MySQL Options on AWS: Self-Managed, Managed, Serverless

Yoav Eilat
Sr. Product Manager, AWS
yeilat@amazon.com

© 2018, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark

Amazon Relational Database Service (Amazon RDS)

Managed relational database service with a choice of popular database engines

Amazon Aurora | MySQL | MariaDB | PostgreSQL | Microsoft SQL Server | ORACLE

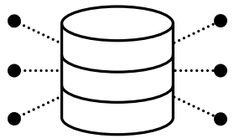

Easy to administer

Easily deploy and maintain hardware, OS, and DB software; built-in monitoring

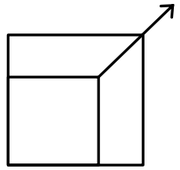

Performant & scalable

Scale compute and storage with a few clicks; minimal downtime for your application

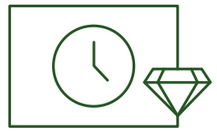

Available & durable

Automatic Multi-AZ data replication; automated backup, snapshots, and failover

Secure & compliant

Data encryption at rest and in transit; industry compliance and assurance programs

Amazon RDS database engines

Commercial	Open source	Cloud native
ORACLE Microsoft SQL Server	MySQL MariaDB PostgreSQL	Amazon Aurora MySQL compatible PostgreSQL compatible
Amazon Elastic Block Store (Amazon EBS)-based storage		Amazon Aurora storage system

Amazon RDS database engines

Commercial	Open source	Cloud native
ORACLE Microsoft SQL Server	MySQL MariaDB PostgreSQL	Amazon Aurora MySQL compatible PostgreSQL compatible
Amazon Elastic Block Store (Amazon EBS)-based storage		Amazon Aurora storage system

Agenda

- Why run MySQL
- Running managed vs. unmanaged MySQL databases
- Amazon RDS and Aurora MySQL
- Aurora Serverless + demo

Why run MySQL?

Why use a managed service?

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

MySQL is the world's most popular database

All respondents

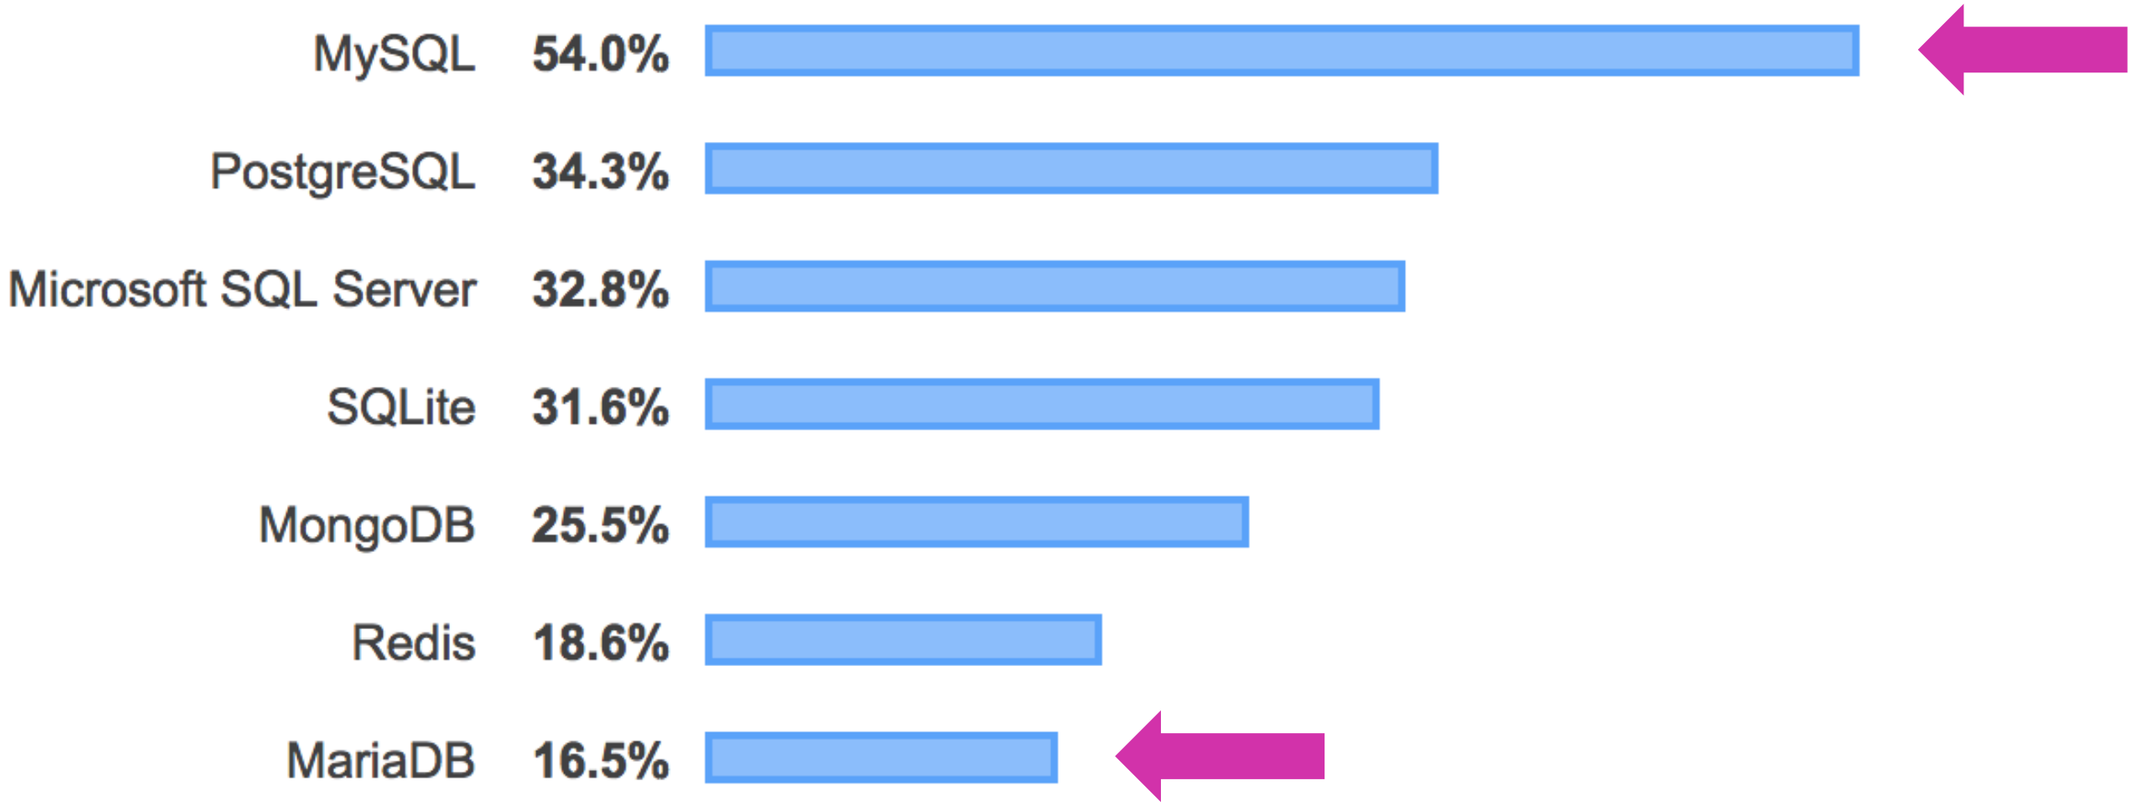

Source: Stack Overflow Developer Survey Results 2019 (https://insights.stackoverflow.com/survey/2019)

MySQL is the world's most popular database

Professional Developers

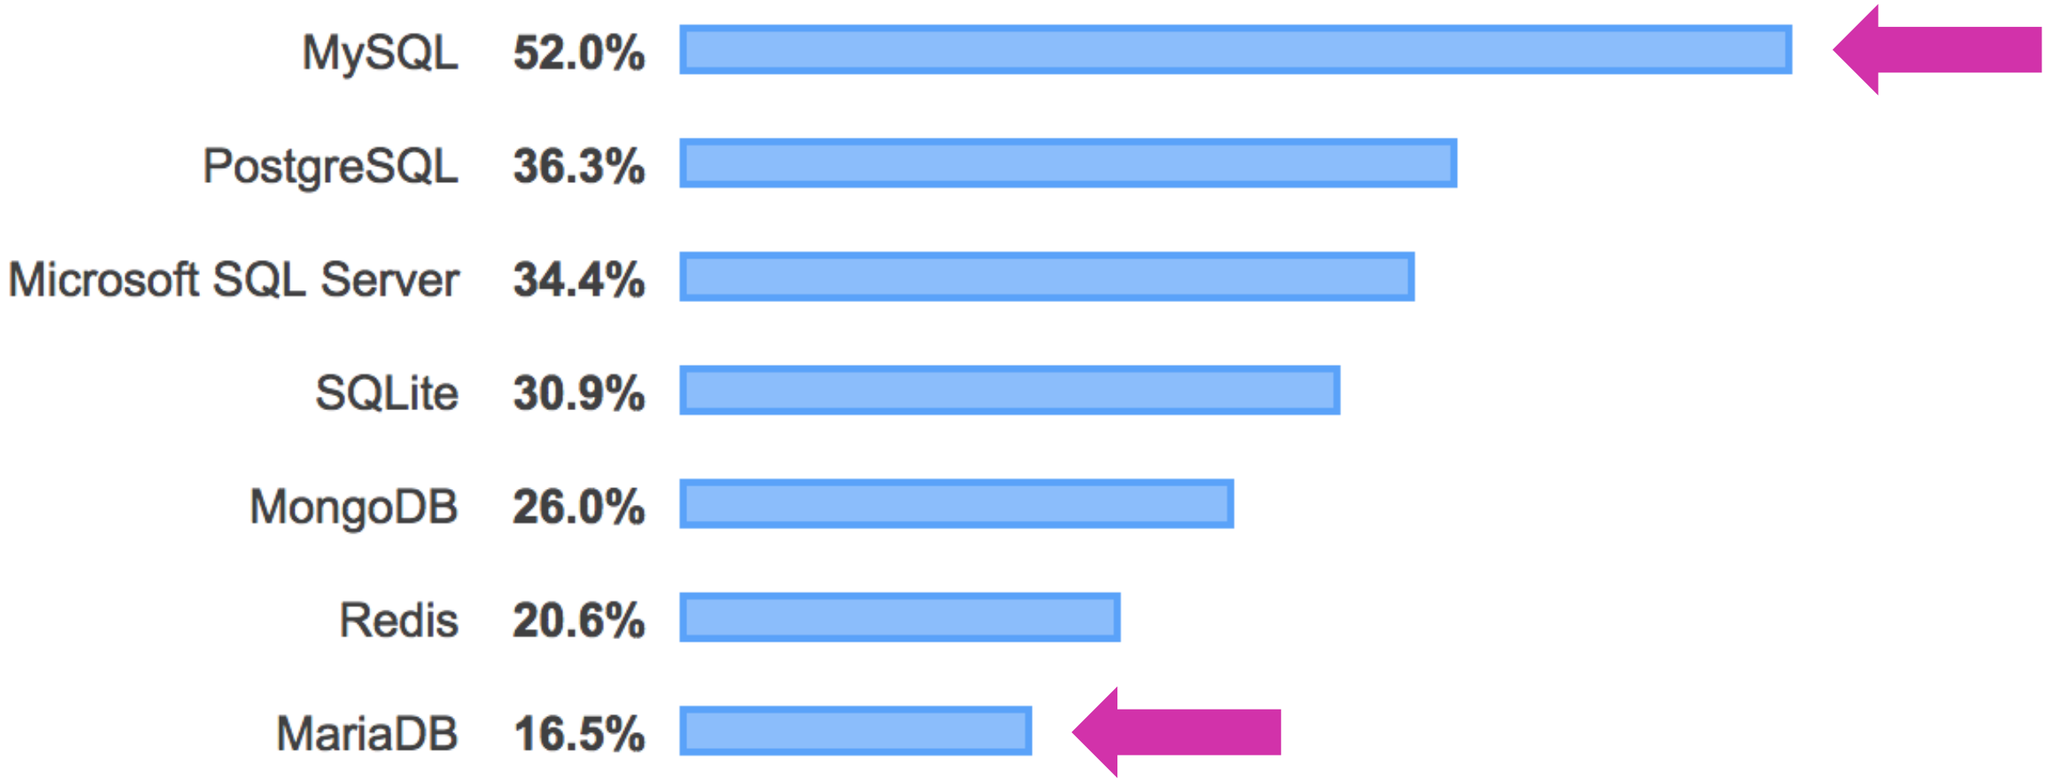

Source: Stack Overflow Developer Survey Results 2019 (https://insights.stackoverflow.com/survey/2019)

“Most popular” buys you . . .

Highly exercised, stable code

Large ecosystem of ISVs, Tools, Implementation and Support Partners

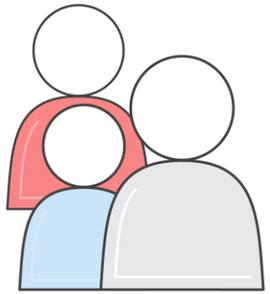

Large community of users and community-driven resources and a larger DBA talent pool

Managed databases: automating administrative tasks

In a managed database, the service provider takes care of your time-consuming database management tasks

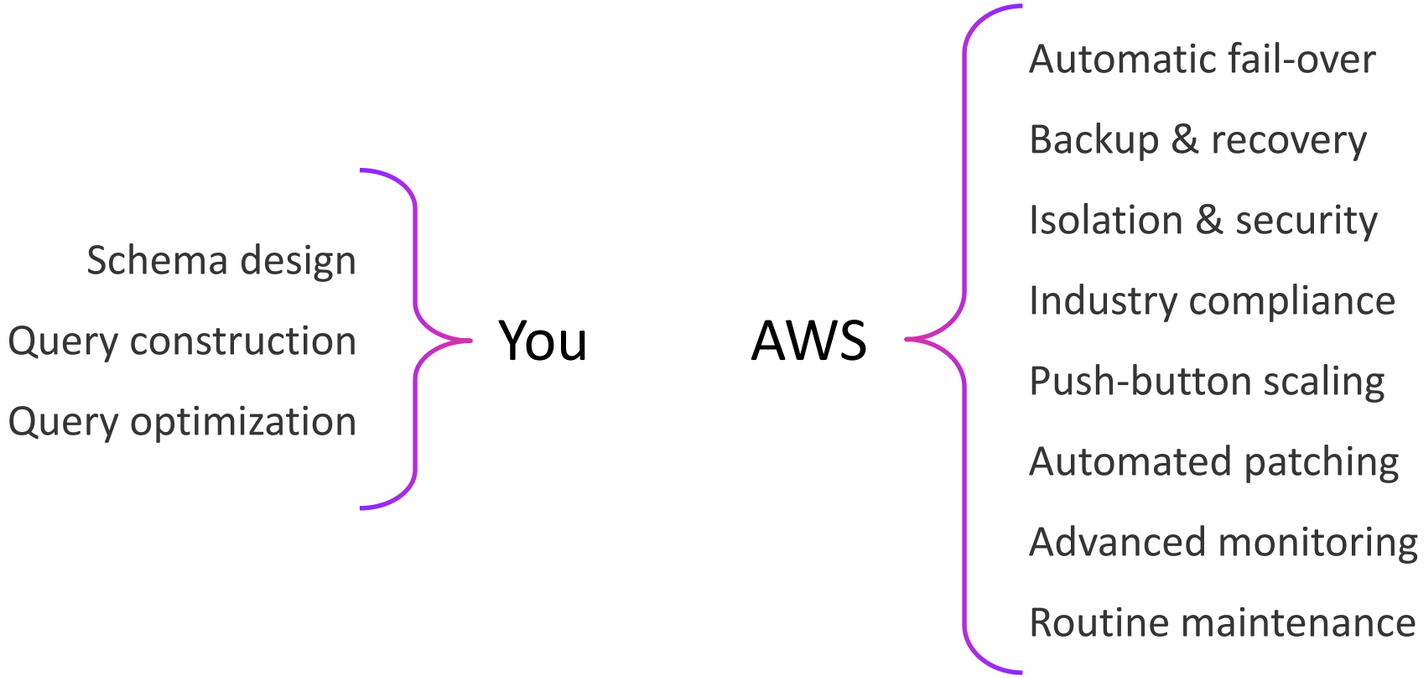

Amazon RDS

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

1. **Popular**
2. Innovative
3. Open

Amazon RDS

Hundreds of thousands of customers

Popular buys you . . .

