WHY DEVELOPERS USING AI IS GREAT FOR DBA JOB SECURITY?

MAYUR@VEEAM

PRAGUE POSTGRES MEETUP

DATABASE COMEDY BLOG

ME@POSTGRES SLACK

PGSCORECARD.COM

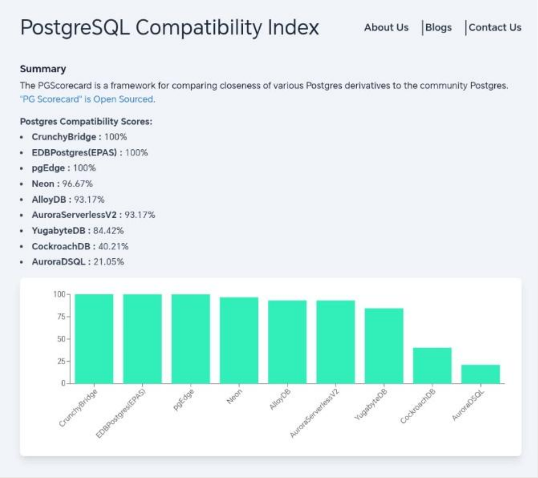

AGENDA

- Brief introduction to AI Hallucinations and LLM lies.
- Vibe coder CEO story
- Concept of Data Cannibalization and limits of AI training.
- What AI can do for your Postgres database (with human DBA oversight)?

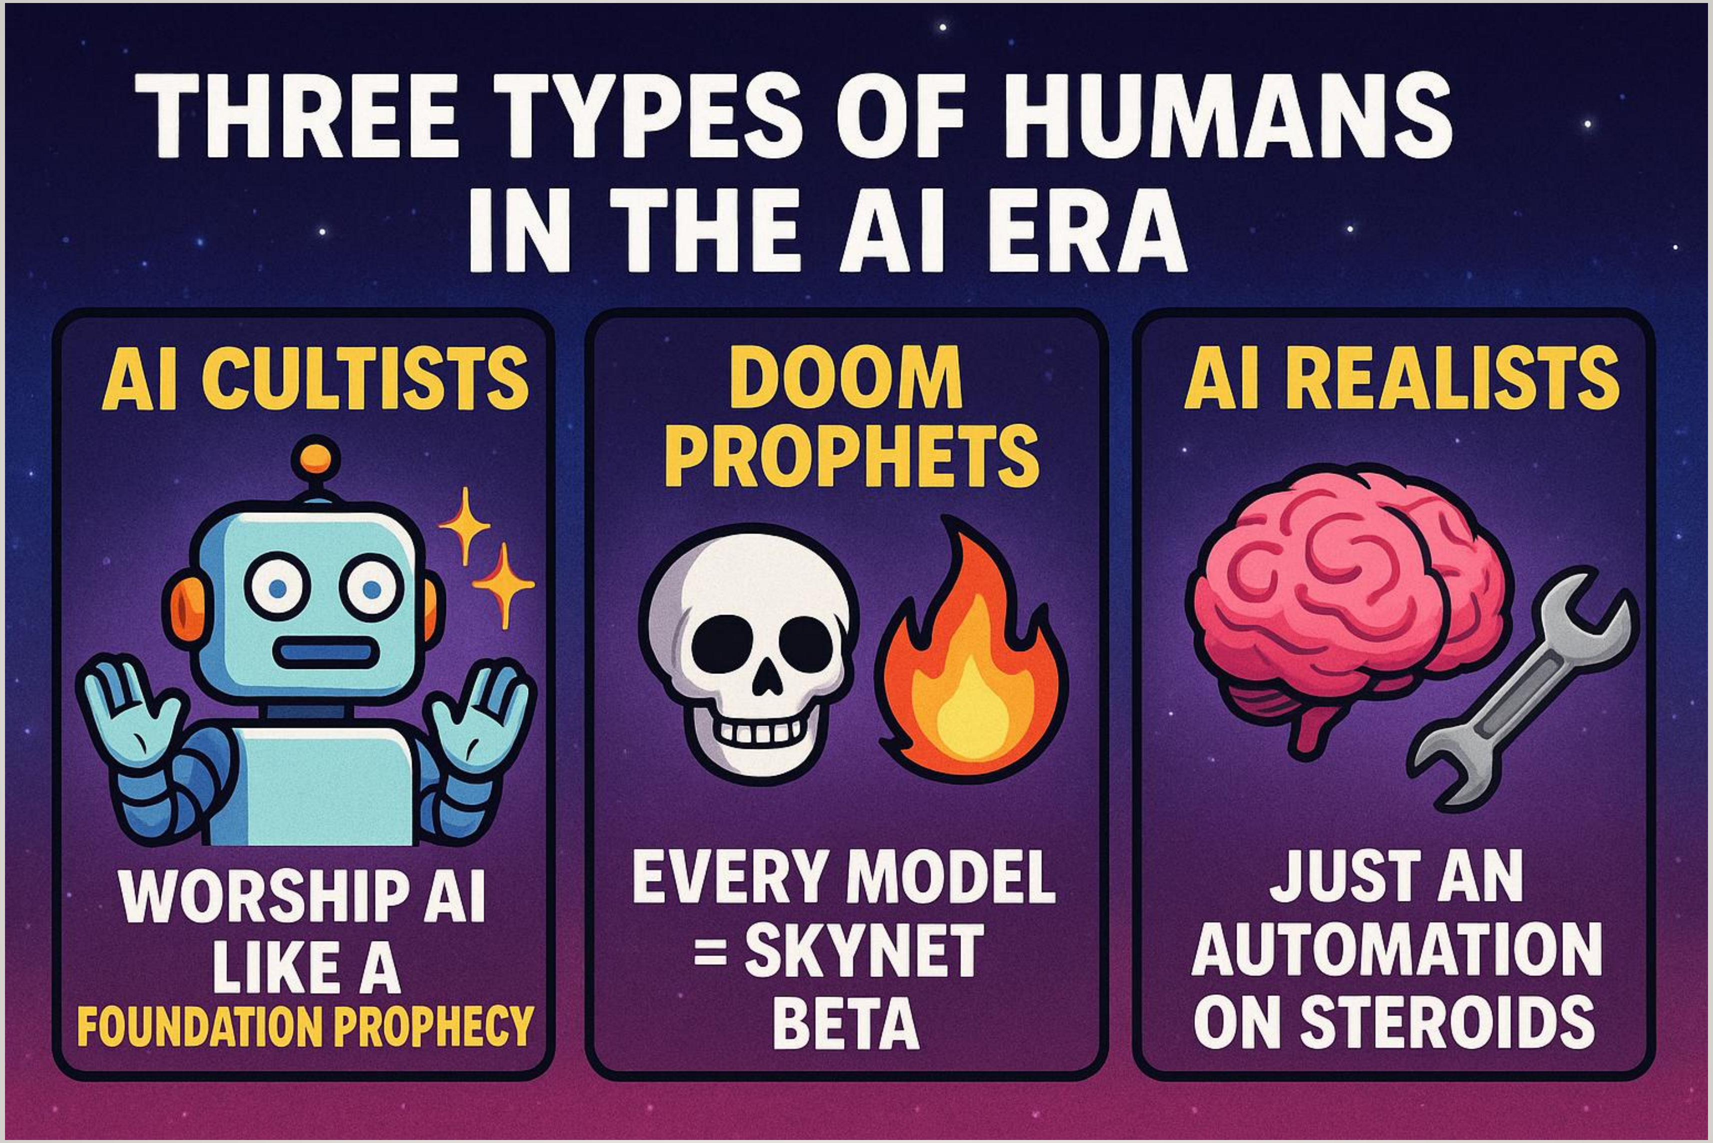
THREE TYPES OF HUMANS
IN THE AI ERA
AI CULTISTS
WORSHIP AI
LIKE A
FOUNDATION PROPHECY
DOOM
PROPHETS
EVERY MODEL
= SKYNET
BETA
AI REALISTS
JUST AN
AUTOMATION
ON STEROIDS

