Cloud Day Warsaw

PGE NARODOWY| 18 SEPTEMBER 2024

DEV208

Optimizing storage price and performance with Amazon S3

Dragoș Mădărășan
Solutions Architect Team Lead
AWS

Agenda

Operating storage at scale

Insights into your Amazon S3 storage and performance usage

Optimizing cost for various data access patterns

Achieving high performance with Amazon S3

Compliance records

Analytics

Geospatial or lunar imagery

Internet of Things (IoT) sensor data

Medical images and records

Data lakes

Customer call-center records

Media master files

Digital record preservation

Mobile sync and storage

Home video recordings

Model checkpoints

Seismic and reservoir simulation data

Pharmaceutical study data

Backups

DNA sequences

Machine learning training data

Surveillance video/closed-circuit television

Media assets

Website hosting

Financial records

Log files

User-generated content

Meteorological and environmental research

Autonomous vehicle data

Oil and gas topography

Data is growing **faster than ever**

Pillars of cost optimization

1

Define your workload requirements

2

Understand your storage and develop insights

3

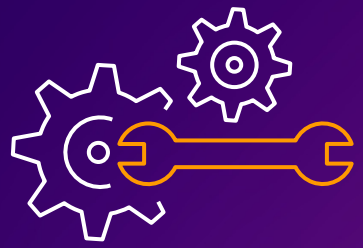

Optimize and measure

Define your workload requirements

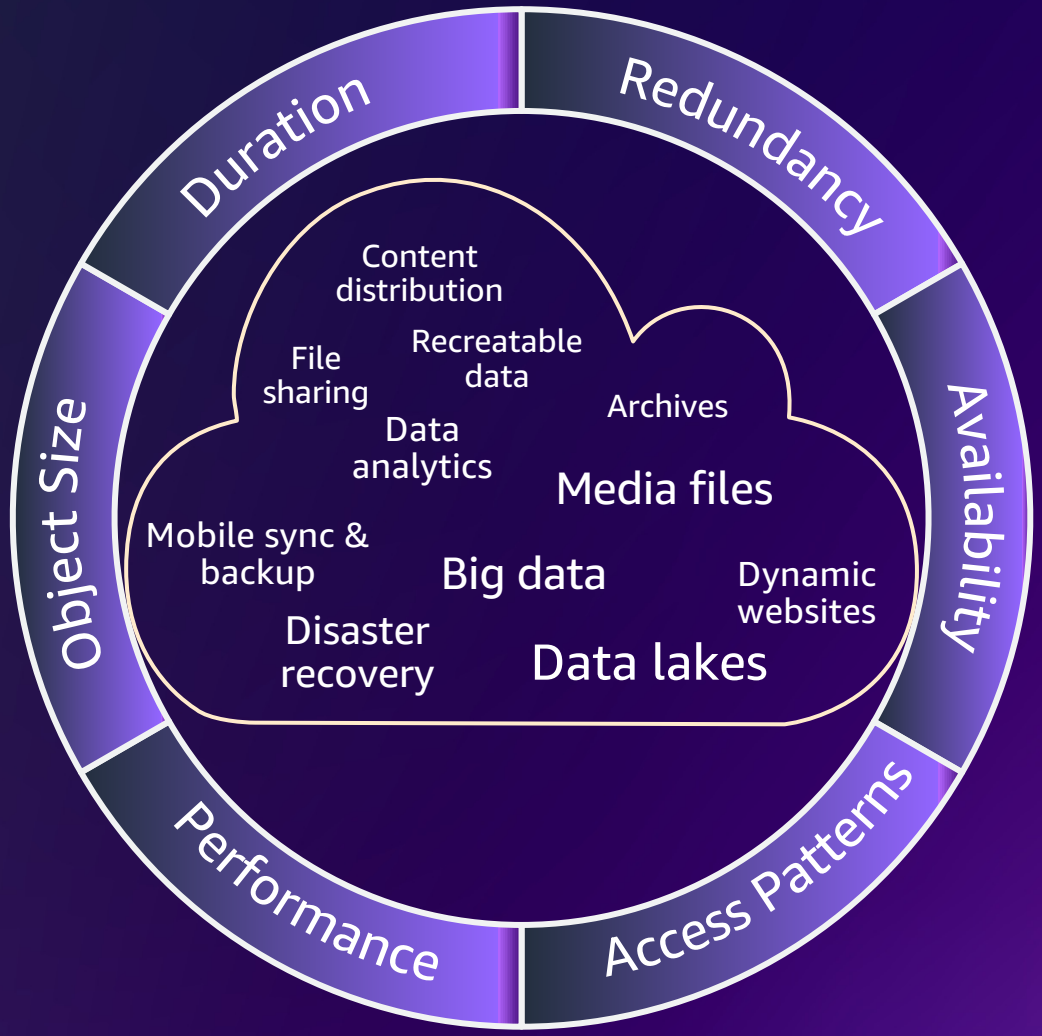

What did customers ask for?

Organization-wide visibility into usage

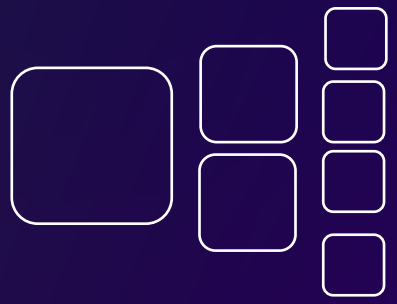

Drill down to gain understanding of usage

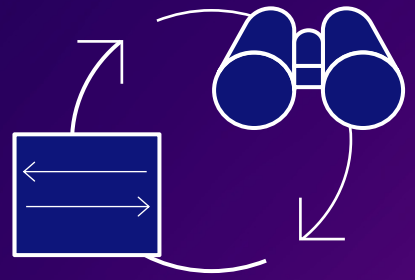

Take action on the insights developed

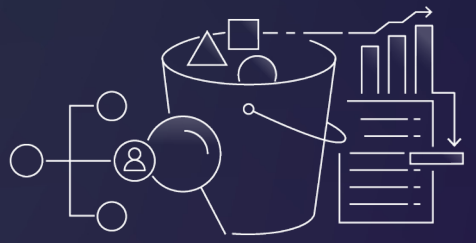

Amazon S3 Storage Lens

Free and paid versions available

- Interactive dashboard experience in the Amazon S3 Console
- Organization-wide visibility into object storage usage
- Drill-down by AWS Region, storage class, bucket, and prefix
- Granular usage & activity metrics

aws

Analyze your storage with S3 Storage Lens

Provides organization-wide visibility into storage usage and activity

COST VISIBILITY "Is my data in the optimal S3 storage class?"

DATA PROTECTION "How do I prevent accidental deletes and overwrites?"

ACCESS CONTROL "Can bucket access for my workloads be improved?"

PERFORMANCE "Am I seeing a rise in 403 authorization errors?"

Customizable aggregation levels in S3 Storage Lens

Object tags

Perform department or application-level cost allocation

Object size

Analyze bucket size distribution to optimize Lifecycle costs

File extensions

Identify the types of data stored in your buckets

Prefixes

Analyze shared multi-tenant buckets like data lakes

Activity and status code metrics aggregated by prefix-level in S3 Storage Lens

Performance improvement

Are my workloads exceeding the per-prefix request rate?

Cost optimization

Which prefixes have gone ‘cold’?