Highly exercised, stable code

Unrivalled operational excellence

1. Popular
2. **Innovative**
3. Open

Amazon RDS

Amazon RDS highlights - AVAILABILITY

- Automated, 0-RPO failover across AZs
- Managed x-region replicas for DR
- Automated backups, manual snapshots
- Point-in-time recovery

Automated, 0-RPO failover across AZs

Each host manages set of EBS volumes with a full copy of the data

Instances are monitored by an external observer to maintain consensus over quorum

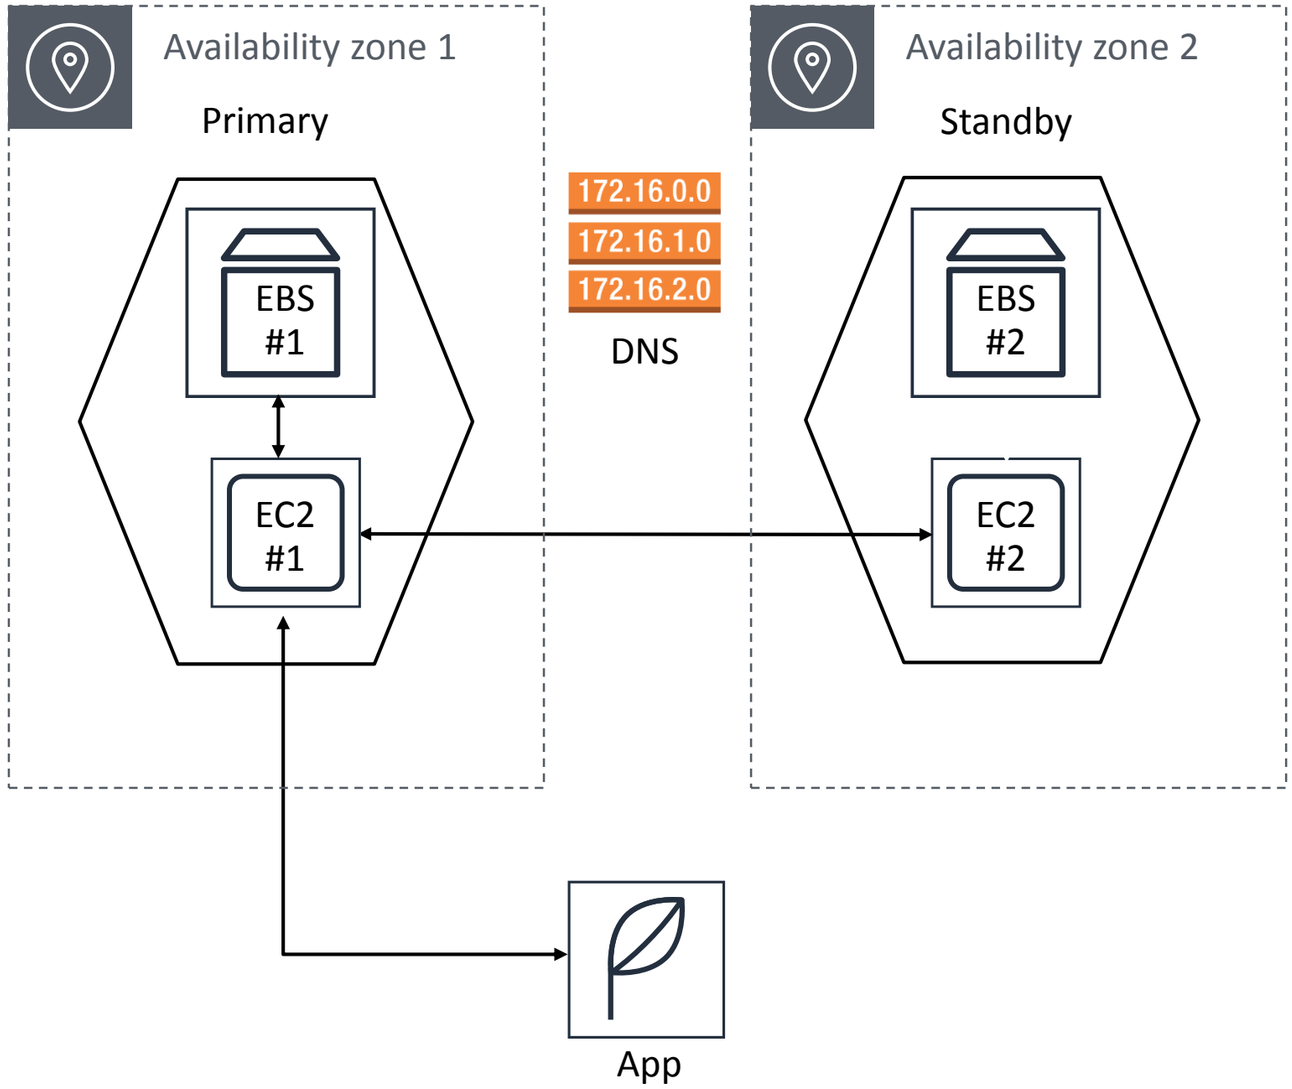

Automated, 0-RPO failover across AZs

Each host manages set of EBS volumes with a full copy of the data

Instances are monitored by an external observer to maintain consensus over quorum

Failover initiated by automation or through RDS API

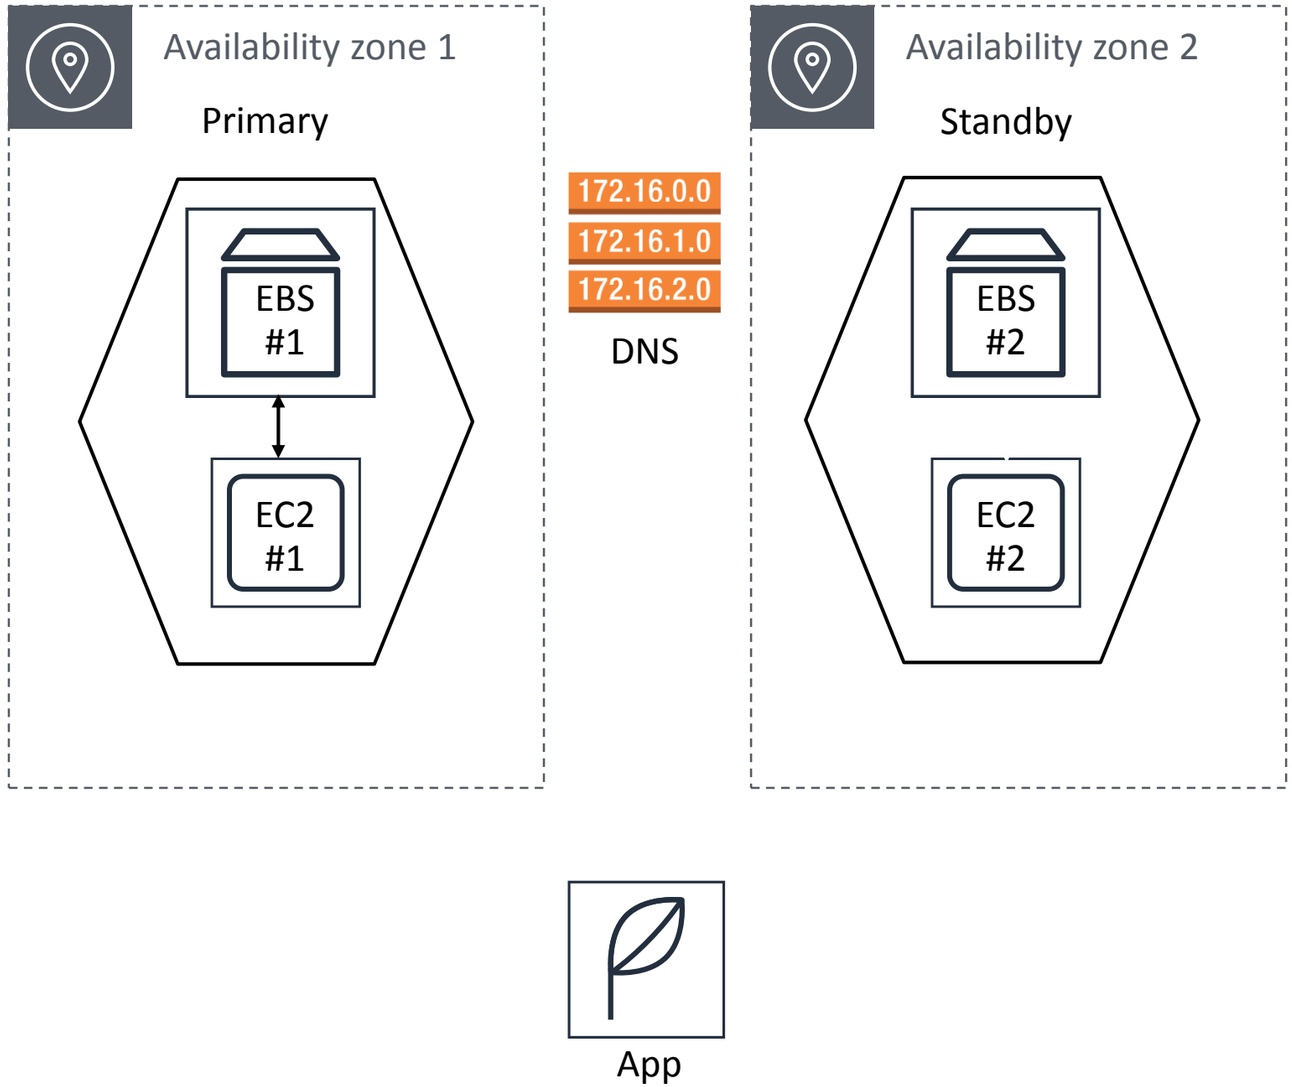

Automated, 0-RPO failover across AZs

Each host manages set of EBS volumes with a full copy of the data

Instances are monitored by an external observer to maintain consensus over quorum

Failover initiated by automation or through RDS API

Redirection to the new primary instance is provided through DNS

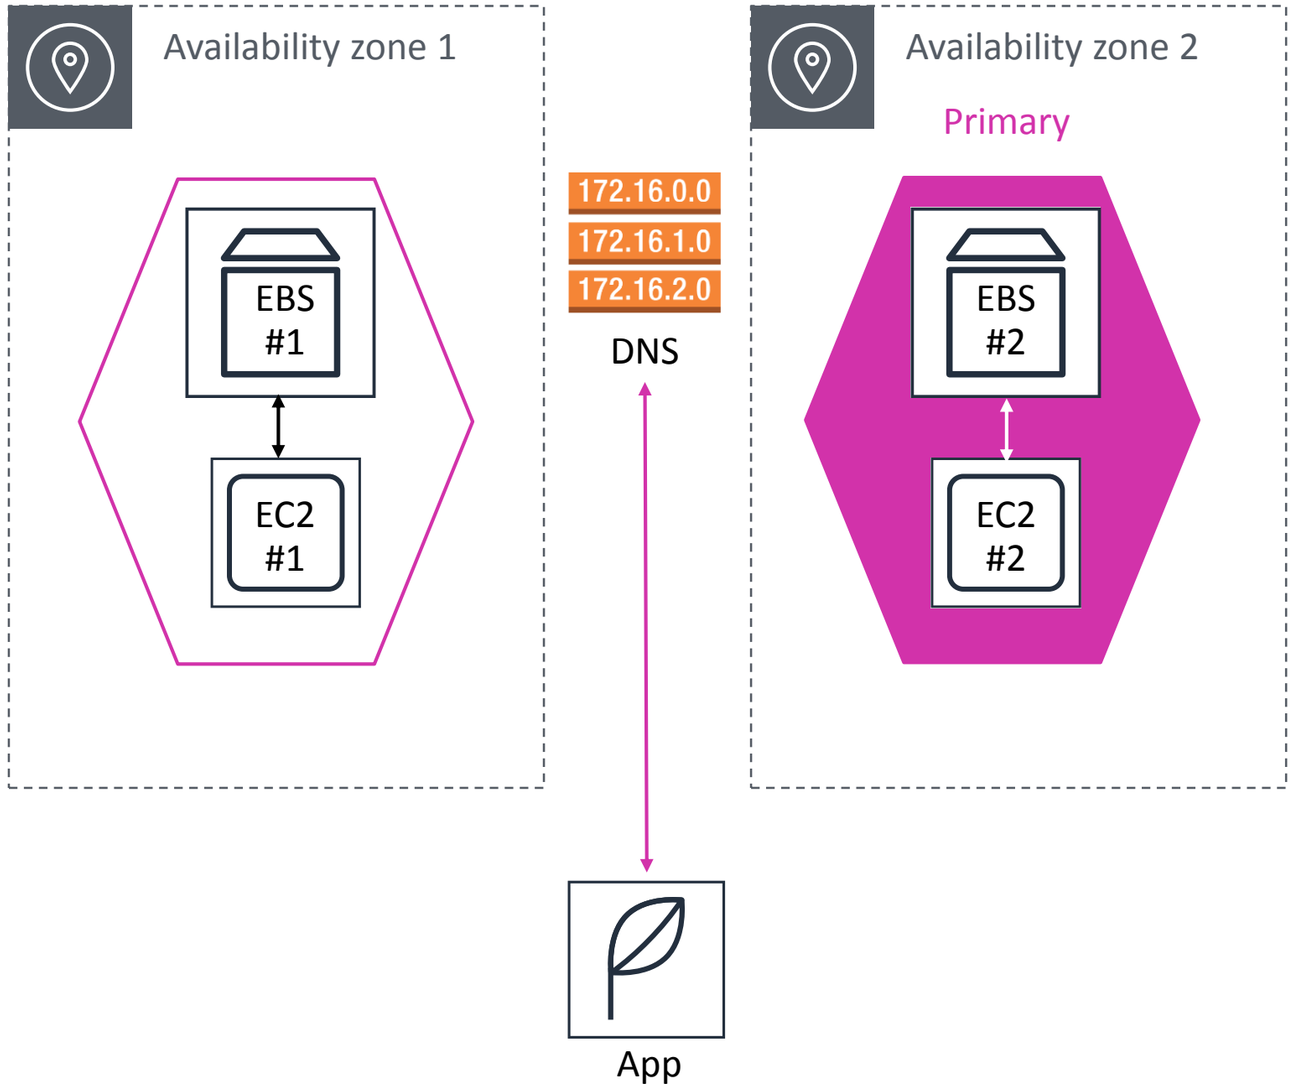

Automated, 0-RPO failover across AZs

Each host manages set of EBS volumes with a full copy of the data

Instances are monitored by an external observer to maintain consensus over quorum

