

dbtune

# A benchmark study on the impact of PostgreSQL server parameter tuning

P2D2 conference, Prague

January 28, 2026



**Luigi Nardi**

Founder & CEO, DBtune

# PG Developer Day Prague training

January 31st, 2023



On the left, a photo of our training session. On the top right three members of the DBtune team and on the bottom the main track.

# About me

Mixed background in industry and academia

Among other things:

- Ph.D. CS at Sorbonne
- Research Staff at Stanford University
- Associate Professor in AI at Lund University

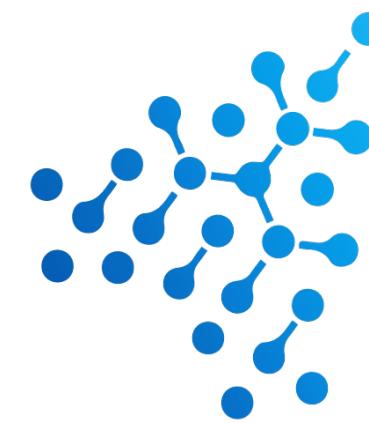
 @luinardi

Since 2020, Founder & CEO at DBtune



Imperial College  
London

 dbtune

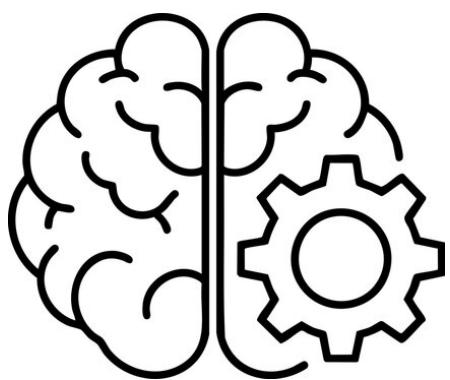


dbtune



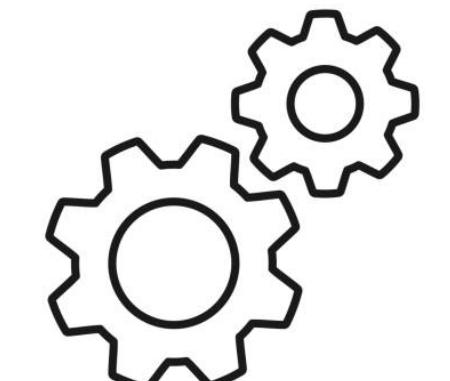
## What

*DBtune is an AI-powered database tuning service*



## Where

Spun out of research at  
Stanford University



## How

Tunes for a specific workload,  
use case and machine

# DBtune now supports CloudNativePG

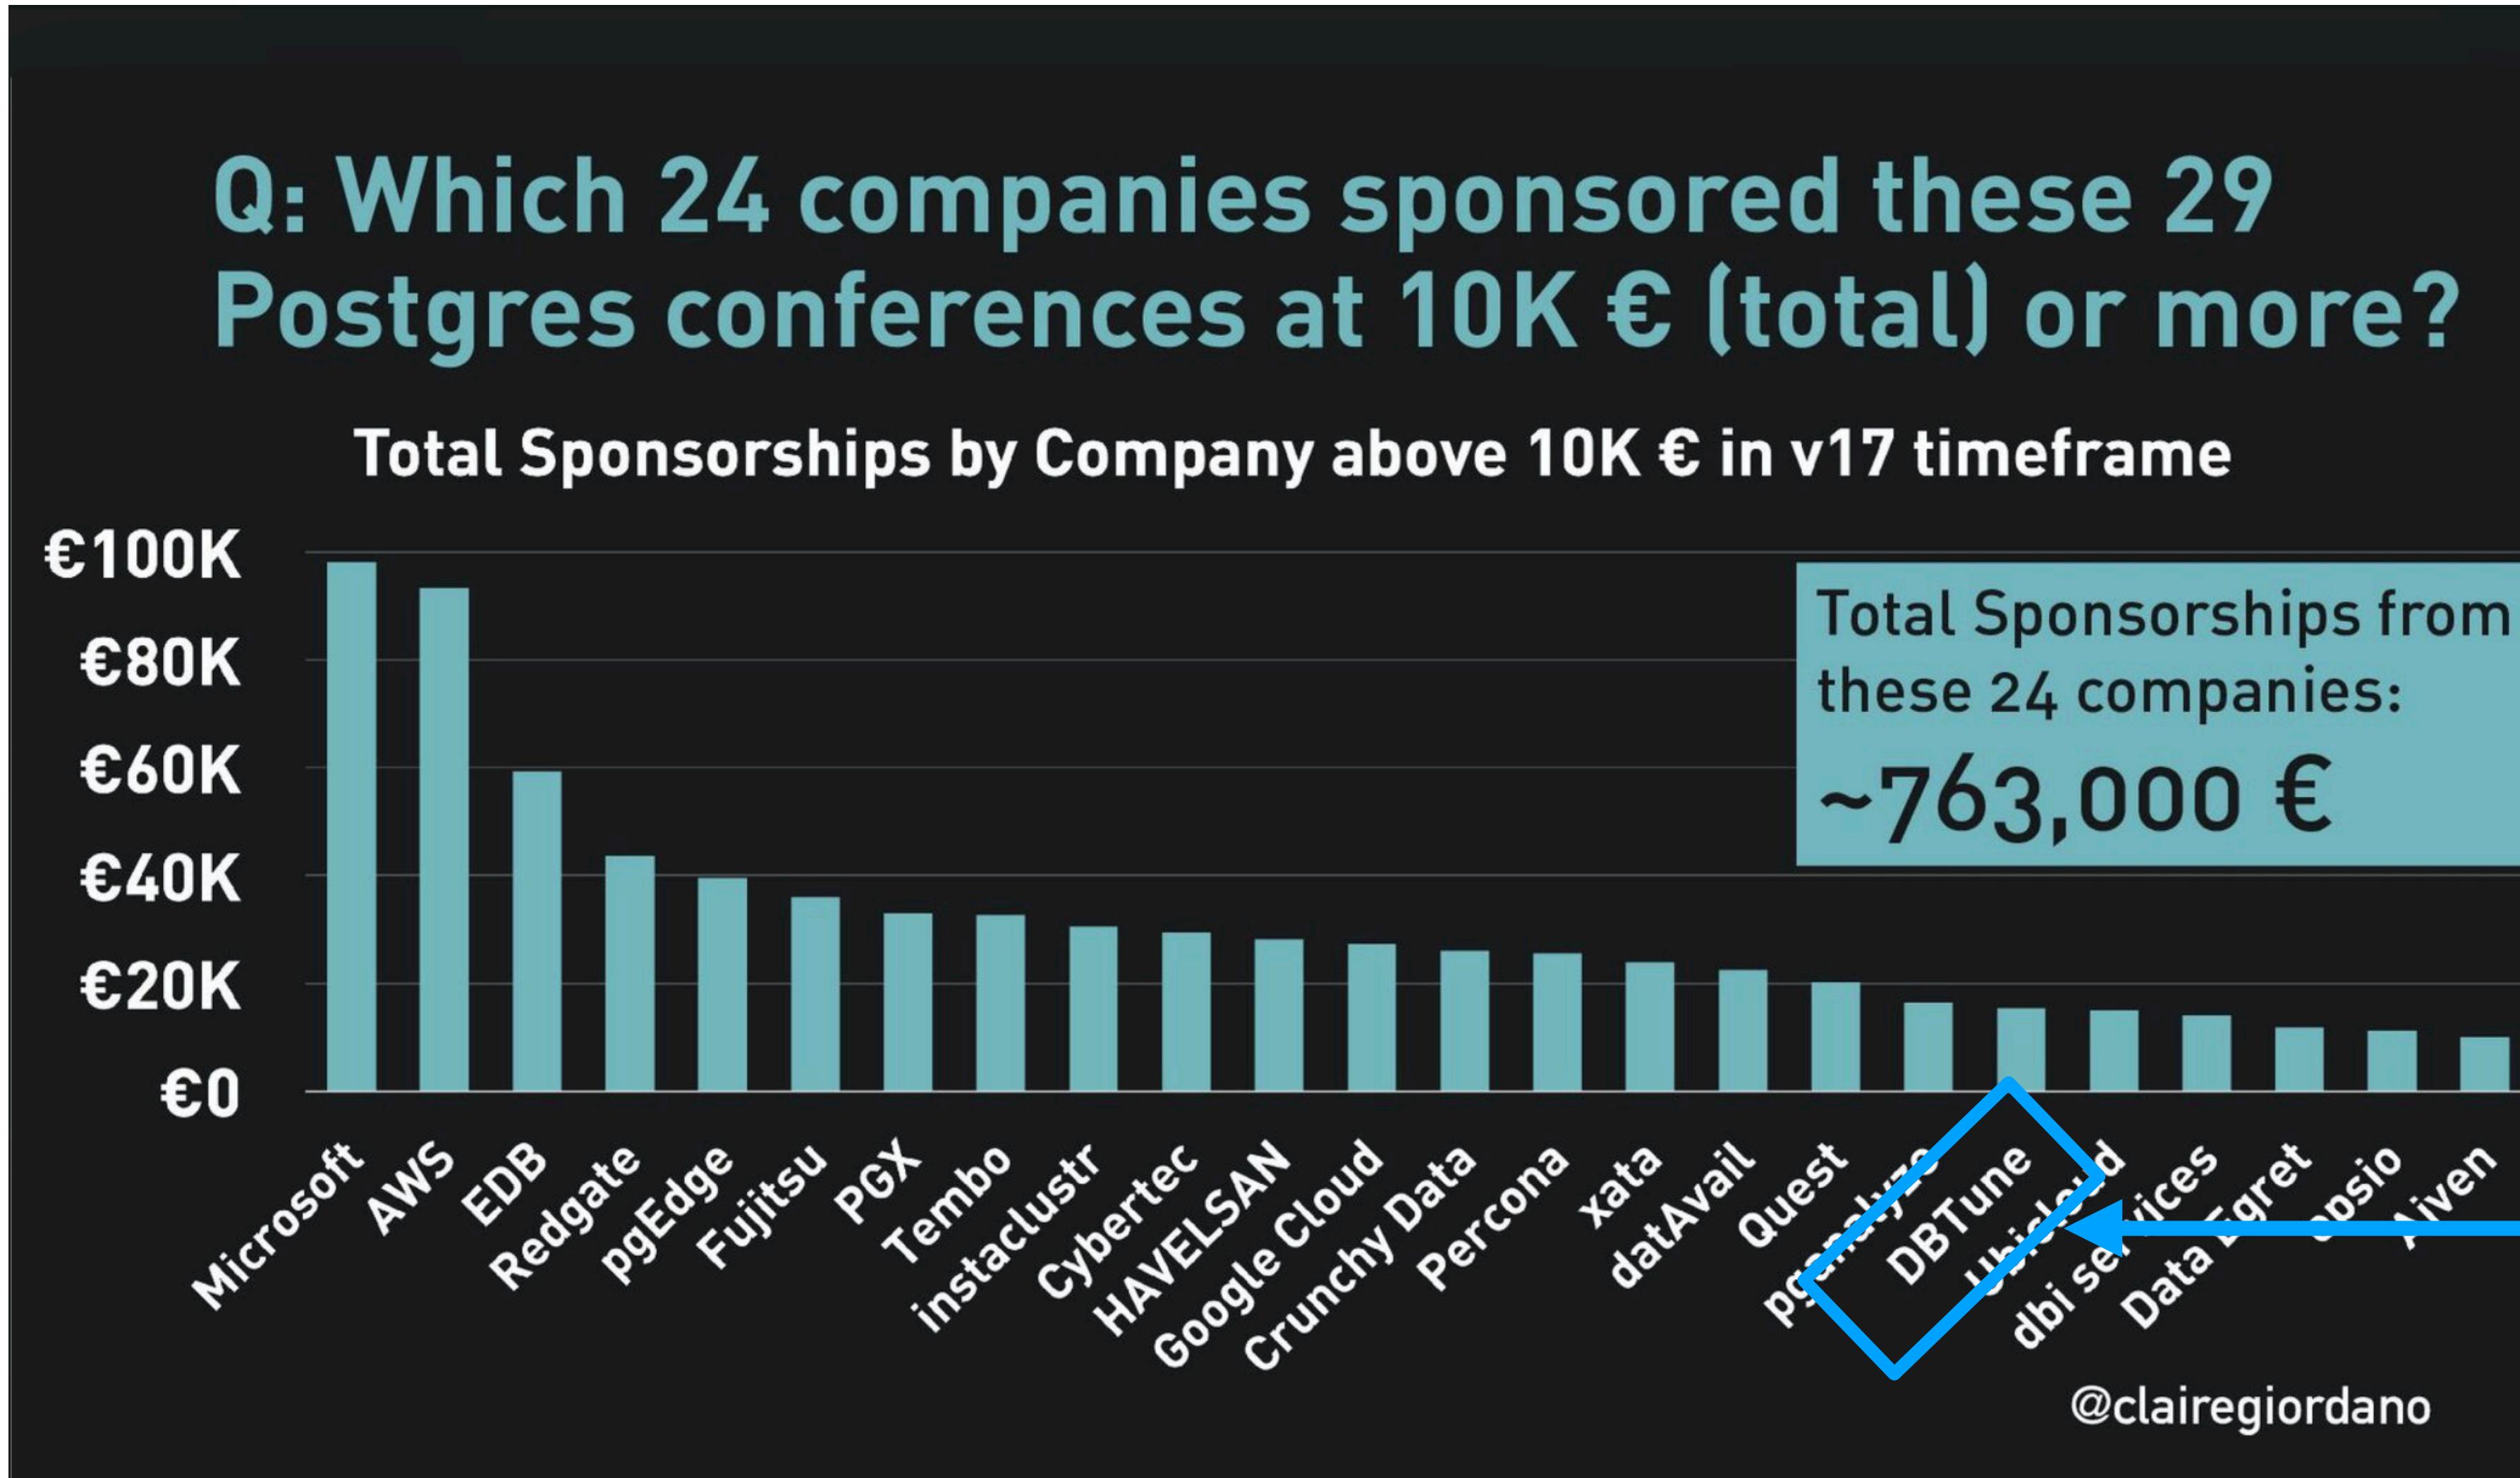
Public release date: January 27, 2026 at P2D2

Tuning CloudNativePG  
without the guesswork



Read the blog post: <https://www.dbtune.com/blog/tuning-cloudnativepg-without-the-guesswork>

# DBtune in the top 20 PostgreSQL sponsors



<https://speakerdeck.com/clairegiordano/whats-in-a-postgres-major-release-an-analysis-of-contributions-in-the-v17-timeframe-claire-giordano-pgconf-eu-2024>

# Malmö PostgreSQL User Group (M-PUG)

## M-PUG organizers



**Ellyne Phneah**  
DBtune

**Luigi Nardi**  
DBtune

**Daniel Gustafsson**  
Microsoft

**Dennis Rilorin**  
Redpill Linpro

- The group is officially recognized by PostgreSQL Europe
- Regular meetups every 4-8 weeks in Malmö — Top speakers
- We are building a vibrant PostgreSQL community in the region

# What is this talk about?

-  Is it worth tuning your server parameters? A benchmark study
-  In this talk you won't learn how to manually tune your PostgreSQL server

# Outline



Introduction to PostgreSQL server parameter performance tuning



Quantitative illustrative examples



How do we solve this today?



Results synthetic and prod workloads



Conclusions

# What is PostgreSQL tuning?

# What is database tuning?

Keeping the database fit and responsive

-  Databases change, grow and slow down
-  Not all workloads and machines are the same
-  **Tuning adapts a database to its current use-case, load and machine**
-  It is a ‘dark-art’ yet an integral part of any DBA and developer’s job
-  Tuning includes query, **server parameters\***, index, OS parameters, etc.