Welcome to S3 Storage Lens

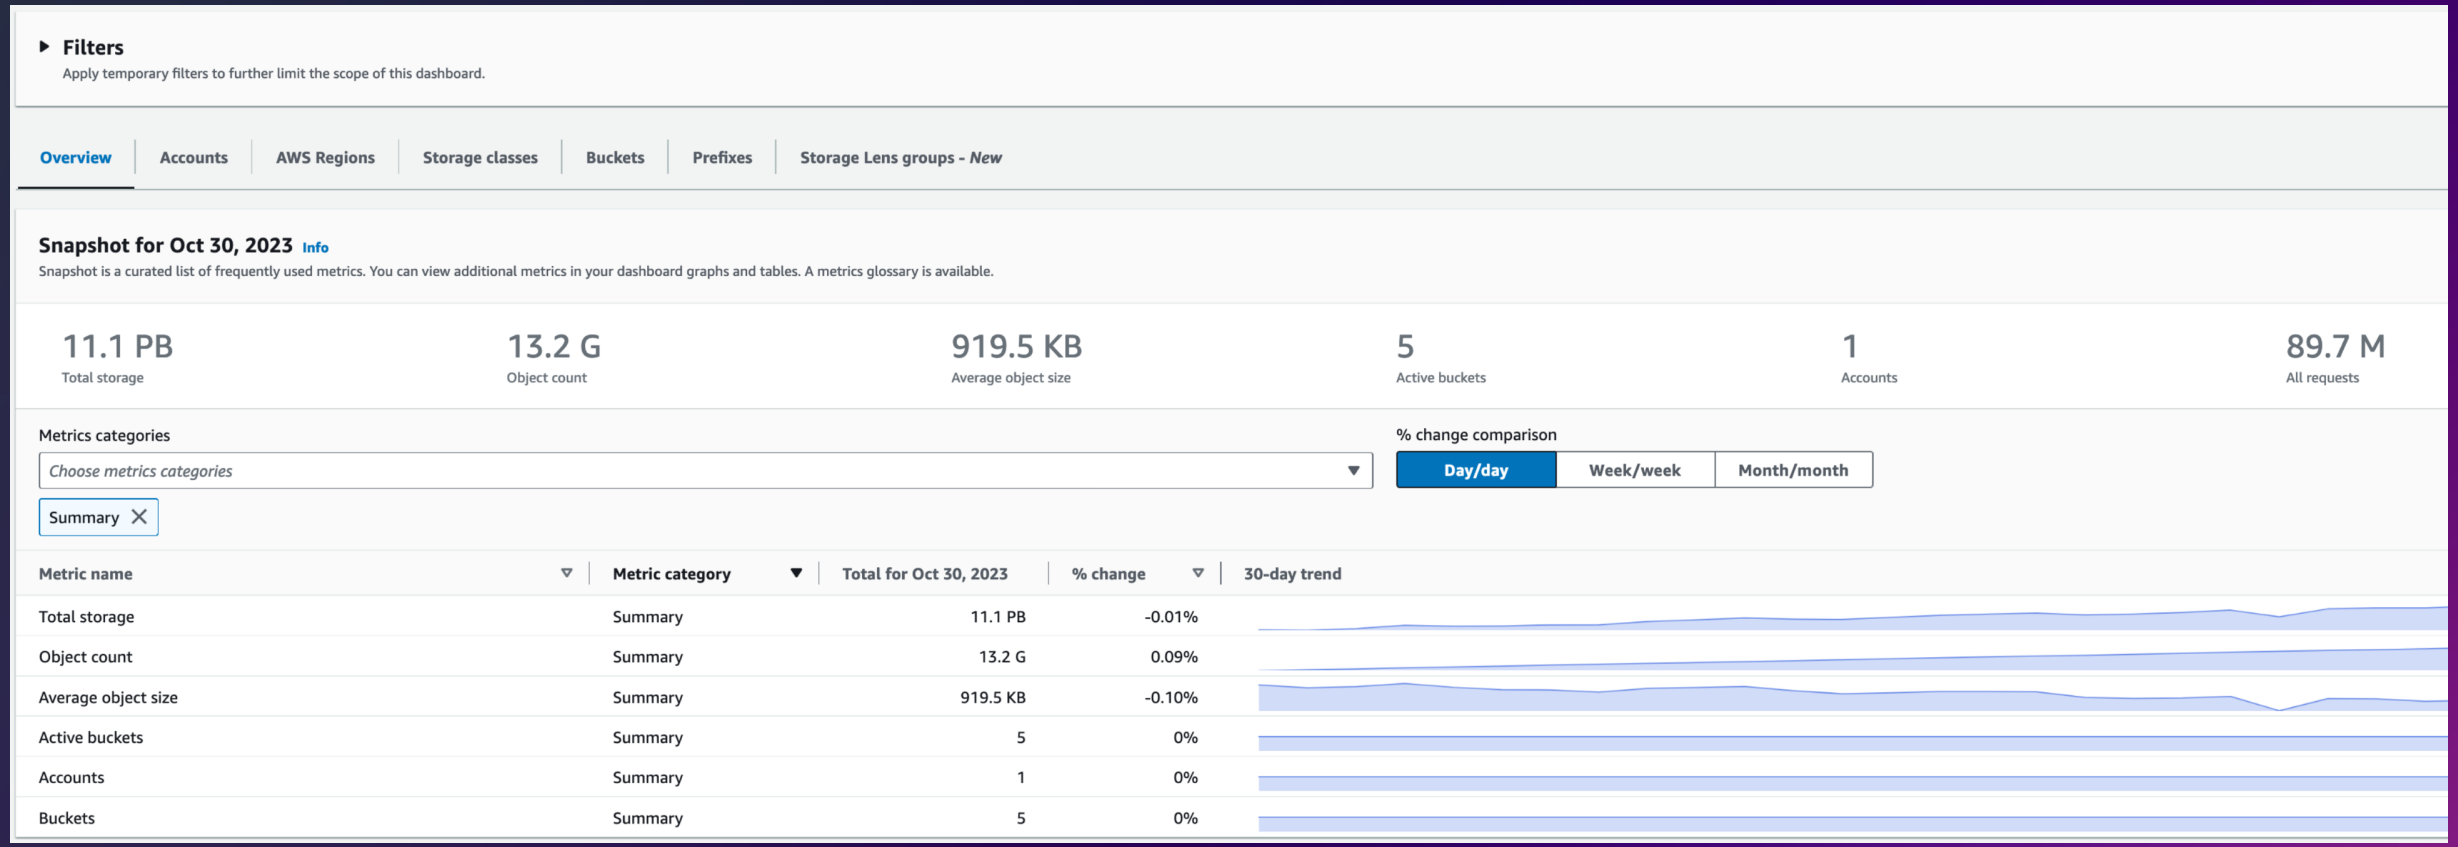

Discover infrequently accessed data

Discover incomplete multipart uploads

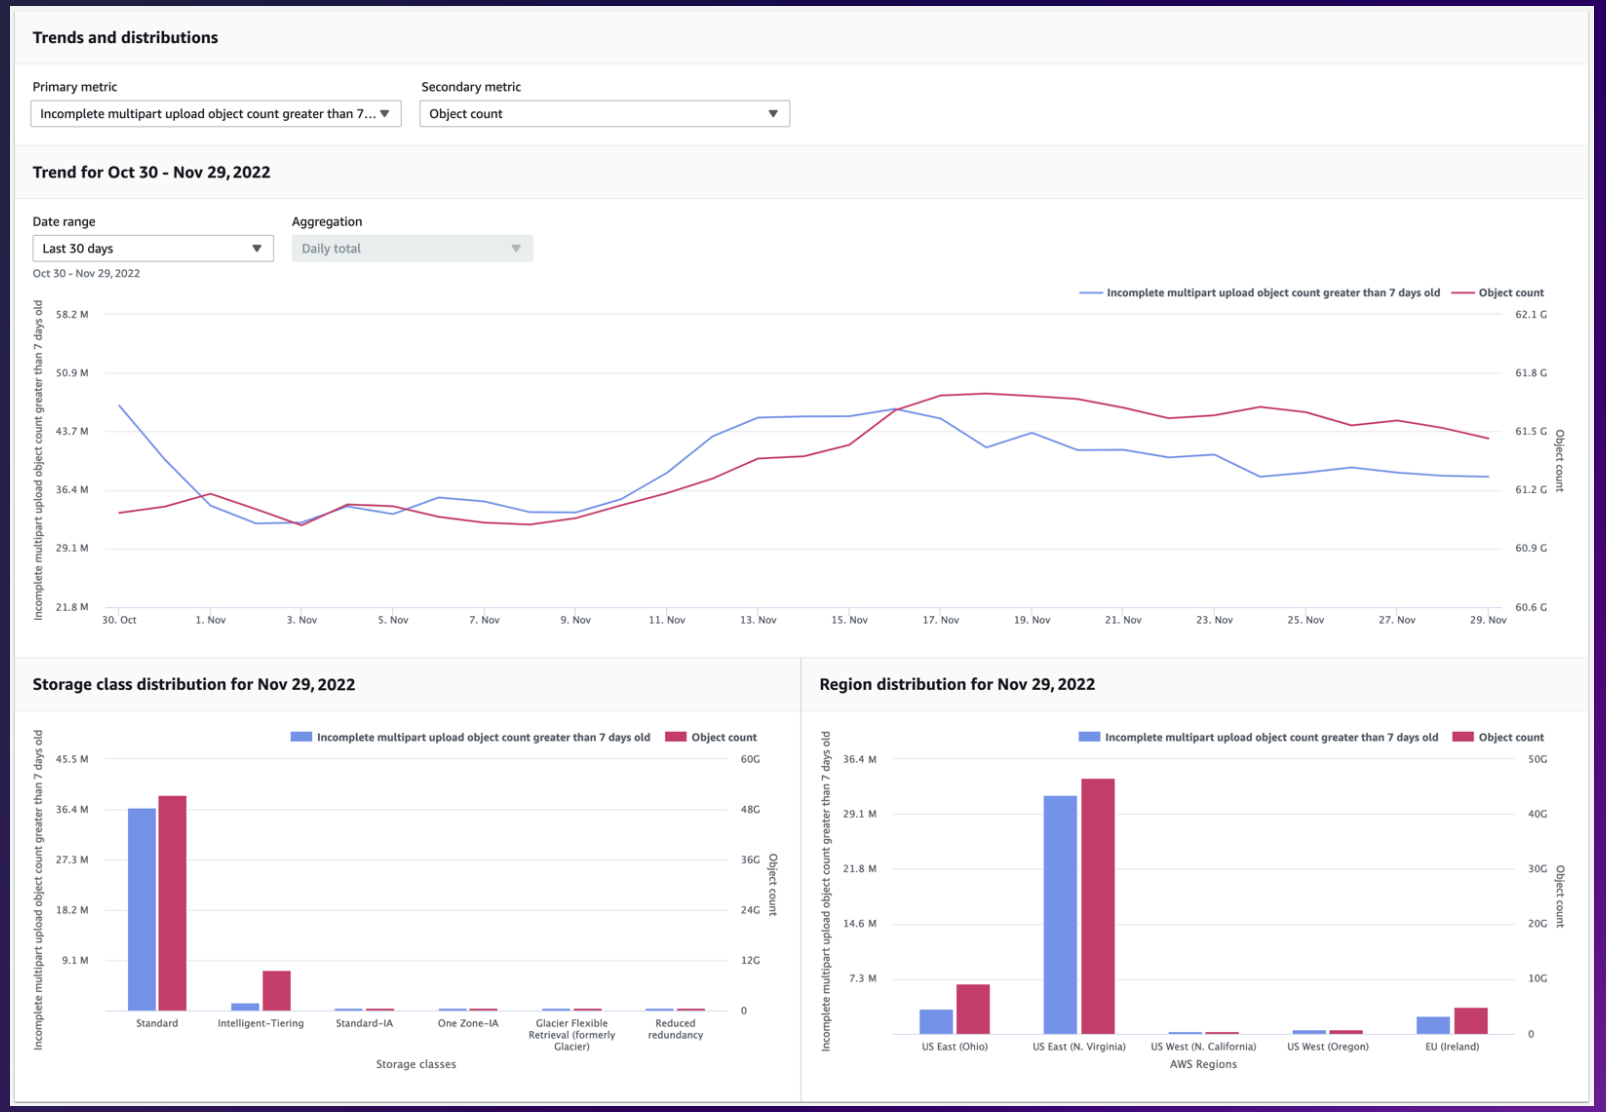

Discover incomplete multipart uploads

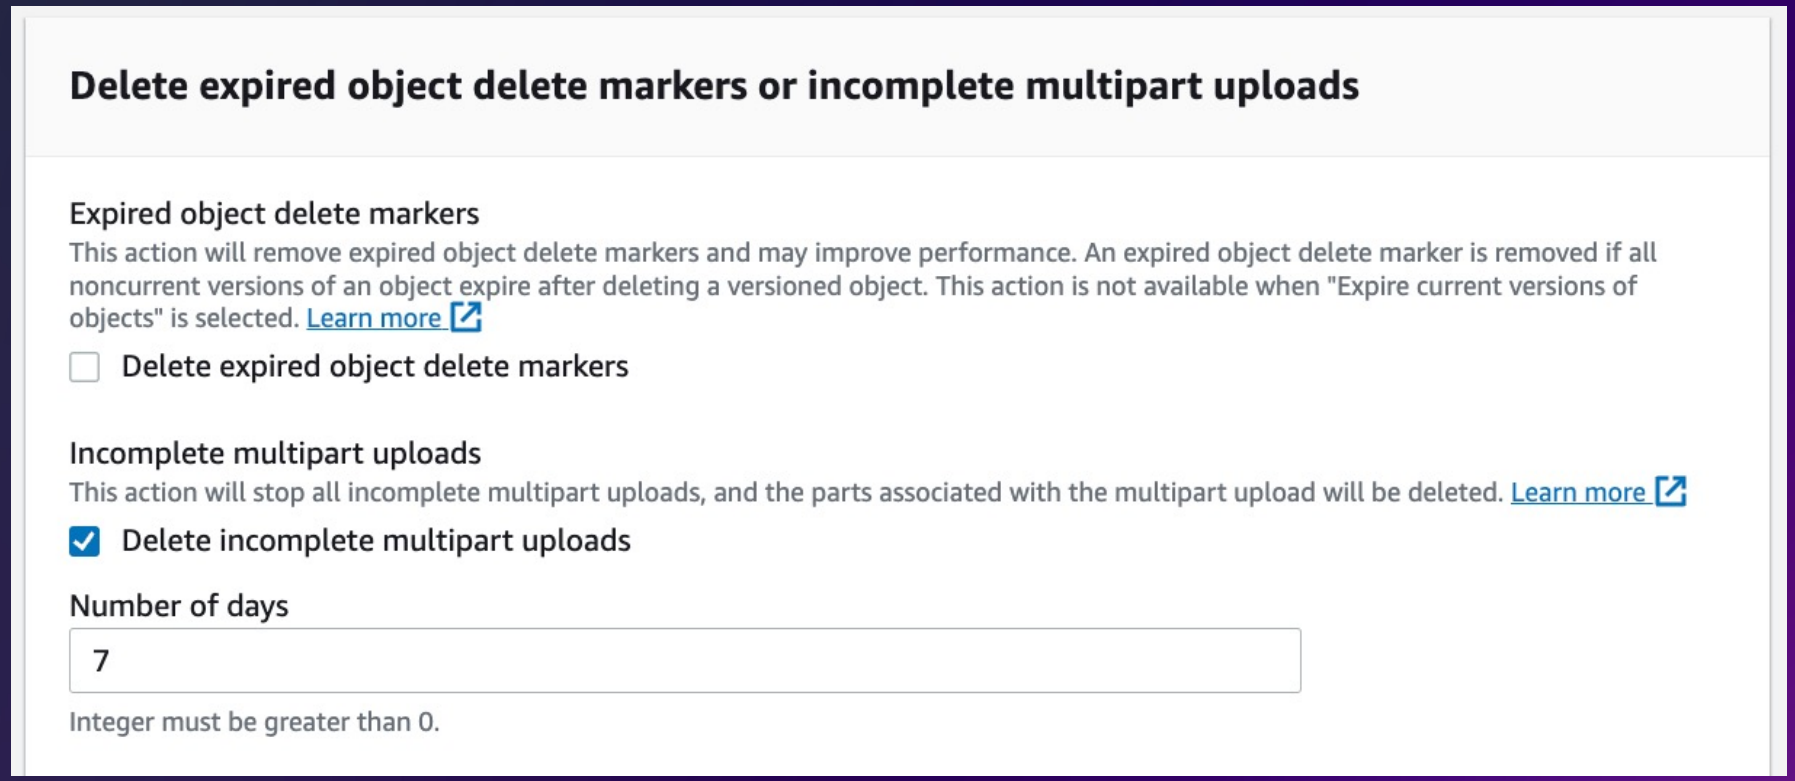

Discover noncurrent versions

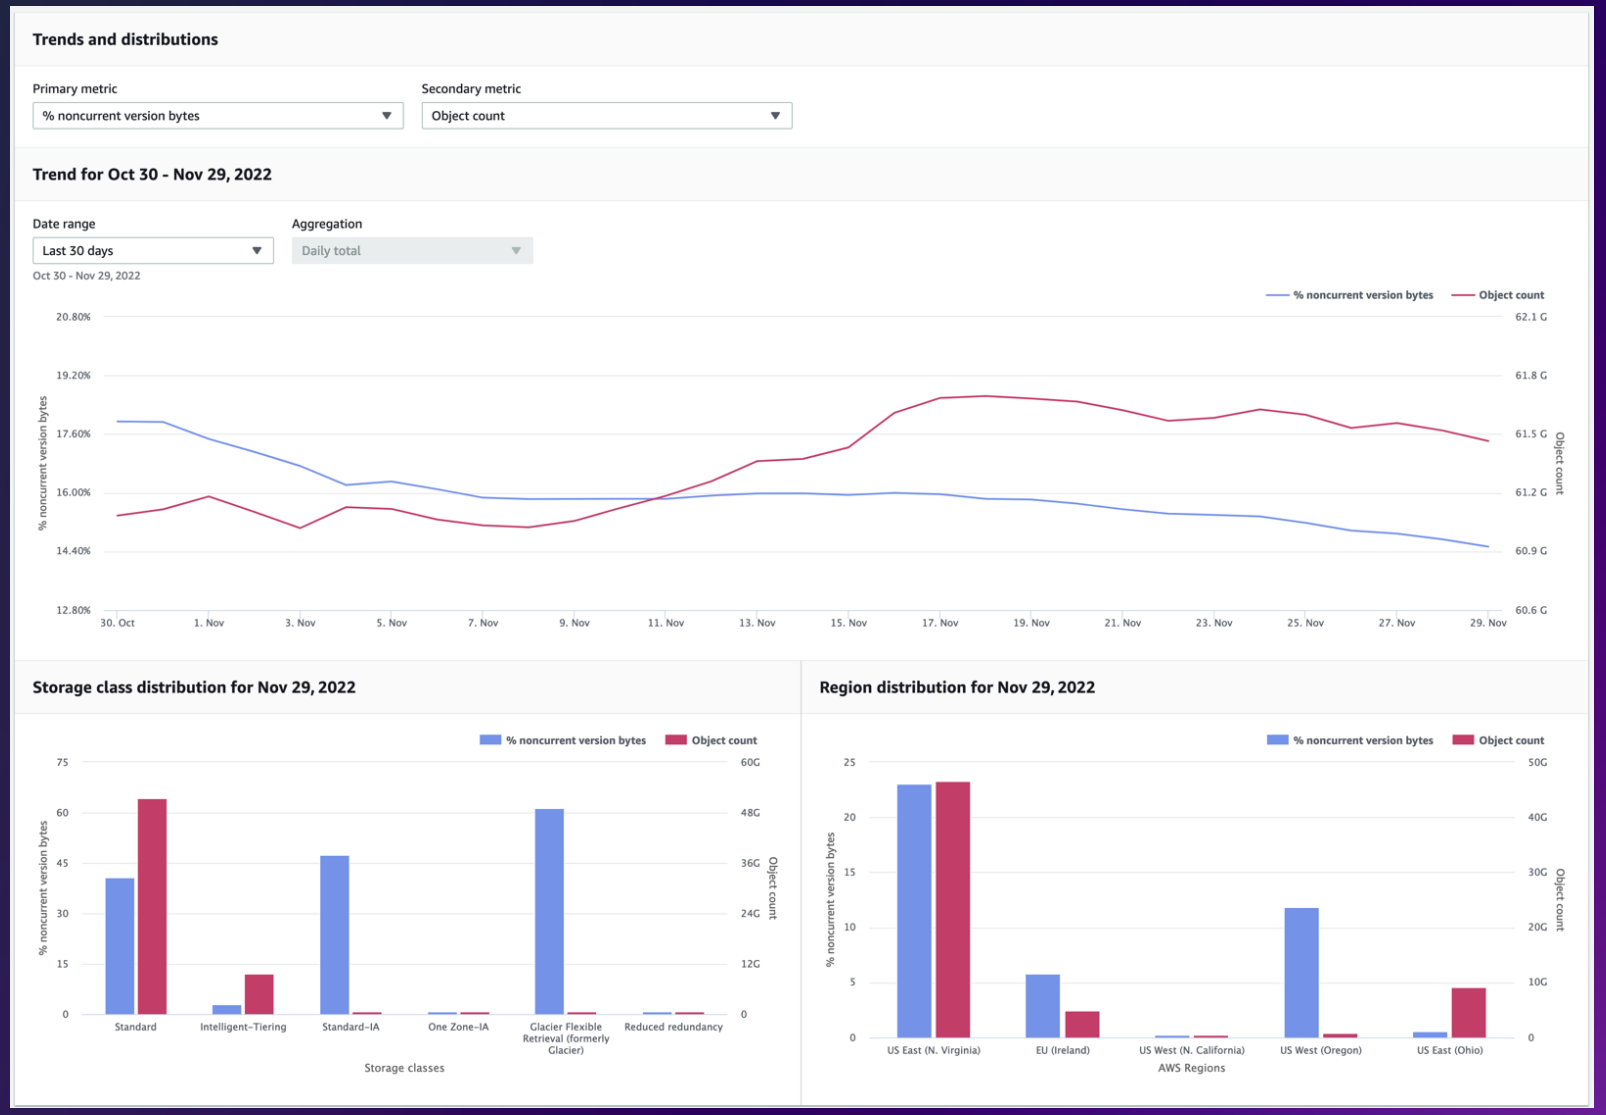

Publish metrics to Amazon CloudWatch

View S3 Storage Lens metrics in CloudWatch

- Centralized, cross-service dashboards
- Alarms and triggered actions
- Anomaly detection

Access metrics via CloudWatch API

- Connect AWS observability partners

S3 Inventory for object-level analysis

S3 Inventory report

