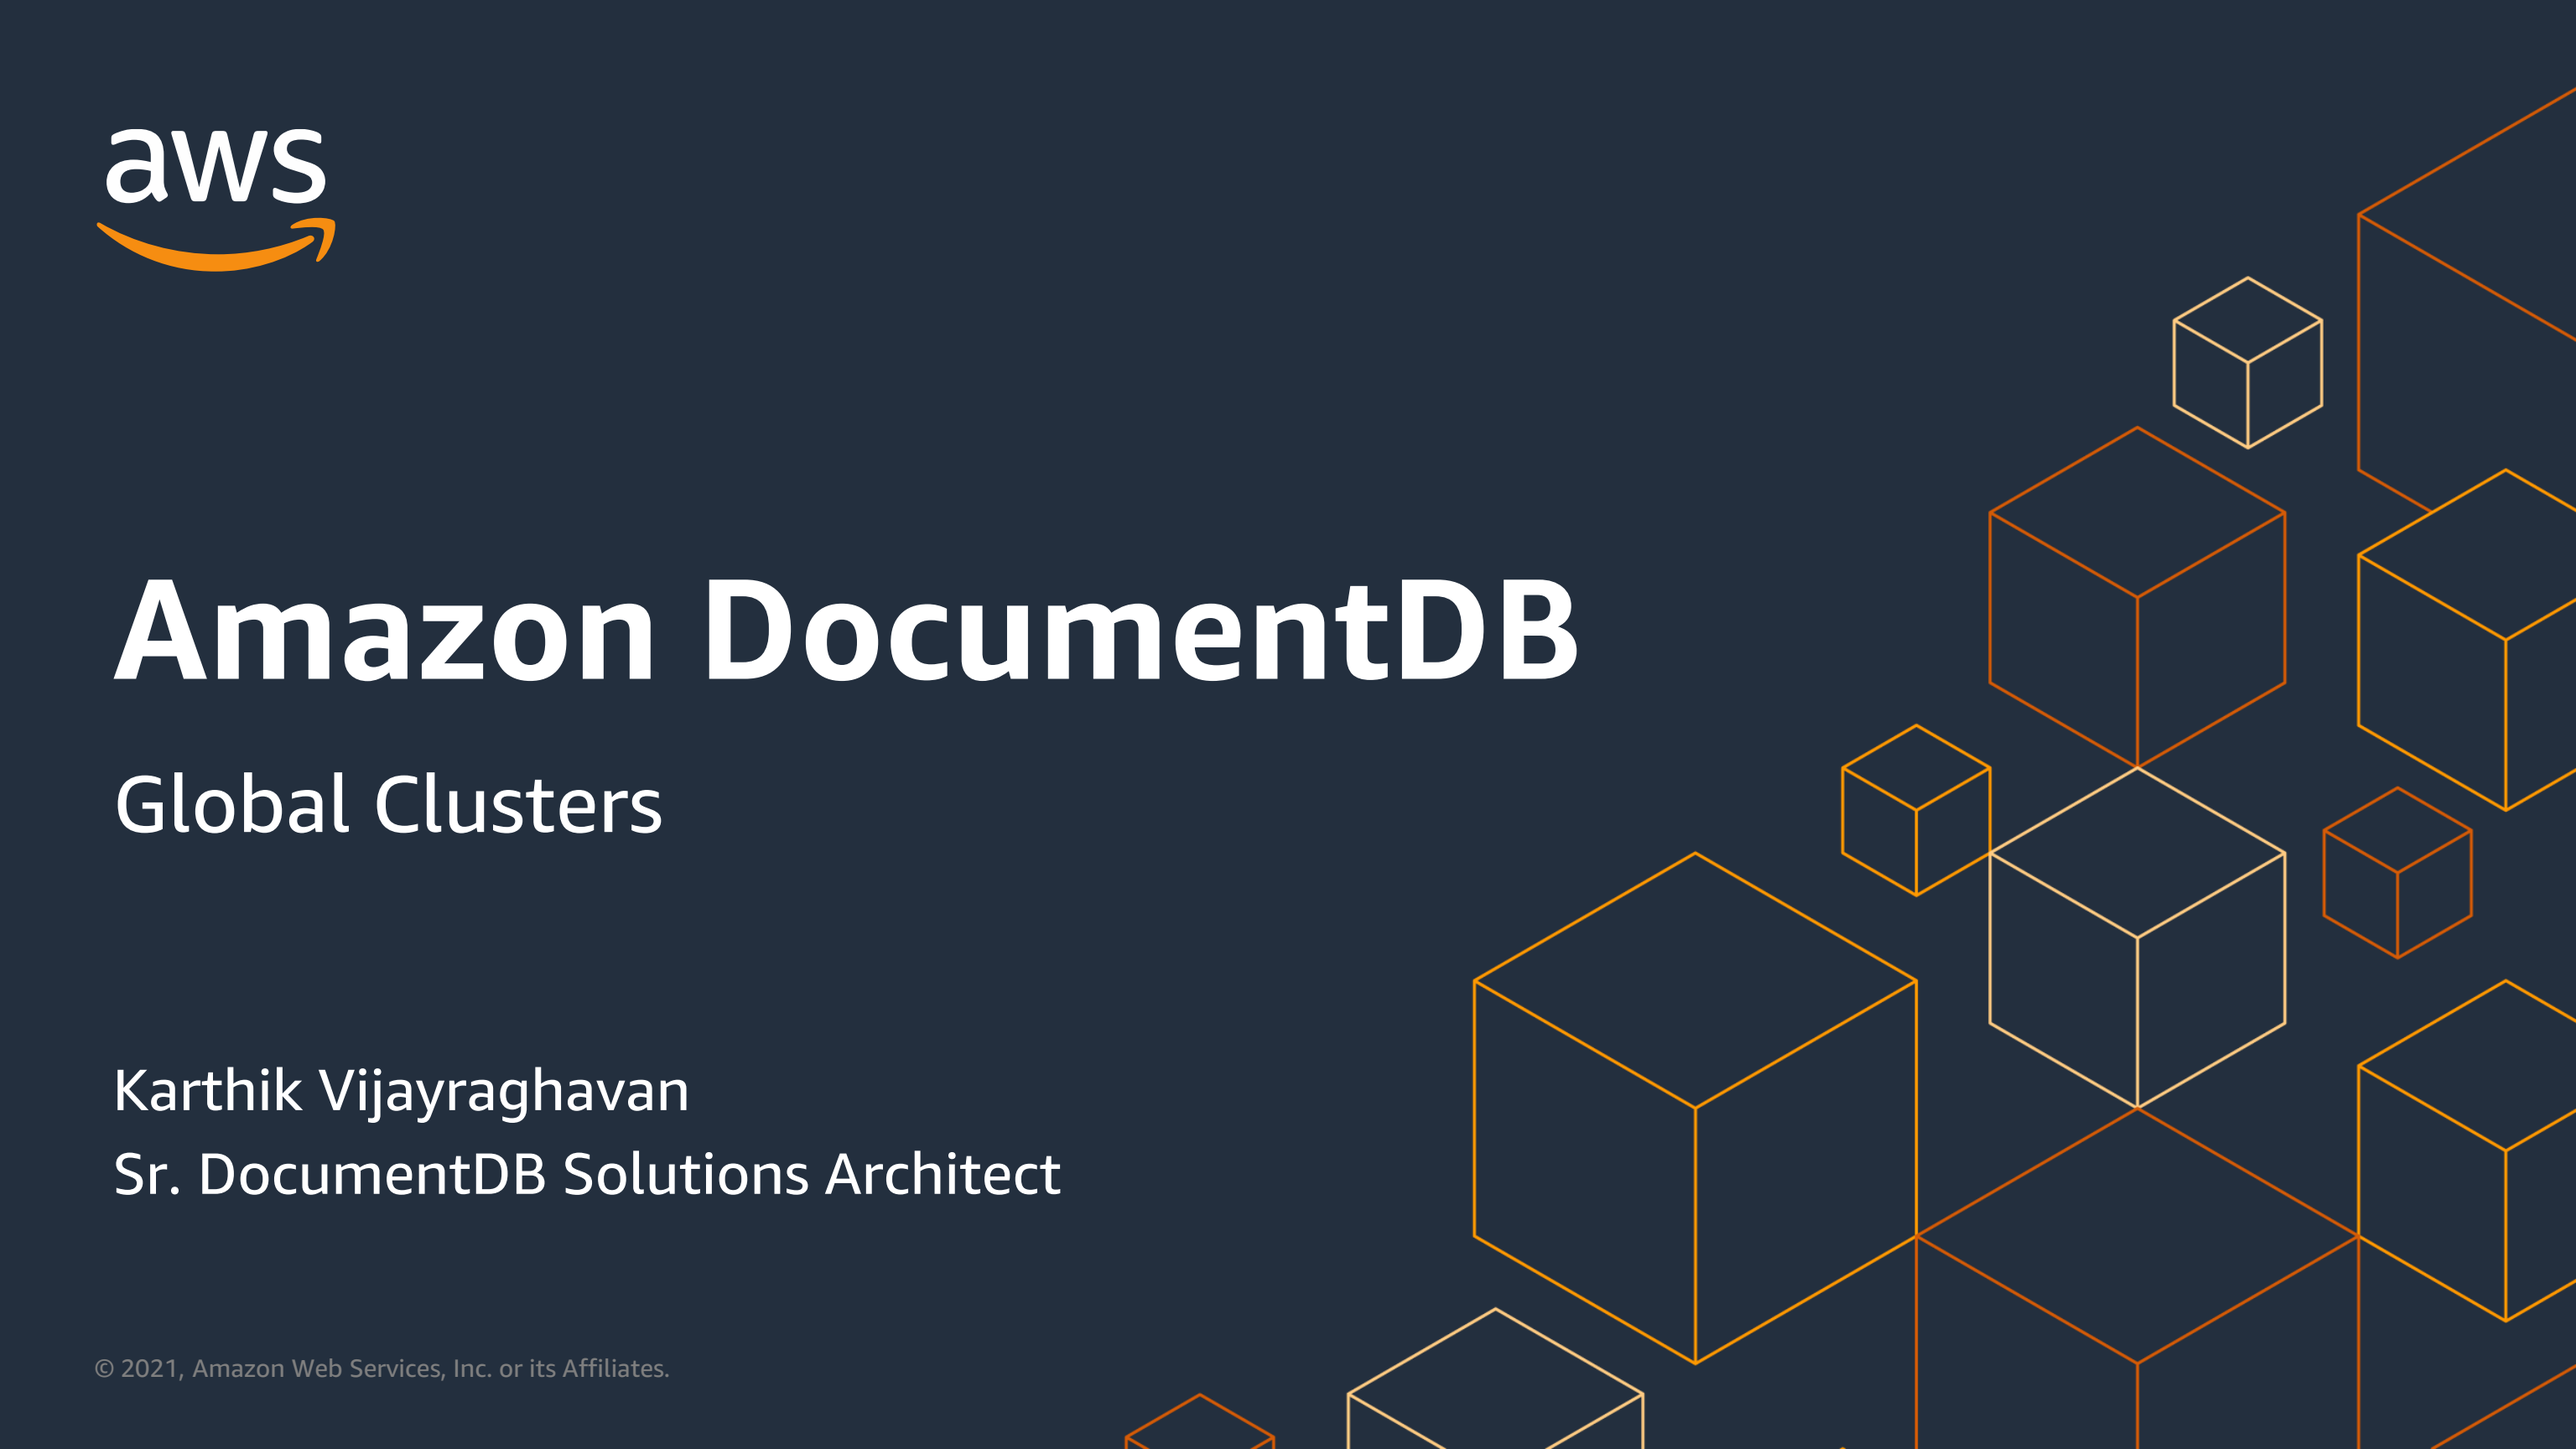

Amazon DocumentDB

Global Clusters

Karthik Vijayraghavan
Sr. DocumentDB Solutions Architect

Table of contents

aws

What is Amazon DocumentDB?

