
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

STG337

Best practices for managing S3 data at scale, with Bridgewater Associates

Shakhi Hali (she/her)

Senior PMT – External Services, Amazon S3
AWS

Rob Wilson (he/him)

Senior Manager, Product Management, Amazon S3
AWS

Robin Anil (he/him)

Director, Research Technology
Bridgewater Associates

Agenda

Overview of Amazon S3

Use case deep dives

- ❑ Securing your data
- ❑ Accessing archive data
- ❑ Monitoring access
- ❑ Managing storage spend

Customer voice: Bridgewater Associates

Recap

Overview of Amazon S3

Operate at scale with Amazon S3

DESIGNED FOR 11 9S OF DURABILITY

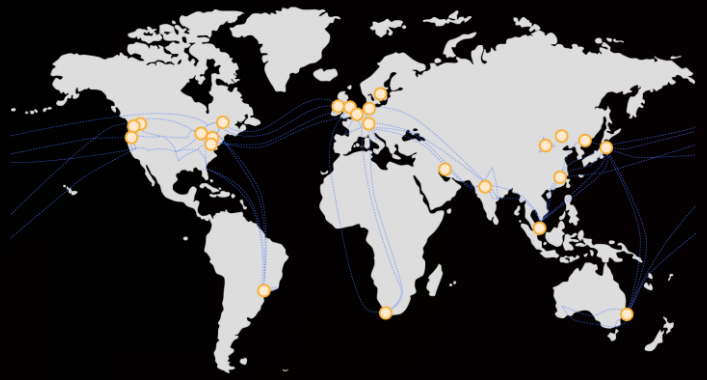

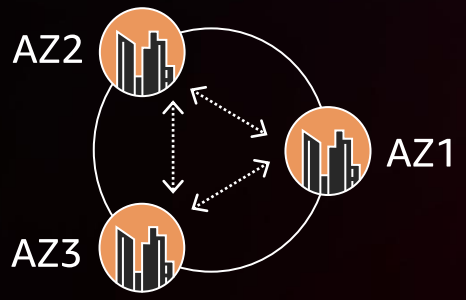

30 Regions

Geographic locations with multiple AZs, separated physically by miles and isolated

96 Availability Zones (AZs)

Objects stored across multiple devices spanning a minimum 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

Industry-leading durability, availability, and performance

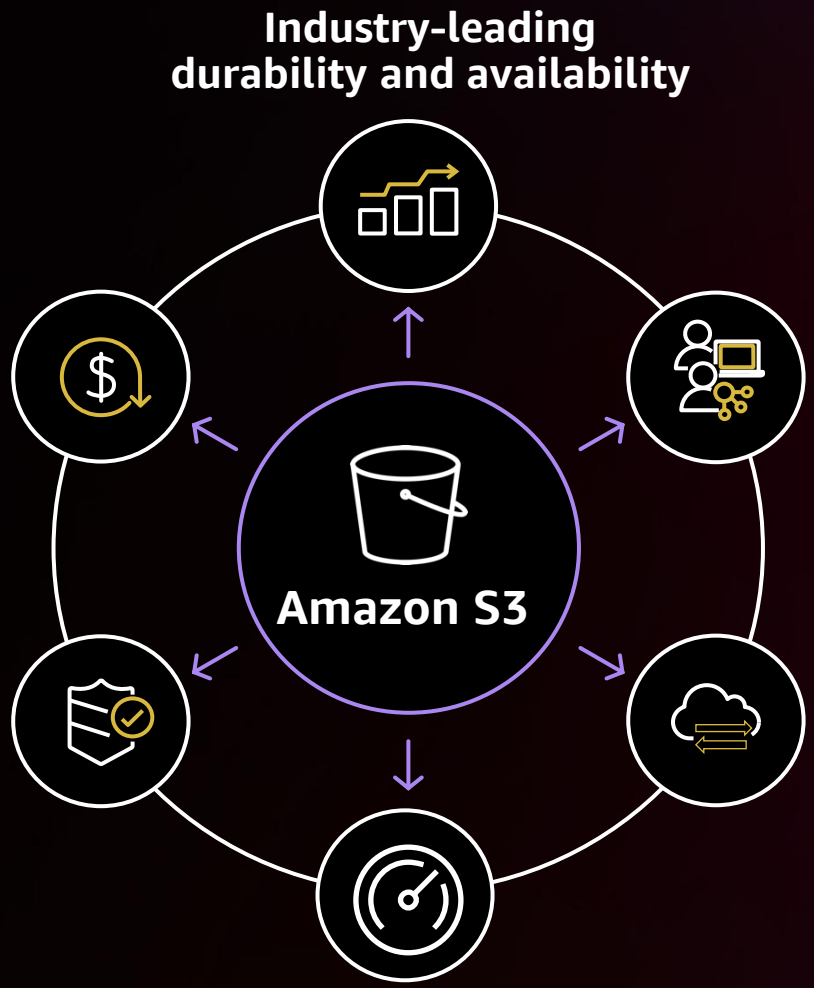

AWS storage customers

Use case deep dives

Securing your data:
Encryption at scale

Amazon S3 encryption support

Amazon S3 default encryption

Bucket-level setting

Automatically encrypt all new objects

Simplified compliance

Supports SSE-S3 and SSE-KMS

Provides Amazon S3 encryption-at-rest support for applications that do not otherwise support encrypting data in Amazon S3

Customer use case 1

The cloud security lead calls a meeting for a detailed security review

The objects stored in Amazon S3 are subject to a wide range of regulatory guidelines

You need to provide a compliance report proving that all the objects stored are encrypted

Default encryption
S3 Inventory
S3 Batch Operations

Approach to encrypting your existing objects

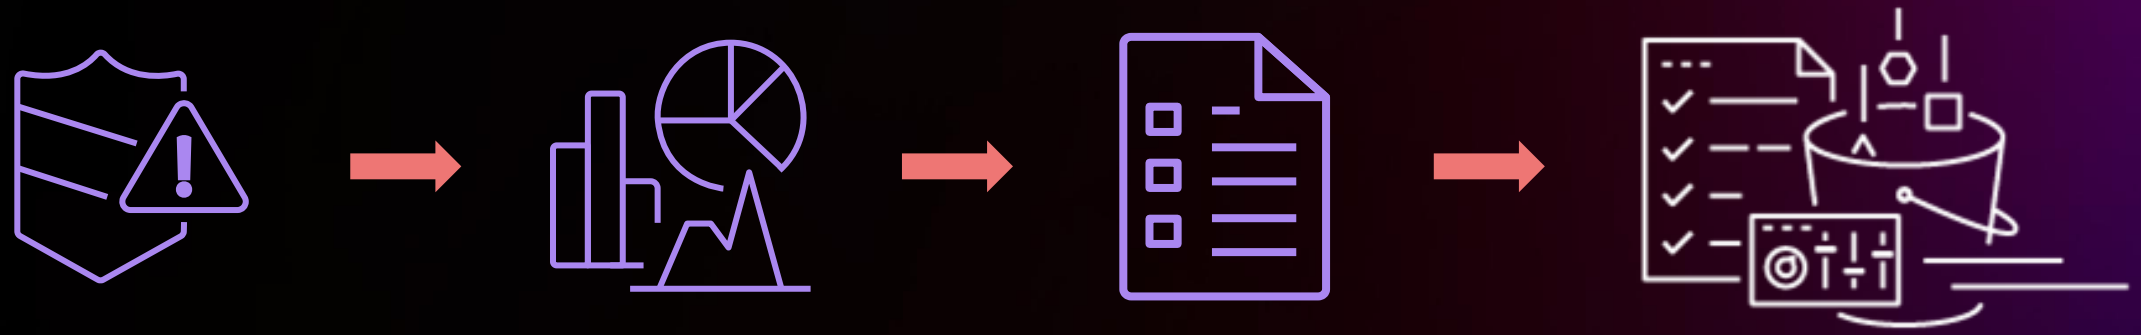

Default encryption

S3 Storage Lens

Inventory

S3 Batch Operations

Encryption insights on your existing objects in S3 Storage Lens

UNDERSTAND YOUR ORGANIZATION-WIDE PERCENTAGE ENCRYPTION BYTES BY REGION AND BUCKETS

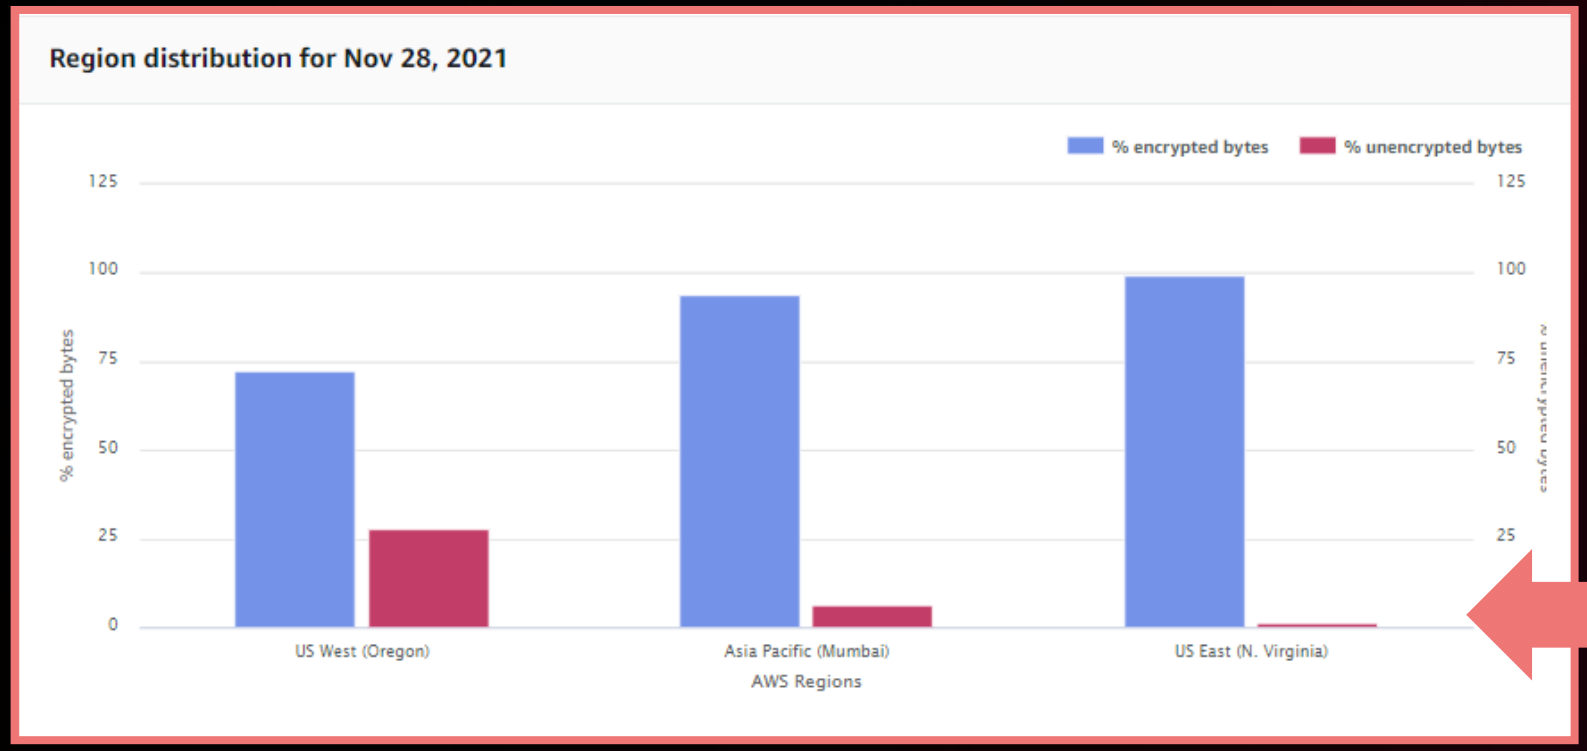

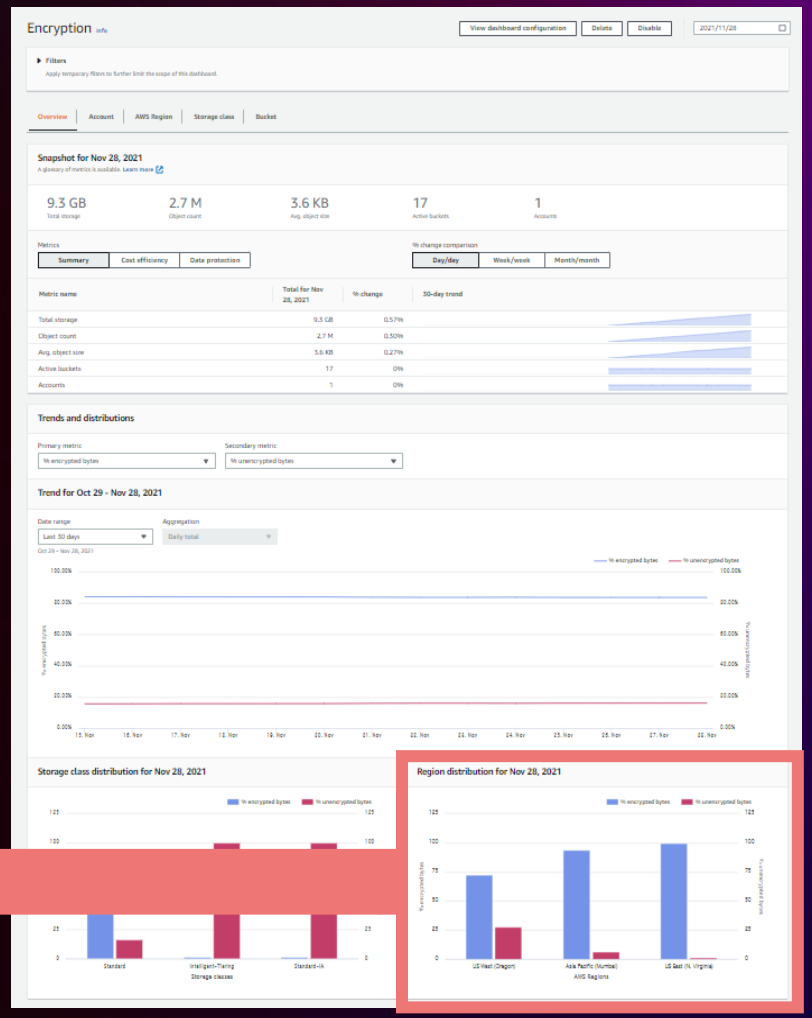

Set up S3 Inventory to identify existing objects

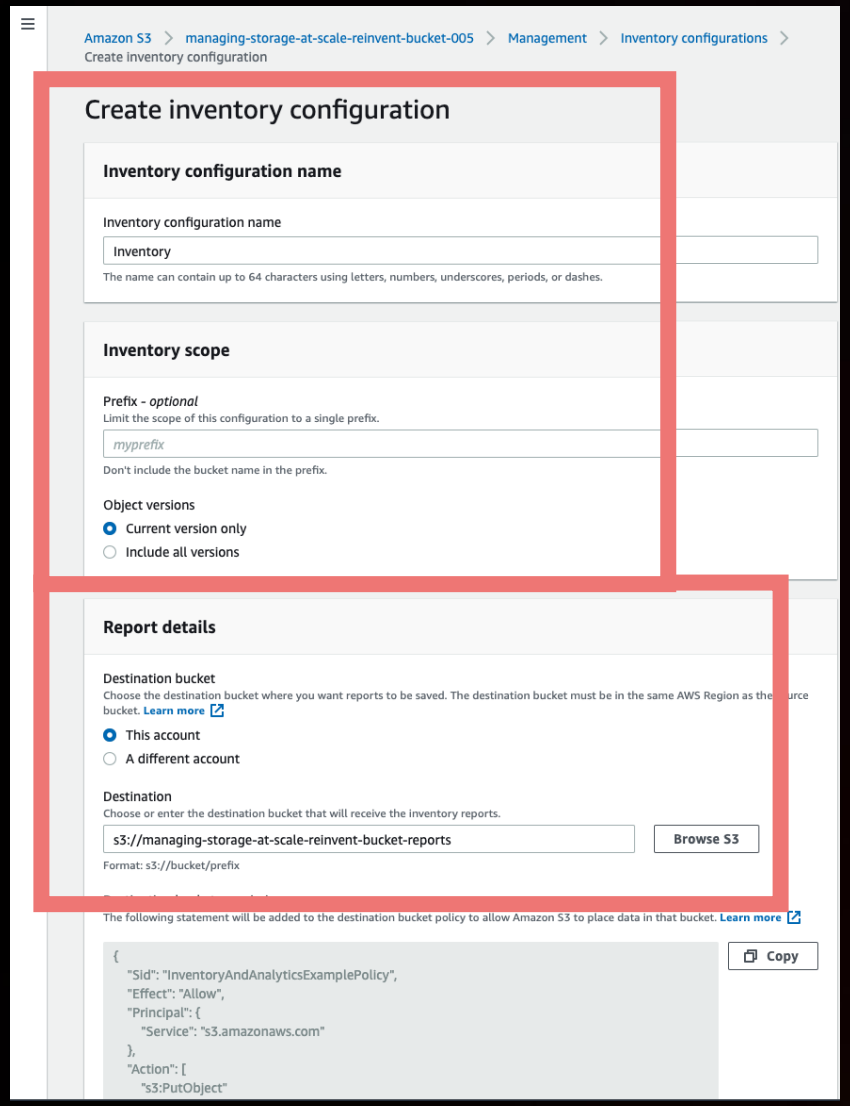

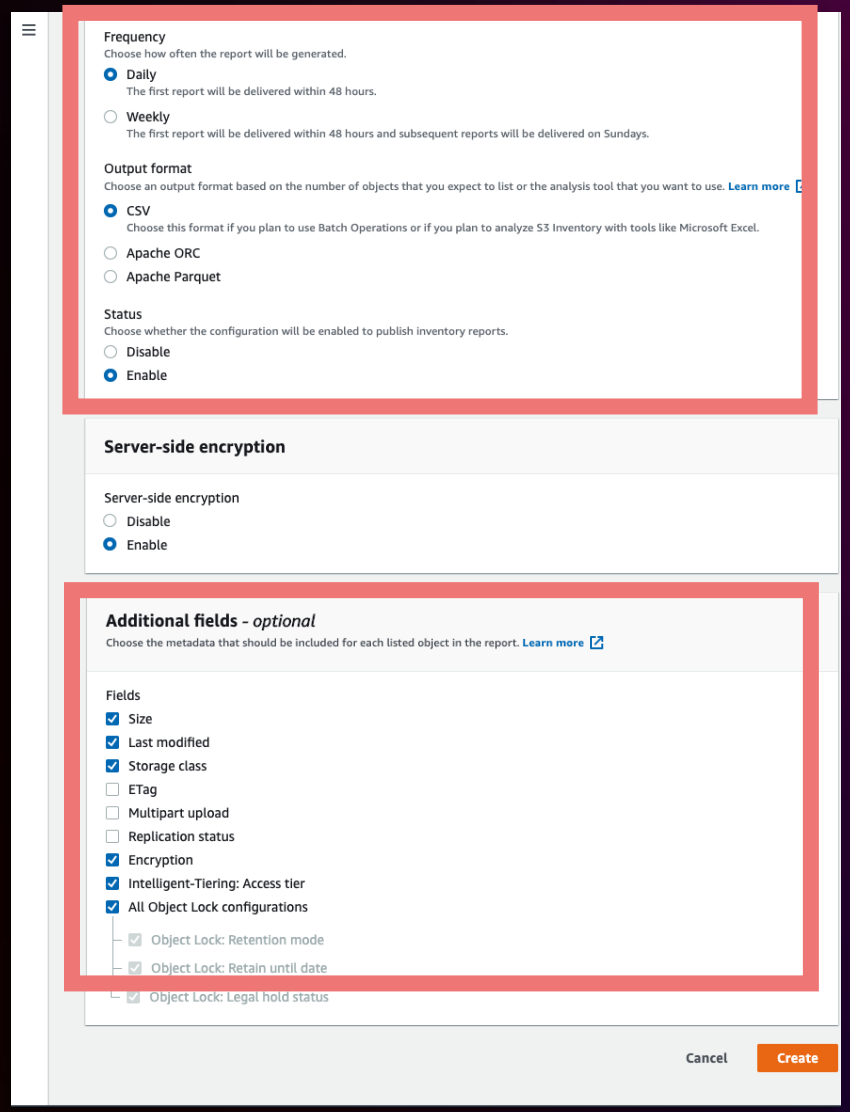

NOT-SSE.

Use Amazon Athena to filter objects “NOT-SSE” in S3 Inventory

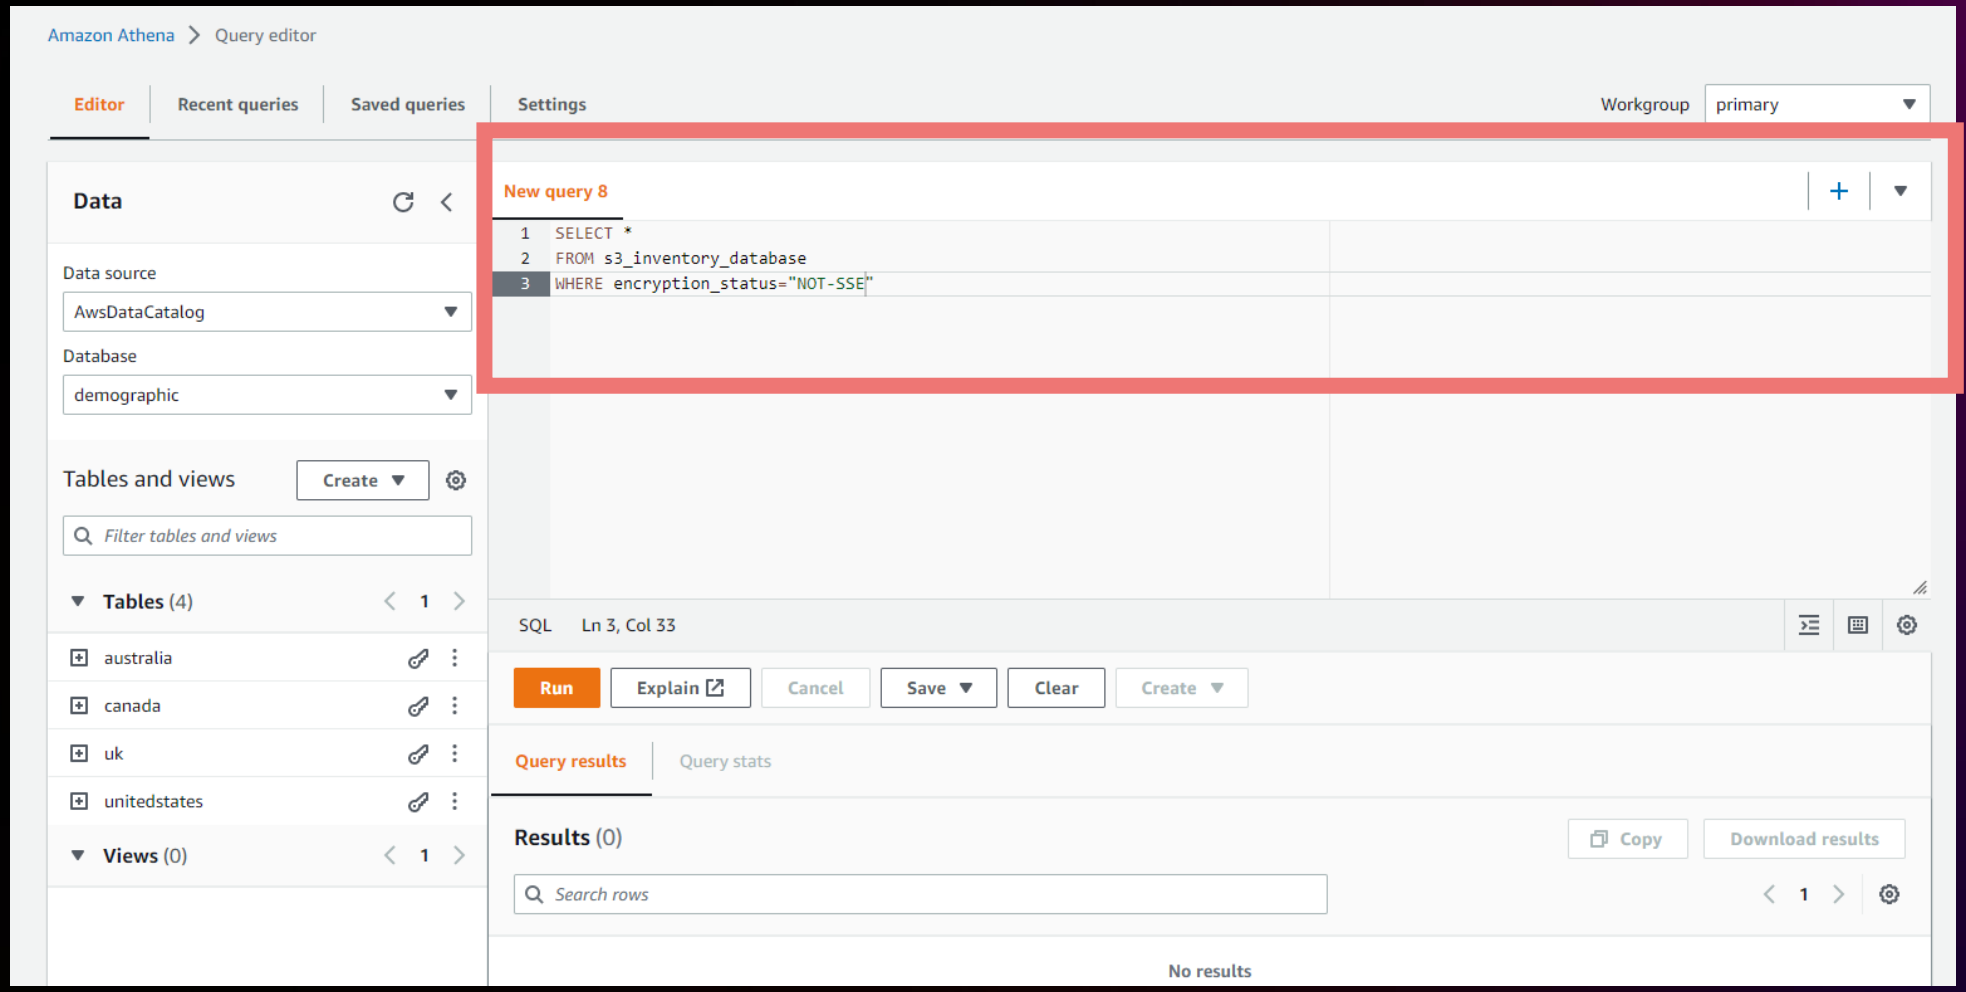

Use S3 Batch Operations to operate on millions and billions of objects

- S3 Inventory report
- CSV list

- Copy
- Invoke Lambda functions
- Replace all object tags
- Replace access control list (ACL)
- Restore from S3 Glacier
- Set Object Lock retention
- Set Object Lock legal hold
- Delete object tags

- Object-level progress
- Job notifications
- Completion report

S3 Batch Operations code for creating encrypted copy of objects

```
aws s3control create-job \

    --region us-west-2 \

    --account-id 12345678\

    --operation '{"S3PutObjectCopy": { "TargetResource": "arn:aws:s3:::destination-bucket" }}' \

    --manifest
'{"Spec":{"Format":"S3BatchOperations_CSV_20180820","Fields":["Bucket","Key"]},"Location":{"ObjectArn":"
arn:aws:s3:::my_manifests/manifest.csv","ETag":"60e460c9d1046e73f7dde5043ac3ae85"}}' \

    --report '{"Bucket":"arn:aws:s3:::bucket-where-completion-report-goes","Prefix":"final-reports",
"Format":"Report_CSV_20180820","Enabled":true,"ReportScope":"AllTasks"}' \

    --priority 42 \

    --role-arn IAM-role \

    --client-request-token $(uuidgen) \

    --description "job Description" \

    --no-confirmation-required
```

Accessing archive data from Amazon S3 Glacier

Why archive to AWS?