Failover initiated by automation or through RDS API

Redirection to the new primary instance is provided through DNS

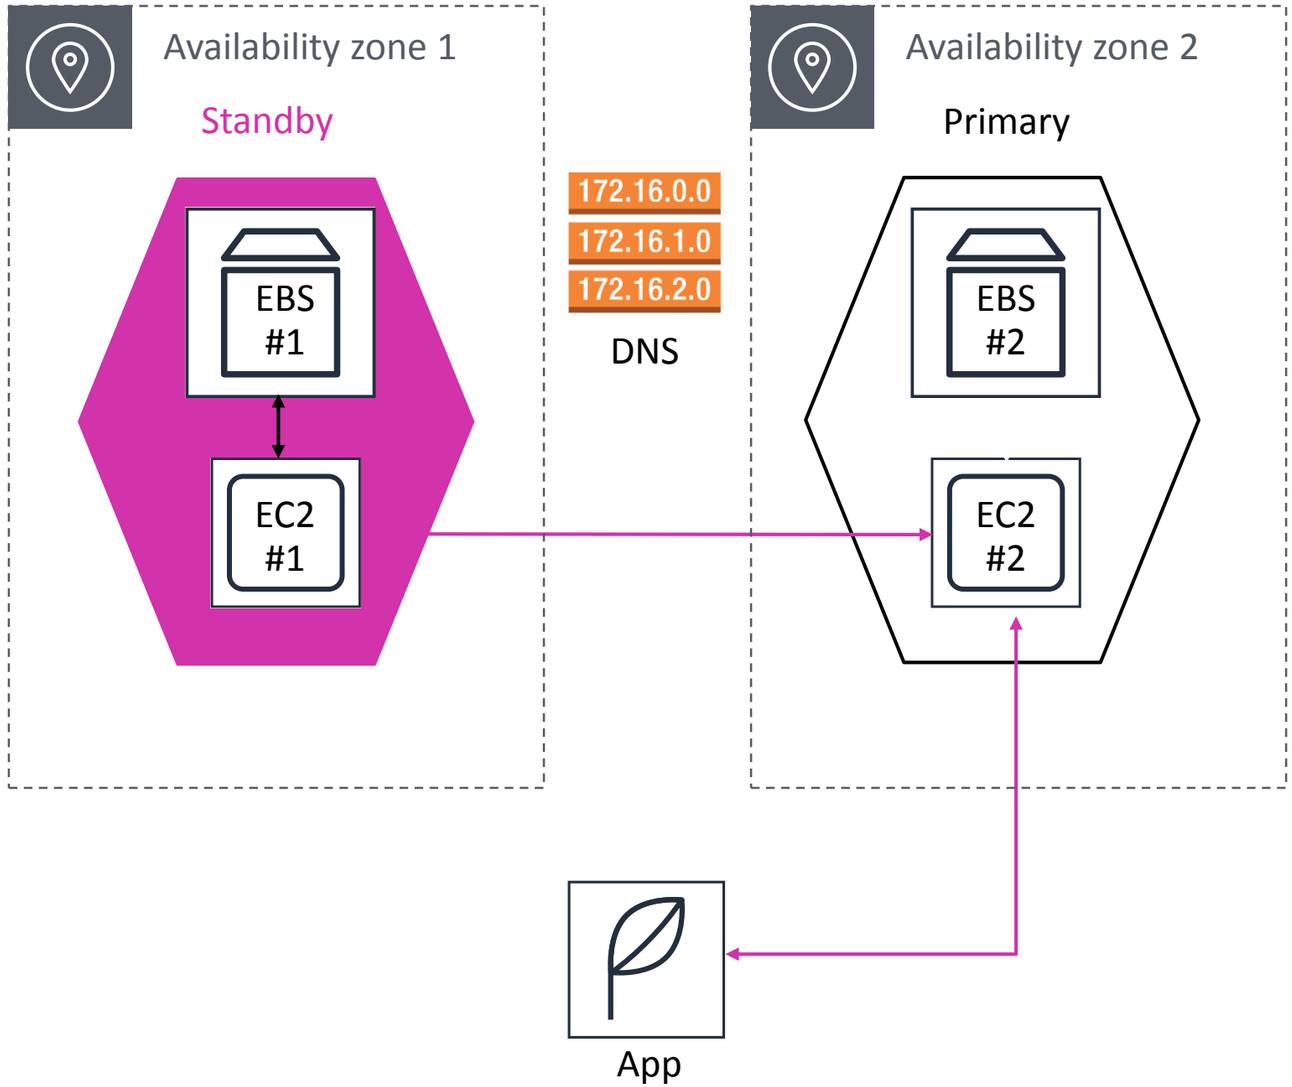

Read Scaling with Read Replicas

Use Amazon RDS read replicas to relieve pressure on your source database with additional read capacity

Create up to five replicas per source database

Monitor replication lag in Amazon CloudWatch or Amazon RDS console

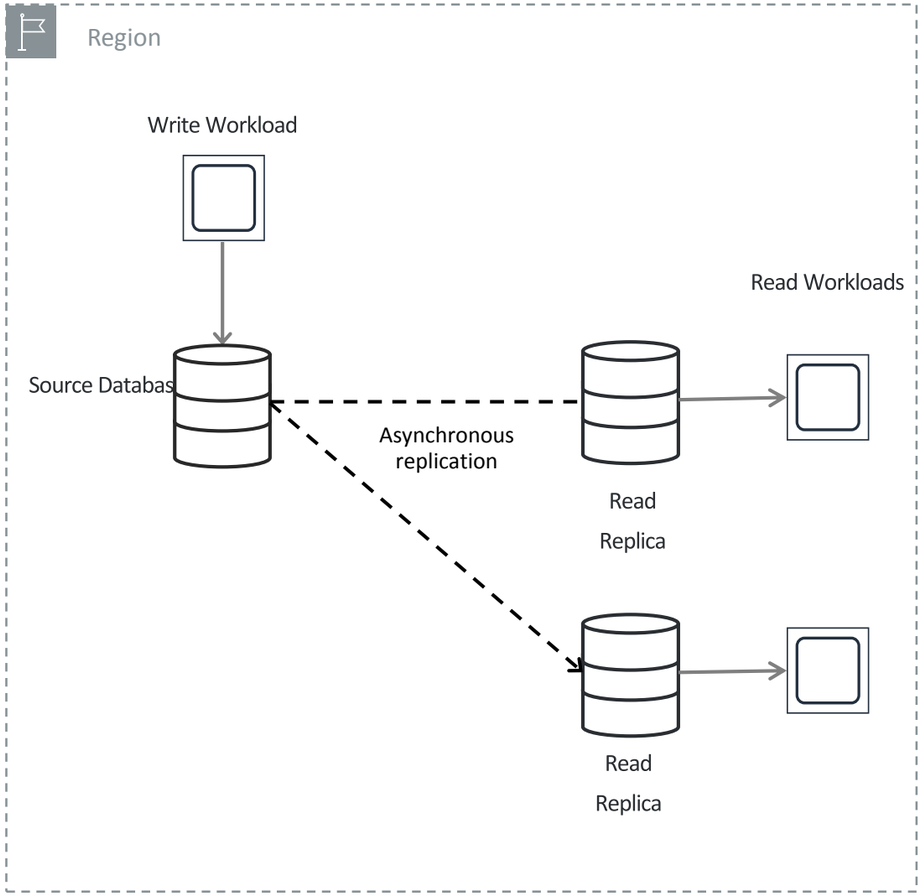

Planning for Disaster Recovery

Use a cross-region read replica as a standby database for recovery in the event of a disaster

Read replicas can be configured for Multi-AZ to reduce recovery time

Can use delayed replication for MySQL to protect from self-inflicted disasters

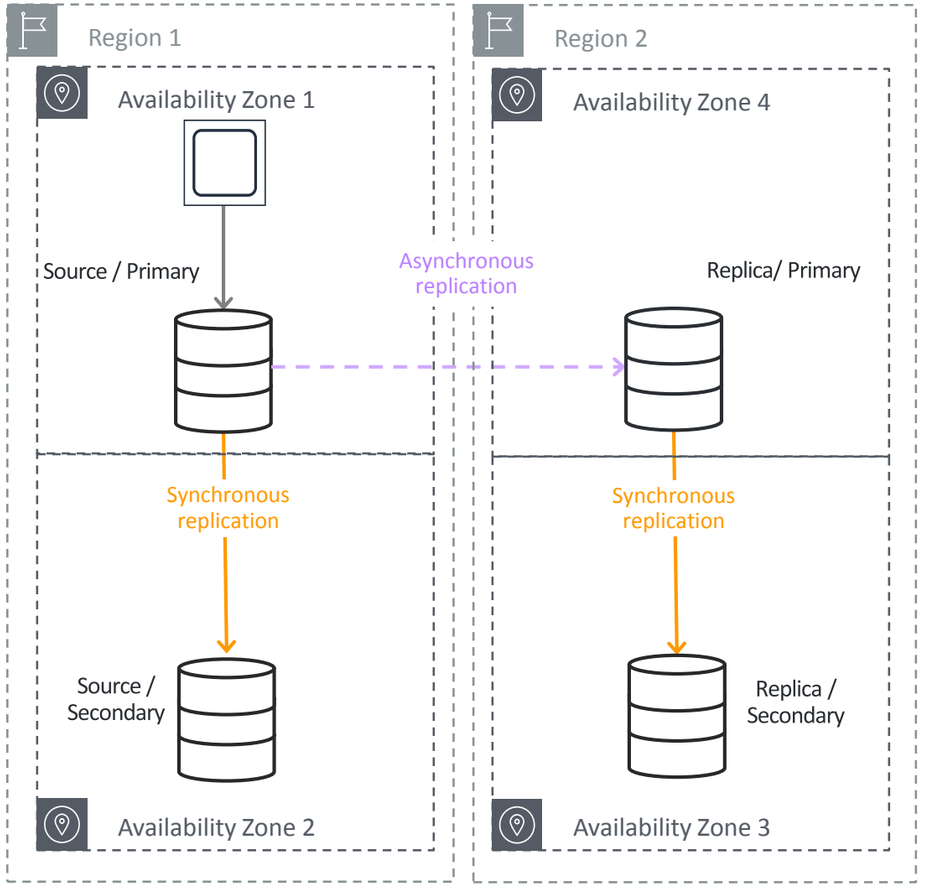

Backups, Snapshots, and Point-in-time restore

Two options – automated backups and manual snapshots

EBS snapshots stored in Amazon S3

Transaction logs stored every 5 minutes in S3 to support Point in Time Recovery

No performance penalty for backups

Snapshots can be copied across regions or shared with other accounts

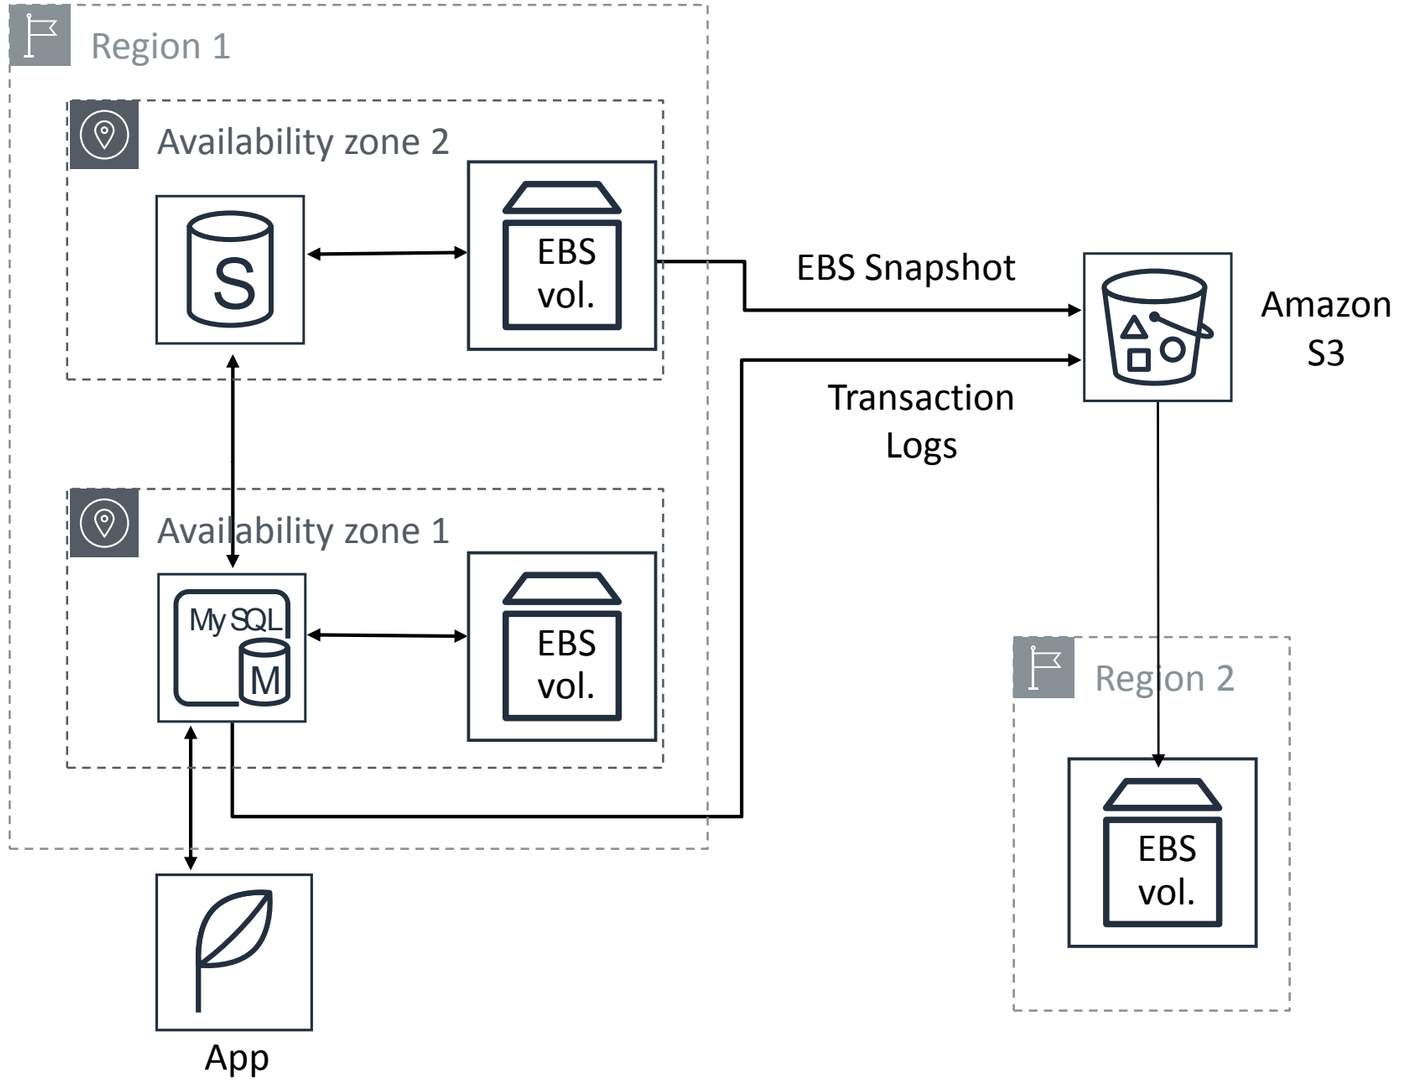

Amazon RDS highlights – SECURITY, MANAGEABILITY

- AWS IAM Database Authentication
- Automated OS and database upgrades
- Push-button scaling
- Managed binlog replication
- Log upload to CloudWatch Logs
- Industry compliance
- Per-second billing **NEW!**

Recommendations

Example issues:

Engine version outdated, Pending maintenance available, Automated backups disabled, Enhanced Monitoring disabled, Encryption disabled

Parameter recommendations:

Non-default custom memory parameters, Change buffering enabled, Logging to table

NEW!

