
AWS
re:Invent

STG201

What's new with Amazon S3

Christoph Bartenstein

Director, Product Management
Amazon S3

AWS re:Invent

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

Agenda

Welcome

Customer-focused innovation

Innovate faster with Amazon S3

What's new

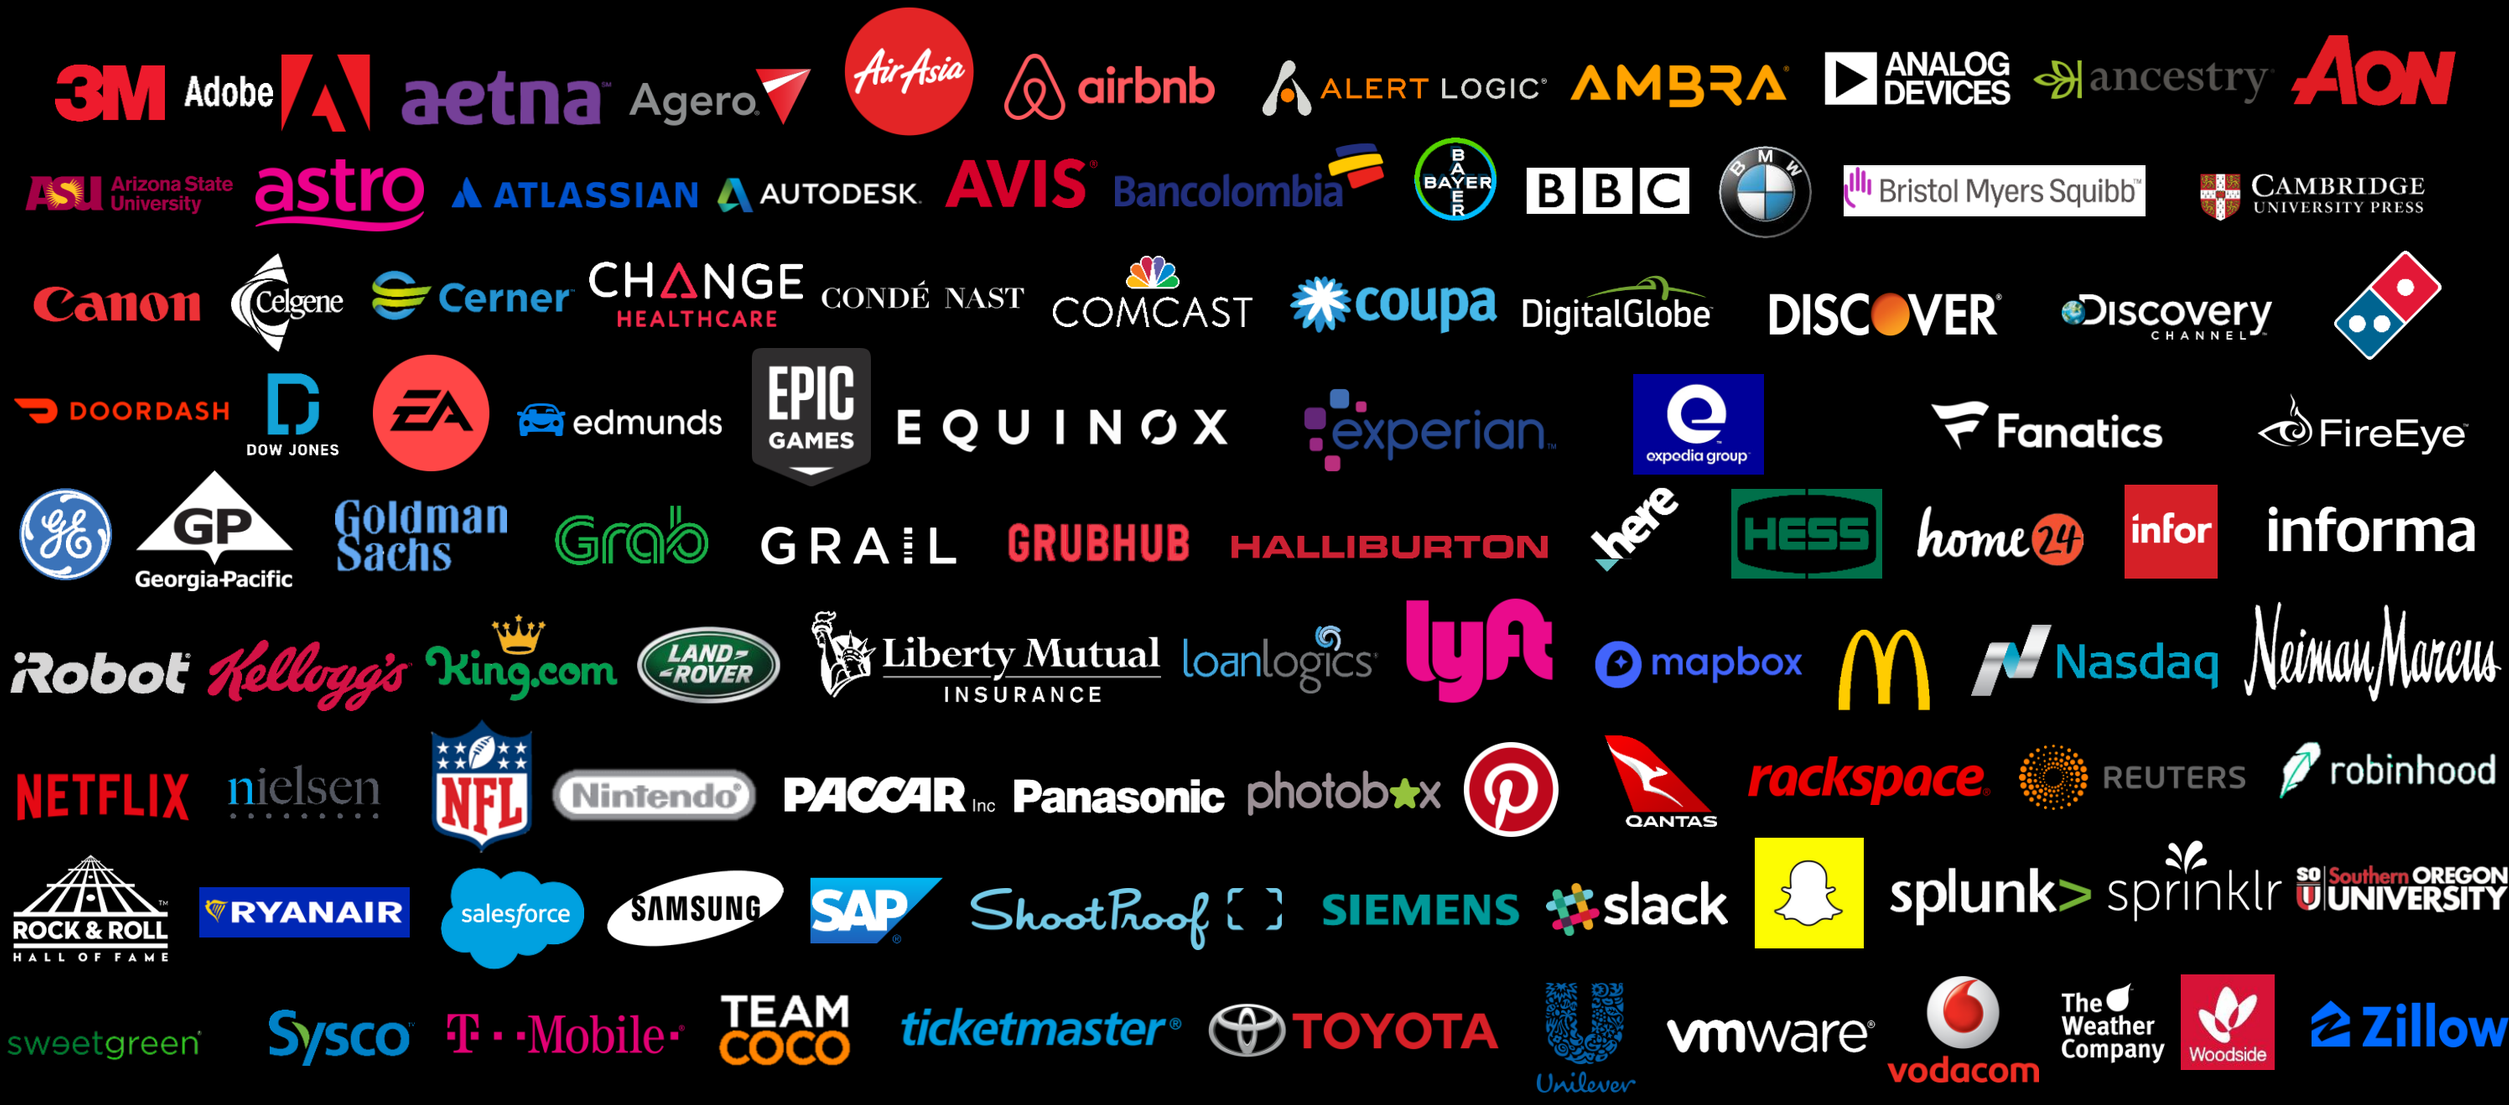
Putting it all together



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

Customer-focused innovation

OUR FEATURE ROADMAP COMES DIRECTLY FROM CUSTOMER FEEDBACK



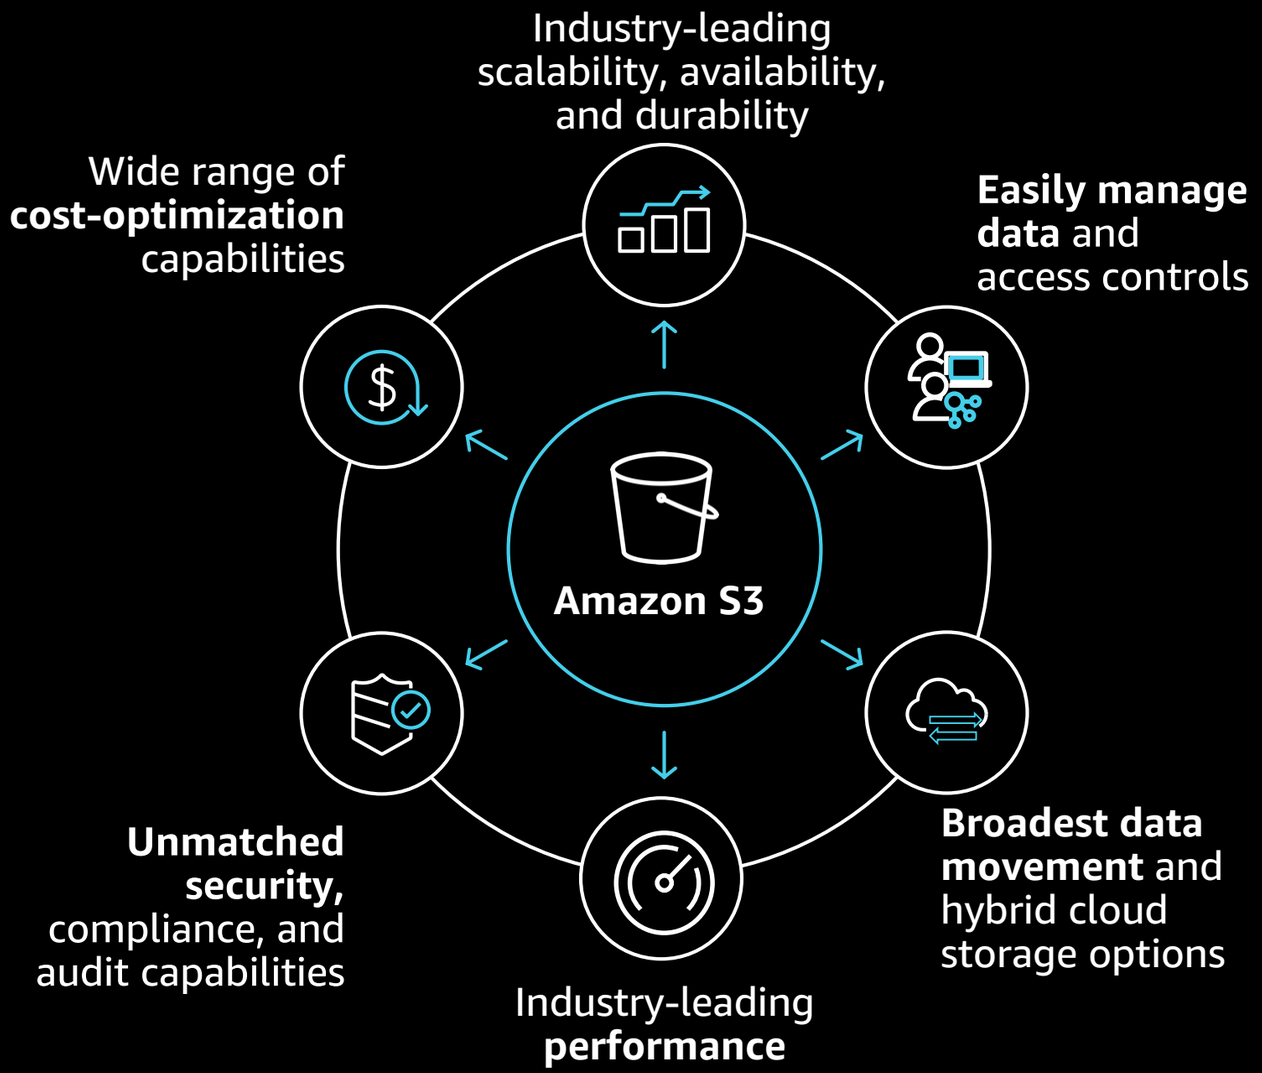


AWS recognized as a cloud leader for the 10th consecutive year

Gartner, Magic Quadrant for Cloud Infrastructure & Platform Services, Raj Bala, Bob Gill, Dennis Smith, David Wright, Kevin Ji, 1 September 2020. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose. The Gartner logo is a trademark and service mark of Gartner, Inc., and/or its affiliates, and is used herein with permission. All rights reserved.

Innovate faster with Amazon S3

STORAGE INNOVATION DRIVES APPLICATION INNOVATION



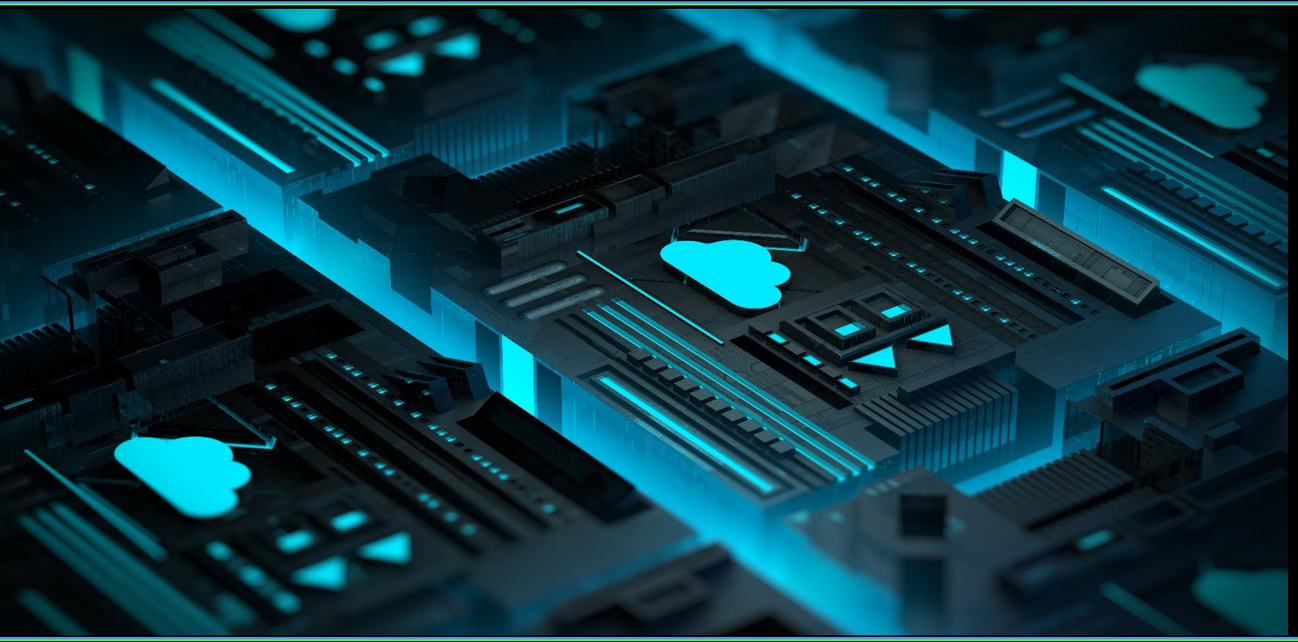
Exabytes of storage and trillions of objects

Tens of thousands of data lakes built on Amazon Simple Storage Service (Amazon S3)

Regularly peak at **millions of requests** per second

S3 Intelligent-Tiering is the only storage class that **automatically saves up to 95% on storage costs**

S3 Glacier Deep Archive offers the **lowest-cost cloud object storage**, at $0.00099 per GB-month



Innovation in Amazon S3

- Cost
- Performance
- Data movement
- Storage management
- Security
- Scalability, availability

Amazon S3 storage classes

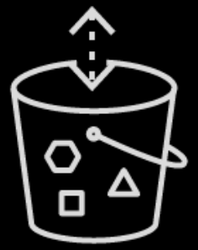