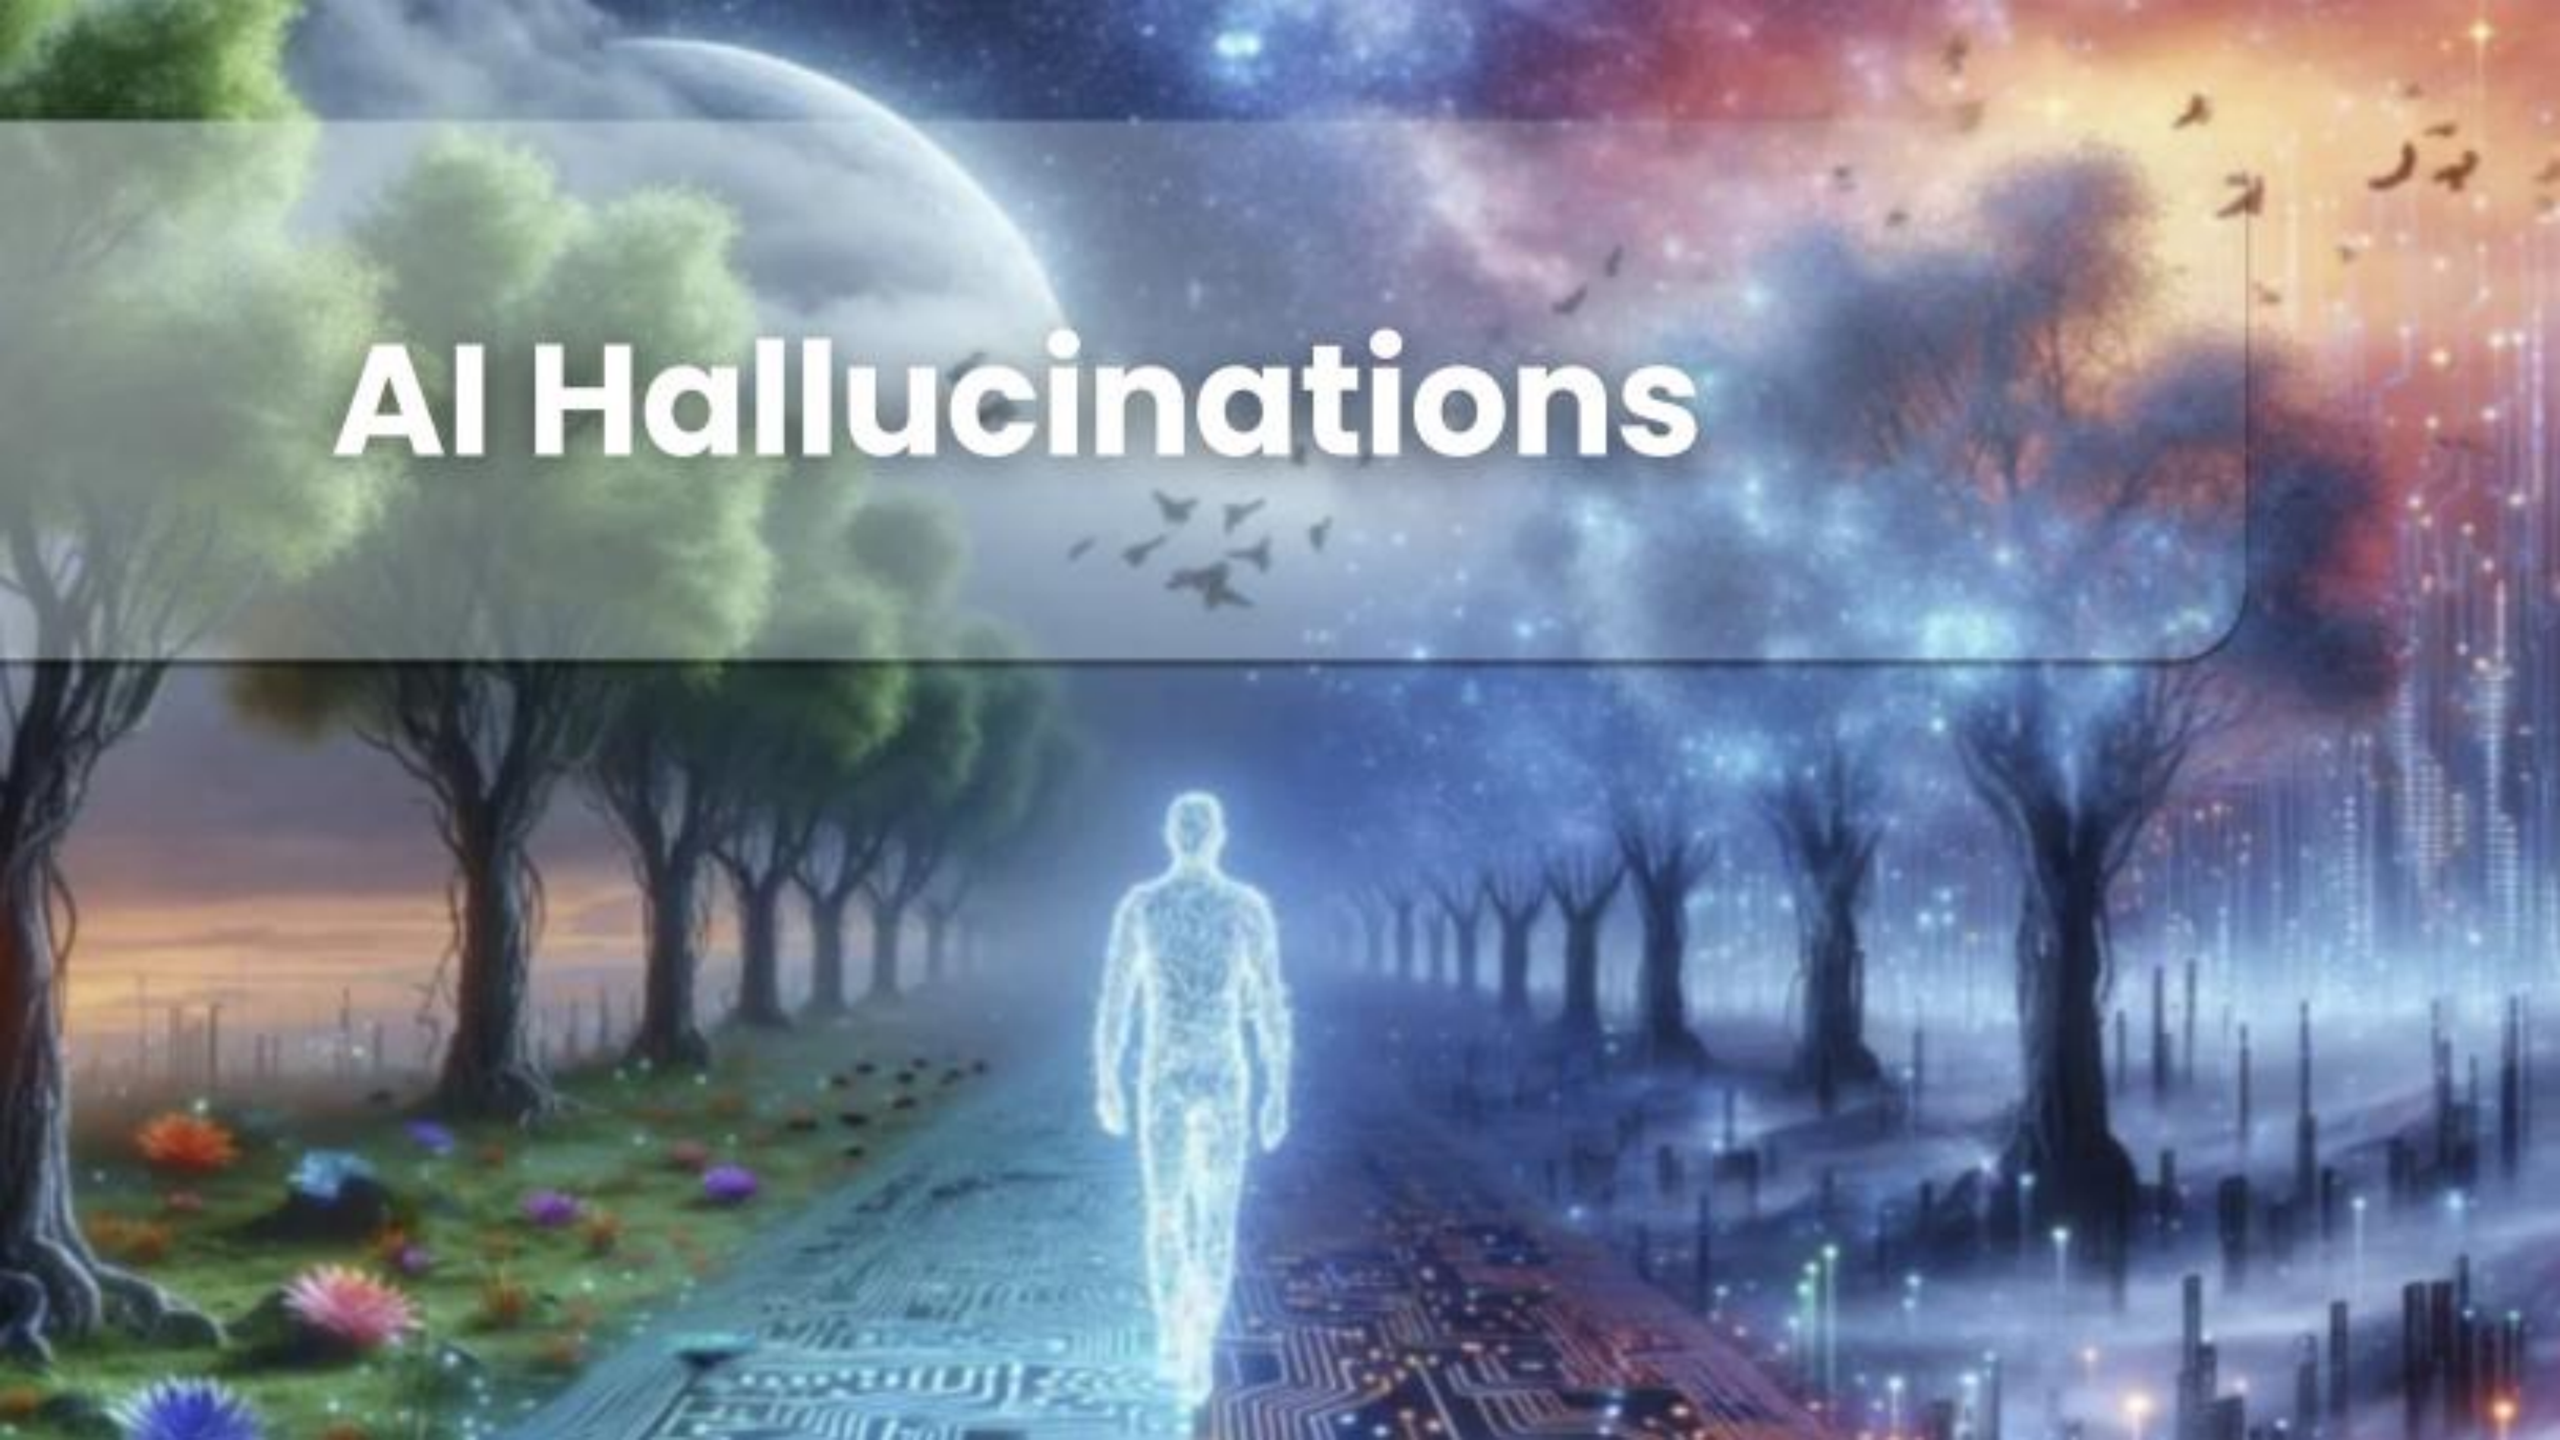
AI Hallucinations

Air Canada ordered to pay customer who was misled by airline's chatbot

Company claimed its chatbot 'was responsible for its own actions' when giving wrong information about bereavement fare

Optimizing for next token, not reality

AI can’t understand context or verify reality like humans

AI HALLUCINATION ARISES FROM LLMS' RELIANCE ON STATISTICAL PATTERNS, PRIORITIZING COHERENCE OVER FACTUAL ACCURACY.

HALLUCINATION IS INTRINSIC BECAUSE LLMS OPTIMIZE FOR TOKEN PREDICTION, NOT TRUTHFULNESS.

NO ABSOLUTE TRUTH: IT CAN’T UNDERSTAND CONTEXT OR VERIFY REALITY LIKE HUMANS.

MIT Technology Review

It’s all hallucination

The takeaway here? It’s all hallucination, but we only call it that when we notice it’s wrong. The problem is, large language models are so good at what they do that what they make up looks right most of the time. And that makes trusting them hard.

MINIMIZE HALLUCINATIONS

Retrieval-augmented generation.

Auto chain-of-thought prompting.

Larger models.

