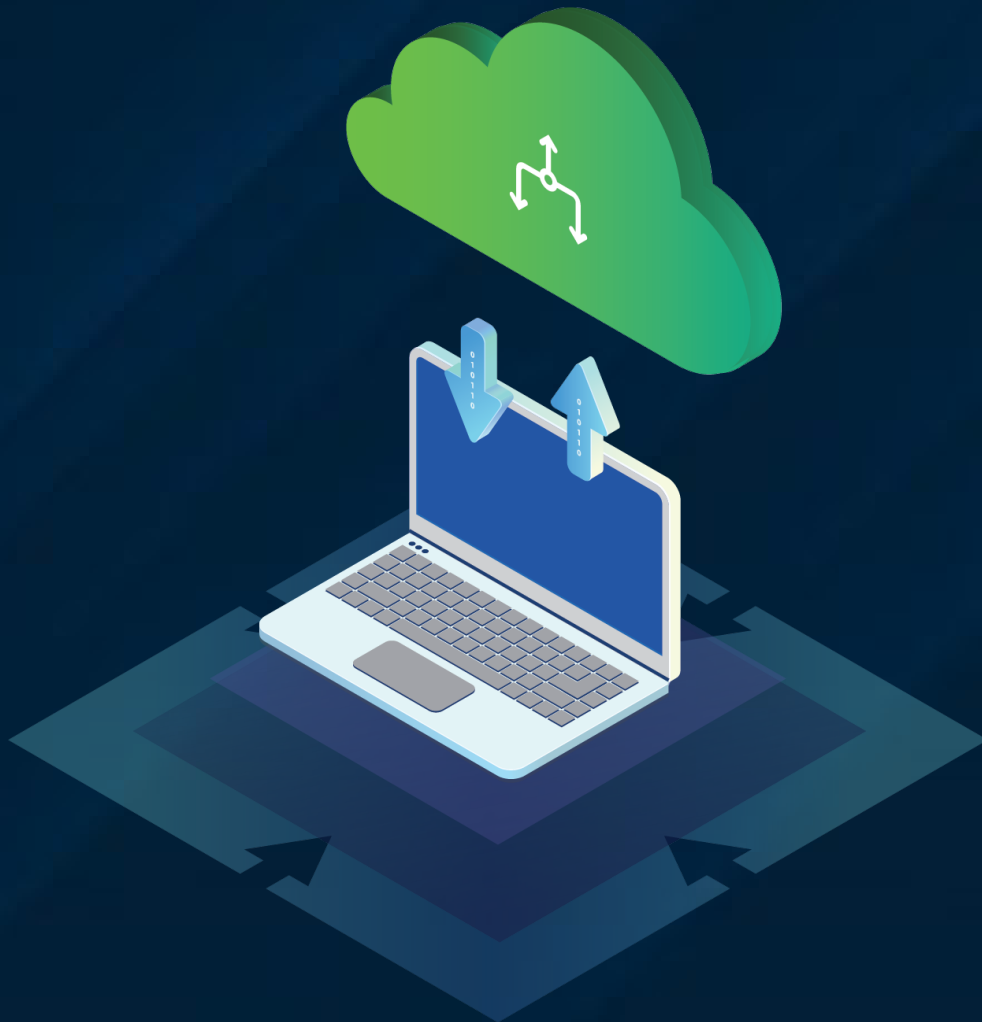


Migrating Your Company to the Cloud

The Complete Guide to Extending or Moving
Your Storage Infrastructure to the Public Cloud



It doesn't make any more sense these days to run your own on-premises storage than it would to run your own electric generating plant.

Storage is quickly becoming a cheap, fast, one-size-fits all commodity. That's why most companies – especially enterprise-scale companies – are moving their data to the cloud. What's driving the widening cost gap between on-premises storage and cloud storage?

There are a lot of hidden costs to running on-premises storage. Some of these are hard costs, such as equipment purchase, data center costs, electricity, and manpower. Other costs are soft costs, such as the distraction factor, the fact that you have IT personnel baby-sitting storage servers instead of working on things that will really move your business forward, or the risk that your on-premises equipment could become obsolete more quickly than planned. And while the cost of cloud storage has historically dropped dramatically over time, once you purchase on-premises storage, you're stuck with the cost.

