INNOVATE

DATA EDITION

19 August 2021

Build scalable, global applications with Amazon NoSQL databases

Sean Shriver
Sr. DynamoDB Solutions Architect
Amazon Web Services

Karthik Vijayraghavan
Sr. DocumentDB Solutions Architect
Amazon Web Services

Agenda

What is NoSQL?

Relational data, aggregations, NoSQL scaling

What is Amazon DocumentDB (with MongoDB compatibility)?

Overview, customers, internal architecture, demo

What is Amazon DynamoDB?

Overview, global replication, customers, demo

What is NoSQL?

All data is relational

IT monitoring

Social graph

Documents

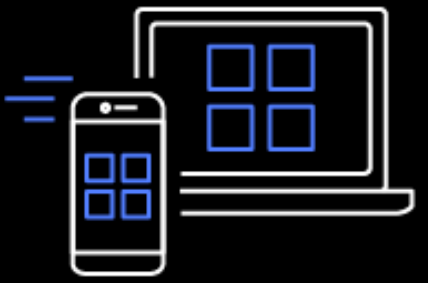

Auth services

Aggregations

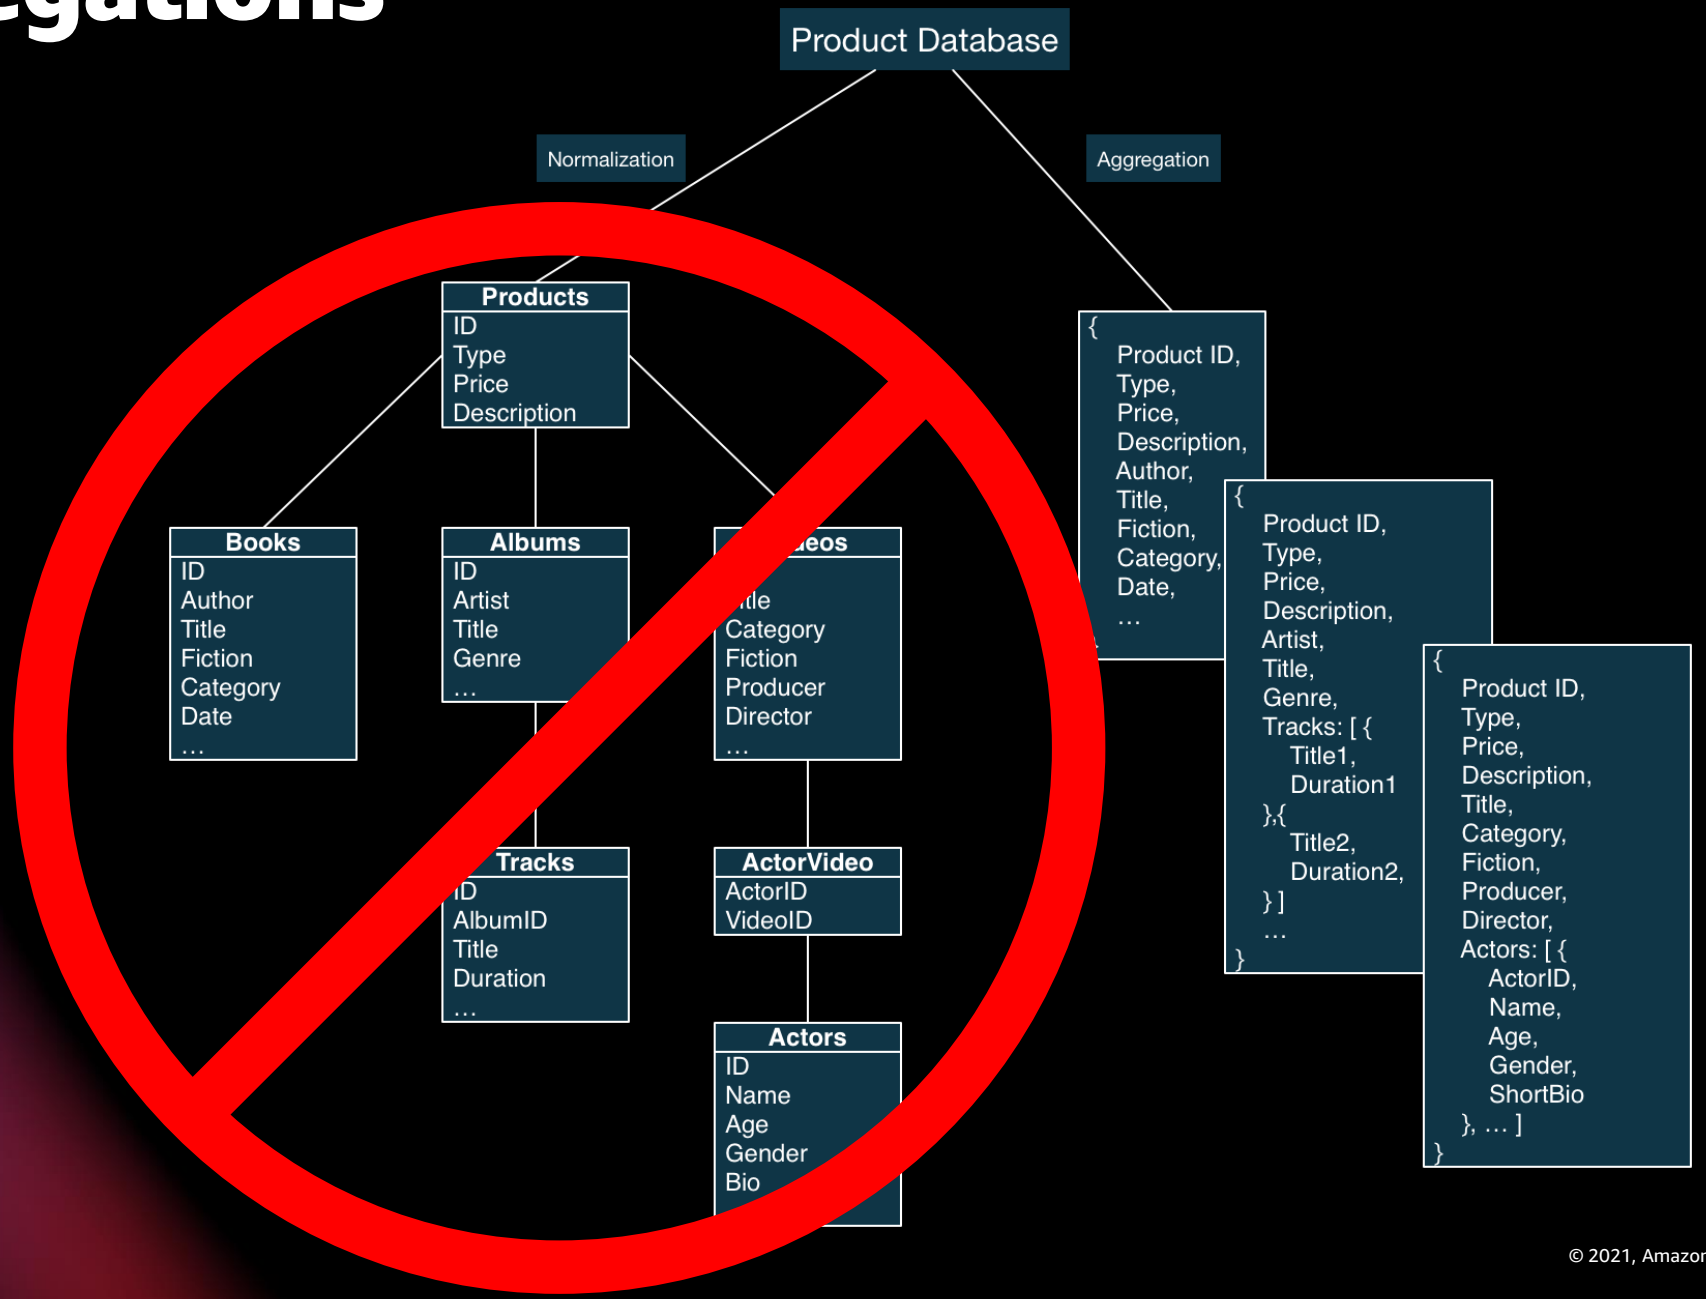

NoSQL scaling

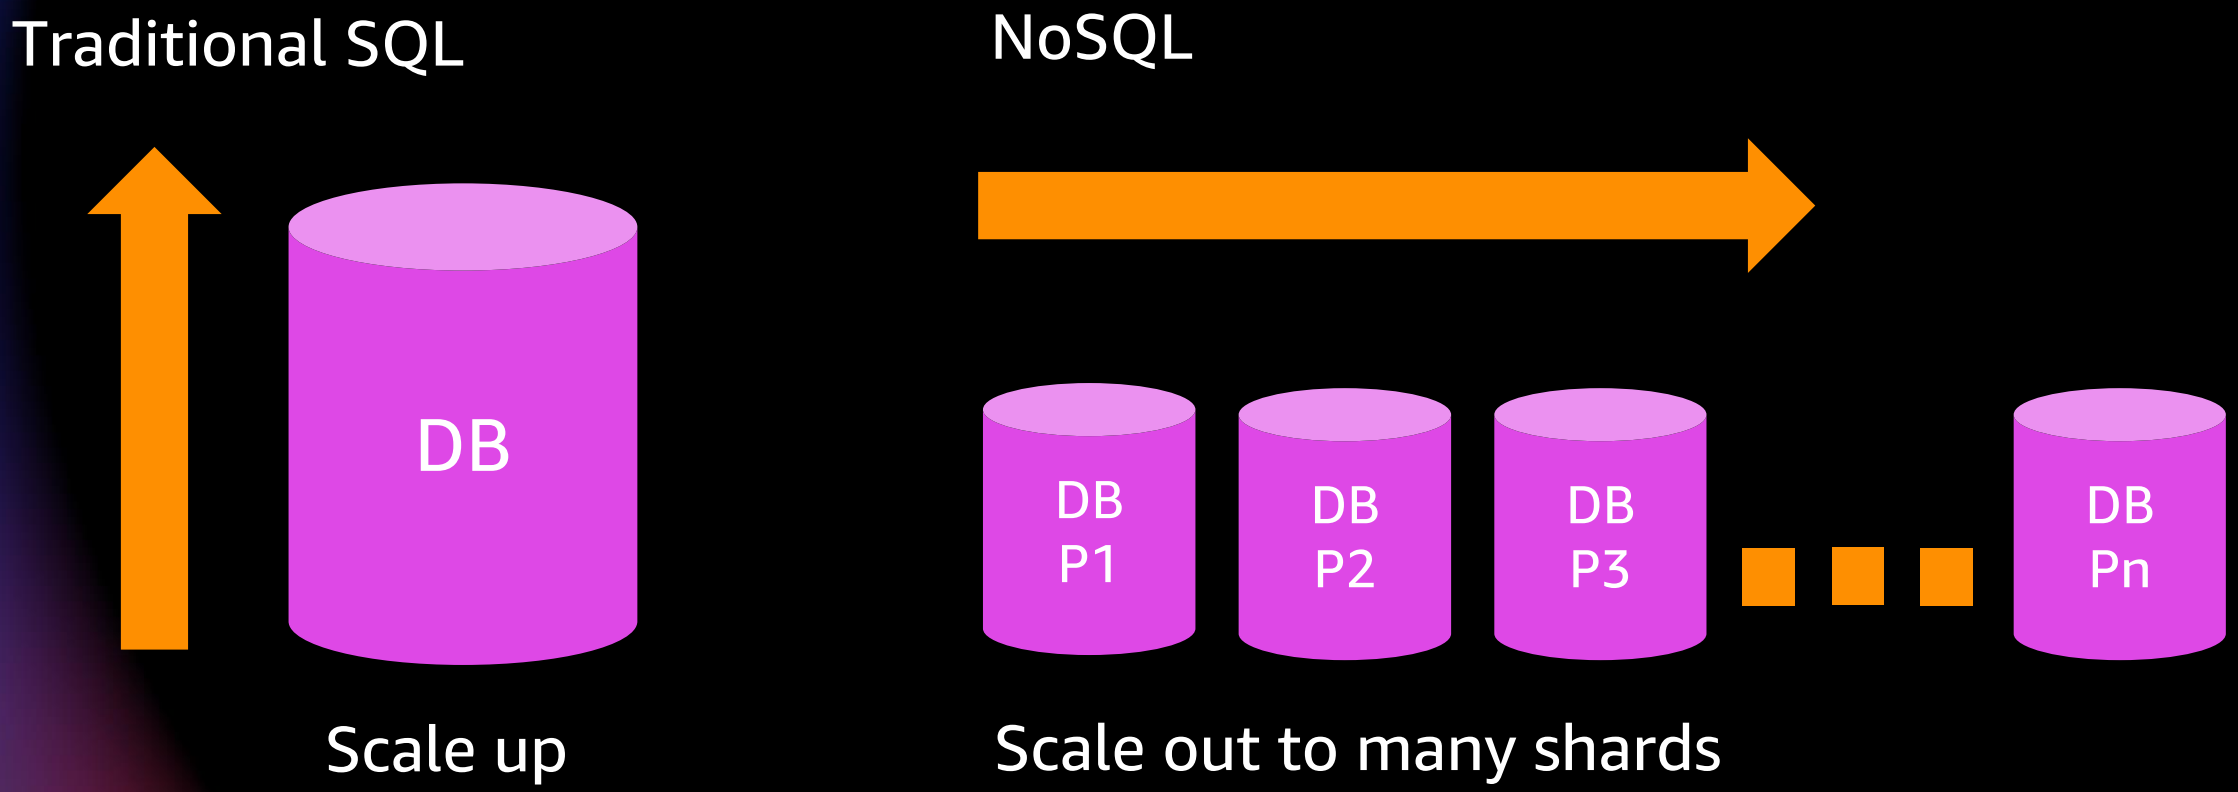

Basic premise: There is a way to design data that's horizontally scalable.

Two NoSQL databases from AWS

Amazon DocumentDB (with MongoDB compatibility)

Enterprise-ready document database service that scales JSON workloads with ease

Amazon DynamoDB

Fast and flexible NoSQL database service for any scale

What is Amazon DocumentDB?

Amazon DocumentDB (with MongoDB compatibility)

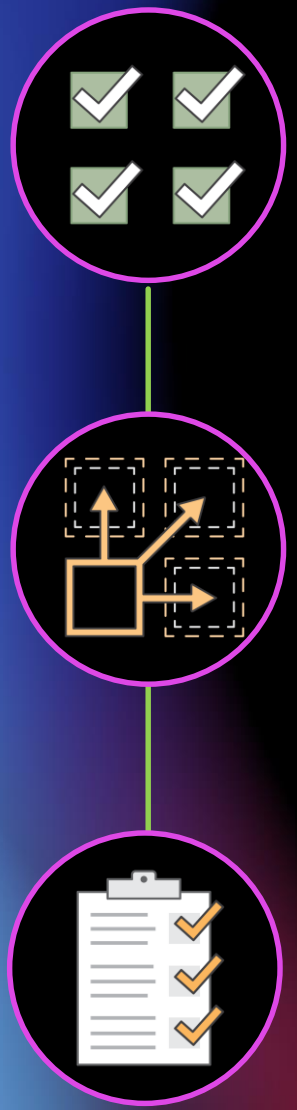

Fully managed

Scalable

MongoDB API compatible

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

朝日新聞

Asahi Shimbun

