

# HIRING DATABRICKS ENGINEERS IN 2026



# TABLE OF CONTENTS

|                                     |   |
|-------------------------------------|---|
| INTRODUCTION                        | 1 |
| BACKGROUND<br>TO DEMAND             | 2 |
| HIRING DATA                         | 3 |
| SALARY<br>BENCHMARK                 | 4 |
| BLOCKERS TO<br>SUCCESSFUL<br>HIRING | 5 |
| FINAL THOUGHTS                      | 6 |

# INTRODUCTION



**DAVE CARROLL**  
Senior Consultant  
KDR Talent Solutions

Databricks continues to be at the centre of enterprise data transformation across the United States. Chicago's strong presence in trading firms, supply chain leaders, and enterprise corporates has created a unique demand for Databricks engineers who can operate across both real-time data and large-scale batch processing environments.

This guide is designed to help hiring managers and data leaders make informed decisions around hiring, especially important when competition for talent is fierce. It will cover salary expectations and recommended hiring timelines as well as technical evaluations and platform maturity. It's a practical breakdown of what's happening, and how to adapt.

So, what's driving the demand for engineers with Databricks? From what we can see here at KDR Talent Solutions, the rise in demand for Databricks engineers seems to be the result of a growing gap between adoption of AI and available skills as well as the foundations required to make AI projects successful.

To speak to this last point, a big driver of Databricks' growth is the mainstream adoption of the Lakehouse architecture. The platform has become synonymous with a new standard for modern data engineering, combining the scalability of data lakes with the reliability of data warehouses. With native support for structured and unstructured data, Delta Lake performance, and integrated governance features, Databricks has effectively positioned itself as the default choice for companies looking to break free from legacy limitations.

Industry analysts such as Gartner have forecast that more than 40% of data teams will shift to a Lakehouse approach by 2027. **Nowhere is this more apparent than in innovation-led metros like New York City, where organisations are building modern, scalable data platforms. Their challenge is now in resourcing that shift.**

## 40%

of Data teams will shift to a  
Lakehouse approach by 2027



# BACKGROUND TO THE CURRENT HIRING DEMAND

In late 2023, Databricks raised \$500 million in Series I funding, bringing its valuation to \$43 billion. The investment round included Nvidia, Capital One Ventures and major cloud players like Amazon Web Services, all of whom are tightly aligned with Databricks' AI-first data strategy. This was a signal to the market and it was a launchpad for a broader ecosystem expansion that's been felt ever since.

Enterprise partnerships with Microsoft Azure, integrations with Unity Catalog for governance, and product acquisitions such as MosaicML have turned Databricks into far more than a data engineering tool. It is now a core part of many organisations' AI and machine learning workflows, further intensifying the demand for engineers who can work across ingestion, transformation, model development and deployment.

# WHAT THE HIRING DATA TELLS US

From a job market perspective, the data is just as compelling. Over the past 12 months we have seen:

Job postings for Databricks roles in the US rising by over 22%, with Chicago emerging as a major hub for Databricks hiring, particularly across enterprise and trading-focused organisations.

On average, each open Databricks role receives tens if not hundreds of applications, yet fewer than five qualified applicants, many of whom are in active interview processes elsewhere.

The contract market is surging, particularly among Series A and B startups building MVP data platforms but lacking in-house capability.

Databricks is now appearing as a core skill requirement, not just a nice-to-have, a shift that has happened within just 18–24 months.

Organisations are waking up to the reality that data maturity isn't possible without engineering maturity and Databricks sits right in the centre of that conversation.





# SALARY EXPECTATIONS FOR 2026

Now we know the demand, what are the salary expectations for 2026? Salary expectations for Databricks engineers have continued to rise, especially in remote-first or hybrid roles anchored in major US cities like Chicago, where demand remains high but cost pressures are more balanced than coastal markets.. From our data; equity and bonus structures are still prevalent among growth-stage companies and can be a key differentiator when competing against more established teams.

