

AWS
re:Invent

STG320-R

Amazon S3 performance and cost optimization

Janko Jerinic

Sr. Software Engineer
Amazon S3
Amazon Web Services

Agenda

Amazon S3 storage classes

Amazon S3 object tagging & lifecycle management

Amazon S3 inventory

Amazon S3 storage class analysis

Summary

Related breakouts

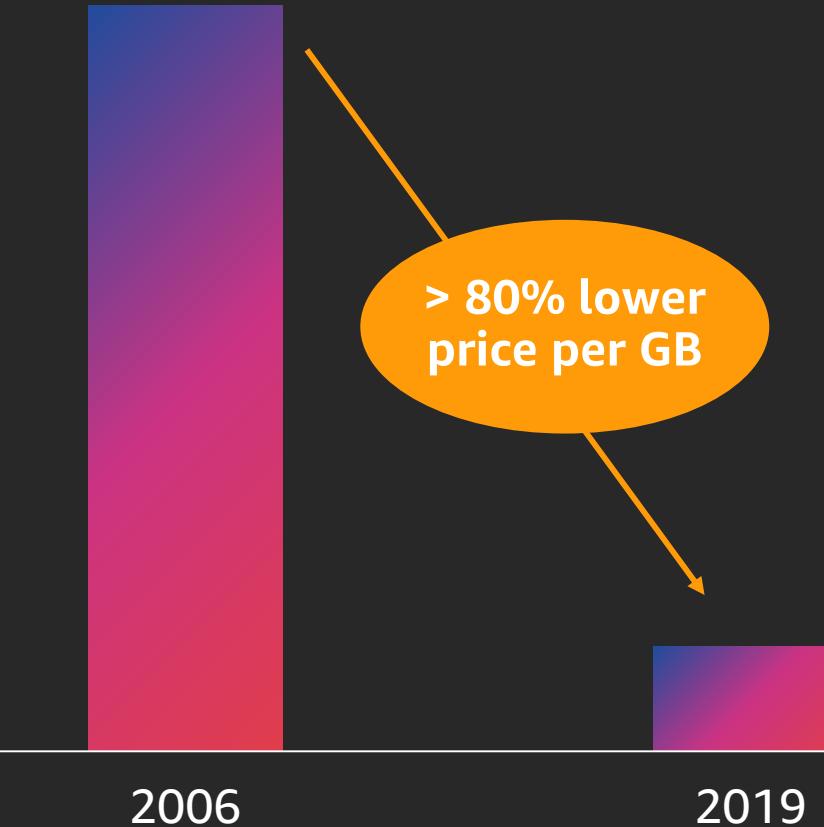
- **STG343 - Optimize your storage performance with Amazon S3, ft. Lyft**
- **STG332 - Guidelines and design patterns for optimizing cost in Amazon S3**