CHATGPT LIES

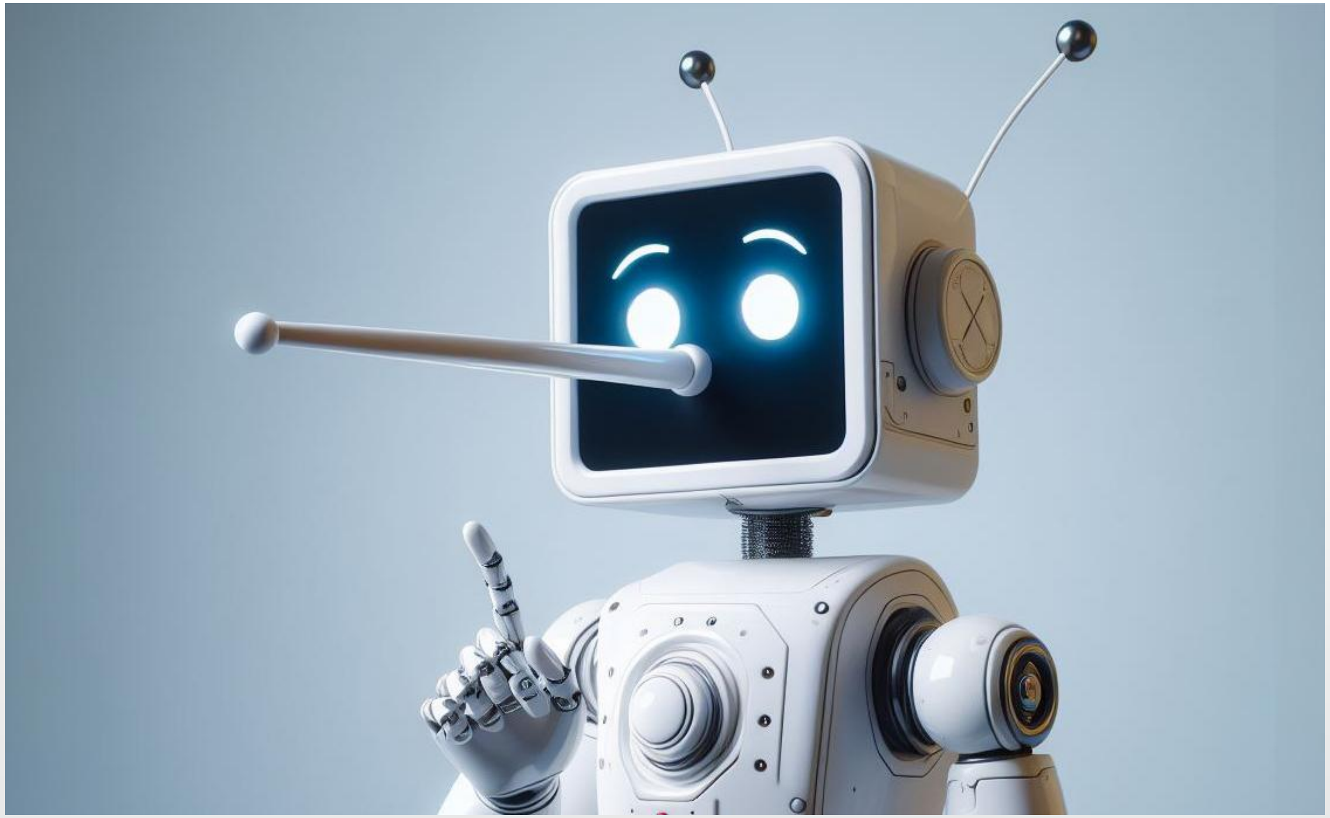

VIBE CEO STORY

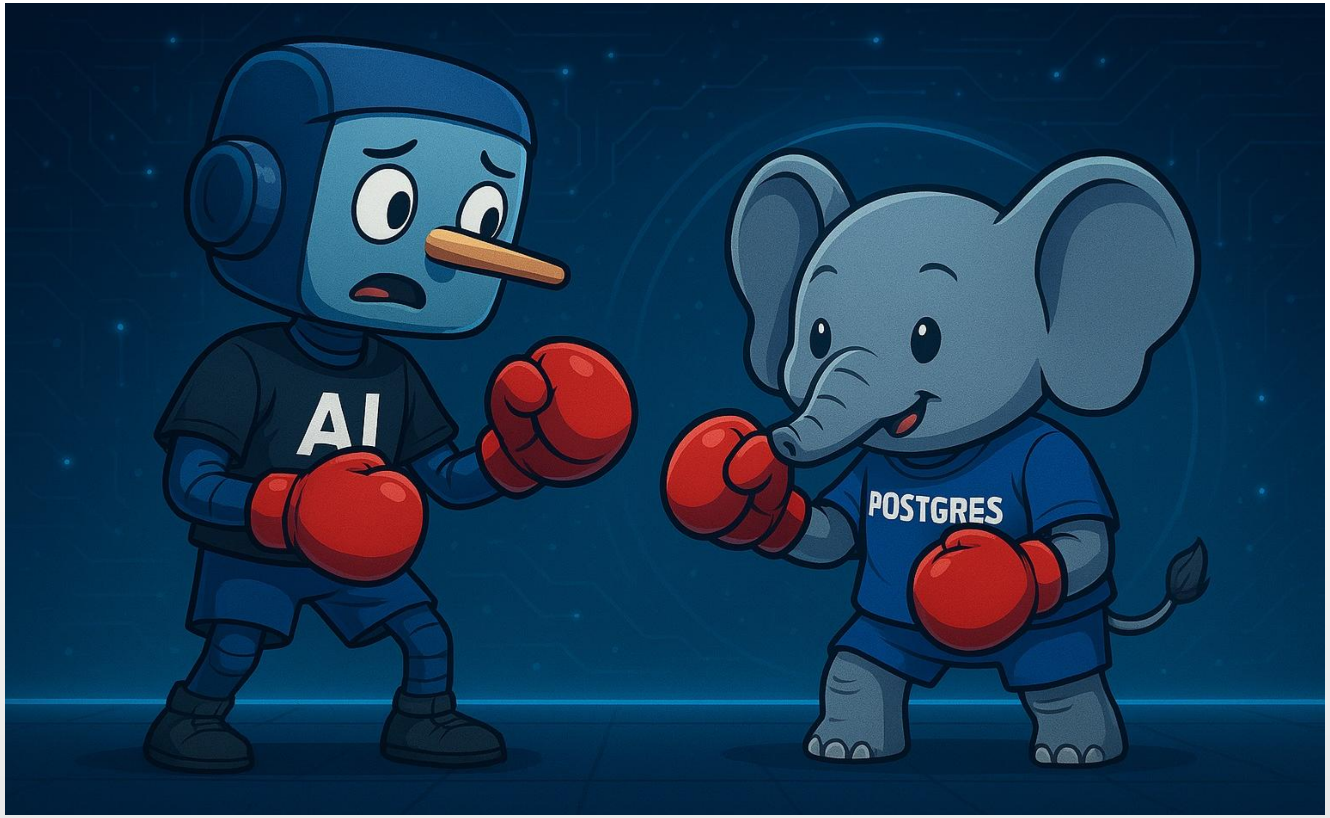

COST REDUCTION INITIATIVE

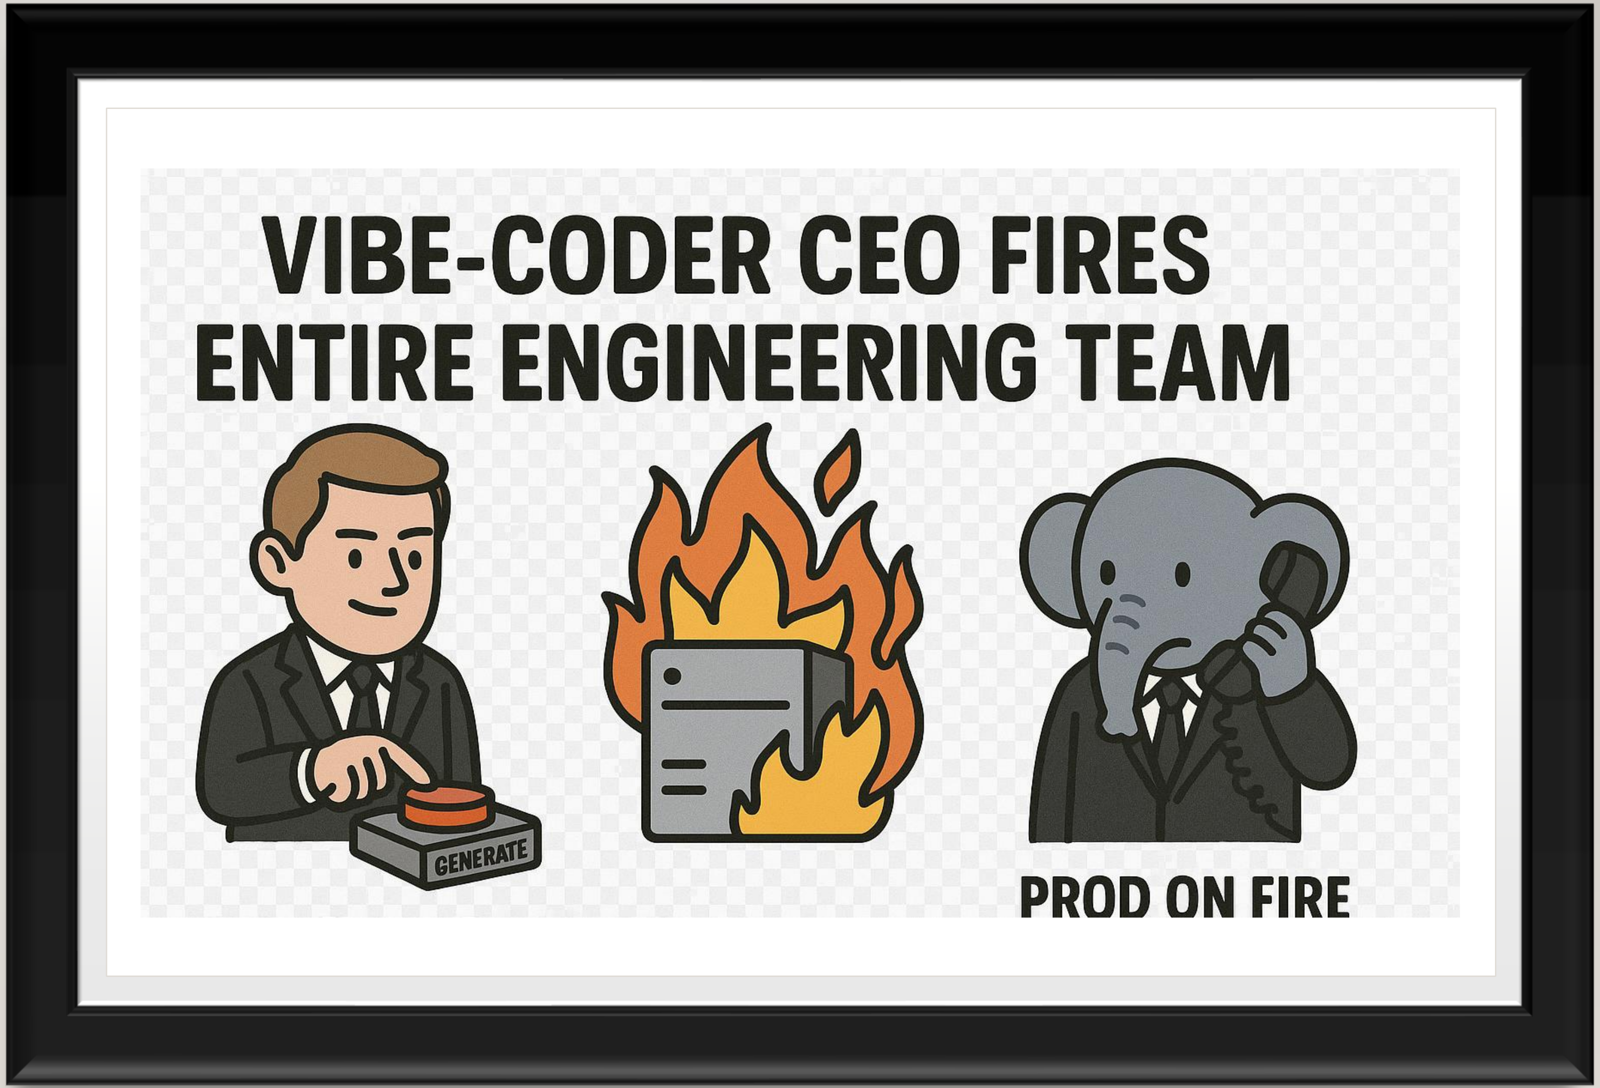

VIBE CODER CEO DISCOVERS ADAPTIVE OPTIMIZER IN POSTGRES.

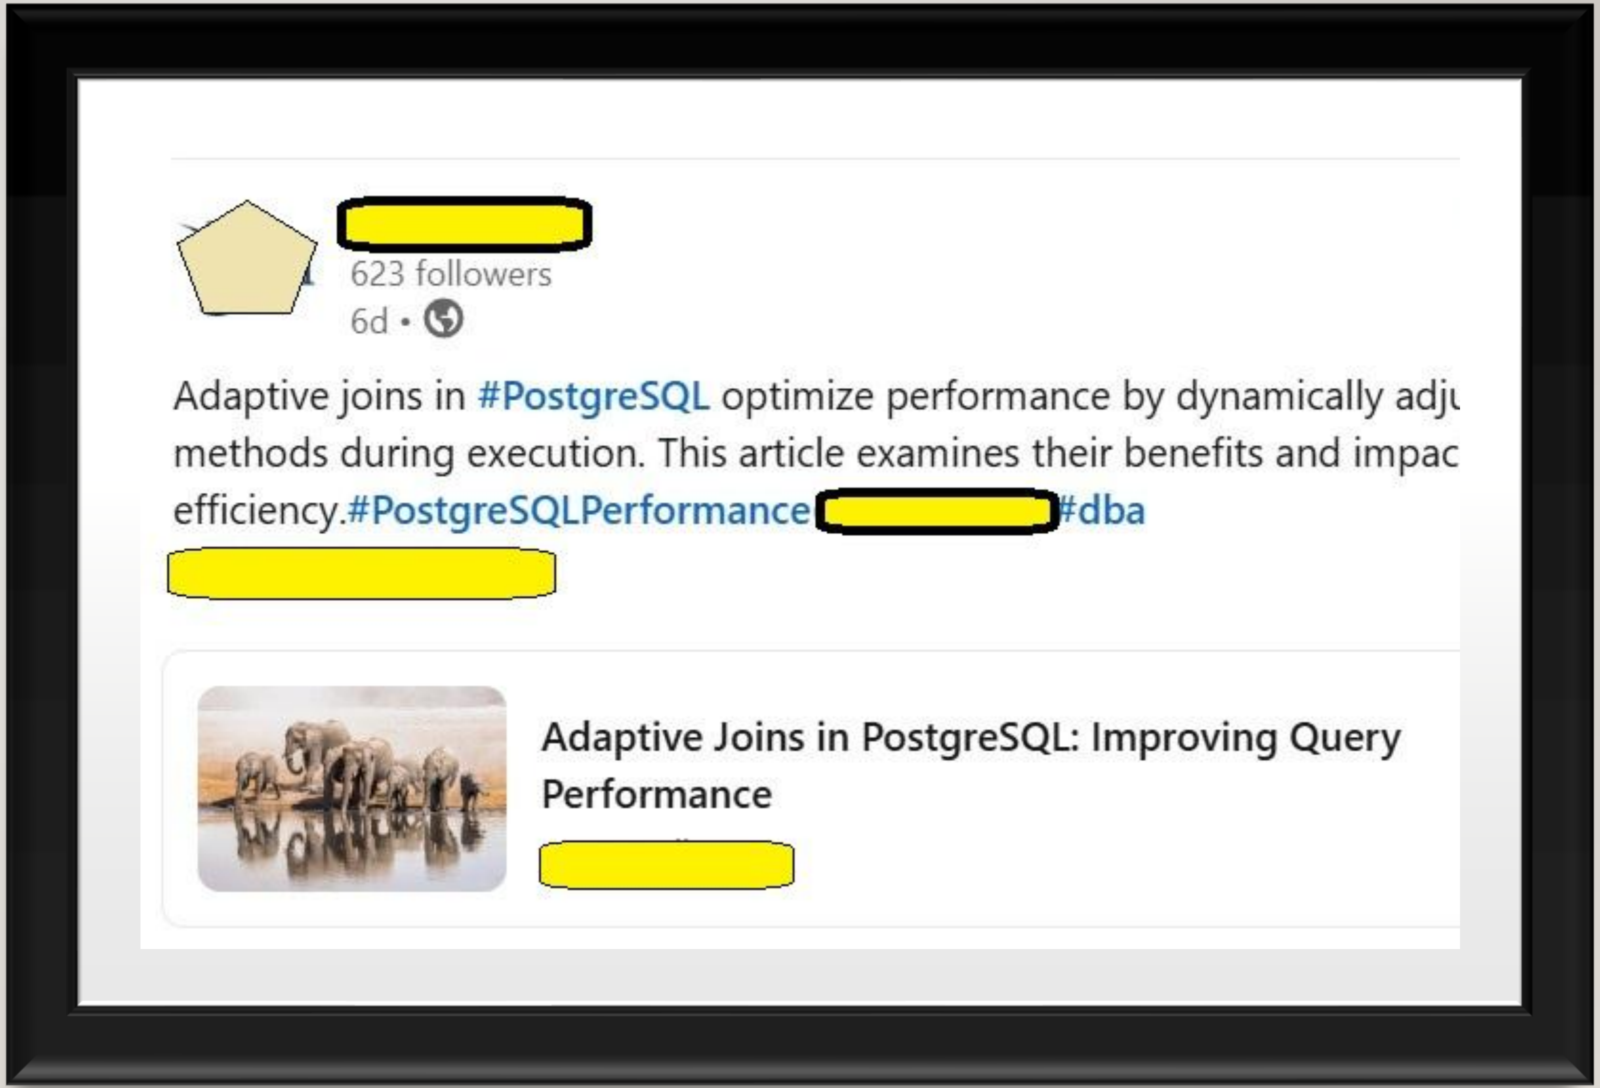

AI TRAINED ON ORACLE DATA PERHAPS

/adaptive-joins-in-postgresql-improving-query-performance/

Traditionally, PostgreSQL would choose a join method, such as a nested loop join, based on its cost estimates before query execution. However, if the join method turns out to be suboptimal due to incorrect cost estimates or changing runtime conditions, it can impact query performance.

With adaptive joins, PostgreSQL can adaptively switch between different join methods based on actual runtime conditions. It starts with an initial join method, such as a nested loop join, and monitors its performance. If the estimated cost doesn't match the actual cost, PostgreSQL dynamically switches to an alternative join method, such as a hash join or a merge join, to improve performance.

Benefits of Adaptive Joins:

VIBE CODER CEO PREPARES FOR A CLIENT DEMO.

VIBE CEO FORGETS TO TUNE QUERIES.

THINKS ADAPTIVE OPTIMIZER OF POSTGRES WILL SAVE THE DAY.

APPLICATION CRASH AND BURN.

