
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

PEX314-R

Building global applications with Amazon Aurora

Arabinda Pani (He/Him)

Principal Database Partner Solutions Architect
AWS

Rajesh Matkar (He/Him)

Principal Database Partner Solutions Architect
AWS

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

- Global application database key requirements
- Amazon Aurora introduction
- Amazon Aurora Global Database
- Amazon Aurora Serverless v2
- Global application database architecture using Amazon Aurora
- Q&A

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Global application database key requirements

- ❑ High availability
- ❑ Lower RTO and RPO SLAs
- ❑ Faster local reads
- ❑ Low replication lag
- ❑ Low performance impact on primary database for replication
- ❑ Low operational overhead for replication management
- ❑ Performance and scalability
- ❑ Cost optimized solution

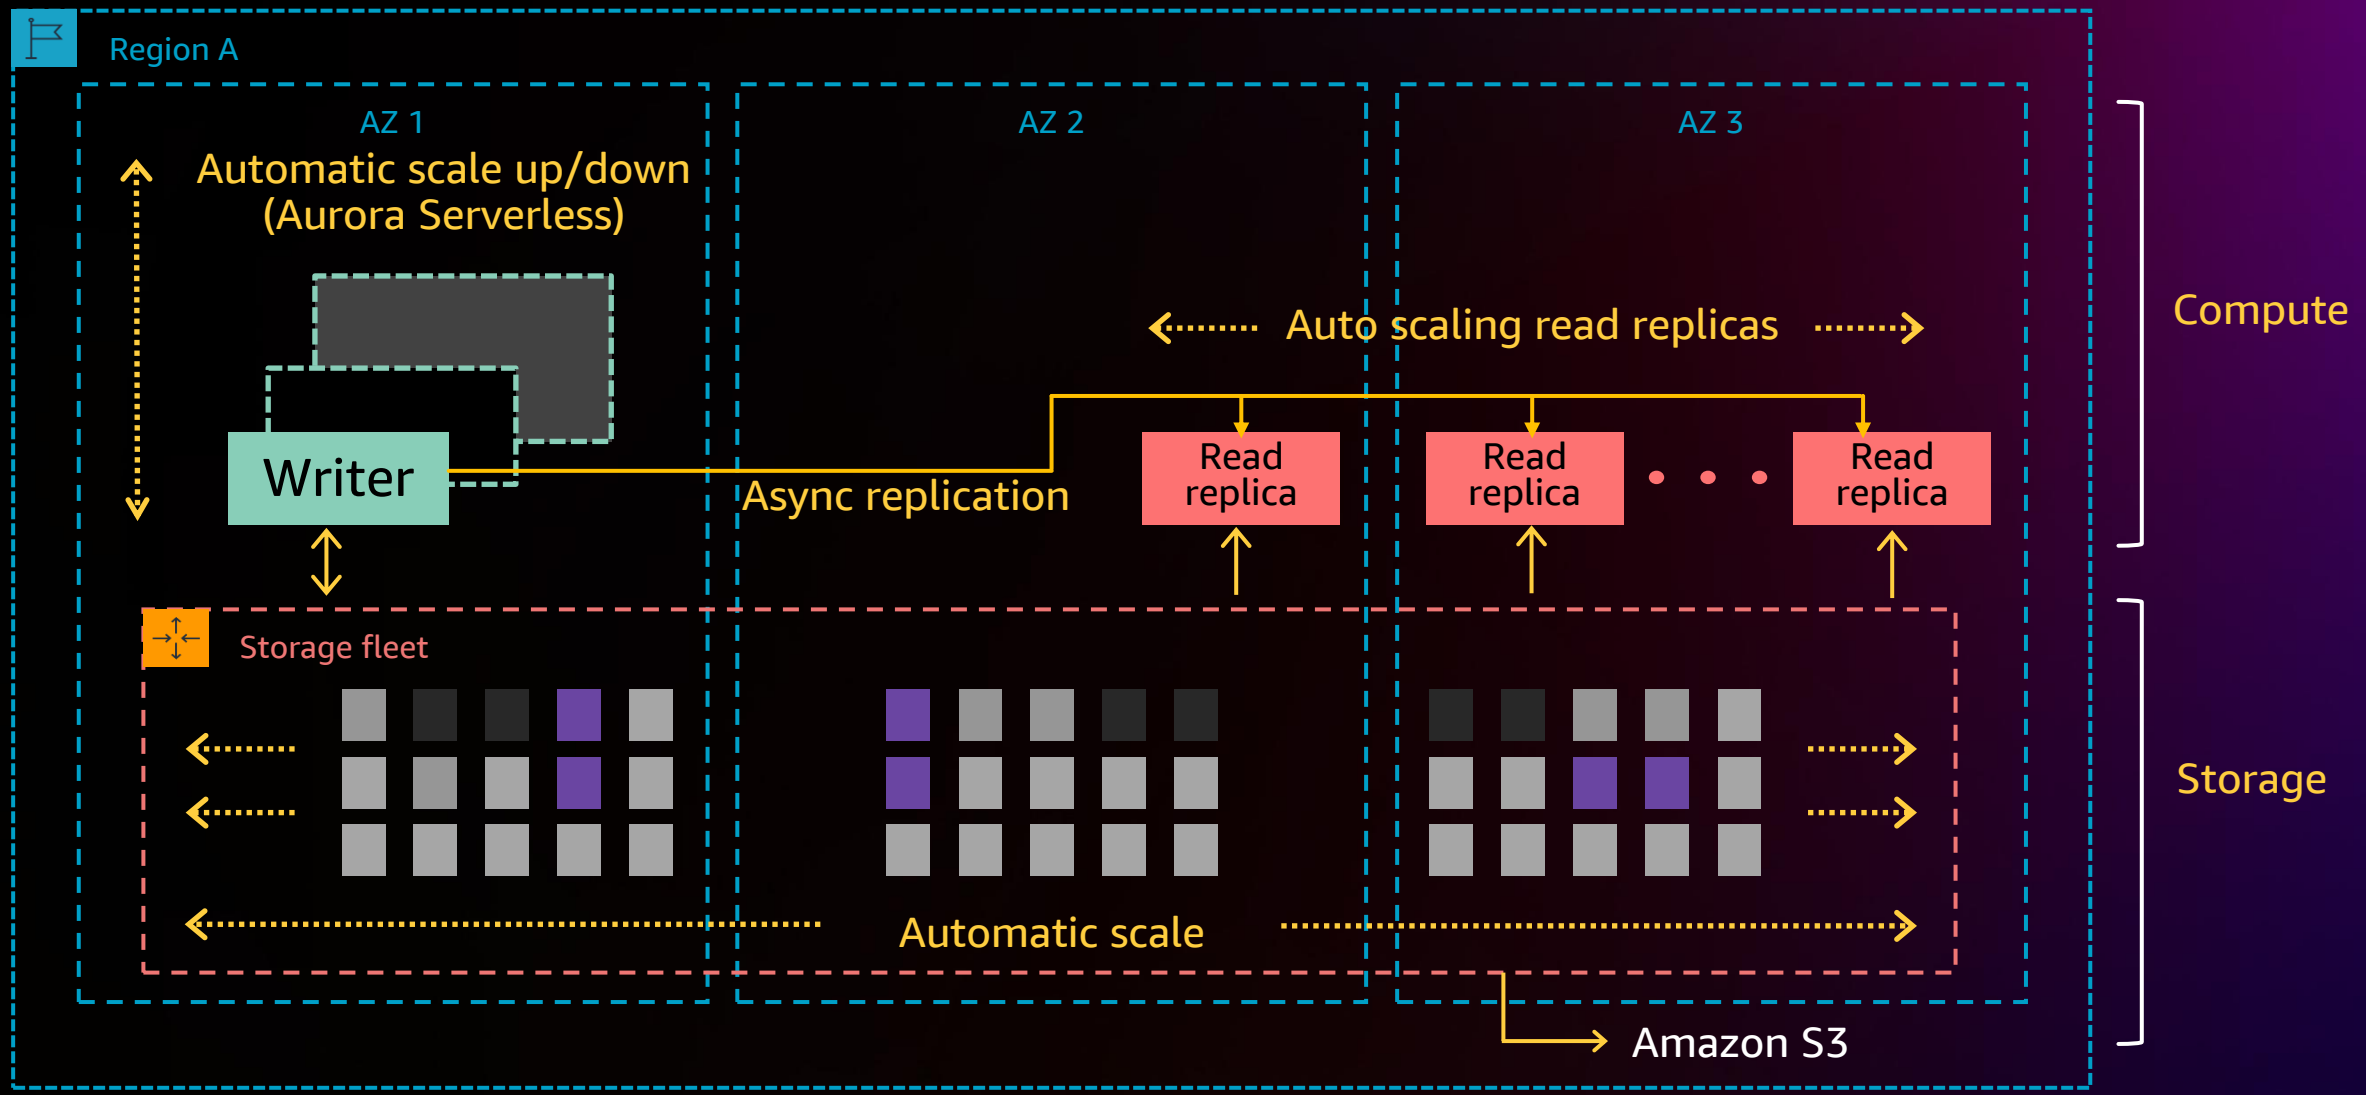


E-commerce | Media streaming | Social media | Online gaming | Shared economy