OPTIMIZE YOUR STORAGE COST BY UTILIZING ALL AMAZON S3 STORAGE CLASSES



S3 Standard



S3 Intelligent-Tiering



S3 Standard-IA



S3 Glacier



S3 Glacier Deep Archive



S3 One Zone-IA

NEW!



S3 Outposts

AWS Region ≥ 3 AZ

- Active, frequently accessed data
- Milliseconds access

- Data with changing access patterns
- Milliseconds access
- No retrieval fees
- Object monitoring fee minimum storage duration
- Minimum storage duration
- Opt in for Automatic Archiving

NEW!

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

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

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

AWS AZ

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

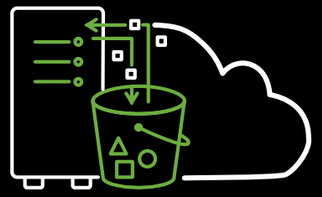
AWS Outposts

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

NEW

Amazon S3 on Outposts

NEW STORAGE CLASS IDEAL FOR DATA RESIDENCY REQUIREMENTS AND LOCAL PERFORMANCE NEEDS



- Object storage on AWS Outposts using the S3 APIs
- Designed to store durably and redundantly across multiple devices and servers
- You can add 48 TB or 96 TB of S3 capacity to an Outpost
- Encrypted by default and S3 security and access control features
- Transfer data between your Outposts bucket to an S3 bucket in an AWS Region by using AWS DataSync

NEW

Amazon S3 Intelligent-Tiering

FIRST AND ONLY CLOUD STORAGE SOLUTION TO PROVIDE DYNAMIC PRICING AUTOMATICALLY BASED ON THE CHANGING ACCESS PATTERNS OF INDIVIDUAL OBJECTS



Only cloud storage that delivers **automated cost savings** by optimizing costs at a granular object level

Moves objects between **four access tiers** for a small monthly monitoring and automation fee

