

SPOTLIGHT ON DATABASES

Ex-BigQuery exec and Motherduck CEO: For some users, the answer is to think small



POST A COMMENT

TIP US OFF

Similar topics

Analytics

Big Data

Database

Send us news

INTERVIEW Jordan Tigani made his name as an engineer leading the team behind BigQuery, Google's data warehouse, which was among a group of systems to transform the market by separating storage and compute. Despite helping win global customers including Vodafone, and taking on big-name rivals Snowflake, Azure Synapse and AWS Redshift, he's begun to think the approach has run out of "magic beans" for some users.

Now founder and CEO of MotherDuck - which built a serverless analytics system based around DuckDB - he finds himself eschewing the virtues of scale-out systems for the benefits a lightweight in-process OLAP database affords. You can catch all of this detail and more in our interview with him below.

Tigani recalled coming across DuckDB while he was chief product officer at SingleStore - his first job after Google.

"We started seeing DuckDB appearing on performance reports and giving us a run for the money. Not beating us, but it was still surprising. Then I started poking at it and I realised it could do some stuff that we couldn't do in SingleStore, that Snowflake couldn't do either, and that was pretty interesting," he says.

The brainchild of academics at Amsterdam's Centrum Wiskunde & Informatica

Oh no, you're thinking, yet another cookie pop-up. Well, sorry, it's the law. We measure how many people read us, and ensure you see relevant ads, by storing cookies on your device. If you're cool with that, hit "Accept all Cookies". For more info and to customize your settings, hit "Customize Settings".

Customize Settings

Free and open source under the MIT License.

Tigani's pitch to co-authors Hannes Mühleisen and Mark Raasveldt was to build a cloud-based serverless product around DuckDB.

The important thing about DuckDB is that it's a scale-up system, contrary to the received wisdom of the last 15 years of analytics, embodied in the first generation of so called big data systems based on the Hadoop Distributed File System. The scale-out approach led to the boom in cloud-based analytics systems like Snowflake, BigQuery, Synapse and Redshift.

MORE CONTEXT

The nodes have it in the Great DB debate: Reg readers pick graph

Cassandra 4.1 promises dev guardrails and pluggable storage

MotherDuck scores \$47.5m to prove scale-up databases are not quakers

Google wants to copy-paste your mainframe applications into its cloud

"At some point, there are no magic beans left. There will be a convergence around the performance of all these systems. Snowflake, Redshift, BigQuery, and Synapse are all probably within a factor of two in performance right now. For the most part, that's not what should drive people to use one of these systems versus the other," Tigani says.

Added to that, the server machinery and infrastructure around these systems can "be really quite cruffy," he adds.

"DuckDB has been able to kind of strip all that away by being an in-process database, and that means that you basically can marshal data in and out of your application, or your data frames, with the minimum of data movements," he says.

DuckDB is not designed to replace company-wide enterprise data warehouse systems like BigQuery but - by courting data scientists, machine learning engineers and developers, it could still have plenty of space to fly. @

Former Google veteran talks to EI Reg about trends in big data of the past decade

Lindsay Clark Tue 21 Mar 2023 11:30 UTC

Government hid anticompetitive ad practices

Aggrieved investors claim tech giant fiddled to inflate stock prices... and it backfired

Personal Tech17 Mar 2023 | 15

Requirement for Google Reader, dead for a decade but not forgotten

SPECIAL REPORT RSS pioneer Dave Winer sees Mastodon as a positive sign

Networks14 Mar 2023 | 50

Google reminds everyone it too can launch a ChatGPT-like chatbot ... waiting list

Meanwhile, Bing can now output images, Adobe touts shiny art platform Firefly

AI+ML21 Mar 2023 | 9

Aerospike targets Java Spring devs with support for the popular framework

Database shortcuts built in popular environment in hopes of scoring devs who want an easy life

Spotlight on Databases22 Mar 2023 | 4

Google suspends top Chinese shopping app Pinduoduo

Alleges it's infected with malware - but not the version in its own digital tat bazaar

Security21 Mar 2023 | 4

The nodes have it in the Great DB debate: Reg readers pick graph

REGISTER DEBATE Industry upstart preferred... for now

Spotlight on Databases10 Mar 2023 | 26

Google taps Fastly to make cookie-free adtech FLEDGE fly

Online ad colossus hopes it can still make money when users want privacy

Personal Tech16 Mar 2023 | 29

What happens when you free DevOps pros from cost management?

Here's how to spot the difference

Sponsored Feature

Now Google to shove its answer to ChatGPT into Gmail, Docs, apps via APIs, more

PaLM springs forth

AI+ML14 Mar 2023 | 9

Google stops selling its biz-grade augmented reality specs

In case of apathy, break Glass

Software16 Mar 2023 | 33

Hong Kong's state-sponsored SEO on national anthem strikes the right note

Can't get Google to finish the job, though

Legal20 Mar 2023 | 11

Don't Be Evil, a gaggle of Googlers tell CEO Pichai amid mega layoffs

Open letter to top brass signed by 1,400+ staff offers 5-point plan to improve redundancy process

OP+Prm20 Mar 2023 | 26

About Us

Contact us

Advertise With us

Who we are

Our Websites

The Next Platform

DevClass

Blocks and Files

Your Privacy

Cookies Policy

Your Consent Options

Privacy Policy

T's & C's

**situation
publishing**

Copyright. All rights reserved © 1998–2023