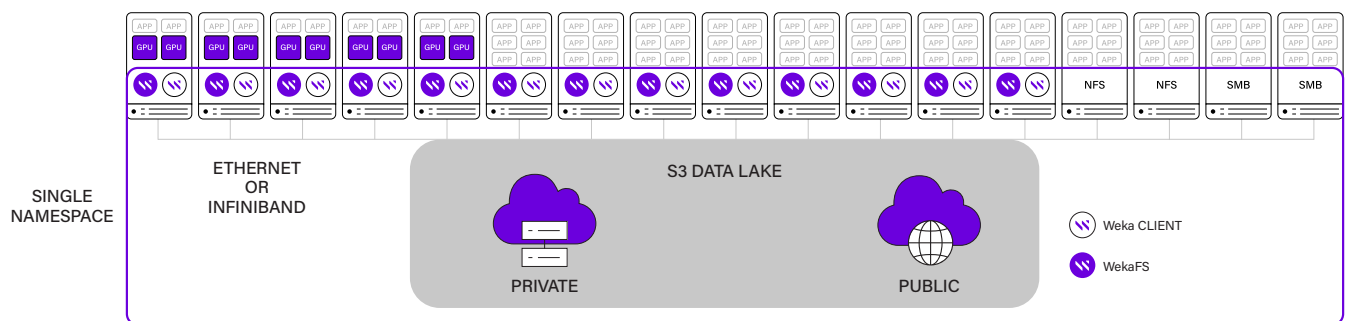


Figure 10: WekaFS in Converged Mode With Compute and Storage on the Same Infrastructure

Dedicated Storage Server

Weka can be configured as a dedicated storage server (Figure 11) when all the resources of the system are dedicated to storage services while applications run on a separate compute infrastructure. This mode is the most popular among customers as it ensures that application disruptions do not impact storage services.

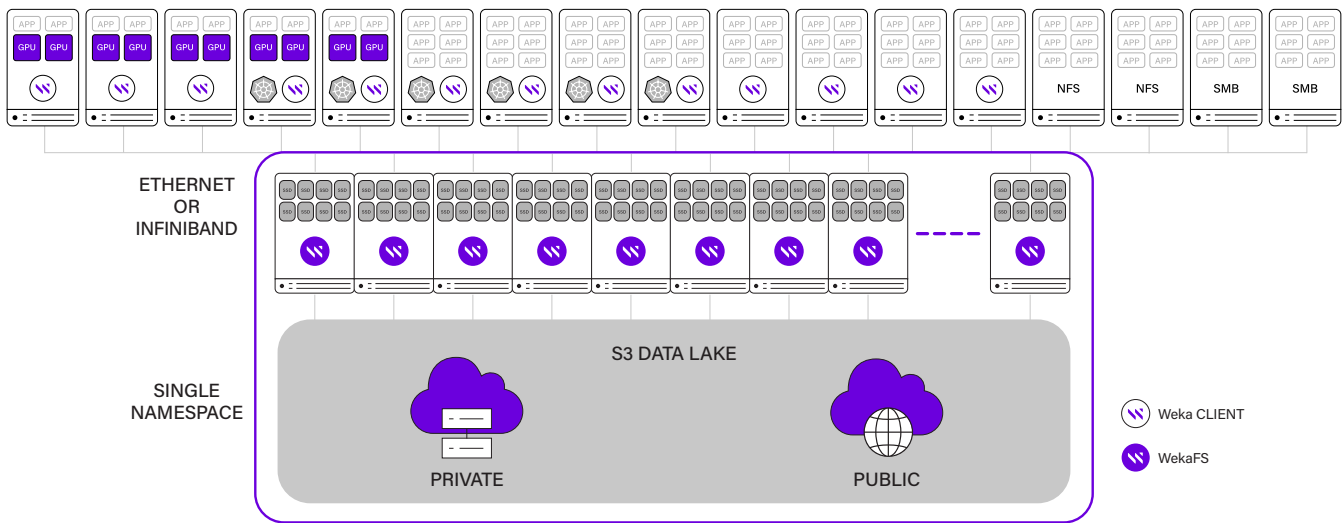


Figure 11 – Dedicated Storage Server with Separate Compute and Storage Infrastructure

Native Public Cloud

Deploy Weka on Amazon just as you would on-premises using EC2 instances with locally attached SSD(s). Weka supports both converged and dedicated storage deployment models in the cloud on a selection of EC2 instance types. For compute-intensive workloads, or to take advantage of GPU-based instances, it is best to use the dedicated storage server mode. If you need additional low-cost, high-capacity storage, leverage Weka automated tiering to Amazon S3. If you want to burst or migrate data to the cloud, you can leverage Weka snapshot-to-object functionality for data movement. For improved resiliency, Weka also supports the spanning of AZs in the public cloud. Figure 12 shows a typical deployment of WekaFS in the AWS public cloud.

For best performance Weka recommends i3en instances with 100Gbit networking and NVMe local storage.