"Since Amazon DocumentDB is a fully managed service, we are freed from operations such as scaling, backup and patch application, which would **lead to reduced development and operation costs**."

aws

Amazon DocumentDB (with MongoDB compatibility)

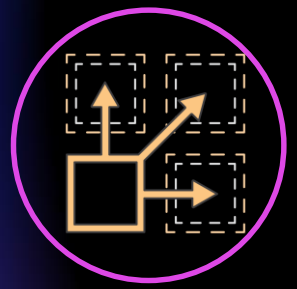

Scalable

Scale compute in minutes

Storage and IO autoscaling

Storage scales to 64TB

Scale out to 15 replicas for millions of reads

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

Amazon DocumentDB (with MongoDB compatibility)

MongoDB API compatible

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

HABBY

"We chose Amazon DocumentDB because of the ease-of-management and its compatibility with MongoDB. We had initially planned to self-manage MongoDB on EC2 but found through testing that it was not convenient and required more work than we were willing to invest."

aws

Amazon DocumentDB customers

BBC

Web publishing

punchh

Personalization

SAMSUNG

IoT

N digital

Customer 360

Retail & marketing

sayurbox

Supply chain management

aws

Case study: Punchh

A LEADING DIGITAL MARKETING SAAS PROVIDER

Problem Statement

- Ingest high data volumes
- Simultaneously run loyalty campaigns with low latency.

Solution advantages

- DB Workload: 2-fold ⬆
- Cost savings: 20% ⬇
- Write performance: 20X ⬆

Requirements

- Store unstructured data with flexible schema.
- Support 100% growth in number of customers.
- Migrate from relational database with no downtime.

