dbtune

PostgreSQL parameter performance optimization

From manual tuning to auto-tuning

Luigi Nardi, Ph.D.

Founder & CEO, DBtune

About me

 @luinardi

 @nardiluigi

B.Sc and M.Sc. Computer Engineering at La Sapienza — Rome (Italy)

- 2006 M.Sc. thesis at LAAS-CNRS — Toulouse (France)
- 2007 Ph.D. Computer Science at Université Pierre et Marie Curie — Paris (France)
- 2011 Research Engineer at Murex SAS — Paris (France)
- 2014 Postdoc Imperial College London (UK)
- 2017 Research Staff at Stanford University (USA)
- 2019 Assistant Professor in Machine Learning at Lund University (Sweden)
- 2021 Founder & CEO at DBtune — Malmö (Sweden)
- 2024 Associate Professor in Machine Learning at Lund University (Sweden)

Introduction

*DBtune is an **AI-powered** database parameter **tuning** service.*

*Spun out of research at Stanford University, DBtune autonomously optimizes the configuration of databases through **machine learning**.*

*It observes, iterates and adapts until converging and delivering the **optimal** settings for any individual workload, use case and machine.*

Outline

- DBtune intro — Who are we?
- Global challenges — Business priorities, FinOps and GreenOps
- Introduction on database tuning
- Machine learning tuning automation
- Tuning impact analysis
- Conclusions and selected real-world use cases
- Demo
- Q & A

The DBtune team

Leadership

Founder & CEO
Dr. Luigi Nardi
Stanford & Lund

Advisor
Dr. Kunle Olukotun
Stanford & Co-founder
SambaNova

dbtune

Senior ML Engineer
Dr. Erik Hellsten
Chalmers, DTU, LTH, Volvo

Frontend Engineer
Aiman Mohsin
Diya, Sia Smtech

Dev

Senior Software Engineer
Muhammad Umair
Freie, Heidelberg, SAP

Backend Engineer
Tahir Masood
FAST, Ibex Global

Strategy Advisor
Kingston Duffie
Serial Entrepreneur

Technology Advisor
Peter Zaitsev
Co-founder & CTO Percona

Tech Lead
Costa Alexoglou
Co-founder VisualEyes, Neo4j

Senior DevOps
Mohsin Ejaz
EDB

Special Consultant
Magnus Hagander
Redpill, PG Core Team

Technology Advisor
Johan Svensson
Co-founder & CTO Neo4j

M & S

Marketing Coordinator
Ellyne Phneah
LTH, ZDNet, Symantec

Director of Sales
Tom Howcroft
VoltDB

Marketing Advisor
Mark Jennings
Techstars, Notch

Sales Advisor
Alan Facey
B2B Sales Leader

External pressures

Challenges facing all enterprises

- Inflation and economic downturn
- Increasing customer expectations
- Security breaches
- Explosion of automation
- Accelerated digital transformation
- Sustainability

Impact on the enterprise

Manifests in adjusted business priorities

- Operational efficiency / cost reduction
- Adoption of open source
- Employee productivity / satisfaction
- Embrace automation with AI / ML
- Focus on security
- Reduce carbon footprint

Opportunity #1: Infrastructure cost optimization

- 35% of all cloud spend is wasted
- 40% of instances are over-provisioned
- **FinOps** departments — **Financial DevOps**, are tasked with finding efficiencies

Estimated annualized committed cloud spend as % of cost of revenue

Source: A16z article *The Cost of Cloud, a Trillion Dollar Paradox*

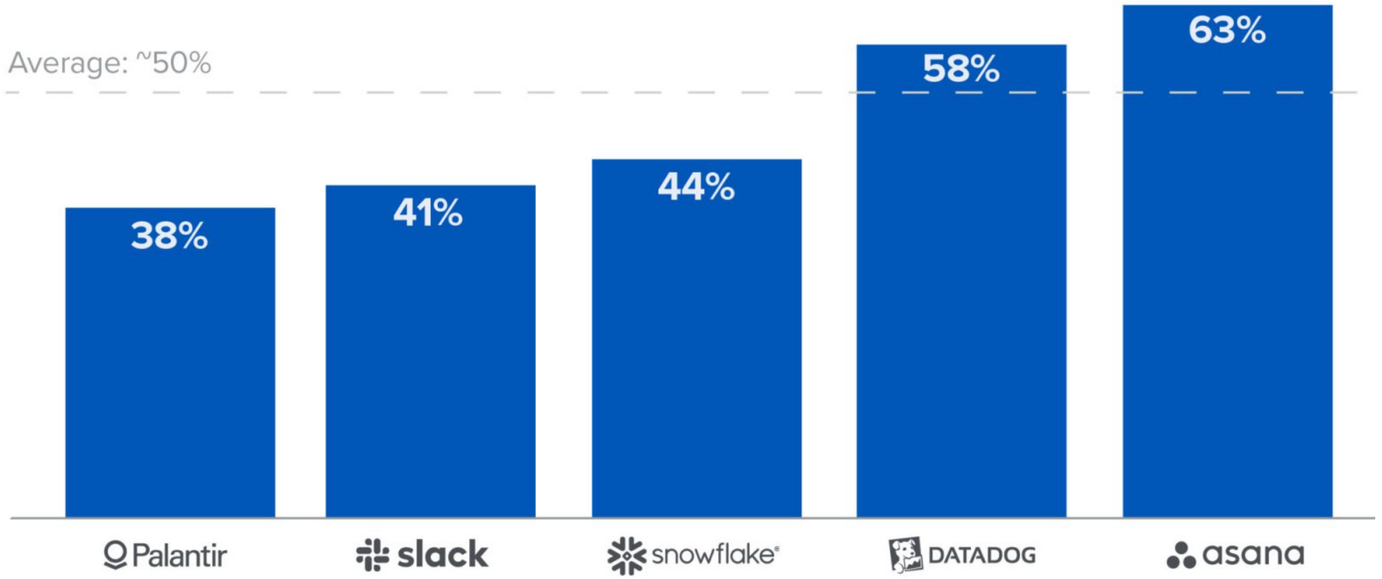

Opportunity #2: The sustainability goal of green computing

Supporting environmental, social and governance (ESG) goals

General companies

Datacenter operators

CSRD

From 2025 the EU Corporate Sustainability Reporting Direction will require companies to report on sustainability

Energy efficiency directive

Reduce energy consumption and greenhouse gasses by promoting energy efficiency

Eco design directive

Set of rules to limit environmental impact of servers & data storage

Data center standard

Publications that define EU standards for development and construction of data centers

- Increased public and regulatory pressure on data infra to reduce CO2 emissions
- Most are recommendation practices at present
- In 2025 CSRD will be mandatory for 2024 CO2 emissions

What is database tuning?
And how can it help us deliver against strategic objectives

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, DBMS parameter, index, OS parameter, etc.

Why does it matter?