Aurora cluster recommendations

NEW!

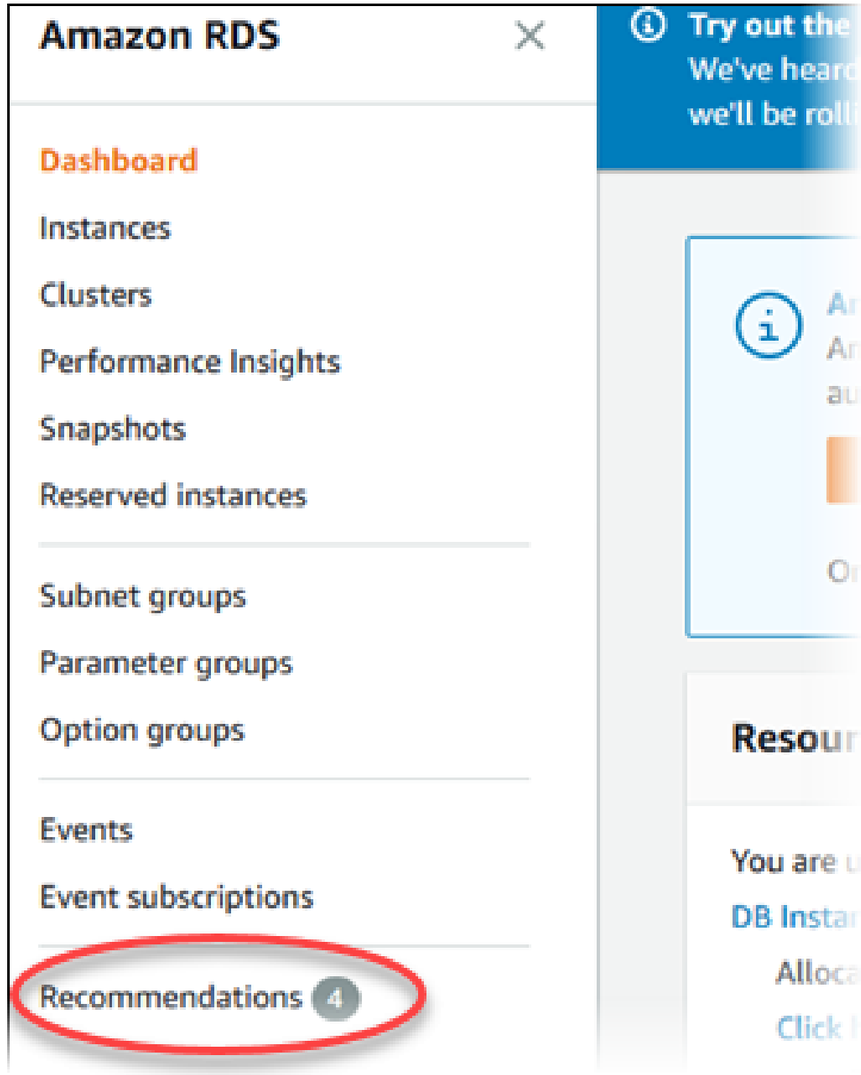

Start and Stop

Solution for development and test environments

Stop and start a running database instance from the console or AWS Command Line Interface (AWS CLI)

Now available for both single-AZ and Multi-AZ DB instances and Aurora DB clusters NEW!

While instance is stopped, you only pay for storage

Backup retention window is maintained while stopped

Instances are restarted after seven days

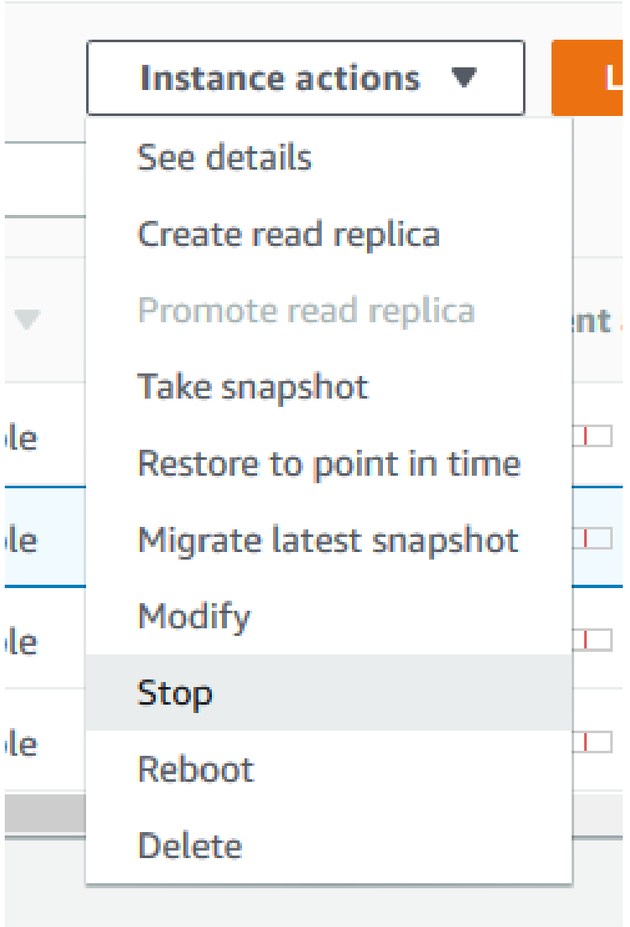

Amazon RDS highlights – PERFORMANCE

- R5, M5, and T3 database instance family **NEW!**
- Elastic volumes up to 64 TB **NEW!**
- Up to 80K Provisioned IOPS **NEW!**

