Getting started with Amazon Aurora Global Database

ROHAN BHATIA

Senior Product Manager
AWS RDS

Agenda

- Aurora Global Database overview
- Working with RDS Proxy and Global Databases
- Demo

Amazon Aurora

MYSQL- AND POSTGRESQL-COMPATIBLE RELATIONAL DATABASE BUILT FOR THE CLOUD

Fully managed: no hardware provisioning, patching, setup, or backups

Performance of commercial databases at 1/10th the price

High availability and cross-Region disaster recovery

Autoscaling compute, storage, and IO

The fastest growing service in the history of AWS

Aurora Global Database overview

FASTER DISASTER RECOVERY AND ENHANCED DATA LOCALITY

- **Global replication:** Up to 5 secondary regions
- **Low replica lag:** Typically < 1 sec cross-region lag
- **Fast recovery:** < 1 min downtime after region unavailability
- **High throughput:** negligible performance impact for writes

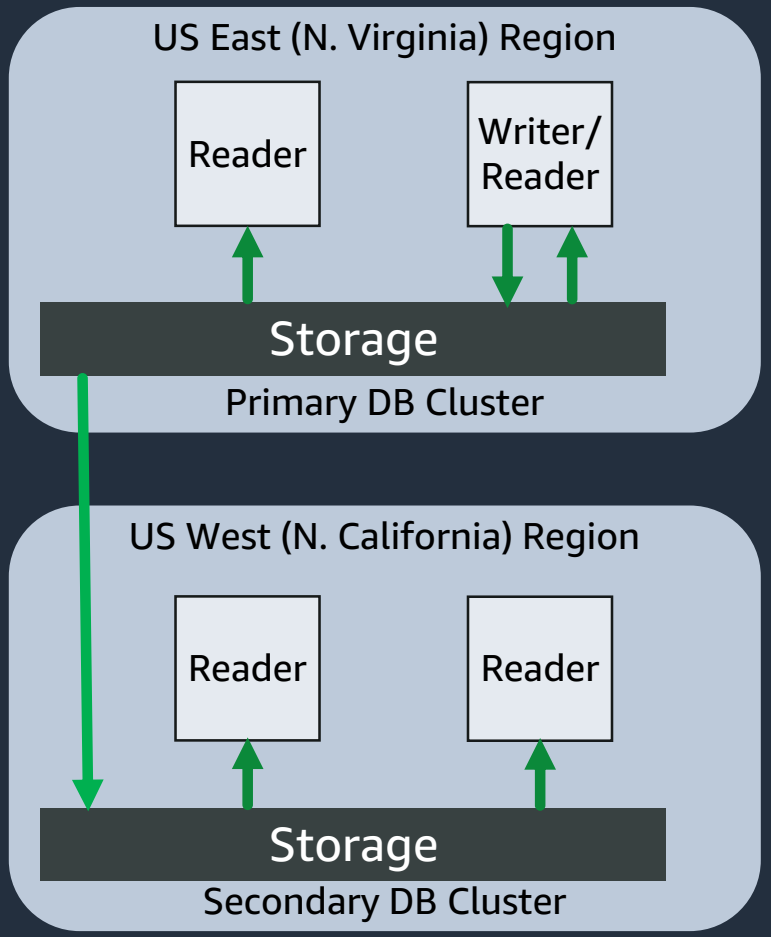

Global Database – Architecture

FASTER DISASTER RECOVERY AND ENHANCED DATA LOCALITY

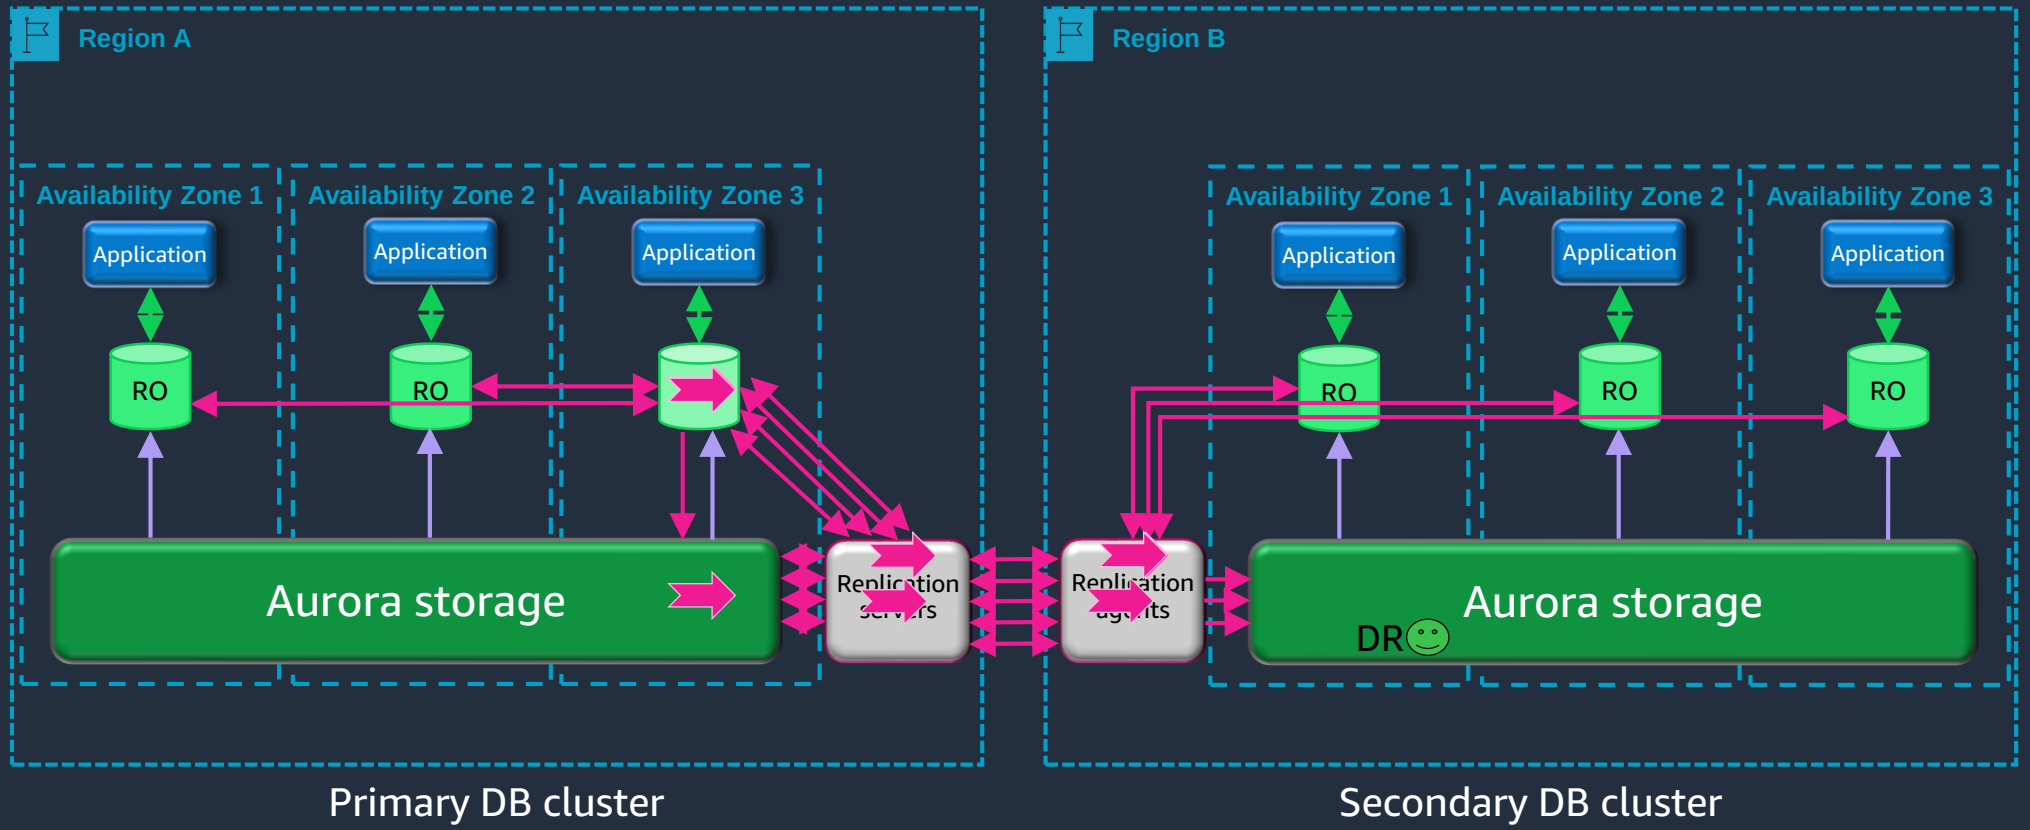

Primary DB cluster

Secondary DB cluster

Global Database use cases

