


# Amazon Simple Storage Service (S3)

Manage and reduce storage costs with S3 Storage Classes

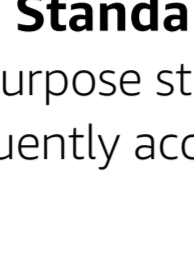


Amazon S3 is the largest and most performant, secure, and feature-rich object storage service. S3 is designed for 99.999999999% durability to protect data from site-level failures, errors, and threats and stores data across a minimum of 3 availability zones (AZs).

S3 has six storage classes purpose-built for varying access needs to help you optimize costs. With the S3 Storage Classes, S3 Storage Class Analysis, and S3 Lifecycle policies, you can enable storage cost efficiencies without impacting availability or performance.


## The Amazon S3 Storage Classes

Designed to support different use cases, data access, and retrieval requirements at different costs




### S3 Standard

General purpose storage for active, frequently accessed data



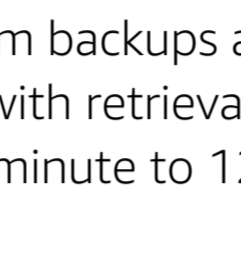
### S3 One Zone-Infrequent Access (S3 One Zone-IA)

Infrequently accessed data in a single AZ for cost savings




### S3 Intelligent-Tiering

Auto-tiering for data with changing or unknown access patterns



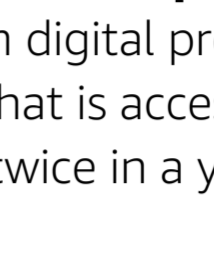
### S3 Glacier

Long-term backups and data archives with retrieval options from 1 minute to 12 hours




### S3 Standard-Infrequent Access (S3 Standard-IA)

Infrequently accessed data that requires rapid retrieval



### S3 Glacier Deep Archive

Long-term digital preservation for data that is accessed once or twice in a year



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.

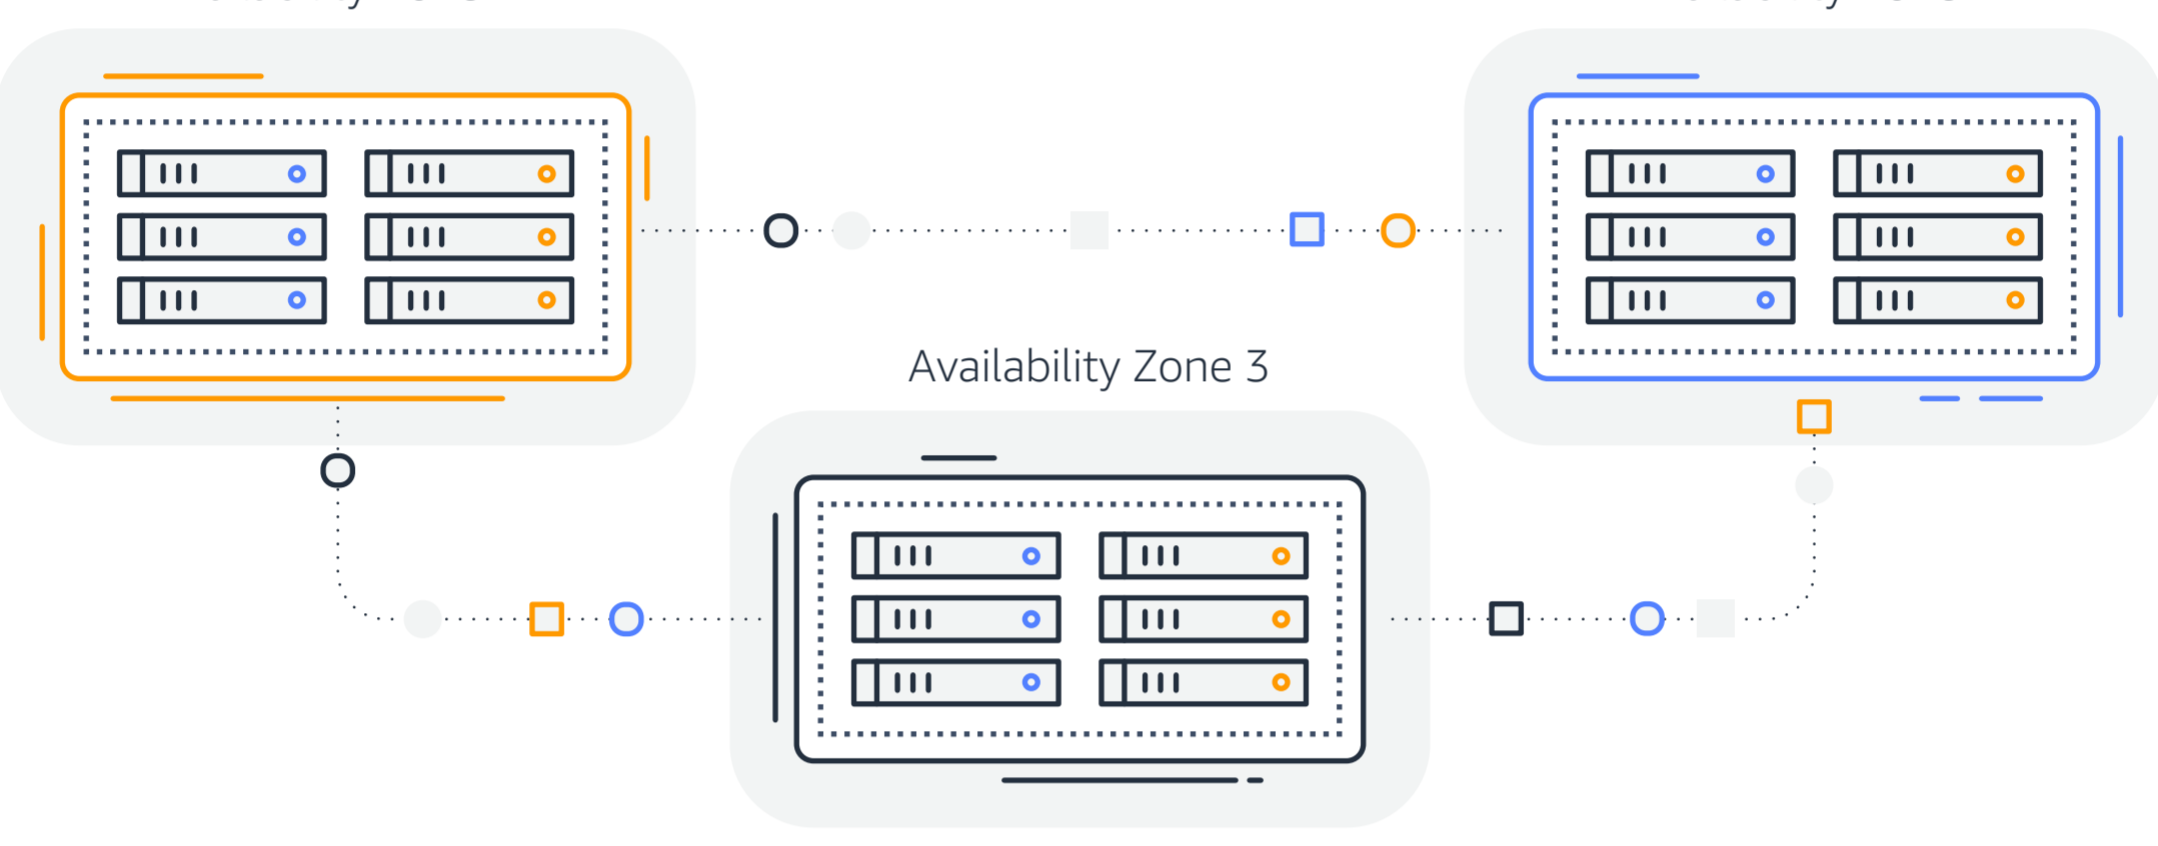
[Learn more about S3 storage classes](#)

