Amazon DocumentDB
(with MongoDB compatibility)

Overview

Jason Plank

Sr. GTM Specialist
jmplank@amazon.com

Douglas Bonser

Sr. Specialist SA
dbonser@amazon.com

Agenda

- Introduction to DocumentDB
  - Why DocumentDB?
  - DocumentDB architecture
  - Service overview
  - Recent feature releases
  - Pricing
  - Q&A

Comprehensive set of services across Databases & Analytics

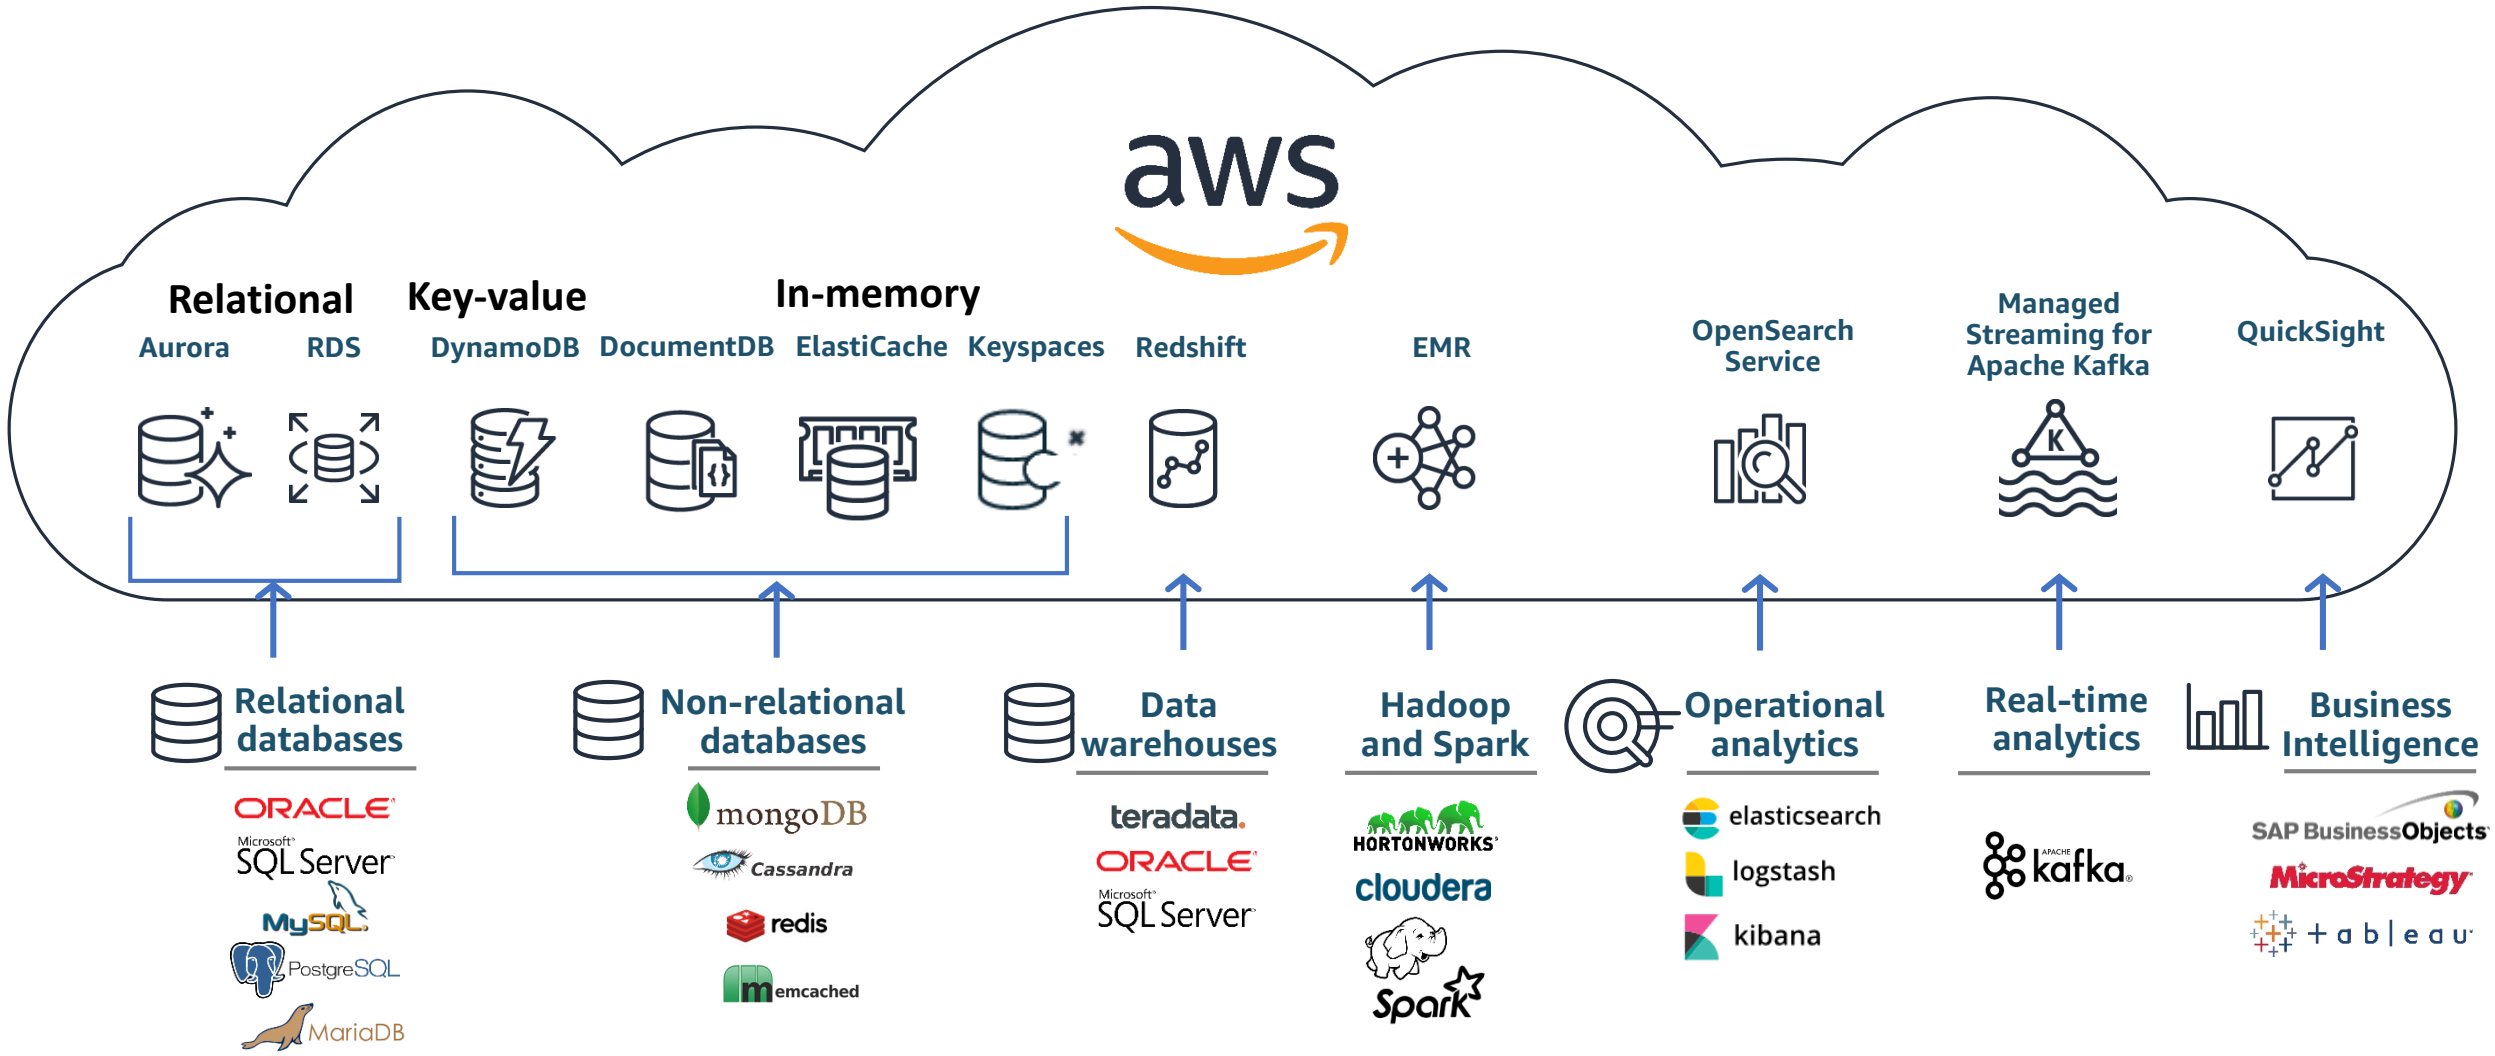

Amazon DocumentDB

A scalable, highly durable and fully managed enterprise database service for operating JSON workloads with ease

Fast and scalable

Supports millions of document read requests per second by scaling compute and storage independently.

Fully managed

Automates hardware provisioning, patching, setup, backups, and other database management tasks.

Enterprise ready

Maintains high availability and durability, built-in security best practices, and low-latency global reads.

MongoDB compatible

Supports hundreds of APIs, operators, and stages; Applications, drivers, and tools can be used with little or no change.

Amazon DocumentDB customers across industries

Media: Content Management

BBC
The Asahi Shimbun
hudl
DFL
The Washington Post
sportradar

eCommerce: Catalogs

freshop
amazon
Rappi

Technology: IoT

Plume

eCommerce: Search Recommendation

Zulily

Marketing Agency: Analytics

N digital

Profile Management

HABBY
DMM GAMES

SaaS: Personalization

punchh

Technology: Various

freshworks
ENKIDOO ai
sandstone TECHNOLOGY
Homage
SUMADI
smartsheet

Technology: Financial Services

Capital One

https://aws.amazon.com/documentdb/customers/

Document Databases

- Why Document DBs
- Use Cases
- Example Structure

Why document databases

- Data is stored in JSON-like documents
- Documents map naturally to how humans model data
- Flexible schema and indexing
- Expressive query language built for documents (ad hoc queries and aggregations)

Documents are first-class objects in the database

Document: Fields: Values

```
{
  id: 1,
  name: "sue",
  age: 26,
  email: "sue@example.com",
  promotions: ["new user", "5%", "dog lover"],
  memberDate: 2018-2-22,
  shoppingCart: [
    {product:"abc", quantity:2, cost:19.99},
    {product:"edf", quantity:3, cost:2.99}
  ]
}
```

Industry Use Case –Examples

Content management

Mobile and web applications

Personalization

Catalogs

IoT

Profile management

Great for complex documents that are dynamic and may require ad hoc querying, indexing, and aggregations

Document Use Case: Gaming User Profile

```
{
  userid: 181276,
  username: "sue1942",
  name: {
    first: "Susan",
    last: "Benoit"
  }
}
```

```
{
  userid: 181276,
  username: "sue1942",
  name: {
    first: "Susan",
    last: "Benoit"
  },
  ExplodingSnails: {
    hi_score: 3185400,
    global_rank: 5139,
    bonus_levels: true
  },
} promotions: ["new user","5%","snail lover"]
}
```

