
AWS
re:Invent

STG205

Amazon S3 foundations: Best practices for Amazon S3

Bijeta Chakraborty
Sr. Product Manager, Amazon S3
AWS

Naveen Nalam
Software Engineer, Infrastructure
Airbnb

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

Storage classes and cost optimization

Data protection

Analytics and insights

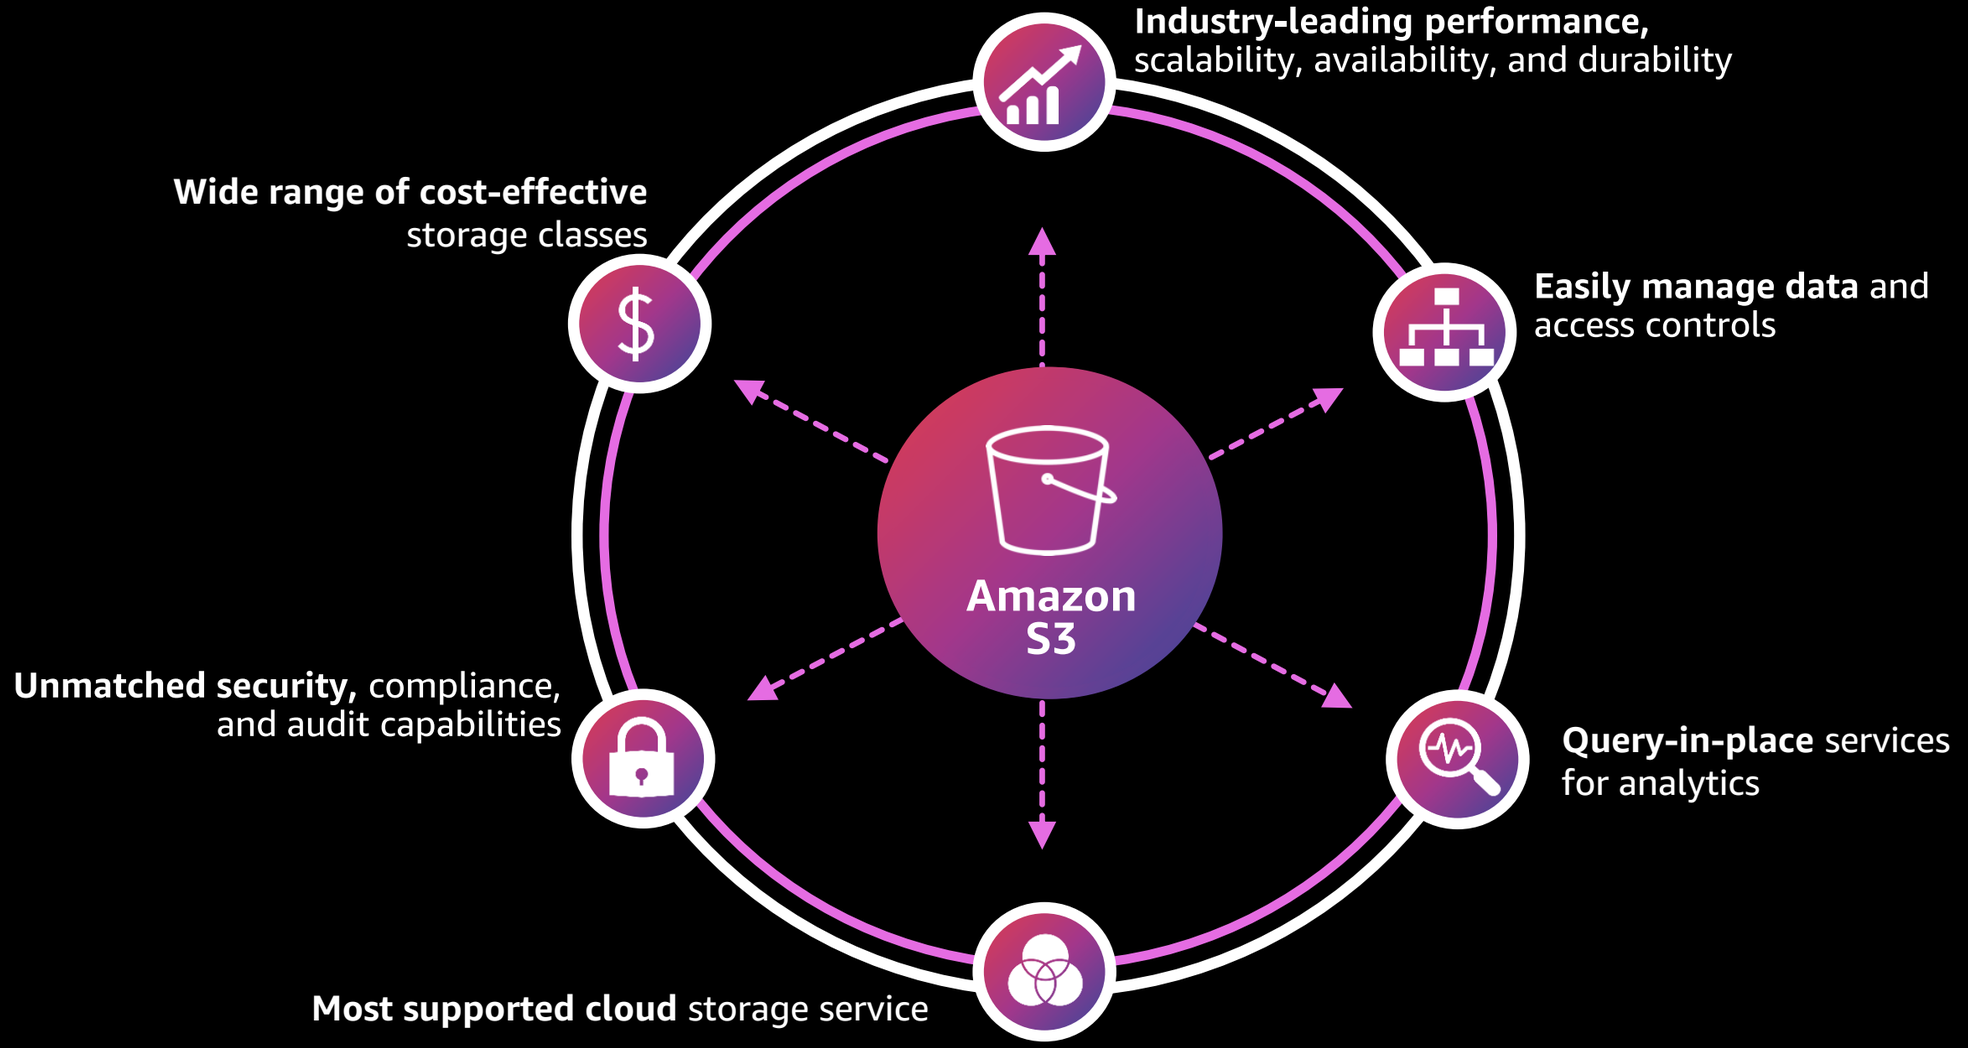
Customer spotlight – Airbnb

Security



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Simple Storage Service (Amazon S3) overview



Storage classes and cost optimization



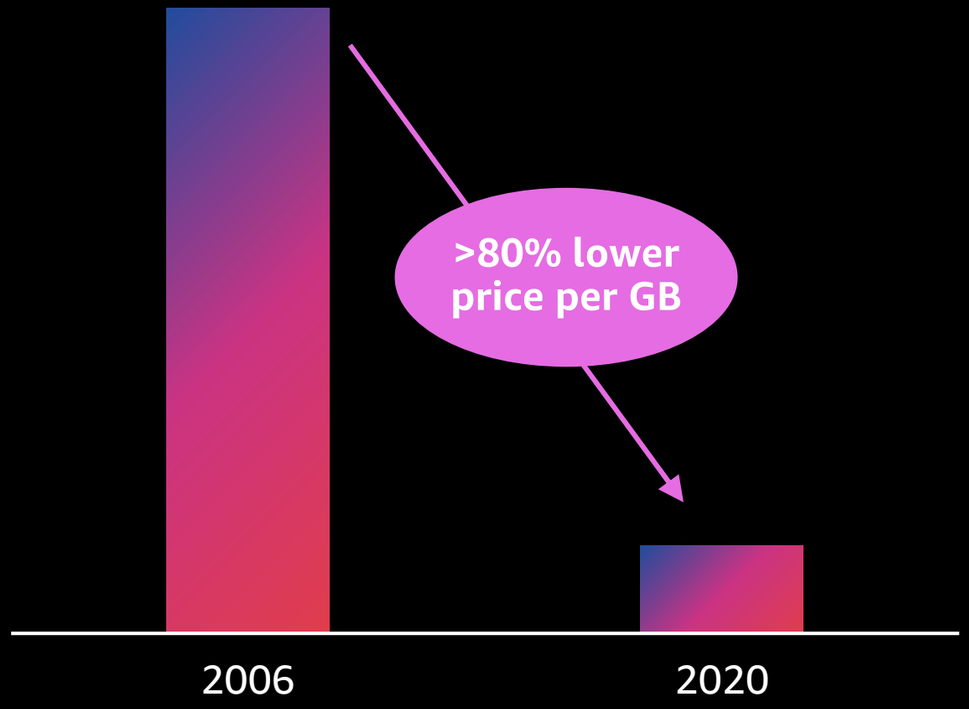
AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

