



# Fireside Chat: Optimize Your Business Continuity Strategy with Amazon Aurora Global Database

Sailesh Krishnamurthy, GM, Amazon Aurora, AWS

Steve Abraham, Principal Solutions Architect, AWS

# Amazon Aurora

Enterprise database at open source price,  
delivered as a managed service



---

**Speed** and **availability** of high-end commercial databases

---

**Simplicity** and **cost-effectiveness** of open source databases

---

Drop-in **compatibility** with MySQL and PostgreSQL

---


Simple **pay as you go** pricing

# Traditional Database Architecture

Databases are all about I/O

Design principles for 40+ years

- Increase I/O bandwidth
- Decrease number of I/Os




# Databases in the Cloud

## Compute & Storage have different lifetimes

- Instances fail and may be replaced
- Instances are shut down
- Instances are scaled up/down
- Instances are added to a cluster to scale out

Compute and Storage are best **decoupled** for scalability, availability, durability



The I/O bottleneck has  
moved to the network

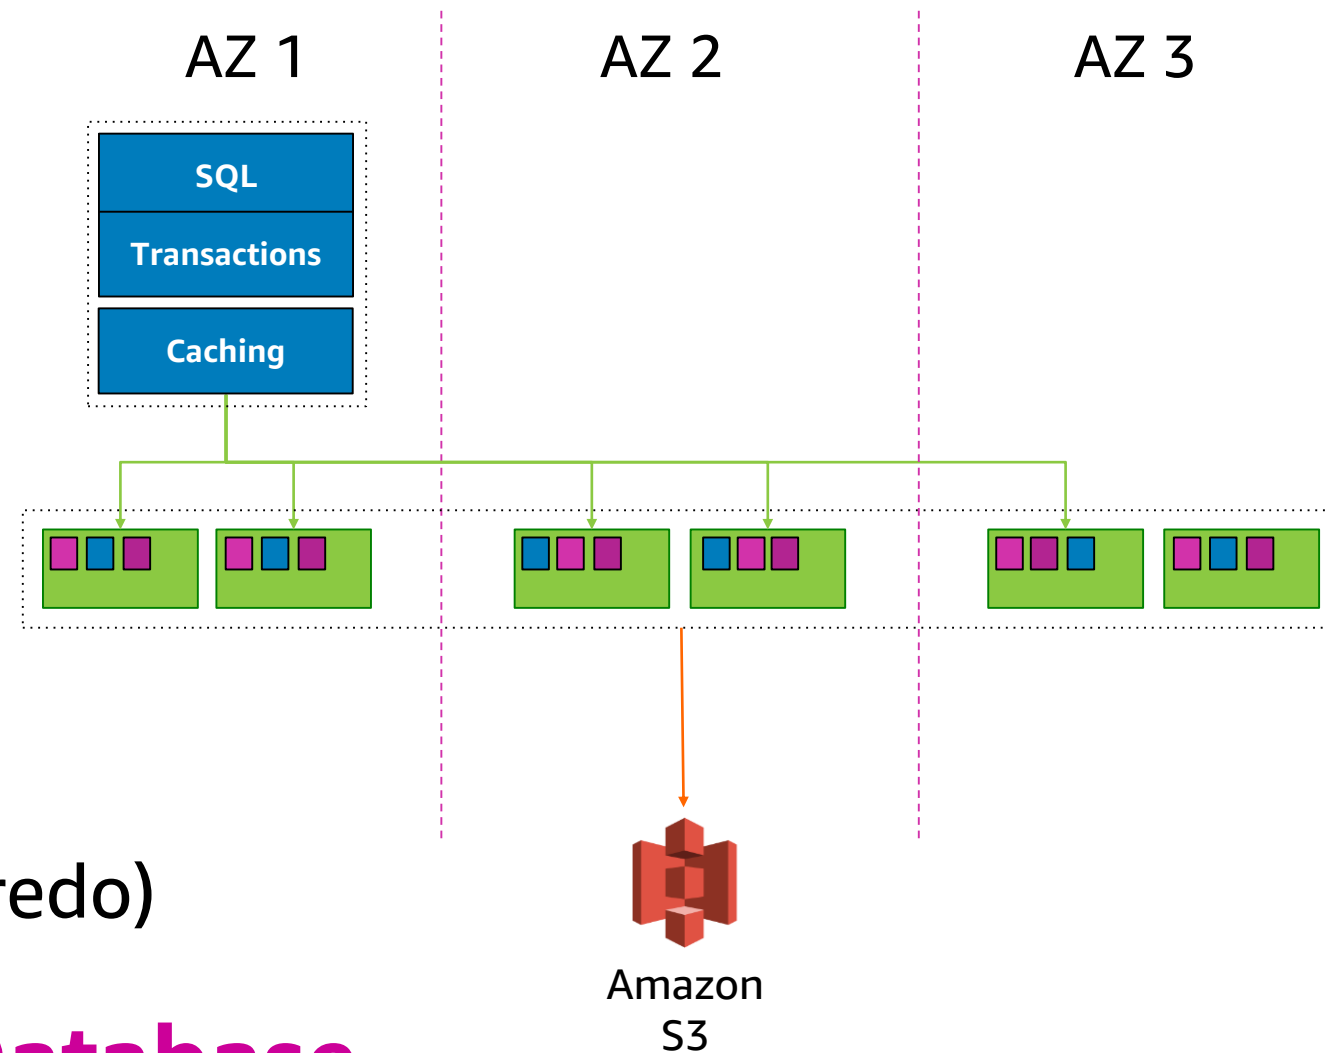
# Aurora: Offload Redo to Storage

## Database Tier

- Writes redo log records on network
- Use 4/6 quorum protocol
- No full data block writes for
  - Checkpointing, cache eviction, bg writes
- Push log applicator to storage

## Storage Tier

- Highly parallel scale out redo processing
- Generate database blocks on demand (redo)
- Materialize database blocks in background (redo)



**The Log is the Database**

# Uncorrelated and Independent Failures

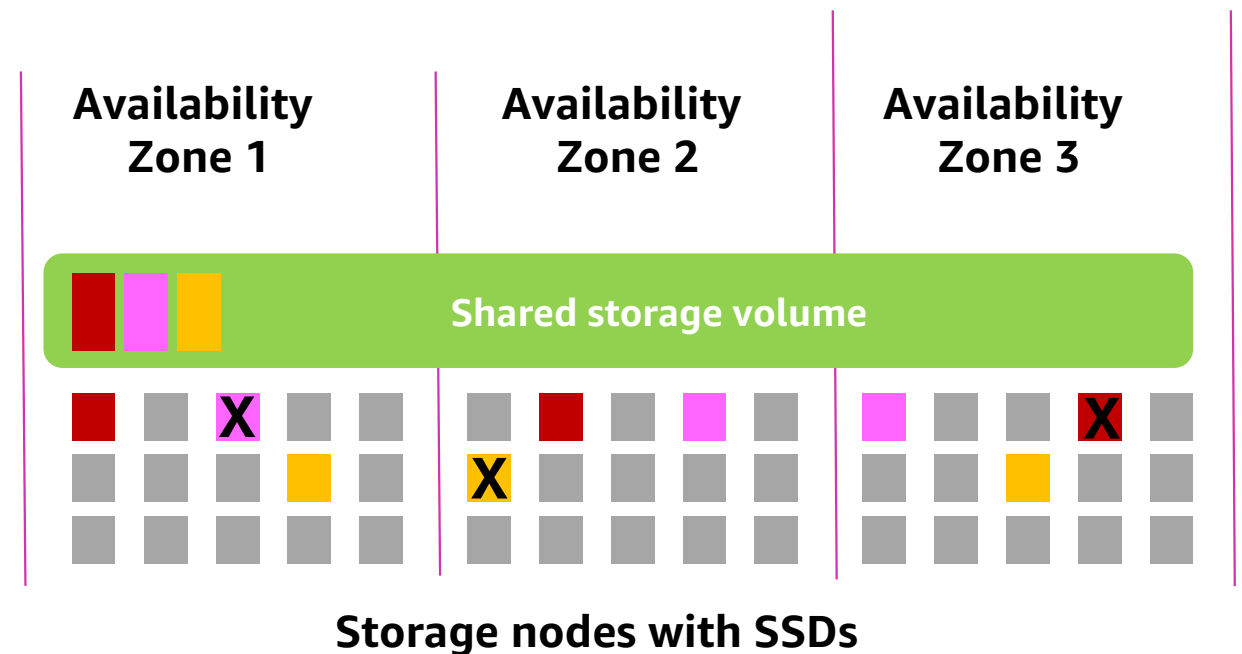
## At scale, continuous independent failures

- Constant **background radiation**
- Nodes, disks, switches all fail

## Replicate storage for resilience

### One common strawman:

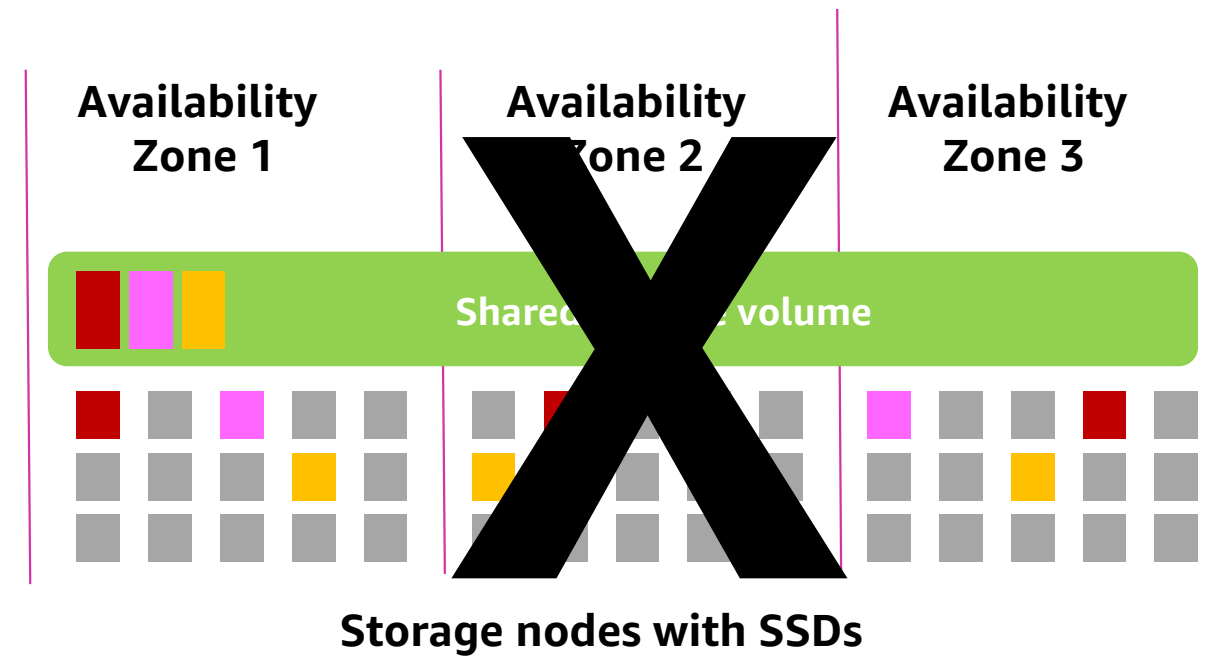
- Replicate 3-ways with 1 copy per AZ
- Use write and read quorums of 2/3