\*We will focus solely on PostgreSQL server parameter tuning

# Why does it matter?

## Technical perspective

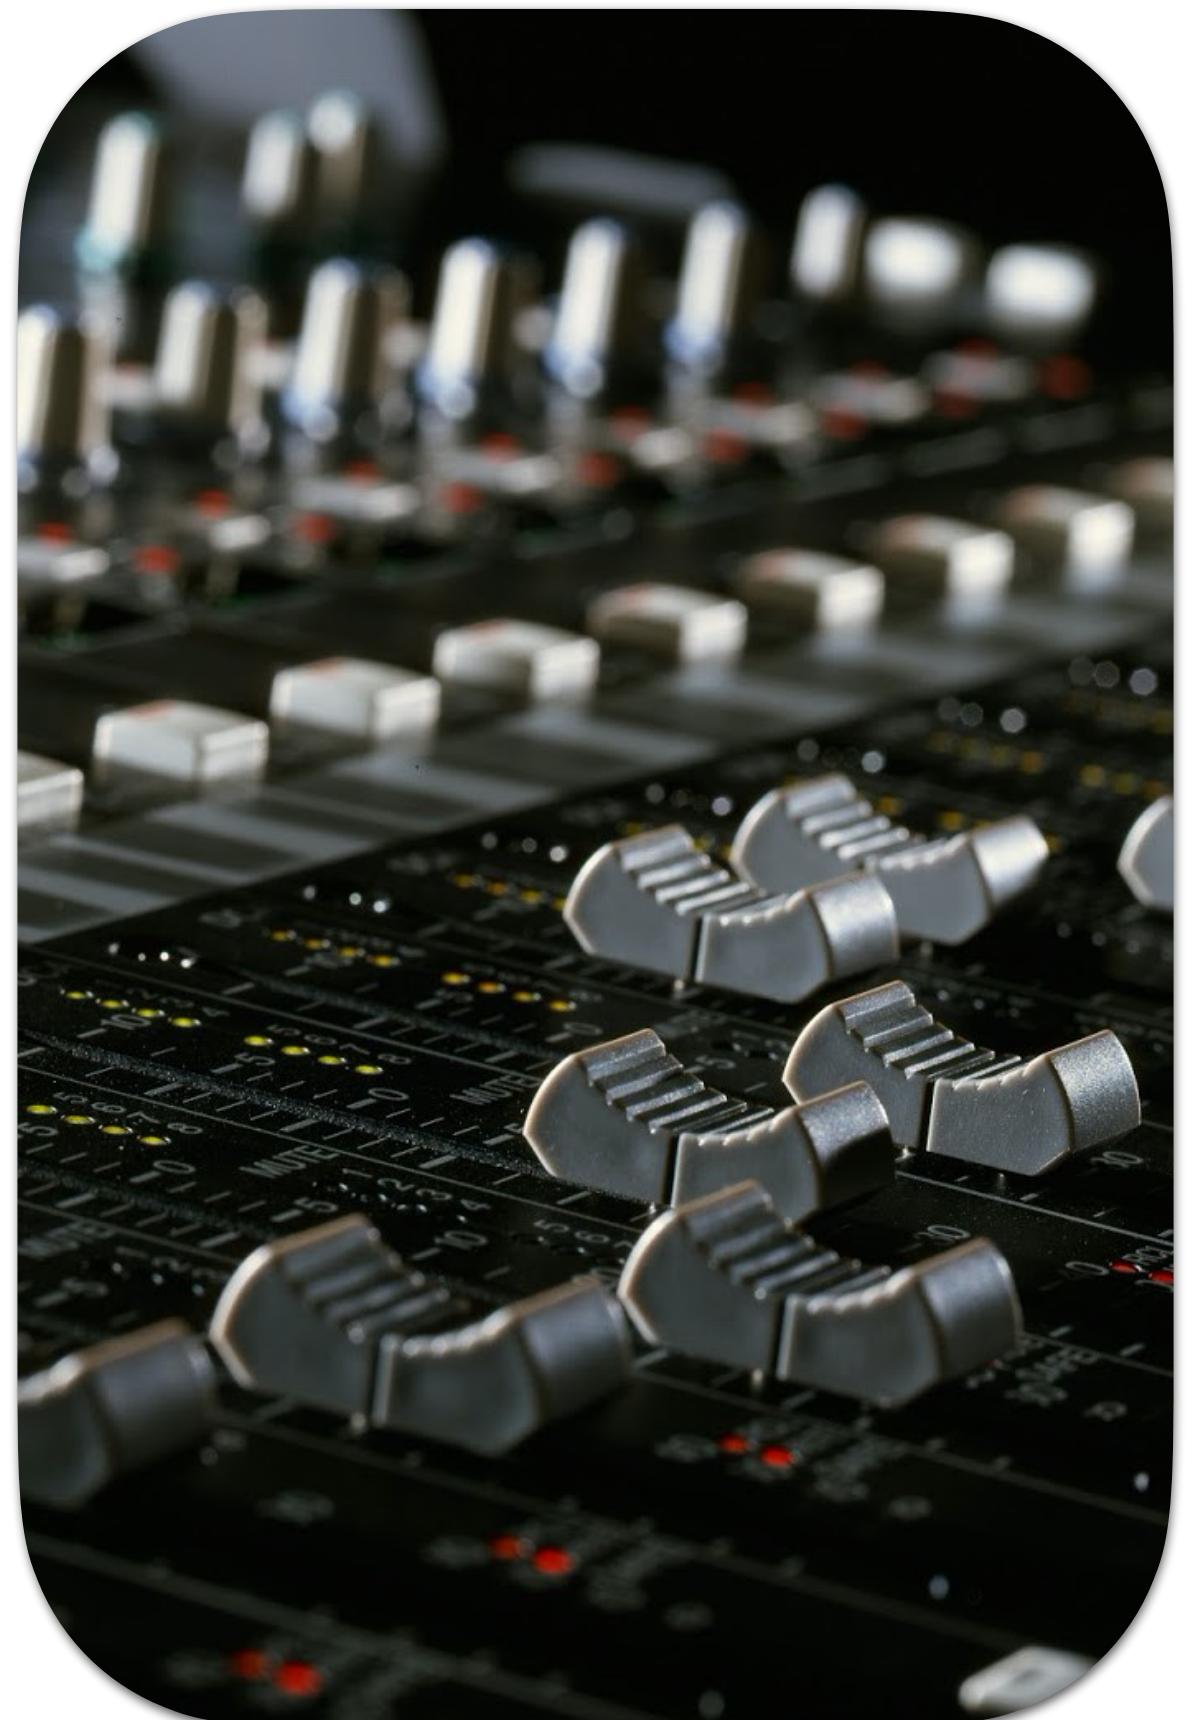
- Impacts system performance
  - Throughput and latency
  - Improves scalability / stability / SLA

## Business perspective

- Decreases infrastructure spend
- Higher end-user satisfaction
- Reduces downtime
- Increases productivity
- Saves energy (ESG)

# PostgreSQL server parameter tuning

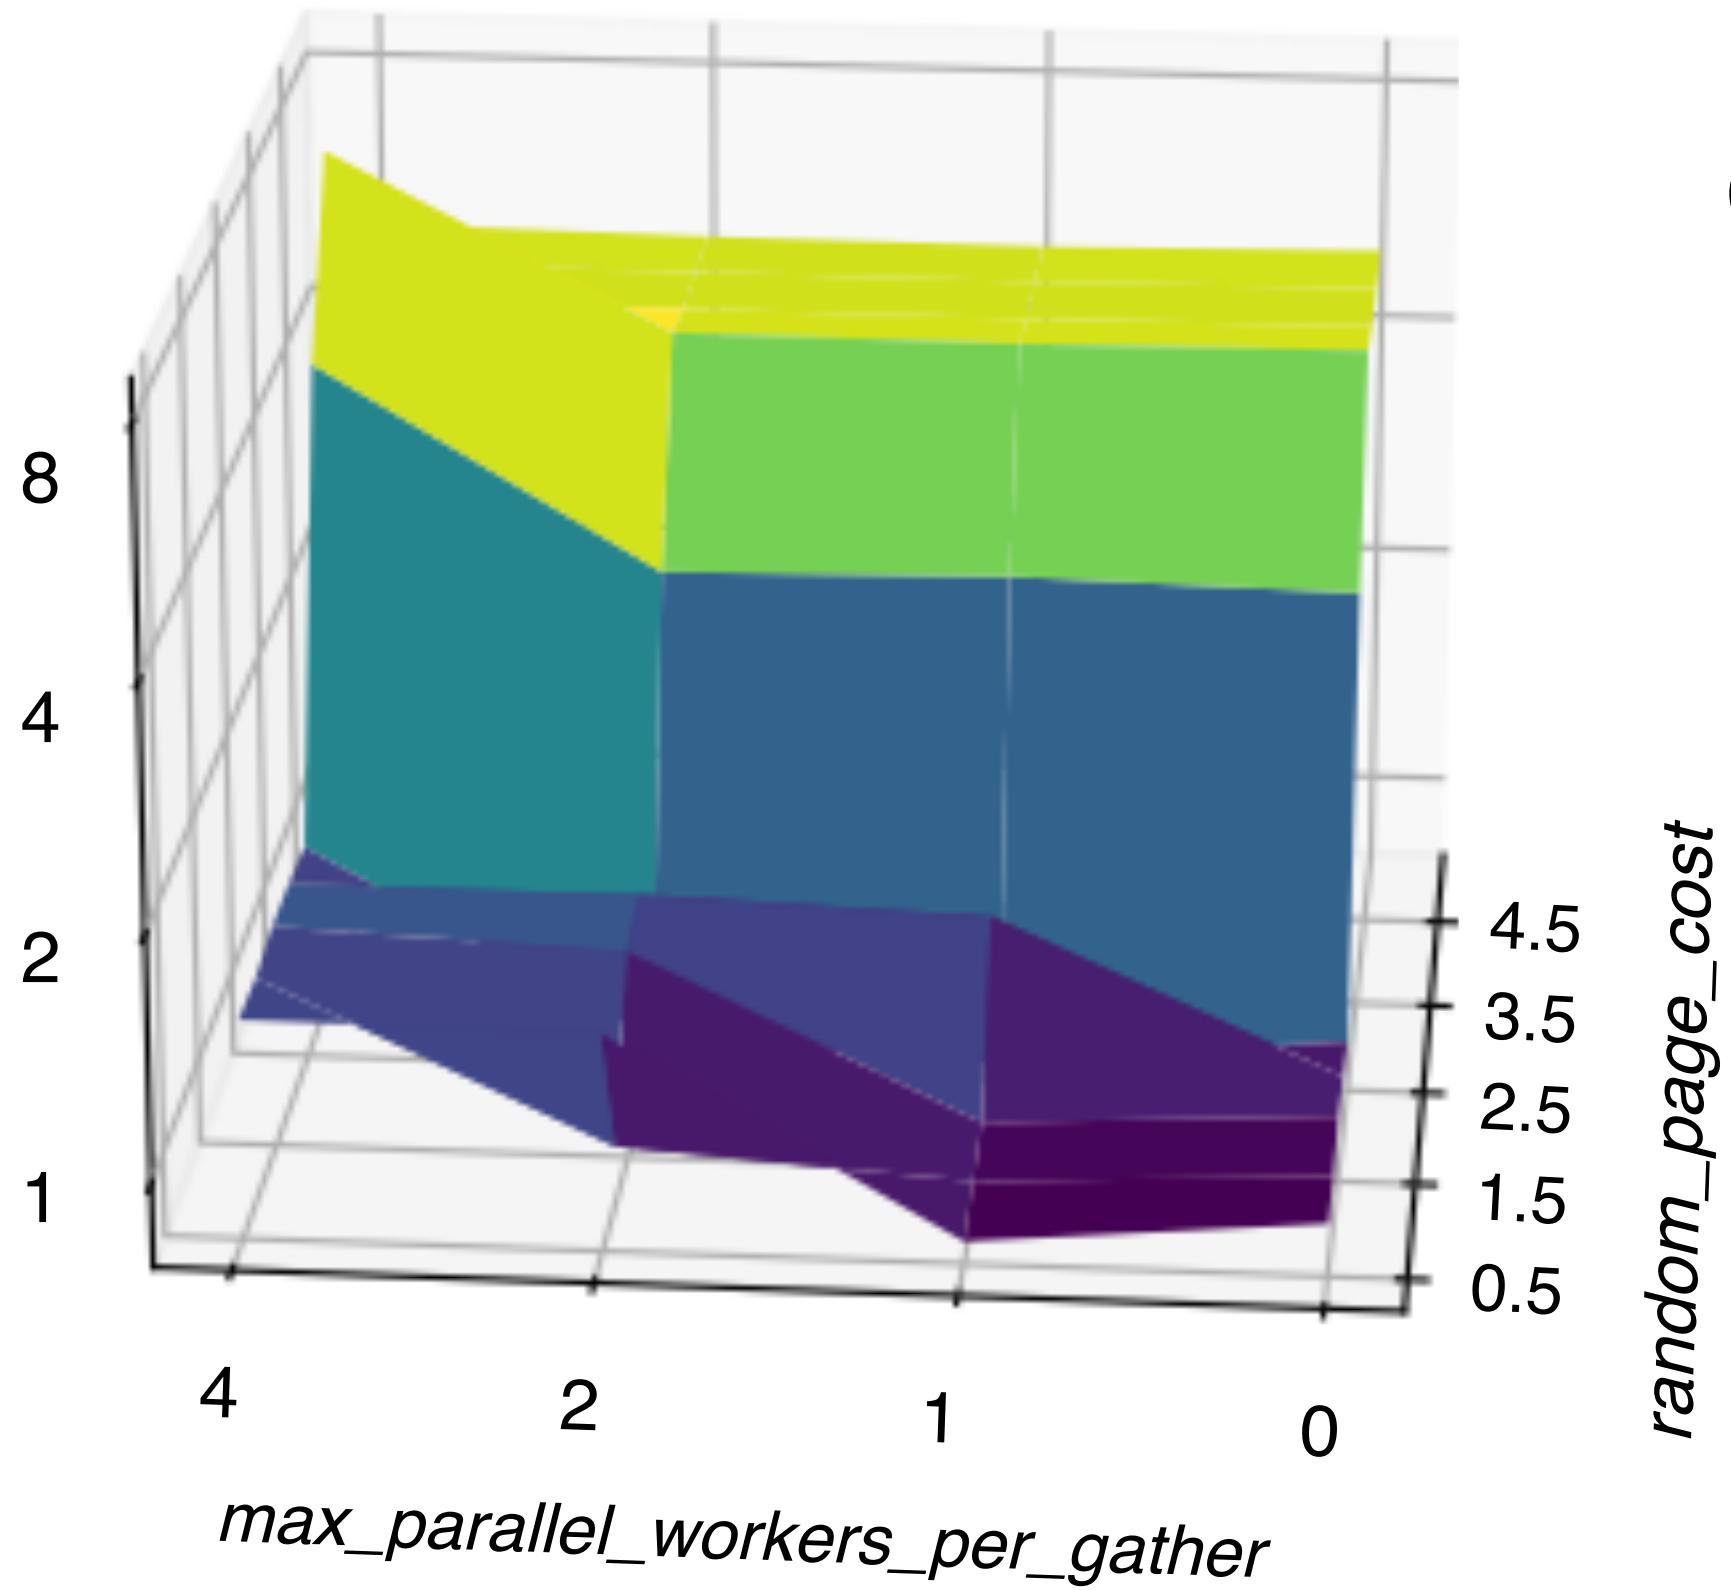
- ✓ Adjusting knobs to best fit the workload
- ✓ PostgreSQL parameters that are typically important: *work\_mem*, *shared\_buffers*, *max\_wal\_size*, etc.
- ✓ Example *max\_parallel\_workers\_per\_gather*:  
Max # of workers started by a Gather or Gather Merge node
- ✓ Example *random\_page\_cost*:  
Planner's cost of a non-sequentially fetched disk page
- ✓ These parameters highly depend on the application



# Average query runtime tuning

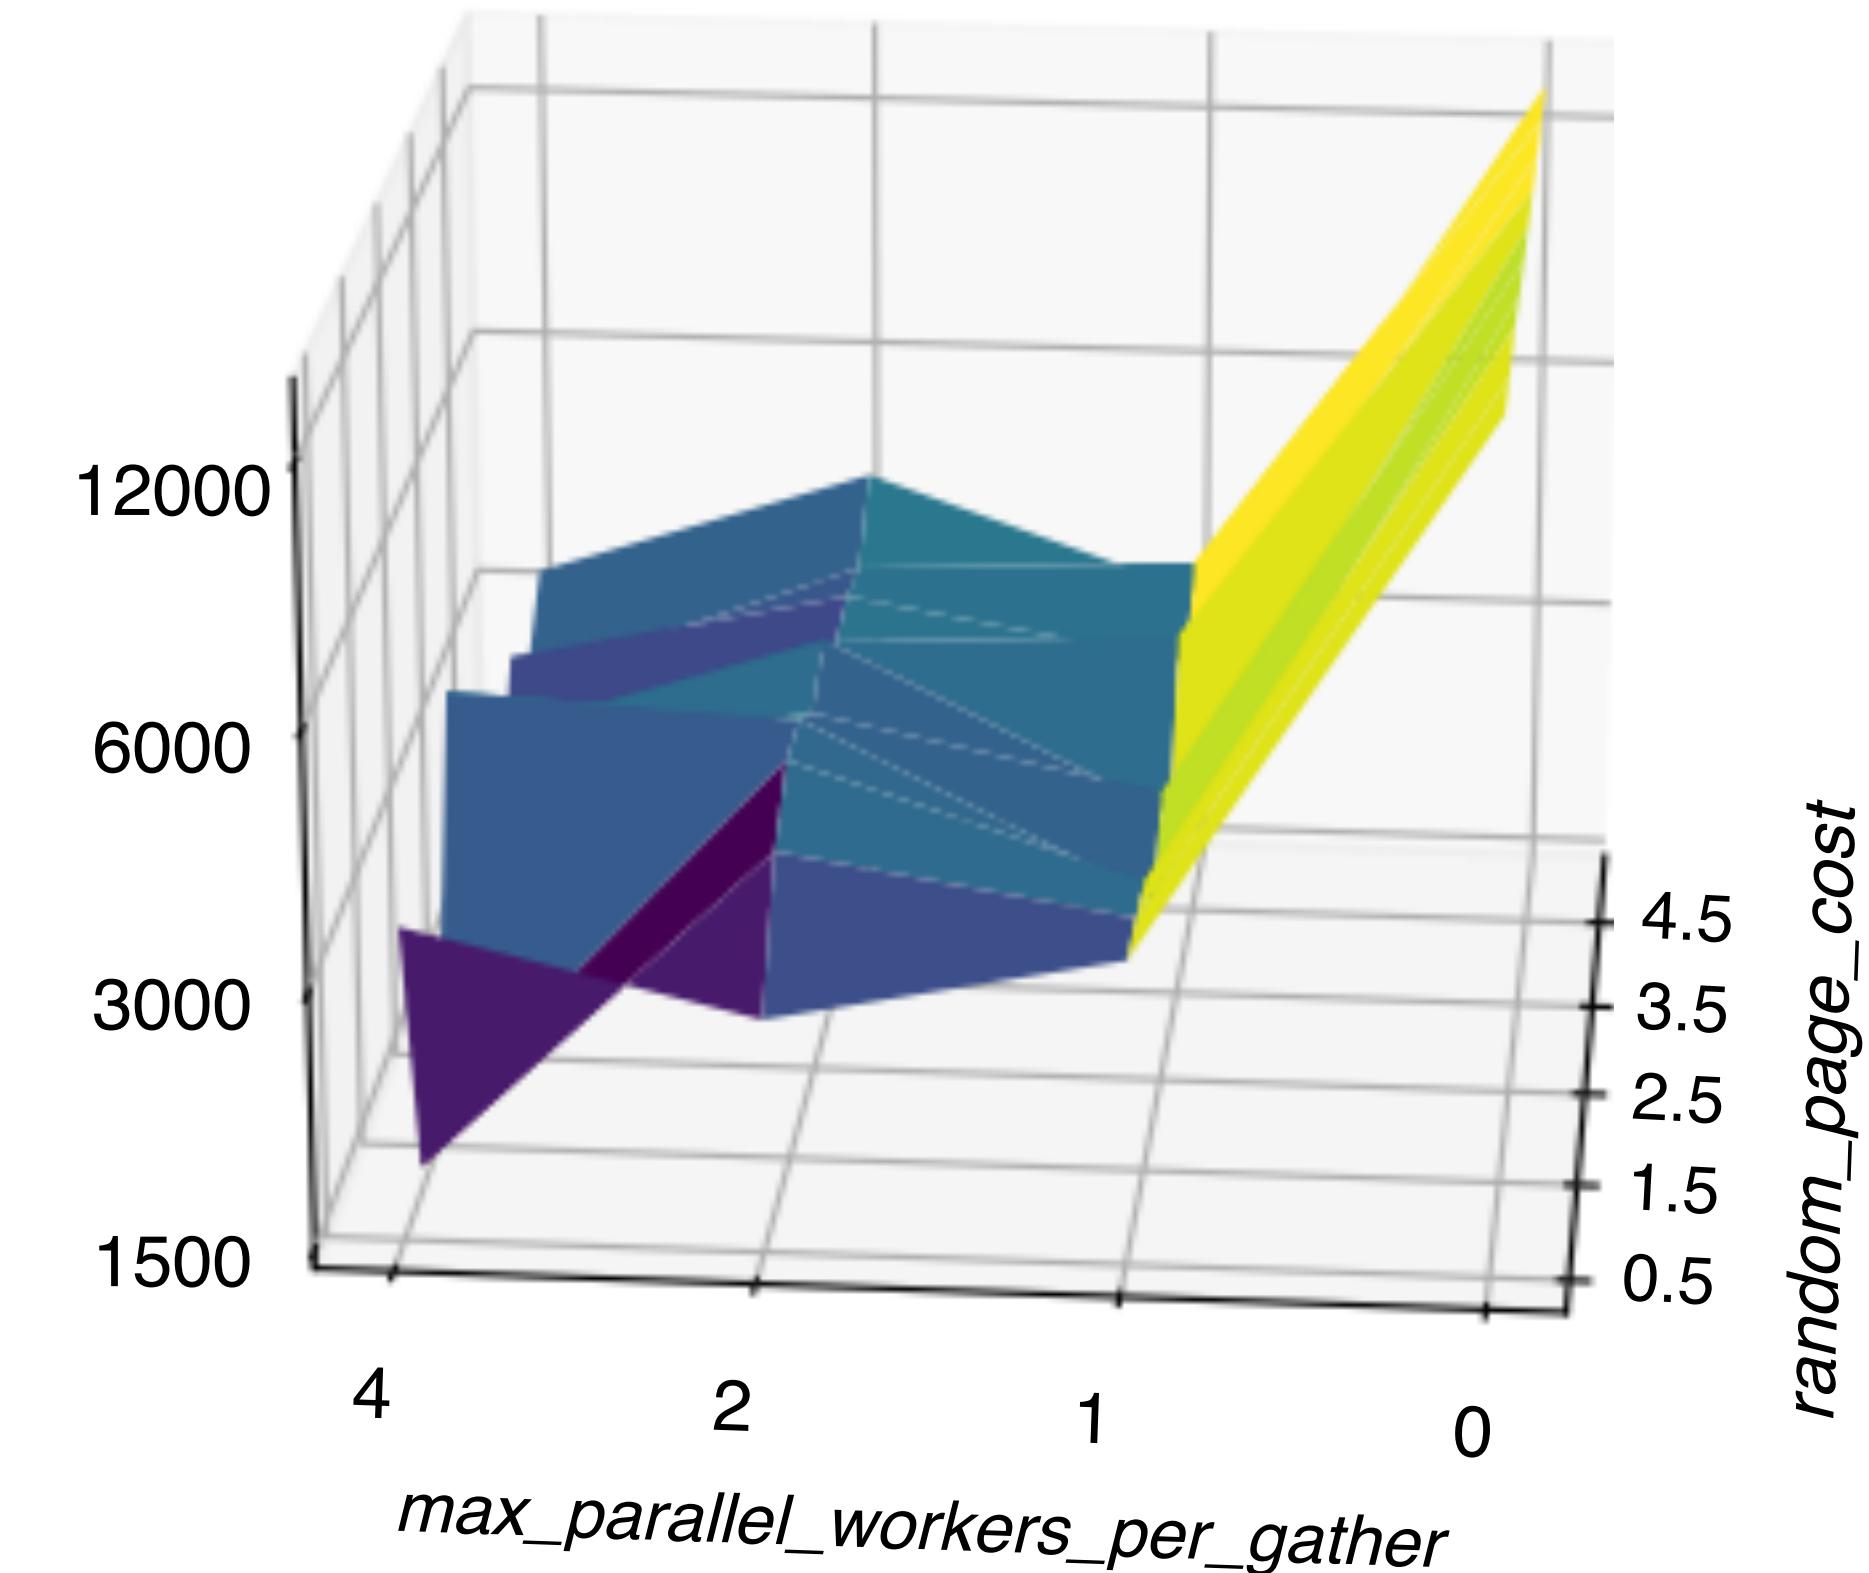
for *max\_parallel\_workers\_per\_gather* and *random\_page\_cost*

