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

STG314-R

Amazon EBS snapshots: Build protection and cost-optimize

Vienna Chen (she/her)

Senior Product Manager Technical, Snapshots
AWS

Kevin McDonald (he/him)

Sr. Storage Specialist Solution Architect
AWS

Vienna Chen (she/her)

Senior Product Manager Technical, Snapshots
Amazon Web Services

Kevin McDonald (he/him)

Sr. Storage Specialist Solution Architect
Amazon Web Services

Shruti Gupta (she/her)

Senior Product Manager Technical, EBS
Amazon Web Services

Dan Booth (he/him)

Enterprise Support Lead
Amazon Web Services

Agenda

Intro to Amazon EBS snapshots

Amazon Data Lifecycle Manager

Recycle Bin for Amazon EBS snapshots and AMIs

Amazon EBS Snapshots Archive

Amazon EBS direct APIs

Hands-on lab

What are Amazon EBS snapshots?

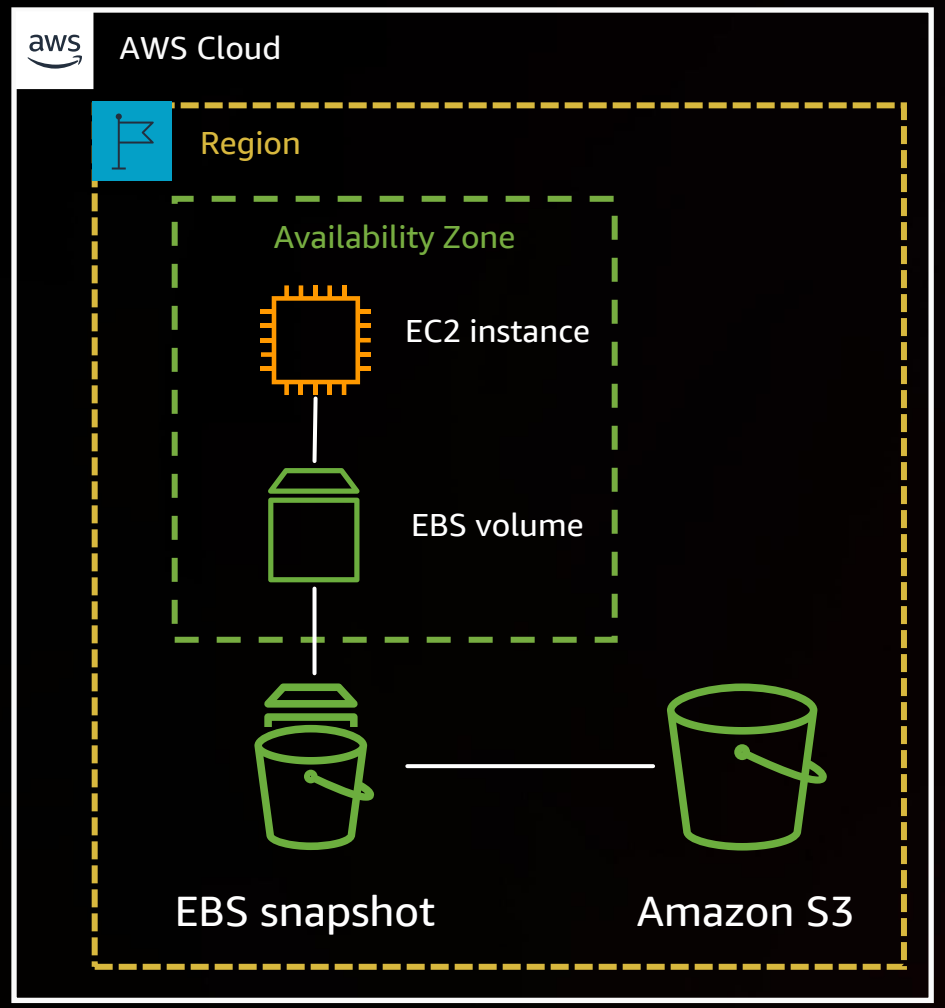

- **Point-in-time backups** of EBS volumes
- Stored on Amazon S3
- Properties
  - **Incremental** – only changed blocks stored
  - **Crash consistent** – completed I/Os persisted in next snapshot
  - Crash-consistent snapshots can be taken with **1 API call for a subset of volumes attached to an instance**
  - Can be **shared and copied** across accounts and AWS Regions

How are snapshots used?

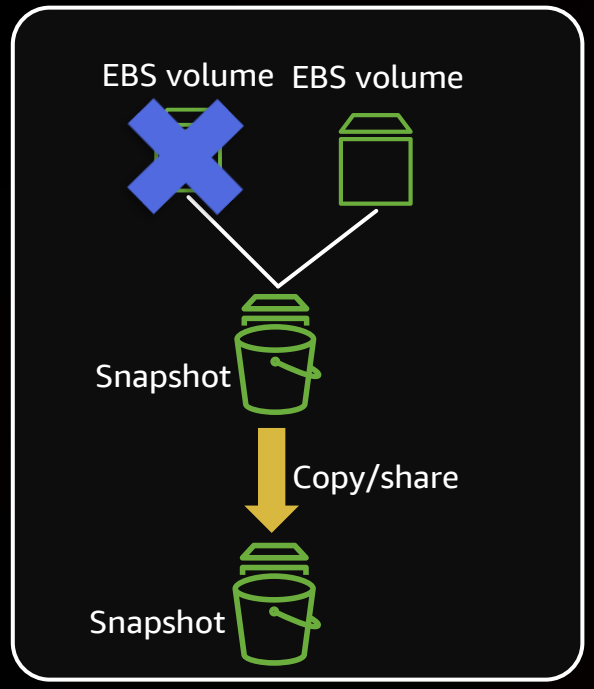

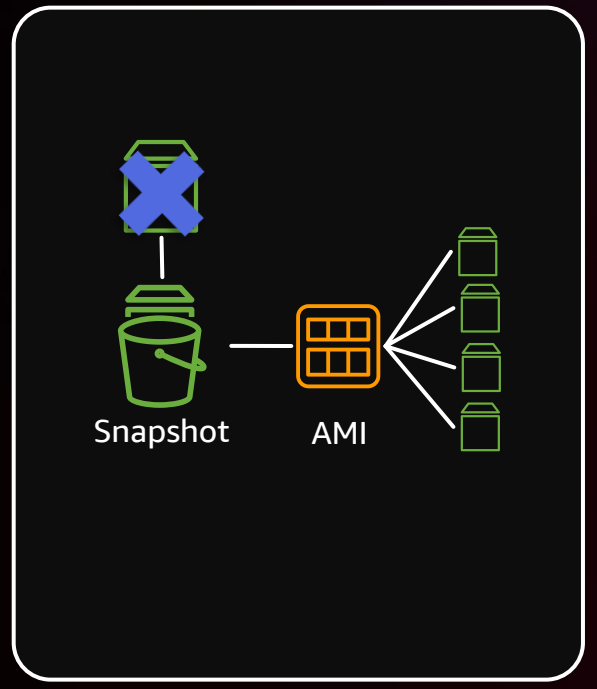

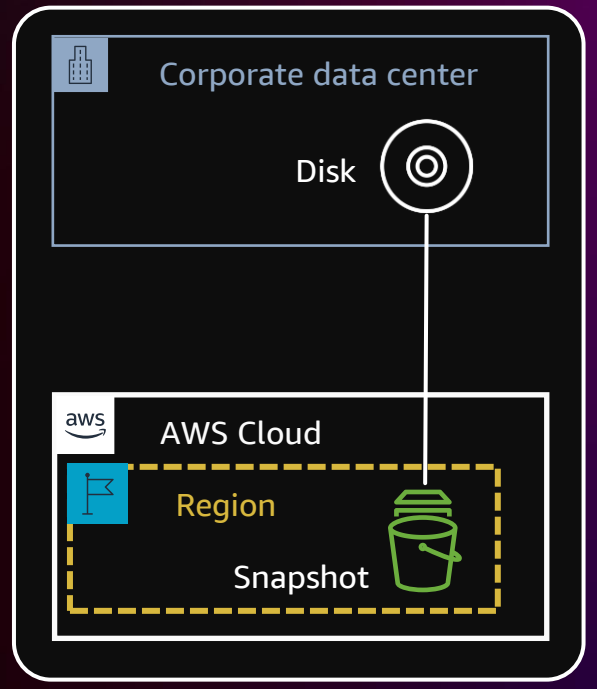

EBS backup and disaster recovery

Refresh, scale-up, data handoff workflows

Non-EBS backup and migration

Snapshots basics

How does an EBS snapshot work?