- **Disaster recovery**

Cross-region disaster recovery

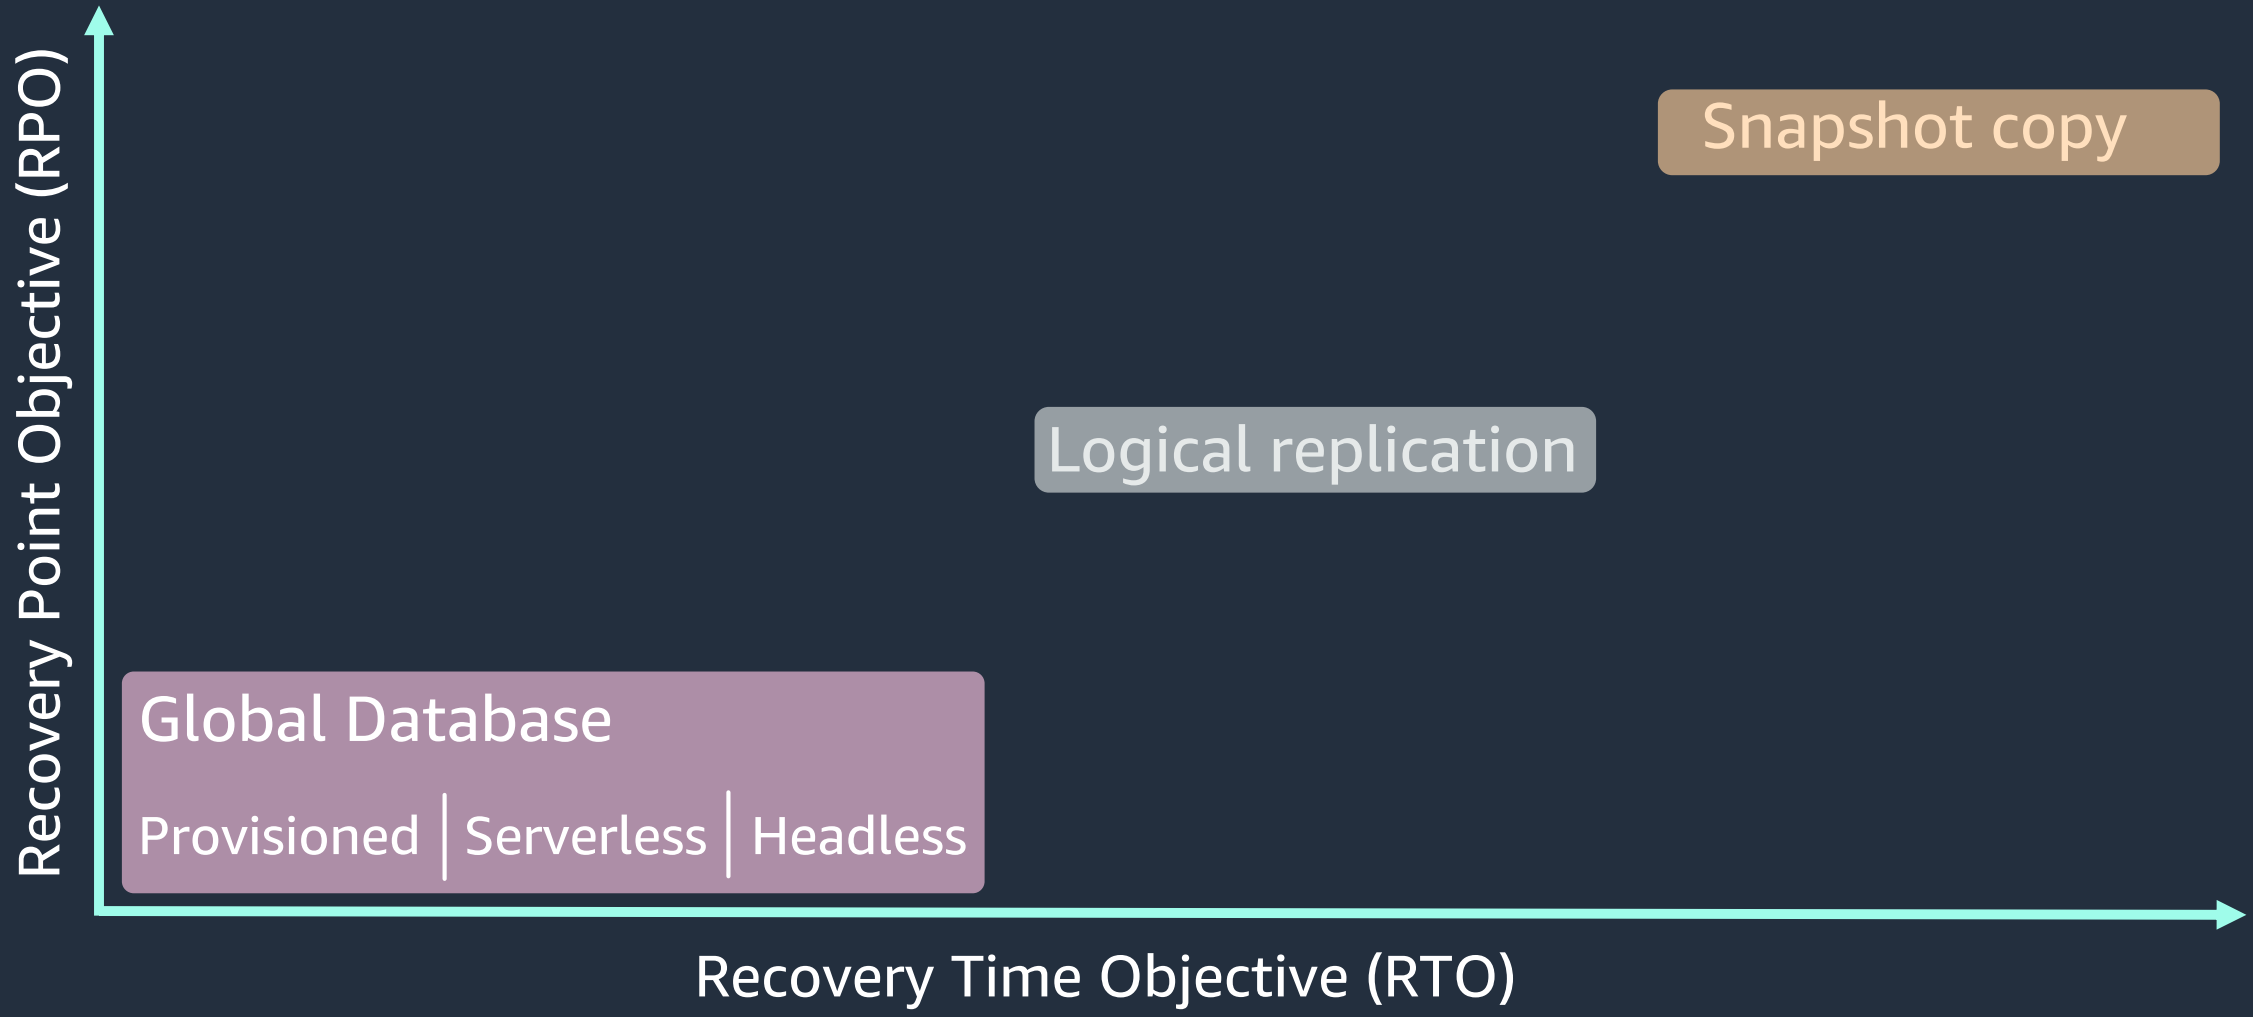

Global Database - Managed planned failover

- Test disaster recovery or make planned maintenance changes
- Automated promotion of a secondary region to primary region
- Maintains the Global Database topology
- No data loss; writes are stopped until the target failover region catches up