VIBE CODER CEO NEEDS TO UPGRADE A PRODUCTION DATABASE

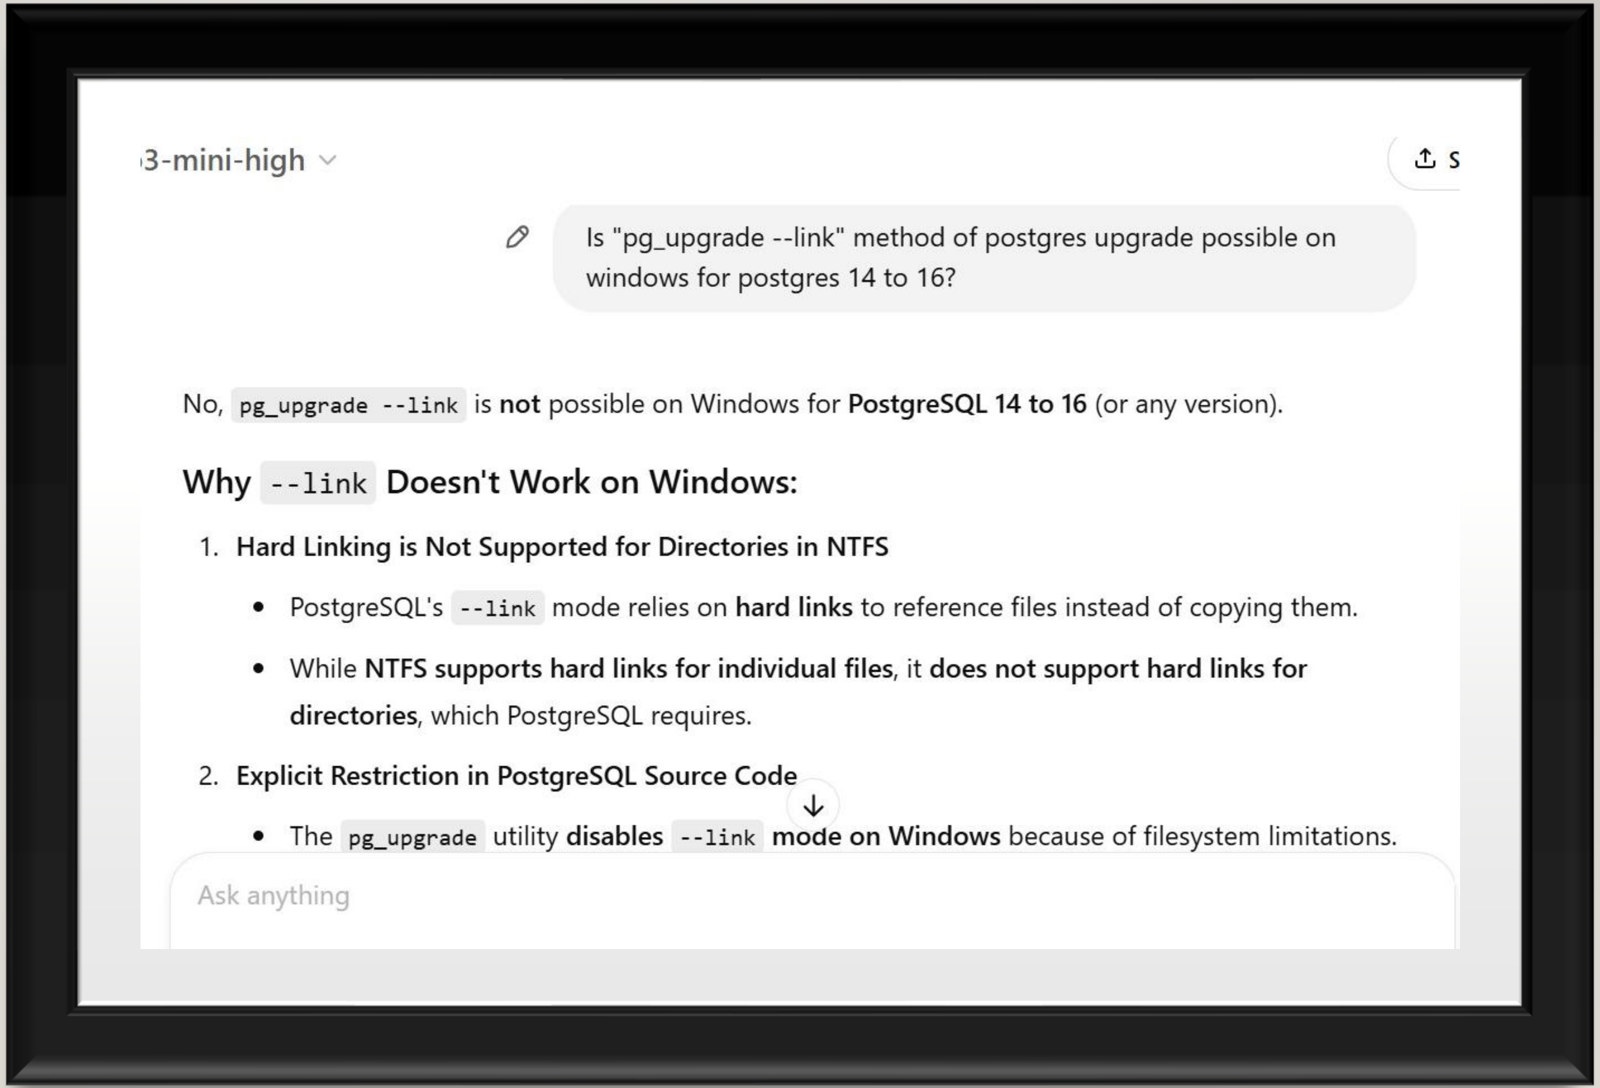

D:\pgdata17

Name	Full Path	Type	Target Path	Created Time
2613	D:\pgdata17\base\1\2613	Hard Link	D:\pgdata\base\1\2613	2/6/2025 2:58:20 PM
2613	D:\pgdata17\base\5\2613	Hard Link	D:\pgdata\base\5\2613	2/6/2025 2:58:20 PM
2613	D:\pgdata17\base\71251465\2613	Hard Link	D:\pgdata\base\71251465\2613	2/6/2025 2:59:35 PM
2683	D:\pgdata17\base\1\2683	Hard Link	D:\pgdata\base\1\2683	2/6/2025 2:58:20 PM
2683	D:\pgdata17\base\5\2683	Hard Link	D:\pgdata\base\5\2683	2/6/2025 2:58:20 PM
2683	D:\pgdata17\base\71251465\2683	Hard Link	D:\pgdata\base\71251465\2683	2/6/2025 2:59:35 PM
71251966	D:\pgdata17\base\71251465\71251966	Hard Link	D:\pgdata\base\71251465\71251966	2/6/2025 3:00:00 PM
71251966_fsm	D:\pgdata17\base\71251465\7125196...	Hard Link	D:\pgdata\base\71251465\71251966_fsm	2/6/2025 3:00:07 PM
71251966_vm	D:\pgdata17\base\71251465\7125196...	Hard Link	D:\pgdata\base\71251465\71251966_vm	2/6/2025 3:00:37 PM
71251974	D:\pgdata17\base\71251465\71251974	Hard Link	D:\pgdata\base\71251465\71251974	2/6/2025 3:00:00 PM
71251975	D:\pgdata17\base\71251465\71251975	Hard Link	D:\pgdata\base\71251465\71251975	2/6/2025 3:00:00 PM
71252328	D:\pgdata17\base\71251465\71252328	Hard Link	D:\pgdata\base\71251465\71252328	2/6/2025 3:00:00 PM
71252331	D:\pgdata17\base\71251465\71252331	Hard Link	D:\pgdata\base\71251465\71252331	2/6/2025 3:00:00 PM
71252332	D:\pgdata17\base\71251465\71252332	Hard Link	D:\pgdata\base\71251465\71252332	2/6/2025 3:00:00 PM
71252338	D:\pgdata17\base\71251465\71252338	Hard Link	D:\pgdata\base\71251465\71252338	2/6/2025 3:00:00 PM
71252344	D:\pgdata17\base\71251465\71252344	Hard Link	D:\pgdata\base\71251465\71252344	2/6/2025 3:00:00 PM
71252345	D:\pgdata17\base\71251465\71252345	Hard Link	D:\pgdata\base\71251465\71252345	2/6/2025 3:00:00 PM
71252349	D:\pgdata17\base\71251465\71252349	Hard Link	D:\pgdata\base\71251465\71252349	2/6/2025 3:00:00 PM
71252352	D:\pgdata17\base\71251465\71252352	Hard Link	D:\pgdata\base\71251465\71252352	2/6/2025 3:00:00 PM
71252353	D:\pgdata17\base\71251465\71252353	Hard Link	D:\pgdata\base\71251465\71252353	2/6/2025 3:00:00 PM
71252357	D:\pgdata17\base\71251465\71252357	Hard Link	D:\pgdata\base\71251465\71252357	2/6/2025 3:00:00 PM
71252371	D:\pgdata17\base\71251465\71252371	Hard Link	D:\pgdata\base\71251465\71252371	2/6/2025 3:00:00 PM
71252376	D:\pgdata17\base\71251465\71252376	Hard Link	D:\pgdata\base\71251465\71252376	2/6/2025 3:00:00 PM
71252377	D:\pgdata17\base\71251465\71252377	Hard Link	D:\pgdata\base\71251465\71252377	2/6/2025 3:00:00 PM
71252408	D:\pgdata17\base\71251465\71252408	Hard Link	D:\pgdata\base\71251465\71252408	2/6/2025 3:00:00 PM
71252414	D:\pgdata17\base\71251465\71252414	Hard Link	D:\pgdata\base\71251465\71252414	2/6/2025 3:00:00 PM

LINK METHOD WORKS.
VIBE CEO JUST WASTES TIME IN PROD.

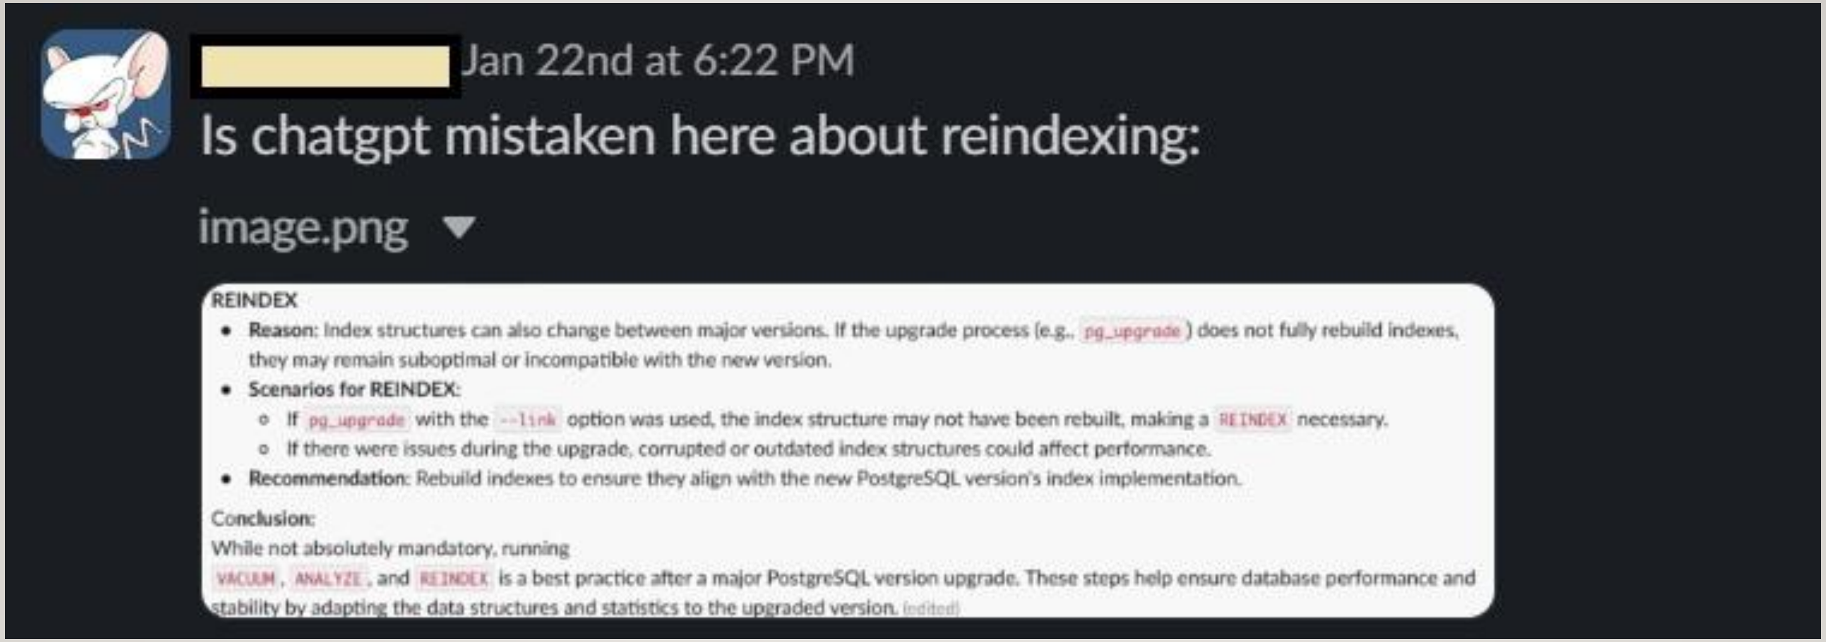

VIBE CEO ADDS MORE DOWNTIME REINDEXING.

EVEN CHATGPT FAILED TO ANSWER THIS!

