
AWS
re:Invent

DAT302

Deep dive on Amazon Aurora with MySQL compatibility

Murali Brahmadesam

Director of Engineering, Aurora Databases & Storage

AWS

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Aurora

ENTERPRISE DATABASE AT OPEN-SOURCE PRICE

Delivered as a **managed** service



Performance 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

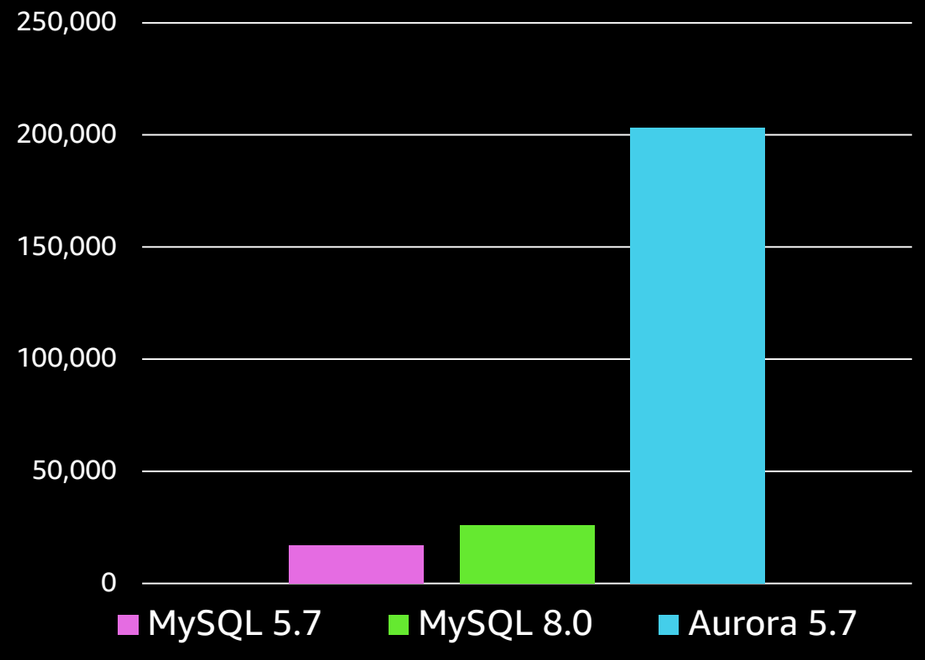
Performance



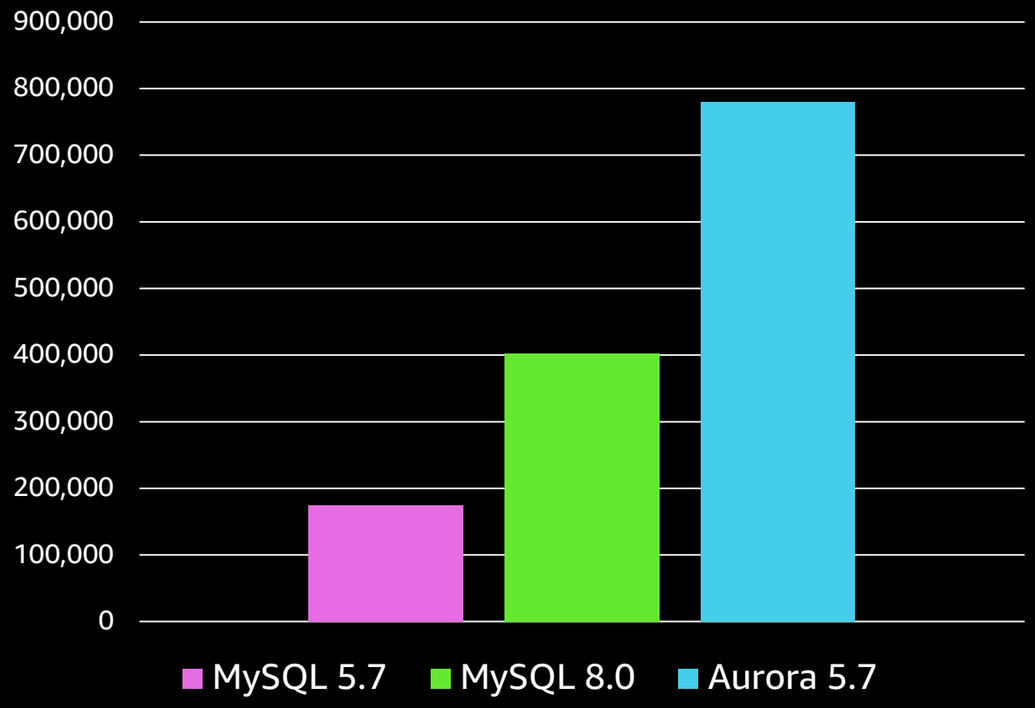
© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Write and read throughput

AMAZON AURORA MYSQL DELIVERS UP TO 5X THE THROUGHPUT OF MYSQL



Write throughput



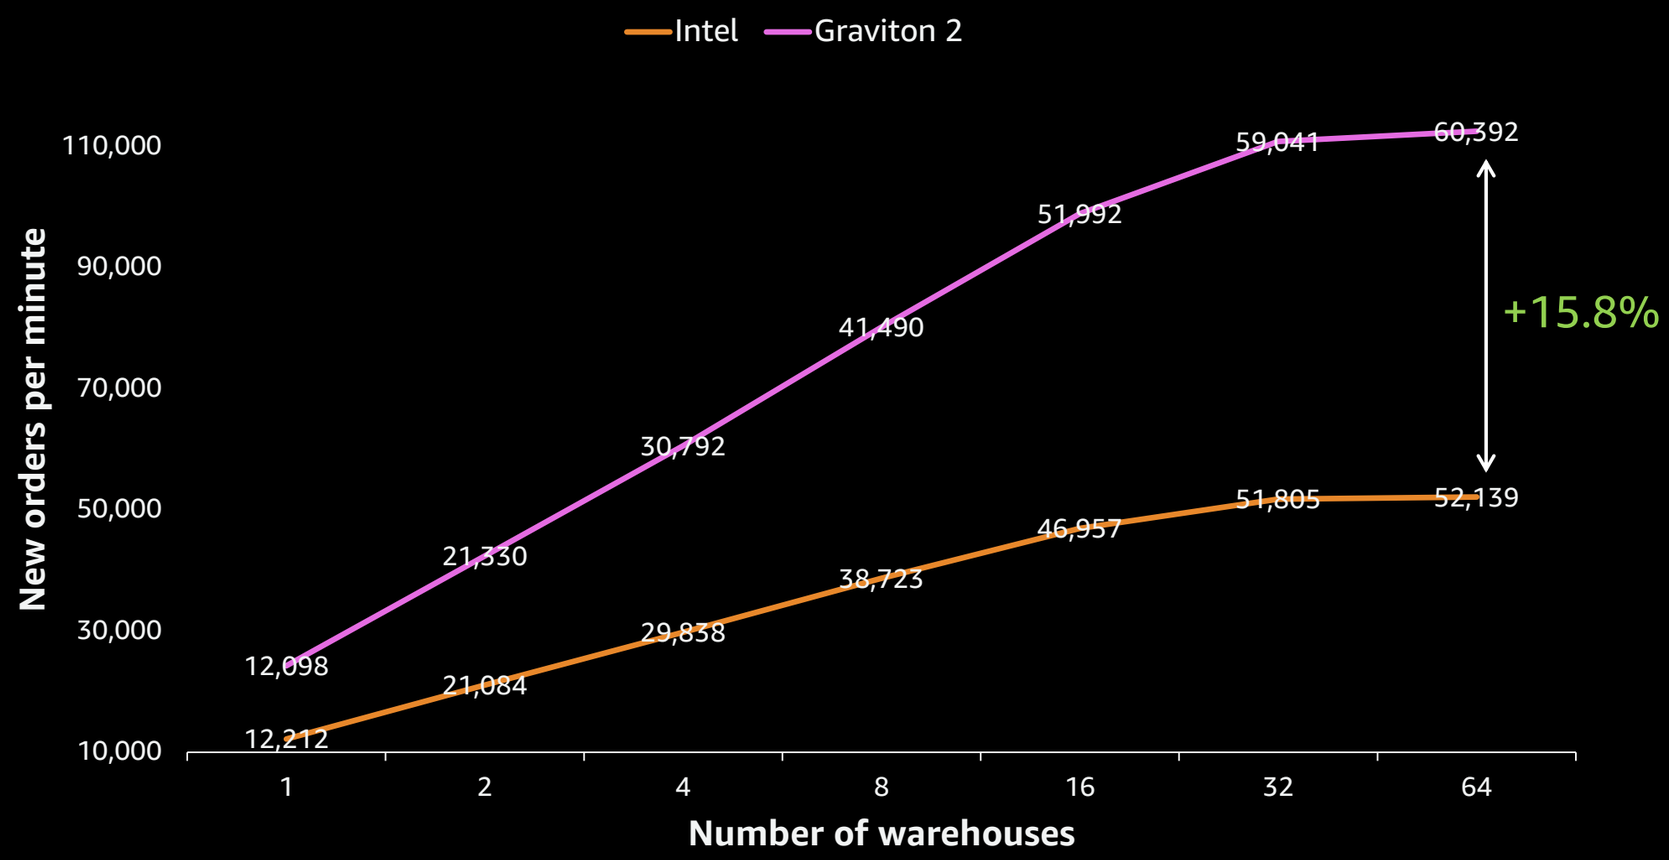
Read throughput

Using sysbench read and write workloads on R5.16XL

Amazon Aurora MySQL support for 8.0 upcoming

AWS Graviton2 instances

AWS'S OWN 64 BIT ARM-BASED EC2 INSTANCES



Write performance variability under load

AURORA IS MORE THAN 100X CONSISTENT



Using sysbench write-only workload on R5-12XL

Write and read throughput

HOW DID WE ACHIEVE THIS?

