AWS re:Invent

FSI309

Relational databases: Performance, scale, and availability

Andrew Shortt
Solutions Architect
Amazon Web Services

Kenneth Jackson
Principal Solutions Architect
Amazon Web Services

AWS re:Invent

aws

Agenda

- The key role of relational databases in Financial Services

- Introducing Amazon Aurora

- Critical durability & availability

- Scale & performance under load

- Audience Q&A

We need databases that provide:

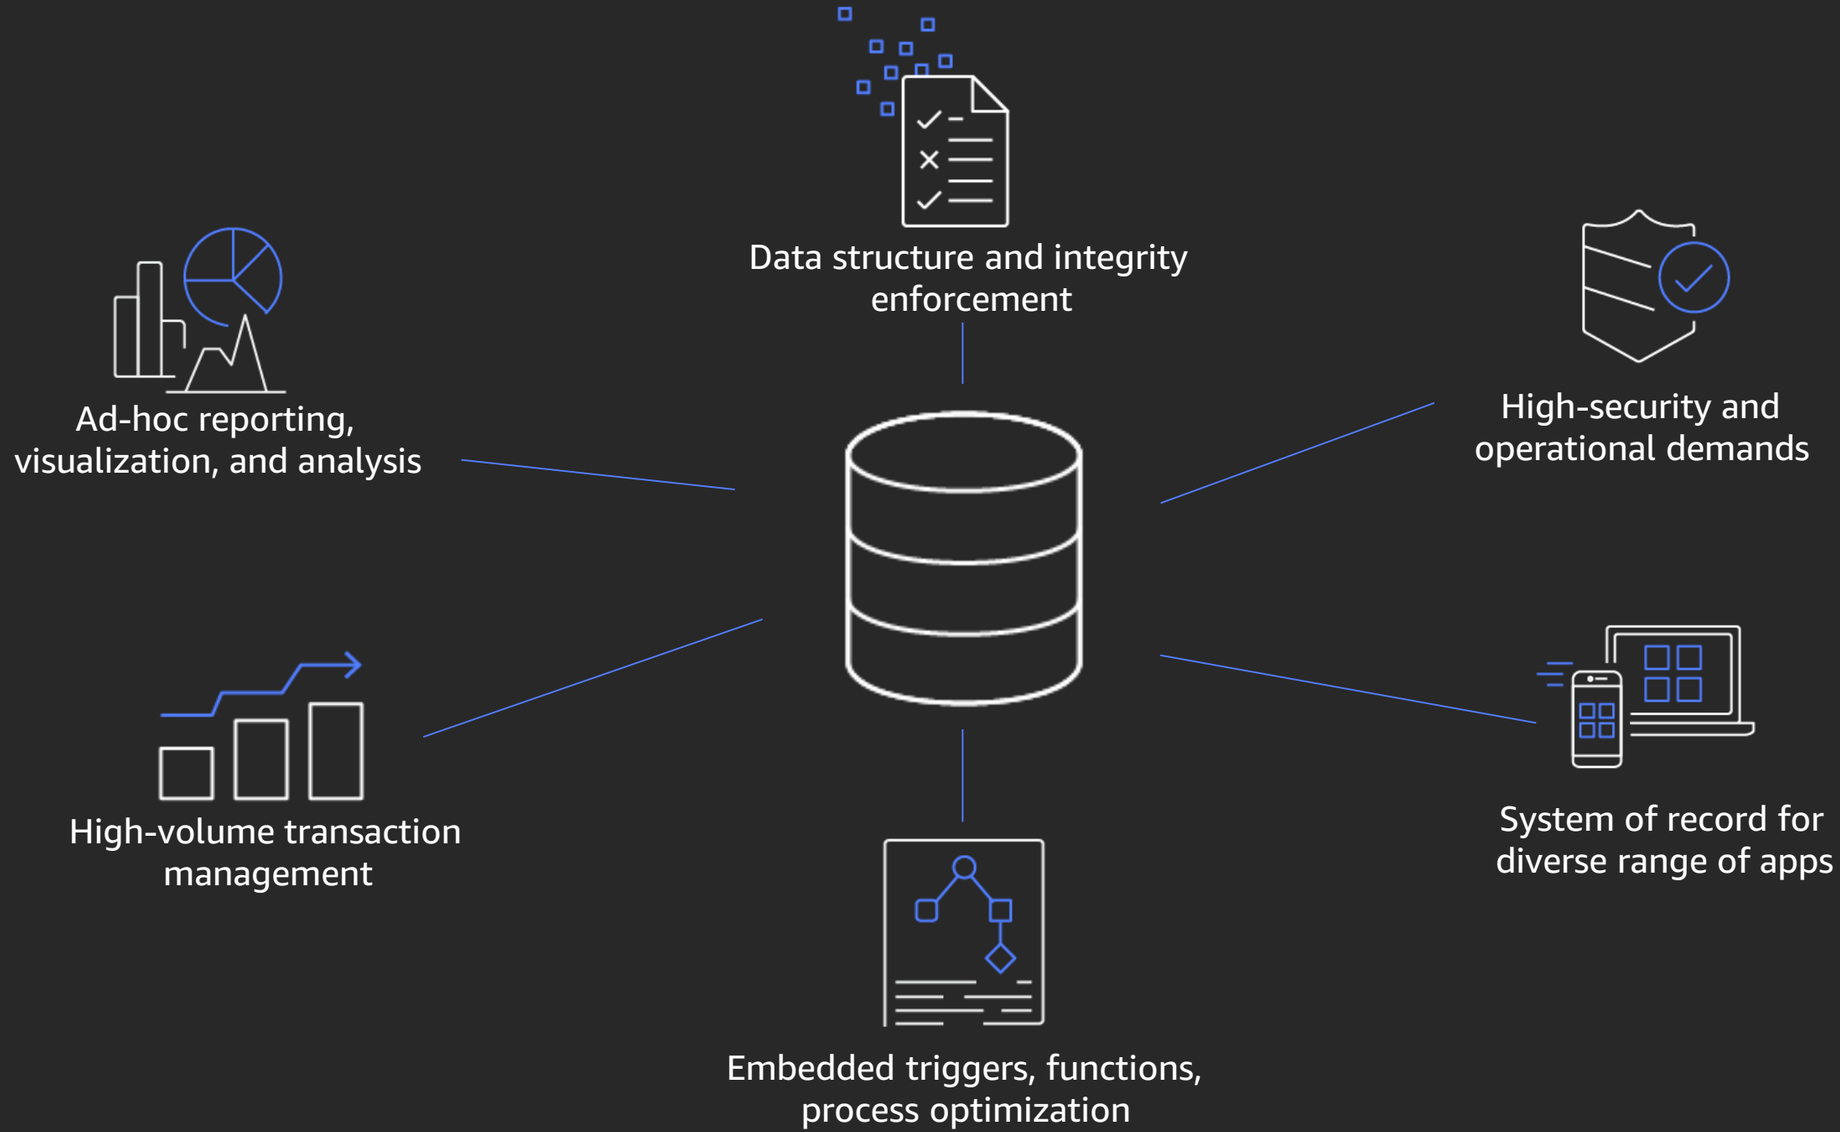

Relational databases

Fast

Capable of thousands of TPS

SQL

Easy to apply query language

ACID

Transaction integrity and consistency

Flexible

Wide variety of use-cases

Durable

Survive multiple physical failures

Amazon Aurora

Enterprise database at open-source pricing

Delivered as a **managed** service

Amazon Aurora

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

Amazon Aurora innovations

Reimagining databases for the cloud

1. Scale-out, distributed, multitenant design
2. Service-oriented architecture leveraging AWS services
3. Automate administrative tasks—fully managed service

Amazon Aurora in Financial Services

Intuit

"Intuit recently migrated their commerce platform to Amazon Aurora MySQL to support increasing global demand. All direct purchases of Intuit's software will go through the Intuit's Commerce Platform running on Aurora, with TurboTax already live to meet traffic demands during tax season. A large portion of our workload involves **low-latency**, read-only access to data. An example is pricing information, which is infrequently updated but needs to be readily available for reads from **coast to coast**. [Amazon] Aurora Global Database, with **sub-second global replication**, enables us to address this business requirement **without performance or latency constraints**. As a financial services company, we also care deeply about **business continuity even in the face of large-scale events**. [Amazon] Aurora Global Database allows us to maintain a strong **disaster recovery** posture by distributing data across AWS Regions with failover typically taking under a minute to complete."

E-commerce and Cloud Platform Engineering

Krishna V.

Durability & availability: traditional RDBMS

- Maintain uptime through hardware, data center, or regional level availability or durability events. Data survives physical destruction.
- Critical Tier 1/0 applications often require recovery times (RTO) of one hour or less and data recovery point (RPO) of less than 30 seconds