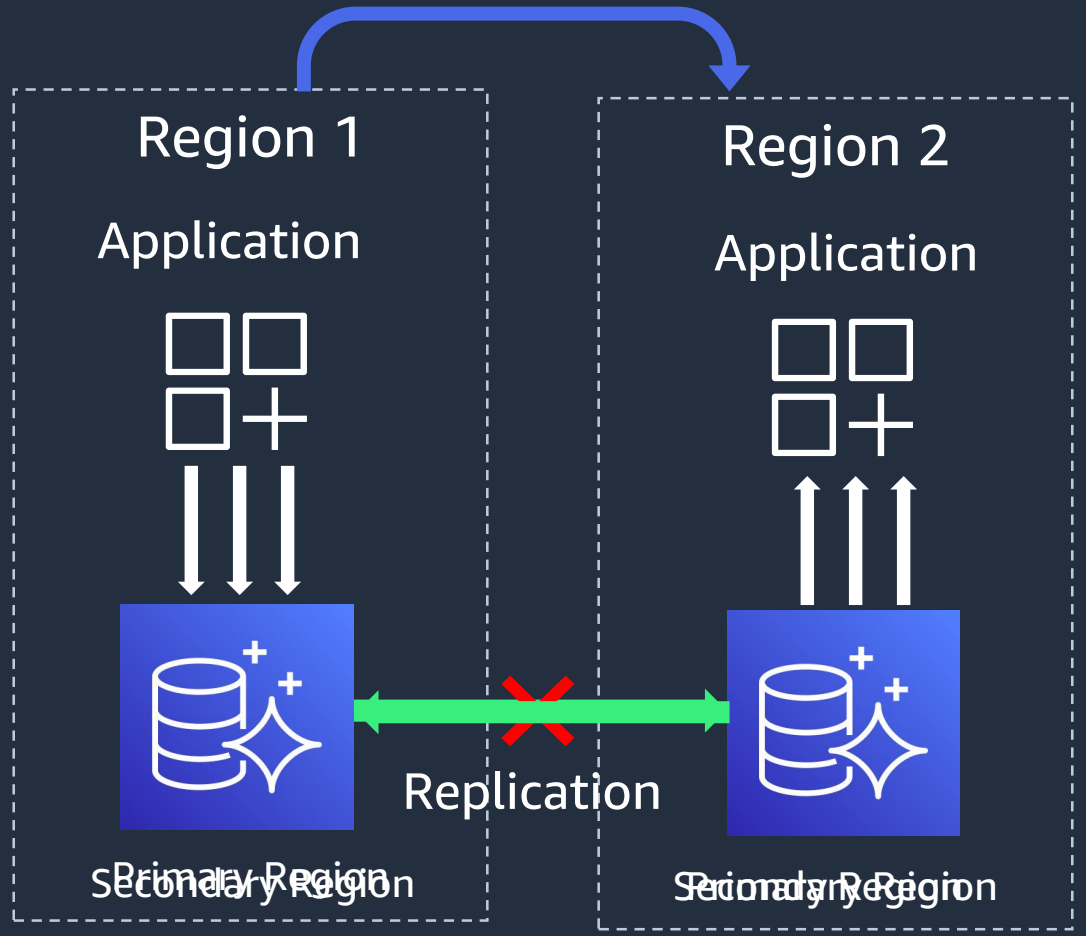

Global Database - Manual unplanned failover

- Recover from unplanned outage in typically under a minute
- Detach and promote a secondary region to a primary region
- Manage recovery point objective (RPO) by specifying a desired window of data recovery
- At least one secondary region stays within the RPO window