S3 Inventory for object-level analysis

S3 Inventory report

Amazon Athena

```
SELECT key FROM your_table_name WHERE is_latest = ‘FALSE’ and size >= ‘10240’
```

Patterns of cost optimization

1. Data with known or predictable access patterns

2. Data with unknown or changing access patterns

Patterns of cost optimization

1. Data with known or predictable access patterns

2. Data with unknown or changing access patterns

Patterns of cost optimization

Data with known or predictable access patterns

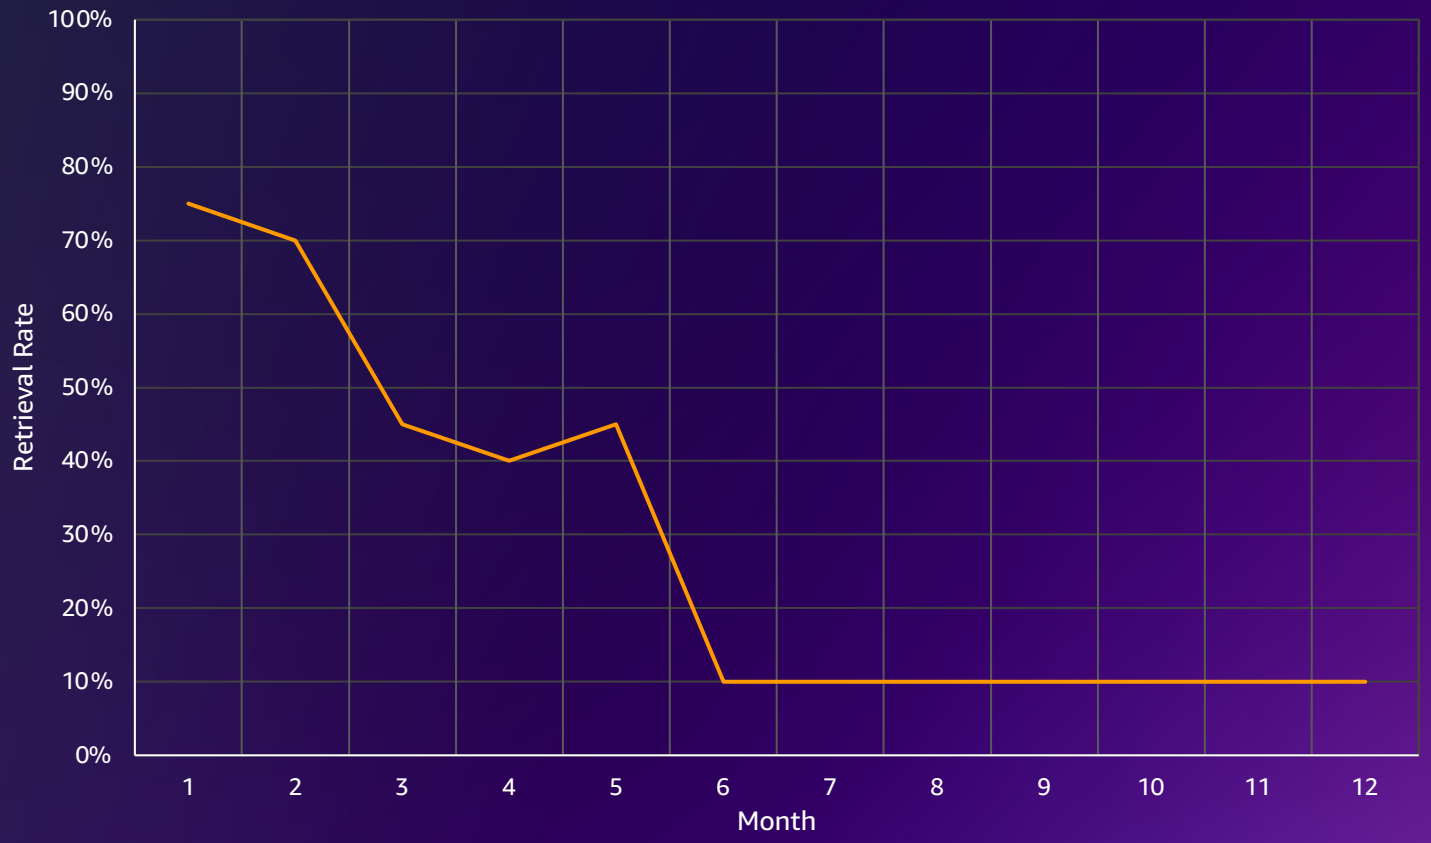

Workloads with predictable patterns often have low retrieval rates for long periods of time after becoming rarely accessed

