<WEBAPPER>
cloud application engineers

webapper.com

GET THE MOST OUT OF YOUR AMAZON S3 BUCKETS

GET THE MOST OUT OF YOUR AMAZON S3 BUCKETS

Amazon Web Services (AWS) provides a reliable, scalable, low-cost cloud infrastructure platform that powers almost one-third of the internet.

One of the cornerstones of AWS is Amazon Simple Storage Service. Amazon S3 is an object storage service that offers scalability, reliability, high performance, and security. S3 buckets house websites, applications, backups, archives, data lakes, and analytics solutions. You can organize data and configure access controls to meet specific organizational requirements.

Amazon S3 is designed for 99.999999999% (11 9's) of durability to store important data. Data can be redundantly stored across facilities and devices. You can store & retrieve as much data as you want using the AWS pay-as-you-go model. Amazon S3 supports data transfer over secured channels and automatic protection of uploaded data. You can choose AWS regions to optimize for latency, minimize costs, or meet regulatory requirements.

Amazon S3 is deeply integrated with other Amazon Web Services services to simplify building solutions using AWS services like Amazon CloudFront, Amazon CloudWatch, Amazon RDS, Amazon DynamoDB, and Amazon VPC. S3 also works with Amazon Lambda, so you can log activities, define alerts, and automate workflows without managing additional infrastructure.

AMAZON S3 STORAGE BASICS

Organizations of any size and in any industry can use Amazon S3 buckets to facilitate public cloud storage. S3 stores objects, which consist of data and descriptive metadata. S3 environments are flat structures — you create a bucket, and then the bucket stores objects (e.g., files, documents, photos, videos) in the cloud. You pay for storing objects in your S3 buckets. Rates depend on your S3 bucket size, how long you stored the objects during the month, and the storage class. There are no retrieval charges.

GETTING STARTED WITH S3

You can create up to 100 buckets in an AWS cloud account. For each bucket, you control access (create, delete, and list objects), access logs, and the geographical region to use for storage. To set up a bucket, you sign in to your AWS account, create your S3 bucket, and select the options such as size & region in the setup. To start using a bucket, you simply add an object — such as a text file, photo, or video — to it.

MANAGING S3

From the console, it's easy to deploy and configure S3 storage. After you create buckets, you can upload objects to Amazon S3. You can also use versioning, storage classes, object locking, batch operations, replication, and tags.

AMAZON S3 SECURITY

By default, users only have access to the S3 resources they create, but you can grant access to other users. The most common ways are to use AWS Identity and Access Management (IAM) to create users and manage their respective access or to use Access Control Lists (ACLs) to enable authorized users to access individual objects.

MONITORING S3

AWS provides tools like Trusted Advisor, CloudWatch, and CloudTrail to monitor Amazon S3. You can monitor buckets manually or configure the tools to do the monitoring for you.

AMAZON S3 STORAGE CLASSES

You can choose from different storage classes to suit your particular performance, resiliency, and cost requirements.

- **Amazon S3 Standard**
  Best for frequent data access where the latency should be low.
- **S3 Standard-Infrequent Access (S3 Standard-IA)**
  When your data is accessed infrequently and stored in a single region.
- **S3 Glacier Flexible Retrieval**
  When your data needs to be archived, and high performance is not required.
- **S3 Glacier Instant Retrieval**
  When your archived data may need immediate access,
- **S3 Glacier Deep Archive**
  For long-term archive and digital preservation with retrieval in hours at the lowest cost storage in the cloud.
- **S3 Intelligent-Tiering**
  Provides automatic cost savings for data with unknown or changing access patterns.

A key to Amazon S3 success is setting up smart S3 bucket policies that are beneficial to intended users but protect data from prying eyes.

AMAZON S3 SECURITY MATTERS

AWS S3 buckets have appeared in headlines many times for not good reasons. In February 2021, Premier Diagnostics, a Utah COVID-19 testing service, leaked patient data records through two publicly available Amazon S3 buckets that lacked any form of password protection or authentication. Over 200,000 images of ID scans were exposed. In May 2021, Twilio, a large cloud communication Platform as a Service company, allowed a bad actor to gain read and write access to a misconfigured S3 bucket. The hackers modified a copy of Twilio's JavaScript SDK that they share with customers. And in August 2021, a misconfigured S3 bucket at SeniorAdvisor exposed details of over 3 million senior citizens, including individuals' names, numbers, and email addresses.

