

Wasabi Delivers High Data Durability & Protection Capabilities

Wasabi is the hot cloud storage company delivering disruptive storage technology that is 1/5th the price of Amazon S3 and faster than the competition with no fees for egress or API requests. You can use Wasabi hot cloud storage for a variety of purposes including primary storage for on-premises or cloud-based workloads, secondary storage for backup and recovery, or archival storage for long-term data retention or on-prem storage offload.

In addition to Wasabi's low price and high-performance capabilities, Wasabi provides a highly durable and reliable storage infrastructure, engineered to preserve data integrity, and ensure high service availability. In the storage industry, durability refers to the ability of a storage platform to provide long-term protection against disk drive failures, bit rot, degradation, or other corruption.

Durability is generally expressed as an annual percentage rate, approaching 100%. The closer to 100%, the greater the durability of the storage platform, and the less likely you are to lose data due to drive failures, bit rot or media corruption.

Historically, first-generation cloud storage vendors achieved high durability by creating two additional copies of the data and storing them in diverse locations. While this method of mirroring data across multiple locations provides durability, it also drives up the cost of data storage because it consumes three times the storage capacity needed for the original data. In recent years, new technologies and methods have been developed to deliver the same durability as mirroring while reducing the cost associated with mirrored data capacity. With modern public cloud object storage services such as Wasabi, the use of disk striping with parity delivers high data durability and increased performance. Because data and parity are striped evenly across multiple disks, no single disk is a bottleneck. Striping allows Wasabi to reconstruct data in case of a disk failure.

Wasabi Delivers High Data Durability

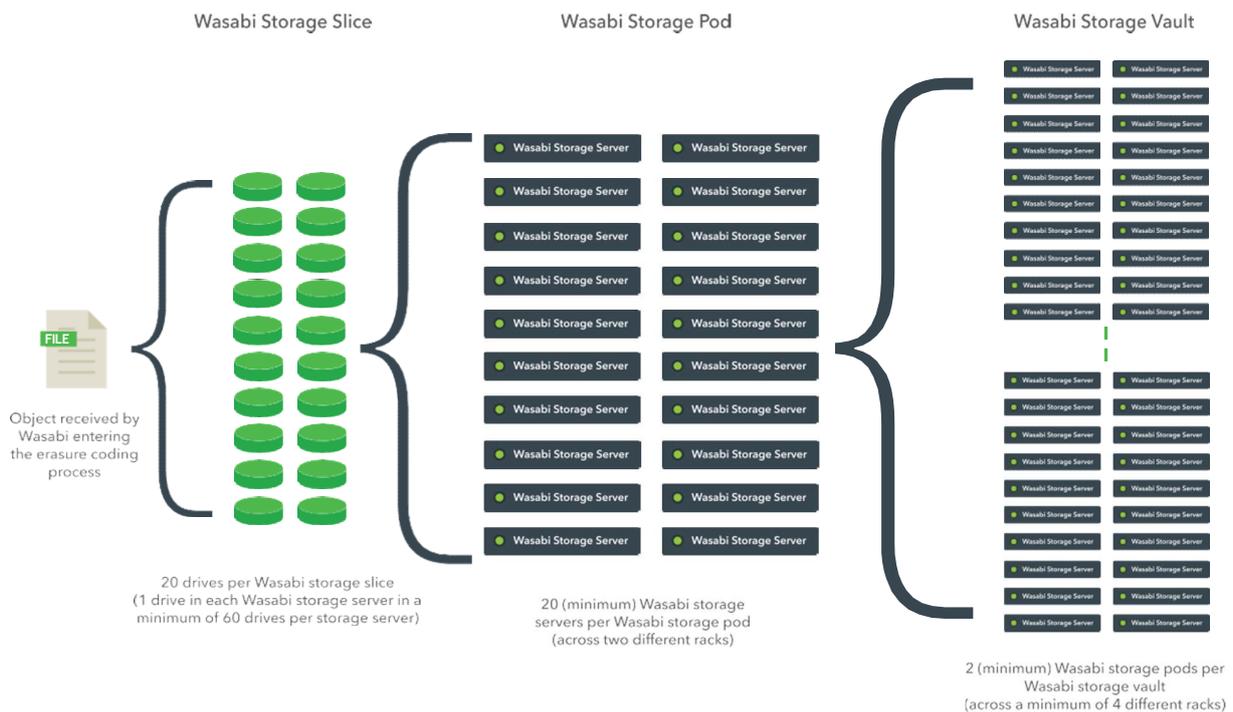
Wasabi uses an advanced, industry-proven process of dividing data into blocks and spreading the data blocks across multiple storage devices, such as a cluster of hard disks or solid-state drives (SSDs) to protect data against drive failures and associated media errors. This allows Wasabi to provide 11 x 9s of data durability, fully protecting customer data without the added capacity requirements of data mirroring over approaches. By data striping, Wasabi can make far more effective use of storage capacity. Wasabi passes the cost savings from efficiencies along to customers in the form of low pricing.

