



WHITE PAPER

How Cloud Storage Retention Minimums Impact Your Budget


Contents

- 03 About Backblaze
- 04 Introduction
- 05 Understanding Minimum Data Retention Periods
- 06 How Minimum Data Retention Periods Are Applied
 - What are delete fees?
- 07 The Financial Impact of Minimum Data Retention Periods
- 08 Conclusion

About Backblaze

Backblaze is the cloud storage innovator delivering a modern alternative to traditional cloud providers. We offer high performance, secure cloud object storage that customers use to develop applications, manage media, secure backups, build AI workflows, protect from ransomware, and more.

Backblaze B2 Cloud Storage is one-fifth the cost of AWS S3 with no minimum retention fees and 3x free egress. With transparent pricing and enterprise grade security, Backblaze helps businesses break free from the walled gardens that traditional providers lock customers into, enabling them to use their data in open cloud workflows with the providers they prefer at a fraction of the cost.

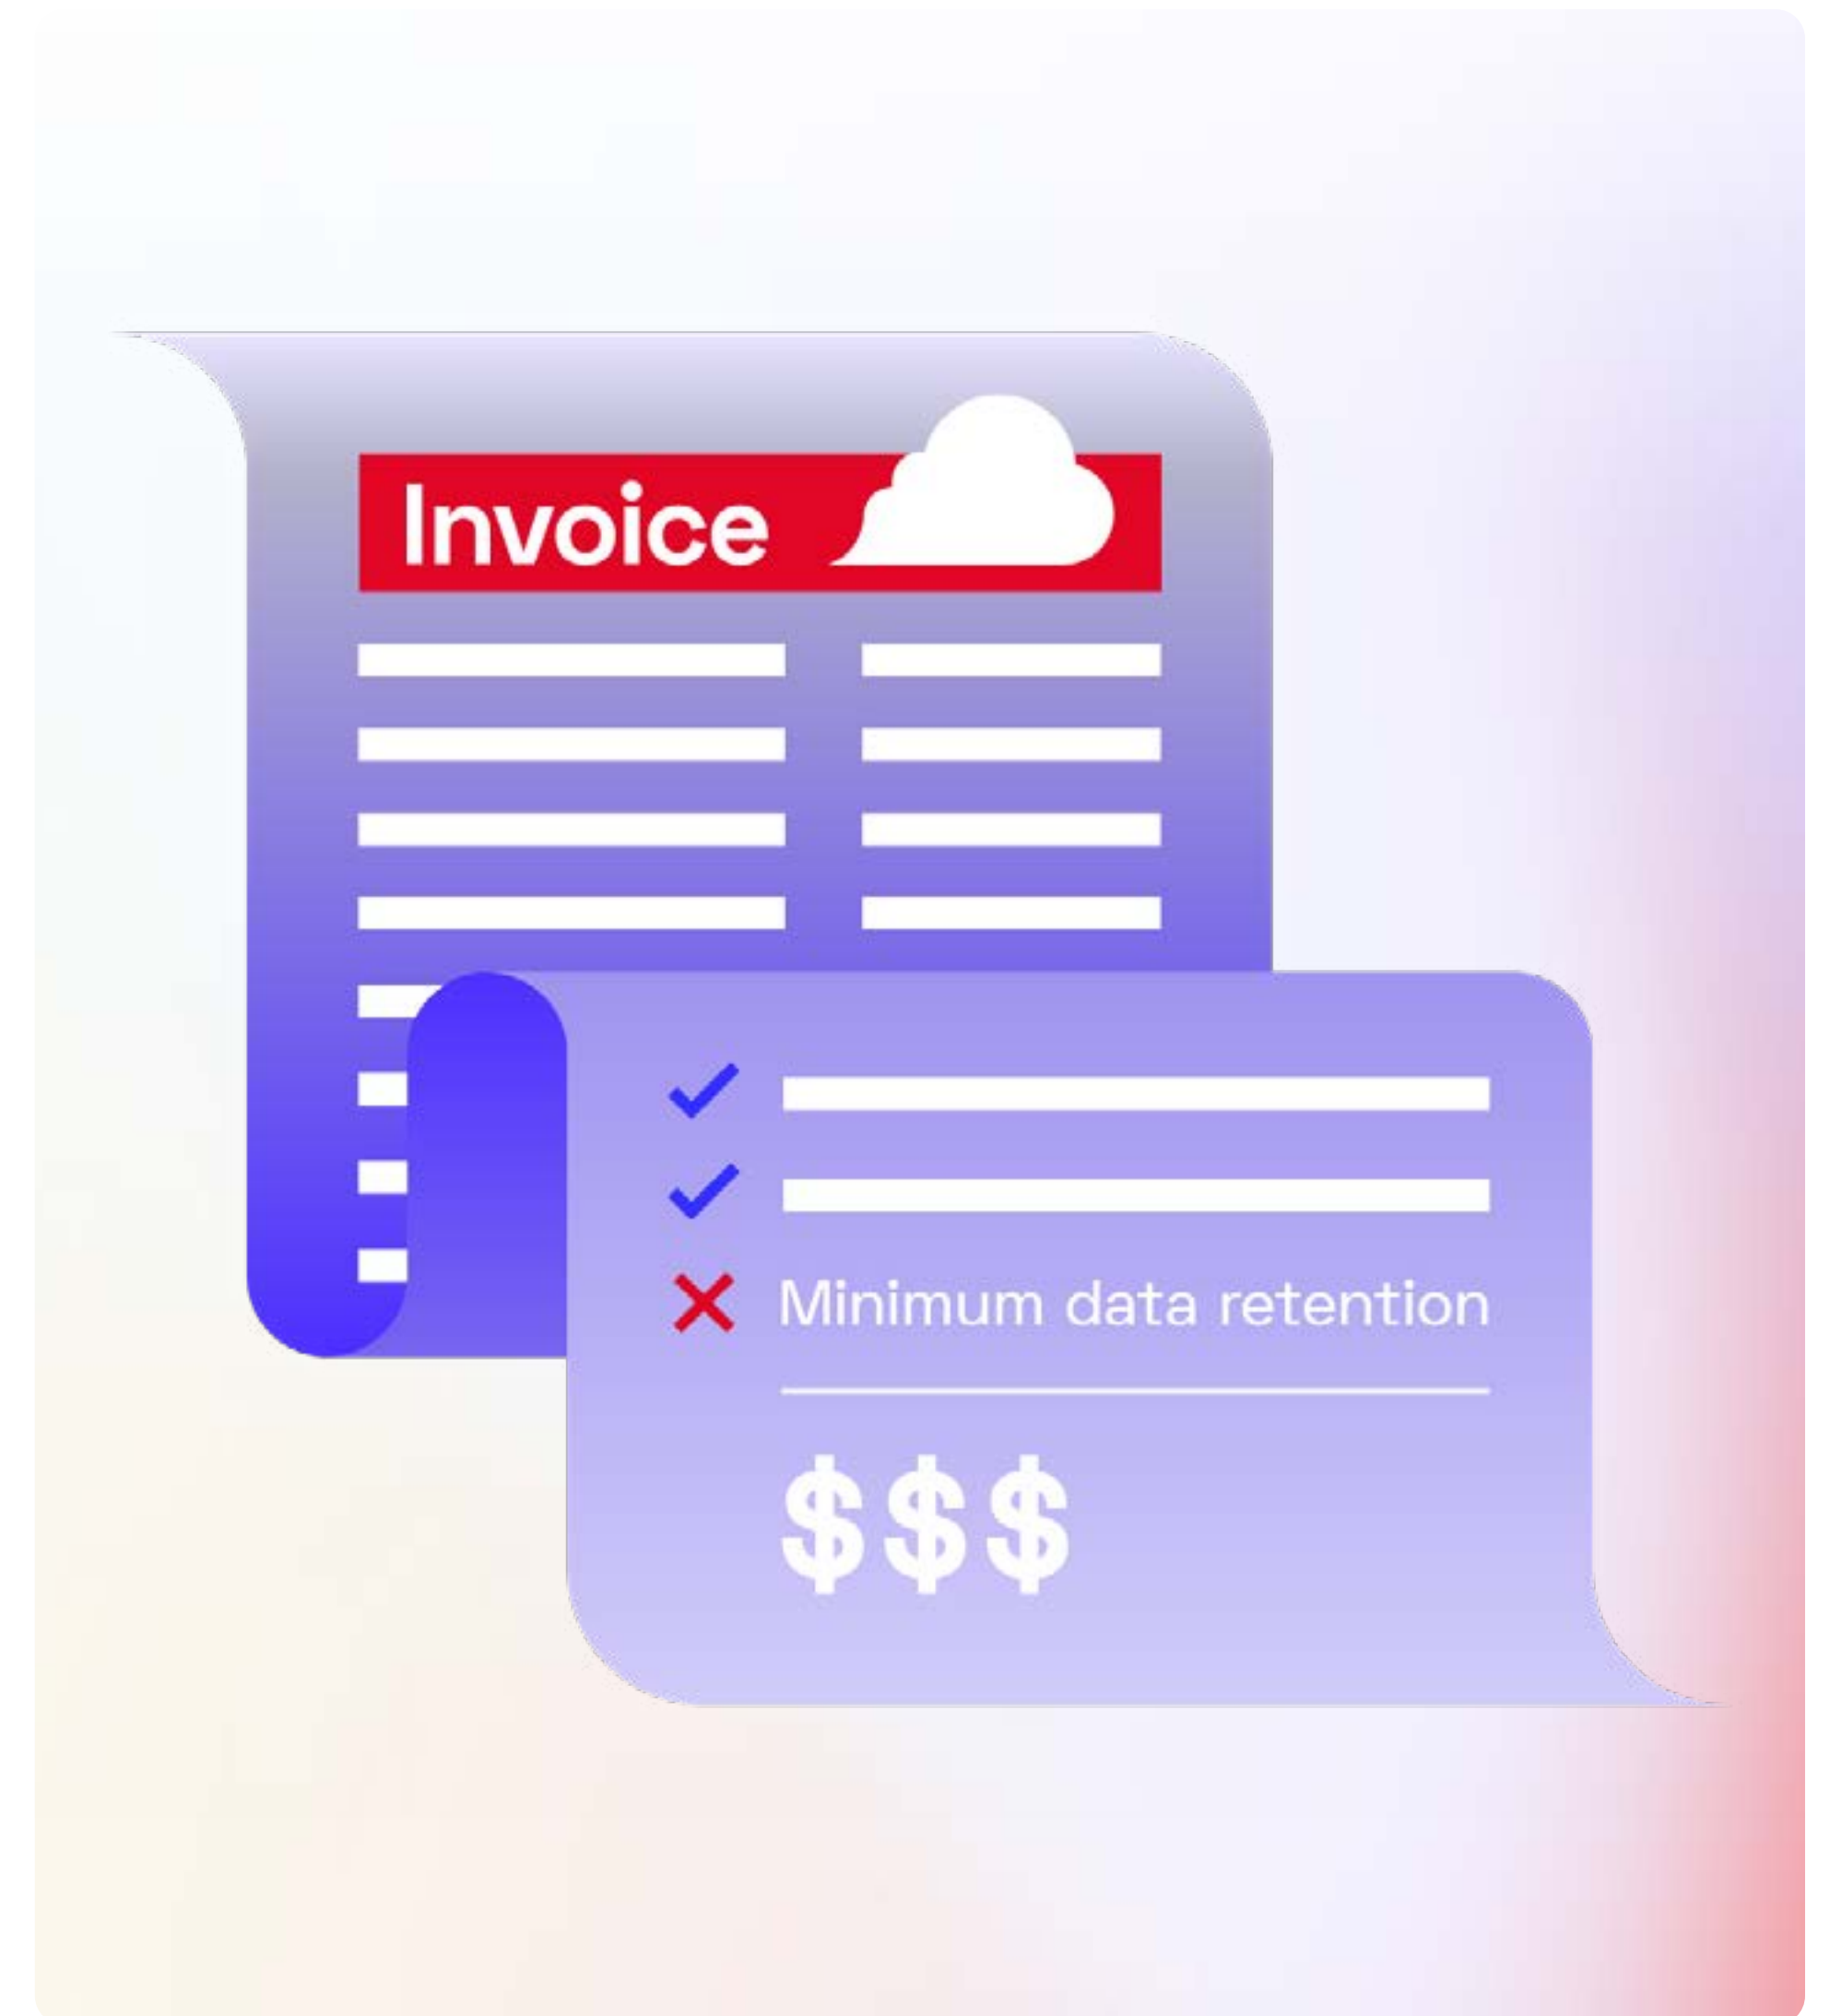


Introduction

Organizations generate and consume massive amounts of data daily. Because of this, cloud storage has become a vital component of modern business infrastructure, offering unprecedented scalability, accessibility, and disaster recovery capabilities.