Luckily, there's a better way...

CHAPTER 01

The Problems with On-Premises Storage

On-premises (on-prem) storage technologies have advanced significantly; and in many cases, the legacy approach to using on-prem storage still holds. Though the advantages of on-premises solutions are vanishing as companies move to a cloud-native or cloud-first approach, there are still some key benefits to storing and computing data on-premise:

- Data security is in the hands of your organization, and with that control can come peace of mind
- Ability for extreme customization
- Control over implementation and upgrade processes

While those can be seen as benefits, depending on your IT and business strategies, they can also be seen as distractions and burdens.

Nonetheless, the amount of data that is now being produced is outgrowing the ability to store it all at a reasonable price using on-prem solutions.

On top of that, companies recognize the increased value of the data that's being generated to their business. From day-to-day operations and logistics information to all of your software and applications, there is no shortage of critically important information being captured every second of every day. So what are the [problems with continuing down the tried and true path of a robust on-prem storage](#) approach?

Cost

On the surface, on-prem storage seems like it would be a cost-effective approach, given that you can specify every aspect of your implementation. But, many of these costs become hidden amongst your organization's larger operating budgets, cemented into the annual overhead with little regard for revisiting or optimizing that spend. CFOs and CTOs tasked with cost management would do well to know, with confidence, where on-prem dollars are going:

Capital Expenditures (CapEx) – With any on-prem approach, there are significant equipment costs upfront. Also, the initial ROI often takes years to recognize due to the upfront costs for equipment, personnel, training, etc.

Operating Expenditures (OpEx) – With any large on-prem setup, you'll have any or all of the following operational costs:

- Recurring power and backup power sources
- Cooling equipment
- Rack space for the servers and other equipment
- Annual hardware and software maintenance and support fees
- Administrative costs and a employee dedicated to administration and upkeep

And when [compared against the cost of cloud](#), its apparent the disadvantage of using on-prem storage will continue to grow as the cloud becomes more affordable, capable and unavoidable as the preferred choice for your infrastructure. The operating efficiencies available to "hyperscale" public cloud providers will continue to drive costs down, and streamline aspects of infrastructure that are simply difficult for other enterprises to address.

Complexity of Configuration & Administration

Legacy storage systems are notoriously difficult to configure and administer. Professionals who work on this equipment and software will need ongoing, specialized training or you'll by paying for expert support from your vendors. As your on-prem system becomes more and more legacy, upkeep requires more specialization and therefore a higher pay scale.

Data Backup and Disaster Recovery

One of the most important roles of your storage infrastructure is your ability to back-up and protect your data in case of a problem – whether it's a small scale file recovery need or a full-blown disaster recovery scenario, failing over to a remote location or straight into the cloud.

As your company grows, this issue becomes more pressing while the solution becomes more difficult. It's likely that some portion of your backups are stored physically offsite as part of your backup/DR strategy, resulting in lost time and expenses for retrieval as they need to be pulled from offsite storage and put into play. Delays of hours if not days are common in these scenarios.

BLOG POST

Why We Need More Visibility and Less Complexity in the Cloud

[READ MORE](#)

Scalability

The lack of scalability is perhaps the biggest drawback to on-prem storage, especially for companies worried about what use cases are coming next.

Your on-prem systems are only as scalable as the hardware and software it's built with. 5-year-old servers aren't nearly as capable of doing what new servers are, and those now-new servers won't be nearly as capable as servers 5 years from today, with edge computing and IoT (Internet of Things) capabilities growing day by day.

Continuous innovation is never going away – both on what your business wants to do, and technology innovations in storage and other areas. In order to remain scalable with an on-prem approach, you'll need to purchase new equipment to keep pace. It's unavoidable.

The further and further your software and hardware gets away from its manufacture date, the more likely it is for vendors to discontinue support plans, retire products, or start charging for the advanced upgrades needed to keep the equipment at least close to on-par with what is currently in the market.

In other words, your scalability ceiling is relative to the time in which your "scalable" systems were deployed, and today's ceiling is in the cloud.

