PRIVATE MARKETS RESEARCH · MARCH 2026

Databricks
The backbone of enterprise AI

A private markets investment report.

$5.4B	$150.8B	>65%	20,000+
REVENUE RUN-RATE	CURRENT VALUATION	GROWTH YOY	ENTERPRISE CUSTOMERS

PRICE PER SHARE	MIN. INVESTMENT	SHARE CLASS	TRANSACTION
$212.80	$5,000	Series L Preferred	Primary

VESTED

01 The Story

Here is a problem every large company has. They have mountains of data sitting in different places, some in a data lake for cheap storage, some in a data warehouse for running reports, and AI models trained somewhere else entirely. None of it talks to each other cleanly. Moving data between systems is slow, expensive, and breaks things.

Databricks was founded in 2013 to fix exactly this. The founders, a team from UC Berkeley who had invented Apache Spark, the open-source engine now used by thousands of companies worldwide, who built what they called a **lakehouse**. The idea was simple: collapse the data lake and the data warehouse into one unified platform. One place where engineers build pipelines, analysts run queries, and data scientists train AI models.

It worked. Today, companies like Adidas, AT&T, Mastercard, Rivian, Block, and Unilever run their data and AI operations on Databricks. So do more than 60% of the Fortune 500. The platform has become the standard infrastructure layer for enterprise AI, the thing sitting

underneath all the models, agents, and dashboards that large organisations are now racing to build.

In February 2026, Databricks announced it had crossed a **$5.4 billion revenue run-rate**, growing more than 71% year over year. The company is free cash flow positive and preparing for an IPO. Vested gives you the opportunity to invest before that happens.

> **One way to think about it:** Databricks is to enterprise AI what AWS was to cloud computing: the infrastructure layer that everyone builds on top of. You might not always see it, but it is almost certainly powering something you interact with every day.

02 The Growth: Numbers That Stand Out

Let's look at the numbers, because they are genuinely exceptional for a company at this scale.

Databricks hit $1 billion in annualised revenue in 2022. Two years later it was at $3 billion. By January 2026 it crossed $5.4 billion. That is more than a 5x increase in three years, while maintaining gross margins in the 74–80% range.

Revenue Run-Rate Growth

Annualised ARR in USD billions (Source: Sacra, Feb 2026)

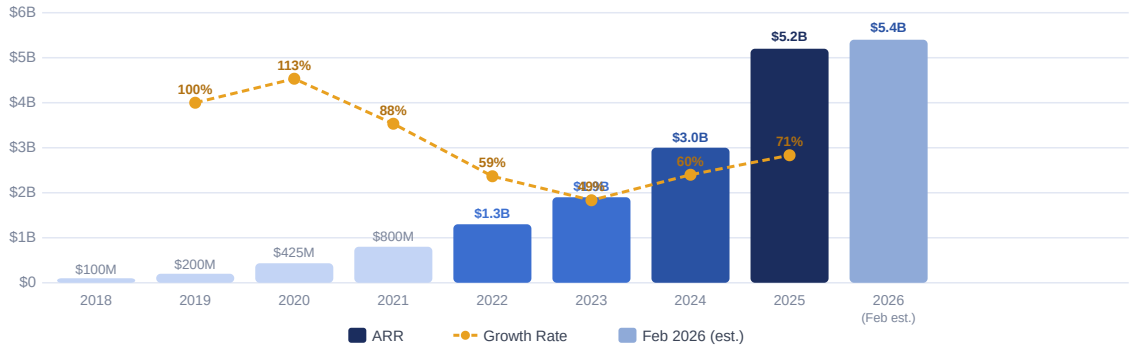

Most enterprise software companies growing at this scale are expanding 20–30% a year. Databricks is growing at more than double that rate. What is driving it?

- **AI workloads:** $1.4 billion in annualised revenue from AI products alone, about 26% of total revenue. Every enterprise trying to build AI on their own data is a potential customer.
- **Databricks SQL:** The data warehousing product went from $100M ARR to $400M ARR in a single year.
- **Lakebase:** Their new AI-native database reached 2x the revenue of Databricks SQL at the same 8-month stage, the fastest product ramp in company history.
- **Enterprise depth:** More than 800 customers spending over $1 million annually. More than 70 spending over $10 million.
- **Customers keep spending more:** Net dollar retention is above 140%. That means for every $100 an existing customer spent last year, they are spending at least $140 this year, without Databricks needing to win any new customers.

What is the $5.4B made of?

Revenue by segment, Jan 2026

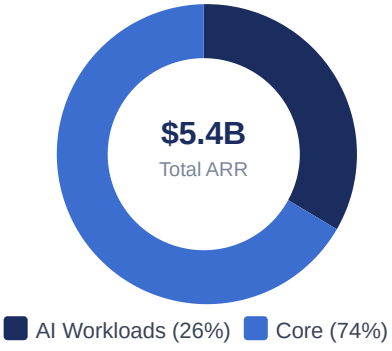

Revenue Growth Rate by Year

Year-over-year %

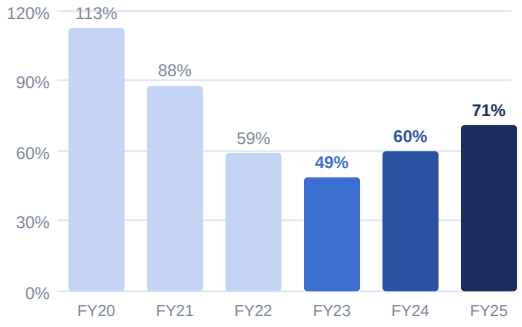

What net dollar retention above 140% actually means: If Databricks signed zero new customers from today onwards, existing customers expanding their usage would alone grow revenue by 40%+ over the next year. That is an unusually strong signal of product-market fit.