S3 BUCKET BASICS

The Shared Responsibility Model means that AWS is responsible for protecting the infrastructure that runs Amazon S3. As an S3 user, your responsibility is to manage access to your data by assigning permissions and access levels.

Amazon S3 bucket policies allow you to grant access to a bucket and the objects (files) it contains. Permissions apply to all objects in the selected bucket. They are critical in securing your S3 buckets against unauthorized access and attacks. As a bucket owner, you are the one who applies a policy to each bucket. Bucket policies are an Identity and Access Management (IAM) mechanism for controlling access to resources. You can add or update a bucket policy using the Amazon S3 console. In addition, you can use the AWS Policy Generator to define a bucket policy for your buckets.

BEST PRACTICES FOR BUCKET POLICIES TO SECURE AWS S3 BUCKETS

- **Bucket Naming**

  When you create a bucket, you choose a name and its AWS Region. Once created, you can't change the name or Region. Names must be between 3 and 63 characters long, consist only of lowercase letters, numbers, dots, and hyphens, and begin and end with a letter or number. It's best to use names that are relevant to you or your organization.

- **Private and Public Buckets**

  While some of your S3 buckets may need to be publicly accessible, most S3 buckets should have restricted access. That is, you should ensure that your S3 bucket is not public unless you explicitly need it to be. You should only grant permissions required to perform specific tasks. Least privilege access is fundamental to reducing your security risks. It's best to use an IAM role to manage temporary credentials for applications or services that access S3. Using a role eliminates distributing long-term credentials (such as a username and password or access keys). IAM roles supply temporary permissions for applications to make calls to AWS resources.

- **Encrypt Your Data**

  You should encrypt data at rest. For server-side encryption, S3 can encrypt your object before saving it and decrypt it when a user downloads the objects. From the client side, you can encrypt data before uploading data to S3. It's also best to enforce encryption of data in transit using HTTPS (TLS) to block eavesdropping or network traffic manipulation.

- **Bucket Versioning**

  If you ever need to go back in time in relation to file changes, you should enable S3 bucket versioning. It protects you from data loss from application issues or human error. With versioning, S3 keeps multiple versions of each object in the bucket. When you upload an object with the same name, S3 stores a new version of the object. If you delete an object, S3 inserts a delete marker. As a side benefit, versioning helps with NIST, PCI-DSS and GDPR compliance. Note, versioning impacts S3 usage: S3 charges are based on storage, requests and data retrievals, data transfer, and data management.

- **Lifecycle Policies**

  You can customize your data retention approach and control storage costs by using object versioning with S3 Lifecycle. With S3 Lifecycle configuration rules, you can tell Amazon S3 to transition objects to less-expensive storage classes, or archive or delete them.

- **Monitoring**

  Audit logging is an important element of your organization's data security. To detect suspicious behavior or spot security incidents, your organization should continuously monitor and audit user activities related to S3 buckets. You can enable CloudTrail data events for all your buckets or for a list of specific buckets. CloudTrail captures a subset of API calls, including calls from the S3 console and code calls to the S3 APIs.

HOW TO CONTROL AMAZON S3 COSTS

It's possible to (needlessly) run up storage costs using AWS. With metered usage including a variety of inputs, it makes sense to look for ways to control Amazon S3 costs without losing the many benefits of using these tools.

AMAZON S3 COST INPUTS

First, it's important to see what the meters track.

- Amount of data stored during the month (GB).
- Size of data retrieved and number of requests.
- How many access operations (e.g., GET, LIST, COPY, POST, SELECT) completed.
- Data transfer fees (bandwidth out from Amazon S3)

The storage class used also impacts costs. It's important to understand the S3 classes and their use cases.

- S3 Standard: for normal storage of regularly accessed data.
- S3 One Zone – Infrequent Access (IA): low-cost storage classes for data that users don't access frequently in one zone.
- S3 Standard — Infrequent Access (IA): low-cost storage classes for data that users don't access frequently.
- S3 Standard Intelligent Tiering: stores data with changing patterns.
- S3 Glacier: to archive and store long-term backup data.
- S3 Glacier Deep Archive: to archive and store (very) long-term backup data.