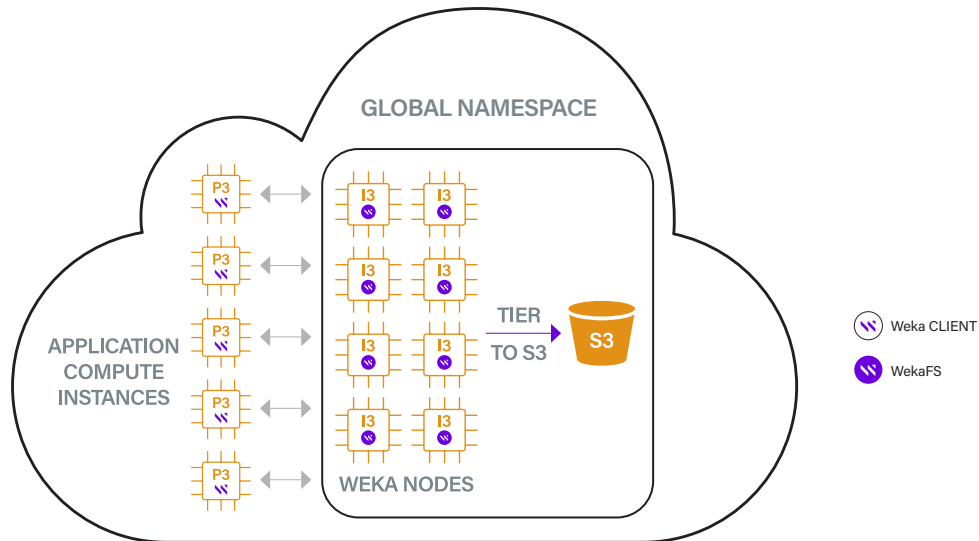


Figure 12 – WekaFS in the AWS Public Cloud

WekaFS PERFORMANCE PROOF POINTS

Until now, IT organizations could only achieve their performance goals for file-based applications by running them on a local file system — also known as direct attach storage (DAS) — or provisioning them over an All-Flash Array (AFA) volume. If shared access to the data set is required by multiple clients, performance is severely compromised because of the performance limitations of Linux NFS or

Windows SMB. Weka delivers important performance benefits over this approach. Our shared POSIX-compliant file system leverages an innovative protocol to deliver file-based semantics that have all the local caching advantages of direct-attach storage and the shareability advantages of NFS or SMB. As a result, Weka allows you to run any workload on a networked shared file system with as much or more performance than DAS can offer. This, coupled with extremely low latency 4K I/O operations to the application, make the need to use complex AFA volumes for a local file system obsolete. The following chart (Figure 13) outlines the actual performance of WekaFS in a production environment with 100Gbit networking, versus All-flash NAS and direct-attached local storage. WekaFS was 3x faster than local storage and 10x faster than NFS-based all-flash NAS. Doubling the network bandwidth to 200Gbit would result in 2x the client throughput, delivering 6x improvement in local file system performance and 20x the performance of NFS.




Figure 13: WekaFS Performance vs. Local NVMe and All-Flash NAS

The Weka performance has been documented publicly on several industry-validated benchmarks. WekaFS set the record for the highest performing file system at Supercomputing 2019, winning both the number 1 overall position as well as the number 1 best metadata performance. WekaFS has set an industry record on SPEC 2014 across every major benchmark in the suite and WekaFS has set multiple records on STAC Council M3 benchmark for financial analytics. Further proof points are outlined at the end of this document.

This document has outlined the architectural capabilities of the Weka file system and the following section provides proof points to demonstrate the performance, scaling, and latency delivered with WekaFS. Select examples are extrapolated from public benchmarks available on industry-validated performance benchmarking institutions and publicly available reference architectures.

Standard Performance Evaluation Corporation: SPEC.org

The **SPEC SFS® 2014** benchmark is the latest version of the Standard Performance Evaluation Corporation benchmark suite measuring file server throughput and response time, providing a standardized method for comparing performance across different vendor platforms.

WekaFS holds the leadership position for databases, Virtual Desktop Infrastructure (VDI), Electronic Design Automation (EDA), and Video Data Acquisition (VDA), with a close second for software builds.

The following figure (Figure 14) demonstrates Weka's performance for databases. WekaFS supported 4480 databases, 2x more than the next closest submission, at 340-microsecond overall response time.