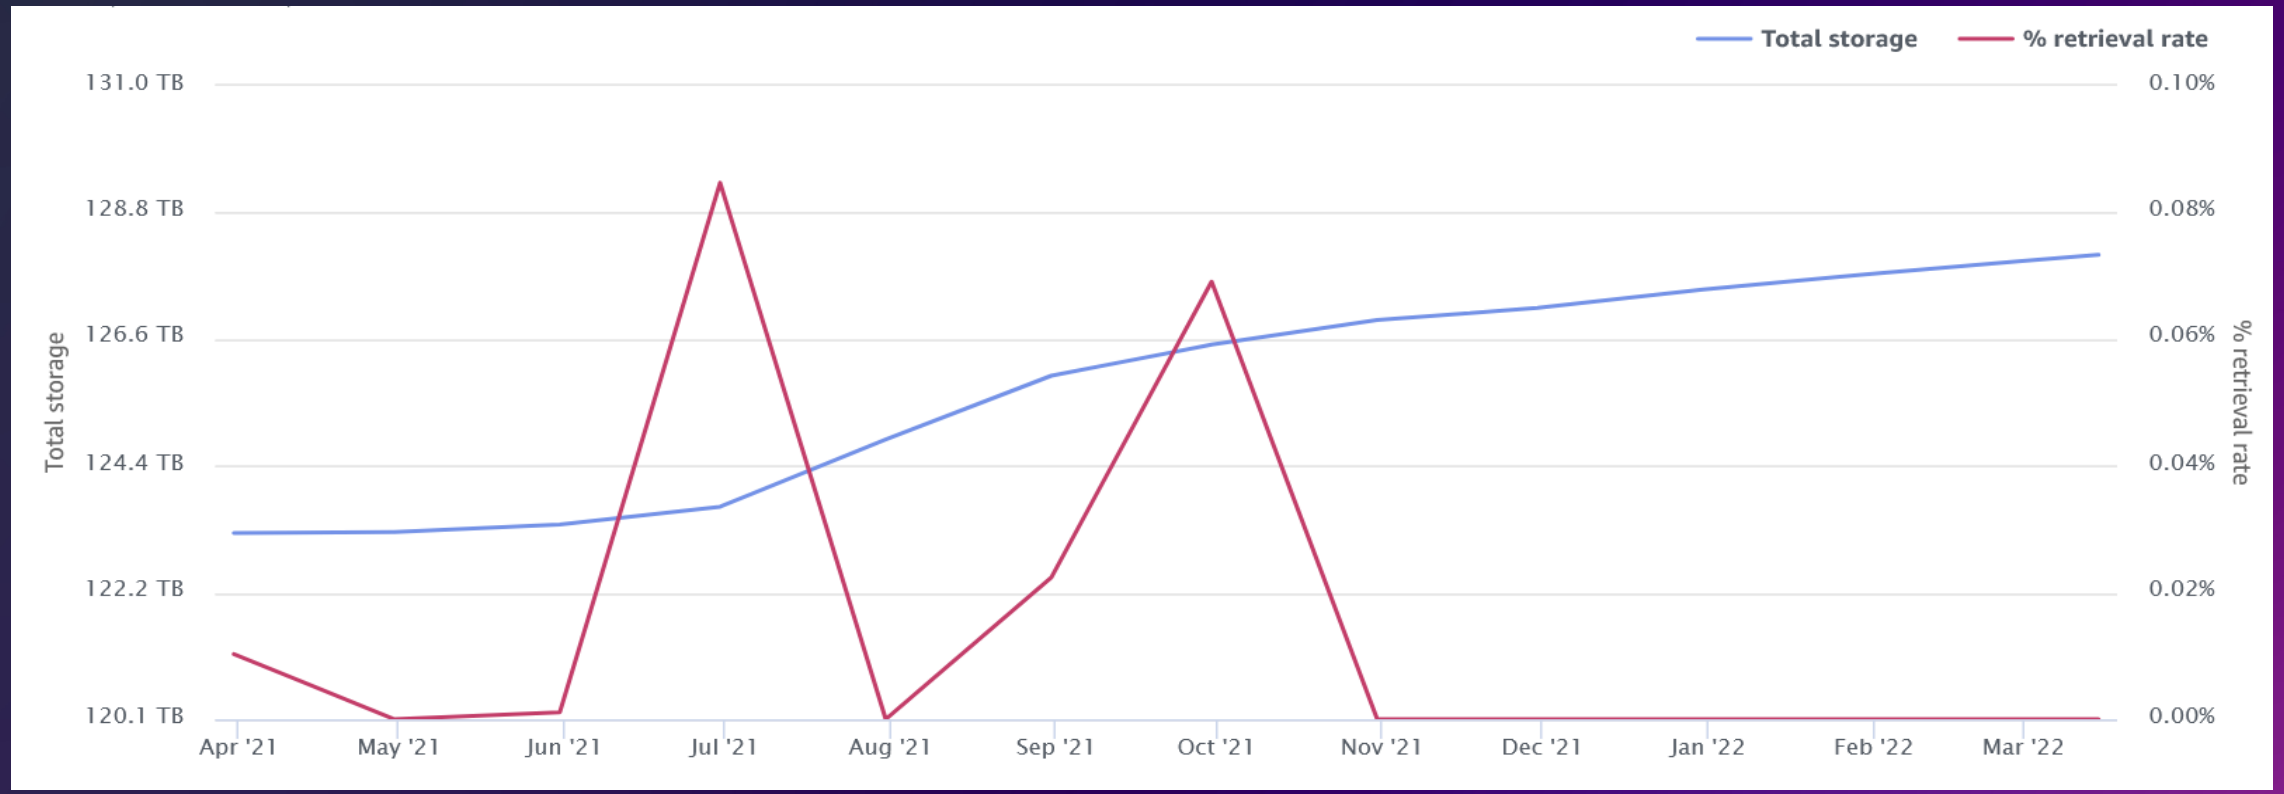


Medical Research Environment v2.0.556

Storage class choice matters at scale

S3 Standard	S3 Standard-IA	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive
				
Frequently accessed data	Infrequently accessed data	Rarely accessed data	Archive data	Long-term archive data

Milliseconds access

Minutes to hours access

Hot

Cold

Storage class choice matters at scale

S3 Standard	S3 Standard-IA	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive
				
Frequently accessed data	Infrequently accessed data	Rarely accessed data	Archive data	Long-term archive data
No min.	Min. 30 days	Min. 90 days	Min. 90 days	Min. 180 days

Milliseconds access

Minutes to hours access

Hot → Cold

Storage class choice matters at scale

S3 Standard	S3 Standard-IA	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive
				
Frequently accessed data	Infrequently accessed data	Rarely accessed data	Archive data	Long-term archive data
No min.	Min. 30 days	Min. 90 days	Min. 90 days	Min. 180 days

Milliseconds access

Minutes to hours access

Hot → Cold

Storage class choice matters at scale

S3 Standard	S3 Standard-IA	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive
				
Frequently accessed data	Infrequently accessed data	Rarely accessed data	Archive data	Long-term archive data
No min.	Min. 30 days	Min. 90 days	Min. 90 days	Min. 180 days

Milliseconds access

Minutes to hours access

Hot → Cold

Storage class choice matters at scale

S3 Standard	S3 Standard-IA	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive
				
Frequently accessed data	Infrequently accessed data	Rarely accessed data	Archive data	Long-term archive data
No min.	Min. 30 days	Min. 90 days	Min. 90 days	Min. 180 days

Milliseconds access

Minutes to hours access

Hot → Cold

Storage class choice matters at scale

S3 Standard	S3 Standard-IA	S3 Glacier Instant Retrieval	S3 Glacier Flexible Retrieval	S3 Glacier Deep Archive
				
Frequently accessed data	Infrequently accessed data	Rarely accessed data	Archive data	Long-term archive data
No min.	Min. 30 days	Min. 90 days	Min. 90 days	Min. 180 days

Milliseconds access

Minutes to hours access

Hot → Cold

S3 Glacier storage classes for archival data

S3 Glacier
Instant Retrieval

S3 Glacier
Flexible Retrieval

S3 Glacier
Deep Archive

Milliseconds to hours

S3 Glacier Flexible Retrieval improves data restore time by up to 85%

FAST RESTORE

Restores from S3 Glacier Flexible Retrieval now begin within minutes

AUTOMATIC

Automatically applies to Standard retrieval tier with S3 Batch Operations

NO ADDITIONAL COST

Available at no additional cost

S3 Glacier Flexible Retrieval – Standard retrieval performance on 250 objects, totaling 25 GB