# What if an entire AZ is lost ?

## Boils down to losing 1 node

- Still have 2/3 nodes
- Can establish quorum
- No data loss 😊

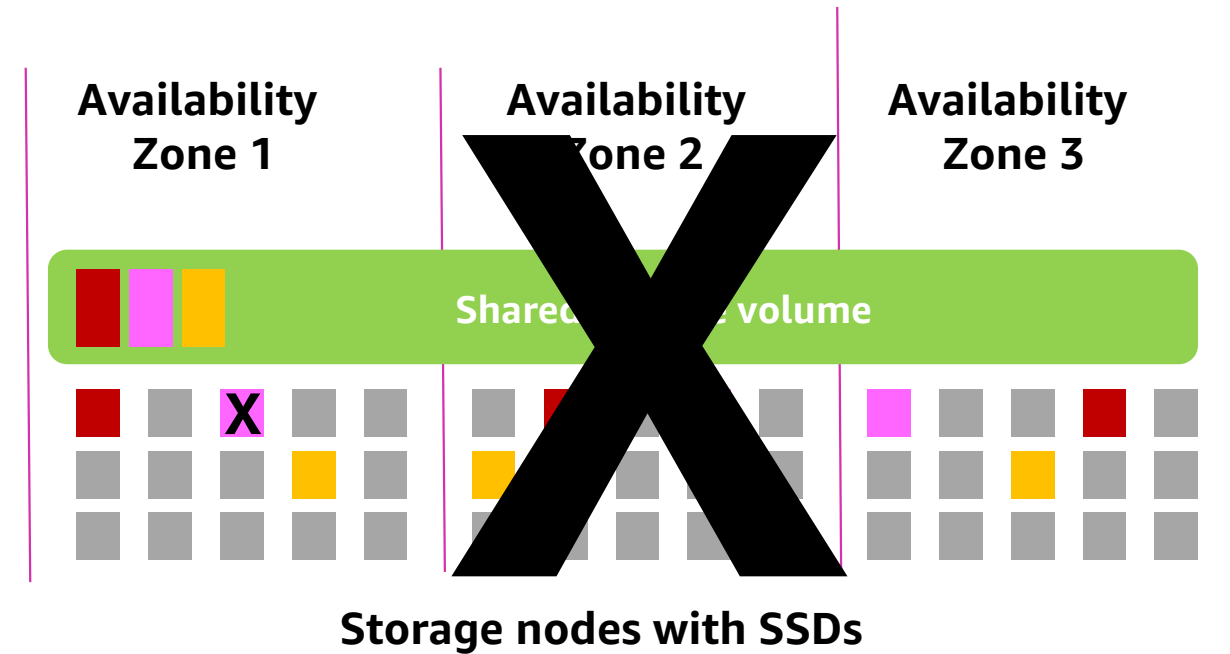




# What if additional node is also lost ?

## AZ+1 model: correlated failures!

- Lose 2/3 nodes
- Lose quorum
- Lose data 😞



# Aurora tolerates AZ+1 failures

## Replicate 6-ways with 2 copies per AZ

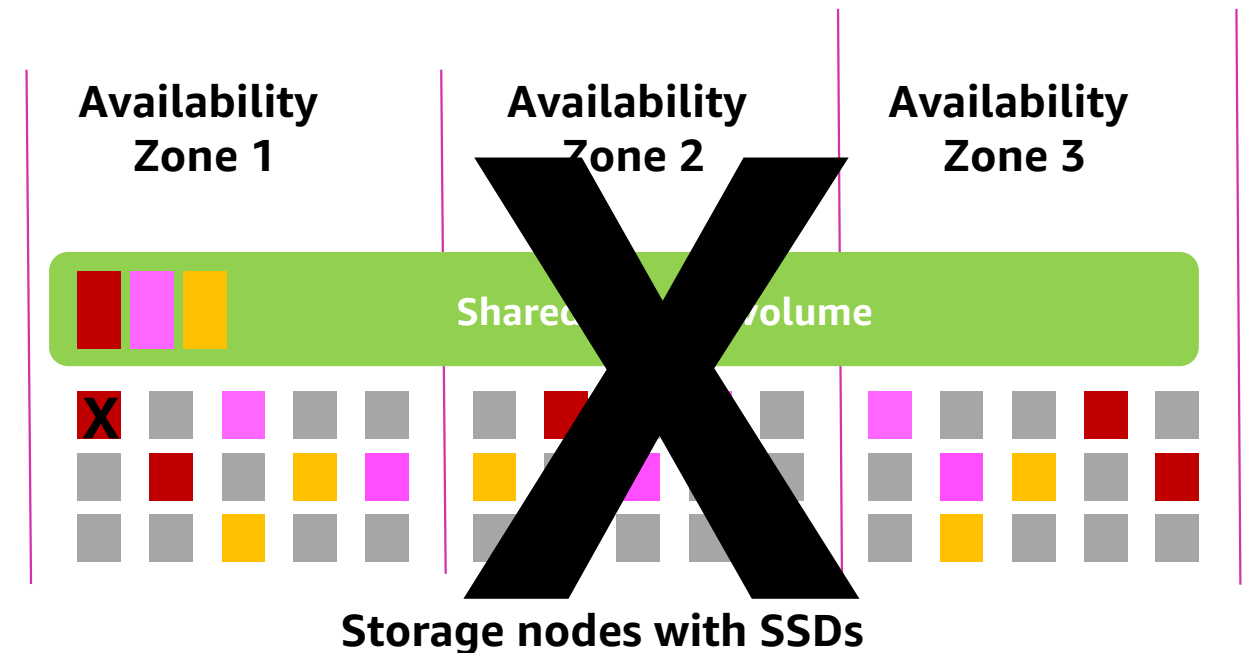
- Write quorum of 4/6
- Read quorum of 3/6 (only for repair)

## What if there is an AZ failure ?

- Still have 4/6 nodes
- Maintain write availability

## What if there is an AZ+1 failure ?

- Still have 3 nodes (read/repair quorum)
- No data loss 😊
- Rebuild failed node by copying from one of other 3
- Recover write availability 😊 😊



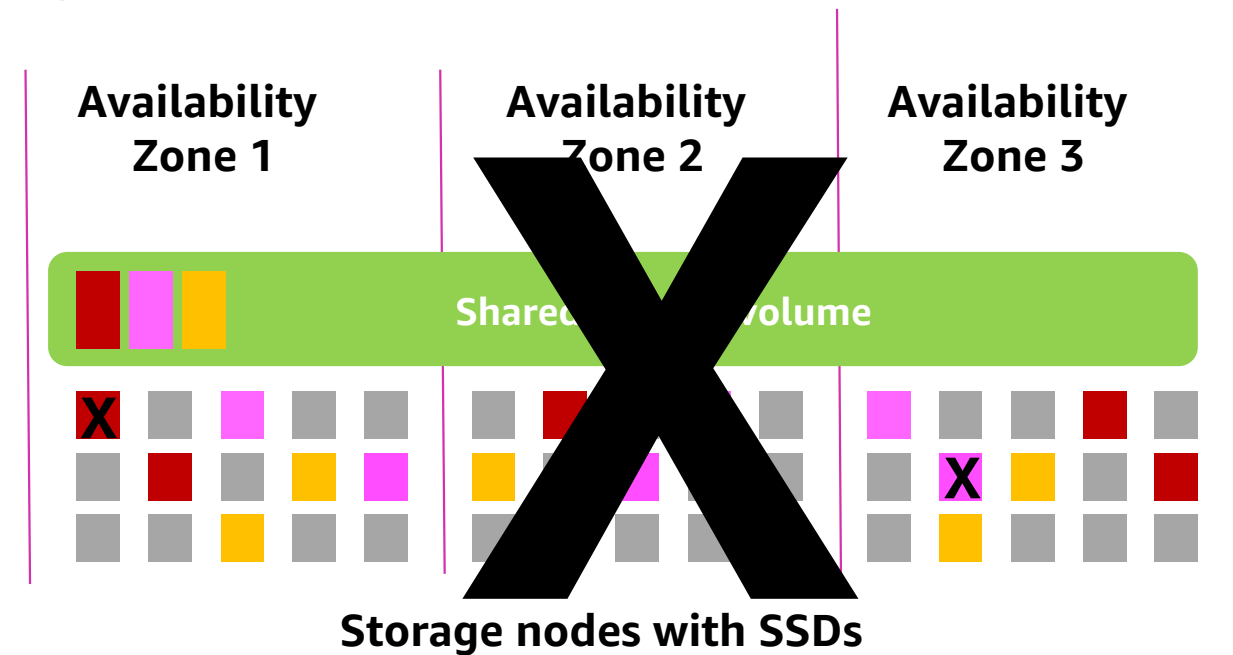
# Is a 4/6 quorum sufficient for AZ+1 ?