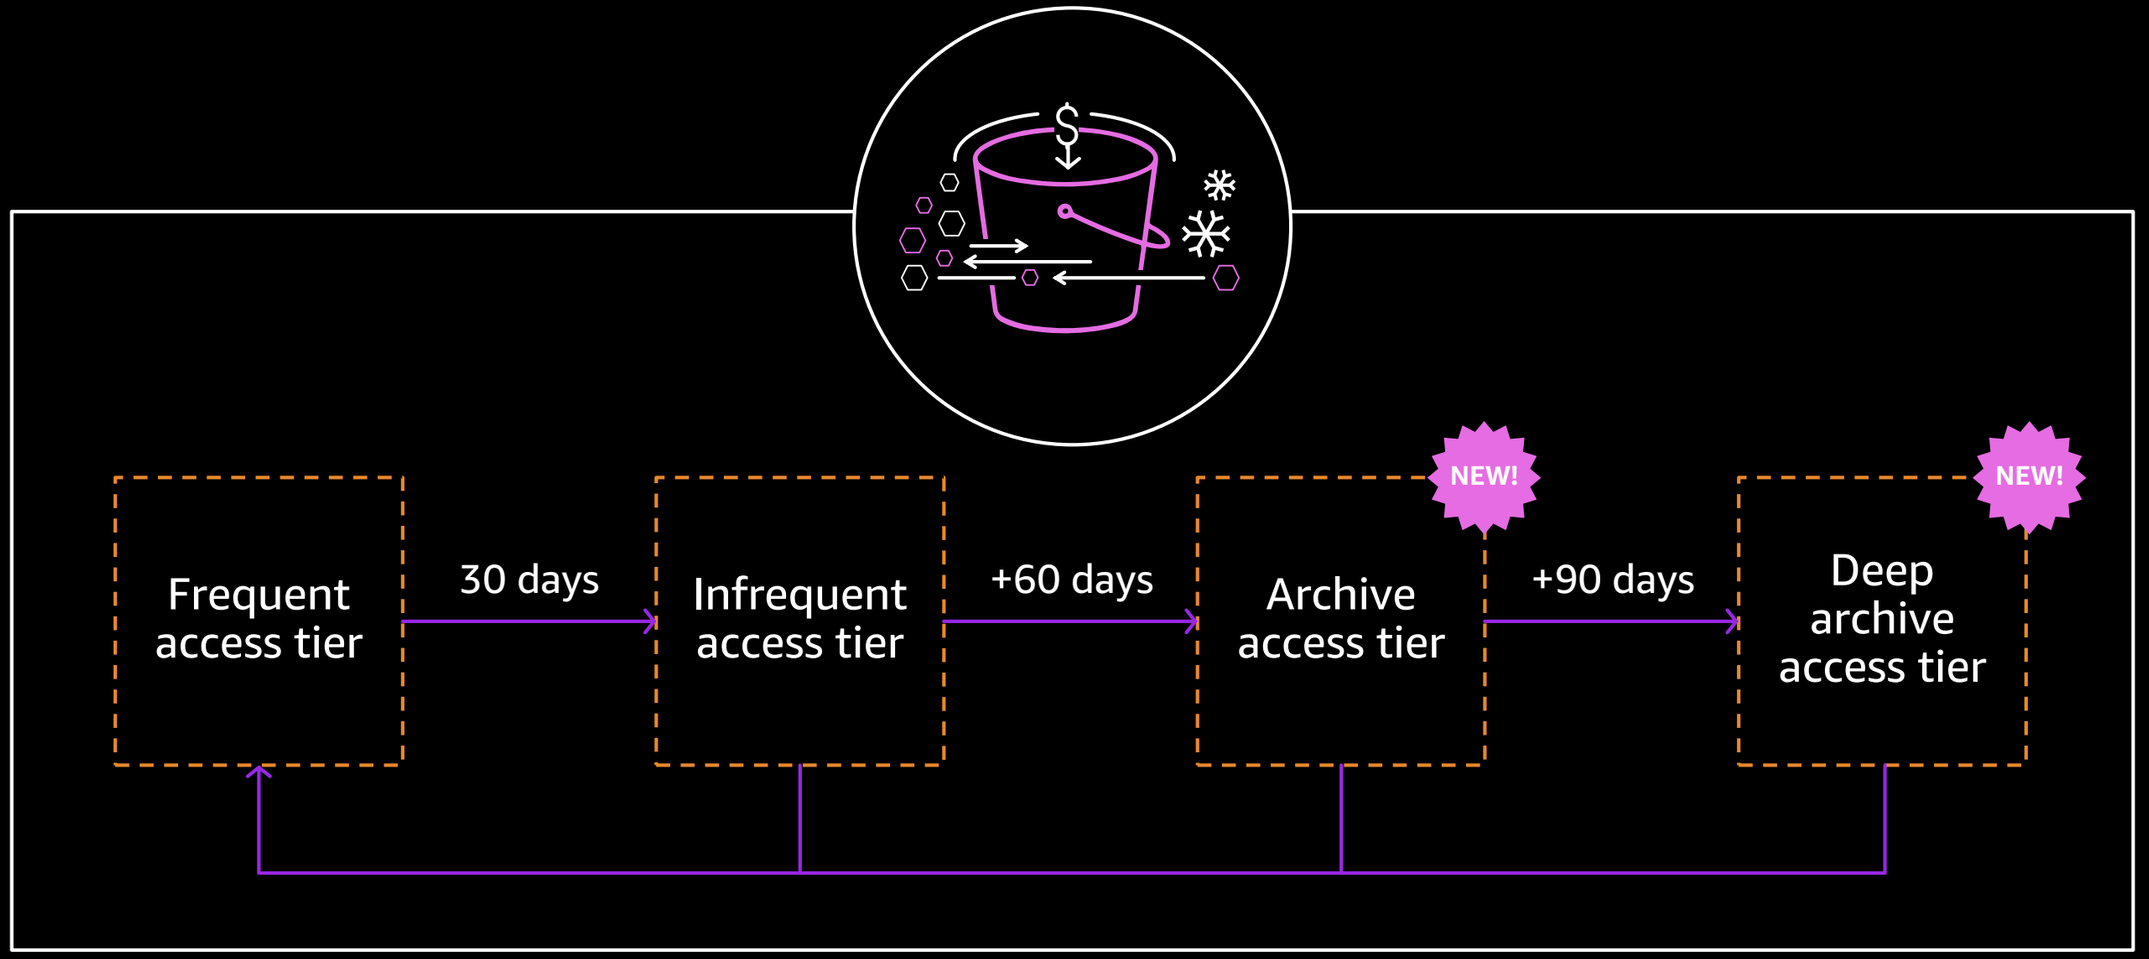
Two low latency access tiers for frequent and infrequent access that **save up to 40%,** and two **new optional archive access tiers** for access in minutes to hours **that save up to 95%** on storage costs