Lowest cost

Durability & resilience

Security & compliance

S3 Glacier storage classes

S3 Intelligent-Tiering

S3 Standard

S3 Standard-IA

S3 Glacier Instant Retrieval

S3 Glacier Flexible Retrieval (formerly S3 Glacier)

S3 Glacier Deep Archive

S3 One Zone-IA

S3 on Outposts

AWS Region | 3 Availability Zones (AZs)

Data with changing access patterns

- Millisecond access
- No retrieval fees
- Object monitoring charge
- Opt-in async archive tiers
- **Archive Instant Access tier**

Frequently accessed data

- Millisecond access

Infrequently accessed data

- Millisecond access

Rarely accessed data

Millisecond access

Archive data

- Retrieval options from minutes to hours
- **Free bulk retrievals**

Long-term archive data

- Retrieval in hours

AWS AZ

Re-creatable, infrequently accessed data

- Millisecond access
- Single AZ resiliency

AWS Outposts

On-premises data

- Millisecond access

Increased throughput (10x) for Amazon S3 Glacier

New!

UP TO 90% FASTER RETRIEVALS FROM S3 GLACIER FLEXIBLE RETRIEVAL, S3 GLACIER DEEP ARCHIVE

Launch blog

Automatically applies to S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive (standard and bulk retrievals) at no additional cost

Supports restore requests at a rate of up to 1,000 transactions per second, per account in an AWS Region

Ideal for restoring backups, responding to audit requests, retraining machine learning models, and performing analytics on historical data

Significantly reduces the restore completion time for datasets composed of small objects

aws

Customer use case 2

Your data science team now wants to access long-term archival datasets to train their machine learning model

How can you provide quick and easy access to archives to your data science teams?

S3 Glacier storage classes
S3 Glacier increase in throughput New!
Restore using S3 Batch Operations
Restore notifications using S3 Event Notifications