Depends on repairing a failed node before AZ+1 becomes AZ+2 (double-fault)

P(AZ+2) in repair interval is function of MTTF

- Can only reduce MTTF & P(AZ+2) so much

Instead try to reduce repair interval (MTTR)



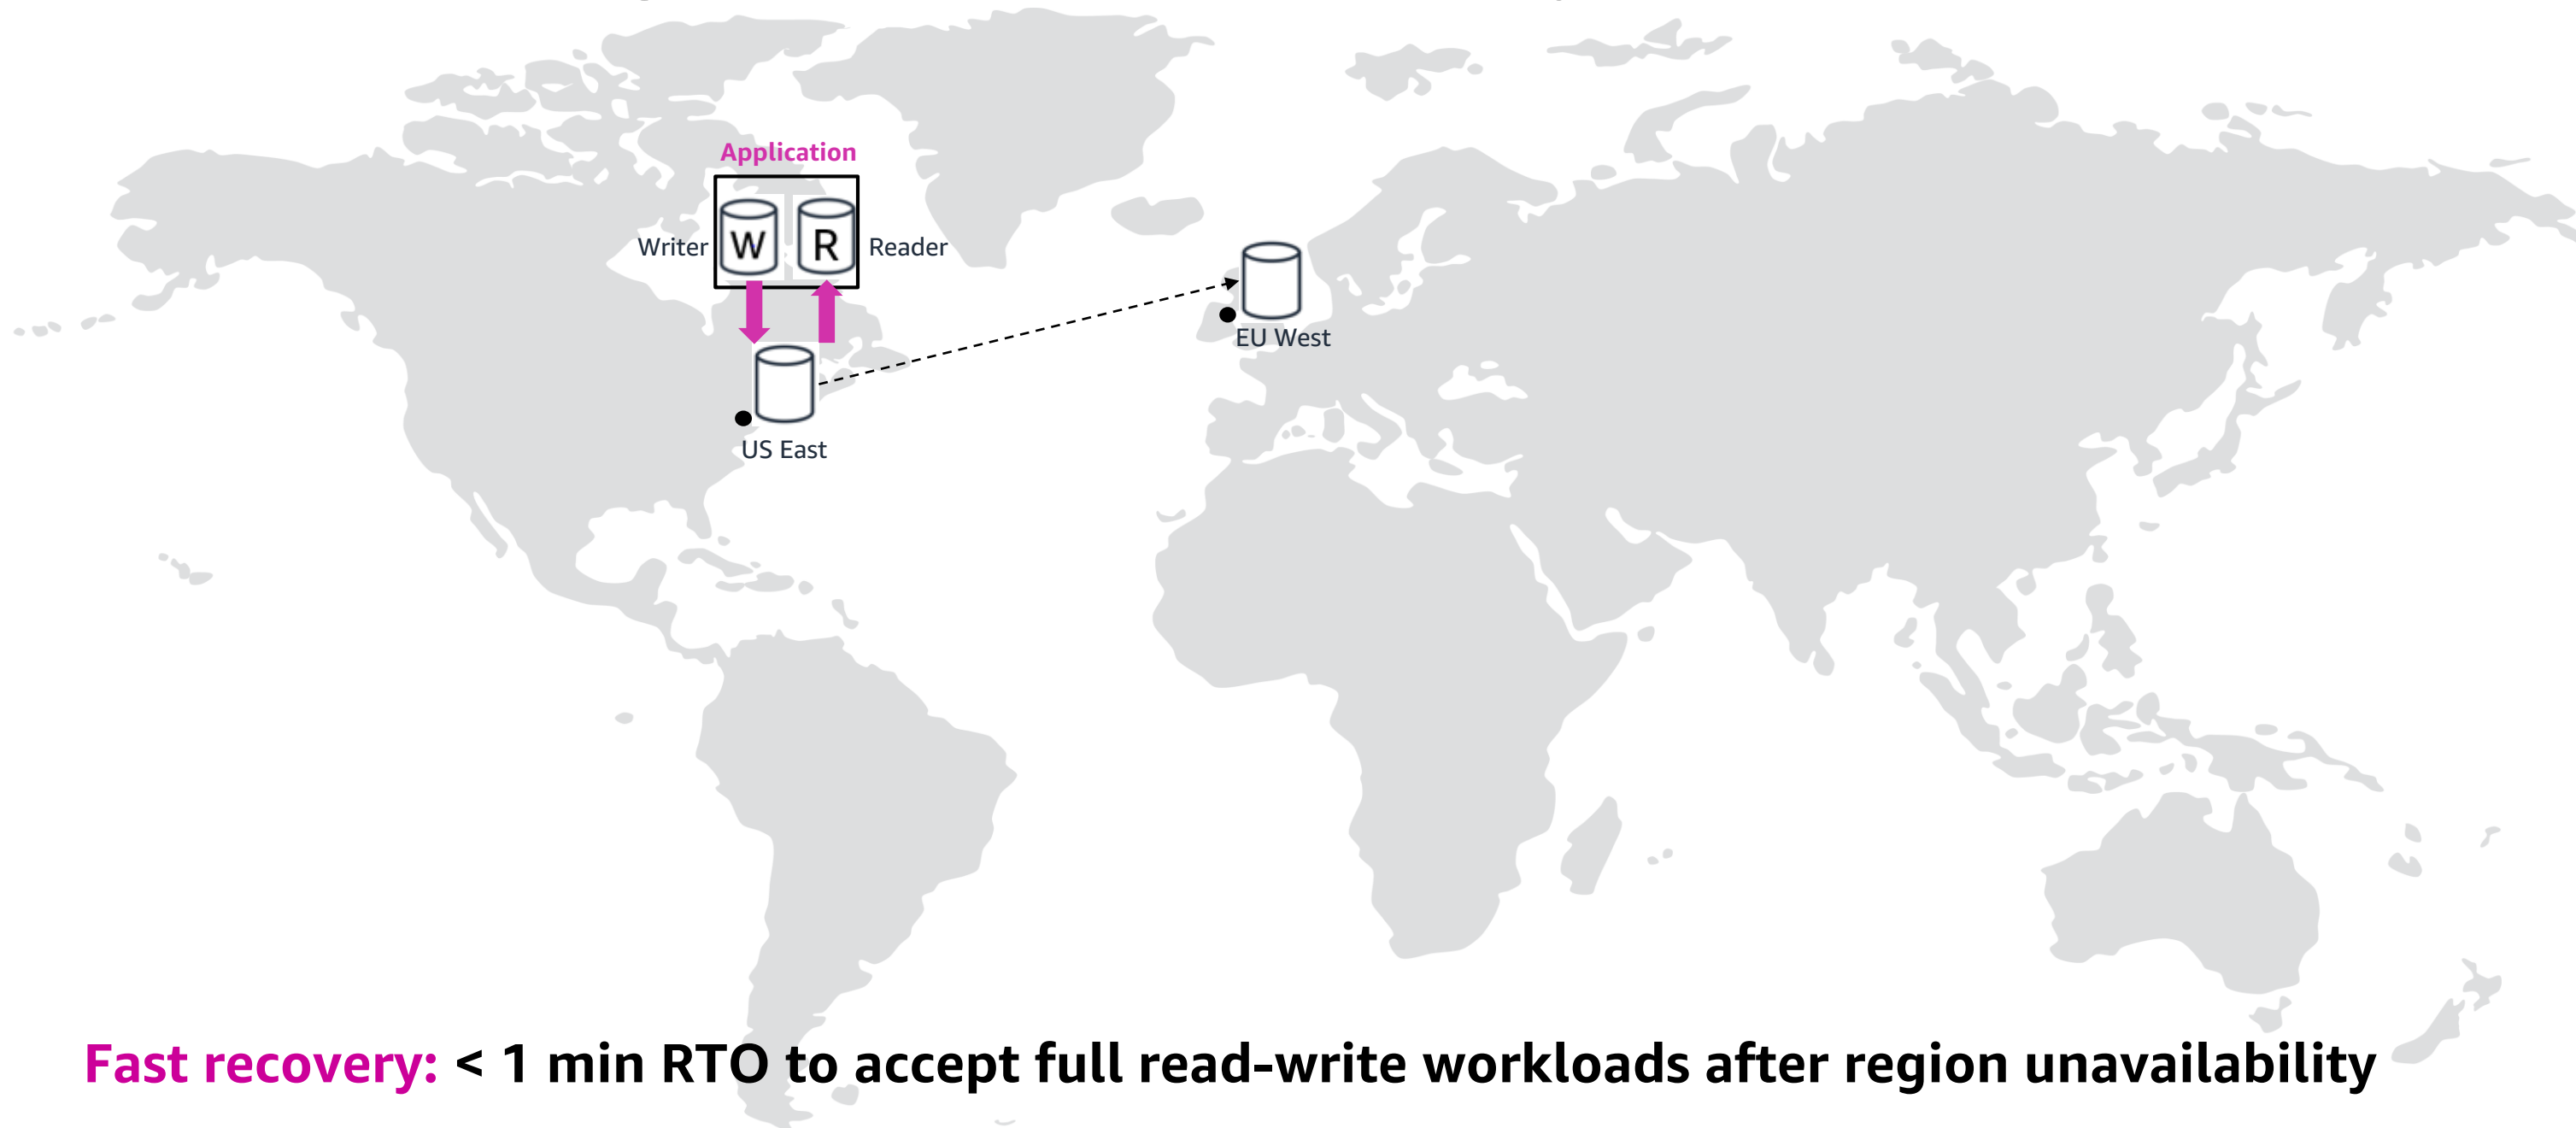
# Aurora Global Database

Faster disaster recovery and enhanced data locality

- Promote remote readers to a master for faster recovery in the event of disaster
- Bring data close to your customer's applications in different regions