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

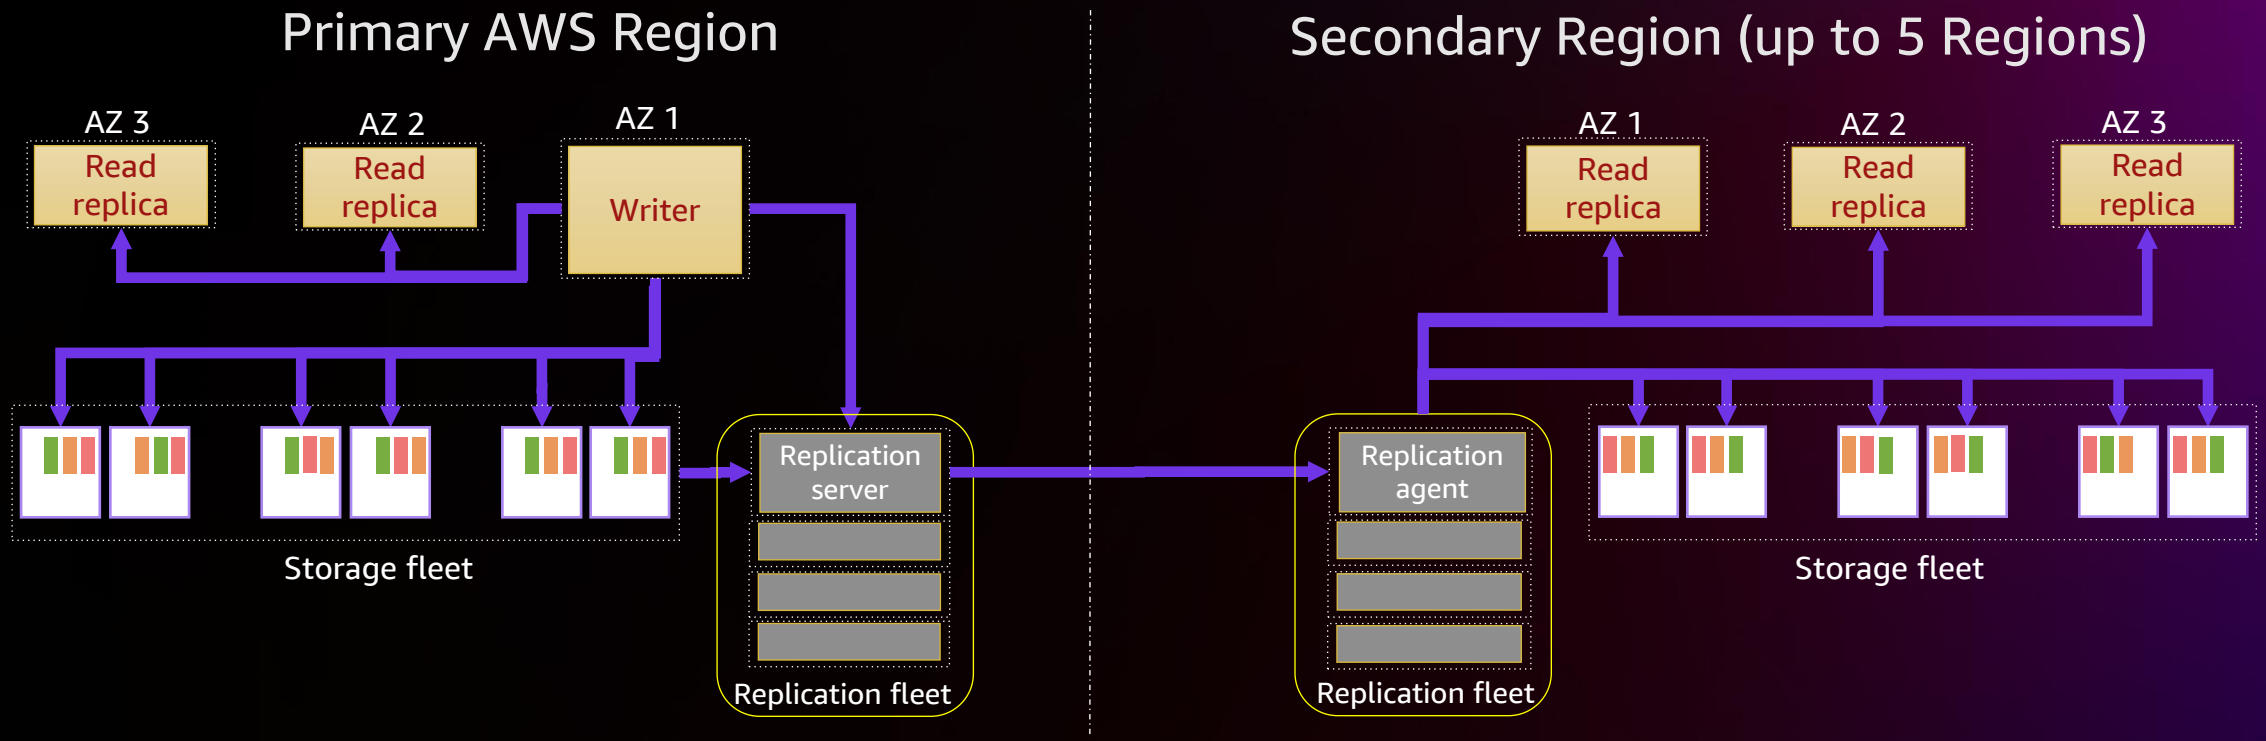
Amazon Aurora architecture



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Aurora Global Database

FASTER DISASTER RECOVERY AND ENHANCED DATA LOCALITY



Global replication: Up to 5 secondary Regions

Low replica lag: Typically <1 sec cross-Region lag

Fast recovery: Typically <1 min downtime after Region unavailability

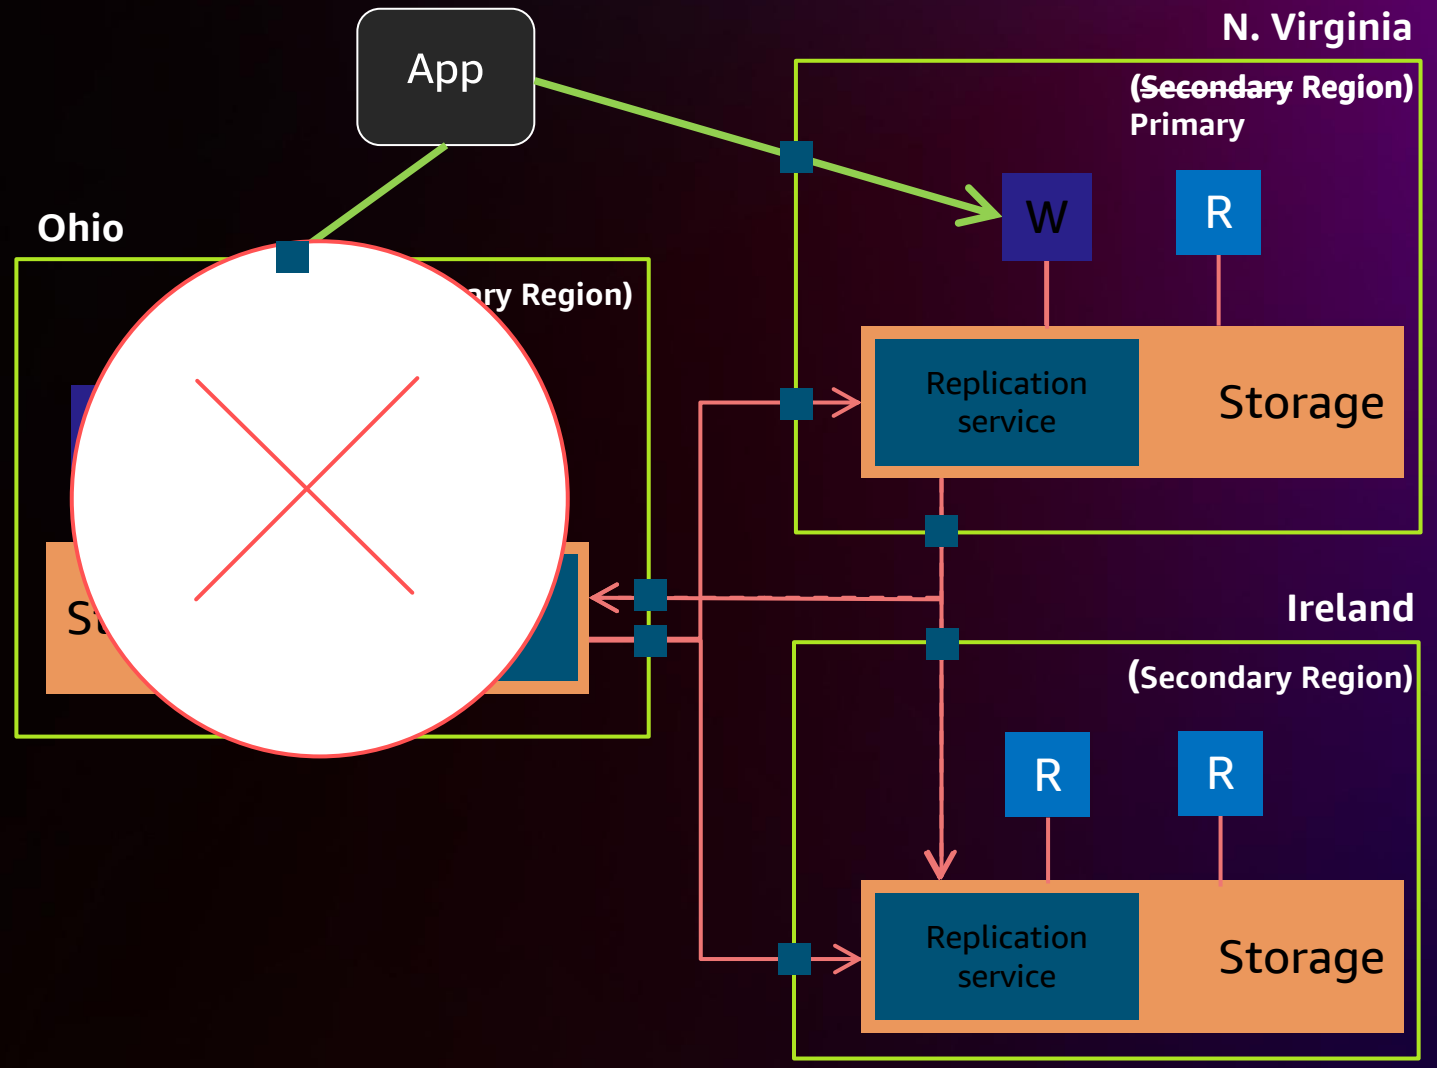
High throughput: Up to 200K writes/sec – negligible performance impact

