AWS re:Invent

DAT350-R1

Migrating open-source databases from Amazon EC2 to Amazon Aurora MySQL

Tiffany Jianto
Sr. Software Dev Engineer
RDS Aurora MySQL

Ahmad Elayat
Sr. Product Manager
RDS Aurora MySQL

Kamal Gupta
Sr. Software Dev Manager
RDS Aurora MySQL

AWS re:Invent

aws

Agenda

Amazon Aurora Overview

Migrating to Amazon Aurora MySQL

- RDS, open source, and commercial engines

Q&A

Amazon Aurora Overview

Relational databases are complex

Our experience running Amazon.com taught us that relational databases can be a pain to manage and operate with high availability

It's expensive and complex to manage administrative functions including regular patching cycles, performance optimization, and backup and disaster recovery—all for constantly changing applications

Amazon Aurora

MySQL and PostgreSQL-compatible relational database built for the cloud

Performance and availability of commercial-grade databases at 1/10th the cost

Performance and scalability

Up to 5x throughput of standard MySQL and 3x of standard PostgreSQL; scale-out up to 15 read replicas

Availability and durability

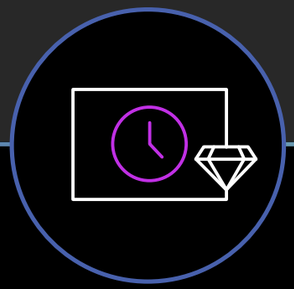

Fault-tolerant, self-healing storage; six copies of data across three Availability Zones; continuous backup to Amazon S3

Highly secure

Network isolation, encryption at rest, compliance and assurance programs

Fully managed

Managed by RDS:
No server provisioning, software patching, setup, configuration, or backups

The compute and storage layers are separated in a scale-out, distributed, and multi-tenant 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