03 What Databricks Actually Does

The product sounds technical, but the business logic is straightforward. Databricks replaces multiple expensive tools with one integrated platform. Every time a customer consolidates another workload onto Databricks, they spend more, which is why net dollar retention is so high.

> *Think of Databricks as the operating system for enterprise data. Every AI model a company trains, every dashboard an executive reads, every data pipeline an engineer maintains; increasingly, it all runs on Databricks.*

The platform today has six main product areas:

- **Lakehouse:** The core platform: unified data lake and warehouse built on Delta Lake (open-source, ACID-compliant). This is what made Databricks famous.
- **Unity Catalog:** A governance layer that tracks who can access what data, with full audit trails. Critical for regulated industries like banking and healthcare.
- **Mosaic AI:** End-to-end machine learning, from training models to deploying them and monitoring their outputs in production.
- **Agent Bricks:** Infrastructure for building autonomous AI agents that reason over a company's own data. Integrates with both OpenAI and Google Gemini models natively.
- **Genie:** A conversational interface that lets non-technical employees ask questions about data in plain English and get answers, no SQL required.
- **Lakebase:** A serverless Postgres database purpose-built for the AI agent era. Built on the acquisition of Neon (May 2025, ~$1 billion).

> **Why this matters for investors:** Each new product expands the amount of a customer's budget that Databricks can capture. A company that starts using Databricks for data pipelines often ends up also using it for SQL analytics, then AI model training, then agent infrastructure. This is why customers keep spending more.

04 Who Is Backing This Company

The investor list matters here, not because famous names guarantee outcomes, but because of what some of these investments signal strategically.

NVIDIA, the company that makes the chips powering the AI revolution, is a strategic investor. Microsoft, which also uses Databricks as a key distribution partner on Azure, invested in the Series L round. Google Cloud partnered with Databricks to bring its Gemini 3 Pro model natively onto the platform. These are not passive bets; they are partnerships from companies whose own AI revenue is tied to Databricks succeeding.

SERIES J · NOV 2024

$10B raised

$55B valuation · Lead: NVIDIA, a16z

SERIES K · SEP 2025

$1B raised

$100B valuation · Lead: a16z, Thrive

SERIES L · DEC 2025

$5B raised

$134B valuation · Lead: Insight, Fidelity

Valuation at Each Funding Round

USD billions (Source: Sacra, Feb 2026)

Total capital raised to date: over $15 billion. Key institutional backers include Andreessen Horowitz, JPMorganChase, Goldman Sachs, T. Rowe Price, Tiger Global, Coatue, NEA, the Qatar Investment Authority, and BlackRock. The Series L round also included Morgan Stanley, Neuberger Berman-affiliated funds, and UBS.

> **The valuation has more than doubled in twelve months**: from $55B in November 2024 to $134B at the time of Series L, driven by a business that kept its growth promises. The current valuation is $150.8B, reflecting secondary market movement since then.

05 The New Bet: Lakebase and the Database Market

For its first decade, Databricks focused on analytical workloads, helping companies store, query, and run AI on historical data. That alone is a large market. But in 2025, the company made a move that signals it is going after something much bigger.

In May 2025, Databricks acquired **Neon**, a cloud-native Postgres database, for approximately $1 billion. Neon's customers include OpenAI, Adobe, and Replit. The product built on this acquisition is Lakebase: a serverless Postgres database designed specifically for AI agents that need to read and write live data in real time.

This is a move into the operational database market, the $100 billion-plus market that companies like Oracle, AWS, and Google have historically dominated. Databricks is not yet the leader there, but the early traction is notable. Lakebase hit 2x the revenue milestone that Databricks SQL achieved at the same 8-month stage, and Databricks SQL itself was considered one of the fastest product ramps in the company's history.

> *AI agents need to do more than read from a database; they need to write to it in real time. Lakebase is Databricks' answer to that problem. If it works, it opens an entirely new category of enterprise spend.*

06 The Valuation: Expensive or Justified?

Let's be direct about this. At $150.8 billion (the current valuation on Vested), Databricks is priced at roughly 28x its annualised revenue. That is expensive by the standards of most public software companies. The question is whether the growth rate justifies it.