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

Designed for **99.9% availability** and 11 9's of durability

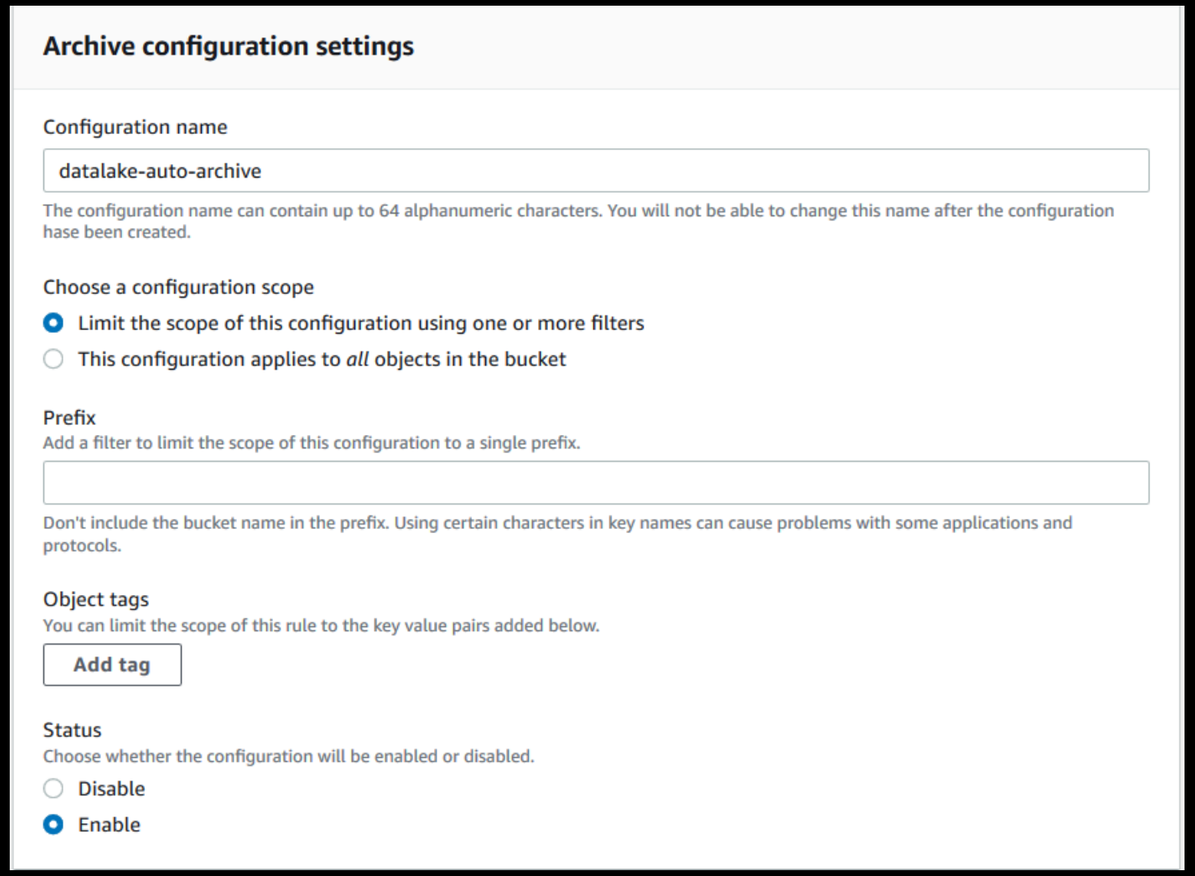
Amazon S3 Intelligent-Tiering

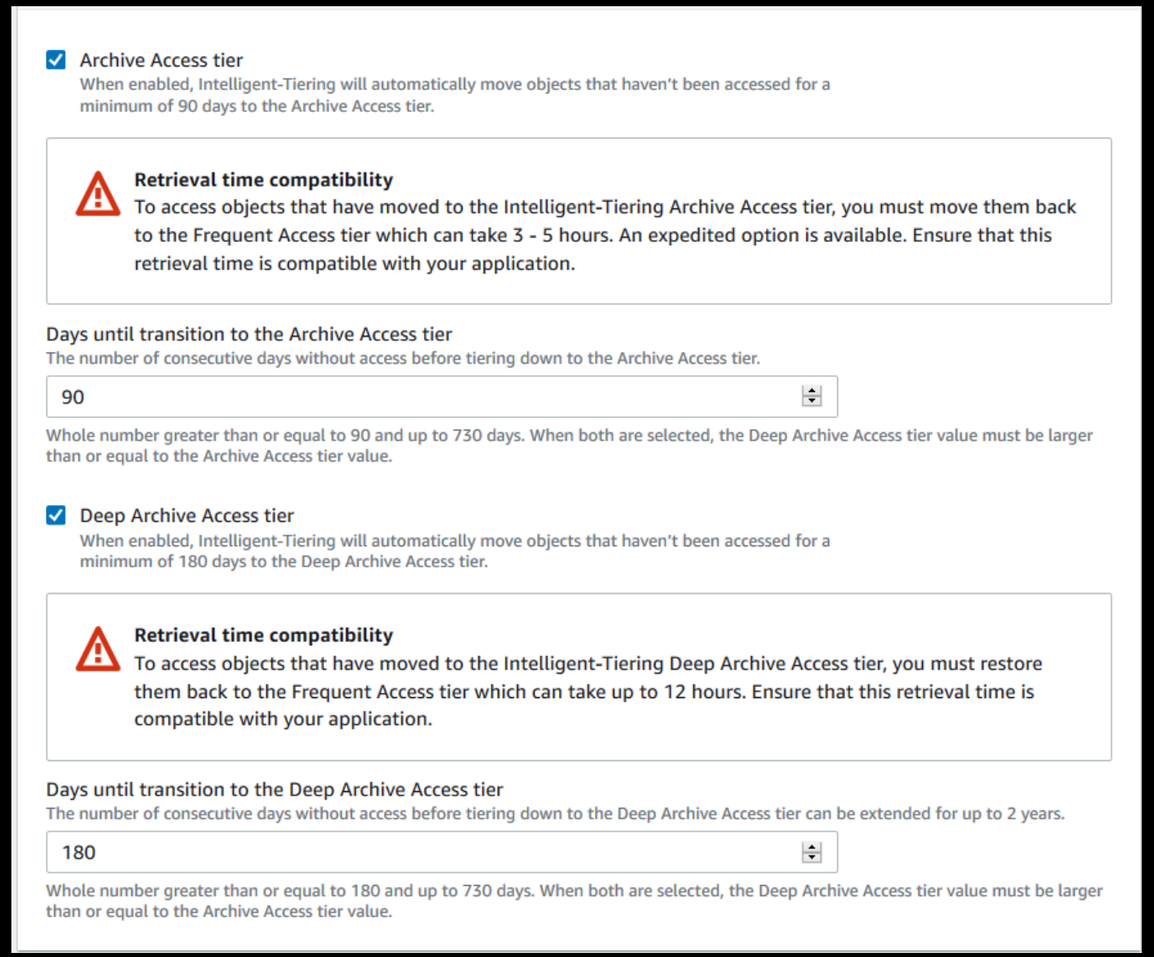
SAVE UP TO 95% ON STORAGE THAT IS AUTOMATICALLY MOVED TO THE DEEP ARCHIVE ACCESS TIER



Amazon S3 Intelligent-Tiering

ACTIVATE ONE OR BOTH OF THE ARCHIVE ACCESS TIERS AT THE BUCKET, PREFIX, OR OBJECT TAG LEVEL





NEW

Amazon S3 Strong Consistency

STRONG READ-AFTER-WRITE AND LIST CONSISTENCY FOR ALL APPLICATIONS



Strong read-after-write and list consistency for **any storage request** at **no additional cost**

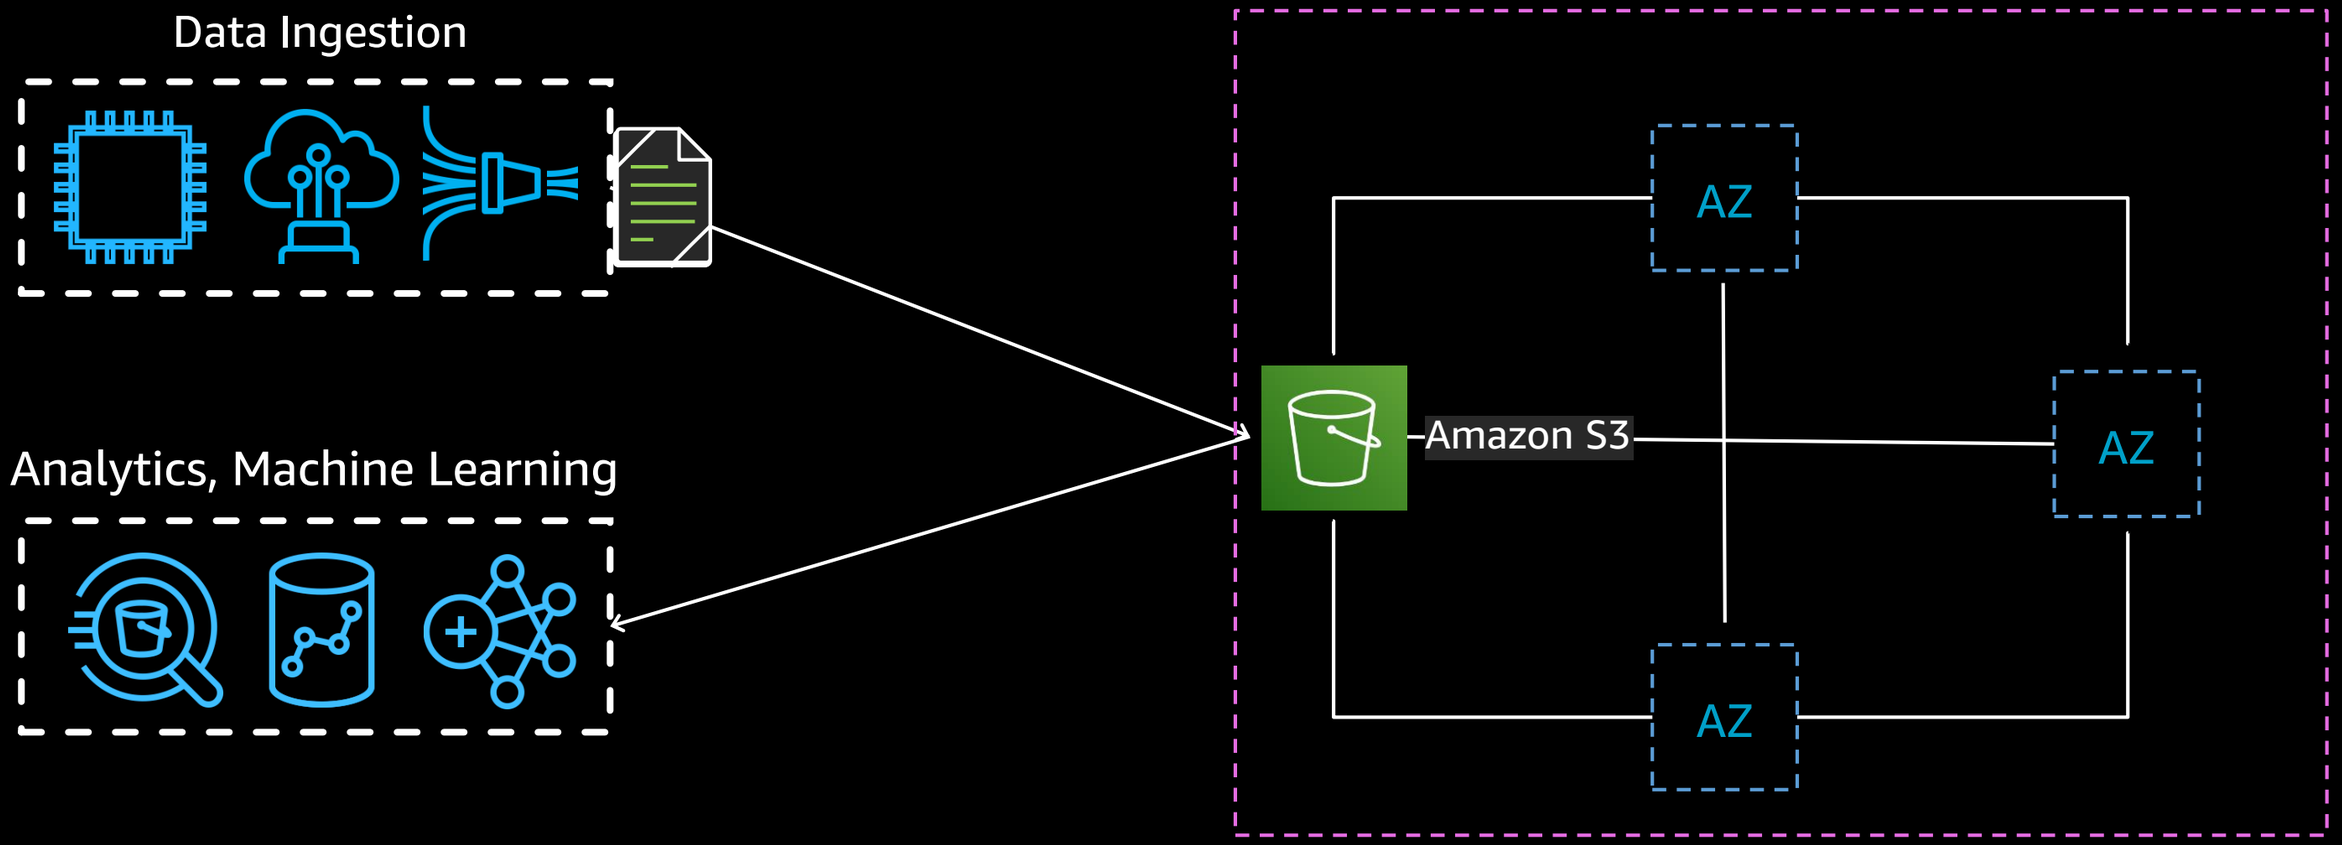
No changes to performance or availability, and without sacrificing regional isolation

Accelerates and **simplifies the migration** of analytics workloads to AWS

Removes the need for code or extra infrastructure to provide strong consistency

Amazon S3 Strong Consistency

ANY STORAGE REQUEST TO S3 STORAGE IS NOW STRONGLY CONSISTENT



Amazon S3 strong consistency

ACCELERATES AND SIMPLIFIES THE MIGRATION OF ANALYTICS WORKLOADS TO AWS

“Strong read-after-write consistency is a huge win for us because it helps drive down the time to complete our data processing pipelines by removing the software overhead needed to deliver consistency, and it also reduces our cost to operate the data lake by simplifying the data access architecture.”