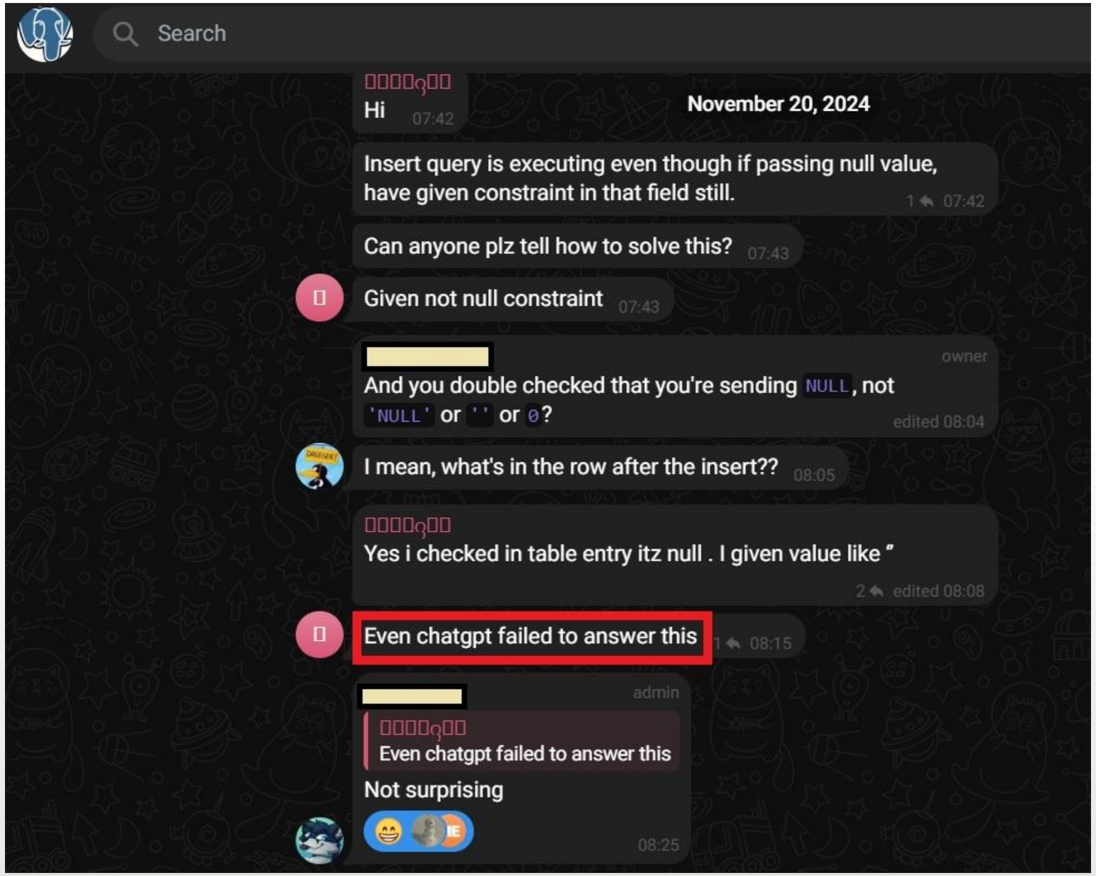

VIBE CEO HIRES PROMPT ENGINEER CTO TO TAKES UP ADVANCED TASK SUCH AS TTL INDEX FOR ARCHIVAL

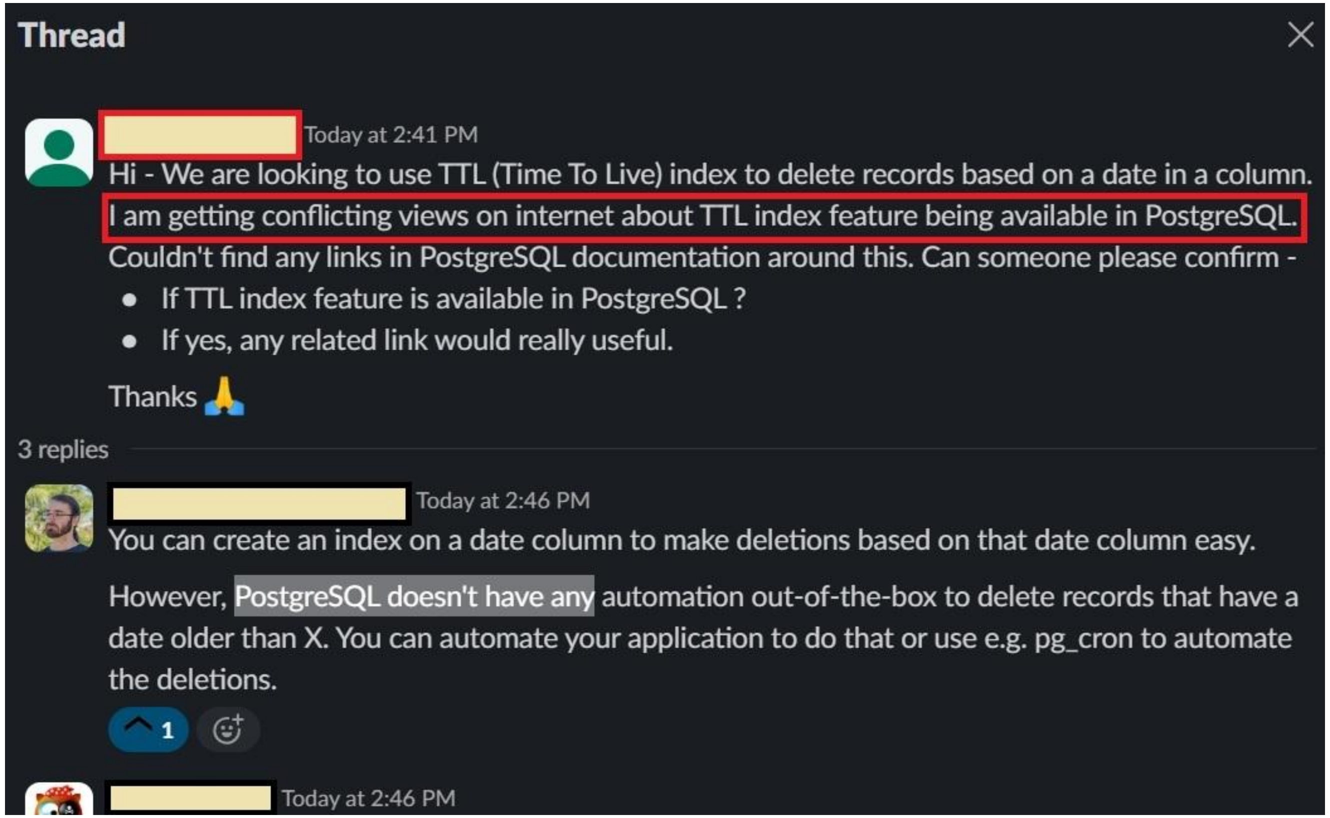

PROMT ENG CTO MOVES ON TO BIGGER CHALLENGES

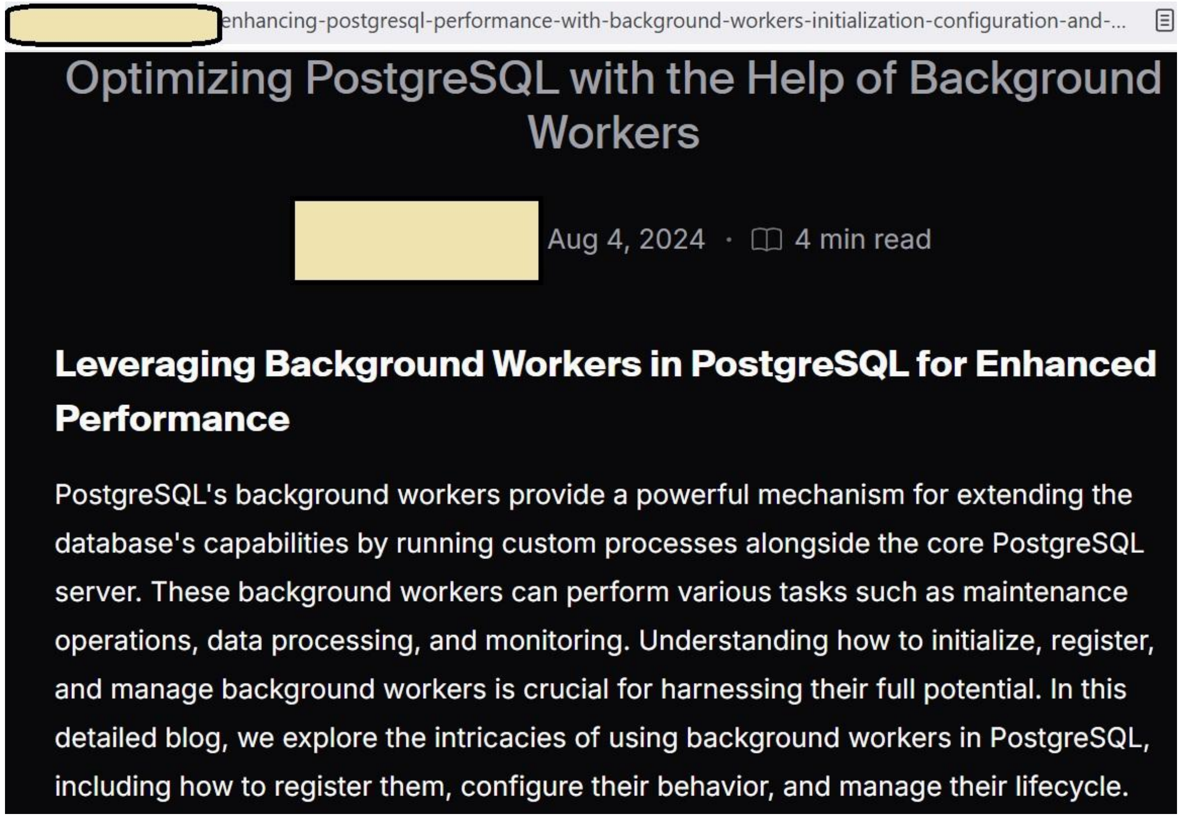

enhancing-postgresql-performance-with-background-workers-initialization-configuration-and-...

Optimizing PostgreSQL with the Help of Background Workers

Aug 4, 2024 · 4 min read

Leveraging Background Workers in PostgreSQL for Enhanced Performance

PostgreSQL's background workers provide a powerful mechanism for extending the database's capabilities by running custom processes alongside the core PostgreSQL server. These background workers can perform various tasks such as maintenance operations, data processing, and monitoring. Understanding how to initialize, register, and manage background workers is crucial for harnessing their full potential. In this detailed blog, we explore the intricacies of using background workers in PostgreSQL, including how to register them, configure their behavior, and manage their lifecycle.

enhancing-postgresql-performance-with-background-workers-initialization-configuration-and-...

Conclusion

Background workers in PostgreSQL provide a flexible and powerful way to extend database functionality. By understanding how to initialize, configure, and manage these processes, you can effectively leverage them to perform a variety of tasks, from maintenance operations to complex data processing. Properly using background workers can significantly enhance the performance and capabilities of your PostgreSQL deployment.

CRASHES AND CORRUPTS PROD DB

VIBE CEO CREATES A 2 NODE PATRONI HA CLUSTER (DCS, PG, PATRONI ALL ON TWO MACHINES)

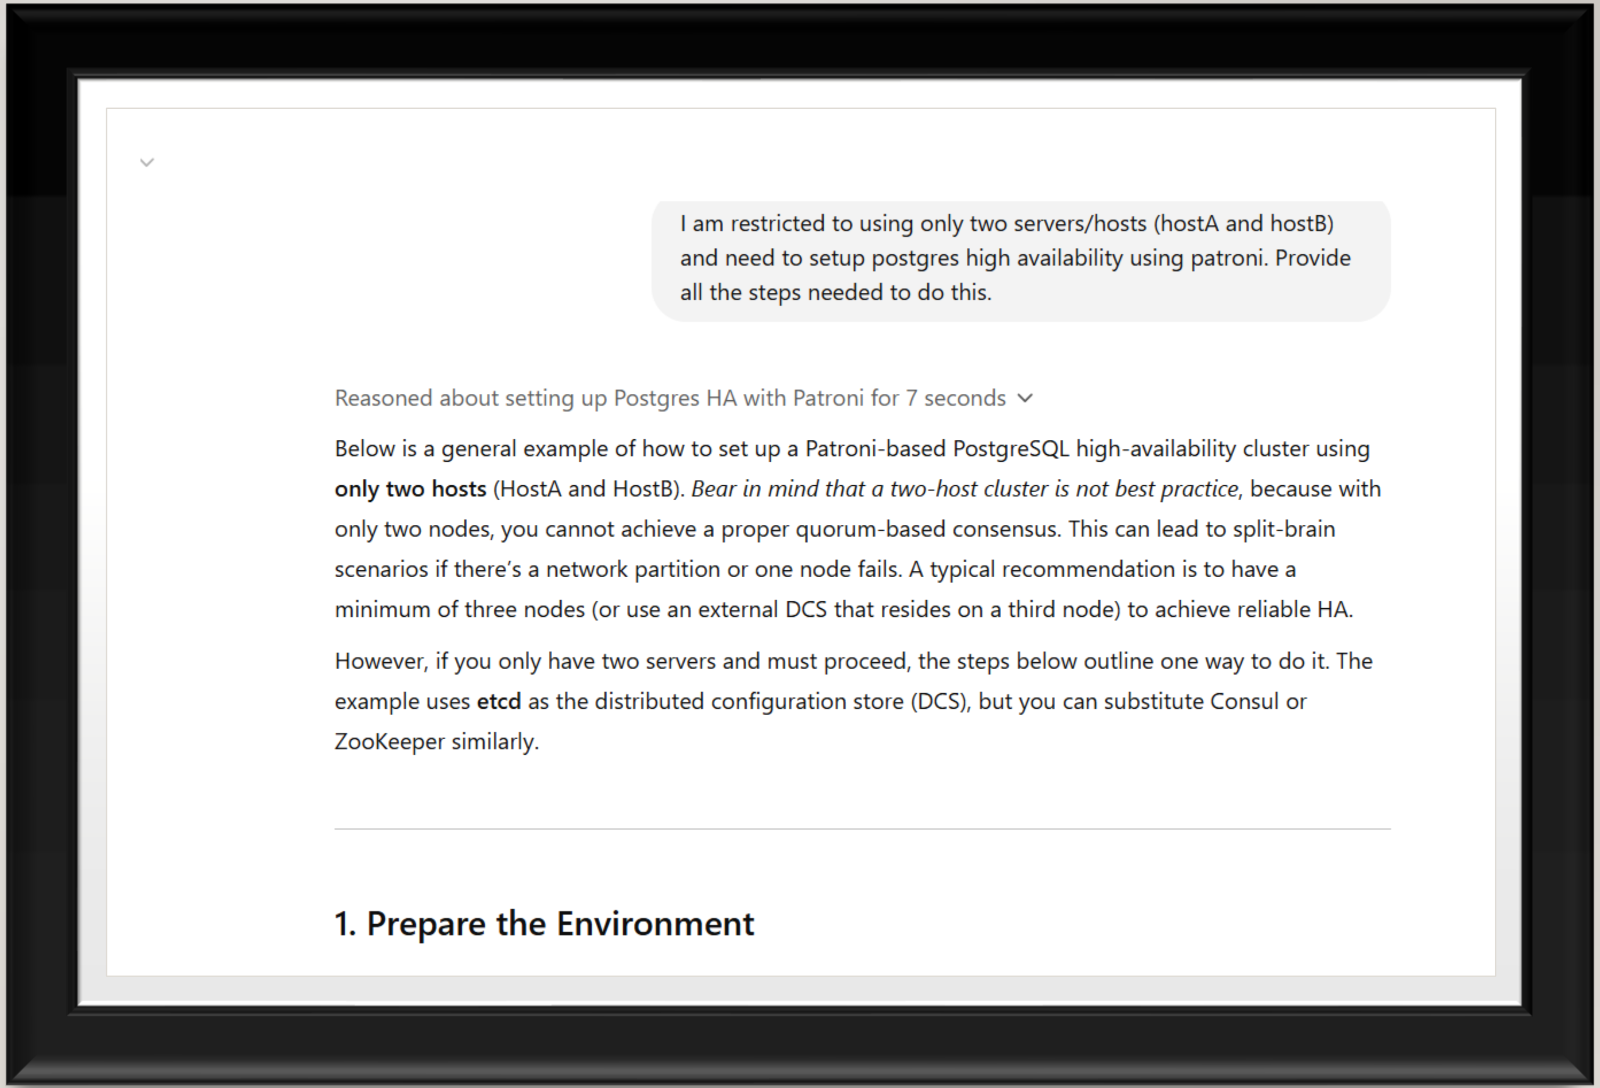

I am restricted to using only two servers/hosts (hostA and hostB) and need to setup postgres high availability using patroni. Provide all the steps needed to do this.

Reasoned about setting up Postgres HA with Patroni for 7 seconds

Below is a general example of how to set up a Patroni-based PostgreSQL high-availability cluster using **only two hosts** (HostA and HostB). *Bear in mind that a two-host cluster is not best practice*, because with only two nodes, you cannot achieve a proper quorum-based consensus. This can lead to split-brain scenarios if there's a network partition or one node fails. A typical recommendation is to have a minimum of three nodes (or use an external DCS that resides on a third node) to achieve reliable HA.

However, if you only have two servers and must proceed, the steps below outline one way to do it. The example uses **etcd** as the distributed configuration store (DCS), but you can substitute Consul or ZooKeeper similarly.

---

1. Prepare the Environment

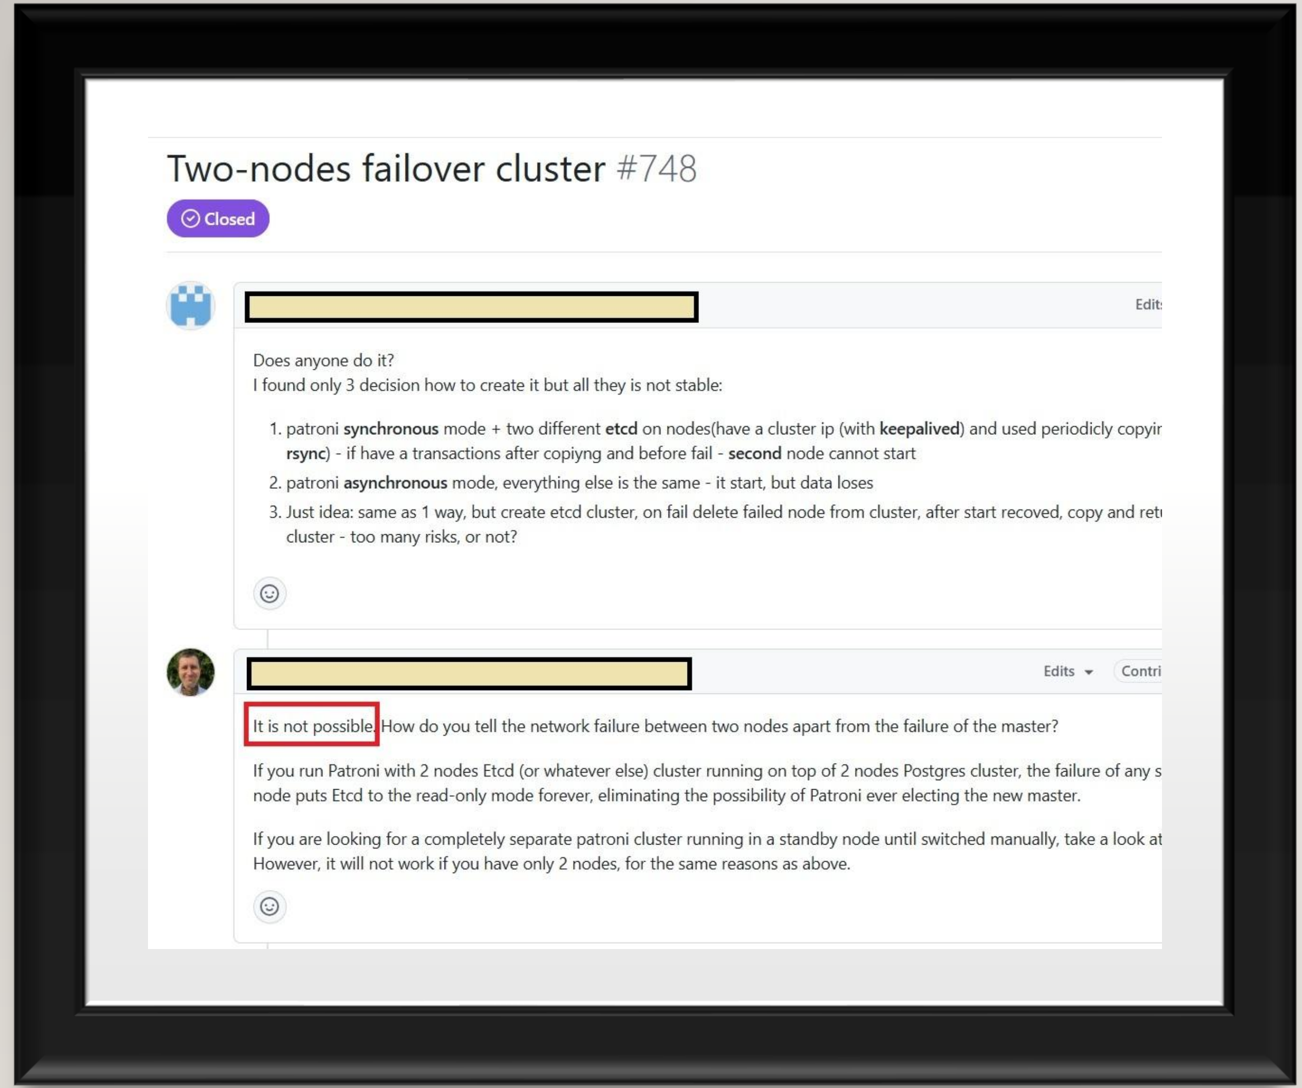

Two-nodes failover cluster #748

Closed

Edits

Does anyone do it?
I found only 3 decision how to create it but all they is not stable:

1. patroni **synchronous** mode + two different **etcd** on nodes(have a cluster ip (with **keepalived**) and used periodicly copyir **rsync**) - if have a transactions after copiyng and before fail - **second** node cannot start
2. patroni **asynchronous** mode, everything else is the same - it start, but data loses
3. Just idea: same as 1 way, but create etcd cluster, on fail delete failed node from cluster, after start recoved, copy and ret cluster - too many risks, or not?