Do less work

- Do fewer IOs
- Minimize network packets
- Cache prior results
- Offload the database engine

Be more efficient

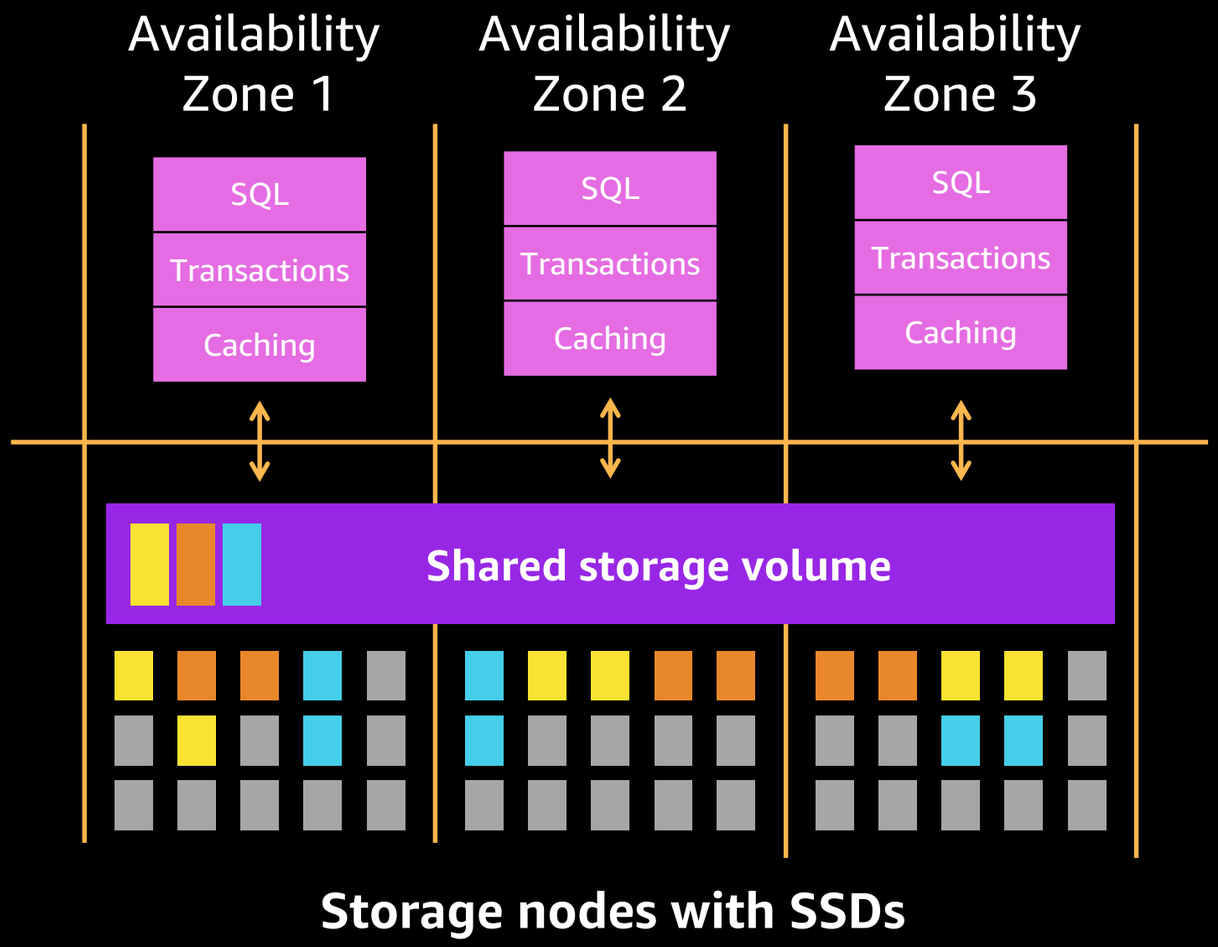
- Process asynchronously
- Reduce latency path
- Use lock-free data structures
- Batch operations together

Write and read throughput

HOW DID WE ACHIEVE THIS?

Log is the database

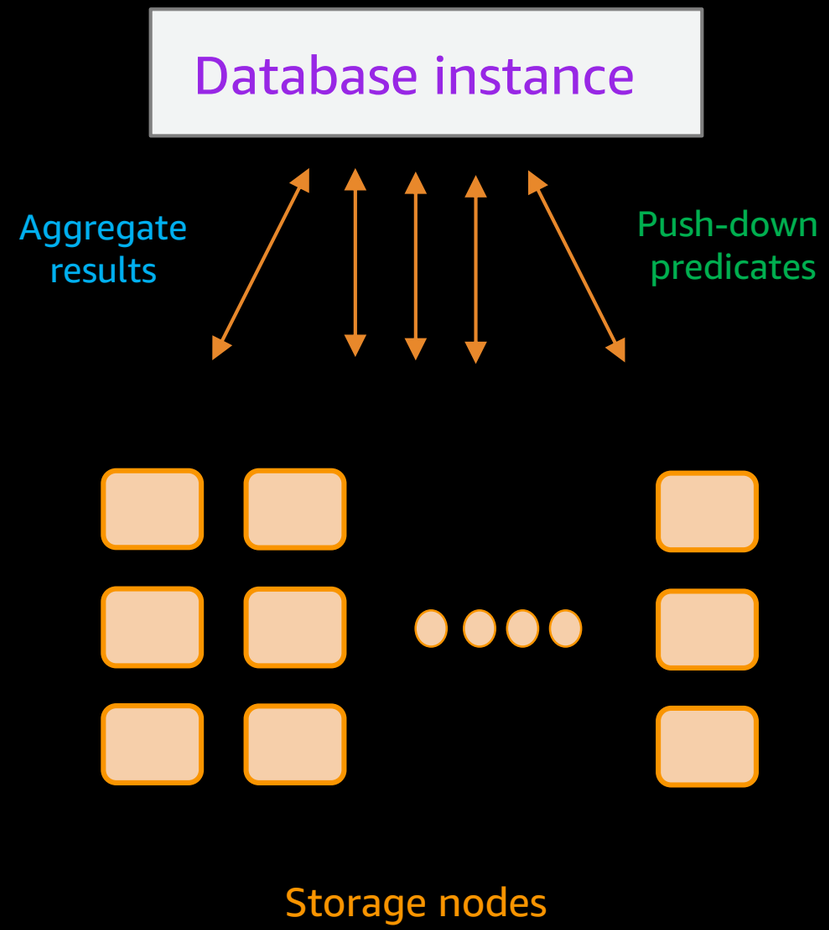
- Offload redo processing to distributed storage that spans thousands of storage nodes
- Write only logs and no database checkpoints
- Use 4/6 write quorum across 3 AZs with 2 copies per AZ



Parallel query (PQ)

USING PUSH-DOWN QUERY EXECUTION

- Aurora database is hosted on top of thousands of storage nodes whose CPUs are used for query processing
- Database engine identifies the portion of the query that can be pushed down to the storage node, which processes the query on the rows that are resident locally
- Storage node returns result sets that are then aggregated on the database node
- Moving processing close to data significantly reduces network traffic and latency



Parallel query (PQ)

WHAT IS PUSHED DOWN?

