

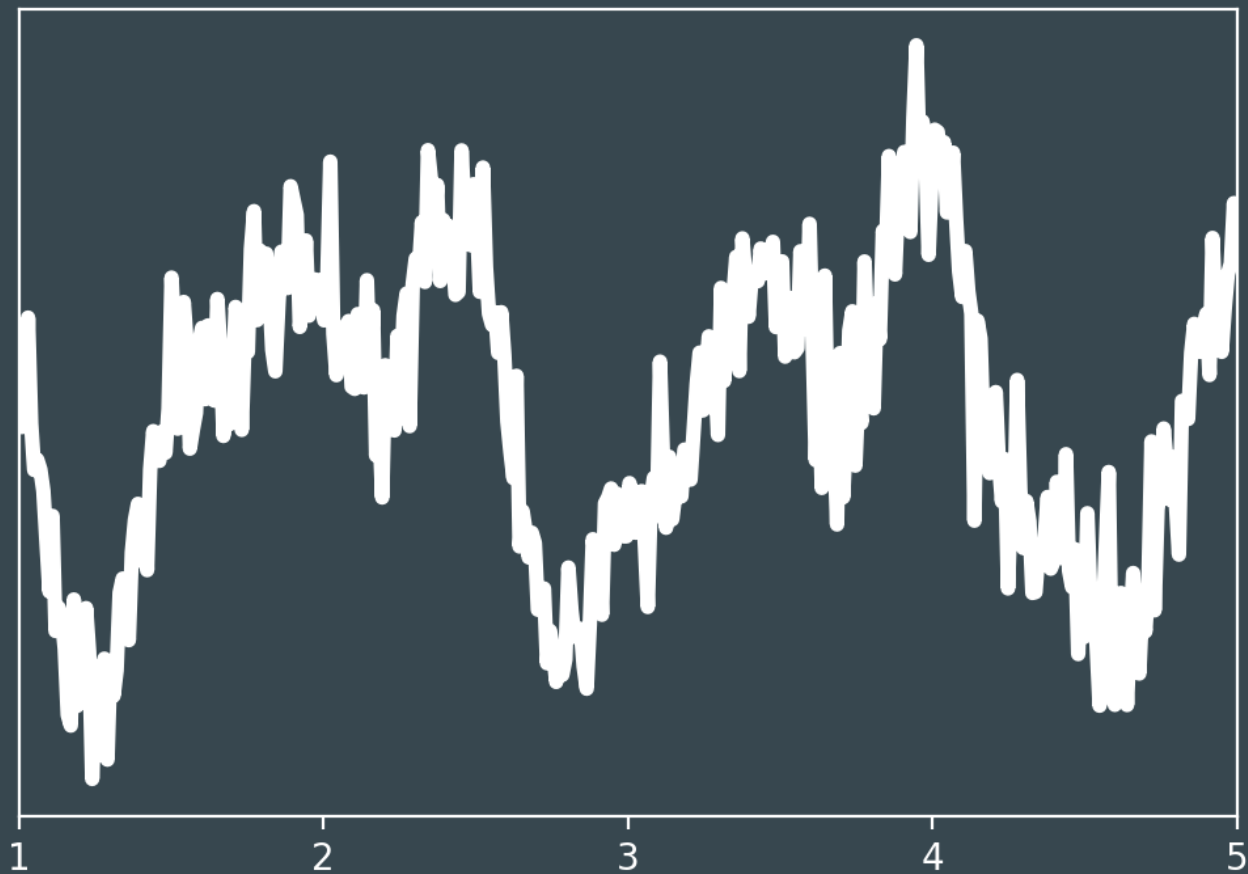
Five Hard Lessons from Five Years of So-Called Serverless Databases



Renato Losio

My (serverless) (db) journey so far

CPU, ACU, My Happiness with Serverless



COM202

Drawing the New York City skyline with Amazon Aurora Serverless v2

Renato Losio (he/him)