Amazon S3 Storage Classes

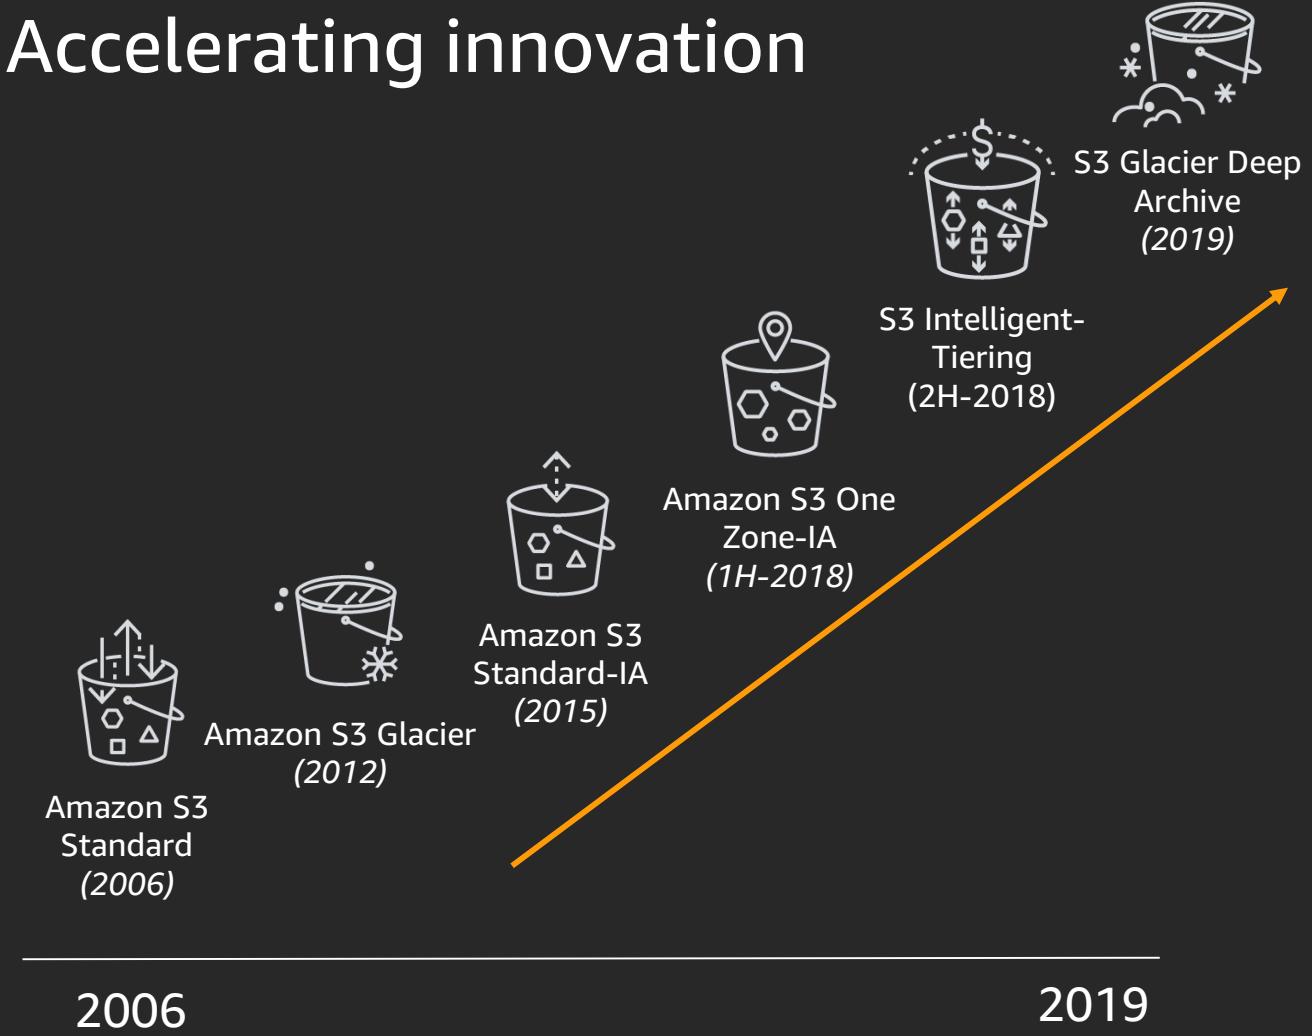
Amazon S3 Storage Classes

Optimize your storage cost by utilizing all Amazon S3 storage classes

Decreasing storage prices



Accelerating innovation



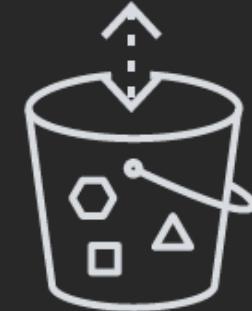
Your choice of Amazon S3 storage classes



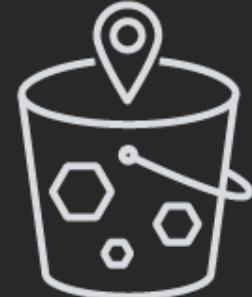
S3 Standard



S3 Intelligent-Tiering



S3 Standard-IA



S3 One Zone-IA



S3 Glacier



S3 Glacier Deep Archive

Frequent

- Active, frequently accessed data
- Milliseconds access
- ≥ 3 AZ
- \$0.0210/GB

- Data with changing access patterns
- Milliseconds access
- ≥ 3 AZ
- \$0.0210 to \$0.0125/GB
- Monitoring fee per object
- Min. storage duration

