

The Hot Cloud Storage Guide to Backup and Recovery



wasabi
hot cloud storage

Table of Contents

Why the “hot cloud storage” guide to backup and recovery?.....	1
Introduction	2
The current state of backup and recovery	3
Growing data deluge	3
Growing compliance and regulation.....	4
Growing costs and scalability issues.....	6
Growing security threats.....	6
Growing complexity	6
Traditional approaches to storage	7
Traditional storage tiers	7
Storage tiers solved the cost vs performance challenge. For a while.....	8
A comparison of on-premises media types	9
Short-term storage: Disks and drives	9
Long-term storage: Magnetic tape.....	11
The bottom line: Yesterday’s storage media can’t handle today’s needs	12
Storage innovation means faster, cheaper, and fewer tiers	13
The rise of cloud storage for backup and recovery	15
Primer: Cloud object storage	15
The 3-2-1 rule: A great place to start	16
The security question: Why your data is more secure in the cloud	17
The approach to cloud storage evolves along with technology.....	17
The benefits of Cloud 1.0 backup.....	18
The challenges of Cloud 1.0 backup	20
The complexity of Cloud 1.0: A look at AWS Simple Storage Service	22
Cloud 2.0 eliminates storage tiering and cost complexity	24
Wasabi hot cloud storage: Fulfilling the promises of Cloud 2.0.....	26
A new approach to cloud storage	26
Why the incumbent Cloud 1.0 storage providers can't catch up.....	27
What puts the sizzle in Wasabi hot cloud storage?	27
Step up to Wasabi hot cloud storage and get more out of your data.....	32
A comparison between on-premises, Cloud 1.0, and Cloud 2.0 storage solutions.....	32
Appendix	34
About the Authors	35

Why the “hot cloud storage” guide to backup and recovery?

Actually, you may be wondering why a “storage” guide at all? After all, isn’t storage a commodity? “Can’t I just pick the lowest-cost option and be done with it?” Sadly, it’s not that easy. There are many different storage options and countless vendors and they don’t make figuring out your true total cost of ownership very easy. Sure, the cost per gigabyte continues to go down, but complexity (*and obfuscation*) is on the rise. Massive data growth is outstripping on-premises storage ability to scale. This plus faster hardware obsolescence is placing IT teams in a constant state of worry about rising costs and compatibility issues with existing equipment.

Cloud storage was supposed to be the answer, but adopters of first-generation cloud solutions are encountering similar challenges with rising costs and complexity. Multiple tiers of object storage services (AWS has six different storage offerings) with their different cost structures, restrictions, and transaction fees force IT teams to establish complicated data lifecycle plans and implement automation software to move data from tier to tier in an effort to control costs. And there’s still no reliable way to predict the total cost of storage.

We agree with you. Storage should be a commodity—a standards-based utility, like electricity. There is only one level of service with one low usage-based fee. Use it how you want and pay only for what you use. That is the thinking behind hot cloud storage, the next generation of cloud storage technology that is driving the Cloud 2.0 revolution. It can help solve for many of the backup and recovery challenges described in this guide.

Introduction

In today's digital world, losing access to your data is the fastest route to failure. Data loss is not only one of the most common causes of business disruption, it's one of the costliest. A recent study by Verizon reported that small instances of data loss (approximately 100 lost or compromised records) can cost businesses an average of \$18,120 to \$35,730, depending upon the size of the company and the value of the data. The same study found that large-scale data loss (100+ million records) costs between \$5 million and \$15.6 million. Some businesses never bounce back from data loss of this magnitude. It is estimated that 1 in 5 small businesses have been forced to close their doors due to data loss caused from a ransomware attack.

It's no wonder business continuity and data backup and recovery are high on the list of IT priorities. If you've downloaded this guide, you're no doubt interested in replacing or upgrading your existing backup and recovery solution. Like many organizations, you may be considering migrating to a third-party cloud backup provider or are planning a DIY approach using one of the major cloud object storage vendors. Perhaps you've invested heavily in on-premises storage and are looking at cloud as a second- or third-copy option for ensuring continuity and resiliency of your business.

Wherever you are in your journey, this guide will help you better understand the latest trends and options for backup and recovery storage, and the pros and pitfalls of each of the storage types available. We'll explore the evolving and growing use cases for cloud storage, and help you determine which cloud is right for you.

The current state of backup and recovery

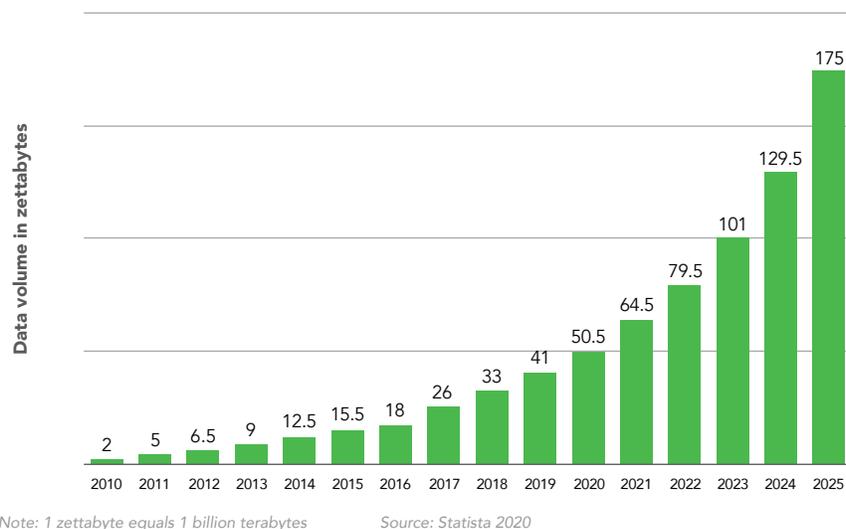
Cloud computing. A.I. Blockchain. Cybersecurity. Data analytics. We don't need to remind you of the evolving digital landscape that is transforming businesses and IT departments everywhere. You're living it every day. However, with all these advancements in IT, many companies are waking up to the fact that their current traditional approaches to backup and recovery are not keeping up. Here are just a few of the reasons why:

Growing data deluge

Businesses are creating, consuming, and storing more data than ever before—not just human-generated, but machine-generated data from sensors, devices, and A.I. algorithms distributed across the enterprise. IT departments all over the world are faced with managing not just hundreds of terabytes (TB), but petabytes (PB) of data. And the data tsunami is accelerating (see Figure 1).

Figure 1 Projected Data Growth Worldwide

The total amount of data created worldwide is forecasted to increase dramatically, reaching 175 zettabytes by 2025.



This deluge of data is creating multiple challenges, not the least of which is the high cost of storage. As companies look to protect all this data, the backups themselves become a part of the problem. As the size of backups grow, so too do the portions of those backups that need to be considered for long-term retention, duplicating data and eating up more storage space. To deal with the problem, businesses are forced to buy more storage capacity or purge data.