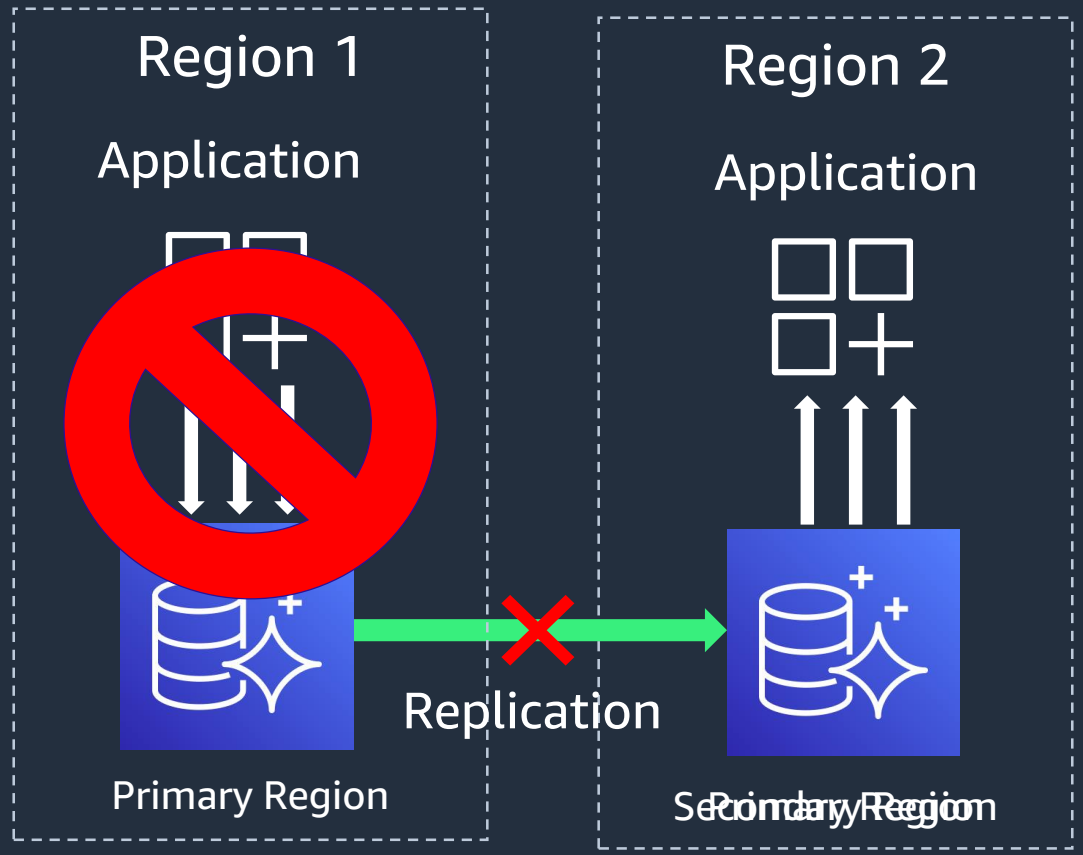

Global Database use cases

- **Low-latency global reads**

Global Database – Read scalability

- Scale database reads across regions and place your applications close to end users
- Read from up to 5 secondary regions (6 total)
- Low-latency (<1 second) reads across regions
- Limited impact to performance and throughput

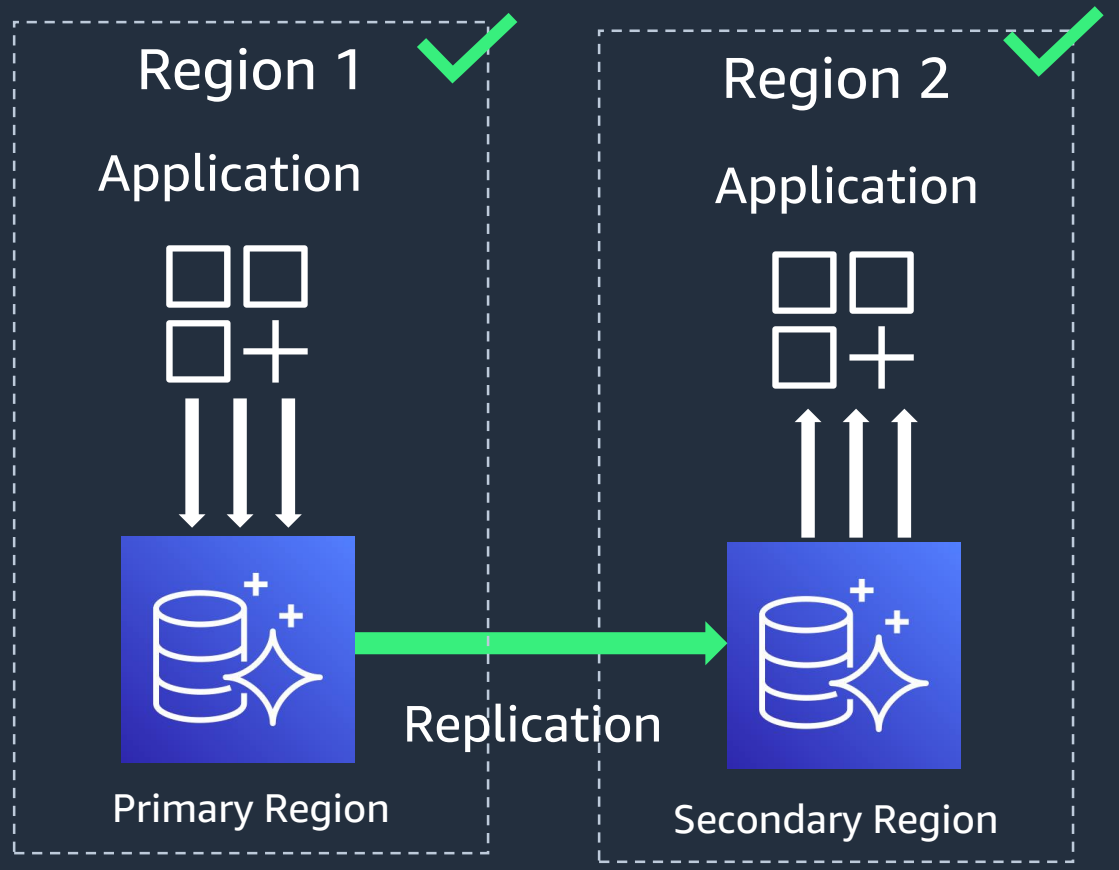

Global Database use cases

- **Global applications with writes from multiple regions**

Global Database - Write Forwarding

- Save time from implementing logic to manage writes from multiple regions
- Build global applications that are agnostic to a single primary region
- Forward writes from the secondary region to the primary region
- Single writer to provide read-after-write consistency in secondary region

