



Optimizing Cost in Amazon S3

Peter Sjoberg
Go To Market Specialist
Amazon S3
Amazon Web Services

Agenda

Amazon S3 overview

Pillars of cost optimization


S3 storage class options

More cost-optimization guidelines and design patterns

Putting it all together

Amazon S3 overview


Broad Portfolio of Storage Services

Block storage 

- General purpose SSD
- Provisioned IOPS SSD
- Throughput-optimized HDD
- Cold HDD
- Elastic volumes


Backup 

- AWS Backup

Data transfer 

- AWS DataSync
- AWS Transfer for SFTP



File storage 

- Amazon EFS Standard
- Amazon EFS Infrequent Access (Amazon EFS IA)
- Amazon FSx for Windows File Server
- Amazon FSx for Lustre

Object storage 

- Amazon S3 Standard
- Amazon S3 Standard-IA
- Amazon S3 One Zone-IA
- Amazon S3 Glacier
- Amazon S3 Intelligent-Tiering
- Amazon S3 Glacier Deep Archive

Amazon S3 offers industry-leading scalability, availability, security, and performance



Exabytes stored across many millions of drives
AmazonS3 stores trillions of objects around the world

We regularly peak at millions of requests per second

In a single region, we process peaks of over 60 tbps in a day


235+ distributed micro-services in 3 or more AZs per Region

Pillars of cost optimization


Pillars of cost optimization



Application
requirements



Data
organization

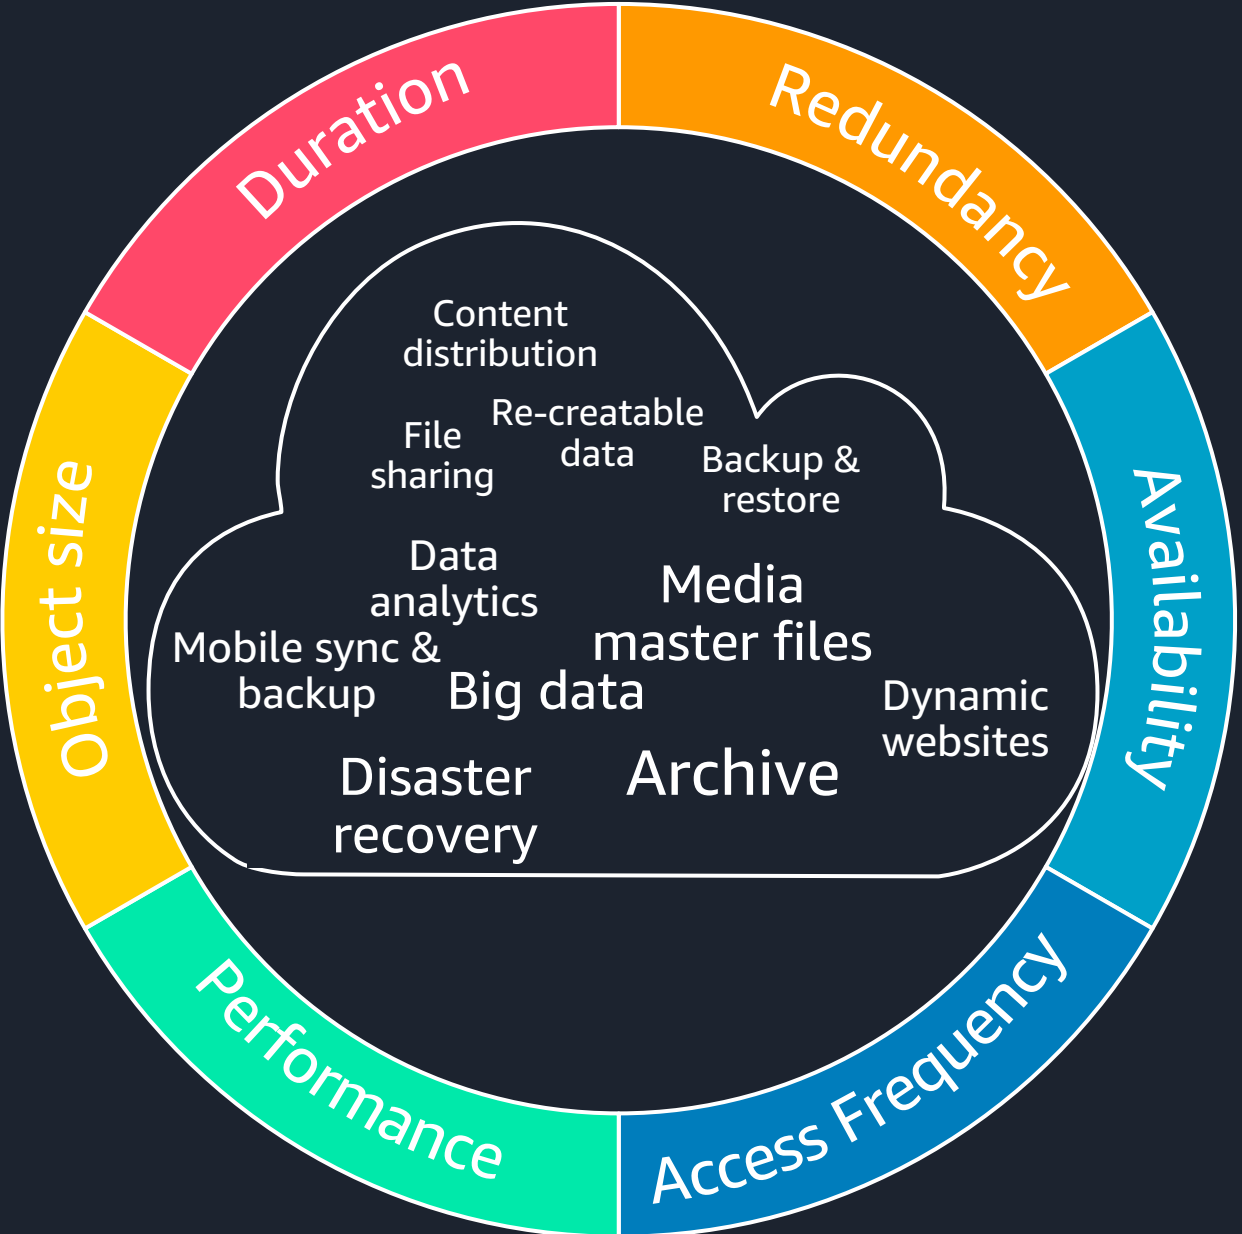


Right
sizing

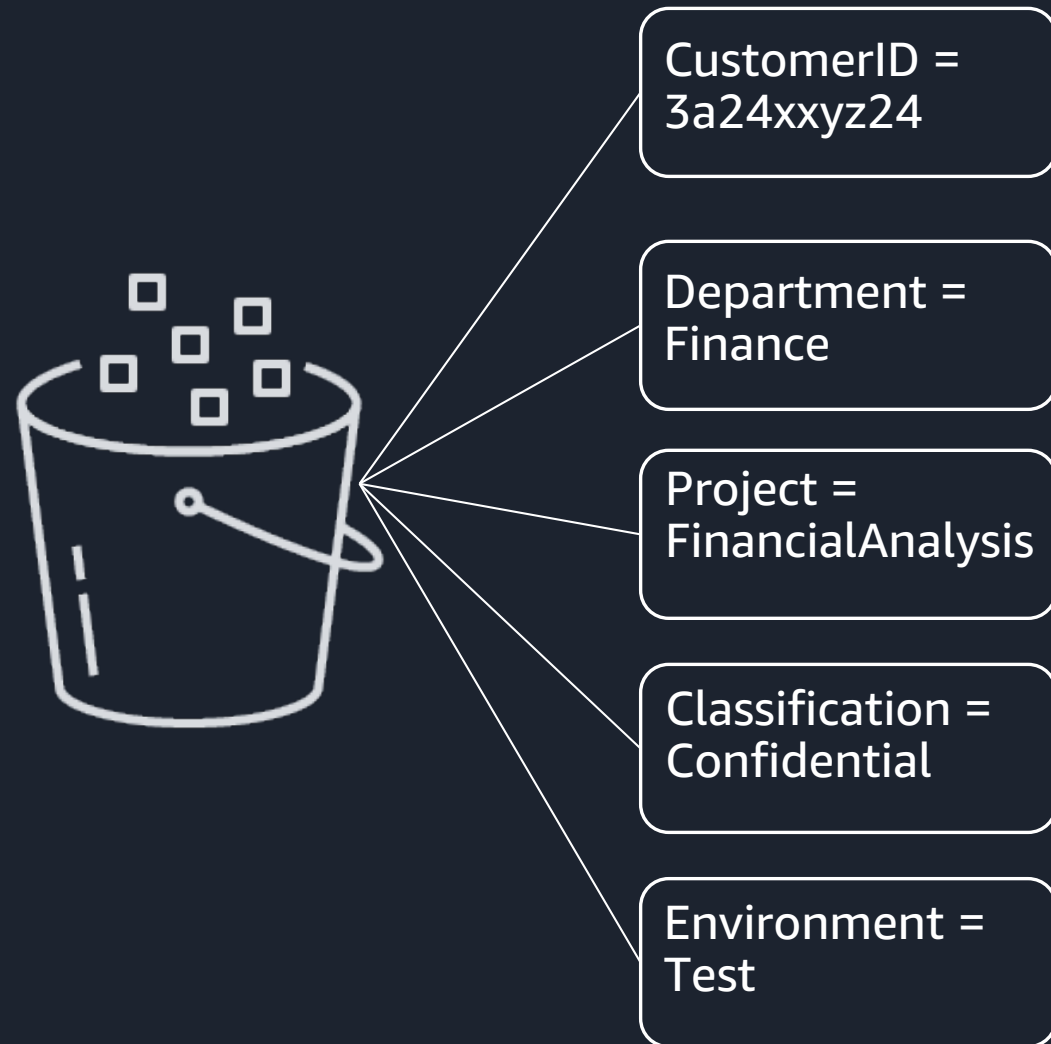



Monitor, analyze,
optimize, and manage

Define application requirements



Organize data with object tags and prefixes ...