Purchasing additional storage is often fiscally prohibitive while “fear of deletion” can paralyze IT decision makers. Clearly some data can be deleted but sifting through terabytes to find the data that may qualify for deletion is a daunting task. Which data could be monetized in the future? Which must be kept for compliance and legal requirements? Fear of deletion is real. For many of us, the voice in the back of our head saying, “Is deleting this data a bad idea?” keeps us up at night.

Growing compliance and regulation

The growing number of government and industry regulatory compliance mandates governing how organizations retain, certify, or destroy data has an enormous impact on your backup and storage decisions. While each regulation has its own distinct set of guidelines, storage compliance typically affects the following three areas: retention, integrity, and security.

Retention dictates how long certain data must be kept in storage. In healthcare, for example, the Health Insurance Portability and Accountability Act (HIPAA) stipulates that documents must be retained for a minimum of six years. But how is a storage administrator to distinguish a HIPAA-related document from any other document type? The people creating the data are not typically required to identify or classify it as such. So, guess what happens when those responsible for managing the data have no way to determine what is governed by compliance standards and what is not? They keep it all.

And it’s not just a matter of throwing everything into inexpensive, long-term storage, either. The stored data must also be easily and quickly retrievable in the event of compliance audits or legal discovery.

Integrity means that the data cannot be changed or lost due to corruption, media failure, or tampering, and is readable regardless of how long it has been stored. In the legal industry, evidence lock or a concept known as “chain of custody” is required by Criminal Justice Information Services (CJIS) and other regulatory bodies. This means that the backup storage they use must provide evidence of a lock functionality, such as immutability, that can protect evidence such as surveillance video or audio from being altered or deleted. If you require this level of assurance and data integrity, look for a storage provider that can provide data immutability. With immutability, data cannot be altered or deleted by anyone—even automatically by the system after the storage retention time is reached. The data can only be unlocked once the immutability configuration has reached the end date.

Security protects your sensitive data from unauthorized access. When assessing third-party backup providers and storage vendors, look for a service built and managed according to modern security best practices and standards for both data and physical data center security:

DATA SECURITY	PHYSICAL SECURITY
<ul style="list-style-type: none"> • Secure Sockets Layer (SSL) protocol for securing internet communications • AES-256 encryption of data at rest and in transit • User authentication and control, including bucket policies and Access Control Lists (ACLs) • Options to connect directly to service provider and bypass public internet 	<ul style="list-style-type: none"> • Tier 4 data centers that are fully redundant and certified for SOC-2 and ISO 27001 compliance • Defense-in-depth network security architecture, protecting both external and internal borders • 24/7 onsite security personnel and surveillance video • Biometric readers and two-factor or greater authentication for building access

When it comes to backing up your data in a third-party cloud, make sure your provider complies with the security requirements set forth by all the major industry and government regulations. Organizations often must comply with multiple regulations simultaneously, and you may not know which ones in advance. Keeping with our legal example, video from a body-worn camera may become evidence for a criminal case (CJIS compliance), recorded in a medical facility (HIPAA compliance), and include information about an individual’s employment or income (IRS 1075 compliance).

Growing costs and scalability issues

The cost of storage may be falling, but the problem is now one of scale not price. In order to accommodate this ever-increasing mountain of data, businesses are forced to increase their storage capacity—and expanding existing backup systems and tape libraries are expensive (see [Appendix on page 34 for cost analysis](#)). Not too long ago, companies tended to purchase storage in five-year periods. Due to the tremendous growth in data in recent years, the typical storage contract is now three years.

Growing security threats

Ransomware and other cybersecurity threats are on the rise. According to a recent survey by US Signal, 83% of organizations have experienced a DDoS attack within the last two years and 38% reported having been impacted by a ransomware attack in the last 12 months.

Ransomware attacks are growing in sophistication, too, targeting backup workloads to encrypt first in order to ensure that their victims have no way to restore affected systems. This worrisome trend makes secure third-party (offsite) cloud-based backup and recovery a critical addition to an organization's cyberdefense and business continuity plans.

Growing complexity

Today's storage environments are a tangled, grown-over-time mix of legacy point products for backup, target storage, and long-term data retention. For backup and recovery, many enterprises rely on multiple data protection solutions: backup servers, appliances, a variety of disk and tape media, as well as cloud storage from multiple vendors. This creates data silos that are increasingly more difficult to manage as each solution is built on proprietary hardware and software with different management tools, upgrade cycles, and maintenance contracts.

As organizational needs change and workloads become increasingly more distributed, businesses should look to consolidate and simplify their data storage.

Traditional approaches to storage

The relationship between price and performance has been the driving force behind storage architecture design since the dawn of the information age. The faster it is to read/write data, the higher the cost. In that respect, data storage is no different than a high-performance automobile. If you want speed, you're going to pay for it. And because high performance costs more, speed is going to come at the expense of capacity.

In order to reduce storage costs as much as possible, tiers or layers of storage were created based on varying levels of performance (and their associated costs). The traditional scale-out approach to storage is categorized into three tiers:



Tier 1 is the performance tier. Used to store all your active, working data, it is the fastest and most expensive tier of storage. Tier 1 includes media types such as solid state drives (SSDs) and high-performance, high-capacity hard disk drives (HDDs). However, faster non-volatile memory express (NVMe)-based flash storage is rapidly taking over this tier. (It's so fast, some storage professionals refer to flash as tier 0.)



Tier 2 is the capacity tier—also referred to as short-term storage. It offers greater balance between price, performance, and availability, making it ideal for large datasets that are less time sensitive, such as back-office applications, ERP, and email, or that are accessed less frequently, such as backups and disaster recovery. Storage options typically include optical media and removable or fixed lower-performance HDDs or arrays. In recent years, cloud storage has become a popular and growing option.



Tier 3 is your archive tier—data that is rarely, if ever, accessed. It is the lowest-cost option, which comes at the expense of both performance and availability.

This long-term storage is for data that must be retained for business and/or regulatory compliance reasons. Archives are the traditional realm of magnetic tape (LTO being the most popular), although cloud is an attractive option here as well.

Tier 1 Performance Tier	Tier 2 Capacity Tier	Tier 3 Archive Tier
<ul style="list-style-type: none"> • Active, working data • Fast but expensive 	<ul style="list-style-type: none"> • Infrequently accessed • Balance between performance, availability, and cost 	<ul style="list-style-type: none"> • Rarely accessed • Low cost, low performance and availability

Storage tiers solved the cost vs performance challenge. For a while.

In the tiered approach, enterprises store their most critical, time-sensitive data and applications—the information that is needed to do business every day—on the fastest (and most expensive) storage media. Data that doesn't need to be accessed as often can go on media that isn't as fast but costs less. Data that is rarely or never used goes into archives, or cold storage, on the cheapest and slowest media available, with little concern about performance.

A comparison of on-premises media types

In this section, we'll take a look at the various on-premises storage media available today and consider the advantages and disadvantages to each.

Short-term storage: Disks and drives

Companies have stored their data backups on disk since the beginning, but technological advancements have resulted in huge improvements in capacity, speed, performance, and accessibility of disk storage over time.

Traditional hard disk drives (HDDs) and arrays

HDDs and arrays, including HDD storage area networks (SANs) and network-attached storage (NAS), have been the go-to medium for enterprise backup for more than 50 years.