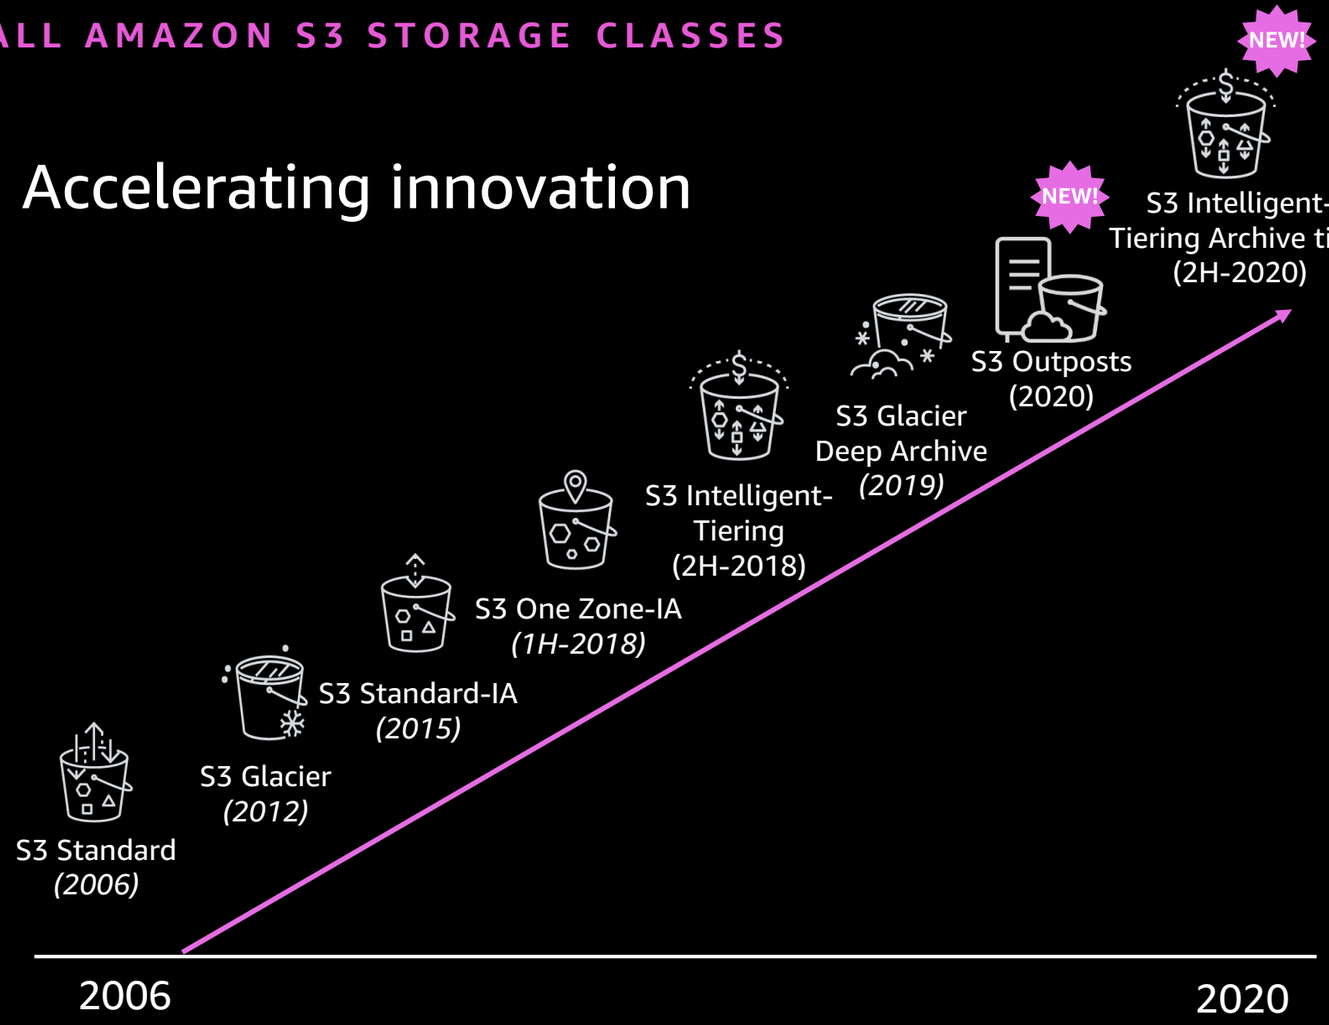
Amazon S3 storage classes

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

- Active, frequently accessed data
- Milliseconds access



S3 Intelligent-Tiering

- Data with changing access patterns
- Automatic archiving
- Milliseconds access
- Archive select hours
- No retrieval fees
- Object monitoring fee minimum storage duration



S3 Standard-IA

- Infrequently accessed data
- Milliseconds access
- Retrieval fee per GB
- Minimum storage duration
- Minimum object size



S3 Glacier

- Archive data
- In minutes and hours
- Retrieval fee per GB
- Minimum storage duration
- Minimum object size



S3 Glacier Deep Archive

- Long-term archive data
- Select hours
- Retrieval fee per GB
- Minimum storage duration
- Minimum object size

AWS Region ≥3 AZ



S3 One Zone-IA

- Re-creatable, less accessed data
- Milliseconds access
- Retrieval fee per GB
- Minimum storage duration
- Minimum object size

AWS AZ



S3 Outposts

- On-premises data
- Milliseconds access
- Encrypted with SSE-S3

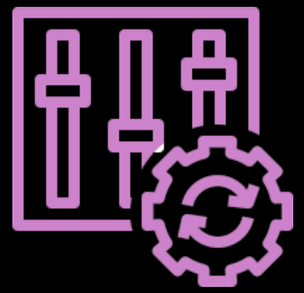
AWS Outposts

Amazon S3 analytics and insights

Amazon S3 Inventory



Amazon S3 Storage Class Analysis

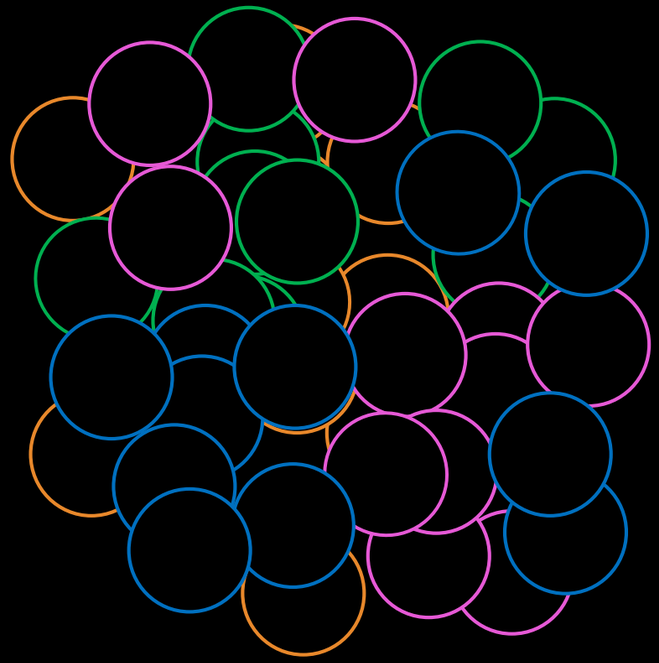


S3 Storage Lens



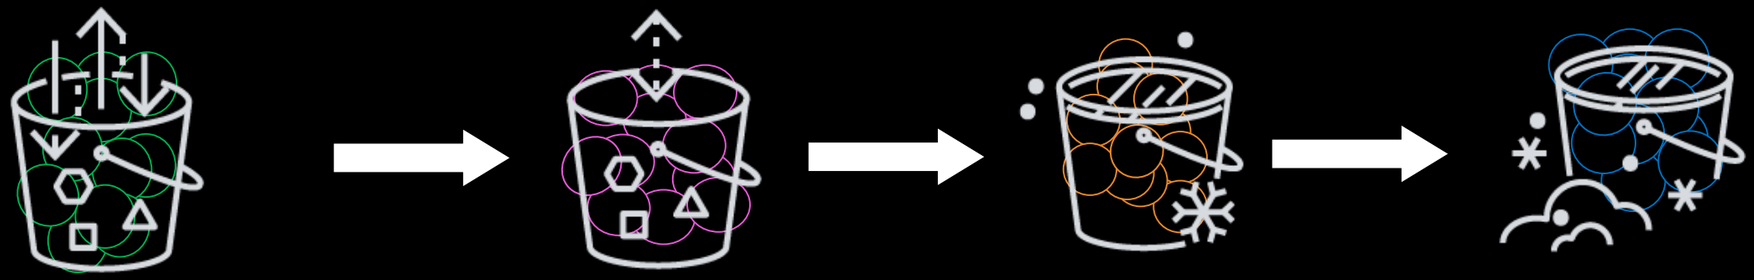
Learn more about your data access patterns

S3 ANALYTICS CAN MAKE STORAGE USAGE SIMPLE TO UNDERSTAND AND OPTIMIZED FOR COST



Learn more about your data access patterns

S3 ANALYTICS CAN MAKE STORAGE USAGE SIMPLE TO UNDERSTAND AND OPTIMIZED FOR COST



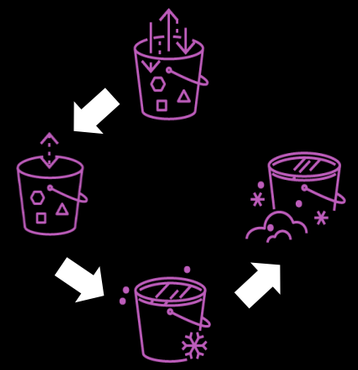
S3 Standard

S3 Standard-IA

S3 Glacier

S3 Glacier
Deep Archive

Choose option that works best for you



S3 Lifecycle to transition to different storage classes



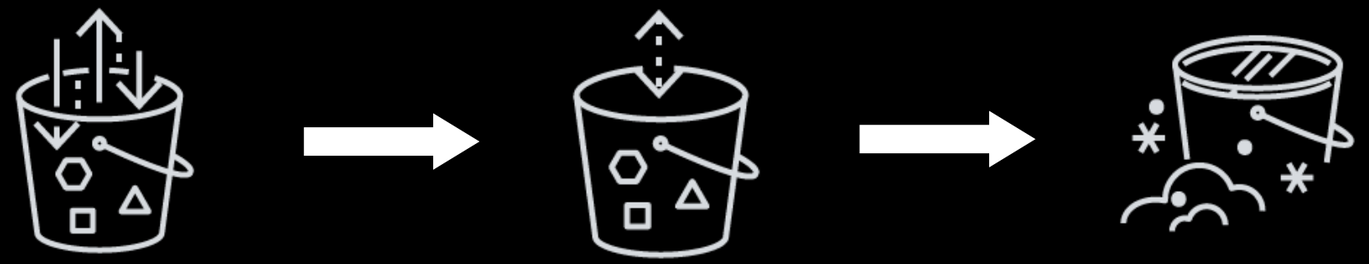
S3 Intelligent-Tiering with Frequent, Infrequent, Archive and Deep Archive access tiers