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

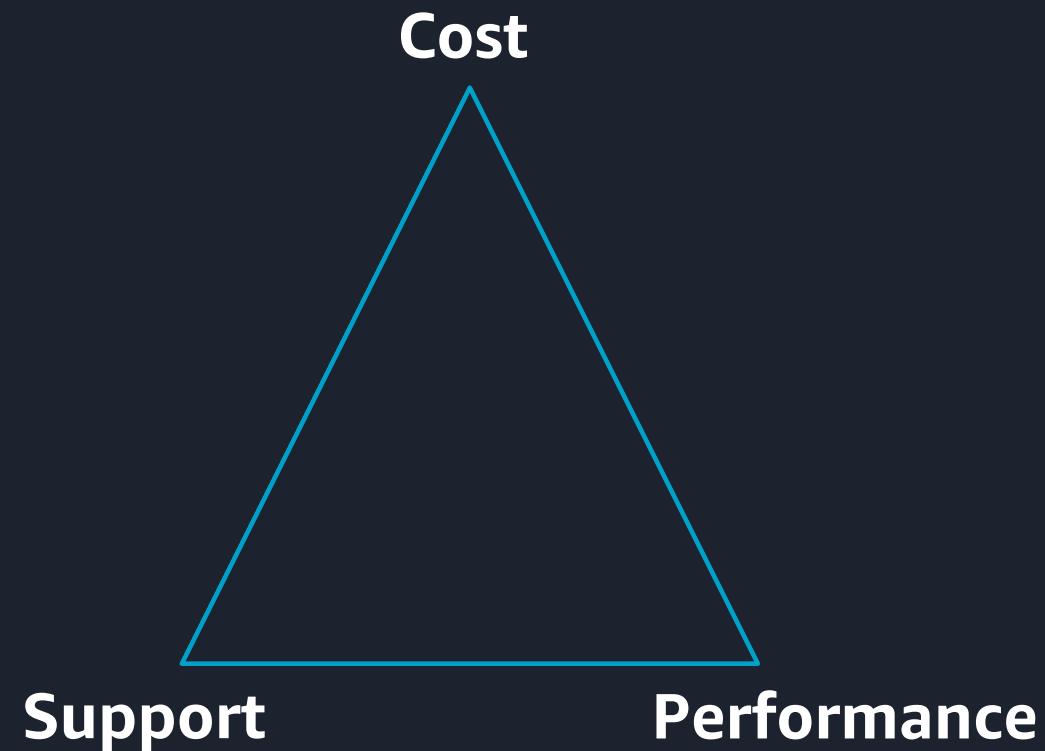

Control access, analyze usage, manage lifecycle policies, and replicate objects

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

Tag objects when created, later, or both

Organize your data from the beginning to enable easy cost optimization

... and decide on your data format



No best data format but tradeoff between cost, performance, and support