PQ current capabilities

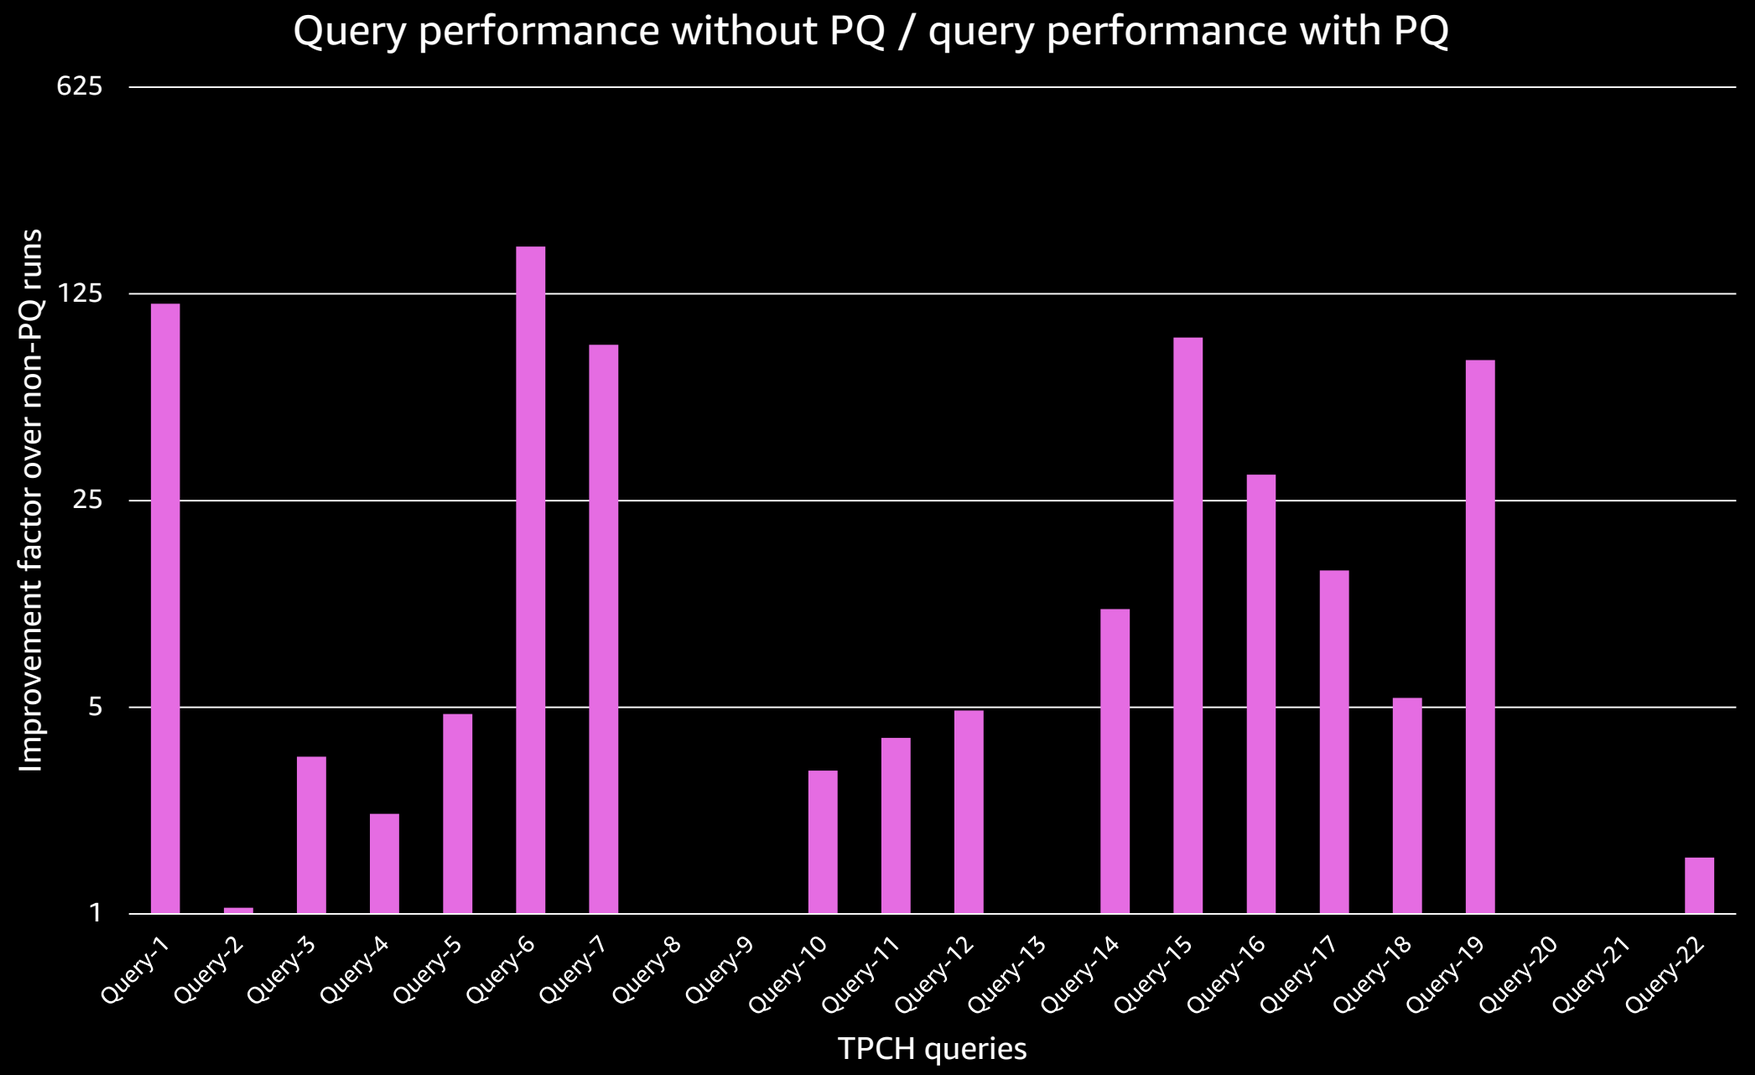
- Filters
- Projections
- Expressions

PQ capabilities (upcoming)

- Aggregations: sum / min / max / avg / count / count(distinct) /
- Group by / distinct
- Support for partition tables
- Support for LOBs
- Support for tables that were modified with Aurora MySQL fast online DDL

Parallel query using push-down query execution

UP TO 180X IMPROVEMENTS IN TPCH QUERY PERFORMANCE



Scalability



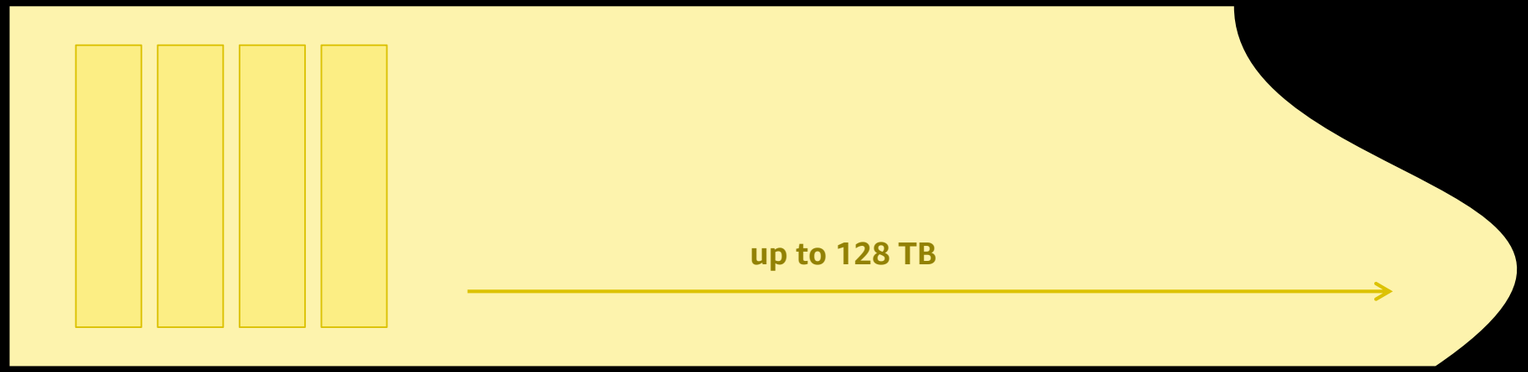
© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Increased database sizes from 64 TB to 128 TB

AUTO-MANAGED STORAGE

- Aurora databases use a purpose-built storage that automatically scales up / down based on database use up to 128 TB, up from 64 TB previously