Epinions



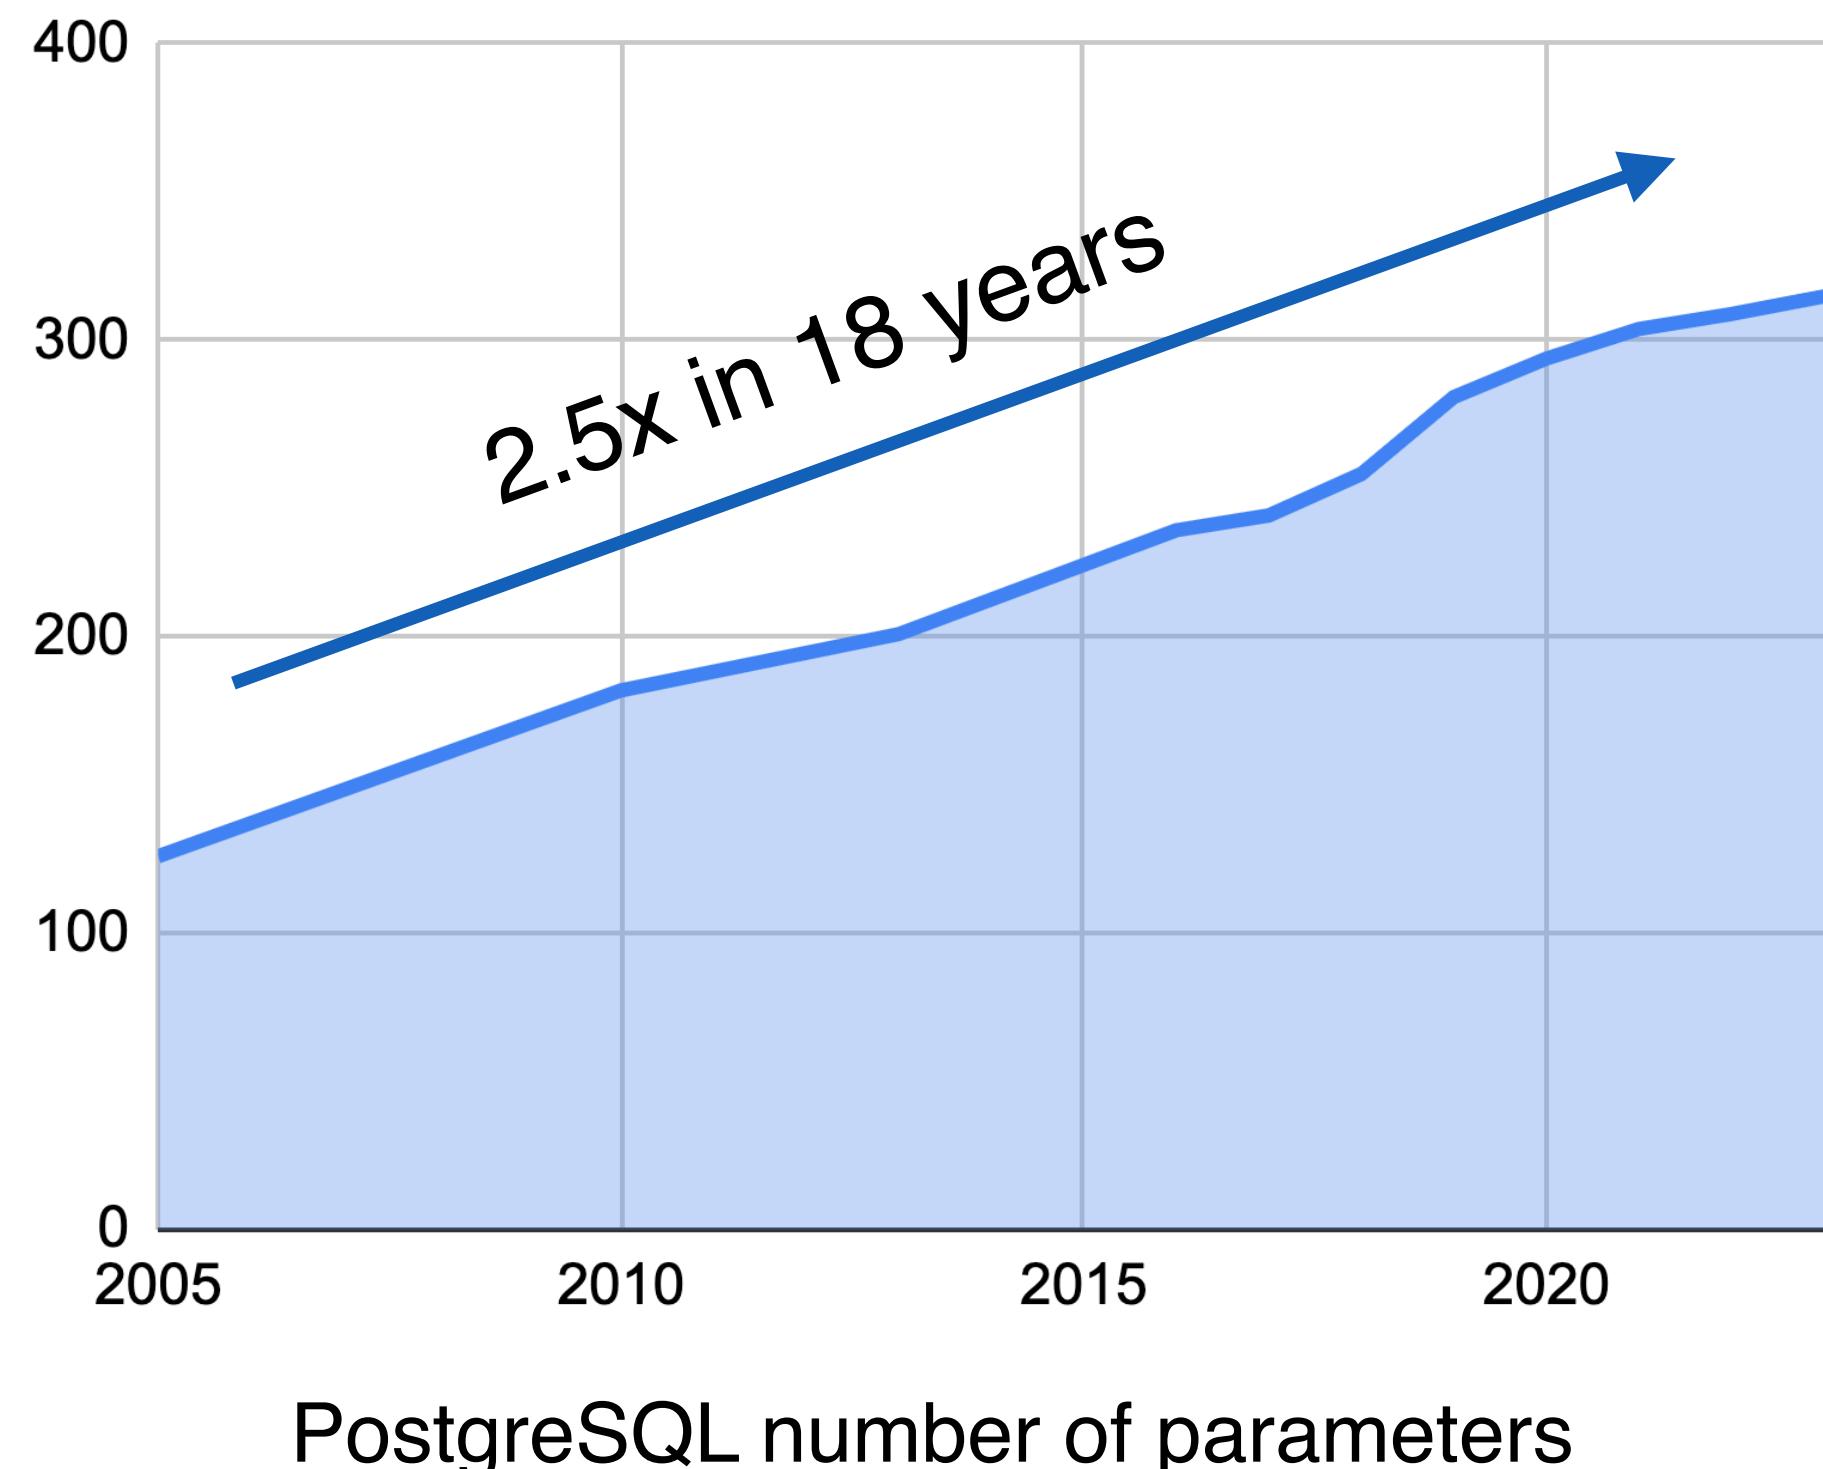
Query runtime in ms  
**Lower the better**