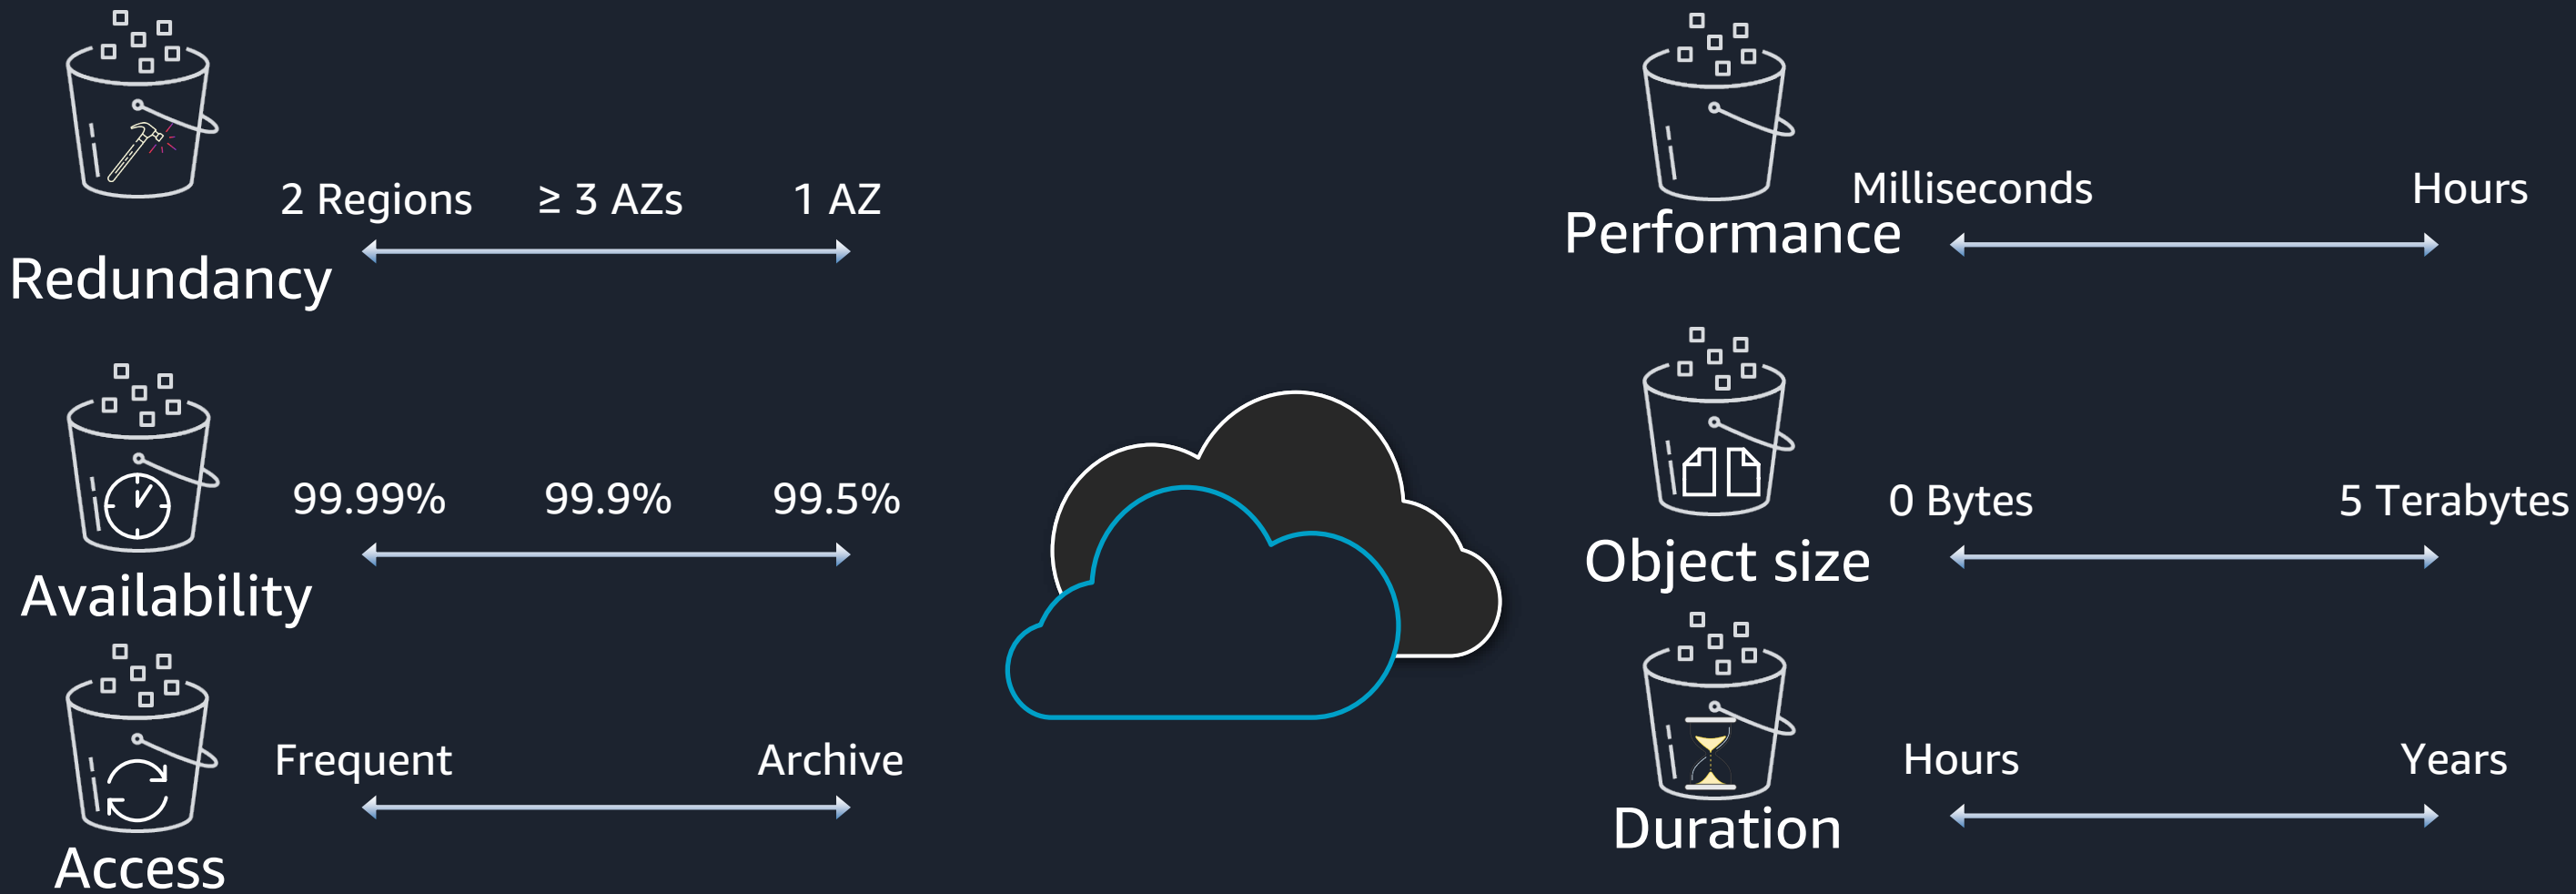
Unstructured formats (e.g., CSV, TSV, JSON, XML) easy to understand and process but less efficient

Columnar compressed formats (e.g., Parquet, ORC, CarbonData) provide lower cost (inherent compression) and more efficient scan and query (self-describing)

Convert raw data into optimized format and keep raw data compressed in archive

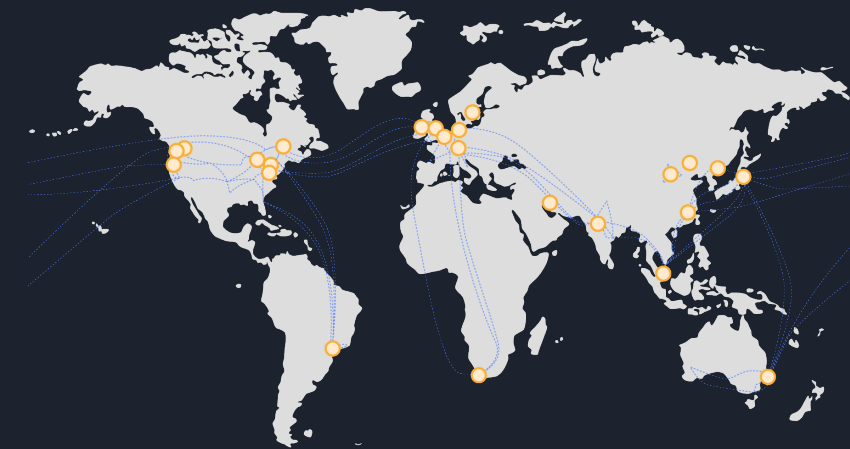
Use efficient data formats to get cost savings and performance gains