Challenges scaling document databases

Traditional Database Architecture Challenges

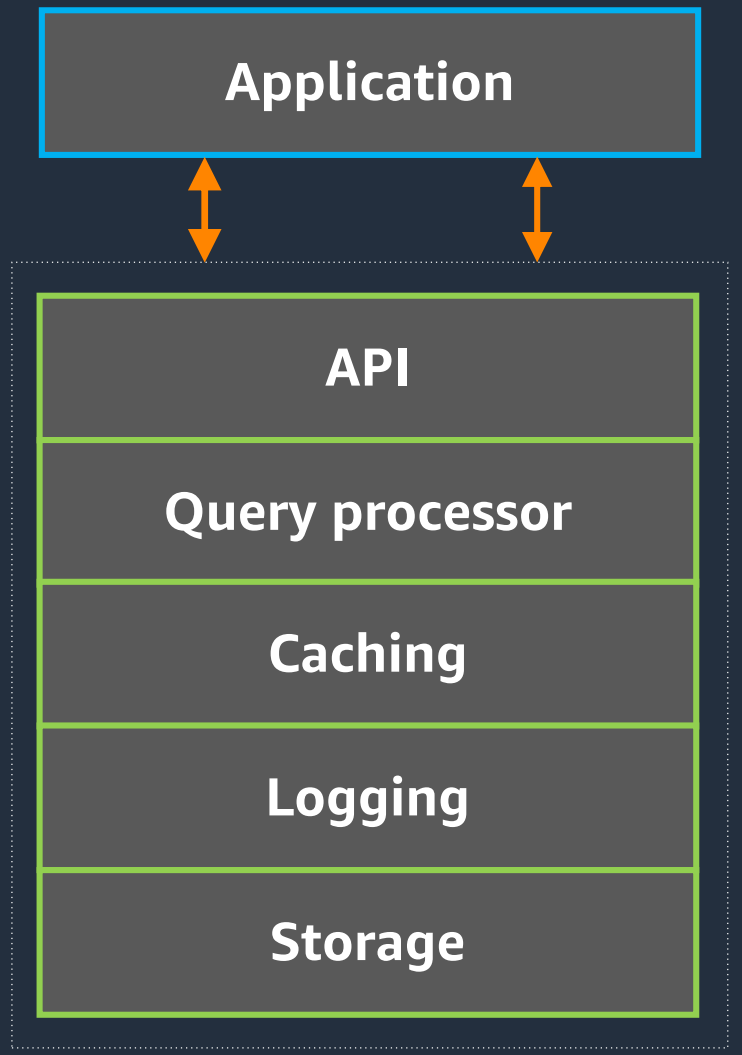

Single monolithic architectures

Scale monolithically
Fail monolithically

Not designed for the cloud

Traditional Database Architecture Challenges

Challenge #1: Add read capacity on-demand

Traditional Database Architecture Challenges

Challenge #1: Add read capacity on-demand

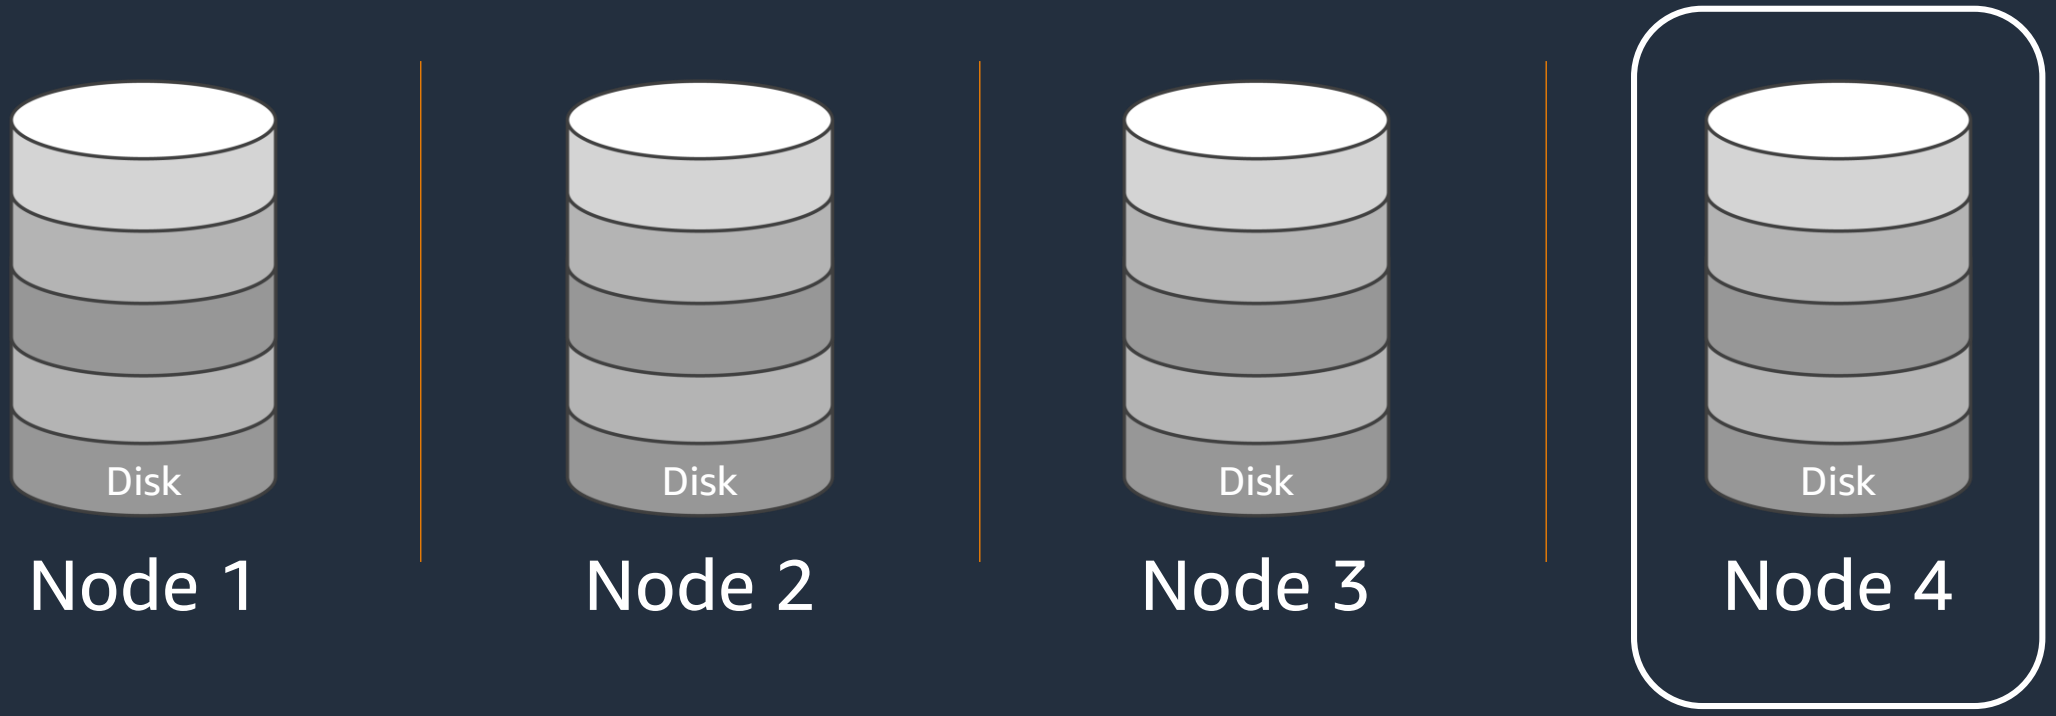

Traditional Database Architecture Challenges

Challenge #1: Add read capacity on-demand

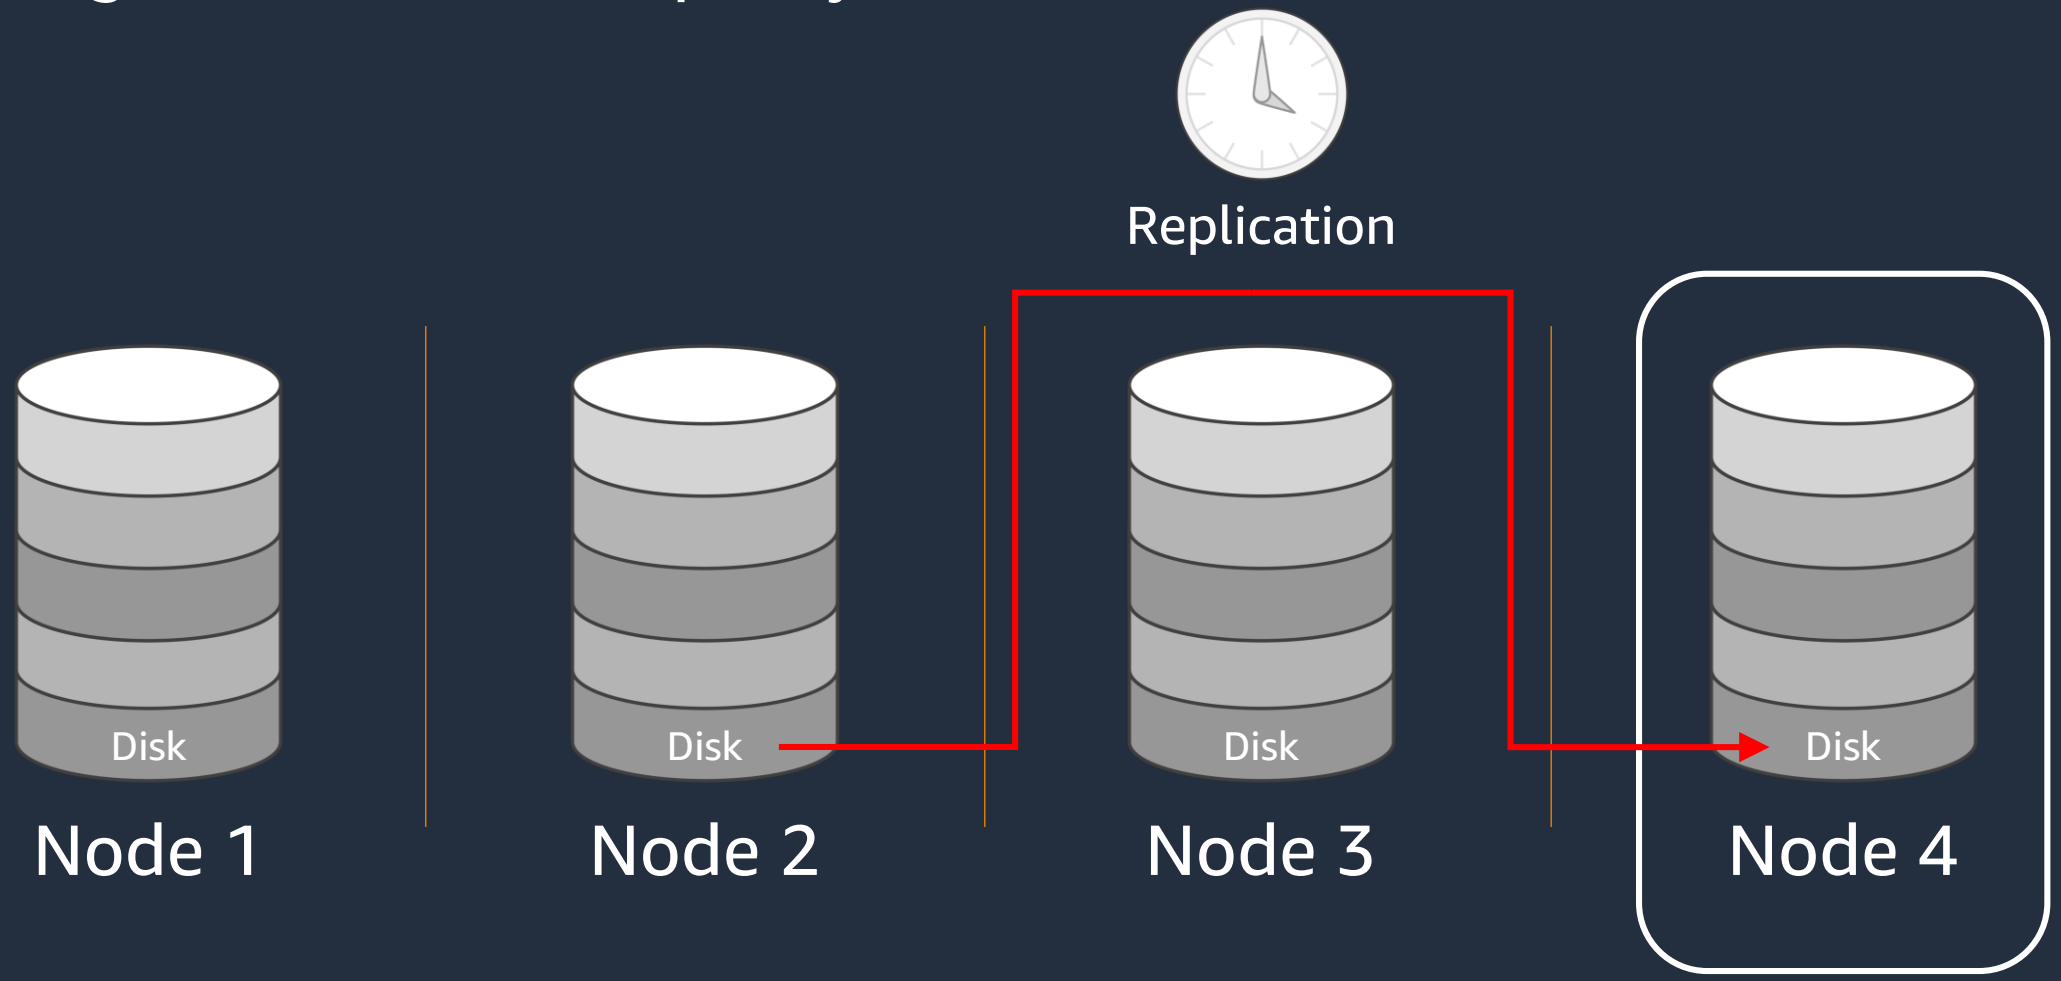

Traditional Database Architecture Challenges

Challenge #2: Recover quickly from node failure

Traditional Database Architecture Challenges

Challenge #2: Recover quickly from node failure

Traditional Database Architecture Challenges

Challenge #2: Recover quickly from node failure

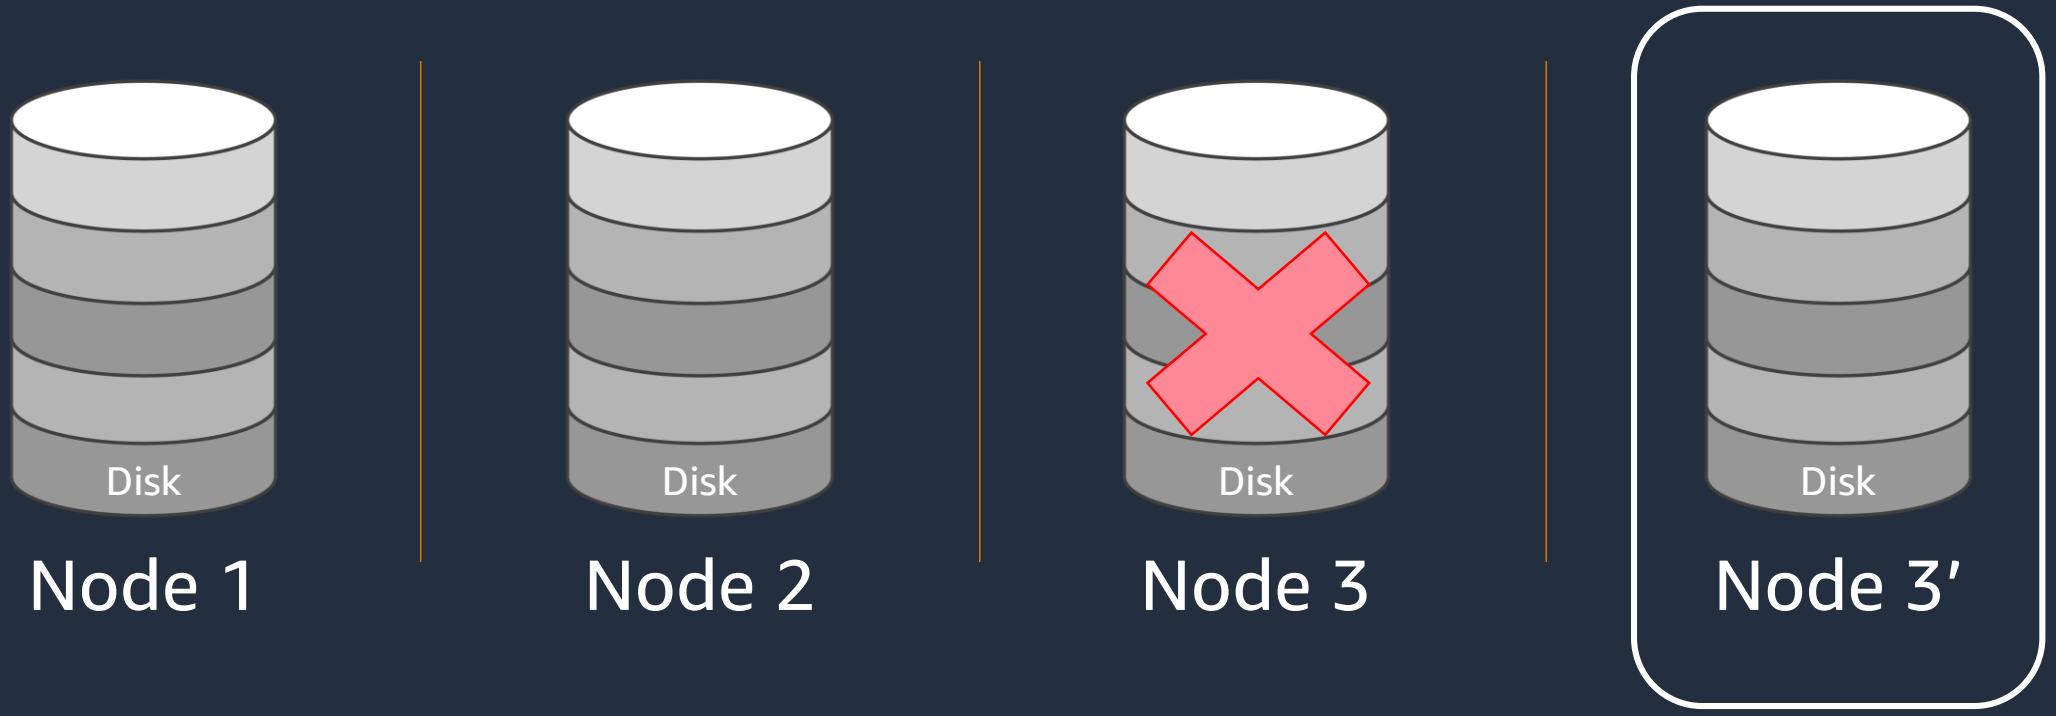

Traditional Database Architecture Challenges

Challenge #2: Recover quickly from node failure

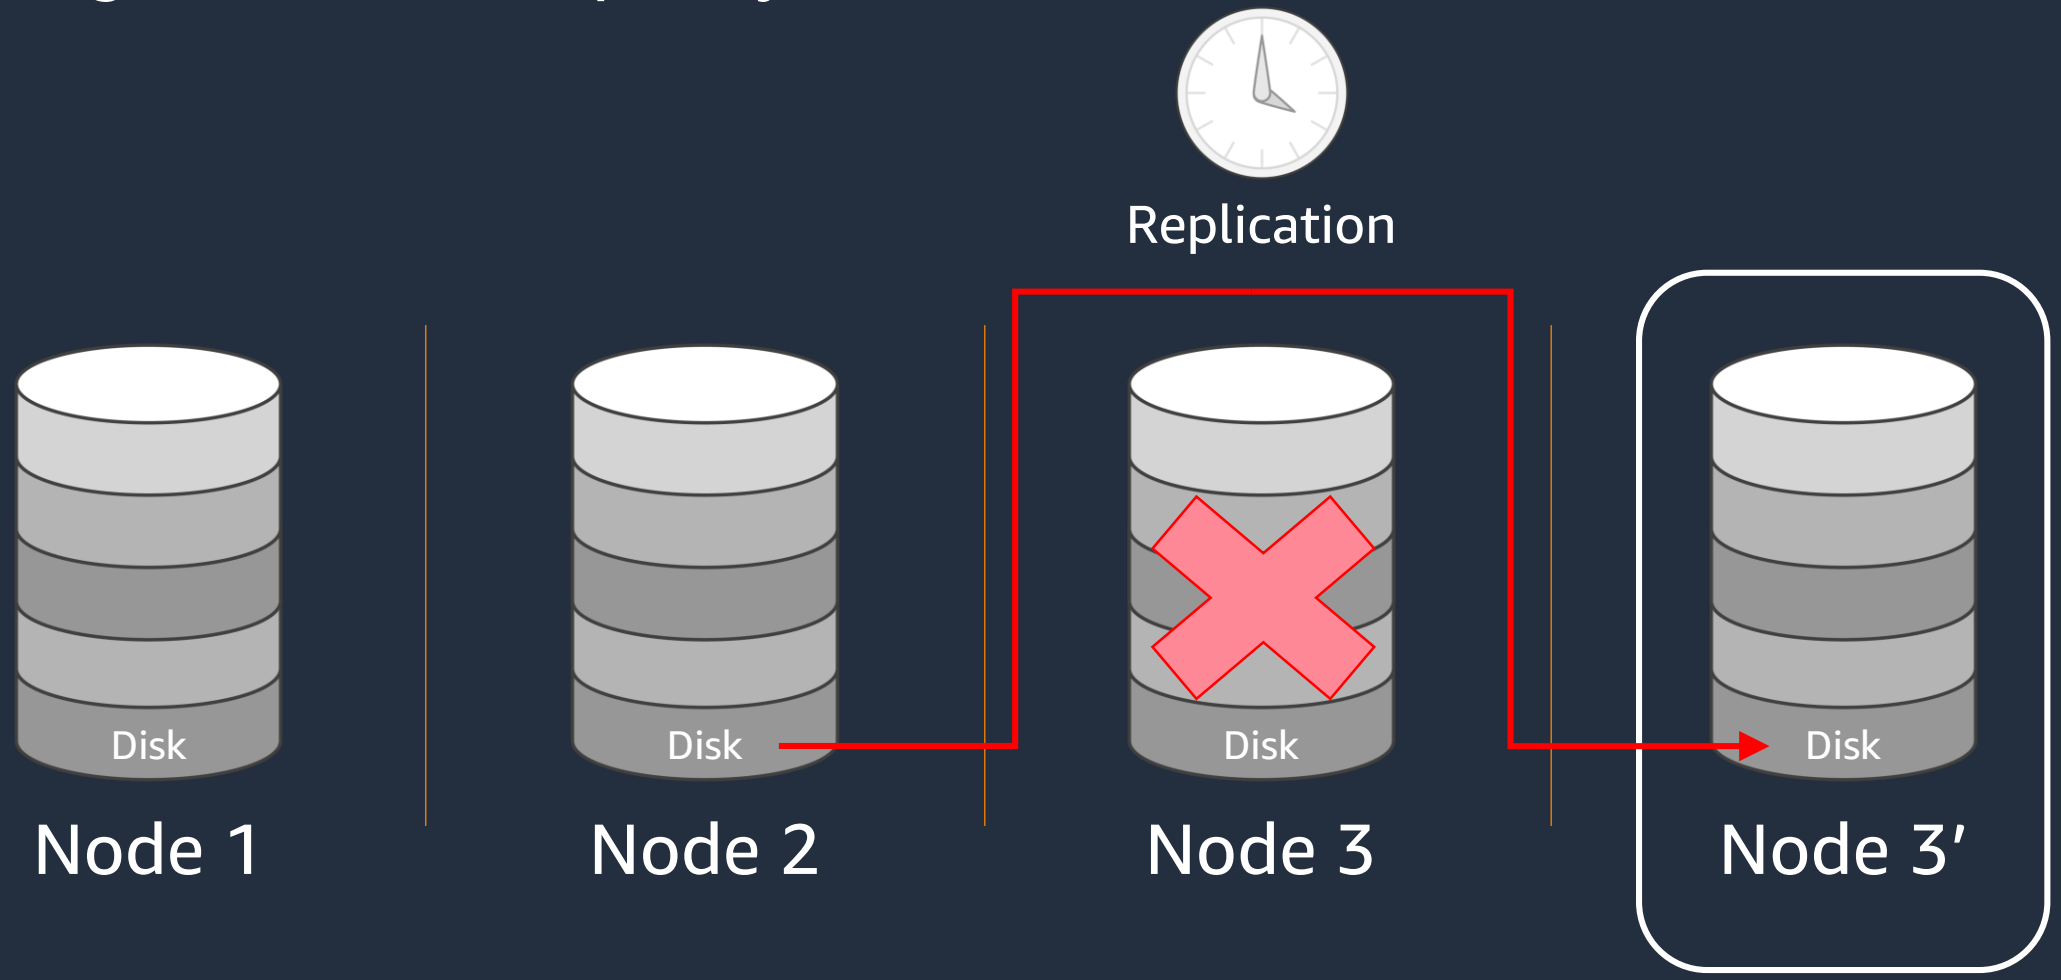

Traditional Database Architecture Challenges

Challenge #3: Scale storage as data grows

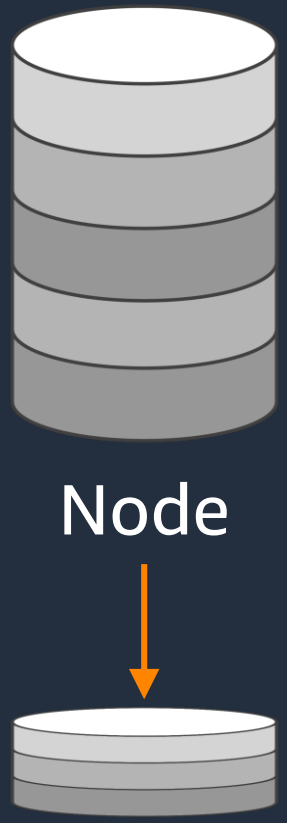