Security

Securing your on-site data can be a bit more challenging if you are doing it yourself. But it's not because your staff isn't capable of configuring your network and data to be secure. It's more about the [cybersecurity landscape as a whole continuously evolving](#). Both financial service and healthcare organizations, for example, need to remain especially in tune with the ever-changing cybersecurity standards in their industry.

Additionally, many internal security initiatives are focused on eliminating access to your systems from outside or foreign devices. But that approach often falls short of completely eliminating malicious breaches that originate internally. Sure, you can block all the hackers in the world from accessing your systems, but sometimes all it takes is one ill-intentioned employee or one lapse in user judgment. When these activities occur inside of your firewalls, it can be devastating.

Having 11 nines of data durability through a data storage solution like Wasabi is critically important for the reliability of the system itself, but companies are realizing that it's not enough to protect from technological issues. In fact, the technological issues are much easier to solve. It's people you should be more concerned about.

This gets to the "fat finger problem." Take a look at some of the recent major data losses and you'll notice they all share the common thread of being caused by human error. When an Oakland, [California police department lost 25% of its body cam videos](#), when [Amazon accidentally took down a large portion of the internet](#), and when [hundreds of companies lost data due to ransomware](#), a human was to blame in each case.

One way to minimize human error as a threat to your data is to make it impossible for an "oops" to turn into a career- and business-limiting issue. [Immutable storage](#) functionality, for example, makes your data read-only once it's written, by whatever policy you decide – the data cannot be deleted or overwritten until the policy rules expire.

Without immutable data, it won't matter how many 9s of durability you have, your system will still be vulnerable to human error – which security studies show happens far more often than system errors or attacks by hackers.

CHAPTER 02

So, Why Migrate to the Cloud?

When you had no other option than to own your own infrastructure, life was quite different than the choices you have today. There are more and more reasons to take the plunge into cloud computing and storage, each becoming more convincing as time passes. CIOs, CTOs, VPs of IT, SysAdmins and other IT decision makers often need to bring a convincing business case for their company to migrate to the cloud. Luckily, the benefits of bringing cloud computing to the business can be easy and exciting for CFOs and CEOs to understand.

Cost

The cost of Cloud 2.0 implementations are significantly less expensive than Cloud 1.0 and legacy on-prem implementations. Cloud 2.0 is a term used to describe the next generation of highly optimized and efficient cloud services. Cloud 2.0 offerings differ from first generation cloud through radically reduced yet enterprise-friendly pricing compared to Cloud 1.0 offerings, while simultaneously taking advantage of technology innovations that result in much higher performance.

REPORT

The New Economics of Cloud Storage

READ MORE

While Cloud 1.0 (the first generation cloud vendors) like Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP) enabled the first wave of migration to the cloud, the various costs associated with on-premises storage are equivalent with [what it costs to leverage the cloud for storing your data](#). The specific cost benefits with Cloud 1.0 and Cloud 2.0 providers include:

Equipment investments are non-existent.

You'll never have to replace outdated hardware, and recurring hardware maintenance fees are eliminated.

Operating expenses are minor.

You don't need an entire team to understand the advanced configuration complexities, nor extra physical office space and utilities to house and power your servers.

While there are more [in-depth cost comparisons between on-premises models, and various cloud service provider models](#), below is a rough breakdown of how the costs of Cloud 1.0, Cloud 2.0 and on-premises storage compare:

	On Premises	Amazon S3 (Cloud 1.0)	Wasabi (Cloud 2.0)
Equipment costs depreciated over 5 years	\$69K	N/A	N/A
Maintenance fees ¹	\$62K	N/A	N/A
Power and networking	\$87K	N/A	N/A
Staffing and overhead	\$52K	N/A	N/A
Storage costs ²	N/A	\$289K	\$73K
Total costs	\$270K	\$289K	\$73K

¹ assuming 20% annual maintenance and support
² assuming average 20% egress download per month

Ease of Maintenance

With the cloud, much of the maintenance and upkeep initiatives stay mostly out of sight, out of mind, and away from your budget. But, it's not because the systems aren't maintained. It's actually the opposite.