S3 Lifecycle policies use rules to manage storage

USE S3 LIFECYCLE POLICIES TO TRANSITION OBJECTS TO ANOTHER STORAGE CLASS

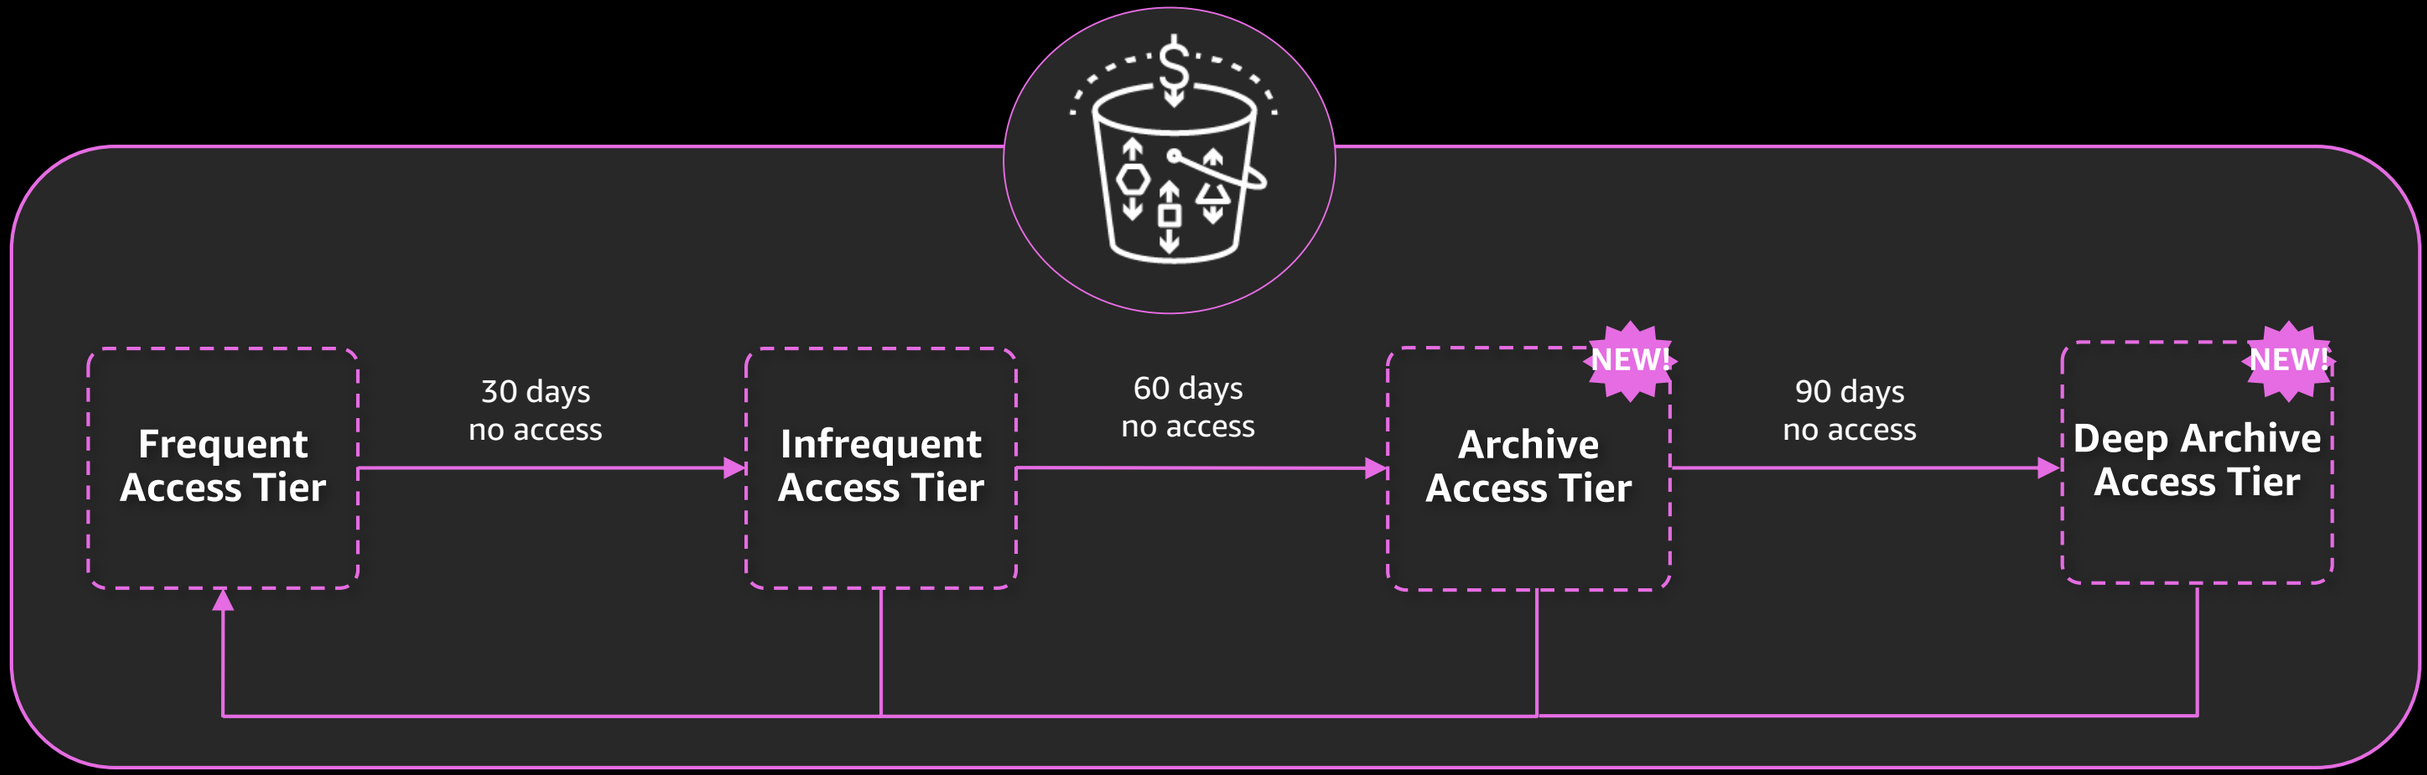
S3 Lifecycle rules take action based on object age; here's an example

1. Move objects older than 30 days to S3 Standard-IA
2. Move objects older than 365 days to S3 Glacier Deep Archive

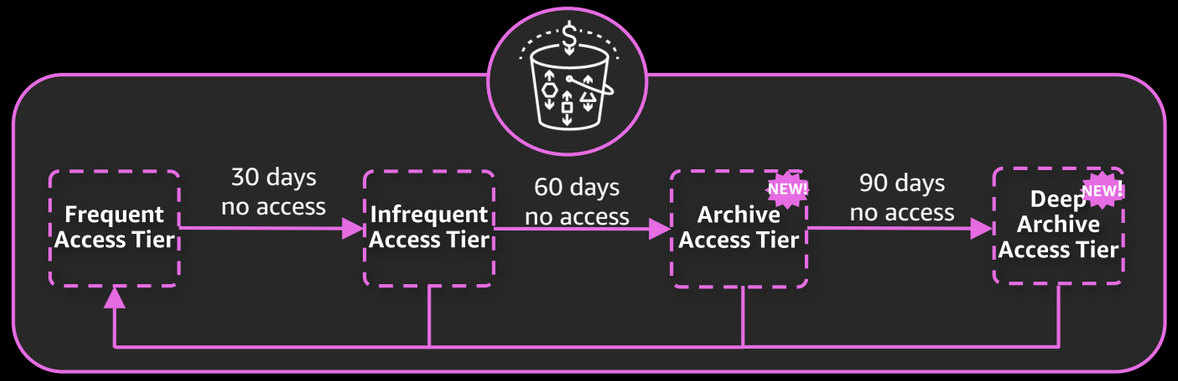


Amazon S3 Intelligent-Tiering – NEW!