In Chicago, we see particularly strong salary competition from trading firms, financial services, and logistics-led organisations, often pushing senior-level compensation toward the top end of these ranges. Top-tier firms (particularly in trading and finance) may exceed these ranges significantly through bonuses and total compensation packages.

| Role                         | Salary        | Location/Trends                                                   |
|------------------------------|---------------|-------------------------------------------------------------------|
| Contract/Freelance           | \$90-\$130/hr | Startup MVPs, migration projects, short-term scale-ups            |
| Databricks Platform Engineer | \$140k-\$170k | Cloud-native teams, especially AWS-heavy                          |
| Lead/Principal Data Engineer | \$165k-\$195k | Enterprise and scale-up environments                              |
| Senior Data Engineer         | \$150k-\$180k | Chicago hybrid roles across FinTech, trading, logistics, and SaaS |
| Mid-Level Data Engineer      | \$120k-\$145k | Remote-first; Central or Midwest time zones preferred             |

# WHAT ARE THE BLOCKERS TO SUCCESSFUL HIRING



Many hiring processes for Databricks roles are still not fit for purpose. We see too many companies losing out on top talent because of lengthy, or misaligned hiring processes. These are the most common mistakes that we see companies are making:

## 21%

Interview Hours per Technical Hire have increased by 21% since 2021

## 14

It takes circa 14 more interview hours to close a technical hire than a business hire.

Ashby Talent Trends Report

## TOO MANY STAGES

Five-round interview processes might seem rigorous, but in a market where good engineers are fielding multiple offers, they simply won't wait. Three steps should be enough to assess skill, mindset and culture fit, anything more is risking drop-off.

## 63%

of candidates reported receiving a job offer in their most recent hiring process, meaning they have choice whilst in process.

Gartner

# WHAT ARE THE BLOCKERS TO SUCCESSFUL HIRING



## MISALIGNED TECHNICAL ASSESSMENT

Many organisations continue to use generic coding tests that don't reflect the day-to-day challenges of Databricks engineering. Live pairing sessions on pipeline optimisation, SQL performance, or cloud orchestration are far more effective and give candidates a better sense of the environment they'd be working in. On this, some companies come up with highly complicated and lengthy technical tests that require many hours of commitment, and for some, on top of the day job this is just too much of an ask.

## DELAYED FEEDBACK

This is a biggie. Time kills momentum and perception of a company. From our experience, candidates should receive feedback within 48 hours of any interview stage. Companies that can move quickly, while still being thoughtful, consistently outperform slower competitors. Candidates can be really put off by companies that just don't feedback, or delay feedback, it feeds into the negative rather than the positive perception bracket. If you set processes up with this in mind it shouldn't ever be an issue. Problems generally arise when processes are ad-hoc.

## NOT SELLING THE DATA TEAM EVP

Data professionals want to know the impact they can have in a role but many companies don't talk about this, the cool projects they are working on or the impact they have.

# WHAT ARE THE BLOCKERS TO SUCCESSFUL HIRING



## HIRE FOR POTENTIAL NOT PERFECTION

One of the most important shifts we recommend to hiring teams is a move away from “stack-matching” and toward capability-based hiring. It’s becoming increasingly rare to find candidates with 4–5 years of direct Databricks experience, especially those who have also worked across AWS, MLflow, Delta Live Tables, and Unity Catalog. But what about those who’ve migrated from Spark, have strong fundamentals in Python, SQL, and cloud infrastructure, and a demonstrable eagerness to learn? These are the hires that will deliver value and grow with your platform.

Databricks offers one of the most accessible learning paths in the modern data space. Candidates with the right foundation and the right mindset can be brought up to speed far faster than most realise. Ignoring this pool is becoming a competitive risk.

# FINAL THOUGHTS



The use cases for Databricks is growing and the engineers who can build on top of it are among the most in-demand professionals in today's market. We think to win in this space, hiring managers need to double down on:

Competitive salaries

Feedback baked into processes from the start

Flexibility in experience requirements

Speed and clarity in feedback

If you're building a data platform or AI function that relies on Databricks, the talent is out there but you need the right strategy to access it.



I speak with US-based Databricks engineers every day. I know what they want, what they expect, and what motivates them to join new teams. If you want real market intelligence, not just a pile of CVs, get in touch.

[dave.carroll@kdrtalentsolutions.com](mailto:dave.carroll@kdrtalentsolutions.com)

Let's connect on LinkedIn