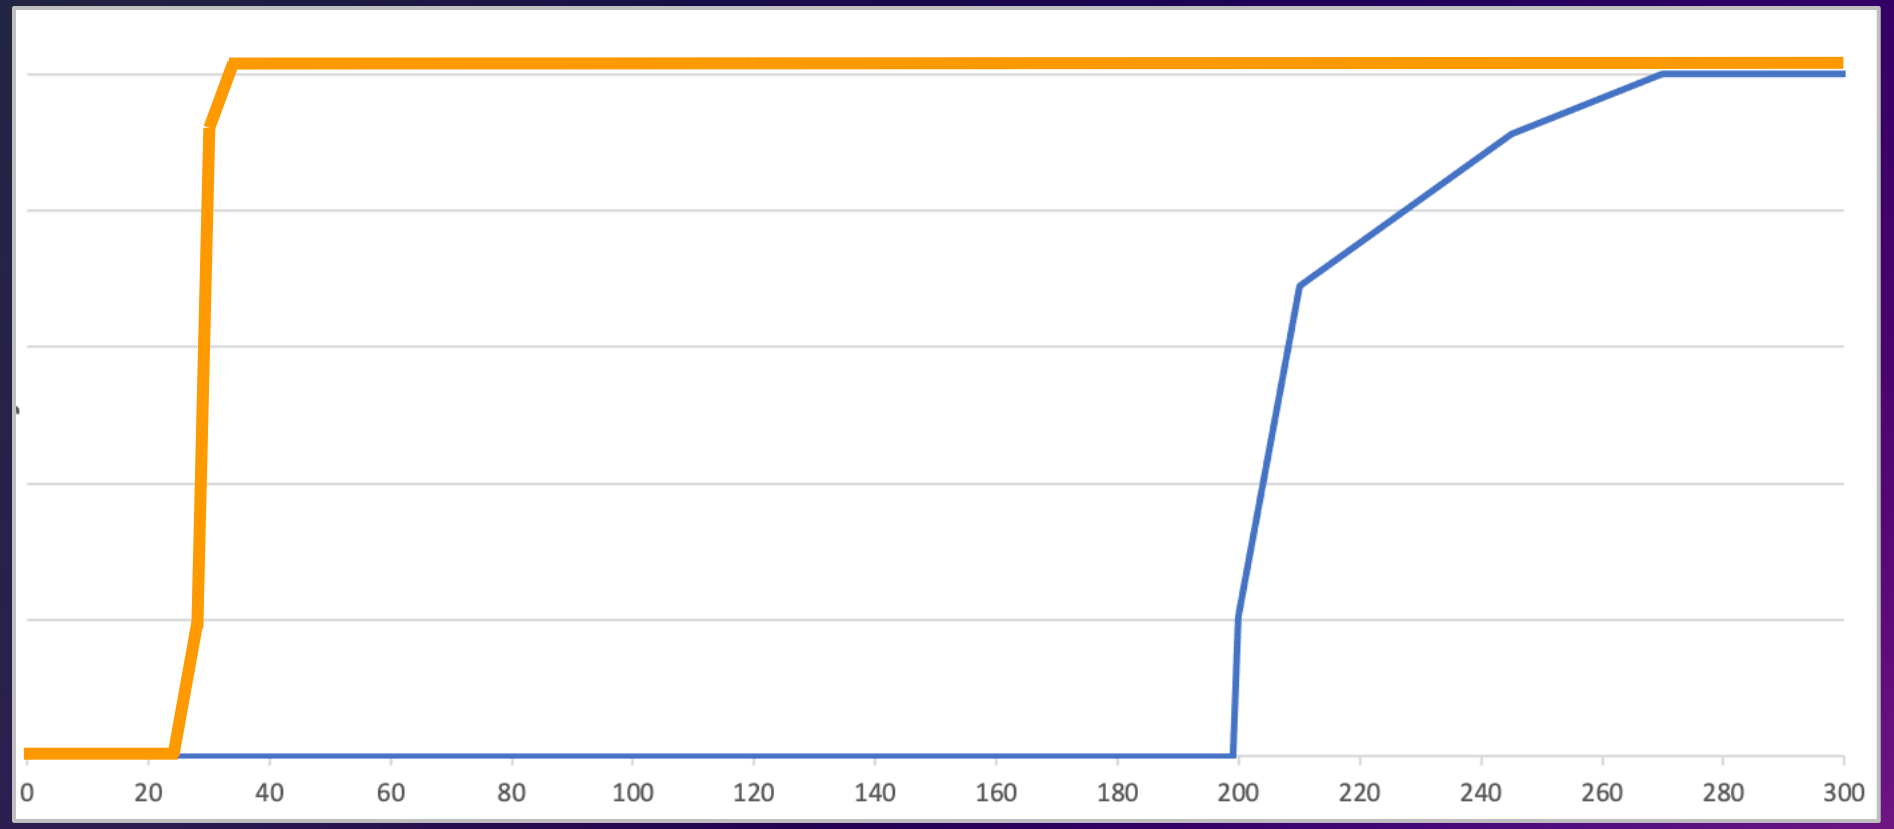

Selecting the right storage class

Frequency of access

Duration of storage

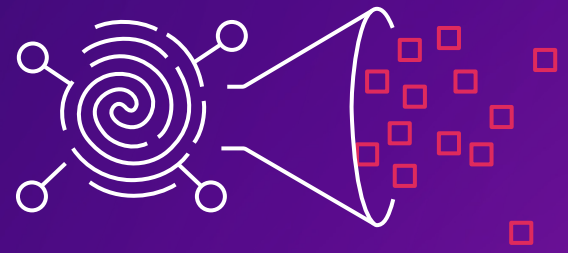

Performance

Impact of object size on Lifecycle cost

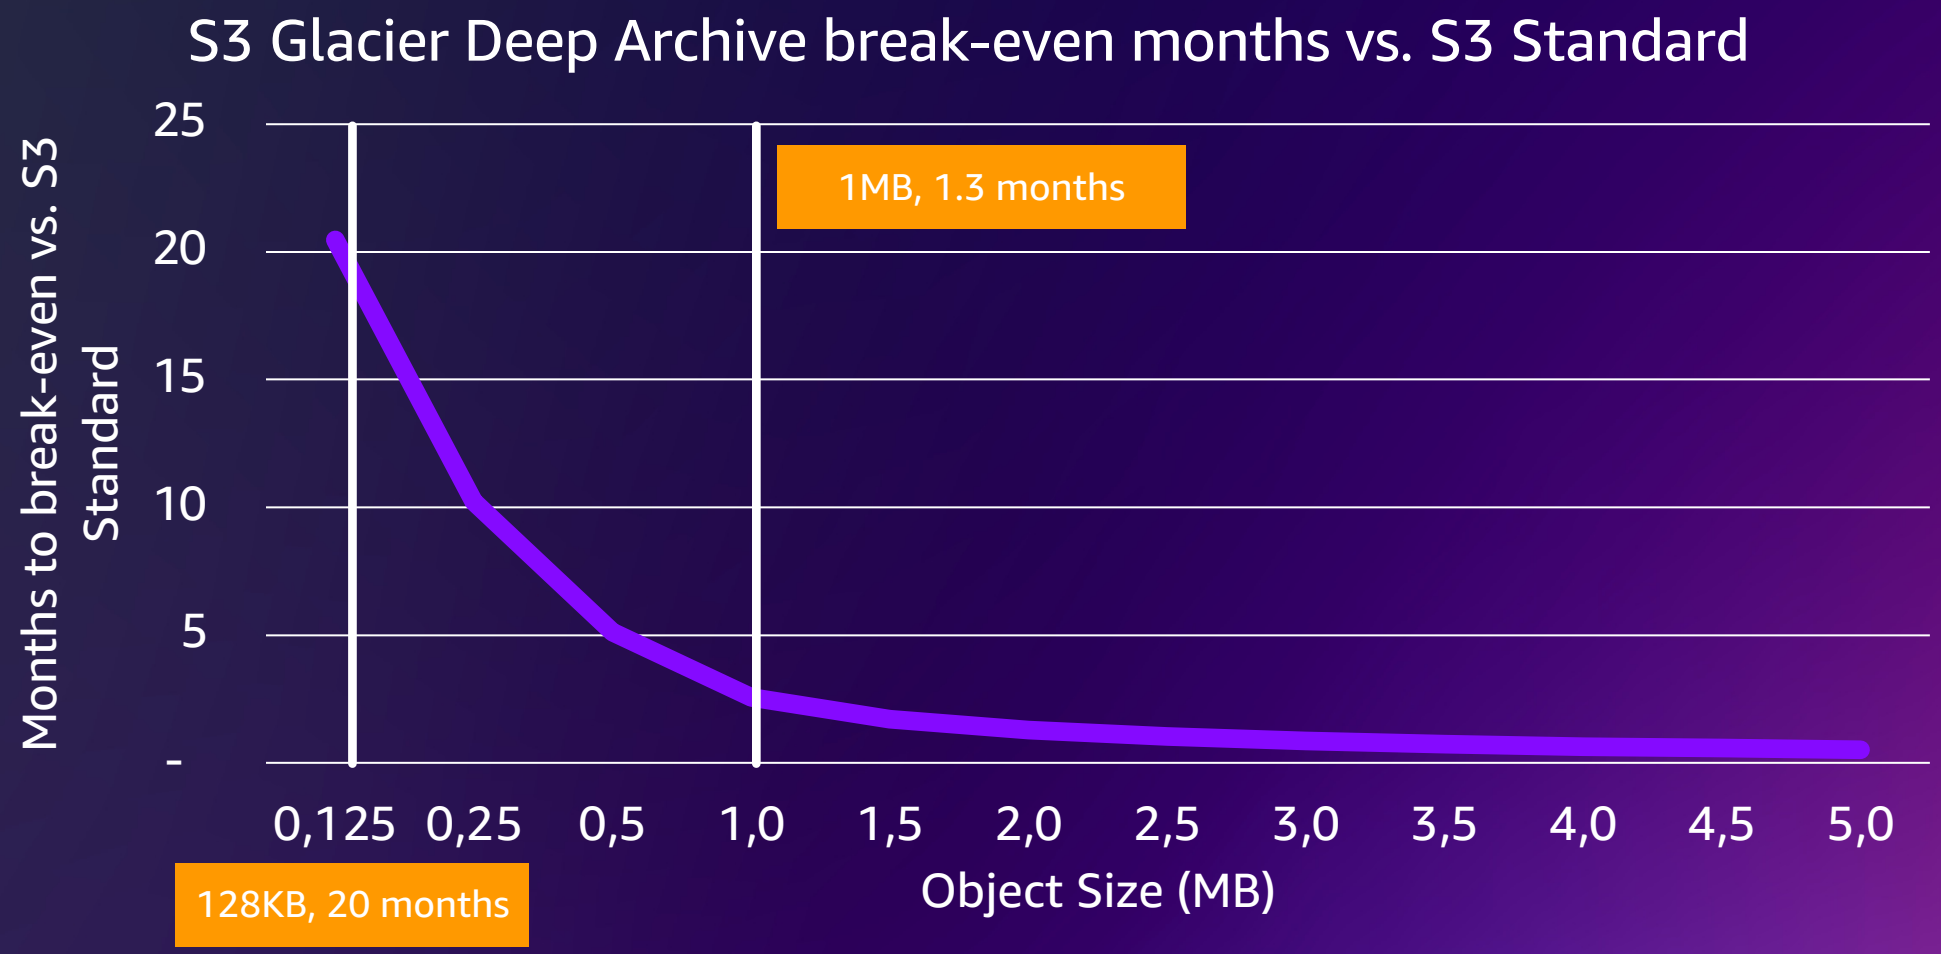

Fine tune your Lifecycle policies

S3 Standard

S3 Glacier
Instant Retrieval

S3 Glacier
Deep Archive

Fine tune your Lifecycle policies

S3 Standard

90 days

S3 Glacier
Instant Retrieval

365 days

S3 Glacier
Deep Archive

Majority of data has **unknown access patterns**

S3 Standard