PROS	CONS
<ul style="list-style-type: none">• Proven technology• Still less expensive than solid state drives (SSDs)• Larger storage capacity than SSDs, with high-capacity drives that offer 10 TB or more	<ul style="list-style-type: none">• Many moving parts increase potential of mechanical failure• Easily damaged if dropped or handled roughly• Potential for corruption, physical damage• Slower data search and retrieval• Larger, heavier than SSDs• Less energy efficient than comparable SSDs

Optical media

Optical media includes CDs, DVDs, and Blu-Ray and BD-R discs. Most rewritable discs are unsuitable for storage and archiving, as their data layers are dye-based and inherently unstable. However, write-once BD-R HTL (High to Low) and some DVD+R discs can last more than 150 years if stored correctly.

PROS	CONS
<ul style="list-style-type: none">• Long media life• Very high data stability• Relatively low cost per bit of storage• Portability	<ul style="list-style-type: none">• Very slow write speeds (approximate max of 21-29 Mbps)• Limited capacity – you’d need 218 DVDs to back up 1 TB• High need for IT management and operation

While HDDs and HDD arrays, optical media—and other traditional options such as SD and micro-SD cards—are still available and used today, their popularity is waning as far more advanced options take their place.

Solid state drives (SSDs) and flash arrays

SSDs have become the “now-standard storage medium for business PCs,” according to a [recent report](#). They are storage devices that don’t have moving parts like traditional HDDs. When SSDs first arrived on the scene, they used RAM and DRAM to store data. Today, SSDs typically use flash technology, which is why even industry experts tend to use the terms “SSD” and “flash” interchangeably. SSDs are faster than HDDs because all of their memory chips are accessible at the same time.

All-flash arrays (AFAs), also called solid-state storage disk systems, are collections of devices that use only flash media to store data. The flash memory provides persistent storage, meaning that even if the device loses power, no data is lost. AFAs are designed largely for primary data storage where speed and time are critical. They’re also ideal for use cases that are read-intensive and have small working datasets, such as ERP or CRM. Simply put, AFAs are more suited to the performance tier than the capacity tier.

PROS	CONS
<ul style="list-style-type: none"> • Excellent performance, latency, and speed compared to HDD • Lower energy costs • Smaller, lighter footprint • Noiseless • Less potential for mechanical failure • Lower IT management requirement 	<ul style="list-style-type: none"> • Higher costs than HDD and other options • Lower capacity than HDD, which can be up to 2.5 times larger • Relatively short lifespan, as cells decay every time data is written or erased • Steep performance decline as the system nears full capacity • Data recovery challenges, such as encryption in onboard controllers—if the controller fails, it will be impossible to retrieve the data • Not suited to capacity or archive tiers

Long-term storage: Magnetic tape

Companies have been storing computing data on magnetic tape since the 1950s. While the market for magnetic tape is not increasing much these days, it's holding surprisingly steady. Large organizations across the world still rely heavily on tape to store massive volumes of archival data. In fact, in 2019, CERN, the European Organization for Nuclear Research, recently predicted its [to-tape archival storage needs](#) would reach 1.3 exabytes—or roughly 1.3 million terabytes—by 2024.

While there have been many types of magnetic tape for backup and storage, the clear winner today is linear-type open, or LTO tape. Developed in the late 1990s, LTO tape quickly edged out competitors thanks to its open source format and significant capacity. It's been enhanced over the decades by Hewlett Packard Enterprise (HPE), IBM, and Quantum.

The 9th generation of LTO tape will offer up to 25 TB of native (uncompressed) capacity and has a confirmed capacity of up to 60 TB of compressed data—double the capacity of LTO-8. It enables transfer rates of up to 708 Mbps and includes read-and-write compatibility with LTO-8 tapes and read compatibility with both LTO-7 and LTO-8. The LTO Organization has released a roadmap for future generations up to LTO-12, which will have a compressed capacity of 480 TB.

Other manufacturers haven't given up on tape yet, despite the popularity of LTO tape. In 2017, Sony and IBM announced the development of a joint product that can accommodate about 330 TB of data within a palm-sized tape cartridge.

PROS	CONS
<ul style="list-style-type: none"> • Super low per-month storage costs • Long-accepted and familiar industry standard • Large storage capacity • Long shelf life • High portability • Strong security due to being stored offline, often called the "air gap" • Ideal for long-term retention, archiving, and cold storage • Many companies have long-time investment in existing tape systems 	<ul style="list-style-type: none"> • Very slow restore speeds • Cumbersome accessibility to data • Expensive scaling—high up-front costs for system equipment • Potential for physical damage, corruption, loss • Obsolescence of older-generation systems, constant need to migrate to newer tape • High maintenance needs—tape rotation, cleaning, replacement, etc. • Eventual magnetic and physical degradation over time

The bottom line: Yesterday's storage media can't handle today's needs

When it comes to backup and recovery, on-premises HDDs, optical media, and magnetic tape don't meet the needs of the modern enterprise. Disks are too expensive to scale and expand capacity and tape's slow restore speeds and general inaccessibility of data increasingly conflict with today's digital business need for agility and immediate access to data.

Beyond these issues of cost, scale, and accessibility, these traditional storage options suffer from:

- **High CAPEX** – significant upfront equipment costs lead to poor ROIs with long payback periods.
- **High OPEX** – recurring power, cooling, and rack space expenses; expensive monthly hardware and software maintenance and support fees; and excessive system administration cost and complexity all lead to high ongoing operations expenses.

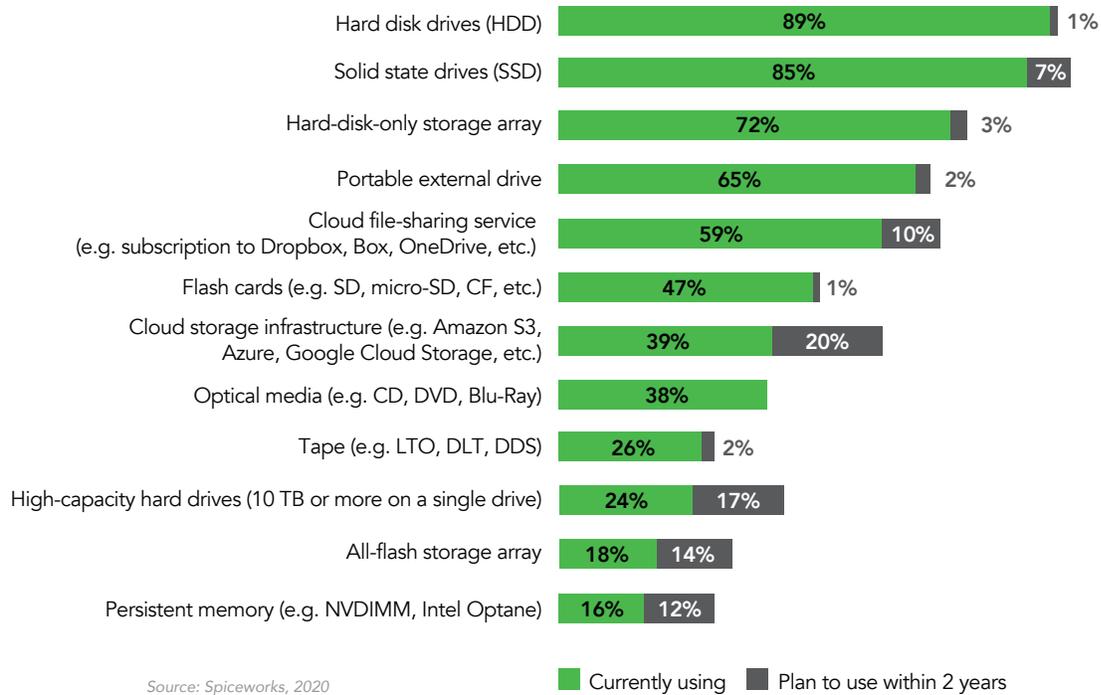
- **Obsolescence concerns** – storage vendors regularly retire products and discontinue support plans, often subjecting customers to costly and disruptive upgrades.
- **High risk** – ensuring business continuity (replicating live data to a secondary data center) is an expensive proposition beyond the reach of most organizations.
- **Complex operations** – legacy storage solutions are difficult to configure and administer and require special training and expertise.