TPC-H

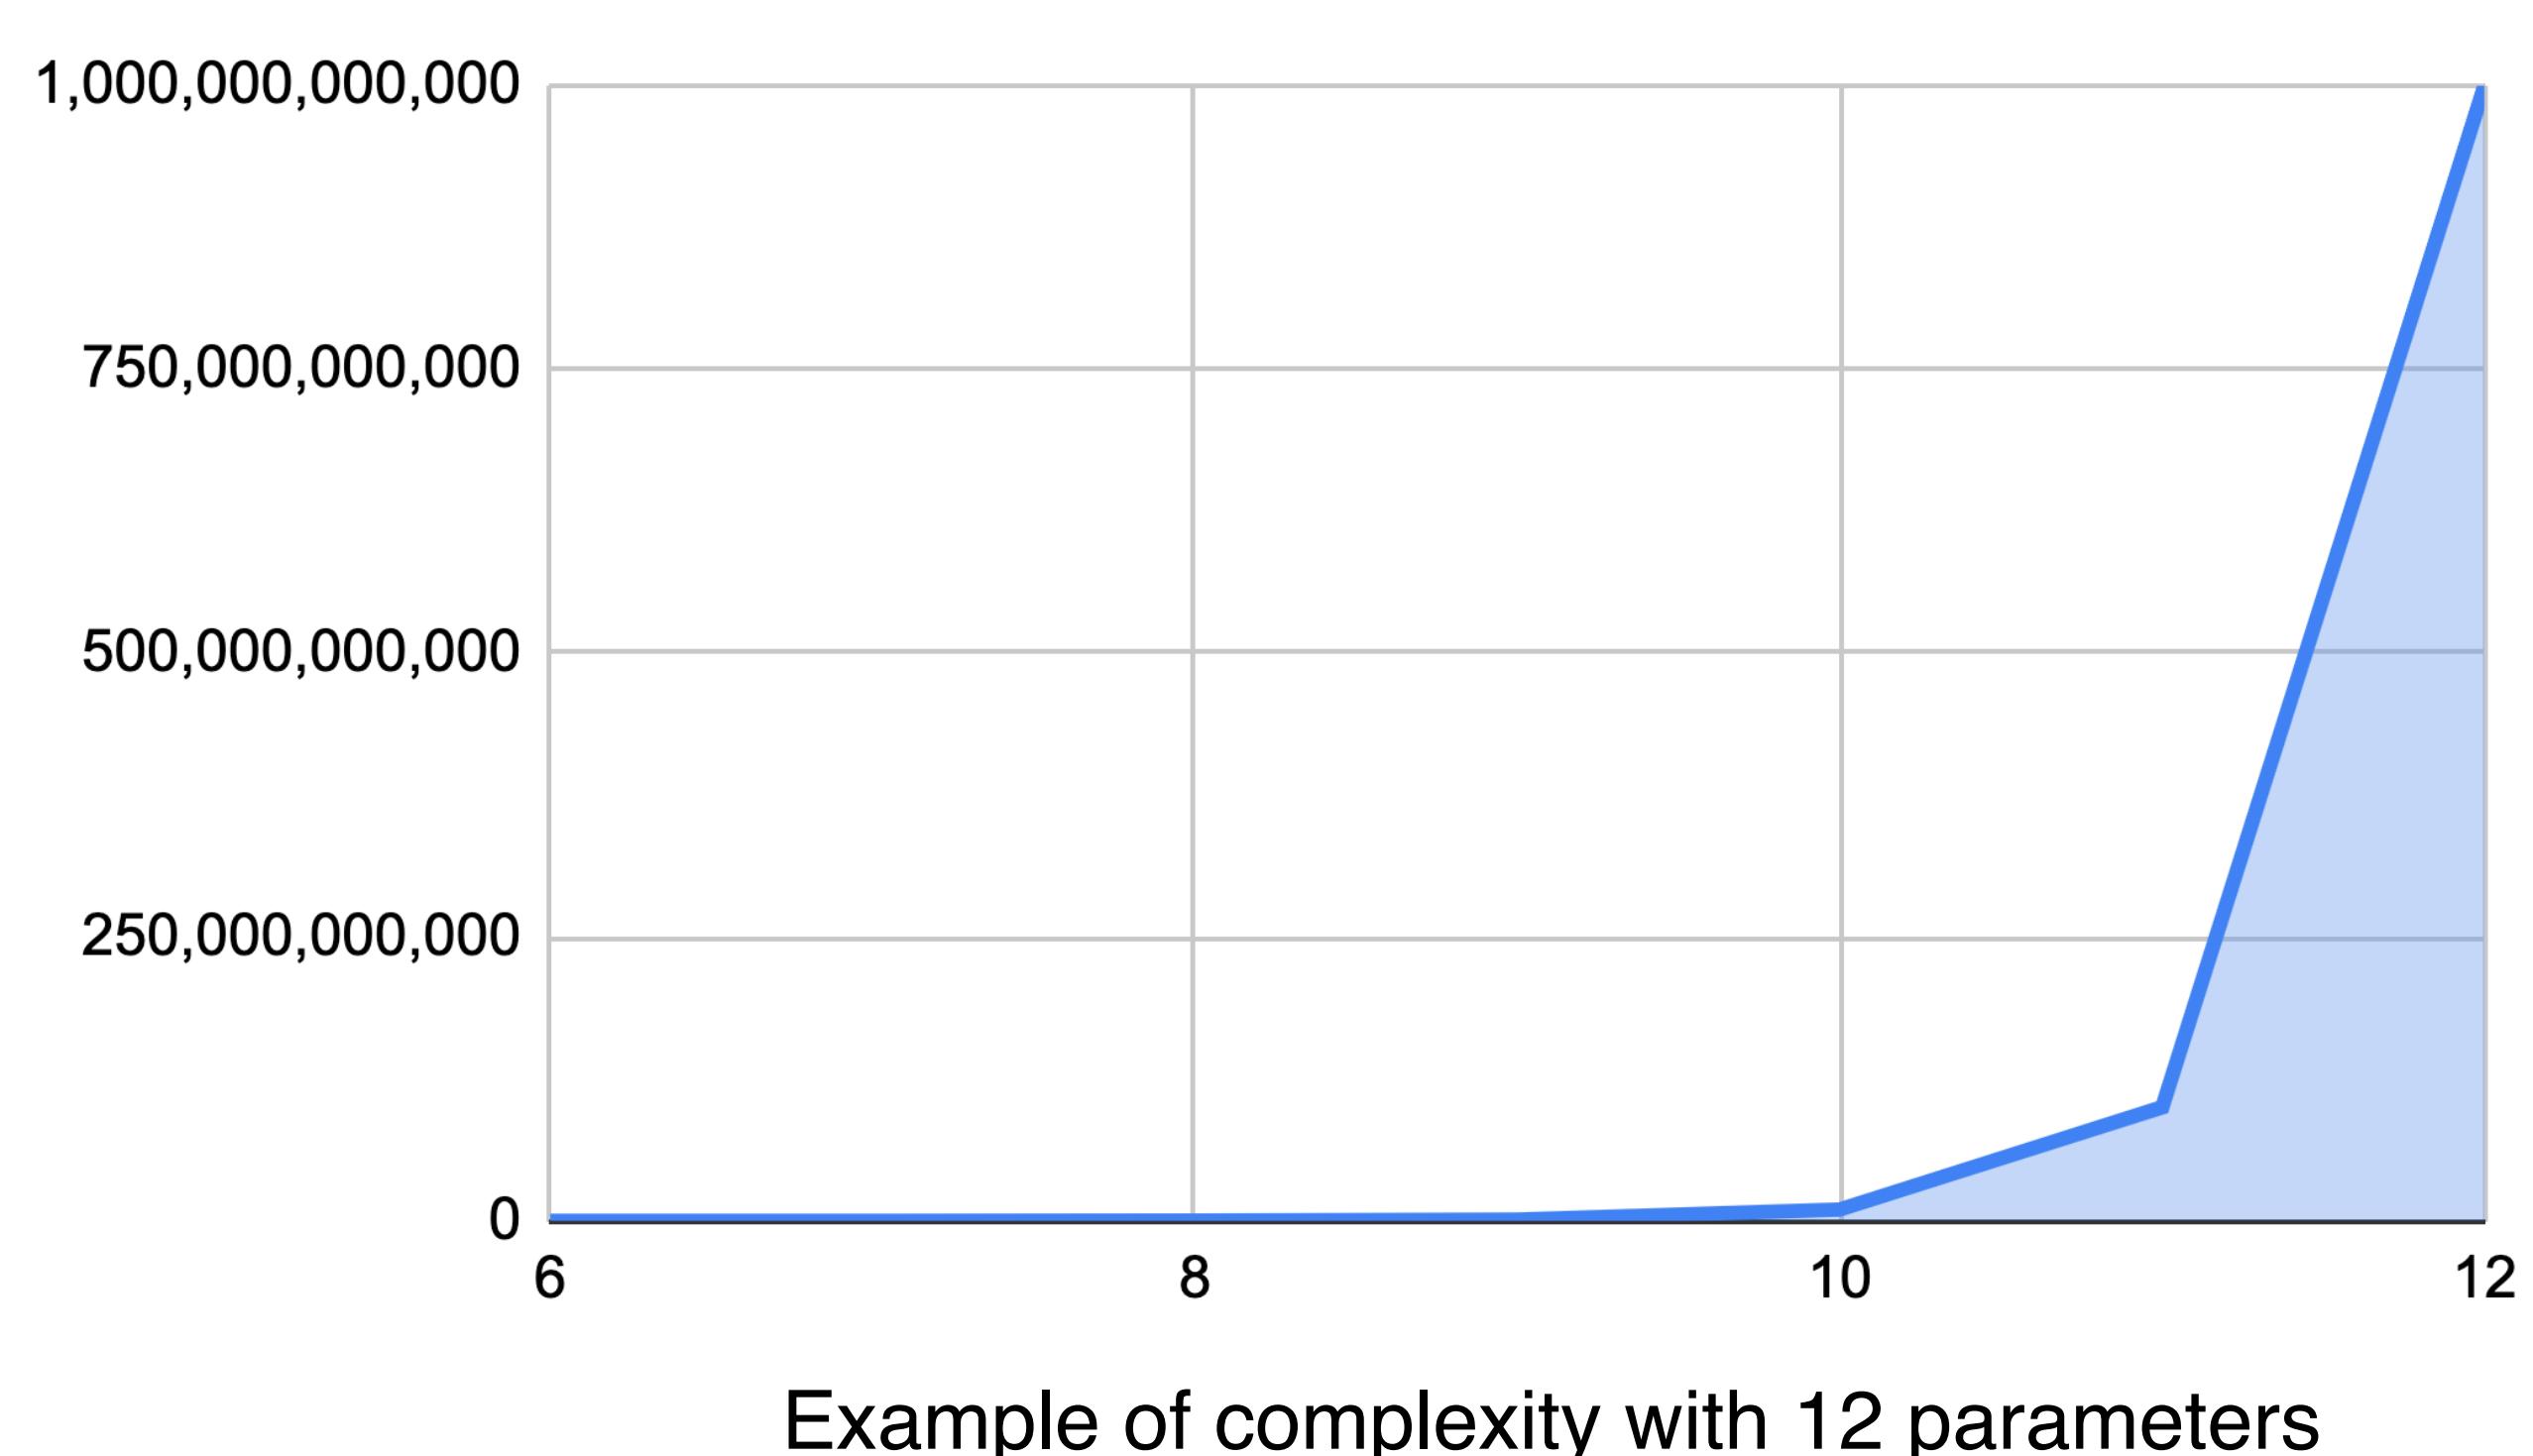


# Complexity is growing over time

The number of parameters  
is growing **linearly**



The number of configurations  
is growing **exponentially**



# How is parameter tuning tackled today by DBAs and developers?



## Manual

Slow

Takes days

Painstaking

Needs high expertise

Ineffective

Tune again in a week

Inadequate

Seasonal workload

Tuning guru

## Heuristics



One-size-fits-all

Uses generic rules

Workload agnostic

Not bespoke

Ineffective

Tune again in a week

Inadequate

Seasonal workload



## AI approach

AI agent that learns by **observation**, adapts to changing workloads, and **autotunes** under minimal supervision

# Heuristic-based server parameter tuning

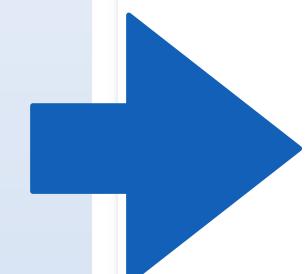
## Heuristics

One-size-fits-all  
Uses generic rules

Workload agnostic  
Not bespoke

Ineffective  
Tune again in a week

Inadequate  
Seasonal workload



**PGTune**

**Parameters of your system**

DB version: 17

OS Type: Linux

DB Type: Web application

Total Memory (RAM):  
Memory size (RAM, required): 64 GB

Number of CPUs: 9

Number of Connections: 41

Data Storage: SSD storage

Generate

**CYBERTEC PostgreSQL Configurator**

Select your version of PostgreSQL: 17

GB of RAM in your server: 64

Number of CPUs (= cores): 9

Disk Type: SSD

Number of disks: 5

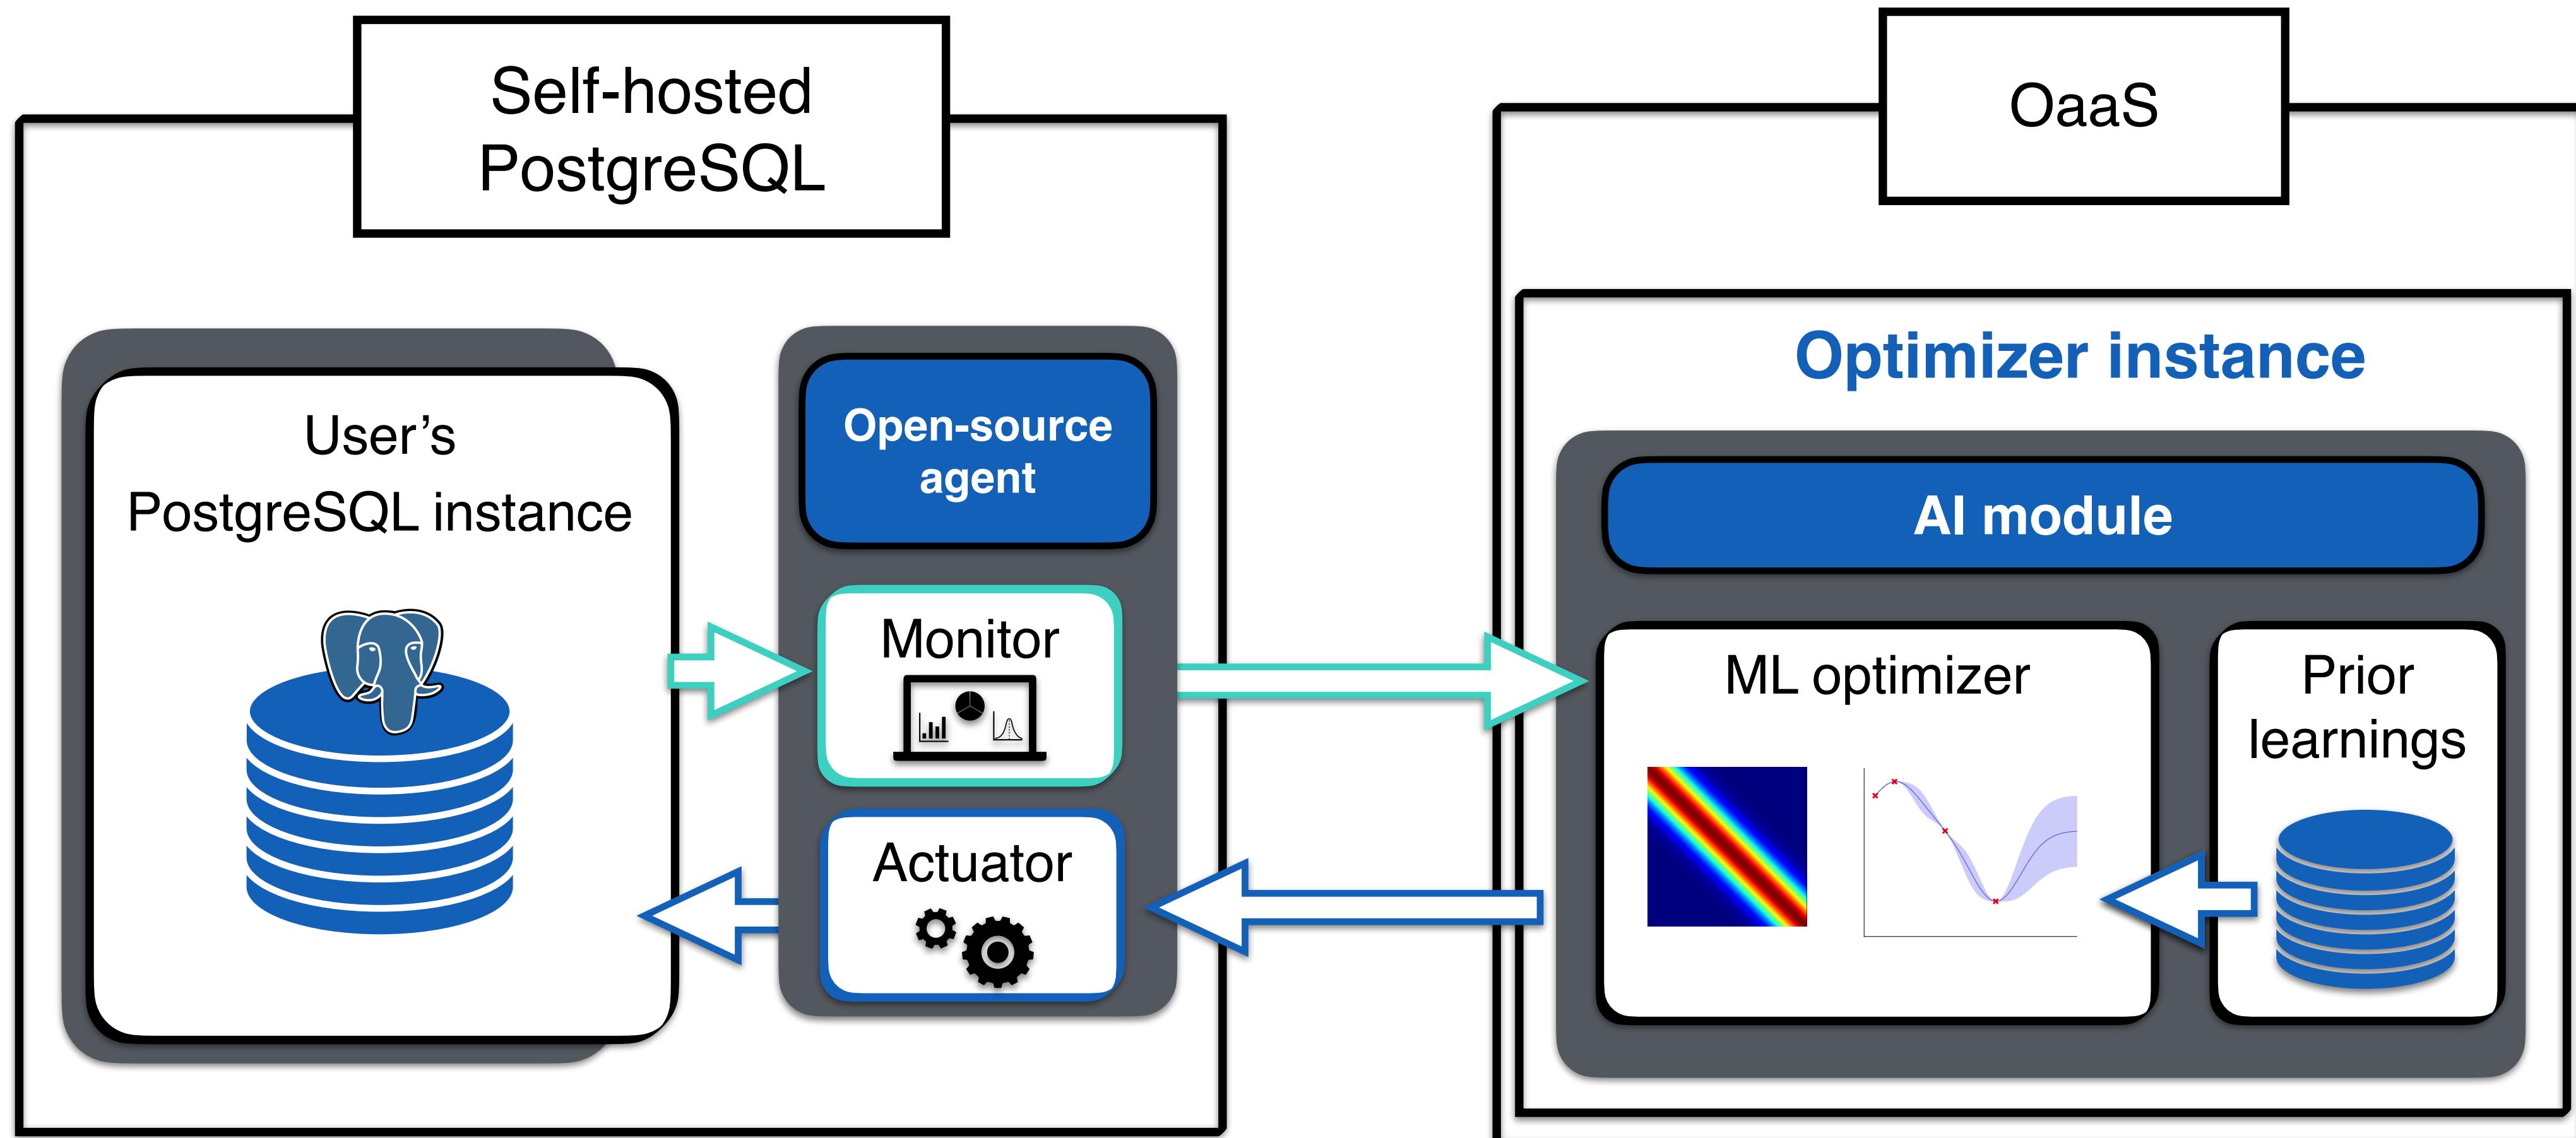
How big is your database? 1TB

How would you describe your workload? Mostly simple short transactions (OLTP)

How many concurrent open connections do you expect? 3020

In this presentation we will use the DBtune free-tier to automate tuning

High-level architecture view for the PostgreSQL Optimizer-as-a-Service (OaaS)



# How often do you tune?

## On-going requirements

-  Scale your cloud instance up or down
-  Change in application workload
-  Change in data distribution and query planner

## Specific moments

-  Migrate from Oracle to PostgreSQL
-  Migrate from on-prem to the cloud – Or vice-versa
-  Migrate PostgreSQL version
-  Experience downtimes

# The reality of how most enterprises treat manual parameter tuning today

- ✓ Tuning is typically **reactive** to something going wrong — Not **proactive**
- ✓ Often engage expensive external resources / experts
- ✓ Different workloads are not treated differently
- ✓ Modus operandi: Throw more hardware / compute at any issue (\$\$\$)

# Example of PostgreSQL parameters tuned in this presentation

## Database reload (15 params)

- ✓ *work\_mem*
- ✓ *max\_parallel\_workers*
- ✓ *max\_parallel\_workers\_per\_gather*
- ✓ *effective\_io\_concurrency*
- ✓ *bgwriter\_lru\_maxpages*
- ✓ *random\_page\_cost*
- ✓ *sequential\_page\_cost*
- ✓ *bgwriter\_delay*
- ✓ *max\_wal\_size*
- ✓ *min\_wal\_size*
- ✓ *checkpoint\_completion\_target*

## Require database restarts (3 params)

- ✓ *shared\_buffers*
- ✓ *max\_worker\_processes*

There is an on-going *shared\_buffers* patch to make it dynamically adjustable (see hackers' list)

# HammerDB TPC-C benchmark suite

Open-source database load-testing and benchmarking tool

## TPC-C benchmark

Wholesale supplier's order and delivery operations across multiple warehouses:

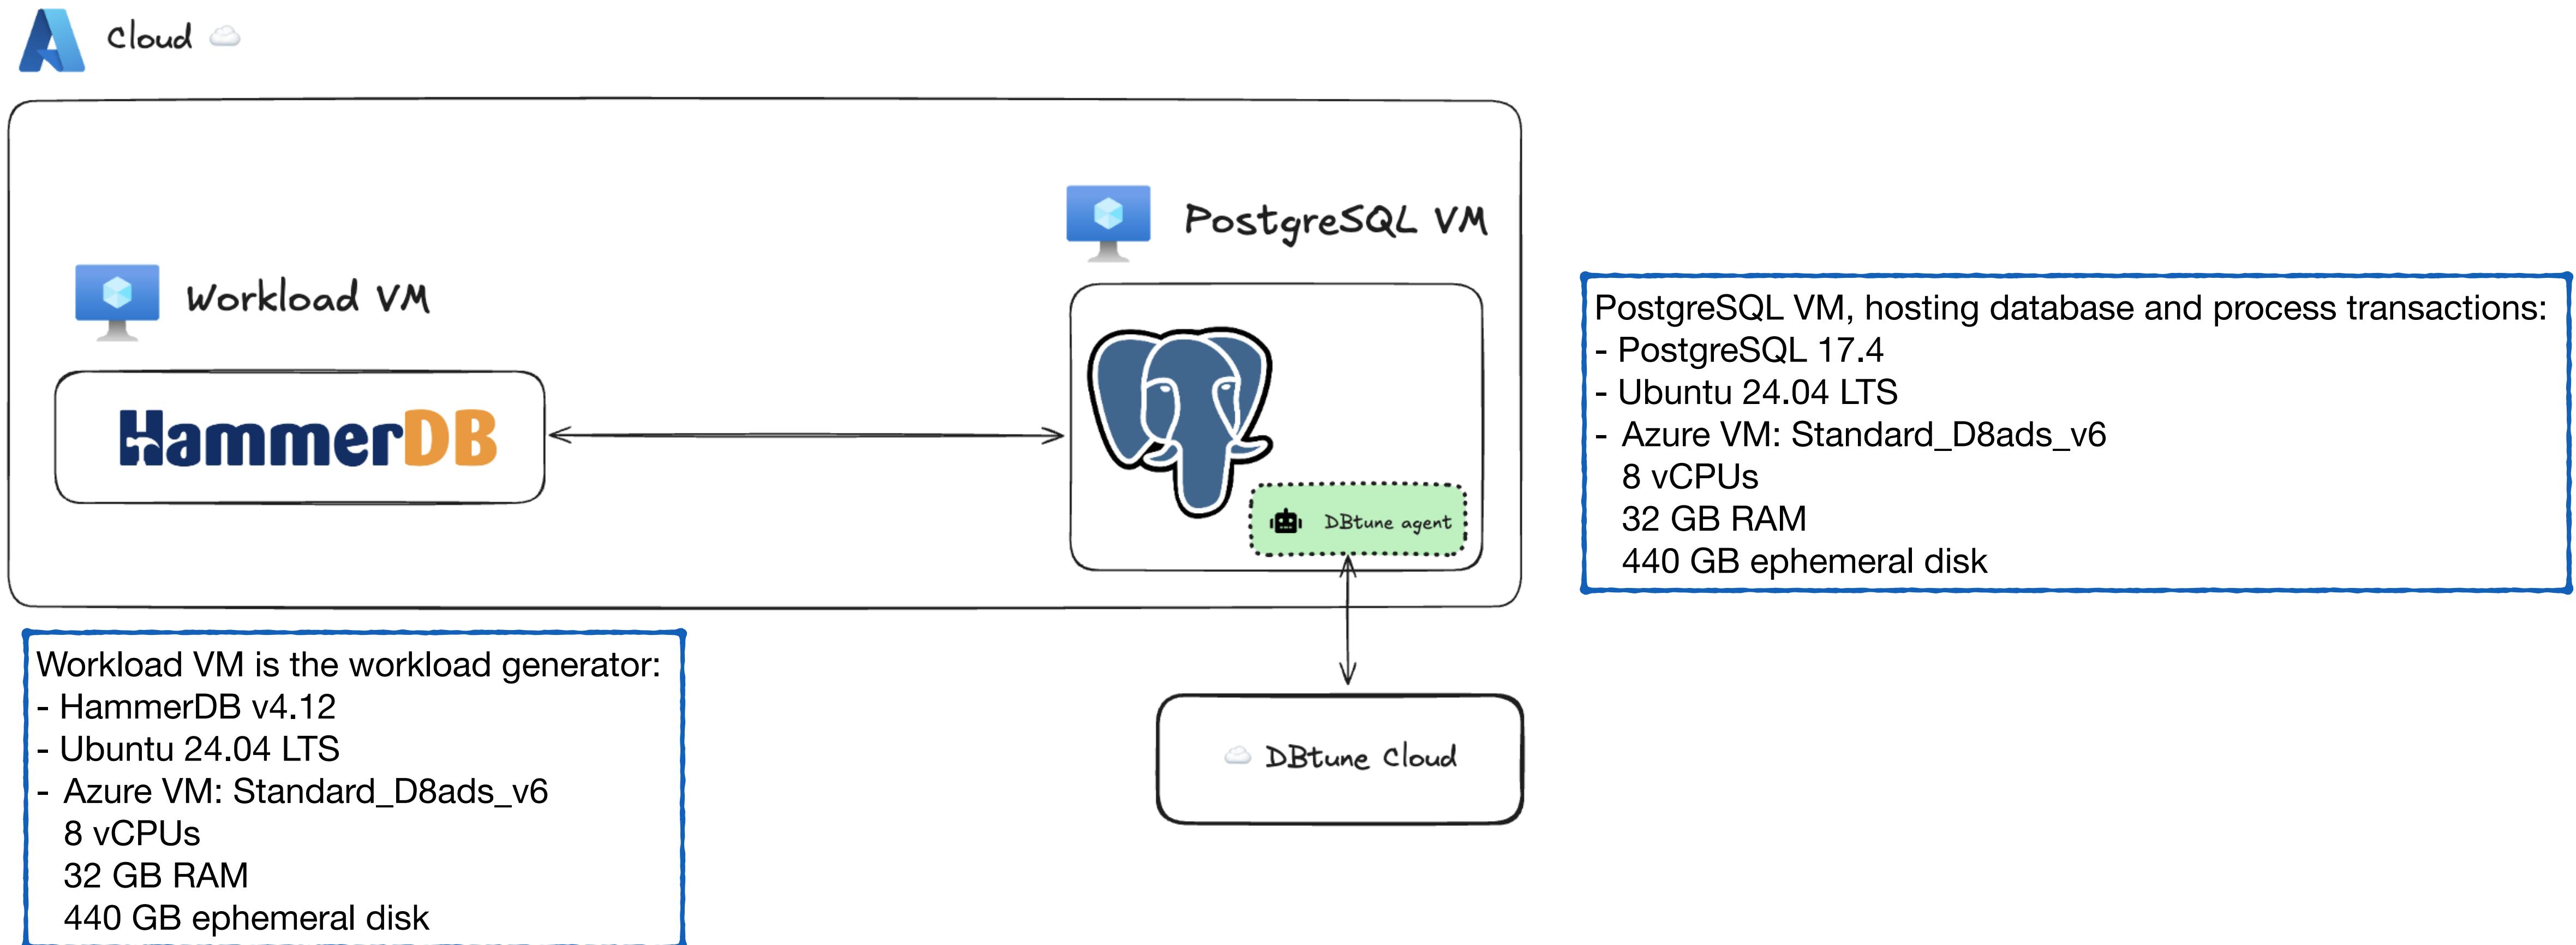
- New order: Creating new customer orders
- Payment: Recording customer payments
- Order status: Querying existing order status
- Delivery: Processing batch deliveries
- Stock level: Checking stock levels

Many users execute transactions concurrently

TPC-C is an OLTP write-heavy workload

# The distributed environment for running PostgreSQL benchmarks

Best practice — Two-node setup one for benchmark runner and one for PG:  
Eliminates contention between the workload generator from the db server



# HammerDB tips and tricks

## Extending run duration: From 5 min to 1 day

```
diset('tpcc','pg_duration','1440')
```

- ✓ Establishing baseline performance and warmup caches (next slide)
- ✓ Observe the system's behavior under sustained load
- ✓ Allow sufficient time to complete a tuning session

## Setting number of warehouses:

- ✓ HammerDB recommends 250-500 per CPU socket: 52 GB of data volume

## Setting virtual users:

- ✓ Set virtual users to 285 

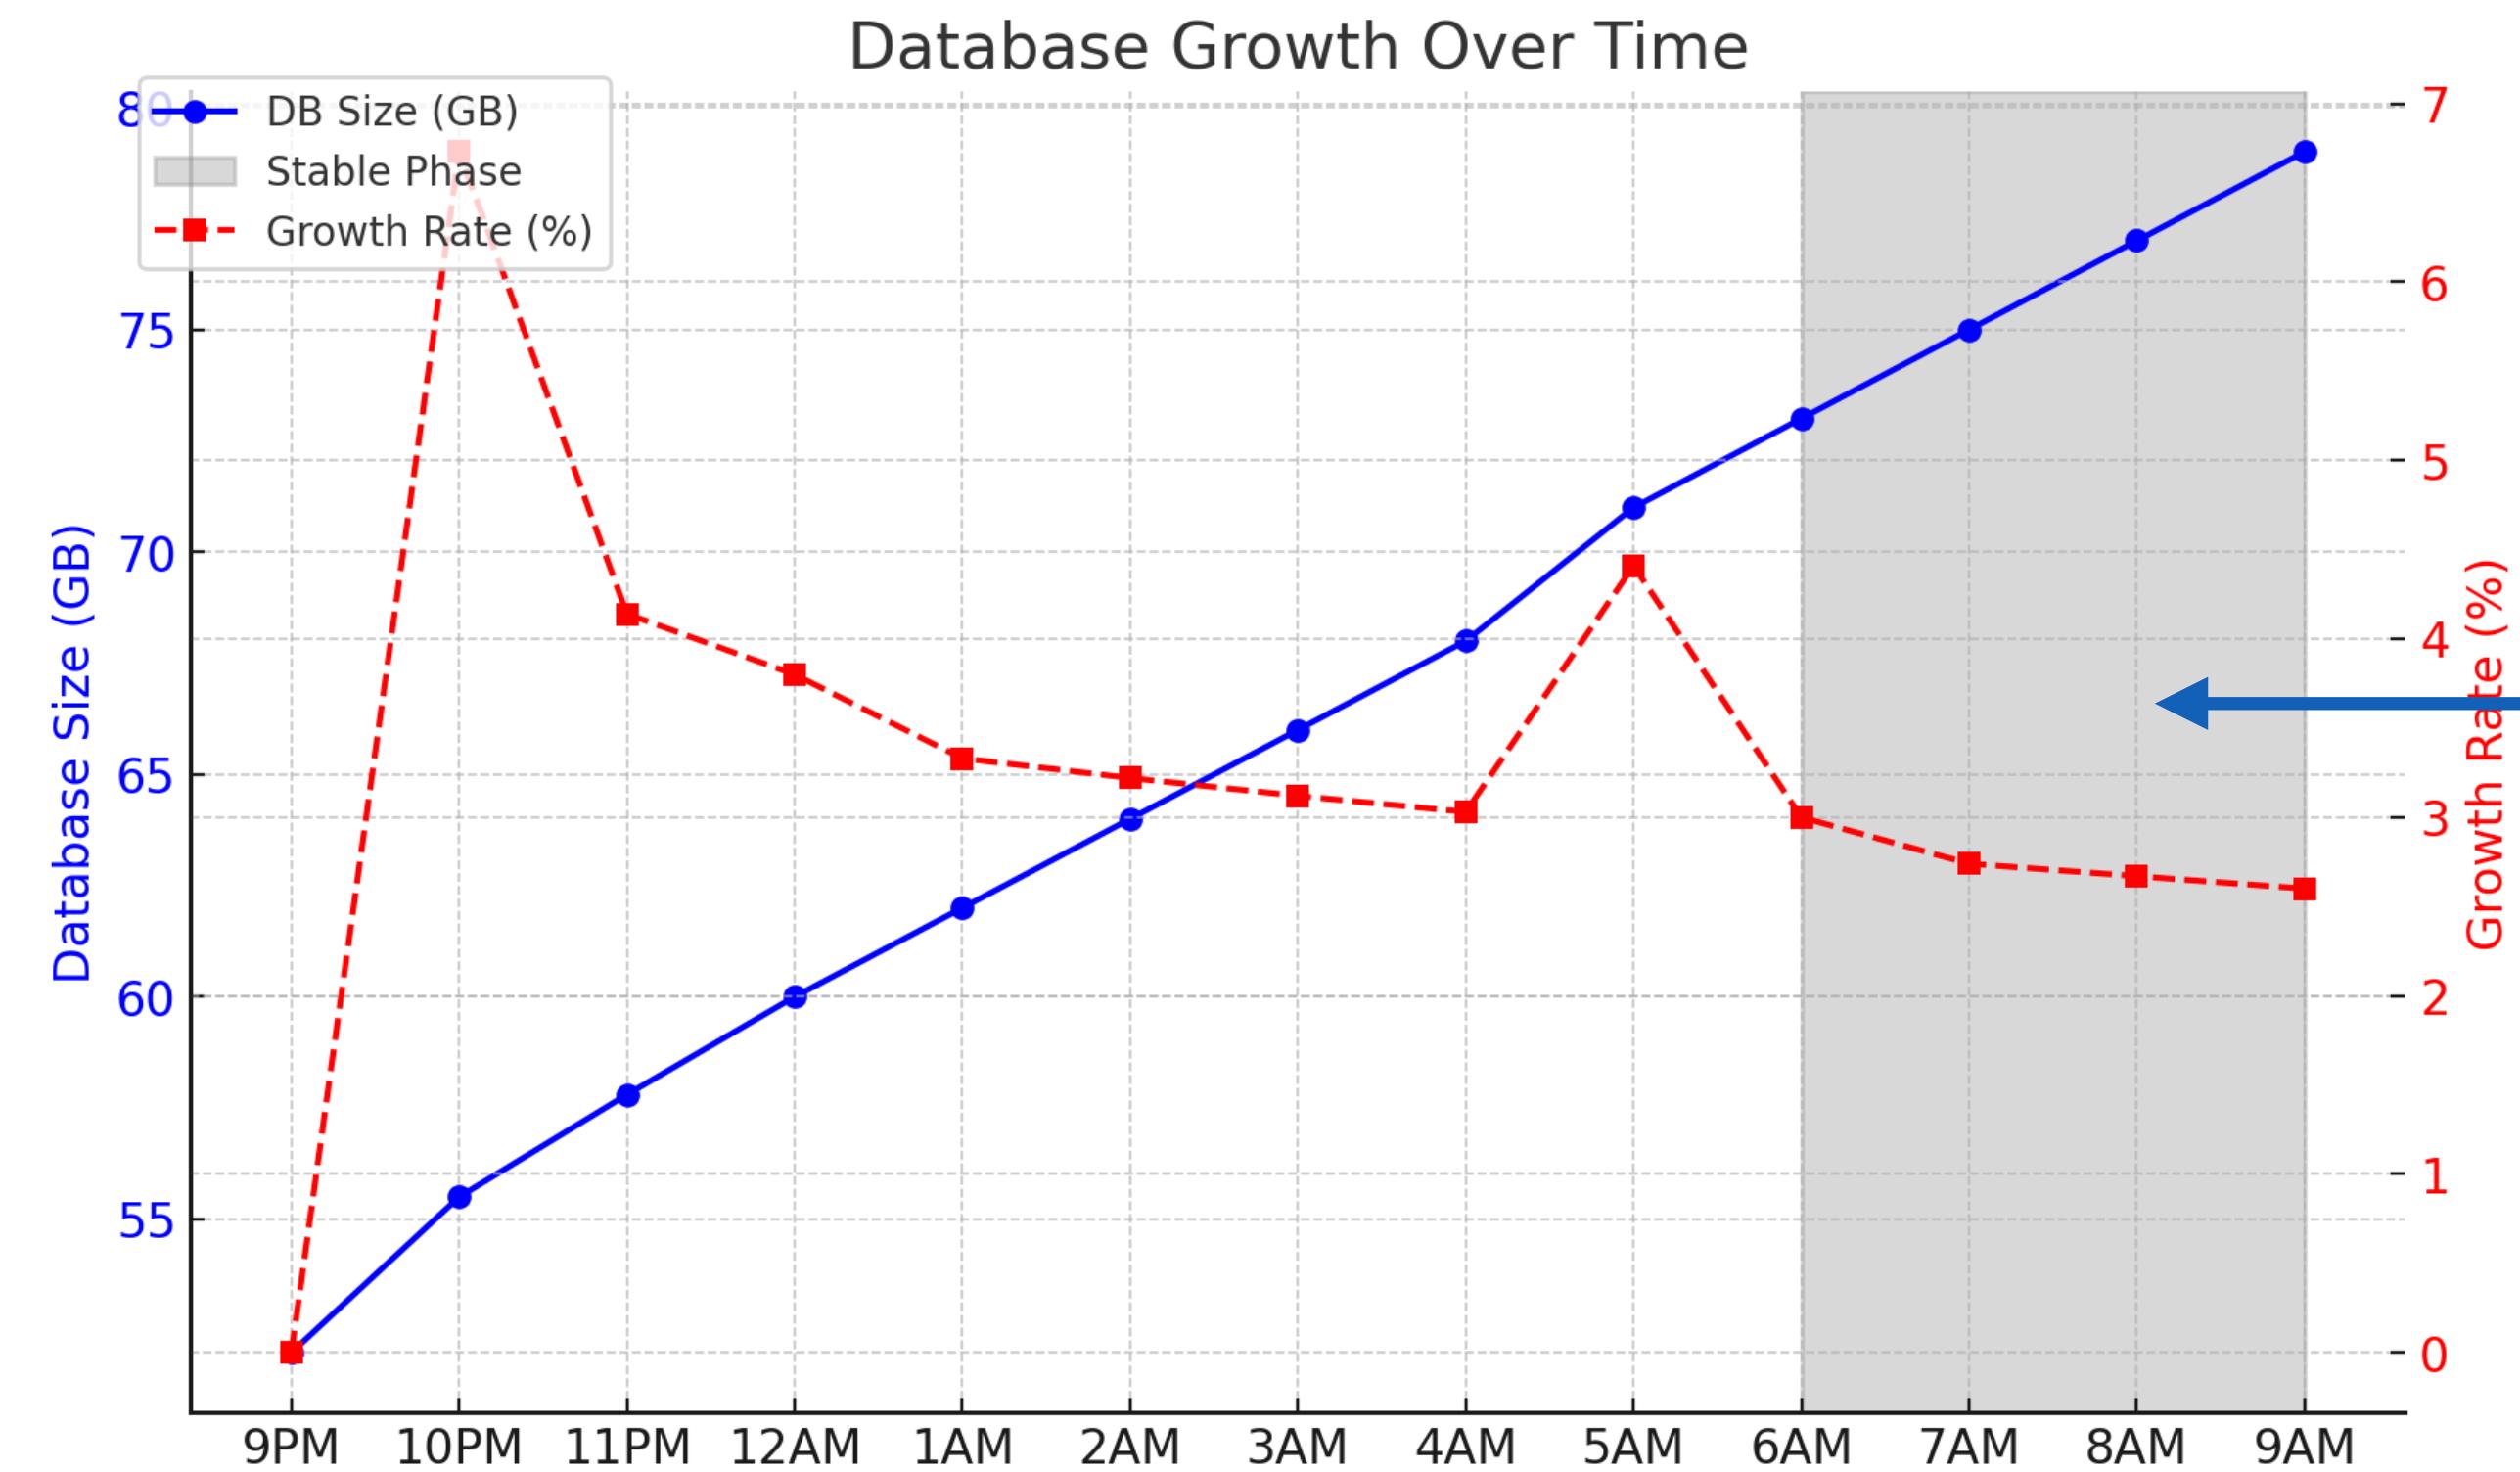
```
vuset('vu',285)
```
- ✓ On 8 vCPUs higher vu creates bottlenecks and a concurrent workload