Solution

- Dedicated primary instance for high-volume data ingestion.
- Read replicas to run multiple campaigns smoothly.
- AWS Glue integration for relational database migration with no downtime.
- DocumentDB change stream to integrate with downstream applications

aws

Amazon DocumentDB architecture

Amazon DocumentDB architecture

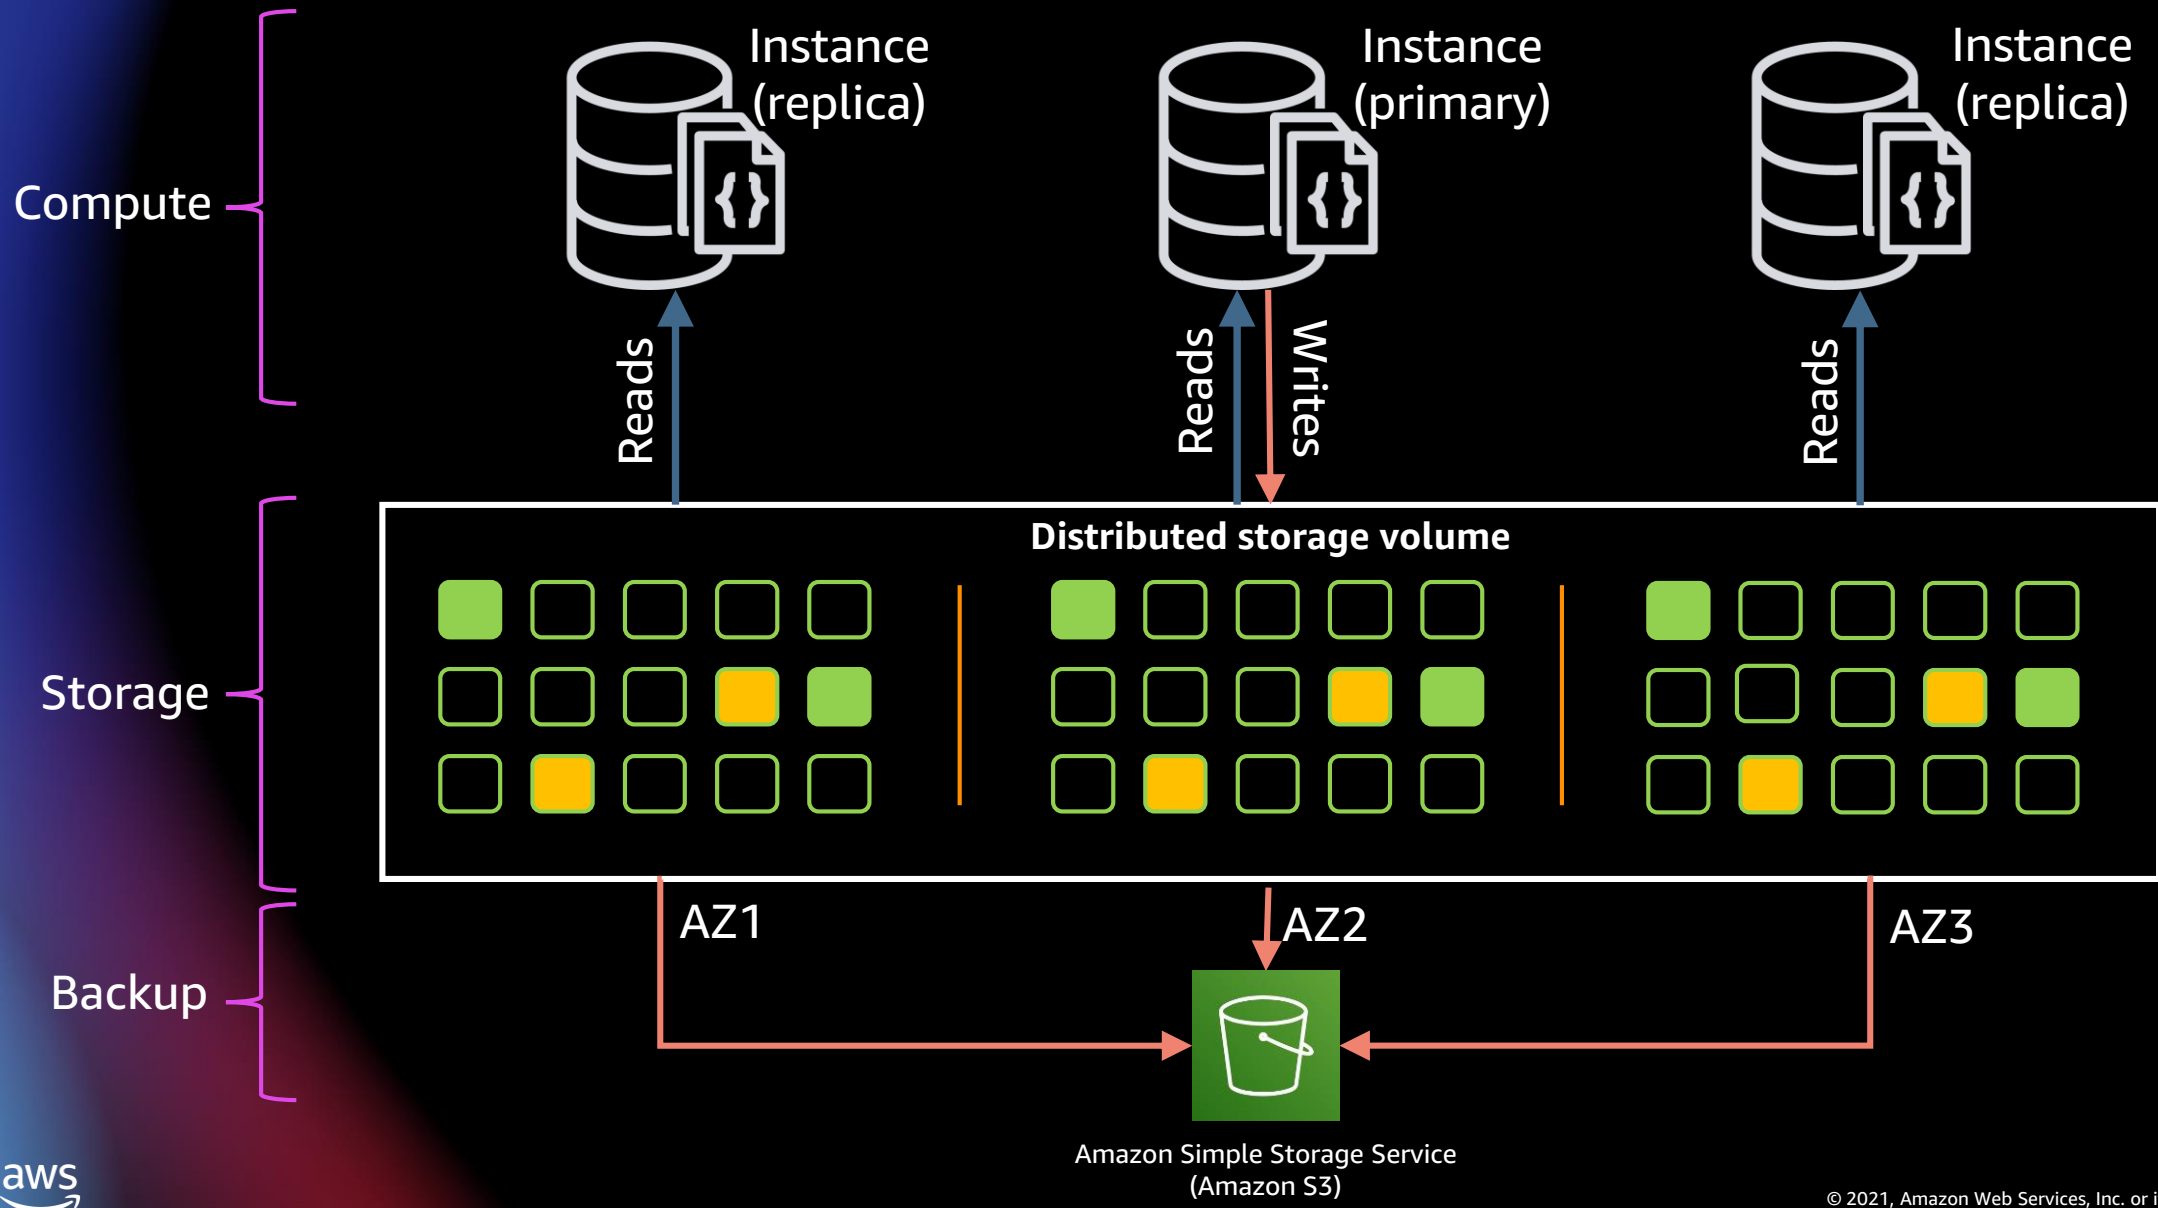

Global clusters

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

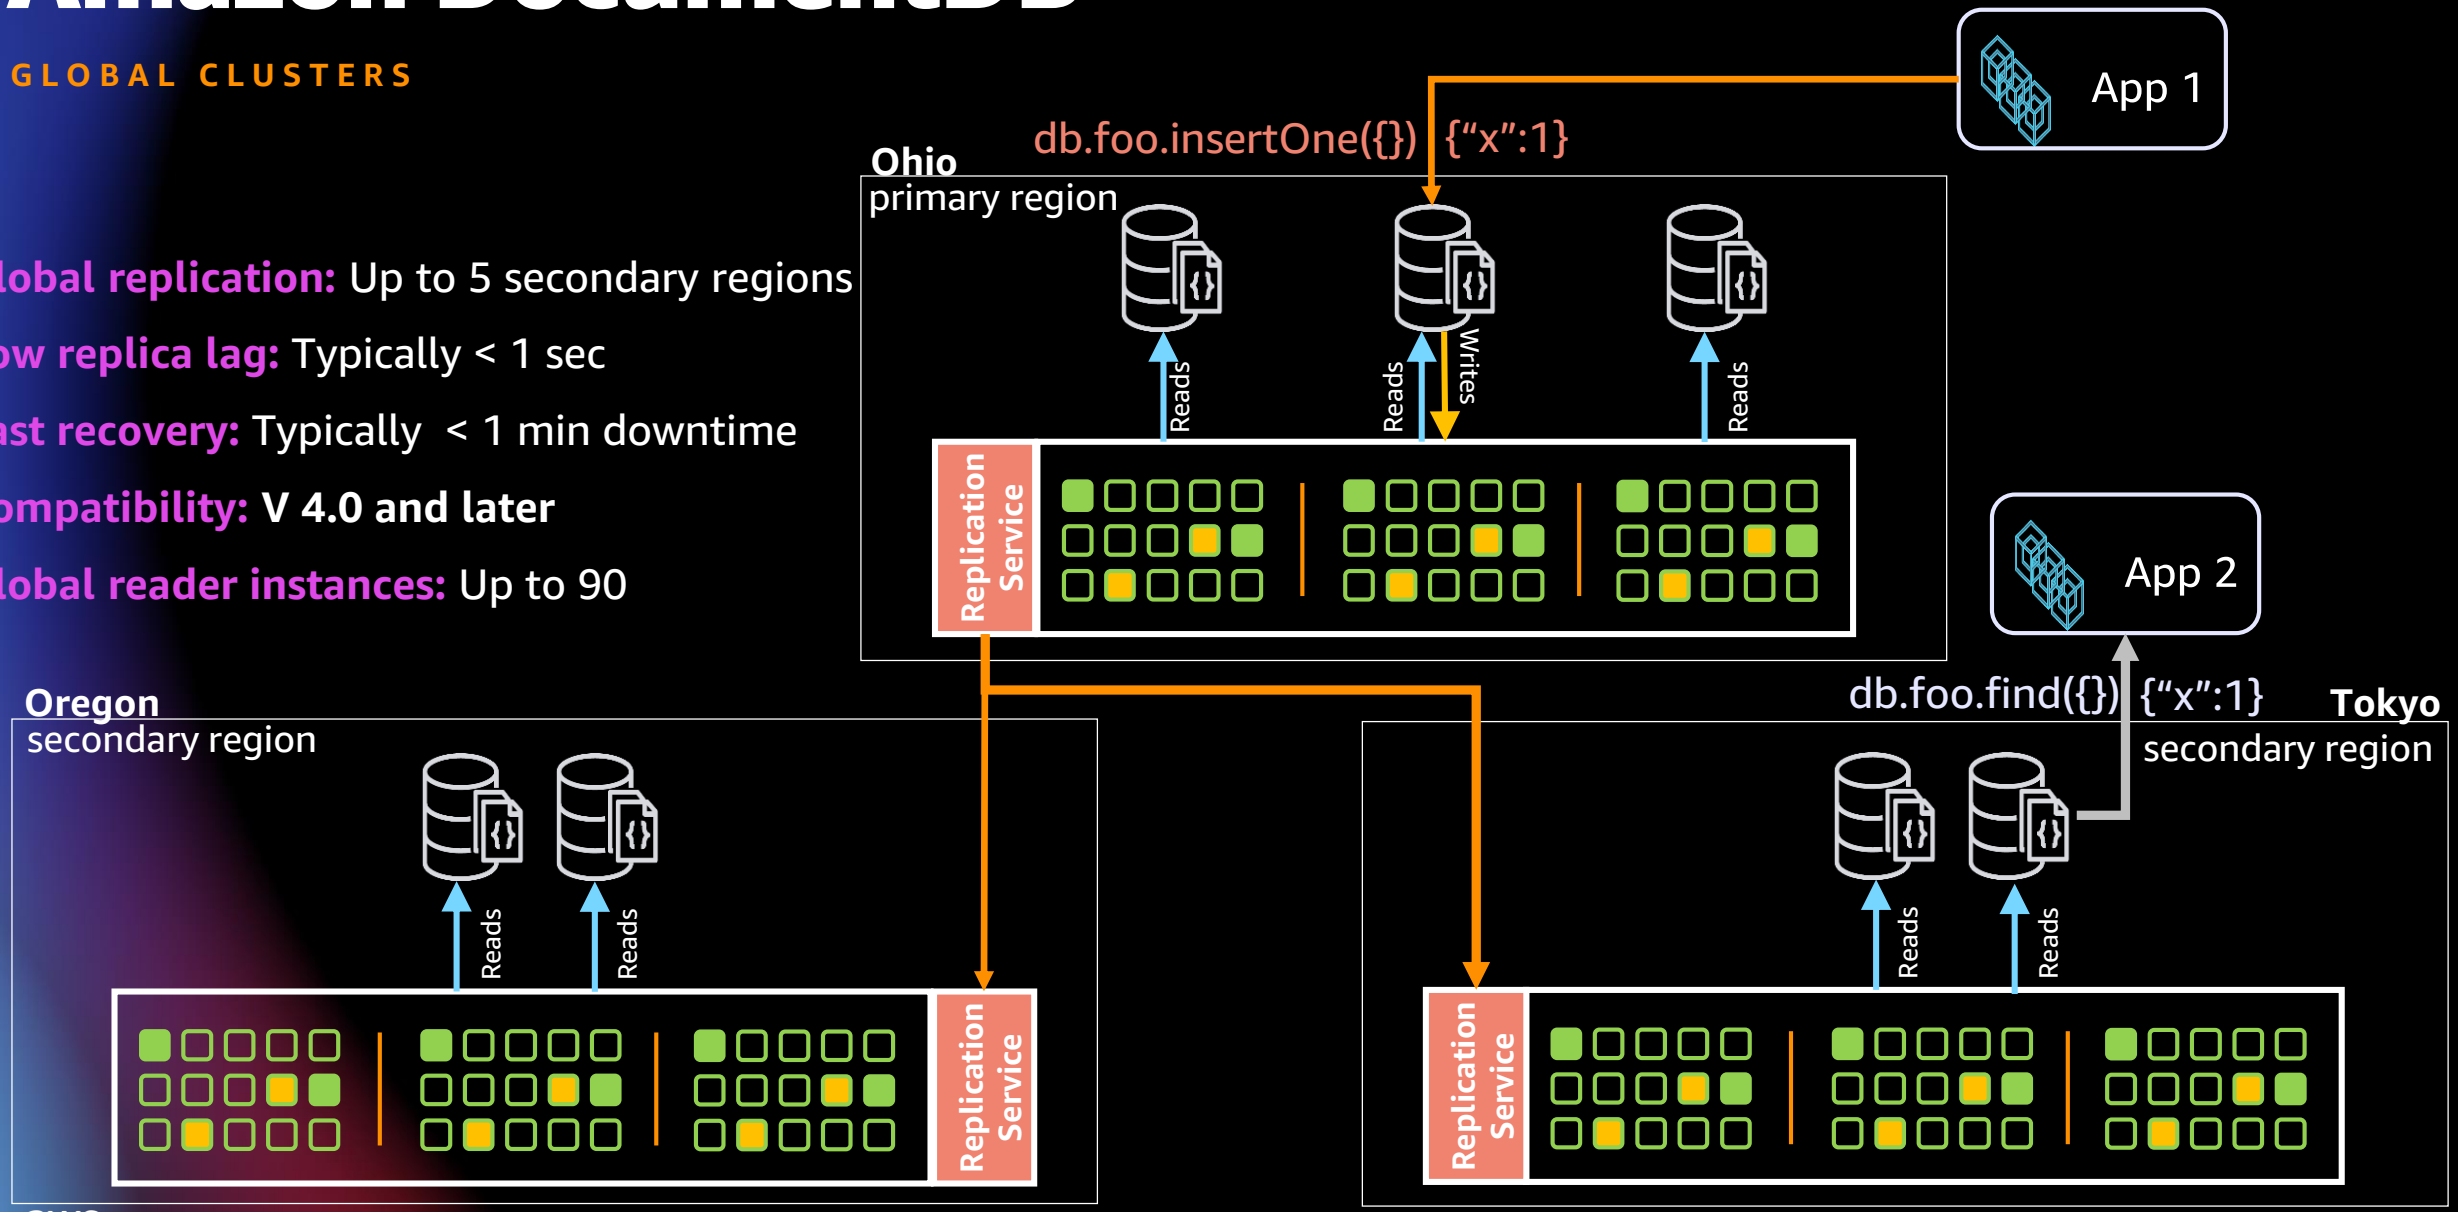

Use cases for global clusters

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

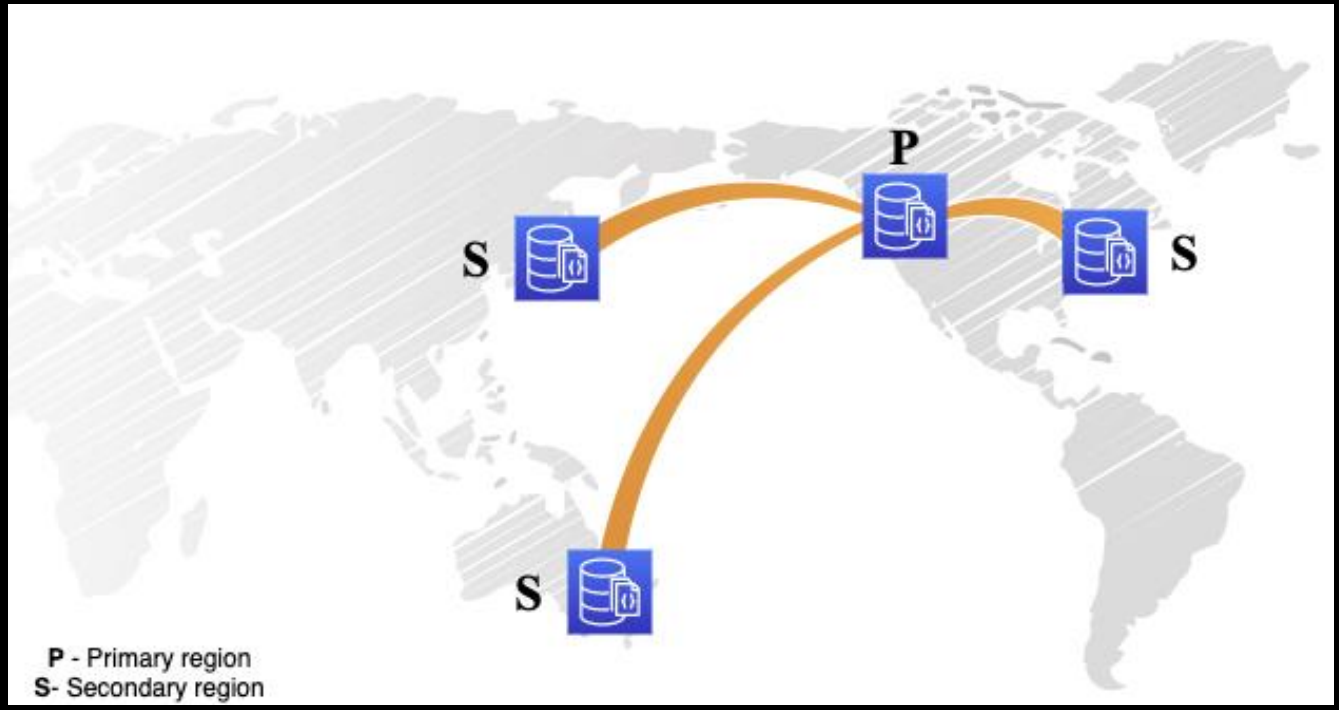

Global clusters demo