S3 Standard-IA

S3 Glacier
Instant Retrieval

S3 Glacier
Flexible Retrieval

S3 Glacier
Deep Archive

Frequent

Infrequent

Archive Instant

Archive

Deep Archive

Automatically optimize each & every object

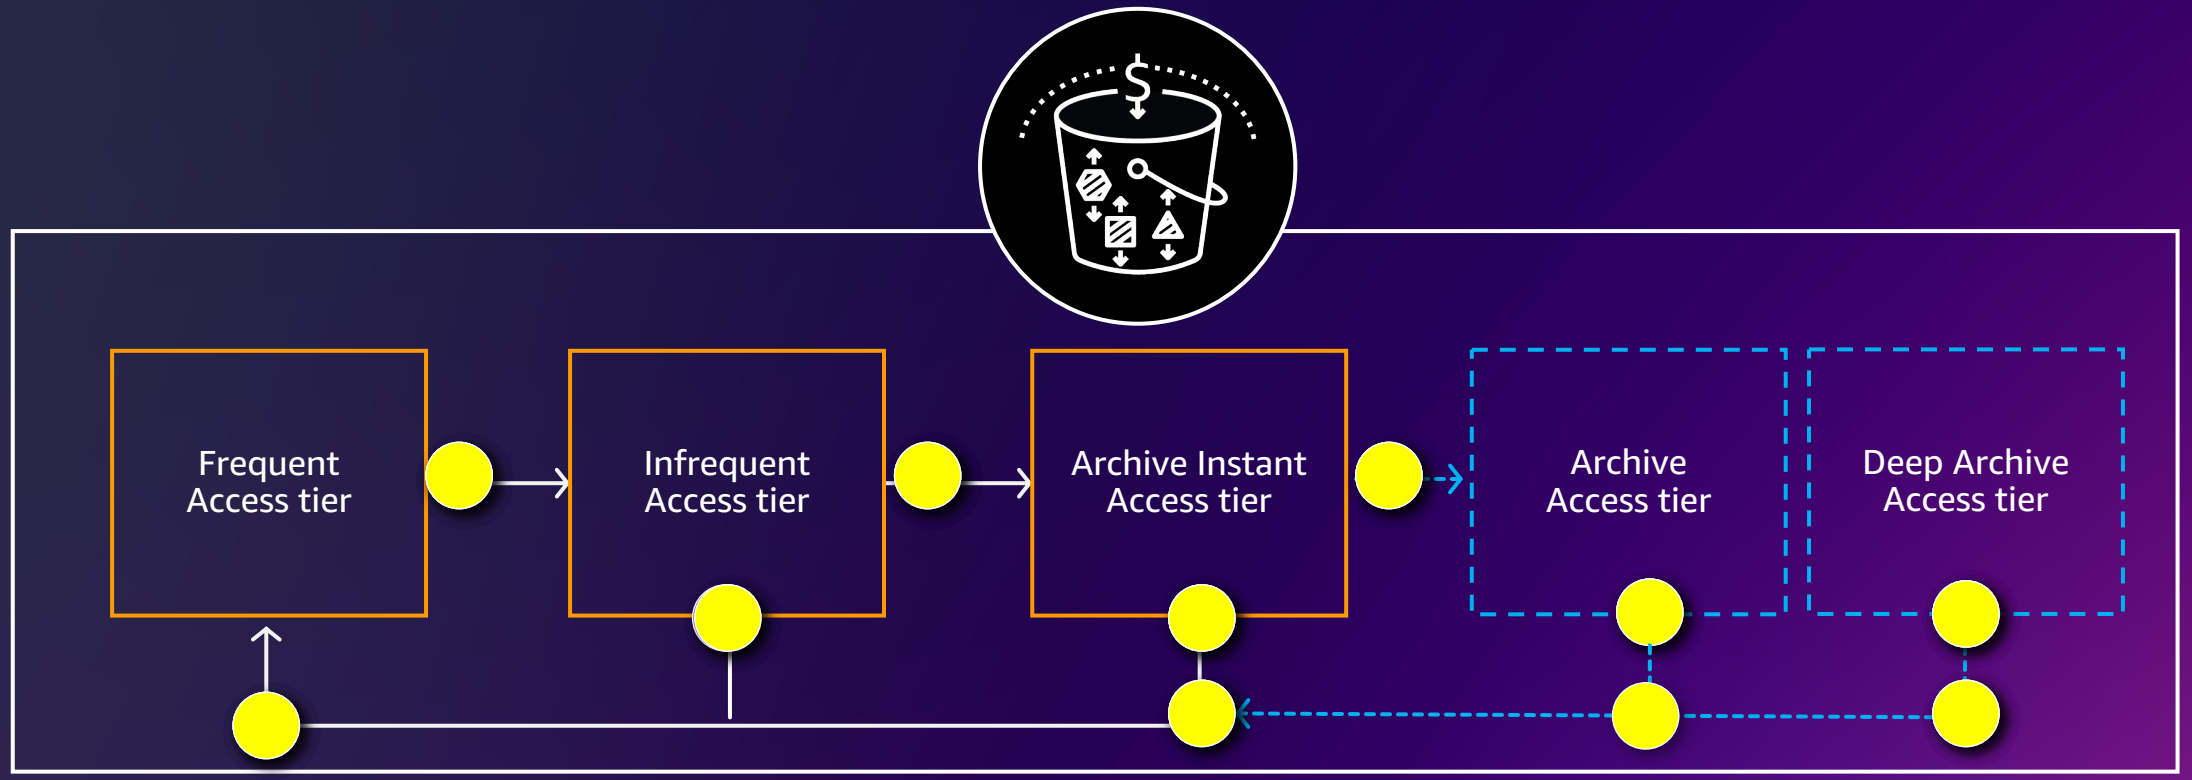

Milliseconds access (automatic)

Minutes to hours (optional)

Automatically optimize each & every object

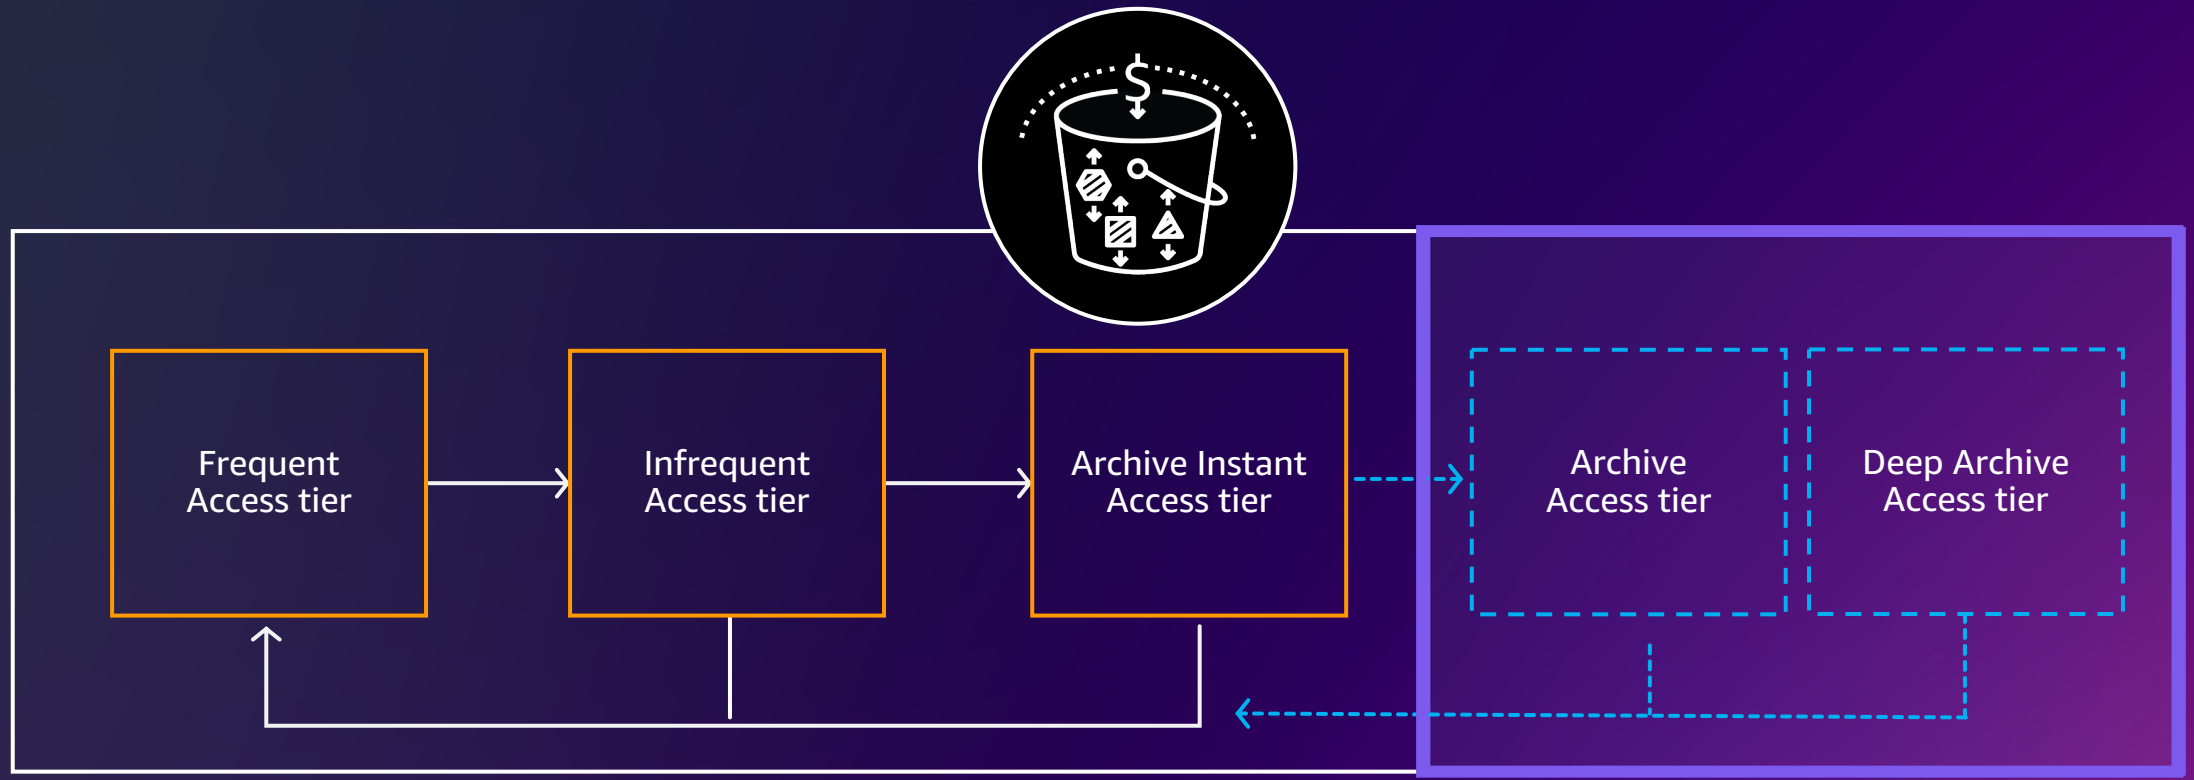

Milliseconds access (automatic)

Minutes to hours (optional)

Use S3 Event Notifications to track objects moved to the asynchronous archive tiers

