

Reference:

[FAQs](#)

Category:

Storage



Amazon Simple Storage Service (Amazon S3)

Complete book:

[Click Here](#)

Created by:

[Ashish Prajapati](#)



What?

- Amazon S3 is a simple key-based object storage built to store and retrieve any amount of data from anywhere.
- Data is stored as objects within resources called “buckets”, and a single object can be up to 5 terabytes in size.
- It is designed to provide 99.999999999% (11 9's) of data durability.

Why?

- You can use a simple web service interface to store and retrieve virtually any amount of data in any format.
- Highly scalable, highly available, fast, inexpensive data storage infrastructure and you only pay for what you use.
- It offers a range of storage classes to choose from based on the data access, resiliency, and cost requirements of workloads.

When?

- You want to store static content, storage backups, want to build data lakes.
- You require version controlled object storage, Multi-Factor Authentication (MFA) Delete capability and selectively grant permissions to users and groups of users.

Where?

- Amazon S3 stores data as objects within buckets. A bucket is created in a Region and requires a globally unique name.
- S3 storage classes provide multi-Availability Zone (AZ) resiliency by redundantly storing data on multiple devices and physically separated AWS Availability Zones in an AWS Region (except S3 One Zone-IA storage class).

Who?

- Upon creation, only you have access to Amazon S3 buckets that you create, and you have complete control over who has access to your data.
- You can use the Amazon S3 Management Console, the AWS SDKs, or the Amazon S3 APIs to interact with it.

How?

- You can get started by creating a bucket in a specific Region and can defining access controls and management options.
- To store an object in Amazon S3, upload the file into a bucket. An object is composed of a file and any metadata that describes that file. When you store data, you assign a unique object key that can later be used to retrieve the data.

How much?

- There are six Amazon S3 cost components to consider when storing and managing your data— storage pricing, request and data retrieval pricing, data transfer and transfer acceleration pricing, data management and analytics pricing, replication pricing, and the price to process your data with S3 Object Lambda.