



WHITE PAPER

The Cost of Cloud Storage

Table of Contents

02 About Backblaze

03 Overview

04 Breaking Down Cloud Costs

06 Storage Costs

- Tiers
- Regions
- Object size
- Upload fees
- Retention periods
- Storage cost strategy

10 Download Costs

- Egress fees
- Retrieval fees
- Download cost strategy

13 Transaction Costs

- API calls
- Object Lock fees
- Miscellaneous additional costs


15 Conclusion: Demystifying Cloud Storage Costs



About Backblaze

Backblaze B2 serves over 100,000 organizations with affordable, secure, easy to understand object storage. When using a service like [Backblaze B2 for your data storage](#), you can trust that your [pricing](#) will be simple and predictable with no tiers and a flat rate no matter the region, object size, method of upload, or how long your data is stored, meaning you can do more with your budget and spend time on important strategic initiatives rather than managing storage.

Backblaze B2's S3-compatible API works with software like Veeam, Commvault, Rubrik, and more to enable Object lock immutability. And with always-hot storage at cold storage pricing, you can fortify your disaster recovery plan and protect against ransomware without any worry about retrieval delays. No tiers, no minimum storage requirements, no hidden fees.



Overview

It's no surprise that enterprise data is growing exponentially. The rise in big data, artificial intelligence, machine learning, the internet of things, and business analytics results in the creation of 120 zettabytes of data each year (for reference, 1 zettabyte=1,000,000 petabytes)*. Accommodating this scale of data puts strain on already tight IT budgets, especially when certain types of data must be retained for several years—either due to compliance requirements like HIPAA, cybersecurity insurance requirements, or other archival needs.

Many organizations maintain a robust on-premises data storage infrastructure to handle this volume of data, but the speed of data growth makes staying fully on-premises unaffordable and unsustainable.

To offset large CapEx investments in additional on-premises infrastructure, many businesses are turning to the cloud. The public cloud offers many tangible benefits, from helping to accelerate development of services, to storing and managing raw datasets, to simplifying workflows, providing backup and archival storage, and more. Indeed, moving data to the cloud can be a great cost-saving measure for organizations... or it can lead to cloud costs spiraling out of control. Better understanding cloud storage costs can help businesses do more with their budgets, even as revenue and cost pressures loom.

In this white paper, we'll take a look at the ways some cloud storage vendors can inflate your data storage costs across three cost categories and steps you can take to contain your expenses. With the right cloud storage provider, you can achieve scalable data storage while still preserving affordability.

[*Source →](#)

Breaking Down Cloud Costs



In short, there are only three cost categories for cloud storage expenses:



01. Storage

The rental for your slice of the cloud.



02. Download

The fees to bring your data out of the cloud.



03. Transactions

Charges for “stuff” you might do to your data inside the cloud.


You probably know someone (you?) that was more than surprised when their cloud storage bill arrived. They have good company, as according to [a recent report](#) by Anodot, a business monitoring company, nearly half of businesses (49%) find it difficult to get cloud costs under control. Hyperscalers like AWS and Azure are common but their pricing is notoriously difficult to unpack and full of “gotchas” you can’t predict. In this white paper, we’re taking a specific look at object storage, but cloud drives like Google Drive and Dropbox have their own [tricky pricing practices](#) as well. That’s not to say all cloud storage providers have bad intentions. Cloud providers can be transparent and offer fair pricing. But as you’re approaching a cloud storage procurement process, it’s helpful to know what to look out for.

We’ll break down each of the categories listed to the left to help you better understand cloud costs so you can make the most of cloud cost efficiencies.



Nearly half of businesses find it difficult to get cloud costs under control.

Storage Costs



Data storage costs continue to pose a significant financial challenge. With the exponential growth of digital information, businesses grapple with the need for expansive storage solutions that don't break the bank. So, what are the key considerations to be aware of when it comes to storage? They include:

- Tiers
- Regions
- Object size
- Upload fees
- Retention periods



Tiers

Tiers define different levels of pricing according to either the type of storage (e.g. hot vs cold cloud storage) or the amount of data stored. In general, the words “**Next**” or “**Over**” on a pricing table are never a good thing. Here’s an example from AWS S3’s Storage Pricing for the U.S. East region as of March 2024:

- First 50TB / Month \$0.023 per GB
- **Next** 450TB / Month \$0.022 per GB
- **Over** 500TB / Month \$0.021 per GB

Those words mean there are tiers in the pricing table which, in this case, means you have to reach a specific level to get better pricing. You don’t get a retroactive discount—only the data above the minimum threshold enjoys the lower price.