Traditional Database Architecture Challenges

Challenge #3: Scale storage as data grows

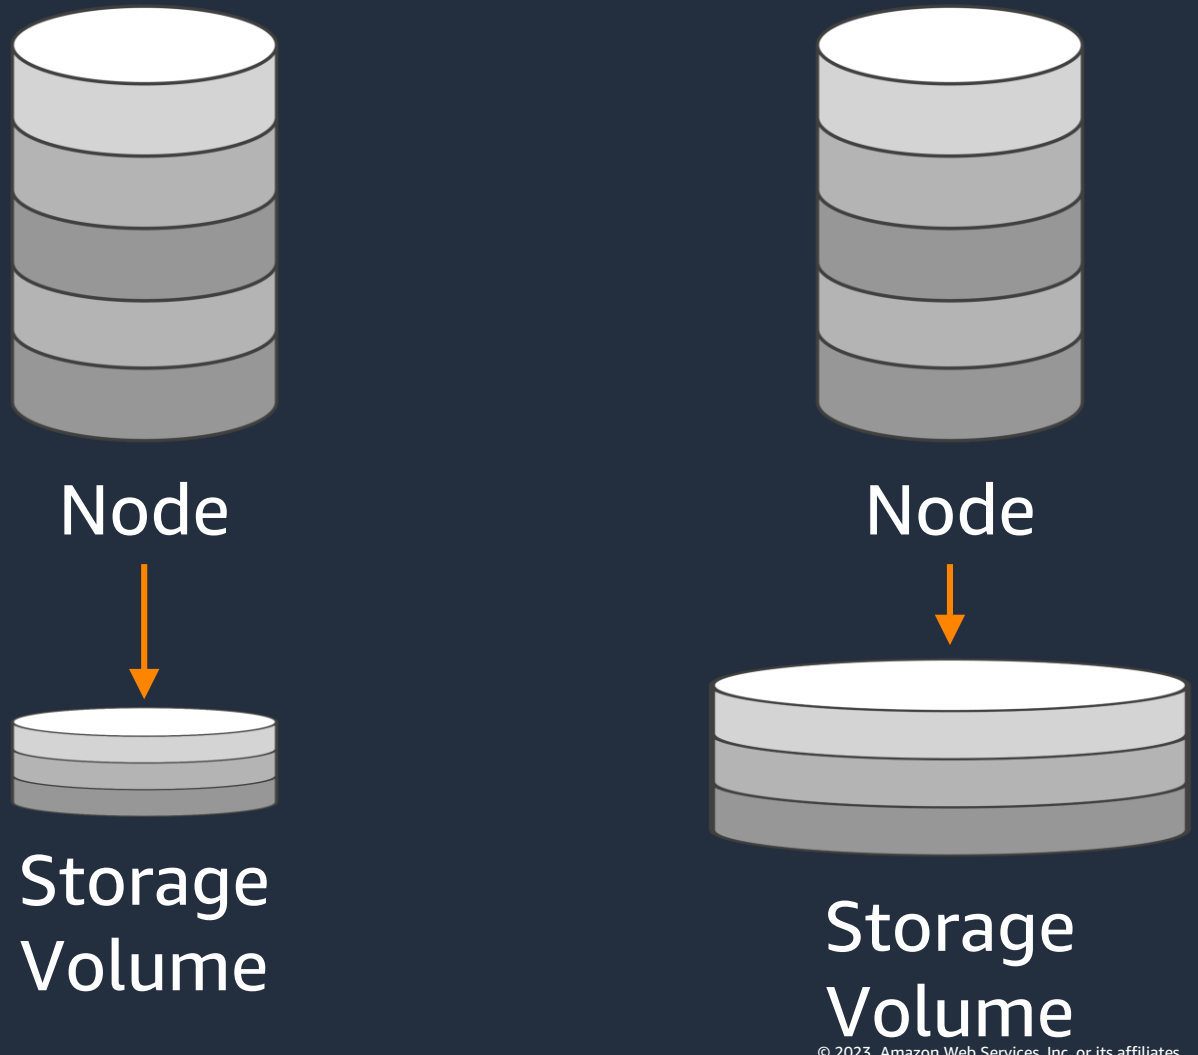

Traditional Database Architecture Challenges

Challenge #3: Scale storage as data grows

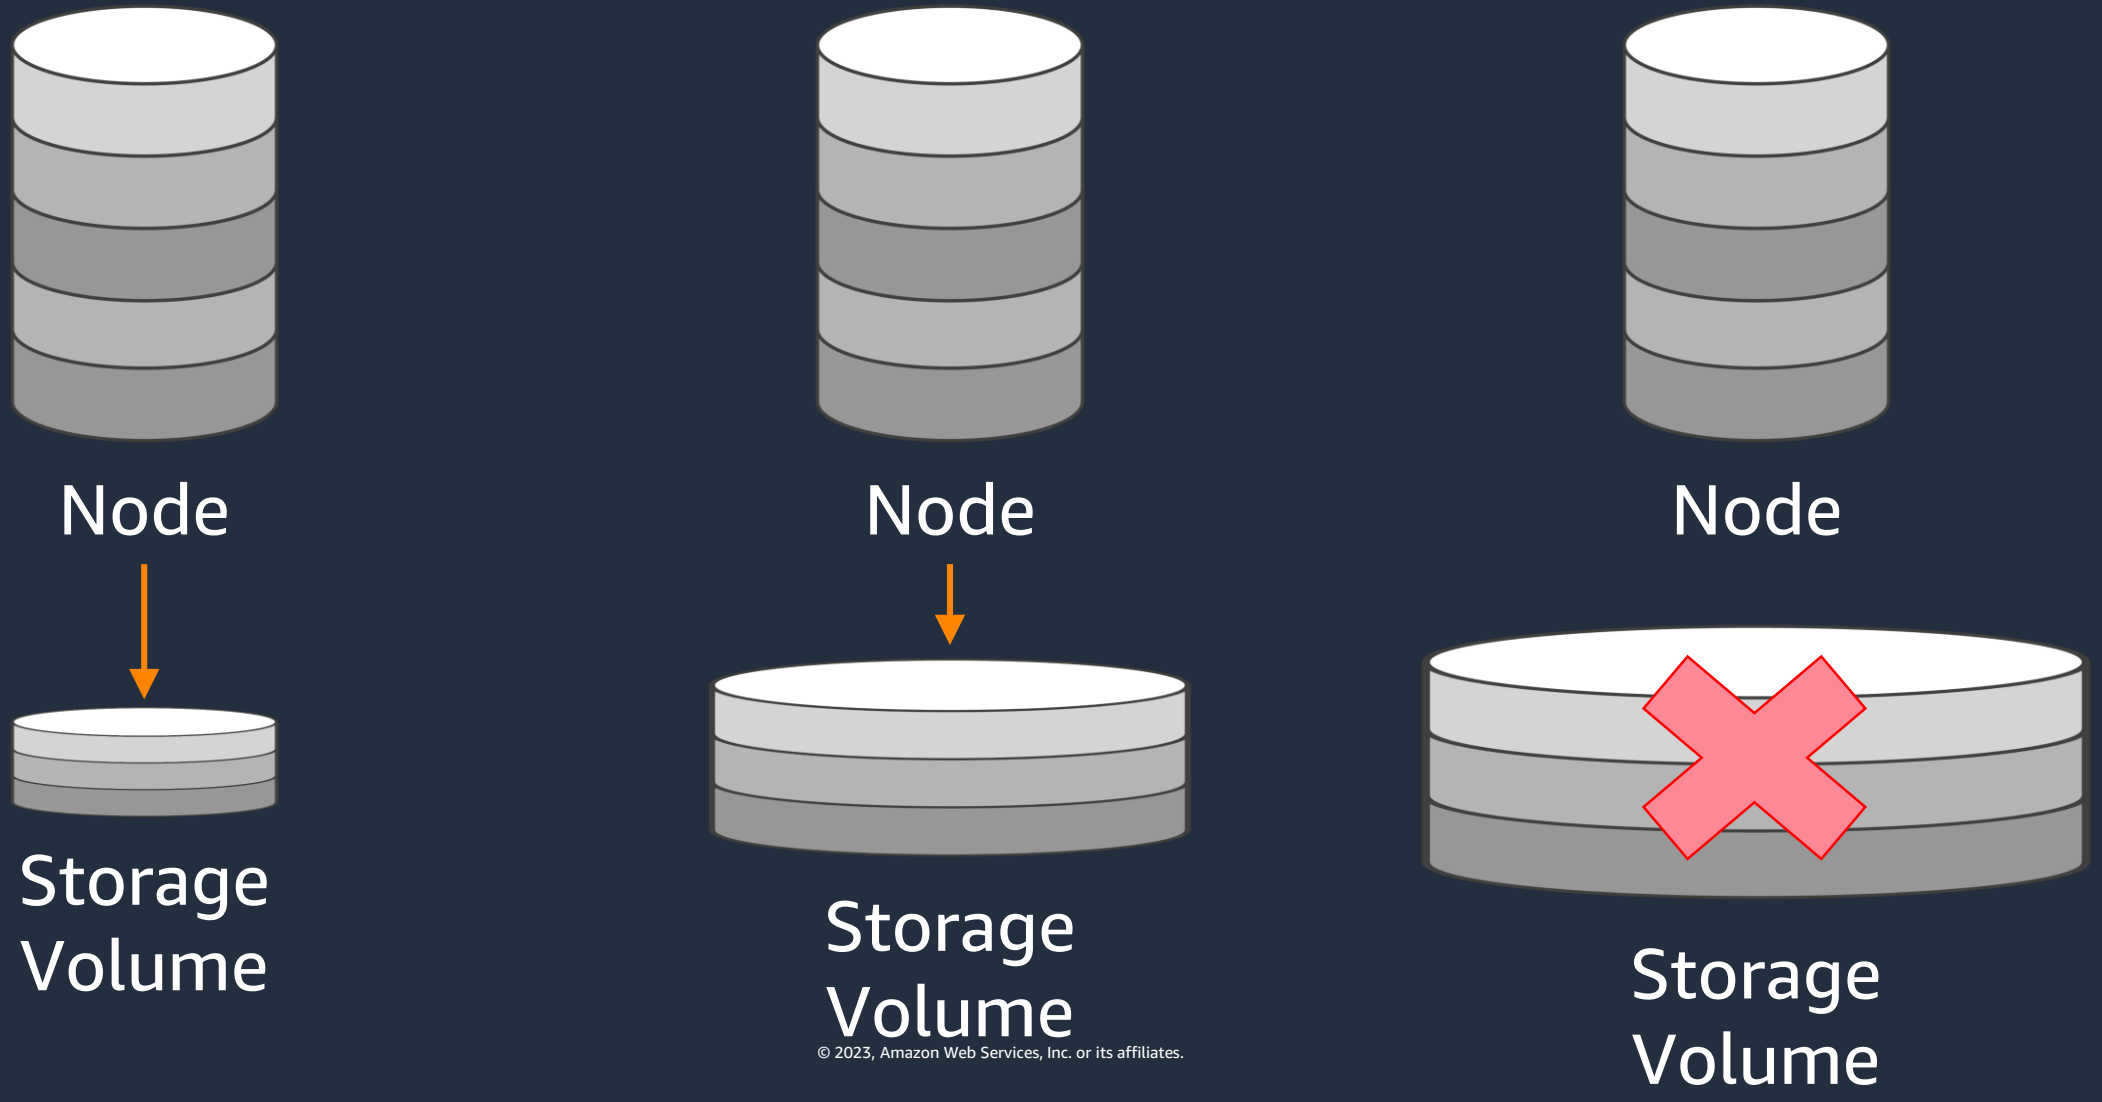

Traditional Database Architecture Challenges

Challenge #4: Backup data without affecting performance

Operational
Node

Traditional Database Architecture Challenges

Challenge #4: Backup data without affecting performance

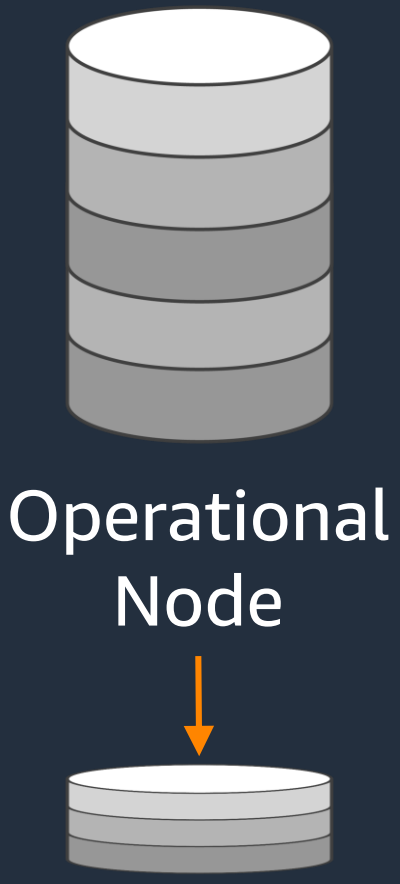

Traditional Database Architecture Challenges

Challenge #4: Backup data without affecting performance

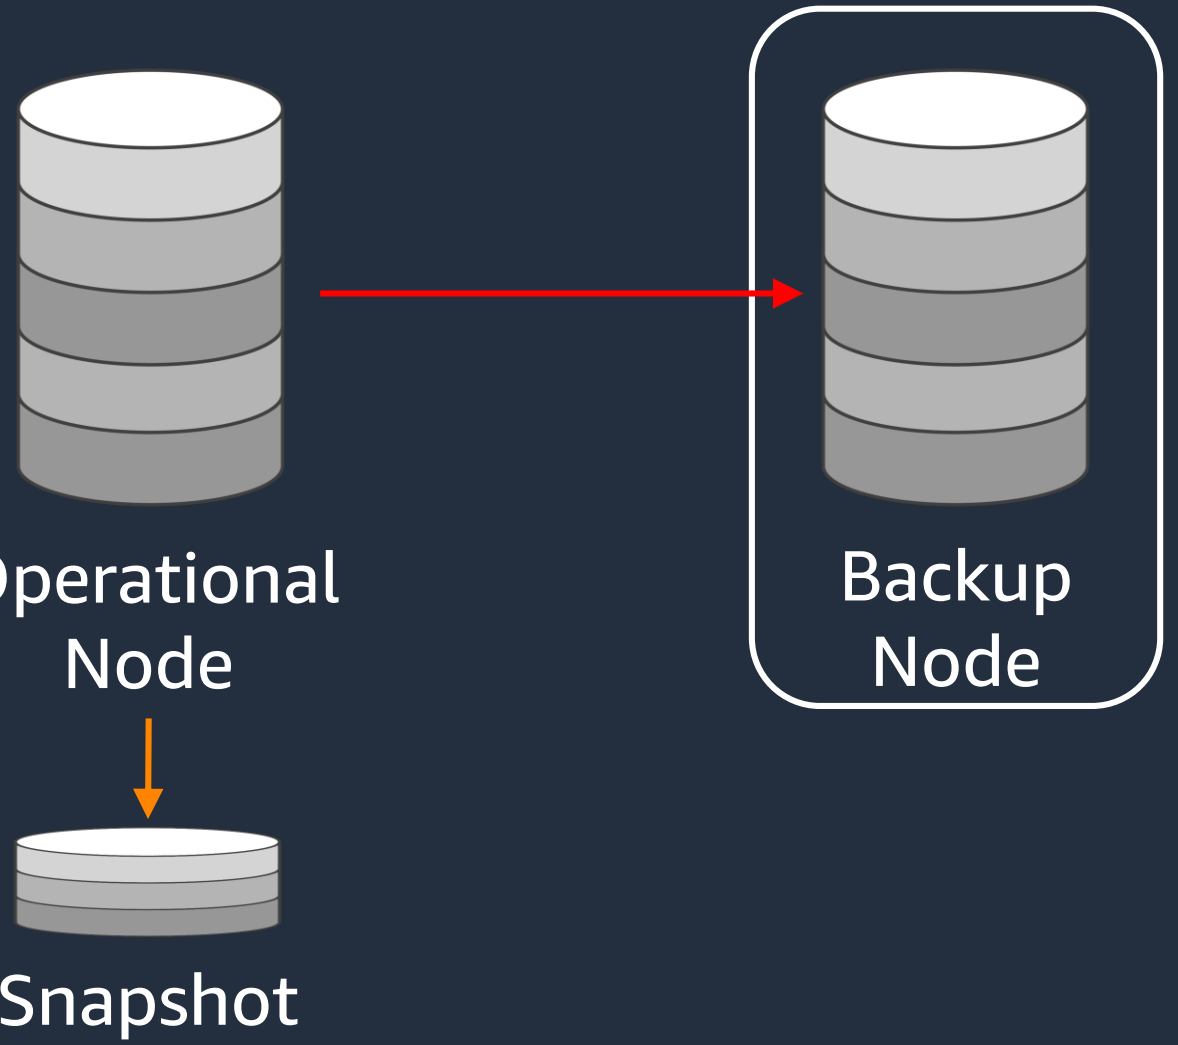

Traditional Database Architecture Challenges

Challenge #4: Backup data without affecting performance

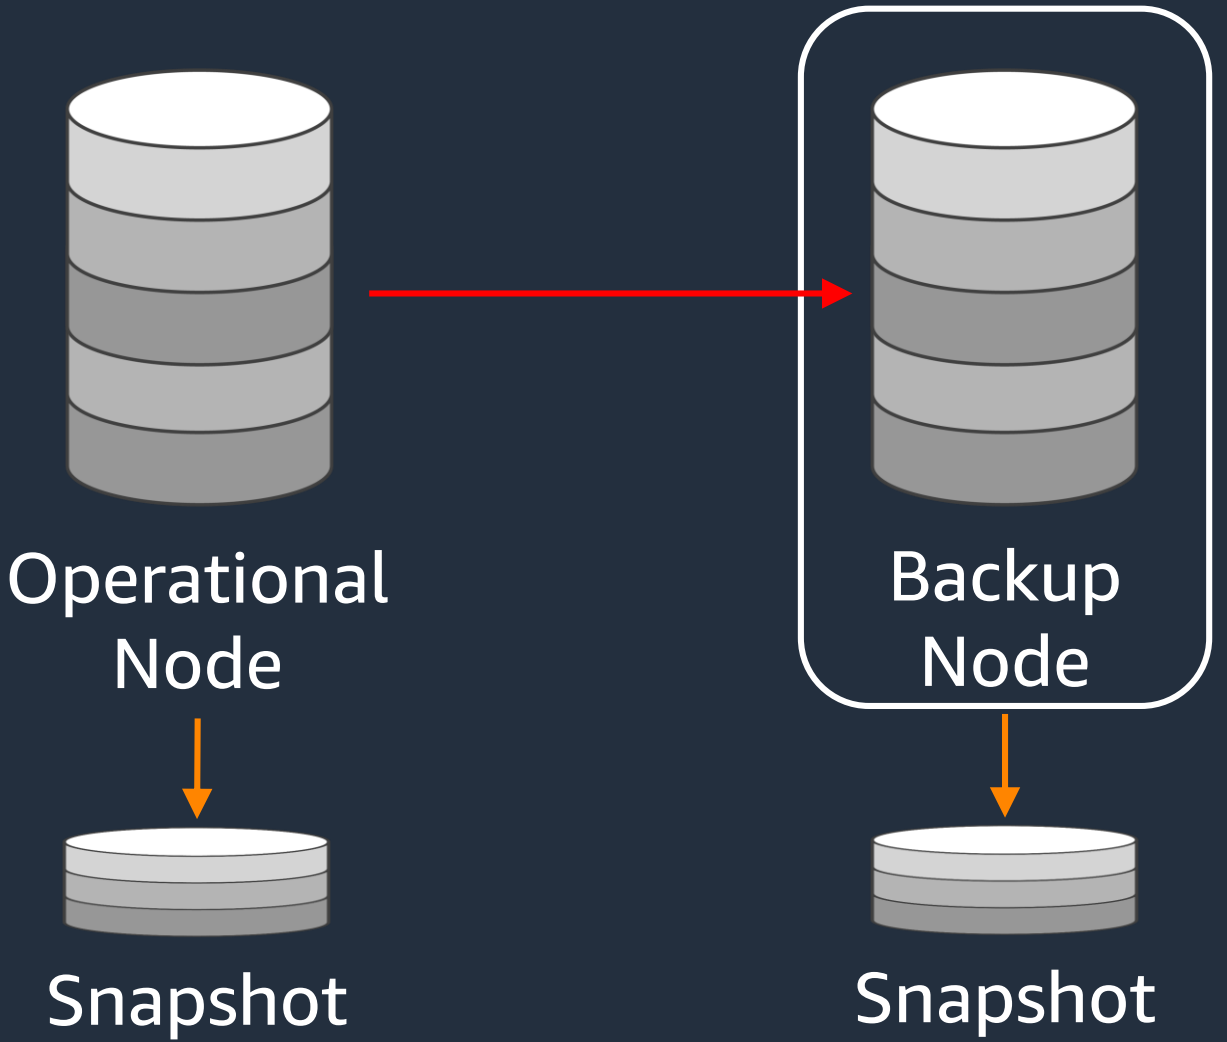

Traditional Database Architecture Challenges

Challenge #5: Data Durability

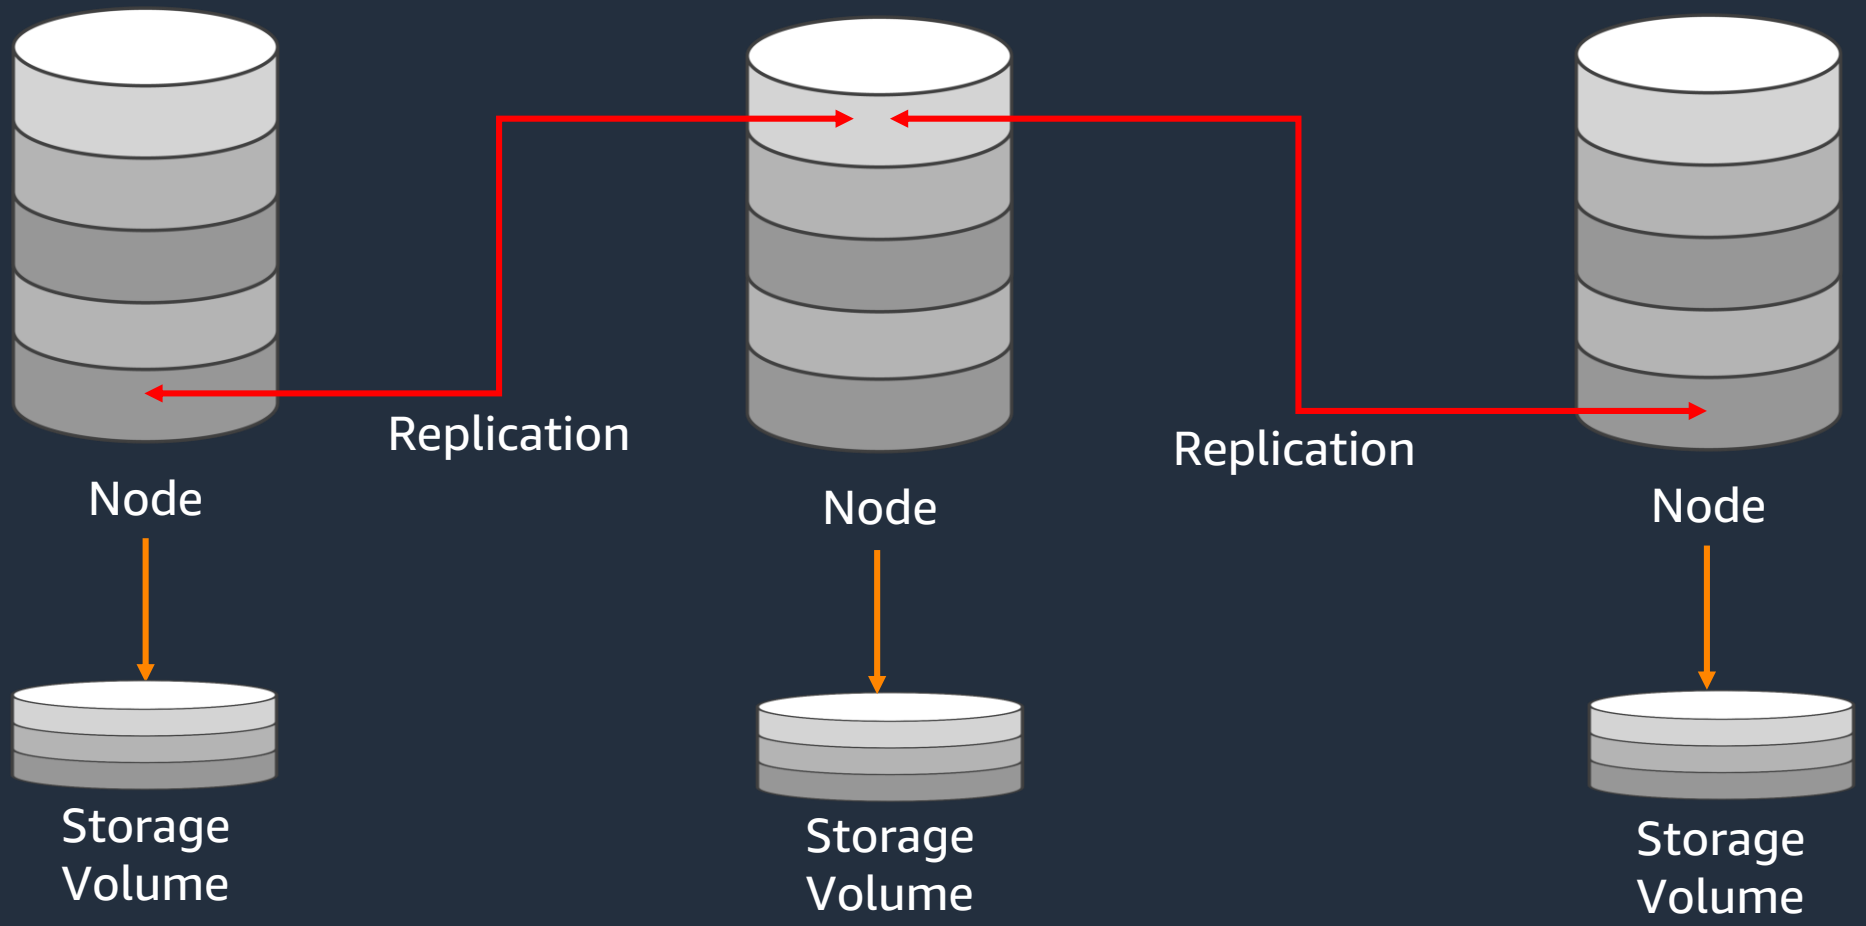

Amazon DocumentDB

Purpose-built and engineered for the cloud

Compute layer

API

Query processor

Caching

Logging

Storage

Decouple compute and storage

Scale compute

Storage layer

Scale storage

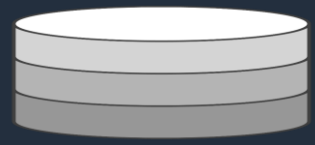

aws