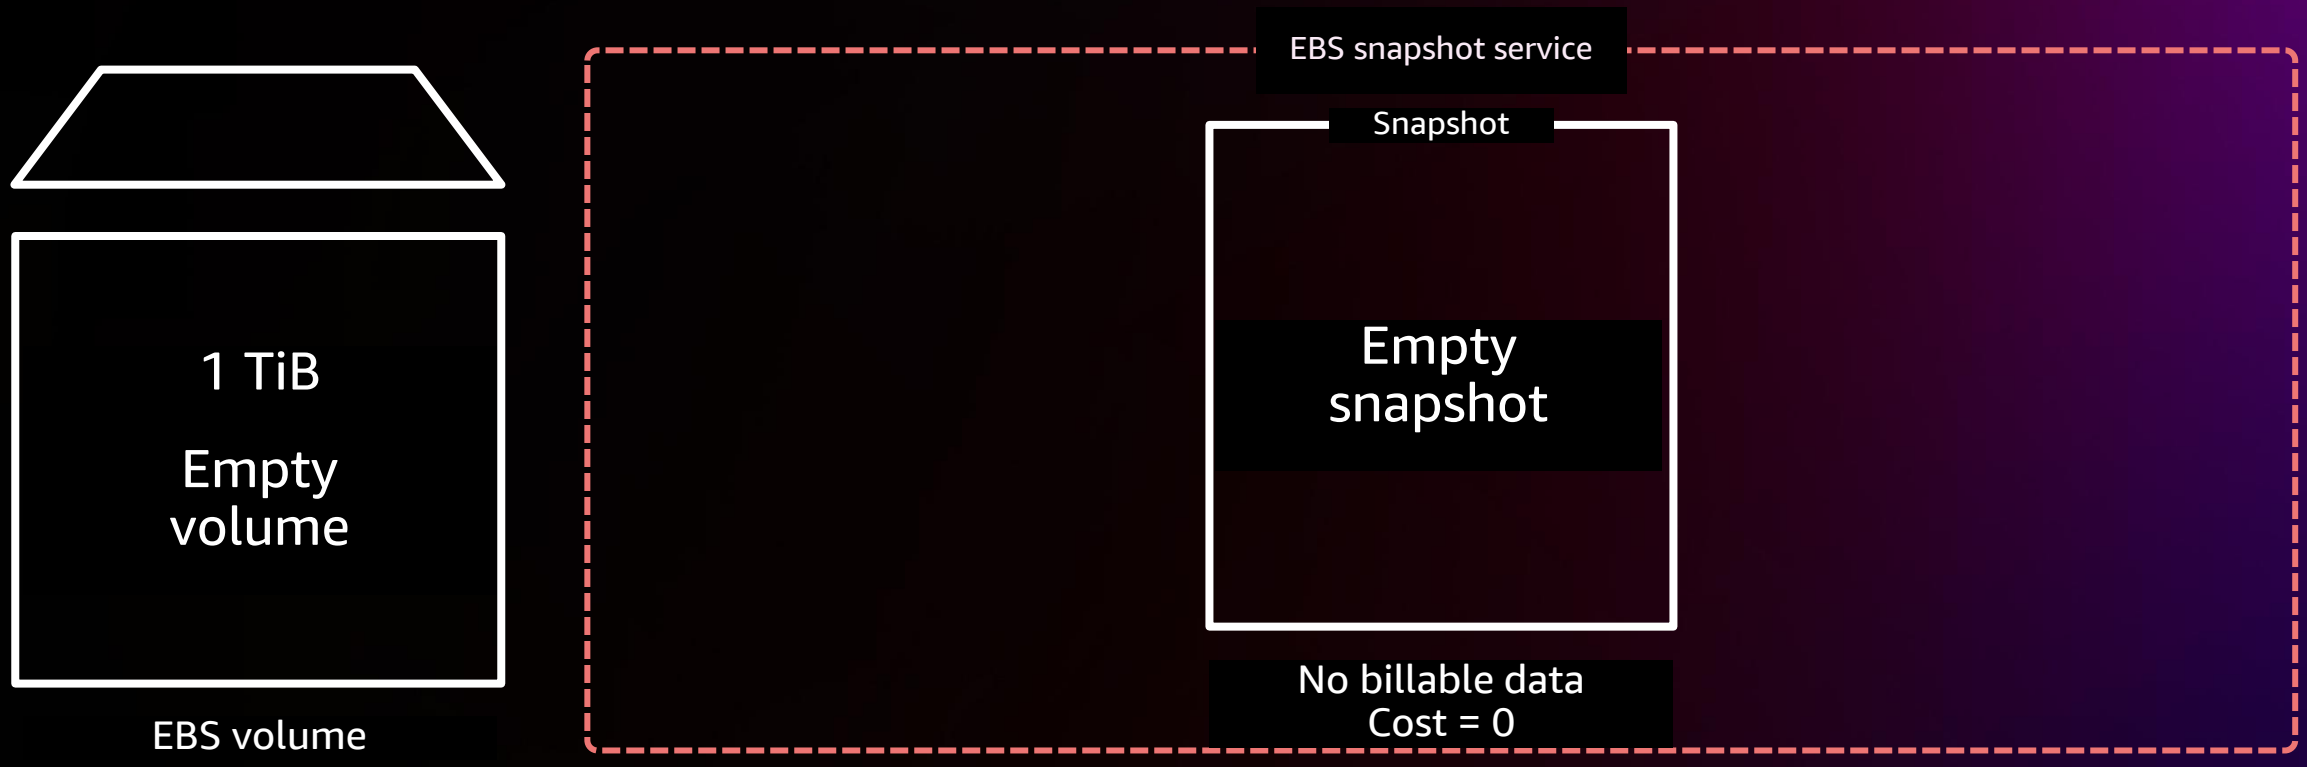

How does an EBS snapshot work?

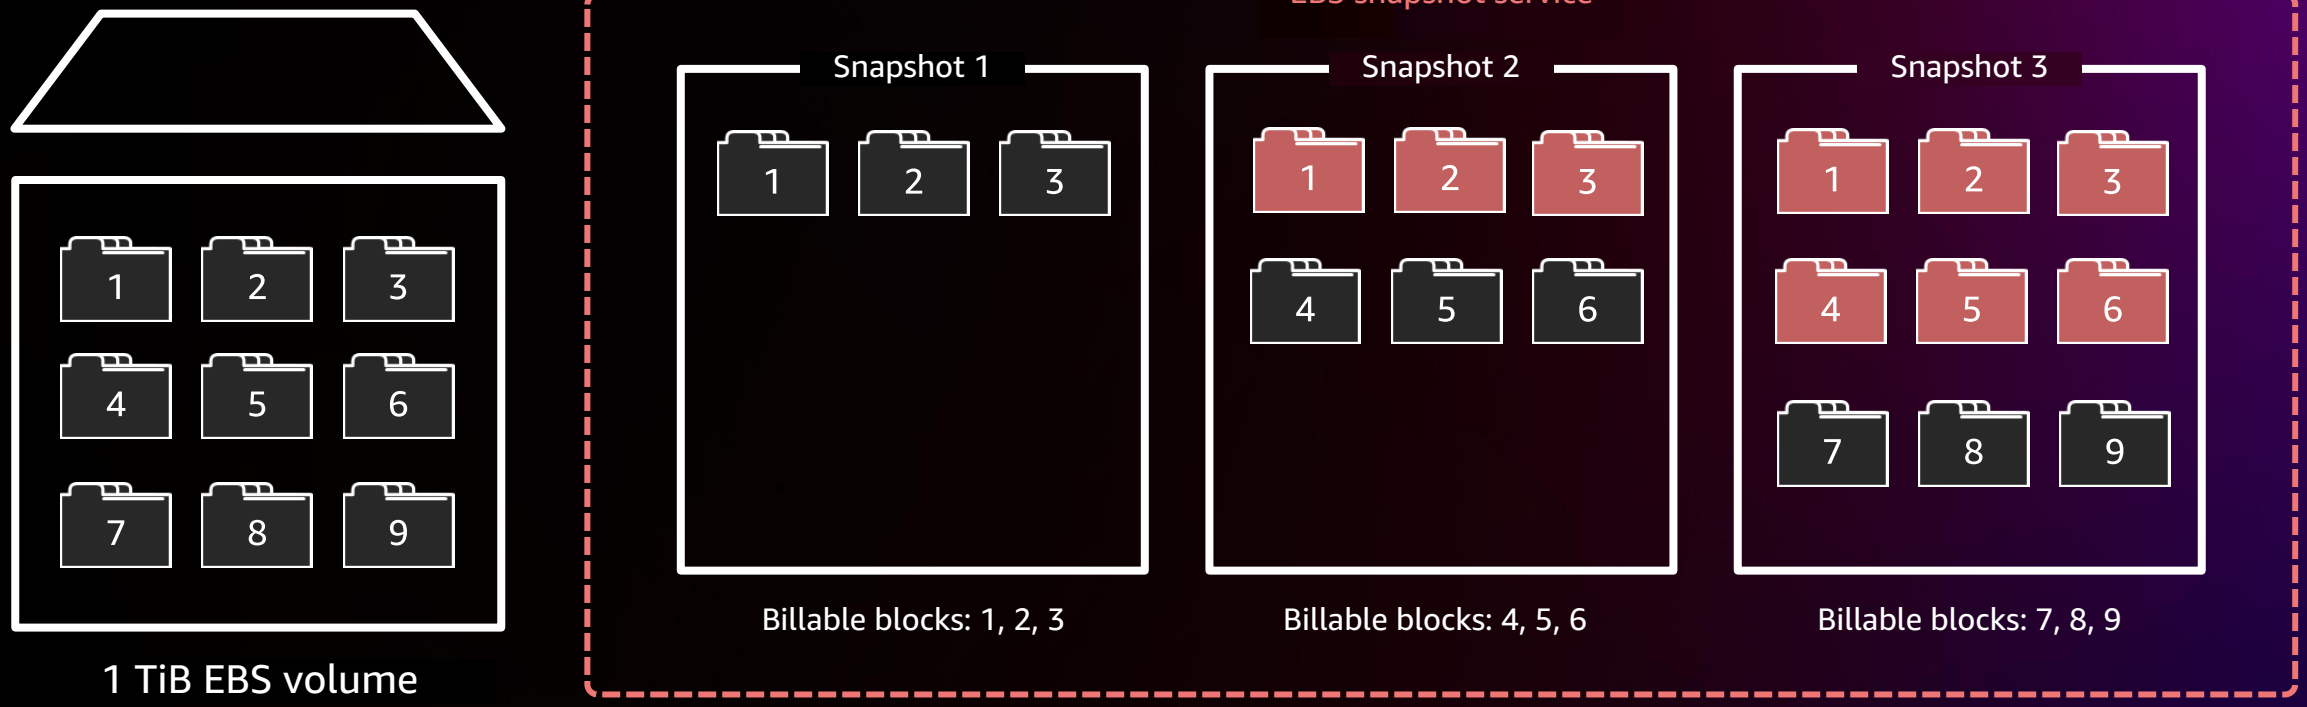

Key things to remember

EBS snapshots:
Full point-in-time backup

EBS snapshots stored incrementally

EBS snapshots charged for written blocks only

Consistency of snapshots

EBS snapshots are crash consistent

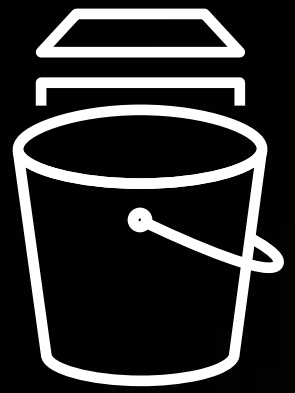

Crash consistency

- Snapshots will contain the blocks of completed I/O operations
- Data not flushed to disk does not exist in the snapshot
- Similar to pulling the power cord of a server

Application consistency

- Application data is flushed to disk prior to snapshot creation
- New writes to application(s) are halted during the snapshot creation process
- Unfreeze/unlock as soon as snapshot creation command is successfully completed