- Ashish Gandhi, Technical Lead Data Infrastructure, Dropbox

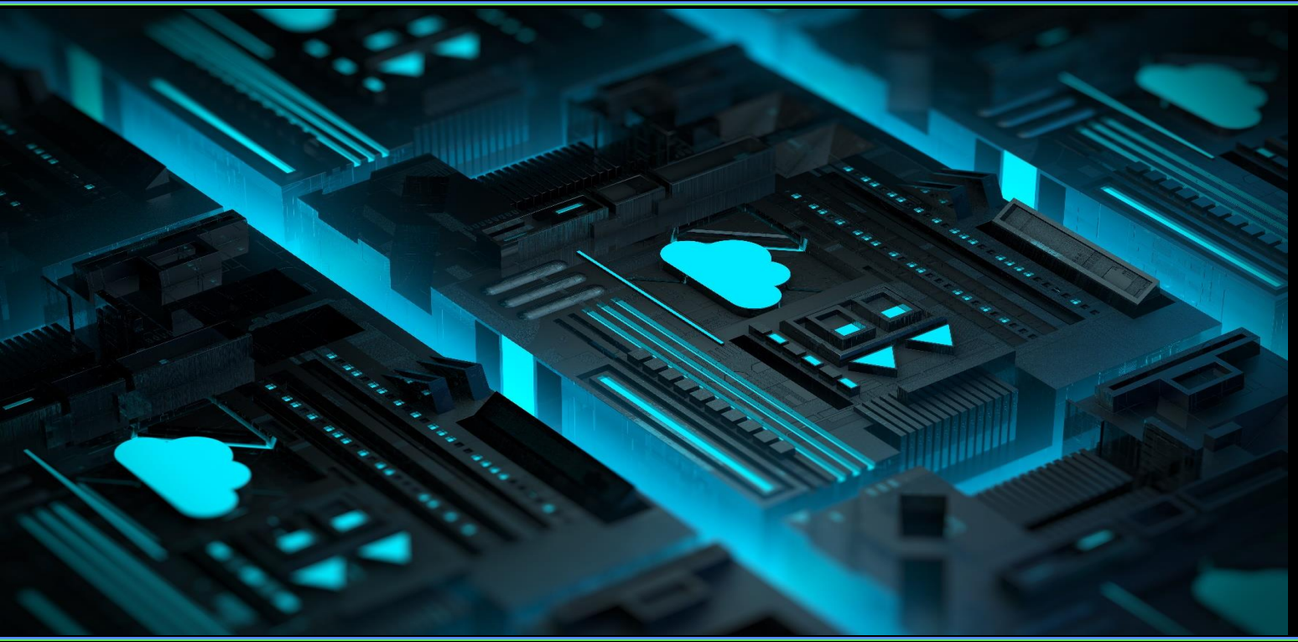
Use Case 1: When modifying (overwrite or delete) an object, you can immediately read a write, and expect that the latest write will be returned

---

Use Case 2: Analytics frameworks like Apache Hadoop or Apache Spark need strong consistency for list operations, as these workflows list contents of a bucket or prefix

---

Use Case 3: Strong consistency for S3 metadata – including S3 ACLs, S3 Object Tag, S3 Object Metadata – helps pipeline workflows that perform “existence checks” on an object



Innovation in Amazon S3

- Cost
- Performance
- Data movement
- Storage management
- Security
- Scalability, availability

AWS portfolio of data transfer services

BROADEST DATA MOVEMENT AND HYBRID CLOUD STORAGE OPTIONS

Offline transfer

Bulk data, files, objects, HDFS, databases



NEW!

AWS Snowcone



AWS Snowball Edge



AWS Snowmobile

Online file and object transfers

Rapid transfers



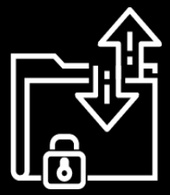
AWS DataSync

Long-distance uploads and downloads



Amazon S3 Transfer Acceleration

File exchanges



AWS Transfer Family

Database and machine migration and recovery



AWS Database Migration Service (AWS DMS)



CloudEndure Migration

Streaming data

Amazon Kinesis family



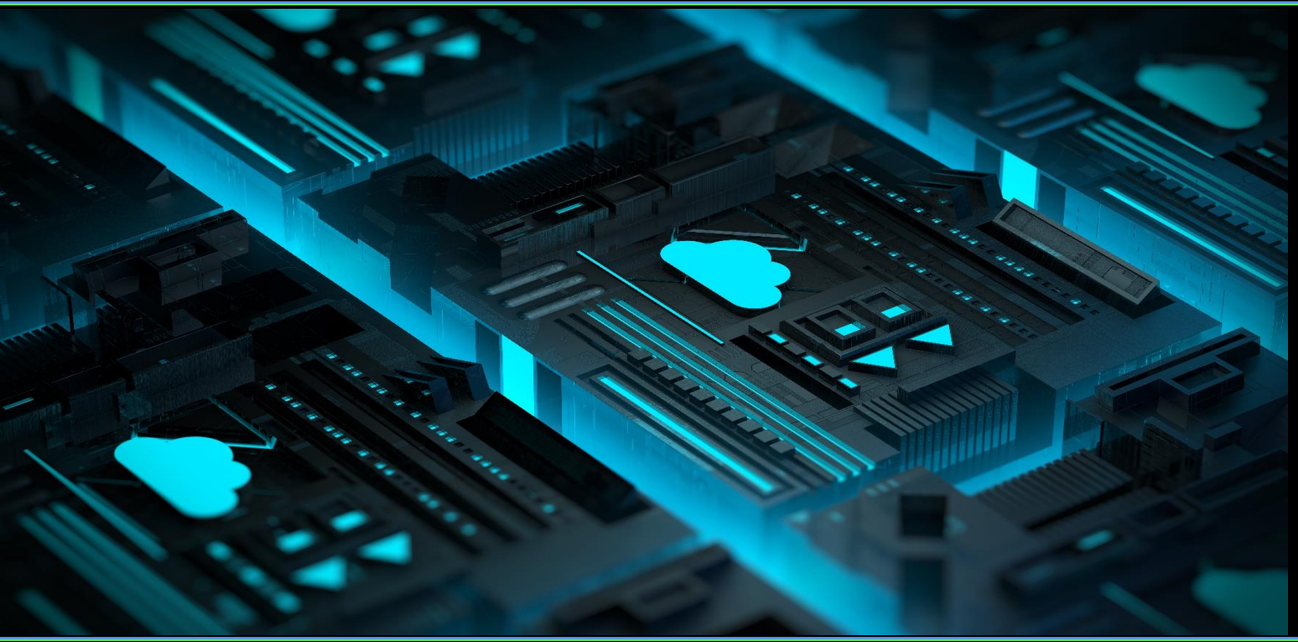
Amazon Kinesis Data Firehose, Amazon Kinesis Data Streams, Amazon Kinesis Video Streams

Hybrid / edge gateways

File, volume, and tape backup storage for on-premises applications



AWS Storage Gateway



Innovation in Amazon S3

- Cost
- Performance
- Data movement
- Storage management
- Security
- Scalability, availability

NEW

Amazon S3 Storage Lens

ORGANIZATION-WIDE VISIBILITY INTO STORAGE USAGE AND ACTIVITY TRENDS



Visibility into storage usage and activity across accounts with drill-downs into the account, bucket, storage class, region, or prefix level