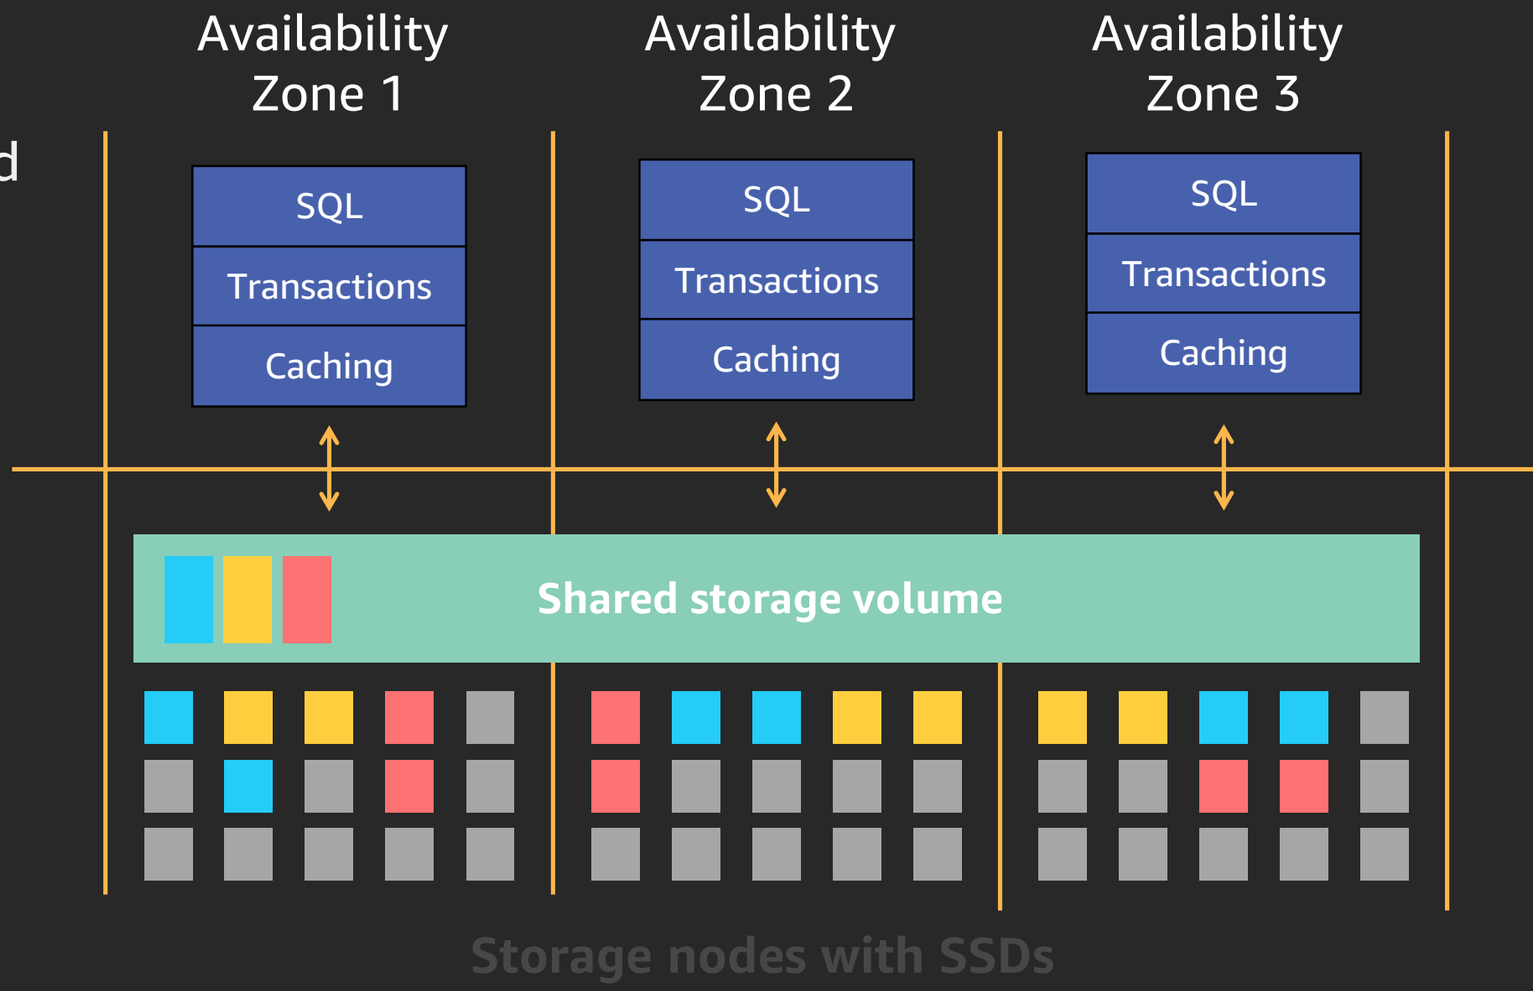

As a fully managed service, Aurora automates administrative tasks

Takes care of your time-consuming database management tasks, freeing you to focus on your applications and business

