TECH BRIEF

Wasabi Extremely High Durability Protects Mission-Critical Data

Executive Overview

Wasabi is fundamentally transforming cloud storage with the industry's most affordable and highest-performing cloud storage service. 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.

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 failures, bit rot, degradation or other corruption. Historically, storage vendors have achieved high durability by replicating data across multiple drives using various RAID (Redundant Array of Independent Disks) schemes.

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.[1]

Durability should not be confused with availability. Availability refers to the ability of a vendor to ensure continuous service in the event of system failures or catastrophes. Service level agreements (SLAs) typically include service availability (uptime) commitments.

Wasabi provides 99.999999999% object durability (eleven 9s) and is backed by a comprehensive SLA.

Advanced Erasure Coding Protects Data Against Disk Failures and Media Corruption

Wasabi uses advanced, industry-proven erasure coding algorithms to protect data against disk failures and media errors. We transform each data object into a series of codes, which are distributed across independent disks for resiliency. In the event of disk failures or data corruption, the original data object can be reconstructed using only a subset of the codes.

Erasure coding is more efficient than and just as reliable as traditional replication-based data protection schemes. It allows us to provide eleven 9s of object durability, fully protecting customer data without the overhead of maintaining duplicate copies on multiple disks. Our approach allows us to make far more effective use of storage capacity

1 To provide some context, if you stored one million 1 GB objects on a storage platform offering eleven 9s of durability, you could expect to lose one object every 659,000 years.

compared to our competitors, while providing the same level of protection. (Amazon, for example, maintains three full copies of all customer data). We pass the cost savings of our efficiencies along to our customers in the form of ultra-low pricing.

Wasabi Distributes Data Across 20 Different Drives for Extreme Resiliency

Wasabi customer data is stored on a Wasabi storage slice composed of 20 disk drives as shown in the diagram below. Each drive is housed in a dedicated physical server to maximize data integrity. (A server failure impacts only a single drive). In addition, the drives are distributed across two racks to protect against cascading power surges or other enclosure-impacting events.

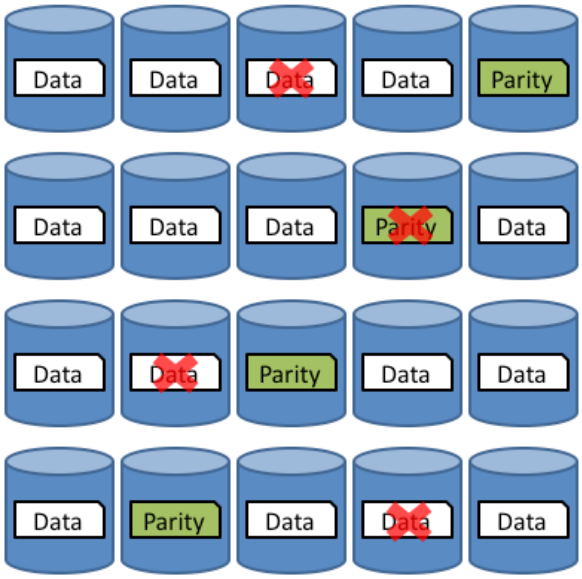

Wasabi Storage Slice with 16+4 Erasure Coding (Logical Depiction)

The coding algorithm converts customer data into a series of data and parity fragments, and stores each fragment on a different disk. When you access an object from storage it is reassembled using the fragments. An object can be fully reconstructed using any 16 of the 20 fragments. In other words, *Wasabi can withstand the failure of up to any four disks within a storage slice, without losing data.*

Erasure coding provides extremely high object durability while making optimal use of storage capacity. It also makes it easy for Wasabi to perform in-service maintenance and upgrades; disks can be swapped out without disrupting service.

Calculating Object Durability

Wasabi calculates object durability by estimating disk drive failure rates and restoration times. To be on the safe side, we use industry-accepted guidelines and assume a

conservative annual disk drive failure rate (AFR) of 5%. (In practice, our observed AFRs are much lower).

We estimate disk drive repair times by multiplying usable disk capacity (14 TB) by write-speed (50 MB/sec), which works out to a mean-time-to-repair (MTTR) of approximately 3.4 days per drive.[2] (In other words it takes about 3.4 days to repair a full disk worth of data).

In order to lose data, five or more disks in a slice must encounter a permanent failure. Assume a disk drive fails and is being rebuilt. The probability (P1) of another disk failing while the first drive is out of service equals:

$AFR \times MTTR = .05/year \times 3.4\ days \times 1/365\ year/days = 4.66 \times 10^{-4}$

The probability of four disks failing during a rebuild (P4) equals:

$(P1)^4 = (4.66x10\text{-}4)^4 = 4.7x10^{-14}$

The annual reliability of all data in a slice equals:

$1 - P4 = 1 - 4.7x10^{-14} = .9999999999999530$ or thirteen 9s of object durability.