Storage innovation means faster, cheaper, and fewer tiers

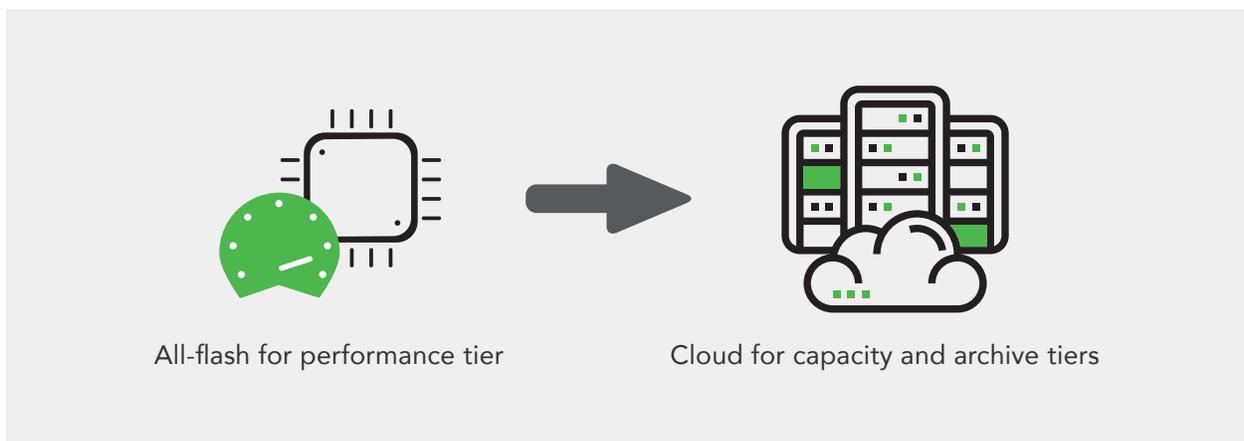
Storage systems are not immune to Moore's Law. Performance and capacity continue to increase even as prices come down. Innovation in both performance storage and capacity storage is simplifying options for enterprise storage buyers.

At the performance tier, NVMe-based flash storage offers greater performance over spinning disks and even SSDs. As prices continue to come down, all-flash arrays are becoming a much more viable and enticing option for on-premises performance-tier backup.

At the capacity tier, the cloud, with its massive economies of scale, among other benefits, is growing rapidly for both short-term and long-term storage. According to a [recent Spiceworks survey](#), 39% of businesses use cloud-based storage and an additional 20% plan to adopt cloud storage by 2022 (see Figure 2 on next page).

Figure 2 Current and Planned Business Adoption of Storage Technologies


As the chart above demonstrates, growth in cloud storage infrastructure (at 20%) is outstripping all other media types, while the planned adoption of all-flash arrays nearly doubled. Some storage experts predict that the industry will soon move to a two-tier model: all flash for the performance tier and cloud for the capacity and archive tiers.



The rise of cloud storage for backup and recovery

As the numbers indicate, more businesses are turning to cloud storage. Just what is cloud storage and why is it growing in popularity?

Primer: Cloud object storage

Object storage is the underlying technology behind all cloud-based storage systems. Unlike the file- and block-level storage technology used by current and mostly proprietary network-attached storage (NAS) and storage area network (SAN) systems, object storage identifies a set of data as an individual entity (or object), instead of a location on a disk. This is what gives cloud its near unlimited capacity and economies of scale, enabling it to dynamically scale across massive, multi-node storage systems.

These systems maintain an independent database of pertinent information about the data itself. Because object storage doesn't place any theoretical limitations on the number or size of objects in storage, it's possible to store a near limitless amount of metadata (information about the data) to enable extremely powerful and flexible data management capabilities, such as index and search, and to establish policy and automate functions such as billing, tiering, and customer access management.

These factors, plus its inherent "offsite" nature, which makes it easy for customers to achieve geographic separation from their primary site (see The 3-2-1 rule: A great place to start on next page), makes cloud the ideal medium for business continuity, backup, and recovery.

THE 3-2-1 RULE: A GREAT PLACE TO START

As the gold standard for optimal protection in data backup, the 3-2-1 rule is supported by practically every backup solution provider that exists. The rule states:

Keep at least 3 copies of your data	Use at least 2 separate storage types	Store at least 1 copy offsite
Multiple copies mean you have a backup if the primary (and even secondary) copy is destroyed.	Protect against different types of hazards, such as an electrical surge that might fry disk drives but leave tape unaffected.	Keep data immune from localized disasters, such as fires, floods, or earthquake damage.

Cloud is changing things

The original 3-2-1 rule was established when backup technology and options were limited. With the cloud as a popular option these days, the 3-2-1 rule is evolving. For instance, if you back up data in the cloud, it's recommended that you have at least two cloud copies that are geographically separated from each other.

Some companies side heavily with cloud: they keep their production data on disk, then put their two other copies in separated clouds. Others prefer a mix of local and cloud-based storage: they have production data on disk, with a copy on another physical medium such as tape, and two copies in the cloud.

There is no backup system that is perfect for every organization. Use the 3-2-1 rule as the minimum starting point and find out what works best for your company.

The security question: Why your data is more secure in the cloud

When cloud first came on the scene, some IT professionals were concerned with security. After all, wouldn't your data be safer and better protected onsite where you can control who has access? Not really. The answer may seem counterintuitive, but unlike you and your busy IT team who have many different roles and responsibilities, cloud storage providers are 100% focused on protecting your data. It is their sole business. When you select a cloud backup provider or a cloud storage vendor over traditional on-premises backup, storage experts at their data centers will:

- Monitor your data 24/7 to protect against data breaches and network failures
- Protect the physical data center against intrusion or disaster
- Automate backups of your data in multiple locations, so that even if one site is compromised, your data is still safe
- Regularly update software and hardware and run integrity checks on data

Backups are also generally more secure and protected against ransomware attacks because they are performed outside of the business network, where backups would be the primary target. Backup data transmitted from the customer site to a cloud backup provider is typically encrypted and remains so while on the provider's storage service system.