Amazon DocumentDB: Cloud Native Architecture

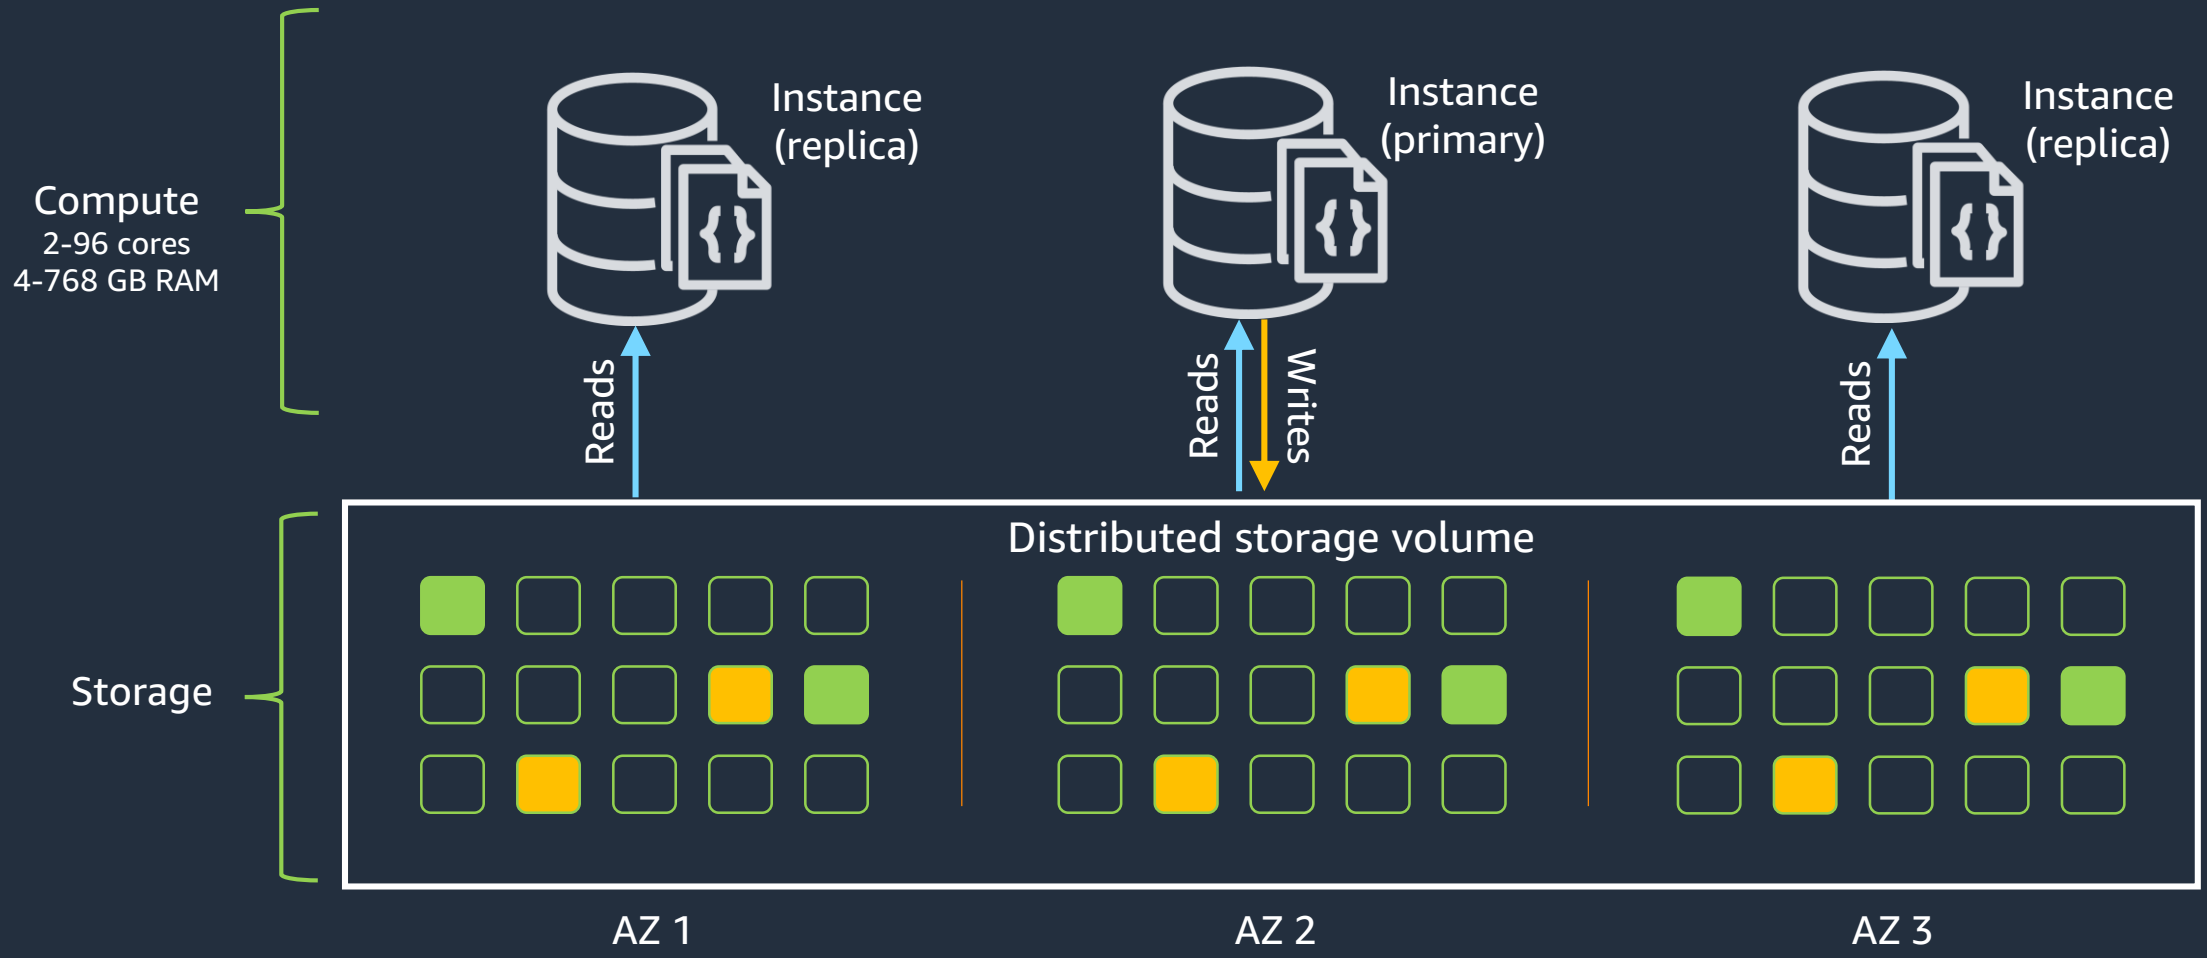

Challenge #1: Add Read Capacity on Demand

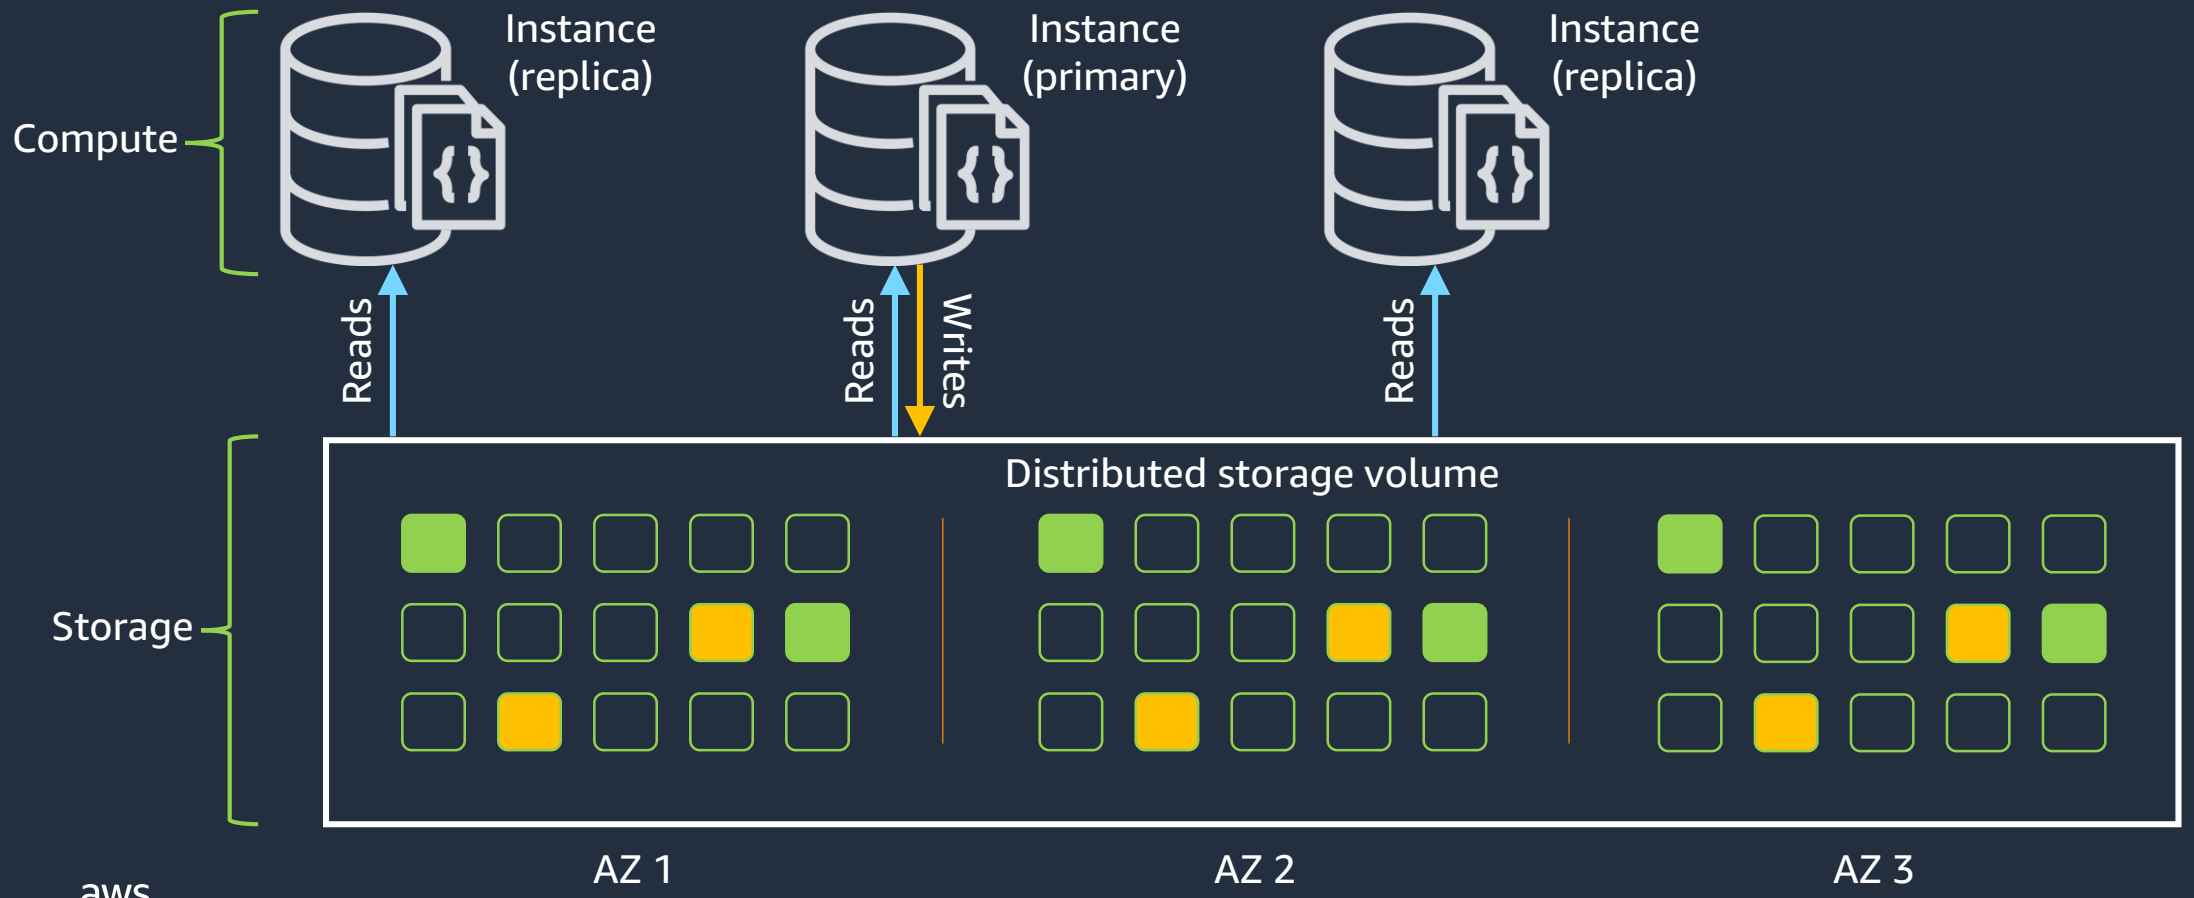

Challenge #1: Add Read Capacity on Demand

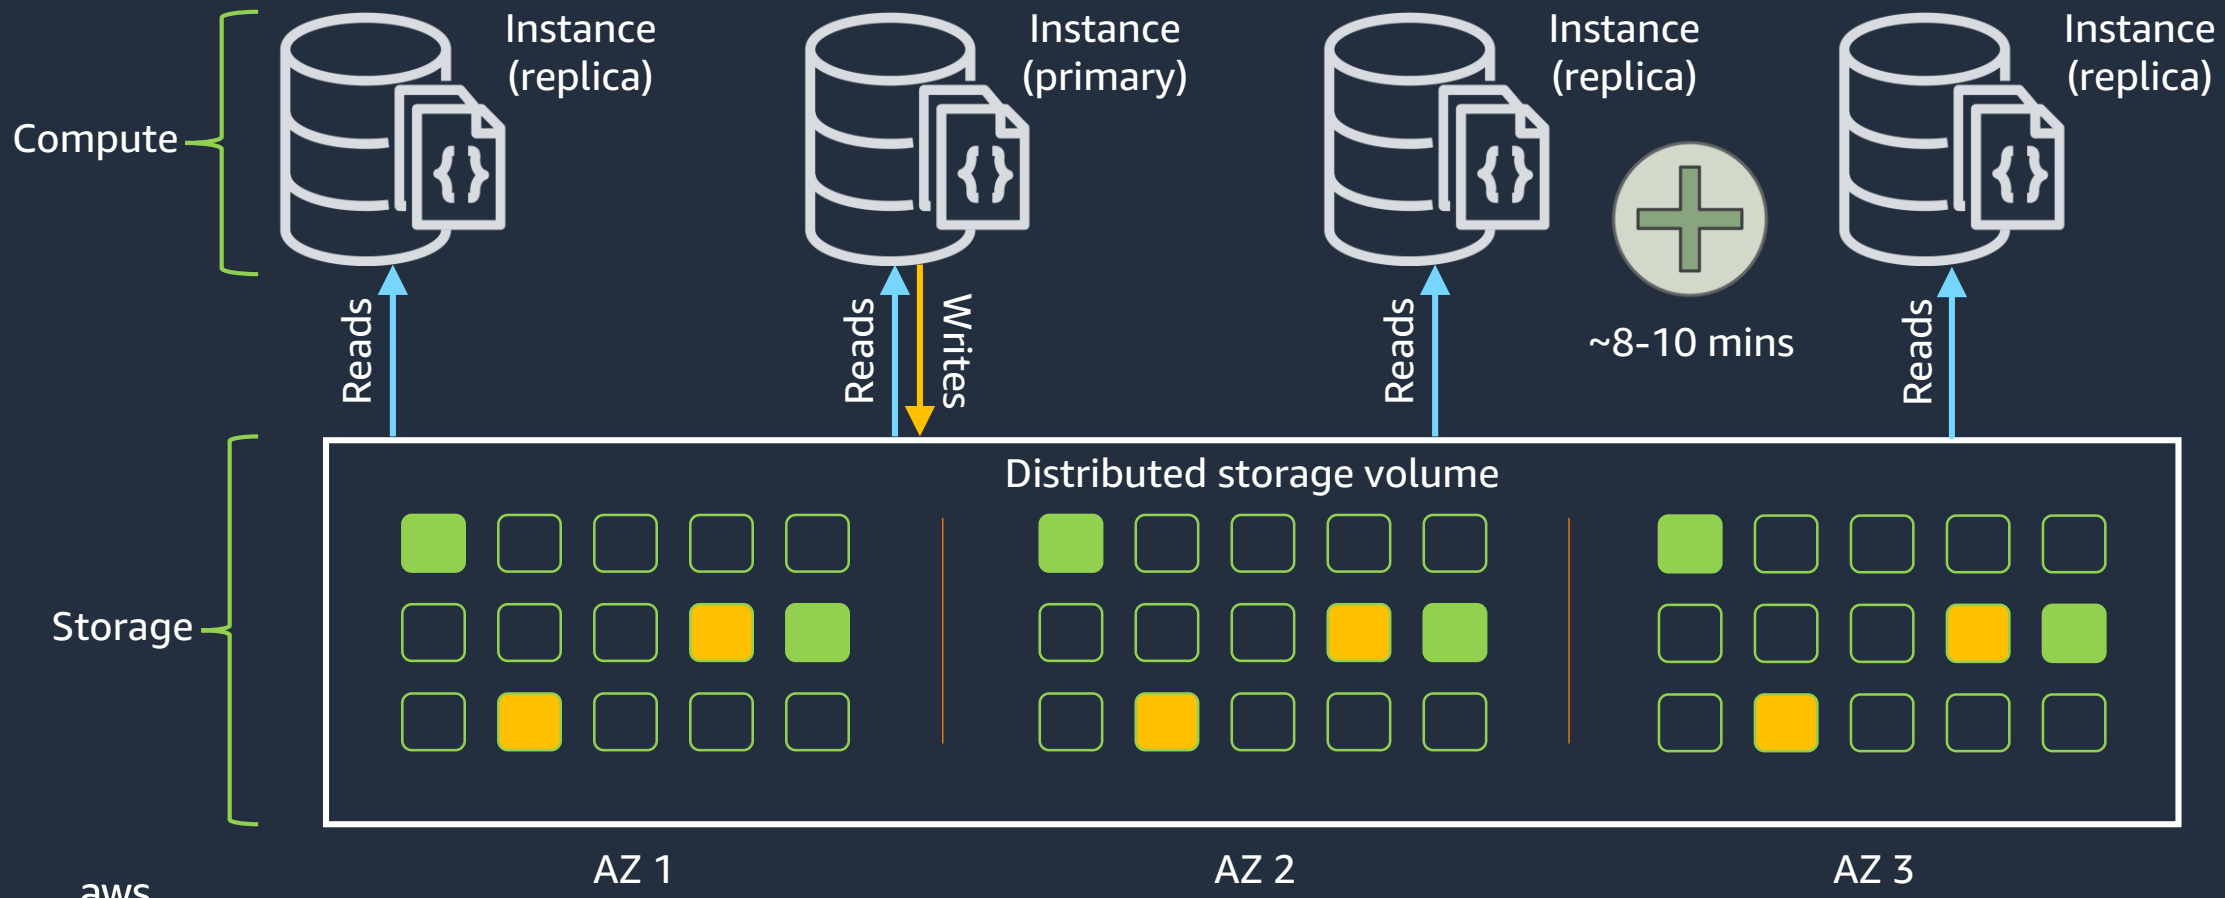

Challenge #2: Quickly Recover From Node Failure

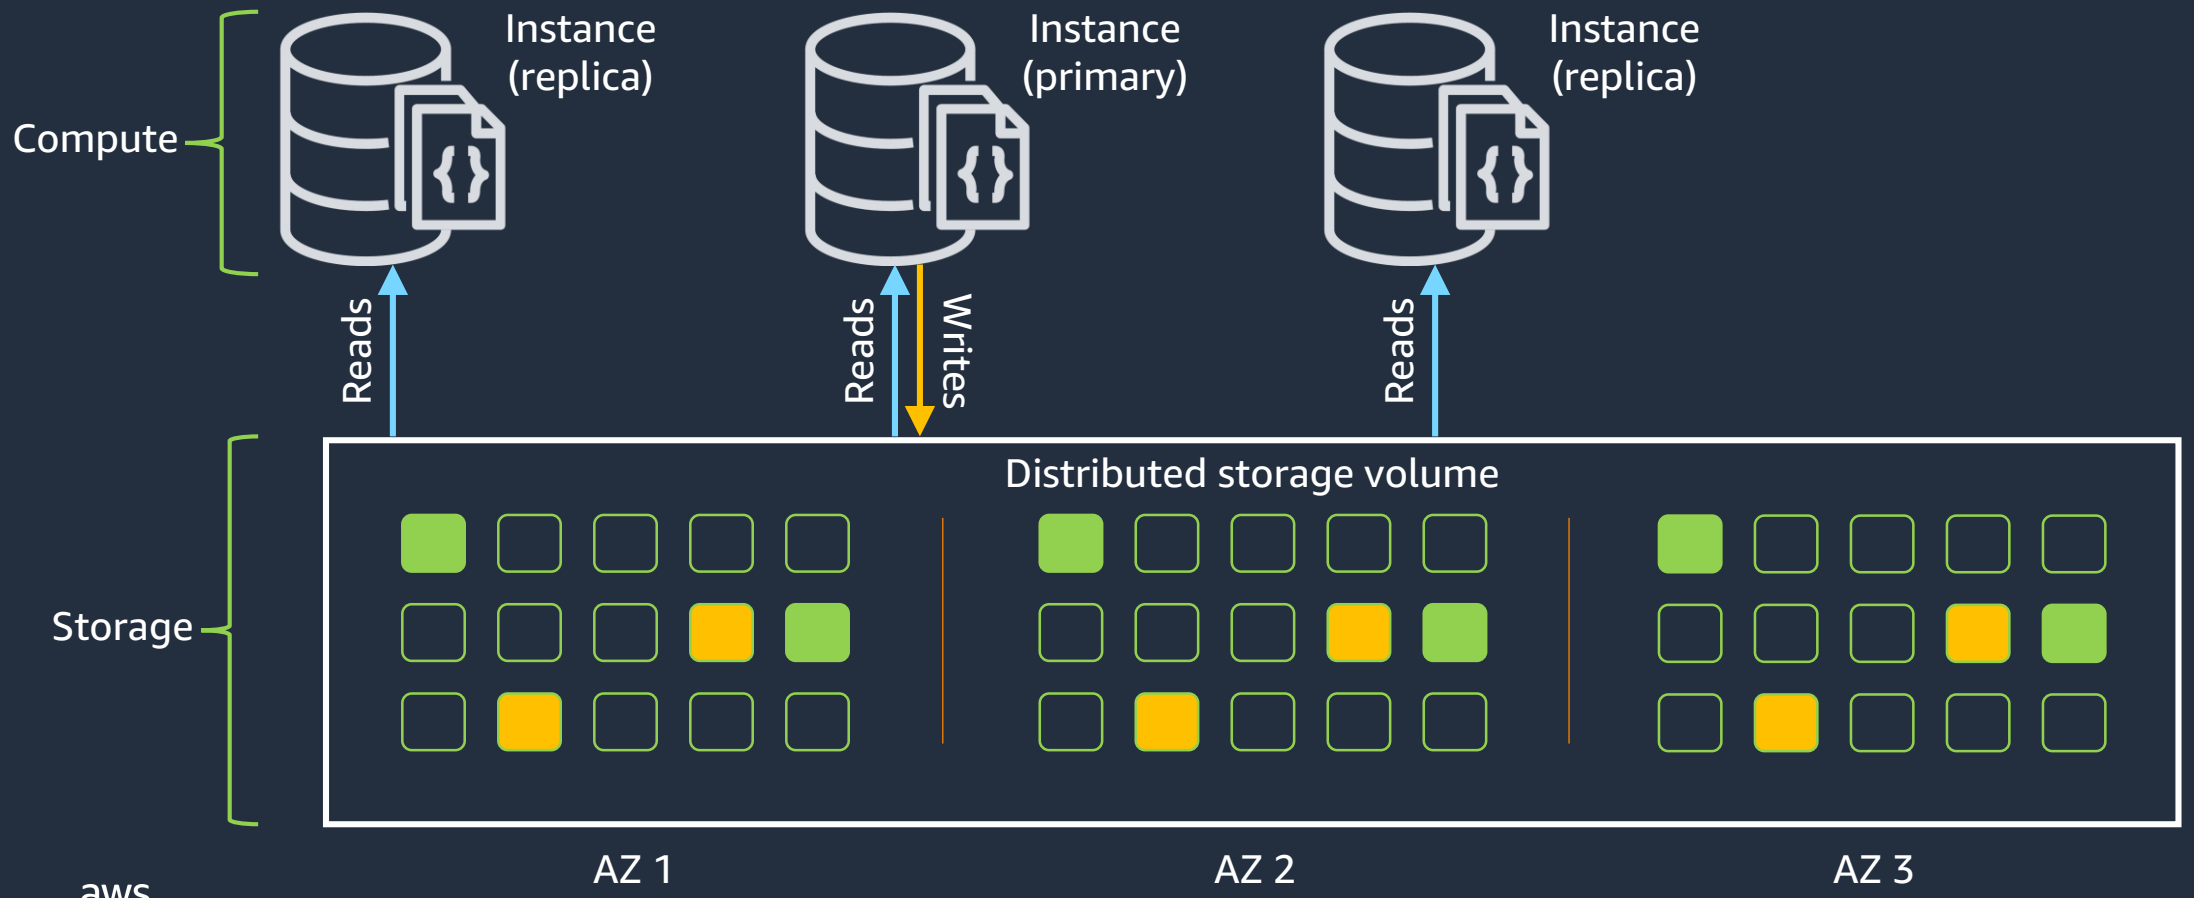

Challenge #2: Quickly Recover From Node Failure

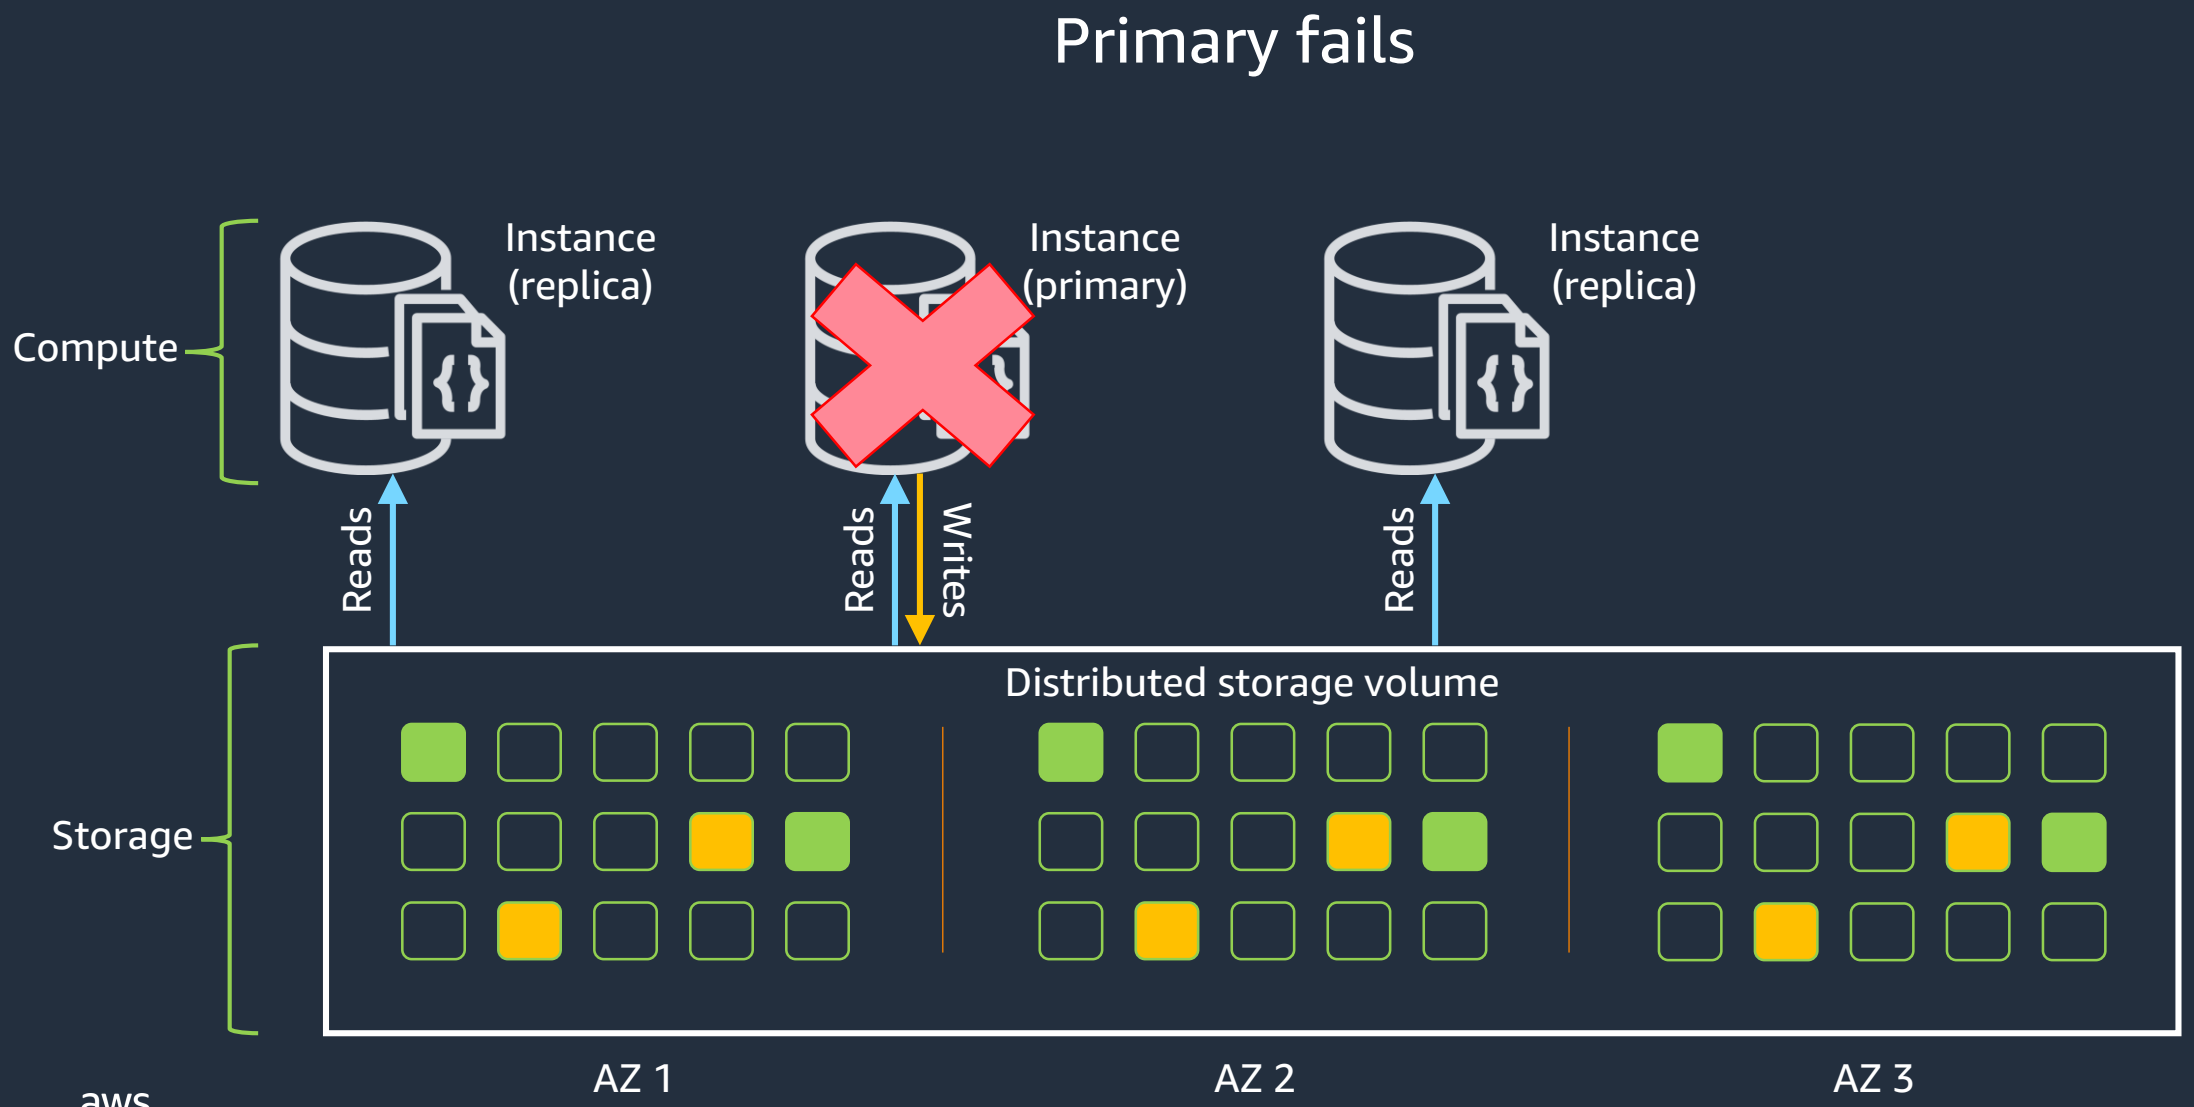

Challenge #2: Quickly Recover From Node Failure

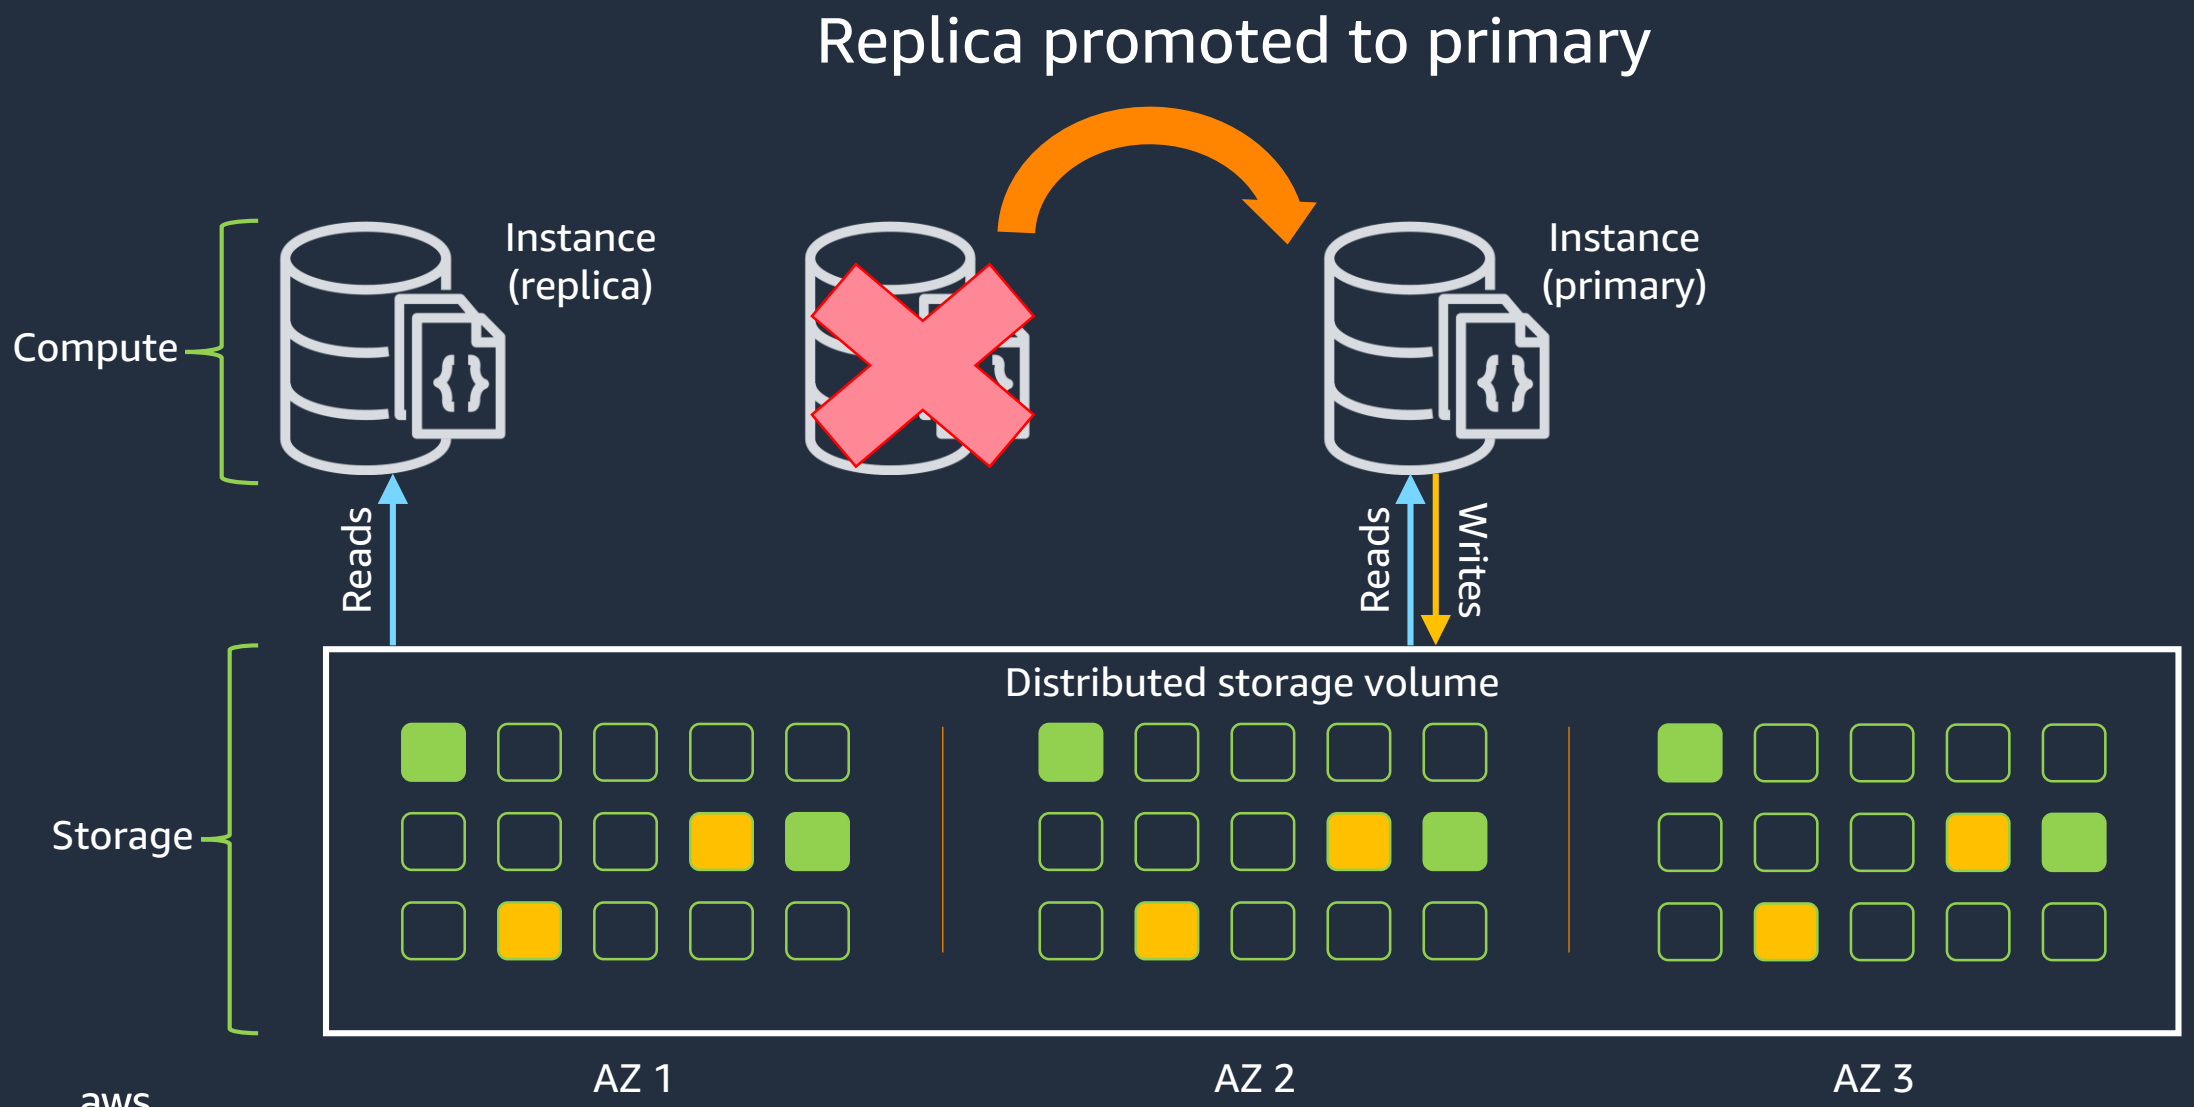

Challenge #2: Quickly Recover From Node Failure

New replica instance created
Automatic return to full strength

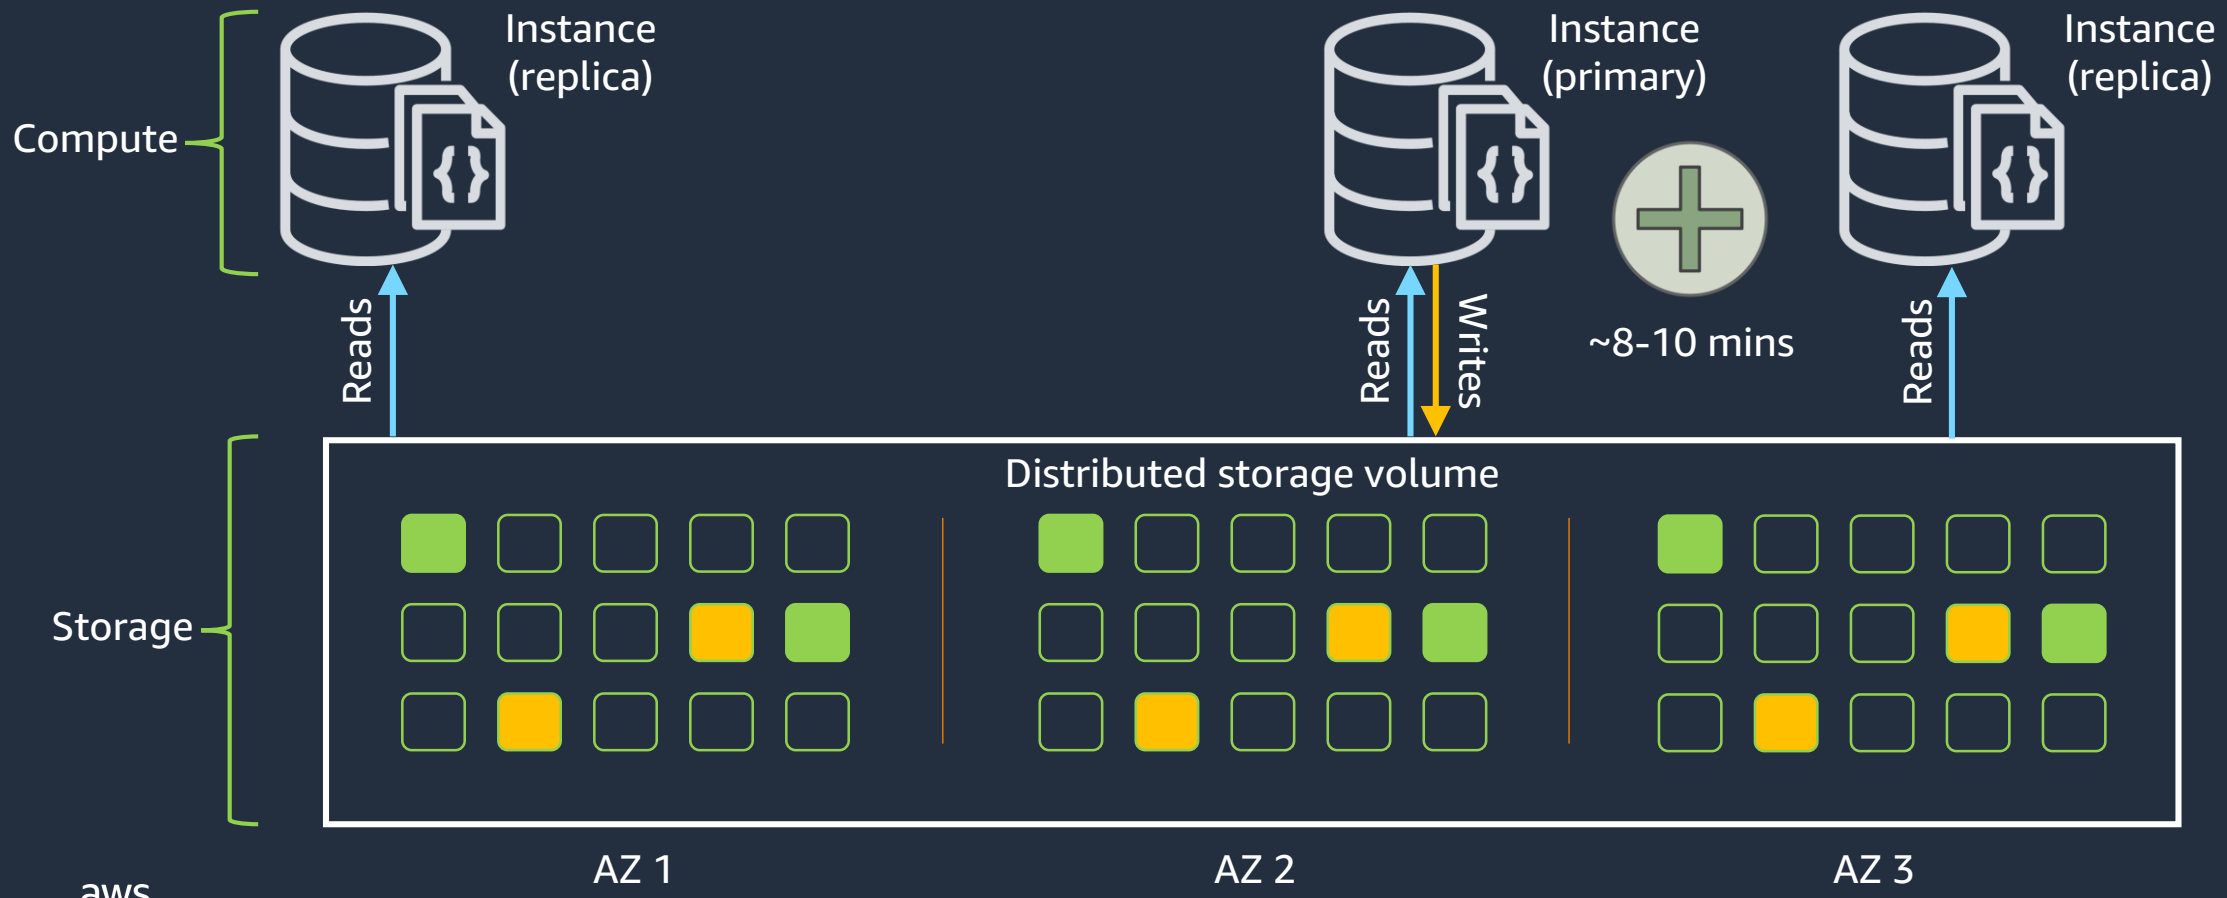

Challenge #3: Scale Storage as Data Grows

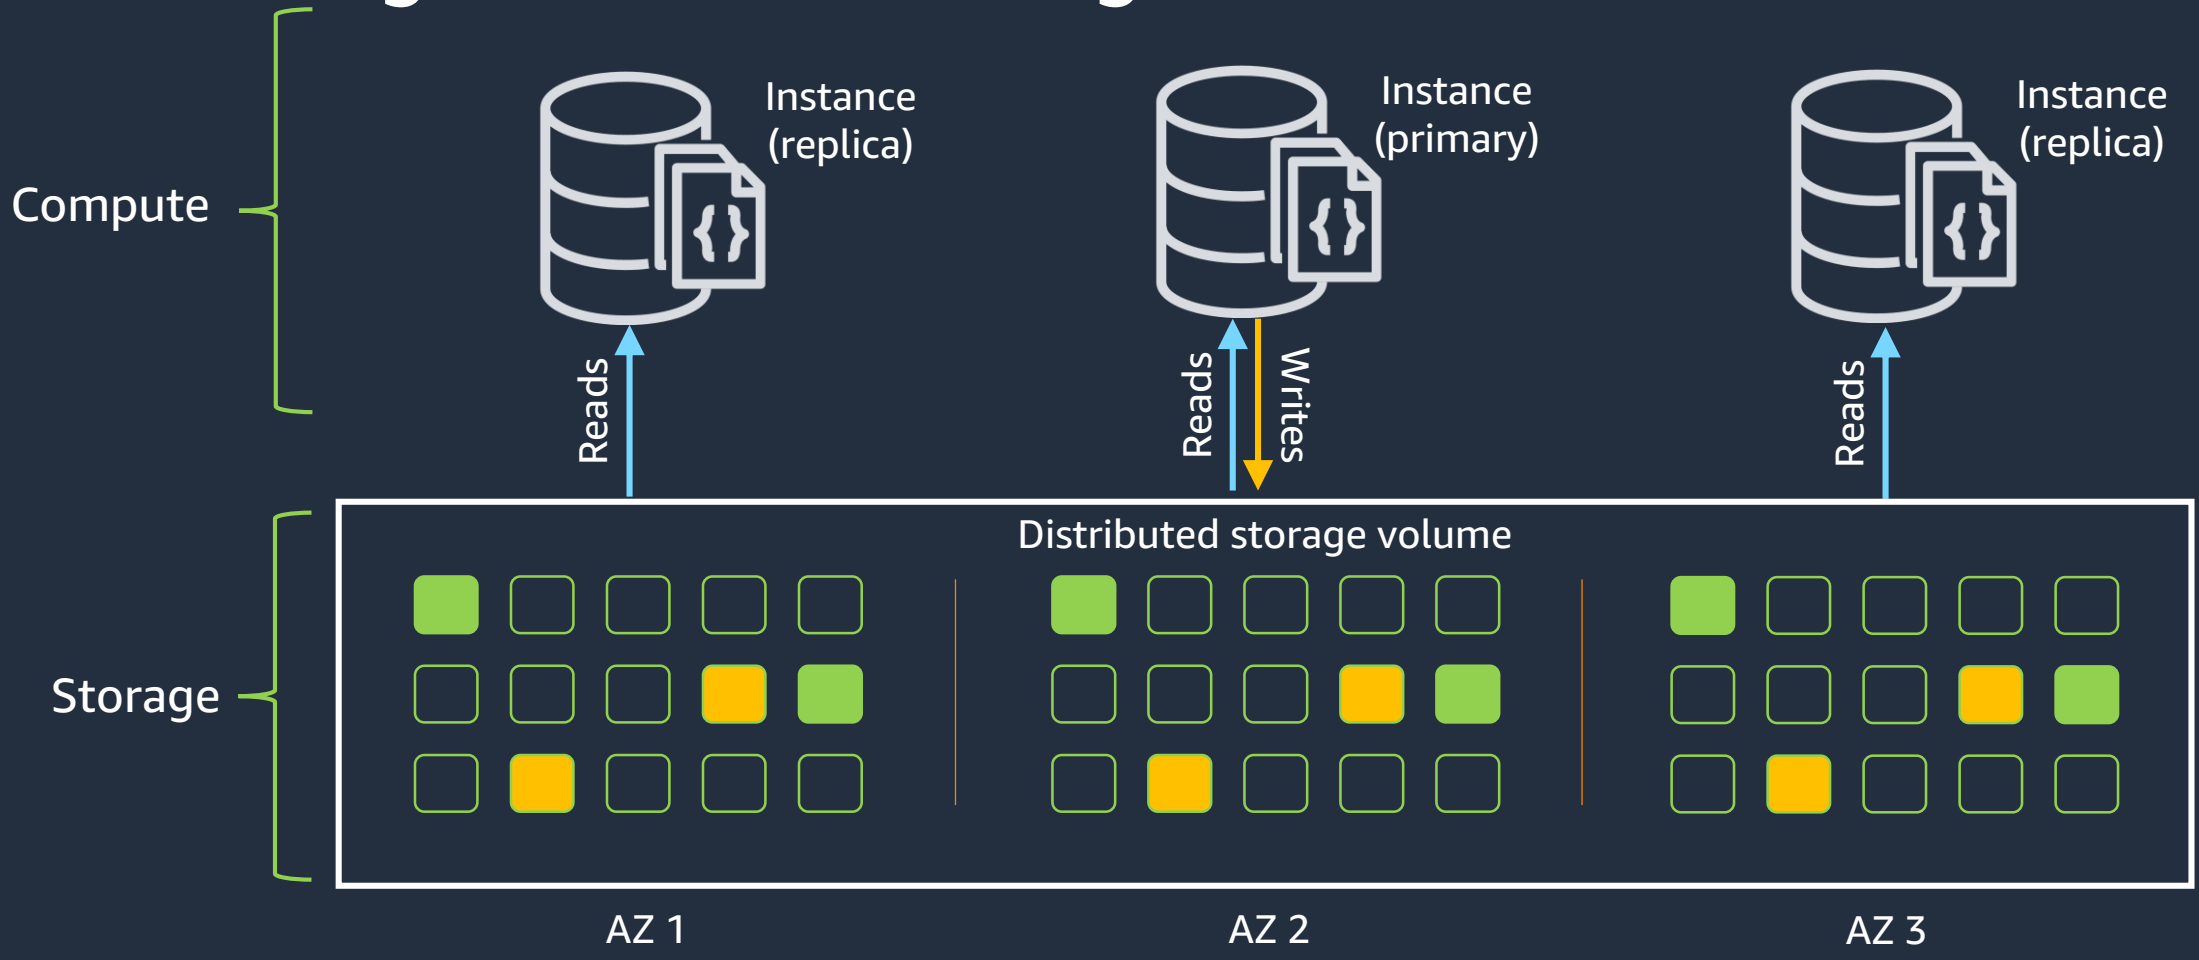

Challenge #3: Scale Storage as Data Grows

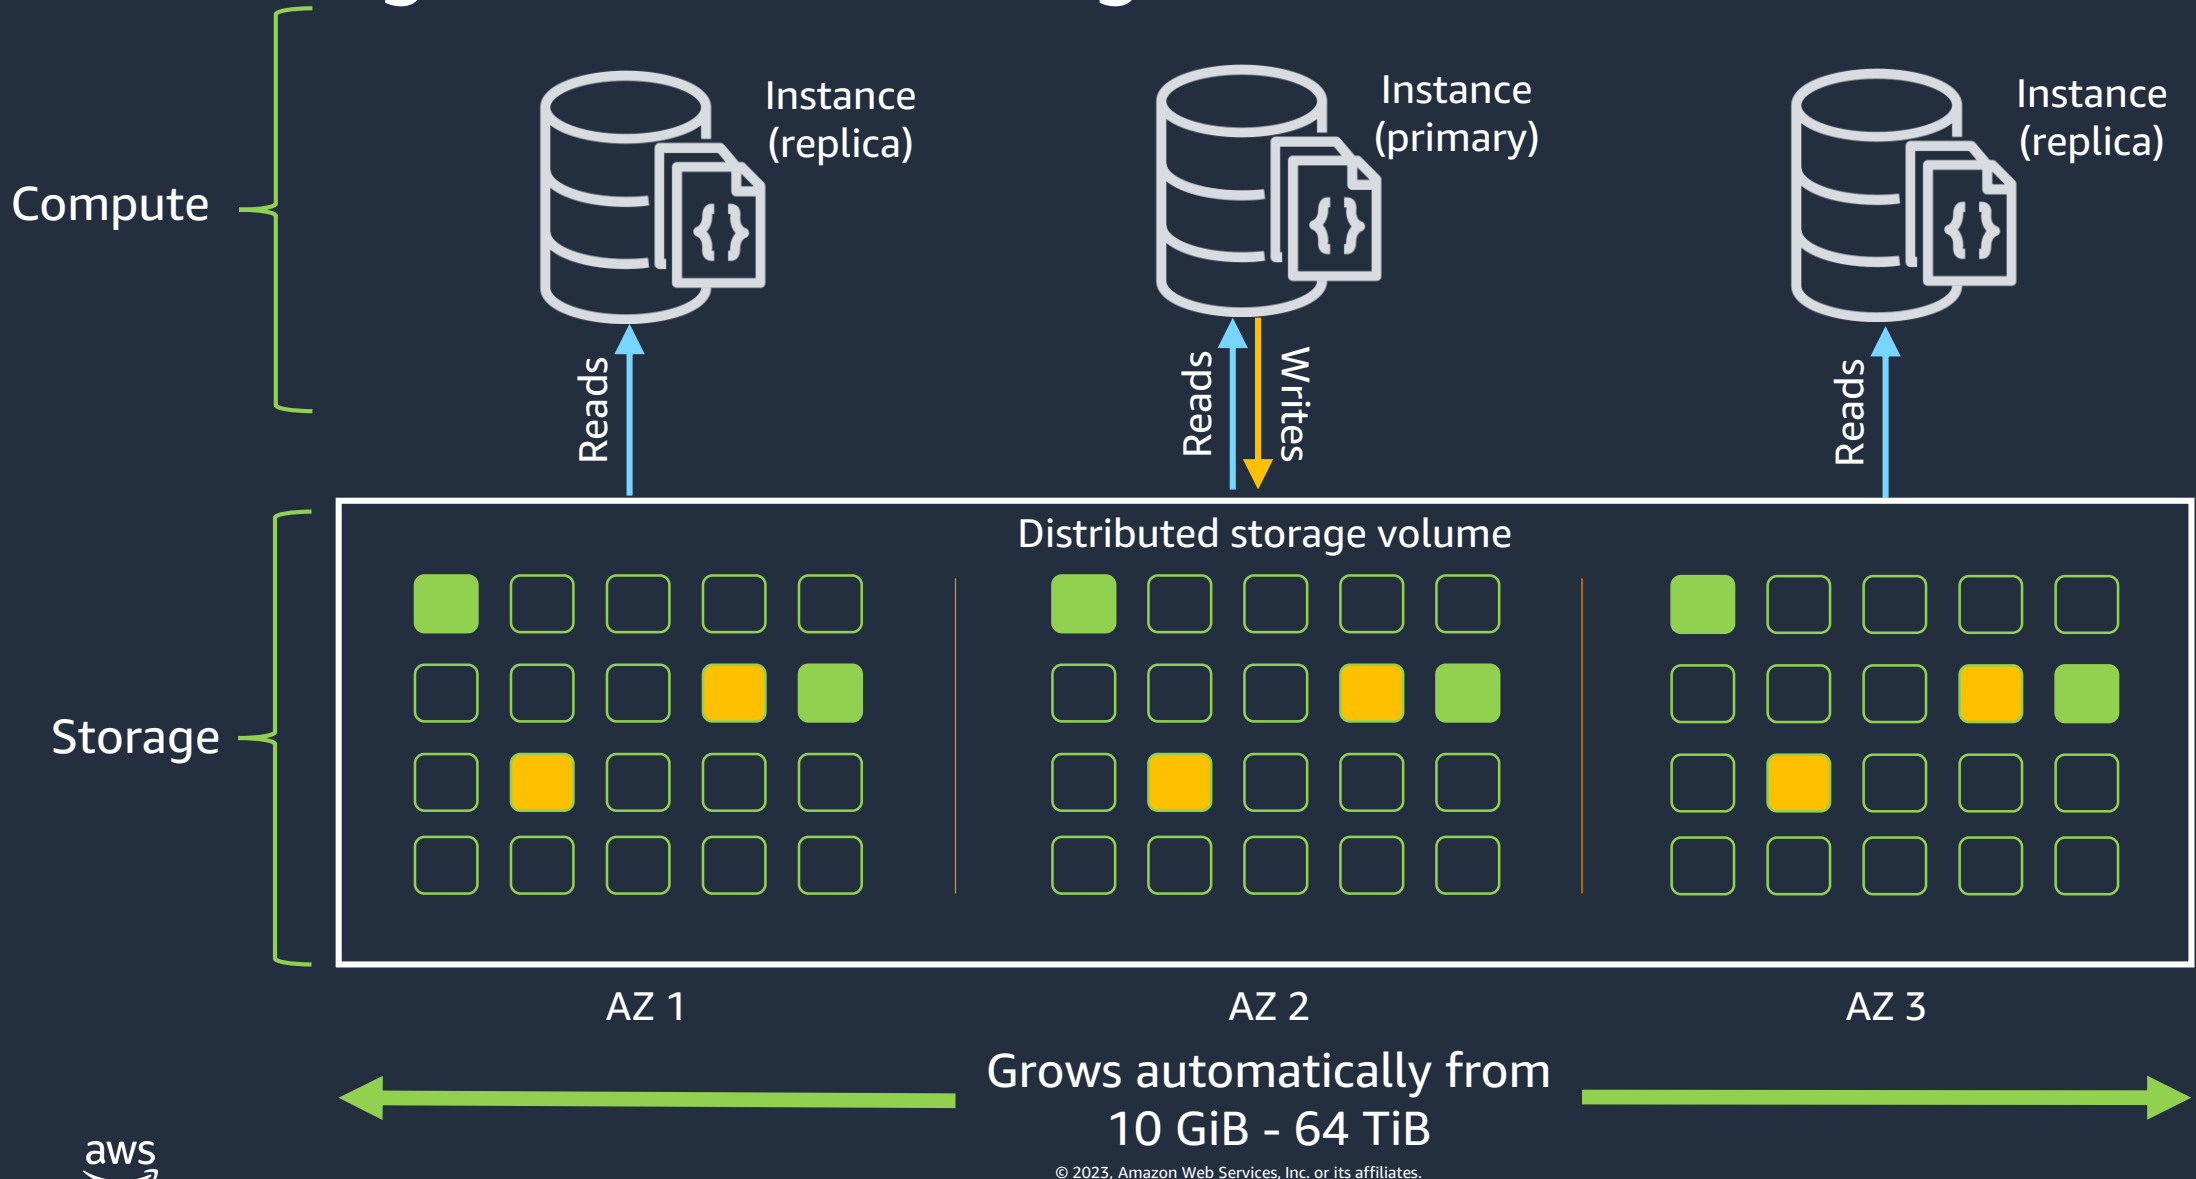

Challenge #4: Backup Data Without Affecting Performance

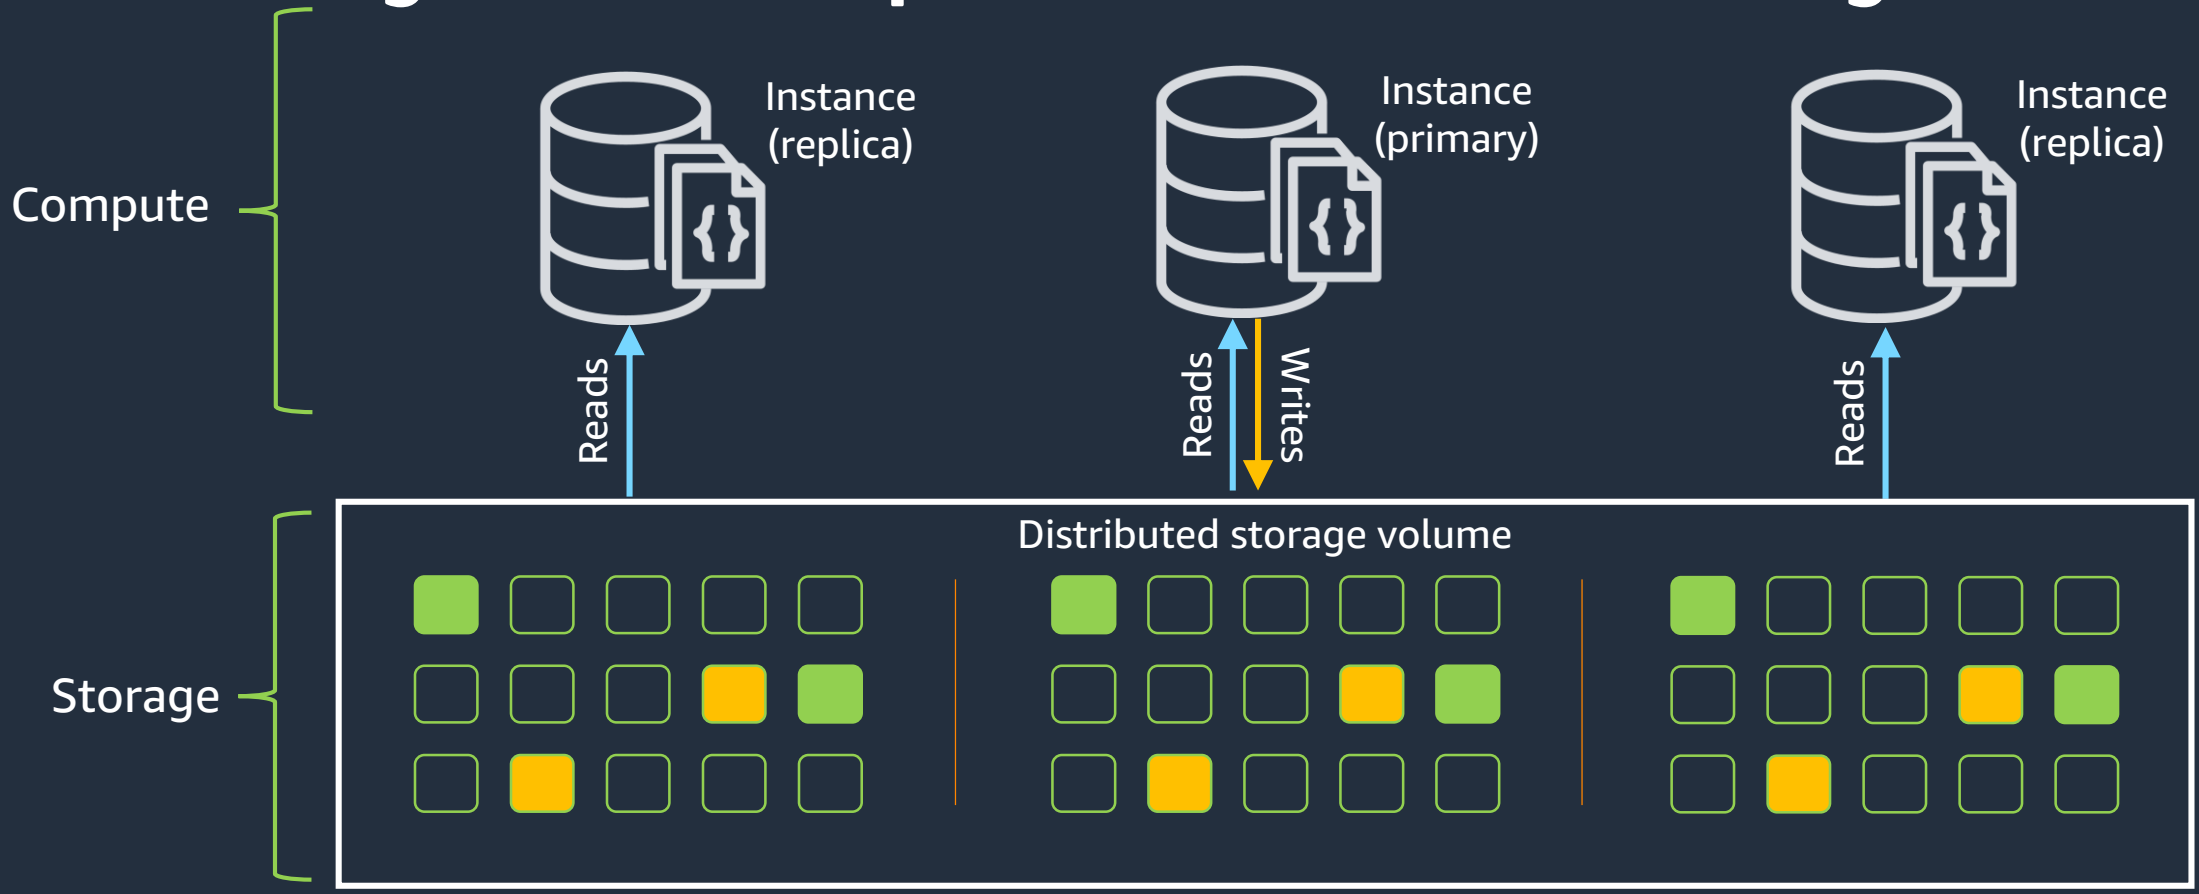

Challenge #4: Backup Data Without Affecting Performance

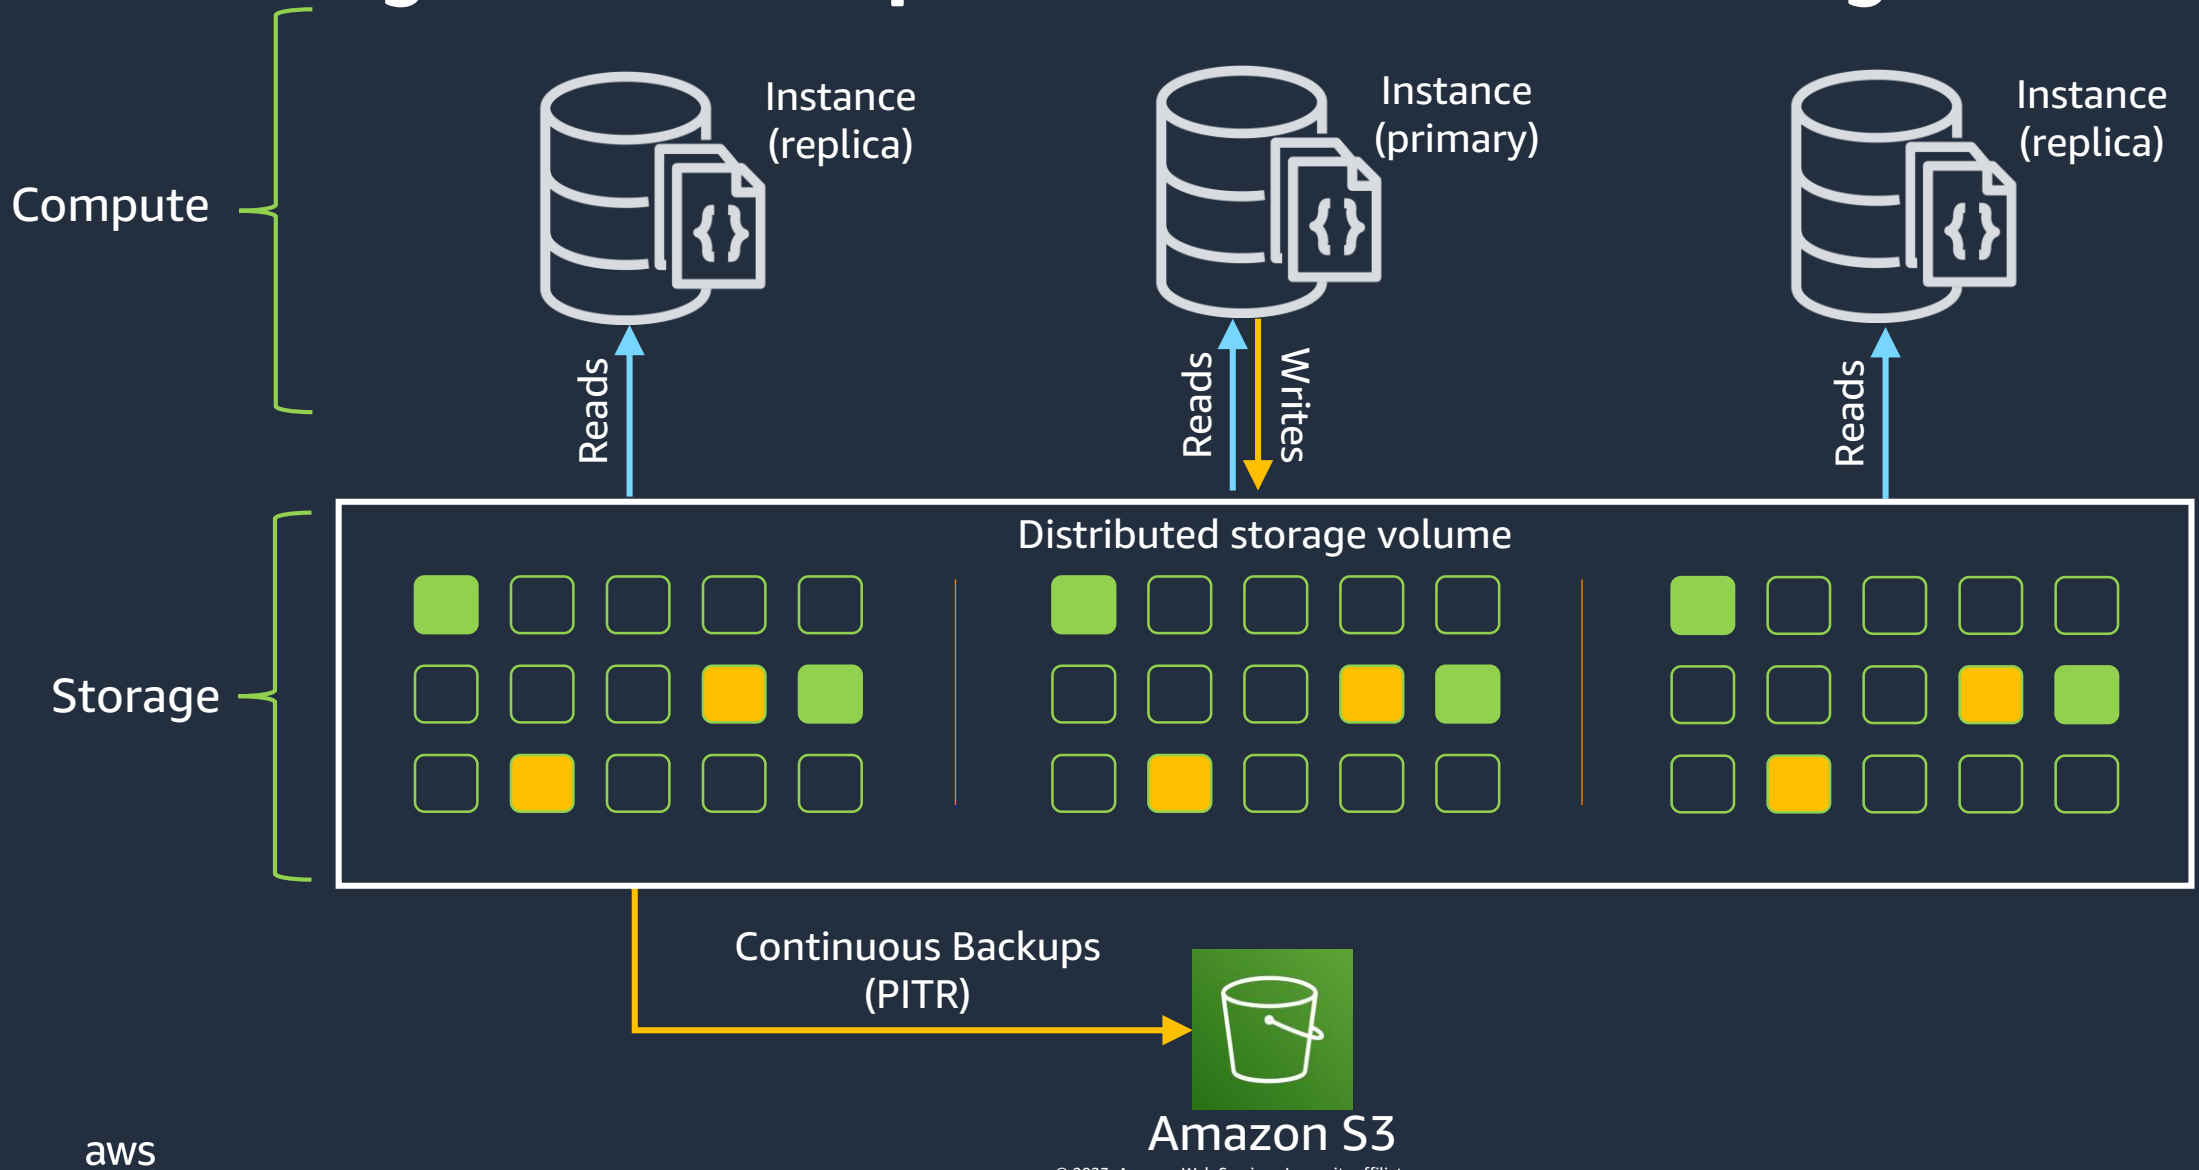

Challenge #4: Backup Data Without Affecting Performance

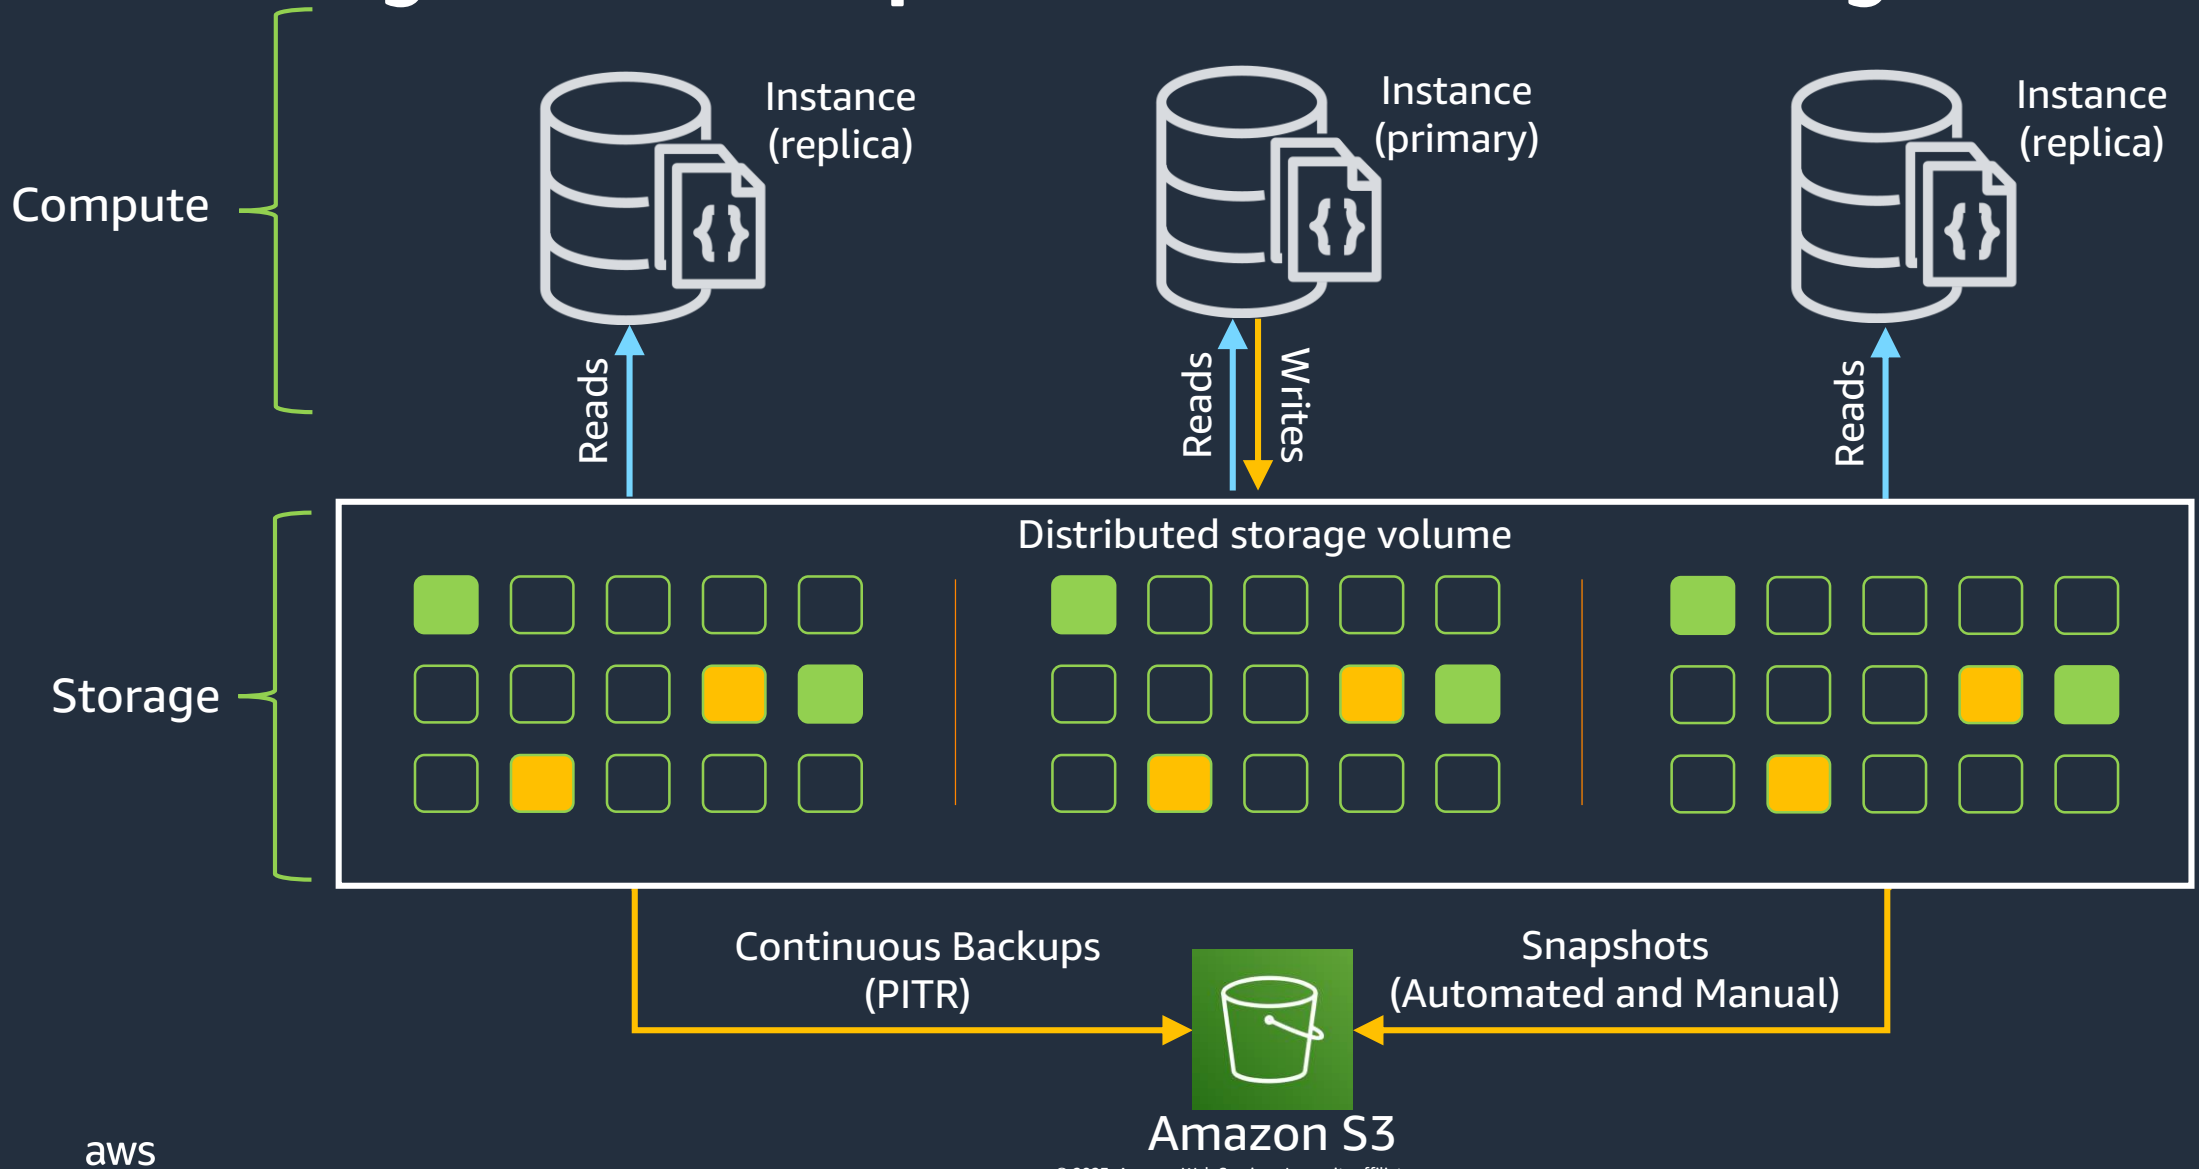

Challenge #5: Data Durability

Instances: 1

