SACRA

MARKET REPORT

Databricks at $2.4B ARR growing 60%

TEAM

Jan-Erik Asplund
Co-Founder
jan@sacra.com

DISCLAIMERS

Published on **Jul 07th, 2024**

Databricks at $2.4B ARR growing 60%

By **Jan-Erik Asplund**

TL;DR: Sacra estimates that <u>Databricks</u> hit $2.4B in annualized revenue in June 2024, up 60% year-over-year, as it retools as an AI company to take on Snowflake (NYSE: SNOW). <u>For more, check out our full Databricks report</u>.
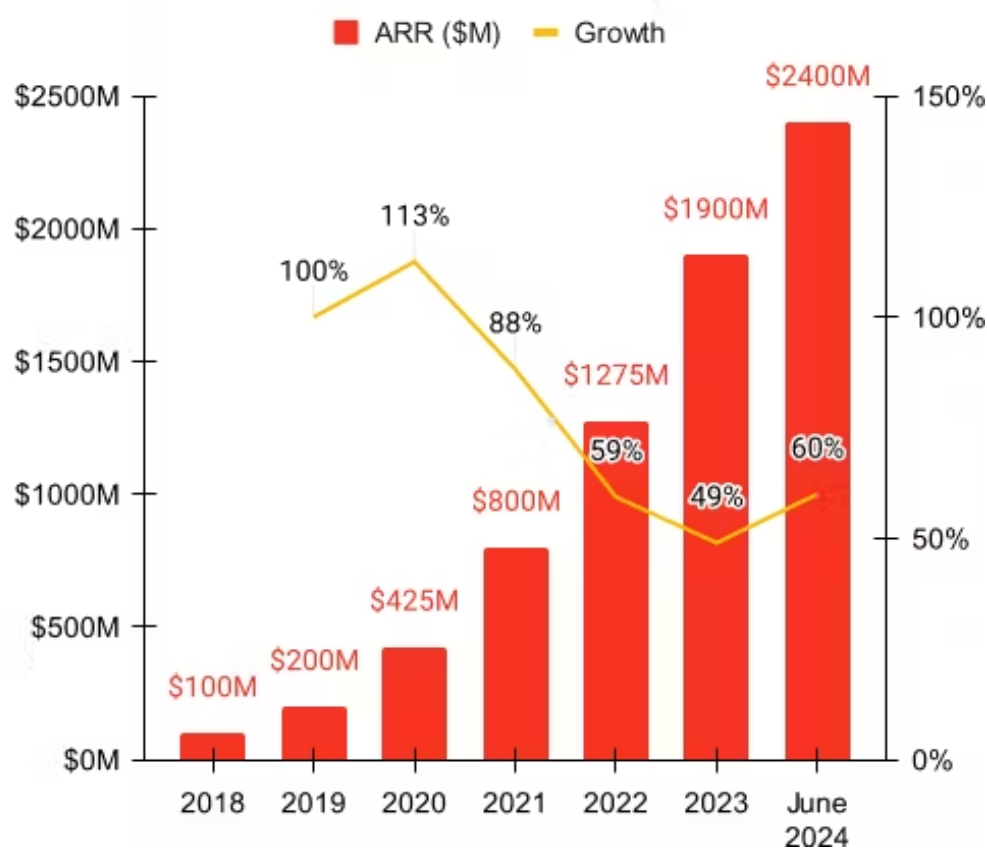
Key points from our research:

- **Sacra estimates that <u>Databricks</u> hit <u>$2.4B in annualized revenue in June 2024</u>, up 60% year-over-year, for a 18x multiple on their $43B valuation (from September 2023), challenging the narrative of a slowdown in enterprise tech spending**. Compare to Snowflake at $3.2B in annualized revenue in Q1'24, up 34%, on a $47.8B market cap for a 15x forward revenue multiple—note that Databricks's Snowflake

competitor, Databricks SQL, hit $400M in annualized revenue in 2024 (up 300% year-over-year).

- **Databricks'** **$1.3B acquisition of MosaicML in 2023 (which trains LLMs at 40% of the price of OpenAI) has further positioned the company as a formidable player in the enterprise AI market, while their open source LLMs are increasingly gaining traction on Hugging Face, from DBRX Instruct (40K+ downloads per month) to Dolly v2 3B (350K+)**. With their push into AI, Databricks is likely unprofitable accounting for their aggressive spend on R&D at 33% of revenue, significantly higher than the 19% average for its public cloud peer group.

- **Databricks'** **Unity Catalog is the wedge product that's become core to their AI strategy, becoming a category leader in the sparsely-populated data-governance-for-AI space with 10,000+ enterprise customers in less than two years after launch**. Unity Catalog provides a unified platform for managing access to datasets—key with the rise of models like Claude 3.5 Sonnet that are empowering non-technical product managers to ship code—and a wedge from which Databricks can cross-sell customers into MosaicML, AI training, and other products.

For more, check out this other research from our platform:

- Databricks (dataset)
- Anthropic (dataset)
- OpenAI (dataset)
- Scale (dataset)
- Hugging Face (dataset)
- OpenAI vs. Anthropic vs. Cohere
- Salesforce, Amplitude, and the fat data layer in B2B SaaS
- Sean Lynch, co-founder of Census, on reverse ETL's role in the modern data stack
- Earl Lee, co-founder and CEO of HeadsUp, on the modern data stack value chain
- Brian Whalley, Co-Founder of Wonderment, on Klaviyo's product-market fit
- George Xing, co-founder and CEO of Supergrain, on the future of business intelligence

- [Conor McCarter, co-founder of Prequel, on the data integration market](#)
- [Julia Schottenstein, Product Manager at dbt Labs, on the business model of open source](#)