The approach to cloud storage evolves along with technology

In the early days of cloud, the concept of "elasticity"—paying only for the storage you use and being able to quickly scale up or down—was the main benefit over traditional storage methods. Those emerging solutions, or what the industry calls Cloud 1.0 solutions, didn't change the way enterprises used data, they simply replicated in the cloud almost the same technology and cost structures that were used in on-premises storage solutions. Including the classic tiered architecture and the strategy of storing as little data as an enterprise could get away with, to save on costs.

Over time, however, data became increasingly critical as analytics solutions enabled enterprises to gain immense value from it. The scarcity mindset of Cloud 1.0 storage—store as little as possible—was no longer a viable business model.

Fortunately, storage technology was advancing at the same time and promised faster read/write times, enhanced reliability, and fewer hardware problems—all at lower cost. Cloud storage providers began to move away from solutions that were simply a “lift and shift” of what was done on premises. They developed Cloud 2.0 storage solutions that took advantage of those technological improvements and that required a new approach to data. Thanks to the increased speeds and lower costs, scarcity was no longer the approach—these new solutions focus on and support data abundance.

In the following sections, we’ll outline how Cloud 1.0 solved some of the challenges of traditional disk and tape solutions, prompting enterprises to move their data storage and backup to the cloud. Cloud 1.0 solutions also present their own unique challenges, however, so we’ll then show you how Cloud 2.0 solutions are now addressing those.

The benefits of Cloud 1.0 backup

Geographic separation and anywhere, anytime accessibility

As we mentioned, one of the best benefits of using the cloud for backup and recovery is the ease in which you can establish geographic separation from your primary data. If your corporate data center goes down due to a local weather disaster, IT could recover the cloud backup to any of the organization’s branch offices without a problem. With the cloud, you can access your data from anywhere, on any device, through the Internet.

Simpler scalability

The cloud is immensely scalable. That means businesses can use as much or as little as they need at any time. Resources are sold on demand and managed by the cloud service provider.

Simpler management

Because cloud service providers handle a lot of the management tasks, cloud backups allow you to protect your data without having to increase IT staff to manage it all. Offloading routine tasks to the service provider frees up IT to spend more time supporting business growth.

Increased security

As previously discussed, many experts consider the cloud to be more secure than most on-premises data centers.

Less potential for manual error and damage

Many on-premises backups are incomplete, corrupted, or damaged due to user error. Disk drives and tapes, for example, can get lost, stolen, or damaged. With redundant backup copies in the cloud, you don't have to worry about accidental backup failures or deletions.

Better data consolidation

It's easier to consolidate all of your backup data in the cloud, because the cloud can accommodate everything, from your primary data center storage systems, to servers and devices in branch offices, to employee devices such as laptops and smartphones. Cloud storage can also help you consolidate data from other clouds, such as Office 365 or cloud-based application workloads.

Lower costs (but not always)

In general, cloud backups are more cost-effective because, unlike building and managing your own on-premises backup infrastructure, you only pay for what you use. Backing up data in the cloud means you can eliminate the upfront costs of the equipment and software needed to maintain tape storage, backup servers, or external drives, as well as the ongoing costs in personnel, power, data center space, license renewals, and media updates.

Of course, more needs to be said about the cost and pricing of cloud backup. While initial costs are typically lower than other storage solutions, there are challenges with current generation cloud storage (Cloud 1.0) that can cause your total costs to escalate quickly.

The challenges of Cloud 1.0 backup

Costs can rise quickly

While initial cost-per-gigabyte (GB) of data stored in Cloud 1.0 solutions seems low, overall costs can add up fast as your data grows. Exacerbating this problem are the number of additional costs and transaction fees placed on the data owner (see Appendix for cost analysis).

Transaction fees make total costs unpredictable and difficult to calculate

Beyond the capacity-based charge for storing your data, depending upon the tier of service you choose, Cloud 1.0 providers make you pay to access or move your data (in the form of egress fees). Again, depending upon your service tier, you could be charged additional transaction fees for every PUT, POST, GET, LIST, and any number of other API requests. These hidden fees not only add substantially to your overall costs, they make it impossible for you to plan and predict your total cost of storage. This unpredictability has forced many enterprises to pay outside consultants just to help them with data backup cost analysis and capacity planning.

You are hit with extra fees at the worst possible time—when you need your data

When disaster strikes and you lose your primary data, you'll be glad to have a second backup in the cloud. But getting hit with costly egress fees to access all of your data after getting hit by a ransomware attack or natural disaster is just compounding your pain.

Backup and recovery speeds can be slow

Because cloud backups require encryption and then transmission over the Internet, the backup process can be slower than saving data to a local disk drive. Speeds can also vary depending on Internet bandwidth and latency. While this may not be a concern when creating a backup, it may be an issue when trying to recover data quickly when you need it.

Vendor lock-in

To get the most out of a cloud service provider's offerings, it's almost a necessity to go "all in" with that vendor's services. While many companies choose to use multiple cloud service providers, they can run into trouble because one provider requires a specific configuration or customization that is incompatible with another provider's platform. Every business wants the flexibility and agility to pivot quickly if they need to. Being locked in with one cloud service provider for your data backup can limit your ability to do that. When exploring cloud backup solutions, look for a cloud storage provider that offers free egress (and compatibility with your existing cloud provider, if you have one). Also look for a cloud backup provider that provides a free "community version" of its software, so you can migrate your backup data to another backup provider without having to pay for it.

Complicated tiers require substantial data lifecycle planning and automation

Cloud 1.0 storage companies have structured their offerings in tiers of service designed to give you price/performance options for active, short-term, and long-term storage. However, this forces IT to come up with policies and procedures for determining how best to classify data and invest in software or services for automating data tiering.

The complexity of Cloud 1.0: A look at AWS Simple Storage Service

While Cloud 1.0 storage providers have structured their offerings to be similar to the traditional concept of performance, capacity, and archive tiers introduced earlier in this guide, Amazon Web Services (AWS) actually offers six different tiers of service (see Figure 3).