AUTOMATIC COST OPTIMIZATION WITH NO PERFORMANCE IMPACT AND NO OPERATIONAL OVERHEAD



S3 Intelligent-Tiering automates cost savings



Automatically **optimizes storage costs** for data with changing access patterns

Activates **new automatic archive** capabilities for objects that become rarely accessed

Stores objects in **four access tiers**, optimized for frequent, infrequent, and rare access

Monitors access patterns and optimizes cost on granular object level

Has no operational overhead, no lifecycle fees, and no retrieval fees

Data protection



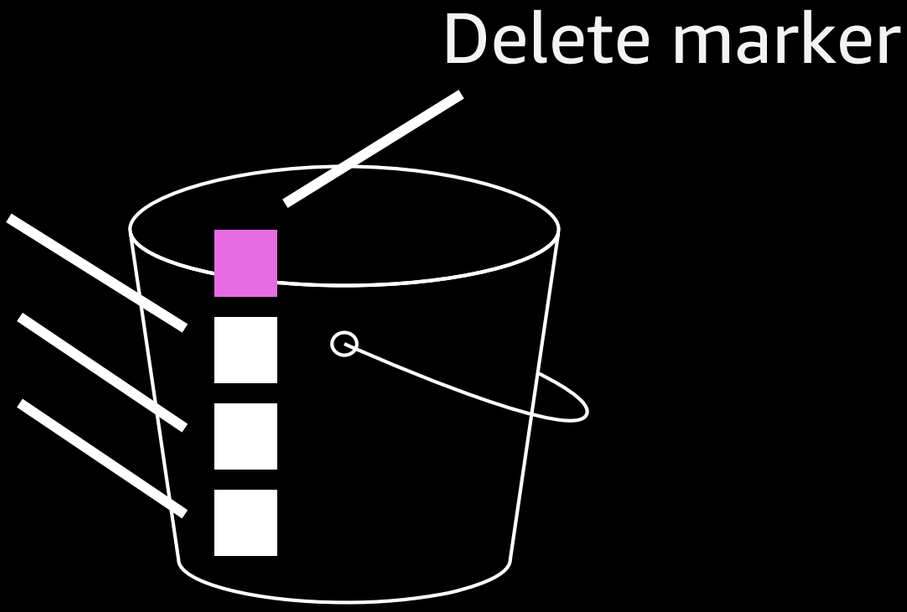
AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Use object versioning to protect your objects

VERSIONING PREVENTS OVERWRITES BY CREATING NEW VERSIONS OF THE OBJECT

DOC-EXAMPLE-BUCKET1 Version c
DOC-EXAMPLE-BUCKET2 Version b
DOC-EXAMPLE-BUCKET3 Version a

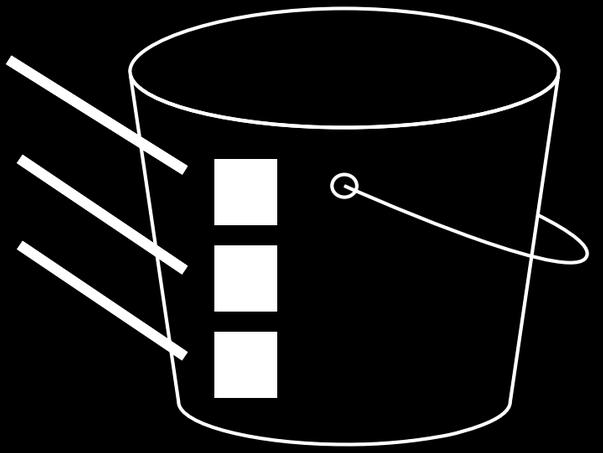


Deleting without a version ID

Use object versioning to protect your objects

VERSIONING PREVENTS OVERWRITES BY CREATING NEW VERSIONS OF THE OBJECT

~~DOC-EXAMPLE-BUCKET1 Version c~~
DOC-EXAMPLE-BUCKET2 Version b
DOC-EXAMPLE-BUCKET3 Version a



Deleting with a version ID

Amazon S3 Replication

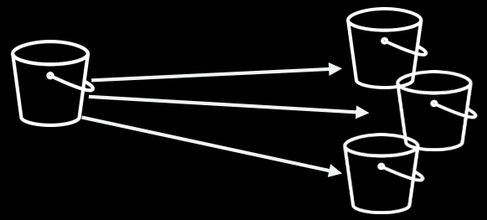
Replicate within the same AWS Region



Replicate to a different AWS Region



Replicate to more than one destination buckets





Replicate to a bucket with object lock retention controls



Replicate using Replication Time Control, for predictable replication, backed by an SLA



S3 Replication Time Control (S3 RTC)

PREDICTABLE REPLICATION TIME, BACKED BY SLA

- Designed to replicate 99.99% of objects within 15 minutes of upload
- Backed by SLA with commitment to replicate 99.9% of objects within 15 minutes
- Entire bucket or filtered set of objects



Amazon S3 Object Lock



Immutable S3 objects

Write-once-read-many (WORM) protection for S3 objects

Retention management controls

- Retention periods
- Legal holds
- Object Lock works only in versioned buckets

Bucket level setting or object level

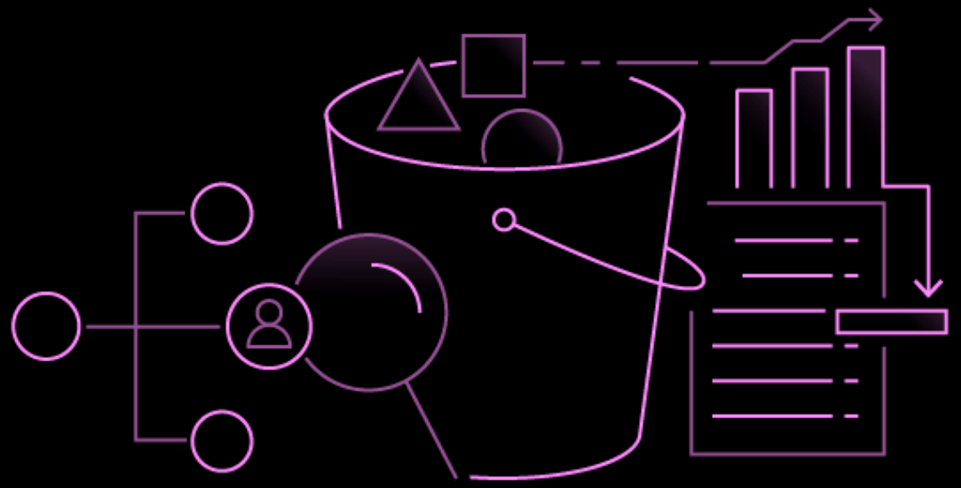
- Can be set up per object
- Default bucket level setting

Analytics and insights



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Introducing S3 Storage Lens – NEW!



Metrics – 15 usage metrics, 14 activity metrics (29 total)

Aggregation levels – organization, account, Region, storage class, bucket, and prefix

Refresh – daily updates (through previous day)

Data retention – up to 15 months

Visualization – interactive dashboard in Amazon S3 console (export to S3 bucket also available)

Recommendations – metrics-driven contextual recommendations to highlight best practices

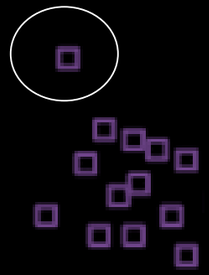
S3 Storage Lens use cases

Summary insights



How much has my storage increased across my entire organization in the past 90 days?

Outliers



Are there any unusual trends or spikes in my storage usage or activity?

Data protection



Are all of my buckets encrypted?