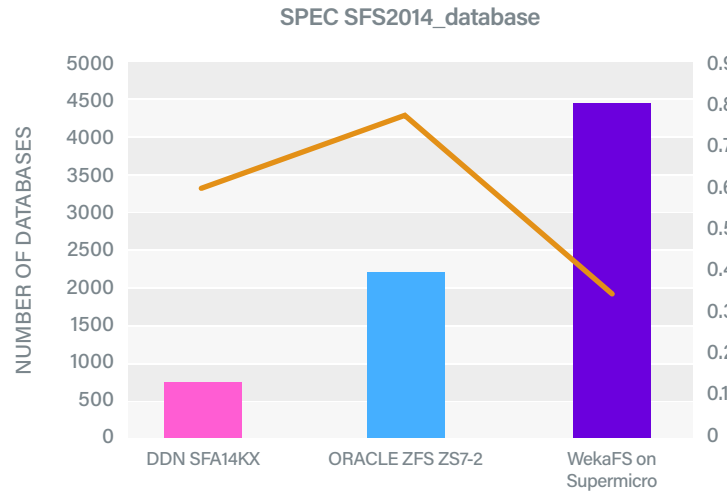


Figure 14: WekaFS Performance for Databases

The next figure (Figure 15) demonstrates the ultra-low latency of WekaFS. While the NetApp AFF A800 supported 8% more builds than Weka, the latency for WekaFS was 260 microseconds compared to 830 microseconds for NetApp. In other words, WekaFS can complete over 3x as many software builds as NetApp in the same time.




Figure 15: WekaFS supports 5700 builds at 260-microsecond latency

Performance Scaling

To demonstrate the linear performance scaling of WekaFS, we executed testing on Amazon Web Services to show how WekaFS scales from a minimum configuration of 8 nodes to 240 storage nodes. The following graph (Figure 16) demonstrates the perfect linear scaling measured on WekaFS. The performance was measured over 10GB network links so it is not representative of the performance that WekaFS can deliver on 100Gbit or 200Gbit InfiniBand networks. Rather, these numbers should be looked at for the linear scalability of the file system.




Figure 16: Perfect Linear Performance Scaling as a Storage Cluster Grows

To understand the performance capability of WekaFS in the AWS public cloud, the following graphs (Figure 17) demonstrate the 4K IOPs and bandwidth performance of WekaFS on i3 instances with 100Gbit networking. WekaFS achieved over 100GB/sec and 5 million IOPs at 250-microsecond latency with only 16 storage nodes. The flat latency graph shows that there is no penalty for scaling, and doubling the storage server resources results in twice the I/O performance.




Figure 17: WekaFS Performance on AWS

Performance to GPU Storage

WekaFS is an ideal file system for GPU-intensive workloads. Weka developed a reference architecture for the NVIDIA® DGX-1™ GPU System. FIO testing provides a baseline measure of the I/O capability of the reference architecture. The performance test was conducted with a single DGX-1 system to establish the performance that WekaFS could deliver with the minimum hardware configuration on a single 100-Gbit InfiniBand link to the host. Figure 18 shows that WekaFS is capable of fully saturating a 100-Gbit link, delivering a peak read performance of 10.8 GBytes/second to a single DGX-1 system. The IOPs performance measurement shows that WekaFS delivered over 250,000 IOPs to a single DGX-1 system on one 100-Gbit network link.




Figure 18: Performance to a Single NVIDIA DGX-1 GPU System over a 100Gbit Network Link

The next set of performance testing measures Weka’s ability to scale from a single GPU server to multiple GPU servers. The following figure (Figure 19) demonstrates how WekaFS maintained perfect linear scaling from 1 NVIDIA DGX-1 to 9 NVIDIA DGX-1 systems. Each NVIDIA DGX-1 had 8 GPUs, for a total scaling to 72 Tesla V100 GPUs.




Figure 19: Performance Scaling from a Single GPU System to 9 GPU Systems

The final set of performance tests (Figure 20) measures Weka’s ability to deliver performance to a single GPU server utilizing GPUDirect Storage. WekaFS scaled performance inside a single NVIDIA DGX-2™ GPU server from one 100Gbit EDR link to 8 EDR links with perfect linear scaling, saturating the network bandwidth.




Figure 20: Performance Scaling Inside a Single NVIDIA DGX-2 GPU Server with GPUDirect Storage

SUMMARY

WekaFS is a POSIX-compliant high-performance clustered, parallel file system that has been built from the ground up to run natively on NVMe based storage. It leverages high-performance networking – either Ethernet or InfiniBand – to fully saturate the network links for maximum performance. It is an ideal solution for performance-intensive applications that demand high I/O and high concurrency to multiple clients. It is widely deployed in Life Sciences, Financial Analytics, GPU-based DL and AI applications, EDA, HPC, and other large bandwidth and I/O-intensive applications that rely on parallel file systems. Weka reduces the cost and complexity of storage, requiring fewer hardware resources compared to traditional solutions. It also fully supports legacy protocols such as NFS and SMB and has a rich set of enterprise-class features.

Weka addresses common IT storage problems by providing a fast, efficient, and resilient distributed parallel file system that is cloud-native and delivers the performance of All-Flash Arrays, the simplicity of file storage, and the scalability of the cloud. Part of Weka's ease of use and cloud-like experience includes rapid provisioning to reduce time to get new workloads deployed, along with elasticity scaling, resiliency, performance, and cost-effectiveness.

To learn more visit www.weka.io.