EBS multi-volume, crash-consistent snapshots

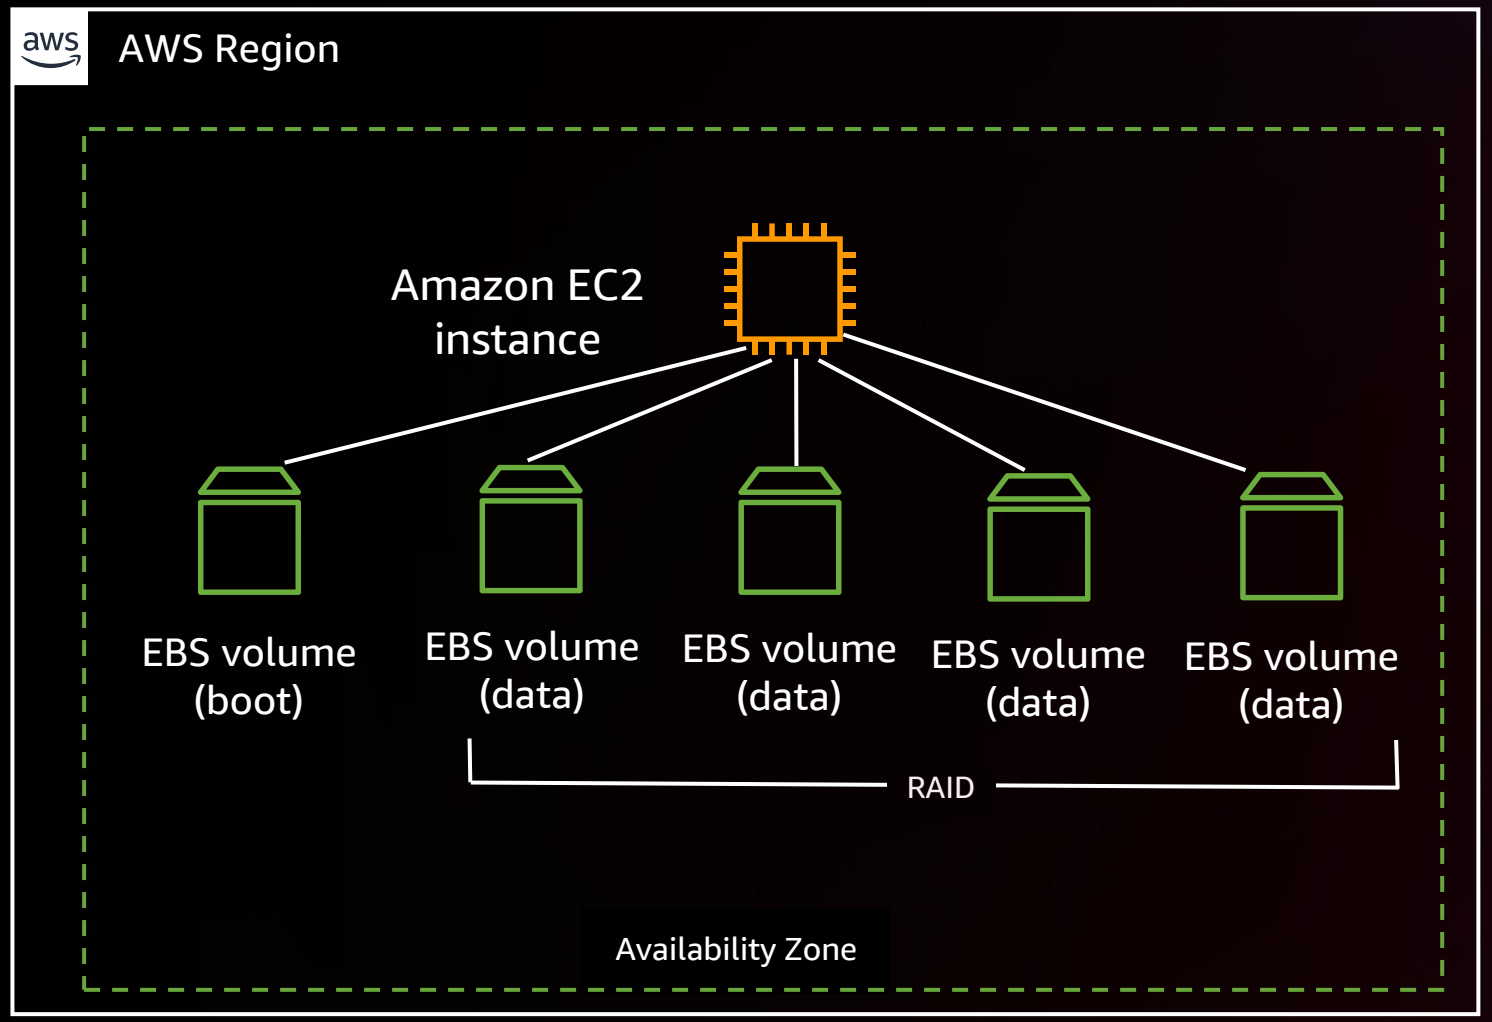

Using the CreateSnapshots API, you can take point-in-time, crash-consistent snapshots across multiple EBS volumes attached to an EC2 instance

Best practice

- Separate boot and data volumes
- Snapshot regularly

https://docs.aws.amazon.com/cli/latest/reference/ec2/create-snapshots.html

New! Snapshot a subset of data volumes

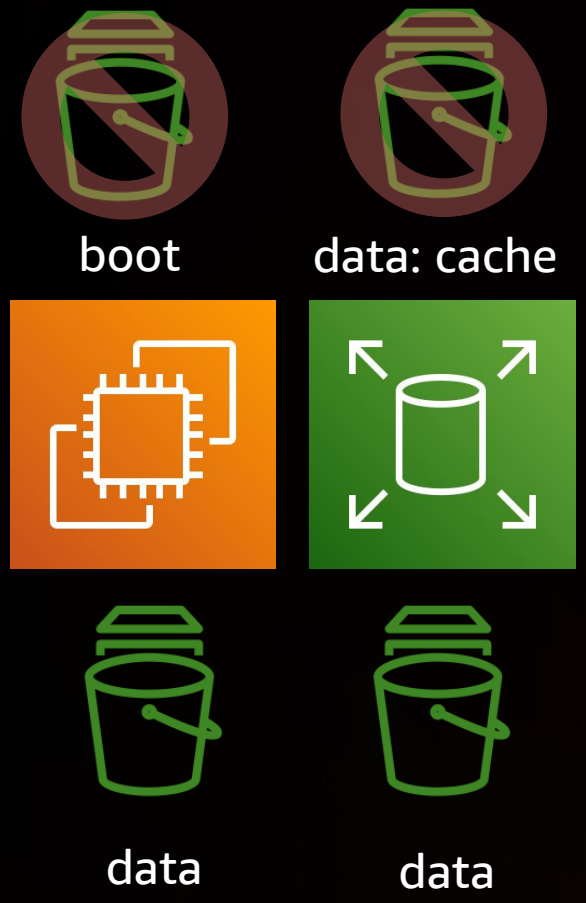

- **Simple** – replaces multiple API calls with a single CreateSnapshots call
- **Cost-efficient** – only snapshot the multi-attached volumes you select
- **Automated** – set tags in Amazon Data Lifecyle Manager policies to specify which volumes to exclude

EBS snapshot encryption

Encryption – Amazon EBS

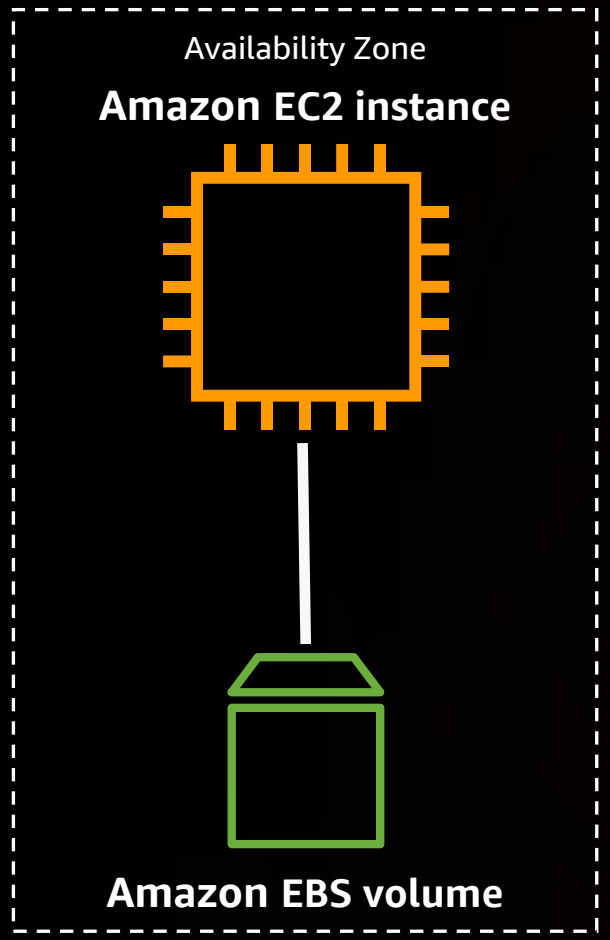

- Integrates with AWS KMS – AES-256 encryption
- Encrypted EBS volume implies the following are encrypted
  - **Data at rest** inside the volume
  - **Data moving** between the volume and instance
  - **Snapshots created** from the volume
  - **Volumes created** from such snapshots

Amazon EBS snapshot encryption

- **Snapshots of encrypted volumes** are automatically encrypted
- **Volumes created from encrypted snapshots** are automatically encrypted
- You can **encrypt an unencrypted snapshot when you copy a snapshot**
- You can **re-encrypt a snapshot you own with a different key when you copy a snapshot**

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html

Snapshots can be

- Shared across accounts
- Copied across accounts
- Copied within accounts
- Copied across AWS Regions
- Create AMIs

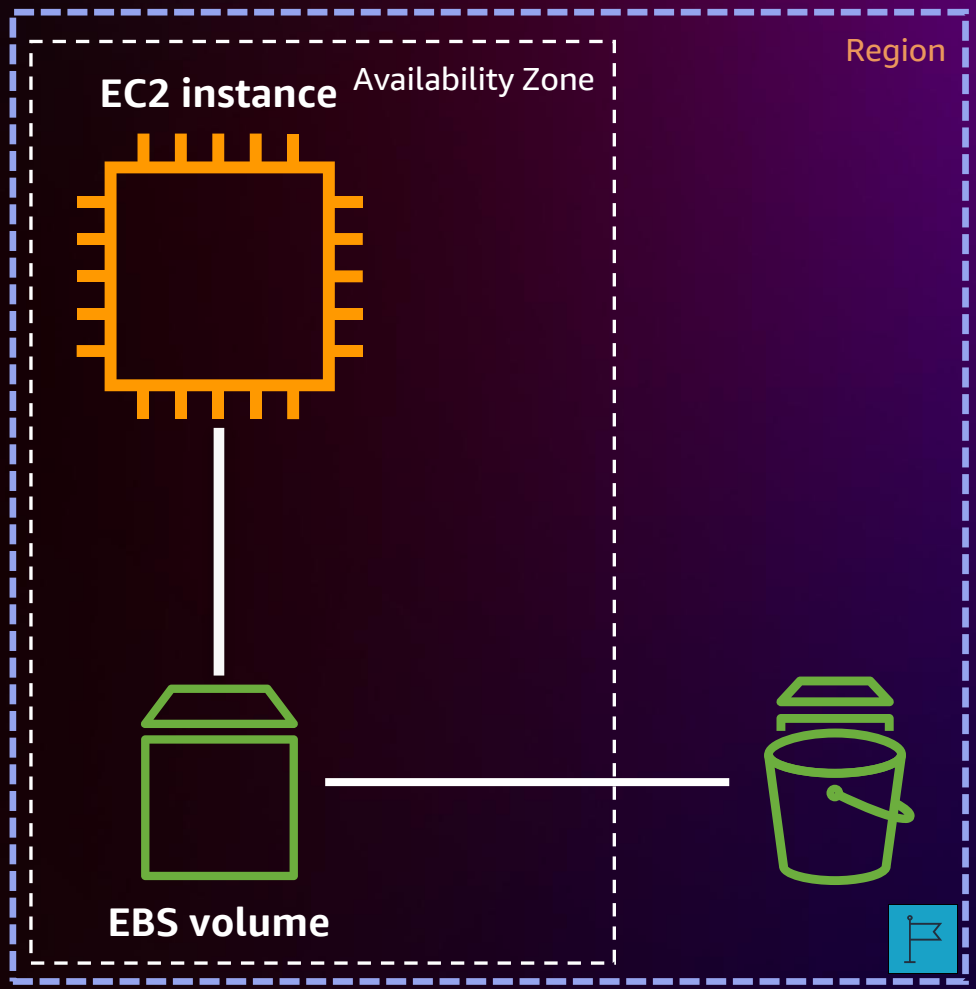

Best practice for sharing snapshots and AMIs

- Public sharing: Reasonable use case for AMIs – AWS Marketplace AMIs
- Share non-AMI snapshots with specific accounts
- **Within the same Region, you can share snapshots and AMIs without being billed for multiple full snapshots**
- To launch a volumes from a snapshot, a snapshot copy must be in the Region