Access frequency

- Infrequently accessed data
- Milliseconds access
- ≥ 3 AZ
- \$0.0125/GB
- Retrieval fee per GB
- Min. storage duration
- Min. object size

- Re-creatable, less accessed data
- Milliseconds access
- 1 AZ
- \$0.0100/GB
- Retrieval fee per GB
- Min. storage duration
- Min. object size

- Archive data
- Select minutes or hours
- ≥ 3 AZ
- \$0.0040/GB
- Retrieval fee per GB
- Min. storage duration
- Min. object size

- Long-term archive data
- Select hours
- ≥ 3 AZ
- \$0.00099/GB
- Retrieval fee per GB
- Min. storage duration
- Min. object size

Amazon S3 Intelligent-Tiering is the world's only cloud storage that automates cost savings



Automatically optimizes storage costs for data with changing access patterns

Stores objects in two access tiers, optimized for frequent and infrequent access

Monitors access patterns and optimizes cost on granular object level

No performance impact, no operational overhead, no retrieval fees

Customers of all sizes and virtually every industry use Intelligent-Tiering and save on cost automatically

Lifecycle policies use rules to manage your storage

Use Lifecycle policies to transition objects to another storage class

Lifecycle rules take action based on object age, here's an example

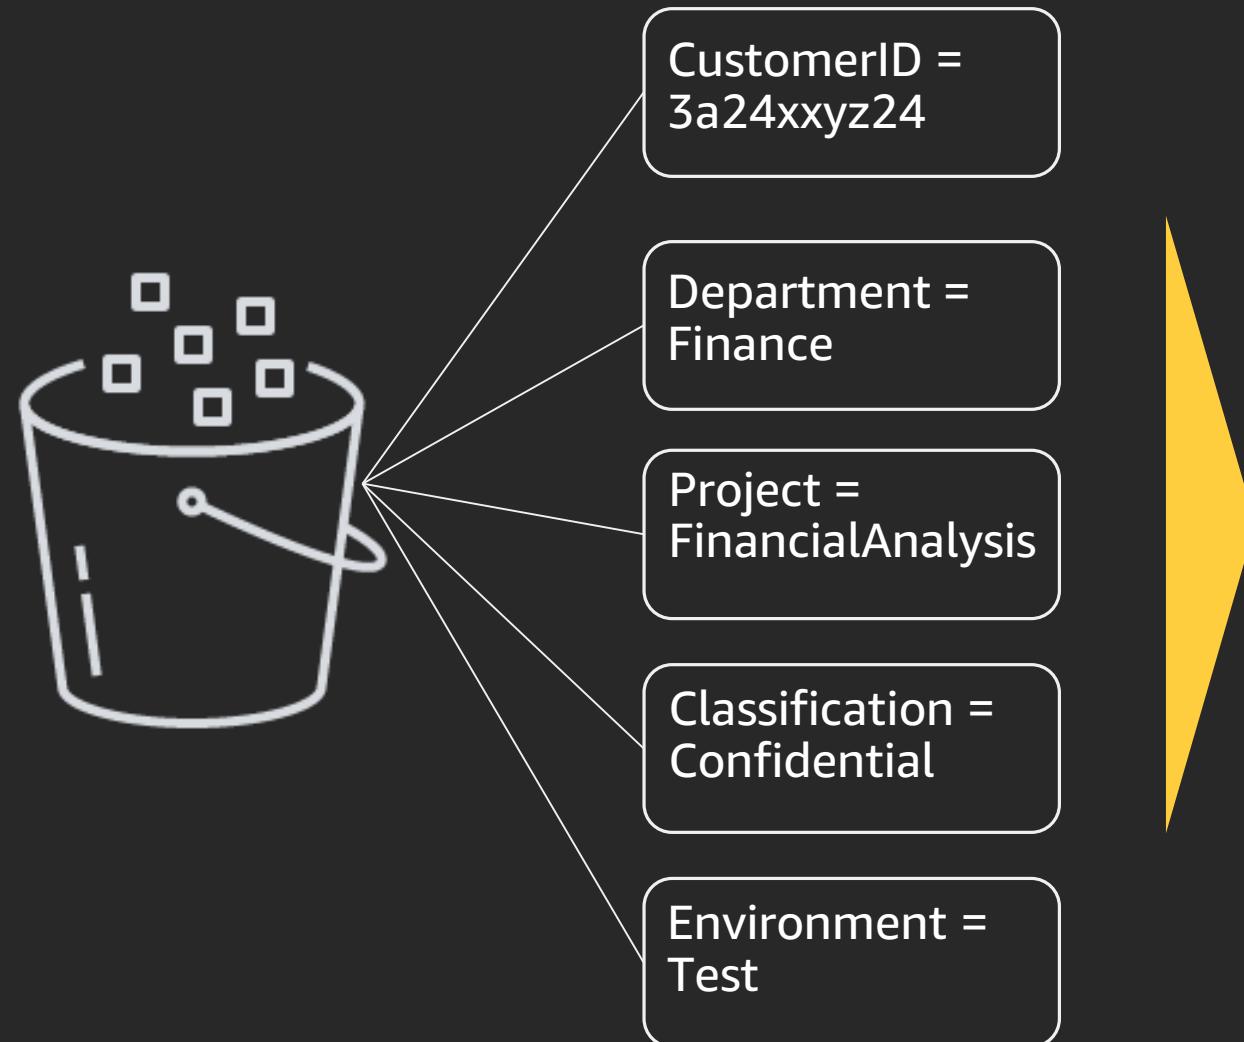
1. Move objects older than 30 days to S3 Standard – Infrequent Access
2. Move objects older than 365 days to S3 Glacier Deep Archive