To error on the side of caution, we advertise eleven 9s of object durability, matching the claims of our competitors.

Additional Data Protection and Availability Capabilities

Wasabi provides a variety of additional capabilities to protect data integrity, including strong security features, an optional data immutability capability and active integrity checking. In addition the service is architected to ensure continuous availability in the event of equipment failures, network problems, power outages or catastrophes.

Data Privacy and Security

Wasabi supports a comprehensive set of data privacy and security capabilities to prevent unauthorized data disclosure. Strong user authentication features tightly control access to stored data. Access control lists and administratively defined policies selectively grant read/write and administrative permissions to users, groups of users, and roles.

Wasabi encrypts data at rest and data in transit to prevent leakage and ensure privacy. All data stored on Wasabi is encrypted by default to protect data at rest. And communications with Wasabi can be transmitted using HTTPS to protect data in transit.

2 Today's disk drives support write-speeds of up to 200 MB/sec. We use 50 MB/sec as a conservative estimate, which also means 3.4 days to repair a full disk worth of data is conservative. Our engineers like to under-promise and over-deliver in the never-ending quest to design for failure and resiliency.

Data Immutability

Wasabi supports an optional data immutability capability that protects data against administrative mishaps or abuse. An immutable object cannot be deleted or modified by anyone—including Wasabi. Wasabi data immutability protects against the most common causes of data loss and tampering including accidental file deletions, viruses and ransomware.

Active Integrity Checking

Wasabi uses three distinct checksum mechanisms to verify the integrity of all stored objects every 90 days. We detect and repair corrupt data automatically to avoid bit rot or other corruption, mitigating risk and uncertainty immediately.

High Availability System Architecture

The Wasabi service 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.

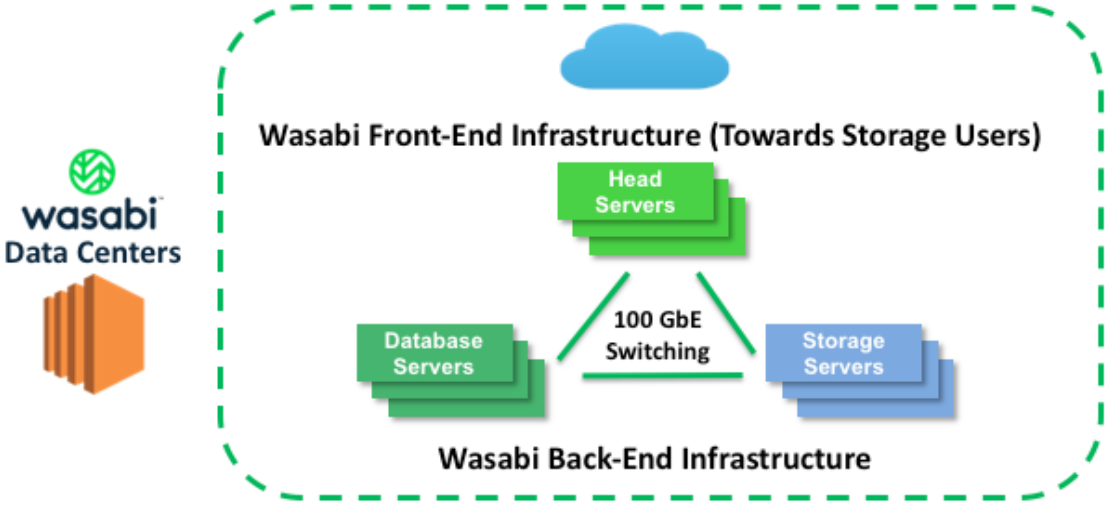

Fully Redundant System Architecture Provides High Availability

Wasabi offers an optional geo-replication capability (customer data is replicated across geographically distributed data centers) to ensure continuous service in the event of a catastrophe or individual data center outage.

Conclusion

Wasabi provides an extremely durable and reliable storage infrastructure, designed from the ground up to preserve data integrity and provide high service availability. We use sophisticated erasure coding algorithms to achieve eleven 9s of object durability, providing the same 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.

Next Steps

- **CONTACT WASABI TODAY.** Learn more about our price, performance and protection benefits.
- **TRY WASABI FOR FREE.** Get up to 1 TB for 30 days.

About Wasabi

Wasabi is the hot cloud storage company delivering low-cost, fast, and reliable cloud storage. Wasabi is 80% cheaper and 6x faster than Amazon S3, with 100% data immutability protection and no data egress fees.

Created by Carbonite co-founders and cloud storage pioneers David Friend and Jeff Flowers, Wasabi is on a mission to commoditize the storage industry. Wasabi is a privately held company based in Boston, MA..

wasabi™

Note: All pricing information presented in this paper was retrieved from cloud service provider websites and was accurate as of July 2017.

Tel 1-844-WASABI-1
Email info@wasabi.com

wasabi™
hot cloud storage

www.wasabi.com