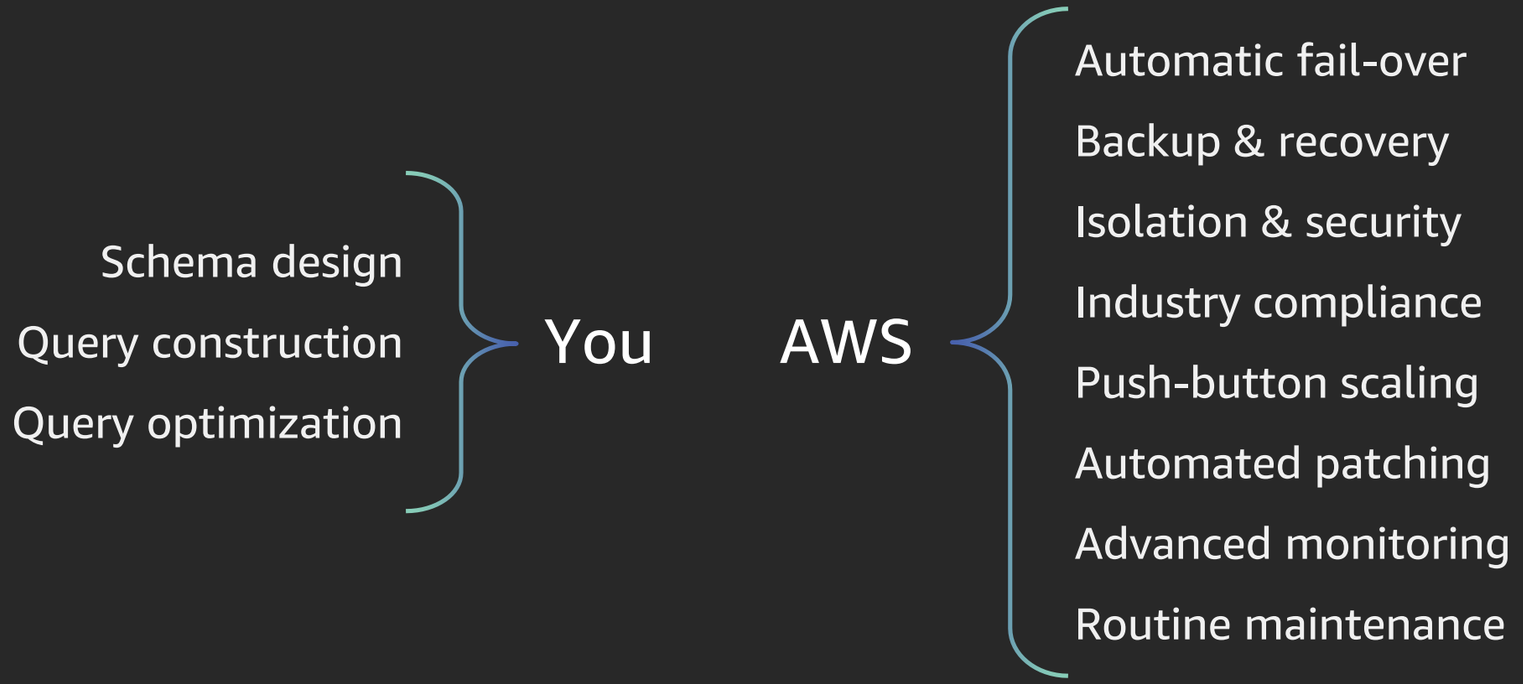

Both customers using open source engines and commercial engines are moving to Aurora

Customers using open source engines

- Higher performance – up to 5x
- Better availability and durability
- Reduces cost – up to 60%
- Easy migration; no application change

Customers using commercial engines

- One tenth of the cost; no licenses
- Integration with cloud ecosystem
- Comparable performance and availability
- Migration tooling and services

Aurora is the fastest growing service in AWS history

Tens of thousands of enterprises and startups run their applications on Aurora

Recently Released RDS Aurora MySQL Key Features

Serverless

Multi-Master

Fast Database Cloning

Database Backtrack

Global Database

Parallel Query Processing

Read Replica and custom end-point

Performance Insights

Autodesk makes software for people who make things like high-performance cars, towering skyscrapers and great films

“Our Account Control Management (ACM) application is a central component for many Autodesk applications, and usage has grown this year. ACM requires a high performance database with low-lag replication, automatic scaling, and MySQL compatibility. Amazon Aurora allowed us to increase our database connections and greatly improve our scalability, response times, and CPU utilization. Our customers’ experience improved in a short time so we’re looking into Aurora for additional applications.”

-Brian Mathews, Platform Engineering Vice President - Autodesk

https://aws.amazon.com/blogs/database/autodesk-builds-on-amazon-aurora/

Migrating to Aurora MySQL

AWS re:Invent

aws

Amazon Aurora migration options

Source database	From where	Recommended option
 	RDS	Console based automated snapshot ingestion and catch up via binlog replication.
   	EC2, on premises	Binary snapshot ingestion through S3 and catch up via binlog replication.
    	EC2, on premises, RDS	Schema conversion using SCT and data migration via DMS.

EC2 MySQL to Aurora MySQL – other options

- Binary snapshot ingestion from Amazon S3 (preferred)
- Native MySQL export and import tools
- AWS Database Migration Service

Binary snapshot ingestion (recommended)

- Managed backup ingestion from Amazon S3
- High performance
- Amazon Aurora and Amazon S3
    - Create Amazon Aurora cluster from Amazon S3 snapshot
    - Import data stored in Amazon S3 bucket into Amazon Aurora database

Self-managed import/export via MySQL tools

- Most compatibility—schema can be migrated as-is
- Both native MySQL and third-party tools supported by well-documented command-line utilities
- Full control over the migration process
- When to use this:
    - Smaller scale migrations
    - Strong network connection between source and target databases
    - Migration performance is not critical
    - Cost of retrying migration is low
    - No intermediate schema or data transformations

Migration via AWS DMS

- Managed service via AWS Management Console
- Supports heterogeneous and homogeneous migrations
- Advantageous when:
  - Data transformations
  - Migrating and replicating multiple database servers into a single Aurora cluster
- Limitations:
  - Does not migrate secondary schema objects—must be migrated or created manually prior to data migration

Best Practices

- Review your database with Amazon Aurora feature overview
- Test migration and workload
- Rollback procedure: configure MySQL read replica
- Upgrading/downgrading between MySQL versions
- Cluster and database parameters – based on workload

Autodesk makes software for people who make things like high-performance cars, towering skyscrapers and great films

“Our Account Control Management (ACM) application is a central component for many Autodesk applications, and usage has grown this year. ACM requires a high performance database with low-lag replication, automatic scaling, and MySQL compatibility. Amazon Aurora allowed us to increase our database connections and greatly improve our scalability, response times, and CPU utilization. Our customers’ experience improved in a short time so we’re looking into Aurora for additional applications.”

-Brian Mathews, Platform Engineering Vice President - Autodesk

https://aws.amazon.com/blogs/database/autodesk-builds-on-amazon-aurora/

Q&A

Thank you!