Demo: Amazon S3 Storage Classes

S3 Object Tagging and Lifecycle Management

Organize data with Object Tags and Prefixes



A single bucket can contain **objects** stored across any storage class

Control access, analyze usage, manage lifecycle policies, and replicate objects

Up to 10 **mutable metadata tags** (key-value pairs) per object

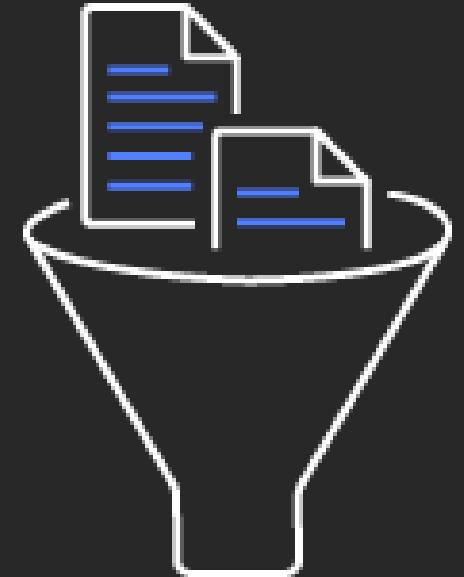
Tag objects when created, later, or both

Object tags work with lifecycle policies

Perform automated actions on a subset of your data with object tags

Lifecycle

- Specify a tag filter to transition or expire objects
- Use S3 Batch Operations to apply object tags at scale
- Ex: Transition all objects tagged “Project : Delta” to S3 Glacier



Use lifecycle policies with object tag filters

Object tag filters simplify lifecycle policies when the same action needs to be performed across multiple prefixes in the bucket

```
<Filter>
  <And>
    <Tag>
      <Key>Project</Key>
      <Value>Delta</Value>
    </Tag>
  </And>
</Filter>
```

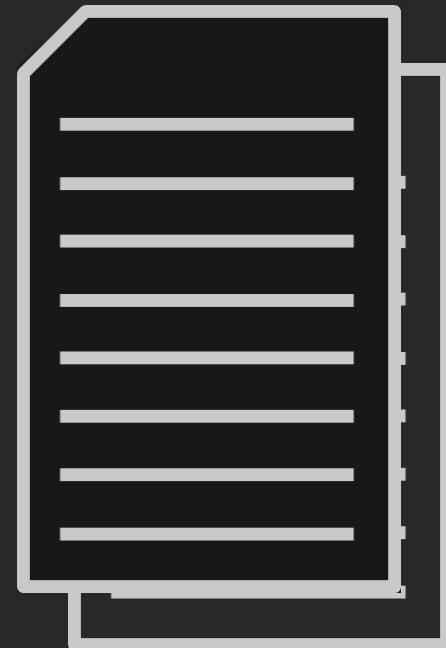
Demo: S3 Object Tagging and Lifecycle Management

Amazon S3 Inventory

Amazon S3 Inventory

A managed alternative to using the LIST API

Regularly generates a list of objects for analytics and auditing.