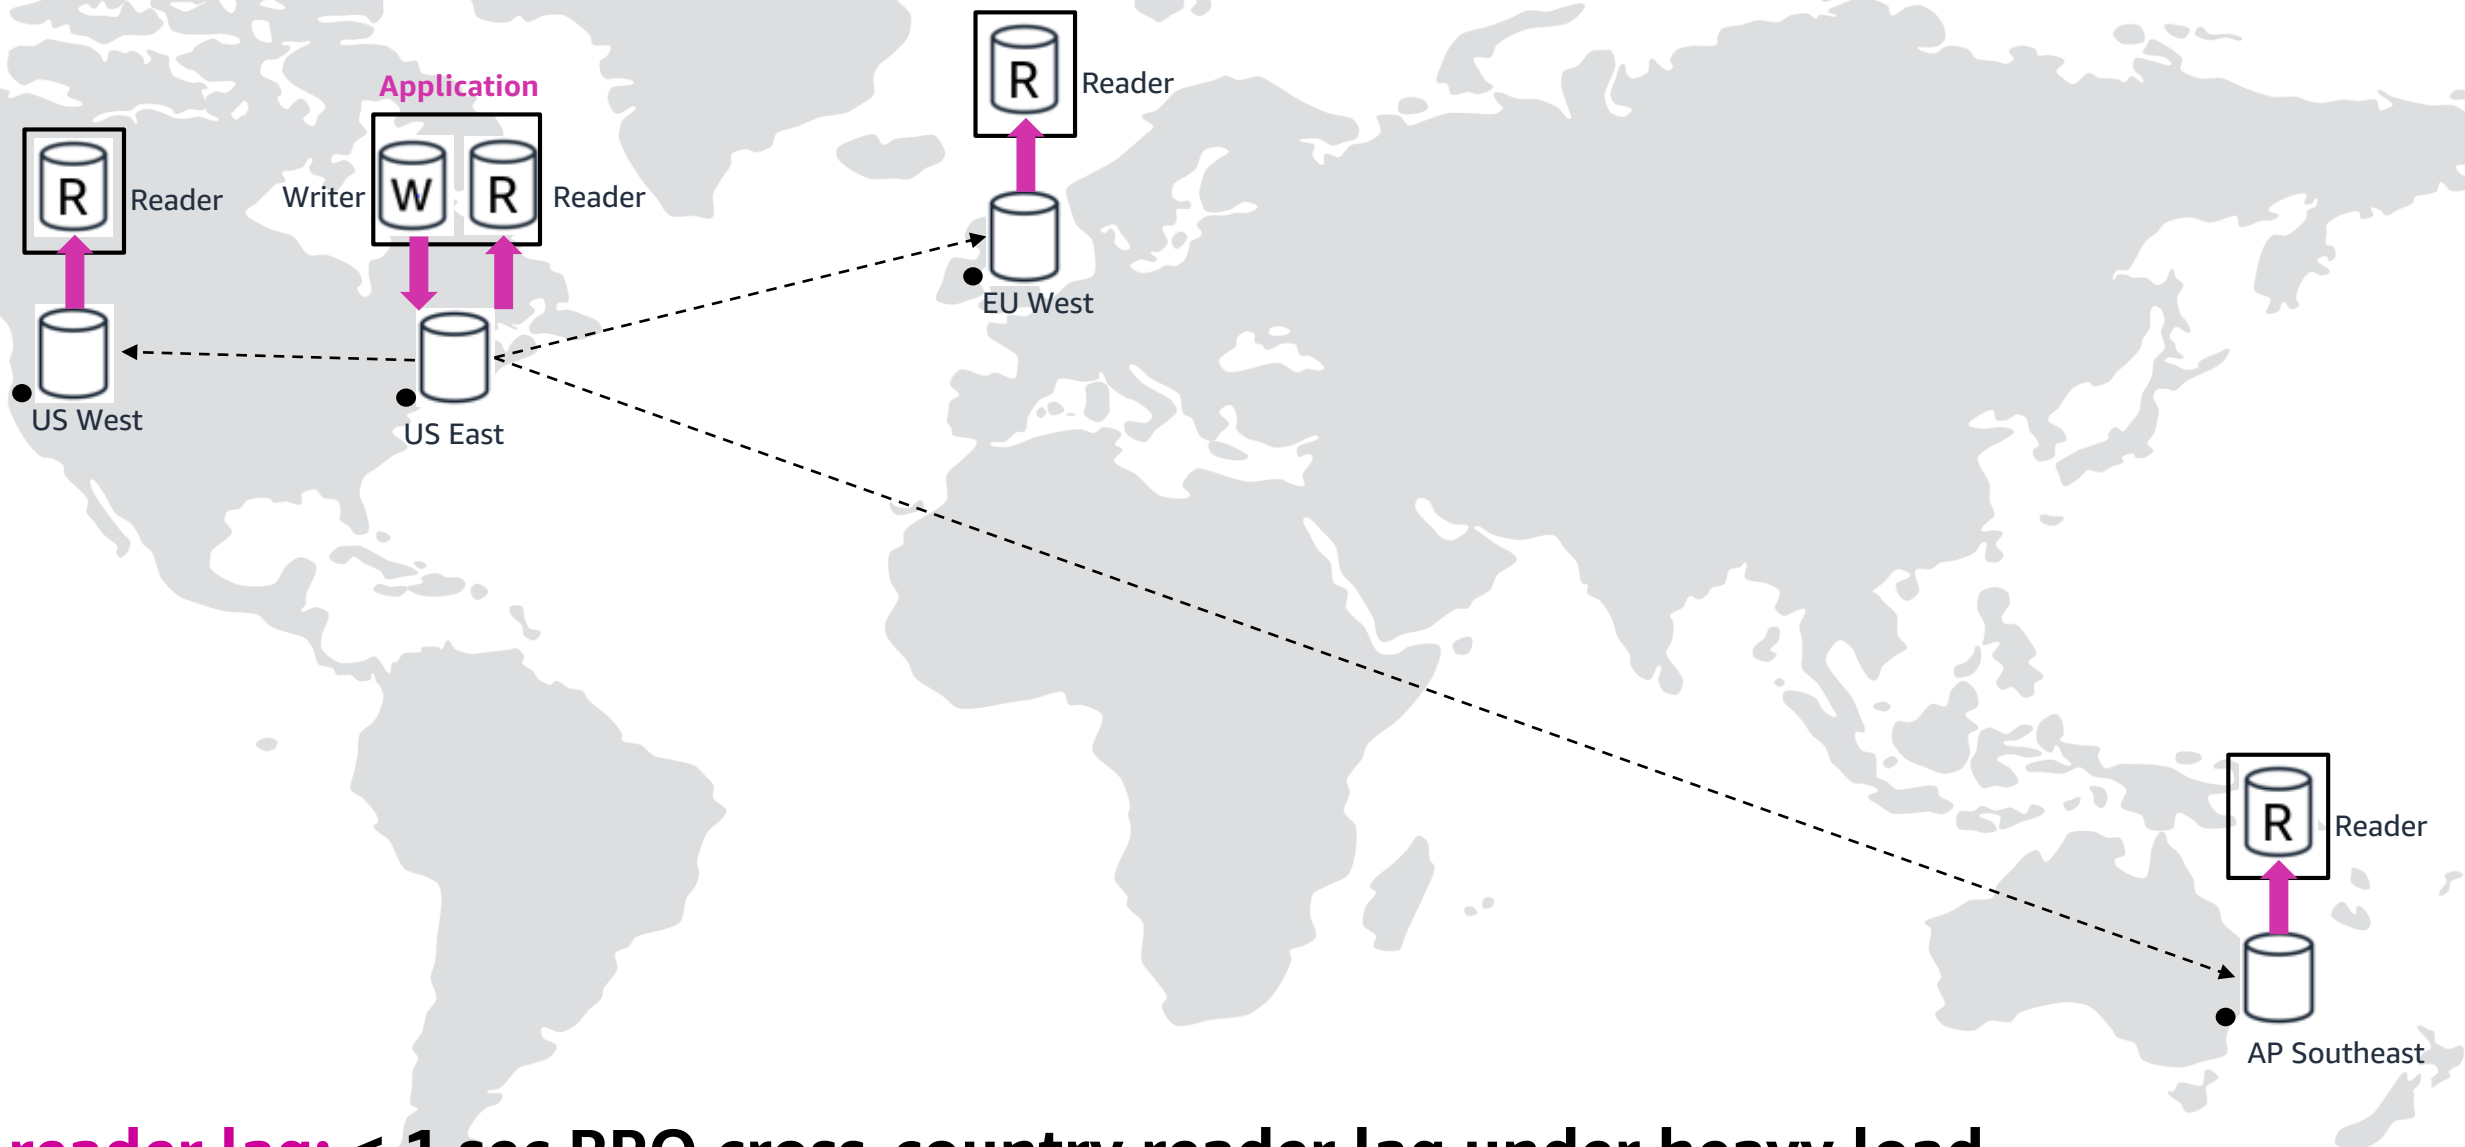


# Fast cross-region disaster recovery



**Fast recovery:** < 1 min RTO to accept full read-write workloads after region unavailability

# Global reads with low replication latency



**Low remote reader lag: < 1 sec RPO cross-country reader lag under heavy load**



# Demo 1: Create an Aurora Global Database

RDS · AWS Console

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

Apps Saved VA Site

aws

Services ▾

Resource Groups ▾

✦

🔔 sabraham @ abrsteve ▾ N. Virginia ▾ Support ▾

Amazon RDS

Dashboard

Databases

Query Editor

Performance Insights

Snapshots

Automated backups

Reserved instances

Subnet groups

Parameter groups

Option groups

Events

Event subscriptions

Recommendations

RDS > Databases

Databases

☒ Group resources

🔄

Modify

Actions ▾

Restore from S3

Create database

< 1 >

⚙️

	DB identifier	Role	Engine	Region & AZ	Size	Status
<input type="radio"/>	<div>wordpress-global-cluster</div>	Global	Aurora MySQL	2 regions	2 clusters	✔️ Ava
<input type="radio"/>	<div>wordpress-source-cluster</div>	Primary	Aurora MySQL	us-east-1	1 instance	✔️ Ava
<input type="radio"/>	<div>wordpress-source-instance</div>	Writer	Aurora MySQL	us-east-1c	db.r5.large	✔️ Ava
<input type="radio"/>	<div>wordpress-target-cluster</div>	Secondary	Aurora MySQL	us-east-2	1 instance	✔️ Ava
<input type="radio"/>	<div>wordpress-target-instance</div>	Reader	Aurora MySQL	us-east-2b	db.r5.large	✔️ Ava

Feedback

English (US)