Environment: dev/test

Availability goal: 99%

Durability: highly durable

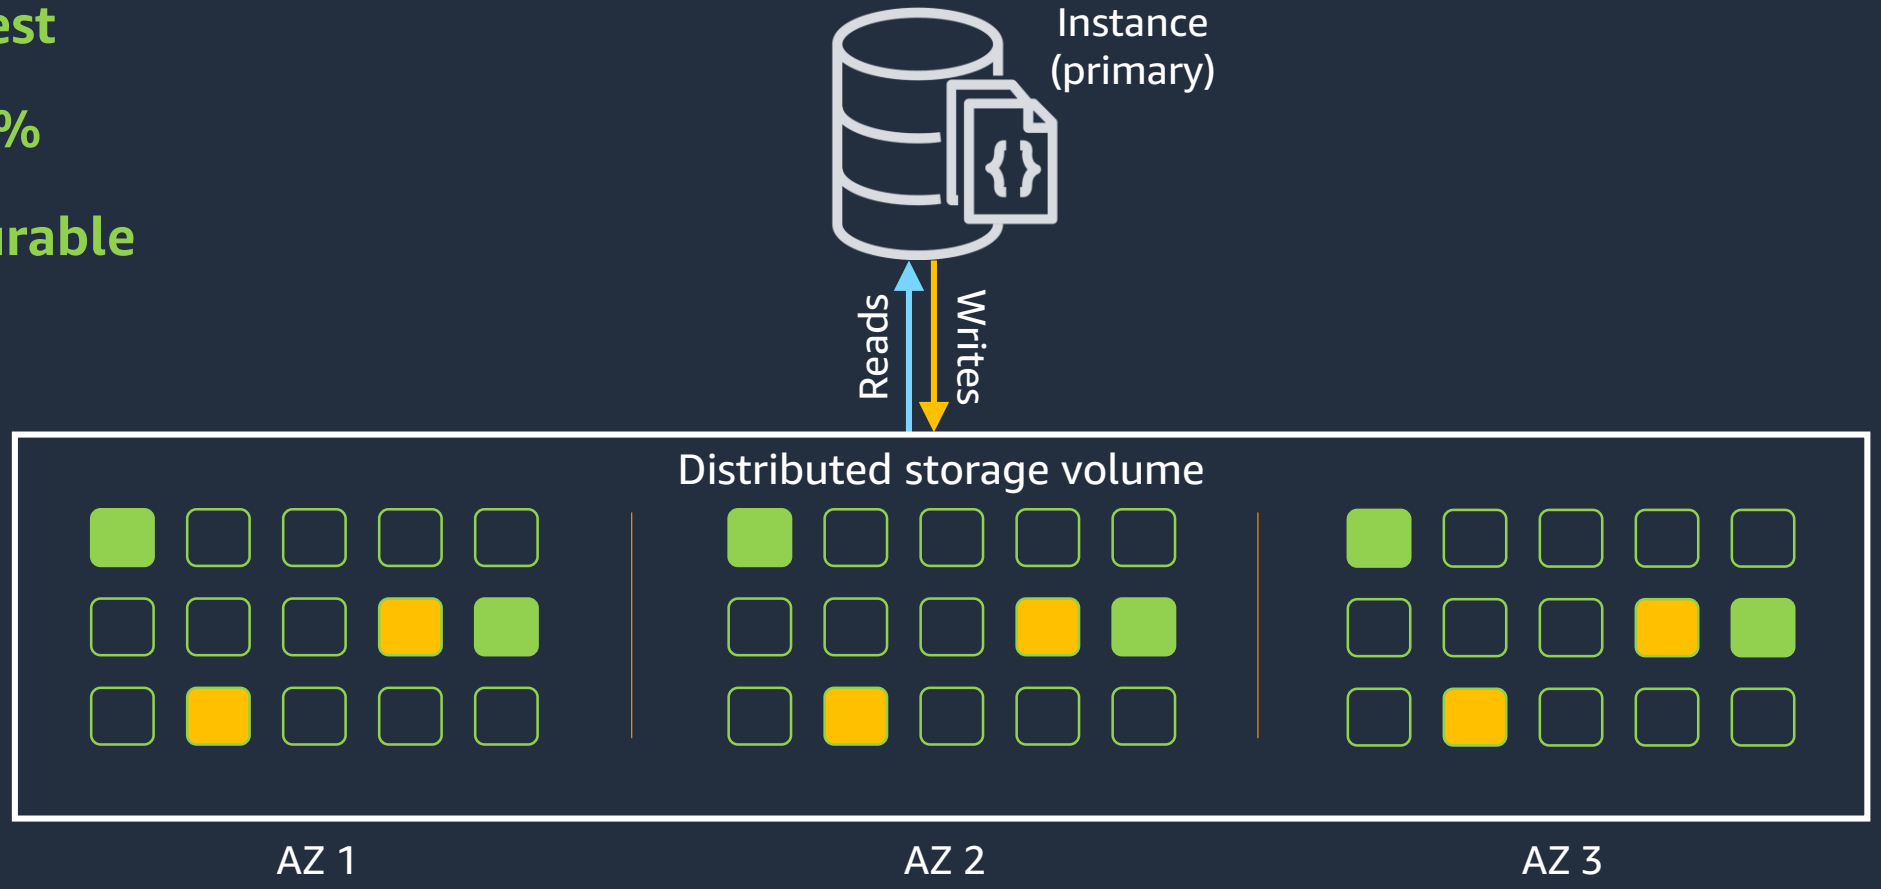

Challenge #5: Data Durability

Instances: 2

Environment: production

Availability goal: 99.9%

Durability: highly durable

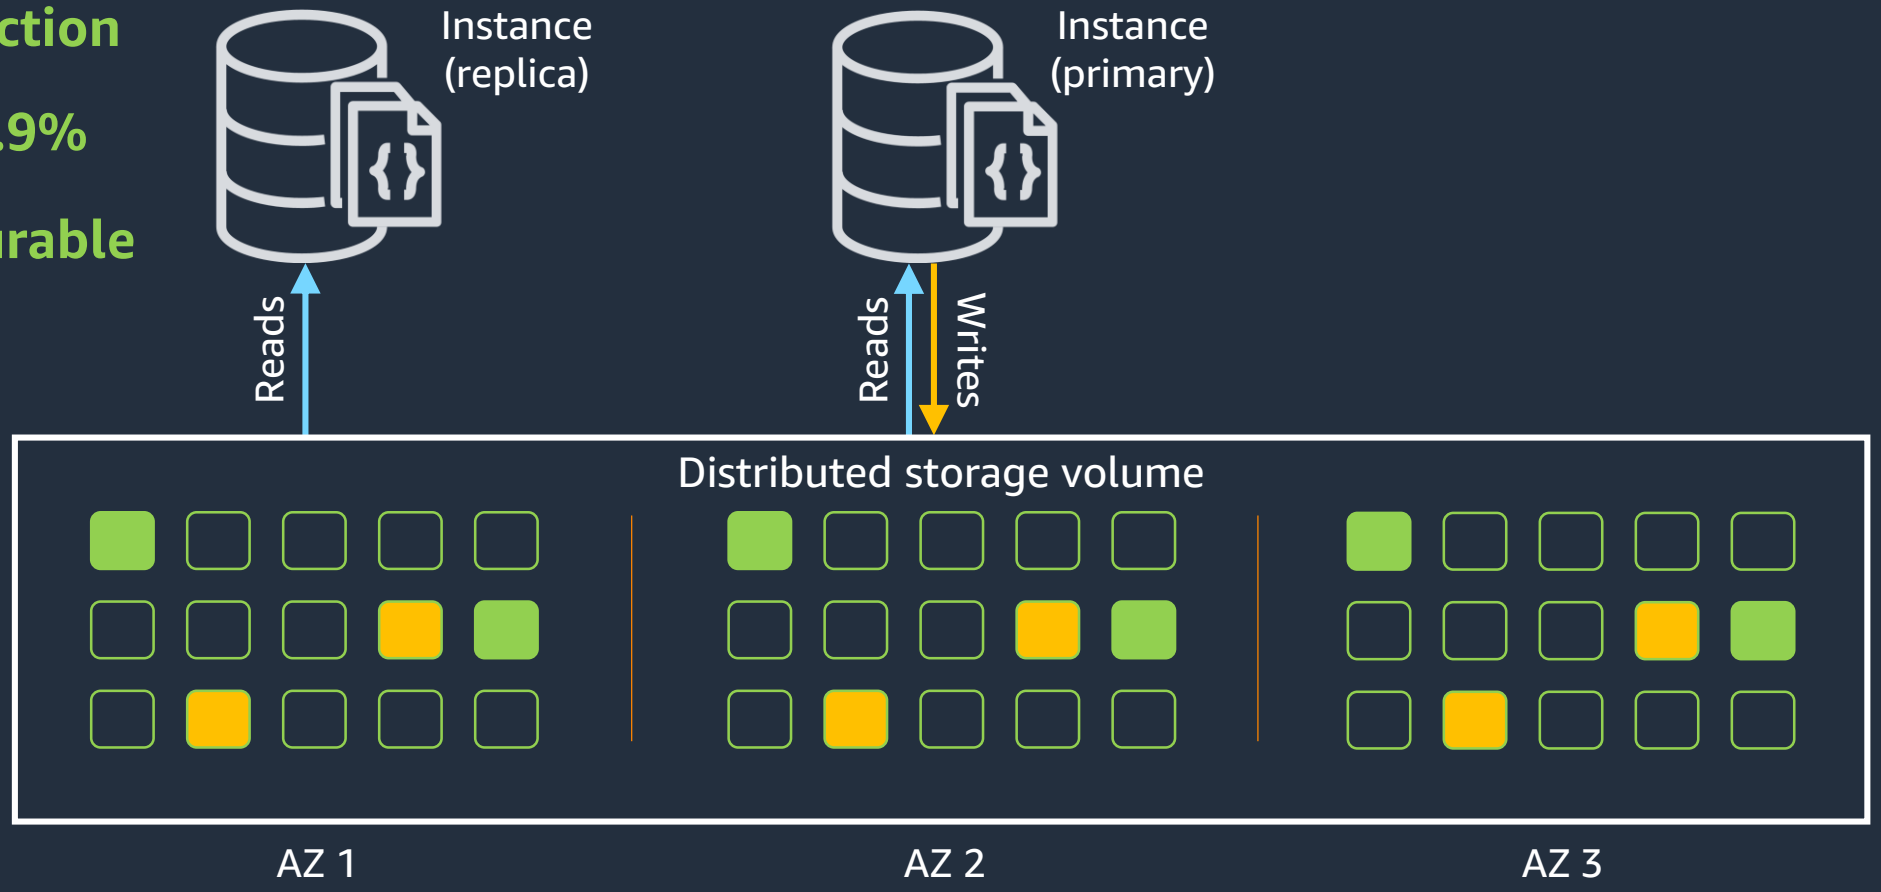

Challenge #5: Data Durability

Instances: 3

Environment: production

Availability goal: 99.99%

Durability: highly durable

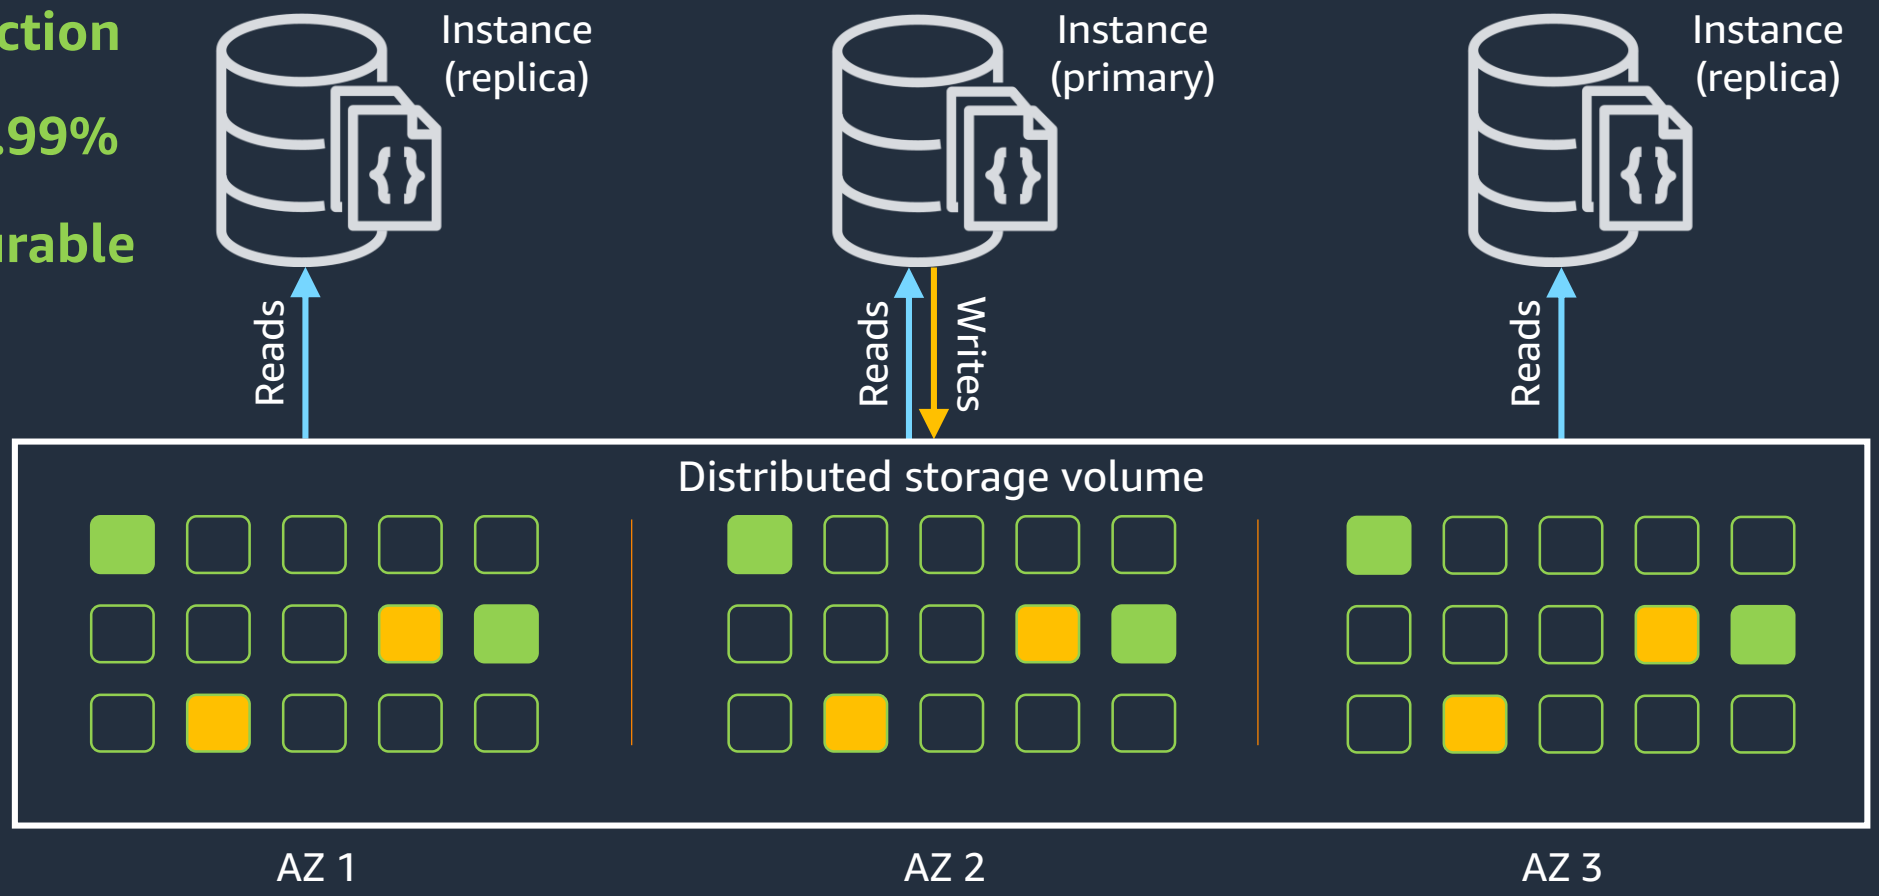

Challenge #5: Data Durability

Instances: 4

Environment: production

Availability goal: 99.99%

Durability: highly durable

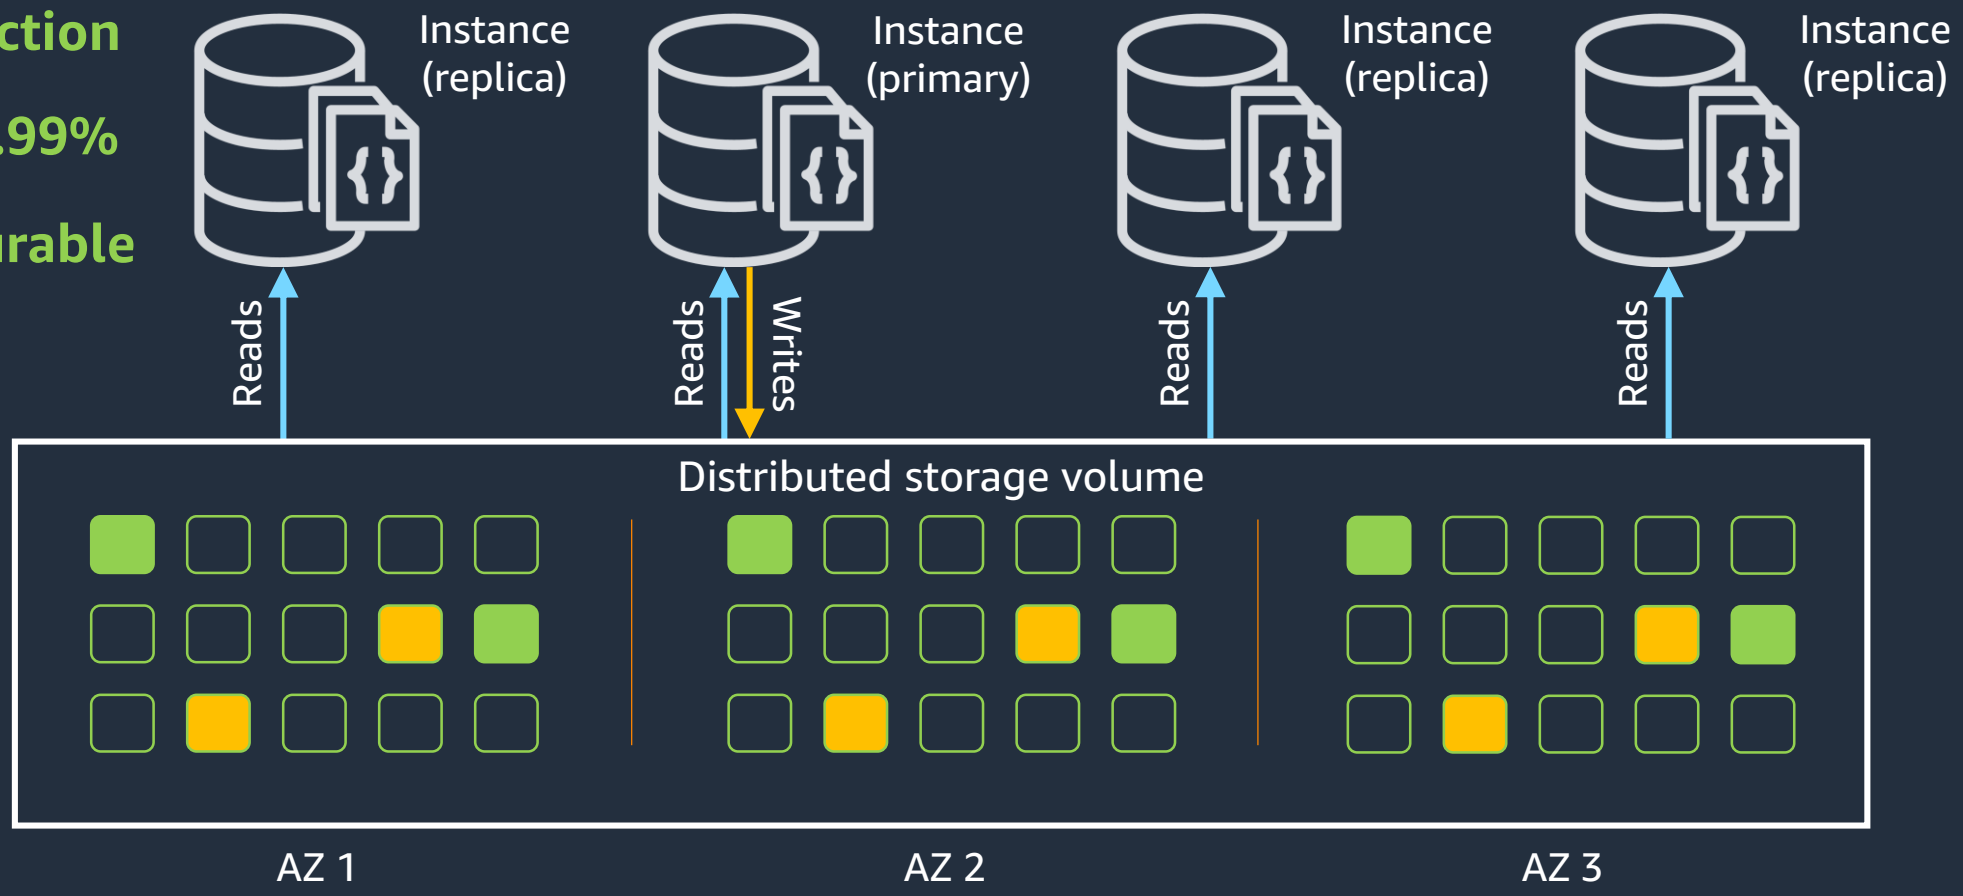

Service overview

Fully Managed

Automatic failure recover and failover

Replicas are automatically promoted to primary

Automatic patching

Up to date with the latest patches

Integrated with AWS services

CloudWatch, CloudTrail, CloudFormation, Secrets Manager, VPC, IAM, CLI

Pay-as-you-go pricing; enterprise grade

Per-second instance billing, no long-term commitments

freshop
Easy Online Groceries

"Our engineering teams now spend less time on operations like backup scripts, scale testing, and managing high availability and instead are able to focus on developing new capabilities for our customers."