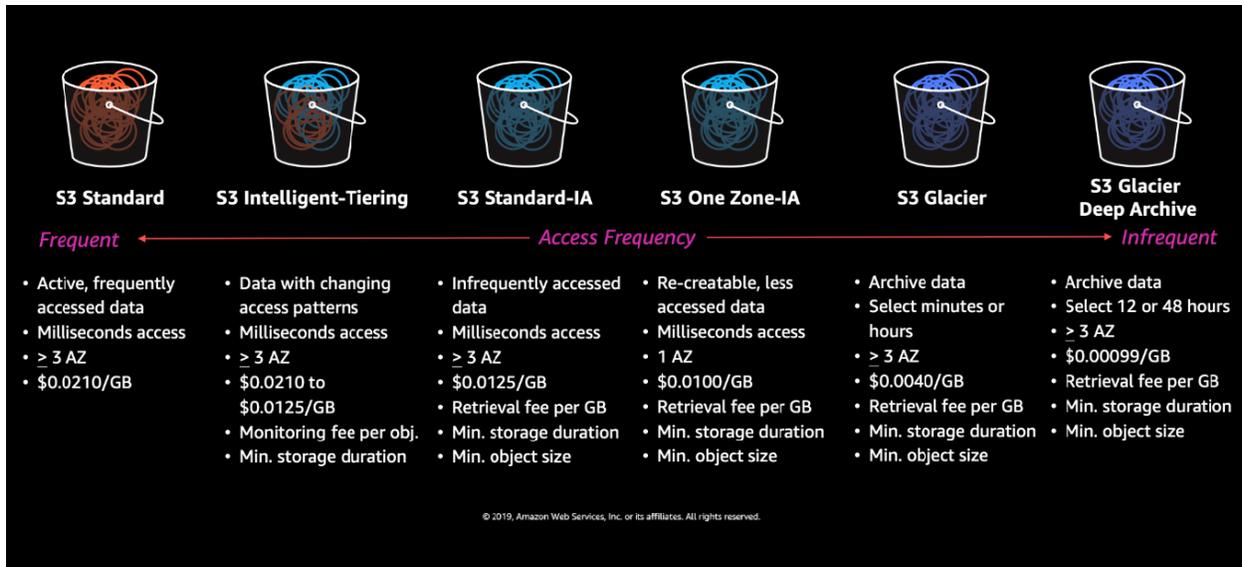
AWS Simple Storage Service (S3) offerings include:

- **S3 Standard** – designed for frequently accessed active data. It has the highest capacity-based charge and the lowest transaction costs with no minimum retention period.
- **S3 Standard-Infrequent Access (IA)** – designed more for backup and recovery. It has a lower capacity-based charge but higher transaction costs, including a data retrieval cost. It also has a minimum storage duration period of 30 days (you will be charged the full 30 days for any objects stored for less than 30 days.)
- **S3 Infrequent Access-One Zone** – a lower-priced option to standard IA. The tradeoff is lower durability and availability as AWS reduces the number of redundant copies of data, and it's only stored in one AWS data center zone (geographic location), making it susceptible to maintenance downtime or localized events such as natural disasters.

Note: There is no difference in read/write speeds until you get to their Glacier offerings.

- **S3 Glacier** – AWS’ original version of cold storage used for archive data. It is significantly slower than the first three offerings and, therefore, comes with an appropriately lower per-GB capacity charge. Because this is “offline” data, objects in Glacier cannot be read and must first be restored into S3. You determine how quickly you need access (between 1–5 minutes and 5–12 hours) and are billed accordingly. It has a 90-day minimum retention period.
- **SE Glacier Deep Archive** – an even lower-cost archive option. The minimum retention is increased to 180 days and retrieval times are increased to 12 or 48 hours. Like S3 Infrequent Access-One Zone, fewer redundant copies of data are kept.
- **S3 Intelligent-Tiering** – a newer service that enables AWS customers to move objects between tiers automatically.

Figure 3 Amazon S3 Storage Classes



Source: Amazon

Cloud 2.0 eliminates storage tiering and cost complexity

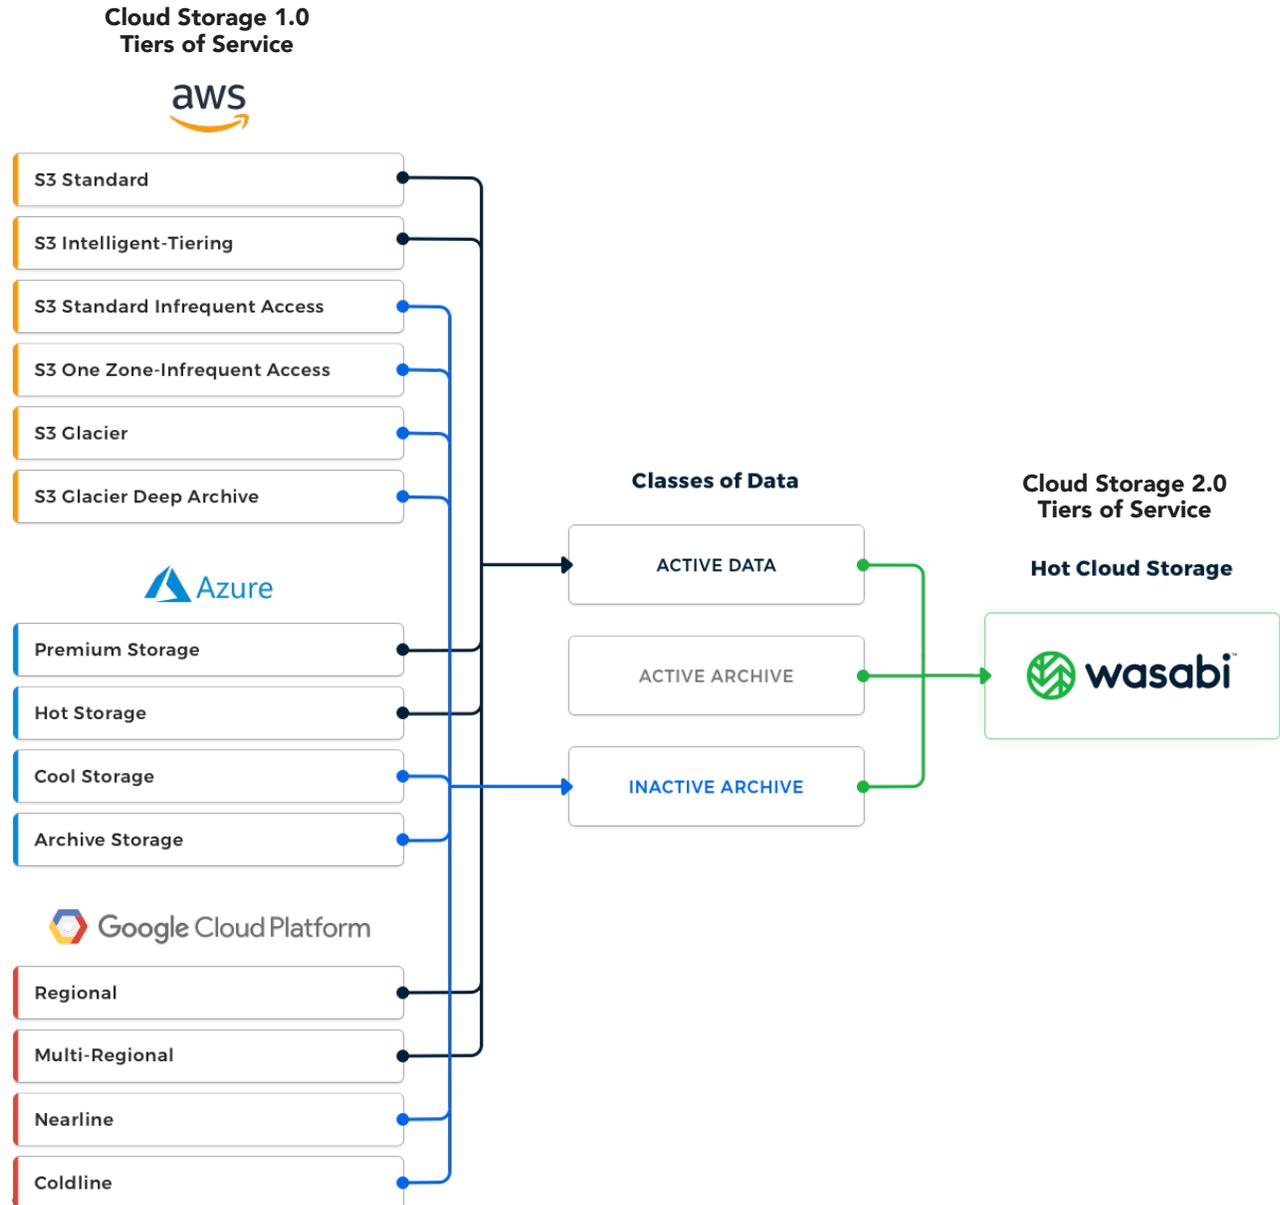
At first glance, the tiering of storage types—with its varied pricing—may seem like a good way to keep costs as low as possible. But as data increases, this adds a high degree of complexity and cost. It can be so complicated to figure out, in fact, that there are consultants working today whose whole job is to help businesses work out storage tiering and pricing in the cloud.