However, as businesses scale their cloud storage usage, many encounter an often overlooked challenge to their budget: **minimum data retention periods**. These policies, implemented by some cloud providers like AWS, Azure, Google Cloud, and Wasabi, can lead to unexpected cost increases and complicated data management strategies. Organizations are forced to develop strategies to manage their data retention policies, balancing cost savings with operational requirements.

This white paper will look into the complexities of cloud storage retention minimums and how they directly impact storage budgets and data management policies.



Understanding Minimum Data Retention Periods

Retention minimums are a key aspect of many cloud providers’ pricing models. These are policies that specify the minimum amount of time that data must be stored in storage before it can be deleted, overwritten, or moved without incurring additional charges. This fee is typically prorated, representing the remaining days in the retention period that the data was meant to occupy in a storage class.

Minimum retention policies ensure that customers cannot frequently move data between storage tiers to exploit lower-cost storage classes for short-term storage. For cloud providers that have a single class of storage, these policies allow providers to stabilize their resource usage and maintain predictable pricing structures.

It’s important to note that minimum retention periods can vary significantly between providers, and even between different storage tiers offered by the same provider. For example, AWS S3 Standard has no minimum retention period but S3 Standard-IA has a 30 day minimum, Glacier has a 90 day minimum, and Deep Archive has a 180 day minimum.

Despite their significance, information about these retention periods is often buried in the fine print of service agreements or technical documentation. Many organizations find these requirements challenging to locate and understand, leading to unexpected costs and complications in their cloud storage strategies. This lack of transparency can make it difficult for businesses to accurately forecast their storage expenses and compare offerings between different providers.

What are delete fees?

Delete fees are a direct consequence of deleting or moving files before the retention minimum is met. Cloud providers charge these fees to ensure that the infrastructure allocated for the data is compensated for the resources it would have otherwise used during the retention period. Note that the terms “delete fees,” “minimum storage duration,” and “minimum retention fees” all refer to a similar policy.

Better understand the three dimensions of cloud costs

Learn more about what to watch out for with regards to cloud storage pricing in this white paper, “The Cost of Cloud Storage.”

Get White Paper ↗



How Minimum Data Retention Periods Are Applied

Early deletion fees can be triggered by various actions, not just the obvious deletion of files. Some examples are below.

- **Moving data from a higher-cost tier to a lower-cost tier before the minimum retention period has been met.**

This scenario often catches organizations off guard when they attempt to optimize costs by transferring infrequently accessed data to a cheaper storage class.

- **Overwriting existing files.** When a file is overwritten, the cloud provider typically treats this as a delete operation followed by a new write operation. If the original file hasn't met its minimum retention period, the organization may be charged for the remaining time, even though they're still using the same amount of storage space.
- **Implementation of automated lifecycle policies.** Many organizations set up rules to automatically move or delete data based on its age or access patterns. However, if these policies don't account for minimum retention periods, they can inadvertently trigger early delete fees on a large scale.