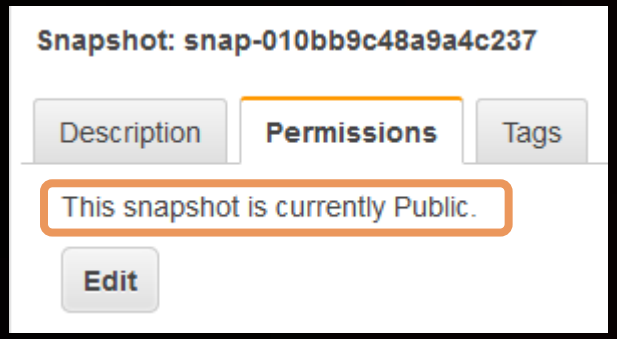


```
snap-010bb9c48a9a4c237 --attribute createVolumePermission
{
    "SnapshotId": "snap-010bb9c48a9a4c237",
    "CreateVolumePermissions": [
        {
            "Group": "all"
        }
    ]
}
```

Copy snapshots across Regions

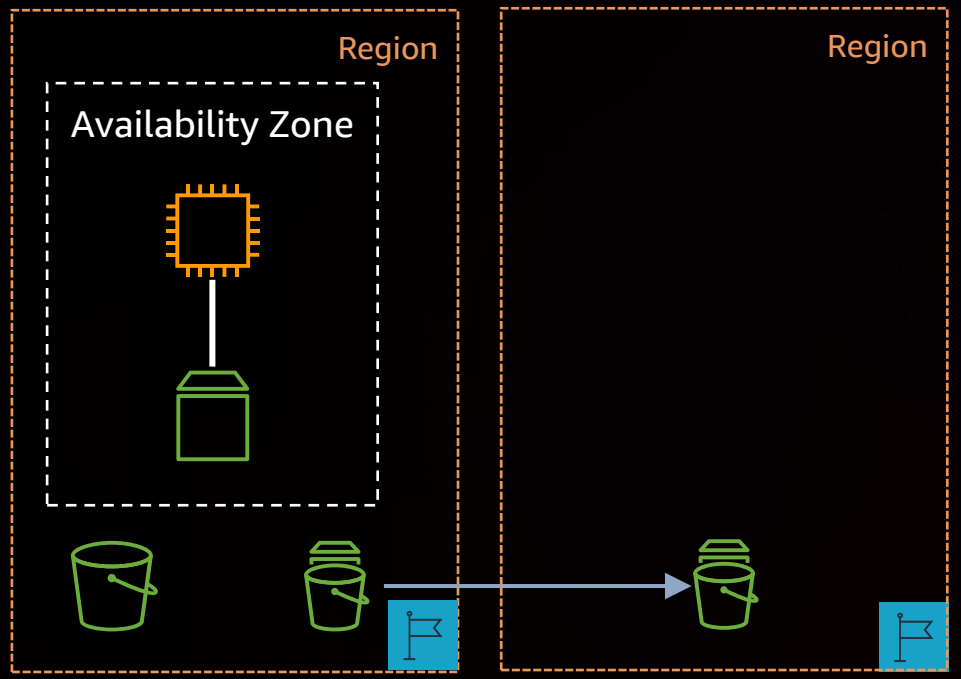

- Copy snapshots across accounts, across Regions
- Lock down resource-level permissions on target snapshot copy
- Multi-Region provides protection against regional events
- Permission lockdown protects against malicious or unintentional deletes of data

How do I enforce encryption?

- EBS encryption by default (EBD) feature
- Enabled on a per account per Region level

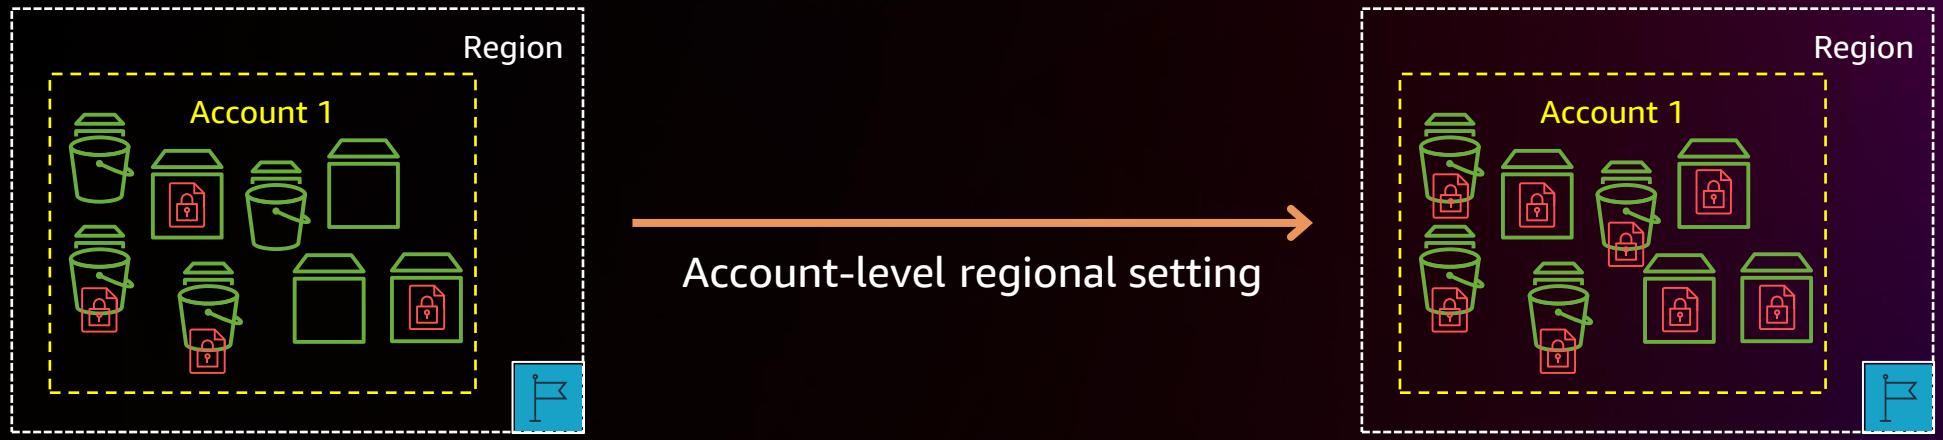

Without changing workflows, newly launched volumes + snapshots are encrypted

aws ec2 enable-ebs-encryption-by-default

Amazon Data Lifecycle Manager (DLM)