Principal Cloud Architect, AWS Data Hero
Funambol



Metrics

Metric

Serverless ... ▼

Statistic

Average ▼

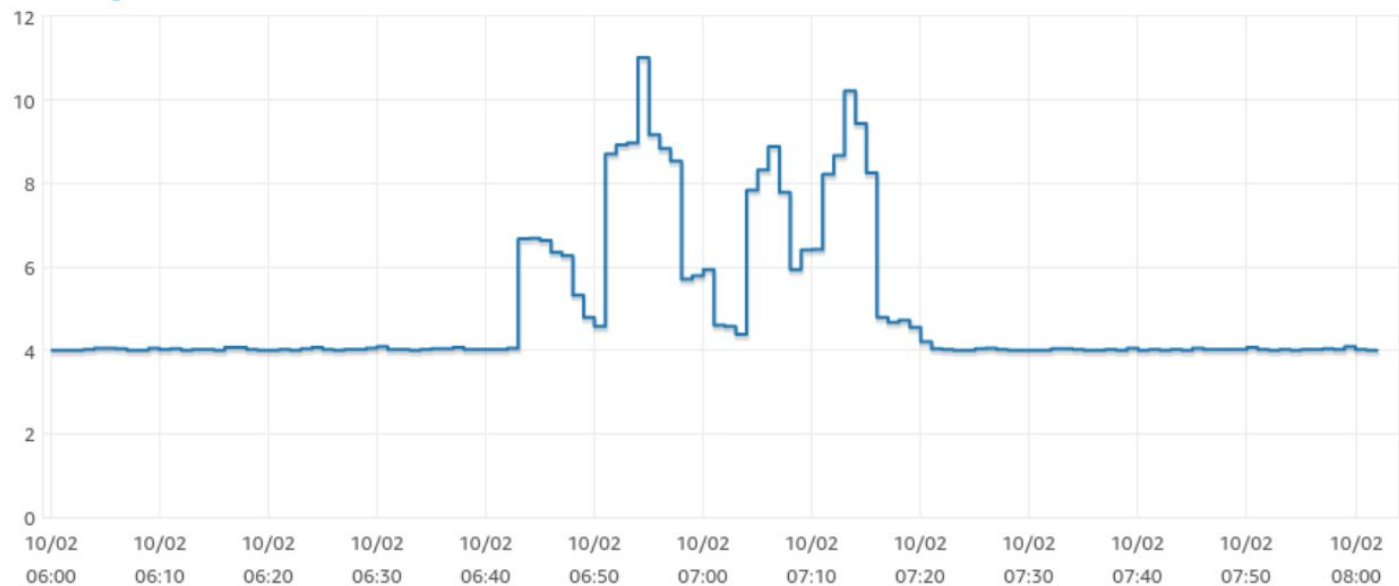
Time Range

Last 24 Ho... ▼

Period

1 Minute ▼

Legend: test-aurora-v2






ConsumedWriteCapacityUnits [🔗](#)

1h 3h 12h 1d 3d 1w Custom Stacked area Actions

Apply time range



Managing auto-scaling of RDS on AWS using the AWS CLI and Bash



Renato Losio

Published in funambol-techblog · 6 min read · Aug 1, 2019



31



1



Cloud services enable you to increase or decrease the capacity of your deployment within minutes, but **how can you scale a relational database automatically keeping predictable performance and costs under control?**

In this post we will show how Funambol manages the vertical scale up and scale down of all its MySQL databases running on RDS and handles peak traffic using the AWS Command Line Interface, AWS CloudWatch alarms and some simple logic that can be even written in a Bash script.

Five Hard Lessons

(serverless databases remove capacity management,
not database engineering)

#1 Serverless is Just an Adjective



Eight Years of Serverless DBs on AWS

- 2018 Aurora Serverless v1 GA
- **2018** **DynamoDB On-Demand**
- 2020 Amazon Timestream
- 2020 Aurora Serverless v2 Preview
- **2022** **Aurora Serverless v2 GA**
- 2022 Redshift Serverless
- 2022 Amazon Neptune Serverless
- 2023 ElastiCache Serverless
- **2025** **Amazon DSQL**
- 2025 DocumentDB Serverless

What about

- Amazon S3
- Amazon SQS
- Amazon Athena


#2 Storage is Underrated

“Aurora Serverless v2 scales instantly to support even the most demanding applications, delivering up to 90% cost savings compared to provisioning for peak capacity”

Storage and Backup now exceed Compute

USD/Month	2022	2023	2024	2025	2026
CPU	3336	4682	6006	6350	4908
Storage	1960	2387	2994	3755	3523
Backup	2274	3300	5840	6952	6650

Usage (GB-Month)



Super simple experiment

- Aurora Serverless: min 0.5 ACU
- Aurora Provisioned: db.r8g.large
- RDS for MySQL: 400 GB + db.r8g.large

Inserting 100K records

- One at a time
- Single thread
- ACID
- Default configuration

```
CREATE TABLE renato_demo
(id bigint(20) NOT NULL AUTO INCREMENT,
datetime TIMESTAMP NULL DEFAULT CURRENT_TIMESTAMP,
value integer DEFAULT NULL,
PRIMARY KEY (id));
```

```
DELIMITER $$
CREATE PROCEDURE renato_demo()
BEGIN
    DECLARE i INT DEFAULT 0;
    WHILE i < 100000 DO
        INSERT INTO renato_demo (value) VALUES (FLOOR(RAND()*1000));
        SET i = i + 1;
    END WHILE;
END$$
DELIMITER ;
```

```
CALL renato_demo();
```

Minutes?