Performance Insights

Dashboard showing database load

- Easy – e.g. drag and drop
- Powerful – drill down using zoom in

Identifies source of bottlenecks

- Sort by top SQL
- Slice by host, user, wait events

Adjustable time frame

- Hour, day, week , month
- Up to 2 years of data; 7 days free

Engine support

- MySQL, MariaDB, and all other RDS database engines

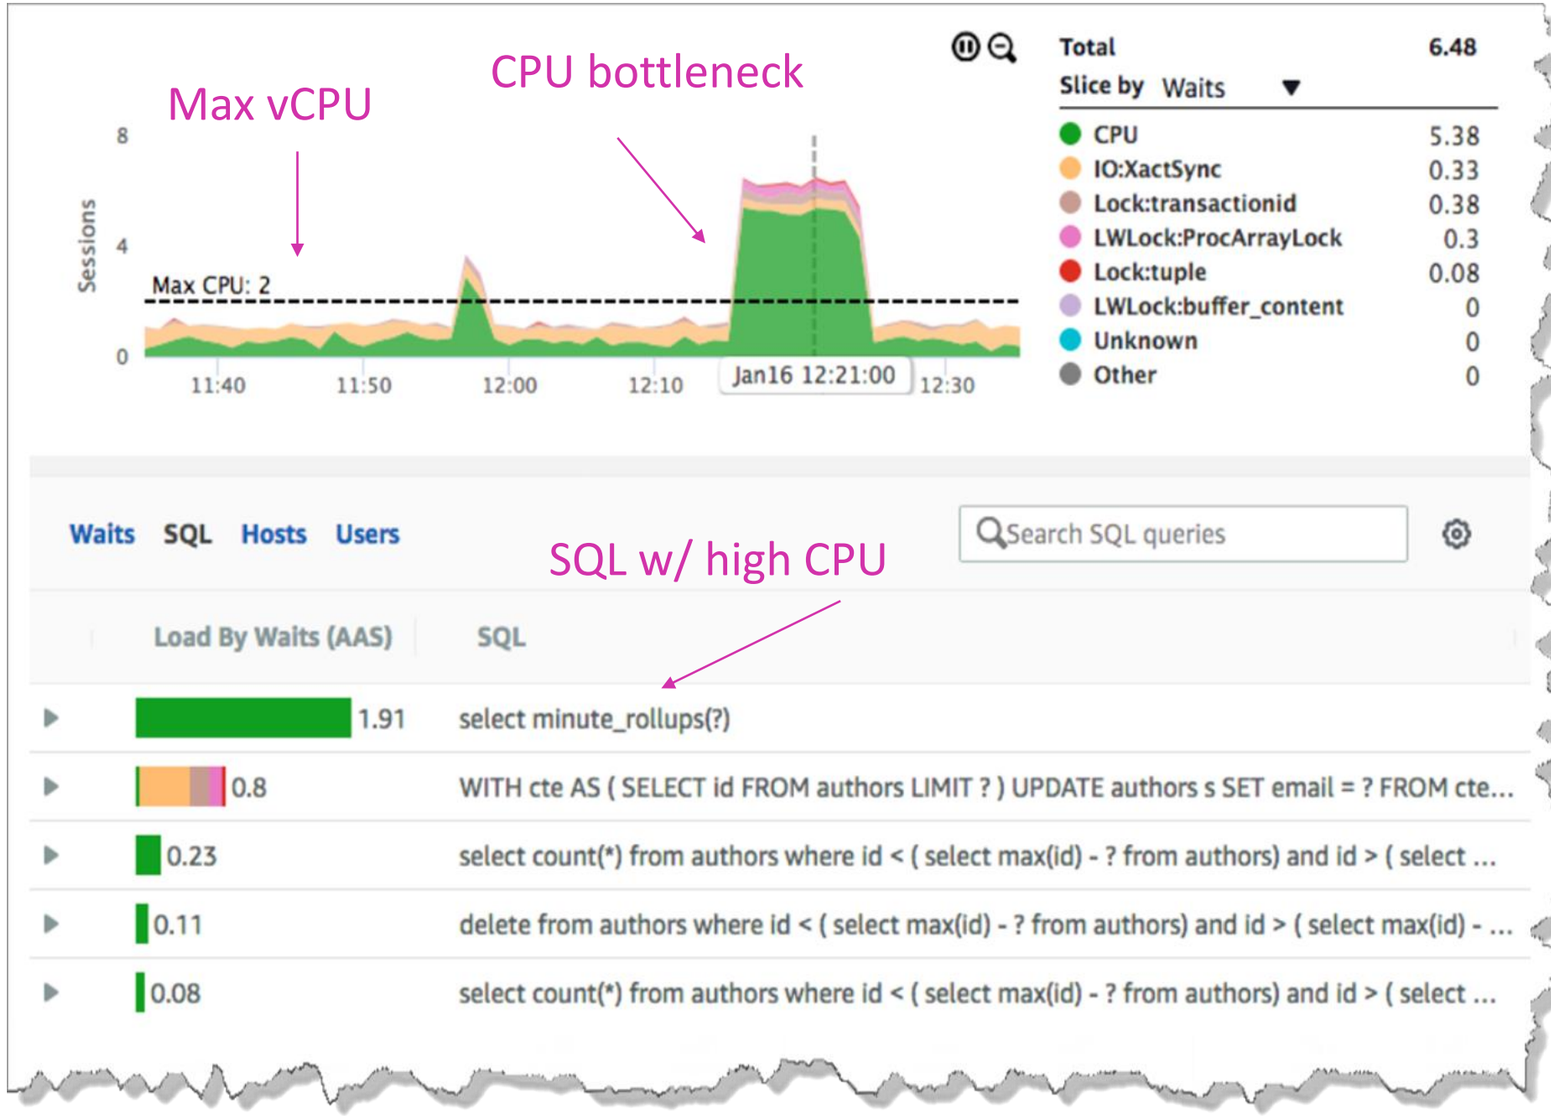

Monitor RDS with CloudWatch

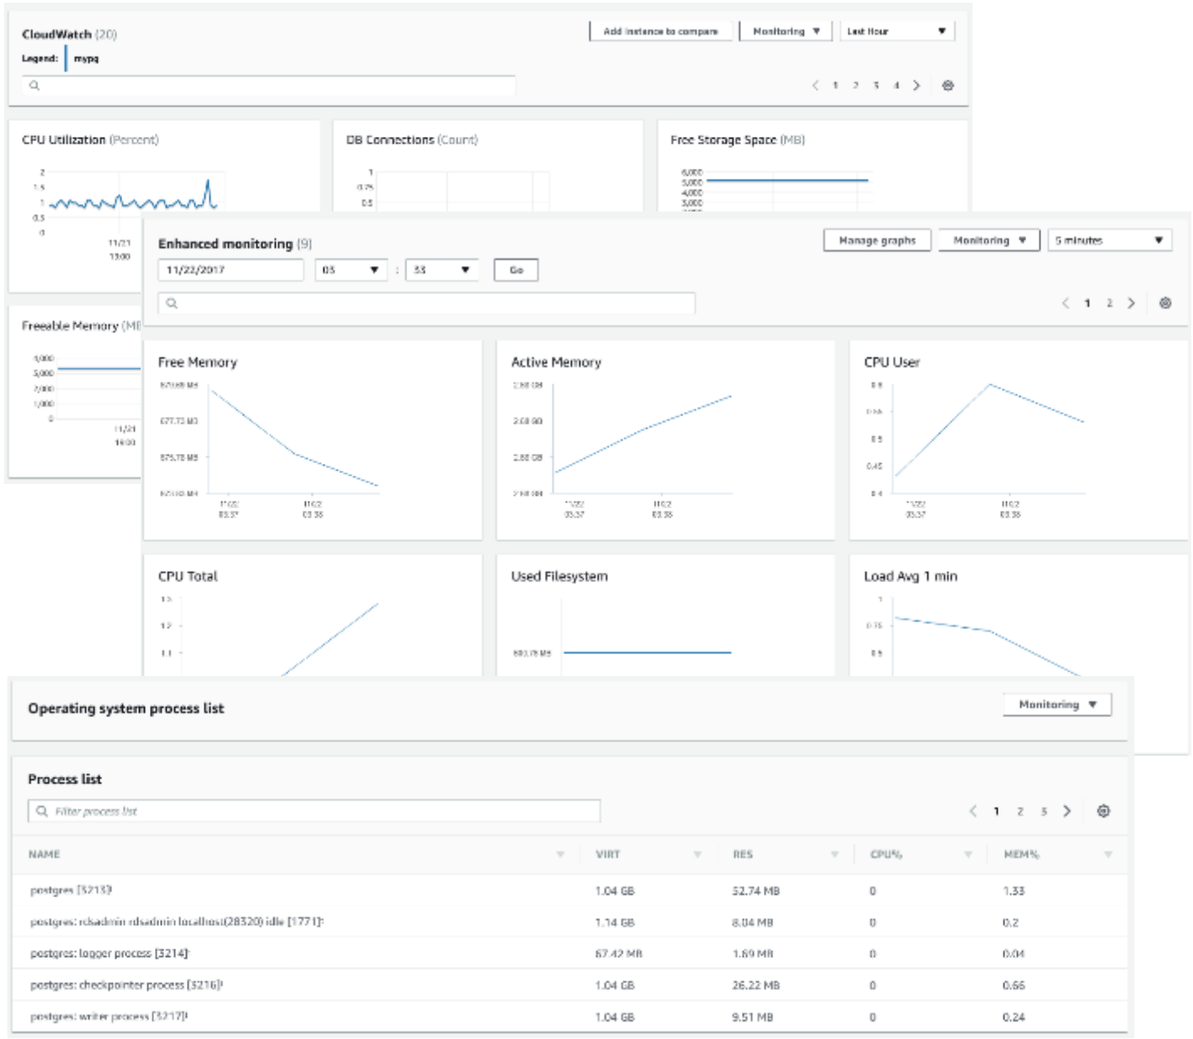

- Amazon CloudWatch metrics
  - CPU/Storage/Memory
  - Swap usage
  - I/O (read and write)
  - Latency (read and write)
  - Throughput (read and write)
  - Replica lag
- Amazon CloudWatch alarms
  - Similar to on-premises monitoring tools
- Enhanced monitoring
  - Access to additional CPU, memory, file system, and disk I/O metrics
  - As low as one-second intervals
- Integration with third-party monitoring tools

Database activity monitoring and insights

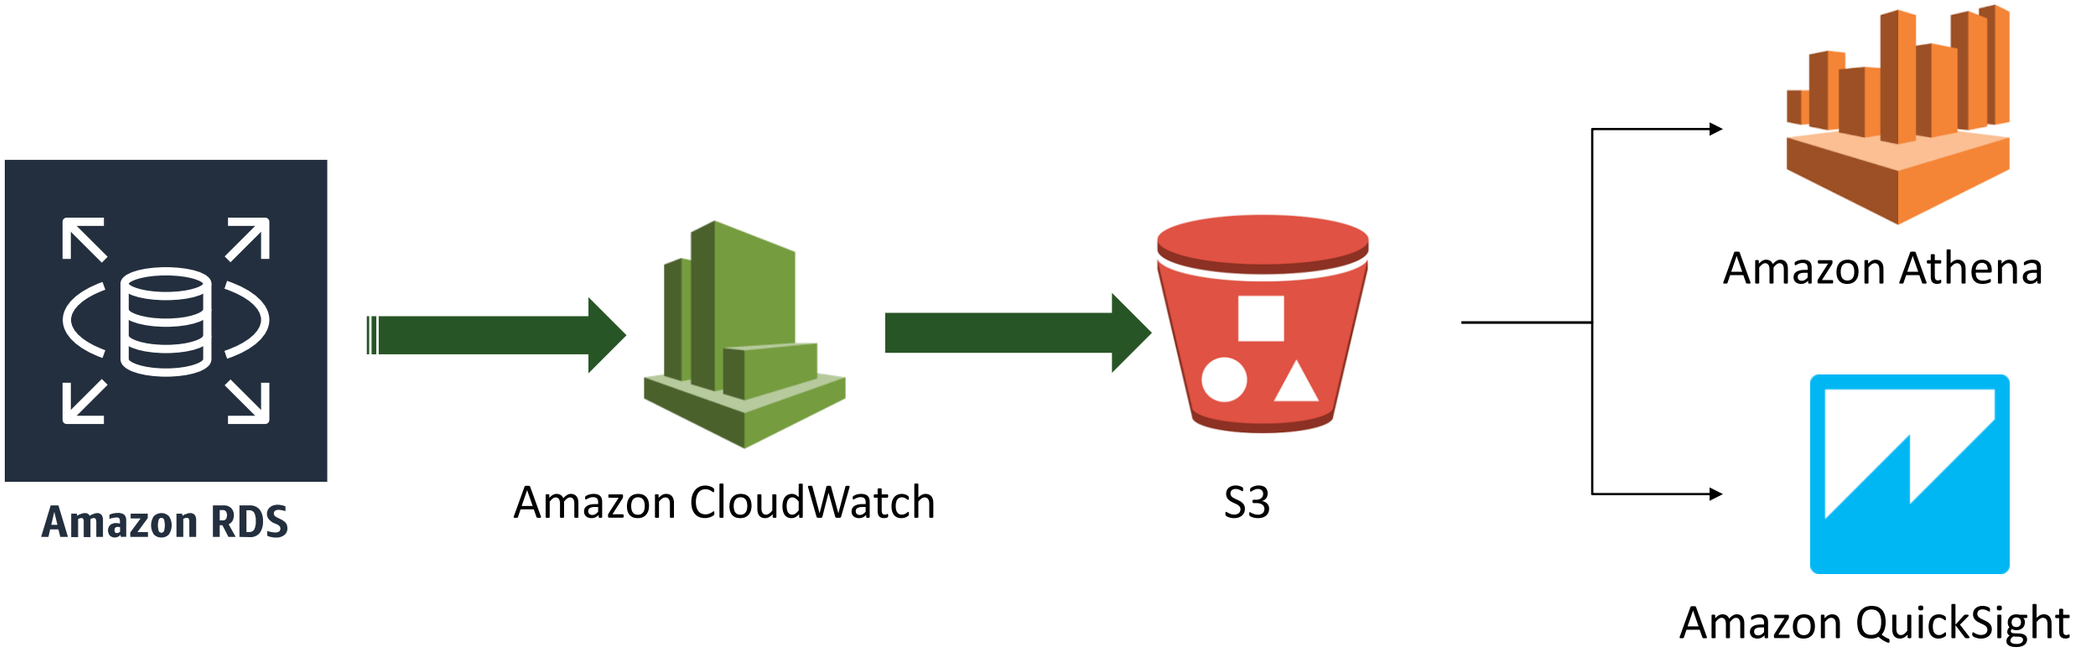

- Continuously monitor activity in your DB clusters by sending audit logs to CloudWatch Logs.
- Export to S3 for long term archival; analyze logs using Athena; visualize logs with QuickSight.

Search: Look for specific events across log files.

Metrics: Measure activity in your Aurora DB cluster.

Visualizations: Create activity dashboards

Alarms: Get notified or take actions

1. Popular

2. Innovative

3. **Open**

Amazon RDS

Freedom to migrate in . . .

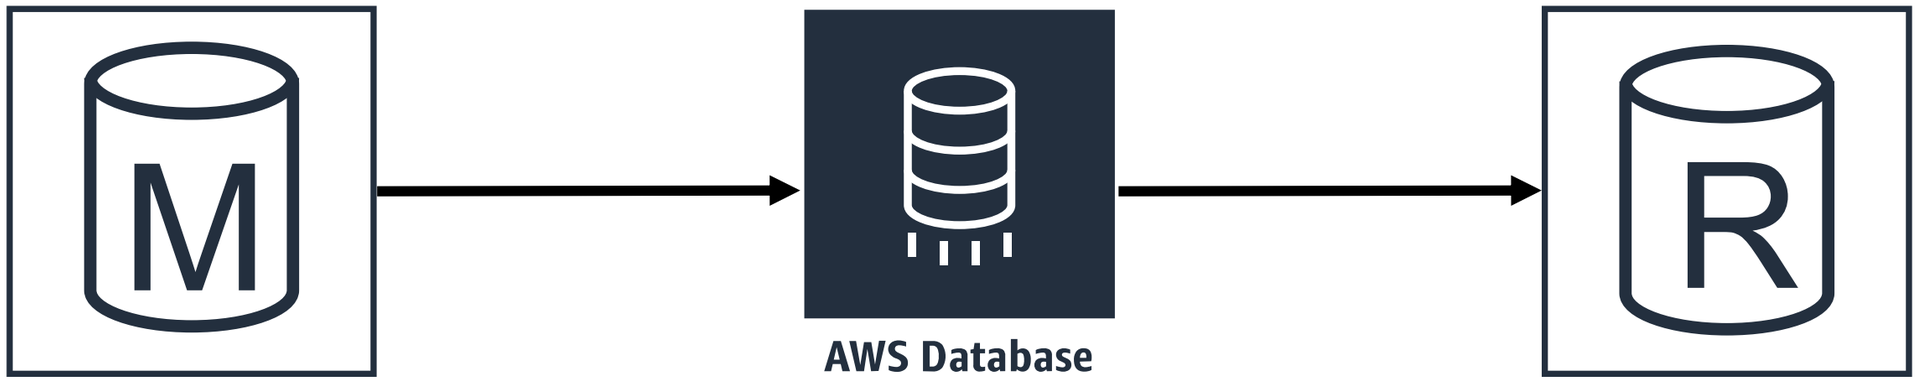

On premises

Amazon Elastic Cloud Compute (Amazon EC2)

Amazon Simple Storage Service (Amazon S3)

Amazon Redshift

Amazon RDS

Freedom to migrate in *and* out

On premises

Amazon EC2

Amazon S3

Amazon Redshift

Amazon DynamoDB

Amazon RDS

Amazon Aurora

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

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

Scale-out, distributed architecture

Purpose-built log-structured distributed storage system designed for databases

Storage volume is striped across hundreds of storage nodes distributed over 3 different availability zones

Six copies of data, two copies in each availability zone to protect against AZ+1 failures

Plan to apply same principles to other layers of the stack

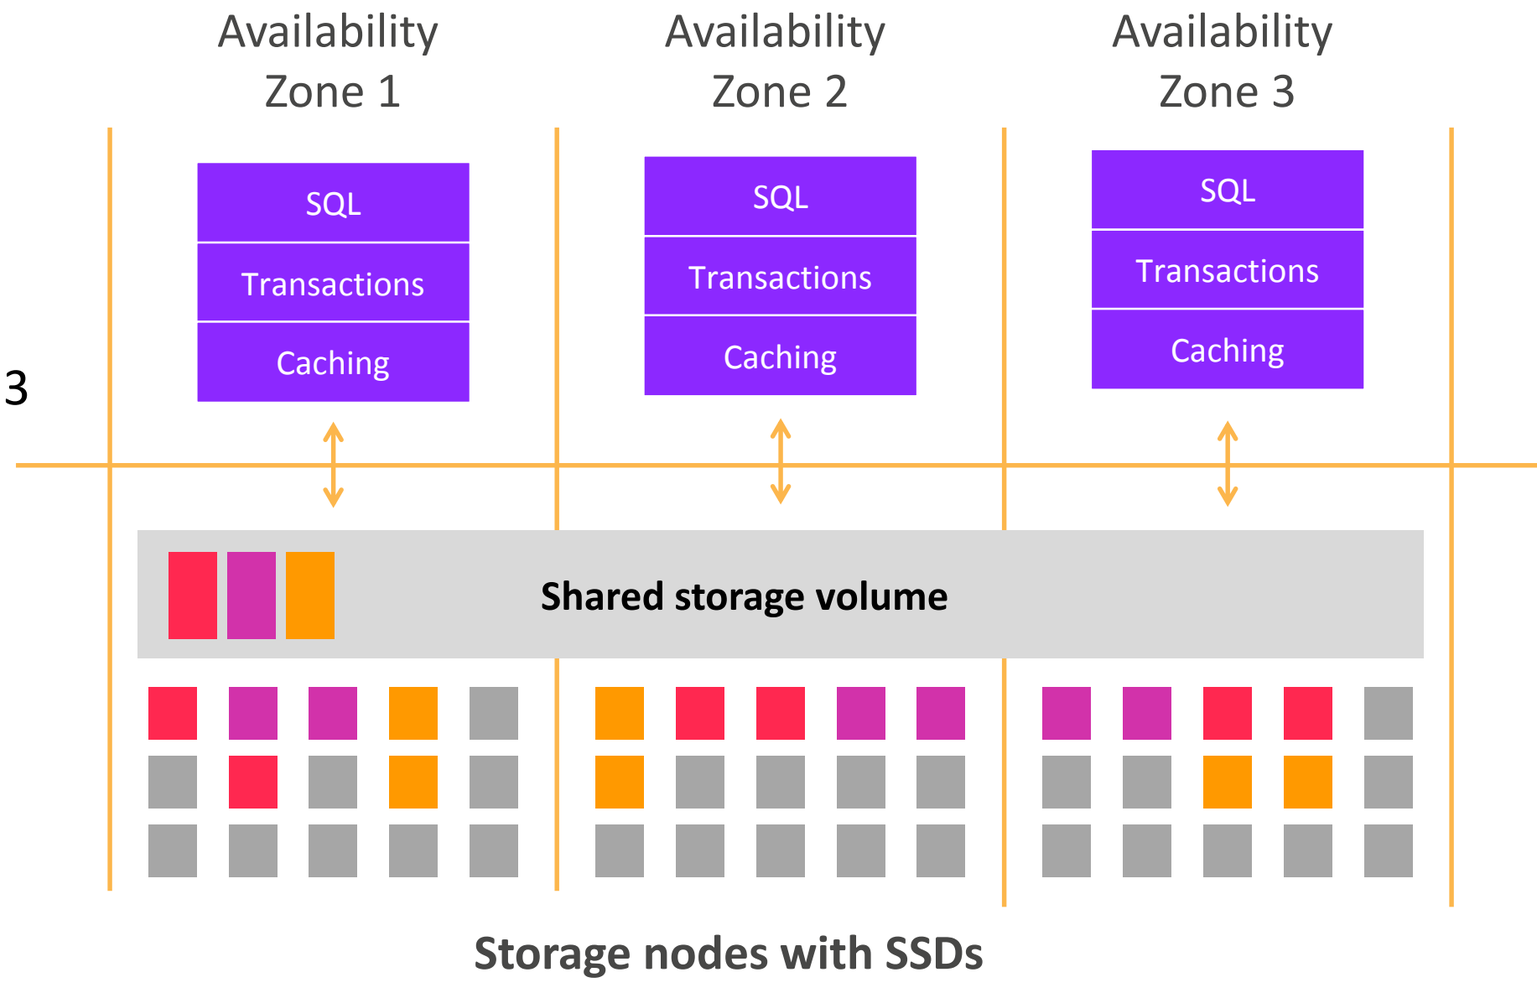

Leveraging AWS services

Lambda function

Invoke AWS Lambda events from stored procedures/triggers

AWS Identity and Access Management

Use AWS Identity and Access Management (IAM) roles to manage database access control