## Setting number of connections:

- ✓ `max_connections` from 100 to 300

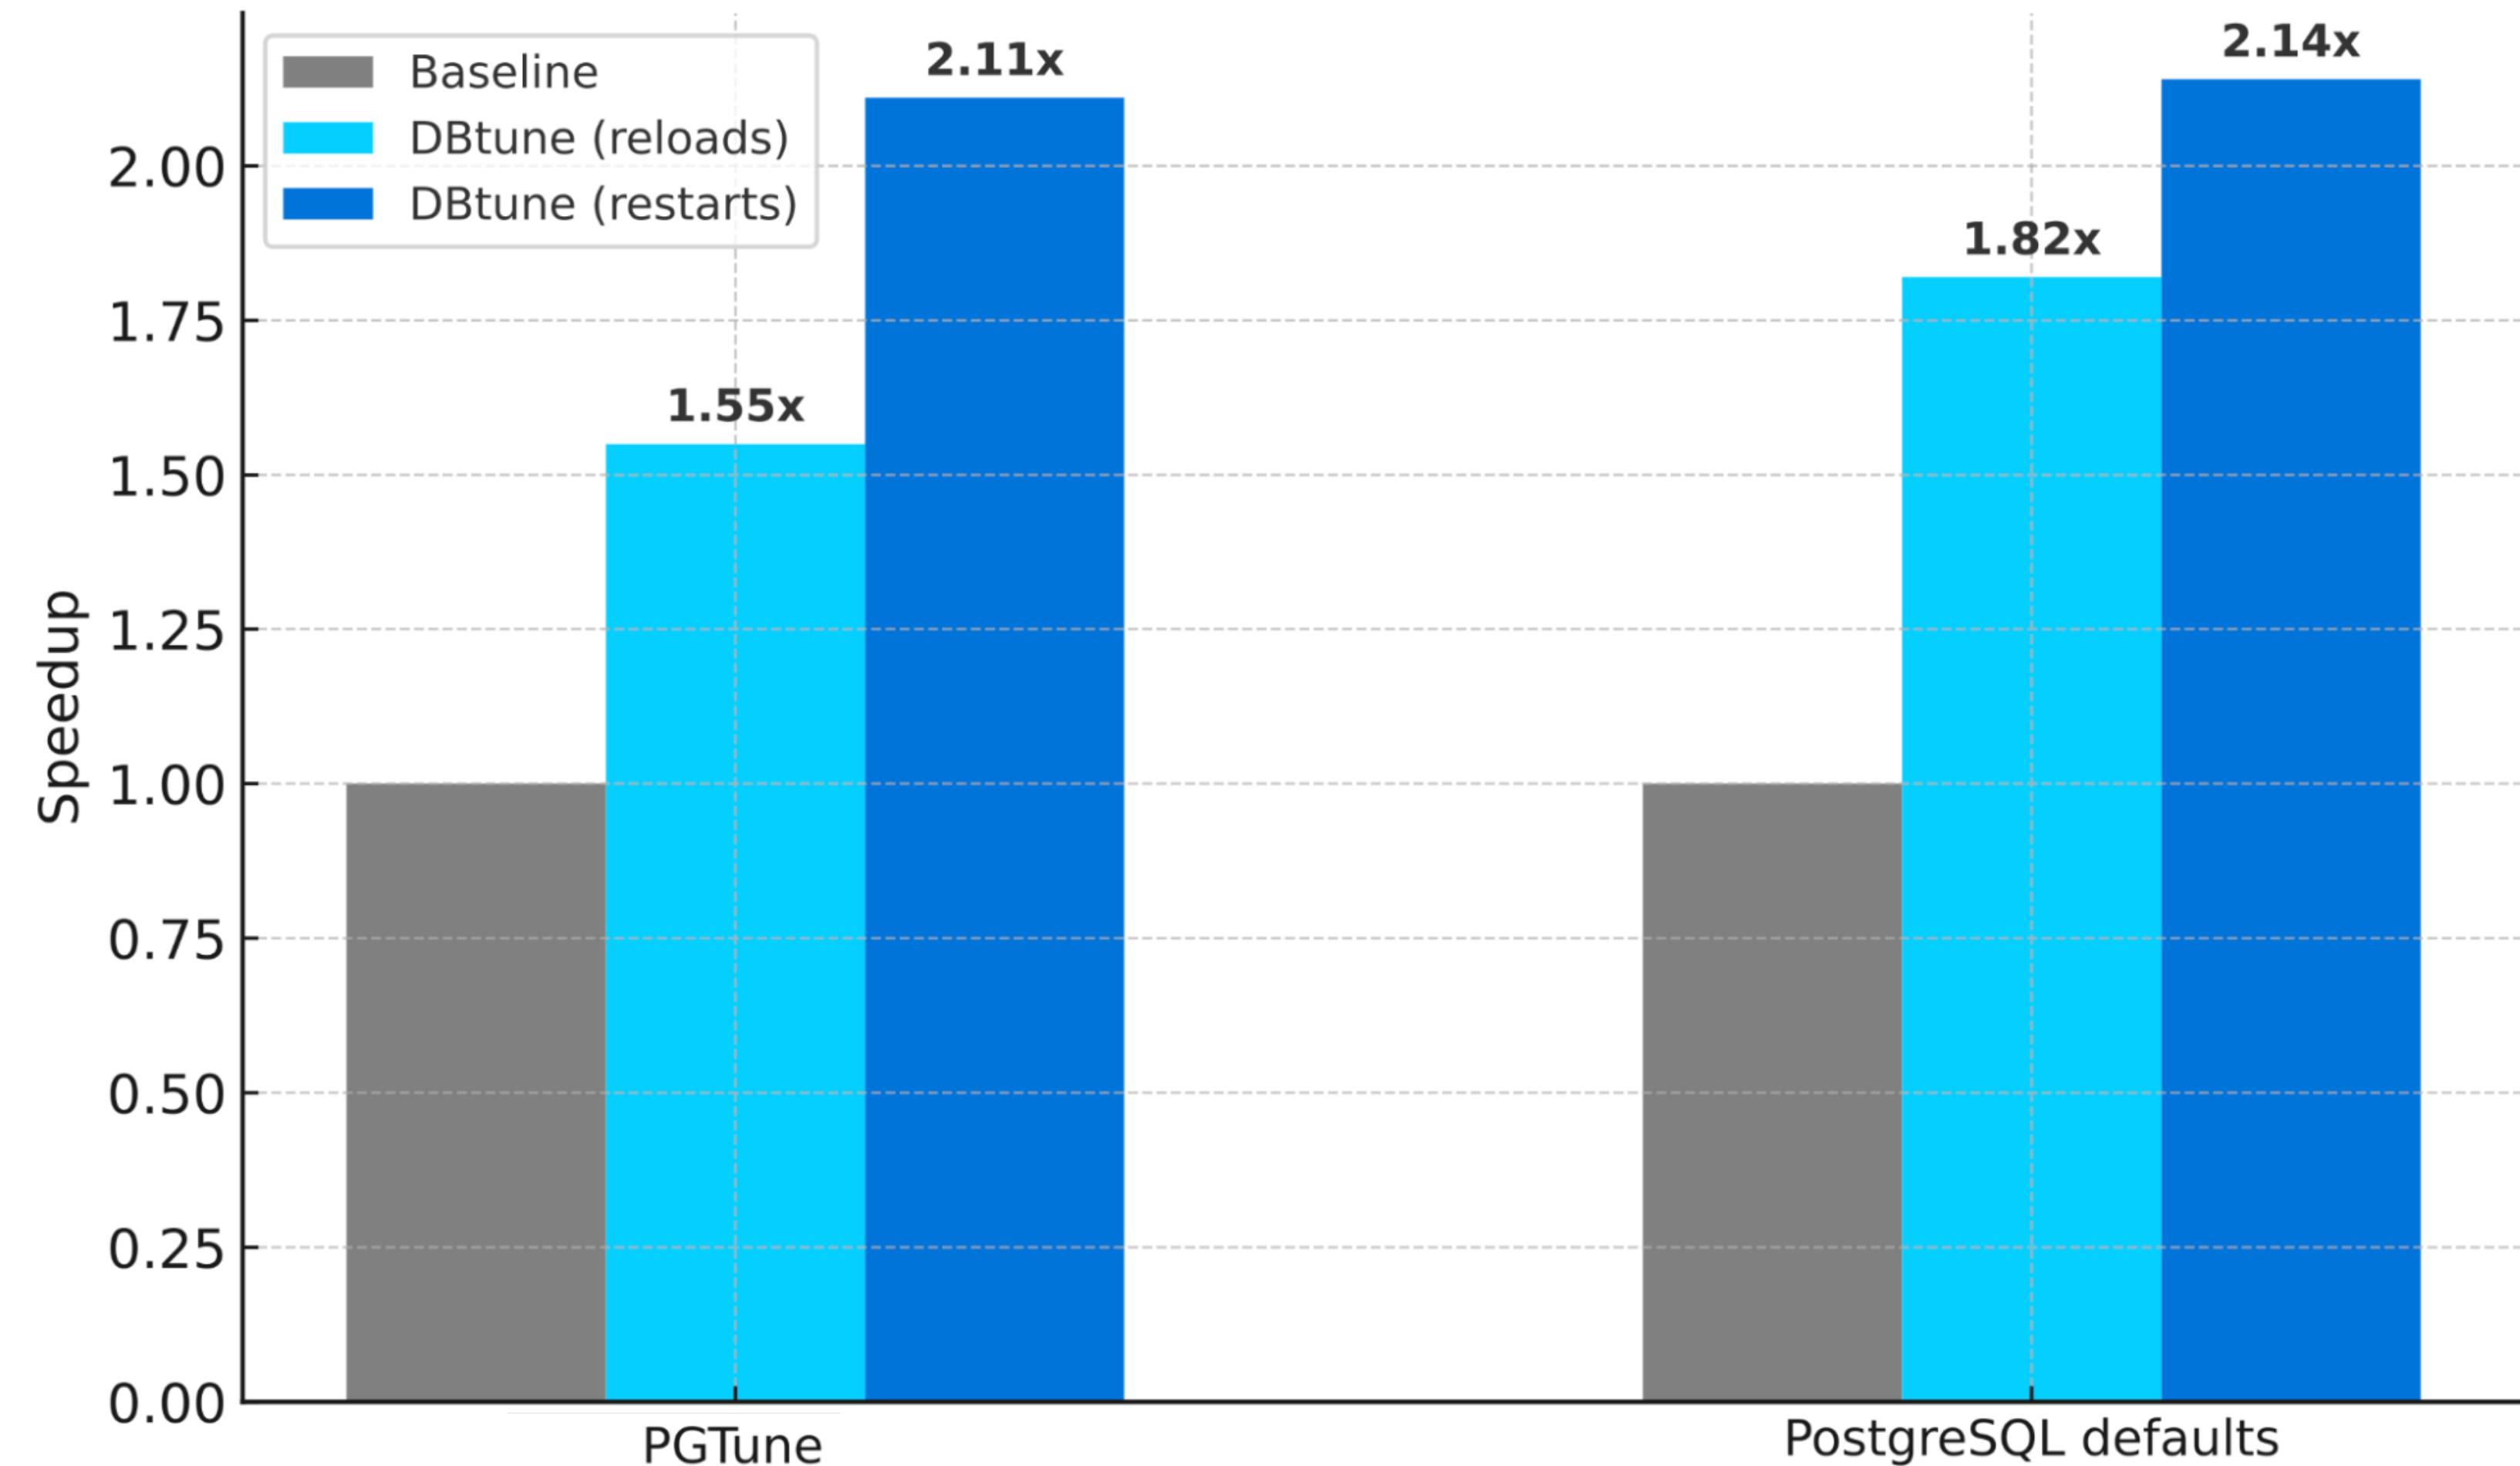
# Establishing a realistic HammerDB TPCC baseline performance

HammerDB performance metrics fluctuate significantly during the initial hours of operation as the system builds up buffer caches, statistics, and query plans



The grey area is closer to a steady state and mimics the real-world more closely

# HammerDB TPCC tuning speedup results on TPS\*



\*The AQR of PG default tuning led to a 4x speedup in the restart scenario

Blog: <https://www.dbtune.com/blog/dbtune-and-hammerdb> by Mohsin Ejaz (DBtune) and Steve Shaw (HammerDB)

# BenchBase benchmarks

## Open-source database benchmark suite

<https://github.com/cmu-db/benchbase>

### ResourceStresser benchmark

Synthetic benchmark that creates isolated contention on system resources:

- The transactions impose some load on 3 specific resources: CPU, I/O, locks

### TPC-H benchmark

Decision support system benchmark:

- Business oriented ad-hoc queries and concurrent data modifications
- Examine large volumes of data, execute high-complexity queries, and give answers to critical business questions

### Epinions.com benchmark

Consumer review website, content management and social media apps:

- Models cross-user interactions like writing product reviews
- 9 transactions: 4 user records, 4 item records, and 1 affecting all tables
- The workload is mixed but primarily read-heavy

### SEATS benchmark

Simulates an airline ticketing system:

- Complex flight searches, reservations, and customer management
- Many concurrent users with a mixed workload and a 60/40 read/write ratio

# BenchBase tips and tricks

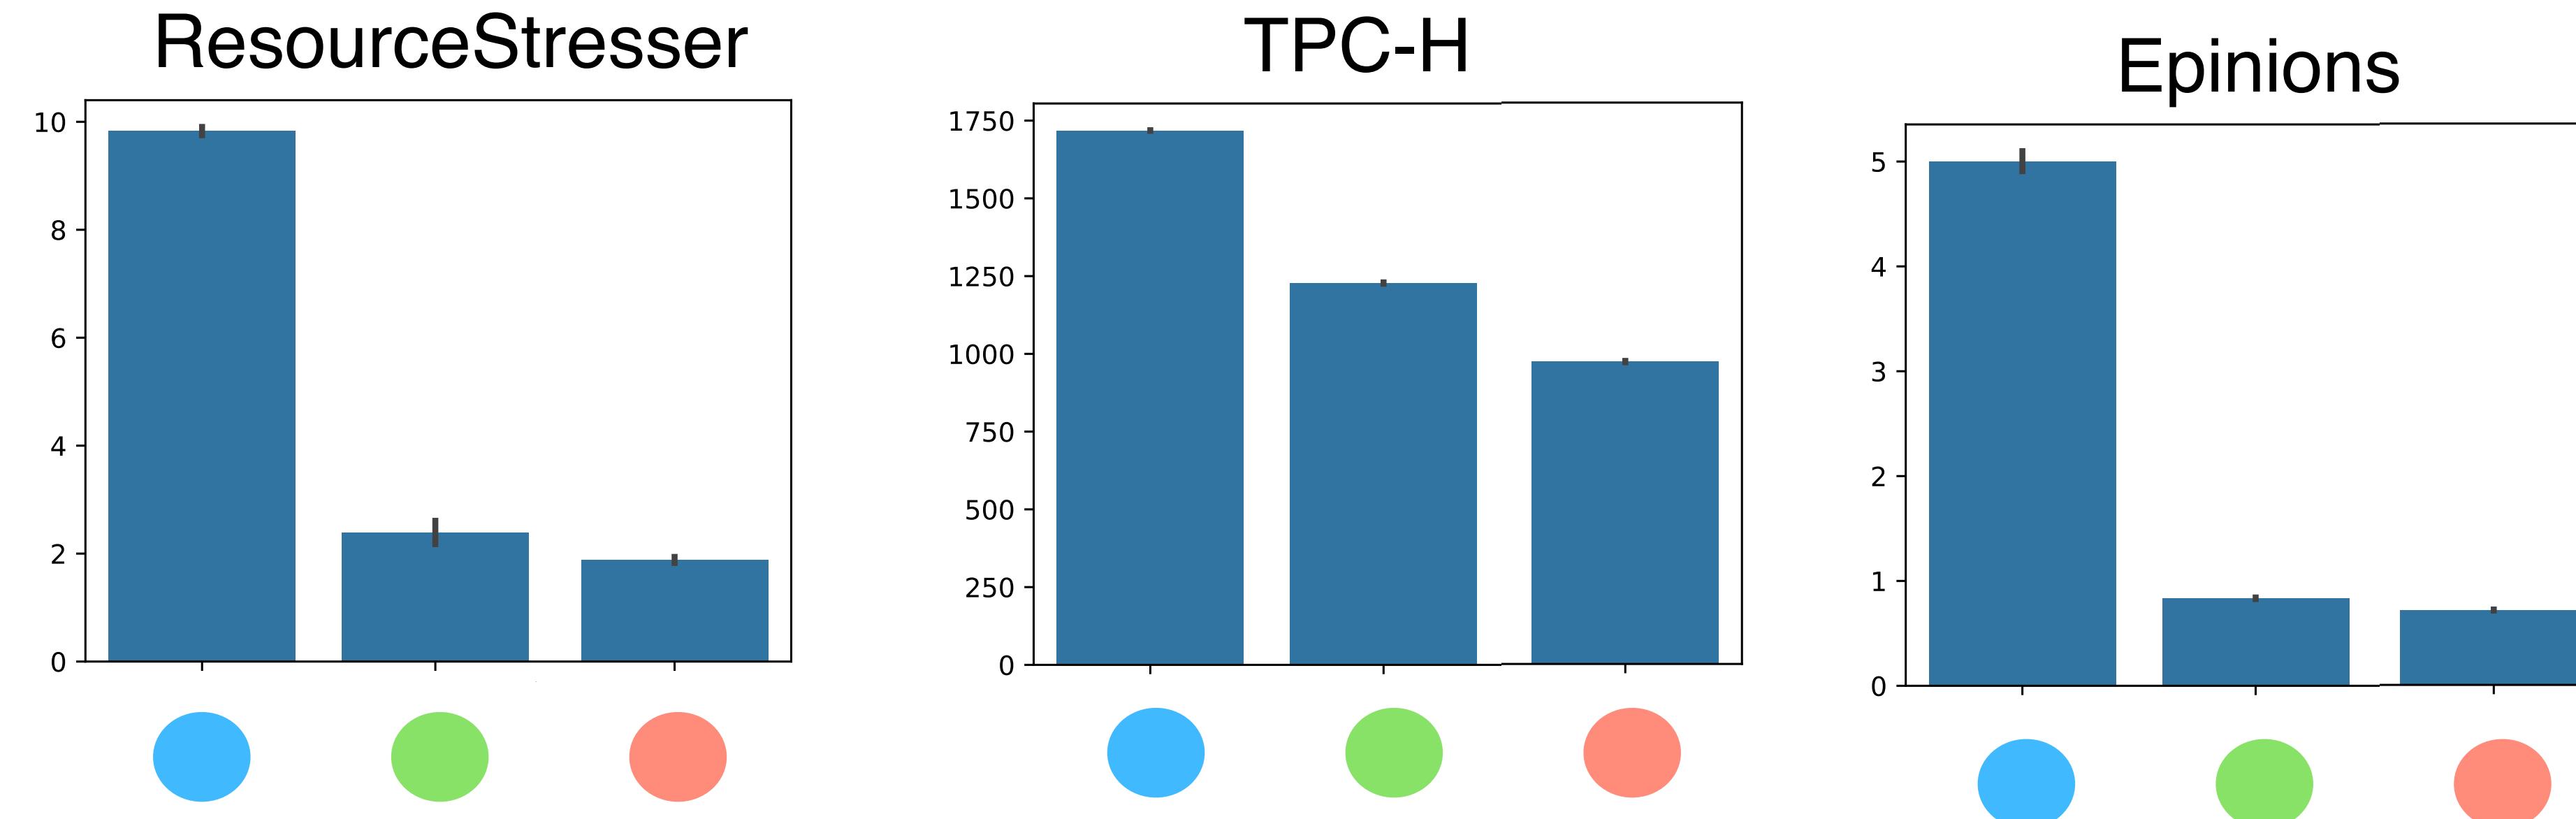
BenchBase benchmarks are configurable:

- **scalefactor**: Sets the size of the database
- **terminals**: Sets the number of connections querying the database
- **time**: Duration time of the benchmark in seconds
- **rate**: Transaction injection rate, determines how many TPS attempts to execute
  - Higher value makes for more intense workloads, commonly set to *unlimited*
- **weights**: Relative frequency of the different query types
- **Etc.**