Maintenance and regular updates are built into the cloud service models. Value-wise, cloud service vendors know that their products are only a key differentiator if they can keep pace with the technology advancing around them. Therefore, to continue serving their customer base without interruption, they've become adept at advanced techniques for implementing large systematic updates and new features without disruption. Even scaling up is on-demand, and you can expand capacity instantaneously and effortlessly.

Reliability & Durability

Your cloud data isn't exactly sitting on a single drive somewhere in Drivertown, USA. Historically, storage solutions have achieved high durability by replicating data across multiple drives using various [RAID](#) (Redundant Array of Independent Disks) schemes. Though it has evolved over time, RAID is an expensive and outdated (30+ years old) technique from both a price and performance perspective. Cloud 1.0 storage solutions all grew out of the state of the art technology available when early clouds were conceived, which was RAID.

With Cloud 2.0 storage, everything has become even more reliable and more durable. 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. With this, Wasabi is able to provide eleven [9s of data durability](#), fully protecting customer data without the overhead of maintaining duplicate copies on multiple disks.

BLOG POST

Availability and Durability – Why You Need to Consider Both When Choosing a Storage Vendor

READ MORE

Alleviating Bandwidth Concerns

One of the initial concerns about cloud storage, especially in the first wave, was the high cost of bandwidth needed to truly leverage it as intended. But, like any [commoditized utility](#), the cost of accessing high speed internet has dropped significantly. And as bandwidth becomes both faster and more affordable, scalability within the cloud become infinite.

New file transfer acceleration techniques also make it possible to squeeze more throughput out of your bandwidth than you may realize. The revolution in connectivity is tightly tied to several trends here that make Cloud 2.0 options significantly different than your choices just 5-10 years ago.

Infinitely scalable?

Sounds like an exaggeration, right? It's actually not far fetched at all.

While Cloud 1.0 data storage options were theoretically infinitely scalable, their performance was still tied to the limitations of on-prem solutions – as Cloud 1.0 storage was built on the same technology as your legacy on-prem options.

Combine that poor performance with the relatively high price of Cloud 1.0 (nearly the same as on-prem), and the hidden costs of getting data back OUT of Cloud 1.0 storage solutions, and infinite scalability becomes infinitely expensive.

Contrast that with Cloud 2.0 storage solutions like Wasabi, where performance is 6x faster than AWS S3 and 80% less expensive, with no hidden and unpredictable fees, and now infinite scalability can truly be yours WITHOUT the infinite headache of explaining to your CFO why your storage costs were twice as much as you'd expected them to be.

Security

Cloud 2.0 technology has laid to rest many of the concerns over data security. Take, for instance, the highly regulated healthcare industry. HIPAA and HITECH regulations represent some of the more stringent compliance requirements around electronic health information.

Like many industries, healthcare regulatory organizations conduct regular and frequent assessments of cloud vendor security systems to ensure that they meet or exceed compliance standards.

An independent compliance assessment firm, Schellman & Company, recently conducted a [thorough audit of Wasabi's Cloud security architecture](#), systems, and practices. After a lengthy evaluation, they confirmed that [Wasabi complies with all HIPAA security and privacy rules](#) for protected health information.

Securing confidential and proprietary data will always be a hot-button issue, and cloud service providers, like Wasabi, remain committed to ongoing research and implementation of ever-improved ways to protect the information that powers your business. Forward-facing cybersecurity initiatives are ingrained within cloud providers and their services are designed to meet stringent security and availability requirements, including:

FERPA

The Family Educational Rights and Privacy Act is a federal law that allows parents access to their children's education records, as well as some control over disclosure of personally identifiable contained within the information.

HIPAA/HITECH

Laws regarding the standards around health care information, electronic billing and other processes, ensuring the protection and confidential handling of protected health information for individuals.

CJIS

The Criminal Justice Information Services makes up a division of the FBI that provides a range of intelligent tools and services to law enforcement, national security and intelligence agencies across the country, and sets certain standards around how this type electronic information is shared and stored.

GDPR

The European based law on data protection and privacy for all individuals within the EU and other European regions, designed to give citizens more control over their personal data and the ability to consent on ways it can and cannot be used. GDPR affects nearly website that collects data from European users.