**In this example, we are using 1TB as equivalent to 1000GB. AWS actually considers 1TB to be equivalent to 1024GB, meaning the actual pricing will be slightly higher than what is shown here.*

The mistake sometimes made is calculating your entire storage cost based on the level for that amount of storage. For example, if you had 600TB of storage, you could wrongly multiply as follows*:

$(600,000 \times 0.021) = \$12,600/\text{month}$

When, in fact, you should do the following:

$(50,000 \times 0.023) + (450,000 \times 0.022) + (100,000 \times 0.021) = \$13,150/\text{month}$

That was just for storage. Make sure you consider the tiered pricing tables for data retrieval as well. Note that there are certainly some cloud service providers that do not have tiers (like Backblaze); they may have a single cost for storage, no matter how much is stored or where it is stored. Be sure to specifically ask if your solution provider has tiers and how they are defined.

Regions

In the AWS S3 example on the previous page, we showed the pricing for the US-East (N. Virginia) region. However, costs can also vary according to the region you choose. If you instead choose to store your data in US-West (Northern California) the cost table now looks like this (as of March 2024):

- First 50TB / Month \$0.026 per GB
- Next 450TB / Month \$0.025 per GB
- Over 500TB / Month \$0.024 per GB

Using the same example for 600TB of storage, you have:

$$(50,000 \times 0.026) + (450,000 \times 0.025) + (100,000 \times 0.024) = \$15,150/\text{month}$$

If you want to intentionally store your data in a region far from your production data for better geographic separation, be sure to check costs for that specific data center.



Cost-effective cloud storage

For Centerbase, a software as a service (SaaS) platform serving high-performing legal practices, nothing is more important than security and performance. While they run a robust, replicated, on-premises data storage system, they found that they could not meet their desired recovery time objectives (RTO) if they faced a disaster that took out both their production and disaster recovery (DR) sites.

Read more about why they chose Backblaze to secure their data off-site and with Object Lock for ransomware protection.

[Read the Customer Story](#)

Object size

Some cloud storage providers (CSPs) require a minimum object size. That could be 4KB, 128KB, or some other minimum requirement. If you think you might be storing lots of small files in your data set, take a look at vendors' minimum file size requirements, often outlined in small print or in an FAQ. Object size requirements mean you will be charged the storage cost for the minimum required size for each object that is below that threshold.

Upload fees

Rarely is there a fee for uploading data to a cloud object storage service. But you should definitely be aware of a couple tricky upload expenses in particular. Cold storage classes, like Amazon's Glacier Instant Retrieval (GIR), do not charge for direct upload but they do charge for PUT requests to upload an object and also for lifecycle transitions that are commonly used to move data from one S3 class to GIR instead. Be aware of these specific fees when considering GIR because they will increase your overall storage rate.


Retention period

The last thing to watch out for is a [minimum retention period](#) requirement, sometimes also called a minimum storage duration requirement. This refers to the length of time you keep the object in cloud storage. If it's required to keep the object in storage for, say, 90 days, and you delete the object before that period of time, you will still be charged for 90 days of storage. Even more important to note is that "delete" in this context also includes overwriting. For instance, if you keep full backups for 60 days and then overwrite them with the newer version, that will be considered deleting the previous file. This particular requirement can significantly increase costs if you're not careful about configuring your backup settings to account for retention requirements.

Storage cost strategy

It can be tempting to simply look for the service with the lowest cost per GB, but you cannot compare apples to oranges when sourcing cloud storage vendors. When comparing CSPs, be sure to carefully examine pricing tables and FAQs to understand how the aforementioned factors can increase the baseline \$ per GB.

Download Costs



Download costs, often referred to as egress fees, are simply the costs for pulling your data back down from the cloud. Businesses often use the public cloud for [backup and archive](#), and it's a common mistake to think that downloads won't be needed. However, there are a few specific scenarios where egress can come into play.

01


First, there's the "break glass in case of emergency scenario." This is when you may download data due to either minor mishaps (e.g. accidental deletion) or major ones (e.g. a tornado wipes out your DR site). This is precisely why you keep a backup in the cloud, right?