© 2008 - 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

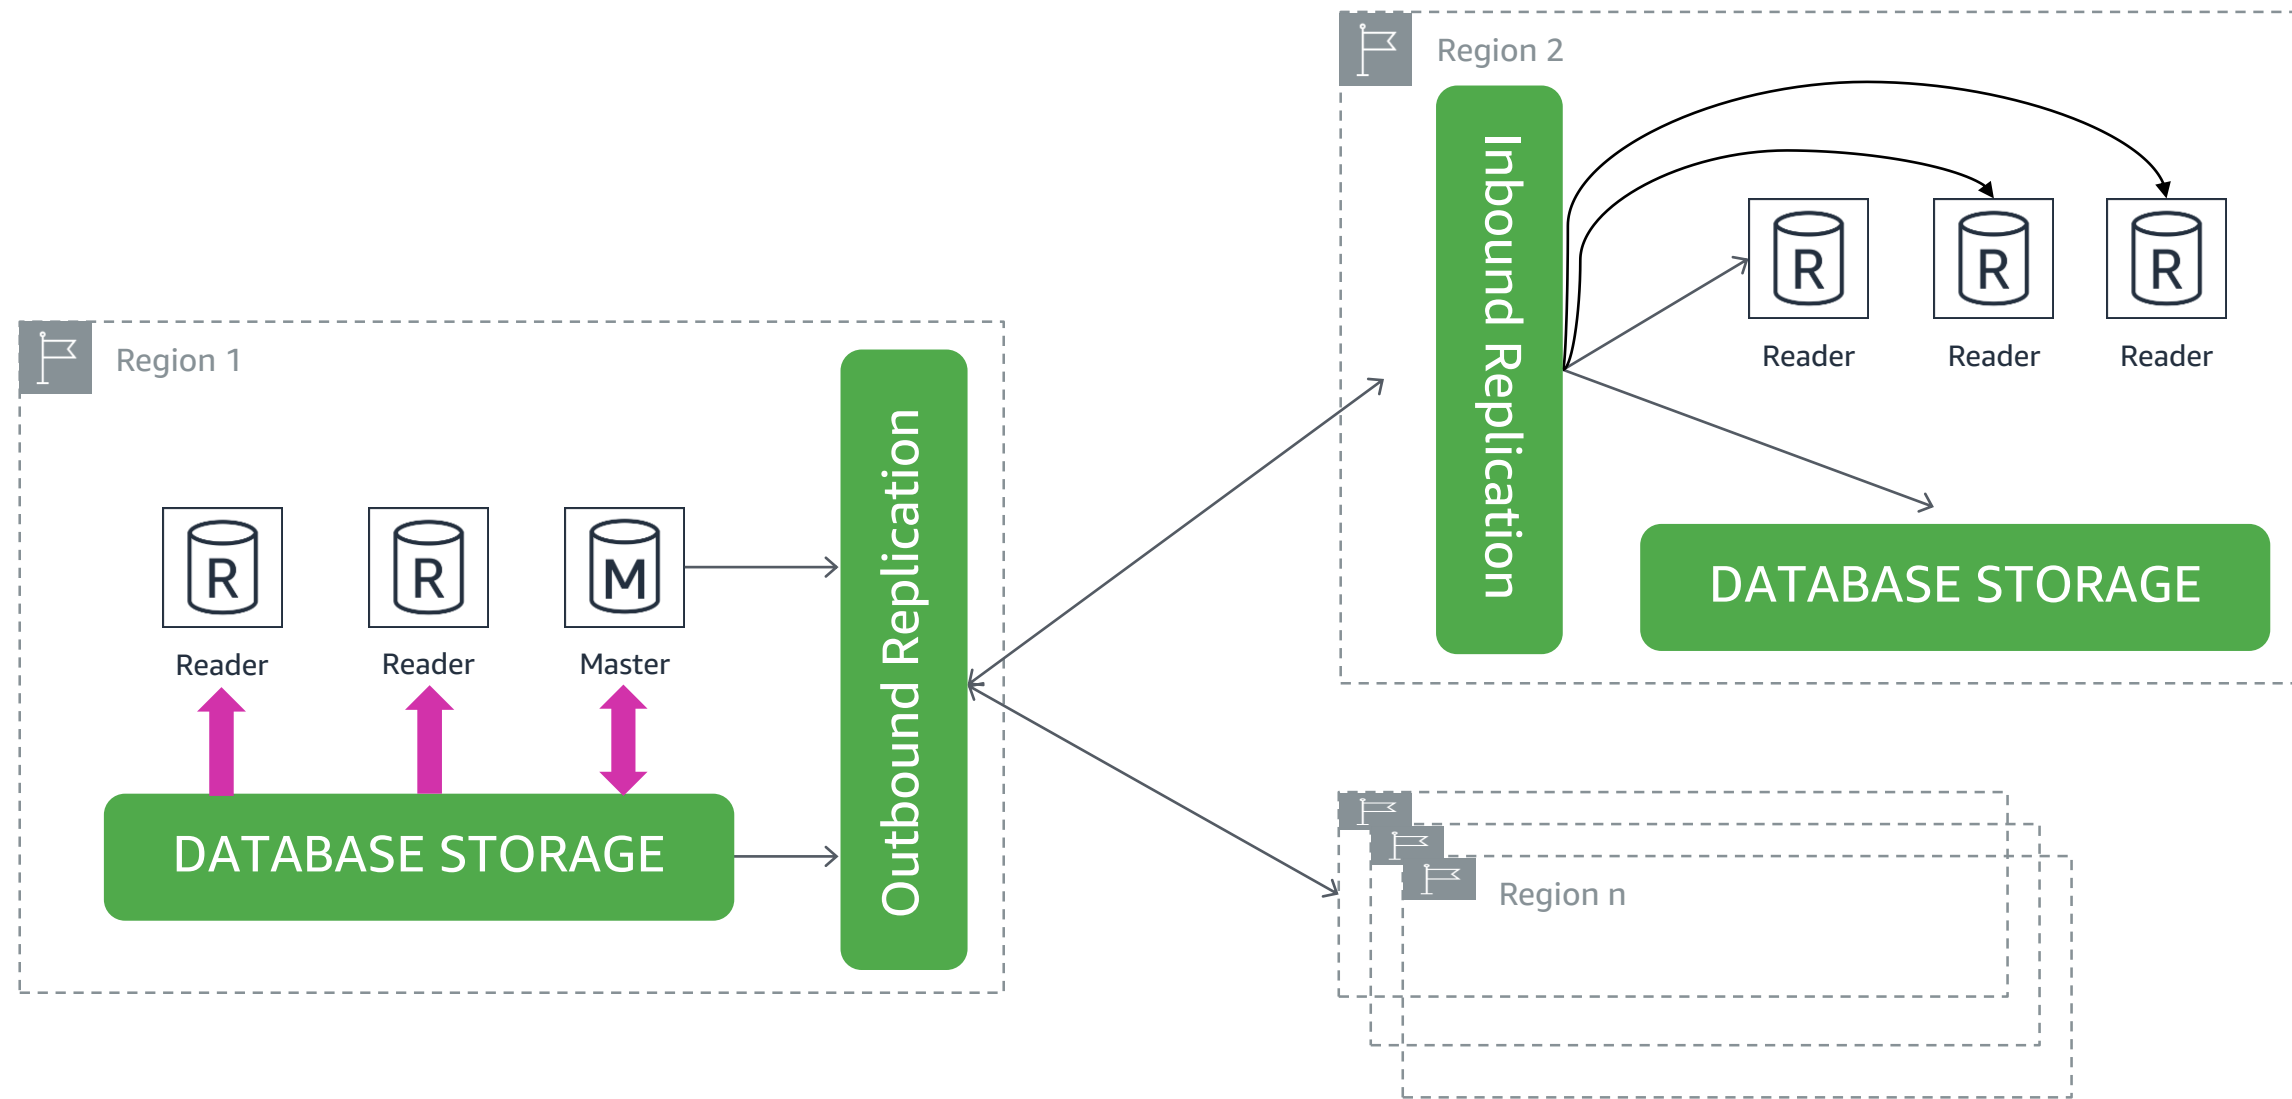
Privacy Policy

Terms of Use






# How does it work?




**High throughput:** Up to 200K writes/sec – negligible performance impact

# Global Database performance benefits

## Logical vs. Physical Replication



Logical replication



Physical replication

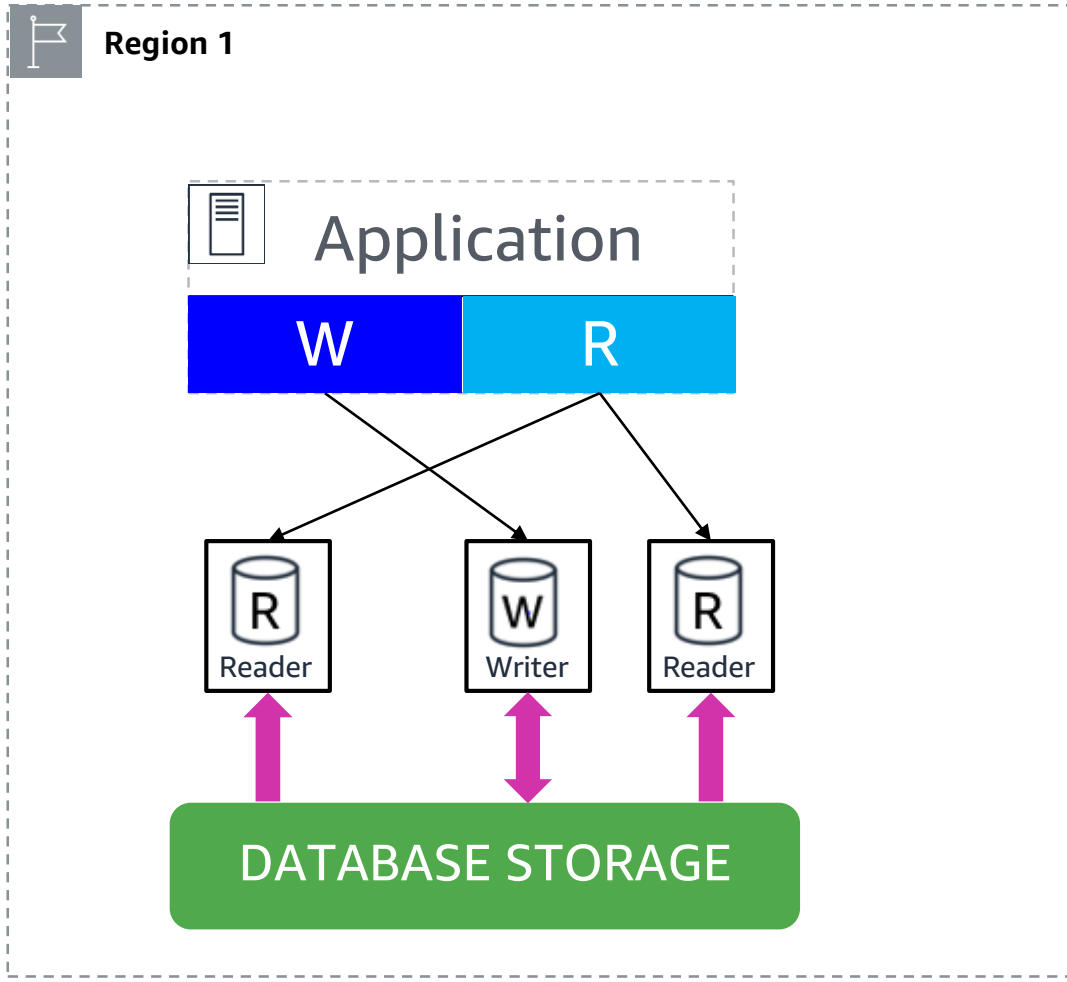
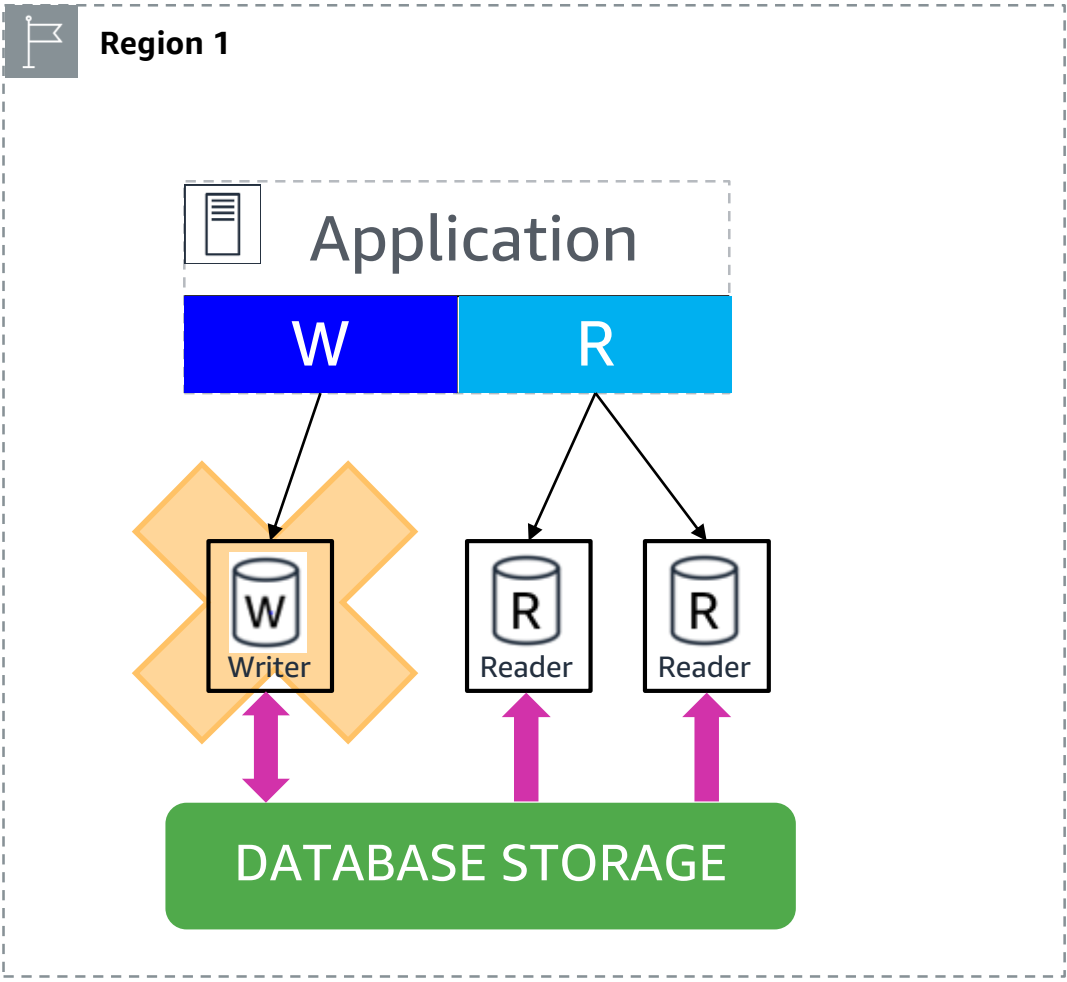