Use S3 Batch Operations to restore your data from S3 Glacier Deep Archive

1. Choose your AWS Region and input (manifest or CSV) to select objects to be restored

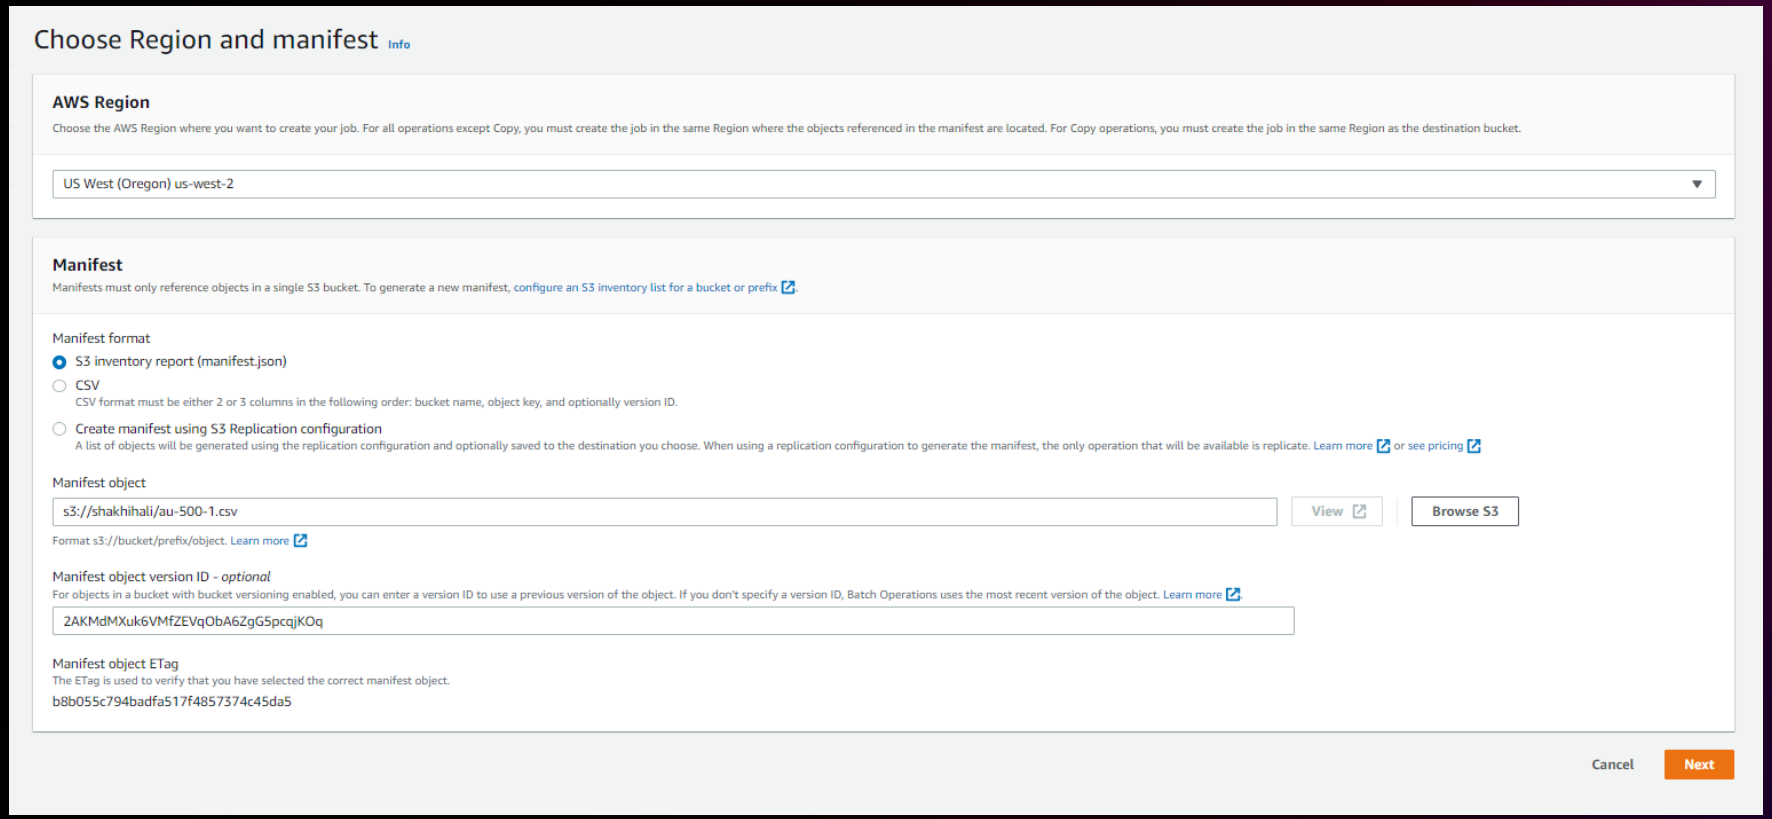

Use S3 Batch Operations to restore your data from S3 Glacier Deep Archive

2. Select **Restore** operation

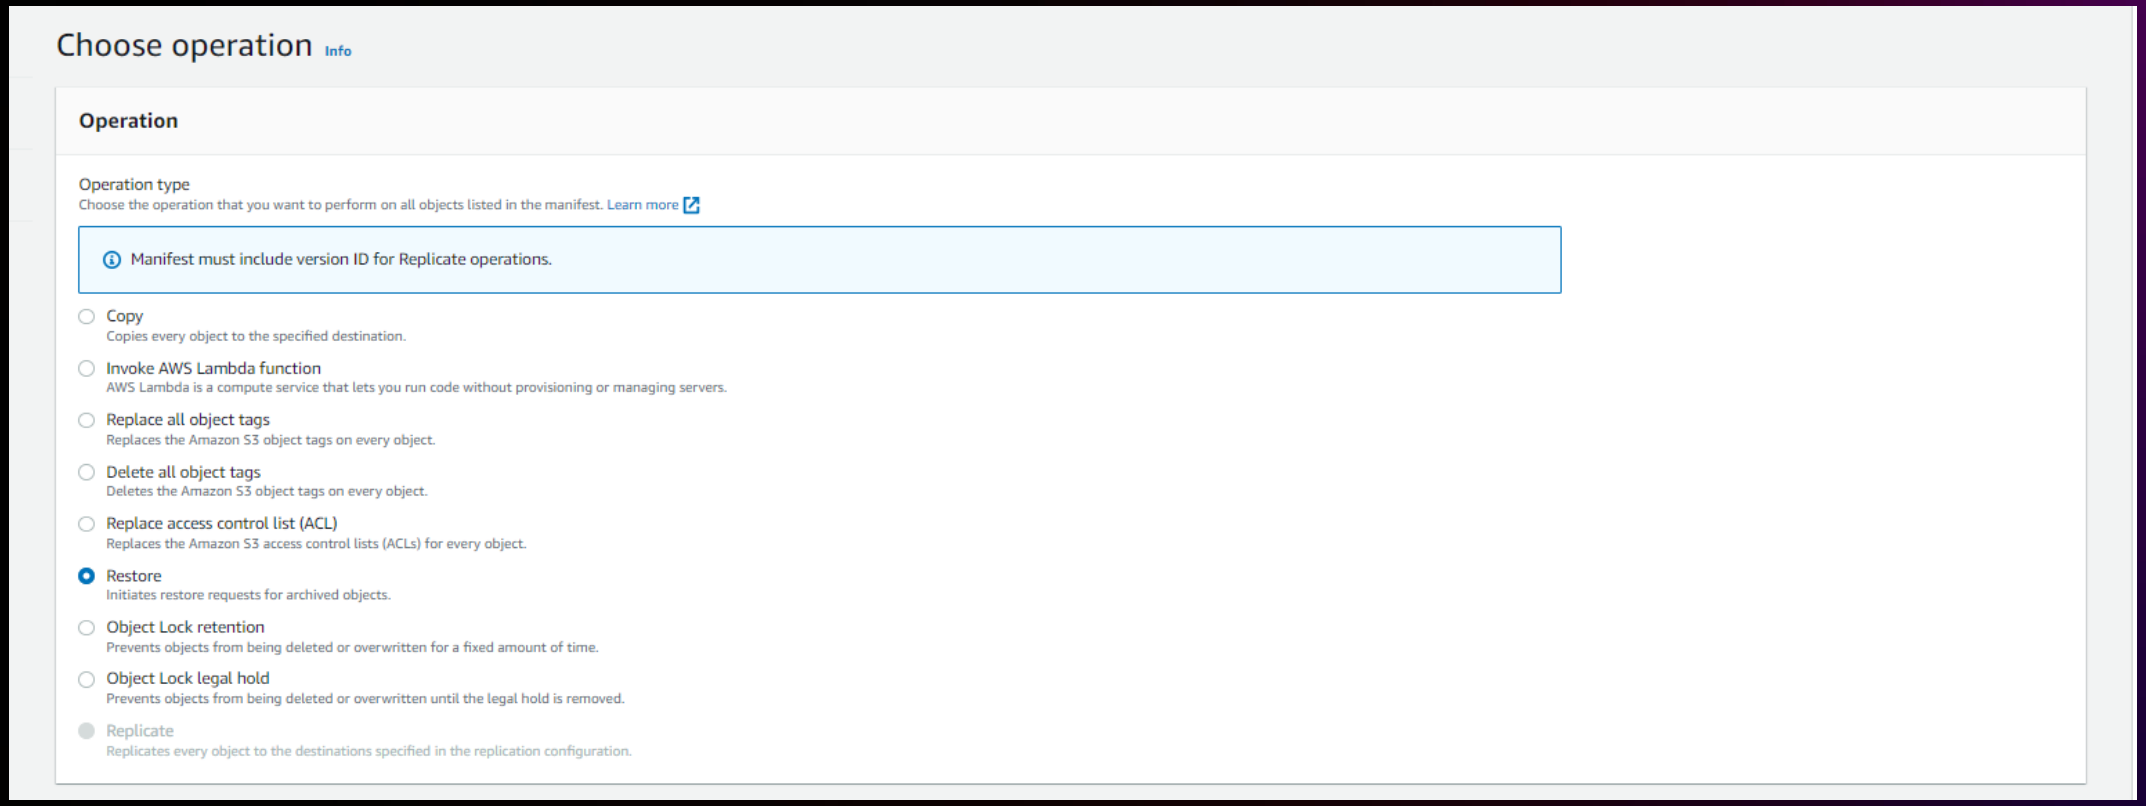

Use S3 Batch Operations to restore your data from S3 Glacier Deep Archive

3. Select restore configurations

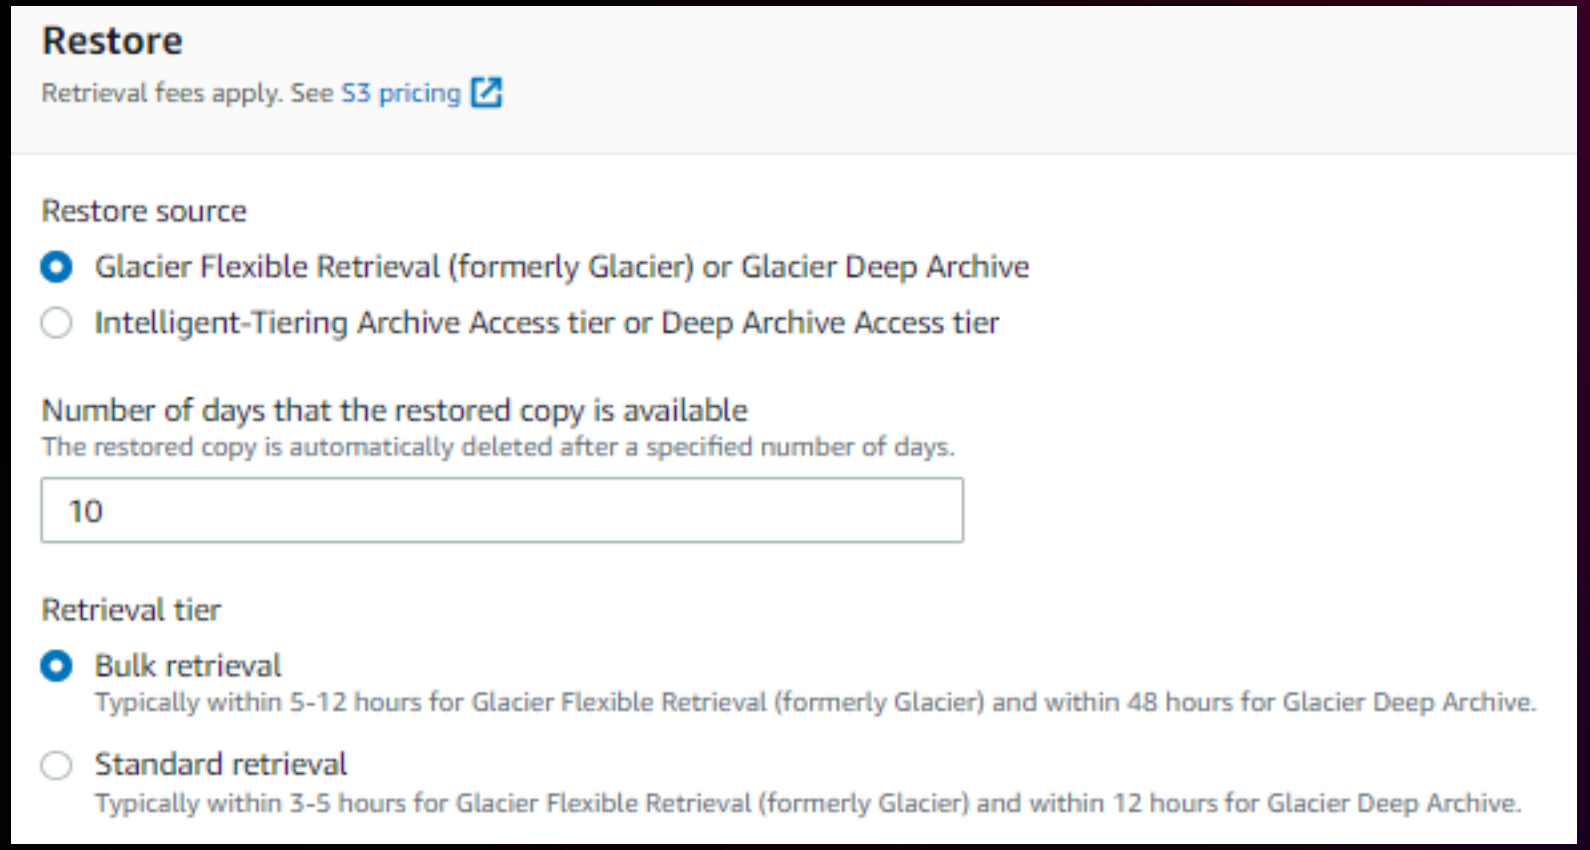

Use S3 Batch Operations to restore your data from S3 Glacier Deep Archive

4. Select additional configurations and create job

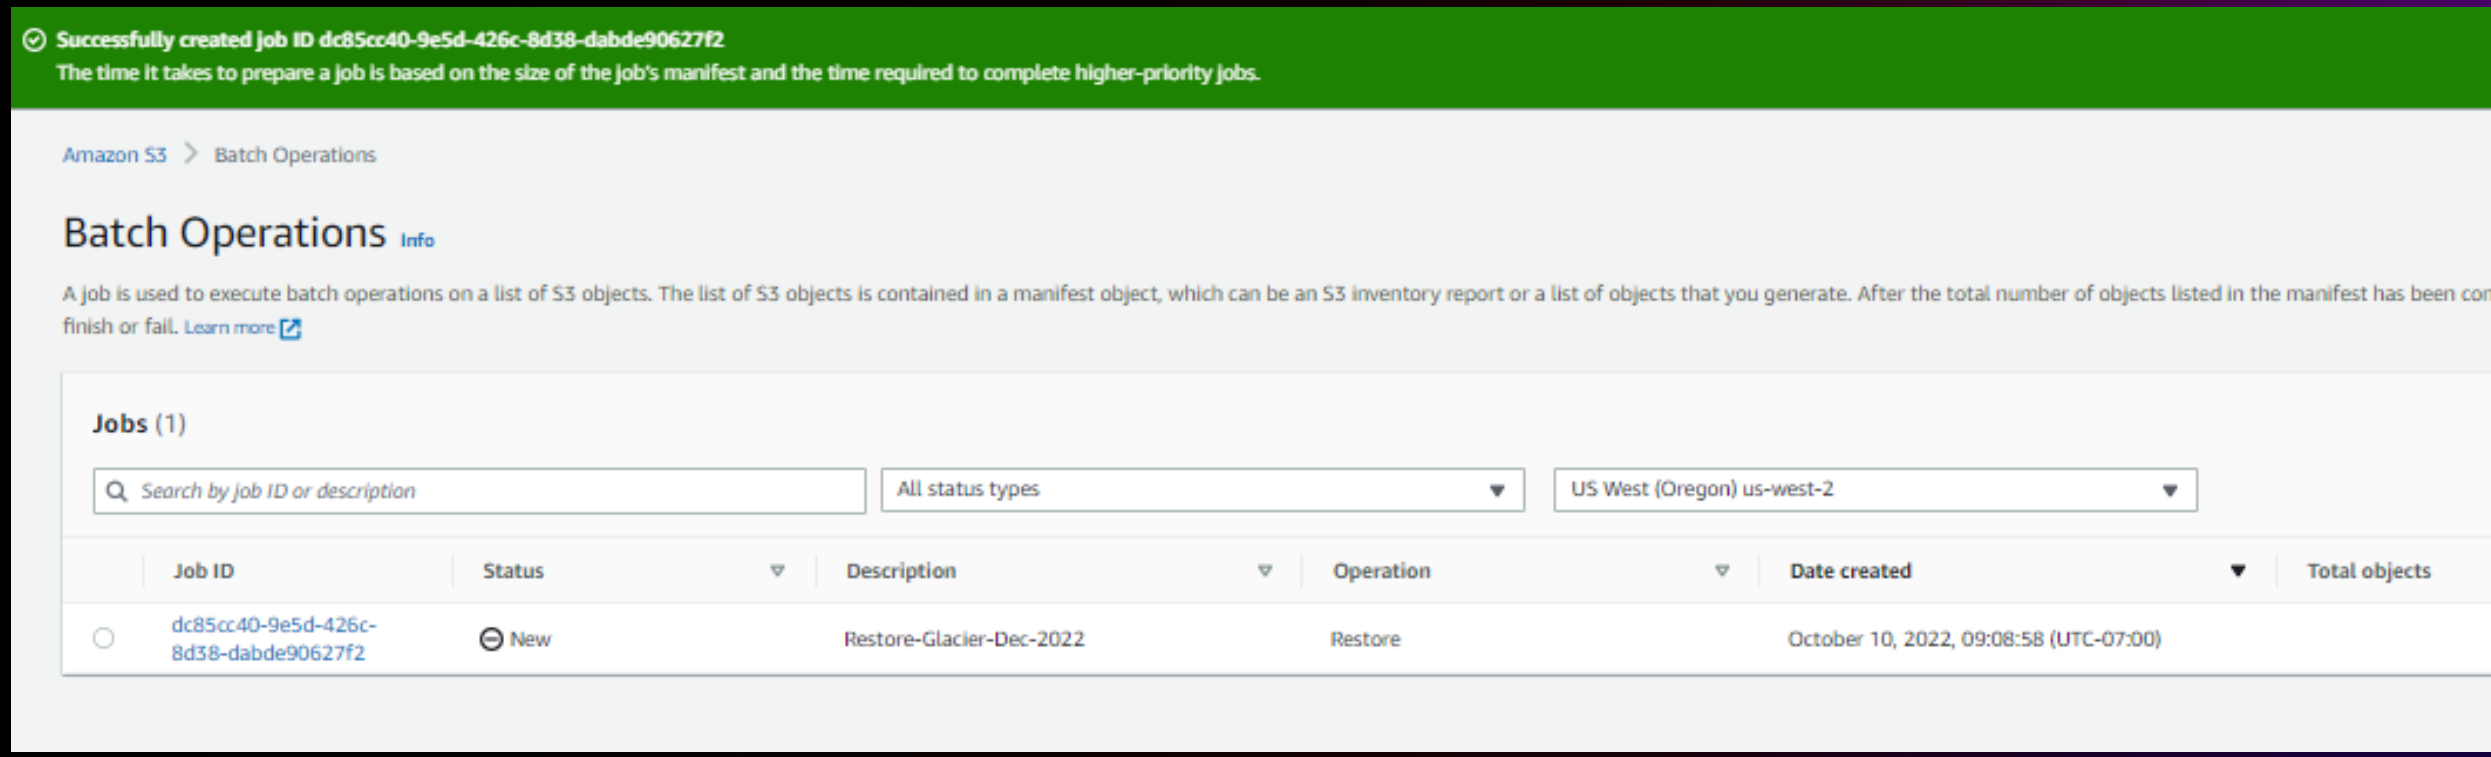

Activate S3 Event Notifications for restores in your bucket properties

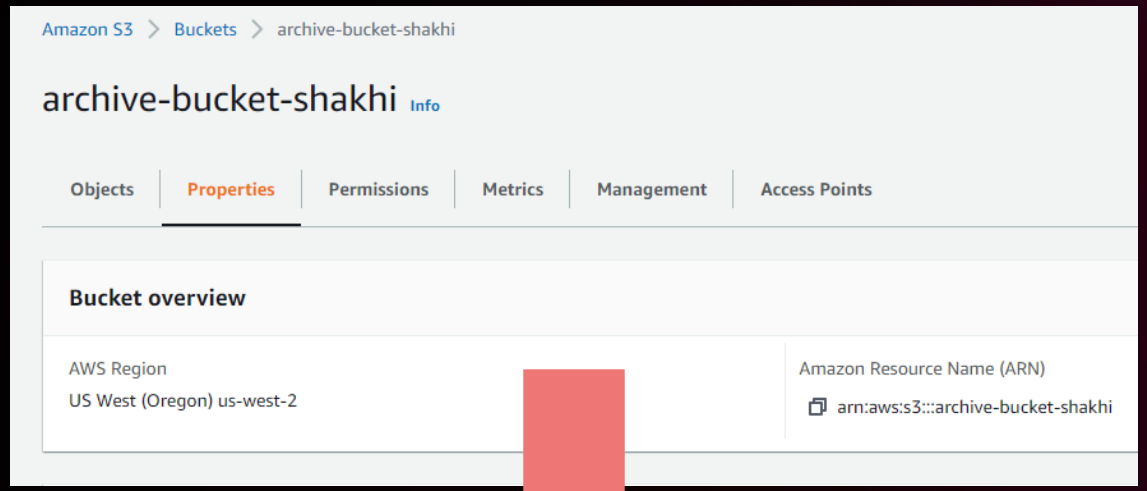

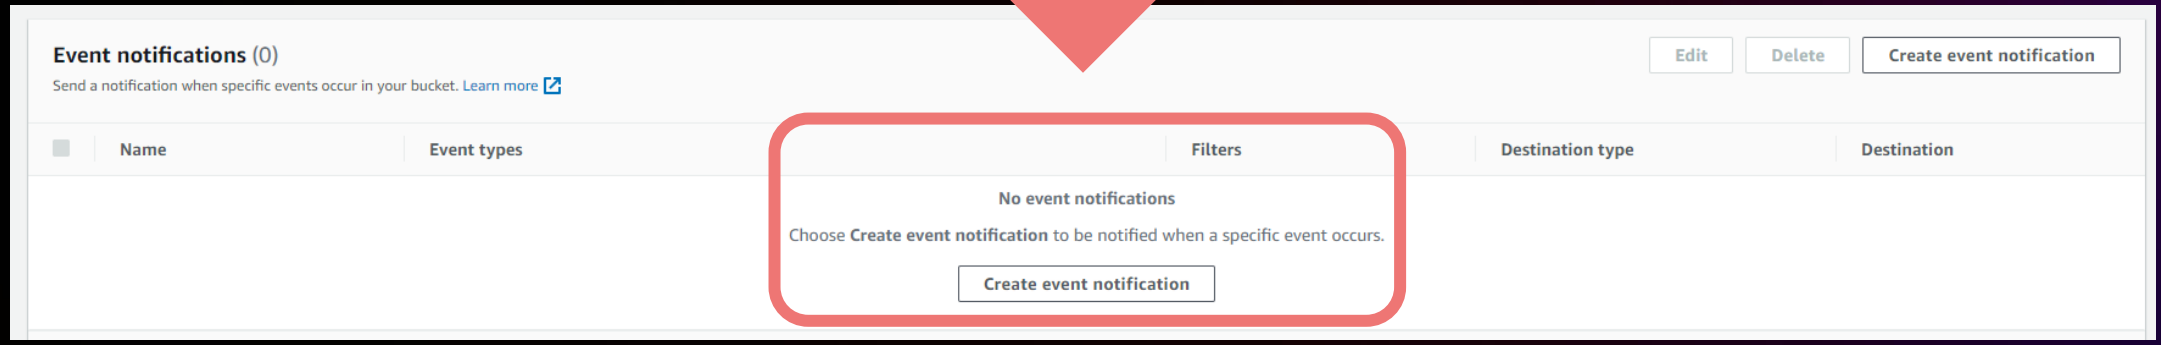

Set up S3 Event Notification configuration and destination

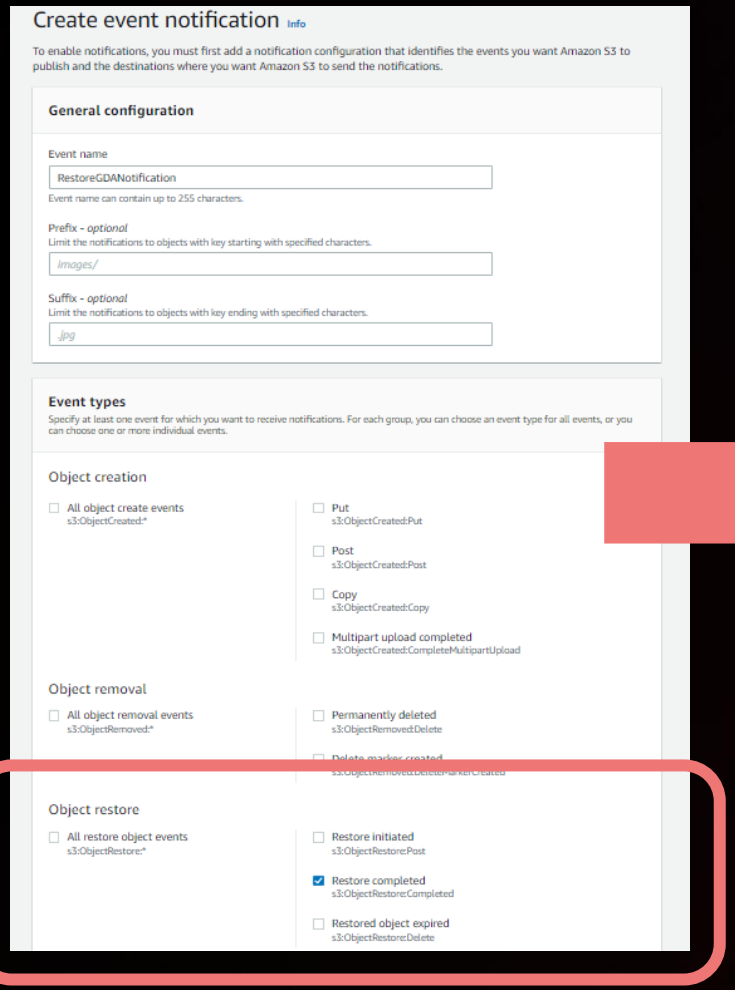

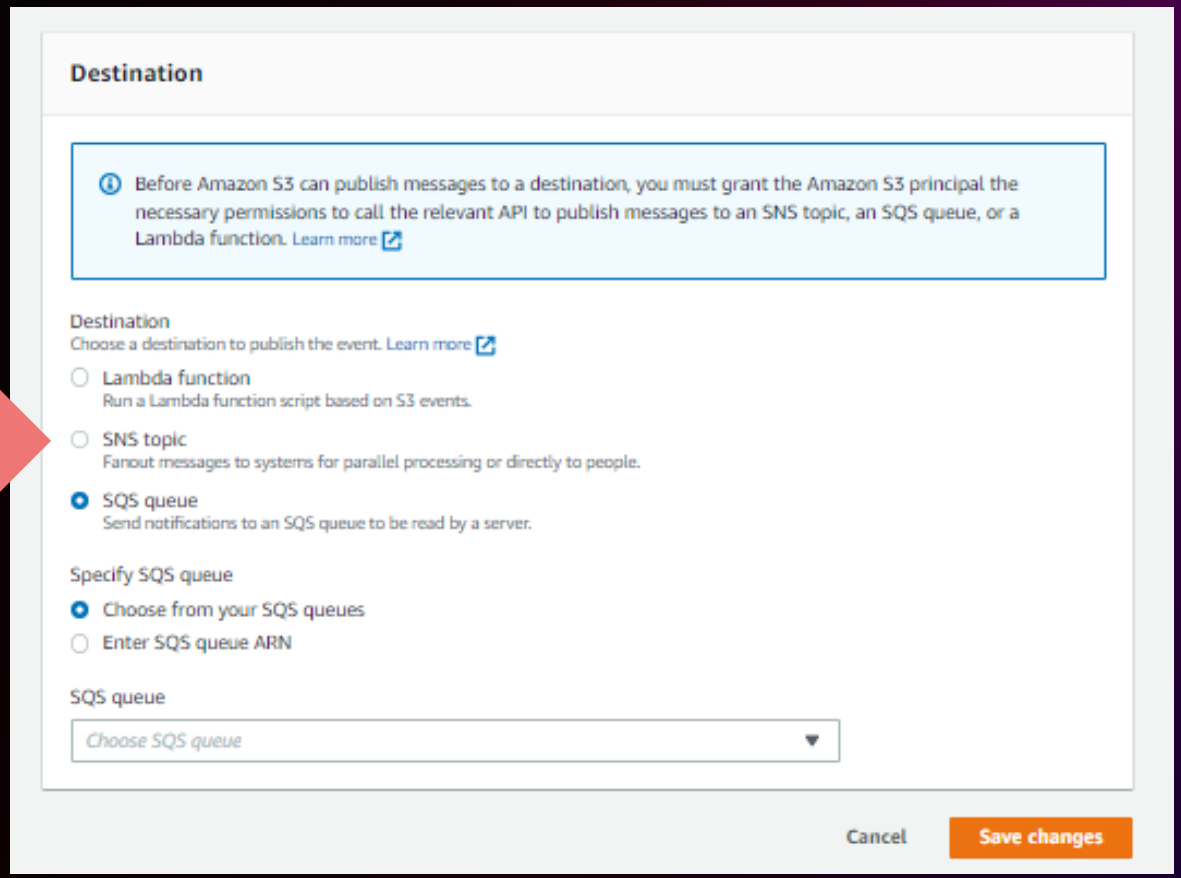

Monitoring access:
Querying request logs

Overview: Insights features in Amazon S3

S3 Storage Lens

Organization-wide visibility into pre-aggregated usage and activity metrics

Free and paid versions available

S3 Inventory

Object-level metadata

S3 storage class analysis

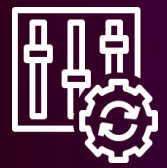

Access patterns by bucket, prefix, tags

S3 CloudWatch Metrics

Monitoring and alarms

S3 server access logs

Detailed request logging

Amazon S3 Storage Lens

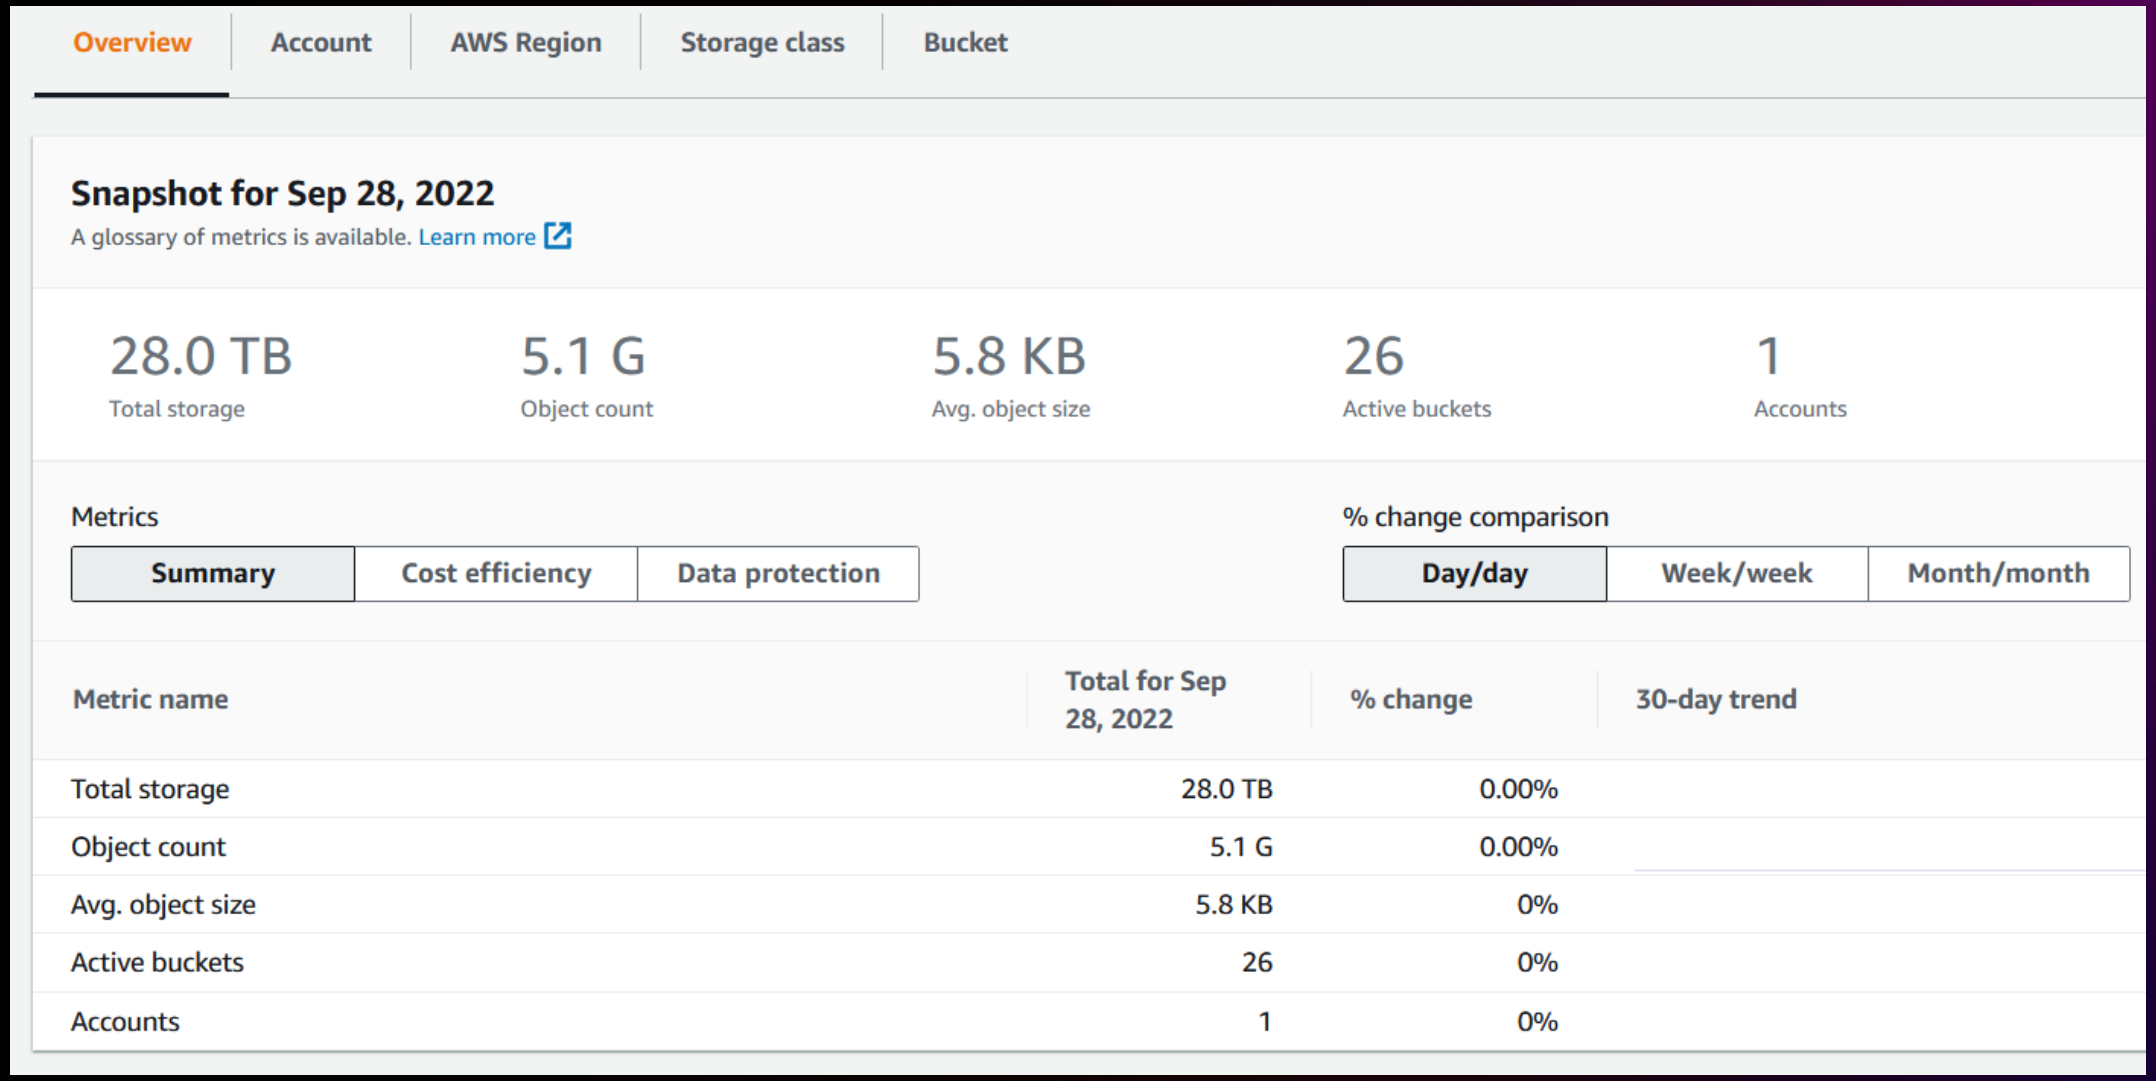

Amazon CloudWatch S3 request metrics

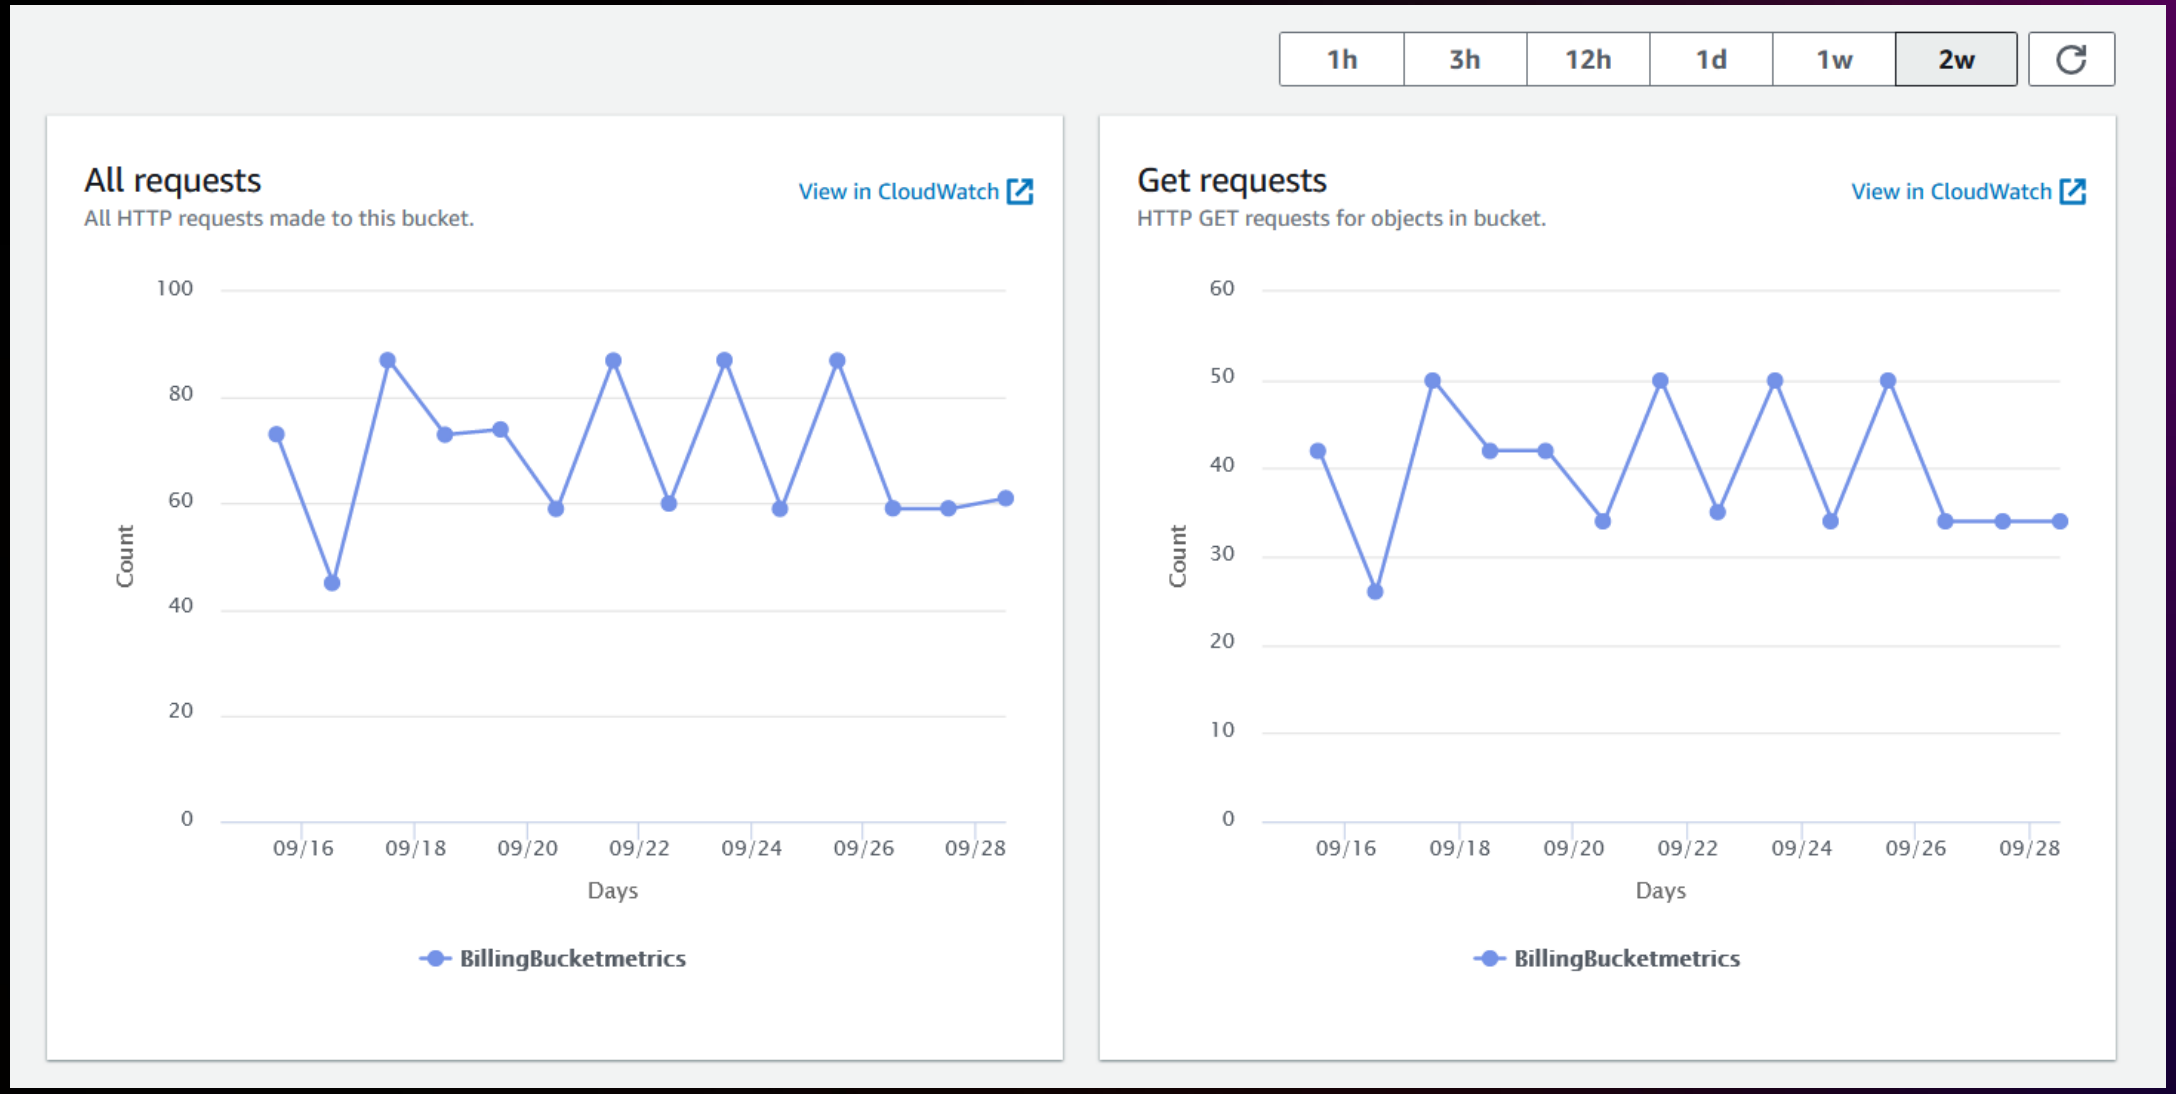

Amazon S3 storage class analysis

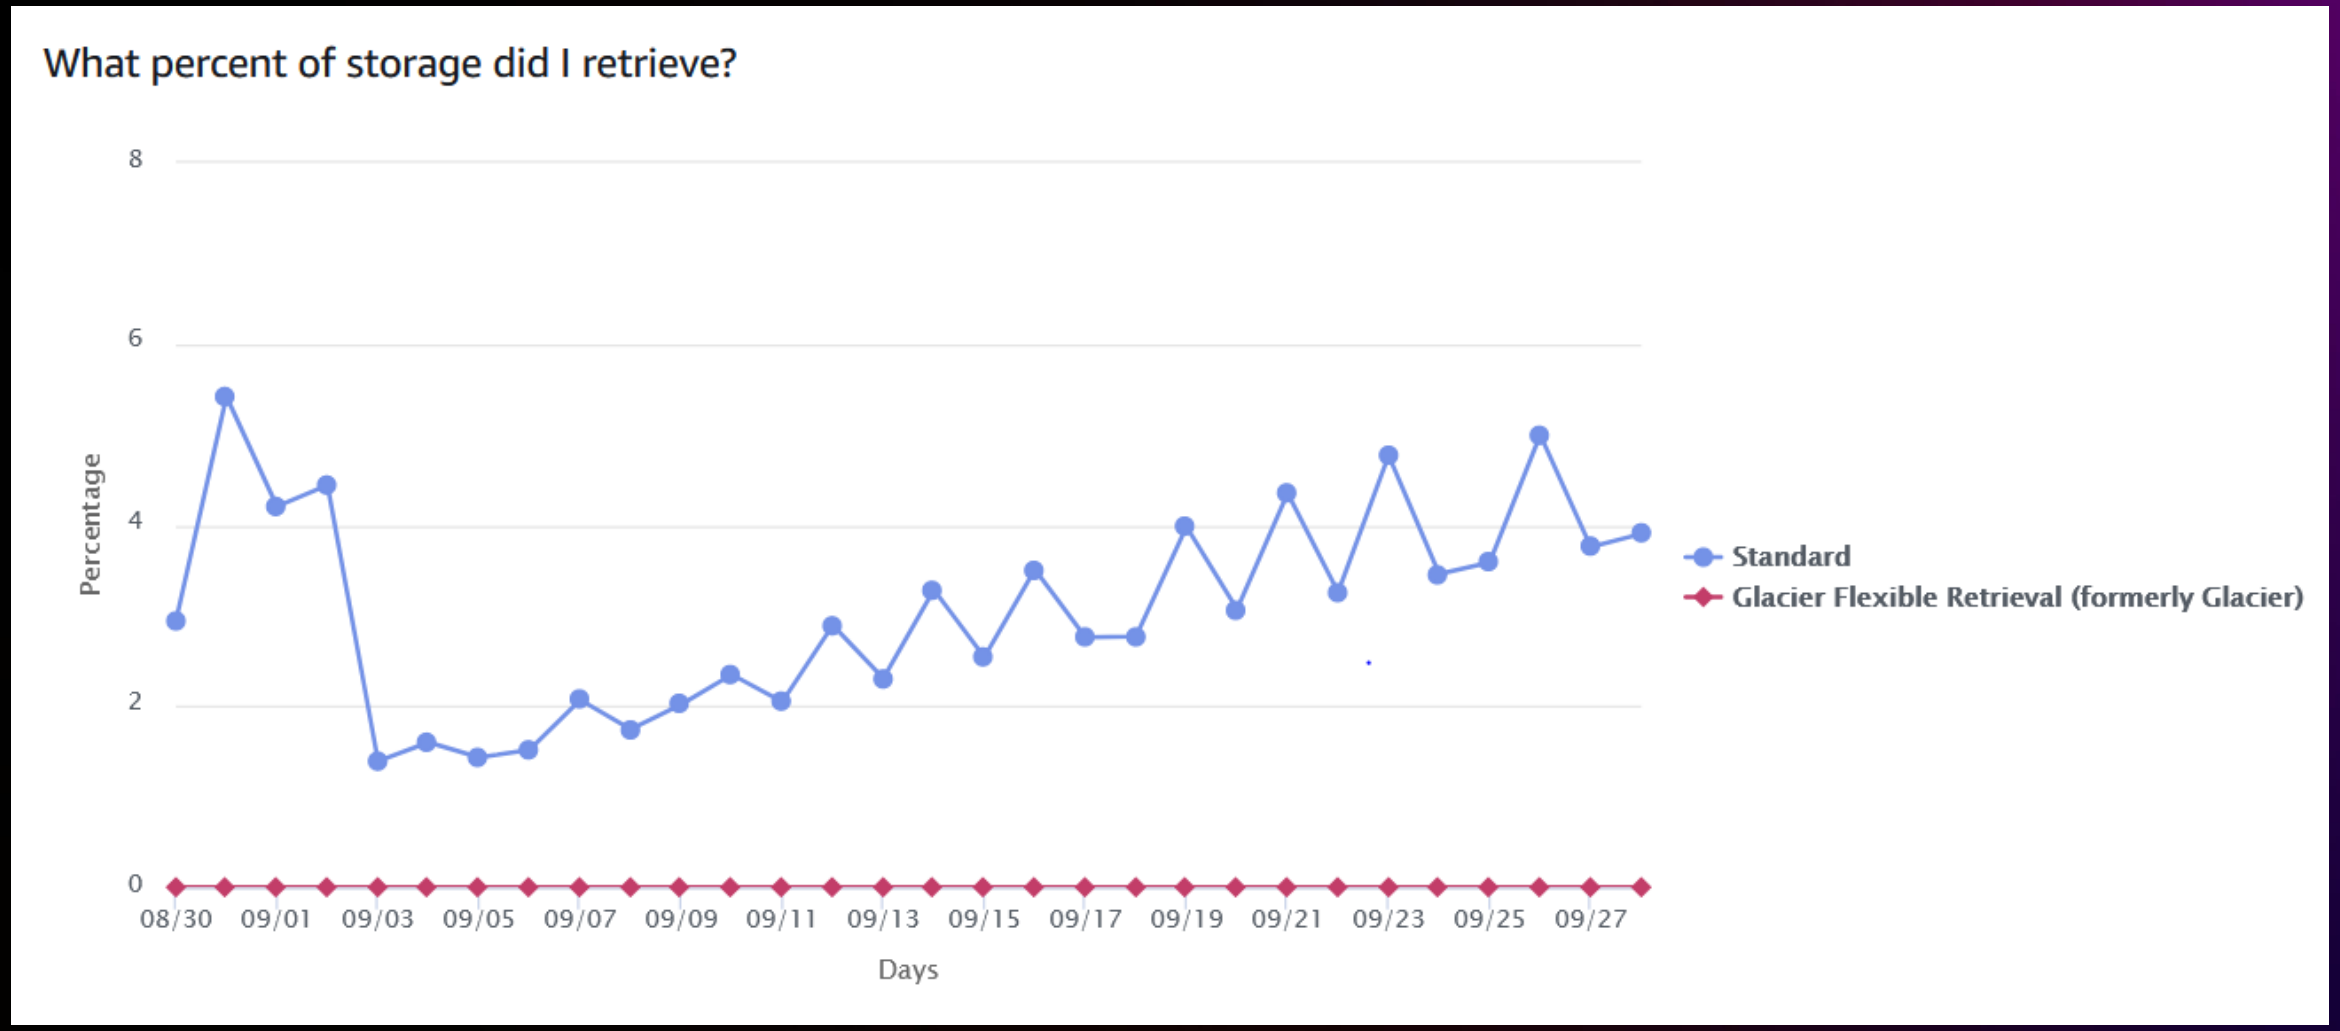

Overview: Options for Amazon S3 logging

Overview: Options for Amazon S3 logging

Amazon S3 server access logs

Requests to a bucket

Space-delineated format

Pay for log storage

Logs for S3 Lifecycle actions

Overview: Options for Amazon S3 logging

Amazon S3 server access logs

Requests to a bucket

Space-delineated format

Pay for log storage

Logs for S3 Lifecycle actions

AWS CloudTrail

Management and data events

JSON format

Pay for log storage, events delivered

Forward to CloudWatch Logs, CloudWatch Events

AWS CloudTrail: Amazon S3 data event