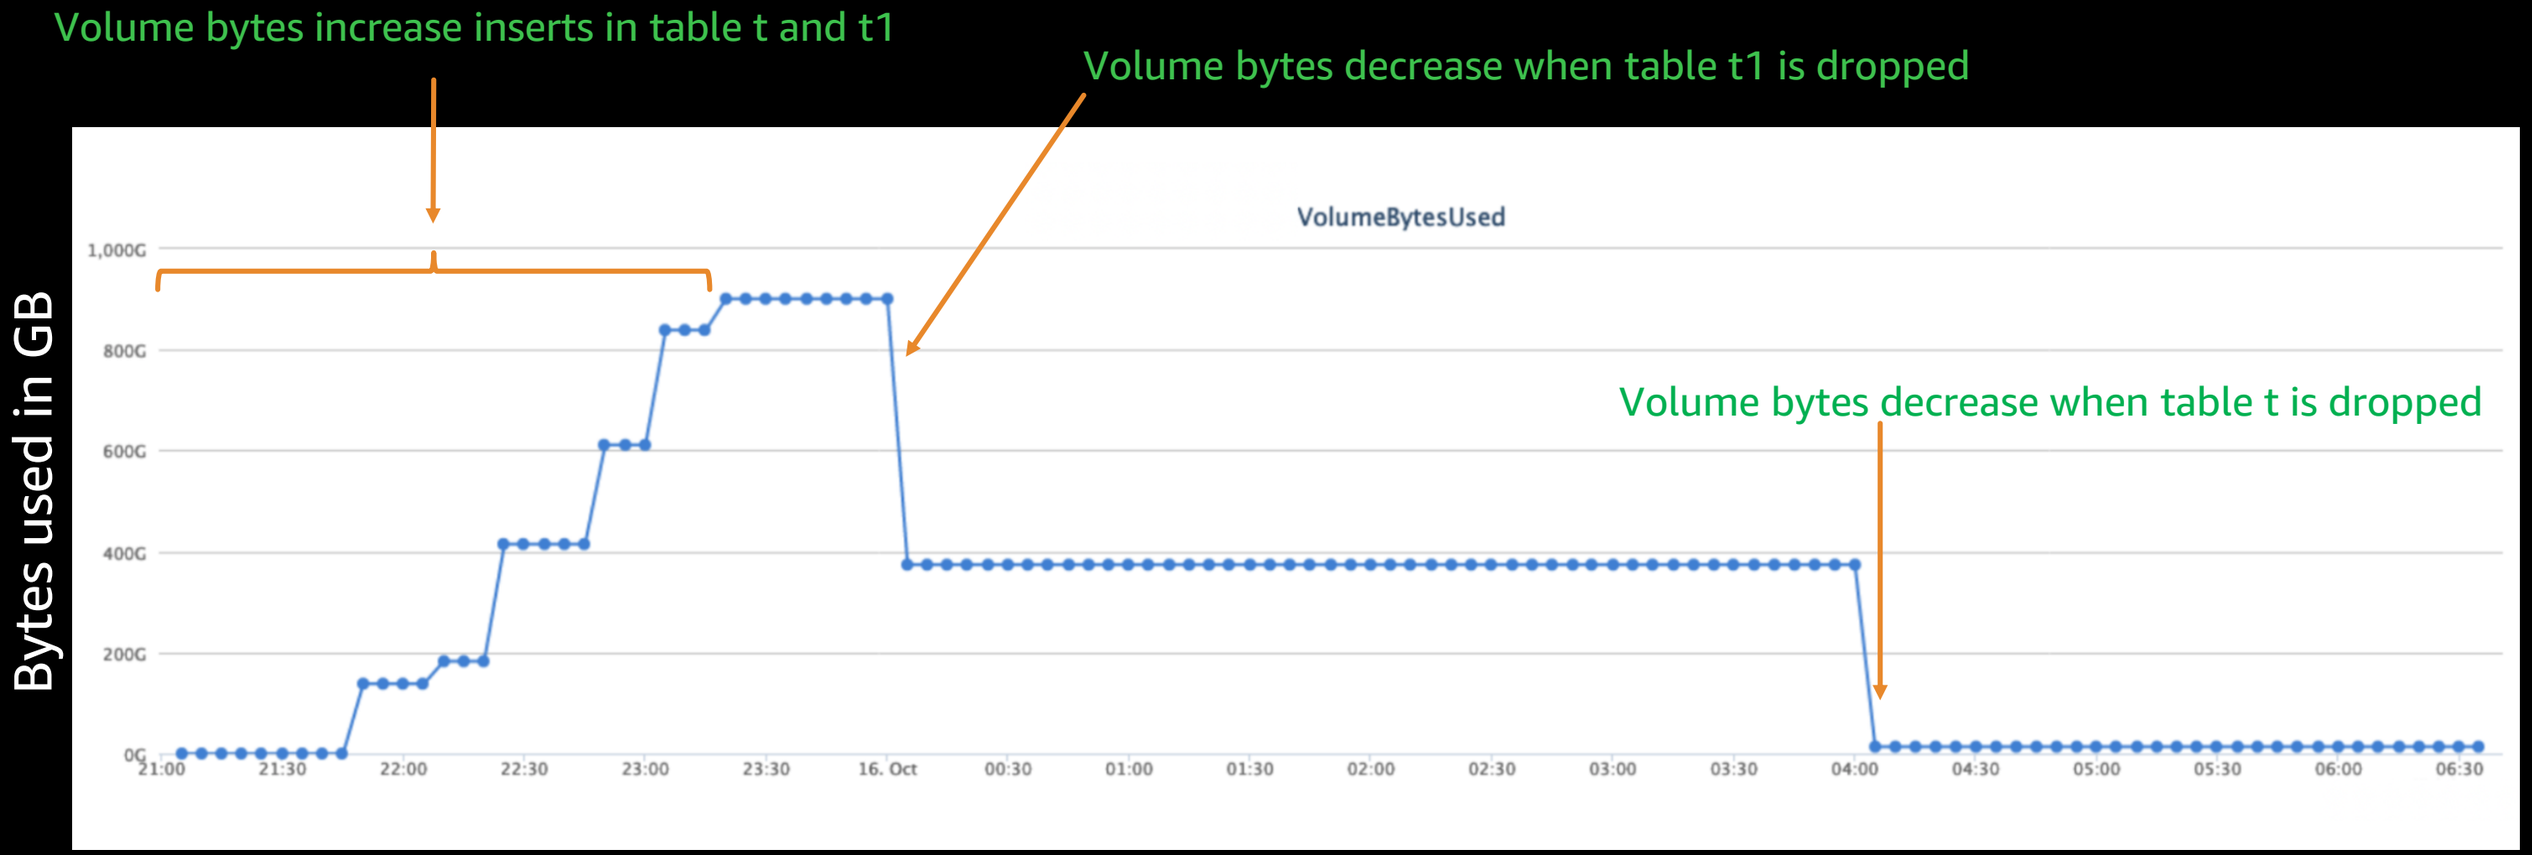
- All Aurora features continue to work with this increased size limit



Auto storage resizing with deletes

AUTO-MANAGED STORAGE

Aurora database engine informs the storage system to free up deleted space without impacting database performance



Availability



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

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

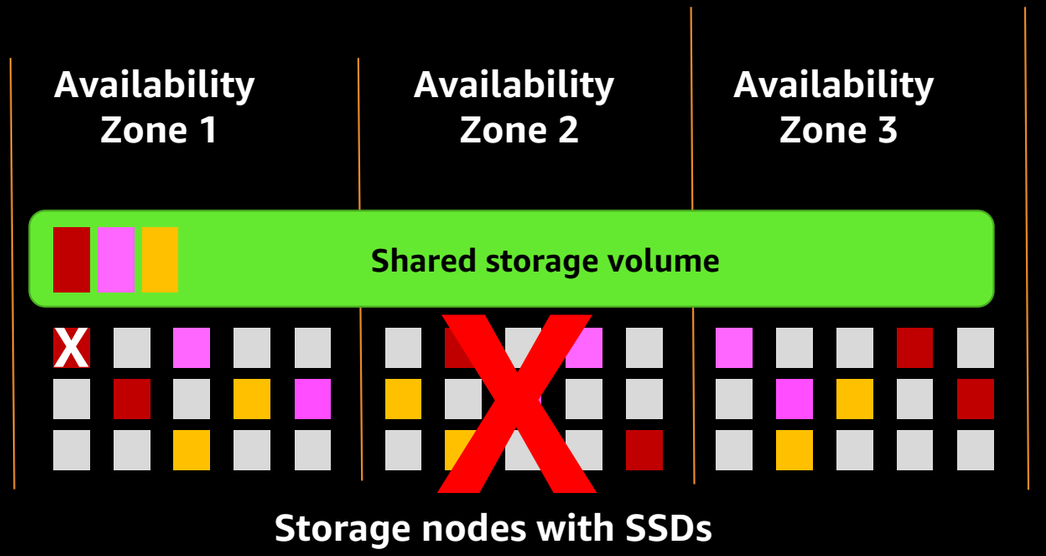
What if there is an AZ+1 failure?

⇒ Still have 3 nodes (read / repair quorum)