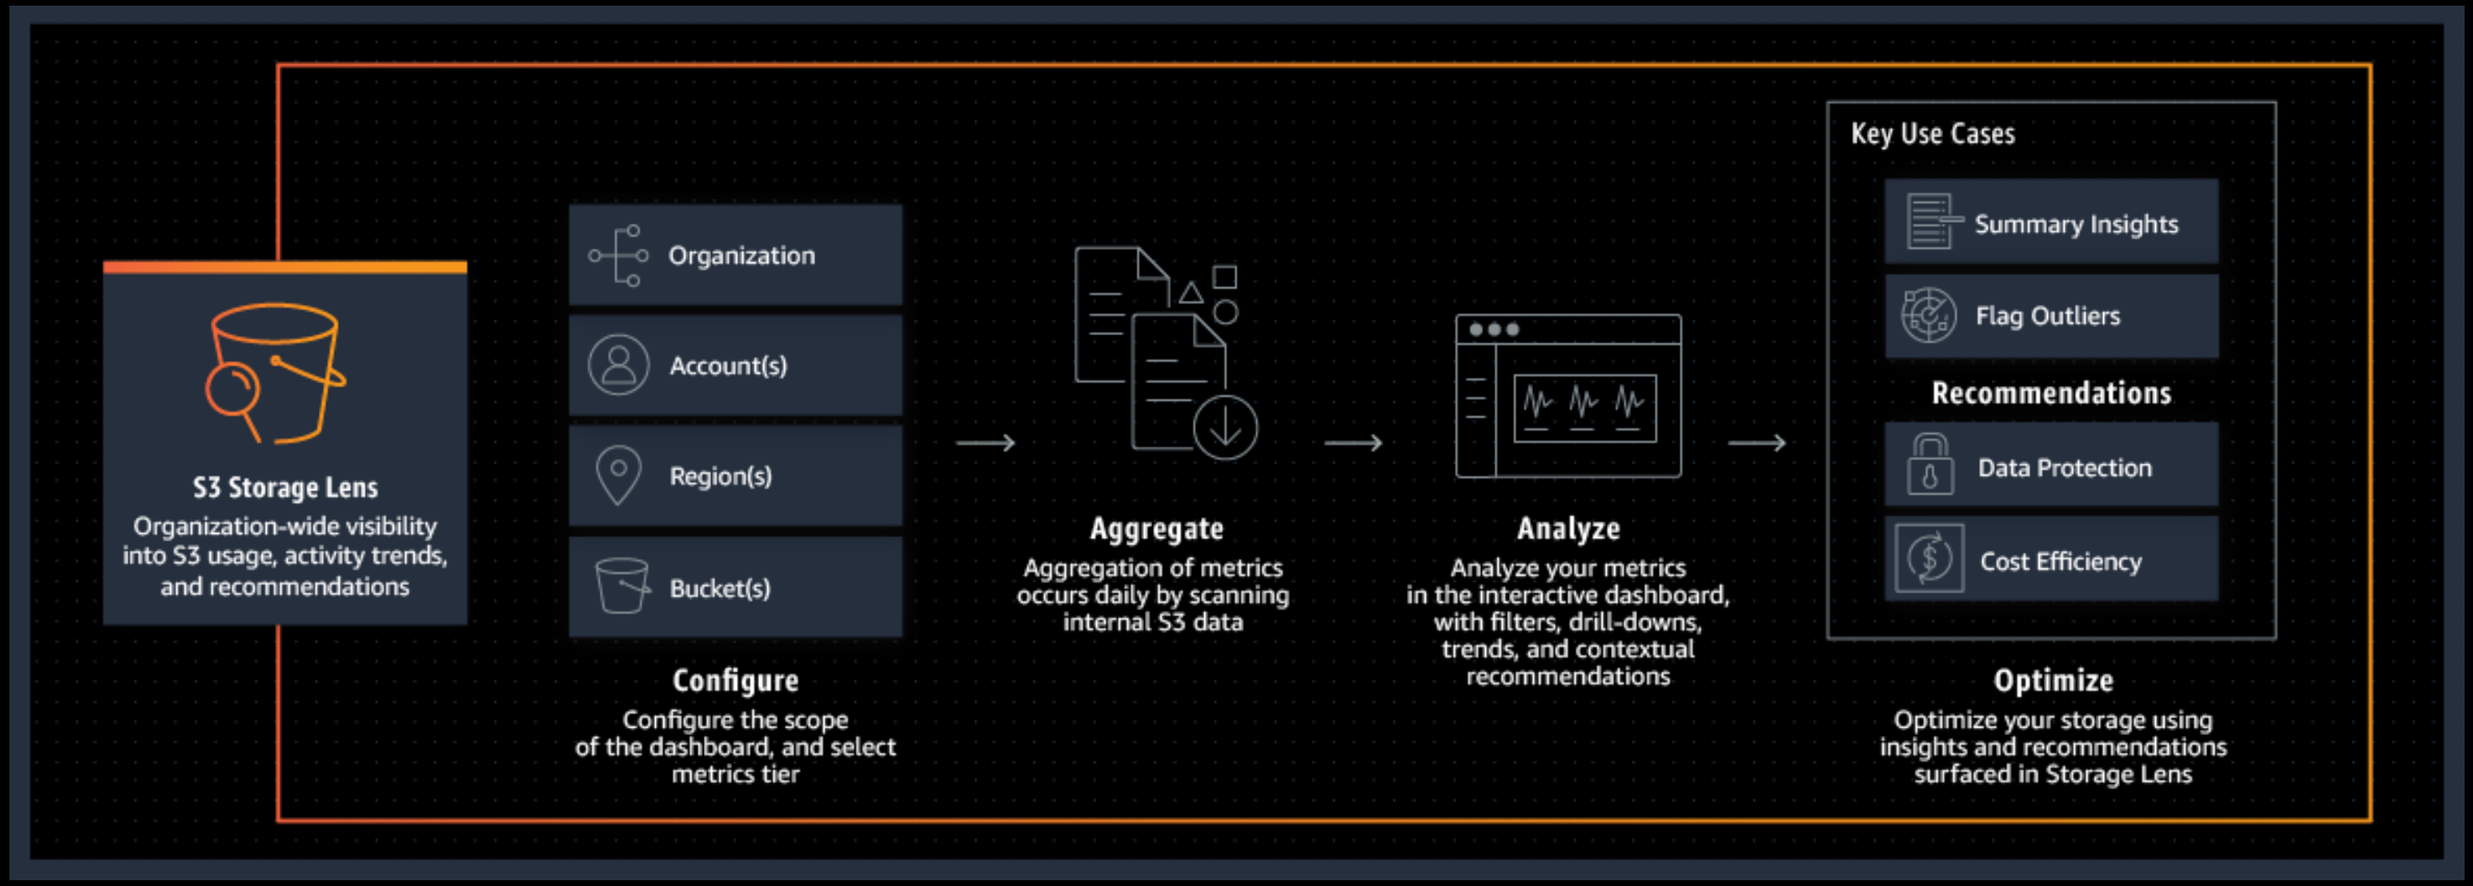
Pre-aggregated storage usage and activity metrics with data retention up to 15 months

Interactive dashboard in the Amazon S3 console, or export metrics in CSV or Parquet format to S3 bucket

Contextual recommendations to find ways to reduce storage costs and apply best practices on data protection

Amazon S3 Storage Lens

DEFAULT DASHBOARD AND CUSTOMIZABLE CONFIGURATION TO ADJUST SCOPE AND METRICS



Amazon S3 Storage Lens

CUSTOMERS USE S3 STORAGE LENS TO MANAGE THEIR ENTIRE STORAGE ESTATE IN A SIMPLE WAY

"The ability to use see metrics based on prefixes will be invaluable to us. We will have better insights into usage for particular prefixes including total storage, object counts, average object sizes and break down by current and non-current versions, which can help us debug and see the trends of operations in an economical way. This wasn't previously achievable before S3 Storage Lens."

Ju-Yi Kuo, Senior Software Engineer - Snowflake

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

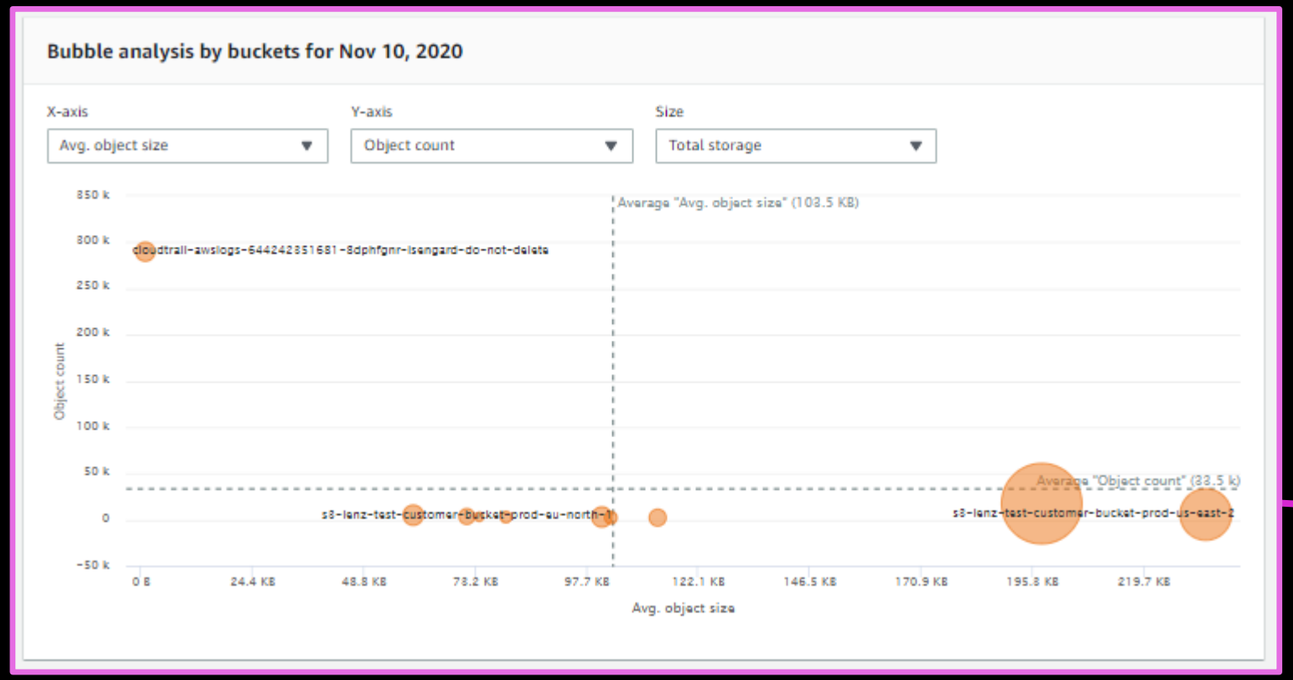
Outliers: Are there any unusual trends or spikes in my storage usage or activity?

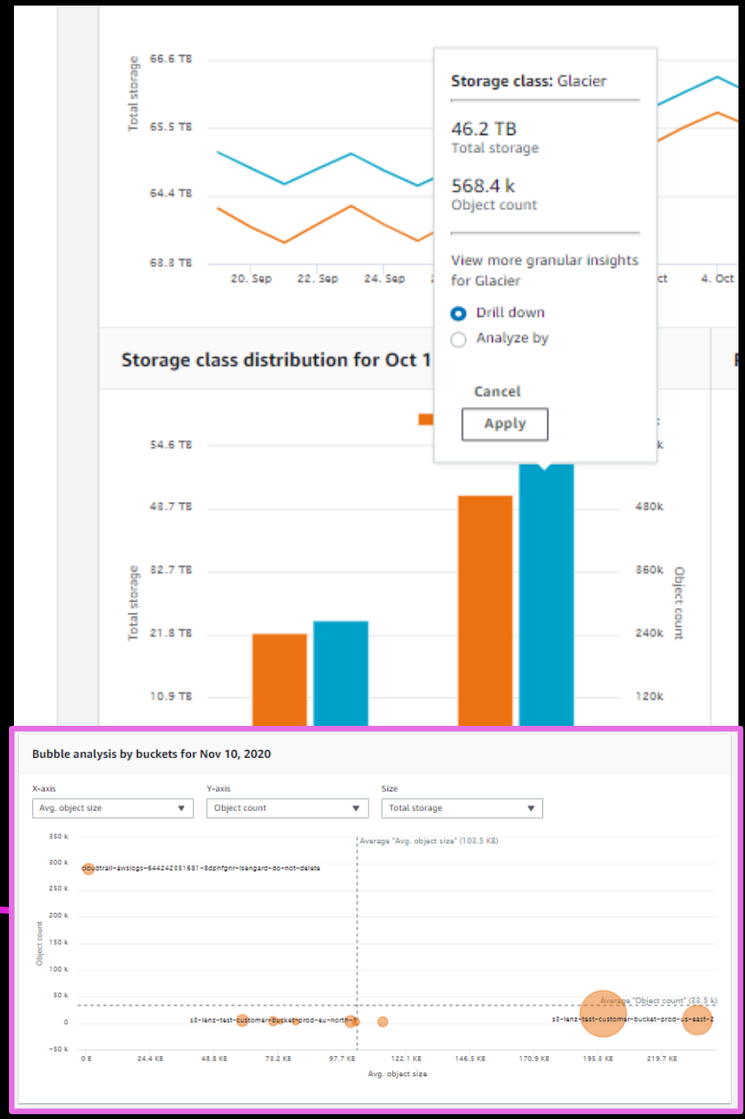
Data protection: Are all of my buckets encrypted and what % of my storage is versioned In a region?