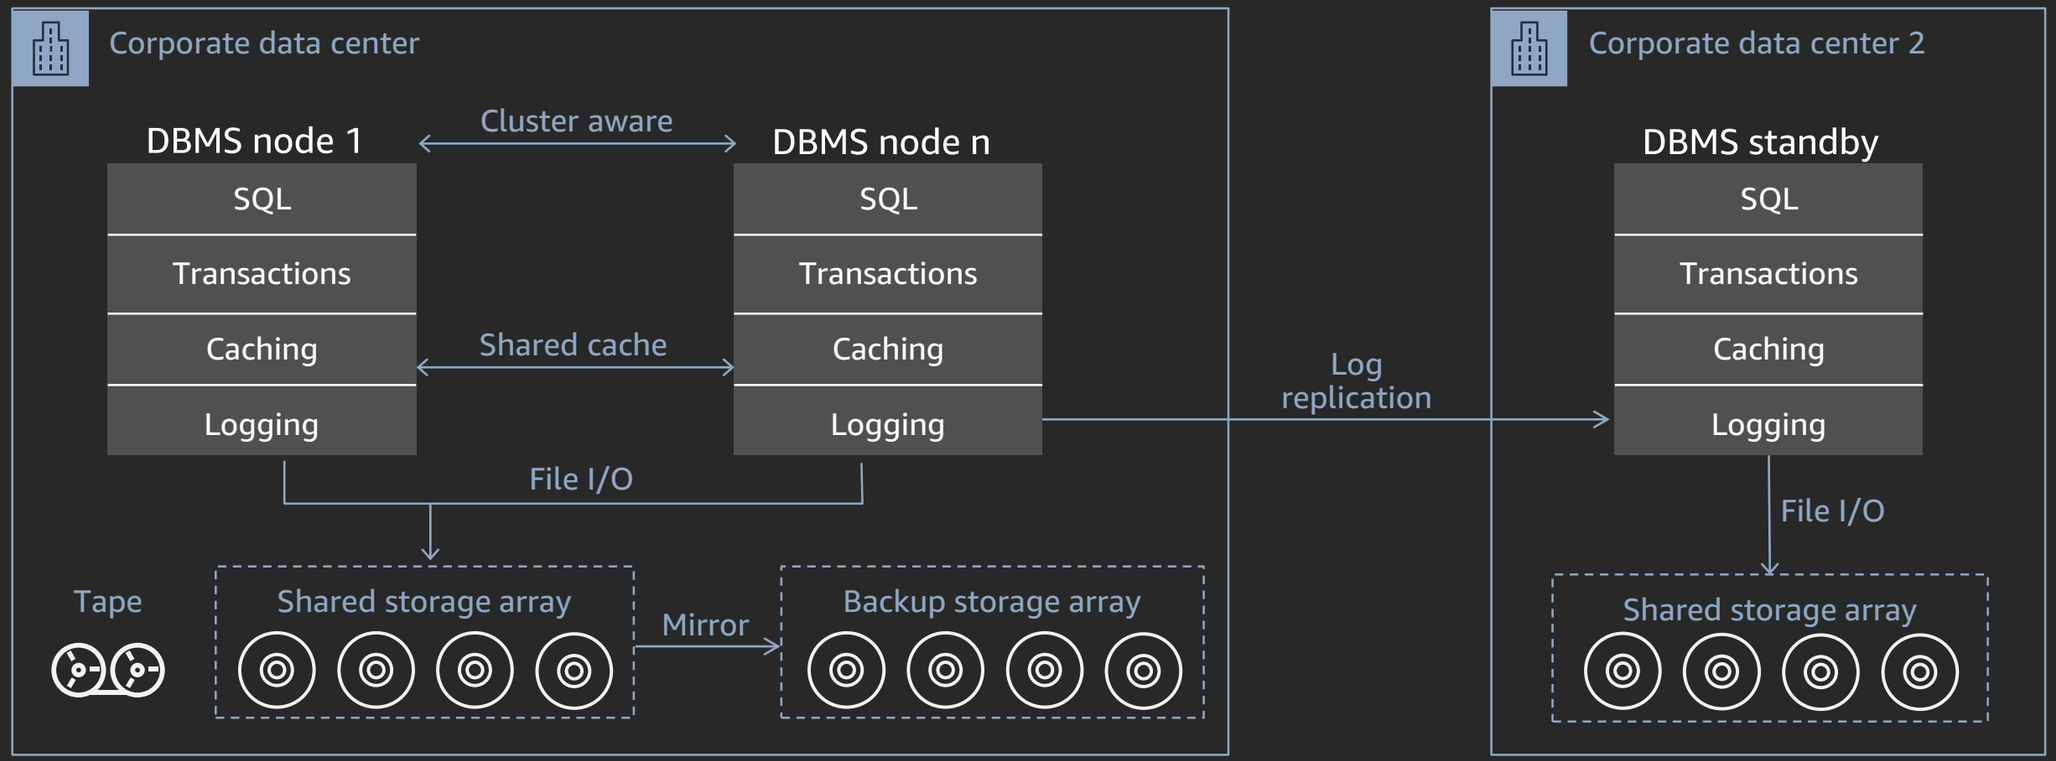

Durability & availability: Amazon Aurora

- Taking full advantage of the AWS global infrastructure, Aurora solves for Tier 1 availability requirements in unique ways