What is Amazon DynamoDB?

Amazon DynamoDB

A FAST AND FLEXIBLE KEY-VALUE DATABASE SERVICE FOR ANY SCALE

Performance at scale

Consistent, single-digit-millisecond response times at any scale. Build applications with virtually unlimited throughput.

Serverless architecture

No hardware provisioning, software patching, or upgrades. Scales up or down automatically. Continuously backs up your data.

Enterprise security

Encrypts all data by default and fully integrates with AWS Identity and Access Management (IAM) for robust security.

Global replication

Build global applications with fast access to local data by easily replicating tables across multiple AWS Regions.

aws

How to build a globally replicated database

OR: ALL THE THINGS YOU WOULD HAVE TO DO YOURSELF

- Size replication capacity
- Implement redundancy for reliability
- Scale elastically
- Implement conflict resolution
- Operate replication service (Handle failures, etc.)

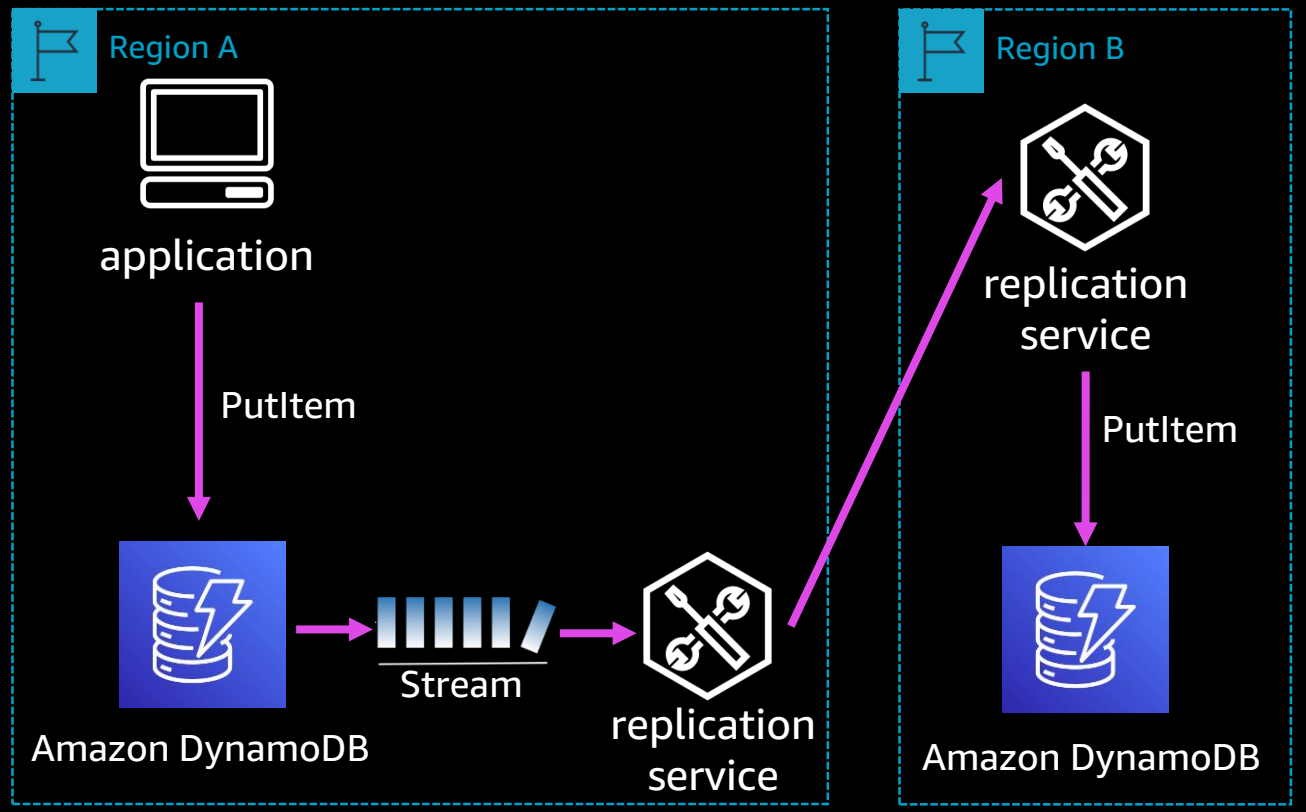

Amazon DynamoDB global tables

AT A GLANCE

- Build high-performance, globally distributed applications
- Low-latency reads and writes to locally available tables
- Multi-Region redundancy and resiliency and 99.999% availability
- Multi-active writes from any Region
- Easy to set up and no application rewrites required