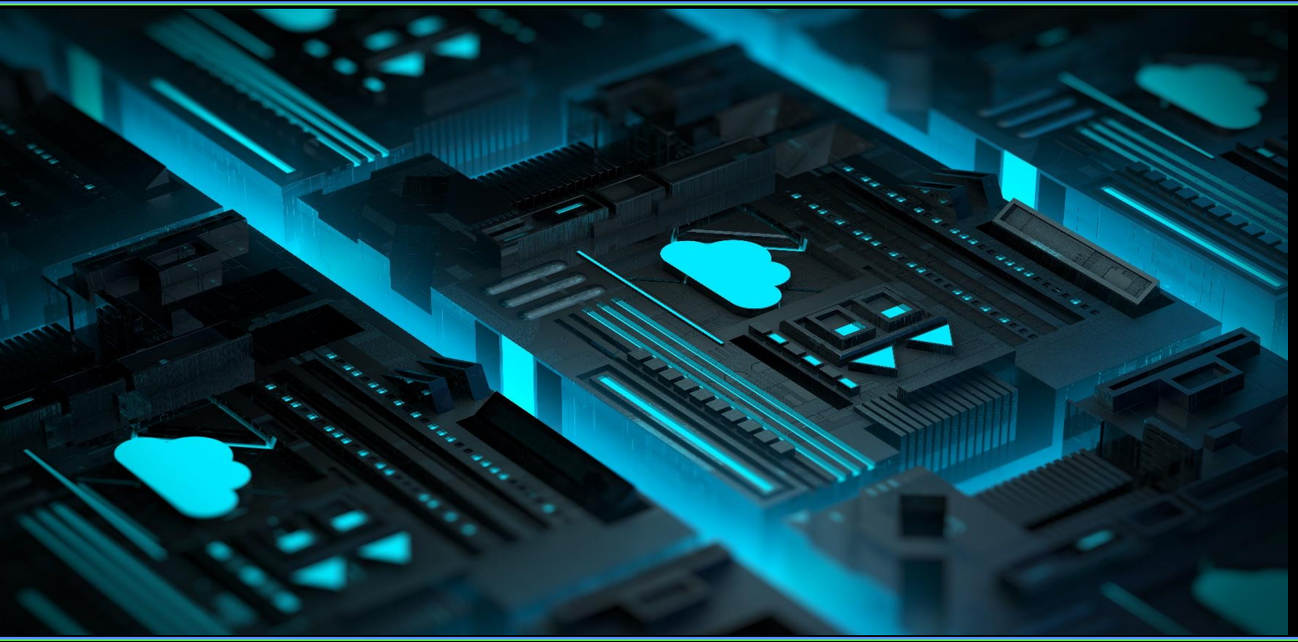
Cost efficiency: Which of my buckets are infrequently accessed?

Amazon S3 Storage Lens

DATA VISUALIZATIONS PROVIDE YOU INSIGHTS ON YOUR STORAGE







Innovation in Amazon S3

- Cost
- Performance
- Data movement
- Storage management
- Security
- Scalability, availability

Amazon S3 Security and Access features

SECURITY IS JOB NUMBER ONE AND WILL ALWAYS BE OUR PRIORITY



Action Last Accessed for Amazon S3

Tighten S3 permissions for your IAM users and roles using access history of S3 actions



Amazon Guard Duty for Amazon S3

New level of visibility into access of data stored in S3



AWS X-Ray for Amazon S3

Trace context propagation enabling customers to view end-to-end requests when using S3



S3 Bucket Owner Condition

Ensures that buckets you use in your S3 operations belong to the AWS accounts that you expect



S3 Object Ownership Overwrite

Enable bucket owners to assume ownership of objects uploaded to their buckets



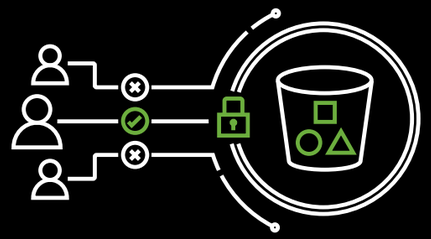
Amazon S3 Bucket Keys

Bucket-level approach to key management for SSE-KMS

NEW

Amazon S3 Bucket Owner Condition

VALIDATE THE CORRECT BUCKET OWNER USING YOUR AMAZON S3 BUCKET



Allows customers to validate the AWS Account ID of the owner of an S3 bucket

Bucket owner condition helps easily verify that the S3 bucket customers interact with are owned by expected AWS Accounts

Helps to prevent accidental interaction with buckets owned by unexpected AWS Accounts

When bucket owner condition is used, S3 API requests will only succeed if the bucket owner matches the account specified

NEW

Amazon S3 Object Ownership Overwrite

ENABLE BUCKET OWNERS TO ASSUME OWNERSHIP OF OBJECTS UPLOADED TO THEIR BUCKETS



You can **control ownership of new objects** that are uploaded to your buckets

Standardize ownership of new objects when you **create shared data sets** with multiple accounts writing to and reading from

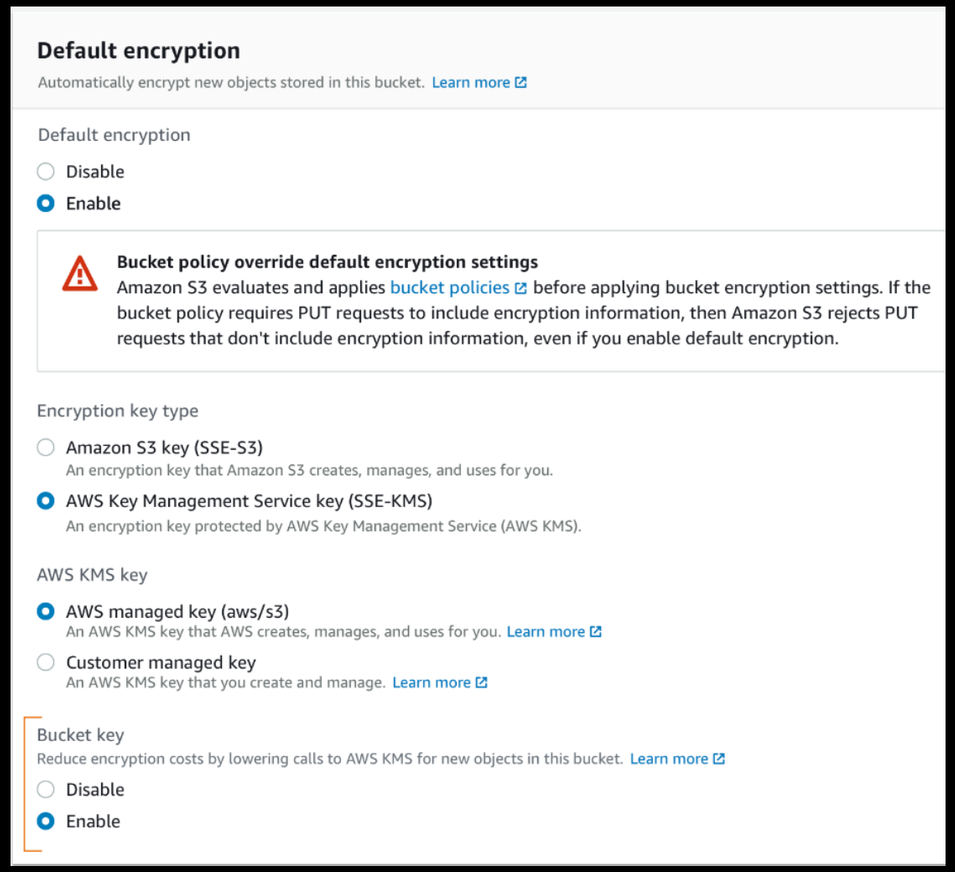
Share and manage access to these objects via **resource-based bucket policies** or Access Points

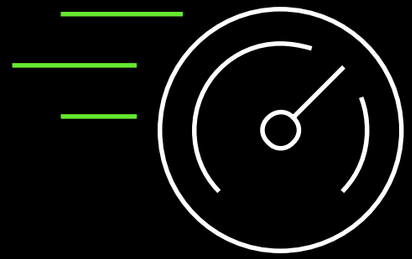
Enforce ownership by adding bucket policy to require all Amazon S3 PUT operations to include the **“bucket-owner-full-control” canned ACL**