Amazon S3

Load data from Amazon Simple Storage Service (Amazon S3), store snapshots and backups in S3

Amazon CloudWatch

Upload systems metrics and audit logs to CloudWatch

Aurora customer adoption

Fastest growing service in AWS history

Aurora is used by ¾ of the top 100 AWS customers

Write and read throughput

Aurora MySQL is 5x faster than MySQL

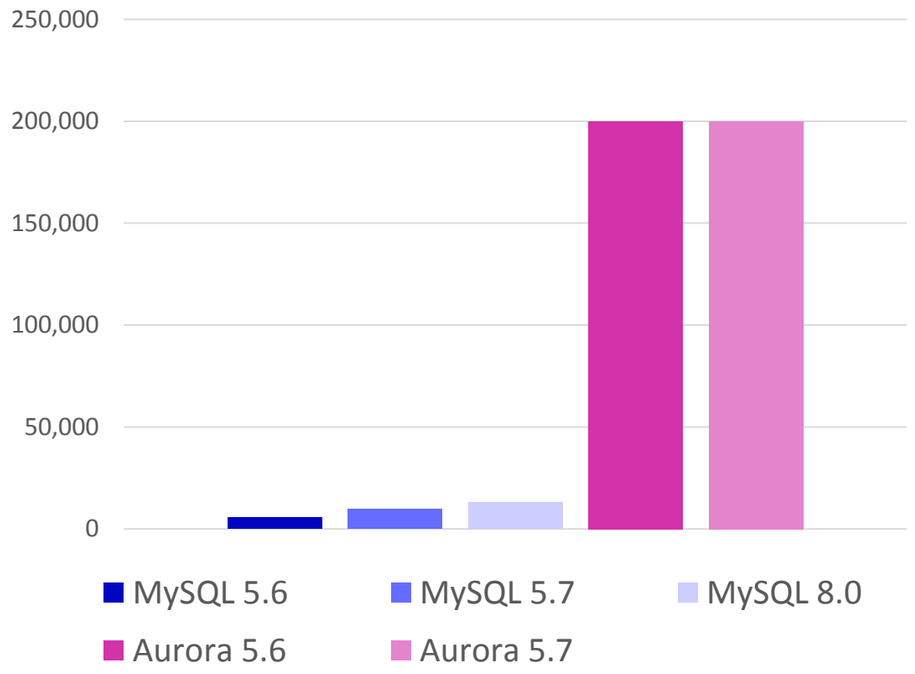

Write Throughput

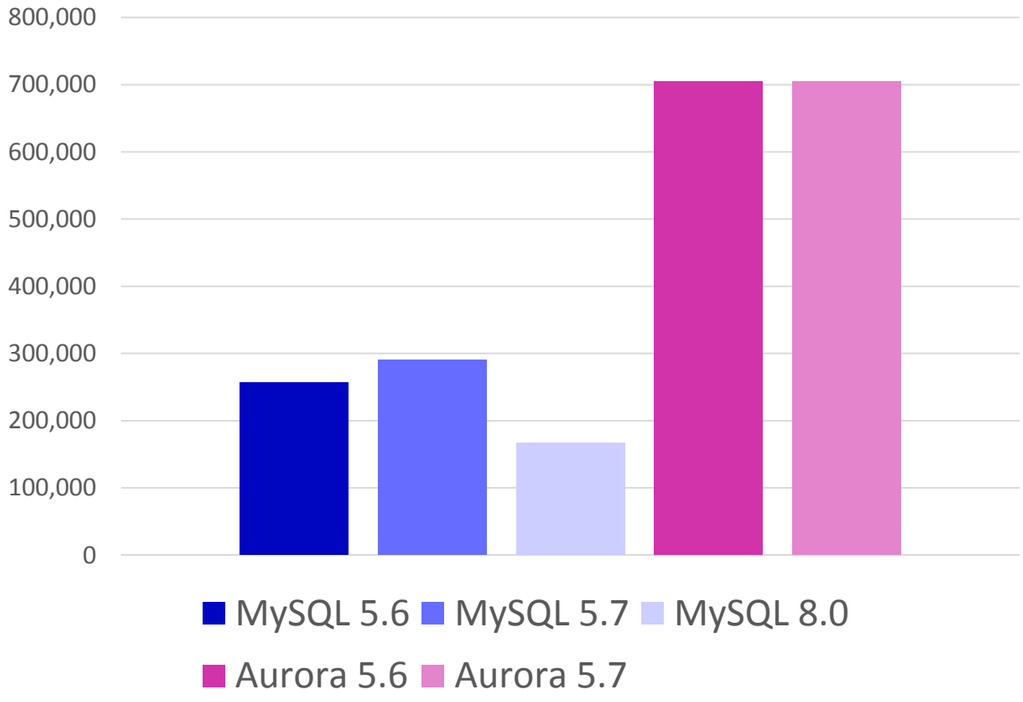

Read Throughput

Using Sysbench with 250 tables and 200,000 rows per table on R4.16XL

Aurora read scaling options

15 promotable read replicas per cluster

Auto-scaling to automatically add & remove replicas

Physical replication across regions (Aurora Global Database)