MPAA

The Motion Picture Association of America maintains a set of compliance standards around the security of how content is created, managed, stored and distributed. The MPAA also controls how movies are rated (based on content) when released to the public.

[Wasabi takes your data security extremely seriously](#), from physical security, to a secure network architecture, policies, procedures and features to enforce your own data privacy and security needs, as well as built in data durability and protection features.

Additionally, Wasabi equips users with [immutable storage](#) capabilities, making your data impossible to delete or alter. Aside from the obvious security reassurances, another key advantage immutable storage is the enhanced auditability it provides your team and/or outside auditors.

If raw data is requested as part of a compliance audit, organizations and cloud service providers alike can provide investigators with not only the unchanged data, but also with the technical specifics of how that data is immutably protected.

BLOG POST

What Does 11 Nines of Durability Really Mean?

[READ MORE](#)

Extend the Life of Previous On-Premises Storage Investments

Just because you are migrating to the “cloud” doesn’t necessarily mean that your data will be 100% stored off-site. Cloud storage is ideal for a variety of applications, including primary storage for on-premises or cloud-based workloads, secondary storage for backup and recovery, and archival storage for long-term data retention. This versatility protects and extends previous storage infrastructure investments, allowing you to migrate to the cloud on your own terms. Use cases might include:

Cap-and-grow

Keep existing applications on legacy on-prem storage and use the cloud as primary storage for new applications like Big Data, analytics and IoT. Transition existing applications to the cloud over time and gradually decommission legacy storage platforms.

Lift-and-shift

Use the cloud as primary storage when moving on-premises applications to the cloud en masse.

Backup and recovery

Protect on-premises storage for a fraction of the cost of alternative solutions.

Archival Storage

Cloud offers affordable and durable long-term data retention. Move infrequently accessed data to cloud-based solutions to free up expensive legacy storage capacity and better align storage costs with data value.

Data center transformation

Cloud solutions help with transitional storage to support corporate mergers and acquisitions, or data center relocation or consolidation initiatives.

CHAPTER 03

Planning Your Migration to the Cloud

Before the initial migration, you'll first choose the type(s) of cloud storage that suits your needs. The type of storage solution you choose should be tied to clear business objectives based on:

- Cost of your storage.
- Life expectancy of your data.
- How often the data will be retrieved or modified.
- How fast the access to the data needs to be.

Based on your organizational objectives and unique data needs, and prior to any migration implementation, you'll need to designate what types of storage "buckets" for which each of your cloud-ready datasets is to be stored.

Cloud Storage Tiers

At the heart of it, there are four main "tiers" of storage:

- Immediate use
- Disaster recovery (a mission critical form of backup)
- Backup (copies that may be hours or weeks behind the most current versions)
- Archival (copies of files that are no longer actively needed, but stored "just in case" it's needed, typically with delays to retrieve the data).

These tiers are often referred to as having retrieval performance characteristics of being either hot (immediate, very fast) or cold (hours or days to get access to).

Hot & Cold Storage Challenges

For the most part, organizations keep infrequently accessed data in cold storage, and everything else is in hot storage.

With on-prem storage, hot storage is typically quite expensive, yet high performance. On-prem cold storage is relatively cheap, and correspondingly, slow to "unthaw" from an offsite location such as a physical storage warehouse, salt mine, etc.

To cover all your bases with on-prem storage options, you need to map out the different data use cases, which ends up creating quite a lot of complexity. Relative to your data, you'll ultimately be determining: What data goes in cold storage? What data goes in hot storage? What data is somewhere in the middle, and how many storage options do I need in-between the two extremes? How and where do users access each of these storage tiers?

Perhaps some of the best news for your cloud migration strategy...

When you're migrating to the cloud, it's now possible to remove some of the old constraints of legacy on-prem storage solutions, and take advantage of the new realities of Cloud 2.0 storage.

Shedding Tiers

With Cloud 1.0 storage providers, you have often have a bewildering variety of tiers to choose from – Standard, frequent access, infrequent access, one zone, multiple zones, intelligent tiering... the pricing pages for these options are quite literally many pages long, and that's for just ONE piece of your cloud computing infrastructure.