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

6 storage classes purpose-built for varying access patterns at corresponding costs

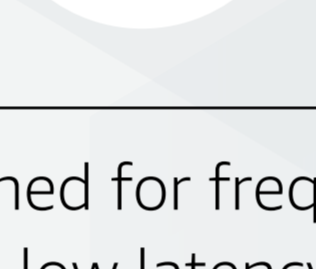


### S3 Standard

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 Intelligent-Tiering

Only cloud storage class with automatic cost savings by moving objects between two tiers



Automatically moves objects between frequent access and infrequent access tiers



Optimizes storage costs based on 30-day access patterns




Data lakes and other data sets with changing or unknown access patterns



### S3 Standard-Infrequent Access (S3 Standard-IA)

High availability and rapid retrieval of S3 Standard—with cost savings



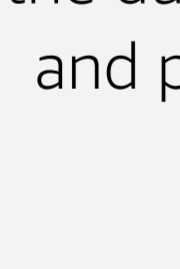
Infrequently accessed data with rapid retrieval and the durability, availability, and performance of S3 Standard



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

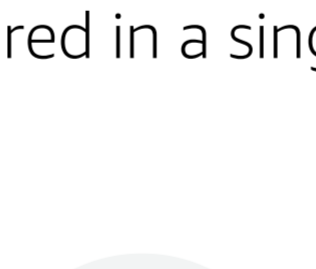


### S3 One Zone-Infrequent Access (S3 One Zone-IA)


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

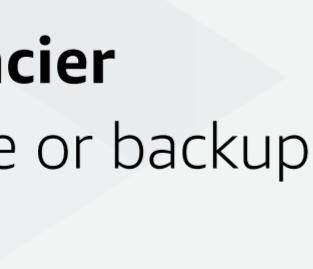


Ideal for secondary backups and workloads with easily re-creatable data

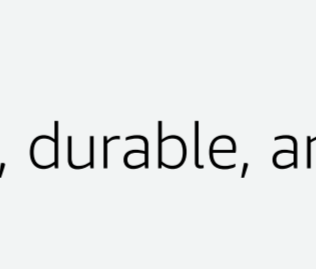


### S3 Glacier

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



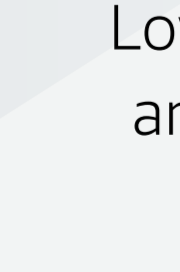
Low-cost, durable archive with low retrieval fees



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 or twice in a year

## Data lifecycle management

Use **S3 Storage Class Analysis** to learn access patterns and discover objects to move to lower-cost classes. Then, use **S3 Lifecycle policies** to transfer objects to lower-cost classes or archive them with **S3 Glacier** and **S3 Glacier Deep Archive** for maximum savings.

[Learn more about Amazon S3 Storage Classes](#)

<https://aws.amazon.com/s3/storage-classes/>