Seconds?

how long?

Milliseconds?

Hours?

Aurora's storage architecture matters more than "serverless"

- Aurora Provisioned (2-3s)
- Aurora Serverless v2 (3-4s)
- Amazon RDS (5m+)



Development

Architecture & Design

AI Infrastructure

Culture & Methods

DevOps

Events

June
10,
2026**Online InfoQ Architect Certification**

The more senior you become, the fewer people pressure-test your decisions. This 5-week cohort gives you that check.
[Register Now.](#)

June
1-2,
2026**QCon AI Boston**

Learn how leading engineering teams run AI in production—reliably, securely, and at scale.
[Register Now.](#)

July
25,
2026**Online InfoQ AI Engineering Certification**

A practical online cohort for senior engineers making decisions around retrieval, agents, evals, and AI infrastructure.
[Register Now.](#)

Nov
16-20,
2026**QCon San Francisco**

Learn what's next in AI and software, from teams already doing it.
[Register Now.](#)

InfoQ Homepage > News > AWS Improves Aurora Serverless: 45% Faster Ramp-Up, 30% Higher Throughput

CLOUD

Architecting for Autonomous Reliability: Embedding AI into Your Observability Stack (Webinar Jun 25th)

AWS Improves Aurora Serverless: 45% Faster Ramp-Up, 30% Higher Throughput

MAY 09, 2026 • 2 MIN READ

[▶ Listen to this article - 4:03](#)

LIKE



by



Renato Losio

Cloud Expert | AWS Data Hero

Write for InfoQ**Feed your curiosity.**

Help 550k+ global senior developers each month stay ahead.

[Get in touch](#)

AWS recently [updated Amazon Aurora Serverless with a new platform version](#) that improves scaling behavior and runtime efficiency. According to the cloud provider, capacity can now scale up about 45% faster during demand spikes, and database performance improves by up to 30% through better resource scheduling and workload-aware scaling decisions.

Platform version 4 introduces improved runtime efficiency and a smarter scaling algorithm, enabling faster capacity scaling and higher database performance for Aurora Serverless clusters.

[Generally available for four years](#), Amazon Aurora Serverless is an on-demand, automatically scaling configuration for Aurora that adjusts database capacity based

RELATED CONTENT**Anthropic Launches Claude Platform on AWS**

MAY 13, 2026

AWS WorkSpaces Now Lets AI Agents Operate Legacy Desktop Applications without APIs

MAY 13, 2026

AWS Interconnect Reaches General Availability with Managed Multicloud and Last-Mile Connectivity

APR 29, 2026

AWS Ends WorkMail and Moves App Runner to Maintenance Mode

APR 26, 2026





#3 Costs Need a Crystal Ball

If you do not know how your workload performs on a serverless DB, **forecasting ACU** by using the baseline of a provisioned cluster is **useless**.



Query 2 : X | Query 3 : X | Query 4 : X | Query 5 : X

(+)

```
1 select count(*) from *****_hourly_cost_and_usage_reports where line_item_resource_id = 'arn:aws:rds:eu-west-1:053652639853:db:****-prod-rds01' ;
```

SQL Ln 1, Col 144



Run Explain Cancel Clear Create

Reuse query results up to 60 minutes ago

Query results Query stats

Completed Time in queue: 83 ms Run time: 887 ms Data scanned: 45.37 MB

Results (1)

Copy Download results CSV

Search rows

< 1 > Settings

_col0

1 1246

“Aurora DSQL is now Generally Available. What’s it cost? Nobody knows, especially AWS. You get charged per DPU, which equates in the documentation to “screw you, benchmark your workloads and find out for yourself.” Of all the pricing strategies from which to choose, I didn’t expect AWS to pick “completely give up.” My Aurora customers are... displeased.”

