

Our Best S3 Cost Optimization Tips for Startups in 2023

Meta Title **The best ways for startups to optimize their Amazon S3 Costs**

Meta Desc **These cost optimization tips for startups will help you make the most of your Amazon S3 installation without breaking the bank.**

Amazon S3, or Amazon Simple Storage Service, is a core product offered by Amazon Web Services (AWS) and is one of the original AWS services, introduced in 2006. S3 provides object storage for your entire AWS implementation. Amazon S3 enables you to actually store and organize the data necessary to the function of your startup. This also makes cost optimization for Amazon S3 one of the most straightforward ways that your startup can [reduce overhead](#).

By properly configuring Amazon S3 for your use case, you can not only improve the reliability of your cloud based app, but also realize some significant cost savings. To help you do this, we have compiled a guide that will walk you through the basics of Amazon S3, and provide you with some simple ways to optimize your deployment and [reduce your AWS costs](#).

What is Amazon S3

Amazon S3 is a cloud-based, serverless and practically infinite object storage service offered by AWS. It enables you to store all the data your startup needs to function, without having to build a physical server infrastructure, saving time, and increasing the scalability of your business.

Amazon S3 data is stored in partitions called buckets. Buckets are the fundamental containers used to store objects, which are the basic entities you work with in Amazon S3. An object refers to data, and can be anything from a small text file to a large video file or a backup.

While Amazon S3 is not a requirement to use AWS, it is one of the most popular AWS services. This is not just because Amazon S3 is a scalable and durable storage solution, rated with 11 nines durability, which means that even with a billion objects you could go a hundred years without losing any single one! In addition, because it integrates so well with a variety of other AWS services, such as [Amazon EC2](#), this makes it easier for companies to build their application, without needing to integrate third party storage services into their cloud implementation.

As with many AWS products, there are a number of ways they can be customized to different workloads, and the easiest way to do this is by using S3 storage tiers, also known as storage classes.

What Are S3 Storage Tiers?

The first thing we need to look at are S3 Storage tiers, also called Object Storage tiers. These are designed to help you optimize your storage costs based on your own specific use-cases. These storage tiers offer different combinations of durability, availability, performance, and cost.

There are seven [storage tiers](#), that can be broken down into two groups:

Main S3 Tiers

Main S3 tiers are designed for frequently accessed data, and offer high performance and low latency at a slightly higher price point than archival storage tiers. These tiers include:

1. **Amazon S3 Standard:** This is the default storage tier and is designed for frequently accessed data. It offers high durability, availability, and performance. This storage tier is ideal for applications that require low latency and high throughput.
2. **Amazon S3 Intelligent-Tiering:** This storage tier is designed for data with unknown or changing access patterns. It uses machine learning to automatically move data between two access tiers, one for frequent access and another for infrequent access. This storage tier can help reduce costs while maintaining high performance. There are many benefits to this tier, and we'll go into more detail about it [here](#).
3. **Amazon S3 Standard-Infrequent Access (S3 Standard-IA):** This storage tier is designed for data that is accessed less frequently but requires rapid access when needed. It offers the same high durability and availability as S3 Standard, but with lower cost per GB stored and a retrieval fee when accessing data.
4. **Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA):** This storage tier is similar to S3 Standard-IA but stores data in a single availability zone instead of multiple zones. This makes it less durable, but also less expensive than S3 Standard-IA.
5. **S3 Glacier Instant Retrieval:** This storage tier delivers low cost storage for objects that are rarely accessed, but still need to be retrieved rapidly when they are required, and the retrieval time is measured in milliseconds.

Archival S3 Tiers

Archival S3 tiers are designed for infrequently accessed data, and are often far cheaper than main storage tiers, but there is more latency before accessing the data stored within.

6. **Amazon S3 Glacier Flexible Retrieval:** This storage tier is designed for data archiving and long-term storage. It offers lower storage cost but with longer retrieval times, typically ranging from minutes to hours.
7. **Amazon S3 Glacier Deep Archive:** This storage tier is designed for long-term data archiving and offers the lowest storage cost of all the S3 storage tiers, but with the longest retrieval times, ranging from hours to one day.

How Should I Select my Amazon S3 tier?

