



Lyve Cloud Object Storage - Understanding Billing



Click here to access an up-to-date online version of this document. You will also find the most recent content as well as expandable illustrations, easier navigation, and search capability.

Contents

1	Frequently Asked Questions	3
	What do I pay for?	3
	How is the amount calculated?	3
	What if I have a pre-paid and/or committed minimum storage level?	3
	Does using Infrequent Access tier change the way billing works?	4
	Examples of usage	4
	• Simple backup example	4
	• Replicated data example	5

Frequently Asked Questions

It is typical in the cloud services world for vendors to provide complex and difficult to understand billing mechanisms for customers. Seagate Lyve Cloud Object Storage offers simple, no nonsense solutions to complex problems. This includes your bills. We provide usage policies and bills which are easy to understand, enabling you to quickly predict your billable storage usage.

What do I pay for?

Customers pay for the data they store in the Lyve system. Yes, it really is that simple!

There are no additional charges other than the storage of data. We do not have API, processing, replication, or egress fees, or any other charges. This means it's easy for you to understand the bill, as it is based purely on the amount of stored data.



If you use our replication functionality and have multiple copies of your data across two or more locations, you pay for each copy of the data. However, there are no processing or replication fees associated with the creation or maintenance of those copies.

How is the amount calculated?

The total storage for an account in the system is calculated from the average storage over the month. Each day the system tracks and reports the usage in your account. At the end of the month, the daily values are used to generate an average total amount of data stored in that month.

The standard class of service has no minimum time for objects to be stored (and thus no ghost data), so the system uses only a total of the data stored. Objects can be created and deleted as desired by customers with no effect on the total usage. Only the overall total is used to generate your bill.

What if I have a pre-paid and/or committed minimum storage level?

Seagate customers often commit to a minimum storage capacity for their account. This removes the concern of fluctuating monthly bills and provides predictability. It does not affect storage usage calculations. The monthly usage level is just compared to the minimum commitment level, and if the usage is below the committed value, then the customer is billed for the minimum commitment. The actual usage

is shown on the bill.

If the amount of the committed storage is exceeded, then a bill is generated for the overage difference. The billable usage amounts will be shown on the bill. If the usage is under the minimum commitment value, then no additional bill will be generated.

Does using Infrequent Access tier change the way billing works?

For customers who sign up for Lyve Cloud Infrequent Access class of storage, the objects saved at that tier have to be stored for at least 180 days. That means that even if you delete objects before the retention period ends, you will still be billed for it (up to 180 days). Objects deleted prior to the end of the retention period are called 'ghost data' in the billing and reporting, and appears in the billing statement as such. Ghost data is counted in the usage amount for Infrequent Access data. This is the only change to the billing calculations for Infrequent Access.

Examples of usage

Below are two examples to show usage calculations of usage and how Lyve Cloud Object Storage both simplifies and reduces your storage bills. Note—The examples assume a 30-day month for ease of calculations and understanding.

Simple backup example

A classic backup usage case has the backup application store data in the cloud and delete data after it is no longer needed. In this example, a customer has various backups which write a total of 1 TB of data per day to the cloud. The application keeps the data for 30 days in the cloud before deleting it.

Assuming the application was set up by a new user on the first day of the month, it will start with no data in the cloud and finish the month with 30 TB of data. In the first month, usage will be the average of this increasing capacity, as on each day the total increased by 1 TB. At the end of the 30 days, the average data stored for the month will be 15.5 TB.

After that month, the data storage will be in a steady state. The application will be adding and deleting 1 TB of data each day, but the average utilization remains the same. The usage for the month will be 30 TB, as that is the average for the month. The number of times the data is accessed or replaced is irrelevant for the usage and billing calculation.



This is a huge saving compared to competitive services, which would charge a minimum retention time for the data (ghost data) as well as potential egress and API fees to list the content of stored objects or to access the objects.

Replicated data example

In this example, a customer has 200 TB of data which is automatically replicated across two sites. The data rarely changes in total, although the customer accesses the data frequently. Their normal usage bill would be for 400 TB, as the customer stores 200 TB in two sites.

However, on the last day of the month, the customer uploaded another 50 TB of data into the account. This adds 100 TB of total data in the account, as the data is replicated to both sites. The average bill for the month is calculated as 29 days where the total capacity was 400 TB, and just one day where the total capacity was 500 TB. The monthly billing statement for the storage utilization would be for 403.3 TB of usage, which is the total averaged across the month.

Assuming the customer does not change the total amount of data stored, in the next month the usage will be 500 TB, because the customer now has 250 TB of data replicated across two sites.