```
"eventVersion": "1.03",
        "userIdentity": {
            "type": "IAMUser",
            "principalId": "111122223333",
            "arn": "arn:aws:iam::111122223333:user/myUserName",
            "accountId": "111122223333",
            "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
            "userName": "myUserName"
        },
        "eventTime": "2019-02-01T03:18:19Z",
        "eventSource": "s3.amazonaws.com",
        "eventName": "ListBuckets",
        "awsRegion": "us-west-2",
        "sourceIPAddress": "127.0.0.1",
        "userAgent": "[]",
        "requestParameters": {
            "host": [
                "s3.us-west-2.amazonaws.com"
            ]
```

Information shown in S3 server access logs

Information shown in S3 server access logs

Bucket owner
Bucket
Time
Remote IP
Requester
Request ID
Operation
Key
Request URI
HTTP status
Error code
Bytes sent
Object size
Total time
Turn-around time
Referrer
User-agent
Version id
Host ID
Signature version
Cipher suite
Authentication type
Host header
TLS version
Access point ARN

Customer use case 3

Phone rings . . .

Colleague calling . . .

Needs help unblocking a demo . . .

Colleague needs detailed error information on requests made to Amazon S3

AWS CloudTrail: S3 data events
Amazon Athena

Common log analyses

1. Finding a specific request ID or error message

Common log analyses

1. Finding a specific request ID or error message

2. Analyzing S3 request latency as you make changes to an application

Common log analyses

1. Finding a specific request ID or error message

2. Analyzing S3 request latency as you make changes to an application

3. Understanding who is accessing data in a bucket

Common log analyses

1. Finding a specific request ID or error message
2. Analyzing S3 request latency as you make changes to an application
3. Understanding who is accessing data in a bucket
4. Investigating access patterns in the bucket

Log-query approach using AWS CloudTrail

Key steps

1. Create a trail in AWS CloudTrail
   a) Specify S3 location for your logs
2. Enable AWS CloudTrail: S3 data events
   a) Select all buckets or filter your data events
3. Set up an Amazon Athena table
4. Run queries