HOW TO CONTROL YOUR AMAZON S3 COSTS

1 Map Your Application Requirements

The first step is an assessment. What do you have? How are you using it today?

2 Analyze Your S3 Pricing Bill

Call this the "before" photo of your weight loss program. Take stock of your application requirements versus what you're paying for.

3 Organize Your Data

Using analysis from #2, you need to understand how your usage relates to your charges.

4 Analyze & Optimize

- Review the cost of individual buckets.
- Delete entire S3 buckets if no longer needed.
- Remove objects that you no longer need.
- Delete incomplete multipart uploads.
- Review your data retrieval costs.
- Review requests made to your bucket.
- Verify (and correct as needed) that you're using the right AWS region for your S3 buckets.
- Compress data before you send it to S3 (fewer + smaller files lowers transfer costs)
- For S3 versioned buckets, use the "lifecycle" feature to delete old versions

5 Continuously Rightsize & Tune

Lather. Rinse. Repeat. Re-run your analysis on a monthly or quarterly basis to ensure you're not overspending. Set up CloudWatch alerts for areas that add to cost (e.g., third party systems that upload objects). Storage cost optimization is not something to set and forget.

ADDITIONAL TOOLS TO CONTROL S3 COSTS

S3 INTELLIGENT TIERING

S3 Intelligent Tiering can transparently manage the tiering aspect. For a small monthly object monitoring & automation charge, it optimizes your storage costs by automatically moving data to the most cost-effective access tier when access patterns change.

S3 STORAGE LENS

S3 Storage Lens helps you can see aggregate usage and activity data so you can optimize storage configurations and costs. The dashboard lets you explore your metrics to see how your storage is being used.

AMAZON S3 STORAGE CLASS ANALYSIS

If you don't use S3 Intelligent Tiering, you can try Storage Class Analysis to help you know when it's wise to transition STANDARD storage to the STANDARD_IA (IA, for infrequent access) storage class.

CLOUDSEE™

SIMPLIFY YOUR S3 STORAGE

CloudSee Drive Untangles Your AWS Buckets Using a Familiar Interface For AWS administrators and end users, an S3 file browser...in your browser.

UPLOAD AND BROWSE IN A BROWSER

See Drive is a simple solution that works for system administrators and end users who need to manage files in AWS S3 buckets. With a familiar interface — like Finder or Explorer — anyone can search, upload, or view files stored in S3 buckets. No AWS console account needed.

ACCESS ANY OF YOUR S3 BUCKETS

CloudSee Drive users can search and view any S3 bucket you enable as accessible. Administrators save time and hassle of supporting the organization's basic needs.

SEARCH FOR FILES

Search and view files stored in S3 with a simple user interface. It's fast, familiar, and easy. Users can get to the digital assets they need immediately, with no technical support.

...AND USE METADATA TOO

Save time searching for files later by adding optional metadata on the way in (during upload). Search by file name, category, ordescription. Sort results by name and date.

CLOUDSEE DRIVE IS POWERFUL AND EASY TO USE

- Easy sign-up and setup.
- Nothing to install.
- 100% browser based.
- Reliably upload & download your files to and from AWS S3.
- Mobile friendly.
- Browse, view, and edit AWS S3 bucket files.
- Add descriptions (searchable too!).
- Reviewed & verified by AWS.

<WEBAPPER>
cloud application engineers

webapper.com

WEBAPPER THRIVES IN THE CLOUD

At Webapper, we have a long history of building software. Starting in the 90s, we built "web 1.0" applications. After the dotcom crash in the early 2000s, Webapper continued building web applications and embraced agile software development. In the early 2010s, we started building and hosting applications on AWS.

CloudSee reflects Webapper's years of experience — we've built tools that we needed for our own cloud journey. Over the past few years, CloudSee has evolved and will continue to do so. We're retooling our original product to be more cloud-native and extensible, and we think you'll love what we deliver.

Webapper brings decades of hosting and development experience, including working with the cloud, to this endeavor. Our team includes certified AWS engineers, folks who have brought numerous products to market, and developers who have walked more than one mile in your shoes. We don't profess to know it all, so we welcome feedback and ideas from our customers. Tell us what you think we should build next!

Webapper Cloud Application Engineers

GET A FREE CLOUD CONSULTATION

Call (970) 670-0169 or visit webapper.com today.