Wasabi Delivers High Data Durability & Protection Capabilities

Wasabi Distributes Data Across 20 Different Drives for High Resiliency

To better understand how Wasabi achieves 11 x 9s of data durability, let's first examine the storage building blocks in Wasabi's system architecture. The most basic element in any storage platform is a disk drive. Wasabi uses 14 drives to form an element known as a Wasabi storage pool. Data is encoded for redundancy across the 14 drives ensuring losses of multiple drives does not impact data integrity. Multiple of these pools are housed in an enclosure, along with spare drives for quickly replacing drives that fail. Each enclosure is connected to 2 different servers through different data paths, and both servers monitor the drives, the enclosure, and each other for maximum resiliency. This means loss of a server does not impact the reachability of the data.

These components are arranged with the supporting infrastructure to form a vault such that loss of rack allows the vault to continue.



Wasabi Erasure Coding Process & Storage System Elements

When an object is written to Wasabi, the Wasabi data striping algorithm converts the object into a series of stripes and parity fragments and stores each data stripe on 14 different drives. When you read an object from storage, Wasabi reassembles the object using the stripes. In the event one or more drives fail (up to 3 drives). This means that Wasabi can withstand the failure of up to any 3 drives within a storage pool without losing data.

Calculating Eleven 9s of Data Durability

Wasabi calculates data durability by factoring in the following metrics:

- Annualized drive failure rates (AFR)
- Mean-time-to-repair a failed drive (MTTR)
- Probability of 5 drives being unavailable at one time

For AFR, Wasabi uses industry-accepted guidelines and assumes a conservative drive AFR of 2.2%. In practice, observed AFRs are much lower. For MTTR, Wasabi estimates 3.64 days per drive, based on the calculation below:

$$MTTR = 18 \text{ TB drive capacity} * 50 \text{ MB/s drive write speed} = 3.64 \text{ days to fully write a replaced 18 TB drive}$$

As stated earlier, for an object to be lost, more than 4 or more drives in a Wasabi storage slice would have to fail. To understand the probability of this, the first step is to understand the probability a single drive failing using the calculation below:

$$Probability \text{ of a 1 drive failing} = AFR (2.2\% \text{ year}) * MTTR (3.64 \text{ days}) * (1/365 \text{ year/days}) = 2.29 \times 10^4$$

The next step would be to understand the probability of four drives failing while another drive in the storage slice was rebuilding. This would be a potential data loss scenario because five drives in a storage slice would not be available (one in rebuild mode, plus four new failures). To calculate this probability, this formula applies:

$$Probability \text{ of 4 drives failing} = Probability \text{ of 1 drive failing} (2.29 \times 10^4) \text{ to the 4th power} = 1.04 \times 10^{11}$$

The final step in calculating data durability is to factor in the probability of 4 drives failing using the following formula:

$$Data \text{ Durability} = 1 - (probability \text{ of 4 drives failing}) = 1 - 1.04 \times 10^{11} = .9999999999$$

As seen in the above calculations, Wasabi storage architecture provides 11 x 9s of durability.

One might ask how data durability of 11 x 9s translates to a real-world file loss scenario. 11 x 9s of data durability corresponds to an average expected loss of .000000001 percent of objects (1 minus .9999999999) per year. If you have 10,000,000 objects stored with Wasabi, then 11 x 9s of durability translates into 1 lost file every 10,000 years.