AWS re:Invent

aws

Write and read throughput

Aurora MySQL delivers up to 5x more throughput than MySQL

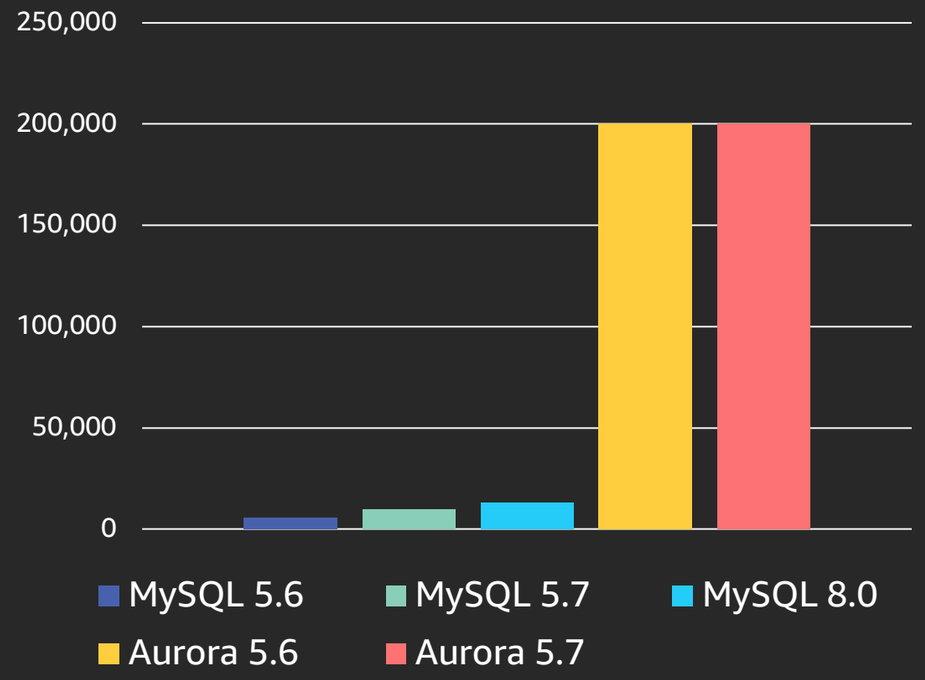

Write Throughput

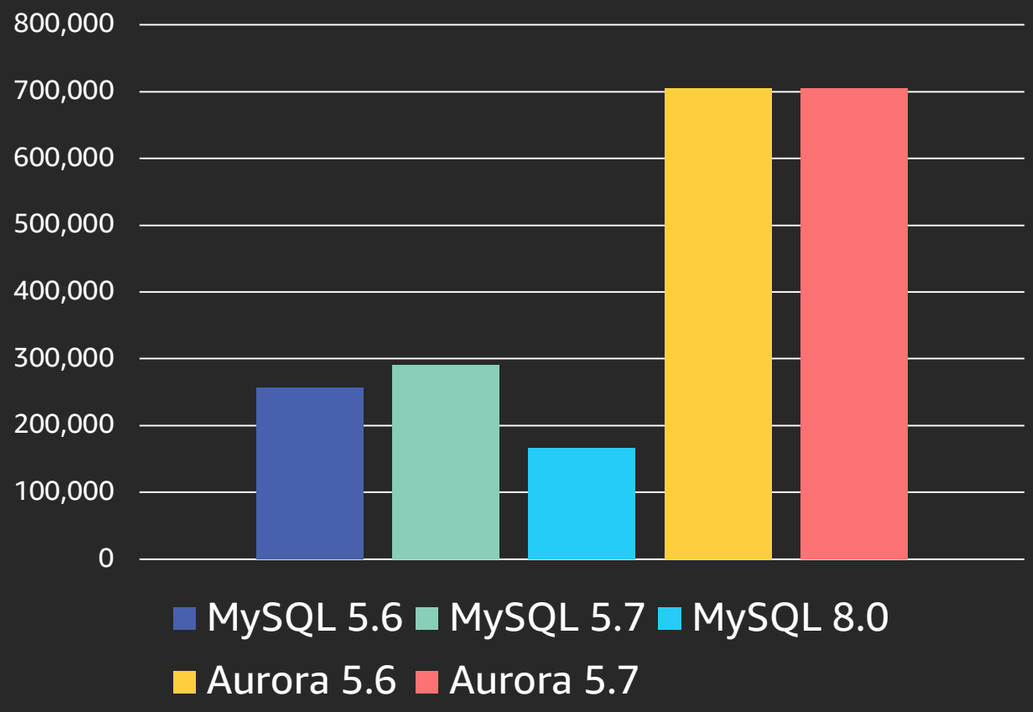

Read Throughput

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