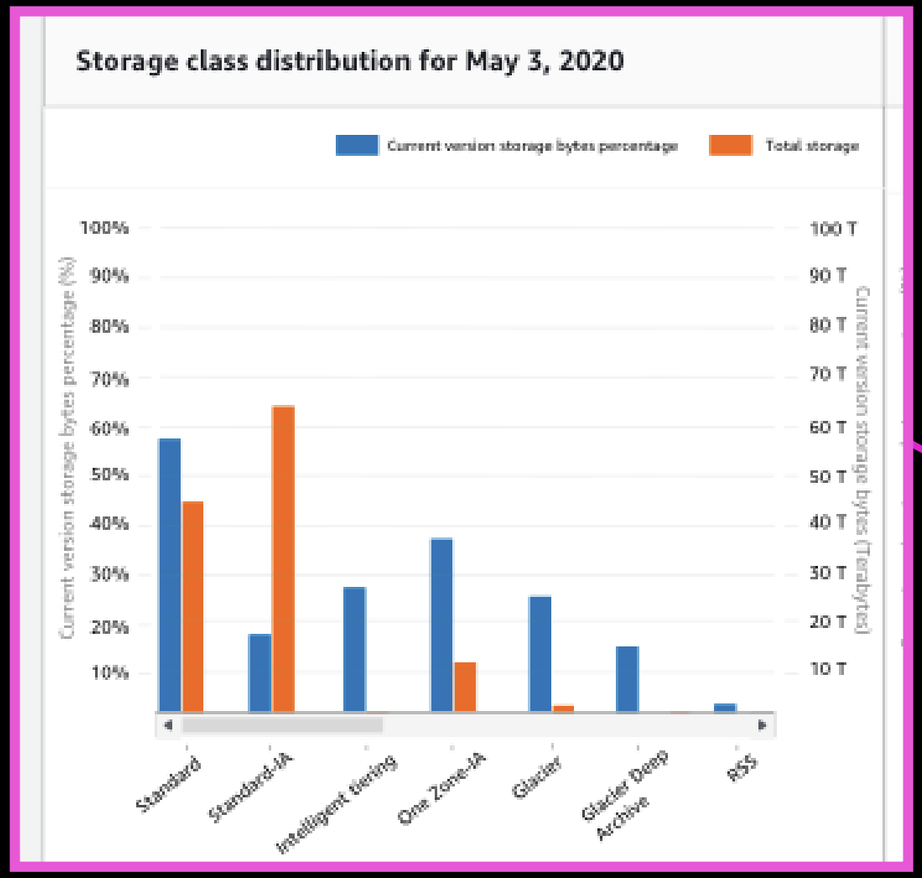
Cost efficiency

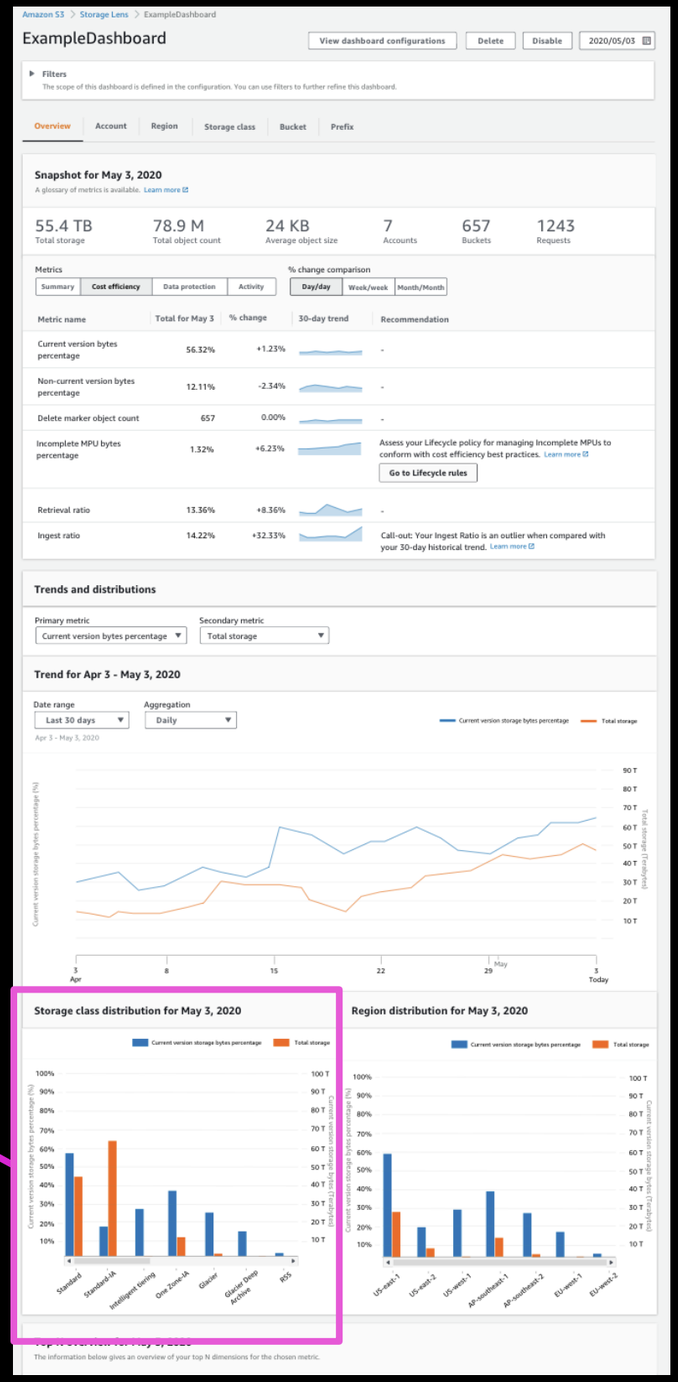


Which of my buckets are infrequently accessed?

Example insights from S3 Storage Lens

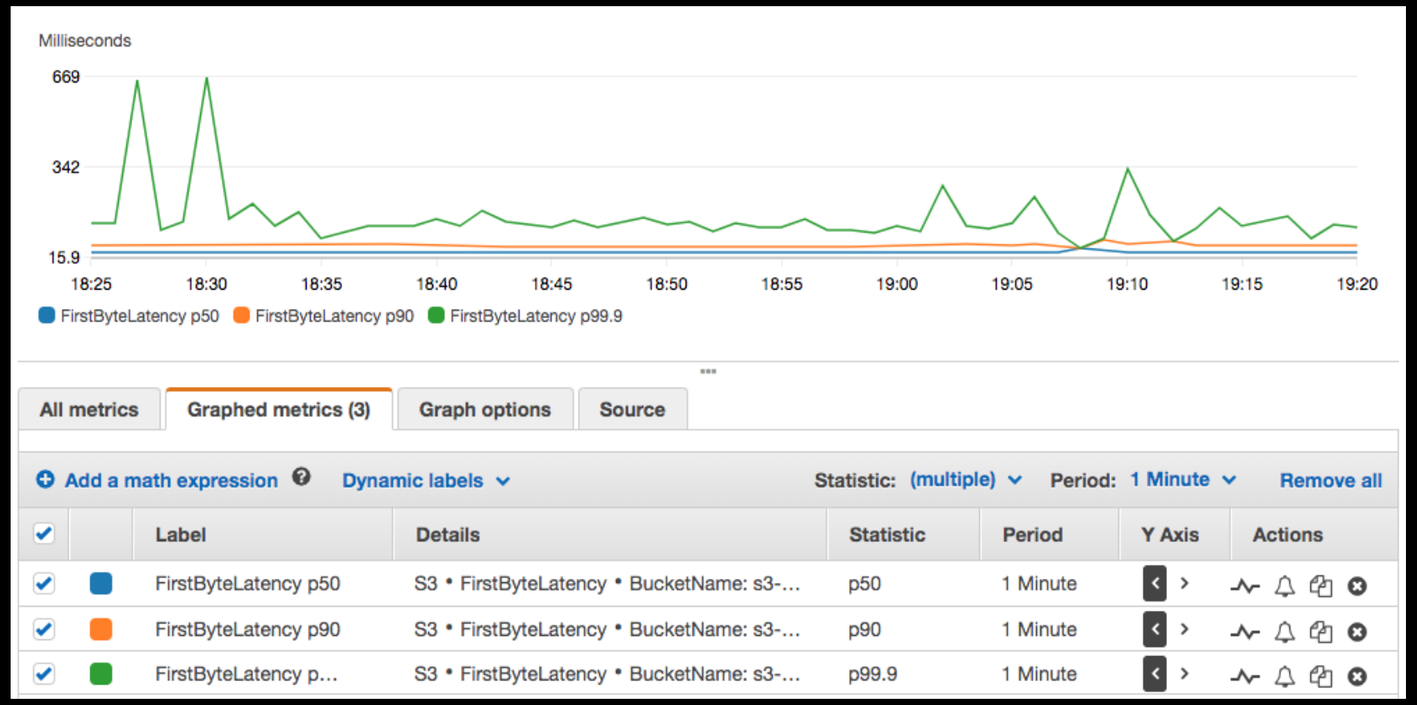
Understand your organization-wide storage allocation by storage class





Amazon CloudWatch metrics for Amazon S3

AMAZON S3 REQUEST METRICS ON ANY PERCENTILE (E.G., P90, P99, P99.9, P100)



- Understand distribution of Amazon S3 request metrics
- Visualize and alarm on any percentile to identify outliers or unusual application behavior
- Avoid false alarms and save time spent monitoring and tracking requests

Customer spotlight – How Airbnb uses Amazon S3 for cost optimization and data protection



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon S3 use cases at Airbnb



User-generated content

Photos of homes, videos of experiences, message attachments



Data warehouse

Storage backend for compute engines, such as Apache Spark and Apache Hive



Search

Storage for search indexes



Logs

AWS CloudTrail logs, CDN access logs, and Amazon S3 bucket access logs



Backup

Storing backups of data stores

Cost optimization at Airbnb



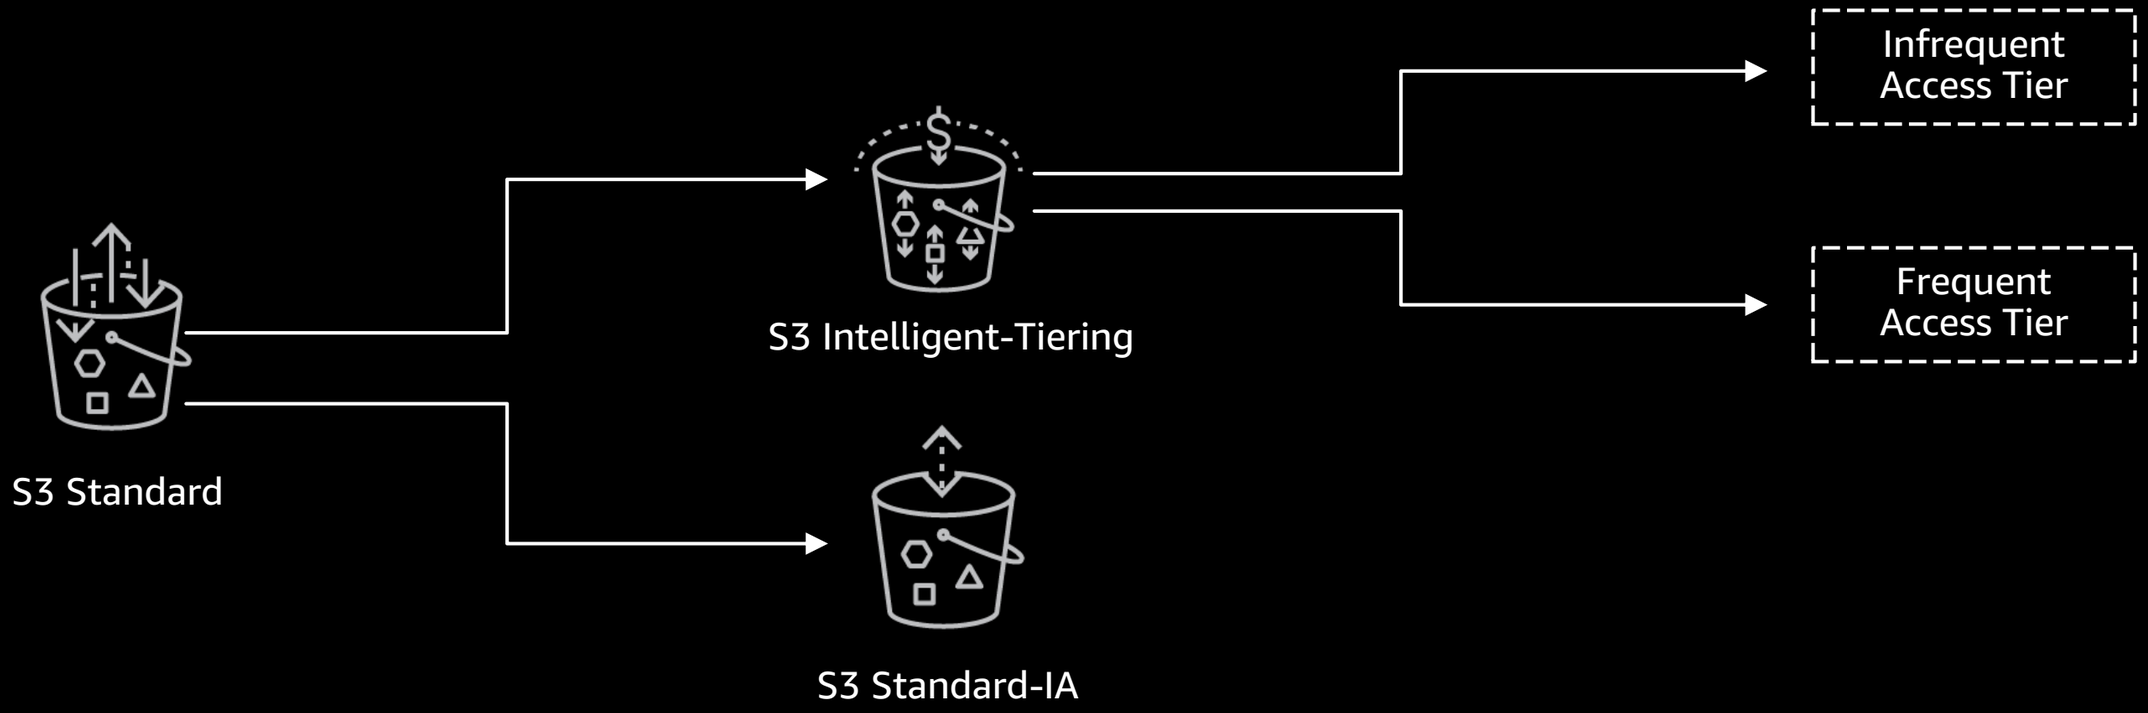
AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