MongoDB compatible

MongoDB 4.0

Compatible with MongoDB Community Edition 4.0

Same drivers, tools

Use the same MongoDB drivers and tools with Amazon DocumentDB

Replica sets

Read scaling is easy with automatic replica set configurations

ENKIDOO ai

"Getting started with DocumentDB was also simple and we migrated our application in a couple of days without needing to make any meaningful code changes. Everything just worked."

Scalable

Scale out in minutes

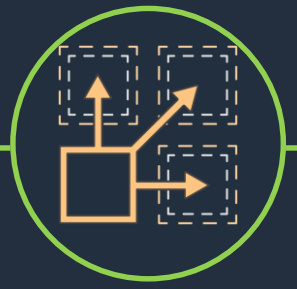

Scale to 15 read replicas

Scale up in minutes

Scale from 4 to 768 GiB of RAM

Autoscaling storage

Storage automatically grows from 10 GB to 64 TiB

Load balancing

Scale reads across replicas

HABBY

"With Amazon DocumentDB, our development team can scale, iterate, and upgrade games quickly, the marketing team can carry out high pertinence promotion activities, and our customer service team can troubleshoot problems from game players efficiently.

Security and compliance

Amazon VPC

Strict network isolation with Amazon Virtual Private Cloud (VPC)

Encryption by default

Encryption at rest with AWS KMS and customer-managed AWS keys; encryption in transit with TLS