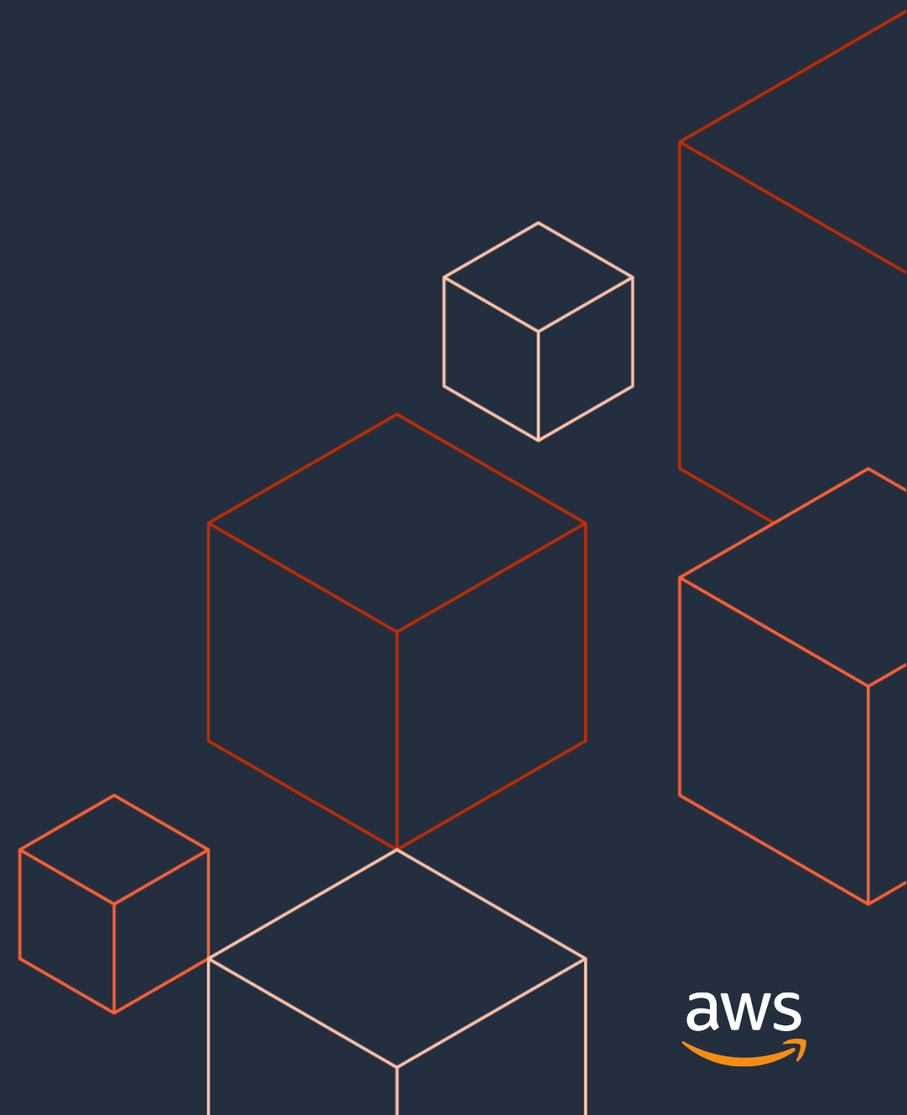

Amazon DocumentDB (with MongoDB compatibility)

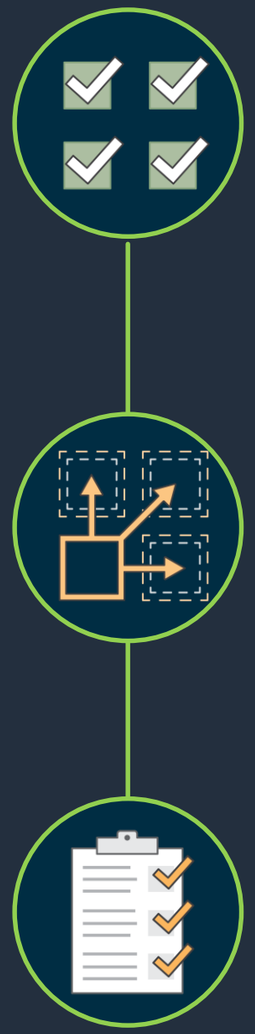

Fully managed and scalable document database service that supports MongoDB workloads

aws

Amazon DocumentDB (with MongoDB compatibility)

Fully managed

Built-in high availability

Backups enabled by default

Durable by default

Security best practices by default

Automatic patching

Monitoring and alerting

FINRA

“Because DocumentDB is a fully managed service, our databases are scalable, highly available, backed up, and encrypted *without any overhead from our engineering teams*”

aws

Amazon DocumentDB (with MongoDB compatibility)

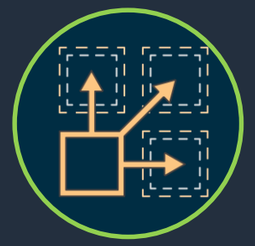

Scalable

Scale compute in minutes

Storage and IO autoscaling

Storage scales to 64TB

Scale out to 15 replicas for millions of reads

Iron.io

"With Amazon DocumentDB, we can add or *scale instances in minutes*, regardless of data size."

aws

Amazon DocumentDB (with MongoDB compatibility)

Applications, drivers, and tools can be used with Amazon DocumentDB with little or no change

Supports hundreds of APIs, operators, and stages

Continually working backward from customers to deliver the capabilities they need

MongoDB API compatible

freshworks

“We love that it’s compatible with MongoDB, so our applications *didn’t require code changes*, and we could easily spin up a DocumentDB cluster to test the MongoDB capabilities that we relied on.”