Edits Contri

It is not possible. How do you tell the network failure between two nodes apart from the failure of the master?

If you run Patroni with 2 nodes Etcd (or whatever else) cluster running on top of 2 nodes Postgres cluster, the failure of any s node puts Etcd to the read-only mode forever, eliminating the possibility of Patroni ever electing the new master.

If you are looking for a completely separate patroni cluster running in a standby node until switched manually, take a look at However, it will not work if you have only 2 nodes, for the same reasons as above.

HUMAN RESPONSE

SPLIT BRAIN

VIBE CEO FINALLY PURCHASES EXPENSIVE SUPPORT CONTRACT FROM EXPERTS LIKE EDB/CYBERTEC/DATAEGRET/PERCONA

DATA CANNIBALISM

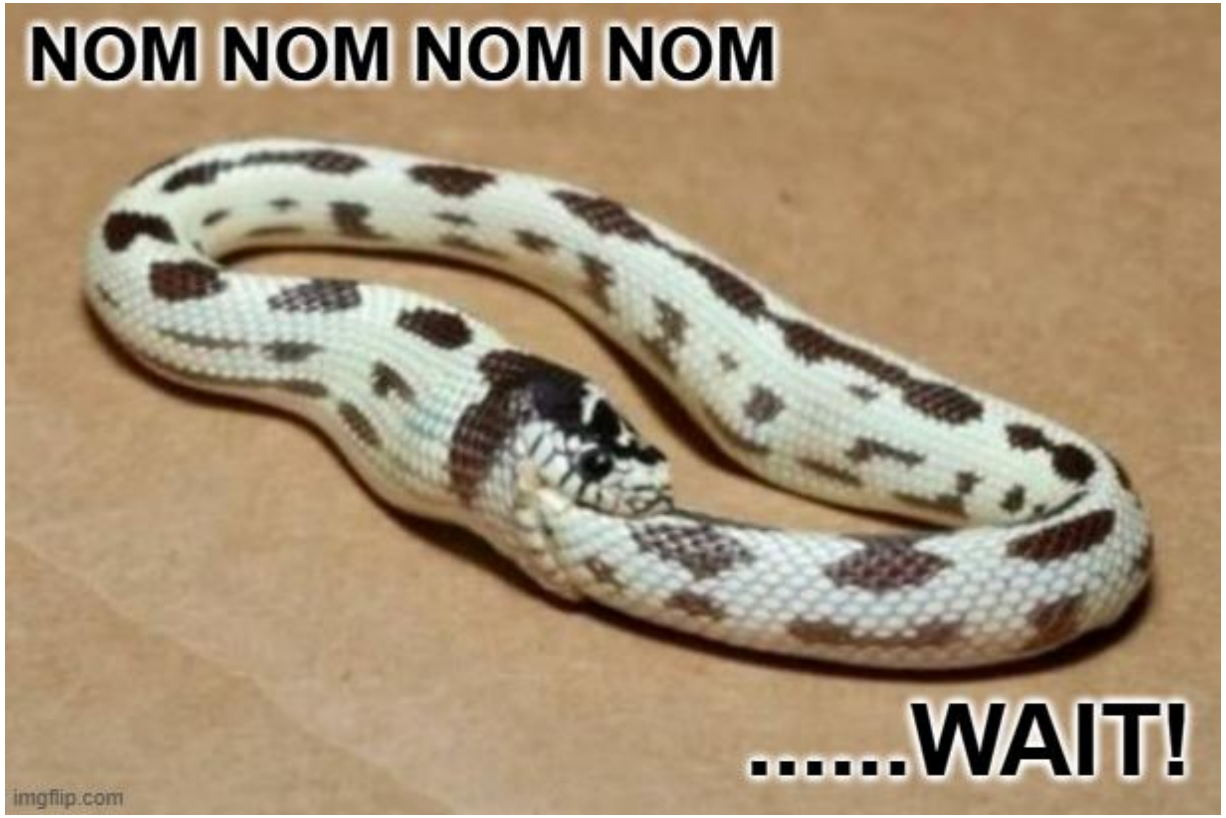

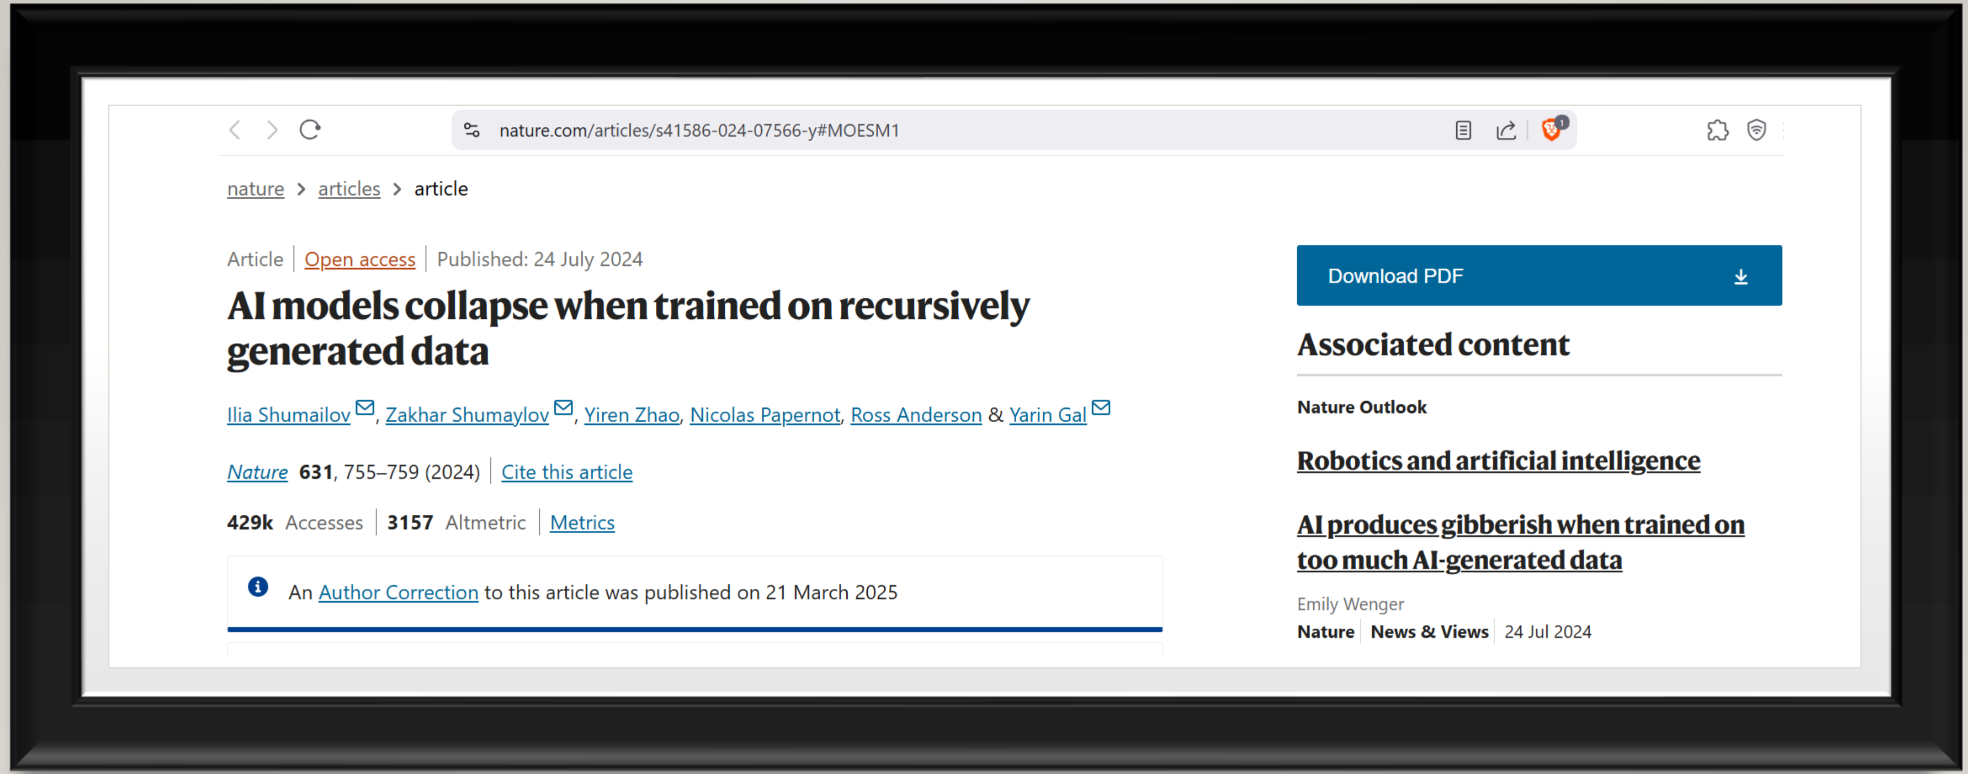

nature.com/articles/s41586-024-07566-y#MOESM1

nature > articles > article

Article | Open access | Published: 24 July 2024

AI models collapse when trained on recursively generated data

Ilia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson & Yarin Gal

Nature **631**, 755–759 (2024) | Cite this article

An Author Correction to this article was published on 21 March 2025

Download PDF

Associated content

Nature Outlook

Robotics and artificial intelligence

AI produces gibberish when trained on too much AI-generated data

Emily Wenger

Nature | **News & Views** | 24 Jul 2024

DATA CANNIBALISM

arXiv > cs > arXiv:2211.04325

Search...

Help | Ad

Computer Science > Machine Learning

[Submitted on 26 Oct 2022 (v1), last revised 4 Jun 2024 (this version, v2)]

Will we run out of data? Limits of LLM scaling based on human-generated data

Pablo Villalobos, Anson Ho, Jaime Sevilla, Tamay Besiroglu, Lennart Heim, Marius Hobbhahn

We investigate the potential constraints on LLM scaling posed by the availability of public human-generated text data. We forecast the growing demand for training data based on current trends and estimate the total stock of public human text data. Our findings indicate that if current LLM development trends continue, models will be trained on datasets roughly equal in size to the available stock of public human text data between 2026 and 2032, or slightly earlier if models are overtrained. We explore how progress in language modeling can continue when human-generated text datasets cannot be scaled any further. We argue that synthetic data generation, transfer learning from data-rich domains, and data efficiency improvements might support further progress.

Subjects: **Machine Learning (cs.LG)**; Artificial Intelligence (cs.AI); Computation and Language (cs.CL); Computer Vision and Pattern Recognition (cs.CV); Computers and Society (cs.CY)

Cite as: arXiv:2211.04325 **[cs.LG]**

(or arXiv:2211.04325v2 **[cs.LG]** for this version)