The key to selecting the right Amazon S3 storage tier is understanding your specific use-case. Just putting all your data into S3 Standard will be prohibitively expensive for most startups, so the key is to determine how much data really **needs** rapid access, and which can be placed in lower cost archival tiers.

This can be done manually, but there are a number of tools that you can use to help your startup optimize its Amazon S3 costs:

1. Take Advantage of Amazon S3 Lifecycle Policies

[Amazon S3 Lifecycle policies](#) are rules that you can set up to automatically manage the lifecycle of your S3 objects. These policies can help you optimize your S3 storage costs, improve data protection, and meet compliance requirements by defining when objects should be transitioned to different storage tiers or deleted.

You can use this to specify when to transfer objects to different storage tiers, for example, by putting data that is infrequently accessed into archival S3 tiers. You can also use it to automatically determine which objects should be expired, or deleted, without requiring user intervention.

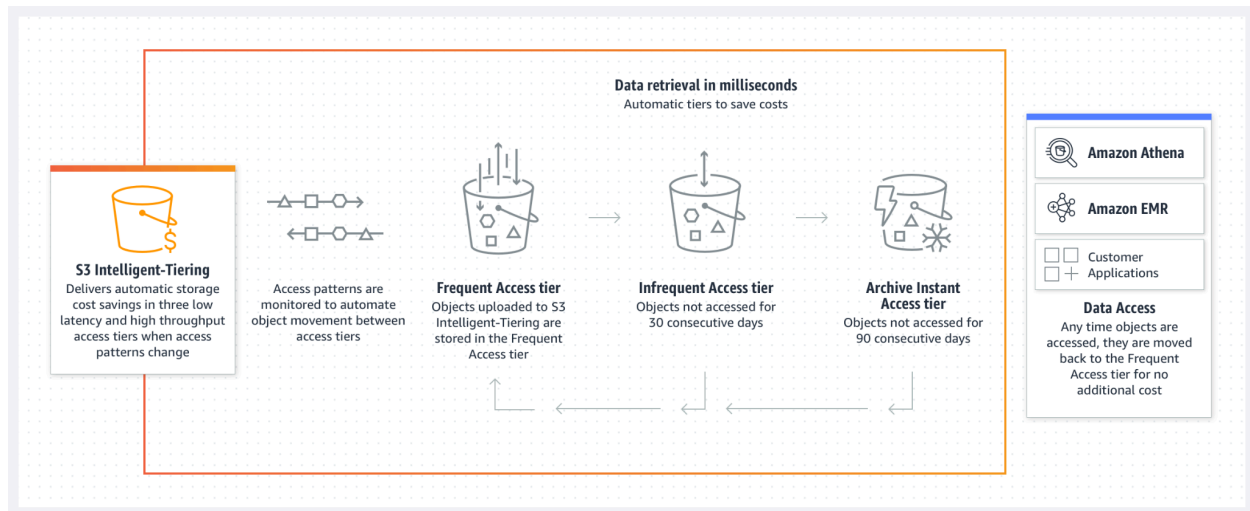
This has two key advantages. Firstly, you will be able to automatically reduce costs as data becomes less relevant to your application or startup. Secondly, you can eliminate the risk of human error, leaving unimportant data in a more expensive tier.

2. Leverage S3 Intelligent-Tiering

[S3 Intelligent-Tiering](#) is a storage tier in Amazon S3 that automatically moves objects between three tiers:

- S3 Frequent Access (similar to S3 Standard) is the default tier where all new objects will be stored.
- After 30 days of no access, objects will be moved to Infrequent Access (similar to S3 Standard Infrequent Access) tier.
- Then after 60 more days of no access, objects will be moved to an even cheaper Archive Instant Access (similar to S3 Glacier Instant Retrieval) tier.

Based on changing access patterns and needs, it essentially allows you to intelligently move data, without needing to make a specific Lifecycle Policy.



A breakdown explaining how S3 Intelligent-Tiering Works

This is very useful for data that you **know** will need to be accessed, and can't risk falling into an archival tier due to a Lifecycle Policy. It enables you to remove any need to manually manage objects, helping you to meet one of your responsibilities under the [Amazon Shared Responsibility Model](#), and reduces the risks of human error.