- ✓ Storage class
- ✓ Creation date
- ✓ Encryption status
- ✓ Replication status
- ✓ Object size, and more
- ✓ S3 Intelligent-Tiering access tier **NEW!**

Use Amazon Athena to filter S3 Inventory Reports

This query selects bucket, object key, and version id for unencrypted objects

```
select s._1, s._2, s._3 from s3object s where s._6 = 'NOT-SSE'
```

Example results:

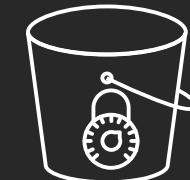
```
batchoperationsdemo,0100059%7Ethumb.jpg,1srtIxksLu0R0ZkYPL.LhgD5caTYn6vu
batchoperationsdemo,0100074%7Ethumb.jpg, sd2M60g6Fdazoi6D5kNARIE7KzUibmHR
batchoperationsdemo,0100075%7Ethumb.jpg,TLYESLn11mXD5c4BwioIinqFrktddkoL
```

Amazon S3 Batch Operations **NEW!**

Save time when performing one-time or recurring actions at scale



- Replace object tag sets
- Change object ACLs
- Restore objects from Amazon S3 Glacier
- Copy objects
- Run AWS Lambda functions



Manage millions or billions of objects with a single request

Automatically handles retries, displays progress, and generates reports

Demo: Amazon S3 Inventory

Amazon S3 Analytics – Storage Class Analysis

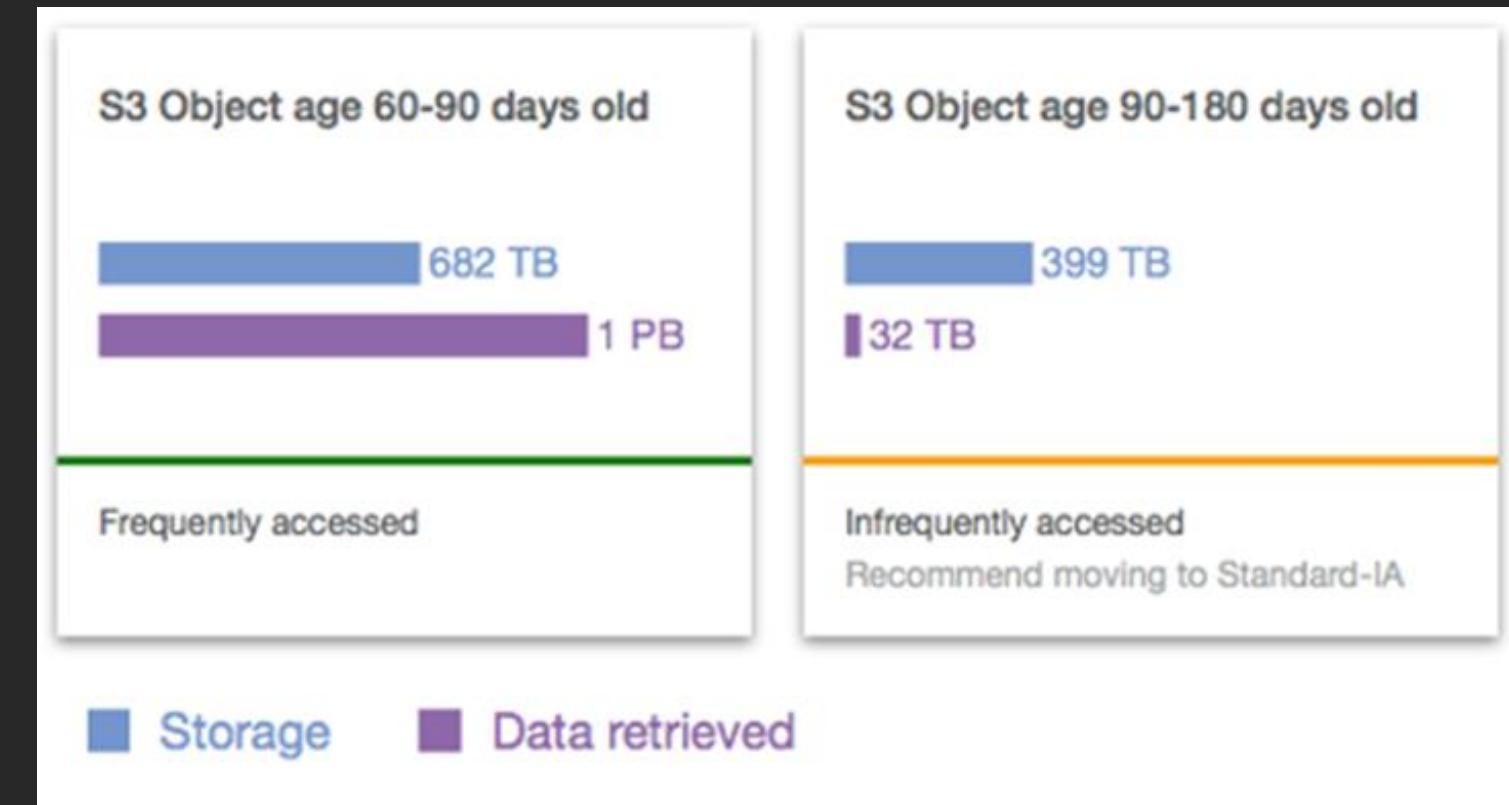
S3 Storage Class Analysis

Provides lifecycle policy recommendations based on access patterns

Monitors access patterns

Classifies data as frequently
or infrequently accessed

Can be filtered by bucket,
prefix, or object tag



Demo: S3 Storage Class Analysis

Summary

Putting it all together

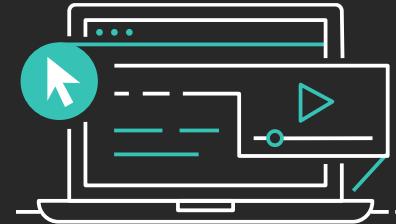
- Take advantage of different Amazon S3 storage classes to optimize cost