This translates into cost savings that are almost no different than on-prem deployments – with the main benefit being that you are shifting from CapEX + OpEX to pure OpEX cost tracking. But based on Cloud 2.0 storage features, whether or not you need to worry about matching your specific data need to multiple tiers of storage at all is really the only question.

With Wasabi, a single tier provides better performance than the fastest options of most Cloud 1.0 storage providers, at a price point that is far less expensive than even their slowest (cold storage) tier.

Why compromise on speed when you can cover all of your hot to cloud storage needs with a single tier?

BLOG POST

Wasabi Wipes Away the Tiers

READ MORE

CHAPTER 04

Hybrid Approaches

Until your company is fully cloud-native, you will need some type of on-prem storage, especially considering the many ways you can optimize to [squeeze the most out of your on-prem storage](#).

A [hybrid or multi-cloud approach](#) is common and cloud providers can usually accommodate businesses who wish to incorporate some form of on-prem storage into their IT architectures.

Why Go Hybrid?

A hybrid approach has several benefits. For one, migrating to the cloud is a big decision for any organization, especially when you've already invested in on-prem hardware. Implementing a hybrid strategy allows you get some extra life out of that expensive hardware, and leverage systems that would be difficult to retro-fit for a cloud-based storage model.

While hybrid clouds and multi-clouds terms are often used interchangeably, there are some key differences to note.



<https://www.youtube.com/watch?v=a43Cw7b2w9A>

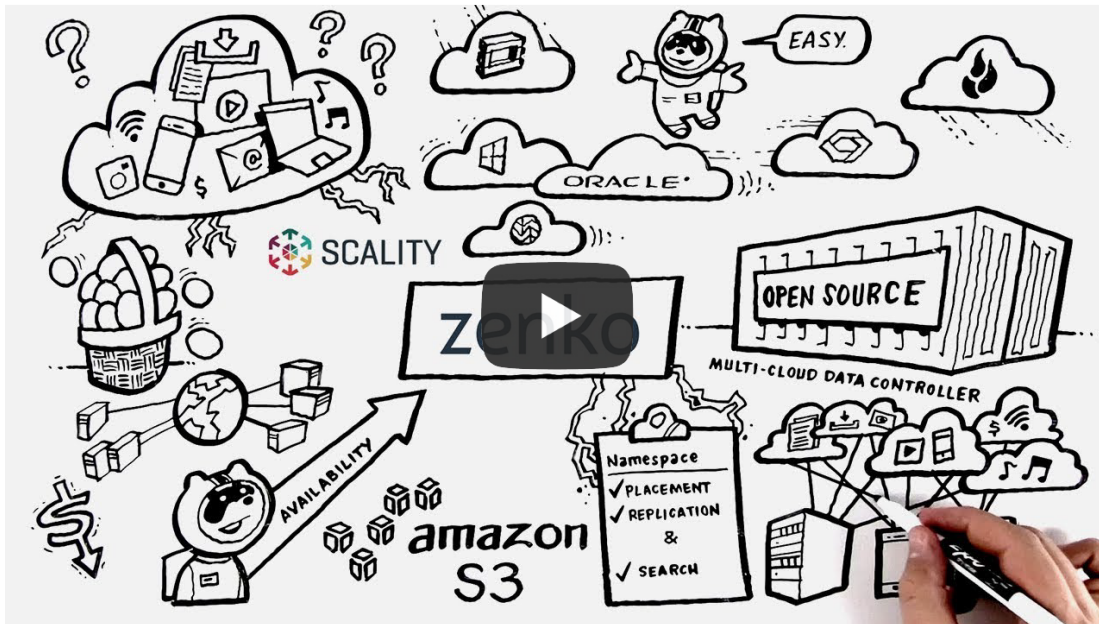
Multi-Cloud

A multi-cloud strategy is the use of more than one cloud provider, usually consisting of 2 or more public clouds. Because different departments, use cases and applications may have different requirements, a multi-cloud strategy provides flexibility in selecting the right cloud for the right job.

Using a multi-cloud also protects your company and its data in the rare instance that one of the clouds fails. It also provides you the safety of not being locked into any single cloud vendor, held ransom by their business model and the hidden fees that make up so many of the big names in public cloud computing.