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Unplanned failover (“Region disaster” scenario)

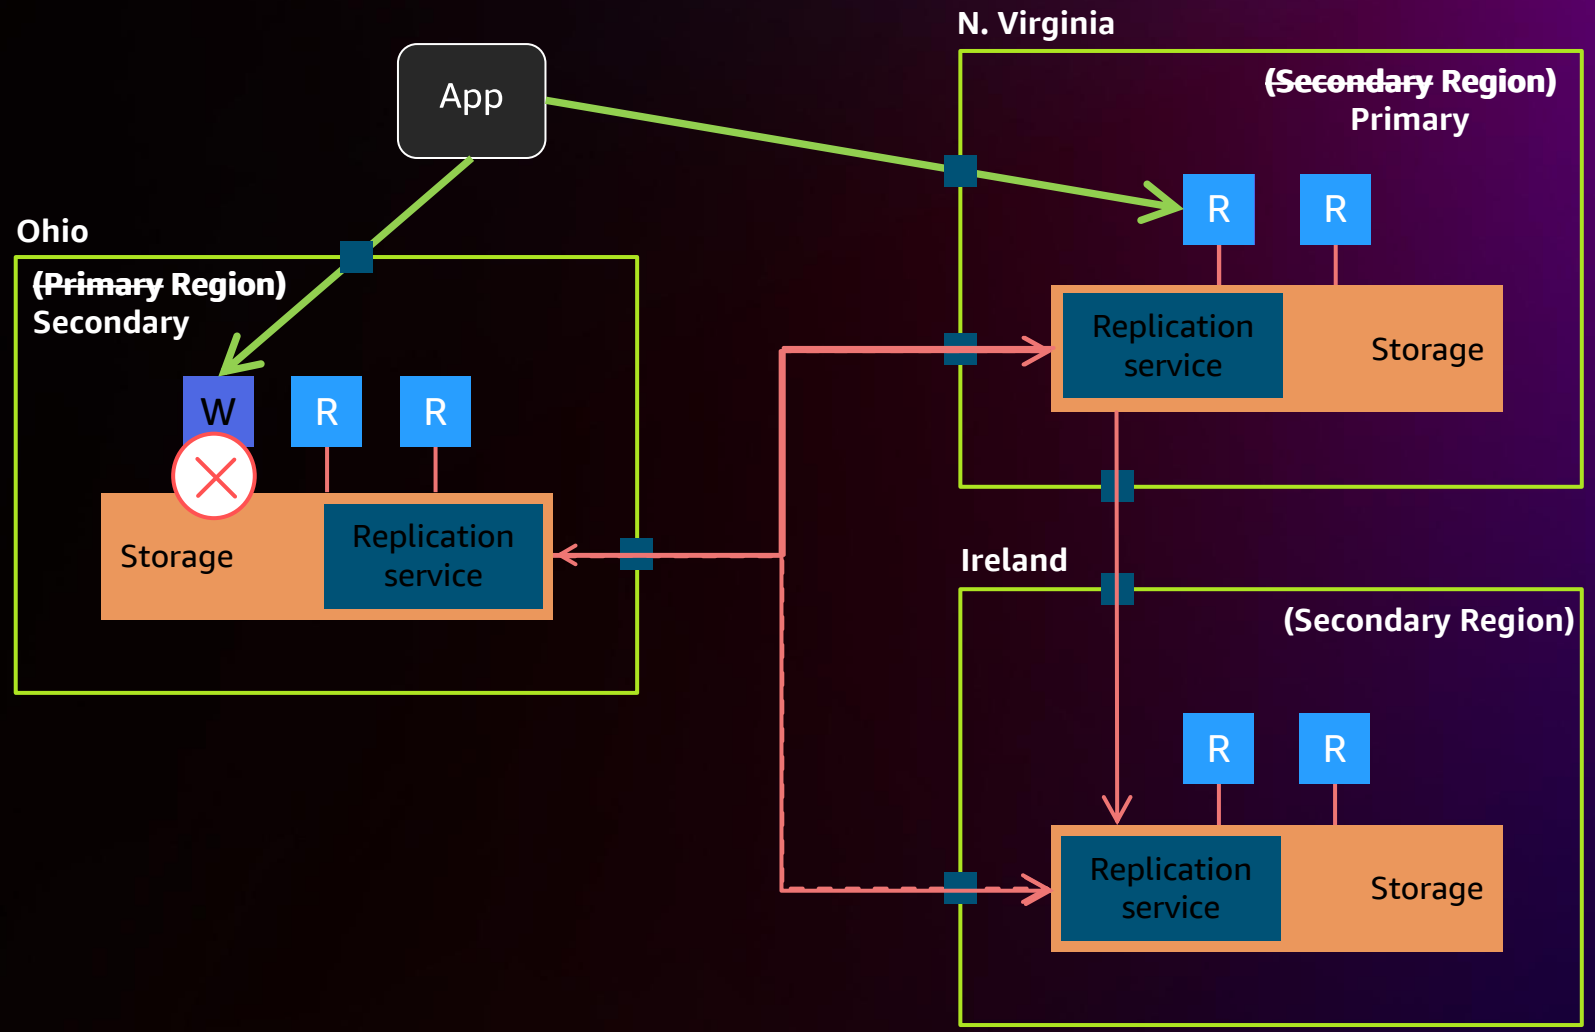
- An outage dictates decision to failover
- Identify Region with least replication lag
- Detach and promote chosen Region to be a “standalone” cluster
- Example: N. Virginia becomes primary
- Once up, application can point to the new endpoint
- Delete stranded clusters in other Regions
- Rebuild secondary Regions when needed/available; fail-back using managed planned failover