⇒ No data loss

⇒ Rebuild failed node by copying from one of other 3

⇒ Recover write availability



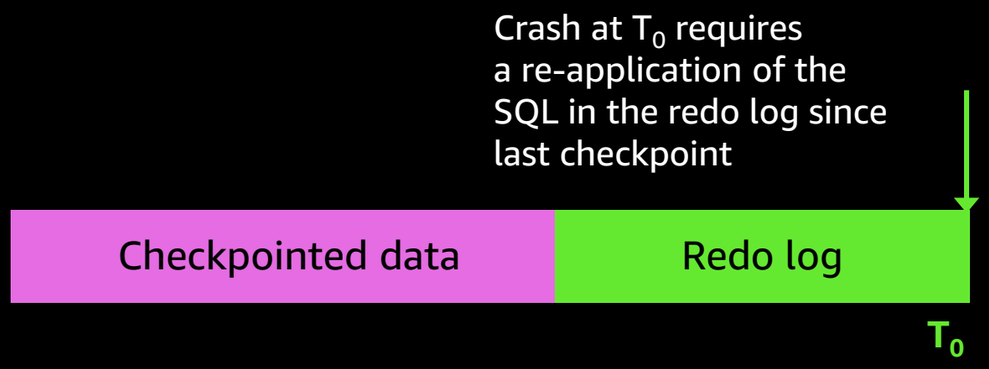
Instant crash recovery

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

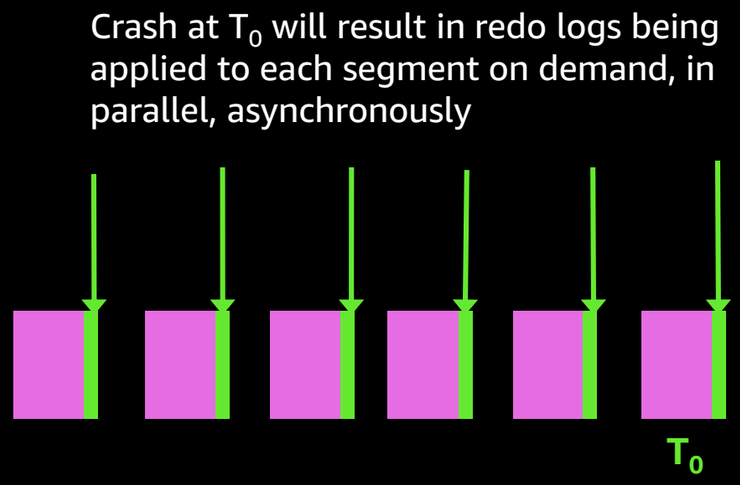


Aurora

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

Parallel, distributed, asynchronous

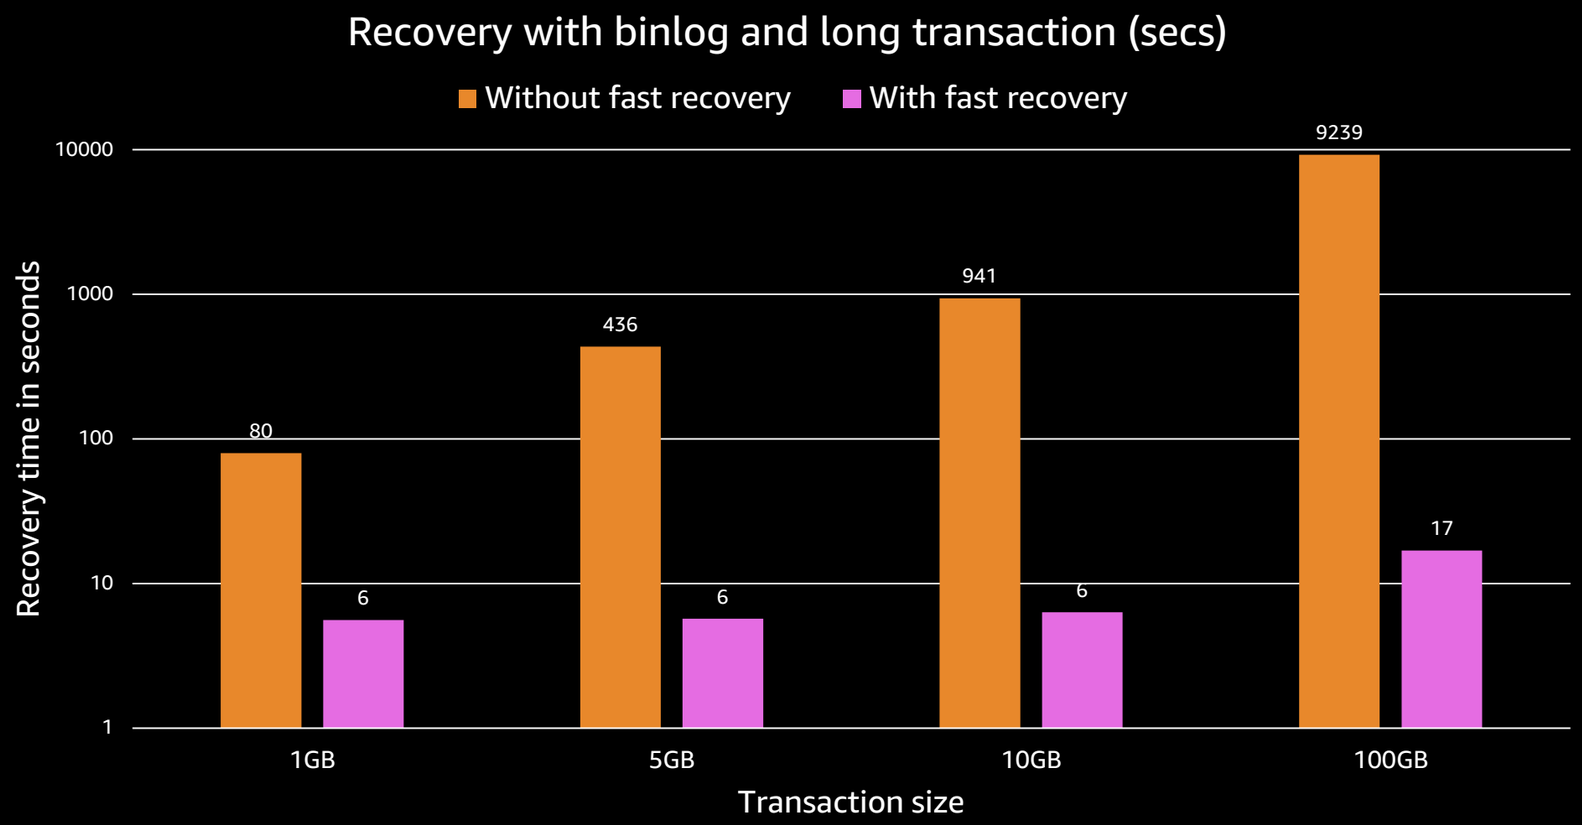
No replay for startup



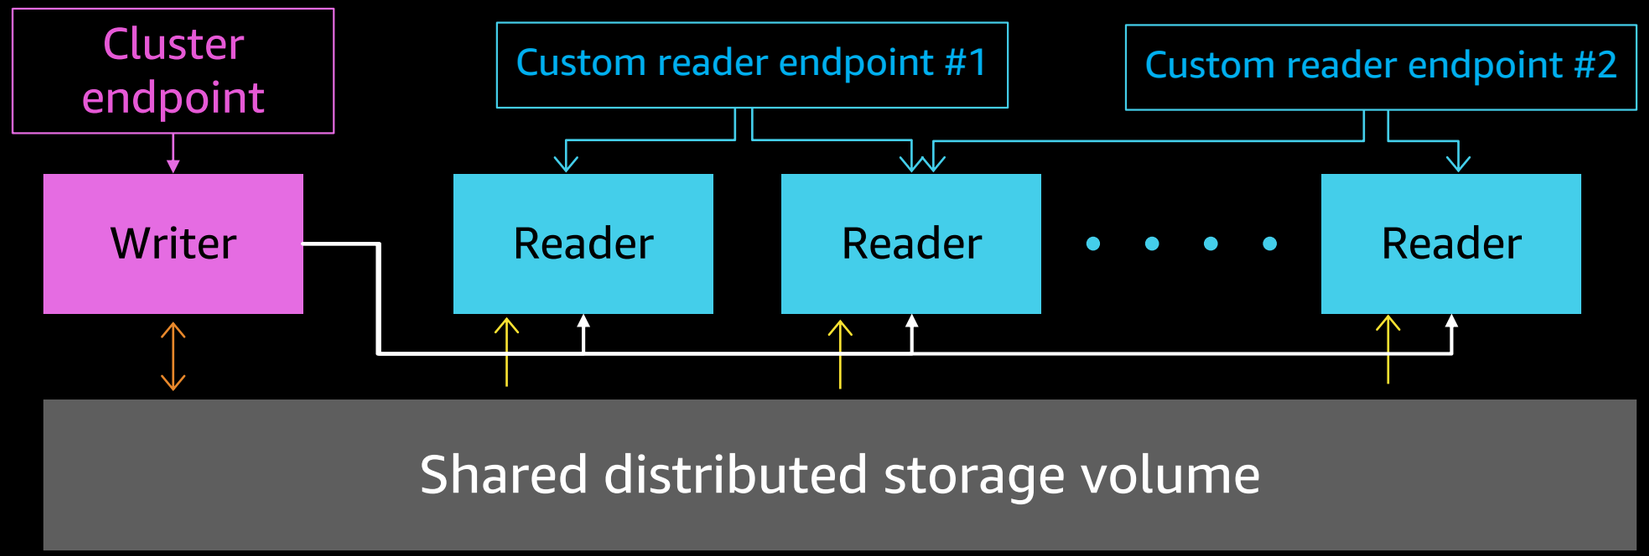
Fast recovery with binlog and long transactions

NEAR CONSTANT TIME RECOVERY

- Binlog events are kept in memory initially until they reach a certain size but then are written directly to durable storage as a temporary binlog file
- Upon commit, the current active binlog file is closed and the temporary binlog file is made the active file without any data copy
- Upon crash recovery, any temporary binlog file in durable storage is just dropped