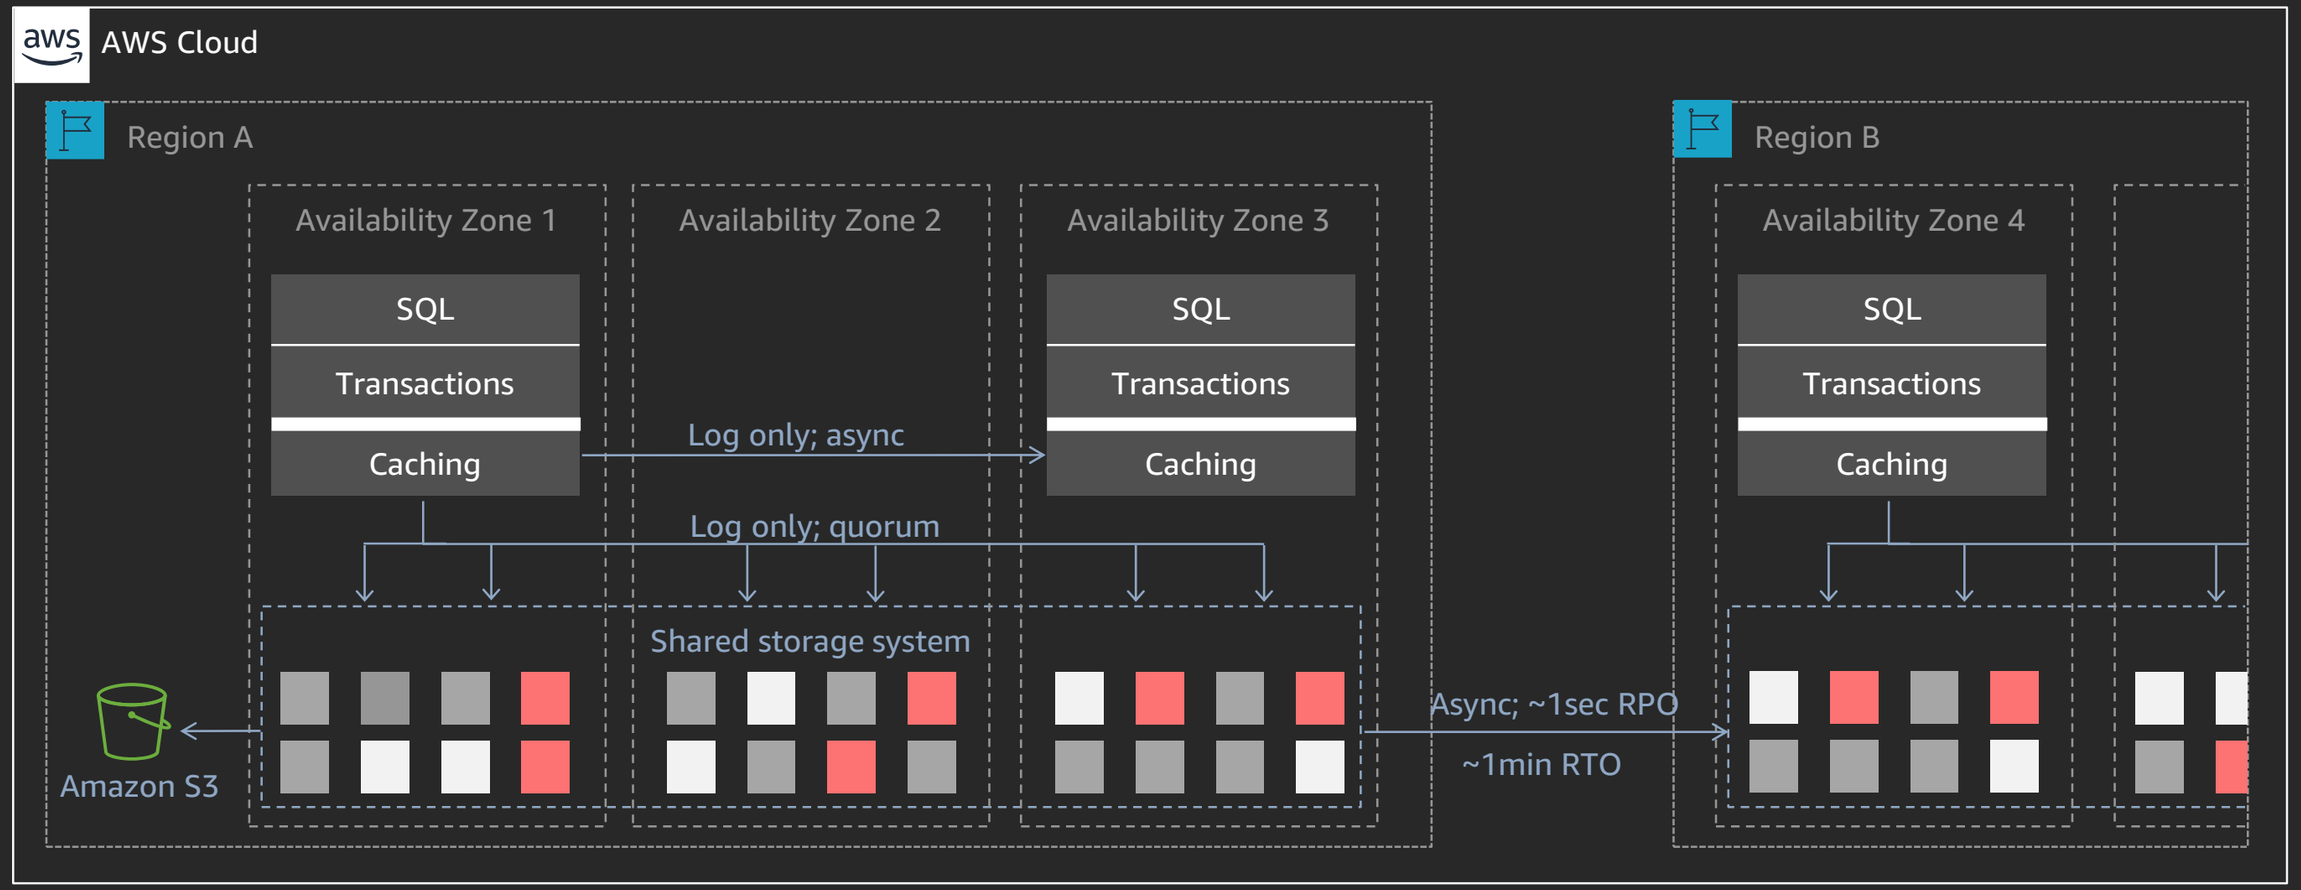

Performance & scale: traditional RDBMS

- Financial Services use-cases put extreme demand on OLTP systems, often experiencing peak periods of thousands of complex transactions per second
- As storage volumes grow, it puts pressure on both the database management engine and the underlying volume to either increase in size or split up the load