Technical perspective

- Directly impacts system performance
  - Transactions per second — Throughput
  - Average query runtime — Latency
- Improves scalability
- Enhances stability / reliability
- SLA

Business perspective

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

Database system parameter tuning

- Adjusting knobs to best fit the workload
- PostgreSQL parameters that are typically important: *work_mem*, *shared_buffers*, *max_wal_size*, etc.
- Example work_mem:
  Memory allocated for each operation within a query
- Example shared_buffers:
  How much memory will be used for page cache
- These parameters highly depend on the application

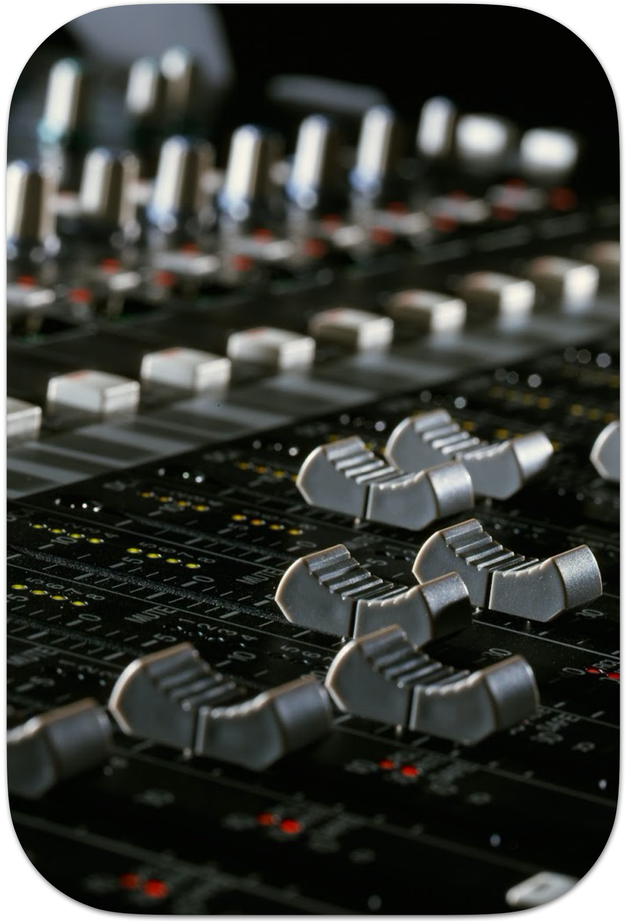

A synthetic example for *shared_buffers* and *work_mem*

ResourceStresser benchmark from the *BenchBase* benchmarking suite

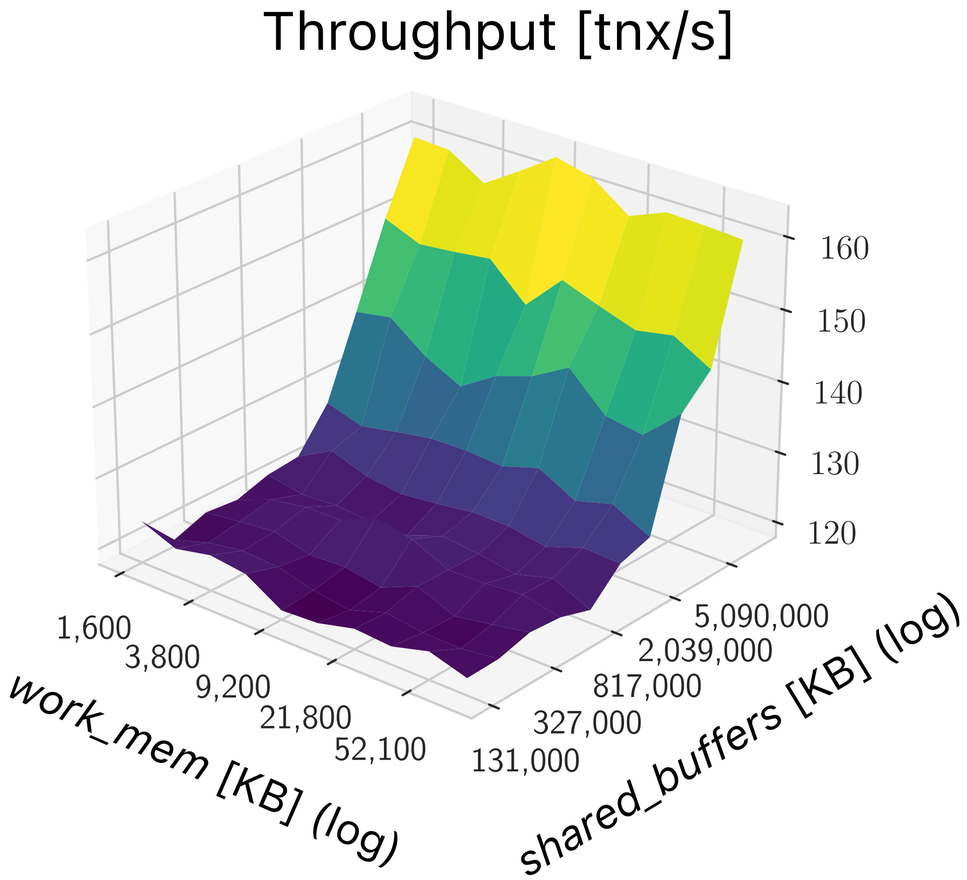

ResourceStresser is disk-bound:

- Increasing shared_buffers is important
- But not work_mem — Queries are simple

Complexity is growing over time

The number of parameters
is growing **linearly**

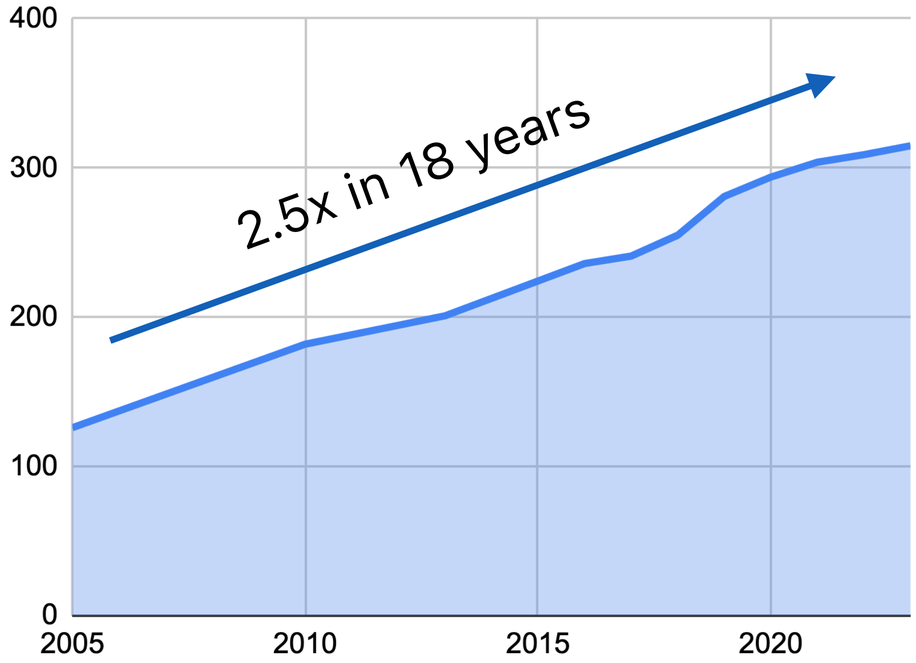

PostgreSQL number of parameters

The number of configurations
is growing **exponentially**

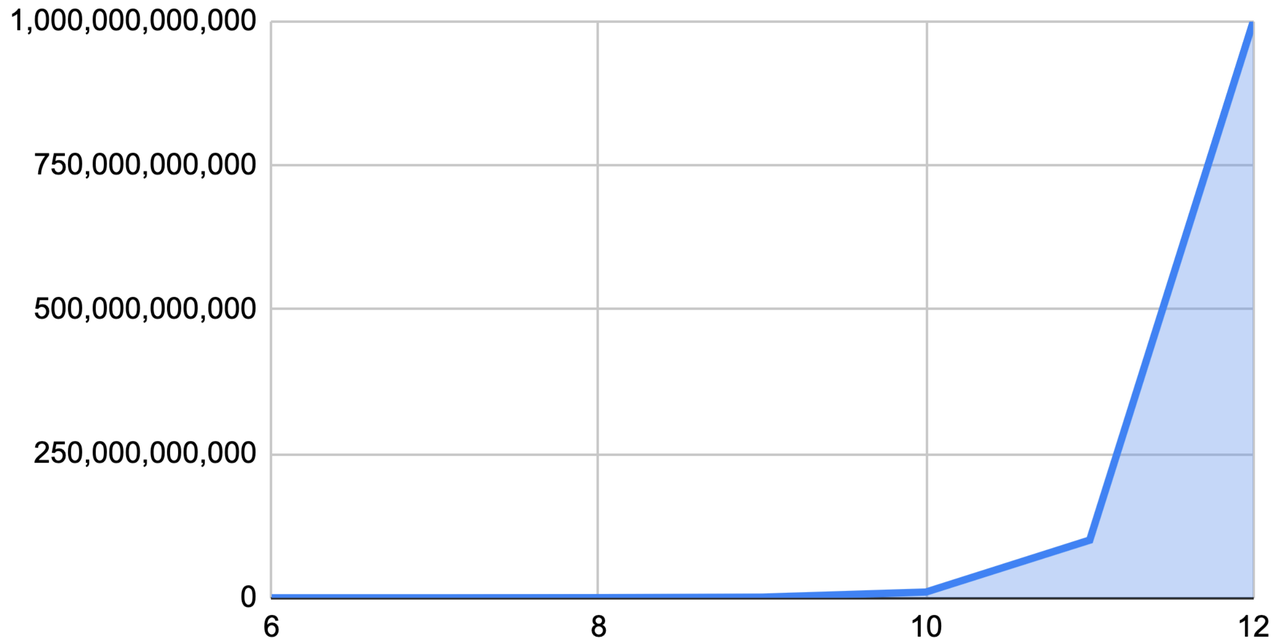

Example of complexity with 12 parameters

How is parameter tuning tackled today by DBAs and developers?