Aurora reader instances

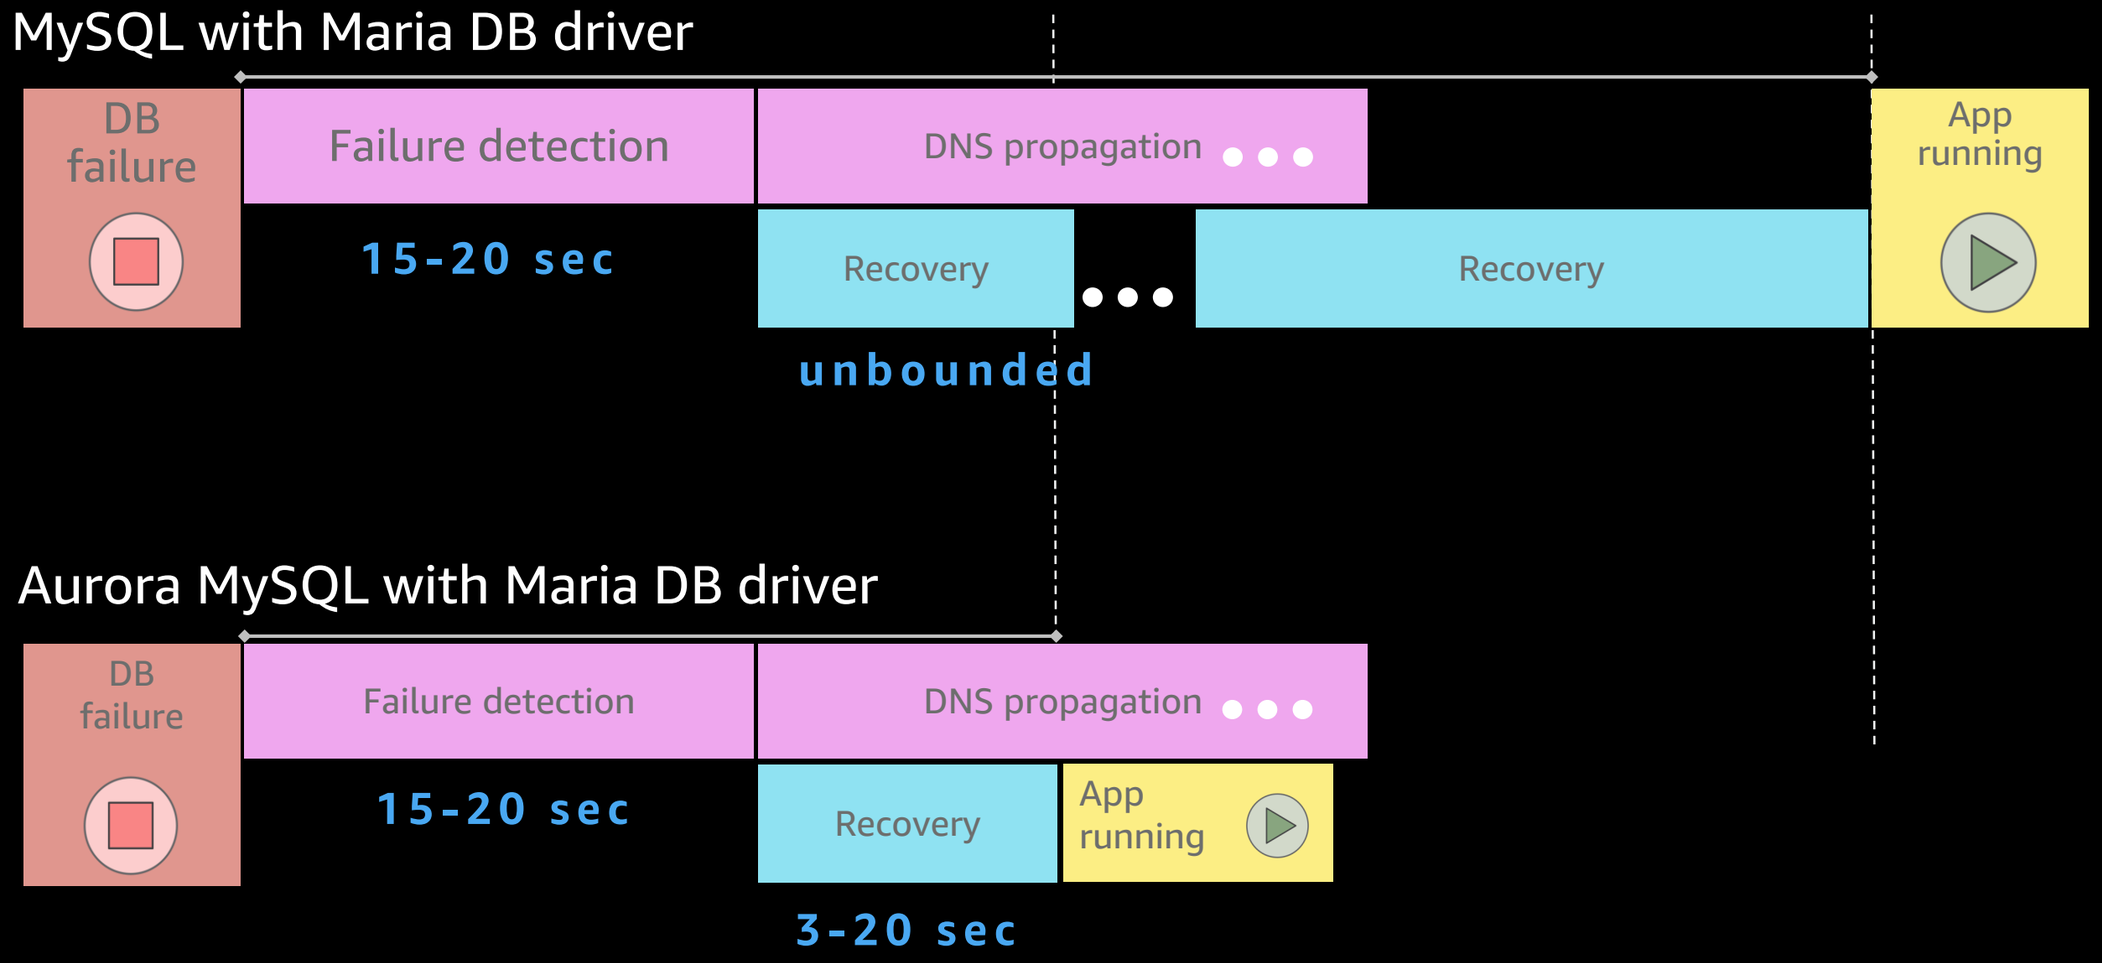


Up to 15 promotable read replicas across multiple Availability Zones

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

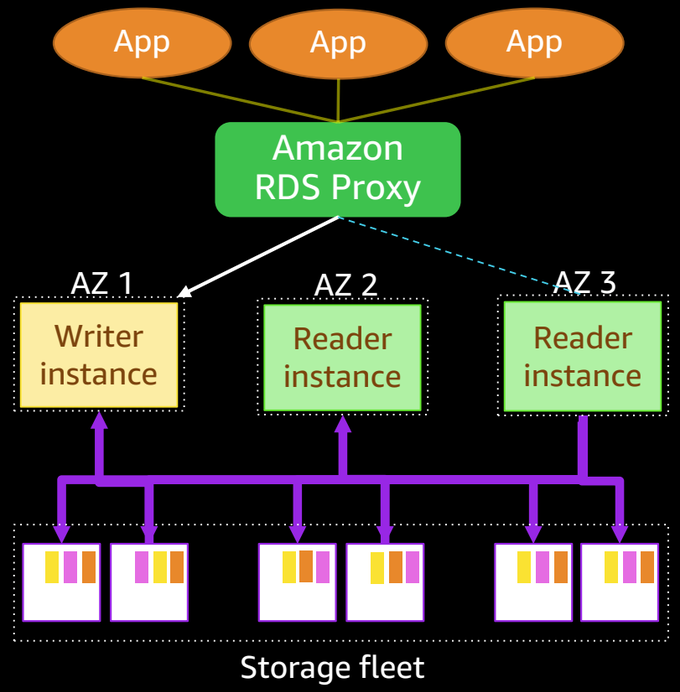
Custom reader endpoint with configurable failover order