NEW

Amazon S3 Bucket Keys

REDUCE REQUESTS COSTS FOR SERVER-SIDE ENCRYPTION (SSE-KMS) BY UP TO 99%



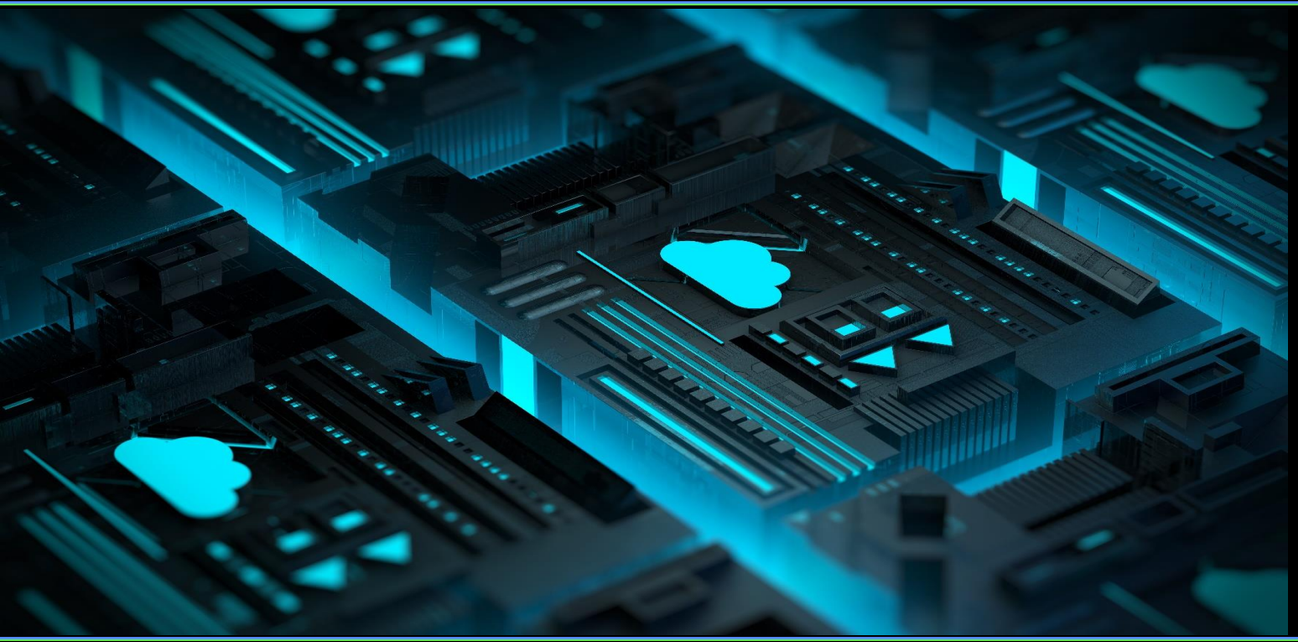


Increase performance for encryption



Reduce request costs for KMS-backed server-side encryption by up to

99%



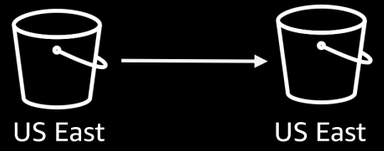
Innovation in Amazon S3

- Cost
- Performance
- Data movement
- Storage management
- Security
- Scalability, availability

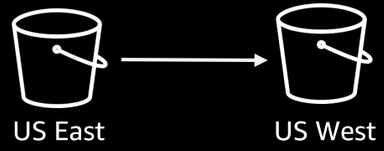
Amazon S3 Replication

S3 REPLICATION IS AN ELASTIC, FULLY MANAGED, LOW COST FEATURE THAT AUTOMATICALLY REPLICATES OBJECTS BETWEEN BUCKETS

Same Region Replication (SRR)
Replicate within the same AWS Region

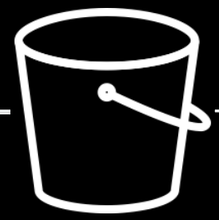


Cross Region Replication (CRR)
Replicate to a different AWS Region



Replication Time Control (RTC)
Predictable replication backed by SLA





Amazon S3

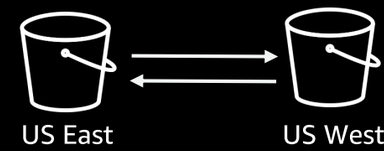
NEW!

Multi-Destination Replication
Replicate to multiple AWS Regions



NEW!

Two-Way Replication
Sync replica and metadata changes



NEW!

Replication Metrics and Notifications
Three metrics for replication destination



NEW

Amazon S3 Replication

MULTI-DESTINATION REPLICATION AND TWO-WAY DIRECTIONAL REPLICATION OPEN UP NEW USE CASES

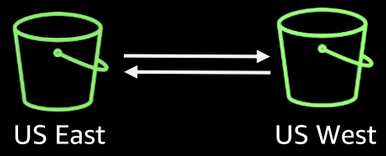


Multi-Destination Replication

Maintain multiple copies of data in the same or different regions

Create shared data sets between global teams and minimize latencies for geographically distributed users

Use CloudWatch metrics to track replication progress



Two-Way Directional Replication

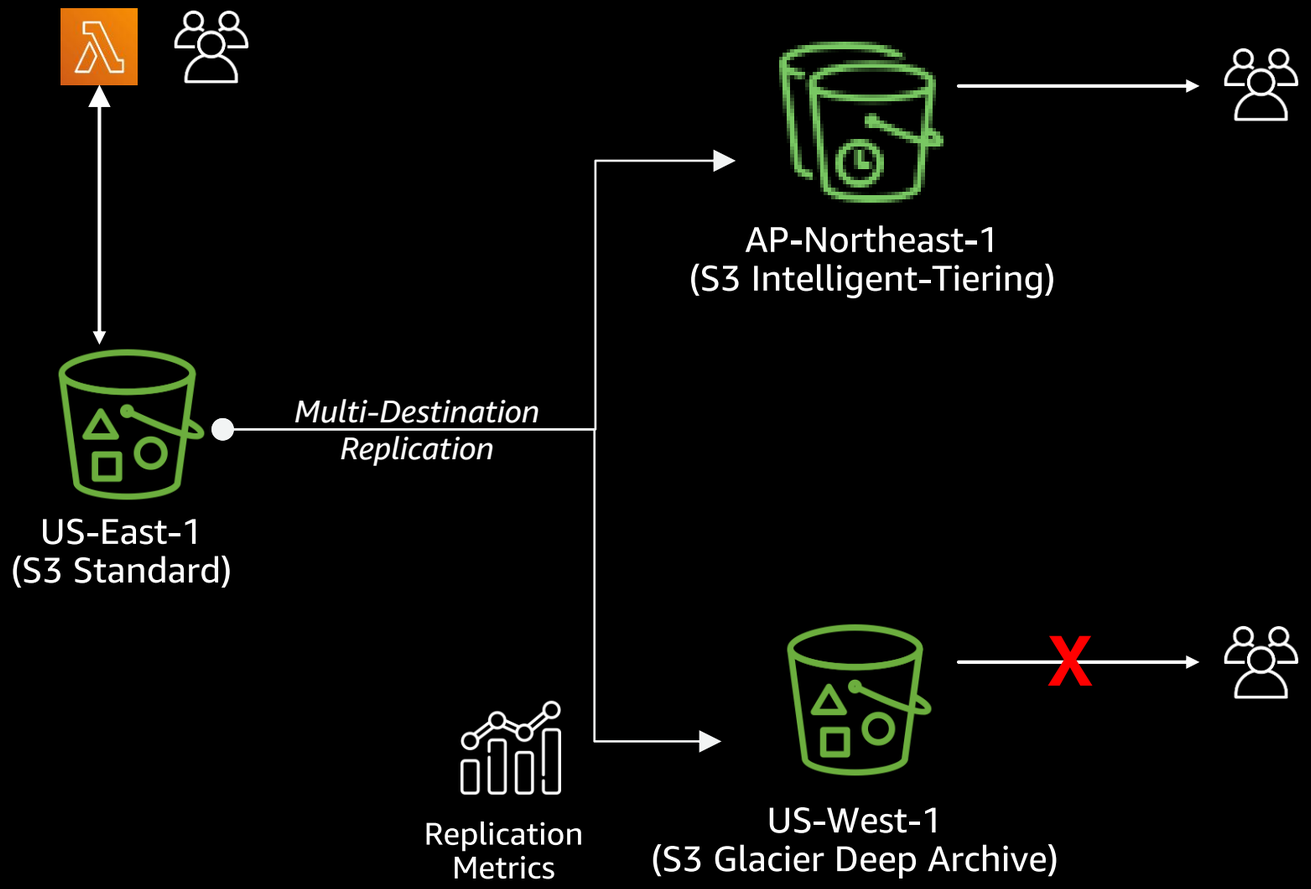
Sync replica changes and replica metadata changes, including object tags and objects ACLs, back to source bucket

You can create shared datasets across multiple regions, and keep all object metadata in sync, to and from the source and destination

Replicate delete markers to one or more buckets with an additional flag in the replication configuration

Amazon S3 Replication

NEW S3 REPLICATION FEATURES HELP YOU MEET EVOLVING BUSINESS REQUIREMENTS



Use case of international financial institution with shared datasets that are read across the world

Reads and writes from source bucket but only occasional reads from one of the destination buckets

Multi-Destination Replication with one bucket for low-latency access and another for bad actor protection

Replication Time Control to meet business requirements

Replication Metrics for the data replicated into different account to protect against bad actor

Use different storage classes to optimize cost and performance

Putting it all together

IT'S ALWAYS DAY 1



S3 Intelligent-Tiering for data lakes with automatic cost savings of up to 95%.



S3 Strong Consistency simplifies the migration of analytics to AWS at no cost
S3 on Outposts for data residency and local performance needs



Broad data movement and hybrid cloud storage options.



Amazon **S3 Storage Lens** provides customers organization-wide visibility into storage usage and activity trends



Security is our number one priority (e.g. Action Last Accessed, Guard Duty, X-Ray, Bucket Owner Condition, Object Ownership Overwrite, Bucket Keys)



Multi-Destination and **Two-Way Directional Replication** are elastic, fully managed features to automatically replicate objects between buckets

Taking it to the next level

STG 304
Lessons from the vanguard: Build modern apps using Amazon S3 or Amazon EBS

STG 202
Now is the time: Move your workloads to AWS storage

STG 203
Breakdown data silos: Build a serverless data lake on Amazon S3

STG 205
Amazon S3 foundations: Best practices for S3

STG 211
Extend Amazon S3 to on-premises environments with AWS Outposts

STG 218
Accelerate your migration to Amazon S3

STG 301
Architecting for high availability on Amazon S3

STG 302
Data lake security in Amazon S3: Perimeters and fine-grained controls

STG 304
Lessons from the vanguard: Build modern applications using S3

Thank you!



AWS re:Invent

© 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.