We leverage S3 storage classes to optimize cost

S3 storage classes for production data

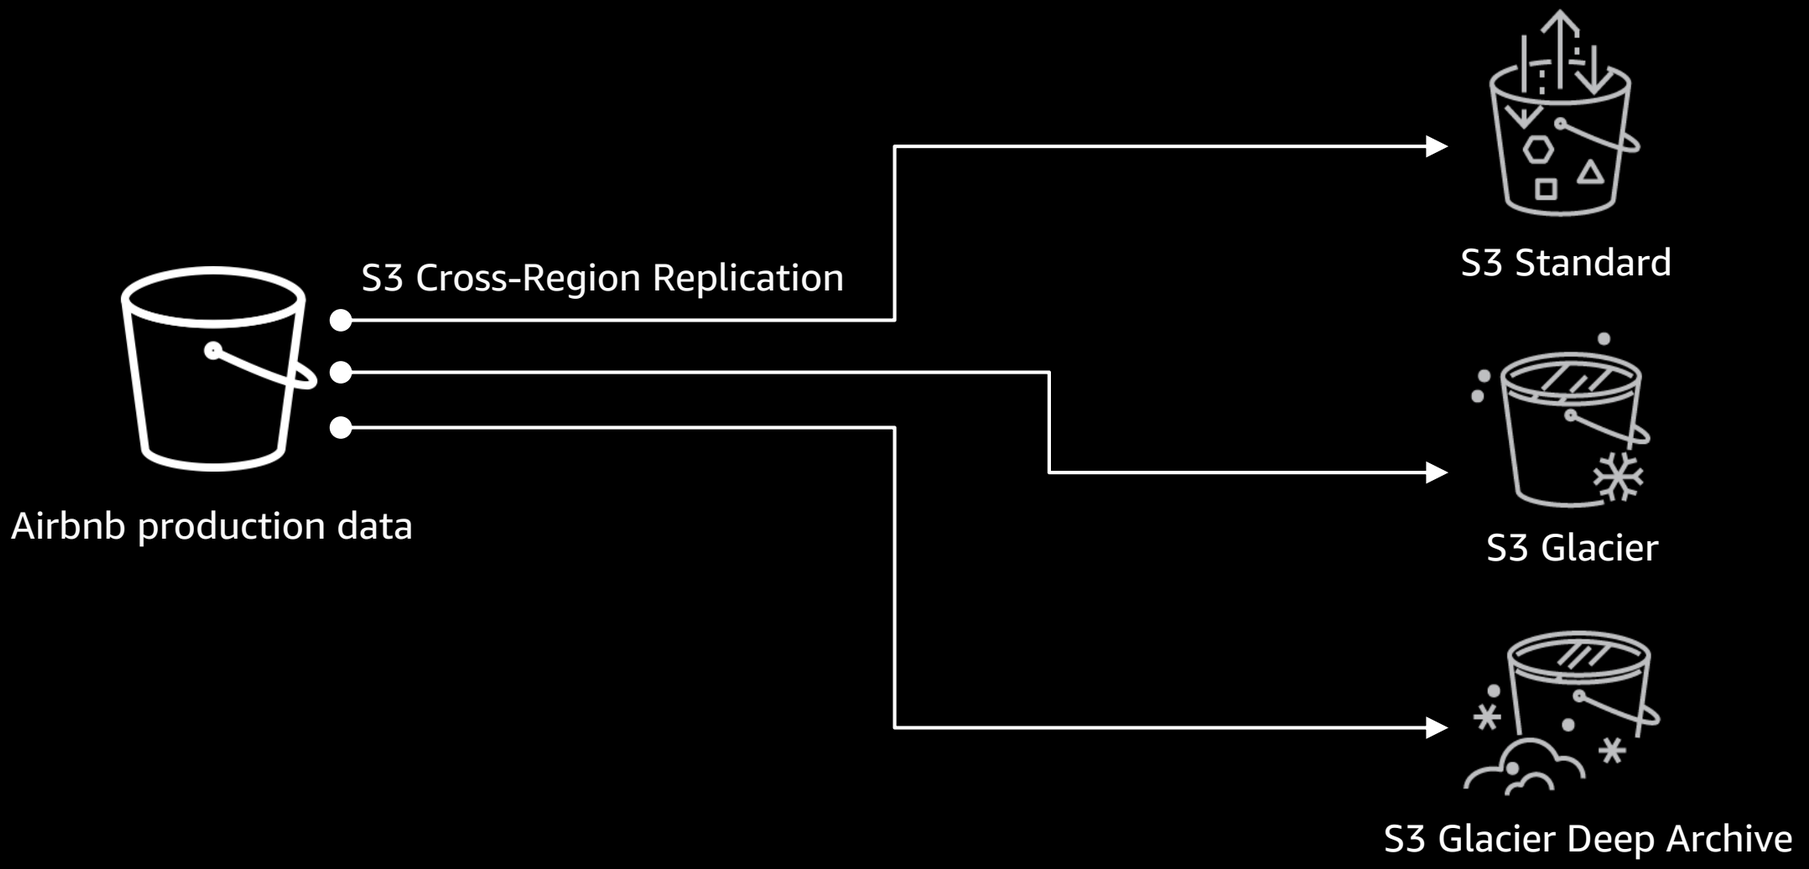
Most objects are in S3 Standard, but we also use S3 Lifecycle to transition to S3 Intelligent-Tiering and S3 Standard-IA



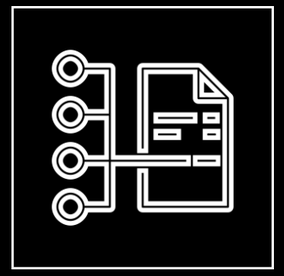
We leverage S3 storage classes to optimize cost

S3 storage classes for backup

Objects are in S3 Standard, S3 Glacier, and S3 Glacier Deep Archive



We use S3 tools to further optimize costs



AWS Cost & Usage Report

- Generated hourly
- Tracking storage costs per bucket



Amazon S3 Inventory

- Generated weekly on our S3 buckets
- Tracking storage costs per object-level prefix



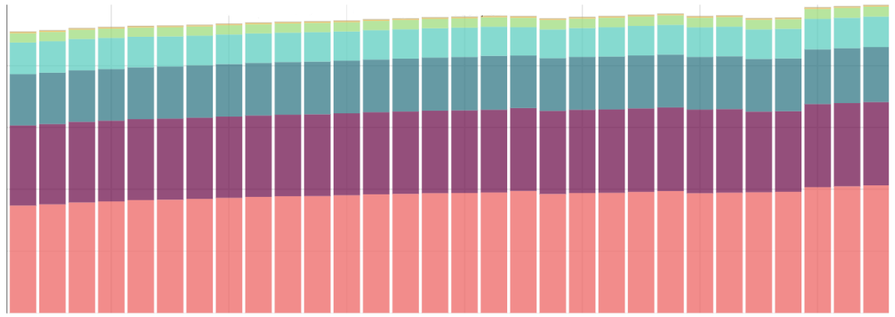
Amazon S3 Server Access Logs

- Enabled on all our S3 buckets
- Helps identify unused data
- Replaced some of our past use of AWS CloudTrail