https://doi.org/10.48550/arXiv.2211.04325

DATA CANNIBALISM

DATA CANNIBALISM

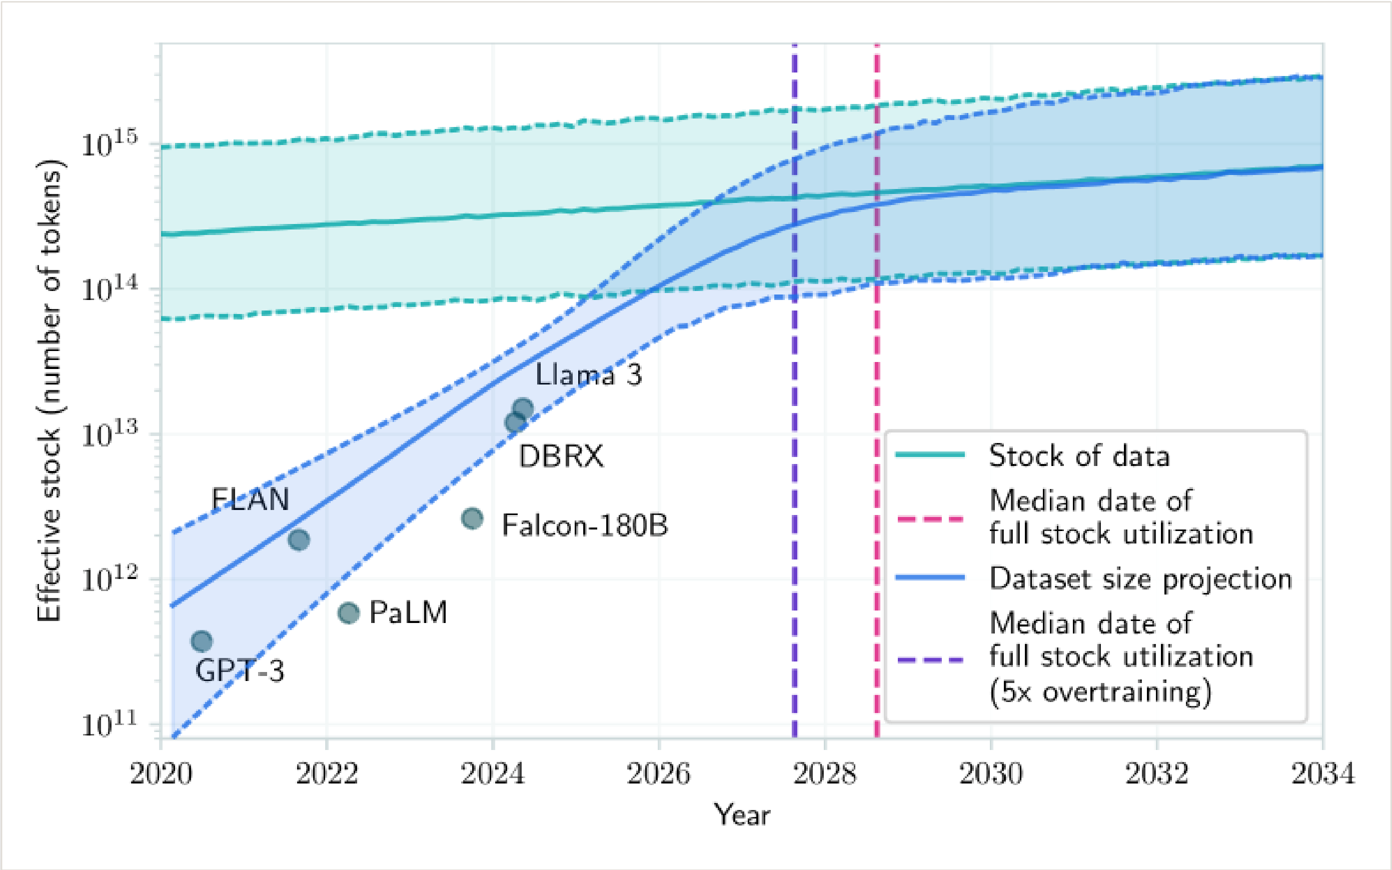

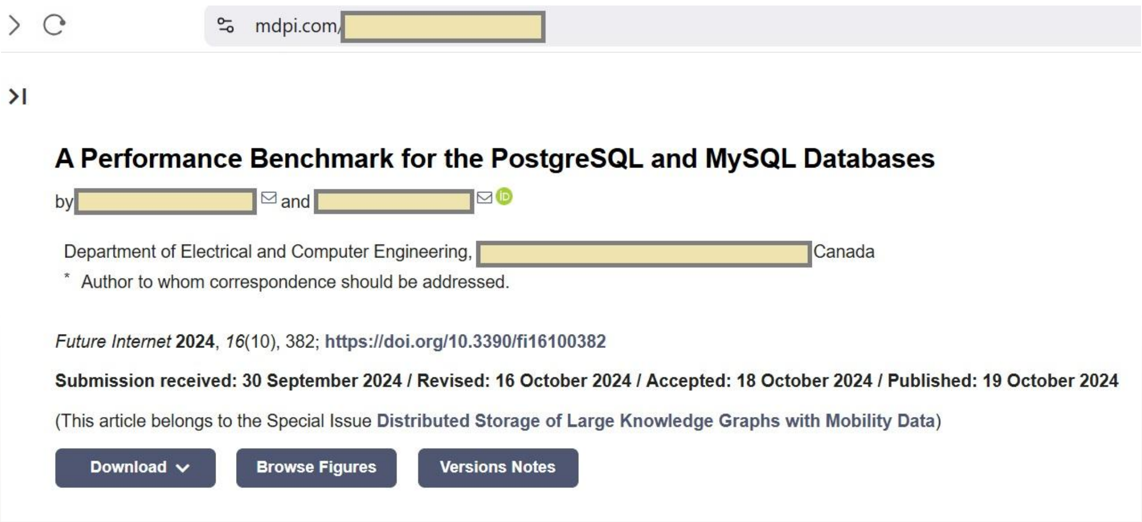

mdpi.com/[redacted]

A Performance Benchmark for the PostgreSQL and MySQL Databases

by [redacted] and [redacted]

Department of Electrical and Computer Engineering, [redacted] Canada

* Author to whom correspondence should be addressed.

Future Internet **2024**, *16*(10), 382; https://doi.org/10.3390/fi16100382

Submission received: 30 September 2024 / Revised: 16 October 2024 / Accepted: 18 October 2024 / Published: 19 October 2024

(This article belongs to the Special Issue **Distributed Storage of Large Knowledge Graphs with Mobility Data**)

Download Browse Figures Versions Notes

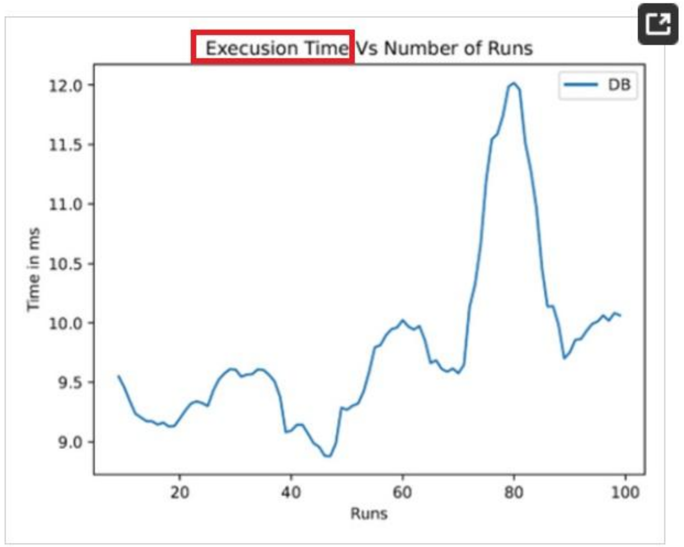

Figure 5. Select query execution time of MySQL for primary experiment one.

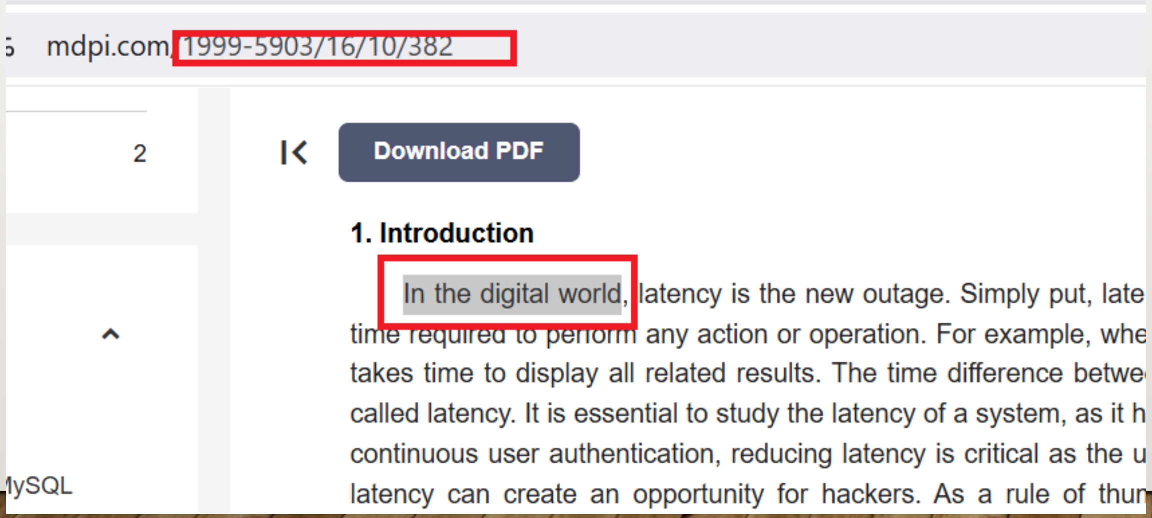

mdpi.com/1999-5903/16/10/382

Download PDF

2

1. Introduction

In the digital world, latency is the new outage. Simply put, late
time required to perform any action or operation. For example, whe
takes time to display all related results. The time difference betwe
called latency. It is essential to study the latency of a system, as it h
continuous user authentication, reducing latency is critical as the u
latency can create an opportunity for hackers. As a rule of thun

lySQL

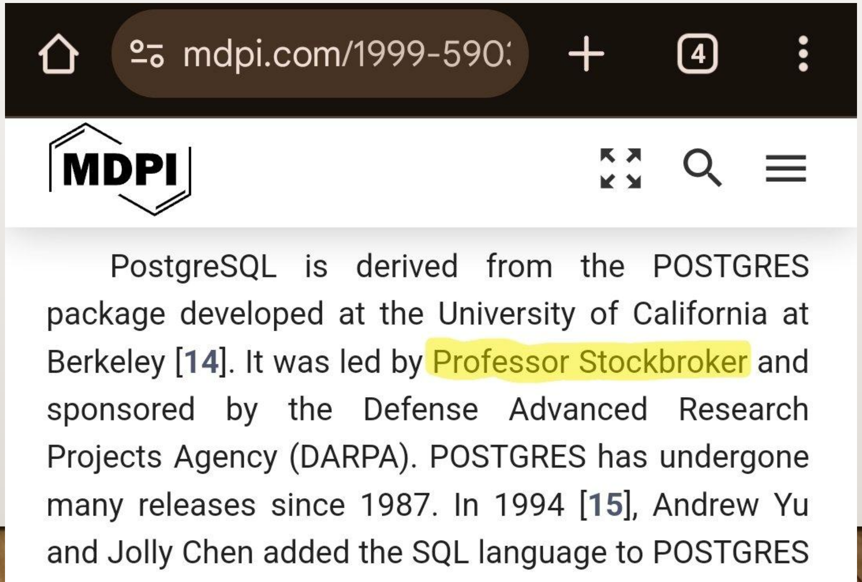

mdpi.com/1999-590:

MDPI

PostgreSQL is derived from the POSTGRES package developed at the University of California at Berkeley [14]. It was led by Professor Stockbroker and sponsored by the Defense Advanced Research Projects Agency (DARPA). POSTGRES has undergone many releases since 1987. In 1994 [15], Andrew Yu and Jolly Chen added the SQL language to POSTGRES


KNOW THE WORK RULES
APPROPRIATE
Scaling is over and LLMs are a dead end
AWWW, YOU'RE SWEET.
INAPPROPRIATE
Scaling is over and LLMs are a dead end
HELLO, HUMAN RESOURCES?!
imgflip.com

NARROW AI WITH LIMITED SCOPE OF CREATIVITY.

- DB Parameter tuning (DBTune)
- Auto-Indexing

  (PGAnalyze)
- Ops Automation (AWS Devops Guru, Xata Agent)

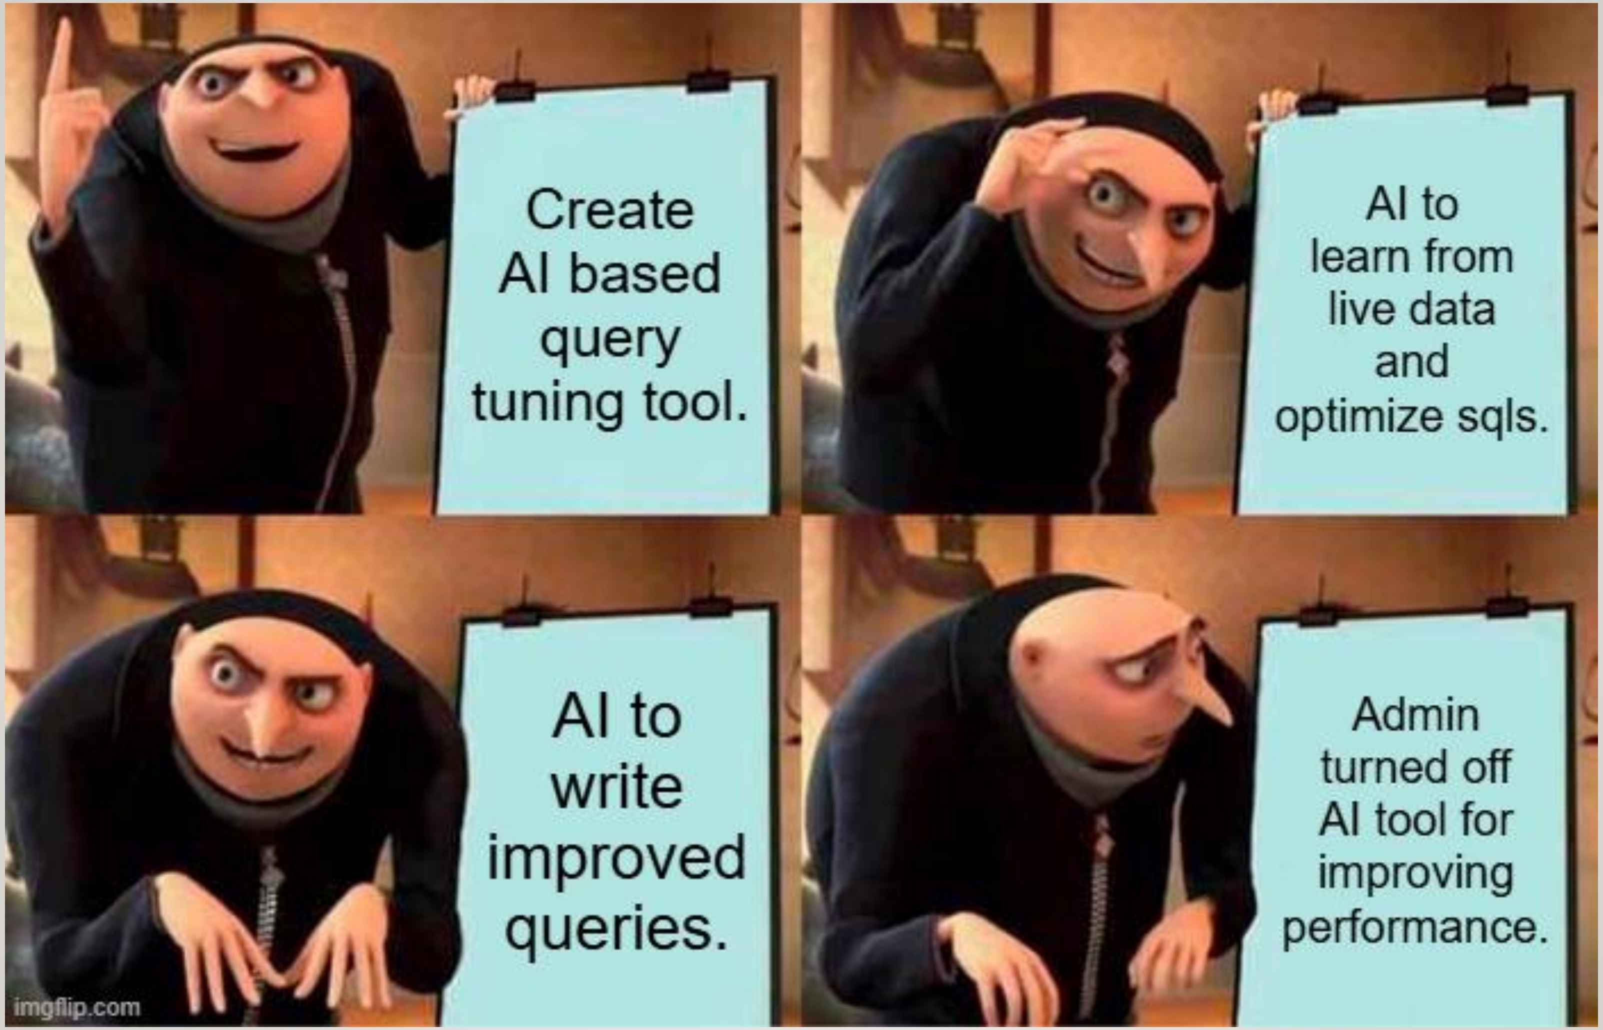
Create AI based query tuning tool.
AI to learn from live data and optimize sqls.
AI to write improved queries.
Admin turned off AI tool for improving performance.
imgflip.com