aws

Amazon DocumentDB Architecture

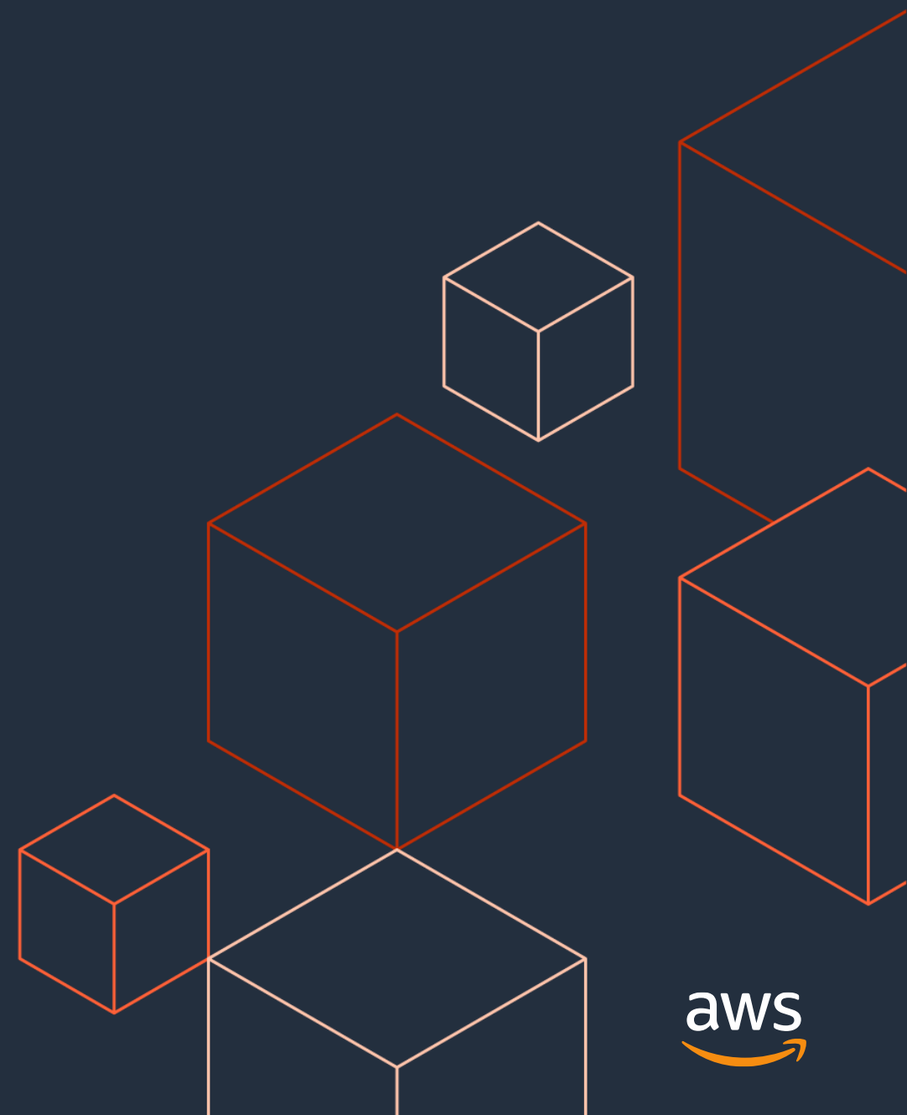

Amazon DocumentDB Architecture

Compatibility

Separation of storage / compute

Replication

Durability

Backup

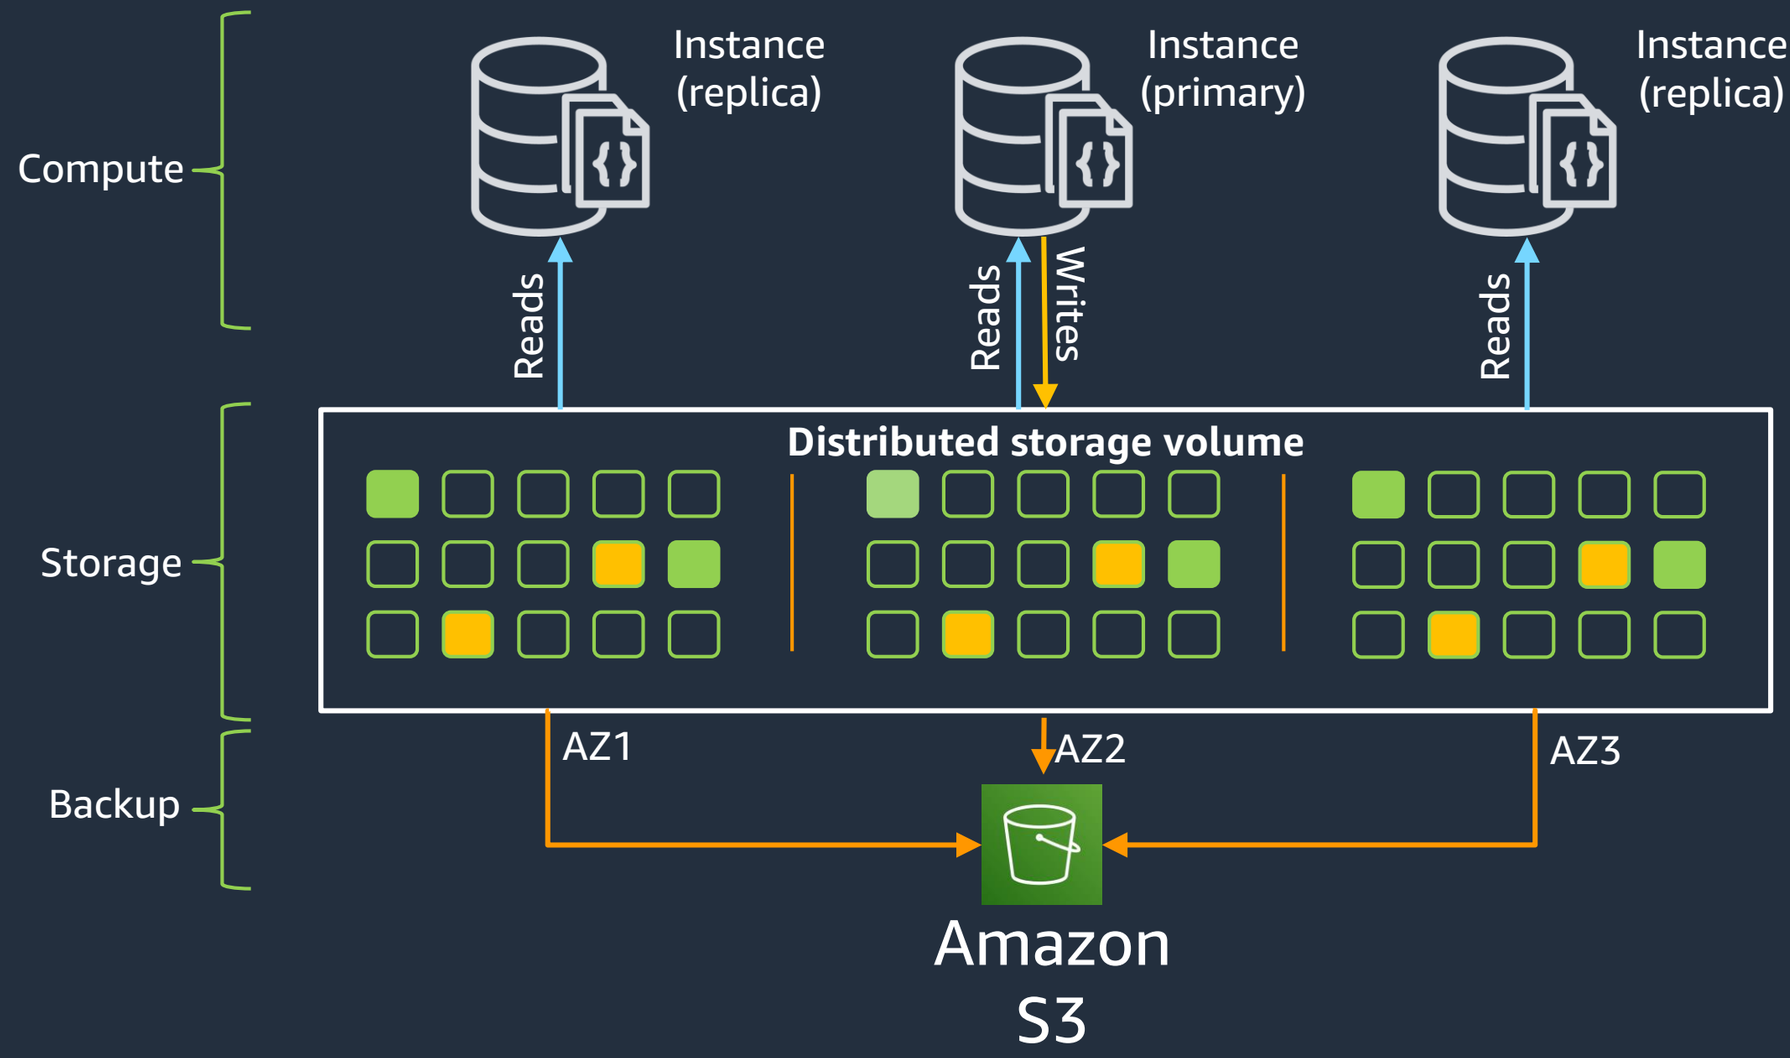

aws

Global Clusters

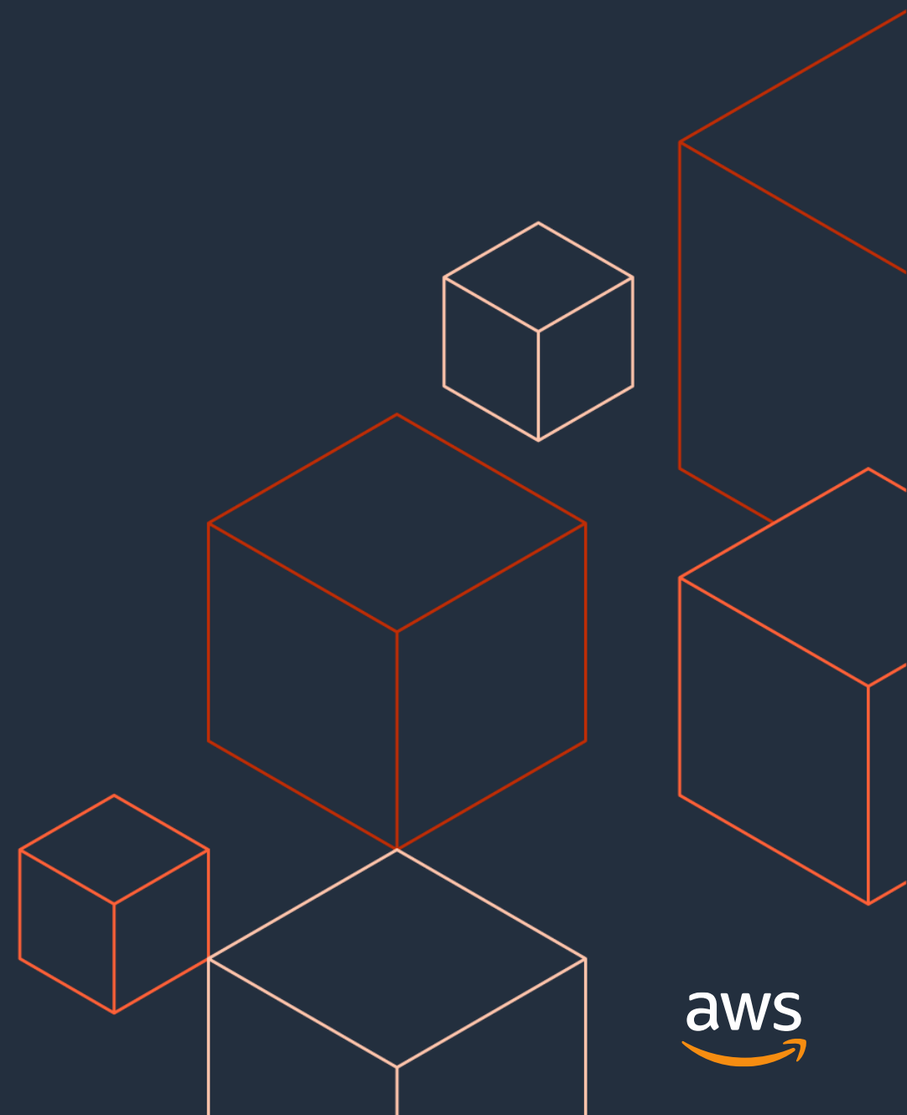