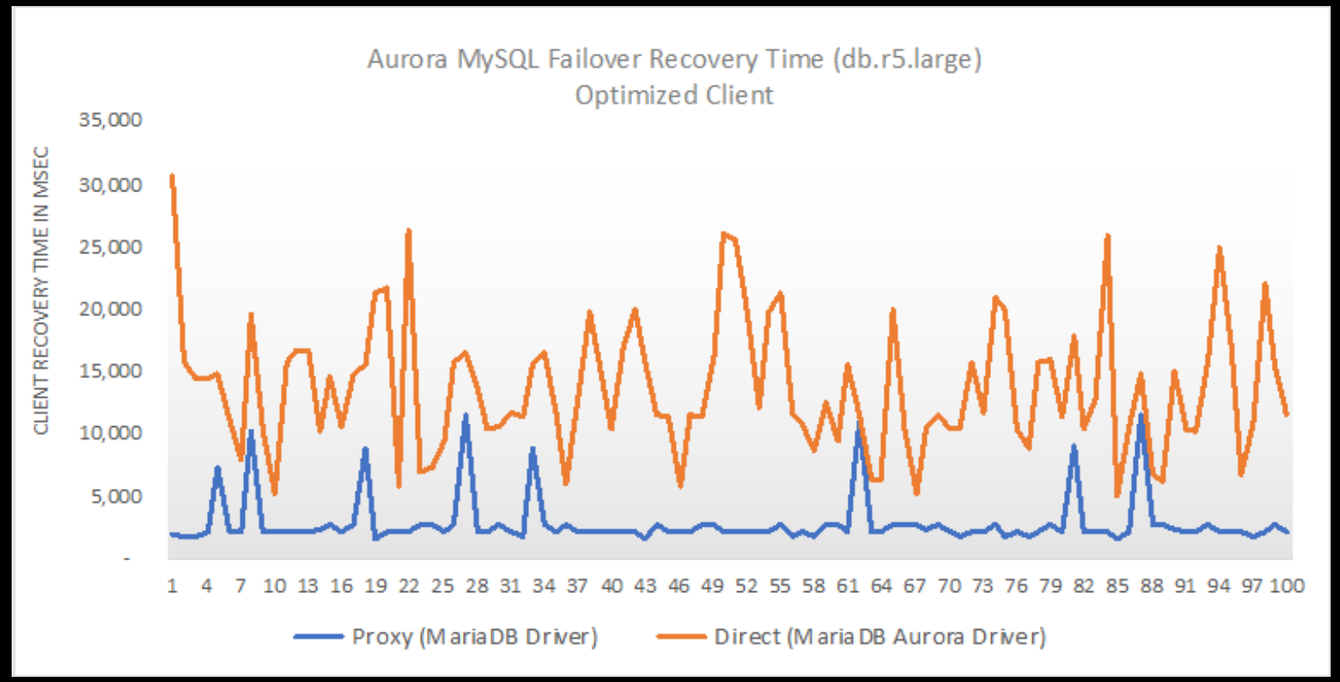
Faster failover



Seamless and faster failovers

AMAZON RDS PROXY





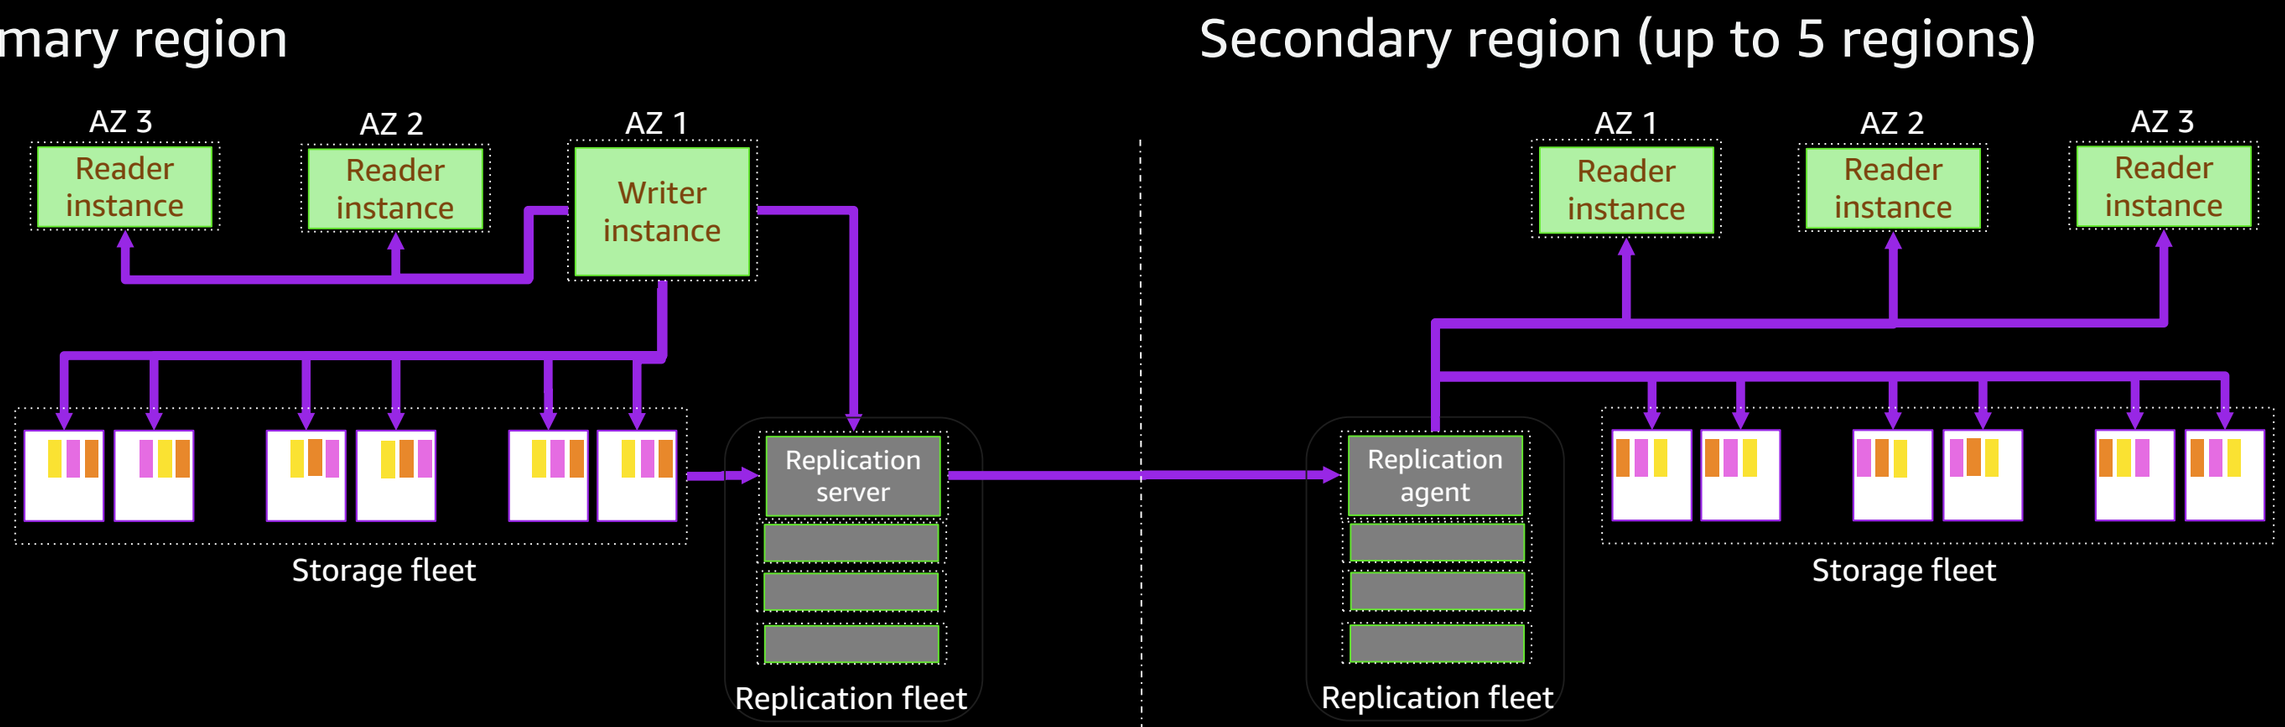
- Application connections preserved during failovers
- Detects failovers and connects to reader instance that become new writer
- Up to 66% faster failover times

Global reader instances: Up to 90

READ SCALABILITY WITH GLOBAL DATABASES

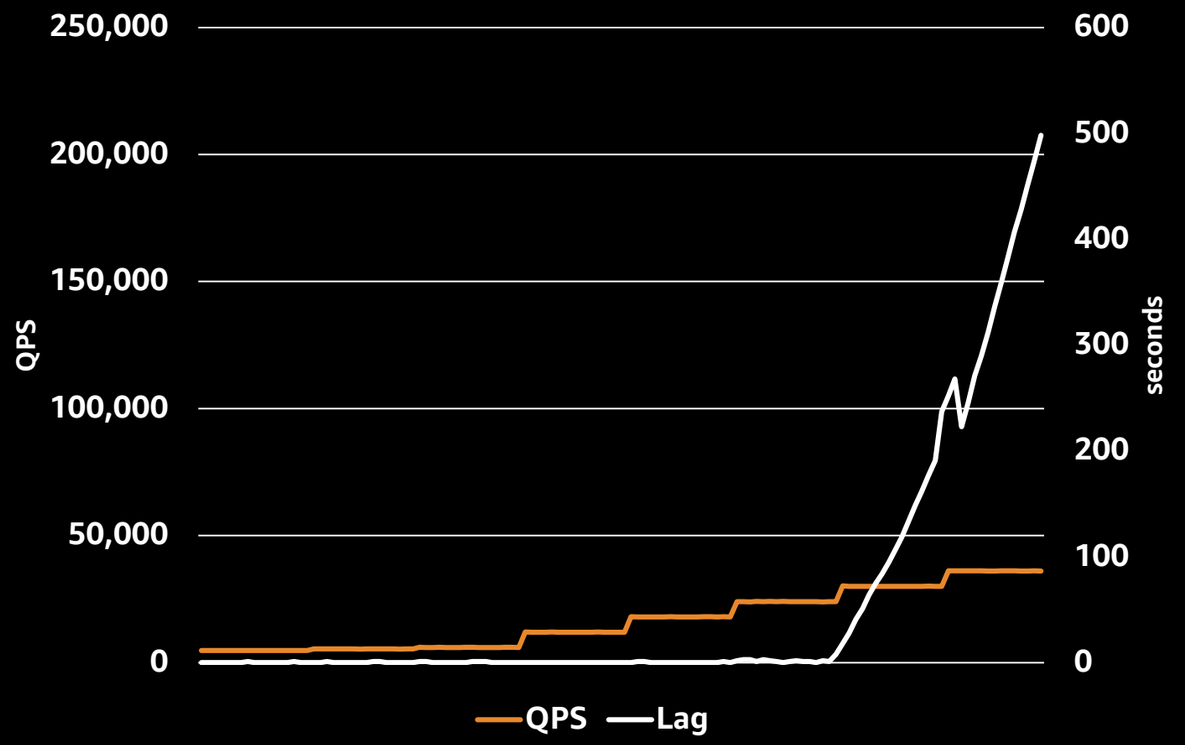
Aurora global database supports up to 90 reader instances spread across 6 regions (15 reader instances per region) without impacting the database writer performance in primary region

Low replica lag: < 1 sec, **fast failover:** < 1 min, **write forwarding:** writes to global reader instances are forwarded to primary region, **failover failback** (upcoming): without having to recreate replication configurations