Safe defaults

Best practices are the defaults

Compliance and support

PCI DSS, ISO 9001, 27001, 27017, and 27018, SOC 1, 2 & 3, HIPAA

sandstone TECHNOLOGY

“Integration with AWS KMS for data encryption at rest, transparent failover with read-replicas and the ability to take incremental database snapshots significantly lowers the operational burden for our team running the platform.”

aws

Backup and recovery

Automatic backups

Automatic, incremental, and continuous backups

No performance impact

Backups do not affect database performance

35 days of PITR

Point-in-time recovery (PITR) for up to 35 days

Archive snapshots

Keep snapshots for as long as you need

Iron

"Adopting Amazon DocumentDB is a game-changer because we offload management, security, and backup of our MongoDB databases to AWS. With Amazon DocumentDB, we can add or scale instances in minutes, regardless of data size. Further, we get automatic backups and point-in-time restore capabilities, which far exceed other managed DB services at less cost."

Innovation in Amazon DocumentDB

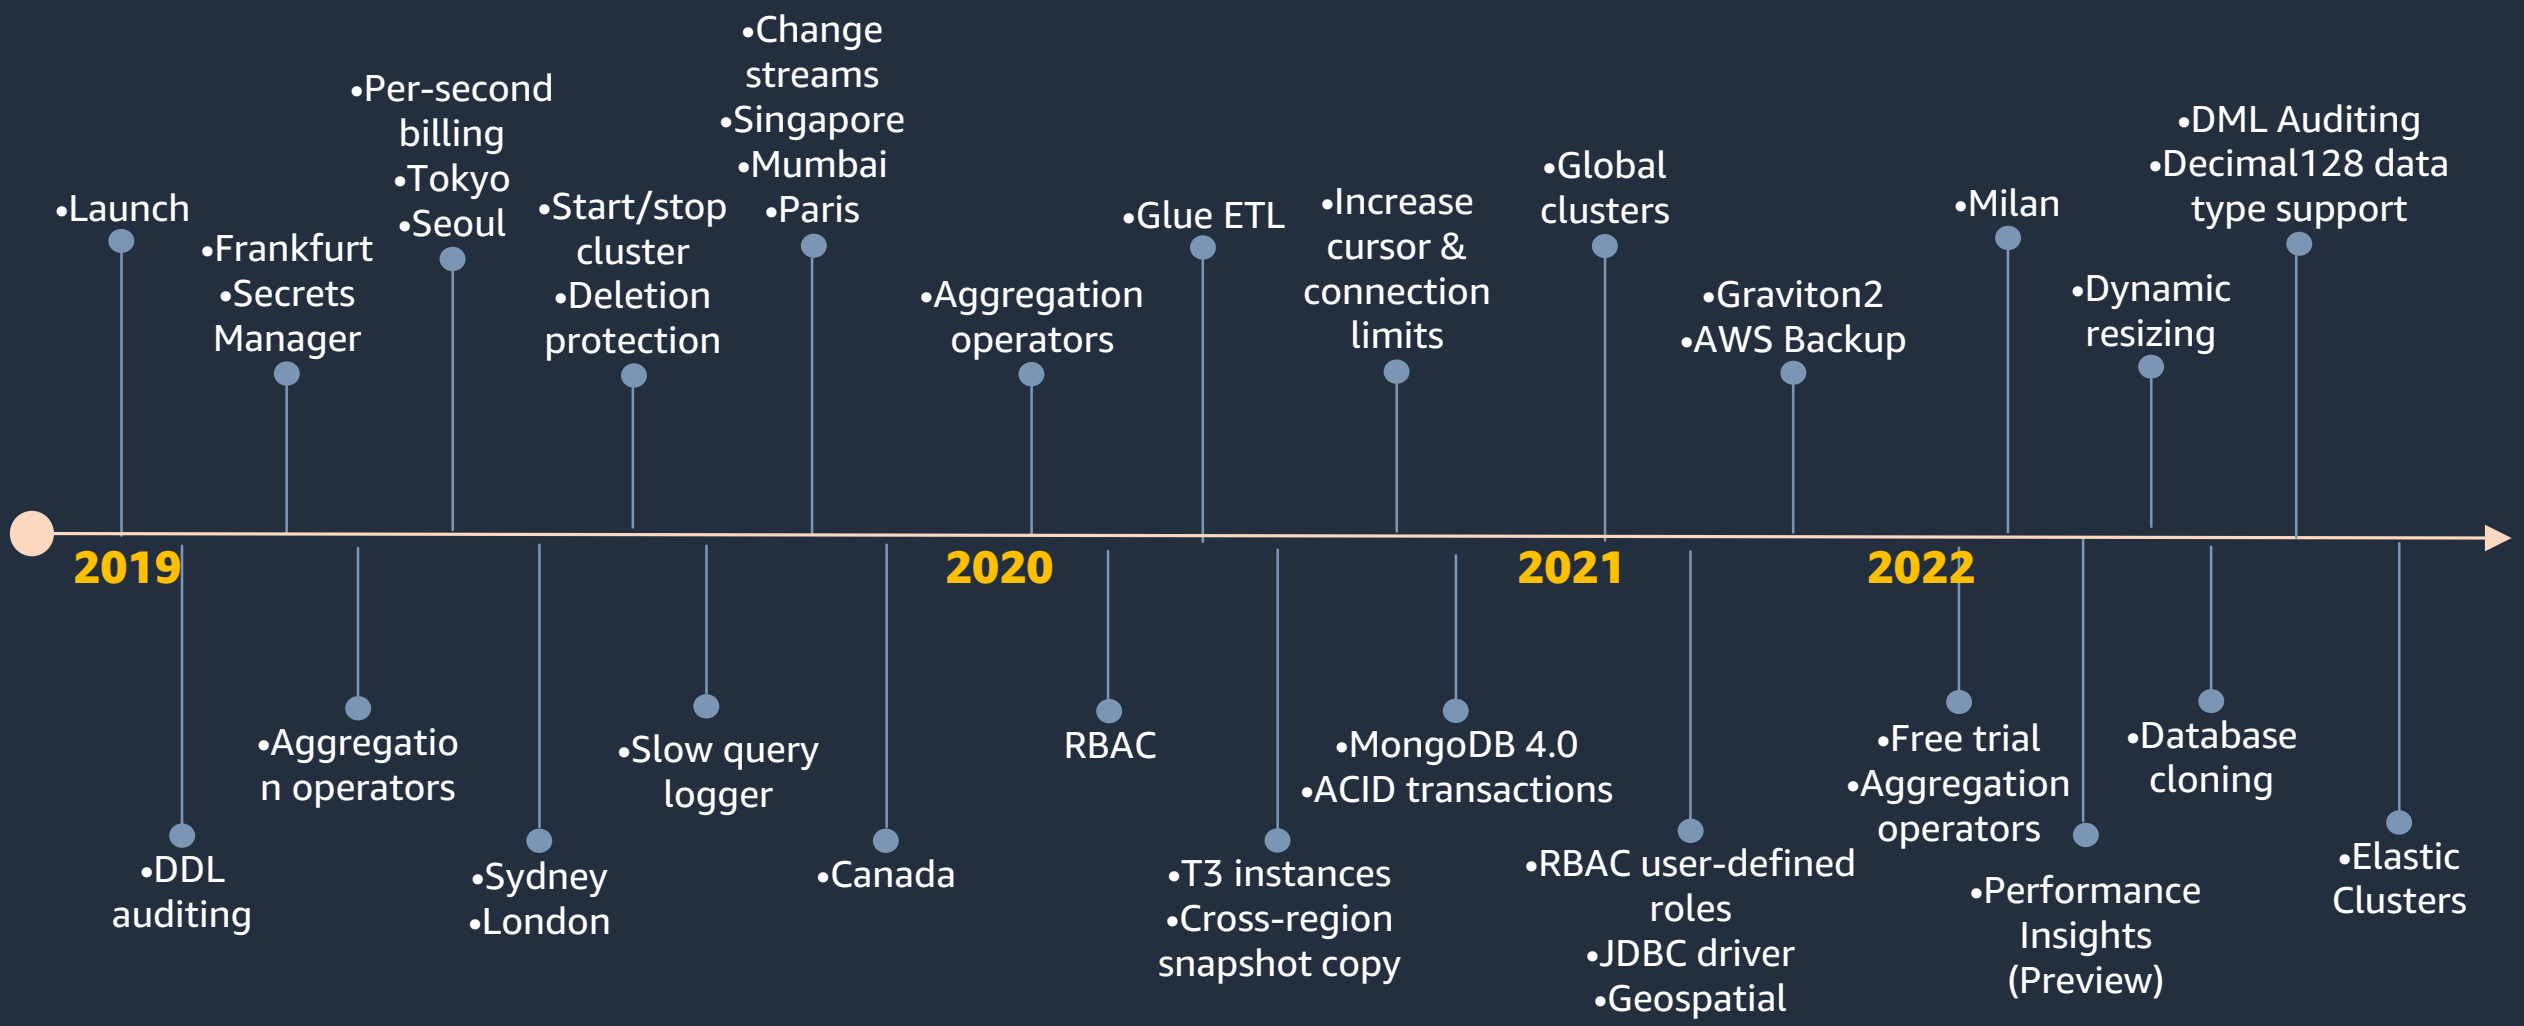

Amazon DocumentDB Elastic Clusters

NEW!