Tuning guru

Manual

Slow
Takes days

Painstaking
Needs high expertise

Ineffective
Tune again in a week

Inadequate
Seasonal workload

Heuristics

PGTune

POSTGRESQL CONFIGURATOR

One-size-fits-all
Uses generic rules

Workload agnostic
Not bespoke

Ineffective
Tune again in a week

Inadequate
Seasonal workload

New approach

dbtune

Ideally a solution that learns by observation and autotunes

A solution that adapts to changing workloads

How often do you tune?

Triggers for tuning

Anytime that

- Your workload changes — Change queries and application
- Your database grows and changes
- You migrate from on-prem to the cloud — Or vice-versa
- You scale your cloud instance — Up or down
- You migrate DBMS — E.g., from Oracle to PostgreSQL
- You upgrade your version of PostgreSQL

The reality of how most enterprises treat manual parameter tuning today

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

We introduce DBtune

A unique AI-powered database tuning cloud hosted service

Machine learning approach

DBtune learns how to solve optimization challenges

Dynamic adaptation

DBtune can tune a database irrespective of its size and complexity

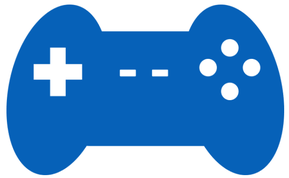

Easy to use

No need for background in AI or database tuning

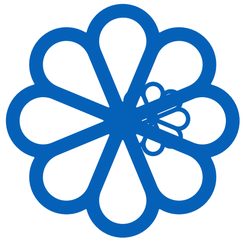

Highly scaleable

DBtune can tune multiple databases in heterogeneous environments

Customer value propositions

DBtune boosts service performance / improves business margins

Reduce cloud / infrastructure costs

Make your service radically faster

Free up your DBAs

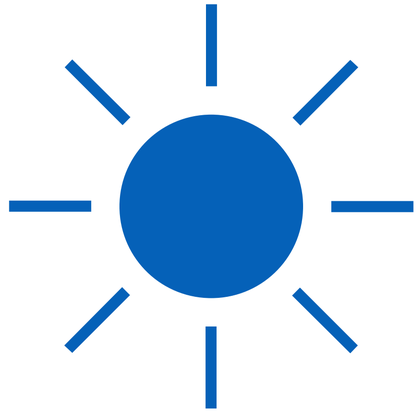

Reduce energy consumption

DBtune architecture for self-managed PostgreSQL

High-level view

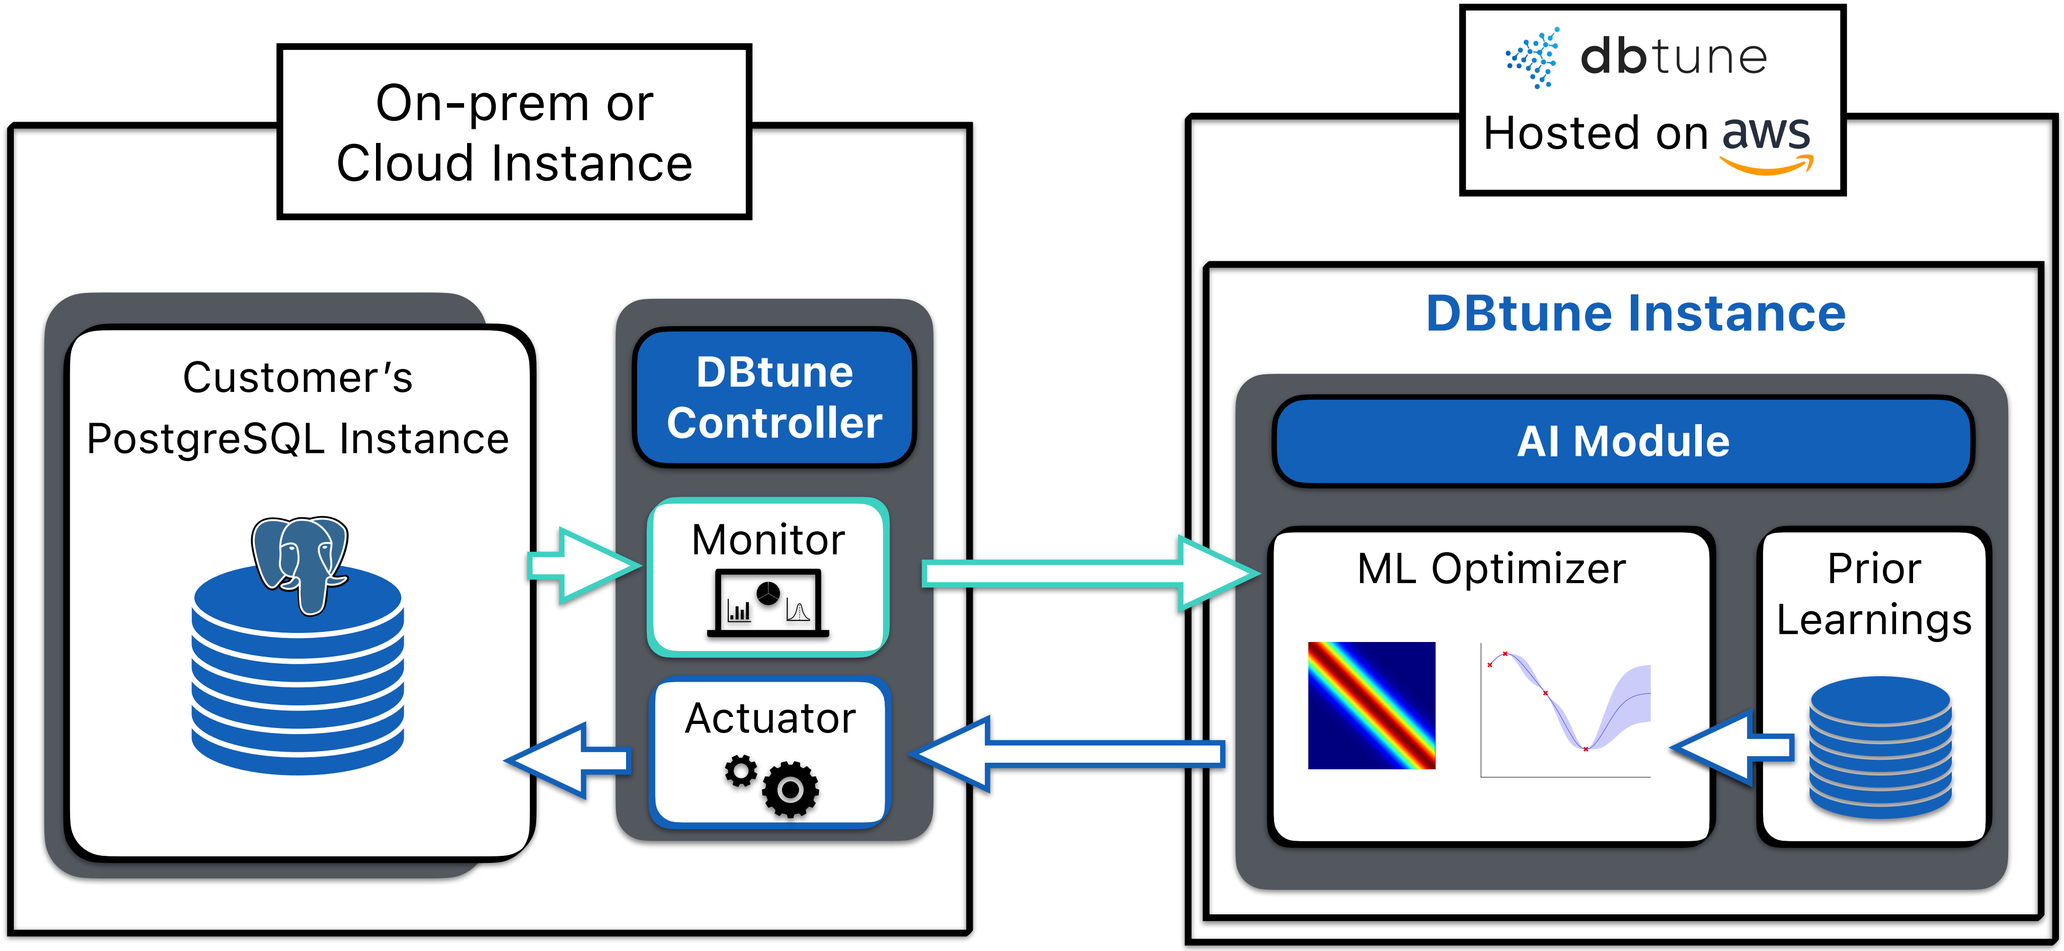