- **Renaming files or folders.** Even seemingly benign actions like renaming files or folders can sometimes be interpreted as delete-and-rewrite operations by certain cloud storage systems, potentially triggering these fees. Note how AWS S3's [documentation](#) explains that renaming adds a delete marker to the original object.

Additionally, in multi-user or multi-team environments, lack of communication about retention policies can lead to unexpected charges. One team might delete or move data without realizing the financial implications for the entire organization.

Understanding the various triggers for early delete fees is crucial for effective cloud storage management. However, knowing what actions can lead to these fees is only part of the equation. To truly grasp the impact on your organization's budget, it's essential to examine the financial implications of these fees in detail.



The costs were getting out of control. We thought Wasabi's pricing was face value and it wasn't.

Cade Webb

Deputy CIO, Cal Poly Humboldt

300 TB+	300	0
stored	servers	surprise bills

The Financial Impact of Minimum Data Retention Periods

Understanding the nuances of minimum data retention periods and their associated costs is crucial for developing an effective and economically sound cloud storage strategy. It enables organizations to make more informed decisions, avoid unexpected expenses, and better align their storage choices with their specific data management needs and budget constraints.

Minimum data retention periods, particularly in cold storage tiers, can have significant impacts on IT budgets. What may have seemed like a cost-saving storage tier can actually increase expenses when operations require frequent deletions or movements of data before the minimum retention period is over. But even in hot storage, these policies can unexpectedly inflate overall costs.

To illustrate the real-world impact of retention minimums, let's examine a few common scenarios:

Backup strategy

Consider a business implementing a 30 day backup strategy for their critical infrastructure. The business hopes to reduce storage costs so they opt for Wasabi object storage. They plan to keep a month's worth of backups in the cloud and will then replace them with the newer backups.

Wasabi's minimum retention policy is 90 days for its Pay as You Go storage (30 days for its Reserved Capacity Storage).

Let's say the business stores an initial 50TB of backups in Wasabi on Day 1. On Day 31, the older backup is deleted and replaced with the newer backup. The business incurs costs for 30 days of Timed Active Storage (50TB) and 60 days of Timed Deleted Storage (50TB). These charges are incurred every time the backup is replaced.

With Wasabi's Pay as You Go storage, their monthly bill will look like this:

$$50\text{TB} \times \$6.99/\text{TB}/\text{month} \times 3^{**} = \$1048.50$$

**Because the 90 day minimum retention policy equals 3 months' time.

Compare this to Backblaze B2 Cloud Storage, which has no minimum retention policy and costs \$6/TB/month:

$$50\text{TB} \times \$6/\text{TB}/\text{month} = \$300$$

The minimum retention policy effectively **triples** the anticipated storage expenses. When scaled across multiple backup sets or extended periods, the impact on the IT budget can be substantial. This pricing mechanism forces organizations to think carefully about how long they plan to store data, introducing a layer of unnecessary complexity to their data management strategy.

Conclusion

Cloud storage has revolutionized how businesses manage their data, offering unparalleled scalability and accessibility. However, as we've explored, the hidden complexities of minimum data retention periods can significantly impact your bottom line. These policies, often buried in the fine print, can lead to unexpected costs and complicate data management strategies across various industries.

Unlike AWS, Google Cloud, Azure, and Wasabi, Backblaze B2 Cloud Storage does not charge minimum retention fees so you can store and manage your data according to your business needs—without worrying about hidden fees. With a single class of always-hot storage and priced at one-fifth the cost of AWS S3, Backblaze provides customers the freedom to architect their ideal data storage strategy so they can focus on leveraging data to drive business value, rather than navigating complex storage policies.

[Backblaze.com](https://backblaze.com) ↗

Stored object

GB

 **Backblaze**

Others

30 day minimum