Amazon Simple Storage Service (S3)

Manage and reduce storage costs with S3 Storage Classes

Amazon S3 is the largest and most performant, secure, and feature-rich object storage service. S3 is designed for 99.999999999% (11 9s) of durability, and stores data for millions of companies all around the world.

S3 offers a range of storage classes that you can choose from based on the data access, resiliency, and cost requirements of your workloads. S3 storage classes are purpose-built to provide the lowest cost storage for different access patterns, and virtually any use case. With the S3 Storage Classes, S3 Storage Lens, S3 Storage Class Analysis, and S3 Lifecycle policies, you can enable storage cost efficiencies without impacting availability or performance.

The Amazon S3 Storage Classes

Purpose-built to provide the lowest cost storage for different access patterns, and virtually any use case

S3 Intelligent-Tiering
Automatic cost savings by auto-tiering data with any access pattern

S3 Standard
General purpose storage for active, frequently accessed data

S3 Standard-Infrequent Access (S3 Standard-IA)
Low cost storage for data accessed monthly, and requires milliseconds retrieval

S3 Glacier Instant Retrieval
Low cost storage for long-lived data, with retrieval in milliseconds

S3 Glacier Flexible Retrieval
Long-term, low-cost storage for backups and archives, with retrieval options from minutes to hours

S3 Glacier Deep Archive
Lowest cost cloud storage for long-term, rarely accessed archive data, with retrieval in hours

S3 One Zone-Infrequent Access (S3 One Zone-IA)
Infrequently accessed data in a single AZ for cost savings

S3 on Outposts
Delivers object storage to on-premises AWS Outposts environments to meet local data processing and data residency needs

S3 Storage Classes support virtually every storage use case from backup and recovery, to archive and digital preservation, to data lakes, business critical applications, and analytics.

Learn more about S3 storage classes

All S3 Storage Classes are designed for **99.999999999% durability**

How is this achieved?

S3 maintains this unmatched durability by creating copies of all objects uploaded to S3 and storing them across multiple systems within a single AWS Region. 11 9s of durability protects your data for all storage classes against site-level failures, errors, and threats. S3 provides even greater resiliency with data replicated to 3+ Availability Zones (AZs) in a single Region that are separated by a minimum of 1 mile, and no more than 100 miles.

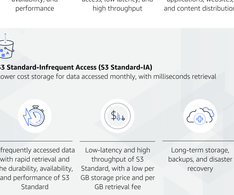

S3 Storage Classes

Storage classes purpose-built for varying access patterns at corresponding costs

S3 Intelligent-Tiering

Only cloud storage class with automatic cost savings by moving objects between cost-optimized access tiers

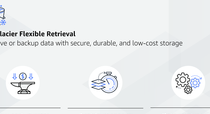

Automatically moves objects between the optimal access tiers

Optimizes storage costs based on access patterns

Data lakes and applications with changing or unknown access patterns

S3 Standard

General purpose storage for active, frequently accessed data with millisecond access

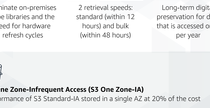

Delivers high durability, availability, and performance

Designed for frequent access, low latency, and high throughput

Data lakes, cloud-native applications, websites, and content distribution

S3 Standard-Infrequent Access (S3 Standard-IA)

Lower cost storage for data accessed monthly, with milliseconds retrieval

Infrequently accessed data with rapid retrieval and the durability, availability, and performance of S3 Standard

Low-latency and high throughput of S3 Standard, with a low per GB storage price and per GB retrieval fee

Long-term storage, backups, and disaster recovery

S3 Glacier Instant Retrieval

Lowest cost storage for long-term archive data, with milliseconds retrieval

Designed for rarely accessed, long-term data that requires immediate retrieval

Save up to 68% on storage costs compared with using the S3 Standard-Infrequent Access storage class

Long-term digital preservation of data that is accessed once per quarter

S3 Glacier Flexible Retrieval

Archive or backup data with secure, durable, and low-cost storage

Low-cost, durable archive with low retrieval fees

3 retrieval speeds: expedited (1–5 mins), standard (3–5 hours), and bulk (12 hours)

Data archive with query-in-place capabilities to learn from data-at-rest

S3 Glacier Deep Archive

Lowest-cost cloud storage for long-term archives at about $1 per TB/month

Eliminate on-premises tape libraries and the need for hardware refresh cycles

2 retrieval speeds: standard (within 12 hours) and bulk (within 48 hours)

Long-term digital preservation for data that is accessed once per year

S3 One Zone-Infrequent Access (S3 One Zone-IA)

Performance of S3 Standard-IA stored in a single AZ at 20% of the cost

Low latency, high throughput storage for infrequently accessed objects that require rapid retrieval

Saves storage costs for data that doesn't require multi-AZ redundancy

Ideal for secondary backups and workloads with easily re-creatable data

Amazon S3 on Outposts

Delivers object storage to your on-premises AWS Outposts environment to meet local data processing and data residency needs

Makes it easy to store, secure, tag, retrieve, report on, and control access to the data on your Outpost

Satisfy demanding performance needs by keeping data close to on-premises applications

Provides on-premises object storage to minimize data transfers and buffer from network variations

Data lifecycle management

If you don't know your data access patterns, or it's unpredictable, you should use S3 Intelligent-Tiering to automatically save on storage costs, without operational overhead, or impacts to performance. If you have well-known and predictable access patterns, then you can use S3 Storage Class Analysis with S3 Lifecycle to define rules to move data to lower-cost storage classes.

Learn more about Amazon S3 Storage Classes
aws.amazon.com/s3/storage-classes/

aws storage