DBtune architecture for instances that are offline

High-level view

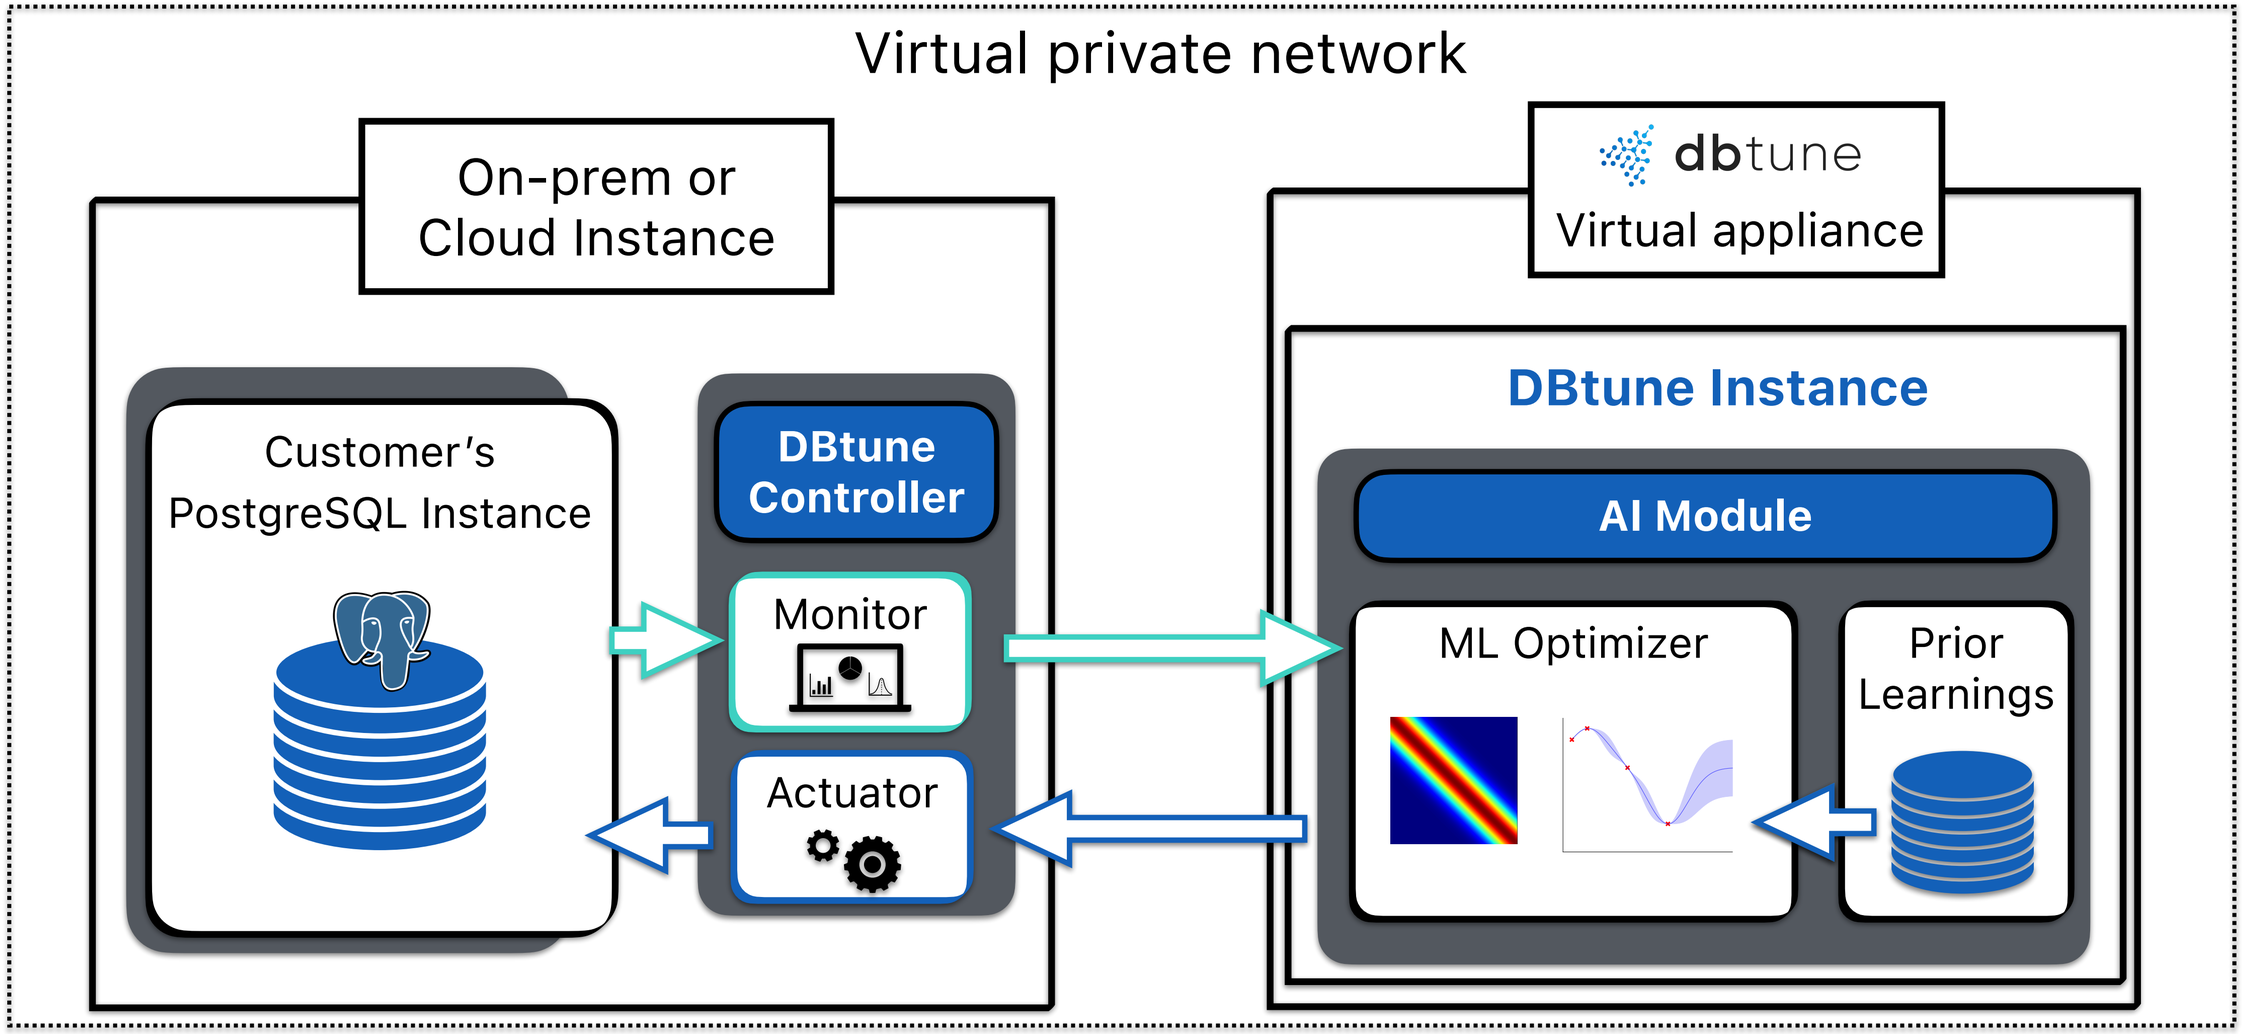

DBtune for Amazon RDS

High-level view

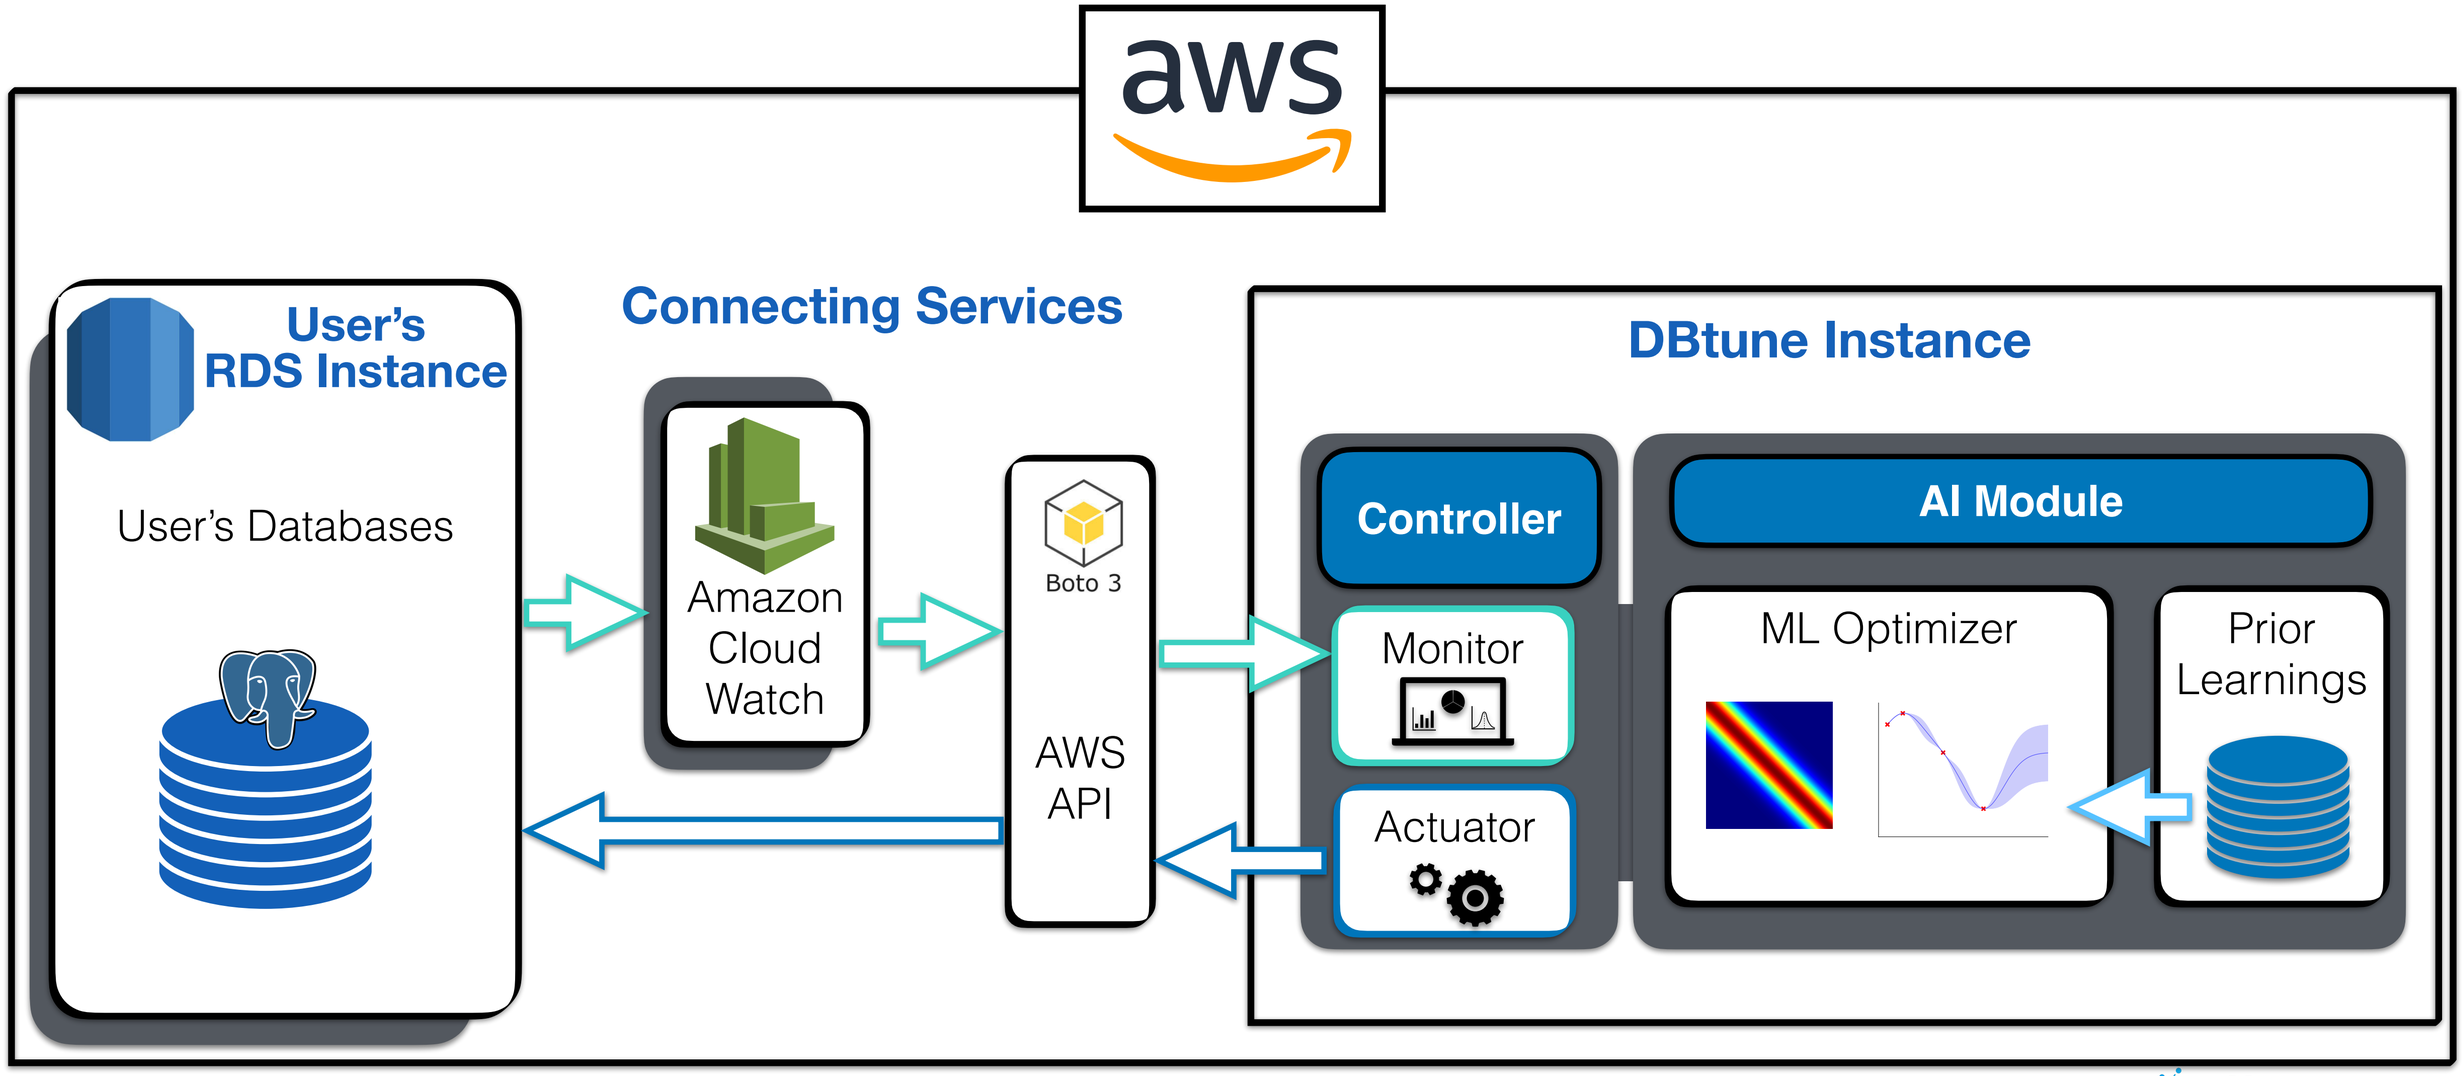

Proof of cost reduction: Comprehensive test results

DBtune doubles the performance of PostgreSQL Amazon RDS

Performance impact of tuning RDS m5.2xLarge cloud instance on the TPCC benchmark

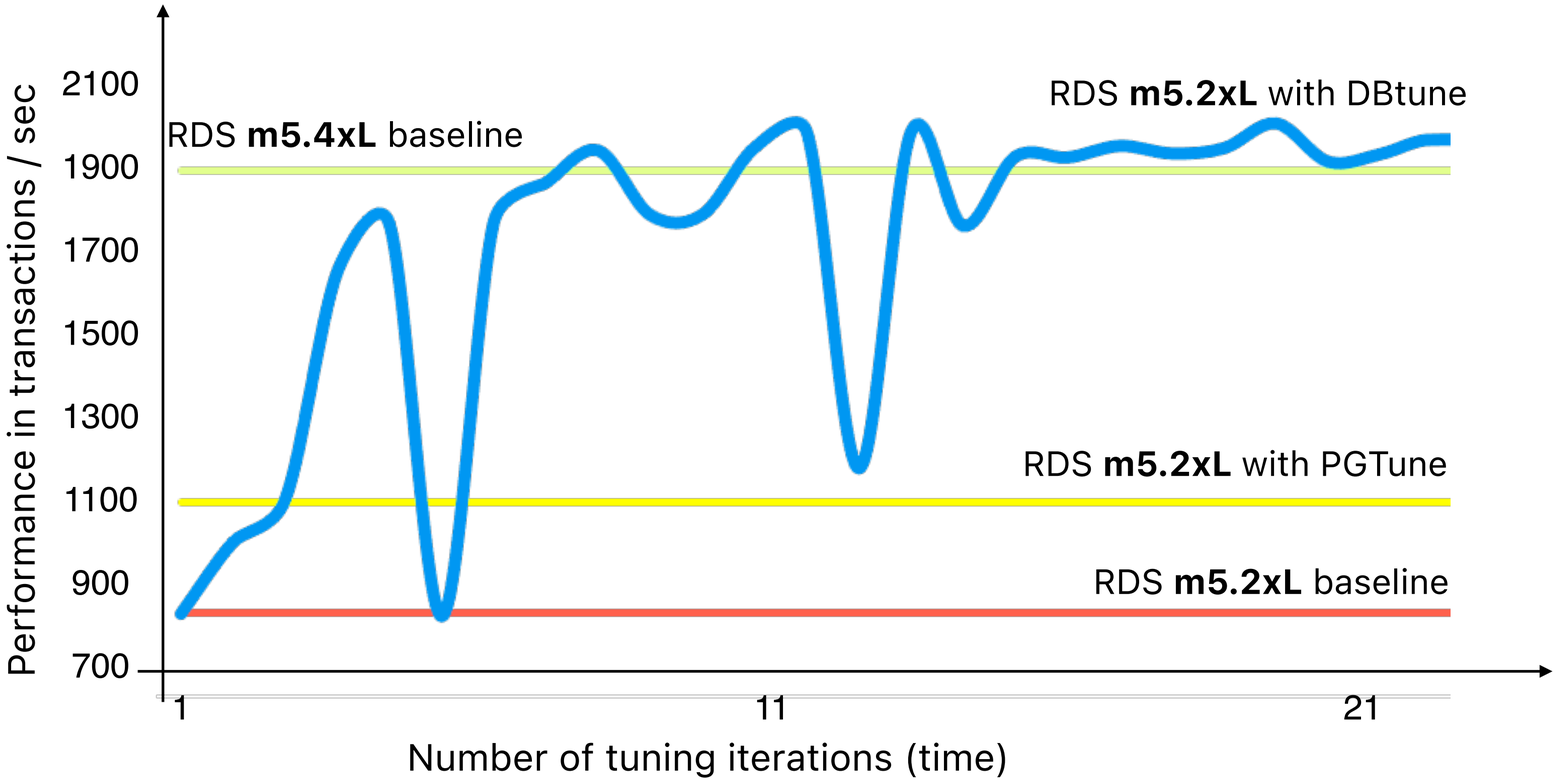

DBtune on the smaller instance type achieves a level performance in excess of that achieved by an instance twice the size

Proof of cost reduction: Detailed cost analysis

DBtune doubles the performance of PostgreSQL Amazon RDS

Hardware				Cost / Year		
AWS RDS Instance Type	Cores	RAM	IOPS	Instance	EBS	Total
db.m5.4xlarge	8	64 GBs	4000	$12,475	$4,800	$17,275
db.m5.2xlarge	4	32 GBs	2000	$6,237	$2,400	$8,637

Per instance savings: $8,638

- DBtune halves RDS cost (50% saving)
- Matches 4xLarge performance on a 2xLarge instance
- Medium and large companies use hundreds* of RDS instances

*A16z article: "The Cost of Cloud, a Trillion Dollar Paradox"

The ESG angle

How much CO2 is on the table?

Database instance size largely impacts data center emissions

Impact on a single database

-50%

CO2 emissions

for a single database with DBtune[1]

Impact on average data center

-32%

CO2 emissions

for an average data center with DBtune[2]

Impact on the EU data center landscape

-7Mt

CO2 savings p.a.

across the EU with DBtune[3,4]

1) DBtune/Teads 2) Borderstep 3) Statista 4) EU digital strategy

Results on PostgreSQL running on AWS EC2

DBtune improvement in throughput and latency

Tested DBtune on a standard simulated DB implementation (**Wikipedia** OLTPBench benchmark)