Challenges with self-managed solutions

Forgotten snapshots incur additional costs

Scripting bugs lead to under- or over-retention

Multi-schedule snapshots management is challenging

Amazon Data Lifecycle Manager

SIMPLE AUTOMATED WAY TO BACK UP DATA STORED ON AMAZON EBS VOLUMES

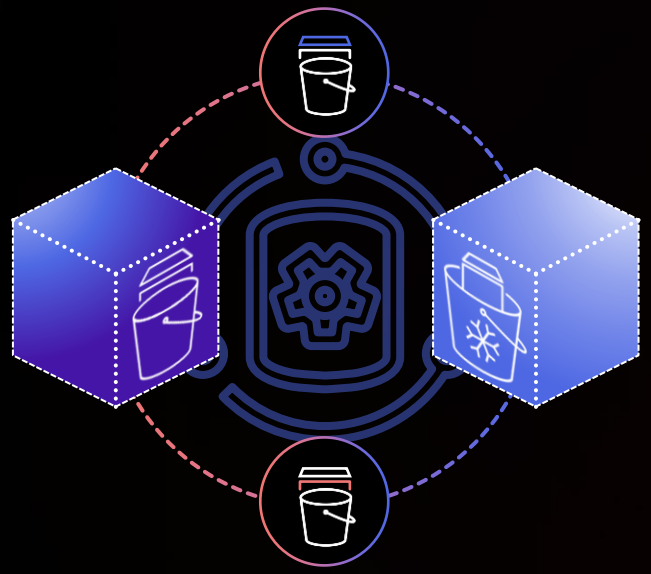

Policy- and tag-based backup solution

Automated backup scheduling

- Automated backup retention management
- Automated archiving of snapshots

Use AWS Identity and Access Management (IAM) to control policy access

- No cost to use

Amazon Data Lifecycle Manager

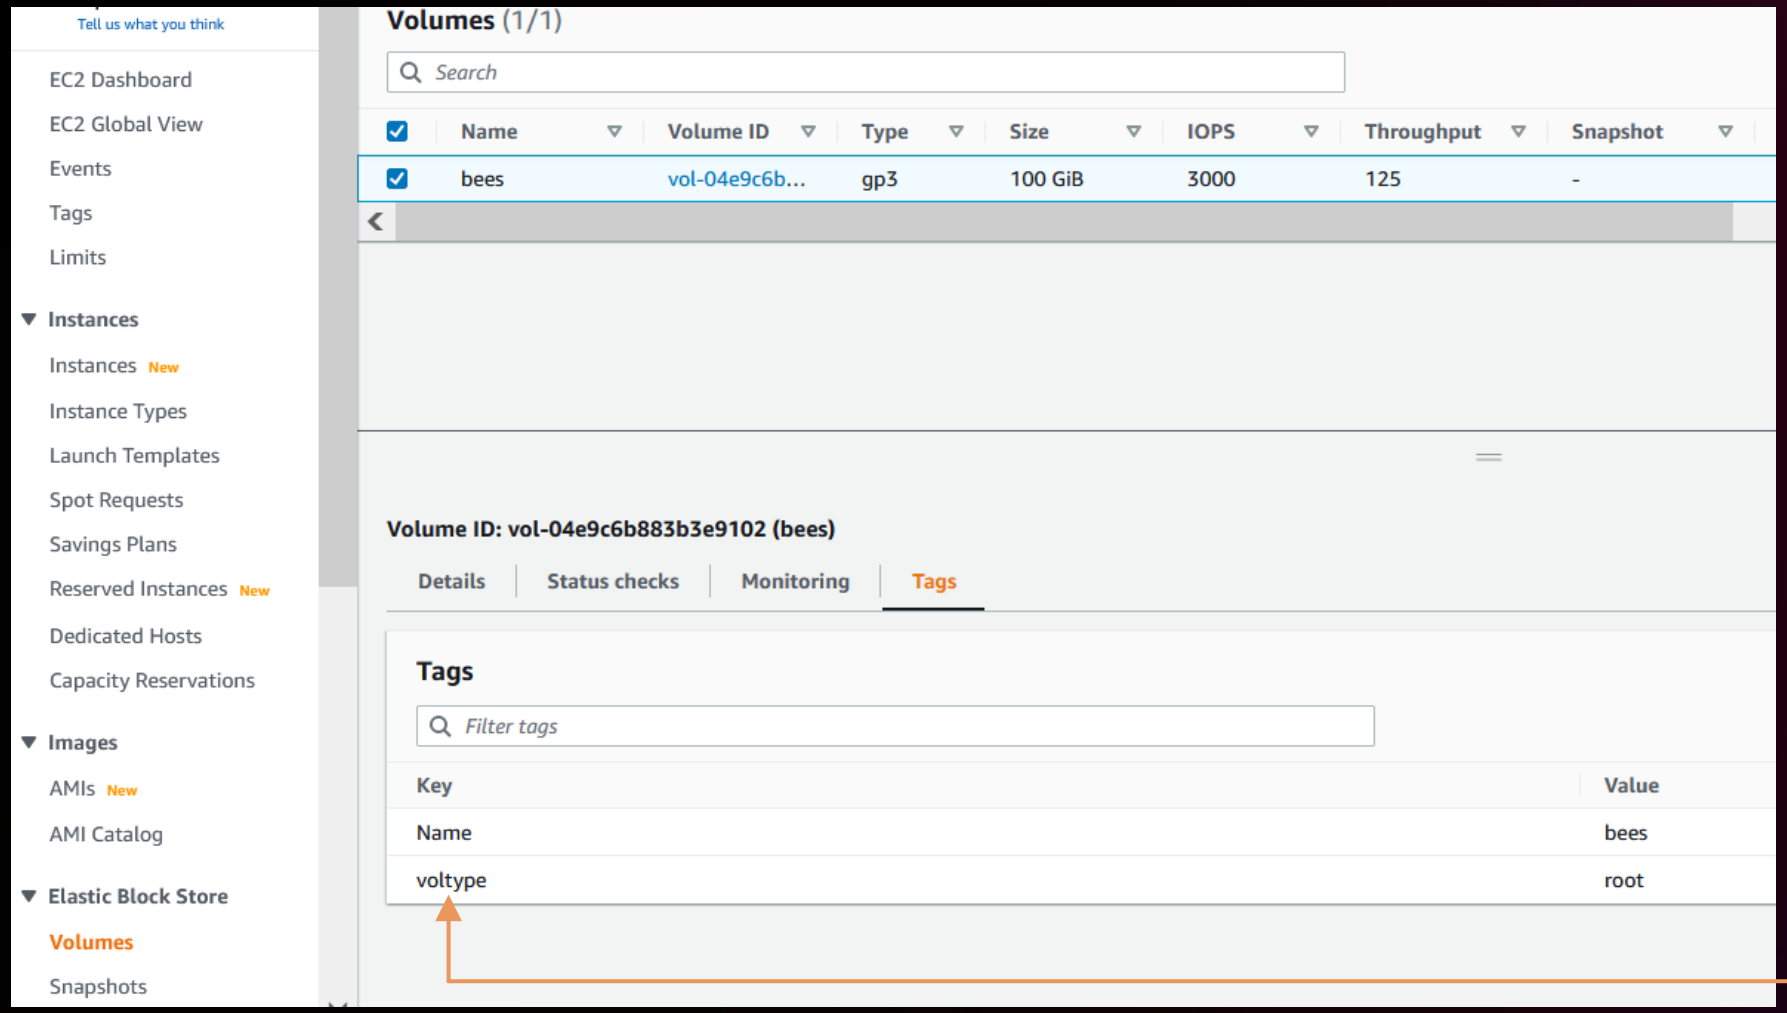

Add tags to the volumes you want to associate with Amazon Data Lifecycle Manager policies