Create a trail

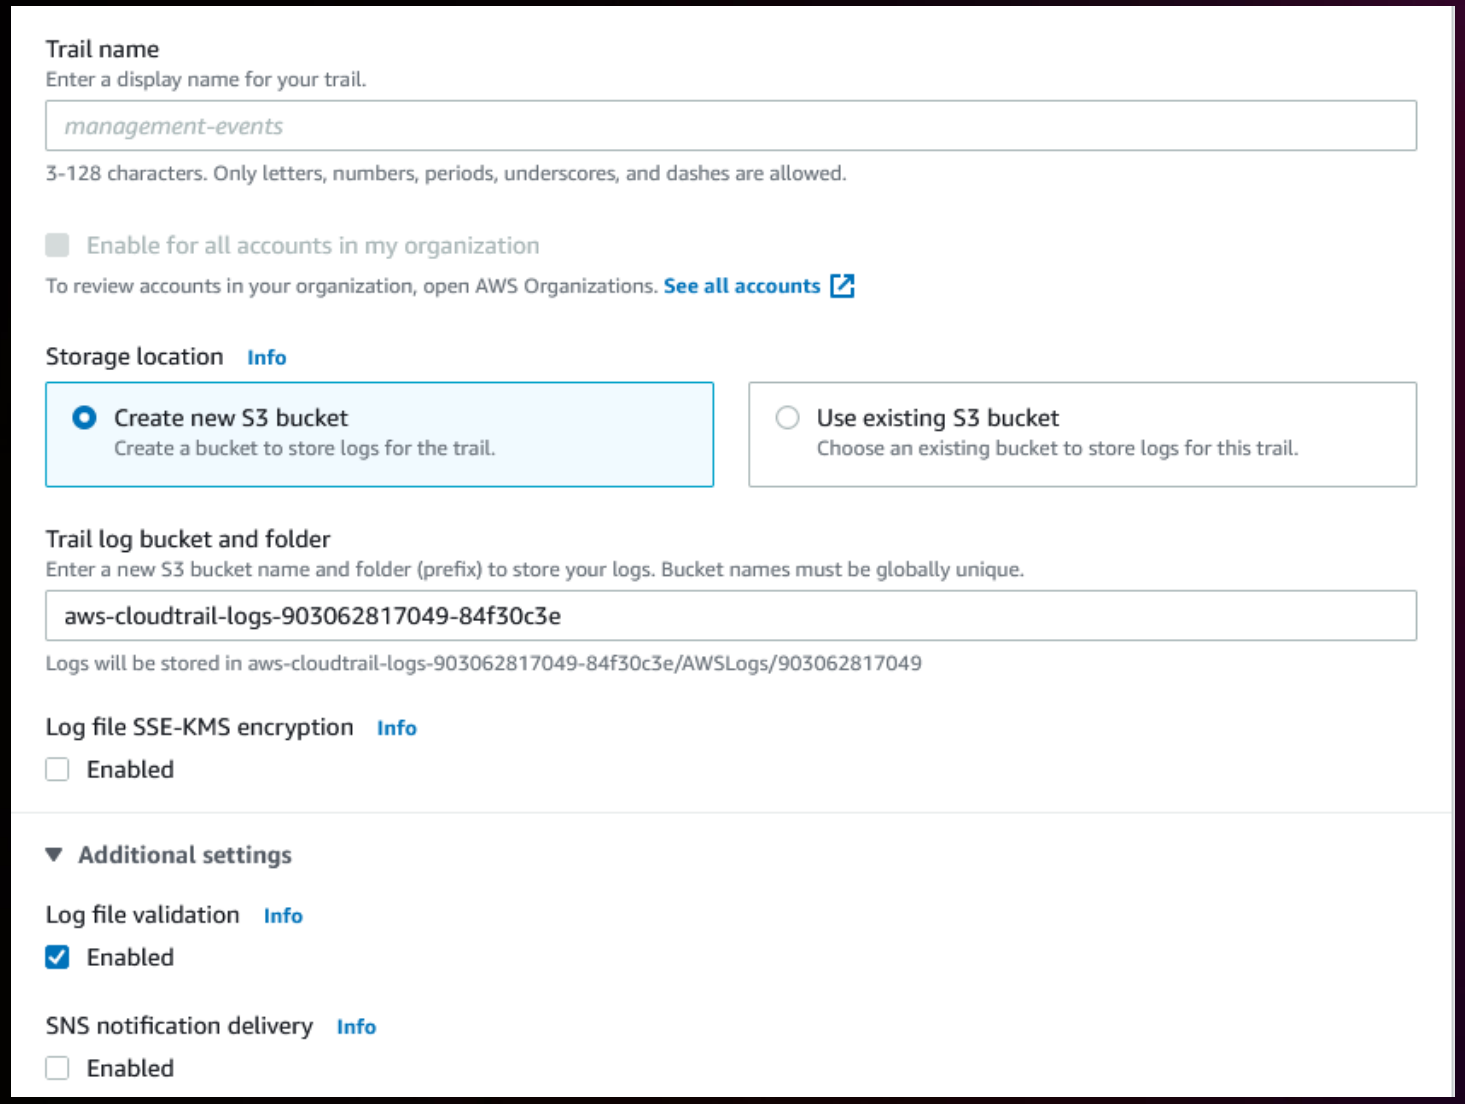

Specify AWS CloudTrail data events

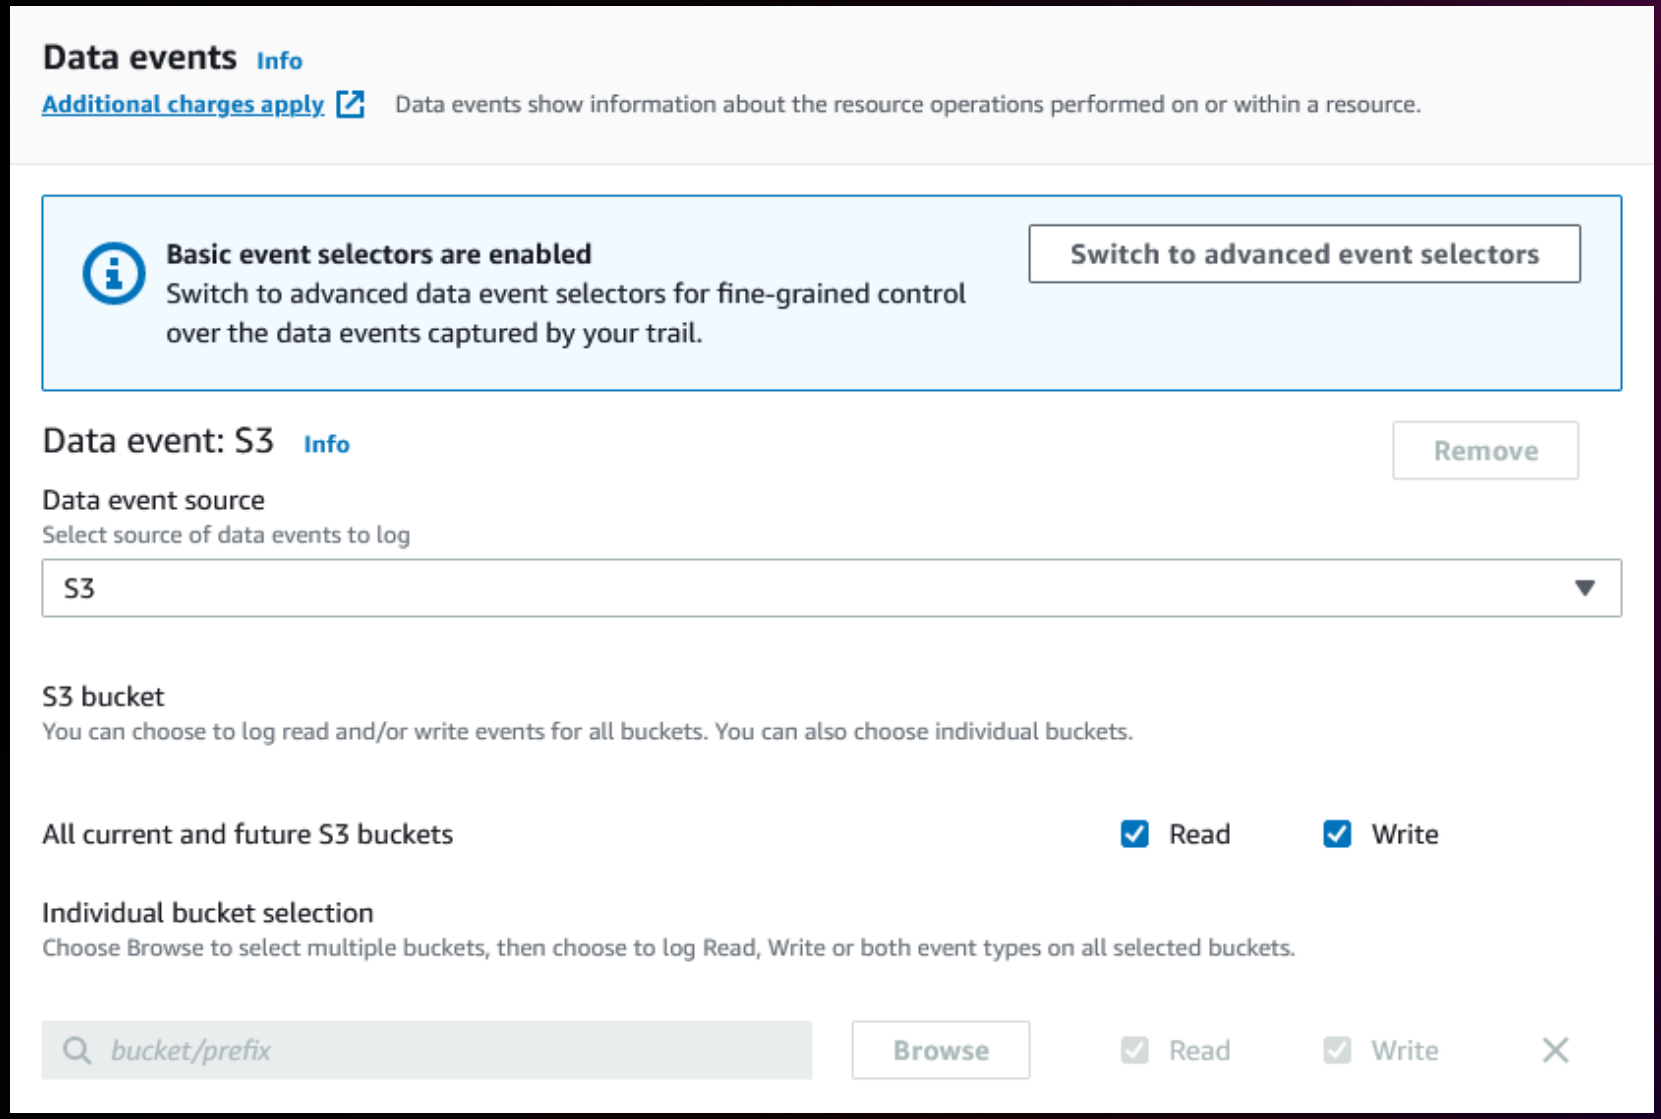

Create a table in Amazon Athena

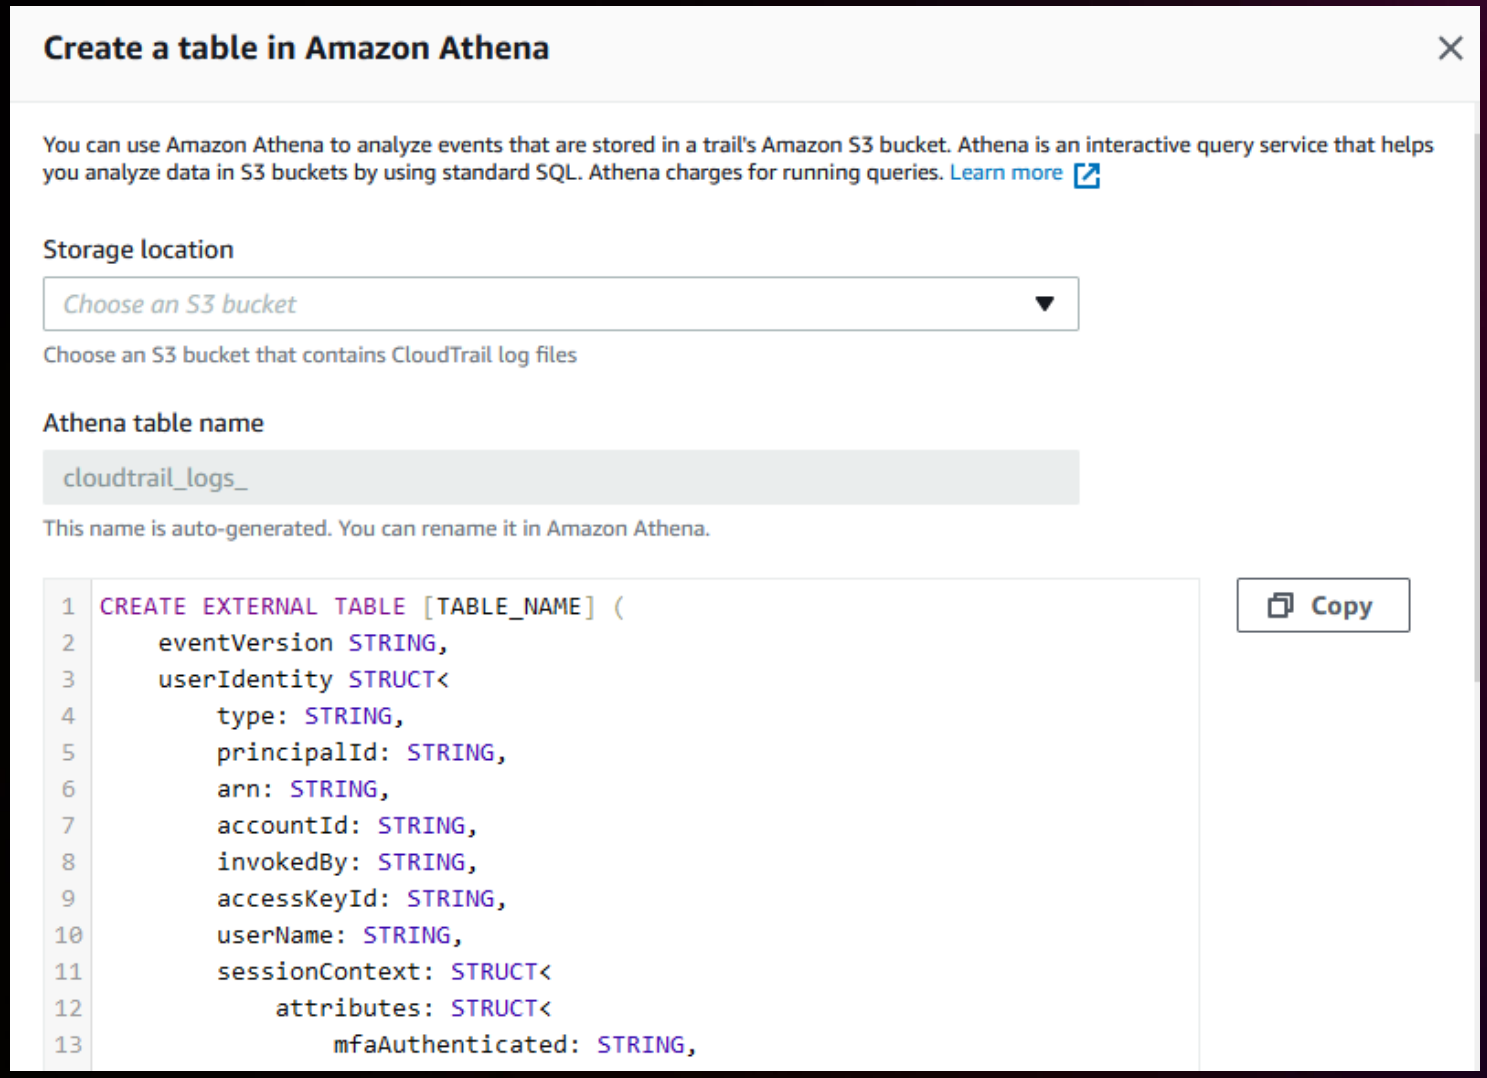

Run query

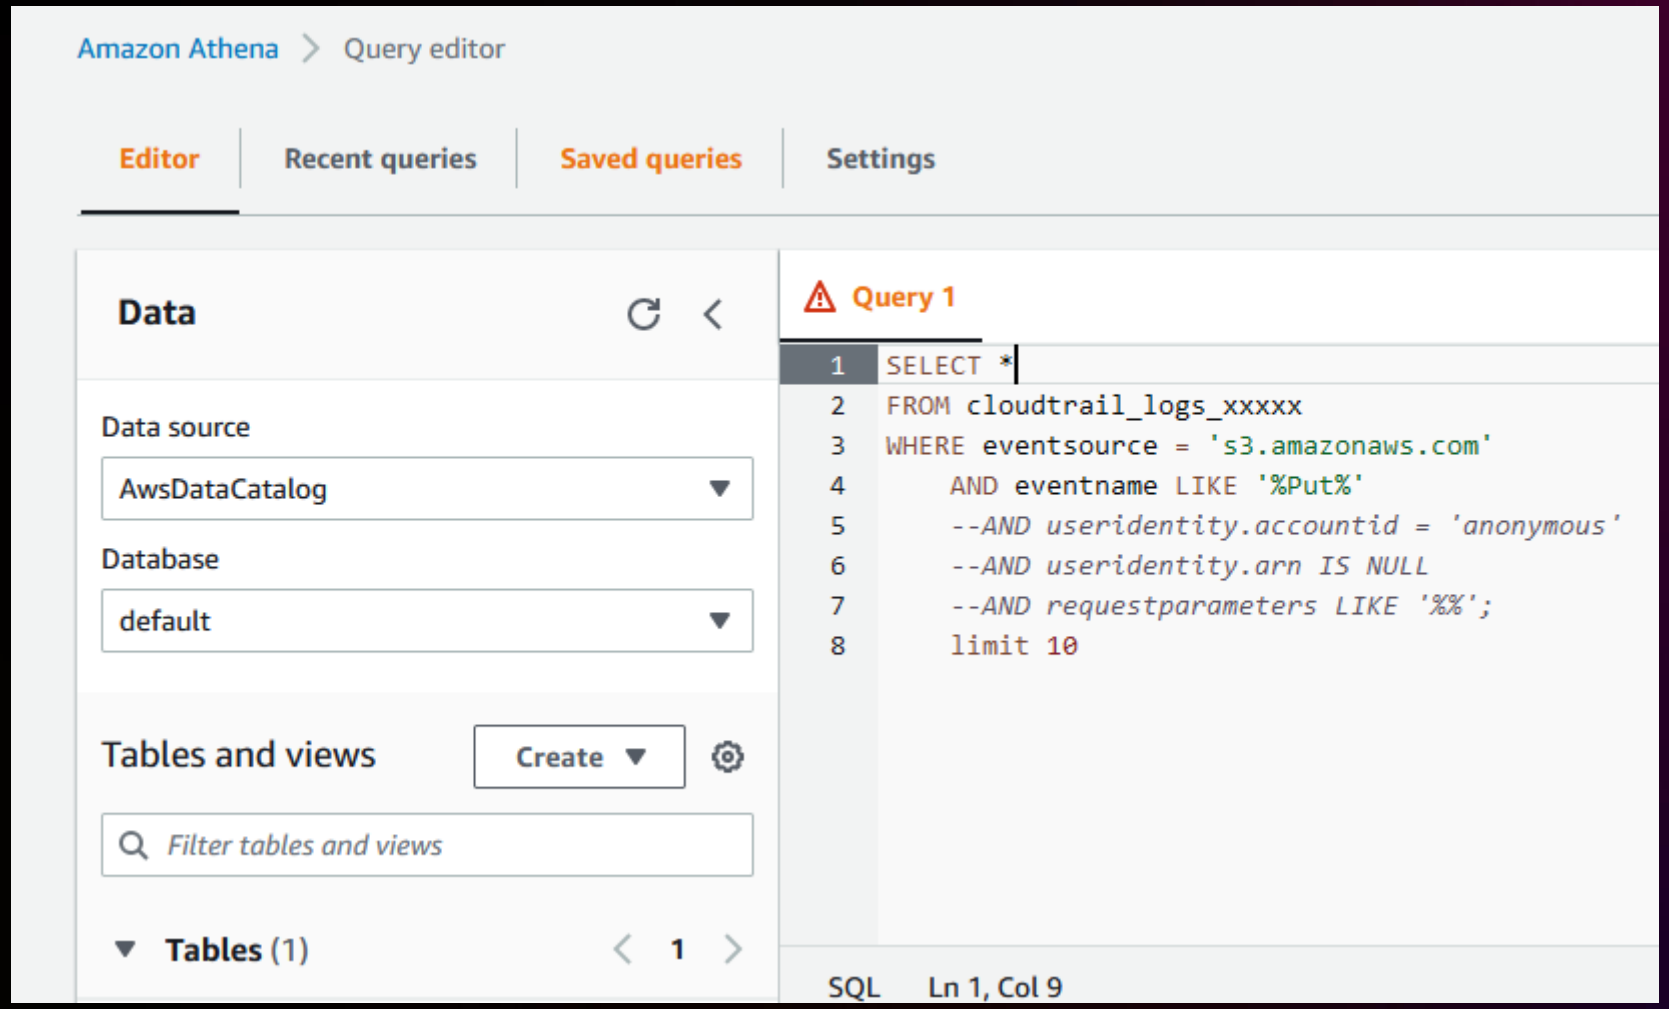

Query example

```
SELECT
    awsregion,
    replace(unnested.resources_entry.ARN,'arn:aws:s3:::') as s3_resource,
    eventname,
    eventtime,
    useragent
FROM cloudtrail_logs t
CROSS JOIN UNNEST(t.resources) unnested (resources_entry)
WHERE unnested.resources_entry.ARN LIKE '%example/datafile.txt'
ORDER BY eventtime
```

Query example

```
SELECT
       useridentity.arn,
       Count(requestid) as RequestCount
FROM s3_cloudtrail_events_db.cloudtrail_awsexamplebucket_table
WHERE eventsource='s3.amazonaws.com'
and json_extract_scalar(additionalEventData, '$.SignatureVersion')='SigV2'
Group by useridentity.arn
```

Query results

eventversion	eventtime	eventsource	eventname	awsregion	sourceipaddress
1.05	2017-10-01T20:34:50Z	s3.amazonaws.com	GetBucketLocation	ap-northeast-1	52.xxx.xxx.xxx
1.05	2017-10-01T20:34:53Z	s3.amazonaws.com	GetBucketLocation	ap-northeast-1	52.xxx.xxx.xxx
1.05	2017-10-01T20:34:52Z	s3.amazonaws.com	GetBucketLocation	ap-northeast-1	52.xxx.xxx.xxx
1.05	2017-10-01T20:34:53Z	s3.amazonaws.com	GetBucketLocation	ap-northeast-1	52.xxx.xxx.xxx
1.05	2017-10-01T20:34:52Z	s3.amazonaws.com	GetBucketLocation	ap-northeast-1	52.xxx.xxx.xxx
1.05	2017-10-01T20:34:51Z	s3.amazonaws.com	GetBucketLocation	ap-northeast-1	52.xxx.xxx.xxx
1.05	2017-10-01T20:34:51Z	s3.amazonaws.com	GetBucketLocation	ap-northeast-1	52.xxx.xxx.xxx
1.05	2017-10-01T20:34:51Z	s3.amazonaws.com	GetBucketLocation	ap-northeast-1	52.xxx.xxx.xxx
1.05	2017-10-01T20:34:52Z	s3.amazonaws.com	GetBucketLocation	ap-northeast-1	52.xxx.xxx.xxx
1.04	2017-08-13T22:15:22Z	s3.amazonaws.com	GetBucketLocation	ap-northeast-1	52.xxx.xxx.xxx

Visualizing log data: CloudWatch Insights

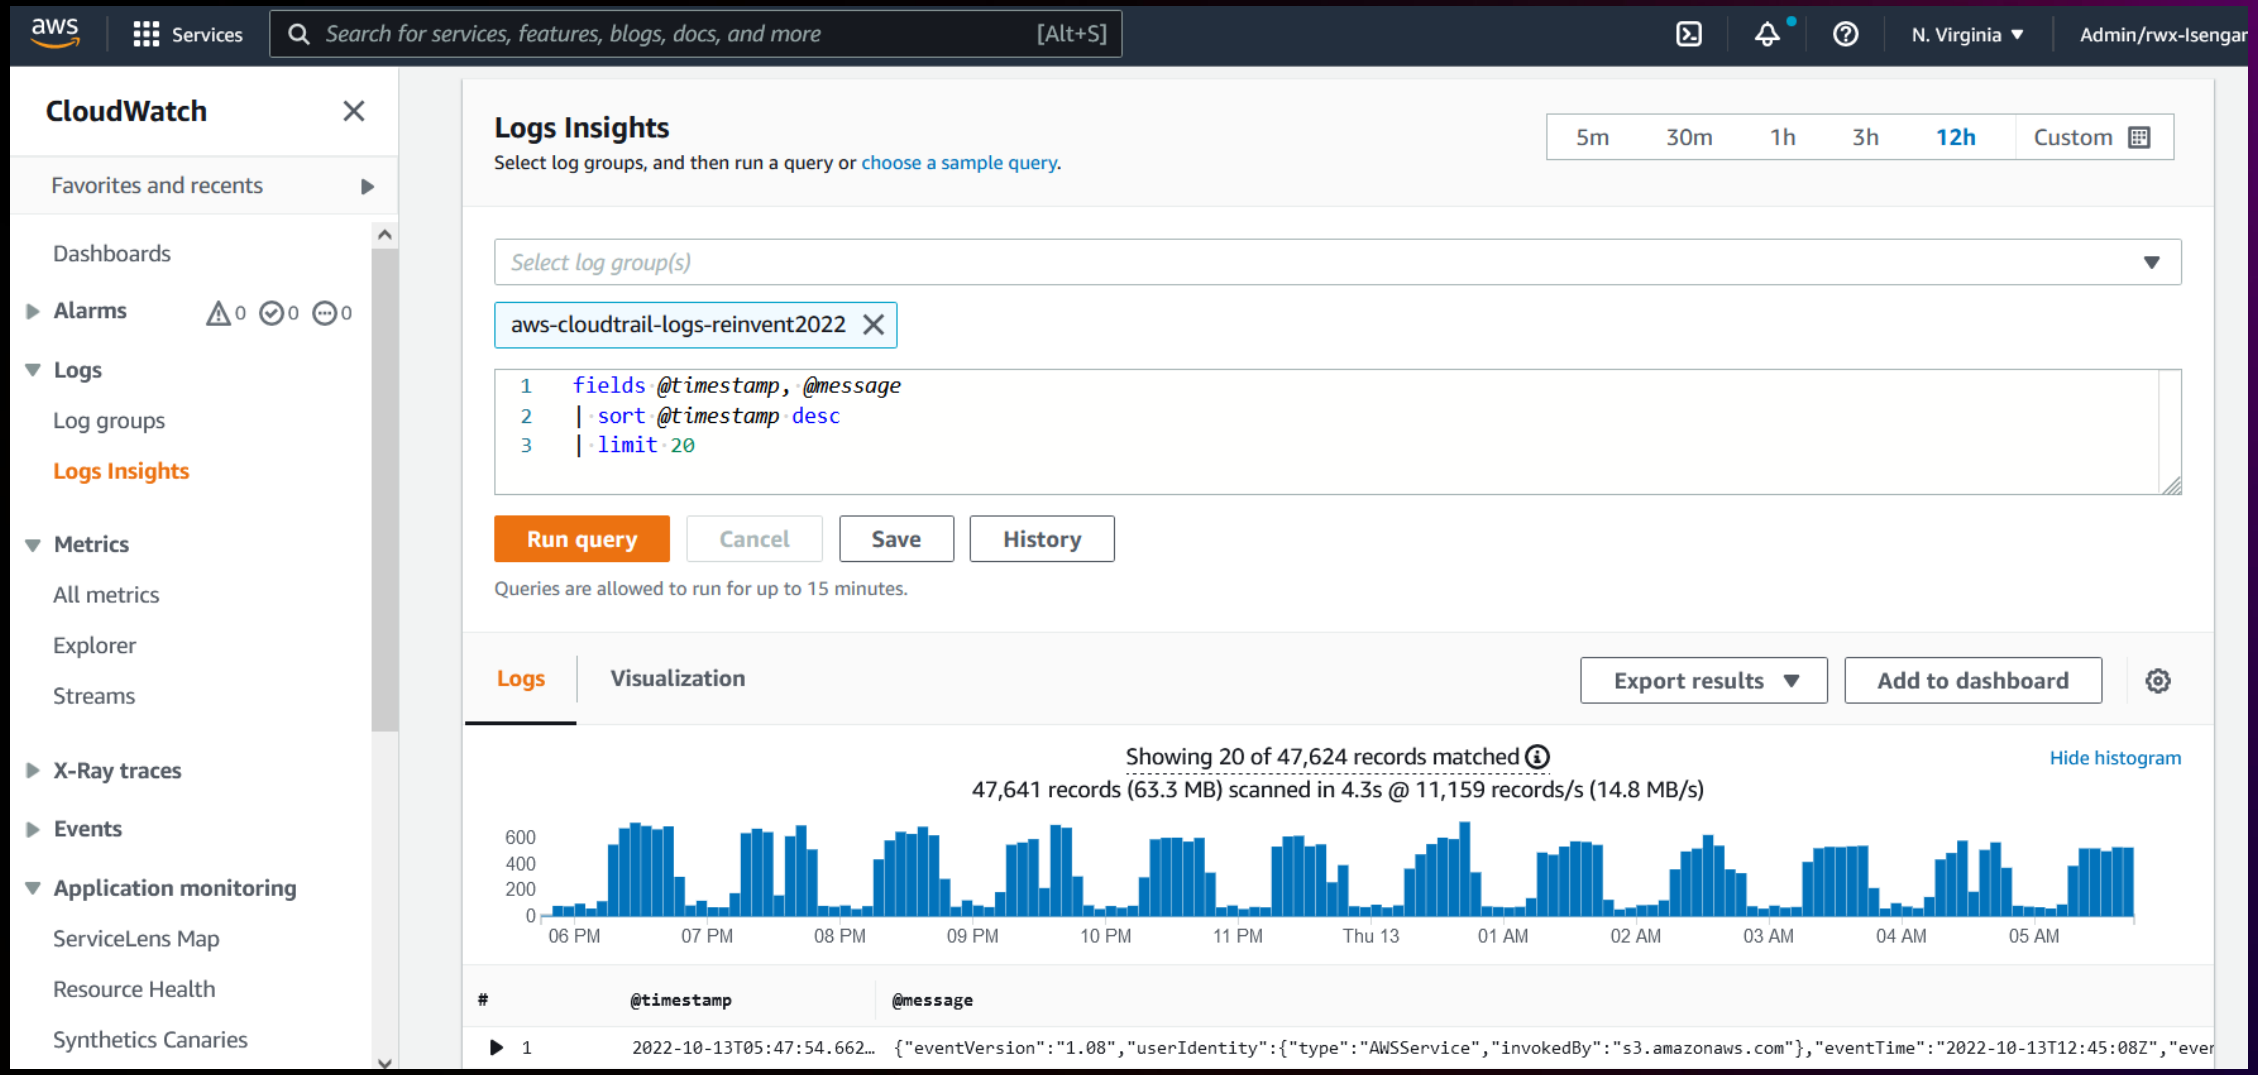

Managing storage spend

Amazon S3 Lifecycle helps optimize storage spend

Transition actions: Define when objects transition to other Amazon S3 storage classes as they age

Expiration actions: Define when objects expire; Amazon S3 deletes expired objects on your behalf

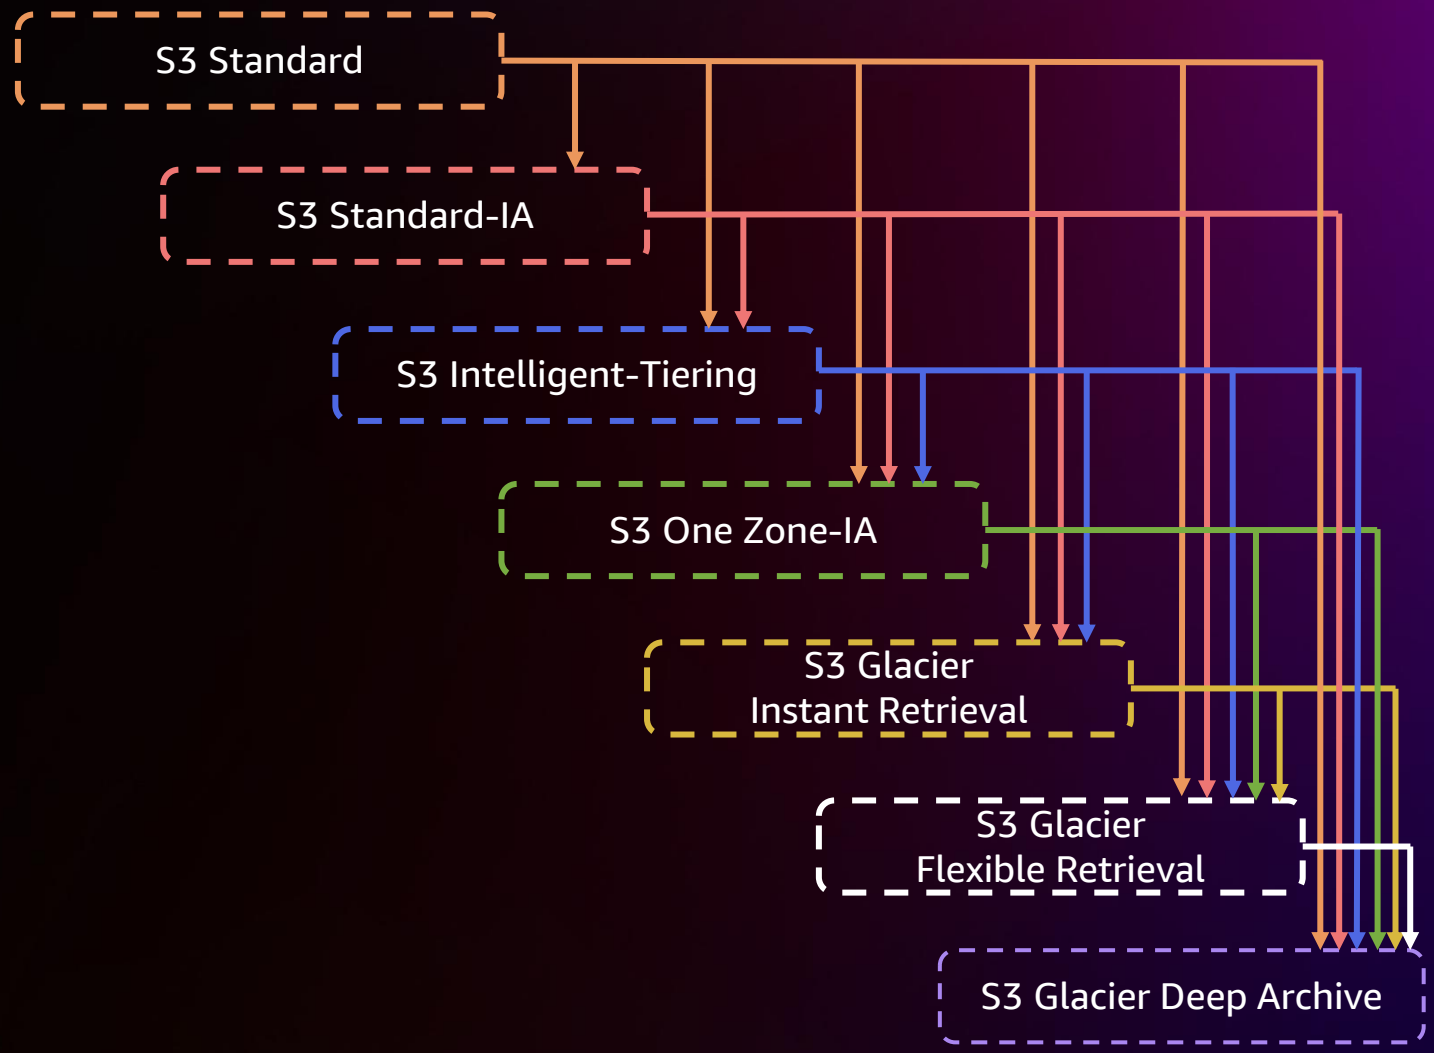

Applying insights to create S3 Lifecycle configurations

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

Anatomy of an S3 Lifecycle configuration

Lifecycle configuration elements

- **ID**
- **Filters**
- **Status**
- **Actions**

```
<LifecycleConfiguration>
  <Rule>
    <ID>ExampleRule</ID>
    <Filter>
      <Prefix>tax/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Transition>
      <Days>365</Days>
      <StorageClass>GLACIER</StorageClass
    </Transition>
    <Expiration>
      <Days>3650</Days>
    </Expiration>
  </Rule>
</LifecycleConfiguration>
```

Use object versioning to protect your objects

Versioning prevents overwrites by creating new versions of the object

Use object versioning to protect your objects

Versioning prevents overwrites by creating new versions of the object

exampleAWSbucket/drafts/novel

Version A

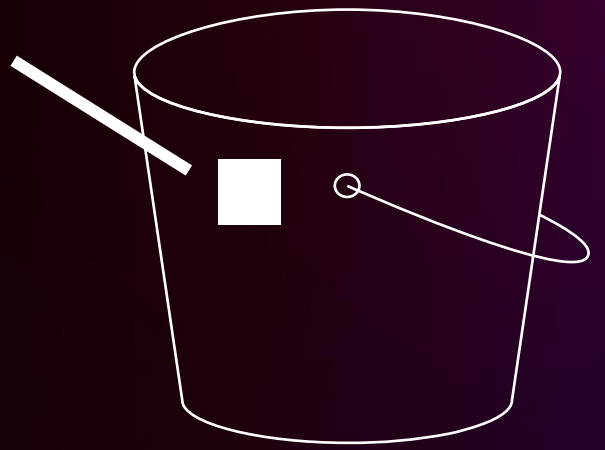

Use object versioning to protect your objects

Versioning prevents overwrites by creating new versions of the object

exampleAWSbucket/drafts/novel Version A
exampleAWSbucket/drafts/novel Version B

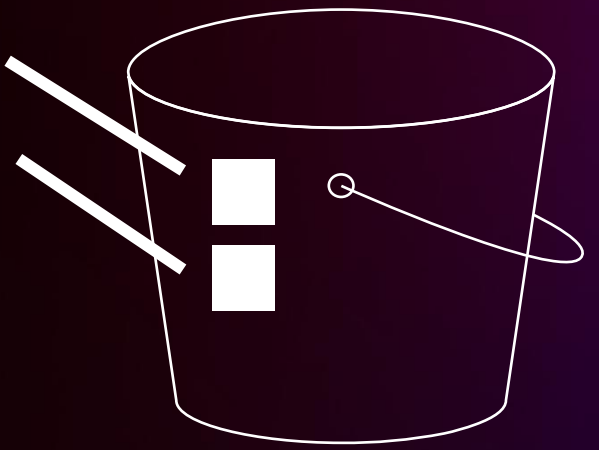

Use object versioning to protect your objects

Versioning prevents overwrites by creating new versions of the object

exampleAWSbucket/drafts/novel Version A
exampleAWSbucket/drafts/novel Version B
exampleAWSbucket/drafts/novel Version C

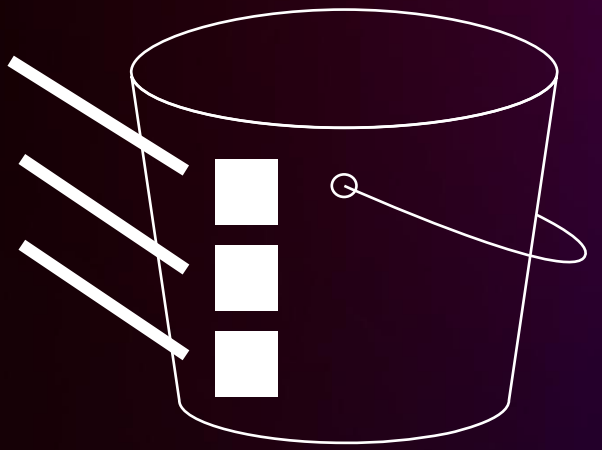

Versioning also provides deletion protection

Versioning can also provide deletion protection

- Deletions without a version ID create a delete marker
- The prior version of the object is retained
- GET requests return a 404, NOT FOUND error

When the version ID is specified in a deletion, that version is deleted

Best practice: Consider an S3 Lifecycle configuration for all versioned buckets

Set retention policy with Amazon S3 Lifecycle

Current version S3 Lifecycle rule

Noncurrent version S3 Lifecycle rule

Expired object delete marker rule

Incomplete multipart upload retention

Customer use case 4

Email from analyst on the finance team . . .