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Managed planned failover

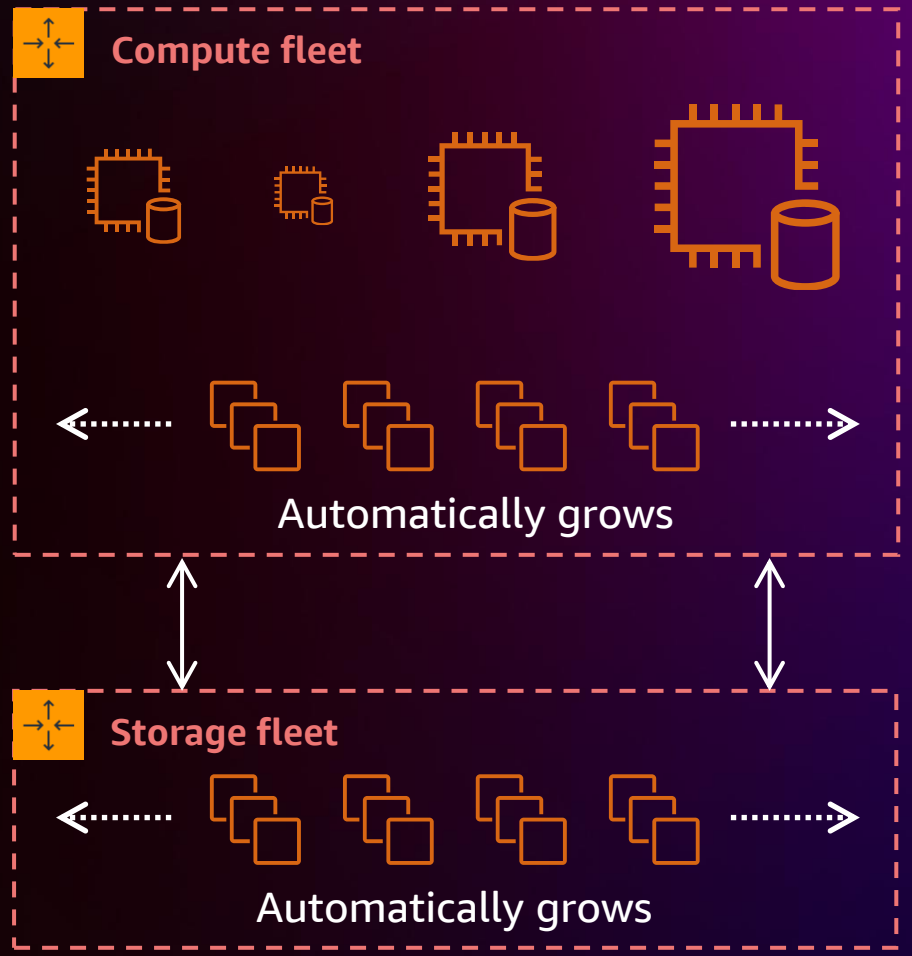
- → We click to promote N. Virginia
- → Ohio writes are stopped
- → Secondary Regions catch up
- → N. Virginia becomes the primary
- → Writes now available in N. Virginia
- → Point application to the new endpoint



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Aurora Serverless v2

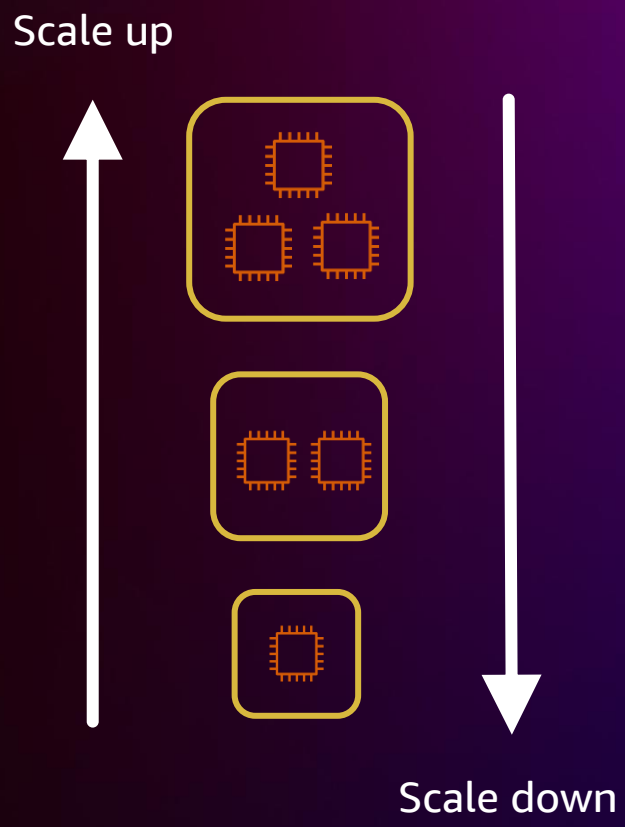
- On-demand and auto scaling configuration
- Scales in place in under a second by adding more CPU and memory resources while preserving state (for example: buffer pool, connections)
- No impact due to scaling even when running hundreds of thousands of transactions
- Compute fleet continuously monitored and scaled horizontally for heat management
- Up to 15x faster scale downs
- Simple pay-per-use pricing per second