Choose the storage class that fits best



Reduce storage cost > 95% by choosing the storage class that best fits your workload

Amazon S3 is designed for 11 9's of durability



24 Regions

Geographic locations with multiple, physically miles separated and isolated Availability Zones



76 Availability Zones (AZs)

Objects stored across multiple devices spanning a min of 3 AZs. Multi-AZ storage classes function normally if an AZ is lost



Highly Durable

Designed to sustain concurrent device failures and data in the event of an entire AZ loss

Operational performance is second only to security


Monitor, analyze, optimize and manage




Monitor with Amazon S3 inventory, Amazon CloudWatch, S3 server access logging



Understand access patterns with Amazon S3 storage class analysis



Perform API actions across thousands, millions, or billions of objects with Amazon S3 Batch Operations



Tier and expire storage manually with Amazon S3 lifecycle policies
or
automatic cost optimization with S3 Intelligent-Tiering



Amazon S3 storage class options

AWS pricing principles



No upfront investment



Pay-as-you-go approach



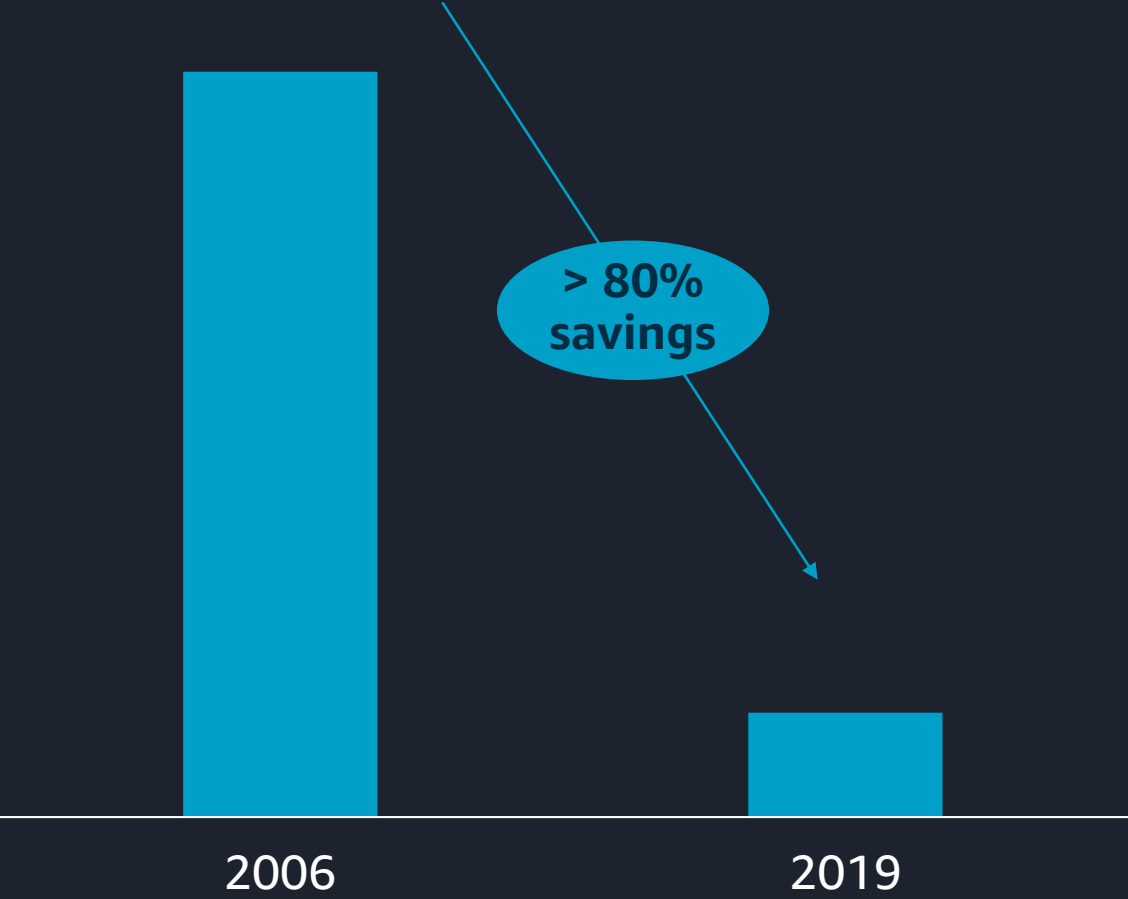
Pay less by using more



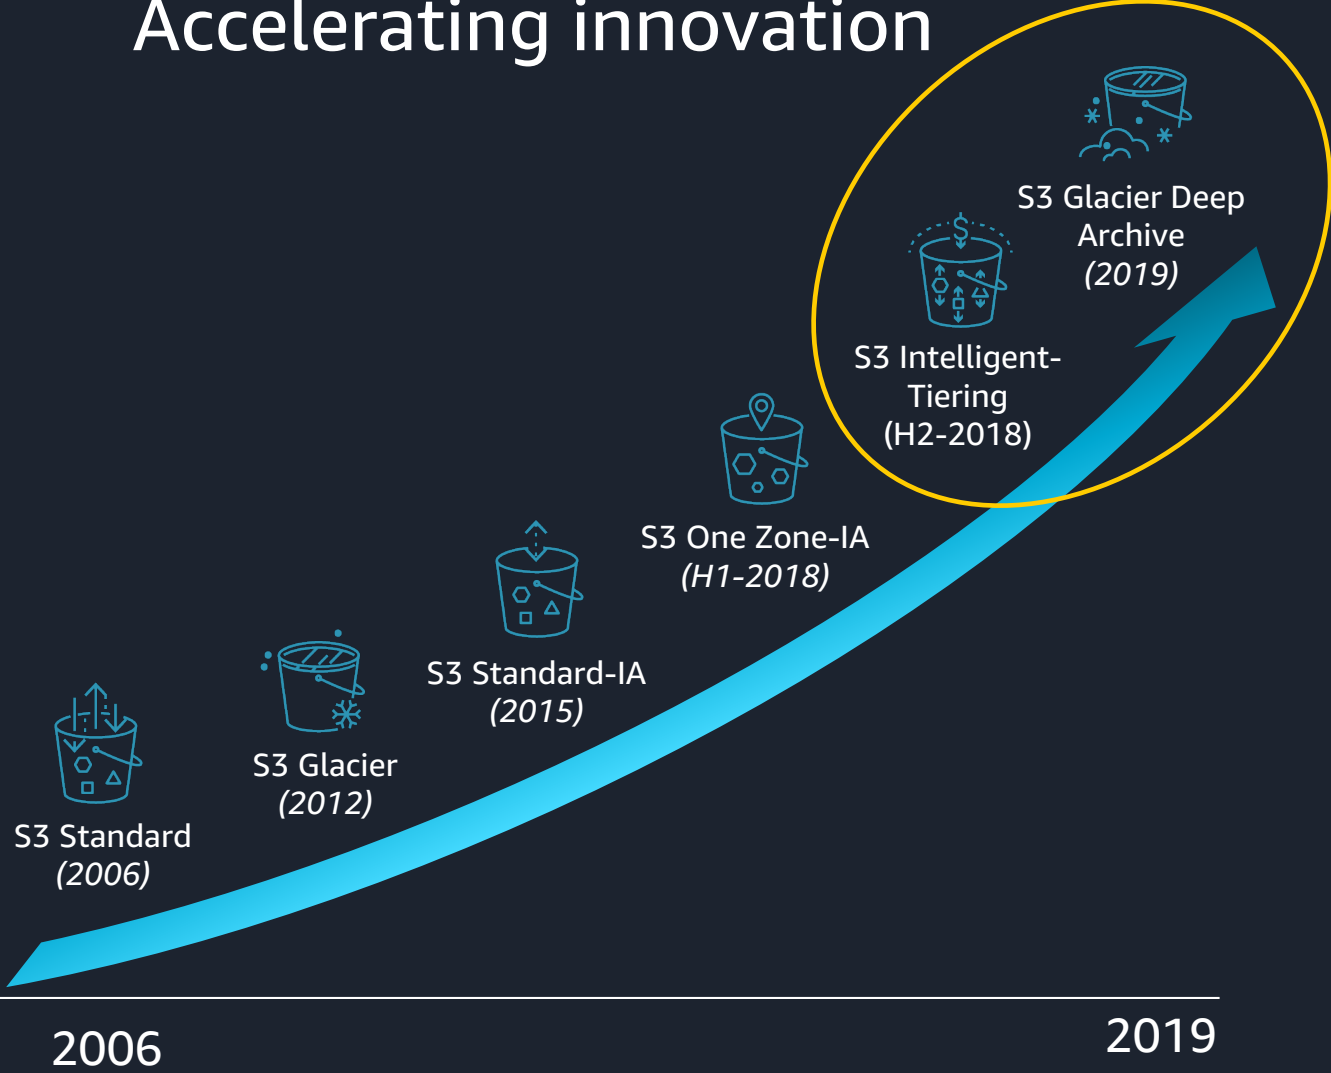
Pay less as AWS grows

Decreasing prices and more storage options


Decreasing storage prices



Accelerating innovation



Your choice of Amazon S3 storage classes



S3 Standard

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



S3 Intelligent-Tiering

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



S3 Standard-IA

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



S3 One Zone-IA

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



S3 Glacier

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



S3 Glacier Deep Archive

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

Frequent

Access frequency

Archive

S3 INT is the world's only cloud storage class 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 INT and save on cost automatically

Ideal use cases for Amazon S3 Intelligent-Tiering



Data Lakes