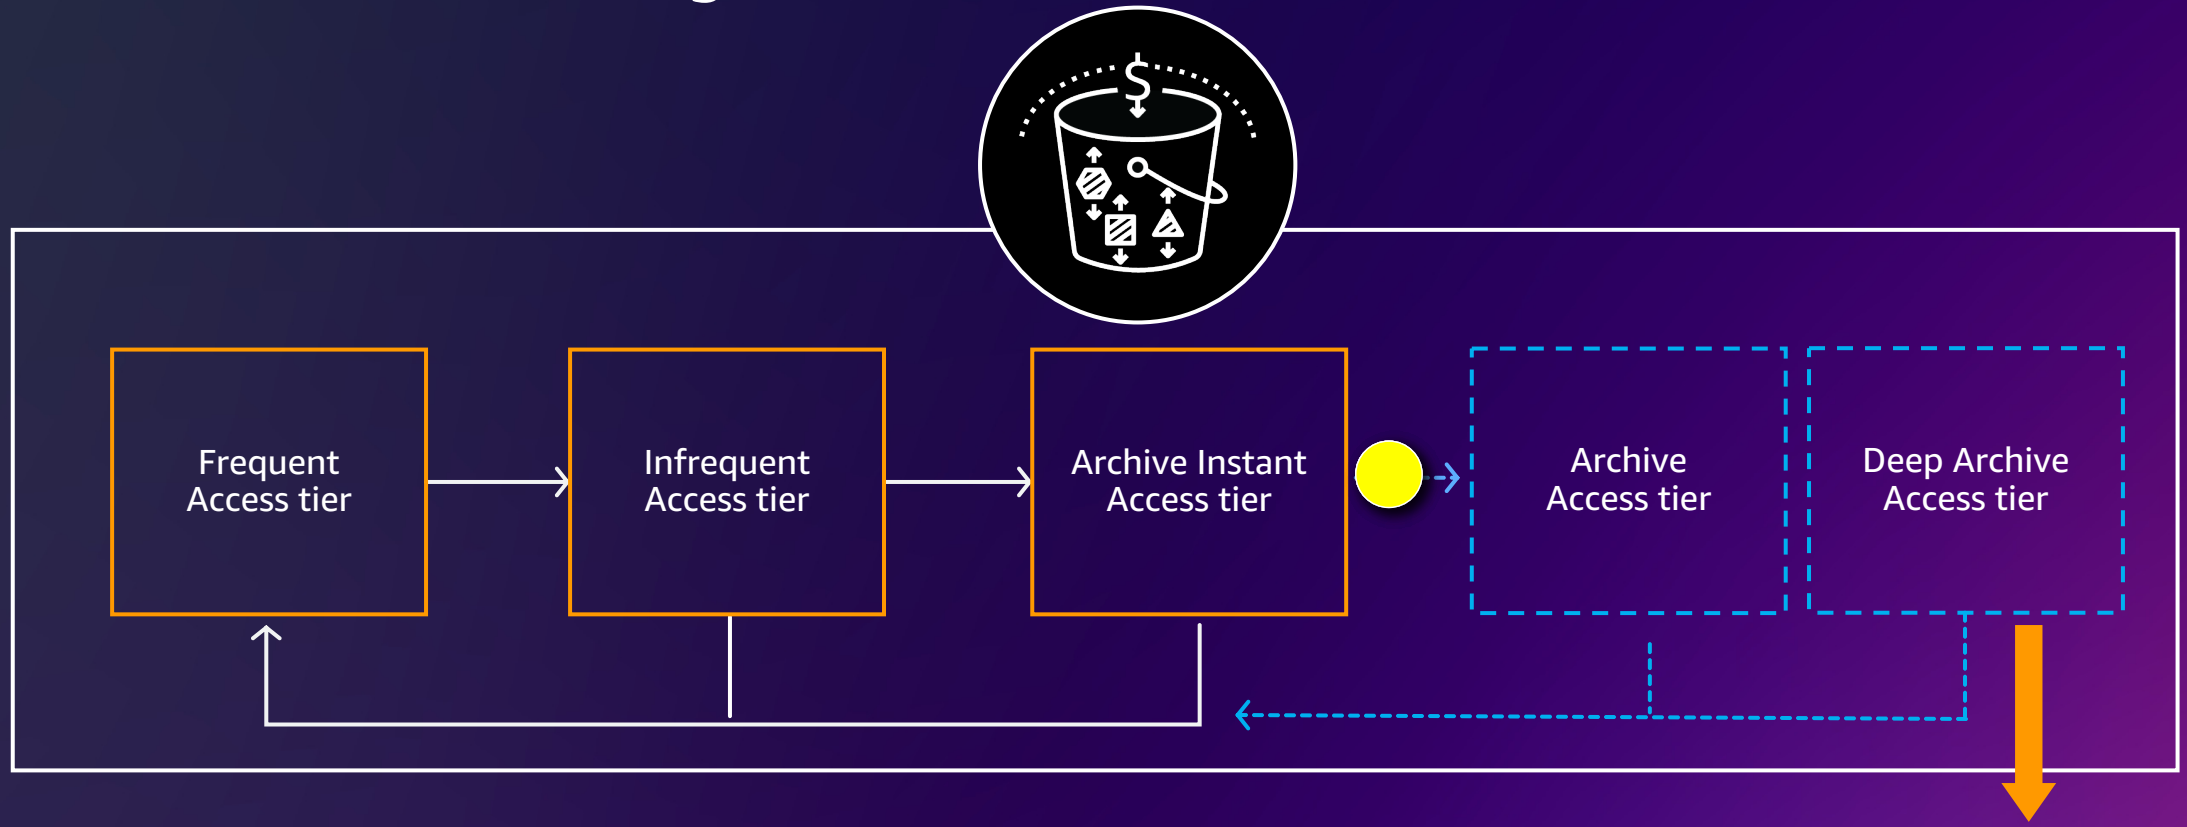

After just 3 months of using S3 Intelligent Tiering, Illumina began to see significant monthly cost savings. For every 1 TB of data, the company saves 60 percent on storage costs. Further, Illumina can provide its customers with near-instant access to thousands of whole genome sequences at a low, competitive cost, helping its customers accelerate their research and development.

illumina®

Scaling to
tens of millions
of requests per second

Characteristics of storage performance

REQUEST RATE

REQUEST LATENCY

THROUGHPUT

Prefix design for scale

3,500

PUT/COPY/POST/DELETE

5,500

GET/HEAD

per prefix

Designing prefixes for “bursty” workloads

/daily-uploads/{date}/{car_id}/drive-data

/daily-uploads/20241127/01/drive-data
/daily-uploads/20241127/02/drive-data
/daily-uploads/20241127/03/drive-data
/daily-uploads/20241127/04/drive-data

Designing prefixes for “bursty” workloads

/daily-uploads/{date}/{car_id}/drive-data

/daily-uploads/20241127/

01/drive-data 02/drive-data 03/drive-data

Designing prefixes for “bursty” workloads

/daily-uploads/{date}/{car_id}/drive-data

/daily-uploads/20241127/

01/drive-data 02/drive-data 03/drive-data

/daily-uploads/20241128/

Designing prefixes for “bursty” workloads

/daily-uploads/{car_id}/{date}/drive-data

Designing prefixes for “bursty” workloads

/daily-uploads/{car_id}/{date}/drive-data

/daily-uploads/

01/20241125/drive-data

02/20241125/drive-data

Measuring 503 slowdown errors

503 Slow Down typically indicates that the number of requests to your S3 bucket is very high

"AmazonS3Exception: Slow Down (Service: Amazon S3; Status Code: 503; Error Code: 503 Slow Down; Request ID: A4DBBEXAMPLE2C4D)"

Amazon CloudWatch metrics

S3 Storage Lens advanced metrics

S3 server access logs

Controlling latencies and maximizing throughput

Tune your timeouts to your environment

Retry slow requests

Parallelize requests

Use multiple connections

Controlling latencies and maximizing throughput

Tune your timeouts to your environment

Retry slow request

Parallelize requests

Use multiple connections

Parallelization of individual operations

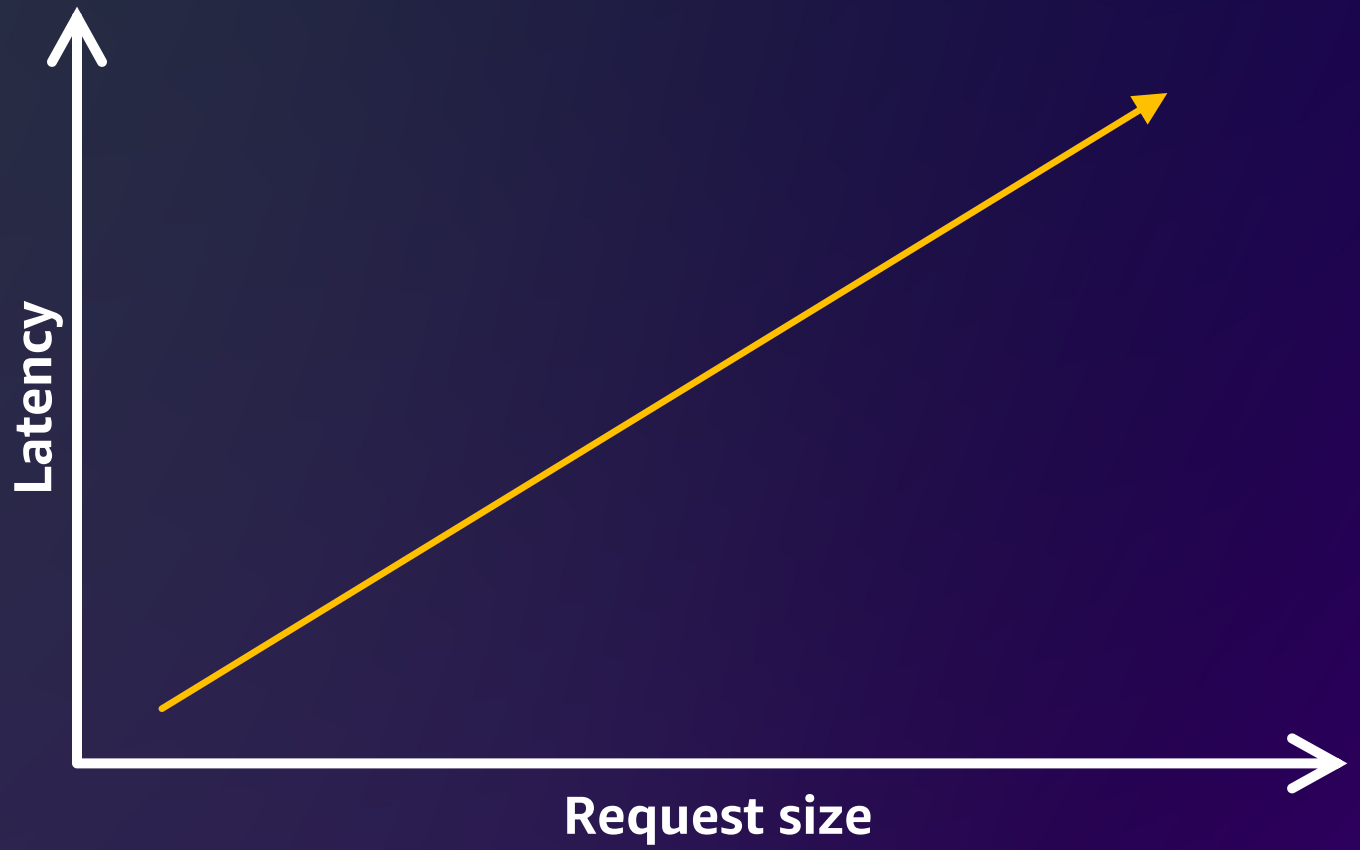

Request time grows as a function of bytes transferred

Parallelize large object uploads and downloads to reduce latency and improve throughput

