

November 16, 2022

Is Big Data Dead? MotherDuck Raises \$47M to Prove It

Jaime Hampton



Big data is dead. Or so says MotherDuck, the builder of a serverless analytics platform based on DuckDB. The company's founders say they learned from real-world users that a vast majority of workloads do not require the high overhead costs of big data distributed computing thanks to recent hardware advances.

"The fact is, 'Big Data' is dead; the simplicity and the ease of making sense of your data is a lot more important than its size," said Jordan Tigani, CEO and co-founder of [MotherDuck](#) in a release.

Tigani's company has just raised \$47.5 million and is partnering with [DuckDB Labs](#) (founded by DuckDB's creators) to build a serverless cloud analytics platform based on DuckDB. The company says the funding will be used to further this collaboration, as well as build out its engineering and GTM teams.



"Laptops today are faster than a data warehouse. With advances in hardware, distributed computation is no longer necessary for most workloads," said Tigani. "Cloud data vendors are focused on the performance of 100TB queries, which is not only irrelevant for the vast majority of users, but also distracts from vendors' ability to deliver a great user experience. We are taking the power of DuckDB and combining it with serverless analytics to help scale up and scale down with ease."

[DuckDB](#) is an open source, in-process database similar to SQLite for analytics workloads. According to MotherDuck, the SQL OLAP database management system has garnered widespread adoption based on its ability to run everywhere (browsers included), query data from anywhere without preloading it, and execute quick analytical queries based on up-to-date academic research. OLAP workloads are complex with long-running queries that process significant portions of a stored dataset, and changes to the data are made with several rows being appended, or large portions of tables being changed or added at the same time, according to DuckDB.



"To efficiently support this workload, it is critical to reduce the amount of CPU cycles that are expended per individual value. The state of the art in data management to achieve this are either vectorized or just-in-time query execution engines. DuckDB contains a columnar-vectorized query execution engine, where queries are still interpreted, but a large batch of values (a "vector") are processed in one operation," says DuckDB's [website](#). "This greatly reduces overhead present in traditional systems such as PostgreSQL, MySQL or SQLite which process each row sequentially. Vectorized query execution leads to far better performance in OLAP queries."

In a company [blog](#), DuckDB Labs commented on the vision of the partnership with MotherDuck: "When the first ideas that eventually led to DuckDB were thrown around, we went against the prevailing wisdom in both industry and research that only massive scale and distributed data processing would be the way forward. From our interactions with data practitioners, we became convinced that while massive datasets exist, they are mostly found in organizations that already have the technological expertise to handle them anyway. We bet on efficient and ergonomic single-node analytics, and we are very happy that the MotherDuck team shares this vision, especially given the team's background."

DuckDB Labs was founded by Hannes Mühleisen and Mark Raasveldt to provide services and development for DuckDB. Mühleisen and Raasveldt were researchers in the Database Architectures research group at Centrum Wiskunde & Informatica (CWI) when they released the first version of DuckDB in 2019.

"DuckDB got its name because I used to have a pet duck," Mühleisen said in a [CWI-authored profile](#) on the company. "Ducks are amazing animals. They can fly, walk and swim, and they are quite resilient to environmental challenges. So, they are the perfect mascot for a versatile and resilient data management system."

Interest in DuckDB seems to be growing. According to MotherDuck, DuckDB's [DB Engines](#) score is growing at 40% each month, while its Python distribution sees 400K downloads in the same time.

MotherDuck's \$47.5 million in funding is comprised of a \$35 million Series A round led by Andreessen Horowitz that follows a \$12.5 million seed round led by Redpoint Ventures, bringing the total valuation of the company to \$175 million. Other investors include Madrona, Amplify Partners, and Altimeter.

"We see tremendous potential in MotherDuck – not just in the market they represent, but in the caliber of talent that is building this game-changing platform," said Tomasz Tunguz at Redpoint Ventures. "We're excited to partner with the team and bring the power of DuckDB to more people than ever before."



Interest in DuckDB is growing, as evidenced by this meme found on Twitter.

Related Items:

[Big Growth Forecasted for Big Data](#)

[Three Ways to Connect the Dots in a Decentralized Big Data World](#)

[The History of Data Science: From Cave Paintings to Big Data](#)



Applications: Data Management, Enterprise Analytics

Technologies: Cloud

Vendors: DuckDB Labs, MotherDuck

Tags: columnar database, DuckDB, DuckDB Labs, MotherDuck, OLAP

THIS JUST IN MOST READ

November 17, 2022

- ▶ [CloudFabrix Launches Composable In-place Search and Log Intelligence](#)
- ▶ [Scality Ships RING9, Software for Hybrid-cloud Data Storage](#)
- ▶ [Open Source Continuous ML Platform Sematic Raises \\$3M Seed](#)
- ▶ [AT&T Accelerates AI Transformation and Speed to Market Up to 70% with H2O.ai](#)
- ▶ [Snorkel AI Accelerates Foundation Model Adoption with Data-centric AI](#)
- ▶ [Grafana Labs Partners with Google Cloud to Advance Cloud Observability](#)

November 16, 2022

- ▶ [Alluxio Reimagines Architecture for Multi-Tenant Environments at Scale](#)

SPONSORED PARTNER CONTENT

[Informatica Named a Leader in Data Integration for 17 Years in a Row](#)

[Cloud First. Data Always.™](#)

[Everyday AI, Extraordinary People](#)

[Dataiku Makes the Use of Data and AI an Everyday Behavior](#)

[Data Fabrics as the best path for Enterprise Data Integration](#)

[Dataiku connects data and does through Everyday AI](#)

[Leaving Legacy ETL Behind](#)

[Streamline Lakehouse Analytics with Matillion and Databricks SQL](#)

[Close the Information Gap: How to Succeed at Analytics in the Cloud](#)

[Who wins the hybrid cloud?](#)

LEADING SOLUTION PROVIDERS



TABOR NETWORK



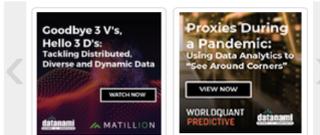
SPONSORED WHITEPAPERS

[3 Ways to Boost Productivity with Hyperautomation](#)

[How to Solve the Top 10 Cloud Data Lake Challenges](#)

▶ [View the Whitepaper Library](#)

SPONSORED MULTIMEDIA



[Goodbye 3 V's, Hello 3 D's: Tackling Distributed, Diverse and Dynamic Data](#)

[Proxies During a Pandemic: Using Data Analytics to See Around Corners](#)

CONTRIBUTORS