Hybrid Cloud

A hybrid cloud is more commonly used to describe an approach a solution that lives partly within your on-prem infrastructure, with connections to one or more cloud providers.



<https://youtu.be/nUKw7lhrSNo>

True multi-cloud storage requires:

- Support for leading and next-generation public cloud storage services, or de facto API standards like S3
- Federated security management to eliminate the complexity and variations in cloud security models
- The ability to store data in native cloud format for open access by cloud apps and services
- Simple policy-based data management

Until recently, multi-cloud strategies weren't really a viable option for most companies due to the different proprietary formats across public clouds and gateway products. Each public cloud speaks a different language with their APIs, and everything from security to authentication is handled differently by each as well. Wasabi has helped lead the charge in multi-cloud technologies and integrations through [strategic partnerships with industry leaders](#), and compliance with the S3 API.

CHAPTER 05

The Cloud Migration Process

How long does migrating to the cloud take?

The process for migrating from on-premises to cloud storage is largely dependent on your data (size, structure, security needs) and your network connection speed. If, for example, you are one of the lucky few companies to have access to 100 gbps network speeds, and the size of your transferring data is less than 1 PB, then you could reasonably expect that the initial migration could occur within a few weeks. A slightly more common, yet still high, 10 gbps would take 10 times longer; and at 1 gbps, you're looking at a pretty lengthy migration period if you're doing the entire data transfer over the public internet.

No one wants the cloud migration process to drag out for several months if there is an option to safely increase the speed. Knowing your options ahead of time will help you to better understand the approach that's best for your company, whether it's over a public network, public network using file acceleration, a direct connection, via a physical transfer, or some combination.

Public Network Connection

With a public network connection, you would simply leverage your existing internet connection to interface with your cloud storage. Unless you're saddled by slow network connectivity or strict data security requirements, the public network connection is suitable for most organizations both during the migration, and for on-going connectivity to your cloud resources.

Public Network Connection + File Acceleration

Legacy file transfer techniques (like FTP and SFTP for example) are inefficient compared to more modern techniques. With much higher bandwidth connections, and much larger files, there are a variety of solutions available that can squeeze much higher throughput out of your connectivity than you may have thought possible.

Direct Connect

Cloud service providers all offer [direct connect options](#) for organizations that require private connectivity directly to the cloud data centers. A migration via direct connect is much faster and more secure than using a similar speed public network connection. You should set up a direct connection if you will also need post-migration direct connectivity between your facilities and your cloud service provider. If a direct connection is not really a necessity other than during the migration, a physical or bulk transfer might be the better approach.

Physical / Bulk Transfer

A physical or bulk migration is sort of like the good old days of saving something to a disk, then taking the disk to another location to then open and access that same data – except on a much larger scale. [The Wasabi Ball Transfer Appliance](#), for example, is a cost-effective and convenient migration option when your network connection speed is lacking. Keep in mind, even if your network connection speed is high, the sheer volume of the data you plan to migrate could lead to a lengthy transfer time. The Wasabi Ball provides the ability to move large datasets quickly and conveniently. With some customers, Wasabi ships multiple Wasabi Balls at a time in order to move large volumes of data per week.

CASE STUDY

IQ Media

[READ MORE](#)

CHAPTER 06

After the Migration

Once the initial migration has ended, you'll continue using either a direct or public network connection to read, write, delete or update your cloud data, whichever your use cases require, and that suit your budget.

Cloud system providers like Wasabi [partner with both innovative cloud technology and software/hardware vendors](#) of all types so that you can create a robust and flexible cloud storage stack personalized to your company's specific needs. If there isn't a solution specific to your sector, or for whatever challenge you might have with your cloud system, then it's probably in development right now.

What's your cloud migration strategy? [Try a 30 day free trial](#) to test the approach that makes sense to you and start your migration today!

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 fees for egress or API calls.

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..



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

©2019 Wasabi Technologies, Inc. All rights reserved. WASABI and the WASABI Logo are trademarks of Wasabi Technologies, Inc. and may not be used without permission of Wasabi Technologies, Inc. All other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s).

 **wasabi**
hot cloud storage
www.wasabi.com