

# Data Retention on Amazon S3 and Amazon Glacier

**Rob Wilson**  
**Senior Product Manager, Amazon S3**

# Benefits of Object Storage on Amazon



## Durable

Designed for 11 9s  
of durability



## Available

Designed for  
**99.99%** availability



## High performance

- Multipart upload
- Range GET
- Parallelize List



## Easy to use

- Simple REST API
- IPv6 support
- Event Notification
- Lifecycle policies



## Scalable

- Store as much as you need
- Scale storage and compute independently
- No minimum usage commitments



## Integrated

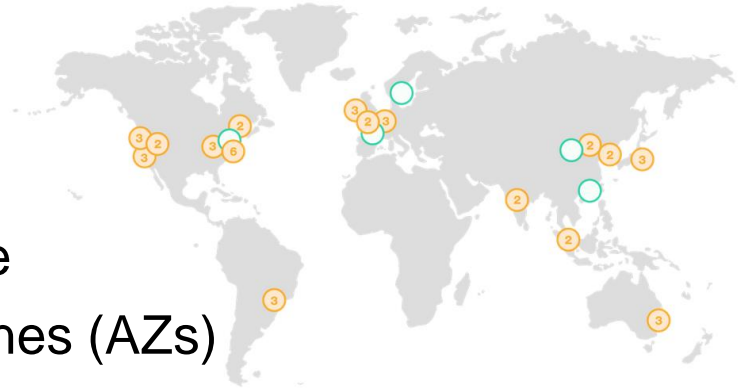
- AWS Elastic MapReduce
- Amazon Redshift
- Amazon DynamoDB
- And more

# Designed for Durability

- AWS Region Design
- Checksum Validations
- Compliance Controls
- Cross Region Replication
- Versioning
- Multi-Factor Authentication

# AWS Region Design


- 16 independent AWS regions worldwide
- Each region has isolated availability zones (AZs)
- Data is stored simultaneously across AZs
- Designed to sustain the concurrent loss of data in two facilities



# Checksum Validations

- Uses CRC, MD5, and SHA256
- Validations performed while data is at rest and in transit
- Data integrity is verified as data is uploaded or downloaded
- Checks performed automatically using the AWS SDK

# Compliance Controls



- Data immutability is enforced on Amazon Glacier using Vault Lock
- Creates an immutable policy
- Can be used to block all deletes or to permit only after a specified time period
- Enables customers to achieve SEC Rule 17a-4(f) (WORM) compliance
- FedRAMP Moderate and FedRAMP High (GovCloud)

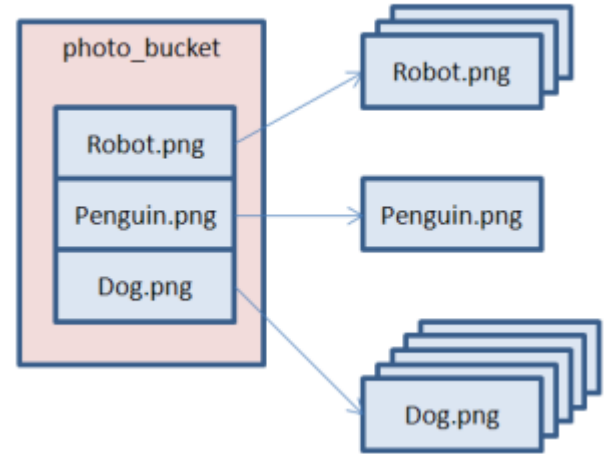
# Cross Region Replication



- Automatic, asynchronous copying of objects across buckets in different AWS regions
- User specifies the target region and bucket
- Replica objects have the same key name and metadata
- Can specify the storage class of the replica, lowering the average cost of storage

# Versioning

- Protects your data from overwrites or accidental deletions, by retaining the original objects
- In both cases, you can easily recover the original object
- Users can access previous versions by specifying the versionID in their requests
- Lifecycle policies allow users to control how much data is retained





# Multi-Factor Authentication (MFA)



- Buckets can be configured for MFA Delete
- Adds another layer of protection for deleting objects or changing the versioning state of the bucket
- Requires two forms of authentication:
  - Your security credentials
  - The concatenation of a valid serial number, a space, and the six-digit code displayed on an approved authentication device

# Questions