Future of AI : Expectation

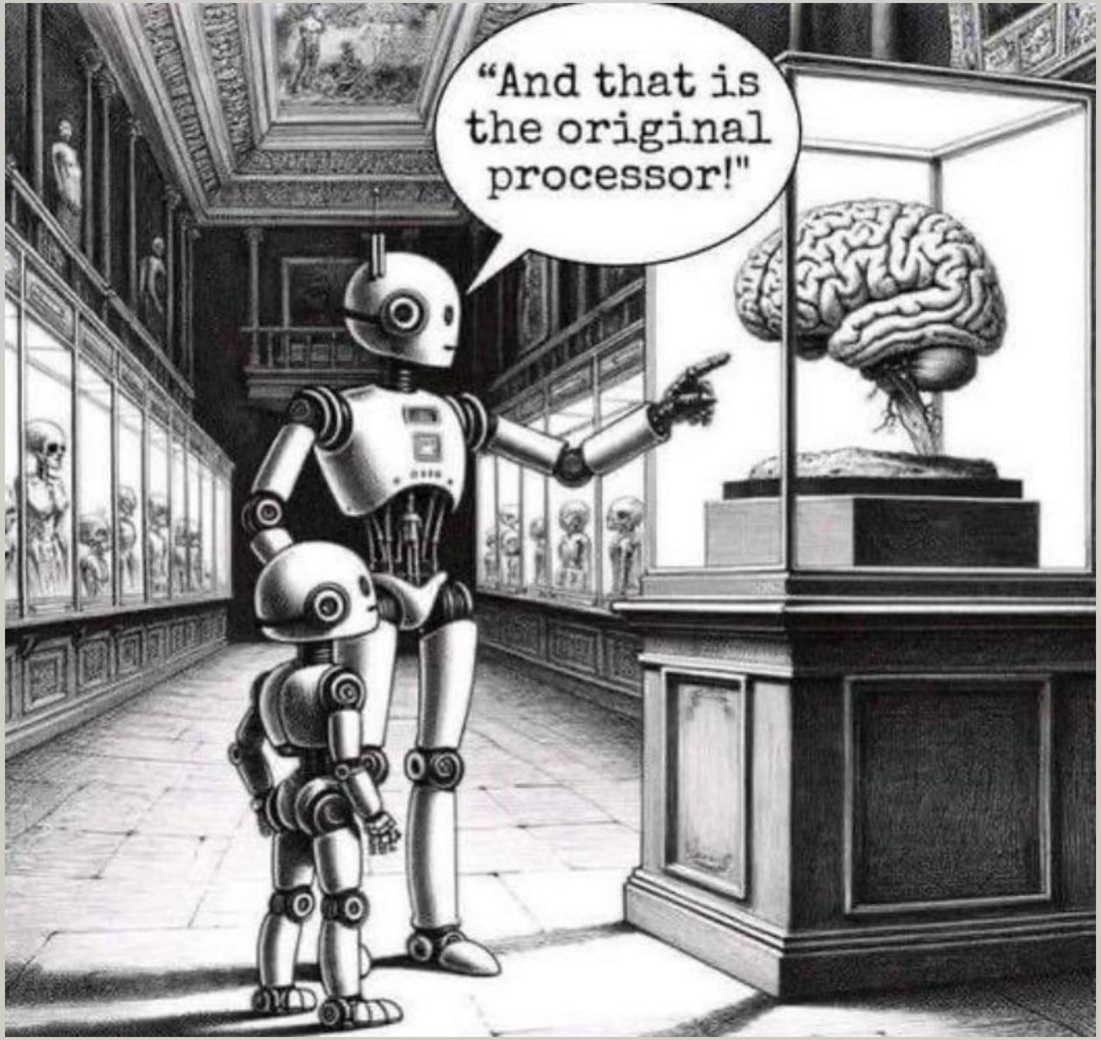

Future of AI : Reality

The Register® | Hewlett Packard Enterprise

Generative AI is not replacing jobs or hurting wages at all, economists claim

'When we look at the outcomes, it really has not moved the needle'

Thomas Claburn   Tue 29 Apr 2025 // 07:18 UTC

Instead of depressing wages or taking jobs, generative AI chatbots like ChatGPT, Claude, and Gemini have had almost no significant wage or labor impact so far – a finding that calls into question the huge capital expenditures required to create and run AI models.

In a working paper released earlier this month, economists Anders Humlum and Emilie Vestergaard looked at the labor market impact of AI chatbots on 11 occupations, covering 25,000 workers and 7,000 workplaces in Denmark in 2023 and 2024.

Many of these occupations have been described as being vulnerable to AI: accountants, customer support specialists, financial advisors, HR professionals, IT support specialists, journalists, legal professionals, marketing professionals, office clerks, software developers, and teachers.

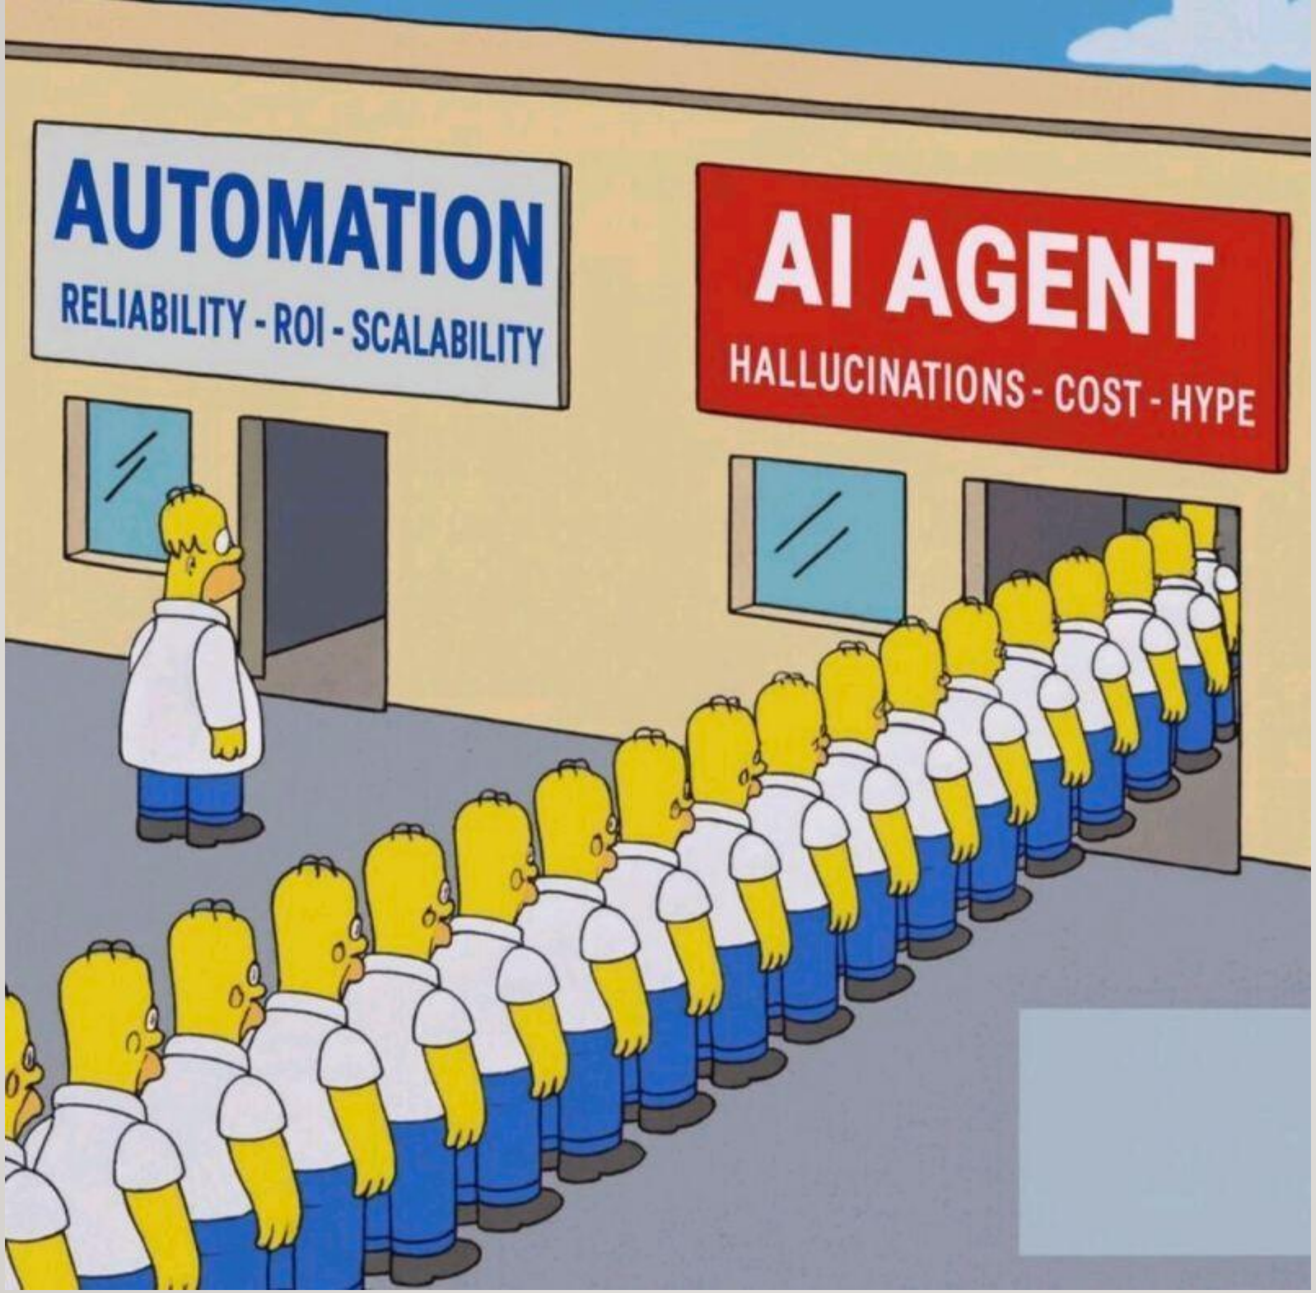
AUTOMATION
RELIABILITY - ROI - SCALABILITY
AI AGENT
HALLUCINATIONS - COST - HYPE

Artificial Intelligence, Machine Learning, Data Scien... Join
Mayank Jain • 3rd+
1d •

📌 The hard truth about AI agents

They're often just glorified 'If > then' statements.

⭕ We're repeating history. Remember 15 years ago when Machine Learning was the answer to everything? Churn prediction, sales forecasting, cross-sells.

The reality? 80% of the value came from simple heuristics:

If user inactive for 14 days > send email.

⭕ If margin too low > stop campaign.

Now, we're doing the same thing with AI Agents. AI is game-changing, but for a majority of use cases, AI agents are:

⏩ Overly complex
⏩ Sources of failure
⏩ Months-long projects that die in the sandbox

⭕ Stop overengineering! The blueprint for real value:

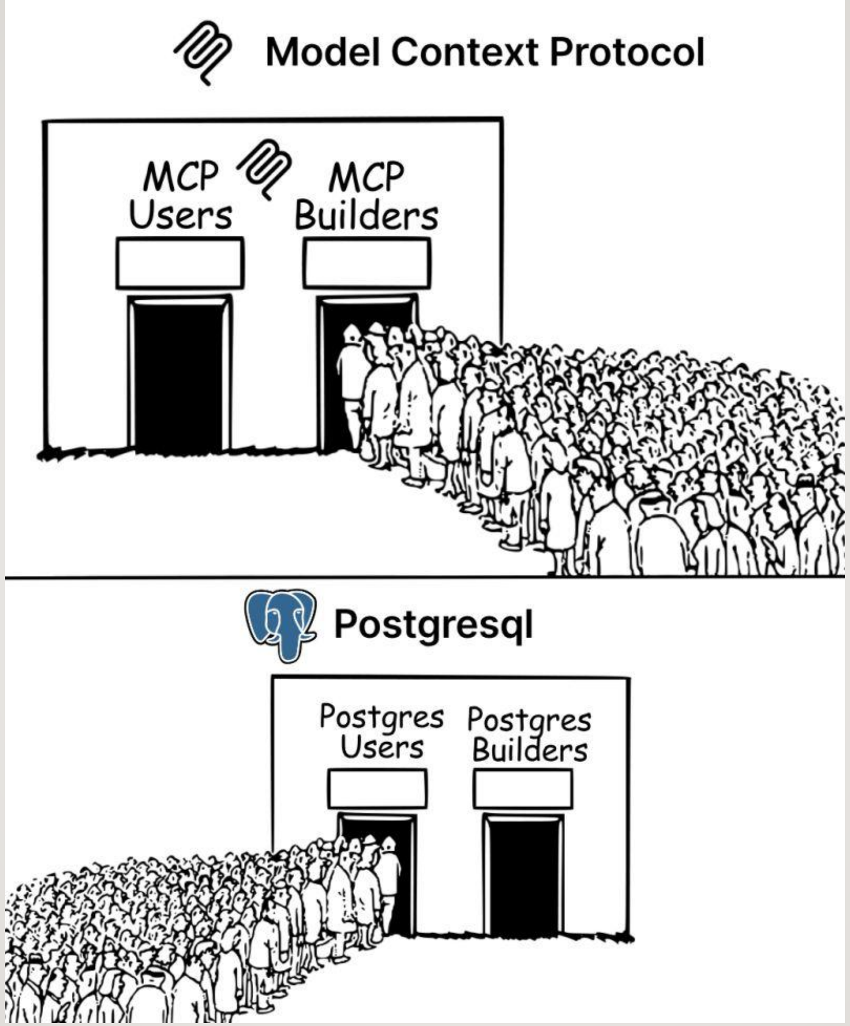
Model Context Protocol
MCP Users
MCP Builders
Postgresql
Postgres Users
Postgres Builders

Who can AI really replace?

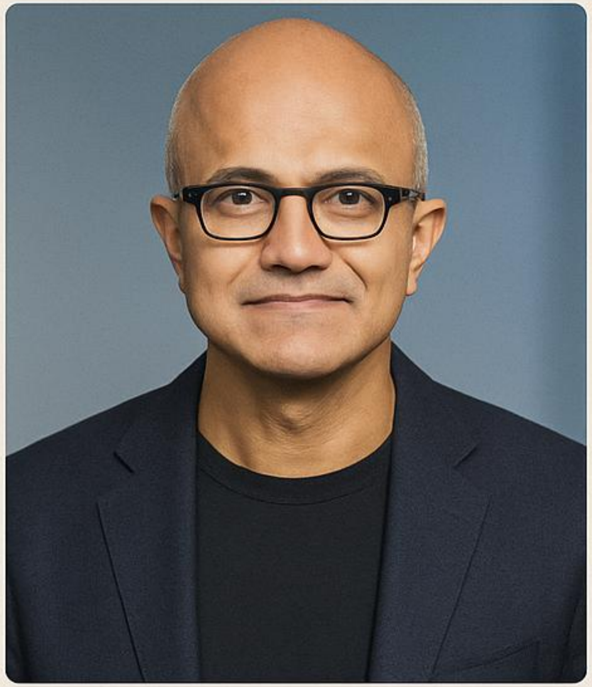

CEOs:
AI can make stuff up

Politicians:
LLMs lie constantly

REFERENCES

- Postgres Slack (Unofficial) https://pgtreats.info/slack-invite
- Postgres Telegram (Unofficial) : https://t.me/postgreschat & https://t.me/pg_sql
- AI Hallucinations - Technology Review Article
- AI Model Collapse Paper, Nature
- AI running out of Human data
- AI Slop and Destruction of Knowledge

THANK YOU