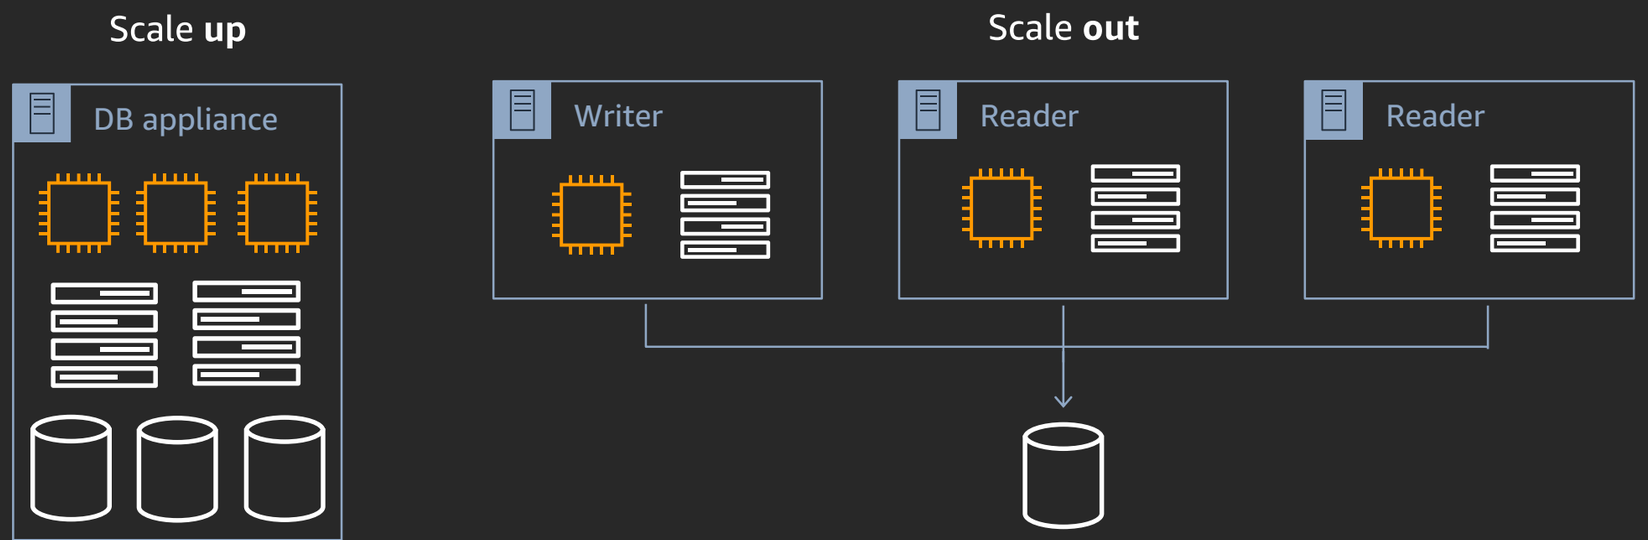

Performance & scale: Amazon Aurora

- Aurora storage layer handles the heavy lifting of storage volume growth and data block materialization
- Optimal cache efficiency and network communication is key for Aurora to perform as expected due to separated storage system