Storage with changing access patterns used by multiple applications



Big Data Analytics

Unpredictable workloads where storage access pattern vary depending on application



Long tail

Constraint on resources and experience to optimize storage themselves

Dynamic cost optimization with no performance impact and no operational overhead

S3 Glacier Deep Archive lowest cost cloud storage



No tape to
manage



Designed for
99.9999999999%
durability



Recover data in
hours vs.
days/weeks



\$0.00099 per GB-
month
Less than 1/4 the cost of
S3 Glacier

Ideal use cases for S3 Glacier Deep Archive



Compliance Archives

Many enterprises like Financial Services and Healthcare must retain regulatory and compliance archives for extended durations

Data Retention

Retain valuable data over the long term for unanticipated future needs

Disaster Recovery

Meet disaster recovery objectives by replicating critical data to a secondary region

Lowest cost storage in the market for long-term data retention requirements

More cost-optimization guidelines and design patterns

Monitor your Amazon S3 costs




AWS Budgets gives you the ability to set custom budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount



Amazon CloudWatch allows you to monitor metrics on your daily S3 storage and minutely request activity or to set up alarms on request counts and data transfer in real time

Lower storage cost by monitoring and staying within your planned budget

Decide on your data protection and retention strategy



Backup and Recovery

Accidental deletion
Existing backup mechanisms
Backup data retention



Data Retention

Data value
Access time
Compliance

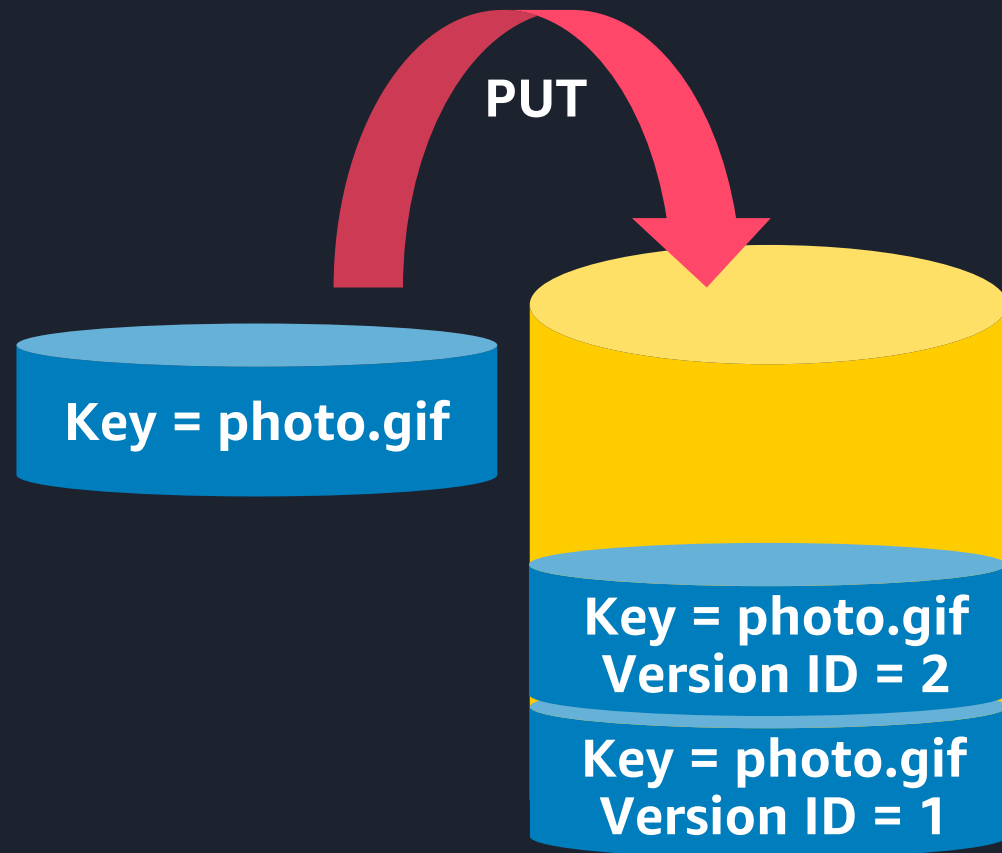


Disaster Recovery

Recovery point objective
Recovery time objective
Geographic redundancy

Understand and define data protection requirements to control costs

Decide on your versioning strategy



Create a **new version with every upload**. Previous versions are retained, not overwritten

Making **delete requests without a version ID** removes access to objects but keeps the data