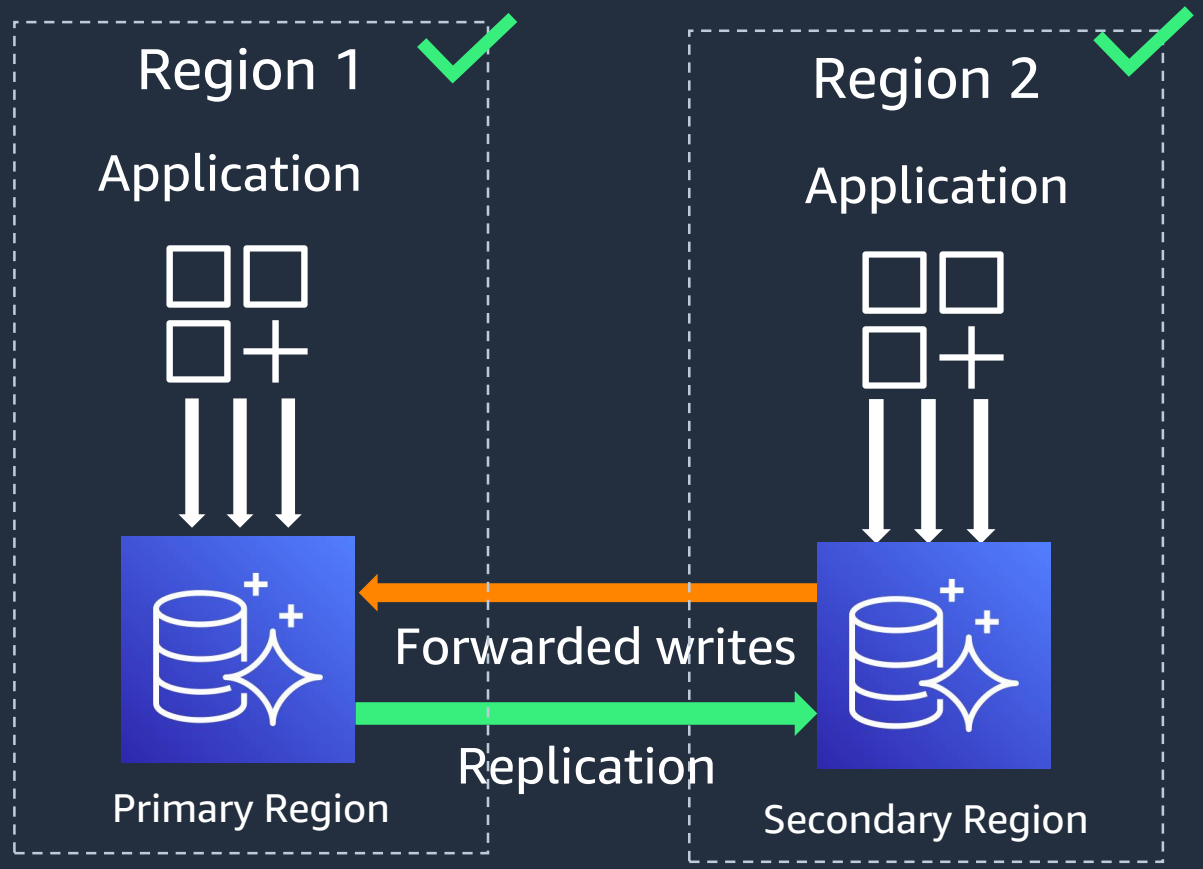

Amazon RDS Proxy

A FULLY MANAGED, HIGHLY AVAILABLE DATABASE PROXY FOR AMAZON RDS AND AMAZON AURORA

Pool and share DB connections for improved app scaling

Increase app availability and reduce DB failover times by up to 66%

Manage app data security with DB access controls

Fully managed DB proxy, compatible with your database

Amazon RDS Proxy supports Aurora Global Database primary and secondary regions
Supports Aurora, RDS MySQL, RDS PostgreSQL, RDS MariaDB, and RDS SQL Server

aws

RDS Proxy benefits with Global Databases

Faster recovery: Speed up recovery time and queue requests during cross-region failovers