Parallelization of individual operations

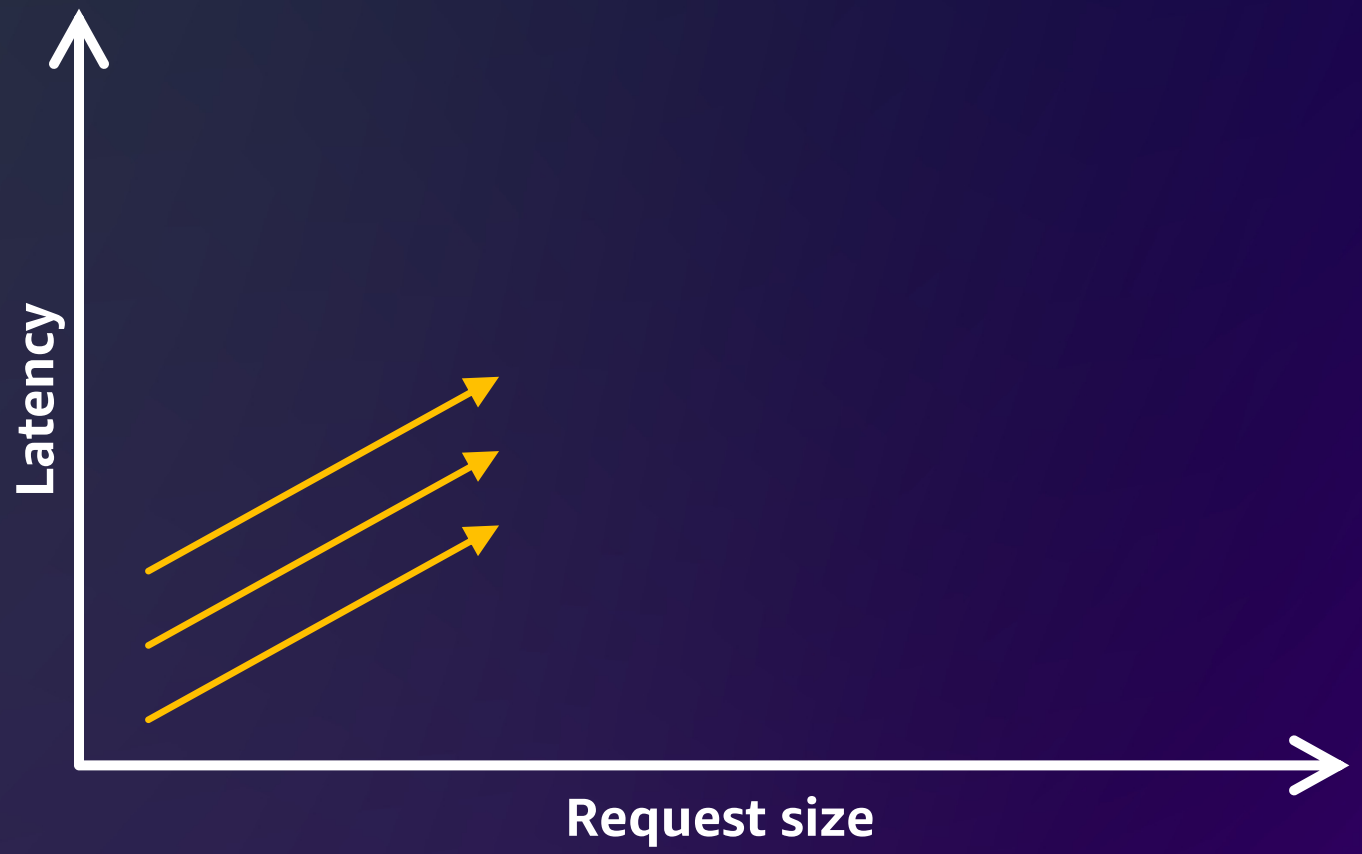

Use multipart upload when uploading large objects

Use byte-range fetches when retrieving them

Parallelize large object uploads and downloads to reduce latency and improve throughput

But what if you don’t want to do any of this?

AWS Common Runtime (CRT)

Implements Amazon S3's performance best practices to deliver high data transfer rates

Timeouts, retries, DNS load balancing, and request parallelization

Out of the box performance optimizations

The Common Runtime in action

Using the AWS CLI to
download 100 GB of data

```
$ time aws s3 sync --quiet s3://s3-crt-
performance-demonstration/100_1GB_files/ ./

real    8m28.668s
user    3m11.200s
sys     1m44.393s
```

After enabling the CRT

```
$ aws configure set
default.s3.preferred_transfer_client crt
$ time aws s3 sync --quiet s3://s3-crt-
performance-demonstration/100_1GB_files/ ./

real    3m52.401s
user    1m23.798s
sys     1m50.750s
```

The Common Runtime in action

Using the AWS CLI to download 100 GB of data

```
$ time aws s3 sync --quiet s3://s3-crt-performance-demonstration/100_1GB_files/ ./

real    8m28.668s
user    3m11.200s
sys     1m44.393s
```

After enabling the CRT

```
$ aws configure set default.s3.preferred_transfer_client crt
$ time aws s3 sync --quiet s3://s3-crt-performance-demonstration/100_1GB_files/ ./

real    3m52.401s
user    1m23.798s
sys     1m50.750s
```

2.2 x faster

The **AWS Common Runtime** on by default

Default performance optimizations when accessing Amazon S3 from Amazon EC2 Trn1, P4d, and P5 instances

Amazon EC2 Trn1, P4d, and P5 instances are ideal for generative AI models, including large language and diffusion models

Mountpoint for Amazon S3 to access access objects using file APIs

High performance open source file client for Amazon S3

Presents Amazon S3 objects as files in the local file system

Translates local file system API calls to REST API calls on Amazon S3 objects

Achieves high single-instance transfer rates, finishing jobs faster

aws

Mountpoint for Amazon S3 optimizes for repeated data access requests

Reduce request cost and accelerate performance for repeat read requests

Cache data in EC2 instance storage, instance memory, or an EBS volume

Reduce request cost and response times for repeated data access requests

Mountpoint for Amazon S3 in action

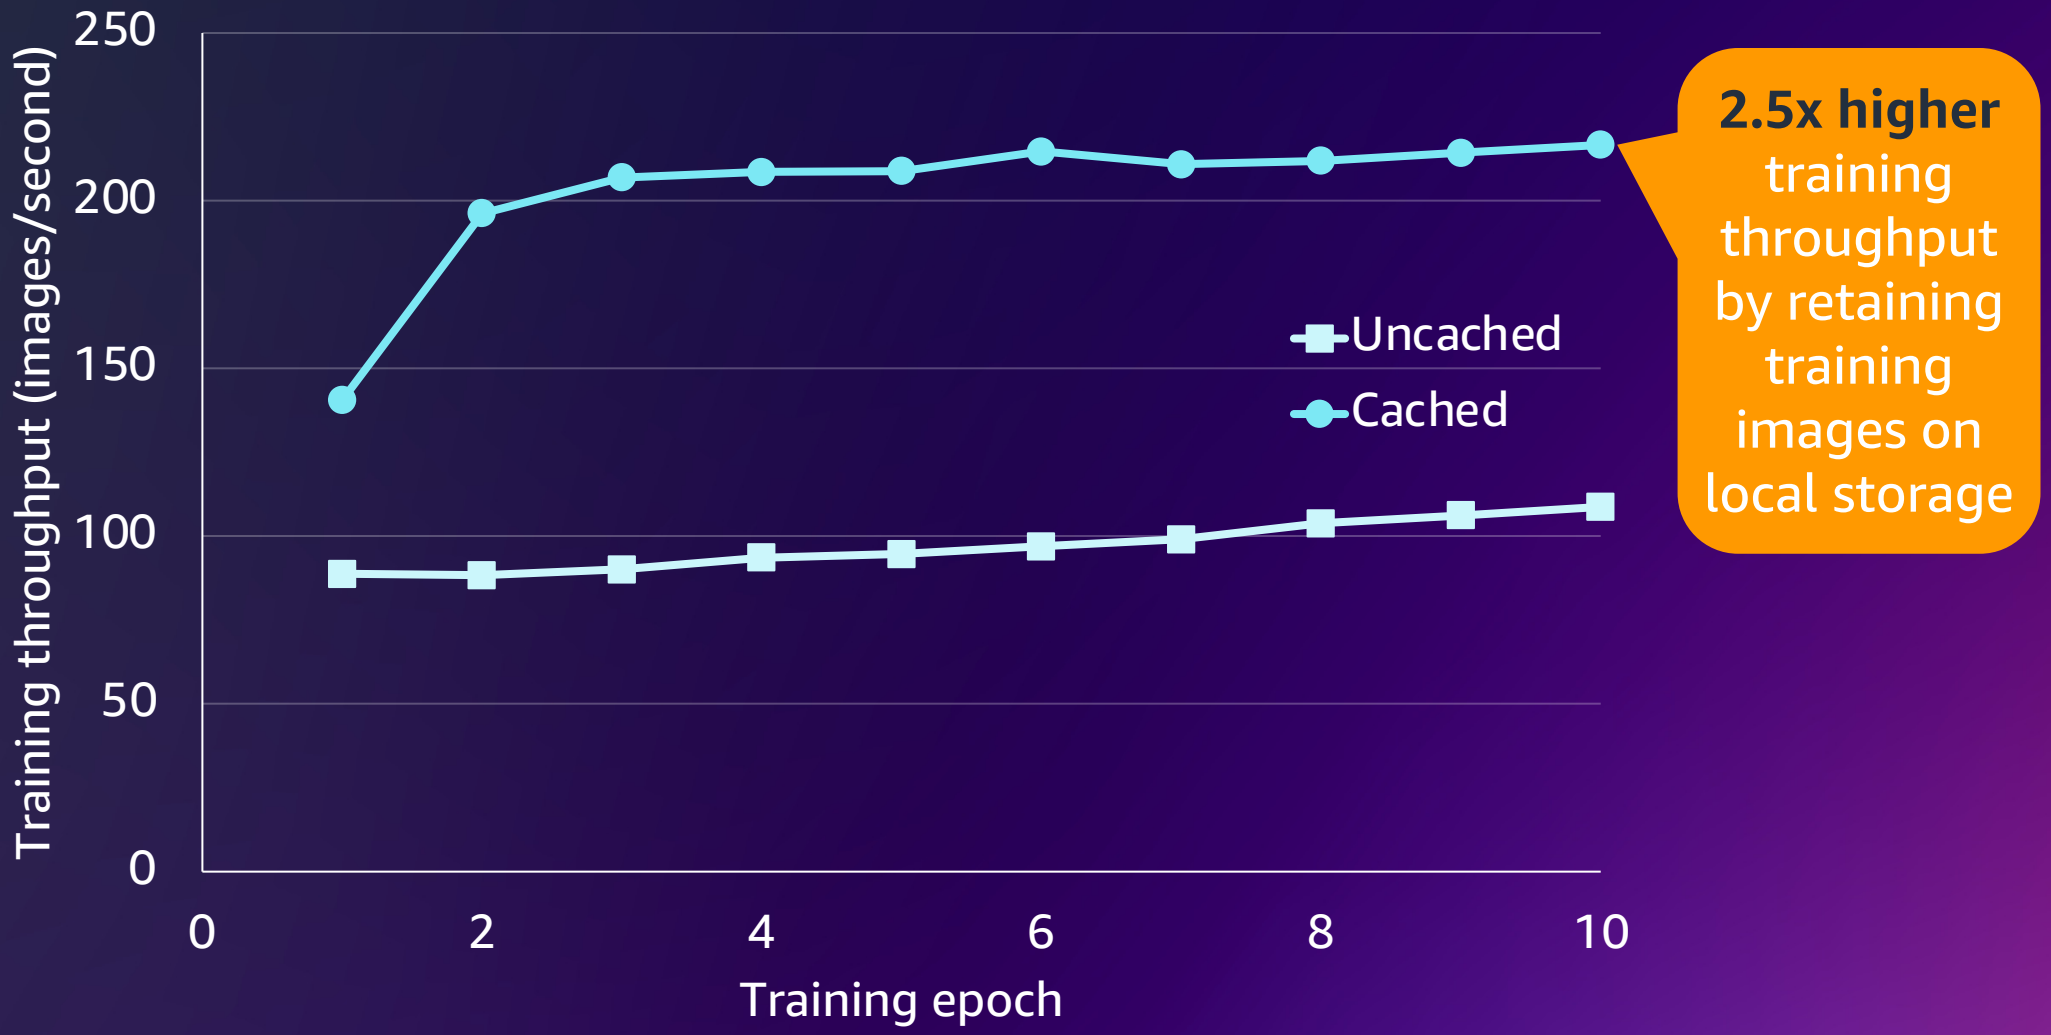

Tips and Tricks

1. Optimize object size
2. Horizontally scale across multiple prefixes
3. The CRT is your friend
4. Think async

Putting it all together

Monitor and manage data with Amazon S3 tools

Optimize storage with purpose-built storage classes

Achieve high performance with Amazon S3

Thank you!

Dragoș Mădărășan

Please complete the session survey in the Events App

App Store

Google Play Store