Logical (binlog) replication to any MySQL database

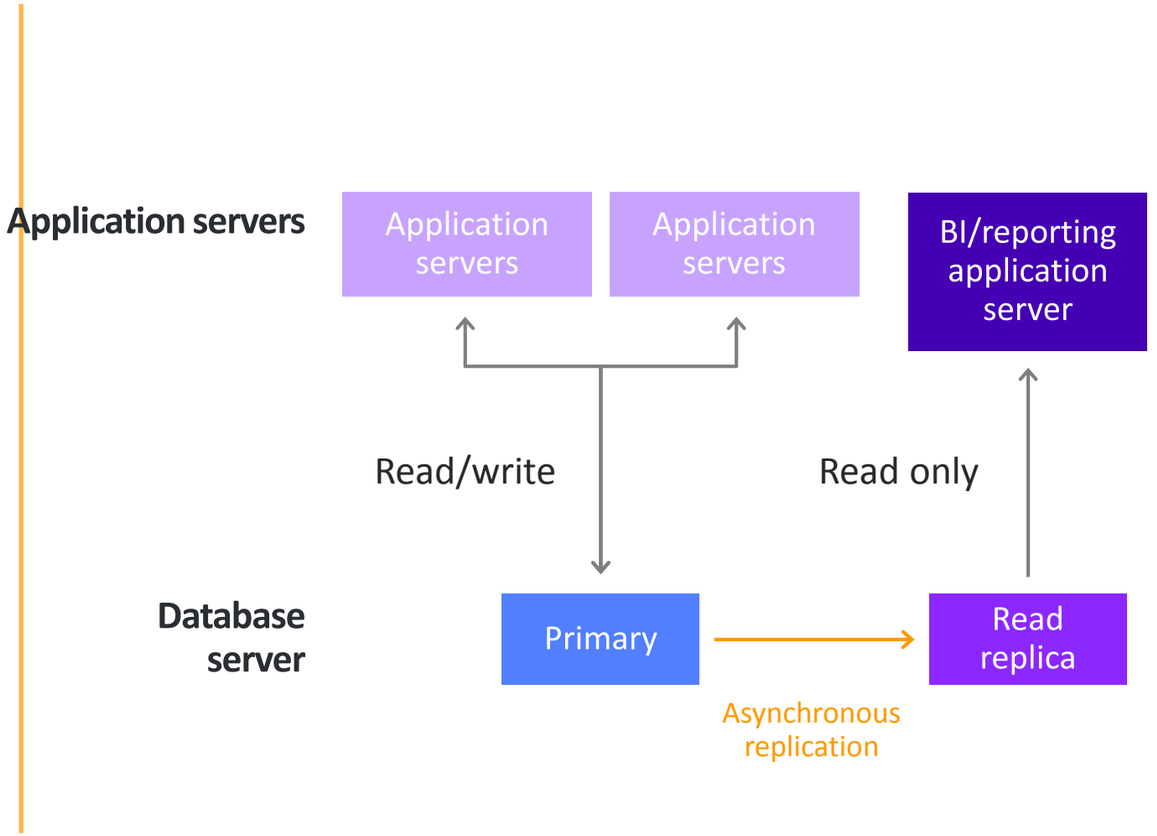

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Aurora read replicas

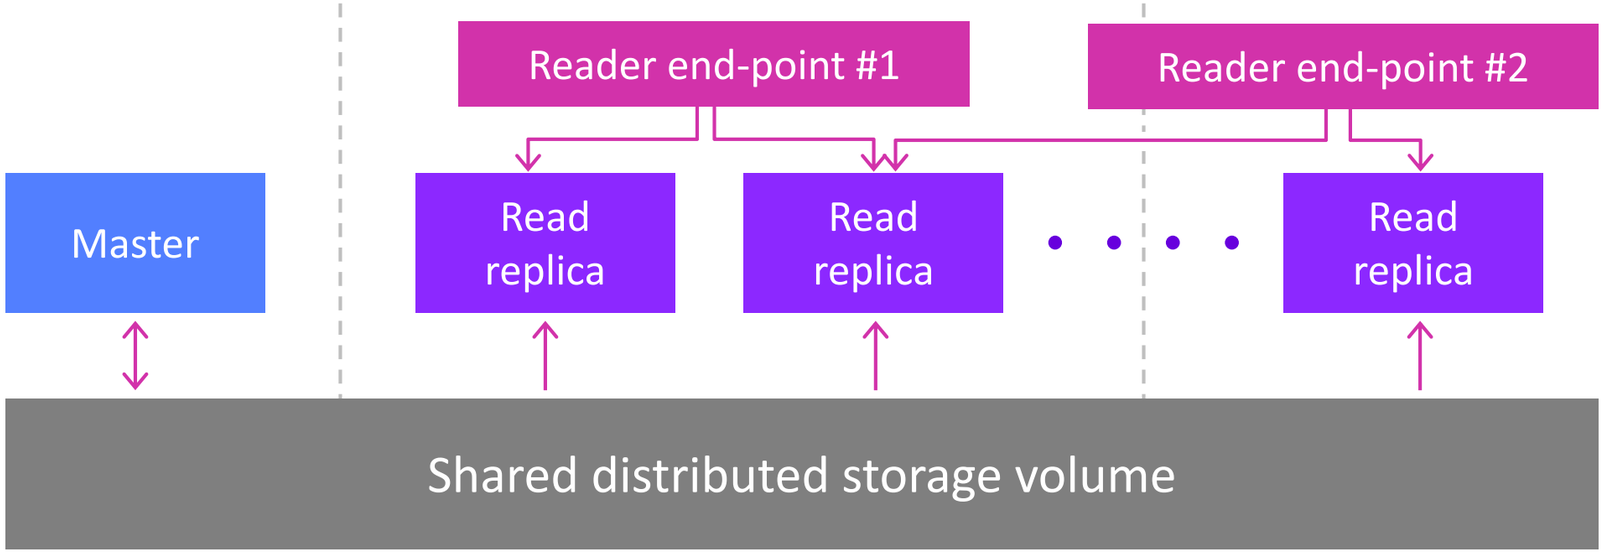

Up to 15 promotable read replicas across multiple availability zones

Redo log based (physical) replication leads to low replica lag—typically < 10ms

Custom reader end-point with configurable failover order

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

Aurora Global Database

Faster disaster recovery and enhanced data locality

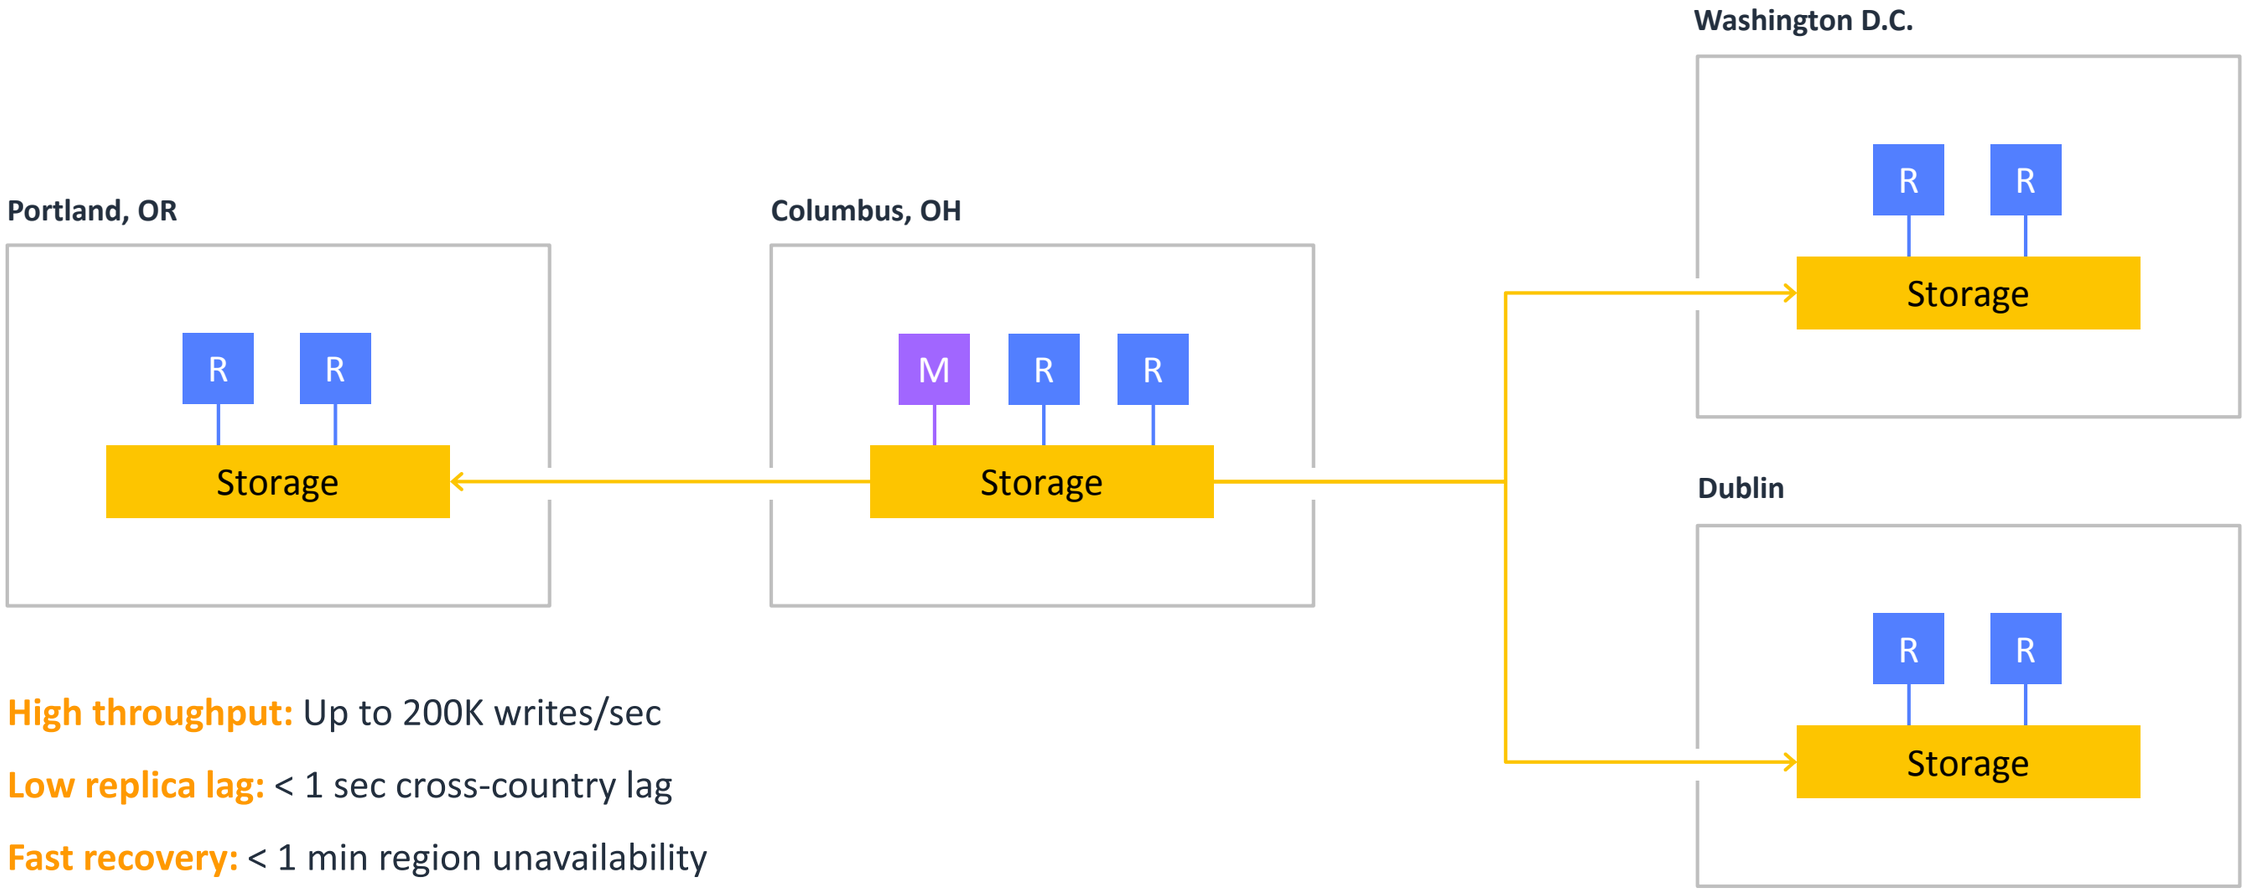

High throughput: Up to 200K writes/sec

Low replica lag: < 1 sec cross-country lag

Fast recovery: < 1 min region unavailability

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

Instant crash recovery

Traditional database

Have to replay logs since the last checkpoint

Typically 5 minutes between checkpoints

Single-threaded in MySQL; requires a large number of disk accesses

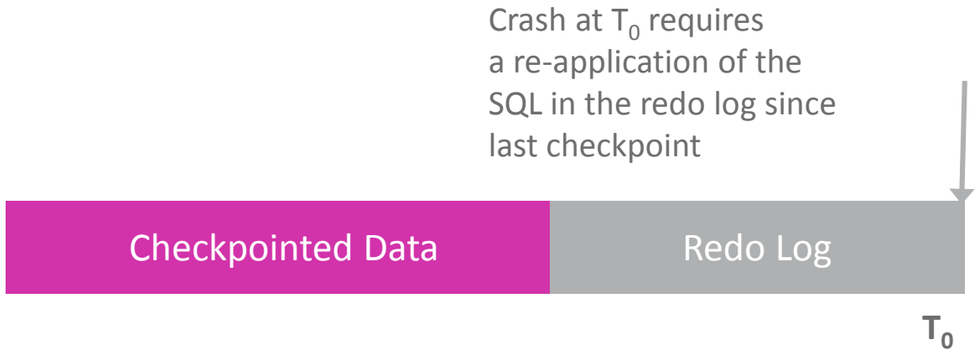

Amazon Aurora

Underlying storage replays redo records on demand as part of a disk read

Parallel, distributed, asynchronous

No replay for startup

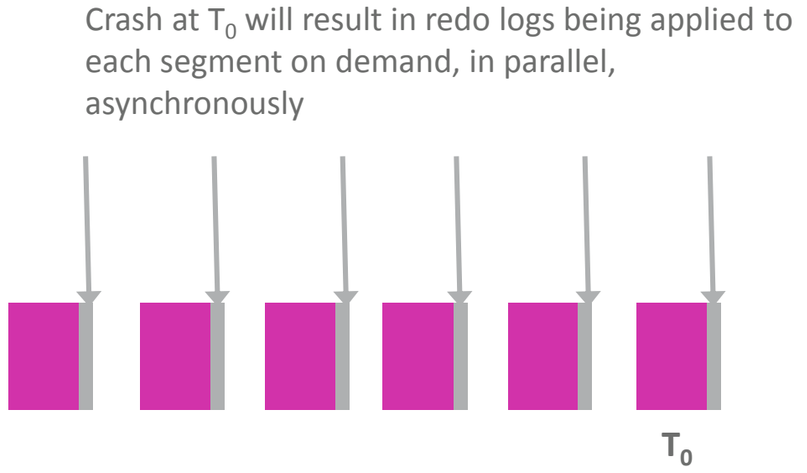

When Database fails – recovery is fast <30 seconds

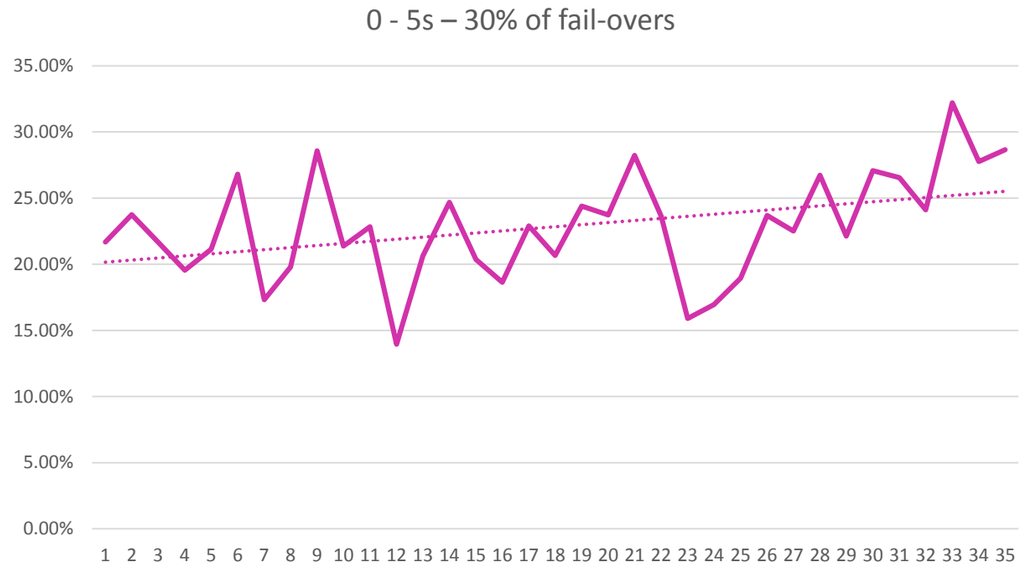

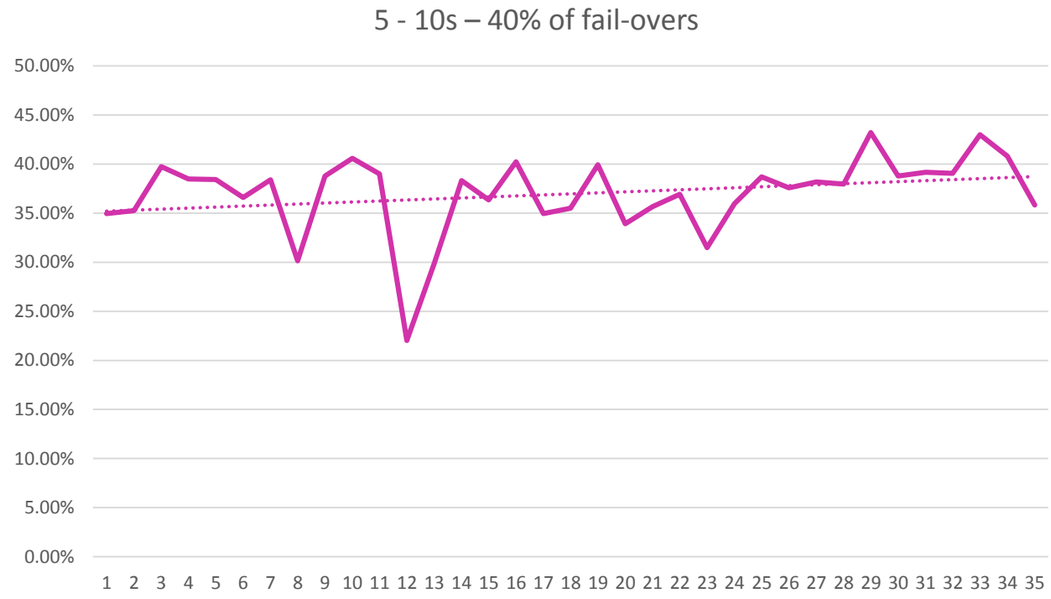

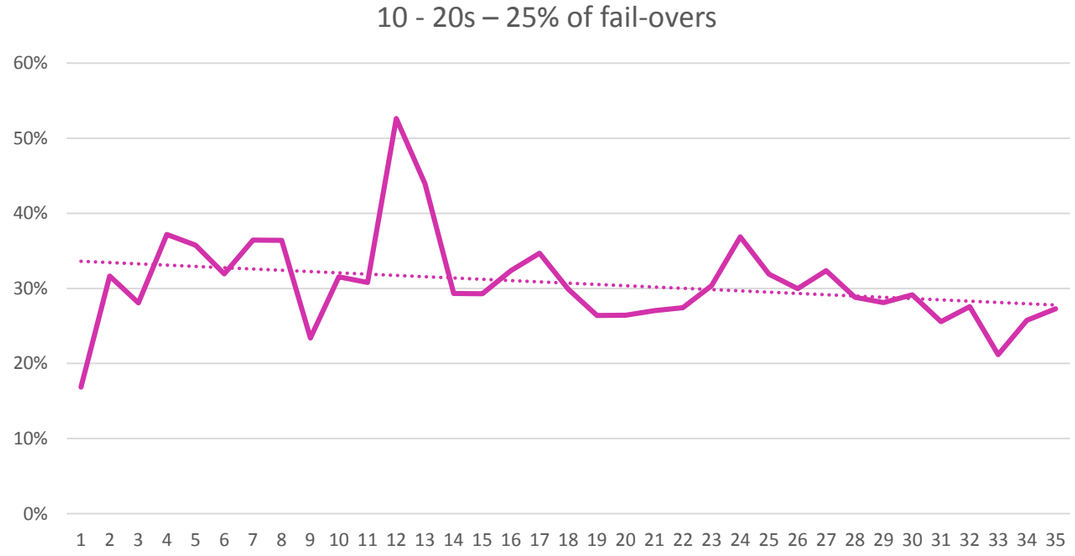

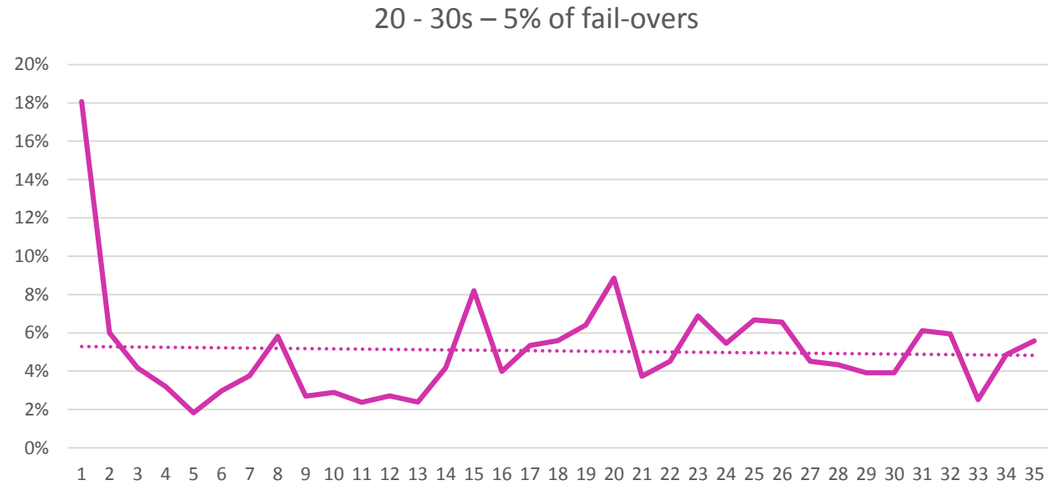

Fast database cloning

Create a copy of a database without duplicate storage costs

- Creation of a clone is nearly instantaneous – we don’t copy data
- Data copy happens only on write – when original and cloned volume data differ

Typical use cases:

- Clone a production DB to run tests
- Reorganize a database
- Save a point in time snapshot for analysis without impacting production system.