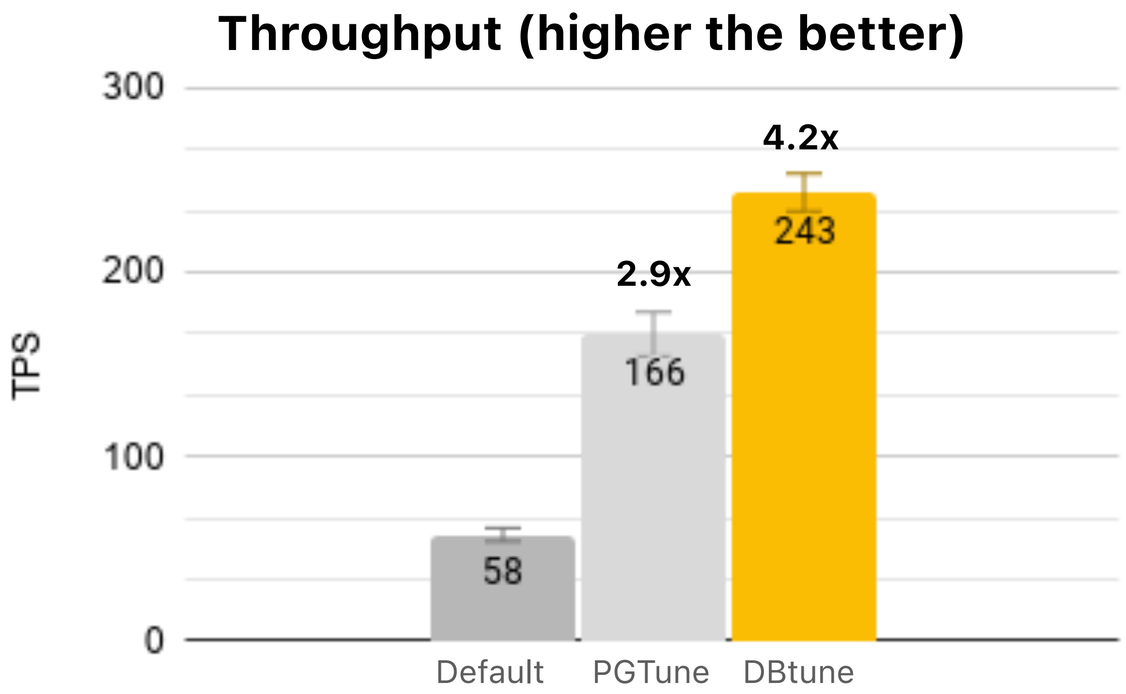

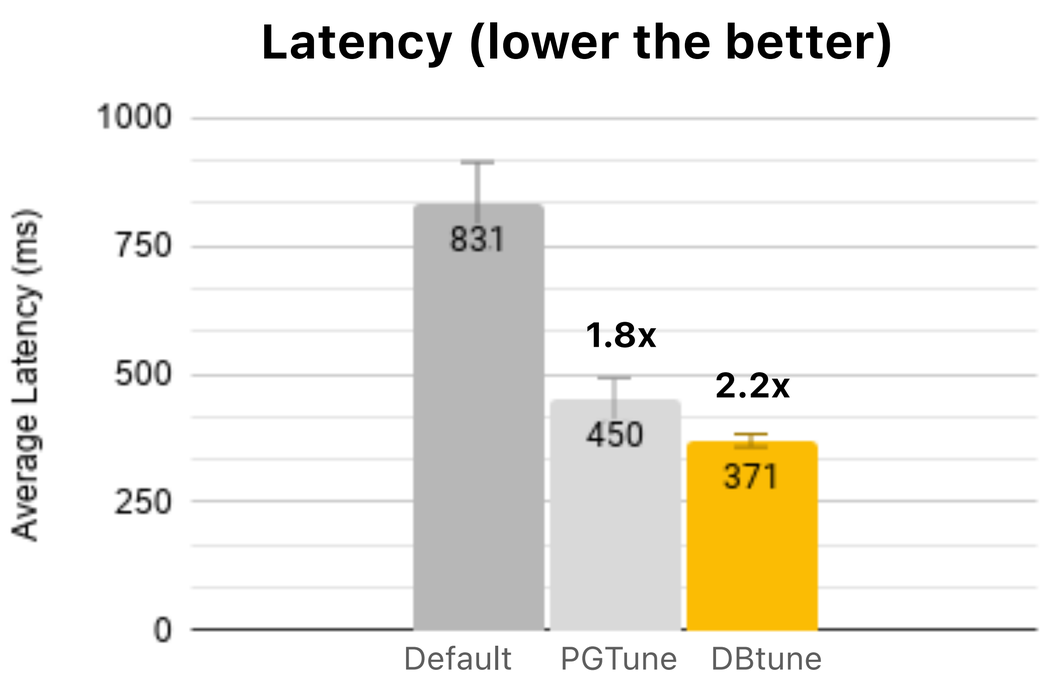

Results on PostgreSQL running on AWS EC2

DBtune improvement in throughput and latency

Performance impact of tuning the CH-benchmark benchmark

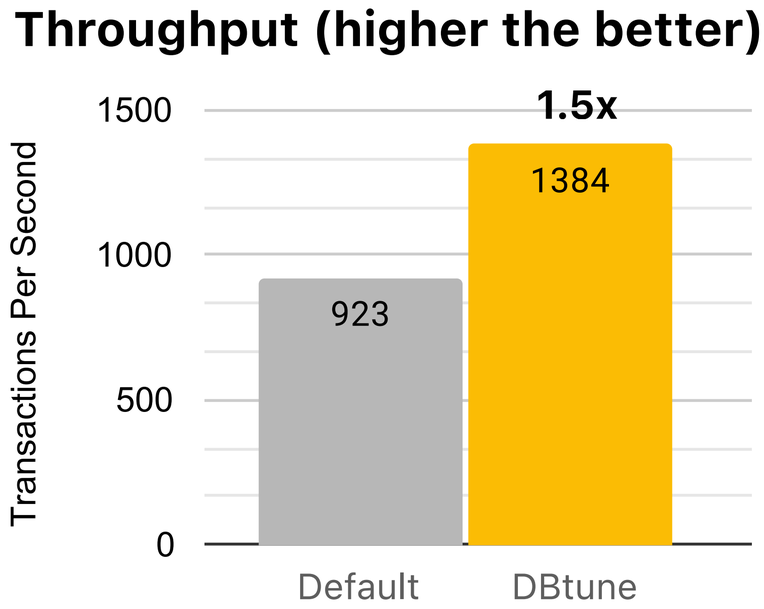

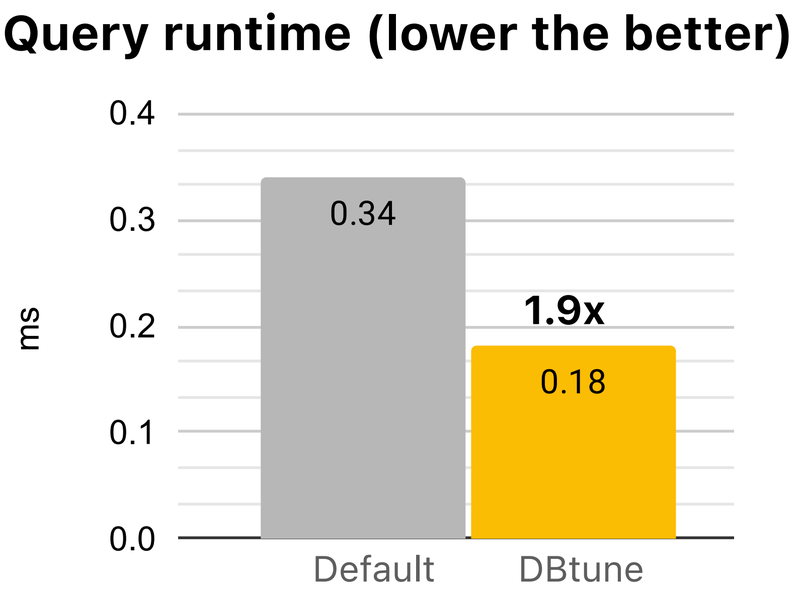

Customer story: Airtel production system optimization

Airtel partnered with DBtune to optimize their infrastructure spend

Airtel is one of the largest communication service providers, globally. Headquartered in India, they serve in excess of half a billion subscribers. Airtel sought a new technology to improve their PostgreSQL database performance.

“DBtune seamlessly integrated into a production system of a mission critical Airtel application. We’ve been impressed by the reliability and robustness of the DBtune product, and the team has enjoyed evaluating the platform.”

Anant Kumar
Airtel CIO digital

Customer story: Helping Anteo to speed up their data operations

Norway-based company, Anteo, offer decision support for sustainable development in the aquaculture industry, as well as real-time monitoring and warning biosafety solutions.