Amazon Data Lifecycle Manager scheduling features

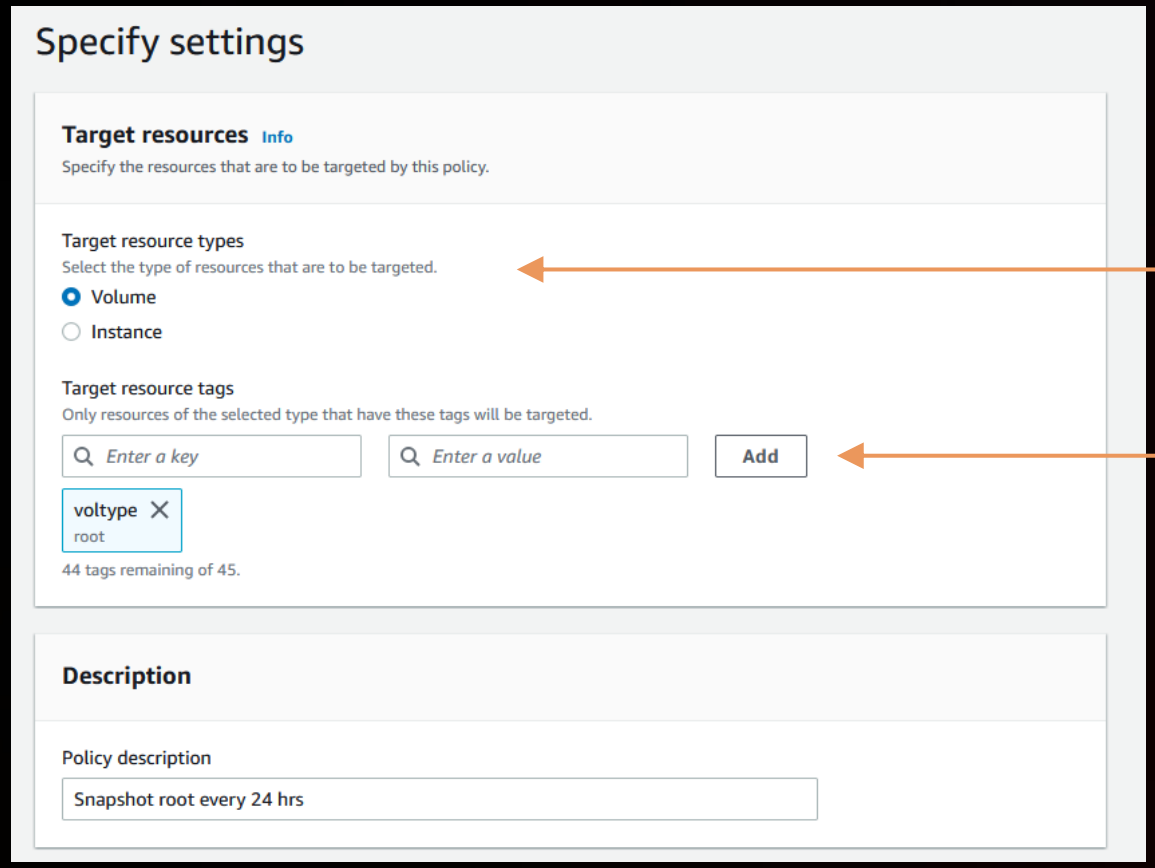

Select whether to target volumes or instances with the policy

Input the appropriate volume tags to target

Amazon Data Lifecycle Manager scheduling features

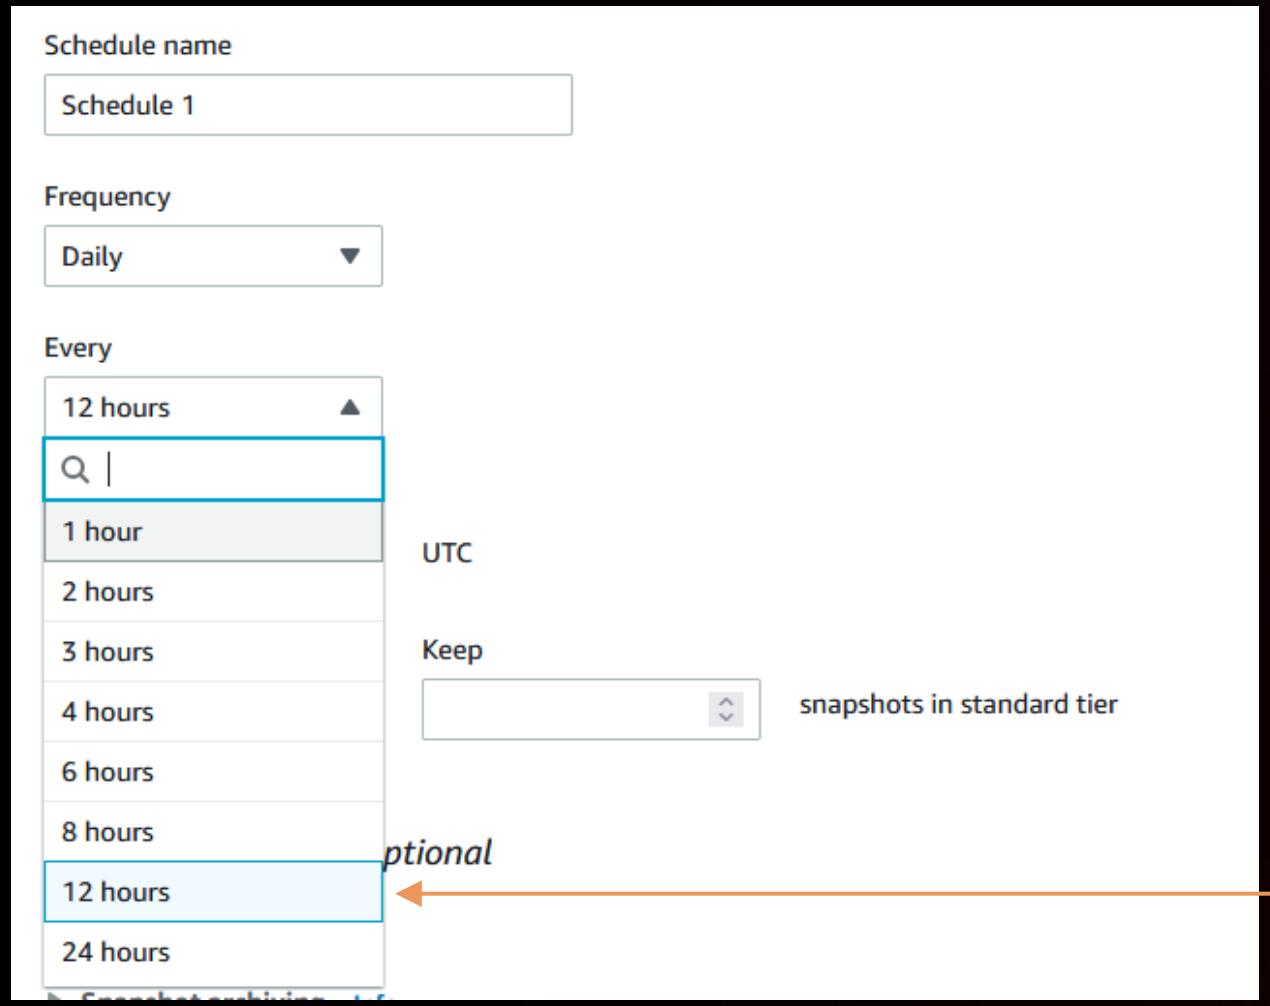

Reduce RPO on your Amazon Data Lifecycle Manager policy up to **1 hour**

Schedule snapshots from every hour to every 24 hours

Amazon Data Lifecycle Manager scheduling features

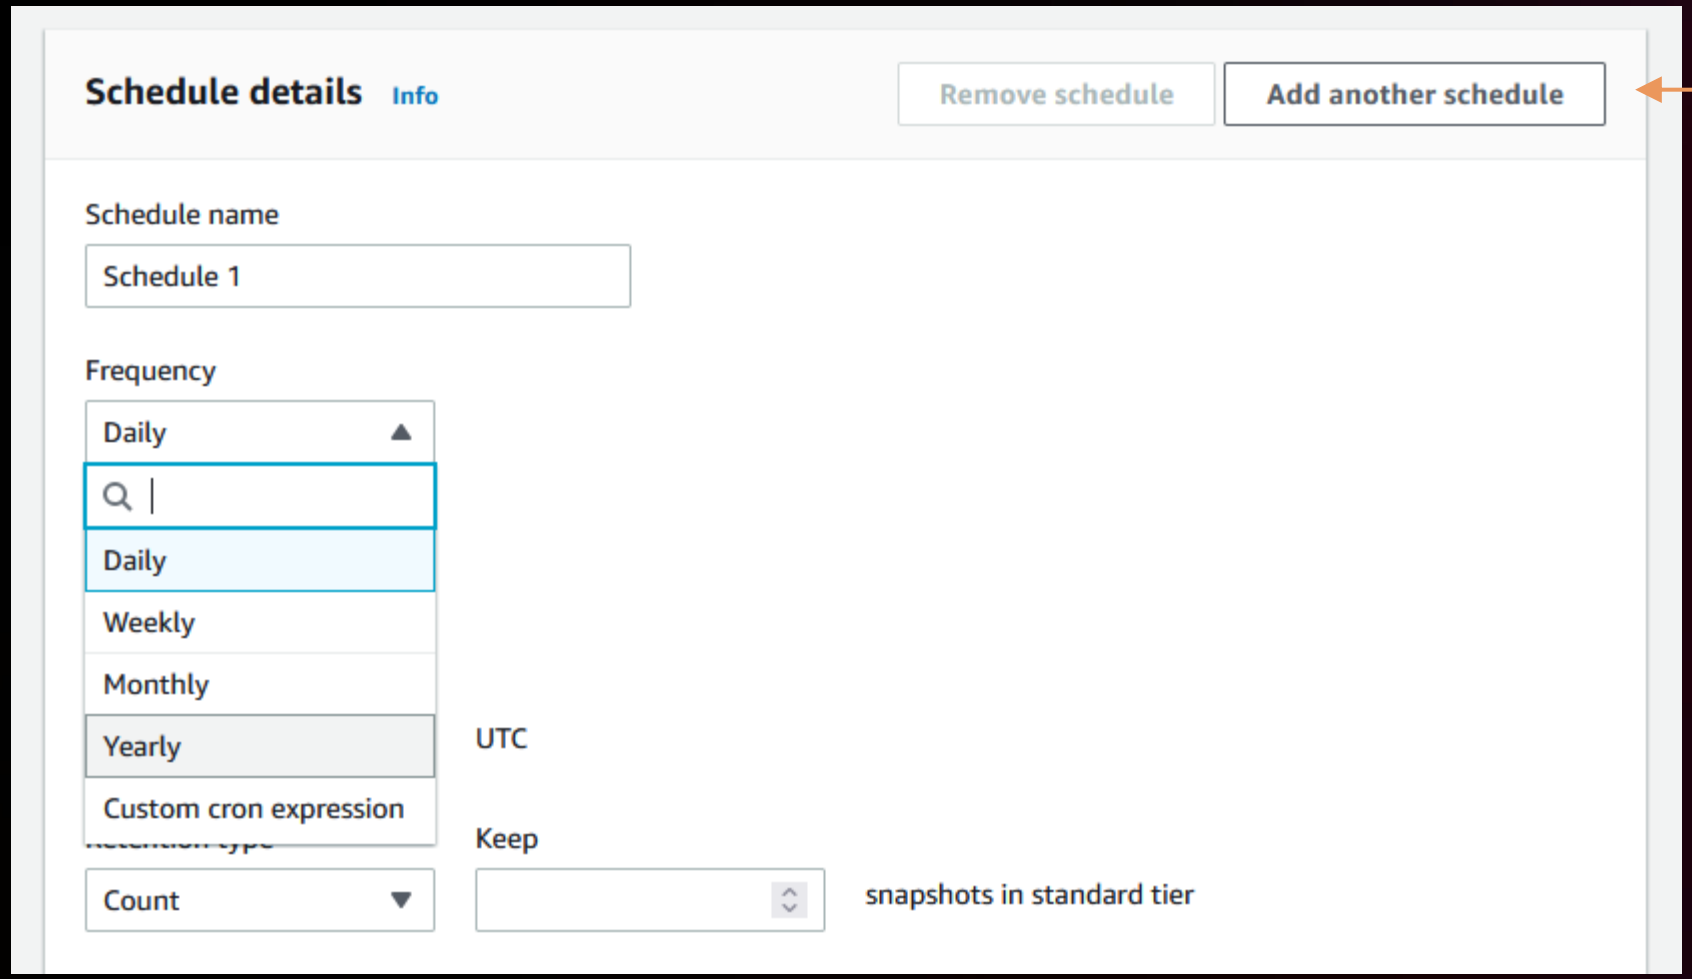

Define multiple schedules in a single policy

Amazon Data Lifecycle Manager retention features

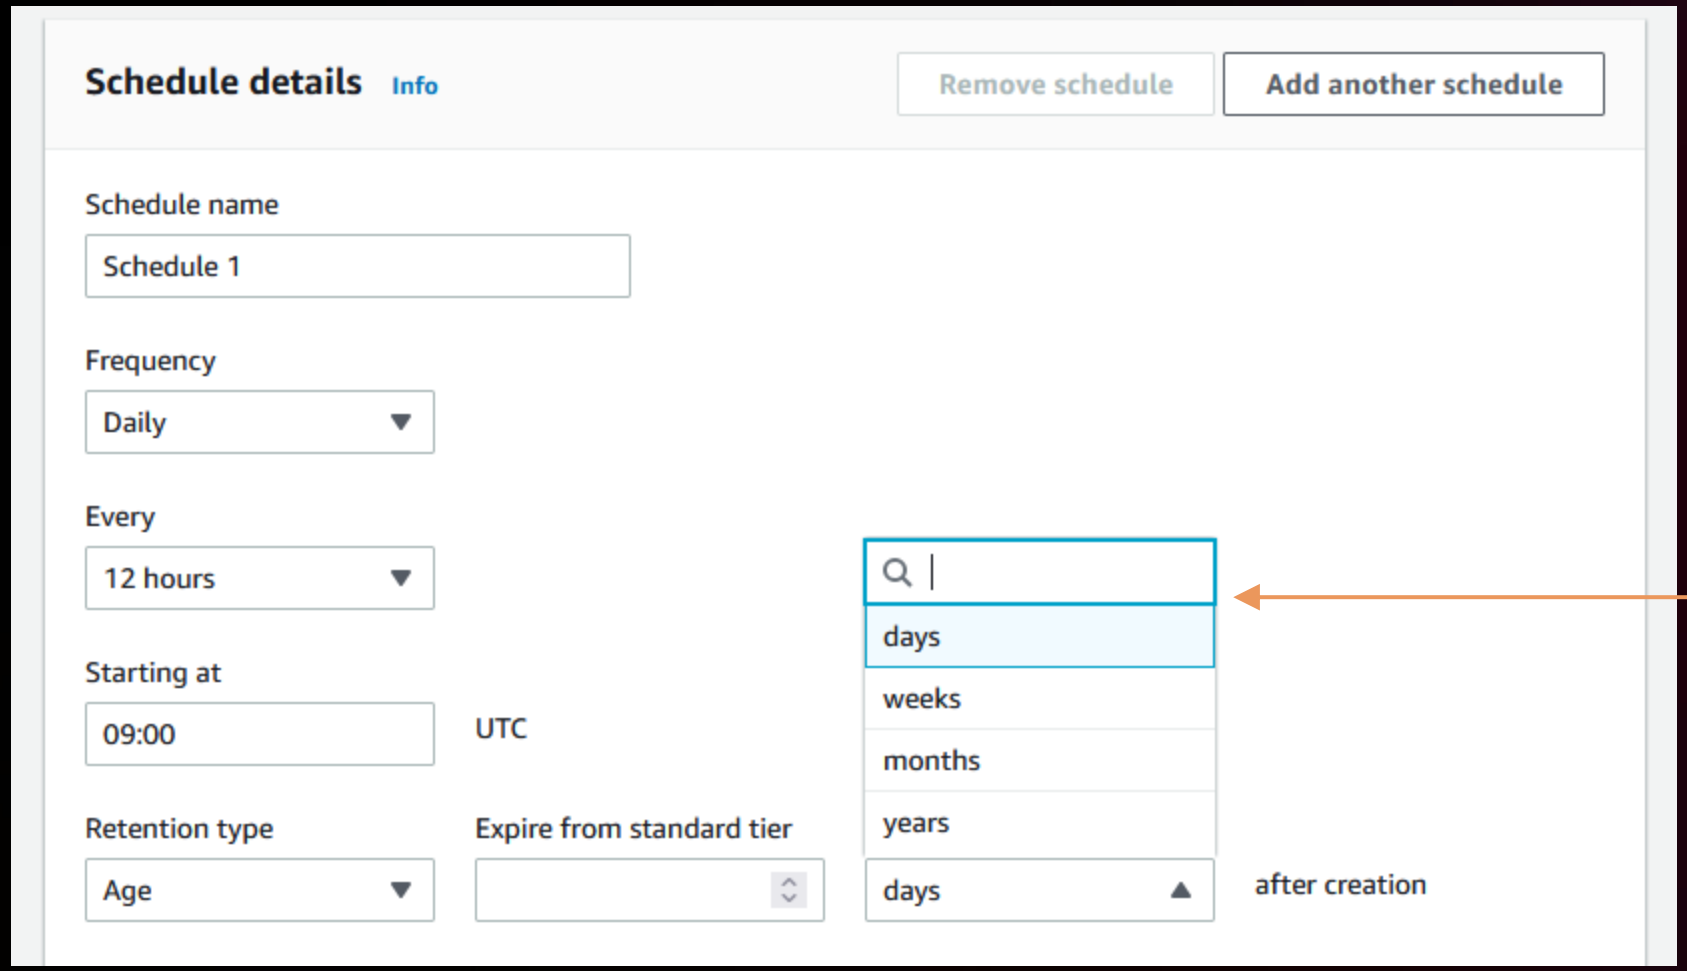

Define retention based on days, weeks, months, or years

Amazon Data Lifecycle Manager cross-Region copy features