Amazon DocumentDB – Global clusters

Global replication: Up to 5 secondary regions

Low replica lag: Typically < 1 sec

Fast recovery: Typically < 1 min downtime

Compatibility: **V 4.0 and later**

Global reader instances: Up to 90

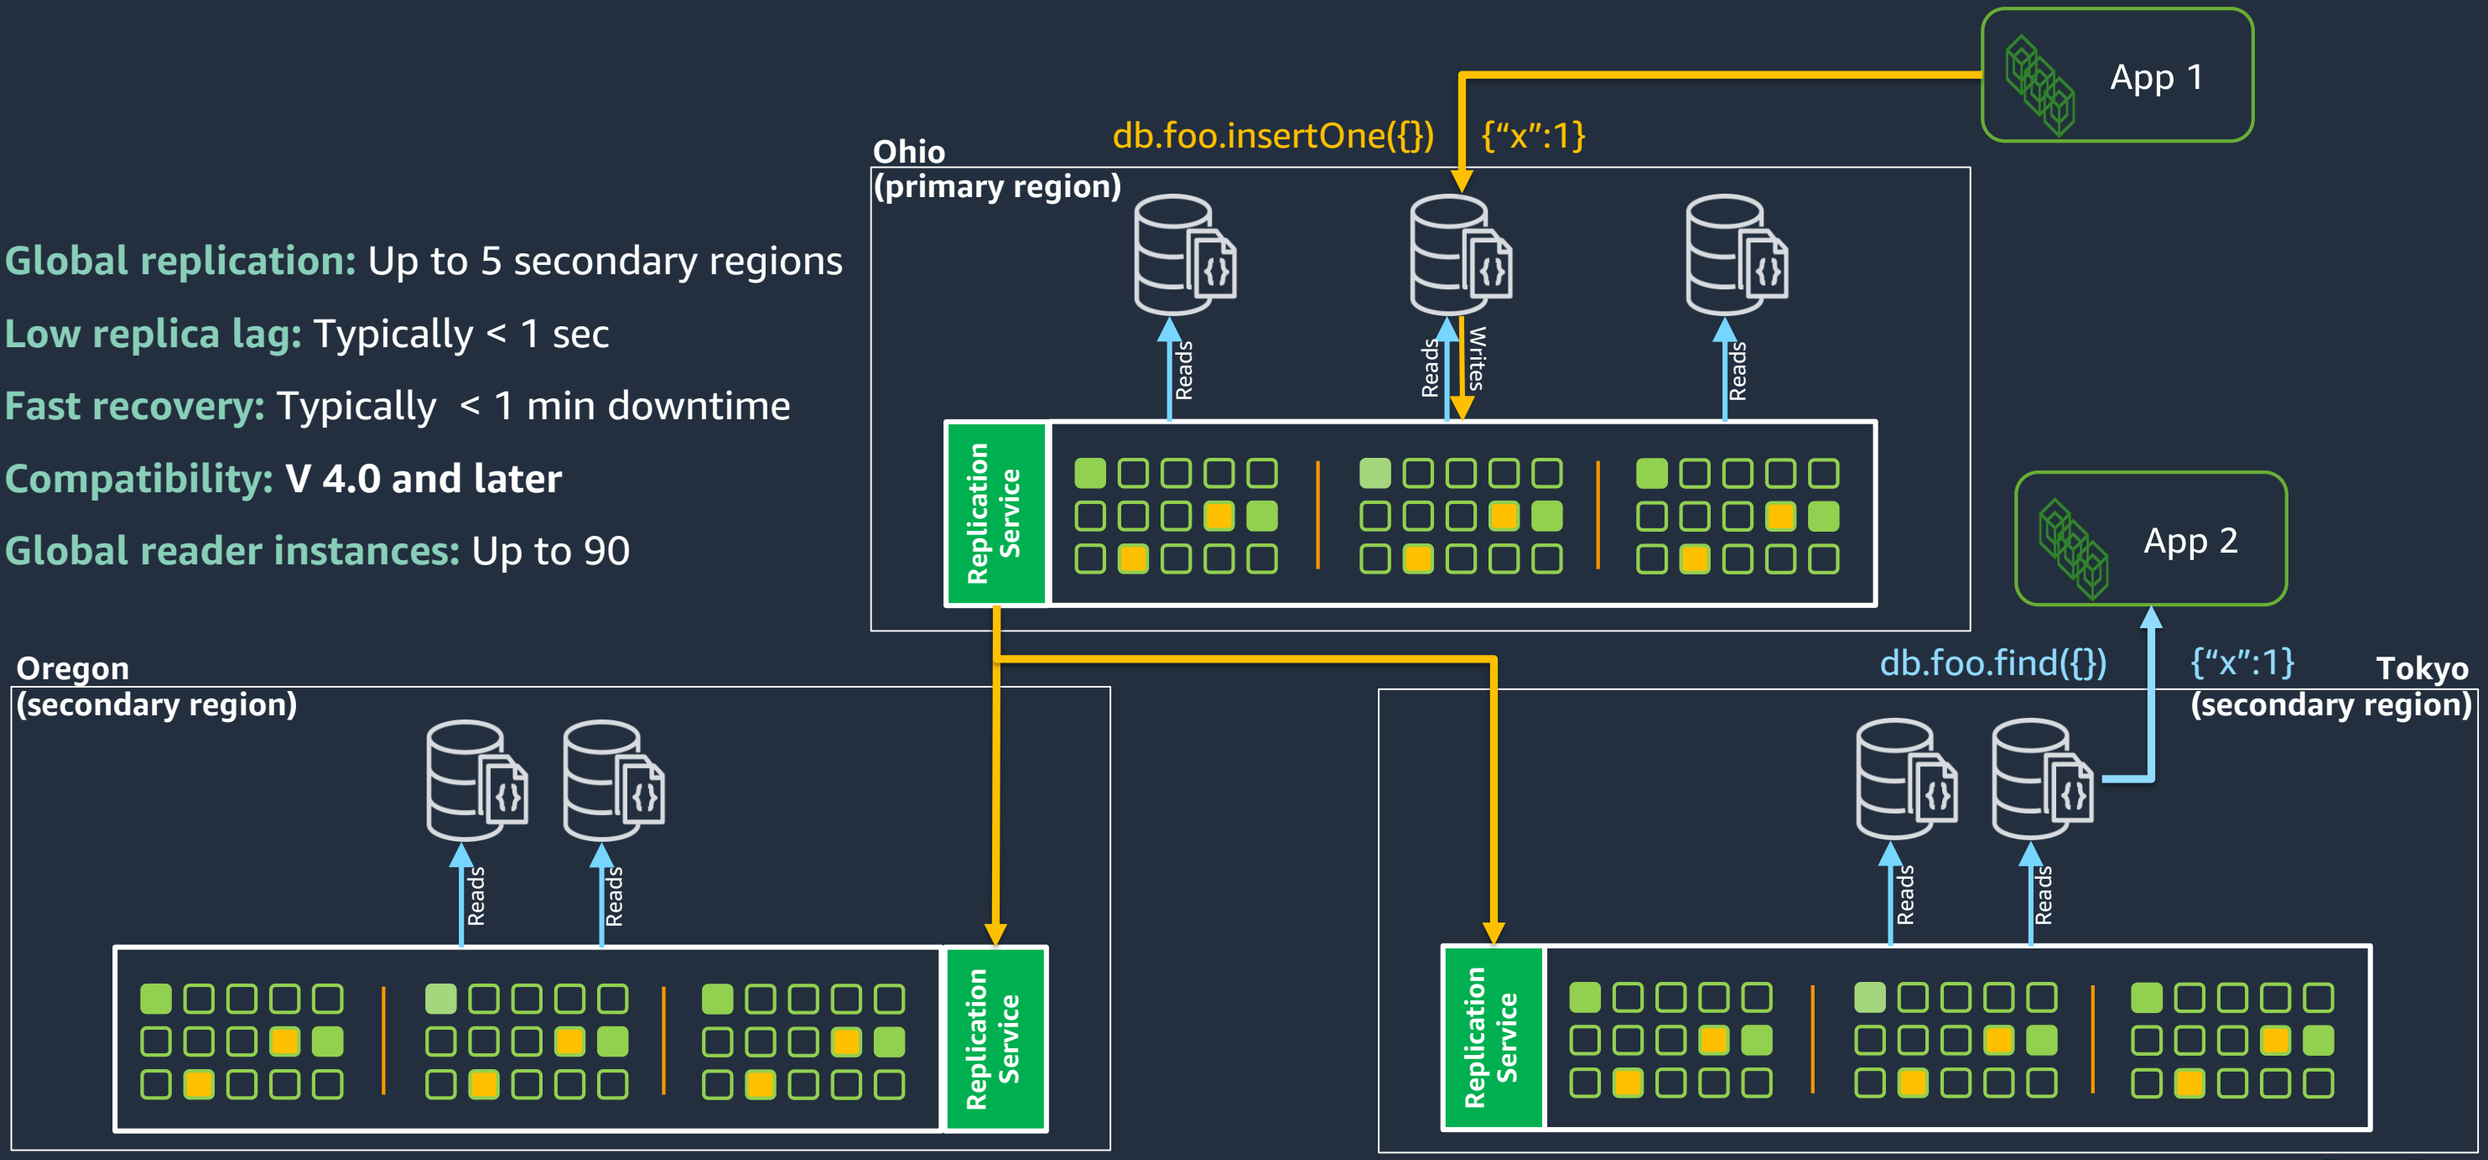

Demo – Create Global Clusters

aws

Global Clusters – Use cases

Disaster recovery:
Promote secondary clusters to primary for faster recovery in the event of regional failures

Data locality:
Bring data closer to users in different regions to enable faster reads for globally distributed applications