COMPANY	REVENUE	GROWTH YOY	REVENUE MULTIPLE
Databricks (private)	$5.4B ARR	>65%	~28x (at $150.8B)
Snowflake (public)	~$4.7B	~29–30%	~13x
Palantir (public)	$4.5B	~56%	~82x
Typical SaaS (public)	Various	20–30%	8–15x

Revenue Growth Comparison: Databricks vs Snowflake

Year-over-year revenue growth rate at comparable revenue scale

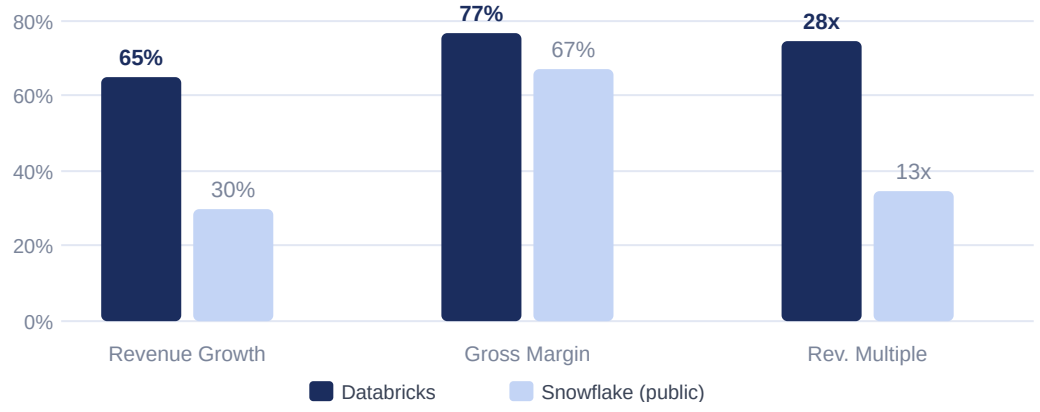

The key difference between Databricks and Snowflake, is the growth rate. Snowflake is growing at around 29–30% at a similar revenue scale. Databricks is growing at more than double that. If you believe the growth gap justifies a valuation premium, the question becomes: how long can Databricks sustain this rate?

There is also the profitability question. Databricks is already **free cash flow positive**, a milestone that many private companies at this scale have not yet hit. Gross margins of 74–80% are consistent with high-quality cloud software businesses. The path to sustained profitability at scale is visible.

> **A simple way to frame it:** If Databricks can sustain 40–50% growth over the next three years and reach $10–12 billion in revenue by 2028, today's valuation does not look unreasonable for a company that is also free cash flow positive and expanding into new markets. If growth decelerates sharply, the valuation math gets harder.

07 What Could Go Wrong

Any honest analysis of Databricks has to include the risks. Here is what to watch for.

What works in your favour

+ Free cash flow positive: not burning through capital
+ Net dollar retention above 140%: existing customers keep expanding
+ Open-source roots (Spark, Delta Lake, MLflow) drive organic adoption
+ No single cloud dependency: runs on AWS, Azure, and GCP
+ Lakebase opens a $100B+ new market
+ Strategic investors (NVIDIA, Microsoft) have vested interest in success
+ IPO expected in 2026: a clear liquidity path

What to watch out for

- 28x revenue is priced for continued high growth: any slowdown hurts
- Snowflake, AWS, Azure, and Google all competing directly
- Open-source tools can be self-hosted: free alternatives always exist
- Large enterprise contracts mean concentrated revenue risk
- Private market: you cannot sell whenever you want
- IPO timing uncertain: macro or market conditions could delay

08 How This Deal Works on Vested

Databricks is a private company, so you cannot buy shares on a stock exchange. Vested gives investors access through a Special Purpose Vehicle (SPV), a pooled structure that holds the shares on behalf of participating investors.

What you are buying

Series L Preferred shares at $212.80 per share. Minimum investment is $5,000. This is a primary transaction; shares come directly from the company.

When can you exit?

When Databricks IPOs (expected 2026), the SPV converts to public shares transferred to your Vested account. You can then hold or sell like any listed stock. Pre-IPO, secondary sales are possible via Vested's platform and partner Monark's Alternative Trading System.

What if there is no IPO?

If Databricks is acquired instead, the SPV receives cash or stock consideration and distributes it to investors. If neither happens for an extended period, your capital remains locked in a private investment with no guaranteed liquidity timeline.

Who is this for?

Investors with a long-term horizon (ideally 3 to 5 years), comfortable with illiquidity, and who understand that private market investments carry the full risk of loss. This is not a fixed-income instrument.

The honest trade-off: You get access to a high-quality private company before the IPO, at the cost of illiquidity. The bet is that the share price at IPO, or after, is meaningfully higher than your entry price of $212.80. Based on Databricks' current trajectory, many investors believe that is a reasonable bet. But it is not guaranteed.

The Bottom Line

Databricks is one of the most important infrastructure companies of the enterprise AI era. The revenue is real, the growth is exceptional, the unit economics are working, and an IPO is on the horizon. The platform is embedded in the workflows of over 20,000 organisations, including more than 60% of the Fortune 500.

The valuation is demanding and execution risk is real. Those are real considerations.

But if you believe enterprise AI is the defining technology investment of this decade, and you want exposure to the infrastructure layer powering it, before the rest of the market gets access at an IPO price, this is one of the most direct ways to do that.

Investment Report | March 2026 | Vested