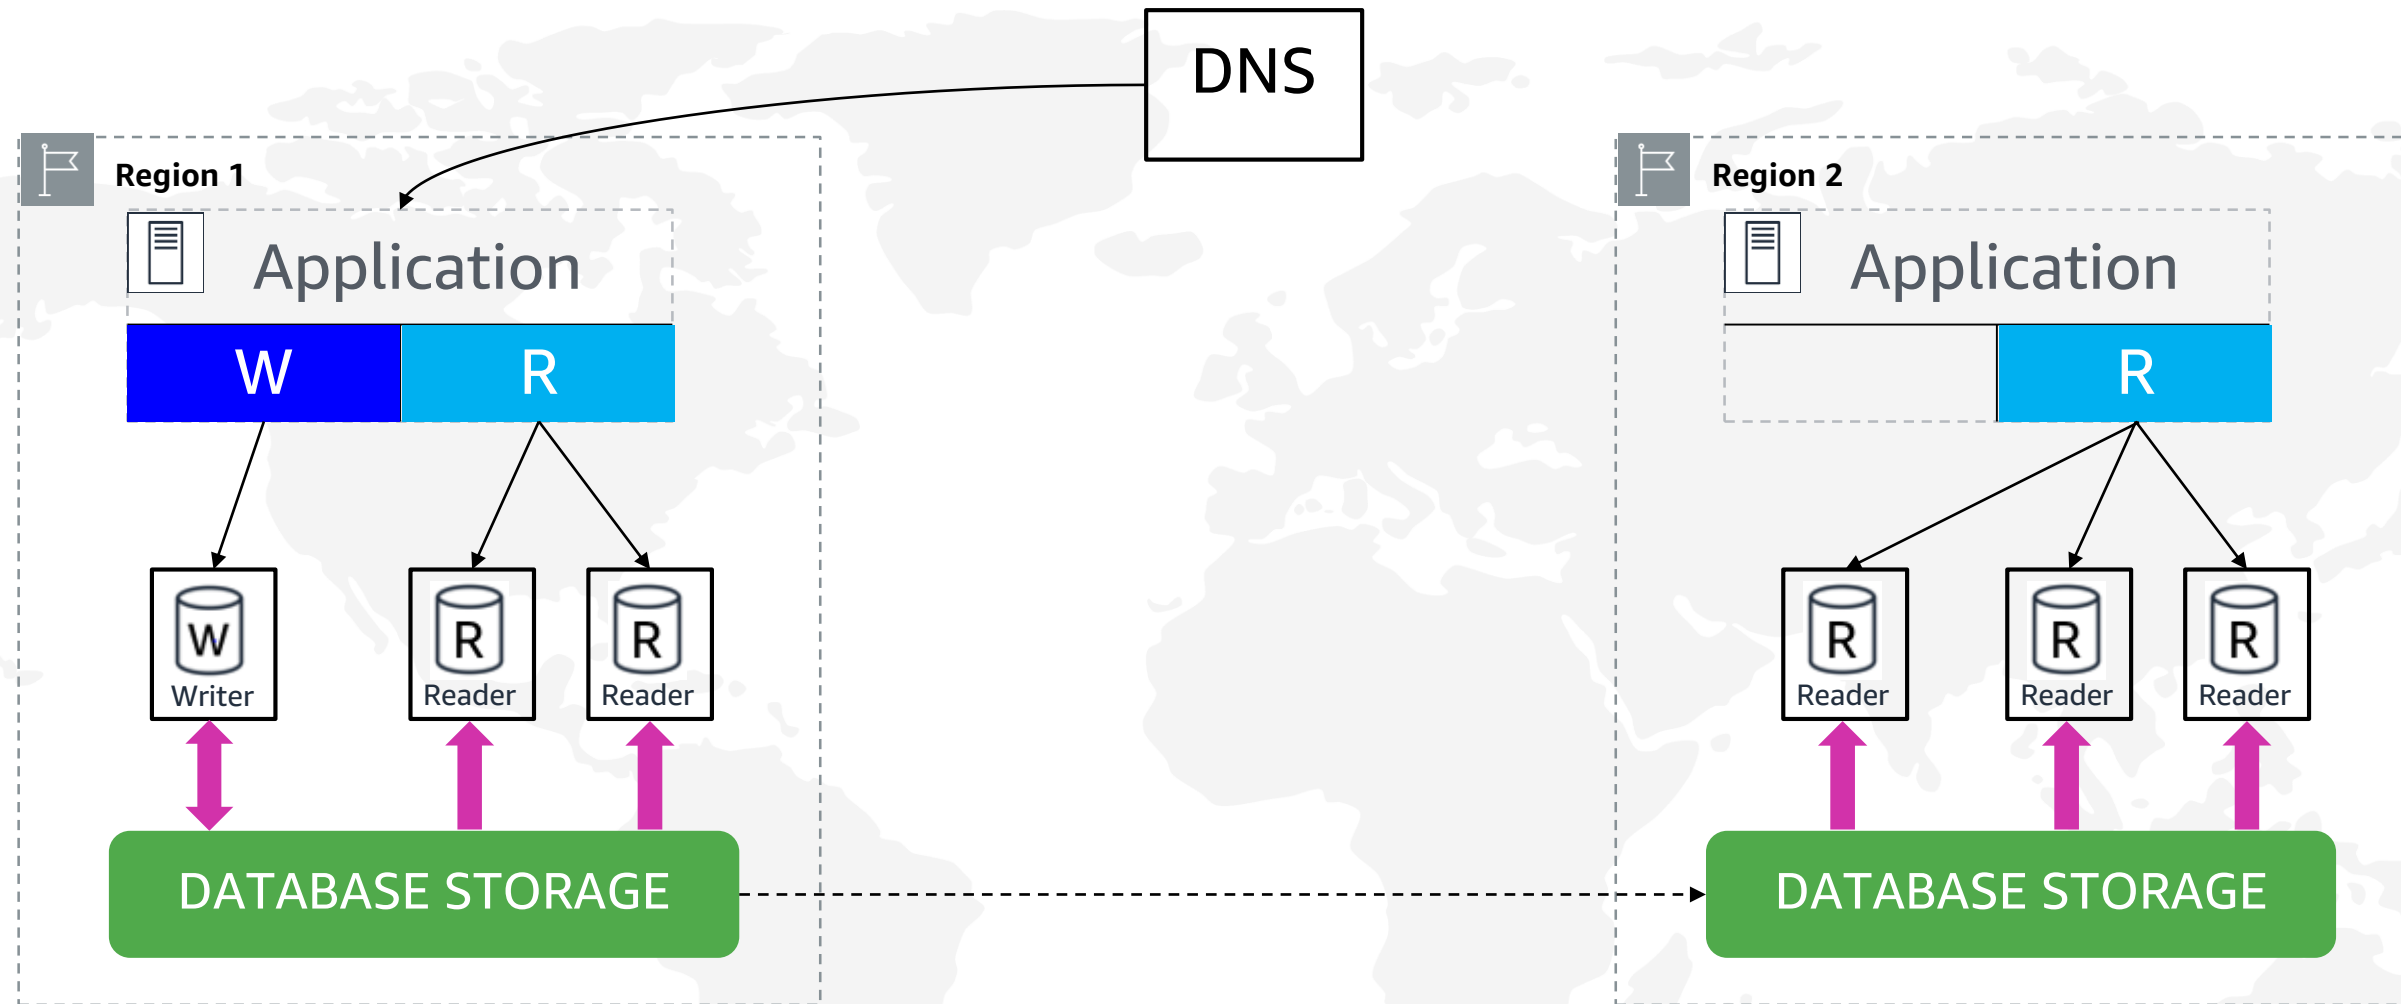
# Demo 2: Performance tests



# Failover within a single region

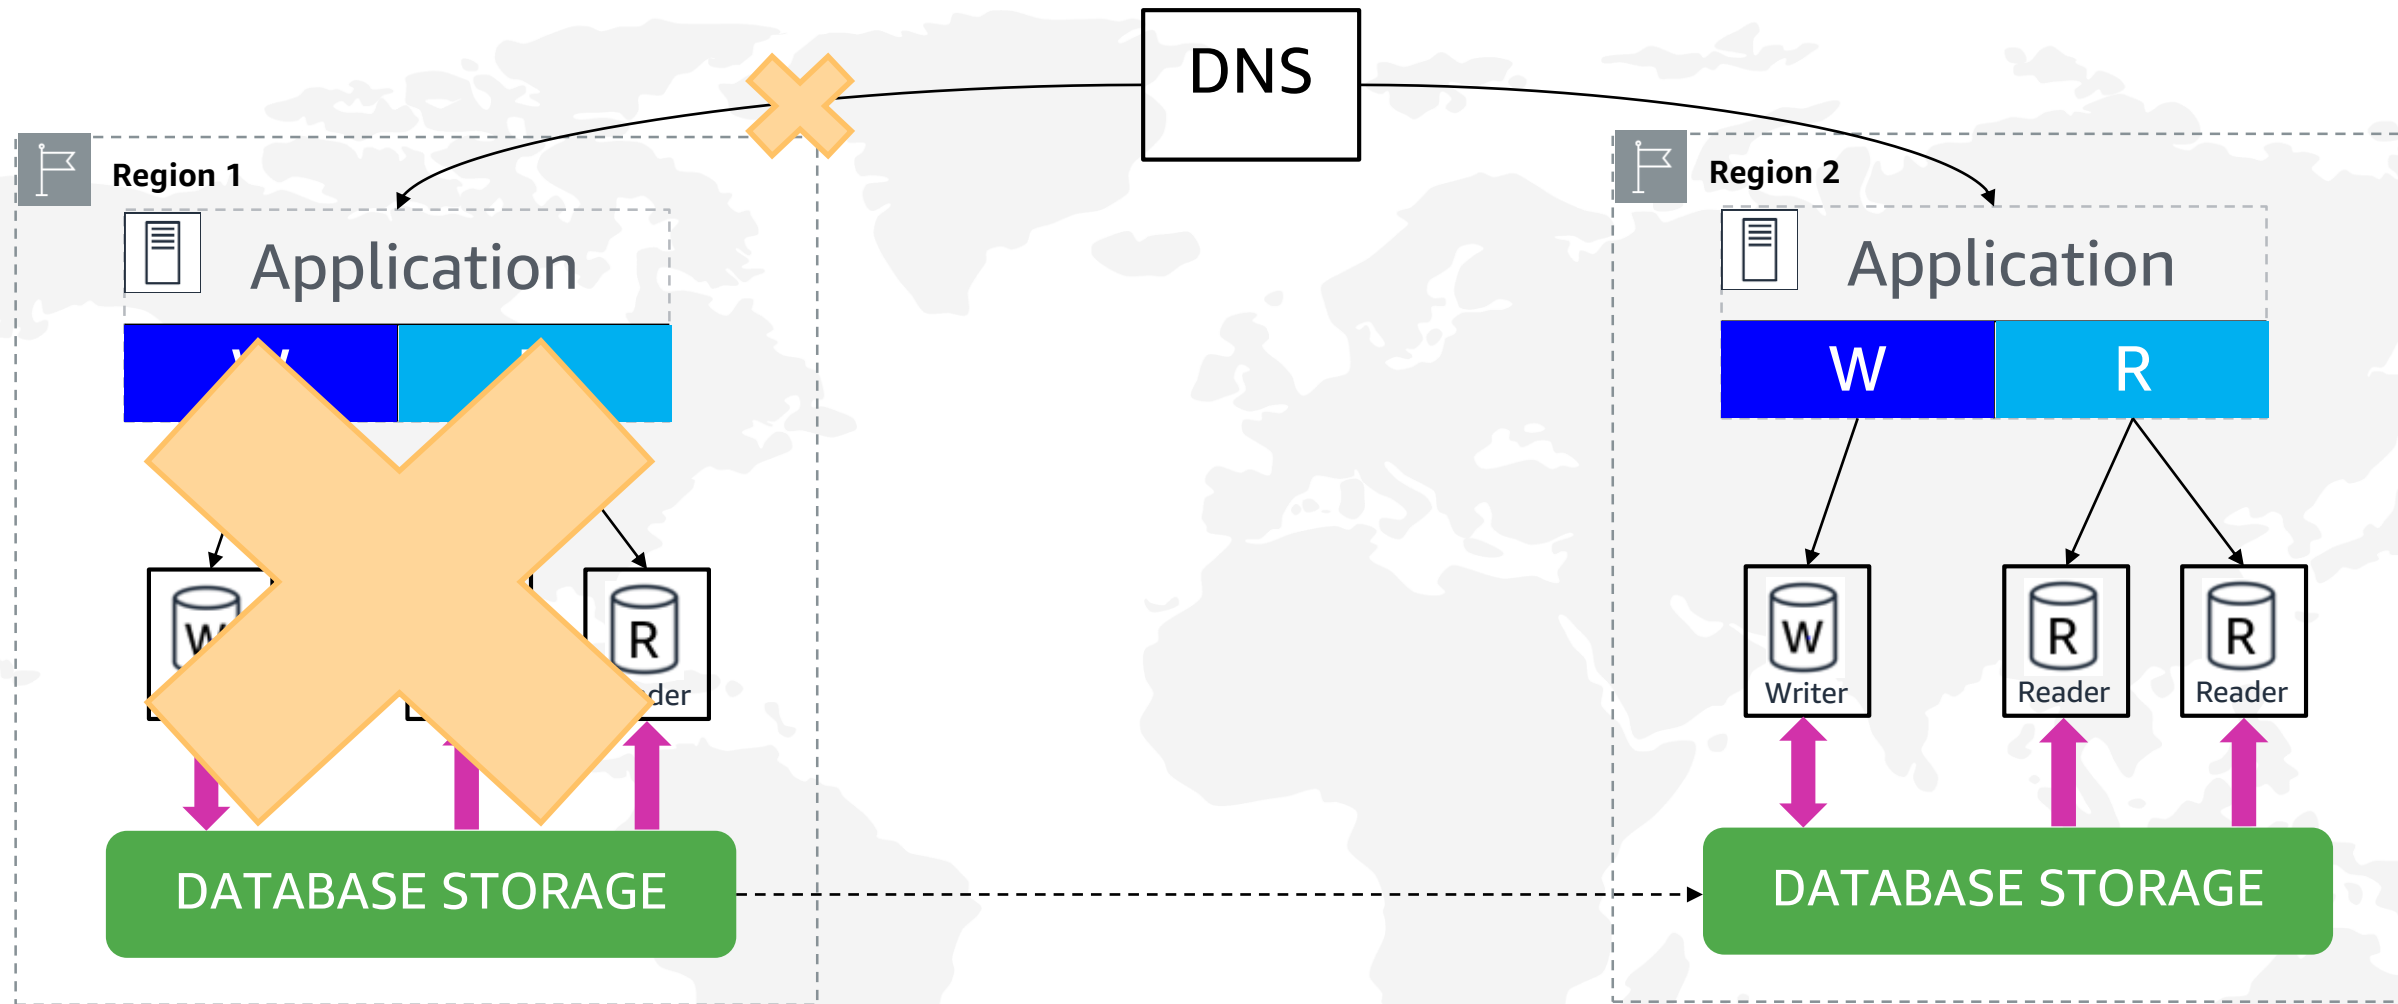


# Active-active cross-region setup with Aurora Global Database



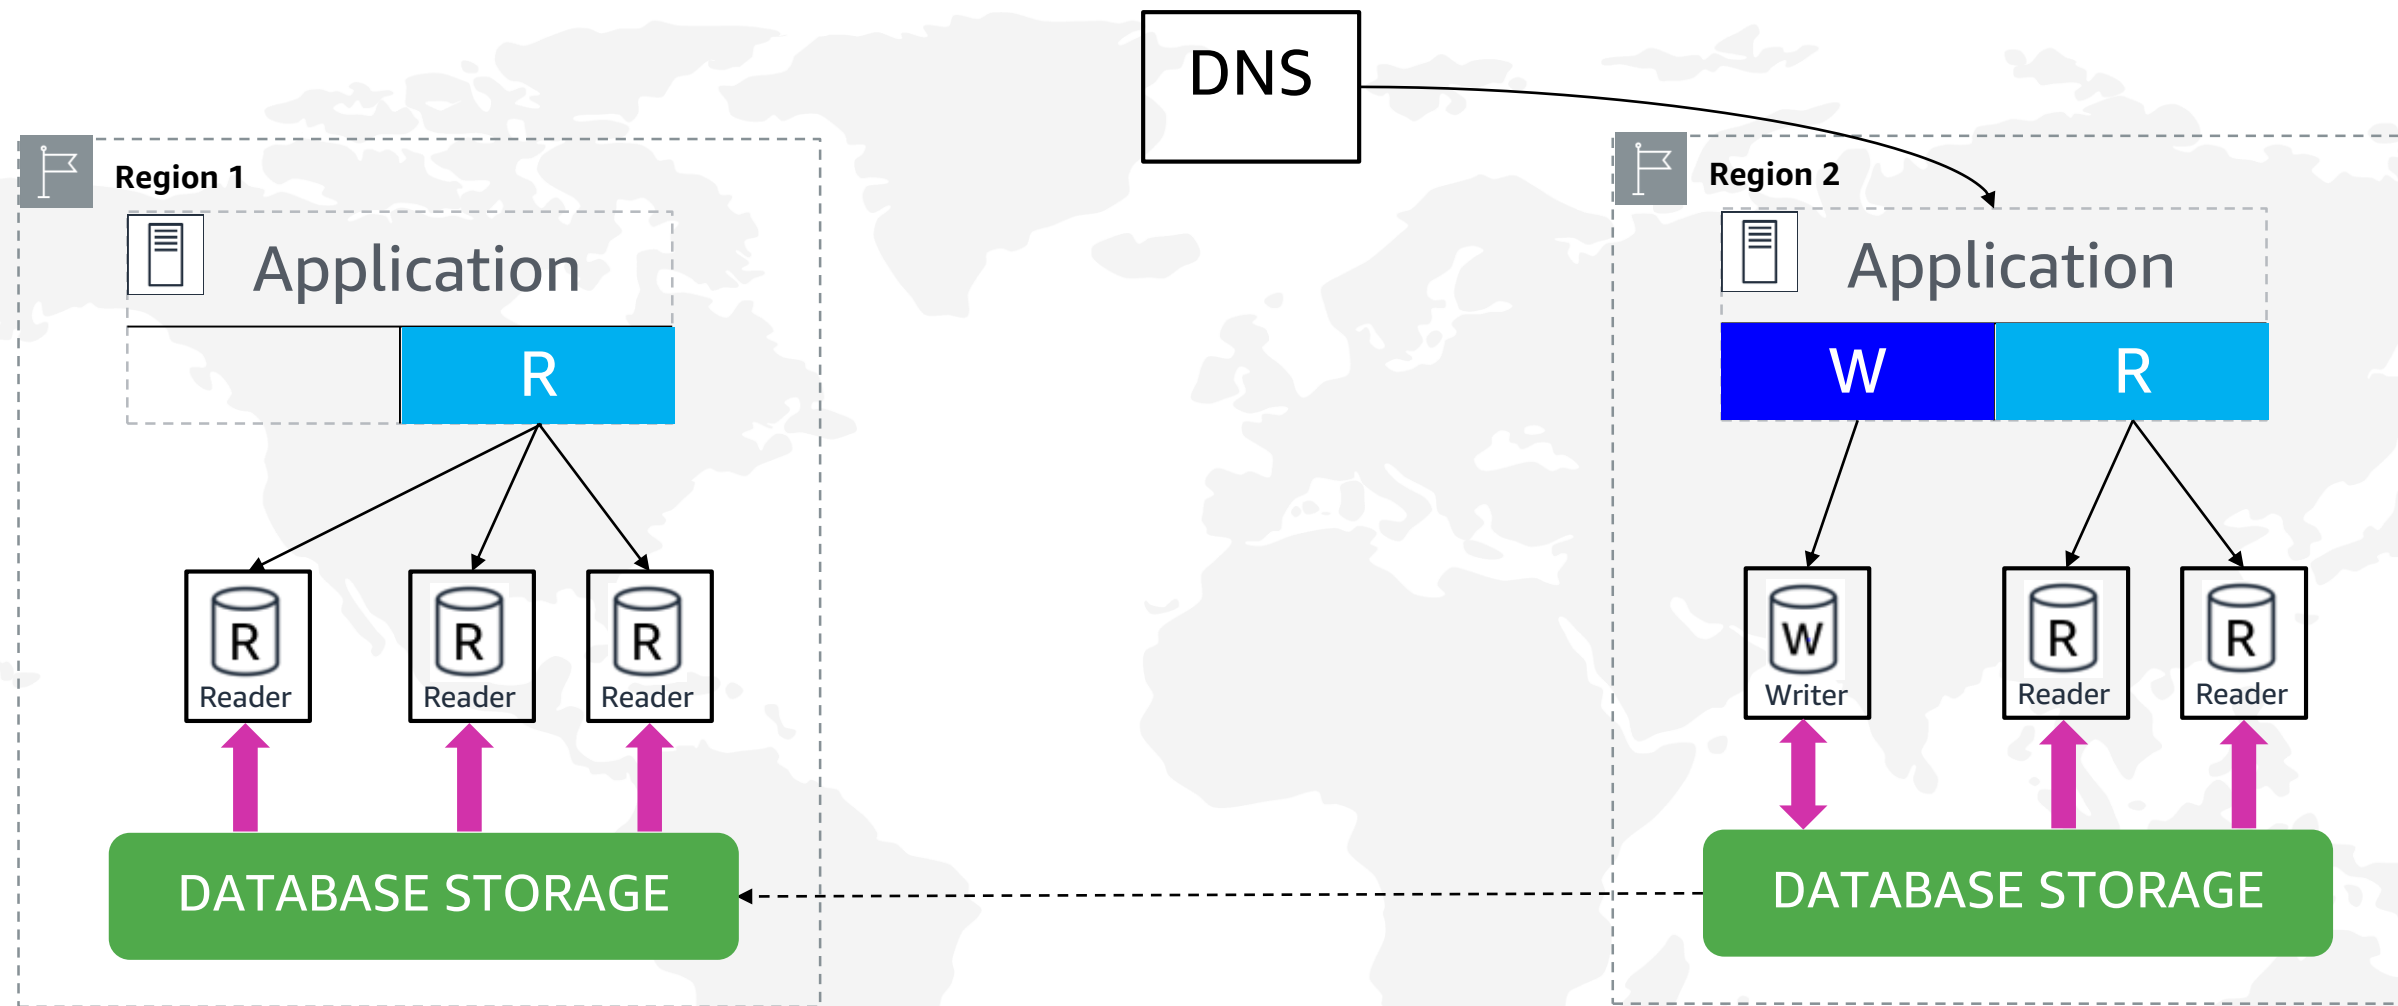
**Version of the application stack is set up in another region  
to provide disaster recovery and serve fast local reads**

# Failover with Aurora Global Database



**Application stack can be spun up in the secondary region and traffic can