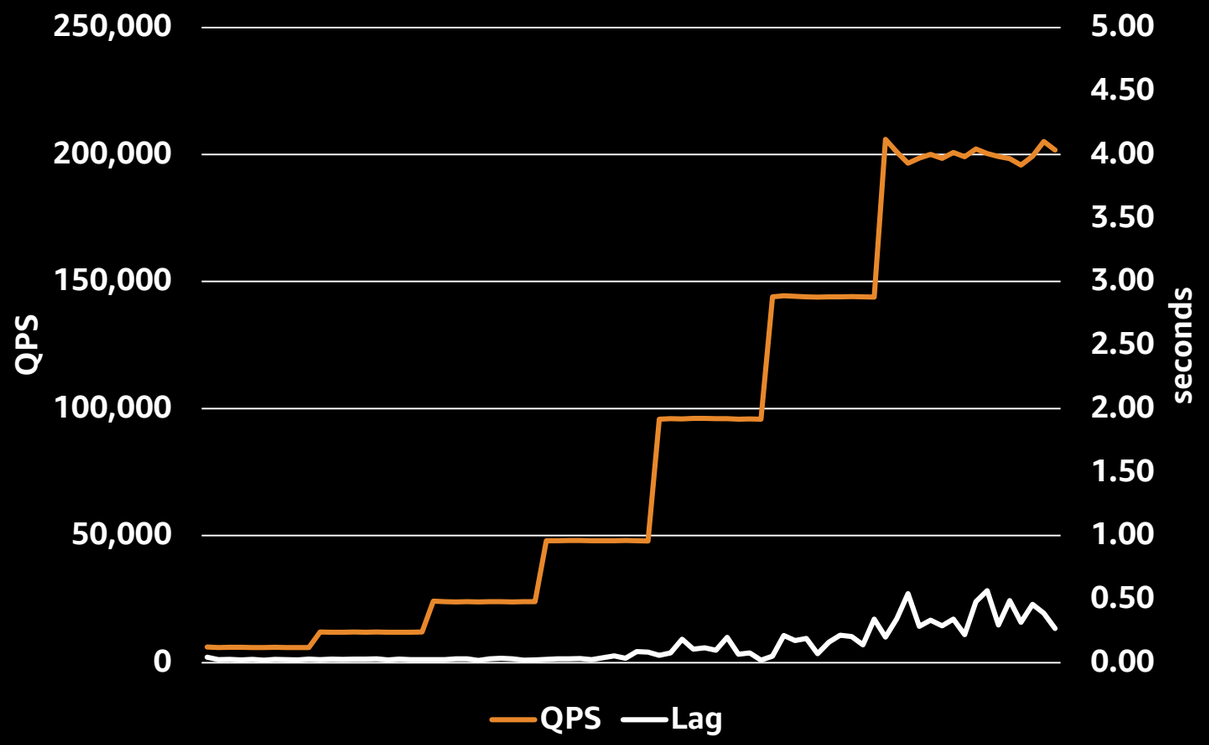


Aurora global database performance

MYSQL BINLOG VS. AURORA GLOBAL DATABASE



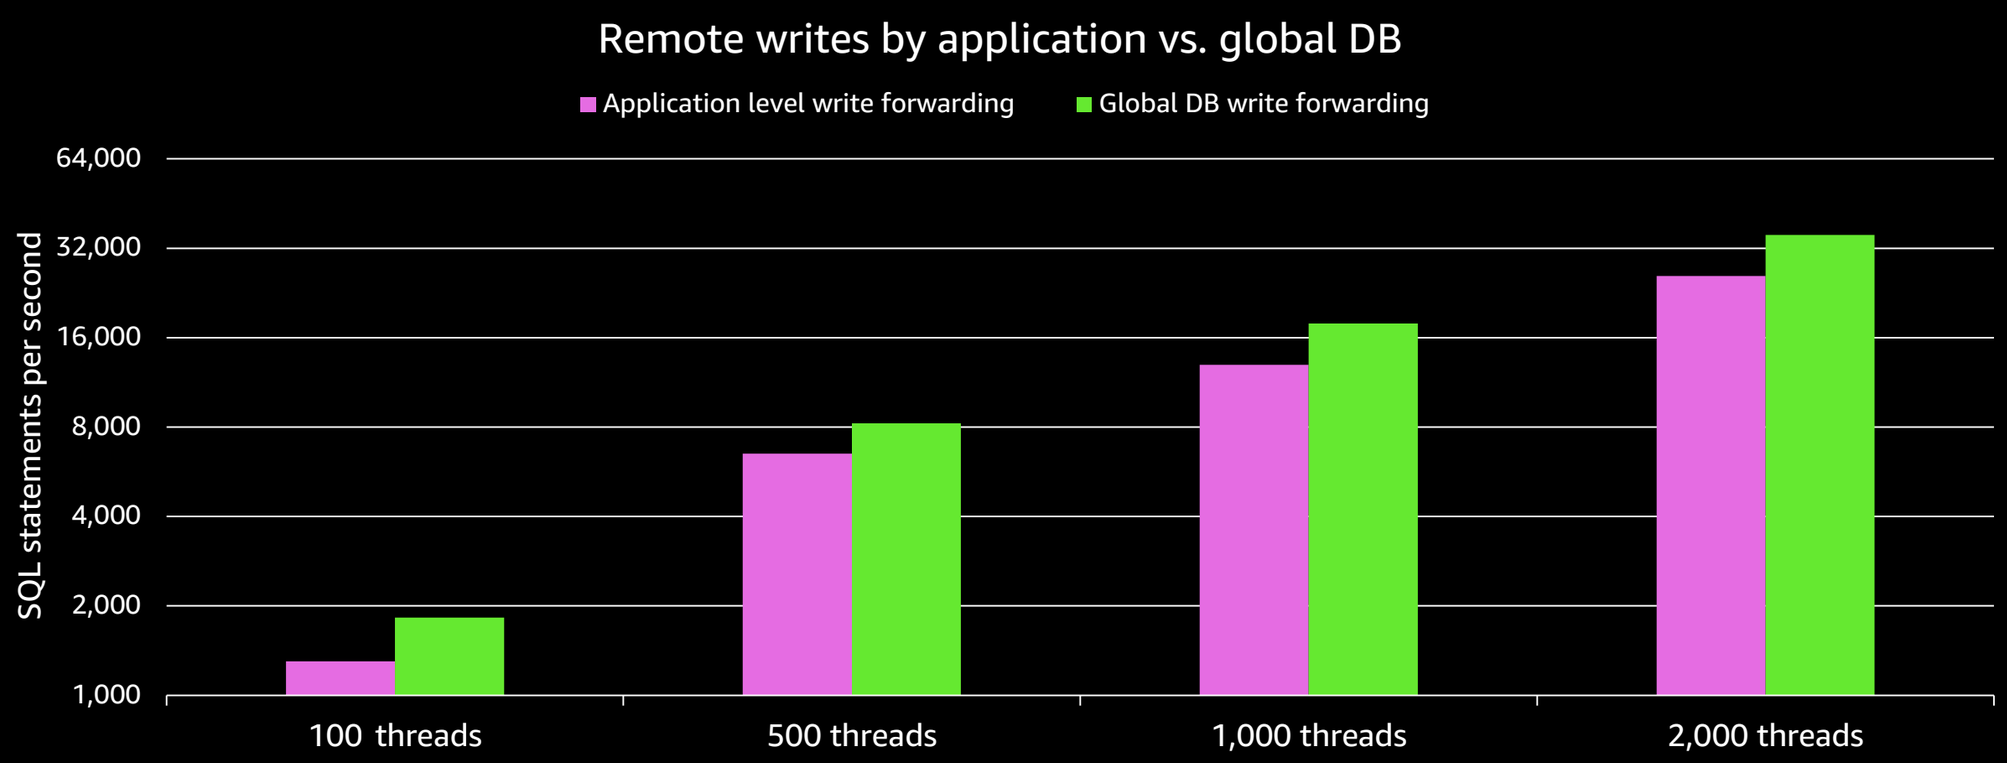
MySQL binlog replication



Aurora global database

Write forwarding from remote Region

GLOBAL DB READER INSTANCE TRANSPARENTLY FORWARDS WRITES TO PRIMARY REGION



Using sysbench OLTP workload executed on a reader instance in Aurora global database remote Region

Version currency

BEING CURRENT IS SIMPLIFIED

- Automated minor version updates
  - We have successfully performed minor version updates to Aurora MySQL fleet for both Aurora MySQL 5.6 and Aurora MySQL 5.7 for those database clusters that have opted for automated updates
  - The updates contained several critical security and availability fixes; we strongly encourage customers to opt for automated updates
  - We released 6 minor versions for Aurora MySQL 5.6 and 7 minor versions for Aurora MySQL 5.7 in 2020
- Major version upgrades
  - In place upgrade from 5.6 to 5.7

Success stories



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Schoology Learning

LEADING K-12 LEARNING MANAGEMENT SYSTEM

Challenge

Support a 400%+ increase in usage caused by COVID-19

Solution

Used Amazon Aurora with MySQL compatibility with 15 reader instances in primary region and 6 reader instances in secondary region using Aurora global database

Result

Schoology Learning team was able to scale to meet the demand of thousands of school districts moving to remote learning

10M+ students online

1.9M queries/second

30M rows read/second

8K writes/second





FINANCIAL MANAGEMENT SOFTWARE FOR CONSUMERS, SMALL BUSINESSES, AND ACCOUNTANTS

Challenge

Required a scalable, high-performance database to power TurboTax application across the U.S. during tax season

Solution

Used **Amazon Aurora** with MySQL compatibility in the US West (Oregon) and US East (N. Virginia) Regions, and enabled the global database feature

Result

The 2018 and 2019 tax seasons went smoothly, with Aurora achieving Intuit’s utilization and latency requirements throughout the U.S.

Related sessions

DAT301

Deep dive on Amazon Aurora with PostgreSQL compatibility

DAT401

Amazon Aurora storage demystified

DAT404

Deep dive on global database for Amazon Aurora



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!

Murali Brahmadesam

brahmade@amazon.com

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Please complete the session survey

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.