Anteo's infrastructure is data intensive. Anteo partnered with DBtune to speed up their PostgreSQL data platform.

"It only took 10 minutes to set up DBtune on our Google Cloud PostgreSQL data platform...The process was easy and pleasant."

Peder Refsnes
Anteo CTO

Customer story: Integration study with the DbVisualizer platform

DbVisualizer is a leading universal database tool for universal database management systems. The company offers a database Integrated Development Environment (IDE) for developers, analysts, and DBAs.

DbVisualizer partnered with DBtune to explore the technical integration with their development platform. The initial pilot validated the technical strength of the DBtune platform.

DbVisualizer

“We see a lot of potential in DBtune’s ability to optimize our customers’ workloads. This is a state-of-the-art optimizing service that is robust and flexible enough to integrate tightly with our platform. DbVisualizer, enhanced with DBtune’s capabilities, would make for a more complete offering for our customers."

Martin Engdahl
DbVisualizer CEO

DBtune technology endorsed by VMware

vRealize Network Insight (vRNI) is a network monitoring tool by VMware that helps build an optimized, highly available and secure network infrastructure across cloud environments. The key-value store FoundationDB database system is at the core of vRNI and its performance.

vRNI's infrastructure is data intensive. VMware partnered with DBtune to speed up their FoundationDB data platform.

"We saw a 34% improvement in our FoundationDB testbed, while we were hoping for a 10% improvement...DBtune exceeded our team's expectations."

Clement Pang, Co-founder & Chief Architect at Wavefront by VMware

"For us, performance is essential, DBtune has overcome the optimization complexity with an innovative solution; they made it simple."

Uday Kurkure
Staff Engineer at VMware

Independent evaluation by Stormatics

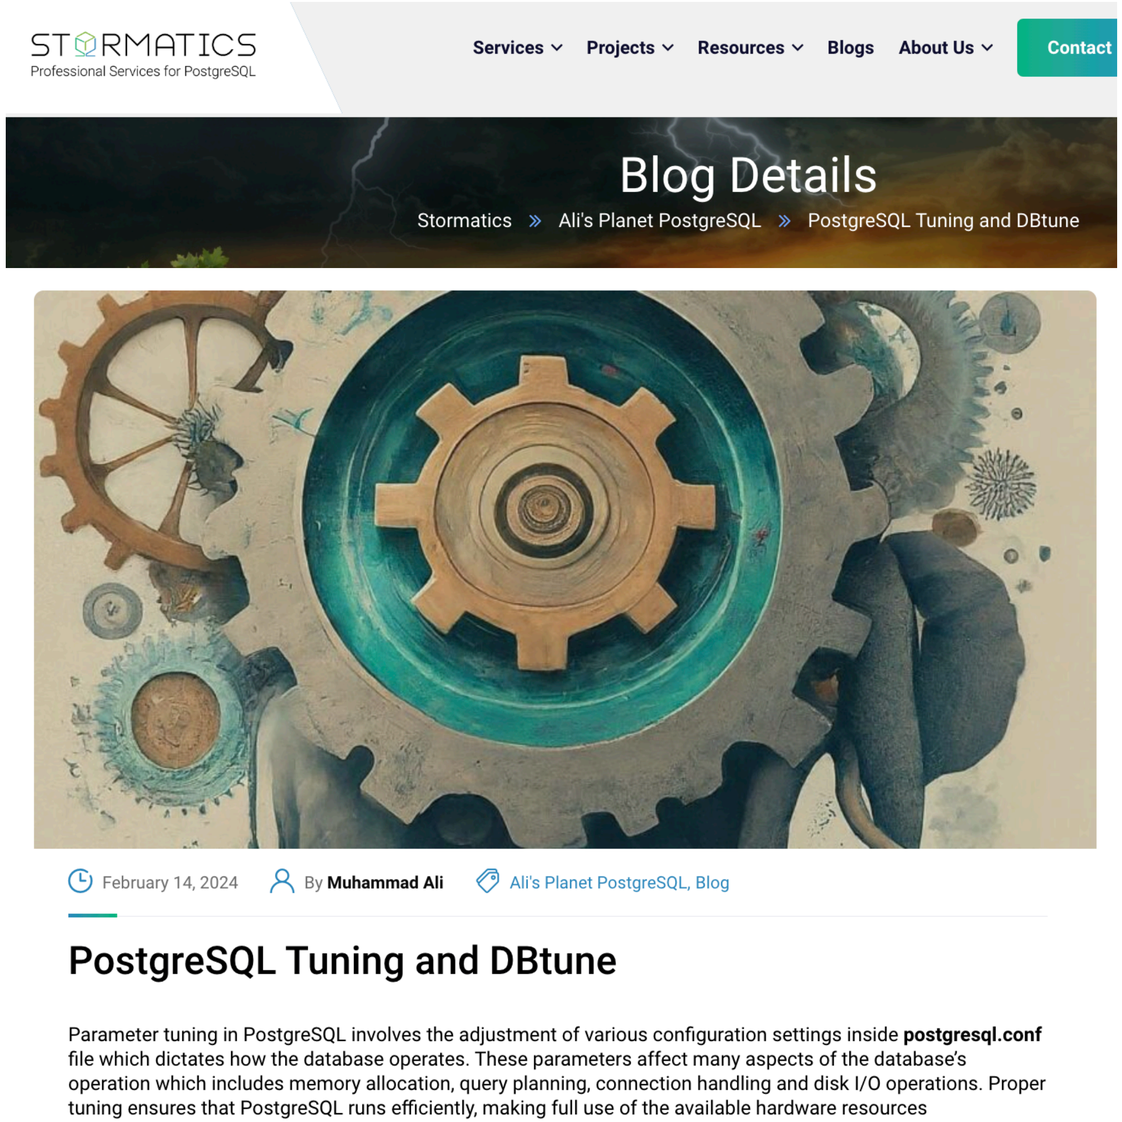

- Across all tests cases DBtune delivered improvement in performance up to 13.6x
- Compared to manual tuning DBtune achieved 2.2x speedup

Blog: https://stormatics.tech/alis-planet-postgresql/postgresql-tuning-and-dbtune

PG Developer Day Prague DBtune training

Live DBtune tuning with 25 attendees (January 31st, 2023)

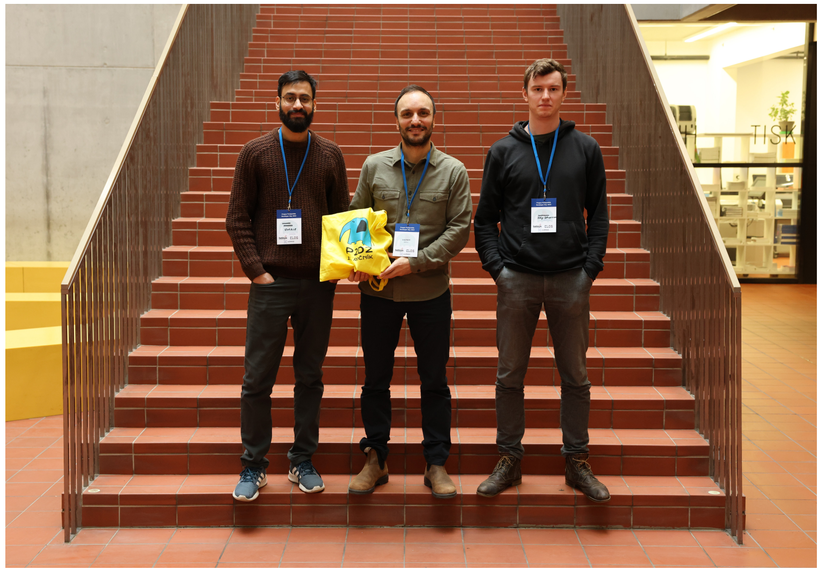

On the left, a photo of our training session. On the top right three members of the DBtune team, Umair, Luigi and Filip, who delivered the training, and bottom the full event.

Sign up today!

app.dbtune.com

Or request a demo

luigi@dbtune.com

dbtune