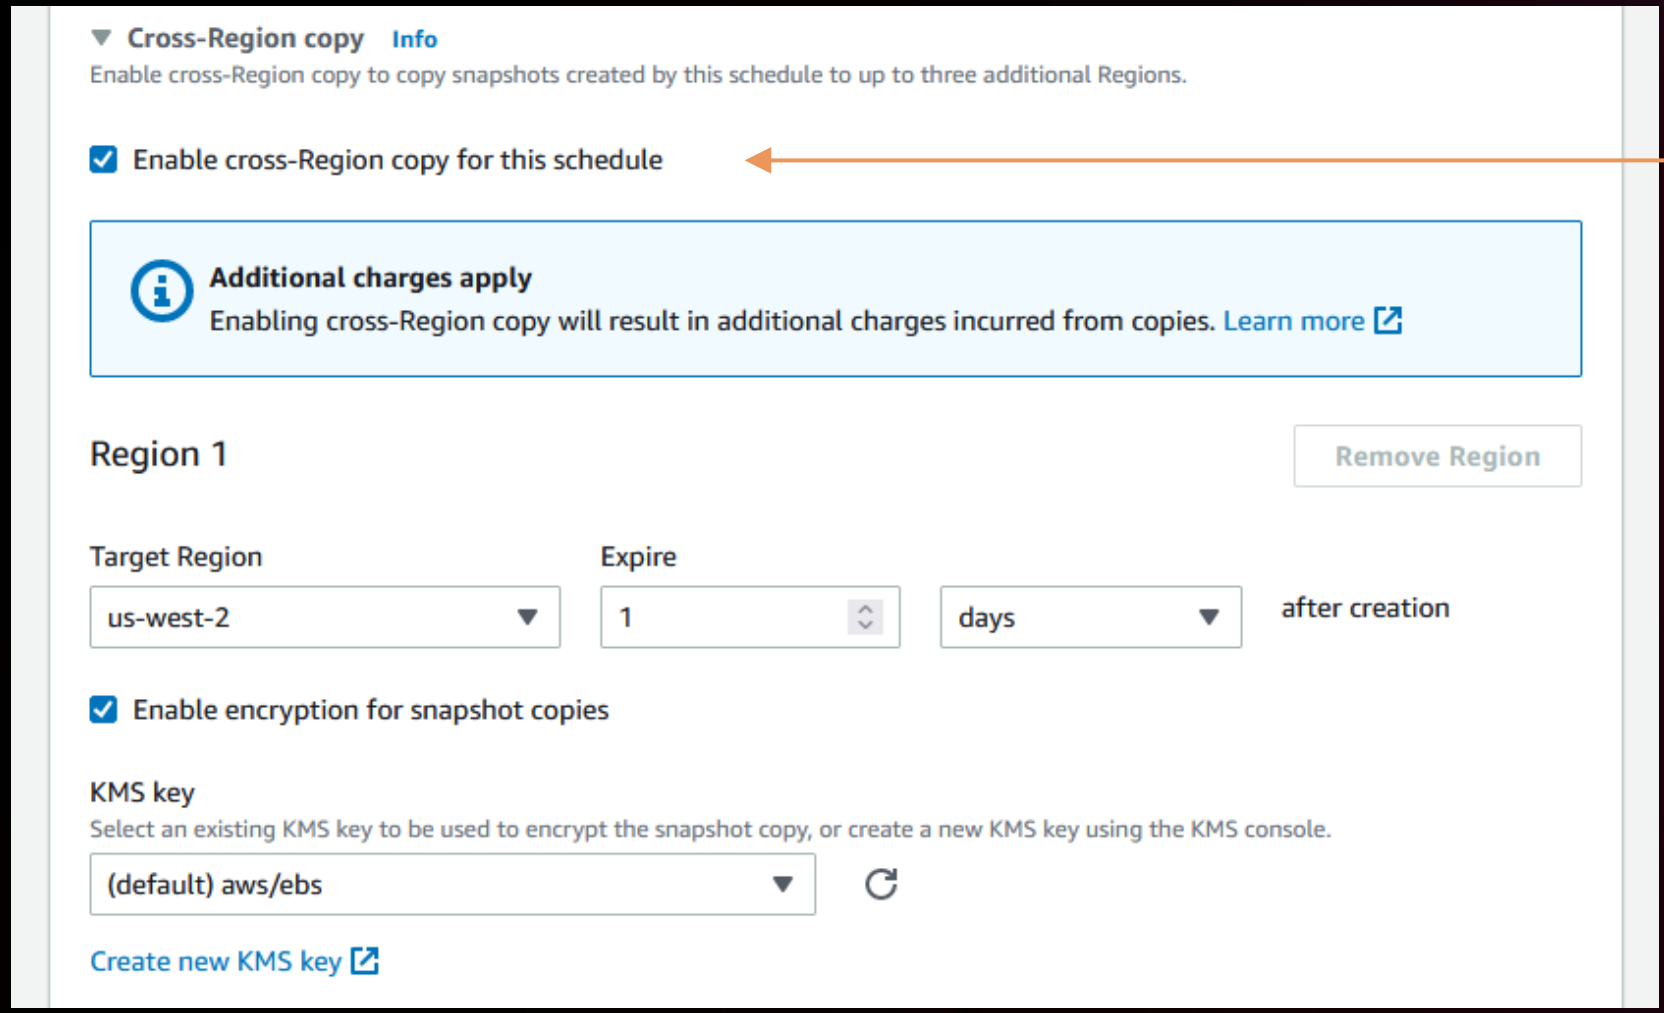

Automate copy across Region

Amazon Data Lifecycle Manager snapshot archive feature

Automatically create and **archive monthly, quarterly, and yearly** snapshots

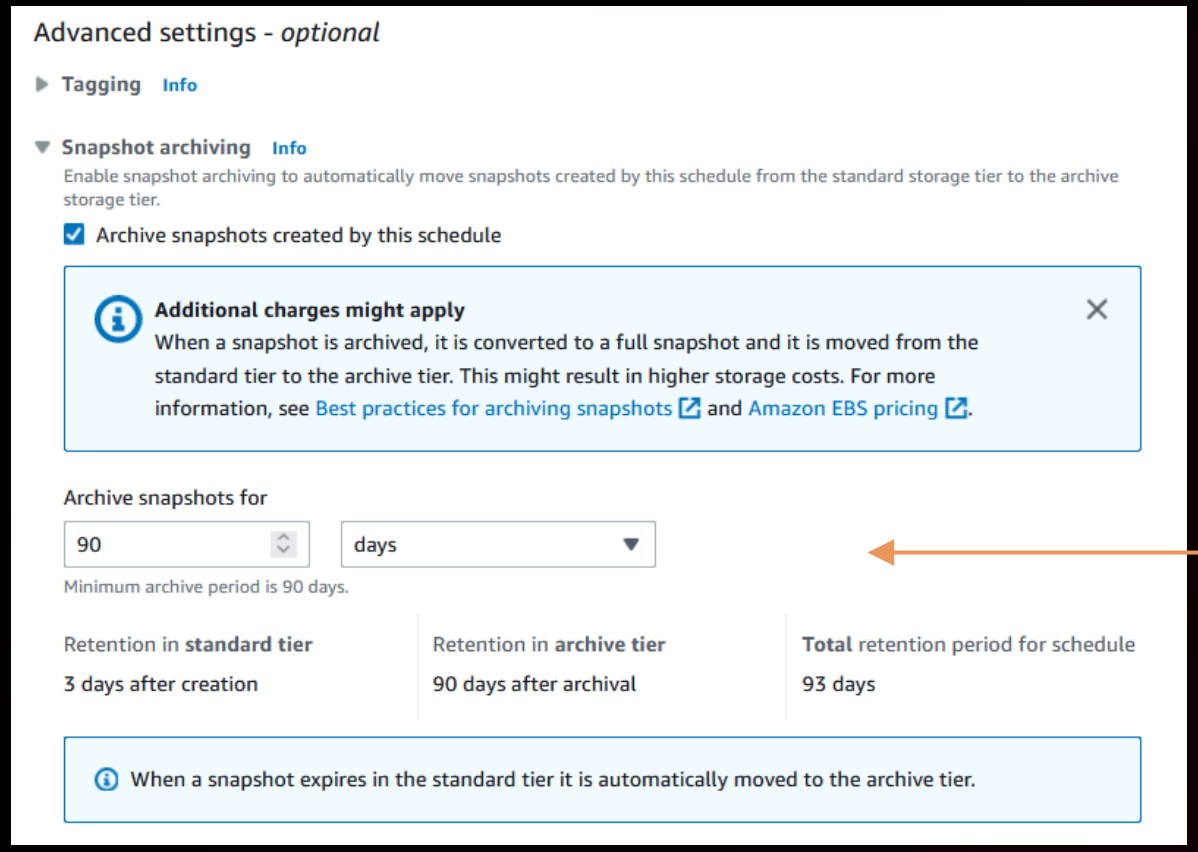

Define retention in archive tier

Recycle Bin for EBS snapshots and AMIs

Causes of accidental snapshot deletion

Simple human errors

Scripting bugs

Accidental retention policy change

Recycle Bin for EBS snapshots and AMIs

ALLOWS FOR QUICK RECOVERY OF DELETED SNAPSHOTS AND AMIS

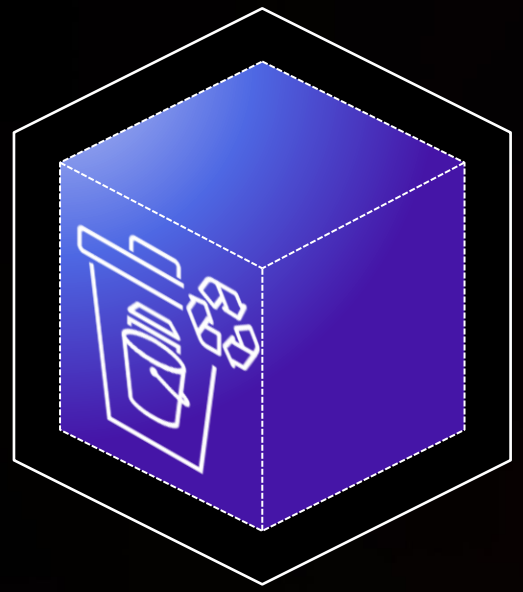

Automatically retain deleted snapshots for a retention period you specify

Use retention rules to specify retention periods for all or some tagged snapshots in your account