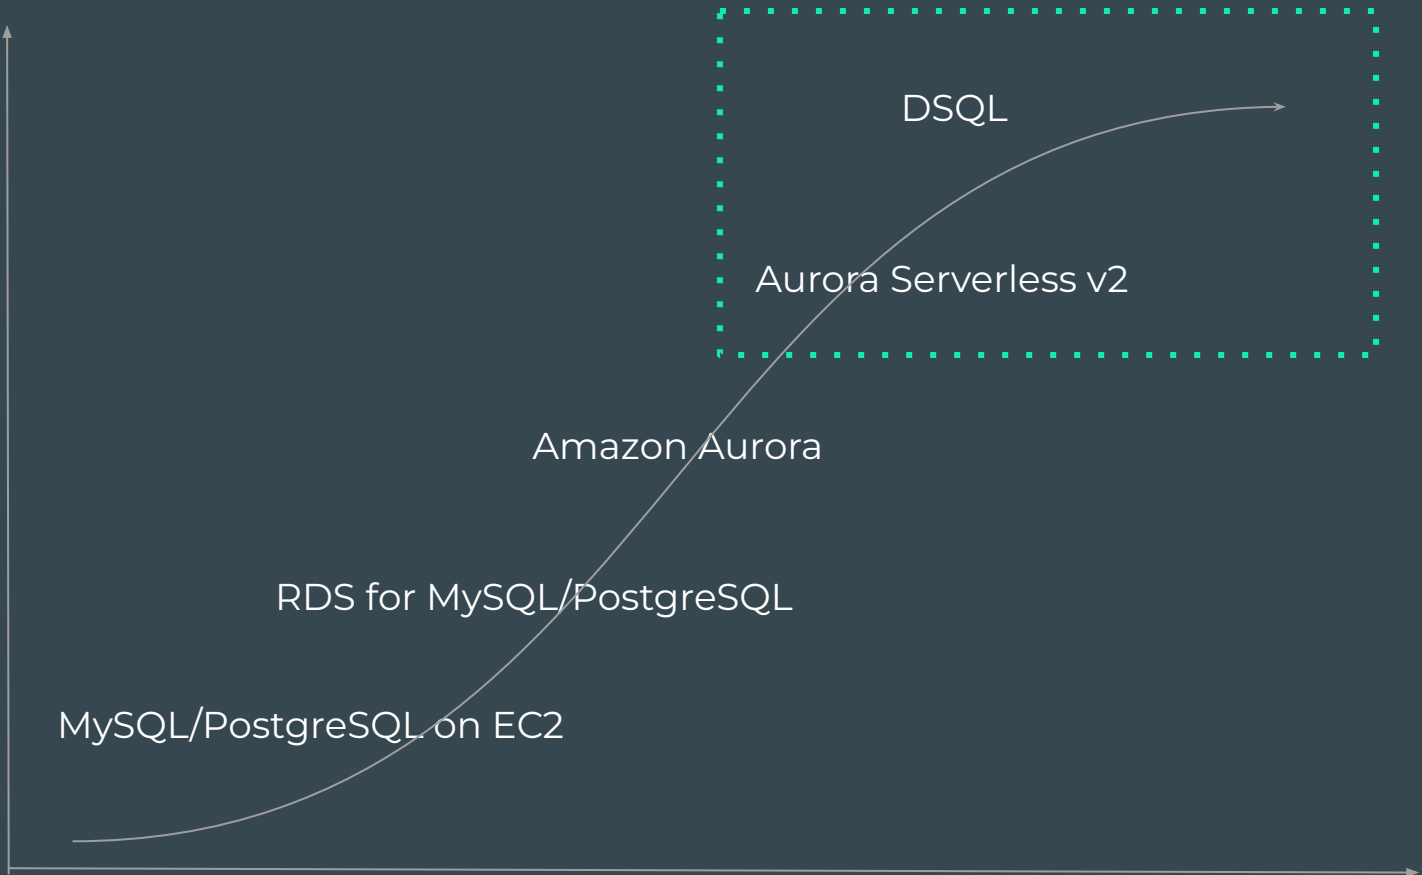
Corey Quinn

One thing for sure, I found Aurora DSQL cheaper than I expected. The pricing is smart to cover AWS's margin for most types of abuse. I'd rather not summarize anything. I think AWS didn't put the pricing in simpler terms due to the nature of relational database queries, as they heavily depend on query plans and index usage.

<https://blog.faridnsh.ninja/unraveling-aurora-dsql-pricing>

#4 Compatibility is Gray Area

managed



compatibility

Compatibility...

- PostgreSQL-compatible until your extension isn't supported.
- MongoDB-compatible until your aggregation behaves differently.
- MySQL wire-compatible but operationally very different.

Compatibility?


Wire (or Client Protocol)

SQL / Query Language

Feature / Behavior

Operational / Ecosystem

Version




ANNOUNCEMENT

Linux Foundation Welcomes DocumentDB to Advance Open, Developer-First NoSQL Innovation



documentdb.io

#5 Going **Horizontal** or Vertical?




DSQL, Dynamo DB



Aurora Serverless (*)

both "Serverless" databases





Is vertical scaling bad “serverless”? **No**

elasticity

- serverless

“If the database still doesn't scale down to the minimum capacity configured, then **stop and restart the database** to reclaim any memory fragments that might have built up over time. Stopping and starting a database results in downtime, so we recommend doing this sparingly.”

(AWS doc)

Leverage managed (**serverless**) database, work on **elasticity**, but understand what runs **under the hood**.

thank you!



Renato Losio

Cloud Expert | AWS Data Hero | InfoQ
Editor | Remote-First Advocate 