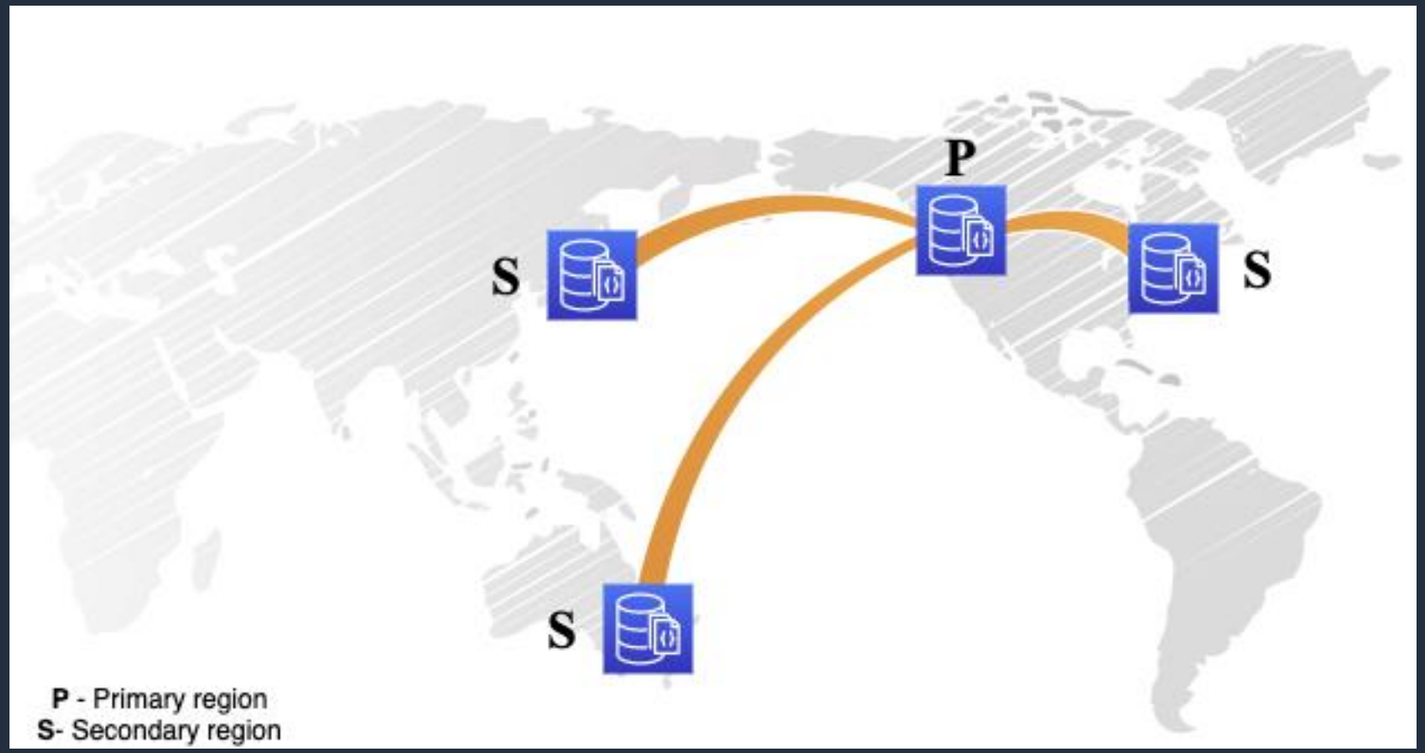

aws

Disaster recovery - Fast global failover to secondary Regions

Low replication lag across Regions

Low recovery time objective

Fast cross-Region migrations

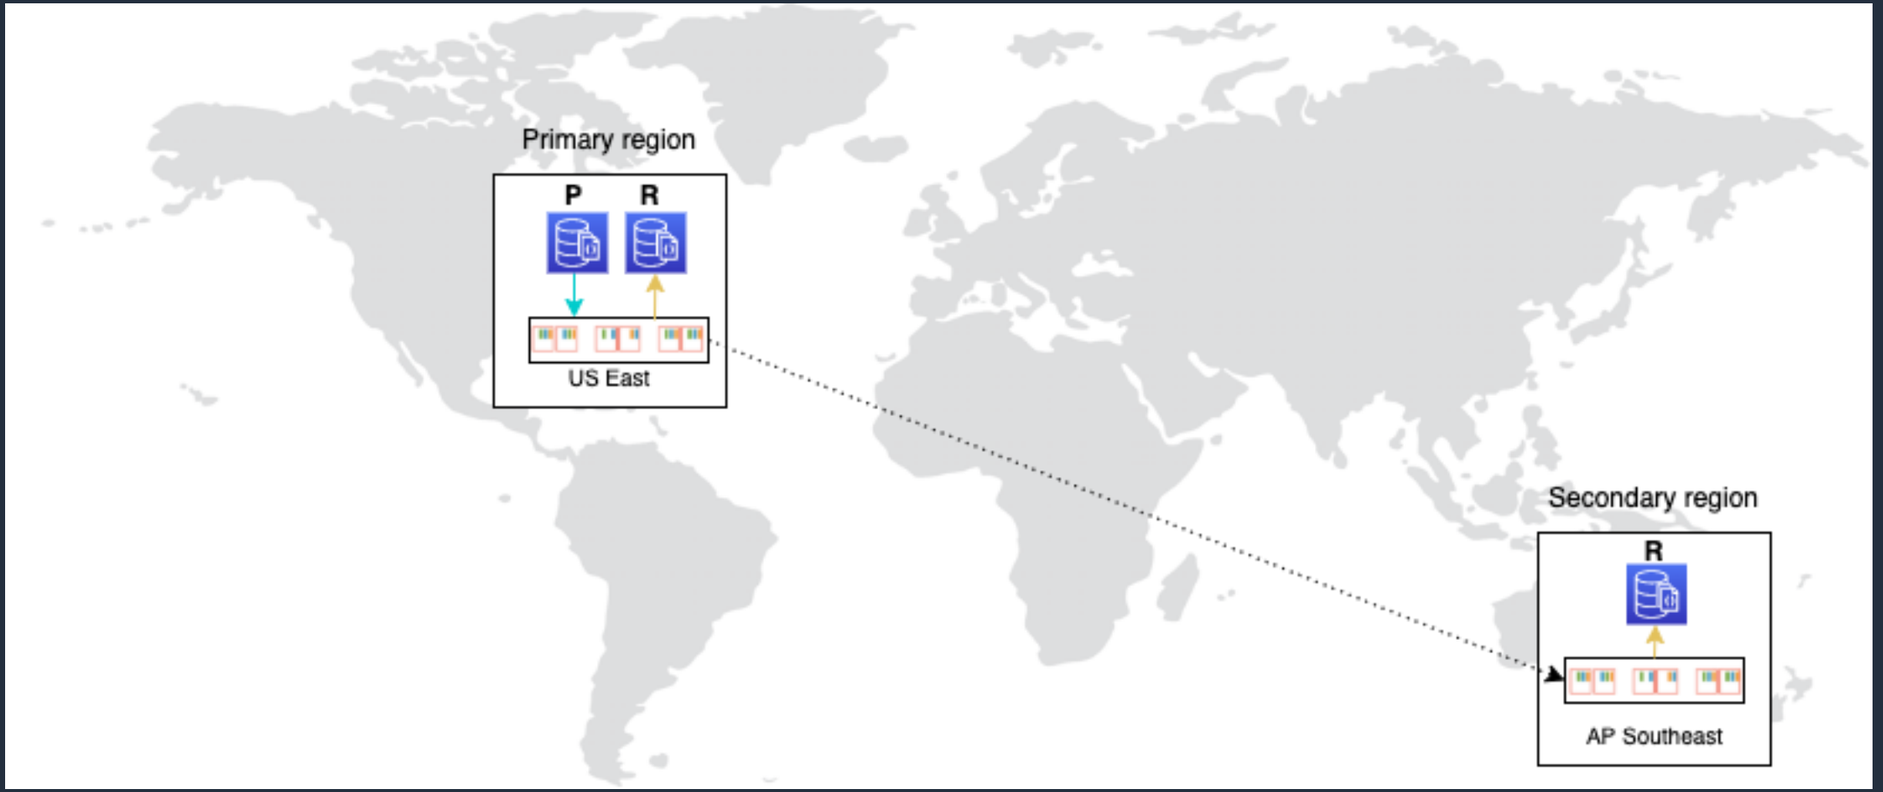

aws