Pay EBS snapshot price for snapshots in Recycle Bin

Restrict access to Recycle Bin rules and resources using administrator IAM role

Recycle Bin: How it works

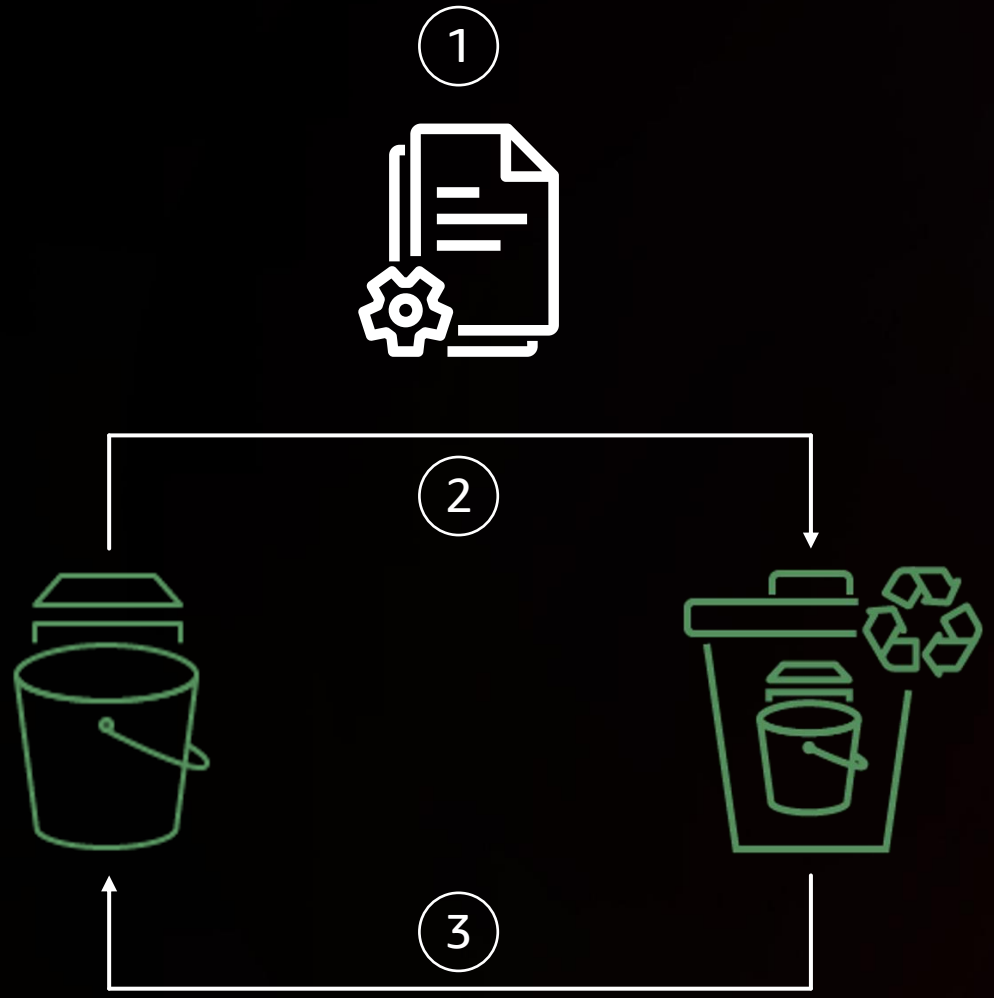

1. Set up retention rules
    a. Account-level rules per Region
    b. Tag-based rules per Region

2. Retain deleted snapshots and AMIs in Recycle Bin

3. Recover deleted snapshots and AMIs before expiry of retention period

AVAILABLE NOW

Introducing Rule Lock for Recycle Bin

PROTECT YOUR RESOURCES FROM ACCIDENTAL OR MALICIOUS DELETION

Lock retention rules to prevent them from being modified or deleted

Protect your EBS snapshots and EC2 AMIs from malicious deletions and account breaches

Maliciously deleted resources are safely retained in the Recycle Bin and can be recovered later

aws

How it works

Create a new locked rule or lock an existing account-level rule

A locked rule can only be unlocked by users with unlock permissions

Unlocked rules stay locked for a configurable period of 7 to 30 days, giving a layer of protection against malicious attempts

EBS Snapshots Archive

EBS Snapshots Archive tier

- **Long-term snapshot retention**
  - 90-day-minimum retention period
- **Full, point-in-time backups**
- **Retrieve snapshot** before use
  - Retrieval time of hours
- **75% lower storage costs at $0.125/GB-mo.**
  - $0.03/GB additional retrieval charges*

EBS Snapshots Archive tier provides low storage cost for long-term retention of rarely accessed EBS snapshots

*For complete pricing information, visit https://aws.amazon.com/ebs/pricing

EBS Snapshots Archive: How it works

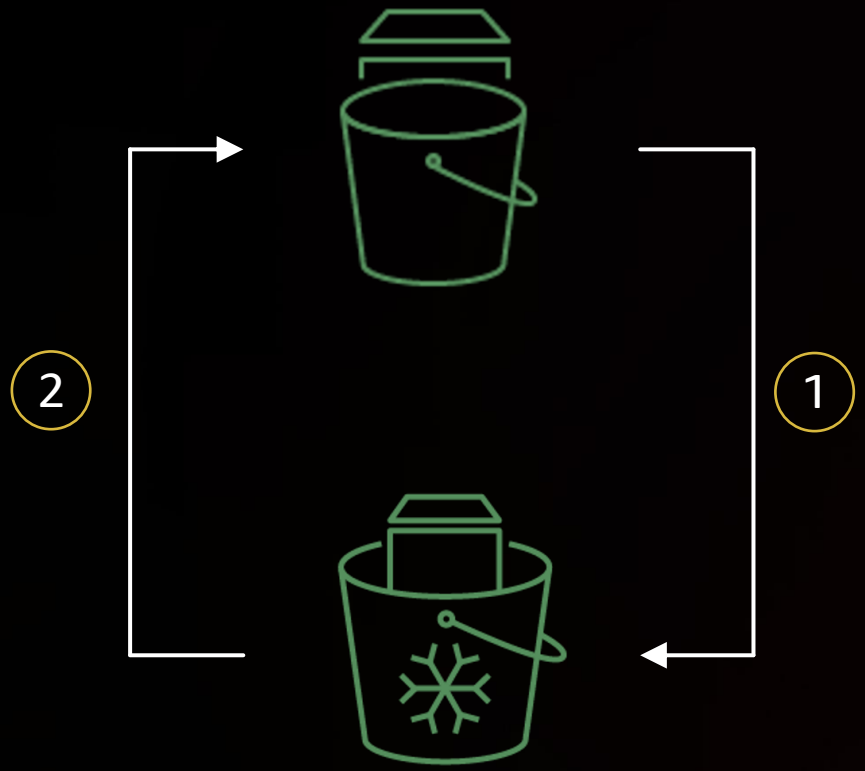

Archive snapshot

1

- Creates full version of snapshot and moves it into archive tier

Restore snapshot

2

- Restores a full version of your archive snapshot to the EBS Snapshot Standard tier
- Can be temporary or permanent

Simple APIs for archiving and restoring snapshots

Restoring snapshots from archive

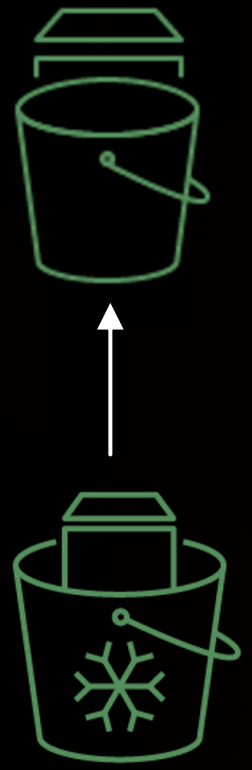

- **Retrieve snapshot from archive for regular use**
- **Retrieval time of hours**
  - Based on snapshot size (24–72 hours maximum)
- **Temporary restore**
  - Restore temporarily for up to 180 days to standard tier
  - Permanent copy stays in archive tier
- **Permanent restore**
  - Change snapshot tier to standard from archive
  - No copy in archive tier
  - Remember 90-day minimum retention in archive

Considerations for archiving snapshots

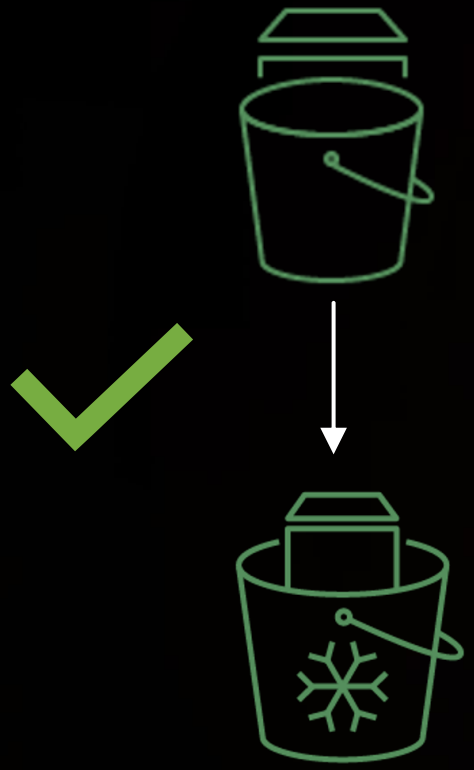

Which snapshots should you consider archiving?

- Last snapshots of inactive lineages
- Full snapshots for compliance
- Snapshots with >25% change rates

Considerations for archiving snapshots

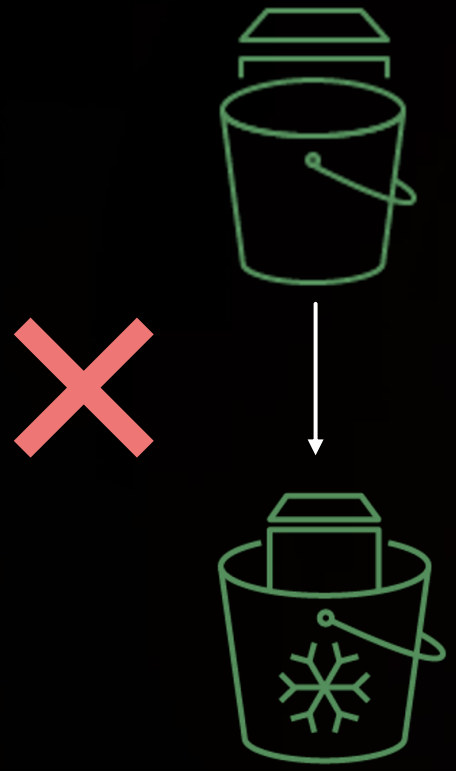

Which snapshots should you **not** consider archiving?

- Snapshots <25% change rates
- First snapshots of active lineages
- Last snapshots of active lineages

EBS direct APIs

Amazon EBS direct APIs: Read and write access to snapshots

With a set of APIs, you can

Write

Create snapshots from any source

Read

- **List blocks** in a snapshot
- **Read blocks** in parallel
- **Compare** snapshots and **read changed blocks**

Thank you!

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