Global read scalability: pool and share database connections for read replicas in secondary regions

Improved security: enforce IAM authentication in both primary and secondary Global Database regions

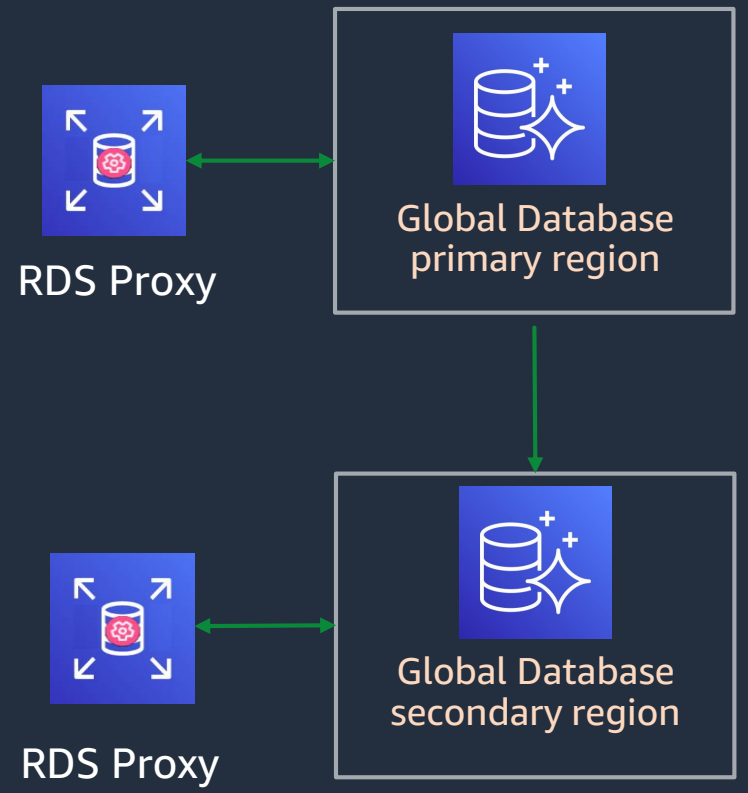

Demo

RDS > US East (N. Virginia)

https://us-east-1.console.aws.amazon.com/rds/home?region=us-east-1#

AWS Tools | aws prfaq | Rialto | Common Networking | Quips | Compliance | Onboarding/Trainin... | Social Media | Risk Mitigation | AWS Common Wikis | Sales Console | Documents | Other Bookmarks

aws | Services | Search [Option+S] | N. Virginia

Resource Groups & Tag Editor

Try the new Amazon RDS Multi-AZ deployment option for MySQL and PostgreSQL

For your Amazon RDS for MySQL and PostgreSQL workloads, improve transactional commit latencies by 2x, experience faster failover typically less than 35 seconds and, get read scalability with two readable standby DB instances by deploying the Multi-AZ DB cluster Learn more

Create database

Or, Restore Multi-AZ DB Cluster from Snapshot

Resources

Refresh

You are using the following Amazon RDS resources in the US East (N. Virginia) region (used/quota)

- DB Instances (5/40)
  - Allocated storage (0.8 TB/100 TB)
  - Increase DB instances limit
- DB Clusters (2/40)
- Reserved instances (0/40)
- Snapshots (23)
  - Manual
    - DB Cluster (5/100)
    - DB Instance (0/100)
  - Automated
    - DB Cluster (2)
    - DB Instance (16)
- Parameter groups (20)
  - Default (20)
  - Custom (0/100)
- Option groups (11)
  - Default (11)
  - Custom (0/20)
- Subnet groups (2/50)
- Supported platforms VPC
- Default network vpc-08268675

Recommended for you

Build RDS Operational Tasks

Watch how to enable users to perform common tasks such as snapshots or restart DB instances in Amazon RDS. Learn more

Time-Series Tables in PostgreSQL

Step-by-step guide to design high-performance time series data tables on Amazon RDS for PostgreSQL. Learn more

Amazon RDS Backup and Restore using AWS Backup

Learn how to backup and restore Amazon RDS databases using AWS Backup in just 10 minutes. Learn more

Test Your DR Strategy in Minutes

Amazon Aurora Global Database now supports planned managed failover, making disaster recovery drills a breeze. Learn more

Thank you!

Rohan Bhatia
@rohanbhatia1