02

The second scenario is [disaster recovery](#) testing. If you want to really craft a rock solid DR plan, you absolutely must test your backups by conducting a restore at least once every six months, if not more often (and we do suggest you do this more often). As you test drawing your backups back down from the cloud you'll incur egress fees, unless your download amount falls into a free egress threshold (more on that in a bit).

03

The last scenario that's less understood but still fairly common is when your backup software (e.g. Veeam, Commvault, etc.) downloads your data from the cloud as part of its workflow. This can happen for various reasons—either during a deduplication or synchronization process or something else. This download amount can be somewhat unpredictable but it is substantial, especially if your CSP doesn't provide any free egress.



Now, with that context, let's take a closer look at how cloud storage egress costs work.

Egress fees

Some cloud providers provide an amount of free egress and customers are only charged if they download more than the free amount. For instance, your CSP might allow something like 3x your storage average in free egress. If you store, say, 100TB of data, you would be allowed to download 300TB each month. If your CSP provides 1x free egress, you would only be allowed to download 100TB for no cost. After the free egress allowance, there is typically a charge for any additional downloads. However, some CSPs may charge for any type of download. You absolutely must take a look at storage vendors' egress fees and consider them as part of overall cost. For businesses using the cloud for backup and archive, a vendor that provides a free egress allowance will often meet the organization's needs.

Retrieval fees

Here's where hyperscalers overcomplicate things. In addition to the download fees, services like AWS Glacier also charge retrieval fees. These have three tiers—Expedited, Standard, and Bulk—and costs vary by tier. Glacier's low cold storage pricing can be very attractive, but you must carefully consider which service and tier to store your data. Using a low-cost service like Deep Archive can look good on paper, but retrieval times run 12–48 hours, and that's simply not a delay you can stomach when facing data disaster. [Here's a good look](#) at the kind of math you'll need to do when trying to predict your Glacier storage costs.

Download cost strategy

Don't dismiss download costs when comparing cloud storage providers. Even for backup and archive use cases, there are scenarios where you will regularly incur download costs—even outside of data disasters. Check for both egress fees and potential retrieval fees. Also be sure to factor in non-financial criteria, like how long it may take you to download your data when you need it—especially if considering cold storage.



Egress costs in action

Bethel School District had been using Amazon S3, but high costs were straining their budget—so much so that they had to shorten needed retention periods. Read more about why they chose Backblaze instead of S3 or Glacier to achieve savings of 75% and a predictable invoice.

[Read the Customer Story](#)

Transaction Costs



For backup and archive use cases, transaction costs are less relevant than storage and download costs, which will make up the bulk of your cloud storage bill. Developers and programmers are more likely to be working with transactions on a regular basis, but we've outlined some important key points about these fees below.

API calls


There are Class A, B, and C transactions. A few cloud storage providers will provide you with a plentiful free amount of transactions per day. This is usually enough for backup and archive use cases. Unsurprisingly, hyperscalers will charge for transactions, including PUT calls that are simply adding objects to a bucket. Be aware of the potential added costs of transaction fees because common backup software tools can issue calls as part of their workflows, which can incur an added expense if you're not allotted a certain amount of free calls or if you go over your free amount.

Object Lock fees

[Object Lock](#) is an essential part of a ransomware prevention strategy—it prevents your data from being encrypted, modified, or deleted. Something to note regarding Object Lock is that some storage providers charge for PUT calls, so if you ever need to renew an object lock on a file (for instance, if your lock expires every 90 days), you may get charged for that call. Your object locks can renew fairly often based on the immutability settings in your software so be sure to ask your CSP if they charge for this.

Miscellaneous additional costs

Cloud storage via hyperscalers has become so complicated that a new role has been created simply to administer and manage the storage. These "Cloud Ops" teams are tasked with overseeing storage, managing costs, and more. If you have a huge footprint that absolutely must be stored in AWS, Azure, or Google Cloud Platform, you may even have this type of role on your team, which is, of course, an additional cost. You'll also want to check how the CSP handles support. Is it included for free? Or is that another nickel and dime that will add to your bill? When comparing vendors, a good question to ask may simply be "What all might I incur a charge for?"



Conclusion: Demystifying Cloud Storage Costs

It's easy to understand why some may shy away from using cloud storage. Keeping your data on-premise can seem simpler, cleaner, and easier to account for. But the right cloud storage provider should make your data storage easier—not more complicated. Businesses need the public cloud to scale responsibly, without oversize investment in storage hardware or expensive DR facilities.