Manage **previous versions with lifecycle**. Transition or expire objects a specified number of days after they are no longer the current version

Manage complex **cleanup workflow with S3 Inventory** (e.g., retaining a fixed number of previous versions)

Lower storage cost by deleting previous versions of objects when no longer needed

Decide on your replication strategy

Amazon S3 same-region replication (SRR)



Amazon S3 cross-region replication (CRR)




Replicate objects to a bucket in the same or different AWS account




Amazon S3

Amazon S3 replication with S3 object lock

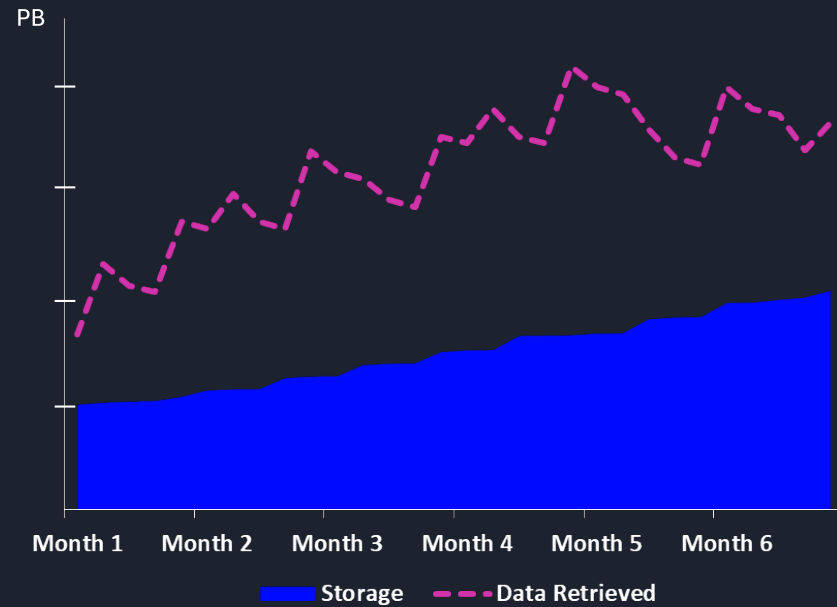


Amazon S3 replication time control



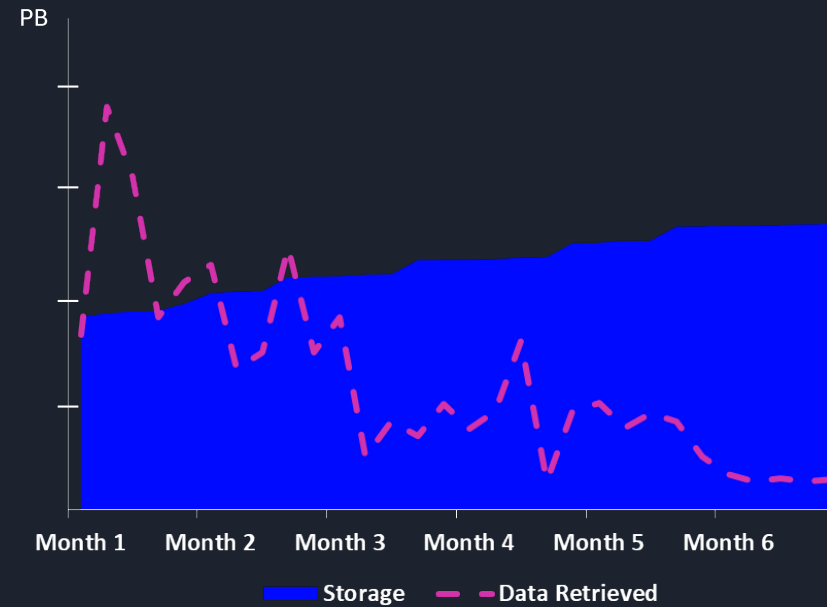
Lower storage cost by leveraging the full breadth of storage classes and flexible policies

Understand how your storage is accessed



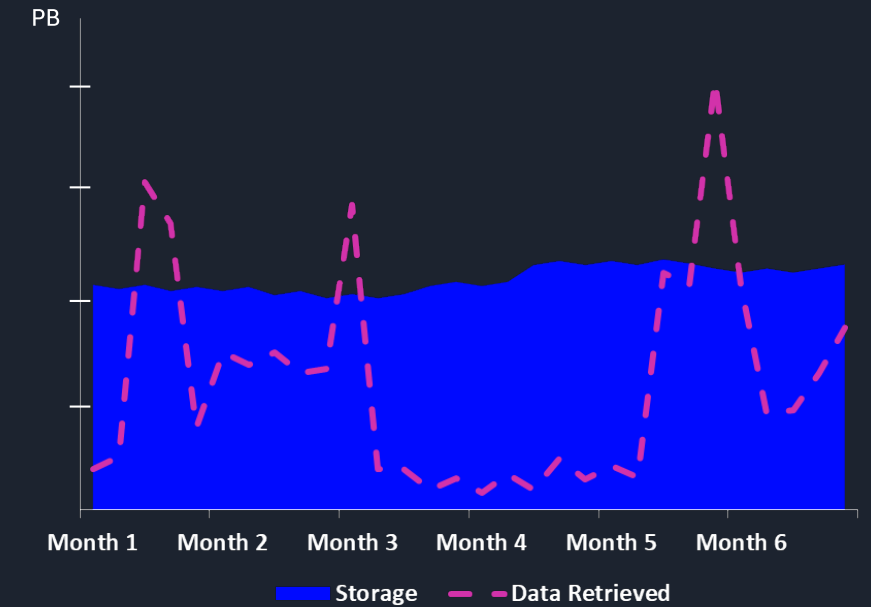
Frequent Access

- Over 100% of storage retrieved per mo
- Storage remains frequently accessed
- e.g.: Big data analytics, dynamic website hosting, IoT sensor data
- Storage classes:
 - S3 Standard, maybe S3 INT



Infrequent Access

- Less than 100% of storage retrieved per mo
- Storage “cools off” over time
- e.g.: Data logs, customer-generated content, mobile sync & backup
- Storage classes:
 - S3 Standard + Lifecycle + S3 S-IA/ Z-IA
 - S3 INT for automated cost savings
 - S3 Glacier/Glacier Deep Archive for archiving




Changing Access

- Changing or unpredictable access patterns; data “warms up” after some time again
- e.g.: Data lakes, ML training data, geospatial data
- Storage classes:
 - S3 INT for automated cost savings

Use S3's building blocks or S3 INT to understand access patterns

Option 1: Optimization with S3's building blocks



- Storage classes analysis classifies data as frequently or infrequently accessed
- Great for predictable workloads in combination with lifecycle management
- Can be filtered by bucket, prefix, or object tag

Option 2: Automatic optimization with S3 Intelligent-Tiering



- S3 Intelligent-Tiering stores objects in two access tiers – frequent access and infrequent access
- Great for unpredictable or unknown access patterns
- Amazon CloudWatch shows bytes by access tier and S3 inventory shows objects by access tier

Set S3 lifecycle policy to tier, archive, and expire storage systematically



Use S3 storage class analysis to identify storage age groups that are less frequently accessed