Meeting underway, many questions about month-over-month changes . . .

Why did storage spend on AWS rise 15% last month in our testing account?

Amazon S3 Storage Lens
Amazon S3 Lifecycle

Let’s start with S3 Storage Lens

Top 3 regions

AWS Region	Total storage	% of total	% change
US East (N. Virginia)	15.4 TB	54.89%	0.00%
US West (Oregon)	12.1 TB	43.26%	-0.00%
US East (Ohio)	528.5 GB	1.85%	0.00%

Top 3 buckets

Bucket	Total storage	% of total	% change
rw	8.9 TB	31.92%	14.95%
rw	7.1 TB	25.53%	0%
rw	5.7 TB	20.24%	0.00%

Let’s start with S3 Storage Lens

Top 3 regions

AWS Region	Total storage	% of total	% change
US East (N. Virginia)	15.4 TB	54.89%	0.00%
US West (Oregon)	12.1 TB	43.26%	-0.00%
US East (Ohio)	528.5 GB	1.85%	0.00%

Top 3 buckets

Bucket	Total storage	% of total	% change
rw	8.9 TB	31.92%	14.95%
rw	7.1 TB	25.53%	0%
rw	5.7 TB	20.24%	0.00%

Let’s start with S3 Storage Lens

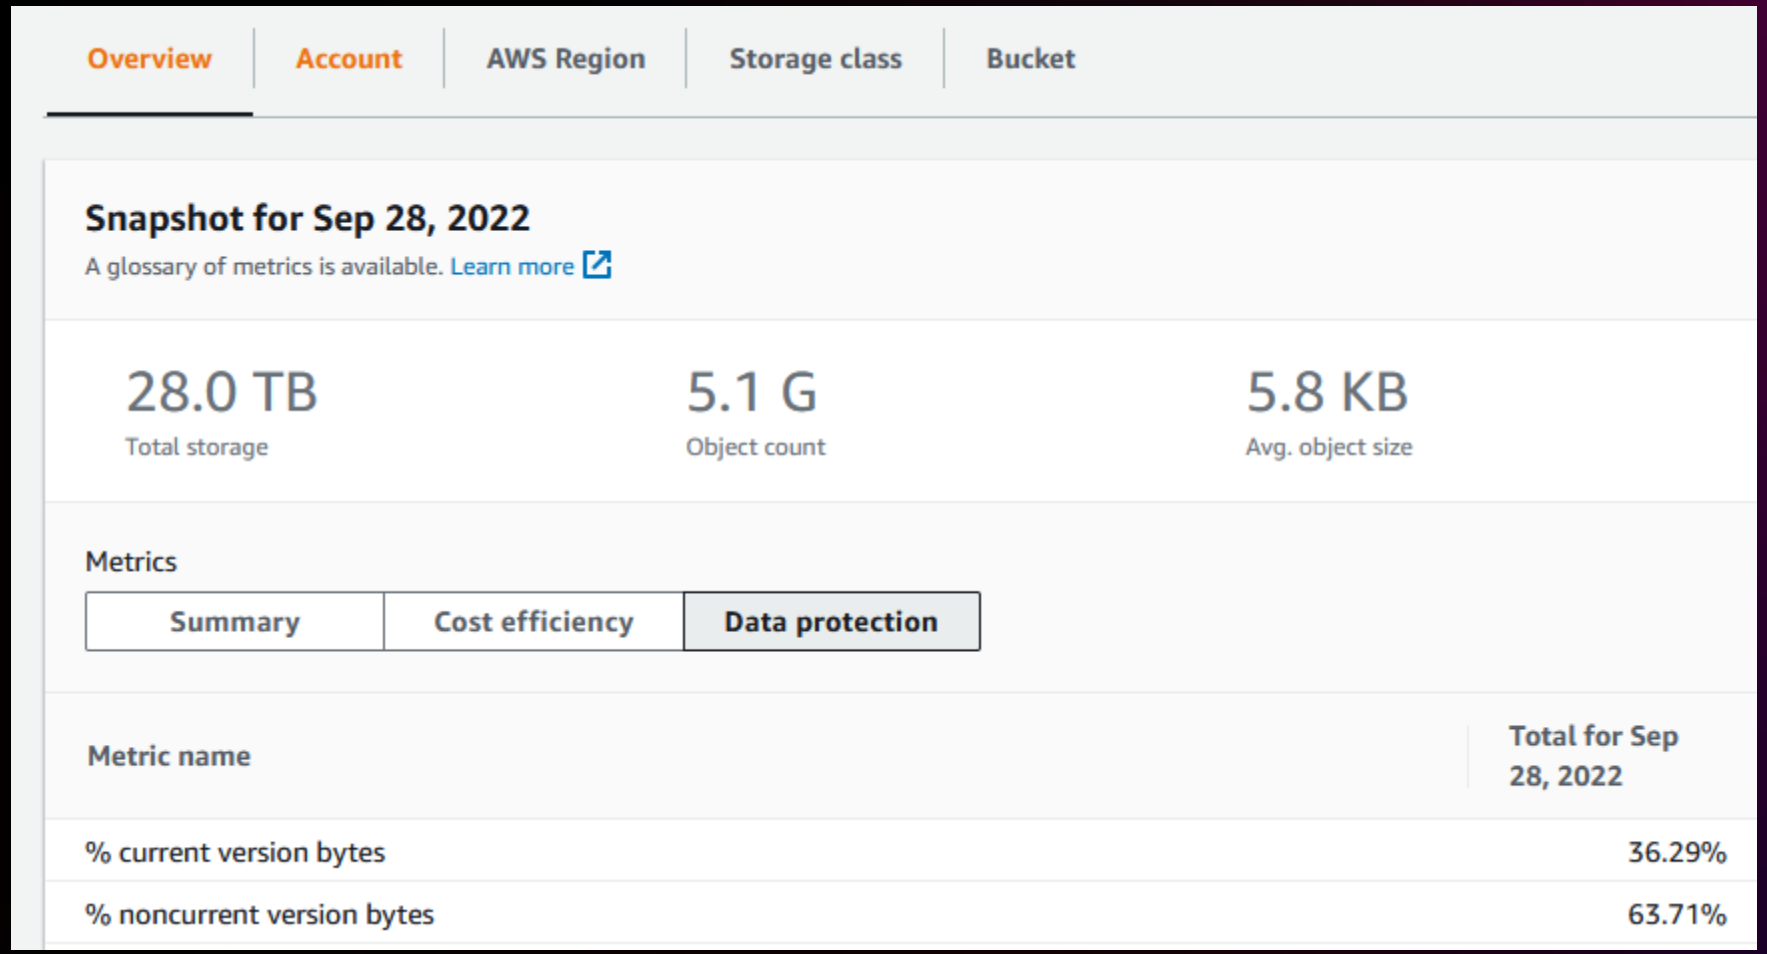

Let’s start with S3 Storage Lens

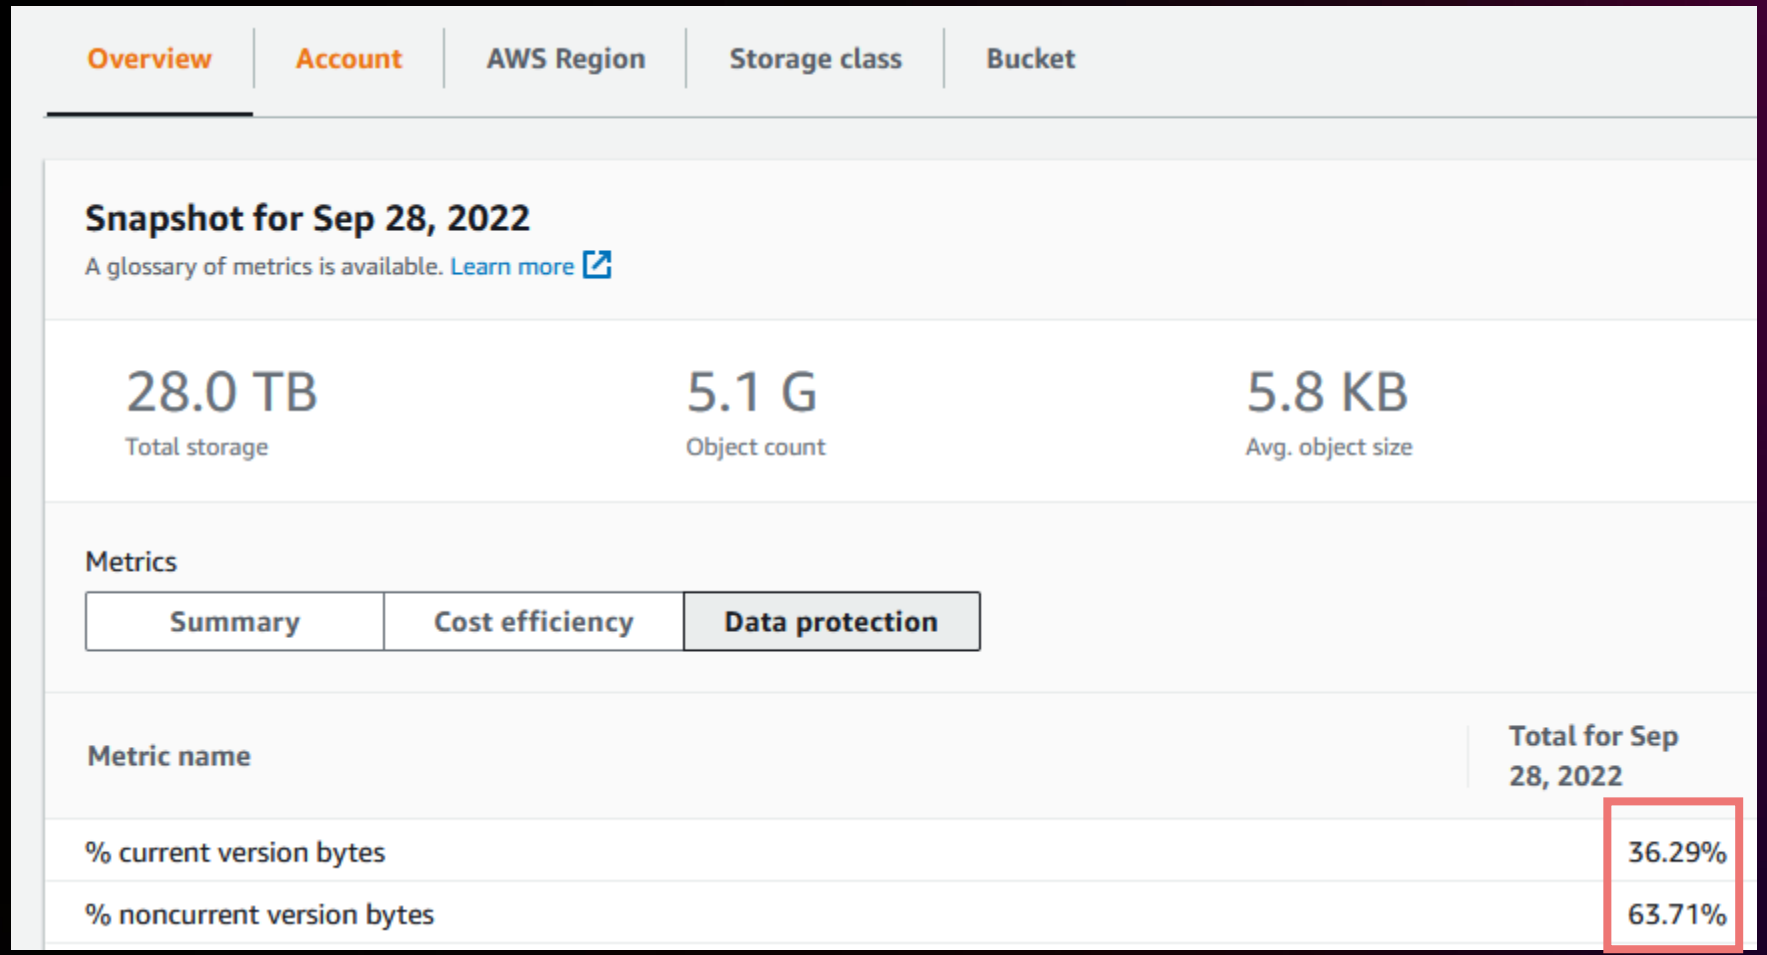

Decide your retention needs

Does this workload/application require

- Retention for a specific period of time?
- Retention of overwritten data?
- Different rules based on object size?

Example configurations: Unversioned buckets

```
<LifecycleConfiguration>
  <Rule>
    <ID>Transition and Expiration Rule</ID>
    <Filter>
       <Prefix>tax/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Transition>
      <Days>365</Days>
      <StorageClass>S3 Glacier Flexible Retrieval</StorageClass>
    </Transition>
    <Expiration>
      <Days>3650</Days>
    </Expiration>
  </Rule>
</LifecycleConfiguration>
```

Example configurations: Unversioned buckets

```
<LifecycleConfiguration>
  <Rule>
    <ID>Transition and Expiration Rule</ID>
    <Filter>
       <Prefix>tax/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Transition>
      <Days>365</Days>
      <StorageClass>S3 Glacier Flexible Retrieval</StorageClass>
    </Transition>
    <Expiration>
      <Days>3650</Days>
    </Expiration>
  </Rule>
</LifecycleConfiguration>
```

Example configurations: Unversioned buckets

```
<LifecycleConfiguration>
  <Rule>
    <ID>Transition and Expiration Rule</ID>
    <Filter>
       <Prefix>tax/</Prefix>
    </Filter>
    <Status>Enabled</Status>
    <Transition>
      <Days>365</Days>
      <StorageClass>S3 Glacier Flexible Retrieval</StorageClass>
    </Transition>
    <Expiration>
      <Days>3650</Days>
    </Expiration>
  </Rule>
</LifecycleConfiguration>
```

Example configurations: Versioned buckets

```
<LifecycleConfiguration>
    <Rule>
        <ID>VersionRetention</ID>
        <Filter>
           <Prefix></Prefix>
        </Filter>
        <Status>Enabled</Status>
        <Transition>
           <Days>30</Days>
           <StorageClass>STANDARD_IA</StorageClass>
        </Transition>
        <NoncurrentVersionTransition>
            <NoncurrentDays>7</NoncurrentDays>
            <StorageClass>S3 Glacier Flexible Retrieval</StorageClass>
        </NoncurrentVersionTransition>
       <NoncurrentVersionExpiration>
            <NewerNoncurrentVersions>5</NewerNoncurrentVersions>
            <NoncurrentDays>365</NoncurrentDays>
       </NoncurrentVersionExpiration>
    </Rule>
</LifecycleConfiguration>
```

Example configurations: Versioned buckets

```
<LifecycleConfiguration>
    <Rule>
        <ID>VersionRetention</ID>
        <Filter>
           <Prefix></Prefix>
        </Filter>
        <Status>Enabled</Status>
        <Transition>
           <Days>30</Days>
           <StorageClass>STANDARD_IA</StorageClass>
        </Transition>
        <NoncurrentVersionTransition>
            <NoncurrentDays>7</NoncurrentDays>
            <StorageClass>S3 Glacier Flexible Retrieval</StorageClass>
        </NoncurrentVersionTransition>
       <NoncurrentVersionExpiration>
            <NewerNoncurrentVersions>5</NewerNoncurrentVersions>
            <NoncurrentDays>365</NoncurrentDays>
       </NoncurrentVersionExpiration>
    </Rule>
</LifecycleConfiguration>
```

Example configurations: Versioned buckets

```
<LifecycleConfiguration>
    <Rule>
        <ID>VersionRetention</ID>
        <Filter>
           <Prefix></Prefix>
        </Filter>
        <Status>Enabled</Status>
        <Transition>
           <Days>30</Days>
           <StorageClass>STANDARD_IA</StorageClass>
        </Transition>
        <NoncurrentVersionTransition>
            <NoncurrentDays>7</NoncurrentDays>
            <StorageClass>S3 Glacier Flexible Retrieval</StorageClass>
        </NoncurrentVersionTransition>
       <NoncurrentVersionExpiration>
            <NewerNoncurrentVersions>5</NewerNoncurrentVersions>
            <NoncurrentDays>365</NoncurrentDays>
       </NoncurrentVersionExpiration>
    </Rule>
</LifecycleConfiguration>
```

Example configurations: Versioned buckets

```
<LifecycleConfiguration>
    <Rule>
        <ID>VersionRetention</ID>
        <Filter>
           <Prefix></Prefix>
        </Filter>
        <Status>Enabled</Status>
        <Transition>
           <Days>30</Days>
           <StorageClass>STANDARD_IA</StorageClass>
        </Transition>
        <NoncurrentVersionTransition>
            <NoncurrentDays>7</NoncurrentDays>
            <StorageClass>S3 Glacier Flexible Retrieval</StorageClass>
        </NoncurrentVersionTransition>
        <NoncurrentVersionExpiration>
            <NewerNoncurrentVersions>5</NewerNoncurrentVersions>
            <NoncurrentDays>365</NoncurrentDays>
        </NoncurrentVersionExpiration>
    </Rule>
</LifecycleConfiguration>
```

Example configurations: Object size

```
<LifecycleConfiguration>
  <Rule>
    <ID>Transition with a prefix and based on size</ID>
    <Filter>
       <And>
          <Prefix>tax/</Prefix>
          <ObjectSizeGreaterThan>131072</ObjectSizeGreaterThan>
       </And>
    </Filter>
    <Status>Enabled</Status>
    <Transition>
      <Days>365</Days>
      <StorageClass>S3 Glacier Flexible Retrieval</StorageClass>
    </Transition>
  </Rule>
</LifecycleConfiguration>
```

Monitor your storage: See the impact

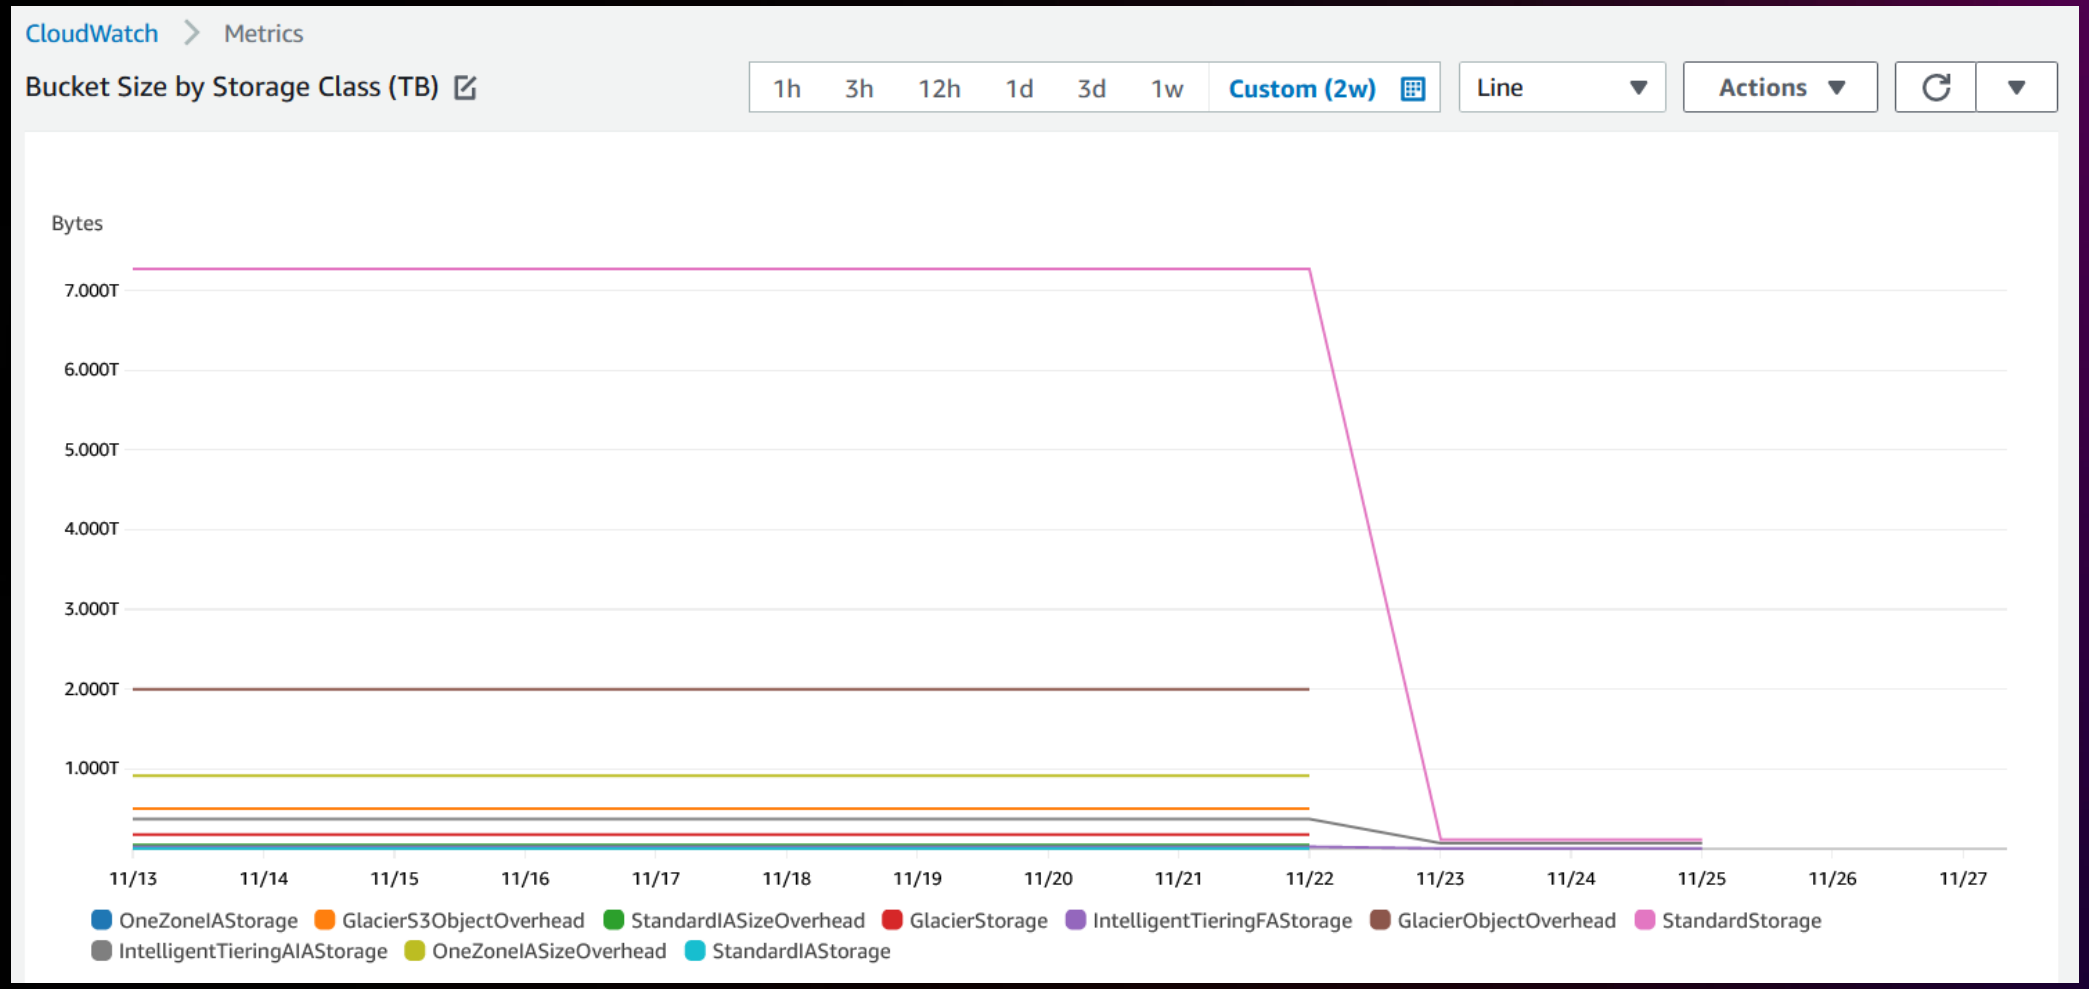

Customer voice:
Bridgewater Associates

About Bridgewater

- Founded by **Ray Dalio**
- **Global macroeconomic investment manager**
- We endlessly pursue the right answer to how economies work
- Our economic model is complex, constantly evolving for over 45 years

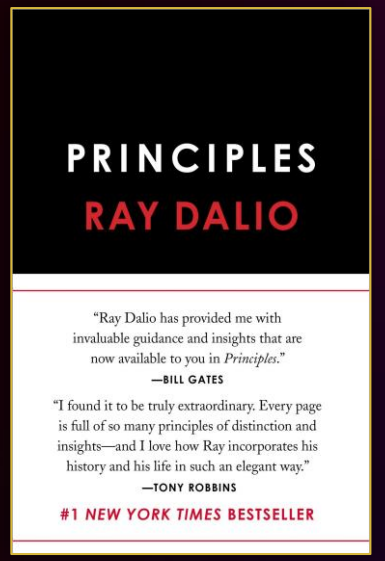

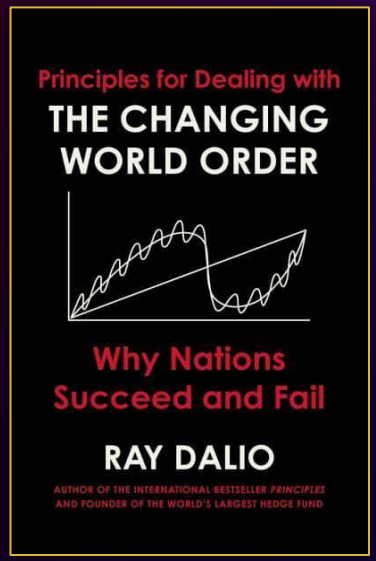

Bridgewater's model of the economy

- Research **timeless and universal** economic and investment principles
- We take these principles and write them as **code**
- We run this code, our **economic model**, continuously