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

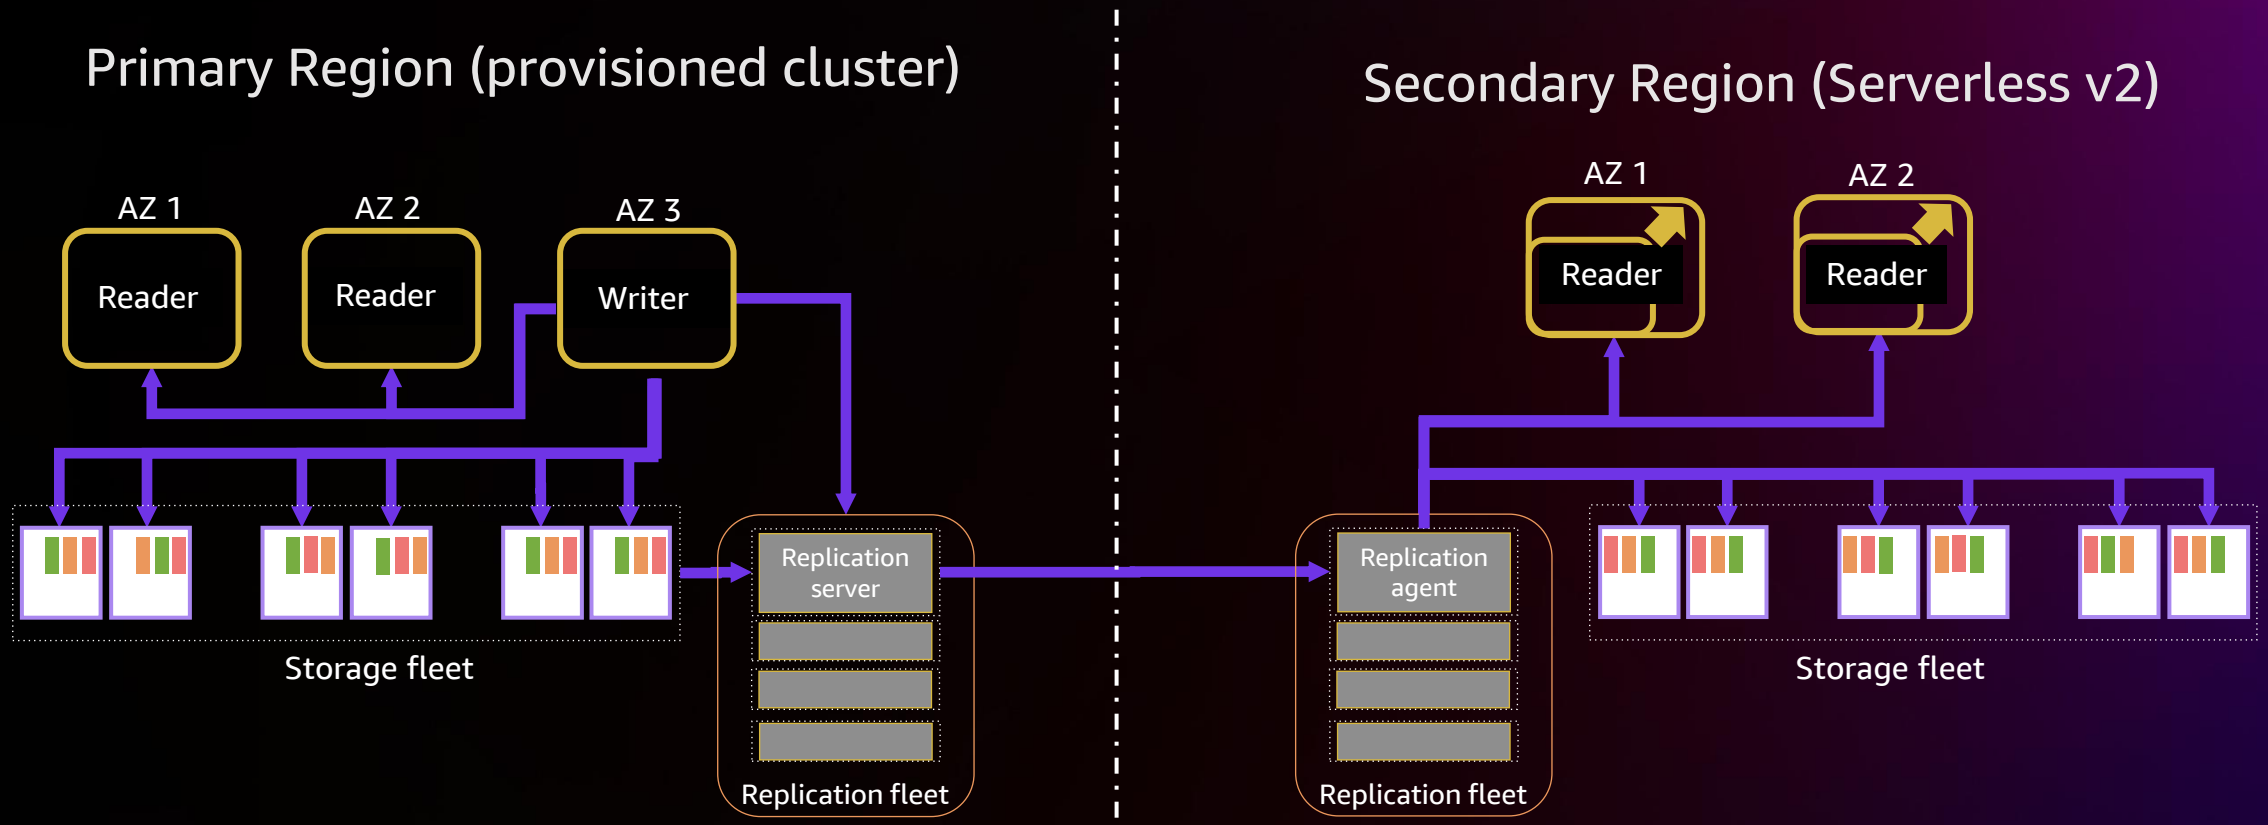
Fine-grained capacity adjustments

- Aurora Serverless capacity is measured in Aurora Capacity Unit (ACU)
- 1 ACU comes with 2 GiB of memory
- CPU and networking similar to what is available in provisioned Aurora instances
- Starting capacity as small as 0.5 ACU (1 GiB)
- Fine-grained scaling with increments as small as 0.5 ACU