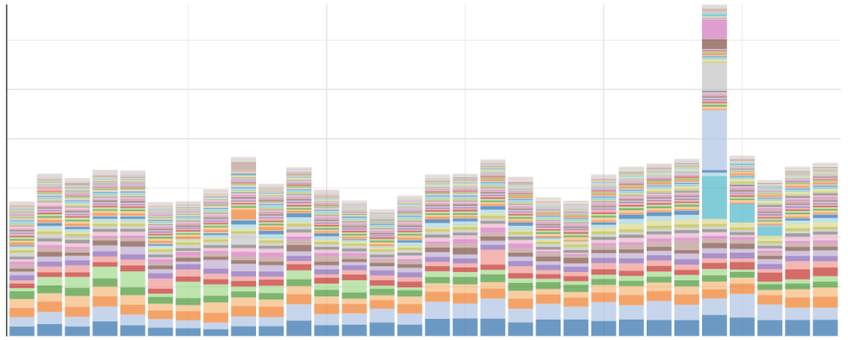
AWS Cost & Usage Report

Airbnb S3 Dashboard generated from AWS Cost and Usage Report

[S3 Storage] Cost Growth by Storage Class, Trailing 90 Days

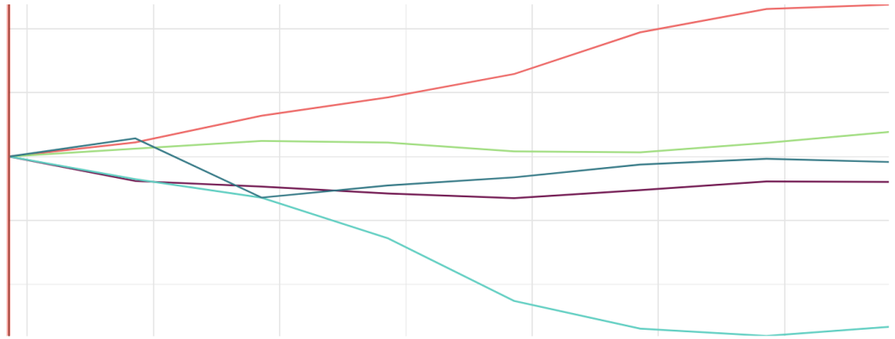


[S3 API Request] Daily Cost by Bucket

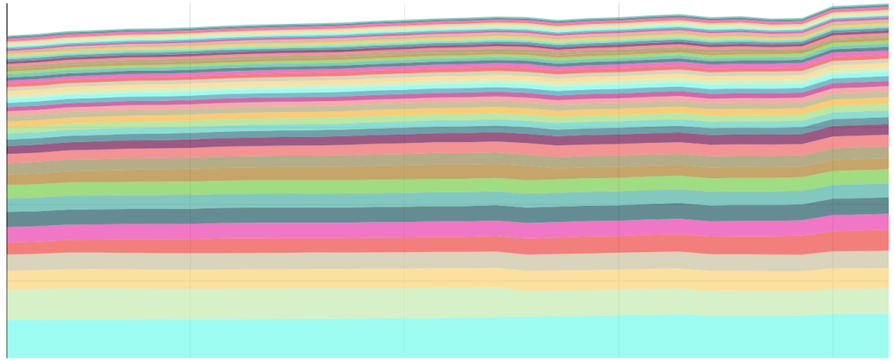


API requests look high for that day

[S3 Storage] Usage Growth by Storage Class, Trailing 2 Months



[S3 Storage] Cost Growth by Bucket, Trailing 90 Days (Top Buckets Only)



We use Amazon S3 Inventory to analyze storage use by prefix

Example query used in a past data cleanup effort

```
SELECT
  SUM(size) / 1073741824 / 1024.0 AS total_size_tb,
  split(key,'/')[1] || '/' || split(key,'/')[2]  AS first_dirs,
  bucket
from aws_cost.s3_inventory
WHERE (bucket = 'xxxx' OR bucket = 'yyyy')
  AND ds = '2020-04-01'
  AND cardinality(split(key,'/')) > 1
GROUP BY 2,3
ORDER BY 1 DESC
```

We use S3 Server Access Logs to identify large data sets that are no longer used

Example query used to check if older data was still being accessed; the results allowed us to set a S3 Lifecycle policy to expire data after 30 days for that prefix

```
SELECT
  inventory.key,
  inventory.last_modified_date,
  *
FROM
  aws_cost.s3_access_log AS accesses
JOIN
  aws_cost.s3_inventory AS inventory ON accesses.key = inventory.key
WHERE
    (inventory.bucket = 'xxxx' OR inventory.bucket = 'yyyy')
    AND inventory.ds = '2020-06-20'
    AND accesses.ds > '2020-06-01'
    AND strpos(inventory.key, 'titanic/monorail/artifacts') != 0
    AND operation != 'S3.TRANSITION_INT.OBJECT'
    AND operation != 'S3.TRANSITION_SIA.OBJECT'
ORDER BY
  inventory.last_modified_date ASC
LIMIT 100;
```

Data protection at Airbnb



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

We replicate data across AWS Regions for data protection

S3 Cross-Region and Cross-Account Replication to protect and isolate data

Example Terraform configuration

```
replication_configuration {
  role = "arn:aws:iam::${data.aws_caller_identity.current.account_id}:role/ReplicateS3Backups"

  rules {
    prefix = ""
    status = "Enabled"

    destination {
      bucket        = var.replica-bucket-arn
      storage_class = "DEEP_ARCHIVE"
    }
  }
}
```



Production data
Account A
AWS US-East

Backup data
Account B
AWS US-West

We set S3 Lifecycle rules to expire noncurrent versions

- Most times we can quickly recover deleted data from production buckets
- Backup buckets have longer retention, but data is mostly in S3 Glacier and S3 Glacier Deep Archive, so recovery takes longer

Example production bucket

```
lifecycle_rule {

  enabled = true

  abort_incomplete_multipart_upload_days = 2

  noncurrent_version_expiration {
    days = 14
  }

  expiration {
    expired_object_delete_marker = true
  }
}
```

Example backup bucket

```
lifecycle_rule {

  enabled = true

  abort_incomplete_multipart_upload_days = 5

  noncurrent_version_expiration {
    days = 60
  }

  expiration {
    expired_object_delete_marker = true
  }
}
```

We use S3 Inventory and S3 Batch Operations to complement replication



S3 Inventory

Ingested into our data warehouse, where we can query it using Apache Hive

We've experienced failed replication when the replication configuration got misconfigured

Example query

```
SELECT bucket,key FROM infra_disaster_recovery.s3_inventory WHERE dt = '2020-10-25-00-00' AND bucket = 'xxxxxxx' AND replication_status = 'FAILED' AND is_latest = true AND is_delete_marker = false
```



S3 Batch Operations

Use CSV output of query as input to S3 Batch Operations

Copy objects in place to re-drive replication

Amazon S3 for Airbnb

Cost optimization

Storage classes

S3 Lifecycle

Data protection

Versioning

Replication

Security



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Security is at the core of everything we do