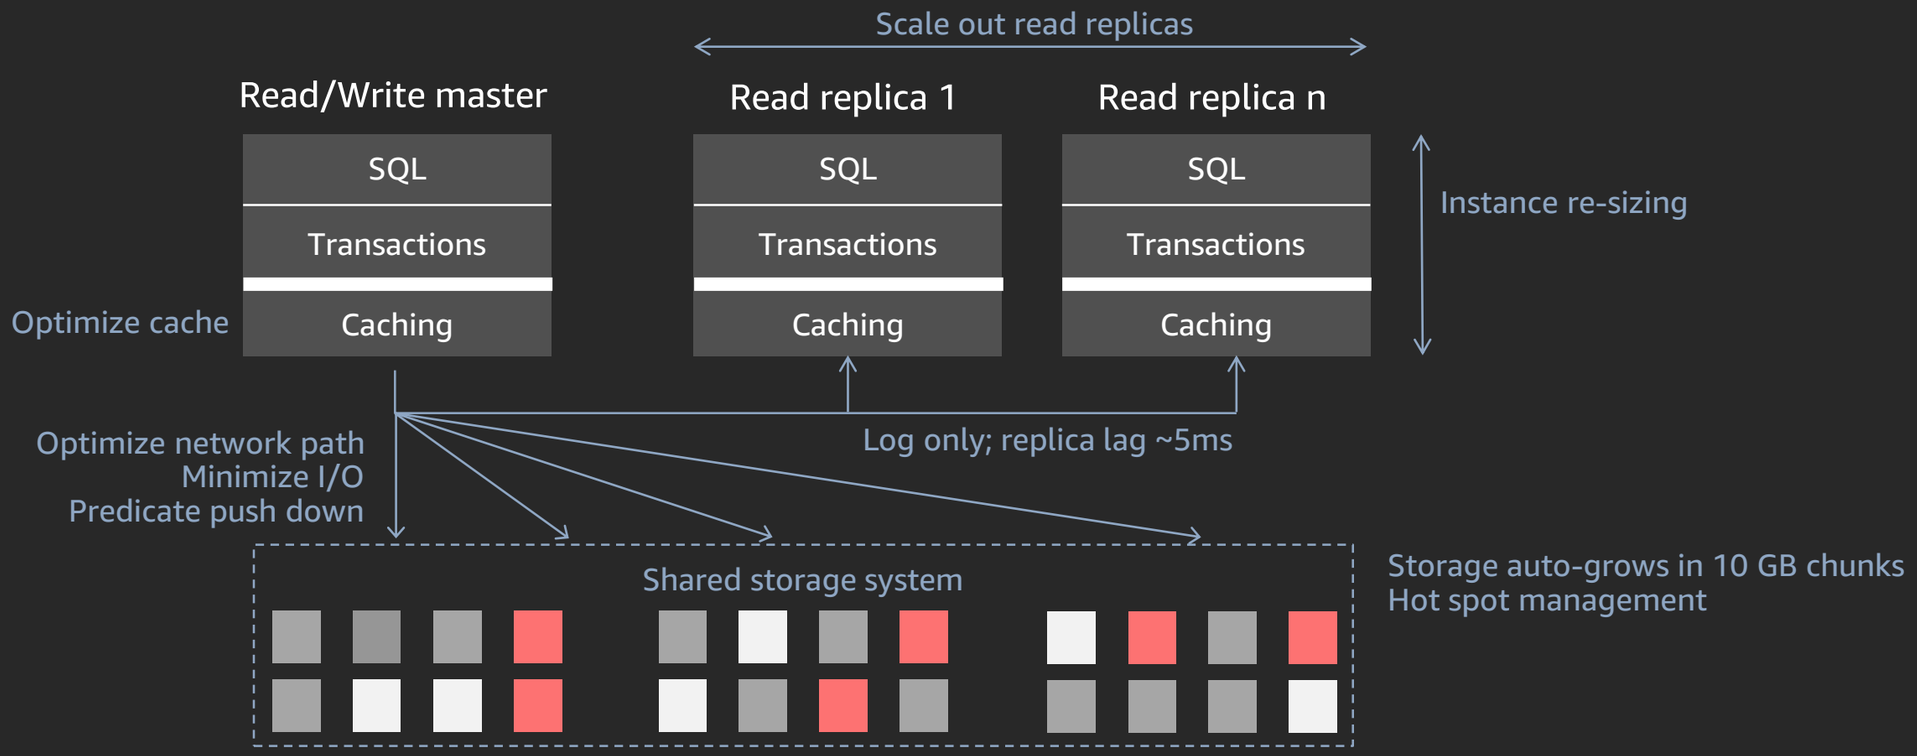

Performance variability under load

Aurora PostgreSQL is ~10x more consistent than PostgreSQL

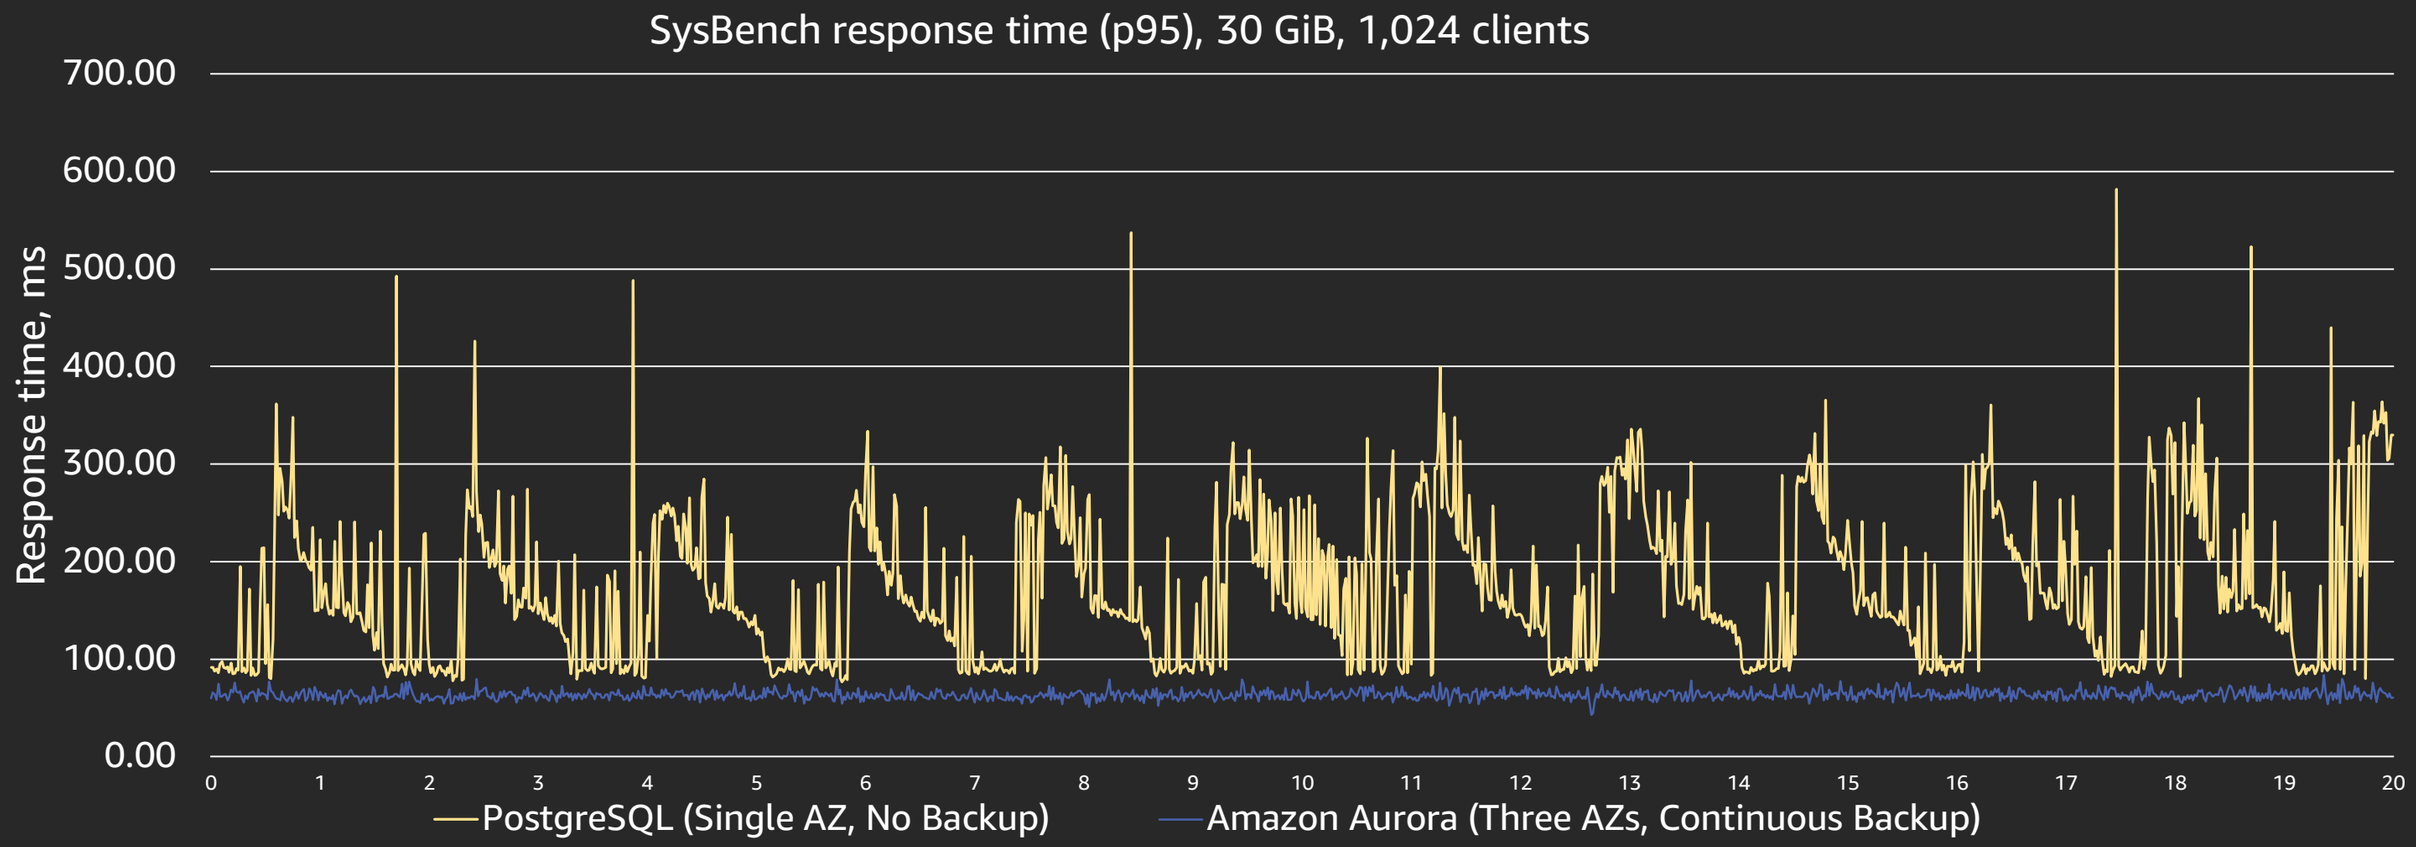

Amazon Aurora—latest features

- Multimaster write scaling
- Custom endpoints
- **Aurora serverless**
- Database backtrack
- Instant clone
- Data API
- Performance insights