Wasabi Delivers High Data Durability & Protection Capabilities

Data durability is just one dimension of data protection and data durability should not be confused with system availability. Availability refers to the ability of a service provider to ensure continuous service in the event of system failures or catastrophes. Service level agreements (SLAs) typically include service availability (uptime) commitments. Wasabi's high system availability allows us to offer [strong SLAs](#) comparable to those offered by our competition.

The Wasabi service utilizes an architecture designed for high system availability and is hosted in top-tier data center facilities that are highly secure, fully redundant and certified for SOC 2 and ISO 27001 compliance. Each individual Wasabi data center is based on a highly scalable, fully distributed architecture with redundant system components, power sources and network connections to ensure high availability. The system design eliminates single points of failure; all system elements are protected using 1:1, 1+1 or N:M redundancy.

In addition to high data durability and system availability, the Wasabi service protects your data by leveraging these features:

Object Lock and Bucket Immutability: Wasabi supports optional data immutability capabilities that protects data against administrative mishaps or ransomware attacks. Object Lock is a data protection feature wherein a user can designate certain files or “objects” to be immutable, meaning they cannot be altered or deleted by anyone until the retention period has expired. This layer of protection is essential for protecting your data from cybersecurity threats and maintaining regulatory compliance. While in compliance mode, a protected file or object can’t be overwritten by any user or Wasabi engineer. When an object is locked in compliance mode, its retention date can’t be shortened.

Alternatively, buckets can be configured to allow for a default retention setting for all objects that are placed in them. For example, if the bucket level policy is set to retain an object for 30 days, the 30-day retention is calculated and applied as each object is added. Therefore, users do not have to set each object’s retention individually. With Wasabi’s immutable buckets, all objects are made immutable according to a uniform set of parameters. All objects in the bucket share the same expiration date. There can be no variation in the retention period between individual objects. Both Object Lock and immutable buckets prevent the most common causes of data loss and tampering.

Bucket Replication: To protect your data beyond the primary region, Wasabi Cloud Sync Manager provides the ability to replicate the objects within your buckets to Wasabi data centers in other regions within the same continent. This feature allows you to implement geo-redundancy by replicating the contents of your us-east-1 bucket to a bucket in the us-west-1 region.

Once you have your data replication configured, all newly added objects in the source bucket will be asynchronously copied to the designated destination bucket in the region of your choice. The Wasabi Bucket Replication feature can also sync deletions between the buckets.

Data Security: Wasabi is secure by default and all data stored in Wasabi hot cloud storage is always encrypted in flight and at rest (even if the data is already encrypted by the storage application prior to sending it to Wasabi). Wasabi follows industry-best security models and security design practices, including the AWS Identity and Access Management (IAM) model, Multi

Factor Authentication (MFA), and enterprise single-sign-on (SSO). Examples of Wasabi security features include:

- HTTPS is supported for the secure upload/download of data
- Buckets are only accessible to the bucket and object creators
- Wasabi supports user authentication to control access to data
- Access control mechanisms such as bucket policies and Access Control Lists (ACLs) can be used to selectively grant permissions to users and groups of users

Conclusion

Wasabi provides a highly durable and reliable storage infrastructure, designed from the ground up to preserve data integrity and provide high service availability. Wasabi uses advanced data slicing and parity features to achieve 11 x 9s of data durability and provide comparable levels of protection as the competition for a fraction of the price. Strong security features, optional data immutability and active integrity checking provide additional protection against bit rot, administrative mishaps and tampering. And a redundant system design ensures high uptime, backed by a comprehensive SLA. Wasabi can help you slash storage cost and complexity, while ensuring the integrity and availability of your mission-critical data.



www.gch-services.com

info@gch-services.com

+44 1628 559980



Wasabi provides simple, predictable and affordable hot cloud storage for businesses all over the world. It enables organizations to store and instantly access an unlimited amount of data at 1/5th the price of the competition with no complex tiers or unpredictable egress fees. Trusted by tens of thousands of customers worldwide, Wasabi has been recognized as one of technology's fastest-growing and most visionary companies. Created by Carbonite co-founders and cloud storage pioneers David Friend and Jeff Flowers, Wasabi is a privately held company based in Boston.

©2022 Wasabi Technologies, LLC. All rights reserved. WASABI and the WASABI Logo are trademarks of Wasabi Technologies, LLC and may not be used without permission of Wasabi Technologies, LLC. All other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s).