begin to be routed there (via DNS) with the promoted writer serving write requests within 1 minute.**

# Active-active cross-region setup can be restored







# Demo 3: Aurora Global Database failover

RDS · AWS Console x user's Blog! – Just another Wordi x Database Error x +

← → ↻ ⓘ Not secure | 52.91.20.77 ★ 🗖 📄 📱 📺 | Ⓢ ⋮

📱 Apps 📁 Saved 🌐 VA Site 🌐 OH Site

🌐 user's Blog! ✍️ Customize 🔄 9 💬 0 ➕ New Howdy, user 🔍


user's Blog! – Just another WordPress site

---

# Hello world!

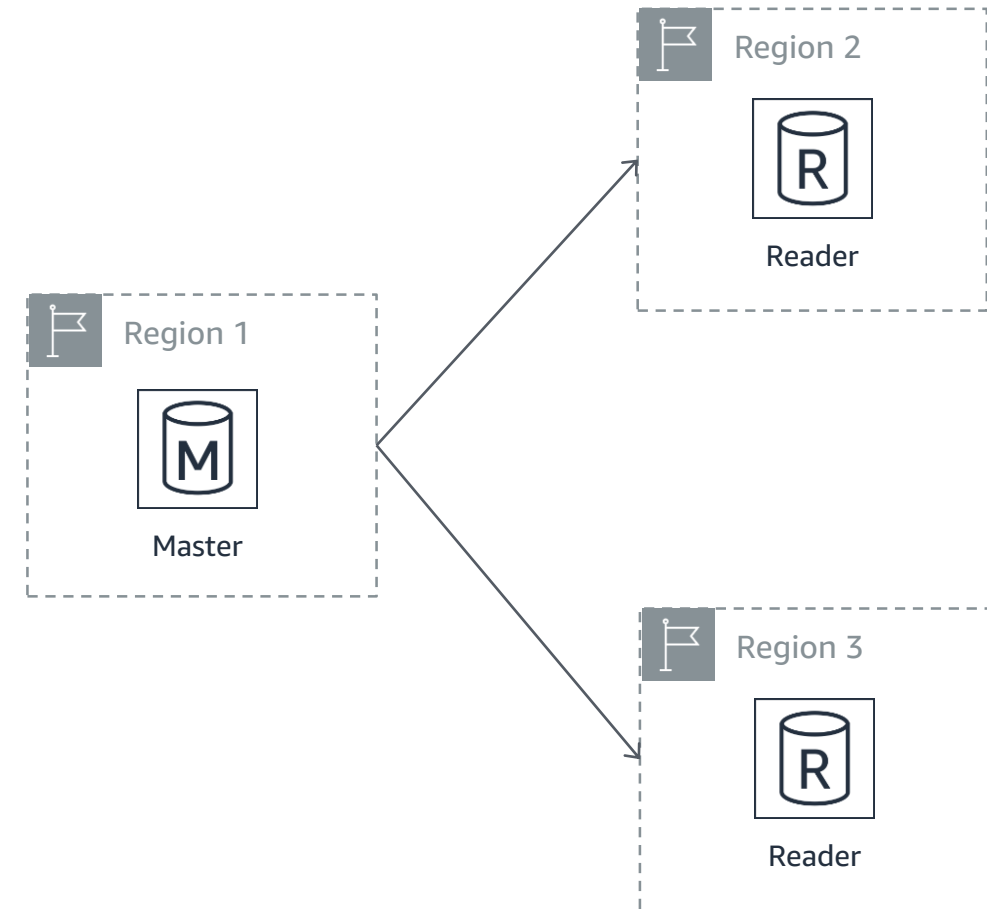
Welcome to WordPress. This is your first post. Edit or delete it, then start writing!

👤 user 🕒 April 30, 2019 📁 Uncategorized 💬 1 Comment ✎ Edit

🖱️  Manage

# 2019 roadmap

- Support in all Aurora regions – Available now
- Performance Insights – Available now
- >2 secondary regions – Coming soon
- Available on Aurora MySQL 5.7 – Coming soon
- In-place conversion (from non-Global DB to Global DB) – Coming soon



# Thank you!

sailesk@amazon.com  
abrsteve@amazon.com