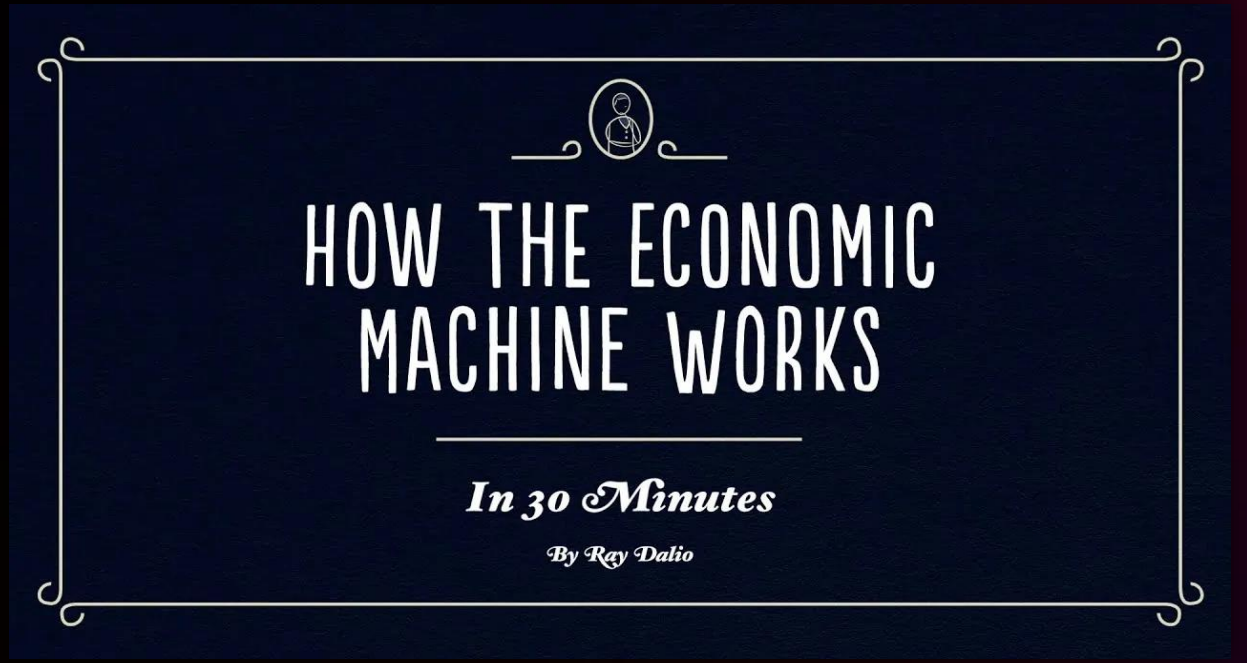

Bridgewater’s economic model runs on AWS

- New data flows in
- Forward-testing every day
- Back-testing all model changes
- Research workflows with hypothetical model changes (what-if scenarios)
- Consistently producing trades in all the markets we operate
- **AWS based architecture allows us to scale out with reliability**