DATA STORED IN AMAZON S3 IS SECURE BY DEFAULT



Amazon S3
Block Public Access



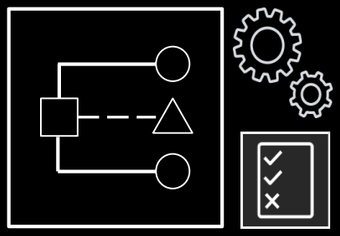
Encrypt data by
default in Amazon S3



Bucket permission
checks



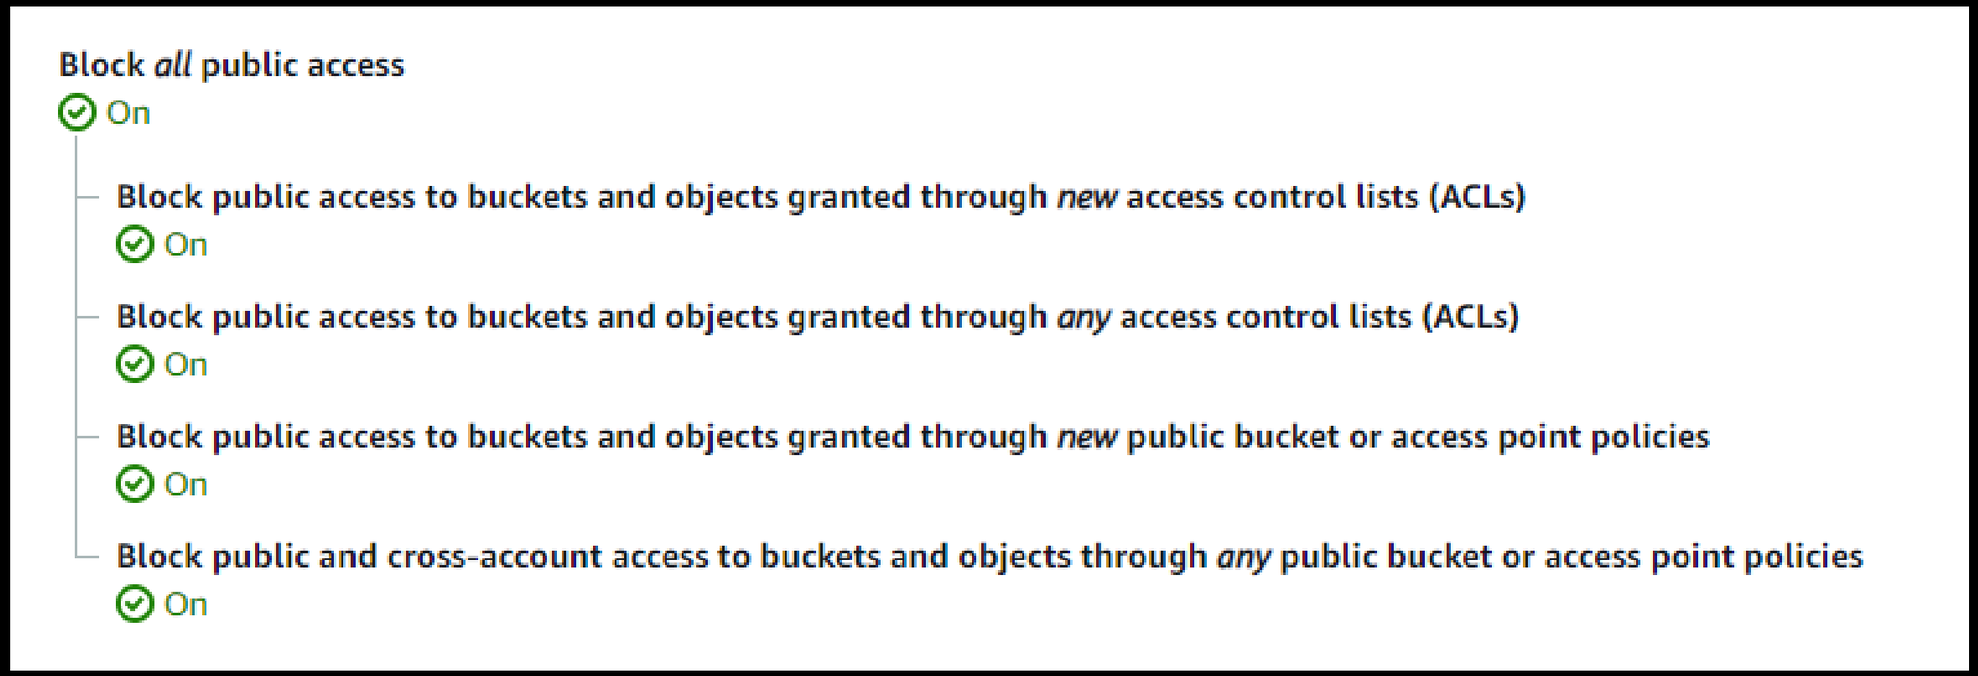
S3 Access
Points



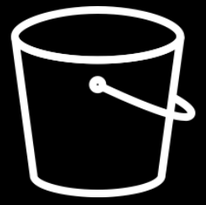
Access Analyzer

Amazon S3 Block Public Access settings

- Applies blanket protection against accidental public access
- Set at the bucket or the account level
- Can be applied to access control list (ACL) access, bucket policy access, or both



AWS S3 default encryption



One-time
bucket-level setup



Automatically encrypts
all new objects

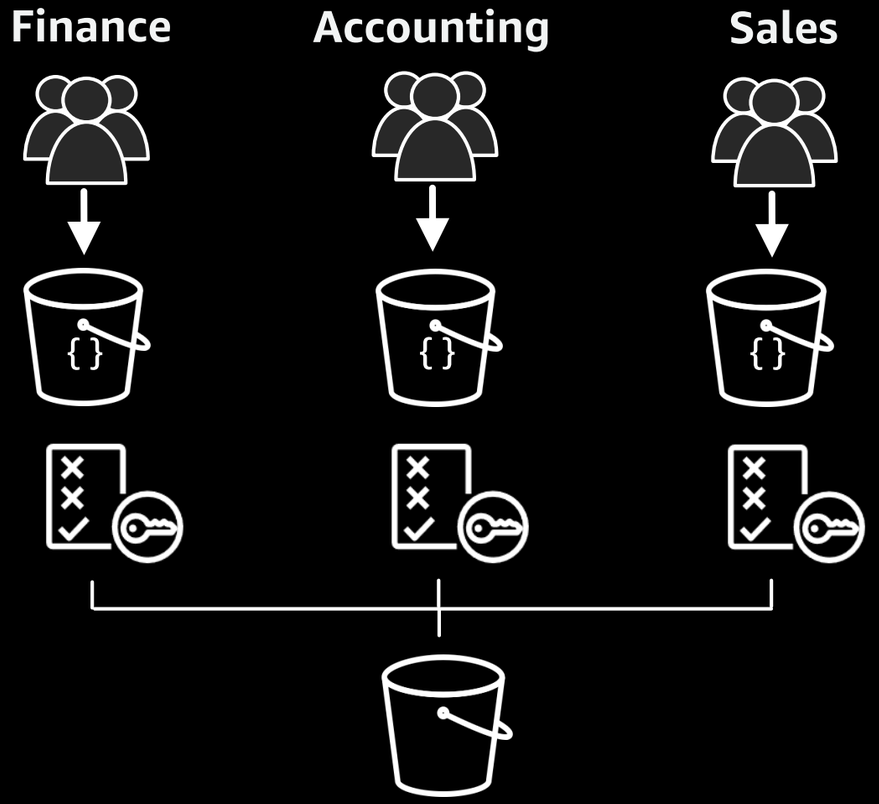


Simplified compliance



Supports SSE-S3
and SSE-KMS

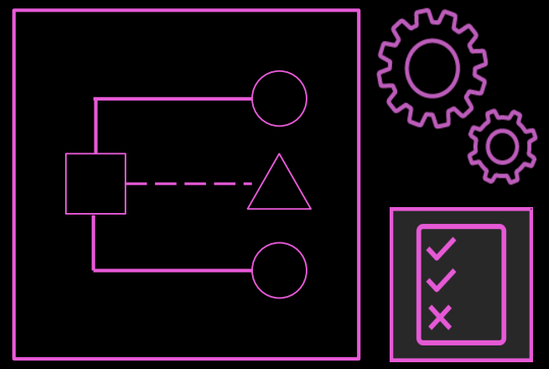
Amazon S3 Access Points



- Segment clients into distinct groups
- Each group gets its own access point
- Apply policy to each access point
- Maintain centralized control over storage

Access Analyzer for Amazon S3 buckets

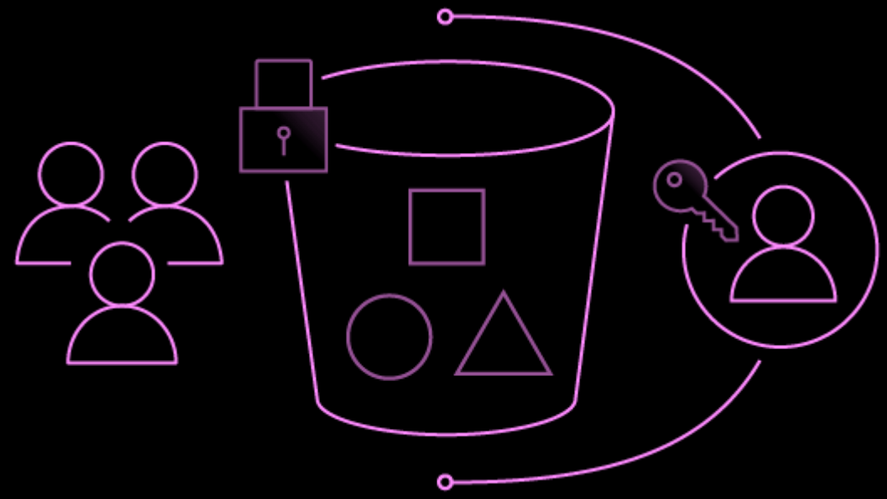
- Continuously monitors and automatically analyzes resources
- Surfaces buckets with public and shared access in AWS Management Console
- Remove public access with a single click
- Acknowledge shared access as intended



Amazon S3 Object Ownership – NEW!

AUTOMATICALLY TRANSFER OWNERSHIP OF OBJECTS TO THE BUCKET OWNER FOR CROSS-ACCOUNT WRITES

- Applies to cross-account writes with bucket-owner-full-control ACL
- Each supports other AWS services delivering data to S3 buckets
- Ensures uniform ownership of objects in bucket
- Enables bucket owners to manage access to objects through bucket policy or access point policy



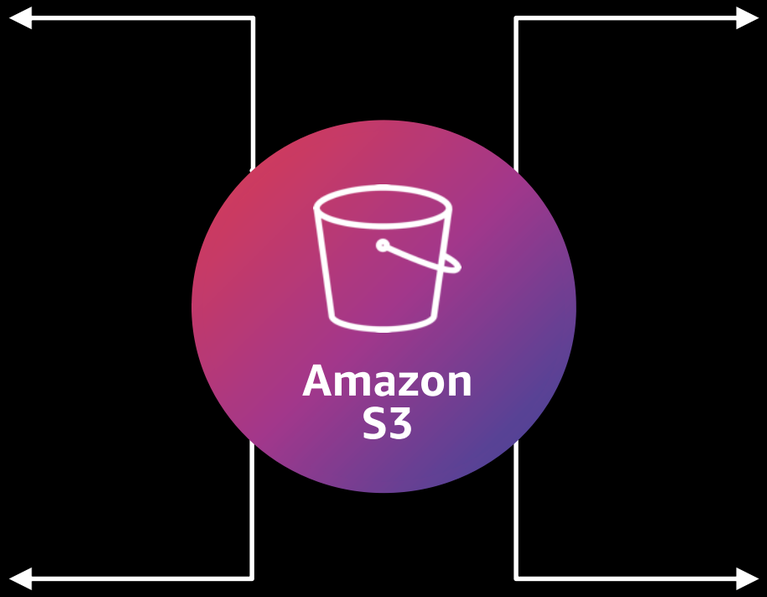
Amazon S3 Best Practices – Quick Recap

Storage classes and cost optimization

- S3 Intelligent-Tiering
- S3 Lifecycle

Security

- S3 Block Public Access
- S3 Encryption
- S3 Access Points
- S3 Object Ownership



Data protection

- S3 Versioning
- S3 Replication

Analytics and insights

- S3 Storage Lens
- Amazon CloudWatch

Thank you!

Bijeta Chakraborty

Sr. Product Manager, Amazon S3
AWS

cbijeta@amazon.com

Naveen Nalam

Software Engineer, Infrastructure
Airbnb

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Please complete the session survey

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.