With Cloud 1.0 storage, you have to figure out which collections of data are more or less critical and how often you will need to access them. In other words, you are forced to decide, up front, not only which data are worth storing, but which data are worth paying more for the privilege of accessing faster. This is the Cloud 1.0 scarcity mindset in action.

Fortunately, it doesn't have to be this way. A new generation of cloud storage has arrived on the scene that doesn't mimic the tiered structure of on-premises storage like first-generation Cloud 1.0 did. Cloud 2.0 simplifies storage by removing tiers altogether (see Figure 4).

Figure 4 Cloud 2.0 Greatly Simplifies Storage

Cloud 2.0 is universal, one-size-fits-all cloud object storage that eliminates confusing storage tiers by satisfying all storage price-performance requirements.



Wasabi hot cloud storage: Fulfilling the promises of Cloud 2.0

As enterprises slowly came to realize the limitations of Cloud 1.0 solutions—costly and confusing service tiers, poor performance, inaccessibility, vendor lock-in, and inability to predict total cost of ownership (TCO)—a Boston-based startup called Wasabi was quietly inventing the future of cloud-based storage based on the Cloud 2.0 mindset of abundance described earlier. Their goal was to increase performance and lower costs while simultaneously simplifying billing and making TCO infinitely more predictable.

With only one tier of service, Wasabi hot cloud storage treats all data as potentially valuable and equally accessible at one remarkably low price.

A new approach to cloud storage

In 2015, Carbonite co-founders and cloud storage pioneers David Friend and Jeff Flowers set out on a mission to make cloud storage a simple commodity and utility, just like electricity. Leveraging the latest advances in storage technology and designing an innovative new file system that made use of all the available space on every disk drive platter, they were able to significantly increase performance, maximize capacity, and lower operating costs. The result was Wasabi hot cloud storage, an enterprise-class, tier-free, instantly available cloud storage service that is 1/5th the price of AWS S3 Standard Service.

Why the incumbent Cloud 1.0 storage providers can't catch up

It may be hard to believe, but first-generation cloud storage is nearly 15 years old. AWS S3 was launched in 2006. A lot of advances in storage technology have emerged since then:

- Higher-density disks for greater capacity
- Helium-filled drives for longer life cycles and lower power requirements for reduced operational costs
- Erasure coding for lower redundancy costs
- Advanced file system software for zettabyte scalability

These technological advances are not plug-and-play compatible with legacy systems. The hardware and software of AWS and other storage vendors that copied AWS' architecture may have been state of the art in 2006, but it can't evolve to take advantage of these innovations. AWS, for example, has exabytes of capacity spread across tens of trillions of objects and millions of drives. Given its scale and its significance to many organizations, change at this point is difficult at best and would require major architectural shifts.

What puts the sizzle in Wasabi hot cloud storage?

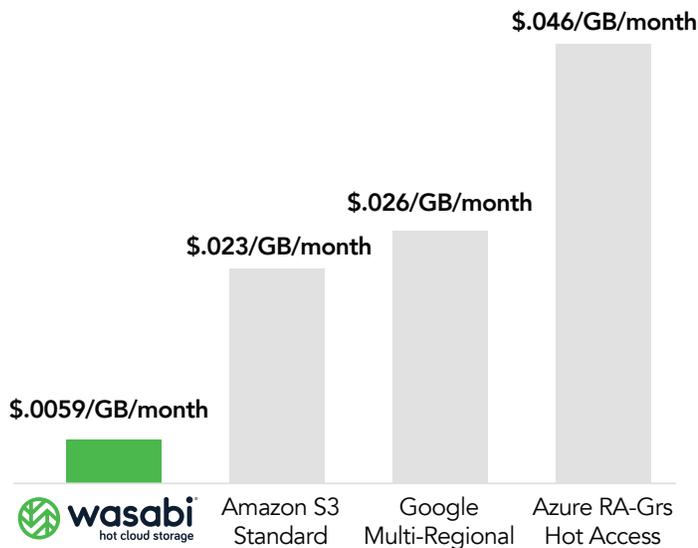
Wasabi's storage architecture allows them to deliver cloud storage that is significantly faster than Cloud 1.0 frequent-access storage at a cost that's less than traditional cold storage. They call it Wasabi hot cloud storage because all data is treated equally and made readily accessible regardless of how you classify it: hot, cool, or cold.

Ultra-low commodity pricing

Wasabi hot cloud storage pricing is straightforward and doesn't include any hidden fees: those small-but-relentless charges for retrieving data, getting a list of stored objects, creating or accessing objects in an application, or copying data from the cloud to other clouds or on-premises systems.

Wasabi doesn't have tiers or fee structures that require specialty consultants to help you make sense of it all. Costs are based on a simple, low-priced per-terabyte fee per month for object storage (see Figure 5). That's it. Scale up or down as needed without hassle.

Figure 5 Wasabi Pricing Compared to Other Leading Cloud Storage Providers



Source: Wasabi

Save even more with Reserved Capacity Storage

Customers who prefer to purchase storage capacity up front can save even more on data storage costs with Wasabi's Reserved Capacity Storage (RCS). RCS delivers true price predictability and simple one-time billing. You can commit to fixed-price terms of up to five years and cut storage costs by up to 27% compared to pay-as-you-go pricing. And you still won't be charged for data egress or API requests.

Figure 6 Cost Comparison of Wasabi RCS Plan vs. Standard Pricing of Typical On-Premises Solution

	Wasabi RCS	Vendor X
Volume	100 TB	100 TB
Duration	5 years	5 years
Price	Less than \$40,000	\$130,000 list price*
Conditions	Includes technical support	Doesn't include costs for maintenance, data center, or personnel

* Price increases to more than \$240,000 when maintenance costs over five years is included.