- Use tags and prefixes to organize your data
- Use lifecycle policies to expire objects and transition them across storage classes
- Use Amazon S3 Inventory and Amazon S3 Storage Class Analysis to monitor and analyze your data
- Use Amazon S3 Intelligent-Tiering to automate cost savings
- Use Amazon S3 Glacier Deep Archive for archival and lowest storage prices

References

- Amazon S3 Storage Classes
 - <https://docs.aws.amazon.com/AmazonS3/latest/dev/storage-class-intro.html>
- Amazon S3 Inventory
 - <https://docs.aws.amazon.com/AmazonS3/latest/dev/storage-inventory.html>
- Amazon S3 Analytics – Storage Class Analysis
 - <https://docs.aws.amazon.com/AmazonS3/latest/dev/analytics-storage-class.html>
- Amazon S3 Object Lifecycle Management
 - <https://docs.aws.amazon.com/AmazonS3/latest/dev/object-lifecycle-mgmt.html>
- Amazon S3 Object Tagging
 - <https://docs.aws.amazon.com/AmazonS3/latest/dev/object-tagging.html>

Learn storage with AWS Training and Certification

Resources created by the experts at AWS to help you build cloud storage skills



45+ free digital courses cover topics related to cloud storage, including:

- Amazon S3
- AWS Storage Gateway
- Amazon S3 Glacier
- Amazon Elastic File Storage (Amazon EFS)
- Amazon Elastic Block Storage (Amazon EBS)



Classroom offerings, like *Architecting on AWS*, feature AWS expert instructors and hands-on activities

Visit aws.amazon.com/training/path-storage/

Thank you!

Janko Jerinic
jerinicj@amazon.com



Please complete the session
survey in the mobile app.