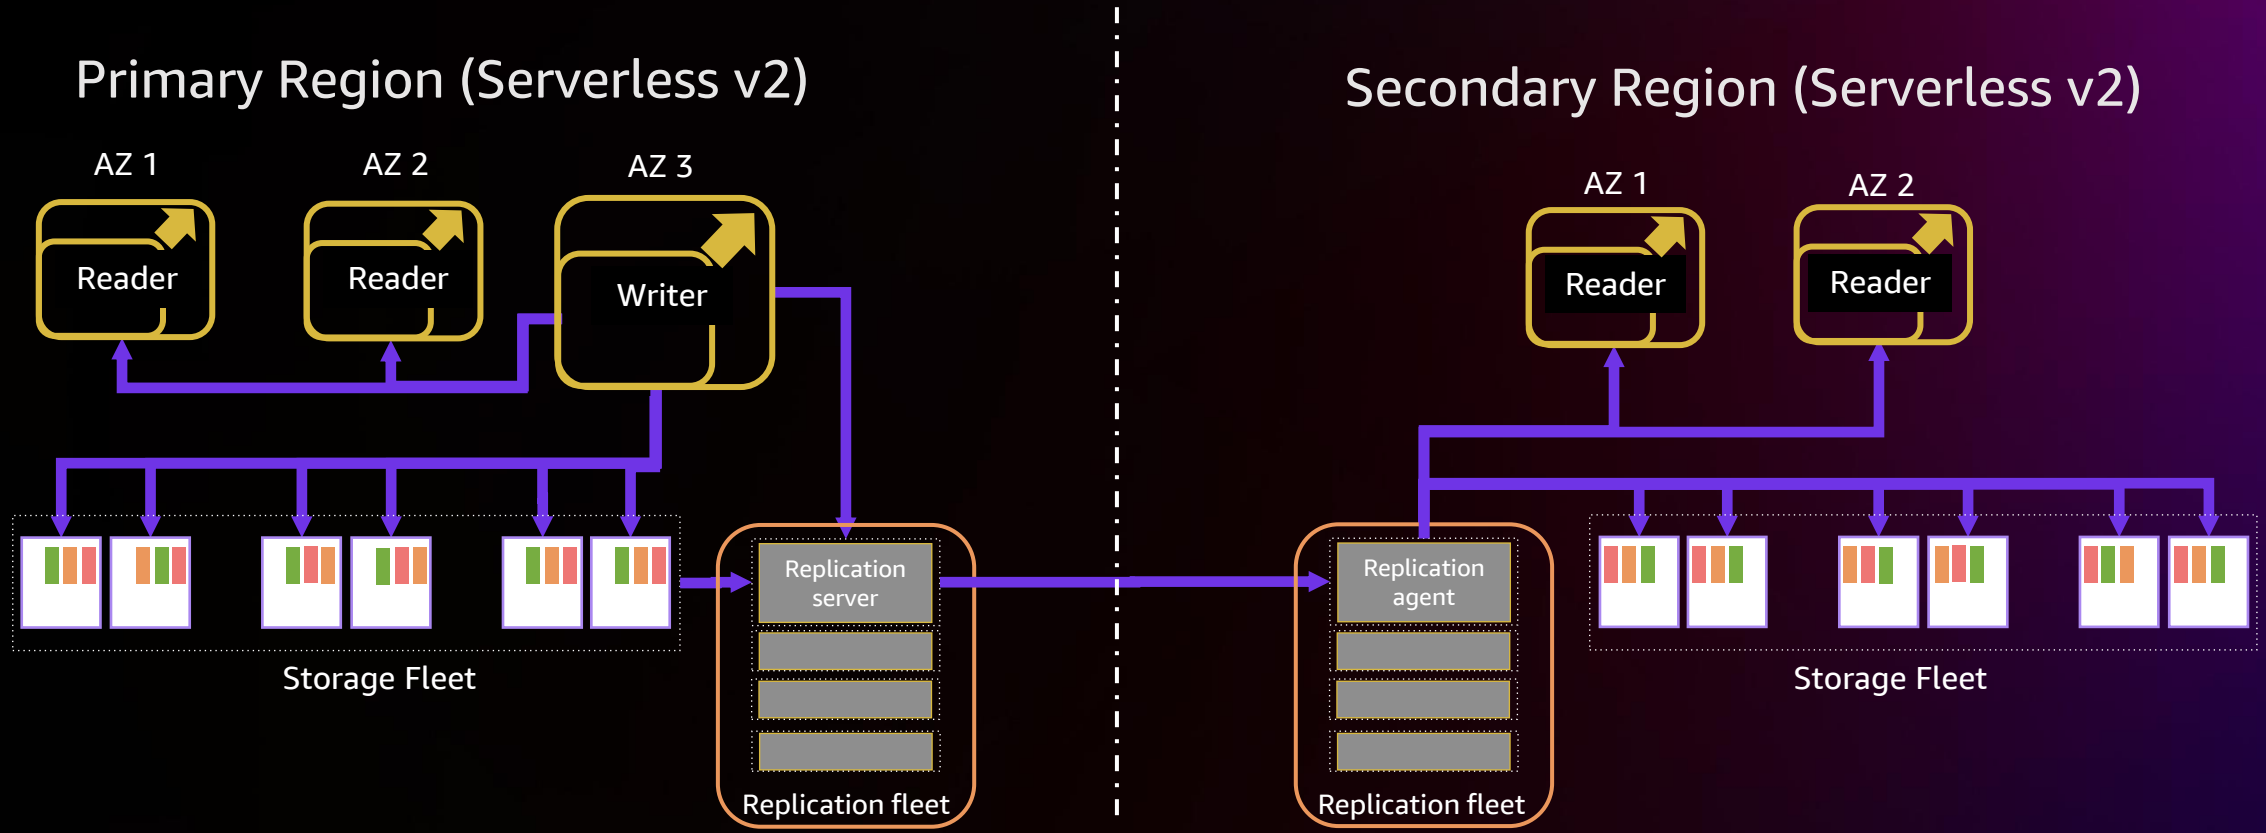
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Aurora Global Database & Aurora Serverless v2



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Aurora Global Database & Aurora Serverless v2