S3 lifecycle policy to **tier to lower cost storage classes** and expire objects. Fine-tune analysis by bucket, prefix, or object tag

For **example**, transition objects older than 180 days to S3 Glacier, and objects older than 365 days to S3 Glacier Deep Archive

Use S3 lifecycle policy to automatically **cancel all incomplete multipart uploads**

Lower storage cost for predictable workloads (object age indicates access frequency)

Manage billions of objects with S3 Batch Operations



Choose Objects

- S3 inventory report
- CSV list

Use case:


Create **S3 inventory report** and run Amazon Athena query to get objects >2MB



Select Operation

- Copy; restore from S3 Glacier; put access control list (ACL); replace object tag sets; run Lambda functions

Run **S3 Batch Operations** to replace object tag sets for objects >2MB



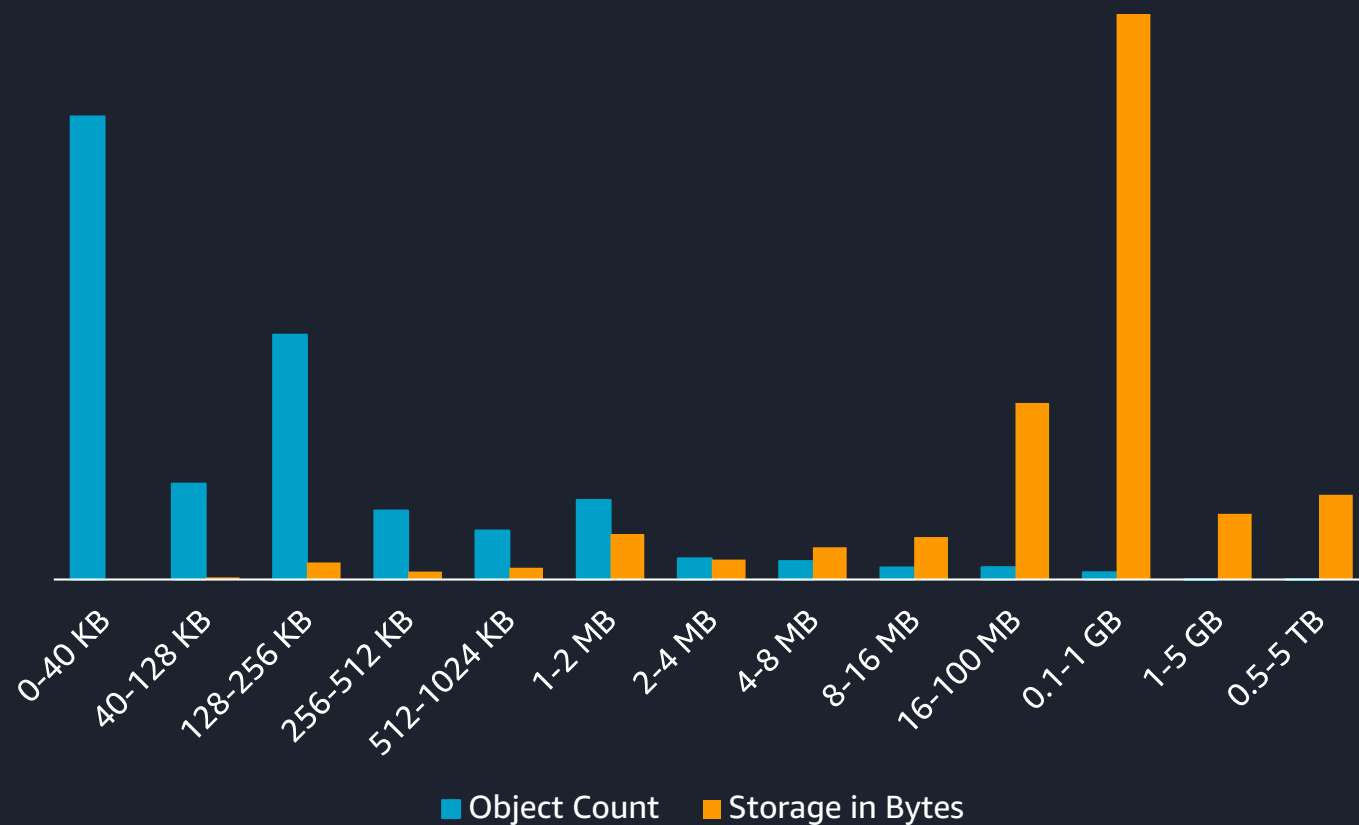
View Progress

- Object level progress
- Job notifications, automatic retries
- Completion report, auditing

Create **S3 lifecycle policy** to transition tagged objects from S3 Standard to INT

Know the characteristics of your dataset

Object Size Distribution



Individual objects can range in size from a min of 0 to a max of 5 terabytes with object part sizes up to 5 gigabytes

Understand object count and storage byte distribution of your storage (min/max, average, by size bins)

Optimize storage cost where majority of storage bytes are and not necessarily where objects counts are cumulated


Compact smaller objects into bigger ones to store more cost efficiently and process faster

Understand your performance requirements

Synchronous
vs
asynchronous



Measure and
iterate



Design key space
and scale horizontally



Monitor and retry
stragglers



Understand
access patterns



Plan for
peak restore



Optimize
object size

[Learn more about how to optimize performance in S3 via AWS Training](#)