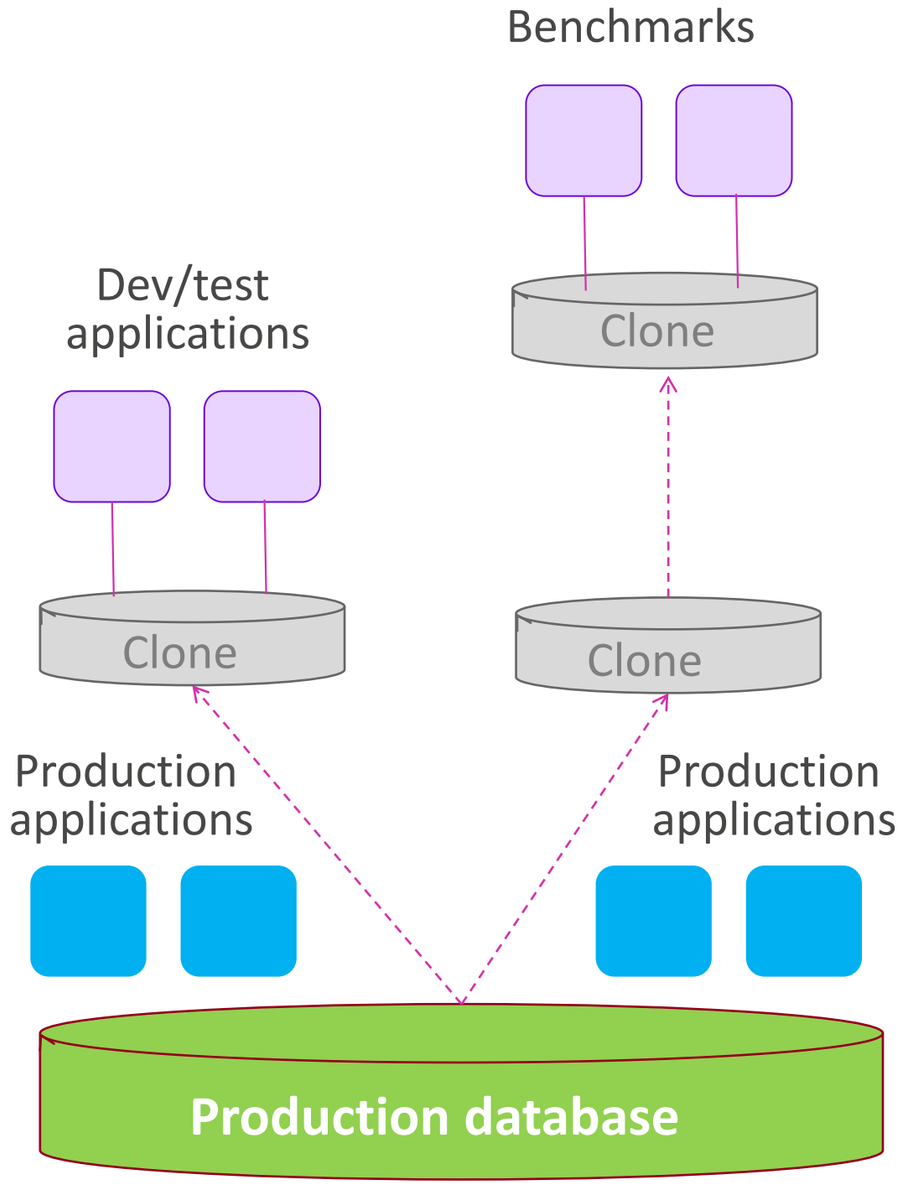

https://aws.amazon.com/blogs/aws/amazon-aurora-fast-database-cloning/

Database Backtrack

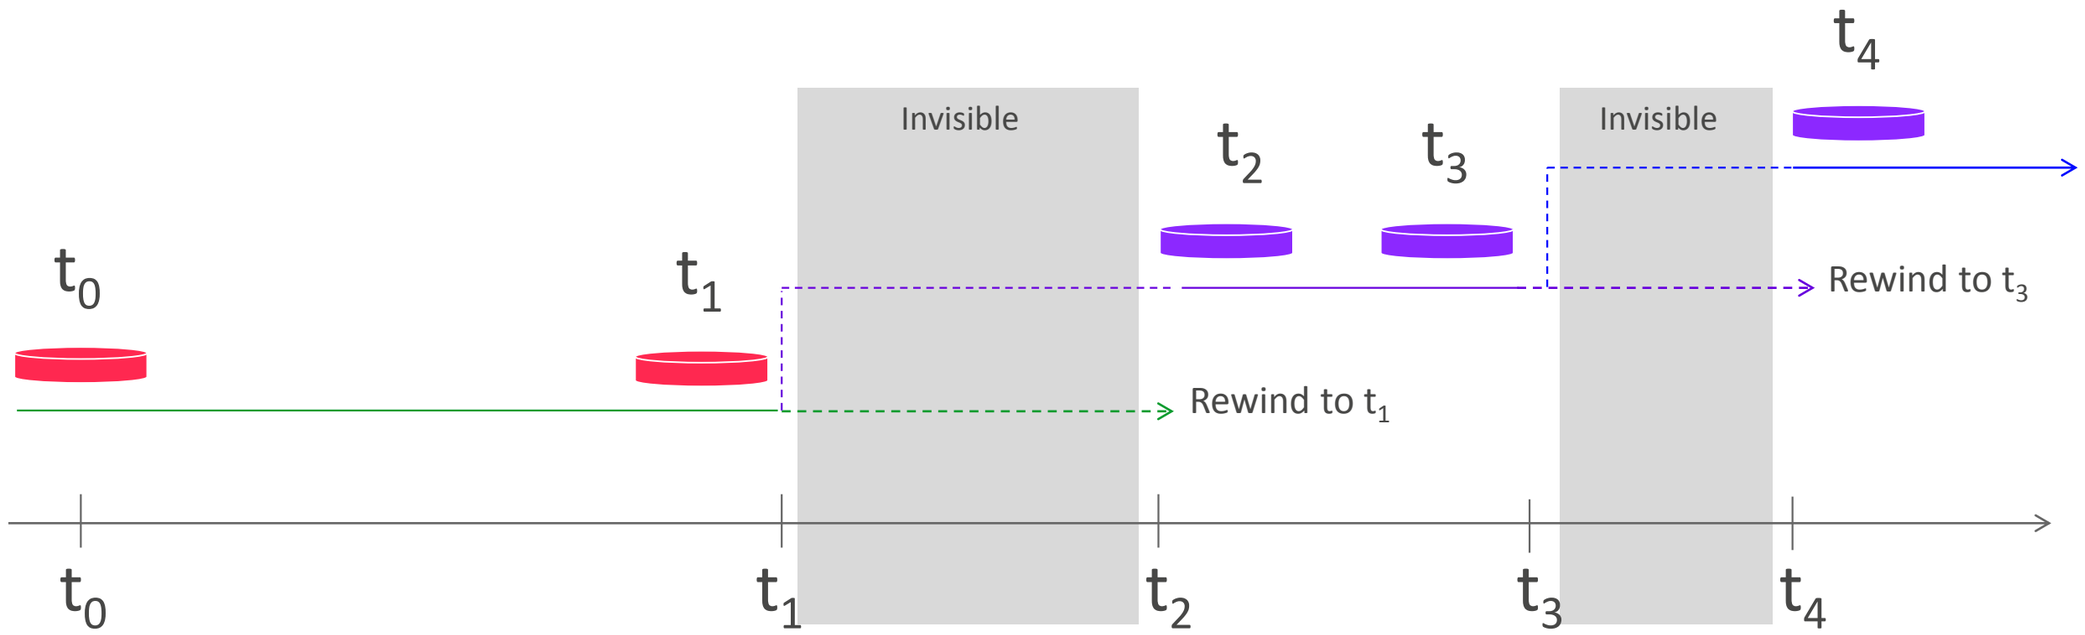

Backtrack brings the database to a point in time without requiring restore from backups

- Backtracking from an unintentional DML or DDL operation
- Backtrack is not destructive. You can backtrack multiple times to find the right point in time
- Also useful for QA (rewind your DB between test runs)

https://aws.amazon.com/blogs/aws/amazon-aurora-backtrack-turn-back-time/

Aurora Serverless

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Aurora Serverless

Starts up on demand, shuts down when not in use

Scales up/down automatically

No application impact when scaling

Pay per second, 1 minute minimum

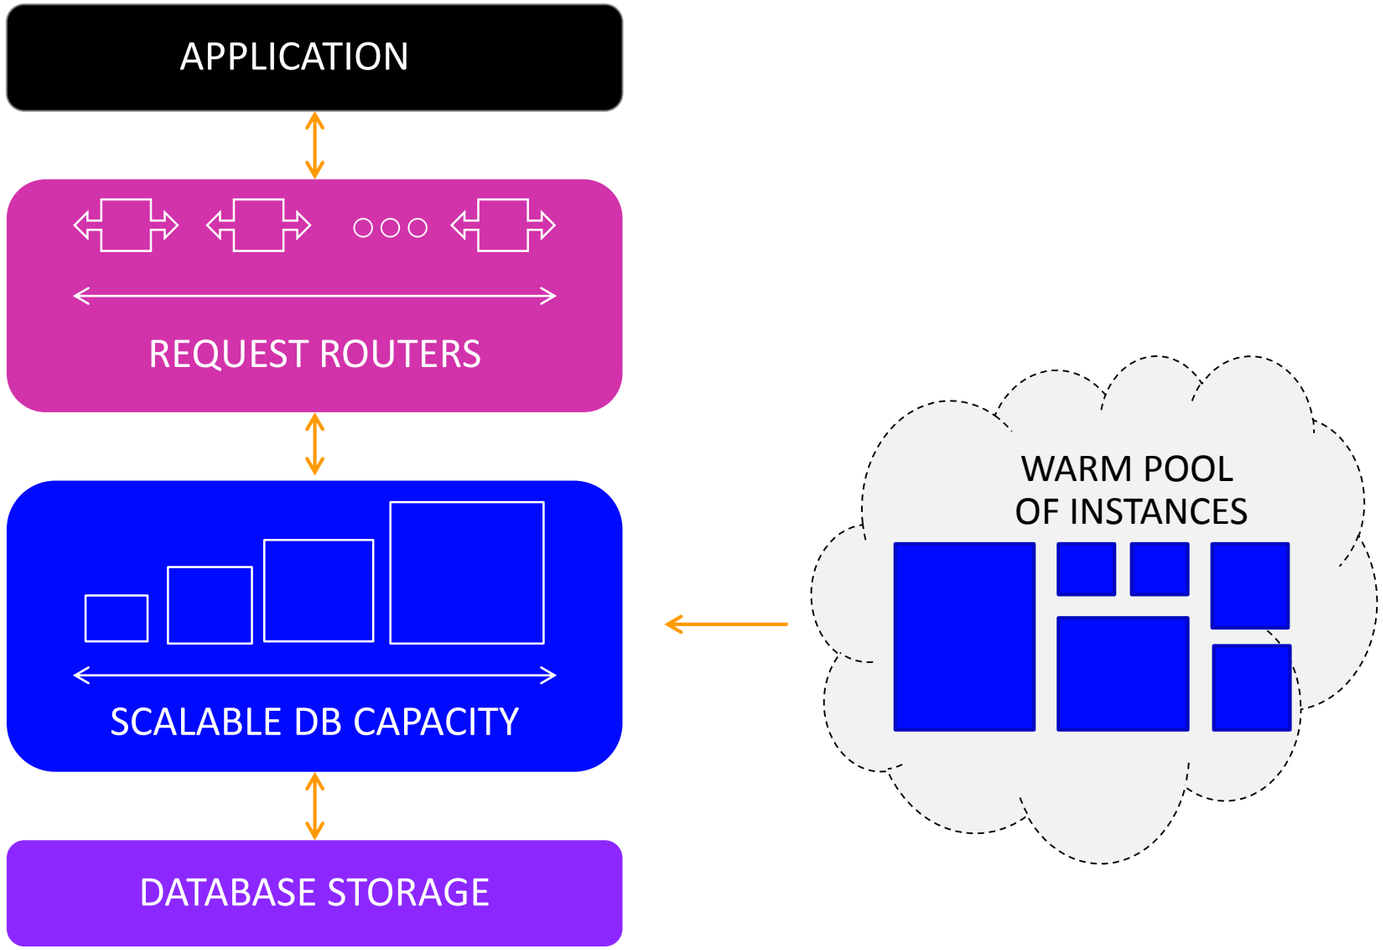

https://aws.amazon.com/getting-started/tutorials/configure-connect-serverless-mysql-database-aurora/
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.html

Amazon RDS Data API for serverless applications

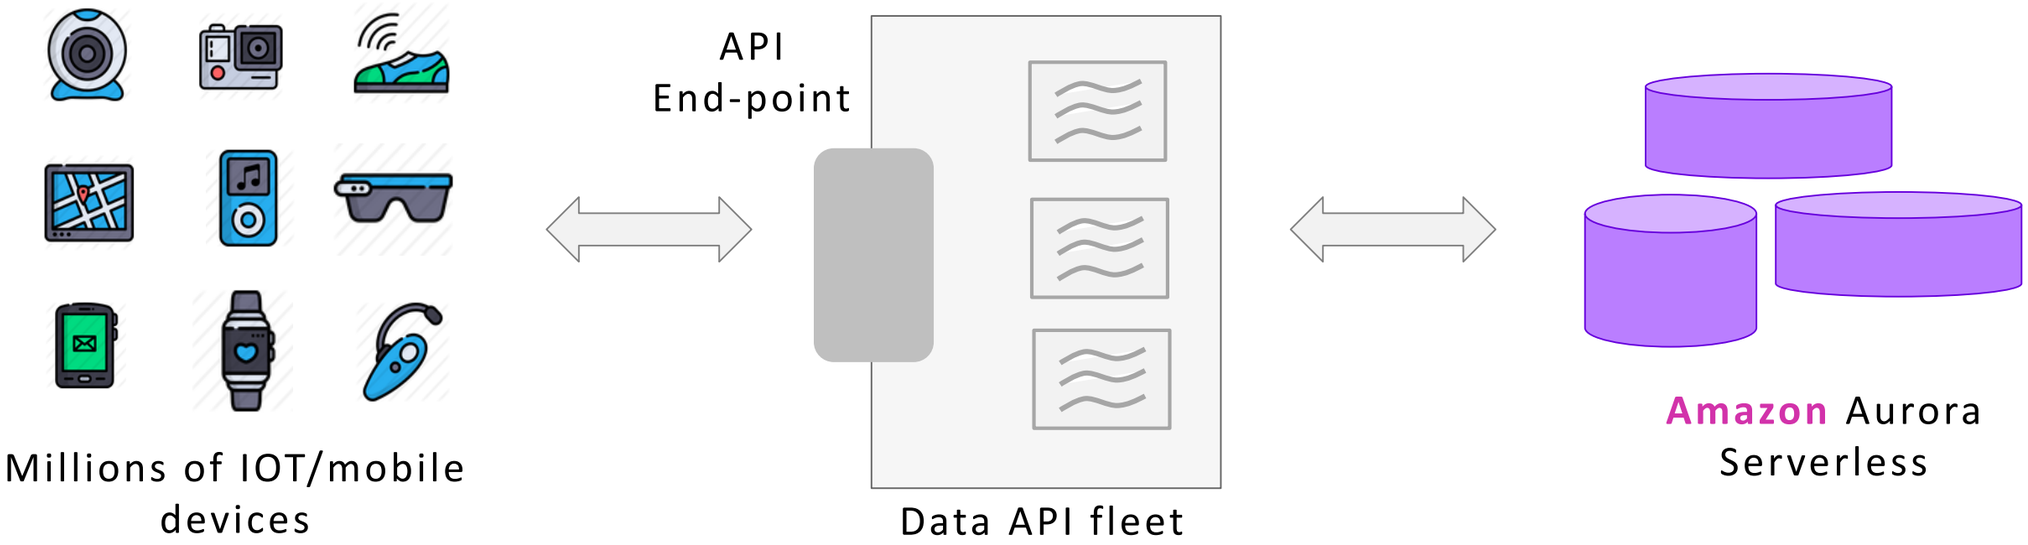

Access through simple web interface

- Public endpoint addressable from anywhere
- No client configuration required
- No persistent connections required

Ideal for Serverless applications (Lambda)

Ideal for light-weight applications (IOT)

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/data-api.html

Aurora Serverless Demo

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

Thank you!

Yoav Eilat
yeilat@amazon.com

© 2018, Amazon Web Services, Inc. or its Affiliates. All rights reserved. Amazon Confidential and Trademark