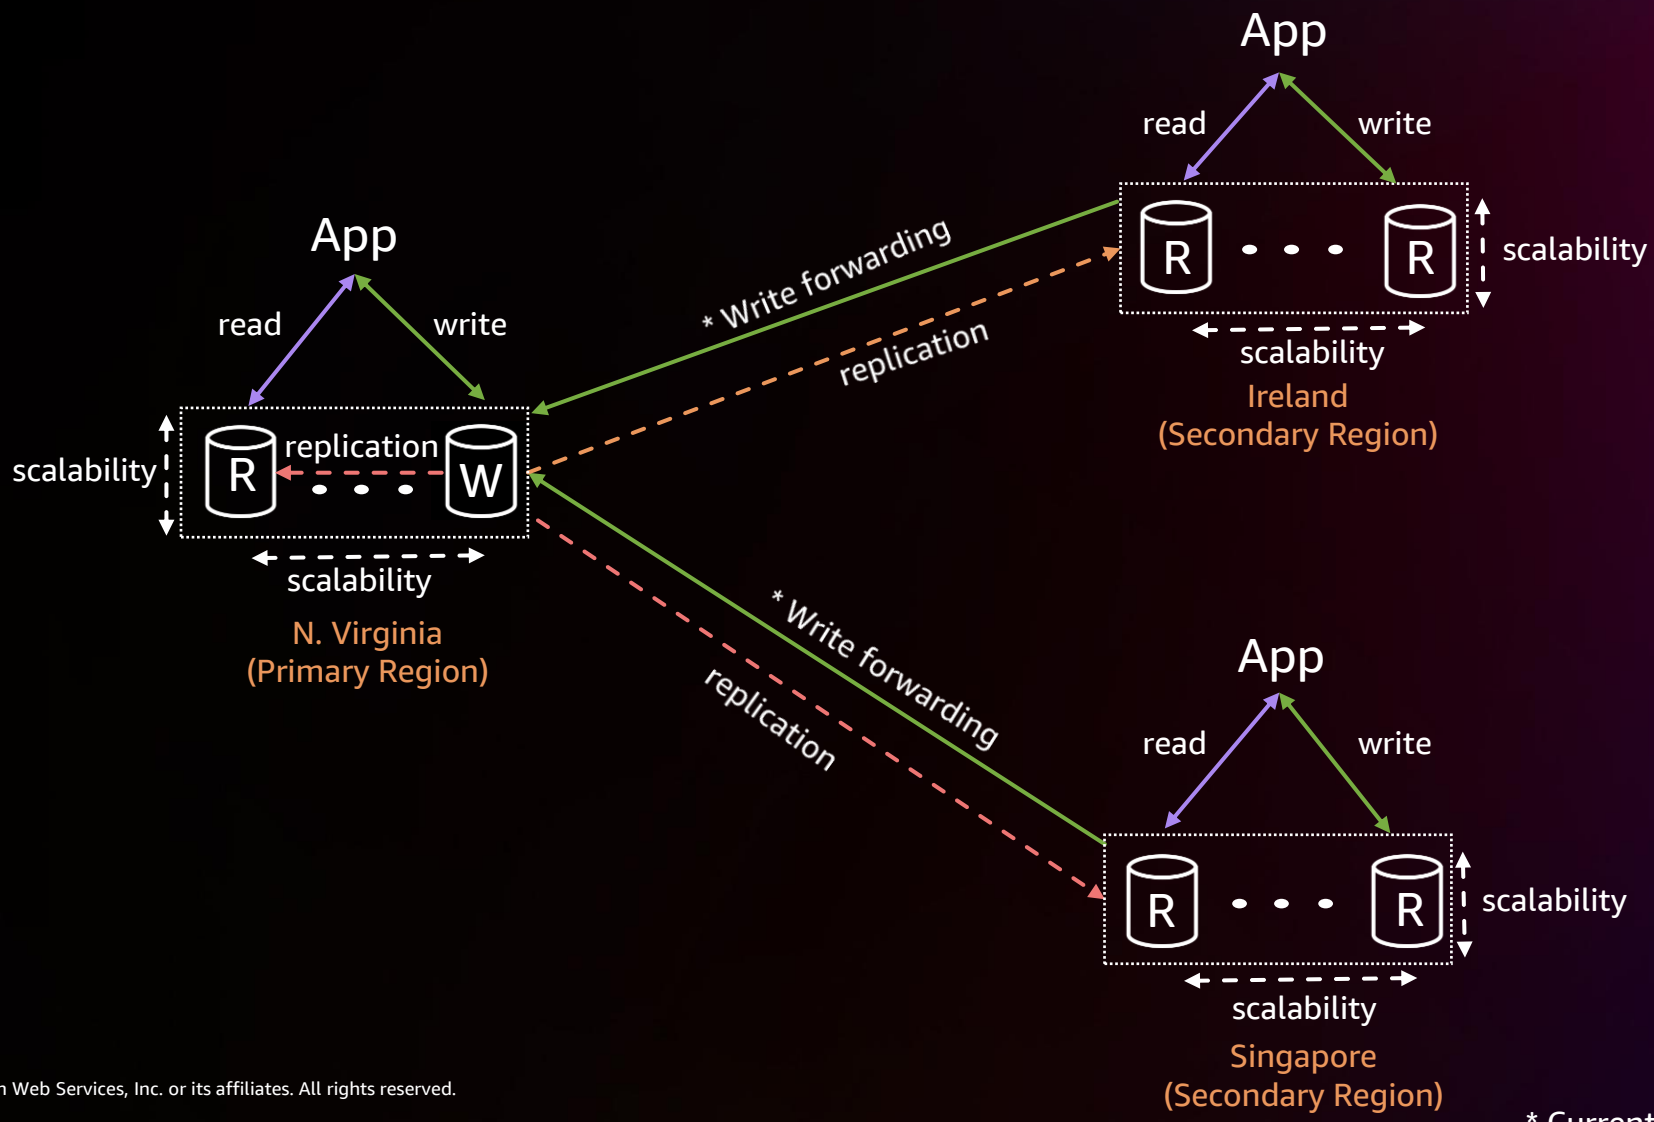
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Aurora for global applications

Key requirements	Amazon Aurora
High availability	Up to 15 replicas and <30 seconds in-Region failover
Lower RTO and RPO SLAs	<1 min RTO and <1 sec RPO with Aurora Global database
Performance and scalability	Built-in performance improvements and dynamic scalability of compute and storage
Low replication lag	<100 milliseconds in-Region and <1 sec cross-Region using Aurora Global database
Low performance impact on primary database for replication	Using physical storage based replication
Low operational overhead for replication management	Offered as part of Aurora managed service
Cost optimized solution	Pay as you use pricing model

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Global application database architecture using Amazon Aurora



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

* Currently available for Aurora MySQL only

Related sessions

DAT310	Tuesday, November 29 2:45 PM - 4:45 PM	Workshop	Develop event-driven applications using Amazon Aurora Serverless v2
FSI301	Tuesday, November 29 4:15 PM - 5:15 PM	Chalk Talk	Optimizing database resilience in a multi-Region universe
DAT317	Wednesday, November 30 2:30 PM - 3:30 PM	Chalk Talk	Deploy AWS global databases for a global user base
DAT326	Wednesday, November 30 2:30 PM - 3:30 PM	Breakout	Deep dive into Amazon Aurora and its innovations
DAT001	Wednesday, November 30 4:00 PM - 4:20 PM	Theater	Go from idea to product quickly with AWS serverless databases

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!

Arabinda Pani
arabindp@amazon.com

Rajesh Matkar
matkarr@amazon.com



Please complete the session survey in the **mobile app**

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.