Source: Wasabi

Superior performance

As technology advances, so does the need for increased cloud storage performance. The read/write speed of Wasabi hot cloud storage is measured in single-digit milliseconds and Wasabi's superior performance has been consistently proven in the laboratory as well as in real-world enterprise applications.

Their hot cloud storage excels even in the most data-intensive, delay-sensitive environments, such as those that support advanced analytics. Even for Amazon Elastic Compute Cloud (EC2) customers, Wasabi proved faster than AWS S3, as shown in the [Wasabi Performance Benchmark Report](#).

Uncompromised reliability

On-premises and first-generation cloud storage providers achieve high durability by replicating data across multiple drives using various RAID (Redundant Array of Independent Disks) schemes. Next-generation cloud object storage services such as Wasabi achieve the same level of reliability, but far more efficiently, through the use of advanced erasure coding algorithms. With erasure coding, each object is transformed into a series of codes that are distributed across independent drives in separate storage servers. In the event of hardware failure or data corruption, the original data object can be reconstructed using a subset of these codes. This more effective use of storage capacity enables Wasabi to deliver eleven nines of durability at a lower cost than traditional RAID schemes—a savings that is passed along to customers in the form of lower prices.

Protection from ransomware and human error

Backing up your data to the cloud adds an additional layer of protection. However, your backup in the cloud can still be vulnerable to cybercriminals and human error. While the vast majority of attacks are initiated on-premises, bad actors have been known to gain access to a victim's cloud credentials and use them to delete previous backups or download them to servers under their control. In other less-malicious cases, misconfigured cloud object buckets (the containers that hold your data) have left private data open to the public and vulnerable to cybercriminals.

When you create a storage bucket in Wasabi, you have the option of making it immutable, which means that any data written to that bucket cannot be deleted or altered in any way—not even by a systems administrator. Wasabi is one of the few cloud service providers capable of providing this immutability feature. It not only prevents encryption by crypto ransomware but can help organizations comply with government and industry regulations such as HIPAA, CJIS, the Financial Industry Regulatory Authority (FINRA), and Markets in Financial Instruments Directive (MiFID), and many others.

Support for advanced analytics

As enterprises realize the value of advanced analytics, some still aren't capitalizing on that value because it's too expensive to store huge volumes of data to feed into analytics systems. With Wasabi, you get fast, inexpensive storage that affordably scales. You no longer have to make the hard decisions about which collections of data to keep, where to store it, or how long to hold on to it.

Wasabi hot cloud storage makes it simple and cost-effective to store all your data, for any purpose, for as long as you need it. It sets the stage for you to take advantage of AI and big data analytics to turn your data into valuable, actionable business information.

What business insights could you glean from a decade's worth of detailed sales or marketing data? How about extensive supply chain or distribution records? Or recorded customer service data? With the ability to parse massive datasets, you could improve productivity, reduce overhead costs, increase customer satisfaction, identify new revenue possibilities, streamline operations, and much more.

Revolutionary archival possibilities

Much like it allows you to store all your data for purposes of analytics, Wasabi hot cloud storage is also revolutionizing archival storage. Now, all your data has value and you no longer have to decide which data to jettison to keep costs low. You can store it all.

Rich partner ecosystem

One of the great things about Wasabi hot cloud storage is that they have partnered with more than 2,500 managed and cloud service providers, technology companies, and systems integrators to provide almost every storage-related service you can imagine, from data protection technologies to backup and recovery specialists to file management and data analytics.

Some cloud storage companies are trying to become all-in-one solutions that deliver it all. But no one can be the best at providing everything. With so many channel partners providing best-of-breed solutions for specific cloud storage use cases, you have the flexibility to choose the right solution for you while Wasabi concentrates on delivering the best cloud storage foundation possible to support those solutions.

Compatibility with legacy systems

With Wasabi hot cloud storage, you can avoid vendor lock-in because it is compatible with legacy cloud storage services and most common management tools and practices. Wasabi hot cloud storage was designed to be 100% compatible with AWS S3 through a simple, standards-based REST API—so all of your existing S3 storage management apps work seamlessly with Wasabi hot cloud storage.

Backup and recovery solutions powered by Wasabi

Backing up to the cloud doesn't have to be a costly and complicated affair. Wasabi hot cloud storage makes it simple to cost-effectively hold on to more data, longer—so you can rest easy knowing that the lifeblood of your business is safe and protected.

- ✓ **Affordable**
- ✓ **Predictable**
- ✓ **Always available**
- ✓ **Sign up in seconds**
- ✓ **Write backups to Wasabi in minutes**

Say no to complicated tiers. No to nickel-and-diming your budget to death with hidden fees. No to the fear that you've accidentally thrown out critical data. No to the difficult decisions on what to store where and for how long. With Wasabi, you get up-front, straightforward, no-nonsense, accessible storage that's less expensive and faster than the other guys. What could be simpler?

For more information and a list of Wasabi-powered backup and recovery solution partners, visit [Wasabi.com/backup-and-recovery](https://wasabi.com/backup-and-recovery)

A COMPARISON BETWEEN ON-PREMISES, CLOUD 1.0, AND CLOUD 2.0 STORAGE SOLUTIONS

	On-Prem	Cloud 1.0	Cloud 2.0
FINANCIAL ATTRIBUTES			
No capital equipment outlays		✓	✓
Pay-as-you-grow, on-demand scalability		✓	✓
No recurring power, cooling, and rack space expenses		✓	✓
No equipment maintenance, admin, and support burden		✓	✓
Easy-to-understand, universal storage solution (no tiers)			✓
Commodity pricing			✓
FUNCTIONAL ATTRIBUTES			
Strong security and control	✓	✓	✓
Rapid read/write speeds	✓		✓
Inherent resiliency		✓	✓
Eleven 9s data durability		✓	✓
Configurable data immutability	✓	✓	✓
Vendor independence (no lock-in)			✓

Appendix

Additional Resources

Does it Pay to Move from On-Premises to Public Cloud Storage?

A simple cost analysis: On-premises storage vs. AWS S3 Standard

The Future of On-Premises Storage: Not So Hot

A look at the total cost of ownership over a 5-year period

Why Cloud Storage Hasn't Killed LTO Tape

The real total cost of ownership of LTO tape storage vs. Cloud 1.0 storage tiers

About the Authors

David Boland, Product Marketing Director, Wasabi

As Wasabi's Director of Product Marketing, David plays an integral role in communicating their mission of making saving data simple, affordable, fast and secure. He has served in various sales, marketing, and product management roles at companies at all stages of growth from startup to Fortune 100, including NetApp, Juniper Networks, Lucent Technologies, and Cabletron Systems.

CONTACT:

dboland@wasabi.com

www.linkedin.com/in/davidlboland

David Ball and Erin Poulson, B2B Writers, Inc

B2B Writers is a leading hi-tech research and thought leadership content firm serving the information technology industry. For more information, visit www.b2bwriters.com.



wasabi
hot cloud storage

wasabi.com