Example on Epinions: scalefactor=10000 (creates an 80 GB database), terminals to 10, time to 86400 (one day), rate=*unlimited*, and keep the weights as is

# Performance downside of non-restart (reload-only) strategy

Average query runtime (in ms) on a community PG instance running on EC2



- Default PostgreSQL configuration & no tuning
- shared\_buffers=25% & reload-only & tuning*
- Restarts allowed & tuning

Query runtime in ms  
**Lower the better**

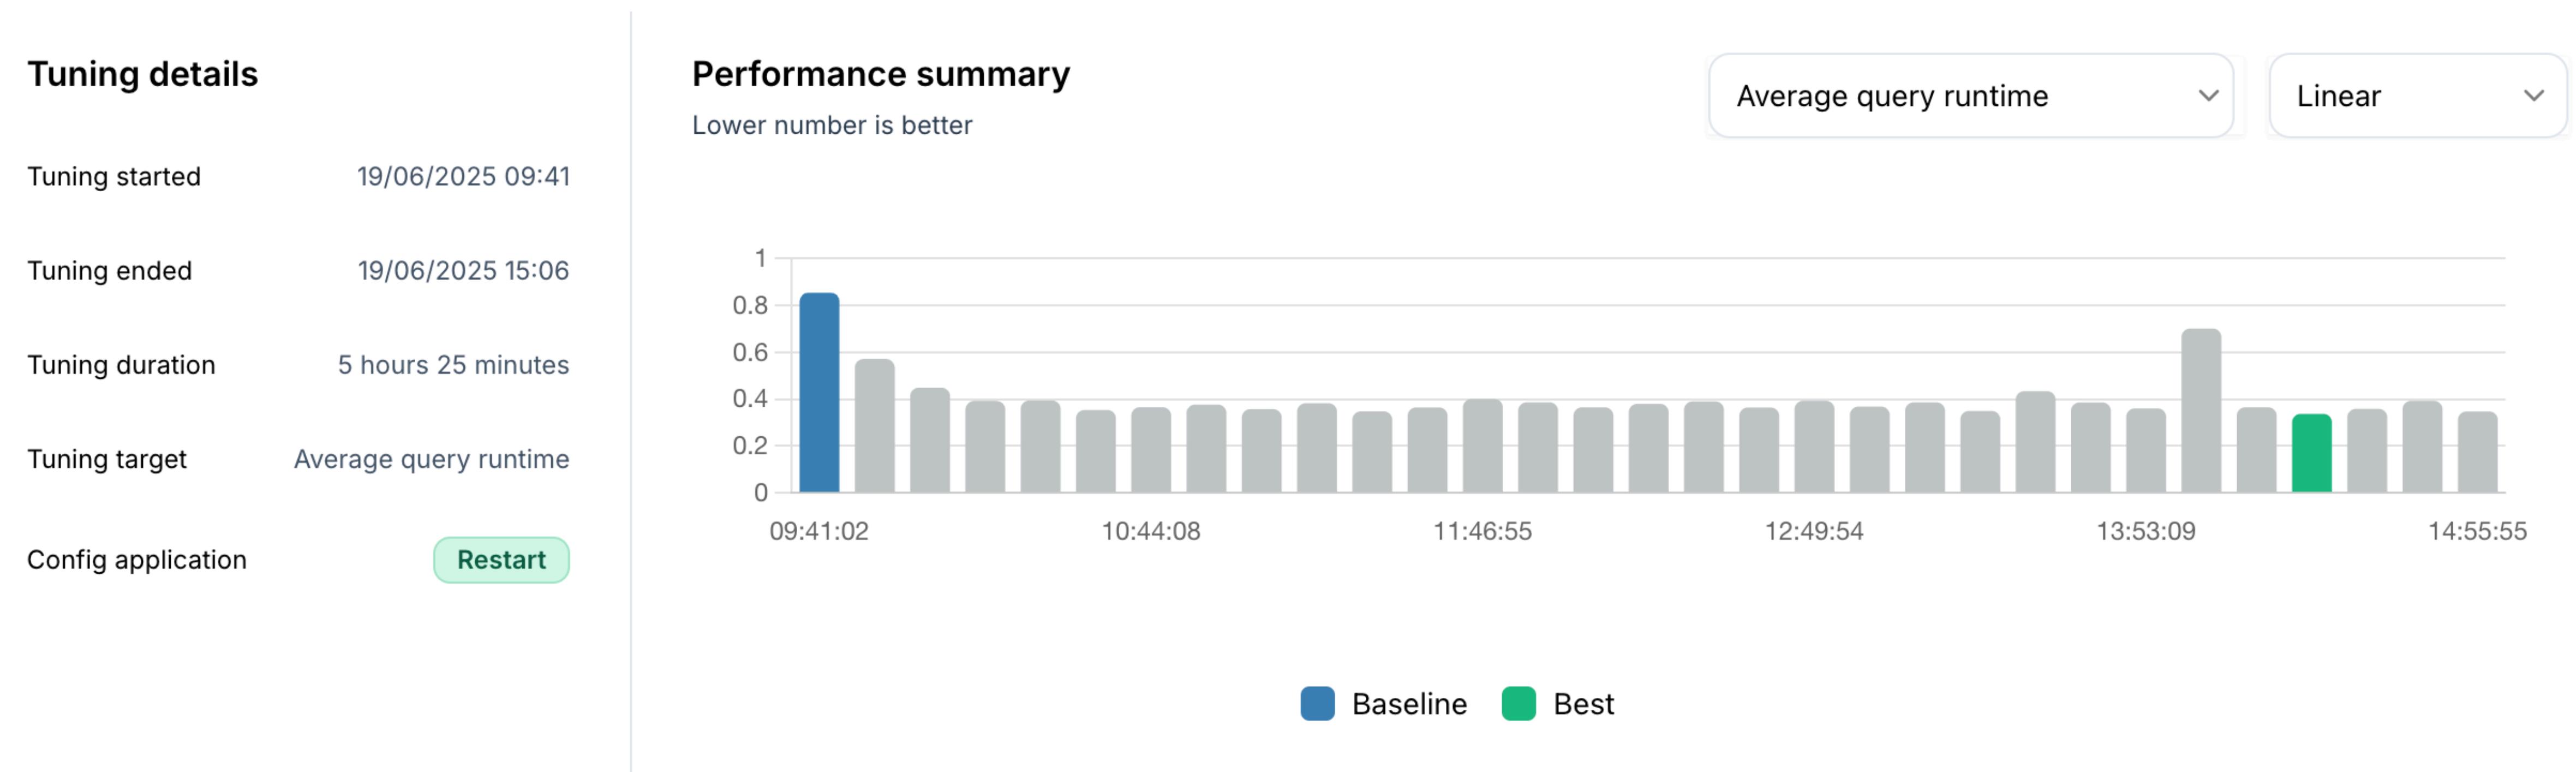


# Insurance application use case study – Customer anonymized data

Environment: 16 vCPU, 32 GB RAM, on-prem, primary instance, PG 15

Manually tuned baseline by expert DBA

Automated tuning with DBtune

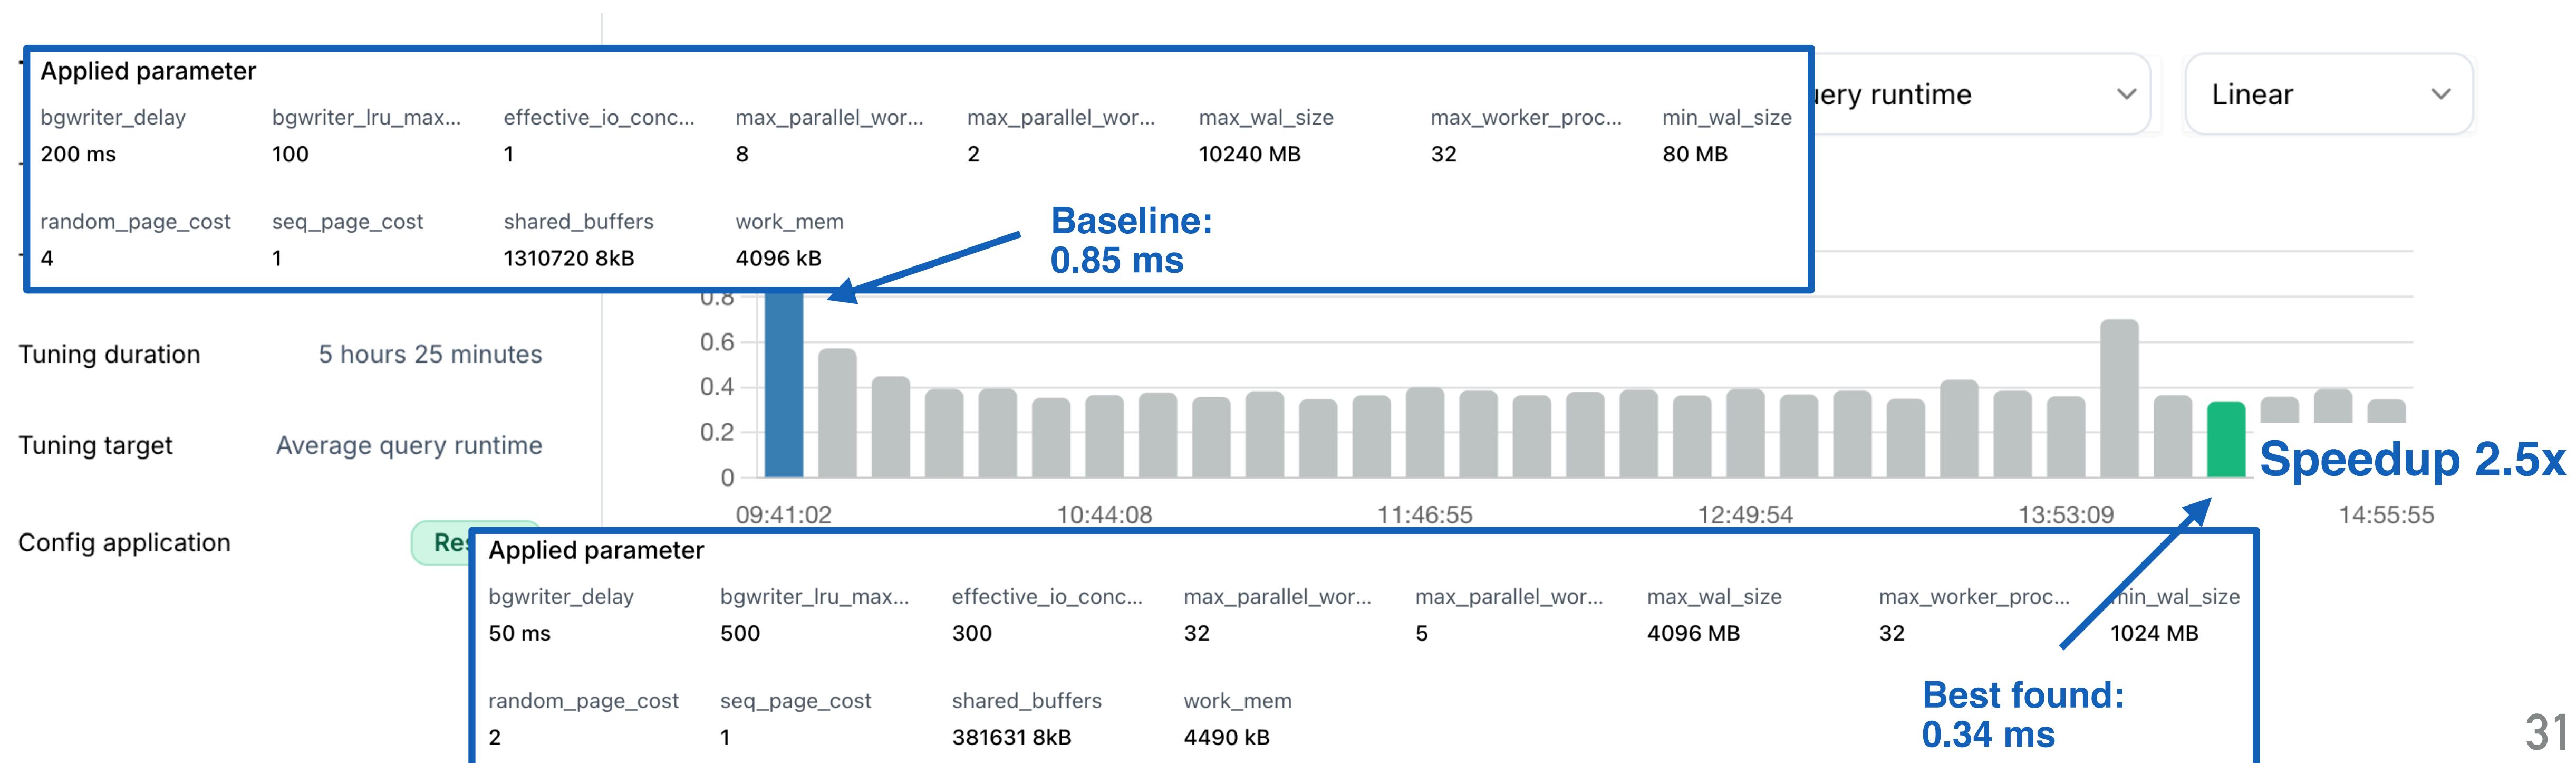


# Insurance application use case study – Customer anonymized data

Environment: 16 vCPU, 32 GB RAM, on-prem, primary instance, PG 15

Manually tuned baseline by expert DBA

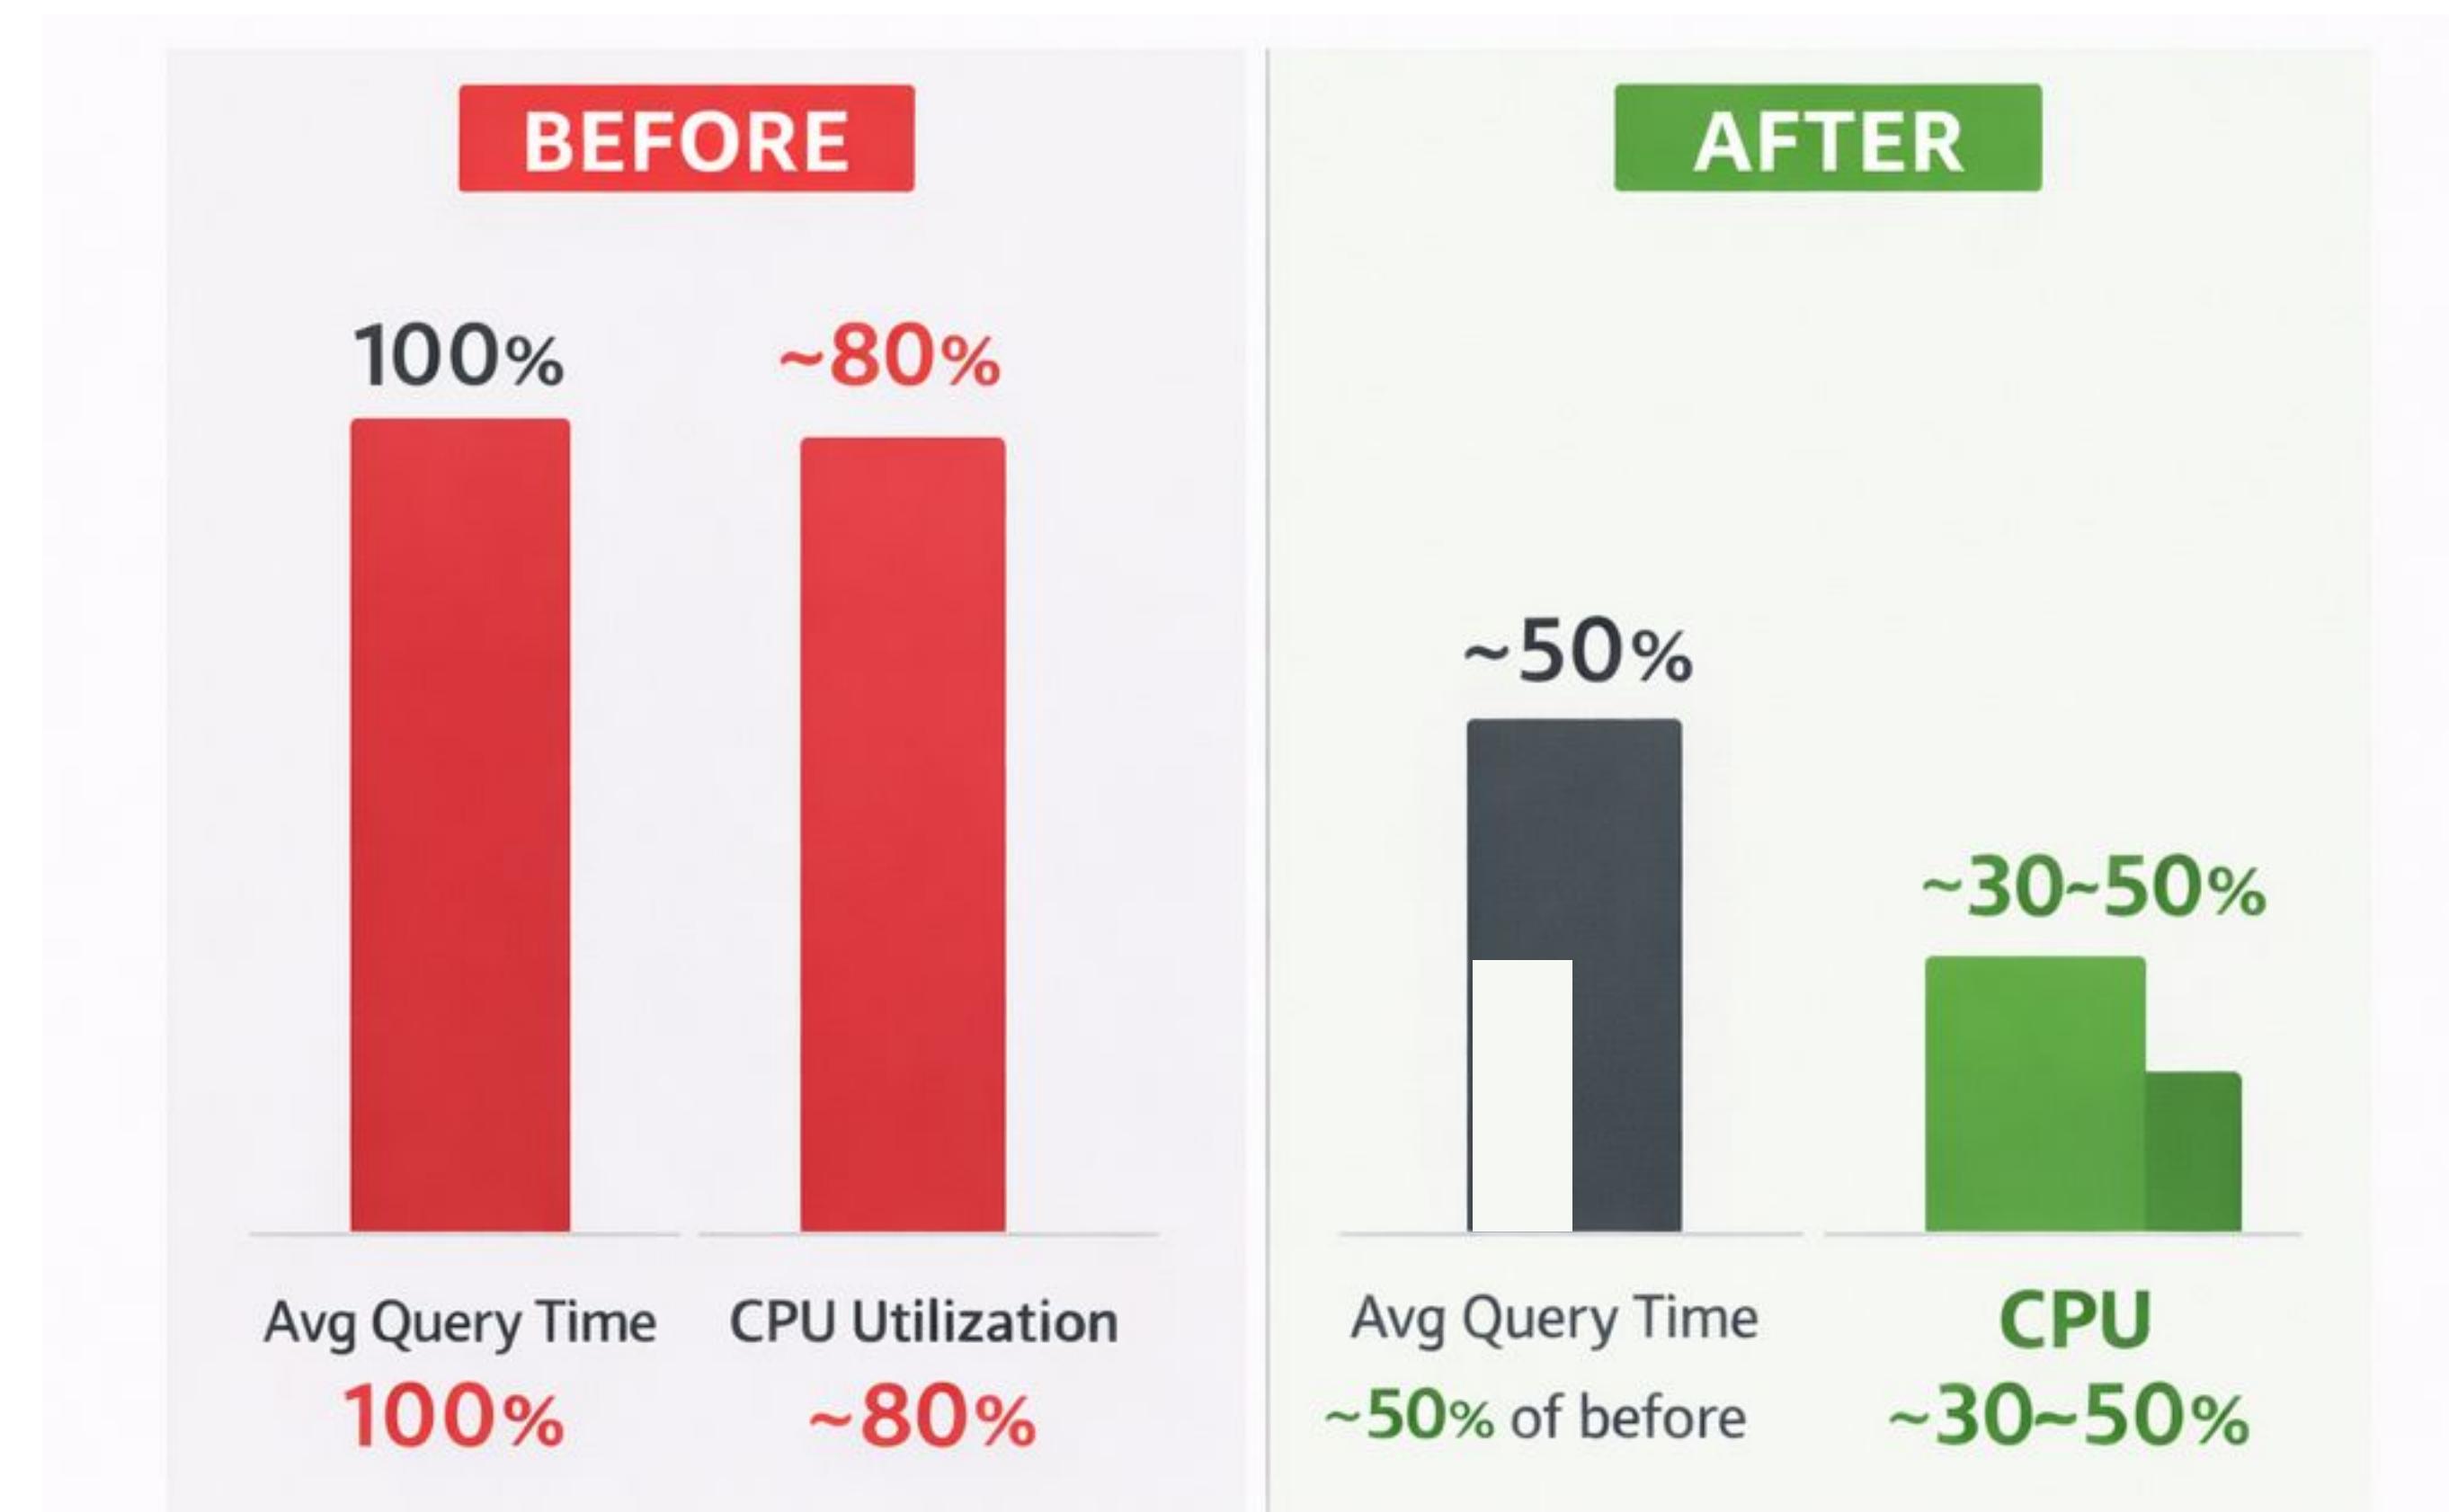
Automated tuning with DBtune



# In production: Workforce management platform by Papershift

Environment: Amazon RDS m5.8xlarge, 32 vCPU, 128 GB RAM, PG 17.6

Baseline by RDS, automated tuning with DBtune



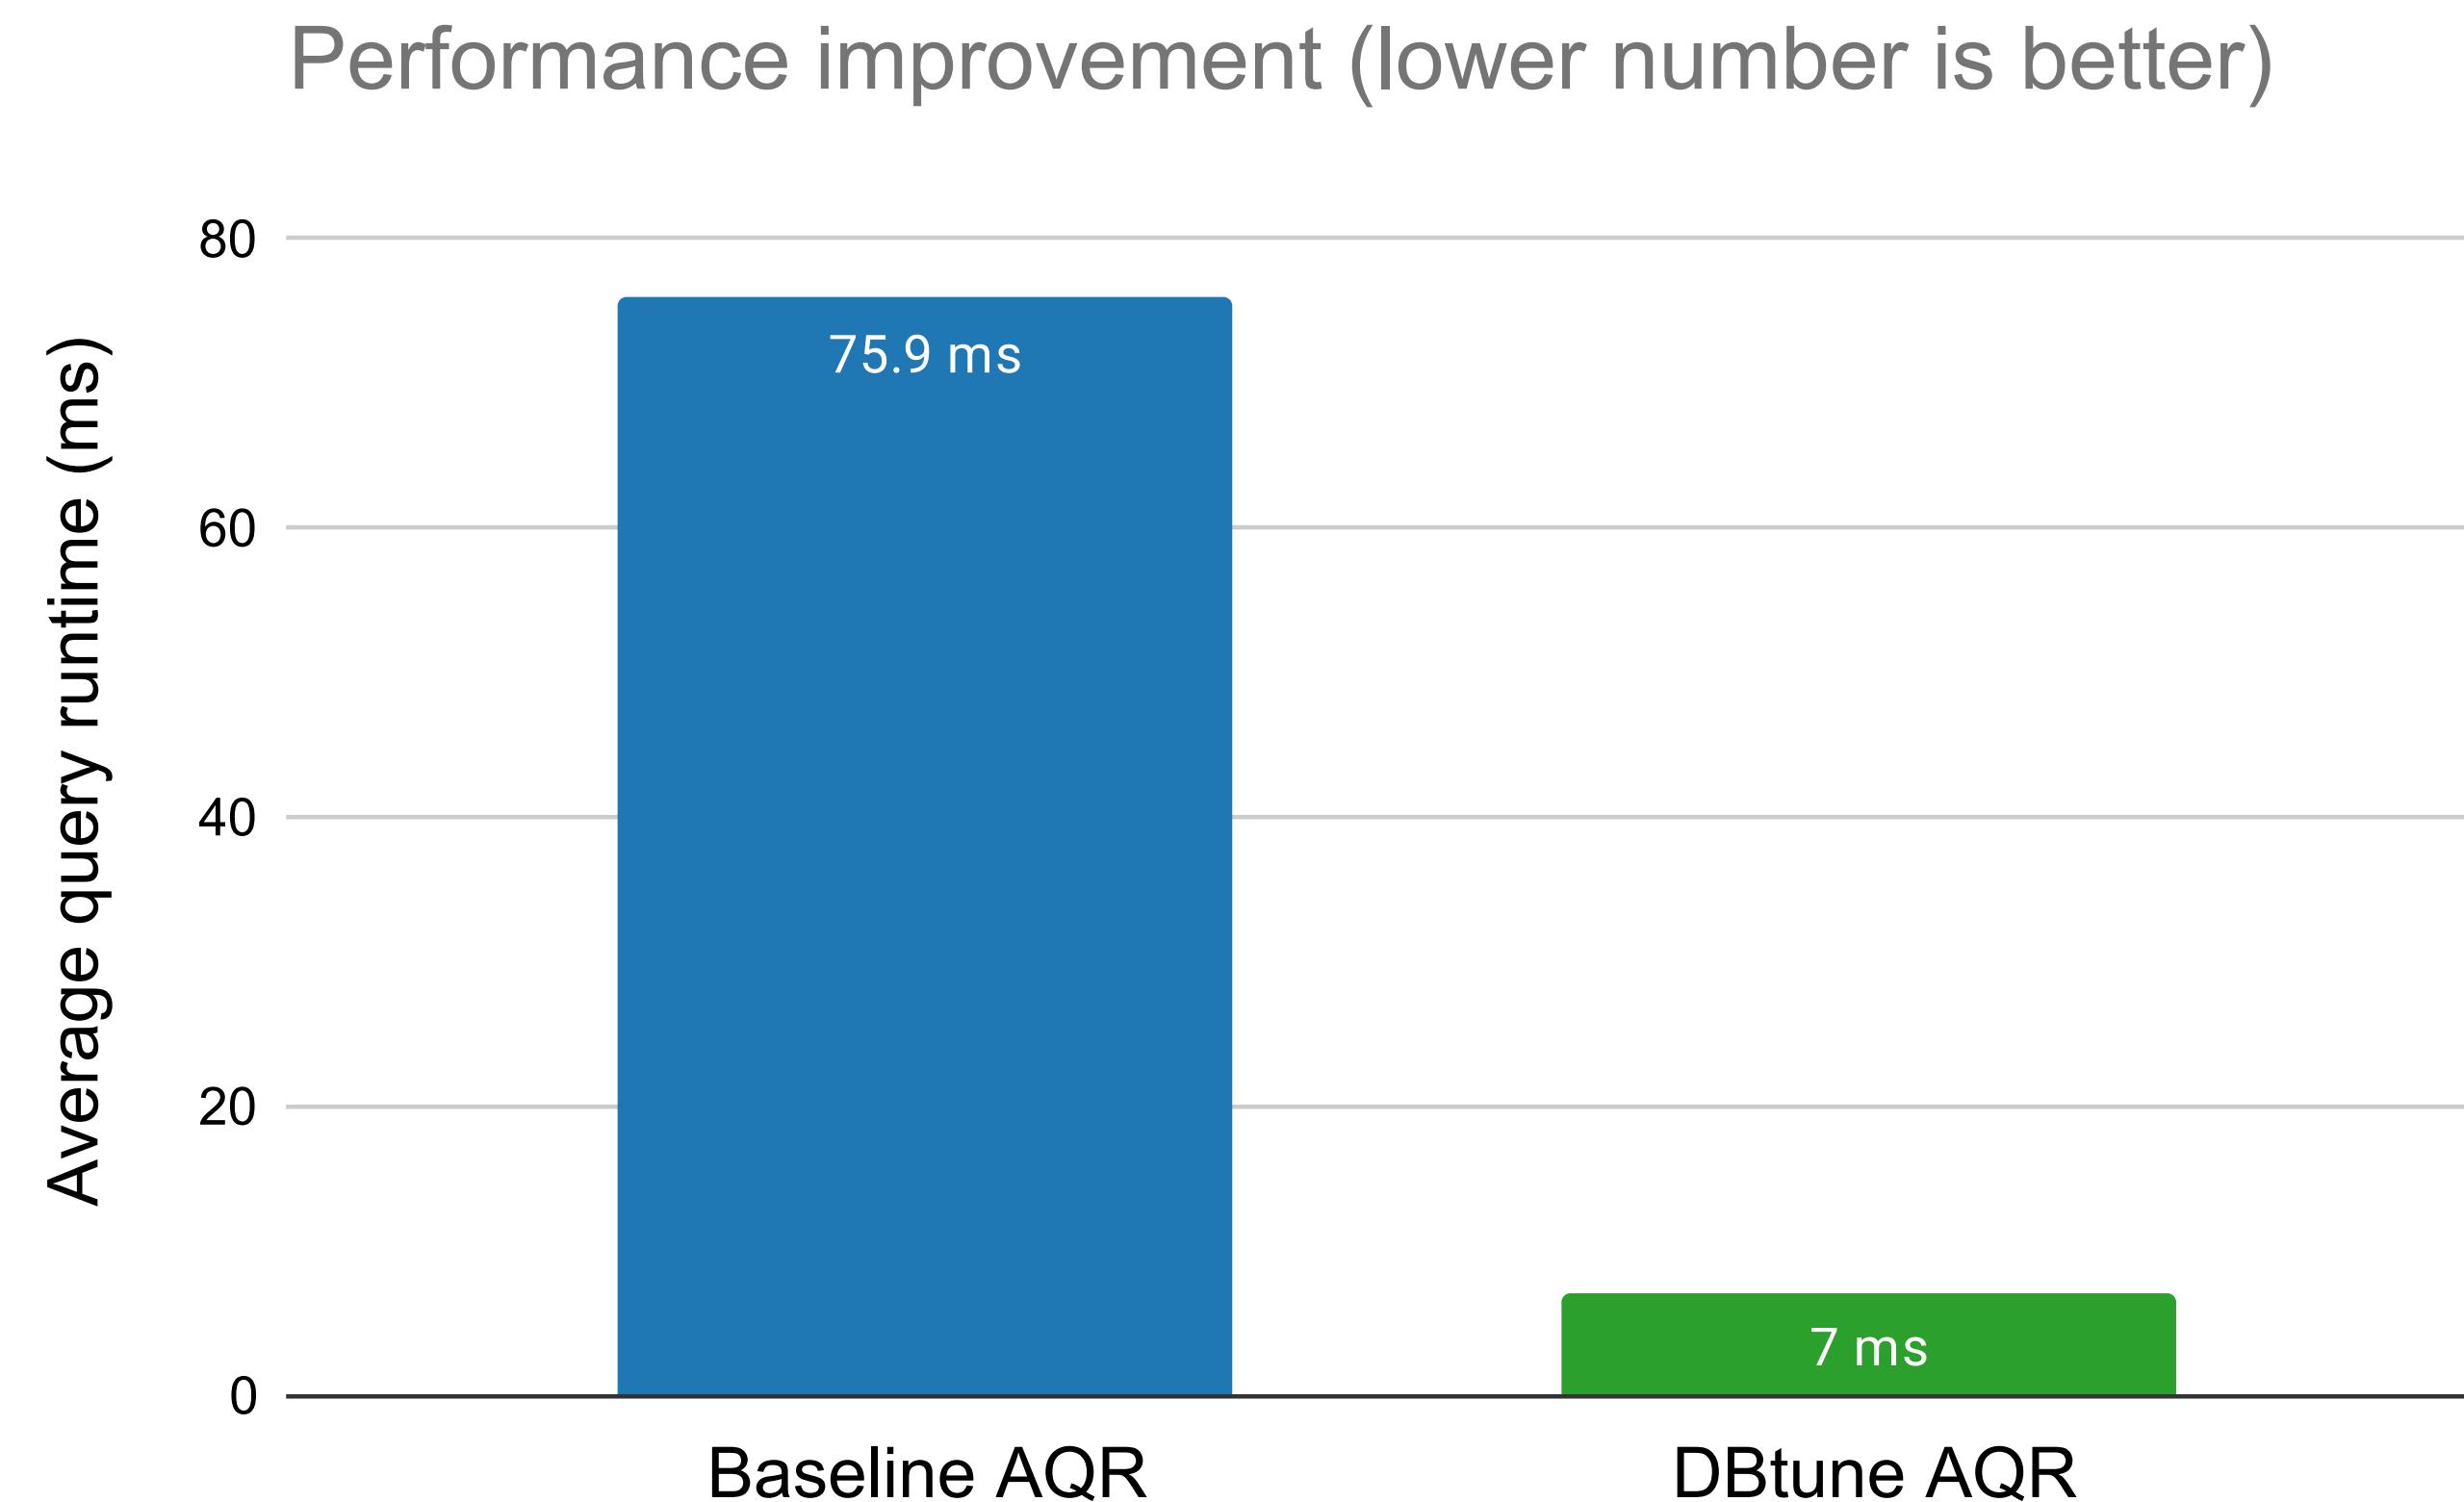
Blog: <https://dev.to/floriansuchan/how-we-used-dbtune-to-cut-our-postgres-query-time-by-50-on-aws-rds-2a5e>

Figure: [https://www.linkedin.com/posts/vibhork\\_postgresql-autonomouspostgres-autonomousdatabase-activity-7408886284879835137-2i\\_m](https://www.linkedin.com/posts/vibhork_postgresql-autonomouspostgres-autonomousdatabase-activity-7408886284879835137-2i_m)

# In production: Digital content service by Midwest Tape

Environment: Amazon RDS r6g.12xlarge, 48 vCPU, 412 GB RAM, PG 14.17

Baseline by RDS, automated tuning with DBtune



# Conclusions and food for thought



Realistic benchmarking:

- Running a benchmark exposes many variables
- These depend on the SKU and on the benchmark application itself
- Benchmarks without configuring them doesn't lead to realistic scenarios



Tuning PostgreSQL is essential for fair database benchmarking



Ultimately the speedup from server parameter tuning depends on:

- How good is your baseline
- How up to date is your baseline – things change and become untuned



Point of view shifts from tuning as a specific task to a maintenance activity

# Questions and additional resources

- Blog: [DBtune and HammerDB: Your guide to fair PostgreSQL benchmarking](#)
  - Useful links: DBtune synthetic workload tutorial [GitHub](#) and [video](#)



@luinardi

luigi@dbtune.com



DBtune newsletter



dbtune