Compute architecture

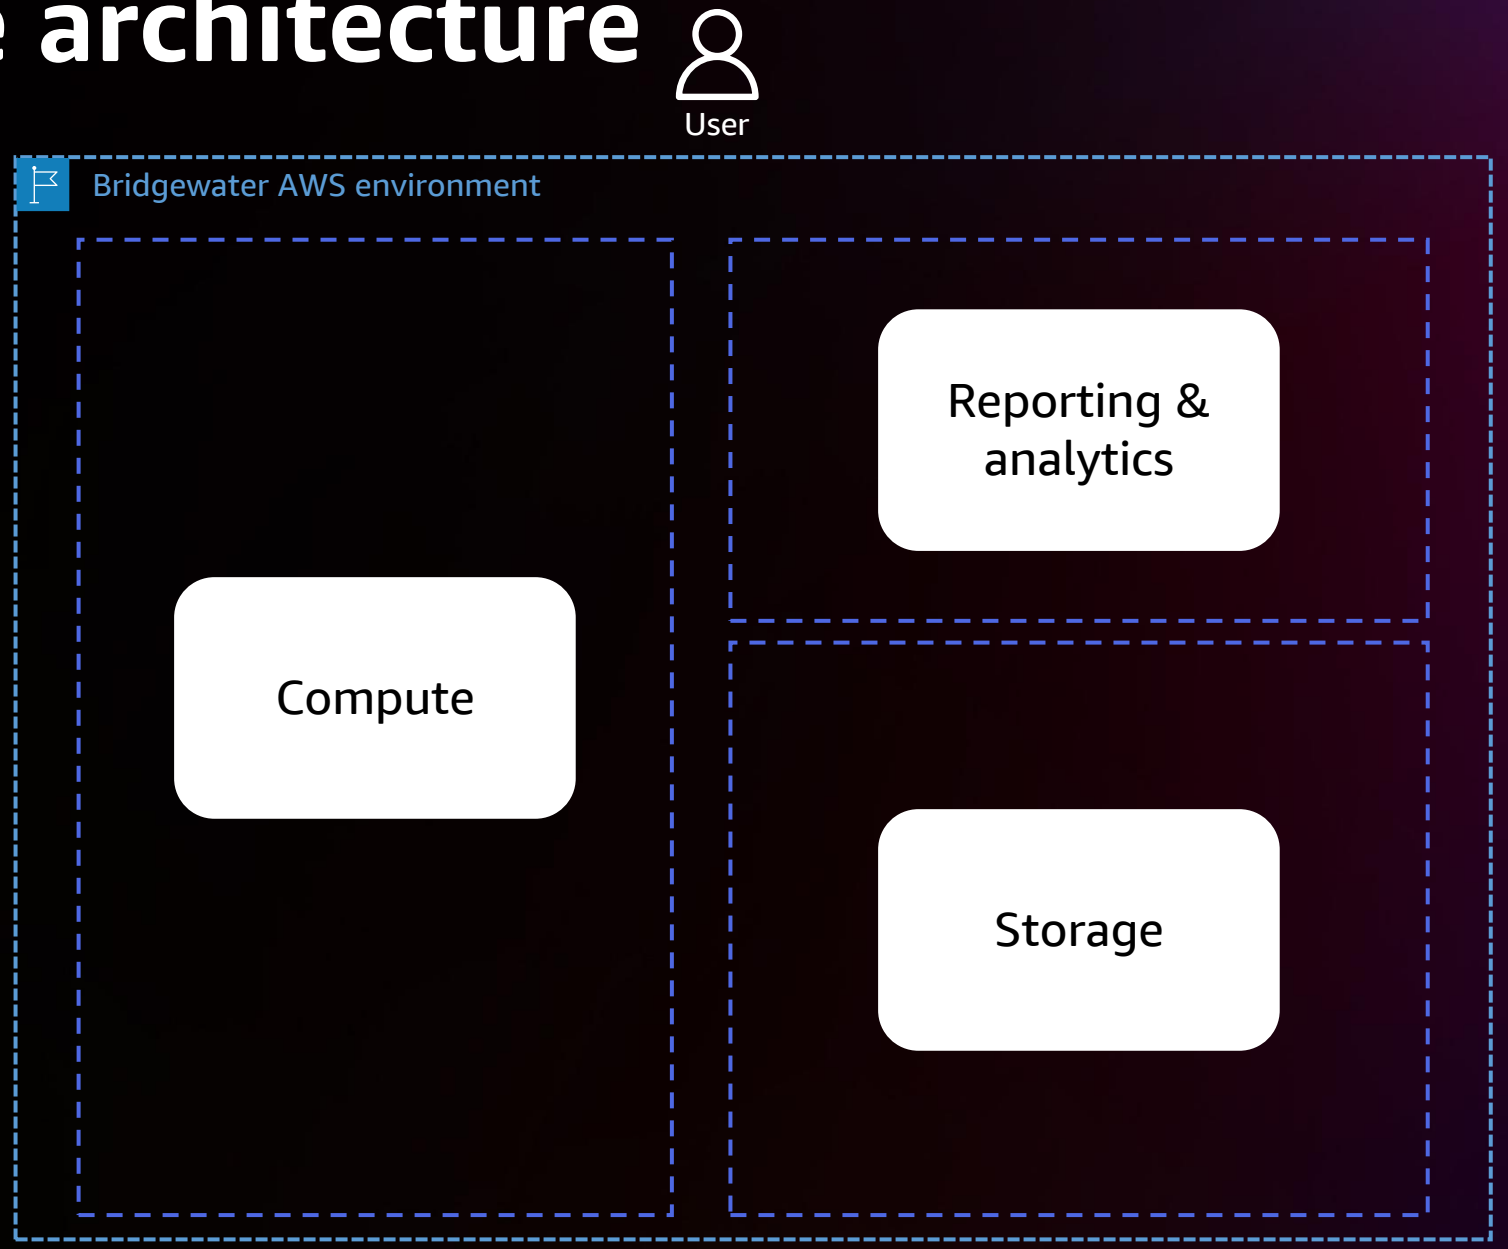

Compute architecture

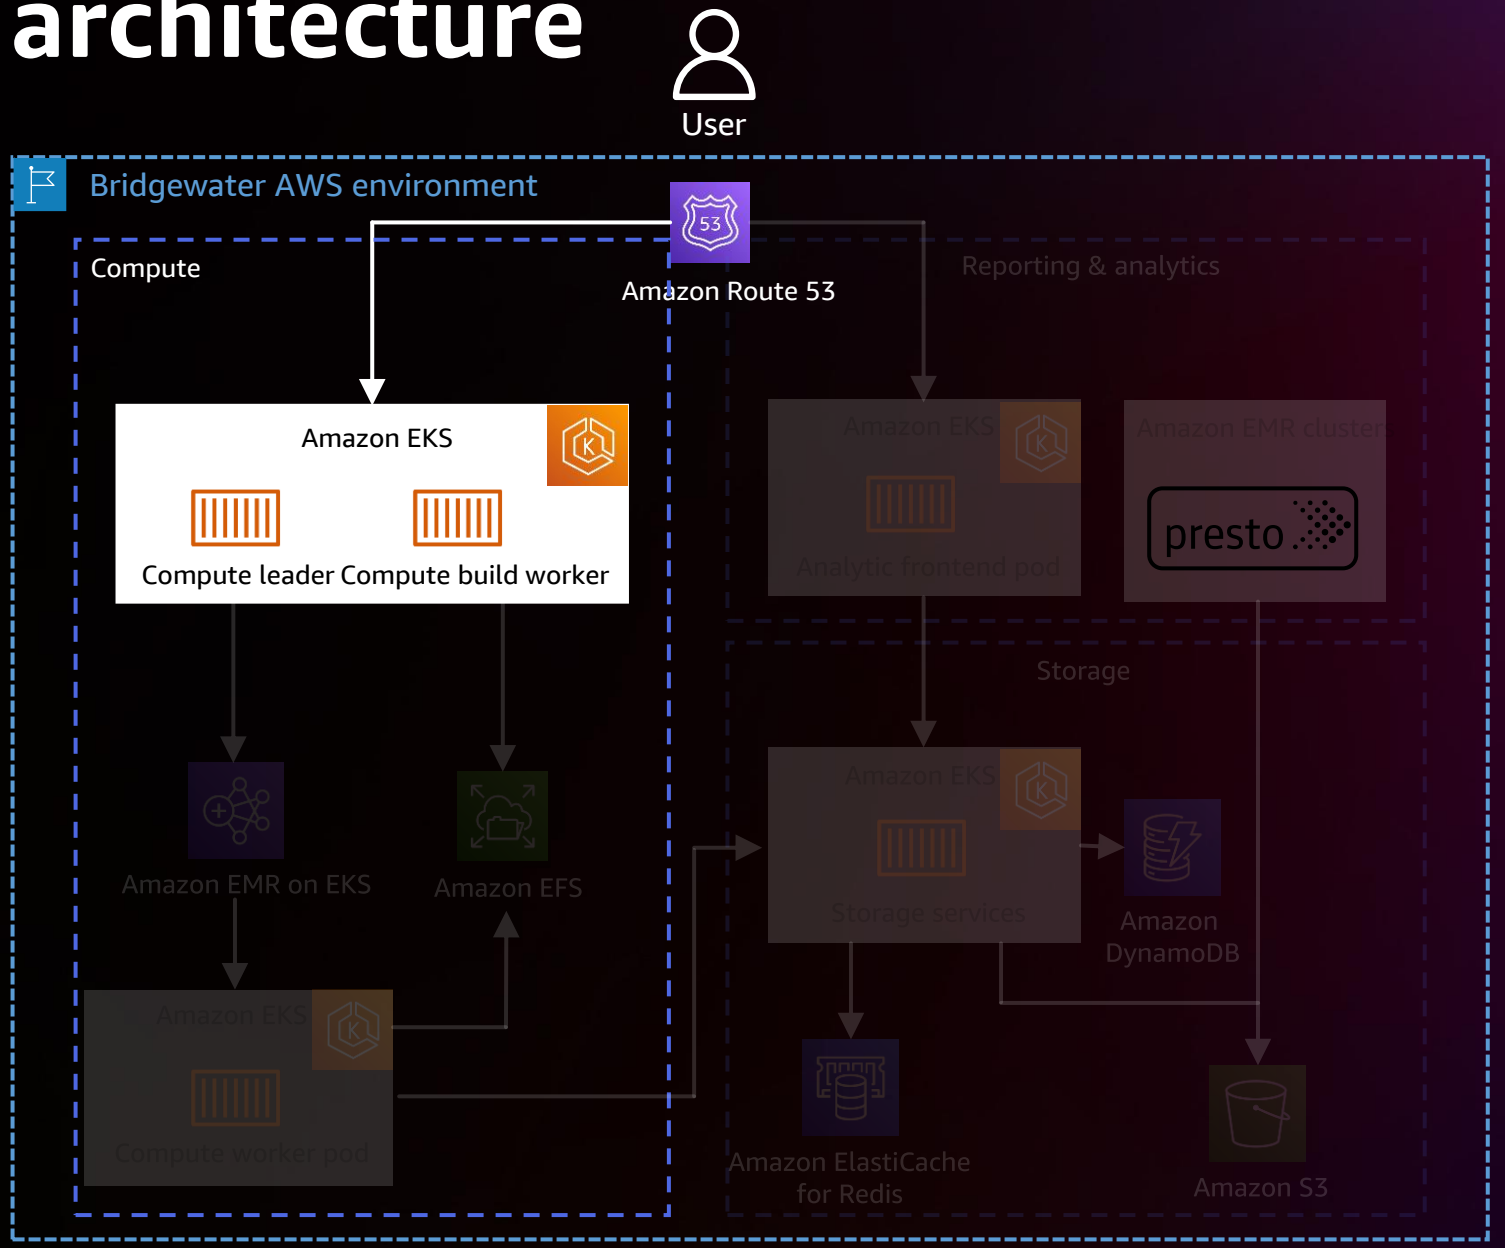

Compute architecture

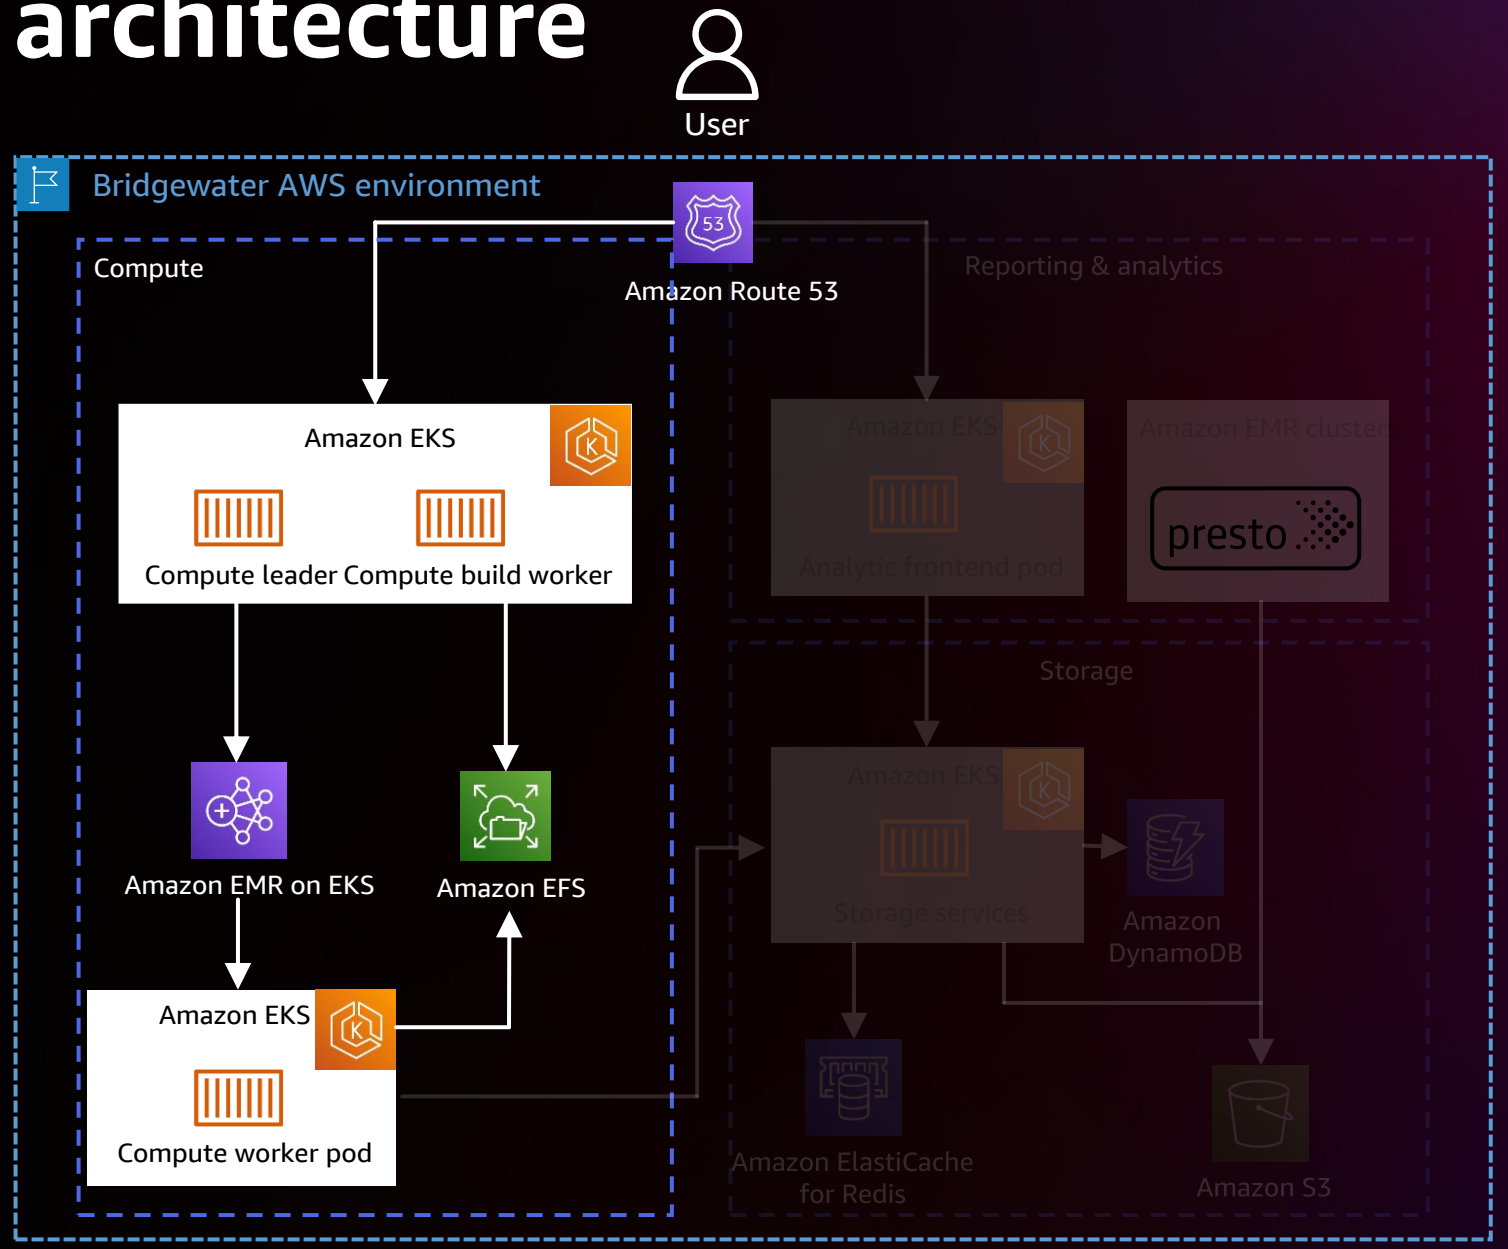

Compute architecture

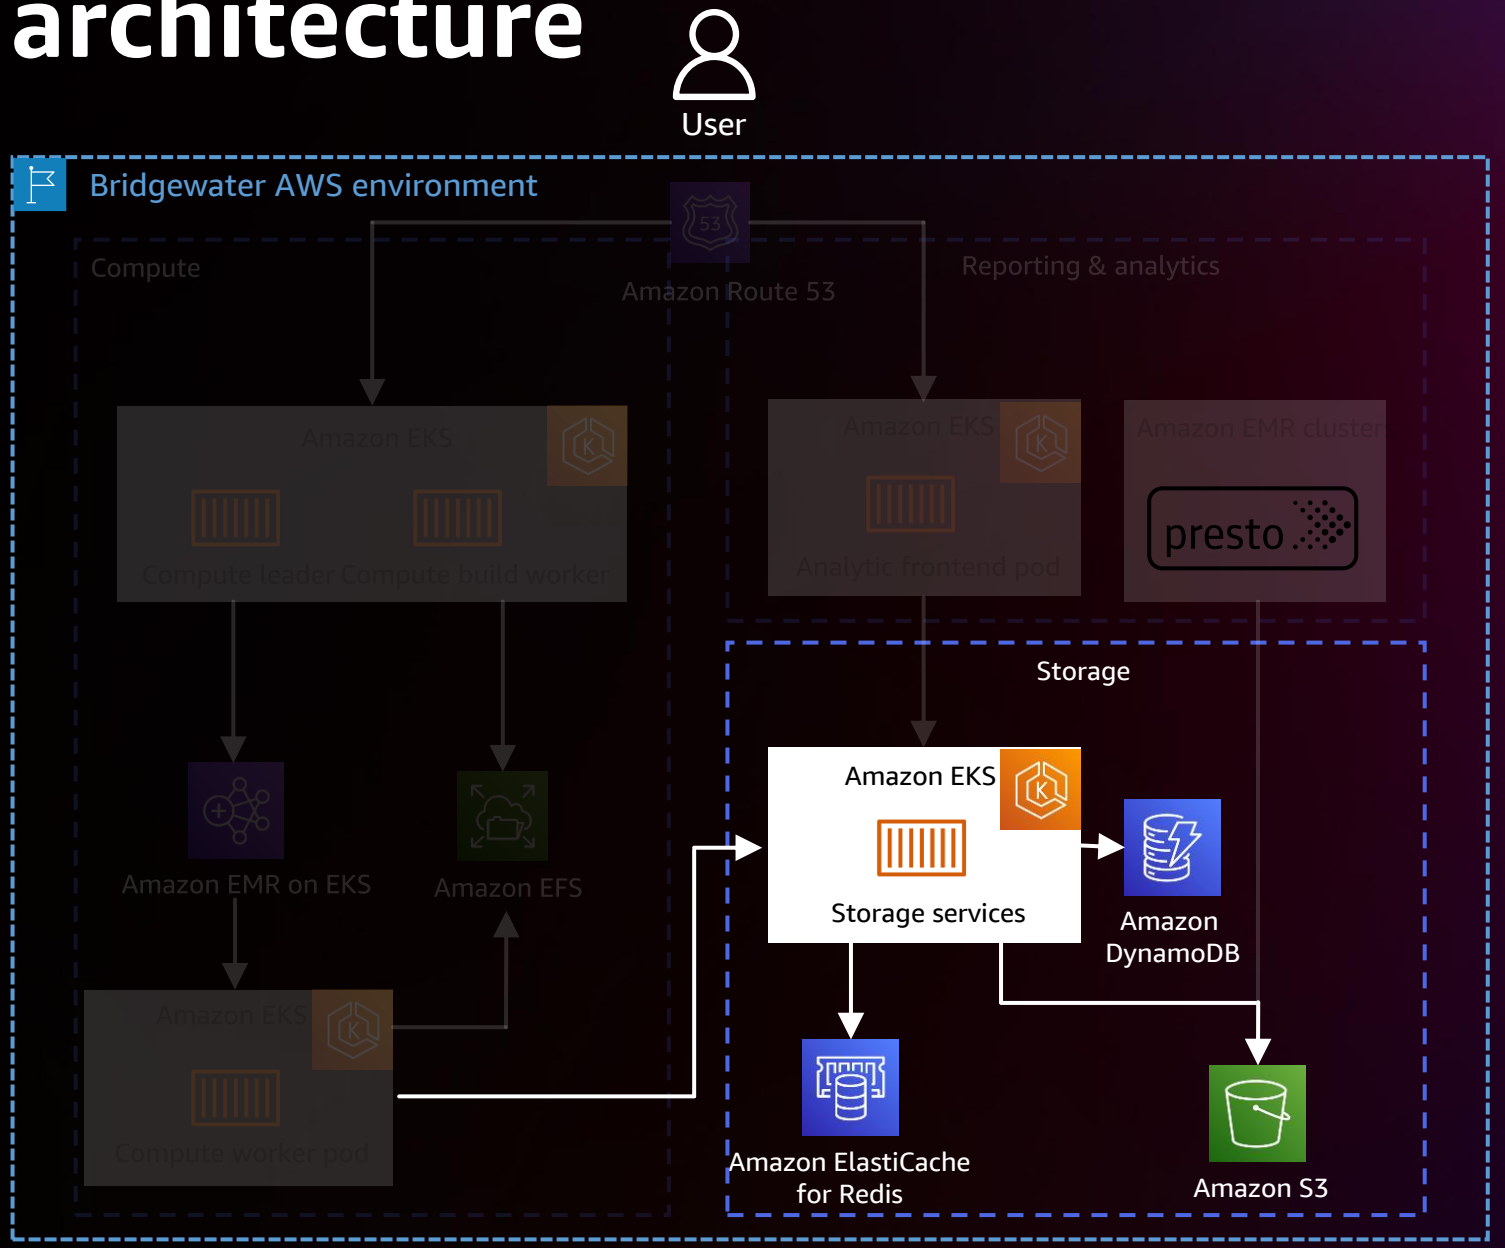

Compute architecture

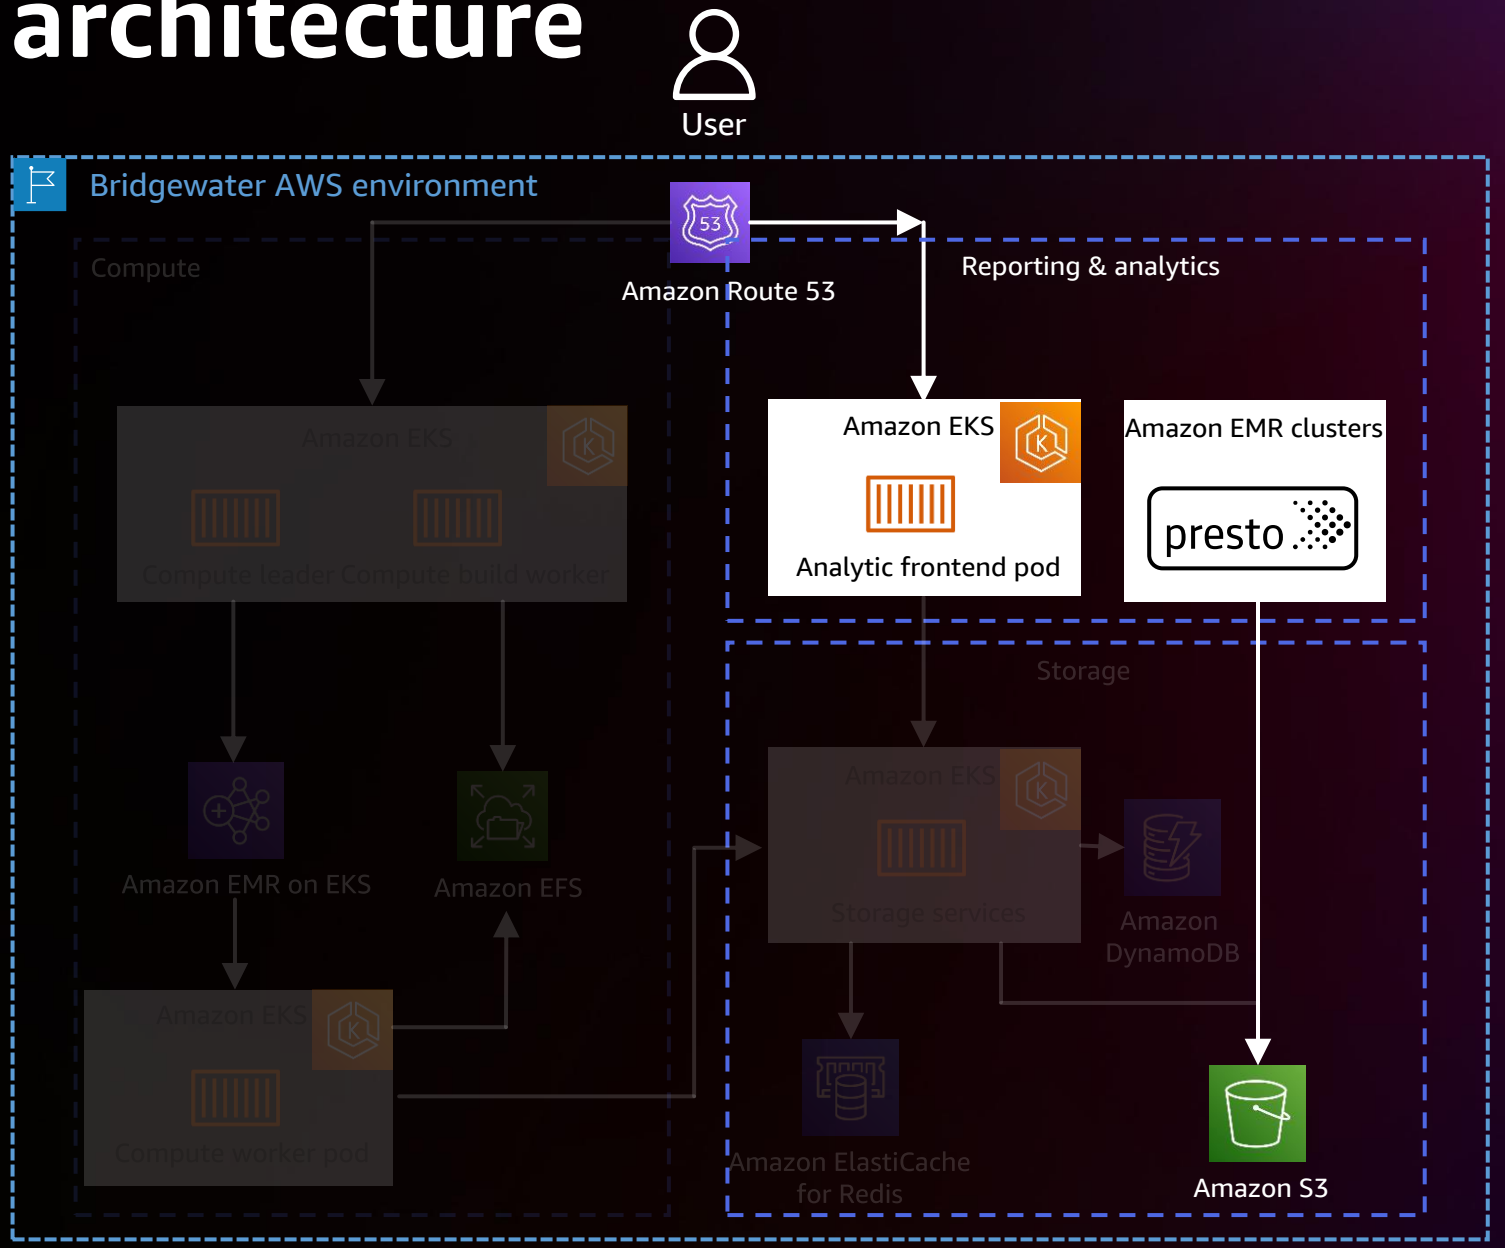

Compute architecture

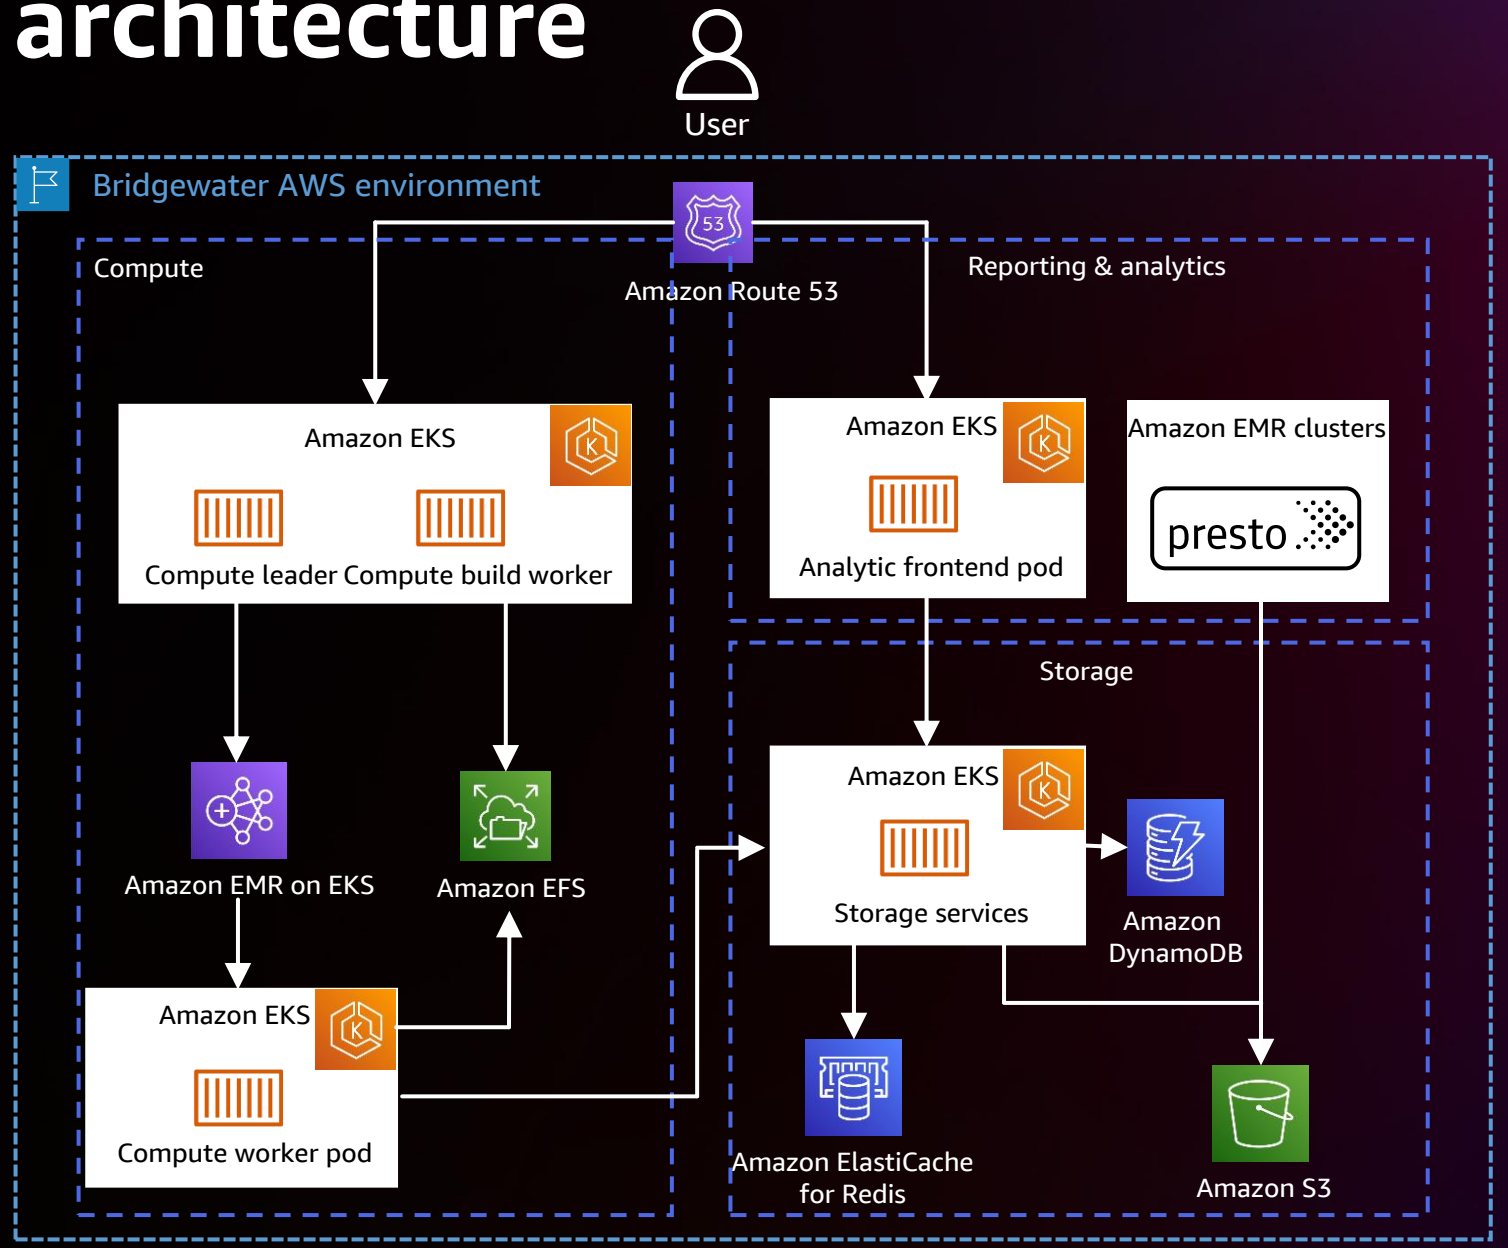

Data storage needs at Bridgewater

Security	Reliability	Cost	Our users need scale

Our storage stack

Proprietary database built on top of **Amazon S3 as the data store**

Amazon DynamoDB as the metadata store

Amazon ElastiCache to accelerate read-heavy workloads

Security and compliance at scale

- Customer managed AWS KMS encryption
- Amazon S3 default encryption
- S3 Object Lock to ensure data is never deleted
- Controls to ensure **security and compliance** as we scale tens of petabytes

Reliability and consistency at scale

- S3 Replication Time Control with S3 Cross-Region Replication
- Consistent reads in S3 (beyond read-after-write)
- S3 Storage Lens for better analytics
- S3 Multi-Region Access Points
- Designed for **full Region outage scenarios**

Our replication happens at the dataset level

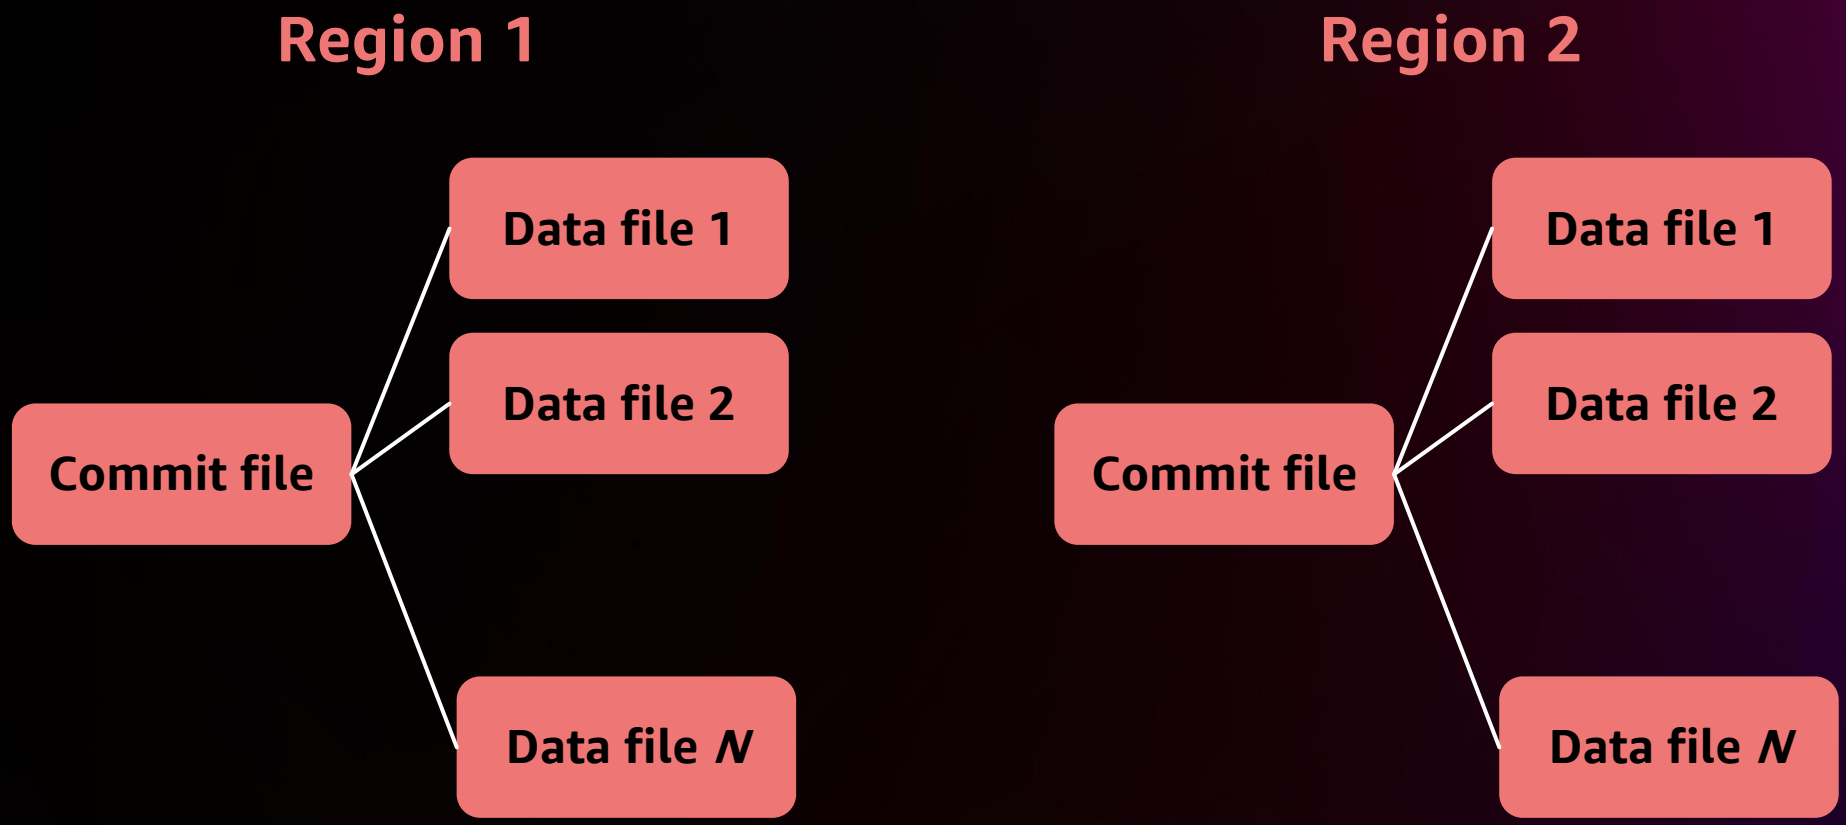

S3 Replication Time Control monitoring

Track the latest timestamp to know when **100%** is replicated

SLA is 15 mins
We usually see the data replicated in 1–3 minutes

Worst case, we only have a **few minutes** of data loss

Managing cost at scale

S3 Intelligent Tiering:

Monthly storage is **up 42%**, while **spend is down 35%**

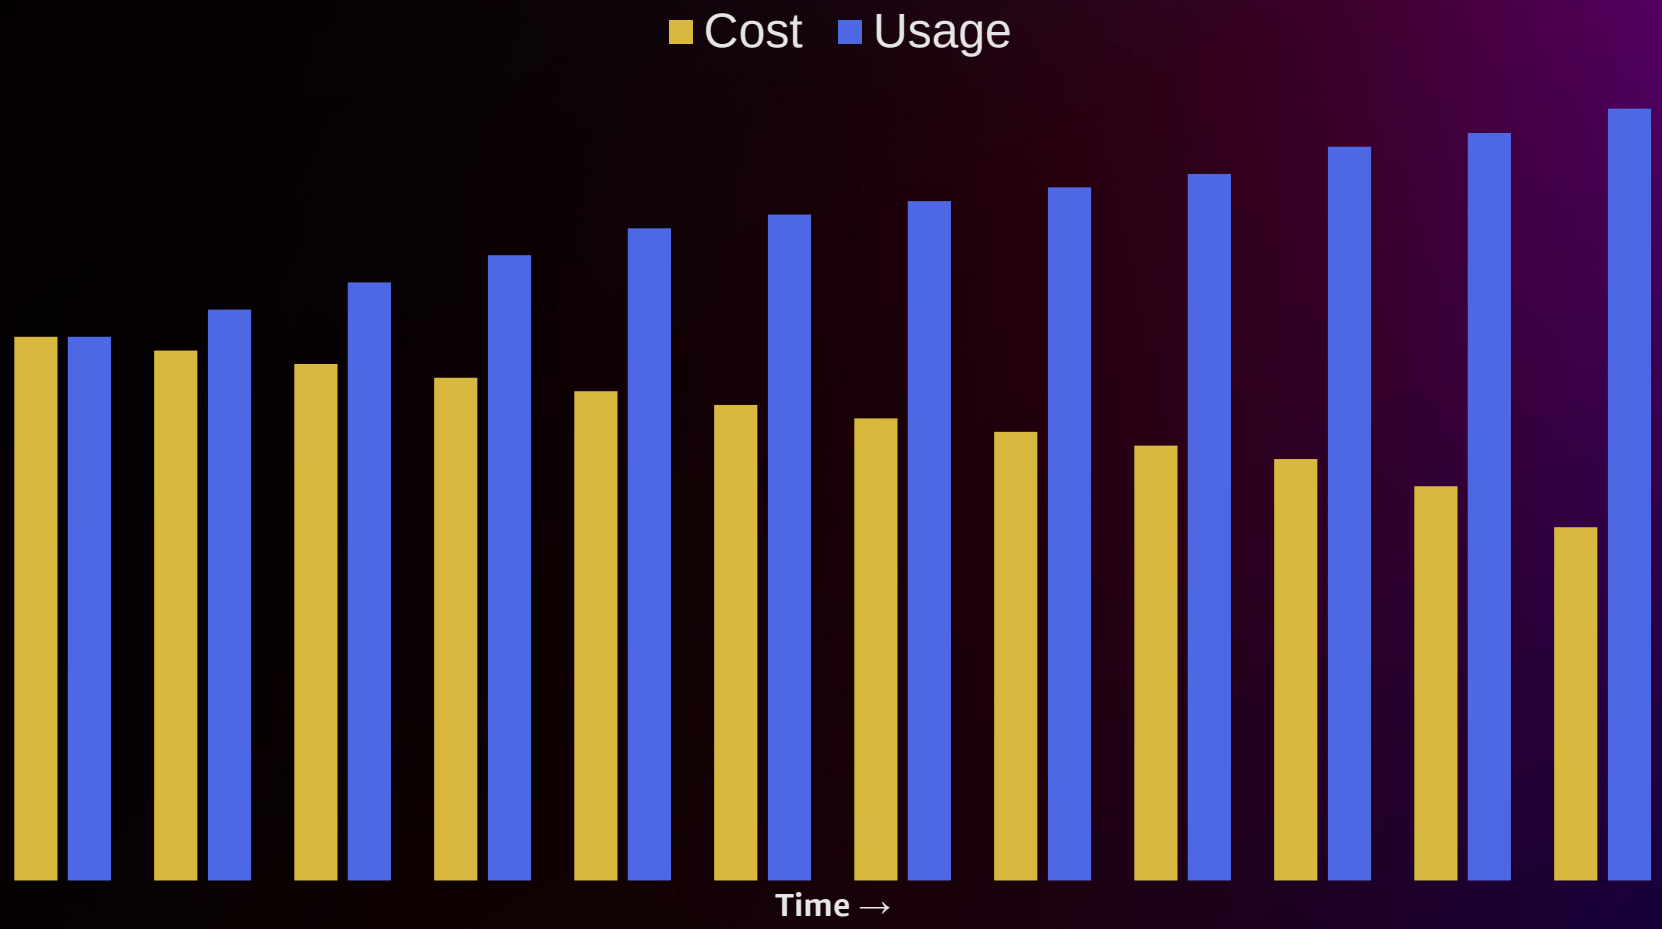

Bridgewater tests the scale of Amazon S3

20+
petabytes

120 B+
objects

2 PB/year
growth

AWS services improve our business

2x
reduction in average runtime

75%+
machine utilization

5x
scale of simulations

Session recap

Securing your data – encrypting objects at scale
Use: S3 Inventory, S3 Batch Operations

Accessing archive data – restoring data
Use: S3 Glacier storage classes, S3 Batch Operations, S3 Event Notifications

Monitoring access – querying request logs
Use: AWS CloudTrail, Amazon Athena

Managing storage spend – configuring retention
Use: S3 Storage Lens, S3 Lifecycle

Customer voice – Bridgewater Associates
Use: S3 Intelligent-Tiering, AWS KMS

Continue your AWS storage learning

Build a learning plan

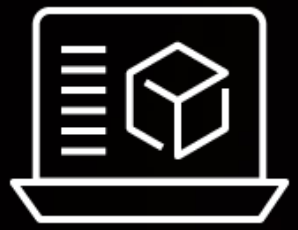

Set your AWS storage learning plans via **AWS Skill Builder**

Increase your knowledge

Use our **Ramp-Up Guides** to build your storage knowledge

Earn AWS Storage badges

Demonstrate your knowledge by achieving **digital badges**

aws.training/storage

Thank you!

Please complete the session survey in the **mobile app**