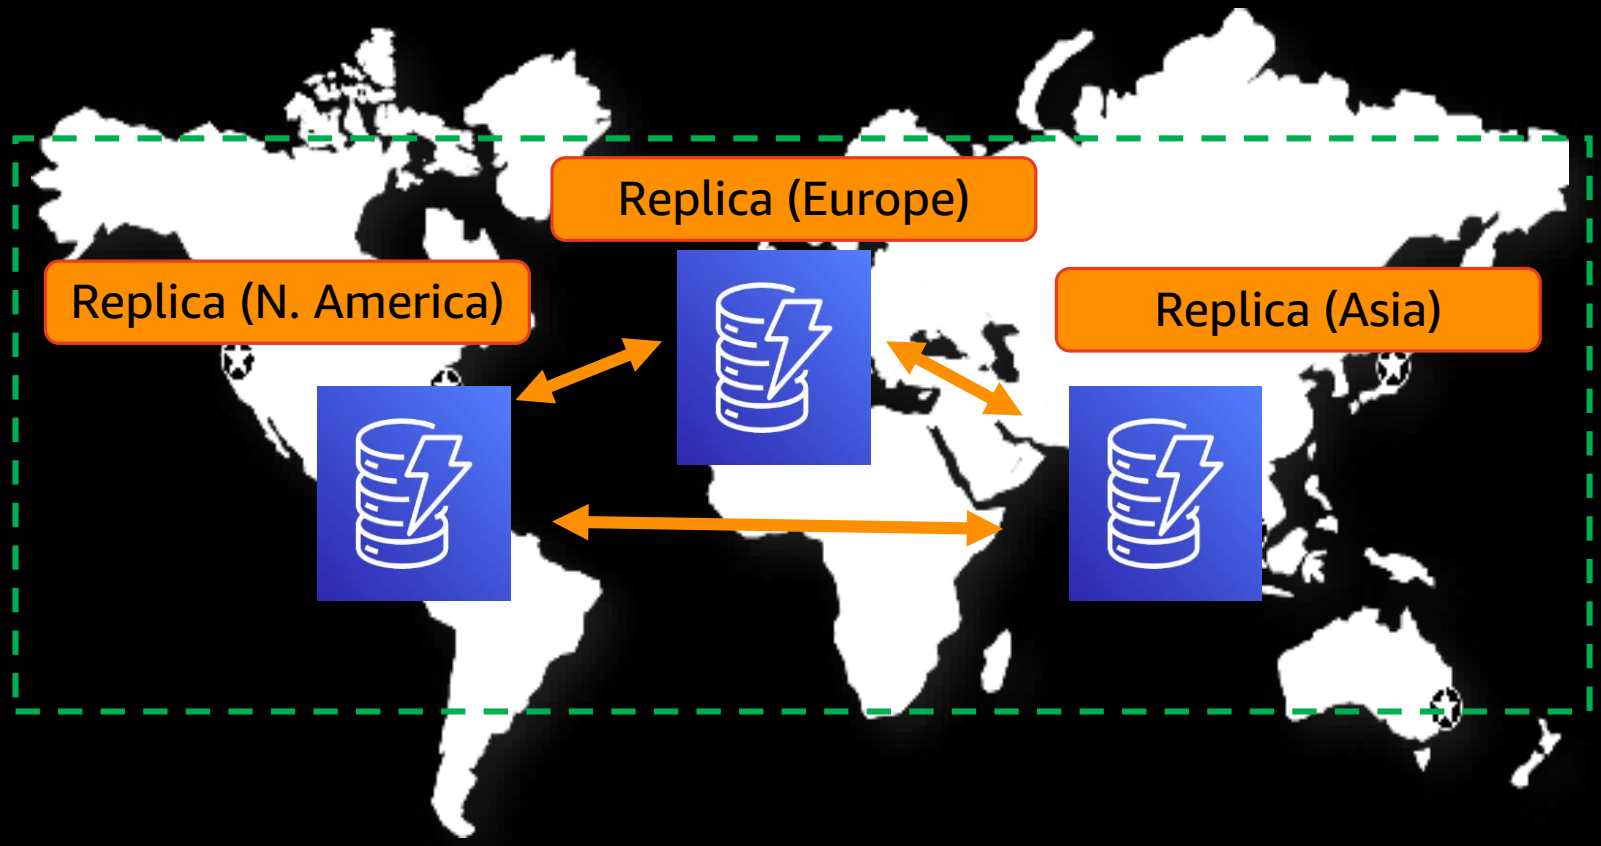

Multi-region replication with Amazon DynamoDB global tables

How global tables resolves conflicts

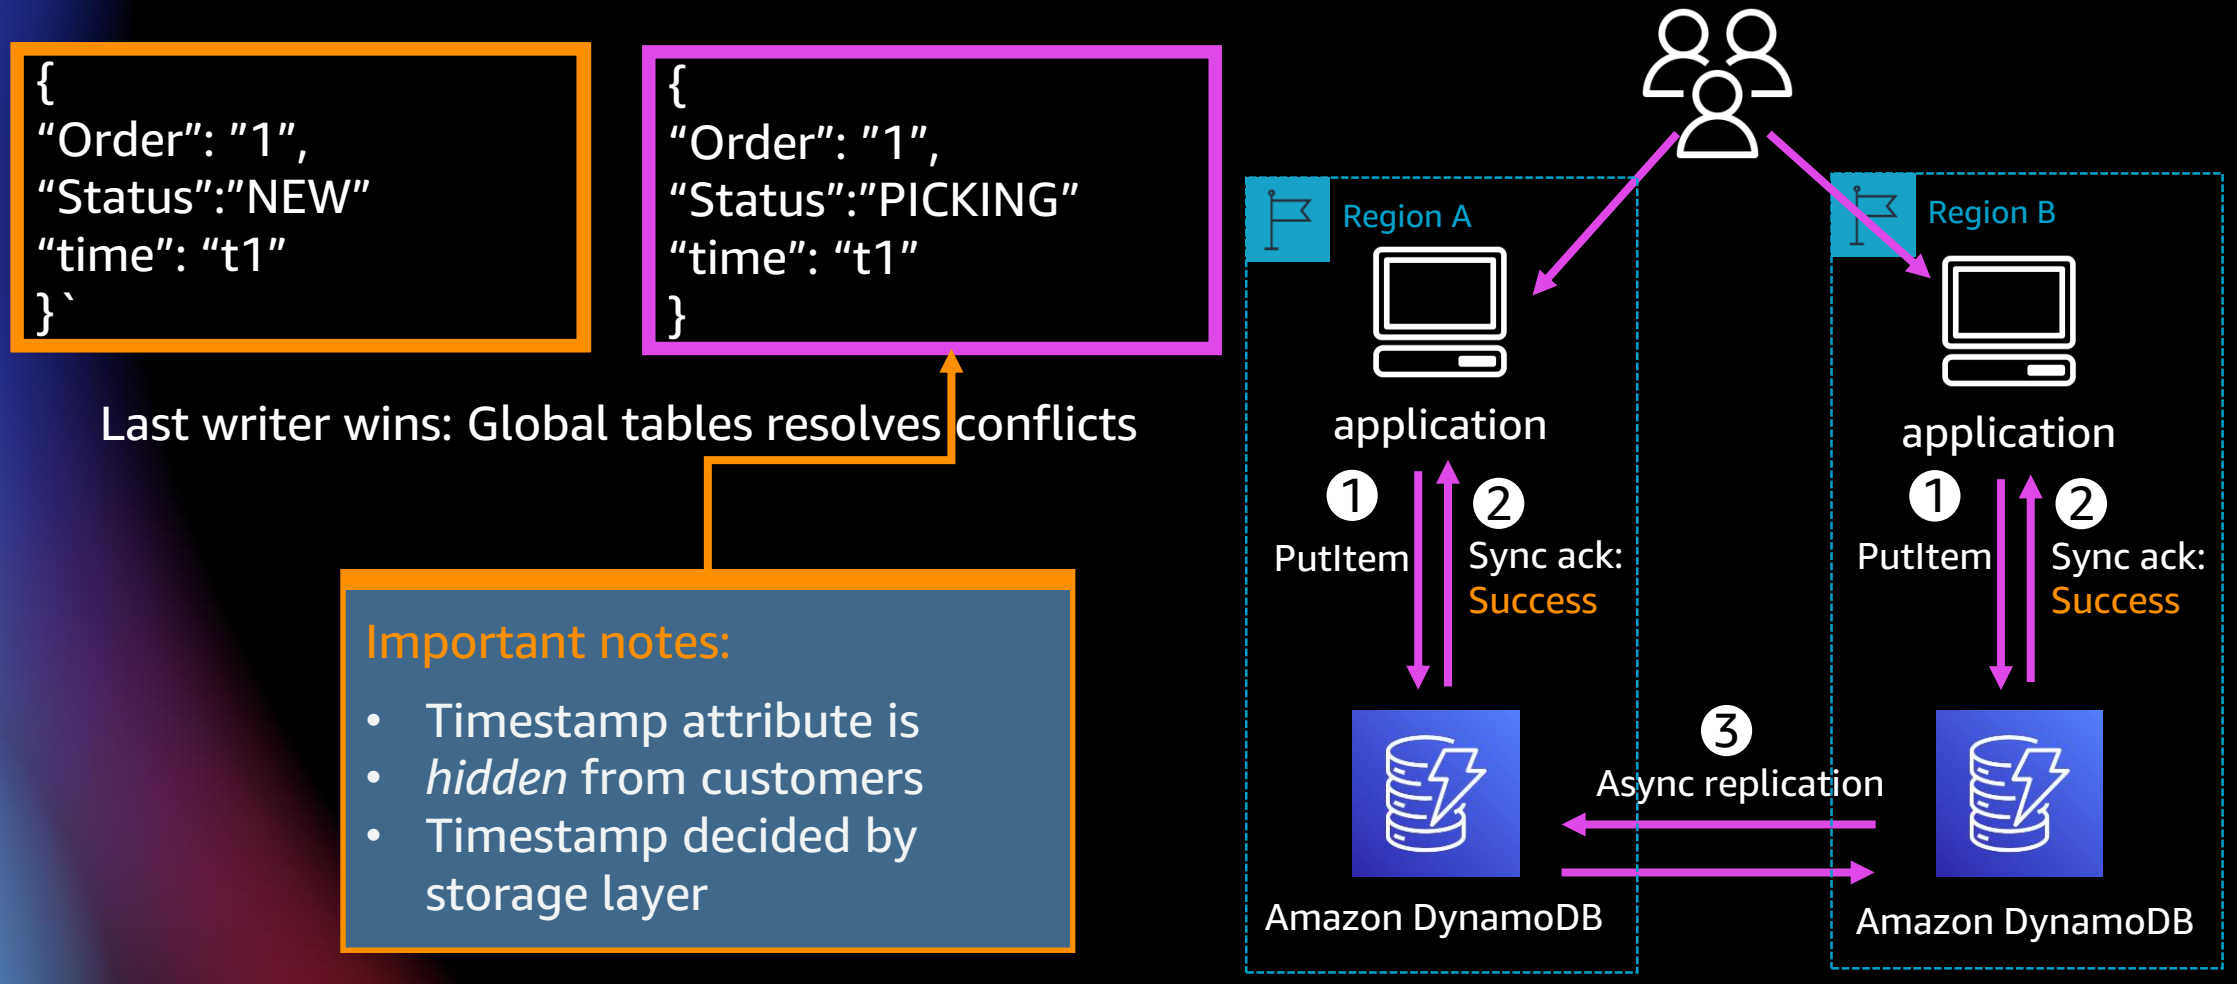

Warm standby without conflicts

SINGLE REGION TO ACCEPT WRITES

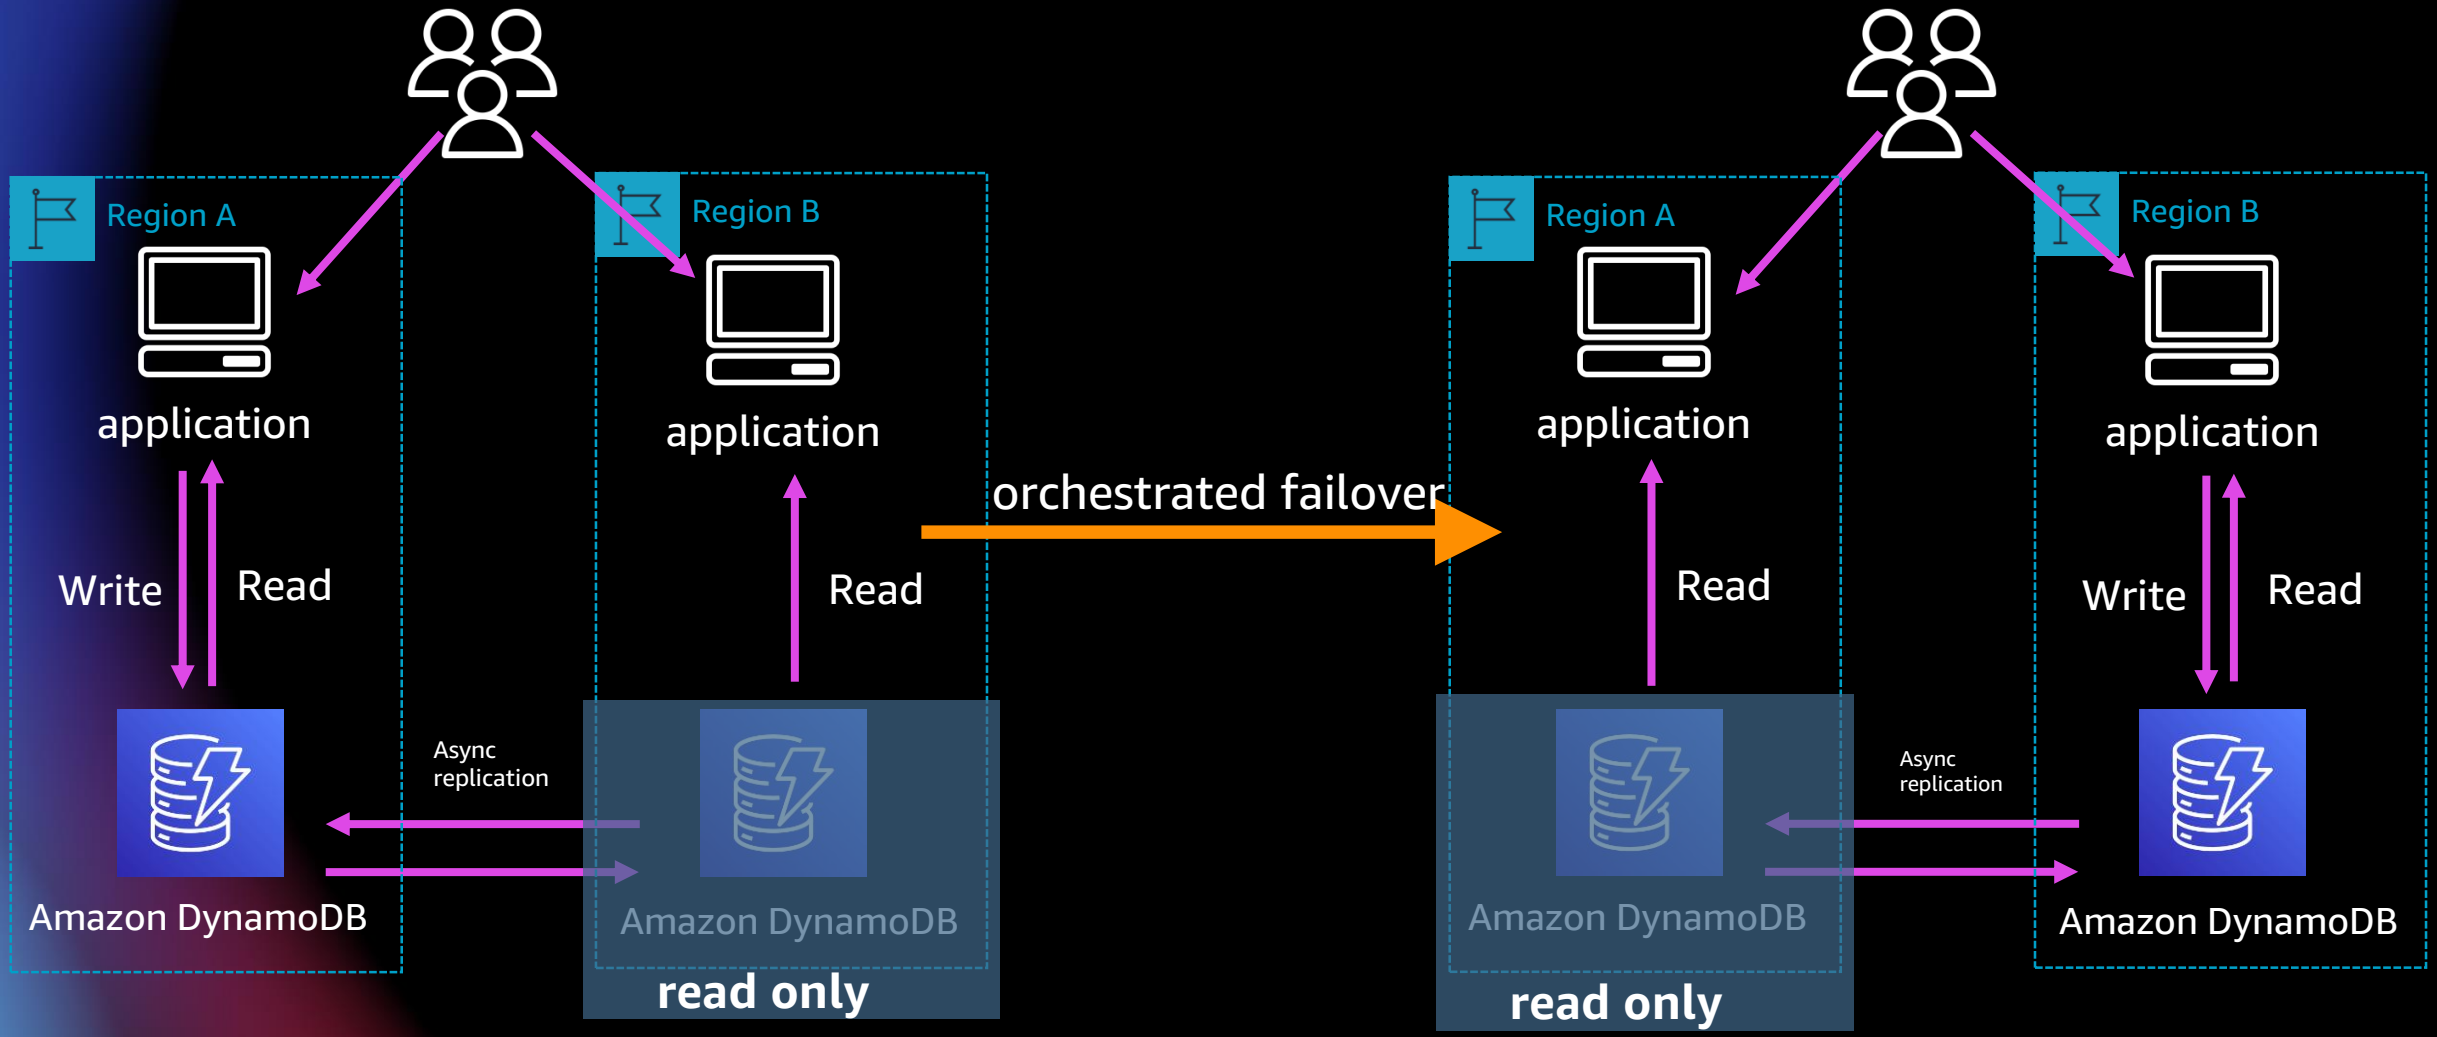

Before: Region A accepts writes

After: Region B accepts writes

aws

Customers around the world

Disney+ on Amazon DynamoDB

HOW DISNEY+ SCALES GLOBALLY ON DYNAMODB

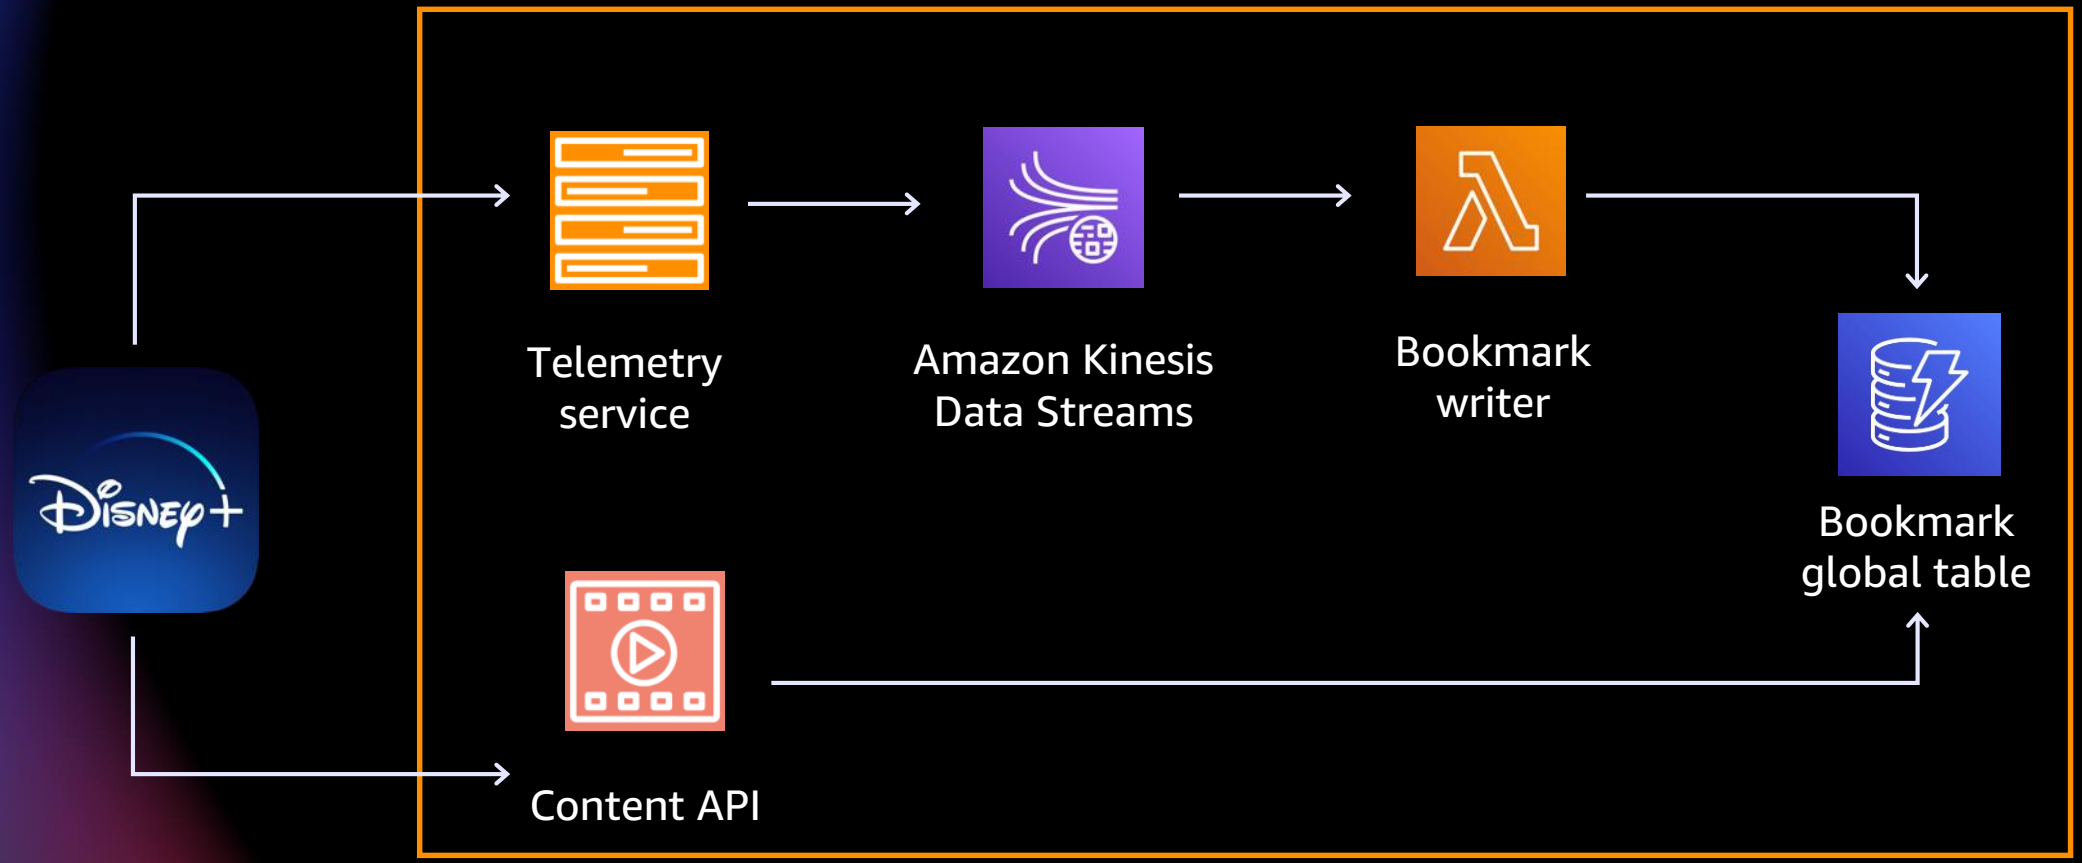

Challenge: Durably store the last viewed location from a video within seconds.
Solution: A telemetry pipeline writing bookmarks to a Amazon DynamoDB global table.

Amazon DynamoDB global tables demo

Global tables demo

1. ~~Create an Amazon DynamoDB global table with AWS CloudFormation~~