Data locality – Faster reads for global applications

Little to no performance impact on your clusters

Physical replication offloaded to storage layer

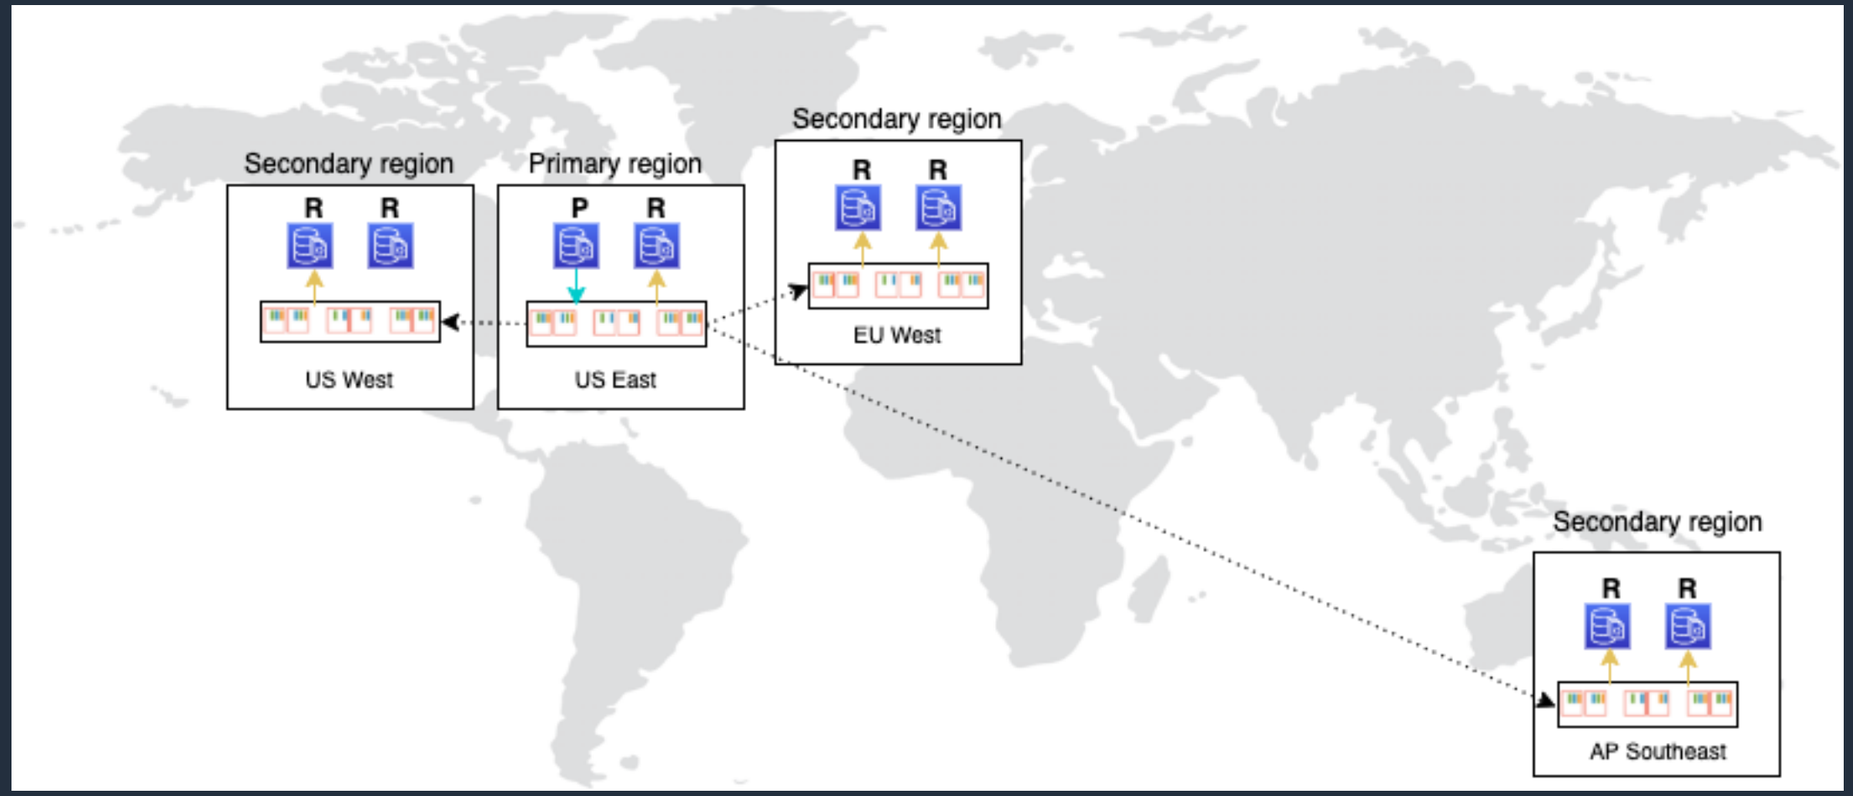

aws

Demo – Use cases

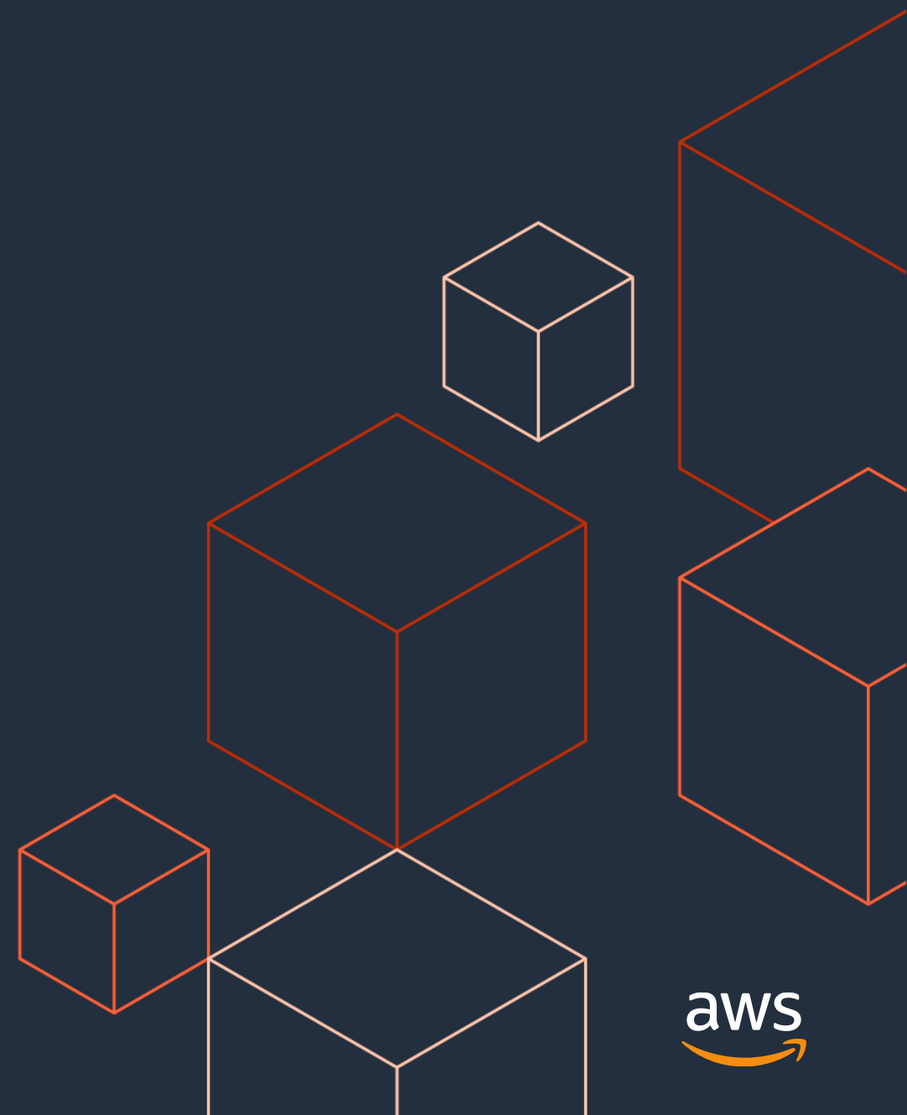

aws

Thank you!