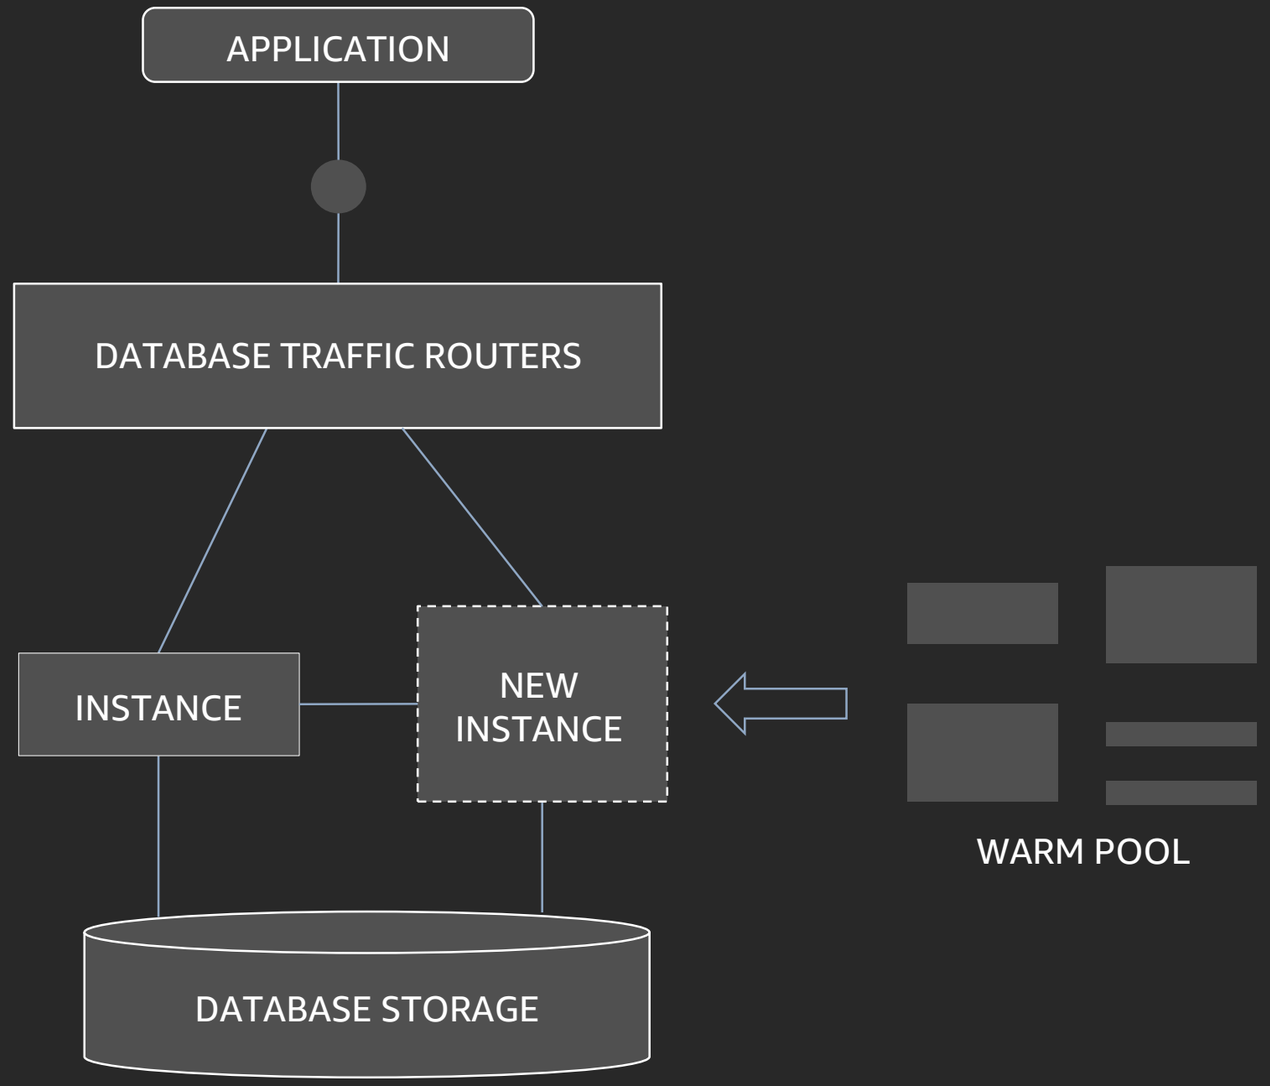

Q&A

Amazon Aurora high-level design

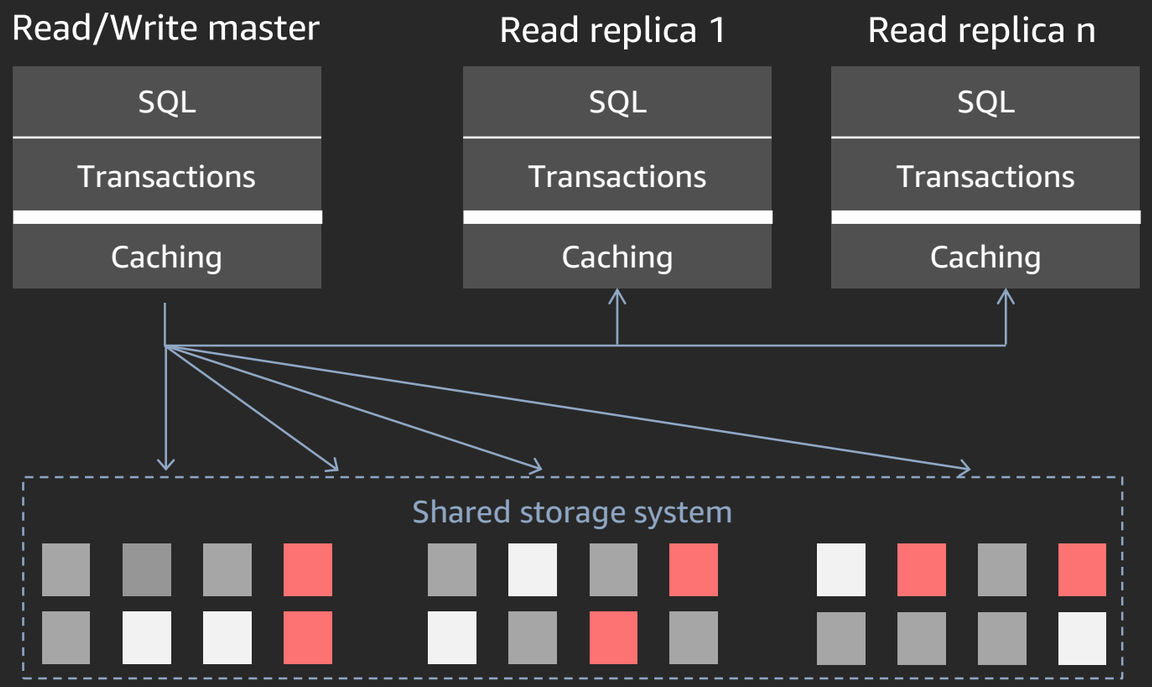

Storage node commit process

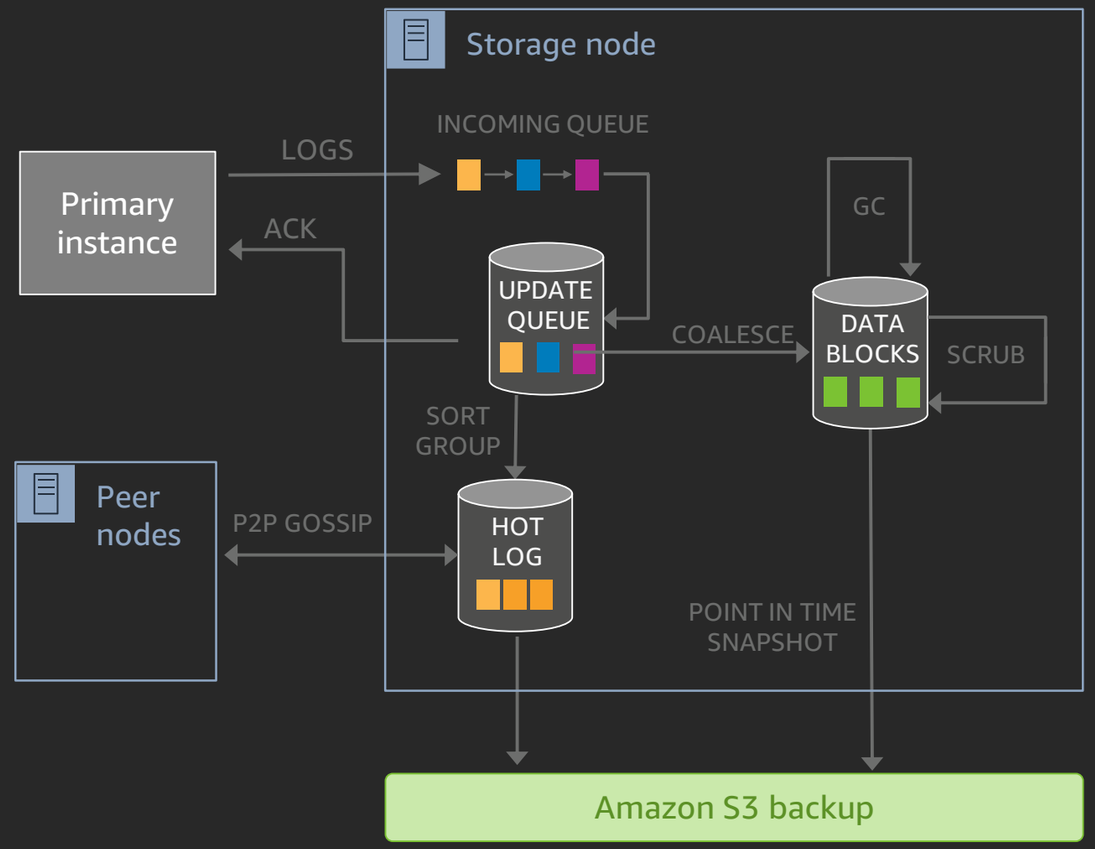

Thank you!

Andrew Shortt
ashortt@amazon.com

Kenneth Jackson
jaknn@amazon.com

AWS re:Invent

aws

Please complete the session survey in the mobile app.