2. Restrict writes to only allow one region to write into Amazon DynamoDB
3. Fail over to new region with enforcement provided by AWS IAM

Next steps

Amazon DynamoDB getting started material
aws.amazon.com/dynamodb/getting-started/

Amazon DocumentDB getting started material
aws.amazon.com/documentdb/getting-started/

aws-apj-marketing@amazon.com

twitter.com/AWSCloud

facebook.com/AmazonWebServices

youtube.com/user/AmazonWebServices

slideshare.net/AmazonWebServices

twitch.tv/aws

Visit the AWS Data Resource Hub

Dive deeper with these resources, get inspired and learn how you can use data to make better decisions and innovate faster.

- Building a winning data strategy
- The new leadership mindset for data & analytics
- Harness data to reinvent your organization
- Put your data to work with a modern analytics approach
- Breaking free from on-premises database constraints
- Cloud storage adoption: From cost optimization to agility & innovation
- A strategic playbook for data, analytics, and machine learning
- … and more!

https://tinyurl.com/aws-data-resource

Visit resource hub

Thank you!

Sean Shriver
@sean_shriver

Karthik Vijayraghavan
@karthikv_vijay

AWS Innovate – Data Edition

We hope you found it interesting! A kind reminder to **complete the survey.**
Let us know what you thought of today's event and how we can improve the event experience for you in the future.


aws-apj-marketing@amazon.com


twitter.com/AWSCloud


facebook.com/AmazonWebServices


youtube.com/user/AmazonWebServices


slideshare.net/AmazonWebServices


twitch.tv/aws