This tier doesn't introduce *object access* costs, and instead charges a monitoring and automation fee of \$0.0025 per 1,000 objects. The actual storage cost will depend upon the tier that it assigns the data to. While it only defaults to only main storage tiers, it is possible to include archival tiers in S3 Intelligent Tiering too.

It's hard to overstate how much money this could save your startup. Many AWS users simply use the default S3 Storage tier, and end overpaying by as much as 70%! S3 Intelligent Tiering can help you avoid this.

3. Take Advantage of the AWS Calculator

The [AWS Pricing Calculator](#) is a web-based tool provided by Amazon Web Services (AWS) that allows you to estimate the cost of using S3 (or any other AWS service). It helps you to calculate the cost of running your applications on AWS by providing an estimation of the monthly bill you can expect to receive.

The AWS Calculator provides a user-friendly interface that lets you input your storage requirements, and the calculator would then provide you with an estimated monthly bill, allowing you to accurately determine what savings you would make by using intelligent tiering or using lifecycle policies and moving data to cheaper tiers.

This enables startups to get an understanding of their storage costs, and play with different pricing methodologies to identify the most appropriate solution for their specific use-case.

4. Use S3 Storage Lens to Identify Potential Pain Points

Amazon [S3 Storage Lens](#) is a cloud-based analytics solution that provides detailed insights and actionable recommendations for optimizing your Amazon S3 storage usage and costs. S3 Storage Lens analyzes your storage usage across multiple accounts and regions to provide a unified view of your S3 environment.

It provides deep data visualization tools and customization options that will help you understand your data in a way that makes sense for your startup. For example, you could set it up to identify the largest or least used objects in your S3 Storage buckets. You can then either delete those objects, or place them in a glacier archive if you think that it may be needed later.

It is even possible to automate alerts from S3 Storage Lens using the Amazon Simple Notification Service (SNS). This will help your team stay on top of large, underutilized objects before they add to your running costs.

5. Use VPC Endpoints To Avoid Expensive External Data Traffic

A common reason startups receive unexpectedly large AWS Data Transfer invoices is because of improperly configured VPC Endpoints. For example, your Elastic Compute Cloud (EC2) service may be pulling large amounts of data from S3. If there is no network route from EC2 to S3 through a VPC Endpoint, EC2 will need to go out of your Virtual Private Cloud (VPC) to the internet. It would then need to resolve the bucket name, get the data, and bring traffic back to the VPC.

You can solve this relatively easy by following these steps:

1. **Create a VPC:** First, create a VPC and subnet in the AWS Management Console. Note down the VPC ID and subnet ID for later use. Optional if you already have a VPC.
2. **Create an S3 bucket:** Create an S3 bucket in the AWS Management Console and upload any files that you want to access from your EC2 instance. Optional if you already have a bucket.
3. **Create a VPC endpoint:** Navigate to the VPC service in the AWS Management Console and create a new VPC endpoint for S3. Specify the VPC ID and subnet ID that you created in step 1. Choose the S3 bucket that you want to access, and configure the route tables to allow traffic between the VPC subnet and S3.
4. **Attach the VPC endpoint to the EC2 instance:** Go to the EC2 service in the AWS Management Console and navigate to the network interfaces section. Create a new network interface and attach it to the EC2 instance that needs to access the S3 bucket. Specify the VPC endpoint that you created in step 3 as the source for traffic.

5. **Configure security groups:** Configure the security groups for the EC2 instance and the S3 bucket to allow traffic between them. Make sure that the security groups allow traffic on the correct ports and protocols for your application.
6. **Test connectivity:** Test connectivity between the EC2 instance and the S3 bucket by attempting to access the files in the bucket from the EC2 instance. If everything is set up correctly, you should be able to access the S3 bucket securely and privately through the VPC endpoint.

6. Work with an AWS Partner Like Cloudvisor!

Another easy way to save money is to work with an expert like Cloudvisor. You will immediately be able to save up to 3% on your AWS Bill, or even donate to charity, by taking advantage of Cloudvisor's AWS [reselling service](#), but there is so much more that we can do to help.

Our team offers a variety of services, [tailored specifically to startups](#), designed to help startups like yours you make the most of AWS. Whether you're looking for a Well Architected Framework review, cost optimization tips, or an AWS migration, our team is standing by to help.

If you want to know more, [reach out to us today](#), and we'll help your startup make the most of AWS Cloud Services.