The power of S3 storage classes



Industry-leading performance, scalability, availability, and durability

Unmatched security, compliance, and audit capabilities

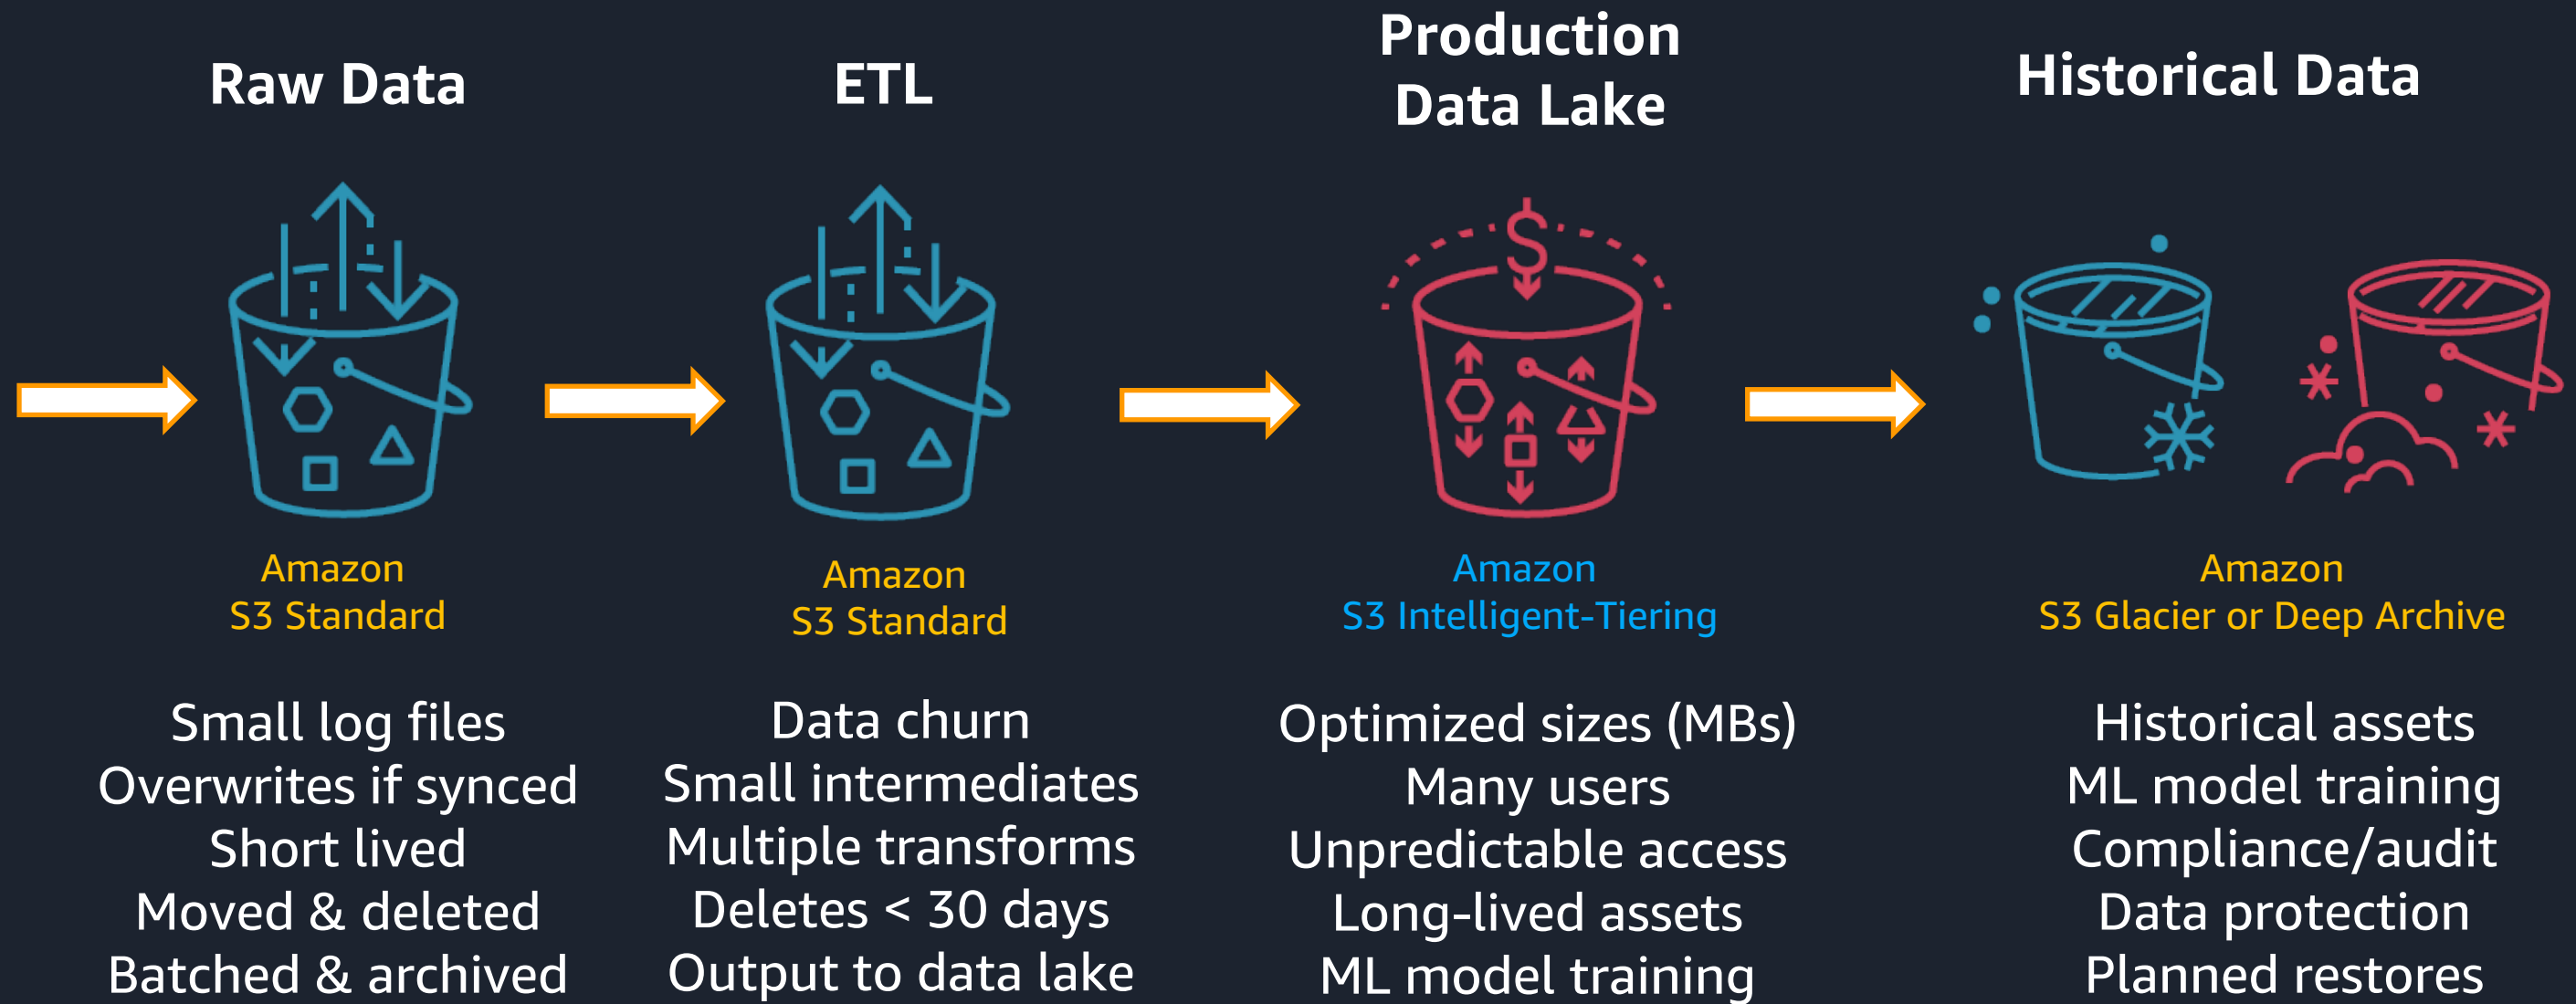
Cost optimization on a granular object-level

Storage classes that fit your workload

Set of analytics services such as S3 Select, Athena, Amazon Redshift, and Amazon EMR

Optimize your storage cost and performance by utilizing all S3 storage classes

Choosing the right storage class for data lakes




Optimize storage cost for all stages of a data lake workflow

Putting it all together

Putting it all together



Understand your application requirements



Use tags and prefixes to organize your data



Cost optimize across all storage classes (per object, tag, prefix)



Use S3 building blocks to monitor, analyze, and manage storage



Store in S3 Intelligent-Tiering for automated cost savings



Archive to S3 Glacier Deep Archive for the lowest storage price

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 System (Amazon EFS)
- Amazon Elastic Block Store (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!