New type of cluster (Elastic Cluster) for Amazon DocumentDB that supports workloads with millions of reads/writes per second and petabytes of storage capacity

Compatible with MongoDB APIs for sharding

Millions of reads and writes

Up to 300,000 connections

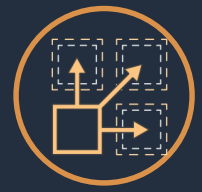

Easy and fast scaling

Elastic Cluster Architecture

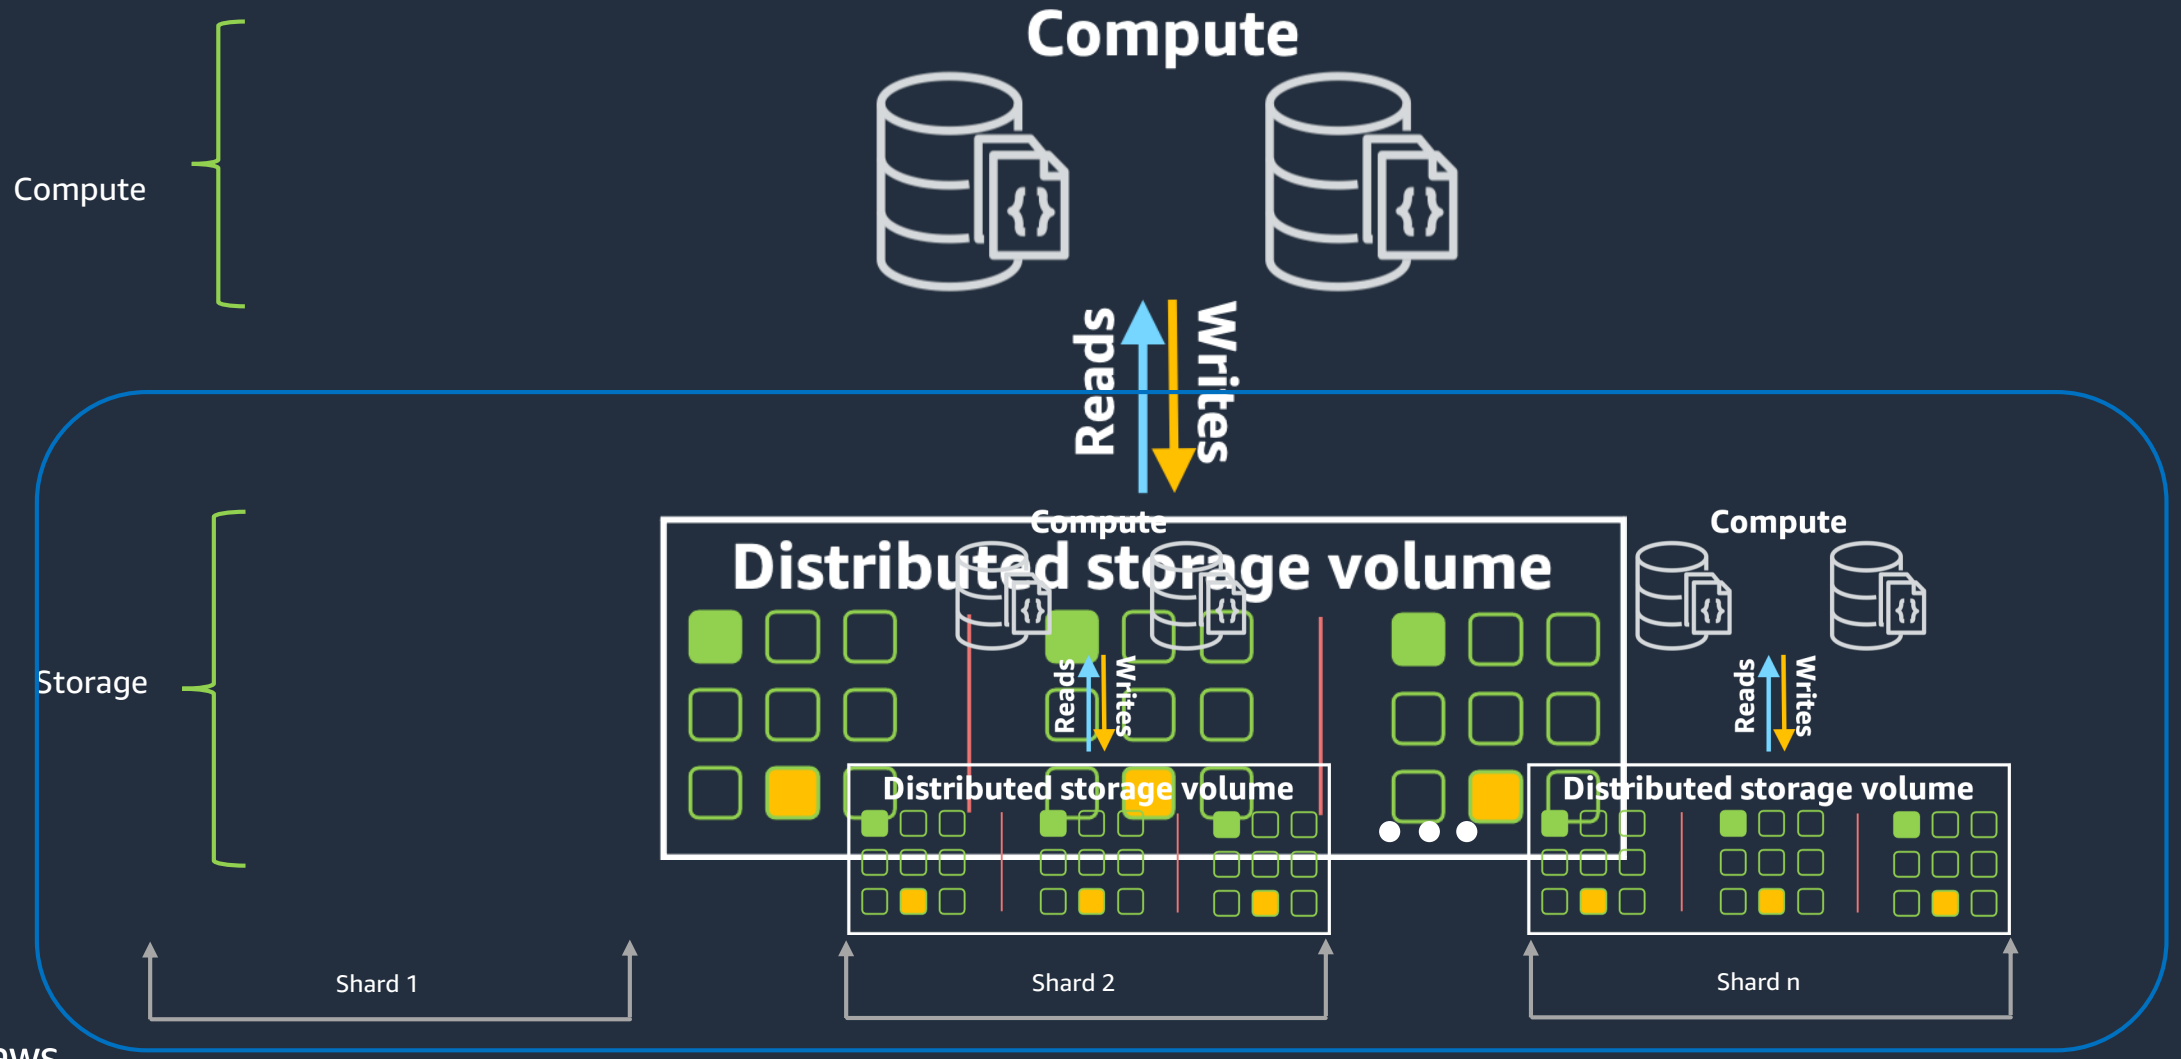

Elastic Cluster Architecture

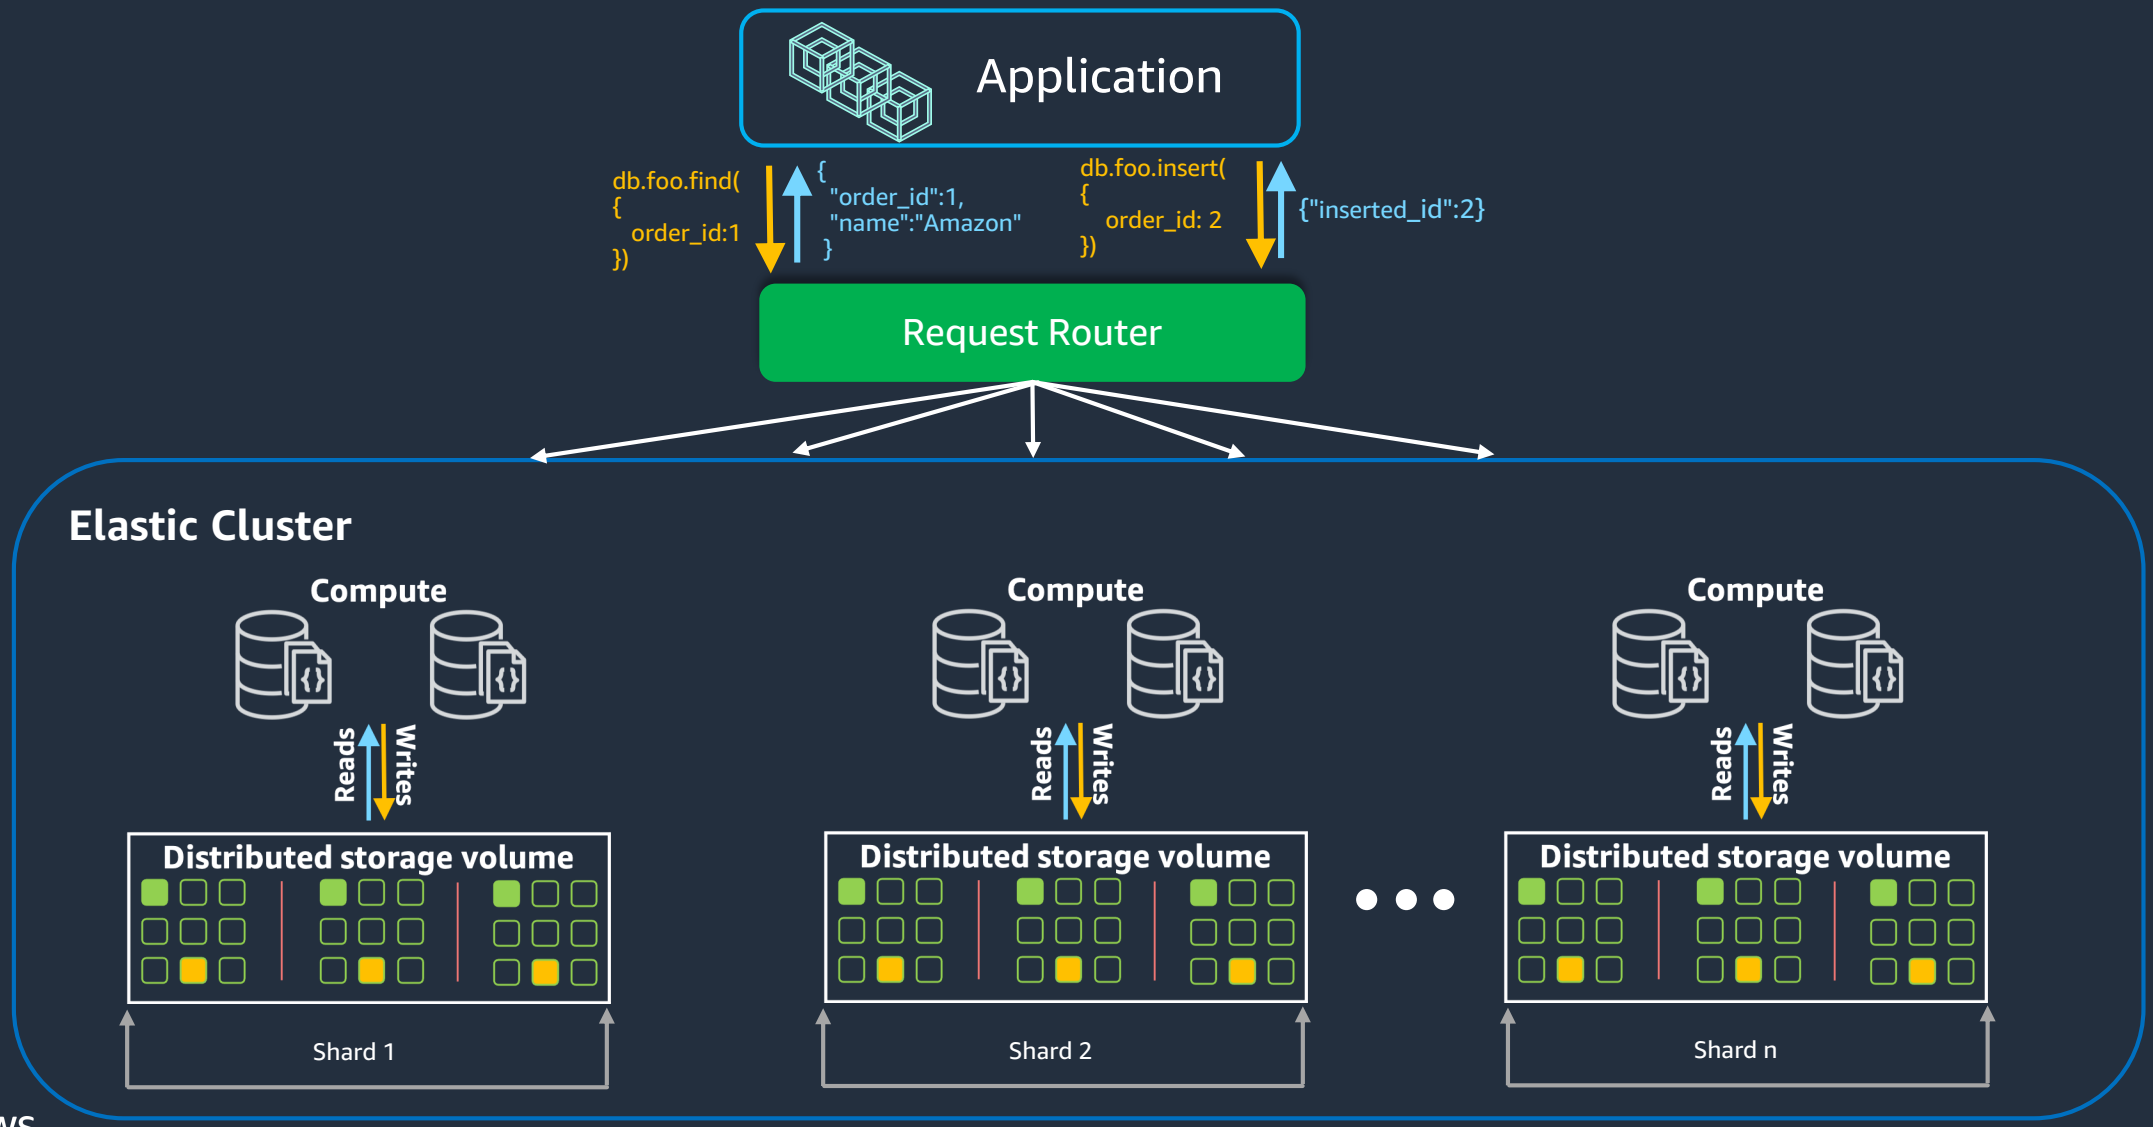

Pricing (us-east-1)

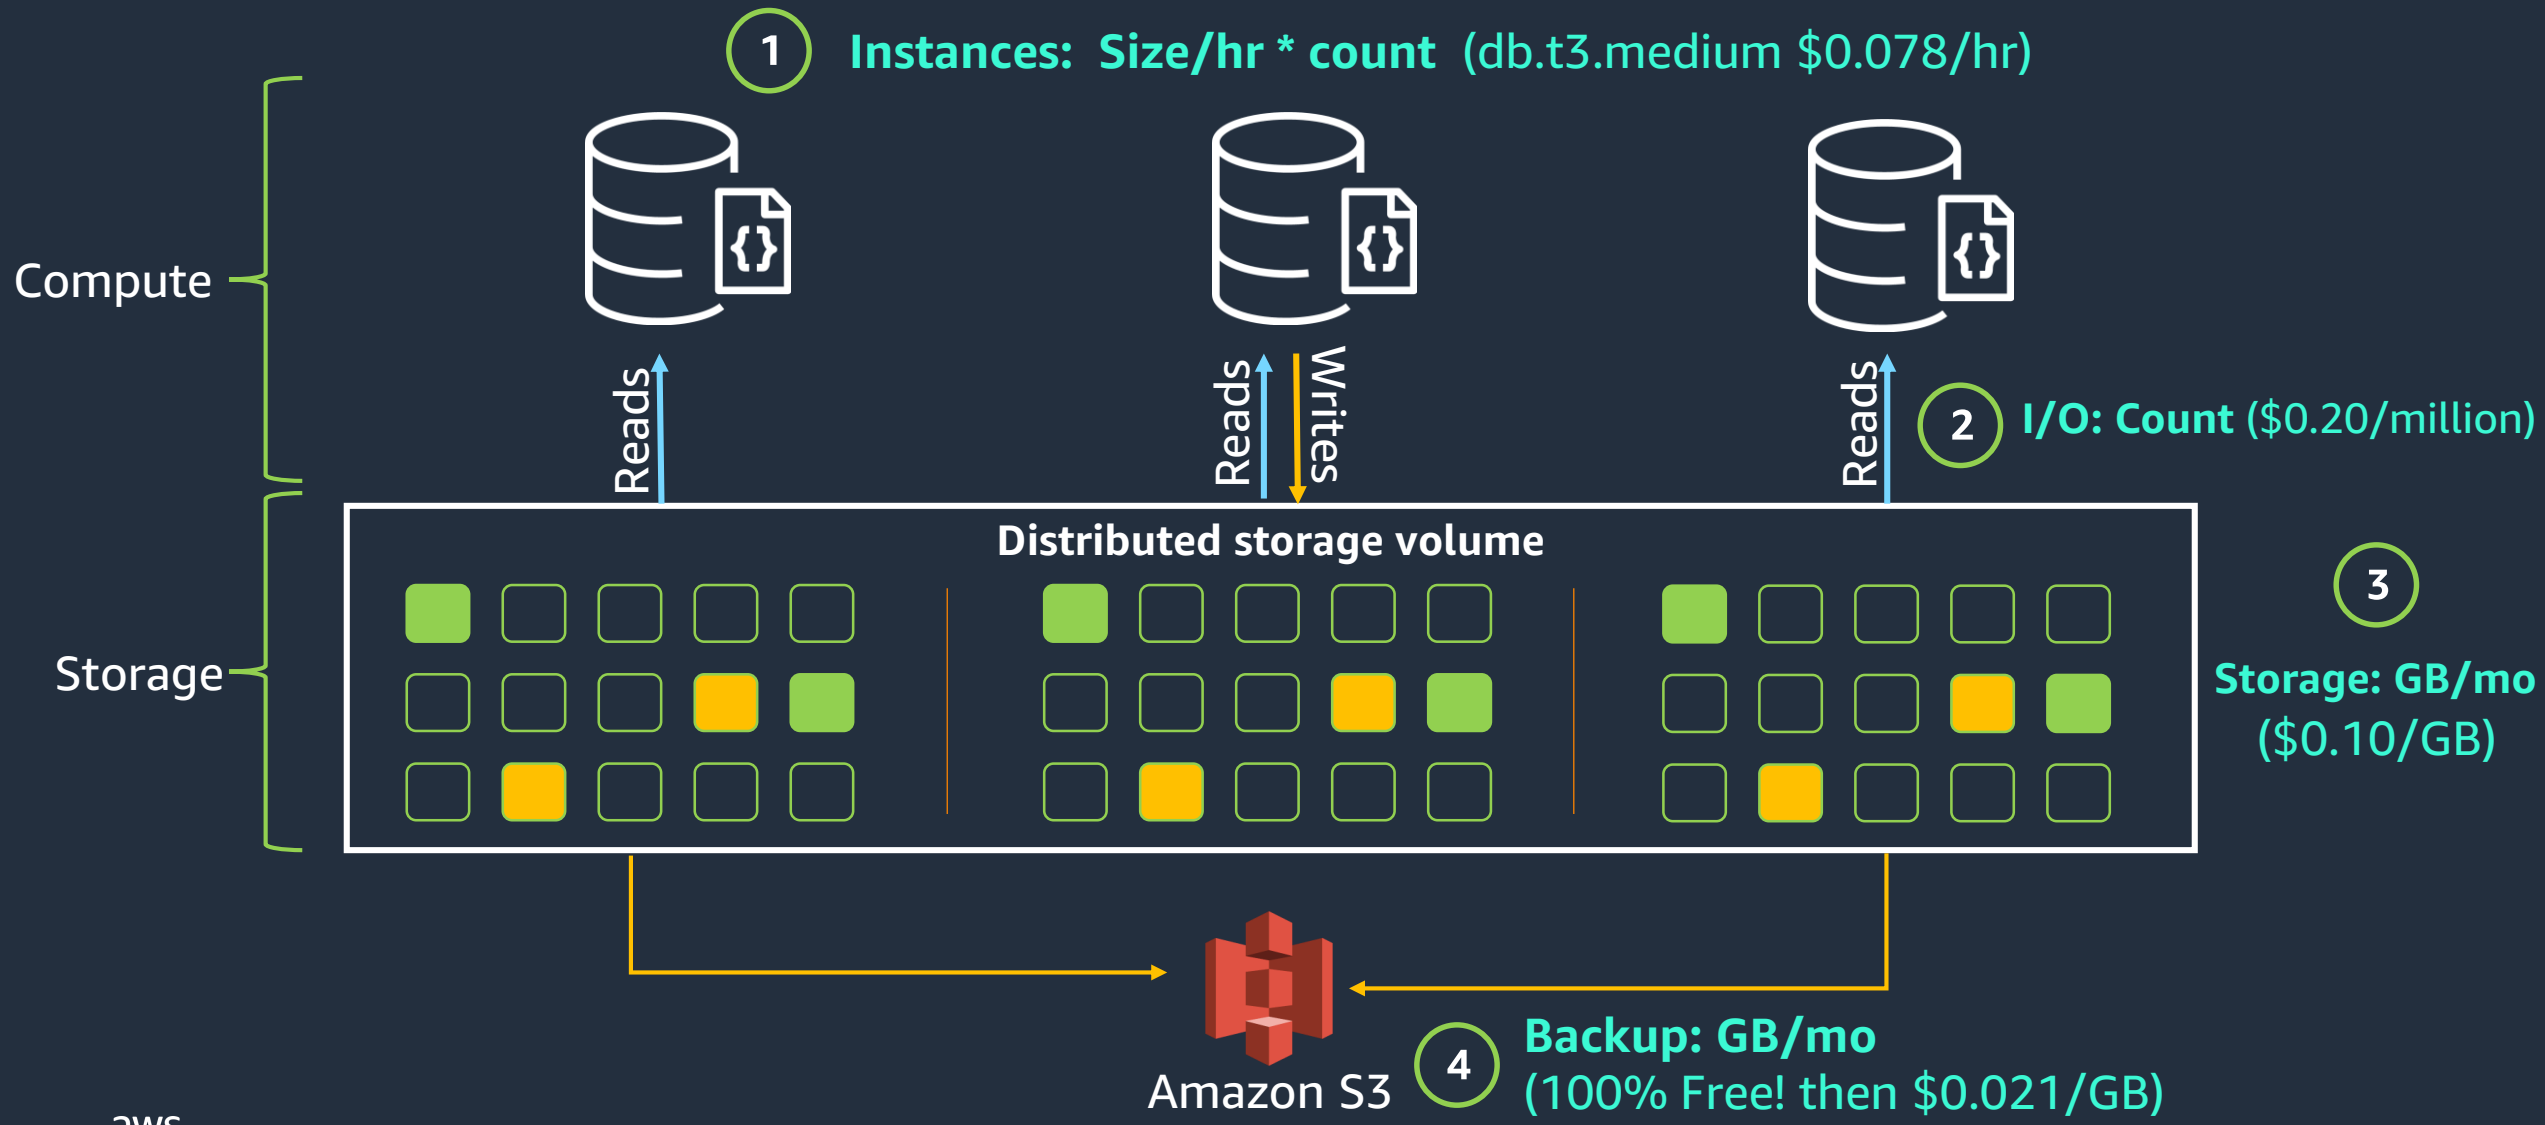

Migration Methods

Sources for both relational and document databases

Offline migration using MongoDB utilities:
Get started quickly, great for proof-of-concepts

Online migration using AWS DMS
Near-zero downtime migration

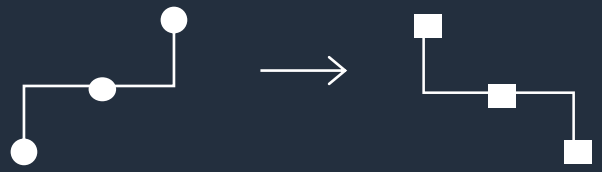

Hybrid migration leverages both solutions

https://docs.aws.amazon.com/documentdb/latest/developerguide/docdb-migration.html

All options support MongoDB on-premises and EC2, for both replica sets and sharded clusters

Amazon DocumentDB Programs/Investments

Program Name	Description	Duration	Cost	Outcome
Cost Analysis	DocumentDB sizing questionnaire using your workload metrics to generate a cluster sizing estimate	Self-serve	None	DocumentDB sizing and pricing estimate
Compatibility Assessment	The DocumentDB compatibility tool will examine log files from MongoDB to determine if your applications use operators that are not supported in DocumentDB. This tool will report use of unsupported APIs, and output the unsupported log lines into a file for future use.	Self-serve	None	MongoDB -> DocumentDB compatibility report
Immersion Day	DocumentDB Immersion Days provide customers with modular content and hands-on labs to learn about DocumentDB use cases, architecture, best practices, migrations, security, monitoring, and more.	1-2 days	None	Enable customers to build a DocumentDB POC, or implement/migrate workload
Well Architected Lens	WAL for DocumentDB assesses customer workloads with a focus on optimizing performance, reliability, security, cost and operational excellence.	Half day	None	Well Architected review report with recommendations to optimize workload.
Springboard	Customized migration game plan that begins with decision makers (sponsor) and diverse stakeholders. Includes 1. current state operations & business commitments 2. specific architectural recommendations on AWS 3. proposed milestones 4. follow up action items	Half day	None	Engagement readout that serves as execution plan.
Data Labs	DocumentDB Data Labs provide technical resources to help customer build tangible deliverables that accelerate data modernization initiatives. Customers who've built DocumentDB POCs are given prescriptive architectural guidance, best practices, and technical roadblock removal by AWS Data Lab engineers and DocumentDB experts. Customers leave the lab with a well-architected prototype, a path to production, and greater knowledge of DocumentDB.	4 days	None	Well architected DocumentDB prototype with path to production
Professional Services	AWS Pro Serve provides hands-on support to review customers' database environments and challenges, requirements, desired outcomes, and success metrics for migrating workloads to DocumentDB. Suitable use cases are identified, and customer gets experience implementing a demo of their solution using sample data. A documented roadmap proposal with next steps to implement migrations to DocumentDB is provided by end of the engagement.	Varies	SOW based	Hands-on support to guide customer from POC->assessment->migration

Thank you!